A Proposal to the
National Science Foundation

MINT: A MULTIPURPOSE INTEGRATED NETWORK TO SUPPORT EXPERIMENTAL COMPUTER SCIENCE RESEARCH

Principal Investigator: Edward M. Riseman

Participating Faculty:

Michael A. Arbib
Andrew G. Barto
David A. Carlson (ECE)
Lori A. Clarke
W. Bruce Croft
Howard Elliot (ECE)
Caxton C. Foster
Robert M. Graham
Allen R. Hanson
Francis S. Hill (ECE)
Walter H. Kohler (ECE)
Wendy G. Lehnert
Victor R. Lesser
David D. McDonald
Robert N. Moll
Krithivasan Ramamritham
Edwina L. Rissland
D. Nico Spinelli
John A. Stankovic (ECE)
David W. Stemple
Harold S. Stone (ECE)
Donald F. Towsley (ECE)
Jack C. Wileden
Steven J. Zeil

(ECE = Electrical and Computer Engineering Department)

PROPOSAL TO THE NATIONAL SCIENCE FOUNDATION
Cover Page

FOR CONSIDERATION BY NSF ORGANIZATIONAL UNIT (Indicate the most specific unit known, i.e. program, division, etc.)	IS THIS PROPOSAL BEING SUBMITTED TO ANOTHER FEDERAL AGENCY? Yes ___ No _x_ : IF YES, LIST ACRONYM(S):
The Coordinated Experimental Research Prog. Computer Science Section	

PROGRAM ANNOUNCEMENT/SOLICITATION NO.:	CLOSING DATE (IF ANY):

NAME OF SUBMITTING ORGANIZATION TO WHICH AWARD SHOULD BE MADE (INCLUDE BRANCH/CAMPUS/OTHER COMPONENTS)

University of Massachusetts/Amherst

ADDRESS OF ORGANIZATION (INCLUDE ZIP CODE)

University of Massachusetts
Amherst, Massachusetts 01003

TITLE OF PROPOSED PROJECT

MINT: A MULTIPURPOSE INTEGRATED NETWORK TO SUPPORT EXPERIMENTAL COMPUTER SCIENCE RESEARCH

REQUESTED AMOUNT	PROPOSED DURATION	DESIRED STARTING DATE
$5,197,337	5 years	June 1, 1983

PI/PD DEPARTMENT	PI/PD ORGANIZATION	PI/PD PHONE NO. AREA CODE:
Computer & Information Science	University of Massachusetts	(413) 545-2746

PI/PD NAME	SOCIAL SECURITY NO.*	DATE OF HIGHEST DEGREE ACHIEVED	MALE*	FEMALE*
Edward M. Riseman	113-34-3547	1969, Ph.D.	x	
ADDITIONAL PI/PD & SIGNATURE				
ADDITIONAL PI/PD & SIGNATURE				
ADDITIONAL PI/PD & SIGNATURE				
ADDITIONAL PI/PD & SIGNATURE				

FOR RENEWAL OR CONTINUING AWARD REQUEST, LIST PREVIOUS AWARD NO.:	IF SUBMITTING ORGANIZATION IS A SMALL BUSINESS CONCERN, CHECK HERE ☐ (See CFR Title 13, Part 121 for Definitions)

* Submission of SSN and other personal data is voluntary and will not affect the organization's eligibility for an award. However, they are an integral part of the NSF information system and assist in processing proposals. SSN solicited under NSF Act of 1950, as amended.

CHECK APPROPRIATE BOX(ES) IF THIS PROPOSAL INCLUDES ANY OF THE ITEMS LISTED BELOW:

☐ Animal Welfare	☐ Human Subjects	☐ National Environmental Policy Act
☐ Endangered Species	☐ Marine Mammal Protection	☐ Research Involving Recombinant DNA Molecules
☐ Historical Sites	☐ Pollution Control	☐ Proprietary and Privileged Information

	PRINCIPAL INVESTIGATOR/ PROJECT DIRECTOR	AUTHORIZED ORGANIZATIONAL REP.	OTHER ENDORSEMENT (optional)
NAME	Edward M. Riseman	Lee D. Beatty	Frederick W. Byron, Jr.
SIGNATURE	Edward M. Riseman	Lee D. Beatty	F. W. Byron, Jr.
TITLE	Professor & Chairman	Director, Office of Grant & Contract Administ.	Dean, Faculty of Natural Sciences & Mathematics
DATE	9/14/82	9/22/82	
Telephone No. Area Code: 413	545-2746	545-0698	545-2627

(5/81)

NOTICE OF RESEARCH PROJECT
SCIENCE INFORMATION EXCHANGE
SMITHSONIAN INSTITUTION
NATIONAL SCIENCE FOUNDATION

PROJECT SUMMARY

SIE PROJECT NO.
NSF AWARD NO.

FOR NSF USE ONLY

DIRECTORATE/DIVISION	PROGRAM OR SECTION	PROPOSAL NO.	F.Y.

NAME OF INSTITUTION (INCLUDE BRANCH/CAMPUS AND SCHOOL OR DIVISION)

University of Massachusetts/Amherst
College of Arts and Sciences

ADDRESS (INCLUDE DEPARTMENT)

Computer and Information Science Department
Graduate Research Center
University of Massachusetts
Amherst, Massachusetts 01003

PRINCIPAL INVESTIGATOR(S)

Edward M. Riseman

TITLE OF PROJECT

MINT: A MULTIPURPOSE INTEGRATED NETWORK to Support Experimental Computer Science Research

TECHNICAL ABSTRACT (LIMIT TO 22 PICA OR 18 ELITE TYPEWRITTEN LINES)

Faculty and students in the COINS Department, in conjunction with colleagues from ECE, are involved in a wide variety of diverse and dynamic research projects. This proposal requests support for an ambitious program comprising coordinated experimental research efforts in:

- Distributed Systems
- Information Systems
- Software Development
- Vision and Robotics.

These four efforts have several striking commonalities, most notably an emphasis on experimentation, based upon the construction and evaluation of large and computationally intensive experimental systems, and a significant interest in distributed computation and graphics.

The Department's present computing facility is inadequate to support our current research program, not to mention the anticipated expansion of our research activities. We therefore propose to build a Multipurpose Integrated NeTwork (MINT), consolidating our current Research Computing Facility and extending it through the addition of substantial new resources. Our goal is to provide additional computing power, increasing significantly the speed with which experimental systems can be constructed and evaluated and allowing us to extend our investigations, particularly in the area of workstation-based systems in a network context. We propose to accomplish this goal by constructing a local area network that can be tailored to the needs of individual research projects while still facilitating the sharing of both hardware and sofware resources among COINS researchers.

1. Proposal Folder
2. Program Suspense
3. Division of Grants & Contracts
4. Science Information Exchange
5. Principal Investigator
6. Off. of Govt. & Pub. Progs.

ABSTRACT

Faculty and students in the COINS Department, in conjunction with colleagues from ECE, are involved in a wide variety of dynamic research projects. This proposal requests support for an ambitious program comprising coordinated experimental research efforts in:

- Distributed Systems
- Information Systems
- Software Development
- Vision and Robotics

These four efforts have several common characteristics, most notably an emphasis on experimentation, based upon the construction and evaluation of large and computationally intensive experimental systems, and a significant interest in distributed computation and graphics.

The department's present computing facility is inadequate to support our current research program, not to mention the anticipated expansion of our research activities. Moreover, despite emergency funding from the University, we also foresee the danger of having a facility that is severely limited, due to an inadequate programming and support staff, and likely to deteriorate if we are unable to continue to meet the high cost of proper equipment maintenance.

This proposal addresses all of the problems plaguing our research computing facility. We propose to build a Multipurpose Integrated NeTwork (MINT), consolidating our current facility and extending it through the addition of substantial new resources. Our goal is to provide a new, qualitatively different style of computing for our researchers. This will allow us to increase significantly the sophistication and complexity of the experimental systems that can be constructed and evaluated. We will accomplish this goal by constructing a local area network that can be tailored to the needs of individual research projects while still facilitating the sharing of both hardware and software resources among COINS researchers. In addition, we are requesting funding for a MINT administrative and support staff consisting of four full-time professionals and a half-time administrative assistant; in addition the University will contribute funds for three full-time professionals. Finally, we are seeking funds to aid us in dealing with the increasingly large costs of operating and maintaining the Research Computer Facility. The University will contribute more than $1.1 million over the five years, while the total request to NSF for the five years is about $5.2 million.

TABLE OF CONTENTS

I. OVERVIEW OF PROPOSAL

I.1. Research

The Department of Computer and Information Science (COINS) is at the center of a rapidly expanding program of computer science research and graduate instruction at the University of Massachusetts (UMass). In conjunction with colleagues in the computer engineering component of the Electrical and Computer Engineering (ECE) Department, COINS researchers are pursuing a wide-ranging and well-funded research program totalling about 2.5 million dollars annually via 48 research grants, including 23 NSF grants. Our research can be broadly categorized into the areas of Distributed Systems, Information Systems, Software Development, and Vision and Robotics. While the individual research projects differ greatly, they have several striking commonalities. Most notable among these are an emphasis on experimentation, based upon the construction and evaluation of large and computationally intensive experimental systems, and a significant interest in distributed computation and graphics.

This proposal requests support for a strongly evolving research program consisting of four coordinated experimental projects. Groups of faculty in each of the four research areas have recently initiated, or are currently initiating, new, major, coordinated experimental research efforts:

- The Distributed Systems group has recently embarked on a number of major research projects, all aimed at understanding organization and control issues in large distributed processing networks. This work, which is experimental in character, involves the building and analysis of sophisticated and computationally intensive software testbeds.

- The Information Systems group is beginning a major collaborative project with DEC on intelligent user interfaces for profession-based systems, integrating work on text and document retrieval, data base management, natural language generation, example generation and explanation, expert systems, and graphics.

- The Software Development group is undertaking the construction of a prototype environment supporting a feedback-directed approach to software development which, when completed, will provide an integrated collection of tools to aid in all stages of the software development process.

- The Vision and Robotics group has recently launched a major research effort in the area of dynamic scene analysis (motion), and the integration of tactile and visual information in the control of robots, while continuing its long-range research effort in image understanding.

Our research environment has several striking commonalities. First, there is a strong orientation toward experimentation. Researchers in each area are actively constructing, employing and evaluating large-scale experimental computational systems: examples include the VISIONS system (Riseman et al.), the Cooperative Distributed Interpretation Testbed (Lesser et al.), the Distributed Software Testbed (Kohler et al.), the ATTEST test data generation system (Clarke et al.), the MUMBLE natural language generation system (McDonald) and the information retrieval system (Croft et al.).

A second common theme among COINS research areas is a strong interest in distributed computation. For instance, the VISIONS group employs distributed processing as a fundamental processing strategy, projects in Distributed Systems focus on a variety of aspects of distributed computing, work on Software Development tools centers on topics pertaining to the development of distributed software systems, and distributed databases are of major interest to the Information Systems group. Furthermore, researchers in both the Software Development and Information Systems areas are concerned with tailoring their systems to the distributed, workstation-based computing facility that will be common in the near future.

Finally, a significant interest in the applications of sophisticated graphics spans all the COINS research areas. Projects in each of the areas are exploring ways that graphics (especially high resolution color graphics) can enhance the interface between a computational system and its users.

Each of the four project areas has already been successful in obtaining multiple grants through the peer review process to support their research efforts. In light of this, our Coordinated Experimental Research proposal does not request support for one new, unproven project. Instead, we are requesting equipment and personnel to support four on-going, large-scale, well-integrated research projects that build upon our research expertise and experience. Since these projects have already obtained significant funding, but lack an adequate computing facility, we have turned to the Coordinated Experimental Research program as one of the few available sources of the necessary funding for equipment, support personnel and maintenance. Without the requested infusion of support for our experimental computing demands, our research program will be severely jeopardized and an enthusiastic evolving research department will be crippled for lack of resources.

I.2. Problems in our Computational Environment

The department's present computing facility, even with the major upgrades that will occur during this year, is inadequate to support the highly interactive and intensive computing style that is necessary for our experimental research. What is necessary is universal access to the interactive, cycle-intensive style of computing that can now be supported by personal workstations (of varying power) equipped with high-resolution graphics and appropriately sophisticated software and hardware.

Despite temporary emergency funding from the University for an expanded staff and maintenance support, we foresee the danger of having a facility that is severely limited due to an inadequate programming and support staff and deterioration caused by an inability to meet the high cost of proper equipment maintenance. Because of economic problems, the University of Massachusetts budget has been declining in real dollars. The demand for growth in a major program such as ours has caused a variety of obvious problems for the UMass administration. Nevertheless, they have responded with a strong commitment (refer to Section IX) and emergency funds of $100,000 for staff and maintenance for the 1982-83 academic year. However, it is obvious to all of us that they are not in a position to meet our future needs in these areas without a significant infusion of external support.

I.3. The MINT Proposal

To overcome the shortcomings in our research computing environment, we propose to build a Multipurpose Integrated NeTwork (MINT), consolidating our current Research Computing Facility and extending it through the addition of substantial new resources. Our goal is to provide additional general and specialized forms of computing power to create a qualitatively different style of computing for our research community. This will significantly increase the speed with which more complex and sophisticated experimental systems can be constructed and evaluated in the four project areas of this proposal. We propose to accomplish this goal by constructing a local area network of personal and specialized workstations that can be tailored to the needs of individual research projects while simultaneously facilitating the sharing of both hardware and software resources among COINS researchers.

Of equal importance is our request for staffing and maintenance support. We require a staff that is large enough to integrate such an infusion of equipment in an effective manner. We also require significant funding to aid us in alleviating the impact of the high maintenance costs, which are needed to keep our research computing facility operational.

The proposed MINT equipment would be phased in over a five-year period with a total expenditure of $2.35 million for equipment during that period. When completed, MINT would be one or more local area distributed computing networks (based on ethernet) providing the following capabilities to our research groups:

1. Access via personal workstations for all members of our research community. These individual Network Access Workstations (NAWs) would provide modest local processing sufficient to support basic text and software preparation activities with high resolution, bit-mapped raster graphics displays, and access via the ethernet to the shared resources.

2. Highly sophisticated, special-purpose workstation for advanced research applications. A pool of about twenty powerful workstations (including our existing VAX 11/750s) with varying capabilities would be shared among the numerous research projects having intermittent need for such facilities.

3. General purpose timeshared computing services supported by a very powerful time-sharing system (to replace our current two VAX 11/780s). In combination with the NAWs it will provide short bursts of intense interactive computing for users not requiring continuous interaction with specialized workstations. It will also serve users who want to execute large-scale production runs in the background.

4. Shared resources of archival storage and hardcopy output. Our existing disk and tape facilities will be greatly expanded (including a video disk) and accessed via network file servers. We will add a

second higher speed laser printer and a camera system for recording color graphics output.

This proposal also includes funding for a half-time administrative assistant and support staff of four full-time professionals and three graduate student research assistants. Their responsibilities will include supervising the evolving operation of the MINT facility and directing the development of network graphics and other specialized software. The University will share significantly in the cost of the support staff over the five year period by contributing about $1.1 million with a request to NSF for $1.4 million in personnel (includes indirect costs). Finally, we are requesting almost $1 million to help defray the cost of operating and maintaining the MINT facility. As we have stressed above, the importance of the support staff and maintenance components of the proposal cannot be overemphasized.

I.4. Relationship of MINT to the Research Areas

The structure of MINT will satisfy the computational needs of the four individual research areas in the following ways:

1. Distributed Systems -- MINT provides a realistic framework within which to build and evaluate testbed systems for investigating such issues as the performance of concurrency control algorithms and distributed task scheduling algorithms. By supporting the clustering of workstations, it also provides sufficient computing cycles for running large and complex distributed simulations (i.e., Distributed Interpretation Testbed and Neural Network Modelling). This capability makes possible the empirical study of

control issues in large distributed processing networks, a major focus of the work in this area.

2. Information Systems -- The high-powered specialized workstations such as the SYMBOLICS 3600 LISP machine, provide the appropriate level of dedicated computing necessary for building large expert (knowledge-based) systems. Time-sharing machines are simply not appropriate because they do not permit efficient use of memory, do not support custom configurations of massive (e.g., > 15 Mbyte) programs, and cannot guarantee real-time response. The powerful workstations have sufficient capability for real-time interaction with users to allow us to carry out experiments with sophisticated user interfaces such as the combination of a natural language generation system with an intelligent help system. Finally, the NAWs and specialized workstations provide a range of computing and graphic capabilities that can serve as prototypes for the hardware systems that will become the norm for profession-based systems in the near future.

3. Software Development -- Work in this area has been hindered by the inability to explore real-time user interaction with sophisticated analysis, description and testing tools. A powerful workstation would, for example, allow program analysis based on symbolic execution to be done rapidly enough to provide useful feedback to an interactive user. MINT provides a workstation network, with the NAWs providing the type of graphics and window management capabilities that will be the prototype for the software development environment of the future. Thus MINT will permit this group to develop their software development tools in a realistic modern environment.

4. Vision and Robotics -- The computing needs in this research are massive, spanning the whole range of proposed equipment. For instance, vision research requires powerful and specialized processors, such as an array processor and a pipelined image processing system, to analyze the high volume of sensor data. Robotics research will require the acquisition of at least one more robot arm and manipulator, with of course some dedicated computers. There is also the need for real-time interaction with the user via high-resolution color graphics in order to do effective algorithm development. Powerful workstations, such as the LISP machine, are needed for the high-level symbolic processing required in image understanding. The robotics work requires dedicated processing for real-time interaction with the device. In short, significant and qualitative changes in results will be achieved by adding more flexible user access via several specialized workstations and far more computing

cycles.

We believe that the nature of the computing needs of our research community places us in an ideal position to derive maximum benefit from the resources available through the NSF Coordinated Experimental Computer Science Research program. In addition to the direct benefits outlined above, MINT will provide a computing environment in which software sharing and interaction among groups will be heightened. Our proposed MINT facility will permit us to pursue important, experimentally-oriented research topics that will otherwise be closed to us. Furthermore, it will be a flexible facility that will make possible the adaptation of our computing environment to new and emerging research areas.

While it might be possible to acquire most of the resources comprising MINT in a piecemeal fashion through individual, small grants, there are serious drawbacks to such an approach. It is obviously a slower and less efficient way to proceed, requiring us to expend more time and energy in obtaining grants rather than working on our research. Also, such piecemeal development threatens the hardware and software compatibility that is one of our primary concerns, and thus, might tend to fragment our research community. Facilities acquired in a piecemeal fashion by individual projects are less likely to be compatible and therefore less likely to facilitate interactions among researchers. Most importantly, a piecemeal approach would make it much more difficult to obtain the widely shared resources that constitute the bulk of our proposed facility; sufficient funding for maintenance and support staff, for example, is very hard to justify in small grants. It is only because the support staff and maintenance for MINT will be shared among so many productive research projects that these resources -- hardware, maintenance and personnel -- can be cost effective. It is precisely because MINT will support such a wide array of research efforts that it is ideally suited for support from the NSF Coordinated Experimental Computer Science Research program.

In the next three sections of this proposal we describe the current status of the COINS department, elaborate on the new resources we are proposing and the justification of their need, and present a detailed acquisition plan and budget for constructing MINT. Sections V and VI contain summaries of our various research efforts and detailed discussions of individual research projects, respectively. We have not included justification of particular equipment needs in the research summaries and detailed descriptions, but instead discuss the needs for all four research areas in Section III.

II. CURRENT STATUS OF THE COINS DEPARTMENT

Like many of the other leading computer science departments in the United States, COINS has been struggling throughout the last few years to maintain and expand quality research and instructional programs. We have been faced with dramatic enrollment pressures, intense competition (particularly from industry) for advanced graduate students and faculty, and fluctuations in funding due to political uncertainties at both the state and federal levels. Unlike many departments in similar situations, COINS has managed to significantly increase its strength and quality during this period:

1. Faculty and Students: We have successfully competed for high quality faculty and students;

2. Research: The scope and funding of our research programs have been expanded substantially and we currently have internationally recognized programs in the areas of AI, Cybernetics, Distributed Computing, Software Development, and Computer Vision;

3. Computing Facility: A first-rate, although already saturated, research computing facility has been assembled;

4. Industrial Affiliations: Strong affiliations with industry have been developed;

5. University Support: Strong university support for the department has solidified with additional faculty positions, space, and a commitment to aid in the development and maintenance of our laboratory.

II.1. Faculty and Students

With the addition of six new faculty positions over the last two years, all six of which we successfully filled despite an intensely competitive job market, COINS has now grown to nineteen full-time faculty and six research computer scientists. The Electrical and Computer Engineering Department has two new computer engineering faculty, bringing their computer engineering faculty to seven. We are now proposing a Center for Computer Science and Computer Engineering that will bring together the personnel from COINS and those members of the University's Electrical and Computer Engineering (ECE) department who have closely related research interests. The Center would serve to enhance interdepartmental computer research at the University and thus will further enrich the COINS research environment. COINS is also establishing closer ties with the Department of Mathematics and all three departments are working on coordinating course offerings.

There are presently about 125 graduate students in COINS, about 40 of whom have passed their Ph.D. qualifying examinations. This year's new class of 32 graduate students are all being financially supported by the department and continues our trend of recent years of attracting very strong graduate candidates. Our recent Ph.D.s have been very successful in the job market, joining the research laboratories of such organizations as DEC, GE, HP, TI, IBM, and SRI.

The following table outlines our graduate degree production since 1977. Note that the figures of the last two years include a second number representing the number of computer engineers graduating from the ECE Department.

	1977-78	1978-79	1979-80	1980-81	1981-82	1982-83
Total Graduate Enrollment	70	83	96	96	115 19	126 42
M.S. Degrees Granted	20	20	22	26 5	30 5	? 20
Ph.D. Degrees Granted	2	4	4	5 1	4 1	10-12 2

The graduate pipeline is just now approaching steady state (not counting further expansion of our program). We have a rather large number of students who are nearing completion of their Ph.D. degrees. Three Ph.D. degrees have been awarded for September, 1982; six more are expected to complete all requirements by February, 1983; nine more are in an advanced state of Ph.D. thesis research. The ECE Department expects to award 2 Ph.D. degrees in computer engineering and has rapidly increased graduate enrollment which is starting to fill their pipeline. In steady state we expect that about 15 Ph.D. students will be turned out each year based on the current structure of the two departments.

It should be noted here that <u>graduate student researchers account for the majority of the usage of our current research computer facility</u>. Obviously the presence of a first-rate computing facility will allow us to continue to attract outstanding students and make it possible for them to perform their doctoral research in a reasonable way.

II.2. Summary of Research Funding

The department has steadily expanded both the scope of its research program and its external research funding which comes from a diverse set of agencies summarized in Table I. The 26 computer science and computer engineering faculty have 48 grants which are active during the period January 1, 1982 onward. The total funding of these grants is $6.9 million. We selected the 12 month period beginning June 1, 1982 to compute the pro-rated yearly funding rate which comes to $2.57 million per year. Much of the research support is in the form of relatively small grants obtained through the competitive, peer-review process, including 23 NSF grants. We are also receiving substantial research support through 11 grants via our industrial affiliations as noted in section II.3.

A summary of the scope of activity in each of the four research areas of this proposal is contained in Table II. Note that the figures involve research by 18 COINS faculty, 6 Ph.D. COINS research associates, and 7 ECE faculty; there are two more faculty who will be arriving in January, 1983.

Note that some faculty contribute to more than one area and thus the personnel add up to more than 31, but each grant was placed in a single area even when there was overlap so that the research dollars per annum will add up to meaningful statistics. Also note that Table II figures total to somewhat more than $2.57 million per year because the grants in each of the areas were not precisely pro-rated over the same 12 month period because of diffferent funding periods for major grants.

Table I. Tabulation of Grant Agencies
(for Grants Active During 1/1/82 Onward)

Agency		Number
NSF		23
DOD		9
DARPA	4	
ONR	2	
ARI	1	
Army CECOM	1	
AFOSR	1	
NASA		1
NIH		2
Industry		11
Foundations		2
Total		48

Table II. Summary of Research Funding by Area.

Yearly funding is computed only for grants active after 6/1/82.

Research Area	# Faculty	# Research Associates	# Grants Funded (pending)	Total $ Funded (pending)	Approximate $/year Funded (pending)
Distributed Systems	7	1	10 (3)	$1,802K ($1,086K)	$500K ($392K)
Information Systems	8	1	8 (4)	$1,162K ($452K)	$445K ($236K)
Software	6	1	7 3	$431K ($716K)	$210K ($326K)
Vision and Robotics	7	3	13 (3)	$1,935K ($1,539K)	$561K ($536K)
Equipment			5	$1,339K	$904K
Miscellaneous			5	$234K	$101K
TOTAL			48 (14)	$6.90million ($3.81million)	$2.72million* ($1.49million)

*The exact figure precisely pro-rated for all funded grants in the period 6/1/82 - 5/31/83 is $2,565,000.

II.3. Industrial Affiliations

COINS has been developing research relationships with DEC in the areas of industrial robotics, profession-based systems, distributed computing, and graphics. The most significant project involves a $750K donation of equipment and $60K of research support by DEC for intelligent user interfaces for profession-based systems. This is a major cooperative research effort between COINS and DEC, integrating approaches from expert systems, natural language processing, data base systems, and office automation.

The vision and robotics group has an ongoing relationship with DEC in the development and use of tactile sensors and vision algorithms for a variety of algorithms, with significant credits towards equipment discounts. The General Electric Corporation recently donated a prototype robot arm and peripheral equipment valued at $75K. Support from DEC and GE have supplemented an NSF grant and made possible the evolution of a serious robotics effort in COINS. In addition, GE has become interested in developing a research relationship in the development of parallel computer architectures. The vision group also has an ongoing relationship with Tufts New England Medical Center and Massachusetts General Hospital in biomedical image processing; the vision research group has been developing a relationship with ASTEC (American Science and Technology), and received funding for efficient transmission and geometric correction of remote sensing images.

In the last three years, three of our graduate students have been awarded IBM Fellowships and last year COINS received an unsolicited research development grant from IBM acknowledging the excellence of our research activities. We are also expanding our ties with CDC in the areas of APL systems, database query languages, and software development environments. Finally, Data General has just donated an Eagle MV-8000 system and peripherals worth $300K to the COINS and ECE Departments for educational use. This is very important since all of the equipment in our Research Computer Facility is restricted to research purposes. Our continuing and growing industrial affiliations, such as those with DEC, CDC, GE, Data General, and IBM are indicative of the relevance and quality of COINS research activities.

II.4. Institutional Commitment

The computer industry is at the heart of the resurgence of the high technology industry in Massachusetts. The Computer and Information Science (COINS) Department is now receiving the strongest support possible from all levels of the University administration. In the face of extreme economic problems and serious budgetary cutbacks for higher education, resources of all kinds are still being reallocated to computer science and computer engineering. Refer to Section IX for the formal institutional statement.

COINS and the computer engineering component of the Electrical and Computer Engineering (ECE) Department are jointly planning their expansion as a cooperative effort in both education and research. The computer systems component of ECE (7 faculty) now resides with the COINS Department (19 faculty) in the same building. Tentative proposals have been drawn up for a Center for Computer Science and Computer Engineering that will serve as an institutional focus for this cooperation.

Six new faculty were hired in COINS in the last two years. ECE has hired two new computer engineering faculty in the last two years. Our long range plans (five years) project a growth to the range of 30-40 faculty, which combined with 10 post-doctoral or visiting researchers, will form a community of 40-50 Ph. D. researchers.

During the past two years the salaries of COINS faculty have been substantially increased to bring them more into line with comparable salaries in other academic and industrial organizations. At the same time the standard teaching load for COINS faculty members has been reduced to three courses per year in recognition of the time commitment required by their extensive research activities.

Significant amounts of space are being made available to accommodate the ongoing expansion of computer science and engineering faculty and graduate students by moving administrative units out of the building. Approximately 4500 square feet of additional space were provided in the last year and about 12,000 square feet are in the planning stages for the next three years of COINS and ECE department expansion. Renovation of 1500 square feet of laboratory space was completed this summer, and another 1300 square feet are in the planning stages in preparation for funding of this proposal.

During the past four years the University of Massachusetts has provided several matching grants totalling $117,000, which have been combined with major grants from NSF, DARPA, and DEC, to allow the acquisition of approximately $2 million of computer equipment for the COINS research facility. As further evidence of administrative support, the COINS Department has been allocated two new professional positions. One is for the laboratory manager of the Research Computer Facility and the other is for an Administrative Assistant to aid the department in its rapid growth and heavy research commitment.

