DOCUMENT RESUME

ED 259 943 SE 045 918

TITLE Division of Computer Research Summary of Awards, Fiscal Year 1984.
INSTITUTION National Science Foundation, Washington, DC. Directorate for Mathematical and Physical Sciences.
REPORT NO NSF-84-77
PUB DATE 84
NOTE 93p.
PUB TYPE Reports - Descriptive (141)

EDRS PRICE MF01/PC04 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Computer Oriented Programs; *Computers; *Computer Science; Databases; *Grants; Higher Education; Mathematics; Program Descriptions
IDENTIFIERS National Science Foundation

ABSTRACT

Provided in this report are summaries of grants awarded by the National Science Foundation Division of Computer Research in fiscal year 1984. Similar areas of research are grouped (for the purposes of this report only) into these major categories: (1) computational mathematics; (2) computer systems design; (3) intelligent systems; (4) software engineering; (5) software systems science; (6) special projects, such as database management, computer-based modeling, and privacy and security of computer systems; (7) theoretical computer science; (8) computer research equipment; and (9) coordinated experimental research. Also included are presidential young investigator awards and awards for small business innovation research. Within each category, awards are listed alphabetically by state and institution. Each entry includes the grantee institution, name of the principal investigator(s), project title, award identification number, award amount, award duration, and description. (This report lists fewer than the actual number of projects currently receiving support because the duration of some grants exceeds 1 year.) (JN)

DIVISION OF COMPUTER RESEARCH

SUMMARY OF AWARDS

Fiscal Year 1984

NATIONAL SCIENCE FOUNDATION
Directorate for Mathematical and Physical Sciences
Washington, DC 20550

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 and minorities to compete fully in any of the research and research-related Programs described in this document.

In accordance with Federal statutes and regulations and NSF policies, no person on grounds of race, color, age, sex, national origin, or physical handicap 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.

NSF has TDD (Telephonic Device for the Deaf) capability which enables individuals with hearing impairments to communicate with the Division of Personnel Management for information relating to NSF programs, employment, or general information. This number is (202) 357-7492.

Preface

The purpose of this report is to provide the scientific community with a summary of those grants awarded by the Division of Computer Research in Fiscal Year 1984. The report lists fewer than the actual number of projects currently receiving support because the duration of some grants exceeds one year.

Similar areas of research are grouped together for reader convenience. The reader is cautioned, however, not to assume these categories represent either 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 purposes 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, awards are listed alphabetically by state and institution.

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

Kent K. Curtis
Division Director
Division of Computer Research

*To get an idea of the scope of possible programs, consult *What Can Be Automated? The Computer Science and Engineering Research Study (COSERS)*, Bruce Arden (Ed.), MIT Press, Cambridge, MA (1980).

Table of Contents

INTRODUCTION

This report provides summaries of awards made by the National Science Foundation (NSF) through the Division of Computer Research in Fiscal Year 1984.

In Fiscal Year 1984, the Division of Computer Research Award Activities were carried out through eight programs. Topics of interest to the Foundation include, but are not limited, to the following:

Computer Systems Design

Principles of computer systems design relating to the structure of computer systems or the process of systems design. Topics include, but are not limited to: computer system architecture; distributed computer systems; integrated hardware/software systems; performance measurement and evaluation; fault tolerant systems; logic design; computer graphics; man-machine interaction; and VLSI design methodology. The scope of this program includes experimental implementation where that is an integral part of the research.

Intelligent Systems

Computer-based systems that have some of the characteristics of intelligence. Relevant areas of inquiry include: knowledge engineering; automated theorem-proving; mechanical inferencing; problem solving; pattern analysis; computer vision; natural language and speech understanding; and areas related to the automatic analysis and handling of complex tasks.

Software Engineering

The structure and design process of computer software, especially verification, testing, portability, reliability, and human interfacing to numeric and non-numeric software systems. Areas of emphasis include: program validation and testing; software tools; and human factors in software design and use. The program also supports research in computationally oriented numerical analysis, the design and construction of high quality portable software for scientific research, and experimental implementation where that is an integral part of the research.

Software Systems Science

Conceptual basis for the specification of future software systems and the necessary experimentation with such systems, including: advanced programming languages and optimizing compilers; functional and relational specification; program transforming systems; systems for the verification and proof of correctness of programs; the study of the concurrency of operations; the discovery of new algorithms and improved measures of effectiveness of known algorithms.

Special Projects

New directions in computer science and applications including computer networks; database management; computer-based modeling; privacy and security of computer systems; and other topics of special interest in computer science.

Theoretical Computer Science

Theories of computation and formal languages; analysis of algorithms; theoretical models for computation; and other theoretical problems concerned with the foundations of computer science.

Computer Research Equipment

Support of special purpose equipment needed by more than one computer research project and difficult to justify on a single project.

Coordinated Experimental Research

Experimental Computer Research — The establishment and enhancement of experimental research facilities, needed technical and professional support personnel, and necessary maintenance of facilities; large multi-investigator projects in experimental computer research (5 year awards expected).
Computer Science Research Network (CSNET) — Build a logical computer-based communication network spanning ARPANET, public packet networks, and PHONENET, a telephone-based relay system; provide communication services, file transfer, and access to remote databases; stimulate new classes of network research activities.

Additional Information

For additional information on any of the projects, contact the principal investigators directly.

Division of Computer Research

Kent K Curtis
Division Director
(202) 357-9747

Harry G. Hedges
Senior Staff Associate
(202) 357-7349

Program or Area	Program Director	Telephone
Computational Mathematics	Martha A. Branstad	(202) 357-7345
Computer Research Equipment	John R. Lehmann	(202) 357-7349
Computer Systems Design	John R. Lehmann	(202) 357-7349
Coordinated Experimental Research	Bruce H. Barnes	(202) 357-7375
CSNET	Bruce H. Barnes	(202) 357-7375
Intelligent Systems	Su-Shing Chen	(202) 357-7345
Software Engineering	Martha A. Branstad	(202) 357-7345
Software Systems Science	Thomas A. Keenan	(202) 357-7375
Special Projects	Lawrence H. Oliver	(202) 357-7375
Theoretical Computer Science	Vacant	(202) 357-7349

November 1984

SUMMARY

	Number of Projects	Value of Awards
Computational Mathematics	28	$934,348
Computer Systems Design	43	$3,331,793
VLSI Design Methodology	6	411,620
Computer System Architecture	24	1,524,017
Fault Tolerance and Reliability	4	385,291
Computer System Performance Measurement and Evaluation	4	294,238
Computer Graphics	5	716,624
Intelligent Systems	55	3,247,500$
Computer Vision and Image Processing	16	772,213
Natural Language and Signal Understanding	11	542,756
Automatic Theorem Proving, Concept Learning, and Inference	8	640,476
Knowledge Representation and Problem Solving	20	1,292,055
Software Engineering	35	$2,412,524
Quality Software	5	360,654
Symbolic and Algebraic Manipulation	4	191,233
Software Studies and Metrics	6	149,042
Program Testing and Verification	4	319,217
Software Tools and Programming Environments	16	1,392,378
Software Systems Science	36	$3,556,868
Programming Languages: Design, Semantics, Implementation	6	395,423
Operating Systems and Concurrency	5	575,670
Data Bases	7	700,736
Programming Methodology	11	1,241,662
Distributed Computing	4	319,860
Verification	3	323,517
Special Projects	21	$918,800
Data Handling, Data Manipulation, Database Research	3	53,323
Privacy & Security	5	302,400
The Societal Impact of Computing	1	99,868
Modelling & Simulation	1	36,702
Networking, Computer Communications, & Distributed Computing	7	382,873
Other Projects, Symposia, Colloquia, and Studies	4	43,634
Theoretical Computer Science	78	$3,600,564
Concrete Complexity and Analysis of Algorithms	29	1,481,902
Parallel, Distributed, and Data Base Models	22	1,322,873
Foundations of Computer Science	21	577,815
Automata and Language-Based Models	6	217,974

	Number of Projects	Value of Awards
Computer Research Equipment	20	$1,389,438
Coordinated Experimental Research	23	$13,500,007
Experimental Computer Research	20	12,738,496
Computer Science Research Network (CSNET)	3	761,511
Presidential Young Investigator Awards	13	512,414
Small Business Innovation Research	3	-

Note: The above amounts are those supported by the Division of Computer Research and do not reflect the total of the awards. Refer to the individual summaries for details.

Computational Mathematics

California Institute of Technology; Heinz-Otto Kreiss; *Numerical Methods for Nonlinear Ordinary Differential Equations (Mathematical Sciences and Computer Research)*; (DMS-8312264); $32,734; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $93,201).

The research will focus on analytical studies and numerical experiments to further develop robust and efficient procedures. The class of systems of ODE's are those that are classified as essentially nonoscillatory stiff ODE's and high oscillatory stiff ODE's. A prototype computer code will be further developed for application to realistic problems such as reentry-roll resonance.

University of California - San Diego; James R. Bunch; *Stability of Matrix Computations for Toeplitz Systems (Mathematical Sciences and Computer Research)*; (DMS-8318412); $14,456; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $28,912).

In this project Professor Bunch will study the stability of matrix computation for Toeplitz matrices, that is, matrices where the elements on each diagonal are all equal. The solution of such matrices arise out of problems involving; time series analysis, linear filtering, image processing, control theory, partial differential equations and approximation theory.

Professor Bunch will study algorithms for $n \times n$ matrices that attempt to improve on the classical work estimate of n^2 operations by using a QR decomposition.

University of California - Santa Barbara; Alan J. Laub; *Algorithms, Analysis, and Software for Riccati Equations*; (ECS-8406152); $55,000; 36 mos. (Joint support with the Systems Theory and Operations Research Program - Total Grant $225,000).

One of the most deeply studied nonlinear matrix equations arising in mathematics and engineering is the Riccati equation. This equation arises naturally in a variety of engineering problems and its use in systems and control theory has been well established. The term "Riccati equation" will mean any of a class of matrix quadratic algebraic or differential or difference equations of symmetric or non-symmetric type arising in the study of continuous-time or discrete-time dynamical systems. One aspect of Riccati equations that has received increasing attention in the recent past, is effective algorithms for their numerical solution in the finite arithmetic environment of a digital computer. The basic thrust of the proposal is the numerical solution of general classes of matrix Riccati equations and their related problems. Algorithms based on certain types of matrix pencils and associated generalized eigenvalue problems are stressed. Further, the results of the proposed research will be implemented as robust mathematical software.

University of Colorado - Boulder, Richard H. Byrd; *Trust Region Methods for Minimization (Computer Research)*; (DCR-8403483); $45,560; 12 mos.

This research is concerned with constrained and unconstrained optimization, solving systems of nonlinear equations, and nonlinear least squares. Four topics will be investigated. First, the development of tensor methods for nonlinear equations and their extension to unconstrained optimization will be continued. These methods show great promise of yielding general purpose algorithms that are highly efficient on singular and nonsingular problems. Second, practical local methods for constrained optimization that are one-step q-superlinearly convergent under realistic assumptions will be developed. Third, global methods for constrained optimization that are practical and satisfy global convergence theory and that perform robustly even from poor starting values will be developed. Finally, algorithms for several important generalizations of nonlinear least squares not addressed by currently available software will be investigated. These include data fitting problems where there are errors in the independent variables, modeling problems where the model is implicit, and curve fitting problems where an orthogonal measure of the distance from the data points to the fitting function is desired.

University of Georgia; Michael Brannigan; *Numerical Methods for Constrained Approximation (Mathematical Sciences and Computer Research)*; (DMS-8319727); $11,217; 24 mos. (Joint support with the Applied Mathematics Program - Total Grant $33,652).

Professor Brannigan will study the theory and numerical realization of constrained Chebyshev approximations. In his study of the discrete, his primary concern will be to develop stable and efficient methods. This will become the basis for good computer software that could be used for solving continuous problems by sequences of such discrete problems. Professor Brannigan will also study linear continuous

Chebyshev approximation. Here he will extend the ideas of "H-sets" to include the imposition of convex constraints on the approximation

Northwestern University; Jorge Nocedal; *Numerical Methods For Nonlinear Optimization*; (DCR-8401903); $41,096; 24 mos.

The project centers on numerical methods for solving nonlinear optimization problems. One part is concerned with using linear programming for solving large sparse sets of nonlinear equations. The goal is to develop a fast algorithm that will maintain the sparsity of the set of equations. Another part involves projected Hessian updating algorithms for nonlinearly constrained optimization. Current techniques for which convergence can be proven are very slow. Convergence analysis for faster algorithms will be conducted. It is expected that convergence can be proven for at least one of these.

University of Illinois - Urbana; Ahmed Sameh; *Parallel Numerical Algorithms*; (DCR-8117010 A02); $59,368; 12 mos.

The goal of this research is the development of high speed and high quality synchronous parallel numerical algorithms. These algorithms primarily concern:

1. Solving sparse symmetric and nonsymmetric systems of linear equations that arise from the finite difference or finite element discretization of partial differential equations, and
2. Solving sparse linear least squares problems in geodesy and image reconstruction. In designing these algorithms we specify the simplest configuration of the processor interconnection network that ensures high speedup without making the algorithm communication intensive. Furthermore, an automatic error analysis package will be used to check and compare the stability of these algorithms.

The project is concerned with both the computational and the error analytic aspects of such shape preserving algorithms. One specific project is to adapt the monotone surface fitting algorithm of the investigator and Z. Ziegler, to local monotonicity preservation. This will then be used to drive a customized contour plotting algorithm which will take advantage of the piecewise, total degree, quadratic nature of the fitted surface. Another project is to develop an adaptive scheme for the interpolation of data with a local convexity preserving surface. This would be done using splines with variable mesh lines.

Purdue University; Walter Gautschi; *Applied Orthogonal Polynomials (Computer Research a. . Mathematical Sciences)*; (DCR-8320561); $37,202; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $53,102)

The principal objective is a study of constructive methods for Gauss-type quadrature rules and, more generally, for orthogonal polynomials, assuming arbitrary weight distributions. An additional goal is the preparation of related mathematical software. The study of Gauss-type quadrature rules is to include Radau- and Lobatto-type formulae, as well as Gauss formulae with multiple nodes. In addition to developing methods for generating general orthogonal polynomials, specific problems connected with modifying the weight distribution are to be investigated.

University of Maryland; Ivo Babuska; *Higher Order Finite Element Methods and Adaptive Approaches (Mathematical Sciences and Computer Research)*; (DMS-8315316); $24,850; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $74,550).

In this project Professor Babuska will investigate adaptive order finite element methods (p and h-p versions) for partial differential equations arising in solid mechanics. This research will be conducted in close collaboration with investigators from the Center of Computational Mechanics, Washington University, St. Louis, Missouri.

Professor Babuska will provide guidance on numerical experiments using a software package, FIESTA (Software System for Static Analysis of Solid Structures Based on the p-Version of the Finite Element Method).

Massachusetts Institute of Technology; Ravindran Kannan; *Computational Complexity of Numerical Algorithms (Computer Research)*; (DCR-8304770 A01); $14,000; 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant, $41,148).

Improved algorithms for computing the Smith and Hermite normal forms of an integer matrix along with necessary empirical improvements are studied. Theoretical improvements and implementations of algorithms for solving the same problems for matrices of polynomials are also studied. The "alternation translation" technique is studied in order to obtain results about nondeterminism and determinism using the polynomial time hierarchy. Finally, several computational problems involving multivariate polynomials are attacked.

Massachusetts Institute of Technology; Roy E. Welsch; *Industry/University Cooperative Research Project: Nonlinear Exponential-Family And Bounded-Influence Regression*; (DCR-8116778 A02); $37,950; 12 mos. (Joint support with the

Division of Industrial Science and Technology Total Grant $37,950).

This is a collaborative research project with John Dennis of Rice University, David Gay of Bell Laboratories, and Roy Welsch of Massachusetts Institute of Technology. The goal of the project is to exploit the common structure of several nonlinear optimization problems. During the past year Gay and Welsch have completed and tested a new nonlinear regression code.

University of Minnesota; J. B. Rosen; *Global Optimization For Large Scale Problems Using Vector Processing*; (DCR-8405489); $50,297; 12 mos.

The objective of this research is to develop and implement efficient and robust algorithms for the computational solution of large scale global optimization problems. The problems considered include the minimization of a concave function subject to linear constraints, the minimization of an indefinite quadratic function subject to linear constraints, and minimum cost network flow problems with some concave costs. All these problems are combinatorial in nature. Large scale problems of this kind typically contain many variables which appear only linearly. In order to solve such large scale problems the algorithms must be designed to take advantage of the linearity by restricting the combinatorial aspects to a subspace of nonlinear variables.

There are many important application areas which can be formulated as problems of the type considered here. These include a wide class of scheduling, investment and planning models. More recently, important aspects of chip design and physical database design problems have been formulated in these terms.

The algorithms will be developed, implemented and texted, using a Class VI vector processing machine, so that thorough testing can be done in a reasonable amount of time. The algorithms and software will be developed to take advantage of the vector processing and parallel capabilities of such machines, so that they can be used efficiently on the new generation of supercomputers.

North Carolina State University - Raleigh; Robert Plemmons and Carl Meyer. *Numerical Linear Algebra (Mathematical Sciences and Computer Research*; (DMS-8219500 A01); $14,652; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $45,347).

Professors Plemmons and Meyer are applied mathematicians working in the area of numerical linear algebra. They are working on a variety of problems that arise in the study of M-matrices - a class of matrices that arise in contexts of solving partial differential equations and queueing theory.

During this support period, Professors Plemmons and Meyers have established necessary and sufficient conditions for the convergence of block Gauss-Seidel iterations associated with M-matrices. They have also extended their prior work on floating point factorizations of M and H-matrices to show that Gaussian elimination without pivoting can normally proceed without undue growth of the elements in the triangular factors for A.

Princeton University; John M. Mulvey; *Solving Generalized Networks With Convex Objectives*; (DCR-8401098); $60,293; 24 mos. (Joint support with the System Theory & Operations Research Program - Total Grant $90,293).

The two primary research objectives of this project are:

1. To integrate recent algorithmic advances in network optimization and nonlinear programming, and
2. To design efficient solution strategies for an important problem-class falling between these two areas.

The focus of attention will be on nonlinear networks with generalized arcs (losses/gains) and non-separable convex objectives —(NLGN).

This model is capable of representing numerous significant real-world problems, including electrical power generation, stochastic networks, air-traffic routing and financial cashflow.

Rutgers University; Michael D. Grigoriadis; *Coordinated Network Optimization System* (DCR-8113502 A02); $52,419; 12 mos.

This research involves the development and implementation of efficient algorithms for solving several classes of large, sparse network optimization problems. Investigation into how to best select and apply these algorithms for solving practical problems will be conducted. Rather than developing general-purpose algorithms for each class of network optimization problems, this approach will employ multiple algorithms, each directed to the most efficient treatment of a particular subclass. The identification of such subclasses of problems, the choice of the most appropriate algorithm and data structures, the generation of good solution strategies and the coordination of these tasks is a complex problem which has not been adequately studied.

University of New Mexico; Robert Russell; *Numerical Methods for Boundary Value Problems in ODEs, with Applications to PDEs (Mathematical Sciences and Computer Research)*; (DMS-8401019); $18,000; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $36,000).

Professor Russell will study a variety of numerical methods for boundary value problems for ordinary differential equations (BVODE's). The underlying approach will be to carefully analyze the associated discretization matrix factorizations. In this framework, he will derive improved stability results for invariant imbedding, study conditions of collocation matrices using spline bases, and re-interpret collocation and initial value methods for singular perturbations.

Professor Russell is a co-author with U. Ascher and J. Christiansen of COLSYS (a collocation software for BVODE's computer program) - a popular software code. In this project, Professor Russell will modify COLSYS to incorporate a more efficient linear system solver and a more robust spline basis representation. The BVODE techniques will be used to study initial/boundary value partial differential equations.

Columbia University; Donald Goldfarb; *Algorithms for Nonlinear Programming*; (DCR-8341408 A01); $8,943; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $17,886).

The purpose of this research is to develop, analyze, and test algorithms for nonlinear programming. Special emphasis will be placed upon the numerical stability of the algorithms and their ability to efficiently solve large sparse problems. As the ultimate goal of this research is the creation of effective and reliable software, implementational considerations will play an important role in the development as well as in the refinement of the algorithms. There are three classes of algorithms that will be investigated: successive quadratic programming algorithms for general nonlinear programs; ellipsoid and related algorithms for convex programs; and quadratic programming algorithms. Under the research opportunities for Small College Faculty Programs, Dr. Ashok Idnani will participate in this research.

Cornell University; Patricia Eberlein; *Norm-Reducing Methods for Algebraic Eigenproblems for Parallel and Micro Computation (Computer Science)*; (IPS-8410001); $6,000; 12 mos. (Joint support with the Division of Research Initiation and Improvement Program - Total Grant $62,141).

This project furthers VPW program objectives which are (1) to provide opportunities for women to advance their careers in engineering and in the disciplines of science supported by NSF and (2) to encourage women to pursue careers in science and engineering by providing greater visibility for women scientists and engineers employed in industry, government, and academic institutions. By encouraging the participation of women in science, it is a valuable investment in the Nation's future scientific vitality.

New York University; Michael L. Overton; *Numerical Methods for Non-Smooth Optimization (Computer Research and Mathematical Sciences)* (DCR-8302021 A01); $20,700; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $36,700).

This project involves investigations into numerical methods for solving certain non-smooth optimization problems. The main problem of interest is that of minimizing a sum of Euclidean norms of affine vector functions. A description is given of a new algorithm which solves dense problems very efficiently. Of particular interest is a convergence result which states that the algorithm is generally quadratically convergent, whether or not the objective function is differentiable at the solution. The goal is to adapt this algorithm to solve certain problems from continuum mechanics, which, when discretized, become large scale sum of norms optimization problems with a sparse strucure. The first example given is smooth but highly nonlinear: finding minimal surfaces. The second example, discussed in some detail, is a non-smooth problem from collapse load analysis. The latter is particularly difficult, because the solution to the underlying infinite-dimensional problem is a characteristic function, constant everywhere except along a curve where it has a jump discontinuity. Adapting the sum of norms algorithm to solve this is quite complicated, because matrix storage cannot be afforded and an elaborate data structure is required to handle the inherent degeneracy. Another topic of research is the matrix inverse eigenvalue problem. This may be formulated as a system of nonlinear equations which is nondifferentiable if the eigenvalues are multiple. In all cases there are three main goals: an improved mathematical framework, efficient and reliable numerical methods, and, ultimately, general-purpose, user-oriented software.

New York University; Olof Widlund; *Fast Iterative Methods for Certain Finite Element and Queuing Network Problems*; (DCR-8405506); $28,700; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $57,400).

This project involves the development of iterative methods for the numerical solution of several types of special large linear algebra problems. Work will be continued on iterative substructuring (domain decomposition) techniques for elliptic finite element problems. The main issue is now the extension of methods which have proven highly sucessful for second order problems in the plane to other elliptic systems. A study will also be made of the related component mode substitution methods, which have been developed by structural engineers as an alternative to traditional finite element eigenvalue methods. Another main part of the project is the application of similar numerical techniques to certain queuing network problems. In

these problems the Kolmogorov balance equations are generated and the steady state probability distribution is computed in terms of the null vector of an often very large, sparse nonsymmetric matrix. This matrix has much in common with those arising in finite difference problems and differs from decomposable problems by matrix of relatively low rank. A systematic study of the use of various iterative methods and special fast solvers and of the underlying continuous problems will be undertaken.

University of Pittsburgh; Werner Rheinboldt; *Computational Methods for Parametrized Equations*; (DCR-8309926); $52,088; 12 mos.

This research concerns the development and study of computional methods for the analysis of the solution manifolds of parametrized equations. Such problems arise frequently in applications as, for instance, in structural mechanics, fluid flow, network analysis, etc. Some of the topics under consideration concern the continued development and further extension of a general program package for the trace of paths on given solution manifolds, an investigation from the viewpoint of numerical analysis of the reduced basis technique pioneered by structural engineers, a study of a new approach combining continuation and discrete variable methods for the solution of differential equations on manifolds, and further work on the development and application of estimates of the discretization error of parametrized equations.

Brown University; Ulf Grenander; *Inference in Markov Processes (Mathematical Sciences and Computer Research)*; (DMS-8406827); $7,500 12 mos. (Joint support with the Statistics and Probability Program - Total Grant $20,000).

The proposed research aims to amplify on problems of statistical inference in the context of patterns.

Rice Univeristy; John Dennis; *Industry/University Cooperative Research Project: Nonlinear Exponential-Family and Bounded-Influence Regression*; (DCR-8116779 A02); $19,018; 12 mos. (Joint support with the Division of Industrial Science and Technology - Total Grant $38,036).

The research involves extending techniques that have proven worthwhile for nonlinear least squares to other members of the exponential family of regression problems and to nonlinear bounded-influence regression problems. These problems all have similar struture: the Hessian (or Jacobian matrix) is the sum of two terms, one of which is readily available. It has proven worthwhile in the case of nonlinear least squares to approximate the remaining term by a quasi-Newton technique and to use an adoptive modeling strategy. This work aims to see how worthwhile such techniques are for these other regression problems. It also includes some research on handling simple bounds on the parameters in these regression problems and on computing convariances and other statistics. One major goal of this project is to write reliable, modular, semi-portable software for solving the various problems addressed.

Southern Methodist University; Andreas Griewank; *Adopting Secant Updates to Large Structured Problems (Mathematical Sciences and Computer Research)*; (DMS-8401023); $9,916; 24 mos. (Joint support with the Applied Mathematics Program - Total Grant $29,748).

Professor Griewank will study large nonlinear systems that arise in the use of finite element and finite difference techniques for differential equations.

Many numerical methods for the solution of such equations employ secant updating techniques to achieve superlinear convergence while avoiding the cost of explicitly evaluating derivative matrices. While such quasi-Newton or Variable-Metric methods are widely used on moderatedly sized problems they have rarely been adopted to structured problems in many variables like those arising from the discretization of differential equations. To overcome these limitations of the classical secant methods, Professors Griewank and Toint have proposed the separate approximation of individual element Jacobians or Hessians, for example, Stiffness matrices in structural engineering.

University of Texas - Austin; David Young and David Kincaid; *Research and Development of Iterative Algorithms and Software for Elliptic Partial Differential Equations (Computer Research and Mathematical Sciences)*; (DCR-8214731 A01); $62,093; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $82,791).

This project involves research and development of iterative algorithms and software for solving large systems of linear algebraic equations with sparse matrices. Particular emphasis will be placed on linear systems arising from the numerical solution of elliptic partial differential equations.

The work will have two aspects, namely, development of research-oriented software and research on iterative algorithms. The first aspect of the work will involve a substantial expansion of the capability of the ITPACK package to include more powerful iterative algorithms and to allow for handling more general classes of linear systems including systems with nonsymmetric matrices. It will also involve continued participation in the ELLPACK project. The second aspect of the work is closely related to the first and involves research on iterative algorithms with emphasis on effective methods for solving nonsymmetric linear problems.

University of Washington; David Rogozin; *Spline Smoothing and Derivative Estimation (Mathematical Sciences and Computer Research)*; (DMS-8308349 A01); $8,000; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $28,300).

Professor Ragozin will investigate numerical, statistical and computational problems related to spline approximation of functions of one and several variables. He will emphasize estimation of derivatives and other smoothness properties of function from finite (empirical) samples which are contaminated by (random) noise.

Specific research topics include 1) analysis of error bounds for function and derivative estimates, 2) examination of the structural properties of a function reflected in a spline transform analog of the discrete Fourier transform, 3) development and analysis of a new method of spline smoothing for noisy data, based on the structure of the discrete K-functional and its relation to the spline transform, and 4) development and testing of algorithms to implement function and derivative estimation procedures.

University of Wisconsin - Madison; Olvi L. Mangasarian, Robert R. Meyer, and Stephen M. Robinson; *Computation and Theory in Nonlinear Programming*; (DCR-8200632 A02); $132,296; 12 mos.

The investigators will conduct research in the computational areas of successive overrelaxation (SOR) methods in mathematical programming, nonlinear programming, and equilibrium algorithms, nonlinear network algorithms, piecewise-linear approximation, exact penalty methods, parallel algorithms and algorithms for the complementarity problem and in the areas of nonlinear programming theory, generalized equations and stochastic programming. Computational testing of the algorithms developed will also be carried out on realistic problems in order to assess efficiently.

University of Wisconsin - Madison; Grace Wahba; *Multivariate and Multiresponse Estimation (Mathematical Sciences and Computer Research)*; (DMS-8404970); $10,000; 12 mos. (Joint support with the Statistics and Probability Program - Total Grant $36,900).

The proposed research is aimed at treating multivariate data where the problems are difficult, the need for good solutions is pressing, and the computational questions challenging.

Computer Systems Design

VLSI Design Methodology

University of Southern California; Melvin Breuer, Alice Parker, and Dennis McLeod; *The Adam VLSI Design System*; (ECS-8310774); $32,500; 12 mos. (Joint support with the Computer Engineering Program - Total Grant $110,000).

This project addresses that research essential to build an Advanced Design Automation System (ADAM). ADAM is to be an experimental system capable of being used for the design, and implementation in silicon, of complex digital computing systems.

Three major research issues need to be dealt with;

1. Overall system architecture of ADAM,
2. The design of a data base system to support ADAM, and
3. Problem specific applications such as, synthesis, verification, layout and testing.

The research to be supported by the NSF grant covers the first two major issues and a single application domain; an expert system for custom layout (CLOUT). CLOUT must be developed as part of the overall design system because physical design (including layout) proceeds in parallel with logical design. Thus, to evaluate CLOUT one needs a total design environment with feedback on system architecture, choice of layout design styles and tradeoffs in performance.

There are three novel aspects to the work;

1. The emphasis on embedding artificial intelligence techniques into a general computer-aided design automation system architecture, not just for a particular application,
2. The design data being structured in a new way which separates target system behavior, structure, control and timing, and physical information, and
3. Much of the design [illegible] and many of the steps in the design process being characterized mathematically which provides a continuum between design and analysis.

Illinois Institute of Technology; Anthony Wojcik; *A Study of the Application of an Automated Reasoning Based System to Design Automation Problems*; (DCR-8317524); $132,553; 24 mos.

The focus of this proposal is on logic verification utilizing a theorem prover approach. In VLSI logic design, the normal approach is to design complex devices by refining the design to include more and more detail as the size of the building block being considered decreases. At this time it is not possible to consider an entire complex design in its finest detail at one time. Dr. Wojcik is using this same strategy to verify complex designs. The approach is to decompose the complex design into a set of simpler designs and a set of verifiable combining rules.

The goal of this research is to continue to develop an LMA (Logic Machine Architecture) based system to be used as a design automation tool. Research will include logic verification between several levels of the design hierarchy, logic synthesis, and the study of complexity issues involved with this approach to verification and synthesis.

University of Illinois - Urbana; Saburo Muroga; *Algorithms and Design Tools for VLSI Design*; (ECS-8405558); $55,000; 12 mos. (Joint support with the Computer Engineering Program - Total Grant $110,000).

A major area in the study of logic design is that of minimization such as the minization of chip area or device pin count in VLSI designs. This project focuses on such minimization and is exploring the following problems, some of which are extensions of previous research efforts by this group.

1. The development of procedures to design compact Programmable Logic Arrays and the variants and procedures to design compact arrays which are NAND networks in 3 or more levels is being pursued. The advantages of these procedures are being explored and compared.
2. The properties of minimal adders with different types of gates and minimal adders with extra functional capabilities (i.e., adders combined with logic operations) are being studied. Based on these properties, parallel multipliers are being designed and the properties of these multipliers explored.
3. Each technology being used for integrated circuit implementation has its own characteristics. Design methodologies must include consideration of these characteristics. The project will focus on minimal networks to be implemented using the MOS technology.

Carnegie-Mellon University; Stephen W. Director, Daniel P. Siewiorek, and Donald Thomas; *Multilevel Computer-Aided Design of VLSI Digital Systems*; (ECS-8207709 A01); $96,567; 12 mos. (Joint support with the Electrical and Optical Communications Program - Total Grant, $193,134).

A hierarchial approach to the design of digital systems is being taken and begins with a behavioral specification of the system. This specification provides a model which accurately characterizes the input - output behavior of the system without reflecting any internal structure. The hierarchial design is supported by a comprehensive set of computer aided design (CAD) methods which take into account the important characteristics of integrated circuit technology.

The research to be performed continues previous work and addresses the following areas;

1. Synthesis of digital VLSI sytems,
2. Multilevel representations for digital VLSI systems,
3. Multilevel optimization and digital system design space exploration,
4. Man-machine environment for CAD of digital VLSI systems, and
5. Special hardware for VLSI CAD.

University of Utah; Richard F. Riesenfeld and Lee A. Hollaar; *A Laboratory for Computer-Aided Design*; (DCR-8121750 AO2); $15,000; 12 mos. (Joint support with the Coordinated Experimental Research Program - Total Grant $30,931).

This award will aid in the establishment of a major experimental facility to support a program of research on Computer-Aided Design (CAD) Systems. The facility will include powerful design workstations based on personal graphics-oriented machines. Included also will be a necessary expansion of general computing facilities to accommodate the expanded research effort. Several specialized 3-D output devices will provide a basis for testing the CAD techniques and tools developed as part of the research.

There are two main thrusts to the current efforts: Computer-aided 3-D shape design and Very Large Scale Integration (VLSI) design. The ALPHA-1 development project, at the heart of the first research thrust, is capable of supporting a wide range of CAD activities. Moving ALPHA-1 to a distributed environment will require research in operating systems and in workstation support and design. The investigators will also work on developing true 3-D display of surfaces and on Numerical-Control Milling Algorithms. Research in portability and software tools will lead to an enhanced software development environment; in database systems it will lead to an efficient uniform means for accessing, modifying, and sharing design data; and in multiprocessor and network operations it will allow the various tasks in the systems to be partitioned and run on processors best suited to their particular requirements. In integrated circuit design, the main emphasis will be on simulation systems improvement, on mapping techniques from high-level representations directly to design, on graphical representations of VLSI designs, and on the testing and diagnosis of integrated circuits.

Defense Advanced Research Projects Agency; Robert S. Cooper; *NSF/DARPA Agreement for Use of DARPA VLSI Implementation Service*; (ECS-8210242 A02); $80,000; 24 mos. (Joint support with the Computer Engineering Program - Total Grant $160,000).

The development of technology making feasible Very Large Scale Integrated (VLSI) chips opens the way for universities to play a major innovative role in the design, construction, and application of novel special purpose hardware and new computer architectures. Considerable research is already in progress in the university community on software and hardware design tools needed to manage the complexity which arises in designing new VLSI computing structures containing tens of thousands of transistors on a silicon chip. Novel designs are being proposed and studied. There is a great need to implement these designs in silicon chips and test their performance.

A joint experiment is being undertaken by the National Science Foundation and the Defense Advanced Research Projects Agency (DARPA) that will allow qualifying universities to use the DARPA fast turnaround VLSI implementation facility for university-based research and educational programs requiring the fabrication of digital designs as integrated circuit chips. As part of DARPA's research in the VLSI area they have developed a capability called "MOSIS" for fast turnaround implementation of VLSI chips. Designs submitted by DARPA researchers in digital form over the ARPANET or Telenet using standard design rules and a standard artwork format, known as CIF, are fabricated and returned to designers in about 4-6 weeks.

