STRATEGIC PLAN
INSTITUTE FOR ADVANCED COMPUTER STUDIES
UNIVERSITY OF MARYLAND, COLLEGE PARK

1 Mission and Strategic Objectives

Mission

The mission of the University of Maryland Institute for Advanced Computer Studies (UMIACS) is to stimulate and nurture interdisciplinary and applications-oriented research programs in computing and to provide fundamental knowledge in computer science and engineering. To carry out these responsibilities, the Institute has established research teams currently led by faculty from 7 academic units from across the College Park campus to work on ambitious projects in areas of national importance such as high-performance computing, software engineering, and intelligent systems. Most of the projects are coordinated closely with scientists from industry and government laboratories. UMIACS has regularly organized workshops and conferences in emerging technological areas and has sponsored summer fellowship programs for High School students from the State of Maryland.

Strategic Objectives

As information technology has become the major enabling technology for providing high-quality and cost-effective services and products, it is crucial that research in computing be pursued at a vigorous pace to sustain and enhance the growth and breadth of the technology, and to shorten its transfer to the marketplace. Given the rich application domains, the role of multi-disciplinary research work with a close industrial linkage is a critical component for the continued success of this technology. UMIACS is well-positioned to achieve the following broad strategic objectives.

I. Address critical technical challenges in the emerging information technology by strenghtening existing multi-disciplinary research programs and building new research initiatives that are closely linked to strategic technology areas.

II. Enhance outreach activities by working closer with industry and government laboratories, thereby enriching the University research environment with real-world applications and helping to transfer technology to the marketplace.

III. Sponsor multi-disciplinary educational programs in computing on the College Park campus, such as leading a new graduate program in computational science and strenghtening our outreach activity for High School students from the State of Maryland.

2 Background

UMIACS has been the focal point for interdisciplinary research activities in computing on the College Park campus since its inception in 1985. The Institute's research programs are led by over 40 faculty from a number of academic units, currently including the Computer Science Department,

Electrical Engineering, Geography, Business and Management, Linguistics, Philosophy, Mechanical Engineering, and the Computer Science Department at UMBC. The research activities at the Institute are enriched by a strong outreach program that includes close collaboration with industry on focused research projects and the organization of topical workshops that bring many of the nation's outstanding scientists to Maryland.

Research Programs

The Institute's faculty conduct research in high performance computing, software engineering, artificial intelligence, systems, combinatorial algorithms, scientific computing, and computer vision.

High-Performance Computing (HPC). In September 1993 UMIACS in conjunction with the Department of Geography was awarded a $ 3M, five year National Science Foundation Grand Challenge grant to develop HPC technology for applications in remote sensing, with an emphasis on land cover dynamics. The NSF Grand Challenge program is one of the most critical components of the highly visible federal program in High-Performance Computing and Communication. The University of Maryland is one of sixteen universities nationally to win a Grand Challenge award. In connection with this project, the Computer Science Department and UMIACS have obtained a $ 1M NSF infrastructure award to purchase a supercomputer in support of the grand challenge research. Another highly successful research program in our HPC activities has been the work conducted through the High Performance Systems Software Laboratory (HPSL) that concentrates on the development of systems software to enable the efficient programming of multiprocessor machines and clusters of workstations. UMIACS faculty have played a leading role in addressing fundamental issues in HPC.

Software Engineering. Software engineering has been identified as one of the critical disciplines to the nation's economy. Our research in software engineering is aimed at supporting an engineering approach based upon empirical studies. The pioneering example of a software engineering research laboratory, and the most successful to date, has been the Software Engineering Laboratory (SEL), established in 1976 as a cooperative effort among the University of Maryland, the NASA Goddard Space Flight Center, and Computer Sciences Corporation. This laboratory has been awarded the first IEEE Software Process Achievement award for its long record of sustained achievement in software process improvement through modeling and measurement. Our research in software engineering has had a dramatic and practical impact on U.S. software engineering practices.

Artificial Intelligence. The internationally recognized UMIACS group consists of faculty from Computer Science, Philosophy, Business and Management, and Linguistics. The AI research interests range from theoretical investigations to simulations and practical applications. Research subareas include (1) logical foundations of AI, (2) cooperative answering and deductive databases (pioneered at Maryland), (3) planning and robotics, and (4) natural language processing. The last subarea, in which we are planning to grow, is led by faculty from Computer Science and Linguistics, and focuses on the development of theories and tools to translate between alternative natural languages such as Arabic, Korean, Japanese and others.

Systems. The activities grouped under systems address problems of management of hardware and software resources and of providing the environment in which applications can function effectively.