In addition, $100,000 of temporary funds for the 1982-83 academic year have been provided to COINS for staff and maintenance on an interim basis until the outcome is determined of this year's efforts in acquiring external support. We are using these funds to hire a full-time software development person to complement the lab manager, assistant lab manager, and graduate students. With the funding of the MINT proposal, the temporary commitment of these funds will solidify into a coherent permanent long-term plan of operation and maintenance of our laboratory. The University will additionally provide COINS with a full-time technician as well as half-time teaching release for the faculty supervisor of the RCF. Section IV, the budget, documents a total University and Departmental commitment of more than $1.1 million over the five year duration of the project. Tha University is prepared to fully support COINS and ECE in developing an outstanding research environment in computer science and engineering and to maintain this environment after completion of the CER grant.

III. JUSTIFICIATION OF PROPOSED NEW RESOURCES

MINT takes a balanced approach to building a computing environment for experimental computer science. It combines fifty relatively low-cost personal workstations that provide for low-end computing and access to shared facilities (hence the name of "NAW", for "Network Access Workstation"); a pool of 20 powerful workstations such as the Symbolics 3600 Lispmachine for cycle intensive computing that requires sophisticated interaction and display capabilities, specialized hardware for our Robotics and Vision research, plus a very high speed timesharing system for short bursts of intense computing, production runs, and base-level service. All of this equipment is to be joined into our local area network (an Ethernet that will support multiple networking protocols) to permit access to existing and requested shared resources (e.g., file servers, print servers, dialup lines, etc.).

We do not see the acquisition and integration of these new resources as a major research endeavor in itself. Rather, we view the construction of MINT primarily as a means of providing an appropriate computing facility to support our current and future research interests. To this end, we propose to purchase off-the-shelf components and software whenever possible and to phase in our acquisitions on a schedule that will minimize the disruption of our research activities. The timing of the phased acquisition should also allow us to take advantage of systems software under development at other sites that are now acquiring similar equipment configurations.

The remainder of this section details the new resources we are requesting. It has six subsections, III.1 detailing our current equipment, III.2 describing the inadequacies of the existing facility and justifying the equipment in terms of its benefits to our research, III.3 specifying the proposed equipment and the acquisition schedule, III.4 defining technical issues involved in the development of the facility, and III.5 and III.6 describing our requests for staff support and maintenance, respectively.

III.1. Current Status of Research Computing Facility (RCF)

Our RCF is now being upgraded as a result of equipment acquisitions made possible by a significant grant from DEC ($750K equipment, $60K research support) for a collaborative research relationship on intelligent user interfaces for profession-based systems; matching funds from DARPA, NSF, and UMass were supplied. We will, by the end of the year, have acquired 9 single-user VAX 11/750 workstations. One of the provisions of the DEC grant is that the 11/750s must be used exclusively as single-user systems. These workstations will be integrated with our existing two large VAX 11/780 time-sharing systems via an Ethernet-II local area network (see Figure 1). Our current time-sharing systems running VMS/EUNICE are networked using a 1 megabyte DECNET link and we will at least initially be using a local area network version of DECNET for basic network services.*

Existing Configuration as of December 1982

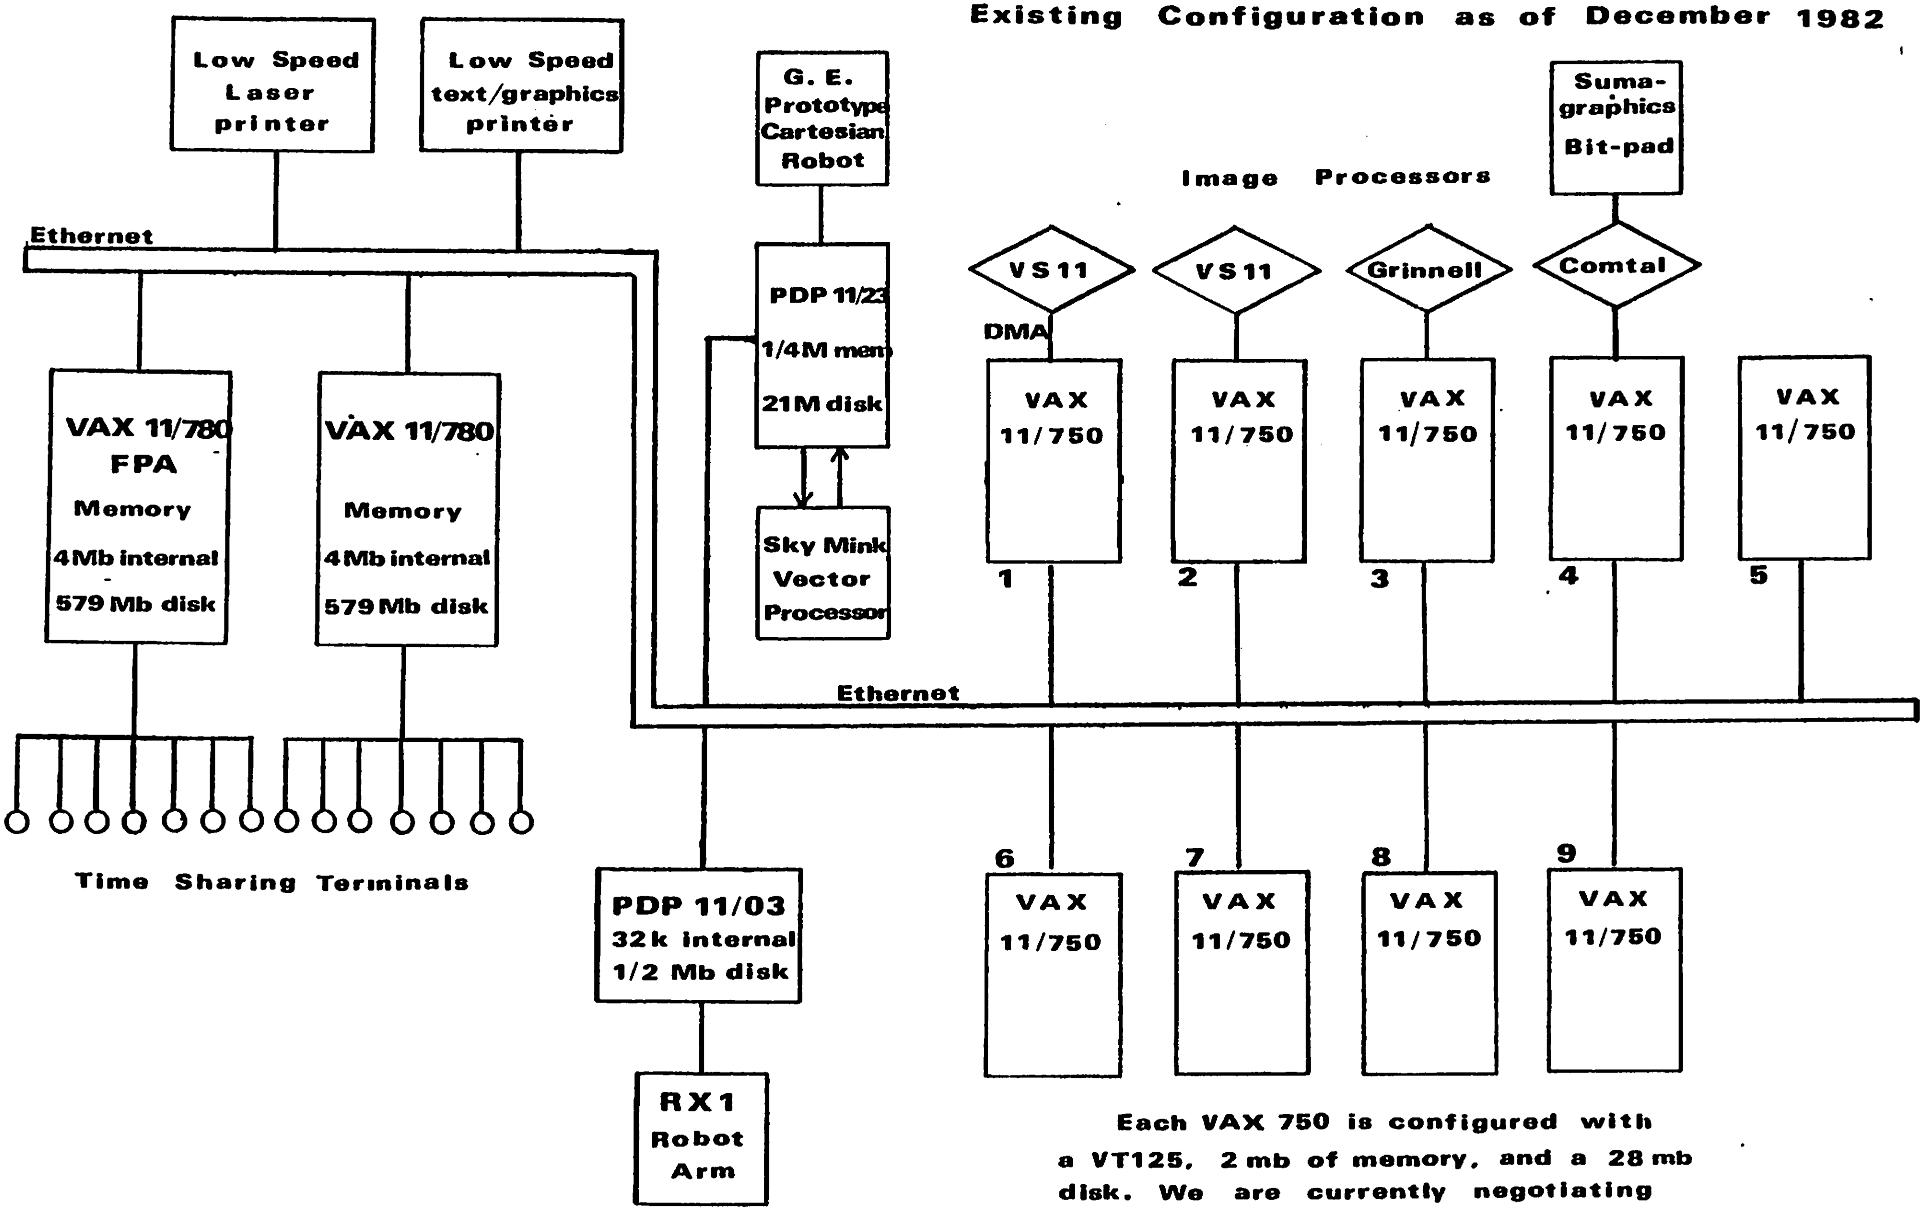

Each VAX 750 is configured with a VT125, 2 mb of memory, and a 28 mb disk. We are currently negotiating for additional mass storage.

Figure 1

The acquisition of the 11/750s is an important upgrade to our computing facility. They represent the first step toward the kind of computing environment that we feel is appropriate for constructing, evaluating, and employing large-scale experimental computing systems. They are a compatible enhancement to our environment that will relieve the immediate pressure on our extremely oversaturated timesharing facilities. However, they are only a short-term solution and do not represent the type of computing environment in terms of computation speed, graphics, and programming support that will be required over the next five years. We expand on the characterization of the VAX 11/750s below, and in particular convey our reasons why they do not totally satisfy our needs.

III.2. Justification

III.2.1. Inadequacies of Existing Facility

Over the last several years we have witnessed a tremendous increase in the number of computing cycles needed in order to effectively build and evaluate our experimental systems. This increased demand has been coupled with a striking change in the style in which researchers interact with their application programs. It is now common for many of our application programs in experimental domains to be

*We are a field test site for this new version of DECNET; this version will permit us to intermix DECNET and non-DECNET protocols on the Ethernet. We specifically are interested in exploring the use of TPC/IP protocols to prepare for integrating non-DEC equipment into the network.

accessed through menus via pointing devices and to display their results through multiple windows and medium resolution color graphics. We believe this need for high-bandwidth, sophisticated interactive graphics and intense computation will become the norm for our program development. In fact this is one of the major directions being pursued by the software development group. This need for a new type of computing environment is becoming increasingly prevalent in academic and industrial research laboratories, as witnessed by the commercial development of sophisticated personal workstations on the model of the SYMBOLICS Lispmachine or the XEROX Dorado.

Our current environment does not support this style of computing, even with the nine new 11/750s. First, there are insufficient access ports that can support high resolution black/white graphics (750 x 1024) or medium resolution (512 x 512 x 8) color graphics; we currently have only four color graphic devices. We also do not achieve interactive graphics because none of our graphic devices can in fact support effective window management. If the common mode of user interaction is through such quality graphics, we estimate a need of at least 50 of these access ports in order to support a user population of over 200 users. Secondly, the VAX 11/750s, which provide less than .75 MIPS, do not supply a sufficient number of computing cycles to effectively build and interact with our larger experimental systems.

There are, however, single-user machines which can meet our need for specialized workstations more effectively than the VAX 11/750. The language-specific design of the Symbolics 3600 Lispmachine, for example, enables it to run at 5MIPS on operations central to lisp such as stack manipulation and function calls; this machine is on the order of 20 times the speed of an 11/750 running LISP. In the 3600 single-user operation can be taken advantage of in optimizing the use of the bit-mapped displays, in the automatic sharing of the entire address space among all processes (which is crucial when interacting systems are being coordinated), and in the simplicity of the interprocess architecture. In addition, the software base for general program development (e.g. coordination between the editor and the debugger, and the sophistication of the window management facilities) is markedly larger and of a higher quality than will be available on the VAX for years to come. Powerful single-user machines specialized to Pascal (and potentially to ADA) are under development as well, notably by Three Rivers, Apollo, and Symbolics.

The need for sophisticated graphics and extremely fast computation are already present in many of our large application programs such as those in natural language processing (McDonald and Lehnert), program analysis and testing (Clarke), distributed problem-solving (Lesser), and high-level vision (Hanson and Riseman). We will briefly discuss the magnitude of the computational requirements of each of these research groups.

The computational problems faced by the information systems, software development, and the distributed processing groups are very similar: (1) large program modules are now being combined into unified systems that must be developed and operated as a coordinated whole, (2) the tasks being performed call for results being available in minutes rather than hours if they are to be of any utility to their users, and (3) the complexity of the information that these systems are to present to their users (not to mention the researchers that develop them) has moved far beyond what can be squeezed onto a conventional 24x80 character terminal.

The information systems group, for example, is currently developing an expert system with a natural language interface on input and output that involves language understanding and language generation components, respectively. These language components are approximately 4 megabytes in size. When these components are coupled with a 6MB applications program (the "expert"), as well as components for composing explanations, key examples, help-tutorials, etc., the resulting expert system can easily grow to at least 12MB; all of this is running on top of 16MB's of LISP and systems' support software. Furthermore, these subsystems are not independent but must now be engineered and debugged as a unit requiring efficient use of time, space, and display resources. It is impossible to develop such sophisticated, and computationally intense programs on a timesharing system; errors that normally should take a few minutes to find would take hours and maybe days.

The software development group is confronted with the same problem in doing symbolic evaluation and test data generation and evaluation. These tools cannot be used in software development today simply because the computation time they require is more than a user is generally willing to pay. Dedicating VAX 11/750's to single users is a step in the right direction, but is still only a step since the 11/750 is unfortunately too slow and does not provide the sophisticated graphics interface that is required to display information allowing exploratory development of complex software system.

Equivalent issues arise in simulation of large distributed processing networks by the distributed processing group. Many important distributed processing issues, involving the performance of large networks, cannot be studied because the current resources severely limit the size of the networks that can be simulated. In addition to the increased computation power and graphics interface, the proposed network will provide a realistic environment for a major focus of the distributed process group's research. Techniques for concurrency control of distributed databases and distributed scheduling can be implemented and empirically evaluated. It also will serve as an ideal experimental setting for investigating distributed system debugging and performance monitoring tools, another focus of this group's work.

The vision and robotics research group has even more stringent requirements for computation. A single static black and white image of reasonable resolution (512x512 pixels) involves 1/4 megabyte of data and a color image of the same resolution requires 3/4 megabyte. Low level relaxation algorithms require the iterative processing of data from a local window about each pixel. This can take several hours of CPU time for each run on a moderately loaded VAX 11/780. Thus, static vision research in the past has pressed our computational environment to the limit in both CPU and storage resources.

Our new research in motion has computational demands of one or two orders of magnitude greater than static image analysis, but has become very important with imminent VLSI hardware breakthroughs. This has led researchers to attempt to construct algorithms on sequential machines despite the extreme difficulties of carrying out meaningful experiments. Motion research requires clips of several seconds at 30 frames per second. Only 10 seconds of B&W video involves processing of 75 megabytes of data. Thus, in these algorithms each of the 75 million pixels may need to undergo iterative local processing in relaxation algorithms, sometimes with storage of intermediate results. This tremendous amount of computing just cannot be accomodated on processors designed for general purpose computation.

The low-level image processing algorithms can be greatly sped up by using specialized processing systems. For example, an image display processor and an array processor will significantly improve the execution time of the low level segmentation algorithms. We recently estimated two orders of magnitude speedup in execution using just the convolution processor in our Comtal image processor, which provides video rate point processing and convolution over 3x3 windows. Without this type of additional hardware, the motion research will only be able to take place with the most judicious conservation of disk and CPU usage.

Although the above observations suggest that the ideal might be to provide each COINS researcher with a sophisticated personal workstation, we feel that such an approach is economically unrealistic. Fortunately, we feel that a good environment can be acquired through a mixture of sophisticated personal workstations and low-cost workstations, backed up by a very high-speed time-sharing system. This is possible because much of the time that COINS researchers spend using a computer typically goes into editing text or program files or using the mail system or other personal assistant functions. For tasks of this kind, the computing power and graphics available in a low-cost workstation such as the SUN is sufficient. Thus, the proposed configuration for MINT will provide the appropriate resources needed by each of the COINS research groups.

III.3. Proposed Equipment

III.3.1. Network Access Workstations

We propose to provide each researcher with a relatively inexpensive workstation, which we call a network access workstation (NAW). Their function is two-fold. First, they will upgrade the quality of terminal interactions, providing high-resolution black and white raster graphics, software for basic window management, control of pointing devices, control of cross-network virtual terminals, and the like. Second, they will provide significant local processing capabilities so that common computing activities, in particular text editing and document production, can be off-loaded from the sophisticated workstations and the timesharing service.

Given the rapidly developing nature of the technology, it is inappropriate to tie down the exact configuration of a typical NAW at this time. A major question will be whether a NAW will need its own mass storage, or whether it will be acceptable for it to either page from a network file server or have sufficient local memory to have its program and data down-loaded at the start of execution. However, we can establish certain minimum specifications. An example of a NAW satisfying our requirements is the SUN Workstation available from SUN Microsystems Inc.. This workstation is based on the 68010 microprocessor: it supports a high-resolution, bit-mapped display, will run BSD32 UNIX, and can be interfaced to an ETHERNET.

Because of the extensive use of graphics in our research, we expect to place the equivalent of the SIGRAPH core graphics package in each workstation to facilitate network graphics with the 750s and the time-sharing system. An early and important use of the NAW will be its integration with the 750s to upgrade them into a prototype of an eventual workstation that will be required by our researchers. We also expect to equip each workstation with a mouse or other comparable pointing device. Beyond these basic capabilities, we anticipate that some NAWs will be adapted for specific applications through the addition of more local memory (the minimum configuration would be one megabyte), color graphics, and specialized interfaces. We are requesting 50 of these NAWs to provide sufficient access points so that all of our 200 users (since a significant number of our researchers spend over 4 hours a day on a machine) can use the NAWs as the basic access path into the network.

We will phase in the purchase of the NAWs so that by the end of year 1 and 2 a significant number of NAWs (i.e., 30) will have been acquired. By having these 30 NAWs in a pooled arrangement, much of the user population will use the NAWs as their standard access path into the network. Additional acquisitions of NAWs in years 4 and 5 will emphasize customized versions of these NAWs adapted to the needs of specific research projects and a more decentralized arrangement so that heavy users can have easier access. These latter acquisitions will also coincide with expected increases in our user population and the more complete adoption by the community of the style of interaction possible with a NAW.

III.3.2. Specialized Workstations

The proposed acquisitions we have described as "specialized workstations" are of two sorts:
(1) specialized equipment for the specific use of our Vision and Robotics groups, e.g., robotic arms (IBM RS1, PUM767), powerful image and array processors (VICOM, COMTAL), and
(2) powerful, single-user processors like the SYMBOLICS 3600 Lispmachine and high-resolution (1024 x 1024 x 16) color graphics display systems (IKONAS, RAMTEK) that will be pooled as a common resource for sophisticated program development and use.

We see the coming availability of specialized language-specific processors (e.g., LISP, PASCAL, ADA machines) with an effective speed of 2-5 MIPS, sophisticated display management, an integrated programming environment and a single-user multiprocessing environment. However, some of these specialized processors have integrated their display management with the basic processor architectures in their initial development. This precludes the effectiveness of the NAW as an access port to these processors via the network. Instead, a researcher will need to be physically located at the specific terminal associated with the processor, thus making the sharing of multiple copies of these processors in a common pool difficult. We have talked with vendors about this issue and we expect that the next generation of many of these specialized processors will permit effective access via a NAW connected to a local area network.

We have planned a gradual phasing in of this equipment over the 5 year period, with a slightly larger proportion of the funds required coming in year one. This will enable us to obtain crucial equipment for Robotics immediately, and also obtain a number of specialized workstations in order to aid the other research groups in acquiring early experience in using them effectively.

The nine VAX 11/750's that will come on-line this year are viewed as specialized workstations in this proposal. As previously discussed, we see these 11/750 single-user systems, particularly as they are currently configured, as an intermediate solution which only partially meets our needs. The use of NAWs as access ports to these stations will remedy some of their deficiencies with respect to graphics. This activity of integrating the NAWs with 750's will be one of the first major software activities undertaken by our software development staff.

The NAW/750 combination should extend the usability of the 750's as a general-purpose single-user workstation for at least 2-3 years. During this period, we also see some of these processors being customized for use as sophisticated controllers for image, array, and graphic processing, and in this way building sophisticated workstations. We also expect to see groups of these processors being configured for both large simulation runs and for experiments by the distributed processing group. During this period, prior to upgrading the VAX 11/780s to an 11/790 (see next section), the 11/750's will still serve as number crunchers which will off-load production runs from the 11/780 time-sharing systems.

In the last year of the grant, we expect to gradually phase out the use of 11/750's as general purpose workstations. We expect that some will be converted to file servers and others customized for specific applications.

III.3.3. General Purpose Timesharing Service

We believe that the timesharing of large mainframes remains the most cost effective means of both providing interactive computing to users with intermittent, but intensive computational demands and also meeting the needs of users with large-scale production runs being done in background. We therefore propose to acquire a much more powerful VAX processor (approximately 10 to 20 times the speed of an 11/780) that we expect to be available from DEC by 1984. Accordingly we have proposed acquisition of this system in 1985; we will finance this purchase in part by the sale of our two existing 11/780 systems. As we acquire more NAWs in our environment and the two year restriction on the use of the 11/750's as single user systems is lifted, the comparative utility of the 11/780's will diminish since it is at most two to five times the speed of those systems. Only a processor that is significantly faster than the other commonly available processors will provide the right balance for those parts of the community with intense but intermittent computational needs.

III.3.4. Communication

The introduction of the NAWs will provide a high degree of connection flexibility in our Research Computing Facility. Prior to the networks of NAWs becoming established in year three, we propose an interim solution of a terminal concentrator with 64 lines. We shall also introduce 8 dial-in/dial-out lines (1200 baud) for communication with other sites. The terminal concentrator will handle both hardwired and dial-in communication and it will increase the efficiency of the timesharing service by relieving it of the responsibility for low-level terminal I/O. It will also permit access to the 750s over the network. The cost of connecting to 11/750s over the network through an 11/780 via a virtual terminal connector for the 11/780 is both computationally expensive and it is difficult to sustain 9600 baud speeds. When more NAWs become available, the main role of the concentrator will shift to handling the dial-in lines.

The hardware we expect to use will consist of four PDP 11/23 machines each handling sixteen 9600 baud lines connected to the Ethernet and running the Ethernet-based terminal concentrator software developed at CMU. A DEC commercial version of this software is expected to be available soon, and we expect there will be other vendors offering a similar product.

III.3.5. <u>Archival Storage</u>

To provide the necessary storage and to make it easily accessible and manageable on our network, we are proposing to purchase four large capacity disks, a video disk system, and two high-speed tape drives.

The video disk system and tape drives will be crucial for handling the large image processing databases and maintaining backups of network files. We plan to use some of the 11/750s as network file servers so as to make these resources available over the network and prevent the overloading of any of the mainframe processors with file management chores.

Expansion of our computer facility and of our research efforts will naturally demand a major addition to our archival storage capabilities. We are already desperately short of disk space due to the needs of our present research program. For example, the VISIONS project requires enormous databases to hold the images to be analyzed, the intermediate results of processing, and the knowledge networks required for interpretation of the images. Likewise, research in information systems requires large databases. For instance, an experimental office automation project exploiting some of MINT's specialized workstations will require substantial storage for document filing and information retrieval applications. Similarly, sophisticated, prototype, network-based software development environments will generate large databases containing software system descriptions. In general, there will be more experimentation possible with the new MINT hardware, resulting in an increase in our storage requirements for experimental data. Even more storage will be needed as the new facilities of MINT become available and additional users come into our environment (we expect the combined faculty of COINS and ECE to increase substantially over the grant period).

III.3.6. Hardcopy Output

In this category we are requesting a second Printronics text/graphics lineprinter plus a medium speed (i.e. 1200baud), letter quality impact printer in year one, a Dunn camera system for recording from our color graphics devices in year two, and a second, higher speed, laser printer in year three. The first three items are all "off-the-shelf" equipment; a "higher speed" laser printer (e.g. 1-2 sec/page) is not yet available in the price range we have specified, however we are expecting competition between manufacturers to change this.

III.3.7. Software

We have a top-level strategy of off-the-shelf acquisition of hardware and software wherever possible in order to make the construction of the facility manageable. We have budgeted approximately $130,000 for the acquisition of software with a major portion of this money coming in the first year. There is a trend of vendors unbundling software costs from hardware, and there is an escalating cost for special purpose software and its maintenance. Thus, we feel an extensive budget for software acquisition is warranted and will markedly improve research productivity.

The following are of some of the major areas in which software products will most probably be acquired:

1. Networking Software - For example, there are already existing commercial products that support TPC/IP protocols under the VAX/VMS operating system. This type of acquisition would help facilitate interfacing our NAWs into the network.

2. Data Base Management System (DBMS) - The acquisition of such a product is especially important for the research in office information systems. If the software is not purchased we will have to reproduce it and waste valuable resources and time. The acquisition of a sophisticated DBMS will also be invaluable for other researchers who have to manage large data bases such as the experimental data in vision and robotics, intermediate results of distributed processing experiments, and the design documentation information that will come out of the tools developed by the software development group.

3. ADA Language Processors - the software development group plans to develop its tools to operate in an ADA environment.

We also plan to acquire a number of other products such as an advanced document production facility that interfaces with our laser printer, statistical and mathematic software packages, etc.

III.4. Technical Issues in the Development of the System

A major issue involves software compatibility among the NAWs and the other processors in the network. Ideally, a NAW could execute all programs residing on the network, possibly in a degraded manner. Currently, however, this level of software compatability is very difficult to achieve, particularly in view of the wide variety of specialized workstations that we intend to acquire. Nevertheless, we will maintain as much software compatibility as possible over the different processors in the network. When there are alternative products of similar capabilities, we intend to purchase those that facilitate compatibility. To assure at least a minimal level of compatibility in the network, we expect that all processors will support Ethernet transactions and that they will be able to both access remote files over the network and directly exchange information with other processors on the network. We have already begun to work on how to integrate workstations using TP/IPC protocols with those based on

DECNET.

We feel it is an inappropriate use of our resources to follow the direction of the SPICE project at CMU, where they are attempting to standardize language processors and operating systems over the network. We recognize, however, the need for standardization of graphics protocols in the network and also the desirability of network protocols that will permit other processors to invoke software that is available only on specialized workstations. We have already addressed the issue of standarized graphics protocols because of the wide variety of graphics hardware in our present facility. We are currently pursuing both high-level approaches to standardization via the DIGRAPH system [DI3079], which implements the SIGRAPH core graphics standards [GSPO79], and low-level approaches via the GUS system [SUTT81] developed locally. We are also currently examining network graphics protocols for our current network configuration in which one VAX controls graphics devices while the other performs general purpose computation.

