SUMMARY OF AWARDS

FISCAL YEAR 1996

Microelectronic Information Processing Systems

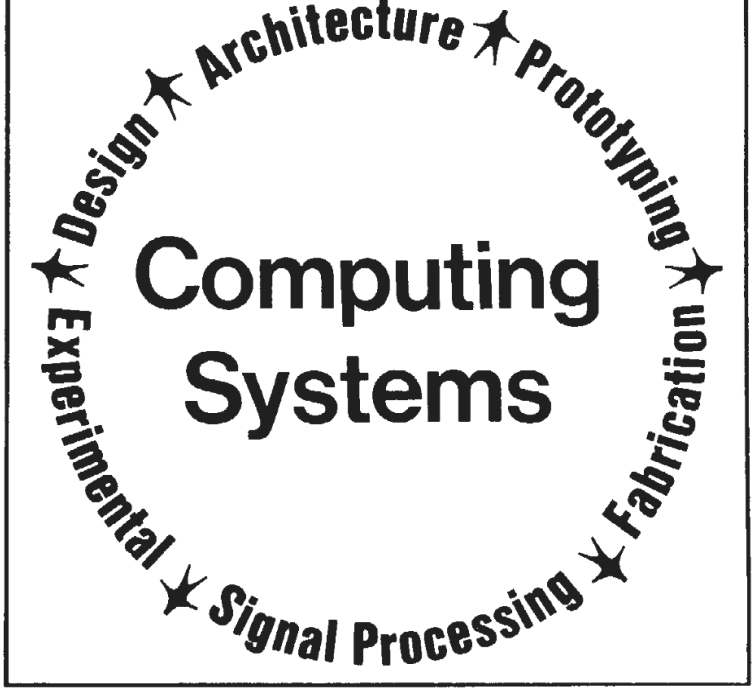

Division of Microelectronic Information Processing Systems
Directorate for Computer and Information Science and Engineering

NSF 97-144

NATIONAL SCIENCE FOUNDATION

The Foundation provides awards for research in the sciences and engineering. The awardee is wholly responsible for the conduct of such research and preparation of the results for publication. The Foundation, therefore, does not assume responsibility for such findings or their interpretation.

The Foundation welcomes proposals on behalf of all qualified scientists and engineers, and strongly encourages women, minorities, and persons with disabilities to compete fully in any of the research-related programs described here. In accordance with federal statutes, regulations and NSF policies, no person on grounds of race, color, age, sex, national origin, or disability shall be excluded from participation in, denied the benefits of, or be subject to discrimination under any program or activity receiving financial assistance from the National Science Foundation.

Facilitation Awards for Scientists and Engineers with Disabilities provide funding for special assistance or equipment to enable persons with disabilities (investigators and other staff, including student research assistants) to work on an NSF project. See the program announcement, or contact the program coordinator at (703) 306-1636.

The National Science Foundation has TDD (Telephonic Device for the Deaf) capability, which enables individuals with hearing impairment to communicate with the Foundation about NSF programs, employment, or general information. To access NSF TDD dial (703) 306-0090; for FIRS, 1-800-877-8339.

Catalog of Federal Domestic Assistance Number 47-070; Computer and Information Science and Engineering.

Preface

The Computer and Information Science and Engineering (CISE) Directorate, under the direction of an Assistant Director, consists of the following six divisions and offices: Advanced Scientific Computing (ACS) Division, Computer and Computation Research (CCR) Division, Cross-Disciplinary Activities (CDA) Office, Information, Robotics and Intelligent Systems (IRIS) Division, Microelectronic Information Processing Systems (MIPS) Division, and the Networking and Communications Research and Infrastructure (NCRI) Division.

The **Microelectronic Information Processing Systems Division (MIPS)** supports research on novel computing and information processing systems including signal processing. Emphasis is on experimental research, technology-related research and particularly the critical link between conceptualization and realization for integrated systems. Technologies include VLSI, ULSI, OPTICAL, OPTO-ELECTRONIC, INTER-CONNECTION and other emerging technologies. The focus is on research pertaining to hardware systems and their supporting software, including: experimental research involving these new systems; infrastructures, environments, tools, methodologies and services for rapid systems prototyping; design methodologies and tools; technology-driven and application-driven systems architectures; and fabrication and testing of systems. For signal-processing systems, research on algorithms and architectures relating to these new technologies that have promise for real-time computing is emphasized.

The purpose of this Summary of Awards for the MIPS Division is to provide the scientific and engineering communities with a summary of those grants awarded in Fiscal Year 1996. This report lists only those projects funded using Fiscal Year 1996 dollars and hence does not list multi-year awards initiated prior to Fiscal Year 1996.

Similar areas of research are grouped together for reader convenience. The reader is cautioned, however, not to assume that these categories represent the totality of interests of each program, or the total scope of each grant. Projects may bridge several programs or deal with topics not explicitly mentioned herein. Thus, these categories have been assigned administratively and for the purpose of this report only.

In this document, grantee institutions and principal investigators are identified first. Award identification numbers, award amounts, and award durations are enumerated after the individual project titles. Within each category, the awards are listed alphabetically by state and institution.

Readers wishing further information on any particular project described in this report are advised to contact the principal investigators directly.

Bernard Chern
Division Director

John R. Lehmann
Acting Division Director

Microelectronic Information Processing Systems Division

Table of Contents

Microelectronic Information Processing Systems

The MIPS Division

The area of Computing Systems, which involves the structure of computers, is central to MIPS. This is a core area of computer science and engineering and in the 1990's encompasses much more than just hardware. Computing systems deals with computer architecture, hardware implementation, system software (operating systems and compilers), networking, and data storage systems. The advent of gigabit networks, high performance microprocessors and parallel systems is dramatically impacting research on systems level architecture of high performance computing systems.

The emphasis in MIPS is on REAL SYSTEMS i.e. physically realizable. Special weight is placed on design, prototyping, evaluation, and novel use of computing systems and on the tools needed to design and build them. This involves technology driven and application related research, experimental research and theoretical studies. The MIPS programs support research on: high level design (design automation and CAD tools); systems level architecture studies; experimental systems research projects which build and evaluate HARDWARE/SOFTWARE SYSTEMS; signal processing algorithms and systems; knowledge of applications; methodologies, tools and packaging technologies for rapid prototyping at the system level; and infrastructure needed to support MIPS' educational and research activities, e.g. MOSIS.

The Programs

Design Automation Program

Supports research aimed at obtaining fundamental knowledge about the complete design cycle for integrated circuits and systems from conception through manufacturing and operational test. Emphasis is on integrating all aspects of the cycle, and automating the design and testing processes.

Prototyping Tools and Methodology Program

Supports research on technologies, tools, and methodologies needed for the prototyping of experimental information processing systems and for Microelectronics Education. Issues that arise in rapid system prototyping are explored, including use of new packaging techniques such as multichip modules, and such systems issues as interfacing and standards. Support is also provided for new prototyping services. Basic research necessary to model, simulate, measure, automate and improve the microfabrication process is supported. Microelectronics Education support includes workshops, conferences, development of curriculum and courseware materials, and educational support services such as those for FPGA's and fabrication (MOSIS).

Computing Systems Research Program

Supports basic research on computing systems and methods for their design. Computing Systems deals with computer architecture, hardware implementation, systems software, networking, and data storage. Research is encouraged on the fundamental aspects of computing systems architectures and scientific design methods that better utilize existing or emerging technologies, support systems software or address important applications whose computational requirements cannot be met by conventional architectures. The program emphasizes physically realizable systems and, when necessary, limited proof-of-concept prototyping.

Signal Processing Systems Program

Supports research on circuit theory and analog and digital signal processing. The emphasis is on modern signal processing, stressing the impact of VLSI, including areas such as: signal representation, filtering, novel algorithms, special-purpose hardware, and real-time computing.

Experimental Systems Program

Supports research experiments that involve building and evaluating information processing and computing systems. These are goal-oriented projects usually undertaken by teams of designers, builders, and users. The building of a system must itself represent a major intellectual effort, and offer advances in our understanding of information systems architecture by addressing significant and timely research questions. The system prototypes being built should be suitable for exploring applications and performance issues.

Basic research in Computer Architecture and Computing Systems is supported by the National Science Foundation primarily through three programs in the CISE Directorate: the Computing Systems Research Program and the Experimental Systems Program in the MIPS Division; and the Computer Systems Program in the CCR Division. The following pie chart shows the relative support through these three programs.

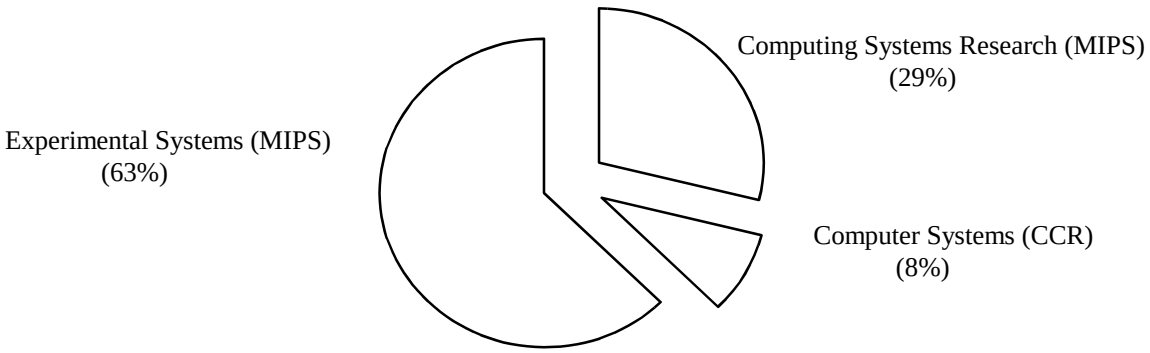

Division of Microelectronics Information Processing Systems

MIPS Staff (October 1996)

Bernard Chern	**Division Director**	**bchern@nsf.gov**
John R. Lehmann	**Deputy Division Director**	**jlehmann@nsf.gov**

Program	**Program Director**	**NET Address**
Design Automation	Robert B. Grafton	rgrafton@nsf.gov
Computing Systems Research	Michael J. Foster (Acting)	mfoster@nsf.gov
Signal Processing Systems	John H. Cozzens	jcozzens@nsf.gov
Experimental Systems	Michael J. Foster	mfoster@nsf.gov
Prototyping Tools and Methodology	Joseph Cavallaro	jcavalla@nsf.gov

MIPS Staff (July 1997)

John R. Lehmann	Acting Division Director	jlehmann@nsf.gov
John R. Lehmann	Deputy Division Director	jlehmann@nsf.gov

Program	**Program Director**	**NET Address**
Design Automation	Robert B. Grafton	rgrafton@nsf.gov
Computing Systems Research	Michael J. Foster (Acting)	mfoster@nsf.gov
Signal Processing Systems	John H. Cozzens	jcozzens@nsf.gov
Experimental Systems	Michael J. Foster	mfoster@nsf.gov
Prototyping Tools and Methodology	Michael J. Foster (Acting)	mfoster@nsf.gov

All persons can be reached on the **W**orld-**W**ide-**W**eb through: www.nsf.gov

The address and telephone number for all of the above:

National Science Foundation
4201 Wilson Blvd.
Arlington, Virginia 22230

(703) 306-1936

MICROELECTRONIC INFORMATION PROCESSING SYSTEMS DIVISION (FY 1996)

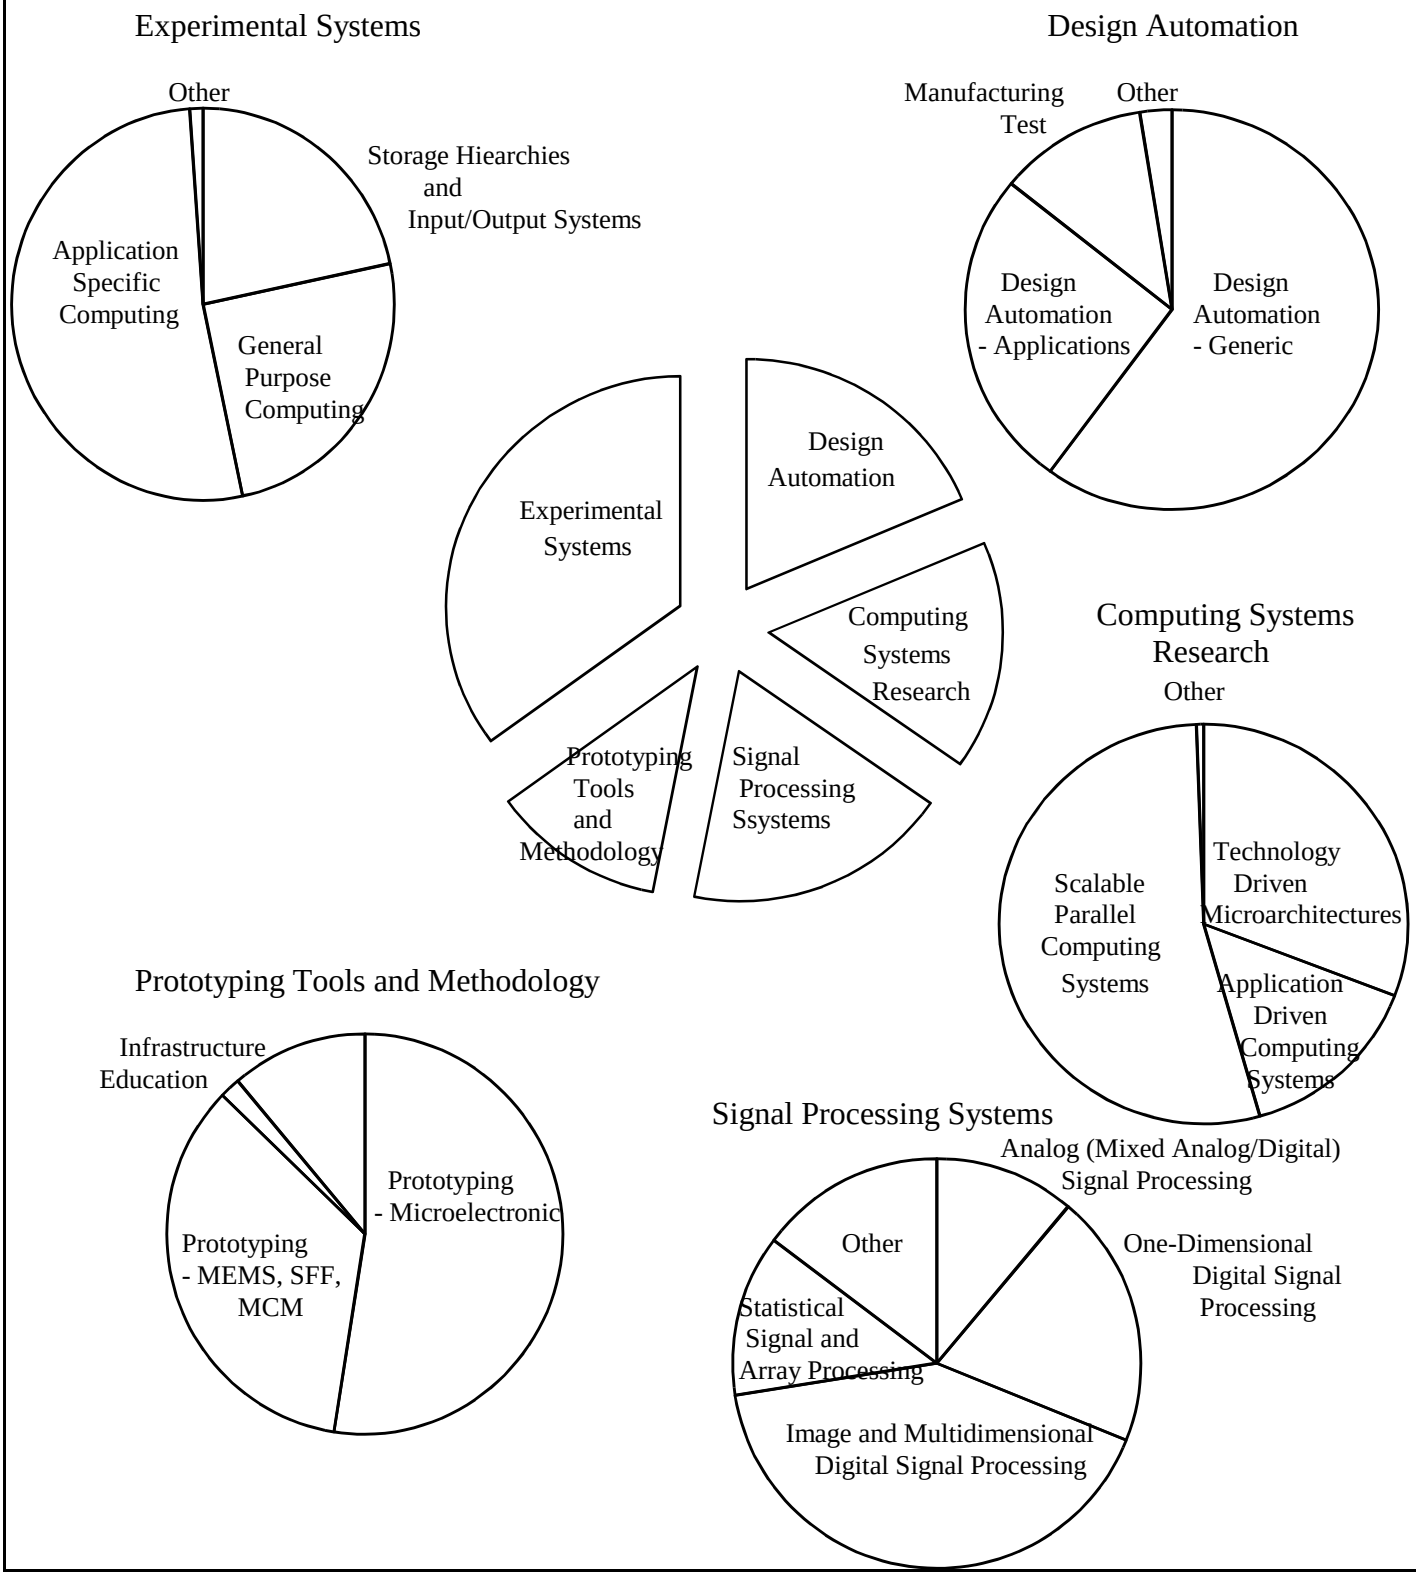

Summary

	Number of Awards	Dollars
Design Automation	57	$4,822,155
Design Automation - Generic	30	$2,884,139
Design Automation - Applications	13	$1,240,502
Manufacturing Test	9	$589,086
Other	5	$108,428
Computing Systems Research	36	$4,227,08
Technology Driven Microarchitectures	14	$1,316,462
Application Driven Computing Systems	6	$612,379
Scalable Parallel Computing Systems	15	$2,278,404
Other	1	$19,835
Signal Processing Systems	55	$4,749,798
Analog (Mixed Analog/Digital) Signal Processing	4	$533,903
One-Dimensional Digital Signal Processing	7	$946,480
Image and Multidimensional Digital Signal Processing	18	$1,959,287
Stistical Signal and Array Processing	11	$595,963
Other	15	$714,165
Experimental Systems	38	$9,059,525
Storage Hiearchies and Input/Output Systems	5	$1,956,829
General Purpose Computing	10	$2,258,447
Application Specific Computing	20	$4,755,568
Other	3	$88,681
Prototyping Tools and Methodology	24	$3,093,033
Prototyping		
Microelectronic	15	$1,626,579
MEMS, SFF, MCM	6	$1,059,990
Education	2	$56,464
Infrastructure	1	$350,000

This summary data includes funds designated for special Foundation initiatives, and equipment matching funds from the Office of Cross-Disciplinary Activities. It does not include program funds use to support Intergovernmental Personnel Act employees, their travel costs, or costs of travel of review panelists and site visitors.

Design Automation
(Formerly Design, Tools and Test)

Dr. Robert B. Grafton, Program Director
(703) 306-1936 rgrafton@note.nsf.gov

The Program

The Design Automation Program supports long-term, basic research in Electronic Design Automation (EDA) and those areas where VLSI design technology is applicable; for example, systems-on-a-chip, embedded systems, and multi-technology (optical, micro-electro-mechanical, etc.) systems. It investigates the scientific methodologies, intellectual processes, abstractions, search paradigms, and information models used in VLSI design. Research covers all phases of the complete EDA design cycle for integrated circuits and systems, from conception through manufacturing test. Goals of the program are:

- Promote discovery, integration, and dissemination of new knowledge in design automation.
- Explore models and methods for design and test in radically new technologies.
- Support research enabling the US to uphold its world leadership in VLSI Design.
- Facilitate a smooth flow of ideas and results to industry and other researchers.

The technology of VLSI circuits changes rapidly. It is now possible to put millions of transistors on a chip and operate at speeds in the hundreds of MHz range. The demand for computing systems of great complexity, performance and trustworthiness is high. VLSI chips and systems of the future will be complex and incorporate technologies ranging from traditional CMOS to optical and mechanical (MEMS). This is causing a re-evaluation of the abstractions on which today's CAD systems are based. Paradigm shifts in design are needed. New abstractions are needed to permit the designer to better manage the complexity of the design. New methodologies need to be developed to cope with technological advances, which make physical phenomena especially important. Design re-use and need for deeper design-space exploration are changing the nature of the design cycle and require new approaches. New tools must be more rapidly disseminated.

> **Of primary importance to research on a new generation of CAD tools is rebuilding the link between CAD researchers and the systems industry. Internships and collaborative research are strongly encouraged.**

Research proposals should contain an explanation of the research strategy and how success will effect the design process, or how it will be measured in the context of a design project. Specifically, proposals should explain the motivation for the work generated by a particular application or innovative target. Related work in industry as well as academia, and the methodology to be used in carrying out and evaluating the research needs to be discussed.

Topics in research can be positioned in a three-dimensional space where the three axes represent the traditional EDA area, the application domain, and the enabling implementation technology. This framework provides a way to assess the potential impact and difficulty of the problems addressed.

Awards

Design Automation - Generic

Auburn University; Richard O Chapman:*CAREER: Formally Verified, Efficient Tools for High-Level Synthesis and Hardware-Software Codesign*; (MIP-9623604); $145,312; 60 months; (Support from: Experimental Program to Stimulate Competitive Research, Design Automation Program - Total Grant $281,354.

This is a program of research and education centered on development of efficient, formally verified CAD tools for high-level synthesis and codesign of digital systems. The notion here is to develop CAD tools which are, themselves, formally verified and guaranteed to synthesize correct designs. The aim is to produce code running on an embedded processor and custom hardware on field programmable gate arrays. The tools are designed to partition a design into hardware and software components, to generate microprocessor code, and to do high-level synthesis. Verification is done by providing semantic models for the representation languages used in the system (behavioral specification, register- transfer level description, gate-level circuit, machine code) and proving that the tools produce designs whose meanings are refinements of the meanings of their specifications. Theoretical and practical classroom courses which compliment this research are being developed.

University of Arkansas; Mitchell A. Thornton; *Design and Implementation of OBDD Variable Ordering/Reordering*; (MIP-96 33085); $57,681; 24 months.

This research is on the design and implementation of ordered binary decision diagram (OBDD) variable ordering and reordering methods. The work explores the use of memory-efficient data structures, called Decision Diagrams (DD), to describe the behavior of Boolean functions. Being developed is a methodology for finding an ordering of the Boolean function variables so that the DD has minimal size. This research exploits the use of conditional probabilities for the purposes of variable ordering and dynamic reordering. Efficient techniques for computation of the spectrum of a Boolean function using its OBDD representation as input are being used. Software is being developed to manipulate OBDDs as well as compute the probability values.

California Institute of Technology; Erik K Antonsson: *Structured Design Methods for MEMS*; (MIP-9529675); $211,584; 24 months; (Support from: Design Automation Program, Prototyping Tools and Methodology Program - Total Grant $260,048).

This research is on developing structured design methods for Micro-Electrical- Mechanical Systems (MEMS). The focus is on MEMS fabrication process simulation and mask-layout synthesis. In analogy to present-day VLSI design, tools to permit rapid, accurate, conservative mask layout synthesis of MEMS are being explored. The objective for the tools is to enable a MEMS designer to specify a micro-mechanical function (e.g. a mechanical spring with particular characteristics), and have a system automatically generate the information (mask- layout, and other fabrication instructions) to create the shape that exhibits the desired function. This includes refining and extending the 3-D anisotropic etch simulation methods. Prototype software design tools are being developed.

Stanford University; Giovanni De Micheli: *Logic Synthesis of Low-Power Circuits*; (MIP-9421129 A001); $69,514; 12 months.

This research is on techniques and tools for automated logic design of low-power, semi-custom circuits. Digital circuits are specified as models in hardware description languages that can be readily compiled into finite-state machines. The latter are described by transition diagrams or by synchronous logic networks. These models are then used to solve logic synthesis problems in encoding sequential circuits, restructuring logic networks, and library binding. Tools for a comprehensive EDA system for low-power design are being developed.

University of California - Los Angeles; Jason Cong*: NYI: Performance Optimization in Layout and Logic Synthesis of VLSI Systems*; (MIP-9357582 A003, A004); $67,500; 12 months.

This research investigates methodologies and efficient algorithms for performance-driven layout design and logic synthesis of high-speed, high-complexity VLSI systems. The research on performance-driven layout focuses on interconnect design and optimization at each stage of layout design

process, including the development of accurate interconnect models which enable efficient layout optimization algorithms, efficient algorithms for circuit partitioning and clustering, interconnect-driven floorplan and placement, interconnect topology design with buffer insertion, simultaneous device and interconnect sizing, and clock layout optimization. The research on performance-driven logic synthesis mainly focus on synthesis and mapping for field-programmable gate-arrays (FPGAs), including optimal or near-optimal algorithms for structural and functional gate decomposition, minimum-delay technology mapping, trade-off of delay and area in technology mapping, combined retiming, technology mapping, and re-synthesis for clock period minimization, and automatic pipelining.

University of California - Los Angeles; Andrew B Kahng: *NYI: Synthesis of High-Speed, High-Complexity VLSI Systems*; (MIP-9257982 A004); $62,500; 12 months.

This research develops a body of knowledge in circuit theory, device physics, and computational techniques for VLSI CAD that address semiconductor industry needs, as well as future research and tool development. Ongoing research includes algorithms for timing/power optimization, layout-driven re-synthesis, and clock distribution for structured-custom, complex ASIC and mixed-signal design methodologies. Other ongoing thrusts are toward topics in data management, e.g., coupling extraction with timing/signal purity analysis tasks; new models of layout tools for constraint-driven layout flows; and search-optimization technologies. Two new topics are also being developed. One is multi-layer constrained area routing, which brings into convergence various technology trends which include: scaling issues related to design complexity and process technology, electro-migration, signal purity, power, thermal and EMI. The second is developing metrics of solution quality and algorithm performance.

University of California - San Diego; Chung-Kuan Cheng: *Research on Circuit Partitioning*; (MIP-9315794 A003); $26,391; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $52,783).

This project is an investigation into a new generation of hierarchical partitioning methods motivated by the need to:

1. cope with high circuit complexity,
2. improve a system's performance, and
3. control intermodule delay.

The principal investigator is extending partitioning research by adopting different circuit models including petri nets, data flows, and state machines, by improving the efficiency and effectiveness of the methods, and by deriving theoretical results on variations of partitioning formulations. Partitioning methods are being applied to netlist mappings for various hardware emulation machines and VLSI design problems.

University of California - San Diego; Kenneth Y Yun: *CAREER: Design Methodologies and Tools for 1 GHz and Beyond*; (MIP-9625034); $215,000; 48 months.

This research explores systematic design methodologies to handle timing and communication in mixed synchronous - asynchronous designs. Formal and automated techniques for managing the design process for independently synchronized domains on a VLSI chip are being developed. Specific research focuses on: mixed- timing design methodologies and components; specification formalism for mixed-timing interfaces, performance driven synthesis and technology mapping techniques for mixed-timing circuits; hierarchical timing verification techniques that check the validity of assumptions used for synthesis and technology mapping, and techniques to estimate and analyze system performance. The education plan provides an environment to nurture students; open-ended problem solving skills; incorporated research activities for undergraduates; opens research opportunities to underrepresented groups; and develops industrial internships in the graduate research program.

University of Southern California; Massoud Pedram: *NYI: Low Power VLSI Design*; (MIP-9457392 A002); $31,250; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $62,500).

