
International Conference on Parallel Computing October 3-4, 1996

On October 3rd and 4th, the Supercomputer Institute, in conjunction with IBM Corporation, sponsored an International Conference on Parallel Computing, held at the Supercomputer Institute. One hundred nineteen registrants from five countries gathered for 17 plenary lectures and nine poster papers on a wide variety of topics in parallel computing. There was also a conference reception and a conference dinner, which featured lively discussion. The exchange of ideas crossed all the traditional boundaries, with academic, industrial, and government labs represented as well as multiple vendors. The talks can be grouped roughly into four areas: fluid dynamics and transport, materials science, physical science, and computer science. Some of the talks are summarized below.

Fluid Dynamics and Transport

Ulrich Trottenberg of the German National Research Center for Information Technology in St. Augustin, Germany discussed the use of the portable CLIC library for solving the 3-D Navier-Stokes equations. One example calculation involved 6,600,000 cells and was solved at 2.8 Gigaflops on a 129-node IBM SP2. Trottenberg singled out fully dynamic adaptive grid structures as the main challenge to efficient parallel computing.

Franz Fiedler of the Institute of Meteorology and Climate Research in Karlsruhe, Germany discussed the use of parallel computing (25 processors) for studying the fate of atmospheric pollutants under various weather conditions. One model he discussed contained 63 substances, 158 reactions, and 21 photolysis reactions coupled to the flow.

Yoichiro Matsumoto of the University of Tokyo presented a Monte Carlo molecular flow simulation that included the rotational motion of individual N_2 molecules.

Denis J. Evans of the Australian National University in Canberra described the use of artificial fields to simulate thermal transport boundary conditions. The new formulation has been tested on a Fujitsu AP10000 parallel computer with 128 processors.

"ICPC..." article *continued*

Positron Emission Tomography Biomedical Image Reconstruction on High-Performance Computing Systems


Biomedical imaging, such as X-ray computed tomography (CT) and magnetic resonance imaging (MRI), has proven to be invaluable in providing physicians and researchers with the ability to observe anatomical structures in living organisms. Another form of imaging technology, positron emission tomography (PET), has extended biomedical imaging to the observation of functional processes in addition to anatomical structures. However, positron emission tomography has been limited by its low photon counting statistics-due to the nature of its signal detection technique-which results in images

with lower resolution and higher noise compared to computed tomography or magnetic resonance imaging.

To improve positron emission tomography image quality and effectively utilize this technology, sophisticated reconstruction algorithms that use substantial computational power are required. The University of Minnesota has great strengths in both high-performance computing and biomedical imaging.

Professor David Lilja and his research group in the Department of Electrical Engineering are capitalizing on these strengths by extending an ongoing interdisciplinary research project between the Departments of Electrical Engineering, Computer Science, and Radiology at the University of Minnesota, and the Positron Emission Tomography Imaging Laboratory at the Veterans Administration Medical Center in Minneapolis. They are focusing on the problem of reconstructing 3-D images of human physiology with data obtained from the PET scanner in the jointly funded Veterans Administration Medical Center/University of Minnesota (VAMC/UM) PET program. In particular, they are implementing a PET image reconstruction algorithm on the Cray T3D multiprocessor system using the standard PVM message-passing library. The outcome of this project is the capability to produce higher quality PET images within a reasonable time frame for clinical research applications. This capability benefits a range of research programs in Radiology, Neurology, Psychiatry, and Cardiology.

figure courtesy of Dr. Jehi-San Liow, Department of Veterans Affairs Medical Center



These positron emission tomography-fluorodeoxyglucose images are from the brain of a normal volunteer. The intensity implies the uptake of fluorodeoxyglucose (FDG) in the brain. FDG is similar to the regular glucose consumed by the tissue except that one of its OH^- is replaced by the positron-emitting isotope F-18 to make it radioactive. The purpose is to get a functional picture of the brain. This method is in contrast to other imaging techniques like CT and MRI which provide only anatomical information.

Workshop on Modeling in Biochemical Engineering October 11-12, 1996

A Conference on Modeling in Biochemical Engineering was held on the Minneapolis Campus of the

University of Minnesota over the weekend of October 11-12, 1996. A select group of leading biochemical engineers and biologists gathered to discuss the current status and future perspectives of modeling and computation in this field. A further purpose of this meeting was to honor University of Minnesota Professor A.G. Fredrickson of the Department of Chemical Engineering & Materials Science for outstanding contributions made in the area of biochemical modeling.

