SC97 CONFERENCE PROCEEDINGS

SC97: HIGH PERFORMANCE NETWORKING AND COMPUTING

SAN JOSE SC97 NOV. 15-21

SPONSORED BY: ACM SIGARCH AND IEEE COMPUTER SOCIETY

ACM SIGARCH Order # is 415972.
IEEE Computer Society Press Order # is RS00160.
ISBN: 0-89791-985-8

Copies may be ordered from ACM Member Services:

1-800-342-6626 (U.S. and Canada)
+1 212-626-0500 (Global)
+1 212-944-1318 (Fax)
Email: orders@acm.org
Online: http://www.acm.org/catalog

SPONSORED BY: ACM SIGARCH AND IEEE COMPUTER SOCIETY

The SC97 Technical Program, presented Tuesday through Friday, November 18-21, includes both plenary and parallel sessions. The parallel sessions have been organized to ensure that topics of interest to a wide range of audiences are available during these times. The parallel tracks are divided into two components as follows:

Presentations by authors of contributed papers. These represent the very best papers submitted to the conference. Each has been through a competitive peer review process and has been selected as an outstanding contribution. There will be a total of 57 technical papers presented at SC97.

Topics of general interest in a wide variety of areas. Sessions will include invited talks on a variety of subjects, presented by an impressive array of experts, as well as panel discussions that are sure to engage attendees' interest.

Tuesday, November 18 -- 10:30am-noon

[Room B1] Communication in Clusters
Session Chair: David Culler, University of California, Berkeley

Abstracts

1. NOMINATED STUDENT PAPER AWARD The Effects of Communication Parameters on End Performance of Shared Virtual Memory Clusters **Angelos Bilas**, Jaswinder Pal Singh, Princeton University

2. FM-QoS: Real-time Communication Using Self-synchronizing Schedules **Kay H. Connelly**, Andrew A. Chien, University of Illinois at Urbana-Champaign

3. NOMINATED STUDENT PAPER AWARD Multi-protocol Active Messages on a Cluster of SMPs **Steven S. Lumetta**, Alan Mainwaring, David E. Culler, University of California, Berkeley

[Room A4] Astrophysics and Hydrodynamics
Session Chair: Subodh Kumar, Johns Hopkins University

Abstracts

1. Highly Portable and Efficient Implementations of Parallel Adaptive N-Body Methods **David Blackston**, University of California, Berkeley; Torsten Suel, Bell Laboratories

2. Optimization and Scaling of Shared-memory and Message-passing Implementations of the Zeus Hydrodynamics Algorithm **Robert A. Fiedler**, Hewlett-Packard Company

3. Scalable, Hydrodynamic, and Radiation-hydrodynamic Studies of Neutron Stars Mergers **Paul E. Saylor**, University of Illinois at Urbana-Champaign; Frank D. Swesty, University of Illinois at Urbana-

Champaign

[Room A3] Performance Tools
Session Chair: Ann Hayes, Los Alamos National Laboratory

Abstracts

1. Implementing a Performance Forecasting System for Metacomputing: The Network Weather Service **Rich Wolski**, Neil Spring, University of California, San Diego; Chris Peterson, Microsoft Corporation
2. NOMINATED STUDENT PAPER AWARD Experiment Management Support for Performance Tuning **Karen L. Karavanic**, Barton P. Miller, University of Wisconsin-Madison
3. Exploiting Global Input/Output Access Pattern Classification **Tara M. Madhyastha**, Daniel A. Reed, University of Illinois at Urbana-Champaign

Tuesday, November 18 -- 1:30-3pm

[Room B1] Compiling Techniques
Session Chair: David Callahan, Tera Computer Co.

Abstracts

1. NOMINATED STUDENT PAPER AWARD Compiling Parallel Code for Sparse Matrix Applications **Vladimir Kotlyar**, Keshav Pingali, Paul Vinson Stodghill, Cornell University
2. Evaluating the Performance Limitations of MPMD Communication **Chi-Chao Chang**, Grzegorz Czajkowski, Thorsten von Eicken, Cornell University; Carl Kesselman, USC/Information Sciences Institute
3. Compiling Stencils in High Performance Fortran **Gerald Roth**, John Mellor-Crummey, Ken Kennedy, Rice University; R. Gregg Brickner, Los Alamos National Laboratory

[Room A4] Computational Fluid Dynamics
Session Chair: Horst Simon, Lawrence Berkeley Laboratory

Abstracts

1. MPP Solution of Rayleigh-Benard-Marangoni Flow **Graham F. Carey**, Christopher Harle, Robert Mclay, Spencer Swift, The University of Texas at Austin
2. A Multi-level Parallelization Concept for High-fidelity Multi-block Solvers **Ferhat F. Hatay**, MCAT, Inc. (presently with HaL Computer Systems, Inc.); Dennis C. Jespersen, Guru P. Guruswamy, Yehia M. Rizk, NASA Ames Research Center: Chansup Byun, Ken Gee, MCAT, Inc.
3. On Parallel Implementations of Dynamic Overset Grid Methods **Andrew M. Wissink**, MCAT, Inc., NASA Ames Research Center; Robert L. Meakin, NASA Ames Research Center

Tuesday, November 18 -- 3:30-5pm

[Room B1] High Performance Fortran

Session Chair: David Bailey, NASA AMES Research Center

Abstracts

1. New Life in Dusty Decks: Results of Porting a CM Fortran-based Aeroacoustic Model to High Performance Fortran **Jeffrey J. Nucciarone**, Penn State University; Yusuf Ozyoruk, Lyle N. Long, Penn State University
2. An Evaluation of HPF Compilers and the Implementation of a Parallel Linear Equation Solver Using HPF and MPI **Konming Gary Li**, Nabil M. Zamel, Saudi Aramco
3. Portable Performance of Data Parallel Languages **Ton A. Ngo**, IBM T.J. Watson Research Center; Lawrence Snyder, Bradford Chamberlain, University of Washington

[Room A4] Visualization Algorithms

Session Chair: Tom DeFanti, EVL, University of Illinois at Chicago

Abstracts

1. Divide and Conquer Spot Noise **Wim de Leeuw**, Center for Mathematics and Computer Science, Kruislaan, Amsterdam, The Netherlands
2. Wavelet-based Image Registration on Parallel Computers **A. El-Ghazaw**, Prachya Chalermwa, The George Washington University; Jacqueline LeMoigne, NASA
3. Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments **Jason Leigh**, EVL, University of Illinois at Chicago

Wednesday, November 19 -- 10:30am-noon

[Room B1] Techniques for Metacomputing

Session Chair: Francine Berman, University of California, San Diego

Abstracts

1. Multi-client LAN/WAN Performance Analysis of Ninf: A High Performance Global Computing System **Atsuko Takefusa**, Umpei Nagashima, Ochanomizu University; Satoshi Matsuoka, Tokyo Institute of Technology; Hirotaka Ogawa, University of Tokyo; Hidemoto Nakada, Satoshi Sekiguchi, Electrotechnical Laboratory; Hiromitsu Takagi, Nagoya Institute of Technology; Mitsuhisa Sato, Real World Computing Partnership;
2. PARDIS: CORBA-based Architecture for Application-level Parallel Distributed Computation **Katarzyna Keahey**, Dennis Gannon, Indiana University
3. Scalable Networked Information Processing Environment (SNIPE) **Keith Moore**, Graham E. Fagg, Al Geist, University of Tennessee; Jack Dongarra, University of Tennessee/ Oak Ridge National Laboratory

[Room A4] Climate
Session Chair: Dennis Gannon, CRPC/Indiana University

Abstracts

1. Optimization of a Parallel Ocean General Circulation Model **Ping Wang**, Daniel S. Katz, Yi Chao, Jet Propulsion Laboratory

2. Parallel Computing at NASA Data Assimilation Office **M P. Lyster**, K. Ekers, J. Guo, M. Harber, D. Lamich, J.W. Larson, R. Lucchesi, R. Rood, S. Schubert, W. Sawyer, M. Sienkiewicz, A. da Silva, J. Stobie, L.L. Takacs, R. Todling, J. Zero, NASA/Goddard Space Flight Center; C.H.Q. Ding, Lawrence Berkeley National Laboratory; R. Ferraro, Jet Propulsion Laboratory

3. FOAM: An Atmosphere-ocean Climate Model for Studying Long-duration Phenomena **Michael Tobis**, Ian T. Foster, Chad M. Schafer, Argonne National Laboratory; Robert M. Jacob, John R. Anderson, University of Wisconsin-Madison

[Room A3] Numerics
Session Chair: David Kahaner, ATIP and NIST

Abstracts

1. Parallel Threshold-based ILU Factorization **George Karypis**, Vipin Kumar, Computer Science, University of Minnesota

2. PLAPACK: Parallel Linear Algebra Libraries Design Overview **Philip Alpatov**, Gregory Baker, H. Carter Edwards, John Gunnels, Greg Morrow, James Overfelt, Robert van de Geijn, The University of Texas at Austin

3. MultiMATLAB: Integrating Matlab with High Performance Parallel Computing **Vijay Menon**, Anne E. Trefethen, Cornell University

Wednesday, November 19 -- 1:30-3pm

[Room B1] Big I/O and Checkpointing
Session Chair: Karim Harzallah, Tandem Computers

Abstracts

1. Optimization and Evaluation of Hartree-Fock Applications' I/O with PASSION **Meenakshi A. Kandaswamy**, Mahmut T. Kandemir, Alok N. Choudhary, Syracuse University/ Northwestern University; David E. Bernholdt, Syracuse University

2. A Checkpointing Strategy for Scalable Recovery on Distributed Parallel Systems **Vijay K. Naik**, Samuel P. Midkiff, Jose E. Moreira, IBM T. J. Watson Research Center

3. CLIP: A Checkpointing Tool for Message-passing Parallel Programs **Yuqun Chen**, Kai Li, Princeton University; James S. Plank, University of Tennessee

[Room A4] Molecular Dynamics
Session Chair: Roy Richter, General Motors Research and Development Center

Abstracts

1. Massively Parallel Simulations of Diffusion in Dense Polymeric Structures **Jean-Loup M. Faulon**, Sandia National Laboratories
2. Molecular Dynamics Simulation of Large-scale Carbon Nanotubes on a Shared-memory Architecture **Deepak Srivastava**, Stephen T. Barnard, MRJ Technology Solutions, NASA Ames Research Center
3. Topology Preserving Dynamic Load Balancing for Parallel Molecular Simulations **David F. Hegarty**, Tahar M. Kechadi, University College Dublin

Wednesday, November 19 -- 3:30-5pm

[Room B1] Real Iron
Session Chair: Forest Baskett, Silicon Graphics, Inc.

Abstracts

1. The Starfire SMP Interconnect **Alan Charlesworth**, Sun Microsystems
2. Tera Hardware-Software Cooperation **Gail Alverson**, Preston Briggs, Susan Coatney, Simon Kahan, Rich Korry, Tera Computer Company
3. Performance Analysis of the T3E Multiprocessor **Ed C. Anderson**, Jeffrey P. Brooks, Charles M. Gassi, Steven L. Scott, Cray Research

[Room A4] Climate and Fusion
Session Chair: Tony Hey, University of Southampton

Abstracts

1. Distributed High Performance Computation for Remote Sensing **Kenneth A. Hawick**, Heath A. James, University of Adelaide
2. Parallel Non-linear Optimization: Toward the Design of a Decision Support System for Air Quality Management **Andrew Lewis**, Griffith University, Australia; Rod Simpson, Griffith University; David Abramson, Monash University, Australia
3. Numerical Tokamak Turbulence Calculations on the CRAY T3E **Vickie E. Lynch**, Jean-Noel Leboeuf, Benjamin A. Carreras, Juan Diego Alvarez, Oak Ridge National Laboratory; Luis Garcia, Universidad Carlos III de Madrid

Thursday, November 20 -- 10:30am-noon

[Room B1] Memory Hierarchies
Session Chair: Daniel Reed, University of Illinois

Abstracts

1. Performance Evaluation of the SGI Origin2000: A Memory-centric Characterization of LANL ASCI Applications **Olaf M. Lubeck**, Yong Luo, Harvey J. Wasserman, Federico Bassetti, Los Alamos National Laboratory

2. A Study of Performance on SMPs and Distributed Memory Architectures Using a Shared-memory Programming Model **Eugene D. Brooks III**, Karen H. Warren, Lawrence Livermore National Laboratory

3. Measuring Memory Hierarchy Performanceof Cache-coherent Multiprocessors Using Micro Benchmarks **Cristina Hristea**, MIT/SGI; Daniel Lenoski, John Keen, SGI

[Room A4] Subsystem Performance
Session Chair: Andrew Chien, University of Illinois and Hewlett-Packard

Abstracts

1. High Performance Software on Intel Pentium Pro Processors or Micro-Ops to TeraFLOPS **Bruce Greer**, Greg Henry, Intel Corporation

2. Parallel Simulation of Parallel File Systems and I/O Programs **Rajive Bagrodia**, Stephen Docy, Andy Kahn, UCLA

3. NOMINATED STUDENT PAPER AWARD A Scalable Mark-sweep Garbage Collectoron Large-scale Shared-memory Machines **Toshio Endo**, Kenjiro Taura, Akinori Yonezawa, The University of Tokyo

[Room A3] Working Parallelism
Session Chair: Jack Dongarra, University of Tennessee and ORNL

Abstracts

1. Parallel Database Processing on a 100 Node PC Cluster: Cases for Decision Support Query Processing and Data Mining **Takayuki Tamura**, Masato Oguchi, Masaru Kitsuregawa, The University of Tokyo

2. Page Replacement Using Marginal Loss Functions **Shamik Das Sharma**, Joel Saltz, University of Maryland; Manuel Ujaldon, University of Malaga

3. Loop Re-ordering and Pre-fetching During Run-time **Suvas Vajracharya**, Dirk Grunwald, University of Colorado-Boulder

Thursday, November 20 -- 3:30-5pm

[Room B1] Tera-scale Computing
Session Chair: Carl Kesselman, Information Sciences Institute, University of Southern Californi

Abstracts

1. QCDSP: A Teraflop-scale Massively Parallel Supercomputer **Norman Christ**, Chulwoo Jung, Adrian Kahler, Steven Kasow, Yubing Luo, Robert Mawhinney, Cheng-Zhong Sui, Pavlos Vranas, Alan Gara, Columbia University; John Parsons, Robert G. Edwards, Anthony D. Kennedy, SCRI, Florida State University; Sten Hensen, Fermilab; Greg Kilcup, Ohio State University; Jim Sexton, Trinity College, Dublin

2. A System Software Architecture for High-end Computing **David S. Greenberg**, Ron Brightwell, Lee Ann Fisk, Sandia National Laboratories; Arthur McCabe, Rolf Riesen, University of New Mexico

3. Performance Characteristics of Gang Scheduling in Multiprogrammed Environments **Morris A. Jette**, Lawrence Livermore National Laboratory

[Room A4] Data Structures
Session Chair: Mary Hall, University of Southern California

Abstracts

1. High-speed Distributed Data Handling for online Instrumentation Systems **William E. Johnston**, Lawrence Berkeley National Laboratory
2. A Common Data Management Infrastructure for Adaptive Algorithms for PDE Solutions **M. Parashar**, J.C. Browne, C. Edwards, K. Klimkowski, University of Texas at Austin
3. Adaptive Blocks: A High Performance Data Structure **Q.F. Stout**; D. deZeeuw, T. Gombossi, C. Groth, H. Marshall, K. Powell, University of Michigan

[Room C2] Bell and Fernbach Awards
Session Chair: Gary Glatzmeier, Los Alamos National Laboratory

Winners of this year's Gordon Bell Prizes and Fernbach Award will present talks on their award-winning work during this session. The Gordon Bell Award was established to reward practical use of parallel processors by giving a monetary prize for the best performance improvement in an application. The prize is often given to winners in several categories relating to hardware and software advancement. Entries are coordinated by IEEE Computer and IEEE Parallel and Distributed Technology magazine, publications of the IEEE Computer Society. The Fernbach award honors Sidney Fernbach, one of the pioneers in the development and application of high performance computers for the solution of large computational problems. It is given to someone who has made "an outstanding contribution in the application of high performance computers using innovative approaches."

Gordon Bell Prize Papers and Abstracts
Fernbach Award Abstract

Tuesday, November 18 -- 10:30am-noon

[Room B1] Communication in Clusters
Session Chair: David Culler, University of California, Berkeley

1. NOMINATED STUDENT PAPER AWARD The Effects of Communication Parameters on End Performance of Shared Virtual Memory Clusters **Angelos Bilas**, Jaswinder Pal Singh, Princeton University

   With the relative maturing of shared virtual memory (SVM) protocols, this work examines the important system bottlenecks that limit effective parallel performance in SVM systems in particular, which parameters of the communication architecture are most important to improve further, which ones are already adequate on modern systems, and how will this change with technology in the future. Through detailed simulations of SMP nodes, we find that the most important system cost to improve is interrupt overhead. Improving network interface bandwidth helps a few bandwidth-bound applications. Surprisingly, neither the processor overhead for handling messages nor the occupancy of the communication interface appear to be major bottlenecks.

2. FM-QoS: Real-time Communication Using Self-synchronizing Schedules **Kay H. Connelly**, Andrew A. Chien, University of Illinois at Urbana-Champaign

   FM-QoS uses feedback-based synchronization, FBS, to provide predictable communication performance for high-speed cluster interconnects. In FBS, network feedback is combined with self-synchronizing communication schedules to achieve synchrony in the network interfaces (NIs). Based on this synchrony, the network can be scheduled to provide predictable performance without special network QoS hardware. We describe a prototype of FBS for Myricom's Myrinet network, which incurs synchronization overheads less than 1% of the total communication traffic. The prototype achieves predictable latencies of 23 ms for a single- switch, 8-node network and 2 KB packets, more than four times faster than the best-effort scheme (104 ms).

3. NOMINATED STUDENT PAPER AWARD Multi-protocol Active Messages on a Cluster of SMPs **Steven S. Lumetta**, Alan Mainwaring, David E. Culler, University of California, Berkeley

   Clusters of multiprocessors promise to be the supercomputers of the future, but obtaining high performance requires an understanding of the multiple levels of interconnection. We present the first multi-protocol implementation of a lightweight message layer Active Messages-II on Sun Enterprise 5000 servers connected with Myrinet. We describe our design and analyzes our implementation with microbenchmarks and applications. Three aspects of the layer are critical to performance: cache-coherence mechanism overhead, concurrent access management, and network state access. Using an adaptive polling strategy, we limit performance interactions between the protocols. Applications within an SMP benefit from this fast communication, but are limited between SMPs by the balance of NICs to processors.

[Room A4] Astrophysics and Hydrodynamics
Session Chair: Subodh Kumar, Johns Hopkins University

1. Highly Portable and Efficient Implementations of Parallel Adaptive N-Body Methods **David Blackston**, University of California, Berkeley; Torsten Suel, Bell Laboratories

   We describe the design of several portable and efficient parallel implementations of adaptive N-body methods, including the adaptive Fast Multipole Method, the adaptive version of Anderson's Method, and the Barnes-Hut algorithm. Our codes are based on a communication and work partitioning scheme that allows an efficient implementation of adaptive multipole methods even on high-latency systems. Our test runs demonstrate high performance and speed- up on several parallel architectures, including traditional MPPs, shared-memory machines, and networks of workstations connected by Ethernet.

2. Optimization and Scaling of Shared-memory and Message-passing Implementations of the Zeus Hydrodynamics Algorithm **Robert A. Fiedler**, Hewlett-Packard Company

   We compare the performance of shared-memory and message-passing versions of the ZEUS algorithm for astrophysical fluid dynamics on a 64-processor HP/Convex Exemplar SPP-2000. Single-processor optimization is guided by timing several versions of simple loops whose structure typifies the main performance bottlenecks. Overhead is minimized in the message- passing implementation through the use of non-blocking communication operations. Our benchmark results agree reasonably well with the predictions of a simple performance model. The message-passing version of ZEUS scales better than the shared-memory one primarily because, under shared-memory, (unless data-layout directives are utilized) the domain decomposition is effectively one-dimensional.

3. Scalable, Hydrodynamic, and Radiation-hydrodynamic Studies of Neutron Stars Mergers **Paul E. Saylor**, University of Illinois at Urbana-Champaign; Frank D. Swesty, University of Illinois at Urbana-Champaign

   We discuss the high performance computing issues involved in the numerical simulation of binary neutron star mergers and supernovae. These phenomena, which are of great interest to astronomers and physicists, can only be described by modeling the gravitational field of the objects along with the flow of matter and radiation in a self consistent manner. In turn, such models require the solution of the gravitational field equations, Eulerian hydrodynamic equations, and radiation transport equations. This necessitates the use of scalable, high performance computing assets to conduct the simulations. We discuss some of the parallel computing aspects of this challenging task.

[Room A3] Performance Tools
Session Chair: Ann Hayes, Los Alamos National Laboratory

1. Implementing a Performance Forecasting System for Metacomputing: The Network Weather Service **Rich Wolski**, Neil Spring, University of California, San Diego; Chris Peterson, Microsoft Corporation

   We describe a system called the Network Weather Service (NWS) that takes periodic measurements of deliverable resource performance from distributed networked resources, and uses numerical models to dynamically generate forecasts of future performance levels. We detail the architecture of the NWS and outline our experience with its implementation. We also include examples of NWS forecasts for resource performance in different metacomputing settings, and we discuss the tradeoff between forecast model accuracy and complexity in a dynamically changing computing environment.

2. NOMINATED STUDENT PAPER AWARD [Experiment Management Support for Performance Tuning] **Karen L. Karavanic**, Barton P. Miller, University of Wisconsin-Madison

   The development of a high performance parallel system or application is an evolutionary process -- both the code and the environment go through many changes during a program's lifetime -- and at each change, a key question for developers is: how and how much did the performance change? No existing performance tool provides the necessary functionality to answer this question. We report on the design and preliminary implementation of a tool that views each execution as a scientific experiment and provides the functionality to answer questions about a program's performance that span more than a single execution or environment.

3. [Exploiting Global Input/Output Access Pattern Classification] **Tara M. Madhyastha**, Daniel A. Reed, University of Illinois at Urbana-Champaign

   Parallel input/output systems attempt to alleviate the performance bottleneck that affects many input/output intensive applications. In such systems, an understanding of the application access pattern, especially how requests from multiple processors for different file regions are logically related, is important for optimizing file system performance. We propose a method for automatically classifying these global access patterns and using these global classifications to select and tune file system policies to improve input/output performance. We demonstrate this approach on benchmarks and scientific applications using global classification to automatically select appropriate underlying Intel PFS input/output modes and server buffering strategies.

Tuesday, November 18 -- 1:30-3pm

[Room B1] Compiling Techniques
Session Chair: David Callahan, Tera Computer Co.

1. NOMINATED STUDENT PAPER AWARD [Compiling Parallel Code for Sparse Matrix Applications] **Vladimir Kotlyar**, Keshav Pingali, Paul Vinson Stodghill, Cornell University

   We have developed a framework based on relational algebra for compiling efficient sparse matrix code from dense DO-ANY loops and a specification of the representation of the sparse matrix. We show how this framework can be used to generate parallel code, and present experimental data that demonstrates that the code generated by our compiler achieves performance competitive with that of hand-written codes for important computational kernels.

2. [Evaluating the Performance Limitations of MPMD Communication] **Chi-Chao Chang**, Grzegorz Czajkowski, Thorsten von Eicken, Cornell University; Carl Kesselman, USC/Information Sciences Institute

   We discuss the major limitations of RPC-based communication for MPMD parallel computing on homogeneous, distributed memory multi-computers. We quantify these limitations using a direct comparison between an MPMD language (CC++) and an SPMD one (Split-C) through a series of micro-benchmarks and three representative applications. Basic RPC performance in CC++ is within a factor of two from those of Split-C and other messaging layers. CC++ applications perform within a factor of two to six from comparable Split-C versions, which represent an order of magnitude improvement over previous CC++ implementations. Results suggest that RPC can be used effectively in many high performance MPMD parallel applications.

3. Compiling Stencils in High Performance Fortran **Gerald Roth**, John Mellor-Crummey, Ken Kennedy, Rice University; R. Gregg Brickner, Los Alamos National Laboratory

   In many Fortran90 and HPF programs performing dense matrix computations, the main computational work is performed by a class of kernels known as stencils. We present a general- purpose compiler optimization strategy that generates efficient code for a wide class of stencil computations expressed using Fortran90 array constructs. This strategy optimizes both single and multi-statement stencils by orchestrating a set of program transformations that minimize both intraprocessor and interprocessor data movement implied by Fortran90 array operations. Our experimental results show that our approach can produce highly optimized code in situations where other compilers do not. In some cases the improvements are as much as several orders of magnitude.

[Room A4] Computational Fluid Dynamics
Session Chair: Horst Simon, Lawrence Berkeley Laboratory

1. MPP Solution of Rayleigh-Benard-Marangoni Flow **Graham F. Carey**, Christopher Harle, Robert Mclay, Spencer Swift, The University of Texas at Austin

   A domain decomposition strategy and parallel gradient-type iterative solution scheme have been developed and implemented for computation of complex 3-D viscous flow problems involving heat transfer and surface tension effects. Details of these implementation issues are described together with associated performance and scalability studies. Representative Rayleigh-Benard and microgravity Marangoni flow calculations on the Cray T3D are presented, and performance results verifying a sustained rate in excess of 16 Gflops on 512 nodes of the T3D have been obtained. Preliminary performance studies have recently been carried out on the T3E and sustained rates above 50 Gflops and 100 Gflops have been achieved on the 512 node T3E-600 and 1024 node T3E-900 configurations respectively.

2. A Multi-level Parallelization Concept for High-fidelity Multi-block Solvers **Ferhat F. Hatay**, MCAT, Inc. (presently with HaL Computer Systems, Inc.); Dennis C. Jespersen, Guru P. Guruswamy, Yehia M. Rizk, NASA Ames Research Center; Chansup Byun, Ken Gee, MCAT, Inc.

   A hybrid domain-decomposition and parallelization concept was developed and implemented into the widely used NASA multi-block Computational Fluid Dynamics solvers employed in ENSAERO and OVERFLOW advanced flow analysis packages. The new parallel solver concept, PENS (Parallel Euler Navier-Stokes Solver), employs both fine and coarse granularity with data partitioning as well as data coalescing to obtain the desired load-balance characteristics on different computer architectures from shared-memory to distributed-memory platforms. Fine-grain partitioning of single-block domains achieves 85% scalable parallel performance. Multi-block simulations of complete aircraft geometries achieve near perfect load-balanced executions using data coalescing and the two levels of parallelism.

3. On Parallel Implementations of Dynamic Overset Grid Methods **Andrew M. Wissink**, MCAT, Inc.,

NASA Ames Research Center; Robert L. Meakin, NASA Ames Research Center

We explore the parallel performance of structured overset CFD computations for multi- component bodies in which there is relative motion between component parts. The two processes that dominate the cost of such problems are the flow solution on each component and the inter- grid connectivity solution. A two-part static-dynamic load balancing scheme is proposed and tested for unsteady calculations on the IBM SP and IBM SP2. Parallel implementation of a new solution-adaption scheme based on structured Cartesian overset grids is also presented.

Tuesday, November 18 -- 3:30-5pm

[Room B1] High Performance Fortran
Session Chair: David Bailey, NASA AMES Research Center

1. New Life in Dusty Decks: Results of Porting a CM Fortran-based Aeroacoustic Model to High Performance Fortran **Jeffrey J. Nucciarone**, Penn State University; Yusuf Ozyoruk, Lyle N. Long, Penn State University

   The High Performance Fortran (HPF) language was developed in part based on experiences with early parallel Fortran compilers such as CM Fortran. We tested the effectiveness of HPF with today's generation of HPF compilers by porting a complex hybrid computational aeroacoustics code originally developed using CMF. We will discuss how different port strategies affected the code performance. We will present how the ported code ran an IBM SP and an SGI Origin 2000. While the current simple port does not match the speed of a CM5, we hope further porting efforts and improved compiler technology will enable us to eventually surpass CM5 performance levels.

2. An Evaluation of HPF Compilers and the Implementation of a Parallel Linear Equation Solver Using HPF and MPI **Konming Gary Li**, Nabil M. Zamel, Saudi Aramco

   In this work, we evaluated the capabilities and performances of two commercially available HPF compilers, xlhpf from IBM and pghpf from the Portland Group. In particular, we examined the suitability of the two compilers for the development of a reservoir simulator. Because of the nature of reservoir simulation, multiple data distributions and data transfer between arrays of different data layouts are of great importance. An HPF compiler that does not provide these capabilities is unsuitable for the development of a parallel reservoir simulator. A detailed comparison of the functionalities of the two compilers and their suitabilities for reservoir simulator development are presented. To test the performance of the compilers, we used several reservoir simulator kernels and a parallel linear equation solver. The solver is based on preconditioned Orthomin and a truncated Neumann series preconditioner. It was first implemented in HPF and later in MPI using hpf_local to improve the performance. Its computational performance was compared for several MPP machines: CM5, IBM SP2, Cray T-3E and Cray Origin.

3. Portable Performance of Data Parallel Languages **Ton A. Ngo**, IBM T.J. Watson Research Center; Lawrence Snyder, Bradford Chamberlain, University of Washington

   A portable program yields consistent performance on different platforms. We study the portable performance of three NAS benchmarks compiled with three commercial HPF compilers on the IBM SP2. Each benchmark is evaluated using DO loops and F90 constructs. Base-line comparison is provided by Fortran/MPI and ZPL. The HPF results show some scalable performance but indicate a considerable portability problem. First, relying on the compiler alone for extensive analysis and optimization leads to unpredictable performance. Second, differences in the parallelization strategies often require compiler specific customization. The results suggest that the foremost criteria for portability is a concise performance model.

[Room A4] Visualization Algorithms
Session Chair: Tom DeFanti, EVL, University of Illinois at Chicago

1. Divide and Conquer Spot Noise **Wim de Leeuw**, Center for Mathematics and Computer Science, Kruislaan, Amsterdam, The Netherlands

   The design and implementation of an interactive spot noise algorithm is presented. Spot noise is a technique that uses texture for the visualization of flow fields. Various design tradeoffs are discussed that allow an optimal implementation on a range of high-end graphical workstations. Two applications are given: the steering of a smog prediction simulation and browsing a very large data set resulting from a direct numerical simulation of turbulence. These applications provide the motivation for the need of interactive visualization techniques.

