AI Magazine Volume 13 Number 2 (1992) (© AAAI)

Autonomous Mobile Robot Research at Louisiana State University's Robotics Research Laboratory

S. S. Iyengar and Daryl Thomas

■ *The Department of Computer Science at Louisiana State University (LSU) has been involved in robotics research since 1985 when the Robotics Research Laboratory (RRL) was established as a research and teaching program specializing in autonomous mobile robots (AMRS). Researchers at RRL are conducting high-quality research in amrs with the goal of identifying the computational problems and the types of knowledge that are fundamental to the design and implementation of autonomous mobile robotic systems. In this article, we overview the projects that are currently under way at LSU's RRL.*

Autonomous mobile robots (AMRs) are synthetic operational systems that can govern themselves while they accomplish given objectives and simultaneously manage their resources and maintain their integrity. The practical importance of research in AMRs has steadily increased over the last several years because of developments in technology that have brought such systems closer to reality. Already, industrial robots are being used in applications that involve monotonous or tedious tasks as well as applications in hazardous environments such as nuclear power plants. Next-generation robots are being planned for applications such as deep-sea mining and salvage operations, servicing-assembly tasks in space, and maintenance activities in toxic environments such as nuclear plants. Continued research in this area is certain to pay large dividends in the near future.

The various unstructured environments in which AMRs are expected to operate can change rapidly. AMR can unexpectedly encounter environmental threats that are hazardous to itself or its environment (for example, the outbreak of a fire or the failure of one of the robot's internal systems). Dynamic and complex environments such as these impose a need for specific capabilities in AMRs, including rapid multimodal sensing and integration, real-time response, real-time interruptability, and fault tolerance.

Facilities at the Robotics Research Laboratory

Much of the primary research of Louisiana State University's Robotics Research Laboratory (RRL) is conducted with the help of a specially modified Denning MRV (mobile robot vehicle) III research robot (figure 1). MRV III is an excellent test bed for robotics research: It has a straightforward user interface, it is easily modified to accept new features, and it has an impressive repertoire of capabilities. A unique feature of MRV III is its omnidirectional design that allows zero turning radius and pure translational motion without the need to rotate the chassis on axis when making turns. Other features include an ultrasonic ranging system, an infrared beacon navigation system, and a machine vision system. Another important feature is that it can be programmed to operate independently because it has a powerful 68000-series processor on board and is completely battery powered. It can also be operated as a slave to another computer through its external ports. Aside from MRV III, the lab is equipped with various robotic manipulators, several Macintosh IIs and PCs, and a HERO robot.

The reliability and accuracy of an automated interpretation system depends largely on the underlying labeling algorithm used.

RRL is continually upgrading its facilities to expand research capabilities. This upgrading is especially true of the Denning MRV III robot. RRL is currently adding several features to MRV III, most notably vision and manipulation subsystems. Also, as part of a joint research effort with the Department of Physics and Astronomy, RRL has access to a new, massively parallel processor for research into parallel computation. This machine is equipped with an impressive array of features, including a high-performance graphics subsystem for real-time imaging of computation results. In all, there is a wide range of equipment available for use by RRL researchers and students.

Current Research Efforts at the Robotics Research Laboratory

RRL investigators are currently exploring several avenues of research, including intelligent machine vision and sensing, spatial planning and navigation, asynchronous production systems, integration of knowledge sources into distributed systems, and architectures for the control of mobile robots. Many of these projects have involved collaboration with the National Aeronautics and Space Administration (NASA), Oak Ridge National Laboratory (ORNL), the Office of Naval Research, and the Jet Propulsion Laboratory (JPL)–California Institute of Technology (Cal Tech). To give the reader a feel for the nature of the work being done at RRL, several of these projects are summarized in the following subsections.

Knowledge-Based Feature Labeling of Oceanographic Satellite Images

This effort involves the design, implementation, and testing of a knowledge-based feature-labeling module intended to play a vital role in the development (by the Remote Sensing Branch of the Naval Oceanographic and Atmospheric Laboratory [NOARL]) of an automated image-interpretation system for satellite-produced oceanographic images.

