SRC TR 87-121

Systems Research Center

by

J. S. Baras

Systems Research Center

JOHN S. BARAS

INTRODUCTION

The Systems Research Center (SRC) at the University of Maryland, College Park, and Harvard University, Cambridge, Mass., is a new forum for fundamental research and education in systems engineering.

Established by a National Science Foundation (NSF) grant in 1985, the SRC is developing advances in design methods and software systems that innovatively address the challenges of productivity and competitiveness facing American industry. At the same time, the Center is training a new generation of systems engineers in an environment that is designed to both expand intellectual frontiers and achieve important research objectives.

SRC Goals and Themes

The impetus behind the activities of SRC is a close and mutually supportive collaboration with industry and government. The Center's programs are designed for two purposes: (1) to make the best possible use of the expertise and interests of an interdisciplinary team of faculty members and available private and government research personnel and facilities, and (2) to meet objectives consistent with the competitive needs of business and industry.

From The Engineering Research Centers: Leaders in Change, National Academy Press, Washington, D.C., 1987.

The Center's research theme is to promote basic study in the applications and implications of advanced computer technology—very large scale integration (VLSI), computer-aided engineering (CAE), and artificial intelligence (AI)—in the engineering design of high-performance, complex automatic control, and communication systems. Its research activities are built around five interrelated focus application areas: chemical process control, expert systems and parallel architectures, manufacturing systems, communications systems and signal processing, and intelligent servomechanisms.

Traditional industries as well as high-technology industries depend in a critical way on automation and information processing systems. As the complexity and demand for these systems has increased dramatically in the past decade, it has become obvious that the modeling and design methodologies of the past are no longer adequate. More emphasis is needed on the modeling and empirical/experimental components of systems science and engineering. In addition, sophisticated system-level design tools are needed to integrate the analytical and computational techniques of control and communication engineering with advances in computer hardware and software. These system-level design tools will increase the productivity and efficiency of engineers and will facilitate teamwork. The Center's programs represent a premier example of the use of advanced computer technology as an "amplifier" of human engineering skills and ingenuity.

The Center's educational goals are to support and enhance education programs and to serve as a source of new courses and material. In the process, the Center is seeking to change the traditional focus of engineering education by placing a new emphasis on both education and training. The Center's broad, interdisciplinary programs, offered in cooperation with Harvard University's Division of Applied Sciences, cut across the boundaries of many engineering and computer science disciplines, and they are designed for interactive participation by America's foremost corporations. The ultimate goals are to gain new knowledge; to train the engineers who can apply this knowledge to a diverse set of complex, real-world problems; and to speed the transfer of research results to the industrial community.

To enhance the interaction among the academic, industrial, and government research communities, an innovative and broad industrial collaboration program has been established. It includes joint research projects, industrial visitors to the university, faculty and student visitors to industry, joint use of laboratories, fellowship programs with industry, intensive short courses and workshops, colloquia, seminars, a software library, and a unique software research "club."

In this paper I will describe the progress that has been achieved to date in the research, educational, and industrial collaboration programs of the Systems Research Center. I will also provide a brief preview of our plans for the second year.

Initial Goals

Let me first briefly recall the goals we set for SRC upon its creation:

- Pursue integrated design of complex automation and information processing systems. Create a computer-aided design (CAD) environment where complete integration can be achieved, from conceptual development to technology selection, hardware implementation, testing, and validation.
- Emphasize modeling, empirical, and experimental issues. It is exactly this component of systems science and engineering that we felt had been neglected over the past 15 years or so. In particular, we wished to emphasize modeling of the new hardware systems used to implement the control or information processing algorithms.
- Develop design theories and tools that incorporate advances in computer technology (hardware and software). This goal was based on the realization that control and communication system design methodologies were not in synchrony with the currently available or planned implementation media.
- Provide real engineering tests for design. We wanted to increase the awareness of both faculty and students about the difficulties inherent in executing and validating a design.
- Provide engineers with system-level design tools that are able to increase design productivity, quality, and efficiency. These tools should include the heuristic/empirical component of the design process, as well as the analytical/numerical sophistication of the new theories.

OVERVIEW OF ACCOMPLISHMENTS

It has been a very exciting year at SRC. Building the organization, infrastructure, and mechanisms needed to support the various programs has been a demanding but intellectually rewarding endeavor for all involved. Let me give an overview of our accomplishments during this beginning period. We initiated a broad, innovative, interdisciplinary research program that I will describe in greater detail in a later section. We undertook several educational innovations and developed advanced design teaching laboratories. We established a broad industrial collaboration program. We began the development of several advanced research laboratories. And we completed the organizational and administrative structure of the Center.

We paid particular attention to developing appropriate policies for faculty and student participation. We firmly believe that early recognition of the significance of these policies is a critical component for the long-range success of such centers. SRC involves 2 universities, 2 colleges, 6 departments, 36 faculty, 50 graduate students, and 30 undergraduate students. It is apparent that without a strong "joint venture" spirit supported by such policies, long-range success may not be possible. The Administrative Council (which includes the chairpersons of the major participating departments, the deans, representative executive officers from industry, and the SRC director) proved to be a very efficient vehicle for formulating and implementing such policies.