This research investigates modeling and estimation of power consumption as well as techniques for minimizing power at the various levels of design abstraction (layout, logic, register- transfer and behavioral levels). Principles and methods to guide the design of power efficient electronic systems are being explored; and the impact of availability of low-power design techniques on chip, module, and system level designs is being assessed. Topics being investigated include: spatio-temporal power estimation; state assignment for low power; power dissipation in boolean networks; common subexpression extraction; and FPGA synthesis for low power.

Colorado State University; Tom Chen, Charles W Anderson, Anneliese VonMayrhauser: *Behavioral Level Design Verifications Using Software Testing Techniques and Neural Networks*; (MIP-9628770); $314,313; 36 months.

A structured methodology for validating and verifying behavioral models is being explored. The methodology is adapted from a set of software test and validation criteria, which are based on control and data flow properties of the behavioral model under test. The test criteria are measured by a set of coverage metrics which indicates the types and the locations of untested portions in the model. The methodology includes an adaptive test data generation technique using neural networks. Given the untested portions in the model, effective test data are generated and identified by neural networks to exercise the untested portions. Tools which will provide designers with a quantitative confidence index at the end of their testing and verification efforts are being developed.

University of Colorado - Boulder; Gary Hachtel, Fabio Somenzi, Michael Lightner: *An Integrated VHDL-based Synthesis and Verification System for VLSI Systems*; (MIP-9422268 A001); $250,000; 12 months.

This project is on synthesis and verification of digital systems. The computational basis for the work is the binary decision diagram (BDD) data structure and extensions. In low power circuit design, new ideas in BDD technology are being used for synthesis algorithms, and to estimate power consumption via probablistic behavior of circuits. Decomposition concepts, such as tearing, to assess properties of very large circuits are being investigated. In verification, approximate exploration ideas are being examined for use in checking equivalence of very large circuits. Hierarchical verification capabilities, where parts of the circuit are modeled - bit level and word level, are being examined. To provide a sound connection between high level synthesis and high level verification, as well as to validate high level VHDL descriptions, refined comparisons of non-deterministic systems, such as bi-simulation equivalence, testing equivalence, and pre-orders are being explored.

University of Colorado - Boulder; Gary Hachtel, Fabio Somenzi: *U. S. -Germany Cooperative Research on Efficient Data Structures for Computer-Aided Design*; (INT-9514775); $3,300; 24 months; (Support from: Western Europe Program, Design Automation Program - Total Grant $9,240).

This award supports Professors Gary Hachtel and Fabio Somenzi, and a graduate student, all from the University of Colorado, to collaborate in computer science research with Professor Christoph Meinel and others of the Department of Computer Science and Complexity Theory and VLSI-Design of the University of Trier, Germany. They are working in the area of computer-aided circuit design, focusing on the synthesis and verification of digital systems using `binary decision diagram` data structures. Together they will explore new classes of these efficient data structures and their behaviors in specific applications. Previous work by the group - Trier on Ordered Binary Decision Diagrams (OBDDs) has focused primarily on the theoretical foundations and the connection with complexity theory. The US group has been motivated by real problems of electrical engineering such as sequential circuit verification, optimization and testing. The plan for collaboration will transfer and combine the complementary knowledge and experience of the two groups. The resulting synergism should result in contributions to both theoretical and practical aspects of circuit design.

University of Delaware; Phillip Christie: *Statistical Analysis of Interconnect-Limited Systems*; (MIP-9414187 A002, A003); $28,749; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $57,499).

The growth in size and complexity of modern computers is matched by the need for more accurate models of their internal wiring structure. The task of modeling this complexity is ideally suited to a recent development in statistical physics called the re-normalization group. It has the power to attack the wiring optimization problem because it takes this one very difficult problem and breaks it up into a very large number of much simpler problems. Rather than attempting to attack the entire range of partitioning and placement requirements from the transistor level to the back-plane, an optimization procedure need only be found for a highly restricted wire length range. Once this has been derived, the procedure is renormalized for all other length scales in a manner which generates a system-wide solution. This project will investigate how such an approach may be used to develop much better estimates of interconnect-limited system performance and how this theory leads to new optimization algorithms. In contrast to other theories

the predictions made by this technique will grow ever more accurate as the systems grow larger and more complex.

University of South Florida; Peter M Maurer: *Improving the Performance of Digital Logic Simulation*; (MIP-9403414 A001); $64,946; 12 months.

This research is on compiled simulation of synchronous and asynchronous circuits. A digital logic simulation algorithm, called the Inversion Algorithm, is being investigated. This is a very efficient algorithm, yet capable of simulating any digital circuit. Basic extensions are being made to the algorithm. These involve including unit and multi-delay timing models, a logic model with unknown value, new gates, new types of functions, as well as bit-parallel simulation. Further, it is being redesigned so as to reduce both the complexity of the simulation and the number of events processed during simulation. Additional activities include extending the existing model to a multiple processor implementation, and to handle asynchronous sequential circuits. Application of the Inversion Algorithm to hierarchical simulation and incremental compilation is being explored.

Northwestern University; Prithviraj Banerjee: *Parallel Algorithms for Synthesis and Test*; (MIP-9696164); $42,0000; 12 months.

This research is on developing circuit design and test algorithms that run on parallel computers. These include efficient, asynchronous, portable parallel algorithms for;

1. synthesis of combinational circuits, and
2. test generation and fault simulation.

Algorithms are being written using an environment that makes it possible to port CAD applications across a wide range of MIMD machines. In addition they are designed to allow a maximum overlap of computation and communication. The algorithms are being tested on several parallel platforms.

Northwestern University; Majid Sarrafzadeh: *Algorithm Design for VLSI Layout*; (MIP-9527389); $130,772; 24 months.

This project is on algorithmic approaches to VLSI. Problems are being investigated are:

1. Literal minimization in logic synthesis using an approach based on term intersection graphs.
2. Gate level placement problems using minimum floating Steiner trees.
3. Budget assignment in timing driven placement.
4. Clock tree design using a gated design which minimizes activity, hence power consumption.

A coherent design system, incorporating above stages, is being developed.

University of Illinois - Urbana-Champaign; C. L. Liu: *Research in Computer-Aided Design of VLSI Circuits*; (MIP-9222408 A004); $55,600; 12 months.

The research is on algorithms for high level synthesis and layout of VLSI CAD designs. Algorithms for complex, large industrial type design problems are being developed. Research topics include:

1. high level synthesis with testability as a goodness measure,
2. timing driven placement algorithms for FPGAs,
3. routing in which cross talk is considered.

For the first problem, the effect of register allocation and functional unit binding on the testability of the circuit is being examined. Then the scheduling step is being examined. Research on the second problem is based on the notion of a "neighborhood graph" which is used in guiding an iterative improvement algorithm that produces placements which satisfy given timing constraints. An integer programming approach is being used for the third problem, to avoid parallel long wires that are close together in the routing solution.

Purdue University; Kaushik Roy: *Logic and Circuit Synthesis for Low-Energy Computing Using Principles of Adiabatic Switching*; (MIP-9633516); $45,146; 18 months.

This research is on adiabatic logic design for very low energy consumption circuits and systems. The primary idea is to use principles of adiabatic switching as the basis for circuit design. Energy recovery using this notion is a relatively new idea, although the concept of reversible logic and zero-energy computing can be tracked back to he early 1970s, the attempt to realize the concept in electronic circuits is a new endeavor. An adiabatic logic, using both fully reversible and partially reversible logic, is being developed. It is based on standard CMOS technology with a switching power supply. Basic logic gates, registers, and memories for recovered energy operations, are being developed. Tools for power management, such as techniques to utilize the redundant signals generated from reversible logic gates, are being developed. The circuits and the power supply are being fabricated and tested.

University of Notre Dame; Edwin H Sha: *High-Level Design Methodologies for Time-Optimal and Memory-Optimal Systems*; (MIP-9501006 A001); $25,000.

This research is on optimization algorithms for synthesis of multi- level loops, which occur in time and memory critical parts of scientific computing applications. The nested loops are modeled as multi-dimensional data-flow Graphs (MDFG); and algorithms taking advantage of the multi-dimensionality are being designed. By considering the multi-dimensional iteration space and the iteration body simultaneously, the transformation and optimization techniques are able to optimize throughput and memory requirement - the behavior level. Research topics include: graph transformation and optimization; data scheduling; and co-design. Polynomial-time algorithms for various graph models are being developed. This avoids exponential integer linear programming approaches.

Massachusetts Institute of Technology; Srinivas Devadas: *NYI: Formal Methods for Hardware and Software Verification*; (MIP-9258376 A005); $62,500; 12 months.

This research is on methodologies for design of large-scale digital systems containing both application-specific and re-programmable hardware, where most (or all) of the circuitry will be on a single chip. A design methodology that optimizes power dissipation of the embedded system, while meeting area and performance constraints, is being developed. Methods to generate small, dense code to reside on the chip, and verification methods that guarantee design correctness are being explored.

University of North Carolina - Charlotte; Dian Zhou: *NYI: Performance-Driven VLSI Designs*; (MIP-9457402 A002); $32,750; 12 months; (Support from: Signal Processing Systems Program, Design Automation Program - Total Grant $65,500).

This research merges two efforts in asynchronous circuit compilation. These are the work of Gopalakrishnan on the language "hopCP" and its use in verification; and the work of Brunvand on asynchronous circuit compilation. The goal is to enhance the expressive power as well as the semantic clarity of concurrent hardware description languages for asynchronous circuits and systems. The research effort includes: finding extensions to the formal basis for compiling from HDL's to circuit designs; exploring formal characterizations and the optimizations used in asynchronous circuit compilation; and studying the performance of implemented circuits with regard to a variety of parameters.

Portland State University; Malgorzata Chrzanowska-Jeske: *Pseudo-Symmetric Binary Decision Diagrams (PSBDDs)*; (MIP-9629419); $183,550; 36 months.

This research is on a data structure, Pseudo-Symmetric Binary Decision Diagrams (PSBDDs), for completely specified Boolean functions. The structure of an OBDD for a totally symmetric function is used as a model for PSBDDs. The Shannon expansion is used to generate the vertices of PSBDDs. There is a join operation which combines two adjacent vertices such that the function is represented as a pseudo-symmetric network instead of a binary tree. The research consists of: a) developing a PSBDD package with four types of symmetry; b) investigating the best heuristic for variable ordering, and ordering sets of symmetric variables; c) comparing mapping results for large functions between PSBDDs and BDDs; d) developing strategies for efficient generation of PSBDDs for incompletely specified functions; e) developing diagrams using Davio I and Davio II expansions; and f) suggesting new architectures for Cellular-Architecture type FPGAs which are design-automation friendly.

Carnegie-Mellon University; Edmund M Clarke: *Automatic Verification of Finite-State Concurrent Systems in Hardware and Software*; (CCR-9217549 A003); $75,000; 12 months; (Support from: Design Automation Program, Software Engineering - Total Grant $150,000).

Logical errors in sequential circuit designs and communication protocols have always been an important problem. The most widely used method for verifying such systems is based on extensive simulation and can easily miss significant errors when the number of possible states is very large. This research deals with developing an alternative approach based on a technique called temporal logic model checking. In this approach specifications are expressed in a propositional temporal logic, and sequential circuits and communication protocols are modeled as state transition systems. An efficient search procedure is used to determine automatically if the specifications are satisfied by the transition system.

Texas A & M University; Dhiraj K Pradhan, Duncan M H Walker, Wolfgang Kunz: *Novel Methods in Computer-Aided Circuit Design and Testing Using Recursive Learning*; (MIP-9406946 A002); $108,731; 12 months.

This research investigates modeling and estimation of power consumption as well as techniques for minimizing power at various levels of design abstraction (layout, logic, register- transfer and behavioral levels). Principles and methods to guide the design of power efficient electronic systems are being explored; and the impact of availability of low-power design techniques on chip, module, and system level designs is being assessed. Topics being investigated include: spatio-temporal power estimation; state assignment for low power; power dissipation in boolean networks; common subexpression extraction; and FPGA synthesis for low power.

University of Texas - Austin; Margarida Jacome: *CAREER: Design Assessment and Reuse*; (MIP-9624231); $220,050; 48 months.

This research is on design reuse and design process assessment. Design reuse can be in any part of the design cycle from initial phases (capture, processing, and organization of design information), to final phases (matching, retrieval and actual reuse of design information). Mechanisms for supporting this activity are being explored. Attention is being paid to maximizing opportunities for reuse, and assessing the impact on productivity of the reuse infrastructure. Issues being addressed include the characterization and representation of design objects and design processes, automatic capture, indexing, retrieval and reuse of design information. The second investigation is developing a framework, including risk assessment and planning tools, for design process assessment. Metrics or measuring management objectives and parameters of the design process are being identified and characterized. Results of the research are being tested and validated in the design of large embedded systems.

University of Utah; Erik L Brunvand, Ganesh Gopalakrishnan: *The Design of Asynchronous Circuits and Systems with Emphasis on Correctness and Proven Optimizations*; (MIP-9215878 A002); $5,000.

This research merges two efforts in asynchronous circuit compilation. These are the work of Gopalakrishnan on the language hopCP and its use in verification; and the work of Brunvand on asynchronous circuit compilation. The research is:

1. enhancing the expressive power and semantic clarity of concurrent hardware description languages for asynchronous circuits and systems
2. extending the formal basis for compiling from HDL's to circuit designs;
3. formally characterizing and improving the optimizations used in asynchronous circuit compilation; and
4. studying the performance of implemented circuits with regard to a variety of parameters.

University of Utah; Ganesh Gopalakrishnan: *A Multi-Paradigm Verification System Tailored for the Design Refinement Cycle*; (MIP-9321836 A002); $15,000.

The design of a VLSI system involves multiple design representations; and the design must go through several iterations aimed at meeting many performance and cost constraints. Verification that the design meets constraints is necessary. This research is developing rigorous verification methods that span multiple design representations, accommodate design revisions, and provide incisive partial verification methods (e. g. verification focussed on the "corners" of the behavioral space) that fit within designers' time budgets. These ideas are being validated by verifying the real asynchronous designs.

University of Utah; Chris J Myers:*CAREER: Design Methods and Tools for Mixed-Timed Systems*; (MIP-9625014); $210,000; 48 months.

This research investigates formalization and automation techniques for design of asynchronous systems. Techniques for managing multiple, independently synchronized domains under a methodology in which timing and communication can be explicitly specified, synthesized, and verified are being explored. Specific tools and techniques being developed for the design of mixed-timed systems include: a performance and power optimizing compiler; a synthesis tool to produce technology-independent circuits; a technology mapping tool that takes the implementation and maps it to high- performance gate libraries; and a timing verification tool. In addition, new design methods are being developed to produce: high-performance mixed-timed data paths, and an interface between synchronous and asynchronous modules.

The PI's education plan includes;

1. curriculum emphasis on communication skills,
2. development of new computer engineering courses,
3. promoting interdisciplinary activities,
4. incorporating research activities into the

curriculum through class and summer projects, and
5. fostering interaction with industry through student summer internships.

University of Virginia - Charlottesville; Gabriel Robins: *NYI: New Directions in High-Performance VLSI Layout*; (MIP-9457412 A003); $62,500; 12 months.

Realistic formulations of performance-driven layout are the focus of this research. Accurate models of circuit delay are being sought. These models include technology parameters, such as capacitance, resistance, inductance, etc. The approach is to study delay-optimal trees to define an envelope of achievable routing performance. Methods for constructing near-optimal layouts are being investigated. Additionally, the routing problem is being recast as one of constructing low-delay routing graphs where cycles are allowed. This can have the advantage of designs being tolerant to certain types of open faults due to manufacturing defects or electro-migration.

University of Washington; Steven M Burns: *NYI: A Design Language for Asynchronous Circuit Synthesis*; (MIP-9257987 A004); $62,500; 12 months.

The long term goal of this project is to find a mechanism for capturing the complete design of an asynchronous digital system. Thus, this research is on mechanisms for representing design decisions, and finding ways to construct transformations of from high-level descriptions to a low-level design. Research includes:
1. developing algorithms for technology mappings of speed-independent circuits;
2. exploring methods to formally verify correctness of design transformations;
3. finding methods for analyzing timing of asynchronous systems; and
4. development of a rapid prototyping facility (MONTAGE) for field-programmable gate arrays.

Design Automation - Applications

University of California-Santa Barbara; Forrest Brewer: *Production Language Based High-Level Synthesis*; (MIP-9320752 A002); $56,002; 12 months.

The overall goal of this research is to create a new class of synthesis tools which address the design of complex controller-data path machines under constraints of pre-defined interfaces. Because the results will be integrated with commercial EDA design tools, a design output format, which can be simulated and allows automated re-design of selected portions of the design, is being developed. In a second task, approximate sequential reachability analysis is being used to design and implement algorithms for optimizing and partitioning controller designs. The third task is to explore scheduling algorithms for both the control and data-path portions of the design. Finally, an optimizing compiler is being built. It contains algorithms which solve encoding issues for high performance designs, and performs re-scheduling of the data-path operations to minimize required resources while maintaining design behavior.

University of California-Santa Barbara; Malgorzata Marek-Sadowska: *Research on Layout and Logic Design*; (MIP-9419119 A001); $103,000; 12 months; (Support from: Prototyping Tools and Methodology Program, Design Automation Program - Total Grant $143,000).

This research is on layout driven synthesis, which is at the intersection of logic synthesis and physical design. The approach is to determine algorithms and methodologies for restructuring logic networks in synthesized digital systems. Three topics, which address the problems inherent in designing and testing in dense, high speed, circuits are being investigated. Topics include:
1. Incremental restructuring of Boolean networks
2. Logic synthesis for testability;
3. Logic synthesis for low power,
4. XOR/AND representation of Boolean functions; and
5. Tools for field programmable gate arrays.

New techniques are being applied to partitioning, routability correction, improving testability, automation of engineering change, free scan in sequential circuits, multilevel synthesis, and FPGA routing.

Georgia Institute of Technology; Giorgio Casinovi: *Computer-Aided Simulation of Optoelectronic Integrated Circuits*; (MIP-9523436); $142,142; 24 months.

The integration of electronic and photonic devices on the same chip offers performance advantages in terms of speed, bandwidth, immunity to interference, power dissipation and interconnection density. This research explores algorithms and tools for simulation of optoelectronic Integrated Circuit (OEIC) designs. The goal is to simulate equally accurately both optical and electronic devices in the same circuit. Both lumped and distributed-constant elements are included in the model. Existing models for optoelectronic devices such as light emitters, detectors, and electro-optic modulators are being refined and incorporated in the simulator. New models are being developed for optical interconnect and for those optoelectronic devices for which no models suitable for time-for domain simulation currently exist. A simulation environment, capable of handling circuits with optical and electronic devices is being constructed.

University of Illinois - Urbana-Champaign; Rajesh K Gupta: *Architecture and Synthesis Techniques for Embedded Systems*; (MIP-9501615 A001); $10,000.

This research is on synthesis techniques for micro-electronics- based "embedded systems". An embedded system is one designed for a specific functionality under stringent constraints on its timing performance and cost. Design automation techniques are being developed for use in a framework where the designer can assess tradeoffs between various embedded system architectures and designs. The main capabilities of this framework are;

1. timing analysis for both execution delay and rate constraints,
2. embeddable software and runtime system generation,
3. software size-performance tradeoffs, and
4. cost- performance analysis of architectural alternatives.

This work is exploring system partitioning and transformation techniques to build system implementations that "guarantee" constraint satisfaction while optimizing system development cost.

University of Illinois - Urbana-Champaign; Naresh Shanbhag:*CAREER: Design of Giga-Scale CMOS Communications Systems: An Integrated Approach*; (MIP-9623737); $220,000; 48 months.

The research objective is to develop methodologies for the design of giga-scale CMOS systems. A focus is high bit rate digital communications. The approach transcends boundaries between algorithmic, architectural, logic, circuit and technological aspects of the VLSI design methodology. The research has two parts. The first is an investigation of fundamental limits on power area speed and reliability in giga-scale computation. Activities include finding achievable bounds on the power dissipation, area, speed and reliability for any given algorithm and implementation technology.

In the second part, algorithms, architectures, and systems for design of high-bit rate communications systems are being investigated. This includes developing power, area, and speed-optimal VLSI algorithms and architectures for high-bit rate digital communications and signal processing applications. As an application, systems for 155.52 Mb/s transmission over twisted-pair wiring for ATM and IMTV, and wireless spread-spectrum systems are being built.

University of Massachusetts Amherst; Israel Koren: *Topological and Physical Layout Design Techniques for Yield Enhancement*; (MIP-9305912 A001); $70,000; 12 months.

This research is on design tools which can compensate for manufacturing yield losses due to the complexity of the logic. Preliminary results have shown that yield can be enhanced by the technique of reducing the sensitivity of the chip to point defects. Yield optimization algorithms and techniques which are applicable to the last two stages of design - topological/symbolic, and physical layout are being explored. Physical layout strategies for yield enhancement in the channel routing and compaction phases for standard cell designs are being investigated. Yield enhancements in topological designs are being illustrated through PLA-based designs.

Michigan Technological University; Jeffrey O Coleman: *RIA: Convex-Programming Design of Signals and Systems*; (MIP-9409686 A001); $2,500; (Support from: Signal Processing Systems Program, Design Automation Program - Total Grant $5,000).

The PI is creating a special-purpose programming language in which optimization problems can be specified in a natural and direct way. The focus is on developing associated translation software to convert such a specification into a particular canonical form for numerical optimization using recently developed, ultra-efficient, interior-point algorithms. The canonical form is a linear-matrix- inequality (LMI) program, where each constraint takes the form of a requirement that a linear (plus a constant) matrix function of the optimization variables be positive definite.

Control theorists have recently demonstrated that a tremendous variety of common (and uncommon) constraints can be put into this generic form. Often, however, the required LMI constraints are not related to the underlying problem in an intuitive way. Software to translate specifications in a "comfortable language" to sets of LMIs is being developed in order to make these powerful optimization techniques easy to apply. The language is being applied to real problems in communication and signal-processing circuit designs.

University of Michigan Ann Arbor; Karem A. Sakallah, Trevor N Mudge, Edward S Davidson: *Timing Issues in the Design of Digital Systems*; (MIP-9404632 A001); $116,678; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $175,017).

This research is on a timing verification and optimization framework for designing an entire digital system (e. g. a microprocessor). The research builds on a widely used model for synchronous timing analysis and an efficient method for estimating gate and wire delays. The model is being extended to include relevant functional information in order to enhance accuracy. Components of the framework are: design decomposition to isolate critical elements; a path delay calculator; algorithms for finding synchronizer components; clock analysis algorithms; a symbolic sequential timing verification component; a hybrid timing-logic simulator; and design optimizers.

Princeton University; Sharad Malik: *NYI: Design Automation for Embedded Systems*; (MIP-9457396 A002); $60,000; 12 months.

High fabrication costs and decreasing time-to-market for products require an increase of the programmable component of integrated circuits. This changes circuit design from just gates in an array, to a combination of instructions running on a core processor coupled with a gate array. Gates, as basic units of computation on silicon, are well understood; but use of embedded instructions as basic units of computation on silicon is not. Three problems are the focus of this research. First is performance analysis of embedded software, with examination of path analysis and micro-architecture modeling. Second is power analysis of embedded software, with development of an instruction level power model. Third is retargetable code generation for digital signal processors, with the application of code generation problems to expression trees for DSPs.

State University of New York - Stony Brook; Michael M Green: *NYI: Improved Circuit Simulation Using Results from Circuit Theory*; (MIP-9457387 A002); $31,250; 12 months; (Support from: Signal Processing Systems Program, Design Automation Program - Total Grant $62,500).

This research is applying the principal investigator's previous work in the area of nonlinear circuit theory to make major enhancements to the way designers simulate analog circuits. In particular, improvements to the continuation methods of solving dc operating points of circuits are being made guaranteeing that all of a circuit's operating points will be found during a single analysis. Moreover, continuation methods are being applied to sensitivity analysis of circuits; for example, by making observations of a continuation curve, a designer could determine whether a circuit is prone to a latch up condition. Erroneous models are thought to be a major source of convergence problems and erroneous results in circuit simulation. Another enhancement to circuit simulation includes checking the accuracy of transistor models by verifying that all models satisfy passivity and the no- gain condition.

University of Utah; Erik L Brunvand, Ganesh Gopalakrishnan, Alan L Davis: *Application-Driven Advancement of Asynchronous Design Methods*; (MIP-9622587); $157,076; 12 months.

This is an investigation into the application of asynchronous circuit design tools to a large, realistic example. The project is being carried out in conjunction with the Avalanche parallel architecture project - the University of Utah. The Cache and Communication Control Unit (CCCU) of the Avalanche multiprocessor, which is designed as a synchronous system, is being implemented as a self-timed system using the original specification for the Avalanche CCCU. To do this, problems in self-timed circuit design are being deeply investigated, with a particular focus on automating the design of high-performance self-timed systems. Chips embodying the designs are being fabricated through MOSIS, tested and evaluated.

University of Virginia - Charlottseville; Joanne Dugan, Barry W Johnson: *Innovative Dependability Analysis Techniques in an Advanced Design Environment*; (MIP-9528258, A001); $209,354; 24 months.

This research explores the relationships between models used to refine the system design, to assess the real-time behavior of the system, and to assess the fault tolerance. These relationships are being combined to develop tools for system- level modeling with dependability analysis of the design. The approach builds on two recently developed tools, ADEPT (Advanced Design Environment Prototype Tool) and DREDD (Dependability and Risk Evaluation using Decision Diagrams). Advanced techniques for dependability analysis of embedded computer systems are being put into the ADEPT design environment. There, a dynamic fault tree model is automatically generated from the ADEPT system design model. The fault tree models uses a combination of Markov techniques, combinatorial approaches and binary decision diagrams. It can model permanent and transient hardware faults, unrelated and related software, automatic recovery and reconfiguration management, sequence dependencies and other dynamic behavior.

University of Washington; Andrew T Yang: *NYI: Modeling and Simulation of Advanced Microelectronics and Optoelectronic Circuits and Systems*; (MIP-9257279 A004); $62,500; 12 months.

This research is on design of deep submicron technology, <0.5 um, circuits. Effects that were previously disregarded are now important in design. For example, reduced power supply voltage and higher operating frequency lead to more noise, signal integrity problems, and devices that appear more analog. Furthermore manufacturing process control has not kept up with reduction in dimensions, leading to designs that are vulnerable to process variations. Topics include:

1. a computer aided design methodology for improving product yield in manufacturing
2. substrate power supply switching noise modeling and simulation; and
3. reduction of interconnect and substrate RC networks from post layout extraction.

Algorithms for solving these problems are being developed.

Manufacturing Test

Stanford University; Edward J McCluskey: *Research on Reliable Computers*; (MIP-9107760 A003); $150,002; 12 months.

This research is on techniques for reducing the occurrence of run-time errors in circuits and systems. Emphasis is on preventing the introduction of faults into the system through verification and synthesis techniques, and on improving the detection and diagnosis of faults causing run-time errors so that the faults can be removed from the system. Topics being pursued are:

1. finding synthesis methods that automatically synthesize a synchronous, register-transfer level (RTL) hardware specification from a behavioral VHDL language specification;
2. determining what coverage of multiple stuck-at faults or bridging faults can be expected or guaranteed by a test for delay faults in an arbitrary circuits;
3. investigating methods for utilizing the use of output waveform characteristics in delay testing; and
4. constructing fault models and methods for detecting intermittent failures.

University of California-Santa Cruz; Frankie J Ferguson: *PYI: Hierarchal Test Pattern Generation for Manufacturing Defects*; (MIP-9158491 A007); $48,530; 12 months.

Two approaches to manufacturing test methodologies are being explored. These are the use of both high level and low level fault models to generate tests that detect virtually all plausible manufacturing defects. The research vehicle is a defect modeling tool, "Carafe" (circuit and realistic fault detector), which determines the most likely faults to occur in a CMOS circuit. First, Carafe is being modified to include algorithms for location of three dimensional defects, and locate shorts that are between nodes on different layers. Second, it is being used to investigate physical design for testability at both the chip and logic cell levels. Third, Carafe is being used to develop algorithms for fault isolation and diagnosis.

University of California-Santa Cruz; Tracy Larrabee; *PYI: Sequential and Combinational Test Pattern Generation for Realistic Faults Using Boolean Satisfiability*; (MIP-9158490 A007); $37,500; 12 months.

This research is based on an automatic test pattern generation (ATPG) system (called Nemesis) that generates a test pattern for a given fault by first constructing a formula representing all possible tests for the fault, and then applying a Boolean satisfiability algorithm to the resulting formula. This method separates the formula extraction from the formula satisfaction thus providing flexibility and generality. A testing system, based on Nemesis, that will generate tests detecting all realistic manufacturing defects in both combinational and sequential ICs is being developed.