2. Wavelet-based Image Registration on Parallel Computers **A. El-Ghazaw**, Prachya Chalermwa, The George Washington University; Jacqueline LeMoigne, NASA

   Digital image registration is very important in many applications such as medical imagery, robotics, and remote sensing. Image registration determines the relative orientation between two or more images. NASA's Mission To Planet Earth (MTPE) program will soon produce enormous Earth data, reaching hundreds of Gbytes per day. Analysis of such data requires accurate and fast registration. We describe a fast algorithm for parallel image registration. Performance of the algorithm is analyzed and experimentally evaluated. It will be shown that the algorithm provides substantial computational savings and the measurements reveal many important characteristics of contemporary high performance computers.

3. Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments **Jason Leigh**, EVL, University of Illinois at Chicago

   CAVERN, the CAVE Research Network, is an alliance of industrial and research institutions equipped with CAVE-based virtual reality hardware and high performance computing resources, interconnected by high-speed networks, to support collaboration in design, education, engineering, and scientific visualization. CAVERNsoft is the collaborative software backbone for CAVERN. CAVERNsoft uses distributed data stores and multiple networking interfaces to provide persistence, customizable latency, data consistency, and scalability that are typically needed to support collaborative virtual reality.

Wednesday, November 19 -- 10:30am-noon

[Room B1] Techniques for Metacomputing
Session Chair: Francine Berman, University of California, San Diego

1. Multi-client LAN/WAN Performance Analysis of Ninf: A High Performance Global Computing System

Atsuko Takefusa, Umpei Nagashima, Ochanomizu University; Satoshi Matsuoka, Tokyo Institute of Technology; Hirotaka Ogawa, University of Tokyo; Hidemoto Nakada, Satoshi Sekiguchi, Electrotechnical Laboratory; Hiromitsu Takagi, Nagoya Institute of Technology; Mitsuhisa Sato, Real World Computing Partnership;

Rapid increase in speed and availability of network of supercomputers is making high performance global computing possible, in which computational and data resources in the network are collectively employed to solve large-scale problems. There have been several recent proposals of global computing including our Ninf system. However, critical issues regarding system performance characteristics in global computing have been little investigated, especially under multi-clients, multi-sites WAN settings. In order to investigate the feasibility of Ninf and similar systems, we conducted benchmarks with different communication/computation characteristics on a variety of combinations of clients and servers in their performance, architecture, etc. under LAN, single-site WAN, multi-site WAN situations.

2. PARDIS: CORBA-based Architecture for Application-level Parallel Distributed Computation
   Katarzyna Keahey, Dennis Gannon, Indiana University

   We describe the architecture and programming abstractions of PARDIS, a system based on ideas underlying the Common Object Request Broker Architecture (CORBA). PARDIS provides a distributed environment in which objects representing data-parallel computation, as well as non-parallel objects present in parallel programs, can interact across platforms and software systems. Each of these objects represents a small, encapsulated application that can be used as a building block in the construction of distributed metaapplications. We will present examples of building such metaapplications with PARDIS, and show their performance in distributed systems combining the computational power of different multi-processor architectures.

3. Scalable Networked Information Processing Environment (SNIPE) **Keith Moore**, Graham E. Fagg, Al Geist, University of Tennessee; Jack Dongarra, University of Tennessee/ Oak Ridge National Laboratory

   SNIPE is a metacomputing system that aims to provide a reliable, secure, fault-tolerant environment for long-term distributed computing applications and data stores across the global InterNet. This system combines global naming and replication of both processing and data to support large scale information processing applications leading to better availability and reliability than currently available with typical cluster computing and/or distributed computer environments. To facilitate this the system supports: distributed data collection, distributed computation, distributed control and resource management, distributed output and process migration. The underlying system supports multiple communication paths, media and routing methods to aid performance and robustness across both local and global networks.

[Room A4] Climate
Session Chair: Dennis Gannon, CRPC/Indiana University

1. Optimization of a Parallel Ocean General Circulation Model **Ping Wang**, Daniel S. Katz, Yi Chao, Jet Propulsion Laboratory

   Optimization of a Parallel Ocean General Circulation Model Ping Wang, Daniel S. Katz, Yi Chao, Jet Propulsion Laboratory

We describe our efforts to optimize a parallel ocean general circulation model. We have developed several general strategies to optimize the ocean general circulation model on the Cray T3D. These strategies include memory optimization, effective use of arithmetic pipelines, and usage of optimized libraries. Nearly linear scaling performance data is obtained for the optimized code, while the speed-up data for the optimized code also shows excellent improvement over the original code. More importantly, the single-node performance is greatly improved. The optimized code runs about 2.5 times faster than the original code, which corresponds to 3.63 Gflops on the 256-PE Cray T3D. Such a model improvement allows one to perform ocean modeling at increasingly higher spatial resolutions for climate studies.

2. Parallel Computing at NASA Data Assimilation Office **M P. Lyster**, K. Ekers, J. Guo, M. Harber, D. Lamich, J.W. Larson, R. Lucchesi, R. Rood, S. Schubert, W. Sawyer, M. Sienkiewicz, A. da Silva, J. Stobie, L.L. Takacs, R. Todling, J. Zero, NASA/Goddard Space Flight Center; C.H.Q. Ding, Lawrence Berkeley National Laboratory; R. Ferraro, Jet Propulsion Laboratory

   The goal of atmospheric data assimilation is to produce accurate gridded datasets of fields by assimilating a range of observations along with physically consistent model forecasts. The NASA Data Assimilation Office (DAO) is currently upgrading its end-to-end data assimilation system (GEOS DAS) to support NASA's Mission To Planet Earth (MTPE) Enterprise. This effort is also part of a NASA HPCC Earth and Space Sciences (ESS) Grand Challenge PI project. Future Core computing, using a modular Fortran 90 design and distributed memory (MPI) software, will be carried out at Ames Research Center. The algorithmic and performance issues involved in the Core system are the main subjects of this presentation.

3. FOAM: An Atmosphere-ocean Climate Model for Studying Long-duration Phenomena **Michael Tobis**, Ian T. Foster, Chad M. Schafer, Argonne National Laboratory; Robert M. Jacob, John R. Anderson, University of Wisconsin-Madison

   We report here on a project that expands the applicability of dynamic climate modeling to very long time scales. The Fast Ocean Atmosphere Model (FOAM) is a coupled ocean- atmosphere model that incorporates physics of interest in understanding decade to century time scale variability. It addresses the high computational cost of this endeavor with a combination of improved ocean model formulation, low atmosphere resolution, and efficient coupling. It also uses message-passing parallel processing techniques, allowing for the use of cost-effective distributed memory platforms. The resulting model runs over 6000 times faster than real time with good fidelity and has yielded significant results.

[Room A3] Numerics
Session Chair: David Kahaner, ATIP and NIST

1. Parallel Threshold-based ILU Factorization **George Karypis**, Vipin Kumar, Computer Science, University of Minnesota

   Factorization algorithms based on threshold incomplete LU factorization have been found to be quite effective in preconditioning iterative system solvers. However, because these factorizations allow the fill elements to be created dynamically, they have been considered to be unsuitable for distributed-memory parallel computers. We present a highly parallel formulation of the ILUT(m, t) threshold-based incomplete factorization algorithm. ILUT employs a dual dropping strategy that is able to control the computational requirements during the factorization as well as during the application of the preconditioner. Our parallel ILUT algorithm utilizes parallel multilevel **k**-way partitioning and parallel independent set computation algorithms to effectively parallelize both the factorization as well as the application of the preconditioner.

2. PLAPACK: Parallel Linear Algebra Libraries Design Overview **Philip Alpatov**, Gregory Baker, H. Carter Edwards, John Gunnels, Greg Morrow, James Overfelt, Robert van de Geijn, The University of Texas at Austin

   The Parallel Linear Algebra Package (PLAPACK) is a maturing fourth generation linear algebra infrastructure which uses a application-centric view of vector and matrix distribution, Physically Based Matrix Distribution. It also uses an ``MPI-like" programming interface that hides distribution and indexing details in opaque objects, provides a natural layering in the library, and provides a straight-forward application interface. We give an overview of the design of PLAPACK.

3. MultiMATLAB: Integrating Matlab with High Performance Parallel Computing **Vijay Menon**, Anne E. Trefethen, Cornell University

   MultiMATLAB is an extension of the popular MATLAB environment to distributed memory multiprocessors. We present a MultiMATLAB architecture that provides performance on multiprocessors while maintaining the functionality and usability of MATLAB. This system will enable users to access high performance parallel routines from within MATLAB, to extend MATLAB with new parallel routines, and to use these routines to develop parallel applications with the MATLAB language. We discuss a general MultiMATLAB architecture and present two implementations built upon MPI. Preliminary results indicate that the MultiMATLAB system can offer the full performance of the underlying multiprocessor to the MATLAB environment.

Wednesday, November 19 -- 1:30-3pm

[Room B1] Big I/O and Checkpointing
Session Chair: Karim Harzallah, Tandem Computers

1. Optimization and Evaluation of Hartree-Fock Applications' I/O with PASSION **Meenakshi A. Kandaswamy**, Mahmut T. Kandemir, Alok N. Choudhary, Syracuse University/ Northwestern University; David E. Bernholdt, Syracuse University

   Many applications tend to perform high-volume data storage, and data retrieval, which demand high performance I/O. We conduct an experimental study of the I/O performed by the Hartree-Fock (HF) method. We use a parallel I/O library to optimize the I/O performance of the application and present extensive experimental results. We rank the optimizations based on their impact on the performance of HF as: I. efficient interface to the file system, II. prefetching, and III. buffering. The results show that within the limits of our experimental parameters, application-related factors are more effective on the overall I/O behavior of this application. We obtained up to 43% improvement in the overall application performance.

2. A Checkpointing Strategy for Scalable Recovery on Distributed Parallel Systems **Vijay K. Naik**, Samuel P. Midkiff, Jose E. Moreira, IBM T. J. Watson Research Center

   We describe a checkpoint/recovery scheme suitable for message-passing parallel applications. The novelty of our scheme is that checkpointed applications can be restored, from their checkpointed state, in reconfigured forms. Using this scheme, applications can quickly recover from partial system failures. A key component of our implementation is the distribution- independent representation of application array data structures in persistent storage. To further optimize the performance, we provide parallel array section streaming operations for distributed arrays. We compare the performance of the reconfigurable checkpoint/restart of parallel applications with that of conventional forms of checkpointing.

3. CLIP: A Checkpointing Tool for Message-passing Parallel Programs **Yuqun Chen**, Kai Li, Princeton University; James S. Plank, University of Tennessee

   Checkpointing is a useful technique for rollback recovery. We present CLIP, a user-level library that provides semi-transparent checkpointing for parallel programs on the Intel Paragon multicomputer. Creating an actual tool for checkpointing a complex machine like the Paragon is not easy, because many issues arise that require careful design decisions to be made. We detail what these decisions are, and how they were made in CLIP. We present performance data when checkpointing several long-running parallel applications. These results show that a convenient, general-purpose checkpointing tool like CLIP can provide fault-tolerance on a massively parallel multicomputer with good performance.

[Room A4] Molecular Dynamics
Session Chair: Roy Richter, General Motors Research and Development Center

1. Massively Parallel Simulations of Diffusion in Dense Polymeric Structures **Jean-Loup M. Faulon**, Sandia National Laboratories

   An original technique to generate close-to-equilibrium dense polymeric structures is proposed. Diffusion of small gases is studied on the equilibrated structures using massively parallel molecular dynamics simulations running on the Intel Teraflops (9200 Pentium Pro processors) and Intel Paragon (1840 processors). Compared to the current state-of-the-art equilibration methods the new technique is faster by some orders of magnitude. The main advantage of the technique is that one can circumvent the bottlenecks in configuration space that inhibit relaxation in molecular dynamics simulations. In agreement with experimental evidence, small gas diffusion coefficients are larger for ethylene-propylene-dimer-monomer (EPDM) than butyl rubber.

2. Molecular Dynamics Simulation of Large-scale Carbon Nanotubes on a Shared-memory Architecture **Deepak Srivastava**, Stephen T. Barnard, MRJ Technology Solutions, NASA Ames Research Center

   Classical molecular dynamics simulations employing Brenner's reactive potential with long range van der Waals interactions have been used in mechanistic response studies of carbon nanotubes to external strains. Elastomechanic response behavior of single and multiwall carbon nanotubes to externally applied compressive strains is simulated and studied in detail. Due to inclusion of non-bonded long range interactions, the simulations show the redistribution of strain and strain energy from sideways buckling to the formation of highly localized strained kink sites. We describe the results and discuss their implication towards the stability of any molecular mechanical structure made of carbon nanotubes.

3. Topology Preserving Dynamic Load Balancing for Parallel Molecular Simulations **David F. Hegarty**, Tahar M. Kechadi, University College Dublin

   We develop a parallel simulation environment to mask the complexity of heterogeneous computing resources. We propose a dynamic load balancing algorithm, Positional Scan Load Balancing which preserves the original ordering of the problem. PSLB allows for dynamic processing powers and is suitable for embedding in a runtime system. A performance model is described and cost prediction formulae derived. We introduce crossover points at which it becomes beneficial to use the algorithm. We examine the algorithms performance on a 256 processor Cray T3D and a shared memory multiprocessor. We present results detailing overhead, performance gain, and speedup of the algorithm.

Wednesday, November 19 -- 3:30-5pm

[Room B1] Real Iron
Session Chair: Forest Baskett, Silicon Graphics, Inc.

1. The Starfire SMP Interconnect **Alan Charlesworth**, Sun Microsystems

   The Starfire Ultra Enterprise 10000 extends the envelope of UNIX SMP systems in several dimensions. Interconnect: It uses four address routers and a 16x16 data crossbar to provide 64 UltraSPARC processors with uniform-memory access at a bandwidth of 10,667 MBps. Flexibility: Starfire can be dynamically reconfigured into multiple hardware-protected operating system domains. Robustness: specially-connected workstation orchestrates service activities. Failing boards can be hot-swapped without interrupting system operation. ECC is carried on both address and data paths. Performance: Starfire has set several TPC-D decision- support benchmark records. It delivers a respectable 21 Gflops on Linpack-parallel.

2. Tera Hardware-Software Cooperation **Gail Alverson**, Preston Briggs, Susan Coatney, Simon Kahan, Rich Korry, Tera Computer Company

   The development of Tera's MTA system was unusual. It respected the need for fast hardware and large shared memory, facilitating execution of demanding parallel application programs. At the same time, it provided a clean machine model enabling calculated compiler optimizations and easy programming as well as novel architectural features necessary to support fast parallel system software. From its inception, system and application needs have molded the MTA architecture. The result is a system that offers high performance and ease of programming by virtue of fast hardware, flat shared memory, and streamlined software systems.

3. Performance Analysis of the T3E Multiprocessor **Ed C. Anderson**, Jeffrey P. Brooks, Charles M. Gassi, Steven L. Scott, Cray Research

   The CRAY T3E is a scalable shared-memory multiprocessor based on the DEC Alpha 21164 microprocessor. The system includes a number of architectural features designed to tolerate latency and enhance scalability. Included among these are stream buffers, which detect and prefetch down small-stride reference streams, E-registers, which allow memory reference pipelining and provide non-unit-stride access capabilities, and a scalable, high-bandwidth interconnection network. We report our experiences with T3E performance. We describe several hardware features, discuss programming implications, and provide related benchmark results. Included are NAS Parallel Benchmark results up to 1024 processors.

[Room A4] Climate and Fusion
Session Chair: Tony Hey, University of Southampton

1. Distributed High Performance Computation for Remote Sensing **Kenneth A. Hawick**, Heath A. James, University of Adelaide

   We describe distributed and parallel algorithms for processing remotely sensed data such as geostationary satellite imagery. We have built a distributed data repository based around the client-server computing model across wide-area ATM networks, with embedded parallel and high performance processing modules. We focus on algorithms for classification, geo- rectification, correlation and histogram analysis of the data. We consider characteristics of image data collected from the Japanese GMS5 geostationary meteorological satellite, and some analysis techniques we have applied to it. As well as providing a browsing interface to our data collection, our system provides processing and analysis services on-demand.

2. Parallel Non-linear Optimization: Toward the Design of a Decision Support System for Air Quality Management **Andrew Lewis**, Griffith University, Australia; Rod Simpson, Griffith University; David Abramson, Monash University, Australia

   We address the optimization component of a decision support system for air quality management, and describes some benchmark studies assessing its effectiveness. Because of the computationally demanding nature of the objective function we are interested in parallelizing the quasi-Newton BFGS algorithm selected for initial study. This is achieved by concurrent evaluation of functions in finite difference approximations to the derivative and a method of interval subdivision in simple bound constrained line searching. In a realistic case study, use of the parallel optimization algorithm is shown to have significant performance gains over other methods of solution.

3. Numerical Tokamak Turbulence Calculations on the CRAY T3E **Vickie E. Lynch**, Jean-Noel Leboeuf, Benjamin A. Carreras, Juan Diego Alvarez, Oak Ridge National Laboratory; Luis Garcia, Universidad Carlos III de Madrid

   Full cross section calculations of ion-temperature-gradient-driven turbulence with Landau closure are being carried out as part of the Numerical Tokamak Turbulence Project, one of the U. S. Department of Energy's Phase II Grand Challenges. Calculations of cylindrical multihelicity ion-temperature-gradient-driven turbulence with previously unattainable resolution were completed on the National Energy Research Scientific Computing Center's 160- processor distributed-memory CRAY T3E parallel computer. Both T3E benchmark results and physics results will be presented. Work is in progress to increase the resolution, improve the performance of the parallel code, and include toroidal geometry in these calculations for the new fully configured, 512-processor, T3E-900 model.

Thursday, November 20 -- 10:30am-noon

[Room B1] Memory Hierarchies
Session Chair: Daniel Reed, University of Illinois

1. Performance Evaluation of the SGI Origin2000: A Memory-centric Characterization of LANL ASCI Applications **Olaf M. Lubeck**, Yong Luo, Harvey J. Wasserman, Federico Bassetti, Los Alamos National Laboratory

   We compare single-processor performance of the SGI Origin and PowerChallenge and utilize a previously-reported performance model for hierarchical memory systems to explain the results. Both the Origin and PowerChallenge use the same microprocessor (MIPS R10000) but have significant differences in their memory subsystems. Our memory model includes the effect of overlap between CPU and memory operations and allows us to infer the individual contributions of all three improvements in the Origin's memory architecture and relate the effectiveness of each improvement to application characteristics.

2. A Study of Performance on SMPs and Distributed Memory Architectures Using a Shared-memory Programming Model **Eugene D. Brooks III**, Karen H. Warren, Lawrence Livermore National Laboratory

   We examine the use of a shared memory programming model to address the problem of portability between distributed memory and shared memory architectures. We conduct this evaluation by extending an existing programming model, the Parallel C Preprocessor, with a type qualifier interpretation of the data sharing keywords borrowed from the Split-C and AC compilers. We evaluate the performance of the resulting programming model on a wide range of shared memory and distributed memory computing platforms using several numerical algorithms as benchmarks. We find the type-qualifier-based programming model capable of efficient execution on distributed memory and shared memory architectures.

3. Measuring Memory Hierarchy Performanceof Cache-coherent Multiprocessors Using Micro Benchmarks **Cristina Hristea**, MIT/SGI; Daniel Lenoski, John Keen, SGI

   Even with today's large caches, the increasing performance gap between processors and memory systems imposes a memory bottleneck for many important scientific and com mercial applications. This bottleneck is intensified in shared-memory multiprocessors by contention and the ef fects of cache coherency. Under heavy memory contention, the memory latency may increase two or three times. Nonethless, as more sophisticated techniques are used to hide latency and increase bandwidth, measuring memory performance has become increasingly difficult. Previous simple methods to measure memory performance can overestimate unipro cessor memory latency and underestimate bandwidth by tens of percent. We introduce a micro benchmark suite that measures memory hierarchy performance in light of both uniprocessor optimizations and the contention and coherence effects of multiprocessors. The benchmark suite has been used to improve the memory system performance of the SGI Origin multiprocessor.

[Room A4] Subsystem Performance
Session Chair: Andrew Chien, University of Illinois and Hewlett-Packard

1. High Performance Software on Intel Pentium Pro Processors or Micro-Ops to TeraFLOPS **Bruce Greer**, Greg Henry, Intel Corporation

   We give a technical discussion of the Intel Pentium Pro processor and optimization strategies used to achieve high performance on scientific applications. We demonstrate these optimizations by characterizing matrix multiplication (DGEMM). We give insight and a model into our efforts on obtaining the world's first TeraFLOP MP LINPACK run (on the Intel ASCI Option Red Supercomputer), based on Pentium Pro processor technology. The importance is carried by the increasing trend of commodity parts in the supercomputing arena.

2. Parallel Simulation of Parallel File Systems and I/O Programs **Rajive Bagrodia**, Stephen Docy, Andy Kahn, UCLA

   Efficient I/O implementations can have a significant impact on the performance of parallel applications.

This paper describes the design and implementation of PIOSIM, a parallel simulation library for MPI-IO programs. The simulator can be used to predict the performance of existing MPI-IO programs as a function of architectural characteristics, caching algorithms, and alternative implementations of collective I/O operations. We describe the simulator and presents the results of a number of performance studies to evaluate the impact of the preceding factors on a set of MPI-IO benchmarks, including programs from the NAS benchmark suite.

3. NOMINATED STUDENT PAPER AWARD A Scalable Mark-sweep Garbage Collectoron Large-scale Shared-memory Machines
   Toshio Endo, Kenjiro Taura, Akinori Yonezawa, The University of Tokyo

   This work describes implementation of a mark-sweep garbage collector (GC) for shared- memory machines and reports its performance. It is a simple 'parallel' collector in which all processors cooperatively traverse objects in the global shared heap. The collector stops the application program during a collection. To achieve scalability, collector performs dynamic load balancing, which exchanges objects to be scanned between processors. However, we observed that the implementation detail affects the performance heavily. For example, large objects, which become a source of significant load imbalance are split into small pieces. With all careful implementation, we achieved 28-fold speed-up on 64 processors.

[Room A3] Working Parallelism
Session Chair: Jack Dongarra, University of Tennessee and ORNL

1. Parallel Database Processing on a 100 Node PC Cluster: Cases for Decision Support Query Processing and Data Mining **Takayuki Tamura**, Masato Oguchi, Masaru Kitsuregawa, The University of Tokyo

   We developed a large scale PC cluster system consisting of one hundred Pentium Pro PCs interconnected by an ATM switch, and examined its performance on data warehouse processing. First, we picked up the most complex query of the standard benchmark, TPC-D, on a 100 GB database. Our PC cluster exhibited much higher performance compared with those in current benchmark reports. Second, we developed a parallel association rule mining algorithm and ran it on the PC cluster. Sufficiently high linearity was obtained. Thus we believe such commodity based PC clusters will play a very important role in large scale database processing.

2. Page Replacement Using Marginal Loss Functions **Shamik Das Sharma**, Joel Saltz, University of Maryland; Manuel Ujaldon, University of Malaga

   We describe a compiler-directed technique to reduce page-faults in multiprocessing systems. Compile-time analysis of access-patterns is coupled with runtime support to characterize access-patterns in the form of marginal-loss functions - these functions describe the extra page faults that would be incurred for an access-pattern if it were given one fewer physical page. The kernel uses these functions to guide its page-replacement decisions by victimizing those processes whose access-patterns are affected the least. We outline how marginal loss functions can be computed for common access-patterns and presents simulation results to demonstrate the technique's effectiveness.

3. Loop Re-ordering and Pre-fetching During Run-time **Suvas Vajracharya**, Dirk Grunwald, University of Colorado-Boulder

   The order in which loop iterations are executed can have a large impact on the number of cache misses that an applications takes. A new loop order that preserves the semantics of the old order but has a better cache data re-use, improves the performance of that application. Several compiler techniques exist to transform loops such that the order of iterations reduces cache misses. We introduce a run-time method to determine the order based on a dependence-driven execution. In a dependence-driven execution, an execution traverses the iteration space by following the dependence arcs between the iterations.

Thursday, November 20 -- 3:30-5pm

[Room B1] Tera-scale Computing
Session Chair: Carl Kesselman, Information Sciences Institute, University of Southern California

1. QCDSP: A Teraflop-scale Massively Parallel Supercomputer **Norman Christ**, Chulwoo Jung, Adrian Kahler, Steven Kasow, Yubing Luo, Robert Mawhinney, Cheng-Zhong Sui, Pavlos Vranas, Alan Gara, Columbia University; John Parsons, Robert G. Edwards, Anthony D. Kennedy, SCRI, Florida State University; Sten Hensen, Fermilab; Greg Kilcup, Ohio State University; Jim Sexton, Trinity College, Dublin

   We discuss the work of the QCDSP collaboration to build an inexpensive Teraflop scale massively parallel computer suitable for computations in Quantum Chromodynamics (QCD). The computer is a collection of nodes connected in a four dimensional toroidial grid with nearest neighbor bit serial communications. An 8192-node computer with a peak speed of 0.4 Teraflops is being constructed at Columbia University for a cost of $1.8 Million. A 12,288-node machine with a peak speed of 0.6 Teraflops is being constructed for the RIKEN Brookhaven Research Center. Other computers have been built including a 50 Gigaflop version for Florida State University.

2. A System Software Architecture for High-end Computing **David S. Greenberg**, Ron Brightwell, Lee Ann Fisk, Sandia National Laboratories; Arthur McCabe, Rolf Riesen, University of New Mexico

   MPP systems can neither solve Grand Challenge scientific problems nor enable large-scale industrial and governmental simulations if they rely on extensions to workstation system software. We present a new system architecture used at Sandia. Highest performance is achieved through a lightweight applications interface to a collection of processing nodes. Usability is provided by creating node partitions specialized for user access, networking, and I/O. The system is glued together by a data movement interface called portals. Portals allow data to flow between processing nodes with minimal system overhead while maintaining a suitable degree of protection and reconfigurability.

3. Performance Characteristics of Gang Scheduling in Multiprogrammed Environments **Morris A. Jette**, Lawrence Livermore National Laboratory

   Gang scheduling provides both space-slicing and time-slicing of computer resources for parallel programs. Each thread of execution from a parallel job is concurrently scheduled on an independent processor in order to achieve an optimal level of program performance. Time- slicing of parallel jobs provides for better overall system responsiveness and utilization than otherwise possible. Lawrence Livermore National Laboratory has deployed three generations of its gang scheduler on a variety of computing platforms. Results indicate the potential benefits of this technology to parallel processing are no less significant than time-sharing was in the 1960s.

[Room A4] Data Structures
Session Chair: Mary Hall, University of Southern California

1. High-speed Distributed Data Handling for online Instrumentation Systems **William E. Johnston**, Lawrence Berkeley National Laboratory

   The advent (and promise) of shared, widely available, high-speed networks provides the potential for new approaches to the collection, organization, storage, and analysis of high- speed and high-volume data streams from high data-rate, on-line instruments. We have worked in this area for several years, have identified and addressed a variety of problems associated with this scenario, and have evolved an architecture, implementations, and a monitoring methodology that have been successful in addressing several different application areas. We describe a distributed, wide area network-based architecture that deals with data streams that originate from online instruments. Such instruments and imaging systems are a staple of modern scientific, health care, and intelligence environments. Our work provides an approach for reliable, distributed real-time analysis, cataloguing, and archiving of the data streams through the integration and distributed management of a high-speed distributed cache, distributed high performance applications, and tertiary storage systems.

2. A Common Data Management Infrastructure for Adaptive Algorithms for PDE Solutions **M. Parashar**, J.C. Browne, C. Edwards, K. Klimkowski, University of Texas at Austin

   We present the design, development, and application of a computational infrastructure to support the implementation of parallel adaptive algorithms for the solution of sets of partial differential equations. The infrastructure is separated into multiple layers of abstraction. This work is primarily concerned with the two lowest layers of this infrastructure: a layer which defines and implements dynamic distributed arrays (DDA), and a layer in which several dynamic data and programming abstractions are implemented in terms of the DDAs. The currently implemented abstractions are those needed for formulation of hierarchical adaptive finite difference methods, hp-adaptive finite element methods, and fast multipole method for solution of linear systems. Implementation of sample applications based on each of these methods are described and implementation issues and performance measurements are presented.

3. Adaptive Blocks: A High Performance Data Structure **Q.F. Stout**; D. deZeeuw, T. Gombossi, C. Groth, H. Marshall, K. Powell, University of Michigan

   We examine a data structure that uses flexible "adaptivity" to obtain high performance for both serial and parallel computers. The data structure is an adaptive grid which partitions a given region into regular cells. Its closest relatives are cell-based tree decompositions, but there are several important differences which lead to significant performance advantages. Using this block data structure to support adaptive mesh refinement (AMR), we were able to sustain 17 GFLOPS in ideal magnetohydrodynamic (MHD) simulations of the solar wind emanating from the base of the solar corona, using a 512 processor Cray T3D at NASA Goddard.

[Room C2] Bell and Fernbach Awards
Session Chair: Gary Glatzmeier, Los Alamos National Laboratory

1. Transient Solid Dynamics Simulations on the Sandia/Intel Teraflop Computer
   Stephen A. Attaway Sandia National Laboratories; Edward J. Barragy, Intel Corporation; Kevin H. Brown, David R. Gardner, Bruce A. Hendrickson, Steven J. Plimpton, Courtenay T. Vaughan, Sandia National Laboratories

We describe our parallelization of PRONTO, Sandia's transient solid dynamics code, via a novel algorithmic approach that utilizes multiple decompositions for different key segments of the computations, including the material contact calculation. This latter calculation is notoriously difficult to perform well in parallel, because it involves dynamically changing geometry, global searches for elements in contact, and unstructured communications among the compute nodes. Our approach scales to at least 3600 compute nodes on problems involving millions of finite elements. We can simulate models using more than ten million elements in a few tenths of a second per timestep.

2. High Performance MP Unstructured Finite Element Simulation of Chemically Reacting Flows
Karen D. Devine, Gary L. Hennigan*, Scott A. Hutchinson, Andrew G. Salinger, John N. Shadid, and Ray S. Tuminaro Sandia National Laboratories, Albuquerque, NM *New Mexico State University, Las Cruces, NM

   We describe the performance of MPSalsa, a MP code that simulates complex systems with strongly coupled fluid flow, thermal energy transfer, mass transfer and non-equilibrium chemical reactions. MPSalsa uses 3D unstructured finite element methods, fully implicit time integration, and general gas-phase and surface-species chemical kinetics to solve the coupled nonlinear PDEs on complex domains. It is designed around general kernels for domain partitioning, unstructured message passing, distributed sparse-block matrix representation of the fully summed global finite element equations, and preconditioned Krylov iterative solvers. Using these techniques, we obtained sustained rates of 210+ Gflop/s for a 3-D chemically reacting flow problem.