Biochemical engineers offer a unique approach to biotechnology by modeling the biological phenomena of cells. These range from simple, non-interacting bacterial cells to complex tissues where individual cells are able to communicate with each other. A fundamental problem is the particulate nature of such systems, as cells normally occur in populations. A rigorous modeling approach is difficult because such models become mathematically and computationally challenging. In addition, biological details are only partially known at the molecular level-although remarkable advances have been made in recent years.

The goal of the workshop was to reevaluate the role of mathematical modeling in biochemical engineering. Compared to other advances in biotechnology, it may appear that mathematical modeling has played a somewhat minor role. The workshop was, therefore, useful in sensitizing the research community to modeling issues. Furthermore, the discussions helped in firming up the foundation for biochemical engineering by delineating the creative role modeling can play in biotechnology.

About 35 scientists and engineers from the United States, Europe, and Japan were invited to participate in this workshop. The four sessions were held over a two-day period. Each session consisted of four speakers who gave brief overview presentations on a broad topic. These served as the basis for extensive discussion periods following the presentations in each session.

The first session of the conference was devoted to the identification and analysis of parameters characterizing biological systems that are used in modeling. Jim Broach of the Department of Molecular Biology at Princeton University opened by giving an overview of the molecular mechanism used by cells in sensing their environment. His talk was followed by Mike Shuler of the Chemical Engineering Department at Cornell University. Mike discussed the Cornell single-cell model he developed and refined over the past decade. Mike Domach from Carnegie Mellon University then presented some of his work using improved NMR techniques to probe for activities in viable cell populations. Friedrich Srienc from the Department of Chemical Engineering & Materials Science and the Bioprocess Technology Institute of the University at Minnesota followed by describing advanced cytometric techniques that allow identification of parameters used in population balance models.

The second session was devoted to corpuscular models which most accurately describe the behavior of cell populations. John Tyson from the Biology Department of the Virginia Polytechnic Institute gave an overview on modeling of the cell cycle and its control mechanisms. He described the status of growth models on the basis of current knowledge of the molecular biology of the eukaryotic cell cycle. He also described the difficulties in extending a single cell model to a population model. Lorenzo Cazzador from the University of Padua, Italy, gave an overview of cell cycle models as they are used in biochemical engineering. John Villadsen from the Biotechnology Institute at the Technical University of Denmark spoke about population balance models and whether they are truly needed in view of the difficulties in using them. Data for the use of such models typically are difficult to obtain and the solutions are computationally intensive. The session was concluded by Doraiswami Ramkrishna from the Department of Chemical Engineering at Purdue University who introduced his recently developed self similarity theory for identifying population balance models from a wider range of multivariate population data.

The second day opened with a session on metabolic models. Greg Stephanopoulos from the Chemical

Engineering Department at Massachusetts Institute of Technology gave an overview of some of their work on analyzing complex pathways. Jim Liao from the Chemical Engineering Department at Texas A & M University gave a talk on how stoichiometric relationships can be used in combination with metabolic control theory to predict improved pathway designs. Jay Bailey from the Biotechnology Institute at the ETH Zurich, Switzerland, gave a presentation on metabolic modeling and how the description of the intracellular reaction network is affected by the expanding volume of growing cells. Dhinakar Kompala, Department of Chemical Engineering at the University of Colorado at Boulder concluded by describing the application of a cybernetic modeling concept that incorporates a dynamic allocation of resources. The model is able to describe sustained oscillations observed in growing yeast cell populations.

The last session of the meeting was devoted to the most complex problem—the modeling of cell tissue. Since comparatively little is known of the handling of cell tissue systems, this session was the most open ended one. Dennis Bray from the Department of Zoology at the University of Cambridge, England, opened by offering insight into how cells might process transmembrane signals. Doug Lauffenburger from the Chemical Engineering Department at Massachusetts Institute of Technology spoke about how cell receptors lead to cell movement. Bernhard Palsson from the Bioengineering Department of the University of California at San Diego presented a talk on the way in which modeling and engineering concepts contribute to designing cell culture devices and how modeling can help in designing strategies for gene therapy. The session was concluded with a presentation by Wei-Shou Hu, Department of Chemical Engineering & Materials Science at the University of Minnesota, on modeling the development of complex tissues such as embryos of plants.

It appeared most natural and appropriate that the conference was held in honor of A.G. Fredrickson of the University of Minnesota, as he made significant contributions in all areas discussed at the Conference. Professor Fredrickson has made pioneering contributions on each of the session topics during the course of his past work. The conference concluded with a banquet in Fredrickson's honor, where an exhibition of a collection of his absorbing photographs titled "Nature and Man-Made Nature," uniquely framed the evening.

