

SYSTEMS RESEARCH CENTER:
THE FIRST YEAR

College of Engineering
University of Maryland
College Park

Division of Applied Sciences
Harvard University
Cambridge

John S. Baras
University of Maryland

Systems Research Center: The First Year

JOHN S. BARAS

INTRODUCTION

The Systems Research Center (SRC) at the University of Maryland and Harvard University is a new and unique forum of fundamental systems engineering research and education.

Established by a National Science Foundation grant in 1985, SRC is developing innovative advances in design methods and software systems which address the basic productivity and competitive challenges 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.

The impetus behind SRC's activities is a close and mutually supportive collaboration with industry and government. Its programs are designed with a twofold purpose: 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 to meet objectives consistent with the competitive needs of business and industry.

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 inter-related 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 tech 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 necessary in order to integrate the sophisticated analytical and computational techniques of control and communication engineering with computer hardware and software advances. 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 where advanced computer technology is used as an "amplifier" of human engineering skills and ingenuity.

Educationally, SRC's goals are to support and enhance educational programs and to serve as a source of new courses and material. In doing so, the Center is seeking to change the traditional focus of engineering education, placing a new emphasis on both *education* and *training*. In cooperation with Harvard University's Division of Applied Sciences, the Center's broad, interdisciplinary programs cut across the boundaries of a great many engineering and computer science disciplines, and 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 it to a diverse set of complex, real-world problems; and to speed the transfer of research results into the industrial community.

To enhance the interaction between 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, software library and a unique software research "club".

In this paper I will describe the progress 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.

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

- Integrated design of complex automation and information processing systems. To create a computer-aided design environment where a 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 was neglected over the last fifteen years or so. In particular we aimed at emphasizing modeling of the new hardware systems used for implementing the control or information processing algorithms.
- Develop design theories and tools that incorporate advances in computer technology (hardware and software). This 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** 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.

It has been a very exciting year at SRC. Building the organization, infrastructure and mechanisms to support the various programs has been a demanding but intellectually very rewarding endeavor for all involved. Let me take a few moments to provide an overview of our *accomplishments during this beginning*. We have initiated a broad innovative, interdisciplinary research program, for which I will say more in detail at a later section. We undertook several educational innovations and developed advanced design teaching laboratories. We established a broad industrial collaboration program. We begun the development of several advanced research laboratories. We initiated a broad industrial collaboration program. We completed the organizational and administrative structure of the Center.

We paid particular attention to develop 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. In particular SRC involves two universities, 6 departments, 2 colleges, 36 faculty, 50 graduate students and 30 undergraduate students. It is apparent that without a strong "joint venture" spirit supported by such policies the long range success may not be possible. The idea of the Administrative Council (which includes the chairpersons of the major participating departments, the deans, representatives from executive officers from industry and the SRC director) proved to be a very efficient vehicle to formulate and implement such policies.

We have made great strides already in diversifying the financial support for the Center. The funding components now in place include funding from other federal and state agencies, the university and industrial sponsors. We are grateful to industry for the spontaneous response both in funds and equipment gifts displayed during our first year of operation. For the first year we received funding of about \$2M from NSF, \$.4M in funds and in excess of \$2M in equipment gifts from industry, \$.3M from the State of Maryland and the 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 this beginning year. From academic faculty it has been enthusiastic; many more people want to participate in the SRC programs 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 with 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 and we had five times as many applicants as we could accommodate. In particular the possibilities for "real" industrial interaction proved to be an extremely attractive reward for students.

THE RESEARCH PROGRAM

As described in further detail in [1-3] 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, 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 in the involvement of undergraduate students in design projects as early as possible in their careers. As emphasized elsewhere [1], in some sense, the subject of fundamental studies carried at SRC is systems science and engineering itself and not its particular application or manifestation in a narrow application area. This in itself justifies the emphasis 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 rewarded us with the discovery of some unsuspected connections between totally unrelated engineering problems. Several key ideas and concepts have already emerged: a pervasive utility of optimization based design in a great variety of engineering systems; the value of AI in the development of systems that can reason and aid in the design process; the superior efficiency and ability of symbolic algebra as an engineer's aid in complex calculations and modeling; the critical importance of representation and

manipulation of engineering databases; the mandatory utilization of AI techniques to handle heuristics in the design process and to enhance and facilitate team work; the mandatory utilization of interactive graphics as the interface mechanism between the systems engineer and the computer; the critical importance of understanding VLSI systems, their architectures and 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 I will highlight certain sample projects to illustrate the nature of our work and our findings. The selection of the projects discussed in some detail, reflects, to a certain degree, my finite memory and capacity to represent the details of many excellent projects underway, but also reflects my perception of projects where intellectual integration across disciplines has occurred at the greatest degree.