University of Iowa; Irith Pomeranz: *NYI: A New Search Strategy for Design Automation*; (MIP-9357581 A003); $41,880; 12 months.

This research explores a design strategy, "Incremental Dynamic Optimization Based Learning". This is a general purpose approach to obtaining high quality solutions to synchronous, sequential VLSI designs. This strategy is based on scaling down a design problem, by reducing circuit size, while retaining critical details. Optimal solutions are found for a sequence of small designs, common features are extracted, and these are scaled up to a solution for the original problem. This method permits exploration of the tradeoff between speed of creating a design and its quality. Additionally, tradeoffs between two methods for test generation and diagnosis are being explored. These are: three valued logic (which is fast) and state enumeration (which is accurate, but slow). Applications to very large circuits are being made.

University of Massachusetts - Amherst; Premachandran Menon: *Delay-Verifiable Combinational and Non-Scan Sequential Circuits*; (MIP-9320886 A001); $29,999; 12 months.

This research explores the concept of a delay-verification test set, which detects the presence of any path delay faults that can affect the timing of the circuit. The approach is to find necessary and sufficient conditions for delay verifiability; i. e. the existence of a complete delay-verification test set. Algorithms for the synthesis of area-efficient, delay-verifiable circuits and the generation of delay verification test sets are being developed. A second topic being explored is a design and test methodology for sequential circuits without scan latches in functional paths. A method, based on clock suppression, is being used to generate transitions on state variables for delay testing. State assignment techniques to improve delay-fault testability without scan and clock suppression are being investigated.

Princeton University; Niraj K Jha: *High-Level Synthesis for Hierarchical Testability*; (MIP-9319269 A002, A003); $81,050; 12 months.

This research explores efficient hierarchical testability techniques for controller-data path systems. Module level test sets, derived for any suitable fault model, are used with high level design synthesis. The resulting test sets are combined into a system level test set which provides complete (or nearly complete) test coverage of all the embedded modules. These testability techniques are being embedded into algorithms for design of scheduling and allocation circuits. Key design criteria are low power and testability.

University of North Carolina - Charlotte; Rafic Z Makki: *New Approaches to the Testing of Digital Integrated Circuits*; (MIP-9406931 A001); $3,580; (Support from: Africa, Near East and South Asia Program, Design Automation Program - Total Grant $7,160).

This research explores test methods for manufacturing test of CMOS integrated circuit (IC) chips based on monitoring of the dynamic portion of the power supply current (iDDT). Since an iDDT pulse is created any time a CMOS circuit switches states, the monitoring of iDDT provides observability

into the switching behavior of the circuit. This, combined with standard test methods, can provide improved testability of manufactured IC chips. New accurate fault models based on actual physical tests as opposed to simulation alone, and new design methods for enhancing iDDT-testability are being investigated. The methods are being incorporated into a design-for-testability tool, and are being evaluated on real designs.

Case Western Reserve University; Daniel G Saab: *An Adaptive Approach to Automatic Test Pattern Generation*; (MIP-9528931); $131,545; 24 months.

This research investigates the generation of manufacturing tests for highly sequential VLSI circuits. The approach is based on combining the best features of deterministic and simulation based techniques. Deterministic techniques are mainly used to identify untestable and redundant faults, and to correct the direction of the simulation.

Simulation based test vector generation examines a set of test vectors to determine a new test sequence to detect some faults. This process, while accurate and fast, generates longer than necessary test vectors. The technique being explored is detecting and correcting the genetic search, using deterministic ATPG. This process involves an intricate exploration of the power of a switch-level logic and fault simulator in an combination with a genetic algorithm and the deterministic ATPG during the test generation process. The technique is being applied to design problems such as, circuit partitioning, initial test set finding, and circuit initialization- structure to determine performance and quality of generated test sets.

University of Wisconsin-Madison; Charles R Kime: *Built-In Self-Test for Random Pattern Resistant Faults*; (MIP-9319742 A002); $65,000; 12 months.

This research is on techniques for testing circuits having random pattern resistant faults. The focus is on methods for synthesizing random pattern testable circuits. The synthesis algorithms being developed include algorithms for design of random pattern testable two level and multi-level combinational logic; extensions to some classes of sequential logic are also being investigated. Two methods being pursued are weighted cellular automaton, and fixed bit biased pseudo-random pattern generator.

Other

Tufts University; Karen P Lentz: *Compression and Interaction Algorithms for Modeling and Simulation Environments*; (MIP-9528194); $8,998; 18 months; (Support from: Design Automation Program, Software Engineering Program - Total Grant $17,996).

This is a research planning grant developing modeling and simulation aigorithms to assess system performance and to recreate events so as to interpret cause- effect relationships. The research planning activities focus on investigating the potential for developing accurate, general purpose modeling techniques that include compressing multiple experiments (variations of effects on the model) while allowing internal observation to expose the perturbations that caused the outcome. The grant resources are being used to contact other researchers involved in related efforts; and to investigate issures involved in expanding existing concurrent algorithms to areas beyond digital logic modeling and experimentation.

University of Cincinnati; Ranganadha R Vemuri, Harold W Carter: *Innovation in VLSI Systems Education: Integration of High Level Synthesis, Analysis and Test Generation Research Results into Undergrad Minor Curriculum in VLSI Systems Engineering*; (CDA-9634462); $65,000; 36 months; (Support from: Design Automation Program, CISE Institutional Infrastructure - Total Grant $249,042).

This CISE Educational Innovation award supports the integration of research in high-level VLSI synthesis, performance analysis and test vector generation into the undergraduate curriculum in VLSI systems engineering. The model curriculum will contain topics on the development of high-level functional and behavioral modeling of VLSI designs, use of high-level VLSI synthesis techniques and tools, developing and using performance models throughout the design process, and generation of design verification tests. Curriculum support material will include mini-books on these topics, term project modules, and, most importantly, free software tools for high-level synthesis, performance analysis

and test generation. All of the curriculum material will disseminated through the Internet and the World-Wide-Web.

University of Virginia - Charlottesville; Gabriel Robins: *Joint Research Between VLSI CAD and MEMS Areas: The Fifth Physical Design Workshop; April 15-17, 1996; Reston, Virginia*; (MIP-9531666); $12,500; 12 months.

The Micro-Electro-Mechanical Systems (MEMS) process technology has matured to the point where it is possible to fabricate a chip with both electrical and mechanical pomponents on it. However, tools and design methodologies are lacking. This workshop brought together design researchers from the VLSI CAD and MEMS design fields to discuss and develop ideas for a new design methodology for combined MEMS/VLSI structures.

Virginia Polytechnic Institute and State University; Dong Ha; *Infrastructure for VLSI Circuit Testing: Development of a Fault Simulator and Test Generators for Industrial Circuits*; (MIP-9632625); 24 months; (Support from: CISE Institutional Infrastructure, Design Automation Program - Total Grant $154,746).

VLSI chip testing tools, produced in research projects - the Virginia Polytechnic University, are being developed and enhanced. The tools are intended to support teaching and research in academia. These tools include, a combinational circuit automatic test pattern generator (ATPG), a sequential fault simulator, a sequential ATPG for both synchronous and asynchronous sequential circuits, and two parsers for analyzing circuit designs in the languages VHDL and EDIF. A benchmark collection of industrial type circuits is being developed. The tools are being developed so that they;

1. are easy to install and use,
2. Are maintenance free,
3. provide a variety of options to support research and teaching activities, and
4. have source code that is easy to read, modify and extend.

The tools and their source code are being distributed freely to universities for research and teaching.

University of Washington; Gaetano Borriello: *Workshop on Future Research Directions in CAD for Electronic Systems: "Putting the 'D' Back in CAD"*; (MIP-9634523); $21,930; 6 months.

This workshop is on evaluating priorities for academic electronic design automation (EDA) research, with the intention that the results of the meeting aid in refocusing the field on its roots. Both future priorities for research and education, as well as areas that are already well-served by the EDA industry and are unlikely to require major new methodological or technological breakthroughs are considered. The workshop report includes:

1. a statement of the most pressing problems to solve within the next decade,
2. recommendations regarding the style and kind of research to be conducted and suggestions to the community as to how to implement these suggestions, and
3. viewpoints on increasing the impact of the EDA field on education in electrical and computer engineering.

Computing Systems Research
(Formerly Microelectronic Systems Architecture)

Dr. Michael Foster. Acting Program Director
(703) 306-1936 mfoster@note.nsf.gov

Program Description

The Computing Systems Research (CSR) program is concerned with developing a fundamental understanding of computing systems, including their design, architecture, implementation, and evaluation. Computing systems include computer architecture, hardware implementation, supporting system software (e.g., compilers and operating systems) for new architectures, interconnection networks, storage and I/O architectures, and novel computing structures and technologies that hold the promise of radically new computing systems for the next century.

Research in computing systems seeks higher levels of performance and more advanced computational or information processing capabilities by enhancing the structure and organization of computing systems with innovative architectural concepts, and by improving the interaction of their underlying hardware and software components. The basic technological domains involved in advanced computing systems include:

- design and evaluation of instruction set architectures and the organization of central processing units, including data paths, arithmetic units, controller design, and processor pipelines;
- memory systems, including addressing, register files, virtual memory support, multilevel cache memories, high performance frame buffers, and pipelined memory structures;
- computer system interfaces to communications networks and other high speed peripherals, including direct memory access mechanisms;
- multilevel storage structures, including disk architecture, disk caches, and mass storage systems;
- interconnection structures among processors, memories, and input/output channels;
- compiler and operating system support for novel architectures;
- fault tolerant and redundant hardware structures, including fault detection and correction mechanisms, design for testability and reliability;
- high-performance input/output systems, including smart sensors, actuators, and special underlying systems for graphics and visualization; (9) the design of heterogeneous distributed computing systems, including global architecture, node architecture, and supporting software; and
- embedded systems, in which hardware, software, and I/O are integrated into an application. In addition, new generations of computer architectures or special purpose computing systems often demand specialized system software and application library co-development, and such specialized system software are part of computing systems research.

Within these technical domains, the CSR program concentrates on research issues that contribute to extending performance, programmability, scalability, and reliability of computing systems. Typical issues which are addressed include: methodologies for system design that map system specifications to physical realizations; design and analysis of components of computing systems; and benchmarking of performance, programmability, scalability, reliability, composability, and predictability. The CSR program encourages studies on the impact of new hardware and software technologies, as well as the impact of new applications and algorithms on computing systems architecture. The styles of research employed include theoretical studies, simulations, limited proof-of-concept prototyping, and measurement of existing systems. The CSR program emphasis is on the design and evaluation of physically realizable systems. Successfully completed research projects are expected to pursue advanced experimental systems design through the Experimental Systems Program.

Technology Driven Microarchitectures

California Polytechnic State University; Joy S Shetler:*CAREER: Computing Systems Research Education and Research*; (MIP-9624967); $200,000; 48 months.

This project establishes an active research effort in Computing Systems Research that enhances the undergraduate curriculum in Computer Engineering, Computer Science and Electrical Engineering, and introduces students to new computer architectures, new technology and microelectronic systems. The ideas developed will be incorporated into the microprocessor and computer architecture curriculum along with new cooperative learning techniques. The focus of the research component is to devise and test Instruction Level Parallelism (ILP) techniques and mechanisms. Programmable Logic Devises, such as FPGAs, are used to implement rapid system prototyping of custom computer designs and are also used to implement custom components in reconfigurable architectures. Engineers and computer Science students are taught with highly participatory "active learning" methods such as cooperative learning which stimulate student interest and learning. By introducing students to teamwork in the engineering classroom, learning is enhanced, fruitful networking relationships are developed and students are able to adapt easily into other team environments. The rapid- prototyping platform developed for this effort is used for several upper division computer architecture and microprocessor courses and the cooperative learning pedagogy serves a national model.

University of California - Davis; Matthew Farrens: *NYI: Award: High Performance Single Chip VLSI Processors*; (MIP-9257259 A005); $61,665; 12 months.

The research is to investigate various configurations of decoupled architectures and to extend the concept into the field of parallel processing. It is anticipated that, with several decoupled processors communicating via architectural queues called Multiple Instruction Stream Computer (MISC), it will function as a type of dynamic superscalar processor, providing significant performance gains. The MISC architecture uses multiple asynchronous processing elements to separate a program into streams that can be executed in parallel, and integrates a conflict-free message passing system into the lowest level of the processor design to facilitate low latency intra-MISC communication. This approach allows for increased machine parallelism with minimal code expansion, and provides an alternative approach to single instruction stream multi-issue machines such as superscalars and VLIWs. The relationship between optimal processor configuration and transistor count is also being investigated. The goal is to define the design points at which a change to multiple processors becomes advantageous.

University of California - San Diego; Chung-Kuan Cheng: *Research on Circuit Partitioning*; (MIP-9315794 A003); $26,392; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $52,783).

This project is an investigation into a new generation of hierarchical partitioning methods motivated by the need to:

4. cope with high circuit complexity,
5. improve a system's performance under I/O pin count constraints, and
6. control intermodule delay in order to optimize system performance.

The principal investigator proposes to extend previous partitioning research by adopting different circuit models including petri nets, data flows, and state machines, by improving the efficiency and effectiveness of the methods, and by deriving theoretical results on variations of partitioning formulations. He plans to apply partitioning methods to netlist mappings for various hardware emulation machines and to find potential applications of these methods to VLSI design problems

University of California - Santa Cruz; Anujan Varma: *NYI: High-Speed Interconnection and Switching Technologies*; (MIP-9257103 A006); $62,500; 12 months.

This research aims to develop architectures and protocols for highspeed interconnection within computer systems. Typical applications to be considered include interconnection of processor subsystems among themselves for multiprocessor configurations, and the interconnection of processors to I/O subsystems. The focus is on high-speed crossbar switches as the interconnection means. The structure of these switches as well as mechanisms for connection setup, routing, and flow-control across multiple cascaded switches are being studied.

Photonic switch architectures obtained by combining dimensions of switching, such as wavelength and space, are under investigation.

University of Southern California; Massoud Pedram: *NYI: Low Power VLSI Design*; (MIP-9457392 A002); $31,250; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $62,500).

This research investigates modeling and estimation of power consumption as well as techniques for minimizing power at the various levels of design abstraction (layout, logic, register- transfer and behavioral levels). Principles and methods to guide the design of power efficient electronic systems are being explored; and the impact of availability of low-power design techniques on chip, module, and system level designs is being assessed. Topics being investigated include: spatio-temporal power estimation; state assignment for low power; power dissipation in boolean networks; common subexpression extraction; and FPGA synthesis for low power.

George Washington University; Hyeong-Ah Choi: *WDM Broadcast Networks and Massively Parallel Computers*; (MIP-9628468); $50,000; 12 months.

This is a career advancement award to apply theoretical methods from the study of interconnection networks to the study of all-optical broadcast networks. The goal of the research is to develop new scheduling algorithms to improve the performance of broadcast networks that use wavelength division multiplexing. The scheduling algorithms will take into account an arbitrary set of system parameters that change with technology. The research will provide insight into the user of multiple wavelengths and tunable devices in optical networks.

Iowa State University; Sachin Sapatnekar: *CAREER: Gate-Level Performance Optimization of Pipelines and General Sequential Circuits*; (MIP-9502556 A002); $5,000.

This project is concerned with performance optimization of high- speed synchronous digital circuits It addresses the application of timing optimizations to pipelines and general sequential circuits in high-performance systems. The core of the research effort is directed towards the method of retiming at the gate level. Retiming is a sequential logic optimization technique that repositions memory elements, namely, edge-triggered flip-flops and level-triggered latches, within a circuits to optimize the timing behavior of the circuit. The use of retiming optimization methods, chiefly device sizing, is also being investigated.

Under the educational part of this project, the PI is incorporating CAD techniques in the undergraduate program from the sophomore level by the use of CAD tools in teaching. At the graduate level, he is involved in developing courses, the supporting existing courses, and in advising graduate students. He is also exposing undergraduates to his research for a better appreciation of the state of the art, and to motivate them towards higher education. he is also contributing towards outreach activities by focusing towards distance education through teaching courses to sites in the state of Iowa via television, and around the country through the medium of satellite.

University of Michigan - Ann Arbor; Karem A Sakallah, Trevor N Mudge, Edward S Davidson: *Timing Issues in the Design of Digital Systems*; (MIP-9404632 A001); $58,339; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $175,017).

This research is on a timing verification and optimization framework for designing an entire digital system (e. g. a microprocessor). The research builds on a widely used model for synchronous timing analysis and an efficient method for estimating gate and wire delays. The model is being extended to include relevant functional information in order to enhance accuracy. Components of the framework are: design decomposition to isolate critical elements; a path delay calculator; algorithms for finding synchronizer components; clock analysis algorithms; a symbolic sequential timing verification component; a hybrid timing-logic simulator; and design optimizers.

Rutgers, The State University; Miles Murdocca: *Development of a Methodology for Implementing Hardware Dataflow Graphs Using Reconfigurable Optical Interconnects*; (MIP-9224707 A002); $10,000.

In this project a system is being built to translate a dataflow language into hardware and implement the hardware on an optical gate array. The system is comprised of three components: a compiler to translate a dataflow language such as ID into a dataflow graph; a method for converting dataflow graphs into hardware described by a hardware description language, along with a simulator; and a synthesis program for mapping HDL descriptions onto an optical gate array. The optical gate array consists of planes of electronic gates, with the planes

interconnected by free-space optics. An advantage of this kind of gate array is quick reconfiguration of the optical interconnections which allows dataflow graphs to be changed quickly. This project is applicable to the development of advanced integrated engineering cells for manufacturing. Because of noise immunity and lack of electromagnetic interference, optical computation is well-suited to the factory floor. Reconfigurable optical computers of the type under development in this project may allow high-speed computation to be inserted into manufacturing processes in new ways.

State University of New York - Buffalo; Chunming Qiao: ***RIA: Reducing Communication Latency with Path Multiplexing in Optically Interconnected Multiprocessor Systems*****; (MIP-9409864 A001); $10,000.**

This project develops an innovative approach to all-optical TDM communications in multiprocessor systems. The proposed approach is an extension of a previously developed connection paradigm for reducing control latency in reconfigurable optical interconnection networks. Using the proposed approach, a connection is established during a set of tome slots along the path which are dependent on each other, so that no time-slot interchanging is required and simple switching devices can be used. As a result, the overall communication latency as well as the complexity of the network can be reduced. Thus, a near term impact of this research is on the design of high-performance, low-complexity optical interconnection networks suitable for multiprocessor systems. The research also has long-term impact on the advancement of technologies such as Asynchronous Transfer Mode (ATM) and Wavelength-Division Multiplexing (WDM) as the emerging technologies for computer communications.

North Carolina State University - Raleigh; Thomas M Conte;***CAREER: New Paradigms for Instruction-Level Parallel Processors*****; (MIP-9625007); $200,000; 48 months.**

This CAREER proposal addresses new paradigms for exploiting and teaching instruction-level parallelism (ILP). Statically-scheduled microarchitectures achieve high performance by relegating to the compiler the responsibility for detecting and scheduling ILP. Drawbacks to this approach include a lack of object code compatibility between generations, the reliance on program profiles, and the erosion of the individual compilation model. This research will develop new paradigms for ILP that remove the drawbacks by enlisting the OS, the linker, the assembler, the instruction encoding and the microarchitecture. The educational plan will introduce the compiler's role in the undergraduate curriculum as the sophomore and senior levels by developing two courses: "Introduction to Machine Language Programming" will be modified to include an awareness of performance issues and the role of the compiler; and a senior level course will be developed to teach back-end ILP compiler optimization. An important feature of this proposal is the training of graduate students in a fundamentally and technologically significant area of ILP. The research and educational activities will use the TINKER Instruction Set Architecture that has been constructed by the principal investigator and his students.

North Carolina State University - Raleigh; Wentai Liu: ***VLSI Delay Vernier Architectures and Applications*****; (MIP-9531729, A001); $214,300; 37 months.**

This project is concerned with two performance enhancing techniques of VLSI systems that have developed from wave pipelining: delay vernier architectures and non-zero skew design methodologies. Many high performance applications require extremely fine timing resolution or very high data sampling rates. Examples include high speed network transceivers, A/D converters, phase-locked and delay-locked loops, and ATE pin electronics. Normally, expensive circuit technologies and very high frequency clock signals are used to satisfy these performance requirements. However, the proposed delay vernier structures are capable of achieving a very high time resolution using CMOS technologies and relatively slow clock signals. This research will focus on refining and extending delay vernier architectures and design techniques so that their usefulness is maximized. The effectiveness of the architecture in a number of applications will be evaluated. Also, some real-world factors, such as process variation and noise, will be investigated. The non-zero skew design approach allows the VLSI systems designer to take advantage of intentional skew to optimize clock distribution for a particular set of propagation delays, and thereby greatly improve system throughput. A critical necessity for these systems in the ability to create precise clock skews. This research will focus very strongly on the delay vernier technique, since it can provide this required ability. Finally, prototypes systems will be built for several promising applications.

Carnegie-Mellon University; David F Nagle:*CAREER: The Architect's Workload Suite Helping Architects Understand Software Systems' Impact on Future Computer Design*; (MIP-9624700); $200,000; 48 months.

Effective computer architecture requires a carefully selected set of workloads that are representative of the target design's software environment. However, it is becoming increasingly difficult for computer architects to obtain workloads that accurately characterize today's complex software systems. The Architect's Workload Suite (AWS) project seeks to address this problem by developing a collection of commercial, public domain, and next-generation software systems that architects can use as a source of representative workloads. Using on-chip performance counters and custom tools, each workload in the AWS will be accompanied by Workload Characterization Worksheet that summarizes its architectural requirements including memory-system, branch-prediction, instruction-level parallelism, course-grained parallelism, and operating system support. Workload summaries will allow researchers to carefully screen the AWS workloads, selecting a set that is appropriate for their work. In addition to the workload suite, this work is developing several hardware/software analysis tools. The first, based on Trap-driven simulation, characterizes and evaluates memory systems. The second, a modified version of DECs ATOm analysis tool, generates accurate software traces of the entire system's activity, including OS and application workloads. Together, the workloads, their characterization worksheets and the tools will help architects determine how best to support future software systems.

University of Pittsburgh; Rami Melhem: *Time Division Multiplexing of Optical Interconnection Networks*; (MIP-9633729); $187,016; 36 months.

The goal of this research is to study, design and evaluate control mechanisms for time-multiplexing optical interconnection networks in multiprocessor applications. Such mechanisms decouple the relatively slow network control from the high optical data transmission bandwidth. The research will assess the advantages, capabilities and limitations of time-division multiplexing. Specifically, models will be developed to specify the communication requirements of applications, to express and measure the communication locality present in communication loads, and to estimate the communication cost in multiplexed interconnection networks. Different control protocols will be simulated and tested on traces collected from actual parallel applications, and on synthetic traces generated with the help of the communication locality model. Centralized, distributed and pipelined control strategies will be implemented and compared. This project is motivated by several observations. First, without multiplexing, the control overhead of establishing end-to-end connections in a network may be prohibitive. Second, applications exhibit communication locality which amplifies the benefits of multiplexing. Third, the high bandwidth of optical paths allows for multiplexing with minimal effect on the transmission rate on individual connections. Finally, the nature of massively parallel programs allows compilers to determine, with current technology, the multiplexing sequences that satisfy most of the communication requirements.

Application Driven Computing Systems

University of Illinois - Urbana-Champaign; Benjamin W Wah: *Global Optimization and Parallel Processing of Nonlinear Problems in Engineering Applications*; (MIP-9632316); $300,000; 36 months.

The objective of this project is to develop an efficient nonlinear optimization method - NOVEL: Nonlinear Optimization Via External Lead - and the supporting hardware/software to solve large-scale nonlinear optimization problems in operations research, digital signal processing, and neural-network learning. Such a method will allow many important application problems in these areas to be solved better as well as faster than existing methods. The method first transforms nonlinear constrained optimization problems using Lagrange multipliers into unconstrained optimization problems. This transformation allows both constrained and unconstrained problems to be solved in a unified framework. The key idea of the method is to use a user-defined trace to pull a search out of a local optimum (a local saddle point in constrained optimization and a local minimum in unconstrained optimization) without having to restart it from a new starting point. The proposed method has a good balance between global search and local search. In global search, NOVEL relies on two counteracting forces: local gradient information that drives the

search to a local optimum, and a deterministic trace that leads the search out of a local optimum once it gets there. The effect of the trace on the search trajectory is expressed in terms of the distance between the current position of the trajectory and that of the trace. Good starting points identified in the global-search phase are then searched more extensively using pure local searches. The proposed research consists of four aspects: refinement of the NOVEL method, parallel processing, extension to mixed-nonlinear optimization problems, and applications of the methods developed to solve problems in operations research, digital signal processing, and neural-network learning.

Massachusetts Institute of Technology; Kai-Yeung Siu; ***NYI: Analysis and Design of Artificial Neural Networks; (MIP-9696144); $31,250*****; 12 months; (Support from: Signal Processing Systems Program, Computing Systems Research Program - Total Grant $$62,500).**

Artificial neural networks present a new model for massively parallel computation and a promising paradigm for solving large scale optimization problems. This research is exploring the advantages of neural network-based models over conventional models for computation, and a novel design of neuromorphic computing architectures for applications in signal and image processing. A theoretical framework is being established to derive tight tradeoffs between the number of elements and the number of layers in neural networks. The results should answer some of the key open questions in the analysis of neural networks using classical mathematical tools such as rational approximation techniques and harmonic analysis.

Western Michigan University; Xiaobo Hu: ***Architectural Design Exploration for Real-Time Embedded Systems*****; (MIP-9612298); $151,177; 36 months.**

Advances in both hardware and software design have enabled system designers to develop large, complex embedded systems. Such systems may contain a mixture of software programs running on general purpose or special-purpose processors, high-speed memory systems, and application specific hardware devices. One of the key tasks in designing these systems is to find a "best" system architecture including determining the "right" partition between hardware and software components and selecting the "right" hardware components. The focus of this project is to develop a viable methodology as well as a set of tools to assist system designers in this architectural design exploration task. The research proposed performs design exploration at the system level. At this level, a system is specified as a collection of task graphs and an implementation is described by an interconnection of hardware components and an assignment of processes to the hardware components. Though the problem shares some similarities with the hardware synthesis problem, many different issues need to be addressed in order for this approach to be effective. One of the issues is identifying and estimating appropriate parameters to be included in the exploration process such that system timing performance can be adequately predicted. The project emphasizes design exploration for embedded systems with real-time constraints. Such systems can be found in many applications, such as navigation and landing control of aircraft, networks and communications. Attributes that are useful for guiding design exploration will be devised so as to efficiently analyze the timing behavior of a large number of implementations with respect to the system timing specifications and hence permit exploration of a wide range of design options. Other issues including sensitivity analysis and component reuse will also be addressed. To properly handle conflicting performance requirements including both functional and non-functional attributes, an approach based on imprecise value functions from utility theory will be pursued. Furthermore, the design exploration problem is formulated as a multi-attribute optimization problem. Techniques from the evolutionary computation field will be investigated in order to tackle the exponential nature of the optimization problem.

University of Rhode Island; Qing Yang: ***Exploring the Design Space for High Performance and Low Cost Memory Hierarchies*****; (MIP-9505601 A001); $10,000.**

The main objective of this research is to investigate the issures related to hardware designs and performance evaluations of a number of innovative cache design techniques. These design techniques are currently being developed for single-chip general purpose processors and multiprocessors. Specifically, the focus of the research includes: 1) minimization of area cost of on-chip caches by CAT--caching address tags: 2) maximization of cache hit ratios by evenly distributing data across cache sets with the help of a few tag bits that change frequently during program executions; 3) investigation of potential impacts of the new CAT cache designs for various cache configurations as well as multiprocessor caches; and 4) devising analytical models, and performing trace-driven simulations and execution-driven simulations for evaluating implementation costs, performances, and design trade-offs of various designs.

Texas A & M University; Nitin H Vaidya: *Bit/Byte Bounded Error Control Codes for Byte-Organized Systems*; (MIP-9423735 A001); $49,952; 12 months.

The principal goal of this project is to design error control codes for "byte-organized" systems using a new error control model. In a byte-organized system, the hardware is partitioned such that each part produced a subset of hits in the output. For example, a memory may be organized as one byte per card. In this case, each card produces a byte in the memory word. For such systems, a new error control model called the bit/byte bounded model is being developed. Briefly, this model can be described using two parameters, say t and u. The codes designed using the bit/byte bounded model can tolerate up to t bit errors confined to a must u bytes. Unlike the existing approaches, the bit/byte bounded model can interpolate between the traditional bit and 1 byte error models. Also, the proposed model can be used to reduce the number of checkbits by trading a small amount of error control capability.