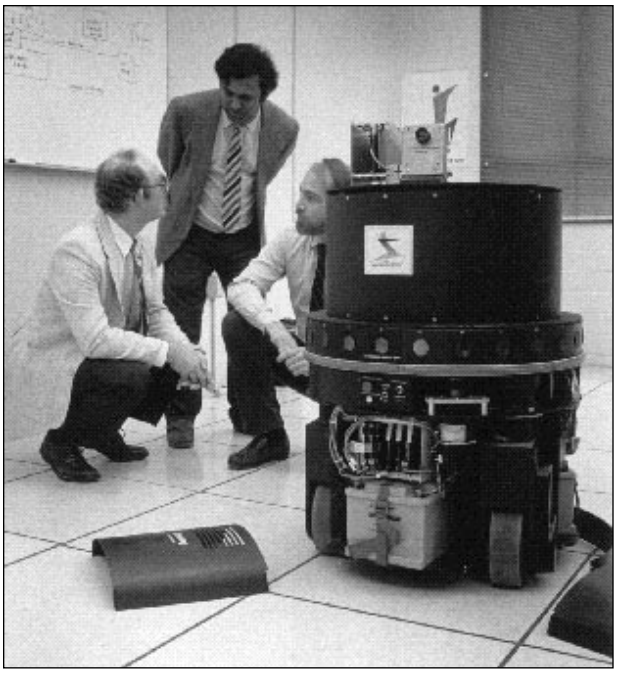

Figure 1. S. S. Iyengar (center) and Daryl Thomas (right) discuss plans for utilizing the MRV III robot's capabilities as a research test bed with a member of Louisiana State University's administrative faculty.

One of the primary missions of this group is to study and develop methods of exploiting satellite data to provide various types of oceanographic information for the United States Navy. From a large volume of satellite data, as well as nonsatellite data, a shipboard user of the image-analysis system must be able to produce operationally useful tactical products. Thus, an objective for NOARL is to transform the research techniques for the interpretation of satellite oceanographic data into operational procedures. To achieve this goal, dramatic improvements in the process of automated image analysis are required. Toward this end, NOARL has initiated the development of automatic detection and interpretation procedures, based on AI techniques, for these oceanic features.

The primary objective of this research effort is to build a powerful automatic image-interpretation system for oceanographic satellite images. To make this difficult problem tractable, it has been divided into two parts: (1) feature labeling and (2) developing an expert system. It is clear that the performance of the labeling algorithm depends heavily on the low-level image-processing algorithms used. Particularly, the output of an edge-detector algorithm plays a major role in the feature-labeling process. In view of this role, a new efficient edge-detector algorithm is used in RRL's feature-labeling technique (Krishnakumar et al. 1990c). It is known that the conventional edge-derivative operators are sensitive to noise and are not suitable for analyzing oceanographic satellite images. The new edge-detector algorithm is based on the gray-level co-occurrence matrix, which is commonly used in image-texture analysis. This algorithm is found to exhibit the characteristics of fine structure rejection while it retains edge sharpness.

Feature labeling (also called region labeling) is an important technique used in pattern recognition and image-understanding systems. It is a process that assigns labels or interpretations to the features obtained after image segmentation. The reliability and accuracy of an automated interpretation system depends largely on the underlying labeling algorithm used. RRL researchers proposed and developed a hybrid architecture that is based on a nonlinear probabilistic relaxation technique and an expert system to solve this difficult problem. The feature-labeling module consists of two submodules: Module one estimates the initial feature probabilities, and module two implements the iterative updating scheme. The initial pixel probabilities are assigned with the help of the previous analysis and the ground truth data. These probability values are then updated using an iterative updating rule that maximizes the consistency and minimizes the ambiguities. The labeled features are then fed to the expert system module.

The objective of this oceanographic expert system is to correctly interpret the dynamics of ocean processes with minimal human interaction. RRL researchers have developed a prototype expert system to generate projections of regional dynamics of mesoscale features in the North Atlantic Ocean (Krishnakumar et al. 1990a). One of the elements of this system is a knowledge database of these features represented in a form suitable for manipulation by the expert system's inference engine (that is, the logical and heuristic procedures for solving problems in the problem domain). This database was constructed based on knowledge gained from discussions with NOARL researchers as well as from the oceanographic literature. The expert system also evaluates the consistency of the labeling process with the help of the knowledge base as well as human input. This evaluation results in the assignment of a confidence factor to each of the features detected. This factor is then fed back to the feature-labeling module, along with the previous data analysis, to improve the consistency of the labeling process.