As new developments in hardware and communication technologies become available, they pose new challenges in the design and implementation of systems software, requiring new systems technology. Two particularly successful subareas have been (1) the design, analysis, and implementation of real-time systems working in distributed environments and providing fault-tolerant operations, and (2) the design and analysis of admission, routing, and scheduling algorithms for computer networks. Industrial interest in this interdisciplinary research area has been growing over the past few years.

Algorithms. The UMIACS algorithms group includes faculty from both EE and CS. It is a very active group with expertise spanning a variety of specialties including parallel algorithms, data structures, combinatorial optimization, and computational geometry. Most of the algorithms research concentrates on discrete algorithmic problems arising in software and applications and is conducted jointly wtih researchers from a variety of areas. The UMIACS work on the fundamental aspects of parallel algorithms has had high external visibility.

Scientific Computing. The research in scientific computing encompasses numerical linear algebra, large-scale parallel computing, and numerical solutions of partial differential equations. Particular applications include (1) models of turbulent flows, (2) computational fluid dynamics and structural mechanics, (3) modeling of semiconductor devices, (4) signal processing, and (5) models for self-replicating nucleotides. The work in scientific computing is led by a very strong research team of faculty from Computer Science, Electrical Engineering, Mechanical Engineering, and the Applied Mathematics Program.

Computer Vision. Research in computer vision is conducted jointly with the Center for Automation Research and the Department of Electrical Engineering. The emphasis of the research is vision systems and applications. Current projects focus on surveillance systems (both for reconnaissance and for monitoring human activities), design of database systems that can index into terabytes of imagery, and object recognition systems that integrate data from visible sensors, radars and infrared sensors. The research is performed in conjunction with both Federal laboratories (Lincoln Laboratories, Wright Patterson Air Force Lab, for example) and industry (McDonnel Douglas, Martin Marietta).

Outreach Activities and Technology Transfer

The Institute has made serious efforts to establish collaborative research programs with industry and government laboratories and to facilitate technology transfer to industry. Examples of such efforts include the outreach activities conducted by the SEL and by the Parallel Computing Laboratory. We have also undertaken educational outreach activities for the State of Maryland High School students ; during their stay at UMIACS, the High School students attended daily seminars and worked with faculty on focused research projects. It should also be noted that UMIACS has been administering a national graduate fellowship program in high performance computing for (D)ARPA.

As stated before, the SEL (Software Engineering Laboratory) is a joint venture between the University of Maryland, the NASA Goddard Space Flight Center, and Computer Sciences Corporation. The SEL has collected extensive amounts of data on software projects which have been used to build predictive models for future projects and to provide a rationale for refining particular processes being used. The Goal/Question/Metric paradigm developed at Maryland is used in several major high technology companies (such as HP, Motorola, IBM) and is the basis for the European

Esprit project AMI. Also, the UM Software Engineering group is one of the founding members of the International Software Engineering Network (ISERN) whose role is to exchange experimental data and replicate experiments around the world.

Another major outreach is conducted through the UMIACS Parallel Computing Laboratory, which provides access to state of the art parallel machines and training for scientists from the College Park campus and from local Government agencies. In fact, the major users of our laboratory have been faculty and scientists not affiliated with UMIACS. This facility has also been used for education throughout the University of Maryland System (Towson, UMBC,UMES). We have received substantial support from (D)ARPA, NSF, NIST, NIH, IBM, and DEC for building and maintaining this laboratory. A related activity is conducted through the HPSL that is dedicated to the development of systems software technology for multiprocessor architectures and clusters of workstations. The tools developed have been ported to most existing multiprocessor platforms and are currently being transferred to industry.

In 1990 UMIACS established a series of international meetings called Workshop on Parallel Algorithms (WOPA), with support from ARPA, NSA, NSF, and ONR. This series became one of the constitutent meetings of the Federated Computing Research Conference (1993 and 1996) which includes some of the most presitigious meetings in computer science and engineering.

The Industrial Associates Program (IAP) sponsored by CS, CFAR, and UMIACS is another indication of the strong ties that exist between industry and computer science at the University of Maryland. Members of the IAP have included high tech companies such as Motorola, Northrop, General Electric, Bellcore as well as international companies such as Siemens, Daimler-Benz, Finseil, Ricoh, Sony, and Panasonic.

During the years 1989-1991, UMIACS conducted a Summer Fellowship Program in computer science for High School students from the State of Maryland. The program was aimed at providing an opportunity for the students to spend 6 weeks at the College Park campus to attend daily seminars and to work with UMIACS faculty on focused research projects. The students received weekly stipends and had access to UMIACS laboratory facilities. Another program aimed at High School Computer Science Teachers was designed but not implemented due to the budget cuts incurred during 1991-92.

