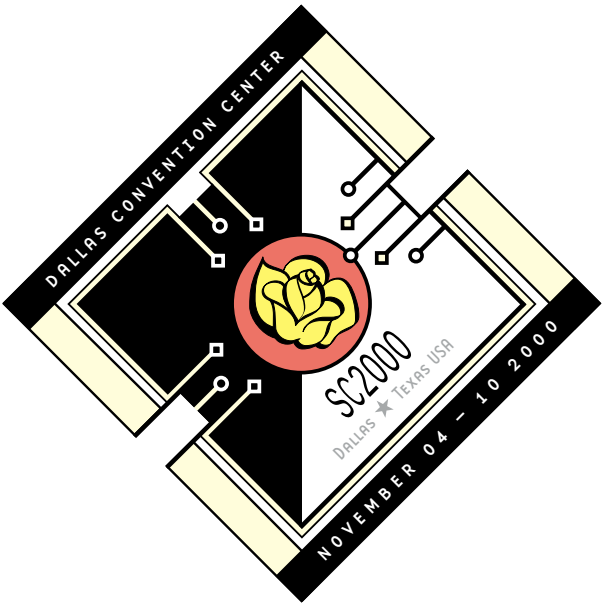

SC2000 Final Program

HIGH PERFORMANCE NETWORKING AND COMPUTING

DALLAS CONVENTION CENTER
Conference: November 04 – 10
Exhibition: November 06 – 09

SC2000 GENERAL INFO
www.sc2000.org
info@sc2000.org

SPONSORED BY
the IEEE Computer Society and ACM SIGARCH

DALLAS TEXAS USA

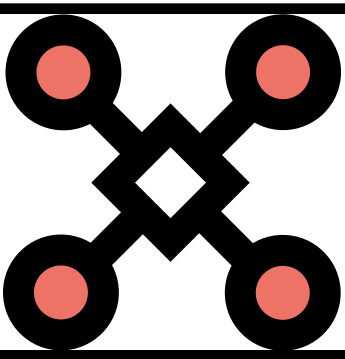

TABLE OF CONTENTS

WELCOME		
	Letter from Conference Chairs	2
CONFERENCE SCHEDULES		
	Conference and Daily Schedules at a Glance	foldouts
TECHNICAL PROGRAM		
	Tutorials	11
	Keynote Address	25
	State-of-the-Field Talks	25
	Masterworks	31
	Webcasts	44
	Technical Papers / Awards	45
	Panels	77
	HPC Games	81
	Education Program	83
	Venture Village / eSCape 2000	89
	Exhibitor Forum	93
	Birds-of-a-Feather	105
NETWORK		
	SCinet 2000	113
EXHIBITION		
	Research Exhibits	123
	Research Gems	145
	Industry Exhibits	155
	Exhibition Map	174
GENERAL INFORMATION		
	Proceedings, Registration, Media Relations, Business and Other Services, Conference Store, First Aid, E-mail/Message Center, Conference Hotels, DART and Dallas Convention Center Maps	175
THANKS / COMMITTEES		
	Acknowledgements SC2000 Committee listings	183

SC2000 welcome

Welcome to Texas and SC2000!

This year's conference promises to be BIG, innovative, and we think you'll enjoy every minute of it. New features added to this year's conference include—the Venture Village, eSCape 2000, a reconfigured HPC Games, and the new networking award. You'll see a strong program of technical presentations and technology displays. Masterworks sessions will highlight the successful application of high performance computing and networking technology to solve real-world problems. Three features new to SC will highlight this year's meeting. The first, called Venture Village, will showcase a collection of entrepreneurial information technology companies, all creating new products to build the infrastructure of tomorrow. The second innovative thrust is eSCape 2000, demonstrating the ability to "escape" from today's technological boundaries and to connect and compute anywhere. The third innovation shows off the SCinet 2000 network, one of the most intense networks on the planet for the duration of the show, through the SC2000 Network Challenge.

11.04 – 11.10
dallas tx usa
www.sc2000.org

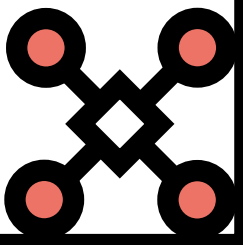

We welcome you to Dallas and to our annual showcase of accomplishments and goals for high performance networking and computing.

Louis Turcotte, *SC2000 Chair*
Betsy Schermerhorn, *SC2000 Vice Chair*
Charlie Slocomb, *SC2000 Deputy Chair*

SC2000 CONFERENCE AT A GLANCE

	SATURDAY 4	SUNDAY 5	MONDAY 6	TUESDAY 7	WEDNESDAY 8	THURSDAY 9	FRIDAY 10
TUTORIALS		8:30AM – 5PM	8:30AM – 5PM				
EDUCATION PROGRAM	1PM – 9PM	9AM – 6:30PM	8:30AM – 7PM	8:30AM – 9PM	8:30AM – 9PM	8:30AM – 2PM	
KEYNOTE ADDRESS				8:30 –10AM: **S. Wallach**			
STATE-OF-THE-FIELD TALKS					8:30AM: **T. Sterling** 9:15AM: **E. Spafford**	8:30AM: **M. Wright** 9:15AM: **J. Browne**	
TECHNICAL SESSIONS *• Masterworks* *• Panels* *• Papers* *• Awards* *• Invited Speakers* *• SC2000 Reception*	**10:30AM – 12NN**			**Masterworks** Computational BioSciences: Genomics	**Masterworks** IEEE Award Winners Seymour Cray Computer Engineering Award Winner Sidney Fernbach Award Winner	**Masterworks** Supercomputing Trends in MCAE I	**8:30AM – 10AM** **Panels** Convergence of the Extremes Computational Grids: A Solution Looking for a Problem?
				Panel Venture Capital Panel **Papers** MPI Numerical Algorithms	**Papers** Applications I Visualization Compiler Optimization	**Papers** Software Tools Data Grid Gordon Bell II	
	1:30PM – 3PM			**Masterworks** Computational Biochemistry and Drug Discovery	**Masterworks** Real World Scalable Computing I	**Plenary Session** Awards and Invited Speaker	**10:30AM – 12NN** **Panels** Open Source: IP in the Internet Era MegaComputers
				Papers Scheduling MPI/OpenMP Potpourri	**Papers** Applications II Networking Hardware-Based Tools		
	3:30PM – 5PM			**Masterworks** Computing Platforms	**Masterworks** Real World Scalable Computing II	**Masterworks** Supercomputing Trends in MCAE II	
				Papers Cluster Infrastructure QoS/Fault Tolerance Biomedical Applications	**Panel** TCP Panel **Papers** Gordon Bell I Parallel Programming	**Panel** Petaflops Around the Corner **Papers** Applications Support Grid Middleware	
EXHIBITOR FORUM				10AM – 5PM	10AM – 5PM	10AM – 4PM	
RESEARCH GEMS, RESEARCH EXHIBITS, INDUSTRY EXHIBITS			7PM – 9PM	10AM – 6PM	10AM – 6PM **Research Gems/eSCape Open House** 10AM –11AM	10AM – 4PM **Research Gems/eSCape Open House** 10AM –11AM	
GALA OPENING & SC2000 RECEPTION			**Gala Opening** 7PM – 9PM			**SC2000 Reception** 6:30PM – 10PM	
BOFS				5:30PM – 7PM	5:30PM – 7PM		
REGISTRATION & STORE	1PM – 5PM	7:30AM – 6PM	7:30AM – 9PM	7:30AM – 5:30PM	7:30AM – 5:30PM	7:30AM – 5PM	8AM – 11AM
MEDIA ROOM	3PM – 5PM	1PM – 4PM	9AM – 6PM	9AM – 6PM	9AM – 6PM	9AM – 5PM	

SUNDAY AT A GLANCE

SATURDAY AT A GLANCE

NOVEMBER 4

REGISTRATION / STORE

1PM – 5PM | **EXHIBIT HALL** | PAGE 175

EDUCATION PROGRAM

1PM – 9PM | PAGES 83 – 88

ALL EDUCATION PROGRAM SESSIONS ARE HELD ON LEVEL 2, C SECTION OF THE DALLAS CONVENTION CENTER UNLESS OTHERWISE NOTED

Teacher Registration
1PM – 5PM | PAGE 84

Basic Skills Tutorial
4PM – 5PM | PAGE 84

Orientation and Laptop Setup
5PM – 6PM | PAGE 84

Education Program Reception
7PM – 9PM | PAGE 84

MEDIA ROOM

3PM – 5PM | **ROOM D 326** | PAGE 177

NOVEMBER 5

SUNDAY FULL-DAY TUTORIALS

S1 THROUGH S5 | 8:30AM – 5PM | PAGES 11 – 14

S1 Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface

S2 An Introduction to High Performance Data Mining

S3 Design and Analysis of High Performance Clusters

S4 High Performance Numerical Computing in Java: Compiler, Language, and Application Solutions

S5 Performance Analysis and Tuning of Parallel Programs: Resources and Tools

SUNDAY HALF-DAY TUTORIALS [AM]

S6A THROUGH S8 | 8:30AM – 12NN | PAGES 14 – 16

S6A Mesh Generation for High Performance Computing Part I: An Overview of Unstructured and Structured Grid Generation Techniques

S7A Introduction to Effective Parallel Computing, Part I

S8 Tools and System Support for Managing and Manipulating Large Scientific Datasets

SUNDAY HALF-DAY TUTORIALS [PM]

S6B THROUGH S10 | 1:30PM – 5PM | PAGES 16 – 18

S6B Mesh Generation for High Performance Computing Part II: Mesh Generation for Massively Parallel-Based Analysis

S7B Introduction to Effective Parallel Computing, Part II

S9 Concurrent Programming with Pthreads

S10 Commodity-based Scalable Visualization

REGISTRATION / STORE

7:30AM – 6PM | **EXHIBIT HALL** | PAGE 175

EDUCATION PROGRAM

9AM – 6:30PM | **LEVEL 2, C SECTION** | PAGES 83 – 88

Registration
7:30AM – 9:30AM | PAGE 84

Welcome and Announcements
9:30AM – 10AM | PAGE 84

Introduction to Computational Science
Overview of Computational Modeling
Overview of Mathematical Modeling
Overview of Visualization
Overview of Curriculum Module Development
10AM – 6:30PM | PAGE 84 – 85

MEDIA ROOM

1PM – 4PM | **ROOM D 326** | PAGE 177

SC2000 NOTES:

MONDAY AT A GLANCE

NOVEMBER 6

MONDAY FULL-DAY TUTORIALS

M1 THROUGH M4 | 8:30AM – 5PM | PAGES 18 – 20

M1 **Performance Analysis and Prediction for Large-Scale Scientific Applications**

M2 **Parallel I/O for Application Developers**

M3 **Framework Technologies & Methods for Large Data Visualization**

M4 **Computational Biology and High Performance Computing**

MONDAY HALF-DAY TUTORIALS [AM]

M5A THROUGH M8 | 8:30AM – 12NN | PAGES 20 – 22

M5A **Application Building with XML: Standards, Tools, and Demos—Part I**

M6A **Parallel Programming with OpenMP: Part I, Introduction**

M7 **High-Speed Numerical Linear Algebra: Algorithms and Research Directions**

M8 **Parallel Programming for Cluster Computers**

MONDAY HALF-DAY TUTORIALS [PM]

M5B THROUGH M10 | 1:30PM – 5PM | PAGES 22 – 24

M5B **Application Building with XML: Standards, Tools, and Demos—Part II**

M6B **Parallel Programming with OpenMP: Part II, Advanced Programming**

M9 **Current and Emerging Trends in Cluster Computing**

M10 **Performance Tuning and Analysis for Grid Applications**

REGISTRATION / STORE

7:30AM – 9PM | **EXHIBIT HALL** | PAGE 175

EDUCATION PROGRAM

8:30AM – 7PM | **LEVEL 2, C SECTION** | PAGES 83 – 88

Modeling, Part 1
8:30AM – 12NN | PAGE 85

Modeling, Part 2
1PM – 4:30PM | PAGE 85

Developing Curriculum Modules
4:30PM – 5:30PM | PAGE 86

Team Meetings
5:30PM – 7PM | PAGE 86

Gala Opening Reception, Exhibit Hall
7PM – 9PM | PAGE 86

MEDIA ROOM

9AM – 6PM | **ROOM D 326** | PAGE 177

GALA OPENING

7PM – 9PM | **EXHIBIT HALLS D & E** | PAGE 174

SC2000 NOTES:

gala opening reception—
industry exhibits,
research exhibits,
research gems

7PM – 9PM
EXHIBIT HALLS D/E
PAGE: 174

SC2000 gala opening

TUESDAY AT A GLANCE

NOVEMBER 7

REGISTRATION / STORE

7:30PM – 5:30PM | **EXHIBIT HALL** | PAGE 175

KEYNOTE ADDRESS

8:30AM – 10AM | **BALLROOM C** | PAGE 25

TECHNICAL PROGRAM

8:30AM – 5PM *(SEE HOURLY SCHEDULE)*

EXHIBITOR FORUM

10AM – 4:30PM | **ROOM D 268** | PAGES 93 – 104

EXHIBITS

10AM – 6PM | **EXHIBIT HALLS D & E** | PAGE 123 – 174

BOFS

5:30PM – 7PM | PAGE 105 – 111

MEDIA ROOM

9AM – 6PM | **ROOM D 326** | PAGE 177

EDUCATION PROGRAM

8:30AM – 9PM | **LEVEL 2, C SECTION** | PAGES 83 – 88

SC Conference Keynote Speaker
8:30AM – 10AM | PAGE 25 & 86

Modeling, Part 3
10:30AM – 12NN | PAGE 86

Modeling, Part 3 (continued)
1PM – 2:30PM | PAGE 86

Team Meetings/Open Lab
Exhibit Hall Guided Tours
2:30PM – 5PM | PAGE 86

Open Lab
7PM – 9PM | PAGE 86

8:30 – 10AM

KEYNOTE ADDRESS—STEVEN J. WALLACH

BALLROOM C | PAGE 25

"Petaflops in the Year 2009"

Steven J. Wallach, CenterPoint Venture Partners/Chiaro Networks

10:30 – 12NN

PANEL	PAPER	PAPER	MASTERWORKS
Venture Capital: Who Wants to be a Billionaire? **ROOM D 262/264** PAGE 77	**MPI** **ROOM D 271/273** PAGES 45 – 46	**Numerical Algorithms** **ROOM D 274** PAGES 47 – 48	**Computational BioSciences: Genomics** **ROOM D 263/265** PAGES 31 – 33

1:30 – 3PM

PAPER	PAPER	PAPER	MASTERWORKS
Scheduling **ROOM D 267** PAGES 48 – 49	**MPI/OpenMP** **ROOM D 271/273** PAGES 50 – 51	**Potpourri** **ROOM D 274** PAGES 51 – 52	**Computational Biochemistry and Drug Discovery** **ROOM D 263/265** PAGES 34 – 35

3:30 – 5PM

PAPER	PAPER	PAPER	MASTERWORKS
Cluster Infrastructure **ROOM D 271/273** PAGES 52 – 54	**QoS/Fault Tolerance** **ROOM D 274** PAGES 54 – 55	**Biomedical Applications** **ROOM D 267** PAGES 55 – 57	**Computing Platforms** **ROOM D 263/265** PAGES 35 – 36

11.07
keynote address
masterworks
panels
papers

SC2000 NOTES:

SC2000 technical program

WEDNESDAY AT A GLANCE

NOVEMBER 8

REGISTRATION / STORE

7:30PM – 5:30PM | **EXHIBIT HALL** | PAGE 175

TECHNICAL PROGRAM

8:30AM – 5PM *(SEE HOURLY SCHEDULE)*

RESEARCH GEMS/ESCAPE 2000 OPEN HOUSE

10AM – 11AM | PAGE 149

EXHIBITOR FORUM

10AM – 4:30PM | **ROOM D 268** | PAGES 93 – 104

EXHIBITS

10AM – 6PM | **EXHIBIT HALLS D & E** | PAGE 123 – 174

BOFS

5:30PM – 7PM | PAGE 105 – 111

MEDIA ROOM

9AM – 6PM | **ROOM D 326** | PAGE 177

EDUCATION PROGRAM

8:30AM – 9PM | **LEVEL 2, C SECTION** | PAGES 83 – 88

Web-based Tools *(concurrent sessions, pick any two of the six)*
ChemViz
Interactivate
Physics Tools
BiologyWorkbench
The Bowfin Project
GirlTech
8:30AM – 12NN | PAGES 86 – 87

Curriculum Integration and Assessment
1PM– 2:30PM | PAGE 87

Team Meetings
3PM– 5PM | PAGE 87

Open Lab
7PM – 9PM | PAGE 87

11.08
invited speakers
masterworks
ieee award winners
panels
papers

8:30 – 10AM

STATE-OF-THE-FIELD TALKS—TOM STERLING, EUGENE SPAFFORD

BALLROOM C | PAGES 26 – 28

COTS Cluster Systems for High Performance Computing—Sterling

A Small Dose of Infosec—Spafford

10:30 – 12NN

PAPER	PAPER	PAPER	MASTERWORKS
Applications I **ROOM D 274** PAGES 57 – 58	**Visualization** **ROOM D 271/273** PAGES 59 – 60	**Compiler Optimization** **ROOM D 267** PAGES 60 – 61	**IEEE Award Winners** **ROOM D 263/265** PAGE 36

1:30 – 3PM

PAPER	PAPER	PAPER	MASTERWORKS
Applications II **ROOM D 274** PAGES 61 – 63	**Networking** **ROOM D 267** PAGES 63 – 64	**Hardware-Based Tools** **ROOM D 271/273** PAGES 64 – 65	**Real World Scalable Computing I** **ROOM D 263/265** PAGES 36 – 38

3:30 – 5PM

PANEL	PAPER	PAPER	MASTERWORKS
Is TCP an Adequate Protocol for High Performance Computing Needs? **ROOM D 262/264** PAGES 77 – 78	**Gordon Bell I** **ROOM D 271/273** PAGES 66 – 67	**Parallel Programming** **ROOM D 274** PAGES 67 – 69	**Real World Scalable Computing II** **ROOM D 263/265** PAGES 38 – 39

SC2000 NOTES:

SC2000 technical program

THURSDAY AT A GLANCE

NOVEMBER 9

REGISTRATION / STORE

7:30PM – 5PM | **EXHIBIT HALL** | PAGE 175

TECHNICAL PROGRAM

8:30AM – 5PM *(SEE HOURLY SCHEDULE)*

RESEARCH GEMS/ESCAPE 2000 OPEN HOUSE

10AM – 11AM | PAGE 149

EXHIBITOR FORUM

10AM – 4PM | **ROOM D 268** | PAGES 93 – 104

EXHIBITS

10AM – 4PM | **EXHIBIT HALLS D & E** | PAGE 123 – 174

MEDIA ROOM

9AM – 5PM | **ROOM D 326** | PAGE 177

EDUCATION PROGRAM

8:30AM – 2PM | **LEVEL 2, C SECTION** | PAGES 83 – 88

Outreach/Mentoring Goals
8:30AM – 9:30AM | PAGE 87

Module Presentations by Teacher Teams
10AM – 11:30AM | PAGE 87

Module Presentations by Teacher Teams
12:30PM – 1:30PM | PAGE 87

Education Program Evaluation Wrap-up
1:30PM – 2PM | PAGE 87

8:30 – 10AM

STATE-OF-THE-FIELD TALKS—MARGARET WRIGHT, JIM BROWNE

BALLROOM C | PAGES 28 – 30

Numbers, Lots of Numbers, And Insight, too: Scientific Computing 2000—Wright

Parallel/Distributed Programming: Research Success—Application Failure?—Browne

10:30 – 12NN

PAPER	PAPER	PAPER	MASTERWORKS
Software Tools **ROOM D 274** PAGES 69 – 70	**Data Grid** **ROOM D 267** PAGES 70 – 71	**Gordon Bell II** **ROOM D 271/273** PAGES 71 – 73	**Supercomputing Trends in MCAE I** **ROOM D 263/265** PAGES 39 – 41

1:30 – 3PM

PLENARY SESSION—AWARDS / GOEFF VOEKLER

BALLROOM C | PAGES 41 – 42 & 73

On the Scale and Performance of Cooperative Web Proxy Caching—Voekler

3:30 – 5PM

PANEL	PAPER	PAPER	MASTERWORKS
Petaflops Around The Corner: When? How? Is It Meaningful? **ROOM D 262/264** PAGE 78	**Scientific Applications Support** **ROOM D 274** PAGES 74 – 75	**Grid Middleware** **ROOM D 271/273** PAGES 75 – 76	**Supercomputing Trends in MCAE II** **ROOM D 263/265** PAGES 42 – 43

6:30 – 10PM

SC2000 CONFERENCE RECEPTION—TECHNICAL REGISTRANTS

Lone Star Park
Buses depart from the Dallas Convention Center to Lone Star Park: 6PM – 10:30PM

11.09
masterworks
panels
papers
awards
invited speakers
sc2000 reception

SC2000 NOTES:

FRIDAY AT A GLANCE

NOVEMBER 10

REGISTRATION / STORE

8AM – 11AM | **EXHIBIT HALL** | PAGE 175

PANELS

8:30 – 10AM	**Convergence at the Extremes: Computational Science meets Networked Sensors** **BALLROOM C 1/2** | PAGE 79 **Computational Grids: A solution looking for a problem?** **BALLROOM C 3/4** | PAGE 79
10:30 – 12NN	**Open Source: IP in the Internet Era** **BALLROOM C 1/2** | PAGES 79 – 80 **MegaComputers** **BALLROOM C 3/4** | PAGE 80

SC2000 EVALUATION FORMS

SC2000 Evaluation Forms accepted at the SC2001 booth.

See you at SC2001!

Denver, Colorado • November 10 – 16, 2001

11.10
panels
evaluation forms

SC2000 TUTORIAL PROGRAM

Tutorials Co-Chair: Valerie Taylor, Northwestern University
Tutorials Co-Chair: Michelle Hribar, Pacific University

This year's tutorials include exciting offerings in new topics, such as mesh generation, XML, parallel programming for clusters, numerical computing in Java, and management of large scientific datasets, along with the return of some of the most requested presenters from prior years, with new and updated materials. In addition, we offer some of the full day tutorials as two half-day tutorials (denoted by the numbering with A and B), thereby increasing the number of half-day tutorials. We have a total of nine full-day and 15 half-day tutorials covering 20 topics. Attendees also have the opportunity for an international perspective on topics through the tutorials on large data visualization, cluster computing, performance analysis tools and numerical linear algebra. Separate registration is required for tutorials; tutorial notes and luncheons will be provided on site (additional tutorial notes will be sold on site). A One- or Two-day Tutorial Passport allows attendees the flexibility to attend multiple tutorials.

SUNDAY FULL-DAY

SUNDAY
NOVEMBER 5
FULL-DAY
8:30AM – 5PM

SI

Using MPI-2: A Tutorial on Advanced Features of the Message-Passing Interface

William Gropp, Ewing (Rusty) Lusk, Rajeev S. Thakur, Argonne National Laboratory

20% Introductory • 40% Intermediate • 40% Advanced

This tutorial will describe how to use MPI-2, the collection of advanced features that were added to MPI (Message-Passing Interface) by the second MPI Forum. These features include parallel I/O, one-sided communication, dynamic process management, language interoperability, and some miscellaneous features.

Implementations of MPI-2 are beginning to appear. A few vendors have completed implementations; other vendors and research groups have implemented subsets of MPI-2, with plans for complete implementations. This tutorial will explain how to use MPI-2 in practice, particularly, how to use MPI-2 in a way that results in high performance. We will present each feature of MPI-2 in the form of a series of examples (in C, Fortran, and C++), starting with simple programs and moving on to more complex ones. We assume that attendees are familiar with the basic message-passing concepts of MPI-1.

SUNDAY
NOVEMBER 5
FULL-DAY
8:30AM – 5PM

S2 An Introduction to High Performance Data Mining

Robert L. Grossman, Magnify, Inc. and University of Illinois at Chicago, Vipin Kumar, University of Minnesota

50% Introductory • 30% Intermediate • 20% Advanced

Data mining is the semi-automatic discovery of 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 the tutorial is to provide researchers, practitioners, and advanced students with an introduction to data mining. The focus will be on basic techniques and algorithms appropriate for mining massive data sets using approaches from high performance computing. There are now parallel versions of some of the standard data mining algorithms, including tree-based classifiers, clustering algorithms, and association rules. We will cover these algorithms in detail and as well as some general techniques for scaling data mining algorithms. In addition, we will give an introduction to some of the data mining algorithms which are used in the recommended systems that are becoming important in e-business. The tutorial will include several case studies involving mining large data sets, from 10-1000 Gigabytes in size. The case studies will be from science, engineering & e-business.

SUNDAY
NOVEMBER 5
FULL-DAY
8:30AM – 5PM

S3 Design and Analysis of High Performance Clusters

Robert Pennington, NCSA, Patricia Kovatch, Barney Maccabe, David Bader, UNM

25% Introductory • 50% Intermediate • 25% Advanced

The National Computational Science Alliance (the Alliance) has created several production NT and Linux superclusters for scientists and researchers to run a variety of parallel applications. The goal of this tutorial is to bring together researchers in this area and to share the latest information on the state of high-end commodity clusters. We will discuss details of the design, implementation and management of these systems and demonstrate some of the current system monitoring and management tools. A wide variety of applications and community codes run on these superclusters. We will examine several of these applications and include details on porting and application development tools on both NT and Linux. We will also discuss how to use these tools to tune the system and applications for optimal performance.

S4

High Performance Numerical Computing in Java: Compiler, Language, and Application Solutions

Manish Gupta, Samuel P. Midkiff, Jose E. Moreira,
IBM T.J. Watson Research Center

15% Introductory • 65% Intermediate • 20% Advanced

SUNDAY
NOVEMBER 5
FULL-DAY
8:30AM – 5PM

There has been an increasing interest in using Java for the development of high performance numerical applications. Although Java has many attractive features—including reliability, portability, a clean object-oriented model, well defined floating point semantics and a growing programmer base—the performance of current commercial implementations of Java in numerical applications is still an impediment to its wider adoption in the performance-sensitive field. In this tutorial we will describe how standard libraries and currently proposed Java extensions will help in achieving high performance and writing more maintainable code. We will also show how Java virtual machines can be improved to provide near-Fortran performance. The proposals of the Java Grande Forum Numerics Working Group, which include a true multidimensional array package, complex arithmetic and new floating point semantics will be discussed. Compiler technologies to be addressed include array bounds and null pointer check optimizations, alias disambiguation techniques, semantic expansion of standard classes and the interplay of static and dynamic models of compilation. We will also discuss the importance of language, libraries and compiler codesign. The impact of these new technologies on compiler writers, language designers, and application developers will be described throughout the tutorial.

S5

Performance Analysis and Tuning of Parallel Programs: Resources and Tools

Barton Miller, University of Wisconsin-Madison, Michael Gerndt, Technical University Munich, Germany, Bernd Mohr, Research Centre Juelich, Germany

50% Introductory • 25% Intermediate • 25% Advanced

SUNDAY
NOVEMBER 5
FULL-DAY
8:30AM – 5PM

This tutorial will give a comprehensive introduction into the theory and practical application of the performance analysis, optimization and tuning of parallel programs on currently used high-end computer systems like the IBM SP, SGI Origin, and CRAY T3E as well as clusters of workstations. We will introduce the basic terminology, methodology, and techniques of performance analysis and give practical advice on how to use these in an effective manner. Next we describe vendor, third party, and research tools available for these machines along with practical tips and hints for their usage. We show how these tools can be used to diagnose and locate typical performance bottlenecks in real-

world parallel programs. Finally, we will give an overview of Paradyn, an example of a state-of-the-art performance analysis tool that can be used for parallel programs of today. The presentation will include the Performance Consultant that automatically locates the performance bottlenecks of user codes. This presentation will be concluded with a live, interactive demonstration of Paradyn.

SUNDAY HALF-DAY—AM

SUNDAY
NOVEMBER 5
HALF-DAY AM
8:30AM – 12NN

S6A Mesh Generation for High Performance Computing Part I: An Overview of Unstructured and Structured Grid Generation Techniques

Steven J. Owen, Patrick Knupp, Sandia National Laboratories

100% Introductory • 0% Intermediate • 0% Advanced

Mesh generation plays a vital role in computational field simulation for high performance computing. The mesh can tremendously influence the accuracy and efficiency of a simulation. Part I of this tutorial will provide an overview of the principal techniques currently in use for constructing computational grids for both unstructured and structured techniques. For unstructured techniques, Delaunay, advancing front and octree methods will be described with respect to triangle and tetrahedral elements. An overview of current quadrilateral and hexahedral methods will be provided, including medial axis, paving, q-morph, sub-mapping, plastering, sweeping and whisker weaving as well as mixed element methods such as hex-tet and h-morph. A survey of some of the unstructured codes currently available will also be provided. For structured techniques, the idea of a mapping from a logical to a physical domain will be discussed. Transfinite interpolation, Lagrange and Hermite interpolation techniques will be described. A one-dimensional problem will be used as an example to introduce basic ideas in grid generation such as grid generation PDE's, optimization, and variational techniques. In addition, basic approaches to two-dimensional grid generation, such as algebraic, conformal mapping, elliptic, and hyperbolic, will be presented. Application of structured grid generation techniques to curves and surfaces as well as adaptive methods will also be described.

SUNDAY
NOVEMBER 5
HALF-DAY AM
8:30AM – 12NN

S7A Introduction to Effective Parallel Computing, Part I

Quentin F. Stout, Christiane Jablonowski, Univ. of Michigan

75% Introductory • 25% Intermediate • 0% Advanced

Effective parallel computing is one of the key solutions to today's computational challenges. This two-part tutorial will provide a comprehensive and practical overview of

parallel computing, emphasizing those aspects most relevant to the user. It is designed for new users, managers, students, and people needing a general overview of parallel computing. The tutorial discusses both hardware and software aspects, with an emphasis on standards, portability and systems that are now (or soon will be) commercially or freely available. Systems examined range from low-cost clusters to highly integrated supercomputers. Part I surveys basic parallel computing concepts and terminology, such as scalability and cache coherence, and illustrates fundamental parallelization approaches using examples from engineering, scientific and data intensive applications. These real-world examples are targeted at distributed memory systems, using the message passing language MPI, and at shared memory systems, using the compiler directive standard OpenMP. Both parallelization approaches will be briefly outlined. The tutorial shows step-by-step parallel performance improvements, and discusses some of the software engineering aspects of the parallelization process. Furthermore, pointers to the literature and web-based resources will be provided. This tutorial can serve as an introduction to specialized programming tutorials.

S8 Tools and System Support for Managing and Manipulating Large Scientific Datasets

Joel Saltz, University of Maryland, Johns Hopkins School of Medicine, Alan Sussman, Tahsin Kurc, University of Maryland

30% Introductory • 60% Intermediate • 10% Advanced

SUNDAY
NOVEMBER 5
HALF-DAY AM
8:30AM – 12NN

This tutorial will address the design, implementation and use of systems for managing and manipulating very large datasets, both on disk and in archival storage. The datasets we target are generated through large scale simulations or gathered by advanced sensors, such as those attached to satellites or microscopes. These datasets are large (hundreds of gigabytes to many terabytes) and typically represent physical quantities, measurements or composited images in a physical or attribute space. Two systems will be described in detail, the Active Data Repository (ADR) and DataCutter. ADR is designed to optimize storage and processing of disk-based large datasets on a parallel machine or network of workstations, while DataCutter is designed to provide support for subsetting and filtering operations on datasets stored in archival (tertiary) storage systems in a Grid environment. The overall design of ADR and the interfaces for customizing ADR for particular data intensive applications, will be explained, and an example application will be used to illustrate the customization. The customization includes storing and indexing datasets into ADR, and

providing user-defined processing functions for the end application. Similarly, the design and current implementation of the DataCutter services will be described, and examples of filter-based applications will be discussed. The relationship of ADR and DataCutter to other systems software for data intensive computing will also be addressed. Such systems include the ISI/Argonne Globus Metacomputing toolkit, the UTK NetSolve network-based computational server and the SDSC Storage Resource Broker (SRB).

SUNDAY HALF-DAY—PM

SUNDAY
NOVEMBER 5
HALF-DAY PM
1:30PM – 5PM

S6B Mesh Generation for High Performance Computing Part II: Mesh Generation for Massively Parallel-Based Analysis

Scott Mitchell, Patrick Knupp, and Timothy Tautges, Sandia National Laboratories

50% Introductory • 50% Intermediate • 0% Advanced

Part II will focus on specific application of mesh generation techniques to high performance computing. Topics discussed will include advanced hexahedral algorithms, mesh quality and mesh generation issues related to massively-parallel based analysis. Advanced hexahedral mesh generation algorithms for meshing assembly geometries will be described; these algorithms are found in the CUBIT Mesh Generation Toolkit and other mesh generation packages. Also discussed will be additional hexahedral mesh generation research ideas. Basic mesh quality requirements will be described for finite element meshes including shape and size metrics for both simplicial and non-simplicial element types. Mesh quality can be improved by various node-movement strategies; their use in mesh sweeping and morphing algorithms will be described. Mesh quality metrics, used to devise discrete objective functions, will also be discussed. Advances in massively parallel and high performance computing have made possible computational simulation at much higher fidelity and finer resolution. Issues specific to larger analysis will be introduced, including techniques for handling model complexity, a team-based approach to generating meshes, tools for generating meshes in pieces and assembling the pieces into a larger mesh, and preparing the mesh for input to massively parallel-based analysis.

S7B Introduction to Effective Parallel Computing, Part II

Quentin F. Stout, Christiane Jablonowski, University of Michigan

50% Introductory • 50% Intermediate • 0% Advanced

SUNDAY
NOVEMBER 5
HALF-DAY PM
1:30PM – 5PM

This tutorial will provide broader and deeper insight into the iterative process of converting a serial program into an increasingly efficient, and correct, parallel program. This tutorial assumes the basic background knowledge of parallel computing concepts and terminology presented in Part I. Using examples from large-scale engineering and scientific applications, Part II will discuss the steps necessary to achieve high performance on distributed memory, shared memory and vector parallel machines. We will give an overview of techniques for code optimization, load balancing, communication reduction and efficient use of cache. The tutorial will include up-to-date performance analysis tools, showing how they can help diagnose and locate typical bottleneck situations in parallel applications and provide hints for tuning. In addition, aspects such as the user's view of system's software and principle life-cycle concerns with parallel software will be addressed. Overall, the tutorial will give an overview of the primary parallelization options available, explaining how they are used in real-world applications and what they are most suitable for. These guidelines will help users make intelligent planning decisions when selecting among the various software approaches and hardware platforms.

S9 Concurrent Programming with Pthreads

Clay P. Breshears, Henry A. Gabb, Kuck & Associates, an Intel Company

65% Introductory • 35% Intermediate • 0% Advanced

SUNDAY
NOVEMBER 5
HALF-DAY PM
1:30PM – 5PM

Multithreading is becoming more prevalent with the increasing popularity of symmetric multiprocessors (SMPs). Multithreading allows programmers to utilize shared memory hardware to its fullest. Pthreads is the POSIX standard library for multithreading and is available on a wide range of platforms. The Pthreads library consists of over 60 functions governing thread creation and management, synchronization, and scheduling. This tutorial will cover design issues involved in concurrent and multithreaded programming, using Pthreads as a practical means of implementation. Before laying a foundation in concurrency, the tutorial will introduce a core of the most useful Pthreads functions. Each function will be discussed in detail with example codes to illustrate usage. Classic models (e.g., monitors, rendezvous, and producer/consumer) will illustrate the use of threads to express concurrent tasks as well as the pitfalls of race conditions and deadlock.

SUNDAY
NOVEMBER 5
HALF-DAY PM
1:30PM – 5PM

S10 Commodity-based Scalable Visualization

Constantine J. Pavlakos, Sandia National Laboratories, Randall Frank, Lawrence Livermore National Laboratory, Patrick Hanrahan, Stanford University, Kai Li, Princeton University, Alan Heirich, Compaq Tandem Labs, Allen McPherson, Los Alamos National Laboratory

40% Introductory • 40% Intermediate • 20%Advanced

The DOE's ASCI Program is constructing massive compute platforms for the purpose of enabling extremely complex computational simulations. Further, the ASCI/VIEWS program is working to develop data management, data exploration, and visualization technologies that are matched to ASCI's compute capabilities. To do so, technologies that scale to the power of thousands of today's highest performing graphics systems must be developed, and super-resolution display systems are needed that enable the visual comprehension of intricate details in high-fidelity data. Cost-effectiveness is also an important pragmatic consideration. This tutorial will motivate the investigation of cluster-based graphics systems, introduce the participant to information regarding the construction of such clusters, address various issues related to the components used, introduce the participant to how parallel rendering can be achieved on the clustered architectures, and provide an overview of results that have been achieved. The status of efforts to develop such systems by the three ASCI national laboratories and certain external partners will also be presented. Upon completing the tutorial, the participant should have a much better understanding of what it takes to construct such a system, what features they offer, whether such systems show any promise for scalable visualization, and what challenging issues remain to be addressed.

MONDAY FULL-DAY

MONDAY
NOVEMBER 6
FULL-DAY
8:30AM – 5PM

M1 Performance Analysis and Prediction for Large-Scale Scientific Applications

Adolfy Hoisie, Los Alamos National Laboratory, Harvey J. Wasserman, Los Alamos National Laboratory

30% Introductory • 50% Intermediate • 20% Advanced

Performance is the most important criterion for a supercomputer. But how do you measure performance? We will present a methodical, simplified approach to analysis and modeling of large-scale, parallel, scientific applications. Various techniques (modeling, simulation, queuing theory), will be discussed so as to become a part of the application developer's toolkit. We will introduce rigorous metrics for serial and parallel performance and analyze the single most important single-processor bottleneck—the

memory subsystem. We will demonstrate how to obtain diagnostic information about memory performance of codes and how to use such information to bound achievable performance. Commonly-utilized techniques for performance optimization of serial and parallel Fortran codes will also be presented. Finally, we will discuss analytical modeling of application scalability using ASCI codes as examples. No particular machine will be emphasized; rather we will consider RISC processors and widely utilized parallel systems, including clusters of SGI Origin2000s, IBM SP2 and CRAY T3E.

M2 Parallel I/O for Application Developers

MONDAY
NOVEMBER 6
FULL-DAY
8:30AM – 5PM

John M May, Lawrence Livermore National Laboratory
50% Introductory • 50% Intermediate • 0% Advanced

This tutorial will present parallel I/O techniques for developers of scientific applications. Because the design of storage devices and file systems profoundly affects I/O performance, the course will begin with a brief review of these topics. It will then proceed to examine the I/O patterns that are common in large scientific applications and show how these patterns affect I/O performance. Next, it will look at a variety of techniques that have been developed to improve performance for common access patterns and discuss their pros and cons. Attendees will learn how to put these techniques into practice using modern I/O interfaces such as MPI-IO and HDF5. We will also discuss two specialized forms of I/O used in parallel computing: checkpointing and data staging for out-of-core problems. The tutorial will conclude with a discussion of current research in the area of scientific data management, including data mining.

M3 Framework Technologies & Methods for Large Data Visualization

MONDAY
NOVEMBER 6
FULL-DAY
8:30AM – 5PM

W. T. Hewitt, University of Manchester, I. Curington, Advanced Visual Systems Inc.
20% Introductory • 70% Intermediate • 10% Advanced

This tutorial will address large data visualization issues in the context of commercial visualization tool development. A review of techniques for multidimensional data visualization will be followed by case studies from CEM, CFD, VLSI, Medicine, and Geophysics. An artifact of this type of visualization is that the visualization task itself becomes a consumer of HPC resources. The second part of the tutorial is concerned with the issues of implementing these techniques in a multiprocessor environment, and improving the performance of current visualization systems. A range of technical areas will be discussed, including experimental research and production algorithm development. Both cur-

rent research and future challenges facing visualization system vendors will be discussed. Attendees at the tutorial will gain an understanding of the issues underlying visualization in a parallel and distributed environment including: familiarity with domain decomposition methods and parallelisation techniques; knowledge of the principles of volume, flow, and multidimensional visualization; ability to use distributed computation to enable accurate and timely visualization of large complex datasets; and familiarity with the latest developments in visualization and HPC systems.

MONDAY
NOVEMBER 6
FULL-DAY
8:30AM – 5PM

M4 Computational Biology and High Performance Computing

Manfred Zorn, Sylvia Spengler, NERSC/CBCG, Horst Simon, NERSC, Craig A. Stewart, Indiana University, Inna Dubchak, NERSC/CBCG

40% Introductory • 40% Intermediate • 20% Advanced

The pace of extraordinary advances in molecular biology has accelerated in the past decade due in large part to discoveries coming from genome projects on human and model organisms. The advances in the genome project so far, happening well ahead of schedule and under budget, have exceeded any dreams by its protagonists, let alone formal expectations. Biologists expect the next phase of the genome project to be even more startling in terms of dramatic breakthroughs in our understanding of human biology, the biology of health and of disease. Only today can biologists begin to envision the necessary experimental, computational and theoretical steps necessary to exploit genome sequence information for its medical impact, its contribution to biotechnology and economic competitiveness, and its ultimate contribution to environmental quality. High performance computing has become one of the critical enabling technologies, which will help to translate this vision of future advances in biology into reality. Biologists are increasingly becoming aware of the potential of high performance computing. This tutorial will introduce the exciting new developments in computational biology and genomics to the high performance computing community.

MONDAY HALF-DAY—AM

MONDAY
NOVEMBER 6
HALF-DAY AM
8:30AM – 12NN

M5A Application Building with XML: Standards, Tools, and Demos—Part I

Bertram Ludaescher, Richard Marciano, UCSD/SDSC

40% Introductory • 40% Intermediate • 20% Advanced

This tutorial will guide participants through the maze of emerging XML standards and focus on practical areas where the use of XML can have immediate benefits for application development. The general theme is "What you

need to know for rearchitecting your HPC with XML under the hood". This tutorial provides a "jump-start" to XML ("what you always wanted to know about XML") and includes a roadmap to the XML universe, coverage of core standards and technology, and how to (re-)architect "XML-enabled" applications.

M6A Parallel Programming with OpenMP: Part I, Introduction

MONDAY
NOVEMBER 6
HALF-DAY AM
8:30AM – 12NN

Rudolf Eigenmann, Ph.D., Purdue University, Tim Mattson, Ph.D, Intel Corp.

75% Introductory • 20% Intermediate • 5% Advanced

OpenMP is an Application Programming Interface for directive-driven parallel programming of shared memory computers. Fortran, C and C++ compilers supporting OpenMP are available for Unix and NT workstations. Most vendors of shared memory computers are committed to OpenMP making it the de facto standard for writing portable, shared memory, parallel programs. This tutorial will provide a comprehensive introduction to OpenMP. We will start with basic concepts to bring the novice up to speed. We will then present a few more advanced examples to give some insight into questions that come up for experienced OpenMP programmers. Over the course of the morning, we will discuss the following: the OpenMP parallel programming model and its specification in Fortran, C and C++; examples of OpenMP programs from scientific/engineering applications; and the status of OpenMP compilers and tools.

M7 High-Speed Numerical Linear Algebra: Algorithms and Research Directions

MONDAY
NOVEMBER 6
HALF-DAY AM
8:30AM – 12NN

Jack Dongarra, University of Tennessee, Iain Duff, Rutherford Lab, Danny Sorensen, Rice University

20% Introductory • 50% Intermediate • 30% Advanced

Present computers, even workstations, allow the solution of very large scale problems in science and engineering. Most often a major part of the computational effort goes in solving linear algebra subproblems. We will discuss a variety of algorithms for these problems, indicating where each is appropriate and emphasizing their efficient implementation. Many of the sequential algorithms used satisfactorily on traditional machines fail to exploit the architecture of advanced computers. We briefly review some of the features of modern computer systems and illustrate how the architecture affects the potential performance of linear algebra algorithms. We will consider recent techniques devised to utilize advanced architectures more fully, especially the design of the Level 1, 2, and 3 BLAS. We

will highlight the LAPACK package which 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. We will give an introduction to this field and guidelines on the selection of appropriate software. We will consider both direct methods and iterative methods of solution, including some recent work that can be viewed as a hybrid of the two. Finally, we address the challenge facing designers of mathematical software in view of the development of highly parallel computer systems. We shall discuss ScaLAPACK, a project to develop and provide high performance scalable algorithms suitable for highly parallel computers.

MONDAY
NOVEMBER 6
HALF-DAY AM
8:30AM – 12NN

M8 Parallel Programming for Cluster Computers

David A. Bader, University of New Mexico, Bruce Hendrickson, Steve Plimpton, Sandia National Laboratories

25% Introductory • 50% Intermediate • 25% Advanced

The price/performance benefits inherent in commodity cluster computers are attracting the attention of a wide range of researchers, many of whom have not traditionally been involved in parallel computing. Clusters are likely to be the computational platforms of choice over the next decade not just for computer scientists, but for disciplinary researchers in fields such as bioinformatics, astrophysics and economics. Unfortunately, these clusters are among the most challenging parallel computers to use effectively. In this tutorial, we will describe the performance challenges inherent to commodity clusters, including poor communication performance, heterogeneity, and 2-level hardware (clusters of 2- and 4-way SMP nodes). We will outline an approach for designing scientific algorithms that works well for clusters and present some case studies. The approach is based on a distributed-memory message-passing model with an emphasis on load-balance, minimal communication, and latency-tolerant algorithms. Finally we'll also highlight tools and libraries currently available for improving parallel programming productivity.

MONDAY HALF-DAY—PM

MONDAY
NOVEMBER 6
HALF-DAY PM
1:30PM – 5PM

M5B Application Building with XML: Standards, Tools, and Demos—Part II

Richard Marciano, Bertram Ludaescher, UCSD/SDSC

20% Introductory • 30% Intermediate • 50% Advanced

This tutorial will include a brief XML overview; however, it will mostly focus on how multiple standards can be used

and woven together to help build useful applications that can benefit HPC application developers. This tutorial naturally extends Part I, however it is packaged independently and is self-contained. If you already know how to design, create, store, and massage your data with XML, XPath, XSLT and the like, then you are ready for "more mileage with XML." This tutorial focuses on applications including website management, directory services and wireless computing with XML.

M6B Parallel Programming with OpenMP: Part II, Advanced Programming

MONDAY
NOVEMBER 6
HALF-DAY PM
1:30PM – 5PM

Rudolf Eigenmann, Purdue University, Tim Mattson, Intel Corp.

10% Introductory • 50% Intermediate • 40% Advanced

OpenMP is rapidly becoming the programming model of choice for shared-memory machines. After a very brief overview of OpenMP basics we will move on to intermediate and advanced topics, such as advanced OpenMP language features, traps that programmers may fall into, and a more extensive outlook on future OpenMP developments. We will also discuss mixing OpenMP with message passing applications written in MPI. We will present many examples of OpenMP programs and discuss their performance behavior. Over the course of the afternoon, we will discuss the following: a brief overview of the OpenMP parallel programming model in C, C++, and Fortran; problems and solutions in OpenMP programming; advanced examples of OpenMP programs and their performance; and future developments for OpenMP.

M9 Current and Emerging Trends in Cluster Computing

MONDAY
NOVEMBER 6
HALF-DAY PM
1:30PM – 5PM

Mark Baker, University of Portsmouth, Rajkumar Buyya, Monash University, Melbourne Australia, Jack Dongarra, University of Tennessee

50% Introductory • 30% Intermediate • 20% Advanced

The commercial success of clusters has pushed them into the vanguard of general purpose computing. They have now permeated all spheres of the computing industry, from the traditional science and engineering field through to the retail and commercial marketplace. This commodity-driven computing platform is advancing at a tremendous pace in terms of both new and emerging hardware and the associated software tools and environments. Clusters are now the platform of choice for providing computing services to a huge range of diverse applications. This tutorial discusses the current and emerging trends in cluster computing. In particular, we detail the current and emerging technologies in areas such as system architecture, networking, software environments, systems configuration

and management tools as well as application libraries and utilities. In the second half of the tutorial, we review four successfully deployed cluster systems that are being used in commerce, industry and research environments. Finally, we summarize our findings, drawing a number of conclusions about current clusters and then briefly consider emerging technology trends and how these will influence clusters of the future.

