ORNL/TM-11810
(CESAR-91/07)

ornl

OAK RIDGE
NATIONAL
LABORATORY

MARTIN MARIETTA

Proposal for Continued Basic Research in Intelligent Machines at the Center for Engineering Systems Advanced Research Fiscal Years 1992–1996

R. C. Mann
M. Beckerman
K. Fujimura
J. P. Jones
E. M. Oblow
F. G. Pin
M. A. Unseren

MANAGED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

ORNL/TM-11810
CESAR-91/07

Engineering Physics and Mathematics Division

PROPOSAL FOR CONTINUED BASIC RESEARCH IN INTELLIGENT MACHINES AT THE CENTER FOR ENGINEERING SYSTEMS ADVANCED RESEARCH FISCAL YEARS 1992 - 1996

R. C. Mann, M. Beckerman, K. Fujimura, J. P. Jones, E. M. Oblow,
F. G. Pin, M. A. Unseren

Prepared for the
Engineering Research Program
Office of Basic Energy Sciences
U.S. Department of Energy

DATE PUBLISHED: MAY 1991

Prepared by the
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
Managed By
MARTIN MARIETTA ENERGY SYSTEMS, INC.
for the U.S. Department of Energy
under Contract No. DE-AC05-84OR21400

TABLE OF CONTENTS

ABSTRACT

This proposal describes the context and technical direction during the next five years for the DOE/ER/BES-sponsored basic research program in intelligent machines at the Center for Engineering Systems Advanced Research (CESAR) at the Oak Ridge National Laboratory (ORNL). Research will address issues related to autonomous systems in unstructured dynamic work environments. Specifically, the work will focus on multiple cooperating robotic systems, combined mobility and manipulation, intelligent sensor systems, machine learning, and embedded high performance computing. Focus for proof-of-principle experiments demonstrating innovative developments will be derived from a number of application areas related to DOE missions.

1. Introduction

CESAR was established during FY 1984 at ORNL for the purpose of addressing fundamental problems and issues arising in the development of intelligent machines. This decision was a result of a growing awareness of the fact that automation-related technologies and intelligent machines can increase productivity and safety in the development and operations of DOE-sponsored systems, and that a well-focused basic research program for intelligent machines was necessary.

The research program at CESAR was initiated based on a commitment to long-term support, pending peer-review, from the Engineering Sciences Program in the Office of Basic Energy Science, Office of Energy Research (DOE/BES/ER). This support has provided the environment necessary to build and maintain a core basic research effort at CESAR.

We characterize intelligent machines as systems that integrate perception, reasoning and action in order to perform tasks under circumstances that can be insufficiently known in advance, and dynamically changing during task performance. Perception, reasoning, and action as well as the integration of these modules into working systems present challenging issues for inter-disciplinary basic research. Sensor-based robots that can execute their missions in unstructured work environments are ideal testbeds for such research. Therefore, part of the approach at CESAR consists in investing a balanced fraction of the resources into the development of an evolving series of mobile robot prototypes HERMIES (**H**ostile **E**nvironment **R**obotic **M**achine **I**ntelligence **E**xperiment **S**eries). These machines serve as testbeds for new methods, and hardware and software developments. They have been used in proof-of-principle experimental scenarios that provide the necessary quantitative data for testing and validation of new approaches, as well as for performance evaluation of different robot system components in integrated systems. HERMIES-IIB and -III are currently operational mobile robots.

The performance record of the CESAR core activity funded by DOE/BES/ER has attracted interest and support from other Offices within the DOE and from non-DOE sponsors. CESAR is a contributor to the National Robotics Technology Development Program (NRTDP) for Environmental Restoration and Waste

Management (ER&WM). For the last four years, CESAR has been providing project management and leadership, and systems integration for a robotics technology base development effort by a consortium of four U.S. universities and industrial partners, sponsored by the Department of Energy's Office of Nuclear Energy (DOE/NE). CESAR is fulfilling a lead role in coordinating the new DOE robotics and related intelligent systems effort for space applications. Non-DOE sponsors include the U.S. Air Force Wright Research and Development Center, the Advanced Concepts and Technology Program of the U.S. Army Material Command, and the U.S. Army Fort Belvoir R&D Center. Figure 1 illustrates the context of the CESAR basic research activity with efforts that focus mainly on needs-driven robot system development.

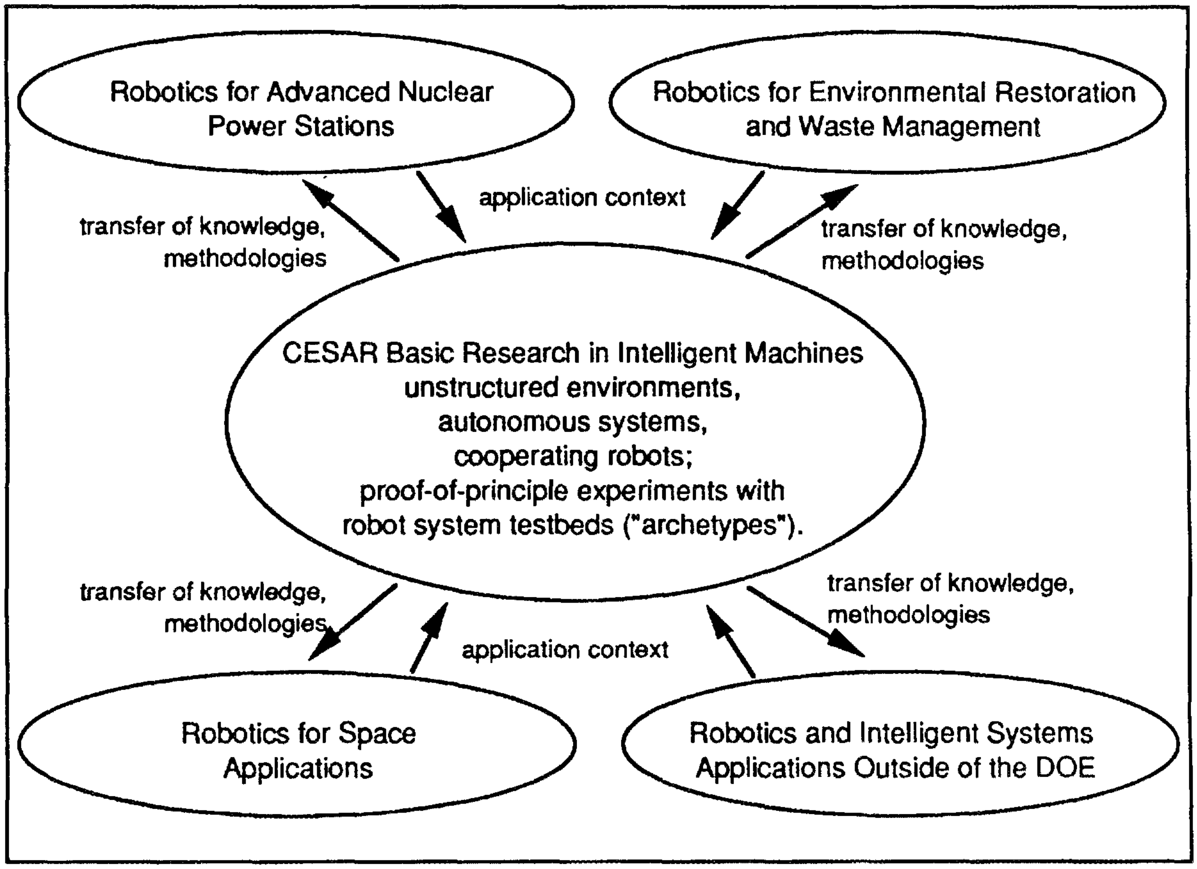

Figure 1: CESAR context with respect to other programs.

Many international connections exist: CEA (robot simulations), France; University of Bochum (neurocomputing), Germany; Aalborg University (robot vision), Denmark; University of Leuven (parallel computing), Belgium.

Experiments with the CESAR robots during the past years have shown that:

- reliable autonomous robots can be built that can navigate autonomously indoors in the presence of moving obstacles based on simple visual information integrated with sonar distance measurements;

- robots can generate precise maps of workspaces that are incompletely known a priori;

- robots can be equipped with massively parallel, general-purpose computers that make it possible to analyze sensor data in a few multiples of the sensor frame rates, and to make decisions accordingly in real time; the parallelism can be made transparent to the applications programmer;

- autonomous robots can make effective use of the redundancy resulting from the combination of mobility degrees of freedom (DOF) with manipulator DOF;

- robots can learn from examples, e.g. they can be trained before entering a workspace to perform tasks similar to the ones they will be confronted with during their mission.

These results were demonstrated in the course of experiments with the HERMIES robots in scenarios that included navigating in an incompletely known, dynamically changing laboratory environment with unexpected events; performing diagnosis, learning and simple manipulation tasks at a mock-up process control panel; and the autonomous removal of a simulated chemical spill in the laboratory.

Detailed technical accounts of the research at CESAR including descriptions of robot demonstrations can be found in over 130 publications (see Appendix B).

An independent committee appointed by DOE/HQ reviewed the CESAR program during the last quarter of FY 1987. It was the unanimous recommendation of this committee panel that the program should be continued for at least an additional five years. FY 1992 marks the end of that period. This proposal describes the scope and subject of research at CESAR through FY 1996.

Summary Roadmap for CESAR Research in Intelligent Machines

1983 - 1987: First Autonomous Mobile Systems

Development, testing and evaluation of the mobile robots HERMIES-I, -II, IIB; development of CESARm redundant robot manipulator; real-time redundancy resolution and position control of CESARm; real-time rule-base system for autonomous navigation in unknown environment based on sonar sensor data; high-precision map building (sonar sensing); new hybrid uncertainty calculus; hypercube concurrent computers integrated in robot control system; concurrent computer vision system applied to robot object recognition, docking and simple manipulation at a process control panel.

1988 - 1991: First Autonomous Mobile Manipulator System

Development, testing and evaluation of the HERMIES-III mobile robot research testbed combining for the first time mobility, dextrous manipulation, advanced multi-sensing, and concurrent computing on-board an integrated robot system; quantitative assessments of robust performance of sensor-guided tasks; library of scalable parallel robot vision algorithms; fusion of sonar and 2-d vision information for robot navigation; robot path planning and navigation based on laser range images; robot learning of actions at process control panel based on experience, using expert system and neural network approaches; random set analysis of machine learning algorithms indicates possible scalability to highly complex, practical tasks; first analytical and simulation results for task-driven planning and redundancy resolution for mobile redundant robotic manipulators (combined mobility/manipulation systems); fast path planning for robots with non-holonomic constraints; custom VLSI fuzzy logic hardware tested onboard new mobile platform for reactive local path planning; leveraging of programs with strong application focus results in first demonstration of robotic emergency response capability during simulated chemical spill clean-up experiment, and in robotic assessment of radiation contamination on irregular surfaces; CESAR evolves into collaborative research facility.

1992 - 1996: First Autonomous Cooperating Systems

Continued proof-of-principle demonstrations with the HERMIES-III robot, and development, testing and evaluation of its successor model; dual manipulator research facility established; mesh-connected concurrent computer system (>320 Mflops) integrated in robot control system; efficient real-time control of mobile manipulator systems; planning and control for dual manipulator systems; motion planning for cooperating robots; robust multi-sensor systems, including fusion of color images, laser range images, and other sensor domains for classes of tasks, including exploration and dextrous manipulation; robot architecture supporting high-level planning and reactive behavior; learning by robotic systems for complex workspaces and classes of tasks; robot experiments focusing on classes of tasks in exploration and characterization of unknown workspaces, and on reliable performance of robotic assembly, construction, maintenance, surveillance, and emergency handling tasks.

Table 1: Summary Roadmap for CESAR Research Agenda

Table 1 provides a summary roadmap of the long-term CESAR research agenda. It shows a capability trail for advanced intelligent machines as demonstrated with the HERMIES robots. The roadmap starts with the first autonomous systems for which the action module encompassed mobility and very limited manipulation (FY 1984 - FY 1987), and proceeds during the second phase (FY 1988 - FY 1991) to the first mobile redundant robotic manipulator system with demonstrated initial capabilities to perform human-scale tasks. During the period covered by this proposal (FY 1992 - FY 1996) we will continue our research towards the next generation intelligent machines with integrated multi-sensor systems, mobile manipulators, embedded high performance computing, and learning capabilities. In addition, we propose to pursue our effort in cooperating multiple robots which we started in FY 1991 with research in multi-robot motion planning, and modeling and control of dual-manipulator systems.

In each of these areas we propose selected high priority research tasks. The combination of these focal points represents the unique integrated approach to basic research in intelligent machines at CESAR:

- generation of minimal world models, and treatment of systematic and random sensor errors through sensor fusion;
- determination of optimal commutation configurations under various constraints to effect optimal combined use of mobility and manipulator DOF;
- parallel processing using mesh-connected distributed-memory processors for applications in real-time control based on imaging sensor data;
- learning of concepts of realistic practical complexity by addressing critical problems in concept representations and algorithmic scalability;
- decentralized motion planning and control of multiple mobile robots with heterogeneous perception and reasoning characteristics;
- position/force control of two interacting manipulators handling a common object.

Although our motivation in selecting these areas is primarily based on the CESAR mission to address fundamental research issues in intelligent machines, our research is also motivated by the long-term technology needs of activities that make use of robotics and intelligent systems (see Fig. 1).

At the end of the five-year period, a number of important advancements in crucial enabling technologies for intelligent machines can be expected from this program. These include advancements in: (1) cooperative problem solving and task performance by multiple robotic systems; (2) robotic control of mobile, redundant manipulators to optimize task-dependent performance criteria; (3) task-driven minimal representations of the robot's work environment based on multiple sensors; (4) learning by robotic systems in environments of practical complexity; (5) mobile robot laboratory testbeds and computational architectures that include parallel processors, and corresponding software environments that support reliable operation of these machines, and portability and re-usability of key elements of the software.

2. CESAR Charter

CESAR was established to address long-range, energy-related basic research in intelligent machine systems. These systems are intended to plan and perform a variety of tasks in incompletely known environments, given only qualitatively specified goals. The Center provides a focal point for interdisciplinary research in machine intelligence, cognitive systems, advanced control, and systems engineering. Research objectives are chosen to address the long-term technology-base requirements for DOE missions that rely on the use of intelligent machines and robotics.

CESAR is intended to be a national resource, and a major objective is to disseminate R&D accomplishments freely and comprehensively. Results and technology advancements are distributed through publications in the scientific literature, through organization of workshops on selected topics, and through the development of prototype systems that demonstrate concepts and new methodologies. CESAR hosts guest researchers from universities, laboratories and industry, serving as a collaborative research facility to provide guests with access to state-of-the-art and often unique equipment in a stimulating research environment.

3. Synopsis of Selected Past Accomplishments

A detailed and comprehensive account of past CESAR accomplishments is beyond the scope of this proposal. This section provides brief summaries of selected achievements that have impact on program directions during the next five years.

3.1 Intelligent Machine Research Facilities

HERMIES-IIB and HERMIES-III (Fig. 2) are the currently operational mobile robots at ORNL/CESAR. HERMIES-IIB stands 1.0 m high and weighs 91 kg. Rechargeable batteries supply 20 watts of power, providing about 20 minutes of untethered running time. Peak movement speed is 0.7 m/s. Sensors include three Sony CCD cameras and 20 Polaroid sonar transceivers mounted on a rotatable turret. The computer architecture consists of a VME rack housing a 68020 CPU, and a variety of I/O boards interfaced via a BIT-3 communication link to an NCUBE, Inc. (Beaverton, OR) hypercube computer featuring 16 nodes with 512 Kbytes RAM each and an Intel 80286 I/O processor, which also serves as a host for the hypercube. Each node processor is a 32-bit microcomputer with on-chip floating point and communications hardware. This gives HERMIES-IIB roughly 16 MIPS in the on-board hypercube. HERMIES-IIB is equipped with two Zenith/Heathkit five degree-of-freedom arms which give the robot severely limited manipulative capability. This has not been a drawback, however, since the robot was not intended for research in manipulation.

HERMIES-III is a mobile robot consisting of an omni-directional wheel-driven chassis, a seven degree-of-freedom manipulator (CESARm), an Odetics laser range camera, multiple cameras (two stereo pan and tilt mechanisms), an array of 32 sonar transceivers, and an on-board computer system that includes five Motorola 68020s in four VME racks, and an NCUBE hypercube computer. CESARm is a compliant, high capacity-to-weight ratio (1/10) arm, with an adjustable gripper, which is equipped with a JR3 force-torque sensor; a LORD tactile sensor pad is being installed. A Fairchild-Weston CCD camera is also mounted on end-effector.

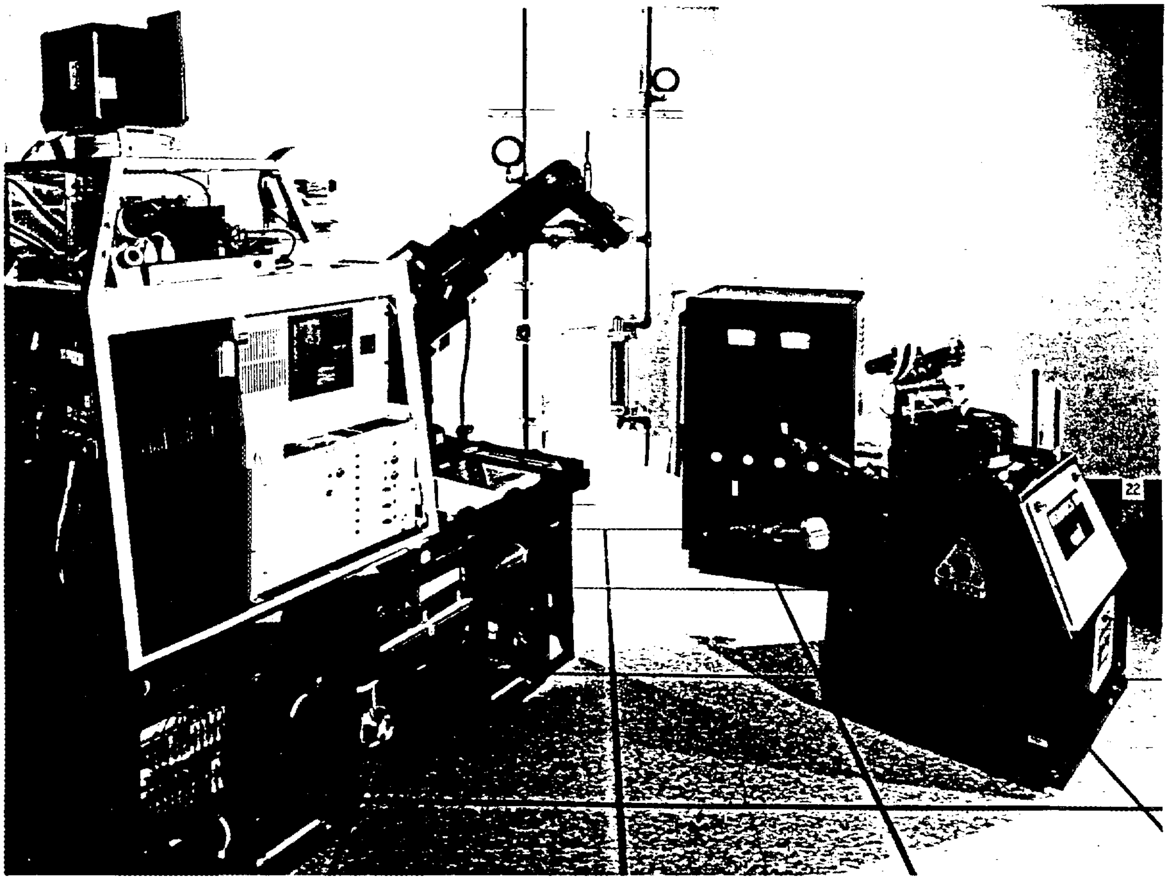

Figure 2: The CESAR mobile robots HERMIES-IIB (right) and HERMIES-III (left).

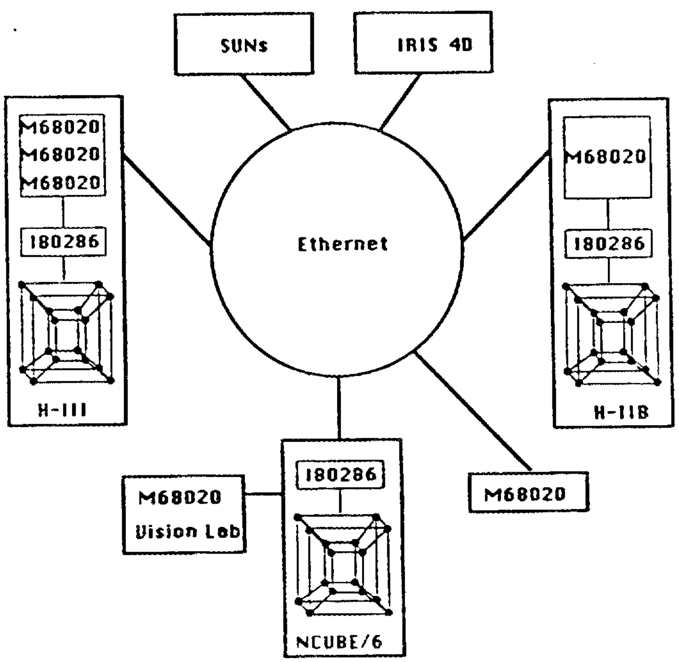

Figure 3: Schematic diagram of the CESAR robot-computer network.

Both robots can be operated completely autonomously, in which case they communicate via RS-232 wireless modems to off-board computers. They can also be interfaced through ethernet to a local area network of computers, as schematically shown in Fig. 3. A diagram of the on-board computer architecture for HERMIES-IIB is shown in Fig. 4.

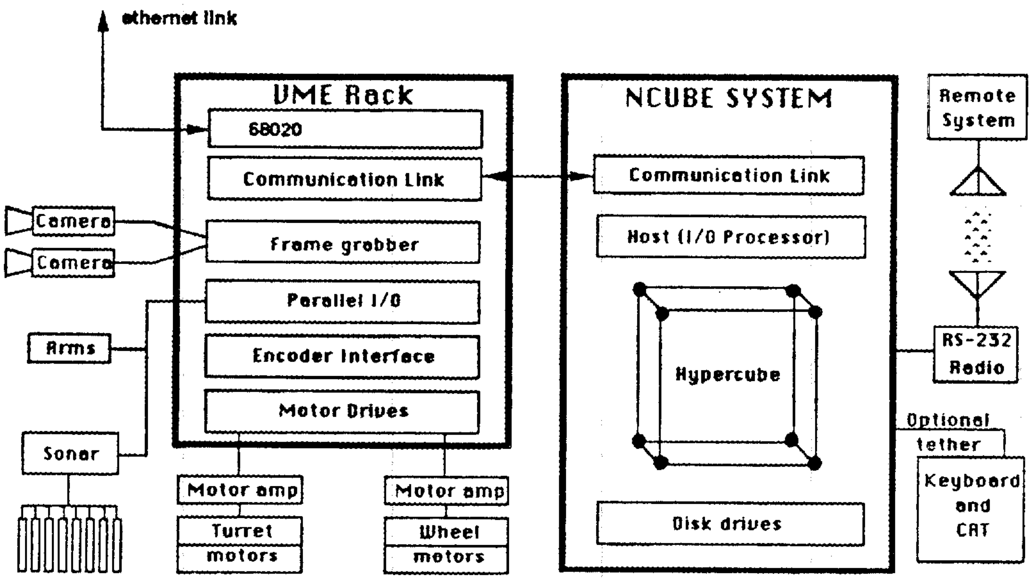

Figure 4: Schematic diagram of HERMIES-IIB on-board architecture.

3.2 Manipulator Kinematics, Dynamics and Control

Research in the kinematics, dynamics and control of kinematically redundant manipulators has been focused on the CESAR research manipulator CESARm [I.12, I.29, II.15, II.28, II.29, II.40, II.41, II.74]* , which was developed as a testbed for control methodologies suitable for applications in unstructured environments. CESARm embodies several unique features. It is a seven DOF

* Numbers refer to CESAR publication list in Appendix C.

kinematically redundant manipulator. It was designed to accommodate high tip (end-effector) speeds (4ft/s), and features a large payload capacity to manipulator weight ratio. CESARm can be operated in a stand-alone, stationary base configuration or as a mobile robot manipulator when mounted on HERMIES-III. To an observer, CESARm appears to be an open chain, serial link mechanism consisting of rigid links connected by single DOF revolute joints. In reality, the manipulator contains two internal four-bar closed loops which are used to actuate the elbow pitch joint. These closed kinematic chains lead to a constrained motion even when the tip/end effector of CESARm is unconstrained. CESARm is a compliant manipulator due to flexibility at the joints and in the cable drive systems, non-rigidity of the links, and backlash.

CESARm is currently mounted on HERMIES-III, which provides the CESAR team with a uniquely powerful testbed for research in dextrous manipulation combined with robot platform mobility [II.81, II.84].

3.3 Multi-Sensor Data Analysis

One of the prerequisites for intelligent behavior in robotic systems is the ability to generate internally self-consistent representations of the environment. In general, this is impossible on the basis of either an isolated sensing operation or a single sensor domain. The term multi-sensor integration denotes the task of combining data and information from more than one robot location, and/or more than one sensor, so that a consistent world model, i.e., a model free of contradiction can be generated, on the basis of which decisions concerning navigation, manipulation, etc. can be made [V.8, V.16]. This objective is particularly difficult to achieve when the robot must operate both autonomously and robustly in an unstructured environment. In particular, systematic errors in the interpretation of the sensor data are likely to occur, producing internal inconsistencies. This class of errors can be treated by carrying out the multi-sensor integration at or below the processing level at which the errors occur. This type of methodology makes use of multiple world models and sensing models of the relevant physical and spatial properties.

Past accomplishments in this area include innovative methods to fuse ultrasound sensor data taken at different points in space and time [I.24], and the integration of data from a 2-D vision sensor (CCD camera) with ultrasound sensor data [II.68, VI.43], in support of world-modeling for robot navigation.

3.4 Robust Performance of Sensor-Guided Tasks

Issues related to reliability and robustness of advanced robotic systems are receiving increasing attention in the robotics research community. The large discrepancy in the level of sophistication of robots that are currently in use in a commercial operational environment, and the robotic devices that are being used in laboratory proof-of-principle demonstrations, is partly due to the fact that systematic studies of reliability and robustness have been largely neglected by the research community.

In an initial experimental investigation [II.68] we studied the performance of a simple integrated system implemented on the HERMIES-IIB robot. Tasks performed by this system included ultrasound-guided navigation, object recognition, vision-guided navigation and docking, and vision-guided manipulation. In each case, closed-loop sensor feedback was used to guide the associated hardware task. Our objective was to measure the performance of each of the subsystems, and then to identify the main problem areas in acquiring and interpreting the sensor information needed to control autonomous locomotion and manipulation tasks.

3.5 Parallel Computing for Intelligent Machines

CESAR has pioneered the use of concurrent multi-computers in the overall architecture for advanced mobile robots. Selected innovative developments by CESAR staff in the area of parallel computing for intelligent machines include:

The Robot Operating System Expert Scheduler (ROSES) [I.8], which provides a near-optimal schedule on the nodes of a hypercube for a static calculation consisting of fixed-length, precedence-constrained tasks, where it is assumed that communication times are negligible compared to task execution times.

A graphical software tool to aid in static scheduling of precedence-constrained tasks on nodes of a hypercube computer [II.58].

A Virtual-Time Operating System Shell was designed, and several functions implemented for the Vertex operating system which runs on the nodes of the NCUBE computer [II.19, IV.4]. This shell provides a number of functions which facilitate time-crucial responses to unpredictable events within the message-passing environment of the hypercube. These functions include: ordering of messages in an input queue by timestamp; correction messages; cancellation messages; the saving of previous input and output messages and state vectors; and roll-back to previous states.

A robot vision system for hypercubes,which includes a number of innovative fundamental methods and algorithms for exploiting hypercube computers for image processing and analysis applications [II.30, II.31]. The core system consists of a set of utilities for distributed multidimensional array allocation, perimeter communication, and I/O operations to imaging sensors and display devices. A variety of distributed image formats are supported, including all basic storage classes, a wide range of image sizes, several methods of distribution, and conversions between these methods, as well as multi-resolution representations [II.52]. Low-level arithmetical, logical, and neighborhood operations on the arrays are specified in a storage-class independent syntax, and the system utilizes the hypercube with good efficiency. The utilities in the core system have been used for a large number of studies of concurrent algorithms for image analysis and applications in mobile robotics.

A discrete numerical algorithm for linear programming was proposed that has common features with non-linear optimization algorithms and with neural network models [I.35, I.38, II.64, II.65]. The "computation energy" of our algorithm is not quadratic, but linear, which greatly simplifies the model. The collective computational properties are obtained by dividing the neurons in two classes: primal and dual. The algorithm is easy to implement in parallel, easier to code than most classical linear programming algorithms, and can be implemented in VLSI.