First the focus areas and the the thrusts within each area are:

- (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 integration of reliability and safety in the computer aided design process.
- (3) **Manufacturing systems.** Major thrusts are the integration of CAD with 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, reliability integration in computer aided design.

From the beginning we have placed emphasis in designing and developing our research laboratories so as to create a sophisticated environment for integrated design. The *constituent laboratories* of SRC is a key concept in our development of SRC. These sophisticated laboratories form the natural home for interdisciplinary groups of faculty and student researchers. In a sense the whole 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 similar and other workstations and computers, so that the engineer can really have the capability to run concurrently several modules of his design software system on different machines. The AI workstation provides the direct interface with the user, the (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, MACSYMA offer a superior medium for design problem definition, conceptualization and implementation. 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 scientist 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 labs 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 in the description of selected research projects below.

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* (Prof. Berenstein and Baras), *hand-eye machine* (Profs. Brockett, Maragos and Wohn), *nonlinear control and robotic manipulators* (Profs. Krishnaprasad, Berenstein, Tits and Abed), *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* 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 multi-body systems. Progress has been achieved in several directions. Professor Krishnaprasad in collaborative work with Professor J. E. Marsden (University of California, Berkeley) and Professor Juan Simo (Stanford University) worked out the Hamiltonian structure of systems of rigid bodies with flexible attachments that obey *finite strain* elastic models and are geometrically exact. These results should have significant impact on the control theory of multibody spacecraft and in the modeling and control of flexible robotic manipulators. They have also been able to establish the Poisson structures underlying a wide variety of problems in interconnected rigid body systems and use these in obtaining stability criteria for interesting 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 MRMS).

On the software side, DYNAMAN has been further developed to include symbolic jacobian calculations. DYNAMAN is a symbolic algebra system, developed at SRC and written in MACSYMA, which automatically generates the dynamical equations of multi-body systems in symbolic form. Further refinements will be implemented and tested using the SRC Lisp Machines.

On the experimental side, Mr. Haven Frank, an M.S. student, has been actively involved in a project to implement a controller for a single link flexible arm. The arm was built at the Westinghouse model shop to our specifications. Preliminary controller designs using DELIGHT.MaryLin are being translated for implementation and testing on an IBM PC/AT with an on-board data acquisition and control system. In this experiment we are learning how to use accelerometers for real time closed loop control with minimal overshoot of the link tip based on hub position and rate sensing. We are in the process of beginning new experiments on the use of piezoelectric materials for sensing and actuation. These investigations will expand towards the study of accelerometer arrays and the development of *neural net* type algorithms for sensor data processing and control.

DELIGHT.MaryLin is an interactive system for optimization-based design of linear control systems. It was developed (and is still under further development) at the University of Maryland, College Park under the direction of Prof. Tits. It is an offshoot of the DELIGHT general purpose interactive system, originally developed at the University of California, Berkeley. It makes extensive use of DELIGHT's optimization facilities, developed jointly at Berkeley and College Park. It also makes use of linear systems analysis routines developed by various groups.

With the DELIGHT.MaryLin tool, a control system designer can take advantage of recent powerful optimization algorithms from the DELIGHT algorithmic library to automatically adjust design parameters of the control system being designed. One of the most striking features of DELIGHT.MaryLin is its extreme flexibility. Using a recently proposed (by Tits and Nye) application-oriented design methodology, the designer may optimize arbitrary performance criteria as well as interactively explore the tradeoffs between multiple competing objectives, while simultaneously keeping several constraint specifications met. Success has been achieved on design problems pertaining to various application areas, including the design of a flight controller for high performance aircraft (in a joint project between SRC faculty Levine and Tits and Grumman Aerospace engineers) and the design of an airborne electro-optical device.

A group of ten students was involved in a design project under the sponsorship of NASA Headquarters, NASA Goddard Space Flight Center and the Systems Research Center. The project is under the supervision of Professor P. S. Krishnaprasad and is one of eighteen projects under way in universities across the nation as a 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 (MRMS) for the Space Station*. Students participating in this project receive academic credit.