MONDAY
NOVEMBER 6
HALF-DAY PM
1:30PM – 5PM

M10 Performance Tuning and Analysis for Grid Applications

Brian Tierney, Lawrence Berkeley National Laboratory, Rich Wolski, University of Tennessee, Dan Gunter, Lawrence Berkeley National Laboratory, Martin Swany, University of Tennessee

25% Introductory • 75% Intermediate • 0% Advanced

For distributed application developers, achieving high-performance in a Grid environment can be especially challenging due to the fact that the bottlenecks can be in any of a number of places, such as the hosts, networks, operating systems, applications, and so on. Therefore one must monitor all system components and instrument all software. This tutorial will discuss what should be monitored, and describe some tools that can be used to perform the monitoring and manage the large volumes of performance information that results. It will also discuss scalable techniques for instrumenting applications and system level resources. Both the problem of scaling end-to-end network performance monitoring and the tradeoff between monitor intrusiveness and monitor accuracy will be discussed. The tutorial will show how to use the monitored data to do performance analysis and to predict resource load and availability dynamically. Further, we will discuss some general techniques for improving application performance in a high-speed WAN environment.

KEYNOTE ADDRESS

"Petaflops in the Year 2009"
Steven J. Wallach, CenterPoint Venture Partners/Chiaro Networks

TUESDAY
NOVEMBER 7
8:30AM – 10AM
BALLROOM C

Steven J. Wallach co-founded Convex Computer Corporation, along with Robert J. Paluck, former chairman and CEO, in 1982. Wallach was the chief designer of the Convex C-Series, the world's first affordable supercomputer, as well as the Exemplar Scalable Parallel Processor (SPP), HP/Convex's.

Wallach is currently an advisor to CenterPoint (www.centerpointvp.com) Venture Partners, Dallas, Texas and Vice President of Technology of Chiaro Networks (www.chiaro.com), Richardson, Texas. Wallach may be best known outside HPCN circles as the Data General engineer who was the principal architect of the 32-bit Eclipse MV superminicomputer series as described by Pulitzer Prize winner Tracy Kidder in The Soul of A New Machine.

Wallach holds 33 patents in various areas of computer design and held a joint appointment in the Graduate School of Management and Brown School of Engineering, Computer Science, Rice University for the 1998 and 1999 academic years. He is a member of the PITAC (Presidential Advisory Board on High Performance Computing, Communications, and Networking) and the advisory committee for the Hybrid Technology MultiThreaded Architecture (HTMT) a US DOD funded project to develop the concepts for a PETAFLOP computer. He is also a member of the National Academy of Engineering.

STATE-OF-THE-FIELD TALKS

Invited Speakers Chair: Patricia J. Teller, University of Texas, El Paso

All four talks will be in Ballroom C

The state-of-the-field talks give conference attendees the opportunity to hear distinguished leaders in four different fields of computer and computational sciences conduct brief tours of the state-of-the-art in these fields and explore possible ways in which the field will change in the coming years.

WEDNESDAY | 8:30AM

Chair: James R. McGraw, LLNL

WEDNESDAY
NOVEMBER 8
8:30AM
BALLROOM C

COTS Cluster Systems for High Performance Computing

Dr. Thomas Sterling
NASA Jet Propulsion Laboratory, High Performance Computing Group, and California Institute of Technology, Center for Advanced Computing Research

In recent years an alternative strategy to achieving high performance, which overcomes the combined problems of cost and architecture variability as well as stability of single suppliers, has emerged. A product of nearly a decade of applied research on workstation clusters, PC clusters, and non-dedicated LAN-connected user desktop and server facilities for cycle harvesting has yielded a rapidly maturing methodology for aggregating and employing low-to-moderate range computer systems in distributed complexes for both capacity and capability workload processing requirements.

COTS clusters are now having significant impact on the realm of processing once reserved to supercomputing. But not all applications are suitable for such loosely-coupled ensembles, and system software environments are still in evolution.

In this talk, Dr. Sterling will explore the history, methodologies, capabilities, and limitations of COTS clusters. He will examine in detail hardware component capabilities and configuration. Sterling will describe software systems and tools for cluster system programming and management, and will present performance and scaling data on successful applications as well as those demonstrating poorer suitability. Finally, Dr. Sterling will discuss examples of near-term research and technology trends that are likely to determine the future directions and capabilities of the next generation of COTS clusters for high-performance computing.

BIOGRAPHICAL SKETCH Dr. Sterling holds a joint appointment with NASA's Jet Propulsion Laboratory (JPL) and the California Institute of Technology (CalTech), serving as Principle Scientist in JPL's High Performance Computing group and Faculty Associate in CalTech's Center for Advanced Computing Research. He received his Ph.D. as a Hertz Fellow from MIT in 1984.

For the last 20 years, Sterling has engaged in applied research in parallel processing hardware and software systems for high-performance computing. He was a developer of the Concert shared-memory multiprocessor, the YARC static dataflow computer, and the Associative Template Dataflow computer concept, and has conducted extensive

studies of distributed shared-memory cache-coherent systems. In 1994, Sterling led the NASA Goddard Space Flight Center team that developed the first Beowulf-class PC clusters. Since 1994, he has been a leader in the national Petaflops initiative. He is the Principal Investigator for the interdisciplinary Hybrid Technology Multithreaded (HTMT) architecture research project sponsored by NASA, NSA, NSF, and DARPA, which involves a collaboration of more than a dozen cooperating research institutions. Sterling holds six patents, and was one of the winners of the 1997 Gordon Bell Prize for Price/Performance.

WEDNESDAY | 9:15AM

A Small Dose of Infosec
Eugene H. Spafford
Purdue University, www.cerias.purdue.edu/homes/spaf

WEDNESDAY
NOVEMBER 8
9:15AM
BALLROOM C

Melissa. ILOVEYOU. Web page defacement. Denial-of-service against e-commerce sites. Theft of laptops from the State Department. Information Warfare. Y2K. Napster. Hacking.

It is no longer possible to avoid stories of information loss, fraud, and compromise. Open any paper or magazine, or listen to the news (on-line as well as in the standard media) and stories relating to information security are sure to be present.

So, what is the current state of information security? Are things getting better, or are they getting worse? And what are the challenges that we are likely to see in the near future?

In this talk, Dr. Spafford will present some highlights of what is happening in information security, and what is yet to happen. He also will include some discussion of the nature of infosec-related challenges that we are likely to face—and few of them are based solely in technology.

BIOGRAPHICAL SKETCH Dr. Spafford is a professor of Computer Sciences at Purdue University, the university's Information Systems Security Officer, and Director of the Center for Education Research Information Assurance and Security (CERIAS). CERIAS is a campus-wide multi-disciplinary center, with a broadly-focused mission to explore issues related to protecting information and information resources. Spafford has written extensively about information security, software engineering, and professional ethics. He has published over 100 research articles and reports, and has written or contributed to over a dozen books; he serves on the editorial boards of most major infosec-related journals.

Spafford is a Fellow of the ACM, Fellow of the AAAS, senior member of the IEEE and is a charter recipient of the

Computer Society's Golden Core award. Among other activities, he is chair of the ACM's US Public Policy Committee, a member of the Board of Directors of the Computing Research Association, and is a member of the US Air Force Scientific Advisory Board. Dr. Spafford regularly serves as a consultant on information security and computer crime to law firms, major corporations, U.S. government agencies, and state and national law enforcement agencies around the world.

THURSDAY | 8:30AM

Chair: Pat Teller, UTEP

THURSDAY
NOVEMBER 9
8:30AM
BALLROOM C

Numbers, Lots of Numbers, And Insight, too: Scientific Computing 2000

Margaret H. Wright
Bell Laboratories, Lucent Technologies

Hamming's famous 1962 dictum, "The purpose of computing is insight, not numbers," has been quoted countless times, and no one doubts its truth. But most insights derived from computing depend on numbers—numbers being produced today at speeds and in quantities that were inconceivable in 1962. And certain crucial numbers, such as sensitivities and uncertainties, allow us to associate precision and measure with otherwise fuzzy insights.

In this talk, Dr. Wright will sketch a few directions in scientific computing that are helping us to obtain, as quickly and reliably as possible, numbers that are as accurate as necessary, to solve as many important problems as possible, all with the ultimate goal of insight and understanding.

BIOGRAPHICAL SKETCH Dr. Wright is head of the Scientific Computing Research Department in the Computing Sciences Research Center, Bell Laboratories, Murray Hill, New Jersey. She received a Ph.D. and M.S. in computer science and a B.S. in mathematics from Stanford University. From 1976–1988, she worked in the Department of Operations Research, Stanford University, and in 1988 she joined Bell Labs. Wright's research interests are numerical optimization, linear algebra, mathematical software, scientific computing, and the application of optimization to real-world problems. During 1995 and 1996, she served as president of the 9,000-member Society for Industrial and Applied Mathematics (SIAM) and in 1997 was elected to the US National Academy of Engineering. Dr. Wright currently chairs the Department of Energy's Advanced Scientific Computing Advisory Committee.

THURSDAY | 9:15AM

Parallel/Distributed Programming: Research Success—Application Failure?

J. C. Browne
The University of Texas at Austin
Department of Computer Sciences, browne@cs.utexas.edu

THURSDAY
NOVEMBER 9
9:15AM
BALLROOM C

Development of parallel/distributed programs that are correct, efficient, scalable, and portable across parallel/distributed execution environments remains the major bottleneck to effective application of parallel/distributed computational environments. Research in parallel programming has generated innovative concepts and potentially effective models for parallel/distributed programming. Some of these concepts and models have been given experimental implementations. A few have demonstrated significant potential for development of correct and efficient parallel/distributed programs.

Little of this research has had any impact upon the practice of parallel/distributed programming. The emerging technologies for compositional development of software have not yet made a major impact on parallel/distributed programming. Current practice of parallel/distributed programming is still dominated by ad hoc extensions to conventional sequential programming systems. We are still trying to fly into space using a horse and buggy.

In this talk, Dr. Browne will give one perspective on why past research has failed to have a major impact on practice, will propose requirements for broadly applicable effective parallel/distributed software system development, and will assess past and current research with respect to these requirements. He will suggest an approach for bringing parallel/distributed programming into the mainstream of software development and suggest a basis for a quantum leap in the practice of parallel/distributed software system development.

BIOGRAPHICAL SKETCH Dr. Browne is Professor of Computer Science and Physics and holds the Regents Chair #2 in Computer Sciences at The University of Texas at Austin (UT-Austin). He earned his Ph.D. in Chemical Physics at UT-Austin in 1960 and taught in the Physics Department at the university from 1960 through 1964. From 1965 through 1968, Browne was Professor of Computer Science and Director of the Computation Laboratory at Queens University in Belfast. In 1968, he rejoined UT-Austin as Professor of Physics and Computer Science. During 1968-69, 1971-75 and 1984-87, he served as Department Chair for Computer Science. Browne is a fellow of the Association for Computing Machinery

(ACM), the British Computer Society, the American Physical Society, and the American Association for the Advancement of Science; he has been Chairman of the ACM Special Interest Group on Operating Systems.

Browne has published approximately 100 papers in computational physics and 200 papers in computer science. He has been active in research on parallel/distributed programming and computation for twenty years, developing parallel/distributed programming systems and using parallel/distributed programming systems. Dr. Browne's research on parallel/distributed programming spans from graphical/visual programming environments (CODE) to libraries for parallel implementation of adaptive solutions of partial differential equations (DAGH, SDDA).

PLENARY / INVITED SPEAKERS MASTERWORKS

SC2000 plenary and invited speakers include the world's leaders in HPNC research and applications. The keynote address promises to be unique and memorable. The four State–of-the-Field talks will continue the tradition of critical assessments about recent progress and future challenges in key areas of computing sciences. This year the focus will be on Cluster Systems, Computer Security, Numerical Methods, and Parallel Programming. SC2000 is also beginning a new tradition, called Masterworks. The Masterworks track features Invited Speakers who will emphasize novel and innovative uses of HPNC in solving challenging problems in the real world. This track will highlight research-quality results that serve practical priorities. This year, the Masterworks track will include presentations on Computational Biology, Scalable Computing/Servers, Computer-Aided Engineering, Large Computing Platforms, and work by the winners of two prestigious IEEE computing awards. We hope that you find this track to be stimulating, and that it expands your understanding of just how broadly HPNC technology can benefit the larger scientific community and the general public.

It is expected that all of the plenary sessions will be webcast over the Internet, using the unique capabilities of SCinet 2000 to deliver high-quality audio and video broadcasts. Webcast events are indicated by this symbol in the Final Program and conference signage.

TUESDAY | 10:30AM – 12NN

TUESDAY
NOVEMBER 7
10:30AM
ROOM D 263/265

COMPUTATIONAL BIOSCIENCES: GENOMICS

Chair: Bill Blake, Compaq

From First Assembly Towards a New Cyber-Pharmaceutical Computing Paradigm

Sorin Istrail, Celera Genomics

The new science of whole genomics will fundamentally change how pharmaceutical companies pursue the vital challenge of developing new and better drugs. Target discovery, lead compound identification, pharmacology, toxicology and clinical trials are likely to merge with the science of bioinformatics into a powerful system for develop-

ing new pharmaceutical agents. It will be possible to simulate the action of new molecules or therapeutic programs against diverse metabolic pathways prior to preclinical testing. Thus, a paradigm of cyber-pharmaceutical testing will be available to the industry, speeding the selection of promising new agents, eliminating products that are likely to exhibit toxicity, and reducing the formidable costs and risks associated with the current paradigm of drug development.

We will report on Celera's design of a whole genome shotgun assembler and its application to the sequencing of the Drosophila and Human genomes. We will also present some of the major emerging computational challenges of the above paradigm in the exciting new areas of proteomics, structural genomics, expression profiling, SNPs and pharmacogenomics.

BIOGRAPHICAL SKETCH Sorin Istrail's work has focused on several areas of computer science and its applications to biology, physics and chemistry. He has worked on linguistics, automata theory, parallel algorithms and architectures, semantics of parallel programming languages, programming logic, complexity theory and derandomization, graph theory, voting theory, genomic mapping, protein folding, biomolecular sequence alignment, biomaterials, combinatorial chemistry, and proteomics. Recently, he resolved a longstanding open problem in statistical mechanics—the Three-Dimensional Ising Model Problem—showing the impossibility of deriving explicit thermodynamic formulas for every three-dimensional model. He is Executive Editor of the Journal of Computational Biology, Founder and General Vice-Chair of the RECOMB Conference Series, and Editor of the MIT Press Computational Molecular Biology Book Series.

Sorin Istrail has a PhD in computer science from University of Bucharest, Romania. After his immigration to the US, he was a visiting scientist at MIT. He joined Sandia Labs in 1992 where he held several positions, including most recently, Principal Senior Member of the Technical Staff. From 1992, he led the Sandia Labs research in genomics and structural genomics within the DOE MICS Computational Biology Project. In April 2000, he joined Celera Genomics, where he is Senior Director of Informatics Research.

Unveiling the Human Genome

Jill P. Mesirov, Whitehead/MIT Center for Genome Research

TUESDAY
NOVEMBER 7
11:15AM
ROOM D 263/265

We are living in an extraordinarily exciting time in the history of research in molecular biology. On June 26, 2000 the Human Genome Project public consortium announced an assembled working draft of the Human Genome. While work continues to complete a "finished" version of the genome with 99.99% accuracy and no gaps or ambiguities, the working draft is already enabling scientist to uncover the genome's mysteries. This work relies heavily on the use of sophisticated computational techniques.

We will discuss a number of the scientific and computational challenges that face the genomics community. These will be drawn from the identification and elucidation of function of the genes in the human genome, the discovery of gene pathways, the mining of gene expression data to aid in the diagnosis and treatment of disease, the search for disease related genes, and the study of human variation.

BIOGRAPHICAL SKETCH Jill P. Mesirov is Director of Bioinformatics and Research Computing at the Whitehead/MIT Center for Genome Research where she is responsible for the informatics, computational biology, and research computing program of the Center. She is also Adjunct Professor of Computer Science at Boston University. Mesirov spent many years working in the area of high performance computing and developing parallel algorithms relevant to problems that arise in science, engineering, and business applications. Her current research interest is the study and development of algorithms for computational biology in such areas as pattern discovery and recognition in gene expression data, genome analysis and interpretation, sequence homology searching, protein secondary structure prediction and classification, molecular dynamics and inverse protein folding. Mesirov received her Ph.D. in mathematics from Brandeis University in 1974. Mesirov came to the Whitehead in 1997 from IBM where she was Manager of Computational Biology and Bioinformatics in the Healthcare/Pharmaceutical Solutions Organization. Mesirov is a member of the Biology and Environmental Research Advisory Committee of the Department of Energy, and a trustee of the Mathematical Sciences Research Institute in Berkeley, California. She has also served as President of the Association for Women in Mathematics, Trustee of the Society for Industrial and Applied Mathematics, and Chair of the Conference Board of the Mathematical Sciences. She is a Fellow of the American Association for the Advancement of Science.

TUESDAY | 1:30PM – 3PM

TUESDAY
NOVEMBER 7
1:30PM
ROOM D 263/265

COMPUTATIONAL BIOCHEMISTRY AND DRUG DISCOVERY
Chair: Eugene Fluder, Merck

Peta-op Computing for Large Scale Biomolecular Simulation
Robert S. Germain, IBM Computational Biology Center

Gaining an understanding of the mechanisms underlying the protein folding process is generally acknowledged as a Grand Challenge because of its scale and importance to biology. Advancing our understanding of these mechanisms is an initial focus area for the scientific portion of the Blue Gene project whose systems component is a blueprint for a 1 petaflop/sec capacity, massively parallel computer architecture that is expected to significantly enhance the state-of-the-art in the realistic modeling and simulation of this class of problems.

This talk will describe some of the challenges in the modeling and simulation of biological processes on the microscopic level, with special emphasis on the protein folding process. A selection of the scientific and algorithmic issues associated with a large-scale simulation effort aimed improving our understanding of the mechanisms behind the protein folding process will be discussed.

BIOGRAPHICAL SKETCH Robert S. Germain manages the Biomolecular Dynamics and Scalable Modeling Group, which is part of the Computational Biology Center at the IBM Thomas J. Watson Research Center. He received his A.B. in physics from Princeton University and his M.S. and Ph.D. in physics from Cornell University in 1986 and 1989, respectively. After receiving his doctorate, Germain joined the Watson Research Center as a Research Staff Member where he has subsequently worked on both scientific and technical problems including the development of novel algorithms to implement a large scale fingerprint identification system. His research interests include the parallel implementation of scientific algorithms and the applications of these algorithms. Germain is a member of the IEEE and the American Physical Society.

TUESDAY
NOVEMBER 7
2:15PM
ROOM D 263/265

Basic Computational Research for Drug Discovery
Simon Kearsley, Merck Research Laboratories

This talk will examine incidents at Merck where HPC has made a difference to the basic research drug discovery process. In addition, the presentation will focus on the newer challenges facing drug discovery and examine how computing resources

must be marshaled to address them in the near future.

BIOGRAPHICAL SKETCH Simon Kearsley is currently the Senior Director responsible for the Molecular Modeling Department at Merck Research Laboratories. His modeling group has a long history at Merck, and has become the nexus for the interaction of several scientific disciplines focused on drug discovery. Before going to Merck he spent several years at Yale University modeling radical reactions within organic crystals. He received his bachelor's, master's and doctor's degrees from Gonville and Caius College, Cambridge.

TUESDAY | 3:30PM – 5PM

COMPUTING PLATFORMS
Chair: Pat Teller, UTEP

TUESDAY
NOVEMBER 7
3:30PM
ROOM D 263/265

Blue Gene
Monty Denneau, IBM

On its target applications, Blue Gene will make the 10-teraflop ASCI White machine look like a Palm Pilot. We tell you how.

BIOGRAPHICAL SKETCH Monty Denneau is the system architect for the Blue Gene project.

Status of the Earth Simulator Project in Japan
Keiji Tani, Japan Atomic Energy Research Institute

TUESDAY
NOVEMBER 7
4:15PM
ROOM D 263/265

The Science and Technology Agency of Japan has proposed a project to promote studies for global change prediction by an integrated three-in-one research and development approach: earth observation, basic research, and computer simulation. As part of the project, an ultra-fast computer, the "Earth Simulator", with a sustained speed of more than 5 teraflops for an atmospheric circulation code, is being developed. The "Earth Simulator" is a MIMD-type, distributed-memory, parallel system in which 640 processor nodes are connected via a fast single-stage crossbar network. Each node consists of eight vector-type arithmetic processors that are tightly connected by a 16 gigabyte shared main memory. The peak performance and the main memory of the total system are 40 teraflops and 10 terabytes, respectively. As part of the development of the basic software system, an operating system service routine called "center routine," is being developed. Since an archival system will be used as main storage for user files, the most important function of the center routine is the

optimal scheduling of not only submitted batch jobs but also the optimal management of user files necessary for them. The design and R&D for both hardware and basic software systems were completed during the last three fiscal years, FY97, 98 and 99. The manufacture of the hardware system and the development of the center routine are underway. Facilities necessary for the Earth Simulator, including buildings, are under construction. The total system will be completed in spring 2002.

BIOGRAPHICAL SKETCH Dr. Keiji Tani graduated from Osaka University, Osaka, Japan, in July 1984 with a Doctor. of Engineering in Nuclear Fusion Research. Tani has worked at the Japan Atomic Energy Research Institute since April 1974, where he has been involved in research in fast-ion confinement in tokamaks using particle-simulation techniques, establishment of a laboratory of advanced photon research, and research and development of the Earth Simulator. With respect to the Earth Simulator project, he is responsible for the development of the basic software system and collaborations necessary for the development of the Earth Simulator.

WEDNESDAY | 10:30AM – 12NN

WEDNESDAY
NOVEMBER 8
10:30AM
ROOM D 263/265

IEEE AWARD WINNERS

Chair: Sally Haerer, Oregon State University

Seymour Cray Computer Engineering Award Winner
Sidney Fernbach Award Winner

WEDNESDAY | 1:30PM – 3PM

WEDNESDAY
NOVEMBER 8
1:30PM
ROOM D 263/265

REAL WORLD SCALABLE COMPUTING I

Chair: Tim Mattson, Intel

Software Applications Accelerate Performance with the VI Architecture; Hot Clusters, & What's Next?
David Fair, Ed Gronke, Giganet

This talk will focus on implementation of clustered solutions, success stories, and clustering challenges for 2001. Over the last year applications have been tuned to exploit the performance benefits of the VI Architecture for database, storage, and backup. In the database arena, for example, Giganet is the backbone for database acceleration; Microsoft's SQL 2000 edition has exhibited 30 percent or more in performance increases utilizing the VI Architecture with off-the-shelf

clustering components. Oracle will be announcing VI-based results at the introduction of their 8.2 release. The roadmap for many of the next suite of VI-enabled applications to be introduced in 2001 will be presented.

BIOGRAPHICAL SKETCH David Fair, VP Solutions, Giganet Inc. Mr. Fair joined Giganet in early 2000 after 15 years with Intel Corporation. Giganet is a leader in VI-based data center networks and Mr. Fair is responsible for working with industry to deliver VI-enabled solutions stacks. Mr. Fair's implementation role at Giganet is a natural transition from his Intel responsibilities, where he was the Director for the VI Architecture Specification in Intel's Server Architecture Lab. Mr. Fair worked in the Intel Supercomputer Systems Division from 1986 through 1993 as the Director of the Research Projects Office.

BIOGRAPHICAL SKETCH Ed Gronke, Principal Solutions Architect, Giganet, Inc. Mr. Gronke joined Giganet in early 2000 after five years with Intel Corporation. He was responsible for delivering VI-enabled applications also has experience in Intel's Supercomputer Systems Division.

From RAIN to Rainfinity
Jehoshua Bruck, Rainfinity

WEDNESDAY
NOVEMBER 8
2:15PM
ROOM D 263/265

The RAIN project was a research collaboration between Caltech and NASA-JPL on distributed computing, communication and data storage systems for future spaceborne missions. The goal of the project was to identify and develop key building blocks for reliable distributed systems built with inexpensive off-the-shelf components. The RAIN software components run in conjunction with operating system services and standard network protocols. Through software-implemented fault tolerance, the system tolerates multiple node, link, and switch failures, with no single point of failure. The RAIN technology has been transferred to Rainfinity, a start-up company focusing on creating clustered solutions for improving the performance, availability and scalability of Internet data centers.

BIOGRAPHICAL SKETCH Jehoshua (Shuki) Bruck is a Professor of Computation and Neural Systems and Electrical Engineering at the California Institute of Technology. His research interests include parallel and distributed computing, fault-tolerant computing, computation theory and neural and biological systems. Dr. Bruck has an extensive industrial experience, including working with IBM for ten years both at the IBM Almaden Research Center and the

IBM Haifa Science center. Dr. Bruck is a co-founder and Chairman of Rainfinity, a spin-off company from Caltech that is focusing on providing software for high performance reliable Internet Servers. Dr Bruck received the Ph.D. degree in Electrical Engineering from Stanford University in 1989. Dr. Bruck is the recipient of a 1997 IBM Partnership Award, a 1995 Sloan Research Fellowship, a 1994 National Science Foundation Young Investigator Award, six IBM Plateau Invention Achievement Awards, a 1992 IBM Outstanding Innovation Award, and a 1994 IBM Outstanding Technical Achievement Award for his contributions to the design and implementation of the SP-1, the first IBM scalable parallel computer. He has published more than a 150 journal and conference papers in his areas of interests and he holds 22 US patents.

WEDNESDAY | 3:30PM – 5PM

WEDNESDAY
NOVEMBER 8
3:30PM
ROOM D 263/265

REAL WORLD SCALABLE COMPUTING II
Chair: Siamek Zadeh, Sun Microsystems

Computing Challenges in the Travel & Transportation Industry
Richard Ratliff, Sabre Holdings, Inc.

This presentation will focus on the use of high-end Sun computers at Sabre, Inc. Sabre's mid-range environment is comprised predominantly of Sun servers, and this equipment is used to run hundreds of different applications for our airline customers worldwide. These applications span a wide range of uses—from large database systems to scientific computing with forecasting and optimization models. Overviews of performance and sizing requirements for selected applications will be discussed. Support, procurement and operational issues will also be addressed for our centrally hosted outsourcing and ASP applications as well as those customers with local hardware implementations.

BIOGRAPHICAL SKETCH Richard Ratliff is the Vice President of Technology and Product Integration at Sabre and is responsible for coordinating the Operations Research model and data exchange among Sabre's airline applications. His primary expertise is in the area of forecasting and optimization modeling. He served for four years as Director—Revenue Management for Ansett Australia and also led Sabre's development and implementation of its advanced yield management systems. Mr. Ratliff has an M.S. in Economics from the Colorado School of Mines.

SETI@home: Internet Distributed Computing for SETI
David Anderson, United Devices, Inc.

WEDNESDAY
NOVEMBER 8
4:15PM
ROOM D 263/265

SETI@home is a radio SETI (Search for Extraterrestrial Intelligence) project that takes a novel approach to the compute-intensive analysis of radio signals. Instead of placing a dedicated supercomputer at the telescope, SETI@home distributes data over the Internet to computers in the homes and offices of volunteers. In SETI@home's first year of operation, over 2 million people have participated, and have contributed 300,000 years of computer time. This approach holds promise for a number of other scientific computing problems.

BIOGRAPHICAL SKETCH Dr. David Anderson is the Director of the SETI@home Project and former visiting scientist at U.C. Berkeley. As the project director, Dr. Anderson was the technical lead—designing and managing the implementation of the client-side and server-side software, database, and the server hardware architecture. Prior to SETI@home, Dr. Anderson was the CTO of Tunes.com where he architected and implemented a database-driven, Web-based system for personalized music discovery and marketing. Prior to these positions Dr. Anderson was the Director of Software Architecture at Sonic Solutions as well as an Assistant Professor, CS Division, EECS Department, at UC Berkeley. Dr. Anderson has authored or co-authored 65 papers in Computer Science and he is the sole inventor on two pending patent applications for technology related to MediaNet and an invention involving 3D interactive television.

THURSDAY | 10:30AM – 12NN

SUPERCOMPUTING TRENDS IN MCAE I
Chair: Ed Turkel, Compaq Corporation

THURSDAY
NOVEMBER 9
10:30AM
ROOM D 263/265

Satisfying CFD Engineering Constraints (with Parallel Processing)
Stephen A. Remondi, James Hoch, Exa Corporation

The field of Computational Fluid Dynamics (CFD) within the engineering development process is changing rapidly. CFD is fast becoming an integral part of the engineering design cycle. In order to meet these real world demands CFD must meet a set of specific requirements. First, CFD must be capable of handling real world complex geometries. Second, CFD must be accurate. The degree of accuracy is defined as being sufficient to enable engineering decisions. Typically, this means in the 3 to 10 percent error range. Third, CFD must be integrated into the engineering process. This means cou-

pling to mechanical CAD and design software as well as to other analysis software packages. Last is performance. CFD must provide these requirements in a timely manner. Engineering answers are not useful if the design has already been signed off. Parallel processing on scalable SMP and clustered systems form the hardware platforms that provide fast turn around times for CFD simulations.

This presentation will review in some detail, the real engineering requirements for CFD. The presentation will then highlight how parallel computing is utilized to meet these requirements. The presentation will conclude with a review of some recent examples of how CFD is being utilized in production engineering.

BIOGRAPHICAL SKETCH James E. Hoch, vice president of software development, was previously systems architect at Maker Communications, a developer of communication processors and associated software for use in the networking industry. Prior to that he was lead architect on several research parallel computer systems and related compiler projects at Sandia National Laboratories. Hoch originally joined Exa in 1993 with extensive experience in compiler technology and parallel computing. Hoch holds both M.S. and B.S. degrees in Computer and Electrical Engineering from Purdue University.

BIOGRAPHICAL SKETCH Stephen A. Remondi is president and CEO and a co-founder of Exa Corporation. Remondi has invested over seven years at Exa leading product development. Before assuming the role of CEO in '99, Remondi began his career at Exa as a director, was promoted to VP of engineering in '95 and then VP of application development in '97. Prior to starting Exa, Remondi worked for Alliant Computer Systems and Data General in the development of parallel computer systems, systems architectures, and ASIC design. Remondi holds a B.S. in Computing & Electrical Engineering from Tufts University and an M.B.A. from Bentley College.

THURSDAY
NOVEMBER 9
11:15AM
ROOM D 263/265

Stochastic Simulation: Breaking the stagnation and fragmentation of contemporary HPC

Jacek Marczyk, EASi Engineering GmbH

Contemporary CAE and HPC are in state of crises. Evidence of this is the excessive fragmentation and fractalization of tools, procedures and the market. Fragmentation is a reflection of a state of crisis. States of crisis appear when a discipline runs out of ideas and the only way to push progress and innovation is via small increments. HPC and CAE are

amidst a state of profound crisis that has been brought about by determinism, reductionism, and obsession with accuracy and optimality. The fact that CAE and HPC are consistently neglecting uncertainty (material properties, load, boundary conditions, geometry) ultimately gives rise to orthodox forms of computing based on brute force increases in finite element number, particle counts and finer time-steps. At the same time, the fundamental issues of confidence and model validity are being embarrassingly omitted, or even silenced, due to their political incorrectness. Stochastic Simulation based on Monte Carlo techniques that have been recently introduced into the industry has the potential of bringing about a reconciliation of CAE/HPC with experimentation, and the capability to deliver robust designs which transcend the simplistic and trendy Multi-Disciplinary Optimization.

BIOGRAPHICAL SKETCH Dr. Jacek Marczyk is the Vice President, Advanced Technologies at EASi Engineering GmbH, which does computer-aided engineering for the European automotive industry. Dr. Marczyk has over 18 years experience in Structural Mechanics, Simulation and Control System Design for the Aerospace, Offshore and Automotive industries. He has a PhD in Civil Engineering from Polytechnic University of Catalonia, Barcelona, Spain and has 28 publications, including four books.

PLENARY SESSION
THURSDAY | 1:30PM – 3PM

AWARDS SESSION
Chair: Sally Haerer, Oregon State University

THURSDAY
NOVEMBER 9
1:30PM
BALLROOM C

On the Scale and Performance of Cooperative Web Proxy Caching
Goeff Voekler, UC San Diego, Computing Research Association Digital Fellow

While algorithms for cooperative proxy caching have been widely studied, little is understood about cooperative-caching performance in the large-scale World Wide Web environment. In this talk, I will describe the work that we have done to explore the potential advantages and drawbacks of inter-proxy cooperation across a wide range of scales in client population. Our work used both trace-based analysis and analytic modeling. With our traces, we evaluated quantitatively the performance-improvement potential of cooperation between 200 small-organization proxies within a university environment, and between two large-organization proxies handling 23,000 and 60,000 clients, respectively. With our model, we extended beyond

these populations to project cooperative caching behavior in regions with millions of clients. Overall, our results indicate that cooperative caching has performance benefits when scaling client populations only within limited population bounds. We also used our model to examine the implications of future trends in Web-access behavior and traffic.

BIOGRAPHICAL SKETCH Geoffrey M. Voelker is an assistant professor at the University of California at San Diego. His research interests include operating systems, distributed systems, and Internet systems. He received a BS degree in Electrical Engineering and Computer Science from the University of California at Berkeley in 1992, and the MS and PhD degrees in Computer Science and Engineering from the University of Washington in 1995 and 2000, respectively.

THURSDAY | 3:30PM – 5PM

THURSDAY
NOVEMBER 9
3:30PM
ROOM D 263/265

SUPERCOMPUTING TRENDS IN MCAE II

Chair: Ed Turkel, Compaq Corporation

Moving a Large Commercial Application from SMP to DMP
David Lombard, MSC Software Corporation

Moving a high-performance FEA application from symmetric multi-processing (SMP) to Distributed Memory Processing (DMP) offers many advantages and challenges. The largest challenge is the re-architecture of the application to suit the demands of message passing—that challenge is recognized and usually understood. Along the way, however, are differences in MPI implementations, issues outside the MPI standard, such as program startup, and platform issues such as communications performance, reliability, and scalability which must all be addressed in order to successfully deploy the application.

THURSDAY
NOVEMBER 9
4:15PM
ROOM D 263/265

LS-DYNA—Application-Driven Strategies for High Performance Computing
Mark Christon, Livermore Software Technology Corp.

LS-DYNA is a multi-physics finite element code that has long been associated with large-scale high performance computing. An application-centric focus coupled with an emphasis on high performance computing has permitted LS-DYNA's continuing feature development to include a comprehensive set of capabilities for problems ranging from automotive crash-worthiness and occupant safety, metal forming and fluid-structure interaction, to heat

transfer, compressible and incompressible flows. Although current computing trends are yielding benefits in some applications areas, they are also challenging the sustained growth of some grid-based solution strategies. To accommodate these trends, LS-DYNA relies on a configurable hierarchical approach to data distribution and ultimately parallelism. At the coarse-grained level, a domain-decomposition message-passing paradigm is used to partition and distribute the data and concomitant computational load across processors. At an intermediate level of granularity, data is distributed across shared-memory processors using directive-based parallelism. Configurable vector/cache blocking is used at the finest level of granularity to achieve an effective processor-level algorithm-to-architecture mapping. Two applications are used to contrast the advantages and potential pitfalls of current trends in scientific computing. The first demonstrates that low-cost clusters and Beowulf systems are making large-eddy simulation accessible to a broad class of users. The second application focuses on optimization with LS-OPT. LS-OPT uses a response surface methodology to construct design rules and conduct optimization requiring a large number of independent parallel LS-DYNA simulations. An example of an automotive crash-worthiness optimization problem demonstrates LS-OPT's embarrassingly parallel use of LS-DYNA simulations on an HP V-Class machine and a network of PC's.

BIOGRAPHICAL SKETCH Mark Christon has been working in the area of acoustic fluid-structure interaction and time-dependent incompressible/low-Mach fluid dynamics with an emphasis on second-order projection methods and large-eddy simulation for the past 10 years. During this period of time he worked at Lawrence Livermore in the Methods Development Group and Sandia National Laboratories in the Computational Physics department. Mark received his PhD in mechanical engineering from Colorado State University in 1990. Currently, Mark is a senior scientist at Livermore Software Technology Corporation and is the primary developer of the incompressible/low-Mach flow solver, fluid-structure capability in LS-DYNA.

SC2000 WEBCASTS

SC2000 will broadcast live over the Internet select presentations from the Technical Program, the Exhibition, Venture Village and eSCape 2000 using the unique capabilities of SCinet 2000 to deliver high-quality audio and video broadcasts.

SC2000 will again be webcasting the keynote address and the plenary sessions. This year, the keynote and plenary sessions will be archived and viewable for the duration of the SC2000 conference. SC2000 will also be adding select industry and research demonstrations from the exhibition hall to the webcast program. In addition, the webcast team will again support a robotic camera to allow viewers to pan the exhibit hall and zoom in on exhibit floor booths.

The SC2000 webcast events will be broadcast in several formats to accommodate users who may not have access to a high performance network to their desktop. Instructions for accessing these webcasts are available from the SC2000 pages at: www.sc2000.org/webcasts/. These pages also feature information about how to configure your computer to view the webcasts, where to download free player software, and sample video clips to test your configuration.

WEBCAST HELP DESK In an effort to help solve any problems our remote viewers may encounter we are providing a webcast help desk. User support will be available by telephone and email shortly before, during, and shortly after all webcast events. Please visit the webcast pages for contact information.

All webcast viewers are encouraged to provide feedback on the quality and scope of the broadcasts, via a form at the above URL. Tell us what you found useful, what you would like to see in the future, and how effective you found the webcasts.

WEBCAST INFRASTRUCTURE The SC2000 webcast infrastructure will include several key components. A PC-class machine located in the presentation area will be used to convert audio and video feeds of the keynote and plenary talks to a digital format. A wireless PC-class laptop will capture events in the Exhibition, Venture Village, and eSCape 2000 areas. Finally, a server-class system located in the SC2000 webcast booth near the SCinet 2000 network operations center will make this content available to remote viewers. SCinet 2000 is providing all supporting network infrastructure. Webcasting components include a high-bandwidth connection out of the convention center and a 100BaseT dedicated connection between the PC encoder in the presentation area and the webcast server.

SC2000 TECHNICAL PAPERS/AWARDS

Conference Vice-Chair, Technical Program: James R. McGraw, Lawrence Livermore National Laboratory

This year's technical program continues the tradition of a diverse set of strong technical papers. The Program Committee selected 62 papers from 179 submissions, including six finalists for Gordon Bell prizes. These papers were chosen based on the significance of their contributions to the field of high performance networking and computing. The papers will be presented in 21 sessions, covering a wide range of topics, including cluster computing, networking, computational girds, data grids, QoS/fault tolerance, and a myriad of applications. Also, don't miss the plenary awards session on Thursday, where you can find out the winners of all of the conference awards (including Best Paper and Best Student Paper) and several national awards.
Notes: First author is presenter, unless another author in the list has an asterisk after his/her name.

TUESDAY | 10:30AM – 12NN

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 271/273

MPI

Chair: Harvey Wasserman, Los Alamos National Laboratory

The Implementation of MPI-2 One-Sided Communication for the NEC SX-5
Jesper Larsson Traff, Hubert Ritzdorf, Rolf Hempel, C&C Research Laboratories, NEC Europe Ltd

We describe the MPI/SX implementation of the MPI-2 standard for one-sided communication (Remote Memory Access) for the NEC SX-5 vector supercomputer. MPI/SX is a non-threaded implementation of the full MPI-2 standard. Essential features of the implementation are presented, including the synchronization mechanisms, the handling of communication windows in global shared and in process local memory, as well as the handling of MPI derived data types. In comparative benchmarks the data transfer operations for one-sided communication and point-to-point message passing show very similar performance, both when data reside in global shared and when in process local memory. Derived data types, which are of particular importance for applications using one-sided communications, impose only a modest overhead and can be used without any significant loss of performance. Thus, the MPI/SX programmer can freely choose

either the message passing or the one-sided communication model, whichever is most convenient for the given application.

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 271/273

Single-sided MPI Implementations for SUN MPI
S. Booth, E. Mourao, EPCC, The University of Edinburgh

This paper describes an implementation of generic MPI-2 single-sided communications for SUN-MPI. Our implementation is layered on top of point-to-point MPI communications and therefore can be adapted to other MPI implementations.

The code is designed to co-exist with other MPI-2 single-sided implementations (for example, direct use of shared memory) providing a generic fall-back implementation for those communication paths where an optimized single-sided implementation is not available. MPI-2 single-sided communications require the transfer of data-type information as well as user data. We describe a type packing and caching mechanism used to optimize the transfer of data-type information.

The performance of this implementation is measured in comparison to equivalent point-to-point operations and the shared memory implementation provided by SUN.

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 271/273

Automatically Tuned Collective Communications
Sathish S. Vadhiyar, Graham E. Fagg, Jack Dongarra,
University of Tennessee, Knoxville

The performance of the MPI's collective communications is critical in most MPI-based applications. A general algorithm for a given collective communication operation may not give good performance on all systems due to the differences in architectures, network parameters and the storage capacity of the underlying MPI implementation. In this paper we discuss an approach in which the collective communications are tuned for a given system by conducting a series of experiments on the system. We also discuss a dynamic topology method that uses the tuned static topology shape, but re-orders the logical addresses to compensate for changing run time variations. A series of experiments were conducted comparing our tuned collective communication operations to various native vendor MPI implementations. The use of the tuned collective communications resulted in about 30 percent to 650 percent improvement in performance over the native MPI implementations.

NUMERICAL ALGORITHMS

Chair: Ricky Kendall, Ames Laboratory

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 274

Landing CG on EARTH: A Case Study of Fine-Grained Multithreading on an Evolutionary Path

Kevin B. Theobald, Gagan Agrawal, Rishi Kumar, Gerd Heber, Cornell Univ., Guang R. Gao, Univ. of Delaware, Paul Stodghill, Keshav Pingali, Cornell Univ.

We report on our work in developing a fine-grained multithreaded solution for the communication-intensive Conjugate Gradient (CG) problem. In our recent work, we have developed a simple, yet very efficient, solution to executing matrix-vector multiply on a multithreaded system. This paper presents an effective mechanism for the reduction-broadcast phase, which is implemented and integrated with the sparse MVM resulting in a scalable implementation of the complete CG application.

Three major observations from our experiments on the EARTH multithreaded testbed are: (1) The scalability of our CG implementation is impressive, e.g., speedup is 90 on 120 processors for the NAS CG class B input. (2) Our dataflow-style reduction-broadcast network based on fine-grain multithreading is twice as fast as a serial reduction scheme on the same system. (3) By slowing down the network by a factor of 2, no notable degradation of overall CG performance was observed.

Parallel Smoothed Aggregation Multigrid: Aggregation Strategies on Massively Parallel Machines

Ray S. Tuminaro, Sandia National Laboratories, Charles Tong, Lawrence Livermore National Laboratory

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 274

Algebraic multigrid methods offer the hope that multigrid convergence can be achieved (for at least some important applications) without a great deal of effort from engineers and scientists wishing to solve linear systems. In this paper we consider parallelization of the smoothed aggregation multigrid method. Smoothed aggregation is one of the most promising algebraic multigrid methods. Therefore, developing parallel variants with both good convergence and efficiency properties is of great importance. However, parallelization is nontrivial due to the somewhat sequential aggregation (or grid coarsening) phase. In this paper, we discuss three different parallel aggregation algorithms and illustrate the advantages and disadvantages of each variant in terms of parallelism and convergence. Numerical results will be shown on the Intel Teraflop computer for some large problems coming from nontrivial codes: quasi-static electric potential simulation and a fluid flow calculation.

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 274

Scalable Algorithms for Adaptive Statistical Designs
Robert Oehmke, Janis Hardwick, Quentin F. Stout, University of Michigan

We present a scalable, high performance solution to multidimensional recurrences that arise in adaptive statistical designs. Adaptive designs are an important class of learning algorithms for a stochastic environment, and we focus on the problem of optimally assigning patients to treatments in clinical trials. While adaptive designs have significant ethical and cost advantages, they are rarely utilized because of the complexity of optimizing and analyzing them. Computational challenges include massive memory requirements, few calculations per memory access, and multiply-nested loops with dynamic indices. We analyze the effects of various parallelization options, and while standard approaches do not work well, with effort an efficient, highly scalable program can be developed. This allows us to solve problems thousands of times more complex than those solved previously, which helps make adaptive designs practical. Further, our work applies to many other problems involving neighbor recurrences, such as generalized string matching.

TUESDAY | 1:30PM – 3PM

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 267

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

Randomization, Speculation, and Adaptation in Batch Schedulers
Dejan Perkovic, Peter J. Keleher, University of Maryland

This paper proposes extensions to the backfilling job-scheduling algorithm that significantly improve its performance. We introduce variations that sort the "backfilling order" in priority-based and randomized fashions. We examine the effectiveness of guarantees present in conservative backfilling and find that initial guarantees have limited practical value, while the performance of a "no-guarantee" algorithm can be significantly better when combined with extensions that we introduce. Our study differs from many similar studies in using traces that contain user estimates. We find that actual overestimates are large and significantly different from simple models. We propose the use of speculative backfilling and speculative test runs to counteract these large over-estimations. Finally, we explore the impact of dynamic, system-directed adaptation of application parallelism. The cumulative improvements of these techniques decrease the bounded slowdown, our primary metric, to less then 15 percent of conservative backfilling.

An Object-Oriented Job Execution Environment
Lance Smith, Rod Fatoohi*, San Jose State University

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 267

This is a project for developing a distributed job execution environment for highly iterative jobs. An iterative job is one where the same binary code is run hundreds of times with incremental changes in the input values for each run. An execution environment is a set of resources on a computing platform that can be made available to run the job and hold the output until it is collected. The goal is to design a complete, object-oriented scheduling system that will run a variety of jobs with minimal changes. Areas of code that are unique to one specific type of job are decoupled from the rest. The system allows for fine-grained job control, timely status notification and dynamic registration and deregistration of execution platforms depending on resources available. Several objected-oriented technologies are employed: Java, CORBA, UML, and software design patterns. The environment has been tested using a CFD code, INS2D.

Towards an Integrated, Web-executable Parallel Programming Tool Environment
Insung Park, Nirav H. Kapadia, Renato J. Figueiredo, Rudolf Eigenmann*, José A. B. Fortes, Purdue University

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 267

We present a new parallel programming tool environment that is (1) accessible and executable "anytime, anywhere," through standard Web browsers and (2) integrated in that it provides tools which adhere to a common underlying methodology for parallel programming and performance tuning. The environment is based on a new network computing infrastructure developed at Purdue University.

We evaluate our environment qualitatively by comparing our tool access method with conventional schemes of software download and installation. We also quantitatively evaluate the efficiency of interactive tool access in our environment. We do this by measuring the response times of various functions of the Ursa Minor tool and compare them with those of a Java Applet-based "anytime, anywhere" tool access method. We found that our environment offers significant advantages in terms of tool accessibility, integration, and efficiency.

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 271/273

MPI/OPENMP

Chair: Jeff Hollingsworth, University of Maryland

Performance of Hybrid Message-Passing and Shared-Memory Parallelism for Discrete Element Modeling

D.S. Henty, Edinburgh Parallel Computing Centre

The current trend in HPC hardware is towards clusters of shared-memory (SMP) compute nodes. For applications developers the major question is how best to program these SMP clusters. To address this we study an algorithm from Discrete Element Modeling, parallelized using both the message-passing and shared-memory models simultaneously ("hybrid" parallelization). The natural load-balancing methods are different in the two parallel models, the shared-memory method being in principle more efficient for very load-imbalanced problems. It is therefore possible that hybrid parallelism will be beneficial on SMP clusters. We benchmark MPI and OpenMP implementations of the algorithm on MPP, SMP and cluster architectures, and evaluate the effectiveness of hybrid parallelism. Although we observe cases where OpenMP is more efficient than MPI on a single SMP node, we conclude that our current OpenMP implementation is not yet efficient enough for hybrid parallelism to outperform pure message-passing on an SMP cluster.