3.6 Path Planning and Navigation

A number of algorithms have been proposed, and are still being developed, for the navigation of mobile robots. These algorithms address the problems of autonomous robot navigation in known and unknown, static and dynamic environments, search for a target, exploration of a region, and so on. Early research at CESAR focussed on experimenting with well-known algorithms, first in computer graphic simulations, then with the HERMIES mobile robots [I.1, I.4, I.9, I.10, I.19]. HERMIES-I and -II were controlled by the off-board real-time expert system PICON. For the more recent CESAR robots HERMIES-IIB and HERMIES-III, navigation is controlled by a production system developed on the expert system shell CLIPS, which runs on the nodes of the NCUBE parallel processor, on board the robot. The rules of the production system can read the sensor data and address the robot's effectors; the rules can also call on C-coded navigation procedures. The production system makes high-level, mostly heuristic, decisions and calls on the coded routines to execute the low-level, algorithmic procedures. This architecture was found to be efficient and flexible for the development and testing of the navigation algorithms: the rule base remains small, fast, and easy to understand and maintain; the C-coded routines are also small and easy to maintain. Modules have been developed and implemented for navigating using edge detection and for searching for a target. These modules are now being refined to increase their robustness and their domain of application.

3.7 Machine Learning

Intelligent machines must learn from and adapt to their environment. For environments with little pre-defined structure, learning is a most difficult, yet essential, task. Most information is sensor-based and active sensing must be possible in order to create and continuously update a model of the robot's surroundings. One major issue in this machine learning task is the development of an appropriate representation for the concepts to be learned. Short of discovering a best representation for each context dependent problem, in-depth research needed to be done to understand what can be learned within the rather generic representations that are currently used to perform pattern classification in human-like intelligent systems.

A part of the CESAR effort, has, therefore, concentrated on this learnability issue for a suitably abstracted version of the discrete representations used in three pattern classification approaches -- genetic algorithms (GA), neural nets (NN), and decision trees (DT). To counter the proven difficulty of learning with any generic representation, the recent probabilistic approach proposed by Valiant has been used as the basis of our research. Within this context, we have focussed on both the practical and theoretical issues of algorithmic implementation of this probabilistic-approximation approach using a random set as our generic representation. Ease of learning characterized by: data sample sizes, levels of representational abstraction, use of positive and negative reward, and fast heuristic schemes, have all been studied in depth. Important general conclusions have been drawn from this work to date. For instance, an exhaustive polynomial search algorithm was proven to be much faster than any of the other learning approaches studied. Also, certain levels of representational difficulty used in GA's and NN's have been shown to be unlearnable in Valiant's theory and might, therefore, be considered unjustified. They can take much more time to achieve results, and can diverge. In the area of data sample sizes, we have improved significantly the upper bounds for effective learnability published by Valiant and others [VI.44]. These and other important results are now being used to investigate specific improvements which can be made to GA, NN, and DT learning schemes. In order to practically demonstrate these research results in the CESAR laboratory several computer programs have been written. The algorithms used in these codes are as general as those employed in many important GA, NN, or DT schemes. The results of published sample problems have been used to compare the new approach to such previously implemented algorithms.

Activities in machine learning with direct focus on the CESAR robot HERMIES-IIB began in 1987 with inductive learning, i.e., learning from examples. Since then, the work has expanded to include artificial neural networks and genetic algorithms, and hybrids consisting of neural nets and rule-based systems [II.75].

3.8 Mobile Robot Experiments

Part of the CESAR robotics effort is dedicated to address crucial issues in systems integration so that research results can be integrated into the

HERMIES prototypes. HERMIES-IIB has been used in experiments showing capabilities in world modeling, autonomous navigation in dynamic environments, handling of contingencies, sensor-guided exploration and goal recognition, robot vision, vision-guided manipulation, and innovative problem-solving based on prior experience [I.33, I.43].

In one of the experiments, the outcome of HERMIES-IIB's navigation process is a proper docking in front of a control panel, such that the robot can read the two analog mA-meters, move either of two slide levers, and press any of the four pushbuttons on the front of the panel. The task is to learn a sequence of actions that turn off a danger light on the front of the panel. The problem space the robot learned comprised 81 panel configurations, each defined by a combination of settings of the two meters (high, middle or low) and the two levers (right, middle or left). These 81 panel states can be grouped into sets of categories defined by the response sequence which terminates the danger light on the panel, with each category having a constant set of defining attributes based on a subset of meter and lever settings. Different categorization schemes were created by using different combinations of attribute settings, each coupled with its unique response string. The learning expert system created for this project [I.32, I.43] consists of three major components: a hypothesis-generating unit which permits the robot to hypothesize about possible correct solutions for previously unseen problems; a response-sequence learning unit which learns sequences of responses on problems for which the robot cannot generate a hypothesis; and an inferencing (category formation) component to generalize from examples, enabling the robot to infer categories of problems.

The expert system learns from experience generated by the robot as it experiments with the environment, and generalizes from that experience to infer categories in which it can classify new problem configurations. As currently designed, the system learns a sequence of responses to alter or shut down a control process. However, the general methodology is applicable to any situation in which a robot needs to learn a sequence of motor operations and then make inferences about a classification scheme. A simulation program running on two PCs was developed to train the robot prior to its entering a hostile or critical environment. Experience gained on the simulator system can be transferred to the robot prior to starting the experiment. This training, coupled with the robot's ability to solve novel tasks by generalizing, eliminates the need to pre-program all possible real world situations. We showed that by presenting a selected set of panel configurations during a training session, the robot can develop the capacity to handle a wide variety of unanticipated panel configurations, making a minimum number of errors. After training on the set of selected panels and inferring categories, the expert system makes no further errors on new panel problems from the same classification scheme. The combination of a robot capable of learning and a simulation system to provide rapid and efficient training seems to be a viable way of creating a robot prepared to cope with a dynamic environment whose characteristics cannot be completely known in advance. The inability to precisely know the robot's environment precludes programming the ability to handle all problems in advance.

The HERMIES robots also serve as testbeds for results from a team effort among four universities (Florida, Michigan, Tennessee, Texas), industry (Odetics, Remotec, Telerobotics International), and ORNL under the DOE/NE Robotics for Advanced Reactors Program [II.38, II.59]. In this program, annual experiments with robot system modules developed by the four university groups and by ORNL are performed in the CESAR laboratory. These demonstrations derived from needs in an applied research program represent results of significant efforts in systems integration supported by the applied program. They show the benefits that result from incorporating accomplishments achieved within the basic research effort supported by BES.

In the latest experimental setup, HERMIES-III was used to demonstrate the capability of robotic clean-up of simulated chemical spills. The robot was given the approximate location of a chemical spill and the instruction to clean up the spill. Based on only partially complete geometric information about the robot's environment, HERMIES-III proceeded to plan a path and navigate, avoiding, obstacles along the way, to the presumed location of the spill. The machine then determined with its vision system the extent and precise location of the spill, and calculated and executed a motion pattern for the CESARm that would result in removal of the spill with a vacuum attachment to the end-effector. This experiment also incorporates results from the collaborative effort among CESAR and four universities which were mentioned previously. Although all operations are performed by the robot autonomously, a human operator can interact with the robot at all times during the experiment through a force-reflecting joystick device with visual information feedback through a Silicon Graphics IRIS 4D workstation. The experiment is made possible by integration of a number of system modules and associated processes that execute simultaneously and asynchronously on a network of several computer systems: Motorola 68020s running OS/9, a DEC microvax running VMS, and a Silicon Graphics IRIS 4D running UNIX.

A new shared memory emulation communications system for such a heterogeneous network was developed. User processes are interfaced to a local shared memory, a replicated copy of which resides on each machine in the network. Sending processes on each machine receive interrupts from user processes when data changes, translate the data into a common format, and broadcast the translated data to the other machines in the network. Receiving processes on each machine read the data from the network, translate the data into local format, and store the results in the local shared memory.

The current experimental program at CESAR involves performance of human-scale tasks, integration of the dexterous manipulator and platform motion in geometrically complex environments, and effective use of multiple cooperating robots. The ongoing research program supports the continued development of autonomous capability for HERMIES-IIB and -III to perform complex navigation and manipulation under time constraints, while dealing with imprecise sensory information.

4. Proposed Research FY 1992 - 1996

This section contains details of our proposed research over the next five years in cooperating robots, mobile manipulator systems, multi-sensor data analysis and fusion, machine learning, and embedded high performance computing. The section concludes with plans for continued experimental facilities for intelligent machine research.

4.1 Cooperating Robots

One of the frontiers in intelligent machine research is the understanding of how constructive cooperation among multiple performing agents can be effected. This research area poses a number of fundamental problems in many robot component technologies, and successful solutions to these problems will have great impact on safety and productivity of operations in a number of applications including flexible manufacturing, nuclear energy facilities, environmental restoration and waste management, future intelligent transportation systems, as well as space applications. Since there is a broad spectrum of fundamental problems in this area, ranging from high-level planning to low-level control, and to important architectural and systems issues, we propose to focus our effort on two problem areas: (1) cooperation by multiple mobile robots in dynamic, incompletely known environments; and (2) cooperating robotic manipulators. In addition, we delineate the scope of the proposed research by including relevant experimental scenarios for proof-of-principle demonstrations derived from application areas currently being addressed or to be addressed in the near future. These include, but are not limited to, environmental restoration and waste management, other nuclear environments hazardous to humans, and space applications for robotics and intelligent systems.

4.1.1 Cooperating Mobile Robots

The purpose of this work is to advance the understanding of how the control of multiple mobile robots can be achieved so that classes of tasks requiring path planning and navigation can be solved cooperatively by all robots involved. Research in mobile robot motion planning has been addressing primarily problems associated with a single mobile robot in work spaces assumed to be relatively benign, i.e., structured and completely known or containing a small number of imprecisely known obstacles (Schwartz and Sharir 1988). In an environment involving multiple mobile robots, the motions of other robots are potentially unknown or unpredictable. The environment is also changing with time, which represents another issue which has not been dealt with sufficiently in the literature. Even for a single robot, as of today, there does not exist an algorithm that can handle an arbitrary number of moving obstacles with complex arbitrary motion patterns. Nevertheless, a number of important results concerning computational complexity (Reif and Sharir 1985; Canny and Reif 1987; Canny 1988), solutions to particular, relatively simple instantiations of the problem (Sutner and Maass 1988; O'Dunlaing 1987), and heuristic methods (Kant and Zucker 1987; Parker 1988; Fujimura and Samet 1989; Lamadrid and Gini 1990; Shih, Lee, and Gruber 1990) have been reported.

4.1.1.1 Background

The scenario to be addressed involves a situation in which multiple mobile robots pursue tasks in a common workspace that can change with time, and their motions have to be planned so that collisions are avoided. Several boundary conditions may affect possible solutions. Among these are whether all robots have identical perception, reasoning and mobility capabilities or not; whether locally reactive behaviors are possible; whether planning is performed globally or locally, i.e., by every robot individually; and whether communications among the robots is possible to implement "negotiations".

Erdmann and Lozano-Perez (1987) assign priorities globally to all robots; first the motion for the robot with the highest priority is planned. Then, the motion for the robot with the second highest priority is determined regarding the first robot as the sole moving obstacle in the environment, and so on. This approach assumes that a satisfactory solution to the problem of moving obstacle avoidance is available. Warren (1990) also adopts this global priority approach, but uses a different obstacle avoidance algorithm. Shih, Lee, and Gruver (1990) also consider the problem of coordinating multiple robots by assigning global priorities. Buckley (1989) uses the priority scheme, but assigns priorities locally, i.e., priorities are assigned to two robots only when their trajectories may potentially intersect. MacNish and Fallside (1990) also use local priority. They use modal temporal logic to detect possible collisions between two mobile robots.

One way of coordinating multiple mobile robots is to specify an arbitrator robot that is responsible for determining motions for all robots (Yap 1984; Freund and Hoyer 1988). Cammarata et al. (1983) consider several schemes to achieve coordination by means of centralization applied to distributed aircraft. Our preliminary studies show that for the case of two mobile robots, additional problems arise when the agents are not assumed to have complete knowledge about their respective plans or capabilities (Fujimura 1991).

Other methods make use of mutually known conflict resolving mechanisms among homogeneous agents (Tournassoud 1986; Saito and Tsumura 1989; Lee and Bien 1990). These methods do not consider negotiation between agents.

Robots moving along planned paths, may encounter obstacles (moving or stationary). Replanning on the fly needs to occur. Solutions depend heavily on the sensory capabilities of the robots. A considerable time for sensory data analysis may be spent to determine obstacle sizes, locations, and motions. Time for planning can be small in comparison.

Fast locally reactive replanning for obstacle avoidance based on sonar sensor data has been described by Borenstein and Koren (1989), for example. Fast obstacle detection based on inverse perspective mappings applied to digitized CCD camera images was reported by Mallot et al (1989). These methods work in real-time (i.e., at sensor frame rates), however, they do not allow analyses of the sensor data that would allow any predictive estimation of the perceived environment, and hence limit the planning options. Although very efficient at particular instants in time when obstacles are detected, these methods fail to produce solutions that can approach globally optimum paths or prevent the robot from becoming trapped in certain, not necessarily pathological obstacle configurations.

Basu (1990) considers the problem of visual navigation among moving obstacles under uncertainty. His formulation may be used for proper evaluation of navigable areas in the environment. Using probabilistic uncertainty models, the problem is to determine a path with the smallest probability of collision with obstacles.

Path planning under incomplete information by a robot equipped with a tactile and a visual sensor has been addressed (Lumelsky and Stepanov,1987; Lumelsky and Skewis,1990). Only stationary obstacles are considered. Extensions of this approach to deal with moving obstacles may be possible.

Regarding task planning for multiple robots, much research has been reported from an artificial intelligence point of view (e.g., Sacerdoti, 1975). Durfee and Lesser (1989) address the issue of decomposing tasks among distributed problem solving agents. They take the partial global planning approach to coordinate agents that need to negotiate for task planning.

In robotics, Roach and Boaz (1987) consider the problem of coordinating two robot arms in the same work space. The key issue is to plan the actions of the robot arms so that they are collision-free. Nagata, Honda, and Teramoto (1988) report on a robot plan generation system for multiple robot domains. Their system consists of two parts: plan generation for assembly tasks and collision detection of two cylindrical robot arms.

Grossman (1988) addresses the control of traffic for a large number of autonomous vehicles which move in a network of roads. He studies a simple grid-like network and proposes a control policy that works well by simulation.

4.1.1.2 Proposed Research

Building and expanding on results summarized in the preceding section and on previous work at CESAR, we propose two topics for basic research in order to realize efficient operations by multiple mobile robots: (1) conflict resolution among mobile agents; (2) cooperative task planning by mobile agents. The first topic is also refered to as weak or negative cooperation, i.e., agents trying to avoid each other (Werner, 1989; Tournassoud, 1986; Saito and Tsumura, 1989; MacNish and Fallside, 1990), whereas the second topic encompasses positively cooperative behavior (Chen et al, 1987; Coupez et al, 1989; Roach and Boaz, 1987).

Initial efforts in the area of weak cooperation will focus on the case of two mobile robots with different mobility characteristics and sensory capabilities, e.g. HERMIES-IIB and HERMIES-III, required to move in a common workspace to their respective destinations in minimal time, in the presence of moving obstacles. The investigations will include scenarios in which robot-robot communication may or may not be possible, and a decentralized control approach will be followed. The studies will assess solid analytical foundations for solutions where possible. Since our work focuses on autonomous agents in unstructured environments, we will place particular emphasis on issues involving the sensory capabilities of the robots, and on experimental verifications of the reliability of weak cooperation schemes. Specific sensors available with the robots include sonar transducers, CCD cameras, and a laser range camera.

For scenarios without direct robot-robot communications, the problem can be viewed as being equivalent to solving navigation problems in dynamic environments for each of the robots involved, taking into account that situations amounting to complete blocking for one or several of the robots may occur. We will investigate crucial trade-offs between the speed of motion and the ability to analyze sensor data to the extent required to support that motion, and the resulting level of capability for weak cooperation among all agents involved.

As in other areas of research covered in this proposal, we will focus our limited resources in selected areas. In this case we will focus on positive cooperation for selected classes of problems (Durfee et al, 1987; Georgeff, 1983; Wesson et al, 1981). The work will initially address the problem of site characterization with multiple robots, e.g. for surveillance or environmental monitoring purposes. The robots are assumed to have different mobility and sensory characteristics, and to communicate with each other. Depending on the amount of a priori information regarding the site to be characterized, a combination of centralized and local planning will be required for this work. Execution of these plans will have to be monitored and eventually modified depending on the information generated by the agents involved.

These investigations will be pursued in close cooperation with the activities described in sections 4.3 and 4.5 of this proposal. They will involve rigorous theoretical analyses, simulations using the CESAR graphics modeling facilities, and experimental evaluations using the HERMIES robots.

4.1.1.3 References

A. Basu, Model-based navigation, PhD Thesis, Computer Science Department, University of Maryland, 1990.

Borenstein, J. and Y. Koren, "Real-time obstacle avoidance for fast autonomous and semi-autonomous mobile robots," Proc. 3rd ANS Topical Meeting on Robotics and Remote Systems, Charleston SC, 1989

S. J. Buckley, Fast motion planning for multiple moving robots, Proceedings of the IEEE International Conference on Robotics and Automation, Scottsdale, AZ, May 1989, 322--326.

S. Cammarata, D. McArthur, and R. Steeb, Strategies of cooperation in distributed problem solving, Proceedings of International Joint Conference on Artificial Intelligence, 767-770, 1983.

J. Canny, Some algebraic and geometric computations in PSPACE, Proceedings of the 20th Annual ACM Symposium on Theory of Computing, Chicago, IL, 460-467, 1988.

J. Canny and J. Reif, New lower bound techniques for robot motion planning problems, Proceedings of the 27th Annual IEEE Symposium on the Foundations of Computer Science, Los Angeles, CA, October 1987, 49-60.

C. L. Chen, C. S. G. Lee, and C. D. McGillem, Task assignment and load balancing of autonomous vehicles in a flexible manufacturing system, IEEE Transactions on Robotics and Automation 3, 6(December 1987), 659-671.

D. Coupez, A. Delchambre, and P. Gaspart, The supervision and management of a two robots flexible assembly cell, Proceedings of IEEE International Conference on Robotics and Automation, Scottsdale, AZ, May 1989, 540-550.

T. Dean and M. Boddy, An analysis of time-dependent planning, Proceedings of AAAI, 1987, 49-54.

E. H. Durfee, V. R. Lesser, and D. D. Corkill, Coherent cooperation among communicating problem solvers, IEEE Transactions on Computers 36, 11(November 1987), 1275-1291.

E. H. Durfee and V. R. Lesser, Negotiating task decomposition and allocation using partial global planning, in Distributed Artificial Intelligence Vol II (Ed. L. Gasser and M. H. Huhns), (1989), 229-243.

E. Davis, Representations of Commonsense Knowledge, Morgan Kaufmann, San Meteo (Calif.), 1990.

A. Elfes, Using occupancy grids for mobile robot perception and navigation, Computer 22, 6(June 1989), 46-57.

M. Erdmann and T. Lozano-Perez, On multiple moving objects, Algorithmica 2, 4(1987), 477-522.

A. Foisy, V. Hayward, and S. Aubry, The use of awareness in collision prediction, Proceedings of IEEE International Conference on Robotics and Automation, May 1990, Cincinnati, OH, 338-343.

K. Fujimura and H. Samet, A hierarchical strategy for path planning among moving obstacles, IEEE Transactions on Robotics and Automation 5, (February 1989), 61-69.

K. Fujimura, Motion planning in dynamic domains, CS-TR-2377, Computer Science Department, University of Maryland, 1989.

K. Fujimura, A reactive model for multiple mobile agents, to appear in IEEE International Conference on Robotics and Automation, Sacramento, CA, April 1991.

M. P. Georgeff, Communication and interaction in multi-agent planning, Proceedings of AAAI, 1983, 125-129.

M. P. Georgeff, A theory of action for multi-agent planning, Proceedings of AAAI, 1984, 121-125.

M. P. Georgeff, A. L. Lansky, and M. J. Schoppers, Reasoning and planning in dynamic domains: An experiment with a mobile robot, SRI International, Technical Note 380, 1987.

D. D. Grossman, Traffic control of multiple robot vehicles, IEEE Journal of Robotics and Automation 4, 5 (1988), 491-497.

J. Halpern, Reasoning about knowledge: An overview, Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge, Monterey, CA, 1986, 1-17.

K. Kant and S. W. Zucker, Toward efficient planning: The path-velocity decomposition, International Journal of Robotics Research 5, (3) (Fall 1986), 72-89.

R. Korf, Real-time A* search, Artificial Intelligence 42, (2-3) (March 1990), 189-211.

J. de Lamadrid and M. Gini, Path tracking through uncharted moving obstacles, IEEE Transactions on Systems, Man, and Cybernetics 20, 4 (November/December 1990), 1408-1422.

J. Lee and Z. Bien, Collision-free trajectory control for multiple robots based on neural optimization network, Robotica 8, 3 (July-September 1990), 185-194.

T. Lozano-Perez and M.A. Wesley, An algorithm for planning collision free paths among polyhedral obstacles, Communications of the ACM 22, (10) (October 1979), 560-570.

V. J. Lumelsky and A. A. Stepanov, Path-planning strategies for a point mobile automaton moving amidst unknown obstacles of arbitrary shape, Algorithmica 2, 4(1987), 403-430.

V. J. Lumelsky and T. Skewis, A paradigm for incorporating vision in the robot navigation function, Proceedings of 1988 IEEE International Conference on Robotics and Automation, April 1988, Philadelphia, PA, 734-739.

C. MacNish and F. Fallside, Temporal reasoning: A solution for multiple agent collision avoidance, Proceedings of 1990 IEEE International Conference on Robotics and Automation, May 1990, Cincinnati, OH, 494-499.

Mallot, H. A., E. Schulze and K. Storjohann, "Neural network strategies for robot navigation," in L. Persinnaz and G. Dreyfus (eds.), Neural Networks from Models to Applications, Paris: I.D.S.E.T., pp. 560-569, 1989

J. S. B. Mitchell, D. M. Mount, C. H. Papademitriou, The discrete geodesic problem, SIAM Journal on Computing 16, 647-668.

T. Nagata, K. Honda, and Y. Teramoto, Multirobot plan generation in a continuous domain: Planning by use of plan graph and avoiding collisions among robots, IEEE Transactions on Robotics and Automation 4, 1(February 1988), 2-13.

C. O'Dunlaing, Motion planning with inertial constraints, Algorithmica 2, 4(1987), 431-476.

C. L. Pape, A combination of centralized and distributed methods for multi-agent planning and scheduling, Proceedings of 1990 IEEE International Conference on Robotics and Automation, May 1990, Cincinnati, OH, 488-493.

L. Parker, A robot navigation algorithm for moving obstacles, MS thesis, University of Tennessee, Knoxville (Tenn.), 1986.

J. Reif and M. Sharir, Motion planning in the presence of moving obstacles, Proceedings of the 26th Annual IEEE Symposium on Foundations of Computer Science, Portland, OR, October 1985, 144-154.

J. W. Roach and M. N. Boaz, Coordinating the motions of robot arms in a common workspace, IEEE Journal of Robotics and Automation 3, 5(October 1987), 437-444.

N. C. Rowe and R. F. Richbourg, An efficient Snell's law method for optimal-path planning across multiple two-dimensional, irregular, homogeneous-cost regions, International Journal of Robotics Research 9, (December 1990), 48-66.

E. D. Sacerdoti, The nonlinear nature of plans, Proceedings of IJCAI, (1975), 206-214.

M. Saito and T. Tsumura, Collision avoidance between mobile robots, Proceedings of IEEE/RSJ International Workshop on Intelligent Robots and Systems, September 1989, Tsukuba, Japan, 473-478.

J. T. Schwartz and M. Sharir, A survey of motion planning and related geometric algorithms, in Geometric Reasoning, (Ed., D. Kapur and J. L. Mundy), Elsevier Science, Amsterdam, The Netherlands, 1988.

C. Shih, T. Lee, and W. A. Gruver, A unified approach for robot motion planning with moving polyhedral obstacles, IEEE Transactions on Systems, Man, and Cybernetics 20, 903-915.

K. Sutner and W. Maass, Motion planning among time dependent obstacles, Acta Informatica 26, (1988), 93-122.

P. Tournassoud, A strategy for obstacle avoidance and its application to multi-robot systems, Proceedings of IEEE International Conference on Robotics and Automation, San Francisco, CA, April 1986, 1224-1229.

C. W. Warren, Multiple robot path coordination using artificial potential fields, Proceedings of IEEE International Conference on Robotics and Automation, Cincinnati, May 1990, 500-505.

C. R. Weisbin, G. de Saussure, J. R. Einstein, F. G. Pin, and E. Heer, Autonomous mobile robot navigation and learning, Computer 22, 6(June 1989), 29-35.

E. Werner, Cooperating agents: A unified theory of communication and social structure, in Distributed Artificial Intelligence, (Ed., L. Gasser and M.N. Huhns), Morgan Kaufmann, San Meteo, CA, 3-36.

R. Wesson, F. Hayes-Roth, J. W. Burge, C. Stasz, and C. A. Sunshine, Network structures for distributed situation assessment, IEEE Transactions on Systems, Man and Cybernetics 11, 1(1981), 5-23.

C. Yap, Coordinating the motion of several discs, TR-105, Computer Science Department, Courant Institute, New York University, February, 1984.

C. Yap, Algorithmic motion planning, in Advances in Robotics Vol. 1, (Ed., J. T. Schwartz and C. Yap), Lawrence Erlbaum, Hillsdale NJ, 1986, Chapter 3.

T. Yeung and G. A. Bekey, A decentralized approach to the motion planning problem for multiple robots, Proceedings of IEEE International Conference on Robotics and Automation, Raleigh, GA, March 1987, 1779-1784.

4.1.2 Cooperating Robot Manipulators

The second main emphasis of our research in cooperating multiple robots is on fundamental problems in cooperating manipulators. Advancements of our basic understanding of how to accomplish efficient multi-arm manipulation will represent benefits for many application areas for advanced robots in unstructured environments.

For example, two manipulators can directly assemble two parts into a rigid product where each part is held by a manipulator. Expensive custom-made fixturing is often required for a single manipulator to perform the same task. Indeed, the fixture and the single robot may be viewed as a pair of hands, where the former is passive and the latter active. Such fixturing is typically not programmable and cannot be reused when the task or production requirements change which frequently occurs in a batch manufacturing environment. Furthermore, such fixturing would be impractical or possibly unavailable in unstructured hazardous environments such as space and for one-of-a-kind tasks where the manipulators themselves would most likely be mobile. With two cooperating manipulators, the fixturing can be reduced or eliminated and the system can be made adaptive to changes in the environment or in the assembly requirements.

Two cooperating manipulators can perform tasks which a single manipulator is not capable of executing, such as transporting an object beyond the load capacity of a single manipulator, transporting an object of odd or unusual geometry such as a long board or automobile axle, and transporting an object beyond the work envelope of a single manipulator (e.g., by directly exchanging it between the robots during its transport).

4.1.2.1 Background

The coordinated motions of two manipulators moving simultaneously in a common workspace are divided into two groups in (Zheng, 1987). The first group is called loosely coordinated motion where the manipulators execute independent work sequences, i.e., they do not come into mutual contact and the main concern is avoiding collisions between the two moving manipulators (Freund and Hoyer, 1988; Roach and Boaz 1987, O'Donnell and Lozano-Perez 1989, Gupta 1990). The second group is called tightly coordinated motion

(Zheng, 1987) in which the motions of the manipulators are interdependent. If one manipulator fails, the task cannot be executed at all; for example, when two manipulators mutually lift and transport a single object. The problems associated with the tightly coordinated motion of two manipulators are very challenging to say the least. This is primarily due to the fact that when two manipulators hold a single object, the overall system forms a single closed chain mechanism and a loss of degrees of freedom occurs (Jamshidi, 1989; Koivo and Bekey, 1988). Strong kinematic and dynamic interactions between the robots result in a constrained system motion. The kinematic coupling causes the positions and orientations of the end effectors to become constrained. Dynamic coupling occurs because the generalized contact forces imparted to the shared load by the manipulators are interacting.

Several control strategies originally developed for the constrained motion of a single manipulator have recently been extended to the case of two manipulators holding a common rigid object. These include, for example, active compliance control (Hayati et al, 1989; Tao et al, 1990; Kim and Zheng, 1989), impedance control (Schneider and Cannon, 1989), damping control (Alberts and Soloways, 1988), operational space control (Khatib, 1987), and hybrid control (Dauchez et al, 1989; Kopf, 1989). However, these approaches do not take into account properly the constrained dynamics of the closed-chain system in the controller design. Indeed, a recent workshop on coordinated multiple manipulators (Koivo and Bekey, 1988) identified the problems of developing constrained dynamical models including reduced order modeling and model-based control methodologies for two or more manipulators holding a common, possibly jointed object as high priority research issues. A comprehensive bibliography of the related literature in this area can be found in (Kreutz and Lokshin, 1988).