This innovative workshop was made possible by a grant from the National Science Foundation and by additional support from the Supercomputer Institute, the Department of Chemical Engineering & Materials Science, and the Bioprocess Technology Institute at the University of Minnesota. The Conference was co-chaired by Doraiswami Ramkrishna, Department of Chemical Engineering, Purdue University and Friedrich Sreng, Department Chemical Engineering & Materials Science and the Bioprocess Technology Institute at the University of Minnesota. They were assisted in the organization by Mike Flickinger, Department of Biochemistry and the Bioprocess Technology Institute, and by Robert Tranquillo and Wei-Shou Hu of the Department of Chemical Engineering & Materials Science at the University of Minnesota.

Evolutionary Algorithms Workshop, October 21-25, 1996

During the week of October 21, the University of Minnesota Supercomputing Institute and Institute for Mathematics and its Applications (IMA) hosted a Workshop on Evolutionary Algorithms, held on the East Bank of the University campus. There were 66 registrants from 20 states and 9 countries, 19 plenary presentations, numerous informal discussion periods, and a banquet. The talks can be grouped roughly into applications and theory.

Applications

The first talk in the applications category was a survey of several genetic algorithm (GA) case studies by David Levine from the Argonne, Illinois offices of the Boeing Company. An example from ligand-receptor docking in biochemistry involved a 297-atom drug candidate binding to a 1564-atom protein. The problem was studied on 101 processors of an IBM SP2 with a parallel speedup of 63. A second example, involving collaboration with Paul Bash of Argonne National Laboratory and L. Ho of Yale University, involved parameter fitting in quantum chemistry to obtain system-specific parameters (SSP-also called SRP for specific range parameters) for an AM1-type ("Austin-method 1") model of small molecules. A third example involved reverse-engineering an integrated circuit. Levine also discussed the use of symmetric multi-processor (SMP) machines for island-based GAs, the advantages of coding with MPI communicators, and a general purpose parallel genetic algorithm library called PGAPack, callable from FORTRAN or C.

David Davis of Tica Technologies Inc. in Cambridge, Massachusetts described applications of GAs to the optimization of telecommunications networks. Shunji Matsumoto of Fujitsu Limited in Chiba, Japan and Peter Ross of the University of Edinburgh in Scotland presented applications of GAs in scheduling problems. A theme that emerged strongly in these talks is that cost-effective GAs have been and can be designed for practical applications by empirical means. "Theory lags behind," said one speaker, but we do not have to wait for it to catch up to solve commercial problems. In addition to scheduling, Ross discussed time tabling, pipe routing, and a host of other applications, including the optimization of catching, delivering, and slaughtering chickens raised by a meat distributor.

Emmanuel Falkenaur of the Free University of Brussels discussed applications of GAs to combinatorial problems. He found theory more useful, especially in its pointing to crossover as the most potent force in a GA. He described the advantages of Grouping GAs and illustrated their use for bin packing, load balancing, and optimizing economies of scale.

Martin Wildberger of the Electric Power Research Institute in Palo Alto, California presented several GA applications, including one in collaboration with Tariq Samad and Steve Harp of Honeywell in Minneapolis for GA optimization of a neural network based on heat rate data from a nuclear power plant.

Rogene Eichler West of the University of Minnesota discussed the application of evolutionary algorithms to parameter selection in models of biologically realistic neurons.

Zbigniew Michalewicz of the University of North Carolina discussed evolutionary algorithms for constrained parameter optimization.

Robert Meyer of the University of Wisconsin showed that GAs are effective for optimal constrained domain decomposition problems that arise in network design and analysis. John Clymer of John R. Clymer and Associates in Placentia, California discussed context-sensitive methodology for complex adaptive analysis of vehicle traffic control.

Theory

The theory talks covered genetic algorithms, evolutionary programming (EP), and evolutionary strategies (ES).

There were five theoretical talks on genetic algorithms. Darrell Whitley of Colorado State University started the workshop with an overview of GAs, focusing on selection, recombination, and mutation strategies. He was the first of many speakers to address the No Free Lunch (NFL) theorem, which states that all search algorithms are equivalent when compared over all possible discrete functions. One participant pointed out in a later discussion that the set of interesting functions has zero measure on the set of all functions, and a theme that emerged from the workshop is that the best algorithm for a given problem can generally be engineered only by building in nonartificial intelligence about the nature of

the application of interest.