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 271/273

A Comparison of Three programming Models for Adaptive Applications on the Origin2000

Hongzhang Shan, Jaswinder P. Singh, Princeton University, Leonid Oliker, Lawrence Berkeley National Laboratory, Rupak Biswas, NASA Ames Research Center

Adaptive applications have computational workloads and communication patterns which change unpredictably at runtime, requiring load balancing to achieve scalable performance on parallel machines. Efficient parallel implementation of such adaptive application is therefore a challenging task. In this paper, we compare the performance of and the programming effort required for two major classes of adaptive applications under three leading parallel programming models on an SGI Origin 2000 system, a machine which supports all three models efficiently. Results indicate that the three models deliver comparable performance. However, the implementations differ significantly beyond merely using explicit messages versus implicit loads/stores even though the basic parallel algorithms are similar. Compared with the message-passing (using MPI) and SHMEM programming models, the cache-

coherent shared address space (CC-SAS) model provides substantial ease of programming at both the conceptual level and program orchestration levels, often accompanied by performance gains. However, CC-SAS currently has portability limitations and may suffer from poor spatial locality of physically distributed shared data on large numbers of processors.

MPI versus MPI+OpenMP on the IBM SP for the NAS Benchmarks
Franck Cappello, Daniel Etiemble, LRI

The hybrid memory model of clusters of multiprocessors raises two issues: programming model and performance. Many parallel programs have been written by using the MPI standard. To evaluate the pertinence of hybrid models for existing MPI codes, we compare a unified model (MPI) and a hybrid one (OpenMP fine grain parallelization after profiling) for the NAS 2.3 benchmarks on two IBM SP systems. The superiority of one model depends on 1) the level of shared memory model parallelization, 2) the communication patterns and 3) the memory access patterns. The relative speeds of the main architecture components (CPU, memory, and network) are of tremendous importance for selecting one model. With the used hybrid model, our results show that a unified MPI approach is better for most of the benchmarks. The hybrid approach becomes better only when fast processors make the communication performance significant and the level of parallelization is sufficient.

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 271/273

POTPOURRI

Chair: Bob Lucas, NERSC

A Wrapper Generator for Wrapping High Performance Legacy Codes as Java/CORBA Components
M. Li, O. F. Rana, M. S. Shields, Cardiff University, UK, D. W. Walker, Oak Ridge National Laboratory

This paper describes a Wrapper Generator for wrapping high performance legacy codes as Java/CORBA components for use in a distributed component-based problem-solving environment. Using the Wrapper Generator we have automatically wrapped an MPI-based legacy code as a single CORBA object, and implemented a problem-solving environment for molecular dynamic simulations. Performance comparisons between runs of the CORBA object and the original legacy code on a cluster of workstations and on a parallel computer are also presented.

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 274

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 274

A Scalable SNMP-based Distributed Monitoring System for Heterogeneous Network Computing
Rajesh Subramanyan, José Miguel-Alonso, José A. B. Fortes, Purdue University

Traditional centralized monitoring systems do not scale to present-day large, complex, network-computing systems. Based on recent SNMP standards for distributed management, this paper addresses the scalability problem through distribution of monitoring tasks, applicable for tools such as SIMONE (SNMP-based monitoring prototype implemented by the authors).

Distribution is achieved by introducing one or more levels of a dual entity called the Intermediate Level Manager (ILM) between a manager and the agents. The ILM accepts monitoring tasks described in the form of scripts and delegated by the next higher entity. The solution is flexible and integratable into a SNMP tool without altering other system components.

A testbed of up to 1024 monitoring elements is used to assess scalability. Noticeable improvements in the round trip delay (from seconds to less than one tenth of a second) were observed when more than 200 monitoring elements are present and as few as two ILM's are used.

TUESDAY
NOVEMBER 7
1:30PM – 3PM
ROOM D 274

ESP: A System Utilization Benchmark
Adrian T. Wong, Leonid Oliker, William T.C. Kramer, Teresa L. Kaltz, David H. Bailey*, Lawrence Berkeley National Laboratory

This article describes a new benchmark, called the Effective System Performance (ESP) test, which is designed to measure system-level performance, including such factors as job scheduling efficiency, handling of large jobs and shutdown-reboot times. In particular, this test can be used to study the effects of various scheduling policies and parameters. We present here some results that we have obtained so far on the Cray T3E and IBM SP systems, together with insights obtained from simulations.

TUESDAY | 3:30PM – 5PM

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 271/273

CLUSTER INFRASTRUCTURE
Chair: Sam Uselton, Lawrence Livermore National Laboratory

PM2: High Performance Communication Middleware for Heterogeneous Network Environments
Toshiyuki Takahashi, Shinji Sumimoto, Atsushi Hori, Hiroshi Harada, Yutaka Ishikawa, Real World Computing Partnership

This paper introduces a high performance communication middle layer, called PM2, for heterogeneous network envi-

ronments. PM2 currently supports Myrinet, Ethernet, and SMP. Binary code written in PM2 or written in a communication library, such as MPICH-SCore on top of PM2, may run on any combination of those networks without re-compilation. According to a set of NAS parallel benchmark results, MPICH-SCore performance is better than dedicated communication libraries such as MPICH-BIP/SMP and MPICH-GM when running some benchmark programs.

Performance and Interoperability Issues in Incorporating Cluster Management Systems Within a Wide-Area Network-Computing Environment

Sumalatha Adabala, Nirav H. Kapadia, José A. B. Fortes, Purdue University

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 271/273

This paper describes the performance and interoperability issues that arise in the process of integrating cluster management systems into a wide-area network-computing environment, and provides solutions in the context of the Purdue University Network Computing Hubs (PUNCH). The described solution provides users with a single point of access to resources spread across administrative domains, and an intelligent translation process makes it possible for users to submit jobs to different types of cluster management systems in a transparent manner. The approach does not require any modifications to the cluster management software. However, call-back and caching capabilities that would improve performance and make such systems more interoperable with wide-area computing systems are discussed.

Architectural and Performance Evaluation of GigaNet and Myrinet Interconnects on Clusters of Small-Scale SMP Servers

Jenwei Hsieh, Tau Leng, Victor Mashayekhi, Reza Rooholamini, Dell Computer Corporation

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 271/273

GigaNet and Myrinet are two of the leading interconnects for clusters of commodity computer systems. Both provide memory-protected user-level network interface access, and deliver low-latency and high-bandwidth communication to applications. GigaNet is a connection-oriented interconnect based on a hardware implementation of Virtual Interface (VI) Architecture and Asynchronous Transfer Mode (ATM) technologies. Myrinet is a connection-less interconnect which leverages packet switching technologies from experimental Massively Parallel Processors (MPP) networks. This paper investigates their architectural differences and evaluates their performance on two commodity clusters based on two generations of Symmetric Multiple Processors (SMP) servers. The perform-

ance measurements reported here suggest that the implementation of Message Passing Interface (MPI) significantly affects the cluster performance. Although MPICH-GM over Myrinet demonstrates lower latency with small messages, the polling-driven implementation of MPICH-GM often leads to tight synchronization between communication processes and higher CPU overhead.

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 274

QoS/FAULT TOLERANCE

Chair: Wu-chun Feng, Los Alamos National Laboratory

MPICH-GQ: Quality-of-Service for Message Passing Programs

Alain J. Roy, University of Chicago, Ian Foster, Argonne National Laboratory and University of Chicago, William Gropp, Argonne National Laboratory, Nicholas Karonis, Northern Illinois University, Volker Sander, Forschungszentrum Juelich GmbH, Brian Toonen, Argonne National Laboratory

Parallel programmers typically assume that all resources required for a program's execution are dedicated to that purpose. However, in local and wide area networks, contention for shared networks, CPUs, and I/O systems can result in significant variations in availability, with consequent adverse effects on overall performance. We describe a new message-passing architecture, MPICH-GQ, that uses quality of service (QoS) mechanisms to manage contention and hence improve performance of message passing interface (MPI) applications. MPICH-GQ combines new QoS specification, traffic shaping, QoS reservation, and QoS implementation techniques to deliver QoS capabilities to the high-bandwidth bursty flows, complex structures, and reliable protocols used in high performance applications—characteristics very different from the low-bandwidth, constant bit-rate media flows and unreliable protocols for which QoS mechanisms were designed. Results obtained on a differentiated services testbed demonstrate our ability to maintain application performance in the face of heavy network contention.

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 274

Scalable Fault-Tolerant Distributed Shared Memory

Florin Sultan, Thu Nguyen, Liviu Iftode, Rutgers University

This paper shows how a state-of-the-art software distributed shared-memory (DSM) protocol can be efficiently extended to tolerate single-node failures. In particular, we extend a home-based lazy release consistency (HLRC) DSM system with independent checkpointing and logging to volatile memory, targeting shared-memory computing on very large LAN-based clusters. In these environments,

where global coordination may be expensive, independent checkpointing becomes critical to scalability. However, independent checkpointing is only practical if we can control the size of the log and checkpoints in the absence of global coordination. In this paper we describe the design of our fault-tolerant DSM system and present our solutions to the problems of checkpoint and log management. We also present experimental results showing that our fault tolerance support is light-weight, adding only low messaging, logging and checkpointing overheads, and that our management algorithms can be expected to effectively bound the size of the checkpoints and logs for real applications.

Realizing Fault Resilience in Web-Server Cluster
Chu-Sing Yang, Mon-Yen Luo, National Sun Yat-Sen University

Today, a successful Internet service is absolutely critical to be up 100 percent of the time. Server clustering is the most promising approach to meet this requirement. However, the existing Web server-clustering solutions merely can provide high availability derived from their redundancy nature, but offer no guarantee about fault resilience for the service. In this paper, we address this problem by implementing an innovative mechanism which enables a Web request to be smoothly migrated and recovered on another working node in the presence of server failure. We will show that request migration and recovery could be efficiently achieved in the manner of user transparency. The achieved capability of fault resilience is important and essential for a variety of critical services (e.g., E-commerce), which are increasingly widespread used. Our approach takes an important step toward providing a highly reliable Web service.

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 274

BIOMEDICAL APPLICATIONS

Chair: Padma Raghavan, University of Tennessee

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 267

Data Access Performance in a Large and Dynamic Pharmaceutical Drug Candidate Database
Zina Ben-Miled, Yang Liu, Indiana University, Purdue University, Dave Powers, Eli Lilly, Omran Bukhres, Indiana University, Purdue University, Michael Bem, Robert Jones, Robert Oppelt, Samuel Milosevich, Eli Lilly

An explosion in the amount of data generated through chemical and biological experimentation has been observed in recent years. This rapid proliferation of vast amounts of data has led to a set of cheminformatics and bioinformatics applications that manipulate dynamic, het-

erogeneous and massive data. An example of such applications in the pharmaceutical industry is the computational process involved in the early discovery of lead drug candidates for a given target disease. This computational process includes repeated sequential and random accesses to a drug candidate database.

Using the above pharmaceutical application, an experimental study was conducted which shows that for optimal performance, the degree of parallelism exploited in the application should be adjusted according to the drug candidate database instance size and the machine size. Additionally, different degrees of parallelism should be used depending on whether the access to the drug candidate database is random or sequential.

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 267

Real-Time Biomechanical Simulation of Volumetric Brain Deformation for Image Guided Neurosurgery

Simon K. Warfield, Harvard Medical School, Matthieu Ferrant, Xavier Gallez, Universitie Catholique de Louvain, Arya Nabavi, Ferenc A. Jolesz, Harvard Medical School, Ron Kikinis, Harvard Medical School

We aimed to study the performance of a parallel implementation of an intraoperative nonrigid registration algorithm that accurately simulates the biomechanical properties of the brain and its deformations during surgery. The algorithm was designed to allow for improved surgical navigation and quantitative monitoring of treatment progress in order to improve the surgical outcome and to reduce the time required in the operating room. We have applied the algorithm to two neurosurgery cases with promising results.

High performance computing is a key enabling technology that allows the biomechanical simulation to be executed quickly enough for the algorithm to be practical. Our parallel implementation was evaluated on a symmetric multi-processor and two clusters and exhibited similar performance characteristics on each. The implementation was sufficiently fast to be used in the operating room during a neurosurgery procedure. It allowed a three-dimensional volumetric deformation to be simulated in less than ten seconds.

TUESDAY
NOVEMBER 7
3:30PM – 5PM
ROOM D 267

Computer Simulations of Cardiac Electrophysiology

John B. Pormann, Craig S. Henriquez, John A. Board Jr., Donald J. Rose, David M. Harrild, Duke University, Alexandra P. Henriquez, North Carolina Supercomputing Center

CardioWave is a modular system for simulating wavefront conduction in the heart. These simulations may be used to

investigate the factors that generate and sustain life-threatening arrhythmias such as ventricular fibrillation. The user selects a set of modules which most closely reflects the simulation they are interested in and the simulator is built automatically. Thus, we do not present one monolithic simulator, but rather a simulator-generator which allows the researcher to make the trade-offs of complexity versus performance. The results presented here are from simulations run on an IBM SP parallel computer and a cluster of workstations. The performance numbers show excellent scalability up through 128 processors. With the larger memory of the parallel machines, we have been able to perform highly realistic simulations of the human atria. These simulations include realistic, 3-D geometries with inhomogeneity and anisotropy as well as highly complex membrane dynamics.

WEDNESDAY | 10:30AM – 12NN

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 274

APPLICATIONS I

Chair: Michael Berry, University of Tennessee

Parallel Algorithms for Radiation Transport on Unstructured Grids
Steve Plimpton, Bruce Hendrickson, Shawn Burns, Sandia National Laboratories, Will McLendon III, Texas A&M University

The method of discrete ordinates is commonly used to solve the Boltzmann radiation transport equation for applications ranging from simulations of fires to weapons effects. The equations are most efficiently solved by sweeping the radiation flux across the computational grid. For unstructured grids this poses several interesting challenges, particularly when implemented on distributed-memory parallel machines where the grid geometry is spread across processors. We describe an asynchronous, parallel, message-passing algorithm that performs sweeps simultaneously from many directions across unstructured grids. We identify key factors that limit the algorithm's parallel scalability and discuss two enhancements we have made to the basic algorithm: one to prioritize the work within a processor's subdomain and the other to better decompose the unstructured grid across processors. Performance results are given for the basic and enhanced algorithms implemented within a radiation solver running on hundreds of processors of Sandia's Intel Tflops machine and DEC-Alpha CPlant cluster.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 274

A Parallel Dynamic-Mesh Lagrangian Method for Simulation of Flows with Dynamic Interfaces

James F. Antaki, Univ. of Pittsburgh Medical Center, Guy E. Blelloch, Omar Ghattas, Ivan Malcevic *, Gary L. Miller, Noel J. Walkington, Carnegie Mellon Univ.

Many important phenomena in science and engineering, including our motivating problem of microstructural blood flow, can be modeled as flows with dynamic interfaces. The major challenge faced in simulating such flows is resolving the interfacial motion. Lagrangian methods are ideally suited for such problems, since interfaces are naturally represented and propagated. However, the material description of motion results in dynamic meshes, which become hopelessly distorted unless they are regularly regenerated. Lagrangian methods are particularly challenging on parallel computers, because scalable dynamic mesh methods remain elusive. Here, we present a parallel dynamic mesh Lagrangian method for flows with dynamic interfaces. We take an aggressive approach to dynamic meshing by triangulating the propagating grid points at every timestep using a scalable parallel Delaunay algorithm. Contrary to conventional wisdom, we show that the costs of the dynamic mesh components (triangulation, coarsening, refinement, and partitioning) can be made small relative to the flow solver.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 274

Self-Consistent Langevin Simulation of Coulomb Collisions in Charged-Particle Beams

Ji Qiang, Robert D. Ryne, Salman Habib, Los Alamos National Laboratory

In many plasma physics and charged-particle beam dynamics problems, Coulomb collisions are modeled by a Fokker-Planck equation. In order to incorporate these collisions, we present a three-dimensional parallel Langevin simulation method using a Particle-In-Cell (PIC) approach implemented on high-performance parallel computers. We perform, for the first time, a fully self-consistent simulation, in which the friction and diffusion coefficients are computed from first principles. We employ a two-dimensional domain decomposition approach within a message passing programming paradigm along with dynamic load balancing. Object oriented programming is used to encapsulate details of the communication syntax as well as to enhance reusability and extensibility. Performance tests on the SGI Origin 2000, IBM SP RS/6000 and the Cray T3E-900 have demonstrated good scalability. As a test example, we demonstrate the collisional relaxation to a final thermal equilibrium of a beam with an initially anisotropic velocity distribution.

VISUALIZATION

Chair: Patricia Crossno, Sandia National Laboratories

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 271/273

Using High-Speed WANs and Network Data Caches to Enable Remote and Distributed Visualization

Wes Bethel, Brian Tierney, Jason Lee, Dan Gunter, Stephen Lau,
Lawrence Berkeley National Laboratory

Visapult is a prototype application and framework for remote visualization of large scientific datasets. We approach the technical challenges of tera-scale visualization with a unique architecture which employs high speed WANs and network data caches for data staging and transmission. This architecture allows for the use of available cache and compute resources at arbitrary locations on the network. High data throughput rates and network utilization are achieved by parallelizing I/O at each stage in the application, and by pipelining the visualization process. On the desktop, the graphics interactivity is effectively decoupled from the latency inherent in network applications. We present a detailed performance analysis of the application, and improvements resulting from field-test analysis conducted as part of the DOE Combustion Corridor project.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 271/273

High Performance Visualization of Time-Varying Volume Data Over a Wide-Area Network Status

Kwan-Liu Ma, David M. Camp, University of California, Davis

This paper presents an end-to-end, low-cost solution for visualizing time-varying volume data rendered on a parallel computer located at a remote site. Pipelining and careful grouping of processors are used to hide I/O time and to maximize processor utilization. Compression is used to significantly cut down the cost of transferring output images from the parallel computer to a display device through a wide-area network. This complete rendering pipeline makes possible highly efficient rendering and remote viewing of high-resolution time-varying data sets in the absence of high-speed network and parallel I/O support. To study the performance of this rendering pipeline and to demonstrate high-performance remote visualization, tests were conducted on a PC cluster in Japan as well as an SGI Origin 2000 operated at the NASA Ames Research Center with the display located at UC Davis.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 271/273

Distributed Rendering for Scalable Displays
Greg Humphreys, Ian Buck, Matthew Eldridge, Pat Hanrahan, Stanford Univ.

We describe a novel distributed graphics system that allows an application to render to a large tiled display. Our system, called WireGL, uses a cluster of off-the-shelf PCs connected with a high-speed network. WireGL allows an unmodified existing application to achieve scalable output resolution on such a display. This paper presents an efficient sorting algorithm which minimizes the network traffic for a scalable display. We will demonstrate that for most applications, our system provides scalable output resolution with minimal performance impact.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 267

COMPILER OPTIMIZATION
Chair: Guang Gao, University of Delaware

Tiling Imperfectly-nested Loop Nests
Nawaaz Ahmed, Nikolay Mateev, Keshav Pingali, Cornell University

Tiling is one of the more important transformations for enhancing locality of reference in programs. Intuitively, tiling a set of loops achieves the effect of interleaving iterations of these loops. Tiling of perfectly-nested loop nests (which are loop nests in which all assignment statements are contained in the innermost loop) is well understood. In practice, many loop nests are imperfectly nested, so existing compilers use heuristics to try to find a sequence of transformations that convert such loop nests into perfectly-nested ones, but these heuristics do not always succeed. In this paper, we propose a novel approach to tiling imperfectly-nested loop nests. The key idea is to embed the iteration space of every statement in the imperfectly-nested loop nest into a special space called the product space which is tiled to produce the final code. We evaluate the effectiveness of this approach for dense numerical linear algebra benchmarks, relaxation codes, and the tomcatv code from the SPEC benchmarks. No other single approach in the literature can tile all these codes automatically.

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 267

Tiling Optimizations for 3D Scientific Computations
Gabriel Rivera, Chau-Wen Tseng, University of Maryland

Compiler transformations can significantly improve data locality for many scientific programs. In this paper, we show that iterative solvers for partial differential equations (PDEs) in three dimensions require new compiler optimizations not needed for 2D codes, since reuse along the third dimension cannot fit in cache for larger problem

sizes. Tiling is a program transformation compilers can apply to capture this reuse, but successful application of tiling requires selection of non-conflicting tiles and/or padding array dimensions to eliminate conflicts. We present new algorithms and cost models for selecting tiling shapes and array pads. We explain why tiling is rarely needed for 2D PDE solvers, but can be helpful for 3D stencil codes. Experimental results show tiling 3D codes can reduce miss rates and achieve performance improvements of 17-121 percent for key scientific kernels, including a 27 percent average improvement for the key computational loop nest in the SPEC/NAS benchmark mgrid.

Improving Fine-Grained Irregular Shared-Memory Benchmarks by Data Reordering

Y. Charlie Hu, Alan Cox, Willy Zwaenepoel, Rice University

WEDNESDAY
NOVEMBER 8
10:30AM – 12NN
ROOM D 267

We demonstrate that data reordering can substantially improve the performance of fine-grained irregular shared-memory benchmarks, on both hardware and software shared-memory systems. In particular, we evaluate two distinct data reordering techniques that seek to co-locate in memory objects in close proximity in the physical system modeled by the computation. The effects of these techniques are increased spatial locality and reduced false sharing.

We evaluate the effectiveness of the data reordering techniques on a set of five irregular applications from SPLASH-2 and Chaos. We implement both techniques in a small library, allowing us to enable them in an application by adding less than 10 lines of code. Our results on one hardware and two software shared-memory systems show that, with data reordering during initialization, the performance of these applications is improved by 12 percent to 99 percent on the Origin 2000, 30 percent to 366 percent on TreadMarks, and 14 percent to 269 percent on HLRC.

WEDNESDAY | 1:30PM – 3PM

APPLICATIONS II

Chair: David Walker, Oak Ridge National Laboratory

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 274

Performance Modeling and Tuning of an Unstructured Mesh CFD Application

William D. Gropp, Dinesh K. Kaushik*, Argonne National Laboratory, David E. Keyes, Old Dominion University, Barry F. Smith, Argonne National Laboratory

This paper describes performance tuning experiences with a three-dimensional unstructured grid Euler flow code from NASA, which we have reimplemented in the PETSc

framework and ported to several large-scale machines, including the ASCI Red and Blue Pacific machines, the SGI Origin, the Cray T3E and Beowulf clusters. The code achieves a respectable level of performance for sparse problems, typical of scientific and engineering codes based on partial differential equations, and scales well up to thousands of processors. Since the gap between CPU speed and memory access rate is widening, the code is analyzed from a memory-centric perspective (in contrast to traditional flop-orientation) to understand its sequential and parallel performance. Performance tuning is approached on three fronts: data layouts to enhance locality of reference, algorithmic parameters and parallel programming model. This effort was guided partly by some simple performance models developed for the sparse matrix-vector product operation.

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 274

Parallel Phylogenetic Inference
Quinn Snell, Michael Whiting, Mark Clement, David McLaughlin, Brigham Young University

Recent advances in DNA sequencing technology have created large data sets upon which phylogenetic inference can be performed. However, current research is limited by the prohibitive time necessary to perform tree search on even a reasonably-sized data set. Some parallel algorithms have been developed but the biological research community does not use them because they do not trust the results from newly developed parallel software. This paper presents a new phylogenetic algorithm that allows existing, trusted phylogenetic software packages to be executed in parallel using the DOGMA parallel processing system. The results presented here indicate that data sets that currently take as much as 11 months to search using current algorithms, can be searched in as little as two hours using as few as eight processors. This reduction in the time necessary to complete a phylogenetic search allows new research questions to be explored in many of the biological sciences.

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 274

Parallel Unsteady Turbo-pump Simulations For Liquid Rocket Engines
Cetin C. Kiris, ELORET/NASA-Ames, Dochan Kwak, NASA-Ames, William Chan, ELORET/NASA-Ames

This paper reports the progress being made towards complete turbo-pump simulation capability for liquid rocket engines. The Space Shuttle Main Engine (SSME) turbo-pump impeller is used as a test case for the performance evaluation of the MPI, hybrid MPI/Open-MP, and MLP versions of the INS3D code. Then, a computational model of a turbo-pump

has been developed for the shuttle upgrade program. Relative motion of the grid system for rotor-stator interaction was obtained by employing overset grid techniques. Unsteady computations for SSME turbo-pump, which contains 101 zones with 31 million grid points, are carried on Origin 2000 systems at NASA Ames Research Center. The approach taken for these simulations, and the performance of the parallel versions of the code are presented.

NETWORKING

Chair: John Mellor-Crummey, Rice University

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 267

The Failure of TCP in High-Performance Computational Grids

W. Feng, Los Alamos National Lab, P. Tinnakornsrisuphap, Univ. of Wisconsin

Distributed computational grids depend on TCP to ensure reliable end-to-end communication between nodes across the wide-area network (WAN). Unfortunately, TCP performance can be abysmal even when buffers on the end hosts are manually optimized. Recent studies blame the self-similar nature of aggregate network traffic for TCP's poor performance because such traffic is not readily amenable to statistical multiplexing in the Internet, and hence computational grids.

In this paper, we identify a source of self-similarity previously ignored, a source that is readily controllable—TCP. Via an experimental study, we examine the effects of the TCP stack on network traffic using different implementations of TCP. We show that even when aggregate application traffic ought to smooth out as more applications' traffic are multiplexed, TCP induces burstiness into the aggregate traffic load, thus adversely impacting network performance. Furthermore, our results indicate that TCP performance will worsen as WAN speeds continue to increase.

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 267

PSockets: The Case for Application-level Network Striping for Data Intensive Applications using High Speed Wide Area Networks

H. Sivakumar, University of Illinois at Chicago, S. Bailey, Infoblox Inc,
R. L. Grossman, University of Illinois at Chicago

Transmission Control Protocol (TCP) is used by various applications to achieve reliable data transfer. TCP was originally designed for unreliable networks. With the emergence of high-speed wide area networks various improvements have been applied to TCP to reduce latency and achieve improved bandwidth. The improvement is achieved by having system administrators tune the network and can take a considerable amount of time. This paper introduces PSockets (Parallel Sockets), a library that

achieves an equivalent performance without manual tuning. The basic idea behind PSockets is to exploit network striping. By network striping we mean striping partitioned data across several open sockets. We describe experimental studies using PSockets over the Abilene network. We show in particular that network striping using PSockets is effective for high performance data intensive computing applications using geographically distributed data.

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 267

Efficient Wire Formats for High Performance Computing
Fabian E. Bustamante, Greg Eisenhauer, Karsten Schwan, Patrick Widener, Georgia Institute of Technology

High performance computing is being increasingly utilized in non-traditional circumstances where it must interoperate with other applications. For example, online visualization is being used to monitor the progress of applications, and real-world sensors are used as inputs to simulations. Whenever these situations arise, there is a question of what communications infrastructure should be used to link the different components. Traditional HPC-style communications systems such as MPI offer relatively high performance, but are poorly suited for developing these less tightly-coupled cooperating applications. Object-based systems and meta-data formats like XML offer substantial plug-and-play flexibility, but with substantially lower performance. We observe that the flexibility and baseline performance of all these systems is strongly determined by their "wire format," or how they represent data for transmission in a heterogeneous environment. We examine the performance implications of different wire formats and present an alternative with significant advantages in terms of both performance and flexibility.

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 271/273

HARDWARE-BASED TOOLS

Chair: Al Malony, University of Oregon

Using Hardware Performance Monitors to Isolate Memory Bottlenecks
Bryan R. Buck, Jeffrey K. Hollingsworth, The Univ. of Maryland, College Park

In this paper, we present and evaluate two techniques that use different styles of hardware support to provide data structure specific processor cache information. In one approach, hardware performance counter overflow interrupts are used to sample cache misses. In the other, cache misses within regions of memory are counted to perform an n-way search for the areas in which the most misses are occurring. We present a simulation-based study and comparison of the two techniques. We

find that both techniques can provide accurate information, and describe the relative advantages and disadvantages of each.

Hardware Prediction for Data Coherency of Scientific Codes on DSM

J.T. Acquaviva, CEA DAM/PRISM, French Atomic Energy Commission, W. Jalby, PRISM, Versailles University

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 271/273

This paper proposes a hardware mechanism for reducing coherency overhead occurring in scientific computations within DSM systems. A first phase aims at detecting, in the address space regular patterns (called streams) of coherency events (such as requests for exclusive, shared or invalidation).

Once a stream is detected at a loop level, regularity of data access can be exploited at the loop level (spatial locality) but also between loops (temporal locality). We present a hardware mechanism capable of detecting and exploiting efficiently these regular patterns.

Expectable benefits as well as hardware complexity are discussed and the limited drawbacks and potential overheads are exposed.

For a benchmarks suite of typical scientific applications results are very promising, both in terms of coherency streams and the effectiveness of our optimizations.

A Scalable Cross-Platform Infrastructure for Application Performance Tuning Using Hardware Counters

S. Browne, Univ. of Tennessee, J. Dongarra, Univ. of Tennessee and Oak Ridge National Laboratory, N. Garner, K. London, P. Mucci*, Univ. of Tennessee

WEDNESDAY
NOVEMBER 8
1:30PM – 3PM
ROOM D 271/273

The purpose of the PAPI project is to specify a standard API for accessing hardware performance counters available on most modern microprocessors. These counters exist as a small set of registers that count "events", which are occurrences of specific signals and states related to the processor's function. Monitoring these events facilitates correlation between the structure of source/object code and the efficiency of the mapping of that code to the underlying architecture. This correlation has a variety of uses in performance analysis and tuning. The PAPI project has proposed a standard set of hardware events and a standard cross-platform library interface to the underlying counter hardware. The PAPI library has been or is in the process of being implemented on all major HPC platforms. The PAPI project is developing end-user tools for dynamically selecting and displaying hardware counter performance data. PAPI support is also being incorporated into a number of third-party tools.

WEDNESDAY | 3:30PM – 5PM

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 271/273

GORDON BELL I

Chair: Rusty Lusk, Argonne National Laboratory

A 1.349 Tflops Simulation of Black Holes in a Galactic Center on GRAPE-6

Junichiro Makino, Toshiyuki Fukushige, Masaki Koga, University of Tokyo

As an entry for the 2000 Gordon Bell performance prize, we report the performance achieved on a prototype GRAPE-6 system. GRAPE-6 is a special-purpose computer for astrophysical N-body calculations. The present configuration has 96 custom pipeline processors, each containing six pipeline processors for the calculation of gravitational interactions between particles. Its theoretical peak performance is 2.889 Tflops. The complete GRAPE-6 system will consist of 3072 pipeline chips and will achieve a peak speed of 100 Tflops. The actual performance obtained on the present 96-chip system was 1.349 Tflops, for a simulation of massive black holes embedded in the core of a galaxy with 786,432 stars. For a short benchmark run with 1,400,000 particles, the average speed was 1.640 Tflops.

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 271/273

92¢ /Mflops/s, Ultra-Large-Scale Neural-Network Training on a PIII Cluster

Douglas Aberdeen, Jonathan Baxter, Robert Edwards, Australian National Univ.

Artificial neural networks with millions of adjustable parameters and a similar number of training examples are a potential solution for difficult, large-scale pattern recognition problems in areas such as speech and face recognition, classification of large volumes of web data and finance. The bottleneck is that neural network training involves iterative gradient descent and is extremely computationally intensive. In this paper we present a technique for distributed training of Ultra Large Scale Neural Networks (ULSNN) on Bunyip, a Linux-based cluster of 196 Pentium III processors. To illustrate ULSNN training we describe an experiment in which a neural network with 1.73 million adjustable parameters was trained to recognize machine-printed Japanese characters from a database containing 9 million training patterns. The training runs with a average performance of 163.3 Gflops/s (single precision). With a machine cost of $150,913, this yields a price/performance ratio of 92.4¢ /Mflops/s (single precision). For comparison purposes, training using double precision and the ATLAS DGEMM produces a sustained performance of 70 Mflops/s or $2.16 /Mflop/s (double precision).

Scalable Molecular Dynamics for Large Biomolecular Systems
Robert K. Brunner, James C. Phillips, Laxmikant V. Kale, University of Illinois at Urbana-Champaign

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 271/273

We present an optimized parallelization scheme for molecular dynamics simulations of large biomolecular systems, implemented in the production-quality molecular dynamics program NAMD. With an object-based hybrid force and spatial decomposition scheme, and an aggressive measurement-based predictive load-balancing framework, we have attained speeds and speedups that are much higher than any reported in literature so far.

The paper first summarizes the broad methodology we are pursuing, and the basic parallelization scheme we used. It then describes the optimizations that were instrumental in increasing performance, and presents performance results on benchmark simulations.

PARALLEL PROGRAMMING

Chair: Barbara Chapman, University of Houston

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 274

A Comparative Study of the NAS MG Benchmark across Parallel Languages and Architectures
Bradford L. Chamberlai, Steven J. Deitz, Lawrence Snyder, University of Washington

Hierarchical algorithms such as multigrid applications form an important cornerstone for scientific computing. In this study, we take a first step toward evaluating parallel language support for hierarchical applications by comparing implementations of the NAS MG benchmark in several parallel programming languages: Co-Array Fortran, High Performance Fortran, Single Assignment C, and ZPL. We evaluate each language in terms of its portability, its performance, and its ability to express the algorithm clearly and concisely. Experimental platforms include the Cray T3E, IBM SP, SGI Origin, Sun Enterprise 5500 and a high-performance Linux cluster. Our findings indicate that while it is possible to achieve good portability, performance, and expressiveness, most languages currently fall short in at least one of these areas. We find a strong correlation between expressiveness and a language's support for a global view of computation, and we identify key factors for achieving portable performance in multigrid applications.

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 274

Is Data Distribution Necessary in OpenMP?
Dimitrios S. Nikolopoulos, Theodore S. Papatheodorou, University of Patras, Constantine D. Polychronopoulos, University of Illinois at Urbana-Champaign, Jesus Labarta, Eduard Ayguade, Technical University of Catalunya

This paper investigates the performance implications of data placement in OpenMP programs running on modern ccNUMA multiprocessors. Data locality and minimization of the rate of remote memory accesses are critical for sustaining high performance on these systems. We show that due to the low remote-to-local memory access latency ratio of state-of-the-art ccNUMA architectures, reasonably balanced page placement schemes—such as round-robin or random distribution of pages—incur modest performance losses. We also show that performance leaks stemming from suboptimal page placement schemes can be remedied with a smart user-level page migration engine. The main body of the paper describes how the OpenMP runtime environment can use page migration for implementing implicit data distribution and redistribution schemes without programmer intervention. Our experimental results support the effectiveness of these mechanisms and provide a proof of concept that there is no need to introduce data distribution directives in OpenMP and warrant the portability of the programming model.

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 274

Extending OpenMP for NUMA Machines
John Bircsak, Peter Craig, RaeLyn Crowell, Zarka Cvetanovic, Jonathan Harris*, C. Alexander Nelson, Carl D. Offner, Compaq Computer

This paper describes extensions to OpenMP which implement data placement features needed for NUMA architectures. OpenMP is a collection of compiler directives and library routines used to write portable parallel programs for shared-memory architectures. Writing efficient parallel programs for NUMA architectures, which have characteristics of both shared-memory and distributed-memory architectures, requires that a programmer control the placement of data in memory and the placement of computations that operate on that data. Optimal performance is obtained when computations occur on processors that have fast access to the data needed by those computations. OpenMP—designed for shared-memory architectures—does not by itself address these issues.

The extensions to OpenMP Fortran presented here have been mainly taken from High Performance Fortran. The paper describes some of the techniques that the Compaq Fortran compiler uses to generate efficient code based on

these extensions. It also describes some additional compiler optimizations, and concludes with some preliminary results.

THURSDAY | 10:30AM – 12NN

SOFTWARE TOOLS
Chair: Al Geist, Oak Ridge National Laboratory

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 274

A Tool Framework for Static and Dynamic Analysis of Object-Oriented Software with Templates
Kathleen A. Lindlan, Janice Cuny, Allen D. Malony, Sameer Shende, CIS, University of Oregon, Bernd Mohr, ZAM, Forschungszentrum Juelich, Reid Rivenburgh, CRAG, LANL, Craig Rasmussen, ACL, LANL

The developers of high performance scientific applications often work in complex computing environments that place heavy demands on program analysis tools. The developers need tools that interoperate, are portable across machine architectures, and provide source-level feedback. In this paper, we describe a tool framework, the Program Database Toolkit (PDT), that supports the development of program analysis tools meeting these requirements. PDT uses compile-time information to create a complete database of high-level program information that is structured for well-defined and uniform access by tools and applications. PDT's current applications make heavy use of advanced features of C++, in particular, templates. We describe the toolkit, focussing on its most important contribution—its handling of templates—as well as its use in existing applications.

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 274

From Trace Generation to Visualization: A Performance Framework for Distributed Parallel Systems
C. Eric Wu, Anthony Bolmarcich, Marc Snir, IBM T.J. Watson Research Center, David Wootton, Farid Parpia, IBM RS6000 Division, Anthony Chan, Ewing Lusk, William Gropp, Argonne National Laboratory

In this paper we describe a trace analysis framework, from trace generation to visualization. It includes a unified tracing facility on IBM SP systems, a self-defining interval file format, an API for framework extensions, utilities for merging and statistics generation, and a visualization tool with preview and multiple time-space diagrams. The trace environment is extremely scalable, and combines MPI events with system activities in the same set of trace files, one for each SMP node. Since the amount of trace data may be very large, utilities are developed to convert and merge individual trace files into a self-defining interval trace file

with multiple frame directories. The interval format allows the development of multiple time-space diagrams, such as thread-activity view, processor-activity view, etc., from the same interval file. A visualization tool, Jumpshot, is modified to visualize these views. A statistics utility is developed using the API, along with its graphics viewer.

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 274

Dynamic Software Testing of MPI Applications with Umpire
Jeffrey S. Vetter, Bronis R. de Supinski, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory

As evidenced by the popularity of MPI (Message Passing Interface), message passing is an effective programming technique for managing coarse-grained concurrency on distributed computers. Unfortunately, debugging message-passing applications can be difficult. Software complexity, data races, and scheduling dependencies can make programming errors challenging to locate with manual, interactive debugging techniques. This article describes Umpire, a new tool for detecting programming errors at runtime in message passing applications. Umpire monitors the MPI operations of an application by interposing itself between the application and the MPI runtime system using the MPI profiling layer. Umpire then checks the application's MPI behavior for specific errors. Our initial collection of programming errors includes deadlock detection, mismatched collective operations, and resource exhaustion. We present an evaluation on a variety of applications that demonstrates the effectiveness of this approach.

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 267

DATA GRID

Chair: Abdullah Almojel, King Fahd University

Computing and Data Grids for Science and Engineering
William E. Johnston, NASA Ames and Lawrence Berkeley National Laboratory, Dennis Gannon, NASA Ames and Univ. of Indiana, Bill Nitzberg, Veridian Systems, PBS Products, Leigh Ann Tanner, Bill Thigpen, Alex Woo, NASA Ames

We use the term "Grid" to refer to a software system that provides uniform and location independent access to geographically and organizationally dispersed, heterogeneous resources which are persistent and supported. While, in general, Grids will provide the infrastructure to support a wide range of services in the scientific environment (e.g. collaboration and remote instrument control) in this paper we focus on services for high performance computing and data handling. We describe the services and architecture of NASA's Information Power Grid ("IPG")—an

early example of a large-scale Grid—and some of the issues that have come up in its implementation.

The MicroGrid: a Scientific Tool for Modeling Computational Grids
H. J. Song, X. Liu*, D. Jakobsen, R. Bhagwan, X. Zhang, University of California, San Diego, K. Taura, University of California, San Diego, University of Tokyo, A. Chien, University of California, San Diego

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 267

The complexity and dynamic nature of the Internet (and the emerging Computational Grid) demand that middleware and applications adapt to the changes in configuration and availability of resources. However, to the best of our knowledge there are no simulation tools which support systematic exploration of dynamic Grid software (or Grid resource) behavior.

We describe our vision and initial efforts to build tools to meet these needs. Our MicroGrid simulation tools enable Globus applications to be run in arbitrary virtual grid resource environments, enabling broad experimentation. We describe the design of these tools, and their validation on micro-benchmarks, the NAS parallel benchmarks, and an entire Grid application. These validation experiments show that the MicroGrid can match actual experiments within a few percent (2 percent to 4 percent).

GORDON BELL II

Chair: Randy Bramley, Indiana University

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 271/273

1.34 Tflops Molecular Dynamics Simulation for NaCl with a Special-Purpose Computer: MDM
Tetsu Narumi, Ryutaro Susukita, Takahiro Koishi, RIKEN, Kenji Yasuoka, Keio University, Hideaki Furusawa, Atsushi Kawai, Toshikazu Ebisuzaki, RIKEN

We performed molecular dynamics (MD) simulation of 9 million pairs of NaCl ions with the Ewald summation and obtained a calculation speed of 1.34 Tflops. In this calculation we used a special-purpose computer, MDM, which we are developing for the calculations of the Coulomb and van der Waals forces. The MDM enabled us to perform large-scale MD simulations without truncating the Coulomb force. It is composed of WINE-2, MDGRAPE-2 and a host computer. WINE-2 accelerates the calculation for wavenumber-space part of the Coulomb force, while MDGRAPE-2 accelerates the calculation for real-space part of the Coulomb and van der Waals forces. The host computer performs other calculations. We performed MD simulation with the early version of the MDM system: 45 Tflops of WINE-2 and 1 Tflops of MDGRAPE-2. The peak performance of the final MDM system will reach 75 Tflops in total by the end of the year 2000.

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 271/273

High-Cost CFD on a Low-Cost Cluster

Thomas Hauser, Timothy I. Mattox, Raymond P. LeBeau, Henry G. Dietz, P. George Huang, University of Kentucky

Direct numerical simulation of the Navier-Stokes equations (DNS) is an important technique for the future of computational fluid dynamics (CFD) in engineering applications. However, DNS requires massive computing resources. This paper presents a new approach for implementing high-cost DNS CFD using low-cost cluster hardware.

After describing the DNS CFD code DNSTool, the paper focuses on the techniques and tools that we have developed to customize the performance of a cluster implementation of this application. This tuning of system performance involves both recoding of the application and careful engineering of the cluster design. Using the cluster KLAT2 (Kentucky Linux Athlon Testbed 2), while DNSTool cannot match the $0.64 per Mflops that KLAT2 achieves on single precision ScaLAPACK, it is very efficient; DNSTool on KLAT2 achieves price/performance of $2.75 per Mflops double precision and $1.86 single precision. Further, the code and tools are all, or will soon be, made freely available as full source code.

THURSDAY
NOVEMBER 9
10:30AM – 12NN
ROOM D 271/273

High Performance Reactive Fluid Flow Simulations Using Adaptive Mesh Refinement on Thousands of Processors

A.C. Calder, University of Chicago, B.C. Curtis, Lawrence Livermore National Laboratory, L.J. Dursi*, B. Fryxell, University of Chicago, G. Henry, Intel Corporation, P. MacNeice, K. Olson, NASA Goddard, P. Ricker, R. Rosner, F.X. Timmes, University of Chicago, H.M. Tufo, University of Chicago/Argonne National Laboratory, J.W. Turan, M. Zingale, University of Chicago

We present simulations and performance results of nuclear burning fronts in supernovae on the largest domain and at the finest spatial resolution studied to date. These simulations were performed on the Intel ASCI-Red machine at Sandia National Laboratories using FLASH, a code developed at the Center for Astrophysical Thermonuclear Flashes at the University of Chicago. FLASH is a modular, adaptive mesh, parallel simulation code capable of handling compressible, reactive fluid flows in astrophysical environments. FLASH is written primarily in Fortran 90, uses the Message-Passing Interface library for inter-processor communication and portability, and employs the PARAMESH package to manage a block-structured adaptive mesh that places blocks only where resolution is required and tracks rapidly changing flow features, such as detonation fronts, with ease. We describe the key

algorithms and their implementation as well as the optimizations required to achieve sustained performance of 238 GFLOPS on 6420 processors of ASCI-Red in 64 bit arithmetic.

PLENARY SESSION
THURSDAY | 1:30PM – 3PM

THURSDAY
NOVEMBER 9
1:30PM – 3PM
BALLROOM C

AWARDS SESSION

Chair: Sally Haerer, Oregon State University

A variety of achievements will be recognized and honored at the Awards Session during SC2000. The second annual IEEE computer Society Seymour Cray Computer Engineering Award will be presented in recognition of innovative contributions to high performance computing systems that best exemplify Seymour Cray's creative spirit. The IEEE Computer Society Sidney Fernbach Memorial Award will be presented for an outstanding contribution in the application of high performance computers using innovative approaches.

In addition to these society awards, the SC2000 Conference will present several other winners. The Gordon Bell Awards were established to reward practical uses of parallel processing and will be given for the best performance improvements in an application within several categories relating to hardware and software advancements. A newly created award for Best Network Application will be presented highlighting the most innovative, bandwidth-intensive application demonstration at SC2000. Awards will also be given for the best technical paper of the conference, the best student technical paper (with a student as principal author), the HPC Games winners, and the best Research Gem.

This session will feature a talk by Prof. Geoffrey M. Voelker, University of California, San Diego. Prof. Voelker is a Computing Research Association Digital Fellow. He will be speaking "On the Scale and Performance of Cooperative Web Proxy Caching."

THURSDAY | 3:30PM – 5PM

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 274

SCIENCE APPLICATIONS SUPPORT

Chair: Bernd Mohr, Research Centre Juelich

Integrating Parallel File I/O and Database Support for High Performance Scientific Data Management

Jaechun No, Rajeev Thakur, Argonne National Laboratory,
Alok Choudhary, Northwestern University

Many scientific applications have large I/O requirements, in terms of both the size of data and the number of files or data sets. Management, storage, efficient access, and analysis of these data present an extremely challenging task. Traditionally, two different solutions are used for this problem: file I/O or databases. File I/O can provide high performance but is tedious to use with large numbers of files and large and complex data sets. Databases can be convenient, flexible, and powerful but do not perform and scale well for parallel supercomputing applications. We have developed a software system, called Scientific Data Manager (SDM), which aims to combine the good features of both file I/O and databases. SDM provides a high-level API to the user and, internally, uses a parallel file system to store real data and a database to store application-related metadata. SDM takes advantage of various I/O optimizations available in MPI-IO, such as collective I/O and noncontiguous requests, in a manner that is transparent to the user. As a result, users can write and retrieve data with the performance of parallel file I/O, without having to bother with the details of actually performing file I/O.

In this paper, we describe the design and implementation of SDM. With the help of two parallel application templates, ASTRO3D and an Euler solver, we illustrate how some of the design criteria affect performance.

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 274

A Framework for Sparse Matrix Code Synthesis from High-level Specifications

Nawaaz Ahmed, Nikolay Mateev*, Keshav Pingali, Paul Stodghill, Cornell Univ.

We present compiler technology for synthesizing sparse matrix code from (i) dense matrix code, and (ii) a description of the index structure of a sparse matrix. Our approach is to embed statement instances into a Cartesian product of statement iteration and data spaces, and to produce efficient sparse code by identifying common enumerations for multiple references to sparse matrices. The approach works for imperfectly-nested codes with dependences, and produces sparse code competitive with hand-written library code for the Basic Linear Algebra Subroutines (BLAS).

A Unified Algorithm for Load-balancing Adaptive Scientific Simulations

Kirk Schloegel, George Karypis, Vipin Kumar, Army HPC Research Center, University of Minnesota

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 274

Adaptive scientific simulations require that periodic repartitioning occur dynamically throughout the course of the computation. The repartitionings should be computed so as to minimize both the inter-processor communications incurred during the iterative mesh-based computation and the data redistribution costs required to balance the load. Recently developed schemes for computing repartitionings provide the user with only a limited control of the tradeoffs among these objectives. This paper describes a new Unified Repartitioning Algorithm that can tradeoff one objective for the other dependent upon a user-defined parameter describing the relative costs of these objectives. We show that the Unified Repartitioning Algorithm is able to reduce the precise overheads associated with repartitioning as well as or better than other repartitioning schemes for a variety of problems, regardless of the relative costs of performing inter-processor communication and data redistribution. Our experimental results show that this scheme is extremely fast and scalable to large problems.