Diversity Issues

Although UMIACS is a relatively new unit on campus, we have been successful in attracting women and minorities who are conducting computer related research. Our record during the past five years can be summarized as follows. We have supported 5 women faculty, 4 women post-doctoral researchers, 3 women visitors, 1 minority faculty, 1 minority post-doctoral fellow, and a good number of female and minority graduate students. One of the ways UMIACS attempts to meet its diversity goal is by having a least one female or minority post-doctoral or visiting researcher each year. Last year, as part of our diversity goal, UMIACS addressed the need of providing cultural awareness training to staff. We held a joint workshop with CFAR on Cross-Cultural Understanding. This year as part of our diversity goal, we plan to hold a follow-up workshop for faculty, researchers and students. In spite of our relative success, we plan to make a more concerted effort to attract additional women and minorities to UMIACS.

3 The Information Technology Environment over the Next Decade

A number of factors are causing a major reorganization of information technology and the emergence of new technical areas that are likely to play a dominant role in computing research during the next decade. These factors include the following:

- The hardware miniaturization of processors and their interconnects continues at a fast rate; such a trend is causing a large shift in the computer industry into small and portable processors that can be easily interconnected and easily supported by compatible software.
- A merging of computing and communications technologies is taking place with the likelihood of major social and economic implications. The highly publicized *National Information Infratructure (NII)* is a critical element in the convergence of these two technologies.
- A substantial change in the climate for Federal Government support of basic research is underway. Social and economic benefits are becoming necessary prerequisites for justifying research expenditures and for setting national funding priorities.

These trends are reflected by the major research initiatives funded by the U.S. Government, including HPCC and NII. As a result, there is a major push towards the support of *team-oriented, multidisciplinary research efforts that are closely tied to applications.* An often cited critical element for such endeavors is a *close working relationship between academia, industry, and government.*

In addition to the structural changes, there are major technical areas that will be critical to the emerging information technology. These areas can be broadly organized under the following categories: (1) *distributed and parallel processing*, (2) *networking*, which includes software support for the seamless coordination of a large variety of heteregeneous systems, (3) *software engineering*, and (4) *multimedia information storage and management including user interfaces.* These technical areas present major technical challenges as well as major opportunities for conducting fundamental research in computer science that will substantially contribute to emerging technology. Advances in these areas are likely to move information infrastructure closer to reality thereby opening up vast new industrial markets and impacting strategic national areas such as health care, manufacturing, commerce, and finance.

4 Goals and Implementation Strategies

Current technological trends strongly favor team-oriented, multi-disciplinary research efforts that are targeted towards specific applications. The organizational structure of UMIACS and the composition of its faculty from 7 academic units is well-suited to carry out such collaborative efforts. In addition, we have considerable technical strength in computer science and in the other disciplines on campus to address fundamental issues in most of the emerging technical areas in computing. We next sketch our strategic goals and our proposed implementations in research, outreach, and education.

RESEARCH INITIATIVES

Our overall strategy in research has two major components. The first is to **maintain and enhance the excellence we have achieved** in the research areas cited in Section 2, and the second is to **develop new research programs** that are at the cutting edge of the emerging information technology.

The Institute faculty consist of leading researchers in their fields. Our faculty will continue to generate the external support needed to sustain and enhance the research conducted in their areas of expertise. UMIACS will continue to provide matching funds and the minimal resources required to support seminar series and workshops.

We next describe the new research initiatives that we will pursue in the next five years. Computer science is a fast moving field in terms of technology and new areas of research. Four major categories of emerging technical areas were mentioned in the previous section. In terms of these categories, two of our initiatives fit under the distributed and parallel processing category, two under the multimedia storage and management including interfaces, and one fits under the networking category.

Center for Language and Computational Studies (CLCS). Interdisciplinary research in computational linguistics has been a focus of *UMIACS* in the past with external funding coming from *ARI, ARPA, and ARO*. Graduate students are also jointly supervised by faculty from Computer Science, Linguistics, and Business and Management. Computational logic and linguistics play a major role in cognitive studies. We propose to develop a CLCS that will be a major part of the campus' initiative in cognitive science.

CLCS will focus its activities in research areas where the faculty already have national and international reputations. Research will focus on the development of computational techniques for the automatic creation of dictionaries for use in parsing and machine translation systems; creation of reusable cross-linguistic processing systems; and exploration of non-symbolic and hybrid approaches to the cognitive function of language understanding and acquisition systems.