We have already made great strides in diversifying the financial support for the Center. In addition to funds from NSF, funds are being provided by other federal and state agencies, the universities, and industrial sponsors. We are grateful to industry for the spontaneous response, both in funds and equipment gifts, that they displayed during our first year of operation. In the first year we received funding of about $2 million from NSF, $0.4 million in funds and more than of $2 million in equipment gifts from industry, and $0.3 million from the state of Maryland and the two universities. In addition, the universities provided space and faculty positions. Eight new faculty have been added at Maryland and one at Harvard as a result of the creation of SRC.

Let me close this section by summarizing the response as I observed it during the first year. From academic faculty the response has been enthusiastic; many more people want to participate in the SRC program than we can currently accommodate. From industry I can characterize the response as extremely supportive and as a "long-awaited dialogue." We have been very encouraged by the convergence of ideas and plans between industrial researchers and our academic faculty on the subject matter of our research. Finally, the response from students has been overwhelming. It has created an unprecedented excitement on campus, so that we have had five times as many applicants as we could accommodate. In particular, the possibilities for real industrial interaction proved to be extremely attractive to students.

THE RESEARCH PROGRAM

As described elsewhere in further detail (see References), the research program of the Systems Research Center has been evolving around five focus application areas. During the first year, serious efforts were undertaken in all areas to introduce computer-aided engineering, along with experimental and empirical design issues, and to cross the traditional boundaries between the multitude of disciplines represented at SRC. In addition, we put particular emphasis on the involvement of undergraduate students in design projects as early as possible in their careers. As emphasized elsewhere (Baras, 1986), the subject of fundamental studies carried out at SRC is in a sense systems science and engineering itself, and not its particular application or manifestation in a narrow application area. This in itself justifies the emphasis that we have placed on system-level design tools for automation and information processing systems.

Our findings throughout the year not only reinforced some of our beliefs and expectations in the value of our approach to systems engineering but also rewarded us with the discovery of some unsuspected connections between totally unrelated engineering problems. Several key ideas and concepts have already emerged: (1) a pervasive utility of optimization-based design in a great variety of engineering systems; (2) the value of AI in the development of systems that can reason and that can aid in the design process; (3) the superior efficiency and ability of symbolic algebra as an engineer's aid in complex calculations and modeling; (4) the critical importance of representation and manipulation of engineering data bases; (5) the mandatory utilization of AI techniques to handle heuristics in the design process and to enhance and facilitate teamwork; (6) the mandatory utilization of interactive graphics as the interface mechanism between the systems engineer and the computer; and (7) the critical importance of understanding VLSI systems, their architectures, and their limitations.

It is impossible in the limited space and time available to describe the multitude of exciting projects currently undertaken by SRC faculty and students. Instead, I will present an overview of the research program and will highlight certain sample projects to illustrate the nature of our work and our findings. The selection of projects discussed reflects, to a certain degree, my finite memory and capacity to represent the details of many excellent projects that are under way and reflects my perception of projects wherein intellectual integration across disciplines has occurred most extensively.

First, the focus areas and the thrusts within each area are as follows:

1. *Intelligent Servomechanisms.* Major thrusts are the design of robust control systems with many sensors and many feedback loops and, in particular, advanced robotic manipulators and flight controllers for advanced aircraft and spacecraft.
2. *Chemical Process Systems.* Major thrusts are modeling and control of industrial processes and the integration of reliability and safety in the computer-aided design process.
3. *Manufacturing Systems.* Major thrusts are the integration of CAD with manufacturing resources planning (MRP), scheduling and resource allocation problems in flexible manufacturing systems, and applications of AI in manufacturing.
4. *Communication and Signal Processing Systems.* Major thrusts are the modeling, design, and control of computer and communication networks, image processing, and speech processing and recognition.
5. *Expert Systems and Parallel Architectures.* Major thrusts are VLSI systems design and architectures for control and communication systems, expert systems for control and signal processing, and reliability integration in computer-aided design.

From the beginning we have emphasized the need to design and develop our specialized research laboratories so as to create a sophisticated environment for integrated design. The constituent laboratories are a key concept in our development of SRC. These laboratories form the natural home for interdisciplinary groups of faculty and student researchers. In a sense the entire research and educational program of SRC evolves around these constituent laboratories. I would like to describe briefly the components of this environment. We have developed the notion of a system engineer's workstation, which is really a design super-workstation combining an AI machine with a graphics engine and a multiprocessor "number cruncher." This super-workstation is part of a network of other workstations and computers (both similar and different) so that the engineer can have the capability of running concurrently several modules of his or her design software system on different machines. The AI workstation provides the direct interface with the user for (often symbolic) problem description and modeling. The multiprocessor number cruncher provides the necessary computing power for almost real-time execution. Finally, the graphics engine provides real-time graphics for simulation, testing, validation, and feedback to the designer.