Access to this VLSI service by universities will be on a controlled basis with usage limited to those institutions that meet requirements jointly established by NSF and DARPA. The integrated circuit chips will be fabricated, packaged, and sent, as soon as practicable, to the university groups submitting the designs.

Computer System Architecture

Yale University; David Gelernter; *Analysis and Simulation of New Architectures for Network Computers*; (DCR-8306836); $81,243; 24 mos.

A network computer is a computer network designed to function and be programmed as a single machine. This project focuses on network computers

designed to support asynchronous distributed programs. A series of analytic and simulation studies designed to lead ultimately to the selection of a good architecture for a major network computer implementation are proposed. In particular, the performances of architectures in two widely different classes are being investigated. Architectures in both classes meet basic requirements imposed in top-down fashion by the execution environment to be supported, and both seem reasonable and interesting. The first class includes architectures based on grids of Ethernet-like contention busses. The second includes architectures in which front-end switches allow the use of a protocol midway between packet and circuit switching; the front-end switches also make it possible to program the communication kernel in such a way that the network reconfigures itself at runtime in response to measured traffic patterns. Analysis and simulation will allow comparison of the two with respect both to performance and manageability.

This is a collaborative project with Professor Hussein Badr at the State University of New York - Stony Brook.

University of Georgia; Jeffrey Smith; *Processors with Parallel Capability for Factoring Integers (Computer Research)*; (DCR-8302877 A02); 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant $52,640).

Special processors for factoring large integers, using the continued fraction algorithm with early exits, are developed. The machines handle extended precision operands and have some parallel capability as computational accelerators. This is a collaborative project with Professor Samuel Wagstaff at Purdue University. Professor Smith is designing the processors, developing their systems software, and evaluating their operation. Professor Wagstaff is analyzing the algorithm with the computational accelerators in mind and writing the application programs.

A Denelcor Heterogeneous Element Processor computer will be used for the two projects. The quadratic sieve integer factoring algorithm, which is the most serious known competitor of the continued fraction method, will be studied and a large-scale test of the Extended Riemann Hypothesis will be made.

University of Illinois - Urbana; David J. Kuck and Duncan H. Lawrie; *Computer System Organization*; (DCR-8300981 A01); $60,330, 6 mos.

Research during this period will focus on studies in two elements of the "Cedar" computer design:

1. Global Memory Switch - Analytical and simulation work on extending the basic omega networks to larger switches, and
2. Global Control Unit (GCU) - The study of how application program graphs perform on the GCU. The speed of the GCU and its effects on system performance are crucial areas of study.

University of Illinois - Urbana; Daniel A. Reed; *Performance Directed Design of Multimicrocomputer Systems (Computer Research)*; (DCR-8417948); $55,825; 24 mos.

This research focuses on developing a high-performance multimicrocomputer system design by exploring the interaction of system design decisions. A multimicrocomputer system consists of a large number of interconnected computing nodes that asynchronously cooperate via message passing to execute the tasks of parallel programs. Each network node, fabricated as a small number of VLSI chips, contains a processor, a local memory, and a communication controller capable of routing messages without delaying the computation processor.

Because there is no global memory, a multimicrocomputer system supporting computations that dynamically create tasks must schedule those tasks on idle nodes using information local to the area where the tasks were created. The system must also support efficient internode communication via packet switching, virtual circuits, or some other message switching strategy. Finally, the system must recover from node and link failures during computations. For a multimicrocomputer system to be successful, scheduling, communication, and fault tolerance must be considered in concert. Via analysis and simulation, the performance of a multimicrocomputer system is being examined as its scheduler, message switching strategy, and checkpointing algorithm interact.

University of Illinois - Urbana; James E. Robertson; *Research in the Area of Digital Computer Arithmetic*; (DCR-8308576 A01); $68,688; 12 mos.

Early in 1980, decomposition procedures for the decomposition of complex structures for addition into simpler structures were developed. The basic building blocks are at a level ranging from a half-adder to a full adder. Studies of the decomposition procedures are continuing at all levels, including establishing a sound theoretical basis for some aspects of the present theory, extension from binary to higher radices, and additional applications of the procedures to practical designs.

Also in 1980, a feasibility study of the design of a variable precision processor module was successfully completed. A module of n bits precision is designed in such a way that kn bits precision can be achieved by using a single module serially k times or by using k modules in parallel. The extension to higher precision requires external registers of the precision desired, but programming changes are unnecessary. The module is

capable of addition, subtraction, multiplication, and division. Studies are underway which are ultimately intended to lead to a detailed design suitable for implementation by VLSI. Thus far, these studies have led to the design of the combinational logic for a radix 16 digit slice, and have provided a useful example of the application of the decomposition procedures for structures.

These studies, including register and gating arrangements, miscellaneous logic for normalization and quotient digit selection, nature of the control hierarchy, selection of an operation code, and detailed control design are continuing.

Purdue University; Samuel S. Wagstaff; *Processors with Parallel Capability for Factoring Integers*; (DCR-8406596); 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant $29,520).

Special processors for factoring large integers, using the continued fraction algorithm with early exits, are developed. The machines handle extended precision operands and have some parallel capability as computational accelerators. This is a collaborative project with Professor Jeffrey Smith at the University of Georgia. Professor Smith is designing the processors, developing their systems software, and evaluating their operation. Professor Wagstaff is analyzing the algorithm with the computational accelerators in mind and writing the application programs.

A Denelcor Heterogeneous Element Processor computer will be used for the two projects. The quadratic sieve integer factoring algorithm, which is the most serious known competitor of the continued fraction method, will be studied and a large-scale test of the Extended Riemann Hypothesis will be made.

University of Iowa; Roger K. Shultz; *Evaluation of Multiprocessor Computer Architectures for Database Support*; (DCR-8403317); $62,326; 12 mos.

Information about multiprocessor support of specialized processing tasks is being collected and evaluated Specifically, an experimental comparison of proposed multiprocessor computer architectures for database support is being conducted by simulating the machines in a controlled database environment. A multiprocessor computer architecture simulator is being used to emulate proposed back-end database machines. The database environment is being established for each experiment by two software systems: A transaction driver and a database generator. The transaction driver produces a controllable transaction mix which is used as input to the emulated machines. The database generator generates experimental databases with known mathematical properties. The simulator, driver and generator allow the design and implementation of experiments for collection of performance statistics.

Many machines have been proposed for the support of database operations. Algebraic analysis has been used to compare various proposed machines and the results of these comparisons are being used to choose the machines to be simulated. Algebraic analysis of computer architecture requires the use of simplifying assumptions about machine and data parameter values. Simplifying assumptions can be reduced through experimental simulation with actual data and traces of application program transactions. Information collected in such experiments will contribute to the general area of specialized multiprocessor computer architecture design.

University of Maryland - College Park; Yaohan Chu; *Prolog Machine Architecture*; (DCR-8320083); $100,000; 24 mos.

This research is focused on the study of computer architectures capable of performing high-speed logic inferences and handling a large knowledge data base. The architectures are based on the language Prolog and incorporate concepts of associative memory and parallelism. Simulation is used to evaluate the various architectures.

University of Maryland - College Park; Satish Tripathi and Ashok Agrawala; *Models for Parallel Software*; (DCR-8405235); $50,000; 12 mos; (Joint support with the Software Engineering Program - Total Grant $116,730).

This project focuses on problems in two related areas of centralized and distributed computer environments, namely load sharing/routing, and system modeling.

Massachusetts Institute of Technology; Jack Dennis; *Data Flow Computer Architecture*; (DCR-7915255 A04); $109,134; 12 mos.

The Computation Structures Group is continuing its fundamental research on program structures and computer architecture for data flow computation. Areas being studied include: program transformations to better match program structure to a hypothetical data flow architecture; machine level program structure and code generation; hardware specification, verification and simulation based on an architecture description language for packet systems; and new methodologies for performance estimation and fault tolerance applicable to packet system architectures.

University of Lowell; Robert J. Lechner; *Harmonic Analysis of Logic Functions I: Encoded Programmable Logic Arrays*; (DCR-8305571 AO1); $21,693; 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant $43,386).

This research is concerned with the development and evaluation of a new unified approach to array logic synthesis using VLSI technology on a scale which applies to software storage as well as to firmware storage and wired control logic within digital systems. Specifically, PLAs are imbedded within linear encoding transformations on their inputs and outputs and combined with ROMs as necessary to define more optimal implementations of logic functions or stored tables than are possible with either ROMs or PLAs alone. In contrast to more general array logic techniques, the imbedded PLA approach retains the homogeneous array device layout and interconnection topology that reduces design and testing costs.

Heuristic algorithms are developed and tested on realistic examples to evaluate the feasibility of the imbedded PLA approach. Theoretical studies are used to evaluate a more rigorous approach to canonical synthesis algorithms that do not require human intervention. A crucial outcome of the study is the identification of that class of functions for which the hybrid (ROM + imbedded PLA) approach to synthesis is more cost-effective than either ROMs or PLAs alone.

University of Massachusetts - Amherst; Victor R. Lesser; *Coordination in Cooperative Distributed Problem Solving Systems (Computer Research)*; (DCR-8300239 A01); $40,282; 12 mos. (Joint support with the Special Projects Program and the Intelligent Systems Program - Total Grant $120,847).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts-Amherst has included the development of a node architecture capable of such sophisticated local decisionmaking. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. The research being carried out will build on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include: organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

University of Michigan - Ann Arbor; Arthur W. Burks; *Languages and Architectures for Parallel Computing with Classifier Systems (Computer Research)*; (DCR-8305830 A01); $102,827; 12 mos. (Joint support with the Intelligent Systems Program - Total Grant $102,827).

A new class of rule-based computing systems, classifier systems, has been developed and tested. While classifier systems are similar to production systems, they are different in basic ways which facilitate the incorporation of learning algorithms in them, and which make it possible for them to be executed in a highly parallel fashion without the usual problems of concurrency and interlock. Classifier systems with learning algorithms have been tested successfully on two model control systems, a game learning program, and a hardware robot project. Progress has been made on the theoretical understanding of how such classifier systems work.

This research will focus on two interrelated areas. (1) The powers and limits of classifier systems will be studied by further computer simulations and theoretical analyses. (2) Novel parallel architectures for the rapid execution of classifier systems will be planned and simulated. On the theoretical side, these studies will contribute to a deeper understanding of concurrency and distributed control, and on the practical side they will lead ultimately to powerful parallel computers which are non-von-Neumann in both programming and architecture.

Duke University, Robert A. Wagner; *Boolean Vector Machine (Computer Research)*; (DCR-8403011); $78,601; 12 mos.

The Boolean Vector Machine (BVM) models the ability of today's computer technology to perform bitwise parallel Boolean operations on huge vectors of zeros and ones. If a BVM's register or vector length is unrestricted, a BVM can solve efficiently (in polynomial time) many combinatorial problems which seem to require exponential time on conventional small-word-size computers. These same problems are of considerable importance for, if they can be solved efficiently in a given computational environment, then so also can several major problems of Artificial Intelligence and Combinatorial Optimization. The research described here seeks to determine the "value" of building a finite version of a BVM, one that costs roughly what a conventional computer of equal memory size would cost but that is capable of solving these combinatorial problems $2^n/n$ times faster than the conventional machine, when both machines have $n \times 2^n$ bits of memory. An essential part of the research endeavor is a design exercise, which will produce detailed designs of the instruction set, interconnection network and packaging of one or more BVM systems. The performance of each of these systems on several

different classes of problems will also be investigated, both in a pencil-paper exercise, where "time" is equated to "instructions executed", and in a real-time fashion.

The design exercise will also involve algorithm redesign, to execute efficiently on various BVM's. The development of new design methods for both hardware and software is an additional expected benefit from this work.

University of North Carolina - Chapel Hill; Kye S. Hedlund; *Wafer Scale Integration of Parallel Processors*; (DCR-8302641 AO1); $69,561.

The objective of this research is to investigate the wafer scale integration of parallel processors. Rather than dice the wafer into individual chips as is usually done, the idea behind wafer scale integration is to assemble an entire system on a single wafer.

This project is focusing on the design of interconnection switches and the construction of prototype building blocks in an environment of several processing elements on a chip size piece of silicon. Future expansion to the wafer size is contemplated. In addition, research on algorithm development is progressing.

State University of New York - Stony Brook; Hussein G. Badr; *Analysis and Simulation of New Architectures for Network Computers*; (DCR-8309958); $56,492; 24 mos.

A network computer is a computer network designed to function and be programmed as a single machine. This project focuses on network computers designed to support asynchronous distributed programs. A series of analytic and simulation studies designed to lead ultimately to the selection of a good architecture for a major network computer implementation is proposed. In particular, the performance of architectures in two widely different classes is being investigated. Architectures in both classes meet basic requirements imposed in top-down fashion by the execution environment to be supported, and both seem reasonable and interesting. The first class includes architectures based on grids of Ethernet-like contention busses. The second includes architectures in which front-end switches allow the use of a protocol midway between packet and circuit switching; the front-end switches also make it possible to program the communication kernel in such a way that the network reconfigures itself at runtime in response to measured traffic patterns. Analysis and simulation will allow comparison of the two with respect both to performance and manageability.

This is a collaborative research project with Professor David Gelernter at Yale University.

Oregon Graduate Center; Richard B. Kieburtz; *The G-Machine: A Fast Graph-Reduction Processor*; (DCR-8405247); $134,058; 24 mos.

Functional languages are markedly superior to imperative languages for development of many types of applications software. They offer advantages of clarity of expression, particularly of nonnumeric algorithms, are amenable to the use of formal verification techniques, and they make possible the use of automatic storage management.

The proposed research is based upon a simple abstract machine model that has recently been proposed by a research group at Chalmers University in Gothenburg, Sweden. Their model has been used to build a functional language compiler that produces code for a conventional processor. The compiler code is remarkably fast, which is strong evidence that the abstract machine model is appropriate to its task.

The emphasis of this research project is on the use of simulation and prototype development to show feasibility of this approach. Certain functional components crucial to performance of the processor, such as the indexed stack (S-stack) to support graph traversal are being implemented and tested. The expected result of this project will be a well-tested exploration of design alternatives and their implications for performance. As the detailed design phase progresses to completion, a single-board implementation using commercially available LSI parts will be constructed in order to test functional behavior and microcode.

Carnegie-Mellon University; Kendall Preston; *Workshop on Multicomputers and Image Processing, Tucson, Arizona, May 22-24, 1984*; (DCR-8315057); $8,580; 12 mos. (Joint support with the Intelligent Systems Program - Total Grant $17,160).

Starting in May 1979, there have been annual workshops involving a small group of scientists who are responsible for most of the major research efforts worldwide in multicomputer systems related to image processing. These systems include the CLIP and DAP systems (Great Britain), PICAP (Sweden), DIP (Holland), AT4 and SYMPATI (France), FLIP (West Germany), PPP and CYBEST (Japan), CYTOS, GLOPR, ZMOB, MPP, PUMPS, PASM (United States). These workshops have acted to stimulate research in this field and this year will be held in May 1984. The workshop will be co-chaired by Professor Preston of Carnegie-Mellon University and by Professor Uhr of the University of Wisconsin. In conjuntion with this particular workshop, these two scholars will conduct pre-workshop benchmark studies on many of the above multicomputer systems.

Pennsylvania State University; Jesse L. Barlow; *Probabilistic Error Analysis of Numerical Computations in Special Computer*

Arithmetics; (DCR-8201065 A01); $28,800; 24 mos. (Joint support with the Software Engineering Program - Total Grant $76,822).

This project involves two separate, but related, problems. The first is the probabilistic error analysis of numerical algorithms. The second is the design of a parallel architecture for solving ordinary differential equations using digital on-line arithmetic. The error analysis of digital on-line arithmetic has already presented problems over and above those of standard floating point arithmetic. The fact that it is a redundant arithmetic system and that its addition and subtraction operations do not automatically normalize is the cause of most of these problems. Thus, it is useful for special purpose devices only. It has the advantage, however, of allowing the hardware designed to incorporate additional parallelism in iterative and recursive algorithms. A particular application of this property that is of great importance is the solution of numerical initial value problems.

Southern Methodist University; David Matula; *Foundations of Finite Precision Rational Arithmetic*; (DCR-8315289); $41,611; 12 mos.

VLSI capabilities are fostering an era wherein complex special purpose computer hardware for specific broad application areas is realizable at acceptable cost. Previous research by Dr. Matula has described special number systems and proposed a new computer arithmetic unit sufficiently competitive in convenience of arithmetic logic design and timing efficiencies to be a serious candidate for a special purpose arithmetic architecture. This research project is further refining the design possibilities for finite precision rational arithmetic by:

1. Investigating the utilization of binary continued fraction representation for values in storage and as input/output of a rational arithmetic unit;

2. Incorporating on-line techniques in chain rational computation and other areas for increased parallelism in a rational arithmetic unit;

3. Utilizing extended range binary floating-slash representation and/or binary continued fraction representation to prescribe a hierarchical precision rational computation system; and

4. Investigating accumulated roundoff error growth for computations hosted in alternative finite precision rational number systems.

University of Texas - Austin; James C. Browne; *Workshop on Resources for Cooperative Research in Parallel Processing*; (DCR-8320669); 6 mos. (Joint support with the Computer Engineering Program and the Division of Industrial Science & Technology - Total Grant $18,369).

It is well established that the difficulties of executing large scale experimental research programs in a university environment has been a major limiting factor in the development and evaluation of concepts for parallel computing. This difficulty in the evaluation of concepts has limited technology transfer to industry and increases the cost and risk of developing computer architectures for parallel computing. These cost and risk factors are believed to be major factors in limiting the rate of growth in the development and application of parallel processing. This workshop will bring together representatives from the computer industry, senior researchers from universities and government laboratories, and representatives from government funding agencies to explore new avenues for cooperation on research and development in parallel processing.

University of Texas - Austin; James C. Browne; *High Performance Parallel Computing*; (DCR-8116099 A02); $138,948; 12 mos.

This research project is an integrated experimental approach to parallel programming and parallel computing. It is driven by the mapping of significant computation problems to the Texas Reconfigurable Array Computer (TRAC). The purpose is to determine the applicability and effectiveness of a reconfigurable network architectured computer system as a vehicle for high performance parallel computing. The research includes design and development of basic software for a reconfigurable multi-processor architecture, programming, execution and measurement of applications on TRAC and extension of the concept base of reconfigurable network architectured computing in the light of the knowledge developed by applications.

University of Wisconsin - Madison; James R. Goodman; *Decoupled Access/Execute Computer Architectures for VLSI/ULSI*; (DCR-8202952 A02); $146,481; 24 mos.

Current large scale scientific computers (supercomputers) are often limited by their inability to sustain a high effective rate of data transfer from memory, when data access patterns become more complicated than the access mechanisms provided by the machine.

A computer system based on a "queue" or "decoupled" architecture is being explored. The architecture has independent processors for data access and computation, thus introducing a new level of parallelism. The project is investigating the design of a system from a relatively small number of very powerful, tightly-coupled access and computation processors. Under program control, the system's access processors acquire data from and store data to central memory, and computation processors manipulate and transform the

data. By this scheme, both vectorizable and complex non-vectorizable sequences of code can be executed efficiently. The access processors cooperate to provide a high sustained rate of data transfer from memory to the computation processors.

University of Wisconsin - Madison; Leonard M. Uhr; *Algorithm-Structured Pyramid and Network Architectures*; (DCR-8302397 A01); $68,537; 12 mos.

This research continues an on-going investigation of multi-computer networks, in an attempt to find a) good algorithm-structured architectures, b) good general network architectures, c) good criteria, both structural and formal, for evaluating specialized and general multi-computer architectures.

Several key issues are being examined related to future VLSI implementation of extremely large multi-computers. These include:

1. How are basic resources best distributed among processors, memories,reconfiguring switches and controllers?
2. How are micro-modular arrays and pyramids best designed?
3. How can fault-tolerant reconfiguring capabilities that take advantage of the micro-modular array/pyramid structure best be added?

Fault Tolerance and Reliability

Stanford University; Edward J. McCluskey; *Reliable Computers*; (DCR-8200129 A02); $99,876; 12 mos.

This research is aimed at developing systematic design techniques for high-availability computer design.

Specific attention will be given to unification of the treatment of:

1. production testing of integrated circuits and inclusion of appropriate circuitry to guarantee that this is economical (Design for Testability),
2. circuitry for concurrent checking of computer operation, and
3. system structure to facilitate restructuring to eliminate failed subsystems from interfering with correct operation.

The novelty of this study stems from the fact that all the aspects of high-availability computer design will be treated in a uniform way. Past studies of this topic have tended to emphasize one aspect of the problem to the exclusion of the others. Modern VLSI techniques make it possible to combine the various techniques into one unified design technique.

University of California - Los Angeles; David A. Rennels and Algirdas A. Avizienis; *Fault-Tolerant Local Networks using VLSI Building Blocks*; (DCR-8307026); $108,399; 12 mos.

This research is aimed at developing and evaluating architectural concepts for fault-tolerant local networks using VLSI building blocks circuits to provide hardware fault-tolerance and a software environment characterized by a high level of security, ease of programming, and run-time checking for a wide range of software errors.

The major areas of study are:

1. an investigation of architectures for large shared memory and computational resources which are self-managing and fault tolerant,
2. an investigation of the applicability of specialized hardware support in intercommunications and synchronization, and
3. a study of smart memory interfaces which automatically map the local address spaces of distributed computers to the virtual address space provided by such a network.

Boston University; Mark G. Karpovsky; *Universal Testing of Computer Hardware*; (DCR-8317763); $112,019; 24 mos.

The time and the cost of testing VLSI devices increases rapidly as the complexity of these devices increases. The research in VLSI-based system fault location concentrates on built-in concepts at the system level, interconnection testing, distributed test control and response analysis, and the use of directed graph models for the analysis and design of diagnostics. This work attempts to provide systematic techniques for cost-effective solutions to portions of the VLSI testing problem which must be solved in order to effectively use VLSI technology.

University of Wisconsin - Madison; Charles R. Kime; *Testing and Diagnosis in VLSI Circuits and Systems (Computer Research)*; (DCR-8206564 A02); $65,000; 12 mos.

Very-large-scale-integrated (VLSI) systems constitute an environment in which the problems of testing and fault location are becoming increasingly difficult as circuit densities and complexities increase. This research addresses two major problem areas: 1) built-in-test approaches for VLSI circuits and 2) fault location techniques for VLSI-based systems. The research in VLSI circuit testing focuses on test parallelism and scheduling, new approaches to exhaustive testing, alternative built-in test structures, analysis of linear data compression structures, fault location in circuits, and

the relationship of built-in-test to structure design methodologies. The research in VLSI-based system fault location concentrates on built-in concepts at the system level, interconnection testing, distributed test control and response analysis, and the use of directed graph models for the analysis and design of diagnostics. This work attempts to provide systematic techniques for cost-effective solutions to portions of the VLSI testing problem which must be solved in order to effectively use VLSI technology.

Computer System Performance Measurement and Evaluation

Stanford University; Donald L. Iglehart (In Collaboration with Gerald S. Shedler, San Jose Research Laboratory, IBM); *Industry/University Cooperative Research: Simulation Output Analysis for Local Computer Networks*; (DCR-8203483 A02); $50,633; 12 mos. (Joint support with the Industry/University Cooperative Program - Total Grant $101,266; 12 mos).

The focus of this research is on the development of rigorous theoretical foundations for discrete event digital simulation. Primary motivation for the specific work is the need for firmly based methods applicable to simulations of local computer networks. Emphasis is on the development of new probabilistic and statistical methods for the design of simulation experiments and the analysis of simulation output. The variability inherent in most local computer networks is such that to make meaningful performance predictions it is necessary to study the evolution of a stochastic system. Computer simulation is required to do the complexity of real systems. Once committed to computer simulation, theoretically sound and computationally efficient methods for carrying out the simulation must be found. Among the issues the simulator must face are the initial conditions for the system being simulated, the length of the simulation run, the number of replications of the experiments, and the length of the confidence interval. The estimation methods being studied rest on results from stochastic process theory.

Duke University; Kishor S. Trivedi; *Analysis of Parallel and Distributed Systems (Computer Research)*; (DCR-8302000 AO1); $69,890.

Recent technological advances and user needs of high computing capacity and reliability have given rise to increased interest in parallel and distributed processing systems. This project focuses on key issues related to evaluation techniques for such systems. The goal of these techniques is to provide accurate and efficient methods for evaluating and designing such systems. Of particular interest are methods for combined evaluation of system performance and reliability for gracefully degradable systems, load-distribution in a distributed processing system, and performance analysis of resource sharing systems where individual transactions possess internal concurrency.

University of Texas - Austin; K. Mani Chandy; *Algorithms For Computer Performance Analysis*; (DCR-8318407); $74,464; 24 mos.

This effort is concerned with the performance analysis of distributed computing systems. The measurement, simulation and mathematical modeling of networks of processors which communicate via messages are being explored. An experimental distributed facility for measurement, an existing distributed simulation language and queueing network analysis tools are being utilized. The major part of the research effort is concerned with developing new techniques for the measurement, distributed simulation and approximate queueing-theoretic modeling which are suited for distributed systems. Attention is focused on distributed resource management.

University of Washington; Edward D. Lazowska and John Zahorjan; *Performance Modelling of Integrated, Locally Distributed Systems (Computer Reseach)*; (DCR-8302383 A01); $99,251; 12 mos.

Queueing network models, because they achieve a favorable balance of accuracy and cost, are important tools in the design and analysis of computer systems. The applicability of queueing network models is being extended to the coming generation of multi-user systems, which are characterized as integrated, locally distributed systems.

In contrast to existing centralized systems, integrated distributed systems will provide good support for personal computing. In contrast to existing locally distributed systems, they will provide good support for the sharing of information and of processing capacity. These characteristics create challenges for the analyst: modelling distributed algorithms, modelling distributed data, and designing and evaluating the algorithms that control this distribution, to name but a few. These are the problems that are being investigated.

Computer Graphics

Princeton University; David P. Dobkin; *Theoretical and Pragmatic Concerns in Computer Graphics*; (DCR-8303926 A01); $93,231; 12 mos.

This research is exploring the implementation of graphics primitives. The study is at the theoretical level but also involves hardware and software implementations. Three problems are being explored in detail. These represent different parts of the graphics process - line drawing, hidden surface elimination and number representations for accurate computation.

Cornell University; Donald P. Greenberg; *Interactive Computer Graphics Input and Display Techniques*; (DCR-8203979 A02); $238,929; 12 mos.

This research project is concerned with the development of new and improved input and display methods for interactive computer graphics. Two general areas are addressed, the graphical creation of geometric object descriptions, and algorithmic processes for realistic and dynamic image synthesis. Specific research topics include color science and equal perception spaces, a polygon hidden surface algorithm, a light reflection model, and investigations into new uses of coherence for high quality image display. Results are potentially useful to a full range of computer-aided-design applications.

Ohio State University; Charles A. Csuri and Edwin Tripp; *Motion Control, Natural Phenomena, and Parallel Architectures for Display Algorithms*; (DCR-8304185 A01); $200,000; 12 mos.

Three areas of research are being investigated. Primary emphasis is on the areas of motion control and parallel architectures for display architectures. Secondary emphasis is on the simulation of natural phenomena. There is no simple or short term solution to any of these tasks. The control of articulated motion presents problems of algorithmic control, knowledge-based animation, movement knowledge base, environmental knowledge base, special problems in kinematics, the application of special purpose motion processors and the generation of complex surfaces for moving figures. Parallel architectures for display algorithms, for the most part, have been described at a theoretical level. There are few case studies to validate various theories. An experimental parallel microprocessor architecture is available on which initial studies on the partitioning of display algorithms is being explored. Simulation of natural phenomena such as clouds, fire, smoke, ice and even trees have not yet been displayed realistically in computer graphics. The problem is mainly concerned with data complexity and computational requirements.

Rice University; Rui J. P. De Figueiredo; *Syntactic/Semantic Methods For 3D Object Reconstruction*; (DCR-8318514); $43,382; 24 mos. (Joint support with Intelligent Systems Program - Total Grant $83,382).

A high-level syntactic/semantic approach to the reconstruction of a surface delimiting a three-dimensional object is being studied. The surface is reconstructed from a set of points on contours present on the surface. Polygonal approximations are represented by attributed strings belonging to an appropriately constructed language. The surface is reconstructed by interconnecting adjoining strings according to their similarities in shape attributes, using advanced concepts from syntactic/semantic pattern recognition theory.

University of Utah; Richard F. Riesenfeld; *Subdivision Algorithms for Curved Surface Representation and Display (Computer Research)*; (DCR-8203692 A02); $141,082; 12 mos.

This project is concerned with the further investigation of subdivision algorithms and their generalizations as applied to computer aided geometric design. The Oslo Algorithm for computing discrete B-splines is a general structure for developing specific algorithms which are optimized and specialized for particular applications. Other work is being devoted to looking for non-tensor product surface formulations and analogues to the Oslo Algorithm. Such algorithms are being applied to designing set operations on freeform surfaces. Continued emphasis is being placed on combining high quality computer graphics and freeform surfaces.

Intelligent Systems

Computer Vision and Image Processing

SRI International; Alex P. Pentland; *Fractal-Based Description of Natural Scenes (Computer Research)*; (DCR-8312766); $121,078; 24 mos.

This research program addresses the problems of (1) representing natural shapes such as mountains, trees and clouds, and (2) computing such a description from image data. The first step towards solving these problems is to obtain a good model of natural surface shapes; such a model, together with the physics of image formation, provides the necessary analytical tools for relating natural surfaces to their images. Fractal functions are a good choice for modeling natural surfaces because many physical processes produce a fractal surface shape, and because fractals are widely used as a graphics tool for generating natural-looking shapes. Encouraging progress has already been achieved. A survey of natural imagery has shown that the 3-D fractal model, transformed by the image formation process, furnishes an accurate description of homogeneous image regions, thus providing a model for both image texture and shading. This characterization of image regions has been shown to be invariant over transformations of scale and linear transforms of intensity.

Stanford University; Paul Switzer; *Statistical Theory and Methods for Processing Spatial Imagery (Mathematical Sciences and Computer Research)*; (DMS-8411300); $13,700; 12 mos. (Joint support with the Statistics and Probability Program - Total Grant $38,600).

The research aims to develop models and algorithms to treat very complex and poorly understood issues connected with reconstructing images from noisy data.

University of California - Berkeley; Alberto Grunbaum; *Reconstruction with Limited and Noisy Data (Mathematical Sciences and Computer Research)*; (DMS-8403232); $8,000; 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $32,000).

The central issue in the data inversion problems to be studied is to determine specific properties of an object from measurements of the way it transmits or reflects acoustic signals, thermal waves, X-rays of other types of incident waves. When the measurements are completely accurate and sufficiently abundant, the desired reconstruction can often be accomplished by elementary means. But in practice the data is often incomplete and corrupted by noise and reconstruction techniques which work for accurate and complete data may fail completely. This research is concerned with the reconstruction problem in such an environment. In particular, the problems of how best to adapt or modify algorithms designed to handle complete and clear data to the case of incomplete and noisy measurements is being investigated.

University of Illinois - Urbana; Thomas S. Huang; *Acquisition, Representation, and Manipulation of Time-Varying Spatial Information (Computer Research)*; (DCR-8206926 A01); $52,125; 12 mos.

This research focuses on basic problems in the modeling and analysis of time-varying spatial information. There are three interrelated aspects of this research: knowledge acquisition, representation, and manipulation. In knowledge acquisition, the emphasis will be on two projects: the estimation of three-dimensional motion parameters and structures of rigid bodies from image sequences taken with a stationary camera; and the determination of occupancy space for a stationary environment using a moving camera. The results of these two projects will be merged to form techniques for acquiring knowledge of time-varying, three-dimensional environments. In knowledge representation and manipulation, the use of octrees to represent three-dimensional data will be explored and computationally efficient algorithms for scene manipulation (e.g., rotation and translation) in terms of these octrees will be developed.

Johns Hopkins University; Joseph O'Rourke; *Computational and Geometric Aspects of Pattern Recognition and Vision (Computer Research)*; (DCR-8304780 A01); $22,915; 12 mos. (Joint support with the Office of Interdisciplinary Research - Total Grant $45,830).

This research will deal with the development of computational geometry for problems typically arising in pattern recognition and computer vision. Four distinct but interrelated problem areas will be explored. The common threads connecting these areas are their shared geometric flavor and concern with computa-

tional issues. The first area is the exploration of dynamic quantization of imagery data in connection with the detection of object motion. The second area is the study of computational complexity of polygon decompositions, a topic of particular interest to the analysis of complex patterns by simpler, more standard patterns. The third area is closely related to decomposition - polygon and polyhedron internal visibility in what is known as the "Art Gallery" problem - a form of developing strategies for maximizing visual contact in a 2-d or 3-d space with specific geometric constraints. The fourth problem concerns the determination of geometric shapes for curves by identifying intrinsic features, called "signatures". Taken together, these four topics will form the basis of a new framework for computational geometry tailored to intelligent systems design. The results of this research will also have indirect impact in application areas which utilize pattern recognition or computer vision technology.

University of Maryland - College Park; Azriel Rosenfeld; *Multiresolution Image Analysis (Computer Research)*; (DCR-8218408 A02); $172,106; 12 mos.

During the past few years there has been increasing interest in the use of multiresolution ("pyramid") image representations for a variety of image processing and analysis tasks. This approach has two principal advantages. First, it often leads to reduced computational costs through the hierarchical structuring of computations (e.g., "divide-and-conquer"). Second, it has suggested some interesting new types of algorithms, involving cooperation among pixels or features at different resolutions, which may have advantages over conventional approaches, in particular by providing a smoother transition between pixel-level and region-level image representations. This research on multiresolution image analysis techniques will be concerned both with the development of new methods and with their incorporation into an integrated image analysis system which allows pixel-level and region-level knowledge to interact through the medium of multiresolution representations.

This grant supports the summer research of Dr. Robert Melter, a small college faculty member at Southampton College, Long Island University.

University of Maryland - College Park; Azriel Rosenfeld; *Workshop on Human and Machine Vision, Montreal, Canada, August 1-3, 1984*; (DCR-8405503); $7,251; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $14,502)

Workshop on Human and Machine Vision will be held in Montreal, Canada on August 1-3, 1984 in conjunction with the International Conference on Pattern Recognition. This workshop will be a sequel to one held in Denver, Colorado in August 1981, which was regarded as extremely successful. The workshop will bring together experts in computer vision and visual perception to discuss important issues of interest to both groups. Tentative topics to be covered include spatial representation, global processes, perceptual reasoning, multiresolution processes, grouping processes, (pre)attentive processes, and perceptual hardware.

Massachusetts Institute of Technology; Whitman Richards; *Natural Computation: A Computational Approach to Visual Information Processing (Information Science and Computer Research)*; (IST-8312240); $57,976 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $173,929).

Natural computation is the study of how man can derive reliable representations of useful aspects of his environment. These representations may be visual, auditory, or consist of programs for touch sensing and motor control. At present, the data under study are the vast arrays of image intensities cast upon the eye or camera. Given these data, the task is to derive rather simple but reliable descriptions of objects and events in a three-dimensional world. The problems involve choice of representation, constraint analysis and explicit implementation (algorithm) to show that reliable and efficient descriptions can indeed be computed.

The research may be divided roughly in four general problem areas:

1. 3D surfaces reconstruction,
2. Property-based representations,
3. Integrating depth maps obtained from stereo, structure-from-motion, and shading; and
4. Computing spatial relations between "parts" and their combination into "object" representations.