The most significant works in this area include (Kreutz and Lokshin, 1988), where the general dynamical equations for several manipulators grasping a common rigid object are presented in several joint and operational space forms. A major contribution of (Kreutz and Lokshin, 1988) is the proof that the number of lost arm configuration DOF due to imposing the closed chain kinematic constraints is equal to the number of DOF gained for controlling the internal stress generalized contact forces exerted on the object by the manipulators which do not induce motion in the payload. In (Luh and Zheng, 1987), kinematic constraints and a reduced order dynamical model having a leader/follower structure are developed for two six-axis manipulators holding three types of objects: a rigid body, a pair of pliers, and a part containing a spherical joint. However, the problems of quantifying the dynamic coupling effects and controlling the generalized contact forces are not discussed. Using a singular systems model for two manipulators holding a rigid object, a suboptimal relation between the joint torques applied by each manipulator is obtained by minimizing a quadratic torque cost in (Carignan and Akin, 1989). The same system is controlled using a parameter adaptive controller in (Carignan, 1990). These works are restricted to the planar case. A model-based hybrid control scheme for the case when a rigid object held by multiple manipulators is also in contact with a rigid environment is discussed in (Yoshikawa and Zheng, 1990). The control of the position and orientation of the object as done in (Yoshikawa et al, 1990) does not utilize all the system DOF available when one or more of the manipulators is kinematically redundant. For the present proposal, recent work at CESAR (Unseren, 1991a,b) on dynamical modeling and control of two manipulators holding two types of objects: (i) a rigid body and (ii) a spherically jointed object will serve as a basis for further study. The final model consists of reduced order equations of motion containing no generalized contact forces, which are calculated separately. Motivated by previous work on control of the hard contact motion of a single manipulator (Kankaanranta and Koivo, 1988; McClamroch and Wang, 1988), a control architecture is suggested which completely decouples the position- and force-controlled DOF in the multi-manipulator system. The previous work on dynamically distributing a load held by two manipulators and controlling the internal stress generalized contact forces (Zheng and Luh, 1988; Orin and Oh, 1981; Walker et al, 1989) focused on the rigid payload case and did not take into account explicitly the kinematic constraints. Our most recent work at CESAR (Unseren, 1991c) suggested a different load distribution technique and derived a rigid body model and control architecture for two manipulators holding (i) a spherically jointed object and (ii) a revolute-jointed object so that the position- and internal stress force-controlled DOF are decoupled.

Additional analytical work remains with the decoupled control architecture method (Unseren 1991a,b,c) including mathematically demonstrating its stability. It should be determined whether this modeling and control framework can be extended to the case where the mutually held payload is also in contact with a rigid structure in the workspace. Furthermore, the approach has not been simulated to determine its effectiveness and robustness to errors caused by modeling inaccuracies. The performance of this control method needs to be compared to those of (Carignan and Akin, 1989 ; Yoshikawa and Zheng, 1990; Lokshin and Kreutz, 1988).

The model-based control schemes proposed in the literature for two manipulators holding a common object all assume that the control designer has perfect knowledge of the nonlinear dynamical terms in the model including the kinematic constraints. They further assume that there are no unexpected force disturbances in the system. (Hayati et al, 1989) warns, however that such control schemes are difficult to implement because there are bounds on the accuracies of the estimated dynamical model terms appearing in the control laws. Future work is needed to investigate the performance of these control schemes when these assumptions are relaxed. Furthermore, there has been only one paper on the problem of adaptively estimating the dynamical terms appearing in the control (Carignan, 1990) and clearly additional research is needed.

The fact that the distribution of a commonly held dynamic load between two manipulators is not uniquely determined makes it a difficult problem which is still in its infancy. Most of the work done in this area has been restricted to mathematical modeling. There is a crucial need to gain additional insight into the behavior and control of the internal stress generalized contact forces, which are quantified by the load distribution. This can be accomplished through extensive dynamical computer simulations of the various methods proposed in the literature, followed by obtaining and analyzing experimental results.

A survey of the literature reveals that there is very little understood about dual-manipulator closed chains when one or both of the manipulators is kinematically redundant (Kreutz and Lokshin, 1988). In fact, most of the dynamical models and control methods proposed are restricted to the case of two six-axis manipulators. The significance and physical interpretation of the the so-called redundant DOF is a crucially important conceptual research area because it may shed light on the advantages of applying two cooperating kinematically redundant manipulators to the transport of an object. It should be mentioned that the literature is also very scarce on the problem of modeling and controlling the constrained motion of a single kinematically redundant manipulator.

It is well known that force control is needed in addition to position control to accomplish successfully tasks involving the constrained motion of a single manipulator, e.g., a manipulator turning a crank, as well as tasks involving the tightly coordinated motions of two interacting manipulators. However, as observed in (Koivo and Bekey, 1988), it is not evident whether the fundamental questions for two interacting manipulators are basically the same as those for a single manipulator, or whether fundamentally new questions arise from the need to coordinate two manipulators. It is felt that establishing a solid experimental base on controlling single and multiple constrained manipulators will shed light on the similarities and differences.

There is abundant literature on two manipulators holding a rigid object, but research on the case where the common load possesses a joint is in its infancy (Luh and Zheng, 1987; Unseren, 1991b,c). With a jointed object, the overall system becomes more complicated and has more DOF when compared to the rigid payload case. Further work is needed to determine if the issues associated with controlling a rigid payload held by two manipulators can be directly extended to jointed loads.

4.1.2.2 Proposed Research

Topic 1.

We propose to perform experimental research on the simultaneous position- and force-control of a constrained single manipulator as a pre-requisite to the position/force control of two interacting manipulators. The seven-axis kinematically redundant CESARm research manipulator will be used to perform tasks such as "peg-in-hole" and "turning-a-crank". The control problem will be solved by carrying through a progression of related stages:

(i) Soft contact control of CESARm in a non-redundant configuration.

(ii) Soft contact control in redundant configuration.

(iii) Hard contact control in non-redundant configuration.

(iv) Hard contact control in redundant configuration.

CESARm can be configured to function as a non-redundant manipulator. Furthermore, it can be set up to operate as a three link planar redundant or non-redundant manipulator. CESARm is equipped with a wrist force/torque sensor and motor encoders to provide force- and position-feedback, respectively.

In the soft contact case, the manipulator remains an open chain mechanism and its joint positions comprise an independent set of generalized coordinates. In the hard contact case, the manipulator and the constraining environment, e.g., a crank it is turning, are viewed as a single closed chain mechanism and the joint positions are interrelated. Intuitively the former case seems easier to implement than the latter; thus, soft contact control will be considered first.

By controlling CESARm in both non-redundant and redundant configurations, the designer will gain insight into the advantages of introducing the kinematic redundancy into robot force control applications.

It should be mentioned that most kinematic redundancy resolution schemes for the unconstrained motion of manipulators are implemented in an open-loop manner to the plant, i.e., the measured joint positions are not fed back to the kinematic optimization scheme. The feasibility of including the redundancy resolution scheme within the closed loop control system will be investigated in the experiments.

It has recently been suggested that there is a dynamic load distribution problem when a single manipulator performs a crank turning task (Unseren 1991d). This proposed load distribution scheme quantifies the internal stress component of the generalized contact forces acting on the crank which are to be controlled in addition to controlling its position. The above described laboratory experiments will include implementing the suggested load distribution scheme. It is envisioned that the manipulator and/or the crank itself may be equipped with a force/torque sensor.

It should be mentioned that the results obtained here may be applicable to the control of the constrained motion of a mobile platform consisting of a pair of steerable, interconnected wheels. This work will involve collaboration with those conducting research on mobile, wheeled platforms, e.g., see (Reister, 1991).

Topic 2.
Dynamic computer simulations of the decoupled control architecture method recently proposed at CESAR (Unseren, 1991a,b,c) for multi-manipulator closed chains will be conducted to investigate its performance and analyze the closed loop system behavior and characteristics when the described controller is applied to a specific configuration. In theory, the position- and internal stress force-controlled DOF are completely decoupled using the proposed controller, but it is assumed the designer has perfect knowledge of the nonlinear dynamical terms in the system model. In practice, there will always be errors associated with estimating the nonlinear terms; thus, they will not be perfectly cancelled by the theoretic control architecture. Simulation is an efficient tool for developing control methods and the effects of errors on controller performance will appear in a properly simulated result. Furthermore, it is felt that considerable insight into the behavior and control of the internal stress contact forces in the closed chain can be obtained by computer simulation. This includes determining desired reference trajectories for the internal stress contact forces.

Typical applications to be investigated by simulation include two structurally dissimilar manipulators holding both rigid and jointed objects. The latter type payload is of particular interest since to our knowledge no one has reported simulation results on such a system.

The performance of the decoupled control architecture method will be compared to that of other control methods proposed in the literature through simulation studies.

Topic 3.

Investigate and assess the physical interpretation of the redundant DOF which arise in dual-manipulator closed chains when one or both of the robots is kinematically redundant. This problem will be addressed by focusing on a specific application, e.g., two manipulators containing six- and seven-axes respectively, holding a rigid object so that there is a single redundant DOF. A configuration space variable needs to identified which characterizes the redundant DOF. Such a variable may be found analytically by considering the following reasoning. If six of the DOF are used to control the Cartesian position and orientation of the held object at its center of mass, the seventh DOF should not contribute to the motion of the object. Instead, it will be determined by mathematical modeling if the redundant DOF can influence and compensate for the mechanical coupling in the closed chain. The significance of the additional equation of motion appearing in the reduced order model (e.g., see (Unseren, 1991 a,b,c)) which describes the dynamic response of the redundant configuration variable will be determined by mathematical modeling. The benefits of introducing the redundant DOF will be demonstrated by dynamic computer simulations.

Topic 4.

Laboratory experiments on controlling two manipulators lifting and transporting a common object. A suggested experiment involves controlling two manipulators to move from their parking positions to hold opposite ends of a long rigid board. The manipulators would lift the board and transport it to another location. Upon releasing the board at the dropoff point, the robots would return to their parking positions. The experiment would be repeated with the board replaced by an object containing a joint.

Each robot would be equipped with a wrist force/torque sensor. The methods developed in topic 1 for controlling a single constrained manipulator would be extended to the multi-manipulator case. Particular emphasis will be placed on identifying and solving the new issues which arise with two interacting manipulators. It is envisioned that the experiments will focus on controlling the position and orientation of the shared load as well as the internal stress generalized contact forces.

Paradigms

The results obtained from the above described research can be applied to a variety of constrained motion tasks requiring one or two robotic manipulators including:

(i) Material handling tasks involving the transport of large, heavy objects or objects of unusual geometry using two manipulators. Applications for material handling with two manipulators occur in industrial manufacturing and in outer space.

(ii) Actuating a valve, lever, or crank-type mechanism. Such tasks frequently occur in the maintenance and servicing of mechanical structures contained in a nuclear power plant.

4.1.2.3 References

Y.F. Zheng,
"Collision Effects on Two Coordinating Robots in Assembly and the Effect Minimization," IEEE Trans. Systems, Man, and Cybernetics, Jan./Feb. 1987, vol. SMC-17, no. 1, pp. 8-12.

E. Freund and H. Hoyer,
"Real-Time Pathfinding in Multi-Robots Systems Including Obstacle Avoidance," Int. J. Robotics Research, Feb. 1988, vol 7, no. 1, pp 42-70.

J.W. Roach and M.N. Boaz,
"Coordinating the Motions of Robot Arms in a Common Workspace," IEEE J. Robotics and Automation, Oct. 1987, vol. RA-3, No. 5, pp. 437-444.

P.A. O'Donnell and T. Lozano-Perez,
"Deadlock-Free and Collision-Free Coordination of Two Robot Manipulators," IEEE Int. Conf. Robotics and Automation, Scottsdale, AZ, USA, May 1989, vol. 1, pp 484-489.

K.K. Gupta,
"Fast Collision Avoidance for Manipulator Arms: A Sequential Search Strategy," IEEE Trans. Robotics and Automation, vol. 6, no. 5, Oct. 1990, pp 522-532.

M. Jamshidi,
"Multi-Arm Robots," Guest Editorial, Robotics and Autonomous Systems, " Dec. 1989, vol. 5, no. 4, pp. 305-306.

A.J. Koivo and G.A. Bekey,
"Report of Workshop on Coordinated Multiple Robot Manipulators: Planning, Control, and Applications," IEEE J. Robotics and Automation, February 1988, Vol. 4, No. 1, pp. 91-93.

S.A. Hayati, K.S. Tao, and T.S. Lee,
"Dual Arm Coordination and Control," Robotics and Autonomous Systems, Dec. 1989, Vol. 5, No. 4, pp 333-344.

J.M. Tao, J.Y.S. Luh, and Y.F. Zheng,
"Compliant Coordination Control of Two Moving Industrial Robots," IEEE Trans. Robotics and Automation, June 1990, vol. 6, no. 3, pp. 322-330.

K.I. Kim and Y.F. Zheng,
"Two Strategies of Position and Force Control for Two Industrial Robots Holding a Single Object," Robotics and Autonomous Systems, Dec.1989, vol. 5, no. 4, pp. 395-403.

S.A. Schneider and R.H. Cannon, Jr.,
"Object Impedance Control for Cooperative Manipulation: Theory and Experimental Results," IEEE Int. Conf. Robotics and Automation, Scottsdale, AR, USA, May 1989, Vol. 2, pp 1076-1083.

T.E. Alberts and D.I. Soloway,
"Force Control in a Multi-Arm Robot System,"
IEEE Int. Conf. Robotics and Automation, Philadelphia, PA, USA, April 1988, Vol. 3, pp1490-1496.

O. Khatib,
"Object Manipulation in a Multi-Effector Robot System," 4th Int. Symp. Robotics Research, University of California at Santa Cruz, 1987, MIT Press Series on Artificial Intelligence, edited by R. Bolles and B. Roth.

P. Dauchez, A. Fournier, and R. Jourdan,
"Hybrid Control of a Two-Arm Robot for Complex Tasks," Robotics and Autonomous Systems, Dec. 1989, Vol. 5, No. 4, pp 323-332.

C.D. Kopf,
"Dynamic Two Arm Hybrid Position/Force Control," Robotics and Autonomous Systems, Dec. 1989, Vol. 5, No. 4, pp 369-376.

K. Kreutz and A. Lokshin,
"Load Balancing and Closed Chain Multiple Arm Control," American Control Conf., Atlanta, GA, USA, June 1988, vol. 3, pp 2148-2155.

J.Y.S. Luh and Y.F. Zheng,
"Constrained Relations Between Two Coordinated Industrial Robots For Motion Control," Int. J. Robotics Research, vol. 6, no. 3, pp 60-70, Fall 1987.

C.R. Carignan and D.L. Akin,
"Optimal Force Distribution for Payload Using a Planar Dual-Arm Robot," ASME J. Dynamic Systems, Measurement, and Control, June 1989, vol. 111, pp 205-210.

C.R. Carignan,
"Adaptive Tracking for Complex Systems Using Reduced Order Models," IEEE Int. Conf. Robotics and Automation, Cincinnati, OH, USA, May 1990, Vol. 3, pp 2078-2083.

T. Yoshikawa and X. Zheng,
"Coordinated Dynamic Hybrid Position/Force Control for Multiple Robot Manipulators Handling One Constrained Object," IEEE Int. Conf. Robotics and Automation, Cincinnati, OH, May 13-19, 1990, vol. 2, pp. 1178-1183.

M.A. Unseren,
"Rigid Body Dynamics and Decoupled Control Architecture for Two Strongly Interacting Manipulators," Accepted for publication in Robotica, (1991a).

M.A. Unseren,
"A Rigid Body Model and Decoupled Control Architecture For Two Manipulators Holding A Complex Object," Oak Ridge National Laboratory Publication, ORNL/TM-11752 CESAR-91/01, Oak Ridge, TN, Jan. (1991b).

R. K. Kankaanranta and H. N. Koivo,
"Dynamics and Simulation of Compliant Motion of a Manipulator," IEEE J. Robotics and Automation, April 1988, vol. 4, no. 2, pp 163-173.

N.H. McClamroch and D. Wang,
"Feedback Stabilization and Tracking of Constrained Robots," IEEE Trans. Automatic Control, May 1988, Vol. 33, No. 5, pp 419-426.

Y.F. Zheng and J.Y.S. Luh,
"Optimal Load Distribution for Two Industrial Robots Handling a Single Object," IEEE Int. Conf. Robotics and Automation, Philadelphia, PA, USA, April 1988, Vol. 1, pp 344-349.

D.E. Orin and S.Y. Oh,
"Control of Force Distribution in Robotic Mechanisms Containing Closed Kinematic Chains," ASME J. Dynamic Systems, Measurement, and Control, vol. 102, June 1981.

I.D. Walker, S.I. Marcus, and R.A. Freeman,
"Distribution of Dynamic Loads for Multiple Cooperating Robot Manipulators," J. Robotic Systems, vol. 6, no. 1, pp 35-47, Jan. 1989.

M.A. Unseren,
"A Rigid Body Model and Decoupled Control Architecture for Two Manipulators Holding a Complex Object," To Appear in Computers and Electrical Engineering, (1991c).

M.A. Unseren,
"Dynamic Coupling Effects in Modeling and Control of Hard Contact Motion of a Manipulator," To be presented at IEEE Int. Conf. Systems Engineering, Dayton, Ohio, USA, August (1991d).

D.B. Reister,
"A New Wheel Control System for the Omnidirectional HERMIES-III Robot," to be presented at the IEEE Int. Conf. Robotics and Automation, Sacramento, CA, USA, April 1991.

4.2 Combined Mobility and Manipulation Systems

Autonomous mobile manipulators (i.e., manipulators mounted on mobile platforms) have attracted significant interest in the industrial, military, and public service communities because of the potential they provide for increased efficiency (over fixed-base manipulators) in material handling and manipulation tasks (such as warehouse management, plant surveillance, maintenance or repair, material transfer and delivery, etc.). They are also of prime interest to the DOE community because of their capability to perform tasks in indoor and/or large-scale outdoor environments unaccessible to or very hazardous for humans (e.g., nuclear plants, radioactive or chemical materials handling, waste management, fire fighting, emergency response, etc.). For high efficiency and speed of operation, such mobile manipulators need to incorporate capabilities in both, as well as in combined, autonomous large-scale mobility and autonomous dexterous manipulation.

An important characteristic of practical mobile manipulators is their particular kinematic redundancy, created by the addition of the DOF of the platform to those of the manipulator (e.g., a 3 DOF platform and 7 DOF manipulator produce a 10 DOF system). This redundancy is quite desirable for dexterous motion, manipulation and transport functions in obstructed environments where fixed-base robots and/or non-redundant robots have limited access or reach. It also allows the systems to be optimally positioned and configured for maximum performance when stringent constraints exist on the robot (e.g., joint limits, maximum actuator torques), the environment (e.g., heavy loads, tight access, reach requirements), and/or the task themselves (e.g., safety requirements, stability, time constraints), thus allowing some tasks to be performed which otherwise would not have been feasible. On the other hand, this particular redundancy represents a challenge from the planning and controls point of view, because the coupling of the mobility and manipulation functions introduces a new order of complexity in the redundancy resolution problem. This complexity arises mainly from the fact that, while the manipulator typically is a serial link system, the platform generally is not (and often includes non-holonomic constraints). From this standpoint, the system is not homogeneous and conventional redundancy resolution methods, developed mainly for fixed-base manipulators, are not directly applicable to global coordinated motion planning and control of the overall system.

Another major consequence of this kinematic heterogeneity, is that the modes of motion of the system (serial link vs. non-serial link vs. mixed mode) and their associated requirements and constraints greatly vary from one task to another, for example, in the following sequence: minimum-time platform motion toward a work area, arm motion toward an object to be grasped under obstacle avoidance requirements, lifting of the object with actuator torque limit constraints, arm motion with the lifted object under platform stability constraint, etc. The major problem resulting from this diversity of motion modes, requirements, and constraints occurs at task commutation, when these parameters change. This is because the final position and configuration reached by the system under a mode of motion may not be suited for, or may not even allow, performance of the next task. In the task sequence mentioned above, for example, the position attained by the platform and the configuration reached by the arm in its motion toward the object to be grasped, may not allow the object to be lifted without exceeding the limit of a joint actuator torque, whereas another position and configuration of the platform and manipulator may have allowed the task to be performed. The key requirement for efficient planning and execution of sequences of widely varying tasks by mobile manipulators is therefore a forecast of the set of possible configurations (and the optimal one if necessary) of the overall system (platform and manipulator) from which the next task is to be initiated, and which assure feasibility of that task under its specific requirements and constraints. This forecasted initial configuration for the upcoming task thus represents the configuration in which the system must be when task commutation occurs (with changes in constraints, requirements, predominant mode motion, etc.), and therefore, also is the final configuration which the system needs to reach in the immediately preceding task (using the redundancy resolution scheme of this finishing task).

To be able to plan and perform widely varying sequences of tasks, three critical subproblems need to be addressed: (1) forecasting the optimal "commutation configurations" for each couple of adjacent tasks, (2) planning in the Cartesian task space the motion of the system for each task, considering not only the requirements and constraints specific to that task but also the two boundary values provided by the initial and final commutation configurations for that task, and (3) given the Cartesian trajectories, calculate and control in articular (or joint) space the motion of the redundant set of DOF. The activity proposed here will focus on the development of the novel methodologies and techniques needed to resolve these classes of problems for the general case of combined N DOF mobility/manipulation systems, and to demonstrate these developed methodologies on proof-of-principle experiments using the HERMIES robots.

4.2.1 Background

Although significant amounts of research and development have been performed in both areas of autonomous navigation for mobile platforms [e.g. 1, 2 and 3, and references therein] and autonomous control of manipulators [e.g. 4-17], very little has been reported on the topic of coupling these two functions for the purpose of developing an autonomous robot with combined large-scale mobility/dexterous manipulation capabilities [18-26]. Several laboratories in the military, NASA, or DOE community have recognized the potential of large-scale mobility and manipulation systems for various battlefield or space exploration support functions, and have designed or fabricated platforms intended as mobile manipulators [20, 21] (e.g., ATLAS, TMAP, Mars Rover). Most of these systems, however, are of little interest here since they are teleoperated, relying exclusively on humans for intelligence, planning, and control, and since none of them incorporates autonomous navigation and manipulation capabilities. CESAR researchers have contributed to the pioneering effort in this area of combined mobility and manipulation (CM&M) through the development of HERMIES-III [18], one of, if not the earliest autonomous and highly redundant experimental mobile manipulators.

Only a few researchers have approached the problem of CM&M, and most of these have focussed on specific control aspects of the subproblem (3) mentioned above. Dubowsky et al. [27] at MIT have analyzed the optimal control of a non-redundant manipulator on a platform, taking into account actuator torque limits and stability of the platform steadied on outriggers. This analysis has recently been extended to the case of a suspended platform [28], however, without incorporating navigation or large-scale motion planning for the platform, i.e., effectively not addressing the major complexity of motion planning for an heterogeneous system. Vachenanos et al. [29] and Lewis [30] have mentioned initial efforts in the control subproblem (3), focussing on an analysis of the dynamic interaction forces between the manipulator and the platform, and aiming at the determination of stability ranges for the manipulator motion for given reaction terms due to the platform.

Few recent reports exist on the position and configuration optimization problem for CM&M, originating mainly from Ohio State University [25, 26] for legged robots, and from Carnegie-Mellon University [19,24] and CESAR [20,22,23] for wheeled robots. All teams have initially focussed on subproblem (1), the development of methodologies for forecasting sets of commutation configurations for CM&M systems. In the proposed approaches, the problem of resolving the redundancy for determining the mobile manipulator commutation positions and configurations was formulated as a constrained optimization in the configuration space of the overall system. A set of constraints representing fixed values or bounds on some of the system variables (e.g., the platform position component in the vertical direction remains zero, the end-effector position is fixed at a given point in space, etc.) are imposed on the system. In one approach, the configuration of the arm is first adjusted for reach and maximum strength, and the platform position is then determined accordingly [25, 26]. In two other approaches, a simple cost function weighing the kinetic energy [25] or the relative magnitudes of the platform and manipulator motions is proposed as the objective function in optimization schemes using gradient descent [19] or simulated annealing [24]. In [22, 23] a multicriteria optimization approach is introduced which allows several requirements which the system seeks to optimize during task performance such as, for example, safety by staying away from obstacles, best utilization of actuator strength, or maintaining some maneuverability in all directions (e.g., by staying away from singular configurations) to be weighted and taken into account in the objective function for the optimization. Note that these optimization problems are formulated independently of time, i.e., they solve for a set of static solutions for the system. This, in the robotic motion planning and control context, is often referred to as a "local" (with respect to time), or "instantaneous" optimization, as opposed to the much more complex "global" optimization problems in which optimization criteria typically are integrals over time-varying trajectories of some functionals of system variables. The class of problems in subproblem (2) mentioned above, is of the latter type, and to our knowledge, has not been approached or published to this date. For clarification of semantics, we should note here that in all types of optimization problems, the terms local vs. global extrema are also utilized to represent suboptimal vs. optimal solutions that may be encountered depending on the form of the objective functional.

Very little work has been published on the problem of global motion planning for CM&M systems [19, 31, 32]. The major complexity arising with mobile manipulators comes from the fact that while the Cartesian variables used to describe the motion of a manipulator in (Cartesian) task space are functions of the joint variables, those describing the motion of a platform in task space are functionals of the joint (or control) variables. When coupling these two subsystems for the purpose of Cartesian path planning, the system becomes non-homogeneous and a unified solution approach is needed. Barraquand, Langlois and Latombe [31] and Zhu and Latombe [32] have proposed the use of guided search techniques in discretized configuration space to resolve this problem. This method, however, is accurate to within a mesh spacing and has to rely on potential functions and heuristics to satisfy the system's constraints and to guide the search. Development of deterministic approaches to the general CM&M path planning problem, where time-dependent constraints and global objective functionals can be satisfied, is needed for application to our robotics domain, and is one of the focus of our proposed research. In our initial activities in this area, we have investigated extensions of Pontryagin Maximum Principles to the problem of deterministic global path-planning for constrained N DOF systems [33].

Concerning the issue of motion control for kinematically redundant systems, a large body of literature exists which deals with redundancy resolution for robotic systems, however, mostly devoted to serial-link systems such as conventional fixed-base manipulators [e.g. 4-17]. The common ground of all these approaches is their use of Moore-Penrose pseudo-inverse [34] techniques to solve the inverse kinematic problem, with the implicit assumption that the least norm of the joint velocity vector or joint torque vector is the primary redundancy resolution criterion. In our recent developments and those by others in this area [35, 36, 37] a more general approach to the inverse kinematic problem for general N DOF redundant systems has been proposed. The approach allows calculation of the entire solution space (an affine space) using generator functions composed of solution vectors projected on the subspaces defined by submatrices of the Jacobian matrix. The approach has the further advantage of being parallelizable and appears promising for application to real-time systems. An alternative criterion to the least norm, reducing the problem to a "Minimax" [23, 38, 39] optimization was also investigated, which provides a much more uniform and realistic joint or torque distribution in the solution.

4.2.2 Proposed Research

Our proposed activities will build upon our initial developments to investigate solution methodologies for the three classes of problems discussed. These are: (1) the determination of the sets of commutation configurations which allow performance of series of tasks with widely varying constraints and requirements (objective functions), (2) the planning of motion trajectories between each pair of commutation configurations, while satisfying time-dependent system and environmental constraints and global objective functions, and (3) given the system trajectories, determining the optimal controls on the system which guarantee appropriate tracking, accuracy, and integrity of the motions.

Within the context of CESAR's commitment to experimental verification of advanced methodologies, our currently available HERMIES-III (10 DOF) mobile manipulator, and possibly its much smaller successor HERMIES-IV, will provide the framework for realistic (sensor-data based) testing and proof-of-principle experiments. Initial experimental paradigms would simulate CM&M operations in highly constrained environments (e.g. navigation, inspection, and manipulation of controls or valves around complex flow pipe systems) and dextrous manipulation in highly obstructed environments with possibly near-limit loads (e.g. stacking or retrieving waste barrels or packages, manufacturing or warehouse management, fire fighting, or component repair in confined and cluttered environments). Other research activities proposed in the areas of perception and advanced computing will provide added support for further developments of the methodologies in problems involving sensor-based CM&M motion planning and control, and real-time CM&M operation in highly dynamic environments. Finally, developments will be pursued for coupling with the progress in the multi-robot cooperation area to address the general paradigm of cooperative activities by multiple mobile manipulators.

4.2.3 References

1. J. Crowley, "Navigation of an Intelligent Mobile Robot," IEEE J. Rob. Auto. (March 1985).

2. C. R. Weisbin, G. de Saussure, J. R. Einstein, and F. G. Pin,"Autonomous Mobile Robot Navigation and Learning," IEEE Computer 22(6), 29--35 (1989).

3. F. G. Pin and H. A. Vasseur, "Autonomous Trajectory Generation for Mobile Robots with Non-Holonomic and Steering Angle Constraints," in Proceedings of the IEEE Workshop on Intelligent Motion Control, Istanbul, Turkey, August 20-22, 1990, pp. 295-9.

4. J. M. Hollerbach and K. C. Suh, "Redundancy Resolution of Manipulators Through Torque Optimization," Proceedings of the IEEE Conference on Robotics and Automation, St. Louis, Missouri, March 1985, pp. 1016-21.

5. J. Baillieul, J. Hollerbach, and R. W. Brockett, "Programming and Control of Kinematically Redundant Manipulators," Proceedings of the 23rd IEEE Conference on Decision and Control, December 1984, pp. 768-74.

6. T. Yoshikawa, "Analysis and Control of Robot Manipulators with Redundancy," Proceedings of the Robotics Research: First International Symposium, ed. M. Brady and R. Paul, Cambridge, Massachusetts, MIT Press, 735-48 (1984).

7. M. S. Konstantinov, M. D. Markov, and D. N. Nenchev, "Kinematic Control of Redundant Manipulators," Proceedings of the 11th International Symposium on Industrial Robots, Tokyo, Japan, 1981, pp. 561-68.

8. H. Hanafusa, T. Yoshikawa, and Y. Nakamura, "Analysis and Control of Articulated Robot Arms with Redundancy," Proceedings of the 8th IFAC World Congress, XIV, August 1981, pp. 78-83.

9. C. A. Klein, "Use of Redundancy in the Design of Robotic Systems," Proceedings of the 2nd International Symposium on Robotic Research, Kyoto, Japan, 1984, pp. 58 -65.