We have also emphasized the integration of symbolic and numerical computation. We are convinced that symbolic languages such as LISP, PROLOG, and MACSYMA offer a superior medium for definition, conceptualization, and implementation of design problems. They are also superior for modeling engineering systems. An additional advantage offered by these languages is their superiority as universal communication tools between engineers and scientists from diverse disciplines and backgrounds.

We have introduced and implemented sophisticated simulation tools. These include analytical, software, and hardware tools. For example, we plan to use critical sampling theory for fast Monte Carlo simulation in computer/communication networks and manufacturing systems. The mathematics justifying such techniques, which reduce simulation times by several orders of magnitude, are quite sophisticated. We also plan to use LISP-based and object-oriented programming for high-level simulation of chemical plants, flexible manufacturing plants, and communication networks. All laboratories will be linked to real-data experiments for testing and validation of proposed designs. We plan to rely heavily on AI and expert systems, to handle the routine and heuristic part of the design automation, and on graphics for the man-machine interface. Examples of the environment we are creating will be given below in the description of selected research projects.

Intelligent Servomechanisms

In the area of intelligent servomechanisms we had the following projects: bifurcation control and multiparameter singular perturbation (with applications to flight control) (Profs. Abed, Krishnaprasad, and Tits); complex analytical methods for design of controllers and signal processing schemes (Profs. Berenstein and Baras); hand-eye machine (Profs. Brockett, Maragos, and Wohn); nonlinear control and robotic manipulators (Profs. Krishnaprasad, Berenstein, Tits, and Abed); and optimization-based CAD (Profs. Tits, Krishnaprasad, Baras, and Levine).

I would like to provide some details about the project on nonlinear control and robotic manipulators, which is under the direction of Prof. Krishnaprasad. Research in this project ranges from specialized hardware design and construction to very sophisticated theories of nonlinear dynamics of multibody systems. Progress has been achieved in several directions. Professor Krishnaprasad, in collaborative work with Prof. J. E. Marsden (University of California, Berkeley) and Prof. Juan Simo (Stanford University), worked out the Hamiltonian structure of systems of rigid bodies with flexible attachments that obey finite strain elastic models and that are geometrically exact. These results should have a significant impact on the control theory of multibody spacecraft and on the modeling and control of flexible robotic manipulators. The researchers have also been able to establish the Poisson structures underlying a wide variety of problems in interconnected rigid body systems and have used these to obtain stability criteria for various equilibria. Certain examples have been simulated and displayed graphically on our IRIS workstation to reveal for the first time the beautiful topological structures of the phase portraits. Our results here should prove useful in the study of control problems for articulated spacecraft and robotic manipulators operating on space platforms (e.g., the Space Station Mobile Remote Manipulator System, or MRMS).

A group of 10 students was involved in a design project under the sponsorship of National Aeronautics and Space Administration (NASA) Headquarters, NASA Goddard Space Flight Center, and the Systems Research Center. The project is under the supervision of Prof. P. S. Krishnaprasad and is 1 of 18 projects under way in universities across the nation as part of the NASA-funded pilot program in advanced space systems design. A primary goal of this NASA program is to encourage and strengthen design-related activities in university curricula at the undergraduate and graduate levels. The goal of the project at the University of Maryland is to design a Mobile Remote Manipulator System for the Space Station. Students participating in this project receive academic credit.

Throughout the year, strong collaboration continued with Harvard in the area of robotic manipulator design. In the project directed by Prof. Brockett, a 9-degree-of-freedom hand is now operating under closed loop position control and 3-degree-of-freedom force control as well. It is currently being used in the Automatix AID600 robot, and software is being written which will make the hand programming easier. Victor Eng's thesis work involves the incorporation of modes into a hand programming language. It is our belief that the control of multi-degree-of-freedom systems will be greatly facilitated by the identification of coordination patterns, or modes, which, while not completely general, allow one to execute frequently occurring motions in a simple way. The situation is analogous to the use of glyphs in graphics—in fact, one can think of modes as "motion glyphs." A closely related idea involves the coordination of sensory information. Here, sensory glyphs are identified with those combinations of sensory output data which occur frequently so as to make it possible to easily set up feedback control loops for regulators. This feature is also being incorporated into the hand control programming environment under development.

A distributed tactile sensor is being developed at the Naval Research Laboratory in support of the distributed sensor work at SRC (Prof. Shamma). This is an all silicon design based on a clever application of piezoresistive and piezoelectric phenomena in conjunction with an amplifier/multiplexer chip for local processing so as to recover estimates of the three components of the applied force field. The sensor can also detect slipping between the manipulator gripper surface and the object. Strips of the material can be attached to the fingers of a manipulator hand, and signal processing is performed by neural-type processor chips.