University of Massachusetts - Amherst; Edward Riseman; *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776); $25,000; 12 mos. (Joint support with the Special Projects Program and the Intelligence Systems Program - Total Grant $160,000).

This grant will support research activity within the Computer and Information Science (COINS) Department in six areas of experimental computer research: computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under twenty-one other grants and contracts, including ten from NSF, this grant will provide partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support will be used for administration and

operation of the COINS research facility, with the funding going towards technical support staff and equipment maintenance.

Michigan State University; George C. Stockman; *Sensing and Matching Surfaces for 3d Object Acquisition (Computer Research)*; (DCR-8304011 A01); $35,066; 12 mos.

This research will focus on the use of illumination techniques to identify and locate industrial objects on a table or conveyor. Past research has given ways of acquiring 3D surface shape from 2D images by using engineered lighting, including the projection of planes or rays of light in controlled directions. Two of the tasks of the current research are the extraction and categorization of elementary surface features and the matching of object features to model features to identify and locate an object. A combination of known shape sensing techniques will be used to extract several patches of an object's surface, each representing local surface shape from a specific point of view. Each surface patch taken from an object is likely to correspond to several possible regions of a model: a novel constraint matching algorithm will be investigated to derive a globally good correspondence between object and model. Aspects of the research deal with the sensing of surface shape, geometrical representation, and constraint satisfaction via clustering. The theoretical significance of the research lies in the approach to matching and constraint satisfaction. Success in shape sensing and matching would have great practical significance for robotics and automation.

University of Rochester; Dana H. Ballard; *Parameter Networks: A Connectionist Theory of Vision*; (DCR-8405720); $41,518; 12 mos. (Joint support with the Office of Interdisciplinary Research - Total Grant $83,035).

This research project will develop a massively parallel, connectionist vision model. This model consists of upwards of 10^6 processors. Each processor stands for an interval in some parameter space. Physical constraints are represented by appropriate interconnections between processors. Each processor has a small amount of state that is updated iteratively as part of a relaxation process. Visual gestalts (or segmentations) correspond to fixed points in processor state space. Previous work has shown that this model can analyze complex motion and shape stimuli in only a few parallel iterations, and that this method of handling constraints, which is based on the Hough transform, is insensitive to noise and occulsion.

This research has three principal objectives:

1. Shape recognition. The objective is to extend the view transform mapping approach to shape recognition, first to connectionist models of 3D polyhedral objects, and next to connectionist models of curved objects.
2. Motion. The objective is to build an integrated connectionist motion machine that starts from raw visual data and extracts egomotion parameters.
3. Extrapersonal space. The objective is to use multiple intrinsic image cues to build an egocentric model of immediate spatial surroundings.

If this basic research project is successful, it could lead to advances in a number of application areas that would help the handicapped.

Three of these areas are:

1. Visual Prostheses. This work in shape recognition and motion would lead to devices for relating this information to blind individuals via non-visual channels.
2. Visual Brain Dysfunction. A detailed characterization of the mechanisms for storing and manipulating visual information would lead to refined assessments and possibly new treatments for stroke-impaired patients.
3. Visual Learning. Knowledge about human shape representation would be invaluable in the study of dyslexia as well as the intellectually handicapped and could lead to new learning methods.

University of Rochester; Christopher M. Brown; *Theory and Techniques for Low-Level Vision: Hough Transform, Color, and Motion (Computer Research)*; (DCR-8302038 A01); $28,033; 12 mos. (Joint support with the Office of Interdisciplinary Research - Total Grant $56,065).

Toward the goal of an advanced integrated computer vision system, this research focuses on two branches of study that are also of interest in their own right. The first is advanced implementations and algorithms using the Hough Transform, a mode-based parameter estimation strategy. The work will involve basic mathematical and statistical analysis, computer simulations, and application to real data. The second is basic studies and computer implementations leading to theories of color and motion perception, especially emphasizing the interaction of color and motion perception with the perception of related, interacting properties such as surface orientation and reflectance. The anticipated results are at the level of algorithms and analysis, and should be applicable to both traditional and massively parallel computational architectures.

Carnegie-Mellon University; Kendall J. Preston; *Logical Transform Research*; (DCR-8311207); $77,871; 18 mos.

A sampled and digitized real function of n variables is a logical function of (n + 1) variables. For example, when the real function $z = \cos(x)$ is sampled, one obtains $z_i = \cos(x_i)$ where $-N \leq x_i \leq +N$. When z_i is

digitized and stored as a b-bit integer, its values will lie within the ranges $-M \leq z_i \leq +M$ where $M = 2^b - 1$. Thus, we may graph all values of z_i within the ranges specified in a (2N) × (2M) array. Each element of the array will have the logical value 0 or 1. Logical transforms may then be performed on the elements of the array to produce useful results. These results may be characterized by non-linear transfer functions. The purpose of this reseach is to study the properties of these transfer functions in both the spatial domain and in the frequency domain and to develop analytical tools for predicting their performance.

Carnegie-Mellon University; Kendall J. Preston; *Workshop On Multicomputers And Image Processing, Tucson, Arizona, May 22-24, 1984*; (DCR-8315057); $8,580; 12 mos. (Joint support with the Computer Systems Design Program - Total Grant $17,160).

Starting in May 1979, there have been annual workshops involving a small group of scintists who are responsible for most of the major research efforts worldwide in multicomputer systems related to image processing. These systems include the CLIP and DAP systems (Great Britain), PICAP (Sweden), DIP (Holland), AT4 and SYMPATI (France), FLIP (West Germany), PPP and CYBEST (Japan), CYTOS, GLOPR, ZMOB, MPP, PUJMPS, PASM (United States). These workshops have acted to stimulate research in this field and this year will be held in May 1984. The workshop will be chaired by Carnegie-Mellon University (Department of Electrical Engineering) and co-chaired by the University of Wisconsin (Department of Computer Science). In conjunction with this particular workshop, these two universities will conduct pre-workshop benchmark studies on many of the above multicomputer systems.

Rice University; Rui J. P. De Figueiredo; *Syntactic/Semantic Methods For 3D Object Reconstruction*; (DCR-8318514); $40,000; 24 mos. (Joint support with Computer Systems Design Program - Total Grant $83,382).

A high-level syntactic/semantic approach to the reconstruction of a surface delimiting a three-dimensional object is being studied. The surface is reconstructed from a set of points on contours present on the surface. Polygonal approximations are represented by attributed strings belonging to an appropriately constructed language. The surface is reconstructed by interconnecting adjoining strings according to their similarities in shape attributes, using advance concepts from syntactic/semantic pattern recognition theory.

University of Washington; Steven L. Tanimoto; *Interactive Parallel Image Processing (Computer Research)*; (DCR-8310410); $60,994; 12 mos.

Parallel computers for image processing have been entering the commercial sector gradually over the last several years. The range of applications of these systems is still relatively restricted because there remains a gap in the understanding of the relationship between the operations that these machines perform (cellular logic, image filtering, etc.), and the algorithms needed to perform useful image analysis. This research will employ a progression of models (for parallel image processing) to study the interactions of cellular logic operations with global image operations. The major goal is to improve the understanding of cellular logic, both conventional and pyramidal, and its interaction with overall image features.

Three tasks will be carried out to accomplish this goal:

1. Development of a capability for representing cellular logic operations in a manner which facilitates human comprehension at both theoretical and empirical levels (the representation is to be supported by an experimental software system),
2. A study of hierarchical methods in image analysis, and
3. The development of some specific algorithms in the domains of industrial automation and document processing for simple, well-defined tasks.

Natural Language and Signal Understanding

Stanford University; K. Jon Barwise; *Computational Aspects of Situation Semantics (Computer Research)*; (DCR-8403573); $9,693; 12 mos. (Joint support with the Theoretical Computer Science Program and the Division of Information Science and Technology - Total Grant $29,079).

Situation semantics is a mathematical theory of linguistic meaning that grew out of artificial intelligence work in natural language processing. The goals of research in situation semantics include developing computationally feasible procedures for natural language understanding systems, developing the mathematical theory of situation semantics, and comparing the efficiency of situation semantics to more standard approaches to the theory of linguistic meaning. This is a collaborative project with Professor John Perry at Stanford University.

Stanford University; John M. Chowning; *An Intelligent System for the Knowledge-Driven Analysis of Acoustic Signals (Computer Research)*; (DCR-8444174); $166,037; 12 mos.

The central concern of the research proposed is the integration of signal processing and context analysis into a goal-directed analysis based system for the understanding of complex structured acoustic signals. There are two aspects to this approach to the problem. First, a control strategy will be designed to employ all levels of analysis, including the signal processing, in a common control mechanism. This is in contrast to previous efforts that have viewed signal processing as a frontend pipeline process. The control strategy will be developed using heuristics and resource allocation algorithms to invoke relevant knowledge sources at appropriate processing stages The second aspect is the guiding of the signal processing by the higher levels of analysis so that the signal processing tools can be invoked to reprocess certain portions of the incoming signals, and the higher levels can feed back parameters to direct the signal processing to a more accurate analysis of the signal. The results of this research are expected to help remove some of the limits computer analysis of complex signals has now reached.

Stanford University; Stanley Peters; *Toward Automated Natural Language Processing: Phrase Linking Grammars for Syntax and Semantics (Information Science and Computer Research)*; (IST-8314396 A01); $17,000; 12 mos. (Joint support with the Theoretical Computer Science Program and the Division of Information Science and Technology - Total Grant $141,744).

This is the third and final year of a research project focused on the characterization and understanding of natural language. An information structure called a linked tree is being used which will allow a syntactic and semantic interpretation of natural language and the process of understanding natural language.

The authors are developing a restricted syntactic theory for the description of natural language, integration of semantics with syntax; and developing a theory of parsing and interpretation based on psycholinguistic evidence.

Stanford University; Ivan A. Sag; *International Workshop on Language Generation, Stanford, California, July 1984*; (DCR-8411612); $6,042; 6 mos. (Joint support with the Linguistics Program - Total Grant $12,083).

Natural-language generation is becoming an increasingly important area of investigation in artificial intelligence. Developing an adequate theory requires one to address many issues in the areas of syntax, semantics, pragmatics, discourse, psychology, and representation of knowledge. Not only is language generation an important applications area, but also it draws upon and contributes to basic research issues in each of these areas. To increase our understanding of these areas of basic research and their relation to language generation and to enable researchers to keep abreast of recent developments in this field, the Center for the Study of Language and Information at Stanford University will hold a three-day workshop to address the theoretical and computational issues involved. The workshop will consider problems in grammatical formalisms, speech-act theory, utterance planning, producing explanations psycholinguistic modeling, production of extended discourse, knowledge representation and modeling of the hearer that are directly related to language generation. Comparisons will be made of the approaches to these issues taken by a number of language-generation systems.

University of Southern California; William Mann; *Constructive Characterization of Text (Information Science and Computer Research)*; (IST-8408726); $59,996; 24 mos. (Joint support with the Division of Information Science and Technology - Total Grant $119,992).

This research seeks to characterize text constructively, i.e. to create automatic methods which produce textual descriptions of numerical data bases.

The project explores production of two kinds of text:

1. Scientific abstracts, and
2. Descriptions of collections of numerical data (such as instrument readings).

It seeks to produce a theory which covers both kinds of text, and to embody this theory in a computer program which is capable of producing fluent English description of collections of numerical data.

The significance of this research lies in the contribution it will make to problems of having computers generate informative written reports from new data.

Duke University; Bruce W. Ballard; *Workshop on Transportable Natural Language Processing - Durham, North Carolina - October, 1984 (Computer Research)*; (DCR-8402803); $6,496; 12 mos. (Joint support with the Division of Information Science and Technology- Total Grant $12,992).

During the 1970's, a number of experimental natural language processors proved themselves capable of productive use by untrained computer users. Unfortunately, each of these systems is quite limited in the "domain" of data it is able to deal with. To combat this problem, many researchers are seeking to devise methods that allow users to adapt an existing natural language system to deal with new types of data. In addition to the obvious practical value of these efforts, many valuable scientific contributions have already been made, with notable progress in the areas of knowledge representation and knowledge acquisition. Given the timely and important nature of the topic, a

workshop is being conducted to consider the problems associated with the design of transportable natural language processors.

Colgate University; Allen B. Tucker; *Knowledge-Based Multilingual Machine Translation*; (DCR-8407114); $32,100; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $47,100).

The early extensive effort in machine translation (MT) which started with Weaver's memorandum (1949) did not produce adequate result, although much was contributed to the development of Artificial Intelligence and Computational Linguistics. The work on MT continued with much less funding and with emphasis mainly on operational, not theoretical issues. The developments gathered momentum with the spectacular improvements in the computational support over the past 20 years, so that at present several MT systems are in operation. Recently, several MT projects have been initiated in the European Economic Community, Japan, Israel and Soviet Union. Significant development of the science and technology of Artificial Intelligence and Knowledge Based Expert Systems has led to the emergence of the proposition of a knowledge based MT system which has a substantial real world knowledge base in addition to knowledge bases of the source and target languages.

A knowledge-based model for multilingual MT is being developed and tested. The model will combine linguistic morphological and syntactic parsing technology in current MT systems with "world knowledge" and "short-term memory" functions of recent Artificial Intelligence research.

Three major goals of the research are to develop and test:

1. A universal intermediate language (IL) for representing the syntactic and semantic information in a text and the knowledge base itself;
2. A knowledge base for a specific subject domain, written in IL, which is well adapted to multilingual MT, and contains linguistic knowledge as well as world knowledge;
3. An objective measure of the quality of a MT system. An English corpus of 10,000 words in machine readable form will be used as a basis for this experiment.

New York University; Ralph Grishman (In collaboration with Lynette Hirschman at the Burroughs Corporation); *Industry/University Cooperative Research: Robust Natural Language Parsing Using Graded Acceptability (Computer Research)*; (DCR-8202373 A02); $33,400; 12 mos. (Joint support with the Division of Industrial Science and Technological Innovation - Total Grant $66,800).

Current computer systems for natural language analysis typically fail to parse 10-20% of their input. Many of these sentences fail because the parser accepts only a "perfect" analysis, namely one that meets a whole set of detailed syntactic and semantic restrictions. The robustness of a natural language system would be considerably improved if the parser accepted the best available analysis, in case no "perfect" analysis were available. We propose to investigate two techniques for preferential parsing which would provide the parser with greater flexibility and robustness. The first technique is restriction relaxation. If the parser can not obtain a perfect analysis, it tries to obtain an analysis where one of a designated set of restrictions is allowed to fail. This technique will also be explored as a debugging aid, to identify the restriction blocking an analysis. The second technique to be investigated involves a weighting of alternative analyses, based on observed frequency occurrence of syntactic and semantic patterns in a parsed sample corpus. The results of parsing with restriction relaxation and parsing with weighting will be compared to the results of the current parsing algorithm.

Burroughs Corporation; Lynette Hirschman (In collaboration with Ralph Grishman at New York University); *Industry/University Cooperative Research: Robust Natural Language Parsing Using Graded Acceptability (Computer Research)*; (DCR-8202397 A03); $9,806; 12 mos. (Joint support with the Division of Industrial Science and Technological Innovation - Total Grant $19,611).

Current computer systems for natural language analysis typically fail to parse 10-20% of their input. Many of these sentences fail because the parser accepts only a "perfect" analysis, namely one that meets a whole set of detailed syntactic and semantic restrictions. The robustness of a natural language system would be considerably improved if the parser accepted the best available analysis, in case no "perfect" analysis were available. We propose to investigate two techniques for preferential parsing which would provide the parser with greater flexibility and robustness. The first technique is restriction relaxation. If the parser can not obtain a perfect analysis, it tries to obtain an analysis where one of a designated set of restrictions is allowed to fail. This technique will also be explored as a debugging aid, to identify the restriction blocking an analysis. The second technique to be investigated involves a weighting of alternative analyses, based on observed frequency occurrence of syntactic and semantic patterns in a parsed sample corpus. The results of parsing with restriction relaxation and parsing with weighting will be compared to the results of the current parsing algorithm.

Carnegie-Mellon University; Raj Reddy, Omer Akin and Ronald Cole; *Knowledge Acquisition and Knowledge Engineering in Speech Understanding (Computer Research)*; (DCR-8205539 A02); $174,708; 12 mos.

Although the area of automatic speech recognition has received increasing attention in recent years, there are still a number of unsolved problems that prevent the development of recognition systems that approximate human performance. Two of the most basic problems are the ability to perform speaker-independent recognition and the ability to perform fine phonetic distinctions. Today's systems require substantial training to adapt to a new speaker, and are unable to deal with confusable vocabularies. For example, the letters "B", "D", "P" and "T" cannot be recognized by most systems with better than 70% to 80% accuracy after training with a single speaker. Error rates of 20% to 30% are simply not good enough to achieve acceptable performance in natural task domains. In addition, adaptive techniques used in present systems are neither generalizable nor extensible to natural task domains. This research deals with a new methodology to find solutions to these problems. Specificially, knowledge acquisition and knowledge engineering techniques will be explored to develop automatic and interactive schemes for adapting a speech understanding system to new speakers, new words and new environments.

University of Pennsylvania; Bonnie L. Webber; *Extended Natural Language Interaction with Database Systems (Computer Research)*; (DCR-8305221 A01); $37,171; 12 mos. (Joint support with the Division of Informatiom Science and Technology - Total Grant $74,342).

This research project is concerned with several major areas in natural language processing dealing with several aspects of cooperative behavior in the context of question-answer systems, and addressing certain key issues in man-machine interactions through natural language. In particular, the project will investigate (1) cooperative responses for correcting misconceptions about the structure of the data base and for dealing with dynamic data bases, (2) the natural language generation of descriptions and explanations in question-answer systems, and (3) the role of mutual beliefs in question-answer systems. Other parts of the research deal with some issues in control of inferences and relationship of grammars to parsing strategies, both of which involve study of some formal systems.

Automatic Theorem Proving and Inference

Stanford University; John McCarthy; *Mechanical Theorem Proving and Development of EKL - An Interactive Proof Checking System (Computer Research)*; (DCR-8206565 A02); $146,391; 12 mos.

The main tool in making a computer prove theorems or follow reasoning is an interactive proof checker. In this research this tool will be used to represent the facts involved in an intellectual problem and check that that representation is adequate to solve the problem. The development of such a tool will be based on an interactive proof-checker, called EKL.

The main distinguishing characteristic of EKL from other existing systems is its flexibility and adaptability to different mathematical environments. The emphasis will be on creating a system which would allow the expression and verification of mathematical facts to be made in a direct and natural way. The development of EKL will be heavily weighted in favor of expressibility and user friendliness as opposed to sophistication in decision procedures. The overall goal of this research, therefore, is to provide a friendly environment for formal manipulation.

Northwestern University; Lawrence J. Henschen; *Logic and Databases (Computer Research)*; (DCR-8306637 A01); $61,611; 12 mos.

The purpose of this research is to enhance the information retrieval capabilities of data bases by incorporating first-order logic into the data base. The relational data model will be used in which the relations themselves already are of the form of elementary logical statements. With this model, it is only necessary to supply the appropriate logical mechanisms for integrating more general logical formulas into the data storage and retrieval processes. It has been shown such a model can be used when the logical formulas represent definitions of new relations, even when such definitions are recursive. This allows the possiblity of greatly reducing storage because such defined relations need not be stored; rather, their information content can easily and efficiently be generated when needed. This can be accomplished at data base creation time, a major step forward. In this study other types of formulas (e.g., constraints) will be examined to determine whether they can be "compiled" at data base creation

time. The research will also look into how functions, which normally lead to infinite deduction paths, can be handled in the specialized domain of data bases.

University of Illinois - Urbana; Ryszard S. Michalski, *Studies in Computer Inductive Learning and Plausible Inference*; (DCR-8406801); $131,384; 12 mos.

This research is concerned with computer inductive learning and formal models of human plausible inference. It encompasses basic theoretical studies as well as the development of algorithms and their computer implementations and experimental applications to selected practical problems. Major topics of the research include:

1. Automated generation of derived descriptors (variables, relations, functions) which are most relevant for constructing inductive assertions.
2. The development of rules that govern the generalization processes and efficient strategies for their application in inductive inference.
3. The study and implementation of techniques for constructing conceptual hierarchies of objects.
4. The development of experimental multi-purpose high performance inductive learning systems which incorporate the methodology for generating derived descriptors.
5. The application of inductive inference to the design of computer architectures and compiler structures.

University of Maryland - College Park; Jack Minker; *Artificial Intelligence, Parallel Logic Programming, and Deductive Databases (Computer Research)*; (DCR-8305992 A01); $47,142; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $94,284).

The major thrust of this research is the application of automated theorem proving and mathematical logic to database investigations. Three types will be treated:

1. Theorem proving by linear resolution with unrestricted selection function using trees ,
2. First-order logic and databases, and
3. Database logic.

Database logic is a recently developed formalism whereby existing database systems can be viewed in terms of logic. It has the prospect of providing a unifying theoretical framework for database systems.

This supplement supports the research of Dr. John Grant, a small college faculty member at Towson State University, Towson, Maryland.

University of Massachusetts - Amherst; Donald Geman; *Applications of Stochastic Relaxation and Simulated Annealing to Problems of Inference and Optimization (Mathematical Sciences and Computer Research)*; (DMS-8401927); $15,000; 24 mos. (Joint support with the Statistics and Probability Program and the Applied Mathematics Program - Total Grant $50,800).

This proposal brings ideas and methods of probability, statistics and computer science to bear on problems of "stochastic relaxation" algorithms. These algorithms have great use in image processing, traveling salesman problems, neural modeling and certain "expert systems". The analogy between the systems involved in the inference and optimization problems of these algorithms with the stochastic physical systems of statistical mechanics brings a rich structure whose exploitation has enormous potential for treating a variety of difficult theoretical issues connected with the algorithms.

University of Massachusetts - Amherst; John Moore and Andrew Barto; *Adaptive Element Models of Classical Conditioning (Psychobiology and Computer Research)*; (BNS-8317920); $12,940; 24 mos. (Joint support with the Psychobiology Program - Total Grant $25,880).

Dr. Moore, a physiological psychologist, and Dr. Barto, a computer scientist, are collaborating in an investigation of the properties of one form of associative learning. The aim of their research is to enlarge the empirical foundation of certain biologically-inspired theories of learning that have been applied to problems of machine learning, cybernetics, and artificial intelligence. The theories at issue describe how conditioned responses unfold in time under complex contingencies and speculate on how classical Pavlovian conditioning and other types of associative learning might be represented within nervous systems. The primary objective of this research project is to investigate the validity of this theoretical approach. Such model-guided studies of animal learning have relevance for artificial intelligence. By studying the conditions under which animals acquire conditioned responses, for example, one may gain insight into the algorithms they have evolved for solving difficult problems of recognition, selection, and motor control. Such algorithms find application in developing efficient computer-based systems that perform these functions within an engineering context.

Carnegie-Mellon University; Peter B. Andrews; *Automated Theorem Proving in Type Theory (Computer Research)*; (DCR-8402532); $82,327; 12 mos.

This project is directed toward enabling computers to construct and to check proofs of theorems of mathematics and other disciplines formalized in type theory (higher order logic) or first order logic, and to assist humans engaged in these tasks.

Previous research has shown that one can search for a proof of a theorem of higher order logic by searching

for an expansion tree proof (ET-proof) of the theorem, and then transforming this into a proof in natural deduction style. The proposed research will focus on developing and implementing a procedure for finding ET-proofs, and on related theoretical questions.

An existing computerized theorem proving system called TPS will be extended to facilitate generating and manipulating examples useful in this research, and to translate back and forth between ET-proofs and natural deduction proofs. TPS will also be enhanced as a practical and convenient tool for constructing and checking proofs interactively, semi-automatically, or automatically.

University of Texas - Austin; Woodrow W. Bledsoe; *Automati. Theorem Proving and Applications (Computer Research)*; (DCR-8313499); $143,681; 12 mos.

This research is a continuation of present work in automatic theorem proving and its applications to program verification and mathematics. Earlier work has included the development of computer programs which: prove theorems as part of a program verification system; prove theorems from intermediate analysis with techniques for instantiating set variables and for handling general inequalities. Other work has been on complete sets of reducers and agenda driven provers. This investigation will continue the work, especially in program verification and complete sets of reducers, and in continuing the effort to build a more powerful general theorem prover.

Knowledge Representation and Problem Solving

Stanford University; Edward A. Feigenbaum and Charles Yanofsky; *MOLGEN—Applications of Artificial Intelligence to Molecular Biology Research in Theory Formation, Testing, and Modification*; (DCR-8310236); $139,215; 12 mos.

The problem of scientific theory formation, modification, and testing is important for both of the domains of artificial intelligence and molecular biology. Using the tryptophan operon model of regulatory genetics developed by Professor Charles Yanofsky as an initial test case, a system to automatically construct and extend such biological models will be built.

A blackboard mechanism will be initially employed in order to allow the many different types of reasoning important in theory formation to be considered. These reasoning types will include model-driven, data-driven and analogical (both to closely related biological systems and to other domains like computer science). Several tools to assist scientists in testing and discriminating among alternative models are being developed. The research will include aspects of discovery systems and automatic knowledge acquisition and will draw upon previous MOLGEN work in experiment design and knowledge base construction.

Stanford University; Edward A. Feigenbaum; *The Mechanization of Formal Reasoning (Computer Research)*; (DCR-8303142 A01); $98,657; 12 mos.

This project, with principal investigator, Edward Feigenbaum, will continue work on the mechanization of formal reasoning for artificial intelligence systems. Based on past research, it will focus on three new and extended components:

1. Continued theoretical research in FOL, a system capable of three different levels of reasoning: Reasoning at the object level, meta-level and self-referring.
2. Experimental development of an expert system using the FOL framework.
3. Exploration of the role of formal reasoning in the design of high-level architecture for artificial intelligence.

University of California - Los Angeles; Walter Karplus and Judea Pearl; *Studies in Heuristics (Computer Research)*; (DCR-8114209 A03); $62,831; 12 mos.

The primary objective of these studies is to continue the analytical investigations of the mathematical properties of heuristics and their influence on the performance of common search techniques. These include: quantifying the relations between the precision of the heuristic models used and the complexity of the search they help guide; comparisons of time-storage tradeoffs for various graph searching algorithms; procedures for comparing, improving, and combining heuristic functions; assessing the complexities and the quality of decisions of game-playing programs. In addition to these, new studies in the areas of mechanical generation of heuristics and the use of parallelism in graph searching will be initiated.

University of California - Los Angeles; Judea Pearl and Moshe Ben-Bassat; *Toward a Computational Model of Evidential Reasoning*; (DCR-8313875); $62,028; 12 mos. (Joint support with

the Decision & Management Science Program - Total Grant $74,028).

This research program is directed at the development of computational models for evidential reasoning. That is, inference mechanisms by which the evidence provided by a set of findings is analyzed in order to gain better understanding of a given situation or phenomenon. Problem solving tasks of this nature include, for instance, medical diagnosis, weather forecasting, corporate assessment, political crisis assessment and battlefield reading. These problems are referred to as situation assessment tasks.

The approach is based on the decomposition of the situation assessment task into evidence integration, hypotheses generation and evaluation, goal(s) setting, information sources evaluation and selection, and sorting of evidence by goals and hypotheses. Inference nets are used for knowledge representation in which the nodes represent observable and inferred events, while a link between nodes E_i and E_j indicates evidential relevancy between the two links. Each link is assigned value(s) that represent the degree of significance for inferring E_i from E_j.

Yale University; Drew V. McDermott; *Spatial Reasoning for Problem Solving (Computer Research)*; (DCR-8407077); $57,425; 12 mos. (Joint support with the Office Interdisciplinary Research - Total Grant $114,850).

The goal of this project is to investigate approximate shape representations to support reasoning about changing systems. Previous work on "fuzzy maps" will be generalized so that the fuzzy numbers are attached to shape representations at different levels of resolution. What it means for an object to have multiple shape representations will be formalized. Surface representations will be treated as primary, and characterize their exactness by two parameters: the "gap size," which bounds the size of holes neglected in abstracting this surface; and the "grain size," which bounds the discrepancies between the idealized surface and the actual surface. Such approximate shape representations can be used to support reasoning about motion, shape change, and fluid flow. They provide a convenient way of abstracting "faces" from arbitrary shapes. These faces can then be the focus of reasoning about degrees of freedom of motion, fluid containment, and so forth.

If handicapped people are to profit from robotics, they will need mobile robots that have the same ability to contemplate and manipulate representations of a changing, incompletely known world. Such robots could be used in two ways: as autonomous servants, and as direct symbionts of a disabled person. In the first mode, the robot would wander about a house or workplace detached from its user. It would respond to commands to go and retrieve objects, or perform simple tasks like unlocking a door to let someone in. In the second mode, the robot would become part of its user, controlling prosthetic legs and hands. This research is directed toward developing the spatial and temporal reasoning that will be needed by such mobile robots.

Illinois Institute of Technology; Shi-Kuo Chang; *An Intelligent System for Crisis Management (Computer Research)*; (DCR-8306282); $64,944; 12 mos.

This project examines the methodology for designing intelligent systems for the management of imprecise data, database alerters, and message filters, with applications to crisis management. To accomplish this, it is first necessary to specify database alerters and message filters based upon the concept of imprecise database.

This specification will then lead to the synthesis and analysis of information systems using the Office Procedure Model(OPM). Specific topic include the following:

1. Generalization of OPM modeling technique to allow for multiple message receptions, which corresponds to k-bounded Petri-net.
2. Investigation of efficient algorithms for message protocols analysis.
3. Development of mathematical model for analysis of system stability.
4. Investigation of message filter design methodology to balance message traffic during peak traffic hours.

Both theoretical investigations and experimental studies will be pursued during the grant period.

Illinois Institute of Technology; Martha W. Evens; *An Expert System Which Can Do Deep Reasoning (Computer Research)*; (DCR-8216432 A01); $71,021; 12 mos.

Research in expert systems capable of deep reasoning using a fast, flexible and powerful theorem prover as the inference engine is proposed. During recent years the design of expert systems has improved greatly in the areas of knowledge representation and planning, but progress in inferencing has been limited. This research will concentrate on inferencing and will take advantage of the current developments in theorem proving. The field of organic synthesis has been chosen as the problem domain since it is a problem well known to require deep reasoning with a long history in artificial intelligence. As part of the research, a series of experiments will be conducted on two of the most powerful theorem provers ever developed (AURA for Automated Reasoning Assistant and LMA for Logic Machine Architecture). The experiments will involve alternative ways of representing molecules and reac-

tion rules as axioms. The result will be a powerful and portable expert system in organic synthesis capable of doing reasoning and making inferences, an important feature not widely available in current expert systems.

Tulane University; Larry H. Reeker; *Methodology and Evaluation of Specialized Information Extraction from Scientific Texts (Information Science and Computer Research)*; (IST-8410510); $30,000; 24 mos. (Joint support with the Division of Information Science and Technology - Total Grant $78,010).

This research studies the general problem of automatic specialized information extraction using as a vehicle reaction information from articles in the chemical literature. Among the problems being studied are the use of information "clues" to indicate documents the system cannot handle directly. The problem of modularization of the system is being studied so that specialized details of discourse can be separated from the more general semantic portion.

Massachusetts Institute of Technology; Benjamin J. Kuipers; *Knowledge Representations for Expert Causal Models (Computer Research)*; (DCR-8417934); $77,000; 12 mos.

The goal of this research is to develop and test a knowledge representation capable of understanding a realistic case description, constructing a detailed model of the internal state of the patient and its evolution, and answering questions about the pathophysiology of the particular case. As AI research, it will focus on a representation for causal relationships in physiological mechanisms that includes a qualitative language for describing the structure of a mechanism and a method for envisioning the possible behaviors of that mechanism. This knowledge representation will then be used as the basis for a program to reason about disorders of sodium and water balance. This limited, but natural and complex medical domain, is well-understood scientifically, yet poses complex clinical problems. The program will assimilate information about a patient in the format of a natural case description, and will determine the diagnosis and formulate a detailed description of the patient's internal state. Finally the behavior of the program will be evaluated according to its ability to answer the types of questions that might be asked of a resident during clinical training. The research method rests heavily on the detailed observation of human experts, both to capture the knowledge used and to determine the representation for that knowledge.

University of Massachusetts - Amherst; Victor R. Lesser; (Coordination in Cooperative Distributed Problem Solving Systems (Computer Research); (DCR-8300239 A01); $60,283; 12 mos. (Joint support with the Computer Systems Design Program and the Special Projects Program - Total Grant $120,847).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts-Amherst has included the development of a node architecture capable of such sophisticated local decision making. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. The research being carried out will build on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include: organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

University of Michigan - Ann Arbor; Arthur W. Burks; *Languages and Architectures for Parallel Computing with Classifier Systems (Computer Research)*; (DCR-8305830 A01); 12 mos. (Joint support with the Computer Systems Design Program - Total Grant $102,827).

A new class of rule-based computing systems, classifier systems, has been developed and tested. While classifier systems are similar to production systems, they are different in basic ways which facilitate the incorporation of learning algorithms in them, and which make it possible for them to be executed in a highly parallel fashion without the usual problems of concurrency and interlock. Classifier systems with learning algorithms have been tested successfully on two model control systems, a game learning program, and a hardware robot project. Progress has been made on the theoretical understanding of how such classifier systems work.

This research will focus on two interrelated areas. (1) The powers and limits of classifier systems will be studied by further computer simulations and theoretical analyses. (2) Novel parallel architectures for the rapid execution of classifier systems will be planned and simulated. On the theoretical side, these studies will contribute to a deeper understanding of concurrency and distributed control, and that on the practical side they will lead ultimately to powerful

parallel computers which are non-von-Neumann in both programming and architecture.

Rutgers University - Busch Campus; Raymond Reiter; *Logical Foundations for the Representation of Incomplete Knowledge in Artificial Intelligence and Database Theory*; (DCR-8203954 A03); $48,655; 12 mos.

A central problem for both Artificial Intelligence and Database Theory is that of representing incomplete information about the real world, and reasoning with this representation. This research project addresses various aspects of this problem of incomplete knowledge, with emphasis on the following issues:

1. The closed world assumption and its relationship to circumscription.
2. A proof theory for a certain class of default theories and its relationship to the design of large parallel hardware semantic networks.
3. Query evaluation for relational databases with null values and disjunctive information.

Rutgers University - Busch Campus; Natesa S. Sridharan; *Exploration of Problem Reformulation and Strategy Acquisition (Computer Research)*; (DCR-8318075); $72,532; 12 mos. (Joint support with the Division of Informatiom Science and Technology - Total Grant $96,710)

The problem solving ability of an AI program is critically dependent on the nature of the symbolic formulation of the problem given to the program. Improvement in performance of the problem solving program can be made not only by improving the strategy of directing search but also by shifting the problem formulation to a more appropriate form. Certain formulations are more amenable to incremental acquisition of problem solving strategy than others. With this in mind, an extensible problem reduction method will be developed that allows incremental strategy construction. The overall objective of this research project is to study strategy acquisition and problem reformulation for planning problems.

Cornell University; Robert L. Constable; *Experiments With A Program Refinement System*; (DCR-8303327); $66,000; 12 mos. (Joint support with the Software Engineering Program and the Software System Science Program - Total Grant $199,308).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project will explore a logic for formal top-down development called refinement logic. During the previous grant period a system, called Program Refinement Logic: PRL, was designed and implemented. During the coming grant period this system will be improved and expanded and considerable experimentation will be conducted.