The purpose of *CLCS* is to provide an environment within *UMIACS*, which will give visibility to *UMIACS* as a national center for computational work in languistics. *CLCS* will attract researchers and students throughout the world to do research in this area. We would also expect a great deal of interaction between *CLCS* and other areas such as cognitive science, data and information retrieval and artificial intelligence in general.

High Performance Computing and Health Care. The technical thrust of this initiative is to develop software architectures capable of supporting user applications which require integrating services provided by a heterogeneous set of parallel and distributed modules. The primary driving applications will be two distinct applications from medical informatics. Both applications will require integration of parallel databases, computerized medical record systems, parallel linear algebra, statistical packages and data mining packages. In particular, a risk scoring medical informatics testbed application will be developed using the Maryland Reproductive Health Database, a clinical database that generates real-time admission notes and discharge summaries for women receiving reproductive health care through the University of Maryland Medical Center system. The other medical application involves the development of more effective methods for delivering health care to geriatric veterans by manipulating the large Department of Veterans Affairs (VA) datasets.

Real-Time Vision Systems. Our goals in this new initiative, to be conducted by faculty from CS, EE, and CFAR, are twofold:

- To combine the strengths of our vision and databases/real time systems groups. This will allow us to move into several expanding areas of computer science research, including heterogeneous databases combining text, image and maps, and high performance (and embedded) real time systems. New programs are planned in both of these areas by NSF and ARPA over the next two years.
- To expand out interdisciplinary research program by pursuing vision systems research in the context of applications. Specifically, we see opportunities in the areas of reconnaissance, surveillance and navigation.

By *reconnaissance* we refer to the general problem of screening large amounts of imagery using automated vision systems to monitor changes in the world. Applications arise in environmental monitoring and modeling, and intelligence. *Surveillance* is an emerging subfield of computer vision that focuses on humans in action - to determine their state of mind through changes in facial expression, to model athletic actions, or to control computer graphics simulations of humans in actions. As for *navigation*, both in the United States and in Europe there are large, coordinated research programs whose goals are autonomous driving and intelligent driver's aids. Applications being studied include completely autonomous operation on highways, collision warning systems, and intelligent cruise controls.

Multimedia Databases. A critical component of the NII is the development of methodologies for designing and accessing heterogeneous, distributed, and multimedia databases. Initial work at the University of Maryland has concentrated on the development of the mathematical foundations of multimedia technology and on the development of techniques for balancing the platform's resources against a sequence of queries for multimedia events. We are planning to significantly enhance this activity by combining a number of currently separate efforts in this general area. In particular, we plan to build multimedia database servers using the theory developed and tying it to real-time systems technology. The goal is to provide a unified platform where an end-user of such a system can create sophisticated multimedia presentations by specifying what is required of the presentation in a high-level language, and then obtain the desired quality of service, subject to the specified real-time constraints.

Mobile Computing Laboratory. The availability of portable, light weight, powerful processors and wireless communication technology permits us to conceive of computer systems in which computers can be moved freely over a wide area with continued efficient access to various computing resources. The enabling support systems are often referred to as mobile computing systems. Our objective is to develop fundamental techniques and software systems technology for supporting mobile computing. Our current work includes mobile-IP, optimizing end-to-end transport protocols with host movement, routing support for ad-hoc networks, new operating system primitives, application layer interface and database design in mobile environment. The current research is supported by IBM (SUR Award and Graduate Fellowship) and TI (database research). Federal funding of research in mobile computing is likely to substantially increase during the next few years.

OUTREACH

In addition to maintaining and enhancing the outreach activities conducted through the Software Engineering Laboratory, the Parallel Computing Laboratory, and the Industrial Associate Program, we plan to vigorously pursue industrial partners in support of our new research initiatives. Our recent success in working out collaborative research arrangements with IBM and DEC is an indication of the industry's strong interest in some of the newer elements of our research program. We now sketch our broad strategies.

In software engineering, we are planning to expand the Software Engineering Laboratory concept towards building an Experience Factory Organization for other NASA sites (JPL, Johnson, Marshall, Langley), and expanding the technologies and organizational structures to a consortium of U.S. industries. Discussions are currently underway with AT&T, HP, Motorola, and Xerox, among others.

