Editor: Daniel Pryor, pryor@super.org

"International Conference on Parallel Computing," D. G. Truhlar, IEEE Computational Science & Engineering **4**/No. 1, 98-100 (1997).

INTERNATIONAL CONFERENCE ON PARALLEL COMPUTING

Donald G. Truhlar

On October 3–4, 1996, over a hundred participants from five countries gathered for 17 plenary lectures and nine poster papers on a variety of topics in parallel computing at the University of Minnesota Supercomputer Institute. Arranged by an international advisory committee chaired by Inder Batra of IBM, the program crossed all the traditional boundaries, aside from its emphasis on parallel computing success stories. Academic, industrial, and government labs were represented, and a variety of work was carried out on machines from multiple vendors. Some of the talks are summarized below, with an emphasis on the sizes of the problems that can currently be tackled, the efficiencies researchers are achieving in the area of parallel speedup, and the applications difficulties encountered.

Fluid dynamics and transport

Ulrich Trottenberg of the German National Research Center for Information Technology discussed use of the portable CLIC library for solving the 3D Navier-Stokes equations. He singled out fully dynamic adaptive grid structures as the main challenge to efficient parallel computing. One example calculation involved 6,600,000 cells and was solved at 2.8 gigaflops on a 129-node IBM SP2. 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 contains 63 substances, 158 reactions, and 21 photolysis reactions coupled to the flow. These talks provided prime examples of parallel computers at work on complex problems that provide a challenge to any computing paradigm.

Materials science

Priya Vashishta of Louisiana State University presented parallel molecular dynamics calculations of mechanical failure, amorphization, and sintering in amorphous silicon nitride and silicon selenide films. The method employed is multiresolution in both space and time, using a curvilinear coordinate transformation, based in part on the work of François Gygi. It employs uniform mesh partitioning in curved space and dynamic load balancing every 50–100 steps. Vashishta reported simulations of up to 135 million particles. Considerable insight can be gained from these simulations. For example, in SiSe2, amorphization proceeds with no change in the density profile; then fracture proceeds along the amorphous-phase–crystal interface.

Alan Needleman of Brown University talked about predicting crack growth in materials like structural stainless steel or polymethylmethacrylate glasses. In many circumstances the processes that determine cracks and cleavage take place on a scale of 10–100 microns. One must calculate the energy relaxation rate on that length scale while ensuring that the cohesive length dominates the discretization length. Needleman attacked the problem in one case via data-parallel implementation on a Thinking Machines CM-5, yielding a factor of 32 speedup for a 32-fold increase in number of processors (from 32 to 1,024). On a 30-processor SP2, he obtained 97 percent efficiency for a problem with over two million degrees of freedom.

Jeremy Broughton of the US Naval Research Laboratory discussed the atomistic simulation of microelectronic systems. He posed the question: Does bulk continuum mechanics fail for a micromechanical device such as a 1 micron ¥ 100 micron sensor? He answered by simulating a system with 2.3 million atoms on 100 processors. With coding in Fortran and MPI he was able to get 98 percent processing and only 2 percent communications, but data storage between runs turned out to be a bottleneck. Using the multiple-time-step algorithm of Tuckerman, Berne, and Martyna, he achieved high efficiency in solving the equations of motion. Nanoscale devices appear to be an area of technology that may revolutionize many branches of science, engineering, and technology. This area of research should be of great interest to readers of *IEEE CS&E* since it is clear that atomistic simulations can guide the design of these devices in ways where obtaining the required information by experiment is clearly infeasible.

Physical science

Gyan Bhanot of IBM Thomas J. Watson Research Center described nine applications for the IBM SP2. A particularly intriguing example was the application of fuzzy text-matching technology, developed for voice recognition, to the problem of molecular

docking at a protein surface. This illustrates how a new interdisciplinary field arises as workers recognize that elements of algorithms developed for one scientific application are precisely what is needed for an entirely different one.

Robert Beck, University of Chicago, surveyed the use of parallel computing in biomedical imaging. He included image-data acquisition, image reconstruction and processing, image storage and distribution, image display and analysis, and the optimization parameters associated with these steps (in the optimal design of detector systems, for instance). It has always been recognized that parallel computing is especially needed for applications where real-time processing is important. The promise of interpreting patient scanning data while the subject is still in the diagnostic center can lend a new dimension to the use of imaging tools in medicine.

Eric Stahlberg of Oxford Molecular Group, Burnsville, Minnesota, discussed the Unichem project for portable, parallel, multiarchitecture electronic-structure calculations on chemical systems. He discussed applications on both the Cray T3D and IBM SP2. One computing strategy in this code suite is the use of dynamically load-balanced distributed-memory nonsynchronous updates of the Fock matrix in the DGauss DFT code.

John Carpenter of Cray Research (now owned by Silicon Graphics) discussed the quest for scalable programs for electronic structure and molecular dynamics. Interestingly, he defined scalable, in practice, as 78 percent efficiency on 128 processors, while in the electronic-structure area he discussed a problem that achieved 77 percent efficiency (1 percent short of goal, John!) on 128 processors for 925 basis functions. He identified the dynamical bottleneck for the algorithm as a global sum. In the molecular dynamics area, he discussed two 128-processor calculations on a 7,142-atom phospholipid biomembrane—a calculation with electrostatic cutoffs that achieved 75 percent efficiency and a calculation with a particle mesh algorithm that achieved 55 percent. This illustrates how different algorithms for the same problem have different potentials for parallel speedup.