10. J. Y. S. Luh and Y. L. Gu, "Efficiency and Flexibility of Industrial Robots with Redundancy," Proceedings of the ASME Winter Annual Meeting, PED 15, Robotics and Manufacturing Automation, ed. M. Donath and M. Lay, 37-41 (1985).

11. R. V. Dubey, J. A. Euler, and S. M. Babcock, "An Efficient Projection Optimization Scheme for a Seven-Degree-of-Freedom Redundant Robot with Spherical Wrist," Proceedings of the IEEE Conference on Robotics and Automation, Philadelphia, Pennsylvania, April 1988, pp. 28-36.

12. P. H. Chang, "A Closed Form Solution for Control of Manipulators with Kinematic Redundancy," Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, California, April 1986, pp. 9-14.

13. C. W. Wampler, "Inverse Kinematic Functions for Redundant Manipulators," Proceedings of the IEEE International Conference on Robotics and Automation, Raleigh, North Carolina, April 1987, pp. 610-17.

14. Y. Nakamura and H. Hanafusa, "Task Priority Based Redundancy Control of Robot Manipulators," Proceedings of Robotics Research - 2nd International Symposium, ed. H. Hanafusa and H. Inoue, Cambridge, Massachusetts, MIT Press, 155-62 (1985).

15. J. Baillieul, "Kinematic Programming Alternatives for Redundant Manipulators," Proceedings of the IEEE Conference on Robotics and Automation, St. Louis, Missouri, March 1985, pp. 722-28.

16. A. A. Maciejewski and C. A. Klein, "Obstacle Avoidance for Kinematically Redundant Manipulators in Dynamically Varying Environments," Robotics Research 4(3), 109-17 (1985).

17. O. Khatib, "Real-Time Obstacle Avoidance for Manipulators and Mobile Robots," Robotics Research 5(1), 90-8 (1986).

18. C. R. Weisbin et al., "HERMIES-III: A Step Toward Autonomous Mobility, Manipulation, and Perception," Robotica 8, 7-12 (1990).

19. W. F. Carriker, P. K. Khosla, and B. H. Krogh "An Approach for Coordinating Mobility and Manipulation," Proceedings of the IEEE International Conference on Systems Engineering, Dayton, Ohio, August 1989, pp. 40-4.

20. S. F. Sousk, F. G. Pin, and C. D. Crane, "Motion Planning for the Universal Self Deployable Cargo Handler (USDCH)," in Proceedings of the 17th U.S. Army Science Conference, Durham, North Carolina, June 12-5, 1990.

21. H. G. McCain et al., "A Hierarchically Controlled Autonomous Robot for Heavy Payload Military Field Applications," in Proceedings of the International Conference on Intelligent Autonomous Systems, The Netherlands, December 8-11, 1986.

22. F. G. Pin and J.-C. Culioli, "Multi-Criteria Position and Configuration Optimization for Redundant Platform/Manipulator Systems," in Proceedings of IROS '90, the Third IEEE International Workshop on Intelligent Robots and Systems, Tsuchiura, Japan, July 3-6, 1990, pp. 103-7.

23. F. G. Pin and J.-C. Culioli, "Optimal Positioning of Combined Mobile Platform-Manipulator Systems for Material Handling Tasks" submitted to Journal of Robotic Systems.

24. W. H. Carter, P. K. Khosla and B. H. Krogh, "The Use of Simulated Annealing to Solve the Mobile Manipulator Path Planning Problem," in Proceedings of the IEEE International Conference on Robotics and Automation, Cincinnati, Ohio, May 1990, pp. 204-9.

25. Y. F. Zheng and Q. Yin, "Coordinating Multi-Limbed Robots for Generating Large Cartesian Force", in Proceedings of the 1990 IEEE Conference on Robotics and Automation, Cincinnati, OH, May 1990, PP 1653-1658.

26. C. SU and Y. F. Zheng, "Task Decomposition for Multi-Limbed Robots to Work in the Reachable but Orientable Space", Proceedings of the 1990 IEEE Conference on Robotics and Automation, Cincinnati, OH, May 1990, pp 1659-1664.

27. S. Dubowsky and E. E. Vance, "Planning Mobile Manipulator Motions Considering Vehicle Dynamic Stability Constraints," in Proceedings of the 1989 IEEE Conference on Robotics and Automation, Scottsdale, Arizona, May 1989, pp. 1271-6.

28. N. A. M. Hootsmans and S. Dubowsky, "Control of Mobile Manipulators Including Vehicle Dynamic Characteristics," in Proceedings of the 1991 ANS Topical Conference on Robotics and Remote Systems, Albuquerque, New Mexico, February 24-28, 1991, pp. 461-70.

29. G. Vachevanos, "Proposed Activities in Mobile Manipulator Research" Personal Communication to F. G. Pin, January 1991.

30. F. Lewis, "On the Dynamic Coupling of Manipulators and Mobile Platform" Personal Communication to F. G. Pin, January 1991.

31. J. Barraquand, B. Langlois, and J. C. Latombe, "Robot Motion Planning with Many Degrees of Freedom and Dynamic Constraints," in Proceedings of the Fifth International Symposium on Robotics Research, August 1989, Tokyo, Japan.

32. D. Zhu and J. C. Latombe, "New Heuristic Algorithms for Efficient Hierarchical Path Planning," Stanford University Technical Report, Department of Computer Sciences, August 1989.

33. D. B. Reister and F. G. Pin, "Global Path Planning of Mobile Manipulators Using Pontryagin Maximum Principles" in preparation.

34. A. Albert, "Regression and the Moore-Penrose Pseudo-Inverse, Academic Press, 1972.

35. P. F. R. Belmans and J.-C. Culioli, "A New Approach to Solve the Kinematics Resolution of a Redundant Robot," ORNL/TM-11435, March 1990.

36. P. F. R. Belmans and J.-C. Culioli, "An Alternative Method to Solving the Kinematic of a Redundant Robot," Trans. Am. Nucl. Soc. 61, pp. 426-7 (1990).

37. S. Hirose and s. Ma, "Redundancy Decomposition Control for Multi-Joint Manipulator", in Proceedings of the 1989 IEEE Conference on Robotics and Automation, Scottsdale, AZ, May 1989, pp 119-124.

38. V. F. Dem'yanov and V. N. Malozemov, Introduction to Minimax, Wiley, New York (1974).

39. S. R. K. Dutta and M. Vidyasagar, "New Algorithms for Constrained Minimax Optimization," Mathematical Programming, 140--55 (1977).

4.3 Multi-Sensor Data Analysis and Fusion

One of the prerequisites for intelligent behavior in robotic systems is the ability to generate internally self-consistent representations of the environment from sensor data. Sensor fusion may be regarded as a strategy for reducing errors and uncertainties, and overcoming sensor and mechanical limitations, to achieve the desired performance goals through the construction of reliable representations of the world. Some of the previous work done in the CESAR laboratory on sensor fusion was described in Secs 3.3 and 3.4. We propose to continue this research, focussing on some of the fundamental problems and issues brought out in our initial studies.

4.3.1 Background

The term "sensor fusion", as used in the literature, denotes a broad area of research in robotics and automation. The activities in this area range from 3-D object recognition (see, for example, Besl and Jain, 1985) to uncertainty calculii, and from high-level, application-independent fusion to the task-driven fusion of data from specific sensors. In more detail, we note that Pau (1982) presented some ideas from the point of view of statistical pattern recognition, refering to high level fusion independent of any particular application. Henderson et al. (1984) introduced the concept of logical sensors and multi-sensor kernel system as a uniform mechanism to deal with data from diverse sensors. In that study no attempt was made to incorporate in the concepts the levels of confidence associated with information from different sensors.

Turning to the task-driven fusion of data from specific sensors, we note that Nitzan et al (1977) and Duda, et al. (1979) fused registered range and reflectance data for scene analysis. Flynn (1985) reported on the combination of ultrasound and infrared data, and Allen and Bajcsy (1985) combined stereo vision and active tactile sensing for recognition of objects. Ruokangas et al. (1986) described a system that integrated 2D vision and acoustic distance measurements for a stationary robot. Magee et al (1985) presented results of experiments with intensity-guided range sensing, and Nandhakumar and Aggarwal (1988) reported results on the combination of thermal and visual data obtained from outdoor scenes.

An important objective in the above-mentioned, low-level fusion studies, as well as in our own investigations, was the reduction of errors. We may recall that sensor errors are of two types: random and systematic. The former include errors due to counting statistics and due to noise. This type of error influences the precision of the extracted information. Systematic errors are all those that are not random in character. These errors, generated primarily by the lack of sufficient information to correctly interpret the output of the low level processing algorithms, determine the accuracy of the extracted information.

We view sensor fusion as a means to reduce both statistical and systematic sensor errors, thereby accomplishing improved reliability and robustness of the integrated robotic system. We adopt the principle that fusion should take place at the lowest level possible in the information processing in order to prevent or reduce the propagation of errors (Beckerman, Farkas and Johnston, 1990).

In Beckerman and Barnett (1991) we reported results of a performance analysis of the visual and ultrasonic sensor subsystems for the HERMIES-IIB mobile robot. The key component in the vision subsystem used in the study was an image segmentation module which first labelled the components of the thresholded binary image, and then did a region analysis to produce a list of geometric properties of each labelled region. In our performance analysis we focused on the areas and height-to-width ratios of the surfaces so identified. We showed that we could decrease many of the errors in these geometric quantities by fusing visual results with ultrasound range information.

A second major theme of our research deals with the control aspects of sensor fusion. Returning to the literature briefly, we find that Aloimonos at al. (1987) noted that ill-posed vision tasks can be made well-posed through the addition of images from different viewpoints and times. The additional data provides the necessary constraints to stabilize edges and other features against noise. Bajcsy (1988) expanded upon these findings and noted that active perception (vision) is more general and includes sensor modeling and control strategies. Control of the acquisition of information was also the subject of studies by Burt (1988), and by Ferrier and Clark (1990). Motion to achieve stability of the low level algorithms was emphasized by Ferrier and Clark (1990). Burt (1988) advocated a selective processing of visual data analogous to peripheral alerting and foveal attention. Bajcsy (1988) introduced the use of loss functions and risk minimization to control the seeking of information, and in our recent study (Beckerman and Barnett, 1991) we pointed out that optimality of a navigation path is often sacrificed in favor of robustness, or stability against sensor errors of various types.

Sensor error models form an important component of any system designed to fuse feature information. They provide needed estimates of the accuracy and precision associated with the extracted features, and should ideally provide estimates for both statistical and systematic errors. The latter, however, are often difficult or impossible to predict. In Beckerman and Oblow (1990) we showed that in some instances these errors can be identified through their patterns of conflict, i.e., they generate recognizable conflicts among initially interpreted data from different measurements. One can then reinterpret the data and remove the errors. This two-stage identification and reinterpretation of the output of low-level sensor processing algorithms is analogous to a feedback process working on the initially interpreted sensor data (Beckerman, Barnett and Killough, 1990). The monitoring and control to achieve internal self-consistency of the sensor-derived world model is an important theme in our sensor fusion research, and extends the ideas put forth by Aloimonos et al (1987) and Bajcsy (1988).

An important goal of our proposed research is to further identify and research those crucial elements needed in intelligent systems for control and monitoring of the data collection, processing and interpretation. In doing so we intend to focus on outstanding questions regarding the treatment of data which are not independent, on construction of useful representations (world models), on developing a common formalism for the various control functions, on methodologies for systematic and for random error reduction, and on techniques for improving the quantity and quality of information acquired from the sensors.

4.3.2 Proposed Research

To achieve the above-mentioned goals we will make use of existing and developing robot systems in the CESAR laboratory. In particular, the HERMIES-

III robot sensor suite will be the subject of experimental and analytical research. Sensors currently available on HERMIES-III, relevant to this work, include a laser range camera, several CCD cameras (including stereo pairs on pan and tilt mechanisms), an array of sonar transducers, wheel encoders to provide initial estimates of the robot's position, as well as a number of sensors associated with the on-board redundant robot manipulator, including joint resolvers, a force/torque sensor, a tactile sensor pad, and an LED proximity sensor. Future additions to the robot sensor package include a color camera and a thermal infrared imaging sensor. The tasks under consideration include robot self- location, exploration and mapping, and sensor-based control of redundant manipulators.

The appropriate characterization of errors associated with the sensors considered is a prerequisite for successful fusion. Specific fusion tasks will include the combination of laser range and reflectance images and features. Contrary to expectations, our preliminary work on characterizing the Odetics laser range camera showed that range and reflectance measurements are not independent. Several types of systematic errors in the range image have been identified, and found to depend on the reflectance image. We propose to generalize earlier results of Gil et al (1983) and Duda et al (1979), and to develop a formal way to treat the lack of independence in these data.

Data representations, also refered to as world models, represent an important issue in sensor fusion. Our work will make explicit that the complexity of these world models needs to reflect the classes of tasks to be performed, and the requirements for monitoring and control. Motivated by our commitment to perform this research in the context of proof-of-principle experiments with the CESAR robots, we will focus on data representations of minimum complexity for a given class of tasks. World models can range from 2, "2 1/2", and 3 dimensional maps, to geometric primitives (2D, 3D line segments, surface patches, etc.), to CAD-based, highly complex models.

Motivated by providing a general framework for fusion of low-level cues and features, particularly for visual data and information, Markov Random Fields (MRFs) have been suggested as a basis for fusion methods (Poggio et al, 1987). Work with MRFs in computer vision areas was inspired by the seminal contribution by Geman and Geman (1984). Recent improvements to the performance, i.e. speed of convergence, of associated optimization methods (e.g., Hiriyannaiah et al, 1989), have made the MRF framework potentially attractive for robotics applications. The preliminary report by Bilbro and Snyder (1988) on fusion of range and reflectance image data using MRFs made oversimplified assumptions regarding the sensor error and scene reflectivity models. With our increased knowledge about the Odetics laser range camera characteristics, we propose to re-evaluate the MRF approach to low level fusion.

Although there is a wide variety of methods for fusing data ranging from direct fusion algorithms to mean square estimation and statistical decision theory (e.g. Hager and Mintz, 1989) to rules of combination, there are underlying common grounds. One of our research tasks will be to clearly delineate these common grounds. We will do so using the CESAR research facilities. We will perform experiments for a number of environments to identify problems and issues, and to perform quantitative assessments of various sensor fusion approaches.

4.3.3 References

Allen, P. and R. Bajcsy, "Object Recognition Using Vision and Touch," Proc. of the 9th Int. Conf. on AI, Los Angeles, August 1985.

Aloimonos, J., I. Weiss and A. Bandyopadhyay, "Active Vision," Proc. 1st IEEE Conf. Comp. Vis. (London), pp. 35-54, 1987.

Ayache, N. and O. D. Faugeras, "Building, Registrating and Fusing Noisy Visual Maps," Int. J. Robotics Res., Vol. 7, pp. 45-65, 1988.

Bajcsy, R., "Active Perception," Proc. IEEE, Vol. 76, pp. 996-1005, 1988.

Beckerman, M. and E. M. Oblow, "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor data for Robotic Navigation", IEEE Trans. Robotics Automat., Vol. RA-6, pp 137-145, 1990.

Beckerman, M., L. A. Farkas and S. E. Johnston, "Treatment of Systematic Errors II: Fusion of Ultrasound and Visual Sensor Data, Oak Ridge National Laboratory Technical Memo No. ORNL/TM- 11349, 1990, and submitted to IEEE Trans. Robotics Automat., 1990.

Beckerman, M., D. L. Barnett and S. M. Killough, "Ultrasound and Visual Sensor Feedback and Fusion," Proc. 1990 Japan-U.S.A. Symposium on Flexible Automation (Kyoto, Japan), pp. 1315-1319, 1990.

Beckerman, M. and D. L. Barnett, "Performance of Visual and Ultrasound Sensing by an Autonomous Robot," Oak Ridge National Laboratory Technical Memo No. ORNL/TM-11733, 1991, and to be submitted to IEEE Trans. Robotics Automat., 1991.

Besl, P. J. and R. C. Jain, "Three-Dimensional Object Recognition," Computer Surveys, Vol. 17, pp. 75-145, 1985.

Bilbro, G. L. and W. E. Snyder, "Fusion of Range and Reflectance Image Data Using Markov Random Fields," Proc. IEEE Int. Symp. Intell. Control, Arlington VA, pp. 154-158, 1988.

Burt, P. J., "Algorithms and Architectures for Smart Sensing," Proc. Image Understanding Workshop (Cambridge, MA), pp. 139-153, 1988.

Duda, R. O., D. Nitzan and P. Barrett, "Use of Range and Reflectance Data to Find Planar Surface Regions," IEEE Trans. Pattern Anal. Machine Intell., Vol. PAMI-1, pp. 259-271, 1979.

Ferrier, N. and J. J. Clark, "Optimal Motion for Active Sensing," Proc 3rd. Int. Symp. Robot. Manuf. (Vancouver, BC), pp. 169-174, 1990.

Flynn, A., "Redundant Sensors for Mobile Robot Navigation," MS Thesis, MIT Cambridge MA, July 1985.

Geman, S. and D. Geman, "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," IEEE Trans. Pattern Anal. Machine Intell., PAMI-6, pp. 721-741, 1984.

Gil, B., A. Mitiche and J. K. Aggarwal, "Experiments in Combining Intensity and Range Edge Maps," Comp. Vis., Graph., Image Processing, Vol. 21, pp. 395-411, 1983.

Hager, G. and M. Mintz, "Task-Directed Multi-Sensor Fusion," proc. IEEE Int. Conf. Robotics Automat. (Scottsdale), pp. 662- 667, 1989.

Henderson, T. C., W. S. Fai and C. Hansen, "MKS: A Multisensor Kernel System," IEEE Trans. Sys. Man, Cybern., SMC-14, pp. 784-791, 1984.

Hiriyannaiah, H. P., G. L. Bilbro, W. E. Snyder and R. C. Mann, "Restoration of Piecewise-Constant Images by Mean Field Annealing," J. Opt. Soc. Am, A, Vol. 6, pp. 1901-1912, 1989.

Magee, M. J., B. A. Boyter, C.-H. Chien and J. K. Aggarwal, "Experiments in Intensity Guided Range Sensing Recognition of Three-Dimensional Objects," IEEE Trans. Pattern Anal. Machine Intell., Vol. PAMI-7, pp. 629-637, 1985.

Nandhakumar, N. and J. K. Aggarwal, "Integrated Analysis of Thermal and Visual Images for Scene Interpretation," IEEE Trans. Pattern Anal. Machine Intell., Vol. PAMI-10, pp. 469-481, 1988.

Nitzan, D., A. E. Brain and R. O. Duda, "The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis," Proc. IEEE, Vol. 65, pp. 206-220, 1977.

Pau, L. F., "Fusion of Multisensor Data in Pattern Recognition," in Pattern Recognition Theory and Applications, J. Kittler, K. S. Fu, L. F. Pau (eds.), pp. 189-201, D. Reidel Publishing Company, 1982.

Poggio, T., et al., "The MIT Vision Machine," Proc. Image Understanding Workshop, Cambridge, MA, pp. 177-198, 1987. Rauch, H. E., Probability Concepts for an Expert System Useed for Data Fusion," AI Magazine, pp. 55-60, Fall 1984.

Ruokangas, C. C., M. S. Black, J. F. Martin and J. S. Schoenwald, "Integration of Multiple Sensors to Provide Flexible Control Strategies," Proc. IEEE Int. Conf. on Robotics and Automation, San Francisco, April 1986.

4.4 Machine Learning

Investigations of many aspects of learning that are important characteristics of systems that can be referred to as intelligent, and that go beyond the simplest forms of parameter estimation and long-standing proven methods of probability theory and statistics, have so far largely failed to show major impact on robotic systems. CESAR research has resulted in successful robot experiments in scenarios with well-understood and constrained complexity. Promising advancements have been achieved in the areas of concept representations and scalability of methods to complex problems. We propose to continue addressing these fundamental issues with special emphasis on bridging the gap between important analytical results and implementations in robotic testbeds.

4.4.1 Background

The field of machine learning (by which we specifically mean induction from examples) has had a long history, characterized by many diverse approaches to concept knowledge acquisition by a machine. Most methods use either a bottom-up approach, relying heavily on sensory data, or a top-down approach, relying more on models and a programmed knowledge base. Of the methodologies receiving the most current attention, neural nets (Grossberg 1982; Rumelhart & McClelland 1989) successfully employ the bottom-up approach, while traditional heuristic search techniques (Pearl 1984, Nilsson 1980) are generally of the top-down variety. Some of the other more practical techniques like decision trees (Quinlan 1986a), genetic algorithm classifiers (Holland 1975, 1986a, 1986b), and parametric and non-parametric statistical methods (Duda & Hart 1973) use both approaches. The chief objective of all these approaches, however, is the same -- to produce effective generalizations (i. e. concepts) from examples.

While significant success has been achieved in many areas, it is generally recognized (although not always explicitly stated), that all of the current methods mentioned fall far short of achieving human inductive skills. Baum (1990) has been one of the few prominent researchers in this field to be openly critical of the current state of affairs. He has stated that despite the apparent success of current theoretical efforts, modern approaches have not been able to solve even the simplest nontrivial problems any better than the methods available 30 years ago.

We strongly agree with Baum's judgment. Despite the current overly optimistic press, inductive learning is still at best only a sophisticated form of parameter estimation. The differences between top-down and bottom-up methods seem to be little more than measures of the relative sizes of the knowledge bases and search spaces involved. This state of affairs leaves many fundamental learning issues related to representation and scalability unsolved and much of the learning field largely unexplored.

Before discussing the proposed effort, we caution that the above assessment should not be taken out of context. While many roadblocks exists, it is equally clear that learning science (i. e. the understanding of our own biological learning capabilities and neurological makeup) is still in its infancy. Our point of departure within this context is that we believe that an important reason for our current (and possibly ultimate) lack of success in this field is the widely overlooked issue of complexity. The long evolutionary development of the human brain has resulted in large numbers of highly specialized structures which interact simultaneously to generate human learning capabilities. This resulting complexity has both architectural and information-theoretic components, both of which are apparently essential for learning in our rich context-dependent environment. Such complexity, we feel is the root cause of the failure of present machine learning efforts to even remotely approach human performance levels.

This assessment forces us to approach our research with the clear possibility that the complexity we are trying to understand cannot be represented by anything less complex than itself. Local approaches to this global problem may not succeed and traditional reductionist approaches used in the physical sciences may also not work here. To reach human levels of performance, it is quite possible that machines might likewise have to evolve over long (possibly unacceptably so) periods of time with exposure to complex human environments and problems. The alternative (which may be even more unpalatable) is to program all context dependent knowledge directly into a machine to solve each specific problem. Either prospect makes the ideal of machine intelligence much farther from our grasp than some would have us believe.

The conclusion we have drawn from this overall assessment is that the complexity issue should be treated as a fundamental problem in machine learning. As such, it must be dealt with directly and systematically. Without dramatic breakthroughs in our understanding of the human brain, this means that we are forced to develop a variety of generic, scalable tools which can accept significant amounts of programmed knowledge and representational information. To do this systematically, we must devote more theoretical attention to the generic aspects of the complexity problems associated with representation and learnability. The time, representational, and storage constraints of a machine algorithm must be quantitatively understood and factored into any learning methodology. Indeed, the inherent speed and storage advantages of machines should be used effectively to offset the lack of global information.

Since the focus of the CESAR learning effort is on autonomous intelligent machines, we must relate these conclusions to the specific task of developing algorithms that give a machine the ability to learn from and adapt to a largely unstructured environment. In this context, most of the relevant learning information is sensor-based and our methodology will, therefore, necessitate a predominantly bottom-up approach centered on pattern classification problems. This makes it imperative that any learning methodology we develop make very efficient use of computational resources.

With these constraints, the three major issues in this machine learning framework that we will focus on are: 1) development of an appropriate generic representation for the concepts to be learned, 2) establishing quantitative understanding of learnability issues using such a representation, and 3) designing efficient and scalable algorithms in such a generic framework. Short of discovering a "best" or most "efficient" algorithm or representation for each context dependent problem, in-depth research needed to be done to understand what can be learned using a widely applicable generic approach. The CESAR effort will, therefore, concentrate on studying learnability for an abstracted version of the representations used in the three most promising inductive methodologies -- genetic algorithm (GA) classifiers, neural nets (NN), and decision trees (DT).

Since the complexity issue has largely been overlooked by others, we propose to make it a central element of our research effort. To address the problem of complexity in representation and scalability, we will draw on the recent seminal work in complexity theory by Garey and Johnson (1979) and its probabilistic interpretation as a theory of learnability proposed by Valiant (1984, 1985) and Blumer, et al., (1989). These works provide a newly developed framework for determining what can be learned with constrained computational resources. These papers have sparked renewed interest in machine learning (see e. g. Haussler & Pitt 1988; Rivest, et al., 1989) by defining a probabilistic basis for approximately learning a wide class of conceptual knowledge. The development of pac-learning (probably approximately correct learning), as this new approach is commonly called, has shed light on ways to solve learning problems that were previously classified as NP-complete.

Despite the considerable theoretical interest generated by Valiant's theory, large-scale implementations for practical discrete- and continuous-space classification problems have not yet appeared. In practice, major discrete-space classifier systems still rely heavily on decision tree (Quinlan 1985), genetic (Holland 1975, 1986a), and stochastic schemes (Duda & Hart 1973). For continuous-space problems, neural networks (particularly the type discussed by Rumelhart 1986) have dominated applications almost entirely.

To bridge this gap and move Valiant's theory toward practical applications, breakthroughs are needed in the algorithmic and representational areas. Our attempt to meet this need, will be to combine pac-learning theory with recent work on uncertain knowledge representation using random set theory (Kendall & Harding 1974; Goodman & Nguyen 1985; Oblow 1987). The relevant random set pac-formalism (Oblow 1990) is quite new, yet it shows promise for utilizing many traditional statistical and measure-theoretical approaches to analyze learnability issues. Its strongest advantage is in providing practical probabilistic measures to gauge learning performance.

Our strong feeling, which we would like to substantiate in this research effort, is that pac-learning with random sets provides a potentially more efficient basis for learning large-scale, sensor-driven, robotics tasks in an advanced computing environment than other current approaches.

4.4.2 Proposed Research

The broad goal of this research program is thus the implementation of Valiant's theory of learnability in a random set framework for use in a robotics environment. Our specific interest is in using Valiant's definition and constructive proof of learnability to develop an efficient computational algorithm for sensor-driven pattern classification applications. This algorithm will be coupled with probabilistic analyses of neural nets, decision trees, and genetic classifiers to cover as wide a spectrum of current learning research areas as possible.

To meet our objectives, we need to develop adaptive, efficient, low-order polynomial algorithms in a random set framework. We propose to accomplish these objectives by using representational transformations, geometric information and statistical methods to find low-order, approximate solutions to hard learning problems. The successful completion of this program should greatly enhance the learning capabilities of machines currently using neural nets, decision trees, and genetic and stochastic algorithms.

In this research effort we will restrict our attention to three specific issues: 1) defining useful classes of low-order polynomial representations, 2) developing algorithms which provably converge to desired performance levels, and 3) improving algorithmic efficiency to levels useful in sensor-driven robotics systems.

In the area of representation, a great deal of effort will be put into establishing classes of practical, low-order, polynomial representations. These classes, will be explored with representational transformations as a means of constructing general learning algorithms. Since these classes can be exhaustively searched rather efficiently, non-convergence with a fixed number of examples can be used to trigger a representational transformation to map this initially hard learning problem into an easier one. Some initial efforts in converting hard GA learning problems into easier ones using affine transformations (Liepins & Vose 1990) give promise to work along these lines. We will need to explore this approach further by estimating how hard a learning problem is while it is being solved. This will allow automated procedures to be developed to determine when to switch representations in order to achieve the best representation for the problem at hand.

A particular class of representations which appears to meet our needs is the class of sets with continuous neighborhoods (a widely applicable representation class for sensor-driven problems). The continuity assumption available in this class allows linear random set algorithms to be developed even for exhaustive search schemes. This approach would have great applicability in visual pattern recognition problems, since use could be made of the already existing geometric neighborhoods of the attribute domain to restrict combinatorial complexity.

In the area of learnability and convergence, our overall approach will be to move away from the current trend of analyzing worst-case pac-learning scenarios. We intend to place more emphasis on average measures of algorithmic performance for particular distributions of learning examples. The emphasis on establishing the complexity order of a learning algorithm with guaranteed probabilistic convergence in worst-case scenarios, has received, we feel, more than enough attention in the literature. More work is needed on adaptive approaches to learning from arbitrary probability distributions of examples, even when these distributions are not known beforehand. The key to this effort will be to establish learnability convergence by on-line estimation techniques.

A major task in this area is to formulate a theory for approximating solutions to a learning problem when the class of representations chosen is inappropriate for even pac-learning with representational transformations. This study is essential if we are going to restrict ourselves to low-order polynomial algorithms. We expect to be able to show limiting, but not necessarily arbitrary, convergence in such cases.

As far as algorithmic efficiency is concerned, the specific technical tasks we intend to pursue all focus on producing pac-learning algorithms in linear or even log-order forms. While this appears to be a formidable task, we feel that many of the stochastic techniques used by genetic algorithms, decision trees and neural nets and other schemes to search unknown concept spaces offer theoretically sound bases for improving algorithmic efficiency. We feel that focusing an exhaustive search on the specific regions of a concept space that are hardest to learn will improve Valiant's approach considerably. Some theoretical work is, however, needed here to both define and estimate hardness in learnability.