3. Large Scale Simulation of Suspensions with PVM
Nhan Phan-Thien, Ka Yan Lee, University of Sydney; David Tullock, Los Alamos National Laboratory

   We describe an application of an indirect boundary element method in the large scale simulation of suspensions in distributed environment with Parallel Virtual Machine (PVM). The method involves re-casting the governing (Stokes or Navier) equations in an integral form which is suitable for iterative solution. Parallelization is achieved by domain decomposition, in which the calculations on a group of particles are farmed out to a workstation in a distributed system. The actual throughput of the distributed farm, consisting of 28 DEC AlphaStations 500/266, based on a wall clock measurement and an actual floating point operation count for the serial version of the program, is 4.69 Gflops/s. This results in a price performance of 65 USD per Mflops/s.

4. Pentium Pro Inside: I. A Treecode at 430 Gigaflops on ASCI Red, II. Price/Performance of $50/Mflop on Loki and Hyglac
Michael S. Warren, Los Alamos National Laboratory; John K. Salmon, California Institute of Technology; Donald J. Becker, Goddard Space Flight Center; M. Patrick Goda, Los Alamos National Laboratory; Thomas Sterling, Caltech/JPL; Grégoire S. Winckelmans, Universite Catholique de Louvain

   We present results from two methods of solving the gravitational N-body problem on ASCI Red. The first method, a trivial O(N^2) algorithm, obtained 635 Gflops for a 1 million particle problem on 6800 Pentium Pro processors. The second method, a treecode which scales as O(N log N), sustained 170 Gflops over a continuous 9.4 hour period on 4096 processors and 430 Gflops on 6800 processors during the initial part of the simulation. We also present two simulations which sustained roughly one Gigaflop on each of two 16 processor Beowulf-class computers constructed entirely from commodity personal computer technology for $50k each in September, 1996.

KEYWORD INDEX

The Role of Techology in Lifelong Learning

The Education Program gives K-12 teachers, administrators, university professors, educational software researchers, community technology leaders, and others a three-day intensive experience with the tools and techniques of high performance computing and high-speed networking as applied in the education domain. Invited speakers, panels, papers, workshops, hands-on labs, and posters are all part of this robust program designed for anyone interested in technology and life-long learning.

The complete program may be viewed below or you will find a summary in the Program at a Glance.. To view abstracts for the various papers, panels, workshops, invited speakers and posters please use the buttons above.

Education Program Reception

SATURDAY, NOVEMBER 15
6-9PM

Room B1, San Jose McEnery Convention Center

Education Program participants can meet one another, along with the members of the SC97 executive committee, at this opening reception for the SC97 educators. Don't miss it!

Invited Speaker

SUNDAY, NOVEMBER 16
8:30-10AM
Room J2, San Jose McEnery Convention Center

Scaffolding Learning & Addressing Diversity: Technology as the Trojan Mouse

Elliot Soloway, University of Michigan

SUNDAY, NOVEMBER 16
10:30AM-NOON

K-12 Teacher Workshops and Hands-on Labs

Chickscope: An Innovative World Wide Laboratory for K-12 Classrooms
Clinton Potter, Umesh Thakkar, NCSA, University of Illinois at Urbana-Champaign
Almaden I Room, San Jose Hilton and Towers

MATLAB in the Science Classroom
Charlotte M. Trout, Maryland Virtual High School; Don Shaffer, Maryland Virtual High School
Almaden II Room, San Jose Hilton and Towers

SUNDAY, NOVEMBER 16
10:30AM-NOON

Education Papers

San Carlos I & II Rooms, San Jose Hilton and Towers

Transforming Teaching for the 21st Century
Lisa Bievenue, NCSA; Donna Cauley, Andalusia High School, Andalusia, AL; Edna Gentry, ASPIRE

Internet in the K-12 Classroom: The Realities of Technology Transfer
Kim Nguyen-Jahiel, Punyashloke Mishra, NCSA

Webwisdom: Architecture of Web-based Training and Education System
Geoffrey Fox, NPAC at Syracuse University

SUNDAY, NOVEMBER 16
NOON-1:30PM

Luncheon Address

Room J2, San Jose McEnery Convention Center

SUNDAY, NOVEMBER 16
1:30-3PM

K-12 Teacher Workshops and Hands-on Labs

Project CERES: Bringing NASA Resources into the Classroom
Stephanie Stevenson, George Tuthill, Dave Thomas, Timothy Slater, Montana State University
Almaden I Room, San Jose Hilton and Towers

GirlTECH
Cynthia Lanius, Milby High School, Houston, TX; Susan Boone, Saint Agnes Academy, Houston, TX
Almaden II Room, San Jose Hilton and Towers

SUNDAY, NOVEMBER 16
1:30-3PM

Education Papers

San Carlos I & II Rooms, San Jose Hilton and Towers

Maryland Virtual High School of Science and Mathematics: Lessons Learned and Future Directions for a Computational Science Internet Project
Mary Ellen Verona, Susan Ragan, Maryland Virtual High School

STEP: A Case Study on Building a Bridge from HPC Technologies to the Secondary Classroom
Kris Stewart, San Diego State University and SDSC; Janet Bowers, San Diego State University

The Network Montana Project: Summary of Materials Development Activities
David A. Thomas, Montana State University

SUNDAY, NOVEMBER 16
3:30-5PM

K-12 Teacher Workshops and Hands-on Labs

Designing Web Pages that are Universally Accessible
Beth Ann Ziebarth, University of Illinois Cooperative Extension Service; Edna Gentry, ASPIRE; Donna Cauley, Andalusia High School, Andalusia, AL
Almaden I Room, San Jose Hilton and Towers

Fractal Analysis From Patterns in Nature: Experimental and Computer Analysis
Elizabeth (Tommi) Holsenbeck, Jefferson Davis High School, Montgomery, AL; Jane Jones, J. O. Johnson High School, Huntsville, AL
Almaden II Room, San Jose Hilton and Towers

SUNDAY, NOVEMBER 16
3:30-5PM

Education Panel

Collaborative Environments and Technology
Beau Fly Jones, Ohio Supercomputer Center; Geoffrey Fox, NPAC, Syracuse University; David Emigh, NCSA
San Carlos I & II Room, San Jose Hilton and Towers

SUNDAY, NOVEMBER 16
7-10PM

Open Labs

Almaden I & II Room, San Jose Hilton and Towers

MONDAY, NOVEMBER 17
8:30-10AM

K-12 Teacher Workshops and Hands-on Labs

GirlTECH
Cynthia Lanius, Milby High School, Houston, TX; Susan Boone, Saint Agnes Academy, Houston, TX
see above

Designing Web Pages that are Universally Accessible
Beth Ann Ziebarth, University of Illinois Cooperative Extension Service; Edna Gentry, ASPIRE; Donna Cauley, Andalusia High School, Andalusia, AL
see above

MONDAY, NOVEMBER 17
8:30-10AM

Education Panel

Standards, Frameworks and Technology Planning for Schools
Ralph Annina, Senior Consultant, Allied Computer Group; Carol Doherty, Director, Professional Development, Center for Innovation in Urban Education
Santa Clara I and II Rooms, San Jose Hilton and Towers

MONDAY, NOVEMBER 17
10:30-NOON

K-12 Teacher Workshops and Hands-on Labs

Chickscope: An Innovative World Wide Laboratory for K-12 Classrooms
Clinton Potter, Umesh Thakkar, NCSA
see above

Project CERES: Bringing NASA Resources into the Classroom
Stephanie Stevenson, George Tuthill, Dave Thomas, Timothy Slater, Montana State University
see above

MONDAY, NOVEMBER 17
10:30-NOON

Education Panel

The Future Impact of the National Partnerships in Computational Infrastructure on Education
Roscoe C. Giles, Boston University; John Ziebarth, NCSA
San Carlos I and II Rooms, San Jose Hilton and Towers

MONDAY, NOVEMBER 17
1:30-3PM

K-12 Teacher Workshops and Hands-on Labs

MATLAB in the Science Classroom
Charlotte M. Trout, Maryland Virtual High School; Don Shaffer, Maryland Virtual High School
Almaden I Room, San Jose Hilton and Towers
see above

Fractal Analysis From Patterns in Nature: Experimental and Computer Analysis
Elizabeth (Tommi) Holsenbeck, Jefferson Davis High School, Montgomery, AL; Jane Jones, J. O. Johnson High School, Huntsville, AL
Almaden II Room, San Jose Hilton and Towers
see above.

MONDAY, NOVEMBER 17
1:30-3PM

Education Papers

San Carlos I and II Rooms, San Jose Hilton and Towers

Instructional Design Issues Surrounding Development of online Training in VRML
Melissa L. Kelly, NCSA

Exploring the Link Between "Real" and "Virtual" Experiences in the Classroom with VRML and 3-D Modeling
E. K. Schroeder, Illinois State Museum, Illinois; T. M. Bulka, Northern High School, Maryland; D. A. Emigh, Quinebaug Valley Community Technical College, Connecticut; B. A. Andersen, Peterson Elementary School, Montana; D. A. Chapman, National Center for Supercomputing Applications, Illinois; R. L. Fixen, Montana State University, Montana; J. A. Leggett, Wilson Middle School, Illinois; U. Thakkar, National Center for Supercomputing Applications, Illinois; R. S. Toomey, Illinois State Museum, Illinois

Toward Defining a SuperWeb
Beau Fly Jones, Ohio Supercomputer Center, Ohio State University

MONDAY, NOVEMBER 17
3:30-5PM

K-12 Teacher Workshops and Hands-on Labs

Authoring Problems by Students for Students: An Advanced Physics High School Project Laboratory
M. E. Hinton, Urbana High School, Urbana, IL; A. J. Fleming, Champaign Centennial High School, Champaign, IL; A. W. Hubler, University of Illinois at Urbana-Champaign; U. Thakkar, NCSA
Almaden II Room, San Jose Hilton and Towers

MONDAY, NOVEMBER 17
3:30-5PM

Education Panel

Development and Assessment of Virtual Reality Materials for Education
Chris Dede, George Mason University; Robert Fixen, Montana State University; Melissa Kelly, NCSA
Almaden I Room, San Jose Hilton and Tower

MONDAY, NOVEMBER 17
3:30-5PM

Education Posters POSTERS (Click for abstracts)

Almaden I Room, San Jose Hilton and Tower

SC97 technical program posters are also open to all Education Program participants.

P1 **Tessellation Tutorials**
Suzanne Alejandre, Math Forum

P2 **The Internet World's Fair by a K-6 School**
Deanna Alexander, Fisher Grade School, Fisher, IL; Melissa Kelly, NCSA

P3 **Online K-12 Earth System Science Lessons that Use Web Resources**
Brian Beaudrie, Timothy Slater, Dave Thomas, Montana State University

P4 **Use of the CAVE in Promoting Safe Street Crossing**
Patricia Brown, Columbia Elementary School, Champaign, IL; Georgette Moore, Yankee Ridge Elementary School, Urbana, IL

P5 **Cyber Mummy VR**
Renee Cooper, Pamela Van Walleghen, Urbana Middle School, Urbana, IL

P6 **Basic Aerodynamics Software Lab**
John D. Eigenauer, NASA Lewis Research Center

P7 **Virtual Reality Technology Enhances Learning for Interdisciplinary Studies on the Missouri River**
Carol Engelman, Ralston High School, Omaha, NE; Elaine Westbrook, Omaha North High School, Omaha, NE

P8 **Fusion of Art and Technology**
Kathy Felling, Karen Ciprari-Murphy, Fulton County Board of Education, Atlanta, GA

P9 **Implementing the National Science Education K-12 Standards Using NASA Web Resources**
Robert Fixen, Timothy Slater, George Tuthill, Stephanie Stevenson, Montana State University

P10 **Entry-level Virtual Reality**
Jane Leggett, Wilson Middle School, Moline, IL

P11 **NASA Lewis Learning Technologies Program**
Beth Lewandowski, NASA Lewis Research Center

P12 **The National Education Supercomputer Program**
Brian Lindow, Lawrence Livermore National Laboratory

P13 **Multiple Representations of the Periodic System of Elements: the Design and Evaluation of a Multimedia Hypertext**
Punyashloke Mishra, NCSA

P14 **Virtual Reality at the Jersey Shore**
Alan Sills, West Essex Regional High School

P15 **Yohkoh Public Outreach Project: Using the Web and HPCC for Lifelong Learning**
Timothy Slater, Montana State University; Marion French, Kansas State University; Donna Governor, Escambia County School District, Pensacola, FL

P16 **Earth System Science High School Classroom Applications Using Scientific Visualization and Distance Learning Experiences**
Robert Smith, Michael Porter, Terry Parker High School, Jacksonville, FL

MONDAY, NOVEMBER 17
7-10PM

Open Labs

Almaden I & II Room, San Jose Hilton and Towers

- TUESDAY, NOVEMBER 18
- 8:30-10AM

Keynote Address

San Jose Civic Auditorium

Vice President Al Gore, who has a long-standing interest in the value and use of high performance networking and computing, particularly in education, has been invited to present the keynote address to the conference.

- TUESDAY, NOVEMBER 18
- 10:30AM-NOON

Invited Speaker

Special Address: Orchestrating Mathematical and Computational Sciences Instruction to the Needs of Science and Society

Richard Tapia, Computational and Applied Mathematics and Center for Research on Parallel Computation, Rice University
Room C1, San Jose McEnery Convention Center

- TUESDAY, NOVEMBER 18
- 1:30-3PM

Education Panel

Presentation of the UCES Awards
Jim Corones, AmesLab
Santa Clara I and II Rooms, San Jose Hilton and Towers

Each year, the Department of Energy has sponsored an awards program for excellence in undergraduate computational engineering and science. During this panel each of the 1997 UCES awardees will make a short presentation on his/her work and receive the UCES award. See http://uces.ameslab.gov/uces/awards for more information.

- TUESDAY, NOVEMBER 18
- 1:30-3PM

Education Papers

Santa Clara I and II Rooms, San Jose Hilton and Towers

Computational Physics: A Course, a Book, and Web Tutorials

Rubin Landau, Physics, Oregon State University

MPI on Line: A Teaching Environment for MPI
Elspeth Minty, Edinburgh Parallel Computing Centre

The Portals Project: Supporting Web-based Collaborations among Students and Scientists
Kallen Tsikalas, Katherine McMillan, Center for Children & Technology, EDC; Peggy Larisch, Silver High School, Silver City, NM; Bruce Bennett, North Polk High School, Alleman, IA

- TUESDAY, NOVEMBER 18
- 3:30-5PM

K-12 Teacher Workshops and Hands-on Labs

Authoring Problems by Students for Students: An Advanced Physics High School Project Laboratory
M. E. Hinton, Urbana High School, Urbana, IL; A. J. Fleming, Champaign Centennial High School, Champaign, IL; A. W. Hubler, University of Illinois at Urbana-Champaign; U. Thakkar, NCSA
Almaden I Room, San Jose Hilton and Towers

- TUESDAY, NOVEMBER 18
- 3:30-5PM

Education Panel

Internet in the K-12 Classroom: The Realities of Technology Transfer
Panelists TBA
San Carlos I and II Rooms, San Jose Hilton and Towers

- TUESDAY, NOVEMBER 18
- 3:30-5PM

Education Posters

Almaden II Room, San Jose Hilton and Towers

See listing Monday, November 17, 3:30-5pm Education Posters.

- TUESDAY, NOVEMBER 18
- 7-10PM

Open Labs

Almaden I & II Room, San Jose Hilton and Towers

- WEDNESDAY, NOVEMBER 19
- 8:30AM-5PM

Local Teachers/Administrators Day

On Wednesday, local area teachers will participate in the Education Program. Two special workshops, a luncheon, and a tour of the exhibit floor will be part of the day's events.

WEDNESDAY, NOVEMBER 19
8:30-10AM

Workshop

Strategic Planning For Integrating Technology into Schools
Carol Doherty, Director, Professional Development Program for Educators and Schools, Center for Innovation in Urban Education, Northeastern University, Boston, MA; Louis Kruger, Professor, Counseling, Psychology and Rehabilitation Northeastern University; Gregory Tutunjian, Independent Software Consultant
San Carlos I and II, San Jose Hilton and Towers

WEDNESDAY, NOVEMBER 19
10:30AM-NOON

Workshop

Strategic Planning: An approach for K-12 Education Organizations
Ralph Annina, Senior Consultant, Allied Computer Group
San Carlos I and II, San Jose Hilton and Towers

SC97 Exhibition Open to SC97 Education Program Participants

Don't miss all that SC97 has to offer! Visit the exhibition, which includes a booth on the SC97 Education Program (Booth R210).

The hours for the SC97 exhibition, which includes research, poster, and industry exhibits, are:

Expanded Exhibits Hours

Monday	November 17 7-9pm	(Gala Opening--All Exhibits)
Tuesday	November 18 10am-6pm	(All Exhibits)
Wednesday	November 19 10am-6:30pm	(includes 5-6:30pm posters reception)
Thursday	November 20 10am-4pm	(All Exhibits)

SC97 Education Program Research Booth-R210

The SC97 Education Program is designed to give K-12 teachers, administrators, university professors, educational software researchers and developers, community technology leaders, and others three days of intense exposure to the tools and techniques of high performance computing and high-speed networking as they can be applied to education. Drop by this exhibit to learn about the SC97 and SC98 Education Programs, and to find out how to become involved in outreach and education programs nationwide.

Technical Program Posters

SC97 Education Program attendees are encouraged to also view the SC97 technical program posters (Click here for abstracts), which are in the Almaden Concourse of the San Jose McEnery Convention Center.

Technical Program Posters Hours

Monday	November 17	7-9pm	(Gala Opening--All Exhibits)
Tuesday	November 18	10am-6pm	(All Exhibits)
Wednesday	November 19	10am-6:30pm	(includes 5-6:30pm posters reception)
Thursday	November 20	10am-4pm	(All Exhibits)

SAN JOSE SC97 NOV. 15-21

Research Exhibits

- The Aggregate
- Ames Laboratory
- Argonne National Laboratory
- Army High Performance Computing Research Center
- ASCI, LLNL, LANL, and SNL
- Boston University
- Brookhaven National Laboratory
- California Institute of Technology
- Center for Research on Parallel Computation
- DoD HPC Modernization Program
- DoE 2000
- Electrotechnical Laboratory
- Emory University
- Esprit HPF+ Project
- Fermi National Accelerator Laboratory
- High Performance Computing Center, Stuttgart
- Indiana University
- Institut National de Recherche en Informatique et en Automatique (INRIA)
- Krell Institute
- Lawrence Berkeley National Laboratory
- Massachusetts Institute of Technology
- MHPCC/HPCERC/AHPCC
- Mississippi State University
- National Aeronautics and Space Administration
- National Aerospace Laboratory of Japan
- National Center for Atmospheric Research
- National Computational Science Alliance
- National Coordination Office for Computing, Information, and Communications
- National Partnership for Advanced Computational Infrastructure
- National Scalable Cluster Project (NSCP)
- NOAA/Forecast Systems Laboratory
- North Carolina State University
- Northwest Alliance for Computational Science and Engineering
- Oak Ridge National Laboratory
- Ohio Supercomputer Center
- Pacific Northwest National Laboratory
- Parallel Tools Consortium
- Pittsburgh Supercomputing Center
- Real World Computing Partnership
- Saitama University
- SC97 Education Program
- SC97 Research Exhibits
- Scalable I/O Initiative
- University of Alaska Fairbanks
- University of California, Santa Barbara
- University of Sao Paulo--Integrated Systems Laboratory
- University of Utah
- University of Virginia
- University of Wisconsin

Abstracts

The Aggregate
Booth R110

Over the past four years, a new model for parallel processing communication has emerged. Neither message-passing nor shared memory, the aggregate function model is based on performing N-way communication functions as single operations. Although this model began at Purdue University as the library for PAPERS (Purdue's Adapter for Parallel Execution and Rapid Synchronization) clusters, this model has become the highly-portable Aggregate Function API (AFAPI), and has spread to a number of other institutions. This exhibit will present AFAPI-related research from several universities, including new hardware and software for clustering of COTS PCs. Multiple parallel systems using AFAPI, including at least one video wall, will be demonstrated.

Ames Laboratory
Booth R101

The cost-effectiveness of commodity components for computing and networking hardware and software has made major strides in the last few years. The level of computing resources that may now be dedicated to departmental level projects is unprecedented, rivaling the capabilities of national centers just a few years ago. Ames Laboratory has considerable experience locating and removing a variety of performance bottlenecks from clusters using ATM and Fast Ethernet. This year's research exhibit will highlight the analysis and performance of a cluster of high performance commodity workstations using commodity software and connected using Gigabit Ethernet technology. A Packet Engines full-duplex repeater will provide the key element of a new low-cost gigabit interconnect. NetPIPE, the network performance application tool developed at Ames Laboratory, will demonstrate the effectiveness of this interconnect as will a selection of Materials Science Grand Challenge applications.

Argonne National Laboratory
Booth R111

Argonne National Laboratory continues to provide portable and scalable software for large-scale computational applications, for the effective use and management of geographically distributed supercomputing resources, for immersive visualization, and for advancing the concept of "national collaboratories." Showcased will be current work in areas including:

- The design and integration of collaborative tools and shared virtual environments. Numerical tools and libraries for large-scale computational applications - ADIC, ADIFOR, PETSc, SUMMA3d. The Optimization Technology Center and the NEOS Server, providing access to the latest optimization techniques over the Internet.
- Romio: a portable implementation of MPI-IO.
- The GLOBUS Project: to develop the software infrastructure for management of the national computational grid.

- Associated Grand Challenge applications. Closely tied with these projects will be an emphasis on collaborations with the ASCI Alliance program, CRPC, the DOE2000 program, the NCSA PACI Alliance, and the Scalable I/O Project.

Army High Performance Computing Research Center
Booth R303

The Army High Performance Computing Research Center (AHPCRC), located at the University of Minnesota in Minneapolis, Minnesota, was established by the US Army. Our mandate is to establish university-Army collaborative research programs to advance the science of high performance computing and its application to critical Army technology issues and to maintain leadership in computing technology vital to national security. The research exhibit will demonstrate the results of AHPCRC research projects that use the TMC CM-5 and the Cray T3D for large simulations. Highlighted will be basic research projects, collaborative research projects involving both AHPCRC and Army researchers, and educational programs. The AHPCRC research exhibit will utilize several methods of information presentation; high performance workstations for research project data visualization demos, a VCR and large-screen monitor for showing narrated videotapes of AHPCRC research project results, and poster-size pictures of selected research project graphics mounted in the exhibit.

ASCI, LLNL, LANL, and SNL
Booth R104

In this coordinated exhibit, three Department of Energy laboratories present demos, posters, and videos related to work at the individual laboratories, as well as the shared DoE Accelerated Strategic Computing Initiative (ASCI) program. Sandia is forcing a revolution in engineering by combining Teraflops-scale simulation with world-class experimental facilities to predict material responses and processes more complex than ever before possible in programs from microelectronics and manufacturing to critical infrastructure surety. LANL will feature simulations of the global climate modeling Grand Challenge application and wildfire predictability. The ongoing software efforts of the POOMA framework and TeleMed collaboration will be demonstrated, as well as the prototype HIPPI-6400 tester. LLNL features information about the laboratory, current tera-scale computational upgrades, and presentations from computational and networking research. Planned demos include Incident Advisory, Python for scientific computing, gang scheduling/parallel tools, and remote visualization over high-speed networks.

Boston University
Booth R105

Boston University's research exhibit features its NSF-funded project MARINER: Metacenter Affiliated Resource In the New England Region. MARINER extends the University's efforts in advanced scientific computing and networking to organizations throughout the region. Users from both public and private sector are eligible to participate in a wide range of programs that offer training in and access to advanced computing and communications technologies. Demonstrations of current research and educational projects developed through the Center for Computational Science and the Scientific Computing and Visualization Group will be shown using graphics workstations and video in the exhibit booth. In addition, we will present distributed supercomputing applications, video animations of recent research, and 3-D visualizations using a stereoscopic display. We will also be announcing our new role as a partner institution in the National Computational Science Alliance under NSF's PACI program.

Brookhaven National Laboratory
Booth R207

A Stereographic Collaborative Visualization Facility and its Applications: with the unmistakable trend of scientific collaborations becoming larger and more geographically dispersed, new emphasis is being placed on the technologies that enable meaningful distance interaction. High-quality visualization facilities are of great importance in the conduct of such collaborations. Even for local staff, visualization is a key to enhanced comprehension and insight into the results of experiments. Brookhaven National Laboratory provides an example of such a facility along with its current and emerging scientific applications. This research exhibit will provide a sampling of the stereographic viewing experience with applications from a variety of scientific disciplines. Among the applications are those in medical imaging, protein chemistry, and the geology of oil exploration and environmental remediation. http://www.ccd.bnl.gov/visualization/examples.html

California Institute of Technology
Booth R310

The California Institute of Technology (Caltech) will highlight ongoing investigations by researchers at Caltech's Center for Advanced Computing Research (CACR) and the Jet Propulsion Laboratory (JPL). Among these projects are Petaflops computing, Beowulf/Clusters of PCs, SF/Express, various Grand Challenge Applications, the Caltech/Hewlett-Packard collaboration, and LARGE‹the Los Angeles Regional Gigabit Environment. CACR's activities in two national-scale, multidisciplinary collaborations will also be featured: the National Partnership for Advanced Computational Infrastructure (NPACI) funded by the NSF and a computational facility for simulating the dynamic response of materials to shock waves, an ASCI Center of Excellence funded by the DoE Accelerated Strategic Computing Initiative (ASCI). Graphics workstations linked to the 256 processor HP Exemplar and an IBM HPSS at Caltech will be used to demonstrate Grand Challenge Applications. A Beowulf cluster in the research exhibit will demonstrate applications and be used to illustrate the construction of Beowulf systems.

Center for Research on Parallel Computation
Booth R214

The Center for Research on Parallel Computation (CRPC), an NSF STC dedicated to making massively parallel computing truly usable, includes researchers at two national laboratories (Argonne and Los Alamos) and five universities (Caltech, Rice, Syracuse, Tennessee, and Texas). The CRPC also has affiliated sites at Boston, Drexel, Illinois, Indiana, and Maryland Universities, University of Houston, and the Institute for Computer Applications in Science and Engineering (ICASE). The CRPC exhibit features technologies the center is transferring and the mechanisms used to transfer them. Technologies demonstrated will include interactive software systems, descriptions of parallel language extensions, and applications developed by CRPC researchers. These include HPF, Fortran D, the D System, CC++, PVM, HeNCE, parallel templates, ADIFOR, and applications, including HPCC technologies in education. Transfer mechanisms include DoD Modernization, PACI, NHSE, HACSC, and CRPC educational programs. Demos, posters, videos, and information about outreach activities, software distribution, technical reports, and knowledge transfer efforts will be available.

DoD HPC Modernization Program
Booth R302

The DoD High Performance Computing Modernization Program (HPCMP) is a multi-year, $1.2B initiative to modernize HPC and advanced networking capabilities for the DoD's research programs. The research exhibit will consist of a research and technology demonstration area and an adjacent area housing the exhibit's computational resources. The research exhibit will be used to highlight various components and initiatives of the HPCMP. The booth will include four major elements: interactive demonstrations that use ATM OC-3 interconnects with computational resources located at remote DoD sites, video reports of large-scale numerical modeling efforts, a technology demonstration using an interactive flight simulator, and a collaboration area containing contributions to DoD's HPC efforts produced by the program's academic and industrial partners.

DoE 2000
Booth R318

DOE2000 is a Department of Energy initiative to develop advanced computing and collaboration technologies for research and development. There are three components: 1) Advanced Computational Testing and Simulation (ACTS): Developing an integrated scientific software toolkit. 2) National Collaboratories: Laboratories without walls that connect researchers, instruments, and computers nationwide. 3) Pilot Projects: Early implementations of virtual laboratories. Demonstrations will include:

- The Materials Micro-Characterization Collaboratory, which links five experimental facilities on-line.
- The Diesel Combustion Collaboratory, using computational models to design the next-generation engine.
- CUMULVS, an application for multiple-site interactive simulations that allows users to modify the parameters of running applications.
- POOMA, an applications framework that makes high performance computing both accessible and portable.
- Real-time remote control of experiments at the Advanced Light Source.
- Virtual reality for user collaborations using a scalable media server.

Electrotechnical Laboratory
Booth R212

This research exhibit presents some results from a number of related research projects in the Electrotechnical Laboratory in the areas of high performance computing and communications. We will present video and live demonstrations of programs on the following research results:

- **EM-X**: We demonstrate a highly parallel computer that can tolerate communication latency by using low latency communication and multithreading. Some real-time parallel benchmark results are demonstrated.
- **Ninf**: Ninf is an ongoing global network-wide computing infrastructure project which allows users to access computational resources including hardware, software and scientific data distributed across a wide area network.
- **Web Information Retrieval (IR)**: This IR system uses speculative execution between processors to achieve high performance. In addition, we will present some results of recent software developments at Electrotechnical Laboratories such as Delegate and Horb, which are of interest to the supercomputing community.