Based on guidelines obtained from the Space Station Reference Configuration Description and other published material on operational requirements, the Maryland team has developed a preliminary design for the MRMS. The proposed MRMS is planned to have a reach of approximately 33 meters and is intended for a variety of tasks ranging from assembly of the Space Station itself to transferring payloads from the Space Shuttle to the Station, assisting in the servicing of disabled satellites, and even the docking of the Shuttle! The kinematic design specifies an eight degrees of freedom arm (including a three degrees of freedom wrist) which travels along the dual keels of the space station. The preliminary sizing of the actuators and the links has been completed. Using a scheme of *enslaving* certain degrees of freedom to

resolve the redundancy, kinematic programs for obstacle avoidance are being written and debugged. The control laws for the joint servos are designed from rational engineering specifications using software (DYNAMAN and DELIGHT.MaryLin) developed at the SRC.

Any design activity of this nature should have a test-bed for trying out various ideas (e.g. control schemes, kinematic programs). To this end, the SRC team has developed a 3-D graphical simulation of the MRMS on a Silicon Graphics Inc. workstation. Soon this workstation will also display a reference configuration Space Station with the MRMS attached to it. The requisite database for this purpose has been supplied by NASA Goddard Space Flight Center. We expect to continue the project in the next academic year possibly including some scale model fabrication of the MRMS. Participation in the MRMS design project has provided the students an excellent opportunity to become aware of the interdisciplinary nature of design activity in general. They have also had an opportunity to explore an exciting world of high technology represented by the NASA Mission in Space Station Automation and Robotics.

Throughout the year strong collaboration with Harvard continued in the area of robotic manipulator design. In the project directed by Prof. Brockett the 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 Automatrix 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.

The Harvard team is also designing a *high resolution magnetic tactile sensor* which is currently envisioned as an array of 64 by 64 magnetic sensors integrated in a CMOS integrated circuit, on a chip whose size is roughly .8 by .9 cm. The individual sensors have dimensions of roughly 90 microns on a side. Each sensing element is a *split drain* MOSFET that is biased so that current flows through the transistor. In the absence of a magnetic field the currents flowing through the two drains of the MOSFET are equal. However, when there is a magnetic field, with a non zero component perpendicular to the plane of the sensor, the electrons flowing in the transistor are deflected, due to the Lorentz force. This deflection results in an imbalance in the current flowing through the two drains of the MOSFET. This current differential is converted to a voltage and amplified by a current mirror. The combination of the split drain MOSFET and the current mirror is equivalent to the input stage of an operational amplifier. The sensor chip contains circuitry for enabling (allowing current to flow through) a given row of sensors and for multiplexing the output voltage by columns, so that the measurements of the individual sensors can be serially output in a raster scan fashion.

The magnetic sensor array integrated circuit will be mounted in the compliant finger developed at the Harvard robotics laboratory. This finger consists of a fluid

filled membrane fixed to a rigid substrate. The inside of the membrane will be coated with a magnetized material, and the sensor array will be mounted on the substrate. When the membrane is deformed, the magnetic field pattern at the sensor plane will be altered. In principle, knowing the field pattern at the sensor plane will allow one to estimate the deformation of the membrane, which will then allow us to obtain some information as to the spatial pattern of the forces applied to the membrane. The current status of the tactile sensor is as follows. The 64 by 64 magnetic sensor array has been designed and laid out using the CAESAR layout utility running on the Data General DS4200 workstations. Critical parts of the circuit have been simulated using SPICE, but the nature of the physics involved in the magnetic sensing process are such that an accurate estimation of the sensitivity of the sensors to changes in the magnetic field can not be obtained. This will have to wait for experiments on the actual circuit.

A distributed tactile sensor is also 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 will be performed by *neural type* processor chips.

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 Professor 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 semi-batch), building research pilot-plants for experimentation and design testing and 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, non-idealities 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 quite difficult for reactor operators to adjust reaction conditions effectively to meet new product specifications.

The research underway 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, polymerization of high impact polystyrene. The mathematical models of these processes will be validated by real time measurement systems and the sophisticated simulation tools being developed for the research pilot plant. Fully digital controllers will be designed utilizing real time measurements such as on-line densitometer, on-line viscometer, and on-line gel permeation chromatograph. 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), 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).