University of Wisconsin-Madison; Parameswaran Ramanathan: *Guaranteed Performance Communication in Distributed Real-Time Systems*; (MIP-9526761); $70,000; 12 months.

This project is exploring new communication mechanisms that go beyond the deadline guarantees of previous work in real-time and non-real-time message streams, using models of real-time communication that can represent tolerance to deadline misses. For real-time message streams, these models suggest new techniques for providing service guarantees on the frequency or clustering of missed deadlines. for non-real-time message streams, the models will be used to develop support mechanisms that minimize average delivery delay subject to the constraints of the real-time streams. Simulation and analytic methods are being used to evaluate competing techniques.

Scalable Parallel Computing Systems

University of California - Los Angeles; Milos D Ercegovac: *Arithmetic Algorithms and Structures for Low-Power Systems*; (MIP-9314172 A002); $98,395; 12 months.

The objective of the proposed research is the development of numerical computing systems which operate with a small amount of electrical power. It concentrates on arithmetic structures and algorithms and models power consumption by the number of signal transitions. These low-power structures are essential for the development of future computing systems. The proposed research will further the understanding of the effect of number representation, algorithm, and implementation on the power dissipation of numerical computation structures. Moreover, methods and technique for the design of low-power structures will be developed, as well as low-power designs for specific operations, such as adders, comparators, multipliers, and dividers.

University of Delaware; Phillip Christie: *Statistical Analysis of Interconnect-Limited Systems*; (MIP-9414187 A002, A003); $28,750; 12 months; (Support from: Computing Systems Research Program, Design Automation Program - Total Grant $57,499).

The growth in size and complexity of modern computers is matched by the need for more accurate models of their internal wiring structure. The task of modeling this complexity is ideally suited to a recent development in statistical physics called the re-normalization group. It has the power to attack the wiring optimization problem because it takes this one very difficult problem and breaks it up into a very large number of much simpler problems. Rather than attempting to attack the entire range of partitioning and placement requirements from the transistor level to the back-plane, an optimization procedure need only be found for a highly restricted wire length range. Once this has been derived, the procedure is renormalized for all other length scales in a manner which generates a system-wide solution. This project will investigate how such an approach may be used to develop much better estimates of interconnect-limited system performance and how this theory leads to new optimization algorithms. In contrast to other theories the predictions made by this technique will grow ever more accurate as the systems grow larger and more complex.

University of Florida; Jih-Kwon Peir: *CAREER: Decoupling and Reforming Tag and Data Arrays for High-Performance Memory Hierarchy Systems*; (MIP-9624498); $200,000; 48 months.

The project investigates issues in memory hierarchy design and proposes solutions to bridge the increasing performance gap between the processor and memory. By observing the unequal time needed

to perform cache tag access/comparison and to access cache data, this path imbalance can be exploited to achieve a more performance optimal cache design. The basic idea is to use an additional tag array record the status and location of the recently used lines in the cache data array. By recording the cache line locations in this subset of the tag array, the data access is decoupled from the tag access/comparison path for a shorter overall cache access time. In a multiprocessor system, cache coherence activities incur extra traffic to the heavily-loaded snooping bus and impose additional latency in accessing the data. By observing that references to the same memory location are often ordered at the software level, frequently rendering cache coherence activities are unnecessary and can be deferred until at the proper synchronization point. The infrastructure required, including a multiprocessor tracing facility, multiple-issue out-of-order execution processor models, and various memory hierarchy models, will be developed in this project. These tools will be enhanced with graphical user interfaces and used for instructional purposes to provide students with hands-on experience in evaluating difference architectural design tradeoffs.

Georgia Institute of Technology; Umakishore Ramachandran: *The Quest for a Zero Overhead Parallel Machine*; (MIP-9630145); $235,442; 36 months.

The primary goal of this project is to demonstrate the viability of cost- effective parallel computing techniques for two application domains that make computing ubiquitous. The first one is parallel databases embodying a clients/server paradigm used in data mining/data warehousing applications. The second one is the physical simulation of virtual multimedia environments that has potential use in several everyday situations. The challenge in this domain is to perform realistic physical simulation in real time of multiple actors (some real and some synthetic), each with several degrees of freedom. This research will investigate the system architecture issues in the design of parallel (cluster) servers engineered with commodity processing and network components for supporting such applications. A shared address space across the cluster is an attractive abstraction from the perspective of programming such applications. However, accessing and caching data from this shared address space can be expensive in terms of observed latencies in a cluster. Therefore, clever techniques that hide the communication overheads for accessing the shared address space are imperative to achieving good performance. The ultimate goals is to make the cluster server appear as though the applications incur zero overhead for communication. Specific research issues to be addressed in this research include (a) mechanisms for realizing shared address space across the cluster, (b) mechanisms for supporting variable granularity, caching, explicit communication, and synchronization (within the framework of shared space) tailored to the needs of these applications, (c) hardware assists necessary to make clusters appealing as cost-effective parallel machines, and (d) detailed performance evaluations of cluster servers under such multimedia and database workloads.

Purdue University; Jose A Fortes: *Data Distribution Independent Parallel Processing*; (MIP-9500673 A002); $91,253; 12 months.

This project is exploring the possibility of a parallel programming paradigm that is data-distribution independent (DDI) in the sense that the user would not be required to program or even invoke data communication routines (hereon called modules). The need for data redistribution would either be eliminated or transparent to the user. The emphasis this work is on the systematic design of computational modules so that either there is no need to redistribute input data or, when this cannot be achieved, the cost of (automatic) redistributionis minimized. In this context, source-to-source program transformations, called modular mappings, and properties that allow commutative parallel processing are being explored as techniques and concepts that enable DDI computation. The extent to which a DDI paradigm could replace existing approaches, complement them or merely apply to special application domains is unclear and this is one of the issues under investigation. In addition, hardware and architecture features that support DDI computation on both general and special purpose aprallel processors are being investigated. The experimental validation is being done in the context of three application areas and program implementations on a fine-grain distributed memory SIMD machine (Maspar MP-1 with 16 thousand processors) and a coarse- grain distributed memory MIMD machine (Intel Paragon with 140 processors). The areas of interest are dense linear algebra, sparse linear algebra and symbolic compute algebra which can be applied to numerous scientific and engineering computing problems. For these three areas, DDI modules are being developed as well as entire programs built of several modules. Performance comparisons are being made between DDI implementation and their non-DDI counterparts.

Iowa State University; Prasant Mohapatra: *Hardware Multicast Routing Techniques for Scalable Parallel Computers*; (MIP-9628801, A001); $167,486; 36 months.

Multicast communication is an essential feature in parallel computers. It benefits several parallel applications and is useful for shared memory, distributed memory, and distributed-shared memory systems. Multicasting also helps in supporting a class of communication patterns known as collective operations. Multicasting can be supported by either hardware or software approaches. However, the hardware approach is significantly faster than the software approach. Most vendors of parallel computers are either unaware of it or have made unsuccessful attempts at implementing deadlock-free hardware multicasting schemes. This project is concerned with efficient hardware multicast routing techniques for meshes and switch-based parallel computers. Several implementation details, such as start-up latency, router delays, and multiple multicasts, will be addressed in this project. The primary objectives of the research will include the minimization of the total communication latency by reducing the number of start-up steps required for multicasting in meshes. For switched-based systems, this project will examine alternative switch designs to support efficient deadlock-free multicasting. Finally, an effort will be made to design low-cost routers that support multicast operation. Two primary constraints will be observed during this research. First, the cost and complexity of the algorithms will be minimized. Second, the PI will focus on techniques that not only support deadlock-free multicast operations but also enhance the performance of node-to-node unicast communications.

Johns Hopkins University; Robert Cypher: *Techniques for Fault-Tolerant Communication in Parallel Computers*; (MIP-9525887); $153,756; 36 months.

The project studies techniques for performing efficient communication in parallel computers that have faulty components. Parallel computers with a large number of identical processors have a great potential for fault-tolerance. The main challenge that must be overcome in order to exploit this potential for fault- tolerance is the fact that faults also affect communications between processors. Hardware faults in processors, routers, and communication links force messages to take new paths that avoid the faulty components. In addition, the mapping of computations from faulty processors to nonfaulty ones changes the pattern of communication. As a result, faults can have a major impact on the performance, and even correctness, of the communication subsystem in a parallel computer. Specifically, the project focuses on the following two key issues; creating efficient, deadlock-free routing algorithms for parallel computers with faulty components, and developing efficient mappings of communication-intensive real-time applications to parallel computers that contain faults. This research will result in more efficient use of parallel computers and a better understanding of how their inherent potential for fault-tolerance can be utilized.

Boston University; Mark G Karpovsky: *Software Implemented Fault Tolerance in Multicomputers*; (MIP-9630096); $220,003; 36 months.

The main objective of this proposal is to develop an integrated approach for increasing the dependability of parallel computers. For detection and diagnosis of hardware faults at the local (node) level, this approach will combine off-line built-in self testing, on-line concurrent checking, and software techniques implementing path and algorithmic-based checks. The aim of this proposal is to develop a systematic method for generating and placing checks into the implementation of an application. For automating the verification process and placement of checks in programs, an interactive tool, VERIFY, will be developed. For estimating the coverage factor provided by these techniques, a software fault-injection tool will be created. Efficient and fast reconfiguration algorithms will be developed for avoiding single and multiple faulty nodes and links. A user transparent software layer will be created for implementing the proposed algorithms. This will be written on top of the existing system software that residents in every node of parallel computing systems.

The Software Object-Oriented Fault-Injection Tool (SOFIT) will be ported on the nCUBE and Tandem SeverNet. The recovery time overhead will be estimated in the presence of single and multiple faults injected by SOFIT into the nodes and links of the nCUBE and SeverNet when programs are being executed. A dependability evaluation study will be performed. The objective of this research is to show that this integrated approach can be used successfully for improving the reliability of the nCUBE and SeverNet by an order of magnitude.

University of Michigan Ann Arbor; Peter M Chen: ***CAREER: Measuring and Improving Memory's Resistance to Software Corruption; Relational Instruction*****; (MIP-9624869); $200,000; 48 months.**

System designers have long assumed that memory cannot reliably store permanent data because of memory's volatility and vulnerability to software corruption. Memory's unreliability slows system performance by forcing some applications to write data through to disk and hurts reliability for forcing applications that care about performance to throw away 30 seconds of data on a crash. The primary research goals of this CAREER plan are 1) to enable memory to reliably store permanent data even in the presence of operating system or database errors; and 2) to measure how often software faults corrupt memory, both with and without extra protection.

The primary education technique of this CAREER plan is relation instruction: spending time and effort relating with students personally. This plan suggests several simple ways, such as group lunches, to do this effectively with large, undergraduate classes; prior experience attests to how caring for students in a personal manner motivates them to learn.

StateUniversity of New York - Geneseo; Rong Lin: ***Reconfigurable Architectures with Shift Switching for Novel Parallel Arithmetic Schemes*****; (MIP-9630870); $110,370; 36 months.**

This research will investigate a novel VLSI arithmetic design approach using reconfigurable architectures with shift switching, aimed at achieving an innovative design methodology which could lead to a substantial improvement on both speeds and areas for several arithmetic devices including parallel counters, parallel compressors, parallel multipliers, precharged CMOS adders/comparators, prefix parallel counters, as well as elementary function evaluations. The proposed technique will employ enhanced reconfigurable buses, i. e. buses containing shift switches with buffers properly inserted. It possesses the following new features. First, when specified digital signals are propagating through shift switching buses, the desired modulo operations can be performed directly, which simplifies the traditional model of this basic arithmetic computation. Second, after passing through each shift switch, the state signals are inverted alternatively in two mutually inverted forms (n and p), which minimizes the load of transistors and maximizes the speed of circuits. The research will focus on the development of enhanced reconfigurable bus-based VLSI schemes for important arithmetic operations. The goals are to investigate:

1. shift switching parallel compressors and conditional compressors;
2. efficient partial product reduction schemes;
3. application-specific array processors for realization of some optimal arithmetic algorithms;
4. various shift switching parallel counter schemes and their applications; and
5. shift switching with precharge CMOS domino logic and its applications in designing arithmetic devices and asynchronous arithmetic devices.

Pennsylvania State University; Chita R Das: ***Application-Driven Network Performance Evaluation*****; (MIP-9634197); $224,496; 36 months.**

Prior research on multiprocessor interconnection networks has primarily focused on the network topology, switching mechanism, and message routing algorithm to maximize network performance. Very little work has been done to study the impact of the communication properties of parallel applications. Communication pattern, message generation frequency, and message size are the three attributes to quantify any communication. The proposed research is aimed at characterizing these communication workloads and analyzing their impact on multiprocessor performance.

The research has two major phases. In the first phase, traffic profiling of a wide range of parallel applications will be conducted by collecting execution traces from parallel machines and also from an execution-driven simulator. The communication traces will be analyzed to determine different types of traffic patterns, rate of communication, and volume of communication. The second phase of the project will use these realistic traffic properties for the performance analysis of interconnection networks via simulation and analytic models. In-depth simulation of known unicast and collective routing algorithms on various topologies will be performed with these workloads. Analytical models capturing wormhole switching, virtual channel flow control, routing mechanism and the realistic workloads will be developed. The main contribution of the project is characterizing application workloads to be used for different architectural and algorithmic research. For example, evaluation with these workloads will quantify the actual performance advantages of adaptive routing algorithms, will identify the potential bottlenecks of existing communication mechanisms, and will provide insight for developing application-specific routing algorithms. Next, the new mathematical tools will be more accurate in predicting the realistic

communication traffic. The techniques and tools developed in this research can be used in understandi ng and evaluating the interplay between parallel architectures and algorithms to maximize the performance of existing machines and to design better machines in the future.

Pennsylvania State University; Ali R Hurson: *Research in Multithreaded Dataflow and Hybrid Architectures*; (MIP-9622836); $161,997; 36 months.

The overall objective of the research is the investigation of architectural features for cost-effective, mainstream computing based on multithreading. In this vein, this research will focus in two directions; investigation of multithreaded dataflow paradigm, and the adaptation of the results to hybrid dataflow-von Neumann architecture. The dataflow model is ideally suited for multithreaded architecture since it facilitates instruction level context switching and fine-grained parallelism. A clear understanding of the issues in supporting multiple threads in dataflow environment will permit us to adapt them to hybrid dataflow/control flow architectures. Hybrid systems present the most interesting opportunities in the area of multiprocessing - they directly address problems that will be faced by future superscalar processors, such as long memory latencies, context switching overheads, multiple active instruction streams, fast and efficient support for task synchronization.

Initially, this project will investigate the use of cache memories for multihreaded dataflow architecture, and compiler techniques for improving the performance of such machines. This project will then study architectural support for multithreading including support for thread management, thread scheduling, communication and synchronization among threads. Simultaneous multithreading in the context of dataflow and hybrid architectures will then be investigated. Finally, this project will explore the suitability of the most significant architectural features for use with convectional RISC technology.

Texas A & M University; Laxminarayan N Bhuyan: *Cache Coherence in Wormhole Networks*; (MIP-9622740); $85,960; 12 months.

This project is investigating the use of wormhole networks in supporting cache- coherent multiprocessors. Cache coherence gtraffic, characterized by fixed length messages that arrive in bursts, is generated by an execution-driven simulator. This traffic is used to evaluate alternative message scheduling schemes, adaptive routing schemes, virtual channels, and interface designs. New coherence and synchronization schemes are being developed and evaluated as part of this research.

The project is divided into three phases. In the first phase the effect of congestion in the interconnection network in performane of cache-coherent multiprocessors is being investigated. The second phase emphasizes the development of techniques to reduce cache coherence traffic and invalidation time. The third phase is the development of new protocols that use wormhole networks effectively.

University of Texas - Arlington; Krishna M Kavi, Behrooz Shirazi: *Research in Multithreaded Dataflow and Hybrid Architectures*; (MIP-9622593, A001); $120,496; 12 months.

The overall objective of the research is the investigation of architectural features for cost-effective, mainstream computing based on multithreading. In this vein, this research will focus in two directions; investigation of multithreaded dataflow paradigm, and the adaptation of the results to hybrid dataflow-von Neumann architecture. The dataflow model is ideally suited for multithreaded architecture since it facilitates instruction level context switching and fine-grained parallelism. A clear understanding of the issues in supporting multiple threads in dataflow environment will permit us to adapt them to hybrid dataflow/control flow architectures. Hybrid systems present the most interesting opportunities in the area of multiprocessing - they directly address problems that will be faced by future superscalar processors, such as long memory latencies, context switching overheads, multiple active instruction streams, fast and efficient support for task synchronization.

Initially, this project will investigate the use of cache memories for multihreaded dataflow architecture, and compiler techniques for improving the performance of such machines. This project will then study architectural support for multithreading including support for thread management, thread scheduling, communication and synchronization among threads. Simultaneous multithreading in the context of dataflow and hybrid architectures will then be investigated. Finally, this project will explore the suitability of the most significant architectural features for use with convectional RISC technology.

University of Washington; Arun K Somani: *Extent and Effects of Error Propagation and Recovery Mechanisms in Cache Memory Systems*; (MIP-9630058); $180,000; 36 months.

Commercial microprocessor systems are being operated at higher and higher clock rates. Faster clocks impact the time that is available to fetch data from cache memory, and increases the probability of transient error occurrence in cache memory systems. Besides data reading error, an error may also occur in the processor subsystem, and it may write erroneous data into the cache memory. Error-correcting codes allow detection and recovery from some of these errors within the code word limits. Fast error detection allows damage containment and reduces the recovery time, which otherwise could be very expensive in time. The goals of this project are: (1) to study the extent of error propagation due to transient faults in computer system when a fault originates either in a processor register or a cache location; and (2) to develop techniques and hardware support needed for early detection and recovery from such errors in computation tasks with low overhead and low performance loss. The proposed techniques will cover a broad spectrum from only detection to full error recovery. This research will identify architectural features which, when provided in commercial microprocessors, will make them suitable for use in fault-tolerant applications. The intent is to keep the performance impact in the normal operation to a minimal.

Other

Princeton University; Douglas W Clark: *Workshop on Critical Issues in Computer Architecture Research, May 21, 1996, Philadelphia, Pennsylvania*; (MIP-9634380); $19,835; 6 months.

The last decade has seen extraordinary technological, intellectual, and market developments that have shifted many assumptions underlying research in computer architecture. Promising avenues for future research will doubtless lead in directions different from those of the past. This project will convene a group of approximately 15 experts from industry and academe, who will discuss this and related issues in a one-day workshop. The proposed date is May 21, 1996, the day before the start of the ACM/IEEE international Computer Architecture Symposium to be held in Philadelphia, PA. The product of the workshop will be a final report published on the World Wide Web.

Signal Processing Systems
(Formerly Circuits and Signal Processing)

Dr. John H. Cozzens, Program Director
(703) 306-1936 jcozzens@nsf.gov
www.cise.nsf.gov/mips/CSPhome.html

The Program

The Signal Processing Systems (SPS) program supports basic research in the areas of digital signal processing, analog signal processing, and supporting hardware and software systems. This research is typically driven by important applications and emerging technologies. Signal processing is a highly active research area, with topics ranging from theory to Very Large Scale Integrated (VLSI) circuit implementations and applications. Although signal processing is typically driven by advances in technology, discoveries in this field also serve as a catalyst for new technological innovations. A taxonomy of research areas, based on signal characteristics, applications, and/or technology, include:

One-Dimensional Digital Signal Processing (1-D DSP) - the representation of time-varying signals (e.g., audio, EKG, etc.) in digital form, and the processing of such signals;

- (adaptive) filtering and equalization
- filter design, theory, and analysis, both linear and nonlinear
- multirate processing and wavelets
- time-frequency representations

Statistical Signal and Array Processing (SSAP) - the use of statistical techniques for the processing of signals that may arise from multiple sources;

- cyclostationary signal processing
- higher order statistics
- (statistical) array processing
- nonstationary and time-frequency

Image and Multi-Dimensional Digital Signal Processing (IMDSP) - the acquisition, manipulation, and display of multi-dimensional data using digital technology;

- image analysis, filtering, restoration, and enhancement
- image and video coding
- vector quantization

Analog Signal Processing (ASP) - the processing of data without conversion to sampled-digital form;

- analog-to-digital conversion
- analog circuits and filters

Other research areas delineated in the SP EDICS listing that are selectively supported by this program are:

- Underwater Acoustics Signal Processing -
- VLSI Signal Processing -
- Emerging Techniques -

Analog (Mixed Analog/Digital) Signal Processing

University of California - Berkeley; Paul R Gray, Robert G Meyer: *Research in High-Frequency Analog Electronic Circuits for Communication Systems*; (MIP-9412940 A001); $136,403; 12 months.

The field of electronic telecommunications has experienced explosive growth in the past ten years, both in terms of its commercial importance and in terms of the research effort devoted to it in academic and industrial laboratories. It appears that radio-based digital communications will play a much larger role than might have been anticipated, complementing data communications on wire pairs with fiber media which are coming into increased use for data communications backbones. As in the past, reducing the cost of electronics associated with VLSI technology will be critical.

This research is directed at exploring new ways to use silicon integrated circuit technology to improve the performance, reduce the power dissipation, and reduce the cost of components for communication systems of various kinds. The primary emphasis of this research is on data communications and Personal Communications Systems (PCS) using radio and free-space optical media. Particular emphasis is placed on circuit techniques applicable to implementation of RF, IF and baseband mixed signal communications circuits with low power and low operating voltage. Specific topics for investigation are: limits-to-phase noise performance in power-optimized monolithic voltage-controlled oscillators; investigation of new parallel-architecture sampling demodulators; limits-to-distortion performance in MOS sampling demodulators; the use of passive-sampling FIR switched capacitor filters for low-power, high-frequency filtering and CDMA despreading; optimization of the sensitivity in fiber optic communication receivers; the application of BiCMOS technology for the particular needs of high-speed communications; and modeling of active and passive IC components for RF design.

Harvard University; Woodward Yang: *NYI: VLSI Design for High Performance Signal Processing and Computation*; (MIP-9257964 A004); $62,500; 12 months.

The focus of this research is on the development of innovative VLSI design methodologies that will facilitate the next generation of high performance signal processing and computing systems. A variety of analog, digital, and mixed signal circuitry will be implemented for use in high performance computing applications including object recognition, smart sensors, adaptive neural networks, and programmable analog filters.

Northeastern University; Lisa McIlrath: *CAREER: Development of a New Program in Integrated Analog Systems Design at Northeastern University*; (MIP-9623907); $210,000; 48 months.

The focus of this Career Development plan is the development of a new broad-based program in the design of integrated analog systems at Northeastern University. The research component of this plan involves the development of analog processing circuits to be used in integrated circuits containing arrays of sensors or display elements. The main focus of this research is on applications related to electronic imaging. The two main thrusts are;

1. to develop high quality imaging arrays with CMOS technology, and,
2. to devise focal-plane computational imagers for extracting useful information from the image data.

The functions performed by the focal plane processors, which include temporal and spatial differencing and spatial filtering, will have applications in fields ranging from image compression and enhancement to machine vision problems such as tracking and inspection.

The educational component of this plan addresses the need to interest and train more engineers in the field of analog design. The major initiatives of this plan include developing an extended analog design project for second and third year undergraduates; incorporating more undergraduates into on-going research; creating a seminar series bringing outside speakers on campus

to increase students' awareness of the importance of analog design in real world systems; and developing an advanced graduate course in analog systems aimed both at our graduate research students as well as at engineers in industry who need to deepen their understanding of design principles and techniques.

Washington State University; Terri S Fiez: *NYI: High Performance Analog Signal Processing for Mixed-Mode Integrated Circuits*; (MIP-9257112 A003); $125,000; 24 months.

The focus of this research is the development of high performance analog architectures and circuits for mixed analog-digital IC's. Emphasis is placed on improving the immunity of the analog circuits to digital switching noise and in obtaining accurate, high speed analog circuit performance with the newly emerging 3.3 Volt power supply standard. (Current-mode circuits will be examined as a way of overcoming these limitations and this research will yield a general understanding of the advantages and limitations of current-mode circuits as compared to voltage-mode circuits. The current-mode circuits will be used to implement area efficient, high-order sigma-delta A/D converters.)

One-Dimensional Digital Signal Processing

University of California - San Diego; James R Bunch: *Stability of Algorithms in Signal Processing*; (MIP-9625482); $131,830; 36 months.

The method of Least Squares has been employed successfully to solve problems in many scientific disciplines, including signal processing. The class of fast Recursive Least Squares Algorithms was introduced in order to address certain desirable performance traits, such as computational complexity, tracking, and misadjustment. Although these algorithms have evolved over many years and are widely used in signal processing, a common approach addressing the issue of algorithmic stability has been lacking. This research provides an analysis that fills this critical void in the research and development of these algorithms. Descriptive definitions of stability are being examined which are consistent with current and previous algorithmic analyses. The effects of the forgetting factor - the extent to which prior information is weighted - with respect to matrix perturbation theory and the control and propagation of arithmetic roundoff errors are being studied. This research is also providing useful and insightful gains in the important areas of filter misadjustment, filter tracking, the conditioning of the auto-correlation matrix, and the usefulness of the filter output.

This research program joins the science of Numerical Analysis with the science of Signal Processing, providing tools by which algorithms for signal processing can be better designed and used, so that they will give more accurate results.

University of Delaware; Gonzalo R Arce, Neal C Gallagher: *Weighted Myriad Filters and Their Applications in Communications*; (MIP-9530923); $105,013; 12 months.

Linear Filtering theory has been largely motivated by the characteristics of Gaussian random processes. In the same manner, research in myriad filtering theory is motivated by the statistical properties of alpha-stable processes which describe an important class of processes, impulsive in nature, that obey a Generalized Central Limit Theorem; thus these processes can arise in practice as a result of physical principles. The foundation of myriad filtering algorithms lies in the definition of the sample myriad as the location estimate of a class of alpha-stable distributions. The main goals of this research are;

1. To study the theoretical properties of the sample myriad and to exploit its geometrical structure for the development of efficient computational algorithms,
2. To solve the optimal "weighted myriad" filtering problem - the problem analogous to the design of the optimal FIR (Wiener) filter and the optimal weighted median filter,
3. To develop adaptive filtering algorithms for the simple design of weighted myriad filters for applications where the statistical characteristics of the underlying signals may be unknown or varying, and,
4. To extend the filter formulation to the case where the signals are complex or multivariate, providing us with the tools needed to develop robust CDMA multiple access detectors.

This work can have a significant impact on applications requiring robust estimation and filtering. This is particularly the case in mobile and personal

communication systems, to be deployed in the near future, where the underlying statistics of the noise and interferences closely follow the models used in this research rather that the traditional models used in practice today. The results of this research can also be applied to a wide range of problems including remote sensing imaging of the environment and non-destructive evaluation of materials.

Ohio State University; Peter Clarkson: ***Robust Parameter Estimation in the Detection of Cardiac Arrhythmias*****; (MIP-9220769 A002); $79,637; 12 months.**

This research is concerned with improving the efficiency and robustness of parameter estimation schemes employed by implantable and portable devices for cardioversion and defibrillation. The goal of the research is to replace the ad-hoc estimation schemes used in present generation devices by robust estimators that produce optimal minimum variance estimates for a wide range of unknown error distributions, and which maintain close-to-optimal performance in the presence of parameter fluctuations. These objectives are pursued through estimation schemes that are based on the application of Order Statistic (OS) operations to the data.

Carnegie-Mellon University; Virginia L Stonick: ***PYI: Practical Approaches to Optimal Adaptive Filtering of Real Time Non-Linear Systems*****; (MIP-9157221 A008); $47,500.**

This research addresses the use of numerical optimization methods to develop real-time adaptive filters for estimating, identifying, or predicting time-varying and potentially nonlinear processes. The first phase of this work will be devoted to the development, analysis, and simulation of an optimal adaptive infinite-impulse response (IIA) filtering algorithm for telecommunications using homotopy continuation methods to perform the necessary nonlinear optimization. This research will increase our understanding of IIA filter structures in time-varying environments, and will ultimately lead to their more widespread use.

