

SAN JOSE **SC97** NOV. 15-21

General Information

Business and Other Services

Federal Express
181 Park Avenue
Park Center Plaza
San Jose, CA
800-463-3339

The Wayback Machine - <https://web.archive.org/web/19980216120239/http://www.supercomp.org:80/sc97/proceedings/G...>



General Information

Business and Other Services

Federal Express
181 Park Avenue
Park Center Plaza
San Jose, CA
800-463-3339

Kinko's
93 E. San Carlos
San Jose
408-295-4336

San Jose McEnery Convention Center Business Center
Located on mail level of convention center
Open 8:00am-5:30pm daily during the conference
408-977-0900

San Jose Convention and Visitors Bureau
333 W San Carlos Street, Suite 1000
San Jose CA
408-295-9600

Coat and Bag Check

Located on the main level outside registration area. Open daily, 7:00am-5:30pm.

Conference Store

SC97 Store Hours

Saturday	November 15	1:00-6:00pm
Sunday	November 16	7:30am-6:00pm
Monday	November 17	7:30am-9:00pm
Tuesday	November 18	7:30am-6:00pm
Wednesday	November 19	7:30am-6:00pm
Thursday	November 20	7:30am-5:00pm
Friday	November 21	8:00am-noon

The conference store phone: 408-288-6820

Location: First Level Lobby of the San Jose McEnery Convention Center, near SC97 registration

The store will be selling:

CD-ROM Proceedings	\$50
IOPADS Proceedings	\$55
Tutorial Notes	
Passport Registrants	\$20
Conference Registrants	\$40
Monday Opening Reception	\$65
Thursday Conference Reception	\$75
Tutorials Luncheon (Sunday or Monday)	\$25

First Aid

Facilities are available to all attendees and exhibitors. The First Aid facility, staffed by City of San Jose paramedics, is located in the San Jose McEnery Convention Center. Emergencies should be reported to security (call 408-277-3500) for immediate assistance.

Media Relations

SC97 Press Room

Located in Room N, lower level
408-288-7540, -7551, -7638
408-288-7647 fax

The SC97 Press Room offers members of the media a full-service support center designed with the reporter's schedule and workload in mind. What you will find in addition to typical telephone and fax equipment are electronic and printed press kits from the exhibitors, computers for accessing e-mail and the Web, a printer, a press room Web site with exhibitor information, message and events board, and a skilled staff of volunteers.

Hours of Operation

Monday	November 17	9:00am-9:30pm
Tuesday	November 18	8:00am-6:00pm
Wednesday	November 19	9:00am-6:00pm
Thursday	November 20	9:00am-5:00pm
Friday	November 21	9:00am-noon

Press Room Services

The SC97 press room staff will be available to assist with press requests, including use of the computer facilities. In addition, staff volunteers will assist in arranging interviews with representatives from the SC97 committees and exhibitors wherever possible. The press room will also distribute invitations to exhibitor-sponsored events to media representatives.

Press Events

Several activities have been scheduled exclusively for members of the press:

- **Press Presentation on SC97 Highlights**
Monday, November 17 5:30-6:30pm
Conference Chair Dona Crawford introduces several of SC97's State-of-the-Field speakers to preview their talks, as well as contributors to the "Science and the Internet: 20 Years" history exhibit.
- **Exhibition Tour**
Monday, November 17 6:30pm
After the Highlights Presentation, members of the SC97 Executive Committee will escort small groups of media representatives on a private tour of the exhibition floor before the opening night reception.
- **Keynote Address**
Tuesday, November 18 8:30-10am
Paul Saffo, director of the Institute for the Future, will speak.

Message Services

Email

408-288-7821

Message Center

408-288-6110

Phone and Room Numbers

Conference Committee Office

408-288-6512

408-288-6515

408-288-6541 fax

Exhibition Management

408-288-6020

408-288-5753 fax

Local Arrangements

408-288-6329

Student Volunteer Office

408-288-6926

408-288-6955 fax

Press Room

408-288-7540

408-288-7647 fax

Registration

408-288-5820

Email

408-288-7821

Message Center

408-288-6110

Audio Visual

408-288-7001

SCinet97/ NOC

408-288-5262

Proceedings/CD-ROM

SC97 proceedings are available on CD-ROM and on the Web at: <http://www.supercomp.org/sc97/proceedings/>. Additional copies of the Proceedings will be on sale at the conference store. To facilitate access to the Proceedings during the conference, the following special services are offered:

Workstations for viewing the proceedings are available in Room J14 at the San Jose McEnery Convention Center, during the following times:

Sunday November 16 1-6:00pm

Monday-Thursday November 17-20 8:00am-7:00pm

Friday November 21 8:00am-noon

Technical registrants will receive one complimentary hardcopy of any individual paper upon request. Additional copies may be ordered at the Conference Store.