Cornell University; John Hopcroft and Alan Demers; *A Program of Research in Robotics*; (ESC-8312096); $40,000; 12 mos. (Joint support with the Computer Engineering Program and the Theoretical Computer Science Program - Total Grant $248,291).

This is a multidisciplinary program of research in robotics.

The research program has two components:

1. A design and implementation component to provide a robotic environment in which to validate new algorithms and ideas on programming language and environment design, and
2. A theoretical program of research aimed at developing the underlying mathematical foundations necessary for robotics.

The design and implementation efforts have centered on developing a system to allow experimentation with a number of different aspects of the robot programming environment. One of the most important of these is the problem of object representation for single objects and generic classes of objects as well. For example, the system should allow a user to describe the concept bolt independent of any specific example. Other examples of areas to be investigated are the problems of motion and task planning algorithms, generic transformations (e.g. bolting, welding), and graphics-based user interfaces. The initial environment will be directed toward automated assembly, but the system will serve much more as a test-bed for new ideas in language and algorithm design than as a production robot programming environment.

New York University; Robert Hummel; *Design Methods for Cooperative Labeling Processes*; (DCR-8403300); $25,036; 12 mos.

Labeling problems occur when a collection of distinguishable objects must be assigned individual interpretations in a consistent manner compatible with initial estimates and evidence. Examples of this very general class of problems occur in computer vision, speech recognition, cryptography, and numerous other domains of information processing. Relaxation labeling methods were introduced in the mid 1970's in order to provide a framework for studying labeling problems. Generally speaking, relaxation labeling is a process which uses local compatibility information to iteratively update an initial assignment of labels to objects so as to change either the assignment or weights representing degrees of assignments, to achieve an improved and less ambiguous labeling.

Professor Robert Hummel and Professor Steven Zucker (McGill University) have recently developed a

model characterizing the goal and the parameters in the relaxation labeling process and a convergence theorem guaranteeing reasonable behavior of the iterative procedure.

The proposal research consists of two categories:

1. Theoretical work involves extensions to the theory, improved convergence results, a comparison with other iterative learning and labeling network schemes, and the development of approximations to exact analytical formulas and algorithms by more "biologically plausible" methods,
2. Empirical investigations will center on applications involving edge interpretation labeling, pixel classification for effective image segmentation and analysis, and matching 3D image data structures.

State University of New York - Buffalo; Shoshana L. Hardt; *Naive and Expert Reasoning about the Physics of Diffusional Processes (Computer Research)*; (DCR-8305249 A01); $23,093; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $45,093)

This project centers on the design and implementation of a knowledge based reasoning system for qualitative reasoning about physical aspects of the process of diffusion in small structures. Typical questions deal with situations in which objects are released from a source and wander randomly in some available space until they encounter a sink which eliminates them. The function of the reasoning system is to estimate the source-sink diffusion time and its dependency on physical parameters such as the size and the shape of the confining space, the size and the shape of the source and the sink and the nature of the medium. Three separate reasoning systems will be developed in an attempt to model three levels of expertise observed in humans who perform this reasoning task. The first system uses only naive physics of everyday (inertia) events to reason about diffusion and therefore gets most of the answers wrong, the second system uses naive physics with exception rules to deal with the relevant special features of diffusion processes, and the third uses a specialized set of core concepts and heuristics which allows it to perform expert reasoning about diffusion. The construction and the investigation of these three computer systems are aimed primarily to explore the functionally significant differences between them and the relevance of these differences to the process of knowledge acquisition from experience.

State University of New York - Buffalo; John Myhill; *An Expert System for Schenkerian Synthesis of Chorales in the Style of J. S. Bach (Computer Research)*; (DCR-8316665); $117,433; 24 mos.

This research is centered upon a knowledge-based expert system for generating four voice chorales in the style of Johann Sebastian Bach. An attempt will be made to describe the Bach chorale style via approximately one hundred and fifty rules written mostly in first order predicate calculus. The rules will be partitioned into three groups which observe the chorale from the harmonic, melodic, and Schenkerian anlaysis points of view, respectively. A program which generates chorales from left to right, all four voices in parallel, and performs intelligent backtracking until a solution satisfying all the constraints is found will be implemented. A substantial number of heuristics will be used for biasing the search toward musical solutions.

Ohio State University; Balakrish Chandrasekaran; *Knowledge Organization and Problem Solving for Diagnostic Reasoning (Computer Research)*; (DCR-8443219); $83,021; 12 mos.

This research is a continuation of investigation, with new emphasis, into knowledge-based problem solving systems, especially those of a diagnostic nature. The proposed methodology is characterized by a decomposition of domain knowledge first into types of problem solving and then, for each problem solving type, into specialists. Each problem solving type implies a particular control scheme, which is embedded in the specialists. Several issues in the coordination of specialists within and among problem solving types will be investigated. Also, the relationship between so-called "deep" reasoning structures and "compiled" knowledge structures will be explored. The field of medical diagnosis has been chosen for exploring the issues, but the underlying principles will apply to other fields as well.

Saint Joseph's University; Ranan B. Banerji; *Knowledge Based Learning and Problem Solving Heuristics (Computer Research)*; (DCR-8217964 A01); $92,881; 24 mos.

During a previous grant period a computer program was developed which could learn general descriptions of sets of related objects and of relations between objects. This learning was done on the basis of examples of such objects and relations. The efficiency and performance of this program depended heavily on the order of presentation of the examples: simple ideas had to be illustrated by examples before more complicated examples could be presented.

While such "graded" learning exemplifies an artificial learning environment (e.g. in schools), "real life" learning (e.g. the development of science) cannot be modelled by such a program. During the current grant period, a program will be developed which can

learn by unstructured examples and develop its own simplifying concepts that could be used for the simplification of subsequent learning.

The application of this technique will be attempted on the improvement of the problem solving ability of computer programs.

Software Engineering

Quality Software

Stanford University; George Dantzig, Richard Cottle and Walter Murray; *Computer Optimization of Complex Systems*; (DCR-7926009 A04); $160,000; 12 mos.

The Systems Optimization Laboratory (SOL), Department of Operations Research, Stanford University, is a center for research in mathematical programming and its applications. Research activities include both algorithmic development and model development with interaction between them, so as to reduce the gap which normally separates these two areas in practice. The topics of research include:

1. Efficient methods for large-scale "staircase" and other specially structured linear programs. Of particular interest are the continuous simplex approach, methods based on decomposition, and factorizations designed for staircase matrices.
2. Several refinements and extensions to the MINOS code for large-scale linearly constrained optimization.
3. Analysis of several topics in nonlinearly constrained optimization.
4. A study of Newton-type, sparse quasi Newton. and conjugate gradient methods for large, sparse, unconstrained optimization. It is anticipated that the results in the unconstrained case should suggest useful approaches to the large sub-problems that arise in constrained optimization.
5. Methods for large, sparse linear least squares, structured quadratic programs, and constrained linear and nonlinear least squares.
6. Evaluation and comparison of algorithms for various classes of optimization problems.

This work will also include an effort to collect sets of test problems that are non-trivial and representative.

University of California - Los Angeles; N. Donald Ylvisaker; *Statistics and Game Theory (Mathematical Sciences and Computer Research*; (DMS-8301587 A01 & A02); $21,000. (Joint support with the Statistics and Probability Program - Total Grant $64,100).

Research is to be carried on in the general areas of statistics and game theory. Professor Jennrich plans to work on a number of problems related to his statistical computing interests. He intends to develop an algorithm for both exploratory and confirmatory factor analysis and to attack the problem of providing standard errors for maximum likelihood estimates in this context. Regression algorithms are to be investigated as well. In the area of survival analysis he will investigate the behavior of tests for the comparison of survival curves under unequal censoring. Ylvisaker's research will be concerned with model-robust design of experiments. Designs and estimates will be studied for the estimation of regression parameters in models which are infinite dimensional approximations to standard finite dimensionl models. Corresponding work is planned for the robust modification of designs and estimates in combinatorial design settings. Professor Ferguson will work on stochastic games with information structure in which players have partial information about which game is being played. This area is a relatively unexplored one and he will consider questions of existence of value, form of optimal strategies, and the connections of this work with Bayesian statistics.

Yale University; John A. Hartigan; *Clustering Algorithms (Computer Research and Mathematical Science)*; (DCR-8401636); $40,400; 12 mos. (Joint support with the Statistics & Probability Program and the Division of Industrial Science & Technology - Total Grant $81,247).

The main thrust of the research is the development, evaluation, and distribution of clustering algorithms. New algorithms for block clustering and for reconstructing evolution from molecular data will be developed. Statistical evaluation takes data to be a random sample from a population, and considers how well various clustering techniques discover clusters in the population. The dip test of multimodality in one dimension will be studied further to establish its power against various alternatives. Extensions to many dimensions based on the minimum spanning tree will be explored. The asymptotic behavior of the minimum spanning tree (connected to single linkage clustering and percolation processes) will be studied. The block clustering algorithm BMDQ3M needs statistical tests to decide when the small blocks may be ignored.

Massachusetts Institute of Technology; Virginia C. Klema; *Software Tools for Computation With IEEE Floating Point Arithmetic*; (DCR-8400246); $91,232; 12 mos.

The recently approved IEEE binary standard for floating point arithmetic and its hardware implementa-

tion in microprocessors provide distinct enhancements for scientific computation. This research is directed toward the design and implementation of 1) software tools to access the floating point arithmetic exceptions and 2) basic linear algebra modules that optimize use of the floating point register stack. These software tools are necessay to support the use of the binary standard for floating point arithmetic.

Pennsylvania State University; Jesse L. Barlow; *Probabilistic Error Analysis of Numerical Computations in Special Computer Arithmetics*; (DCR-8201065 A01); $48,022; 24 mos. (Joint support with the Computer Systems Design Program - Total Grant $76,822).

This project involves two separate, but related, problems. The first is the probabilistic error analysis of numerical algorithms. The second is the design of a parallel architecture for solving ordinary differential equations using digit online arithmetic. The error analysis of digit online arithmetic has already presented problems over and above those of standard floating point arithmetic. The fact that it is redundant arithmetic system and that its addition and subtraction operations do not automatically normalize is the cause of most of these problems. Thus, it is useful for special purpose devices only. It has the advantage, however, of allowing the hardware designed to incorporate additional parallelism in iterative and recursive algorithms. A particular application of this property that is of great importance is the solution of numerical initial value problems.

Symbolic and Algebraic Manipulation

University of Illinois - Chicago Circle; Vera Pless; *Computer Algorithms in Finite Groups and Combinatorial Structures (Mathematical Sciences and Computer Research)*; (DMS-8201311 A02); $12,000; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $51,000).

This research is in the areas of error-correcting codes, computer algorithms, combinatorics, development of a group-theoretic and combinatorial computer system.

Pless will continue some difficult open cases of primes which can divide the group orders of certain extremal codes. She will investigate the theoretical aspects of "contracted codes," the possibility of classifying doubly-even codes of minimum weight 8 or more, cyclic codes from a global point of view, and also the interconnections between projective geometries, designs, and self-dual codes. Leon will continue to develop algorithms for computing with groups and combinatorial structures, including graph isomorphism and automorphism algorithms and prepare portable standard versions of several of his algorithms. He will work on the classification of 28-dimensional Hadamard matrices. He will investigate questions relating to the complexity of group-theoretic and combinatorial algorithms.

Leon and Pless will continue to improve their computer system CAMAC (a unified system for algebraic and combinatorial computation). They are adding new capabilities such as the computation of the automorphism group of Hadamard matrices and will incorporate programs to test whether 2 Hadamards or error-correcting codes are equivalent. They will continue to distribute CAMAC to potential uses. In addition, they are developing a Pascal CAMAC which will be easily transportable to most (including mini) computers, uses structured programming and will handle graphs and designs as well as codes and groups.

New York University; Jacob Schwartz; *Conference on Computer Algebra as a Tool of Research in Mathematics and Physics, New York, New York; April 12-13, 1984 (Mathematical Sciences and Computer Research)*; (DMS-8406681); $2,000; 12 mos. (Joint support with the Mathematics Special Projects Program - Total Grant $6,000).

A conference will be held in the general area of applications of computer algebra in pure and applied mathematics as well as mathematical and theoretical physics. It will offer a forum for interaction between leading experts in the development and implementation of algebraic computional systems and techniques with those persons in the mathematics and physics communities.

The conference aims to expand the horizon of existing algebraic computation systems and their applications leading, it is hoped, to tools of mathematical research that are useful, accessible, and effective. The conference setting will encourage discussion and suggestions which will lead to strengthening of the mathematical power of new or evolving computer-algebra system.

Southern Methodist University; David Y. Yun; *Algorithms and Systems For Symbolic Algebraic Computation*; (DCR-8314600 A01 & A02); $92,366; 12 mos. (Joint support with the Division of Industrial Science & Technology - Total Grant $170,795).

This project involves the study of algorithms for symbolic and algebraic computation. The effort will embrace pure mathematical and algorithmic research as well as implementation and testing of algorithms with the goal of achieving efficiency and practicality. The research is to be conducted by B. David Saunders and S. K. Abdali, faculty members in the Department of Mathematical Sciences at Rensselaer Polytechnic Institute, Richard D. Jenks, and Barry M. Trager. of the Mathematical Sciences Department, IBM Research, Yorktown Heights, New York, and David Y. Y. Yun of the Department of Computer Science and Engineering, Southern Methodist University.

The objectives of the research program are to discover, analyze, implement, and test algorithms related to the automatic solution of equations in closed form.

The research program consists of four parts:

1. Investigation of algorithms for solving Diophantine equations over unique factorization domains,
2. The development of a subsystem in computer algerbra for linear algebraic computations over non-numerical fields,
3. Investigation of efficient algorithms for solving systems of polynomial equations, and
4. Study of Liouvillian theories and development of algorithms for solving linear ordinary differential equation.

University of Washington; Martin Tompa; *The Combinatorial Structure of Computations and Symbolic Manipulation (Computer Research)*; (DCR-8301212 A02); $13,804; 12 mos.(Joint support with the Theoretical Computer Science Program - Total Grant $27,609).

The success of theoretical computer science is due largely to the process of abstracting clean, tractable problems from the morass of realistic detail. Yet there is much left to be desired in the understanding of the structure of computations. For instance, it is not known how to exploit the highly parallel systems that will soon be real. Neither is it known how to exploit randomness (in the sense of probabilistic, or Monte Carlo, computation) to derive efficient algorithms. Even in the setting of ordinary sequential computation, where a reasonbly good repertoire of algorithms exists, there are few techniques to show that those algorithms are optimal.

Part of the problem is that not enough of the detail is abstracted away. General computational models like Turing machines or random access machines have sufficient detail to defy analysis.

In the project, computations are abstracted further in order to lay bare the underlying combinatorial structures. With all inessential details stripped away, it is often simpler to discover new efficient algorithms, and demonstrate their optimality (at least within this structured framework). The areas studied are synchronous parallel complexity, sequential time and space complexity, and probabilistic complexity.

University of Wisconsin - Madison; George E. Collins; *Algebraic Algorithms and Systems*; (DCR-8408107); $84,867; 12 mos.

This project involves the study of five related research areas in symbolic and algebraic manipulation. These topics are:

1. Polynomial factorization,
2. Cylindrical algebraic decomposition,
3. Polynomial residue class rings,
4. Linear diophantine equations, and
5. Development of the SAC-2 symbolic manipulation system.

Software Studies and Metrics

University of Arizona; Mike Szilagyi; *High Performance Electrostatic Lenses for Ion Beam Lithography*; (ECS-8317485); $7,120; 9 mos. (Joint support with the Solid-State and Microstructures Engineering Program - Total Grant $71,209).

The use of charged particle beams in the semiconductor area is widespread and encompasses ion implantation techniques as well as surface studies such as SIMS. For future VLSI circuits, focussed ion beams may be used for direct writing of patterns and focussed ion droplets can be used for direct writing of metallic interconnects and mask repair. Present ion beam optics suffer deficiencies which appear as abberations limiting the focussed spot size and ion collection efficiency. The proposed research addresses this problem by considering new ways to develop electrostatic lenses for application to ion beam lithography. Specifically the approach will be computer based and will utilize approximation for the electrostatic potential distribution; the original electrode configuration can then be restructed from this. The approximation and calculation will be handled using the artificial intelligence language LISP in conjunction with an algebraic language MACSYMA. Ultimately an EXPERT system will be developed which will have the full capability of ion optical synthesis and electrode designs.

Yale University; Elliot M. Soloway; *Mapping Between Programmers' Conceptualizations And Programming Language*

Constructs (Computer Research and Information Science); (DCR-8302382 A01); $12,048. (Joint support with the Division of Information Science and Technology - Total Grant $18,048).

Expert programmers have and use rich mental conceptualizations that capture information beyond the syntax and semantics of programming languages, such as when to use a construct, how to integrate one construct with another. The hypothesis to be explored is: "If a programming language construct mirrors a programmer's conceptualization of the task, then that construct can be used more effectively than one that does not." Thus, the mappings between programmers' mental models and programming language constructs, for 3 basic programming notions: variables, conditionality, and looping will be examined.

Three empirical techniques will be employed:

1. In construct comparison studies a construct that matches a programmer's mental model better than an existing language construct will be created; it is conjectured that a group using the new construct should out perform (in accuracy, speed) a group using the existing construct. Tasks will be program generation, modification, and debugging.
2. Bug patterns and frequencies, and
3. Video-taped interviews will be used as corroborating evidence.

By examining how programmers comprehend specific problem solving and programming notions, and how they realize these notions in programming languages, it should provide important information to the designers of programming languages, environments, and methodologies.

University of Maryland - College Park; Satish K. Tripathi and Ashok K. Agrawala; *Modelling of Parallel Software*; (DCR-8405235); $66,730; 24 mos. (Joint support with the Computer Systems Design Program - Total Grant $116,730).

The objective of this research is to model concurrent software to determine optimal algorithm partitions, evaluate programming languages concurrency constructs, model computer networks applications layers, and identify software bottlenecks. In prior work models were built specifically for a particular program, and thus could not be applied to other concurrent programs. Or they handled only simple interprocess communication/sychronization.

An initial model of parallel software in which each constituent process executes asynchronously on a dedicated uniprogrammed processor and periodically exchanges information with another process under an assumption of exponential state occupancy times has been designed. The code of each process is abstracted mechanically into a Markov chain consisting of four types of states: computing, waiting, making information available to another process, and obtaining information from another process.

The research plan consists of;

1. further developing this initial model,
2. exploring alternatives model:
3. relaxing the assumptions, and
4. model validation.

Michigan Technological University; Karl J. Ottenstein; *Program Translation For Multiprocessors*; (DCR-8404463); $41,001; 24 mos.

"Supercomputers" based on multiprocessors offer the fastest means known for solving many problems of importance to science and society. Converting existing large-scale programs for execution on multiprocessors is a very machine-dependent task that would be time consuming and expensive if done manually. Automated conversion would not only lower cost, but also provide the speed benefits of multiprocessors to the scientific community in a shorter time frame. This research will study this conversion task, particularly the detection of implicit parallelism in programs. New transformations will be developed to improve program performance on particular multiprocessors. Metrics on dependence patterns in programs will be defined and data gathered in order to gain a better understanding of the nature of programs and the complexity bounds of a variety of program analysis and transformation algorithms.

Michigan Technological University; Karl J. Ottenstein; *Program Translation for Parallel Systems (Computer Research)*; (DCR-8203487 A01); $8,339.

Computers are being designed and built which permit a high degree of concurrency. Some of these MIMD (multiple-instruction, multiple-data stream) architectures are control-driven in the conventional sense while others are data-driven—that is, instructions execute according to the availability of data. This project will investigate techniques for translating "conventional" languages (such as FORTRAN), having no special extensions for the expression of parallelism, into high-level or machine languages having those features. Average case complexity bounds for the designed transformations will be constructed based on data to be gathered and analyzed for data-dependency metrics.

University of Colorado - Boulder; Leon Osterweil; *Evaluation and Continued Developments of Software Tools and Environments*; (DCR-8403341); $150,000; 12 mos.

This research centers on the continued development and evaluation of tools and environments to support the production of software. The research will build upon work done in the past few years on the Toolpack project. This project has resulted in the creation of a set of tools and integration software capable of providing considerable support for software development, maintenance, and testing. The architecture of the tools and integration software has been carefully designed to accommodate additions and changes to the toolset and command language. Thus this software seems to provide an excellent testbed for experimentation with tools, as well as substantive experimentation with the flexibility and extensibility of the Toolpack architecture itself.

University of Massachusetts - Amherst; Lori A. Clarke, Debra J. Richardson, and Steven Zeil ; *The Development of a Scientific Testing Method*; (DCR-8404217); $109,174; 12 mos.

There are a number of research projects addressing the problem of software testing. For the most part, recent research has moved for developing tools that gather information about programs to developing techniques that actually apply this information. There has been little work, however, on understanding the comparative strengths and weaknesses of existing techniques or their potential interactions. In this project several testing techniques will be evaluated to gain a better understanding of their strengths, weaknesses, and overlap. This evaluation should lead to insights on how to extend and improve some of these techniques and how to integrate a set of these techniques into a powerful testing method. The goal is to assure that the error detection capabilities of this integrated method are well-understood and clearly stated and so we refer to this as scientific testing.

University of Massachusetts - Amherst; Krithivasan Ramamritham; *Synthesis of Resource Controllers for Distributed Systems*; (DCR-8403097); $35,043; 12 mos.

Devoting a controller process to manage each shared resource in a distributed system limits the types of interactions between processes and paves the way for the modular design of distributed systems. Such a resource controller accepts access requests from user processes and services them based on criteria such as the invariant properties of the resource and the required priority and fairness in service.

Efforts from this group so far have led to (1) a language for specifying the properties of resources and their controllers, (2) schemes for verifying extant resource controllers - both when control is exercised in a centralized manner as well as in a distributed manner, and (3) the design of a system for automatically synthesizing code for resource controllers given the specifications of the resources and their controllers.

Investigations will continue along related directions. The first involves the implementation of a prototype synthesis system so as to study the practical issues involved in synthesizing resource controllers and to build the heuristic support required to synthesize efficient controllers. The present synthesis algorithm is geared to handle situations in which a user's view of a resource can be realized via a single physical resource and its controller. The second goal of this research will be to overcome this limitation. Towards this end, the synthesis of concrete resources controlled by a number of interacting controllers is being investigated in order to realize a given abstract resource accessed via operations defined on that resource.

University of Massachusetts - Amherst; Edward Riseman; *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776); $25,000; 12 mos. (Joint support with the Special Projects Program and the Intelligence Systems Program - Total Grant $160,000).

This grant will support research activity within the Computer and Information Science (COINS) Department in six areas of experimental computer research: computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under twenty-one other grants and contracts, including ten from NSF, this grant will provide partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support will be used for administration and operation of the COINS research facility, with the funding going towards technical support staff and equipment maintenance. Over the next three years, this grant will cover approximately 40%, 25%, and 15%, respectively, of the total costs.

University of Arizona; Christopher W. Fraser; *Faster and More Thorough Object Code Optimizations*; (DCR-8320257); $136,503; 24 mos.

Peephole optimization improves inefficient patterns in object code. Many optimizing compilers include peephole optimization because it corrects inefficiencies hard to correct in earlier compilation phases. Classical peephole optimizers require hand written, machine-specific patterns and are hard to make complete. Recent machine-directed peephole optimizers avoid these deficiencies but are slower and less flexible. These optimizers are still developing.

Research is proposed to advance machine-directed optimizers. They will be made faster by automatically generating optimizations at compile time. They will be made more general by a new replacement strategy that should eliminate narrow peepholes, and by including generalizations that subsume traditional peephole optimizations not included earlier. Increased regularization and formalization of a notoriously idiosyncratic class of optimizations should result. Experiments will be run to evaluate the speed and power of these optimizations.

University of Arizona; Ralph E. Griswold; *High-Level Programming Language Facilities for Data Structure Processing*; (DCR-8401831); $73,058; 12 mos.

The focus of this investigation is the study of high-level programming language facilities for processing nonnumerical data, such as strings and lists.

Previous work has concentrated on two efforts: (1) the design of new facilities for characterizing and processing nonnumerical data, and (2) the implementation of these facilities so that they can be evaluated and made available to the computing community. There is a close interaction between these two efforts: the value of new facilities often can be determine only by their application to real problems. Conversely, the use of new facilities frequently suggests new possibilities and neglected areas of application.

This research will build on past results and focus on two areas of investigation:

1. the development of a laboratory for the study and design of high-level pattern-matching facilities for strings, structures, and data bases, and
2. the design of programming language facilities for manipulating sequences as data objects.

University of Arizona; David R. Hanson; *Programming Language Toolbox*; (DCR-8302398 A01); $48,017; 12 mos.

Despite the increasingly large number of diverse programming language that have been introduced or otherwise described in recent years, most language fall into one of two classes: the "Pascal-like" language, such as Pascal and Ada, and the "very high-level" languages, such as APL, SETL, and some of the applicative languages. Much of the work on the Pascal-like language is directed toward refinement and perfection of known languages, per se, as is research in very high-level languages.

Research in very high-level languages require extensive experimentation. The enormous amount of work required for this experimentation is one reason there are so few well established such languages (e.g. APL, SNOBOL4). Compiler generation tools, which have greatly simplified the implementation of Pascal-like languages, are much less useful for languages with rich run-time semantics or for experimental languages whose semantics undergo constant change.

This research is directed toward simplifying the design and implementation of very high-level languages via a programming language toolbox. The emphasis is on linguistic mechanisms with late, and even adjustable, binding times. The major goal is to promote futher experimental research in very high-level programming languages by providing a means of constructing prototype implementations from a toolbox of "semantic components". The major expected result is a toolbox from which experimental languages can be constructed quickly and modified easily.

University of California - Irvine; Peter A. Freeman; *Reusable Software Engineering (Computer Research)*; (DCR-8304439 A01); $113,143; 12 mos.

The goal of this project is the development of effective reusable software engineering (defined as activities that produce desired software engineering workproducts and information that can be reused later in producing a new workproduct). This work is a continuation and expansion of current work and falls into two categories of experimentally based work: Software Construction Using Components and Advanced Models and Tools. The first category is a continuation of work on the Draco approach to the use of software components that utilizes program transformations and domain-specific languages. The second category addresses the key problem of representation of system development information and is exploring the reuse of graphical models of system design information.

University of Southern California; Robert M. Balzer and David Wile; *Transformation-Based Maintenance (Computer Re-*

search); (DCR-8304190 A01); $111,500; 12 mos. (Joint support with the Software Systems Science Program - Total Grant $223,000).

For several years this project has been creating the technology needed to support an alternative software development paradigm. This paradigm is based on capturing the entire program development process—including specification, design, implementation and maintenance—within the computer and suppporting it with automated tools. The investigators believe this paradigm will ultimately result in dramatic improvements in both software productivity and reliability because the human developer's role will be reduced to decision making and guidance, while the computer's will be expanded to manipulation, analysis and documentation. To this end, the primary research focus has been on recording, structuring, and formalizing the program development process, making it amenable to machine aid and manipulation.

The investigators expect that the leverage provided by this paradigm will derive more from its support of the maintenance process than from its support of the initial implementation, wherein the increased formalization of the development process requires more involvement by the human developer. On the other hand, maintenance involves small changes. Increased formalization of these changes is insignificant if the already formalized development can be adapted and reused. Therefore, the project will focus on supporting such reimplementations by strengthening the formal representation of developments, by supporting the evolution of specifications through the same mechanisms that support maintenance, and by developing a formal language of modification for both specifications and developments.

University of Connecticut; Taylor L. Booth; *Design Methods and Design Tools to Aid the Development of High Performance Software (Computer Research)*; (DCR-8303326 A01); $57,227; 12 mos.

The research is directed toward the development of tools that can be used by software engineers to include performance considerations as an integral part of the software design process. As a step in this direction, a software system, PASS (Performance Analysis Software System), has been designed and implemented. During the coming year, the system will be expanded and experimentation with the system will be continued.

Brandeis University; Jacques Cohen; *Software Tools for the Analysis of Programs*; (DCR-8317892); $119,931; 24 mos.

This project involves the study of algorithms and the corresponding software to estimate the performance of computer programs. Research will be conducted in the following areas:

1. Studying specific algorithms,
2. Extending the analysis to parallel and real-time programs,
3. Determining the probabilities with which conditions are evaluated as true or false, and
4. Using program analysis as a guide to program optimization.

Massachusetts Institute of Technology; Gerald J. Sussman, Charles Rich, and Richard C. Waters; *Abstraction, Inspection and Debugging in Programming*; (DCR-8117633 A02); $148,360; 12 mos.

Software Engineering has much to learn from other mature engineering disciplines because the problem solving behavior of engineers in different disciplines has many similarities. Three key ideas in current artificial intelligence theories of engineering problem solving are:

1. Abstraction - using a simplified view of the problem to guide the problem solving process. An example of using abstraction in programming is to first implement the desired behavior of a program on typical data before worrying about exception handling.
2. Inspection - problem solving by recognizing the form ("plan") of a solution. In programming, a plan may be a particular control strategy with unspecified primitive actions or an abstract data structure with an unspecified implementation.
3. Debugging - incremental modification of an almost satisfactory solution to a more satisfactory one. In programming, this evolution is triggered as often by a change in the problem specification as by an error in the initial solution. These three techniques are typically used together in a paradigm called AID (for Abstracting, Inspection, and Debugging). First an abstract model of the problem is constructed in which some important details are intentionally omitted. In this simplified view inspection methods are more likely to succeed, yielding the initial form of the solution. Further details of the problem are then added one at a time with corresponding incremental modifications to the solution. This project will investigate this model of program construction.

New Mexico Institute of Mining & Technology; Allan Stavely; *Analysis Tools for Design-Level Assessment of Software Systems*; (DCR-8315028); $73,286; 24 mos.

The goal of the project is to produce software analysis tools which help the designer to see the implications of design decisions and to assess the acceptability of a

proposed design. A set of prototype tools for analyzing designs of software systems has been constructed. These include the sets of states that the system can reach and the posssible sequences of significant events, such as state changes and message transfers, that can occur during the system's execution. The capabilities of the design language and software tools to provide support for the entire process of software design, from the first consideration of a software system's internal organization to the classical stepwise refinement of individual components will be extended. The project will include the continued study of techniques for identifying greatly simplified abstractions of a complex system which contain sufficient information to perform a given analysis, so that it will be computationally feasible to analyze large, real-world designs.

Columbia University; Cyrus Derman; *On Debugging Procedures (Mathematical Sciences and Computer Research)*; (DMS-8405413); $20,000; 24 mos. (Joint support with Statistics and Probability Program - Total Grant $70,400).

The proposed research relates to an increasingly important area, computer software reliability, which has yet to attract substantial attention from probabilists and statisticians. The study of fundamental and theoretical properties of 'debugging' procedures is potentially of great value.

Cornell University; Robert L. Constable; *Experiments with a Program Refinement System*; (DCR-8303327); $67,308; 12 mos. (Joint support with the Software Systems Science Program and the Intelligent Systems Program - Total Grant $199,308).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project will explore a logic for formal top-down development called refinement logic. During the previous grant period a system, called Program Refinement Logic: PRL, was designed and implemented. During the coming grant period this system will be improved and expanded and considerable experimentation will be conducted.

Cornell University; Ray Teitelbaum: *Syntax-Directed Programming Environments (Computer Research)*; (DCR-8202677 A02); $133,262; 12 mos.

The goal of this research is the design and implementation of highly interactive environments for computer programming that will significantly improve programmer productivity. The scope of the research includes instructional as well as professional programming systems, end-user environments as well as software tools that facilitate the creation of such systems, algorithm design as well as system implementation.

The research includes the design, implementation, and distribution of:

1. The Synthesizer Generator, a language-independent constructor of syntax-directed editors able to enforce context-sensitive, as well as context-free, structural constraints;
2. The Semantics Processor, a language independent constructor of syntax-directed environments providing execution, source-level diagnostic interpretation and debugging, flow analysis, and program anomaly detection
3. A Transformation Specification System for constructing environments providing collections of program transformations for program refinement, optimization, modification, and maintenance.

Cornell University; Kenneth G. Wilson and David J. Gries; *The Gibbs Project*; (DCR-8312319); $109,071, 12 mos.

This program of research and development is aimed at providing the scientist—initially the physicist—with a programming environment that is at least an order of magnitude more effective than existing ones. The research will take advantage of work in a number of computer science areas, such as programming environments, language theory, program transformations, and desk-top workstations with sophisticated graphics facilities. The direct collaboration of physicist and computer scientists is a key element of the project.

University of Oregon - Eugene; Stephen Fickas; *The Mechanization and Documentation of Software Specification*; (DCR-8312578); $65,000; 24 mos.

This research project is concerned with the mechanization of the software specification process. The current non-mechanized construction of formal specifications present problems of both completeness and consistency. The former is rooted in the large amount of low level common sense knowledge that must be made part of a complete specification; the latter in the complex side-effects that a change in one part of a specification can have on the rest. Further, the problem solving steps involved in specification construction are not documented. Because the specification process is done outside of the computer (except for the lowest level editing steps), integration of the specification process with other software development processes (design, coding, maintenance) is difficult. The goal of this research is to address these problems.

The research will build on two prior efforts in the area of mechanization of software development:

1. The formal specification language Gist, and
2. Glitter, a system for automating the transformational development of software.

The general approach centers on two key ideas:

1. The effort expended in constructing a complete, consistent formal specification can be at least partially reused by maintaining a library of specification schema from which to build, and
2. The specification construction steps can be mechanized, and hence automated and documented.

William Marsh Rice University; Kenneth W. Kennedy and Robert T. Hood; *The Computation and Use of Interprocedural Information in a Programming Environment*; (DCR-8303638 A01); $87,510; 12 mos.

A programming environment provides an intergrated collection of tools to aid programming teams in entering and testing their programs, along with sophisticated facilities for managing all the codes in the various program components. An important part of such an environment is the project data base, in which is recorded all the information that the environment knows about the programs and modules that comprise the project.

Once the concept of such a data base is accepted, many applications for the centralized store of information present themselves. For example, the system can record and use information about the variables each module uses or defines. A sophisticated optimizing compiler can then use this information to produce code that is substantially more efficient than it would be if nothing were known about the behavior of separately compiled subroutines. There are many other interesting, and sometimes surprising, applications.

This investigation is concerned with ways to creatively, effectively, and efficiently use a centralized project data base to pass information among the various tools of a programming environment.

Virginia Polytechnical Institute and State University; H. R. Hartson and R. W. Ehrlich; *The Structure of Human-Computer Interaction (Information Science and Computer Research)*; (IST-8310414); $29,202; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $116,809).

The goal of this research is to develop a theory of the dialog which takes place between an information system and its user and to embody that theory in a Dialog Management System, which will serve as a test bed for the theory's validation. The theory will be based on a model which the investigator will develop to provide a framework for discussion of all aspects of a user-system dialog without regard to the specific context of that dialog or to the specific technologies by which it might be implemented.

Software Systems Science

Programming Languages: Design, Semantics, Implementation

University of Southern California; Edward K. Blum; *Algebraic Approaches to Programming Concepts and Their Semantics*; (DCR-8406920); $16,218; 12 mos.