Michel Dupuis of Pacific Northwest National Laboratory discussed the use of Global Array Tools in the NWChem quantum chemistry and molecular dynamics software. The PNNL group, including Jeff Nichols, Robert Harrison, Rick Kendall, and Tjerk Straatsma, is using this software for various problems in environmental chemistry. In the programming area, Dupuis emphasized the advantages of interrupt-driven communication. He discussed the case of a four-index transformation that achieved 78 percent efficiency on 128 nodes. Michael Colvin of Sandia National Laboratories presented a parallel ab initio quantum chemistry program based on C++ objects. He reported good scaling of Hartree-Fock calculations on 14 quad-processor Pentium Pros connected by fast ethernet (63 percent efficiency on 24 processors). He also reported a 2,356-basis-function calculation with no symmetry on phosphoramide amide mustard; the calculation required 12 hours per SCF iteration on a 512-processor Intel Paragon. In these two talks we are beginning to see the emergence of a powerful new generation of codes written specifically for the new computing models, rather than retrofitting old Cray-1-type single-processor vector codes. See http://www.emsl.pnl.gov:2080/emslweb/orgs/hpcc.html and http://midway.ca.sandia.gov/~mecolv/ for more information on these exciting new developments in parallel software engineering.

Computer science

Yousef Saad of the University of Minnesota discussed the challenges for developing sparse iterative methods on distributed-memory computers and listed three: preconditioners, preconditioners, and preconditioners. Can we define good preconditioners specifically for distributed processing? Saad reviewed the pros and cons of several choices, including a new method that filters high-frequency components. He also discussed the solution of large sparse eigenvalue problems that arise in electronic-structure problems in materials science. He touched on the intriguing idea of computing an electronic density without computing eigenvectors (by using projection operators). Andreas Stathopoulos of the University of Minnesota, collaborating with Saad and Jim Chelikowsky, also of the University of Minnesota, discussed porting electronic-structure calculations to the SP2 using the master-slave paradigm. Stathopoulos discussed the necessity of using MPI collective operations such as `allreduce` and `alltoallv` to obtain good speedups. Both Saad and Stathopoulos discussed the advantages of a preconditioned Davidson algorithm for the computationally demanding eigenvalue step.

David Ceperley of the University of Illinois at Urbana presented a scalable pseudorandom number-generator library for parallel Monte Carlo calculations and illustrated it by application to a 200-particle quantum path integral calculation. He raised interesting questions on how to efficiently satisfy the demand for independent strings of high-quality random numbers on multiple processors while minimizing communications.

Nanoscale devices may revolutionize many branches of science, engineering, and technology.

Geoffrey Fox of Syracuse University closed the conference by presenting his vision of petaflop/terabyte computing in the 2005–2010 time period. He gave reasons why Java will have replaced Fortran and C++ by then (good graphics, Web linkage, object structure).

With youngsters learning Java as their first programming language, he suggests a "bottom-up revolution" will sweep over the world when they reach graduate school and inform their research advisors that they won't accept Fortran in place of Java. Fox discussed the design and building of Java wrappers to existing Fortran 77 and Fortran 90 code, making me feel more comfortable that I may not have to throw away all my legacy code after all. He also discussed the increasing necessity of addressing memory hierarchies. As for hardware 10 years down the pike, he presented a detailed comparison of the prospects for superconducting machines versus worldwide metacomputers. Of course, time will tell which of his visionary scenarios will come true, but it left everybody with considerable food for thought. ♦

Donald G. Truhlar is Institute of Technology Professor of chemistry and director of the Supercomputer Institute at the University of Minnesota, Minneapolis. He is also an IEEE CS&E *area editor for applications.*

High-Performance Computer Applications in the Behavioral Sciences

Lynne K. Edwards

This symposium was a follow-up to the 1995 US National Science Foundation conference on Advanced Computing for Psychology, which examined the increasing importance of high-performance computers, including workstation clusters, in the behavioral sciences. This May 10–12, 1996, conference at the Minnesota Supercomputer Institute included such topics as computer-intensive simulation methods, virtual reality, neural networks, neural magnetic resonance imaging, information processing, large-data handling, perception and human vision, and graphics visualization.

Psychologists Richard Shiffrin and Peter Nobel, Indiana University, discussed methods for building mathematical models of human memory. They spoke about the function of high-performance computing in building and testing models. These talks illustrate the emergence of scientific computation as a cross-disciplinary field in its own right, since many of the problems and experiences presented were similar to those encountered in model-building endeavors in other disciplines.

Richard Golden, University of Texas at Dallas, described a Markov random field probabilistic model where subjects recall events in a text from memory. James Cutting, Cornell University, used paintings to show that the human visual system exhibits tolerance for some deviations from a Euclidean representation of space, and an intolerance for others. In the pictures perception cues were provided by occlusion, height in the visual field, relative sizes and densities, binocular disparities, motion perspectives, and the like.

Daniel Kersten, University of Minnesota, presented 3D animation tools that produce realistic movies with strong perceptual cues. The goal of his work is to increase understanding of the human visual system by observing the cues and responses of viewers. Kersten has applied Bayesian analysis to understanding the identity of objects and their spatial relationships by cues such as cast shadows. In the afternoon, Mary Kaiser, NASA Ames Research Center, discussed how the visualization of a particular data set (the Digital Terrain Model of Mars, which was derived from Viking Orbiter imagery) can be optimized using our knowledge of human perception. Jeffrey Mulligan, also from Ames, demonstrated the use of a video camera to track human eye movements. He noted that the real-time digitization and storage of large video imagery is a major bottleneck, but that recent developments in video compression hardware have made it easier and cheaper to manage these tasks. Images from the retina and the pupil can be analyzed with basic image-processing tools such as filtering, correlation, and thresholding, all of which are well-suited for implementation on vectorizing supercomputers.

Sam Williamson, New York Univer-