Registration

On-site Registration

On-site registrants will pay late registration fees. Forms of payment accepted on site: credit cards (MasterCard, VISA, American Express, Diners Club), traveler's checks, company and personal checks drawn on U.S. banks, and payable to the "IEEE Computer Society," cash, and purchase orders from U.S. Government offices only (original purchase order must be presented on site).

Member and Student Discounts

To qualify for member discounts, include your current ACM, ACM SIGARCH, IEEE, or IEEE Computer Society membership number(s) on the registration form. To qualify for a student discount, attach a copy of a valid full-time student identification card to the registration form.

Changes/Substitutions

A letter on your company's letterhead is required from the original registrant stating the conditions of the substitution and the name of the replacement, and should be delivered to the Special Assistance desk at SC97 Registration.

Lost Badge Fee

SC97 will assess a \$40 processing fee to replace badges for those attendees who lose their badges during the conference.

Registration and Store Hours

Saturday	November 15	1:00pm-6:00pm
Sunday	November 16	7:30am-6:00pm
Monday	November 17	7:30am-9:00pm
Tuesday	November 18	7:30am-6:00pm
Wednesday	November 19	7:30am-6:00pm
Thursday	November 20	7:30am-4:00pm
Friday	November 21	8:00am-noon

Media Registration and Information

Members of the media who need to register on site should go to the Special Assistance counter at the SC97 registration area. For more information on press services at SC97, see [Press Registration](#).

Special Registrations Needs

We will attempt to accommodate participants' special registration needs. Please contact the Special Assistance desk in the registration area for assistance.

Tutorials

Tutorials are being offered on Sunday and Monday, November 16 and 17. Tutorial fees are not included in the technical program registration fees and require separate registration.

A special Tutorials Passport for one or two days is available. A ONE-DAY Passport (either Sunday or Monday) entitles you to attend any of the tutorials offered on that day, a luncheon, and your selection of notes for one full-day tutorial or two half-day tutorials taking place that day.

A TWO-DAY Passport gives you approximately a 20% discount on the price of two, One-Day Passports. You will be able to attend any of the tutorials offered on both days, a luncheon each day, and your selection of notes as you specified on the registration form. **PLEASE NOTE:** there is no half-day tutorial passport option. You must select one full-day or two half-day notes with your daily passport.

Seating at the tutorials is on a first-come, first-served basis. You are welcome to sit in on as many tutorials as you like for the day(s) you are registered, but you will receive only the notes that you reserve in advance as specified on the registration form.

Tutorial Notes

An additional benefit of the Tutorials Passport is that you can purchase additional tutorial notes at a significantly reduced rate. Notes for the tutorials can be picked up at the SC97 Conference Store upon presentation of the notes ticket(s) provided in your registration packet. You will receive notes only for those tutorial(s) you selected on your registration form. Extra copies of tutorial course notes will be on sale at the Conference Store following the completion of the course, and orders can be placed there for additional tutorial notes to be delivered after the conference.

Exhibits Only Badges

Exhibits Only Badges (\$50 a day) allow entrance to the exhibit floor for one of the three days of the exhibition: Tuesday, Wednesday, or Thursday. You may also buy an Exhibits Only Badge for each day, however the price remains \$50 a day (\$100 for two days or \$150 for three days). Exhibits Only Badges will go on sale beginning Tuesday morning at 7:30 am. SC97 is offering Exhibits Only registration to children 12 to 16 years of age for whom Exhibits Only Badges (\$50 a day) must be

purchased. Children in this age group must be accompanied by an adult at all times while on the exhibit floor. Children under age 12 are not permitted on the exhibit floor at any time.

What You Get With Your Registration

Special Transportation Assistance

San Jose offers the physically challenged visitor up-to-date information for special transportation requirements. Transportation is available with 24 hours notice through MV Transportation at 408-292-3600.