The feature-labeling software was implemented on a VAX 8300 using the I²S image-processing system at the NOARL Remote Sensing Laboratory. The software modules are written in C. The results of this research were published in two international journals (Krishnakumar et al. 1990c; Wang, Iyengar, and Patnaik 1989) and were also included as a chapter in *Advances in Artificial Intelligence* (Krishnakumar et al. 1990a). Various results are also presented in Krishnakumar et el. (1990b, 1989).

In the future, RRL researchers propose to make the oceanographic expert system learn from its past experience in analyzing the satellite imagery data and investigate the possibility of implementing a parallel relaxation-labeling algorithm to speed the labeling process. As a part of RRL's ongoing efforts, researchers have made an exhaustive survey of similar techniques used for automated interpretation in the areas of remote sensing, medical diagnosis, astronomy, and oceanography.

This research project was supported by a three-year grant (U.S. Navy, number N00014-88-K-6002) from NOARL at Stennis Space Center. The project will be supported for an additional three years to improve the labeling performance.

Participants: S. S. Iyengar, Ronald Holyer, Matthew Lybanon, Narayanan Krishnakumar, and S. Kumar

Computational Neural Learning Theory for Perceptual Manipulation

This project addresses a fundamental problem in computational AI: developing a new class of massively paral-

lel, computational neural learning algorithms for robustly abstracting complex nonlinear transformations (for example, functional, spatial, temporal, and statistical invariants) from representative samples in real time. The provision of such a capability is crucial to the solution of many difficult problems in robotics, signal processing, remote sensing, and adaptive control.

In contrast to existing dynamic neural learning formalisms, RRL researchers have proposed models that encode information as singular, rather than regular, solutions to neurodynamics using the notion of terminal attractors. The infinite local stability resulting from such encoding provides dramatic speedups during the learning process. RRL researchers have extended this framework to allow problem-dependent, adaptive evolution of network topologies. Extensive simulation results have demonstrated that RRL's model outperforms state-of-the-art neural learning formalisms by two to three orders of magnitude.

Although it is true that explicit mechanisms exist for encoding problem-specific constraints during the learning process, neural networks trained using such a priori skewing constructs can capture only limited aspects of the transformation of interest. Each different run-time interest requires retraining the network. To avoid excessive training and retraining computational costs, RRL researchers have exploited advances in renormalization group theory to present a methodology for modulating network response to a multiplicity of run-time constraints. This methodology does not require additional retraining or disturbing of the learned synaptic structure of the network.

The bulk of the computational cost incurred during any gradient–descent-based learning is expended in computing the parametric sensitivities, that is, dE/dp_μ. Existing learning algorithms require a system of N equations to be solved for each parameter p_μ, where N denotes the number of neurons in the network. Using the notion of adjoint operators from nonlinear sensitivity theory, RRL researchers have shown how to derive efficient algorithms wherein all the derivatives can be obtained by solving a single set of N appropriately constructed linear equations. RRL researchers have also analyzed the scalability of theoretical neural networks based on their model to large-scale embodiments in neural hardware.

Dynamic system formulations to neural network models are strewn with parameter-decay constants, response gains, and other constants. Until now, however, their selection has largely been accomplished by heuristics based on anecdotal exploration. RRL researchers conducted a systematic analysis of these parameters on the neurodynamics (for example, throughput and fault tolerance) during both concurrent simulations and implementation in concurrent very large-scale integrated (VLSI), optical, and optoelectronic hardware. For the first time, dynamic diagnostics, for example, Lyapunov exponents, are used to formally characterize the widely observed dynamic instability in neural networks as emergent computational chaos and broadband white noise. Using contracting operators and non-

constructive theorems in fixed-point theory, RRL researchers rigorously derived the necessary and sufficient conditions for eliminating all oscillatory and chaotic behavior. RRL researchers also derived neural algorithms for conditioning Cohen-Grossberg-Hopfield (additive-type) networks to operate under true concurrent asynchrony. Extensive empirical testing with arbitrarily large neural networks (over 100-million interconnects) demonstrated methodological robustness, even in the presence of large signal-propagation time delays.