The random set approach offers some significant advantages in pursing this goal. Several adaptive approaches to achieving probabilistic convergence that are potentially easy to implement are based on random set statistical theorems. We intent to explore, in particular, the techniques of: Robbins (1968), to estimate unobserved events from observed ones, and Robbins (1944), to estimate directly the probabilistic measure associated with a random set. These approaches, in fact, should allow many general random set problems to be treated analytically.

Another efficiency issue we will study is the use of control strategies for learning from various sources of examples. Since one or another of these sets of examples will produce an easier distribution of examples to learn from, we expect to be able to increase efficiency considerably by being able to choose the easiest learning set as examples are presented. Research is again needed here to find ways to estimate learning difficulty in an online algorithm. In addition, we would like to study the dynamics of such a control strategy. It is possible to obtain complex dynamics when switching from one set of examples to another.

Several other peripheral, yet important issues related to practical implementation of the random set approach might also be explored if time and funding resources permit. One of these would be the use of sigmoid functions to handle continuous attribute problems so that we could directly couple this approach to the one used in neural nets. This implementation would complement and augment the work of Glover (1990) on hybrid rule-based/neural-net learning systems which has generated recent theoretical interest. Another issue would be implementing random set algorithms on a massively parallel computer such as the high performance systems available at CESAR and elsewhere in the Engineering Physics and Mathematics Division at ORNL. The random set methodology appears to be ideally suited for parallelization and the result could be a state-of-the-art approach to scaling up learning theory to the realm of really practical human-scale problems.

4.4.3 References

Baum, E.B. (1990). On learning a union of half spaces. Journal of Complexity, 6, 67-101.

Blumer, A., Ehrenfeucht, A., Haussler, D., & Warmuth, M.K. (1989). Learnability and the Vapnik-Chervonenkis dimension JACM, 6, 929-965.

Duda, R.O. & Hart, P.E. (1973). Pattern classification and scene analysis. New York: John Wiley & Sons.

Garey, M.R. and Johnson, D.S. (1979). Computers and Intractability. New York, NY: W.H. Freeman and Co.

Glover, C.W., M. Silliman, M. Walker, P. F. Spelt and N. S. V. Rao (1990). Hybrid neural network and rule-based pattern recognition system capable of self-modification. SPIE Proceedings 1293, pp 290-300

Goodman, I.R. and Nguyen, H.T. (1985). Uncertainty Models For Knowledge-Based Systems. New York, NY: North-Holland Press.

Grossberg, S. (1982). Studies of Mind and Brain. Boston, MA: D. Reidel Publishing Co.

Haussler, D. & Pitt, L. (Eds.) (1988). Proceedings of the 1988 Workshop on Computational Learning Theory. Boston, MA: Morgan Kaufmann.

Haussler, D. (1988a). Quantifying inductive bias. Artificial Intelligence, 36, 177-221.

Holland, J.H. (1975). Adaptation in natural and artificial systems. Ann Arbor, MI: University of Michigan Press.

Holland, J.H. (1986a). Induction: Processes of Inference, Learning, and Discovery. Cambridge, MA: MIT Press.

Holland, J.H. (1986b). Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R.S. Michalski, et al., (Eds.), Machine Learning: An Artificial Intelligence Approach. Los Altos, CA: Morgan Kaufmann.

Kendall, D.G. (1974). Foundations of a theory of random sets. In Harding, E.F. & Kendall, D.G. (Eds.), Stochastic Geometry. New York, NY: John Wiley & Sons.

Liepins, G.E. & Vose, M.G. (1990). Representational issues in genetic optimization. J. Exp. & Theor. AI, 2, 4-30.

Lippmann, R.P. (1987). An introduction to computing with neural nets. IEEE ASSP Magazine.

Nilsson, N.J. (1980). Principles of Artificial Intelligence. Palo Alto, CA: Tioga Publishing Co.

Oblow, E.M. (1987). O-theory: A hybrid uncertainty theory. Int. J. Gen. Syst., 13, 95-106.

Oblow, E.M. (1990). Implementation of Valiant's theory of learnability using random sets, Technical Report ORNL/TM-11512. Oak Ridge, TN: Oak Ridge National Laboratory.

Pearl, J. (1984). Heuristics: Intelligent Search Strategies For Computer Problem Solving. Reading, MA: Addison-Wesley.

Quinlan, J.R. (1986). Induction of decision trees. Machine Learning, 1, 81-106.

Rivest, R., Haussler, D., & Warmuth, M.K. (Eds.) (1989). Proceedings of the 1989 Workshop on Computational Learning Theory. Santa Cruz, CA: Morgan Kaufmann.

Robbins, H.E. (1944). On the measure of a random set. Annals of Mathematical Statistics, 15, 70-75.

Robbins, H.E. (1968). Estimating the total probability of the unobserved outcomes of an experiment. Annals of Mathematical Statistics, 39, 256-257.

Rumelhart, D.E., & McClelland, J.L. (1986) Parallel Distributed Processing. Cambridge, MA: MIT Press.

Valiant, L.G. (1984). A theory of the learnable. Comm. ACM, 27, 1134-1142.

Valiant, L.G. (1985). Learning disjunctions and conjunctions. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, 560-566. Los Angeles, CA: Morgan Kaufmann.

4.5 Embedded High Performance Computing

Parallel computing has probably been one of the most rapidly progressing technologies during the last decade. Massively parallel hardware systems (> 1000 processors) have become commercially available. They are being used to solve many problems that were previously regarded as intractable because of the number of computations required (e.g., climate modeling, fluid dynamics simulations, etc.).

These applications have in common that algorithms for the approximate solution of the problem at hand are known, and that performance beyond today's supercomputers is required to deal with the enormous amount of computations in reasonable time.

Parallel computing for intelligent machines provides additional and different challenges, since rigorous justifications (i.e., theories) and corresponding algorithms for the solutions of many problems do not exist, e.g., real-time dynamic vision and scene understanding, optimal path planning under various constraints, real-time control of redundant and compliant robots, task planning and allocation.

In contrast to the steady fast pace of development of parallel computing hardware, the advancements in algorithms and software environments that can make optimal use of the new hardware have been sporadic at best and seriously lagging behind. This represents a major obstacle to achieving computing performance close to the hardware design specifications.

The CESAR effort in this area will continue to concentrate on fundamental issues related to concurrent computing on-board mobile robots. These include: parallel algorithms for multi-sensor information processing, e.g. 3-D vision, multi-resolution representations, real-time sensor-based control of mobile robots, and computing environments for distributed intelligent control.

4.5.1 Background

Sensor-driven autonomous robotics presents a variety of computation-intensive problems. Chief among these are the necessity to process and interpret large volumes of sensor data, the requirement to generate commands to mechanical control systems at high frequency, and the desire to integrate these to perform stable real-time, closed-loop control of mechanical systems with large numbers of degrees of freedom. In many cases of interest, this cannot be achieved simply because our computers are not fast enough.

There are fundamentally two strategies for improving performance. The first strategy is to make improvements at the algorithmic level. For example, replacing an $O(n^2)$ sorting algorithm with an $O(n \log n)$ algorithm would improve performance in applications where sorting dominated computation time. The second strategy is to make improvements at the implementation level, by replacing the current computer with a faster one. As long as we restrict our attention to the von-Neumann model of computation, algorithmic issues and implementation issues are largely separate.

However, with the advent of high-performance parallel computers, there has been a proliferation of models of computation [1], and the distinction between algorithmic-level and implementation-level issues has become more complex. To optimize performance we must recognize that the the criteria by which we evaluate algorithms (the expression of algorithmic complexity) must respect the opportunities offered, and the constraints imposed, by computational models. It is only possible to decide upon an optimal algorithm if the underlying model of computation is specified.

Therefore, we propose to address issues which arise specifically in the application of distributed-memory message-passing MIMD parallel processors to sensor-based robotic control. We assume the underlying model of computation is the special case of the "Bulk Synchrony" [2] model supported by available machines, and we regard the systolic model of computation [3] as a special case of message-passing. We recognize that other alternatives, such as SIMD machines [4] and various special-purpose computers, are available, but we will leave the detailed investigation of these alternatives to others.

It has long been recognized that parallel computation offers potentially dramatic performance improvements. A speedup of 1000 means that a computation which would take 10 years on a sequential computer could be done in about 3 1/2 days on a parallel computer. It has been enormously difficult to achieve this potential. The seemingly simple idea of letting 1000 processors work together on a problem gives rise to challenging problems in hardware design and engineering, operating systems, algorithm design and analysis, programming languages, models, and tools, and techniques suitable for various classes of applications. Thus, parallel computing covers the entire spectrum of theoretical and applied computer science, and continues to change and mature at a rapid pace. A comprehensive review would be well beyond the scope of this proposal, but [5] provides a convenient source of information.

Here we will focus on those aspects of parallel computing which are particularly relevant to robotics. Robotics can be distinguished from other endeavors, at least as regards its computing requirements, principally by the need for real-time capabilities, high-bandwidth communications with unusual I/O devices, and hard engineering constraints on space and power supply. In addition to the ordinary criteria of scalability, time optimality, etc., upon which most issues in parallel computing are decided, these extraordinary criteria will help govern our choice of problem selection and guide our evaluation of algorithmic, implementation, and experimental methods.

We note that prior applications of parallel computing to robotics have addressed mainly regular domains (as encountered in low-level computer vision), and synchronous applications. We propose to continue to expand this limited knowledge base by addressing selected crucial issues in problems that lead to dynamic irregular domains, and asynchronous processing.

4.5.2 Proposed Research

In this section we motivate specific issues by proposing a sensor-based robotic control system, considering in detail an architecture for its implementation, and isolating the issues involved for parallel computing. We consider point-to-point motion in a three-dimensional initially unknown environment. This paradigm was chosen for several reasons. It represents in itself an important unsolved problem. It also represents a specific example of a wider class of problems that arise through the integration of perception, reasoning and action in intelligent machines, and results and "lessons learned" are therefore widely applicable. The CESAR experience base in work related to this paradigm is considerable, and experimental facilities crucial to this work exist in the CESAR laboratory.

We make a series of assumptions about the paradigm which are motivated by a desire to conform to the actual performance characteristics of realistic instrumentation (specifically, instruments currently available in the CESAR laboratory: the Odetics laser range camera and the HERMIES-III robot). In particular, we assume a non-point, rigid robot body, the 3D shape of which can be adequately represented as a collection of planar polyhedral surfaces. We assume the number of obstacles, and their shapes, attitudes, and motions are not known in advance. We assume the obstacles have surface properties which are not specular or absorptive at the wavelengths used by the ranging sensor, and that the obstacles cannot move at velocities greater than the maximum velocity of the robot. We assume the robot is equipped with low-level controllers that can accept changes in its commanded translational and rotational velocities at a relatively high frequency (about 30 Hz).

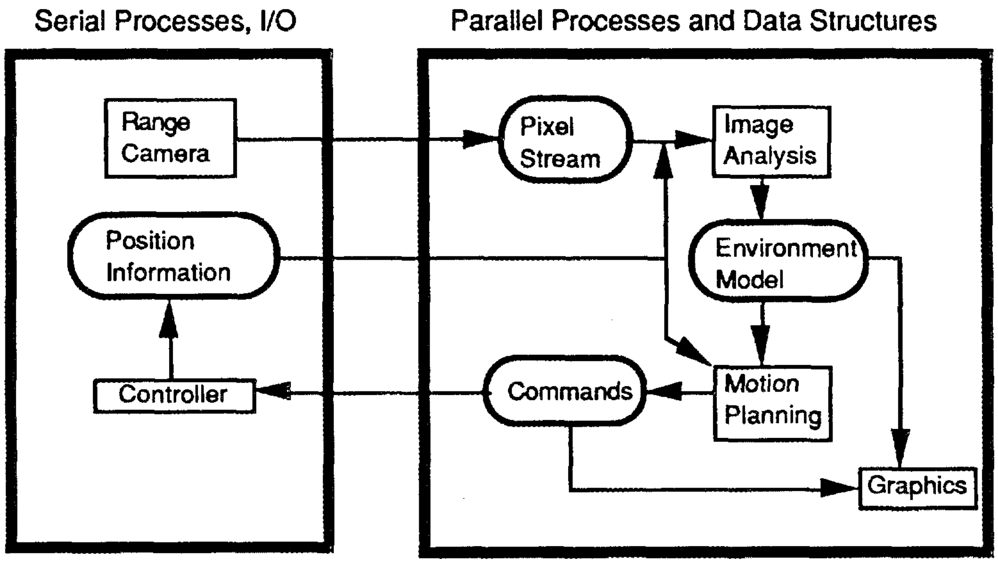

Figure 5: Architecture of a parallel navigation system.

Figure 5 shows a block diagram for a software architecture which can solve this problem. The architecture consists of a number of processes which communicate via a number of data structures. Range data are acquired through a range imaging system, producing a stream of range estimates along with enough positional information about the sensor to calculate the sensor's position in the world coordinate system at the time the data were acquired. A sensor analysis program computes a representation of the robot's immediate surroundings suitable for use by a path planning module, which produces commands to a motion controller. An optional graphics module can display either the environment model, the planned path, or both. Sequential programs for each of these components have been implemented and tested in the laboratory [6].

We consider in greater detail the computational tasks related to this particular instance of integrated perception, reasoning and action. After that, we will consider the implications of performing these tasks with a parallel computer.

We consider first the perception task. We desire to process range sensor data to construct a useful representation of the environment. In this instance, "useful" is interpreted as useful for navigation and path planning algorithms, where it is sufficient to know if a particular volume of space is occupied or unoccupied. A simple scheme which satisfies our utilitarian objective is an occupancy representation. This representation is merely a tessellation of 3-D space into elementary volumes, called "voxels". We characterize the state of a voxel by labelling it as *"Free Space"* ,*"Obstacle"*, or *"Unknown"*. The representation can be encoded as a 3-dimensional array, or as an oct-tree.

The data for this representation can, in principle, be acquired from a variety of sources, but we will consider a range imaging sensor. Traditionally, range imaging sensors are regarded as devices which operate "instantaneously", that is, they acquire an "image" over an arc-segment where the pixel values of the image are proportional to some function of the range from the sensor to the nearest object. We propose to think about the range camera in a slightly different way, specifically, as a device which produces a continuous stream of range values. This method of thinking about the system (a) is consistent with the behavior of other ranging sensors such as sonars, (b) eliminates the problem of compensating for sensor motion during the time required to snap a "frame", and (c) is encouraged by the incorporation in imaging sensors (as in the Odetics camera) of "pixel streaming" modes of operation. In the Odetics camera, the range value associated with each reading is made available to the controlling processor as soon as it is acquired, at a rate of one reading every 32 microseconds.

The computational problem of computing an occupancy representation from range data in a dynamic environment can then be formulated as a ray-tracing calculation. We assume that the "pixels" of the incoming data stream have associated corollary data sufficient to specify a ray in the world coordinate system. One endpoint of the ray is rooted at the location in world coordinates occupied by the sensor, the other on the surface of the sensed object.

With the above description of the problem we can now describe the fundamental issues for parallel computing. It is our fundamental objective in any parallel program to balance the computation while minimizing the communication. With the traditional way of looking at range images, it is readily apparent how one would construct a parallel program to do 3D ray-tracing into a volume. First we "filter" the image to throw out any bad values. Then we "balance" the computation across all of the processors by equalizing the sum of the lengths of all the vectors (this is most readily accomplished by sorting by length and preferentially communicating the shortest-length rays). Then we do the ray-tracing proper without any communication. Finally we communicate any voxels to the appropriate destination processor (determined by the domain decomposition strategy we have chosen for the environmental representation). This is best accomplished on hypercubes using the "index" algorithm of Fox and Furmanski [7].

However, regarding the imaging system as a source of a pixel stream, it is obvious that the above strategy is irrelevant. That is, not all of the rays are available at any given time, so it is impossible to construct synchronous communications routines for the load balancing phase and the final communications phase. Instead, at any given time there will be a number of "pixels" (data structures consisting of a range value, a direction in which the range reading was taken, and a position and the instantaneous extrinsic calibration data of the sensor) which are ready for prefiltering, a number of pixels ready for coordinate system transformation, a number of rays ready for clipping, and a number of rays ready for tracing. The bulk of the computational work is associated with tracing the rays, and the amount of work for each ray is proportional to its length. The amount of work to be done is therefore not known until the rays have been clipped, and changes from moment to moment.

This situation suggests that a load balancing strategy is necessary. It is insufficient to use traditional load balancing techniques for this purpose [8]. Static load balancing techniques [9] assume that the amount of work to be done is constant as a function of time. Dynamic load balancing techniques do not make such an assumption, therefore we must use dynamic methods. It is possible to decompose the processing necessary for each ray into a sequence of steps where the computation time required for each step is known immediately prior to executing the step. Therefore is it possible to maintain a queue of work to be done on each processor, and (crudely) by adding up the contents of the queue determining the "workload" of each processor. It is then possible to equalize the workload on all of the processors by allowing the work from the most heavily loaded processors to "diffuse" into the queues of the more lightly loaded processors.

At this point, we have not yet specified the embedding of the occupancy representation into the parallel processor. The decision on this embedding can be made only after considering the topology of the interconnection network, the computation to communication ratio, the dimensionality of the data structure, and, most importantly, the properties of the downstream process which will consume the representation. It is impossible to state what the optimal domain decomposition is without simultaneously stating the algorithm which uses the decomposition.

If we were considering sensor data analysis only, we could be content if the results were available in an arbitrary location in the memory of the parallel processor. However, as we discuss integrating sensor data analysis with reasoning, i.e. path planning in this specific case, we need to assure that the information is readily available to effect planning and control at sensor frame rates.

A number of unresolved generic issues for parallel processing will have to be addressed: What is the best way to inject the workload into the parallel processor, as it is generated by the sensor? We could conceive of a "master" processor which does the prefiltering and assigns work to slaves. This strategy would likely create a bottleneck at the master processor, and as the number of processors grows would starve the array. We could distribute the incoming data stream over all of the processors, but we run the risk of consuming extra communication bandwidth unnecessarily. Can a "multiple masters" policy be both scalable and efficient? If so, what variables determine the optimal ratio of masters and slaves and how does one assign masters and slaves to processors in the network so that communication bandwidth is conserved?

What is the best scheduling for dynamic load balancing when the amount of work is known (after clipping)? If we consider only ray-tracing as a source of work (and it is the dominant source) we can maintain a queue on each processor and allow the work from the more heavily loaded processors to "diffuse" into neighboring more lightly loaded processors. Only synchronous non-adaptive strategies have been actually applied to date. One can easily imagine asynchronous adaptive strategies which generate messages from interrupts to determine the load and to balance the load.

What is the best dynamic load balancing strategy for distributed multiple queues? Each queue contains different kinds of work? Should we rely on round-robin scheduling, or is there a better way? Note that this problem does not necessarily reduce to the problem of single queues because naive treatment of the contents of each queue may result in unnecessary data movement.

As the steps in N-DOF path planning [e.g. 10] include distance transformations and searches in configuration space to optimize certain objective functions, the following fundamental issues emerge: What is the best domain decomposition and load balancing strategy for distance transformations on distributed memory machines? Is there an incremental algorithm for distance transformations which would eliminate the necessity of doing complete distance transformation at every step? We already know that chamfering algorithms will probably not parallelize well. What is the best algorithm for searching when it is impossible to expand each node with the information available only on a single processor. What is the best way to do collision detection when the necessary data is distributed? Where should the data be distributed? What is the best way to plan paths in an environment where we can assume an "anytime" controller. The hard constraint on the controller is that we must have a motion command available at the frequency that the controller requires its input. A traditional "obstacle avoidance" algorithm is probably not the best we can do, since it does not make full use of the available information, yet we probably do not have to plan a complete path at every iteration. We are looking for a method which will at each iteration choose the "best" command to send to the controller given the available information.

We will implement this control system on an Intel iWarp concurrent computer. The iWarp is a mesh-connected, distributed-memory parallel processor with systolic processing capability. Each processor has a peak performance of 10 Mflops, and 1.5 Mbyte memory. The system was selected based on performance, integrated systolic processing capability, compatibility with existing VME-based computer systems at CESAR, and expected sufficient operating system support for the proposed research. The iWarp is being procured in FY 1991.

4.5.3 References

1. R. Duncan, "A survey of parallel computer architectures", IEEE Computer, 23, 5-17, 1990.

2. L.G. Valiant, "A bridging model for parallel computation", Comm. ACM, 33, 103-111,1990.

3. H.T. Kung, "Computational models for parallel computers", Phil. Trans. R. Soc. Lond. A, 326, 357-371, 1988.

4. W. D. Hillis, "The Connection Machine", MIT Press, Cambridge, MA, 1985.

5. D. W. Walker and Q. F. Stout, Proceedings of the 5th Distributed Memory Computing Conference, IEEE Computer Society Press, Washington DC, 1990.

6. C. S. Anderson, C. B. Madsen, J. J. Sorensen, N. O. S. Kirkeby, J. P. Jones, and H. I. Christensen, "Navigation using range images on a mobile robot", Computers & Electrical Engineering, in press, 1991.

7. G. C. Fox and W. Furmanski, "Optimal communications algorithms for regular domain decompositions on the hypercube", Proc. 3rd Conference on Hypercube Concurrent Computers and Applications, 648-713, 1988.

8. G. C. Fox, M. A. Johnson, G. A. Lyzenga, S. W. Otto, J. K. Salmon, D. W. Walker, "Solving problems on concurrent processors", Prentice-Hall, 1988.

9. K. M. Dragon and G. L. Gustafson, "A low-cost hypercube load-balance algorithm", Proc. 4th Conf. on Hypercube Concurrent Computers and Applications, 583-589, 1989.

10. J. C. Barraquand, J. C. Latombe, "Robot motion planning: a distributed representation approach", Stanford University Technical Report, 1989.

4.6 Intelligent Machine Research Facilities

Excellent experimental facilities have played a major role in the success of the CESAR effort. Continued development of and experiments with the HERMIES robots provide focus for the research, and allow to perform rigorous testing and evaluation of new concepts in a state-of-the-art laboratory testbed.

The HERMIES robots have evolved from a relatively simple mobile platform equipped only with sonar sensors with all computing performed off-board to the multi-purpose mobile robot HERMIES-III, equipped with a redundant manipulator arm, multiple sensors (CCD cameras, sonar, laser range finder), and an on-board multi-computer system that includes a hypercube multi-processor. HERMIES-III is designed to support a second manipulator arm, and additional sensors (e.g. a thermal infrared camera). It has omni-directional steering capability, and is designed to hold batteries that allow it to operate as a completely self-contained, autonomous system for about one hour.

With appropriate modifications, up-grades, and additions to the computer systems, sensor suite, and manipulator(s), HERMIES-III can continue to serve as a uniquely powerful experimental facility for the period covered in this program plan.

HERMIES-IIB has been used for experiments longer than its initial design called for. It is expected that rising incidence of failures and associated repairs will make it necessary to schedule a phase out of HERMIES-IIB during FY 1992. HERMIES-IIB and -III will be used for initial experiments related to the CESAR research in cooperating robots.

A new experimental platform based on an innovative wheel design is currently being tested at CESAR [II.77]. This new development can form the basis for a HERMIES-IV robot, which will represent a replacement for HERMIES-IIB during the second half of the period covered in this program plan. This new robot will further support the work in cooperating agents, and distributed intelligent control.

Proposed acquisitions include a commercially available redundant robot manipulator. This manipulator will replace CESARm on-board HERMIES-III. CESARm will continue to be used as a stationary system in dual-arm experiments with the new arm on HERMIES-III. We propose to replace CESARm later with a new modular robot manipulator.

CESAR was a beta-test site for the first generation NCUBE hypercube computer. This technology has progressed very rapidly, and the current hypercube systems at CESAR represent out-dated technology. The use of these systems will diminish during the period covered by this proposal, and the new iWarp computer (acquisition of initial configuration during FY 1991) will be the main parallel computer system at CESAR. In addition, the present CESAR computer network will have to be upgraded with new workstations as this technology progresses.

ORNL resources have been approved for the design during FY 1992 of a new building, the Intelligent Machine Research Facility, that will provide office space and a high-bay area to allow for safe and expanded experiments with the CESAR mobile robot systems.

5. Milestones and Estimated Costs

Research Areas	FY 1992	FY 1993	FY1994	FY 1995	FY1996
4.1 Cooperating Robots 1. Multiple mobile agents 2. Cooperating manipulators	*92.1a* *92.1b*	*93.1a* *93.1b*	*94.1a* *94.1b*	*95.1a* *95.1b*	*96.1a* *96.1b*
Funding (K$)/Person-years	350/1.8	400/2.0	420/2.0	440/2.0	460/2.0
4.2 Combined Mobility and Manipulation	*92.2*	*93.2*	*94.2*	*95.2*	*96.2*
Funding (K$)/Person-years	190/1.0	200/1.0	210/1.0	220/1.0	230/1.0
4.3 Multi-Sensor Data Analysis and Fusion	*92.3*	*93.3*	*94.3*	*95.3*	*96.3*
Funding (K$)/Person-years	285/1.5	360/1.8	420/2.0	550/2.5	460/2.0
4.4 Machine Learning	*92.4*	*93.4*	*94.4*	*95.4*	*96.4*
Funding (K$)/Person-years	190/1.0	200/1.0	210/1.0	220/1.0	230/1.0
4.5 Embedded High Performance Computing	*92.5*	*93.5*	*94.5*	*95.5*	*96.5*
Funding (K$)/Person-years	285/1.5	300/1.5	315/1.5	330/1.5	460/2.0
4.6 Intelligent Machine Research Facilities	*92.6*	*93.6 a* *93.6 b*	*94.6*	*95.6*	*96.6*
Funding (K$)/Person-years Capital Equipment	95/0.5 380	100/0.5 200	105/0.5 150	110/0.5 150	115/0.5 200
Total Resources (K$) Operating Capital Person-years	 1395 380 7.3	 1560 200 7.8	 1680 150 8.0	 1870 150 8.5	 1955 200 8.5

Table 3: Resource plan; no. s in italics refer to milestones on next pages.

List of proposed milestones: (J : publication; S: computer simulation; E: robot experiment)

92.1a Speed of motions versus sensory data analysis tradeoffs analysis (J, E)
93.1a Decentralized control of weak cooperation (J)
94.1a Planning of time-optimal motions for classes of exploration tasks (J, S)
95.1a Cooperative task planning for site assessment and mapping tasks (J, E)
96.1a Cooperative performance of site assessment tasks by heterogeneous robots (J, E)

92.1b Simultaneous force- and position-control of soft contact constrained motion of single redundant robotic manipulator (E, J)
93.1b Dynamic computer simulation of decoupled control architecture method for controlling two interacting manipulators (S, J). Experiments on hard contact motion of redundant manipulator (E, J).
94.1b Physical and quantitative interpretation of redundant DOFs in multi-manipulator closed chains (S, J).
95.1b Experimental evaluation of position- and internal stress force-control methods for two manipulators holding a rigid object, using the new CESAR dual arm system (E, J).
96.1b Continuation of laboratory experiments on two interacting manipulators, with emphasis on handling complex, jointed payloads (E, J).

92.2 Demonstrate redundancy resolution with secondary criteria, e.g. obstacle avoidance, torque minimization (E, J)
93.2 Configuration optimization with respect to reach and dexterity (J, E)
94.2 Sensor-based planning and control of CM&M system (E)
95.2 Demonstrate CM&M capability for manipulation of valve in tight workspace (E, J)
96.2 Demonstrate CM&M system for material handling tasks (E, J)

92.3 Adaptive fusion of laser range/reflectance, sonar and vision data to effect safe robot motion in unstructured environment (J, E)
93.3 Fusion of sensor data for self-location and exploration (J, E)
94.3 Integration of color image sensor into fusion system (J, E)

95.3 Sensor fusion for site assessment and characterization by multiple robots (J, E)

96.3 Active control of multiple robots based on sensory task requirements (J, E)

92.4 Comparative analysis of performance of RST, GA, and NN learning approaches (J, S)

93.4 Control strategies for learning from various sources of examples (J, S)

94.4 Implementation and evaluation of most efficient learning approach in multi-robot motion planning scenario of constrained, known complexity (E, J)

95.4 Develop criteria for choices of representations to effect most efficient learning (J, S)

96.4 Evaluation of algorithms for learning under time constraints (J, S)

92.5 Initial testing of elements of parallel navigation system using iWarp parallel computer (S, J)

93.5 Implementation and testing on HERMIES-III (E, J)

94.5 Generalization to control of N DOF robots (S, J)

95.5 Evaluation of generalization in experiments with mobile manipulator (E)

96.5 Generalization to coordinated control of multiple robots (J, E)

92.6 Acquisition of dual redundant manipulator system

93.6a Integration of HERMIES-IV testbed

93.6b Acquisition of color imaging sensor

94.6 iWarp-equivalent expansion of on-board parallel processor

95.6 Procurement of new laser range imaging sensor

96.6 Procurement of new modular robot manipulator

Appendix A

Vitae of CESAR Principal Investigators

Martin Beckerman

Martin Beckerman has been a member of the Engineering Physics and Mathematics Division (EPMD) at the Oak Ridge National Laboratory (ORNL) since 1987. He received the B.S. and M.S. degrees in physics from the University of Florida in 1964 and 1966, respectively, and the Ph.D. degree in physics from the University of Miami in 1970. Prior to joining EPMD he was an Associate Professor of Physics at the University of Tennessee through an appointment with the Joint Institute for Heavy Ion Research at Oak Ridge, and was a Research Scientist and Principal Investigator in the Laboratory for Nuclear Science of the Massachusetts Institute of Technology. He has conducted research at the University of Rochester and the Weizmann Institute of Science, and served as a consultant to the Physics Division at ORNL and to the Kellogg Radiation Laboratory at Caltech. He has authored over 80 publications in robotics, and in experimental and theoretical nuclear physics. His current research interests include adaptive sensing, sensor data processing and sensor fusion.