SC97 Conference Hotels

Airport Inn International
13555 N. Fourth St.
408-453-5340

Best Western
455 S. Second St.
408-298-3500

Crowne Plaza
Park Center Plaza
282 Almaden Blvd.
408-998-0400

Doubletree
2050 Gateway Place
408-453-4000

Fairmont
Fairmont Plaza
170 S. Market St.
408-998-1900

San Jose Hilton & Towers
300 Almaden Blvd.
408-287-2100

Hotel De Anza
233 W. Santa Clara St.
408-286-1000

Hyatt Sainte Claire
302 S. Market St.
408-885-1234

Hyatt San Jose
San Jose Airport
1740 N. First St.
408-993-1234

Radisson Plaza
1471 N. Fourth St.
408-452-0200

Westin
5101 Great America Pkwy
408-986-0700

Wyndham
1350 N. First St.
408-453-6200

SC97 Housing Bureau
408-295-2265, ext. 426
408-293-3705 fax

SC97 Shuttle/Bus Service

Shuttle bus transportation will be provided to the San Jose Convention Center from the following hotels:

- Airport Inn International
- Doubletree
- Hyatt San Jose
- Radisson Plaza
- Westin Santa Clara
- Wyndham

The shuttle will depart from the above hotels to the San Jose McEnery Convention Center every hour on the following schedule

Sunday	7:00am to 11:00am
November 16	3:00pm to 10:30pm
Monday	7:00am to 11:00am
November 17	3:00pm to 10:30pm
Tuesday	7:00am to 7:30pm
November 18	
Wednesday	7:00am to 7:30pm
November 19	
Thursday	7:00am to 6:00pm
November 20	6:00pm to 11:30pm for Great America Event
Friday	7:00am to 1:00pm
November 21	

Special Transportation Assistance

San Jose offers the physically challenged visitor up-to-date information for special transportation requirements. Transportation is available with 24 hours notice through MV Transportation at 408-292-3600.

SC97 Social Activities

Monday, November 17 7-9:00pm

The Gala Opening for SC97 provides technical program attendees with their first glance at the exhibition, including research, poster, and industry exhibits. Tickets for this event will be on sale in the store.

Wednesday, November 19 5-7pm

The posters reception offers attendees the opportunity to discuss research results with poster presenters. The reception takes place in the Almaden Concourse of the San Jose McEnery Convention Center.

Thursday, November 20 6:30-9pm

The Thursday Conference Reception will take place at the Great America Paramount Park. We will have buffet dinner in the Pavilion from 6:30 to 9:00pm and specialty stations throughout the park from 8 to 10pm. The following areas and attractions will operate from 6:30 to 10:30pm.

- Carousel Plaza
 - Carousel Columbia
 - The Skytower
 - Two screenings of the IMAX film "Special Effects"
 - Appearances by costumed Star Trek characters
- Orleans Place
 - Top Gun Jetcoaster
 - Rue Le Dodge Bumper Cars
 - One Performance of Karaoke Connection
- Hometown Square
 - Vortex
 - Triple Play

The buses will run from 5:30-11pm. See [SC97 Shuttle/Bus Service](#) for more information.



SC97 CONFERENCE PROCEEDINGS

SC97: HIGH PERFORMANCE NETWORKING AND COMPUTING

	 SC97 Proceedings Table of Contents
	 SC97 Final Program Table of Contents
	 Instructions on Using the CD-ROM
SAN JOSE SC97 NOV. 15-21	 Welcome to SC97



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>

-  [Technical Papers](#)
-  [Education Program](#)
-  [Research Exhibits](#)
-  [Poster Abstracts](#)
-  [Tutorials](#)
-  [HPC Challenge](#)
-  [Gordon Bell](#)
-  [Panels](#)
-  [Industry Exhibitors](#)
-  [Invited Speakers](#)
-  [Fernbach Award](#)
-  [Committee Lists](#)
-  [SC98 Call For Participation](#)



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.



SESSIONS



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



1.  [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.  [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



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-

[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
2.  [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.



1.  [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



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



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



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



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



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



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



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



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.



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



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



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 Performance of 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



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.  [A Scalable Mark-sweep Garbage Collector on 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



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 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
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



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](#)




ABSTRACTS


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

[Room B1] Communication in Clusters

Session Chair: David Culler, University of California, Berkeley



1. [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. [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 analyze 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. [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. [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.,

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 on 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 Chalermwua, 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; Mitsuhsa 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 an 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 straightforward 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 algorithm's 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