Chemical Process Systems

In the area of chemical process systems we had the following projects: automation of dynamic process simulation (Prof.

Cadman); modeling, simulation, and control of chemical reactors (Prof. Choi); application of expert systems to distillation column control (Prof. McAvoy); and knowledge-based expert systems for chemical plant operations (Prof. Modarres).

I would like to give some details on the project directed by Prof. Choi on the development of advanced control strategies for a class of industrially important polymerization processes. The research in this project ranges from developing detailed mathematical models for free radical and condensation polymerization reactors (for various reactor configurations: continuous, batch, and semibatch), to building research pilot plants for experimentation and design testing, to implementing advanced closed-loop control based on the sophisticated theories of stochastic control and nonlinear estimation. Many industrial polymerization processes are characterized by high release of reaction heat, complex polymerization kinetics, nonidealities in micromixing and macromixing, nonlinear reactor dynamics, and a lack of adequate on-line sensors to measure the progress of reaction and important polymer properties.

High productivity, precisely controlled polymer properties, improved reactor safety, and flexibility in reactor operation are the major objectives in highly profitable polymerization processes. Many industrial polymerization processes are rather custom designed, and it is difficult for reactor operators to adjust reaction conditions effectively to meet new product specifications.

The research under way at SRC is aiming at the development of efficient software tools for process analysis and control, based on enhanced understanding of intrinsic characteristics of polymerization systems. A combination of sophisticated theoretical and experimental work has been undertaken. Current interest is focused on the following polymerization systems: high-conversion suspension polymerization of methyl methacrylate with high monomer/water charge ratios, solution copolymerization of styrene and acrylonitrile, precipitation and solution polymerization of acrylamide to produce water soluble polymers, and polymerization of high-impact polystyrene. The mathematical models of these processes will be validated by real-time measurement systems and by the sophisticated simulation tools being developed for the research pilot plant. Fully digital controllers will be designed that utilize real-time measurement devices such as on-line densitometers, on-line viscometers, and on-line gel permeation chromatographs. A particularly innovative feature of the project is the use of advanced real-time estimation algorithms and sophisticated control system design techniques.

Other exciting developments in this area include the development of expert systems for chemical reactor safety analysis and safety incorporation in design (Prof. Modarres) and the development of AI/optimization-based, multi-time-scale scheduling and planning tools for chemical plant operation and design (Profs. McAvoy, Asbjornsen, Tits, and Baras).

Manufacturing Systems

In the area of manufacturing systems we had the following projects: flexible manufacturing cell (Profs. Anand, Kirk, and Nau); computer-integrated manufacturing (Profs. Harhalakis and Mark); discrete-event dynamical systems and manufacturing automation (Profs. Ho, Makowski, and Baras); manufacturing of thermoplastics (Profs. Azarm, Choi, Hammar, Mechlenburg, Pandelidis, Pecht, and Smith), and printed wiring board design and manufacturing (Profs. Pecht and Palmer).

I would like to provide brief descriptions of some well-integrated projects in this area. In a project directed by Prof. Nau, new methodologies are developed for knowledge representation and reasoning for process planning. The interdisciplinary team also includes Profs. Anand, Kirk, and Harhalakis. This research is based on a new frame-based approach to knowledge representation called "hierarchical knowledge clustering." In most frame-based reasoning systems, the data manipulated by the system are represented by using frames, but the problem-solving knowledge used to manipulate these data is represented as production rules. However, this is not always the best approach. Production rules are not always a natural way to represent knowledge—and in addition, rule-based systems containing large knowledge bases may require excessive computation to determine which rules are applicable. Hierarchical knowledge clustering provides ways to address these problems, yielding a more natural way to represent knowledge as well as improved computational efficiency.

A prototype system using hierarchical knowledge clustering was implemented in PROLOG, in a system called SIPP. An improved version is being implemented in LISP, in a system called SIPS (Semi-Intelligent Process Selector). SIPS is being adapted for use in process planning in the Automated Manufacturing Research Facility (AMRF) project at the U.S. National Bureau of Standards (NBS). Further research in this area will involve extending the approach used in SIPS to develop a practical AI-based process planning tool. Ideally, this tool will be capable of producing process plans for complex objects completely from scratch, using only the specification of the part to be produced and knowledge about the intrinsic capabilities of each manufacturing operation. The planned research requires the development of ways to integrate solid modeling techniques with AI reasoning and problem-solving techniques (e.g., ways to extract meaningful features from solid models), as well as the development of more sophisticated ways to reason about the properties of three-dimensional objects; work is also under way on a new approach to solid modeling.

The advent of computers in almost every manufacturing corporation, together with the plethora of relevant software packages aiming at increased efficiency and profitability, has produced an uncontrollable situation. Attainable benefits evaporate due to the unprecedented multiplication of input, maintenance, and output and the amount of money and manpower required to implement and coordinate all these systems.