In the area of **manufacturing systems** we had the 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, Smith), *printed wiring board design and manufacturing* (Profs. Pecht, Palmer).

I would like to provide brief descriptions of several well integrated projects in this area. In a project directed by Professor Nau new methodologies are developed for *knowledge representation and reasoning for process planning*. The interdisciplinary team also includes Professors 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 is represented using frames, but the problem solving knowledge used to manipulate this 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 in order to determine which rules are applicable. Hierarchical knowledge clustering provides ways to address these problems, yielding both a more natural way to represent knowledge and 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. 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), and development of more sophisticated ways to reason about the properties of three-dimensional objects, and work is underway 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 have produced an uncontrollable situation. Attainable benefits evaporate due to the unprecedented multiplication of input/maintenance and output and the amount of money and manpower expended to implement and keep all these systems under some sort of coordination.

In recognition of this problem we have started up an integration project headed by Professors Harhalakis and Mark, which will eventually lead to minimization of data transfer and of the burden to run such a variety of "data-vehicles" and data processors. A core system, Manufacturing Resources Planning (MRP II), is suggested to host Computer Aided Design (CAD) as a first step towards 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:

- 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
- Ability to retrieve and enquire upon pictorial and textual information on parts and assemblies at every level of the organization

The basic problem in homogeneous database integration is the required initial design of the global schema which is the union-without redundancy and internal conflicts - of the schemata of the databases to be integrated. We propose the interoperability approach to database integration based on the philosophy that conflicts and redundancy must be accepted, but controlled. We define the global schema as the concatenation of the existing schemata and we add a set of updated dependencies to control the redundancy and conflicts between them. Update dependencies monitor operations changing a database and propagate the changes to other databases by remote operation calls

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 one-single Computer Integrated Production (CIP) system.

As an extension of this project, and in recognition of the critical significance of engineering databases in engineering design we have initiated a new major project on *engineering information systems*, under the direction of Professor Roussopoulos. In response to SRC projects the database 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 which are based on these technologies: database management, AI and distributed processing systems. The environment of an EIS is naturally distributed. Therefore, all the concurrency and consistency control of distributed databases 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 Database Management System* to support EIS. More specifically : (a) an Object Oriented Data Model is being developed for defining engineering objects, (b) the database protocols needed for concurrent access and update of multiple version objects, and (c) access methods and update protocols of distributed EIS architectures will be developed.

A well integrated, interdisciplinary effort is being undertaken under the direction of Prof. Pecht for the development of a *Reliability and Maintainability Computer Aided Design (RAMCAD) workstation*. Litton Amecom industries will be an integral part of the project with funding pending from the AF. The construction of the workstation will be performed by Litton. This research effort aims at the integration of reliability and maintainability of electronic circuits in the computer aided design process.

This project involves the development of a second level electronic packaging design process consisting of tools and procedures with reliability design capabilities. The goal is to reduce the lag time of the design process. Presently, schedule constraints often compromise the effectiveness and efficiency of the design process. For example, reliability analysis is generally conducted late in the design process, often when it is too late to make significant (or the most appropriate) design changes. The goal of this effort is to make R&M an integral, "real-time" and upfront part of the design process, such that designs do not have to be re-submitted. In other words, we don't want to have to re-design a completed design because reliability engineering represents the last phase of analysis in a project.

The concept behind this effort is an innovative sharing of resources within a RAMCAD workstation. This is accomplished by establishing two modes of operation, a foreground mode and a background mode, which operate in parallel and which share data. In foreground mode, the design engineer uses data capture tools and design rules to firm-up (but not necessarily lock-in) a design. At the same time, the foreground acts as an advisor which looks over the shoulder of the engineer and may provide design recommendations. The foreground process runs in an online, highly interactive mode. It will present analysis and design violation results in graphical and/or tabular form from the background mode. The background mode operates in parallel with the foreground mode, taking design information from the foreground as input. The background mode consists of the analysis routines and the tools required for parameter selection and prediction, trade-off studies. The analysis routines will include a mix of standard software packages and software packages and software designed specifically for the RAMCAD electronics design workstation. For example, we will incorporate a PCB routing routine and a micro-circuit fault grading program into RAMCAD. While such software could be executed independently, inputs are typically provided from the data capture, and outputs should be used as inputs for analysis and reliability prediction.