Selected Recent Publications

M. Beckerman and D. L. Barnett, "Performance of Visual and Ultrasound Sensing by an autonomous Robot," Oak Ridge National Laboratory Technical Memo ORNL/TM-11733 (CESAR-90/54), January, 1991, and to be submitted to the IEEE Transactions on Robotics and Automation.

M. Beckerman, L. A. Farkas and S. E. Johnston, "Treatment of Systematic Errors II: Fusion of Ultrasound and Visual Sensor Data," submitted for publication to the IEEE Transactions on Robotics and Automation, July, 1990. (CESAR-90/45).

M. Beckerman and E. M. Oblow, "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor Data for Robotic Navigation," IEEE Transactions on Robotics and Automation, Vol. RA-6, pp. 137-145, 1990.

M. Beckerman, D. L. Barnett, M. Dickens and C. R. Weisbin, "Performance of Multiple Tasks by an Autonomous Robot Using Visual and Ultrasound Sensing, " Robotics and Manufacturing, Vol. 3, pp. 389-395, 1990.

M. Beckerman, D. L. Barnett and S. M. Killough, "Ultrasound and Visual Sensor Feedback and Fusion," in Proceedings of the 1990 Japan-U.S.A. Symposium on Flexible Automation, Kyoto, Japan, July 9-13, 1990, Vol. III, pp. 1315-1319, 1990.

M. Beckerman, L. A. Farkas, J. P. Jones, R. C. Mann and C. W. Glover, "World Modelling and Multi-Sensor Integration for a Mobile Robot," in Proceedings of the ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, March 13-16,1989, 11-2, pp. 1-8, 1989. (CESAR-88/42).

M. Beckerman, J. P. Jones, R. C. Mann, L. A. Farkas and S. E. Johnston, "Spatial Reasoning in the Treatment of Systematic Sensor Errors," in Proceedings of the SPIE Conference on Sensor Fusion: Spatial Reasoning and Scene Interpretation, Cambridge, MA, pp. 338-349, 1988. (CESAR-88/25).

Kikuo Fujimura

Education

BS	1983, University of Tokyo, Information Science
MS	1985, University of Tokyo, Information Science
PhD	1989, University of Maryland, Computer Science

Research Interests

planning in dynamic environment,
sensory data processing for robotics,
distributed artificial intelligence,
data structures, computer graphics.

Professional Activities

Session chairman, 1991 IEEE Int. Conf. of Robotics and Automation
Reviewer, IEEE Transactions on Robotics and Automation,
Reviewer, IEEE Transactions on System, Man, and Cybernetics,
Member of ACM, IEEE, AAAI

Selected Recent Publications:

Motion planning in dynamic environments.
Monograph to be published by Springer-Verlag, 1991.

A reactive model for multiple mobile agents. to appear in Proceedings of the IEEE Int. Conf. of Robotics and Automation, 1991.

Motion planning in a dynamic domain. Proc. of the IEEE Int. Conf. of Robotics and Automation, Cincinnati (Ohio) 324-330, 1990

A hierarchical strategy for path planning among moving obstacles IEEE Transactions on Robotics and Automation 5, 1 61-69, 1989

Judson P. Jones

Judson P. Jones received a B.S. from Vanderbilt University in 1980, and a Ph.D. in Anatomy from the University of Pennsylvania in 1985. His work centered on experimental electrophysiological methods for multidimensional linear and non-linear systems identification in the mammalian visual system. He joined the Center for Engineering Systems Advanced Research in 1986, where his research interests expanded to include signal processing and image analysis, robot motion planning, algorithms and environments for parallel computing, and software systems integration. He is the principal investigator for concurrent computing and computer vision in CESAR, and since 1988 has served as the technical team leader for software systems integration in the DOE's Robotics for Advanced Reactors project, a university/industry/national laboratory collaboration. He has published diversely in parallel computing, computer vision, visual neurophysiology, and linguistics. He is a member of IEEE, IEEE-CS, IEEE-R&A, and ACM.

Selected Recent Publications

J.P. Jones, "On the design of a concurrent image analysis system." Proc. 3rd International Symposium on Robotics and Automation, July 18-20, Vancouver, B.C., pp. 205-210, 1990

H. Embrechts, J.P. Jones ,"An input-output algorithm for M-dimensional rectangular domain decompositions on N-dimensional hypercube multicomputers." Proc. 5th Conference on Distributed Memory Concurrent Computers, D.W. Walker & Q.F. Stout, ed.s, pp. 876-882, 1990.

J.P. Jones, M. Beckerman, R.C. Mann, "Design and implementation of two concurrent multi-sensor integration algorithms for mobile robots." Sensor Fusion II: Human and Machine Strategies, P.S. Schenker, ed., Proc. SPIE 1198, pp. 301-312, 1990.

H. I. Christensen, J. P. Jones, "Concurrent multi-resolution image analysis", Proceedings of the 4th Conference on Hypercubes, Concurrent Computers, and Applications, Monterey, CA, 1989.

J. P. Jones, "A concurrent on-board vision system for a mobile robot", Proceedings of the 3rd Conference on Hypercubes, Concurrent Computers, and Applications, pp 1022-1032, 1988.

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, 2, pp18-27, 1987.

J. P. Jones, L. A. Palmer, "The two-dimensional spatial structure of simple receptive fields in cat striate cortex", Journal of Neurophysiology, 58, 1187-1121, 1987.

J. P. Jones, R. A. Stepnoski, L. A. Palmer, "The two-dimensional spectral structure of simple receptive fields in cat striate cortex", Journal of Neurophysiology, 58, 1212-1232, 1987.

J. P. Jones, L. A. Palmer, "An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex", Journal of Neurophysiology, 58, 1233-1258, 1987.

Reinhold C. Mann

Reinhold C. Mann received a Diplom-Mathematiker degree (M.S. in mathematics) in 1977, and a Dr. rer. nat. degree (PhD) in physics in 1980 from the Johannes Gutenberg University in Mainz, Federal Republic of Germany. From 1978 until 1980 he was a research associate in the Biophysics Department at Mainz University, and a consultant with the Laser and Optics Group at Battelle Institute in Frankfurt, F.R.G., in the areas of digital image analysis and pattern recognition. In 1980 he joined the Image Analysis Group at the Fraunhofer Institute for Data and Information Processing in Karlsruhe, F.R.G.

He was awarded a Feodor-Lynen Fellowship by the Alexander von Humboldt Foundation in Bonn, F.R.G., which allowed him to spend 1981 and 1982 as a Visiting Scientist in the Biology Division at Oak Ridge National Laboratory (ORNL), working on biomedical applications of pattern recognition and image analysis. He was a staff member in the Biology Division from 1983 until 1986 when he joined the Engineering Physics and Mathematics Division to work on multi-sensor systems for intelligent machines and mobile robots. He was leader of the Advanced Computing and Integrated Sensor Systems Group from 1987 until 1989, and has been head of the Intelligent Systems section and Director of the ORNL Center for Engineering Systems Advanced Research (CESAR) since 1989. He is an Adjunct Associate Professor in the Computer Science Department at the University of Tennessee in Knoxville. His research interests include computer vision, multi-sensor integration, pattern recognition, and concurrent computing. He is a member of IEEE, ANS, ACM, AAAS, and the International Neural Network Society.

Selected Recent Publications

R. C. Mann , "Multi-sensor integration using concurrent computing", SPIE 782, pp 83-90, 1987.

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

D. K. Wehe, J. C. Lee, W. R. Martin, R. C. Mann, W. R. Hamel, "Intelligent robotics and remote systems for the nuclear industry", Nuclear Engineering and Design, 113, pp 259-267, 1989.

H. P. Hiriyannaiah, G. L. Bilbro, W. E. Snyder, R. C. Mann, "Restoration of piece-wise constant images by mean field annealing", Journal of the Optical Society of America, 6 (12), pp 1901-1912, 1989

G. L. Bilbro, W. E. Snyder, R. C. Mann, "The mean field approximation minimizes relative entropy", submitted to Journal of the Optical Society of America, 1990.

Edward M. Oblow

Education

B.S., Engineering, City College of New York, 1965
M.S. and Ph.D., Engineering, Columbia University, 1970

Research Interests

Sensitivity and uncertainty theory, computer calculus, artificial intelligence, learning theory, parallel computing

Selected Recent Publications

E. M. Oblow, "O-theory - a hybrid uncertainty theory", Int. J. of Gen. Syst. 13, pp 95, 1986.

E. M. Oblow, "Supertracks, supertrack functions and chaos in the quadratic map", Physics Letters a, 128, 8, 1988.

E. M. Oblow, "Foundations of O-theory: measurements and relation to fuzzy set theory", Int. J. of Gen. Syst. 14, PP357, 1988.

M. Beckerman and E. M. Oblow, "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor Data for Robotic Navigation," IEEE Transactions on Robotics and Automation, Vol. RA-6, pp. 137-145, 1990.

E. M. Oblow, "Implementation of Valiant's theory of learnability using random sets", Technical Report ORNL/TM-11512. Oak Ridge, TN: Oak Ridge National Laboratory, 1990.

Francois G. Pin

Dr. Francois G. Pin is the Group Leader of the Autonomous Robotic Systems Group at the Oak Ridge National Laboratory and a principle investigator of the Center for Engineering Systems Advanced Research (CESAR) Program. He earned the Maitrise de Mecanique (1976) from the Universite de Nancy 1, France, and the Diplome National d'Ingenieur Electro-mecanicien (1977) from the Ecole Nationale Superieure d'Electricite et Mecanique de Nancy, France. His M.S. (1978) and Ph.D. (1982) in Mechanical Engineering and Aerospace Sciences are from the University of Rochester, New York. He joined the research staff of the Oak Ridge National Laboratory in 1982 at which time he conducted the major part of his research work in the area of Mathematical and Numerical Modeling. His current research work and interests include Methodologies for Intelligent Machines and in particular for the Planning, Reasoning, Learning, and Decision Making of Autonomous Mobile Robots, Manipulation Systems, and Man-Machine-Synergistic Systems. He is on the Editorial Board of the International Journal of Robotics and Mechatronics and the Japanese Journal of Advanced Automation Technology, and is a guest editor of the Computers and Electrical Engineering Journal.

Selected Recent Publications

Pin, F. G. and J.-C. Culioli, "Optimal Positioning of Combined Mobile Platform-Manipulator Systems for Material Handling Tasks," submitted to Journal of Robotic Systems.

Pin, F. G., P. F. R. Belmans, S. Hruska, C. Steidley, and L. E. Parker, "Robot Learning from Distributed Sensory Sources," accepted for publication in IEEE Transactions on Systems, Man, and Cybernetics (1991).

Vasseur, H. A., F. G. Pin, and J. R. Taylor, "Navigation of Car-Like Mobile Robots in Obstructed Environments Using Convex Polygonal Cells," accepted for publication in Computers and Electrical Engineering 17(3) (1991).

Iyengar, S. S., A. S. Sabharwal, F. G. Pin, and C. R. Weisbin, "Asynchronous Production System for Control of an Autonomous Mobile Robot in Real Time

Environment," accepted for publication in Journal of Applied Artificial Intelligence (1991).

Pin, F. G. and J.-C. Culioli, "Optimal Positioning of Redundant Manipulator-Platform Systems for Maximum Task Efficiency" in Robotics and Manufacturing, Vol. 3, eds. M. Jamshidi and M. Saif, ASME Press, 1990, pp. 489-495.

Killough, S. M. and F. G. Pin, "A Fully Omnidirectional Wheeled Assembly for Robotic Vehicles," Trans. Am. Nucl. Soc. 61 425-426 (1990).

Vasseur, H. A. and F. G. Pin, "Trajectory Generation for Car-Like Robots," Trans. Am. Nucl. Soc. 61 421-422 (1990).

Weisbin, C. R., G. de Saussure, J. R. Einstein, F. G. Pin, and E. Heer, "Autonomous Mobile Robot Navigation and Learning," IEEE Computer 22(6), 29-35 (June 1989).

Sabharwal, A., S. S. Iyengar, C. R. Weisbin, and F. G. Pin, "Asynchronous Production Systems," Knowledge-Based Systems 2(2), 117-127 (1989).

Spelt, P. F., G. de Saussure, E. Lyness, F. G. Pin, and C. R. Weisbin, "Learning by an Autonomous Robot at a Process Control Panel," IEEE Expert 4(4), 8-16 (1989).

Weisbin, C. R., W. R. Hamel, D. P. Kuban, S. A. Meacham, and F. G. Pin, "The Robotics and Intelligent Systems Program at ORNL," Robotica 7, 101-111 (1989).

Michael A. Unseren

M.A. Unseren received the PhD degree in Electrical Engineering from Purdue University in August 1989, the Master of Engineering degree from the University of South Florida in April 1984, and the B.S.E.E. degree in Electrical Engineering from the Illinois Institute of Technology in August 1989. Dr. Unseren has been a robotics research engineer at the Oak Ridge National Laboratory since October 1989. His research interests include coordinated multiple manipulators, constrained (hard contact) motion of manipulators, and kinematically redundant robots. He is currently a member of the IEEE.

Selected Recent Publications

M.A. Unseren, "A Rigid Body Model and Decoupled Control Architecture for Two Manipulators Holding a Complex Object," accepted for Publication in Computers and Electrical Engineering journal, 1991.

M.A. Unseren, "Rigid Body Dynamics and Decoupled Control Architecture for Two Strongly Interacting Manipulators", accepted for Publication in Robotica, 1991.

M.A. Unseren and A.J. Koivo, "Reduced Order Model and Decoupled Control Architecture for Two Manipulators Holding an Object," 1989 IEEE International Conference on Robotics and Automation. Scottsdale, AZ, pp 1240-1245.

M.A. Unseren and A.J. Koivo, "Kinematic Relations and Dynamic Modeling for Two Cooperating Manipulators in Assembly," 1987 IEEE International Conference on Systems, Man, and Cybernetics. Vol. 1, pp 798-802. Alexandria,VA.

Appendix B

Guest Researchers at CESAR, 1985 - 1991

University/Industry Research Collaborators or Consultants

A. Scientific Staff

G. L. Bilbro, North Carolina State University, Raleigh, NC
W. J. Book, Georgia Institute of Technology, Atlanta, GA
A. Cossic, CEA-CEN, Saclay, France
E. Depiante, MIT, Cambridge, MA
E. Heer, Heer Associates, LaCanada, CA
S. S. Iyengar, Louisiana State University, Baton Rouge, LA
D. Kincaid, University of Texas, Austin, TX
A. I. Khuri, University of Florida, Gainesville, FL
B. Oldham, Texas Technical University, Lubbock, TX
N. Rao, Old Dominion University, Norfolk, VA
W. E. Snyder, North Carolina State University, Raleigh, NC
S. L. Tosunoglu, University of Texas, Austin, TX
H. Watanabe, University of North Carolina, Chapel Hill, NC

B. Technical Support

D. M. Jollay, The University of Tennessee, Knoxville, TN

C. Graduate Students

H. I. Christensen, Aalborg University, Aalborg, Denmark
H. D. Embrechts, Katholieke Universiteit, Leuven, Belgium
J. Graham, Louisiana State University, Baton Rouge, LA
S. Gulati, Louisiana State University, Baton Rouge, LA
V. G. Hedge, Louisiana State University, Baton Rouge, LA
H. P. Hiryannaiah, North Carolina State University, Raleigh, NC
A. S. Sabharwal, Louisiana State University, Baton Rouge, LA
R. S. Sawhney, The University of Tennessee, Knoxville, TN

D. Great Lakes College Association Faculty Research Participation

D. S. Christiansen, Albion College, Albion, MI
R. R. Winters, Denison University, Granville, OH

E. Oak Ridge Associated Universities - Faculty Research

K. Bowyer, University of South Florida, Tampa, FL
C. R. Dalton, University of Florida, Gainesville, FL
P. F. Spelt, Wabash College, Crawfordsville, IN
C. W. Steidley, Central Washington University, Ellenburg, WA

F. Oak Ridge Associated Universities - Post-Doctoral Program

P. F. R. Belmans, Institut National de Recherche en Informatique et en Automatique, Creeil, France
J. F. Culioli, Schlumberger Montrouge Research Park, France
J. H. Han, University of Maryland, College Park, MD
N. Toomarian, Technion, Israel

G. Oak Ridge Associated Universities - Undergraduate Research Program

D. P. Katzman, University of Pennsylvania, Philadelphia, PA
J. L. Painter, East Tennessee State University, Johnson City, TN

H. Oak Ridge Associated Universities - Science & Engineering Research

T. K. Arledge, Ohio Northern University, Ada, OH
V. K. Bansal, Cornell University, Ithaca, NY
R. A. Cacheiro, Tennessee Technological University, Cookeville, TN
N. D. Drochak, Miami University, Oxford, OH
L. A. Farkas, Cornell University, Ithaca, NY
K. H. Gaiser, Michigan Technological University, Houghton, MI
T. Heywood, University of Florida, Gainesville, FL
A. M. Jijina, Illinois Institute of Technology, Chicago, IL
E. S. Lyness, Iowa State University, Ames, IA
M. S. Madanat, Cheyney University of Pennsylvania, Cheyney, PA
K. L. Marrs, University of Missouri-Rolla, Rolla, MO
P. Marquis, Tufts University, Medford, MA
J. W. Palmer, State University of New York-Buffalo, Buffalo, NY
J. A. Sandberg, Utah State University, Logan, UT

I. Oak Ridge Associated Universities - Practicum Assignment

S. I. Hruska, Jacksonville State University, Jacksonville, AL

J. Oak Ridge Associated Universities - Professional Internship Program

M. S. Silliman, University of Missouri-Rolla, Rolla, MO

K. U. S. Military Service Academy Research Associates

J. R. Dundas, U. S. Naval Academy, Annapolis, MD
D. J. Durn, U. S. Naval Academy, Annapolis, MD
T. M. Mathis, U. S. Naval Academy, Annapolis, MD
C. A. McGougan, U. S. Military Academy, West Point, NY
D. McVay, U. S. Military Academy, West Point, NY

Appendix C
CESAR Publication List

CENTER FOR ENGINEERING SYSTEMS ADVANCED RESEARCH
PUBLICATIONS LIST

I. BOOKS AND REFEREED JOURNALS

I.1. S. S. Iyengar, C. C. Jorgensen, S. V. N. Rao and C. R. Weisbin, "Robot Navigation Algorithms Using Learned Spatial Graphs," ROBOTICA, Vol. 4, pp. 93-100 (1986).

I.2. E. M. Oblow, "O-Theory: A Hybrid Uncertainty Theory," International Journal of General Systems, 13, 2, pp. 95. CESAR-86/02.

I.3. J. Barhen and J. F. Palmer, "The Hypercube in Robotics and Machine Intelligence," Computers in Mechanical Engineering, 4, No. 5, pp. 30-38 (March 1986). CESAR-86/03.

I.4. C. C. Jorgensen, W. R. Hamel and C. R. Weisbin, "Exploring Autonomous Robot Navigation," BYTE, pp. 223-235 (January 1986). CESAR-86/04.

I.5. E. M. Oblow, "A Probabilistic-Propositional Framework for the O-Theory Intersection Rule," Int. Journal of General Systems, Vol 13, pp. 187-201 (October 1987). CESAR-86/10.

I.6. J. Barhen, "Hypercube Ensembles: An Architecture for Intelligent Robots," in Special Computer Architectures for Robotics and Automation, J. Graham, Editor, Chapter 8, pp. 195-236, Gordon and Breach, New York (1987). CESAR-86/11.

I.7. E. M. Oblow, F. G. Pin and R. Q. Wright, "Sensitivity Analysis Using Computer Calculus: A Nuclear Waste Isolation Application," Nuclear Science and Engineering, Vol. 94-1, pp. 46-65 (September 1986). CESAR-86/19.

I.8. J. Barhen and E. C. Halbert, "ROSES: An Efficient Scheduler for Precedence - Constrained Tasks on Concurrent Multiprocessors," Hypercube Multiprocessors, '86, M. T. Heath, ed., Chapter 11, pp. 123-148, SIAM, Philadelphia, PA, (1986). CESAR-86/21.

I.9. C. R. Weisbin, G. de Saussure and D. W. Kammer, "Self-Controlled: A Real-Time Expert System for an Autonomous Mobile Robot," Computers in Mechanical Engineering, Vol. 5, 2, pp. 12-19 (September 1986). CESAR-86/25.

I.10. S. V. N. Rao, S. S. Iyengar, C. C. Jorgensen and C. R. Weisbin, "Robot Navigation in an Unexplored Terrain," Journal of Robotic Systems, Vol. 3, 4, pp. 389-407 (1986). CESAR-86/28.

I.11. J. R. Einstein and J. Barhen, "Virtual-Time Operating-System Functions for Robotics Applications on a Hypercube," Hypercube Multiprocessors, '87, M. T. Heath ed., pp. 100-107, SIAM, Philadelphia (1987). CESAR-86/34.

I.12. M. G. Forest-Barlach and S. M. Babcock, "Inverse Dynamics Position Control of a Compliant Manipulator," IEEE Journal of Robotics and Automation, Vol. RA-3, No. 1, pp. 75-83 (February 1987). CESAR-86/08.

I.13. A. Guez, V. Protopopescu and J. Barhen, "On the Stability, Storage Capacity, and Design of Nonlinear Continuous Neural Networks," IEEE Trans. on Systems Man & Cybernetics, Vol. 18, No. 1, pp. 80-87 (January/February 1988). CESAR-86/49.

I. BOOKS AND REFEREED JOURNALS (cont'd)

I.14. R. C. Mann, "A Comparison of Algorithms for the Analysis of Images Generated by Two-Dimensional Electrophoresis," submitted to Electrophoresis (1987). CESAR-87/01.

I.15. N. S. V. Rao, S. S. Iyengar, and C. C. Jorgensen, "Terrain Model Acquisition by a Finite-Sized Mobile Robot in Plane," 1987 IEEE Robotics and Automation Conference, Raleigh, N.C., Proceedings, Vol. 1, pp. 283-288 (1987). CESAR-86/60.

I.16. C. R. Weisbin, "Intelligent-Machine Research at CESAR," AI Magazine, Spring 1987, Vol. 8, No. 1 (1987). CESAR-87/16.

I.17. C. C. Jorgensen and C. Matheus, "Catching Knowledge in Neural Nets," AI Expert, Vol. 1, No. 4, pp. 32-38 (December 1986). CESAR-86/31.

I.18. J. Barhen, N. Toomarian, V. Protopopescu, "Optimization of the Computational Load of a Hypercube Supercomputer Onboard a Mobile Robot," Applied Optics Special Issue on Neural Networks, Vol. 26, 23, pp. 5007-5014 (July 1987). CESAR-87/26.

I.19. B. L. Burks, G. deSaussure, C. R. Weisbin, J. P. Jones, and W. R. Hamel, "Autonomous Navigation, Exploration and Recognition," Winter 1987 Issue of IEEE Expert, pp. 18-27. CESAR-87/25.

I.20. A. Guez, V. Protopopescu, and J. Barhen, "On the Stability, Storage Capacity and Design of Nonlinear Continuous Neural Networks," IEEE Journal (1987). CESAR-87/35.

I.21. J. Barhen, W. B. Dress and C. C. Jorgensen, "Applications of Concurrent Neuromorphic Algorithms for Autonomous Robots," Publisher: Springer-Verlag, Book Chapter, Chapter 6, pp. 321 (January 1988). CESAR-87/45.

I.22. M. Beckerman, L. A. Farkas, and S. E. Johnston, "Treatment of Systematic Errors II: Fusion of Ultrasound and Visual Sensor Data," submitted to IEEE Transactions on Robotics and Automation.

I.23. N. Toomarian, E. Wacholder and S. Kaizerman, "Sensitivity Analysis of Two-Phase Flow Problems," Nuclear Science and Engineering, 99, pp. 53-81 (1988). CESAR-87/56.

I.24. M. Beckerman and E. M. Oblow, "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor Data for Robotic Navigation," Vol. 6, No. 2, pp. 137-145, IEEE Journal of Robotics and Automation (April-1990) CESAR-88/07.

I.25. E. M. Oblow, "Supertracks, Supertrack Functions and Chaos in the Quadratic Map," Physics Letters A, Vol. 128, No. 8, pp. 406-412 (March 1988). CESAR-88/10.

I.26. G. Bilbro, R. C. Mann, T. K. Miller, W. E. Snyder, D. E. VandenBout, M. White, "Simulated Annealing Using the Mean Field Approximation," Science (1988). CESAR-88/11.

I.27. E. Wacholder, J. Han, R. C. Mann, "A Neural Network Algorithm for the Multiple Traveling Salesman Problem," Biological Cybernetics, Vol. 61, pp. 11-19 (1989). CESAR-88/21.

I. BOOKS AND REFEREED JOURNALS (cont'd)

I.28. E. M. Oblow, "Foundations of O-Theory: Measurements and Relation to Fuzzy Set Theory," Int. Journal of General Systems, Vol. 14, pp. 357-378 (May 1988). CESAR-88/32.

I.29. R. V. Dubey, J. A. Euler and S. M. Babcock, "Real-Time Implementation of a Kinematic Gradient Projection Optimization Scheme for Seven-Degree-of-Freedom Redundant Robots with Spherical Wrists," submitted to Journal of Robotics and Automation. (In review at Clemson University with J. Luh-June 1989). CESAR-88/36.

I.30. C. R. Weisbin, W. R. Hamel, D. P. Kuban, S. A. Meacham and F. G. Pin, "The Robotics and Intelligent Systems Program at ORNL," Robotica, Vol. 7, No. 2, pp. 101-111 (April 1989). CESAR-88/40.

I.31. Carlos March-Leuba and R. B. Perez, "Adaptive Optimal Control of Uncertain Nonlinear Systems: On-Line Microprocessor-Based Algorithm to Control Mechanical Manipulators," submitted to the Journal of Robotic Systems (August 1988). CESAR-88/19. See VI.21 - ORNL/TM-10764.

I.32. P. F. Spelt, E. Lyness, G. deSaussure, "Development and Training of a Learning Expert System for an Autonomous Robot," Simulation, Vol. 53(5), pp. 223-228 (Nov. 1989). CESAR-88/46.

I.33. C. R. Weisbin, G. de Saussure, J. R. Einstein, E. Heer and F. G. Pin, "Autonomous Mobile Robot Navigation and Learning," Computer, Vol. 22, No. 6, pp. 29-35 (June 1989). CESAR-88/59.

I.34. A. Sabharwal, S. S. Iyengar, C. R. Weisbin and F. G. Pin, "Asynchronous Production Systems," Knowledge-Based Systems, Vol. 2, No. 2, pp. 117-127 (1989). CESAR-88/67.

I.35. J. C. Culioli and V. Protopopescu, "An Algorithm for Linear Programming That is Easy to Implement," Applied Mathematics Letters (2)2, pp. 125 (1989). CESAR-89/08.

I.36. A. Sabharwal, S. S. Iyengar, F. G. Pin and C. R. Weisbin, "Asynchronous Production Systems for Control of an Autonomous Mobile Robot," Accepted for publication in the Journal of Applied Artificial Intelligence (Dec. 1990). CESAR-89/10.

I.37. A. Guez, V. Protopopescu and J. Barhen, "On the Stability, Storage Capacity and Design of Nonlinear Continuous Neural Networks," Submitted to IEEE Transactions on Systems, Man and Cybernetics, Vol. 18, No. 1 (January and February 1988) CESAR-86/49. See VI.10, ORNL/TM-10329.

I.38. J.-C. Culioli, "On Solving 0-1 Optimization Problems with Neural Networks," Submitted to Neural Networks (March 1989). CESAR-89/18.

I.39. J-C. Culioli and G. Cohen, "Decomposition/Coordination Algorithms in Stochastic Optimization," Vol. 28, Issue 6, SIAM Journal on Control and Optimization (11/90). CESAR-89/20.

I. BOOKS AND REFEREED JOURNALS (cont'd)

I.40. E. Wacholder, "Neural Network Algorithm for the Weapon-Target Assignment Problem," Accepted for publication in the ORSA Journal on Computing. See also VI.27. CESAR-88/66.

I.41. C. R. Weisbin, B. L. Burks, J. R. Einstein, R. R. Feezell, W. W. Manges and D. H. Thompson, "HERMIES-III: A Step Toward Autonomous Mobility, Manipulation and Perception," Robotica Vol. 8 pp. 7-12 (1990). CESAR-88/70.

I.42. D. K. Wehe, J. C. Lee, W. R. Martin, R. C. Mann, W. R. Hamel and J. Tulenko, "Intelligent Robotics and Remote Systems for the Nuclear Industry," Nuclear Engineering and Design, Vol. 113, pp. 259-267 (1989). CESAR-89/29.

I.43. P. F. Spelt, G. deSaussure, E. Lyness, F. G. Pin and C. R. Weisbin, "Learning by an Autonomous Robot at a Process Control Panel," IEEE Expert, Vol. 4(4), pp. 8-16, Winter Edition (1989). CESAR-87/55.

I.44. H. Hiriyannaiah, G. L. Bilbro, W. E. Snyder and R. C. Mann, "Restoration of Piecewise Constant Images Via Mean Field Annealing," Journal of the Optical Society of America A Vol. 6, No. 12 (December 1989). CESAR-90/01.

I.45. G. deSaussure, C. R. Weisbin and P. F. Spelt, "Navigation and Learning Experiments by an Autonomous Robot," Robotics and Manufacturing Journal Vol. 6, No. 4, pp. 295-301 (1990) CESAR-90/04.