GRID MIDDLEWARE

Chair: Wolfgang Gentzsch, Sun Microsystems

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 271/273

The AppLeS Parameter Sweep Template: User-Level Middleware for the Grid

Henri Casanova, Graziano Obertelli, Francine Berman, University of California, San Diego, Rich Wolski, University of Tennessee, Knoxville

The Computational Grid is a promising platform for the efficient execution of parameter sweep applications over large parameter spaces. To achieve performance on the Grid, such applications must be scheduled so that shared data files are strategically placed to maximize re-use, and so that the application execution can adapt to the deliverable performance potential of target heterogeneous, distributed and shared resources. Parameter sweep applications are an important class of applications and would greatly benefit from the development of Grid middleware that embeds a scheduler for performance and targets Grid resources transparently.

In this paper we describe a user-level Grid middleware project, the AppLeS Parameter Sweep Template (APST), that uses application-level scheduling techniques and various Grid technologies to allow the efficient deployment of parameter sweep applications over the Grid. We discuss several possible scheduling algorithms and detail our software design. We then

describe our current implementation of APST using systems like Globus, NetSolve and the Network Weather Service, and present experimental results.

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 271/273

Requirements for and Evaluation of RMI Protocols for Scientific Computing

Madhusudhan Govindaraju, Aleksander Slominski, Venkatesh Choppella, Randall Bramley, Dennis Gannon, Indiana University

Distributed software component architectures provide a promising approach to the problem of building large scale, scientific Grid applications. Communication in these component architectures is based on Remote Method Invocation (RMI) protocols that allow one software component to invoke the functionality of another. Examples include Java remote method invocation (Java RMI) and the new Simple Object Access Protocol (SOAP). SOAP has the advantage that many programming languages and component frameworks can support it. This paper describes experiments showing that SOAP by itself is not efficient enough for large scale scientific applications. However, when it is embedded in a multi-protocol RMI framework, SOAP can be effectively used as a universal control protocol, that can be swapped out by faster, more special purpose protocols when large data transfer speeds are needed.

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 271/273

Expressing and Enforcing Distributed Resource Sharing Agreements

Tao Zhao, Vijay Karamcheti, New York University

Advances in computing and networking technology, and an explosion in information sources has resulted in a growing number of distributed systems being constructed out of resources contributed by multiple sources. Use of such resources is typically governed by sharing agreements between owning principals, which limit both who can access a resource and in what quantity. Despite their increasing importance, existing resource management infrastructures offer only limited support for the expression and enforcement of sharing agreements, typically restricting themselves to identifying compatible resources. In this paper, we present a novel approach building on the concepts of tickets and currencies to express resource sharing agreements in an abstract, dynamic, and uniform fashion. We also formulate the allocation problem of enforcing these agreements as a linear-programming model, automatically factoring the transitive availability of resources via chained agreements. A case study modeling resource sharing among ISP-level web proxies shows the benefits of enforcing transitive agreements: worst-case waiting times of clients accessing these proxies improves by up to two orders of magnitude.

SC2000 PANELS

Panels Chair: Philip Papadopoulos, SDSC

Panels offer SC2000 attendees the opportunity to become more involved in the conference, providing a question-and-answer session along with the opportunity to hear from the leaders in HPNC. This year, we are putting a strong emphasis on debate and differing opinions, to highlight the panel subject matter. And, as a tradition, the Friday panels serve as anchors to entice you to stay through to the very end of the conference. Don't miss the chance to participate in these discussions relevant to our industry and help make the sparks fly!

TUESDAY | 10:30AM – 12NN

TUESDAY
NOVEMBER 7
10:30AM – 12NN
ROOM D 262/264

VENTURE CAPITAL: WHO WANTS TO BE A BILLIONAIRE?

Moderator: Steve Wallach, CenterPoint Ventures/Chiaro Networks
Panelists: Matt Blanton, Startech
Jackie Kimzey, Sevin Rosen Funds
Scott Grout, Chorum

Okay, let's play! For $100, what is an IPO?

A. Immediate Public Outcry
B. Illogical Policy—Ongoing
C. Internal Political Overthrow
D. Initial Public Offering

The right answer is actually E: an Incredible Panel Opportunity to hear four of the best venture capital "life-lines," talking about turning innovative ideas into business success. Key questions to be addressed include: how do you select the right strategic partners? And what are the most common potential stumbling blocks?

WEDNESDAY | 3:30PM – 5PM

WEDNESDAY
NOVEMBER 8
3:30PM – 5PM
ROOM D 262/264

IS TCP AN ADEQUATE PROTOCOL FOR HIGH PERFORMANCE COMPUTING NEEDS?

Moderator: Hilarie Orman, Novell
Panelists: Jamshid Mahdavi, Novell
Volker Sander, Central Institute for Applied Mathematics (ZAM)
John von Neumann Institute for Computing (NIC)
Wu-chun Feng, Los Alamos National Laboratory and Purdue Univ.
Stuart Bailey, Infoblox

Lawrence Brakmo, Compaq
Deepak Bansal, MIT LCS
Brian L. Tierney, Lawrence Berkeley National Laboratory

TCP/IP is the standard for data transmission in the web era. This fundamental protocol has undergone many changes to continually improve its dynamic range for performance and stability. TCP itself is essential for stability of today's wide-area networks. However, this protocol isn't for everyone. For example, high performance cluster builders dispense with TCP as being too heavyweight and tuning TCP for high performance (e.g. Gigabit) WAN connections is very time consuming.

This panel will address issues of how the high performance community and the massive consumer Internet-at-large community can converge on a protocol that satisfies all emerging networking needs? Some questions include: where are the fundamental problems and can they be overcome? What is the current level of performance that can be attained and how difficult was it to tune the network?

THURSDAY | 3:30 PM – 5PM

THURSDAY
NOVEMBER 9
3:30PM – 5PM
ROOM D 262/264

PETAFLOPS AROUND THE CORNER: WHEN? HOW? IS IT MEANINGFUL?

Moderator: Neil Pundit, Sandia National Laboratories
Panelists: Marc Snir, IBM Research
Bill Camp, Sandia National Laboratories.
Thomas Sterling, JPL/NASA/Caltech
Paul Messina, DOE HQ
Rick Stevens, Argonne National Laboratory
Pete Beckman, Turbolabs

Teraflops-capable machines are being deployed now and a petaflops-capable computing complex is just around the corner. This panel will address questions about the necessity of computers with this level of power, what are the competing design issues to attain a petaflop, and can applications actually make good use of such an extreme capability by the end of this decade. Expert panelists represent competing design schemes, extreme application drivers and dissenting opinions about the need for this level of power in this decade.

FRIDAY | 8:30AM – 12NN

CONVERGENCE AT THE EXTREMES: COMPUTATIONAL SCIENCE MEETS NETWORKED SENSORS

FRIDAY
NOVEMBER 10
8:30AM – 10AM
BALLROOM C 1/2

Moderator: David Culler, UC Berkeley
Panelists: Deborah Estrin, UCLA
Larry Arnstein, University of Washington
James Demmel, UC Berkeley
Paul J. McWhorter, Sandia National Laboratories

The SC conference has ridden the technology waves of the "killer micro" and the "killer network" over the past decade. Two new "killer technologies" are emerging that hold the potential to radically change the nature of computational science. Micro-electromechanical systems (MEMS) make it possible to incorporate tiny sensors into all sorts of computational devices. Low-power CMOS radios make it possible to communicate with minute devices deeply embedded in the physical environment. Not only do these technologies fundamentally change the degree of instrumentation possible in experiments, they provide a unique means of linking computational science and real-time measurement in novel environments. This panel will explore the potential of these emerging technologies, what really is new, and what difference it all makes.

COMPUTATIONAL GRIDS: A SOLUTION LOOKING FOR A PROBLEM?

FRIDAY
NOVEMBER 10
8:30AM – 10AM
BALLROOM C 3/4

Moderator: Jennifer Schopf, Northwestern University
Panelists: Ian Foster, Argonne National Laboratory
Cherri Pancake, Oregon State University
Marc Snir, IBM Research
Geoffrey Fox, Syracuse University

The computational grid, or simply "The Grid" is the latest buzz in computing. Among all the hype are real problems and issues that need addressing before the Grid becomes a ubiquitous and transparent service. Questions that this panel will address include: Will Grid computing "make it" and be successful? What will the Grid really look like? What's the most important problem still to solve to make it successful? And, how do Grids differ from what's being done already in distributed computing, clusters, and OS design?

OPEN SOURCE: IP IN THE INTERNET ERA

FRIDAY
NOVEMBER 10
10:30AM – 12NN
BALLROOM C 1/2

Moderator: Robert Borchers, NSF
Panelists: Susan Graham, UC Berkeley
Richard Gabriel, Sun Microsystems
Todd Needham, Microsoft Research
José Muñoz, Dept. of Energy

Open source operating systems, tools and community scientific codes represent a fundamental shift in the ownership and intellectual property of software. Of particular importance is the fact that open source can provide continuity of the high-performance software stack as underlying hardware technology undergoes dramatic changes. However, open source gives up control of critical intellectual property and has become a two-edged sword, igniting passions from all communities that focus on software

This panel will address the following questions: Can the open source paradigm be applied to other disciplines and should the term become open intellectual property? Can open intellectual property be a component of a profitable business plan? What constitutes truly open source code/intellectual property? And, what software needs to be open source to serve the needs of the high-end/scientific computing user community?

FRIDAY
NOVEMBER 10
10:30AM – 12NN
BALLROOM C 3/4

MEGACOMPUTERS

Moderator: Larry Smarr, UC San Diego
Panelists: Andrew Chien, Entropia
Ian Foster, Argonne National Laboratory
Thomas Sterling, JPL
David Anderson, United Devices
Andrew Grimshaw, University of Virginia

We are nearing a major discontinuity in the evolution of high-performance computing, creating a third era of supercomputing. The first era, which ended with the Cray 1, sought performance by building the fastest single processor possible. The second era, starting with the Illiac IV and continuing to today's tightly coupled computing complexes (PC Superclusters, MPPs, SMPs, and DSMs) sought performance through tens to thousands of identical processors. The third era couples very large numbers of heterogeneous computers on networks to create a virtual supercomputer. This third era is now going through an important phase change as numerous academic and private companies move from the intranet to the Web itself, harnessing tens of thousands to millions of PCs on the Net.

Questions to be addressed by the panel include: What is the software architecture model? Can general-purpose computing machines be created in this manner? What are the classes of applications most likely to move to this new computing fabric? And what new computer science research frontiers are important in this more "biological" effervescent style of architecture?

SC2000 HPC GAMES

HPC Games Co-Chairs:
Eleanor Anne Schroeder, Naval Oceanographic Office
James Arthur Kohl, Oak Ridge National Laboratory

"Speed has always been important, otherwise we wouldn't need the computer." These are the words of Seymour Cray, the father of supercomputing. In the spirit of Seymour Cray's vision, this year's HPC Games is thrilled to introduce a new and exciting challenge: The $10K Computer Challenge. Participants are being asked to build their own high-performance computing machine (or machines) for under $10K, and bring these machines to the SC2000 exhibition floor to compete, running a series of predetermined benchmarks. The benchmarks will exercise a balanced assortment of metrics covering CPU performance, disk performance and network performance, and various combinations of these metrics as seen in typical high performance applications.

Scoring and qualifying rules, as well as information on benchmarks, can be found at the web site www.sc2000.org/games.

Remember that bigger is not always better, and some pretty good things do come in small packages. Do you DARE to accept the HPC Games Challenge? The 19th century gave us Charles Babbage. The 20th century gave us Seymour Cray. Who knows? The winner of this contest may become the 21st century's computing visionary…

SC2000 EDUCATION PROGRAM

Education Vice Chair: Jeffrey C. Huskamp, East Carolina University
Education Program Co-Chairs: Lisa Bievenue, NCSA and Edna Gentry, University of Alabama at Huntsville

SC2000 offers high school teachers an opportunity to learn computer modeling and simulation, and their application to science and mathematics curricula. After the conference the selected teachers will receive ongoing support from SC2000 staff and continue their training at regional forums. At SC2001 the teachers will share their experiences with a new group of teachers and will be expected to be leaders in their school systems and region for a wider adoption of modeling and simulation by classroom teachers.

Teams of participating teachers were selected in July 2000. Teachers local to the Dallas area and other teachers who don't need accommodations or travel support are also welcome to attend on a space available basis (advanced registration is required; send email to education@sc2000.org). SC2000 Technical Program Registrants are welcome to attend the Education Program, although participation in the hands-on sessions and individual team meetings will be limited to the pre-selected teacher teams. No pre-registration in the Education Program is needed for SC2000 Technical Program Registrants to attend the Education Program sessions.

EDUCATION PROGRAM SPEAKERS

Richard Allen, Albuquerque High Performance Computing Center
Jonathan Barnes, Maui High Performance Computing Center
Lisa Bievenue, NCSA
Ken Flurchick, Ohio Supercomputing Center
Edna E. Gentry, University of Alabama in Huntsville
Robert R. Gotwals, Jr., Shodor Education Foundation
Barb Helland, Krell Institute
Jeffrey C. Huskamp, East Carolina University
Eric Jakobbson, NCSA
Scott A. Lathrop, NCSA
Cynthia Lanius, Rice University
Robert M. Panoff, Shodor Education Foundation
Helen M. Parke, East Carolina University
Susan Ragan, Maryland Virtual High School
Deanna Raineri, University of Illinois, Urbana-Champaign
Todd Veltman, West Leyden High School
Mary Ellen Verona, Maryland Virtual High School
Bill Wiecking, Maui High Performance Computing Center

EDUCATION PROGRAM SPONSORS

National Science Foundation
Association for Computing Machinery
Compaq Computer Corporation
High Performance Systems
IEEE Computer Society
Microsoft Corporation
National Aeronautics and Space Administration
SC Conference
Shodor Education Foundation
Wolfram Research, Inc.

—ALL SESSIONS ARE HELD ON LEVEL 2, C SECTION OF THE DALLAS CONVENTION CENTER UNLESS OTHERWISE NOTED—

SATURDAY, NOVEMBER 4

SATURDAY
NOVEMBER 4
4PM – 9PM
LEVEL 2
C SECTION

1PM – 5PM **Teacher Registration**

4PM – 5PM **Basic Skills Tutorial**

5PM – 6PM **Orientation and Laptop Setup**

7PM – 9PM **Education Program Reception**

SUNDAY, NOVEMBER 5

SUNDAY
NOVEMBER 5
9AM – 6:30PM
LEVEL 2
C SECTION

9AM – 9:30AM **Registration**

9:30AM – 10AM **Welcome and Announcements**

10AM – 6:30PM **Introduction to Computational Science**
Computational science is the process of doing science with the aid of computer-based simulation, modeling, and visualization technologies. This introduction will present the role of computation in scientific method and discovery, the process of computational science, and several modeling and visualization tools used by scientists.

Overview of Computational Modeling
Instructors in this session will demonstrate the process of computational modeling using Stella modeling software. Participants will have opportunities to modify models on their own computer.

Overview of Mathematical Modeling
Instructors in this session will demonstrate the process of computational modeling using Mathematica modeling software. Participants will have opportunities to modify models on their own computer.

Overview of Visualization
Instructors in this session will demonstrate the process of scientific visualization. Participants will have opportunities to create visualizations on their own computer.

Overview of Curriculum Module Development
Instructors will present the process of developing curriculum modules that incorporate the use of simulation, modeling and visualization. Teams will also begin discussing their team module.

MONDAY, NOVEMBER 6

During each of the Modeling Sessions (Modeling Part 1, 2, and 3) participants will be divided into smaller groups with each group rotating through the featured modeling tools: Stella, Microsoft Excel, and Mathematica. During these sessions, participants will receive the tool and will build several models using that tool.

Stella The STELLA software lets users build, simulate, and communicate models of dynamic processes. Students use the software to render, test, and share their mental models of everything from how a bowl of soup cools to how a galaxy expands. (High Performance Systems, http://www.hps-inc.com/edu/stella/stella.htm)

Microsoft Excel is a spreadsheet software application. With the aid of Excel's library of functions and data manipulation tools, many kinds of mathematical models can be constructed. (Microsoft Corporation, http://www.microsoft.com/office)

Mathematica is a sophisticated mathematical modeling and visualization software application. "It is famous for its high-quality, three-dimensional graphics, its ability to handle arbitrary-precision arithmetic, [and] its symbolic-processing abilities." (Ian Sammis, *MacAddict*; December 1, 1999) (Wolfram Research, http://www.wolfram.com)

8:30AM – 12NN **Modeling, Part 1**

1PM – 4:30PM **Modeling, Part 2**

MONDAY
NOVEMBER 6
8:30AM – 7PM
LEVEL 2
C SECTION

4:30PM – 5:30PM **Developing Curriculum Modules**
Instructors will present the requirements and goals for developing the team curriculum modules. Teams will continue discussing their team module.

5:30 – 7PM **Team Meetings**

MONDAY
NOVEMBER 6
EXHIBIT HALL

7PM – 9PM **Gala Opening Reception, Exhibit Hall**

TUESDAY, NOVEMBER 7

TUESDAY
NOVEMBER 7
8:30AM – 9PM
LEVEL 2
C SECTION

8:30AM – 10AM **SC Conference Keynote Speaker**

10:30AM – 12NN **Modeling, Part 3**

1PM – 2:30PM **Modeling, Part 3** (continued)

2:30PM – 5PM **Team Meetings/Open Lab**
Exhibit Hall Guided Tours

7PM – 9PM **Open Lab**

WEDNESDAY, NOVEMBER 8

WEDNESDAY
NOVEMBER 8
8:30AM – 9PM
LEVEL 2
C SECTION

8:30AM – 12NN **Web-based Tools**
(concurrent sessions, pick any two of the six)

ChemViz
ChemViz includes two web-based chemistry visualization tools that help students understand and "experiment" with atomic and molecular structures. Other chemistry visualization tools will also be demonstrated.

Interactivate
Project Interactivate is a suite of interactive web-based middle school level mathematics lessons, discussions, and activities.

Physics Tools
Several interactive web-based physics tools will be demonstrated, including GalaxSee (galaxy formation), SimSurface (computational techniques in the minimization of potential energy), and The Pit and the Pendulum (simple machines).

BiologyWorkbench
The Biology Workbench is a computational interface and environment that can be used to search many popular protein and nucleic acid sequence databases. Database searching is integrated with access a wide variety of analysis and

modeling tools, all within a point and click interface that eliminates file format compatibility problems.

The Bowfin Project
The USS Bowfin Submarine Museum & Park and the Maui High Performance Computing Center have joined forces to celebrate the 100th anniversary of the US Submarine Navy. A submarine-science/education model describes a historical timeline of submarine development by highlighting five US submarines and examining one or more physics concepts and associated scientific principles for each. Mathematical models are developed to illustrate the various physics concepts, many of which can be implemented in Excel; others require the use of high-performance computing.

GirlTech
GirlTECH, a teacher professional development program, attacks the problem of girls' under-use of computers with a two-week program that first makes teachers aware of the underrepresentation of women in IT, then promotes strategies for teachers to encourage girls to become part of the IT forces.

1PM – 2:30PM **Curriculum Integration and Assessment**
In this session, instructors and participants will discuss the interdisciplinary nature of curriculum modules, introduce the process for loading modules into a database for web access, and present an overview of the assessment procedure for collecting evidence of student learning and teacher professional growth.

3PM – 5PM **Team Meetings**

7PM – 9PM **Open Lab**

THURSDAY, NOVEMBER 9

8:30AM – 9:30AM **Outreach/Mentoring Goals**
This session will present leadership opportunities and strategies for outreach to other teachers.

10AM – 11:30AM **Module Presentations by Teacher Teams**

12:30 – 1:30PM **Module Presentations by Teacher Teams**

1:30PM – 2PM **Education Program Evaluation Wrap-up**

Next year the SC2001 Education Program will again offer high school teachers an opportunity to participate in the

THURSDAY
NOVEMBER 9
8:30AM – 2PM
LEVEL 2
C SECTION

National Computational Science Leadership Program. The program will focus on developing a core group of teachers who are prepared to utilize computational science to enhance science and math education and to share this knowledge and expertise with other teachers. Twenty-five teams of four teachers will be selected to attend SC2001 in Denver, Colorado and a two-week summer workshop. Teacher teams will develop learning modules that will become a permanent part of the national online computational science resource repository.

If you are a high school math or science teacher and are interested in finding out more about this unique program visit the SC2001 Education Program website: www.sc2001.org/education/

Applications are due on May 31, 2001.
Questions: education@sc2001.org

VENTURE VILLAGE

Conference Vice Chairs, Venture Village:
Dennis Duke, Florida State University
Steve Wallach, CenterPoint Ventures/Chiaro Networks

An innovative new exhibit joins SC2000 to showcase the best of the best! The Venture Village will showcase a collection of entrepreneurial information technology companies, which are creating new products for the 21st century. The Venture Village will be located strategically on the Main exhibit floor, and will offer a village atmosphere for visiting with the venture companies and/or your colleagues.

www.sc2000.org/venture

eSCape 2000

eSCape 2000 Co-Chairs: Stephen Jones, U.S. Army ERDC
John West, WorldCom

SC2000 will explore the potential effects of the convergence of personal computing devices, the software infrastructure necessary to support them, and wireless connectivity and what it means to do "HPC Anywhere" in an event called eSCape 2000. The conference will provide a wireless infrastructure as part of SCinet 2000 to serve as a showcase for demonstrations of novel interfaces to HPC resources, data and applications. The tone of the event is that of a demonstration—not a competition—and the goal is to provide a forum for stimulating broad community discussions about how these new technologies will impact HPC. Participants are invited to demonstrate applications and technologies that they feel best illustrate how being able to reach "HPC Anywhere" will impact supercomputing and the people and industries that depend on supercomputing.

The event will be organized as a combination forum/permanent exhibit. The eSCape 2000 events will be held on the exhibit floor within a reserved section co-located with Research Gems. A presenter from each project will be available Wednesday and Thursday from 10AM – 11AM during exhibit hours for a presentation and demonstration of their research project. In addition, each team will be provided space within the eSCape 2000 venue to set up a permanent poster describing the project. The poster exhibit will remain up for the duration of the exhibit to provide those who cannot attend one of the presentations the opportunity to learn about the project.

www.sc2000.org/escape2000

eSCape 2000 Demonstrations
SDSC GridPort Toolkit Applications
Stephen Mock, SDSC/UCSD
Mary Thomas, Scientific Computing Department, SDSC

We will be demonstrating wireless access to applications that have been built using the SDSC GridPort Toolkit (https://hotpage.npaci.edu). These applications include the NPACI HotPage (https://hotpage.npaci.edu) and other computational science web portals. Access will allow users to view current HPC system information, including machine and job status.

The PUNCH Computing Portal: Integrating Grid Services and Web Technologies for High Performance Computing Anytime, Anywhere
Jose Fortes, Rudolf Eigenmann, Mark Lundstrom, Purdue University, Valerie Taylor, Northwestern University, Miron Livny, University of Wisconsin-Madison, Sumalatha Adabala, Renato Figueiredo, Nirav Kapadia, Purdue University, Jose' Miguel-Alonso, UPV, Spain

The demonstration will highlight computing portal and wide-area network-computing technologies that allow seamless management of high-performance applications, data and machines distributed across wide-area networks. These technologies allow users to access and run applications (even unmodified commercial applications) from anywhere via standard Web browsers. The presentation will outline the manner in which state-of-the-art technologies can be used to build end-to-end solutions for high-performance, ubiquitous supercomputing on Intranets and the Internet. The demonstration will be based on PUNCH, a network computing system that has been operational for five years. PUNCH currently provides computing services to about 800 users across ten countries; 50 engineering software packages developed by 13 universities and six vendors are available. PUNCH can be accessed at www.ece.purdue.edu/punch. The presentation will include an overview of the utility and benefits of such technologies in the context of several multi-university projects for collaborative engineering and distance education.

Speech-to-Text & Speech-to-Sign eSCapes Hearing Barriers
Gottfried E. Zimmermann, Gregg Vanderheiden, University of Wisconsin-Madison Trace R&D Center, Dan Deignan, Personal Captioning

If you can't hear, it can be hard to participate in time-critical discussions. This can be because you are on a noisy factory floor, or because of functional impairment in hearing.

Networked services, and particularly mobile unWired services could deliver more pervasive access to help in this area. Last year we demonstrated using Internet2 to make remote sign language interpreting feasible. This year we are experimenting with options for real-time transcription of speech to text and translation from speech to sign for collaborative participation of people with hearing impairments. In addition to exploring the capabilities of various strategies in and around the convention, we will demonstrate the techniques we are investigating to visitors at the eSCape 2000 area in the exhibition hall.

Middleware and the eSCaped Web

Alfred S. Gilman, Gottfried E. Zimmermann, University of Wisconsin-Madison, Trace R&D Center

Network middleware to adapt web content for mobile devices is a very dynamic area. New players are emerging on all sides. There has been work in this area for the purpose of disability access for several years, and the W3C Note on guidelines for mobile-usable documents relies heavily on the Web Content Accessibility Guidelines 1.0 Recommendation developed for disability access purposes. We will be demonstrating how various middleware strategies perform against two usability criteria: fitness for use with small mobile devices, and fitness for use by people with disabilities.

The IU Enotebook

Dennis Gannon, Randall Bramley, Madhu Govindaraju, Benjamin Temko, Ken Chiu, Juan Villacis, Indiana University

The IU Enotebook allows access to and control of Grid resources from any browser-equipped device using only standard Web protocols. Features of the Enotebook include:

1. Authentication via GSSAPI and Globus certificates of authority
2. A Python scripting mechanism for locating, instantiating, connecting, and controlling Common Component Architecture HPC components across the grid
3. A dual-level user model which includes script developers and end-users who select and run scripts for experiments
4. Java servlet-based framework to allow users to disconnect and later reattach to running experiments rom different IP addresses
5. Tools for automatically handling files in a distributed

computation, and providing a user with a full record of their experiments
6. Component managers for rapid encapsulation and use of existing binaries which cannot be recompiled or directly linked with a component framework
7. Support for multiple languages and communications protocols.

	TUESDAY	WEDNESDAY	THURSDAY
10AM – 10:30	**Implementing Intelligent Infrastructures Using SAN Appliances** Robert Woolery, VP of Corporate Development and Strategic Planning, DataDirect Networks,Inc. **PAGE 95**	**The Challenge of Next Generation High Performance Networks** John Freisinger, Essential Communications **PAGE 98**	**Use of Juniper Routers in Research and Education Networks** John Jamison, Juniper Networks **PAGE 102**
10:30 – 11AM	**Tarantella—Web Enable Any Application on Any Platform and Access from Anywhere** Ron Green, Unicomp, Inc. **PAGE 95**	**"Teaming" Technology —STAGGERING PERFORMANCE POWER THROUGH SIMPLICITY** Lynn West, CTO, Times N Systems **PAGE 98**	**TurboLinux Cluster Cockpit** Peter H Beckman, TurboLinux **PAGE 102**
11AM – 11:30	**Supercomputing— High Bandwidth and High Performance** David Parks, NEC–HNSX Supercomputers **PAGE 95**	**New capabilities and plans for High Performance Storage System** Ramin Nosrat (tentative), High Performance Storage System (HPSS) **PAGE 99**	**Enabling HPC Research in DoD** Tim Campbell, Logicon **PAGE 102**
11:30 – 12NN	**Liberating Design Teams Via the Internet** Tommy Steele, Alibre, Inc. **PAGE 96**	**Legion—An Applications Perspective** Andrew Grimshaw, Applied MetaComputing, Inc. **PAGE 99**	**How Clusters Fill a Vendor Void** David Rhoades, High Performance Technologies, Inc. **PAGE 102**
12:30 – 1PM	**Open Source Community** John H. Terpstra, TurboLinux **PAGE 96**	**Deep Blue Update** David Turek, IBM Vice President, Deep Computing, Web Server Division **PAGE 99**	**Internet Computing: Distributed Computing Leads Third Generation Internet Applications** Jim Madsen, President and CEO, Entropia **PAGE 103**
1PM – 1:30	**Building a 30 TeraOPS Supercomputing System** Richard Kaufmann, Compaq Computer Corporation **PAGE 96**	**High Performance Alpha Clusters** Tom Morris, API **PAGE 100**	**New Compiler Features for High Performance Computing** Doug Miles, PGI **PAGE 103**

	TUESDAY	WEDNESDAY	THURSDAY
1:30 – 2PM	**Objectivity/DB The High Performance Database Engine** Leon Guzenda, CTO, Objectivity **PAGE 97**	**Linux on the Bleeding Edge, CPlant and FSL** Greg Lindahl, High Performance Technologies, Inc. **PAGE 100**	**StorageTek** **PAGE 103**
2PM – 2:30	**Fujitsu's High Performance and Highly Scalable Supercomputers and Servers** Kenichi Miura, Fujitsu **PAGE 97**	**Remote Management of a Storage Network** Derek Gamradt, StorNet, Inc. **PAGE 100**	**MAN/LAN/WAN— Blurring Boundaries for Switching and Routing** Earl Ferguson, Chief Technical Officer of Foundry Networks **PAGE 103**
2:30 –3PM	**Concept-to-Production: End to End High Performance Network Monitoring and Performance Analysis Solutions** Erik Plesset, Spirent Communications **PAGE 97**	**The Art of Internet Computing** Dr. Steve Armentrout, CEO, Parabon **PAGE 101**	**Optimization and Parallelization with the Intel Compilers and the KAI Tools** Joe H. Wolf, Intel **PAGE 104**
3PM – 3:30	**High Bandwidth Supercomputers of Today and Tomorrow** Dave Kiefer & Burton Smith, Cray, Inc. **PAGE 98**	**Delivering a New World of Flexibility for HPC** Ben Passarelli, SGI **PAGE 101**	**Large-Scale Distributed Computing: The Benefits and Advantages of Compute Farms** Rich Edelman, Blackstone Technology Group **PAGE 104**
3:30 – 4PM	**Cost Effective High Performance Computing on Clusters of Commodity Processors** Robert D. Bjornson, Scientific Computing Associates, Inc. **PAGE 98**	**Real-World Performance Analysis** Karl Solchenbach, Pallas **PAGE 101**	**SRC Computers Inc: Background, Status, and Plans** **PAGE 104**
4PM – 4:30	**Collaborative Visualization Solutions for HPC Involving Reality Center and Immedsive Visualization Products** Linda Jacobson, SGI **PAGE 98**	**InfiniBand™ Management Solutions** VIEO (formerly Power Micro Research) **PAGE 101**	**Sun Microsystems** **PAGE 104**

SC2000 EXHIBITOR FORUM

Exhibitor Forum Chair: Joe McCaffrey, Mississippi State University

TUESDAY | NOVEMBER 7
ROOM D 268

10AM – 10:30AM

Implementing Intelligent Infrastructures Using SAN Appliances

Robert Woolery, Vice President of Corporate Development and Strategic Planning, DataDirect Networks, Inc.

With massive bandwidth and storage capacities readily available, and servers, desktops and supercomputers linked like never before, new applications and rich media delivery methods promise to bring real world data to individuals and groups worldwide.

This paper will explore the use of cutting edge, intelligent network technology that enables lower latency, high-speed data access within a new extended network architecture—providing high performance computing environments with high Quality of Service (QoS) and greater user experience. Extensive discussion will be devoted to creating flexibility and scalability in these new architectures with "data aware" network technology deemed necessary in building next generation infrastructures.

This discussion will detail how data-aware technology will enable scalable lower latency, data access with an intelligent network infrastructure device developed to manage and transport vast quantities of text, numeric data, audio and video images at very high speeds to and from multiple servers and clustered workstations to sheared storage resources.

10:30 – 11AM

Tarantella—Web Enable Any Application on Any Platform and Access from Anywhere

Ron Green, Unicomp, Inc.

Tarantella lets you use the web to deliver any application immediately to any client—without any client software installation. This revolutionary solution ensures users have access to the latest applications and data, wherever they are on the 'net. In addition, the innovative Tarantella Adaptive Internet Protocol (AIP) ensures optimal network performance, even over low-bandwidth connections. It offers a powerful and cost-effective method for deployment over the web using standard browser technology.

With Tarantella, your Windows users can have access to your Unix or Linux applications without using X-Windows managers. Your Linux and Unix desktop users can remotely run any Windows applications as if they were sitting in front of the PC. All of this with no rewrites and no client installs.

Come see the future of web-enabled application delivery!

We are a Caldera Business partner, providing solutions, training, and consultation in the field of high performance computing.

11AM – 11:30AM

Supercomputing—High Bandwidth and High Performance

David Parks, NEC–HNSX Supercomputers

Industrial users around the world continue to rely on the performance provided by parallel vector supercomputers for their mission critical applications. NEC, as the leading provider of parallel vector systems in the world today, continues to

improve the technologies and make them more cost effective.

The session will discuss the differences between commodity scalar processors and vector processors, how that effects performance on various applications, and how both scalar and vector processors are evolving through enabling technologies.

11:30AM – 12NN

Liberating Design Teams Via the Internet

Tommy Steele, Alibre, Inc.

The appeal of the Internet is based on the lower requirement for hard assets, a direct pipeline to customers, and a new level of speed and operational efficiency. For the mechanical design and manufacturing industries in particular, the Internet represents the next wave of value innovation through a new business model—the Application Service Provider (ASP).

The ASP business model for the MCAD market is compelling because it removes the barriers associated with traditional mechanical design products. The high cost of ownership—software, hardware, IT supports and system administration—is significantly reduced. Distributed computing power and storage capacity made available via the Internet is potentially unlimited. In addition, the Internet also enables real-time team collaboration between best-in-class experts, suppliers, contractors, engineers and others—allowing companies to get products to the market faster and cheaper, with reduced development time, fewer design changes and less travel expense.

12:30PM – 1PM

Open Source Community

John H. Terpstra, TurboLinux

The Open Source Software community is turning much of the information technology world upside down. This presentation identifies the key developments that will impact the supercomputing environment, looks at who and what is driving the key projects as well as the companies that are trying to gain maximum benefit from these initiatives.

The spectrum of concerns and opportunities are elucidated and each is colorfully articulated from both the perspective of the Open Source community as well as from that of the end user. Business and Organizational users need to see how the Open Source community addresses the process of problem identification and resolution.

1PM – 1:30PM

Building a 30 TeraOPS Supercomputing System

Richard Kaufmann, Compaq Computer Corporation

On August 22, 2000 the U.S. Department of Energy's (DOE) National Nuclear Security Administration (NNSA) selected Compaq to build the world's fastest and most powerful supercomputer, a 30+ TeraOPS system code-named 'Q'—the latest advancement in the Accelerated Strategic Computing Initiative (ASCI). This presentation reviews the pronounced technical challenges in developing and building a system to meet the ASCI requirements for a 30 TeraOPS code development, execution and operational environment.

1:30PM – 2PM

Objectivity/DB—The High Performance Database Engine

Leon Guzenda, CTO, Objectivity

Objectivity/DB is the leading high performance, distributed, and scalable database engine with unrivaled support for mixed language development and mixed hardware environments. Objectivity's distributed architecture provides fault tolerance and data replication, boosts developer productivity and shortens development time. Objectivity provides the ideal platform for mission critical applications requiring continuous performance and adaptability to future technologies.

2PM – 2:30PM

Fujitsu's High Performance and Highly Scalable Supercomputers and Servers

Kenichi Miura, Fujitsu

The presentation is about the Fujitsu 64-bit architecture VPP5000 supercomputers and the SPARC64-based servers, PRIMEPOWER. It will cover the major features of the VPP5000 such as the vector and scalar units, and the crossbar network. It will also include the performance of various commercially available third party application software on the VPP5000 and how the VPP5000 is used by various customers worldwide.

The presentation will also cover the features of the PRIMEPOWER HPC servers such as Solaris compatibility as well as enhancement features such as the Cross-bar Switch, the Partition Feature, Fault Tolerant Technologies, etc.

2:30PM – 3PM

Concept-to-Production: End to End High Performance Network Monitoring and Performance Analysis Solutions

Erik Plesset, Spirent Communications

Spirent Communications test solutions are enabling the development and deployment of new Internet and communications technologies worldwide. A comprehensive range of tests are used to address functionality, reliability, performance, and conformance to standards etc., of networks and network devices. These abilities are exemplified by their use in the SC2000 Demonstration Network, and the presentation will explain how these lessons can be used in any high performance network management strategy.

3PM – 3:30PM

High Bandwidth Supercomputers of Today and Tomorrow

Dave Kiefer & Burton Smith, Cray Inc.

Cray Inc. designs and builds supercomputers that have high global bandwidth, making them particularly well suited for communications-intensive problems in science and engineering. Cray supercomputers have two essential ingredients: custom latency tolerant processing hardware to enable global bandwidth, and high performance interconnection networks to actually deliver it. Cray will continue to innovate in high bandwidth system architecture, and this talk will describe some new ideas and product directions in both scalar and vector supercomputers.

3:30PM – 4PM

Cost Effective High-Performance Computing on Clusters of Commodity Processors

Robert D. Bjornson, Scientific Computing Associates, Inc.

High-performance computing on cost effective clusters of commodity processors is of growing importance in many industries. Scientific Computing Associates, Inc. (SCA), a pioneer in cluster computing, is partnering with leading hardware and software vendors to help industry leverage this technology. This presentation describes several recent SCA initiatives which make clusters more accessible to the broad user community.

Scientific Computing Associates, Inc. is the foremost commercial developer of cluster software tools, middleware, and applications. Products from Gaussian (the ab initio quantum chemistry program) to Crystal Ball Turbo (an Excel-based parallel Monte Carlo simulation tool) depend on SCA's technology. SCA's focus is on tools that help in the development of cluster enabled applications.

4:00 – 4:30PM

Collaborative visualization solutions for HPC involving Reality Center and immedsive visualization products

Linda Jacobson, SGI

WEDNESDAY | NOVEMBER 8
ROOM D 268

10AM – 10:30AM

The Challenge of Next Generation High Performance Networks

John Freisinger, Essential Communications

John will discuss what Essential has learned in providing gigabit per second technologies to the industry and the obvious hurdles that face the next generation of multi-gigabit networking technologies. The challenges include CPU utilization, buffering to enable multiple speeds within the same network, cable plant limitations and bridging to legacy technologies. Solutions that will be discussed in the talk include TCP/IP offload, interrupt coalescing, OS bypass, and multiprotocal backboning.

10:30AM – 11AM

"Teaming" Technology—STAGGERING PERFORMANCE POWER THROUGH SIMPLICITY

Lynn West, CTO, Times N Systems

TNS has designed an innovative approach to harness the collective computational power of industry standard components providing a scalable, expandable infrastructure to support multiple parallel and non-parallel applications. This technology is especially well-suited for today's fast-paced Internet environment. By unifying—or Teamin—the disk, memory, and processors of individual computers via a unique means of sharing memory we deliver system-wide, the speed and power necessary to alleviate network bottlenecks and system latency problems. The ability to span local disk and memory resources allows the construction of multi-terabyte

hard drives and multi-gigabyte RAM drives.

The TNS technology can be integrated into specific existing compute infrastructures to maximize performance efficiencies and balance scaling of processing power, Disk I/O bandwidth, and memory bandwidth for up to 128 processors.

11AM – 11:30AM

New Capabilities and Plans for High Performance Storage System

Ramin Nosrat (tentative), High Performance Storage System (HPSS)

11:30AM – 12NN

Legion—An Applications Perspective

Andrew Grimshaw, Applied MetaComputing, Inc.

Legion is a reflective metasystem project at the University of Virginia designed to provide users with a transparent interface to resources in a wide-area system, both at the programming interface level as well as at the user level. Legion addresses issues such as parallelism, fault-tolerance, security, autonomy, heterogeneity, resource management and access transparency in a multi-language environment.

While fully supporting existing codes written in MPI and PVM, Legion provides features and services that allow users to take advantage of much larger, more complex resource pools. With Legion, for example, a user can run a computation on a supercomputer at a national center while dynamically visualizing the results on a local machine. As another example, Legion makes it trivial to schedule and run a large parameter space study on several workstation farms simultaneously. Legion permits computational scientists to use cycles wherever they are, allowing larger jobs to run in shorter times through higher degrees of parallelization. Key capabilities include the following:

- Legion eliminates the need to move and install binaries manually on multiple platforms. After Legion schedules a set of tasks over multiple remote machines, it automatically transfers the appropriate binaries to each host. A single job can run on multiple heterogeneous architectures simultaneously; Legion will ensure that the right binaries go to each, and that it only schedules onto architectures for which it has binaries.
- Legion provides a virtual file system that spans all the machines in a Legion system. Input and output files can be seen by all the parts of a computation, even when the computation is split over multiple machines that don't share a common file system. Different users can also use the virtual file system to collaborate, sharing data files and even accessing the same running computations.
- Legion's object-based architecture dramatically simplifies building add-on tools for tasks such as visualization, application steering, load monitoring, and job migration.
- Legion provides optional privacy and integrity of communications for applications distributed over public networks. Multiple users in a Legion system are protected from one another.

12:30PM – 1PM

Deep Blue Update

David Turek, IBM Vice President, Deep Computing, Web Server Division

IBM's Deep Blue technology has spawned a number of advances in supercomputing. This presentation will highlight the emerging areas of research which are being enabled

by this technology such as Life Sciences and will discuss the emerging LINUX phenomenon.

1PM – 1:30PM

High Performance Alpha Clusters

Tom Morris, API

The marriage of cluster computing paradigms with the performance of the Alpha microprocessor is allowing researchers to address problems on a grand scale without investing in expensive specialized systems. Achieving optimal price/performance in computer system design requires balance among the various components of the system as well as a good fit between the application demands and the system design. We'll look at the measured performance characteristics of different elements such as system interconnects, processor speeds, and L2 cache speeds & sizes to show how they impact real world application performance. We'll also look at future developments in processors, packaging, and interconnect to understand how they affect delivered application performance.

1:30PM – 2PM

Linux on the Bleeding Edge, CPlant and FSL

Greg Lindahl, High Performance Technologies, Inc

Linux is used today in many bleeding edge applications, such as Sandia's CPlant cluster (2,000 cpus total), a traditional supercomputer installation with extremely high reliability at the Forecast Systems Lab in Boulder, CO, and commercial clusters at Google (5,000 cpus) and Incyte (3,000 cpus). What makes Linux adequate—or superior—for such systems?

2PM – 2:30PM

Remote Management of a Storage Network

Derek Gamradt, StorNet, Inc.

Many companies today are faced with managing distributed storage infrastructures across multiple locations. However, the lack of availability of storage-centric systems administrators has limited the ability of companies to effectively manage their data storage networks.

This presentation will provide an overview of the challenges involved with managing and monitoring a Storage Management Infrastructure remotely. We will discuss the various elements of a storage management solution, including centralized disk storage, highly available data applications, disaster protection and storage networking. The presentation will compare the various methods that are available to allow a data center to keep up with the exploding amount of data, while centralizing the storage management function. The core approach of this presentation is the importance of defining a storage management availability goal and developing a long range plan for achieving it. StorNet has been designing enterprise storage applications since 1990 and has helped hundreds of organizations deploy and manage mission critical storage management strategies. StorNet is an independent storage solutions provider and can develop a storage management plan without the "one-size-fits-all" approach common to systems manufacturers.

2:30PM – 3PM

The Art of Internet Computing

Dr. Steve Armentrout, CEO, Parabon

The new field of Internet computing harnesses the excess capacity of Internet-connected computers, and makes this massive resource available for scientific research and commercial R&D. In addition to unprecedented power and superior accessibility, this new computational model offers relief from the fixed capacity, rapid obsolescence, and rising maintenance costs of traditional high-performance computing.

Dr. Armentrout, CEO of Parabon Computation, will describe the core technological infrastructure required to create a stable and secure Internet computing platform suitable for general use. He will also discuss the realities of computing with unreliable resources and doing so amidst varied security threats.

Dr. Armentrout will conclude with a discussion of the problem domains and solution strategies that are well suited for Internet computing. Illustrative case studies that include benchmarks will be drawn from the biotech, financial, and film rendering industries.

3PM – 3:30PM

Delivering a New World of Flexibility for HPC

Ben Passarelli, SGI

With the introduction of SGI™ Origin™ 3000 series servers and SGI™ Onyx® 3000 series graphics systems, technical and creative professionals are given the same modularity, freedom of choice, and ease of upgrade that consumers have appreciated and benefited from in such areas as home-entertainment systems. Build-to-suit HPC is now possible with systems from SGI.

NUMAflex™ technology delivers breakthrough capability and flexibility by using modular bricks to add specialized capacities in graphics, central processing, storage, PCI expansion or I/O capacity. Whether SSI or cluster, systems are easily built, modified or upgraded to fit customer requirements and changes over time.

NUMAflex is the suite of benefits provided by SGI's innovative and flexible implementation of NUMA architecture for multiprocessor computers. SGI™ NUMA (nonuniform memory access) exceeds the capabilities of SMP (symmetric multiprocessing) architecture and delivers superior scalability and performance. Leveraging the company's long history and experience in leading-edge computing, SGI is the only computer manufacturer capable of offering this robust third-generation of NUMA, and paving the way for a new world of modular computing.

3:30PM – 4PM

Real-World Performance Analysis

Karl Solchenbach, Pallas

To be useful "in the real world," performance analysis tools must meet stringent requirements: scalability in processes and time, support for shared-memory clusters, ease of use and partially automatic operation. The Pallas tools (Vampir and Vampirtrace) are continually enhanced to meet these requirements, with the latest developments including support for ASCI-scale scalability and for hybrid OpenMP/MPI applications.

4PM – 4:30PM

InfiniBand™ Management Solutions

VIEO (formerly Power Micro Research)

THURSDAY | NOVEMBER 9
ROOM D 268

10AM – 10:30AM

Use of Juniper Routers in Research and Education Networks

John Jamison, Juniper Networks

Juniper Routers are in use in the networks of over 20 high performance research networks, supercomputer centers, universities and research labs. This presentation will describe how these networks and institutions use Juniper Routers to support: traffic engineering, high data rate applications (both unicast and multicast), and security enhancing line rate packet filtering. It will also include a discussion of how future Juniper software upgrades will enable users to implement explicit routing solving the "Fish Problem" that has been the bane of R&E networks since the days of the NSFnet.

10:30AM – 11AM

TurboLinux Cluster Cockpit

Peter H Beckman, TurboLinux

Over the last decade, clusters have grown from handfuls of nodes to thousands. For example, one of the world's most popular Internet search engines is a Linux cluster with thousands of nodes. While the complexity of managing these ever-larger clusters has been steadily increasing, the software tools and practices have not kept pace. Part of the problem is that supercomputers are almost always compared on the basis of peak theoretic megaflops and price, not manageability. TurboLinux has created a new and unique software framework for managing clusters. The framework is based on component technology, with modules that can be loaded or unloaded on the fly, and standardized interfaces for basic module functionality. Furthermore, XML technology is used to describe many components and data. During this session, the component architecture of Cluster Cockpit as well as the issues for managing Linux clusters will be presented.

11AM – 11:30AM

Enabling HPC Research in DoD

Tim Campbell, Logicon

Headquartered in Herndon Virginia, Logicon is a leader in advanced information technologies, systems and services. Logicon has annual sales of nearly $1.5 billion and has approximately 12,000 employees worldwide. Areas of expertise include systems integration, storage management, training and simulation, science and technology, information systems, command and control, and systems and support services.

User needs drive our solutions.

Demonstrations include a 3D visualization of a gas-turbine swirling spray combustion simulation, a hands-on demonstration of an operational distance learning system, and examples of successful systems integration projects and processes.

11:30AM – 12NN

How Clusters Fill a Vendor Void

David Rhoades, High Performance Technologies, Inc.

Most large vendor offerings require programmers to redesign their code to take advantage of the computing architecture. Straightforward MPI is not enough, as the SMP on the 4-, 8, or 16-processor node requires explicit programming techniques to be

fully utilized. This talk will address the various vendor approaches to parallel computing, and explore why many researchers are responding by building their own clustered systems. Here are some issues that we will address. Are single- or dual-node clusters more efficient and cost-effective? What changes have occurred in the clustered systems over the past year, and what changes will we likely see in the next? Will clustered systems exist five years from now? Who should consider a cluster, and who should not?