Finally, RRL researchers provided insights for exploiting this powerful repertoire of adaptive learning formalisms to form an enabling core for addressing a fundamental problem in computational robotics: the design of autonomous robots to perform tasks in unstructured and unpredictable environments. Leveraging some recent results in task analysis and dynamic modeling, RRL researchers proposed a perceptual manipulation architecture. Conceptualized within a perceptual framework, the architecture was shown to be well beyond state-of-the-art model-directed robotics. A technical critique of the proposed architecture was presented, juxtaposing it with existing robot architectures (Gulati, Barhen, and Iyengar 1991). For a stronger physical interpretation of these implications, the discussions were embedded in the context of a novel system concept for automated space operations.

This research was supported by a grant from JPL–Cal Tech, number 958309 MOD3.

Participants: S. S. Iyengar, Sandeep Gulati, and Jacob Barhen

Event-Driven Expert Systems for Real-Time Applications

One area of AMR research showing particular promise is the *asynchronous production system* (APS), a concurrent, rule-based inference engine capable of monitoring and processing real-time information. APSs allow AMRs to make optimal judgments and take appropriate actions in response to uncertain input from the environment. They also allow AMRs to improve their performance over time by learning from experience.

Autonomous, unmanned vehicles and systems are subject to a set of unique and exacting computational requirements because of the absence of continuous human interaction. This fact necessitates the development of intelligent, self-sufficient inference systems to operate the system with minimum remote control. These computational requirements are especially stringent when the autonomous system is a mobile robot designed for hazardous terrains. For example, the on-board knowledge-based system on such a robot must be large enough to be applicable to a variety of operating environments; it must also be able to monitor and react to dynamic, unexpected events and provide guaranteed, intelligent responses to such events while it makes optimal use of limited on-board resources.

One area…showing particular promise is the asynchronous production system…a concurrent, rule-based inference engine…

The success of such a self-contained, intelligent computational system is intimately related to the development of a uniform, high-level programming environment that is integrated into all the real-time system components. RRL has contributed to this project by developing the concept of APSs and simulating an APS-based control architecture for the ORNL robot (Iyengar et al. 1992). This architecture includes a mechanism that provides for the distribution of monitored data throughout the system along with mechanisms that use heuristic-based specifications to interpret data from remote sources so that appropriate action can be taken when necessary. This system also includes modules for sensor preprocessing and integration, emergency handling and knowledge-based recovery, and assumption-based parallel planning for using metaknowledge for task allocation and resource utilization (Sabharwal et al. 1988).

The development of this system addressed numerous issues, including (1) the integration of the rule-based inference system with the real-time monitoring system; (2) the development of an interrupt-driven, concurrent inference engine; and (3) the development of intelligent simulation to respond to real-time events. As an adjunct to their work at ORNL, RRL researchers are also making use of the Lab's Denning MRV III robot to investigate these topics (Sabharwal et al. 1989).

This research was supported by a grant from the U.S. Department of Energy, through ORNL, number 19X-55902V.

Participants: S. S. Iyengar, F. G. Pin, Chuck Weisbin, Arvind Sabharwal, Vinayak Hegde, Jeff Graham, Phil Graham, and Shiva Subramanim

Theory of Fault-Tolerant Signal Integration for Dynamic Distributed Sensor Networks

A *distributed sensor network* (DSN) is a set of spatially scattered, intelligent sensors designed to derive data from the environment, abstract relevant information from the data gathered, and derive appropriate inferences from the information gained. DSNs depend on multiple processors to simultaneously gather and process information from many sources. Interest in these systems stems from a realization of the limitations imposed by relying on a single source of information to make decisions.

In recent years, there has been increasing interest in the development of dynamic DSNs for information gathering. This increased interest is in part because of the availability of new technology that makes them economically feasible to implement and in part because of the increasing complexity of today's information-gathering tasks to which they are applied. These tasks are usually time critical and depend on the reliable delivery of accurate information for their completion. To meet these requirements, DSN must be able to dynamically respond to fault conditions, reconfiguring itself as necessary to compensate for disturbances. Thus, the search for efficient, fault-tolerant architectures for DSNs has become an important area in computer science.