Emory University
Booth R215

CCF (Collaborative Computing Frameworks) is a suite of software systems and tools, communications protocols, and methodologies that enable collaborative, computer-based cooperative work. CCF constructs a virtual work environment on multiple computer systems connected over the Internet, to form a collaboratory. Participants may interact, share applications and data repositories or archives, collectively create and manipulate documents and spreadsheets, perform computational transforms, and conduct a number of other activities via telepresence. CCF is an integrated framework that consists of multiple coordinated infrastructural elements, each of which provides a component of the virtual collaborative environment. A prototype of a complete collaboration system will be exhibited with live demonstrations, and will include the underlying reliable multicast protocol suite, application sharing and X-multiplexer systems, shared dataspace and filesystem, the computing harness, and the clearboard, audiotool, multiway chat, and video conferencing tools. http://ccf.mathcs.emory.edu/

Esprit HPF+ Project
Booth R307

The Esprit IV project "HPF+: Optimizing HPF for Advanced Applications" aims to improve the HPF language and related compilation technology extending the functionalities of HPF and developing compilation strategies based on the requirements of a set of advanced applications. The purpose of the exhibit is to demonstrate the results achieved in HPF+ with focus on benchmark development in HPF+ based on real scientific applications, the Vienna Fortran Compiler (VFC) and an evaluation of the effectiveness of HPF+ and its implementation. The compiler and runtime technology developed in the VFC includes the general block distributions, INDIRECT distributions, ON clauses and schedule reuse as well as the Inspector-Executor parallelization strategy. The evaluation of HPF+, by means of the performance tool MEDEA, provides a comparison of the developed technology with HPF-2 and message passing approach.http://www.par.univie.ac.at/hpf+

Fermi National Accelerator Laboratory

Booth R216

High Energy Physics (HEP) requires massive computational resources for data acquisition; design, control, and analysis of experiments; and, for theoretical analyses. To improve our understanding of fundamental physics, the CPU requirement will increase by more than an order of magnitude in the next two years. The increasing power of commodity computing hardware and network technology allows scientists to use clusters of relatively inexpensive computers for HEP applications. Fermi National Accelerator Laboratory will demonstrate two prototype clusters under consideration for production systems, both based on Intel microprocessors, running the Linux operating system. One receives, analyzes, and filters data at 100 MBytes/sec from the CDF Collider Experiment (http://www-cdf.fnal.gov/) using ATM and Fast Ethernet technologies. The second, a prototype PC FARM (http://www-ols.fnal.gov/pcfarms/) analyzes HEP data using (http://fnhppc.fnal.gov/farms/farmspage.html) loosely-coupled parallel processing techniques developed at FNAL. Both systems execute FNAL codes ported to Linux, and both demonstrate the scalability of such commodity-based computing for HEP tasks. http://www.fnal.gov/

High Performance Computing Center, Stuttgart
Booth R304

The High Performance Computing Center Stuttgart (HLRS, a division of RUS, the computing center of Stuttgart University) is a National HPC Center for Germany. Embedded in industrial cooperations, it focuses on engineering applications. We will showcase our work under the G7 GIBN initiative, done with the Pittsburgh Supercomputing Center (PSC) and Sandia National Laboratory (SNL). We have set up a metacomputing environment that links up 512 processor Cray T3E's at HLRS and PSC and a Paragon at SNL, via a transatlantic ATM link, vBNS and ESnet. This is complemented by collaborative visualization in a distributed virtual environment, and applied to: fluid dynamics simulations of the re-entry of a space vehicle; the entry of meteorites into the Earth's atmosphere as well as their impact on the Earth; and ab-initio calculations of magnetic properties of alloys. Results will be visualized in a distributed virtual environment between HLRS, PSC, and SNL.

Indiana University
Booth R315

Indiana University has a vigorous program in HPC research and instruction. Three projects of note will be presented in this exhibit:

- **SCAAMP**: This project enables research advances in computational astronomy, chemistry, geology, mathematics, physics, and theater design by providing an infrastructure that allows application scientists to move their computations from PCs and workstations up to geographically distributed arrays of supercomputers without massive investments in programming time at the upper end. The hardware component of the computing environment consists of an array of shared-memory multiprocessors, high performance 3-D visualization hardware, and fast networks to connect the components. This system will be an element in a national array of supercomputers, coupled to solve the most challenging problems currently considered in simulations. http://www.cica.indiana.edu/scaamp
- **Extreme**: The Extreme Computing group is working on a number of projects including HPC++, a specification and compilers for portable parallel programming using the HPC++ programming language; the Sage compiler preprocessor toolkit; the TAU tools (Tuning and Analysis Utilities); high performance

parallel Java; and Problem Solving Environments, a multi-institutional, multidisciplinary research project focused on symbolic computation, user interfaces and collaborative technologies for parallel object-oriented programming. The PSEware project is application driven, the initial set of driving applications includes numerical tokamak, cosmology, environmental models, and signal processing. http://www.extreme.indiana.edu

- **VR and HPN**: Initiatives begun recently at IU in domestic and international high performance networking and in virtual environments will be described and key projects in these initiatives will be discussed. http://www.cica.indiana.edu/vbns, http://www.cica.indiana.edu

Institut National de Recherche en Informatique et en Automatique (INRIA)
Booth R305

This exhibit covers the following topics of INRIA's research in high performance computing:

- Software prototypes for high performance architectures. We present TSF, an extension of SIMULOG's Foresys system that integrates tools for assisting the process of porting applications to high performance architectures, the PACHA run-time system, COBRA, based on CORBA and SCI for networks of workstations, and C++// and Java//, two libraries for parallel, distributed, and concurrent object-oriented programming.

- Compiler technology. We demonstrate Salto, a retargetable system for assembly language transformations and optimizations such as assembly-code schedulers (e.g., software pipelining), as well as profiling and tracing tools. Salto is available for Sparc, Alpha, TriMedia, and Texas TMSC620. Interactive environments for object-oriented languages Eiffel, Eiffel//, and Java are also presented.

http://www.inria.fr/welcome-eng.html
http://www.irisa.fr/caps/PROJECTS/Salto/
http://www.irisa.fr/caps/PROJECTS/TSF/
http://www.inria.fr/sloop/c++ll
http://www.inria.fr/sloop/javall
http://www.inria.fr/sloop/prosit
http://www.inria.fr/croap/eiffel-ll

Krell Institute
Booth R312

The research at the Krell Institute focuses on applying modern computing, communication, and information technologies to educational, environmental, and energy-related priorities. The SC97 research booth will feature examples of employing information technologies to enhance computational science curriculum that spans middle school through undergraduate studies. The K-12 Adventures in Supercomputing (AiS) program was designed to introduce computational science in the context of a project-based, learner-centered curriculum that integrates mathematics and science with technology. The Undergraduate Computational Engineering and Science Project (UCES) was designed to encourage and support the infusion of computational science into the undergraduate curriculum. The UCES project is made up of three components: a web-based archive of education materials, authoring tools, and a professional recognition program. The 1997 UCES awards will be presented at SC97. These activities are supported by the Department of Energy.

http://www.krellinst.org

Lawrence Berkeley National Laboratory
Booth R200

When Ernest Orlando Lawrence founded America's first national laboratory in Berkeley, he redefined the way scientific research is conducted by fostering multidisciplinary research teams. Today the Computing Sciences organization at Lawrence Berkeley National Laboratory is redefining the way scientists will work in the next century by developing new tools for computation and collaboration. Demonstrations will include:

- Interactive visualizations using the National Energy Research Scientific Computing Center's (NERSC's) Cray T3E-900/512, which has the largest storage configuration ever assembled by SGI/CRI.
- Remote control of scientific instruments via the Energy Sciences Network (ESnet).
- Progress on several of the DoE's Grand Challenge projects.
- Geographically distributed systems for management and storage of massive data objects.
- Collaborations with the University of California, Berkeley, Computer Science Division, including the world's most powerful non-commercial network of workstations, a prototype system that may have a major impact on large-scale computing for years to come.

Massachusetts Institute of Technology
Booth R102

Project Bayanihan studies the idea of volunteer computing, which enables people to join a large parallel computation by simply visiting a Web site. Because volunteering requires no prior human contact and very little technical knowledge, it becomes very easy to build very large computing networks. This creates exciting new possibilities. With true volunteer systems, one can reach new heights in performance by using many thousands of anonymous volunteer nodes around the world. On a smaller but more practical scale, companies or institutions can use forced volunteer systems to pool together their internal computing resources with minimal administration costs. Paid volunteer systems are also possible, wherein volunteers are somehow compensated for their participation, or allowed to engage in the barter trade of processing cycles. Currently, we have developed a flexible software framework using Java and HORB, and are using it to study issues such as adaptive parallelism, security, and user-interface design. http://www.cag.lcs.mit.edu/bayanihan/

MHPCC/HPCERC/AHPCC
Booth R306

The High Performance Computing Education and Research Center (HPCERC), a strategic research center at the University of New Mexico, initiates and coordinates education and research programs in high performance computing. A major program of HPCERC is the Maui High Performance Computing Center (MHPCC), a leader in scalable computing and application technology. MHPCC supports the transition of projects from initial concept through production for Department of Defense, government, commercial, and academic users. Another key program of the HPCERC is the Albuquerque High Performance Computing Center (AHPCC), which provides an high performance computing environment for basic research and education at the University of New Mexico. This booth will highlight key projects of MHPCC and AHPCC. This includes image enhancement research to support the Air Force's electro-optical telescopes and the development of parallel models to support disaster planning. There will also be a demonstration of the Maui Scheduler, an advanced systems software tool developed to address requirements of MHPCC's large IBM SP system.

Mississippi State University
Booth R107

The ERC is a multi-disciplinary center that puts high performance computers to work in a number of ways. This booth illustrates a cross-section of these activities, and emphasizes practical uses of existing and emerging computational and visualization technology. In all, software techniques, algorithms, and applications are the emphasis, not the raw hardware, which can do little without the added value of the technology to be displayed.

National Aeronautics and Space Administration
Booth R319

The National Aeronautics and Space Administration's SC97 research exhibit highlights some of the most significant high performance networking and computing research projects underway at five field installations: Ames Research Center, Goddard Space Flight Center, Jet Propulsion Laboratory, Langley Research Center, and Lewis Research Center. Coordinated under the booth theme "Next Generation Supercomputing," the exhibit features real-time and interactive demonstrations of advanced networking applications; advanced computing architectures, including distributed heterogeneous systems; advanced operating systems and software tools for parallel systems; and advanced applications, including multidisciplinary aerodynamic design and optimization, global climate modeling, and planetary and cosmological phenomena simulation. A stereo theater graphically depicts various elements of NASA's next-generation supercomputing research, and shows other segments such as high definition rendering of the Martian landscape. A virtual reality workbench demonstration plus videos and static displays of other research round out the exhibit.

National Aerospace Laboratory of Japan
Booth R100

National Aerospace Laboratory(NAL) of Japan is a national research institute devoted to aerospace technology. NAL has been one of the most advanced HPC users in Japan. In 1993, NAL developed a special-purpose high performance computer: Numerical Wind Tunnel (NWT) jointly with Fujitsu. NWT's main target is to serve as a flow solver of Computational Fluid Dynamics (CFD) applications for national aerospace development projects. It is one of the top high performance computers, consisting of parallelly connected 166 vector processors each of which has performance of 1.7 GFLOPS and 256 MB memory. The distinctive architecture is most fit to efficiently process multiple vector do loops that are found in Navier-Stokes solvers. Thus sustained performance of 111 GFLOPS was realized in a 3-D compressor blade analysis code which was awarded the Gordon Bell Prize in 1996. Recently joint researchers with universities and industries and international collaborations are increasing in number. Easy and high speed access to NWT from outside is requested. In the exhibit, Remote NWT Access Utility, which runs on a Web client, will be demonstrated. NWT system and recent CFD accomplishment will be also on display.

National Center for Atmospheric Research
Booth R219

The National Center for Atmospheric Research (NCAR) exhibit will highlight investigations that use a broad range of supercomputing technologies to gain insight into real-world phenomena. These insights often lead to a better understanding of the phenomenon itself, and may provide the means to develop predictive capabilities based on fine tuning of the models executed on NCAR's supercomputers. Putting supercomputers to work on problems that may result in immediate benefit for humanity is a major objective of the scientific and research staff at NCAR. The computers used are some of the most sophisticated and fastest machines available. Model output, when combined with high-end visualization capabilities, frequently yield stunning images that further assist in the interpretation and analysis of the results of the model's execution. For 1997, NCAR's research exhibit will focus on several research projects that required long-running simulations on NCAR's supercomputers. In addition, this year will mark the debut of the Scientific Computing Division's Visualization Theater to display in 3-D format the visualized output of these (and other) models. The exhibit will feature four supercomputing components: NCAR's Climate Simulation Laboratory's (CSL) simulation of a 125-year build up of carbon dioxide in the atmosphere, a revised and terrain-enhanced visualization of a coupled atmosphere-fire model, real-time execution of NCAR's CCM2 (or MM5) climate model(s) on the HP 2000 Exemplar machine, and turbulence over mountainous terrain.http://www.scd.ucar.edu/info/SC97/

National Computational Science Alliance
Booth R114

The National Computational Science Alliance (the Alliance) is a partnership of individuals and institutions prototyping the nation's advanced computational infrastructure for the 21st Century. The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign anchors the Alliance, which is funded by the National Science Foundation's Partnerships for Advanced Computational Infrastructure program. The Alliance is organized into four teams: Application Technologies (AT), Enabling Technologies (ET), Regional Partners, and Education, Outreach and Training (EOT). AT teams will demonstrate how the Alliance attacks large-scale problems of science and engineering and drives technology development in cosmology, chemical engineering, environmental hydrology, molecular biology, nanomaterials and scientific instrumentation. ET teams will show how tools and infrastructure are being developed to benefit science. The booth will also include information about Regional Partners, who help distribute the resources of the Alliance nationally, and EOT teams, who bring new technologies to schools, the industrial sector, and under-served populations.

National Coordination Office for Computing, Information, and Communications
Booth R106

The National Coordination Office for Computing, Information, and Communications (NCO/CIC) coordinates multi-agency CIC R&D activities in computing, information, and communications. These activities are organized into five Program Component Areas: High End Computing and Computation; Large Scale Networking; High Confidence Systems; Human Centered Systems; and Education, Training, and Human Resources. Our booth will highlight NCO activities and CIC R&D conducted by 12 participating organizations‹DARPA, NSF, DoE, NASA, NIH, NSA, NIST, ED, VA, NOAA, EPA, and AHCPR‹through exhibition of representative results from agency-sponsored projects. NCO publications, including the High Performance Computing and Communications (HPCC) Program's FY 1998 Annual Report to Congress (the Blue Book), FY 1998 CIC brochure, and FY 1998 HPCC Implementation Plan, will be distributed. The booth will also highlight the activities of the recently established Presidential Advisory Committee on HPCC, Information Technology, and the Next Generation Internet. Materials from the Committee's first two meetings will be available. Materials from activities of other organizations supported by the NCO, such as the Applications Council and the new Technology Policy Subcommittee (TPS), will also be available. (The CIC R&D Subcommittee, Applications Council, and the TPS report to the Committee on Computing, Information, and Communications (CCIC), under the National Science and Technology Council (NSTC)).

National Partnership for Advanced Computational Infrastructure
Booth R112

NPACI is a new national organization supported by the NSF's Partnerships for Advanced Computational Infrastructure initiative. NPACI builds on the foundation of the San Diego Supercomputer Center and brings together more than three dozen academic, industrial, and research institutions in 18 states to make the world's most advanced computational resources and applications available to the nation's scientists and engineers. NPACI's efforts are concentrated into nine major areas: Enabling Technologies, Data-Intensive Computing Enabling Technologies, Interaction Environments Enabling Technologies, Adaptable, Scalable Tools and Environments High Performance Computing Resources Applications, Molecular Science Applications, Neuroscience Applications, Earth Systems Science Applications, Engineering, and Education, Outreach, and Training. The booth exhibits will demonstrate new tools and applications being developed by the cooperating partners, will present NPACI resources in action through "transparent supercomputing," and will show how the partnership's applications, research projects, and services are meeting real needs of the computational science community. http://www.npaci.edu

National Scalable Cluster Project (NSCP)
Booth R301

We will demonstrate software tools and applications that have been developed by the National Scalable Cluster Project (NSCP) and the National Data Mining Laboratory (NDML) related to data mining and data-intensive computing. The NSCP is a collaboration between research groups at the University of Illinois at Chicago (UIC), the University of Pennsylvania (UPenn), and the University of Maryland at College Park (UMD), which focuses on technologies related to cluster computing. The NDML is a collaboration headquartered at UIC focusing on data mining, data-intensive computing, and related technologies. During SC 97, we will demonstrate a variety of applications involving data mining and data-intensive computing running on local- and wide-area clusters of workstations. Applications include mining scientific, engineering ,and medical data. Workstation clusters at the exhibit and at the sponsoring institutions will be connected using ATM and vBNS links.

NOAA/Forecast Systems Laboratory
Booth R103

The mission of NOAA's Forecast Systems Laboratory (FSL) is to evaluate and transition technology to the operational weather services. FSL has been evaluating the appropriateness of different computing architectures for their use in real-time numerical weather prediction. To that end, FSL/ACB has developed a software toolbox known at the Scalable Modeling System (SMS) to parallelize numerical weather prediction (NWP) models. FSL has parallelized several NWP models using SMS. The focus of our research exhibit will be the running of real-time NWP models on a number of architectures. The forecasts will be run on our SGI Origin 2000 and Intel Paragon at FSL as well machines on the SC 97 exhibit floor. The forecasts will be visualized in our booth using simple animated 2-D displays of temperature and precipitation as well as animated 3-D displays of more sophisticated phenomena.

North Carolina State University
Booth R108

The purpose of the exhibit is to present the content and technology that supports the Regional Training Center for Parallel Processing (RTCPP). RTCPP was established by the North Carolina State University (NCSU) with NSF support. RTCPP provides an advanced network-based environment for learning about parallel computing, high performance networking, and computational sciences in general. The environment has facilities for construction of customized lessons and courses based on a collection of re-usable lesson elements (objects), as well as real-time capture of the materials. RTCPP is also one of the principal NC State University participants in the development of the regional GigaPOP and Internet 2 Test-Bed. The RTCPP education and training library contains material in a variety of formats, from videotapes, to Web-based streaming media material, to MPEG-2 based material. RTCPP will showcase the latest version of its Internet-2 oriented Web Lecture System. http://renoir.csc.ncsu.edu/RTCPP
http://renoir.csc.ncsu.edu/WLS

Northwest Alliance for Computational Science and Engineering
Booth R203

Want to learn high performance computing without becoming a computer scientist? NACSE provides Web-based training materials and tools aimed at scientists, engineers, and students. Demonstrations will include:

- **Remote Access to Scientific Databases**. HyperSQL is an interoperability layer that makes it possible for non-computer-scientists to build Web query interfaces to remote databases--without having to become an expert in SQL, HTML, or database administration. We will show off the work of several scientific groups who have adopted our software to "open up" their research databases.

- **Customizable Web-to-database Interfaces**. Query Designer allows the user to develop or customize query interfaces to remote databases, using simple point-and-click operations.

- **Learning to Speak**. Machine learning techniques developed on HPC machines are used to "train" a program to pronounce English sentences more accurately. A voice synthesizer illustrates, step-by-step, how machine learning improves the program's performance.

- **Hands-on Computational Physics**. A highly interactive set of Java-based experiments makes computational physics fun as well as challenging.

Oak Ridge National Laboratory
Booth R300

The ORNL exhibit features the application of high performance computing, networking, and storage to real-world problems. Of special interest are the latest results from ongoing research in large-scale distributed computing. In collaboration with SNL and PSC, we are using PVM software and ATM hardware over ESnet to run a huge problem on geographically distributed multiple supercomputers. Along with the outstanding high performance computing environment provided by the Department of Energy (DoE) through our Center for Computational Sciences (CCS), ORNL provides leadership in Grand Challenge computing. ORNL high performance computing activities include: developing accurate mathematical models of complex phenomena, creating scalable algorithms for their solution, developing and operating high performance computing and storage environments, generating data management and analysis methods, developing software tools, and implementing communications technologies. The Computational Center for Industrial Innovation, a DoE National User Facility, provides an effective mechanism for collaborative research between industry and ORNL. http://www.ccs.ornl.gov/SC97/SC97.html

Ohio Supercomputer Center
Booth R204

The Ohio Supercomputer Center (OSC) is Ohio's technology leader. A statewide focus on networking initiatives is moving Ohio to the national forefront in communications and the Center continues to provide state-of-the-art scalable resources for computer modeling and simulations. The Center will present an interactive exhibit at SC97 highlighting diverse network research and development applications including projects from OCARNet, a seven-university ATM testbed. The Center also will showcase the work of Ohio researchers such as droplet deformation and breakup using Lattice-Boltzmann Method (LBM), which will enhance understanding of multi-phase flow dynamics; and an endoscopic sinus surgery simulation that promises to improve medical training. Best practices in Web-based education and training, arising from OSC's involvement with Ohio universities and in the DoD HPC Modernization Program and the NSF's Partnerships for Advanced Computational Infrastructure (PACI) program as part of the National Computational Science Alliance (NCSA) effort will be featured. Center activities also will be highlighted in the DoD and the NCSA booths.http://www.osc.edu

Pacific Northwest National Laboratory
Booth R317

Productive use of the massive parallel advanced computing resources available to research scientists requires not only a revolution in computational methods for efficient use on these systems, but also a corresponding revolution in software tools for managing complex computational experiments, remote communication access, analyzing, tracking, and managing large complex input and output data sets. The Environmental Molecular Sciences Laboratory (EMSL) is coming on-line at the Pacific Northwest National Laboratory (PNNL) in Richland, Washington. This new national user facility and collaboratory is focused on environmental molecular science. EMSL has a new 512 node IBM SP system along with and EMASS Data Storage system located in the Molecular Science Computing Facility (MSCF). We have been developing software tools to effectively and efficiently utilize this new resource as well as other computational resources made available by the Department of Energy.

Parallel Tools Consortium
Booth R201

Recent investigation has shown that, despite increasingly vocal demands from the HPC community for software support, parallel tool use within that community remains disappointingly low. The Parallel Tools Consortium (Ptools) brings together researchers, developers, and users from the federal, industrial, and academic sectors to improve the responsiveness of parallel tools to user needs. Ptools has assumed a leading role in the definition, development, and promotion of parallel tools that meet the specific requirements of users who develop scalable applications on a variety of platforms. Additionally, the High Performance Debugging Forum (HPDF) is a recent effort sponsored by Ptools. The HPDF goal is to define standards relevant to debugging tools for HPC systems. The HPDF effort, as well as other Ptools activities and prototypes, will be the highlight of the Ptools research exhibit at SC97. Members of the Ptools Consortium Steering Committee (which includes national laboratory technical staff, university researchers, software tool developers, and computer hardware vendors) will be available to discuss the Ptools mission and activities and to further promote the dialog between tool users and tool developers.http://www.ptools.org/

Pittsburgh Supercomputing Center
Booth R213

Booth The Pittsburgh Supercomputing Center (PSC) is an independent national supercomputing center dedicated to providing academic, industrial, and government researchers with access to state-of-the-art high performance computing and communication resources. In addition, we provide our researchers with a flexible environment that is conducive to solving today's largest and most challenging computational science problems. Our research exhibit will demonstrate the capabilities of our resources, which include the latest massively parallel systems (Cray T3E/LC512 and T3D/MC512) as well as traditional parallel vector platforms (Cray C916 and J90s). We will feature a wide variety of demonstrations designed to showcase award-winning PSC research; particular areas of focus include weather modeling, seismology, atmospheric science and computational biomedical research such as structural biology, pathology and fMRI. We will also highlight our participation in an interesting metacomputing project in collaboration with Sandia National Laboratory and the High Performance Computing Center at Stuttgart, Germany.

Real World Computing Partnership
Booth R218

We have been developing a parallel programming environment on workstation and PC clusters. A programming language MPC , an operating system SCore-D, and a low-level communication library PM have been co-designed to achieve time-sharing high performance computing environment on clusters. Our latest PC cluster, consisting of 64 Pentium PRO processors connected by Myrinet, will be brought to the exhibits. We will not only demonstrate the ability of our system software but also show our parallel applications. One of our parallel applications is an integrated system for protein structure analysis. This system provides researchers of structural biology with a framework of parallel programming in MPC on SCore-D for their particular protein structure analysis. The demonstration shows, as an example, the performance of searching thousands of proteins in Protein Data Bank (PDB) for a homologous (i.e. similar) protein structures and/or sequences to a query protein, on the PC Cluster.

Saitama University
Booth R202

The main features of the research exhibit of Saitama University are: educational program using visualization tools such as AVS and VRML, research activities using vector and parallel supercomputers, and activities in scientific visualization and networking. The 3-D graphics and virtual reality systems are intensively regarded as advanced tools applied to educational programs and research. Applications such as AVS, VRML, and others are used in actual exercises for first and second grade students, where it was found that the younger students are ready to accept the virtual reality system or 3-D environment. This may be due to their popular experiences on 3-D video games at home. High performance computing has been carried out by using vector and parallel computers in the fields of scientific and engineering studies, and has been shown to give detailed understandings, particularly when visualizations are included. Analysis using motion pictures has been examined and applied in the fields of sporting education and others such as welfare.

SC97 Education Program
Booth R210

Booth The SC97 Education Program is designed to give K-12 teachers, administrators, university professors, educational software researchers and developers, community technology leaders, and others three days of intense exposure to the tools and techniques of high performance computing and high-speed networking as they can be applied to education. Drop by this exhibit to learn about the SC97 and SC98 Education Programs, and to find out how to become involved in outreach and education programs nationwide.
http://www.supercomp.org/sc97

SC97 Research Exhibits
Booth R314

Drop by to learn how you can participate in a research exhibit at SC98, which will take place November 7-13, 1998 in Orlando, Florida, USA.
http://www.supercomp.org/sc97
http://www.supercomp.org/sc98

Scalable I/O Initiative
Booth R308

To achieve balance between compute power and I/O capability, the Scalable I/O Initiative has adopted a system-wide perspective to analyze and enhance the diverse, intertwined system software components influencing I/O performance. Using the PABLO performance analysis environment (http://www-pablo.cs.uiuc.edu/Projects/Pablo/pablo.html), application developers and computer scientists have cooperated to determine the I/O characteristics of a comprehensive set of I/O-intensive applications. These characteristics have guided development of parallel I/O features for related system software components: compilers, runtime libraries, parallel file systems, high performance network interfaces, and operating system services. The resulting software, built upon results from ongoing research projects sponsored by DARPA, DoE, NASA, and NSF, will be demonstrated using workstations and an Intel Paragon. Participating institutions include:

- Argonne National Laboratory, http://www.mcs.anl.gov/scalable/scalable.html
- University of Arizona, http://www.cs.arizona.edu/sio
- California Institute of Technology, http://www.cacr.caltech.edu/SIO/
- Carnegie-Mellon University, http://www.cs.cmu.edu/afs/cs/project/pdl/WWW/SIO.html
- University of Illinois at Urbana/Champaign, http://www-pablo.cs.uiuc.edu/Projects/IO/sio.html
- Northwestern University, http://www.cat.syr.edu/passion.html
- University of Maryland, http://www.cs.umd.edu/projects/hpsl/io/io.html
- Princeton University, http://www.cs.princeton.edu/sio
- Rice University, http://www.cs.rice.edu/~mpal/SIO/index.html
- University of Washington, http://www.cs.washington.edu/homes/wolman/sio-net.html

http://www.cacr.caltech.edu/SIO

University of Alaska Fairbanks
Booth R211

The Arctic Region Supercomputing Center (ARSC) supports the computational needs of academic, industrial, and government scientists and engineers with HPC resources, programming talent, technical expertise, and training. Scientific research performed at ARSC includes ocean modeling, atmospheric sciences, climate/global change, space physics, satellite remote sensing, and civil, environmental, and petroleum engineering. ARSC provides computational resources for research with emphasis on the high latitudes and the Arctic. ARSC, located at the University of Alaska Fairbanks, operates a CRAY Y-MP M98, a CRAY T3E and a network of SGI workstations in three visualization laboratories across the campus. Exhibit displays include: Arctic Ocean circulation and ice movement models, Alaska fly-by virtual reality based on SAR and satellite data, atmospheric modeling for weather prediction and contaminant transport, modeling of ionospheric dynamics during magnetic storms, and simulation of permafrost stabilization through subsurface cold air circulation.
http://www.arsc.edu
http://www.uaf.edu

University of California, Santa Barbara
Booth R109

Javelin: Internet-Based Parallel Computing Using Java. We present our research on a distributed, heterogeneous, high performance infrastructure for running coarse-grained parallel applications on Intranets or potentially even the Internet. Our approach is based on recent advances in Internet connectivity, and environments that allow for the safe execution of untrusted code such as the Java VM. We propose an architecture that uses supply and demand and market mechanisms to motivate individual users to offer their resources. Our approach has the potential for running parallel supercomputing applications involving thousands of anonymous machines on the Internet. We have implemented a prototype based on Java that clearly shows the feasibility of our approach. Our system‹called Javelin‹is built on Internet software that is interoperable, increasingly secure, and essentially ubiquitous: it requires participants to have access only to a Java-enabled Web browser. A Javelin-based parallel raytracer running on a heterogeneous network will be demonstrated.
http://www.cs.ucsb.edu/~schauser/papers/96-superweb.ps
http://www.cs.ucsb.edu/~schauser/papers/97-javelin.ps

University of Sao Paulo--Integrated Systems Laboratory
Booth R316

SPADE-2 is a parallel machine based on commodity processing nodes (PN) and high-speed interconnection networks (IN). Some of the research topics investigated in this project are: architecture support for both message passing and shared memory models, NUMA/CC-NUMA/COMA shared memory architectures, SW-DSM, light-weight communication protocols, design and implementation of high speed low latency INs, fast and scalable collective operations implemented in software and with hardware support, and operating systems for CC-NUMA multiprocessors. Several prototypes are being constructed in order to demonstrate the key ideas of the project: i) efficient implementation in network of workstations of shared memory model through SW DSM, and message passing through user-level light-weight protocols; ii) design and implementation of a low-cost high performance interconnection network with support for shared memory model and message passing; iii) low-cost HPC computing system implemented with commodity SHV components. The project also includes the development of several high performance computing applications.
http://www.lsi.usp.br/hpcac/spade2.html

University of Utah
Booth R205

The exhibit will showcase research in HPC and visualization at the University of Utah. A number of research centers or groups are involved, including:

- Center for Simulation of Accidental Fires and Explosions (C-SAFE), recently funded through the DoE Accelerated Strategic Computing Initiative (ASCI), and whose focus is on providing state-of-the-art, science-based tools for the numerical simulation of accidental fires and explosions.

  Henry Eyring Center for Theoretical Chemistry, whose research involves fundamental theoretical studies of the dynamics of complex, condensed matter systems.

- Scientific Computations and Imaging Laboratory (SCI), involved in both designing efficient and accurate tools for scientific computing, and in using these tools for scientific applications.