We believe that by implementing software conventions for graphics and for the remote invocation of specialized software we can deal with most of the issues involved in the sharing of software among research groups without going to standardized language processors and operating systems.

We also recognize that we will eventually need to provide more sophisticated mechanisms supporting the control and scheduling of clusters of workstations dedicated to a single task. The general solution to this problem, which involves a network operating system allowing dynamic process migration, is not critical for our purposes. We believe that, at least initially, the applications that need the added power of multiple workstations can be accommodated by dedicating a set of workstations to the task and then using the basic remote call mechanisms and virtual terminal facilities that will be available on the network. This approach will satisfy most of our immediate needs.

We feel confident that we have the expertise needed to carry out the software modifications necessary for this task. We have already built a number of large software systems for the VAX including a PASCAL system superior to that developed by DEC, a full-scale LISP system, a variety of low- and high-level graphics support tools, and will have developed significant network expertise due to our work this year on integrating the 11/750s into our environment.

III.5. Management Organization

The COINS Department has a solid management organization for its current Research Computer Facility (RCF). It has been in place for two years with representation from all levels of our environment. The 1982/83 operating budget of the RCF will be about $220K with significant University funding and the remainder supplied by grant funds.

Until this year the COINS Department ran its dual-VAX facility for 120 researchers at an embarrassing support level with a single full-time staff member assisted by two half-time graduate student assistants. We now have three full-time professionals supported by three graduate student assistants. This level of support staff has only been possible this year due to $100,000 of temporary emergency funds by the UMass administration; refer to section III.6 on maintenance for further discussion on the need and use of these emergency funds.

As depicted in Figure 2, the current management organization includes a faculty Laboratory Director (Prof. Victor Lesser), a laboratory manager (Mr. Arthur Gaylord, with the equivalent of an M.S. in Computer Science), an assistant laboratory manager (B.S. in Computer Science), a systems software development person (M.S. in Computer Science), and three departmental research assistants. There is a six-member faculty/student RCF committee overseeing the lab operation, setting policy, and deciding upon major equipment acquisitions. There is also a graduate student User Council which makes recommendations to the Lab Manager and the RCF committee. The Chairman of COINS (Prof. Edward Riseman) oversees the entire operation.

Current Management of the RCF

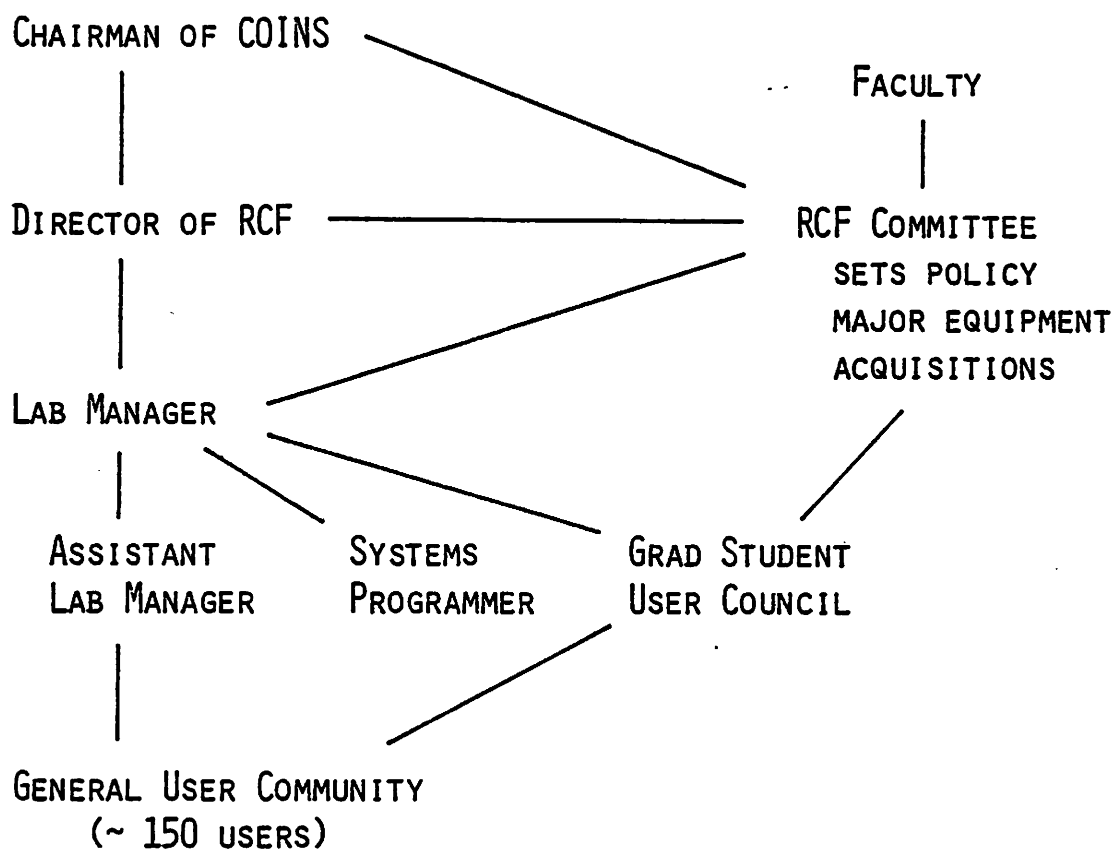

Figure 2

The management structure outlined above has been evolving and now functions in a satisfactory manner. However, we feel that significant additional staffing will be required in order to effectively deal with procurement, software, and maintenance issues involving such a large collection of hardware. We are requesting NSF support for two research associates, two systems programmers, three research assistants and 1/2 support for an administrative assistant. Though this will be a significant addition of personnel in our operation of the RCF, we believe that we have already dealt with the major problems in constructing an effective management structure. Therefore, the organization to integrate the new systems programmers, research associates and research assistants of the MINT support staff outlined in Figure 3 is very similar to the existing structure.

Let us provide a brief discussion of the duties of the particular staff members. The management of a project of this size is a major undertaking, which we feel cannot be accomplished solely through a committee structure. Thus, we have budgeted one-half release time for a senior member of the faculty who will be responsible for the overall project. Aiding the project manager in these duties will be an administrative assistant and a laboratory manager. The administrative assistant will be responsible for routine administrative duties and bookkeeping and serve as project secretary. The laboratory manager will be responsible for the day-to-day operation of the facility. The laboratory

Introduction of New Staff

Into Management Structure

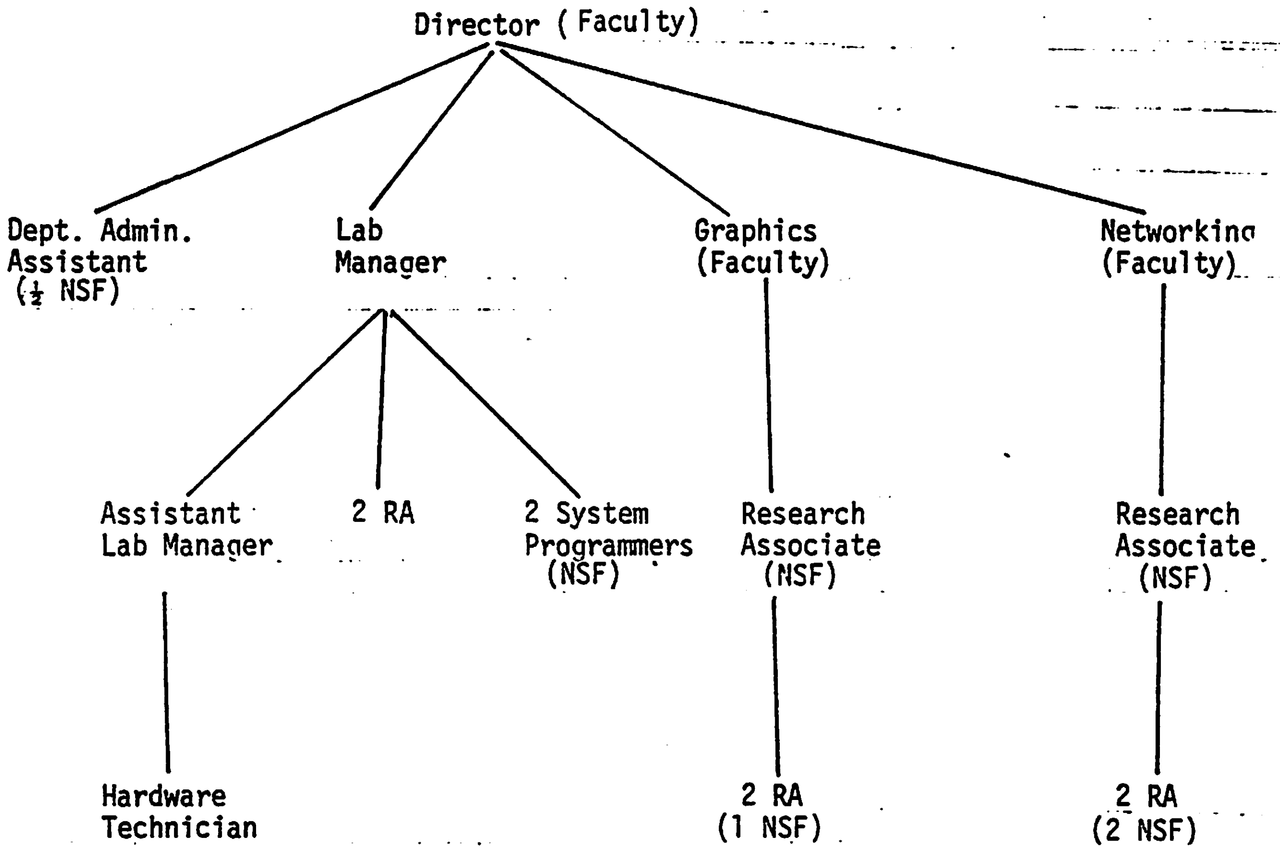

Figure 3

manager's salary will not be charged against this grant since the manager's position is funded by the University. Reporting to the laboratory manager will be the necessary staff to support procurement and installation of the proposed equipment as well as staff to develop the necessary software base.

From our experience in building our current Research Computing Facility, we have found that the task of hardware procurement can become a full-time job. We have recently added an assistant laboratory manager, whose prime responsibility is the procurement and maintenance of hardware. Reporting to this person will be a hardware technician. We feel a technician on the project will facilitate the installation of the wide variety of specialized equipment that we propose to acquire and will significantly decrease maintenance costs.

The laboratory manager's software staff will consist of two systems programmers and two research assistants. This staff will be responsible for general maintenance of the software base and implementation of low-level utilities. We see this as a non-trivial activity, since we anticipate that there will be a non-homogeneous hardware and software base in the network. Thus, there will be a number of different versions of operating systems, language processors, and specialized output devices that will have to be maintained, along with the interfaces necessary for interconnecting these systems.

We also propose to set up two other software support groups, one for developing networking software and the other for developing an extensive system of graphics software. We see both of these activities as major development activities over the entire length of the project. Both groups will be headed by a research associate due to the more research-oriented nature of these software development activities. In addition to a research associate, each group will consist of two research assistants.

The networking group will be responsible for modifications to the specialized workstation software so that these stations can be utilized effectively in the network. For instance, we expect this group to conduct research into the network software necessary to coordinate a number of specialized workstations to provide more processing power for a single task. This is especially appropriate for many of our computationally demanding experimental systems since they are already structured as logically distributed systems. This group will also be responsible for developing the software for the network access workstations and for making modifications to the general purpose timesharing software so that the appropriate division of labor is established among the different processors.

An extensive system of graphics software tools is a common requirement among the research projects. Thus, we feel that a software group responsible for planning and developing an integrated software system providing a wide range of support tools for graphics will be of substantial benefit to our environment. This group will be responsible for designing and, when necessary, implementing the network graphics protocols, the appropriate graphics software for the NAWs, the software for the high- and low-level graphical displays, and the interfaces to the specialized graphics devices in the network.

In our calculation of the personnel budget, we have assumed that the entire complement of people previously outlined will be needed throughout the five-year project. We have also included annual inflation/salary upgrades for years two through five. The budget also contains funds for travel ($5,000/yr.). The travel funds will be used mainly for site visits and conferences involving hardware acquisition, and for one trip per year for the research associates.

III.6. Maintenance Costs

Until this year, the COINS Department faculty have raised sufficient funds for computer costs through their individual research grants to maintain the Research Computer Facility. However, when the facility expanded with the second VAX 11/780, additional mass storage, and other peripherals, the maintenance budget began to dominate the entire financial resources of the COINS Department.

With the arrival of the 9 VAX 11/750s currently taking place, we expect our maintenance budget will jump to about $120,000 per year. This estimate is based upon the hope that the nine 11/750s will only cost us about $50,000 instead of the $100,000 for the standard maintenance contract. We are negotiating a new type of maintenance contract for the 11/750s with DEC that will involve low-priority maintenance and board-swapping. Because there are nine identical units and each unit is only serving a single user, down-time on one or two machines is tolerable.

The equipment maintenance outlined above is part of an estimated $220,000 operational budget which includes personnel and supplies. The University administration has responded with $100,000 of temporary emergency funds this year to supplement our external grant support. About $75,000 is intended for maintenance and $25,000 towards personnel. Without these funds the evolution of our environment would have been severely disrupted.

We are in a difficult situation, even without the acquisition of any new equipment. The University administration has been attempting to respond to our needs during a period of several years where the overall University budget has been declining in terms of inflation-adjused dollars. Thus, we foresee serious impact to our still evolving research program (and even possible disaster) if we do not get some external support to deal with the maintenance problem. This proposal requests $883,800 for maintenance over the five years, which includes

$50,000 per year for the nine 11/750s.

IV. BUDGETS

IV.1. Budget Summary

	Equipment	Maintenance/ Operations	Salaries/ Fringe/ Travel	Indirect Costs	Total
Year 1	$ 689,000	$114,200	$172,347	$128,946	$1,104,493
Year 2	447,500	148,550	182,688	149,057	927,795
Year 3	472,000	214,350	193,649	183,600	1,063,599
Year 4	370,000	241,950	205,268	201,248	1,018,466
Year 5	372,000	272,750	217,584	220,650	1,082,984
TOTAL	$2,350,500	$991,800	$971,536	$883,501	$5,197,337

EQUIPMENT, MAINTENANCE, OPERATION, AND ASSOCIATED INDIRECT COSTS	$3,788,610
PERSONNEL, FRINGE BENEFITS, TRAVEL COSTS, AND ASSOCIATED INDIRECT COSTS	1,408,727
TOTAL	$5,197,337

IV.2. MINT Equipment and Maintenance Budget

	Price	Maintenance/yr.*
A. Network Access Workstations		
1. 50 workstations powerful microprocessor (e.g. 68010), high density BW display, 1-2 Mbytes main memory, interface to the Ethernet; operating system with virtual memory, multiprocessing, window management, and networking software (e.g. BSD32 UNIX); phased acquisition: $15-10,000 each	$600,000	$30,000
	$600,000	$30,000
B. Specialized Processors		
1. For Vision and Robotics a) Image processing station: pixcel level operations, high memory bandwidth, (e.g. VICOM, COMTOL) b) Array Processor (FPS) c) CAM d) Robotics hardware	$360,000	$28,800
2. For Sophisticated, Interactive Programs a) High-Resolution Color Graphic Stations (e.g., Ikonas and Ramtek) b) High Performance Language-Specific Machines (e.g., LISP machines: Symbolics and Xerox; Pascal machines: enhanced Perqs and Appollos)	640,000	51,200
	$1,000,000	$80,000
C. General Purpose Timesharing Service		
1. High-Speed Processor VAX 11/790 processor, expected to be available in 1984; replaces the two existing 11/780's; net cost: 500K cost - 250K trade-in, maintenance calculated on cost.	$250,000	$40,000
	$250,000	$40,000

MINT Equipment and Maintenance Budget (Cont.)

	Price	Maintenance/yr*
D. Shared Resources		
1. Storage		
a) 4 RU81 Disk drives (486mb) ($26,000 each)	$104,000	$10,400
b) 2 TU78 tape drives + controller	79,500	7,950
c) Read/write archival video disk	50,000	5,000
	$233,500	$23,350
2. Hardcopy Output		
a) High-speed Laser printer	$50,000	$5,000
b) Printronix P-300 text/graphics printer	8,000	800
c) Graphics Hardcopy (e.g.,DUNN camera system with 35mm and Polaroid)	22,000	2,200
	$80,000	$8,000
3. Software		
a) Network Support		
b) Relational Database Management System		
c) IMSL - Mathematics and Statistics Subroutine Library		
d) BLISS-32 language license		
e) VAX/VMS source license		
f) Ada language license		
g) Document production system		
h) General Software	$130,000	$10,400
	$130,000	$10,400

MINT Equipment and Maintenance Budget (Cont.)

	Price	Maintenance/yr.*
D. Shared Resources (Continued)		
4. Communications		
a) Network Terminal Concentrator (e.g., 4 PDP 11/23)	$ 50,000	4,000
b) 8 dial-ins (e.g., VADIC 300/1200 baud answer/originate)	7,000	0
	$ 57,000	$ 4,000
	$500,500	$45,750

E. Facility Operating Costs

1. Disc Packs
2. Paper
3. Telephone line charges for networking and remote access
4. Tapes

$108,000

FIVE YEAR TOTAL FOR EQUIPMENT	$2,350,500
FIVE YEAR TOTAL FOR OPERATING COST	108,000

*Maintenance has been estimated at 10% for mechanical devices, 8% for software and one-of-a-kind processors, and 5% for multiple processors for which we can tolerate some down time and thus expect to negotiate a reduced maintenance rate.

IV.3. MINT Yearly Equipment and Maintenance Budget

	Price	Maintenance/yr	Other
Year 1			
1. 15 Network Access Workstations	$220,000	$ 11,000	
2. Specialized Processors	300,000	24,000	
3. Terminal Concentrator	50,000	4,000	
4. Software	50,000	4,000	
5. TU78 Tape drive plus controller	54,000	5,400	
6. Printronics	8,000	800	
7. Dial-ins	7,000	0	
8. Maintenance on VAX 750s	0	50,000	
9. Operating Costs	0	0	15,000
YEAR 1 TOTALS	$689,000	$ 99,200	$15,000
Year 2			
1. 15 Network Access Workstations	$180,000	$ 9,000	
2. Specialized Processors	200,000	16,000	
3. TU78 tape drive	25,500	2,550	
4. Graphics Hardcopy	22,000	2,200	
5. Software	20,000	1,600	
6. Operating Costs	0	0	$18,000
	$447,500	$ 31,350	$18,000
Accumulated Maintenance		99,200	
YEAR 2 TOTALS	$447,500	$130,550	$18,000
Year 3			
1. VAX 11/790 processor	$250,000	$ 40,000	
2. Specialized Workstations	100,000	8,000	
3. High-speed Laser printer	50,000	5,000	
4. 2 RU81 disk drives	52,000	5,200	
5. Software	20,000	1,600	
6. Operating Costs	0	0	$24,000
	$472,000	$ 59,800	$24,000
Accumulated Maintenance		130,550	
YEAR 3 TOTALS	$472,000	$190,350	$24,000

MINT Yearly Equipment and Maintenance Budget (Cont.)

	Price	Maintenance/yr	Other
Year 4			
1. 10 Network Access Workstations	$ 100,000	$ 5,000	
2. Specialized Processors	200,000	16,000	
3. Archival Video Disk	50,000	5,000	
4. Software	20,000	1,600	
5. Operating Costs	0	0	$24,000
	$ 370,000	$ 27,600	$24,000
Accumulated Maintenance		190,350	
YEAR 4 TOTALS	$ 370,000	$217,950	$24,000
Year 5			
1. 10 Network Access Workstations	$ 100,000	$ 5,000	
2. Specialized Processors	200,000	16,000	
3. Software	20,000	1,600	
4. 2 RU81 Disk drives	52,000	5,200	
5. Operating Costs	0	0	$27,000
	$ 372,000	$ 27,800	$27,000
Accumulated Maintenace		217,950	
YEAR 5 TOTALS	$ 372,000	$245,750	$27,000
5 YEAR TOTAL	$2,350,500	$883,800	$108,000

TOTAL FOR EQUIPMENT, MAINTENANCE, AND OPERATIONS	$3,342,300
INDIRECT COSTS ON MAINTENANCE AND OPERATIONS	446,310
TOTAL FACILITY COST OVER FIVE YEARS	$3,788,610

MINT Yearly Equipment and Maintenance Budget (Cont.)

IV.4. Personnel Budget

	NSF Costs	University	Department Grants
Salaries (Year 1)			
MINT Director (senior faculty) 1/2 release time (@$20,000)	$ 0	$ 20,000	$ 0
Lab Manager	0	32,000	0
Assistant Lab Manager	0	0	20,000
2 Research Associates (@$30,000)	60,000	0	0
1 Technician	0	20,000	0
2 Systems Programmers (@$25,000)	50,000	0	0
6 Research Assistants (@$8,000)	24,000	0	24,000
1 Administrative Assistant	9,000	0	0
Total Salaries	$ 143,000	$ 72,000	$ 44,000
Additional Personnel Costs (Year 1)			
Fringe Benefits @20.46%	$ 24,347	$ 14,731	$ 4,092
Travel	5,000	0	0
Indirect Costs @45%	77,556	39,029	21,641
YEAR 1 TOTALS	$ 249,903	$125,760	$ 69,733
Direct and Indirect Costs for 5 years			
First Year Direct Costs	$ 172,347	$ 86,731	$ 48,092
Second Year Direct Costs	182,688	91,935	50,978
Third Year Direct Costs	193,649	97,451	54,037
Fourth Year Direct Costs	205,268	103,298	57,279
Fifth Year Direct Costs	217,584	109,496	60,716
TOTAL DIRECT COSTS OVER FIVE YEARS	$ 971,536	$488,911	$271,102
TOTAL INDIRECT COSTS OVER FIVE YEARS	437,191	220,010	121,996
5 YEAR TOTAL FOR PERSONNEL	$1,408,727	$708,921	$393,098

TOTAL UNIVERSITY AND DEPARTMENT COMMITMENT $1,102,019

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

Year 1

ORGANIZATION University of Massachusetts/Amherst	FOR NSF USE ONLY PROPOSAL NO.	DURATION (MONTHS) Proposed	Granted
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Edward M. Riseman	AWARD NO.		

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $20,000)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. ( ) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES				60,000	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass:$20,000				50,000	
3. (6) GRADUATE STUDENTS @$8,000, (UMass: $24,000)				24,000	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL - Administrative Assistant (1/2 time)				9,000	
6. (2) OTHER Lab Manager (UMass: $32,000), Asst. Lab Mgr. (UMass: $20K)				0	
TOTAL SALARIES AND WAGES (A+B)				143,000	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS) 20.46%				24,347	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				167,347	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See detailed budget section of proposal					
TOTAL PERMANENT EQUIPMENT				689,000	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				5,000	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS 1. STIPENDS $ 2. TRAVEL 3. SUBSISTENCE 4. OTHER					
TOTAL PARTICIPANT COSTS				0	
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				15,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER - Maintenance				99,200	
TOTAL OTHER DIRECT COSTS				114,200	
H. TOTAL DIRECT COSTS (A THROUGH G)				975,547	
I. INDIRECT COSTS (SPECIFY) 45% Modified Total Direct Costs (Equipment not included)					
TOTAL INDIRECT COSTS				128,946	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				1,104,493	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$1,104,493	$

PI/PD TYPED NAME & SIGNATURE* Edward M. Riseman *Edward M. Riseman*	DATE 9/14/82	FOR NSF USE ONLY INDIRECT COST RATE VERIFICATION
INST. REP. TYPED NAME & SIGNATURE* Lee D. Beatty, Director, Grants & Contracts	DATE	Date Checked / Date of Rate Sheet / Initials - DGC / Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

Year 2

FOR NSF USE ONLY

ORGANIZATION	PROPOSAL NO.	DURATION (MONTHS) Proposed	Granted
University of Massachusetts/Amherst			
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR	AWARD NO.		
Edward M. Riseman			

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $21,200)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. (1) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES Research Associates				63,600	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass: $21,200				53,000	
3. (6) GRADUATE STUDENTS @8,480 each (UMass: $25,440)				25,440	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL Administrative Assistant (1/2 time)				9,540	
6. (2) OTHER Lab Manager (UMass: $33,920), Asst Lab Mgr (UM: $21,200)				0	
TOTAL SALARIES AND WAGES (A+B)				151,580	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				25,808	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				177,388	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See Section IV (Budgets) of the proposal for detailed list.					
TOTAL PERMANENT EQUIPMENT				447,500	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				5,300	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS 1. STIPENDS $ 2. TRAVEL 3. SUBSISTENCE 4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				18,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER -- Maintenance				130,550	
TOTAL OTHER DIRECT COSTS				148,550	
H. TOTAL DIRECT COSTS (A THROUGH G)				778,738	
I. INDIRECT COSTS (SPECIFY) 45% of Modified Total Direct Costs (Total Direct Costs less Equipment)					
TOTAL INDIRECT COSTS				149,057	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				927,795	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$ 927,795	$

PI/PD TYPED NAME & SIGNATURE*	DATE	FOR NSF USE ONLY		
Edward M. Riseman *Edward M. Riseman*	9/14/82	INDIRECT COST RATE VERIFICATION		
INST. REP. TYPED NAME & SIGNATURE*	DATE	Date Checked	Date of Rate Sheet	Initials - DGC
Lee D. Beatty, Director Grant & Contracts				Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

Year 3

ORGANIZATION: University of Massachusetts/Amherst	FOR NSF USE ONLY — PROPOSAL NO.	DURATION (MONTHS) Proposed	Granted
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR: Edward M. Riseman	AWARD NO.		

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $22,472)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. (1) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES Research Associates				67,416	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass:$22,472				56,180	
3. (6) GRADUATE STUDENTS @$8,989 each (UMass: $26,966)				26,966	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL Administrative Assistant (1/2 time)				10,112	
6. (2) OTHER Lab Manager (UMass:$35,955) Lsst Lab Mgr (UM: $22,472)				0	
TOTAL SALARIES AND WAGES (A+B)				160,674	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				27,357	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				188,031	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See Section IV.2 of the proposal for detailed list.					
TOTAL PERMANENT EQUIPMENT				472,000	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				5,618	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS 1. STIPENDS $ 2. TRAVEL 3. SUBSISTENCE 4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				24,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Maintenance				190,350	
TOTAL OTHER DIRECT COSTS				214,350	
H. TOTAL DIRECT COSTS (A THROUGH G)				879,999	
I. INDIRECT COSTS (SPECIFY) 45% Modified Total Direct Costs (Total Direct Costs less Equipment) TOTAL INDIRECT COSTS				183,600	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				1,063,599	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$ 1,063,599	$

PI/PD TYPED NAME & SIGNATURE*: Edward M. Riseman *Edward M. Riseman*	DATE: 9/14/82	FOR NSF USE ONLY — INDIRECT COST RATE VERIFICATION
INST. REP. TYPED NAME & SIGNATURE*: Lee D. Beatty, Director, Grants & Contracts	DATE	Date Checked / Date of Rate Sheet / Initials - DGC / Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

Year 4

	FOR NSF USE ONLY	
ORGANIZATION: University of Massachusetts/Amherst	PROPOSAL NO.	DURATION (MONTHS): Proposed / Granted
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR: Edward M. Riseman	AWARD NO.	

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $23,820)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. (1) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES Research Associates				71,461	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass:$23,820				59,551	
3. (6) GRADUATE STUDENTS @$9,528 each (UMass: $28,584)				28,584	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL Administrative Assistant (1/2 time)				10,719	
6. (2) OTHER Lab Manager (UMass:$38,112) Asst Lab Mgr (UM:$23,820)				0	
TOTAL SALARIES AND WAGES (A+B)				170,315	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				28,998	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				199,313	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See Section IV.2 of the proposal for detailed list.					
TOTAL PERMANENT EQUIPMENT				370,000	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				5,955	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS 1. STIPENDS $ 2. TRAVEL 3. SUBSISTENCE 4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				24,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Maintenance				217,950	
TOTAL OTHER DIRECT COSTS				241,950	
H. TOTAL DIRECT COSTS (A THROUGH G)				817,218	
I. INDIRECT COSTS (SPECIFY) 45% Modified Total Direct Costs (Total Direct Costs less Equipment) TOTAL INDIRECT COSTS				201,248	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				1,018,466	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$ 1,018,466	$

PI/PD TYPED NAME & SIGNATURE*	DATE	FOR NSF USE ONLY		
Edward M. Riseman *Edward M. Riseman*	9/14/82	INDIRECT COST RATE VERIFICATION		
INST. REP. TYPED NAME & SIGNATURE*	DATE	Date Checked	Date of Rate Sheet	Initials - DGC
Lee D. Beatty, Director Grants & Contracts				Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

Year 5

FOR NSF USE ONLY	
PROPOSAL NO.	DURATION (MONTHS) Proposed / Granted
AWARD NO.	