12:30PM – 1PM

Internet Computing: Distributed Computing Leads Third Generation Internet Applications

Jim Madsen, President and CEO, Entropia

Internet and Intranet Computing is changing the way people think about high power computing. Entropia's solutions complement existing computing resources by breaking compute ceilings with on-demand scalable power, delivering results in a fraction of the time and at a fraction of the cost normally associated with big iron. Behind the firewall, solutions leverage ROI on sunk costs in normal desktops within companies, and massive compute power can be imported to meet peak demand, much like the electric grid distributes power where needed. This truly "organic to the internet" application is leading the way in a third generation of peer-to-peer applications. Entropia, with its distributed computing grid running on the internet for three years now, continues to lead the way in this new space. Entropia will discuss the milestones achieved and challenges ahead, and a future in which internet computing is pervasive.

1PM – 1:30PM

New Compiler Features for High Performance Computing

Doug Miles, PGI

The Portland Group (PGI) is moving into the 21st century with cutting edge methods to assist scientific and technical developers with the latest in compiler technology. PGI provides new solutions to take advantage of the latest x86 processor features—32- and 64-bit Streaming SIMD Extensions, cache prefetch instructions, and large cache sizes. In addition, PGI supports parallel applications via full OpenMP capability as well as automatic parallelization. PGI also provides a solution for cluster computing with the Cluster Development Kit (CDK), which allows users to harness the power of workstations networked together. Applications take advantage of both distributed-memory and SMP parallelism on Linux clusters. PGI has bundled its compilers, tools, and a variety of open-source cluster products into a single turn-key package.

This talk will describe how users can take advantage of these new and exciting features offered by PGI.

1:30PM – 2PM

StorageTek

2PM – 2:30PM

MAN/LAN/WAN—Blurring Boundaries for Switching and Routing

Earl Ferguson, Chief Technical Officer of Foundry Networks

Gigabit Ethernet, wire-speed switching, and next generation routing are fueling the latest generation of network technologies. As well, emerging players promote Layer 4-7 switching as the solution to yet unheard of

Web-based networking problems. Terabit Routing companies offer advice on the removal and replacement of older routers connecting to the WAN. Does increasing the speed of access to the Internet via "next generation" Internet Routers really solve our problems? What should your concerns be? How much will it cost to implement these solutions? What are the pros and cons of these various technologies? Are they separate or should they be combined together? Will they help or hinder network design? Do specific technologies live only in one portion of a network deployment or do the newly emerging product solutions cross traditional networking boundaries? This presentation will discuss the various approaches to "next generation" technologies as vendors rush to combine them together to solve current and future networking problems.

2:30PM – 3PM

Optimization and Parallelization with the Intel Compilers and the KAI Tools

Joe H. Wolf, Intel

We will discuss the advanced optimization techniques employed by the Intel® C++ and Fortran Compilers V5.0 for the IA-32 architecture, including the Intel® Pentium® 4 processor. It also includes the parallelization technology from Kuck and Associates (KAI), an Intel company. Learn how using the Intel compiler's vectorization and parallelization features, including OpenMP*, in conjunction with the KAI tool set, gives maximum application performance in a multiprocessor system. Each of the optimization techniques and tools are illustrated with sample code and on-stage demonstrations.

3PM – 3:30PM

Large-Scale Distributed Computing: The Benefits and Advantages of Compute Farms

Rich Edelman, Director of Software Development, Blackstone Technology Group

With companies that require massive amounts of compute power increasingly turning toward the distributed computing model, compute farms are emerging as a practical alternative to supercomputers. While both compute farms and supercomputers require specific skills in programming and set-up, compute farms offer distinct advantages in both design and functionality.

This presentation will detail the benefits and advantages of the compute farm solution, which include access to unlimited compute power on-demand, and incremental scalability that allows for handling demand in peak-use periods. Compute farms eliminate the need for costly and time-consuming queuing systems, and allow companies to concentrate on their core competencies, rather than spending time managing their mainframe. Also discussed in-depth will be features of the compute farm infrastructure, including service levels, security, monitoring and network performance.

3:30PM – 4PM

SRC Computers Inc: Background, Status, and Plans

Michael J. Henesey, SRC Computers, Inc.

4PM – 4:30PM

Sun Microsystems

SC2000 BIRDS-OF-A-FEATHER SESSIONS

BOFS Chair: Rajeev Thakur, Argonne National Laboratory

Birds-of-a-Feather sessions, or BOFS, are informal get-togethers for conference attendees to discuss topics of mutual interest. Below are descriptions of BOFS scheduled when this program went to press. A BOF notice board with final dates, times, room numbers and any additional sessions will be posted daily in the convention center. This information is also available through the SC2000 website.

TUESDAY | 5:30PM – 7PM

High Speed Interconnects: Status and Peek into the Future

Organizer: Markus Fischer, University of Mannheim

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 263/265

Cluster Computing is one of the most important research platforms for parallel computing in recent years. While standard processors offer high computation performance at low cost, the performance gap to IO is steadily increasing. While Gigabit/s media exist to transfer data, the current bottleneck is the interface between CPU and NIC. This is about to change with InfiniBand™ The presentation of the InfiniBandA™ specification will be the starting point for a fruitful discussion of the future of IO. Other High Speed Interconnects will be presented, giving an overview of their peek into the future. This BOF will be the continuation of the big success of the last BOF at SC99.

Topics of interest:

- InfiniBandA™—the next generation of System IO
- Available High speed interconnects
- Software models such as MPI-1, MPI-2, PVM, VIA, with details on their NIC specific implementation including zero-copy, DMA copy / PIO
- News in MPI-2
- User Level I/O devices

Using Clusters for Visualization

Organizer: Arthurine Breckenridge, Sandia National Laboratories

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 268

As simulation and modeling moves to cluster-based technology, visualization is also being performed on clusters. Come and share the progress everyone has made since the BoF last year. Topics of interest include parallel rendering, panoramic virtual reality, tiled screen displays, cross-platform peripherals, etc.

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 274

OpenMP and its Future Developments
Organizer: Tim Mattson, Intel Corp

OpenMP is the standard application-programming interface (API) for writing parallel programs for shared memory computers. It "was developed" and "is maintained" by the OpenMP Architecture Review Board—or the ARB for short. The goal of the ARB is to promote OpenMP and make sure it evolves to meet the changing needs of the parallel programming community.

This BOF provides a major forum for discussion between members of the ARB and the parallel programming community. We have three goals for the session: (1) to present our latest language specification—OpenMP 2.0 for Fortran, (2) discuss our plans for OpenMP 2.0 for C/C++, and (3) hold an open discussion on the future of OpenMP. Members of the ARB will be at the BOF as will vendors who supply OpenMP products. If you have opinions about OpenMP and how it should evolve, come to this BOF and share them.

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 267

The ASCI Simulation and Computer Science Technology Prospectus
Organizer: Jo Sander, U.S. Dept. of Energy

ASCI's technology prospectus is a set of comprehensive strategic roadmaps (previously known as "ASCI curves") driving ASCI research and development in computational and computer science. These representations depict planned progress and critical barriers that must be overcome for ASCI to meet Stockpile Stewardship objectives. Examples of technology areas include visualization, networking, storage, interconnection networks for ultra-scale computers, and simulation development environments. These roadmaps were developed through a multi-laboratory process where ASCI's experience with applications and platforms efforts led to a clearer understanding of new technology requirements. Furthermore, the introduction of several new ASCI programs, focusing on distance computing, visualization, and data management, required a comprehensive integrated approach to manage interrelated programs. This technology prospectus is an essential part of that integration strategy. This session will provide a forum for presenting and discussing some of these strategic roadmaps that up until now have mostly been developed within the internal ASCI community.

Windows 2000 High Peformance Computing Users BOF
Organizer: David Lifka, Cornell Theory Center

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 271/273

Windows 2000 clusters are rapidly gaining acceptance as a practical way to do high performance computation science on industry standard hardware and software. Already this year NCSA and the Cornell Theory Center have demonstrated that Windows 2000 is scalable and reliable to perform up to the standards of traditionally big iron solutions by making the Top 500 Super Computer List (see http://www.top500.org). A lot of work has gone into developing the necessary tools and programming environments to make the transition from Unix-based systems over to the Windows 2000 environment. In this BOF we will focus on administration experiences and user experiences on Windows 2000 clusters. Topics will include compilers, tools, systems management tools, user experiences, and performance results. Experts from NCSA, CTC, UCSD and others will be on hand to share their experiences. A number of cluster hardware, network and software vendors are also expected to be in attendance.

The ACTS Toolkit Development
Organizer: Leroy Drummond, Lawrence Berkeley National Laboratory

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 270

The ACTS (Advanced Computational Testing and Simulation) Toolkit is a set of tools mainly developed at DOE labs that make it easier for programmers to write high performance scientific applications for parallel computers. The purpose of this BOF is to gather the tool developers and discuss past and current activities under ACTS as well as future services to the scientific computing community. Feedback from other SC2000 attendees is also welcome.

The ACTS tools have been successfully used by many scientists to address their computational needs. The ACTS toolkit actively promotes High Performance Computing and has provided a framework for educating users of the HPC technology, and we are eagerly working to reach out to other areas in computational sciences. The issues to be discussed during the proposed BOF are tool interoperability, tool evaluation, user education, the on-line information center, success stories, feedback from users and future research under the ACTS umbrella.

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 262/264

Parallel Tools Consortium Advancing Tools for Parallel Computing
Organizer: Timothy H. Kaiser, Parallel Tools Consortium

The Parallel Tools Consortium brings together representatives from the federal, industrial, and academic sectors addressing factors that inhibit tool use and usability on parallel computers.

Current parallel tools do not respond to the specific needs of scientific users. They vary widely across current platforms, so the steep learning curve must be repeated each time a user migrates to a new machine.

The lack of specialized support for heterogeneous or scalable applications also deters users from investing the effort needed to parallelize scientific libraries or applications to be used by third parties.

Our mission is to take a leadership role in defining, developing, and promoting tools that meet the specific requirements of users who develop applications on parallel platforms.

We will discuss successful tools that have been developed from efforts of the Parallel Tools Consortium, status of current projects, and explain how users can become involved in continuing efforts.

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 269

Research Experiences for Undergraduates Program
Organizer: Ange Mason, SDSC

The Research Experiences for Undergraduates (REU) program of the NSF has been highly successful for many years. It is important to build a community of REU program students, mentors, and organizers, to strengthen the program and to help develop nationwide interrelationships that will benefit the participants. The EOT-PACI program, which is part of the NSF's Partnerships for Advanced Computational Infrastructure program, proposes to hold a Birds of a Feather session at SC2000 on REU programs to bring together participants and organizers to discuss successes, areas in need of improvement, and plans to build an REU community. Top researchers from the Alliance and NPACI will be present to meet with students and mentors and hear about their research experiences. Additionally, EOT-PACI is inviting those who manage other REU programs to share in the exchange of ideas and information.

Using & Administering GPFS
Organizer: Terry Jones, LLNL

TUESDAY
NOVEMBER 7
5:30PM – 7PM
ROOM D 261

GPFS is IBM's high performance parallel file system. It is notable for it's scalability and performance: file systems of over 10 TBytes with bandwidths of several GBytes/sec to a single file are deployed and in daily production use. This BOF will provide an informal forum for both scientific app writers and system administrators to discuss issues and findings relevant to using GPFS. Information on the newest release (version 1.3 which was released in July 2000) will be presented. Come with your experiences and/or questions!

WEDNESDAY | 5:30PM – 7PM

Cluster Computing and the IEEE Task Force on Cluster Computing
Organizer: Tim Mattson, Intel Corp.

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 271/273

The TFCC is an international forum that promotes cluster computing in research, industry and education. You can learn more about this group at http://www.ieeetfcc.org. In this BOF, we briefly explore the activities of the TFCC. We then move on to a discussion of the key technical issues in modern cluster computing. The format will be informal and consist of panels of experts "on stage" and "in the audience." Topics include programming environments, algorithms and the high performance networks required to build effective clusters.

TOP500 Supercomputers
Organizer: Eric Strohmaier, University of Tennessee

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 263/265

The TOP500 is a project tracking supercomputer installations in the world since 1993. The sixteenth TOP500 list will be published in November 2000 in time for SC2000. Various experts will present detailed analyses of the TOP500 and discuss the changes in the HPC marketplace during the last years. The number of systems installed, the installed performance, the locations of the various supercomputers, the architectures of HPC systems, and also the applications for which these HPC systems are used, will be analyzed. Such statistics provide a better understanding of the high performance market and can facilitate the exchange of data and software. This BoF will also give background information and explanation about the TOP500 project itself and is also meant as an open forum for discussion and feedback between the TOP500 authors and the user community.

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 267

Promoting High Performance Computing Literacy at Low Cost
Organizer: Don Morton, University of Montana

Many organizations have potential interest in HPC, but lack the physical and intellectual infrastructure to enter this realm. Collaborators at the Arctic Region Supercomputing Center, North Carolina Supercomputer Center, University of Texas Health Sciences Center at San Antonio, and University of Montana are launching a promotional program, targeted at researchers in various disciplines who may not have immediate access to HPC resources. This program will emphasize training and development at a grassroots level, on low-cost clusters with a primary goal of helping researchers "get started" in HPC. It is important that these activities take place in a "portable" manner so researchers may seamlessly move their activities from clusters to supercomputers, and vice versa. The organizers of this program request your presence at this initial BOF, so that we can better understand what you, the researchers, need from us in order to accomplish these goals.

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 262/264

Global Grid Forum
Organizer: Charlie Catlett, Argonne National Laboratory

The Grid Forum is a community-initiated forum of individual researchers and practitioners working on distributed computing, or "grid" technologies. Grid Forum focuses on the promotion and development of Grid technologies and applications via the development and documentation of "best practices," implementation guidelines, and standards with an emphasis on rough consensus and running code.

Wide-area distributed computing, or "grid" technologies, provide the foundation to a number of large-scale efforts utilizing the global Internet to build distributed computing and communications infrastructures. As common Grid services and interoperable components emerge, the difficulty in undertaking these large-scale efforts will be greatly reduced and, as importantly, the resulting systems will better support interoperation.

Initially a primarily North American effort, Grid Forum is joining with similar efforts in Europe (www.egrid.org) and Asia-Pacific to form the Global Grid Forum. This BOF will provide an update on Grid Forum activities and objectives.

Web100: Automatically Tuning TCP for High Performance Networks
Organizer: Basil Irwin, NCAR

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 268

Unless the TCP transmit buffer is hand-tuned, TCP sessions across long distance, high performance networks routinely yield extremely poor results for bulk data transfer. The Web100 project seeks to automatically tune the TCP transmit buffer using congestion feedback information available in real time from the TCP stack.

Tuning algorithms will be implemented at the user level, using a newly defined TCP-MIB that continuously monitors the progress of individual TCP sessions via about 3 dozen metrics, and is made available to user processes via an API. Such information may be read by any process, thereby providing a basis for developing TCP diagnostic and monitoring tools. However, processes with root privilege will be able to alter TCP tuning parameters such as each TCP session's TCP transmit buffer.

Portable Batch System (PBS) Birds-of-a-Feather
Organizer: James P. Jones, PBS Products/Veridian Systems

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 274

Veridian invites all current and potential users of the Portable Batch System (PBS) to this Birds-of-a-Feather session. Join the PBS support and development staff for a discussion of OpenPBS(tm) v2.3 and PBS Pro(tm) 5.0 capabilities and features as well as planned product enhancements. Share your ideas and experiences with other PBS users. The Portable Batch System, originally developed for NASA, is the leading workload management system for supercomputers and the de facto standard for linux clusters.

SCICOMP, The IBM SP Scientific Computing User Group
Organizer: Bronis R. de Supinski

WEDNESDAY
NOVEMBER 8
5:30PM – 7PM
ROOM D 272

The IBM SP Scientific Computing User Group, SCICOMP, is an international organization of scientific/technical users of IBM SP systems. The purpose of SCICOMP is to share information on software tools and techniques for developing scientific applications that achieve maximum performance and scalability on SP systems. The SC2000 SCICOMP BOF will review the organization's recently drafted by-laws as well as terchnical concerns raised during the August SCICOMP meeting at SDSC. In addition, SP users and IBM staff will lead technical discussions on topics such as future SP directions and advanced development efforts by IBM's Advanced Computing Technology Center.

SCINET 2000

Vice Chair for Information Architecture:
Bill Kramer, NERSC, LBNL

A significant part of the technology showcased at each year's SC conference is the SCinet network that supports the conference and makes it (for the duration of the show) one of the most intense networks on the planet.

This year, SCinet 2000 is working with Qwest Communications International, Cisco, Nortel Networks, Juniper, Marconi, Foundry, Extreme, Spirent and others to establish flexible, wide area connections to the show floor. Using Qwest's fiber infrastructure in the Dallas area and Qwest SONET, ATM and IP backbones nationwide, the wide area network will feature multiple OC-48 (2.5 gigabits per second) connections, several OC-12 connections and possibly other connections, using the very latest technology and protocols. The total connectivity between SC2000 and the outside world will be 10.5 to 11 gigabits per second—a new record for the SC conferences! In addition to commodity Internet access, WAN connection links to SCinet 2000 will include:

- **ESnet**
- **Abilene**
- **HSCC**
- **vBNS**

SCinet plans to deploy and support IPv4, IPv6, ATM, and Packet over SONET connections, Myrinet, Quality of Service demonstrations, and advanced network monitoring. Other types of connections might be possible based on discussions with requestors.

www.sc2000.org/scinet

SCinet will install and operate more than 40 miles of fiber optics throughout the conference areas, including the types of connections listed below. SCinet 2000 is planning to support an all-ST-terminated, all-fiber show floor network to interconnect booths using switched technologies.

- **100BaseFX, Fast Ethernet**
- **1000BaseSX, Gigabit Ethernet**
- **1000BaseLX, Gigabit Ethernet**
- **OC-3c ATM**
- **OC-12c ATM**
- **OC-48c ATM**
- **1.28 and 2.0 Gigabit Myrinet**

In order to support the complex logistics and requirements of an exhibition on the scale of SC2000, SCinet is deploying four overlapping networks within the Dallas Convention Center. A diagram of the network is shown on page 120. Watch the web pages for updated information. They are all interconnected, but can operate independently of each other. At the lowest level, several days before the show starts, SCinet deploys a commodity Internet network to connect show offices, the conference Education Program, and the show's e-mail facilities. At the next level, there is a production network, provisioned with leading-edge hardware from various vendors. This year, this network will feature Gigabit Ethernet and OC-48 ATM.

The Network Operations Center is also developed from scratch just before the show starts. This year, in addition to the traditional functions of supporting the network equipment, providing a help desk, and work areas for the network engineers, the NOC will also house a variety of displays and information. Spirent is providing their "SmartBits" technology to monitor aspects of

SCinet. The SCinet "bit-o-meter" will display aggregate network traffic. Specific applications and events will be monitored throughout the show. Scinet will also use the "Bro" package from LBNL to monitor network traffic for intrusion. Further displays such as from the Netflow package will also be viewable.

One of the most impressive things about SCinet is that every year, it brings together the best network professionals from across the country to help create the entire network and support the multiple aspects of the program. This year, staff from many organizations are supporting SCinet, including:

Aaronsen Group, Argonne National Laboratory (DOE), Army Research Laboratory (DOD), Avici Systems, Caltech, CISCO Systems, the Dallas Convention Center, the Dallas Convention and Visitor's Bureau, Extreme Networks, Foundry Networks, GST Telecom, Internet2, Juniper Networks, Lawrence Berkeley National Laboratory (DOE), Lawrence Livermore National Laboratory (DOE), Marconi, MCI, MITRE Corporation, National Center Scientific Applications (NSF), Northeast Regional Data Center/University of Florida, Nichols Research/CSC, Nortel Networks, Oak Ridge National Laboratory (DOE), Oregon State University, Pacific Northwest Laboratory (DOE), Qwest Communications, Sandia National Laboratory (DOE), SBC Communications, Spirent, Texas A&M University, U.S. Army Engineer Research and Development Center Major Shared Resource Center (DOD), University Corporation for Advanced Internet Development, University of Tennessee/Knoxville, the Very high performance Backbone Network Services–vBNS (NSF).

These people are the critical ingredients who make SCinet work. Each year they attempt to top the previous year as well as give the best possible service and support the show. They worked year around planning and implementing SCinet, and then will spend more than three weeks in Dallas building and running SCinet. Without the people, all the fiber, all the routers and all the infrastructure would not pass one bit of information.

For the first time SCinet will offer tours of the NOC and other equipment for attendees and exhibitors. A sign-up sheet will be available at the NOC for a limited number of tours. It is a chance to see the equipment from behind the scenes and get to talk to some of the vendors and volunteers who put together this most intense network.

WIRELESS

Working with Cisco Systems Inc., SBC Communications and other vendors, SCinet is creating a large 11 Mbps wireless network on the exhibit floor, in the Education Program area, and other locations throughout the conference space, possibly the entire SC 2000 conference area. This wireless network will support the Education Program and the eSCape2000 activities, among other things.

Wireless connectivity is planned for attendees as well. A standards-based 802.11b network with DHCP service will cover the exhibit floor. Attendees with laptops equipped with standards-compliant wireless Ethernet cards, and an operating system which will configure network services as a client of DHCP should have immediate connectivity. A selection of cards and operating systems known to work are listed on the SCinet web page along with links to vendors, drivers, and instructions. SCinet personnel will not be able to provide direct support to attendees who have trouble connecting.

SCinet will not be providing wireless cards for individual systems. SCinet does not support setup, configuration and/or diagnosis of individual systems, but will provide links to information about these subjects at the web site.

The priority areas supported for wireless are the exhibit areas, education area, convention center lobby, meeting rooms and other spaces. If limits are necessary, we will attempt to indicate range and limits with signage. SCinet discourages people/groups/exhibits from bringing their own base stations, because of issues with base station conflicts. SCinet also reserves the right to disconnect any base station that interferes with the SCinet network.

XNET

Xnet is the leading edge, technology-development showcase segment of SCinet.

Since exhibitors, users and attendees become more and more dependent on SCinet to provide robust, high-performance, production-quality networking, it has become increasingly difficult for SCinet to showcase bleeding edge, potentially fragile technology. Simultaneously, vendors have sometimes been reticent about showcasing bleeding-edge hardware in SCinet as it became a production network.

Xnet provides the solution to this dichotomy. It provides a context which is, by definition, bleeding-edge, pre-standard, and in which fragility goes with the territory. It thus provides vendors an opportunity to showcase network gear or capabilities which typically does not exist in anyone's off the shelf catalog.

This year, Xnet will demonstrate early, pre-production 10-Gigabit Ethernet equipment connecting several show floor booths.

In order to encourage the demonstration of bandwidth-intensive applications on this unique, once-a-year network, SCinet is sponsoring innovative, bandwidth intensive application demonstrations. These are applications that will both stress the capabilities of the SCinet network infrastructure and deliver innovative application value. The most impressive of these applications will be recognized with special awards at the SC2000 Awards Program session on Thursday, Nov. 9. More information about this can be found at the SCinet 2000 web page, along with the latest details about the network.

SC2000 NETWORK CHALLENGE

The SC conference has long been a place where high performance computers and high-speed networks meet. At SC2000, the SCinet team is creating a particularly exciting network and has spent a lot of effort to find real applications that have been waiting for such a high powered network to be feasible. SCinet solicited proposals for the demonstrations of innovative (especially bandwidth-intensive) applications as a way to challenge the community to show that this unique network can be used for exciting applications. This effort has turned into the SC2000 Network Challenge, sponsored in part by Qwest Communications.

The 11 applications selected will both fully utilize (and some claim overwhelm!) the SCinet network infrastructure and deliver innovative application value. A team of judges will review the performance of the applications at the show and make awards during the award session on Thursday. Qwest has pledged award funding for this Challenge as well. The challenge will be the first step in a sustained network-oriented prize that will be featured at SC2000 and future SC conferences.

SC2000 NETWORK CHALLENGE ENTRIES

A DATE MANAGEMENT INFRASTRUCTURE FOR CLIMATE MODELING RESEARCH

A. Chervenak, C. Kesselman, Univ. of Southern California/Information Sciences Institute, I. Foster, S. Tuecke, W. Allcock, V. Nefedova, D. Quesnel, Argonne National Laboratory, B. Drach, D. Williams, Lawrence Livermore National Laboratory, A. Sim, A. Shoshani, Lawrence Berkeley National Laboratory

Description: We will demonstrate our infrastructure for secure, high-performance data transfer and replication for large-scale climate modeling data sets. Climate modeling data sets typically consist of many files, ranging to many gigabytes in size. These files may be replicated at various locations. When a climate modeling researcher requests a particular view of the data, we initiate a secure transfer of the relevant files from the data replica that offers the best performance.

Our application includes several components. First, a user specifies at a high level the characteristics of the desired data (for example, precipitation amounts for a certain time period and region). A metadata infrastructure maps between these high-level attributes and file names. Next, we use a replication management infrastructure to find all physical locations of the desired files. We select among these physical locations by consulting performance and information services such as the Network Weather Service and the Globus Metacomputing Directory Service to predict relative performance of transfers from each location. Once a particular physical replica is selected, we initiate secure, high performance data transfer between the source and destination sites. Finally, the desired data is presented graphically to the user.

This project is joint work by three groups. Researchers at Lawrence Berkeley National Laboratory created a request manager that calls low-level services and selects among replicas. Scientists at Lawrence Livermore National Laboratory provided the user interface and visualization output for the application as well the metadata service that maps between high-level attributes and files. Finally, the Globus project team at USC Information Sciences Institute and Argonne National Laboratory provided basic grid services, including replica management, information services, and secure, efficient data transfer.

PROJECT DATASPACE

R. Grossman, G. Reinhart, E. Creel, M. Mazzucco, S. Connelly, A. Turinsky, H. Sivakumar, J. Jamison, Univ. of Illinois at Chicago, B. Hollebeek, P. Proropapas, Univ. of Pennsylvania, C. Rocke, T. Arons, Univ. of California Davis, Y. Guo, S. Hedvall, Imperial College, London, P. Milne, G. Williams, ACSys, Canberra, M. Cornelson, P. Hallstrom, Magnify, Inc.

Description: Project DataSpace will link 14 sites across five continents to demonstrate a new infrastructure to handle (1) remote data access, analysis, and mining and (2) distributed data analysis and mining. Led by researchers at the University of Illinois at Chicago, the team will demonstrate a variety of tools using its new Data Space Transfer Protocol (DSTP) to publish, access, analyze, correlate and manipulate remote and distributed data. The team hopes that the DSTP infrastructure will provide the same ease of use for distributed data analysis and data mining that HTTP provided for viewing remote documents.

We will showcase DSTP Servers, DSTP Clients, and a variety of DSTP applications. The applications also

use the Predictive Model Markup Language (PMML), an emerging standard for statistical models. For example, in the Sky Survey application, the DSTP Client downloads stellar object catalog data from a DSTP Server, creates a machine learning model based on PMML, and scores large amounts of data at high rates using high performance DSTP applications we have developed. Last year, we were able to move 250 Mbits/sec (~113 GB/hr) from our lab at the University of Illinois (Chicago) to the SC99 showroom floor in Portland with no network tuning. We expect even higher rates at this year's SC2000 using a new release of our software.

The Network Storm application will demonstrate the flexibility of the DSTP Protocol. Locations will be set up on three continents to collect network traffic data. DSTP Servers are already installed at those locations and any person utilizing the DSTP protocol can download data and view the state of the network. Ultimately the group plans to build an infrastructure that will predict network storms and allow for improved network traffic management. We will also demonstrate several other high performance DSTP applications.

INTREPID NETWORK COLLABORATOR (INC)

L. Keely, Numerical Aerospace Simulation Division, NASA

Description: A scientific visualization application is instrumented to provide real-time images, which are multicast onto the SCinet network. The application may reside on a host on the conference floor or on a host at the Numerical Aerospace Simulation Division at Ames Research Center. participants on the conference floor will be able to display these images in network browsers through the use of a "thin" client application. Additionally, participants will be able to steer the scientific visualization application through a TCP connection from the client. A prototype system with compressed images and a frame of about 8 frames/sec puts data on a 100Mbs Ethernet at approximately 15 Megabits/sec.

GIGABYTE PER SECOND FILE TRANSFER IN A CLUSTERED COMPUTING ENVIRONMENT

T. Pratt, J. Naegle, L. Martinez, M. Barnaby, Sandia National Laboratory

Description: Cluster computation and high performance networks have opened the way for network file transfers that compete with the performance of local disk systems. We intend to demonstrate file transfer rates of 1 gigabyte/sec transfer using parallel transfer methods developed by ASCI's DISCOM program.

Terascale scientific simulations create terabytes of data. This data generation capability can quickly swamp any size local filesystem. The result of this is that the computer performance can end up being paced by its ability to transfer data out to other file systems.

RESERVOIR SIMULATION AND HISTORY MATCHING—GRID BASED COMPUTING AND INTERACTIVE DATASET EXPLORATION

J. Saltz, Univ. of Maryland/Johns Hopkins Univ., A. Sussman, Univ. of Maryland, T. Kurc, U. Catalyurik, Univ. of Maryland/Johns Hopkins Univ., M. Wheeler, S. Bryant, M. Peszynska, Univ. of Texas

Description: Modeling of flow and reactive transport arises in many physical applications. For example, environmental quality modeling, (atmospheric, bays and estuaries, groundwater, wetlands), reservoir modeling, and in medical applications such as blood flow.

Reactive transport can involve twenty or more components, thousands of

time steps, hundreds of multiple realizations, and fine grid resolution leading to terabytes of data. Local access is critical in efficient interpretation of this data. On the other hand complicated three-dimensional heterogeneous flow fields can be computationally intense and runtime depends strongly on the availability of optimized libraries.

We will demonstrate the use of MetaChaos software to support the flow and the reactive transport components of the UT Austin subsurface modeling code PARSIM. The flow components will run on Blue Horizon at SDSC and the reactive transport components will run on a local Beowolf cluster at SC2000. Data will be stored in the Active Data Repository (ADR) on the local cluster. As new data is generated and stored in ADR, ADR will support clients that (1) carry out interactive exploration of new datasets, (2) compare features seen in the new dataset with analogous portions of multiple previously computed datasets, and (3) allow the use of this data exploration and visualization capability to carry out computational steering of PARSIM.

TELEMISCROSCOPY OVER IPV6

M. Hadida, National Center for Microscopy and Imaging Research, San Diego Supercomputer Center, A. Durand, Computation Center, Osaka Univ., Y. Kadobayashi, Computation Center, Osaka Univ., T. Hutton, San Diego Supercomputer Center, B. Fink, Lawrence Berkeley National Laboratory, M. Ellisman, National Center for Microscopy and Imaging Research, UCSD

Description: Our telemicroscopy system makes use of advanced networks to enable researchers at remote sites to interactively control a high power electron microscope located at the University of California, San Diego (UCSD).

At SC2000, we envision showing live remote control of the UCSD microscope from the showroom floor. Further, we plan to collaborate in real-time with our colleagues attending the Society for Neuroscience conference in New Orleans, which coincides with SC2000.

BANDWIDTH THIRSTY PARTICLE PHYSICS EVENT COLLECTION ANALYSIS AND VISUALIZATION USING OBJECT DATABASES AND THE GLOBUS GRID MIDDLEWARE

J. Bunn, H. Newman, J. Patton, California Institute of Technology, Caltech, K. Holtman, CERN

Description: We will demonstrate high bandwidth retrieval of particle physics event collections from object databases hosted at Caltech and at CERN. The application we will use is based on Globus middleware and the Objectivity ODBMS, and will rely heavily on optimized network paths between Dallas and both Caltech and CERN.

VISAPULT—USING HIGH-SPEED WANS AND NETWORK DATA CACHES TO ENABLE REMOTE AND DISTRIBUTED VISUALIZATION

W. Bethel, J. Shalf, S. Lau, Visualization Group, Lawrence Berkeley National Laboratory, D. Gunter, J. Lee, B. Tierney, Data Intensive and Distributed Computing Group, Lawrence Berkeley National Laboratory, V. Beckner, Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, J. Brandt, D. Evensky, H. Chen, Networking Security and Research, Sandia National Laboratory, G. Pavel, J. Olsen, B.H. Bodtker, Advanced Communications and Signal Processing Group, Electronics Engineering Technologies Division, Lawrence Livermore National Laboratory

Description: Visapult is a prototype application and framework for remote

visualization of large scientific datasets. We approach the technical challenges of tera-scale visualization with a unique architecture that employs high speed WANs and network data caches for data staging and transmission. High throughput rates are achieved by parallelizing i/o at each stage in the application, and by pipelining the visualization process. Visapult consists of two components; a viewer and a back end. The back end is a parallel application that loads in large scientific data sets using a domain decomposition, and performs software volume rendering on each subdomain, producing an image. The viewer, also a parallel application, implements. Image Based Rendering Assisted Volume Rendering using the imagery produced by the back end.

On the display device, graphics interactivity is effectively decoupled from the latency inherent in network applications. This submission seeks to make two high-water marks: remote visualization of a dataset that exceeds 1 terabyte in size, and application performance that will exceed 1.5 gigabytes per second in sustained network bandwidth consumption.

WORLD WIDE METACOMPUTING

M. Mueller, High Performance Computing Center Stuttgart, S. Sanielevici, Pittsburgh Supercomputing Center, A. Breckenridge, Sandia National Laboratory, S. Sekiguchi, Electrotechnical Laboratory, Japan, J. Brooke, Manchester Computing Centre, F-P. Lin, National Center for High-Performance Computing, Taiwan, T. Hirayama, Japanese Atomic Energy Research Institute, Japan

Description: We will show a broad range of applications from the participating partners, among them are:

• URANUS: The CFD-code URANUS (Upwind Algorithm for Nonequilibrium Flows of the University of Stuttgart) has been developed to simulate the reentry phase of a reusable space vehicle in a wide altitude velocity range. Simulation is the only way to get an idea what exactly happens during the reentry phase of a space vehicle, because measuring is very limited and expensive. Due to the role of chemistry and chemical reactions near the hot surface of the space vehicle, even wind tunnel experiments are not comparable to the flows and reactions appearing during a real reentry.

• P3T-DSMC: This is a Direct Monte Carlo Simulation that was developed to simulate granular matter. Granular materials are ubiquitous in nature, industrial processing and everyday life. Examples range from small scale particles in dust, cement or flour over medium-sized plastic granulates to the planetary rings on the astrophysical scale. Similarly broad are the physical phenomena controlling their behavior in transport, storage and processing.

However, despite their importance, continuum or other large-scale modeling still shows severe deficiencies and the understanding of the mesoscopic physics in these systems, as exemplified by fragmentation, dissipative effects, pattern formation, etc. is incomplete since many theoretical methods otherwise applicable to many-particle systems do not apply.

Often, large-scale computation is the only way to deepen our insight because typical granular systems consist of million of particles and most phenomena are only visible after long time scales.

• Electronic structure simulation: First-principles electronic structure simulation of magnetic alloys, using an order-N algorithm. Running on 1480 T3E-1200 processors, this code was the first to sustain over 1 teraflop and won the 1998 Gordon Bell Prize. See www.psc.edu/science/wang.html for details.

• Jodrell Bank Radio-Telescope De-dispersion code: The search for fast (milli-second) pulsars has important implications for our knowledge of the large scale structure of our galaxy and for fundamental physics. However the pulsar signal is often below the level of background noise in the signal of the radio-telescope. Specialist hardware used to be deployed to enhance the pulsar signal above this background, but this can now be done more effectively using fast Fourier transforms on a supercomputer. We intend to show two applications, both very demanding in terms of network bandwidth.

1. A search of data previously gathered from the radio-telescope using the metacomputer to correct for the effects of signal dispersion by the interstellar medium.
2. Using a known dispersion measure, to perform real-time processing of the signal from the telescope.

The challenge is to adjust the signal processing algorithms to the differing data transfer rates across the networks. The computing challenge is to ensure that the processing algorithms maintain the balance between processing speed and network bandwidth.

GIGABIT/SEC HIGH DEFINITION TV OVER IP

C. Perkins, L. Ghari, A. Mankin, T. Gibbons, USC Information Sciences Institute, D. Richardson, University of Washington, G. Goncher, Tektronix, Inc.

Description: The application driving this demonstration is transport of uncompressed high definition TV signals in SMPTE-292 format over an IP network. This will be the first time transport of studio quality uncompressed HDTV signals over IP has been demonstrated. The uncompressed SMPTE-292M media stream comprises an RTP/UDP/IP flow at approximately 1.5 Gbps. Media encoding and packetization is implemented in special purpose hardware, with the control plane on a standard PC.

More information is available from http://www.east.isi.edu/projects/NMAA/

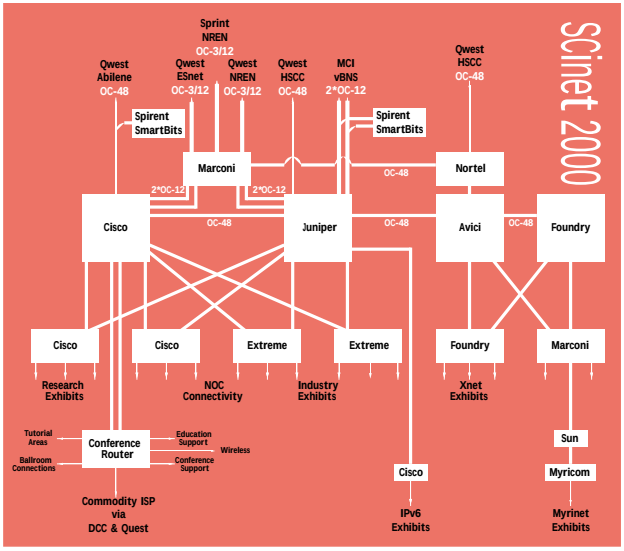

GIGABIT/SEC HIGH DEFINITION TV OVER IPT

Gibbons, E.P. Love, C. Perkins

SCALABLE HIGH-RESOLUTION WIDE AREA COLLABORATION OVER THE ACCESS GRID

B. Olson, T. Disz

SCALABLE HIGH-RESOLUTION WIDE AREA COLLABORATION OVER THE ACCESS GRID

Lisa Childers, Terry Disz, Bob Olson, and Rick Stevens, Futures Laboratory, Argonne National Laboratory

Description: The Access Grid is the ensemble of resources used to support group-to-group human interaction across the grid. It consists of large-format multimedia displays, presentation and interactive software environments, interfaces to grid middleware, and interfaces to remote visualization environments. The Access Grid is designed to support large-scale distributed meetings, collaborative teamwork sessions, seminars, lectures, tutorials and training. The Access Grid design point is small (three to 20 people per site) but promotes group-to-group collaboration and communication. Large-format displays integrated with intelligent or active meeting rooms are a central feature of Access Grid nodes, and are primary new features extending collaboration technology from the desktop. Access Grid Nodes are "designed spaces" that explicitly support the high-end audio and visual technology needed to provide a high-quality compelling and productive user experience.

For this demonstration, select Access Grid nodes will use inexpensive or public domain JPEG encoders and decoders to deliver multiple high resolution, high frame rate video streams in addition to the traditional h.261 video streams used by the Access Grid today. The goal of the demonstration is to show both the future of scalable wide area collaboration and to stress today's networks with the requirement of high bandwidth low latency delivery of information. The demonstration will show multiple live group to group interactions as well as interactions with computer simulations via manipulation within the video windows. A typical Access Grid session may contain several dozens of live or computer-generated video streams.

QOS ENABLED AUDIO TELEPORTATION

C. Chafe, CCRMA Stanford University, S. Shalunov, B. Teitelbaum, Internet2, M. Gröger, Deutsche Telekom, R. Roberts. Stanford Networking, S. Wilson, D. Chisolm, R. Leistikow, G. Scavone. CCRMA

Description: Real-time Internet transmission of CD-quality sound will be demonstrated between the SC2000 floor and Stanford University. Uncompressed live audio streams are possible by employing network enhancements that support minimal delay and low-jitter packet delivery over WAN. Applications include two-way communication (full-bandwidth voice and music) and "surround sound" multi-channel eavesdropping on Stanford spaces. And the Internet itself can be listened to as a vibrating acoustic medium, as if it were a guitar string, with a new technique for generating sound waves on the Internet from real-time echoes (SoundWIRE). This auditory "ping" is used as a tool for evaluation of network constancy. Network quality of service (QoS) for this demonstration consists of marking application traffic for Expedited Forwarding (EF), shaping and policing it at the network edge, and sending it over the Stanford University, CalREN2, and Abilene backbone net-

work, where EF-marked traffic is preferentially serviced. The QoS network design in this demo reflects the architecture of the Internet2 QBone Premium Service. Heavy congestion is created at one or more points near the edge and effective protection of application traffic is demonstrated. For comparison purposes, QoS configuration is dynamically enabled and disabled via the Globus GARA tools and application quality without QoS protection is demonstrated as well.

HIGH-RESOLUTION VISUALIZATION PLAYBACK ON TILED DISPLAYS

M. Papka and R. Stevens, Futures Laboratory, Argonne National Laboratory

Description: High-resolution movies are needed to visualize the time-dependent results of large-scale computer simulations. This is due to the fact that the amount of data generated does not allow for real-time visualization because today's graphics systems are incapable of rendering the output in real-time. High-resolution output is necessary for scientists to gain understanding and insight from the results of their simulations. Therefore, it is common for movies to be made in an offline batch mode for playback and review at a later time. Building from this premise and others, we have constructed systems for playback of movie files on high-resolution tiled displays. These displays are constructed out of numerous individual desktop systems to form an integrated high-resolution display. Most use of these systems today is with local playback of the movie files.

The goal of the demonstration is the playback of a high-resolution tiled visualization movie over the network. A 6-tile mural will be part of Argonne's booth at Supercomputing, the tiled display will be capable of ~3072 x 1536 of pixel resolution. A 24bit-uncompressed movie would require approximately 500 MB/s for smooth animation. Currently several bottlenecks exist in a system of this sort, the goal here is to demonstrate that next generation networking infrastructure, such as what is being deployed at SC2000, will remove networking as a bottleneck.

SC2000 RESEARCH EXHIBITS

Research Exhibits Chair: Virginia To, High Performance Technologies, Inc.

AIR FORCE RESEARCH LABORATORY

Booth Number: R1215

www.afrl.af.mil | hannaj@rl.af.mil

The CCCP comprises n Athlon-processor nodes connected via switched Fast Ethernet configured in a star topology. It utilizes readily available commercial components to assess the capabilities of Beowulf-class supercomputers that are not subject to export controls. Optimal performance will be attained by optimizing compilers and hand-tuned libraries. The CCCP aims to provide a parallel Java design environment and development support for the Common HPC Software Support Initiative. It will be used to support the Joint Battlespace Infosphere development project at the Air Force Research Laboratory. The cluster will also provide a test bed for the evaluation and integration of new and emerging COTS (Commercial Off-The-Shelf) PCI-based multiprocessor add-on hardware including multiple embedded processors and reconfigurable FPGA (Field Programmable Gate Array) architectures. The primary goal for the CCCP is the research, development, and application of advanced computing technologies relevant to the Air Force Information Technology mission. This mission includes high performance computing architectural design and simulation, advanced parallel hyperspectral imaging algorithm design and evaluation, parallel surveillance and tracking algorithm development and characterization, and advanced mixed technology (MEMS, MEFS, electro-optical, etc.) system-level design and simulation.

AMES LABORATORY, SCALABLE COMPUTING LAB

Booth Number: R298

www.scl.ameslab.gov | halstead@ameslab.gov

The Scalable Computing Laboratory in the DOE Ames Laboratory will be showcasing work on assessing and improving communication of real-world parallel applications, on large cluster computer systems. Included in this work is the option of trading latency performance for bandwidth by using data compression. This has particular relevance to cluster computing in which compute cycles are cheap, but intersite communications are limited. In addition, research into improving real-world application performance with lightweight message passing and OS bypass techniques will be presented together with parallel resource management tools.

ARCTIC REGION SUPERCOMPUTING CENTER

Booth Number: R1086

www.arsc.edu | ljevans@arsc.edu

The Arctic Region Supercomputing Center supports the computational needs of researchers within the Department of Defense High Performance Computing Modernization Program, the University of Alaska Fairbanks, other academic institutions and government agencies by providing high performance computing, visualization and networking resources, programming and technical expertise, and training. Areas of specialty supported by ARSC include ocean modeling, atmospheric sciences, climate/global change, space physics, satellite remote sensing, and civil, environmental and petroleum engineering. ARSC collaborates in a number of partnerships, including a joint effort with the U.S. Army Engineer Research and Development Center

and SGI to build and evaluate a 512-processor SGI Origin 3800 single-system image. The Arctic Region Supercomputing Center operates a Cray T3E and a Cray SV-1, with visualization resources including a Pyramid Systems ImmersaDesk and a network of SGI workstations located in a video production/training lab and three additional access labs on campus.

ARGONNE NATIONAL LABORATORY

Booth Number: R186

www.mcs.anl.gov | freitag@mcs.anl.gov

Researchers at Argonne National Laboratory are developing powerful collaborative tools and technologies that will advance the state of the art in large-scale computing and will make scientists more productive. The exhibit showcases work in the following areas: numerical libraries for large-scale computational applications, parallel programming tools, collaborative tools, scalable superclusters, advanced visualization environments, software infrastructure for the national computational grid, servers enabling problem solving over the Internet, and associated scientific computing applications in such areas as computational chemistry and climate modeling. Closely tied with these projects is an emphasis on collaborations, including the ASCI program and the NCSA PACI Alliance.

BOSTON UNIVERSITY

Booth Number: R892

www.scv.bu.edu | glenn@bu.edu

Boston University's research exhibit features its NSF-funded project, MARINER (Mid-level Alliance Resource In the NorthEast Region). MARINER is a partner in the National Computational Science Alliance and extends the university's efforts in advanced scientific computing and networking to organizations throughout the region. 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, posters, and videos in the exhibit booth. We will also be demonstrating distributed computing, collaboration, and visualization software with our Alliance and other partners.

BROOKHAVEN NATIONAL LABORATORY

Booth Number: R388

www.bnl.gov | peskin@bnl.gov

Brookhaven National Laboratory proposes to exhibit new computational science developments in four research areas: 1) The Center for Data Intensive Computing (CDIC), initiated in 1998, is now fully established and pursuing research in high energy and nuclear physics, biological and environmental studies, hydrodynamics, material science and chemistry, and advanced scientific computing environments, 2) The Relativistic Heavy Ion Collider (RHIC), whose experiments require massive computational facilities for data collection and analysis, has been commissioned and has produced its first collisions, 3) The RIKEN QCD-SP machine, winner of a Gordon Bell Prize in 1998, will be succeeded by a follow-on 10 Teraflop supercomputer, whose architecture will be described, and 4) The Brookhaven Data Visualization group will demonstrate and highlight advances in the three areas above, along with other new applications in the physical and life sciences and in distance learning. A portable version of Brookhaven's unique stereographic visualization facility will be in place for demonstrations.

CALIFORNIA INSTITUTE OF TECHNOLOGY, CENTER FOR ADVANCED COMPUTING RESEARCH (CACR)

Booth Number: R686

www.cacr.caltech.edu | chip@cacr.caltech.edu

The Center for Advanced Computing Research (CACR) at the California Institute of Technology was established to foster advances in computational science and engineering. Therefore, CACR focuses on enabling breakthroughs in computational science and engineering by: a) following an applications-driven approach to computational science and engineering research, b) conducting multidisciplinary research on leading-edge computing facilities, c) providing a rich, creative intellectual environment which cultivates multidisciplinary collaborations, and d) harnessing new technologies to create innovative large-scale computing environments. To achieve these goals, CACR conducts multidisciplinary application-driven research in computational science and engineering through collaborations with Caltech, JPL faculty and staff, and R&D establishments throughout the world. Results from these collaborations will be featured in the research exhibit.