- Utah Tomography and Modeling/Migration (UTAM) consortium at the Department of Geophysics, whose mission is to develop innovative tomography, modeling and migration methods for exploration geophysics. Other participating research groups include the Medical Imaging Research Laboratory at the School of Medicine's Department of Radiology, the Mathematical Biology Group, Avalanche (Scalable Parallel Processor Project), the Cosmic Ray Group, the Biomagnetic Group, the SGI Supercomputing Visualization Center, the Quantum Chromodynamics group and the Cooperative Institute for Regional Weather Prediction. Computational resources at the University of Utah include a 64 node IBM SP, a 60 CPU SGI Origin 2000 with 8 Infinite Reality Graphics Engines, and a 16 processor SGI PowerChallenge. These resources are managed by the Center for High Performance Computing, which provides hardware, software, and consulting services to support HPC at the University of Utah.

http://www.chpc.utah.edu
http://www.chpc.utah.edu/sc97/

University of Virginia
Booth R206

The Legion Project at the University of Virginia is building metacomputing software to join thousands of machines, millions of users, and billions of objects. This software will join together diverse administrative domains, including NSF PACI participants (NCSA, NPACI), Universities (Caltech, UT Austin, UVa, etc.), DoE National Labs (Sandia, Los Alamos, and Livermore). Legion is also being deployed as part of the DoD Modernization project. We will demonstrate Legion version 1.0, running scientific computations on a metacomputer composed of systems from the above-mentioned sites. Our demonstrations will include distributed interactive applications, Web browsers for the Legion object space, and computationally intensive applications.

University of Wisconsin
Booth R313

Paradyn is a tool for measuring and analyzing the performance of parallel and distributed programs. Paradyn can measure large, long-running programs and provides facilities for helping to automatically find performance problems in parallel programs. Paradyn operates on executable files by dynamically inserting measurement code while the program is running. Paradyn can measure programs running on a variety of operating systems and platforms, or heterogeneous combinations of these systems. Paradyn can handle PVM and SP2 MPL and MPI. We'll exhibit Paradyn with a variety of applications and platforms, including:

- Various engineering and numeric codes developed at Wisconsin and running on our various Cluster, SMP, and MPP platforms.

- A Grand Challenge satellite imagery code developed at Maryland.

- Applications running on IBM PowerParallel exhibit.

- Demonstration of new tools built with the Paradyn DynInst API. This API allows tool builders to easily use our on-the-fly program instrumentation technology.

Proceedings Table of Contents

Final Program Table of Contents

SAN JOSE SC97 NOV. 15-21

Poster Exhibits

Poster exhibits report results, tools, and experiences related to high performance networking and computing. They offer researchers an excellent opportunity to present their results in an informal setting that encourages one-on-one interaction. Education posters provide insights into current issues in computing and education. All poster sessions are open to all attendees.

SC97 technical program posters are located in the Almaden Concourse of the San Jose McEnery Convention Center, along with the SC97 Education Program posters. All poster sessions are open to all SC97 attendees.

Technical Program Posters Hours

Monday	November 17	7-9pm	(Gala Opening--All Exhibits)
Tuesday	November 18	10am-6pm	(All Exhibits)
Wednesday	November 19	10am-6:30pm	(includes 5-6:30pm posters reception)
Thursday	November 20	10am-4pm	(All Exhibits)

Education Posters Hours

Monday	November 17	3:30-5pm
Tuesday	November 18	3:30-5pm (All Exhibits)
Wednesday	November 19	5-6:30pm (posters reception)

Education Posters	Systems Posters
Applications Posters	System/Runtime Tools Posters
Software Systems Posters	Benchmarking Posters
Visualization Tools Posters	Performance and Debugging Tools Posters

Training Poster

Education Posters

Please see abstracts in Education Program section.

P1 **Tessellation Tutorials**
Suzanne Alejandre, Math Forum

P2 **The Internet World's Fair by a K-6 School**
Deanna Alexander, Fisher Grade School, Fisher, IL; Melissa Kelly, NCSA

P3 **Online K-12 Earth System Science Lessons that Use Web Resources**
Brian Beaudrie, Timothy Slater, Dave Thomas, Montana State University

P4 **Use of the CAVE in Promoting Safe Street Crossing**
Patricia Brown, Columbia Elementary School, Champaign, IL; Georgette Moore, Yankee Ridge Elementary School, Urbana, IL

P5 **Cyber Mummy VR**
Renee Cooper, Pamela Van Walleghen, Urbana Middle School, Urbana, IL

P6 **Basic Aerodynamics Software Lab**
John D. Eigenauer, NASA Lewis Research Center

P7 **Virtual Reality Technology Enhances Learning for Interdisciplinary Studies on the Missouri River**
Carol Engelman, Ralston High School, Omaha, NE; Elaine Westbrook, Omaha North High School, Omaha, NE

P8 **Fusion of Art and Technology**
Kathy Felling, Karen Ciprari-Murphy, Fulton County Board of Education, Atlanta, GA

P9 **Implementing the National Science Education K-12 Standards Using NASA Web Resources**
Robert Fixen, Timothy Slater, George Tuthill, Stephanie Stevenson, Montana State University

P10 **Entry-level Virtual Reality**
Jane Leggett, Wilson Middle School, Moline, IL

P11 **NASA Lewis Learning Technologies Program**
Beth Lewandowski, NASA Lewis Research Center

P12 **The National Education Supercomputer Program**
Brian Lindow, Lawrence Livermore National Laboratory

P13 **Multiple Representations of the Periodic System of Elements: the Design and Evaluation of a Multimedia Hypertext**
Punyashloke Mishra, NCSA

P14 **Virtual Reality at the Jersey Shore**
Alan Sills, West Essex Regional High School

P15 **Yohkoh Public Outreach Project: Using the Web and HPCC for Lifelong Learning**
Timothy Slater, Montana State University; Marion French, Kansas State University; Donna Governor, Escambia County School District, Pensacola, FL

P16 **Earth System Science High School Classroom Applications Using Scientific Visualization and Distance Learning Experiences**
Robert Smith, Michael Porter, Terry Parker High School, Jacksonville, FL

Applications Posters

P21 High Performance Parallel Sparse LU Factorization with Partial Pivoting

Xiangmin Jiao, Cong Fu, Tao Yang, University of California, Santa Barbara

We present a high performance algorithm for Gaussian-elimination-based sparse LU factorization with partial pivoting. A recent performance study of this newly revised algorithm considers memory constraints. Using the new techniques, we achieved up to 6.878 GFLOPS on 128 T3E nodes. This is the highest known performance for this challenging problem.

P22 On Improving the Performance of Parallel Sparse Matrix-vector Multiplication

James White, Ohio Supercomputer Center; P. Sadayappan, Ohio State University

We describe modifications to the compressed-sparse-row format for unstructured sparse matrices that improve the performance of sparse matrix-vector multiplication by increasing the availability of fine-grained parallelism for modern pipelined, superscalar processors. The modifications prove particularly effective for matrix-vector systems that fit into cache. We demonstrate dramatic superlinear speedup for realistic sparse matrices.

P23 Large-scale Material Science Calculations on the Cray T3E900

Andrew M. Canning, NERSC

The Plane Waves, LDA method is an extreme test of the communication network of parallel machines. It demands clever communication algorithms to run on large machines. We will present results of codes that use this method running on up to 512 Cray T3E900 processors that demonstrate the processor speed and bandwidth performance.

P24 Parallelization of 3D Spectral Laser Plasma Interaction Code using High Performance Fortran

Vadim Elisseev, University of Alberta

We present the results from a parallelization of sequential code used for investigating parametric instabilities in large-scale laser produced plasmas. A detailed quantitative performance analysis of the parallel code (speedup, scalability, and load balancing) has been carried out for an implementation on an IBM SP2 using the XL HPF compiler.

P25 Stencil Methods on Distributed and High Performance Computers

Heath James, Craig Patten, Ken Hawick, University of Adelaide

A distributed computational infrastructure for applying stencil or kernel operators on arbitrarily sized images is described. The system allows for deferred application of a compound stencil operator constructed from a sequence of primitive operations. Performance has been measured on traditional parallel supercomputers and less tightly coupled systems.

P26 Large-scale SVD Computations in Linear Geophysical Inverse Problems

Osni Marques, Lawrence Berkeley National Laboratory

We focus on the main issues related to the solution of SVD problems arising in geophysics by means of an eigenvalue approach on a Cray T3E-900, and strategies to tackle very large cases. Although sparse, such problems are difficult to deal with since several thousand singular values and singular vectors are required.

P27 Parallel Monte Carlo Transport Modeling in the Context of a Time-Dependent, 3-D Multi-Physics Code

Richard Procassini, Lawrence Livermore National Laboratory

This implementation of the MONACO Monte Carlo transport package utilizes explicit data communication (message passing) between domains. The resultant parallel transport package will experience non-deterministic communication patterns. The communication of particles between subdomains during a Monte Carlo time step may require a significant level of optimization to achieve a high parallel efficiency.

P28 Nonlinear Nonparametric Population Pharmacokinetic Modeling on a Supercomputer

Michael Van Guilder, Alan Schumitzky, Xin Wang, Roger Jelliffe, USC Laboratory of Applied Pharmacokinetics, Los Angeles; Robert Leary, SDSC; Alexander Vinks, The Hague Hospital

We extended our nonparametric EM (NPEM) PC software to include nonlinear systems, incorporating a differential equation solver and parallelization. For the Cray T3E, speedup is 24-fold for 32 nodes. On 8 nodes, a 5-parameter Michaelis-Menten NPEM model of piperacillin, using 80,000 grid points, took 1.2 hrs.

P29 **Nonlinear Nonparametric Population Pharmacokinetic Modeling on a Supercomputer Modeling of Touch Mode Capacitive Sensors and Diaphragms**
Qiang Wang, Wen Ko, Case Western Reserve University

The finite element modeling of touch mode capacitive diaphragms aimed to derive computer design aid for these diaphragms is presented. By geometrically nonlinear computing and using GAP element in ABAQUS, the deflection, stress and strain of a diaphragm in touch mode can be calculated. The simulation results can be used to predict sensor performance.

P30 **A Parallel Lanczos Method for Symmetric Eigenvalue Problems**
Kesheng Wu, Horst Simon, Lawrence Berkeley National Lab/NERSC

This poster presents a parallel version of LANSO, pLANSO, which implements a Lanczos iteration for symmetric generalized eigenvalue problems. It uses the Omega-recurrence to detect loss of orthogonality among the Lanczos vectors. In most of the cases tested, pLANSO runs with the same or better efficiency as the parallel matrix-vector multiplication.

P31 **Modeling Groundwater Flow through Heterogeneous Porous Media on Massively Parallel Computers**
Steven Ashby, Chuck Baldwin, William Bosl, Steve Carle, Robert Falgout, Richard Hornung, Reed Maxwell, Nina Rosenberg, Dan Shumaker, Steve Smith, Andrew Tompson, Carol Woodward, Lawrence Livermore National Laboratory

A sophisticated, scalable simulation code for modeling flow and multicomponent transport through 3-D heterogeneous porous media is presented. The simulator includes scalable subsurface modeling capabilities, a fast flow solver, and accurate component transport schemes. The simulator runs on a variety of computing platforms, from PCs to massively parallel computers.

P32 **A Parallel, Adaptive Refinement Model for Groundwater Flow and Transport**
Alan Stagg, Joseph Schmidt, U.S. Army Waterways Experiment Station

The USAE Waterways Experiment Station has developed a parallel groundwater model in response to the DoD's need to conduct large-scale simulations of groundwater remediation scenarios. The Adaptive Hydrology finite element model has been targeted for distributed memory systems supporting a message-passing paradigm. Flow solutions for groundwater sites and performance results are presented.

P33 **Parallel Solution of a Streamline Upwind Petrov-Galerkin Model of the Shallow Water Equations**
Steve Bova, Mississippi State University

The parallel performance of a streamline upwind Petrov-Galerkin finite element approximation to the 2-D, depth-averaged shallow water equations is presented. Time-iterative solutions are obtained using enhanced stability, multi-stage Runge-Kutta methods on unstructured triangular grids. Results are computed for a portion of the upper Mississippi River.

P34 **A Heterogeneous SPMD Adaptive Finite Element Code for Simulating Multiphase Flow Through Porous Media-Linux and Cray T3E Implementations**
Don Morton, University of Montana/Arctic Region Supercomputing Center

A research code for the exploration of parallel adaptive finite element methods has been developed to simulate two-phase oil-water flow through porous media. The poster outlines the modular, heterogeneous approach taken and presents examples of the application area and performance results.

P35 **Three-Dimensional Simulations of Compressible Turbulence on Leading-Edge Parallel Platforms**
Arthur Mirin, Ronald Cohen, William Dannevik, Andris Dimits, Ronald Eastman, Donald Eliason, Brendan McNamara, Oleg Schilling, Lawrence Livermore National Laboratory; Steven Orszag, Cambridge Hydrodynamics, Inc.; David Porter, Paul Woodward, University of Minnesota

A parallelized, 3-D hydrodynamics code based on the Piecewise Parabolic Method is used to examine compressible fluid turbulence in three dimensions. We consider both Direct Numerical Simulation and Large Eddy Simulation; in the latter case the idea is to run at coarser resolution through Subgrid-scale parameterization. Performance results are presented.

P36 **A Portable and Scalable 3-D Hydrodynamic Simulator for Parallel Computers**
Jianping Zhu, Purushotham Bangalore, David Huddleston, Anthony Skjellum, Mississippi State University; Billy Johnson, US Army Engineer Waterways Experiment Station

We will discuss the development of a parallel version of a 3-D, hydraulic simulation code which has been used to model various coastal and estuarien phenomena. Timing results for simulating New York Bay will be presented to demonstrate significant reduction of the execution time.

P36 **A Portable and Scalable 3-D Hydrodynamic Simulator for Parallel Computers**
Jianping Zhu, Purushotham Bangalore, David Huddleston, Anthony Skjellum, Mississippi State University; Billy Johnson, US Army Engineer Waterways Experiment Station

We will discuss the development of a parallel version of a 3-D, hydraulic simulation code which has been used to model various coastal and estuarien phenomena. Timing results for simulating New York Bay will be presented to demonstrate significant reduction of the execution time.

Software Systems Posters

P37 **HPF for Practical Scientific Algorithms**
Chris Ding, Berkeley National Laboratory

Scaling of HPF codes is not as good as MPI codes. HPF codes performances are highly inconsistent, i.e., a minor change of code results in order of magnitude changes in performance. Many peculiarities of HPF coding will be discussed.

P38 **High Performance Parallel Programming in Java**
Vladimir Getov, Sava Mintchev, University of Westminster; Susan Flynn-Hummel, IBM T.J. Watson Research Center

This piece of research introduces a way of successfully tackling the difficulties in binding both scientific and message-passing libraries to Java. We have created a tool (JCI) for automating the creation of portable interfaces to native existing libraries in Fortran77 and C. Evaluation results on the IBM SP2 at Cornell are presented using the IBHPCJava compilertes a native code the RS60000 architecture.

P39 **Combining Interprocedural Compile-time and Run-time Parallelization**
Mary Hall, Sungdo Moon, Byoungro So, University of Southern California

This project focuses on combining high-quality interprocedural analysis with information from the run-time environment to push the boundaries of automatic compiler parallelization technology. With the Stanford SUIF compiler group, we have demonstrated a Specfp95 ratio higher than any reported results on an 8-processor Digital AlphaServer.

P40 **DISCWorld: A Distributed Information Systems Control Environment**
Ken Hawick, Heath James, Craig Patten, Francis Vaughan, University of Adelaide

A general computational framework to enable application domain experts to construct distributed, high performance programs is described. User programs in the DISCWorld environment are assembled and monitored using Web technology. Programs can operate on large data sets encapsulated within the DISCWorld's hierarchical storage environment.

P41 **Parallel Make for the IBM SP2**
Richard Johnson, Lawrence Livermore National Laboratory

We report on a new version of Pmake that is based on MPI. The master node performs the usual makefile processing, builds the dependency trees, and then acts as a taskmaster for the other nodes. This version of Pmake has demonstrated nearly perfect scalability.

P42 **Compositional C++**
Carl Kesselman, Mei-Hui Su, University of Southern California

P43 **High Performance Distributed Computing**
Carl Kesselman, Steve Fitzgerald, Bob Lindell, Karl Czajkowski, Soonwook Hwang, University of Southern California

P44 **CCF: Collaborative Computing Frameworks**
Vaidy Sunderam, Emory University

CCF (Collaborative Computing Frameworks) is a suite of software systems and tools, communications protocols, and methodologies that enable collaborative, computer-based cooperative work. CCF constructs a virtual work environment on multiple computer systems connected over the Internet. A prototype of a complete collaboration system will be exhibited.

Visualization Tools Posters

P45 **Use of an Interactive VRML Visualization Tool by Educators**
Thomas Bulka, Northern High School; David Emigh, Quinebaug Valley Community Technical College; Erich Schroeder, Illinois State Museum; Ray Plante, Dee Chapman, Umesh Thakkar, NCSA/University of Illinois at Urbana-Champaign

The VRML Server is a tool that allows researchers, educators, and students to interactively create 3-D visualizations of astronomical images. The visualizations are returned to the user as VRML files. The presentation will illustrate learning and teaching in classrooms using this tool. http://imagelib.ncsa.uiuc.edu/imagelib/VRMLServer/

P46 **DT - A Scalable Visual Tool**
Diana Keen, Lawrence Livermore National Laboratory

Running parallel jobs adds a new challenge to performance analysis tools in deciding how to show many processors and nodes. We present a scalable interface that is adapted to SMPs and allows viewing at least 4,096 processors. The user may view at the processor level or node level.

P47 **Collaborative Scientific Data Visualization**
Scott Klasky, Byeongseob Ki, Northeast Parallel Architectures Center

We present a collaborative scientific visualization package (Scivis). Its two key features are that the filters are user definable, and the system is fully collaborative between multiple researchers. Our system is being used in the NSF Grand Challenge, "The Binary Black Hole Grand Challenge," an alliance among 10 universities.

P48 **When High Performance Computing Environments and Visual Arts Meet: New Modes of Visualization and Interaction**
Rosemarie McKeon, Dru Clark, Michael Bailey, University of California, San Diego and SDSC; Richard Marciano, SDSC , University of California, San Diego

As computer technologies rapidly shift, a major question is whether artists will respond with work that makes creative use of these changes. Conversely, can new multidisciplinary collaborative projects help scientists gain fresh insights into techniques they use to communicate results? These concerns are addressed in the context of high performance computing environments.

Training Poster

P49 **HPC Training at CEWES**
John Eberle, Nichols Research

The US Army Corp of Engineers Waterways Experiment Station Major Shared Resource Center provides training in high performance computing to a geographically diverse set of users. This presentation will describe the technology used to provide training via on-site classes, broadcasts to remote locations, and archives of recorded lectures.

Systems Posters

P50 **Embeddable Supercomputer Technology**
Jay Block, Jeff Koller, Jeff Draper, Craig Steele, Claire Lacour, University of Southern California

Advanced packaging enables scalable systems with very high densities, unprecedented performance, and novel architectures. We present three interrelated systems that are being built: Packaging-Driven Scalable System (PDSS), using area I/O for communication; the Integrated Thermal Management (ITeM) system, achieving 1 Teraflop per cubic foot; and ASNT, a multiple-network machine.

P51 **Single System Image: Need, Approaches, and Supporting Systems**
Rajkumar Buyya, Centre for Development of Advanced Computing

For high performance computing on proprietary or commodity hardware to be practicable, it is important that systems provide a single system image at any one (or more) of the following levels: hardware, operating system, message-passing interfaces, language/compiler, or tools. This poster presents the needs and approaches for building a single system image.

P52 **A Reconfigurable SIMD Machine for Parallel Genetic Algorithms Based on Farming Model**
Tomio Inoue, Masahiko Sano, Yoshizo Takahashi, The University of Tokushima

A reconfigurable SIMD machine for parallel genetic algorithms (PGA), named OCTOPUS, is developed. OCTOPUS is constructed as an oct-tree to meet the farming model known as one of the most promising PGA models. Each PE is implemented in a field programmable gate array. Node PEs are connected to a shared memory matrix.

P53 **Global/Mobile Universal Modular Packaging System**
Jeff LaCoss, John Granacki, University of Southern California

The DARPA-sponsored GUMPS project is developing architectures and new packaging methods for mixed-mode (combined analog and digital) systems. We are delivering hand-held computing nodes and high performance RF modules designed to facilitate research and experiments in nomadic computing and wireless networking and communication.

P54 **Performance Analysis of COMPaS: A High Performance Pentium-Pro PC-based SMP Cluster**
Mitsuhisa Sato, Yoshio Tanaka, Motohiko Matsuda, Kazuto Kubota, Makoto Ando, Jens Gerlach, Tsukuba Research Center, Japan

We have built a cluster of SMPs, COMPaS (Cluster Of MultiProcessor System), which consists of eight nodes of 4-way Pentium-Pro (200MHz) SMPs connected by a Myrinet High-speed Network and 100Base-T Ethernet. To use SMP-node locality, we integrate multithreaded programming with Solaris threads for intra-nodes and remote-memory based message passing for inter-nodes.

System/Runtime Tools Posters

P55 **Generating Distributed Job Scheduling Information**
Heath James, Ken Hawick, Kevin Maciunas, University of Adelaide

A model for generating the placement and scheduling information for jobs within a heterogeneous group of distributed, high performance computers is described. The system addresses selection and dynamic discovery of computational and storage services, minimization of data movement costs, and optimization of job placement with respect to a cost metric.

P56 **Maui Scheduler, an Advanced System Software Tool**
Margaret Lewis, Lyn Gerner, MHPCC

The Maui Scheduler attempts to maximize SP node utilization and to make available various "qualities of service" (priorities), with rapid reconfigurability. Better use is achieved through the use of a single queue, using a scheduling algorithm with several configurable parameters to control the flow of the workload.

P57 **RSVP: Reserving Resources on the Internet**
Bob Lindell, University of Southern California

RSVP is a setup reservations protocol that creates flow-specific reservation state in routers and hosts, extending the Internet architecture to provide inherent support for quality of service. RSVP provides robust, efficient, flexible, and extensible service for multicasts and unicasts, and is an approved Internet Proposed Standard.

P58 **The Design and Implementation of a Disk Cache Manager for Multi-Node Shared Disk**
Joo Man Kim, Electronics and Telecommunications Research Institute

The design and implementation of the Disk Cache Manager, which operates on top of the microkernel-based parallel operating system, MISIX, and enhances the I/O performance of the shared disks is presented. To increase the shared disk I/O performance, the DCM employs Seven and Sodd Disks for each physical shared disk.

Benchmarking Posters

P59 **Performance and Scalability of the NAS Parallel Benchmarks on the Cray T3E**
Steven C. Caruso, Cray Research/SGI; William C. Saphir, Lawrence Berkeley National Laboratory

This poster presents the results of running the NAS Parallel Benchmarks (NPB) versions 1 & 2 on a T3E-900 with over 1000 processors. We will demonstrate the performance and scalability of this benchmark on the T3E, and provide analysis to explain the observed behavior.

P60 **New NAS Parallel Benchmarks Results**
Maurice Yarrow, Alex Woo, Rob Van der Wijngaart, NASA Ames Research Center; William Saphir, National Energy Research Scientific Computing Center

We present the latest performance results for the NAS Parallel Benchmarks (NPB) version 2 on many architectures, including: IBM P2SC/120 MHz, SGI Origin 2000, Cray Research T3E, Berkeley N.O.W. (UltraSPARC cluster), Sun Ultra Enterprise 4000, and Intel PentiumPro Cluster. These results provide an insightful comparison of the performance of high performance computers.

P61 **Parallel Performance Evaluation and Modeling in the DoD Modernization Program**
Jack Dongarra, Phil Mucci, Erich Strohmaier, University of Tennessee; Gina Goff, Charles Koelbel, CRPC, Rice University

Performance evaluation is done using the ParkBench codes. We are extending the ParkBench low-level, kernel, and compact applications to include examples representative of DoD applications and HPF codes. Our long-term goal is to use ParkBench results as inputs to performance models, which can in turn estimate performance of full applications.

P62 **The Effects of Cache Performance and Memory Bandwidth Effects in Densely Strided Floating Point Vector Operations**
Philip Mucci, Erich Strohmaier, Jack Dongarra, University of Tennessee

As part of the DoD HPC Modernization program's Performance Evaluation Team, we are studying the effects of the memory hierarchy on floating-point kernels with two custom benchmarks: cachebench and blasbench. Analysis of the performance of the O2K's R10000 and the SP2's P2SC 120, and their effectiveness to sustain maximal floating-point performance is presented.

Performance and Debugging Tools Posters

P63 **A Common Parallel Tools Environment for the DoD HPC Modernization Program**
Shirley Browne, Jack Dongarra, Kevin London, University of Tennessee; Chuck Koelbel, CRPC, Rice University

This poster reports on our preliminary efforts toward providing a consistent set of debugging and performance tools across DoD Majored Shared Resource Center computing platforms. To help provide a common parallel debugging environment, we are participating in the High Performance Debugging Forum (HPDF).

P64 **POEMS-PerfOrmance End-to-end Modeling System**
Jim Browne, The University of Texas at Austin; Vikram Adve, Rice University; Rajive Bagrodia, UCLA; Elias Houstis, Purdue University; Olaf Lubeck, Los Alamos National Laboratory; Pat Teller, Richard Oliver, The University of Texas at El Paso; Mary Vernon, The University of Wisconsin

The POEMS project will provide a cohesive framework for the end-to-end performance modeling of complex parallel and distributed systems including applications, runtime software, and hardware. The goal of the project is to accomplish system performance evaluation at 10% of the cost of prototyping and in advance of availability of hardware.

P65 **SIS: SCALA Instrumentation System**
Mario Pantano, University of Vienna

This poster presents the SCALA Instrumentation System (SIS), a compiler-integrated performance measurement tool that allows an automatic and selective instrumentation of irregular codes parallelized by the Vienna Fortran Compiler (VFC). SIS supports the measurement of important code regions. SIS is fully integrated into a restructuring compilation system.

P66 **S-Check: A Sensitivity Analysis Tool for Detecting Bottlenecks in Parallel Programs**
Robert Snelick, National Institute of Standards and Technology

We present a novel tool for identifying performance bottlenecks in parallel and networked programs. S-Check is a highly-automated sensitivity analysis tool that extends benchmarking and conventional profiling. A case study is presented that compares and contrasts sensitivity analyses of the same program on different architectures and offers solutions for performance improvement.

P67 **High MFLOP Rates for Out of Core Stencil Calculations Using Time Skewing**
Dave Wonnacott, Jaime Spacco, Haverford College; Tina Shen, Bryn Mawr College

Empirical evidence is presented that shows that time-skewing is faster than the traditional method as the machine balance increases. The speed of the code generated by the Omega Calculator is within a few percent of a hand-optimized version. Both of these versions are orders of magnitude faster than the code without time-skewing.

Performance Prediction Posters

P68

Predicting the Performance of Parallel Algorithms on Parallel Architectures

Zahira S. Khan, Bloomsburg University

Performance of the parallel hash-join algorithm is examined on the CM-2. The results obtained from experiments using different combinations of data size, data distributions, and join selectivities explain the behavior of the hash and join phases of the algorithm for the observed levels of these three factors.

P69

Predicting the Performance of Parallel Algorithms on Parallel Performance Prediction using Stochastic Values

Jennifer Schopf, Francine Berman, UC San Diego

It is difficult to make predictions of application performance in a dynamic distributed computing environment. We address this problem by allowing characteristic application and system data to be represented in performance models as an interval of possible values, and an interval of possible values and associated probabilities, respectively.

Proceedings Table of Contents

Final Program Table of Contents

SAN JOSE SC97 NOV. 15-21

Tutorials

Sunday Monday

Sunday

S1 **CANCELLED**

S2 Sunday, November 16
Half Day 8:30am-noon

Parallel Solution of PDEs Using the PETSc Software

Satish Balay, William D. Gropp, Lois Curfman McInnes, Barry F. Smith, Argonne National Laboratory
Level: 25% Beginner, 50% Intermediate, 25% Advanced

It is widely accepted that porting a full PDE application code to the distributed-memory parallel environment carries a cost of multiple person-months. However, parallel software libraries can ease this transition for porting codes as well as for constructing new applications. The Portable, Extensible Toolkit for Scientific Computation (PETSc) http://www.mcs.anl.gov/petsc/petsc.html, developed at Argonne National Laboratory, has evolved rapidly during the past few years, and many serial codes can now be ported to distributed-memory architectures using this library. Much of the application-specific portions of numerical simulations, representing the highest value of legacy serial codes in Fortran77 or C, can be carried over without serious modification.

S3 Sunday, November 16
Half Day 1:30pm-5pm

A Rapid and Practical Introduction to the Message Passing Interface (MPI)

William C. Saphir, Lawrence Berkeley National Laboratory (NERSC)

Level: 40% Beginner, 40% Intermediate, 20% Advanced

MPI has taken hold as the library of choice for portable high performance message passing applications. The complexity of MPI presents a short but steep learning curve. The large amount of overlapping functionality presents a confusing array of choices for non-experts, often leading to incorrect or inefficient code. This tutorial will provide a rapid introduction to MPI, to motivate and describe its core functionality. The emphasis will be on how to obtain high performance and on understanding what is useful and what is not. It is appropriate for application programmers who are thinking of using MPI or who have started using MPI but want to learn more.

S4 Sunday, November 16
Half Day 1:30pm-5pm

A Tutorial Introduction to High Performance Data Mining

Robert L. Grossman, University of Illinois at Chicago
Level: 50% Beginner, 30% Intermediate, 20% Advanced

Data mining is concerned with uncovering patterns, associations, changes, anomalies, and statistically significant structures and events in data. Traditional data analysis is assumption driven in the sense that a hypothesis is formed and validated against the data. Data mining in contrast is discovery driven in the sense that patterns are automatically extracted from data. The goal of this tutorial is to provide researchers, practitioners, and advanced students with an introduction to data mining. The focus will be on algorithms, software tools, and system architectures appropriate for mining massive data sets using techniques from high performance computing. There will be several running illustrations of practical data mining, including examples from science, business, and computing. We will describe several architectural frameworks for high performance data mining systems and discuss their advantages and disadvantages. Finally, we will cover several case studies involving mining large data sets, from 10-500 Gbytes in size.