In all stages of the design process, the consequences of foreground design decisions will be "examined" in the background mode to indicate violations, check guidelines, etc. The specific background tasks should include but are not limited to:

- (1) geometric constraints and tolerances
- (2) conformance to derating
- (3) part application guidelines including:
 - (a) reverse voltages in tantalum capacitors
 - (b) latch-up conditions (signal before power) in CMOS

subcomponent. Finally, the transpose subcomponent consists of an array of shift registers that execute a pipelined transpose operation satisfying the same I/O requirements as the previous subcomponents. The clock period that can be used for reliable operation of each cell is 40 ns, yielding a clock frequency of 25 MHz. At this frequency with our fully pipelined approach, we can compute 30,000 DFTs per second with 16 bits of input/output precision.

Professors Baras and Ja'Ja' in a joint project with engineers from Sperry are studying VLSI architectures for linear and nonlinear signal processing. Professor Baras has been developing the *DELPHI* 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, VLSI complexity and architectures, in a directly useable 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 with 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 MIP 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.

I would like to close this section with a few remarks regarding the interaction and integration of the problems and disciplines represented at 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 printed circuit board automated 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 multi-body dynamics and distributed sensor fusion. Computer and communication networks and flexible manufacturing factories 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 very large VLSI chips and the integration of CAD with CAM need 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-

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. In any practical situation, where the source and channel encoders are of finite complexity and hence suboptimal, the channel encoder affects the performance of the source encoder. Therefore, it should be beneficial to design the source and the channel encoders jointly. This is in contrast with the traditional arguments in information theory which suggest a separation of the source and channel encoders. We have considered the analysis and design of minimum mean-squared-error zero-memory quantizers for operation on memoryless sources. Here, the channel is modeled as a binary symmetric channel. An algorithm for optimal design of the quantizer and the optimal codeword assignment is devised. It is shown that substantial performance improvements can be obtained over conventional schemes which design the source and channel coders separately. For example for a memoryless Laplacian source and a channel with crossover probability 0.01 at transmission rate 3 bits/sample, the signal-to-noise ratio of the optimal system is 17.23 dB, while that of the suboptimal scheme is only 7.44 dB - a 10 dB performance improvement! Analysis and design of an optimal block transform image coding scheme for transmission over noisy channels has been completed. The system is based on the two-dimensional discrete cosine transform. The quantizer structure, the codeword assignment and the bit allocation are optimized. Noticeable objective and subjective performance improvements are obtained through this optimized combined source-channel coding scheme. Work is under way to generalize these results to adaptive block transform schemes for image and speech coding over noisy channels. Also, more sophisticated channel models such as channels with fading and intersymbol interference are being considered.

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. For Gauss-Markov and Laplace-Markov sources the performance of the optimal system at different bit rates and for different block sizes is obtained. It is shown that substantial improvements can be gained over conventional block transform coding with Lloyd-Max quantization. Extensions to image coding problems have been considered. It is shown that especially at low bit rates the subjective performance improvements are very pronounced. 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 NIH and 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 IC fabrication of recording microelectrode arrays with CMOS amplifiers and multiplexers to be used for further acquisition of experimental data necessary for the above modeling efforts.

In the first stage, we generate spatio-temporal representations of speech sounds that preserve the fine temporal information at the outputs of the cochlear filters. Various signal parameters (e.g. harmonic spectra, formants, and pitch) are encoded in these two-dimensional patterns as robust morphological features (e.g. edges) that are insensitive to sound intensity and channel nonlinearities (i.e. saturation due to limited dynamic range). Edge extraction algorithms are implemented as models of non-recursive lateral inhibitory neural networks (LIN) to derive the spectral parameters. In the second stage, massively interconnected inhibitory neural networks are simulated to perform the pattern recognition (phoneme classification tasks). The networks are designed first to undergo a learning phase in which "templates" are established in the connection profiles through repeated application of the processed signal patterns. Once the appropriate connections are made, classification of unknown inputs is performed by iterative mappings of the network to specific fixed points which correspond to predefined input classes. These algorithms are currently applied successfully to various vowel and other steady or slow varying sounds, and we are attempting to extend these concepts to the recognition of dynamic phonemes (e.g. stop consonants). Finally, neurophysiological data to support the above modeling efforts will be obtained from guinea pigs via Gold-on-Silicon microelectrode arrays with on-board microelectronic circuitry to amplify and multiplex the signals from several electrodes. Extensive software is being developed to facilitate the automatic detection and processing of the recorded neural signals.