The UMIACS Parallel Computing Laboratory has just received a substantial boost resulting from the IBM Shared University Research (SUR) award and the cooperative research agreement with DEC. The purpose of the IBM SUR award is to establish and strengthen focused collaborative research efforts between the faculty at Maryland and scientists at IBM. Our work in HPC (run-time libraries and parallel I/O) has been chosen by IBM to receive the major share of this award. On the other hand, the cooperative research agreement with DEC involves the exchange of HPC technology in software systems and parallel I/O. The DEC agreement has allowed UMIACS and IPST to team up to acquire a cluster of 40 very powerful processors. The resulting peak performance is nearly 40 billion floating point operations per second. With the acquisition of the IBM machine and the DEC cluster, our laboratory is one of the best equipped HPC laboratory in the nation. Our plan for the next five years is to enhance the already substantial outreach activities conducted under this laboratory.

A number of outreach activities are associated with our National Science Foundation Grand Challenge program. A semiannual workshop is organized at which researchers working on the project present the results of their work both to one another and to visitors from government agencies and laboratories. A hands-on workshop is being planned for 1996 in which environmental scientists will be brought to the College Park campus to learn the high performance computing tools being developed on the project.

A strong outreach activity of UMIACS has been the organization and sponsoring of topical workshops on emerging technology areas. We are planning to continue this activity. We mention two such workshops related to HPC held in 1994. The purpose of the first workshop, held in March 1994 and supported by NSF, was to develop a computer science agenda for high peformance computing. This workshop attracted leading scientists from academia, industry, and government, and led to the publication of a book describing future directions and applications for HPC. The second workshop, being supported by NSF and co-organized with CRPC (Center for Research on Parallel Computation, an NSF Science and Technology center) and Berkeley, will be held in December 1994 and is related to our health care initiative. The purpose of this workshop is to produce an outline of ways in which results of HPCC research may be profitably applied to leading edge health care applications. The result of this workshop will be used to inform NSF as it structures future research initiatives.

EDUCATION

The educational activities of UMIACS will be focused on (1) initiating and sponsoring graduate

interdisciplinary programs in computing, (2) establishing a graduate fellowship program to attract the most talented graduate students to work on interdisciplinary research projects, and (3) to re-establish our summer fellowship program in computer science for talented State of Maryland high-school students. We briefly describe specific strategies to achieve the goal stated in (1); we have already been successful in running the summer fellowship program and we plan to broaden it to attract more students. We should note that UMIACS faculty have been actively involved in various educational initiatives including the establishment of the Electrical Engineering Honors program and the Computer Engineering program to be run jointly by Computer Science and Electrical Engineering.

A new graduate program in *Computational Science* is being developed by faculty from UMIACS and Computer Science. The main goal of the program is to enhance the computational skills of students working toward regular advanced degrees in the sciences and engineering. Students will be admitted to both the Computational Science program and to one of the degree granting programs of the University. The goals of this program are:

1. to educate students from the various disciplines in the common computational elements.
2. to develop courses that unify computational techniques across several disciplines (for example, computational fluid dynamics).
3. to foster interdisciplinary research between graduate students and faculty.
4. to foster interaction with industry and external laboratories.

There is strong support for starting the proposed computational science program which initially will involve departments primarily from CMPS and Engineering, subject to the availability of required resources.

5 Resources

We have outlined our strategic goals in research, outreach, and education, and have outlined specific implementation plans to reach these goals. Our overall strategy is based on (1) maintaining and strengthening our current research activities, (2) starting five new research initiatives, and (3) establishing two educational programs. We believe that, if implemented, such a strategy will have a substantial positive impact on campus activities in computing and will generate internationally visible programs that will be very attractive to industry and government. However, the Institute's current resources will not allow us to push forward this plan in its entirety. While most of the actual costs associated with the research initiatives will be generated through external funding opportunities, some seed money is needed to start these initiatives. Since our current research programs are externally well-funded, we estimate that we will be able to support one or two of the new initiatives through anticipated increases in DRIF funds. The costs associated with each of the remaining initiatives are as follows.

We estimate that each new initiative will require $ 75K-100K (partial support for postdocs, graduate students, workshops, and equipment) and roughly 1,000 square feet for laboratory space, offices for postdocs and graduate students. In addition, these initiatives need the active participation of faculty from outside the Computer Science Department. We estimate the need to increase

our faculty lines by 3 to accomplish the necessary technical breadth. Concerning the computational science program, we estimate that a $ 100K annual budget will be needed to support a quality program. As for the Summer Fellowship program, we are willing to share the costs with the University.

To summarize, our current state budget and external funds will allow us to push forward the major elements in our strategic plan. With the current funding level, we will be able to maintain the excellence we have achieved in research and outreach activities as well as start one or two new initiatives. To carry out our overall strategic plan, the following additional resources will be needed.

1. A 2% increase of the UMIACS State budget per year over the next five years.
2. An additional 3 faculty lines to attract faculty from outside the Computer Science Department.
3. An additional 5,000 square feet of space.