ORGANIZATION: University of Massachusetts/Amherst

PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR: Edward M. Riseman

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $25,249)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. (1) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES Research Associates				75,748	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass:$25,249				63,124	
3. (6) GRADUATE STUDENTS @$10,100 each (UMass: $30,300)				30,300	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL Administrative Assistant (1/2 time)				11,362	
6. (2) OTHER Lab Manager (UMass:$40,398) Asst Lab Mgr (UMass:$25,249)				0	
TOTAL SALARIES AND WAGES (A+B)				180,534	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				30,738	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				211,272	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See Section IV.2 of the proposal for detailed list.					
TOTAL PERMANENT EQUIPMENT				372,000	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				6,312	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS					
1. STIPENDS $					
2. TRAVEL					
3. SUBSISTENCE					
4. OTHER					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				27,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER				245,750	
TOTAL OTHER DIRECT COSTS				272,750	
H. TOTAL DIRECT COSTS (A THROUGH G)				862,334	
I. INDIRECT COSTS (SPECIFY) 45% Modified Total Direct Costs (Total Direct Costs less Equipment)					
TOTAL INDIRECT COSTS				220,650	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				1,082,984	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$ 1,082,984	$

PI/PD TYPED NAME & SIGNATURE*	DATE	FOR NSF USE ONLY		
Edward M. Riseman *Edward M. Riseman*	9/14/82	INDIRECT COST RATE VERIFICATION		
INST. REP. TYPED NAME & SIGNATURE*	DATE	Date Checked	Date of Rate Sheet	Initials - DGC
Lee D. Beatty, Director Grants & Contracts				Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

(SEE INSTRUCTIONS ON REVERSE BEFORE COMPLETING)

SUMMARY PROPOSAL BUDGET

TOTAL FIVE YEAR BUDGET

ORGANIZATION: University of Massachusetts/Amherst	FOR NSF USE ONLY — PROPOSAL NO.	DURATION (MONTHS): Proposed / Granted
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR: Edward M. Riseman	AWARD NO.	

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	NSF FUNDED PERSON-MOS. CAL.	ACAD	SUMR	FUNDS REQUESTED BY PROPOSER	FUNDS GRANTED BY NSF (IF DIFFERENT)
1. Project Director - 1/2 time release (UMass: $112,741)				$ 0	$
2.					
3.					
4.					
5. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
6. (1) TOTAL SENIOR PERSONNEL (1-5)				0	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. (2) POST DOCTORAL ASSOCIATES Research Associates				338,225	
2. (3) OTHER PROFESSIONALS (TECHNICIAN, PROGRAMMER, ETC.) UMass:112,741				281,855	
3. (6) GRADUATE STUDENTS (UMass: $135,290)				135,290	
4. ( ) UNDERGRADUATE STUDENTS					
5. (1) SECRETARIAL-CLERICAL Administrative Assistant (1/2 time)				50,733	
6. (2) OTHER Lab Manager (UMass:$180,385) Asst Lab Mgr (UM:$112,741)				0	
TOTAL SALARIES AND WAGES (A+B)				806,103	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				137,248	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				943,351	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM EXCEEDING $1,000; ITEMS OVER $10,000 REQUIRE CERTIFICATION) See Section IV.2 of the proposal for detailed list.					
TOTAL PERMANENT EQUIPMENT				2,350,500	
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				28,185	
2. FOREIGN					
F. PARTICIPANT SUPPORT COSTS 1. STIPENDS $ ___ 2. TRAVEL ___ 3. SUBSISTENCE ___ 4. OTHER ___					
TOTAL PARTICIPANT COSTS					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				108,000	
2. PUBLICATION COSTS/PAGE CHARGES					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Maintenance				883,800	
TOTAL OTHER DIRECT COSTS				991,800	
H. TOTAL DIRECT COSTS (A THROUGH G)				4,313,836	
I. INDIRECT COSTS (SPECIFY) 45% Modified Total Direct Costs (Total Direct Costs less Equipment)					
TOTAL INDIRECT COSTS				883,501	
J. TOTAL DIRECT AND INDIRECT COSTS (H + I)				5,197,337	
K. RESIDUAL FUNDS (IF FOR FURTHER SUPPORT OF CURRENT PROJECTS GPM 252 AND 253)					
L. AMOUNT OF THIS REQUEST (J) OR (J MINUS K)				$ 5,197,337	$

PI/PD TYPED NAME & SIGNATURE*: Edward M. Riseman *Edward M. Riseman*	DATE: 9/14/82	FOR NSF USE ONLY — INDIRECT COST RATE VERIFICATION
INST. REP. TYPED NAME & SIGNATURE*: Lee D. Beatty, Director Grants & Contracts	DATE	Date Checked / Date of Rate Sheet / Initials - DGC / Program

NSF FORM 1030 (8-80) SUPERSEDES ALL PREVIOUS EDITIONS

*SIGNATURES REQUIRED ONLY FOR REVISED BUDGET (GPM 233)

V. RESEARCH SUMMARIES

This section contains summaries of our research efforts, grouped into the four major areas. More detailed descriptions of the projects appear in section VI.

V.1. Distributed Systems

This research area involves eight faculty members, A. Barto, W. Kohler, V. Lesser, N. Spinelli, J. Stankovic, H. Stone, D. Towsley, and J. Wileden, and research associate D. Corkill. The funding for this research area comes from a variety of sources, AFOSR, ARMY, DARPA, DEC, NSF, and totals approximately $500,000 per year.

A common theme among many of these research projects is the study of coordination in distributed computing. This work is concerned with how to organize and control large distributed computational structures. Of special concern is development of methodologies for decentralized control and cooperation which are effective and robust in face of a dynamically changing task load and hardware environment. Integrally tied to this development is the construction of decentralized problem-solving and control algorithms that can function effectively even though they have incomplete and possibly inconsistent views of the state information necessary for their computations. Such an approach to decentralized coordination is of crucial importance both to the extension of distributed processing to new application areas such as distributed sensor networks and distributed command and control systems, and to the construction of networks of hundreds of nodes in which there is significant interaction (through cooperative problem-solving or resource sharing) among nodes.

Under the various grants, coordination among distributed computing elements is being investigated at different levels, including the application level, the distributed database level, the operating system level, the communications level, and a level that makes contact with neurophysiology. Most of this work is based on empirical studies using testbeds or simulation models. Lack of computer power for these empirical studies has significantly limited the range of experiments; especially with large networks. Only in such large networks do the problems of cooperation and coherence and the effects of organizational structure dominate.

The fundamental scientific issues being dealt with by these research projects are:

1. The development of a methodology for partitioning and assigning tasks. The goal is to organize tasks in a distributed system to maximize parallelism, to minimize communication, and to permit graceful degradation in the face of individual component failure. As part of this methodolgy, we are seeking to understand what are the task and hardware characteristics that dictate when a particular organizational structure is most appropriate. We are also studying ways to redistribute tasks in a distributed system dynamically so as to maintain high performance and reliability as computational demands, processing resources, and communication resources vary.

2. The development of techniques for effective cooperation among nodes. We are pursuing the development of strategies that permit nodes to cooperate in order to stop the spread of incorrect results generated by nodes with inconsistent or incomplete information. We are

also exploring whether these same strategies for cooperation can handle incorrect results caused by faulty processing nodes or communication channels.

3. The development of mechanisms for obtaining organized behavior using decentralized control. Inherent in the control of a distributed system is the problem of maintaining an accurate model of the entire system when there is no global mechanism for gathering and acting on the current state of system performance. We are developing control algorithms that operate satisfactorily (e.g., are stable and sufficiently accurate) even when the available information about the system is inconsistent or incomplete. We are also studying ways to address reliability questions as part of the general issue of decentralized control.

4. The development of formal specification and analysis techniques for distributed control. We are developing formal approaches to the specification of behavioral abstractions of concurrent activities and the use of these specifications as a debugging tool for distributed systems.

The approaches to these scientific issues vary greatly among the individual researchers. Lesser, from an empirical and application level perspective, is addressing these issues by developing a node architecture that permits sophisticated decentralized decision-making; this form of complex local control allows a node to balance its own perception of appropriate problem solving activity with activities deemed important by other nodes. 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 [LESS80, 81a, 81b].

Using a cooperative distributed computational paradigm, Spinelli and Barto are studying, largely by simulation, how networks of complex neuron-like adaptive elements (approaching the complexity of simple computers) can form a substrate for adaptive problem solving in such domains as vision and motor control [ARBI79b,81a,81b, SPIN79, SUTT81a]. Kohler, Stankovic, and Towsley are also pursuing an empirical approach to issues in distributed databases and have built a software testbed on a local-area network* [KOHL80,81]. This testbed is used to test the actual performance of concurrency control and recovery algorithms for distributed transaction systems [GARC81].

Other researchers are approaching these issues from a more formal and domain-independent perspective. Stankovic [STAN78,80,81] and Stone [STON77,78a,78b] are using techniques from probability theory and graph theory, respectively, to develop scheduling techniques for distributed operating systems. Stankovic is developing large simulation models for several decentralized scheduling algorithms. Wileden is approaching concurrency control issues through the development of a formal language for the specification of behavioral abstractions of concurrent activity [WILE79]. He is using this language to provide a framework for the debugging and performance monitoring of distributed systems [BATE81]. Towsley is analyzing various communication protocols using queueing theory and

*Kohler is currently using a distributed processing network developed at DEC to run the testbed.

simulation.

While the different research projects in distributed computing are interrelated, they also tie very closely into the research projects of Croft and Stemple on databases, Clarke and Wileden on the pre-implementation phase specification of concurrent software, and Hanson and Riseman on the use of logically distributed computational structures for computer vision.

V.2. Information Systems and User Interfaces

In this section we will discuss research in a variety of areas including expert systems and natural language, office information systems, text retrieval systems and database systems. The faculty involved in this research are B. Croft, W. Lehnert, V. Lesser, D. McDonald, R. Moll, J. Orenstein, E. Rissland and D. Stemple, and research associate S. Kulikowski. Their work is currently being supported by grants from NSF, the Sloan Foundation and Digital Equipment Corporation totalling about $445,000 annually.

Although the several areas grouped above have for the most part been regarded as distinct, they deal with aspects of the same problem; that is how to store, manipulate, retrieve and display information for different populations of users with different requirements. The design of intelligent, friendly interfaces for these systems is a central theme for the research in our environment and it is the basis of several joint projects. Some of the topics in information systems and user interfaces that are studied by the faculty are as follows:

Expert Systems
- Natural language interfaces
- Knowledge representation
- The structure, use and generation of examples
- Adaptive control strategies and machine learning
- Handicapped communication for the non-vocal severely disabled

Office Information Systems
- Intelligent assistance based on office procedures
- Filing and retrieval systems for office text

Text Retrieval
- Probabilistic models of retrieval and indexing
- Adaptive search strategies

Database Systems
- Optimization of relational query execution
- Database schema design and translation
- Human factors in query languages
- Implementation of relational databases
- Distributed databases

The design of sophisticated office information systems requires the integration of techniques from a number of areas in information systems and this application area is being used both to coordinate current research efforts and generate new areas of interest. A major result of this concentration on a common topic is the large project (involving six faculty) funded by Digital Equipment Corporation to develop a system which will provide intelligent assistance to office workers in the execution of their tasks as represented by office procedures [CROF82c]. The procedures are defined using a formalism which provides the system with both procedural and goal-oriented information about the user's actions [HUFF82]. The issues involved in the design of this system center around being able to do interpretation and planning in an environment where the user's actions may be partially defined with procedures, but where there are also many exceptions and ambiguities. As part of providing a flexible and supportive environment, the intelligent assistant will have a natural language input facility and an output facility using both natural language and examples. Another goal of this project is to apply both document retrieval and database techniques to the design of an "electronic filing cabinet" in order to achieve effective access to the data and text in the office.

More generally, in the area of expert systems McDonald and Lehnert are working in complementary areas of natural language processing. McDonald is developing general techniques for the implementation of natural language interfaces as an extension of his work on natural language generation [MCDO80]. His work will be used to generate descriptions and explanations of the procedures being used by the office assistant system. Lehnert's research involves the design of internal memory representations for text understanding systems, the role of inference processes in natural language and interfaces for database systems [LEHN81, 82a, 82b]. She also has extensive experience in the design of natural language interfaces for database systems. In another area related to user interfaces, Rissland has focused on studying the structure of knowledge in complex domains such as law and the generation and use of examples meeting specified constraints [RISS80a, 80b, 80c]. Her example generation system will be used in cooperative projects in vision, software develoment and office systems.

Other work in the expert systems area is done by Moll and Kulikowski. Kulikowski is in the process of developing an intelligent user interface for severely disabled non-vocal clients such as cerebal palsy and stroke victims. The system resides on a microcomputer, has a user-defined knowledge base and allows an individual who can control a switch via a single muscle to generate text and speech output in a relatively effective manner [KULI81]. Moll, in conjunction with DEC, is using production systems for the task of UNIBUS configuration.

In the area of text retrieval, Croft is working on systems for filing and retrieving text documents by content in the office environment and in a library environment [CROF82a, 82b]. These systems will enable users with little experience to locate, file and classify documents much more effectively than a conventional system using Boolean queries. The most recent development in providing flexible and effective access to the text is an adaptive system which is based on work done by Barto on adaptive networks [CROF81a].

Another area of concentration is database systems. In this area, Stemple is working on algorithms for optimizing the execution of relational queries taking into account the implementations of relational algebra primitives. He is also considering the human factors involved in comparing procedural and non-procedural query languages [WELT81], and the design of conceptual database schemas, particularly for database translation. Orenstein, who will be joining the faculty in

January, is studying the data structures needed for the implementation of relational databases [OREN82].

V.3. Software Development

The faculty involved in software development are D. Carlson, L. Clarke, R. Graham, K. Ramamritham, J. Wileden, and S. Zeil, and research associate D. Richardson. They have been investigating a wide range of issues concerning the software development process. Currently, members of this group receive support from NSF and NASA, as well as some industrial grants from GE and Hewlett-Packard, totaling about $210,000 per year. As indicated by our faculty hiring over the past several years, we consider software development to be an important, emerging research area of COINS.

The software development group has been primarily focusing on two central themes:

- The design of suitable languages for representing software development during the various phases of the lifecycle;
- The design, implementation, and evaluaton of analysis and testing techniques that provide feedback to the user on the consistency, completeness, and correctness of an evolving software system.

The language design aspect of our research is concerned with the development of constructs supporting particular concepts in software development. For example, Clarke, Ramamritham, and Wileden are working on projects that are evaluating and developing notations for representing concurrent behavior. While the underlying assumptions differ with each project, these researchers are working closely together. In addition, Clarke and Wileden are investigating constructs for expressing real-time constraints, exceptions and exception handlers, and visibility control [CLAR80b]. While this language design effort is particularly concerned with the pre-implementation phases of the software lifecycle, consistent, appropriate abstractions are being sought which facilitate transitions between representations for all phases of the lifecycle.

Our work in language design is strongly influenced by the belief that analysis is of foremost importance in the software development process. Our conviction that the feedback provided from analysis techniques will increase the reliability, while decreasing the cost, of a software system under development has led us to address the issue of software analysis on many fronts. Wileden, in collaboration with Professor G. Avrunin in the Mathematics department, is developing a modelling scheme and an algebraic technique for describing and evaluating special classes of concurrent systems [AVRU81]. Ramamritham is investigating a model of concurrency that focusses on processes that manage shared resources and is developing analysis techniques for this model that are based upon temporal logic. Carlson is also investigating a formal model of resource management. Clarke, Richardson, and Zeil are examining several analysis techniques to support the software testing process. Clarke and Richardson have been investigating symbolic evaluation methods [CLAR80a, CLAR81b], as well as developing an experimental symbolic evaluation system that supports automatic test data generation [CLAR76, 78]. This work has lead to the development of the partition analysis method [RICH81], which incorporates information from the software design, as well as the implementation, and utilizes this information in the judicious selection of test cases. Zeil has been investigating some of the theoretical limits of certain testing methods [ZEIL81] and is currently working on a system to empirically study the relative power of a number of proposed testing strategies. Clarke, Richardson, and Zeil intend to seek support to examine methods for integrating and evaluating several approaches to testing [CLAR82b].

In addition to the various research projects mentioned above, which are investigating particular descriptive concepts and analysis techniques, Clarke, Graham, and Wileden are seeking support for their work on the development of a prototype software development environment. It is the intent of this project to address issues of integration and uniformity while concentrating on notation and analysis techniques best suited for the pre-implementation phase of software development [CLAR81a]. Thus, the project will focus on developing tools for use in the pre-implementation phases, providing the user with a knowledgeable and uniform interface, and establishing a consistent internal repository of information from which the various analysis tools can enter and retrieve information.

The work in software development is continually benefiting from interactions with the other researchers in our environment. The work of Lesser has provided insight into concurrent software development problems. Moreover, the development of an integrated environment will draw directly from the expert systems research of Rissland and Lesser and from the research on information systems and database management of Croft and Stemple.

V.4. Vision and Robotics

This section provides a brief overview of current research in the general areas of machine vision (utilizing both static and temporally varying images), robot sensor development, mechanisms for robot control and visuomotor coordination, image transmission over narrow bandwidth channels, analysis of remotely-sensed imagery, the structure and modelling of biological visual systems, and distributed computing. The faculty involved in this research are Arbib, Elliott, Foster, Hanson, Hill, Riseman, and Spinelli with Research Associates Lawton, Overton, and Reynolds. This research is currently funded by grants from ONR, NSF, NIH, AFOSR, Digital Equipment Corporation, Tufts New England Medical Center, Sloan Foundation, and ASTEC, totalling over $600,000 per year.

The Vision and Robotics group is composed of an integrated group of researchers focussing on a set of common problems of fundamental scientific importance. They include the extraction and description of multimodal sensory data, the creation and maintenance of environmental models, the structure of cooperating control systems, and the development of knowledge structures necessary for integrating diverse sources of sensory data into a comprehensive whole. We are also concerned with the neurophysiological correlates of early visual processing and with the corresponding methods developed in biological systems; e.g., Spinelli's pioneering work [SPIN70,72,79,80] on the effects of early visual experience on neural connections and feature detection systems.

More specifically, the work of Hanson and Riseman [HANS78a,78b, PARM80, WEYM82, WESL82] on the VISIONS Image Understanding System is aimed at the development of an experimental testbed for investigating the construction of integrated computer vision systems. The goal is to provide an analysis of color images of outdoor scenes, from segmentation through the final stages of symbolic interpretation. The output of the system is intended to be a symbolic representation of the three-dimensional world depicted in the two-dimensional image. This involves the determination of object labels for image regions and an approximate placement of objects in three-dimensional space [HANS78b, PARM80, WEYM82]. The VISIONS system has been successfully applied to a variety of image types and practical problems, including our work with outdoor scenes, biomedical domains (retinal studies for tracking glaucoma, analysis of melanoma), industrial images, and remotely sensed imagery.

Concern with the sensory guidance of motion and the extraction of environmental depth information in images has led to the current work of Lawton, Hanson, and Riseman on machine processing of image sequences. The focus of the motion research is on methods for directly constructing environmental surface information by means of optic flow fields derived from a sequence of images in which both the sensor and objects in the scene are in motion. Techniques recently developed by Lawton [LAWT81, LAWT82a,82b,82c] have successfully extracted optic flow fields and corresponding environmental depth information from a variety of complex images without recourse to point-to-point matching in image sequences. They are more robust than currently available methods since they are not sensitive to problems caused by occlusion due to camera or object motion. We have also been investigating hierarchical computation schemes applicable to motion analysis. These may be of importance in eventual real-time implementation of a wide class of motion algorithms [GLAZ82].

Overton, Arbib, and Riseman's work on tactile sensation reflects our interest in methods for integrating multi-modal sensory data into a form useful for robot manipulator control. This work is focussed on the development of tactile sensors and on the construction of surface descriptions and object representations for manipulation tasks [OVER81]. A prototype tactile array sensor (currently 16x8 active elements) has been developed by Overton in cooperation with Digital Equipment Corporation. This sensor is being installed on a manipulator attached to a cartesian arm donated by General Electric. Control of the arm will be via a distributed network of processors supplied by DEC. This is closely integrated with Arbib's work on visuomotor coordination [ARBI75,78,81c], with Lesser's work on distributed control architectures, and with the work on VISIONS. This work is leading towards substantial industrial involvement in the development of a comprehensive robotics laboratory and a partial focus on problems of industrial and commercial importance.

The work of Foster [FOST76, WAL80, BON81, WEE82] on content addressable memories (CAMs) and associative processors is directly related to a concern of our research effort: the suitability of current von Neumann architectures for image processing tasks. In these tasks (particularly when motion is involved), massive amounts of low-level sensory data must be processed, often repeatedly. Foster has recently proposed the hardware development (using VLSI technology) of a CAM as a type of parallel array processor that is characterized by a single instruction stream operating on multiple data streams. This architecture matches nicely the conceptually parallel structure of the low-level processing techniques developed in the group.

As an example of the advantages of this architecture for image processing, a 3x3 convolution applied to a 512x512 image takes roughly 100 microseconds and we have a rough estimate that we can resolve the parameters of motion (using Lawton's techniques) from approximately 256 recognizable image points in under 50 milliseconds or at the rate of about 20 frames per second [LAWT82d]. This means that an appropriately configured CAM processor will provide the parameters required for real-time navigation at about video rates.

An area of collaborative research, combining the research interest of Elliott, Hanson, Hill, Reynolds, and Riseman, is in the analysis of remotely sensed imagery, particularly satellite imagery. The work of Elliot [ELLI80, 81, 82a, 82b] on processing images with low signal to noise ratios (using estimation techniques) directly confronts the problems of information extraction from multiple sensors, in the presence of large amounts of observation noise, in some optimum or near optimum manner. It is directly related to efforts in the VISIONS group to reliably measure features in images as a prelude to further processing. Hill has been examining methods for the progressive transmission of image data over narrow bandwidth channels. Such techniques promise to provide users of remotely sensed images with more efficient methods for examining images and spotting key information quickly, before the entire image has been transmitted in detail. A somewhat related effort is Hill's work on efficient storage and display of maps of on-board aircraft navigation systems. The emphasis of Reynold's work is on matching images to maps and on the development of an expert knowledge-based system for analysis of remotely-sensed data.

VI. DETAILED RESEARCH DESCRIPTIONS

VI. 1. Distributed Systems

VI. 1. 1. Coordination in Cooperative Distributed Problem-Solving Systems

V. Lesser is the principal investigator of this project with J. Wileden as an associated faculty and D. Corkill as a Research Associate. This project is being funded by joint NSF (MCS-8006327) and DARPA (NR049-041) grants totaling $385,000 for the two-year period ending September 30, 1982; continuation of DARPA funding for an additional three years has been approved and a proposal has been submitted for continuation of NSF funding.

Cooperative distributed problem solving systems are distributed networks of semi-autonomous nodes that are capable of sophisticated problem solving and cooperatively interact with other nodes to solve a single problem. Our research has emphasized designing such networks for applications in which there is a natural spatial distribution of information and processing requirements among the nodes but insufficient information for any node to make completely accurate processing and control decisions without an unacceptably excessive amount of internode communication [LESS80, LESS81].

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation [CORK80, CORK82b, SMIT81]. This problem is difficult because limited internode communication (which restricts each node's view of network problem solving activity) and network reliability issues (which require that the network's performance degrades gracefully if a portion of the network fails) preclude 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.

Our current research included the development of a node architecture capable of such sophisticated local decisionmaking [CORK82a, LESS82]. 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. Our new research 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:

1. organizational self-design;
2. distributed load-balancing;
3. negotiation among nodes;
4. local planning of internode communication strategies.

We are also planning to work on distributed diagnosis of network hardware and problem solving errors, on distributed debugging, on increasing the sophistication of the local planner of a node, and on developing an understanding of the effectiveness of particular organizational structures in different task environments.

This research will provide concrete information about the nature of coordination required for effective distributed problem solving system performance where there are large numbers of nodes operating in a dyamically changing environment. The understanding gained from this research will also have an important impact on research in knowledge-based Artificial Intelligence systems and large distributed processing systems.

As part of a broader study of strategies for distributed computation, Wileden has been developing tools to support the debugging of distributed systems. One of these tools, called the Event Description Language (EDL), permits a user to construct and monitor a high-level behavioral abstraction of the concurrent activity in the system [BATE81]. We are currently constructing a version of this tool for use in the interpretation testbed.

VI. 1. 2. Concurrency Control and Recovery Algorithms for Distributed Transaction Processing Systems

W. Kohler is the principal investigator of this project with D. Towsley and J. Stankovic as co-principal investigators. This project is currently funded by the Computer Engineering program of NSF. Preliminary work was supported by the Corporate Research Group, Digital Equipment Corporation, Hudson, Massachusetts, where Kohler is a consultant.

The problems of concurrency control of access to distributed databases and of the consistent recovery of these databases in case of hardware or software failure are well known and many algorithms have been proposed as "good" solutions [KOHL80,81]. However, faw have been implemented in real systems and next to nothing is known quantitatively about their relative performance. In fact, there is currently little experimental evidence to support the hypothesis that the choice of concurrency control or recovery protocols has a major impact on the overall performance of a transaction processing system, or how that impact compares with other design choices, e.g., database buffer size or allocation of database files to hosts. Because of this, system designers cannot effectively choose from among the proposed algorithms. We intend to evaluate and compare individual algorithms in a real environment for different workload and system characteristics.

More specifically, the objectives of this research are to understand and quantify the impact of different concurrency control and recovery mechanisms on the performance of real distributed transaction processing systems. By understandng the design tradeoffs of known mechanisms and developing improved mechanisms where possible, the projected results of this research can be used by distributed system designers to improve the performance of their systems.

To achieve these objectives, we plan to address three subgoals, which are briefly stated below.

1. Develop both a methodology for evaluating the performance of distributed transaction processing systems and a tool for exercising the systems and gathering performance data.

2. Using the methodology and tool developed above, evaluate and compare the performance of a distributed transaction processing "testbed" system [GARC81] as a function of workload and system design choices.

3. Perform experiments to validate or invalidate proposed analytic models of subsystem behavior. Models will be validated by monitoring the "testbed" system and comparing the results with those predicted by the models. Once validated they can be used as part of the methodology for designing new systems.

The CARAT distributed processing testbed [GARC81] that is being used for this research is currently operating on a local area network located at DEC in Maynard, Massachusetts, and some preliminary results have already been obtained [KOHL]. Because of the experimental nature of the testbed and associated hardware, it is very difficult to run this testbed from a remote site. Thus, the ability to duplicate this testbed on MINT would significantly speed up the time to design and evaluate testbed experiments.

VI.1.3. Adaptive Neural Networks

D.N. Spinelli is the principal investigator of this project with A. Barto as co-principal investigator. This project is being supported by AFOSR contract F33615-80-C-1088 ($383,909 for the three-year period ending June 1983).

This project is addressed to the study of pattern recognition perception and learning, both in animals and machines. It must be emphasized at the start that

1) even the "simplest" vertebrate brain is a structure of awesome complexity, thus progress in uncovering function and structure of neural circuits has gone hand in hand with the technological and computational capability available to the experimenter,

2) modelling of neuron-like circuits, while not dependent in principle on computers, has in fact produced serious dynamic realizations only with the advent of computers.

Even so, neuron-like network simulations to date, <u>because of computational limits</u> both in terms of speed and power, are made up of an unrealistically small number of elements and of unrealistically simple elements. For example, the classic binary model (the McCulloch-Pitts neuron) has proved inadequate for the analysis of neurophysiological data on membrane potentials and spike trains. Instead, we must model each neuron by a nonlinear differential equation. Moreover, in studies of adaptive networks, individual synapses each require complex models involving many variables to keep track of both the short-term and long-term processes that are involved in changes in synaptic effectiveness.