S5 Sunday, November 16
Full Day 8:30am-5pm

Performance Evaluation of Parallel Applications: Measurement, Modeling, and Analysis

Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory
Subhash Saini, NASA Ames Research Center
Erich Strohmaier, University of Tennessee
Patrick H. Worley, Oak Ridge National Laboratory
Level: 25% Beginner, 50% Intermediate, 25% Advanced

This tutorial will give a comprehensive introduction to the methodology and usage of Performance Evaluation of Parallel Applications. Various issues related with workload-driven evaluation and characterization of applications and systems will be discussed. These issues will be illustrated with real application codes. The team will also present the latest benchmark results including NPB (vendor optimized, HPF based and MPI based) results for the state-of-the-art vector-based as well cache-based supercomputers. The methods for analytical performance modeling and prediction will be introduced in detail, and their application will be shown in case studies. Different statistical methods for the analysis of benchmark results and benchmark suites will be introduced and their application to standard benchmark suites such as NPB will be shown.

S6 Sunday, November 16
Full Day 8:30am-5pm

How to build a Beowulf: Assembling, Programming, and Using a Clustered PC Do-it-yourself Supercomputer

Thomas Sterling, California Institute of Technology
Level: 25% Beginner, 50% Intermediate, 25% Advanced

It has recently become possible to assemble a collection of commodity mass market hardware components and freely available software packages in a day and be executing real world applications by dinner time to achieve a sustained performance at greater than 1 Gflops at a total cost of around $50,000. Furthermore, on almost a daily basis, these numbers are improving. This full-day tutorial will cover all aspects of system assembly, integration, software installation, programming, application development, system management, and benchmarking. Demonstrations with actual hardware and software components will be conducted throughout the tutorial. Participants will be encouraged to closely examine and manipulate elements of a Beowulf at various stages of integration with strong Q&A interaction between presenters and attendees. The presenters will include David Bailey (applications benchmarking), Don Becker (networking), Jack Dongarra (applications benchmarking), Al Geist (PVM), Ewing (Rusty) Lusk (MPI), John Salmon (applications), and Thomas Sterling (system structure).

S7 Sunday, November 16
Full Day 8:30am-5pm

Shared Memory Parallel Programming

Robert H. Kuhn, Kuck & Associates, Inc.
Rudolf Eigenmann, Purdue University
Level: 50% Beginner, 25% Intermediate, 25% Advanced

Shared Memory Parallel (SMP) systems are well established. They are considered mainstream now by all workstation vendors. On the frontiers, at the low end, hardware trends make it almost as easy to put four Intel processors in a system as one. At the high end, the trend is toward NonUniform Memory Access (NUMA) systems. This tutorial will give a novice in parallel programming the insight, tools, and techniques one needs to program today's SMP systems. The tutorial covers three perspectives one needs to know: First, as orientation, today's SMP systems and the trends for tomorrow are surveyed. Second, practical, available programming models for SMP parallelism are reviewed and contrast. Third, a parallel software engineering methodology is covered. That is, the instructors will walk through the steps of parallelizing an application and at each point describe some of the tools and techniques that are available. Practical examples from scientific engineering applications are provided to demonstrate key points and stimulate attendees to discuss the problems they face in parallelizing applications.

S8 Sunday, November 16
Full Day 8:30am-5pm

High Performance I/O Systems: Architectures, Software and Applications

Rajesh R. Bordawekar, California Institute of Technology
Alok Choudhary, Northwestern University
Level: 30% Beginner, 50% Intermediate, 20% Advanced

Large-scale computing includes many application areas with intensive I/O demands, including scientific computing, databases, mining, decision support, and multimedia. For high performance, it is critical to improve I/O performance of these applications. There are many solutions at different levels to address I/O performance problems. This tutorial will present architecture and software issues in designing scalable and efficient parallel I/O systems as well as I/O requirements, characteristics and examples from application domains of science and engineering, databases, data mining and multimedia systems. The instructors will discuss interactions between compute node architectures, hierarchical storage systems, and system software. The discussion will also include examples of commercial systems as well as software research projects. Finally, applications from various domains listed above will be discussed with respect to their I/O requirements and their impact on architectures and system software.

S9 Sunday, November 16
Full Day 8:30am-5pm

High Performance Fortran‹Practice and Experience

John Towns, NCSA
Subhash Saini, NASA Ames Research Center
Level: 15% Beginner, 85% Intermediate

Throughout 1997, NCSA has conducted a series of hands-on workshops designed to teach effective strategies for using parallel computers. An HPF component in these workshops has enabled NCSA users to migrate several large-scale applications from the CM-5 to NCSA's currently available SGI Origin and Hewlett-Packard Exemplar platforms, as well as IBM SP2 and CRAY T3E systems at other NSF centers. It has also enabled many users new to parallel computing to get up and running quickly on NCSA's parallel systems. Also, in February 1997, the first HPF User's Group meeting was held in Albuquerque, NM. A series of high-quality talks were presented by HPF users working on large-scale industrial HPF applications, benchmarks, libraries, and programming tools. This tutorial will combine material from NCSA's HPF Programming module with selected presentations from HPF users to provide a thorough introduction to HPF applications development for existing and prospective users of the HPF language.

Monday

M1 Monday, November 17
Half Day 8:30am-noon

Tuning MPI Applications for Peak Performance

Ewing (Rusty) Lusk, William Gropp, Argonne National Laboratory
Level: 50% Intermediate, 50% Advanced

MPI is now widely accepted as a standard for message-passing parallel computing libraries. Both applications and important benchmarks are being ported from other message-passing libraries to MPI. In most cases it is possible to make a translation in a fairly straightforward way, preserving the semantics of the original program. On the other hand, MPI provides many opportunities for increasing the performance of parallel applications by the use of some of its more advanced features, and straightforward translations of existing programs might not take advantage of these features. New parallel applications are also being written in MPI, and an understanding of performance-critical issues for message-passing programs, along with an explanation of how to address these using MPI, can provide the application programmer with the ability to provide a greater percentage of the peak performance of the hardware to his/her application. This tutorial will discuss performance-critical issues in message passing programs, explain how to examine the performance of an application using MPI-oriented tools, and show how the features of MPI can be used to attain peak application performance. It will be assumed that attendees have an understanding of the basic elements of the MPI specification. Experience with message-passing parallel applications will be helpful but not required.

M2 Monday, November 17
Half Day 1:30-5:00pm

MPI-2

William C. Saphir, Lawrence Berkeley National Laboratory (NERSC)
Level: 20% Beginner, 60% Intermediate, 20% Advanced

MPI-2 is a set of extensions to the MPI (Message Passing Interface) standard. It was finalized by the MPI Forum in June 1997. This tutorial will describe the new features in MPI-2, including I/O, one-sided communication, dynamic process management, and support for C and Fortran 90. Since MPI-1 already provides core functionality, many MPI-2 features are not appropriate for all programs. The tutorial will help participants sort out what is useful and when it is useful.

M3 Monday, November 17
Full Day 8:30am-5pm

Linear Algebra Algorithms and Software for Large Scientific Problems

Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory
Iain S. Duff, Rutherford Appleton Laboratory
Danny C. Sorensen, Rice University
Henk vander Vorts, Utrecht University
Level: 20% Beginner, 50% Intermediate, 30% Advanced

This tutorial will discuss a variety of algorithms for solving linear algebra problems indicating where each is appropriate and emphasizing their efficient implementation. In particular, the development of vector and parallel computers since the late 1970s led to a critical review of mathematical software. Many of the sequential algorithms used satisfactorily on traditional machines fail to exploit the architecture of advanced computers. The LAPACK package will be highlighted. This package provides a choice of algorithms mainly for dense matrix problems that are efficient and portable on a variety of high performance computers. For large sparse linear systems, the situation is more complicated and a wide range of algorithms is available. An introduction and guidelines on the selection of appropriate software will be given. The tutorial will consider both direct methods and iterative methods of solution, including some recent work that can be viewed as a hybrid of the two. In the case of direct methods, emphasis will be on frontal and multifrontal methods, including variants performing well on parallel machines. For iterative methods, the discussion will include CG, MINRES, SYMMLQ, BiCG, QMR, CGS, BiCGSTAB, GMRES, and LSQR. For large (sparse) eigenproblems, the discussion will be on some of the most widely used methods, such as Lanczos, Arnoldi, and Davidson. Particular attention will be given to efficient implementation of Arnoldi's method and the Implicitly Restarted Arnoldi Method, along with guidelines for their use, preconditioning, and hints for the selection of these algorithms.

M4 Monday, November 17
Full Day 8:30am-5pm

Introduction to Effective Parallel Computing

Marilynn Livingston, University of Oregon
Quentin F. Stout, University of Michigan
Level: Level: 50% Beginner, 50% Intermediate

This tutorial will provide a comprehensive overview of parallel computing, focusing on the aspects most relevant to the user. Throughout, the emphasis will be on the iterative process of converting a serial program into an increasingly efficient, and correct, parallel program. The tutorial will help people make intelligent planning decisions concerning parallel computers, and help them develop efficient application codes for such systems. It will discuss hardware and software, with an emphasis on systems that are now (or soon will be) commercially available. Program design principles such as load balancing, communication reduction, and efficient use of cache will be illustrated through examples selected from engineering, scientific, and database applications. While there is base material that appears in all of the instantiations of the tutorial, the tutorial will be tailored to the target audience.

M5 Monday, November 17
Full Day 8:30am-5pm

Hot Chips and Hot Interconnects for Hot Computer Systems

Subhash Saini, NASA Ames Research Center
Level: 25% Beginner, 50% Intermediate, 25% Advanced

R10000/12000, which is used in the SGI Origin2000; the HP PA-RISC 8000/8500, which is used in the Convex Exemplar SPP2000; the PowerPC 604/620, which is used in an IBM SMP; the Intel Pentium Pro Processor (P6), which is used in the first DOE ASCI system and in P6-based clusters; the DEC Alpha 21164, which is used in the Cray T3E; a Tera proprietary processor, which is used in Tera Multi-threaded Architecture; a Fujitsu proprietary processor, which is used in the Fujitsu VPP700; an NEC proprietary processor, which is used in NEC SX-4; and a Hitachi proprietary processor, which is used in Hitachi SR2201. The architectures of these processors will first be presented, followed by interconnection networks and a description of high-end computer systems based on these processors and networks. The performance of various hardware/programming model combinations will then be compared, based on latest NAS Parallel Benchmark results (MPI and HPF), thus providing a cross-machine and cross-model comparison. Performance of various networks used in these high-end computing systems including network of workstations and P6-based clusters will also be presented. The tutorial will conclude with discussions of new models of computation, which include parallel, swarm, and quantum computing as well as new physical mechanisms for computation, which includes DNA computing, cellular engineering, and neural networks.

M6 Monday, November 17
Full Day 8:30am-5pm

Building Metacomputing Systems and Metacomputing Applications

Ian Foster, Argonne National Laboratory
Carl Kesselman, USC/ISI
Level: 40% Beginner, 40% Intermediate, 20% Advanced

Recent advances in high-speed networking technology made it possible to develop applications that execute on networked virtual supercomputers, or metacomputers. These applications use local and wide-area networks to combine resources such as massively parallel computers, high performance network-connected mass storage systems, large-scale information bases, and high-end display devices. Thus metacomputing applications must address issues found in both distributed and parallel computing. The goal of this tutorial is to introduce metacomputing and show why it is important, to provide a basic understanding of the problems that face the developers of metacomputing systems and applications, and finally to introduce tools that are currently available to build metacomputing applications. Throughout the tutorial, the instructors will illustrate concepts and techniques with real examples based on the Globus system. Globus software technology was used as the basis for the highly successful I-WAY experiment at the SC 95 conference, and is currently being used to construct an international metacomputing environment.

M7 Monday, November 17
Full Day 8:30am-5pm

A Practical Guide to Java

Ian G. Angus, Boeing Information and Support Services
Level: 50% Beginner, 25% Intermediate, 25% Advanced

In the span of two years Java has risen from an experimental language to be the Internet programming language du jour. The question naturally arises: What is Java and how is it used? This tutorial will go beyond the hype. The instructor will: introduce the Java programming language and explain the concepts (and buzzwords) that define it with careful coverage given to both its strengths and weaknesses; show with simple examples how Java can be used on the WEB and how the participants can take advantage of it; and demonstrate Java's greater potential with a discussion of selected non-WEB applications for which Java provides unique capabilities.

M8 Monday, November 17
Full Day 8:30am-5pm

Implementation Methods for Parallel Adaptive Mesh Refinement Solutions of Partial Differential Equations

Jim Browne, Manish Parashar, University of Texas at Austin
Level: 20% Beginner, 50% Intermediate, 30% Advanced

This tutorial will give the participants an operational familiarity with the principles of and methods for parallel implementation of adaptive mesh refinement algorithms. The focus of the tutorial will be practical. The participants will be organized into teams that will actually carry through development of a parallel implementation of an adaptive algorithm for a simple application. That a parallel implementation of an adaptive solution for even a very simple application can be developed in one day is an important statement of the rapid development of this field. The tutorial will begin with an introduction to adaptive mesh refinement computational methods. It will introduce the HDDA/DAGH (PAR95) infrastructure for parallel implementation of adaptive mesh refinement algorithms, and it will survey some of the other infrastructures for parallel implementation of adaptive mesh refinement algorithms such as Kelp and A. It will give operational familiarity with adaptive methods and the HDDA/DAGH infrastructure. Furthermore, the tutorial will consist of interleaving of lectures and laboratory sessions where participants will construct and execute a parallel adaptive solution for a simple 3D application. The laboratory sessions will be supported by two developers of the methods in addition to the two lecturers.

Questions should be addressed to:

Ann Hayes, SC97 Tutorials Chair
Los Alamos National Laboratory
505-665-4506
tutorials97@mail.supercomp.org

Proceedings Table of Contents

Final Program Table of Contents

SAN JOSE SC97 NOV. 15-21

HPC Challenge

The SC97 HPC Challenge is targeting two problems this year: a race to solve a known, hard cryptographic problem, and a race to identify a complicated scene using ray-tracing methods. Information about participating teams and their exhibit booths will be posted throughout the San Jose McEnery Convention Center.

The Challenge takes place using the networking and computing resources available on the exhibit floor and through the SCinet97 Internet connection. The competition begins on Tuesday at 1:30pm in full view of attendees and ends at noon on Thursday. Awards will be given at the awards session on Thursday at 1:30pm.

The SC97 HPC Challenge features teams competing to successfully address one of two problems: a race to solve a known, hard cryptographic problem, or a race to identify a complicated scene using ray-tracing methods. The competition begins on Tuesday at 1:30pm in full view of attendees and ends at noon on Thursday. Awards will be given at the awards session on Thursday at 1:30pm.

Teams will be comprised of 2-8 participants and an industry or research organization supporter providing consultation and equipment for each team to connect to SCinet97 and the display equipment necessary to provide visual results.

Gold, silver, and bronze medals will be awarded to participating teams. Awards will be considered in the following categories for both challenges:

- the solution using the most systems
- the solution that is the most geographically distributed
- the solution or team that is most creative, innovative, or unique

Additional award categories for the ray-tracing challenge include:

- the fastest rendering that yields acceptable fidelity
- the highest fidelity image produced in a predefined time

An additional award category for the hard cryptographic problem will be:

- the fastest solution

Bell and Fernbach Awards

Winners of this year's Gordon Bell Prizes and Fernbach Award will present talks on their award-winning work during this session. The Gordon Bell Award was established to reward practical use of parallel processors by giving a monetary prize for the best performance improvement in an application. The prize is often given to winners in several categories relating to hardware and software advancement. Entries are coordinated by IEEE Computer and IEEE Parallel and Distributed Technology magazine, publications of the IEEE Computer Society. The Fernbach award honors Sidney Fernbach, one of the pioneers in the development and application of high performance computers for the solution of large computational problems. It is given to someone who has made "an outstanding contribution in the application of high performance computers using innovative approaches."

Thursday, November 20 -- 3:30-5pm

[Room C2] Gordon Bell Prize Winners
Session Chair: Gary Glatzmeier, Los Alamos National Laboratory

Abstracts

- Transient Solid Dynamics Simulations on the Sandia/Intel Teraflop Computer
Stephen A. Attaway Sandia National Laboratories; Edward J. Barragy, Intel Corporation; Kevin H. Brown, David R. Gardner, Bruce A. Hendrickson, Steven J. Plimpton, Courtenay T. Vaughan, Sandia National Laboratories

- High Performance MP Unstructured Finite Element Simulation of Chemically Reacting Flows
Karen D. Devine, Gary L. Hennigan*, Scott A. Hutchinson, Andrew G. Salinger, John N. Shadid, and Ray S. Tuminaro Sandia National Laboratories, Albuquerque, NM *New Mexico State University, Las Cruces, NM

- Large Scale Simulation of Suspensions with PVM
Nhan Phan-Thien, Ka Yan Lee, University of Sydney; David Tullock, Los Alamos National Laboratory

- Pentium Pro Inside: I. A Treecode at 430 Gigaflops on ASCI Red, II. Price/Performance of $50/Mflop on Loki and Hyglac
Michael S. Warren, Los Alamos National Laboratory; John K. Salmon, California Institute of Technology; Donald J. Becker, Goddard Space Flight Center; M. Patrick Goda, Los Alamos National Laboratory; Thomas Sterling, Caltech/JPL; Grégoire S. Winckelmans, Universite Catholique de Louvain

SAN JOSE SC97 NOV. 15-21

Panels and Town Hall Meetings

Participate in SC97! The Panels and Town Hall meetings are designed for attendee involvement. This is your opportunity to ask questions of experts in high performance networking and computing, or to make your voice heard about future SCing conferences. Education panels offer perspectives on the use of HPC in the classroom and beyond, for life-long learning.

Sunday	Monday	Tuesday
Wednesday	Thursday	Friday

Sunday, November 16

3:30-5pm

Education Panel: Collaborative Environments and Technology
Participants: Beau Fly Jones, Ohio Supercomputer Center; Geoffrey Fox, NPAC, Syracuse University; David Emugh, NCSA

This panel will address the issue of the development of collaborative learning environments, the associated technologies, and hidden problems in their use.

Monday, November 17

8:30-10am

Education Panel: Standards, Frameworks, and Technology Planning for Schools
Participants: Ralph Annina, Senior Consultant, Allied Computer Group; Carol Doherty, Director, Professional Development, Center for Innovation in Urban Education

This panel will address the issues surrounding the impact of national math, science, and technology education standards on technology planning for schools.

10:30am-noon

Education Panel: The Future Impact of the National Partnerships in Computational Infrastructure on Education
Participants: Roscoe C. Giles, Boston University; John Ziebarth, NCSA; Kris Stewart, SDSU, NPACI

This panel will focus on the next decade of HPC and its impact on education and society. HPC has had a significant impact in the past; we are now at a transition point developments to be seen over the next 10 years.

3:30-5pm	**Education Panel: Development and Assessment of Virtual Reality Materials for Education** **Participants**: Chris Dede, George Mason University; Robert Fixen, Montana State University; Melissa Kelly, NCSA

This panel will evaluate the application of virtual reality technology in the classroom and other educational settings.

Tuesday, November 18

10:30am-noon	**Panel: Digital Libraries** **Moderator**: Sally Howe, National Coordination Office for Computing, Information, and Communications **Panelists**: Alexa McCray, National Library of Medicine; Terence Smith, University of California at Santa Barbara; Howard Wactlar, Carnegie Mellon University **Room A1**

"Digital libraries" is a rapidly advancing field that will place heavy demands on high performance computing and communications technologies and infrastructure. The panelists will describe the demands emanating from the technologies they are developing to organize, find, and extract knowledge from digital information:

- Terry Smith is principal investigator of the Alexandria Project, one of six NSF/DARPA/NASA Digital Libraries Research Initiative projects. Alexandria allows users to locate maps and other spatially-index material in geographically dispersed digital libraries and databases.

- Howard Wactlar is principal investigator of the Informedia project, another of the six Initiatives. That project is integrating speech, image, and natural language understanding to create and explore digital video libraries.

- Alexa McCray leads NLM in developing advanced technologies for extracting knowledge from biomedical informatics resources including the Library's Unified Medical Language System, Internet Grateful Med, and Visible Human databases.

The panelists will also discuss the potential impact of technologies developed in a possible Digital Libraries II initiative, whose proposed focus will be on developing tools that make the libraries easy to use by their highly diverse user communities and on content.

1:30-3pm	**Panel: High Performance I/O** **Moderator**: Jim Pool, Center for Advanced Computing Research **Panelists**: Paul Messina, Jim Pool, Center for Advanced Computing Research; David Greenberg, Sandia National Laboratories; Reagan Moore, NPACI; Marc Snir, IBM Thomas J. Watson Research Center; Tom Ruwart, University of Minnesota **Room A1**

While transferring data among the components of a computing system has always been a challenge to system architects, continuing increases in processor speed now make high performance I/O essential to achieving potential system performance. Quantitative investigations of I/O-intensive scientific and engineering applications on large scalable parallel systems have demonstrated dramatic imbalances between compute power and I/O, in addition to mismatches between system design and application practice. Increased use of scalable parallel systems for data-intensive applications-both commercial and emerging scientific and engineering applications-introduces additional problems including network I/O and transfers to and from archival storage systems. Processor speeds have generally outpaced the performance of other system components, making sequential approaches to I/O inadequate and parallelism in I/O a necessity. Since I/O involves many hardware and software components, introducing parallel I/O and avoiding possible bottlenecks is a difficult design problem. Moreover, effective use of parallel I/O currently imposes a burden on the application developer since minor changes in parameters or options can have major impacts on I/O performance. Panelists representing large-scale applications, system vendors, and computer science research projects will review requirements for high performance I/O and speculate on possible approaches to pragmatically satisfying these requirements.

1:30-3pm

Education Panel:Presentation of the UCES Awards
Moderator: Jim Corones, AmesLab

Each year, the Department of Energy has sponsored an awards program for excellence in undergraduate computational engineering and science. During this panel each of the 1997 UCES awardees will make a short presentation on his/her work and receive the UCES award. http://uces.ameslab.gov/uces/awards

3:30-5pm

Panel: Building Computational Grids
Moderator: Ian Foster, Argonne National Laboratory
Panelists: Carl Kesselman, USC Information Sciences Institute; Larry Smarr, NCSA; Songnian Zhou, Platform Computing Corporation
Room A1

Developments in networking and software technologies are enabling the integration of geographically distributed compute, storage, network, and other resources into high performance distributed computing systems or computational grids. These computational grids promise to revolutionize high performance computing by enabling both the construction of new tools providing qualitatively new capabilities, and wider access to unique resources such as supercomputers, scientific databases, and instruments. This panel will discuss a variety of issues relating to the construction of computational grids, including: How can computational grids be expected to change science and engineering? What will they look like? What are the major obstacles to creating and using grids? Where are research advances must urgently required? Can we expect the required advances to come from the commercial sector, or do research problems remain? The panelists include Carl Kesselman, co-leader of the Globus research project, which is investigating advanced grid technologies; Larry Smarr, director of the National Computational Science Alliance and a leading advocate of the grid concept; and Songnian Zhou, chief technology officer at Platform Computing Corporation, which develops grid software.

3:30-5pm

Education Panel: Internet in the K-12 Classroom: The Realities of Technology Transfer
Participants: TBA

This panel will deal with specific issues concerning the integration of technology into the curriculum and the classroom.

5-6:30pm

Town Hall Meeting with the Presidential Advisory Committee on HPCC, IT and NGI
Moderator: Ken Kennedy, CRPC, Rice University
Room A1

In February, President Clinton established the Presidential Advisory Committee on HPCC, IT and NGI, headed by Ken Kennedy of the Center for Research on Parallel Computation and Bill Joy, co-founder and current vice-president for research at Sun Microsystems. The committee provides guidance and advice on all areas of high performance computing, communications, and information technologies as they relate to the administration's efforts to accelerate development and adoption of information technologies that will be vital for American prosperity in the 21st century. In this Town Hall meeting, SC97 participants will have an opportunity to ask questions of several representatives of this committee. Attendees' questions submitted prior to the Town Hall meeting will be answered first, followed by an open mike session for questions from the floor. The committee will also be looking for input on their areas of focus and future actions.

Wednesday, November 19

10:30am-noon

Panel: Confidence in Simulation
Moderator: Vic Reis, Department of Energy
Panelists: Tom Adams, Los Alamos National Laboratory; Richard Hodur, Naval Research Laboratory; Greg Shubin, The Boeing Company; Sam Thurman, Jet Propulsion Laboratory
Room A1

This panel will explore the challenges, new opportunities, and limitations of high-fidelity numeric simulation. The panelists will address how simulation results are compared with actual physical phenomena and how results are projected when physical tests are not possible. With the variables imposed by the accuracy of the model, the accuracy of the measurements, and the fundamental physical limitations, how well can we do--and when do we decide that the results are good enough? The software validation for the entry, descent, and landing for the Mars Pathfinder, along with simulations related to ASCI, numerical weather prediction, and aircraft design and manufacture will be discussed.

1:30-3pm

Panel: Breakthroughs and Challenges ahead in Computer Architecture
Moderator: David Culler, University of California, Berkeley
Panelists: James H. Gray, Microsoft; Greg Papadopoulos, Sun Microsystems Computer Company; Burton Smith, Tera Computer; John L. Hennessy, Stanford University
Room A1

In recent years we have seen a technological turnover in high performance computer architecture, followed by increasing levels of commoditization and consolidation in the industry. Meanwhile, we have approached the brink of new technological epochs, such as processor and DRAM on chip, multi-gigabit fiber links, ten thousand processor systems, lithography limits, and practical speed of light constraints. This panel will explore views on where computer architecture is headed, what breakthroughs may lie around the corner, and what are likely to be the critical challenges ahead.

3:30-5pm

Panel: The Future of Software Development Environments for Parallel/Distributed Computing
Moderator: Barton Miller, University of Wisconsin

Panelists: Dennis Gannon, Indiana University; Ken Kennedy, Rice University; Doug Pase, IBM Power Parallel Division; Dan Reed, University of Illinois
Room A1

The parallel and distributed community has suffered from the lack of powerful, ubiquitous programming tools. While there have been some notable successes, on the whole, most programmers find limited choices on most platforms. The question affects compilers, libraries, profilers, debuggers, and visualizers. The panel will present their views on this issue, addressing questions that include: What is your vision for the route to ubiquitous quality tools? What platforms should we be targeting? What is the future of UNIX vs. Windows/NT vs. custom systems? Is portability important or even possible? Is there a best language or programming model? Is there a commodity programming language or programming model? The panelists include Dennis Gannon, leader of the HPC++/SAGE++ project; Ken Kennedy, director of the Center for Research on Parallel Computation (CRPC); Doug Pase, Tools Architect for IBM RS/6000 Power Parallel Division; and Dan Reed, leader of the Pablo project.

5-6:30pm

Town Hall Meeting on the Internet: Then, Now, and Tomorrow
Moderators: Dan Lynch, Founder of Interop; Katie Hafner, Newsweek, Silicon Valley Correspondent
Invited Participants: Irving Wladawsky-Berger, IBM; Vint Cerf, MCI; Craig Partridge, GTE/BBN; Don Nielson, SRI International; Jordan Becker, AOL/ANS; Steve Wolff, Cisco; Al Weis, Advanced Network Services; Brian Reid, Digital Equipment Corporation; Bill Joy, Sun Microsystems; Bob Kahn, Corporation for National Research Initiatives; Larry Roberts, ATM Systems; Paul Mockapetris, Software.Com; George Strawn, National Science Foundation; Len Kleinrock, University of California, Los Angeles; Larry Smarr, NCSA; Van Jacobson, Lawrence Berkeley National Laboratory; Dave Clark, Massachusetts Institute of Technology; Larry Landweber, University of Wisconsin; Dave Farber, University of Pennsylvania
Room A1

Spend an hour with any or all of the founders of the Internet. A brief presentation by four of the originals will be followed by a free-ranging open microphone session for questions, reminiscences, and dreams of the future.

Thursday, November 20

10:30am-noon

Panel: Networking Initiatives: Internet 2 and Next Generation Internet
Moderator: Larry Landweber, University of Wisconsin
Panelists: David J. Farber, University of Pennsylvania; George Strawn, National Science Foundation; Jim Gray, Microsoft; Steve Wolff, Cisco Systems
Room A1

Internet 2 (I2), a project of more than 100 U.S. research universities, and Next Generation Internet (NGI), a multi-agency program of the U.S. government, are separate but related initiatives. Their goals include the development of new network-based applications and the prototyping of networks that incorporate new pre-competitive network technologies to support the different service levels required by these applications. NGI also includes a research component that will look further into the future. In addition, to support the various components, both NGI and I2 have as a goal the connection of U.S. research organizations to backbones at significantly higher data rates than is now the case.

10:30am-noon

Town Hall Meeting on SC Conferences
Moderator: Cherri Pancake, Oregon State University

On Thursday, November 20, from 5-6:30pm, attendees will have the chance to interact with committee members from SC97, SC98, SC99, and SCXY (the SC conference series' steering committee), discuss this year's conference, and make suggestions for future years' conferences. This is also the ideal opportunity for attendees to express their interest in serving on one of the conference's diverse committees. Plan to attend and make your voice heard.

3:30-5pm

Panel: Visualization
Moderator: Paul Woodward, University of Minnesota
Panelists: Pat Hanrahan, Stanford University; Paul Woodward, University of Minnesota; Chris Johnson, University of Utah; Tom DeFanti, University of Illinois at Chicago; Philip Heermann, Sandia National Labs
Room A1

DoE's ASCI program, NSF's PACI program, the DoD HPC modernization program, and NASA's grand challenge program are accelerating the pace of supercomputer hardware development. As a result, researchers are running much larger simulations and are therefore generating much larger data sets for visualization. This panel will discuss the present bottlenecks that are limiting the rate at which visualization systems and software can accelerate in order to keep pace with the new, faster growth in demand for data visualization.

Friday, November 21

8:30-10am

Panel: PACI
Moderator: Bob Borchers, National Science Foundation
Panelists: Sid Karin, National Partnership for Advanced Computational Infrastructure; Peter Kollman, University of California, San Francisco; Larry Smarr, National Computational Science Alliance; Bob Sugar, University of California, Santa Barbara
Room A1

In 1985 the National Science Foundation (NSF) established a group of supercomputing centers to further research using computational methods. From 1985 to 1997, partly through the contributions of these centers, the US saw computation become a widely used tool in many disciplines, an explosive growth of the Internet, the birth of the Web, and tremendous increases in computing power. In 1995, in response to the changing technological landscape, the NSF initiated a competition to establish Partnerships for Advanced Computational Infrastructure (PACI). With PACI, starting in October of 1997, the focus will shift from centralized facilities to nationwide partnerships, reflecting the growing importance to computational science of computer science, high-speed networking, distributed computing, and remote collaboration. Two partnerships, the National Computational Science Alliance led by the University of Illinois at Urbana-Champaign and the National Partnership for Advanced Computational Infrastructure led by the University of California, San Diego, have joined dozens of research organizations across the country in large-scale efforts to improve the technological environment available to computational scientists. This panel will outline the current status of the PACI program, its resources, and future plans from both the perspectives of the principal investigators and those of scientists whose research depends on the use of computing resources.