I.46. A. J. Koivo and M. A. Unseren, "Modeling Closed Chain Motion of Two Manipulators Holding a Rigid Object," Mechanism and Machine Theory Vol. 25, No. 4, pp. 427-438. (1990). CESAR-90/10.

I.47. D. B. Reister, "Time Optimal Trajectories for a Two Wheeled Robot," Submitted to Robotics Research Journal (1990). CESAR-90/13.

I.48. P. F. Spelt, G. deSaussure, G. Oliver and M. Silliman, "Concept Formation and Generalization Based on Experimentation by an Autonomous Robot," IEEE Expert (May 1990). CESAR-90/20.

I.49. K. Fujimura, "Planning a Time-Minimal Motion Among Moving Obstacles," submitted to Algorithmica CESAR 90/40.

I.50. Hubert A. Vasseur, Francios G. Pin, and Jack R. Taylor, "Navigation of Car-Like Mobile Obstructed Environments using Convex Polygonal Cells," submitted to Computers and Electrical Engineering, Vol. 16(4), (December 1990), CESAR 90/42.

I.51. M. A. Unseren, "Rigid Body Dynamics and Decoupled Control Architecture For Two Strongly Interacting Manipulators," submitted to Robotica, CESAR 90/28.

I. BOOKS AND REFEREED JOURNALS (cont'd)

I.52. Francois G. Pin and Jean-Christophe Culoili, "Optimal Positioning of Combined Mobile Platform-Manipulator Systems for Material Handling Tasks," submitted to Robotic Systems, CESAR 90/53.

I.53. Claus S. Andersen, Claus B. Madsen, Jan J. Sorensen, Neils O. S. Kirkeby, Judson P. Jones and Henrik I. Christensen, "Navigation Using Range Images on a Mobile Robot," submitted to Computers and Electrical Engineering, CESAR 90/53.

I.54. M. A. Unseren, "Dynamic Coupling Effects in Modeling and Control of Hard Contact Motion of a Manipulator," submitted to IEEE International Conference on Systems Engineering, August 1-3, 1991, Dayton, Ohio, CESAR-90/56.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED

II.1. J. Barhen, S. M. Babcock, W. R. Hamel, E. M. Oblow, G. N. Saridis, G. de Saussure, A. D. Solomon and C. R. Weisbin, "Basic Research on Intelligent Robotic Systems Operating in Hostile Environments: New Developments at ORNL," The 1984 National Topical Meeting on Robotics and Remote Handling in Hostile Environments, Proceedings, pp. 105-116 (April 1984).

II.2. S. M. Babcock and J. Barhen, "Real-Time Algorithms for Robotics Control of Teleoperators," Robots-8 Conference, Detroit, MI, Proceedings, pp. 1972-1987 (June 4-7, 1984).

II.3. C. R. Weisbin, G. de Saussure, J. Barhen, E. M. Oblow and J. C. White, "Minimal Cut-Set Methodology for Artificial Intelligence Applications," The First Conference on Artificial Applications, sponsored by the IEEE Computer Society in Cooperation with American Association for Artificial Intelligence, Proceedings, pp. 465-469 (December 5-7, 1984).

II.4. C. R. Weisbin, J. Barhen, T. E. Swift, G. de Saussure, C. C. Jorgensen and E. M. Oblow, "HERMIES-I: A Mobile Robot for Navigation and Manipulation Experiments," The Robots-9 Conference, Detroit, MI, Proceedings, Vol. 1, pp. 1-41 (June 3-6, 1985).

II.5. M. G. Forrest and S. M. Babcock, "Control of a Single Link, Two-Degree-of-Freedom Manipulator with Joint Compliance and Actuator Dynamics," The 1985 International Computers in Engineering Conference, Boston, MA, Proceedings, pp. 189-197 (August 4-8, 1985).

II.6. J. Barhen, "Robot Inverse Dynamics on a Concurrent Computation Ensemble," 1985 International Computers in Engineering Conference, Boston, MA, Proceedings, pp. 415-429 (August 4-8, 1985).

II.7. S. S. Iyengar, C. C. Jorgensen, S. V. N. Rao and C. R. Weisbin, "Learned Navigation Paths for a Robot in Unexplored Terrain," The Second Conference on Artificial Intelligence Applications, Miami Beach, FL, Proceedings, pp. 148-155 (December 11-13, 1985).

II.8. M. G. Forrest-Barlach and S. M. Babcock, "Inverse Dynamics Position Control of a Compliant Manipulator," The 1986 IEEE International Conference on Robotics and Automation, San Francisco, CA, Proceedings, Vol. 1, pp. 196-205 (April 8-10, 1986). CESAR-86/08.

II.9. S. V. N. Rao, S. S. Iyengar, C. C. Jorgensen and C. R. Weisbin, "Concurrent Algorithms for Autonomous Robot Navigation in an Unexplored Terrain," The 1986 IEEE International Conference on Robotics and Automation, San Francisco, CA, Proceedings, Vol. 1, pp. 1137-1144 (April 8-10, 1986).

II.10. C. R. Weisbin, "Research in Intelligent Machines at ORNL." Appeared in the ORNL Review Special Issue - Basic Physical Sciences at ORNL-1986, Number Three, Volume 19, pp. 130-143, Proceedings of The Second World Conference on Robotics Research, MS86-772, Scottsdale, Arizona (August 19-21, 1986). CESAR-86/12.

II.11. J. Barhen, E. C. Halbert and J. R. Einstein, "Advances in Concurrent Computers for Autonomous Robots," The Second World Conference on Robotics Research, Proceedings, MS86-771, Scottsdale, Arizona (August 19-21, 1986). CESAR-86/14.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.12. C. R. Weisbin, G. de Saussure and D. W. Kammer, "A Real-Time Expert System for Monitoring the Performance of an Intelligent Robot and Diagnosing Unexpected Occurrences," SPIE - The International Society for Optical Engineering, Proceedings, 729, pp. 49-58 (1986). CESAR-86/44.

II.13. C. C. Jorgensen, "Neural Network Modeling of Early Retina Vision Processes," presented to Neurcom, San Jose, CA. CESAR-86/48.

II.14. M. Goldstein, F. G. Pin, G. de Saussure and C. R. Weisbin, "3-D World Modeling Based on Combinatorial Geometry for Autonomous Robot Navigation," 1987 IEEE International Conference on Robotics and Automation, Proceedings, Vol. 2, pp. 727-733. CESAR-86/51.

II.15. S. M. Babcock, W. R. Hamel and S. M. Killough, "Advanced Manipulation for Autonomous Mobile Robots," International Topical Meeting on Remote Systems and Robotics, Pasco, WA, Proceedings, pp. 290-297 (March 29-April 2, 1987). CESAR-86/52.

II.16. N. S. V. Rao, S. S. Iyengar and C. R. Weisbin, "On Autonomous Terrain Model Acquisition by a Mobile Robot," JPL (NASA) Conference on Telerobotics, Proceedings 1, pp. 283-288 (1987). CESAR-86/59.

II.17. N. S. V. Rao, S. S. Iyengar, C. C. Jorgensen and C. R. Weisbin, "On Terrain Acquisition by a Finite-Sized Mobile Robot in Plane," The Workshop on Space Telerobotics, Proceedings, Vol. 3, pp. 1314-1319, Pasadena, CA (1987). CESAR-86/60.

II.18. M. Goldstein, F. G. Pin, G. de Saussure and C. R. Weisbin, "3-D World Modeling with Updating Capability Based On Combinatorial Geometry," JPL (NASA) Conference on Telerobotics. The Workshop on Space Telerobotics, Proceedings, Vol. 1, pp. 273-282 (1987). CESAR-87/02.

II.19. J. R. Einstein, J. Barhen and D. Jefferson (UCLA), "A Virtual-Time Operating System Shell for a Hypercube in Robotics Applications," SPIE Cambridge on Optical and Optoelectronic Engineering, Cambridge, Massachusetts, October 21-31, 1986. Proceedings, Vol. 726, Society of Photo-Optical Instrumentation Engineers. CESAR-86/61.

II.20. L. E. Parker and F. G. Pin, "A Methodology for Dynamic Task Allocation in a Man-Machine System," Second International Symposium on Methodologies for Intelligent Systems, Charlotte, North Carolina, October 14-17, 1987. Methodologies for Intelligent Systems, pp. 488-495 (1987). CESAR-87/14.

II.21. L. E. Parker and F. G. Pin, "Dynamic Task Allocation for Man-Machine Symbiotic System," 1987 Goddard Conference on Space Allocation of Artificial Intelligence and Robotics, May 13-14, 1987, NASA/Goddard Space Flight Center, Greenbelt, Maryland. Proceedings, Paper 14, Robotics Section (1987). CESAR-87/08.

II.22. C. C. Jorgensen, "Neural Network Representation of Sensor Graphs for Autonomous Robot Navigation." IEEE International Conference on Neural Networks, San Diego, California, June 21-24, 1987. Proceedings, pp. IV507-IV515 (1987). CESAR-87/04.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.23. J. Barhen, V. Protopopescu, N. Toomarian, and M. Clinard, "Concurrent Neuromorphic Algorithms for Optimization of the Computational Load of a Hypercube Supercomputer." IEEE International Conference on Neural Networks, San Diego, California, June 21-24, 1987, Proceedings, pp. IV687-IV696 (1987). CESAR-87/06.

II.24. M. Goldstein, F. G. Pin, and C. R. Weisbin, "Range Image Segmentation and Surface Parameter Extraction for a Mobile Robot," 1988 IEEE International Conference on Robotics and Automation, Philadelphia, Pennsylvania, April 25-29, 1988. CESAR-87/34.

II.25. M. Goldstein, F. G. Pin, and C. R. Weisbin, "3-D World Modeling for An Autonomous Robot," 1987 SPIE Cambridge Symposium, Intelligent Robots and Computer Vision Conference, Cambridge, Massachusetts, November 1-6, 1987. CESAR-87/36.

II.26. N. S. V. Rao, S. S. Iyengar, and G. de Saussure, "The Visit Problem: Visibility Graph-Based Solution," Proceedings 3, pp. 1650-1655, 1988 IEEE International Conference on Robotics and Automation, Philadelphia, Pennsylvania, April 25-29, 1988. CESAR-87/38.

II.27. B. L. Burks, D. L. Barnett, S. M. Killough, R. C. Mann, and F. G. Pin, "Dealing with Unanticipated Events in Autonomous Navigation," 26th IEEE Conference on Decision and Control, Los Angeles, California, Proceedings, Vol. 3, pp. 1808-1815 (December 1987). CESAR-87/40.

II.28. R. V. Dubey, J. A. Euler, and S. M. Babcock, "Optimization of Performance Criteria with Bounded Joint Velocities for Robots with One Degree of Redundancy," The USA-JAPAN Symposium on Flexible Automation, July 18-20, 1988, Minneapolis, Minnesota, Proceedings, pp. 115-122 (1988). CESAR-87/41.

II.29. R. V. Dubey, J. A. Euler, and S. M. Babcock, "An Efficient Gradient Projection Optimization Scheme for a Seven-Degree-of-Freedom Redundant Robot with Spherical Wrist," 1988 IEEE International Conference on Robots and Automation, April 25-29, 1988, Philadelphia, Pennsylvania, Proceedings, Vol. 1, pp. 28-36. CESAR-87/42.

II.30. J. P. Jones, R. C. Mann, "Concurrent Algorithms for a Mobile Robot Vision System," Applications of Artificial Intelligence VI, Mohan Trivedi, Editor, Proceedings, SPIE 937, pp. 497-503 (1988). CESAR-88/06.

II.31. J. P. Jones, "A Concurrent On-Board Vision System for a Mobile Robot," Third Conference on Hypercube Concurrent Computers and Applications, Proceedings, Vol. II, pp. 1022-1032, Pasadena, CA (January 19-20, 1988). CESAR-88/13.

II.32. C. W. Glover, "Multi-Sensor Integration on the NCUBE Hypercube Computer," Third Conference on Hypercube Concurrent Computers and Applications, Proceedings, V.II, pp. 1236-1246, Pasadena, CA (January 19-20, 1988). CESAR-88/14.

II.33. N. Toomarian, "A Concurrent Neural Network Algorithm for the Traveling Salesman Problem," The Third Conference on Hypercube Concurrent Computers and Applications, Pasadena, CA (January 19-20, 1988), Proceedings, pp. 1483-1490 (1988). CESAR-88/13.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.34. D. M. Jollay and R. E. Ricks, "Sensor Fusion for Robot Navigation," to be presented at the 26th Annual Conference of the Southeast Region, Association for Computing Machinery, Mobile, Alabama (April 21-22, 1988). CESAR-88/15.

II.35. M. Goldstein, N. Toomarian and J. Barhen, "A Comparison Study of Optimization Methods for the Bipartite Matching Problem (BMP)," The IEEE Annual International Conference on Neural Networks, San Diego, CA (July 24-27, 1988), Proceedings, pp. 267-273 (1988). CESAR-88/18.

II.36. J-C. Culioli and V. Protopopescu, "A Linear Programming Method Inspired by the Neural Network Framework," IEEE International Conference on Neural Networks, San Diego, CA, July 24-27, 1988. CESAR-88/22.

II.37. M. W. Noakes and D. B. Reister, "Design of an Obstacle Avoidance System for AIMS," Proceedings of SPIE, Vol. 1006, pp. 124-131, Space Station Automation IV, Cambridge, Mass, November 6-11, 1988 (1988). CESAR-88/26.

II.38. R. C. Mann, W. R. Hamel and C. R. Weisbin, "The Development of an Intelligent Nuclear Maintenance Robot," 1988 International Conference on Robotics and Automation, Philadelphia, PA, April 25-29, 1988. Proceedings, Vol. 1, pp. 621-623 (1988). CESAR-88/05.

II.39. E. Wacholder, J. Han and R. C. Mann, "An Extension of the Hopfield-Tank Model for Solution of the Multiple Traveling Salesman Problem," The IEEE Annual International Conference on Neural Networks, San Diego, CA, July 24-27, 1988, Proceedings, pp. 305-323 (1988). CESAR-88/24.

II.40. S. M. Babcock, R. V. Dubey, J. A. Euler, R. L. Glassell, R. L. Kress, W. R. Hamel, "Telerobotic Control of the Seven-Degree-of-Freedom CESAR Manipulator," NATOl Advanced Research Workshop, "Robots with Redundancy: Design, Sensing and Control," June 27-July 1, 1988, Salo, Lago di Garda, Italy. CESAR-88/35.

II.41. J. A. Euler, R. V. Dubey and S. M. Babcock, "Self Motion Determination Based on Joint Velocity Bounds for Redundant Robots," The Second International Symposium, Robotics and Manufacturing: Recent Trends in Research, Education and Applications, November 16-18, 1988, Albuquerque, New Mexico, Proceedings, pp. 291-298 (1988). CESAR-88/43.

II.42. G. deSaussure, C. R. Weisbin and P. F. Spelt, "Navigation and Learning Experiments by an Autonomous Robot," The Second International Symposium on Robotics and Manufacturing: Recent Trends in Research, Education and Applications, November 16-18, 1988, Albuquerque, NM, Proceedings, pp. 935-948 (1988). CESAR-88/44.

II.43. R. C. Mann, J. P. Jones, M. Beckerman, C. W. Glover, L. Farkas, J. Han, E. Wacholder, J. R. Einstein, "Advanced Computing and Integrated Sensor Systems," The 1988 IEEE International symposium on Intelligent Control, August 24-26, 1988, Arlington, Virginia. CESAR-88/48.

II.44. D. B. Reister, "Design of an Obstacle Avoidance System for Aims," 1988 SPIE Cambridge Symposium, Advances in Intelligent Robotics Systems, Cambridge, Massachusetts, November 6-11, 1988, Proceedings, pp. 124-131 (1988). CESAR-88/17.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.45. L. E. Parker and F. G. Pin, "Man-Robot Symbiosis: A Framework for Cooperative Intelligence and Control," The 1988 SPIE Cambridge Symposium, Advances in Intelligent Robot Systems, Cambridge, Massachusetts, November 1-6, 1988, Proceedings, pp. 94-103 (1988). CESAR-88/27.

II.46. F. G. Pin, G. deSaussure, P. F. Spelt, S. M. Killough and C. R. Weisbin, "Recent CESAR Research Activities in Sensor Based Reasoning for Autonomous Machines," The IEEE International Workshop on Intelligent Robots and Systems, October 31-November 2, 1988, Tokyo, Japan, Proceedings, pp. 417-421 (1988). CESAR-88/30.

II.47. R. C. Mann, J. P. Jones, M. Beckerman, C. W. Glover, L. Farkas, J. Han, E. Wacholder and J. R. Einstein, "An Intelligent Integrated Sensor System for the ORNL Robot," The 1988 IEEE International Symposium on Intelligent Control, August 24-26, 1988, Arlington, VA. CESAR-88/48.

II.48. R. C. Mann, J. P. Jones, M. Beckerman, C. W. Glover, L. Farkas, G. L. Bilbro and W. Snyder, "Multi-Sensor Integration for Autonomous Robots in Nuclear Power Plants," IEEE 1989 International Conference on Robotics and Automation, Scottsdale, AZ, May 14-19, 1989. CESAR-88/57.

II.49. Carlos March-Leuba and R. B. Perez, "Optimal Control of Uncertain Nonlinear Systems: An Application to a Two-Link Robotic Arm," The 27th IEEE Conference on Decision and Control, December 7-9, 1988, Austin, TX. CESAR-88/58.

II.50. G. de Saussure, P. F. Spelt, S. M. Killough, F. G. Pin, and C. R. Weisbin, "Automatic Learning by an Autonomous Mobile Robot," The ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, March 13-16, 1989, Proceedings, Section 6-1, pp. 1-5 (1989). CESAR-88/62.

II.51. C. R. Weisbin, B. L. Burks, J. R. Einstein, R. R. Feezell, W. W. Manges, and D. H. Thompson, "Research in Autonomous Robotics at ORNL Using HERMIES-III," Quality Use of the Computer: Computational Mechanics, Artificial Intelligence, Robotics, and Acoustic Sensing, Vol. 177, pp. 165-169, The 1989 ASME Pressure Vessels and Piping Conference, July 23-27, 1989, Honolulu, Hawaii (1989). CESAR-89/09.

II.52. H. I. Christensen, J. P. Jones, "Concurrent Multi-Resolution Image Analysis," The 4th Conference on Hypercubes, Concurrent Compliance and Applications, Monterey, CA, March 6-8, 1989. CESAR-89/12.

II.53. H. I. Christensen, "Concurrent Spatio-Temporal Image Analysis," The 4th Conference on Hypercubes, Concurrent Compliance and Applications, Monterey, CA, March 6-8, 1989. CESAR-89/13.

II.54. L. E. Parker and C. R. Weisbin, "Review of the 1988 Workshop on Human-Machine Symbiotic Systems," Workshop on Integration of AI and Robotic Systems, Scottsdale, AZ, May 19, 1989. CESAR-89/14.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.55. R. C. Mann, J. P. Jones, M. Beckerman, C. W. Glover, L. Farkas, J. R. Einstein, "Sensor Integration Using Concurrent Computing On-Board the ORNL Mobile Robot," Workshop on Parallel Algorithms and Architectures for Robotics, IEEE Conference on Robotics and Automation, Scottsdale, AZ, May 1989. CESAR-89/15.

II.56. M. Beckerman, D. L. Barnett, J. R. Einstein, J. P. Jones, P. D. Spelt and C. R. Weisbin, "Robust Performance of Multiple Tasks by an Autonomous Robot," Workshop on Integration of AI and Robotic Systems, Scottsdale, AZ, May 19, 1989. CESAR-89/17.

II.57. J-C. Culioli, C. W. Glover, J. P. Jones and C. Roe, "NCUBE Implementation of Some Heuristics and an Optimal Algorithm for Large-Scale Assignment Problems," The Fourth Conference on Hypercubes, Concurrent Computers and Applications, March 1989, Monterey, CA. CESAR-89/21.

II.58. J. R. Einstein, C. W. Glover and M. Dickens, "A Tool to Aid in Mapping Computational Tasks to a Hypercube," The Fourth Conference on Hypercubes, Concurrent Computers and Applications, March 1989, Monterey, CA. CESAR-89/22.

II.59. R. C. Mann, C. R. Weisbin and F. G. Pin, "Research in Mobile Robotics at ORNL/CESAR," The IEEE International Conference on Systems Engineering, Proceedings, pp. 367-371, August 24-26, 1989. CESAR-89/25.

II.60. J. C. Culioli, V. Protopopescu and R. C. Mann, "Chaotic Behavior in a New Class of Parallel Optimization Algorithms," The Seventh Symposium on Energy Engineering Sciences, June 19-21, 1989, Argonne, IL.

II.61. F. G. Pin, M. Beckerman, P. F. Spelt, J. T. Robinson and C. R. Weisbin, "Autonomous Mobile Robot Research Using the HERMIES-III Robot," The 1989 RSJ International Workshop on Intelligent Robot and Systems '89, Proceedings, pp. 251-256, "The Autonomous Mobile Robot and Its Application," September 1989, Tsukuba, Japan.

II.62. D. L. Barnett, S. S. Iyengar, F. G. Pin, and A. Sabharwal, "Robot Navigation Research at CESAR," The 28th IEEE Conference on Decision and Control, Proceedings, Vol. 2, pp. 1162-1166, Tampa, Florida, December 13-15, 1989. CESAR-89/16.

II.64. J. C. Culioli, V. Protopopescu, C. Britton and N. Ericson, "Neural Network Models for Linear Programming," International Joint Conference on Neural Networks, January 15-19, 1990, Washington, DC. CESAR-89/32.

II.65. J. C. Culioli, V. Protopopescu, C. Britton and N. Ericson, "A Neural Network for Bounded Linear Programming," International Joint Conference on Neural Networks, January 15-19, 1990, Washington, DC. CESAR-89/33.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.66. N.S.V. Rao and C. W. Glover, "Symbolic Pattern Processing System for Planar Spatial Structures," 10th International Conference on Pattern Recognition, Atlantic City, New Jersey, June 17-21, 1990, Proceedings, pp. 351-365 (1990). CESAR-89/39.

II.67. P. F. Spelt, G. deSaussure, E. Lyness and M. Silliman, "Concept Formation and Generalization Based on Experimentation by an Autonomous Mobile Robot," IEEE Robotics and Automation Conference, May 13-18, 1990, Cincinnati, OH. CESAR-89/41. (rejected).

II.68. M. Beckerman, D. L. Barnett and M. Dickens, "Robust Performance of Multiple Tasks by a Mobile Robot," IEEE Conference on Robotics and Automation, Cincinnati, OH, May 13-18, 1990. CESAR-89/42.

II.69. P. F. Spelt, G. deSaussure, G. Oliver and M. Silliman, "Experimentation and Concept Formation by an Autonomous Mobile Robot," Japan-USA Symposium on Flexible Automation, July 9-11, 1990, Kyoto, Japan. CESAR-89/45. (rejected).

II.70. M. Beckerman, D. L. Barnett and S. M. Killough, "Ultrasound and Visual Sensor Feedback and Fusion," Japan-USA Symposium on Flexible Automation, Kyoto, Japan, July 9-11, 1990. Proceedings, Vol. 3, pp. 1315-1319 (1990). CESAR-89/46.

II.71. C. W. Glover, "Design and Implementation of a Multi-Sensor Fusion Algorithm on a Hypercube Computer Architecture," SPIE's Advances in Intelligent Robotics, Philadelphia, PA, November 5-10, 1989. CESAR-89/47.

II.72. J. P. Jones, M. Beckerman and R. C. Mann, "Design and Implementation of Two Concurrent Multi-Sensor Integration Algorithms for Mobile Robots," SPIE Symposium on "Advances in Intelligent Robotic Systems," Philadelphia, PA, November 5-8, 1989. Proceedings, Vol. 1198, pp. 301-312 (1989). CESAR-89/49.

II.73. P. Belmans, "An Approach to Elemental Task Learning," SPIE's Applications of Artificial Intelligence VIII, Orlando, FL, April 17-19, 1990, Proceedings, Vol. 1293, pp. 634-644 (1990). CESAR-90/02.

II.74. P. Belmans and J-C. Culioli, "An Alternative Method to Solving the Kinematics of a Redundant Robot," American Nuclear Society, 1990 Annual Meeting, Nashville, TN, June 10-14, 1990, Proceedings, Vol. 61, pp. 426-427 (1990) CESAR-90/03.

II.75. C. W. Glover, M. Silliman, M. Walker, P. F. Spelt and N.S.V. Rao, "Hybrid Neural Network and Rule-Based Pattern Recognition System Capable of Self-Modification," SPIE's Applications of Artificial Intelligence VIII Conference, Orlando, FL, April 17-19, 1990. Proceedings, Vol. 1293, pp. 290-300 (1990). CESAR-90/07.

II.76. C. W. Glover and P. F. Spelt, "Hybrid Intelligent Perception System: Intelligent Perception Through Combining Artificial Neural Networks and an Expert System," First Workshop on Neural Networks: Academic/Industrial/NASA/Defense, Auburn University, Auburn, AL, February 5-6, 1990. CESAR-90/08.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.77. S. M. Killough and F. G. Pin, "A Fully Omnidirectional Wheeled Assembly for Robotic Vehicles," American Nuclear Society, 1990 Annual Meeting, Vol. 61, pp. 425-426, Nashville, TN, June 10-14, 1990. Proceedings, Vol. 61, pp. 425-426 (1990). CESAR-90/05.

II.78. H. A. Vasseur and F. G. Pin "Trajectory Generation for Car-Like Robots," American Nuclear Society, 1990 Annual Meeting, Nashville, TN, June 10-14, 1990. Proceedings, Vol. 61, pp. 421-422 (1990). CESAR-90/06.

II.79. K. Fujimura and H. Samet, "Motion Planning in a Dynamic Domain," 1990 IEEE International Conference on Robotics and Automation, Cincinatti, OH, May 13-18, 1990. Proceedings, Vol. 1, pp. 330-342 (1990). CESAR-90/09.

II.80. Hugo Embrechts and J. P. Jones, "An Input/Output Algorithm for M-Dimensional Rectangular Domain Decompositions on N-Dimensional Hypercube Multicomputers," 5th Conference on Distributed Memory Concurrent Computers, April 8-12, 1990, Charleston, SC. Proceedings, Vol. II, pp. 876-882 (1990). CESAR-90/18.

II.81. F. G. Pin and J.-C. Culioli, "Multi-Criteria Position and Configuration Optimization for Redundant Platform/Manipulator Systems," Proceedings of the IROS '90, the 1990 IEEE International Workshop on Intelligent Robots and Systems, Tsuchiura, Ibaraki, Japan, July 3-6, 1990. Proceedings, pp. 103-107 (1990). CESAR-89/51.

II.82. N.S.V. Rao and C. W. Glover, "Computational Aspects of Similarity Measures Based on Linearization of Lattice Points," IAPR Workshop on Syntactic and Structural Pattern Recognition, Murray Hill, New Jersey, June 13-15, 1990. CESAR-90/19.

II.83. M. Beckerman, D. L. Barnett, M. Dickens and C. R. Weisbin, "Performance of Multiple Tasks by an Autonomous Robot Using Visual and Ultrasound Sensing," 3rd ISRAM International Symposium on Robotics and Manufacturing, British Columbia, Canada, July 18-20, 1990. Proceedings, Vol. 3, pp. 389-395 (1990). CESAR-90/22.

II.84. F. G. Pin and J. C. Culioli, "Optimal Positioning of Redundant Manipulator-Platform Systems for Maximum Task Efficiency," To be published in the Proceedings of the ISRAM '90 Third International Symposium on Robotics and Manufacturing, Vancouver, B.C., Canada, July 18-20, 1990, Proceedings, pp. 489-495 (1990). CESAR-89/48.

II.94. J. P. Jones, "On the Design of a Concurrent Image Analysis System," Proceedings of the ISRAM '90 Third International Symposium on Robotics and Manufacturing, Vancouver, B.C., Canada, July 18-20, 1990. Proceedings, Vol. 3, pp. 205-210 (1990). CESAR-89/50.

II.95. F. G. Pin and H. A. Vasseur, "Autonomous Trajectory Generation for Mobile Robots with Non-Holonomic and Steering Angle Constraints," To be published in the Proceedings of the IEEE International Workshop on Intelligent Motion Control, Istanbul, Turkey, August 20-22, 1990. Proceedings, pp. 295-299 (1990). CESAR-90/26.

II.96. R. C. Mann, "A Hardware/Software Environment to Support R&D in Intelligent Machines and Mobile Robotic Systems," To be presented at the DARPA Workshop on Dstributed Intelligent Control Systems, July 17-19, 1990, Pacifica, CA. CESAR-90/30.

II. PROCEEDINGS OF CONFERENCES - FULL PAPER REVIEWED (cont'd)

II.97. Darren Dawson and M. A. Unseren, "On the Control of Two Manipulators Holding a Rigid Object In the Presence of Uncertainty," submitted to 1991 American Control Conf., Boston, MA, June 26-28, 1991. CESAR-90/33.

II.98. D. B. Reister, "A New Wheel Control System for the Omnidirectional HERMIES-III Robot," Submitted to 1991 IEEE International Conference on Robotics and Automation, Sacramento, CA, April 7-12, 1991. (To be published in the Proceedings.) CESAR 90/34.

II.99. J. P. Jones, D. B. Reister, P. L. Butler, M. Beckerman and F. J. Sweeney, "DEMO 89 - The Initial Experiment with the HERMIES-III Robot," Submitted to the 1991 IEEE International Conference on Robotics and Automation, Sacramento, CA, April 7-12, 1991. CESAR 90/35.