Drexel University; Athina P Petropulu: ***PFF: Signal Reconstruction and Applications to Communications, Ultrasound Image Processing, and Earthquake Engineering*****; (MIP-9553227, A001); $105,000; 12 months.**

The goal of the project is twofold: research in the area of signal reconstruction, and teaching in signal processing, both components emphasizing the interdisciplinary aspect of signal processing. The main thrust of the research is the reconstruction of a signal with unknown statistical description that propagates through a system with unknown characteristics, solely from the system's output (blind deconvolution). Higher Order Statistics are powerful mathematical tools being employed in this problem, capable of revealing a wealth of information about the signal/process (e.g., phase, presence of nonlinearities), and at the same time suppressing noise. The theoretical developments of this work will be applied in the areas of digital communications, medical diagnostics and earthquake engineering, and are expected to bring new knowledge to these fields. In the area of teaching the goal is to develop and maintain a new systems curriculum, highlighting principles that unite different engineering disciplines. Towards this goal, a state-of-the-art multimedia signal processing laboratory will be developed, and new courses will be designed.

University of Pittsburgh; Patrick Loughlin; ***CAREER: Joint Density-Based Methods of Applied Nonstationary Signal Processing*****; (MIP-9624089); $205,000; 48 months**.

While the Fourier transform is an indispensable tool in signal processing and Linear Time-Invariant (LTI) systems analysis, as well as a staple of the electrical engineering curriculum, many natural and man-made processes are not time-invariant but rather exhibit frequencies that change over time (e.g., FM communication systems, the Doppler effect, speech and other biomedical signals). The conventional short-time, or "quasi-stationary," extensions of LTI concepts for studying such processes is often inadequate; for example, depending upon the short-time analysis interval length selected, dramatically different results can be obtained for the same process, particularly when the signal contains both transients and harmonics. Accordingly, there is a need, both in engineering practice and pedagogy, for the development of new methods for time-varying (or nonstationary) signal processing. This research involves the development and application of a general method of nonstationary signal processing that surmounts the limitations of methods based on the extension of LTI concepts to time-varying situations. The principle objectives of this research are to:

1. Develop new joint density-based methods for nonstationary signal processing (e.g., scale);
2. Apply these new methods to challenging, practical problems in biomedical signal analysis and machine health monitoring (which fall under the biotechnology and manufacturing Federal Strategic Areas); and,
3. Develop a general educational framework for nonstationary signal processing based on the new methods, via new instructional laboratory courses and research opportunities for college undergraduate and graduate students.

William Marsh Rice University; Richard G Baraniuk: ***NYI: Signal Analysis and Processing in Matched Coordinate Systems*****; (MIP-9457438 A002); $62,500; 12 months.**

This research aims to extend current methods of time-frequency and time-scale analysis by developing a general theory for signal analysis and processing in alternative coordinate systems. Specific tools under investigation include optimal, signal- dependent time-scale representations, information measures for time-frequency and time-scale analysis, and operator-based, generalized coordinate systems. Test signals are being drawn from problems in machine health monitoring, magnetic resonance imaging, and dispersive signal processing.

Image and Multidimensional Digital Signal Processing

Stanford University; Robert M Gray, Richard A Olshen: ***Tree-structured Image Compression and Classification*****; (MIP-9311190 A004); $80,054; 12 months.**

Tree-structured vector quantization is an approach to image compression that applies ideas from statistical clustering algorithms and tree-structured classification and regression algorithms to produce compression codes that trade off bit rate and average distortion in a near optimal fashion. This research is examining the explicit combination of these two forms of signal processing, compression and classification, into single tree-structured algorithms that permit a trade off between traditional distortion measures, such as squared error, with measures of classification accuracy such as Bayes risk. The intent is to produce codes with implicit classification information, that is, for which the stored or communicated compressed image incorporates classification information without further signal processing. Such systems can provide direct low level classification or provide an efficient front end to more sophisticated full-frame recognition algorithms. Vector quanitization algorithms for relatively large block sizes are also being developed with an emphasis on multiresolution compression algorithms. In order to improve the promising performance found in preliminary studies or combined compression and classification, it will be necessary to use larger block sizes or, equivalently, more context. Multiresolution or hierarchial quantizers provide a simple and effective means of accomplishing this. Other related issues are being explored, including improved prediction methods for predictive vector quantization and image sequence coding.

University of California - San Diego; Pamela C Cosman: ***CAREER: Multiresolution Region-Based Image Compression*****; (MIP-9624729); $222,500; 48 months.**

Image Compression reduces the number of bits required to represent an image, and thereby allows image and video information to be stored more efficiently, or to be transmitted much faster from one place to another. Some compression methods degrade the quality of the image. The research portion of this CAREER plan involves designing compression methods which examine the content of an image in order to control or eliminate the degradation for some portions of the image. The intent is to increase the visual quality or scientific utility of the images. This work also involves the design of multiresolution methods for region-based image and video compression, and the development of novel methods for quality evaluation for such codes. Dividing an image into regions can allow a compression algorithm to employ different encoding strategies in different portions of the image, and can allow some regions to be represented at higher quality or finer resolution than others. Methods for segmentation, bit allocation, and code selection at different levels of resolution will be examined in the course of our investigations.

The main objects of the education plan are to promote interactive and hands-on learning, to be highly accessible and supportive to students whom I advise, and to facilitate interdisciplinary as well as academic/industrial discourse.

Georgia Institute of Technology; Ronald W Schafer, S. J McGrath, Mark A Clements, James H McClellan, Thomas P Barnwell: *Infrastructure and Research Program for Signal Processing in Multimedia Systems*; (MIP-9205853 A006); $171,177; 12 months; (Support from: Signal Processing Systems Program, CISE Institutional Infrastructure - Total Grant $295,895).

The objective of this proposal is to establish an infrastructure for signal processing in multimedia systems. Equipment will be acquired, installed, and integrated into an existing research environment consisting of networked UNIX workstations and minicomputers. This equipment will provide the capabilities to acquire multimedia signals, including image, video, speech, audio, text, and fax; store and access these signals on high-speed, high-capacity disk storage subsystems; process them using custom high-speed, real-time, DSP microcomputer-based processors, hosted by UNIX workstations; communicate a variety of these signals between workstations over a high-speed fiber network; and output the processed signals at their intended destinations. Special-purpose systems - an HDTV workstation, and a custom DSP multiprocessor - will be integrated into the infrastructure to allow for real-time processing of high-resolution video sequences and images. The goal of this infrastructure is to build an environment in which the DSP algorithms, which will play a major role in the growth of multimedia technology, may be developed in such a fashion as to lead directly to real-time implementations of those algorithms. The companion part to this infrastructure is a program of research based on this environment, to be conducted in areas with immediate and urgent application to the problems facing multimedia information systems. This includes real-time video encoding, concentrating on the problems of inter- and intra-frame redundancy elimination. Advances in this area will lead to affordable applications in a host of areas such as video mail, video messaging, real-time video communications, and interactive learning and presentations.

University of Illinois at Urbana -Champaign; Yoram Bresler: *PYI: Statistical Techniques in Inverse Problems*; (MIP-9157377 A004, A005); $100,000; 12 months.

This research falls into four broad areas: image reconstruction, reconstruction of time-varying distributions, sensor array processing, and visualization of multiparameter data. In the area of image processing, the principal objective is to develop the theory and associated computational algorithms for superresolution image reconstruction from partial and noisy data, using statistical models. For the second area, the goal is to develop optimum signal acquisition schemes subject to physical or economic constraints, and the associated efficient reconstruction algorithms, for imaging spatial data that is time varying during the acquisition process. In the area of sensor area processing, several issues are being addressed including the design of computationally efficient algorithms for the (sub)optimal solutions of model fitting problems, wideband source location, and imaging with sensor arrays. Finally, in the last area, the goal is to address the effective fusion, display, and visualization of multi-parameter spatially-related data, such as is acquired in multispectral, or multi-modality remote sensing and diagnostic imaging.

University of Illinois - Urbana-Champaign; Michael T Orchard: *NYI: Optimal Motion Compensation for Video Compression*; (MIP-9357823 A003); $62,500; 12 months.

Efficient compression of video sequences should exploit the high interframe redundancy due to the smoothness of motion fields in typical scenes. Current video coding algorithms use very simplistic motion models, limiting the degree to which motion-induced redundancy can be exploited in video coding. This research is investigating improved methods for estimating motion in video sequences, and compensating for that motion to achieve increased coding efficiency. Methods are being considered which estimate motion at the encoder, requiring transmission of motion overhead, and those which estimate motion directly at the decoder. The goal is to provide a unifyied framework for these two approaches to motion estimation, and to develop hybrid algorithms taking advantage of the best features of both approaches.

University of Notre Dame; Ken D Sauer: ***Model Based Tomography: A Comprehensive Approach to Iterative Image Reconstruction*****; (MIP-9300560 A002); $70,201; 12 months.**

Low dosage transmission medical imaging, emission medical imaging, and nondestructive testing of materials are all examples of tomographic reconstruction problems which can benefit greatly from improved reconstruction techniques. This research introduces a strategy for the development of computationally efficient reconstruction algorithms which directly search for the statistically optimal fit to measured data. This approach exploits the wide availability of digital computation and storage to substantially improve reconstruction in demanding applications. The research has three essential components;

1. a general measurement system model which characterizes the physical measurement apparatus for both the transmission and emission problems, and can be extended to include nonlinear effects such as scattering,
2. a new class of computationally tractable image cross section models which preserves image detail while suppressing noise artifacts, and,
3. A fast numerical algorithm, known as Gauss-Seidel, which serves as a basis for efficient and accurate multiscale reconstruction methods.

Algorithmic techniques are being evaluated using data collected from the Nondestructive Evaluation Section at the Lawrence Livermore National Laboratories.

Johns Hopkins University; Jerry L Prince: ***PFF: Three-Dimensional Image Processing*****; (MIP-9350336 A002); $100,000; 12 months.**

This program has two major components: research and teaching in image processing, both having a focus on three-dimensional data. The research component is divided into two major thrusts:

1. the development of new methods for estimating motion in three dimensions from three-dimensional data sets. The first aim involves optimizing data acquisition methods given a prior stochastic description of the motion and assuming the use of a certain optical flow method. The second aim explores the theory of three-dimensional vector tomography and its implementation using a magnetic resonance scanner.
2. The estimation of shape from three- dimensional data sets. This topic focuses on active surface methods: the definition of new active surfaces, the algorithms and their convergence properties, and the use of active surface methods for 3-D image registration.

The teaching component of this grant has two main objectives:

1. to develop a state-of-the-art signal and image processing laboratory, and,
2. to develop new courses and textbooks on image reconstruction.

The overall goal is to provide a comprehensive educational program at the cutting edge of signal and image processing

Northeastern University; Eric Miller; ***CAREER: An Integrated Approach to the Study of Inverse Methods in Electrical Engineering;*** **(MIP-9623721); $204,092; 48 months.**

The objective of this CAREER award is to develop a fully integrated career plan designed to meet the educational needs of students in the Electrical and Computer Engineering (ECE) Department of Northeastern University. The underlying theme behind all aspects of this work is the study of inverse problems, a topic which is all but absent at the undergraduate level and not often considered even in the graduate school despite the ubiquity of these problems in many fields of engineering and the applied sciences. In the area of educations activities, two projects are envisioned. At the undergraduate level, a capstone design course in the general area of signal processing is being developed and implemented. At the graduate level, a course devoted to the study of the theoretical and practical issues surrounding linear inverse problems is being developed. In the area of research, we are exploring the use of statistical and multiscale signal processing methods for solving non-linear electromagnetic inverse scattering problems. These algorithms will incorporate scattering models based upon a Green's function formulation of the vector Helholtz equation. A statistically-based, multi-grid-type algorithm for extracting image and target information (number of targets, locations, and classifications) is being developed. The incorporation of detailed, physical models directly into the data processing algorithms makes these techniques directly applicable to a wide range of problems including radar signal processing for image formation and target detection/ classification, remote sensing, microwave medical imaging, and geophysical exploration.

Medico-Technological Research Institute; John J Wild: *Detection of Differential Sonic Energy Reflection by Soft Tissues*; (MIP-9634543); $49,950.

The objective of this work is to provide a capability of obtaining a significantly increased resolution for abnormal foci (tumors and knots) in cellular biological tissues by ultrasonic detection. An ultrasonic transducer array with special electroding is being designed, tested and its operation evaluated. This work will look at the electrode patterns and the measurement of the actual beam characteristics using a biological phantom. Electronic circuitry is being developed for increased sensitivity in data acquisition and detection.

Stevens Institute of Technology; Alan L Stewart, Roger S Pinkham: *Self-Adjoint Operators and Models of Space-Variant Visual Acuity*; (MIP-9405081 A003); $78,524; 12 months.

Current theories of visual processing are theories of local responses. Even models of space-variant acuity concentrate on properties of the local receptive fields, with the rationalization that the visual system is approximately homogeneous within any small neighborhood. The response of the entire visual field is pieced together from models of local responses.

The theory of integral operators allows the experimentla study of human acuity to be united with computational model is visual processing. Their use in this context leads to simpler and more intuitive proofs of key theorems which relate threshold assessment to eigenvalue problems. At the same time, everyday experimental concepts, such as threshold sensitivity, take on new elegance when placed within the theory of integral operators.

Rensselaer Polytechnic Institute; William A Pearlman: *New Techniques for Audio, Image, and Video Compression*; (NCR-9523767); $184,546; 36 months; (Support from: Signal Processing Systems Program, Communications Research, Network Infrastructure - Total Grant $369,092).

The goal of this research is to investigate more effective techniques for digital audio, image, and video data compression. For audio, the intention is to employ more efficient entropy-constrained, conditional and unconditional, quantization in the coding of the critical subbands of the MPEG audio standard. The dependence of the subbands is next exploited using set partitioning in hierarchical trees (SPIHT) as done originally in embedded, wavelet, zerotree coding of images. The resulting embedded coding allows delivery of different qualities of service with a single compressed bit stream. A newly developed SPIHT procedure, which obtained probably the best-to-date monochrome, still image results (lossy or lossless) and with extremely fast encoding and decoding, is utilized for audio and also for color image and video coding. The extremely fast execution of the SPIHT coding procedure makes it especially suitable for video, where a requirement of real-time coding becomes an insurmountable obstacle for other procedures. Both lossy and lossless coding are investigated, as they fit naturally into the SPIHT framework. Other investigations involve the use of information theoretic criteria to create optimal adaptive subband compositions and the use of trellis coded quantization (TCQ) in block discrete cosine transform (DCT) and lapped orthogonal transform (LOT) coding.

Rensselaer Polytechnic Institute; Badrinath Roysam; *Real-Time Processing and Multi-Spectral Imaging Equipment for Intraocular Image Processing;* (MIP-9634206); $21,521; 6 months.

This research involves two components. The first component is an extension of prior work under NSF grant MIP-9412500 on the development of algorithms for automatic generation of retinal vasculature maps, and real-time feature-based determination of retinal position from intra-ocular digital video image sequences. Specifically, a TMS320C80 video signal processor chip will be used to achieve fully transparent image tracking by reducing the time required to establish a tracking lock to a small number of video frame intervals. The second component to this research is to modify the fundus imaging system acquired under the previous award to achieve simultaneous imaging in the near infrared and visible bands using a pair of low-light video cameras. This instrumentation will be used to investigate real-time algorithms for artificially superimposing computed vasculature landmarks on late-phase Indocyanine Green angiograms. This research will result in the development of real-time image processing hardware and software technology for enabling the development of novel instrumentation for laser retinal surgery and related applications. It is motivated by the need to reduce the currently high failure rate (approaching 50%) of current procedures. This project is an application of high-performance computing technology to enable new advances in ophthalmology and related problems in biology and medicine. In addition to reducing the failure rate, this technology is expected to reduce the cost of laser retinal surgery by eliminating the need for repeated corrective procedures.

Rensselaer Polytechnic Institute; John W Woods: *Motion Compensated Object Based Video Coding*; (MIP-9528312); $119,126; 12 months.

This research program investigates an advanced video compression approach for applications such as multimedia, video on networks, video conferencing, and video on demand. It is a novel combination of object-based coding usually used for very low rates and 3-D subband/wavelet waveform-based coding, usually reserved for higher rates.

In this method, the motion field is first segmented into objects which make up the image frame, and move at different velocities. The boundaries of these objects are coded using a type of contour coding which preserves the shape of the edge. This contour coding is performed as an update on the object boundaries found for the previous frame. Finally the interior of these space-time objects is coded using the waveform coding method based on motion compensated 3-D subband filtering. The overall goal is an efficient unified set of video compression algorithms scalable from video phone up to HDTV.

Rochester Institute of Technology; Edward R Dougherty: *Model-Based Design of Optimal Nonlinear Filters for Binary Images*; (MIP-9520139); $55,240; 12 months.

All translation-invariant binary image filters possess morphological representations. Hence, finding an optimal filter involves finding the structuring elements that yield a filter having minimum error - in our case, mean-absolute error (MAE). This observation leads to automated design procedures for restoration and enhancement filters when treating both the observed (corrupted) and ideal images as random set processes. A key design aspect is postulation of image models for the ideal and observed processes. This research concerns filter robustness relative to modeling assumptions: to what degree is filter performance degraded relative to deviations from the design assumptions? Our approach is to assume parameterized ideal and observed image processes and to consider the effects of applying an optimal filter for a given parameter vector to images corresponding to different parameter vectors. The design of enhancement algorithms is an integral part of the manufacture of electronic office equipment. This time-consuming (up to a year), and performance is limited by the ability of a human being to consider only a relatively small number of possible algorithms. Recently, enhancement algorithms have been automatically designed by higher-level algorithms that use statistics in conjunction with mathematical models of images and imaging systems, thereby significantly cutting design time while at the same time producing better enhancement. The present research investigates design from various image models. It also addresses the "robustness" question: since a designed enhancement algorithm is based on modeling assumptions, practical application requires that we have a reasonable idea of how it will perform on image data that differs to varying degrees from the data on which it has been designed.

University of North Dakota; Richard R Schultz: *CAREER: Integration of Image Processing Education and Research*; (MIP-9624849); $225,690; 48 months.

The research component of this program has two major thrusts:

1. an investigation of the mathematical and computational aspects of the multiframe, bilinear image- recovery problem; and
2. the development and application of solutions for problems in the areas of high-resolution transmission electron microscopy, and ground- and space-based astronomy.

Theoretical advances in this program are providing a unified approach to many imaging problems in which intensity data are recorded, including those for which satisfactory solutions have not been previously discovered. Collaborations with materials scientists on problems such as the study of atomic scale processes in solid batteries are providing exciting applications for which new methods developed in this program are being used to enhance the resolving and information-gathering powers of electron microscopes.

Through the teaching component, an educational program is being established with emphasis on:

1. the development of a multidisciplinary laboratory in which faculty and students from many departments jointly explore problems in remote sensing and imaging;
2. the development of multidisciplinary courses in remote sensing and imaging in which students from many departments are included, and the educations goals for each student - ranging from theoretical foundations to specific applications - are tailored to his or her individual needs; and
3. the use of a strong research-teaching exchange in which cutting-edge technology is used in the classroom and laboratory as a tool to excite, motivate, and retain enthusiasm in the students

University of Pennsylvania; Eero Simoncelli; *CAREER: Visual Information Processing*; (MIP- 9624855); $210,000; 48 months.

This CAREER award provides support for the construction of a research and education program in visual information processing. This field is inherently interdisciplinary, and includes portions of computer science, electrical engineering, applied mathematics, psychology, and neuroscience. Yet, educational curricula often do not cross these disciplinary boundaries. One objective of this research is the construction of an educational curriculum, starting next year with the introduction of two innovative foundational courses. The first, an undergraduate lecture course entitled "Foundations of Vision: Perception and Computation," will cover a broad range of material in vision science, from perceptual psychology, to single-cell physiology, to computational theories of visual processing. The second, a graduate lecture/laboratory course entitled "Fundamentals of Image Processing and Computer Vision," will cover basic tools, both theoretical and applied, that are common to these two fields. The second objective of this research is the study of a set of basic and applied research issues in visual image processing. The research will be based on the family of "steerable pyramid" image representations, and will include the use of this representation in an adaptive context, rotation-invariant pattern matching, noise removal and enhancement, representation of texture and orientation content, and differential recovery of three-dimensional scene structure.

University of Washington; Eve A Riskin: *NYI: Vector Quantization Codebook Processing and Organization*; (MIP-9257587 A004); $62,500; 12 months.

New ways to use Vector Quantization (VQ) other than strictly for data compression are being investigated, and are being applied to applications such as image processing, halftoning, progressive transmission, and immunity against communication channel noise. In many applications, both VQ and many image processing operations are applied to small subblocks of an image. The image processing step can be applied ahead of time to each vector in a VQ codebook, with the processed vectors stored along with the codebook. If the computational complexity of the VQ encoder is lower than that of the image processing step, this reduces the computational complexity. This approach is being applied to halftoning, edge detection, and histogram equalization. In a progressive transmission system, the received image is reconstructed as an increasingly better reproduction of the transmitted image as bits arrive. Ways to organize and order a VQ codebook so that it can be used for direct progressive transmission of full search VQ are being studied. In an ordered VQ, the VQ codeword index is correlated with the codeword location in the input space. This ordinal mapping feature of clustering codewords with similar indexes to obtain additional reproduction vectors for the decoder is being exploited. Extentions to progressive transmission of ordered VQ indexes over noisy communication channels are being included.

University of Wisconsin - Madison; Truong Nguyen; *The Structure, Design and Application of M-band Symmetric Wavelets and Filter Banks*; (MIP-9626563); $151,666; 12 months.

The wavelet transform with symmetric basis functions is emerging as a next-generation standard for still-image compression. Wavelet methods outperform the current JPEG standard at low bitrates, and degrade gracefully with decreasing bitrate. The new M-channel symmetric wavelets are used in image and video coders, with particular attention to quantization strategies and motion compensation. This research is addressing the theory, design, and application of orthogonal and biorthogonal linear-phase M-channel filter banks and M-band symmetric wavelets. Multirate filter banks are the fundamental building blocks for subband and wavelet decompositions, as well as frequency-domain/time-domain multiplexing. Theoretical questions being addressed by the research include the structure of linear-phase M-channel filter banks and wavelets, including a comprehensive theory of permissible lengths, minimal structures for parametrization and fast computation, and mappings for two-dimensional signal processing. Criteria for wavelet filter design are also being addressed, including vanishing moments, smoothness and nonnegativity of the scaling function, and unequal length filters. The investigators are studying applications of the above to perceptually-based data compression of image, video and seismic data, multicarrier modulation, and multiscale feature detections. They are also investigating the use of the M-band wavelet transform for multitone modulation. The frequency partitioning properties of the M-channel filter bank enable high efficiency for channels such as twisted-pair copper wire and hybrid fiber-coax, with reduced susceptibility to narrow-band interference.

Statistical Signal and Array Processing

Stanford University; Iain M Johnstone, David L Donoho; *Mathematical Sciences: Adaptive Estimation: New Tools, New Settings*; (DMS-9505151 A001); $70,000; 12 months; (Support from: Signal Processing Systems Program, Computational Mathematics, Statistics - Total Grant $190,000).

This research is developing statistical theory and computational tools in the general area of adaptive methods for representing and analyzing signals, images and other objects, and is showing how to tune them so they are noise-cognizant and stable. Underlying the approach are:

1. the idea of oracles, which know perfectly well how to ideally adapt representations;
2. the idea that the goal of adaptation in the presence of noisy data is to quantify how closely realizable procedures (which do not have privileged information about the object) can mimic an oracle; and,
3. the design of procedures coming as close as possible to the oracle.

The research is also developing methods for comparing different adaptation schemes by comparing oracles of different kinds, for example time-frequency oracles and time-scale oracles. This is an outgrowth of our earlier results on wavelets, where this approach was used to show that wavelets have a property of being nearly-ideally spatially adaptive. In addition a computational environment is being developed for implementing and systematically testing such approaches.

As a further outgrowth of earlier work on wavelets, a number of improvements and extensions of wavelet shrinkage, for example in the directions of classification, confidence bands, correlated data and selection of orthogonal bases, are being explored. These efforts may have two spin-offs. First, some of the results may be stimulating and/or useful to the community of "inventors of adaptive procedures" in signal, image, speech, and time/frequency, and related communities. Second, the theoretical work may stimulate statisticians to take more interest in making further contributions in such directions.

University of Colorado-Boulder; Clifford T Mullis, Louis L Scharf; *Statistical Signal Processing and Communication Within Subbands*; (MIP-9529050); $61,789; 12 months.

The efficient use of bandwidth for multimedia information transmission depends on the efficient use of imperfect channels, such as coaxial cables and optical fibers. This means redundant information must be removed and what remains must be coded for reliable transmission. This research is addressing the problems of redundancy removal (source coding) from the point of view of multichannel filterbanks; it is also addressing the problem of efficiently implementing filterbanks in digital electronics.

Specifically, we are studying applications of undersampling operators, oversampling operators, and tight frames to source coding, channel coding and halftoning. The aim of this research is to develop efficient state variable realizations for tight frames and related maps, and to determine the computational complexity and performance characteristics of these realizations. Of particular interest are sum of all-pass realizations for multichannel filterbanks. The results will produce insights and design rules for implementing source coders and channel coders in state variable filterbanks that have good numerical properties.

University of Florida; Jian Li; *NYI: SAR Image Formation and Processing Techniques for Environmental Monitoring*; (MIP-9457388 A002); $45,347; 12 months.

This research shows the advantages of using Synthetic Aperture Radar (SAR) to detect, analyze, and quantify environmental changes. There are three foci.

1. Improve SAR image formation with a phased array airborne or spaceborne radar. Efficient SAR image formation techniques are being developed by appropriately designing the transmitted waveforms of the phased array radar. The trade-offs between nonparametric and parametric techniques are being studied.
2. Research on the SAR image understanding and ground truth evaluation. Statistical clustering algorithms and various feature extraction schemes that adequately incorporate electromagnetic phenomena are being developed and evaluated.
3. Research on the detection, analysis, and quantification of environmental changes through repeated SAR imaging of critical regions including the environmentally fragile Florida wetlands. Change detection techniques are being developed to quantify and document subtle environmental changes from these images.

Purdue University; Michael D Zoltowski; *Closed-Form Angle Estimation with Circular Arrays/Apertures for Mobile/Cellular Communications and Surveillance Radar*; (MIP-9320890 A002, A003); $28,770.

The digital communications industry is currently investing enormous resources towards the development and experimental verification of prototype antenna arrays to be deployed on mobile communication vehicles, including the commercial automobile of the future, as a means of discriminating amongst signals co-located in frequency based on their respective spatial locations. Given the small aperture on a mobile communications unit, the Uniform Circular Array (UCA) geometry is ideal due to its rotational invariance with respect to azimuth. This research is based on a recent development of a simple, closed form algorithm, UCA-ESPRIT, for use in conjunction with a UCA that provides automatically paired source azimuth and elevation angle estimates. To date, the algorithms for 2D arrival-angle estimation have required expensive spectral searches, iterative solutions to multidimensional optimization problems, or ad-hoc schemes for pairing direction cosine estimates with respect to each of a number of different array axes. UCA-ESPRIT is fundamentally different from ESPRIT in that it is not based on a displacement invariance array structure but rather is based on phase mode excitation and hinges on a recursive relationship between Bessel functions.

A theoretical performance analysis of UCA-ESPRIT is being conducted. This proves extremely useful for predicting its performance in a mobile communications environment. Novel strategies for incorporating mutual coupling effects are also being developed. The real world performance of UCA-ESPRIT is being assessed with experimental data from a prototype circular antenna array currently being built at the Polytechnic University of Madrid for mobile sea communications with the INMARSAT satellite system. Adaptations of UCA-ESPRIT for filled circular arrays are being developed.

University of Minnesota; Kevin M Buckley; *PYI: Digital Signal Processing for Hearing Aids and Source Localization*; (MIP-9057071 A004); $37,500.

In the general area of Source Localization Estimation (SLE), this research will;

1. continue SLE algorithm performance evaluation(s),
2. investigate algorithm improvements based on considerations of analytical performance expressions derived during the course of this project, and,
3. address the combined issue of robust estimation and high resolution in the presence modeling errors.

In the area of acoustical/biomedical digital signal processing, specifically hearing aids, effort will be directed towards the specification of algorithm constraints and appropriate cost functions, incorporating robustness and real- time testing both in the laboratory and in the field. Adaptive methods will be examined for use in active noise cancellation of undesired nonstationary noise.

University of Pennsylvania; Saleem A Kassam; *Arrays for High Resolution Imaging and Efficient Digital Filtering*; (MIP-9321856 A003); $64,462; 12 months.