The main objective of this project is to analyze and develop the relations connecting operational semantics to algebraic semantics of programming concepts and programming languages. This entails a study of the applicability of algebraic and more general axiomatic methods to the semantics of programming languages and systems. One area of emphasis will be the implementation of data types by an algebraic method recently initiated by the investigators. This method will be developed further and its practicality will be tested on the I.E.E.E. floating-point type and the proposed Fortran array type, both of which provide real test cases for studying the relation of algebraic to operational semantics. A second area of emphasis is the proper mathematical framework for recursive derived operations which arise in implementations of data types and general computations over algebras. A third area of emphasis is the algebraic and operational semantics of concurrent systems.

Yale University; Joseph A. Fisher; *The Extension and Application of Global Compaction Techniques for Horizontal Scientific Code (Computer Research)*; (DCR-8308988 A01); $87,860; 10 mos.

This project is to complete the designing and writing of a compiler for the ELI-512, a very parallel, scientific attached processor with a 500 + bit instruction word. It will execute ten to thirty RISC-level operations per cycle.

One would not consider designing the ELI-512 without the development of Trace Scheduling. Trace Scheduling finds large amounts of global parallelism for statically scheduled, fine-grained architectures. Because the task of scheduling thousands of unrelated operations is not humanly possible, Very Long Instruction Word (VLIW) architectures cannot be effectively hand-coded.

During the first one and a half years of this project, much work has been done towards eliminating all the conceptual difficulties that once prevented the use of VLIWs for parallel computation. A Trace Scheduling compiler now works; it generates code for an idealized ELI-like machine. The compiler, called BULLDOG, has demonstrated that one can generate very parallel code from a large range of ordinary programs.

This project covers the completion of work on the compiler. Work remains on the ELI-512 generator, the FPS-164 code generator, various front ends, the array anti-aliasing system, and other tools and utilities.

Yale University; Paul Hudak; *Distributed Combinator Reduction (Computer Research)*; (DCR-8302018 A01); $21,042.

Combinator reduction provides a simple execution environment for functional programming languages. Work in this area will be extended to distributed processing systems, by developing strategies for effectively distributing the combinator reduction process. By using combinators as a basis, the need for an environment to support bound variables is obviated, and we are left with a powerful yet simple graph reduction scheme amenable to analysis and manipulation. A constant effort of this research will be the development of general principles of distribution, that will allow combinator reduction to occur naturally and automatically, and be applicable to a variety of systems. The research is divided into two phases; the first to develop tools to observe distributed combinator behavior, the second to use these tools to develop effective distribution strategies.

University of North Carolina - Chapel Hill; Richard T. Snodgrass; *The Formal Specification and Implementation of a Temporal Query Language*; (DCR-8402339); $67,774; 12 mos.

The goal of this project is two-fold: to develop a formal semantics for a query language for relational databases containing historical information, and to implement a historical database management system utilizing this language. The tuple calculus semantics, already defined for a subset of the language will be extended to handle aggregate functions and incomplete information. The implementation will be based on a

temporal relational algebra that supports the incremental updating of historical relations. This algebra will be formalized and proved to preserve the semantics defined by the tuple calculus. Optimization strategies for historical databases will also be investigated.

Columbia University; Rodney Farrow; *A Software Laboratory for Research in Attribute Grammars*; (DCR-8310930); $113,999; 24 mos.

Attribute grammars are a formalism for specifying the semantics of context-free languages. Translator-writing systems based on attribute grammars have been maturing for the last decade. Some techniques of these systems are effective enough to be used in commercial compilers. However, few comparisons of the many competing techniques have been done. This project will establish a collection of attribute grammars and software tools for building, analyzing, and transforming them. Quantitative experiments will be performed to evaluate the relative effectiveness and applicability of these techniques.

Syracuse University; John C. Reynolds and Nancy J. McCracken; *The Design, Definition, and Implementation of Programming Languages*; (DCR-8017577 A04); $88,530; 12 mos.

Research will be carried out on the design of programming languages and the mathematical definition of their semantics, with emphasis on the design of Algol-based languages and the use of category theory as a definitional tool. Specific topics include: the application of category-sorted algebras to treat implicit conversions and generic operators; the use of a non-functional Cartesian closed category to provide an abstract model of the store; extension and modelling of a language in which interference can be determined syntactically, and development of a variant of specification logic for proving programs in such a language; and models of languages with type binding.

Operating Systems and Concurrency

University of California - Irvine; Rami Razouk; *Modeling and Analysis of Communications Protocols Using Timed Petri Nets*; (DCR-8406756); $129,970; 24 mos.

Communications protocols exhibit characteristics which distinguish them from other software systems: They require close interaction between hardware and software; they are truly concurrent since they support communication between peers which reside on two or more processors; their behavior is timing sensitive, and they are designed to tolerate failures in underlying communication services. These characteristics lead to difficulties in specifying and verifying protocols.

This research project seeks to develop modeling techniques which can describe concurrent, time-sensitive software/hardware systems, and analysis techniques which yield results relating to performance and correctness. In particular, this work builds on current research using Timed Petri Net models and seeks to apply the new techniques (supported by tools) to complex communications protocols such as DOD's Transmission Control Protocol (TCP).

University of Colorado - Boulder; Michael G. Main; *Semantics of Nondeterminism and Nondeterministic Recursive Equations*; (DCR-8402341); $60,671; 24 mos.

About 15 years ago the order-theoretic semantics was introduced with the goal of providing a systematic method for solving general recursive equations occurring in the formal study of programming languages. Shortly after this, a number of people proposed methods to extend the order-semantics to nondeterministic recursive equations which occur in nondeterministic and concurrent programming languages. This project will continue the study of nondeterministic recursive equations, focusing on the case of linear nondeterminism. Linear nondeterminism comes from the idea of call-time-choice, proposed by Hennessy and Ashcroft. The key idea is that any nondeterminism in a function's argument is decided at the time of a function call. This is in contrast to the run-time-choice of many studies of nondeterminism. Three particular directions of study are to be investigated. First, a unification of the three classical powerdomain constructors, using linear algebra. This should provide order-theoretic techniques for solving recursive equations in semiring modules. Second, a study of the Arbib and Manes pattern-of-calls approach to solving equations, concentrating on linear nondeterminism. Third, solving nondeterministic domain equations and giving models of nondeterministic calculus with call-time-choice.

University of North Carolina - Chapel Hill; J. Dean Brock; *Scenarios: A Formal Model of Concurrent Computation*; (DCR-8406850); $88,514; 24 mos.

The scenario model of computation is a semantic theory for non-determinate systems in which processes and networks of processes are represented by an input-output relation augmented with a causality ordering placed on individual communication events. In addition to having a relatively simple characterization of

non-determinate behavior, scenario theory also has straightforward rules for deriving the characterization of a network of processes from the characterizations of its components.

Presently, the application of scenario theory is restricted to programming languages implementable with static dataflow graphs communicating with streams (simple sequences) of values. This project will extend the theory to dynamic dataflow graphs which communicate with a richer set of values, such as streams of streams (multi-level sequences) or "lazy" streams (streams with "holes"). The relation of scenario theory to computational models more directly concerned with the specification, verification, or implementation of non-determinate systems will also be investigated.

University of Pittsburgh; John P. Kearns and Mary L. Soffa; *The Implementation of Concurrency (Computer Research)*; (DCR-8405006); $119,118; 24 mos.

The growing importance of concurrency as a natural programming paradigm, particularly in embedded real-time systems, renders the efficient implementation of concurrency imperative. There currently exists no reasonable rigorous methodology for expressing the demand placed upon an execution vehicle by a program which consists of a number of cooperating tasks. The development of such a methodology would seem to be a necessary prerequisite for the rational design of language implementations which support concurrency.

This project will develop the methodology mentioned above as the initial step of an in-depth study of the implementation of concurrency. This study would address performance evaluation techniques for implementations of concurrent control, the design of language primitives which express concurrency in a matter conducive to their efficient implementation, and the determination of the impact that the execution system has on the concurrency implementation.

Washington State University; David B. Benson; *Algebraic Reasoning About Concurrent Computer Programs*; (DCR-8402305); $177,397; 24 mos.

Many problems in science and engineering are solved using linear and multilinear algebra. One major problem area in computer science and software engineering is to understand and specify the behavior of computer programs. A variation of linear and multilinear algebra provides the basis for solution techniques. The variation is to replace the numbers used in traditional mathematics by Boolean and other symbolic quantities. The resulting mathematics is easy to use and explicates other forms of reasoning such as Dijkstra's weakest preconditions. These ideas will be extended to the questions of concurrent computation - indeed some progress has already been made - and to the related questions of processes.

Data Bases

Boston University; Nathan Goodman; *Recovery, Concurrency Control, and Component Integration in Database Systems (Computer Research)*; (DCR-8317888); $127,095; 12 mos.

To help systematize and simplify the engineering of database systems, this project will analyze formal models of database system recovery and concurrency control. Specifically, it will analyze: the correctness of recovery algorithms; distributed recovery problems caused by replication; the performance of concurrency control algorithms; models of multilevel concurrency control algorithms; concurrency control algorithms for dynamic search structures; and probabilistic concurrency control. Work will also begin on a new system project, Database Kit — a modular, reconfigurable database system.

University of Minnesota; David Hun Du; *Studies in Associative Retrieval (Computer Research)*; (DCR-8405498); $80,552; 24 mos.

If a query in an information retrieval system is allowed to specify conditions dealing with multiplicity of the keys, the query is considered as an associative retrieval.

On the file design problem for associative retrieval, this project will study new file structures under both static and dynamic environments. To improve the response time, a file can be stored on a multiple independently accessible disk system. The file allocation problem is to find a way to allocate all buckets in a file onto m ($m > 1$) disks such that the maximum disk access concurrency is achieved. In this research the file allocation problems for two types of Cartesian product files (Binary Cartesian product files and power-2 Cartesian product files) will be studied.

Two general models which serve as the theoretical bases for the file design and file allocation problems for associative retrieval have been considered. The solutions to these two models can be also applied to several related problems. It has been shown that both models

are computationally intractable. Therefore, the project will mainly concentrate on the heuristics (especially low order polynomial time heuristics) of these two models. The concurrency control algorithms for associate retrieval will also be studied.

State University of New York - Stony Brook ; David S. Warren and Edward Sciore; *An Integrated Prolog-Database System*; (DCR-8407688); $94,595; 12 mos.

There is a close and fundamental relationship between the logic programming language Prolog and relational database query languages. Previous Prolog development has concentrated on efficient execution of relatively small recursive programs running entirely in memory. Relational database systems, on the other hand, have developed efficient ways to handle very large amounts of data stored on disk; they also handle multiple concurrent users, allowing update of shared data. This project involves research into how a single integrated system can efficiently perform all these functions.

The research includes the development of a prototype Prolog system that efficiently executes both recursive programs and retrievals from large databases stored on disk. It also includes theoretical work on how to include concurrency and database update cleanly and logically in the Prolog framework. The prototype system will serve as a testbed for this theoretical work.

Case Western Reserve University; Z. Meral Ozsoyoglu; *A Statistical Database Management System: Design, Implementation and Experiments (Computer Research)*; (DCR-8306616 A01); $60,959; 12 mos.

A Statistical Database Management System (SDBMS) provides statistical data analysis capabilities as well as capabilities to store, model and manipulate data suitable for statistical operations. The statistical database (SDB) security problem is to limit the use of the SDB in order to prevent users from inferring protected information in the database. An SDBMS should also have SDB security enforcement capabilities.

Recently, the research team at Case Western Reserve University completed a preliminary design of a SDBMS, called a System for Statistical Databases (SSDB) with the following major features:

1. It has a semantic data model with tools to model data for statistical operations,
2. It has a screen oriented query language to manipulate summary and raw data, and
3. It is a general purpose SDBMS useful in a variety of applications and has the capability to enforce SDB security mechanisms.

This research project will study SDBMS's and consists of two parts. In the first part, the research team will study SDB problems within the framework of the design and implementation issues of SSDB. The second part of the project consists of experimenting with a test-bed implementation of SSDB in order to evaluate efficiency, effectiveness and usefulness, and carrying out a comparative performance evaluation of the proposed SDB security mechanisms in real-world applications.

University of Texas - Austin; Abraham Silberschatz and Donald Fussell; *New Applications and Techniques of Database Concurrency Control*; (DCR-8104017 A01); $155,511; 24 mos.

Research on controlling the concurrent interaction of transactions has led to a small number of basic mechanisms for assuring database consistency, while at the same time allowing all concurrent transactions to terminate. Past research has shown how to extend one of these mechanisms, locking, to allow for greater concurrency by using more information about the behavior of transactions than had been considered previously. Means have also been investigated for assuring deadlock-freedom whenever possible, and for removing deadlocks efficiently otherwise. This project will extend these results to allow even greater concurrency by incorporating new types of information into the concurrency control mechanism. The use of operations such as rollback in different concurrency control schemes to achieve serializability and deadlock-freedom without locking, and general means for correct database restructuring in addition to data modification will also be investigated. New basic mechanisms for controlling transaction interactions will be sought, and the impact of distributed environments on concurrency control schemes will be examined. The performance of all mechanisms developed will be evaluated in order to ascertain their applicability to real database systems.

University of Virginia; John L. Pfaltz; *Analysis of Multi-Attribute Retrieval using Indexed Descriptor Files (Computer Research)*; (DCR-8302654 A01); $67,938; 12 mos.

Indexed Descriptor Files are a promising method of general multi-attribute file access. They

1. offer superior multi-attribute retrieval performance,
2. are conceptually simple,
3. require low maintenance overhead, and
4. may be constructed in either a "top-down" or "bottom-up" fashion.

Their conceptual simplicity has permitted detailed, and accurate, derivations of the expected retrieval and storage costs.

The latter three points above indicate that this file organization may be ideal for distributed data base organizations. In particular,

1. Entire subsets of a data base, with the relevant portions of the index, may be effectively transmitted to other nodes of the network,
2. The descriptors, that are used to search subsets of a file, may also be used to implement concurrency locks with low overhead.

This project will refine existing work on Indexed Descriptor File performance; but with particular emphasis on their utility in distributed networks, and concurrency control in any environment.

University of Wisconsin - Madison; Michael J. Carey; *The Performance of Algorithms for Shared Relational Database Systems*; (DCR-8402818); $114,086; 24 mos.

This project will investigate the performance of various algorithms for relational database management. The initial phase of this research will focus on concurrency control algorithms, continuing simulation studies of the relative performance of a number of proposed algorithms. CCSIM, a simulator for investigating concurrency control algorithms for centralized database systems will be used to examine conversational transactions, concurrency control algorithms for index structures, performance improvements due to reduced consistency requirements, and main-memory database environments. CCSIM will also be extended to model distributed database systems so that a representative subset of the many distributed concurrency control algorithm proposals can be studied.

The second phase of this research will involve experimentally studying the performance of several types of algorithms for distributed databases. A testbed system will be constructed for investigating the performance of algorithms for relational query processing, concurrency control, and recovery in a high-speed local computer network environment. This testbed will be based on porting WiSS, a low-level data management system, to the Crystal multicomputer system. Query processing issues to be examined include tradeoffs between static and dynamic planning, the performance of existing algorithms, and dynamic query processing algorithms which utilize load balancing techniques. Concurrency control and recovery issues to be examined include performance tradeoffs between centralized and decentralized control and the cost of maintaining multiple copies of data.

Programming Methodology

SRI International; Richard J. Waldinger; *The Automatic Synthesis of Computer Programs*; (DCR-8105565 A02); $98,442; 12 mos.

Program synthesis is the automatic derivation of a computer program to meet a given specification. The specification expresses the purpose of the desired program, but does not need to give any indication of the method by which that purpose is to be achieved. The ultimate goal of program synthesis is to automate the programming process. Interim goals include conceptual and mechanical aids for program derivation and improved methods for program verification, program transformation, and the semantic description of programming languages.

A deductive approach to program synthesis has been developed in which to construct a program is regarded as a problem of finding a proof of a mathematical theorem. This approach has been developed for the synthesis of "applicative" programs—LISP-like programs that return an output value but produce no side effects. The focus of the present research is to extend the deductive approach to the synthesis of nonapplicative programs, which may produce side effects as well as return output values. For this purpose, the deductive system in which proofs are conducted is extended to employ a situational calculus, which allows explicit mention of the states of a computation, so that side effects—changes of state—may be easily described. The specification of the desired program and the description of the constructs of the target programming language can then be expressed in terms of the operators of situational-calculus, and the program can be extracted from a proof of the appropriate situational-calculus theorem. The approach is amenable to machine implementation, and the same system can be applied to the verification and transformation of computer programs, as well as to their synthesis.

Stanford University; Donald E. Knuth; *Analysis of Algorithms (Computer Research)*; (DCR-8300984 A01); $239,410; 12 mos.

This research is oriented towards an improved quantitative understanding of the fundamental algorithms of computer science, and it comprises five major activities:

1. Development of new, efficient methods for use on computers.
2. Analysis of the performance of important computer techniques.
3. Analysis of the computational complexity of problems.

4. Studies of programming languages and software methodology.
5. Combinatorial theory to support the above activities.

Stanford University; Zohar Manna; *Interactive Program Synthesis (Computer Research)*; (DCR-8214523 A01); $91,248; 12 mos.

The goal of research on program synthesis is to allow the problem-solving ability of the computer itself to assist in the programming task. A program synthesis system accepts from its user a specification which gives no hint of a method by which that purpose can be achieved; the design of the program is the responsibility of the system.

The special thrust of this project is the design and implementation of an experimental interactive system, which would carry out the routine details of the program's derivation, but be guided in its major decisions by a human collaborator. The techniques to be used are based on the "deductive approach" of Manna and Waldinger, in which the problem of constructing a program is regarded as a task of proving a mathematical theorem.

University of Southern California; Robert M. Balzer and David Wile; *Transformation-Based Maintenance (Computer Research)*; (DCR-8304190); $111,500; 12 mos. (Joint support with the Software Engineering Program - Total Grant $223,000).

For several years this project has been creating the technology needed to support an alternative software development paradigm. This paradigm is based on capturing the entire program development process—including specification, design, implementation and maintenance—within the computer and suppporting it with automated tools. The investigators believe this paradigm will ultimately result in dramatic improvements in both software productivity and reliability because the human developer's role will be reduced to decision making and guidance, while the computer's will be expanded to manipulation, analysis and documentation. To this end, the primary research focus has been on recording, structuring, and formalizing the program development process, making it amenable to machine aid and manipulation.

The investigators expect that the leverage provided by this paradigm will derive more from its support of the maintenance process than from its support of the initial implementation, wherein the increased formalization of the development process requires more involvement by the human developer. On the other hand, maintenance involves small changes. Increased formalization of these changes is insignificant if the already formalized development can be adapted and reused. Therefore, the project will focus on supporting such reimplementations by strengthening the formal representation of developments, by supporting the evolution of specifications through the same mechanisms that support maintenance, and by developing a formal language of modification for both specifications and developments.

University of Illinois - Urbana; Nachum Dershowitz and David A. Plaisted; *Term Rewriting Systems and Computer-Aided Deduction (Computer Research)*; (DCR-8307755); $156,488; 24 mos.

A rewrite system is a non-deterministic pattern-directed program that takes a term as input and returns a term as output. Rewrite rules have long been used for ad-hoc computation in symbol manipulation systems, for simplifying in theorem provers and in conjunction with abstract datatypes. There are four aspects to verifying the correctness of rewrite systems:

1. Soundness - that terms are only rewritten to equal terms,
2. Termination - that there are no infinite derivations
3. Output correctness - that the output term satisfies given requirements,
4. Confluence - that for each input term there is at most one output term.

A rewrite system is canonical for an equational theory if it is sound, terminating and confluent. All four properties are in general undecidable.

This project will study the verification of rewrite programs, including partial correctness, termination and confluence. General methods will be developed for combining termination orderings, for proving termination of systems containing associative-commutative operators, and for proving termination of the completion procedure. REVE and OBJ are specific program systems which can be used to implement rewrite systems. This project also intends to enhance REVE and OBJ with incorporation of a decision procedure for real polynomials and associative-commutative unification.

Indiana University - Bloomington; David S. Wise and Steven Johnson; *Methods and Architectures for Applicative Programming*; (DCR-8405241); $81,937; 12 mos.

This project will advance the study of applicative languages as a basis for general purpose programming. This research has two main aspects: the formulation of constructs and methods required for robust programming endeavors, and the construction of hardware that supports the tenets of programming that we espouse. In previous investigations a small number of language constructs have been isolated that serve to specify a surprisingly broad class of applications, including those that commonly arise in systems programming. "Data recursion" and indeterminate constructor have

been added to the applicative vocabulary. With data recursions cyclic behavior can be represented by infinite data structures. The indeterminate constructor is used to address concurrency and real-time behavior. These constructs are related. Each enforces the programmer's view of activity as an attribute of data and thus its treatment in spatial terms.

This research has focused on the individual programmer, and yields methods that are claimed to improve productivity. The crucial question is whether these methods work "in the large", for example, in applications involving many participants. To explore such issues an existing, purely applicative programming language will be enhanced into a programming system.

In its present form the language implements "suspending construction", an operational model of computation that was proposed in the mid 1970s. This model can serve as a basis for a multiprocessing host that supports efficient execution of applicative programs. A sequence of prototype architectures will explore this model.

Cornell University; Robert L. Constable; *Experiments With A Program Refinement System*; (DCR-8303327); $66,000; 12 mos. (Joint support with the Software Engineering Program and the Intelligent Systems Program - Total Grant $199,308).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project will explore a logic for formal top-down development called refinement logic. In an earlier project a system, called Program Refinement Logic: PRL, was designed and implemented. In this project, PRL will be improved and expanded and considerable experimentation will be conducted.

Cornell University; David Gries and Fred B. Schneider; *Programming Methodology (Computer Research)*; (DCR-8320274); $123,799; 12 mos.

This project will continue research in five areas:

1. A theory of predicate transformers.
2. Synthesizing proof obligations for concurrent programs.
3. Investigation of trace logics.
4. A script-model editor, and
5. Sequential programming.

State University of New York - Stony Brook; Jieh Hsiang and Mandayam Srivas; *Theory and Applications of Term Rewriting Systems*; (DCR-8401624); $169,500; 24 mos.

The notion of term rewriting systems was first used in the 60's by Knuth & Bendix for solving word problems in universal algebra. The last decade has witnessed a surge of interest in term rewriting systems triggered mainly by an increasing understanding of the theory, and identification of new areas of application such as symbolic manipulation, specification and verification of data types, program generation, and programming language design. Ironically, the theory had largely been ignored by the theorem proving community because it was applicable only to equational theories. Recently, a way has been found to extend rewrite methods to first order predicate calculus. Several complete strategies have been discovered and experimental results are encouraging. This project will continue exploring the use of term rewriting systems for deductive reasoning in first order theories.

Program transformation is another area where the term rewriting method can be applied. Rewrite-rule based inference techniques for reasoning about the transformation of applicative programs are being studied. With the term rewriting mechanism, the method being used can also operate in the inductive theory while other approaches cannot.

Oregon Graduate Center; Richard B. Kieburtz and Robert G. Babb; *Software Templates (Computer Research)*; (DCR-8303927 A01); $57,427; 12 mos.

Software Templates is a name coined to describe high-level specifications of computer programs. High-level specifications have several advantages over explicit programs; they are easier to read and comprehend, and they are not specialized to a particular programming environment. Thus a software template needs to be written only once to specify an algorithm, and can be re-used in many different programs and for many different applications.

However, it has not been possible to automatically translate high-level specifications of algorithms into machine-executable programs. For this reason, high-level specifications have not been used to a great extent in practice. A concrete program, given in one of the current programming languages for which a complier or interpreter is available, has the advantage of reality, i.e. it can be executed on a machine, which an abstract specification does not.

This project will endeavor to find out whether recent research results provide sufficient base to develop a technology for translating high-level specifications directly into programs. There is some risk that the research may fail; if it does, it will nevertheless provide some indication of what remains to be done. If it succeeds, it will provide a basis for further development of a technology of automatic programming.

Brown University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8308806);

$45,911; 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant $91,822).

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis and empirical studies to verify the analysis. At best this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. Even if these goals are not achieved, the research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complex analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research will develop the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

Distributed Computing

University of Arizona; Gregory R. Andrews and Richard D. Schlichting; *Concurrent and Fault-Tolerant Programming*; (DCR-8402090); $90,217; 12 mos.

This project is concerned with the design of language notations for concurrent systems, and the development of techniques for the construction of fault-tolerant distributed software. In particular, it is investigating two specific topics: the design and implementation of a family of languages for concurrent programming called Darwin, and the use of fail-stop processors as a basis for fault-tolerant programming.

The goal of the Darwin project is to develop a collection of languages for the different levels of systems programming. The languages will contain a common core, and each family member will contain mechanisms specifically suited for programming one of four levels: kernels and interrupt handlers, shared-memory systems, distributed systems, and fault-tolerant systems.

A fail-stop processor is a virtual processor that suffers only one type of failure: a "crash". This project will investigate the development of software for execution of (an approximation of) such a machine. The basis for this software will be fault-tolerant actions, a language construct that allows recovery protocols to be applied to arbitrary sequences of statements.

The two projects are related in that the highest-level member of Darwin will include mechanisms for fault-tolerant programming such as fault-tolerant actions.

Yale University; Michael J. Fischer; *Theory of Algorithms and Distributed Systems (Computer Research)*; (DCR-8405478); $10,469; 12 mos. (Joint support with the Theoretical Computer Science Program - Total Grant $46,064).

Theory is developed in the area of algorithms and distributed systems. The goals are to apply theoretical methods and tools to practically-motivated problems arising in applied areas of computer science, particularly in distributed computing. Topics studied include methods of achieving fault-tolerance in unreliable systems, a logic of knowledge for describing protocols, and other problems in the area of algorithms and efficient data structures.

Georgia Institute of Technology; Martin S. McKendry; *An Operating System Architecture for Reliable Distributed Computing (Computer Research)*; (DCR-8316590); $75,372; 12 mos.

The physical attributes of distributed systems indicate that it should be possible to construct programs that tolerate failures better and run faster than they would on single machines. However, despite advances in hardware technology, especially in local area networks, this is not currently possible. Why?

This project is based on the belief that the primary reason is that the architectural foundations of centralized systems cannot be extended to the distributed environment without also carrying over the limitatic of centralized systems. In particular, conventional architectures do not account for failures. Means of supporting distributed programs that provide mechanisms needed to manage failures have been studied. The techniques are under development in the Clouds group at Georgia Tech, and are based on the notions of objects and atomic actions. Objects (instances of abstract data types) are a means of encapsulating data and guaranteeing data invariants. Thus, they provide a basis for localizing the effects of failures. Actions are a generalization of the notion of a transaction. They provide

the basis for synchronization, and they delimit the scope of failures. Actions consist of operations on objects.

The techniques to be used offer extremely high benefits. Suport for recovery is integrated into the lowest levels of the system. Thus, all system components can be designed to survive failures. Reliable computations can then be constructed to reconfigure automatically to accomodate failures. The facilities provided will simplify both application systems and the operating system itself.

Massachusetts Institute of Technology; Barbara H. Liskov; *Issues in Fault-Tolerant Distributed Programs (Computer Research)*; (DCR-8203486 A02); $143,802; 12 mos.

This research is concerned with problems that arise in fault-tolerant distributed programs. While interests within this area are very broad, ranging from distributed algorithms to program structure to program specification and verification, the project is focussed on defining a programming language and system, called Argus, that supports the construction and execution of fault-tolerant distributed programs. Argus will be used to study the other problems, both at the level of user programs, and at the level of system support.

The main recent achievement has been the design of Argus. This language is designed to satisfy a number of requirements that have been identified by studying application programs. Issues that arise in the implementation of Argus have been examined and some algorithms in this area have been developed.

The project will continue work on the language design and will do a realistic, distributed implementation of Argus. Theoretical work to enhance the understanding of new concepts that are arising will also be carried out.

Verification

Artificial Intelligence Research Institute - Texas; Frank M. Brown; *Verification Of Advanced Language (Computer Research)*; (DCR-8402412); $49,603; 12 mos.

The overall goal of this project is an integrated deductive system applicable to program verification starting with an existing system called SYMEVAL. To reach this goal many difficult problems will need to be solved, including the proof by induction of fully quantified theorems and the formalization of nondeterministic programming-language features.

William Marsh Rice University; Robert S. Cartwright; *Program Specification and Testing within an Integrated Programming Environment*; (DCR-8403530); $172,267; 24 mos.

Despite the growing importance of program specifications in programming methodology, there are virtually no practical software tools that exploit them. In fact, most programming enviroments do not accommodate program specifications at all. Programming environment research has largely neglected program specifications in favor of issues such syntax directed editors, source level debuggers, and program management systems.

In contrast, we believe that formal specifications should play a central role in the programming process and that programming environments should actively assist the programmer in this regard. In the presence of formal specifications, a much more systematic approach to software development and validation is possible. Fortunately, most of the technology required to construct such an environment already exists. Fifteen years of vigorous research in program semantics and verification have laid most of the necessary theoretical groundwork. The primary goal of this research project is to adapt and extend that work to form a comprehensive practical discipline of program specifications and testing.

The research program breaks down into three parts. First, to deepen our understanding of the mathematical foundations of program specifications and use that knowledge to improve the design of specification languages. The success of a specification and testing system hinges on the quality of the underlying specification language. Of particular interest are the special problems posed by floating point arithmetic and concurrency. Second, to develop more sophisticated methods-employing a judicious combination of verification and testing-for certifying that programs implement their formal specifications. Finally, to evaluate the effectiveness of ideas, an experimental program specification and testing system will be built as an extension to the Programming Environment already under development at Rice.

University of Texas - Austin; Robert S. Boyer and J. Strother Moore; *Mechanizing the Mathematics of Comput. Program Analysis*; (DCR-8202943 A02); $101,647; 12 mos.

The aim of this project is to make practical the use of computers to prove facts about computer programs. A practical, mechanical method for certifying computer programs would substantially increase their reliability and decrease their cost. A logic convenient for specify-

ing properties of computations has been developed and a computer system using that logic has been applied to prove facts about FORTRAN programs, functional programs, machine code programs, and real time control programs. The system has even been used to prove the soundness of some additions to itself. The development of the logic and construction of the system was supported in part by NSF Grant MCS-7904081. To enhance the logic of this system, the class of acceptable induction principles will be extended to permit the expression of schematic concepts with bound variables. Several improvements of the system's theorem-proving power will be made, including: the use of more decision procedures, the exploitation of facts about congruence relations, more effective control of forward and backward chaining, and the use of our theorem-prover by our compiler. The interface between the system and the user will also be improved.

Special Projects

Data Handling, Data Manipulation and Database Research

University of California - Berkeley; Randy H. Katz; *A Data Management System for VLSI Design Data (Computer Research)*; (DCR-8406123); $25,568; 8 mos.

This research emphasizes a database management system approach to managing the information about Very Large Scale Integrated Circuit (VLSI) designs. The database component, based on relational database techniques, will provide an interface to stored data. A design management system, that supports the hierarchical construction of a VLSI circuit design from primitive cells, will be built on top of the database component and will be responsible for organizing data about alternative design representations and versions. Finally, programs that provide a tailored interface to design data, based on its semantic content, can be built on top of the design data manager. The system should simplify the rapid construction of new tools for VLSI design by relieving the tools designer of the burden of managing the design data.

University of California - Santa Barbara; Donna J. Peuquet; *Study of Some Current Problems in the Design and Implementation of Raster-Based, Geographic Information Systems (Computer Research)*; (DCR-8309188); $53,323; 24 mos. (Joint support with Geography and Regional Science - Total Grant $78,323).

While the field of spatial data handling has certainly advanced since its inception, this advance has not kept pace with the introduction of computer applications in other areas of science, engineering and commerce. Part of this delay arises from the special, and largely unsolved, problems inherent in handling large volumes of two and three-dimensional coordinate data in the digital domain, but a significant amount can be traced to hesitation on the part of many analysts to turn to approaches which, while more efficient for computer processing, are not direct representations of the familiar point, symbol, line and area representations used in the traditional analog storage and display media, the map.

Many of the problems currently experienced in efficient computer handling of cartographic data are also encountered in such fields as image processing. This research will explore new alternative methods of representing spatial data in digital form and undertake comparative analysis of specific algorithm/data structure combinations. Empirical analyses will be carried out through the construction and use of a software testbed system.

New York University; Zvi Kedem; *Theory and Applications of Database Concurrency Control*; (DCR-8416422); $7,465; 7 mos.

Recent research has led to the development of non-two-phase locking protocols for controlling concurrency which assure correct results in a database without causing deadlocks. Such systems make no use of information about the operations performed by transactions on data entities, although they do use a priori information about the order in which transactions may access the entities. If information of the former type is also used, concurrency may be increased, but deadlocks are more difficult to avoid. The investigators are studying non-two-phase locking protocols for such "partially interpreted" systems, with the goal of using both semantic and syntactic information to increase concurrency and possibly also avoid deadlocks. They are also studying mechanisms for removing deadlocks when avoidance is impossible or too costly. Methods for conditioning the structure of a database to assure proper accessing behavior of the transactions so that such protocols can be used, as well as new general means for allowing concurrent dynamic modifications of the structure of the database, are being sought. The performance of all mechanisms developed will be evaluated in order to ascertain their applicability to real database systems.

Privacy and Security

SRI International; Dorothy E. Denning; *Multilevel Secure Database Systems (Computer Research)*; (DCR-8313650); $176,000; 24 mos.

This research project involves the study of fundamental questions concerning the design, verification and implementation of a multilevel secure database

system that would allow users with different security clearances to access a database containing data of different classifications. A model of a multilevel database system in which the classified objects are database views will be employed. A view constitutes a complete context, which is essential for classifying data by association, for specifying access and integrity constraints based on context, for creating sanitized reports, for relating an aggregate of data to its constituent parts, and for stating and proving properties about inference. The research goals include developing a methodology for defining views, for classifying them, and for proving the security requirements are satisfied. Special emphasis is given the inference problem, since solving this problem is crucial to preventing unauthorized disclosure of classified data.

University of California - Berkeley; Manuel Blum; *Transaction Protection Protocols (Computer Research)*; (DCR-8204506 A02 & A03); $63,300; 12 mos.

The objective of this research is to develop, study and experiment with protocols for protecting electronic business transactions. A number of such protocols have already been developed for sending certified mail, signing contracts, and exchanging secrets. A problem in each case is to achieve the equivalent of PARALLELISM through SEQUENTIAL exchange of messages. Signing contracts, for example, normally requires that the two parties to the contract be present at the same place and at the same time for signing. Signing contracts by telephone is difficult, even with computers, since whoever signs first obilgates him/herself first. It is the task of the protocol to achieve the EQUIVALENT of simultaneous signing. This is possible using techniques of number theory and complexity theory that have been developed recently for work on public-key cryptography and signature schemes. The investigators intend to develop protocols for various problems such as these and to further improve those already available in hopes of making them practical. With a few notable exceptions, the published literature does not give proofs of correctness and security, though these are starting to be developed. The investigators expect to contribute to the development of such proofs.

Drexel University; Neal R. Wagner; *Database Security*; (DCR-8403350); $54,100; 24 mos.

Professor Wagner will study several approaches to database security, including the use of cryptography in a fundamental way and the use of fingerprinting. Current work in database security has not incorporated cryptography into the basic design of the database, but has at most included it in an external framework. This project will employ techniques that use cryptography to create a database with enhanced security, though without the full functionality of a modern database management system. Preliminary versions of the basic techniques in question have been described by Professor Wagner and others in recent publications.