David Goldberg of the University of Illinois at Urbana presented his perspective on the issue of efficiency. The basic problem of GAs that he addressed is that population sizes must grow exponentially for hard problems. He advocated the empirical design of genetic algorithms to find the "sweet spot" in the 2-D strategy parameter space of mixing probability (achieved, for example, by crossover) vs. selection pressure. Goldberg placed special emphasis on the balance between critical innovation and takeover time; one tries to allow enough time for linkage recognition before the entire population is taken over by individuals optimized with the linkages recognized so far. One hopes to develop competent GAs by creative combinations of building blocks.

John Holland of the University of Michigan emphasized several issues that one should take into consideration, including time to discovery (how long do we have to wait for the first occurrence of a genotype?), NFL, premature convergence due to hitchhiking, the danger of relying on inversion if less than 10,000 generations are sampled, and the possibility that the error cone diverges exponentially in which case almost all offspring trajectories may be below average. The latter phenomenon was related to the "Petersburg Paradox," and to the best of the knowledge of this reporter, none of the participants suggested calling it the Inverse Lake Wobegon Phenomenon.

Michael Vose of the University of Tennessee gave a mathematical perspective on GAs, and Kenneth De Jong of George Mason University in Fairfax, Virginia itemized the choices one faces when parallelizing GAs. De Jong concluded that no clear winner has emerged from studies of parallelization strategy.

Heinz Mühlenbein of the German National Research Center for Information Technology presented a theory of the breeder GA, which is designed according to the science of livestock breeding. Critical new elements in the theory include response to selection and realized heritability.

The theory of EP was presented by David Fogel of Natural Selection Inc. in La Jolla, California. Fogel emphasized the simulation of evolution as a learning process to generate artificial intelligence. The goal is to optimize the variation operator by looking at the distribution function.

There were two back-to-back talks on the ES perspective, by Hans-Paul Schwefel and Thomas Bäck of the University of Dortmund and Informatik Centrum Dortmund in Germany. Schwefel emphasized learning on the fly by optimizing genetically controlled strategy parameters as well as objective variables. He illustrated this with self learning of a variable metric controlling linearly correlated mutations, which is analogous to regulatory genes controlling mutability. He pointed out that if K represents the upper limit for the generational life span of a parent, the optimum value is empirically bigger than one and less than infinity, but the former is better than the latter. Schwefel also emphasized the desirability of multiple-criteria decision making. Bäck continued with a discussion of self-adaption of strategy parameters as applied to GAs.

Concluding Remarks

The Institute is grateful to Darrell Whitley, Michael Vose, David Davis, and Ken De Jong for organizing the Workshop.

The workshop was a successful mixing of theoretical and practical cultures and it engendered a higher than average amount of discussion during and after the talks. The workshop also gets high marks in that only one speaker (who shall remain nameless) succumbed to making jokes about Minnesota weather. The next joint workshop of the Supercomputing Institute and the Institute for Mathematics and its Applications will be focused on mathematical issues in drug design and will take place in April, 1997 (see Rational Drug Design Workshop for details).

Workshop on Rational Drug Design, April 7-11, 1997

The Supercomputer Institute and the Institute for Mathematics and its Applications are sponsoring a Workshop on Rational Drug Design. Please note that the dates for this conference have been changed to April 7-11, 1997.

The goal of this workshop is to bring together physical and computational scientists and mathematicians to examine the current state of this broad field of research, and to identify the areas where cross-fertilization of ideas and collaborative research might most effectively advance the field.

**For further information,
please refer to the
Workshop on Rational Drug Design World Wide Web announcement.**

Finite Element Modeling of Shallow Water Flows

Kazuo Kashiya
Department of Civil Engineering
Chuo University
Tokyo, Japan

Finite element computations of shallow water flows can be applied to many practical problems: design of river, coastal and offshore structures, disaster prediction, and other applications related to hydrodynamics, thermal and chemical transport behavior in oceans, lakes and rivers.

In this lecture, recent advances in finite element modeling of shallow water flows were explained and two finite element methods were presented. In addition, numerical examples were used to demonstrate that the strategies presented are applicable to large-scale computations of various shallow water flow problems.

Atomistic Force Fields for Large Molecules

Grant Smith
Department of Chemical Engineering
University of Missouri
Columbia, Missouri

Finite element computations of shallow water flows can be applied to many practical problems: design of river, coastal and offshore structures, disaster prediction, and other applications related to hydrodynamics, thermal and chemical transport behavior in oceans, lakes and rivers.

The utility of quantum chemistry calculations in the parameterization of atomistic force fields for molecular mechanics applications has been thoroughly demonstrated for relatively simple molecules such as alkanes and ethers. For aromatic and other large conjugated molecules, it is a computationally demanding task. Quantum chemistry codes developed specifically for workstation platforms have been successfully employed for deriving force fields for large molecules. Several of these systems, and issues involved with large-scale quantum chemistry calculations, including use of parallel architectures, were discussed in this seminar.