Many imaging systems use arrays of individual elements to sense the propagating field produced by an object. Through signal processing techniques such as beamforming, an image of the object may then be formed. In addition to such passive arrays, imaging systems may use active arrays to both illuminate an object with a radiated field and to record the reflected field. Examples of imaging array systems occur in radio astronomy, sonar, microwave imaging and ultrasound imaging.

This research is aimed at developing fundamental new results for array design and associated signal processing, based on recent developments on the characterization of array performance in linear imaging. Active imaging systems (e.g., ultrasound imaging arrays) are a particular focus of this investigation, although some of the proposed work also addresses passive arrays. Using the idea of the "coarray", this work will consider how to deploy array elements (and associated hardware) in the most efficient way to obtain large array apertures and high resolutions. The results will allow minimum redundancy active arrays and minimum complexity active arrays to be specified. This research on arrays includes a study of their characteristics under real operating conditions and will extend to some experimental work with an acoustic array system.

University of Texas - Austin; Guanghan Xu, Dim-Lee Kwong; *Development of Advanced Signal Processing Algorithms for On-Line Temperature Profile Measurement in Semiconductor Manufacturing*; (MIP-9400732 A003); $79,980; 12 months.

This research is directed towards the development, implementation, and demonstration of advanced model-based signal processing algorithms for real-time measurement of wafer temperature profile in Rapid Thermal Processing (RTP). RTP

cluster tools are strategically important for submicron semiconductor manufacturing because of trends towards reduced thermal budget and tightened process control requirements on large diameter silicon wafers. Despite its significant advantages, commercial versions of RTP modules for various chemical vapor deposition applications are not available.

This project is developing advanced model-based signal processing algorithms, which when coupled with the acoustic thermometry and acoustic/pyrometer approaches, accurately measure the wafer temperature profile at fast acquisition rates and with a minimum number of sensors. This entails algorithm development, implementation, and validation using real data from commercial RTP tools at UT-Austin and SEMATECH.

University of Texas at Austin; Guanghan Xu; *CAREER: Development and Implementation of Antenna Array Processing Techniques for Wireless Communications*; (MIP-9502695 A001); $41,301.

Array signal processing techniques were traditionally limited to military applications. Researchers recently found that these techniques could be applied to significantly expand channel capacity and improve quality of wireless communication systems through exploitation of spatial diversity. Despite significant research activities in array processing during the last decade, conversion of military technology to commercial technology has not been cost-effective. This research is developing innovative antenna array processing techniques and fast implementation schemes for various wireless communication systems. The program scope is not limited to algorithm development, theoretical analysis, and simulation studies; our ultimate objective is to implement the resulting algorithms in real hardware.

Industries specializing in digital signal processing and its applications in telecommunications have experienced rapid growth due to a high demand for information access and processing. To respond to such a growth, new courses in these areas and innovative teaching methodologies must be developed so students can quickly adapt to various challenges. Plans include;

1. the development of several new courses in these areas including undergraduate laboratory courses,
2. the continued involvement of undergraduate students in industry sponsored projects,
3. the design of more computer projects,
4. the creation of opportunities for minority students and students with disabilities, and,
5. the development of a multimedia teaching method drawing on advanced technology, e.g., live audio, video, and graphical illustrations and real hardware demonstrations.

William Marsh Rice University; Don H Johnson; *Adaptive Receivers for Uncertain, Time-Varying Channels*; (NCR-9628236); $60,000; 36 months; (Support from: Signal Processing Systems Program, Communications Research - Total Grant $186,542).

Analysis based upon the "typicality" of the data has become one of the primary theoretical and practical concepts in modern information theory. The so-called "Type" (appropriate histogram estimators of the amplitude distribution) have not only been shown to be sufficient statistics for both classification and compression, but have also been shown to converge exponentially fast to the true underlying distribution. Because type-based detectors make no a priori assumptions about the channel characteristics and measure the quantities needed to provide asymptotically optimal reception, these detectors can theoretically achieve the same exponential error rate as the clairvoyant receiver that knows the channel characteristics perfectly and uses them optimally.

The goal of this research is to develop a nonparametric adaptive communications receiver strategy by exploiting results from the theory of universal classification and data compression. The approach will be extended to deal with intersample dependencies and colored noise by using channel measurements in an optimally effective way. Two techniques drawn from data compression work, context trees and Lempel-Ziv (universal) coding, are being examined to determine which represents channel-induced dependencies most efficiently and which is most adaptable. By merging systems based on these two theories, a receiver that adapts to unknown, time-varying channels will be developed. It will be demonstrated that this receiver can be used in multiuser channels, both wireless and optical fiber, with little modification. The research will develop the theoretical underpinnings of this adaptive receiver strategy, and will develop a complete software implementation of the receiver, using as test data actual communication channel recordings.

Brigham Young University; A. L Swindlehurst; *Analysis and Development of Algorithms for Antenna Array Based Communications Systems*; (MIP-9408154 A001); $49,343; 12 months.

Communication systems employing antenna arrays are required in situations where multiple co-channel signals are present simultaneously. The spatial discrimination provided by multiple antenna elements allows spectrally overlapping signals to be individually extracted. Such systems have typically been proposed in the context of military applications, but important commercial applications have recently gained attention. It has been proposed that antenna arrays be used in land-based mobile radio systems to provide enhanced spatial discrimination and hence increased capacity. Such arrays could potentially provide frequency reuse in adjacent cells, or within a communication cell itself.

This research is examining signal copy, interference cancellation, and source localization using antenna arrays, with emphasis on scenarios that might be encountered in mobile radio communications (e.g., multipath fading environments, signal formats such as analog FM IS-3 AMPS, GSM, IS-54 offset QPSK, IS-95 CDMA, etc.). This research focuses on the development and analysis of algorithms that exploit all available temporal and spatial information. Conventional techniques for Direction Finding (DF) and Signal Copy (SC) rely on either spatial or temporal structure in the data, but typically not both. Three different techniques that exploit both spatial and temporal signal structure for improved DF/SC are being studied;

1. iterative blind-least-squares,
2. decision directed algorithms array-based equalization of multipath channels, and,
3. optimal joint DF/SC with parametric array uncertainty.

Finally, special emphasis is being placed on the use of real mobile cellular radio data to test the effectiveness of the algorithms developed.

George Mason University; Geoffrey C Orsak; *Adaptive Receivers for Uncertain, Time-Varying Channels*; (NCR-9628294); $57,471; 36 months; (Support from: Signal Processing Systems Program, Communications Research - Total Grant $149,995).

Analysis based upon the "typicality" of the data has become one of the primary theoretical and practical concepts in modern information theory. The so-called "Type" (appropriate histogram estimators of the amplitude distribution) have not only been shown to be sufficient statistics for both classification and compression, but have also been shown to converge exponentially fast to the true underlying distribution. Because type-based detectors make no a priori assumptions about the channel characteristics and measure the quantities needed to provide asymptotically optimal reception, these detectors can theoretically achieve the same exponential error rate as the clairvoyant receiver that knows the channel characteristics perfectly and uses them optimally.

The goal of this research is to develop a nonparametric adaptive communications receiver strategy by exploiting results from the theory of universal classification and data compression. The approach will be extended to deal with intersample dependencies and colored noise by using channel measurements in an optimally effective way. Two techniques drawn from data compression work, context trees and Lempel-Ziv (universal) coding, are being examined to determine which represents channel-induced dependencies most efficiently and which is most adaptable. By merging systems based on these two theories, a receiver that adapts to unknown, time-varying channels will be developed. It will be demonstrated that this receiver can be used in multiuser channels, both wireless and optical fiber, with little modification. The research will develop the theoretical underpinnings of this adaptive receiver strategy, and will develop a complete software implementation of the receiver, using as test data actual communication channel recordings.

Other

University of Arizona; Richard J Greenberg, Robert G Strom; *Image Processing Teaching: Curriculum Materials Development*; (ESI-9553841); $75,000; 12 months; (Support from: Instructional Materials Development, Signal Processing Systems Program - Total Grant $399,774).

Image processing has been shown to be a powerful medium for engaging diverse learners, whose needs have not been met in the past, in the study of science and mathematics. The technique is an effective and engaging way to study the applications of science and mathematics to real-world situations as represented by digital imagery. It brings science research to the classroom; science is learned as science is done. In this project, students and teachers are lead into exploration, inquiry and analysis using real data sets and state-of-the-art computer tools. These supplemental materials, useful in middle and high schools are based on

successful materials developed in teacher enhancement programs and provide a context for changing traditional classroom structure to a learning community. Teachers are supported by on-line services and a network of professional development opportunities. Subject specific CDROMs are produced in physics, biology and Earth science. The standards-based activities are designed to meet specific learning goals and lead to creative student-generated products for performance-based assessment of student achievement. Students also learn about mathematical concepts that underlie image processing such as coordinate systems, slope and intercept, measurement and statistics.

University of California - Berkeley; Kjell A Doksum; *Mathematical Sciences: Topics in Nonparametric Analysis and Model Building*; (DMS-9625777); $30,000; 12 months; (Support from: Signal Processing Systems Program, Statistics - Total Grant $72,000).

With the advent of computer data bases of unprecedented size and complexity and with the dramatic increase in computer power, it has become increasingly more desirable and possible to develop more flexible models, concepts, and procedures that can be used to study relationships between variables and to construct models without relying on rigid global assumptions. Much of the recent work in statistics have addressed this need for more general and flexible methods.

This research addresses the question of how intuitive and concise linear model concepts and techniques can be extended to nonparametric settings. Nonparametric counterparts of such commonly used linear model ideas as regression coefficients, correlation coefficients, coefficient of determination, and principal components are considered. The research also studies nonparametric techniques for model diagnostics that can be used for dimensionality reduction and to address the question of adequacy of particular models. Both asymptotic and finite sample properties are studied, and the problem of developing reliable data-based methods for smoothing parameter selection for functionals and curves is addressed.

University of California - Los Angeles; A. N. Willson; *Circuits and Systems for Signal Processing*; (MIP-9632698); $294,216; 36 months.

A major research area, one evolving out of the principal investigator's long-term work in nonlinear circuit theory, is the development of a rigorous stability theory for the operating points of dc transistor circuits. A complete solution (necessary and sufficient conditions for dc operating point stability in nonlinear circuits) seems almost within grasp, and this research project seeks to obtain it. Research on low-power circuits encompasses the design of new low- power clock buffers, schemes for implementing circuits to achieve charge recovery on a databus and the investigation of techniques for the design of low-power practical circuits such as Viterbi decoders for CDMA digital cellular applications.

Another project concerns the development of new techniques for multiplierless digital filter design, including the design of novel multiplierless adaptive filters. Work will continue on new techniques for IIR filter pipelining, a topic that shows great promise in providing more economical structures (in terms of lower power-dissipation and IC-area requirements) for DSP systems. In addition, research will continue in the area of video compression, where a technique for rate-distortion optimization is being developed. Finally, recent circuit design experience has made evident the need for advances in logic simulation. The principal investigator's research has already produced several simple ways to improve simulation efficiency and it is expected that major improvements will be forthcoming through further vigorous pursuit of this project

Colorado State University; Richard A Davis, Peter J Brockwell, Murray Rosenblatt; *Mathematical Sciences: Time Series Models and Extreme Value Theory*; (DMS-9504596 A001); $10,000; 12 months; (Support from: Signal Processing Systems Program, Statistics - Total Grant $70,000).

The research is concerned with problems of estimation and research for time series models whose theory is not yet fully understood. The standard linear Gaussian models for time series data inadequately describe many of the time series observed in practice. It is therefore important to develop techniques for estimation and prediction based on more general models. Efficient estimation procedures and prediction techniques for non-causal and non-invertible ARMA processes are being developed, and a study of the theory and application of continuous-time linear and non-linear ARMA processes is being made. Extreme value theory for linear and non-linear models is also investigated.

University of Colorado- Boulder; Delores M Etter; *Workshop on Future Research Directions for Signal Processing Research, August, 1996, Durango, CO*; (MIP-9614719); $56,740; 12 months.

Today, all areas of technology use signal processing. The applications range from adaptive coding of speech signals, to analysis of medical images, to echo- canceling in cellular telephones. The successful implementation of signal processing algorithms in both software and hardware is a result of research into various aspects of signal processing, including analysis, modeling, synthesis, algorithm development, tool development, and implementation. This workshop in intended to bring together a distinguished group of research scientists and educators to explore the expanding role of signal processing within applications that include the telecommunications infrastructure and the expanding information highway.

The workshop will meet for 3 1/2 days to assess and document recent advances and current and future research directions in signal processing. During this time we expect to prepare recommendations for research activities and to initiate a debate on the future directions of the field from among the research community. Participants will be drawn from academia, industry, and companies whose products include significant signal processing components. The goal of the meeting will be to develop a vision statement for the field and determine areas which have the potential for high payoff in terms of application areas. The participants will also discuss ways of integrating signal processing research results and applications into the undergraduate engineering curricula - "the integration of research and education. " Another important topic will be the discussion of ways to promote public awareness of the pervasive role of signal processing.

Iowa State University; Dianne H Cook; *Using CAVE Technology to Explore High-Dimensional Data*; (DMS-9632662); $15,000; 24 months; (Support from: Signal Processing Systems Program, Computational Mathematics, Statistics - Total Grant $50,000).

Virtual Reality offers the construction and realization of an artificial 3-Dimensional (3-D) world that has important potential use in scientific areas. Statistics is the science of making sense of complicated data. For example, typical environmental data sets may contain thousands of data values on hundreds of variables. Graphical methods are needed for exploring and determining relationships among variables such as tree crown health, topography, temperature, population density and atmospheric pollutants.

This research explores visualization methods for examining high-dimensional data with low-dimensional projections with CAVE Virtual Reality technology. Dynamic graphical methods for exploratory data analysis are examined in a 3-D immersible environment. The methods are based on generating motion graphics from continuous sequences of 3-D projections of high-dimensional data. Interactive user control is considered in the form of controlling the path of the 3-D sequences in the data space.

University of Maryland - College Park; K. J. Ray Liu; *NYI: High Performance Computing for Signal Processing*; (MIP-9457397 A002); $31,250; 12 months; (Support from: Prototyping Tools and Methodology Program, Signal Processing Systems Program - Total Grant $62,500).

There are three major architectural models used in high-performance signal/image processing:

1. VLSI- signal processing - high- throughput VLSI architectures for low-cost application-specific implementations used in applications such as communication systems, speech, video/HDTV, and radar;
2. parallel signal processing on massively parallel computers - parallel algorithms for complex signal/imaging systems used in computer vision, medical imaging, and the processing of vast amounts of data in deep space exploration;
3. distributed signal processing on high-speed networks - used in applications such as document image processing, multimedia, automatic signal processing in manufacturing, and medical signal/image processing.

This research will focus on the development of efficient algorithms and architectures for each architectural model, and in comparative studies of the advantages and disadvantages of these different computing schemes. The goal is to investigate which signal/image processing problems can be carried out optimally under different computing and communication schemes.

Massachusetts Institute of Technology; Alexander Samarov; *Topics in Nonparametric Analysis and Model Building*; (DMS-9626348); $20,000; 12 months; (Support from: Signal Processing Systems Program, Statistics - Total Grant $59,000).

With the advent of computer data bases of unprecedented size and complexity and with the dramatic increase in computer power, it has become increasingly more desirable and possible to develop more flexible models, concepts, and procedures that

can be used to study relationships between variables and to construct models without relying on rigid global assumptions. Much of the recent work in statistics have addressed this need for more general and flexible methods.

This research addresses the question of how intuitive and concise linear model concepts and techniques can be extended to nonparametric settings. Nonparametric counterparts of such commonly used linear model ideas as regression coefficients, correlation coefficients, coefficient of determination, and principal components are considered. The research also studies nonparametric techniques for model diagnostics that can be used for dimensionality reduction and to address the question of adequacy of particular models. Both asymptotic and finite sample properties are studied, and the problem of developing reliable data-based methods for smoothing parameter selection for functionals and curves is addressed.

Massachusetts Institute of Technology; Kai-Yeung Siu; *NYI: Analysis and Design of Artificial Neural Networks; (MIP-9696144); $31,250*; 12 months; (Support from: Signal Processing Systems Program, Computing Systems Research Program - Total Grant $$62,500).

Artificial neural networks present a new model for massively parallel computation and a promising paradigm for solving large scale optimization problems. This research is exploring the advantages of neural network-based models over conventional models for computation, and a novel design of neuromorphic computing architectures for applications in signal and image processing. A theoretical framework is being established to derive tight tradeoffs between the number of elements and the number of layers in neural networks. The results should answer some of the key open questions in the analysis of neural networks using classical mathematical tools such as rational approximation techniques and harmonic analysis.

TERC Inc; Susan J Russell; *Investigations in Number, Data, and Space: An Elementary Mathematics Curriculum*; (ESI-9050210 A008); $600; 12 months; (Support from: Instructional Materials Development, Signal Processing Systems Program - Total Grant $549,820).

TERC will develop a comprehensive mathematics curriculum for grades K-6 based on investigations in number, data, and space, and emphasizing depth and understanding. The curriculum will stress mathematics as a pattern-finding science, make the discipline more accessible to both students and teachers, and build teachers' knowledge of how students learn mathematics. TERC will design, evaluate, and disseminate ten curriculum modules for each grade level. Each module will focus on a set of related and developmentally appropriate investigations which introduce key mathematical content within a compelling context. The teaching is built into the print materials (through scenarios depicting analytic work in classrooms), and into videotapes of effective mathematics teaching in action. Evaluation of the project will include formative research, as well as an examination of the impact of the curriculum on students and teachers. Investigations will be disseminated through publications aimed specifically at elementary school educators. The curriculum will be published by Dale Seymour Publications, who has made a major financial commitment to the project.

Michigan Technological University; Jeffrey O Coleman; *RIA: Convex-Programming Design of Signals and Systems*; (MIP-9409686 A001); $2,500; (Support from: Signal Processing Systems Program, Design Automation Program - Total Grant $5,000).

The PI is creating a special-purpose programming language in which optimization problems can be specified in a natural and direct way. The focus is on developing associated translation software to convert such a specification into a particular canonical form for numerical optimization using recently developed, ultra-efficient, interior-point algorithms. The canonical form is a Linear-Matrix-Inequality (LMI) program, where each constraint takes the form of a requirement that a linear (plus a constant) matrix function of the optimization variables be positive definite.

Control theorists have recently demonstrated that a tremendous variety of common (and uncommon) constraints can be put into this generic form. Often, however, the required LMI constraints are not related to the underlying problem in an intuitive way. Software to translate specifications in a "comfortable language" to sets of LMIs is being developed in order to make these powerful optimization techniques easy to apply. The language is being applied to real problems in communication and signal-processing circuit designs.

University of Minnesota; Keshab K Parhi; *NYI: Dedicated VLSI Digital Signal and Image Processors*; (MIP-9258670 A004); $62,500; 12 months.

Research efforts are directed towards the design of dedicated, high-performance digital signal and image processors. The emphasis is on real-time processing, where samples are processed as they are

received from the source, as opposed to being stored in buffers and then processed in batch. Design of algorithm topologies for recursive signal processing algorithms were once considered a major challenge. Using the relaxed look-ahead technique, new concurrent algorithms and topologies for adaptive LMS and lattice filters, cascade and lattice recursive digital filters, and predictive speech and image coders have been developed. Design of concurrent topologies for wave digital filters, decision-feedback equalizers, and adaptive differential vector quantizers are being pursued. The decoding speed in Huffman and arithmetic coders (used for lossless compression) is limited due to the feedback. For the Huffman decoder, the codeword length multiplicity constraint is being exploited to design codes where multiple bits can be simultaneously decoded in parallel. The performance of these decoders is further improved by the use of conditional coding. Novel approaches for design of parallel arithmetic coders are also being pursued.

State University of New York - Stony Brook; Michael M Green; *NYI: Improved Circuit Simulation Using Results from Circuit Theory*; (MIP-9457387 A002); $31,250; 12 months; (Support from: Signal Processing Systems Program, Design Automation Program - Total Grant $62,500).

This research is applying results in the area of nonlinear circuit theory to enhance the simulation of analog circuits. Improvements to the continuation methods of solving dc operating points of circuits are being made to guarantee that all circuit operating points are found during a single analysis, and are being applied to sensitivity analysis of circuits. Erroneous models are thought to be a major source of convergence problems and erroneous results in circuit simulation. Another enhancement to circuit simulation includes checking the accuracy of transistor models by verifying that all models satisfy passivity and the no-gain condition.

Earmark Inc; Elisabeth Perez-Luna; *Intersections: Art, Science and Technology -- A Radio Project*; (ESI-9634183); $8,000; 12 months; (Support from: Informal Science Education, Prototyping Tools and Methodology Program, Signal Processing Systems Program - Total Grant $25,000).

This project will examine the implications of the intersection of art, science, and technology as revealed by the events occurring around the celebration of the fiftieth anniversary of the first general electronic computer, ENIAC-50. The finished product will be radio programming and audio material that presents the coming together of scientists, artists, and other participants in Philadelphia. It will build on the previous research and explorations in both of these fields toward finding commonalties and ways in which each has influenced the other. The finished product will be broadcast for the general public and also can be made available to scientists and educators in the field as resource material for courses, seminars, and lectures as they explore the intersections of science and art and implications this has for research and education.

Brown University; Harvey F Silverman; *Parallel Architectures for Speech Recognition: Nonlinear Optimization of Expectation-Maximization (EM) Training of Hidden Markov Models (HMMs) in a Reconfigurable Environment*; (MIP-9509505 A001); $45,859; 12 months.

This research focuses on one of the strategic areas of national concern, that of high performance computing. It involves the development, construction and testing of new architectures for high-speed computing. The idea is to combine general-purpose, RISC-based processing nodes, each with multiple Field-Programmable-Logic-Array (FPLA) based coprocessor systems. This kind of system has the hardware properties of a general-purpose computer, but the advantage of performance more akin to that of a special-purpose engine. In particular, the 20-node Armstrong III system has been built, is operational, and provides nearly two-orders of magnitude improvement over current advanced workstations. The hardware system, combined with its configuration compiler, a program that automatically translates a C function subroutine to both machine code and to the hardware design of the reconfigurable coprocessors, are the basic building blocks for this class of architectures.

This project is unique in that not only has the hardware/software system been built, but it is also being tested on real applications. The training of a modern hidden markov model based speech recognition system requires hundreds of hours, even with recent variants developed at Brown and elsewhere. The reduction of this training time to ten or so minutes allows progress in this area to be made at a faster rate and/or that expensive nonlinear optimization techniques may now be applied to the problem. Also, this means that data from a less cumbersome sensor system (the microphone array systems in place and being developed at Brown) can be suitably incorporated into a more robust speech recognition system.

Experimental Systems

Dr. Michael J. Foster, Program Director
(703) 306-1936 mfoster@nsf.gov

The Program

The Experimental Systems program supports research projects that involve building, evaluating, and experimenting with a computer or information-processing system. These are goal-oriented projects generally undertaken by teams of designers, builders, and users. The building of the system must itself represent a major intellectual effort, and offer advances in our understanding of information systems architecture. A system supported by the Experimental Systems program will usually include both hardware and software components.

Research on information processing systems involves interaction among diverse elements such as hardware architectures, computational models, compilers, operating systems, applications, performance evaluation tools, and user interfaces. Building and evaluating real experimental systems is the only way to understand these interactions in large systems; other techniques, such as simulation and analysis, have only limited uses in understanding the system issues in such a complex environment. Software simulators, for instance, do not provide the computing speed needed for large experiments, nor the needed performance incentives for porting large application systems for experimentation. Without real experimental systems, important areas of information systems architectures cannot advance.

A successful proposal to the Experimental Systems program should demonstrate the feasibility and utility of the project. Feasibility can be shown by describing prior proof-of-concept prototypes or simulation studies that indicate that the proposed system can be built and will meet its design goals. Utility can be shown by demonstrating that building the system will provide substantial advances in computer system architecture, or that the system is inherently useful. Details of the measurement and evaluation procedures that will demonstrate the benefits of the system in an application should be given in the proposal.

The system to be built must be novel in some way, and the impact of the novel aspects of the system upon its architecture must be evaluated during the course of the research. To justify construction, the new system must be potentially superior to existing systems in the chosen application area. Ideally, building the system would provide new knowledge of systems architecture, open up new application areas, and/or contribute to our knowledge about system building techniques. An appropriate project might be a system built using a new architecture or technology, which addresses an application in a new way. An inappropriate project would be one in which the research uses, simply as a platform, a special purpose machine whose design, fabrication, and evaluation are straightforward. The novel aspects of an experimental system may fall into several different areas; the system might feature application of a new technology, new architecture, or new techniques for performance measurement and evaluation to a computationally stressing problem. Examples of technological innovation are massively parallel analog systems, or applications of opto-electronics. Architectural innovations might include new parallel l/O structures, hardware-software codesign, or limited modifications to commodity processors. New evaluation techniques might include instrumentation for performance evaluation or debugging. These innovations might be applied to produce high-performance computers, intelligent sensors, or signal processing architectures, for example.

To justify support under this program, a proposal should show that system building is necessary for answering significant and timely research questions. The research issues should be such that the best way to address them is to build the proposed system and measure its performance. Building for its own sake is discouraged; analysis and simulation should be performed in sufficient detail before a proposal is sent to the Experimental Systems program. Furthermore, off-the-shelf hardware should be employed in the building stage whenever the research goals do not require custom construction.

By encouraging the design, construction, test, and evaluation of novel information processing systems, NSF hopes to achieve several goals:

* Settle major research issues and add to fundamental knowledge in information processing;
* Guide university research in computer science and engineering toward meaningful problems of industrial interest;
* Strengthen the system-building expertise in our research institutions;
* Educate a new generation of researchers in experimental systems research.

Potential applicants are encouraged to discuss their research ideas with the program director prior to formal submission.

Optical Science and Engineering

During the 1996 Fiscal Year the Experimental Systems program participated in an initiative on Optical Science and Engineering (OS&E). Optical Science and Engineering is an enabling field with demonstrated contributions to many technologies and diverse areas of science and engineering: in the use of optics and optoelectronics in information and communications, in the development of novel materials and devices, in the control of chemical reaction pathways, in the elucidation of quantum processes with laser light, in metrology for precision manufacturing processes, in optical instrumentation advances and their profound influence in the physical, biological, and environmental sciences, and in the development of an educated and skilled workforce.

Awards

Storage Hierarchies and Input/Output Systems

California Institute of Technology; Paul Messina: *Scalable I/O Initiative*; (CDA-9415797 A001); $50,000; 12 months; (Support from: Experimental Systems Program, CISE Institutional Infrastructure, Operating Systems and Systems Software - Total Grant $300,000.

This award supports the infrastructure for the Scalable I/O national project. This national project will investigate many of the issues surrounding the scalability of Massively Parallel Processing (MPP) System input and output systems. The national project consists of 18 application groups involving more than 30 researchers.

The equipment provided by this award consists of I/O nodes for an MPP and a small MPP for crashable experimentation. The infrastructure will be available to all of the application groups for their experimental research needs.

University of Southern California; Michel Dubois: *Hardware Prototyping of Shared-Memory Multiprocessor Architectures on RPM*; (MIP-9633542); $794,566; 48 months.

This project is conducting research in multiprocessor architecture by using a rapid prototyping engine for multiprocessors (RPM) that consists of commodity processors and memories interconnected with programmable logic. The main purpose of RPM is to demonstrate the feasibility, complexity, performance, and cost-effectiveness of various selected multiprocessor architectures, by allowing accurate hardware prototypes to be constructed. The prototype can include input/output, and can run production operating systems and application codes. Prototype architectures evaluated as part of this project include cache-coherent NUMAs and cache-only machines with varying cache and memory protocols, and several software distributed shared memory machines with fine-grain hardware support. Expected outcomes of the research include accurate measurements of the performance of various architectures, as well as the demonstration of a new methodology for rapid and cost-effective prototyping of computer systems.

North Carolina A & T State University; Kelvin S Bryant: *Performance Tools and Compiler Support for Parallel I/O Systems*; (MIP-9625083); $200,000; 48 months.

This award supports a unified program of research and education in parallel input-output systems for high-performance computing. Research activities will include development of tools for I/o performance evaluation, applying the tools to characterize I/O performance of scientific codes, and provision of compiler support for parallel I/O. Educational activities span the needs of graduate and undergraduate students, and include outreach to K-12 and the community. These include development of a graduate course, augmentation of undergraduate curricula, and workshops for participants outside the university.

Texas A & M University; A. L. Narasimha Reddy: *Scalable Storage Server for Multimedia and Scientific Applications*; (MIP-9624310); $200,000; 48 months.