In recognition of this problem we have initiated an integration project headed by Profs. Harhalakis and Mark that will eventually lead to minimization of data transfer and of the burden of running such a variety of "data vehicles" and data processors. A core system, Manufacturing Resources Planning (MRP II), is suggested to host computer-aided design as a first step toward integration. MRP II is by definition addressing all facets of industrial business, from marketing planning through engineering to manufacture, final inspection, and shipment. CAD is meant to assist the front end of the product life cycle and to focus on engineering, design, and drafting related activities.

The integration will be founded on a data base level. Sample features of the proposed integration include the following:

- automatic part master record generation and single-level product structures on completion of a new CAD drawing,

- engineering change control via checks performed at inventory and order levels and through status messages transmitted to MRP II and CAD screens, and
- the ability to retrieve and query pictorial and textual information on parts and assemblies at every level of the organization.

It is estimated that a large number of companies already using or planning to use MRP II and CAD will benefit substantially from such an integrated set, which ensures a smooth and effective flow of information. Future plans include the establishment of more links between MRP II and computer-aided manufacture (CAM), computer-aided testing (CAT), and others—all of them aiming at building ultimately a single computer-integrated production (CIP) system.

As an extension of this project, and in recognition of the critical significance of engineering data bases in engineering design, we have initiated a new major project on engineering information systems under the direction of Prof. Roussopoulos. As a result of SRC projects, the data base group of the Computer Science Department was increased by two new faculty members. All computer-aided design/engineering activities will be supported by Engineering Information Systems (EIS) which are based on the following technologies: data base management, AI, and distributed processing systems. The environment of an EIS is naturally distributed. Therefore, all the concurrency and consistency control of distributed data bases is present. Furthermore, an EIS has additional distribution requirements that are distributed by the presence of tools interacting with it. The basic research undertaken here is for the development of an object-oriented data base management system to support EIS. More specifically: (a) an object-oriented data model for defining engineering objects is being developed, as are (b) the data base protocols needed for concurrent access and update of multiple-version objects and (c) access methods and update protocols of distributed EIS architectures.

Communication and Signal Processing Systems

In the area of communication and signal processing systems we had the following projects: performance evaluation and design of queuing networks (Profs. Makowski, Baras, Ephremides, and Tripathi), multiuser channels with uncertain statistics (Prof. Narayan), link performance in the presence of co-user interference

(Profs. Geraniotis, Ephremides, and Narayan), mobile radio networks (Profs. Ephremides and Geraniotis), design and analysis of data compression schemes for image and speech signals (Profs. Farvardin and Shamma), speech analysis and recognition (Profs. Shamma, Peckerar, and Farvardin), mathematical methods for spectrum estimation (Prof. Benedetto).

I would like to describe now briefly some well-integrated projects. We have initiated the development of a computer and communication network laboratory, where sophisticated simulation, performance evaluation, optimization, and design tools are being developed. In a joint project with AT&T Bell Laboratories Prof. Makowski is analyzing and enhancing the Performance Analysis Workstation (PAW) performance evaluation system. The system employs direct graphical input for queuing systems representing computer systems, communication networks, and manufacturing production lines. It has the capability of computing various statistics of the network. In a broader effort Profs. Baras, Ephremides, Geraniotis, and Makowski have initiated the development of advanced simulation and optimal design tools. The key idea is to develop distributed LISP-based, object-oriented programming tools for simulation and control of complex, variable networks. We also plan to link these systems with more traditional discrete event simulators and optimization-based design of adaptive, distributed network control algorithms and protocols. This effort is a perfect example of the laboratory/design environment that we are creating at the SRC. In a related effort Profs. Davisson, Ephremides, Farvardin, and Geraniotis are developing a sophisticated hybrid (software and hardware combination) channel simulator.

Professor Farvardin has been studying intensively combined source channel coding and optimum entropy-constrained block transform coding. The first topic addresses the design and analysis of data compression schemes when the output of the data compressor is to be transmitted over a noisy channel. Here, the channel is modeled as a binary symmetric channel. An algorithm for optimal design of the quantizer and the optimal code word assignment is devised.

The second topic includes the analysis of a block transform coding system in which the average squared error is minimized subject to an overall constraint on the output entropy of the encoder. An algorithm for the optimal design of the quantizers, including the optimal entropy assignment, is developed. Research is under way to replace entropy-constrained quantization by permutation coding which has a block structure and thus does not suffer from problems associated with variable-length coding. Future work will address VLSI implementation of these schemes.

Professor Shamma, in collaboration with scientists and engineers from the National Institutes of Health (NIH) and the Naval Research Laboratory (NRL) has been investigating issues related to the front-end processing and recognition of speech phonemes using models of the mammalian auditory nervous system. The work involves three phases. (a) The analysis stage: the development of biophysical models and digital implementations of the auditory periphery to generate new, richer, and more robust representations of speech. (b) The recognition stage: neural network models mimicking the parallel distributed architecture of the central nervous system are used to perform various phoneme classifications. (c) The experimental stage: the integrated circuit fabrication of recording microelectrode arrays with CMOS (complementary metal oxide semiconductor) amplifiers and multiplexers to be used for further acquisition of experimental data necessary for the above modeling efforts.