Scheduling From the Perspective of the User

Francine Berman
Department of Computer Science and Engineering
University of California at San Diego
La Jolla, California

Fast networks have made it possible to aggregate distributed central processing units, memory, and storage resources to provide the potential for application performance superior to that achievable on a single system. These aggregate resources enable the implementation of problems which could not previously be solved and constitute a new generation of high-performance programs.

However, achieving the performance potential of such applications can be difficult. Since resources of the metasystem often reside in distinct administrative domains, and are under the control of distinct local schedulers, it is not feasible to rely on a single system scheduler to deliver application performance. Metacomputing applications must be scheduled carefully using application-specific and dynamic system information in order to leverage heterogeneity, minimize the effects of contention for shared resources, and satisfy the developer's performance criteria.

The speaker introduced application-level scheduling, a scheduling paradigm in which all resources in the metasystem are evaluated in terms of their impact on an application. The principles of application-level scheduling are being used to guide the development of AppLeS Application-Level Scheduling agents for individual metacomputing applications. The AppLeS architecture and description of experiments with AppLeS prototypes was also outlined.

Vortex Shedding and Three-Dimensionality in Bluff Body Wakes and Their Role in Wake Cavity Modeling

S. Balachandar
Department of Theoretical and Applied Mechanics
University of Illinois at Urbana
Urbana, Illinois

This talk began with an investigation of instability and the transition to three-dimensionality in cylinder wakes of different cross-sections. Results from Floquet analysis and direct numerical simulations on cylinders of circular, elliptic, and square cross-sections were presented. These results show that different modes of instability are present and that their interaction can be quite complex, resulting in significantly distorted three-dimensional flow even at modest Reynolds numbers. The possibility of a period doubling scenario arising from spanwise subharmonic instability was also discussed. Consequences of the appearance of three-dimensionality were elucidated.

High-Performance Distributed Computing: The I-WAY Experiment and Beyond

Ian Foster
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, Illinois

Recent developments in networking make it feasible to construct high-performance computations that couple geographically distributed computing, information, and display resources. The I-WAY was an ATM-based wide area computing environment designed to support experimentation in this area. Its use was demonstrated with over 60 applications at the Supercomputing conference in December 1995. This researcher described the I-WAY experiment and the lessons that have been learned from it. Also

discussed was ongoing research that seeks to develop a software infrastructure, called Globus, to support applications such as those deployed on the I-WAY.

Simulation of Dendritic Solidification Using the Phase-Field Method

William Boettinger
Metallurgy Division
National Institute of Standards and Technology
Gaithersburg, Maryland

Under many circumstances a material solidifies (crystallizes) from its melt in a complex geometrical pattern-called a dendrite-which separates the liquid and solid regions during the cooling process. This microscopic tree-like structure occurs because it efficiently disposes of the heat of solidification for pure materials, and excess solute for alloys. Recently, considerable progress has been made with the simulation of these patterns using the phase-field method. The formulation of the model for pure materials and alloys was described in this talk. Results showing the development of segregation patterns within the solid dendrites were highlighted as well.

On the Initial Photochemical Event in Photosynthesis: A Theorist's View

Michael Zerner
Departments of Chemistry and Physics
The University of Florida
Gainesville, Florida

Photosynthesis is the process by which light energy is converted into chemical energy to synthesize sugars and starches from atmospheric carbon dioxide. It is the principal source of the biomass on Earth and is responsible for creating an oxidizing atmosphere above the surface of the Earth.

This seminar described the electronic spectroscopy of chlorophyll and explain why chlorophyll is a particularly good choice to be the central molecule in this process. It also described the excited states of the "special pair," a dimer of chlorophyll, and showed why this again is a unique structure, guaranteeing that the energy is transferred from the antenna chlorophyll to the reaction center.

Kinetic Partitioning Mechanism as a Unifying Theme in the Folding of Proteins and RNA

Devarajan Thirumalai
Institute of Physical Science and Technology and Department of Chemistry
University of Maryland
College Park, Maryland

Conflict in the interactions on short and long length scales in biomolecules leads topological frustration. A consequence of topological frustration is that biomolecular folding takes place by distinct mechanisms which can be succinctly summarized in terms of the kinetic partitioning mechanism (KPM).

In this seminar, estimates of the various dominant time scales in the KPM were provided. Experimental evidence of KPM for kinetics of refolding of proteins and RNA was also provided.