As part of RRL's work in this area, researchers are developing an advanced theory of integration for sensor control and information sharing between sensor units (Prasad et

al. 1991; Jayasimha, Iyengar, and Kashyap 1991). The primary goal of this work is the development of a new mathematical theory of sensor integration for networking spatially distributed sensors. This work includes the development of efficient algorithms for data combination, noise removal, and information abstraction as well as methods to process multispectral signals to isolate, retrieve, and enhance desired information, minimizing or compensating for background noise. This effort also involves improving the integration of multispectral sensor information in terms of internode coupling and data distribution, developing efficient algorithms for signal routing in distributed signal-processing networks, and providing fault-tolerant capabilities within the signal-processing architecture.

This research was supported by grants from the Louisiana Educational Quality Support Fund, number LEQSF-RD-040-JCP-03, and the Office of Naval Research.

Participants: S. S. Iyengar, R. L. Kashyap, S. V. N. Rao, Andrew Hoppe, Bush Jones, Lakshman Prasad, and Daryl Thomas

Fast Parallel Algorithms for a Special Class of Graph-Related Structures

As serial processors approach their limits, the push to develop parallel processors has intensified. This fact has made the development of new, parallel algorithms to run on these machines a top priority. Consequently, the design of fast parallel algorithms has increased in importance as a research area over the last several years.

Although parallel processing is currently more expensive to perform than sequential processing, with improved architectures and declining costs, it is certain to become the method of choice for almost all data-intensive, time-critical computations. Clearly, parallel algorithms are potentially much faster than their sequential counterparts. Unfortunately, industry's ability to design faster and cheaper parallel processors has far outpaced the ability to use them effectively. One reason is that there is no simple, straightforward method for converting a given serial algorithm into a parallel one. The development of a parallel algorithm to solve a particular problem often requires extensive analysis to determine the most advantageous way to decompose the problem for parallel solution. One must consider load balancing, communication costs and delays, processor use, and many other variables. Another reason for low use is the great diversity of parallel architectures, each requiring its own method of programming and data organization. On first inspection of a problem, it is not always apparent which architecture is the most appropriate on which to implement a parallel solution.

In light of this discussion, it should be clear that the development of parallel algorithms is a nontrivial task. It is in this context that RRL researchers are conducting research both to develop new, fast parallel algorithms for certain useful computational problems on special classes of perfect graphs and to analyze the complexity of the algorithms developed (Iyengar 1990). RRL researchers are also endeavoring to classify the developed parallel algorithms into possibly new complexity classes. A collateral goal of this effort is to gain a deeper insight into, and a better understanding of, the structure of perfect graphs.

Participants: S. S. Iyengar, Sridhar Radhakrishnan, and Rajanarayanan Subbiah

Solid Modeling for Robotics Applications

The focus of this project is to explore and develop methods of generating solid model images of the environment for use in robot navigation. Image generation involves both shading computation to determine the appearance of the visible objects and hidden-surface computation to determine which objects in the simulated environment are visible and which are obscured.

RRL researchers are currently working with ray-tracing methods to model solid objects using combinatorial geometry; that is, solid objects are modeled as compositions of primitive solids and combined using the Boolean set operators union (+ &) and difference (-). Examples of primitive solids are blocks, spheres, cones, and cylinders. A solid composition is represented as a binary tree in which the leaves are primitives.

At the heart of the ray-tracing process is the computation of the intersection of imaginary light rays emanating from the viewer with the set of objects in the sense or world model. The ray-tracing process is used in several different contexts in the image-generation process. For example, in hidden-surface computation, rays are traced from the viewer, through the viewer's picture plane, and into the three-dimensional environment to determine the front-most visible objects. These rays are often called viewer rays, hidden-surface rays, or first-level rays. The drawback of the ray-tracing process is that it is computationally intensive: It would take hours of central processing unit time on a serial processor to display any scene of reasonable complexity. Consequently, RRL researchers are attempting to reduce the computation costs by distributing the process of intersection computation over several processors.

A basic routine for intersection computation was written for ray-spear (primitive object) intersection. RRL researchers also developed routines to display a single spherical object using the diffuse reflection model. Routines are currently being written to display a scene consisting of several spheres. The constructive solid geometry routines will be written on top of these routines. To display images with greater realism, properties of reflection and transmission are being included, making the process even more intensive. A graphic interface based on the machine-dependent software tool kit is also under development for use on DEC workstations.