CORNELL THEORY CENTER

Booth Number: R1014

www.tc.cornell.edu | cal@tc.cornell.edu

The Cornell Theory Center (CTC) will showcase a number of computational science and engineering applications that are running on our Velocity and Velocity+ high-performance clusters. Among applications featured will be multi-scale materials modeling and LOOPP (Learning, Observing, and Outputting Protein Patterns), software developed and released through CTC's NCRR-funded parallel resource for biomedical scientists. CTC will demonstrate a Domain Monitoring tool developed at the center and the Cornell Multitask Toolbox for MATLAB. SciCentr, a unique informal education site in virtual world space will also be shown. Acquired with the support of members of CTC's Advanced Cluster Computing Consortium (AC3), the Velocity clusters are built from industry-standard components and run the Windows 2000 operating system.

DEPARTMENT OF DEFENSE, HIGH PERFORMANCE COMPUTING MODERNIZATION PROGRAM

Booth Number: R397

www.wes.hpc.mil/index.htm
stinsod@wes.army.mil

The High Performance Computing Modernization Office (HPCMO) will demonstrate its support to the DoD and the warfighter via an interactive poster slide show, hard copy posters, and some limited demonstrations. We will emphasize how the technology employed by the High Performance Computing Modernization Program "trickles down" to help the warfighter. We will show the Computational Technology Areas supported by the program, where these areas are supported and highlight some of the "Challenge Projects" for the past year. We will also describe the Defense Research and Engineering Network, as well as our software initiative known as CHSSI. Information about the Programming Environment and Training (PET) program will also be provided.

DEPARTMENT OF ENERGY, ACCELERATED STRATEGIC COMPUTING INITIATIVE (ASCI)

Booth Number: R1124

www.llnl.gov/asci | chapman@lanl.gov

The U.S. Department of Energy's (DOE) Accelerated Strategic Computing Initiative (ASCI) was created to develop and deploy the supercomputing environments (hardware and software) required to support three-dimensional, full-system simulations of U.S. nuclear weapon systems. In the absence of underground nuclear tests of those systems, numerical simulation will be the key to ensuring the safety and reliability of the U.S. nuclear deterrent. To meet this challenge, DOE scientists will need computers that achieve a speed of at least 100 Teraflop by 2004. The ASCI program has assembled powerful partnerships with DOE's three national security laboratories (Sandia, Los Alamos, and Lawrence Livermore National Laboratories), some of the world's leading computer manufacturers and several of America's great universities. In this booth, the University of Minnesota's InTENsity PowerWall will be used to demonstrate the accomplishments of the programs by focusing on the overarching theme of the ASCI roadmap. This will include the application accomplishments and the ASCI technology development timelines and challenges.

EAST CAROLINA UNIVERSITY

Booth Number: R914

www.ecu.edu/si | creaseyw@mail.ecu.edu

East Carolina University's (ECU) exhibit showcases Embracing Visualization Methods, Progressive HPC Applications, and Innovative NGI Initiatives. In collaboration and partnership with the Ohio Supercomputer Center, state of the art visualization applications developed for the leading edge Reconfigurable Advanced Visualization Environment (RAVE) will be demonstrated on an ImmersaDesk. Presentations include a walk-through of a Space-based Low Earth Orbit Hospitality Facility, Bioactive Scorpion Venom Protein Simulation, a 3D simulation prototype of Blackbeard the Pirate's flagship, research results in quantum chemistry, and Monte Carlo simulations in molecular biology utilizing massively parallel processing techniques. Other demonstrations include quality of service techniques for assessing high bandwidth telemedicine applications which utilize real-time high fidelity pediatric heart sound transmission. As an NCSA affiliate, ECU will host the Education Outreach and Training Partnership for Advanced Computational Infrastructure (EOT-PACI) at our booth to provide information about national education projects and outreach initiatives.

EDINBURGH PARALLEL COMPUTING CENTRE

Booth Number: R597

www.epcc.ed.ac.uk | m.sawyer@epcc.ed.ac.uk

EPCC is one of the leading HPC centers in Europe. It was established ten years ago as a focus for the University of Edinburgh's HPC work, and today has 45 full-time staff members. The center's mission is to promote the uptake of HPC in academia, industry, and commerce. EPCC is involved in providing service, in-depth support and training for UK scientific researchers, and works on technology transfer projects with industry. The center runs a successful visitor program (TRACS) for European researchers, and is involved in international HPC collaborations including the Java Grande Forum and the European Grid Forum. EPCC hosts several HPC platforms, ranging from a 344-processor Cray T3E to an 18-

processor Sun Microsystems E6500 and an in-house built 16-processor Beowulf-class cluster.

ELECTROTECHNICAL LABORATORY

Booth Number: R792

www.etl.go.jp | yoshio@etl.go.jp

Electrotechnical Laboratory (ETL), a Japanese research institute, performs research to produce creative science and technology. Tsukuba Advanced Computing Center (TACC), operated by Agency of Industrial Science and Technology, supports the computational needs of HPC researchers. This exhibit demonstrates the following recent research topics in HPC: 1) Global computing using Ninf: A network enabled server system for global computing and 2) High performance computing with supercomputers including Hitachi SR8000, IBM RS6000/SP, 256-node Alpha Cluster, etc. provided by TACC. TACC plays an important role as a leading computing center for computational grids in the world. The basic concept and status of a global computing infrastructure around the Asia-Pacific area called Asia-Pacific Grid (ApGrid) will be introduced in this exhibit.

EUROPEAN CENTER FOR PARALLELISM OF BARCELONA (CEPBA-UPC)

Booth Number: R1214

www.cepba.upc.es | torres@cepba.upc.es

The booth will present the developments and results achieved by CEPBA in research and development projects during past years. The main project will be Paraver, a visualization and analysis tool for MPI, OpenMP, and Java programs. Other projects in the booth will be Nanos (cooperation between OpenMP compiler and OS scheduling on multiprogrammed multiprocessors) and Dimemas (a simulator of DIstributed MEmory MAchines that is being successfully used in tuning MPI applications). We intend to show that a careful design of different tools enables their integrated use, supporting methodologies and practices that lead to very high productivity of parallelization activity. The people interested in these topics who come to the booth will be able to see demonstrations of the different projects and will obtain explanations about CEPBA developments and activities.

FERMI NATIONAL ACCELERATOR LABORATORY AND STANFORD LINEAR ACCELERATOR CENTER

Booth Number: R500

www.fnal.gov | www.slac.stanford.edu

ecs@lightlink.com

The world's physicists probe the laws of nature at Fermilab and SLAC, the nation's principal experimental high-energy physics facilities. Graphical displays representative of the billions of collisions measured by current experiments will introduce strategies for petabyte-scale data storage and access in this joint exhibit. SLAC will demonstrate its 150-terabyte-and-growing Objectivity DB database used by hundreds of physicists. Fermilab will demonstrate systems that will manage and analyze even larger volumes. Progress on the Particle Physics Data Grid's High-Speed File Replication Service will be shown. A demonstration of PingER will show worldwide network monitoring and data from six years. Physics analysis techniques and results will be described using the BaBar and CMS experiments as examples. The Sloan Digital Sky Survey will illustrate growing collaboration with other practitioners of data-dominated science. Fermilab (operated by URA Inc.) and SLAC (operated by Stanford

University) are funded by the US Department of Energy.

HIGH PERFORMANCE COMPUTING CENTER STUTTGART (HLRS)

Booth Number: R1142
www.hlrs.de | mueller@hlrs.de

The High Performance Computing Center Stuttgart (HLRS) is a national HPC center in Germany for research. In addition, together with debis Systemhaus GmbH and Porsche, it has formed a joint company to provide access to supercomputers for research and industry. At SC2000 HLRS will demonstrate its activities in the field of Grid Computing or Metacomputing. Our presentation will show the main building blocks HLRS is working with. Several projects highlight how these blocks are put together. This includes computers, networks, middleware and applications. To show the potential of the resulting framework HLRS will integrate computing resources at Tsukuba Advanced Computing Center (TACC)/ Japan, Japan Atomic Energy Research Institute (JAERI)/Japan, Sandia National Labs (SNL)/USA, Pittsburgh Supercomputing Center (PSC)/USA, Manchester Computing Center (MCC)/UK, the National Center for High Performance Computing (NCHC)/Taiwan, and HLRS/Germany. Several applications in different fields will use this framework with results being visualized by our own collaborative visualization tool COVISE.

INDIANA UNIVERSITY, PURDUE UNIVERSITY, AND THE UNIVERSITY OF NOTRE DAME

Booth Number: R701
www.indiana.edu/~rindiana
dhart@indiana.edu

Research@Indiana is a collaboration among Indiana's three R1 universities—Indiana University, Purdue University, and the University of Notre Dame—to showcase research and development accomplishments in high performance computing, high performance networking and visualization. The display will showcase computer science developments, including cluster technology, collaboration, grid technology, and massive data storage systems, as well as applications in areas such as astronomy, bioinformatics, chemistry, engineering, medicine, and physics. Indiana has become increasingly important as a center of information technology research, development, and commerce. Home to the Abilene and TransPAC NOCs, Indiana's universities are consistently represented in the Top500 list, and Indiana's computer scientists are developing important new software technology. Much as the research activities of Indiana's three R1 universities cover a great diversity of disciplines, so do accomplishments of Indiana-based researchers making use of HPCC applications.

INTERNET2

Booth Number: R800
www.internet2.edu | ekl@internet2.edu

Internet2 is a consortium of over 170 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies, accelerating the creation of tomorrow's Internet. Internet2 is recreating the partnership among academia, industry, and government that fostered today's Internet in its infancy. The primary goals of Internet2 are to:

- Create a leading edge network capability for the national research community
- Enable revolutionary Internet applications
- Ensure rapid transfer of new network services and applications to the broader Internet community.

JAPAN ATOMIC ENERGY RESEARCH INSTITUTE, CCSE

Booth Number: R1248

www.guide.tokai.jaeri.go.jp/ccse/
hirayamt@koma.jaeri.go.jp

CCSE of the Japan Atomic Energy Research Institute was established in April 1995 by governmental guidance to promote computational science and engineering among the national and other semi-governmental research organizations. CCSE has the largest, best equipped complex parallel computer system consisting of nine sets of parallel computers of different machine architecture, five of which are vector processors and four are scalar ones. The computers are located at four sites in Japan. With a distinctive feature of complex computer systems, CCSE has put special emphasis on the R&D activities of parallel computing technologies and developed a common technological basis of parallel processing such as: a) distributed parallel computing environment, b) parallel computation algorithms, c) parallel processing tools, and, as a specific application of these technologies, d) the development of techniques for numerical experiments. We report on our activities regarding these issues.

JAPAN MARINE SCIENCE AND TECHNOLOGY CENTER (JAMSTEC)

Booth Number: R1108

www.jamstec.go.jp | otsukak@jamstec.go.jp

Japan Marine Science and Technology Center (JAMSTEC), a general oceanographic research institution, was established in October 1971. JAMSTEC installed a NEC SX-4 in spring 1996 and an SX-5 in November 1999 for studying global change. These supercomputer systems are indispensable to understand and predict phenomena such as El Niño events, global warming, weather disasters, and tectonic structure around plate boundaries. In 1997, the Earth Simulator project was started as the cooperative project among JAMSTEC, JAERI, and NASDA under the direction of the Science and Technology Agency (STA) of Japan. The Earth Simulator is a distributed memory parallel supercomputer composed of 640 processor nodes, and each node consists of eight vector processors. The total peak performance and main memory capacity are 40 teraflops and 10 terabytes, respectively. The Earth Simulator will be in operation in the first quarter of 2002 in Yokohama. The Earth Simulator is expected to implement a coupled atmosphere-ocean general circulation model with high resolution, which is being developed by the Frontier Research System for Global Change (FRSGC).

JAPAN SCIENCE AND TECHNOLOGY CORPORATION (JST)

Booth Number: R1109

www.jst.go.jp | minmin@tokyo.jst.go.jp

Japan Science and Technology Corporation (JST), a semi-governmental organization, reports to the Science and Technology Agency. Since 1996, it has been operating the JST Super Computer Complex (SCC). The system is being used for two projects; one is "Electronic Structures Data System" in the material science field and the other is "Gene Finding Tools" in the life science field. In 1998 JST started "Research and Development for Applying Advanced Computational Science and Technology (ACT-JST)," a grant program that promotes research and development using computational science and technology and networks. There are 25 projects ongoing under the program and

some of them are running tools on SCC. In the exhibit, some systems that are being developed on ACT-JST program and SCC are shown.

JOHN VON NEUMANN INSTITUTE FOR COMPUTING

Booth Number: R1142

www.fz-juelich.de/nic | n.attig@fz-juelich.de

The John von Neumann Institute for Computing (NIC), mainly carried out by the Central Institute for Mathematics (ZAM), is one of three national HPC Centers in Germany. Its task is to support and further develop scientific computing in Germany in cooperation with other centers, universities, and research institutes by providing supercomputer resources nationwide, developing computational methods and conducting interdisciplinary research. Using a software system developed within the government-funded UNICORE Plus project, we will showcase the capabilities of uniform access to different supercomputers in Germany. R&D work on SMP-cluster systems, like porting applications, performance analysis, programming models and tools for hierarchical systems, and especially the interaction of shared memory parallelism and distributed memory parallelism, will also be presented. We will also demonstrate recent activities in the performance analysis of parallel programs, the design of parallel algorithms, and the simulation and visualization of complex applications.

KRELL INSTITUTE

Booth Number: R1136

www.krellinst.org | kilmer@krellinst.org

The Krell Institute booth will showcase research conducted in two fellowships administered by Krell, and will demonstrate computer-based computational science training for K-12 teachers using materials from the Adventures in Supercomputing pilot program. The fellowships on which the Krell booth will focus are first, the Department of Energy's Computational Science Graduate Fellowship (CSGF) and second, the computer science fellowship sponsored by Los Alamos National Laboratory and Lawrence Livermore National Laboratory. The CSGF fellows carry out research in a wide variety of resource intensive computational science areas including turbulent combustion, aeroelastic system simulation, and adaptive algorithms used for sensory perception, planning, and control. Computer science fellows concentrate their research in the high performance computing areas of scalable operating/run-time systems, hierarchical program systems, compiler design, networking research, performance modeling, and component architectures.

LAWRENCE BERKELEY NATIONAL LABORATORY

Booth Number: R302

www.lbl.gov | rlboucher@lbl.gov

Lawrence Berkeley National Laboratory (LBNL), home to the Department of Energy's National Energy Research Scientific Computing Center (NERSC) and Energy Sciences Network (ESnet), has long been a leader in computing and networking research. Today, the Lab's HPC and networking capabilities and facilities are helping transform DOE's research programs. Berkeley Lab will highlight its expertise and achievements in computational science—NERSC users will present results achieved using CRAY T3E and IBM SP supercomputers.

• Scientific applications: LBNL, a leader in Adaptive Mesh Refinement, will release the latest AMR software and visualization packages and will demonstrate AMR's capabilities. DeepView, a scalable system for distributed microscopy and informatics, will be shown.
• Networking: In addition to spotlighting ESnet's capabilities, LBNL will also demonstrate its BRO reactive firewall application, Quality of Service efforts, and the VisaPult visualization framework for distributed volume rendering using a workstation with graphics hardware.
• Grid technologies: Demonstrations will show the availability of computing resources and software for accessing remote datasets.

LOS ALAMOS NATIONAL LABORATORY

Booth Number: R298

www.lanl.gov | chapman@acl.lanl.gov

Scientific and technological research and development remain core to Los Alamos. Ongoing research ranges from innovative biological research to predictive modeling of the global climate, wildfire, and transportation systems, and from developing the scientific visualization tools of tomorrow to implementing one of the most powerful computers installed in the world. Los Alamos software teams are responding to the challenges of taming software complexity by building an integrated software infrastructure for scientific simulation development. Demos, posters, and an interactive video presentation will be shown that demonstrate these capabilities.

MAUI HIGH PERFORMANCE COMPUTING CENTER AND THE ALBUQUERQUE HIGH PERFORMANCE COMPUTING CENTER

Booth Number: R196

www.mhpcc.edu | www.arc.unm.edu
cshirley@mhpcc.edu

The Maui High Performance Computing Center (MHPCC) and the Albuquerque High Performance Computing Center (AHPCC) are national supercomputing centers managed by the University of New Mexico (UNM). Established under a Cooperative Agreement with the Air Force Research Laboratory (AFRL), MHPCC is a leader in scalable computing technologies and uniquely chartered to support the Department of Defense (DoD), government, commercial, and academic communities. AHPCC provides an environment for research and education in advanced high performance computing, interdisciplinary applications, and state-of-the-art communications. MHPCC is a Distributed Center of the DoD High Performance Computing Modernization Program (HPCMP), and both MHPCC and AHPCC are SuperNodes of the National Computational Science Alliance (Alliance). Projects featured at SC2000 include: advanced image enhancement research, new material design, development of Linux clusters and collaborative environments for performance analysis of scientific and engineering applications, and industrial design projects with private industry.

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

Booth Number: R695

www.nasa.gov | pelson@mail.arc.nasa.gov

NASA's research exhibit demonstrates how NASA meets its goals using high-performance computing and net-

working with projects from five field installations. A variety of real-time and interactive demonstrations feature the latest research in:

- Computational applications serving NASA's aerospace, earth science, and space science missions
- Remote collaboration and use of virtual reality
- Software tools for developing, debugging, converting, monitoring, and optimizing code in grid environments
- Learning technologies, and
- High-end networking.

A large collection of workstations, interactive theaters, and virtual reality devices are used to display the research and encourage visitor interaction.

NATIONAL AEROSPACE LABORATORY OF JAPAN AND STA HPC GROUP

Booth Number: R1208

www.nal.go.jp | nahirose@nal.go.jp

National Aerospace Laboratory (NAL) is a government research institute under Science and Technology Agency (STA) of Japan. NAL has been leading in HPC among various STA organizations which organize the STA-JAPAN HPC Group. In SC2000, NAL joins its efforts with sister organizations to present STA-HPC Group activities. NAL developed and has been operating the Numerical Wind Tunnel (NWT) since 1993. NWT is the main computing engine of the NAL Numerical Simulator (NS) system. Its contribution to Japanese aerospace research and development projects such as NEXST: supersonic civil transport project, and HOPE: unmanned space shuttle between the Space Station and Japan, is significant. NAL also promotes fundamental research in fluid dynamics and computational sciences. In the exhibit, the CFD Research Center and Multidisciplinary Simulation Concept on the next generation NWT will be shown. Present status of WANS: web access to NS system, and UPACS: unified CFD software package, will also be demonstrated connecting NWT and the booth.

NATIONAL CENTER FOR ATMOSPHERIC RESEARCH (NCAR), SCIENTIFIC COMPUTING DIVISION

Booth Number: R812

www.scd.ucar.edu | don@ucar.edu

NCAR's Scientific Computing Division presents The Virtual Earth System, a science and technology exhibit focused upon our large-scale simulation efforts in a variety of societally important research domains. This presentation uses a combination of virtual reality, digital media, and high-bandwidth networking to explore and display results from a broad spectrum of research efforts including: climate and global change, high-resolution severe weather, wildfires, atmospheric chemistry, terrestrial and solar turbulence, and more. The challenges and demands of these simulation efforts and the subsequent visualization of results provide insight into the importance and difficulty of developing a better understanding of our planet's processes. Progress in the Earth System Grid project, a joint effort with DOE in the area of wide-area large-scale data analysis, will also be demonstrated.

NATIONAL CENTER FOR HIGH-PERFORMANCE COMPUTING (NCHC), TAIWAN

Booth Number: R806

www.nchc.gov.tw | fplin@nchc.gov.tw

The National Center for High-performance Computing (NCHC) is one of the national laboratories under the

National Science Council (NSC) in Taiwan. It is the only research center for high-performance computing applications in Taiwan. Recently, the center was also made the center for the next generation research network of Taiwan. NCHC has conducted various research applications regarding high-performance computing and networking. In the research exhibition, we will use immersive and collaborative virtual reality to showcase the following programs: 1) Plastic extrusion modeling and mold design using numerical wind tunnel, 2) A virtual GIS-based 3D hydro-dynamic model of the Tamshui river, 3) The crashworthiness of the newly developed Yulong vehicle during a frontal impact, and 4) The structure-based drug design model of a transmembrane endothelin receptor and its antagonist. Moreover, we will participate in "Bandwidth-Intensive Network Applications at SC2000" collaboratively with five supercomputing centers from Germany, US, Japan, and the UK to showcase global meta-computing applications.

NATIONAL COMPUTATIONAL SCIENCE ALLIANCE

Booth Number: R804

www.alliance.ncsa.uiuc.edu
kareng@ncsa.uiuc.edu

The National Computational Science Alliance (Alliance) is a partnership of more than 50 institutions working to prototype the next century's advanced computational and information infrastructure. The National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign anchors the Alliance, which is funded by the NSF's Partnerships for Advanced Computational Infrastructure (PACI) program. At SC2000, Alliance teams will demonstrate how the Alliance attacks large-scale problems of science and engineering and drives technology development. They will also show how scientists are using the emerging PACI Grid, a prototype national technology grid, to conduct long-distance collaborative research. The use of cluster technologies in scientific research—specifically clusters of processors running on Intel's new Itanium 64 chip—will also be highlighted. Many applications have been ported to the cluster platform, and performance often matches or exceeds performance on more traditional supercomputing platforms.

NATIONAL COORDINATION OFFICE FOR COMPUTING, INFORMATION AND COMMUNICATIONS (NCO)

Booth Number: R906

www.ccic.gov | vandamme@ccic.gov

The National Coordination Office (NCO) for Computing, Information, and Communications exhibit will feature demonstrations and displays about Federal information technology R&D. This years exhibit will highlight SuperNet applications—SuperNet is DARPA's gigabit speed Next Generation Internet (NGI) testbed. These multi-agency applications will include:

- Digital Amphitheater: A demonstration of the protocols and toolsets for multi-site (up to eight sites) collaborations.
- Matisse: A scientific microscope providing reliable tools (including ultra-high resolution motor control and stroboscopic LED illumination) for in-situ visualization of the motions of internal structures in MEMS.
- Digital Earth: Open, distributed, scalable, multi-resolution 3D representations of the Earth into which massive quantities of geo-referenced

information can be embedded.

- Land-Speed Record: A demonstration of multiple HDTV streams, designed to break the "land speed record," for maximum data flow rates in the range of 1 GB/second.

Additional information will be available about multi-agency Federal IT R&D efforts, the president's Information Technology Advisory Committee, and the role of the NCO.

NATIONAL PARTNERSHIP FOR ADVANCED COMPUTATIONAL INFRASTRUCTURE (NPACI)

Booth Number: R904

www.npaci.edu | mgannis@sdsc.edu

NPACI is a NSF-supported consortium of four dozen premier academic, industrial, and research institutions, led by SDSC at UC San Diego. Its mission is to advance science by creating a national computational infrastructure through

- Capability computing: Providing compute and information resources of exceptional capability to enable scientific discovery at scales not previously achievable
- Discovery environments: Developing and deploying integrated, easy-to-use computational environments to foster scientific discovery in traditional and emerging disciplines
- Computational literacy: Extending the excitement, benefits, and opportunities of science to a diverse population.

NPACI's exhibit will showcase "Alpha projects" in bioinformatics infrastructure, protein folding, telescience, multicomponent environmental modeling, scalable visualization, biological fluid dynamics, and cellular microphysiology. We will demonstrate new tools and applications being developed by the cooperating partners, present NPACI resources in action through "transparent supercomputing," and show how the partnership's activities, products, and services are meeting real needs of the computational science community.

NORTHWESTERN UNIVERSITY

Booth Number: R598

www.ece.nwu.edu/EXTERNAL/vtaylorlab/
zlan@ece.nwu.edu

Efficient execution of applications requires insight into how the system features impact the performance of the application. For distributed systems, the task of gaining this insight is complicated by the complexity of the system features. This insight generally results from significant experimental analysis and possibly the development of performance models. This research exhibit will present the Prophecy Project, an infrastructure that aids in gaining this needed insight based upon one's experience and that of others. Prophecy consists of three major components: profiling and instrumentation, relational databases and data analysis.

OAK RIDGE NATIONAL LABORATORY

Booth Number: R587

www.csm.ornl.gov | rileyba@ornl.gov

ORNL highlights scientific discoveries in climate, genomics, materials, molecular biology, and nuclear medicine made possible by advances in mathematical methods and high performance computing. Learn how performance evaluations of early systems are used to develop specialized techniques to optimize applications for terascale systems. See how M3C and C3 tools and Harness/PVM help build fault-tolerant clusters (NT, Linux, or NOW) that can be administered via a web browser. Check out plug-ins for going unplugged and the CCA

(Common Component Architecture)—the next best thing to cut and paste for developing large-scale multi-disciplinary simulations. Visit the ORNL booth and explore:

- Will tomatoes grow in Tennessee in 2099?
- Can I have a cool car in California?
- Can I take that shipment through Nevada?
- How many ways can a virus vibrate?
- Why do I have this craving for cheese?
- How can I scale my cluster administrator (without leaving scars)?
- How fast is fast enough?

OHIO SUPERCOMPUTER CENTER (OSC)

Booth Number: R1006

www.osc.edu | kkelley@osc.edu

Over the years, OSC has made a place for itself in the world of High Performance Networking and Computing. This year, we plan to showcase how the work we have done has made a difference. Our theme this year is "Expanding the HPC Frontier." One important focus of OSC is our development of a working prototype system for the virtual simulation of temporal bone dissection. This scientific visualization has the potential to strongly impact medical procedures and instruction. Another focus is our work with Internet2 technology. ITEC-Ohio led by OARnet, a division of OSC, and a consortium of Ohio universities and corporate partners, is one of two national testbeds for Internet2 testing and evaluation of emerging Internet technologies. Finally, OSC's collaborations with research institutions such as East Carolina University have opened doors to expanded technological applications. For example, OSC participated in an Internet-supported military exercise where medical doctors from Columbus provided diagnostic and treatment advice for U.S. military personnel injured in battlefields halfway around the globe.

OSAKA UNIVERSITY, CYBERMEDIA CENTER

Booth Number: R900

www.rd.center.osaka-u.ac.jp
date@rd.center.osaka-u.ac.jp

Analysis of brain data is too time-consuming, in spite of recent processor development, because the analysis is often performed on a single processor workstation or personal computer. Our goal is to seamlessly acquire, analyze, and visualize brain data, which are essential to understanding brain function, on the Internet. For this purpose, we have been developing MEG (magnetoencephalography) data analysis system using emerging grid technologies. MEG is a promising brain functional imaging modality. We will show that our system can dramatically reduce analysis time and has the potential to improve the current situation of brain science. In detail, MEG data analysis is performed using remote high-performance computers. After that, results of analysis are immediately visualized.

PACIFIC NORTHWEST NATIONAL LABORATORY

Booth Number: R998

www.pnl.gov/cse | nlee.prince@pnl.gov

Pacific Northwest combines computational, mathematical and engineering sciences and a performance computing technology to solve key scientific problems in:

- Applied mathematics—complex, multi-spatial-scale problems
- Atmospheric science—local, regional, and global models
- Computational biology—modeling

micro and macro-scale systems with large data sets
- Chemistry—management and remediation of legacy nuclear wastes
- Computer science—visualization, problem-solving environments, and artificial intelligence products
- Engineering—modeling effects of stress, heat, and impact on materials
- Reactive chemical transport—coupling atomic scale with continuum models

Pacific Northwest scientists are developing modeling and simulation applications in these domains. Together with high performance tools and libraries, these applications run on massively parallel computer systems to efficiently manage and analyze very large data sets. Pacific Northwest will exhibit ongoing research in these domains.

PITTSBURGH SUPERCOMPUTING CENTER (PSC)

Booth Number: R508

www.psc.edu | hackworth@psc.edu

The Pittsburgh Supercomputing Center (PSC) is a NSF national terascale supercomputing center also receiving funding from the Department of Energy, the National Institutes of Health and the State of Pennsylvania. PSC is dedicated to providing government, academic and industrial users with access to state-of-the-art high performance computing and communication resources. Our educational mission features an internship program that provides participants with real experience in a high technology environment. Above all, we strive to provide a flexible environment conducive to solving today's largest and most challenging computational science problems. Our research exhibit will demonstrate the capabilities of our resources, which include a Cray T3E/LC512 and other HPC platforms. We will feature a variety of demonstrations designed to showcase PSC research. Particular areas of focus include materials science, high energy physics, weather modeling, and computational biomedical research such as bioinformatics, structural biology, computational pathology and functional MRI.

PURDUE UNIVERSITY

Service Providers

Booth Number: R198

www.purdue.edu | kapadia@purdue.edu

The booth will present and demonstrate computing portal and wide-area network-computing technologies that allow seamless management of high-performance applications, data, and machines distributed across wide-area networks. These technologies allow users to access and run applications (even unmodified commercial applications) from anywhere via standard Web browsers. The exhibit will highlight the manner in which state-of-the-art technologies can be used to build end-to-end solutions for high performance, ubiquitous supercomputing on Intranets and the Internet. The demonstration will be based on PUNCH, a network computing system that has been operational for five years. PUNCH currently provides computing services to about 800 users across ten countries; 50 engineering software packages developed by 13 universities and six vendors are available. The exhibit will include presentations on several multi-university projects for collaborative engineering and distance education that utilize PUNCH as an enabling technology.

REAL WORLD COMPUTING PARTNERSHIP (RWCP)

Booth Number: R403

www.rwcp.or.jp | ishikawa@rwcp.or.jp

Real World Computing Partnership (RWCP), funded by the Japanese government, will show: 1) network architectures and 2) system software and applications on seamless parallel and distributed computing environments. Two network architectures will be presented: i) Comet, the clustering technology using IP over OC48c POS (2.5 Gbps) for information grids and ii) RHiNET, a local area system network for high performance parallel computing. The following system software and applications will be shown: i) Score cluster system software, ii) cluster-enabled Omni OpenMP compiler for SMPs and PC clusters, iii) PROMISE programming environment for regular and irregular scientific applications, iv) SPST Programming Tool for heterogeneous parallel and distributed systems, and v) parallel data mining system. To demonstrate the system software and applications developed at RWCP, RWC PC Cluster III, consisting of 128 PC's connected by Myrinet, three Fast Ethernets, and Gigabit Ethernet, will be brought to SC2000. You will realize the impact of the network bandwidth on applications.

RESEARCH ORGANIZATION FOR INFORMATION SCIENCE & TECHNOLOGY (RIST)

Booth Number: R602

www.tokyo.rist.or.jp/ | wada@tokyo.rist.or.jp

Research Organization for Information Science & Technology (RIST) was established in 1995 under the umbrella of STA (Science and Technology Agency). Since then, RIST, located in the center of Tokyo, in accordance with STA's guidance, has been making endeavors to advance the frontier of computational science and technology. One of its major missions is to support integrated computational environments focusing on earth science and its related areas. Exhibitions are mainly on GeoFEM (parallel FE solid earth simulation code) and Foo-Jing (a framework for the next generation atmospheric model).

SAITAMA UNIVERSITY

Booth Number: R992

www.cent.saitama-u.ac.jp
ido@poti.fms.saitama-u.ac.jp

Computational studies using high performance computers and simulation codes are applied in the fields of engineering and material sciences. Three-dimensional visualization is developed for research and education. Educational systems are developed in these visualization environments for undergraduate students and even younger students in the middle school in Saitama.

STANDARD PERFORMANCE EVALUATION CORPORATION (SPEC)

Booth Number: R1000

www.spec.org/hpg
eigenman@ecn.purdue.edu

The booth will present benchmarking activities of the SPEC High-Performance Group. The exhibit pursues two goals. First, it will present SPEC's high-performance computing benchmarks, SPEChpc, and the new SPEComp suite. These benchmarks are a service to the HPC community, where they can be used for machine procurement, to improve existing computer systems and for research on software and hardware components of HPC systems. Second, the booth will present several research

projects that are closely related to the SPEC effort. These projects define performance evaluation methodologies, characterize computational applications and evaluate candidate benchmarks. We will present a number of such efforts from several participating organizations. One particular highlight of this year's exhibit will be the new SPEComp benchmark suite. SPEComp provides new benchmarks written in the emerging parallel programming standard OpenMP, which is now supported on all major HPC platforms.

STANFORD UNIVERSITY, CENTER FOR COMPUTER RESEARCH IN MUSIC AND ACOUSTICS

Booth Number: R400

www.ccrma.stanford.edu
cc@ccrma.stanford.edu

A research project for evaluating quality of service (QoS) using "Sound Waves on the Internet from Real-time Echoes." New, no compromise, TCP-based computer applications for audio will be demonstrated using a simplified approach for high quality music and sound streaming over IP networks. Previously existing systems for streaming digital audio involve a number of trade-offs. Because of transmission bandwidth limitations and best effort delivery, audio signal compression of one form or another is typical. Buffering of data, which often delays a signal by seconds, safeguards against delivery uncertainties. Audio is an unforgiving test of networking—one data packet arrives too late and we hear it. Compromises of signal quality and interactivity have been necessary to avoid this basic fact. Along with our new professional audio applications we have developed SoundWIRE, a utility which affords an intuitive way of evaluating transaction delay and delay constancy. Its final form is an enhanced "ping" that uses actual sound reflection. A musical tone, such as a guitar pluck, can be created by repeatedly reflecting a digital acoustic signal between two hosts. Using the network delay between these reflections to substitute for a guitar string creates a tone whose stability represents perfectly regular service and whose pitch represents transmission latency. The ear's ability to discern minute differences makes this an unforgiving test of network reliability.

THE AGGREGATE

Booth Number: R1148

www.aggregate.org | hankd@engr.uky.edu

The Aggregate refers to a collection of researchers and the technologies that they use to make the components of a parallel computer work better together. Since before our first Linux PC work in the PAPERS project, we have been considering all aspects of Compilers, Hardware Architectures, and Operating Systems (KAOS) together, optimizing system performance rather than performance of the individual parts, focusing primarily on improving Linux PC clusters. This year, led by the University of Kentucky, The Aggregate will again present research from several universities. Demonstrations will focus on applications developed using the new technologies created by The Aggregate, especially a CFD (Computational Fluid Dynamics) code that is currently a Gordon Bell Price/Performance award finalist. New technologies being presented include 3DNow!-enhanced libraries, Flat Neighborhood Network (FNN) design tools, and a variety of enhancements on earlier research developments including PAPERS and our video wall library.

THE INSTITUTE OF PHYSICAL AND CHEMICAL RESEARCH (RIKEN)

Booth Number: R1092

www.riken.go.jp/engn/index.html
atsushi@atlas.riken.go.jp

We are developing the Molecular Dynamics Machine (MDM), a heterogeneous computer system for MD simulations. It accelerates calculation of Coulomb force using two special-purpose hardware systems—MDGRAPE-2 and WINE-2. The full system will be complete at the end of this year, and its peak performance will be about 75 teraflops. We exhibit the building blocks of MDGRAPE-2 and WINE-2. We also present live MD simulations on a 128 gigaflop subset of the MDM system.

UNIVERSIDADE DE SAO PAULO

Booth Number: R715

www.lsi.usp.br/~spade | kofuji@lsi.usp.br

A prototype of a next generation "global information infrastructure" will be demonstrated, based on high performance, high availability and with high security servers, and small, portable, wearable, and wireless computers. Instead of "Grand Challenge" class applications, these servers are going to store, process, and retrieve infomation to a large number of endpoint "terminals." To demonstrate the concept we are going to demonstrate a virtual collaborative secure work environment and several kinds of remote monitors.

UNIVERSITY OF GREENWICH

Booth Number: R1121

www.gre.ac.uk | P.Leggett@gre.ac.uk

The Computer Aided Parallelisation Tools (CAPTools) is an interactive toolkit for the semi-automatic parallelization of serial Fortran code into a message passing SPMD style parallel code. The parallel code produced by CAPTools is portable, efficient, and very similar to the original serial, making it easy to understand and allowing for further manual optimisation of the generated code, if required. CAPTools has been developed by the Parallel Processing Research Group (PPRG) at the University of Greenwich, England for over a decade. The group's focus has been on the development of manual parallelization strategies for Computational Mechanics codes (structured and unstructured mesh) with the subsequent embedding of these strategies within CAPTools. CAPTools was exhibited at SC99 and has been installed and used at many sites in the U.S. including NASA, DoD and various regional supercomputing centres.

UNIVERSITY OF HOUSTON

Booth Number: R1219

www.cs.uh.edu/~chapman
chapman@cs.uh.edu

Many research and development projects currently aim to facilitate the use of computational grids for the execution of supercomputing applications. Such environments promise an improvement in the utilization of existing computational resources, as well as faster time to completion for the individual user's job. However, an efficient utilization of grid resources is currently impeded, from both points of view, by the manual effort involved in resource selection and job submission. These tasks are supported at a high level for certain classes of applications only. Our goal is to permit the majority of grid users to specify the needs of their job in a convenient manner, and to support them in the task of detecting and selecting computational resources that are likely to

meet these needs. In this exhibit, we display on-going work taking place in this area at the University of Houston.

UNIVERSITY OF ILLINOIS AT CHICAGO, NATIONAL CENTER FOR DATA MINING/NATIONAL SCALABLE CLUSTER PROJECT/LABORATORY FOR ADVANCED COMPUTING

Booth Number: R798

www.ncdm.uic.edu | grossman@uic.edu

Project Data Space will link 14 sites across five continents to demonstrate a new infrastructure to handle: 1) the remote access, analysis, and mining of data and 2) the distributed analysis and mining of data. Researchers at the University of Illinois at Chicago will be joined by over 15 participating institutions and organizations and demonstrate a variety of DataSpace tools and applications. The exhibit will feature a high performance, wide-area, data-intensive testbed called the Terabyte Challenge, developed by the National Scalable Cluster Project (NSCP). The NSCP is collaboration led by research groups at the University of Illinois at Chicago and the University of Pennsylvania, and focused on wide area cluster computing.

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, CENTER FOR SUPERCOMPUTING RESEARCH AND DEVELOPMENT

Booth Number: R891

www.csrd.uiuc.edu | scarroll@uiuc.edu

The CSRD booth will feature the PROMIS compiler, a multilingual, retargetable, optimizing, and restructuring, parallelizing compiler. Also featured will be the Scalable Java Server built on nanoThread technology to provide better interrupt handling and scalability.

UNIVERSITY OF MANCHESTER, MANCHESTER COMPUTING

Booth Number: R1236

www.man.ac.uk | j.m.brooke@man.ac.uk

Manchester Computing is Europe's premier university computing facility supporting world class research and teaching in all disciplines. It is used by the UK academic community and, increasingly, by many overseas higher education institutions. It is also a major node in the Eu-sponsored EUROGRID project and is a member of the eGRID forum. The major components are: CSAR, the flagship HPC service for UK academia and commerce, and MIIMAS, providing on-line access to major datasets for the social and physical sciences. Manchester is a major network operations center on SuperJanet, the UK academic network, and is taking a leading role in establishing connectivity with partners in Internet2 projects. MC provides computing services to the University of Manchester, through the Manchester Research Centre for Computational Science and the Manchester Visualization Centre. It is an international center for HPCN and visual supercomputing, with a recently installed Virtual Reality center specializing in Virtual Medicine. More than 25,000 users from over 150 UK institutions use Manchester Computing. Highlight: Demonstration of a global metacomputer linking the US, Europe, and Japan. We demonstrate real-time data processing of the signal from the Jodrell Bank radio telescope searching for new fast pulsars that can test the fundamental theories of physics.

UNIVERSITY OF TENNESSEE, INNOVATIVE COMPUTING LABORATORY

Booth Number: R202

www.icl.cs.utk.edu | swells@cs.utk.edu

The Innovative Computing Laboratory (ICL), which is part of the Computer Science Department at the University of Tennessee, carries out research in high performance computing focused around four main areas: distributed network computing, numerical linear algebra, software repositories, and performance optimization and benchmarking. Under the leadership of Jack Dongarra, this research has resulted in creative and original solutions to many significant problems inherent in high performance computing applications and architectures. Among the widely known ICL projects and technologies to be presented are NetSolve, Automatically Tuned Linear Algebra Software (ATLAS), Repository in a Box (RIB), the Top500 Report, and the Performance counter Application Programmer Interface (PAPI). This ICL research receives government support from the NSF, DOD, DOE and NASA. We also collaborate closely with companies such as Sun, IBM, SGI, and Cray in order to address the demanding high performance computing challenges confronted by the research and engineering community.

UNIVERSITY OF TOKYO, GRADUATE SCHOOL OF FRONTIER SCIENCES, ADVENTURE PROJECT

Booth Number: R1115

www.adventure.q.t.u-tokyo.ac.jp
yoshi@q.t.u-tokyo.ac.jp

ADVENTURE is one of the projects in the Computational Science of developing an advanced general-purpose computational mechanics system, which can analyze a model with 10-100 million DOFs within 1 hour to 1 day using the world's fastest computer with 30-100 TFLOPS in 2002. To efficiently deal with such an ultra-large scale model, massively parallel algorithms are embedded in the pre-, main-, and post-processes. Neural networks, GAs, and VR then enable ideal user-friendliness. A distributed object-oriented design and programming method is employed such that new theories and algorithms can be easily implemented in various heterogeneous computer environments. In the exhibition, an overview of the project and key technologies such as the Hierarchical DDM are described, together with some numerical examples of a full scale 3D model of nuclear pressure vessel with 100 millions DOFs.

UNIVERSITY OF TOKYO, GRAPE

Booth Number: R713

www.grape.c.u-tokyo.ac.jp/grape/
fukushig@provence.c.u-tokyo.ac.jp

GRAPE (GRAvity piPE) is a special-purpose computer for astronomical N-body problem. These are machines specialized for evaluation of gravitational interaction between particles, and work as backend processors attached to general-purpose computers. The GRAPE computers handle the calculation of the interaction, and the front-end general-purpose computers handle everything else, such as time integration and I/O. Our projects were awarded 1995, 1996 and 1999 Gordon Bell Prizes. We are now developing GRAPE-6, whose peak speed will be over 100 teraflops. Our simulation using a small configuration of the GRAPE-6 system is selected as this year's finalist for the Gordon Bell prize.

UNIVERSITY OF UTAH, CENTER FOR HIGH PERFORMANCE COMPUTING

Booth Number: R1114

www.chpc.utah.edu | julia@chpc.utah.edu

The Center for High Performance Computing provides large-scale computer resources to facilitate advances in the field of computational science at the University of Utah. The projects supported by CHPC come from a wide array of disciplines requiring large capacity computing resources, both for calculating the solutions of large-scale, 2-D and 3-D problems, and for graphical visualization of the results.

UNIVERSITY OF VIRGINIA, LEGION

Booth Number: R1098

www.legion.virginia.edu
humphrey@cs.virginia.edu

Legion is an object-based software project designed to be the meta-operating system for computational grids. Legion runs on top of host operating systems such as UNIX and Windows and provides the illusion of a single virtual machine from a collection of heterogeneous underlying physical machines ranging from workstations to supercomputers. From the project's beginning, in late 1993, the Legion Research Group's goal has been a highly useable, efficient and scalable system founded on solid principles. We have been guided by our own work in object-oriented parallel processing, distributed computing, scheduling, and security, as well as by decades of research in distributed computing systems. Our system addresses key issues such as scalability, parallelism, programming ease, faults tolerance, security, and site autonomy. At SC2000, the newest features of Legion will be showcased in a Legion network of over 400 hosts, emphasizing ease of use, security, and performance.

UNIVERSITY OF WISCONSIN-MADISON & UNIVERSITY OF MARYLAND

Booth Number: R1091

www.cs.wisc.edu/paradyn
and www.cs.umd.edu/projects/dyninstAPI/
bart@cs.wisc.edu

Paradyn is a scalable tool particularly suited to measure, efficiently and effectively, the performance of long-running programs on large-scale parallel/distributed machines, such as SMPs and (heterogeneous) clusters of workstations. Novel techniques allow instrumentation of a program while it is running—automatically controlling the instrumentation to collect only the information needed to find current problems. Dynamic Instrumentation directly instruments unmodified applications during their execution, greatly reducing the total amount of performance data collected: a machine-independent interface, known as dyninstAPI, is used as the basis of a wide range of research and commercial tools. Decision support provided by the Performance Consultant isolates bottlenecks to specific causes and parts of a program with an automated search and measurement process. Additional research includes extending dynamic instrumentation to work on commodity operating system kernels, on-the-fly code tuning, and managing and exploiting execution data from experiments and revisions throughout the evolutionary development of a parallel application.

UNIFIED PARALLEL C (UPC)

Booth Number: R1221

www.hpc.gmu.edu/~upc | tarek@gmu.edu

This research exhibit will describe the underlying concepts of UPC, an explicitly parallel extension of ANSI C designed to provide both good performance and ease of programming for high-end parallel computers. UPC provides a distributed shared-memory-programming model and includes features that allow programmers to specify and exploit memory locality. Such constructs facilitate explicit control of data and work distribution among threads so that remote memory accesses are minimized. Thus, UPC maintains the C language heritage of keeping programmers in control of and close to the hardware. Among the advanced features of UPC are shared and private pointers into the shared and private address spaces, shared and private data, efficient synchronization mechanisms including nonblocking barriers, and support for establishing different memory consistency models. In addition to its original open-source implementation, UPC has gained acceptance from several vendors who are producing exploratory compilers. Additional information can be found at hpc.gmu.edu/~upc.

WASEDA UNIVERSITY

Booth Number: R991

www.waseda.ac.jp

muraoka@muraoka.info.waseda.ac.jp

This is the third year of exhibiting this project at SC. The Japanese government, through the Japan Society, supports this project for the promotion of science. The project's objective is to develop a parallelizing, restructuring compiler as well as related tools for parallel and heterogeneous distributed computing environments. The project puts equal emphasis on both practical and theoretical sides. As an outcome of the project, we have developed a software system called GCI for distributed computing. The GCI will be used as part of a Japanese version of GRID, which is planned to be interconnected to a global GRID.

SC2000 RESEARCH GEMS

Research Gems Chair: Jennifer Schopf, Northwestern University

Research Gems provide an opportunity for researchers at SC2000 to present insights into the solutions of specific research problems in high performance networking and computing. This category of participation replaces the traditional Poster Exhibits, and will be prominently placed in the Exhibition Hall.

This year, the Research Gems will have special Open Houses on Wednesday, November 8 and Thursday November 9, from 10AM–11AM. Don't miss this chance to discuss the posters with the authors.

Adding OpenMP to an Existing MPI Code: Will It be Beneficial?

Joseph D. Blahovec, Air Force Research Lab
Keith L. Cartwright, Air Force Research Lab

ICEPIC (Improved Concurrent Electromagnetic Particle In Cell), developed at the Air Force Research Laboratory, is a 3-D particle-in-cell code specifically designed for parallel high performance computing (HPC) resources. ICEPIC simulates collision-less plasma physics phenomena on a Cartesian grid. ICEPIC has several novel features that allow efficient use of parallel architectures. It is written in ANSI C and uses the MPI message passing standard to provide portability to a variety of HPC systems. Currently, we are determining if adding OpenMP to the existing MPI will increase our performance on the new IBM SMP.

Advantages of Multi-block Curvilinear Grid and Dual-Level Parallelism in Ocean Circulation Modeling

Phu V. Luong, University of Texas, Austin, ERDC MSRC Clay P. Breshears, Rice University, ERDC MSRC Andy Haas, LOGICON, NAVO MSRC
http://www.ticam.utexas.edu/~vphu/On-Site/Onsite.html

A multi-block curvilinear grid technique is used in ocean modeling to eliminate problems inherent in the traditional one-block rectangular structured grid. A dual-level parallel technique is included to improve the performance of the Multi-block Grid Princeton Ocean Model (MGPOM) ocean circulation model. This technique involves the use of Message Passing Interface (MPI) in which each grid block is assigned to a unique processor. Since not all grid blocks are of the same size, the workload varies between MPI processes. To alleviate this, OpenMP dynamic threading is used to improve load balance.