A second part of the project will study specific uses of fingerprints on data (that is, traceable alterations) to detect security breaches. Here there are promising statistical methods that have not yet been exploited.

The research will also incorporate other known database security techniques (such as the restriction of statistics and the addition of noise) into a specific prototype implementation oriented toward secure personal workstations.

University of Wisconsin - Milwaukee; George I. Davida; *Data Security (Computer Research)*; (DCR-8112345 A03); $9,000; 12 mos.

This research project involves data security and concentrates on the topics of operating systems security, encrypted relational databases, and key management problems. The researchers are interested in designing the encryption algorithms necessary to support the security issues being addressed. They have installed a VLSI design station that is operational and are in the process of designing VLSI encryption systems. They are able to display the designs, however, they have been unable to get a hard copy of the designs. The grant is being supplemented to permit the purchase of an ACT printer/plotter that interfaces to the VECTRIX terminal they have. This will allow the generation of the hard copy directly from the screen without any software modifications.

Societal Impact of Computing

University of California - Irvine; Rob Kling; *Social Dynamics of Computer Use in Complex Organizations*; (DCR-8117719 A03); $99,868; 12 mos.

This study investigates three compelling questions that face social analysts of computing:

1. How does the social organization of computing development and delivery influence the kind and quality of computing services available to users?
2. Which theoretical perspectives provide the most explanatory power for analyzing computing developments?

3. How is the interplay between complex information systems and complex social settings best understood?

This project will help answer these questions through a set of theoretical and empirical studies which examine the development and use of a specific type of complex computerized information system - material requirements planning (MRP).

Social analysts of computing often rely upon one of four theoretical perspectives in conceptualizing the social and organizational environment of computing: formal-rational, structural, interactionist, and organizational politics. Each of these perspectives casts a different analytical light on the character of the local computing package. This project will examine the relative explanatory power of these four perspectives in analyzing the development and use of MRP systems.

Four organizations will be studied in depth. They vary considerably in the organizational complexity of their computing arrangements and in the technical complexity of the computing development and use which vary with these complexities.

Modelling and Simulation

Wayne State University; Bernard P. Zeigler; *Distributed Simulation of Hierarchical, Multicomponent Models*; (DCR-8407230); $36,702; 12 mos.

Efforts will be continued to elaborate a methodology in which multicomponent discrete event models may be simulated employing multiprocessor architectures. A major part of the research concerns itself with the mapping of hierarchical multicomponent models onto distributed simulators so that correct and efficient simulation results. Test classes of models will be generated from case studies of computer networks as well as employing properties expressable within the discrete event formalism. Procedures for simplification of multilevel models so as to fit the constraints of multiprocessor architectures will also be developed. Theoretical concepts and results will be cross-checked with experiments on models of distributed simulators expressed in conventional simulation languages and, when it becomes available, on a commercial network-in-a-box multiprocessor system.

The results of the research will 1) contribute to the computer science area by increasing the understanding of the utilization of a potentially powerful class of architectures, and 2) contribute to systems science because of the extensions that will be made to the theory of system representation so that it applies to the enlarged domain of distributed computer simulation.

Networking, Computer Communications, and Distributed Computing

Purdue University; Douglas E. Comer and J. Tim Korb; *High-Level Network Protocols (Computer Research)*; (DCR-8219178); $100,000; 12 mos. (Joint support with the Coordinated Experimental Research Program - Total Grant $238,054).

This research centers on three major problems ocurring in distributed and computer networks. The first area of research is on algorithms and data structures to support electronic communication. As electronic mail services provide more information flow, we will have to find ways to manage this information. The second project concerns integrating the personal workstation into a general network where the computational resources of the entire network are available to the user. The final project will deal with the problem of managing data concerning project development in a general distribut . environment. As automation aids become more available, we must manage and coordinate the activities of many users in a network which has both shared and independent computational and storage resources.

Massachusetts Institute of Technology; Robert G. Gallager and Pierre A. Humblet; *The Dynamics of Data Network Research*; (ECS-8310698); $33,000; 12 mos. (Joint support with the Electrical and Optical Communications Program - Total Grant $122,744).

This group has approached their research to communication networks by developing, in a fundamental way, new models and new algorithmic and analytical techniques. This study invoives numerous problems the solution of which will reduce the ad-hoc procedures employed in the design of networks. Their

approach requires that each individual investigator have an understanding of communication and information theory, computer science, control theory, and operations research. Their trademark is a thorough, indept investigation which results in a fundamental contribution. As a consequence of this approach, their research commitments are long-range.

University of Massachusetts - Amherst; Victor R. Lesser; *Coordination in Cooperative Distributed Problem Solving Systems (Computer Research)*; (DCR-8300239 A01); $20,282; 12 mos. (Joint support with the Intelligent Systems Program and the Computer Systems Design Program - Total Grant $120,847).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts-Amherst has included the development of a node architecture capable of such sophisticated local decisionmaking. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. The research being carried out will build on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include: organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

University of Massachusetts - Amherst; Edward M. Riseman; *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776); $110,000; 12 mos. (Joint support with the Intelligent Systems Program and the Software Engineering Program - Total Grant $160,000).

This grant will support group research activity within the Computer and Information Science (COINS) Department in six areas of experimental computer research: computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under twenty-one other grants and contracts, including ten from NSF, this grant will provide partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support will be used for administration and operation of the COINS research facility, with the bulk of the funding going towards technical support staff and equipment maintenance.

Michigan State University; Herman D. Hughes; *Analysis of Priority Schemes for Local Networks*; (DCR-8318529); $56,726; 12 mos.

The advent of large scale integration (LSI) and very large scale integration (VLSI) devices provides inexpensive high-bandwidth communication between computers within a limited geographic territory, and is a significant event in the history of computing. The new technology has given rise to a proliferation of local area networks (LANs) during recent years, and this trend seems likely to continue during the next decade. As LANs and their applications (e.g. office automation, real-time processes, research and development, educational activities, etc.) increase in the United States and elsewhere, it becomes very important that they are designed to meet certain performance goals for a wide range of services (voice, video, facsimile, etc.). Three important design parameters for LANs are:

1. The transmission media,
2. The topologies, and
3. The channel access protocols.

These parameters determine whether or not a given LAN can achieve certain performance goals.

This study will examine the design parameters for LANs recommended by the IEEE Standards Committee in regard to their performance (i.e., channel throughput, queueing delays) for integrated services, and publish reports on how these parameters may be altered (e.g., changing baseband to broadband, assigning priority to packets) to allow a single communication network to yield acceptable performances.

Cornell University; Sam Toueg; *Routing, Broadcasting and Deadlock-Prevention in Packet-Switching Networks (Computer Research)*; (DCR-8303135 A01); $42,865; 12 mos.

Problems related to efficient and reliable data communication in packet-switching networks are being investigated. In particular, the study involves failsafe distributed algorithms for routing, broadcasting, and deadlock prevention in failure-prone networks. Particular attention is being given to the minimization of certain complexity measures related to the performance of distributed algorithms. The main complexity measures that are considered are the computational and communication complexity, the convergence time, and the recovery time following a change in the network configuration.

University of Texas - Austin; Simon Lam; *Protocols for Communication and Resource Allocation (Computer Research)*; (ECS-8304734); $20,000; 12 mos. (Joint support with the Electrical and Optical Communications Program - Total Grant $72,343).

Resource allocation techniques and their implementation in a distributed network environment are being investigated. Networks based upon virtual channels are currently the dominant packet networks. A new model is used to study the performance of these networks. Specifically, problems of congestion control, routing, aid-to-aid flow control, and the effect of their interaction on network performance are being examined. Techniques for synchronizing network status information in both broadcast and non-broadcast networks, and scheduling in broadcast networks with limited synchronization are being sought. A hierarchial model is being used to develop methods to analyze communication protocols.

Other Projects, Symposia, Colloquia and Studies

National Academy of Sciences; Gary E. Clark, *Core Funding for the Computer Science and Technology Board of the National Research Council, National Academy of Sciences*; (DCR-8212683 A01); $25,939; 12 mos.

Approximately five years ago, the National Research Council established the Computer Science and Technology Board to consider disciplinary and scientific issues in computer science and relevant problems in associated technological areas. A group of highly qualified individuals from industry, government and academia comprises the Board and carries out its mandate. Results of its studies are published in reports by the National Academy of Sciences.

The Board has established panels that focus on individual issues and problems and meets three to four times per year to review the on-going work of the panels and consider new issues of concern.

This award enables the Board to continue addressing crucial problems in computer science and technology. The core funding needed is independent of the support of the work of the various panels.

Thinking Machines Corporation; Stephen Wolfram; *Workshop on Theories of Complexity: Common Frontiers of Physics, Biology, and Computation (Information Science and Computer Research)*; (IST-8410691); $5,000; 8 mos. (Joint support with the Division of Information Science and Technology, the Office of Interdisciplinary Research, and the Mathematics Special Projects Program - Total Grant $44,900).

The study of large systems of weekly interacting simple parts which exhibit complex and interesting behavior of the whole is a common feature of physics, biology, and computation. This Workshop provides a forum for investigating theories of such systems, and especially theories that are relevant to the understanding of a wide range of phenomena including physical systems, the functioning of the brain, biological life, and thinking machines. While the relevant disciplines of physics, biology, and computation bring their own distinctive point of view and methods to the study of the central problem, there are important models, such as cellular automata, spin glasses, and threshold networks, that arise in all these fields and already suggest great commonality and universality. Computation theory provides a mathematical framework as yet largely unexploited in physics and biology. Reciprocally, physics suggests approaches untried in biology and computation, and biology offers models the most intriguing and complex types of self-organizing behavior. There are many opportunities for cross-fertilization and unification. The general theory of complexity that might ultimately result would be a major scientific achievement.

University of Maryland - College Park; Carl Smith; *A Special Year for Logic and Related Aspects of Computer Science at the University of Maryland, College Park, 1984-1985*; (DCR-8413498); $10,695; 12 mos.

The Mathematics Department of the University of Maryland will hold a Special Year in Logic during the 1984 - 1985 academic year. The purpose of the year is to support the study of various aspects of logic. The organizing committee from the Mathematics Department is staffed by D. Keuker, E.G.K. Lopez-Escobar, and J. Owings. Since substantial overlap of research interests between the Computer Science and Mathematics Departments was perceived by the organizing committee, it has been decided to use the Special Year, in part, to promote research interaction between mathematicians and computer scientists. Mathematicians realize that computation is a fundamental aspect of pure math and logic and are interested in finding out what the problems and issues are. The computer scientists seek the expertise of the mathematicians and ways to bring this expertise to bear on their problems. The Special Year will also foster discussion of fundamental issues in the fields of semantics and logics of programs, recursion theory, automath, and non monotonic logics.

The plans call for a week devoted to each of the four areas which are of interest to logicians and computer scientists.

University of Oregon - Eugene; Eugene Luks; *Regional Conference on Computational Complexity Theory, University of Oregon, Eugene, Oregon, August 20-24, 1984*; (DMS-8443965); $2,000; 6 mos. (Joint support with the Mathematics Special Projects Program - Total Grant $22,532).

The regional conference planned to be held at the University of Oregon in Computational Complexity Theory is an exceptionally oversubscribed one owing to the intense interest in the programatic content.

Theoretical Computer Science

Concrete Complexity and Analysis of Algorithms

Stanford University; Andrew C. Yao; *Studies in Algorithms and Computational Co· · 'exity (Computer Research)*; (DCR-8308109 A01 & .A02); $130,959; 12 mos.

Research is conducted along two general directions: the development and analysis of complexity models for classes of computational problems, and the probabilistic analysis of algorithms. Three objectives are emphasized:

1. The development of general unifying models which permit the study of many individual problems within the same framework;
2. The proof of upper and lower bounds for individual problems;
3. The development of novel mathematical techniques for solving problems.

Specific topics investigated include the complexity of data structure problems, generating nonuniform random numbers, distributed computing, VLSI computations, probabilistic behavior of algorithms, and probabilistic complexity.

University of California - Berkeley; Michael Harrison; *Theory of Computing (Computer Research)*; (DCR-8311787); $104,678; 12 mos.

Several research directions are pursued. These include operating system protection, formal language theory questions, and string rewriting systems.

In operating system protection, a general model for a protection system is formulated so that "safety" questions are mathematically precise and so that security policies can be formulated and tested. Additionally, the addition of cryptographic techniques to operating systems will be investigated.

In formal language theory and string rewriting systems questions related to generalized pumping lemmas, grammatical transformations, the complexity of parsing and recognition, and equivalence questions for DOL systems are studied.

University of California - Berkeley; Ernest S. Kuh; *Theory and Algorithms for Layout Design (U.S.-China Cooperative Research)*; (INT-8318395); $2,500; 12 mos. (Joint support with the Division of International Programs and the Electrical and Optical Communications Program - Total Grant $19,819).

The objective of this proposal is to develop basic theoretical knowledge and efficient algorithms for the placement and interconnection of the large number of components and subsystems in computers designed to given specifications (layout design).

University of California - Berkeley; Eugene L. Lawler and Charles U. Martel; *Combinatorial Algorithms (Computer Research)*; (DCR-8311422); $82,068; 12 mos.

Efficient algorithms for sequencing and scheduling are studied in the general context of combinatorial optimization. A variety of important problems are studied including machine shop scheduling (important for the efficient utilization of scarce resources), the travelling salesman problem (important, for example, for the siting of refineries), and shortest path problems for large networks (important for computer communication problems). The goal is to develop fast algorithms for special cases of the above problems. Additionally, the theoretical framework for the generalized study of such problems is investigated.

University of Southern California; Leonard Adleman; *Computation Complexity and Its Relationship to Number Theory (Computer Research)*; (DCR-8022533 A03); $47,280; 12 mos.

Research is being conducted in three areas: Computational complexity, cryptography, and VLSI.

1. Computational Complexity. Several aspects of computational complexity will be investigated with particular emphasis on its relationship to number theory. Much will center on the P = NP problem.
2. Cryptography. Public security and privacy.
3. VLSI. The VLSI related problems are; layout, design aids, design paradigms, and complexity.

University of Colorado - Boulder; Harold N. Gabow; *Design and Analysis of Algorithms for Deterministic Scheduling and Related Problems (Computer Research)*; (DCR-8302648 A01); $27,384.

The goal of this work is to explore techniques and principles for designing algorithms; the criteria concerned are correctness and efficiency in the worst-case, asymptotic sense. The setting for the work is problems in combinatorial optimization; more specifically problems for scheduling jobs on processors, and related problems involving graphs and matroids. These do-

mains are rich and varied enough so that the fundamental tools should have applications beyond them.

One major area of investigation is the critical path or highest-level-first method of scheduling. This approach has been used successfully to solve a number of scheduling problems involving unit length jobs, precedence constraints, and more than one processor. Applications of the method, both in approximation algorithms for problems that are known (or suspected) to be NP-complete, and also exact algorithms for some open problems, are explored. Examples of the latter are uniform processor problems, about which little is known for nonpreemptive schedules.

Several one-processor scheduling problems investigated are examples of matroids. The general matroid versions of these problems are studied. The simplicity and generality of the matroid model allows algorithms to be developed that are both efficient and have broad applicability. Past work has shown this to be a profitable approach. Of particular interest is the development of augmenting path algorithms for the matroid parity problem.

Northwestern University; Der-Tsai Lee; *Concrete Computational Complexity (Computer Research)*; (DCR-8202359 A01); $22,750; 12 mos.

The main purpose of this research is to develop efficient algorithms for solving problems in computational geometry. Results in this new and growing research area have various applications in operations research, statistics, computer-aided design, computer graphics, etc. Specific topics studied are:

1. The construction of the weighted Voronoi diagrams and the applications to proximity-related problems;
2. The shortest path problems with barriers that are represented as polygons or circles;
3. The off-line insertion-retrieval problem and its applications; and,
4. Design of parallel algorithms for geometric problems and their implications.

More emphasis will be placed on worst-case analysis of the algorithms, but average-case analysis is also studied.

University of Chicago; Stuart A. Kurtz; *Computational Complexity: Relativization and Measure (Computer Research)*; (DCR-8305207 A01); $21,358; 12 mos.

Problems in computational complexity theory, especially relativizations, machine dependent tradeoffs, and measure theory, are studied.

There exist oracles relative to which:

1. the Meyer-Stockmeyer hierarchy exists;
2. the Berman-Hartmanis conjecture holds;
3. the Berman-Hartmanis conjecture does not hold.

Theorems will be sought which do not relativize in the hope of finding techniques which might resolve the P vs NP problem. Machine dependent tradeoff results will be examined since these techniques do not always relativize. Finally measure theoretic techniques applied to problems in complexity theory are studied.

Purdue University; Greg N. Frederickson; *Graph Algorithms and Data Structures (Computer Research)*; (DCR-8320124); $44,772; 12 mos.

This research focusses on three interrelated areas: data structures for update problems, algorithms for graph problems, and complexity issues in distributed computing. The common theme will be the investigation of efficient computation in these areas. The generation of improved proof techniques and the characterization of significant complexity relationships are the goals of this research.

Louisiana State University - Baton Rouge; Duncan A. Buell; *An Investigation of the CPS Factoring Method (Computer Research)*; (DCR-8311580); $33,125; 12 mos. (Joint support with the National Security Agency - Total Grant $58,125).

A new method of factoring, CPS, due to C. P. Schnorr, is studied. This has been shown to be theoretically very good, in terms of computations in quadratic class groups, given certain assumptions about quadratic class numbers. Some early statistical results indicate that these assumptions are indeed justified, and some asymptotic heuristics agree with the statistics. Problems have arisen, however, in attempting to implement the CPS method, primarily because the class group computations themselves are complex. The validity of the theoretical assumptions necessary for effective factoring by the CPS method is examined and new methods for class group computation that might be applicable in the implementation of this special case of class group computation are studied. The possible changes in computer architecture that would speed up these computations enough to make this method practical for factoring are also investigated.

Massachusetts Institute of Technology; Ravindran Kannan; *Computational Complexity of Numerical Algorithms (Computer Research)*; (DCR-8304770 A01); $27,148; 12 mos. (Joint support with the Software Engineering Program - Total Grant, $41,148).

Improved algorithms for computing the Smith and Hermite normal forms of an integer matrix along with necessary empirical improvements are studied. Theoretical improvements and implementations of algorithms for solving the same problems for matrices of polynomials are also studied. The "alternation translation" technique is studied in order to obtain results

about nondeterminism and determinism using the polynomial time hierarchy. Finally, several computational problems involving multivariate polynomials are attacked.

Massachusetts Institute of Technology; Albert R. Meyer; *Logic of Programs and Computational Complexity (Computer Research)*; (DCR-8010707 A06); $118,977; 12 mos.

The relations among several programming logics in the literature are considered and a study of their comparative expressive power is undertaken. These logics include dynamic logic (Pratt), algorithmic logic (Engeler), algorithmic logic (Salwicki), programming logic (Constable), and logic of effective definitions (Tiuryn). In computational complexity the investigators will consider the relation between the geometry and accessibility of memory structure, and speed of computation..

Massachusetts Institute of Technology; Ronald L. Rivest; *Concrete Computational Complexity (Computer Research)*; (DCR-8006938 A03 & A04); $126,677; 12 mos.

This research is concerned with the study of algorithms and computational complexity in the area of VLSI design and the theory of cryptology. The approach to VLSI design involves developing component packing algorithms and wire routing algorithms subject to the constraints imposed by VLSI manufacturing techniques. The approach to cryptology research involves the development of a conceptual framework that allows the measurement of the security of cryptosystems. This conceptual framework rests upon the development of a theory of "finite" computational complexity.

Johns Hopkins University; S. Rao Kosaraju; *Applications of Foundations of Computing (Computer Research)*; (DCR-8205167 A02); $57,390; 12 mos.

The objective of this research is the investigation of the applicability of principal theoretical tools developed in foundations of computing to more practical problems in computer science. The problems selected are from the design of sequential and parallel algorithms, mesh networks, VLSI, redistribution of computations in algorithms, and coding theory.

In the area of parallel algorithms, the research will focus on problems pertinent to VLSI such as the comparative power of networks and embeddings leading to efficient layout, and problems derived from the unbounded parallelism model. Research in mesh networks will emphasize algorithm design and properties of network topologies. Mesh designs are extremely useful for VLSI. In addition, the investigator will pursue the advantages of redistributing computations to increase the efficiency of algorithms. Finally, coding theory will be investigated as a technique in algorithm design.

University of Minnesota; Oscar H. Ibarra; *Problems in Computational Complexity (Computer Research)*; (DCR-8304756 A01 & A02); $81,575; 12 mos.

Research is conducted in the following areas:

1. Characterizations and computational complexity of some models for VLSI,
2. Automata-based complexity theory: complexity measure tradeoffs, nondeterminism versus determinism, hierarchies of computation, characterizations of complexity classes, and
3. Decision problems in automata theory, formal languages, and other areas of computer science.

General as well as specific problems are being studied and possible approaches are being investigated.

University of Minnesota; Sartaj K. Sahni; *Algorithmic Study of Combinatorial Problems (Computer Research)*; (DCR-8305567 & A01); $78,978; 12 mos.

Combinatorial problems are problems that deal with collections of objects in order to obtain information about the structures of those collections. In theoretical computer science the combinatorial problems are abstracted from problems that are interesting in the sense that they solve real world problems. An example of such a problem is a wire routing problem for VLSI circuits where it is desirable to connect circuit elements together minimizing the longest wire length between circuit elements. Typically combinatorial problems are computationally difficult to solve exactly and thus approximate solutions are sought.

University of Minnesota; Shankar M. Venkatesan; *An Investigation of Network Flows (Computer Research)*; (DCR-8402045); $49,632; 24 mos.

Prof. Venkatesan is investigating network flow problems. The main focus is on the closely related problems of finding maximum flows, circulations, shortest paths, optimum solutions for linear systems, and related problems.

The research involves:

1. developing faster serial and parallel algorithms for the above problems;
2. studying the relations between network flows and shortest paths (at least in the planar case);
3. extending the theory on planar networks and graphs;
4. studying applications of this theory to the above problems, and to problems in layout design;

5. studying the optimization of simple linear systems by reducing such systems to network flow problems.

The proposed solutions are new, fast, and easily parallelizable. New theoretical tools which provide a fresh insight into many problems in computer science are also developed.

Columbia University; Zvi Galil; *The Design of Efficient Algorithms and Proving Lower Bounds (Computer Research)*; (DCR-8303139 A01); $79,637; 12 mos.

The research consists of two parts:

1. The Design of Efficient Algorithms. The researcher and his students will study the best known algorithms for various combinatorial problems such as maximizing network flow, finding maximal (weighted) matching in bipartite (or general) graphs, graph isomorphism in general graphs (or graphs with bounded degree), a membership test for permutation groups and string matching. New efficient parallel algorithms for these problems are sought.
2. Proving Lower Bounds. The researcher and Pavol Duris continue their joint effort in deriving lower bounds for various restricted models of computation.

Cornell University; Juris Hartmanis; *Computational Complexity and Structure of Feasible Computations (Computer Research)*, (DCR-8301766 A01); $90,324; 12 mos.

During the last five years, research in computational complexity has revealed an unexpectedly rich structure of the feasible computations; it has yielded new research techniques and pinpointed new problems. The structure of the classic complexity classes, such as LOGTAPE, NLOGTAPE, PTIME, NPTIME, PTAPE, etc., is studied. In the first part of this research the relations between the existence of sets of different densities in these classes and the collapse of the corresponding higher deterministic and nondeterministic complexity classes are investigated. For example, there exist polynomially sparse sets in NP-P if and only if EXPTIME = NEXPTIME. This result has a variety of applications and has suggested new structural properties and relations between the classic complexity classes which will be investigated. The classic separation problem of these complexity classes by means of quasi-homomorphic reductions is also studied. Finally, the computational complexity of finite objects is studied in order to give a feasible and quantitative refinement and extension of the Kolmogorov-Chaitin randomness concepts.

Cornell University; John Hopcroft and Alan Demers; *A Program of Research in Robotics*; (ECS-8312096); $40,000; 12 mos. (Joint support with the Computer Engineering Program and the Intelligent Systems Program - Total Grant $248,291).

This is a multidisciplinary program of research in robotics.

The research program has two components:

1. A design and implementation component to provide a robotic environment in which to validate new algorithms and ideas on programming language and environment design, and
2. A theoretical program of research aimed at developing the underlying mathematical foundations necessary for robotics.

The design and implementation efforts have centered on developing a system to allow experimentation with a number of different aspects of the robot programming environment. One of the most important of these is the problem of object representation for single objects and generic classes of objects as well. For example, the system should allow a user to describe the concept bolt independent of any specific example. Other examples of areas to be investigated are the problems of motion and task planning algorithms, generic transformations (e.g. bolting, welding), and graphics-based user interfaces. The initial environment will be directed toward automated assembly, but the system will serve much more as a test-bed for new ideas in language and algorithm design than as a production robot programming environment.

State University of New York - Albany; Victor Y. Pan; *The Computational Complexity of Arithmetic Problems (Computer Research)*; (DCR-8203232 A02); $27,024; 12 mos.

Arithmetic computational problems such as matrix multiplication and inversion, DFT, and convolution have important applications. The efficiency of the algorithms for such problems can be measured by their arithmetic complexity, bit-operation complexity, stability, and the complexity of their logical structure. The object of this study is the efficiency of arithmetic algorithms in terms of the criteria listed above and the relationship among the different criteria. Despite the progress in recent studies in this area (which was particularly rapid in the study of arithmetic complexity), several important problems remain open. For instance, it is not known if Gaussian elimination is optimal for matrix inversion in terms of bit-operation complexity although it is known to be non-optimal in terms of arithmetic complexity. Recently introduced successful techniques and concepts are used as the starting point of this investigation. The results of the study are also applicable to coding and combinatorial problems.

State University of New York - Stony Brook; Alan Tucker; *Perfect Graphs, Circular-Arc Graphs, and Associated Algorithms (Mathematical Sciences and Computer Research)*; DMS-8301934); $5,000; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $38,700).

In his work on perfect graphs, circular-arc graphs and associated algorithms, Professor A. Tucker has created a unique blend of the mathematical study of combinatorics and its applications to theoretical computer research. The Strong Perfect Graph Conjecture is one of the most important open problems in combinatorial theory today. It concerns a fundamental graph theoretic duality involving cliques (sets of mutually adjacent vertices) and stable or independent sets(set of mutually non-adjacent vertices). The conjecture has been of particular interest to mathematical and computer scientists because of its relation to the development of perfect codes.

University of Oregon - Eugene; Eugene Luks; *The Computational Complexity of Algebraic Problems and Graph Isomorphism (Computer Research)*; (DCR-8403745 & A01); $29,849; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $41,841)

The research focuses on the computational complexity of fundamental algebraic problems, particularly those involving permutation and matrix groups. A major objective is a subexponential algorithm for graph isomorphism using new techniques which are linking such algorithmic problems to recent advances in finite group theory. This would simultaneously resolve many of the remaining complexity questions for permutation groups. Many of the results are extendible to matrix groups and algebras, where similar algorithmic questions are independently motivated. However, it is expected that some of these are provably hard and evidence to support that viewpoint is compiled.

University of Oregon - Eugene; Eugene Luks; *Regional Conference on Computational Complexity Theory, University of Oregon, Eugene, Oregon, August 20-24, 1984 (Mathematical Sciences and Computer Research)* (DMS-8404540); $10,776; 9 mos. (Joint support with the Mathematics Special Projects Program - Total Grant $21,532).

The proposal is for one of a number of regional conferences recommended for funding by the NSF, with supporting services for the processing of the proposals and for the conduct of the conferences to be provided by the Conference Board of the Mathematical Sciences (CBMS) through the mechanism of a contract with NSF.

Criteria for the recommendation of funding of each proposal were:

1. Principal Speaker. The principal speaker is a world leader in current research in the topic of the conference, as well as an effective and able lecturer and expositor.
2. Subject Matter. The topic of the conference is of real significance for both mathematics and computer science, and is one in which there exists considerable research activity.
3. Regionality. The topic of a conference was chosen from a field in which there is a need for stimulation and expansion of research activity in a broad geographic area (perhaps 300-500 miles in radius) surrounding the host institution.
4. Host Institution. It was deemed desirable for the host institution to have a modicum of competence in the subject area of a proposed conference in order that substantial benefits may accrue to the host institution's faculty. The principal speaker is not a member of the faculty of the host institution, nor is he (or she) from a neighboring institution.

St. Joseph's University; Ranan B. Banerji; *Conference in Algorithm Design and Computational Complexity, Philadelphia, Pennsylvania, July 16 - 20, 1984 (Computer Research)*; (DCR-8403626); $17,443; 12 mos.

The intent of this regional conference is to transfer technology from the theoretical studies of algorithms to the design, analysis, and implementation of algorithms on specific architectures. The conference is timely in that a body of literature on algorithm analysis has been built up, several principles of algorithm analysis are beginning to be understood, and the benefits of algorithm analysis in the environment of software development are recognized.

University of Pittsburgh; Errol L. Lloyd; *Heuristics for NP-Hard Optimization Problems (Computer Research)*; (DCR-8103713 A01); $47,672; 24 mos.

NP-hard optimization problems arise frequently in computer science. Optimization problems are problems whose solution requires the optimization of some cost criteria given a particular set of conditions and inputs. Two examples follow:

1. Given a set of crates of varying dimensions and a supply of identical railroad cars, pack the crates into a minimum number of cars.
2. Given a graph find a way to color vertices so that no two adjacent vertices have the same color, using the minimum number of colors. Many optimization problems are NP-hard. NP-hard problems are such that it is unlikely that any computationally feasible algorithm exists to solve the problems.

Although exact solutions to NP-hard optimization problems most likely require infeasible amounts of computation time, approximate solutions can frequently be obtained with algorithms that are computationally feasible. Professor Lloyd proposes to search for effective heuristics for approximate solutions to NP-hard problems in three areas: scheduling, PLA folding and deadlock recovery. These are important problem areas in operations research and computer science, VLSI design, and operating system theory respectively.

Brown University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8308806); $45,911; 12 mos. (Joint support with the Software System Science Program - Total Grant $91,822.

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis, and empirical studies to verify the analysis. At best this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. Even if these goals are not achieved, the research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complex analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research will develop the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

University of Washington; Faith E. Fich; *Concrete Complexity in Parallel Models of Computation (Computer Research)*; (DCR-8402676); $20,000; 12 mos.

Upper and lower bounds for specific problems in various models of parallel computation are studied. For concurrent-read, concurrent-write, parallel random access machines, the effect of different write conflict resolution schemes are investigated. When concurrent-read, but not concurrent-write, is allowed, attention is focussed on the relationship between adaptive and nonadaptive algorithms. Finally, the tradeoff between depth and size for parallel prefix circuits is examined.

University of Wisconsin - Madison; Samuel W. Bent; *Tree Algorithms (Computer Research)*; (DCR-8402402); $20,000; 12 mos.

Trees are a frequently occurring data structure in combinatorial computing. Algorithms that operate on trees, or that use trees to represent other combinatorial objects, can make a significant difference in the efficiency of applications algorithms. The general goals of research on tree algorithms are to improve the efficiency of current algorithms, to add new tree schemes to the repertoire of the algorithm designer, to unify the theories of similar classes of trees so as to provide a clearer basis for comparison, and to understand better the power and limitations of trees as a method for storing information.

The particular problems studied involve biased search trees (a scheme for storing objects with dynamically changing non-uniform access frequencies), balanced trees (schemes for distributing the information content uniformly throughout the tree), and ranking and enumeration (methods for ordering and generating all trees of a given size).

Parallel, Distributed and Data Base Models

Stanford University; Christos H. Papadimitriou; *Algorithms, Complexity and Database Theory (Computer Research)*; (DCR-8320000); $94,513; 12 mos.

Professor Papadimitriou is studying a number of areas in the fields of the theory of algorithms and computational complexity. These include linear programming problems, linear decision tree models, decision making under uncertainty, graph theory, and data base theory. Two types of results are sought: foundational and problem specific. The former is illustrated by the search for a proper framework to study the linear programming problem. In one sense this problem is "polynomial" yet it can be argued that the polynomial solution is really "exponential". Satisfactory resolution of this question is sought. An example of a prob-

lem specific research area is the study of graph algorithms where the complexity of questions such as "Does a digraph have an even cycle?" are investigated.

Stanford University; Jeffrey D. Ullman; *Theoretical Investigations into Very Large Scale Integrated Circuit Technology (Computer Research)*; (DCR-8203405 A02); $67,233; 12 mos.

Concepts from theoretical computer science are applied to problems of integrated circuit technology. There are a variety of tools that the designer of VLSI circuits requires. The complexity of the algorithms used in such tools (e.g., simulators, design rule checkers) is investigated in the hope of developing superior algorithms or at least understanding the tradeoffs between speed and generality of input that must be made. Another area ripe for the design of tools is automatic routing. While the general problem is intractable, simple cases have been solved recently, and these solutions influence the way man and machine should interact to design circuits. There are also many algorithm design questions that are influenced by VLSI technology. For example, the design of systolic algorithms becomes very important, as does the study of parallelism in general. Complexity takes on new meaning, with issues such as pipelining, off-chip communication, and replication of inputs joining time and space as interacting criteria of goodness for algorithms. Hence, techniques for the development of good VLSI-oriented algorithms, with consideration given to the fundamental limits implied by the new models, are investigated.

University of Southern California; Seymour Ginsburg; *Mathematical Foundation of Data Bases (Computer Research)*; (DCR-8318752); $36,884; 12 mos.

This study is concerned with theoretical aspects of data bases. The relational database model is extended in order to take advantage of naturally occuring semantic constraints. The algebraic properties of these constraints are studied and applications to database efficiency and the physical design of database machines are explored.

Yale University; Michael J. Fischer; *Theory of Algorithms and Distributed Systems (Computer Research)*; (DCR-8405478); $35,595; 12 mos. (Joint support with the Software Systems Science Program- Total Grant $46,064).

Theory is developed in the area of algorithms and distributed systems. The goals are to apply theoretical methods and tools to practically-motivated problems arising in applied areas of computer science, particularly in distributed computing. Topics studied include methods of achieving fault-tolerance in unreliable systems, a logic of knowledge for describing protocols, and other problems in the area of algorithms and efficient data structures.

University of Georgia; Jeffrey W. Smith; *Processors with Parallel Capability for Factoring Integers (Computer Research)*; (DCR-8302877 A02); $52,640; 12 mos. (Joint support with the Computer Systems Design Program - Total Grant $52,640).

Special processors for factoring large integers, using the continued fraction algorithm with early exits, are being developed. The machines handle extended precision operands and have some parallel capability as computational accelerators. This is a collaborative research project with Professor Samuel Wagstaff at Purdue University. Professor Smith is designing the processors, developing their systems software, and evaluating their operation. Professor Wagstaff is analyzing the algorithm with the computational accelerators in mind and writing the application programs.

A Denelcor Heterogeneous Element Processor computer will be used for the two projects. The quadratic sieve integer factoring algorithm, which is the most serious known competitor of the continued fraction method, will be studied and a large-scale test of the Extended Riemann Hypothesis will be made.

Purdue University; Samuel S. Wagstaff; *Processors with Parallel Capability for Factoring Integers (Computer Research)*; (DCR-8406596); $29,520; 12 mos. (Joint Support with the Computer Systems Design Program - Total Grant $29,520).