Unlike earlier adaptive network studies, our approach begins with adaptive elements that are much more complex than elements studied in the past. We have developed adaptive elements that are capable of learning to solve, by themselves, nontrivial problems involving the control of unstable dynamical systems. These learning problems are difficult due to the low specificity and high uncertainty of the feedback available to evaluate system performance. Learning under these conditions requires sophisticated algorithms for generating trials and assigning credit. In experimenting with elements of this complexity we are attempting to develop adaptive elements having the minimal complexity permitting them to function as interconnected, cooperating nodes of networks capable of surpassing the learning capabilities of their components. The early adaptive network studies greatly underestimated the complexity of the tasks that individual adaptive elements face as components of networks. Our goal is to demonstrate that the resulting type of distributed adaptive system can be useful for solving problems that do not lend themselves either to classical mathematical control theory or to knowledge-based artificial intelligence approaches. We have demonstrated some of these capabilities using computer simulations of route findings in spatial environments [BART81, 82a] and of the pole-balancing control problem [BART82b].

VI. 1. 4. Decentralized Control of Job and Process Scheduling in Distributed Processing Systems

J. Stankovic is the principal investigator of this project. This project is funded by NSF grant MCS-8104203 and an Army CECOM grant DAAB07-82-K-J015 for a total of $169,629 for the three-year period ending September 30, 1984.

This research concentrates on a specific class of distributed processing systems [JENS78, LELA77, STAN78] that are in the early research stage. These systems are characterized by decentralized system-wide control of resources for the cooperative execution of application programs [CASE77, STAN80, OUST80]. Moreover, these systems must operate with missing, incomplete, delayed, or erroneous state information. Specifically, the primary research effort of this proposal is to develop, compare, and evaluate job and process scheduling algorithms [BOKH79, CHOW79, STON77, STON78a] for such a demanding environment. Heuristics based on statistical decision theory and stochastic learning automata are being employed. Testing is being done by extensive simulation studies. A formal model of decentralized control has also been developed and used in conjunction with the simulation results.

The potential advantages of this research are to reduce the amount of centralized control required in a distributed system, thus eliminating critical entities (resources), and provide better overall performance. Furthermore, due to the elimination of critical resources and the development of algorithms that can operate in the presence of missing or erroneous data we hypothesize that improved reliability is also possible.

VI.1.5. Control and Analysis of Distributed Computations

H. Stone is the principal investigator of this project. This project is funded under NSF grant MCS-8106191 for a total of $123,802 for the two-year period ending June 30, 1983. This grant funds continuation of research initially supported under NSF grant MCS-7805295.

Techniques for the optimal control of distributed systems have been treated in the literature in a theoretical context, but in most cases have not been reduced to practice. There have been some notable exceptions to this statement in routing control for communication networks [GERL78, RUDI76]. Several researchers have examined the question of task assignments to processors in distributed systems [LAWS78, GYLY76, STON77, BOKH79, CHU69, CHU80, MORG77]. This is the problem of assigning data blocks and program execution to specific sites in a distributed system, and it is a crucial factor in the performance of the system. While the approaches cited here range over a large number of techniques, none of the writers claim that their method is suitable in its present form for real-time control of distributed systems. There appears to be a significant gap between what control algorithms have been formulated and what is needed in practical implementations.

Prior research has developed simple control algorithms based on commodity-flow graphs for controlling distributed computations [STON77, 78b]. Attempts to extend the simple models in various ways have shown that the extensions are typically of the "hard problem" (NP-complete) class, for which no efficient algorithms for finding solutions in the worst case have been discovered. Recently we have discovered that random graphs lead to notably simple optimal control strategies because of a strong clustering property.

We are currently investigating the statistical properties of distributed programs to determine if they exhibit the strong clustering property of random graphs. If so, we believe that near-optimal real-time control of distributed processes is achievable in spite of the NP-complete nature of the problems on worst-case situations. The research will attempt to identify the statistical behavior of distributed computations through mathematical analysis and simulation. These results will be validated on a small testbed of a distributed system. If the results show a strong clustering property much like the working set phenomenon in virtual memory systems, then a practical algorithm can presume the form of a final solution immediately, and perturb the system slightly around this assumed form to find good near-optimal controlled states.

VI.2. Information Systems

VI.2.1. Expert Systems and Natural Language

Rissland, McDonald, Lehnert and Moll are the main faculty involved in the research on expert systems. They are supported by grants from the NSF and the Sloan Foundation totalling more than $150,000 annually.

McDonald is the principle investigator of the project entitled "Natural Language Interfaces to Expert Systems" which was begun during the spring of 1981 with support from NSF (IST-8104984 with M. Arbib) and the Sloan Foundation (as part of the Cognitive Science Training Program). The project has two foci: (1) basic research on the process of natural language production, and (2) the application of this research to provide output interfaces for other COINS projects, particularly example generation, VISIONS, and Croft and Lesser's work on scripts, enabling us to carry on joint research on the explanation of expert reasoning and communication via natural language.

The basic research is a continuation of the work of [MCDO80] and [CONK79]. It breaks new ground in the areas of planning rhetorical goals and of interactions between conceptual and linguistic knowledge and conventions during processing. Linguistic rigor is strongly emphasized by employing theoretically motivated grammars and by paying close attention to the available psycholinguistic data; this has lead to very close ties with the Linguistics

Department.

Lehnert's previous research in natural language processing includes a control structure for question answering processes, designs for integrated memory modification, and various strategies for memory representation [LEHN78, 82b]. Her recent theoretical work has been on the problem of narrative text summarization [LEHN82a]. It appears that summaries can be naturally generated from a level of memory representation based on plot units, which encode the structural features of the text in a cyclic graph stucture. She has implemented a program to generate plot unit graphs, and a number of experiments have found evidence of psychological validity for this approach [LEHN81].

At the same time, she has also been involved with the development of natural language interfaces for existing data bases, so that users who have no technical knowledge of computers or query languages can access information using natural language queries. Her work with natural language query systems has been conducted in the domains of well data for oil exploration and inventory control for business management.

The research of Rissland is focused on three topics: (1) the structure of knowledge in complex domains like mathematics and law; (2) the structure, use and generation of examples, and (3) machine learning. She is interested in these topics with respect to issues such as the design of knowledge bases to capture expert knowledge for use in expert or tutoring systems, the role of examples in understanding and explanation, and the relationship between learning and teaching systems.

Currently, she is investigating the structure, use and generation of examples in law, linguistics, mathematics and computer science under a grant from the National Science Foundation (IST-80-17343, $147,000). This grant has also supported joint work by colleagues in the Linguistics Department under the auspices of the Cognitive Science program.

Rissland's past work has included research on the understanding of mathematics [RISS78], the design of intelligent CAI systems (with Professor Elliot M. Soloway of Yale University under grant MDA903-80-C-0508 from the Army Research Institute) and the tailoring of a conceptual framework for high school level mathematics (under grant NIE-G-80-0096). Her research group has built a system to generate examples meeting specified constraints [RISS80a,81a]. The CEG (Constrained Example Generator) system is being used not only to study the example generation process but also to study issues in learning [RISS81a,81b]; examples are critical to any system -- man or machine -- that learns and a system that generates examples ought to be able to learn how to improve its capabilities. One of Rissland's thesis students is using ideas from the CEG research to investigate the role of scenarios in tactical planning.

Rissland, together with McDonald and others, hopes to integrate a CEG-like system into a software environment and use it to illustrate explanations that the environment provides its users. Rissland and McDonald plan to build a facility that will be able to generate natural language explanations that incorporate illustrative examples. This work may also be applied to research in vision where new scenes can often be generated by modifying already existing scenes, and to medical and legal knowledge bases where the knowledge structure is complex and in which examples play a central role.

Moll is working with Digital Equipment Corporation on the development of an expert system for automatically configuring VAX-11/780 computer systems when supplied with a customer order. The most complicated and troublesome segment of DEC's current system (R1) is the unibus configuration. Moll's prototype system performs these configurations much faster and more accurately than R1. This prototype is currently being installed in the DEC production system.

Kulikowski is working on a prosthetic system which enables a person with extreme physical disabilities to control a standard microcomputer and make the decisions needed to achieve a dignified and independent lifestyle. The Speak-Easy Communication System is designed for both general and specific communication by physically disabled people who are nonvocal and almost completely paralyzed--that is, their muscle abilities are such that usual methods of communication like speech, writing, or signing are not feasible. The system develops a knowledge base about the client's environment and expected subjects of desired communication. This project joins techniques of artificial intelligence, that is the construction of computers to perform intelligent tasks, with the decreasing costs of microcomputers. The result will be an intelligent communication prosthesis that can be flexibly adapted to the individual needs of each client. The project involves the development of software for environmental control, drawing tools, curricula for various educational purposes, recreational devices for personal and social uses, and a graphics system for vocational programming needs.

VI. 2. 2. Office Information Systems

This subject is the basis of a large project funded by Digital Equipment Corporation (DEC) for a total of approximately $1,000,000. The principal investigators for the project are Croft and Lesser, although all of the faculty mentioned in the previous section are involved in various capacities.

There are two systems being developed under this project - TESS and POISE. The TESS (TExt Storage and Search) system is described in more detail in the next section. POISE (Procedure Oriented Interface for a Supportive Environment) provides a means of defining and supporting office activities in terms of procedures. Although office procedures can be difficult to specify [FIKE80], they have been used both to analyse offices and to partially automate higher level activities such as processing an order rather than activities tied directly to tools such as editing a document or sending mail [ELLI80, ZISM78]. The aim of POISE is to provide intelligent assistance, which means being able to recognize user actions, correct or warn about possibly incorrect actions based on system predictions, inform users about the current state of their actions (over a number of terminal sessions and procedures) and perform a number of actions (such as filing copies of forms) automatically [CROF82c]. To make the system as flexible as possible and to allow for exceptions to the "typical" procedure for an office activity, POISE includes a goal-based description of the activities as well as the procedure description.

The interface to POISE will be designed with the office worker in mind and will make extensive use of natural language and examples. The formalism used to describe the office procedures is based on the EDL language developed here by Wileden and Bates [BATE81]. One important problem that will be studied is the design of an interface to EDL that will make it easy for the user to define and modify procedures. The POISE approach is also being used to develop an intelligent software development environment [HUFF82].

VI.2.3. Text Filing and Retrieval Systems

In this area the main faculty member is Croft. His research is currently supported by an NSF grant (IST-8111108, $121,240) and the DEC grant just mentioned.

The work on text filing and retrieval centers on the design and implementation of strategies based on probabilistic models of word occurrences in text [CROF82b]. This approach leads to a very flexible user interface and search strategies which are more effective in locating relevant documents than a Boolean query/free text system. A recent project extended the probabilistic models of retrieval to include information about the frequency of occurrence of words within documents [CROF81b]. This model has been tested experimentally with large collections of documents and shown to significantly improve the effectiveness of retrieval [CROF82f].

The current NSF project deals with the design of a text retrieval system using an adaptive control strategy [CROF81a]. This system, which also incorporates a sophisticated representation of statistical information about the text documents, will automatically choose one of a variety of search strategies in order to get the best results for a particular user and query type. The control strategy is based on the adaptive network developed at the University of Massachusetts by Barto and others [BART81].

As part of the office information system project, Croft is working on a prototype system (TESS) for the filing and retrieval of documents in the business environment. This system must deal with a variety of document types, such as letters, memos, reports and electronic mail, which have different characteristics than the documents in typical bibliographic systems [CROF82e]. In addition to the file cabinet/drawer/folder facility used in many systems for personal filing, TESS will provide uniform and simple access by content to documents in any location. Other TESS facilities will include mail classification, indexing of offline and non-text material and selective dissemination of information. A first version of this system was developed as a COINS facility for handling personal and departmental bibliographies [CROF82d].

VI.2.4. Database Systems

The faculty in this area are Stemple, Carlson and Orenstein (in January 1983). Proposals for funding are being written.

Stemple's research is in the area of database management systems architecture, query execution optimization and human factors of database interfaces. The motivation for the different aspects of his work is the desire to find cost-effective techniques for building and maintaining large-scale data based models.

While a number of groups are working with the concept of database machines, Stemple has been mainly concerned with restructuring the way in which database schemas are processed, whether in host machines or backend processors. He proposed [STEM76] an approach in which schemas are translated into procedures in order to meet the opposing demands of high levels of data independence and efficient processing. This work was supported in part by a CDC research grant (MC1OCA). In this work, being taken up again after a four-year hiatus, he intends to exploit the proposed technique in developing a data abstraction approach to schema processing and to investigate ways in which this technique might be used in CAD/CAM and geographic information applications.

A second case of efficiency versus high-level specification deals with relational query execution optimization. This is a very important area which has received much attention. Numbers of valuable results and techniques have been reported, most notably in [YAO79, SELI79, SMIT75, BECK80]. Stemple is attempting to generalize and add to these approaches, while studying problems such as the sensitivity of certain heuristics to the accuracy of database statistics, cost functions which account for buffer overflow, and data description techniques to support query optimization. With Ramamritham he is studying the use of port-based operating systems architecture in query processing and as a basis for the interface between a query processor and applications programs.

Efficiency in processing queries can be of little use if the query language is too hard to learn and use. Stemple (and his Ph.D. student, Welty) have run experiments testing the learnability of two query languages differing in procedurability [WELT81]. This work was supported by NSF (grant MCS-78-07616). A proposal is currently being made to run futher experiments to find other ways of improving query languages. For example, he is looking into using the graphic feedback possible with an incremental, procedural language such as TABLET [REIS81].

Carlson's research is related to computer-aided instruction (CAI) and involves designing and evaluating database organizations and programming language tools for a CAI environment to be employed by the Air Force Systems Training Command. This effort is both developmental and research oriented, with the intent being to propose and analyze database organizations for specific CAI installations and more general-purpose applications. The Ada programming language is being evaluated for its suitability in a CAI environment, and database portability issues are also being addressed.

VI. 3. Software Development

VI. 3. 1. Tools Supporting Pre-Implementation Development of Flight Control Software

Wileden and Clarke are undertaking a project, funded by NASA (NAG1-115, 1/1/81 - 2/28/83, $134,721), aimed at enhancing pre-implementation software development support for flight control software. This project entails a preliminary design directed toward defining and producing a set of tools to support the requirements and design activities that are central to the pre-implementation phase of software development.

The difficulty of developing correctly-functioning software is a major concern in every area of computer applications, but it is particularly problematic for flight control software. Because numerous interrelated aspects of an aircraft's control system may demand attention simultaneously, synchronization of concurrent activities and stringent constraints on response times are an inherent part of flight control. Thus, such software naturally incorporates concurrency and real-time constraints, two features that significantly complicate the task of producing reliable software. In a very real sense then, development of flight control software typifies the enterprise of software development.

We have identified five topics that we believe to be of primary importance for the design of tools supporting pre-implementation development of flight control software. These topics are:

. Concurrency and synchronization
. Real-time constraints
. Exceptions and exception handling
. Formal specification methods
. Analysis techniques

We are currently focusing our attention on these five topics and conducting detailed studies of each. This effort is directed at discerning the major issues of importance with respect to each topic, discovering what relevant work has been done on each topic (whether or not that work was done in the context of tools for pre-implementation software development), and recommending what impact each topic should have on the design of pre-implementation tools.

VI.3.2. Design of Software for Distributed Systems

K. Ramaritham is the principal investigator of this project, which is being funded by NSF (MCS 82-02586, 4/1/82 - 3/31/84, $61,383).

In the development of distributed systems, two important issues are the interactions among the different processes and the distribution and management of shared resources [ROTH80, OUST80]. Devoting one process to manage each shared resource limits the type of interactions between processes and paves the way for modular design of distributed systems. This model for a distributed system is shared by other researchers as well [BRIN78, HOAR78, KELL78]. Given this model, resources and the processes that manage them play a critical role in the functioning of distributed systems.

This research is aimed at developing specification, design and implementation techniques for distributed system software. The focus is on the processes that manage shared resources. Topics that are being investigated include techniques for specifying the properties of shared resources and the processes that share the resources as well as strategies for synthesizing code for processes that manage the resources. In addition, techniques for verifying the correctness of extant mechanisms for resource management are also investigated in order to understand and develop the principles behind synthesis. Specific goals include:

1) Provision of a flexible specification language for the design of distributed systems. This language should enable the specification of shared resources, their managers and user processes. Fairness in serving requests from multiple processes should also be specifiable.

2) Developing algorithms for automating the construction of managers of the resources. Since correctness considerations are normally taken into account during the development of software, it should be possible to embed such considerations into a framework for the synthesis of code from the specifications obtained in the previous step [RAMA82a].

3) Developing mechanisms for verifying the correctness of distributed system software. Verification will also include the analysis of resource managers and resource users for possibilities of deadlock in the system. While previous approaches [ATK79, GRIF74] have mainly dealt with safety properties, our endeavour [RAMA81, RAMA82b] has been to consider both safety and liveness properties using the temporal logic formalism.

By focusing attention on the specification, veification and automated code generation of the server components, our attempts will be to improve the reliability of distributed systems. We use the formalism provided by temporal logic as the framework for our investigations. This makes it possible to treat both invariant and dynamic properties of software systems uniformly.

Currently the scope of our research is being extended to take into consideration protection of shared resources. This is being done in collaboration with D. Stemple. A proposal is being submitted to the NSF for conducting research in this area. Towards achieving protection of shared objects, we extend the use of ports, which typically have been utilized as communication channels between cooperating processes [RASH81]. For this purpose, a port is viewed as an abstract data type and protection is achieved by restricting the operations that are available to processes for manipulating ports. We are examining the design issues involved in a such a paradigm of communication and protection. In particular, we are investigating primitive port operations, port connection establishment, synchronization and cooperation among processes mediated by ports, message structure and failure recovery. This investigation should reveal those questions that need to be answered during the design of systems which use ports to achieve communication and protection. Based on this research, we intend to design a prototype system in which use of resources is synchronized by resource managers and protection is achieved through the use of ports.

VI.3.3. Toward Feedback-Directed Development of Complex Software Systems

In a project related to the NASA project described above, Clarke, Graham, and Wileden are requesting support from DARPA to investigate feedback-directed development of complex software systems [CLAR81], with particular emphasis on supporting the pre-implementation phases of software development. Specifically, we plan to develop appropriate languages for the specification and design phases of the life cycle and to consider improved methods for incorporating analysis into the development process. We intend to explore these issues guided by our goal of providing an integrated software development environment.

Where the few previous projects directed toward the pre-implementation phases have focused primarily upon notational issues, we intend to strongly emphasize analysis. Although some existing specification and design systems do offer certain analysis capabilities [CAIN75, TEIC77, HEBA79, ROBI77], the analysis has generally taken a secondary role and has not been used actively to guide the software development process. Analysis methods for most of these systems have tended to be relatively simple, consisting primarily of reformatting methods.

We believe that more powerful analysis capabilities are possible and that analysis should be an integral aspect of software development at every phase. We therefore plan to study the necessary modifications to a number of existing analysis techniques as well as develop new techniques, specifically for use during the pre-implemenation phase. Since most existing techniques (see [REIF77, ZELK78]) were developed to analyze programs, they are not directly applicable to more abstract notations. There is a strong similarity, however, between programs and the abstract representations applicable to specifications and designs. Therefore, many of these techniques such as data flow analysis [OSTE76] and symbolic execution [CLAR76] can be adapted for use with abstract representations to provide feedback on the consistency and completeness of the evolving design. Additionally, existing analysis techniques (e.g., [GRAH73]) must be made applicable to data abstraction, concurrency, and real-time language features. Some of the new techniques we are considering, such as algebraic techniques for analyzing concurrent systems and partition analysis for evaluating different levels of description of a system, are currently being investigated and are described further in this section.

VI.3.4. Algebraic Techiques for the Analysis of Concurrent Systems

In collaboration with G. Avrunin of the Mathematics Department, Wileden has developed algebraic techniques for the analysis of certain special classes of concurrent systems [AVRU81] based on a description of the systems in a suitable modelling scheme [WILE80]. Both the modelling scheme and these analytic techniques have been specifically tailored for use in the design stage of concurrent software system development. Wileden and Avrunin have recently submitted funding proposals for continued work in this area.

The descriptive notation being used as a basis for this work is the Dynamic Process Modelling Scheme (DPMS). This notation was originally developed for studying distributed systems with dynamic structure [WILE78]. It evolved from the PPML formalism [RIDD79] that served as the foundation for the DREAM software development system [RIDD78]. One component of DPMS is a modelling language, called DYMOL, that can be used to formulate precise, high-level, procedural descriptions of constituent processes in a distributed system [WILE80]. A second component of the modelling scheme, called <u>constrained expressions</u> [WILE82], is a closed form, non-procedural, representation for all the possible behaviors that could be realized by some distributed system. For an important subset of dynamically-structured distributed systems these two components of DPMS are related by an effective procedure for deriving the constrained expressions representing the potential behavior of a given system described in DYMOL.

For the purpose of analysis, the possible behaviors of a system modelled in DPMS are regarded as a set of strings of symbols representing events involving the internal computations of the component processes of the system and the transmission of messages between those processes. To analyze a design for a distributed system expressed in DPMS, we determine whether a particular symbol, or pattern of symbols, appears in a string representing a possible behavior of the system. The symbols in question may correspond to some desirable property of the system, such as graceful degradation, or may represent a pathology, such as deadlock.

Our analysis techniques begin with a collection of rules which are used to iteratively generate inequalities involving the numbers of occurrences of particular symbols that can appear in various segments of a string. These rules are based on the underlyng semantics of DPMS and on the particular symbols in question. If the assumption that a certain pattern of symbols occurs in a string representing a possible behavior leads, at any stage of the iterative process, to an inconsistent system of inequalities, we have reached a contradiction. We may then conclude that our assumption is incorrect and the given pattern does not occur in a behavior. Otherwise, we continue to generate inequalities until we have enough information to construct a behavior use.

We believe that this approach is well-suited for use in a systematic, iterative refinement style of distributed software system development. It facilitates the production of the incomplete and abstract descriptions that are appropriate during early stages of the development process. Moreover, it provides a means for rigorously analyzing these incomplete and abstract descriptions. Thus, it offers the prospect of a development process guided, from its earliest stages, by continual assessment of the evolving design. Such a carefully guided development process could dramatically increase the productivity of distributed software system developers.

We are presently working on increasing the applicability of the proposed modelling scheme by extending the analytic techniques to a wider range of systems and developing the collection of techniques into an effective procedure for determining possible system behavior. In addition, questions related to the efficient implementation of these techniques as a design aid and certain associated theoretical issues are being considered.

VI.3.5. The Partition Analysis Method to Demonstrate Program Reliability

Clarke and Richardson are the investigators in this project funded by NSF (MCS 8104202, 7/1/81 - 6/30/83, $81,723). This project was an outgrowth of previous work by Clarke, supported by AFOSR, to design, develop, and evaluate a symbolic execution and test data generation system, called ATTEST. ATTEST [CLAR76,78] is an experimental system that can either augment user-selected test data or indepenently generate test data sets. Although ATTEST will not guarantee correctness of a program, it does surpass current testing techniques in that it offers a systematic method of selecting test data to achieve a specified measure of program coverage. Work on ATTEST has provided an impetus for general work on testing, including an evaluation of domain testing [CLAR82a], a comparison of and extension to symbolic evaluation techniques [CLAR81b,c], and the development of the partition analysis method [RICH81].

The partition analysis method assists in program testing and verification by evaluating information from both a design and an implementation. This method employs symbolic evaluation techniques to partition the set of input data into subdomains, where the elements of each subdomain are treated uniformly by the design and processed uniformly by the implementaton. This partition divides the problem domain into more manageable units. Information related to each subdomain is used to guide in the selection of test data and to verify consistency between a design and an implementation. Moreover, the test data selection process, called partition analysis testing, and the verification process, called partition analysis verification, are used to enhance each other, and thus increase program reliability.

We are currently working on extending the partition analysis method to a wider class of problems, improving and formalizing the analysis techniques involved, and experimentally evaluating the method's effectiveness in improving program reliability. A preliminary evaluation of the method [RICH82] provided very favorable results on the effectiveness of partition analysis. Moreover, this study has lead to some interesting insights on the feasibility of integrating several testing strategies, such as domain and computation testing. We intend to continue exploring methods for effectively integrating testing strategies.

VI.3.6. <u>A Testbed for Evaluating Software Testing Methods</u>

S. Zeil is the principal investigator of this project, which is funded by NSF (MCS-8210084, 6/1/82 - 5/31/84, $39,917).

Research into program testing over the last several years has suggested new testing strategies, many of which involve automation of at least part of the testing procedure. Few of these methods, however, have been accepted into general practice. One reason for this may be the difficulty of evaluating and comparing the power and utility of these widely differing strategies. Indeed, there is no clear consensus as to what criteria should be employed for such comparisons. Theoretical definitions of correctness are so much stronger than most proposed testing methods that they do not serve to differentiate between opposing methods.

Empirical studies of testing strategies have not been without their problems also. Besides the problems of finding appropriate test programs, the choice of evaluation criteria that would be simultaneously applicable, informative, and theoretically defensible has proven difficult. A new measure, which offers significant improvements in these aspects, is based on the number of <u>blindness expressions</u> left behind by a test method. Blindness expressions arise from the assignment statements and conditional expressions encountered during execution. If, for example, along some path up to a program statement the variable X has been assigned a value f(Y), then the expression "X - f(Y)" must evaluate to zero. Consequently, "X - f(Y)" can be added to any expression in that program statement where it is permitted by the syntax of the language.

It can be shown that, frequently, the set of all possible errors that would escape a set of test data can be formed from combinations of a finite number of blindness expressions. The number of expressions required for this purpose can be used as a measure of the size of the set of possible errors that could be missed using the chosen test data. Because the blindness expressions arise from the functions assigned to each variable, this measure effectively represents the degree to which the test data exercises all distinctively different computations available in the program.

This project involves the construction of a testbed system to compute the blindness expressions of test data for FORTRAN programs. This system will be based upon Clarke's ATTEST symbolic execution system [CLAR78]. In addition to performing the symbolic execution of the FORTRAN programs being tested, the manipulation capabilities of ATTEST will be used to set up and simplify the systems of equations required to find the blindness expressions. These very large, ill-conditioned systems will then be solved to obtain the final measure of test data effectiveness.

This testbed will be employed to empirically study the relative power of a number of proposed testing strategies. The proposed measure will permit a more uniform evaluation of a wider range of testing strategies than has been possible in previous studies. Such an experiment should not only provide evidence of which methods are stronger, but also give some insight into what types of errors and program characteristics make testing more difficult.

VI.4. Vision and Robotics

VI.4.1. VISIONS: Segmentation and Interpretation of Natural Scenes

Co-Principal Investigators: A. Hanson and E. Riseman
Funding Period: terminates February 28, 1983
Average Funding: $173,000 per year up to 9/1/82
currently $50,000 per year plus major AFOSR grant pending
Agencies: ONR, Tufts New England Medical Center, DEC, Univ. of Mass. Biomedical Research Grant, and Massachusetts General Hospital

The VISIONS Image Understanding System is an experimental testbed for investigating the construction of integrated computer vision systems. The goal is to provide an analysis of color images of outdoor scenes, from segmentation through the final stages of symbolic interpretation. The output of the system is intended to be a symbolic representation of the three-dimensional world depicted in the two-dimensional image. This involves the determination of object labels for major image regions and an approximate placement of objects in three-dimensional space. The theme of this research is discussed in the set of papers appearing in Computer Vision Systems (Academic Press, 1978) [HANS78c].

Our research efforts in segmentation and interpretation have been (somewhat artificially) separated into several major efforts. A low level segmentation system is responsible for decomposing the original image into easily manipulatable visual primitives, such as regions and edges, and the extraction of visual features, such as color, texture, size, shape, etc. An interpretation is created by grouping the visual primitives in the appropriate ways and linking them to semantic labels under the constraints imposed by world knowledge contained in Long Term Memory (LTM). This process is accomplished by applying sequences of knowledge sources (KSs), which are modular processes governing the transformation of data between particular levels of representation. The KS application takes place under the guidance of a control strategy, and extends a partially constructed interpretation resident in Short Term Memory (STM).

Segmentation

Reliable segmentation is one of the major problems in scene analysis and has confounded traditional approaches in the field. Over the last several years a variety of image segmentation algorithms have been developed in the VISIONS group. In particular, three major approaches to low-level processing are currently being developed and employed:

1. a boundary formation process which aggregates local edges into continuous boundaries based on local contextual cues [PRAG80, HANS80a];

2. a region formation process based upon detection and labelling of peaks (or clusters) in color and texture feature distributions [NAGI79b, NAGI81a, NAGI81b, NAGI81c, NAGI82, KOHL82a];

3. a surface segmentation which provides distances to environmental surfaces, based on motion analysis [PRAG79, WILL80, WILL81, GLAZ81, LAWT81a, LAWT81b, LAWT82a, b, c].