In the area of **expert systems and parallel architectures** we had the projects: *integrated CAD of real-time nongaussian 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 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 useable 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 multi-layer 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 which 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 to describe his 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 plan. 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 non-optimal designs even for some simple and natural tasks. We have been studying several fundamental problems that have to be resolved before such optimized tools 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 in all of these problems, some of which will be outlined next. 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 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. There are two basic problems that have to be tackled in this approach. The first consists of manipulating the given functions into an optimized form. The second has to 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 tools 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 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 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 a fully pipelined bit-serial systolic architectures. The class of problems for which such decompositions exist include filtering, convolution and computing the Discrete Fourier Transform. 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. We now give a brief overview of this design.

The structure has two major components, one for computing the 15-point DFT and the second for the 16-point DFT. Each of these components consist of three sub-components: the summation the scaling, and the transpose. The summation subcomponent consists of an array of bit serial adders, subtractors or delay elements. The input is processed in a fully pipelined bit skewed fashion. The scaling subcomponent consists of an array of bit serial multipliers (2's complement) to perform the required multiplications. The input/output requirements are identical to the summation

subcomponent. Finally, the transpose subcomponent consists of an array of shift registers that execute a pipelined transpose operation satisfying the same I/O requirements as the previous subcomponents. The clock period that can be used for reliable operation of each cell is 40 ns, yielding a clock frequency of 25 MHz. At this frequency with our fully pipelined approach, we can compute 30,000 DFTs per second with 16 bits of input/output precision.

Professors Baras and Ja'Ja' in a joint project with engineers from Sperry are studying VLSI architectures for linear and nonlinear signal processing. Professor Baras has been developing the *DELPHI* 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, VLSI complexity and architectures, in a directly useable 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 with 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 MIP 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.

I would like to close this section with a few remarks regarding the interaction and integration of the problems and disciplines represented at 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 printed circuit board automated 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 multi-body dynamics and distributed sensor fusion. Computer and communication networks and flexible manufacturing factories 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 very large VLSI chips and the integration of CAD with CAM need 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 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 artificial intelligence 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 currently under development *Applied AI Laboratory*.

The SRC is rapidly developing plans for a specialized program of short courses which 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, in order to facilitate the educational function of the center.

In what follows, I would like to briefly describe the highlights of what has been achieved todate. Further details can be found in [2-3].

- 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 twenty engineering courses with particular emphasis on VLSI, AI and CAE technologies.
- We sponsored a variety of interdisciplinary systems colloquia, including weekly ones 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 co-advising 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 on 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 academic advisor.
- We facilitated student visits and residency at industrial sites. Seven graduate and four undergraduate SRC students will be spending part of this summer on internships with industrial or government R&D labs.

THE INDUSTRIAL COLLABORATION PROGRAM

We are greatfull 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 todate 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 in [1-3] and put primarily 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 todate.

- 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 initiated the development of the SRC Research Advisory and Administrative Councils.
- We have received support from industry in 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 [1-3], with its three tier structure. We anticipate that by September 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 ten corporations at the basic *associate* level.
- We have initiated strong collaborations with NBS, NRL, NASA, NIH on a variety of interdisciplinary research projects.
- The following corporations have provided support to the SRC todate:
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, Cincinatti Milacron, Control 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: On the research program, more intellectual integration and the formal initiation of the software "club"; On the educational program, short courses with industry, increased industrial participation in teaching, initiate the lecture note series and the technical magazine, provide new "integrated" design courses and projects; On financial support, diversify further from other federal, state and industrial sources; On facilities, move into new 26,000 net 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, Advanced Robotics Lab.

REFERENCES

- [1] J.S. Baras, *Systems Research Center*, in "The New Engineering Research Centers: Purposes, Goals, and Expectations", National Academy Press, 1986, pp.61-74.
- [2] Systems Research Center, Report for the period May 1, 1985 to January 31, 1986, submitted to the NSF, November 15, 1985.
- [3] Systems Research Center, First Annual Report, August 1986.

Acknowledgements: 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 presentation.

For further information regarding results presented here, including technical papers and reports, software etc. please contact: *Ms. Pari Noskin, Assistant to the Director for Information Dissemination, Systems Research Center, University of Maryland, College Park, Building 099, College Park, Maryland 20742, Tel: (301) 454-6167*