This project is building a high-performance distributed storage system using communication switches and workstations. Data to be stored will be distributed across the workstations and will be retrieved using the communication network between the workstations. Research issues to be addressed include identification of hardware and software bottlenecks, limits to system scaling, scheduling of multimedia I/O in a distributed system, and the implications of I/O on switch architecture. Education is a part of this project; plans include a graduate course on I/O systems, an undergraduate laboratory in computer architecture, and use of the machine developed during the research as a testbed for the graduate course.

William Marsh Rice University; Willy Zwaenepoel: *Reliable Main-Memory Storage Systems*; (MIP-9521386); $712,263; 24 months.

This project investigates the use of non-volatile semiconductor memory (such as flash memory) as a storage system to replace some of the functions currently assigned to disks. A board (eNVy) incorporating a large (128MB) flash memory is being designed at Rice University; concurrently, a log-structured file system (Rio) that can use the flash memory board as a cache between main memory and disk is being developed at the University of

Michigan. In later stages of the project these research prototypes will be integrated. Research issues in hardware development include space reclamation (garbage collection), and tradeoffs between memory size, performance, and memory lifetimes. Research issues in file system development include optimization of the sizes of cache and main memory, algorithms for compression of data in the cache, protection of memory from software faults, and exploration of file system index structures. Integration of the prototypes will provide insight into the protability of these techniques to new systems.

General Purpose Computing

Stanford University; Michael J Flynn, Bruce A Wooley, S. Simon Wong, Giovanni De Micheli, R Fabian W Pease: *Sub-Nanosecond Arithmetic II*; (MIP-9313701 A002); $330,000; 12 months.

This project is attempting to speed up computer arithemtic by several orders of magnitude using a combination of algorithmic, circuit, and packaging techniques. CAD tools to automate the application of these techniques are also being developed under the project. Specific research problems include the development of a package capable of passing large numbers of signals with 100 picosecond rise times, and the integration of wave pipelined data paths into an overall system.

University of California - Berkeley; David E Culler: *A Next-Generation Infrastructure for Integrating Computing and Communications*; (CDA-9401156 A004); $200,000; 12 months; (Support from: Experimental Systems Program, CISE Institutional Infrastructure - Total Grant $263,630.

This award provides support for the development of Titan, a computing system consisting of an integrated ensemble of computing and communication elements, organized to provide the user with a number of services. These services will include multimedia capabilities in delivery vehicles; storage and communication; large computing power; large storage space; innovative parallel languages, debuggers, and libraries; and high accessibility from both mobile and fixed locations. The experimental facilities requested include workstations and servers constituting the backbone of the distributed system providing cycles to the user, ATM switches for linking workstations and servers, a video editing system, a massive storage unit, and equipment for linking the network to the currently available CM-5 parallel computer.

The proposed research projects fall into three areas: network and communications; distributed supercomputer projects which are mainly concerned with providing parallel computing to every user through a combined architecture, operating systems, and programming language effort; and multimedia services which requires integrating systems support, software support, and artificial intelligence tools to create, store, play, edit, search, input, and output multimedia objects. The networking, multimedia, and computing aspects of Titan form will form the infrastructure for a number of computationally intensive applications.

University of Illinois - Urbana-Champaign; Josep Torrellas: *NYI: Increasing the Performance of Scalable Shared-Memory Multiprocessors*; (MIP-9457436 A002); $87,500; 12 months.

Scalable shared-memory machines are a promising way of attaining large-scale multiprocessing without surrendering much programmability. Achieving high performance from these machines, however, is challenging because many complex architecture and architecture-software implementation issues that have been only partially studied considerably impact the performance of the machines. The objective of this research is to contribute in three areas to help make shared-memory multiprocessors the preferred source of computing power. The three thrusts of this project are to: study the optimal division of responsibilities among the hardware, compiler, and operating system to maintain cache coherence in scalable share-memory multiprocessors; design algorithms and hardware to effectively support multiprogramming of parallel programs in scalable shared-memory multiprocessors; and optimize the management of memory hierarchies and the interaction of the operating system with the architecture.

Princeton University; Kai Li, Margaret R Martonosi, Douglas W Clark, Edward W Felten, Richard J Lipton: *SHRIMP: Architectural and Systems Support for Inexpensive, High-Performance Multicomputers*; (MIP-9420653 A001); $450,000; 12 months.

This project is building a high-performance multiprocessor from commodity desktop computer systems and off-the-shelf interconnects. Commercial Intel Pentium workstation boards, each with attached memory, disk, and I/O, are attached to a Paragon backplane. Communication uses a new mechanism called virtual memory-mapped communication, which disguises interprocessor communication as write operations to memory. The node interface maps physical pages in the memories of individual nodes to each other, so that a write to one mapped page results in messages to other nodes that share the mapped page. The operating systems on the individual nodes use their ordinary virtual memory mechanism to support virtual page mapping. In addition to this word-by-word communication, DMA transfers are available, with control registers located in the address space of individual processes. This allows high bandwidth communication that maintains user-level protection. Research to be addressed in the project includes the achievement of high-bandwith low-latency communication between processes, the structure of an I/O system supported by the new communication mechanism, and performance evaluation of the resulting system.

New York University; Zvi M Kedem: *Collaborative Research: High Performance Parallel Processing: Fault-Tolerant Computing on a Network of Workstations*; (CCR-9411590 A001); $25,000; 12 months; (Support from: Experimental Systems Program, Computer Systems Architecture - Total Grant $71,233).

This project deals with both theoretical and prototyping research in parallel processing to harness the power of workstation clusters. The primary techniques will be based on aggressive scheduling, evasive memory and dispersed data management. The formal work expands the previous results of the PIs in asynchronous computing in several directions:

4. Extend the models to fine-grained programs.
5. Permit the inclusion of architecture specific characteristics into the theoretical framework.
6. Techniques to distinguish inconsistencies in memory states as viewed by different processors resulting from time-outs, interrupts and failures (or indefinite postponement of a computation).
7. Techniques to allow non-deterministic executions.
8. Model input/output operations of fault-tolerant programs.

The experimental part of the project will build a scaled version to execute parallel programs on clusters of workstations. The prototype will address issues related to the eager scheduling of threads, the implementation of evasive memory and data management. The prototype will attempt to evaluate the effectiveness of the environment in terms of efficiency.

This project is conducted jointly by P. Dasgupta of Arizona State University and Z. Kedem of New York University.

University of Washington; Susan Eggers; *Simultaneous Multireading*; (MIP- 9632977); $200,859; 24 months; (Support from: Experimental Systems Program, rototyping Tools and Methodology - Total Grant $519,164).

This is a project to develop and extend support for multithreaded computation on superscalar processors. The intention is to provide high processor utilization despite increasing memory latency. In simultaneous multithreading, multiple independent threads will issue instructions to a superscalar processor's functional units in a single cycle. The objective of this project is to fully define the low-level architecture for simultaneous multithreading, including scheduling, register architecture, pipeline design, and speculative execution. The project is developing compiler and operating systems support for multithreaded execution on the new processor architecture. A primary goal is to maximize multiple thread performance without increasing single thread latency. Most measurements are from detailed simulations, using simulators and compilers provided by industrial partners.

University of Washington; Lawrence Snyder, Carl Ebeling: *Chaotic Routing: Study and Implementation*; (MIP-9213469 A005); $1,000; 18 months.

The chaotic router for multiprocessor systems avoids congestion in message routing by derouting packets chosen at random at congested nodes of a network. The routers can thus adapt to varying message traffic. In this project, the router is being implemented and its performance is being measured within a testbed that approximates a real multiprocessor.

University of Wisconsin - Madison; Max G Lagally: *DNA Computation on Surfaces*; (CCR-9613799); $50,000; 12 months; (Support from: Experimental Systems Program, Theory of Computing, Computational Biology Activities, ARPA - Total Grant $300,000).

1. This multi-disciplinary project concentrates on developing computation at the molecular level, via manipulation of DNA strands on surfaces. An experimental demonstration of DNA-based computation by Adleman
   a. has changed the view of what computation is,
   b. offers the potential for unprecedented computing power at the molecular level, and
   c. raises fundamentally new research problems in chemistry, computer science, and material science.

   This project represents a concerted attack on these problems which involves close collaboration between chemists, material scientists, and computer scientists.
2. This research will help answer two major questions on DNAcomputation. First, does computation at the molecular level have the potential to provide computing power that is orders of magnitude greater than foreseeable extrapolations of current technology? Second, is this type of computation suited to solving NP-hard problems (problems that are well beyond the limits of current technology?
3. The potential for huge computing power rests on the use of DNA strands to encode information and the manipulation of these strands in a massively parallel fashion, involving as many as 2¬70 (2 to the 70th power) distinct strands. The premise of this project is that surface-based chemistry is a critical technology in approaching this scale. With this approach, DNA strands are immobilized on a surface, thus allowing a much greater degree of control in chemical processes that manipulate the DNA than is achievable via the test tube based methodology of Adleman. While surfacebased chemistry is the basis for recent strides in combinatorial chemistry, this project is the FIRST to fully exploit surface-based manipulation of DNA strands for the purposes of DNA computation.
4. Scaling up current surface chemistry to approach the requirements of DNA-based computation requires high-quality research at the interface of materials science and chemistry. Improvements in the nanoscale morph ology and chemical makeup of the surface is key to ensuring that a high density of information can be obtained and reliable chemical manipulations performed. Good surface attachment chemistry and control of chemical ``operations" or enzymatic processes are also extensively developed. This project should result in significant advances in the state of- the-art in surface chemistry and should provide a solid basis for predicting the limits of surfacebased DNA computation.
5. An eventual goal of DNA-based computation is to perform massivelyparallel searches for optimal solutions of NPhard problems. However, the differences between DNAbased computation and conventional technology require that radically new algorithms be developed for this paradigm. Tempering the potential for unprecedented parallelism (up to 2¬70 simultaneous operations) is extremely slow and errorprone nature of the operations themselves. Novel strategies for generating and searching solution spaces for NPhard problems are investigated. These strategies embody sound algorithmic principles that can be applied to a range of problems. Using a combination of analysis and simulations, the effectiveness of these strategies are tested on selected applications. This work provides new ways of attacking NPhard problems that, while designed for a hypothetical DNA based computer, are valuable for massively parallel computing paradigms, other than the DNA-based paradigm studied here.
6. Thus this comprehensive program will provide a deep understanding of the potential and limitations of DNAbased computation, from both the chemical and algorithmic viewpoints.

University of Wisconsin - Madison; Gurindar S Sohi, James E Smith: *Prototyping Multiscalar Processors*; (MIP-9505853 A001); $229,962; 12 months; (Support from: Experimental Systems Program, Computer Systems Architecture - Total Grant $309,962.

This project is evaluating a new architectural paradigm that can extract and exploit the parallelism in sequential code. This new approach uses both software scheduling in the compiler, as in VLIW, and hardware scheduling at run-time, as in superscalar architectures. The compiler segments code into large blocks of instructions that form subgraphs of the control-flow graph, though not necessarily basic blocks. The compiler appends synchronization information to each block that describes which registers must be shared with blocks. The compiler appends synchronization information to each block that describes which registers must be shared with other blocks. The blocks or tasks are passed to separate identical parallel execution units, each of which executes its task sequentially. Tasks are scheduled optimistically, so that some tasks may be executed by a unit when they would not be executed

in a sequential system; in such cases, the tasks are "squashed," which means that their results are not written to memory or registers. At the end of each task, it blocks until the system determines that the task will not be squashed; at that time results are committed. The collection of execution units appears logically to be one unit, with a single register file. Shared register values are passed on an inter-execution-unit network, and tasks that need shared values block until the values are produced.

During the first two years of this award is devoted to compiler design for the new architecture, to comparison with other architectures using simulation, and to conceptual refinement of the architecture. During later years, work is expected to focus on implementation.

University of Wisconsin - Madison; David A Wood, Pei Cao, James R Larus, Mark D Hill: *Tornado: Fine-Grain Distributed Shared Memory for SMP Clusters*; (MIP-9625558); $684,126; 24 months.

This project is exploring the cluster approach to symmetric multiprocessing, in which bus-based multiprocessors are interconnected using a more scalable network. Using the Tempest parallel programming substrate, the project is building clusters that extend the shared memory model with fast references to recently used locations across a cluster. The key issue is achieving good performance without resorting to extensive custom hardware.

The project encompasses design and some implementation of a high-speed Tempest system.

Application Specific Computing

University of California - San Diego; Henry D. I. Abarbanel: *Synchronization and Communication in Nonlinear Optical Systems*; (NCR-9612250); $315,800; 36 months; (Support from: Experimental Systems Program, Networking and Communications Research Program, Optical Communications Systems, Office of Multidisciplinary Activities - Total Grant $1,069,701).

The is a joint effort through subcontracts, with Raj Roy, Georgia Institute of Technology, and Steven Strogatz, Cornell University

The researchers will investigate synchronization of chaotic lasers for purposes of using chaos as a carrier for new modes of communication including multiplexing information bearing signals and private communications using chaotic carriers as well as methods for optical channel equalization using nonlinear properties of the fiber. The program will include experiments on Nd:YAG and Erbium Doped lasers, analysis of synchronization methods and innovative communications techniques, and mathematical analysis of synchronization of a few and assemblies of chaotic oscillators with application to optical communications.

A "group meeting" workshop will be held semi-annually and each year we will invite experts in optics, optical engineering, and optical communications to provide additional research expertise to broaden the horizons of the students, postdoctoral fellows, and faculty involved in this project

University of California - San Diego; Michael J Bailey, Ramesh C Jain: *Making Rapid Prototyping Viable for Remote Use on the National Information Infrastructure*; (MIP-9420099 A001); $318,305; 12 months.

This project is working to make rapid prototyping and reverse engineering facilities available over wide-area computer networks. It includes several research thrusts along with an integration effort to ensure that the facilities will be usable remotely with little on-site intervention. The project has four major thrusts:

1. development of algorithms for checking consistency of geometric descriptions;
2. development of new languages based on constructive solid geometry for communicating geometric descriptions;
3. connection of rapid prototyping equipment to a wide rear network, with device drivers and server software to allow remote access;
4. connection of a 3D scanner to the network for remote access.

University of California - Santa Cruz; Richard Hughey, Kevin Karplus: *Multi-Purpose Parallel Process for Biosequence Analysis*; (MIP-9423985 A001, A002); $176,576; 12 months.

This project is building an application-specific computer system for biosequence comparison. The architecture is suitable for a wide range of sequence comparison methods, including the Smith- Waterman algorithm, BLAST, profile searches and dictionary methods. In addition, the architecture is being integrated with software for the statistical analysis of sequences using techniques such as hidden Markov models.

The core of the architecture is a linear array of SIMD processing elements, each with a small local memory. A single chip can contain 64 of these elements, so a board with 20 chips will hold 1280 processors. The processing elements have been tuned to sequence comparison by incorporating a single cycle add-and-min instruction and a data-path for quickly recording the results of the instruction. Software, I/O, and algorithms have been considered in the design of this architecture; the resulting balanced implementation should result in high performance at low cost.

National Academy of Sciences; Ronald D Taylor: *Assessment of Optical Science and Engineering*; (ECS-9414956 A002); $6,500; 12 months; (Support from: Graduate Fellowships, Experimental Systems Program, Spectroscopy, Metals, Ceramics, And Electronic Materials, Quantum Electronics, Plasmas, and Electromagnetics, Instrumentation and Instrument Development - Total Grant $47,500).

The Board on Physics and Astronomy, in collaboration with the National Materials Advisory Board, has established a committee to carry out a major study of the field of optical science and engineering (OS&E), a diverse interdisciplinary area of scientific and technological activity that is playing an increasingly area of scientific and technological activity that is playing an increasingly important role in driving economic growth. OS&E is emerging as a key enabling element for a host of applications that address critical national needs such as manufacturing, information technology, health care, transportation, environment, and education. It is also a vital area of frontier scientific research. The purpose of the study is to define the field of optical science and engineering, assess the state of the art in research and technology, project the future impact of OS&E on national needs, identify institutional and educational innovations needed to optimize the contributions of OS&E to national needs, identify how public policy influences the ability of the field to meet these needs, and examine trends in private and public research activities as compared to those in other countries. The study is being supported jointly by NSF, NIST, and ARPA. The Foundation of the report will be based on the conduct of 5 major workshops; each workshop would be organized according to national needs and would address issues that cross OS&E scientific and technical disciplinary boundaries.

University of Hawaii; Eun Sok Kim, Mehrdad G Nejhad, Junku Yuh, Stephen Y Itoga: *Rapid Prototyping: Virtual Rapid Prototyping System for Piezoelectric Micro Systems*; (DMI-9420378 A001); $110,000; 12 months.

This award is to develop a virtual rapid prototyping system for piezoelectric micro systems (PMS). This is an area of Micro Electromechanical Systems (MEMS), where actuator or sensor devices are built to the scale of a human hair. PMS's move and either are driven by or produce electrical signals. Some examples are microphones, pumps, or motors. The rapid prototyping system that will be developed used computer simulation to design the PMS devices, design the special processes to make them, and do performance tests on the virtual prototype. As part of the program this interdisciplinary team will build these devices using this system and compare their virtual system predictions with experimental data. The technological impact of this virtual rapid prototyping system will be reduced cost and lead time by providing the engineering tools to design both the PMS devices and the processes to make them. This means new competitive PMS-based products can be designed for commercial and industrial applications. This award will also build the research and human resource infrastructure to be competitive in this area of technology.

Massachusetts Institute of Technology; Robert Berwick; *High Performance Computing for Learning*; (IRI-9217041); $20,000; 12 months; (Support from: Robotics and Machine Intelligece, Knowledge Models & Cognitive Systems, Linguistics, Human Cognition & Perception, Experimental Systems Program, CISE Institutional Infrastructure - Total Grant $636,944).

The Grand Challenge Application Groups competition provides one mechanism for the support of multidisciplinary teams of scientists and engineers to meet the goals of the High Performance Computing and Communications (HPCC) Initiative in Fiscal Year 1992. The ideal proposal provided not only excellence in science: focussed problem with potential for substantial impact in a critical area of science and engineering) but also significant interactions between scientific and computational

activities, usually involving mathematical, computer or computational scientists, that would have impact in high-performance computational activity beyond the specific scientific or engineering problem area(s) or discipline being studied. In the award to Berwick, Bizzi, Bulthoff, Jordan, Wexler, Poggio, Rivest, Winston, and Yang at MIT, the research project - High Performance Computing for Learning - has been designed explicitly to push the High Performance Computing algorithmic and architectural envelope via a CM-5 and VLSI testbed and to address many of the HPCC goals. It will advance new algorithms and software for a broad class of optimization and learning problems, tested on and directly driving operating system and architectural changes on the CM-5 (working with one of the CM-5's key architects). The learning problems addressed are essentially an entire class of modeling/optimization problems that intersect with nearly all HPCC Grand Challenge Problems.

Massachusetts Institute of Technology; John L Wyatt, Ichiro Masaki: *Cost-Effective Hybrid Vision System for Intelligent Highway Applications*; (MIP-9423221 A001); $1,247,000; 24 months.

The goal of this project is a new cost-effective architecture for machine vision, which will be evaluated for intelligent highway applications. Components of three smart image sensors are being developed and integrated into a heterogeneous vision system. The smart sensors use analog, digital, and mixed signal techniques to perform 3 dimensional measurement for adaptive cruise control, lane detection, and time-to-collision measurements. All of these tasks are intended for machine vision systems within intelligent vehicles, and will be tested in intelligent vehicle applications within the later stages of the award. MIT is performing basic research on the first versions of all of these cameras, while later development, integration, and testing will be carried out by industrial partners.

University of Michigan - Ann Arbor; A. Galip Ulsoy, Yoram Koren, Kang G Shin: *Hierarchical Controller for Real-Time Quality Control in Machining*; (DMI-9313222 A003); $30,000; 12 months; (Support from: Robotics and Machine Intelligence, Experimental Systems Program, Manufacturing Machines and Equipment - Total Grant $151,628).

The award is for the design and development of an experimental, modular, open-architecture machining system controller for prismatic and rotational parts that enhances part quality by an order of magnitude while sustaining high machining rates. The approach would enable integration of underutilized results of the prior work of various researchers as well as generating new results. The hypothesis of the research is that information embedded in the control-loop error signals can be effectively used to improve part quality and machining system performance. A hierarchical controller structure is proposed consisting of compensators at the machine tool servo, process, and part inspection levels. Intelligent integration of hardware and software components of the controller is proposed here by combining expertise in machine tools, machining processes, sensors, control theory, software architecture, and real-time computing.

The hierarchical, open-architecture controller is expected to have a major impact on manufacturing practice and research. Potential benefits for industrial users include reduction of part manufacturing cycle times, reduction of indirect inspection costs, controller improvement costs, and the cost of using sensors. Implementation of future machining research results would also be facilitated by the availability of such a controller.

Cornell University; Herbert B Voelcker: *Massively Parallel Computation for Mechanical Manufacturing and Design*; (MIP-9317620 A002); $402,326; 12 months.

Solid modeling is a critical enabling technology for mechanical CAD/CAM because it provides geometrically complete representations of parts and products, and enables important manufacturing processes to be modeled. Today's industrial systems operate far below the technology's potential because solid modeling requires enormous computing resources, and because current algorithms and representations cannot handle several important applications.

This project is designing new implementations of ray-casting representations for the solid models used in mechanical CAD, and extending the models to new applications. New implementations for ray-casting engines are being developed, using both custom hardware and software running on massively parallel machines. Using the new hardware, new applications of ray representations are being explored, including the solution of boundary-value problems, computation of medial axis transforms, methods for representing mechanical tolerances, and representation of solid objects.

State University of New York - Stony Brook; Arie E Kaufman: *Scalable Architecture for Real Time Volume Rendering*; (MIP-9527694 A001); $194,063; 12 months.

This research deals with the development of a new scalable volume visualization architecture and its associated algorithms. The architecture, called Cube-4, exploits parallelism and pipelining to achieve real-time rendering of high-resolution images from volume data. The architecture is based on an algorithm for ray casting of a volume buffer of voxels which is stored as a skewed distributed memory to support conflict-free access to voxel structures. It performs interpolation of sampled points, shading, and compositing to generate the pixel values. Computations are done using limited communication between processors, so that the architecture is scalable over a wide range of performances and image resolutions. This project encompasses algorithm development, architecture research, and construction of a reduced-resolution prototype.

University of North Carolina - Chapel Hill; Henry Fuchs, John W Poulton: *ImageFlow: Real-Time Image-Based Rendering*; (MIP-9612643); $483,131; 12 months.

Depth images are used as rendering primitives in the high-performance graphics engine that is being constructed in this project. A depth image is a two-dimensional image that includes depth relative to a viewpoint, in addition to color and other properties. From a small number of depth images, each of which represents an arbitrarily complex scene from a single viewpoint, an image can be computed for any viewpoint within the neighborhood of the original viewpoints. This image-based rendering approach offers advantages in realism and distributing the rendering computation among processors that generate depth images and those that reproject them to new viewpoints.

The project is using the existing PixelFlow graphics engine to test the image- based rendering ideas. During the first part of the project, algorithms for acquiring and pre-processing depth images will be investigated. At the same time, the PixelFlow machine will be enhanced for image acquisition and storage. This will permit software for image-based rendering to run in real-time on the PixelFlow machine, performing the following operations: determination of reference images to be used as sources of depth pixels, warping of reference image pixels to the screen space, evaluation color, area, and other parameters for each pixel, and blending of pixels to form the final reference image. In later stages of the project, hardware support may be developed and evaluated for some of these operations.

University of North Carolina - Chapel Hill; John W Poulton, Henry Fuchs: *Scalable Graphics: From Personal to Supercomputer Visualization Engines*; (MIP-9306208 A007); (Support from: ARPA - Total Grant $803,641).

The object of this project is to build and experiment with a new graphics engine that will eliminate the current limits to scalability in commercial graphics systems. The work centers on a new graphics engine architecture called image composition, which is radically different from the organization of today's commercial systems. In image composition, rendering is distributed over a number of identical processors. Each renderer generates a full-screen image, but for only a fraction of the primitives in the scene. The system then merges these images over a high-speed network to form a single image of all primitives. Since each subimage is independent, and since the images can be merged on a distributed network whose throughput scales linearly with the number of subimages, performance of the entire system can be scaled up arbitrarily by adding more processors.

University of Oregon - Eugene; Zary Segall, Stephen F Fickas: *Collaborative Research: Architecture, Design and Implementation of Mobile Computers*; (MIP-9403573 A001); $194,615; 12 months.

This is a joint effort between two universities for rapid prototyping of mobile computers. The projects involve teams of students who over the course of a semester design the hardware, software, and packaging of mobile computer systems, and fabricate prototypes by combining standard electronic parts with custom fabricated cases and interconnect harnesses. This project focusses primarily on mobile computing for the diagnosis of telecommunications networks. Mobile computer systems that include stationary and mobile computers, interface devices, and software, are being designed and fabricated for these applications. The project goal is to systematize the prototyping process in several ways: by providing an integrated design tool that can simultaneously represent electronic, thermal, and mechanical constraints; by providing modular template architectures for mobile computers; and by reporting on experience in the co-design of software, mobile hardware, and stationary hardware in a mobile computing system.

Carnegie-Mellon University; L. Richard Carley, David J Allstot, Rob A Rutenbar, Donald E Thomas: *Design of Ultra-Low Power IC's*; (MIP-9408457 A001); $300,000; 12 months.

This project combines analog circuit design and CAD to produce digital circuits that can operate at substantially reduced voltages, which will result in lower power consumption. The project has introduced the QuadRail logic family, which allows voltage control of transistor thresholds to maintain constant logic thresholds. Use of this family permits lower supply voltages because of the tighter control of device thresholds. However, use of this family requires optimization of individual digital cells, new device layout methods, and new strategies for floorplanning, placement, and routing. The research in this project is exploring all of these optimization areas, and is producing a battery powered demonstration system for speech signal processing.

Carnegie-Mellon University; Susan Finger, Lee Weiss, Daniel P Siewiorek, Andrew P Witkin: *Rapid Design Through Virtual and Physical Prototyping*; (MIP-9420396 A001, A002); $159,962; 12 months; (Support from: Experimental Systems Program, Integration Engineering - Total Grant $510,211).

This project is creating an experimental system using the Internet that will allow students in design courses to use rapid prototyping services. Three participant institutions, Carnegie-Mellon University, the University of California at Berkeley, and Stanford University are participating in the project. Each has developed individual technologies for virtual and physical prototyping, which currently stand alone. Bringing these technologies together will result in exciting new capabilities. Partnerships with industrial partners and a Federal laboratory are also anticipated. This project addresses key issues in prototyping, and is creating a network of interconnected services to support the rapid design, test, and manufacture of mechanical, electro-mechanical, and electronic products.

Carnegie-Mellon University; Ralph Hollis, Mahadev Satyanarayanan, Mark H Kryder: *HPCC: A Distributed Architecture for Rapidly Reconfigurable Assembly Systems*; (DMI-9527190 A001); $290,000; 12 months; (Support from: Experimental Systems Program, Manufacturing Machines and Equipment, Integration Engineering - Total Grant $539,953).

This research aims to improve assembly in product manufacturing. Assembly is a difficult and time-consuming process to automate, especially when small tolerances are necessary. This architecture for agile assembly is a strategic framework using agent-based robotic elements that are precise, modular, and extensible. These elements form sets of leased self-contained software/hardware modules that can be programmed and operated over the Internet, and can be brought together to form miniature agile factories for assembly. Such a scheme requires a combination of intelligent networked communication, distributed computing resources using high-performance processors, and distributed sensor/actuator subsystems of novel design. Design and construction of a prototype miniature factory according to the principles of agile assembly architecture will be carried out to validate the research and provide a unique and powerful reconfigurable platform for assembly research and evaluation by industry. Partial assembly of magnetic storage disk drives will serve as a testbed.

This work aims to advance knowledge in three main areas:

1. seamless wide area and local area networking of distributed agents emphasizing high quality of service;
2. modular precision robotic assembly elements containing high-performance embedded processors; and
3. a comprehensive software environment for modeling, simulation, and programming of the miniature factories.

The results of the research could allow manufacturers to develop a capability for geographically distributed design and deployment of assembly systems while providing drastically reduced changeover times, higher quality, and a new level of manufacturing system portability.

Carnegie-Mellon University; Daniel P Siewiorek: *Collaborative Research: Architecture, Design and Implementation of Mobile Computers*; (MIP-9403473 A002, A003); $210,000; 12 months.