All of these algorithms were developed them within a simulated parallel architecture, called the "processing cone", which is hierarchically organized into a layered system. An extensive image operating system has been implemented around this (simulated) architecture on our VAX 11/780 [KOHL82a,b].

Interpretation

There are few projects which have attempted to deal in a general experimental manner with the complexities inherent in the interpretation of natural scenes [HANS78b]. In any system where real-world sensory information is being manipulated, the problems of ambiguous, errorful, and incomplete data arise. One goal of the research in image interpretation is the construction of more reliable systems by integrating the processing from multiple sensory modes (as in vision and tactile), from analysis of different aspects of the visual data (such as color, texture, and motion), and from overlapping local contextual cues. This has led us to develop distributed system architectures, composed of many different kinds of knowledge sources (processes), within which intermediate results from using different knowledge or constraints of the physical world can be integrated into a consistent model. This research heavily overlaps the issues that Lesser has examined in his work on the Hearsay speech system [ERMA77, ERMA80] and more recently in distributed problem solving.

Our research on image interpretation is concerned with methods and techniques for aggregating the preliminary environmental region, boundary, and surface sensory data into a structured description of the scene. The research questions being addressed here are those of <u>knowledge representation</u> and <u>use</u>; these may be loosely categorized as:

1. <u>Knowledge-directed Interpretation via Schemas [PARM80, WEYM81, WEYM82]</u>. One of the principles of the schema model of perception is that perception is guided by expectation. We are concerned with how stored visual knowledge of various types and the physical constraints of the real world can be used to link the visual primitives extracted from the sensory data to high level semantic concepts.

2. <u>Control of the Interpretation Process [WESL82, WEYM81, WEYM82]</u>. We are also examining a range of <u>control strategies</u> for applying this information in the process of interpretation of visual and tactile data. Issues here include accessing relevant schemas based on prominent features, focus of attention mechanisms for selecting worthwhile portions of the sensory field for analysis, and the ways of decomposing knowledge hierarchically so that partial matching can be effective.

3. <u>Knowledge Source Development</u>. In order to build cooperating control strategies capable of interpreting images one must have a rich set of such KSs. Modular processes have been implemented for analyzing and matching 2D shapes, hypothesizing objects based upon the match of color and texture attributes of regions and stored object attributes, specific object matching routines based upon relative locations, analysis of perspective cues to place objects in space, matching of the size of stored objects with the hypothesized size of surfaces and volumes, shadow merging strategies, and resegmentation mechanisms based upon high-level analysis, among others. We are exploring schema control in the generation of hypotheses by these processes, as well as the use of inference networks [DUD76, LOW82, WES82] to provide a framework for distributed control as hypotheses are formed and partial models of the scene begin to evolve.

Recent applications of VISIONS include work in conjunction with Dr. Bernard Schwartz at Tufts New England Medical Center on quantification of diagnostic and disease history variables relevant to the treatment of glaucoma, Dr. Calvin Day of Massachusetts General Hospital on the prognosis of melanoma, and Dr. Robert Huguenin in Physics and Astronomy and the Center for Remote Sensing at the University of Massachusetts on processing of satellite imagery. We have also begun work with General Electric/Schenectady on industrial applications of robotics and motion and have an ongoing relationship with Digital Equipment Corporation in the general area of industrial robotics.

VI.4.2. Processing Dynamic Images from Camera Motion

Principal Investigator: E. Riseman
Co-Principal Investigator: A. Hanson
Funding Period: June 1, 1982 - May 31, 1984
Amount: $420,000
Agency: DARPA

Motion analysis has become one of the most important research areas in the development of computer vision systems. This reflects, in part, the fundamental nature of motion itself --- many of the physical events in the task domain of computer vision applications involve the observable motion of elements of the environment. Motion processing supplies important environmental information without the intervention or use of high-level semantic knowledge.

We are investigating several basic issues necessary to develop computer vision systems for terrestrial and airborne motion. From a sequence of complex images obtained from a sensor in motion, we have demonstrated the feasibility of determining the changes in the sequence of images and establishing a consistent environmental model over time for particular sensor motions [LAWT81a, LAWT81b, LAWT82a,b,c]. The key scientific issue being addressed is the recovery and effective representation of information concerning the physical environment, such as surface distance, extent, and orientation, relative to a moving observer.

We are using series of controlled environments in which to develop theories and experimentally test their applications. The <u>first</u> domain involves a computer vision system in terrestrial motion, as in the navigation of an automobile. In this case the sensor motion is initially constrained to a line, then a plane, and later to less constrained three-dimensional motion (as in driving over hilly terrain). The <u>second</u> involves airborne motion, as in the landing of an aircraft upon the runway of an airfield; here the sensor motion is also highly constrained, and there is a natural extension to arbitrary smooth motion.

The general issues being addressed are:

- determining environmental depth and velocity maps from the image sequences using a variety of image descriptors;

- structuring these depth and velocity maps into surfaces;

- determining the presence and describing the motion of independently moving objects;

- determining extremal surface boundaries, particularly occlusion boundaries, and predicting the behavior of such boundaries over time;

- representing the environment in a manner suited to real-time control of flying and driving behavior, utilizing the relative location of environmental surfaces, the direction of motion, and potential collisions;

- predicting image events from an established environmental model;

- developing a hierarchical representation of the motion of image and environmental events so that appropriate levels of abstaction can be chosen for different types of goals;

- understanding the differences in processing required for solving the start-up problem (i.e., initiating the analysis) and the updating problem as further image sequences are received;

- understanding the changes in surface highlights, shadows, and texture that are induced by motion;

- developing mechanisms to handle image registration of actual sequences;

- examining the effectiveness of algorithms for real-time motion processing on state-of-the-art hardware; this would include both array processors and display hardware;

- integrating motion information processing into a more complete image interpretation system such as the VISIONS system at UMass.

Vision research on dynamic images produces enormous computational requirements that make experimentation time consuming and cumbersome. Each image frame requires extensive computation and one can easily imagine two orders of magnitude increase in computational requirements over static scene analysis.

VI.4.3. Visuo-Tactile Coordination for Robot Control

Co-Principal Investigators: M. Arbib and E. Riseman
Funding Period: 3 years commencing January 15, 1982
Amount: $317,758
Agency: NSF with support from General Electric and DEC

We are now in the process of establishing a research facility for the study of robot sensing and control. In this research we intend to draw upon both our AI work in computer vision and in our brain-theoretic analysis of vision and the control of movement. The Intelligent Machines Lab directed by Joe Gibbons at General Electric's corporate research facility in Schenectady, New York has donated a cartesian manipulator, a group of servo amps, and an Anorad controller. DEC has provided the computing equipment needed for our work. We are presently combining all of this equipment to form a unified research environment.

The acquisition, processing, and use of sensory information by humans performing delicate manipulative tasks remains quite complex compared to the capabilities of current robots. We intend to experiment with integrated visual and tactile guidance of robot arms and hands, and to construct theories and algorithms describing the use of sensory information in adaptive control structures. We have just completed construction of a tactile sensor with the cooperation of the industrial automation group at DEC and have already begun to process force data derived from this sensor [OVER81].

The areas to be studied during this research include: static and dynamic vision, static and dynamic tactile sensation, the interactions of multimodal sensory data in the creation and maintenance of environmental models (schema-assemblages), and adaptive control structures using a knowledge-directed approach.

One major area of our work involves the use of machine vision for robot control. This research will build upon the low-level vision algorithms currently in use, particularly the bottom up segmentation algorithms within the VISIONS systems (cf. VISIONS, Section VI.4.1). The processing of dynamic visual images to determine the changes in a sequence of images allows the inference of a model of the environment [GLAZ81, LAWT81b, LAWT82a,b,c]. The environmental model in turn can be used to predict and interpret ongoing image transformations (c.f. motion research, Section VI.4.4).

The second major area involves tactile sensation [OVER81]. Given the internal model of an object, the exact location and orientation in space of a surface of the object can be derived from tactile sensor arrays. The array of force values may be thought of as comprising a set of control points for a surface patch [YORK80, YORK81a, YORK81b]. Our research goal is to develop effective algorithms for comparing a constructed surface patch with the representation of the surface of the object, thus allowing determination of the location and orientation of the object.

The final area concerns a schema approach to sensory control of a robot and the formation of an integrated perception and action cycle [ARBI75, 78, 81c]. We are examining the representation of the environment in relation to the robot, the updating of that representation as the robot moves and as new tactile and visual data are gathered, and the use of the representation by programs which control the movement. We will attempt to decompose functions into a family of simultaneously active processes, and model the formation and updating of the internal representation as a distributed process.

VI.4.4. Parallel Architectures

The Vision and Robotics group has maintained a commitment to the develoment of parallel algorithms for early processing of perceptual data [HANS80b]. This commitment is due to the large amounts of data which must be repeatedly processed in order to extract useful visual primitives and our belief that close to real time processing of perceptual data will require specialized parallel hardware. Concurrent with this work, Foster's group has been developing hardware realizations of content addressable memories (CAMs) [FOST76]. CAMs are a type of parallel array processor that is characterized by a single instruction stream operating on multiple data streams.

Over the past three years, Foster's architecture group has attained several important milestones:

1. A small (320 word) CAM has been developed using conventional discrete hardware technology, and a number of parallel algorithms for this experimental device have been developed [LEVI78, FOST79, WALL80, BONA81, HOUG82].

2. Expertise in CAD/CAM design of VLSI chips has been developed.

3. A small (4x4) VLSI CAM chip has been designed. This chip is being fabricated at Digital Equipment Corporation's new facilities in Hudson, Massachusetts.

4. A chip with sixty-four CAM cells arranged as an 8x8 grid has been designed. 4096 such chips on 64 printed circuit cards would produce a memory of 512x512 or a quarter of a million words. The design is such that increases in device density on a chip will allow us to increase word length and/or increase the number of cells per chip without major redesign [WEEM81]. Funding for actual fabrication of this chip is being sought.

While Foster's work is being driven by local interest in image processing and motion sensing, the design is very general and it will be useful in many applications. In particular, we expect our design to be useful for:

image identification in real time,

sonar and radar signal processing,

tactile image processing for robots,

representation of knowledge,

and motion analysis.

In particular, we estimate that our design will be able to perform a nine-cell-convolution of an entire picture in about one hundred microseconds. Since convolution is one of the operations basic to almost all picture processing and recognition schemes, this ability offers a potential speedup of image processing of up to 100,000.

In the area of motion detection we have a rough estimate that we can resolve the rotation and translation of up to 256 recognizable points in under 50 milliseconds or at a rate of 20 frames per second. This means that a CAM will allow real-time navigation at about video input rates [LAWT82d].

Specifically, Foster proposes to design and construct a Content Addressable Memory consisting of 2**18 cells of thirty-two bits each. These cells will be arranged on a square grid 512 cells on a side. Each cell will be able to communicate with its four immediate neighbors and with the central control unit. This memory will be attached to a VAX 11/780 computer system.

The General Electric Corporation has recently expressed an interest in helping to support this effort. Their interest lies in the utility of the proposed architecture to practical problems in industrial inspection and robotics. Thus, the CAM architecture effort provides a focal point for vision, motion, robotics, and industry. It also represents a serious effort on the part of COINS to promote VLSI activities within the department.

VI.4.5 Remote Sensing

Co-Principal Investigators: H. Elliott, A. Hanson, F. Hill, G. Reynolds, E. Riseman
Funding Period: Two years
Average Funding: $80,000
Agencies: AFOSR, ASTEC, ONR

An emerging focal point for the application of the research expertise and techniques developed by the Visions and Robotics group is in the general area of processing remotely sensed data, primarily satellite data. Portions of this work are being done in cooperation with Duncan Chesley, director of the Digitial Image Analysis Laboratory, and Dr. Richard Hugenin, Department of Physics and the Five College Radio Astronomy group, both at the University of Massachusetts.

The research goals with respect to remote sensing include:

1. development of analytic and knowledge based approaches to automatic geometric correction;

2. construction of expert systems for use by the remote sensing community;

3. development of techniques for the progressive transmission of hierarchically organized images to more closely approach real- time transmission over narrow bandwidth lines;

4. interfacing geographic information systems, digital terrain models, and remote sensing data; and

5. development of techniques for the extraction of information from low signal-to-noise images produced from multiple sensors.

Reynolds has been coordinating the remote sensing activities and has begun developing analytic and knowledge-based approaches to (semi-) automatic geometric registration and correction of satellite imagery. He is investigating the applicability of the techniques developed for motion analysis to this type of imagery. The advantage of this approach is that the motion techniques do not require point-to-point matching. Hence it is not necessary to ensure that the same points are selected in both images (i.e. ground-truth and image to be corrected). The steps involved are:

1. Selection of ground control points.

2. Matching of known ground control points of images in ground-truth database to given image.

3. Correct given image using polynomial warping techniques.

Methods for dealing with partially occluded images (e.g. cloud cover, snow cover, etc.) and images exhibiting regular seasonal variations are being developed.

Elliott [ELLI80,81,82a,b] has been investigating techniques for processing and analyzing low signal-to-noise images such as those obtained from satellites. Problems with noise can arise because of low signal levels in the image, e.g. detection of thermal gradients in ocean currents, mineral deposits in regions with vegetation, military targets in clutter or cloud shadows, or it could be introduced by the sensor or communication channel.

These algorithms deal with boundary extraction and image segmentation and have been designed for use at low signal-to-noise ratios. The current emphasis of the work is the extension of these algorithms for use with remotely sensed imagery. We have had particular success in the use of Markovian process models, and maximum a-posteriori probability (MAP) estimation procedures.

Hill has been examining techniques for the progressive transmission of large images over narrow bandwidth channels. Progressive transmission, first proposed by Knowlton, is an alternative to standard scan line transmission schemes; it is a method for transmitting images in such a way that progressively better approximations emerge as the transmission proceeds. This has implications for the development of low-cost interactive workstations operating over standard phone lines.

An important limitation of the original scheme was the requirement for four look-up tables (whose size is a function of the gray level resolution at each pixel) for encoding and decoding the image. Satellite images often contain 16 bits per pixel, which would require each table to have over 8 billion entries, and would clearly prevent the use of this method. However, Hill has found a way to replace the tables with efficient procedures that require only 3 compares and 3 additions to effect each encode or decode operation. Efficient methods exist for roaming and zooming over large areas of an image.

VI.4.6. Neurophysiology of Vision

Principal Investigator: D.N. Spinelli
Duration: terminates June 30, 1983
Funding Level: $153,300 per year
Agencies: NIMH, AFOSR

Principal Investigator: M.A. Arbib
Duration: 1/1/79 - 12/31/83
Funding: $50,000 per year (in direct costs)
Agency: NIH

These projects are addressed to the study of pattern recognition, perception, learning, and visual control of movement both in animals and machines. There are three facets to the overall effort: neurobiological investigations into the development of the mammalian visual system, the modelling of adaptive networks, and neural net simulations of visuomotor coordination and control of locomotion (discussed under Distributed and Cooperative Computation).

Neurobiology of the Develoment of Behavior and Vision

Spinelli studies mammalian visual system from the retinal input to brain areas involved in visual processing. On-line use of computers for neural monitoring enabled us to produce highly precise maps of functional properties of single cells and to discover how higher levels in the system control lower levels. We discovered that functional properties of single cells in the visual cortex are a direct result of visual experience during development [SPIN72] and that experience shapes the functional properties not only of visual elements but also of sensory-motor ones [SPIN79]. These functional effects are the result of actual structural changes of how neural elements develop [SPIN80].

We have been investigating visuo-motor behavior, and how the function and structure of nerve cells in the sensory-motor cortex change. We find that visual feature detectors become permanently tuned to features of a danger stimulus, and that cells in the sensory-motor cortex become specialized to the motor behavior required of the kittens [SPIN81].

An important element in the way in which experience restructures neural circuits [VIAN81, JENS79] has to do with the clustering of dendrites, which in turn produces a functional clustering of related, important feature selective cells. A model of this mechanism [SPIN82] provides novel insights into how parallel processing, clustering, and experience driven self-organization take place in naturally intelligent systems.

Neural Mechanisms of Visuomotor Coordination

Since the inception of the department, we have had a program in neural modelling and Brain Theory, which has maintained strong connections with work in machine vision and in distributed and cooperative computation. As documented elsewhere, this work has provided a seedbed for our current work in robotics (VI. 4. 2) as an outgrowth of a concern with sensorimotor coordination, and in processing dynamic images (VI. 4. 2) as an outgrowth of studies of human and animal motion perception. Current work uses the frog brain as a focus for analyzing neural mechanisms in visuomotor coordination at a level of detail which allows design and simulation of detailed experiments in neurophysiology, neuroanatomy and animal behavior. We study cooperative computation of tectum with other brain regions in subserving such behaviors as facilitation and habituation, pattern-recognition, size-constancy, stereopsis, and detour behavior. Our initial modelling used 56 neurons, each characterized by a nonlinear differential equation. Current modelling uses a 2-dimensional model of tectum and pretectum with approximately 500 neurons. But this is still only a pilot system. As we add models of the retina, brainstem and other brain regions, and as we address experimental data with a finer grain of spatial resolution, we will have to go to networks of 10,000 neurons, described by a set of 10,000 simultaneous nonlinear differential equations.

VI.4.7. Distributed Computing in Vision, Brain, and Language

M. Arbib is the principal investigator of this project and is being supported by grants from NIH (No.5 RO1NS14971, a five-year project ending December 1983, totaling $239,000 in direct costs), NSF (No.MCS-8005112 with E. Manes on semantics of programming languages, $80,000 for two years, ending June 1982; and No. IST-8104984 with D. McDonald on dynamics of vision and language, $174,000 for three years ending June 1984), and Sloan* Foundation. This project emphasizes a form of distributed computing, called cooperative computation, in which the different processes are embodied in dedicated resources and cooperate to achieve some global goal.

The three "computation-rich" subprojects in this effort are:

1. Cooperative computation for dynamic machine vision, tactile processing, and sensory-motor coordination in robot control [ARBI78]. (This project is conducted in collaboration with the Hanson-Riseman VISIONS group.)

2. Cooperative computation in neural circuitry subserving visuomotor coordination and control of movement; models of neural development and plasticity [ARBI81a, 81b]. (Many of the simulation problems here are common with those of Barto and Spinelli's Adaptive Networks group.)

3. Cooperative computation in linguistics [ARBI79b]. One subproject (in collaboration with McDonald) studies the interaction of semantic representations of visual scenes with linguistic representations in the production of scene descriptions. A second models cognitive-syntactic interactions in language acquisition. A third studies the interaction of phonetic, syntactic, and pragmatic representations, using a simulation so structured that the effects of subsystem removal can be tested to analyze a variety of hypotheses about the way in which brain damage causes the various aphasias.

One emphasis of this research is on low-level cooperative computation, involving high parallelism of many copies of a few types of processes. The coordination technique among processes is based on the neural control mechanisms of excitation, inhibition, thresholding, and decay. There is also beginning work on high-level cooperative computation structures as a model for interaction among brain regions.

---

*The Sloan support is part of a large grant (No. 80-6-13: $490,000 for three years ending August 1983, spread over four departments with twelve co-principal investigators for a "Training Program in Cognitive Science" and related research.

VII. BIBLIOGRAPHY

[ARBI75] Arbib, M.A., "Artificial Intelligence and Brain Theory: Unities and Diversities," Annual Biomedical Engineering, 3, (1975), pp.238-274.

[ARBI78] Arbib, M.A., "Segmentation, Schemas, and Cooperative Computation," in Studies in Mathematical Biology, Part I (S.A. Levin, Ed.), MAA Studies in Mathematics, 15, (1978), pp.118-155.

[ARBI79a] Arbib, M.A. and Manes, E.G., "Partially-Additive Monoids, Graph-Growing and the Algebraic Semantics of Recursive Calls," Graph-Grammars and Their Application to Computer Science and Biology, Lecture Notes in Computer Science 73, Springer-Verlag, (1979), pp.127-138.

[ARBI79b] Arbib, M.A. and Caplan, D., "Neurolinguistics Must Be Computational," Behavioral and Brain Sciences, 2, (1979), pp.449-483.

[ARBI81a] Arbib, M.A., "Visuomotor Coordinaton: From Neural Nets to Schema Theory," Cognition and Brain Theory, 4, (1981), pp.23-39.

[ARBI81b] Arbib, M.A., "Rana Computatrix: An Evolving Model of Visuomotor Coordination in Frog and Toad," Machine Intelligence, 10, (D. Michie and J. Hayes, Eds.), Ellis Horwood, (1981), (in press).

[ARBI81c] Arbib, M.A., "Perceptual Structures and Distributed Motor Control," Handbook of Physiology, Section on Neurophysiology, Vol.III on Motor Control, (V.V. Brooks, ed.), The American Physiological Society, (1980).

[ATKI79] Atkinson, R.R. and Hewitt, C.E., "Specification and Proof Techniques for Serializers," IEEE Transactions on Software Engineering, SE-5, (January 1979), pp. 10-23.

[AVRU81] Avrunin, G. and Wileden, J., "Algebraic Techniques for the Analysis of Concurrent Systems," COINS Technical Report 81-11, University of Massachusetts, Amherst, (May 1981).

[BART81a] Barto, A.G., Sutton, R.S.; Brouwer, P.S., "Associative Search Network: A Reinforcement Learning Associative Memory," Biological Cybernetics, 40, (1981), pp.201-211.

[BART81b] Barto, A.G. and Sutton, R.S., "Landmark Learning: An Illustration of Associative Search," Biological

Cybernetics, 42, (1981), pp. 1-8.

[BART82a] Barto, A. G., Anderson, C. W., and Sutton, R. S., "Synthesis of Nonlinear Control Surfaces By A Layered Associative Search Network," Biological Cybernetics, 43, (1982), pp. 175-185.

[BART82b] Barto, A. G., Sutton, R. S., and Anderson, C. W., "Neuron-Like Adaptive Elements That Can Solve Difficult Learning Control Problems," COINS Technical Report 82-20, University of Massachusetts, Amherst, (in preparation), (1982).

[BATE81] Bates, P. C. and Wileden, J. C., "EDL: A Basis for Distributed System Debugging Tools," COINS Technical Report 81-7, University of Massachusetts, Amherst, (1981).

[BECK80] Beck, L. L., "A Generalized Implementation Method for Relational Data Sublanguages," IEEE Transactions on Software Engineering, (March 1980), pp. 152-162.

[BERN80] Berstein, P. A. and Goodman, N., "Fundamental Algorithms for Concurrency Control in Distributed Database Systems," Technical Report CCA-80-05, Computer Corporation of America, (February 1980), 254 pages.

[BOEH73] Boehm, B., "Software and Its Impact: A Quantitative Assessment," Datamation, 19, 5, (May 1973), pp. 48-59.

[BOKH79] Bokhari, S. H., "Dual Processor Scheduling with Dynamic Reassignment," IEEE Transactions on Software Engineering, SE-5, 4, (July 1979).

[BONA81] Bonar, J. G. and Levitan, S. P. "Real Time LISP Using Content Addressable Memory," International Conference on Parallel Processing, Bellaire, Michigan, (August 1981).

[BRIN78] Brinch, H. P, "Distributed Processes: A Concurrent Programming Concept," CACM, 21, 11, (November 1978), pp. 934-971.

[BURT75] Burt, P. J., "Computer Simulations of a Dynamic Visual Perception Model," International Joint Man-Machine Studies, 7, (1975), pp. 529-546.

[CAIN75] Caine, S. and Gordon, E., "PDL - A Tool for Software Design," Proceedings National Computer Conference, (1975), pp. 271-276.

[CASE77] Casey, L. and Shelness, N., "A Domain Structure for Distributed Computer Systems," Proceedings of the Sixth ACM Symposium on Operating Systems Principles, (November 1977), pp.101-108.

[CHOW79] Chow, Y.-C. and Kohler, W.H., "Models for Dynamic Load Balancing in a Heterogeneous Multiple Processor System," IEEE Transactions in Computers, C-28, 5, (May 1979).

[CHU69] Chu, W.W., "Optimal File Location in a Multiple Computing System," IEEE Transactions in Computers, C-18, 10, (October 1969), pp.886-889.

[CHU80] Chu, W.W., et al., "Task Allocation in Distributed Data Processing," Computer, 13, 11, (November 1980), pp.57-69.

[CLAR76] Clarke, L.A., "A System to Generate Test Data and Symbolically Execute Programs, IEEE Transactions on Software Engineering, SE-2, 3, (September 1976), pp.218-222.

[CLAR78] Clarke, L.A., "Automatic Test Data Selection Techniques", Infotech State of the Art Report, Software Testing, 2, Infotech International, (September 1978), pp.43-65.

[CLAR80] Clarke, L.A., Wileden, J.C., and Wolf, A.L., "Nesting in Ada is for the Birds," Proceedings of the Ada Symposium, Boston, (December 1980).

[CLAR81a] Clarke, L.A., Graham, R.M., and Wileden, J.C., "Thoughts on the Design Phase of an Integrated Software Development Environment," Proceedings of the Fourteenth Hawaii International Conference on System Science, Honolulu, (January 1981), pp.11-17.

[CLAR81b] Clarke, L.A. and Richardson, D.J. "Symbolic Evalution Methods for Program Analysis," in Program Flow Analysis: Theory and Application, S. Muchnick and N. Jones, eds., Prentice Hall, (1981), pp.264-300.

[CLAR81c] Clarke, L.A. and Richardson, D.J. "Symbolic Evaluation Methods -- Implementation and Applications," Computer Program Testing, Chandrasekaran and Radicchi, eds., North-Holland, (1981), pp.65-102.

[CLAR82a] Clarke, L.A., Hassell, J. and Richardson, D.J. "A Close Look at Domain Testing," IEEE Transactions on Software Engineering, (July 1982), pp.380-390.

[CLAR82b] Clarke, L.A. and Richardson, D.J. "Reliable Test

Data Selection Strategies — An Integrated Approach," Fourth Israel Conference on Quality Assurance, Israel, (October 1982).

[CONK79] Conklin, J.C., "A Scene Description System: Preliminary Survey," Master's Project, Department of Computer and Information Science, University of Massachusetts, Amherst, (1979).

[CORK80] Corkill, D.D., "An Organizational Approach to Planning in Distributed Problem Solving Systems," COINS Technical Report 80-13, University of Massachusetts, Amherst, (May 1980).

[CORK82a] Corkill, D.D., Lesser, V.R., and Hudlicka, E., "Unifying Data-Directed and Goal-Directed Control: An Example and Experiments," to appear in Proceedings of the Second National Conference on Artificial Intelligence, (August 1982).

[CORK82b] Corkill, D.D., "A Framework for Organizational Self-Design in Distributed Problem Solving Networks," Ph.D. Thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, (September 1982), (in preparation).

[CROF81a] Croft, W.B., "Incorporating Different Search Models Into One Document Retrieval System," Proceedings of the 4th International ACM SIGIR Conference, Berkeley, (1981).

[CROF81b] Croft, W.B., "Document Representations in Probabilistic Models of Retrieval," Journal of the American Society for Information Science, (to appear).

[CROF82a] Croft. W.B. and Pezarro, M., "Text Retrieval Techniques for the Automated Office," Office Information Systems, N. Naffah, Ed., North-Holland, (1982).

[CROF82b] Croft, W.B,. "An Overview of Information Systems," Information Technology: Research and Development, (to appear), (1982).

[CROF82c] Croft, W.B. and Lefkowitz, L., "POISE - An Intelligent Assistant for Profession-Based Systems," COINS Technical Report, University of Massachusetts, Amherst, (1982)

[CROF82d] Croft, W.B. and Ruggles, L., "The Implementation of a Document Retrieval System," ACM/BCS Conference on Research and Development in Information Retrieval, Berlin, (1982).

[CROF82e] Croft, W. B., "Experiments with Automatic Text Filing and Retrieval in the Office Environment," SIGIR Forum, Vol. 16, (1982), pp. 2-9.

[CROF82f] Croft, W. B., "Experiments with Representation in a Document Retrieval System," Information Technology, (to appear), (1982).

[DEV75] Dev, P., "Perception of Depth Surfaces in Random-Dot Stereograms: A Neural Model," Internatonal Joint Man-Machine Studies, 7, (1975), pp. 511-528.