Participants: S. S. Iyengar, Mahesh Dandapani, Ramesh Phatak, V. Sridhar, Nitin Naik, and Wu Wang

Trajectory Planning of Robot Manipulators Using Stochastic Automata

The main thrust of this project is to develop new trajectory planning methods for multilink, sensor-based robot manipulators operating in noisy work spaces. When it is said that a work space is *noisy*, it means that all sensing operations performed in the work space (to determine the position and orientation of the manipulator and its effector) are subject to measurement error because of background interference. When a manipulator is to operate in a noisy environment, the task of moving the

manipulator's effector from some initial position, say, P_i, to some final position, say P_f, is not a trivial one because the determination of P_f must be done in the presence of noise-induced measurement error.

For the purposes of this project, RRL researchers are interested in developing algorithms for the operation in an environment where manipulators have prismatic, revolute joints. The algorithms that are under development to solve this problem are based on the use of learning automata and do not require the estimation of parameters, the computation of inverse kinematics, or the use of hardware-oriented feedback control (Oommen, Andrade, and Iyengar 1991). The method involves the use of a learning automaton at every joint of the manipulator, which means maintaining a finite-state machine at each joint to dictate the motion of the joint. In this way, the motion of the joint can be controlled using only repeated (possibly noisy) observations of P_f . Further, these automata operate in parallel.

Participants: S. S. Iyengar, B. J. Oommen, Nicte Andrade, and Vinayak Hegde

Computational Paradigms for Robot Vision Problems

Machine vision is a central feature of any autonomous robot that must monitor and react in real time to dynamic and unexpected events in the environment. In this context, the major task of a machine vision system is to process image data and produce abstract interpretations of objects present in the scene. Although much work has been done on machine vision, few methods seem to be satisfactory when applied to real-time applications. In an effort to address certain of these problems, RRL is expanding its research in this area (Wang, Iyengar, and Patnaik 1989).

With the arrival of LSU's new MAS PAR parallel processor, the RRL research area will be broadened to include development and testing of new parallel algorithms for vision processing. In addition, there are plans to acquire state-of-the-art vision equipment to provide the lab's Machine Vision Group with the capability to generate and process high-resolution stereo color images. In addition to this equipment, plans include the acquisition of a sophisticated laser ranging device. This equipment will be integrated into a multimodal system used to investigate computational problems and techniques related to the interpretation of information from multiple sensors in the context of dynamic environments (that is, isomorphic matching of representations, merging of different representations, and automatic generation of higher-level interpretations). RRL researchers will also work on developing new software architectures for the integration of these processes.

In particular, RRL researchers plan to conduct research into new techniques for object recognition and perception by machine vision systems, information-integration strategies for machine vision systems, knowledge-based systems for robotic planning and spatial databases, and visualization systems for spatial databases.

Participants: S. S. Iyengar, Charles Harlow, Ramesh Phatak, Nitin Naik, and Krishna Kumar

New Techniques for Sensor-Based Autonomous Navigation with Learning

This project explores methods of robot navigation that require no prelearned model for use in unknown, dynamic terrains. These methods involve making maximal use of available information, recording and synthesizing information from multiple journeys, and applying concepts of learning that allow for continuous transition from local to global path optimality (Rao et al. 1986).

The problem of robot navigation has been studied extensively and reported in the literature. These works mainly fall into two categories: (1) navigation in known terrains and (2) navigation in unknown terrains. *Navigation in known terrains* is characterized by the fact that the complete model of the terrain is known a priori for the robot. Navigation algorithms exist for a robot required to navigate in a known terrain. The robot is initially given the world model, and path planning is done using the world model. For safety, it can be provided with proximity sensors. Navigation in known terrains can lead to easy and efficient solutions (Rao and Iyengar 1990), but their application is limited.

Algorithms for navigation in unknown terrains operate on sensor information and, thus, are considerably dependent on the performance of the robot's sensory system. Many of these algorithmic structures (Rao and Iyengar 1990) are based on rather strong assumptions, such as infinite sensor range and absolute sensor accuracy. In real-life applications, these assumptions prove invalid. This work involves the unification of environment learning with path planning under the realistic assumption that sensors have limitations that must be accounted for (Rao and Iyengar 1990).