Expert Systems and Parallel Architectures

In the area of expert systems and parallel architectures we had the following projects: integrated CAD of real-time non-Gaussian signal processors (Prof. Baras); an expert system for stochastic nonlinear control and filtering (Prof. Blankenship); VLSI systems (Profs. Ja'Ja' and Nakajima).

Professor Blankenship, in collaboration with J. P. Quadrat from INRIA (Institut National de Recherche en Informatique et Automatique, Le Chesny, France), has been developing an expert system based on symbolic manipulation programs for the analysis of stochastic control and nonlinear filtering problems. This software system brings to the practicing engineer, in directly usable form, such sophisticated techniques as Bellman's dynamic programming. Professor Nakajima has been investigating several problems in VLSI layout and silicon compilation. The ultimate goal is to develop a hierarchical layout design system for VLSI circuits. Recently, the development of efficient algorithms for the topological aspect of the circuit layout problem on a single layer was completed. Research on the multilayer layout problem has been initiated. In particular, we have developed an efficient channel-routing algorithm for three layers. Experimental results show that this algorithm produces, in most cases, a channel-routing pattern that requires a smaller number of horizontal tracks than previous algorithms. In the area of via minimization, a polynomial time algorithm has been developed for testing whether all nets of two or three terminals can be connected without using any via. Additional work is focusing on the development of a silicon compiler.

Professor Ja'Ja' has been investigating problems related to the complexity, architecture, design, and fabrication of VLSI chips with applications to signal processing problems. The automated generation of optimized circuit layouts from a high-level description is currently considered to be one of the most challenging problems in VLSI research. The ultimate goal is to relieve the user from all low-level details and to allow him or her to describe the design in a very high level language. The resulting layouts should be regular, compact, fast, and reliable. Recent research efforts have concentrated on a few less ambitious, general methods such as gate arrays, standard cells, and fixed floor plans. While these tools have been used successfully in the past few years, they all suffer from the fact that intermediate manual intervention is required in different phases of the design process and that they will generate highly nonoptimal designs even for some simple and natural tasks. We have been studying several fundamental problems that must be resolved before such optimized tools can exist. These problems include mapping logical functions into optimal layouts, placement and routing for special structured environments, and mapping structures represented by graphs into optimized layouts. Significant progress has been made for all of these problems.

A new approach for laying out logical functions has been developed using partially symmetric functions. A new software system called SYMBL (SYMmetric Boolean Layout) has been written to implement our approach. SYMBL is based on a strategy that first partitions the set of input variables into equivalence classes such that the given functions are symmetric with respect to each equivalence class. It turns out that this step can be implemented quite efficiently. The second main step is to determine a near-optimal "cover"—i.e., a set of appropriate subfunctions whose logical sum produces the function. We use a decomposition tree whose leaves are symmetric functions of the partitioned variables, and the partitions are combined as we go up the tree. Finally, the last phase consists of placement and routing routines that optimize the layout structure. The user can introduce his or her design in a high-level language which is then converted into a truth table. The truth table is handled by SYMBL, which produces the final layout without any intervention from the user.

We are also exploring the possibility of mapping functions into a general array type called Weinberger arrays. Two basic problems must be tackled in this approach. The first consists of manipulating the given functions into an optimized form. The second must place and route the Weinberger cells corresponding to the logical form obtained. This second problem can be formulated as a purely graph theoretic problem for which combinatorial tóols are very useful. We have developed a set of good heuristic algorithms that work well for almost all cases. Our next step will be to implement these tools and to try them on real-world cases.

Several architectures have been proposed in the literature for handling basic signal processing computations. These architectures are highly regular and allow a good degree of concurrent processing. However, most of the implementations have considered the standard algorithms and have mapped them into these architectures. We have introduced fully pipelined structures that are based on a novel strategy consisting of decomposing a computation into a set of subcomputations that can be executed in parallel. A problem of size n will be roughly decomposed into $\sqrt{n}$ subproblems, each of size roughly $\sqrt{n}$, such that all these subproblems can be solved in parallel on fully pipelined, bit-serial systolic architectures. The class of problems for which such decompositions exist include filtering, convolution, and computing the discrete Fourier transform (DFT). We have shown that these structures can be implemented quite efficiently with compact hardware. As a matter of fact, we have designed a 25-MHz chip for computing the 240-point DFT that can handle up to 30,000 such computations per second.