[DI3079] DI-3000, Users Guide, Precision Visuals, Inc., 250 Araphoe, Boulder, Colorado, 80302, (1979).

[DUDA76] Duda, R. O., Hart, P. Z., and Nilsson, N. J., "Subjective Bayesian Methods for Rule-Based Inference Systems," Proceedings of the National Computer Conference, (1976).

[ELLI80] Elliot, H., Cooper, D. B., Cohen, F., Reiss, L., and Symosek, P., "Stochastic Boundary Estimation and Object Recognition," Computer Graphics and Image Processing, (April 1980), (also a chapter in Image Modelling, Academic Press).

[ELLI81] Elliot, H. and Srinivasan, L., "An Application of Dynamic Programming to Sequential Boundary Estimation," Computer Graphics and Image Processing, Vol. 17, No. 6, (December 1981).

[ELLI82a] Elliot, H., Cooper, D. B., Cohen, F., and Symosek, P., "Implementation, Interpretation, and Analysis of a Suboptimal Boundary Finding Algorithm," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-4, No. 2, (March 1982).

[ELLI82b] Elliott, H, and Hansen, F. R., "Image Segmentation Using Simple Markov Field Models," Computer Graphics and Image Processing, Vol. 19, No. 4, (August 1982).

[ELLI79] Ellis, C. A., "Information Control Nets: A Mathematical Model of Office Information Flow," ACM Proceedings of the Conference Simulation, Modelling and Measurement of Computer Systems, (1979), pp. 225-240.

[ERMA77] Erman, L. and Lesser, V., "A Retrospective View of the Hearsay-II Architecture," Proceeding of the Fifth IJCAI, (August 1977), pp. 790-800.

[ERMA80] Erman, L., Hayes-Roth, F., Lesser, V., and Reddy, D., "The Hearsay-II Speech-Understanding System:

Integrating Knowledge to Resolve Uncertainty," Computing Surveys, 12(2), (June 1980), pp. 213-253.

[ESWA76] Eswaran, K.P., Gray, J.N., Lorie, R.A., and Traiger, I.L., "The Notions of Consistency and Predicate Looks in a Database System," Communications of the ACM, 19, 11, (November 1976), pp. 624-633.

[FOST76] Foster, C.C., Content Addressable Parallel Processors, Van Nostrand Reinhold Company, (publishers), New York, (1976).

[FOST79] Foster, C.C., Levitan, S., and Wall, R., "A Postage Stamp Database: An Exercise in Content Addressable Memory Systems for Novice User Query Languages," Department of Computer and Information Science, University of Massachusetts, Amherst, (unpublished), (1979)

[GARC81] Garcia-Molina, H., Germano, F., and Kohler, W., "Architectural Overview of a Distributed Software Testbed," Technical Report, Corporate Research Group, Digital Equipment Corporation, Maynard, Massachusetts, (August 1981).

[GERL78] Gerla, M. and Mason, D., "Distributed Routing in Hybrid Packet and Circuit Data Networks," COMPCON 78 Fall, IEEE Computer Society, (1978).

[GLAZ81] Glazer, F., "Computing Optic Flow," Proceedings of IJCAI-7, (August 1981), pp. 644-647.

[GRAH73] Graham, R., Clancey, G., and DeVaney, D., "A Software Design and Evaluation System," Communications of ACM, 16, 2, (February 1973), pp. 110-116.

[GRAY79] Gray, J.N., "Notes on Data Base Operating Systems," in Operating Systems: An Advanced Course, R. Bayer, R.M. Graham, and G. Seegmuller, editors, Springer-Verlag, (1979), pp. 393-481.

[GRIF74] Griffithms, P., "SYNVER: A System for the Automatic Synthesis and Verification of Synchronization Processes," Technical Report TR 22-74, Harvard University, (1974).

[GSPC79] GSPC, "Status Report of the Graphics Standards Planning Committee of ACM/SIGGRAPH," Computer Graphics, 13, 3, (August 1979).

[GYLY76] Gylys, V.B. and Edwards, J.A., "Optimal Partitioning of Workload for Distributed Systems," COMPCON Fall 76, IEEE Computer Society, (September

1976), pp.353-357.

[HANS78a] Hanson, A.R. and Riseman, E.M., "Segmentation of Natural Scenes," Computer Vision Systems, (A.R. Hanson and E.M. Riseman, eds.), Academic Press, (1978).

[HANS78b] Hanson, A.R. and Riseman, E.M., "VISIONS: A Computer System for Interpreting Scenes," Computer Vision Systems, (A.R. Hanson and E.M. Riseman, eds.) Academic Press, (1978).

[HANS78c] Hanson, A.R. and Riseman, E.M., Computer Vision Systems, Academic Press, New York, (1978).

[HANS80a] Hanson, A.R., Riseman, E.M., and Glazer, R.C., "Edge Relaxation and Boundary Continuity," Consistent Labeling Problems in Pattern Recognition (R. Haralick, Ed.), Plenum Press, 1980.

[HANS80b] Hanson, A. and Riseman, E., "Processing Cones: A Computational Structure of Image Analysis," in Structured Computer Vision (S. Tanimoto, ed.), Academic Press, New York, (1980). (Also COINS Technical Report 81-38, University of Massachusetts, Amherst, (December 1981)).

[HEBA79] Hebalkar, P. and Zilles, S., "TELL: A System for Graphically Representing Software Design," Proceedings COMPCON Conference, San Francisco, (1979).

[HIRS70] Hirsch, H.V.B. and Spinelli, D.N., "Visual Experience Modified Distribution of Horizontally and Vertically Oriented Receptive Fields in Cats," Science, 168, (1970), pp.869-871.

[HOAR78] Hoare, C.A.R., "Communicating Sequential Processes," Communications of the ACM, 21, (August 1978), pp.666-677.

[HOUG82] Hough, A.A., "Tune Recognition Using a Content Addressable Memory," Masters Project, Department of Computer and Information Science, University of Massachusetts, Amherst, (1982).

[HUFF82] Huff, K.E. and Lesser, V.R., "Knowledge-Based Command Understanding: An Example for the Software Development Environment," COINS Technical Report 82-6, University of Massachusetts, Amherst, (1982).

[JENS78] Jensen, D., "The Honeywell Experimental Distributed Processor -- An Overview of its

Objectives, Philosophy and Architectural Facilities," IEEE Computer, 11, 1, (January 1978).

[JENS79] Jensen, F.E. and Spinelli, D.N., "Early Experience Effect on Dendritic Bundles," Society for Neuroscience Abstracts, 5, (November 1979).

[KELL78] Keller, R.M., "Sentinels: A Concept for Multiprocess Coordination," Technical Report UUCS-78-104, University of Utah, (June 1978).

[KOHL80] Kohler, W.H., "Overview of Synchronization and Recovery Problems in Distributed Databases," Proceedings of COMPCON 80/Fall, (September 1980), pp.433-442.

[KOHL81] Kohler, W.H., "Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems," Computing Surveys, (June 1981).

[KOHL82a] Kohler, R., "Integrated Syntactic Knowledge for Segmentation," Ph.D. Thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, (in preparation), (1982).

[KOHL82b] Kohler, R. and Hanson, A., "The VISIONS Image Operating System," to appear in Proceedings of the Sixth International Conference on Pattern Recognition, Munich, Germany, (October 1982).

[KULI81] Kulikowski, S., Hawkins, B., and Riseman, E., "The Speak Easy Communication System for the Non-Vocal Severely Disabled," The First Annual Johns Hopkins Contest on Microcomputers for the Handicapped, (1981).

[LAWS78] Lawson, J.T., "Distributed Data Processing System Design — A Look at the Partitioning Problem," COMPSAC 78 Proceedings, IEEE Computer Society, (1978).

[LAWT79] Lawton, D.T. and Prager, J.M., "The Computation of Optic Flow and Its Use in the Control of a Mobile Robot," Workshop on Computer Analysis of Time-Varying Imagery, Philadelphia, (1979), pp.30-31.

[LAWT80] Lawton, D.T., "Constraint-Based Inference From Image Motion," Proceedings of the AAAI Conference, Stanford University, (1980), pp.31-34.

[LAWT81a] Lawton, D.T., "Optic Flow Field Structure and Processing Image Motion," Proceedings of IJCAI-81, Vancouver, British Columbia, (1981).

[LAWT81b] Lawton, D. T., "Processing Dynamic Images and the Control of Robot Behavior," Ph.D. Dissertation, Department of Computer and Information Science, University of Massachusetts, Amherst, (in preparation), (1981).

[LAWT82a] Lawton, D., "Processing Translational Motion," to appear in Computer Graphics and Image Processing, (1982).

[LAWT82b] Lawton, D., "Motion Analysis via Local Translational Processing," to appear in Proceedings of the IEEE Workshop on Computer Vision: Representation and Control, Rindge, New Hampshire, (August 1982).

[LAWT82c] Lawton, D., "Applications of Translational Motion Processing to Robotics," to appear in Proceedings of the SPIE Conference on Robotics and Industrial Inspection, (August 1982).

[LAWT82d] Lawton, D. and Steenstrup, M., "Analysis of CAM-Based Motion Algorithms," COINS Technical Report, University of Massachusetts, Amherst, (in preparation).

[LELA77] Le Lann, G., "Distributed Systems — Towards a Formal Approach," Proceedings of the IFIP Congress, Toronto, North Holland Publishing Company, (August 1977), pp.155-160.

[LEHN78] Lehnert, W., The Process of Question Answering, Lawrence Erlbaum Associates, Hillsdale, New Jersey, (1978).

[LEHN81] Reiser, B., Lehnert, W., and Black, J., "Recognizing Thematic Units in Narratives," Proceedings of the Third Annual Cognitive Science Conference, (1981).

[LEHN82a] Lehnert, W., "Plot Units: A Narrative Summarization Strategy," Strategies for Natural Language Processing, Lehnert and Ringle, eds., Lawrence Erlbaum Associates, Hillside, New Jersey, (1982).

[LEHN82b] Lehnert, W., "BORIS - An Experiment in In-Depth Understanding of Narratives," Journal of Artificial Intelligence, (in press), (1982).

[LESS80a] Lesser, V.R. and Erman, L.D., "Distributed Interpretation: A Model and Experiment," IEEE Transactions on Computers, Vol.C-29, No. 12, (December 1980), pp.1144-1163.

[LESS80b] Lesser, V.R. and Erman, L.D., "Distributed Interpretation: A Model and Experiment," IEEE Transactions on Computers, C-29, (December 1980), pp.1144-1163.

[LESS81a] Lesser, V.R. and Corkill, D.D., "Functionally Accurate, Cooperative Distributed Systems," IEEE TRansactions on Systems, Man, and Cybernetics, Vol.SMC-11, No.1, (January 1981), pp.81-96.

[LESS81b] Lesser, V.R., et al., "A High-Level Simulation Testbed for Cooperative Distributed Problem-Solving," COINS Technical Report, University of Massachusetts, Amherst, Massachusetts, (June 1981).

[LEVI78] Levitan, S.P. and Foster, C.C., "CAMEOS: A Content Addressable Memory Enhanced Operating System," Department of Computer and Information Science, University of Massachusetts, Amherst, (unpublished), (March 1978).

[LOME77] Lomet, D.G., "Process Structuring, Synchronization, and Recovery Using Atomic Actions," Proceedings of the ACM Conference on Language Design for Reliable Software, SIGPLAN Notices, 12, 3, (March 1977), pp.128-137.

[LOWR82] Lowrance, J., "Dependency-Graph Models of Evidential Support," Ph.D. Thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, (October 1982).

[MCDO80] McDonald, D.D., "Natural Language Production as a Process of Decision-Making Under Constraints," Ph.D. Dissertation, Massachusetts Institute of Technology, (1980).

[MORG77] Morgan, H.L. and Levin, K.D., "Optimal Program and Data Locations in Computer Networks," Communications of the ACM, 20, 5, (May 1977), pp.315-322.

[NAGI77] Nagin, P., Hanson, A., and Riseman, E., "Region Extraction and Description Through Planning," COINS Technical Report 77-8, University of Massachusetts, Amherst, (May 1977).

[NAGI79a] Nagin, P., Kohler, R., Hanson, A., and Riseman, E., "Segmentation, Evaluation, and Natural Scenes," Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing, Chicago, (1979), pp.515-522.

[NAGI79b] Nagin, P., "Studies in Image Segmentation

Algorithms Based on Histogram Clustering and Relaxation," COINS Technical Report 79-15, University of Massachusetts, Amherst, (September 1979).

[NAGI81a] Nagin, P.A., Hanson, A.R., and Riseman, E.M., "Variations in Relaxation Labelling Techniques," Computer Graphics and Image Processing, (to appear), (1981).

[NAGI81b] Nagin, P.A., Hanson, A.R., and Riseman, E.M., "Studies in Global and Local Histogram-Guided Relaxation Algorithms," submitted to IEEE Transaction on Pattern Analysis and Machine Intelligence, (1981).

[NAGI82] Nagin, P., Hanson, A., and Riseman, E., "Studies in Global and Local Histogram-Guided Relaxation Algorithms," IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-4, No.3, (May 1982), pp.263-277.

[NAGI81c] Nagin, P.A., Hanson, A.R., and Riseman, E.M., "Region Relaxation in a Parallel Hierarchical Architecture," Real-Time Medical Image Processing, (K. Preston, Ed.), Plenum Press, (1981).

[OSTE76] Osterweil, L. and Fosdick, L., "DAVE - A Validation Error Detection and Documentation System for FORTRAN Programs," Software - Practice and Experience, 6, (1976), pp.473-486.

[OUST80] Ousterhout, J.K., et al., "Medusa: An Experiment in Distributed Operating System Structure," Communications of the ACM, 17, 7, (February 1980), pp.92-105.

[OVER81] Overton, K.J. and Williams, T., "Tactile Sensation for Robots," Proceedings of IJCAI-81, Vancouver, British Columbia, (August 1981), pp.791-795.

[PARM80] Parma, C.C., Hanson, A.R., and Riseman, E.M., "Experiments in Schema-Driven Interpretation of a Natural Scene," Nato Advanced Study Institute on Digital Image Processing, (R. Haralick and J.C. Simon, eds.), Bonas, France; also COINS Technical Report 80-10, University of Massachusetts, Amherst, (1980).

[PRAG77] Prager, J.M., Hanson, A., and Riseman, E., "Extracting and Labelling Boundary Segments in Natural Scenes," COINS Technical Report 77-7, University of Massachusetts, Amherst, (May 1977).

[PRAG79] Prager, J.M., "Segmentation of Static and Dynamic

Scenes," Ph.D. Dissertation, COINS Technical Report 79-7, University of Massachusetts, Amherst, (1979).

[PRAG80] Prager, J.M., "Extracting and Labelling Boundary Segments in Natural Scenes," IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-2, (January 1980), pp.16-27.

[RAMA81] Ramamritham, K. and Keller, R.M., "Specifying and Proving Properties of Sentinel Processes," Proceedings of the 5th International Conference on Software Engineering, (March 1981).

[RAMA82a] Ramamritham, K., "Degrees of Fairness and Their Realization", manuscript, Department of Computer and Information Science, University of Massachusetts, Amherst, (submitted for publication), (May 1982).

[RAMA82b] Ramamritham, K., "Proof Techniques for Resource Controller Processes", COINS Technical Report, University of Massachusetts, Amherst, (submitted for publication), (March 1982).

[RASH81] Rashid, R. and Robertson, G., "Accent: A Communication Oriented Network Operating System Kernal", Carnegie-Mellon University, Department of Computer Science Technical Report, (April 1981).

[REED79] Reed, D.P., "Implementing Atomic Actions on Decentralized Data," Proceedings of the 7th ACM Symposium on Operating System Principles, (December 1979) (Abstract only. The full paper is to appear in a future special issue of CACM.).

[REIF77] Reifer, D.J. and Trattner, S., "A Glossary of Software Tools and Techniques," Computer, 10, 7, (July 1977), pp.52-60.

[REIS81] Reisner, P., "Human Factors Studies of Database Query Languages: A Survey and Assessment," ACM Computing Surveys, 13, 1, (March 1981), pp.13-31.

[RICH81] Richardson, D.J. and Clarke, L.A., "A Partition Analysis Method to Increase Program Reliability," 5th International Conference on Software Engineering, San Diego, (March 1981), pp.241-253.

[RICH82] Richardson, D.J. and Clarke, L.A., "On the Effectiveness of the Partition Analysis Method," IEEE Sixth International Computer Software and Applications Conference (COMPSAC 82), (November 1982).

[RIDD78] Riddle, W., Wileden, J., Sayler, J., Segal, A., and Stavely, A., "Behavior Modelling During Software Design," IEEE Transactions on Software Engineering, (July 1978), pp. 283-292.

[RIDD79] Riddle, W., "An Approach to Software System Modelling and Analysis," Journal of Computer Languages, (1979), pp. 49-66.

[RISE79] Riseman, E.M., "Mechanisms for Parallel Processing of Local Contexts with Uncertainty," draft of working paper, SRI, (January 1979).

[RISS80a] Rissland, E.L. and Soloway, E.M., "Overview of an Example Generation System," Proceedings First National Conference on Artificial Intelligence, Stanford, (August 1980).

[RISS80b] Rissland, E.L. and Soloway, E.M., "Generating Examples in LISP: Data and Programs," COINS Technical Report 80-07, University of Massachusetts, Amherst, Proceedings International Workshop on Program Construction, France, (September 1980).

[RISS80c] Rissland, E.L., "Example Generation," Proceedings Third National Conference of the Canadian Society for Computational Studies of Intelligence, Victoria, B.C., (May 1980).

[RISS81a] Rissland, E.L., Selfridge, O.G., and Soloway, E.M., "The Role of Experiences and Examples in Learning Systems," Proceedings Third Annual Meeting of Cognitive Science Society, Berkeley, California, (August 1981).

[RISS81b] Rissland, E.L., and Soloway, E.M., "Constrained Example Generation: A Tastbed for Studying Issues in Learning," Proceedings IJCAI-81, Vancouver, B.C., (1981).

[RISS81c] Rissland, E.L., "The Structure of Knowledge in Complex Domains," Proceedings NIE-LRDC Conference on Thinking and Learning Skills, Lawrence Erlbaum Associates, (in press), (1981).

[RISS82] Rissland, E.L., "Examples, Adaptation and Learning," to appear in Proceedings NATO ARI, Plenum, (1982).

[ROBI77] Robinson, L. and Levitt, K., "Proof Techniques for Hierarchically Structured Programs," Communications ACM, 20, 4, (April 1977), pp. 271-283.

[ROSE78] Rosenkrantz, D. J., Stearns, R. E., and Lewis, P. M., "System Level Concurrency Control for Distributed Database Systems," ACM Transactions on Database Systems, 3,2, (June 1978), pp. 178-198.

[ROTH80] Rothnie, et al., "Introduction to a System for Distributed Databases (SDD-1)," ACM Transactions on Database Systems, 5,1, (March 1980), pp. 1-17.

[RUDI76] Rudin, H., "On Routing and 'Delta Routing': A Taxonomy and Performance Comparison of Techniques for Packet-Switched Networks," IEEE Transactions on Communication, (January 1976).

[SELF81] Selfridge, O. G., Rissland, E. L. and Arbib, M. A., (eds.), "Adaptive Control and Ill-Defined Systems," Plenum, (in preparation), (1981).

[SELI79] Selinger, P. G., Astrahan, M. M., Chamberlin, D. D., Lorie, R. A., and Price, T. G., "Access Path Selection in a Relational Database Management System," Proceedings of ACM SIGMOD Internatonal Conference on Management of Data, Boston, (1979), pp. 23-34.

[SKEE81] Skeen, D. and Stonebraker, M., "A Formal Model of Crash Recovery in a Distributed System," Proceedings of the Fifth Berkeley Workshop on Distributed Data Management and Computer Networks, (February 1981), pp. 129-142.

[SMIT75] Smith, J. M. and Chang, P. Y. T., "Optimizing the Performance of a Relational Algebra Database Interface," Communications of ACM, 18,10, (October 1975), pp. 568-579.

[SMIT81] Smith, R. G. and Davis, R., "Frameworks for Cooperation in Distributed Problem Solving," IEEE Transactions on Systems, Man, and Cybernetics, SMC-11, No. 1, (January 1981), pp. 61-70.

[SPIN72] Spinelli, D. N., Hirsch, H. V. B., Phelps, R. W., and Metzler, J., "Visual Experience as a Determinant of the Response Characteristics of Cortical Receptive Fields in Cats," Experimental Brain Research, 15, (1972), pp. 289-304.

[SPIN79] Spinelli, D. N. and Jensen, F. E., "Plasticity: The Mirror of Experience," Science, 203, (1979), pp. 75-78.

[SPIN80] Spinelli, D. N., Jensen, F. E., and Viana Di Prisco, G., "Early Experience Effect on Dendritic Branching in Normally Reared Kittens," Experimental Neurology, 68, (1980), pp. 1-11.

[SPIN81] Spinelli, D.N., "Cluster Analysis in the Visual System of Animal Brains," Special Session of the American Mathematical Society, (October 1981).

[SPIN82] Spinelli, D.N. and Jensen, F.E., "Brains, Machines and Crystals of Knowledge," Tenth International Machine Intelligence Workshop, Case Western University, (1982).

[STAN78] Stankovic, J.A., et al., "Issues in Distributed Processing," IEEE Computer, 11, 1, (January 1978).

[STAN80] Stankovic, J.A., "A Comprehensive Framework for Evaluating Decentralized Control," Proceedings of the 1980 International Conference on Parallel Processing, (August 1980).

[STAN81] Stankovic, J.A., "The Analysis of a Decentralized Control Algorithm for Job Scheduling Utilizing Bayesian Decision Theory," Proceedings of the 1981 International Conference on Parallel Processing, (August 1981).

[STEM76] Stemple, D.W., "A Database Management Facility for Automatic Generation of Database Managers," ACM Transactions on Database Systems, 1,1, (March 1976), pp.79-94.

[STON77] Stone, H.S., "Multiprocessor Scheduling with the Aid of Network Flow Algorithms," IEEE Transactions on Software Engineering, SE-3, (January 1977).

[STON78a] Stone, H.S., "Critical Load Factors in Distributed Computer Systems," IEEE Transactions on Software Engineering, SE-4, (May 1978), pp.254-258.

[STON78b] Stone, H.S. and Bokhari, S., "Control of Distributed Processes," Computer, 11,7, (July 1978), pp.97-106.

[SUTT81a] Sutton, R.S. and Barto, A.G., "Toward a Modern Theory of Adaptive Networks: Expectation and Prediction," Psychological Review, 88, (1981), pp.135-170.

[SUTT81b] Sutton, R.S. and Cromarty, A.S., "GUS: A Device-Independent Graphics Utility Subsystem," Preliminary Reference Manual, (February 1981).

[TEIC77] Teichroew, D. and Hershey, E., "PSL/PSA: A Computer-Aided Technique for Structured Documentation and Analysis of Information Processing Systems," IEEE Transactions on Software Engineering, SE-3, 1, (January 1977), pp.41-48.

[VIAN81] Viana Di Prisco, G. and Spinelli, D.N., "Observations on Cortical Plasticity Using HRP," Experimental Neurology, (accepted for publication), (1981).

[WALL80] Wall, R.S., "Decryption of Simple Substitutation Cyphers with Word Divisions Using a Content Addressable Memory," Cryptologia, Vol.4, No.2, (April 1980), pp.109-116.

[WEEM81] Weems, C., "Life is a C.A.M. Array, Old Chum," (unpublished), (1981).

[WEEM82] Weems, C., Levitan, S., and Foster, C., "Titanic: A VLSI Based Content Addressable Parallel Array Processor," International Conference on Computer Circuits, New York, New York, (September 1982).

[WELT81] Welty, C. and Stemple, D.W., "Human Factors Comparison of a Procedural and Nonprocedural-Query Language," ACM Transactions on Database Systems, (in press), (1981).

[WESL82] Wesley, L., "Control as an Inference Process: A Model for Making Resource Allocation Decisions in the VISIONS System," COINS Technical Report, (in preparation), University of Massachusetts, Amherst, (1982).

[WEYM81] Weymouth, T., "Experiments in Knowledge-Driven Interpretation of Natural Scenes," Proceedings of IJCAI-7, (August 1981), pp.628-630.

[WEYM82] Weymouth, T., "Schema-Guided Interpretation," Ph.D. Thesis proposal, Department of Computer and Information Science, University of Massachusetts, Amherst, (in preparation), (1982).

[WHIT80] White, L.J. and Cohen, E.I., "A Domain Strategy for Computer Program Testing," IEEE Transactions on Software Engineering, SE-6, 3, (May 1980), pp. 247-257.

[WILE78] Wileden, J. "Modelling Parallel Systems with Dynamic Structure," COINS Technical Report 78-4, University of Massachusetts, Amherst, (January 1978).

[WILE79] Wileden, J.C., "DREAM - An Approach to Designing Large Scale, Concurrent Software Systems," Proceedings of the 1979 ACM National Conference, Detroit, Michigan, (October 1979), pp.88-94.

[WILE80] Wileden, J., "Techniques for Modelling Parallel Systems with Dynamic Structure," Journal of

Digital Systems, 4,2, (Summer 1980), 177-197.

[WILE82] Wileden, J. "Constrained Expressions and the Analysis of Designs for Dynamically-Structured Distributed Systems," Proceedings of the 1982 International Conference on Parallel Processing, Shanty Creek, Michigan, (August 1982).

[WILL79] Williams, T.D., "Region Analysis for a Moving Scene," Workshop on Computer Anmlysis of Time-Varying Imagery, Philadelphia, (1979), pp.17-19.

[WILL80] Williams, T.D., "Depth From Camera Motion in a Real World Scene," IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-2, (1980), pp.511-516.

[WILL81] Williams, T.D., "Computer Interpretation of Dynamic Images From a Vehicle in Motion," Ph.D. Dissertation, Department of Computer and Information Science, University of Massachusetts, Amherst, (1981).

[YAO79] Yao, S.B., "Optimization of Query Evaluation Algorithms," ACM Transactions Database Systems, 4,2, (June 1979), pp.133-155.

[YORK80] York, B.W. and Riseman, E.M., "A Surface Representation For Computer Vision," Proceedings of the IEEE Workshop on Picture Data Description and Management, (August 1980).

[YORK81a] York, B.W., "Shape Representation in Computer Vision," Ph.D. Thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, (February 1981).

[YORK81b] York, B.W., Hanson, A.R., and Riseman, E.M., "3D Representation and Matching with BpSplines and Surface Patches," Proceedings of IJCAI-81, Vancouver, British Columbia, (August 1981).

[ZEIL81] Zeil, S.J. and White, L.J., "Sufficient Test Sets for Path Analysis Testing Strategies," 5th International Conference on Software Engineering, San Diego, (March 1981), pp.184-191.

[ZELK78] Zelkowitz, M.V., "Perspectives on Software Engineering," Computing Surveys, 10,2, (June 1978), pp.197-216.

[ZISM78] Zisman, M.D., "Use of Production Systems for Modelling Asynchronous Concurrent Processes," in Pattern-Directed Inference Systems, R.Findler

(Ed.), Academic Press, (1978).

VIII. CURRENT AND PENDING SUPPORT

CURRENT SUPPORT:

1. M. Arbib, Principal Investigator, NIH (R01 NS14971 SSS) "Networks: Networks Developmental Studies; Visuomotor Activity," $349,961, 1/1/79 - 12/31/83.

2. M. Arbib and eleven Co-Principal Investigators, Sloan (80-6-13) "A Program in Languge and Information Processing," $490,000, 7/1/80 - 8/31/83.

3. M. Arbib and E. Manes, Co-Principal Investigators, NSF (MCS-8005112) "Partially-Additive Semantics of Programming Languages," $79,997, 7/15/80 - 12/31/82.

4. M. Arbib and D. McDonald, Co-Principal Investigators, NSF (IST-8104984) "Dynamic Interaction of Multiple Representations: Vision and Language," $174,416, 6/1/81 - 5/31/84.

5. M. Arbib, Principal Investigator, NSF (INT-8100477) "U.S.-Japan Cooperative Seminar: Competition and Cooperation in Neural Nets/Kyoto, Japan/February, 1982," $19,266, 10/31/81 - 10/31/82.

6. M. Arbib and E. Riseman, Co-Principal Investigators, NSF (ECS-8108818) "Visuo-Tactile Coordination for Robot Control," $317,758, 1/15/82 - 12/31/84.