Participants: S. S. Iyengar, N. S. V. Rao, and Gili Mendel

Design Automation Algorithms for Very Large-Scale Integrated Applications

Efficient chip design is an important facet of semiconductor lithography because it affects both the performance and the economics of the finished product. In VLSI design, one aspect of efficiency is determined by how well the routing problem can be solved, that is, realizing a particular interconnection among different modules in as small an area as possible. RRL researchers are currently conducting research into improved algorithms for solving such VLSI routing problems (Ho and Iyengar 1991). This work could lead to significant advancements in the production of VLSI circuitry, which, in turn, might affect the development of AMRs.

There are many different routing strategies for finding efficient interconnections among different modules of a VLSI chip. One of the most important routing strategy forms is *channel routing*, which allows the routing problem to be reduced to a collection of simpler problems. A channel router is designed to route nets that interconnect terminals on two opposite sides of a rectangular region called a channel. Typically, a virtual grid is assumed, and a Manhattan model is adopted; that is, all the horizontal segments are routed in one layer and all the vertical segments in another. Because the interconnection area represents 65 to 80 percent of the total area in a typical polycell integrated circuit design, the primary goal of a channel router is to minimize the area by limiting the number of tracks used. The number of vias and the length of the nets are also important in evaluating the

quality of the routing.

The algorithm being developed at RRL approaches the solution systematically in a greedy way based on simple concepts (Ho, Iyengar, and Zheng 1991). The current focus is on working with the Manhattan model. The advantages of the RRL algorithm are its simplicity and generality. The only data structure used is a vertical constraint graph together with column density information and spans of nets. Because of its general nature, the algorithm treats cyclic and noncyclic problems identically. For this reason, it can easily be expanded to function as a switch-box router, a three-layer channel router, or a multilayer channel router. Additionally, it can be applied in an overlapped routing environment without extensive modification.

The key idea underlying this algorithm is track-by-track routing. First, all possible routings for each net in a certain track are calculated, a process that takes $O(nc)$ time, where n is the number of nets, and c is the number of columns used. Depending on the channel density (d_{max}) and the structure of the vertical constraint graph, nets are then chosen that minimize either d_{max} or the length of the directed longest path of the vertical constraint graph while they maximize the total length of significant horizontal wires in a track based on some priority. The performance of the algorithm is outstanding. It successfully routed both Burstein's switch-box problem and Deutsch's example in the Manhattan model without any backtracking.

Participants: S. S. Iyengar, Tai-Tsung Ho, and Si-Qing Zheng

Future Directions

RRL's research addresses problems that are fundamental to the implementation of autonomous intelligent machines. This research is comprehensive, covering most aspects of intelligent systems. Research at RRL will continue to be directed along these lines, with expansion into related areas as appropriate. A natural extension of this research would be to explore the application of AI techniques to other problem areas. One such area is using expert systems to diagnose real-time process control systems. Another area of planned research includes the study of robotics applications for individuals who are motorically disabled. The potential benefits to these individuals from the application of advanced robotic technology are underexplored.

Acknowledgments

Funds for programs at RRL have been provided by the state of Louisiana; JPL–Cal Tech; the U.S. Navy, Naval Ocean Research and Development Activity; the U.S. Department of Energy, ORNL (Martin Marietta); the Office of Naval Research; NASA; the National Science Foundation (Laser Program); and Apple Computer. In particular, the authors would like to thank Chuck Weisbin and Jacob Barhen of JPL–Cal Tech and R. Mann and F. G. Pin of ORNL.

References

Gulati, S.; Barhen, J.; and Iyengar, S. S. 1991. Neuro Computing Formalism for Computational Learning and Machine Intelligence. In *Advances in Computers*, volume 33, ed. M. C. Yovits, 173–245. San Diego, Calif.: Academic.

Ho, T., and Iyengar, S. S. 1991. New Results on Channel Routing. In Proceedings of the VLSI Design '91 Conference, 174–179. Washington, D.C.: IEEE Computer Society.

Ho, T.; Iyengar, S. S.; and Zheng, S. 1991. A General Greedy Algorithm for Channel Routing Problem. *IEEE Computer Aided Design* 10(2): 204–211.