II.100. Stephen M. Killough and Francois G. Pin, "Design and Control of a Fully Omnidirectional and Holonomic Platform for Robotic Vehicles," Submitted to the 1991 IEEE International Conference on Robotics and Automation, Sacramento, CA, April 7-12, 1991. CESAR 90/36.

II.101. Hubert A. Vasseur, Francois G. Pin and Jack R. Taylor, "Navigation of a Car-Like Mobile Robot Using a Decomposition of the Environment in Convex Cells," Submitted to the 1991 IEEE International Conference on Robotics and Automation, Sacramento, CA, April 7-12, 1991. CESAR 90/37.

III. PROCEEDINGS OF CONFERENCES - SUMMARY REVIEWED

III.1. C. R. Weisbin, G. de Saussure and J. Barhen, "Automated Planning for Intelligent Machines in Energy-Related Applications," Conference on Intelligent Systems and Machines, Proceedings, pp. 13-15, Oakland University, Rochester, MI (April 24-25, 1984).

III.2. J. Barhen, C. R. Weisbin and G. de Saussure, "Real-Time Planning by an Intelligent Robot," ASME Conference on Computers in Engineering, Proceedings, pp. 358-360 (August 1984).

III.3. C. R. Weisbin, G. de Saussure, J. Barhen, T. E. Swift and J. C. White, "Strategy Planning by an Intelligent Machine," First World Conference on Robotics Research: The Next Five Years and Beyond, LeHigh University, PA, Proceedings, MS84-501 (August 14-16, 1984).

III.4. J. Barhen and S. M. Babcock, "Parallel Algorithms for Robot Dynamics," First World Conference on Robotics Research: The Next Five Years and Beyond, LeHigh University, PA, Proceedings, MS84-499 (August 14-19, 1984).

III.5. C. R. Weisbin, J. Barhen, G. de Saussure, W. R. Hamel, C. C. Jorgensen, J. L. Lucius, E. M. Oblow and T. E. Swift, "Machine Intelligence for Robotics Applications," 1985 The Conference on Intelligent Machines, Oakland University, Rochester, MI, Proceedings, pp. 47-57 (April 23, 1985).

III.6. E. M. Oblow, "A Hybrid Uncertainty Theory," The Fifth International Workshop: Expert Systems and Their Applications, Avignon, France, Proceedings, pp. 1193-1201 (May 13-15, 1985).

III.7. W. R. Hamel, S. M. Babcock, M. C. G. Hall, C. C. Jorgensen, S. M. Killough and C. R. Weisbin, "Autonomous Robots for Hazardous and Unstructured Environments," The Robots-10 Conference, Chicago, IL, Proceedings, pp. 5-9 through 5-29 (April 20-24, 1986).

III.8. D. W. Kammer, G. de Saussure and C. R. Weisbin, "A Real-Time Expert System for Robotic Task Planning, Monitoring and Diagnosis of Unexpected Occurrences," 1986 Conference on Intelligent Systems and Machines (April 29-30, 1986). CESAR-86/20.

III.9. E. M. Oblow, "O-Theory: A Hybrid Uncertainty Theory," Int. J. General Systems, Vol. 13, pp. 95-106 (1987).

III.10. J. R. Einstein, J. Barhen and D. Jefferson, "Intelligent Operating System for Autonomous Robots: Real-Time Capabilities on a Hypercube Supercomputer," The International Topical Meeting on Remote Systems and Robotics in Hostile Environments, American Nuclear Society, LaGrange, IL (March 1987). CESAR-86/57.

III.11. G. de Saussure, D. W. Kammer and C. R. Weisbin, "A Real-Time Expert System for the Control of Autonomous Robot Navigation in the Presence of Moving Obstacles," Pasco, WA (1987). CESAR-86/50.

III.12. A. S. Sabharwal, S. S. Iyengar, C. R. Weisbin and F. G. Pin, "Asynchronous Production Systems for Real-Time Expert Systems," Eighth International Workshop on Expert Systems and Their Applications, Avignon, France, Proceedings, pp. 113-146 (May 1988). CESAR-88/08.

III. PROCEEDINGS OF CONFERENCES - SUMMARY REVIEWED (cont'd)

III.13. R. V. Dubey, J. E. Euler, R. L. Glasswell and S. M. Babcock, "Real-Time Implementation of a Kinematic Optimization Scheme for Seven-Degree-of-Freedom Redundant Robots with Spherical Wrists," The American Control Conference, Atlanta, GA (June 15-17, 1988), Proceedings, pp. 1376-1378 (1988). CESAR-87/49.

III.14. G. deSaussure, P. F. Spelt, S. M. Killough, F. G. Pin, and C. R. Weisbin, "Automatic Learning by an Autonomous Mobile Robot," Third Topical Meeting on Robotics and Remote Systems, Charleston, South Carolina, March 13-16, 1989. CESAR-88/37.

III.15. D. L. Barnett, G. de Saussure and R. A. Lexvold, "Using an Expert System for Intelligent Control in Autonomous Robot Navigation," To appear in the Proceedings of the American Nuclear Society, November 11-16, 1990, Washington, D.C. CESAR-90/25.

III.16. P. F. Spelt, "A Hybrid Artificial Intelligence Architecture as an Intelligent Interface in Hierarchical Supervisory Control", submitted to Second Workshop on Neural Networks (WNN-AIND'91) Auburn University, Auburn, Alabama February 11-13, 1991. CESAR-90/57.

IV. INVITED TALKS

IV.1. C. R. Weisbin, J. Barhen, G. de Saussure, W. R. Hamel, C. C. Jorgensen, J. L. Lucius, E. M. Oblow and T. E. Swift, "Artificial Intelligence for Energy-Related Applications," 1985 WATTEC Conference, Knoxville, TN (February 12, 1985).

IV.2. C. R. Weisbin, J. Barhen, G. de Saussure, W. R. Hamel, C. C. Jorgensen, J. L. Lucius, E. M. Oblow and T. E. Swift, "Artificial Intelligence for Concurrent Computation for Robotic Applications," Transactions of American Nuclear Society Annual Meeting, Boston, MA, Vol. 49, pp. 310-311 (June 9-14, 1985).

IV.3. J. P. Jones and M. A. Scudieri, "Image Processing on a Hypercube Concurrent Multiprocessor," RIMSIG Symposium, Santa Fe, New Mexico (October 22-23, 1986). CESAR-86/39.

IV.4. J. R. Einstein and J. Barhen, "Robot Control with a Hypercube: Virtual-Time Functions for the Node's Operating System," RIMSIG Symposium, Santa Fe, New Mexico (October 22-23, 1986). CESAR-86/41.

IV.5. R. C. Mann, "Quantitative Analysis of Two-Dimensional Protein Patterns Using an NCUBE Hypercube," RIMSIG Symposium, Santa Fe, New Mexico (October 22-23, 1986). CESAR-86/42.

IV.6. R. C. Mann, "Multi-Sensor Integration for a Mobile Robot Using Concurrent Computing," The Federal Republic of Germany (May 11-13, 1987). CESAR-86/53.

IV.7. William R. Hamel and Charles R. Weisbin, "Human-Scale Experiments In Mobile Autonomous Robotics," The Federal Republic of Germany (May 11-13, 1987). CESAR-86/58.

IV.8. B. L. Burks, D. L. Barnett, J. P. Jones and S. M. Killough, "A Demonstration of Autonomous Navigation and Machine Vision Using The Hermies-IIB Robot," The 1987 Southeastern Regional Media Leadership Council (SRMLC) Conference on Microcomputers and Artificial Intelligence in Communication Technologies, UT at Chattanooga, October 29-31, 1987. Also accepted for publication in Winter Edition of IEEE Expert: Intelligent Systems and Their Applications. CESAR-87/46.

IV.9. D. Jollay and R. E. Ricks, "Robot Navigation Using Simple Sensor Fusion," Spring 1988 Decus, Cincinnati, Ohio, May 16-20, 1988. CESAR-88/02.

IV.10. M. Beckerman and E. M. Oblow, "Nonlocal Treatment of Systematic Errors in the Processing of Sparse and Incomplete Sensor Data," AAAI Workshop on Uncertainty, St. Paul, Minn. (August 19-21, 1988). CESAR-88/16.

IV.11. J-C. Culioli and V. Protopopescu, "Sigmoidic Linear and Nonlinear Optimization Algorithms," International Neural Network Society, Annual Meeting, September 6-10, 1988, Boston, Mass (1988). CESAR-88/23.

IV. INVITED TALKS (cont'd)

IV.12. B. L. Burks, P. F. Spelt, "Ongoing Research Using HERMIES - The Hostile Environment Robotic Machine Intelligence Experiment Series." The Materials Handling FOCUS '88, September 13-14, 1988. Georgia Institute of Technology, Atlanta, Georgia. (Loose-leaf Proceedings, no page numbers.) CESAR-88/50.

IV.13. B. L. Burks and P. F. Spelt, "Ongoing Research Using HERMIES - The Hostile Environment Robotic Machine Intelligence Experiment Series," DOE/ANL Training Course on The Potential Safety Impact of New and Emerging Technologies on the Operation of DOE Nuclear Facilities, August 29 - September 1, 1988, Idaho Falls, Idaho (1988). (Loose-leaf Proceedings, no page numbers.) CESAR-88/54.

IV.14. A. Sabharwal, A. El-Amawy, S. S. Iyengar, G. deSaussure, F. G. Pin and C. R. Weisbin, "Asynchronous Production Systems: Implementation Issues," COMCON 88: Conference on Advances in Controls and Communications, October 19-21, 1988, Baton Rouge, LA. CESAR-88/69.

IV.15. Martin Beckerman, "Sensor Data Fusion for Mobile Robots," IMS Special Topics Meeting on Industrial Statistics, Philadelphia, Pennsylvania, June 10-12, 1991. CESAR-91/05A.

V. INVITED PAPERS

V.1. J. Barhen, "Hypercube Concurrent Computation and Virtual Time for Robotic Applications," IEEE Workshop on Special Computer Architectures for Robotics and Automation, pp. 5.1-5.30. International Conference on Robotics Research, St. Louis, MO (March 25, 1985).

V.2. W. R. Hamel, "Manipulators in Teleoperation," American Nuclear Society Conference on Remote Operations and Robotics in the Nuclear Industry, Pine Mountain, GA (loose-leaf notebook, no page numbers) (April 21-24, 1985).

V.3. J. Barhen, M. C. G. Hall and J. R. Einstein, "Robotics Applications on an Advanced Hypercube Multiprocessor: Initial Developments," IEEE Workshop on Special Computer Architectures for Robot Control, 1986 International Conference on Robotics and Automation, San Francisco, CA (April 7-10, 1986). CESAR-86/13.

V.4. C. C. Jorgensen, "Control Algorithms for Autonomous Robot Navigation," Proceedings of the International Topical on Advances in Human Factors in Nuclear Power Systems, Knoxville, TN (April 21-24, 1986), pp. 100-106 (April 1986).

V.5. J. Barhen, "Combinatorial Optimization of the Computational Load Balance for a Hypercube Supercomputer," The Fourth Symposium on Energy Engineering Sciences, Argonne National Laboratory, Argonne, IL CONF-86-05122, Proceedings, pp. 71-80 (1986). CESAR-86/30.

V.6. W. R. Hamel and S. M. Babcock, "Manipulator Technology: The Critical Element of Useful Autonomous 'Working' Machines," Intelligent Autonomous Systems, An International Conference, Amsterdam, The Netherlands (December 8-11, 1986). CESAR-86/56.

V.7. J. R. Barhen, J. R. Einstein and C. C. Jorgensen, "Advances in Concurrent Computation for Machine Intelligence and Robotics," Second International Conference on Supercomputing, Santa Clara, CA (May 3-8, 1987). CESAR-87/10.

V.8. R. C. Mann, "Multi-Sensor Integration Using Concurrent Computing," SPIE Symposium on Infrared Sensors and Sensor Fusion, May 18-22, 1987, Orlando, FL, Proceedings, Vol. 782, pp. 83-90 (1987). CESAR-87/18.

V.9. E. M. Oblow, "O-Theory: A Probabilistic Alternative to Fuzzy Set Theory," Invited Paper, to be published in the series, "Lecture Notes in Computer Science," by Springer-Verlag, (Vol. title: "Uncertainty in Knowledge-Based Systems,") eds. Yager and Buchon (1987). CESAR-87/17.

V.10. W. R. Hamel, "Experiments in Autonomous Robotics," The 5th Symposium on Energy Engineering Sciences, Argonne National Laboratory, June 17-19, 1987. CONF-8706187, Proceedings, pp. 123-130 (1987). CESAR-87/21.

V.11. L. E. Parker, and F. G. Pin, "Architecture for Dynamic Task Allocation in a Man-Robot Symbiotic System," Proceedings, Vol. 851, pp. 95-102, 1987 SPIE Cambridge Symposium, Advances in Intelligent Robot Systems, November 1-6, 1987, Cambridge, Massachusetts. CESAR-87/37.

V. INVITED PAPERS

V.12. E. Wacholder, J. Han, R. C. Mann, "An Extension of the Hopfield-Tank Model for Solution of the Multiple Traveling Salesmen Problem," IEEE Annual International Conference on Neural Networks, San Diego, CA, July 24-27, 1988. CESAR-88/24.

V.13. C. R. Weisbin, B. L. Burks, J. R. Einstein, R. R. Feezell, W. W. Manges and D. H. Thompson, "HERMIES-III: A Step Toward Autonomous Mobility, Manipulation and Perception," The NASA Conference on Space Telerobotics, January 31-February 2, 1989, Jet Propulsion Laboratory, Pasadena, CA (1989). CESAR-88/65.

V.14. W. W. Manges, W. R. Hamel, C. R. Weisbin, J. R. Einstein, B. L. Burks, D. H. Thompson, R. R. Feezell, S. M. Killough, "Development of the HERMIES-III Mobile Robot Research Testbed at Oak Ridge National Laboratory," The ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, March 13-16, 1989, Proceedings, Session 2, pp. 2-3 (1989). CESAR-88/68.

V.15. M. Beckerman, J. P. Jones, R. C. Mann, L. A. Farkas, S. E. Johnston, "Spatial Reasoning in the Treatment of Systematic Sensor Errors," The SPIE Conference on Sensor Fusion: Spatial Reasoning and Scene Interpretation, Cambridge, MA, November 6-11, 1988, Proceedings, SPIE Vol. 1003, pp. 338-349 (1988). CESAR-88/25.

V.16. M. Beckerman, J. P. Jones, R. C. Mann and C. W. Glover, "World Modelling and Multi-Sensor Integration for a Mobile Robot," The ANS Third Topical Meeting on Robotics and Remote Systems, Charleston, SC, March 13-16, 1989, Proceedings, Section 11-2, pp. 1-8 (1989). CESAR-88/42.

V.17. K. W. Bowyer, J. P. Jones and C. H. Lake, "Computing the Hough Transform on an MIMD Hypercube," The 6th Scandinavian Image Analysis Conference in Oulu, Finland, June 19-22, 1989. CESAR-89/23.

VI. WORKSHOP PROCEEDINGS AND ORNL REPORTS

VI.1. Select Panel on Research Goals and Priorities in Intelligent Machines, Proceedings of a Workshop (November 1984).

VI.2. ORNL/TM-9923, CESAR-86/01. Proceedings of a Workshop: Planning and Sensing for Autonomous Navigation, Los Angeles, CA, August 17-18, 1985 (Edited by E. S. Howe and C. R. Weisbin) (February 1986).

VI.3. ORNL/TM-9759, CESAR-86/02. "O-Theory - A Hybrid Uncertainty Theory," E. M. Oblow (October 1985).

VI.4. ORNL/TM-10031, CESAR-86/05. "Modelling Early Stages of Human Vision," C. C. Jorgensen, R. Gawronski and F. Holly (August 1986).

VI.5. ORNL/TM-9732, CESAR-86/06. "Research and Development Program Plan for the Center for Engineering Systems Advanced Research (CESAR)," C. R. Weisbin, W. R. Hamel, J. Barhen and M. C. G. Hall (February 1986).

VI.6. ORNL/TM-9782, CESAR-86/07. "Robot Navigation Algorithms Using Learned Spatial Graphs," S. S. Iyengar, C. C. Jorgensen, S. V. N. Rao and C. R. Weisbin (December 1985).

VI.7. ORNL/TM-9987, CESAR-86/09. "ROSES: Robot Operating System Expert Scheduler," J. Barhen, P. C. Chen and E. Halbert (February 1986).

VI.8. ORNL/TM-9983, CESAR-86/10. "Foundations of O-Theory I: The Intersection Rule," E. M. Oblow (March 1986). (NOTE: see I.6)

VI.9. ORNL/TM-10232, CESAR-86/47. "Research and Development Program Plan for the Center for Engineering Systems Advanced Research," C. R. Weisbin, J. Barhen and W. R. Hamel .

VI.10. ORNL/TM-10329, CESAR-86/49. "On the Stability, Storage Capacity and Design of Nonlinear Continuous Neural Networks," A. Guez, V. Protopopescu and J. Barhen. See I.40.

VI.11. ORNL/TM-10107, CESAR-87/05. "Extension of O-Theory to Problems of Logical Inferencing," E. M. Oblow.

VI.12. ORNL/TM-10481, CESAR-87/20. "Navigation Planning Using Quadtrees," R. C. Fryxell.

VI.13. ORNL/TM-10493, CESAR-87/24. "Foundations of O-Theory II: Measurements and Relation to Fuzzy Set Theory," E. M. Oblow. Published June 30, 1987.

VI.14. ORNL/TM-10550, CESAR-87/39. "Supertracks: The Building Blocks of Chaos," E. M. Oblow.

VI.15. ORNL/TM-10558, CESAR-87/44. "Analysis of Tasks for Dynamic Man/Machine Load Balancing in Advanced Helicopters," C. C. Jorgensen.

VI.16. ORNL/TM-10633, CESAR-87/48. "Laboratory Telerobotic Manipulator Prototype Singularity Analysis," R. V. Dubey, J. A. Euler, S. M. Babcock and J. N. Herndon.

VI. WORKSHOP PROCEEDINGS AND ORNL REPORTS (cont'd)

VI.17. C. W. Glover and R. C. Mann, "Hypercube VLSI Supercomputers in Ballistic Missile Defense: Advanced Methodologies for SBKEW Fire Control," Final Report on FY 1987 Activities at ORNL (November 1987).

VI.18. ORNL/TM-10591, CESAR-87/57. "A Comparison of Algorithms for the Analysis of Images Generated by Two-Dimensional Electrophoresis." R. C. Mann, D. J. Riese, B. K. Mansfield, J. K. Selkirk,

VI.19. ORNL/TM-10679, CESAR-88/04. "A Computer Vision System for a Hypercube Concurrent Ensemble," J. P. Jones, R. C. Mann and E. M. Simpson (February 1988).

VI.20. ORNL/TM-10726, CESAR-88/09. "Path-Planning in a Known Environment with Dynamic Obstacles: Applications to the AATR," N. Sreenath (May 1988).

VI.21. ORNL/TM-10764, CESAR-88/19. "Adaptive Optimal Control of Uncertain Nonlinear Systems: On-Line Microprocessor-Based Algorithm to Control Mechanical Manipulators," Carlos March-Leuba and R. B. Perez (May 1988). See I.34.

VI.22. ORNL/TM-10763, CESAR-88/07. "Treatment of Systematic Errors in the Processing of Wide Angle Sonar Sensor Data for Robotic Navigation," M. Beckerman and E. M. Oblow (June 1988).

VI.23. ORNL/TM-10799, CESAR-88/49. "A Neural Network Algorithm for the Multiple Traveling Salesman Problem," E. Wacholder, J. Han, R. C. Mann (October 1988).

VI.24. ORNL/TM-10929, CESAR-88/52. "A New Algorithm for Linear Programming that is Easy to Implement: Application to the Transportation Problem." J. C. Culioli and V. Protopopescu.

VI.25. ORNL/TM-10935, CESAR-88/55. "Object Recognition for the HERMIES Robot," Kevin Bowyer.

VI.26. ORNL Internal Memo, CESAR-88/56. "The Fire Control Problem: Formalism and Possible Solutions," J. C. Culioli.

VI.27. ORNL/TM-11025, CESAR-88/66. "A Neural Network-Based Optimization Algorithm for the Weapon-Target Assignment Problem," E. Wacholder. See also I.40. (February 1989).

VI.28. ORNL/TM-11051, CESAR-89/01. "A Concurrent Approach to Model Based Motion Detection," H. I. Christensen (January 1989).

VI.29. ORNL/TM-10976, CESAR-88/33. "Bifurcating Optimization Algorithms and Their Possible Application," J.-C. Culioli, V. Protopopescu (November 1988).

VI.30. L. E. Parker and C. R. Weisbin, "1988 Workshop on Human-Machine Symbiotic Systems," Executive Summary published in the Proceedings of the 1988 Workshop on Human-Machine Symbiotic Systems, pp. 3-7, December 5-6, 1988, Oak Ridge, TN, ORAU 89/C-140. CESAR-89/19.

VI. WORKSHOP PROCEEDINGS AND ORNL REPORTS (cont'd)

VI.31. ORNL/TM-11198, CESAR-89/24. "Non-Linear Surface Fitting of Laser Range Images Using a Hypercube Concurrent Computer," Michael Finley (May 1989).

VI.32. ORNL/TM-11261, CESAR-89/43. "Matrix Reduction Algorithms for Gress and Adgen," J. E. Horwedel.

VI.33. ORNL/TM-11373, CESAR-89/44. "Sensitivity Analysis of AIRDOS-EPA Using Adgen with Matrix Reduction Algorithms," J. E. Horwedel, R. J. Raridon and R. Q. Wright (November 1989).

VI.34. ORNL/TM-11308, CESAR-89/34. "Job Planning and Execution Monitoring for a Human-Robot Symbiotic System," Lynne E. Parker (November 1989).

VI.35. ORNL/TM-11042. "Replacement Cost Integration Program Model Overview," D. B. Reister (February 1989).

VI.36. ORNL/TM-11132. "Replacement Cost Integration Program Model Description," D. B. Reister (March 1989).

VI.37. ORNL/TM-11124. "The Replacement Cost Integration Program User's Guide," D. B. Reister (March 1989).

VI.38. ORNL/TM-11046. "Validation of the Replacement Cost Integration Program," D. B. Reister and D. S. Christiansen (February 1989).

VI.39. ORNL/TM-11435. "A New Approach to Solve the Kinematics Resolution of a Redundant Robot," P.F.R. Belmans and J-C. Culioli (March 1990). CESAR-89/53.

VI.40. ORNL/TM-11530. "Laser Range Camera Modeling," Kai Storjohann (April 1990). CESAR-90/16.

VI.41. ORNL/TM-11510. "Time Optimal Trajectories for a Two-Wheeled Robot," D. B. Reister (May 1990). CESAR-90/12.

VI.42. ORNL/TM-11518. "Proceedings of the 1989 CESAR/CEA Workshop on Autonomous Mobile Robots," May 30-June 1, 1989. CESAR-90/15.

VI.43. ORNL/TM-11349. "Treatment of Systematic Errors II: Fusion of Ultrasound and Visual Sensor Data," M. Beckerman, L. A. Farkas, S. E. Johnston (June 1990). CESAR 89/38.

VI.44. ORNL/TM-11512. "Implementation of Valiant's Learnability Theory Using Random Sets," E. M. Oblow. (June 1990). CESAR-90/59.

VI.45. ORNL/TM-11635. "The Part Task Trainer for Airborne Weapons Systems: Human Factors Evaluation of the User Interface," Michele Terranover, Diane Hartley (September 1990). CESAR 90/43.

VI. WORKSHOP PROCEEDINGS AND ORNL REPORTS (cont'd)

VI.46. ORNL/TM-11661. "Adapting Sensory Data For Multiple Robots Performing Spill Cleanup," Kai Storjohann, Eric Saltzen. (September 1990). CESAR 90/43.

VI.47. ORNL/TM-11592. "Team-Computer Interfaces in Complex Task Environments," Michele Terranova. (September 1990). CESAR 90/49.

VI.48. ORNL/TM-11752. "A Rigid Body Model and Decoupled Control Architecture for Two Manipulators Holding a Complex Object," M. A. Unseren. (January 1991). CESAR 91/01.

VI.49. ORNL/TM-11741. "A Neural Network-Multiple Sensor Based Method for Recognition of Gene Coding Segments in Human DNA Sequence," E. C. Uberbacher, R. C. Mann, R. C. Hand, Jr., R. J. Mural. (February 1991). CESAR-91/06.

VII. STUDENT PAPERS

VII.1. Thomas Heywood (University of Florida), "HELIX: A Shared Memory Emulation System for Heterogeneous Multicomputing," Prepared in partial fulfillment of the requirements of the Oak Ridge Science and Engineering Research Semester under the direction of R. C. Mann and J. P. Jones, August 1989. CESAR-89/31.

VII.2. Sadruddin (Suds) S. Khawaja (Coe College), "Implementation of Fuzzy Reasoning Algorithms for the Navigation of an Autonomous Mobile Robot in an Unknown Environment," Prepared in partial fulfillment of the requirement of the ORNL-GLCA/ACM Science Semester under the direction of F. G. Pin, December 15, 1989. CESAR-89/56.

VII.3. Srirajasekhar (Bobby) Koritala (Coe College), "Fuzzy Inferencing Studies Based on a CMOS VLSI "Fuzzy Logic Inference Processor" Chip," Prepared in partial fulfillment of the requirement of the ORNL-GLCA/ACM Science Semester under the direction of F. G. Pin, December 15, 1989. CESAR-89/57.

VII.4. Jerry Kirk (Memphis State University), "A Control Panel Search and Docking Algorithm for the HERMIES-IIB Robot," Prepared in partial fulfillment of the requirements of the SCUU Science Semester under the direction of G. de Saussure, December 8, 1989. CESAR-89/54.

VII.5. Jabe Sandberg (Utah State University), "Object Recognition from Laser Range Camera Data Using Neural Network Classifiers," Prepared in partial fulfillment of the requirements of the Oak Ridge Science and Engineering Research Semester under the direction of C. W. Glover, March 1989 - September 1989. CESAR-90/41.

VII.6. Jamie Coleman (Albion College), "A Simulator For Robots in a Laboratory Environment," Prepared in partial fulfillment of the requirement of the ORNL-GLCA/ACM Science Semester under the direction of David Reister, December 14, 1990. CESAR-90/58.

INTERNAL DISTRIBUTION

1. B. R. Appleton
2. J. E. Baker
3. A. L. Bangs
4-8. M. Beckerman
9. R. J. Carter
10. J. R. Einstein
11-15. K. Fujimura
16. C. W. Glover
17. J. N. Herndon
18-22. J. P. Jones
23. H. E. Knee
24. G. E. Liepins
25-29. R. C. Mann
30-34. E. M. Oblow
35-39. F. G. Pin
40. D. B. Reister
41. J. C. Schryver
42. P. F. Spelt
43. F. J. Sweeney
44-48. M. A. Unseren
49-53. R. C. Ward
54-55. Laboratory Records Dept.
56. Laboratory Records, ORNL-RC
57. Document Reference Section
58. Central Research Library
59. ORNL Patent Section
60-61. EPMD Reports Office

EXTERNAL DISTRIBUTION

62. Office of Assistant Manager for Energy Research and Development, U.S. Department of Energy, Oak Ridge Operations Office, P. O. Box 2001, Oak Ridge, TN. 37831-8600
63. Dr. Peter Allen, Department of Computer Science, 450 Computer Science, Columbia University, New York, NY 10027
64. Dr. Wayne Book, Department of Mechanical Engineering, J. S. Coon Building, Room 306, Georgia Institute of Technology, Atlanta, GA. 30332
65. Professor Roger W. Brockett, Wang Professor of Electrical Engineering and Computer Science, Division of Applied Sciences, Harvard University, Cambridge, MA. 02138
66. Dr. James S. Coleman, Director, Division of Engineering and Geosciences, Office of Basic Energy Sciences, U.S. Department of Energy, Washington, DC 20585

67. Professor John J. Dorning, Department of Nuclear Engineering and Physics, Thornton Hall, McCormick Road, University of Virginia, Charlottesville, VA. 22901

68. Dr. Steven Dubowsky, Massachusetts Institute of Technology, Building 3, Room 469A, 77 Massachusetts Avenue, Cambridge, MA. 02139

69. Dr. Avi Kak, Department of Electrical Engineering, Purdue University, Northwestern Avenue, Engineering Mall, Lafayette, IN. 47907

70. Dr. James E. Leiss, 13013 Chestnut Oak Drive, Gaithersburg, MD. 20878

71-78Dr. Oscar P. Manley, Division of Engineering and Geosciences, Office of Basic Energy Sciences, U.S. Department of Energy, Washington, DC 20585

79. Professor Neville Moray, Department of Mechanical and Industrial Engineering, University of Illinois, 1206 West Green Street, Urbana, IL 61801

80. Dr. Michael J. Rabins, Department of Mechanical Engineering, 100 Engineering Physics Building, The Texas A&M University, College Station, TX. 77843-3123

81. Dr. Wesley E. Snyder, Department of Radiology, Bowman Gray School of Medicine, 300 South Hawthorne Drive, Winston-Salem, NC 27103

82. Professor Mary F. Wheeler, Department of Mathematical Sciences, Rice University, P. O. Box 1892, Houston, TX. 77251

83-92.Office of Scientific Technical Information, P. O. Box 62, Oak Ridge, TN. 37831