7. D. Carlson, Principal Investigator, NSF (ECS-8204894) "Analysis of Resource Tradeoffs," $47,988, 6/1/82 - 11/30/84.

8. L. Clarke and D. Richardson, Co-Principal Investigators, NSF (MCS-8104202), "A Partition Analysis Method to Demonstrate Program Reliability," $81,723, 7/1/81 - 6/30/83.

9. B. Croft, Principal Investigator, NSF (IST-8011605) "A Generalized Model of Document Retrieval," $61,740, 7/1/80 - 6/30/82.

10. B. Croft, Principal Investigator, NSF (IST-8111108) "A Framework for User-Oriented Adaptive Document Retrieval Systems," $121,240, 1/1/82 - 12/31/83.

11. H. Elliot, Co-Principal Investigator, Hewlett-Packard, Fort Collins, Colorado, "HP Software Development Program," $50,000, 9/81 - 8/82.

12. H. Elliot, Principal Investigator, NSF (ECS-8110042) "Hybrid Adaptive Control of Multivariable Systems," $70,000, 11/1/81 - 10/31/83.

13. H. Elliot, Principal Investigator, ONR (NR 609-009) "Stochastic Image Processing," $98,818, 11/1/81 - 10/31/84.

14. R. Graham, Principal Investigator, NSF (SER-8013829) "Real Time Computing Laboratory," $18,083, 10/15/80 - 9/30/83.

15. F. Hill and L. Franks, Co-Principal Investigators, NSF, "Simplified Implementation of Digital Signal Processors Based on Delta Modulation Encoding," $81,000, 6/80 - 6/82 (plus extension).

16. F. Hill and C. S. Chana, Co-Principal Investigators, NSF, "Microcomputer-Based Enhancement of Soil Mechanics Course," $29,950.

17. F. Hill, Principal Investigator, General Electric Co., "User-Friendly Videodisc Interactions," $15,000, 9/1/82 - 8/31/83.

18. F. Hill, Principal Investigator, ASTEC Corporation, "Progressive Transmission of Remote-Sensing Images," $16,000, 5/82 - 5/83.

19. W. Kohler, J. Stankovic and D. Towsley, Co-Principal Investigators, NSF (ECS-81-20931) "Concurrency Control and Recovery Algorithms for Distributed Transaction Processing Systems," $128,000, 5/15/81 - 4/30/84.

20. V. Lesser, Principal Investigator, NSF (MCS-8006327) and DARPA (NR049-041) (joint grant), "Cooperative Distributed Problem-Solving," $385,000, 10/1/80 - 9/31/82.

21. V. Lesser and E. Riseman, Co-Principal Investigators, NSF (MCS-8204251) "Computer Science and Computer Engineering Research Equipment," $175,000, 6/15/82 - 5/31/83.

22. V. Lesser and B. Croft, Co-Principal Investigators, Digital Equipment Corporation, personnel support for "Office Information Systems," $60,000, 6/82 - 9/82.

23. V. Lesser, Principal Investigator, DARPA, "Coordination in Distributed Problem-Solving," $379,945, 10/1/82 - 9/30/85.

24. V. Lesser, Principal Investigator, DARPA (equipment grant), $296,000, 6/1/82 - 3/31/83.

25. V. Lesser and B. Croft, Co-Principal Investigators, DEC, "Equipment Donation to Support the Development of Intelligent User Interfaces," $750,000, 6/1/82 - 5/31/84.

26. K. Ramamritham, Principal Investigator, NSF (MCS-8202586) "Resource Controllers for Distributed Systems," $61,383, 4/15/82 - 4/14/84.

27. E. Riseman and A. Hanson, Co-Principal Investigators, ONR (N00014-75-C-0459) "Semantically Directed Vision Processing," $140,000, 4/1/80 - 3/31/84.

28. E. Riseman and A. Hanson, Co-Principal Investigators, NSF (MCS79-18209) "A Computer System for Visual Interpretation of Natural Scenes," $373,800, 9/1/79 - 2/28/83.

29. E. Riseman and A. Hanson, Co-Principal Investigators, DARPA (N00014-82-K-0464) "Processing Dynamic Images from Camera Motion," $420,000, 6/1/82 - 5/31/84.

30. E. Riseman and A. Hanson, Co-Principal Investigators, ASTEC, "Geometric Correction of Remote Sensing Images," $19,000, 2/1/82 - 12/31/82.

31. E. Riseman, Principal Investigator, Tufts New England Medical Center, "Biomedical Image Analysis Applied to Ophthalmology," $6,000, 7/1/81 - 6/30/82.

32. E. Riseman, Principal Investigator, NSF, "Specialized Research Equipment: Sensory Processing System and Mass Storage for Visuo-Tactile Robot Control," $48,000 9/15/82 - 2/29/84.

33. E. Riseman and M. Arbib, Co-Principal Investigators, General Electric Corporation, "Equipment for Robotics Laboratory," $70,000, 1/15/82 - 6/30/83.

34. E. Riseman, Principal Investigator, unsolicited IBM Departmental Research Support Grant, $25,000, 9/1/81 - 8/31/82.

35. E. Riseman, Principal Investigator, UMass BRSG from NIH (RR07048-16) "Biomedical Image Analysis Applied to the Prognosis of Malignant Melanoma," $4,360, 4/1/81 - 3/31/83.

36. E. Riseman, Principal Investigator, DEC, "Support of Computer Vision and Parallel Architectures," $10,000, 6/1/82 - 5/31/83.

37. E. Riseman, Principal Investigator, DEC, "Computer Vision Applications for Industrial Automation," $109,173, 9/1/80 - open.

38. E. Rissland, Principal Investigator, Army Research Institute (MDA903-80-C-0508) "The Development and Evaluation of Instructional Strategies for an ICAI System," $78,161, 7/1/81 - 6/30/82.

39. E. Rissland, Principal Investigator, NSF (IST-8017343) "The Structure of Examples," $147,848, 1/1/81 - 12/31/82.

40. E. Rissland, Principal Investigator, NSF (IST-8212238) "Conference on Intelligent User Interfaces," $28,818, 4/1/83 - 3/31/84.

41. D. Spinelli and A. Barto, Co-Principal Investigators, AFOSR (F33615-80-C-1088) "Decision Making by Adaptive Networks of Goal-Seeking Components," $383,909, 6/1/80 - 6/30/83.

42. J. Stankovic, Principal Investigator, NSF (MCS-8104203) "Decentralized Control of Scheduling in Distributed Processing Systems," $38,799, 7/1/81 - 6/30/83.

43. J. Stankovic, Principal Investigator, DEC Applied Technology Center (BD-749339) "Design and Develoment of a Distributed Computer-Based Instructional System," $47,350, 7/1/81 - 6/30/82.

44. J. Stankovic, Principal Investigator, U.S. Army CECOM (DAAB07-82-K-J015) "Decentralized Control of Scheduling in Distributed Processing Systems," $130,830, 12/15/81 - 2/15/85.

45. H. Stone, Principal Investigator, NSF (MCS-7805295) "Control of Distributed Processes," $123,802, 12/31/81 - 12/31/83.

46. D. Towsley and J. Wolf, Co-Principal Investigators, NSF (ECS-7921140) "On the Analysis and Synthesis of Computer Communications Networks," $164,944, 6/1/80 - 11/30/83.

47. J. Wileden and L. Clarke, Co-Principal Investigators, NASA Langley Research Center (NAG1-115) "Tools Supporting Pre-Implementation Software Development," $134,721, 3/1/81 - 2/28/83.

48. S. Zeil, Principal Investigator, NSF (MCS-8210084) "A Testbed for Evaluating Software Testing Methods," $39,917, 6/1/82 - 5/31/84.

PROPOSALS PENDING:

1. M. Arbib and E. Manes, Co-Principal Investigators, NSF, "Algebraic Semantics of Programming Languages," $192,530, 1/1/83 - 12/31/85.

2. A. Barto, Principal Investigator, AFOSR, "Adaptive Problem-Solving with Networks of Goal-Seeking Components," $540,016, 6/15/83 - 6/14/86.

3. A. Barto and J. Moore, Co-Principal Investigators, NSF, "Experimental Tests of a Model of Expectation and Anticipation in Adaptive Networks," $180,000, 6/1/83 - 5/31/85 (in preparation; to be submitted within 60 days).

4. L. Clarke, R. Graham and J. Wileden, Co-Principal Investigators, DARPA, "Toward Feedback Directed Development of Complex Software Systems," $420,000, 6/1/83 - 5/31/85 (in preparation; to be submitted within 60 days).

5. C. Foster, Principal Investigator, ARO, "Construction of a Content Addressable Memory," $300,000, 3 years (in preparation; to be submitted within 60 days).

6. A. Hanson and E. Riseman, Co-Principal Investigators, AFOSR, "Representation and Control in the Interpretation of Complex Scenes," $1,133,711, 1/1/83 - 12/31/86.

7. F. Hill, Principal Investigator, AFOSR, "Compaction of Map Data Bases for Rapid On-Board Map Displays," $115,000, 2 years.

8. S. Kulikowski, Principal Investigator, NSF, "Organization of Knowledge-Based Language Production in Systems with Few User Inputs," $98,171, 1/1/83 - 12/31/84.

9. S. Kulikowski and E. Riseman, Co-Principal Investigators, INA Foundation and Raytheon Corporation, "Development of a Microcomputer Prosthetic for the Severely Physically Handicapped," $69,169, 9/1/82 - 8/31/84.

10. S. Kulikowski and E. Riseman, Co-Principal Investigators, Massachusetts Rehabilitation Commission, "Independent Living Use of Microcomputers for the Severely Physically Handicapped," $8,812, 10/1/82 - 9/30/83.

11. W. Lehnert, Principal Investigator, NSF, "Narrative Text Summarization," $226,074, 3/1/83 - 2/28/85.

12. V. Lesser, Principal Investigator, NSF, "Coordination in

Distributed Problem Solving," (coordinated with DARPA contract -- see #23 above in Current Funding), $366,000, 2/1/83 - 1/31/86.

13. D. Stemple, Principal Investigator, NSF, "Strategy Enhancement in Query Language Learning," approximately $50,000, 6/1/83 - 12/1/84 (in preparation; to be submitted within 60 days).

14. J. Wileden and G. Avrunin, Co-Principal Investigators, ONR, "Algebraic Techniques for Analyzing Concurrent Systems," $103,413, 6/1/83 - 5/31/85.

IX. STATEMENT OF RESPONSIBILITY FOR EQUIPMENT

Professor Edward Riseman, Chairman of the COINS Department, and Professor Victor Lesser, Director of the Research Computer Facility, will be jointly responsible for the acquisition and maintenance of the proposed equipment. There is a well-organized management structure for the RCF that has been described in detail in Section III.5 of this proposal.

UNIVERSITY OF MASSACHUSETTS

AMHERST

TO: Whom It May Concern

FROM: Dean, Faculty of Natural Sciences and Mathematics

DATE: September 21, 1982

SUBJECT: University Commitment to the Department of Computer and Information Sciences

The general area of Computer Science and Engineering is clearly of tremendous importance to the economic future of our country in addition to being of great scientific significance. In the Commonwealth of Massachusetts, the importance of this subject is enhanced by the profound dependence of our State's economy on high technology industries. It is not suprising, therefore, to discover that the University of Massachusetts has made a commitment to this research area which goes far beyond any academic commitments made in at least the past decade.

The Department of Computer and Information Sciences (COINS) added three new faculty last Fall and three additional faculty this Fall. Its sister Department of Electrical and Computer Engineering (ECE) has also added several new computer faculty. Furthermore, COINS has been allocated an administrative assistant for the present Chairman. In addition, the University has in recent years provided $117,000 in matching funds for three NSF equipment grants to COINS.

Perhaps the most precious resource which a university possesses after its faculty is space, and here our commitments have been extensive. Roughly 4,500 square feet of additional space has been allocated to our computer science and engineering effort during the past year. More significantly, we are committed to a shared research environment for COINS and those ECE faculty in the computer area. To this end, we are preparing to move our Graduate School from the Graduate Research Center into new quarters to allow adequate space for the joint growth of these two collaborating units. This will mean an additional 12,000 square feet.

Finally, we have allocated some of the University's Research Trust Funds to the Department in the amount of 15% of the indirect costs generated via departmental grants from external agencies. This comes to over $50,000 per year in extremely flexible funds which can be spent for any research-related purpose. Much of this money has been used to support the Department's Research Computing Facility. Should the present CER proposal to the NSF be funded, we are also prepared to fund a laboratory manager and a technician from Research Trust Funds at a total cost of about $60,000 per year.

As far as the future is concerned, clearly the Department cannot continue to grow indefinitely at a rate of three new faculty per year. However, our present long-range planning does contemplate an additional five or six faculty for COINS which will be added during the lifetime of the present proposal. In addition, significant growth in the Computer Engineering section of ECE is planned so that a total computer research community of 30-40 faculty is virtually assured.

Why are we making the extensive commitments described in this memorandum? Firstly, because we think we know an exciting, indeed revolutionary, research area when we see one. Secondly, because as the Commonwealth of Massachusetts' land-grant institution we have a special responsibility to be concerned with benefitting the State in our research and teaching activities while at the same time remaining faithful to the scholarly principles of a major university. Our commitment to the area of computer science and engineering will help ensure our ability to respond to the Commonwealth's needs in the vital area of high technology, and at the same time help develop fundamental knowledge in this enormously promising scientific discipline.

F. W. Byron, Jr.

F. W. Byron, Jr.
Dean, Faculty of Natural Sciences
and Mathematics

Date of this Report 9/15/82

METRICS FOR RESEARCH ENVIRONMENT
(Mark unavailable data as N/A)

Except where noted, these numbers reflect COINS statistics, only.

DEPARTMENTAL FACTORS

Number of Full Time Faculty 19 COINS/ 7 ECE

Number of Technical Support Staff 3

Number of Clerical Support Staff 10

Budget

Internal instructional* 27.6 %

Internal research* 5.8 %

External research 66.6 %

Other ______ %

Departmental Budget per Faculty ($/person) $172,526

Capitalization per Faculty** $105,263

Space per Faculty

Office 714

Laboratory 212

Total 926

Append faculty curriculum vitae***

Append a list of faculty publications 1977-1982***

Append a list of full time faculty for each year 1972-82, indicating additions and departures.

* As defined by the institution.

** Major equipment only, market value cost basis.

*** Only data <u>not</u> included in proposal.

STUDENT QUALITY

Except where noted, these numbers reflect COINS statistics, only.

Undergraduates

Number

Lower division 38

Upper division 90

Percentage attending graduate studies 40%

Percentage taking employment Not available at this time

Average salary ____

Average number of offers ____

Graduates

Number

M.S. (part-time) ____

M.S. (full-time) 81

Ph.D 45

Average entering test scores(GRE,etc.)

1325 out of 1600

Undergraduate GPA 3.40

Ratio number accepted to number enrolled 1.583

Number holding Fellowships 7

Research Assistantships in dept. 46

other dept. 18

Teaching Assistantships 35

Append list of first positions of Ph.D. graduates*

Append list of dissertation titles*

---

* From beginning of program.

All figures represent joint COINS Department and ECE Department statistics.

EXTERNAL SUPPORT

Append a list of current active grants and contracts*

Append a list of pending proposals*

Number of grants 48 Number per faculty 1.846

Number of contracts ______ Number per faculty ______

Total annual grant and contract budget 2.565 Million

Budget per faculty $98,654

Number of agencies/foundations providing support 9

Number of companies providing support 4

Number of corporate and other gifts 1

* Only data not included in proposal.

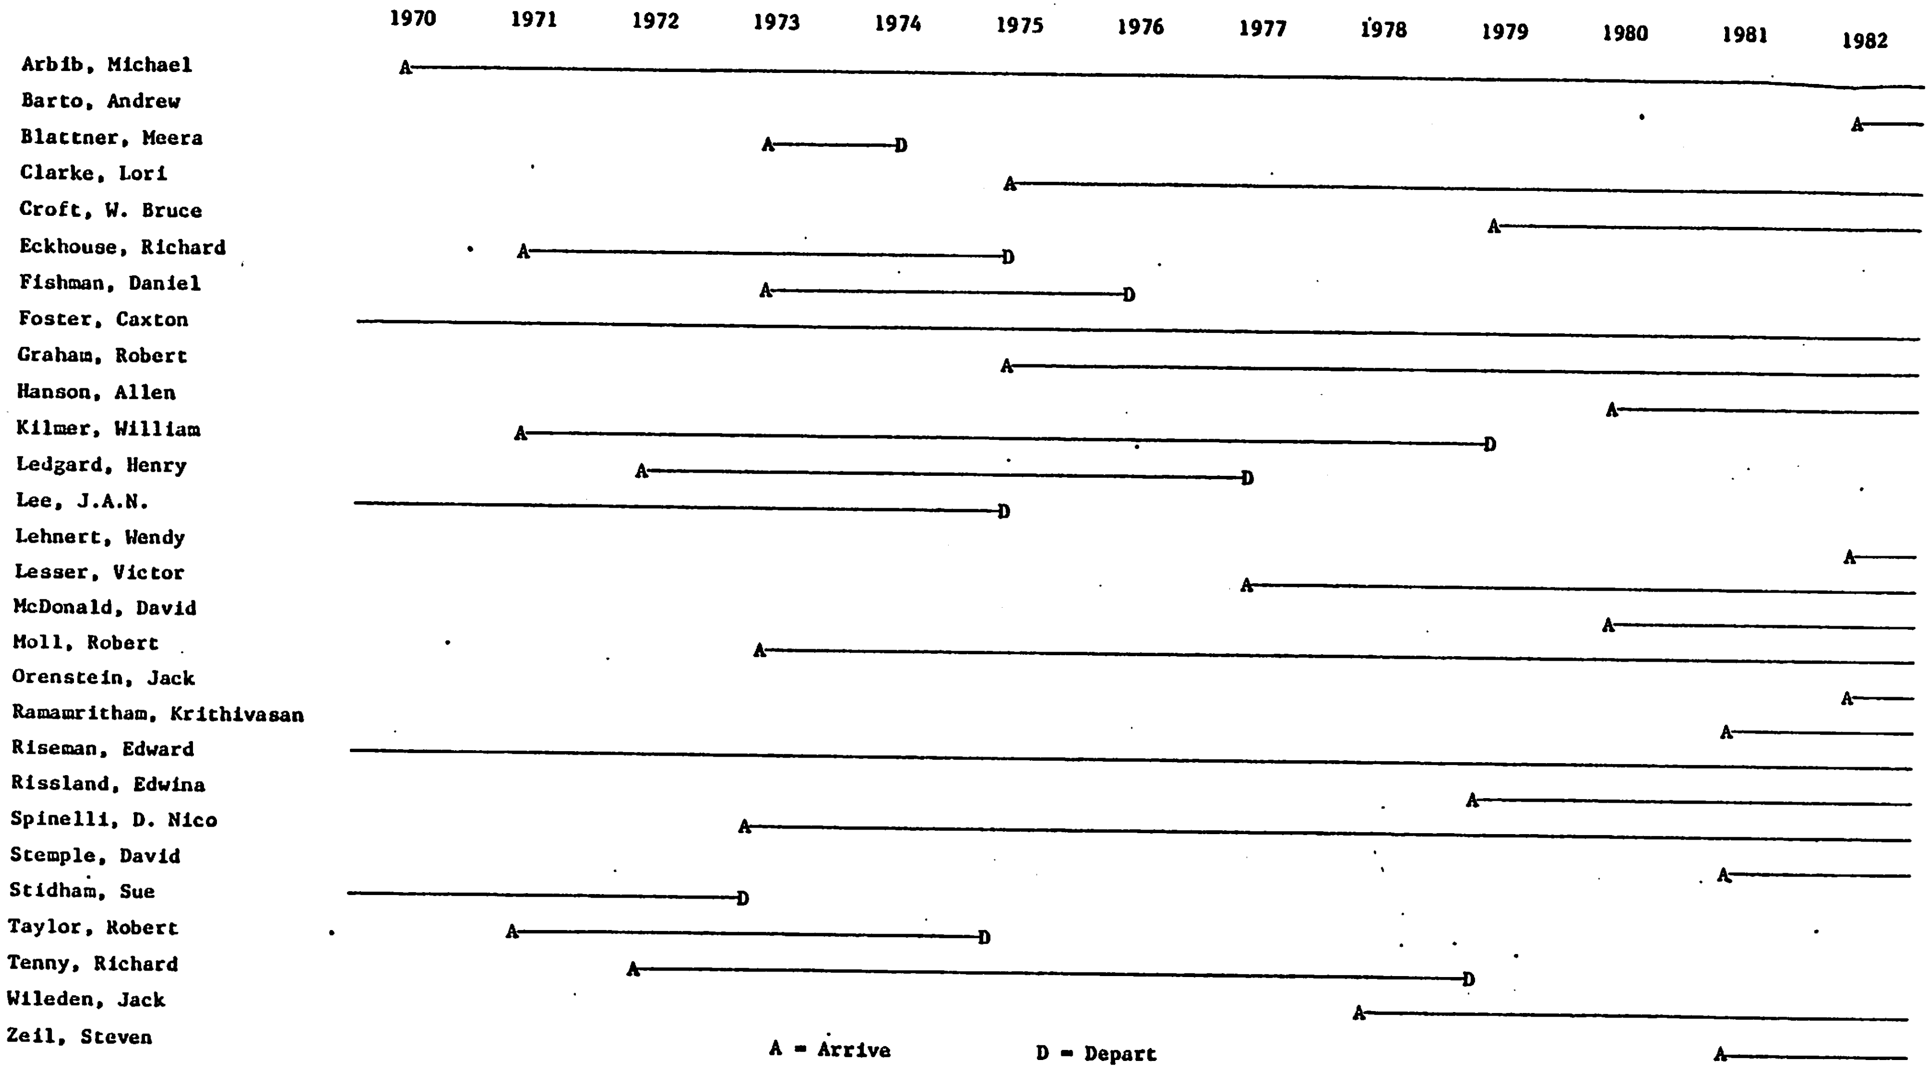
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
Arbib, Michael
Barto, Andrew
Blattner, Meera
Clarke, Lori
Croft, W. Bruce
Eckhouse, Richard
Fishman, Daniel
Foster, Caxton
Graham, Robert
Hanson, Allen
Kilmer, William
Ledgard, Henry
Lee, J.A.N.
Lehnert, Wendy
Lesser, Victor
McDonald, David
Moll, Robert
Orenstein, Jack
Ramamritham, Krithivasan
Riseman, Edward
Rissland, Edwina
Spinelli, D. Nico
Stemple, David
Stidham, Sue
Taylor, Robert
Tenny, Richard
Wileden, Jack
Zeil, Steven
A = Arrive
D = Depart

COINS Ph.D. Theses

Suad Alagic (2/74): Algebraic Aspects of Programming and Formal Languages.

Fred Lenherr (6/75): Effects of Inter-Ocular Coherence and Audio-Visual Correlations on the Development of Visual Cortex in Kittens.

Edward Fisher (6/76): The Use of Context in Character Recognition.

James Stanley (9/76): Network Models of Habituation.

Fanya Montalvo (9/76): Visual Feature Organization and Interaction as Modelled in Neural Networks.

Peter Burt (9/76): Stimulus Organizing Processes in Stereopsis and Motion Perception.

Stephen Hegner (2/77): Applications of Topological Vector Spaces in Linear System Theory.

David Stemple (2/77): A Database Management Facility and Architecture for the Realization of Data Independence.

Frederic Richard (9/77): Compiling Technique for Associative Processors.

Yuan-Chieh Chow (9/77): Queueing Models for Distributed Computer Systems.

Elliot Soloway (9/78): Mechanisms for Knowledge-Directed Learning.

Arthur Karshmer (9/78): The Use of Frame-to-Frame Differences in Encoding Computer Graphics Data in a Network Environment.

John Prager (5/79): Segmentation of Static and Dynamic Scenes.

Charles Welty (5/79): A Comparison of a Procedural and a Nonprocedural Query Language: Syntactic Metrics and Human Factors.

Paul Nagin (9/79): Studies in Image Segmentation Algorithms Based on Histogram Clustering and Relaxation.

Andrew Singer (9/79): Formal Methods and Human Factors in the Design of Interactive Languages.

Ruth Maulucci (2/80): Kinetics and Optimality in Quadraped Locomotion.

John Woods (5/80): Path Selection for Symbolic Execution Systems.

Thomas Probert (2/81): A Model of Renewable Resource Dynamics in Exploited Ecosystems.

Alan Morse (5/81): Computer Graphics and Modelling: Using Effective Data Displays to Enhance Understanding of Simulation Results.

Scott Reed (5/81): Using Numerical Search in Managing Ecological Systems: Handling Nonlinearities and Time Delays.

Thomas Williams (5/81): Computer Interpretation of Dynamic Images from a Vehicle in Motion.

Bryant York (5/81): Shape Representation in Computer Vision.

Balakrishnan Krishnamurthy (9/81): Examples of Hard Tautologies and Worst-Case Complexity Results for Propositional Proof Systems.

John Tan (9/81): An Empirical Approach to Computer Performance Improvement.

Debra Richardson (9/81): A Partition Analysis Method to Demonstrate Program Reliability.

Rolando Lara (2/82): Neural Models of the Visuomotor System of Amphibia.

Helen Gigley (9/82): Neurolinguistically Constrained Simulation of Sentence Comprehension: Integrating Artificial Intelligence and Brain Theory.

Jane Hill (9/82): A Computational Model of Language Acquisition in the Two-Year-Old.

John Lowrance (9/82): Dependency-Graph Models of Evidential Support.

First Positions of Ph.D. Graduates

Suad Alagic: Teaching Position, Electro Technology Department, University of Sarajevo, Yugoslavia.

Fred Lenherr: Staff Associate, Center for Systems Neuroscience, University of Massachusetts, Amherst, Masssachusetts.

Edward Fisher: Research and Development Group, Pattern Analysis and Recognition, Inc., Rome, New York.

James Stanley: Post-Doctoral Researcher, Neurobiology and Anatomy Department, University of Texas Health Science, Houston, Texas.

Fanya Montalvo: Industrial Consultant, Computer Science Department, Lawrence Berkeley Lab, University of California, Berkeley, California.

Peter Burt: Post-Doctoral Researcher, Psychology Department, New York University, New York, New York.

Stephen Hegner: Systems Analyst, Lawrence Livermore Lab, University of California, Livermore, California.

David Stemple: Associate Director for Research, Computing Center, University of Massachusetts, Amherst, Massachusetts.

Frederic Richard: French Army.

Yuan-Chieh Chow: Assistant Professor, Computer Science, Wright State University, Dayton, Ohio.

Elliot Soloway: Post-Doctoral Researcher, Computer and Information Science, University of Massachusetts, Amherst, Massachusetts.

Arthur Karshmer: Assistant Professor, Computer Science, New Mexico State University, Las Cruces, New Mexico.

John Prager: IBM Scientific Center, Cambridge, Massachusetts.

Charles Welty: Associate Professor, Department of Computer Science, University of Maine, Portland, Maine.

Paul Nagin: Assistant Professor, Department of Opthamology, Tufts New England Medical Center, Boston, Massachusetts.

Andrew Singer: Vice-President for Research, E & L Instruments, Derby, Connecticut.

Ruth Maulucci: Post-Doctoral Researcher, Brain Research Lab, American Oncologic Hospital, Philadelphia, Pennsylvania.

John Woods: Consultant, Government Systems, Control Data Corporation, Sunnyvale, California.

Thomas Probert: EDIS/CEAS/MEAD, Washington, D.C.

Alan Morse: Senior Software Engineer, Applicon, Lexington, Massachusetts.

Scott Reed: Owner of a software company, Amherst, Massachusetts.

Thomas Williams: Manager of Manufacturing Automation Project, Digital Equipment Corporation, Maynard, Massachusetts.

Bryant York: Technical Staff Member, IBM Research Laboratory, San Jose, California.

Balakrishnan Krishnamurthy: Research Staff, General Electric, Schenectady, New York.

John Tan: Has not yet accepted a position.

Debra Richardson: Visiting Assistant Professor, Computer and Information Science, University of Massachusetts, Amherst, Massachusetts.

Rolando Lara: Faculty Member, Centro de Investigaciones en Fisiologia Celular, Universidad Nacional Autonoma de Mexico, Mexico City, Mexico.

Helen Gigley: Assistant Professor, University of New Hampshire, Durham, New Hampshire.

Jane Hill: Assistant Professor of Computer Science, Smith College, Northampton, Massachusetts.

John Lowrance: Research Scientist, SRI International, Menlo Park, California.