Iyengar, S. S. 1990. Fast Parallel Algorithms for Processing Graph-Related Data Structures. Presented at the Fourth Annual Symposium on Parallel Processing, 30 April–4 May, Fullerton, California.

Iyengar, S. S.; Sabharwal, A.; Pin, F.; and Weisbin, C. 1992. Asynchronous Production Systems for Control of an Autonomous Mobile Robot. *Applied Artificial Intelligence*. Forthcoming (May issue).

Jayasimha; Iyengar, S. S.; and Kashyap, R. L. 1991. Information Integration and Clock Synchronization in Distributed Sensor Networks. *IEEE Transactions on Systems, Man, and Cybernetics* 21(5). Forthcoming.

Krishnakumar, N.; Iyengar, S. S.; Hoyler, R.; and Lybanon, M. 1990a. An Expert System for Interpreting Mesoscale Features in Oceanographic Satellite Images. In *Advances in Artificial Intelligence: Applications and Theory*, volume 27, ed. James C. Bezdek, 342–355. Singapore: World Scientific Series.

Krishnakumar, N.; Iyengar, S. S.; Hoyler, R.; and Lybanon, M. 1990b. An Expert Systems for Interpreting Mesoscale Features in Oceanographic Satellite Images. In Application of Artificial Intelligence, 184–194. Bellingham, Wash.: The Society of Photo-Optical Instrumentation Engineers.

Krishnakumar, N.; Iyengar, S. S.; Hoyler, R.; and Lybanon, M. 1990c. A Technique for Labeling Mesoscale Features in Oceanographic Images. *Journal of Image and Vision Computing* 8(2): 142–147.

Krishnakumar, N.; Iyengar, S. S.; Hoyler, R.; and Lybanon, M. 1989. A Technique for Feature Labeling in Infrared Oceanographic Images. Presented at the Fifth International Conference on Intelligent Processing Systems for Meterology, 11–16 December, Anaheim, California.

Oommen, J.; Andrade, S.; and Iyengar, S. S. 1991. Path Planning of Robot Manipulators in Noisy Workspaces Using Stochastic Automata. *International Journal of Robotic Research* 10(2): 135–148.

Prasad, L.; Iyengar, S. S.; Kashyap, R. L.; and Madan, R. 1991. Functional Characterization of Sensor Integration in Distributed Sensor Networks. *IEEE Transactions on Systems, Man, and Cybernetics* 21(5). Forthcoming.

Rao, N. S. V., and Iyengar, S. S. 1990. Autonomous Robot Navigation in Unknown Terrains: Incidental Learning and Environmental Exploration. *IEEE SMC* 20(6): 1443–1449.

Rao, N. S. V.; Iyengar, S. S.; Jorgensen, C. J.; and Weisbin, C. 1986. Robot Navigation in an Unexplored Terrain. *Journal of Robotic Systems* 3(4): 389–407.

Sabharwal, A.; Iyengar, S. S.; Weisbin, C.; and Pin, F. 1989. Asynchronous Production Systems for Real-Time Expert Systems. *Journal of Knowledge-Based Systems* 2(2): 117–127.

Sabharwal, A.; Iyengar, S. S.; Weisbin, C.; and Pin, F. 1988. Asynchronous Production Systems for Real-Time Expert Systems. Presented at the Eighth International Workshop on Expert Systems and Their Applications, 6 June, Avignon, France.

Wang, W.; Iyengar, S. S.; and Patnaik, L. M. 1989. Memory-Based Reasoning Algorithm for the Recognition of Binary Images. *Journal of Pattern Recognition* 22(5): 505–518.

S. S. Iyengar, professor and chairman of the Computer Science Department, directs the Robotics Research Laboratory at Louisiana State University. His research addresses problems that are fundamental to the implementation of autonomous intelligent machines. This research is comprehensive, covering most aspects of intelligent systems. He has published over 150 research papers and conducted important work on the problems of intelligent robot navigation, intelligent sensing, world modeling using acoustic sensors, and real-time expert systems for the control of mobile robots.

Daryl Thomas is currently a research assistant at the Robotics Research Laboratory (RRL) at Louisiana State University, pursuing a Ph.D. in computer science. He is also laboratory manager for RRL. His interests include machine vision and AI.