Special processors for factoring large integers, using the continued fraction algorithm with early exits, are being developed. The machines handle extended precision operands and have some parallel capability as computational accelerators. This is a collaborative research project with Professor Jeffrey Smith at the University of Georgia. Professor Smith is designing the processors, developing their systems software, and evaluating their operation. Professor Wagstaff is analyzing the algorithm with the computational accelerators in mind and writing the application programs.

A Denelcor Heterogeneous Element Processor computer will be used for the two projects. The quadratic sieve integer factoring algorithm, which is the most serious known competitor of the continued fraction method, will be studied and a large-scale test of the Extended Riemann Hypothesis will be made.

Harvard University; Michael O. Rabin; *Randomized Algorithms and Concurrent Computations (Computer Research)*, (DCR-8121431 A03); $125,362; 12 mos.

Research is concentrated on three interrelating themes centering around randomized algorithms on the one hand and concurrent processes on the other hand:

Randomized algorithms for hash functions and pattern matching problems; algorithms for concurrency control and synchronization; concurrent operations on very large data-structures. The implementation of these algorithms in VLSI is also studied.

Harvard University; Leslie G. Valiant; *Parallel Computation (Computer Research)*; (DCR-8302385 A01); $138,200; 12 mos.

This project is concerned with investigating the fundamental potentials and limitations of highly parallel computers.

Emphasis is placed on the following topics:

1. Efficient algorithms for implementing basic systems functions, such as data communication, on highly parallel distributed machines,
2. Algorithms for compiling sequential algorithms into efficient parallel algorithms, and
3. Developing proof techniques for showing that for certain problems there are inherent logical impediments to computing them fast in parallel.

Massachusetts Institute of Technology; Nancy Lynch, *High-Level Models for Reliable Distributed Computing (Computer Research)*; (DCR-8302391 A01); $147,708; 12 mos.

The specific goals of the project are:

1. To evaluate the nested transaction model as a choice of programming model for distributed systems,
2. To isolate and understand the many problems involved in the implementation of nested transactions, and
3. To improve on the usual presentation of distributed database concurrency control theory by viewing it as a special case of the theory of nested transactions.

University of Lowell; Robert J. Lechner; *Harmonic Analysis of Logic Functions I: Encoded Programmable Logic Arrays (Computer Research)*; (DCR-8305571 A01); $21,693; 12 mos. (Joint support with the Computer Systems Design Program - Total Grant $43,386).

This research is concerned with the development and evaluation of a new unified approach to array logic synthesis using VLSI technology on a scale which applies to software storage as well as to firmware storage and wired control logic within digital systems. Specifically, PLAs are imbedded within linear encoding transformations on their inputs and outputs and combined with ROMs as necessary to define more optimal implementations of logic functions or stored tables than are possible with either ROMs or PLAs alone. In contrast to more general array logic techniques, the imbedded PLA approach retains the homogeneous array device layout and interconnection topology that reduces design and testing costs.

Heuristic algorithms are developed and tested on realistic examples to evaluate the feasibility of the imbedded PLA approach. Theoretical studies are used to evaluate a more rigorous approach to canonical synthesis algorithms that do not require human intervention. A crucial outcome of the study is the identification of that class of functions for which the hybrid (ROM + imbedded PLA) approach to synthesis is more cost-effective than either ROMs or PLAs alone.

Princeton University; Richard J. Lipton; *Resource Trade-off Models (Computer Research)*; (DCR-8308827 A01 & A02); $71,125; 12 mos.

As fundamental limits on the speed of processors are approached, new ways of increasing computing power must be sought. One approach is to develop distributed computing systems wherein computing problems are automatically broken into pieces, distributed to separate computers for solution, and recombined into a final solution. This approach allows massive data bases to be decentralized, allows inexpensive processors to be effectively used in conjunction with mainframe computers, and allows difficult computing tasks to be completed in a reasonable amount of time.

An important theoretical problem that arises in the design of distributed computer systems is the consensus problem. This problem occurs because no one processor in a distributed system has full knowledge of the entire system. Thus information must be communicated between processors in order to be able to reach a "consensus" concerning the state of the computation.

Professor Lipton noted a connection between this problem and some deep results in combinatorics involving extensions of Ramsey's theorem. He was able to derive some non-trivial lower bounds for the problem (i.e., lower bounds on the amount of information that must be exchanged to reach a consensus). He also identified several important open problems and proposes to continue the research to attempt to solve those problems.

City University of New York - Brooklyn College; Kenneth McAloon; *Arithmetic Theories and Incompleteness Phenomena (Computer Research)*; (DCR-8304788 A01); $23,949; 12 mos; (Joint support with the Topology and Foundations Program - Total Grant $33,949).

The incompleteness phenomenon is an important characteristic of first order arithmetical systems. Given any mechanizable axiom system for arithmetic, this phenomenon consists of the existence of statements for which neither the statements nor their negations have proofs within the axiom system. These statements are generally complex self-referential statements.

New York University; Richard Cole and Chee-Keng Yap; *Techniques for Geometric Retrieval and Related Problems in Computational Geometry (Computer Research)*; (DCR-8401633); $57,500; 12 mos.

Computational techniques for geometric retrieval problems are investigated. Three techniques are investigated - invariant regions, coherence, and a form of divide and conquer. Problems studied include k-null problems in computational geometry and motion planning problems in robotics.

New York University; Colm O'Dunlaing and Chee-Keng Yap; *Motion Planning Problems in Robotics: Algorithmic Issues (Computer Research)*; (DCR-8401898); $84,500; 12 mos.

Motion planning is a central problem in robotics. As a purely algorithmic problem, it has recently been investigated extensively by a number of researchers.

This work is continued and extended in some new directions, notably:

1. Development of the retraction method. This powerful technique has given rise to the most efficient algorithms known in each case studied. Although its general applicability to motion planning problems seems clear, this remains to be proved.
2. Motion with bounded acceleration. Some initial work has begun by considering the problem of finding a motion with minimal maximum acceleration sufficient to avoid a given set of moving obstacles.
3. Coordination of motion. Near optimal algorithms for the cases of coordinating two or three discs have been given. Practical refinements of these algorithms and the coordination of two Arms in space are among the work to be done next.
4. Experimental issues. Implementation of the retraction algorithm for a moving disk is under way. Two issues are studied: the "robustness" of the method and practical methods to shorten the path obtained through the Voronoi diagram.

New York University; Paul Spirakis; *Distributed Systems, Parallel Graph Algorithms and Performance of Concurrency Control (Computer Research)*; (DCR-8300630 A01); $34,200; 12 mos.

Research topics studied include (a) synchronization in distributed systems, (b) expected time complexity of parallel graph algorithms, and (c) performance analysis of concurrency control in databases. Topic (a) includes (1) the use of sophisticated probabilistic distributed coordination techniques to improve the performance of real-time algorithms. A near-optimal probabilistic resource allocation scheme is proposed. Aloha-type techniques, for adaptive selection of lengths of random waits for efficient synchronization are outlined. The question of how much one can relax the restrictions on relative speeds of processors and still get a real-time solution is investigated. Topic (a) includes (2) research for lower-bounds in the performance of deterministic distributed synchronization techniques in which processes are viewed as communicating finite automata, (3) applications of the real-time techniques to database concurrency control and to the implementation of multitasks in Ada, and (4) the possible effect of amount of knowledge about relative speeds of processors in differences of efficiency between synchronous and asynchronous systems. Topic (b) includes research on the expected time complexity of parallel graph algorithms (for connectivity, transitive closure, isomorphism, shortest paths, etc.) when the input is a random graph. In (c) a performance analysis methodology of concurrency control is proposed. It includes simulations, empirical modelling, and analysis of the models based on Markov processes, occupancy problems and hierarchical decomposition. It is applied to two-phase locking.

New York University; Uzi Vishkin; *Synchronous Parallel Computation-Efficiency of Resources (Computer Research)*; (DCR-8318874); $15,769; 12 mos.

Future ultra-large computers will be built with a design space (i.e., the model for which programs are written) which is very permissive. A certain permissive model of parallel computation is studied. The implementation space will have to meet, however, the limitations of available technologies. The interrelation between the design space and several suggestions for both the implementation space and simulation of the design space into the implementation space is analyzed. Algorithms in the design space concentrating on the complexity of the time, number of processors, and communication facility are studied.

Carnegie-Mellon University; Merrick L. Furst; *Resource Tradeoff Models (Computer Research)*; (DCR-8308805); $54,255; 12 mos.

The main research topic addressed concerns lower time bounds on multi-party protocols. Such protocols arise in distributed systems in which a number of parties (usually system processes) must reach a consensus under the constraints that individually none of the parties has all the information necessary to reach the consensus, yet the sum of the information available to all the parties is sufficient to reach the consensus. Such situations arise, for example, in trying to determine if two processes are simultaneously trying to access the same data items within a distributed environment.

Furst's contribution to this problem was to note a connection between the problem and some deep results in combinatorics involving extensions of Ramsey's theorem. He was able to derive some non-trivial lower bounds for the problem (i.e., lower bounds on the amount of information that must be exchanged to reach a consensus). He has identified several open problems and proposes to continue the research to attempt to solve those problems.

Brown University; John Savage; *VLSI Algorithms and Analysis*; (ECS-8306812); $40,835; 12 mos. (Joint support with the Computer Engineering Program - Total Grant $81,670).

The design of semiconductor chips with thousands or million of active elements is a task in which the management of complexity is a central problem that must be addressed and one that requires attention to methodology and to performance of algorithms.

The P.I.'s research concentrates on two principal topics:

1. Heuristic algorithms for silicon layout, and
2. Theoretical limits on the performance of parallel and VLSI algorithms.

The first topic seeks to improve a methodology for silicon compilation that produces chip layouts from functional descriptions.

The second topic seeks to understand limits on performance of parallel and VLSI algorithms where area, space, time, external space, and inequalities are sought.

Vanderbilt University; John J Grefenstette; *Parallel Algorithms (Computer Research)*; (DCR-8305693 A01); $68,292; 12 mos.

This project involves the design, analysis, and implementation of new parallel algorithms. Two areas are investigated. The first area of investigation concerns the average case behavior of parallel search algorithms, focusing on parallel branch-and-bound techniques. The tradeoffs between reducing the communication costs and efficiently pruning the search space are quantified. The second area of investigation concerns parallel numerical solutions for large sparse problems. The focus of this part of the project is the relationship between network topology and problem structure.

University of Washington; Richard E. Ladner; *Distributed Algorithms and Complexity (Computer Research)*; (DCR-8402565); $87,295; 12 mos.

Problems in the theory of distributed computing are investigated. These problems include the problem of coordinating a multiple access channel and the development of a game theoretic component as part of a mathematical framework for distributed computing theory.

University of Washington; Walter Ruzzo; *Computational Complexity Theory For Highly Parallel Systems*; (ECS-8306622); $22,000 12 mos. (Joint support with the Computer Engineering Program - Total Grant $44,000).

The technological capability exists to build very highly parallel computer systems. The technology to effectively use such systems is very much less well-developed. The goals of this research project are to help develop theoretical foundations for the design and analysis of algorithms for parallel computer systems, and to apply these principles to help elucidate appropriate parallel computer architectures and algorithms.

The five specific areas in which research is proposed are:

1. Models of parallel computers and their interrelationships,
2. Interconnection networks for parallel machines,
3. Lower bounds or other evidence that specific problems cannot be efficiently solved by parallel machines,
4. Parallel algorithms for a variety of natural problems, and
5. Computational issues arising in VLSI systems.

University of Washington; Martin Tompa; *The Combinatorial Structure of Computations and Symbolic Manipulation (Computer Research)*; (DCR-8301212 A01); $13,805; 12 mos. (Joint support with the Software Engineering Program - Total Grant $27,609).

The success of theoretical computer science is due largely to the process of abstracting clean, tractable problems from the morass of realistic detail. Yet there is much left to be desired in the understanding of the structure of computations. For instance, it is not known how to exploit the highly parallel systems that will soon be real. Neither is it known how to exploit randomness (in the sense of probabilistic, or Monte Carlo, computation) to derive efficient algorithms. Even in the setting of ordinary sequential computation, where a reasonbly good repertoire of algorithms exists, there are few techniques to show that those algorithms are optimal.

Part of the problem is that not enough of the detail is abstracted away. General computational models like Turing machines or random access machines have sufficient detail to defy analysis.

In the project, computations are abstracted further in order to lay bare the underlying combinatorial structures. With all inessential details stripped away, it is often simpler to discover new efficient algorithms, and demonstrate their optimality (at least within this structured framework) than it is with the details included. The areas studied are synchronous parallel complexity, sequential time and space complexity, and probabilistic complexity.

Stanford University; K. Jon Barwise; *Computational Aspects of Situation Semantics (Computer Research)*; (DCR-8403573); $9,693; 12 mos. (Joint support with the Intelligent System Program and Division of Information Science and Technology - Total Grant $29,079).

Situation semantics is a mathematical theory of linguistic meaning that grew out of artificial intelligence work in natural language processing. The goals of research in situation semantics include developing computationally feasible procedures for natural language understanding systems, developing the mathematical theory of situation semantics, and comparing the efficiency of situation semantics to more standard approaches to the theory of linguistic meaning. This is a collaborative project with Professor John R. Perry at Stanford University.

Stanford University; Zohar Manna; *Temporal Verification and Synthesis of Concurrent Programs (Computer Research)*; (DCR-8111586 A02); $56,505; 12 mos.

Temporal logic has been found to be a useful and powerful formalism for the analysis of concurrent programs. This research is aimed at developing the temporal methodology into a practical verification and synthesis tool. The ultimate goal is to construct an experimental Temporal logic verifier for concurrent programs, and to establish a temporally-based methodology for the synthesis of concurrent programs.

Stanford University; Stanley Peters; *Toward Automated Natural Language Processing: Phrase Linking Grammars for Syntax and Semantics (Information Science and Computer Research)*; (IST-8314396 A01); $13,000; 12 mos. (Joint with the Division of Information Science and Technology and Intelligent Systems Program - Total Grant $141,744).

This research project is focused on the characterization and understanding of natural language. An information structure called the "linked tree" is being used as a representation which will allow a syntactic and a semantic interpretation of natural language as well as assisting in the process of understanding natural language.

The investigators are particularly interested in the development of a restricted syntactic theory that is adequate to allow a revealing description of any natural language; an integration of semantics with syntax a formulation of a theory of parsing and interpretation that is consistent with available psycholinguistic evidence; and insight into the exploitation of parallelism to study the complexity of language.

Stanford University; Vaughan R. Pratt; *Logical Methods for Program Analysis (Computer Research)*; (DCR-8205451 A03); $41,618; 12 mos.

The research concerns the application of logical methods to the automatic analysis and management of storage, types, statistics, communication, dynamics, and checkpointing of dataflow programs. The work consists of the development of decision methods for fragments of logic tailored to the applications.

Wesleyan University; Kevin J. Compton; *Computational Problems in Finite Model Theory and Combinatorics (Computer Research)*; (DCR-8404233); $29,263; 24 mos.

The principal investigator, Kevin Compton, is investigating two types of problems. The first concerns questions in finite model theory which are closely related to computational problems. The major components of this research will be consideration of combinatorial properties expressible in logics more appropriate to asymptotic problems than first order and monadic second order logic; consideration of classes of structures which have figured prominently in combinatorics and computer science; and development of Ehrenfeucht game and analytic techniques for these logics and classes. The second type of problem concerns algorithms for the computation of fixed points in ordered semirings. This may be seen as a generalization of the well-known path problems in combinatorics. Results here should have applications to probabilistic computations (in particular, for Markov chains), formal language theory, and combinatorics.

Wesleyan University; Susan Landau; *Algebraic Algorithms and Computational Complexity (Computer Research)*; (DCR-8402173); $9,532; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $17,632).

One of the major mathematical results of the 19th century was the development of a method to determine the solvability of polynomials by Galois. The algorithms for this problem suffered from the defect that they required exponential computation time The PI, Professor Landau, developed an algorithm that solved this problem and which required only polynomial time.

Professor Landau is now studying algorithms for determining the order of Galois groups and for determining the generators of solvable Galois groups. An ultimate goal of this research is a polynomial time algorithm for determining the Galois group regardless of solvability.

Boston University; Peter Gacs; *Reliable Computation in a Homogeneous Medium and Algorithmic Information Theory (Computer Research)*; (DCR-8405270); $30,000; 17 mos.

Professor Gacs proposes to study two, mostly independent, topics. The first is reliable computation. This work builds upon a problem initially formulated by von Neumann. Given an array of automata, each of which may generate an incorrect result with probability e, can one build an equivalent array of automata so that computations by the entire array are correct with probability 1 - e (hence errors don't accumulate). This was partially solved by von Neumann using physically unrealizable components. Professor Gacs is extending one dimensional results he has already obtained in order to solve this problem fully. Such a solution would be both theoretically important and have practical implications for fault tolerant computing.

The second topic of study is algorithmic information theory. This topic rests on a notion of randomness formulated by Kolomogorov which measures the randomness of a sequence by the length of programs that generate the sequence. Professor Gacs is studying several problems in this area including a topic that links computational complexity and cryptography to algorithmic information theory.

Boston University; Leonid A. Levin; *Informational Complexity and Computational Efficiency (Computer Research)*; (DCR-8304498 A01); $48,800; 12 mos.

Informational Complexity Theory (ICT) is an area started from the observation by A. N. Kolmogorov and R. J. Solomonoff that there exists an invariant way to define the notion of Complexity K(x) of a word x - the length of the shortest program producing x. It was found to be fundamentally related to the ideas of randomness, a priori probability, information, inductive inference and others. Research in this area, with emphasis on the questions of computational efficiency, is continued. Examples of exciting recent achievements related to both areas are the results of M. Blum, S. Micali and A. Yao on pseudo-random number generators. Three particular topics are studied. The first - Linear Programming - does not really need ICT for its justification. It is included because many computational difficulties in ICT are related to handling convex functions. The second problem is the reformulation of ICT concepts in a "limited time environment". The third problem involves an application of ICT to VLSI models.

Harvard University; Harry R. Lewis; *Logic, Circuits and Oracles (Computer Research)*; (DCR-8402489); $25,469; 12 mos.

Three related areas of computational complexity are studied. First, the study of computational problems in logic is pursued, both decision problems relevant to other computational domains and specific algorithmic problems such as the unification computation. Second, the constant-depth circuit model is studied, both as it relates to logical complexity and with an eye towards extending recently obtained negative results on its computational power. Third, recursion theoretic techniques are applied to relativized versions of some of the basic problems of sequential and parallel computational complexity.

Smith College; Joan P. Hutchinson; *RUI: Separator Theorems for Graphs (Computer Research)*; (DCR-8411690); $17,000; 12 mos.

Graph theory and theoretical computer science have enjoyed a rich cross-fertilization of ideas for both disciplines study the same structures sometimes for the same reasons, sometimes for different ones. One area of mutual benefit has been in the development of separator theorems for graphs and in the design of related recursive algorithms, using the technique of "divide and conquer". A separator is a small set of vertices of a graph whose removal disconnects the graph and leaves all components small and of roughly equal size. When such a separator exists and can be found efficiently, other properties of the graph (e.g., a coloring or a maximum independent set) can be determined for the smaller components and pieced together in the original graph. This research focuses on families of graphs that have good separators, on variations in the kinds and characteristics of separators and on the applications of the resulting separator theorems. The research adds a new dimension to that currently being pursued in the mathematics and computer science departments of Smith College. Students majoring in either discipline will be selected for summer projects on the algorithmic and theoretical aspects of the research.

Williams College; Kim B. Bruce; *Models of Typed and Untyped Lambda Calculus (Computer Research)*; (DCR-8402700); $26,775; 12 mos.

Research is continued on the semantics of typed and untyped lambda calculus. The definition of environment models of the second-order lambda calculus and the related soundness and completeness theorems is refined. Combinatory models for this language are investigated. Similar work is done with the typed lambda calculus. Special effort is put into models of the typed lambda calculus where recursive domain equations must be satisfied. Concrete algebraic methods for solving recursive domain equations are also developed. Finally type assignment for the second-order lambda calculus is investigated.

University of Maryland; Carl Smith; *Automatic Program Synthesis and Structures Common to Programs and Data (Computer Research)*; (DCR-8301536 A01); $20,000.

This research is in two parts. The first involves a theoretical investigation of the process of how computers learn by example. Inductive inference can be formalized to yield a model where an algorithmic device inputs, over time, the graph of a recursive function and while doing so outputs programs intended to compute the input function. The emphasis of the first part of this research is on the study of inductive inference as a process used by computers to synthesize programs given examples of their input/output behavior.

In particular, the following problems are studied:

1. Program synthesis by teams of machines
2. Inference of approximate explanations
3. Complexity of inductive inference

During the course of previous research, the principal investigator developed a new model of computation wherein recursion is expressed implicitly and elegantly without the invocation of a "recursion theorem". The resulting model is the subject of the second part of the research. Problems concerning parallelism and general time/space tradeoffs will be investigated.

City University of New York - Baruch College; L. A. S. Kirby; *Models of Arithmetic (Mathematical Sciences & Computer Research)*; (DCR-8308412); $10,330; 12 mos. (Joint support with the Topology and Foundations Program - Total Grant $20,330).

Models of axiomatic systems are mathematical structures which satisfy the axioms and rules of inference of those axiomatic systems. For many axiomatic systems there exists more than one mathematical structure that satisfies the axioms. The intended structure (the one that the axioms were supposed to describe) is known as the standard structure or model; the others are known as non-standard models.

The PI is investigating models of the axioms of Peano arithmetic (that is, axioms intended to describe ordinary arithmetic). The study of these non-standard models leads to natural statements equivalent to Gödel's incompleteness theorem for arithmetic. The study also leads to results concerning the provability of important computational complexity problems, such as the P = NP question.

City University of New York - Brooklyn College; Rohit Parikh; *Logics of Programs (Computer Research)*; (DCR-8304959 A01); $51,673; 12 mos.

Research on logics of programs is pursued. Particular emphasis is placed on

1. Applications of process logic to actual examples, specifically to parallel computation,
2. Developing a logic for probabilistic computations,
3. Studying the relationship and similarity between dynamic logic and denotational semantics, and
4. Studying the computational complexity of programs and the strength of the logics used to prove their termination properties.

Clarkson College of Technology; James F. Lynch; *Problems in Finite Model Theory (Computer Research)*; (DCR-8402206); $18,126; 12 mos.

The objective of this project is to obtain further results in two areas of finite model theory. The first area is concerned with the asymptotic probability that a random finite model satisfies a given first-order sentence. The principal investigator intends to study the applicability of several new combinatorial methods for proving that sentences in certain first-order languages always have asymptotic probabilities.

The second area is definability of properties on finite models: given a finite model for each natural number n, and a set of relations on n, how large must a first-order sentence be to define the set of relations on the model? The principal investigator will study definability on models with a successor relation.

Columbia University; Joseph F. Traub; *Information and Complexity (Computer Research and Information Science)*; (DCR-8214322 A01); $56,103; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant, $71,103).

Research is continued on the solution of problems which cannot be exactly solved at finite cost, that is problems which can only be solved with uncertainty. The power of the approach stems from analysis at the information level rather than at the algorithm level.

Most problems can only be solved with uncertainty. This is true of problems from fields as diverse as physical science, biological science, economics, engineering, statistics, statistical decision theory and mathematics. Even the problems which could be solved exactly are sometimes solved probabilistically or approximately, that is, with uncertainty, to decrease their complexity.

Among the new topics studied are: Average Case Analysis Parallel Algorithms, Value of Information, Distributed Computation, Smoothness and Complexity, Complexity of Partial Differential Equations, Complexity of Linear Programming, New Applications.

University of Pittsburgh; Robert P. Daley; *Computational Complexity and Inductive Inference (Computer Research)*; (DCR-8402002); $26,655; 12 mos.

Three topics are investigated:

1. An axiomatic approach to the computational complexity of inductive inference,
2. The complexity of concrete procedures for inductive inference, and trade-offs between the complexity of such procedures and various criteria for successful inference, and
3. The differences in complexity between non-deterministic and deterministic computations which use little or no workspace.

University of South Carolina - Columbia; Jerrold Griggs and William Trotter; *Research in Combinatorics (Mathematical Sciences and Computer Research)*; (DMS-8401281); $6,000; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $44,400).

The research topics to be covered during the period of the award span a wide sprectrum of combinatorics and include problems in finite partially ordered sets, graph theory, Ramsey theory, discrete geometry, combinatorial algorithms, and theoretical computer science.

University of Washington; Victor L. Klee; *Convexity, Complexity and Combinatorics (Mathematical Sciences and Computer Research)*; (DMS-8116069 A02); $9,000; 12 mos. (Joint support with the Algebra and Number Theory Program - Total Grant $29,800).

Professor Klee will continue research in the areas of convex sets, optimization, computational complexity, and combinatorics. In particular, his attention will fall on questions related to sign quasistability for real square matrices; diameters of random bipartite graphs; relative minima among triangles containing a given convex polygon; the d-step conjecture in the combinatorial structure of convex polytopes; decision-tree complexity; computational complexity for convex polyhedra; tilings of infinite-dimensional normed linear spaces; uniform properties of convex bodies; preference relations and utility functions; and minimum graphs. Much of the work will be done in conjunction with students and colleagues.

University of Washington; Paul R. Young; *Mathematical Theory of Computation (Computer Research)*; (DCR-8319218 A01 & A02); $73,200; 12 mos.

Research is pursued in three areas:

1. Independence results for computer science.
2. The power of nondeterminism in the computation of discrete functions.
3. Analytic computational complexity.

The first two areas are concerned with discrete computational complexity while the third is concerned with the computation of real valued functions. Strong independence results would suggest that P = NP is dependent on the model of computation chosen, solutions in the second area would refine such computational models, while solutions in the third area would extend complexity resuits to computable real numbers and real analysis.

University of Wisconsin - Madison; Deborah A. Joseph; *Computational Complexity and Geometric Algorithms (Computer Research)*; (DCR-8402375); $38,766; 12 mos.

Several problems relating to the foundations of computer science and the design and analysis of algorithms are studied. These problems fall into two research areas, the goals of which are to provide additional insight into or answer the following very general questions:

1. What are reasonable formal systems and proof techniques for resolving separation questions concerning complexity classes?
2. In what situations can efficient algorithms be found for motion planning and collision avoidance problems that arise in robotics?

Automata and Language-Based Models

University of California - Santa Barbara; Ronald V. Book; *Computational Complexity Theory (Computer Research)*; (DCR-8312472); $40,279; 12 mos.

The major open question in theoretical computer science is the P = NP question. The P = NP question asks whether problems solvable by nondeterministic programs whose running times increase as a polynomial of the length of their inputs (NP) are also solvable by deterministic programs with polynomial running times (P). The effect of nondeterminism is to allow alternatives in a problem to be explored simultaneously. While no proof has yet been found, most researchers believe that P ≠ NP.

The reason the P = NP question is important is twofold:

1. The problems in P are considered to be "computationally feasible".
2. Representative problems in NP (known as NP-complete problems) constitute an important class of optimization problems.

All direct approaches to proving P = NP have failed. These approaches have concentrated on using "diagonalization" arguments developed by Cantor in

the 19th century in an attempt to construct a problem in NP that is different from all problems in P. One of the reasons for the failure of this method is the fact that if a diagonalization proof could show that P = NP, then the same proof could be used to show that P = NP in the presence of "oracles" where an oracle is an extension of the computational power of the programs used to solve the P and NP problems. In fact there are oracles such that both P = NP and P = NP are provable.

This research concentrates on studying the P = NP and related questions in the presence of oracles. The ultimate goal of the research is to gain an understanding of the computational power of nondeterminism.

University of California - Santa Barbara; Ronald V. Book; *Applications of Combinatorics on Words (Computer Research)*; (DCR-8314977); $31,949; 12 mos.

Term rewriting systems have been used in a number of areas of computer science, particularly in automated theorem proving and programming language research. Such systems transform strings of characters to strings of characters by rewriting part or all of the original string according to simple rules.

Professor Book proposes to investigate several applications of term rewriting systems. These include studying restricted systems with computationally easy properties, studying applications to context free languages, studying applications to tree-manipulation systems, and investigating network communication protocol applications.

University of Colorado - Boulder; Andrzej Ehrenfeucht; *New Directions in Language Theory (Computer Research)*, (DCR-8305245 A01); $35,995; 12 mos.

Research is conducted in several areas of formal language theory. The research topics studied include combinatorial properties of formal languages, the Post correspondence problem, the theory of codes, selective substitution grammars, and graph grammars. The main object of the research is to gain further insight into the basic structure of formal languages - addressing both generation questions and decomposition questions. The research on graph grammars may also have important applications in other areas of computer science.

Georgia Institute of Technology; Richard DeMillo and Kimberly King; *Models of Computation and Algorithms (Computer Research)*, (DCR-8103608 A03); $17,883.

Theoretical research will be performed on models of computation. In the area of parallel models and concurrency, two projects will be undertaken; one project introduces the notion of stochastic synchronization and will clarify properties of contention-based synchronization, while the other project is a comparative study of models of parallel computation. In the area of more conventional models a number of problems are addressed; computational complexity of alternating machines, logical independence, VLSI algorithms, and combinatorial graph embeddings.

Iowa State University; Alan L. Selman; *Complexity of Feasible Computations (Computer Research)*; (DCR-8402033); $48,810; 12 mos.

The computational complexity of feasible computations is studied. The objectives are to reveal the rich structure that individual complexity classes, such as NP, appear to have, and to increase understanding of the structural relations between low-level complexity classes. To these ends several techniques are applied.

Properties of polynomial time-bounded reducibilities on NP, and consequences of P = NP and similar hypotheses, are investigated and compared. Relativization techniques are applied, and fundamental issues are raised concerning the import of results about classes relativized to oracles.

Applications of these studies are made to problems in cryptocomplexity. Appropriate complexity measures are identified according to which public-key cryptosystems are to be hard, and then investigations are made to determine whether hard public-key cryptosystems can exist.

State University of New York - Albany; Harry B. Hunt; *The Complexity of Multi-Set Descriptors, Difference Equations, and Algebraic and Recursively Presented Structures (Computer Research)*; (DCR-8403014); $43,058; 12 mos.

Professor Hunt is investigating the following problems:

1. Multi-set descriptors, unambiguous language descriptors, and difference equations;
2. The complexity of algebra on structures such as lattices, semi-rings, rings and regular algebras; and
3. The complexity of analyzing recursively presented structures.

Particular problems to be considered include the following:

In area 1, finding minimal representations, testing equivalence of Multi-set descriptors, the algorithmic analysis and solution of difference equations;

In area 2, the development of relational algebras and relational calculi for logics other than two-valued logic, the determination of particular algebraic structures for which good algorithms exist for solving systems of simultaneous linear equations; and

In area 3, the complexity of very elementary and basic problems for various classes of recursively presented structures, especially that of determining extraneous productions or state-transitions.

Computer Research Equipment

University of California - Irvine; George S. Lueker; *Equipment for Computer Science Research*; (DCR-8404898); $123,809; 12 mos.

A distributed minicomputer and workstation system is being obtained in support of research in the following areas:

1. Software technology.
2. Architecture and operating systems,
3. Theoretical computer science,
4. Artificial intelligence, and
5. The applications and impacts of computing.

Georgia Institute of Technology; Raymond E. Miller; *Computer Research Equipment (Computer Science)*; (DCR-8405020); $207,330; 12 mos.

Several large minicomputers and several small microcomputers are being obtained to support research in the following areas:

1. Reliable operating system architecture,
2. The ProNet language for reliable distributed processing,
3. Integration of desk-top computers in a total computer environment, and
4. Studies of user interaction.

Boston University; Joyce Friedman; *Computer Research Equipment (Computer Science)*; (DCR-8405037); $26,036; 12 mos.

Additional memory and disk drives are being obtained for attachment to an existing minicomputer in support of research in the following areas:

1. Computer studies in formal linguistics,
2. Transportable natural language database update, and
3. Statistical data base security.

Harvard University; Harry R. Lewis; *Computer Research Equipment (Computer Science)*; (DCR-8405079); $66,744; 12 mos.

A minicomputer and workstation system is being obtained in support of research in the following areas:

1. Logic programming, software prototyping, and databases,
2. Program development systems, and
3. Knowledge-based object modelling.

University of Massachusetts - Amherst; David D. McDonald and Victor R. Lesser; *Equipment for Computer Research*; (DCR-8404899); $69,716; 12 mos.

The memory of existing large minicomputers is being upgraded and several workstations added in support of research in the following areas:

1. Distributed problem solving,
2. Intelligent interfaces,
3. Natural language processing, and
4. Vision processing.

Michigan State University; Richard C. Dubes; *Computer Research Equipment (Computer Science)*; (DCR-8404935); $85,000; 12 mos.

An image processing system is being obtained in support of research in the following areas:

1. Three-dimensional object acquisition,
2. Texture analysis, and
3. Two-dimensional grammatical inference.

Michigan Technological University; Karl J. Ottenstein; *RUI: Equipment for Computer Science Research*; (DCR-8404909); $45,140; 12 mos.

A large minicomputer system is being obtained in support of research in the following areas:

1. The program dependence graph in a software development environment,
2. Translation for multiprocessors, and
3. Measurement-based program and architecture optimization.

University of Minnesota; Sartaj K. Sahni; *Computer Research Equipment (Computer Science)*; (DCR-8403806); $82,500; 12 mos.

Several workstations are being obtained in support of research in the following areas:

1. Combinatorial problems,
2. Large sparse systems,
3. Large scale global optimization.
4. Workstation based queueing theoretic and performance models, and
5. Analysis of time-varying imagery.

Mississippi State University; Bradley D. Carter and Lois C. Boggess; *Computer Research Equipment (Computer Science)*;

(DCR-8405022); $9,425; 12 mos. (Joint support with the Office of Interdisciplinary Research - Total Grant $18,850)

Minicomputer-based workstations are being obtained in support of research in the following areas:
1. A text manipulation system for the handicapped,
2. Spatial modeling of locative prepositions, and
3. Models and methodology in automated data base reorganization.

North Dakota State University; Kenneth I. Magel; *Equipment for Computer Science Research*; (DCR-8404505); $15,795; 12 mos.

A network interconnected collection of microcomputers is being obtained in support of research in the following areas:
1. Ethernet implementation and experiments with different transmission patterns,
2. User-definable interfaces,
3. UNIX on the IBM PC,
4. Comparisons of approaches to distributed directory management, and
5. Examination of methods for processing distributed data base queries.

University of Nebraska - Lincoln; Roy F. Keller; *Computer Research Equipment (Computer Science)*; (DCR-8404902); $53,808; 12 mos.

A minicomputer system including an optical scanner, digitizer, high resolution printer, and high resolution graphic display is being obtained in support of research in the following areas:
1. Subspace selection for projection methods,
2. Hierarchical representation of optically scanned documents,
3. Digital image registration
4. Data base techniques for remote sensing and geo-information, and
5. Cartographic labeling.

City University of New York - Brooklyn College; Kenneth McAloon; *RUI: Computer Research Equipment (Computer Science)*; (DCR-8405477); $79,403; 12 mos.

A large minicomputer is being obtained in support of research in the following areas:
1. File system performance analysis: application to fault-tolerant systems,
2. B-trees, multiway search trees, and related structures,
3. Parallel processing architecture for real-time simulation of communication systems, and
4. Adaptive protocols for local area networks with intermittent real-time requirements.

Colgate University; Allen B. Tucker; *RUI: Computer Research Equipment (Computer Science)*; (DCR-8404552); $68,556; 12 mos.