Professors Baras and Ja'Ja', in a joint project with engineers from Sperry Corporation, are studying VLSI architectures for linear and nonlinear signal processing. Professor Baras has been developing the IDELPHI expert system for integrated design of VLSI chips for nonlinear, real-time signal processing. This software system has several modules: signal model development and validation, computation of sufficient statistics, architecture selection, and chip design. It will be intelligent enough, when fully developed, to understand the level of user expertise. It brings to the practicing engineer a sophisticated array of techniques and methodologies from stochastic systems, communication engineering, numerical mathematics, and VLSI complexity and architectures in a directly usable form. Several open problems in real-time sequential estimation and detection have been resolved by an innovative combination of sophisticated numerical techniques and VLSI architectures. Professor Baras and students LaVigna and Simmons are currently completing the design of a special purpose VLSI chip, the Zakai I chip, which provides a real-time solution to the celebrated nonlinear filtering problem. A printed circuit board prototype will be finished soon, and then the fabrication of the large (about 140,000 transistors) 22-MIPS chip will be undertaken. Our research here has revealed a major weakness of currently used signal models for communication and control: they are not properly structured for real-time processing. Planned research includes the investigation of massively parallel architectures (like connection machines), neural net-type architectures, and applications in adaptive array processing and speech processing for reduced bit-rate transmission.

The Value of the Approach

I would like to close this section with a few remarks regarding the interaction and integration of the problems and disciplines represented within the SRC. Our research has revealed that systems science and engineering is a powerful unifying, interdisciplinary approach to engineering design problems. Our focus on automatic control and communication systems amplifies this point further. We found that chemical plants and automated printed circuit board factories need the same hierarchical, multi-time-scale decision aids for scheduling, planning, operations, and even design. Large, flexible space structures and highly dexterous advanced robotic arms and hands need the same laboratory/design environment and can benefit from sophisticated theories and software systems analyzing multibody dynamics and distributed sensor fusion. Computer and communication networks, on the one hand, and flexible manufacturing factories, on the other, can benefit from the system-level design tools for modeling, simulation, and performance evaluation that we are developing. Polymerization reactor control and sequential target discrimination rely heavily on our ability to design digital, real-time, nonlinear estimators. Design of unusually large VLSI chips and the integration of CAD with CAM both require the object-oriented data base management and control schemes that we are developing. As we continue our efforts at the SRC we expect this vast cross-fertilization to guide us in the creation of the design tools for tomorrow's systems.

THE EDUCATIONAL PROGRAM

The educational program of the SRC is aimed at developing undergraduate and graduate curricula with emphasis on the five focus technical areas of the Center. This program complements the research activities of the Center and reflects our commitment to developing an extensive and continuous exchange of educational information with other universities and research institutions, private industry, and government research and development (R&D) laboratories.

Of particular concern in the development of the educational programs of the SRC is the planning and timing of specific courses targeted at bringing the advancements in AI, VLSI, and CAE to many undergraduate and graduate students. This is critical in order to create the necessary "technologically literate" student core for the SRC programs. Special purpose courses on AI will be offered in the fall semester of 1986 in four separate sections (in the electrical, mechanical, chemical engineering, and computer sciences departments) at the sophomore and junior levels. In addition, a graduate course/seminar on AI tools will be given from the Applied AI Laboratory, which is currently under development.

The SRC is rapidly developing plans for a specialized program of short courses that will bring state-of-the-art research results to industrial research scientists. This program will be an SRC-wide extension of the very successful short course on chemical process control offered by the Chemical Engineering Department. These courses will be sponsored by SRC and will bring, as speakers, authorities on various subjects of interest to SRC, both from faculty affiliated with SRC and elsewhere.

In addition, we have initiated the sponsoring or cosponsoring of colloquia, workshops, satellite video conferences, and symposia to facilitate the educational function of the Center.

I would like to briefly describe the highlights of what has been achieved to date. Further details can be found elsewhere (Systems Research Center, 1985, 1986).

- We initiated a shift to AI language programming from FORTRAN. As explained earlier, we believe that AI languages offer many advantages for the description and resolution of engineering design problems.
- We introduced or modified some 20 engineering courses, with particular emphasis on VLSI, AI, and CAE technologies.
- We sponsored a variety of interdisciplinary systems colloquia, including weekly colloquia for SRC students.
- We emphasized undergraduate research projects and facilitated the necessary matching between faculty and students. Students were strongly encouraged to build real systems.
- We established procedures for coadvising and joint teaching, breaking several cross-departmental boundaries.
- We established an advanced design laboratory for college-wide classroom use. Optimization-based design and other advanced software tools will be introduced to students.
- By insisting that the students actually go through the design process, we have increased student awareness of implementation and digital design issues.
- We created a distinguished graduate and undergraduate SRC fellowship program, with explicit industrial connections. We try to provide the student with an industrial mentor, in addition to his or her academic adviser.
- We facilitated student visits and residency at industrial sites. Seven graduate and four undergraduate SRC students spent part of summer 1986 on internships with industrial or government R&D labs.