This is a joint effort between two universities for rapid prototyping of mobile computers. The projects involve teams of students who over the course of a semester design the hardware, software, and packaging of mobile computer systems, and fabricate prototypes by combining standard electronic parts with custom fabricated cases and interconnect harnesses. This project focusses primarily on mobile computing for the diagnosis of telecommunications networks. Mobile computer systems that include stationary and mobile computers, interface devices, and software, are being designed and fabricated for these applications. The project goal is to systematize the prototyping process in several ways: by providing an integrated design tool that can simultaneously

represent electronic, thermal, and mechanical constraints; by providing modular template architectures for mobile computers; and by reporting on experience in the co-design of software, mobile hardware, and stationary hardware in a mobile computing system.

Carnegie-Mellon University; D. Lansing Taylor, Scott E Fahlman: *High Performance Imaging in Biological Research*; (DBI-9217091 A007, A009); $100,000; 12 months; (Support from: Knowledge Models and Cognitive Systems, Robotics and Machine Intelligence, Neuroscience, Experimental Systems Program, Cell Biology - Total Grant $600,000).

The Grand Challenge Application Groups competition provides one mechanism for the support of multidiscipinary teams of scientists and engineers to meet the goals of the High Performance Computing and Communications (HPCC) Initiative in Fiscal Year 1992. The ideal proposal provided not only excellence in science: focussed problem with potential for substantial impact in a critical area of science and engineering) but also significant interactions between scientific and computational activities, usually involving mathematical, computer or computational scientists, that would have impact in high-performance computational activity beyond the specific scientific or engineering problem area(s) or discipline being studied.

This is a project to research and develop an Automated Interactive Microscope (AIM). The AIM will combine the latest technologies in light microscopy and reagent chemistry with advanced techniques for computerized image processing, image analysis, and display, implemented on high-performance parallel computers. This combination will produce an automated, high-speed, interactive tool that will make possible new kinds of basic biological research on living cells and tissues. While one milestone of the research will be to show the proof-of-concept of AIM, the on-going thrust will be continued development as new technologies arise and the involvement of the biological community.

Pennsylvania State University; Robert M Owens, Mary J Irwin: *Architecture, Algorithms and Software in the Design of a Massively Parallel, Fine Grain Processor*; (MIP-9408921 A001, A002); $180,326; 12 months.

This project is investigating the design, implementation, and use of a family of family of massively parallel computers that use one processor per digit. A first generation design, completed in an earlier project, contains 16,384 fine grain processors on a single board. The processors contain a small number of parallel logic circuits and a configuration memory that controls the logic and interconnect. The entire machine is programmed for a problem by setting all configuration registers to transform the machine into a special-purpose computer for the problem. By improving the density of processors on chips and the speeds of interconnect and interface, this project is increasing the operation rate of the machine by a factor of 4 to 10. In addition to the design and implementation of architectures, this project is developing both high and low level programming tools for the machines. Finally, algorithms for solving compute intensive problems on these machines are being developed and tested.

Brown University; Harvey F Silverman: *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A005); $16,964.

This is a joint project between Brown and Rutgers to build a large microphone array with associated processors for beam forming. Applications are direction finding, echo cancellation, and speaker differentiation in telconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

Other

National Institute of Standards and Technology; Judson French: ***Purchase of Optoelectronic Prototypes by the U. S. Broker - Interagency Transfer*****; (ECS-9530730 A001); $50,000; 12 months; (Support from: Experimental Systems Program, Quantum Electronics, Plasmas, and Electromagnetics - Total Grant $150,000).**

The Joint Optoelectronic Project (JOP) is a joint program entered into by the U. S. and Japanese Governments to provide prototyping services for experimental optoelectronic devices and modules. The goals of the project are to establish a joint research project in optoelectronics to improve the availability of novel prototype optoelectronic materials, devices, circuits and modules by providing users with access to leading edge optoelectronic fabrication facilities, to stimulate research activities in optoelectronics for computing in both the U. S. and Japan and to encourage effective commercialization. The funds provided in this award will be used to help support U. S. user participation in the JOP. It is anticipated that the novel, custom-designed optoelectronic devices and modules obtained using the broker facility will advance device and systems research activities within the academic community.

Salk Institute for Biological Studies; Terrence J Sejnowski: ***Workshop on Neuromorphic Engineering; June 25, 1995 - July 8, 1995; Telluride, Colorado*****; (IBN-9511637 A001); $5,000; 12 months; (Support from: Interactive Systems, Experimental Systems Program, Neuroengineering, Computational Neuroscience, Computational Biology Activities - Total Grant $48,880).**

Recently a new field of engineering has emerged, referred to as neuromorphic engineering, that is based on the design and fabrication of artificial neural systems, such as vision chips, head-eye systems, and roving robots, whose architecture and design principles are based on those of biological nervous systems. A two-week workshop on neuromorphic engineering will be held to bring together young investigators and more established researchers from academia with their counterparts in industry and national laboratories, working on both neurobiological as well as engineering aspects of sensory systems and sensory-motor integration. Formal lectures will be given, but the primary focus will be hands-on experience with research tools for all participants. The workshop will serve as a bridge between the engineering world of artificial neural systems and the neuroscience community. The workshop will provide an environment for intensive interactions between members of these two communities - merging engineering principles with experimental results from neuroscience. The interaction of these two disciplines should have significant impact both on the development of new technologies (new artificial neural systems) and on our understanding of how the nervous system is designed.

University of Washington; Lawrence Snyder: ***Conference on Experimental Research in Computer Science, June 20-21, 1996, Arlington, Virginia*****; (MIP-9634389); $33,681; 6 months.**

This award supports a conference to assess the current status and future prospects for experimental computer science research in the United States. The goals are to highlight progress in experimental research, largely using selected projects funded by the ES program, to discuss issues of importance to the fields practitioners in panel sessions, and to further an awareness of experimentation among policy makers. Researchers, policy makers, administrators, and leaders from academia, industry, and government will attend. A proceedings will be published.

Prototyping Tools and Methodology
(Formerly System Prototyping and Fabrication)

Dr. Michael Foster. Acting Program Director
(703) 306-1936 mfoster@note.nsf.gov

The Program

Vision and Goals

The Prototyping Tools and Methodology (PTM) program addresses rapid prototyping technologies. The goal is to develop and integrate the tools and technologies needed for rapid and efficient design and fabrication of products, processes and systems. The PTM program consists of three principal elements. The first (systems prototyping) supports fundamental research on rapid system prototyping methodologies, tools, environments, etc. with particular interest in the informational infrastructure needed for prototyping systems (virtual prototyping). The second (microfabrication) supports research related to advancing the state of the art in the modeling and control aspects of the fabrication (physical prototyping). The third element (education) provides assistance to microelectronics education through support of MOSIS and administrative oversight for its involvement in microfabrication for educational institutions.

Strategies

Strong interaction exists with the educational VLSI community through the MOSIS service. Every year approximately 5000 students in over 180 classes design custom VLSI chips to be fabricated through the MOSIS service. Special workshops are conducted at NSF, such as the 1996 Workshop on Integration of Education and Research in Microelectronics, to monitor and refocus the educational MOSIS program. Additionally, a series of workshops on emerging research in advanced manufacturing are sponsored. The 1995 NSF Workshop on Structured Design Methodologies for Micro-Electromechanical Systems (MEMS) brought together outstanding researchers from university and industry to address research directions.

Key Research Problems

Initiatives in this program include research in rapid prototyping for advanced manufacturing. New tools and technologies for virtual prototyping coupled with innovative services and an updated infrastructure that allows distributed rapid prototyping over high speed networks are of particular interest. Projects include new languages for machine and process design, and CAD/CAM integration, as well as projects encompassing modeling, simulation, model validation, and design tools and techniques. Research dealing with the application of useful VLSI design paradigms to SFF (Solid Free-form Fabrication) and MEMS (Micro-ElectroMechanical Systems) are also addressed.

Awards

Prototyping

Microelectronic

University of California - San Diego; Chung-Kuan Cheng; *Research on Rapid Prototyping Systems Using Field Programmable Devices*; (MIP- 9529077); $175,479; 12 months.

Rapid prototyping systems composed of programmable gate arrays present many technical challenges affecting system utilization and performance. The objective of this research is the development of engineering methodologies which address these currently unresolved technical issues. Current rapid prototyping systems suffer several shortcomings, including engineering bottlenecks caused by time- consuming mapping processes, inadequate performance for real-time emulation of high-speed systems, and high cost resulting from poor system utilization due to the severe /IO limit of programmable devices. The proposed research targets these key areas. Both hardware and CAD algorithm issues will be addressed. The proposed project will contribute to identify the bottlenecks of current prototyping systems and discovering novel solutions. This research will contribute to faster mapping processes, low cost - high utilization systems, and real-time emulation.

University of California - Santa Barbara; Malgorzata Marek-Sadowska: *Research on Layout and Logic Design*; (MIP-9419119 A001); $40,000; 12 months; (Support from: Prototyping Tools and Methodology Program, Design Automation Program - Total Grant $143,000).

The proposed research is on layout driven synthesis, i. e. the intersection of logic synthesis and physical design. The focus is on restructuring logic networks in synthesized digital systems. Four topics, which meet the goals of improving routing efficiency or power consumption, are being investigated. These are:

4. Incremental logic resynthesis to control wiring,
5. Coupling wiring with logic restructuring and finding optimizations to eliminate wiring overflows.
6. Use of generalized Reed-Muller forms to analyze logic as an aid to: - designing cell libraries and for technology mapping, - developing new multi level optimization techniques, - designing networks of provably good testability.
7. Develop new methods for power optimization, at the technology independent and technology dependent levels in logic synthesis, and also find better routing tools to handle power constraints.

University of Florida; Mark E Law:*PFF: A Multi-disciplinary Approach to IC Process Modeling Using the SUPREM-IV Modeling Tool*; (MIP-9253735 A004); $100,000; 12 months.

This multidisciplinary research focuses on the development of silicon models for point defect behavior, which are vital to understanding dopant diffusion. Models are being developed and parameterized for the effect of silicidation and stress on point defect kinetics. These models are then implemented in SUPREM-IV, a standard integrated circuit process modeling tool that utilizes advanced finite element techniques. This is a Presidential Faculty Fellow (PFF) Award initially awarded in Fiscal Year 1992.

Southern Illinois University - Carbondale; Ralph Etienne-Cummings:*CAREER: VLSI Implementation of Computation Sensors for Visual Information Processing*; (MIP-9624141); $200,000; 48 months.

This proposal describes the design and implementation of VLSI vision systems for real-time dynamic tasks. The two part career plan develops smart application specific and general purpose visual computational sensors which may be used with other intelligent systems to realize solutions to various robot vision and image processing problems. The three part research tasks are:

1. the development and implementation of compact VLSI algorithms for image acquisition, edge detection and motion detection;
2. applications of these algorithms to specific problems such as image velocitometry, target tracking and rudimentary line-following auto-navigation and
3. development of a general purpose VLSI computational sensor which performs image acquisition, edge detection and 2D optical flow measurement in a single compact system.

The goals of the educational component of this proposal include inclusion of VHDL tools into introductory undergraduate digital design courses, and the development of VLSI courses in imaging technology and mixed digital/analog signal processing with fabrication and testing requirements.

University of Illinois - Urbana-Champaign-Urbana; Farid N Najm:*CAREER: Reliability Engineering for Integrated Circuits*; (MIP-9623237); $200,000; 48 months.

A methodology is to be developed which allows designers to do reliability prediction and reliability budgeting using a high-level functional description of the design. Prediction models at the gate and function level will be developed and verified using observed data. Under reliability budgeting a specified failure rate, or mean-time-to-failure, for the whole chip may be partitioned to provide reliability targets for smaller pieces and thus lower levels of abstraction. The work leads to the development of a new course in "Reliability Engineering for Integrated Circuits", which will provide a software environment where ICs may be subjected to stress test experiments in virtual environment at various design stages using observational data.

University of Notre Dame; Jay B Brockman: *CAREER: Concurrent Optimization of Cell Libraries and Fabrication Processes*; (MIP-9625152); $200,000; 48 months.

This research proposes to develop techniques for concurrent optimization of both the circuit cell libraries and the fabrication process. A key advantage of the proposed approach is that it allows both the circuit designer and fabrication process engineer to work towards the common goal of maximizing yield, which will in turn reduce product cost. A second advantage is that through concurrent engineering, this approach can significantly reduce product development time. Provisions have been made for validating proposed techniques in an industrial setting. The research component of this proposal is also complemented by an extensive education plan for integrating studies in concurrent design into the undergraduate curriculum.

University of Kentucky; Charles E Stroud: *RIA: Self-Testing, Diagnostic and Repair Configurations for Field Programmable Gate Arrays*; (MIP-9409682 A001); $10,000.

This research is directed at system level testing by developing algorithms and circuit configurations which can be loaded into FPGA's to provide self-testing, fault diagnostic and repair capabilities. Once faults in an FPGA have been detected and identified, reconfiguration algorithms would then be used to remap the system function into the FPGA while avoiding the existing faults. It is conjectured that by using this approach, development of application-specific system diagnostic software for testing the FPGA system function can be avoided.

University of Maryland - College Park; K. J. Ray Liu: *NYI: High Performance Computing for Signal Processing*; (MIP-9457397 A002); $31,250; 12 months; (Support from: Prototyping Tools and Methodology Program, Signal Processing Systems Program - Total Grant $62,500).

There are three major architectural models used in high-performance signal/image processing:

1. VLSI- signal processing - high- throughput VLSI architectures for low-cost application-specific implementations used in applications such as communication systems, speech, video/HDTV, and radar;
2. parallel signal processing on massively parallel computers - parallel algorithms for complex signal/imaging systems used in computer vision, medical imaging, and the processing of vast amounts of data in deep space exploration; and
3. distributed signal processing on high-speed networks - used in applications such as document image processing, multimedia, automatic signal processing in manufacturing, and medical signal/image processing.

This research focuses on the development of efficient algorithms and architectures for each architectural model, and in comparative studies of the advantages and disadvantages of these different computing schemes. The goal is to investigate which signal/image processing problems can be carried out optimally under different computing and communication schemes.

University of Maryland - College Park; Linda Milor: *A Statistical Modeling Methodology for Submicron MOS Devices and Circuits*; (MIP-9501912 A001); $10,000.

The performance of an integrated circuit depends on the designers' choices and the process engineers' decisions during the manufacture of the device. Typically the choices and decisions are independent of each other with little awareness of the interplay of the two engineering domains.

This research develops methods for raising the designers' awareness of the effects of circuit implementation parameters and to sensitize process engineers to the impact of processing decisions on the performance of the final circuit. A statistical

model is developed to correlate device and process parameters, primarily to establish circuit performance variations on spatial and process dependencies in replicated circuits. Circuit simulators are used to solve the resultant (huge) set of nonlinear differential equations.

Michigan Technological University; Ashok K Goel, Esther T Ososanya: *Experimental Validation of Interconnection and Transistor Delay Models for the GaAs-Based Integrated Circuits*; (MIP-9223989 A004); $10,000; 12 months.

During the last few years, GaAs technology has emerged rapidly from basic research to device and circuit development. It is crucial to know the expected propagation delays in an integrated circuit before it is fabricated. To meet this objective, numerical models have been developed that address crosstalk and propagation delays in the parallel and crossing VLSI multilevel interconnections as well as for the transverse delays in the GaAs MESFETs and GaAs/AlGaAs MODFETs. In addition to determining the crosstalk and propagation delays, the models can be utilized to achieve the optimization of the device and interconnection dimensions and other parameters for minimum crosstalk and delays. Validation of these numerical models by comparison of the modeling results with actual experimental observations is critical if they are to be incorporated into GaAs CAD tools. This research effort focuses on the following set of objectives:

1. design and fabrication of several GaAs-based logic circuits to retain the ability to alter the various design parameters;
2. application of the interconnection and the GaAs MESFET delay models recently developed for the determination of propagation delays in these GaAs-based logic circuits;
3. experimental measurements of propagation delays in these circuits and comparison with developed delay models;
4. modification of the interconnection and transistor delay models, as required; and
5. experimental validation of the final models.

Cornell University; Geoffrey M Brown: *Data Acquisition Systems with User Configurable Hardware*; (MIP-9530811, A001); $162,295; 24 months.

This project develops a new class of data acquisition systems with user configurable hardware. The primary development effort is the creation of software tools which will enable scientists and engineers who are not skilled hardware designers to utilize such systems. Some of the primary research issues are the development of techniques for specifying and analyzing timing behavior between fixed and configurable hardware and the development of techniques for designing robust interfaces.

A typical application of this work is a reconfigurable data acquisition card consisting of a fixed analog section and a configurable digital section built from one or more field programmable gate arrays (FPGAs) and static RAM chips. Existing data acquisition cards are generally used by scientists and engineers who are not skilled hardware designers; however, the currently available tools for FPGAs require significant digital design experience. For a reconfigurable data acquisition card to be useful, it is necessary that the configuration tools accept as input high level programming language which is accessible to a competent programmer. Any custom hardware necessary to interface the analog sections or host CPU must be provided in the form of libraries which can be conveniently linked to the user's code.

State University of New York - Binghamton; Jiayuan Fang: *NYI: Analysis and Modeling of High-Speed Interconnects in Electronics Packaging*; (MIP-9357561 A003); $74,000; 12 months.

This research is concerned with the analysis and modeling of electrical performance of high-speed interconnects in electronics packaging. The finite-difference time-domain (FDTD) method, which is a full-wave solution of Maxwell's equations in three dimensions, is used to simulate signal propagation through interconnects. Topics pursued are:

1. Development of a computational scheme for conformed finite-difference grid to model complex-shape interconnects. The objective of this scheme is to enhance the resolution and accuracy of numerical solutions while maintaining the computation efficiency associated with the regular rectangular finite-difference grid.
2. Analysis and modeling of electrical properties of interconnection discontinuities in electronics packaging. Issues involved in this topic include: modeling of electrical characteristics of interconnection discontinuities over the frequency range from dc to tens of gigahertz; evaluation of impacts of parasitics associated with interconnection discontinuities on the propagation of high-clock rate signals; and development of design guidelines for typical interconnection discontinuities in high performance electronics packaging.

North Carolina State University - Raleigh; Paul D Franzon: *NYI: Advanced Interconnect and Display Approaches*; (MIP-9357574 A003); $62,500; 12 months.

The primary focus of this work is to resolve issues dealing with the design and implementation of high bandwidth reconfigurable interconnect systems based on Micro ElectroMechanical Systems (MEMS) which are commonly referred to as Micromachines. Different guided-wave optical and holographic free-space optical switch elements are being implemented and compared in terms of bandwidth, switch reconfiguration rate, and range of programmability. Application to data switching, and programmable interconnect devices for rapid prototyping are being addressed with attention to both technological and system-wide performance/cost design issues. Comparisons are made with conventional technologies. Also being investigated is the application of some of the optical MEMS elements to advanced image projection.

University of North Carolina - Charlotte; Dian Zhou: *NYI: Performance-Driven VLSI Designs*; (MIP-9457402 A002); $32,750; 12 months; (Support from: Prototyping Tools and Methodology Program, Design Automation Program - Total Grant $65,500).

Three research in high-performance VLSI system design are being addressed:

1. how to relate the system performance function, characterized by electrical parameters, to the geometrical parameters of the VLSI physical design;
2. how to model performance driven VLSI physical designs based on given technology and computational capability; and
3. how to characterize the fundamental computational aspects of modeled problems and develop effective algorithms for solving them.

A distributed RLC circuit model for interconnects is being designed. It considers: non-monotone circuit response, coupling effects among the signal lines, and low energy consumption. Efficient computation methods that solve time-varying Maxwell equations using the adaptive wavelet collocation method are being devised. The algorithms and methods are being included in a CAD system.

University of Washington; Susan Eggers; *Simultaneous Multithreading*; (MIP- 9632977); $3318,305; 24 months; (Support from: Experimental Systems Program, Prototyping Tools and Methodology - Total Grant $519,164).

This is a project to develop and extend support for multithreaded computation on superscalar processors. The intention is to provide high processor utilization despite increasing memory latency. In simultaneous multithreading, multiple independent threads will issue instructions to a superscalar processor's functional units in a single cycle. The objective of this project is to fully define the low-level architecture for simultaneous multithreading, including scheduling, register architecture, pipeline design, and speculative execution. The project is developing compiler and operating systems support for multithreaded execution on the new processor architecture. A primary goal is to maximize multiple thread performance without increasing single thread latency. Most measurements are from detailed simulations, using simulators and compilers provided by industrial partners.

MEMS, SFF, MCM

University of California - Berkeley; Carlo H Sequin, Paul K Wright: *Rapid Prototyping Interface for 3D Solid Parts*; (MIP-9632345); $301,123; 24 months.

It is proposed to develop a simple and clear language for use as a digital interface in Solid Free-Form Fabrication. This language is called "SIF" for "Solid Interchange Format" inspired by CIF, the Caltech Intermediate Form for LSI layout description. The role of SIF would be to describe the ideal solid part desired, in a fabrication-process-independent way. It is further proposed to create a foundation for some of the key tools needed with such an exchange language, in particular, a robust slicing tool to create 2.5D layers needed for process planning, as well as checker programs that will verify that the shape specified in the SIF file is indeed a closed solid, has the required topological and

geometric properties, and specifies only realizable features containing certain minimum feature sizes and separations.

Finally, it is proposed to develop test structures of various kinds that will test the robustness of analysis and conversion tools working from SIF input for different SFF fabrication processes and to monitor the quality of individual production runs.

University of California - Los Angeles; K J Pister: *CAREER: MEMS for the Masses*; (MIP-9624124); $135,000; 48 months; (Support from: Prototyping Tools and Methodology Program, Microelectromechanical Research - Total Grant $200,000).

A method is proposed for the decoupling of design and manufacturing of micro-electromechanical systems (MEMS) and devices. A design infrastructure is to be developed that will allow non-experts to quickly design and simulate MEMS devices and systems. The philosophy follows the VLSI paradigm, where there is a clear separation between the wafer fabrication and the design effort that creates the pattern which is to be implemented.

MEMS software tools, device models, and educational materials will be made available on the Internet. These tools will include schematic editing, layout synthesis, parameter extraction and simulation.

The proposal also contains a strong education plan extending the PI's successes in recruiting undergraduates, including women and under-represented minorities, to become research students.

University of Central Florida; Aravinda Kar, Amar Mukherjee: *Design Methodology for Laser-Aided Direct Rapid Prototyping (LADRP)*; (MIP-9625752); $269,187; 24 months.

The proposed study is concerned with a one-step direct process to fabricate three-dimensional parts using a laser beam. This process is referred to as the Laser-Aided Direct Rapid Prototyping (LADRP). Laser light can melt and/or vaporize virtually any material and laser tooling leads to the unification of various conventional tools because a laser beam can be used to carry out many different machining operations such as cutting, drilling, welding and milling. Such process integration is very important for the rapid production of three- dimensional parts directly in near net shape.

The success of the proposed rapid prototyping process depends on the development of systematic design methodologies and developing suitable digital interfaces between design and manufacturing. The proposed research will carry out experimental studies and develop mathematical models for the LADRP process, and will develop design rule characterization of the LADRP process so that lower- level CAD tools for process planning simulation and verification may be implemented. A layered description language will be developed that can generate the slices of 2.5-D layers for the part from an input file specified in some higher level description language such as the STL or ACIS format.

Carnegie-Mellon University; Gary K Fedder:*CAREER: Schematic Design for Microelectromechanical Systems*; (MIP-9625471); $200,000; 48 months.

The proposed research deals with developing tools for the design of suspended micromechanical systems. A Schematic design methodology is to be devised which complements and builds upon existing design tools, chiefly commercial CAD packages that handle 3D visualization, finite-element analysis, macromodeling and simulation. The work includes: general design methodology; schematic design, including validation; identifying reusable components; generalized component interface; creating a set of lumped-parameter component models and design, fabrication and test of prototype components.

University of Pennsylvania; Ruzena K Bajcsy, Dimitri Metaxas, Vijay Kumar, Daniel K Bogen: *Rapid Prototyping of Rehabilitation Aids for the Physically Disabled*; (MIP-9420397 A002, A003); $92,180; 12 months; (Support from: Prototyping Tools and Methodology Program, CISE Institutional Infrastructure - Total Grant $342,180).

The methods of rapid prototyping are ideally suited to rehabilitation devices. Because each person requires unique performance and function in a rehabilitation device, devices specific to each person must be rapidly designed and produced. This project is investigating a completely integrated approach to the design and prototyping of passive mechanical rehabilitation devices. The approach involves: the quantitative assessment of the form and performance of human limbs; the design of the assistive device; evaluation of the device using virtual prototyping; feedback from the consumer and therapist; actual prototyping of the device; evaluation of the function and performance of the device; and redesign based on performance. The contributions of the product

include: the development of new computer-based tools for the assessment of human performance; a manufacturing technique for a new class of hyperelastic materials; the integration of tools into a rapid prototyping system for rehabilitation devices; and development of mechanisms for systematic evaluation of the final product.

West Virginia University; Lawrence A Hornak: ***NYI: Cointegrated Polymer Waveguide Optical Interconnections for Wafer-Level MCM Systems*****; (MIP-9257101 A004); $62,500; 12 months.**

Motivated by the need for robust polymers optically superior to polyamides yet suitable for cointegration of optical interconnection waveguides directly with the active CMOS interconnection substrate of advanced multi-chip modules (MCMs), the research seeks to fabricate the first optical waveguides with polyphenylsilsesquioxane (PPSQ), a spin castable, low temperature processed silicon ladder polymer used as a thin and thick film dielectric for microelectronics. This material potentially offers the low loss obtainable with less stable optical polymers while offering thermal stability and process latitude (patterning, wet, dry etching) exceeding that of polyamide. This research assesses the suitability of PPSQ together with planarizing and superstrate polymer layers for providing a multilayer system supporting fabrication of high density waveguide arrays directly over the Si devices of emerging active MCM substrate.

Education

California Institute of Technology; Erik K Antonsson: ***A Workshop on Structured Design Methods for MEMS; November 12-15, 1995, Pasadena, California*****; (MIP-9529977); $48,464; 12 months.**

This is a proposal to hold a workshop involving about forty leading researchers from academia and industry to discuss:

1. the importance and value of Structured Design Methods for MEMS;
2. the research opportunities in this area, and,
3. the likely increase in effectiveness of MEMS design utilizing these methods.

Earmark Inc; Elisabeth Perez-Luna: ***Intersections: Art, Science and Technology -- A Radio Project*****; (ESI-9634183); $8,000; 12 months; (Support from: Informal Science Education, Prototyping Tools and Methodology Program, Signal Processing Systems Program - Total Grant $25,000).**

This project will examine the implications of the intersection of art, science, and technology as revealed by the events occurring around the celebration of the fiftieth anniversary of the first general electronic computer, ENIAC-50. The finished product will be radio programming and audio material that presents the coming together of scientists, artists, and other participants in Philadelphia. It will build on the previous research and explorations in both of these fields toward finding commonalties and ways in which each has influenced the other. The finished product will be broadcast for the general public and also can be made available to scientists and educators in the field as resource material for courses, seminars, and lectures as they explore the intersections of science and art and implications this has for research and education.

Advanced Research Projects Agency; Robert F. Lucas; NSF/ARPA Agreement for Use of ARPA VLSI Implementations; (MIP-9419682 A001); $350,000; 12 months; (Joint support with the Microelectromechanical Research Program, the Manufacturing Machines and Equipment Program, and the DUE Course and Curriculum Program - Total Grant $800,000).

A 1994 Memorandum of Understanding (MOU) between the National Science Foundation (NSF) and the Advanced Research Projects Agency (ARPA) extended a three-year joint program supporting VLSI (Very Large Scale Integration Fabrication) by MOSIS (Metal Oxide Semiconductor Implementation Service) for qualifying educational institutions. The continuation of the MOU expands the original program to accelerate critical capabilities for Microsystems Design and Prototyping in U.S. educational institutions. This includes expanding the original services and technologies available to schools authorized to use MOSIS as they become generally available and cost-effective (e.g., semi-custom and gallium arsenide chip fabrication), stressing VLSI education, especially undergraduate education needed for designing future electronic system; exploring new fabrication services designed specifically for the research and education community's desire for cost effective experimentation of state-of-the-art technologies (e.g., design of advanced sensor and micro-mechanical devices); and rapid prototyping methodologies, tools, and services needed for complete systems.

Index of PYI, NYI, CAREER, and PFF Investigators

CAREER	Faculty Early Career Development
NYI	NSF Young Investigator
PYI	Presidential Young Investigator
PFF	Presidential Faculty Fellow

Index of Principal Investigators

Index of Institutions

SPINE PRINTING