A large minicomputer system including a color graphics display is being obtained in support of research in the following areas:
1. Design of processors for very large AI data bases,
2. Knowledge-based multilingual machine translation,
3. Efficient heuristics for production planning with uncertainty,
4. Cognitive science, and
5. Digital music synthesis.

Cornell University; Robert L. Constable; *Computer Research Equipment (Computer Science)*;(DCR-8406052); $83,556; 12 mos.

A large minicomputer is being obtained in support of research in the following areas:
1. Automated reasoning,
2. Intelligent systems,
3. Programming methodology,
4. Programming logic, and
5. Programming languages and environments.

New York University; Jacob T. Schwartz and Olof Widlund; *Computer Research Equipment (Computer Science)*; (DCR-8405004); $56,500; 12 mos.

A hard copy color output device, logic analyzer, and wide bed color plotter are being obtained in support of research in the following areas:
1. Robotics and computer vision,
2. Design studies for a very high-performance parallel computer, and
3. Computational many-body theory.

State University of New York - Buffalo; Patricia J. Eberlein; *Computer Research Equipment (Computer Science)*; (DCR-8405024); $35,969; 12 mos.

An image processing graphics system, a workstation, and Ethernet connection equipment are being obtained in support of research in the following areas:
1. Computational vision,
2. Denotational semantics, and
3. A graph-based programming language with built-in self reference.

State University of New York - Buffalo; Hinrich R. Martens; *Computer Research Equipment (Computer Science)*; (DCR-8404242); $26,132; 12 mos.

A minicomputer s, ·t is being obtained in support of research in the following areas:

1. Network operating system concepts utilizing shared I/O operation substructure,
2. An advanced multi-processor based simulation facility for the evaluation of simulation techniques and microprocessor based control algorithms, and
3. Design and analysis of local area network protocols.

Ohio State University; Bruce W. Weide; *Computer Research Equipment (Computer Science)*; (DCR-8405029); $65,960; 12 mos.

Three workstations are being obtained in support of research in the following areas:

1. The control of articulated motion,
2. Simulation of natural phenomena, and
3. Graphical programming languages for real-time systems.

University of Houston; Willis K. King; *Comp · Research Equipment (Computer Science)*; (DCR-840447); $.J3,059; 12 mos.

A distributed minicomputer system is being obtained in support of research in the following areas:

1. The experimental evaluation of a fuzzy set based approach to estimating the correctness of complex programs,
2. Experimental research in distributed data bases,
3. A method for displaying the dynamic behavior of real-time software systems,
4. Distributed algorithms and disk controller algorithms, and
5. Grantor-controlled authorization - a distributed implementation.

Virginia Polytechnic Institute and State University; Dennis G. Kafura; *Computer Research Equipment (Computer Science)*; (DCR-8404214); $85,000; 12 mos.

A large minicomputer system is being obtained in support of research in the following areas:

1. An interactive environment for constructing and executing functional programs,
2. Production of a discrete event simulation model development environment,
3. Experiments in information retrieval,
4. Development of software structure analysis tools and techniques, and
5. Research in logic programming and its applications.

Coordinated Experimental Research

Experimental Computer Research

University of Arizona; David R. Hanson; *A Programming Systems Laboratory*; (DCR-8320138); $883,000; 12 mos.

This project will aid in the establishment of a major experimental facility for research in programming languages and systems. The goal is to develop a comprehensive computing laboratory to support projects in programming languages, software validation, distributed computing, and programming environments.

Research in programming languages will be directed towards accelerating the design and implementation of very high-level languages, including the Icon programming language and new functional programming languages. Validation research will concentrate on numerical software testing and on integrated testing tools based on program mutation techniques. In distributed computing research, efforts will focus on operating systems, concurrent languages, fault-tolerant systems, and design and measurement tools and techniques. Finally, research on the environmental aspects of programming systems will include studies on the automatic generation of editing systems and user interfaces and on the further development of an integrated command and programming language.

University of California - Los Angeles; Algirdas Avizienis, Gerald Estrin, Leonard Kleinrock, Gerald Popek, Walter Karplus, and Terrence Gray; *An Advanced Network Environment for Distributed Systems Research*; (DCR-8121696 A05); $728,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support a program of research on advanced distributed computing systems and their applications. The facility will include a local-area computer network containing a small number of powerful computers, a large number of flexible minicomputers, and special purpose peripheral equipment.

This experimental facility is intended to immediately support investigations of the applications of distributed systems, as well as studies of the fundamental architectural issues of such systems. Results in both these directions will be returned to the facility itself in the form of evolutionary enhancements. The facility is at once both an object of and a tool for advanced experimental computer science research.

A variety of research projects are planned and underway. The proposed research program encompasses: distributed operating systems and distributed databases; performance of new distributed architectures and systems; distributed simulation of network architectures; innovative architectures and fault tolerance; performance evaluation of multiprocessor systems; system design tools and user application interfaces.

Yale University; Martin Schultz, Alan Perlis, and Roger C. Schank; *An Attached Processor Systems Laboratory*; (DCR-8106181 A03); $437,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support research in numerical analysis, scientific computing, algorithms, array processing, computer architecture, artificial intelligence, natural language processing, and systems programming. These new resources will enhance their computational capability and overcome address-space limitations which restrict their natural language research. They will also add a powerful array processor to support scientific computing and specialized processors for the LISP programming language. All these resources will communicate over a high bandwidth network, while terminal access will be provided through a switching network.

The completed local network of sophisticated computing resources will support cooperative research among three major research groups: the systems programming group, the scientific computing group, and the artificial intelligence group. The scientific computing group will concentrate on new algorithms and architectures, such as attached processors, to find ways to dramatically increase the scale of problems which can be considered for computer solution. The artificial intelligence group will pursue studies of integrated understanding and memory by modelling human memory organization and search strategies. Besides pursuing interests in distributed computing, the systems programming group will make major contributions to the others by bringing their experience to bear on the problems of supporting attached processors and implementing interactive problem understanding environments.

University of Illinois - Urbana; Charles W. Gear; *A Compute and Software Systems Design Laboratory*; (DCR-8105896 A04); $676,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support research and development of computer aids to the design of computer and software systems. The project will result in the development of methodologies, techniques and software for supporting various aspects of computer science research. Joint studies by computer science researchers of various backgrounds will be facilitated.

The facility will expand on current departmental resources adding three major computational resources and a supporting high bandwidth network to interconnect the current and new resources. Highly interactive access will be provided through user interfaces and a switching network. Terminals can be directly connected to a particular resource through the switching network while user and process communication is maintained over the high speed network.

A variety of research projects are planned and underway in four major areas: theoretical foundations, algorithm implementation, support software, and runtime systems. T' e topics include: theory of computation, creation of scientific software, VLSI technology, computer architecture, compilers, databases, inductive inference and conceptual data analysis, and software systems research. The enhanced facility will serve as a catalyst for cooperative and productive research.

Purdue Univers ty; Douglas E. Comer, J. Tim Korb and Walter Tichy; *High-L ' Network Protocols (Computer Research)*; (DCR-8219178); $138,054; 12 mos. (Joint support with the Special Project Program - Total Grant $238,054).

This research centers on three major problems ocurring in distributed and computer networks. The first area of research is on algorithms and data structures to support electronic communication. As electronic mail services provide more information flow, we will have to find ways to manage this information. The second proje t concerns integrating the personal workstation into a general network where the computational resources of the entire network are available to the user. The final project will deal with the problem of managing data concerning project development in a general distributed environment. As automation aids become more available, we must manage and coordinate the activities of many us .s in a network which has both shared and inde endent computational and storage resources.

University of Maryland - College Park; Gilbert W. Stewart, Ashok K. Agrawala, Victor Basili, Jack Minker, and Azriel Rosenfeld; *A Testbed For Parallel Algorithm Development*; (DCR-8219507 A01); $1,025,000. 12 mos.

This project will aid in the establishment of a major experimental facility for the development and testing of parallel algorithms. The facility will be developed around the highly-parallel, ring-structured, multicomputer ZMOB now being delivered to the Computer Science Department. The project will include the development of necessary hardware interfaces, hosts, and software systems to support the ZMOB machine.

The principal areas of research on parallel algorithms will be computer vision, problem solving, and numerical analysis; some research in distributed systems and language will also be conducted in addition to the creation of a software environment in a number of research areas, particularly in artificial intelligence and numerical analysis, and will involve collaborative efforts among a number of research groups within the department.

Cornell University; Robert L. Constable, Richard W. Conway, David J. Gries, and Alan J. Demers; *A Laboratory for Experiments on the Programming Process*; (DCR-8105763 A03); $550,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support research on the programming process at Cornell University. The investigators will consider a variety of topics including programming environments, the discipline of programming, and language design and implementation. While having distinct goals and clearly separate approaches the high degree of interaction among the researchers should provide results which strongly impact our understanding of programming.

The facility will include a variety of current and new computational resources interconnected by a local area network. While each researcher may access specialized resources for a particular project, all resources will be available to each user of the network. This environment can provide excellent communication between research projects while matching the appropriate resource to an individual project. Cornell has been characterized as being engaged in highly individualized computer science research in an atmosphere of close rapport. This facility will enhance the research opportunities while maintaining this highly productive environment.

New York University; Olof Widlund and Jacob T. Schwartz; *A Laboratory for Advanced Computational and Software Techniques*; (DCR-8320085); $511,000; 12 mos.

This project will aid in the establishment of a major experimental facility for research, centering on robotics as a focus, uniting key and demanding problems in computer science and engineering and in applied mathematics. The robotics research draws upon and ties together three other areas of current research: software technology, super-speed parallel computation, and scientific computation/numerical analysis.

Research in robotics and computer vision is evolving from initial theoretical studies on collision-free paths, motion planning, hybrid position/force robot control and various problems in which geometry and force both play roles. A basic robotics and vision laboratory has been developed and cooperation with other robotics activities at Cornell and IBM has been established. Research enabled by this project will concentrate on servo control and planning algorithms, techniques for processing robot data, and design and development of languages for robotics.

State University of New York - Stony Brook; Jack Heller: *A Data-Oriented Network System*; (DCR-8319966); $580,000; 12 mos.

This project will aid in the establishment of a major experimental facility for research on operating systems, programming environments, programming languages, and databases. The design and construction of a data-oriented network system within a network environment as a research facility will be a large part of the project. The data-oriented network system will exploit the conceptual simplicity of the relational data model to benefit end-users and developers of applications systems by providing a high-level, uniform view of data. The system's facilities will include a relational operating system environment, a relational editor, support for logic programming, and a software development environment. Applications of this facility to graphics, VLSI, office automation, and natural language processing are planned. The network software will provide location-transparent access to distributed relations and operating system support for distributed programs. The system will be a testbed for research into such issues as resource allocation, concurrency control, and multicast communication.

University of Rochester; Jerome A. Feldman; *A Testbed for the Study of Parallel Computation*; (DCR-8320136); $795,000; 12 mos.

This project will aid in the establishment of a major experimental facility for research on parallel computation centered on a 128-processor multicomputer. The goal is to develop a comprehensive computing laboratory to support projects on massively parallel models, computer vision, knowledge representation, natural language understanding, computer architectures, and systems for parallel computation.

A tightly coupled, parallel machine will be purchased and integrated into the Rochester research network. With a large machine having adequate reliability and software support, it will be possible to experiment with truly parallel algorithms. Current research at Rochester is a particularly good base for pursuing this line of inquiry. The researchers have been studying computational models which involve millions of active processing elements. Although these models are not yet practical, they can be effectively simulated on machines with highly parallel structures. The motivation for this research comes from a continuing interest in artificial intelligence, particularly in computer vision. Currently, research in natural language and knowledge representation does not appear so naturally parallel as vision research but preliminary results with distributed systems show promise that algorithms for tightly coupled processors could be found. This environment will be an excellent testbed for systems and architecture research and the knowledge gained in a truly parallel environment will impact on these research areas.

Duke University; Donald Rose; *VLSI Computing Structures, Design Methods, and Interactive Computer Graphics*; (DCR-8309911 A02); $308,000; 12 mos.

This award will aid in the establishment of a major experimental facility to be shared between researchers at Duke University and the University of North Carolina at Chapel Hill. The facility will support a highly integrated program of research centering on VLSI (very-large-scale-integrated) computing structures and design methods and on interactive computer graphics. The project will also be closely tied to the state-funded Microelectronics Center of North Carolina.

The facility will include a variety of computing resources such as workstations, graphics equipment, and general purpose computers. The Microsystems Laboratory at the University of North Carolina, used heavily by the Department of Computer Science at both Duke University and the University of North Carolina-Chapel Hill for design and testing of VLSI components and systems, will be expanded.

The research will be conducted at three levels in a hierarchy beginning with fundamental VLSI design support and tools, through specialized and general purpose VLSI-based architectures, to the integration of new architectures with software in an interactive computer graphics environment. Each of these areas is closely related. Research in graphics, software engineering, man-machine interfaces, and simulation will provide new ideas for the development of VLSI tools and computer-aided design (CAD) systems. The VLSI CAD systems will support research on new highly-parallel architectures and the new architectures for computer graphics.

University of North Carolina - Chapel Hill; Frederick P. Brooks *VLSI Computing Structures. Design Methods. and Interactive Computer Graphics*; (DCR-8219306 A01); $770,000; 12 mos.

This award will aid in the establishment of a major experimental facility to be shared between researchers

at Duke University and the University of North Carolina at Chapel Hill. The facility will support a highly integrated program of research centering on VLSI (very-large-scale-integrated) computing structures and design methods and on interactive computer graphics. The project will also be closely tied to the state-funded Microelectronics Center of North Carolina.

The facility will include a variety of computing resources such as workstations, graphics equipment, and general purpose computers. The Microsystems Laboratory at the University of North Carolina, used heavily by the Departments of Computer Science at both Duke University and the University of North Carolina-Chapel Hill for design and testing of VLSI components and systems, will be expanded.

The research will be conducted at three levels in a hierarchy beginning with fundamental VLSI design support and tools, through specialized and general purpose VLSI-based architectures, to the integration of new architectures with software in an interactive computer graphics environment. Each of these areas is closely related. Research in graphics, software engineering, man-machine interfaces, and simulation will provide new ideas for the development of VLSI tools and computer-aided design (CAD) systems. The VLSI CAD systems will support research on new highly-parallel architectures and the new architectures for computer graphics.

University of Pennsylvania; Aravind K. Joshi, Ruzena Bajcsy, and Peter Buneman; *Modeling Interactive Processes: Flexible Communication with Knowledge Bases. Computer Interaction in Three Dimensions*; (DCR-8219196 A01); $1,140,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support two main areas of research: flexible communication with knowledge bases and computer interaction in three dimensions. The two areas share several technical aspects and are united by the general theme of modeling an interactive process, whether it is a person interacting with a machine, or a machine with its environment. The latter is relevant to robotics.

Research in these topics involves many established areas in computer science: natural language interfaces to knowledge bases; database systems and methods for interacting with heterogeneous, distributed databases; graphic interfaces for databases; three-dimensional object recognition; representation of visual and other sensor information, in particular, tactile information; integration of sensory information from different modalities; architecture of special purpose machines; movement description and simulation; and a number of related areas.

The facility will include a variety of computing resources including general and special purpose devices. This equipment and the associated personnel support, together with the supporting environment currently existing, will enable the investigators to carry out these long-range research programs of technological and scientific significance, and will also provide a superb environment for training in experimental computer science.

Brown University; Robert Sedgewick, Andries van Dam, Thomas Charniak, Thomas W. Doeppner, Paris C. Kannellakis, Steven P. Reiss, John E. Savage, Jeffrey S. Vitter, and Peter Wegner; *An Integrated Environment for Research in Computer Science*; (DCR-8121806 A02); $444,873; 12 mos.

This award will aid in the establishment of a major experimental facility to support a program of research in computer science. The facility will include a local-area computer network containing a powerful mini computer, a large number of personal computers, and special purpose peripheral equipment.

The investigators will develop a computing environment for the effective use of this network of powerful personal computers. At the core of this environment will be a powerful set of graphics tools. These tools will allow each researcher to customize a graphics-oriented user interface to a specific research project. A database management framework and system will underlie and support the graphics tools.

Research in graphics, databases and operating systems will be key to developing a stable and useful system. The investigators will provide the graphics tools to the larger research community for assistance in guiding refinements and revisions. Other research by the investigators in algorithms, artificial intelligence, compilers, document preparation, learning environments and other areas has made good use of existing tools. The interaction between this research and that connected with the facility itself will also lead to improvements and new research directions.

University of Texas - Austin; Alfred G. Dale, James C. Browne, Donald Good, Elaine Rich, and Abraham Silberschatz; *An Experimental Computing Facility for the Design and Analysis of Reliable High-Performance Distributed Systems*; (DCR-8122039 A02); $911,000; 12 mos.

This award will aid in the establishment of a major experimental facility that will incorporate mid-scale computers with large virtual addressing capabilities and powerful individual workstations linked over a high-speed local network. The facility, together with a professional support staff, will provide an environment for the development and experimental use of software for mechanical program and theorem proving, performance measurement and distributed computing. This integrated system design and analysis facility (SDAF) will:

1. Permit the consolidation of methods for mechanical program verification and performance analysis into an integrated approach to the analysis of distributed systems,
2. Support extensive experimental validation of formal analysis techniques,
3. Support novel applications of formal analysis methods to intelligent user-system interfaces and VLSI design.

William Marsh Rice University; Kenneth W. Kennedy, Robert Cartwright, John Dennis, and J. Robert Jump; *An Experimental Computer Network to Support Numerical Computation*; (DCR-8121884 A02); $500,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support a program of research on advanced distributed computing systems and numerical computing. The Rice Numerical Network, or R .n -, will consist of approximately 24 single-user numerical machines equipped with high-resolution bit-mapped screens, a 32-bit central processor, and vector floating point hardware. It will also include several specialized server nodes and various peripheral devices including a gateway to the CSNET communications network linking the nation's major computer science research centers.

The new facility will support a coherent research program in software systems, computer architecture, and quality numerical software, directed at creating a modern interactive environment for numerical computation.

University of Utah; Richard F. Riesenfeld; *A Laboratory for Computer-Aided Design*; (DCR-8121750 A02 & A03); $634,931; 12 mos. (Joint support with the Computer Systems Design Program - Total Grant $649,931).

This award will aid in the establishment of a major experimental facility to support a program of research on Computer-Aided Design (CAD) Systems. The facility will include powerful design workstations based on personal graphics-oriented machines. Included also will be a necessary expansion of general computing facilities to accommodate the expanded research effort. Several specialized 3-D output devices will provide a basis for testing the CAD techniques and tools developed as part of the research.

There are two main thrusts to the current efforts: Computer aided 3-D shape design and Very Large Scale Integration (VLSI) design. The ALPHA-1 development project, at the heart of the first research thrust, is capable of supporting a wide range of CAD activities. Moving ALPHA-1 to a distributed environment will require research in operating systems and in workstation support and design. The investigators will also work on developing true 3-D display of surfaces and on Numerical-Control Milling Algorithms. Research in portability and software tools will lead to an enhanced software development environment; in database systems it will lead to an efficient uniform means for accessing, modifying, and sharing design data; and in multiprocessor and network operations it will allow the various tasks in the systems to be partitioned and run on processors best suited to their particular requirements. In integrated circuit design, the main emphasis will be on simulation systems improvement, on mapping techniques from high-level representations directly to design, on graphical representations of VLSI designs, and on the testing and diagnosis of integrated circuits.

University of Utah; Richard F. Riesenfeld and Lee A. Hollaar; *1984 NSF CER Conference, University of Utah, Salt Lake City, Utah, February 23-34, 1984 (Computer Research)*; (DCR-8406682); $38,638; 6 mos.

The 1984 NSF Coordinated Experimental Research Conference will bring together two or more representatives from each of the fourteen CER grantees along with guests from academia, industry, and government. The purpose of the meeting is to review the progress of computer research at each of the CER sites. The conference will be organized with general and parallel sessions devoted to both research and research management issues. It is planned that each CER grantee will make two or three research presentations and participate in panels and forums on problems facing academic experimental computer research.

University of Washington; Jerre D. Noe, Hellmut Golde, Edward D. Lazowska, and G. T. Almes; *A Functionally Integrated Environment for Distributed Computation*; (DCR-8004111 A04); $648,000; 12 mos.

This award will aid in the establishment of a major experimental research facility at the University of Washington. The investigators will design and build a powerful new kind of computing environment to facilitate communication and cooperation among researchers in the Department of Computer Science. It will be based on a "building sized" highly-parallel computer consisting of a functionally integrated collection of physically distributed personal computers cooperating through a local network. The primary focus of the research will be the development and exploration of programming methodologies, adapted to a distributed concurrent environment, which give each user access to the power and diversity of the entire system.

University of Wisconsin - Madison; Lawrence H. Landweber, Robert P. Cook, David J. Witt, and Raphael Finkel; *A Software-*

Partionable Multicomputer: A Testbed for Distributed Processing; (DCR-8105904 A04); $1,020,000; 12 mos.

This award will aid in the establishment of a major experimental facility to support research in databases, distributed processing, operating systems, multicomputer architecture, and numerical analysis. The investigators will implement and study a computer architecture that will exploit recent advances in computer and communication technology. The facility will include a multicomputer, a small number of high-performance personal computers, and a support center. The multicomputer will implement the target research environment. The personal computers will serve as prototypes for program development research projects. The support center computers will provide program preparation, simulation, file, and device services.

The multicomputer will be a research tool and a research object for a variety of projects. The implementation of the architecture and kernelized operating system will itself be a major effort. Once a prototype system is constructed, it will allow the multicomputer to be logically partitioned for independent simultaneous use by different researchers. A large number of projects can then be supported to examine research issues in the respective areas.

Computer Science Research Network (CSNET)

University Corporation for Atmospheric Research; Leonard C. Romney; *Computer Science Research Network Host*; (DCR-8304539 A01); $700,000; 12 mos.

The Computer Science research NETwork, CSNET, has been under development by contractors to the NSF for several years. CSNET is a logical network comprised of several physical computer communications networks including ARPANET, Phonenet, and X.25-based commercial common carriers such as GTE-Telenet. The objective of CSNET is to make available uniform network services to the entire computer research community with network costs tied to performance and level of service. Among the services offered are electronic mail, file transfer, and virtual terminal access (computer-controlled remote access). The project has concentrated on the development of software implementations of network protocols and on the expansion of Phonenet capabilities. Phonenet is a telephone-based mail relay system which provides electronic mail and file transfer services. The CSNET Project will be managed by the University Corporation for Atmospheric Research (UCAR), who will house and support the CSNET Executive Director and the CSNET Executive Committee. Technical operations and the Communications and Information Center (CSNET-CIC) will be managed by Bolt, Beranek, and Newman under a subcontract. Substantial support for this project will be provided through user fees and dues.

Bolt, Beranek, and Newman, Inc.; Richard D. Edmiston; *CSNET Coordination and Information Center*; (DCR-8202846 A03); $35,463; 6 mos.

This project is part of a cooperative effort among computer scientists to establish a computer-based communication network for computer science research groups in universities, industry, and government. A logical network spanning several existing computer networks, including ARPANET, public packet networks, and Phonenet, a telephone-based relay system, will be established. Communication services to be provided initially include message services, file transfer, and access to remote systems.

Bolt, Beranek and Newman, Inc. will assist the entire project by establishing a Coordination and Information Center (CIC). The CIC will act as the organization responsible for continuing network management with policies set by a Board of Directors chosen from among leaders in the computer science research community. Other functions to be provided by the CIC include bulk purchase of communication services, network accounting and billing, and dissemination of information on network access and services.

Objectives of CSNET are to help computer science research flourish through resource sharing and to stimulate new classes of research activities. The CIC is expected to be the focal point for continuing CSNET development and service so as to unite computer science researchers in the United States.

Defense Advanced Research Projects Agency; Robert E. Kahn; *ARPANET Support for NSF Computer Networking Activities*; (ASC-8116606 A04), $26,048; 12 mos. (Joint support with the Office of Advanced Scientific Computing Total Grant $411,046).

This is an interagency transfer of funds to cover the costs of providing ARPANET access to Cornell University, Purdue University, the University of Delaware, the University of Washington, and the University of Wisconsin-Madison. A goal of the activity is to support continuing development of the computer science research network, CSNET, as well as to provide more general access for scientific research under the NSF Office of Advanced Scientic Computing.

Presidential Young Investigator Awards

Stanford University; David R. Cheriton; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352048); $50,000; 12 mos.

Communication is the key component of distributed and parallel computer systems, being the basic facility that allows the multiple components of the systems to cooperate in sharing information and hardware resources.

This project will study:

1. Use of communication in highly parallel and distributed computations with particular interest in the benefits of "group" communication paradigms, such as those supported at the hardware level of broadcast networks.
2. The design of communication primitives for parallel computations. Starting from the V kernel IPC design, it is hypothesized that some modest but crucial changes will be required to meet the needs of parallel program structures.
3. The efficient implementation of interprocess communication both on a multiprocessor machine and across a local network joining multiple processors.

The feasibility of a special-purpose IPC chip for this purpose as well as the design of intelligent network interfaces for local network communication will be investigated.

Stanford University; John L. Hennessy; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351269); $24,951; 12 mos.

It is anticipated that new computer architectures will use instruction sets that decompose further than existing architectures. Such computers will require improved optimization techniques in compilation so as to produce efficient programs. This project will investigate techniques that can shift computations from run-time to compile-time and from inner loops to outer loops.

These same optimization techniques can be applied to the efficient compilation of applicative programming languages. The efficiency of such languages can be improved by a two part approach, program transformations to eliminate redundant computations and application of optimization techniques to programs during compilation.

Stanford University; Ernst W. Mayr; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351757); $55,000; 12 mos.

The following topics are studied:

1. Paradigms for the design of good parallel algorithms for a variety of parallel architectures,
2. The inherent parallel complexity of algorithmic problems,
3. Tradeoffs between parallel time and communication complexity,
4. Automatic or semi-automatic program transformations for computer aided design of efficient parallel architectures,
5. Heuristics for general parallel search problems.

Stanford University; Brian K. Reid; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352015); $24,963; 12 mos.

This project is directed toward the use of programming language techniques to solve problems in manufacturing automation. The long term objective is a language in which engineers can specify manufacturing processes. Such specifications should be archivable and portable between properly automated factories. The available theoretical background for the intellectual support of such a language is limited. Specific research areas are nondeterministic semantics and mechanisms for representing non-hierarchical structures.

Johns Hopkins University; Joseph O'Rourke; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351468); $62,500; 12 mos.

Research is progressing along five lines:

1. Algorithms for approximating (by e.g. circumscribing) polygons and polyhedra,
2. Coordinating the movement of several objects in a cluttered environment,
3. Terrain navigation on an undulating surface,
4. Computational geometry on curved manifolds; and
5. The shape of non-inferior regions in multi-objective programs. A monograph is being written on "Art Gallery Theorems and Algorithms."

University of Maryland - College Park; Dana S. Nau; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351463); $62,500; 12 mos.

This research is concerned with several topics involving reasoning and problem solving methods in the field of Artificial Intelligence. In particular it will:

1. Explore an approach to knowledge-based problem solving, which is based on a generalization of the set covering problem. This approach has potential application in the development of expert systems for diagnostic problem solving and other areas.
2. Investigate the applicability of expert computer system techniques to a number of automated manufacturing activities including process selection, stock selection, error classification and diagnosis, and the kinds of spatial reasoning that are required for manufacturing planning activities.
3. Investigate game tree searching. The topics to be investigated are the existence of pathology in game trees, the factors governing why pathology occurs, and various alternatives to minimaxing which may avoid the occurence of pathology.

Massachusetts Institute of Technology; Frank T. Leighton; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352349); $25,000; 12 mos.

Problems involving very large scale integration (VLSI) are studied. Two specific topics are investigated:

1. Techniques for solving design-related problems, and
2. Strategies for efficiently utilizing systems of microprocessors.

Problems investigated include: optimal layouts for the shuffle-exchange graph layout theory, three-dimensional VLSI, wafer-scale integration of systolic arrays, channel routing, methodologies for designing fault-free processor arrays, algorithms and networks for parallel computation, and gate array routing.

Rutgers University; Tom M. Mitchell; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351523); $35,000; 12 mos.

This research project is concerned with two areas of Artificial Intelligence: knowledge-based expert systems and machine learning. One research thrust focusses on developing a knowledge-based consultant system to aid in the design of digital VLSI circuits. This research involves extending current methods for developing rule-based expert systems, by applying them to design tasks, and by integrating rule-based methods with algorithmic methods for dealing with different aspects of the design task. The second research area, machine learning, has focussed on developing the LEX system which learns problem solving heuristics by generating and solving practice problems. This research has led to development of a domain-independent technique (goal-directed learning) for generalizing from examples. Research on machine learning over the coming year will include an attempt to apply these techniques to the task of learning rules for VLSI design.

Columbia University; John R. Kender; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351352); $25,000; 12 mos.

The principal goal will be to create and examine representations and algorithms for the detection and confirmation of visual surfaces by a mobile camera. The foundation of the work will be the inversion of the mathematics of perspective projection, as applied to modelable textures. Such "shape from texture" work will be pursued experimentally by mounting a camera on a robot arm, digitizing and processing the images, and transforming the two-dimensional data into inferences about local and global three-dimensional surface properties.

Among the investigations to be included are the following:

1. Optimal algorithms for edge detection based on psychophysical observations (Laplacian of Gaussian operations),
2. Cooperation of various low-level visual operations to arrive at an intelligent consensus of surface properties,
3. Inference mechanisms using both optimal interpolatory techniques and experimental control structures (simulated annealing), and lastly
4. Investigations of the applicability of these algorithms to special purpose multiprocessors.

University of Rochester; James F. Allen; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351665); $25,000; 12 mos.

Natural language communication is enabled only by the conversants' rich body of knowledge about the topic being discussed, knowledge about the conventions of language itself, and specific knowledge about the other conversants. This project is investigating some basic issues in knowledge representation that are crucial to understanding how communication occurs. In particular, formal models of a naive person's knowledge about actions, space, time, and how they interact in plans are studied. An important subclass of actions are communicative actions.

The results of the work will be published in the artificial intelligence literature. In addition, as the formal models are developed, they will be incorporated as

appropriate in our ongoing projects to build systems to model natural language communication in dialogues, and problem solving in a dynamic, changing world.

Carnegie-Mellon University; Merrick L. Furst; *Presidential Young Investigator Award (Computer Research)*: (DCR-8352081); $62,500; 12 mos.

Four areas are studied:

1. Computational experiments to discover applications on improvements of the shortest vector algorithms of Lenstra, et.al.,
2. Upper and lower bound proofs for multi-party communication protocols,
3. Semantics-directed code generation,
4. Computational aspects of group theory.

University of Washington; Martin Tompa; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352093); $30,000; 12 mos.

This research involves the study of the computational complexity of various algorithms.

The following topics will be explored:

1. Probabilistic algorithms,
2. The parallel complexity of algebraic & number-theoretic problems,
3. Complexity issues in cryptosystem security,
4. Properties of 2-person pebbling, a combinatorial model of parallel computation, and
5. Algebraic simplification.

University of Washington; John Zahorjan; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352098); $30,000; 12 mos.

The purpose of the work proposed here is to investigate the feasibility of, and mechanisms and policies to support load sharing. The first step involves evaluating the conflicting influences of response time improvement achieved by balancing the load and response time degradation caused by task migration overhead. Because the questions to be answered are broad and the details poorly defined, this problem is best addressed by modeling. The goal here is to determine system characteristics (e.g., workload profile, communication channel speed, processor speeds) necessary for load balancing to be a reasonable approach. Once this has been done, implementation of specific mechanisms and policies will be used to find more detailed answers to the more detailed questions that can be posed. The goal here is to evaluate the relative performance of techniques for deciding when to move tasks, which tasks to move, and where to move them.

Small Business Innovation Research

Computer Corporation of America; Umeshwar Dayal; *Knowledge-Oriented Database Management*; (DCR-8360576); $35,000; 6 mos.

Existing database managment systems (DBMSs) can efficiently handle only formatted alphanumeric data. However, many applications (e.g., office automation, industrial automation, computer-aided design, military command and control, and intelligence) also require access to other types of information such as text, maps, diagrams, photographs, images and signals which cannot be efficiently stored and accessed effectively by current DBMSs. As a result, applications suffer unacceptable errors and delays while information is manually compiled from a variety of information sources. Another shortcoming of existing DBMSs is that they can store only limited kinds of semantic knowledge (viz., object types, relationships, and simple constraints) about data in the database. All other knowledge must be built into the application programs. However, this makes it extremely expensive to develop and change the programs, especially as knowledge evolves. The objective of this effort is to develop an advanced DBMS with facilities to input, store, retrieve, and output multiple information types and facilities for declaring and manipulating general knowledge. The feasibility of this challenging objective rests on a key technical insight: it is necessary to enhance existing data models with (a) dimensional (space and time) aspects, which are the common characteristic of non-record information; and (b) recursive predicates and queries which will provide deductive capabilities equivalent to those of many knowledge-based systems. The goal of Phase I is to substantiate this insight, and, if successful, to develop plans for a new knowledge-oriented DBMS architecture.

C. Abaci, Inc.; Edward L. Battiste; *Scientific Software for Micro-Computer Users: Quality Control*; (DCR-8313661); $45,585; 6 mos.

The use of micro-computers is migrating from specific to general purpose. Hardware and software technology are evolving rapidly. No significant software effort is taking place so as to achieve the goals of (1) providing robust software kernels and problem solving environments for scientific users and (2) applying pressure toward the acceptance of standards for languages and software which will allow positive evolution of usage of micro-computers. The advances in software of the last decade are not transferring to micro-computer users.

The Phase I study investigated the feasibility of producing scientific software adequate for use by the (new) micro-computing audience. Problems were isolated which increase difficulty in, and cost of, production. Many will vanish with technology evolution. In Phase II the major problem encountered, quality control in a multi-computer, multi-language environment set with limited external storage facilities, will be investigated. Embedded in that problem is a fundamental arithmetic problem. Most computer releases present incorrect arithmetic and elementary function facilities, to the extent that scientific use may be dangerous. Techniques which isolate such errors will be studied. The quality control effort will utilize a host-target computer concept in a local network. This can lead to non-local network implementations which will serve the dispersed using audience with robust scientific software.

Information Research Associates; Douglas M. Neuse; *Graphical Programming for Simulation Models of Computer Systems*; (DCR-8360779); $35,000; 6 mos.

This project will explore the feasibility of developing computer system simulation programs directly from graphical representations of software systems and computer systems resource/device configurations. The concept is to extend and utilize a set of existing graphical and declarative representations of computer system execution. This set appears to be a nearly adequate basis of representation, and includes: a) the execution-graph model of software system execution behavior proposed by Smith and Browne; b) a queueing network representation of a computer system that includes passive resources; c) information processing graphs, which represent the flow of work between devices on a computer system; d) the PAWS simulation language, in which the execution behavior of given workload elements on a given resource configuration is specified in an entirely declarative as opposed to procedural fashion. These representations appear to be sufficiently similar and to have sufficiently explicitly defined relationships so that automation of the transformation between the representations can be accomplished by standard compilation translation techniques. The research will be used to establish data base schema representations of execution graphs, queueing network models, information processing graphs and PAWS programs and to define transformations between these representations that will allow automated progression from the graphical specifications of the hardware and software systems to a FORTRAN simulation program.