8:30-10am

Panel: I/O and Interconnects
Moderator: Bill Boas, Abba Technologies
Panelists: Greg Chesson, Silicon Graphics; Dave Follett, Giganet; Randy Rettberg, Sun Microsystems; Justin Rattner, Intel; Chuck Seitz, Myricom
Room A3

This panel will be a survey of the current technologies and architectures in use or in development today for both internal SMP interconnects as well as for clustering multiple processors in shared and non-shared memory architectures. The discussion will focus on design trade-offs, limiting bandwidth factors, latency, applications and operating system interfaces, driver structures, media technologies, parallel vs. serial transmission, ASIC densities and limitations, speed, and number of gates. Technologies discussed will include Craylink, HiPPI-6400, SCI, Gigaplane, Myrinet, ATM, PCI evolution, and the emerging Virtual Interface Architecture.

10:30am-noon

Panel: ASCI Strategic Alliances
Moderator: Alex Larzelere, Department of Energy
Panelists: Michael T. Heath, University of Illinois;Daniel Meiron, California Institute of Technology;David Pershing, University of Utah; William Reynolds, Stanford University;Robert Rosner, University of Chicago
Room A1

In 1995 the U.S. Department of Energy (DOE) established the Accelerated Strategic Computing Initiative (ASCI). Its purpose is to bring together the resources of three DOE Defense Program National Laboratories (Los Alamos, Lawrence Livermore, and Sandia National Laboratories) in a consolidated effort to achieve significant advances in computational modeling and simulation capabilities available to meet national goals for nuclear stockpile stewardship. One of the five key ASCI strategies is to "Encourage Strategic Alliances and Collaborations." The ASCI Academic Alliances strategy was established to enhance the overall ASCI goals by establishing technical interactions between the ASCI Laboratories and leading-edge academic R&D in the United States. These alliances will take the form of Centers of Excellence working on key multi-disciplinary applications and high performance computing requiring major advances in large-scale computer-based simulation approaches. This panel will discuss the major goals and objectives of the Academic Alliances program, including the establishment and validation of the scientific methods of large-scale modeling and using simulation and computation as a cornerstone of the scientific methodology in SBSS-related applications requiring coupled complex simulation sequences.

10:30am-noon

Panel: The Role of High Performance Computing and Informatics in Pharmaceutical Research and Development
Moderator: Frederick H. Hausheer, BioNumerik Pharmaceuticals, Inc.
Panelists: Frederick H. Hausheer, BioNumerik Pharmaceuticals, Inc.; Babu Venkataraghavan, American Cyanamid Company; Eugene Fluder, Merck & Co.
Room A3

This panel will discuss current applications of computing and informational technologies being applied in the pharmaceutical industry to help discover and develop new drugs in less time. It is clear that new therapeutic targets will be discovered by the identification of genomic sequencing. The panel will discuss how high performance computing and information technologies may play a role in the translation of potential new targets into new drugs. The panelists are recognized industry experts with extensive experience in pharmaceutical discovery and development and the application of high performance computing and information technologies. Frederick H. Hausheer, is the CEO of BioNumerik Pharmaceuticals, Inc.; Babu Venkataraghavan, is director of Biomedical Research Computing at American Cyanamid Company, Medical Research Division, Lederle Laboratories; Eugene Fluder is senior investigator, Molecular Design and Diversity, at Merck.

SAN JOSE SC97 NOV. 15-21

Industry Exhibits

Hosting SC97 in Silicon Valley, with its more than 3,600 technology companies, has drawn new exhibitors and interest this year. In recognition of the proximity and prominence of industry to this year's conference, exhibit hours have been increased:

Expanded Exhibits Hours

Monday	November 17 7-9pm	(Gala Opening--All Exhibits)
Tuesday	November 18 10am-6pm	(All Exhibits)
Wednesday	November 19 10am-6:30pm	(includes 5-6:30pm posters reception)
Thursday	November 20 10am-4pm	(All Exhibits)

Don't miss the exhibits or the Exhibitor Forum!

Exhibitor	Location
Active Tools	327
ACM SIGARCH	Lobby
Advanced Network Applications	753
Avalon Computer Systems	138
CISCO Systems	1237
Compaq Computer	135
Computer Network Technology	445
Computers in Physics/American Institute of Physics	836
Digex, Inc.	843
Digital Equipment Corp.	707

Dolphin Interconnect Solutions	446
DPC Technologies	930
DRS Precision Echo, Inc.	242
Elsevier Science	441
Essential Communications	1035
FORE Systems	938
Fujitsu	907
Genroco	1134
Gigabit Ethernet Alliance	1136
GigaLabs, Inc.	745
Hewlett-Packard/Convex	117
Hitachi, Ltd.	921
HNF	1135
HPC Asia '98	134
HPCC Week	931
HPCwire	331
HPSS-Oak Ridge National Lab	835
IBM	719
IEEE Computer Society	Lobby
Impactdata	447
Instrumental, Inc.	343
IPT Corporation	329
Kingston Technology Company	1133
Madge Networks	132
MAXSTRAT	728
MCI	1336
The MIT Press	1228
Morgan Kaufman Publishers	1029
MPI Software Technology, Inc.	142
MRJ Technology Solutions	841
Mutech Solutions	948
National Transparent Optical Network	468
NEC/HNSX Supercomputers, Inc.	217

Nichols Research Corporation	228
Northrop Grumman Corp.	935
Numerical Algorithms Group, Inc.	743
On-Line Power	234
Pacific-Sierra Research	929
Pallas GmbH	834
Platform Computing	129
The Portland Group	421
Raytheon/E-Systems	336
RCI, Ltd.	1128
SCALI	1231
SC97 History of the Internet Exhibit	1048
Sage Science Press	1132
Scientific Computing Associates	1129
Scientific Computing & Automation	341
SCinet97	1143
Sequent Computer Systems, Inc.	442
SIAM	Lobby
Silicon Graphics/Cray Research	417
Silicon Graphics Open MP	420
Sony Electronics, Inc.	829
Standard Microsystems Corporation	1131
Storage 2000, Inc.	1031
StorageTek/Network Systems Corp.	335
Sun Microsystems	1107
Tera Computer Company	1229
Texas Memory Systems	118
Unitree Software Inc.	233
Visual Numerics/Applied Parallel Research	229
YEM	131

Keynote Address, Invited Speakers, and State-of-the-field Talks

SC97 offers several presentations by invited speakers, starting with Paul Saffo keynote. Other invited speakers include leaders and pioneers in networking, astrophysics, climate modeling, medical libraries, financial data mining, and entertainment. And new this year, the State-of-the-field Talks are designed to provide attendees with a valuable source of technical insight that would be nearly impossible to find anywhere else.

TUESDAY, NOVEMBER 18
8:30-10AM

Keynote Address

San Jose Civic Auditorium

Paul Saffo is a director at the Institute for the Future, where his work focuses on the long-term social and commercial impacts of new information technologies. Saffo also writes on emerging technology issues and their long-term implications. He has contributed an occasional column to Wired, and his essays have appeared in a variety of publications, including The Harvard Business Review, The Los Angeles Times, The New York Times, and Fortune. A book of his essays, Dreams in Silicon Valley, is available in Japan. Saffo holds degrees from Harvard College, Cambridge University, and Stanford University. He was recently named by the World Economic Forum as one of the Forum's 100 1997 "Global Leaders for Tomorrow."

TUESDAY, NOVEMBER 18
1:30-3PM

The Role of Supercomputers and Advanced Networking in Solving Upcoming Computer Science Challenges

Session Chair: Greg Papadopoulos, Sun Microsystems
Herbert Schorr, University of Southern California

As we move forward, the focus of computer science changes. Many problems of the past are no longer frontier problems; however, many tend to resurface in new forms. For example, as we delve into sub-micron technologies, the assumption of a separation between logical and physical design may no longer hold. New approaches to high-level design are required if anything other than consumer-oriented microprocessors will be affordable in the future. Other areas are being driven by new research on the Internet and in multimedia. For example, digital libraries bring the subject of information retrieval once more in the foreground. Information visualization is perhaps the next frontier that builds on the success of scientific visualization.

Metacomputing and middleware are the next areas of focus for operating systems research. Meanwhile, after many years of effort, slow progress is finally being made on parallel compiler technology. For the Internet itself, the questions of electronic commerce, the relationship to ATM, and in general voice and wireless remain open questions. In artificial intelligence, agents, very large knowledge bases, data mining, and better human computer interfaces, including natural language, appear to be the next frontiers while software engineering (with a few notable exceptions) seems to be searching for new directions. This talk will explore some of these topics and the role or non-role of supercomputers and networking in these areas.

Herbert Schorr has been the executive director of the University of Southern California Information Sciences Institute since 1988. He graduated from City University of New York and received his Ph.D. in electrical engineering from Princeton in 1963. He was an instructor of electrical engineering at Princeton during the 1961 academic year. During 1962-63 he was an NSF Postdoctoral Fellow at Cambridge University, Cambridge, England.

Upon returning from England, he became an assistant professor at Columbia University for the year prior to joining IBM. Schorr joined IBM as a research staff member at Yorktown in 1965. In 1966 he became manager of systems architecture and programming for the Advanced Computing System project in Menlo Park, California, and later served on the Corporate Technical Committee in Armonk. In 1968 he rejoined the Research Division as director of computer sciences and in 1973 was named vice president, product and service planning for the Advanced Systems Development Division of IBM. In August 1977 he again rejoined the Research Division as department manager of systems technology, now called the Systems Laboratory, and in November 1981, he was named vice president, systems, Research Division. In October 1984 he was named group director, Products and Technology, IS&SG. In 1987 he was named group director, Advanced Systems, Information Storage and Systems Group and was responsible for the introduction of new, advanced technology. In particular, he was the artificial intelligence and image champion for the IBM Corporation and managed project offices in both areas responsible for products, marketing, and internal applications of these technologies.

A System Perspective of Supercomputing: Software Engineering, Security, Reliability, and Risk Avoidance

Session Chair: Greg Papadopoulos, Sun Microsystems
Peter G. Neumann, SRI International

The supercomputer field has in the past shown a tendency to resist becoming a part of the mainstream. Old-timers may recall such things as overdependence on special-purpose hardware, avoidance of modern software technology, use of assembly languages for `"efficiency," and avoidance of mainstream operating systems. I have even detected some unwillingness to believe that supercomputing is integrally related to the rest of the computer field and that as such it could take good advantage of advances in computer science. However, I am now encouraged to believe that some of these tendencies have been diminishing in recent years-perhaps in response to economic realities and some lowest-common denominator effects of the marketplace, perhaps because of the influx of some younger or more practical minds. Nevertheless, I have no desire to flog a straw herring in the foot by dwelling on the past.

I have long held my own set of peeves with the computer field as a whole for unduly ignoring security, reliability, fault tolerance, and system survivability-especially in distributed systems and networks-as well as interoperability and easy evolvability, often traded off in the belief that those attributes are incompatible with performance. Based on the years of involvement with risks, I also have many strong views with respect to risk management and the avoidance of problems such as system failures and malicious misuse. I will consider some of these views, as they relate to supercomputing.

Peter G. Neumann received AB, SM, and Ph.D. degrees from Harvard in 1954, 1955, and 1961, respectively. In 1960 he received a Dr rerum naturarum from the Technische Hochschule, Darmstadt, Germany, where he was a Fulbright scholar for two years.

In the Computer Science Lab at Bell Labs throughout the 1960s, he was involved in research in computers and communications; during 1965-69, he participated extensively in the design, development, and management of Multics, jointly with MIT and Honeywell. He was a visiting Mackay Lecturer at Stanford in 1964 and Berkeley in 1970-71.

At SRI since 1971 (where he is now a principal scientist), he has been concerned with computer systems having requirements for security, reliability, human safety, and high assurance. He served on the NRC study (Cryptography's Role In Securing the Information Society, http://www2.nas.edu/cstbweb).

For the ACM, he was founder and editor of the SIGSOFT Software Engineering Notes (1976-1993) and is now associate editor for the RISKS material; Chairman of the ACM Committee on Computers and Public Policy (since 1985); and a Contributing Editor for CACM (since 1990) for the monthly 'Inside Risks' column. In 1985 he created, and still moderates, the ACM Forum on Risks to the Public in the Use of Computers and Related Technology, which is one of the most widely read of the online computer newsgroups. His RISKS-derived book (Computer-Related Risks, Addison-Wesley, 1995) explores the benefits and pitfalls of computer-communication technology and suggests ways of avoiding risks in critical systems.
http://www.csl.sri.com/neumann.html; Neumann@CSL.sri.com

TUESDAY, NOVEMBER 18
3:30-5PM

Building and Testing a Regional Telemedicine Testbed for the Pacific Northwest: from Bench to Bedside and Beyond

Session Chair: Sally Howe, National Coordination Office for CIC
Sherrilynne Fuller, University of Washington, School of Medicine

The University of Washington Academic Medical Center has undertaken the creation of a regional telemedicine testbed across five states (Washington, Wyoming, Alaska, Montana and Idaho) to evaluate the current and future impact of the National Information Infrastructure on healthcare. Telediagnosis, teleconsulting and telecollaboration elements of the testbed will allow distant providers to "see" patients and each other, and enable clinicians at sites across the region to better communicate, collaborate, and gain access to UW Academic Medical Center specialists, knowledge resources, and services. Intensive development of infrastructure components is currently underway.

The presentation will focus on design features, current status, and challenges of critical infrastructure components including: authentication, encryption, authorization and physical linkages; uniform graphical user interface to bring together, from the caregiver's perspective, required databases including the electronic medical record, databases, decision support tools, and services; database architecture to support transformation of basic genetics research data into useful clinical decision support tools; secure clinical electronic mail for enabling communications among providers and with their patients; and image management. Our ultimate vision includes creation of a comprehensive regional integrated medical network providing easy and secure access to quality, cost-effective healthcare for citizens of the Pacific Northwest an Academic Medical Center without Walls where research findings are transformed rapidly into improvements in clinical care and public health.

Sherrilynne Fuller serves as acting director, Informatics, School of Medicine; coordinator, Health Sciences Information Systems Integration, Health Sciences Center; and director, Health Sciences Libraries and Information Center, University of Washington. Her other responsibilities at UW include: director, National Network of Libraries of Medicine, Pacific Northwest Region, and assistant director of libraries. She is associate professor, Department of Medical Education and Adjunct Associate Professor, Graduate School of Library and Information Science, and Department of Health Services, School of Public Health and Community Medicine. Fuller has a BA degree in biology and a master's in library science from Indiana University and a Ph.D. from the University of Southern California in library and information science. Her areas of research include integrated health sciences information systems design and evaluation. She is the principal investigator of several large health information research projects including UW Health Sciences Center's Integrated Advanced Information Management Systems program; From Bench to Bedside and Beyond, to create and evaluate a regional telemedicine testbed; and a collaborator in the creation of clinical genetics knowledge bases to transform basic genetic research data in useful clinical data to respond to the information needs of healthcare professionals.

Fuller is a member of the Board of Regents of the National Library of Medicine, the Board of Directors of the American Medical Informatics Association, and the President's Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet.

Multinational Corporate Metacomputing

Session Chair: Sally Howe, National Coordination Office for CIC
Colin Upstill, Parallel Applications Center, University of Southampton

Successful business innovation results from the appropriate combination of information technologies, business processes, and human resources. Internet-based technologies can be exploited to enable businesses to create, collate, and share information in ways previously not possible, and as a result to define new business practices. High performance computing enables engineers and scientists to simulate large, complex systems more cost-effectively than ever before. The time for business-critical calculations is reduced from days, weeks or months to overnight, over lunch, or even over coffee. By using geographically distributed metacomputing and Internet technology, massive computing resources can be harnessed to solve the most demanding problems. Advanced distributed multimedia information systems and software agent technologies can support the management and mining of diverse sources of data in large, geographically distributed corporations. There have now been a significant number of cluster and meta-computing application projects and experiments.

Intelligent resource management has enabled a pan-European aerospace consortium to use 375 processors in four countries to perform sensitivity analyses using a wide range of engineering software. Novel intranet architectures have enabled deployment of advanced multimedia and information retrieval technologies to deliver 'information on demand' in the chemical and pharmaceutical sector. These examples demonstrate that it is now feasible to use non-dedicated distributed computing resources on a large scale to solve otherwise intractable problems without huge IT investments.

Colin Upstill is the director of the Parallel Applications Centre (PAC) in Southampton, UK, one of Europe's leading advanced IT technology transfer organizations. He has more than 20 years experience of high performance computing. Until the mid-1980s, he undertook wide-ranging research in computational physics and mathematics at the Universities of Bristol and Cambridge. He then moved to R&D in industry, and from 1987 to 1991 was chief engineer at Plessey Research, Roke Manor, one of the UK's foremost electronic systems and software laboratories. He has been director of the PAC since its foundation in 1991.
http://www.pac.soton.ac.uk/

WEDNESDAY, NOVEMBER 19
1:30-3PM

Pflops, Box Office Hits, and the Human Singularity-Will we Remain Human Long Enough to Collect Social Security?

Session Chair: Margaret Simmons, NPACI, University of California, San Diego
David Brin, Astrophysicist and Science-fiction Author

David Brin has a triple career as scientist, public speaker, and author. His novels, such as Startide Rising, have made the New York Times Bestseller List and have won multiple Hugo and Nebula awards. As a lecturer, Brin is in demand to give talks about various aspects of technology and the human future. His 1989 novel, EARTH, is widely credited as one of the best previews of the worldwide data network of tomorrow, as well as exploring the shape of environmental problems in the decades ahead. Another novel, the Postman, is being made into a film by Warner Brothers Studios, with Kevin Costner both directing and starring in the lead role.

As a scientist, Brin was a fellow at the California Space Institute and later a research affiliate at the Jet Propulsion Laboratory. He also participated in interdisciplinary activities at the UCLA Center for the Study of Evolution and the Origin of Life and has published papers in fields from optics to sociobiology. He now lives in San Diego County with his wife, three small children, and about a hundred very demanding trees.

TBA

Session Chair: Margaret Simmons, NPACI, University of California, San Diego
Invited speaker TBA-information to be posted on SC97 Updates page

WEDNESDAY, NOVEMBER 19
3:30-5PM

Mapping (and Computing) the Universe

Session Chair: David Tennenhouse, MIT Laboratory for Computer Science
Margaret Joan Geller, Harvard-Smithsonian Center for Astrophysics

During the last 15 years, surveys of the nearby universe have revealed a rich tapestry of structure on scales as large as several hundred million light years. Computational n-body models reproduce some of the salient features of these data, but fundamental questions about the origin and evolution of the patterns remain. Some of these questions may be answered as large telescopes on the ground come into operation: these instruments along with the Hubble Space Telescope enable direct observation of most of the history of the universe.

Margaret Joan Geller is an astrophysicist who studies the spatial distribution of galaxies; these studies give us a new understanding of the structure of the universe. Geller has also made two award-winning films about her work: Where the Galaxies Are and So Many Galaxies... So Little Time.

Geller is Professor of Astronomy at Harvard University and Senior Scientist at the Smithsonian Astrophysical Observatory. She is a member of the National Academy of Sciences and of the American Academy of Arts and Sciences. She has received a number honors including a MacArthur Fellowship (1990-1995), the Newcomb-Cleveland Prize of the American Association for the Advancement of Science (1990), the Hogg Lectureship of the Royal Astronomical Society of Canada (1993), the Klopsteg Award of the American Association of Physics Teachers (1996) and the Bethe Lectureship at Cornell University (1996).

Operational Prediction of Thunderstorms: Turning Vision into Reality with Massively Parallel Processors

Session Chair: David Tennenhouse, MIT Laboratory for Computer Science
Kelvin K. Droegemeier, School of Meteorology and Center for Analysis and Prediction of Storms, University of Oklahoma

The Center for Analysis and Prediction of Storms (CAPS), an NSF Science and Technology Center at the University of Oklahoma, is redefining the notion of local weather forecasts by developing techniques for the numerical prediction of individual thunderstorms and their related weather. In this presentation, I describe the scientific and technological challenges associated with this effort, and focus on the program's key element--a new scalable parallel weather model, called the Advanced Regional Prediction System (ARPS), that has been designed specifically for broad classes of parallel architectures. The ARPS utilizes new coding and documentation strategies to ensure simplicity of use, readability, rapid interfacing with other physical models, and quick optimization. In 1997, these achievements received national recognition when CAPS was awarded both the Discover Magazine Award for Technology Innovation and the Computerworld Smithsonian Award.

To demonstrate the practical capabilities of the ARPS, I present results from recent operational tests in which the forecast system was run in real time, on a daily basis, on massively parallel computers in collaboration with the National Weather Service and NOAA Storm Prediction Center. Additionally, I discuss a new three-year research and development project between CAPS and American Airlines that seeks to transfer this emerging technology to commercial aviation.

Kelvin K. Droegemeier received a B.S. with special distinction in meteorology in 1980 from the University of Oklahoma, and M.S. and Ph.D. degrees in atmospheric science in 1982 and 1985, respectively, from the University of Illinois at Urbana-Champaign under the direction of R. Wilhelmson. He joined the University of Oklahoma in September, 1985. He was co-founder of the NSF Science and Technology Center for Analysis and Prediction of Storms (CAPS), and became director on July 1,1994. In 1992, Droegemeier spent a six-month sabbatical as a senior research fellow at the Army High Performance Computing Center, University of Minnesota. As director of the CAPS model development project for five years, he managed the creation of a multi-scale numerical prediction system that is envisioned as national prototype for operational implementation early in the next century. This computer model was a finalist for the 1993 National Gordon Bell Prize in High Performance Computing. In 1997, he was given the Discover Magazine Award for Technology Innovation (computer software category), and also in 1997 CAPS won the Computerworld Smithsonian Award (science category).

Drogemeier has been a major force behind the development and application of high performance computing systems both at OU and across the US. In 1995 he created as principal investigator, and now directs, the $1.4 million NSF/OU Environmental Computing Applications System. Droegemeier's research interests lie in thunderstorm dynamics and predictability, variational data assimilation, mesoscale dynamics, computational fluid dynamics, massively parallel computing, and aviation weather.

THURSDAY, NOVEMBER 20
1:30-3PM

Awards Plenary and Invited Speaker

Session Chairs: Margaret Simmons, NPACI, University of California at San Diego; Greg Papadopoulos, Sun Microsystems

The award session will consist of two parts. At this session, the winners of the Best Paper Award, Best Student Paper Award, the Sid Fernbach Award, the Gordon Bell Award, and the HPC Challenge awards will be announced. The session will also include an invited talk by Paul Saffo of the Institute of the Future. In the second session, Thursday from 3:30-5pm, winners of this year's Gordon Bell Prizes and Fernbach Award will present talks on their award-winning work.

TBA

Vic Reis, Department of Energy

THURSDAY, NOVEMBER 20
3:30-5PM

Parallel Large Eddy Simulation Technique for Industrial Applications

Session Chair: Phil Papadopoulos, Oak Ridge National Laboratory

Chisachi Kato, Mechanical Engineering Research Laboratory, Japan

This paper describes massively parallel LES (Large Eddy Simulation) technique for viscous incompressible flow applied to several industrial problems. The scheme combines finite elements and moving boundaries, with parallelization based on domain decomposition and the RGB (Recursive Graph Bisection) algorithm. Estimates of parallel computing performance show good scalability on the Hitachi SR2201. The numerical examples for the turbulent wake of a circular cylinder and a pantograph cover for the high speed train show good agreement with the experimental data, and allow prediction of aerodynamic sound. Another example for the flow in a pump with the impeller-casing interaction shows good agreement with experimental results for the flow-induced forces on the impeller. We will also show a computer-generated video of the application of this to Japan's Shinkansen (bullet train).

TBA

Session Chair: TBA
Invited speaker TBA-information to be posted on SC97 Updates page

Proceedings Table of Contents

Final Program Table of Contents

Bell and Fernbach Awards

Winners of this year's Gordon Bell Prizes and Fernbach Award will present talks on their award-winning work during this session. The Gordon Bell Award was established to reward practical use of parallel processors by giving a monetary prize for the best performance improvement in an application. The prize is often given to winners in several categories relating to hardware and software advancement. Entries are coordinated by IEEE Computer and IEEE Parallel and Distributed Technology magazine, publications of the IEEE Computer Society. The Fernbach award honors Sidney Fernbach, one of the pioneers in the development and application of high performance computers for the solution of large computational problems. It is given to someone who has made "an outstanding contribution in the application of high performance computers using innovative approaches."

Thursday, November 20 -- 3:30-5pm

[Room C2] Gordon Bell Prize Winners and Fernbach Award
Session Chair: Gary Glatzmeier, Los Alamos National Laboratory

HIGH PERFORMANCE COMPUTATIONAL ENGINEERING: A PERSONAL PERSPECTIVE

Charbel Farhat
University of Colorado at Boulder
Department of Aerospace Engineering Sciences
Campus Box 429
Boulder, CO 80309-0429
charbel@boulder.colorado.edu

Extended Abstract:

Have you ever wondered why the F117 Stealth fighter has flat, angular surfaces and looks like a flying diamond, while the B-2 Stealth bomber has aerodynamic rounded surfaces? Simply because the F117 fighter was designed around 1975, when computers did not have a sufficiently large memory and a

powerful CPU to allow three-dimensional designs, or rounded shapes, and therefore when the radar cross section of an aircraft configuration could be accurately calculated only if it was in two dimensions. On the other hand, the B-2 bomber was designed later using three-dimensional computations on supercomputers capable of hundreds of megaflops [1]. This is perhaps the most exciting example of the impact of high performance computing on engineering design, but is certainly not the only one. After the unfortunate Challenger incident, thousands of CRAY X-MP CPU hours have contributed to the redesign of the Solid Rocket Booster [2]. Today, high performance computing platforms are very visible in the automotive industry and play an important role in structural integrity, noise, vibration, and harshness studies.

Nevertheless, the use of computers in engineering applications has traditionally lagged hardware advances. This has been so at the two extremes of the spectrum: microcomputers and vector supercomputers. Software tools have to be provided and their functionality accepted by the user community at large before an application base emerges. Vector supercomputers made their debut in government laboratories devoted to complex physics calculations. Microcomputers controlled laboratory experiments, word processors, and games. Today, high performance workstations and vector supercomputers are routinely used by large industrial corporations, and personal computers are used by thousands of engineers for tasks that challenged a mainframe two decades ago. On the other hand, parallel processors are still thought of as exotic machines. Currently, only a fraction of the engineering community is engaged in applying the parallel computing technology. Incorporation of these new machines in the mainstream of large scale computations has proved to be more challenging than that of vector supercomputers, but will eventually have a greater impact on computational engineering power. However, moving engineering applications to parallel processors continues to face significant obstacles that center on hardware architecture, numerical methods, (sometimes languages), and education. During the last twelve years, the author has dedicated a significant amount of his energy to address some of these issues, in the context of finite element simulation based engineering designs.

Back in 1984, the first-generation hypercube parallel processors were essentially devoid of software. Partitioning a mesh into pieces to be loaded in every processor seemed inevitable and lead the author to develop perhaps one of the first and simplest mesh decomposition algorithms [3],[4]. This algorithm, known today as the Greedy algorithm, performs surprisingly well even by computer scientists standards. Currently, the impact of mesh partitioning on computational science and engineering is reflected in the increased number of publications and the impressive number of technical sessions devoted exclusively to this topic in most of the conferences on numerical analysis, parallel processing, and computational engineering. Unfortunately, most of the effort in this area seems to focus on mesh partition interface sizes, interprocessor communication costs, and other so-called "quality" criteria that have little if any relevance to real applications, whereas the issues with the greatest impact on parallel engineering computations are more complex and subtle, and call for a two-step partitioning paradigm [5],[6]. However, the good news is that software packages such as TOP/DOMDEC [7],[8] from the University of Colorado,

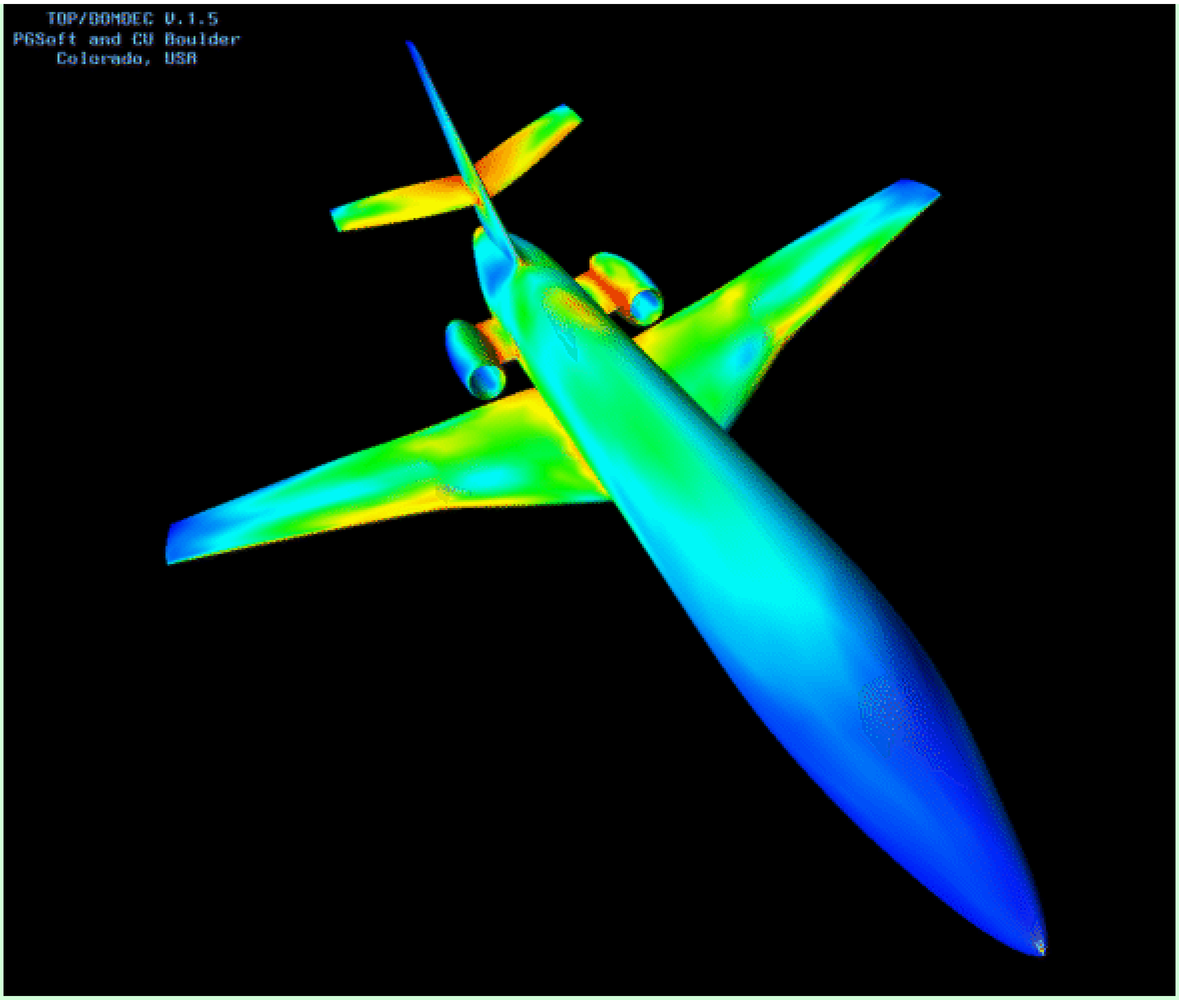

CHACO [9] from the Sandia National Laboratories, and METIS [10] from the University of Minnesota have been and are now available to assist the engineering community in embracing the parallel processing technology. As a matter of fact, these packages are now routinely used in several aerospace, automotive, computer, and third-party software industries, and in many research facilities and government laboratories.