THE INDUSTRIAL COLLABORATION PROGRAM

We are grateful to many corporations and government laboratories for a plethora of technical exchanges and joint projects, as well as for support received in the form of funds for faculty and student activities, unrestricted funds, and equipment gifts. We are pleased with the developments to date and look forward to even more interactions during the second year. From the description of the research program given above it should be obvious that we are developing a long-range, deep collaboration program with industry and government. The students have played a key role here. We basically followed our plans delineated elsewhere (Baras, 1986; Systems Research Center, 1985, 1986) and put primary emphasis on the development of strong technical ties between SRC researchers and industry and government scientists and engineers. In the limited time and space available here, I would like to describe the highlights of the program to date.

- We had numerous technical contacts, ranging from full-day meetings to shorter visits. As a result, many joint projects have been initiated.
- We created an endowed SRC fellowship program that proved to be extremely attractive to industry. Sponsorship of students provides an extremely robust and strong link with industry.
- Several technical advisory committees have been established, and we formed the SRC Research Advisory and Administrative Councils.
- We have received support from industry by means of a relatively unstructured and flexible program, given the time delays necessary to establish more formal schemes. We are currently formalizing many of these relationships along the lines of our planned industrial affiliates program (see Baras, 1986), with its three-tier structure. We anticipate that by September 1986 we will have formal agreements with at least three corporations at the highest (sustaining partner) level, with at least two corporations at the middle (sponsor) level, and with at least 10 corporations at the basic (associate) level.
- We have initiated strong collaborations with NBS, NRL, NASA, and NIH on a variety of interdisciplinary research projects.
- To date, the following corporations have provided support to the SRC. *Affiliates:* Applied Technology Inc., ARCO Chemical Corporation, Control Data Corporation, Digital Equipment Corporation, E.I. Dupont de Nemours & Co. Inc., Eastman Kodak, Exxon Corporation, Foxborough Inc., Grumman Aerospace Corporation, Martin Marietta Corporation, Mobil Oil Corporation, Rexnord Inc., Sperry Corporation, Texas Instruments Inc., Westinghouse Electric Corporation. *Equipment donors:* Application Engineering Corporation, AT&T, Cincinnati Milacron, Con-

trol Data Corporation, Data General Corporation, Digital Equipment Corporation, Intel Corporation, Kinetic Systems Corporation, Lisp Machine Inc., Moldflow Australia, Silicon Graphics, Sun Microsystems, Symbolics Inc., Texas Instruments Inc., Valid Logic Systems.

A PREVIEW OF THE SECOND YEAR

Our plans for the second year call for continuation of our efforts along the same lines with emphasis on the following. In the research program, pursue more intellectual integration and the formal initiation of the software "club"; in the educational program, develop short courses with industry, increase industrial participation in teaching, initiate the lecture note series and the technical magazine, and provide new integrated design courses and projects; in financial support, diversify further with other federal, state, and industrial sources; in facilities, move into a new 26,000-sq.-ft. building and continue development of the SRC constituent laboratories—VLSI Systems Design Lab, VLSI Systems Testing Lab, Process Systems Lab, Computer Integrated Manufacturing Lab, Intelligent Servomechanism Lab, Applied AI Lab, Signal Processing Lab, Computer and Communication Network Lab, Computer-Aided Design Lab, and Advanced Robotics Lab.

(For further information regarding results presented here, including technical papers and reports, software, etc., please contact: Mr. Timothy McGraw, Assistant to the Director for Information Dissemination, Systems Research Center, University of Maryland, College Park, Building 093, College Park, Maryland 20742; telephone: (301) 454-6167.)

REFERENCES

Baras, J. S. 1986. Systems Research Center. Pp. 61-74 in *The New Engineering Research Centers: Purposes, Goals, and Expectations*. Washington, D.C.: National Academy Press.

Systems Research Center. 1985. Report for the period May 1, 1985, to January 31, 1986, submitted to the National Science Foundation November 15, 1985.

Systems Research Center. 1986. First Annual Report, August.

ACKNOWLEDGMENTS

I would like to thank all SRC faculty, students, and staff whose accomplishments I described here. In particular I would like to thank my colleagues R. Brockett, K. Y. Choi, N. Farvardin, G. Harhalakis, J. Ja'Ja', P. S. Krishnaprasad, M. Pecht, S. Shamma, and A. Tits for their help in the preparation of this paper. Special thanks are also due to SRC staff members Pamela Harris and Pari Noskin for their tireless and efficient efforts in the preparation of this paper and the symposium presentation.

DISCUSSION

One listener expressed concern that foreign-born students educated in the Center might return to their home countries, taking ERC technology and ideas with them. Dr. Baras expressed his opinion that this would not be a problem as long as the United States remained competitive and provided a challenging and attractive working environment for top engineers. He said that the program of international collaboration being run by his Center with scientists and engineers from other countries was a productive, two-way cooperation. To another question regarding the range of apparently disparate applications of Center projects, Dr. Baras emphasized that the objective is to develop a common set of system-level design tools that can be applied to virtually any application. "We are trying," he said, "to investigate systems science and engineering itself, not the particular manifestation peculiar to a given application."