Application-level Implementation of Asynchronous Methods in a CORBA-based Distributed Object Database

Milorad Tosic, Rutgers, The State Univ. of New Jersey Helen Berman, Rutgers The State Univ. of New Jersey, Department of Chemistry , John Westbrook, Rutgers The State Univ. of New Jersey, Department of Chemistry
http://pdb.rutgers.edu:4000/lima/

In this poster we consider a concurrency model for method calls in a CORBA-based implementation of a distributed object database for the topology based sub-similarity search of chemical structures. Each object in such databases typically has both computationally intensive and light-weight methods. Applying the concurrency model at the

ORB/BOA level in OMG CORBA enforces the same model for all of the methods of an object. We propose a wrapper class for the C++ multi-threading library that enables a concurrency model choice at the method level. This abstraction is developed in a design pattern style for managing threading issues in any application.

Applications of Parallel Process HiMAP for Large-scale Multidisciplinary Problems

Guru Guruswamy, NAS Division, Ames Research Center, David Rodriguez, Eloret Inc, Ames Research Center, Mark Postdam, Eloret Inc, Ames Research Center
http://george.arc.nasa.gov/gguruswa/home.html

A modular process to simulate coupled multi-physics interactions of flexible vehicles using high fidelity equations is developed. The process is designed to execute on massively parallel processors (MPP). Computations of each discipline are spread across processors using IEEE standard message passing interface (MPI) for inter processor communications. Disciplines can run in parallel using a middleware MPIRUN developed based on MPI and C++. In addition to discipline parallelization and coarse-grain parallelization of the disciplines, embarrassingly parallel capability to run multiple parameter cases is implemented using a script system. The poster will show the development and applications of the process. Results that highlight load balancing and portability issues will be shown.

APPMAP: A New Way to Predict Application Performance in Minutes

John Gustafson, Sun Microsystems Rajat Todi, Ames Laboratory/Iowa State University Don Heller, Ames Laboratory/Iowa State University
http://www.scl.ameslab.gov/Projects/APPMAP/index.html

Our research proposes a convolution technique using hardware and software signatures to accurately predict the system's performance in a relatively short time. Hardware signatures are obtained by actually running the broad-spectrum benchmark HINT or derived from the machine characteristics using analytical HINT (AHINT). Application signatures are extracted from the application as a black box using hardware performance counters and dynamic tracing. In this poster we will present the model and show the predicted result of SPEC '95. We will validate the model using statistical analysis such as correlation, rank correlation, and linear prediction.

Automatic TCP Socket Buffer Tuning

Jian Liu, National Laboratory for Applied Network Research, National Center for Supercomputing Applications, Jim Ferguson, National Laboratory for Applied Network Research, National Center for Supercomputing Applications
http://dast.nlanr.net/Projects/projects.html#autobuf

Relying on the static system socket buffer size for TCP data transfers seriously limits the efficient use of high-speed network bandwidth and system memory. The most popular application run across high-speed networks is FTP, whose performance is directly affected by the TCP window size. We propose an Automatic Buffer Tuning method at the application level as an attempt to solve this problem. An Automatic Buffer Tuning enabled FTP client demonstrates that dynamic adjustment of TCP socket buffer size at the TCP connection setup time causes enhanced FTP application performance, improved utilization of network bandwidth, flexible reservation

of system memory, and requires no kernel modification.

Core AS (Autonomous System) Internet Topology

Bradley Huffaker, CAIDA/SDSC/UCSD, Theresa Ott Boisseau, CAIDA/SDSC/UCSD

http://www.caida.org/analysis/topology/as_core_network/

When the Internet was in its infancy, monitoring traffic was relatively simple. However, after experiencing phenomenal growth in the 1990's, tracking connectivity has become a daunting task. Recently, CAIDA researchers have attempted to strip away lesser connected autonomous systems (or ASes) in order to find out how Internet connectivity is distributed among ISPs. This graph, showing peering richness and geographic information, clearly reveals the highly 'core connected' topology of ASes based in North America. All except one of the top 15 ASes are based in the U.S., and there are few links directly between ISPs in Asia and Europe.

Dynamic Load Balancing Techniques for Improving Adaptive Mesh Refinement

Zhiling Lan, ECE Department, Northwestern University, Valerie Taylor, ECE Department, Northwestern University, Gregory Bryan, MIT

http://ece.nwu.edu/~zlan/enzo

ENZO is one of the successful implementations of AMR (Adaptive Mesh Refinement) on parallel and distributed systems. One of the key issues related to this implementation is dynamic load balancing which allows large-scale and adaptive cosmology applications to run efficiently on parallel systems. We present a new technique for dynamic load balancing of AMR applications. We describe our grid-splitting technique and present some results illustrating the performance gains for the ENZO applications executed on parallel systems. Several metrics for measuring the quality of dynamic load balancing will be used.

Efficient Process Migration Mechanisms for a Heterogeneous Distributed Computing Environment

Kasidit Chanchio, Illinois Institute of Technology, Xian-He Sun, Illinois Institute of Technology

http://www.cs.iit.edu/~scs/SNOW/

This poster presents a solution for process migration in a heterogeneous distributed computing environment. Our approach solves three fundamental problems in transferring execution, memory, and communication states of a process. To migrate the execution state, we propose a technique using source code annotation. To migrate the memory state, we propose a graphical model and its associated runtime system. Finally, we propose process migration and communication protocols to migrate the communication state. We have tested our migration mechanisms on several sequential and parallel programs. Experimental results advocate the correctness and practicality of our approach.

Enabling Technologies for High-Performance Computing Portals: The PUNCH Approach

Nirav H. Kapadia, Purdue University, Sumalatha Adabala, Purdue University, Renato J. Figueiredo, Purdue University, Dolors Royo, UPC, Spain, Jose' Miguel-Alonso, UPV, Spain, Jose' A. B. Fortes, Purdue University, Mark S. Lundstrom, Purdue University

http://www.ece.purdue.edu/punch

Network-centric computing promises to revolutionize the way in which

computing services are delivered to the end-user. Analogous to the power grids that distribute electricity today, computational grids will distribute and deliver computing services to users anytime, anywhere. PUNCH, the Purdue University Network Computing Hub, is a network-computing system that allows seamless management of applications, data, and machines spread across administrative domains and wide-area networks. Users can access any application from anywhere via standard Web browsers. PUNCH already provides computing services to 800 users across 10 countries. The poster will focus on the core technologies that make up the PUNCH infrastructure.

Extending Quality of Service on the Grid: Gara and PBS

Thomas Milford, Computer Science Department, Northwestern University, Jennifer M. Schopf, Computer Science Department, Northwestern University
http://anchor.cs.nwu.edu/gara

Gara is the Quality of Service component of the Globus Toolkit. Currently, it provides a means for making reservations through Globus resources including disk bandwidth, CPU time and fork. We have extended Gara to allow communication with PBS thus taking advantage of the PBS reservation features. By managing the reservations in Gara, we add the ability to allow the users to define callback procedures (ie. when the reservation is about to go 'live'). We also have a central point where we can track all of the reservation status information.

Extending the Portable Batch System with Preemptive Job Scheduling

Gabriel Mateescu, National Research Council Canada
http://www.sao.nrc.ca/~gabriel/presentations/sc2000

The Portable Batch System (PBS) provides robust and adaptable job and resource management. Unfortunately, PBS does not support job preemption. In environments with multiple user groups that have dynamic resource requirements, the lack of preemption causes throughput and resource utilization penalties. We propose a preemptive scheduling scheme that avoids these pitfalls. We employ a preemption manager, along with PBS hooks for stopping and resuming jobs. Based on resource requirements, system utilization, and scheduling policies, the manager determines when and what jobs to preempt, then instructs PBS to preempt the jobs. Preempted jobs become eligible for resumption after the requests that caused preemption are satisfied.

General Portable SHMEM Library for High Performance Computing

Krzysztof Parzyszek, Ames Laboratory, Jarek Nieplocha, Pacific Northwest National Laboratory, Ricky A. Kendall, Ames Laboratory
http://www.scl.ameslab.gov/Personnel/rickyk/gpshmem

GPSHMEM is a communication library that follows the interface of the original Cray SHMEM library, but attempts to achieve full portability. It is implemented on top of the Aggregate Remote Memory Copy Interface (ARMCI) and a message-passing library (currently MPI). The functionality of GPSHMEM covers most of the SHMEM functionality available on the Cray T3D. The Generalized Portable

SHMEM library will benefit CRI T3D/T3E users by providing them a mechanism to use SHMEM on different parallel supercomputers, basically any system that has an ARMCI port. Tests show that the GPSHMEM programming model can be as efficient as as the underlying ARMCI memory access implementation.

GridSearcher

S. R. Melody, Computer Science Department, Northwestern University, Jennifer M. Schopf, Computer Science Department, Northwestern University
http://anchor.cs.nwu.edu/GridSearcher/

This poster presents our work in developing a cross platform tool for resource discovery on the Grid. It expedites the process of searching for a machine and makes the Grid more accessible to application scientists. It attempts to create a quick, simple, and powerful way to find what you are looking for. In order to achieve this goal and maintain as much portability as possible, the software package has two modules, the first searches an MDS tree to create an XML/DOM tree and the second creates a table and generates an easy to search representation.

HDF5: High Performance Science Data Solution for the New Millennium

Albert Cheng, NCSA/University of Illinois, Michael Folk, NCSA/University of Illinois
http://hdf.ncsa.uiuc.edu

The Hierarchical Data Format (HDF) has been used for scientific data management since 1988. NCSA, in collaboration with NASA, the ASCI project and others have developed a new format and library, HDF5, that addresses the needs of vastly expanded computational and storage systems. I/O performance, especially parallel I/O, plays a critical role in the HDF5 project. HDF has also incorporated the Grid Forum access features into the HDF5 library for remote resource access in the Virtual Machine Room environment. An XML-DTD specification of the HDF5 format is created to support exchange of HDF5 data for Web-based applications.

Model-Based Integration of Heterogeneous Neuroscience Data Sources

Bertram Ludaescher, San Diego Supercomputer Center, UCSD, Amarnath Gupta, San Diego Supercomputer Center, UCSD, Maryann. E. Martone, Department of Neuroscience, UCSD, Ilya Zaslavsky, San Diego Supercomputer Center, UCSD,
http://www.npaci.edu/DICE/Neuro/SC2000/

We present a novel wrapper-mediator approach called Model-Based Mediation in which views are defined and executed at the level of conceptual models rather than at the structural (XML) level. Another novel feature of our architecture is the use of Domain Maps, a kind of semantic net used to mediate across sources from "multiple worlds"—in our case different Neuroscience data sources. As part of registering a source's conceptual model with the mediator, the wrapper creates a "semantic index" of its data into the Domain Map. A prototype establishing the viability of the approach is operational. Our knowledge-guided mediation is used within the "Federation of Brain Data" project under NPACI.

WEDNESDAY & THURSDAY
EXHIBIT HALL | 10 – 11AM
RESEARCH GEMS/ESCAPE 2000 OPEN HOUSE

Platforms: An Accurate Arithmetics Approach

Yun (Helen) He, Lawrence Berkeley National Laboratory, Chris H.Q. Ding, Lawrence Berkeley National Laboratory

www.nersc.gov/research/SCG/ocean/NRS

Numerical reproducibility of large-scale scientific simulations, especially climate modeling, on distributed memory parallel computers are becoming critical issues. In particular, global summation and dot products of distributed arrays are very susceptible to rounding errors. We analyzed several accurate summation methods and found that two methods are particularly effective to improve (ensure) reproducibility: Kahan's self-compensated summation and Bailey's double-double precision summation. We provide an MPI operator MPI_SUMDD to work with MPI collective operations to ensure a scalable implementation on large number of processors. The final methods are particularly simple to adopt in practical codes.

Parallel Programming in Java with OpenMP-like Directives

J. Mark Bull, Edinburgh Parallel Computing Centre, Mark E. Kambites, University of York, Jan Obdrzalek, Masaryk University, Brno

www.epcc.ed.ac.uk/research/jomp

We present a specification of directives, library routines, and system properties (in the spirit of OpenMP) for shared memory parallel programming in Java. A prototype reference implementation consisting of a compiler and a runtime library is described. The compiler translates Java source code with directives to Java source code with calls to the runtime library, which in turn uses Java threads to implement parallelism. The whole system is pure Java, and can be run on any Java virtual machine. The performance of the system is compared to hand coded Java threads and to commercial C and Fortran OpenMP implementations.

Policy Specification and Restricted Delegation in Globus Proxies Specification and Restricted Delegation in Globus Proxies

Babu Sundaram, University of Houston, Christopher Nebergall, Western Illinois University, Steven Tuecke, Argonne National Laboratories

http://www.cs.uh.edu/~babu/poster.html

Grid Security gets complicated when users need a single log-on to distributed resources that have heterogeneous local policies. Policy languages with rich features and functionalities are needed to specify security policies in proxies. This approach is examined for the Globus toolkit. Also, facilities for site administrators to specify local policies were considered. Classified Advertisements, from the University of Wisconsin Condor project, were used to specify and evaluate policies as attributes. The appropriate attributes have been identified. Also, the authentication and the authorization processes are successfully implemented by modifying the Globus Proxy initiation, Gatekeeper, and Job-manager. This provides fine-grained control and more protection against stolen proxies.

Rapid Design Realization of 3D Woven Composites by HPC

Seung Jo Kim, Department of Aerospace Engineering, Seoul National University, Chang Sung Lee, Department of Aerospace Engineering, Seoul National University, Heon Shin, Department of Aerospace Engineering, Seoul National University, Jeong Ho Kim, KORDIC Supercomputing Application Lab

http://aeroguy.snu.ac.kr

A design realization of 3D woven composites is tackled by a full-scale FEA (not by the limited "unit cell approach") through High Performance Computing (HPC) techniques. An accurate finite element modeling, which includes warp yarns, filler yarns, stuffer yarns and resin regions and reveals these complex geometrical characteristics is first prepared. An efficient direct solver, parallel multi-frontal solver and a parallel explicit dynamic solver with central difference scheme are used. An optimal design pattern of 3D orthogonal woven composites, which has a required stiffness in-plane directions and a good through-the-thickness strength without local stress concentrations, can be achieved.

A Storage Broker for the Globus Environment—A ClassAd Based Implementation

Sudharshan S. Vazhkudai, The Univ. of Mississippi, Steven Tuecke, Argonne National Laboratory
www.globus.org

An increasing number of scientific applications ranging from high-energy physics to computational genomics require access to large amounts of data (to the tune of tera and peta bytes) with varied quality of service requirements. This diverse demand has contributed to the proliferation of storage system capabilities, thus making storage devices, an integral part of the Grid environment. In order to access and utilize storage systems efficiently in the Globus Grid environment, we need to: provide an efficient architecture (a Storage Broker) for querying and selecting them based on application requirements; determine their properties and expose them to the interested Grid services and high-end applications. Our work illustrates the architecture of such a storage broker for the Globus Grid environment.

STWAVE: A Case Study in Dual Level Parallelism

Rebecca A. Fahey, Engineer Research and Development Center, Major Shared Resource Center
http://chl.wes.army.mil/research/wave/wavesprg/numeric/wtransformation/stwave.htp

This case study explores a dual-level parallel implementation of the Steady-State Spectral Wave Model (STWAVE) code, a near-shore, wind-wave, growth and propagation simulation program with two natural levels of parallelism. The embarrassingly parallel calculations for processing multiple wave runs are distributed via MPI, resulting in a linearly scalable code. However, the achievable speedup is limited by the number of wave runs. To obtain additional speedup, loop-level parallelism is exploited with OpenMP. For STWAVE, the speedup attained with a dual-level approach surpasses the speedup possible with either method alone.

Two Levels of Parallelism for Models of Tracer Particle Transport

Vickie E. Lynch, Oak Ridge National Laboratory, Benjamin A. Carreras, Oak Ridge National Laboratory, Nathaniel D. Sizemore, Oak Ridge National Laboratory
http://www.ornl.gov/fed/theory/avalanche.html

Avalanche transport may be a dominant transport mechanism in magnetically confined plasmas. Models used range from a sandpile model to 3-D turbulence models. To understand what happens to transport, we evolve test particles simultaneously with the evolution of the macroscopic field. These models require

long time evolutions and large grids to resolve the different scales. With the IBM SP capable of two levels of parallelism, we want to use the shared memory processors to do calculations of a partition of the grid using OpenMP and then run different bunches of particles on the distributed memory nodes using MPI.

Utilizing Idle Workstations in a Scheduled Parallel Computing System

Scott Hansen, Brigham Young University,
Quinn Snell, Brigham Young University,
Mark Clement, Brigham Young University
http://ncl.cs.byu.edu/yrm

Increasing demand for parallel computational resources has motivated many researchers to use idle desktop systems in an attempt to increase computational power with minimal cost increase. This presentation describes the BYU Resource Manager (YRM) that, in conjunction with the Maui Scheduler, supports non-dedicated resources. We show that by effectively handling the complexities of non-dedicated resources, a resource management system can increase job throughput and job turnaround time. We also show that the YRM provides a resource management framework necessary for dynamic process management in the next generation of parallel programming libraries such as MPI-2.

VG Cluster: Large Scale Visual Computing System for Volumetric Simulations

Shigeru Muraki, Electrotechnical Laboratory,
Kazuro Shimokawa, Electrotechnical Laboratory, Masato Ogata, Mitsubishi Precision Co., LTD., Kagenori Kajihara, Mitsubishi Precision Co., LTD., Kwan-Liu Ma, University of California-Davis, Yutaka Ishikawa, Real World Computing Partnership
http://www.etl.go.jp/~visual/

This research aims at the development of a parallel visual computing system for large-scale scientific simulations. Our initial design, named VG cluster, is based on a Linux PC cluster with a high-speed switch, the specially tuned system software, and the distributed volume rendering technique. We subdivide a large simulation space into sub-volumes of similar size and distribute the simulation and visualization computations to individual PCs. A 3-d neuron excitement simulation using confocal microscope data from real neurons will be used for our study on the VG cluster. The test results and our future plan of implementing the rendering process in hardware will be presented.

Visual MD: An Innovative Approach to Molecular Simulation

Jennifer Hare, U.S. Army Research Laboratory,
Betsy Rice, U.S. Army Research Laboratory,
Jerry Clarke, Raytheon Systems Company,
Margaret Hurley, CCM PET ARL MSRC/OSC,
William Mattson, Department of Physics,
University of Illinois, Urbana Champaign

Computational molecular modeling is recognized as an efficient step of modern developmental programs. We are developing an integrated suite of user-friendly, highly scalable molecular simulation codes to study reactive and non-reactive processes in solid

materials. The software is designed with flexibility, portability and ease of use in mind, to meet the ever changing user needs and changes in computational platforms. These codes allow both the non-expert in molecular simulation to exercise the codes and an expert to customize them. Unlike monolithic solutions, this integrated suite of software tools includes a GUI, scientific visualization, and the ability to add custom functionality.

Visualization of Events from the Relativistic Heavy Ion Collider

Michael D. McGuigan, Brookhaven National Lab, Stephen Murtagh, Carnegie-Mellon University
http://www.ccd.bnl.gov/visualization

This Research Gem discusses a high performance visual display of events from the Relativistic Heavy Ion collider at Brookhaven National Lab used to search for a new state of matter not seen since the big bang. A comparison is made of several hardware and software combinations representing the events. Images and animations corresponding to an OpenInventor application displaying the collisions and detector responses are presented. Conclusions are drawn from comparisons of fitted track data directly from the visualizations.

Visualization of Extra Dimensions from String Theory

Michael D. McGuigan, Brookhaven National Lab, Stephen Murtagh Carnegie-Mellon University
http://www.ccd.bnl.gov/visualization

We address the use of advanced visualization techniques (implicit functions, algebraic geometry and topology) for high performance visualizations of the internal dimensions predicted to exist by String Theory, the leading candidate for a theory beyond the Standard Model of particle physics. One complication is that the theory is only consistent in ten dimensions. Since we live in four dimensions, three space and one time, the extra six dimensions are assumed to be curled up in small space that can't be directly seen. In this research we magnify the internal dimensions by a factor of 10^34 so they can be analyzed visually using a powerful visualization computer.

Workload Characterization and Similarity Analysis of SPECcpu95 Benchmarks

Abdullah I. Almojel, King Fahd University of Petroleum and Minerals, Ali S. AlSwayan, King Fahd University of Petroleum and Minerals
www.ccse.kfupm.edu.sa

Quantifiable methods can help establish the basis for scientific design of parallel computers based on application needs, to optimize performance to cost. This research work is based on the Vector Space model of workload representation and similarity. It is shown that this model can effectively represent real-life workloads. The model also lends itself to many practical applications. These applications include the design of effective Benchmark Suites, Experimental Parallel Computer Design, and Performance Prediction of real-life applications on real machines. This research work investigates the architectural characteristics and the similarity between the SPECcpu95 benchmarks in order to remove the redundant workloads (exercise the machine similarly) from the benchmark suite.

SC2000 INDUSTRY EXHIBITS

Vice Chair, Exhibits: Joe McCaffrey, Mississippi State University
Exhibition Management: Don Collier, DC Expositions, Inc.

ADIC

Booth 224
P.O. Box 97057; Redmond, WA 98073
Phone: 425-881-8004
Fax: 425-881-2296
Email: info@adic.com
URL: www.adic.com

With an installed base of more than 65,000 automated libraries, innovative storage management software, Fibre Channel SAN solutions, and NAS appliances, ADIC is a leading device-independent storage solutions provider to the open systems marketplace. The company offers a broad range of products designed to enhance organizations' ability to store, protect, manage and use their rapidly growing network data. Call to speak to a Territory Manager.

ALIBRE, INC.

Venture Village A
1701 N. Greenville Ave, Suite 702
Richardson, TX 75081
Phone: 972-671-8492
Fax: 972-671-8493
Email: info@alibre.com
URL:www.alibre.com

Using a combination of Web-based technology and a revolutionary new business model, Alibre makes state of the art design, collaboration and data management technology accessible to the entire global supply chain. It's affordable, easy-to-obtain, easy-to-use and best of all—it helps you design better products, faster and more profitably.

www.sc2000.org/industry

API (ALPHA PROCESSOR, INC.)

Booth 608
Guy Ludden, Marketing Manager
130C Baker Ave. Ext.; Concord, MA 01742
Phone: 978-318-1100
Fax: 978-371-3177
Email: info@alpha-processor.com
URL: www.alpha-processor.com

API, (Alpha Processor, Inc.) is a leading developer of high performance solutions for high bandwidth and compute-intensive applications dedicated to the development of Linux solutions for the Internet and High Performance Computing (HPC) markets. Addressing the demand for increased memory and I/O bandwidth to power rapidly expanding compute-heavy applications, API's products are ideal for today's high-performance clustering, rendering, Web server and e-commerce solutions. API's industry-recognized design team is committed to working with leading companies to deliver the highest performance, scalable systems available for VARs and system integrators.

APPLIED METACOMPUTING, INC.

Booth 413
Lisa Romano
300 Preston Avenue, 5th Floor
Charlottesville, VA 22902
Phone: 804-951-0154
URL: www.appliedmeta.com

Applied MetaComputing's flagship product, Legion, is a scaleable, fully-integrated, highly secure software platform for next-generation distributed computing applications. Legion enables users to share resources, such as CPU cycles, storage and databases, while also resolving massive computational challenges, harnessing vast storage capacity, and building advanced P2P and distributed computing applications.

THE ASSOCIATION FOR COMPUTING MACHINERY (ACM)

Booth Registration/Lobby Area
ACM Membership Department
1515 Broadway, 17th Floor
New York, NY 10036
Phone: 212-626-0500 or 800-342-6626
Fax: 212-944-1318
Email: orders@acm.org
URL: www.acm.org

ACM is the premier computing society. Members receive numerous benefits including: Communications of the ACM; discounts on all ACM publications, as well as conference registrations; and the option to subscribe to ACM's Digital Library, a collection of over 20 ACM publications including journals, magazines and proceedings. Join ACM today!

ATIPA

Booth 468
Brandon Fuhr
4700 Belleview, Ste. 300; Kansas City, MO 64112
Phone: 800-360-4346
Fax: 816-595-3001
Email: bfuhr@atipa.com
URL: www.atipa.com

Atipa provides scalable, high performance Alpha, AMD, and Intel-based Beowulf clusters. Atipa is dedicated to providing the most cost-effective open source solutions available today. For those considering Open Source and Linux based systems, Atipa offers pre-installation business analysis and consulting services, optimized Linux-based Web, email, FTP, clustering platforms, software, workstations, and an array of installation, service, and support plans.

BLACKSTONE TECHNOLOGY GROUP, INC

Booth 220
Susan Barry
100 Grove Street; Worcester, MA 01605
Phone: 508-793-2162
Fax: 508-793-2962
Email: barry@computefarm.com
URL: www.computefarm.com

Blackstone Technology Group, Inc. is a pioneer in developing and implementing compute farm solutions that deliver secure, scaleable and uninterrupted compute power on demand. The company offers ComputeFarm Advantage, a comprehensive suite containing all the tools and services needed to access, provision and scale compute power. Sun, Celera, Qualcomm, Lucent and other leading innovators rely on Blackstone for compute farm solutions.

CHIARO NETWORKS

Venture Village B

CLEANSCAPE SOFTWARE INTERNATIONAL

Booth 646
Jennifer Nance
1530 Meridian Ave., Ste. 200
San Jose, CA 95125
Phone: 408-978-7000
Fax: 408-978-7002
Email: jennifer@Cleanscape.net
URL: www.cleanscape.net

Cleanscape is a leading innovator of automated software development and testing solutions that simplify build environments for enabling customers to shorten development cycles, optimize resource utilization, increase product quality, and quicken return on investment. Products include: automated test & report generation systems, debuggers, and build management systems for Windows, Unix, Linux.

COMPAQ COMPUTER CORPORATION

Booth 268
Kathleen Hudson
200 Forest St., Mailstop MRO1-3/K14
Marlboro, MA 01752

Phone: 508-467-1499
Fax: 508-467-1961
Email: kathleen.hudson@compaq.com URL: www.compaq.com/hpc

Over the past year, Compaq has been selected to build the world's most powerful supercomputers, solidifying Compaq's strong emergence as the world's leading supplier of supercomputing systems.

- World's most powerful supercomputer: ASCI Q
- World's most powerful supercomputer for public research: PSC/NSF
- Europe's most powerful supercomputer: CEA
- World's most powerful Linux system: Sandia Cplant
- Most powerful supercomputer in the film and video industry: Blue Sky Studios
- World's leading supplier of supercomputers for genomics research

Exhibits and presentations on these projects will be featured in the Compaq booth.

COMPUTER NETWORK TECHNOLOGY

Booth 640

COMPUTER SCIENCES CORPORATION

Booth 960
Cindy Christopher
4090 South Memorial Parkway
Huntsville, AL 35802
Phone: 256-971-7437
Fax: 256-971-7491
Email: cindy@asc.edu
URL: www.csc.com

CSC provides High Performance Computing services to U.S. Federal and State agencies through its High Performance Computing Center of Excellence unit, and to commercial customers through its Supercomputing Operations unit. CSC provides strategic and operational IT solutions that keep pace with the rapid advances in computing technologies.

CRAY INC.

Booth 724
Ellen Roder
411 First Avenue South
Seattle, Washington 98104
Phone: 651.683.7277
Fax: 651-683-5599
Email: ellen@cray.com
URL: www.cray.com

Newly formed Cray Inc. is the world's leading provider (only U.S. provider) of high-bandwidth, low-latency supercomputers capable of tackling the most economically important problems of science and industry. A multinational research team used a Cray T3E system to set the current world record for sustained speed on a real-world, 64-bit application, at 1.02 teraflops. Cray plans to provide customers with future supercomputers with unrivaled sustained performance for addressing new and previously intractable problems. Visit the Cray Inc. booth to learn more about our company, solutions and customer achievements.

CSP INC., MULTICOMPUTER DIVISION

Booth 708
Brad Stamp
40 Linnell Circle; Billerica, MA 01821
Phone: 978-663-7598, ext.3100
Fax: 978-663-0150
Email: sales@cspi.com
URL: www.cspi.com or www.fastcluster.com

Based in Billerica, Massachusetts, CSP Inc., MultiComputer Division helps its customers solve high performance computing problems in the medical imaging and defense markets by sup-

plying VxWorks and Linux-based multiprocessing platforms with powerful I/O capabilities that require minimal physical space and power. These systems are used in a broad array of applications, including sonar, radar and surveillance signal processing.

DATADIRECT NETWORKS, INC.

Booth 746
Brad Winett
9201 Oakdale Ave.; Chatsworth, CA 91311
Phone: 818-700-7600
Fax: 818-700-7601
Email: sales@datadirectnet.com
URL: www.datadirectnet.com

Storage Area Networking appliance leader DataDirect Networks provides Rich Media communities with high bandwidth, virtualized, scalable and flexible network infrastructures, reducing TCO, enabling higher productivity with increased ROI. The SAN DataDirector, allowing powerful and quick SAN deployment for high performance computing will be exhibited.

DELL COMPUTER CORPORATION

Booth 974
One Dell Way; Round Rock, TX 78682
Phone: 800-727-1100
Fax: 800-926-7595
URL: www.dell.com

Dell is the world's leading direct computer systems company, based on revenues of $28.5 billion for the past four quarters, and is a premier provider of I.T. products and services required for customers to build their Internet infrastructures. Dell designs, manufactures and customizes products and services to customer requirements, and offers an extensive selection of software and peripherals. Information on Dell and its products can be obtained on the World Wide Web at www.dell.com.

DESIGN SCIENCE, INC.

Booth 732
Mike Norwood
4028 Broadway; Long Beach, CA 90803
Phone: 562-433-0685
Fax: 562-433-6969
Email: sales@mathtype.com
URL: www.mathtype.com

MathType is a powerful interactive tool for Windows and Macintosh that lets you create mathematical notation for word processing and desktop publishing documents, web pages, presentations, and for TeX & LaTex and MathML documents. MathType is the professional version of Equation Editor. If you are familiar with our Equation Editor which is included in Microsoft Word, Corel, WordPerfect, AppleWorks and many other products, then you will really appreciate MathType's extra features.

EMC

Booth 172

ENTROPIA, INC.

Venture Village C
Tim Cusac
4350 Executive Drive, Suite 200
San Diego, CA 92121
Phone: 858-623-2840
Fax: 858-623-2847
Email: tcusac@entropia.com
URL: www.entropia.com

Entropia, Inc., is creating the world's largest Internet computing service by converting the unused CPU cycles, storage, memory and Internet connections of commodity PCs into large-scale compute power. Entropia sells this power to biopharmaceutical, financial and entertainment companies as a complement to their existing big iron and as a stand-alone solution. Customers pay a fraction of the cost and get results

in a fraction of the time normally associated with high power computing. Entropia's grid, now running for three years delivering peer reviewed, published results, recently passed the 60 exaflop milestone. Entropia is privately held, headquartered in San Diego.

ESSENTIAL COMMUNICATIONS

Booth 954

Pamela Heater
1551-A Mercantile NE; Albuquerque, NM 97107
Phone: 505-344-0080
Fax: 505-344-0408
Email: pheater@esscom.com
URL: www.esscom.com

Essential develops and manufactures products based on the highest-speed standards-based networking technologies available. The current high-speed networking weitches and adapter cards offered by Essential are based on Gigabyte System Network (GSN) and High Performance Parallel Interface (HIPPI) technologies. GSN reliably moves data at up to 10 gigabits per second, per direction on each data port.

ETNUS

Booth 854

Mary Kay Bunde
651 Crimson Leaf Trail; Eagan, MN 55123
Phone: 651-994-4564
Fax: 651-905-1019
Email: mkay@etnus.com
URL: www.etnus.com

Etnus is a leading provider of software development tools for developers of complex parallel and distributed applications in High Performance Computing as well as the technical commercial or enterprise markets. The award winning debugger and analyzer, TotalView, supports major parallel programming models and eases the debugging task for highly complex and scalable codes.

FAKESPACE SYSTEMS

Booth 466

Jeff Brum, Marketing Manager
809 Wellington St. N.
Kitchener, ON Canada
N2G 4J6
Voice: 519-749-3144
Fax: 519-749-3151
Email: jbrum@fakespacesystems.com
URL: www.fakespacesystems.com

Fakespace Systems is today's leading worldwide provider of devices, structures and environments for advanced interactive visualization technologies. Working to accelerate insight and enable truly interactive collaboration, Fakespace builds on a core of innovative visualization and interactive technologies to provide you with better ways to create, communicate and collaborate.

FOUNDRY NETWORKS, INC.

Booth 166

Karen Smith
2100 Gold Street; PO BOX 649100
San Jose CA 95164, USA
Phone: 408-586-1700
Fax: 408-586-1900
Email: info@foundrynet.com
URL: www.foundrynet.com

Foundry Networks is a leader of high-performance, strategic end-to-end switching and routing solutions including Internet routers, Layer 3 switches and Layer 4-7 Internet traffic and content management switches for enterprises and ISPs. Foundry products are installed in the world's largest ISPs including AOL, EarthLink, AT&T WorldNet, and Cable & Wireless.

FUJITSU

Booth 538

Ron Matlock
2500 CityWest Blvd., Suite 300
Houston, TX 77042
Phone: 713-267-2331
Fax: 713-267-2305
Email: scg-sales@fujitsu.com
URL: www@fujitsu.com

Fujitsu offers a comprehensive range of HPC products. Servers are based on vector-parallel or scalable architecture to best suit application requirements. Software products include high performance compilers and libraries, debugging tools, and system operation utilities. Also supported are hardware/software combinations for high-speed I/O. With these offerings, the Possibilities are Infinite.

GENROCO, INC.

Booth 856

Donald D. Woelz
255 Info Hwy; Slinger, WI 53086
Phone: 262-644-2505
Fax: 262-644-6667
Email: don@genroco.com
URL: www.genroco.com

GENROCO, Inc. is the premier supplier of ultra high performance system and storage area network components for workstations and superservers including Compaq, Fujitsu, IBM, Silicon Graphics, Sun, and other platforms. Along with custom offerings, the 25 year old firm produces high bandwidth Gigabyte System Network, Fibre Channel, Gigabit Ethernet, and HIPPI host bus adapters, routers, bridges, and switches.

HEWLETT PACKARD

Booth 572

HIGH PERFORMANCE STORAGE SYSTEM (HPSS) COLLABORATION

Booth 474

Bob Coyne
IBM Corp.
1810 Space Park Drive; Houston, TX 77058
Phone: 281-335-4040
Fax: 281-335-4231
Email: coyne@us.ibm.com
URL: www.clearlake.ibm.com/hpss

High Performance Storage System (HPSS) is hierarchical storage software meeting the highest capacity and throughput requirements of supercomputers and high-speed networks. HPSS was developed by a collaboration of U.S. Government laboratories and IBM. HPSS now supports IBM AIX and Sun Solaris platforms and Ampex, IBM, and StorageTek tapes.

HIGH PERFORMANCE TECHNOLOGIES, INC.

Booth 173

Nancy S. Coleman, Marketing Manager
4121 Wilson Boulevard, Suite 200
Arlington, VA 22203
Phone: 703-682-5318
Fax: 703-522-5512
Email: ncoleman@hpti.com
URL: www.hpti.com

High Performance Technologies, Inc. (HPTi) is a systems integration and professional services firm specializing in the creative application of leading-edge information technologies. Our projects combine large High Performance Computing (HPC)/parallel servers, advanced networking/intranets, multi-media data bases, and On-Line Analytic Processing (OLAP)/visualization. We enjoy the challenge of making these technologies work to the competitive advantage of our clients in industry and government.

HITACHI
Booth 422

HNF
Booth 851

HPCWIRE
Booth 752

Jennifer Tabor
8445 Camino Santa Fe; San Diego, CA 92121
Phone: 858-625-0070
Fax: 858-625-0088
Email: jennifer@tgc.com
URL: www.tgc.com

HPCwire is the publication of record for technical and commercial high performance computing. It is delivered worldwide each Friday via internet Email: and the World Wide Web, and is the most complete and for many the sole source of information about this fast-paced and important industry. HPCwire will be publishing LIVEwire, daily coverage of SC2000 in Dallas, Texas.

IBM
Booth 562

Parker Blevins, Ph.D.
RS/6000 Scientific and Technical Computing Program Manager
11501 Burnet Road, MD 9263
Austin TX 78758
Office: 512-823-7162 or TL 793-7162
FAX: 512-838-8561 or TL 838-8561
Email: blevinsp@us.ibm.com
URL: www.rs6000.ibm.com

IBM offers a full range of high performance computing products and solutions. The company's Scientific and Technical Computing Group supports government, higher education, petroleum,and manufacturing customers. IBM Research runs the new Deep Computing Institute and the Advanced Computing and Technology Center.

IBM, along with business partners and customers, will demonstrate Deep Computing projects accelerated by the IBM RS/6000 technology. Real-time exhibits and a theater in the booth will both be utilized to demonstrate the science/engineering and solutions. Current and experimental computing technology will be on display.

IEEE COMPUTER SOCIETY
Booth Registration/Lobby Area

Marian Anderson
10662 Los Vaqueros Circle
Los Alamitos, CA 90720-1314
Phone: 714-821-8380
Fax: 714-821-8380
Email: manderson@computer.org
URL: www.computer.org

The IEEE Computer Society strives to be the leading provider of technical information and services to the world's computing professionals. The IEEE Computer Society offers its nearly 100,000 members a comprehensive program of publications, meetings, and technical and educational activities, fostering an active exchange of information, ideas, and innovation. Headquartered in Washington, DC, the society serves its members from offices in Los Alamitos, CA, Tokyo and Brussels. The society is the largest technical society within the Institute of Electrical and Electronics Engineers (IEEE). See more about us at http://computer.org.

INSTRUMENTAL, INC.
Booth 415

Ken Stone
1450 Energy Park Drive, Suite 101
St. Paul, MN 55108
Phone: 651-603-0294
Fax: 651-603-0781
Email: ken@instrumental.com
URL: www.instrumental.com

Instrumental delivers professional services and system management tools focused on high-end UNIX solutions. Our work spans large commercial IT departments to supercomputer centers. Instrumental is uniquely qualified to provide: System architecture design and validation; benchmarking, testing, modeling and simulation for performance analysis and capacity planning; systems integration, and more.

INTEGRATED SENSORS, INC. (ISI)

Booth 831

Harold F. Messias
502 Court Street; Utica, NY 13502
Phone: 315-798-1377
Fax: 315-798-8950
Email: messias@sensors.com
URL: www.sensors.com

ISI will display RTExpress, a compiler and run-time software environment that compiles MATLAB script to run on parallel computers. New RTExpress software will be introduced, that allows parallel MATLAB applications to run across networks of SUN or PC Linux processors. Also, ISI will be introducing the RTExpress Application Server; a scalable Linux PC hardware/software platform used for high performance parallel MATLAB computing.

INTEL/KAI

Booth 849

Tom Rozanski
1906 Fox Drive; Champaign, IL 61820
Phone: 217-356-2288
Fax: 217-356-5199
Email: tom.rozanski@intel.com
URL: www.kai.com

Two technology leaders have combined forces to better serve the HPC community. Visit KAI, now an Intel subsidiary, at Booth #849 to see the latest products for HPC including KAP/Pro Toolset for OpenMP on all high performance platforms, Cluster Computing, plus our newest service offering: the Parallel Applications Center.

JUNIPER NETWORKS, INC.

Booth 981

Jim Jackson, Marketing Manager
1194 No.Mathilda Ave.; Sunnyvale, CA 94089
Phone: 408-745-5020
Fax: 408-745-8908
Email: jjackson@juniper.net
URL: www.juniper.net

Juniper Networks, Inc. converts optical bandwidth into scaleable, differentiable IP services using our new class of integrated silicon and software-based routing systems. These Internet backbone routers are designed and built by industry experts who have broad experience in IP infrastructure applications, as well as in depth understanding of the future direction of IP infrastructures.

KNOWLEDGEPORT ALLIANCE, INC.

Venture Village D

Ken Welch
2004 Fox Drive Suite 2A; Champaign, IL 61820
Phone: 217-355-4110
Email: kenwelch1@kwlport.com

KnowledgePort Alliance, Inc. exists to assist creators of academic software to commercialize their technologies. KPAI acts as the business and devlopment entity shielding creators from the difficulty of transitioning to a business entity. Creators may choose to do this without resigning from their academic positions, without losing ultimate control of their technologies and may participate in the commercialization effort if they so desire.

LINUX NETWORX
Booth 510
Kathy Herrmann, Tradeshow Manager
8689 South 700 West; Sandy, UT 84070
Phone: 801-562-1010
Fax: 801-568-1010
Email: kherrmann@lnxi.com
URL: www.linuxnetworx.com

The Evolocity 4500 from Linux Networx is a fully integrated cluster system that is designed to have the look, feel, and better performance than supercomputer or enterprise level systems. The added benefits of the Evolocity 4500 over traditional hardware include scalability, uptime and ability to upgrade system hardware in the future.

LODGEPOLE TECHNOLOGY INC. WITH SOFTWARE CONCEPTS INC.
Venture Village H
Steve Barnhart
10242 189th CT NE; Redmond, WA 98052
Phone: 425-427-5589
Fax: 425-427-5582
Email: sbhart@ricochet.net
URL: www.lodgepole.com

LTI-SCI is a joint venture partnership of combined talents that brings opensource cluster computing together with software re-engineering and parallel code conversion of legacy Fortran applications. In partnership, LTI-SCI provides a single source for both scalable cluster hardware and software re-engineering through use of exciting, new patented tools and processes.

LOGICON, A NORTHROP GRUMMAN COMPANY
Booth 738
Doug Bove
Phone: 703-866-7432
Fax: 703-866-7536
Email: dbove@logicon.com
URL: www.logicon.com

Headquartered in Herndon Virginia, Logicon is a leader in advanced information technologies, systems and services. User Needs Drive Our Solutions!

Demonstrations include a 3D visualization of a gas-turbine swirling spray combustion simulation, a hands-on demonstration of an operational distance learning system, and examples of successful systems integration projects.

MARCONI
Booth 864

MCIWORLD COM
Booth 833

MICROWAY, INC.
Booth 382
Ann Fried
Research Park Box 79; Kingston, MA 02364
Phone: 508-746-7341
Fax: 508-746-4678
Email: info@microway.com
URL: www.microway.com

Since 1982, Microway has designed high speed numerics products including transputer- and i860-based accelerator cards; compilers; and Alpha and Pentium workstations. Microway currently manufactures cost effective, powerful Linux Beowulf clusters based on Alpha, Intel and Athlon with optional Alteon, Myrinet, or Wulfkit interconnects. Our RuggedRak and Quadputer chassis yield the most compact packaging of motherboards and CPUs with optimal cooling. Microway is the largest provider of Alpha-based systems in North America following Compaq.

MORGAN KAUFMANN PUBLISHERS

Booth 480

Marketing Dept., Morgan Kaufmann Publishers
340 Pine Street 6th Fl.; San Francisco, CA 94104
Phone: 781-359-2405
Fax: 415-982-2665
Email: mkp@mkp.com
URL: www.mkp.com

Morgan Kaufmann publishes distinguished books for your high performance computing needs, including professional books, textbooks, and collected volumes. Two new titles include Parallel Programming in OpenMP by Rohit Chandra, Ramesh Menon, Leo Dagum, David Kohr, Dror Maydan, and Jeff McDonald. And Parallel I/O for High Performance Computing by John May.

MPI SOFTWARE TECHNOLOGY, INC

Booth 716

Amanda Henry
101 South Lafayette Street, Suite 33
Starkville, MS 39759 USA
Phone: 662-320-4300
Fax: 662-320-4301
Email: Amanda@mpi-software.com
URL: www.mpi-softtech.com

MPI Software Technology, Inc. (MSTI) provides "middleware" solutions for high-performance parallel computing. MSTI's products include MPI/Pro™, a commercial-grade implementation of MPI (message passing interface), Cluster Controller®, a job scheduler product for Windows clusters, and VSI/Pro®, a math library. MSTI also provides consulting and continuing education services at the customer's request.

MYRICOM, INC.

Booth 712

Nan Boden
325 N. Santa Anita Avenue
Arcadia, CA 91006
Phone: 626-821-5555
Fax: 626-821-5316
Email: nan@myricom
URL: www.myri.com

Myricom produces Myrinet, the leading interconnect for high-performance and high-availability clusters. Myrinet-2000 products, with 8-microsecond latency, 2+2 Gigabit/s full-duplex copper and fiber links, and unlimited scalability, are on display. Myrinet is supported on all popular platforms, and is available from major computer vendors and from a worldwide network of resellers.

NAG

Booth 565

NEC – HNSX SUPERCOMPUTERS INC.

Booth 532

Philip Tannenbaum
NEC – HNSX Supercomputers Inc.
4200 Research Forest Drive Suite 400
The Woodlands, Texas 77381-4257
Phone: 281-465-1610
Fax: 281-465-1699
Email: ptannenbaum@sx.nec.com
URL: www.hstc.necsyl.com

NEC–HNSX Supercomputers provides the NEC SX-5 Series, the most powerful parallel vector supercomputer available today. SX-5 Series systems incorporate 10 gigaflops individual processors to deliver up to 5 teraflops peak performance, up to 8 terabytes of main memory, and over 40 terabytes/sec of memory bandwidth. The booth will showcase the technologies necessary to produce true supercomputers, and will introduce a new range of scalable 16 processor SMP servers based on IA-64 architecture.

NETGEN LEARNING SYSTEMS

Venture Village E

NORTEL
Booth 980

OBJECTIVITY, INC.
Booth 843

Nancy Robbins
301B East Evelyn Ave; Mountain View, CA 94041
Phone: 650-254-7181
Fax: 650-254-7181
Email: nancy.robbins@objectivity.com
URL: www. Objectivity.com

Objectivity/DB is the leading high performance, distributed, and scalable database engine with unrivaled support for mixed language development and mixed hardware environments. Objectivity provides the ideal platform for mission critical applications requiring continuous performance and adaptablity to future technologies.

PALLAS
Booth 862

Werner Krotz-Vogel
Hermuelheimer Str. 10
Bruehl D-50321 Germany
Phone: +49 2232 1896-0
Fax: +49 2232 1896-29
Email: info@pallas.com
URL: www.pallas.com

Pallas is Europe's leading software company specializing in High Performance Computing. Pallas provides quality and portability, tools and services, around MPI, PVM, HPF, OpenMP. Vampir and Vampirtrace parallel profiling and visualization products represent state-of-the-art performance analysis tools. Pallas is a European distributor for The Portland Group compilers, Etnus TotalView debugger, KAI C++ and KAP/Pro-toolset.

PARABON COMPUTATION, INC.
Booth 786

Pat O'Boyle
3930 Walnut Street; Suite 100
Fairfax, VA 22030-4738
Phone: 703-460-4100
Fax: 703-460-4101
Email: sales@parabon.com
URL: www.parabon.com

Whether you're using a workstation, server, or supercomputer, Parabon delivers an easy, affordable path to more compute. Frontier™ draws on the vast, untapped power of Internet-connected computers, boosting computational power on demand. Launch mammoth jobs without the usual challenges of massive-scale distributed processing. And focus on what really counts —your application.

PARALOGIC,INC.
Booth 1076

George Palmer
130 Webster Street; Bethlehem, PA 18015
Phone: 610-814-2800 Ext 207
Fax: 610-814-5844
Email: info@plogic.com
URL: plogic.com

Paralogic engineers and supports turn-key commodity Linux supercomputer clusters (Beowulf). Our clusters are custom designed and optimized for for production computing environments. We offer Intel, AMD, and Compaq processors, 100/1000 Ethernet, Giganet, Myrinet, and storage solutions in 1U, 2U, and 4U rackmount packaging. Our software solutions include PGI Compilers, BERT 77 (Fortran conversion), MPI-PRO, PBS-PRO, GFS, DEEP-MPI, TotalView, and Linda. As an authorized reseller, we provide full manufacturer support for all software from the application to the Linux kernel. We also offer/sup-

port turn-key software applications in CFD, Computational Chemistry, Seismic Processing, and many others.

PLATFORM
Booth 380

QUADRICS
Booth 682
One Bridewell Street
Bristol
BS1 2AA
UK
Tel +44 117-907-5375
Fax: +44 117-907-5395
URL: www.quadrics.com

Quadrics is the leading European provider of supercomputer technology. The company specialise in hardware and software components for building scalable systems. The QSnet interconnect is a high performance commodity processor interconnect capable of connecting many 100s of processing nodes. Quadrics software provides solutions for managing supercomputer class systems.