Solving a large-scale system of equations is another important issue for many engineering simulations. Today, at some engine companies, there is a pressing need for solving structural problems with as many as five million degrees of freedom on a daily basis. This is rather impressive given that such problem sizes used to be encountered in aerospace and mechanical engineering contexts mostly, if not exclusively, in flow simulations. The corresponding memory and CPU requirements are such that many commercial finite element software houses have already included, and others are currently considering including, iterative solvers in their packages, which is a noticeable step given that breaking the monopoly of sparse direct solvers in such finite element structural codes was unthinkable a decade ago. In this area, domain decomposition based iterative solvers are attractive because, in addition to their potential for combining the best of the direct and iterative worlds, they are ideal for high performance parallel computing. The family of FETI methods [11--16] developed by the author and his colleagues fit into this category of solvers and have gained wide acceptance in finite element production (and a few commercial) codes, particularly because they are both numerically and parallel-wise scalable. Their convergence is nearly independent of the problem size and the number of subdomains, and their performance on most parallel hardware architectures has delivered close to linear speedups. Today, the impact of FETI can be measured, among others, by the large number of papers proposing variants of this methodology, by its worldwide acceptance as a fast parallel solver [17], and by its selection for the Salinas ASCI project at the Sandia National Laboratories.

For many people, the words Grand Challenge, which unfortunately have become out-of-fashion today, have a "science" connotation. Nevertheless, engineering grand challenges exist, and high performance computing has allowed breaking grounds for several of them. Indeed, the investigation of many engineering grand challenge problems incurs multidisciplinary numerical simulations targeted at improving the understanding and prediction of interaction phenomena such as fluid/structure/thermal coupled processes. The storage and computational requirements of such simulations are beyond conventional computers. During the last five years, under the sponsorship of the National Science Foundation, we have tackled at the University of Colorado one such Grand Challenge problem: the nonlinear aeroelastic simulation of a complete aircraft.

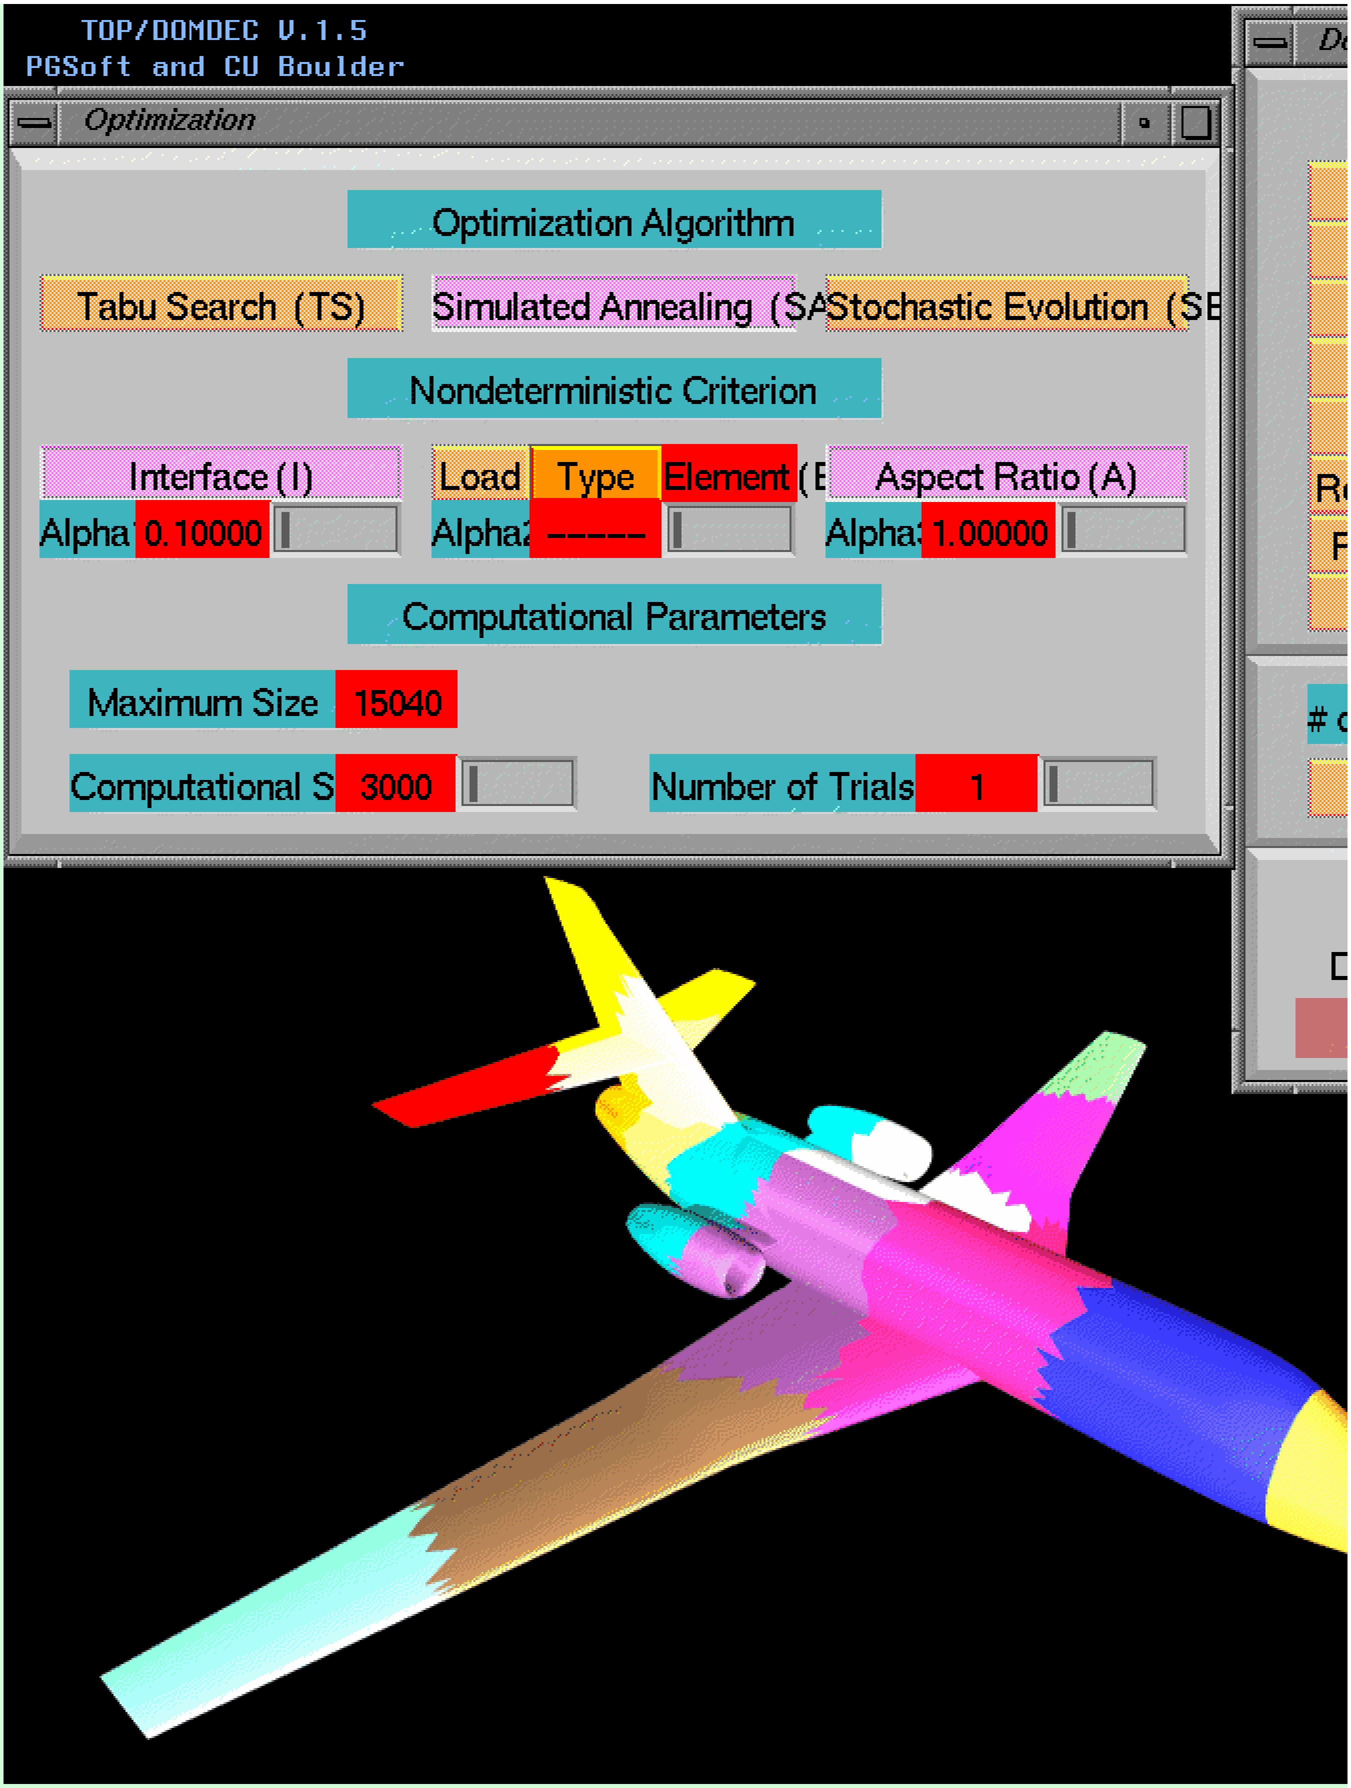

Under the sponsorship of the Air Force Office of Scientific Research we are currently working on the full flight simulation of an aircraft --- that is, on integrating structures, fluids, and controls. High performance computing may not be a sufficient condition for succeeding in this endeavor, but it is certainly a necessary one.

Predicting the future of high performance computing is risky, and goes beyond the scope of this abstract. What is certain is that working in this area has been intellectually rewarding and a pleasure.

References

1. B. R. Rich and L. Janos, **Skunk Works**, Little, Brown and Company.

2. N. G. Christensen, **Supercomputing gives a boost to the shuttle solid rocket motor redesign**, in: Proc. Fourth International Symposium on Science and Engineering on CRAY Supercomputers, A Cray Research, Inc. book (1988)

3. C. Farhat, **Multiprocessors in computational mechanics**, Ph.D. Thesis, University of California at Berkeley, 1986.

4. C. Farhat, **A simple and efficient automatic FEM domain decomposer**, Comput. and Struc., Vol. 28, No. 5, pp. 579-602, 1988.

5. D. Vanderstraeten, R. Keunings and C. Farhat, **Beyond conventional mesh partitioning algorithms and the minimum edge cut criterion: impact on realistic applications, Parallel Processing for Scientific Computing**, ed. by D. Bailey et. al., SIAM, pp. 611-614, 1995.

6. D. Vanderstraeten, C. Farhat, P. S. Chen, R. Keunings and O. Zone, **A retrofit and contraction based methodology for the fast generation and optimization of mesh partitions: beyond the minimum interface size criterion**, Comput. Meths. Appl. Mech. Engrg., Vol. 133, pp. 25-45, 1996.

7. C. Farhat, S. Lanteri and H. D. Simon, **TOP/DOMDEC, A software tool for mesh partitioning and parallel processing**, J. Comput. Sys. Engrg., Vol. 6, No. 1, pp. 13-26, 1995.

8. C. Farhat and M. Lesoinne, **Automatic partitioning of unstructured meshes for the parallel solution of problems in computational mechanics**, Internat. J. Numer. Meths. Engrg., Vol. 36, No. 5, pp. 745-764, 1993.

9. B. Hendrickson and R. Leland, **The Chaco User's Guide: Version 2.0**, Tech. Report SAND94--2692, 1994.

10. http://www.cs.umn.edu/~karypis/metis/metis.html

11. C. Farhat and F. X. Roux, **A method of finite element tearing and interconnecting and its parallel solution algorithm**, Internat. J. Numer. Meths. Engrg., Vol. 32, pp. 1205-1227, 1991.

12. C. Farhat and F. X. Roux, **Implicit parallel processing in structural mechanics**, Computational Mechanics Advances, Vol. 2, pp. 1-124, 1994.

13. C. Farhat, J. Mandel and F. X. Roux, **Optimal convergence properties of the FETI domain decomposition method**, Comput. Meths. Appl. Mech. Engrg., Vol. 115, pp. 367-388, 1994.

14. C. Farhat, P. S. Chen and J. Mandel, **A scalable Lagrange multiplier based domain decomposition method for implicit time-dependent problems**, Internat. J. Numer. Meths. Engrg.,

Vol. 38, pp. 3831-3854, 1995.

15. C. Farhat and J. Mandel, **The two-level FETI method for static and dynamic plate problems - Part I: an optimal iterative solver for biharmonic systems**, Comput. Meths. Appl. Mech. Engrg., 1997.

16. C. Farhat, P. S. Chen and F. X. Roux, **The two-level FETI method for plate problems - Part II: extension to shells, parallel implementation and performance results**, Comput. Meths. Appl. Mech. Engrg., 1997.

17. **EUROPORT-1: Survey and results in fluid dynamics and structural mechanics - Parallel codes ready for industrial use**, European Commission, Directorate-General III, Industry, Esprit Initiative.

SC97 Thanks/Committee Lists

Acknowledgement

SC97 gratefully acknowledges the many supporters of this year's conference, including representatives from academia, industry, and the government. The resources pulled together each year for the SCing conference are impressive, and too numerous to list. But the conference would not be a success without the donations of time, energy, and resources from throughout the high performance networking, computing, and education communities, and we thank you.

SC97 Committees

SC97 Committee

Conference Chair
Dona Crawford
Sandia National Laboratories
510-294-2628
dona@ca.sandia.gov

Conference Office

Bernie Marx
Sandia National Laboratories
510-294-2629
bimarx@ca.sandia.gov

Conference Vice Chair
Dave Cooper
LLNL
510-422-4010
dcooper@llnl.gov

Conference Deputy Chair
Dennis Duke
SCRI
904-644-0175
dduke@scri.fsu.edu
NPACI, UCSD
619-534-8381
dduke@NPACI, UCSD.edu

Program Co-chair
Greg Papadopoulos
Sun Microsystems
415-786-7378
gregp@corp.sun.com

Program Co-chair
Margaret Simmons
NPACI, UCSD
970-925-5774
mls@NPACI, UCSD.edu

Deputy Program Chair
Joanne Martin
IBM Corporation
914-433-8493
jmartin@vnet.ibm.com

BOF Chair, Tutorials Chair
Ann Hayes
LANL
505-665-4506
ahh@acl.lanl.gov

Education Co-chair
Roscoe Giles
Boston University
617-353-6082
roscoe@bu.edu

Education Co-chair
Bryant York

Northeastern University
617-373-2177
york@ccs.neu.edu

Networking Chair
Bob Borchers
NSF
703-306-1970
borchers@nsf.gov

Industrial Liaison Chair
Bill Boas
Essential Communications
505-344-0080, Ext. 307
bboas@esscom.com

Exhibits Chair
Cherri Pancake
Oregon State University
541-737-2109
pancake@cs.orst.edu

Registration Chair,
Committee Coordination Chair
Karen Friedman
NCAR
303-497-1276
karen@ncar.ucar.edu

Finance Chair
Pam Howard
LLNL
510-423-6099
phoward@llnl.gov

Conference Center Chair
John Ranelletti
LLNL
510-424-6975
johnr@llnl.gov

Local Arrangements Chair
Mary Amiot
Cray Research, Inc.
612-683-3524
mary.amiot@cray.com

Publicity Chair
Ann Redelfs
NPACI, UCSD
619-534-5032
redelfs@NPACI, UCSD.edu

Graphic Designer
Mo Viele
Mo Viele, Inc.
607-272-4172
mv12@cornell.edu

Exhibition Management
Don Collier
DC Expositions, Inc.
972-423-4286, 1-888-980-5488
dcexpo@worldnet.att.net

SIGARCH Liaison
Lisette Burgos
ACM
212-626-0165
burgos@hq.acm.org

IEEE Computer Society Liaison
Anne Marie Kelly
IEEE Computer Society
202-371-1013
a.m.kelly@computer.org

SC97 Program Committee

Program Committee Co-chair
Greg Papadopoulos
Sun Microsystems

Program Committee Co-chair
Margaret Simmons
NPACI, UCSD

Deputy Program Chair
Joanne Martin
IBM Corporation

Panels/Workshops Chair
Francine Berman
UC San Diego

Posters Chair
Pat Teller
University of Texas at El Paso

Survey Talks Chair
Jim McGraw
LLNL

Proceedings Chair
Dan Dwyer
Cornell Theory Center

David Bailey
NASA, Ames Research Center

Forest Baskett
Silicon Graphics, Inc.

Eric Brewer
UC Berkeley

David Callahan
Tera Computer Co.

Andrew Chien
University of Illinois, Hewlett-Packard

David Culler
UC Berkeley

Tom Defanti
EVL, University of Illinois at Chicago

Jack Dongarra
University of Tennessee, ORNL

Ian Foster
Argonne National Laboratory

Dennis Gannon
CRPC, Indiana University

Gary Glatzmaier
LANL

Mary Hall
Uiversity of Southern California

Steve Hammond
NCAR

Karim Harzallah
Tandem Computers

Ann Hayes
LANL

Tony Hey
University of Southampton

Mari Horowitz
Stanford University

Sally Howe
National Coordination Office for CIC

David Kahaner
NIST

Carl Kesselman
Caltech

Subodh Kumar
SC96 Best Student Paper Winner

Anne Lavery
Sun Microsystems

Kai Li
Princeton University

Bart Miller
University of Wisconsin

Mike Norman
NCSA

Yoshio Oyanagi
University of Tokyo

Philip Papadoupolos
ORNL

Dan Reed
University of Illinois

John Reynders
LANL

Roy Richter
General Motors R&D Center

Richard Shaginaw
Bristol-Myers-Squibb

Horst Simon
Lawrence Berkeley Laboratory

Per Stenstrom
Chalmers University of Technology

David Tennenhouse
MIT Laboratory for Computer Science

David Wood
University of Wisconsin

Paul Woodward
University of Minnesota

SC97 Tutorials Committee

Tutorials Chair
Ann Hayes
LANL

Tutorials Vice-chair
Rajeev Thakur
Argonne National Laboratory

Richard Allen
Sandia National Laboratories

Donald Austin
National Coordination Office for CIC

David Bailey
NASA Ames Research Center

Jeff Brown
LANL

Marc Christon
Sandia National Laboratories

Dennis Duke
SCRI, NPACI, UCSD

Marty Itzkowitz
Silicon Graphics, Inc.

Tom Kitchens
US Department of Energy

Allen Malony
University of Oregon

Sam Milosevich
Eli Lilly and Company

Ken Miura
Fujitsu America

Tom Morgan
Argonne National Laboratory

Nicholas Nystrom
Pittsburgh Supercomputing Center

Rod Oldehoeft
US Department of Energy

Doug Pase

IBM Corporation

Dan Pryor
Center for Computing Sciences

Richard Pryor
Sandia National Laboratories

John Ranelletti
LLNL

Pat Teller
University of Texas at El Paso

Linda Torczon
Rice University

Steve Wallach
Hewlett-Packard

Mary Zosel
LLNL

SC97 Industrial Liason Committee

Industrial Liaison Chair
Bill Boas
Essential Communications

HPC Challenge Chair
Louis Turcotte
USAE Waterways Experiment Station

HPC Challenge Committee

HPC Challenge Chair
Louis Turcotte
USAE Waterways Experiment Station

Bob Pearson
Essential Communications

Steve Jones
USAE Waterways Experiment Station

John West
USAE Waterways Experiment Station

SC97 History Exhibit Committee

SC97 History Exhibit Chair
Bill Boas
Abba Technologies

Jordan Becker
AOL/ANS

Gwen Bell
The Computer Museum

Tim Berners-Lee
W3 Consortium

Jacquie Boas
Abba Technologies

Vint Cerf
MCI

Dave Clark
Massachusetts Institute of Technology

Ellen Cline
LVB Design

Andrew Cochran
Cochrane Communications

Jan Hauser
Sun Microsystems

Van Jacobson
Lawrence-Berkeley Nat'l Labs

Dennis Jennings
University College Dublin

Bill Joy
Sun Microsystems

Robert Kahn
Corporation for National Research Initiatives

Jon Keegan
Media Circus

Len Kleinrock
University of California, Los Angeles

Dan Lythcott-Haims
Lythcott-Haims Design

Merry Maisel
NPACI, UCSD

Paul Mockapetris
Software.Com

Don Nielson
SRI International

Stacy Oliver
SV.MARCOM

Craig Partridge
GTE/BBN

Brian Reid
Digital Equipment Corporation

Larry Roberts
ATM Systems

Larry Smarr
NCSA University of Illinois

George Strawn
National Science Foundation

Al Weis
Advanced Network Services

Irving Wladawsky-Berger
IBM Internet Division

Steve Wolff
Cisco Systems

SC97 Research Exhibits Committee

Research Committee Chair
Sally Haerer
NCAR

Mary Kay Bunde
Silicon Graphics, Inc., Cray Research, Inc.

Joan Francioni
University of Southwestern Louisiana

B. Horner-Miller
Jet Propulsion Lab/Caltech

Walter Wehinger
University of Stuttgart

SC97 Education Committee

Education Co-chair
Roscoe Giles
Boston University

Education Co-chair
Bryant York
Northeastern University

Student Volunteer Chair
Juan Meza
Sandia National Laboratories

Teacher Travel Coordinator
Edna Gentry
University of Alabama-Huntsville

Gypsy Abbot
University of Alabama at Birmingham

Andrew Bernat
University of Texas at El Paso

Ginger Caldwell
NCAR

Don Coleman
Howard University

Harvey Gould
Clark University

Charles Hardnett
Spelman College

Arthur Karshmer
New Mexico State University

Roseann Krane
San Ramon Valley Unified School District

Cynthia Lanius
Houston School District

Tom Marchioro
University of Washington

Josie Morgan
Alabama Supercomputing Authority

Judy Olsen
NCSA

Brahm Rhodes
Onyz Sciences Corp

Kris Stewart
San Diego State University, NPACI, UCSD

Barbara Summers
ORNL

Valerie Taylor
Northwestern University

Umesh Thakkar
NCSA

Dave Thomas
Montana State University

Mary Ellen Verona
Maryland Virtual High School

Beth Ann Ziebarth
University of Illinois Cooperative Extension Service

John Ziebarth
NASA, Ames Research Center

SC97 Networking Committee

Networking Chair
Bob Borchers
NSF

Guy Almes
Advanced Network Services

Kathy Benninger
Pittsburgh Supercomputing Center

Bridget Borchers
Non Stop Computing

Jim Brandt
Sandia National Laboratories

Bob Bryant
LLNL

Wayne Butman
LLNL

Jerry DeLapp
LANL

Rich Gay
Sandia National Laboratories

Shris Gee
Pacific Bell

Ken Goodwin
Pittsburgh Supercomputing Center

Ken Hays
Florida State University

Russ Hobby
UC Davis

John Jamison
MCI

Merritt Jones
The Mitre Corporation

Michael Lambert
Pittsburgh Supercomputing Center

John Laroco
Sandia National Laboratories

Bill Lennon
LLNL

Rick Maurer
Sandia National Laboratories

Jim McCabe
NASA Ames Research Center

David Meyer
University of Oregon/Network Services

Eric Montoya
Essential Communications

Randy Nicklas
MCI

Kevin Oberman
Energy Sciences Network

Ben Peek
GST Information Services

Tom Pratt
Sandia National Laboratories

Bob Reschly
Army Research Laboratory

Steve Tenbrink
LANL

Tim Toole
MCI

Scott Whyte
Network Evolutions, Inc.

Warren Williams
Pacific Bell

Dave Wiltzius
LLNL

SC97 Exhibits Committee

Exhibits Chair
Cherri Pancake
Oregon State University

Exhibitor Forum Coordinator
Joan Winter
Oregon State University

Exhibition Management
Don Collier
DC Expositions, Inc.

Exhibits-Show Contractor
Darryl Monahan
GES Exposition

Mary Kay Bunde
Silicon Graphics, Inc./Cray Research Inc.

Hugh Caffey
Hewlett Packard

Tim Mattson
Intel

Sherry Pittam
Oregon State University

SC97 Conference Center Committee

Conference Center Chair
John Ranelletti
LLNL

SC97 Signage Chair
Heidi Peterson
LLNL

SC97 Publicity Committee

Publicity Chair
Ann Redelfs
NPACI, UCSD

Graphic Designer
Mo Viele
Mo Viele, Inc.

Gail Bamber
NPACI, UCSD

Mike Bernhardt
The Bernhardt Agency

Cheryl Converse-Rath
NPACI, UCSD

Laura Doiron
NPACI, UCSD

Amy Finley
NPACI, UCSD

Dave Hart
NPACI, UCSD

Nancy Jensen
NPACI, UCSD

Barbara Richmond
NPACI, UCSD

Mika Wheeler
NPACI, UCSD

SCXY Steering Committee

SCXY Chair
Beverly Clayton
Pittsburgh Supercomputing Center

Marisa Chancellor
Cisco Systems, Inc.

David Cooper
LLNL

Dona Crawford
Sandia National Laboratories

Dennis Duke
SCRI

Ann Hayes
LANL

Sid Karin
UC San Diego

George Michael
Retired

Cherri Pancake
Oregon State University

Daniel Pryor
Center for Computing Services

Xiaodong Zhang
University of Texas at Austin

Ex-officio members

Lisette Burgos
ACM

Anne Marie Kelly
IEEE Computer Society

Proceedings Table of Contents
Final Program Table of Contents

This year marks the 10th anniversary of the SCXY conference series, returning to the city of the first conference.

Come participate in one of the world's most prestigious conferences! We invite you to contribute to the following program elements:

Technical papers
Tutorials
Roundtables and Panels
Education program
Poster Exhibits
Research Exhibits
SCinet98

Other components of the conference will include:

Keynote Address
Invited Presentations and Panels
Industry Exhibits
Presentations of Awards

You can participate in SC98 in many ways. Write for additional information about:

Technical Papers
papers98@mail.supercomp.org

Education Papers
education98@mail.supercomp.org

Tutorials
tutorials98@mail.supercomp.org

Research Exhibits
research_exhibits98@mail.supercomp.org

Poster Exhibits
posters98@mail.supercomp.org

SCinet98
scinet98@mail.supercomp.org

Exhibitor Forum
forum98@mail.supercomp.org

Birds-of-a-Feather
bof98@mail.supercomp.org

ORANGE COUNTY CONVENTION CENTER

http://www.citizens-first.co.orange.fl.us/onetgov/dept/conven/index.htm
9800 International Drive; Orlando, FL, USA 32819
1 407 345 9800

SC98 GENERAL INFORMATION

www.supercomp.org/sc98
sc98@mail.supercomp.org
1 888 778 sc98 (toll free)

SC98 EXHIBIT INFORMATION

www.supercomp.org/sc98/exhibits
dc@dcexpo.com
1 888 980 5488 (toll free)

SC98 GENERAL ADDRESS

San Diego Supercomputer Center
10100 Hopkins Drive
La Jolla, CA, USA 92093

Proceedings Table of Contents Final Program Table of Contents

Final Program Table of Contents

SAN JOSE SC97 NOV. 15-21

Welcome to SC97

We are pleased to welcome you to SC97: High Performance Networking and Computing. The conference is sponsored by the Association for Computing Machinery (ACM) Special Interest Group on Computer Architecture (SIGARCH) and the IEEE Computer Society Technical Committees on Supercomputing Applications and Computer Architecture.

Amidst many other highlights, we are particularly proud of the technical track and education program. We received more technical paper submissions than any previous year, and the papers selected comprise all aspects of high performance networking and computing.

The education program focuses on the role of technology in lifelong learning, and through it we will explore new models for teaching and learning that make use of emerging computing technologies.

This year's 17 tutorials cover a wide range of topics, including do-it-yourself supercomputers, MPI, parallel adaptive mesh refinement solutions, data mining, performance evaluation, metacomputing systems, Java, and more. And the exhibition, featuring research, poster, and industry exhibits, will be the largest in the history of the SCing conference series.

As in past years, the infrastructure that comes together for this week-long event draws on the resources and expertise of our nation's best in communications and support. The volunteers who pull together to make the infrastructure--and in fact the entire conference--a success are what makes this annual meeting so dynamic and successful.

SC97 promises to be an unsurpassed opportunity for our community to interact and discuss the future of our industry, the innovations that will fuel it, and our responsibility to help educate the next generation. Welcome!

Dona Crawford, Sandia National Laboratories Conference Chair
Dave Cooper, Lawrence Livermore National Laboratory Conference Vice Chair