QWEST COMMUNICATIONS
Booth 986
Wesley K. Kaplow
1306 Concourse Drive; Linthicum, MD 21090
Phone: 410-694-4803
Fax: 410-694-4898
Email: wesley.kaplow@qwest.com
URL: www.qwest.com

Qwest is a facilities-based leading-edge data network services provider. Using the latest in Dense Wavelength Division Multiplexed fiber-optic transport technology, Qwest provides high-performance IP, ATM, SONET, and Wavelength services. Via Qwest's CyberCenters®, Qwest provides secure on-backbone hosting space and operations for mission-critical applications and storage.

RACKSAVER
Booth 734
David Driggers
7100 Convoy Court; San Diego, CA 92111
Phone: 858-874-3800
Fax: 858-874-3838
Email: davidd@racksaver.com
URL: www.racksaver.com

RackSaver specializes in Custom and Turn-key "High Density" Beowulf Clusters utilizing AMD and/or Intel Processors. Our Interconnects include 10/100/1000 or Fiber LAN, our High-speed interconnects utilize Dolphin/Scali or Myrinet GigE. No cluster is too large or too small! We love custom design projects so please present us with your requirements!

RAYTHEON
Booth 762
Craig Stair
1200 S. Jupiter Road; Garland, TX 75042
Phone: 972-205-7677
Fax: 972-205-6144
Email: Craig_R_Stair@raytheon.com
URL: www.raytheon.com/igs

Raytheon is a recognized leader of advanced high performance computing. By configuring and integrating a mix of computation, network and storage products, Raytheon offers a comprehensive system-level simulation approach. Whether your task is achieving the most from commercial components or solving pioneering challenges, Raytheon has the experience and the expertise you need.

SC2000 WEBCASTS
Booth 886
Greg S. Johnson
University of California San Diego/SDSC
9500 Gilman Drive MC 0505
La Jolla, CA 92093-0505
Phone: 858-534-8367

Fax: 858-534-5152
Email: johnson@sdsc.edu
URL: www.sc2000.org/webcasts

This year, SC2000 will broadcast live over the Internet select presentations using the unique capabilities of SCinet 2000 to deliver high-quality audio and video over the world wide web. SC2000 will webcast a number of technical and exhibitor presentations, including plenary sessions, selected exhibitor demonstrations and other conference events. In addition—back by popular demand—a robotic camera will allow remote viewers to pan the exhibit hall and zoom in on exhibit floor booths.

SCALI AS / DOLPHIN INTERCONNECT SOLUTIONS

Booth 966
Jjur Fjellbirneland
Olaf Helsets vei 6
Postboks 70 Bogerud
Oslo, Norway N-0621
Phone: +47 22 62 89 54
Fax: +47 22 62 89 51
Email: sf@scali.com
URL: www.scali.com

SCIENTIFIC COMPUTING ASSOCIATES, INC.

Booth 470
Marilyn Riordan
One Century Tower; 265 Church Street
New Haven, CT 06510
Phone: 203-777-7442
Fax: 203-776-4074
Email: riordan@sca.com
URL: www.sca.com

Scientific Computing Associates, Inc. is the leading commercial provider of middleware products, technology, training and consulting services for high performance computing on supercomputers and clusters of PC's and workstations. Scientific Computing Associates helps clients deliver on the promise of parallel and distributed computing in fields like bioinformatics, finance, engineering analysis and more!

SCIENTIFIC COMPUTING & INSTRUMENTATION

Booth 579
Matt Lally
301 Gibraltar Drive
Morris Plains, NJ 07950
Phone: 973-292-5100
Fax: 973-539-3476
Email: mlally@cahners.com
URL: www.scimag.com

Scientific Computing & Instrumentation is the leading publication covering technical computing news, products, and applications for science and research. Visit our exhibit for a free copy of our magazine and to start your free subscription.

SCINET 2000

Booth 880
William Kramer
NERSC - MS 50B-4230
One Cyclotron Road
Berkeley, CA 94720
Phone: 510-486-7577
Fax: 510-486-4300
Email: scinet@sc2000.org
URL: www.sc2000.org/scinet

A significant part of the technology showcased at each year's SC conference is the SCinet network which supports the conference and makes it (for the duration of the show) one of the most intense networks on the planet. This year, SCinet 2000 is working with Qwest Communications International, Cisco, Nortel Networks, Juniper, Marconi, Foundry, Extreme, Spirent and others to establish flexible wide

area connections to the show floor. Using Qwest fiber in Dallas and Qwest SONET, ATM and IP backbones nationwide, the wide area network will feature multiple OC-48 (2.5 gigabit per second) connections, several OC-12 connections and possibly other connections, using the very latest technology and protocols. The total connectivity between SC2000 and the outside world will be 10.5 to 11 Gigabits per second—a new record for the SC conferences! In addition to commodity Internet access, WAN connection links to SCinet 2000 will include:

- ESnet
- Abilene/Internet 2
- HSCC
- vBNS
- NREN (NASA's Research & Education Network)

SCinet plans to deploy and support IPv4, IPv6, ATM, and Packet over SONET connections, Myrinet, Quality of Service demonstrations, and advanced network monitoring. Other types of connections might be possible based on discussions with requestors. SCinet will install and operate more than 40 miles of fiber optics throughout the conference areas, including the types of connections listed above.

SGI

Booth 550
Weston Bielby
1600 Amphitheatre Parkway; Mail Stop 846
Mountain View, CA 94043
Phone: 650- 960-1980
Fax: 650-933-0316
Email: west@engr.sgi.com
URL: www.sgi.com

Consistently acknowledged for leadership and innovation in technical and high performance computing, SGI technology is employed wherever unleashing creativity and insight is a priority. The newly-released family of SGI 3000 systems allows modular deployment of advanced computing and visualization solutions for technical enterprises and scientific, engineering and creative professionals throughout the world, enabling consistent accomplishment and keeping you "One Step Ahead." Discover the benefits of NUMAflex™ and our breakthrough advances in modular computing by visiting our wide range of HPC solutions in Booth 550.

SIAM—SOCIETY FOR INDUSTRIAL AND APPLIED MATHEMATICS

Booth Registration/Lobby Area
Michelle Montgomery
SIAM
Marketing Manager
3600 University City Science Center
Philadelphia, PA 19104
Phone: 215-382-9800 x368
Fax: 215-386-7999
Email: Montgomery@siam.org
URL: www.siam.org

Visit the SIAM booth to check out our new publications. There are new books in almost every SIAM series. Recently published books and information about our newest book series will be on hand. You will see sample journal issues of SIAM's renowned journals, plus membership information and applications will be available for those individuals interested in becoming a SIAM member. Don't forget to pick up a copy of SIAM News for the road.

SOFTWARE AG, INC.

Booth 581
Tai Dunson-Strane
2613 Camino Ramon, Ste. 110
San Ramon, CA 94583

Phone: 925-242-3784
Fax: 925-242-3800
Email: tai.dunson-strane@softwareag.com
URL: www.softwareagusa.com

TAMINO (Transaction Architecture for the Management of Internet Objects)
Tamino, from Software AG, is the world's first information server based on data storage and retrieval in a pure, standard XML format. More than just a new-generation database management system, Tamino adds advancement functionality to a new, fast, and internal XML-based DBMS. The result is a powerful integration engine that is ideal for mission-critical applications. The information server stores and retrieves self-describing XML information without conversion to other formats. With Tamino, enterprises have instant access to any data across business boundaries, using Internet technologies.

SPIRENT COMMUNICATIONS

Booth 982
26750 Agoura Road; Calabasas, CA 91403
Phone: 818-676-230
Fax: 818-676-2700
Email: info@spirentcom.com
URL: www.spirentcom.com

Spirent, the world leader in performance analysis solutions, manufactures the award winning products of Adtech, DLS, GSS, SmartBits and TAS product lines. Every time you access the Internet or corporate network, your data passes through systems that have been designed, manufactured or evaluated with Spirent Communications' industry-standard solutions.

SRC COMPUTERS, INC.

Booth 278
Mike Henesey/Patricia Miranda
4240 North Nevada Ave.
Colorado Springs, CO 80907
Phone: 719-262-0213
Fax: 719-262-0223
Email: michael.henesey@srccomp.com
URL: www.srccomp.com

SRC Computers, Inc. will have the prototype SRC-6 powered on in booth 278. See the machine and meet with key SRC personnel and partners as we prepare for the launch of "the next big thing" in computing architectures: Reconfigurable Computing. We are a start-up company with sound financial backing, an extraordinary technology roadmap crafted by Seymour Cray himself, and a team of dedicated professionals focused on turning the dream into reality. Join us as we MAP™the future of HPC. See you at the show!!

STARTECH BUSINESS DEVELOPMENT

Venture Village B

STORAGETEK

Booth 416
Patty Bank
One StorageTek Dr.; Louisville, CO 80028-4341
Phone: 303-661-2972
Fax: 303-661-5478
Email: murphpf@louisville.stortek.com
URL: www.storagetek.com

StorageTek is the only global information storage provider to use "Open, Intelligent and Integrated"™ solutions to promote the efficient use of critical information and to remove the information management constraints that inhibit business growth. Our competencies—storage automation, virtual storage and open SANs (storage area networks)—effectively manage the

explosion of information, save time and money, and maximize personnel resources.

STORNET, INC.

Booth 845
Steve Sweetin
1109 Saunders Ct.; West Chester, PA 19380
Phone: 610-692-8400
Fax: 610-692-8355
Email: info@stornet.com
URL: www.stornet.com

SUN MICROSYSTEMS, INC.

Booth 520
Rich Brueckner
901 San Antonio Rd.; Palo Alto, CA 94303
Phone: 800-555-9786
Fax: 800-277-7724
Email: rich.brueckner@west.sun.com
URL: www.sun.com

Since its inception in 1982, a singular vision, "The Network Is The Computer," has propelled Sun Microsystems, Inc., to its position as a leading provider of high performance hardware, software, and services, enabling network-centric supercomputing and expanding the power of the Internet. Sun is pleased to showcase innovation across the HPC community and demonstrate a breadth of solutions necessary for "Massive Scalability."

SYSKONNECT, INC.

Booth 863
Al Piscione
1922 Zanker Road
San Jose, CA 95112
Phone: 408-437-3800
Fax: 408-437-3866
Email: info@syskonnect.com
URL: www.syskonnect.com

SysKonnect is presenting the latest addition to their gigabit product family, the SK-NET GE-T adapters for CAT5 UTP cabling. Available now, the SK-NET GE-T 1000BaseT adapters are designed for 64bit/66MHz PCI bus systems and offer both single port and dual port fail-over LAN connectivity for crucial server connections. SysKonnect SK-NET GE NIC's are also available for multi-mode (SX), and single-mode (LX) fiber applications.

TEXAS MEMORY SYSTEMS

Booth 569
Holly Frost
11200 Westheimer, Suite 1000
Houston, Texas 77042
Phone: 713-266-3200
Fax: 713-266-0332
Email: hf@texmemsys.com
URL: www. superSSD.com

RAM-SAN
The SAM-500/520 (RAM-SAN) is a solid-state storage appliance which eliminates all limitations inherent to disk based SANs. Its internal bandwidth of 2.8 gigabytes/sec easily handles 15 Fibre Channels running at full speed with random accesses at 500,000 I/O sec. The RAM-SAN installs quickly, providing your users with instantaneous SAN access.

THE MIT PRESS

Booth 748
John Costello
5 Cambridge Center; 4th Floor; NE25-4062
Cambridge, MA 02142-1493
Phone: 617-258-5764
Fax: 617-253-1709
Email: jcostell@mit.edu
URL: http://mitpress.mit.edu

The MIT Press publishes cutting edge computer science books in core areas: programming languages and software, computer systems, artificial intelligence and intelligent systems, and theoretical computer science,

including books on specialized computer technologies, advanced topics in computer design, and innovative applications of computing. Our Scientific and Engineering Computation Series is of particular interest to the SC2000 audience. Please visit http://mitpress.mit.edu/computer-science/ for more information.

THE PORTLAND GROUP INC

Booth 754

Douglas Miles
9150 SW Pioneer Ct, Suite H
Wilsonville, OR 97070
Phone: 503-682-2806
Fax: 503-682-2637
Email: miles@pgroup.com
URL: www.pgroup.com

The Portland Group (PGI) provides optimizing HPF, F90, F77, C++ and C compilers for High Performance, parallel systems based on Intel (IA-32) Architectures. PGI also provides High Performance Fortran (HPF) for many of the large parallel systems currently available. PGHPF is a retargetable compiler for both shared and distributed-memory systems.

TIMES N SYSTEMS, INC.

Booth 1172

Paige Thurgood, Marketing Comm. Manager
1908 Kramer Lane, Bldg. B, Suite P
Austin, Texas 78758
Phone: 512-997-9800
Fax: 512-997-9805
Email: paiget@timesn.com
URL: www.timesn.com

Times N Systems' Teaming technology harnesses the computational power of industry-standard computers providing a scalable, expandable foundation to support today's fast-paced technological environment. We offer extreme performance advantages, alleviating systems bottlenecks and latency problems while delivering system-wide performance increases at a fraction of the costs.

TURBOLINUX, INC.

Booth 482

8000 Marina Boulevard, Suite 300
Brisbane, CA 94005
Phone: 650-228-5000
Fax: 650-228-5001
URL: www.turbolinux.com

TurboLinux is the leader in high-performance enterprise Linux, building reliable e-commerce, Internet infrastructure and clustering solutions. The first distribution to port to Intel's new IA64 processor, the TurboLinux OS is available factory-installed on IBM Netfinity servers. Headquartered near San Francisco, TurboLinux is the world's fastest-growing Linux provider, shipping more than four million units of our products globally.

UNICOMP, INC.

Booth 464

Walt Brainerd
7660 East Broadway Blvd, Suite 308
Tucson, AZ 85710
Phone: 520-298-7212
Fax: 520-298-7074
Email: walt@uni-comp.com
URL: www.uni-comp.com

Unicomp sells HPC software, books, test and benchmarking suites, and tools such as Foresys and Tarantella. We provide training in Linux, Fortran 90/95, HPF, Visual Fortran, MPI, and OpenMP. Unicomp develops courseware for specialized topics in HPC. We provide a code conversion service, upgrading and modernizing Fortran 77 codes.

UNITED DEVICES, INC.

Booth Venture Village G

information@ud.com
11675 Jollyville Rd, Suite 100
Austin, TX 78759
Phone: 512-331-6016
Fax: 512-331-6235
Email: information@ud.com
URL: www.ud.com

United Devices is an Internet distributed computing company that develops and manages the infrastructure required to aggregate idle computation, storage and bandwidth resources on the Internet. United Devices delivers this unprecedented power in the form of application services for a number of industries, including Bioinformatics, Web Testing and Monitoring, and Content Search and Management.

UNITREE SOFTWARE

Booth 644

Robert Clark
4 Crow Canyon Ct #110
San Ramon, CA 94583
Phone: 925-552-4990, 1-888-UNITREE
Fax: 925-552-4991
Email: sales@unitree.com
URL: www.unitree.com

UniTree Software solutions are the answer to managing all your online, nearline, offline and backup needs. UniTree offers an integrated software environment for managing data in large filesystems and databases. Our software is designed for data intensive applications, where performance, scalability, and long-term data integrity are critical requirements.

VERIDIAN

Booth 782

James Patton Jones
2672 Bay Shore Parkway, Suite 810
Mountain View, CA 94043
Phone: 650-967-4675
Email: jjones@pbspro.com
URL: www.veridian.com

Veridian—a 20–year provider for trusted solutions in high performance computing industry.

PBS—supercomputing workload management tool, and de facto standard resource management system for linux clusters.

VAST—automatic parallelization /vectorization for shared and distributed memory systems.

DEEP—development environment for profiling/analysis of HPC programs for message passing, OpenMP and node level cache performance.

VIEO, INC.

Booth 1056

Steve Harriman
12416 Hymeadow Drive, Suite 200
Austin, TX 78750
Phone: 512-257-3031
Fax: 512-257-3101
Email: sharriman@vieo.com
URL: www.vieo.com

VIEO, Inc., formerly Power Micro Research, was established in 1994 to develop high performance I/O solutions for a range of customers including the world's major system platform vendors. VIEO is partnering with leading high technology companies to deliver products that enable, integrate and manage InfiniBand‘ I/O fabrics.

ViON CORPORATION AND HITACHI DATA SYSTEMS

Booth 178 and 179

Cindy Sheets
1055 Thomas Jefferson St., NW, Suite 406
Washington, DC 20007
Phone: 202-467-5500
Fax: 202-342-1404
Email: sheets@vion.com
URL: www.vion.com or www.hds.com

ViON Corporation, established in 1980, is a highly successful provider of information technology specializing in storage and Storage Area Networks in Business-to-Government. We offer solutions for both the high-end enterprise and departmental systems. ViON has partnerships with Hitachi Data Systems, Brocade, Crossroads, McData, and Veritas to supply these solutions for your storage/SAN.

WOLFRAM RESEARCH

Booth 638

Kristin Kummer
100 Trade Center Drive
Champaign, IL 61820
Phone: 217-398-0700
Fax: 217-398-0747
Email: kristink@wolfram.com
URL: wolfram.com

Mathematica is the world's only fully integrated technical computing system, combining unrivaled computational power with unprecedented ease of use. Used by over a million scientists, engineers, programmers, and students. Unmatched ease of computation (numeric and symbolic), data visualization, and programming. Fully customizable, cross-platform, publication-quality electronic and printed documents.

YOTTAYOTTA, INC.

Booth 1066

Deanna Hagerich
4030 Lake Washington Blvd NE, Suite 201
Kirkland, WA 98033
Phone: 425-739-9191 x228
Fax: 425-739-0115
Email: dhagerich@yottayotta.com
URL: www.yottayotta.com

YottaYotta, Inc., the Yottabyte NetStorage Company, is the only company to deliver a truly scalable storage solution for the next generation of network-based corporations in the major WAN, LAN, SAN, PAN, BAN and M2AN markets. YottaYotta's NetStorage Cube increases data storage as well as performance and reliability.

SC2000
DALLAS CONVENTION CENTER | NOVEMBER 4 – 10
HALLS D & E

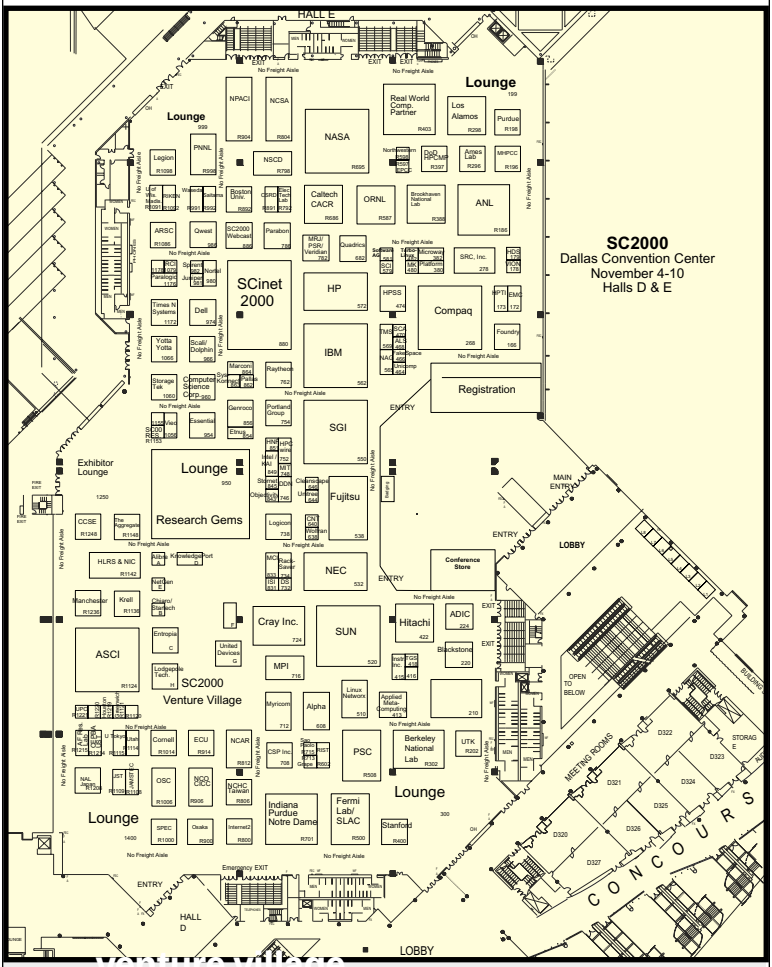

venture village
research gems
research exhibits
industry exhibits
scinet 2000
registration

PROCEEDINGS/CD-ROM

Attendees registered for the technical program will receive a ticket that they may use to redeem a copy of the SC2000 Proceedings CD-ROM at the Conference Store. Additional copies of the Proceedings will be available for sale at the Conference Store for $50.00 each.

REGISTRATION

Exhibit Hall, 214-853-8001

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," and cash. U.S. Federal Government Purchase Orders and wire transfers will not be accepted at on-site registration.

Member, Retired Member and Student Discounts

To qualify for discounted registration rates, include your current ACM, ACM SIGARCH, IEEE, or IEEE Computer Society membership number(s). Students are required to show a copy of a valid full-time student identification card or a letter from the university confirming full-time student status at the time of registration.

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. This letter should be delivered to the Special Assistance desk in the registration area.

Lost Badge Fee

SC2000 will charge a $40 processing fee to replace badges for those attendees who lose their badge during the conference.

Registration Needs

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

REGISTRATION AND STORE HOURS

Saturday, Nov. 4	1:00PM—5:00PM
Sunday, Nov. 5	7:30AM—6:00PM
Monday, Nov. 6	7:30AM—9:00PM
Tuesday, Nov. 7	7:30AM—5:30PM
Wednesday, Nov. 8	7:30AM—5:30PM
Thursday, Nov. 9	7:30AM—5:00PM
Friday, Nov. 10	8:00AM—11:00AM

TUTORIALS

Full-day tutorials are being offered on Sunday and Monday, November 5th and 6th. These 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 tutorials offered on that day, a tutorial 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 both days, and your selection of notes as you specified on the registration form. An additional benefit of the Tutorials Passport is that you can purchase additional tutorial notes at a significantly reduced rate.

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

Tutorial notes can be picked up at the SC2000 Conference Store upon presentation of the notes ticket(s) provided in your registration packet. You will only receive notes for the tutorial(s) you selected on your registration form. Extra copies of tutorial course notes will be on sale at the Conference Store starting Wednesday, November 8. Orders for additional tutorial notes can be placed at any time in the Conference Store for delivery after the conference.

EXHIBITS ONLY BADGES

SC2000 is offering Exhibits Only registration to children 12 to 16 who must purchase an Exhibits Only Badge ($80 a day) and be accompanied by an adult at all times while on the exhibit floor. The SC2000 Planning Committee is pleased to welcome family members of SC2000 attendees to the exhibit floor on Wednesday, November 8 from 4PM to 6PM. Spouses, children, partners and other interested immediate family members of badged attendees will be admitted free of charge during this period. See additional information on Family Members Day on page 179. Children under age 12 are not permitted on the exhibit floor at any time other than during the Family Members Day hours of Wednesday, November 8 from 4PM to 6PM.

MEDIA ROOM

ROOM D 326, 214-743-2696

The SC2000 Media Room will be located in Room D 326 near the main lobby of the Dallas Convention Center.All media representatives are asked to check in at the Media Room upon arrival. If you have preregistered by Friday, Oct. 20, you can pick up your registration in the Media Room. If you are planning to register at the conference, you'll need to have your credentials verified at the Media Room before your registration can be completed. See the Press Registration Guidelines on the SC2000 Web site for credentialing information.

The SC2000 Media Room will provide:

- Telephones—the main Media Room phone number is 214-743-2696
- Computers (both Mac and PC) with Internet connections for email and Web access;
- A printer
- Exhibitor media kits
- On-line exhibitor information
- Message board
- Calendar of events showing daily highlights
- Refreshments

The Media Room is available (see hours below) to help you with facilities, schedule interviews, or any other activities to assist your coverage of SC2000. Stop by during Media Room hours or contact Media Relations Coordinator Jon Bashor at JBashor@lbl.gov.

Media room hours are:

Saturday, Nov. 4	3PM – 5PM
Sunday, Nov. 5	1PM – 4PM
Monday, Nov. 6	9AM – 6PM
Tuesday, Nov. 7	9AM – 6PM
Wednesday, Nov. 8	9AM – 6PM
Thursday, Nov. 9	9AM – 5PM

BUSINESS CENTER SERVICES

214-741-9090; 214-743-2494 fax

The Dallas Convention Center Business Center/Gift Shop is located in Lobby C, lower level. Hours of operation are as follows:

- **Sunday, Nov. 5** 10AM – 2PM
- **Monday, Nov. 6–Thursday, Nov. 9** 9AM – 5PM
- **Friday, Nov. 10** 9AM – 2PM

CONFERENCE HOTELS AND DART MAP

	SC2000 HOTELS	SINGLE	DOUBLE	TRIPLE	QUAD
1	**Adams Mark Hotel** 400 North Olive 214-922-8000	$140	$150	NA	NA
2	**Adolphus Hotel** 1321 Commerce Street 214-742-8200	$156	$176	NA	NA
3	**Fairmont Hotel** 1717 N Akard Street 214-720-2020	$155	$155	NA	NA
4	**Hampton Inn West End** 1015 Elm Street 214-742-5678	$119	$119	$129	$129
5	**Hyatt Regency Dallas** 300 Reunion Blvd *(near Convention Ctr)* 214-651-1234	$155	$175	NA	NA
6	**Ramada Plaza Hotel–DCC** 1011 Akard Street *(near Convention Ctr)* 214-421-1083	$119	$129	NA	NA

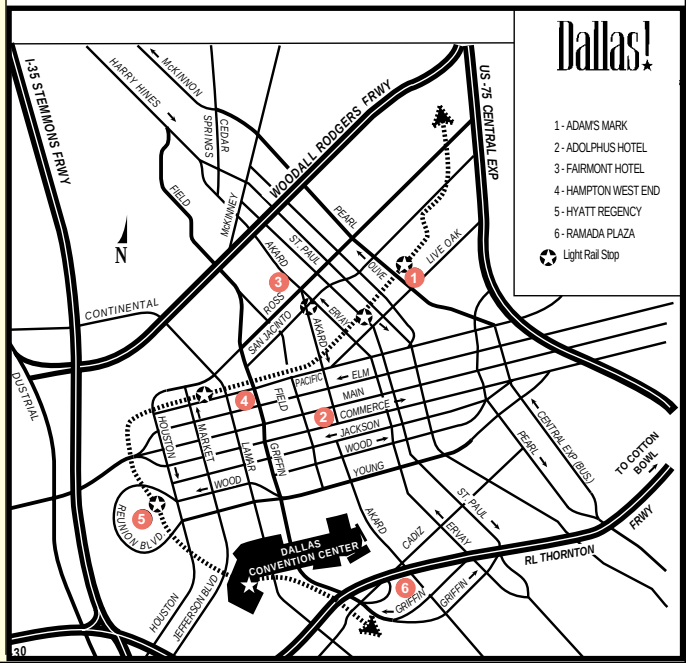

FAMILY MEMBERS DAY | WEDNESDAY, NOV. 8

4PM – 6PM

The SC2000 Planning Committee is pleased to welcome family members of SC2000 attendees to the exhibit floor on Wednesday, November 8th from 4 PM to 6 PM. Spouses, children, partners and other interested immediate family members of badged attendees will be admitted free of charge during this period.

No badge is required for these guests. The following guidelines have been established for this policy, which is exclusive to the SC2000 Conference and will not be repeated in future SC conferences.

- Guests must be accompanied at all times by a badged participant
- Every two children under 16 years of age must be accompanied by an adult at all times
- SC2000 (and/or its agents) will not be responsible or liable for any damages resulting from the misconduct of guests
- Guest sponsors (badged exhibit participant) will be responsible for the conduct of their guests.

E-MAIL/MESSAGE CENTER

ROOM D 321, 214-743-2686

Sun workstations with Internet access for e-mail have been provided courtesy of Sun Microsystems, Inc for conference attendees at the convention center in room D-321. In addition, there will be 100BaseT connections available for attendees who would like to use their own laptop computers for accessing e-mail.

E-mail Room Hours

- **Sunday, Nov. 5** 1PM – 6PM
- **Monday, Nov. 6–Thursday, Nov. 9** 8AM – 7PM
- **Friday, Nov. 10** 9AM – 12NN

DART

Dallas Area Rapid Transit (DART) provides a light rail system which connects the downtown area and conference hotels with the Dallas Convention Center. Please see map on page 178 for dropoff locations.

- Operating hours: Seven days a week— 5:30AM – 12:30AM
- Trip Frequency: Every 5-10 minutes during peak drive times—6 – 9AM, 3 – 6PM
- Every 10-20 minutes during off-peak drive times— 5:30 – 6AM, 9AM – 3PM, 6PM – 12:30AM

DALLAS CONVENTION CENTER | LEVEL 2

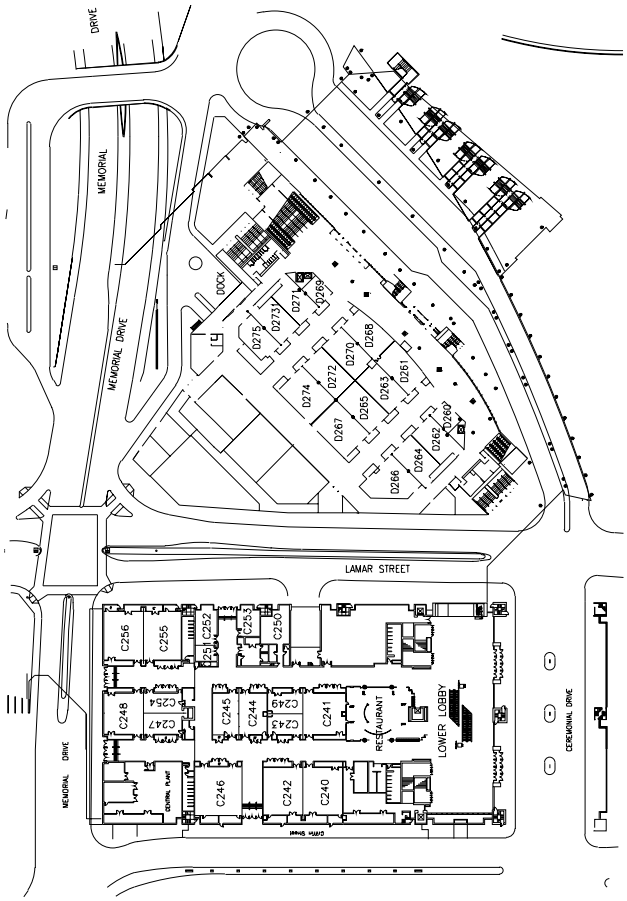

DALLAS CONVENTION CENTER | LEVEL 3

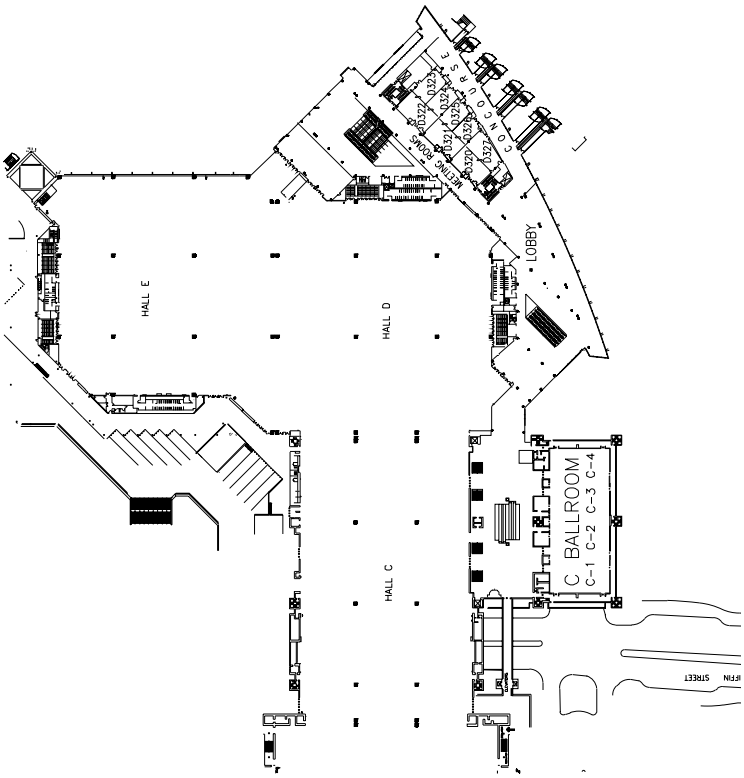

COAT AND BAG CHECK

A coat and bag check area will be provided for conference attendees beginning on Sunday morning and operating every day through 2PM on Friday.

Coat and Bag Check Hours

- **Sunday, Nov. 5** 7:45AM – 6PM
- **Monday, Nov. 6** 7:45AM – 9:15PM
- **Tuesday, Nov. 7–Thursday, Nov. 9** 7:45AM – 6PM
- **Friday, Nov. 10** 7:45AM – 2PM

FIRST AID

First aid stations are located in the Dallas Convention Center as follows:

- **Arena First Aid Station**—located on the fourth level of the arena: 214-939-2780 or 2781
- **C Hall First Aid Station**—located on the lower C lobby next to Room C 250: 214-939-2715
- **E Hall First Aid Station**—located inside E Hall on the back wall: 214-939-2953 or 2954

The community which comes together to put on the annual SC conference is strong, resourceful, innovative, and definitely the best collection of researchers, IT workers, administrators, and support staff imaginable. The SC2000 committee gratefully recognizes the hours of dedication and significant resources that were brought to bear to make SC2000 the best in this conference series. We also thank the organizations listed here for their support of the individuals who gave so freely of their talents to SC2000.

SC2000 thank you

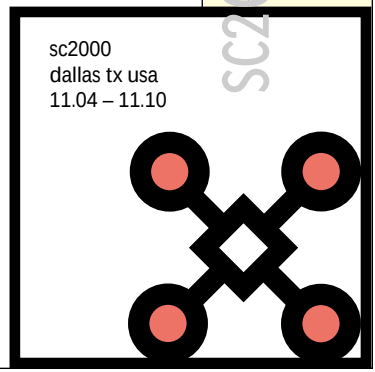

GENERAL CHAIR
Louis Turcotte, *Rose-Hulman Institute of Technology*

GENERAL VICE CHAIR
Betsy Schermerhorn, *Fermi National Accelerator Laboratory*

GENERAL DEPUTY CHAIR
Charlie Slocomb, *Los Alamos National Laboratory*

CONFERENCE VICE CHAIR, FINANCE
Evelyn Turcotte

COMMITTEE RELATIONS CHAIR
Wendy Duke

IEEE COMPUTER SOCIETY REPRESENTATIVES
Anne Marie Kelly and Mary-Kate Rada

ACM REPRESENTATIVE
Alisa Rivkin

SC2001 LIAISON
Debby Thompson, *Los Alamos National Laboratory*

CONFERENCE ARRANGEMENTS
Conference Vice Chair, Mary Amiot, *Northstar Event Management*

Conference Arrangements Chairs

Registration Chair, Tony Baylis, *National Center for Supercomputing Applications*
Kristine Weiand, *IEEE Computer Society*
Sheryl Reeder, *NCSA*
Conference Store Mgr, Gloria Montoya-Rivera, *Los Alamos National Laboratory*
Corinne Fresquez, *Los Alamos National Lab*
Dale Reeder, *NCSA*
Jean Soliday, *NCSA*
Annabelle Valdez, *Los Alamos National Lab*
Michele Bianchini-Gunn, *Lawrence Livermore National Laboratory*
Patrice Turcotte, *Alltel*
Conference Office Manager, Christie Salazar, *Los Alamos National Laboratory*
Bridget Borchers
Dolores Roybal, *Los Alamos National Laboratory*
Catering, Jan Hull, *Los Alamos National Lab*
Heidi Peterson, *Lawrence Livermore National Lab*
Housing, Debbie Huszar, *Task Handlers, Inc.*

CONFERENCE INFRASTRUCTURE
Conference Vice Chair, John Cobb, *Oak Ridge National Laboratory*

Conference Infrastructure Chairs

Equipment Rental Chair, Space Czar, John Cobb, *Oak Ridge National Laboratory*
Conference Equipment & PCs Chair, Robert R. Borchers, *National Science Foundation*
Signage Chair, Becky J. Verastegui, *Oak Ridge National Laboratory*
Email/Web Room Chair, Chip Watson, *Jefferson Lab*
A/V Chair, James "Trey" White, *Oak Ridge National Laboratory*

EDUCATION
Conference Vice Chair, Jeff Huskamp, *East Carolina University*

Education Chairs

Program Co-Chair, Lisa Bievenue, *National Center for Supercomputing Applications*
Program Co-Chair, Edna Gentry, *University of Alabama, Huntsville*
Technology Czar, Richard Bach, *Merck and Co.*
Publications Chair, Ginger Caldwell, *NCAR*
Systems Integration Chair, Barb Helland, *Krell Institute*
Research Exhibit Booth Chair, Scott Lathrop, *National Center for Supercomputing Applications*
Student Volunteer Chair, Bruce Loftis, *National Center for Supercomputing Applications*

Education Committee

Ernie Marshburn, *East Carolina University*
Helen Parke, *East Carolina University*
Ann Redelfs, *San Diego Supercomputer Center*
Ken Flurchick, *Ohio Supercomputer Center*
Robert Gotwals, Robert Panoff, *Shodor Education Foundation*
Pam Joop, *NCSA*
Marcie Tolley, *North Carolina Super.Center*
Wendy Creasey, *East Carolina University*
Donna Pettersen, *NCAR*

Advisory Committee

Norma Boakes, *Oakcrest H.S*
Jason Painter, *Farmville Central H.S.*
Eric Pitcher, *Silicon Graphics, Inc.*
Kris Stewart, *San Diego Supercomputer Center*
Paul Tisdel, *SATEC*
Chuck Swanson, *University of Minnesota*
Tom Lennon, *SGI*

EXHIBITS

Conference Vice Chair, Joe McCaffrey, *Mississippi State University*

Exhibits Chairs

Research Exhibits Chair, Ginny To, *HPTI*
Dev Prabhu, *ARL*
Andrew Mark, *ARL*
Exhibitor Outreach Chair, Ginger Caldwell, *NCAR*
Exhibition Management, Don Collier, *DC Expositions, Inc.*
Show Services, Darryl Monahan, *Freeman Decorating*
Security, Elizabeth Bogard, *Hewlett Packard*

INFORMATION ARCHITECTURE

Conference Vice Chair, Bill Kramer, *Lawrence Berkeley National Laboratory/NERSC*

Information Architecture Chairs

SC2001 SCinet Chair, Steve Tenbrink, *Los Alamos National Laboratory*
Deputy, Timothy Toole, *Sandia National Labs*
Experimental Network Chair, William "Bill" Wing, *Oak Ridge National Laboratory*
Production, Chuck Fisher, *Oak Ridge National Lab*
Physical Infrastructure, Jeff Mauth, *Pacific Northwest Lab*
Network Mgmt/Monitoring, Martin Swany, *University of Tennessee, Knoxville*
Info Manage/Customer Support, Doug Luce, *Aaronsen Group*
Committee Networking, Rex Duncan, *Oak Ridge National Laboratory*
Network Security, Eli Dart, *Sandia National Labs*
App Evangelist, Ian Foster, *Argonne National Lab*

Information Architecture Committee

Adam Duke, *FSU*
Jim Rogers, *CSC/Nichols*
David Crowe, *Oregon State University*
Bill Nickless, *Argonne National Laboratory*
Jim Ross, *Sandia National Laboratories*
Rick Maurer, *Sandia National Laboratories*
Warren Birch, *Army Research Lab*
Tom Dunlop, *DCC*
Dave Koester, *MITRE Corp*
Tom Kile, *Army Research Lab*
John Dugan, *NCSA*
Ed Kempe, *Dallas Visitor's Bureau*
Paul Love, *Internet2*
Larry Flourney, *Internet2, Texas A&M University*
Bill Lennon, *Lawrence Livermore National Lab*
Wendy Huntoon, *Pittsburgh Supercomputing Center*
Jalal Haddad, *Oregon State University*
Steve Corbato, *Internet2 UCAID*
Hal Edwards, *Nortel Networks*
Paul Daspit, *Nortel Networks*
Ben Peek, *GST Telecom*
Kevin Oberman, *Lawrence Berkeley National Lab*
John West, *WorldCom*
Steve Jones, *CEWES*
George Miller, *MCI*
Larry Dunn, *CISCO*
Greg Almes, *Internet2*
James Patton, *Caltech*
John Jamison, *Juniper Networks*
Bryan Bodker
Howard Walter, *Lawrence Berkeley National Lab*
Roberta Bourcher, *Lawrence Berkeley National Lab*
Jason Hasse, *CISCO*
David Wheeler, *NCSA*
Linda Winkler, *Argonne National Laboratory*
Pascal Boudreau, *Terabeam Networks*
Philip Schreher, *Qwest*
Robert Spenser, *Qwest*
Zaid Albanna, *MCI*
Glen Smith, *Qwest*
Wesley K. Kaplow, *Qwest*
Greg Goddard, *University of Florida*
Willard Ostrander, *TeraBeam Networks*
Janae Tinsley, *Smart City Networks*
Stacy Eubanks, *DCC*
Patrick Dorn, *NCSA*

PUBLIC RELATIONS

Conference Vice Chair, Kate Carey, *Ohio Learning Network*

Public Relations Chairs

Webmaster, Blaise Barney, *Lawrence Livermore National Laboratory*
Media Relations Chair, Jon Bashor, *Lawrence Berkeley National Laboratory*
Publications Chair, L.J. Evans, *Arctic Region Supercomputing Center*

Public Relations Committee

Webcasts Chair, Greg Johnson, *San Diego Supercomputer Center*
Webcasting, David Barber, *OhioLINK*
David Hart, *San Diego Supercomputer Center*
Harry Ammons, *SDSC/UCSD*
Mark Davenport, *RP Interactive*
Betsy Toole
Graphic Designer, Kim Iles, *Ilesnet Design*

TECHNICAL PROGRAM

Conference Vice Chair, Jim McGraw, *Lawrence Livermore National Laboratory*

Technical Program Chairs

Proceedings Chair, Jed Donnelley, *Lawrence Berkeley National Laboratory*
Proceedings Backup, Tom DeBoni, *Lawrence Berkeley National Laboratory*
Information Coordinator, Steve Elbert, *Entropia*
Awards Chair, Sally Haerer, *Oregon State University*
Tutorials Co-Chair, Michelle Hribar, *Pacific Univ.*
eSCape Co-Chair, Steve Jones, *U.S. Army ERDC*
HPC Games Co-Chair, James Arthur Kohl, *Oak Ridge National Laboratory*
Gordon Bell Award Chair, Ewing "Rusty" Lusk, *Argonne National Laboratory*
BoFs Assistant Chair, John May, *Lawrence Livermore National Laboratory*
Papers Chair, Rod Oldehoeft, *Los Alamos National Laboratory*
Panels Chair, Philip Papadopoulos, *San Diego Supercomputer Center*
Research Gems Chair, Jennifer Schopf, *Northwestern University*
HPC Games Co-Chair, Eleanor Schroeder, *NAVO*
Tutorial Co-Chair, Valerie Taylor, *Northwestern Univ.*
Invited Talks Chair, Pat Teller, *Univ. of Texas, El Paso*
BOFs Chair, Rajeev Thakur, *Argonne National Laboratory*
eSCape Co-Chair, John West, *WorldCom*

Papers Sub-Chairs

Apps/Numerics/Visualization, Randy Bramley, *Indiana Univ.*
Arch/Networks/Dist. Computing, Bob Lucas, *Lawrence Berkeley National Laboratory*
Performance Optimizaiton, Dan Reed, *University of Illinois*
Prog. Models and Tools, Ron Perrott, *Queens University, Belfast*

Technical Papers Committee

Abdullah Almojel, *King Fahd University*
Gail Alverson, *Cray Inc.*
Martin Audet, *National Res Council of Can.*
Robbie Babb, *University of Denver*
Fran Berman, *University of California, San Diego*
Michael Berry, *University of Tennessee*
Bill Carlson, *IDA Center for CS*
Barbara Chapman, *University of Houston*
Helen Chen, *Sandia National Laboratories*
Pat Crossno, *Sandia National Laboratories*
David Culler, *University of California, Berkeley*
George Cybenko, *Dartmouth College*
Frederica Darema, *National Science Foundation*
Karsten Decker, *Swiss Center for Sci Comp.*
James Demmel, *University of California, Berkeley*
Thomas Fahringer, *University of Vienna*
Wu-Chun Feng, *Los Alamos National Laboratory*
Kyle Gallivan, *Florida State University*
Dennis Gannon, *University of Indiana*
Guang Gao, *University of Delaware*
Al Geist, *Oak Ridge National Laboratory*
Wolfgang Gentzsch, *Genias Software*
Andrew Grimshaw, *University of Virginia*
Mary Hall, *Info Sciences Institute*
Michael Heroux, *Sandia National Laboratories*
Jeff Hollingsworth, *University of Maryland*
Bill Johnston, *Lawrence Berkeley National Lab*
Steve Karmesin, *Turbolabs*
Ricky Kendall, *Ames Laboratory*
Carl Kesselman, *Info Sciences Institute*
David Keyes, *Old Dominion University*
Peter Kogge, *University of Notre Dame*
Kai Li, *Princeton University*
Tara Madyastha, *University of California, SC*
Al Malony, *University of Oregon*
John Mellor-Crummey, *Rice University*
Kenichi Miura, *Fujitsu Corporation*
Bernd Mohr, *Research Centre Juelich*
Yoichi Muraoka, *Waseda University*
Hillarie Orman, *Novell*
Jean-Louis Pazat, *IRISA*
Padma Raghavan, *University of Tennessee*
John Rice, *Purdue University*
Diane Rover, *Michigan State University*
Joel Saltz, *University of Maryland*
Dave Schneider, *Cornell University*
Karsten Schwan, *Georgia Institute of Tech*
Karl Solchenbach, *Pallas GmbH*
Mark Stalzer, *HRL Laboratories*
Thomas Sterling, *California Institute of Tech*
Brian Tierney, *Lawrence Berkeley National Lab*
Donald Truhlar, *University of Minnesota*
Sam Uselton, *Lawrence Livermore National Lab*
David Walker, *Oak Ridge National Laboratory*
Steve Wallach, *CenterPoint Ventures/Chiaro Networks*
Harvey Wasserman, *Los Alamos National Lab*
Harry Wijshoff, *Leiden University*
Rich Wolski, *University of Tennessee*
Katherine Yelick, *University of California, Berkeley*

Tutorials Committee
Thomas Canfield, *Argonne National Laboratory*
Bradley M. Comes, *U.S. Army, MSRC*
John Grosh, *DoD HPCMP*
Mary Beth Hribar, *Tera Computer Co.*
Alice Koniges, *Lawrence Livermore National Laboratory*
Joe Mambretti, *Northwestern University*
Armin R. Mikler, *University of North Texas*
Reagan Moore, *San Diego Supercomputer Center*
Krish Namboodiri, *National Coordination Office*
John Sopka, *Sun Microsystems*
George K. Thiruvathukal, *DePaul University*
Mateo Valero, *University In Spain*
Michael Warren, *Los Alamos National Laboratory*
Jerry Yan, *Nasa Ames*
Pen-Chung Yew, *University of MN*
Mary Zosel, *Lawrence Livermore National Laboratory*

VENTURE VILLAGE
Conference Vice Chair, Dennis Duke, *Florida State University*
Conference Vice Chair,Steve Wallach, *CenterPoint Ventures/Chiaro Networks*

KEYNOTE ADVISORS
Al Brenner, *IDA*
Dave Cooper, *Lawrence Livermore National Laboratory*
Dennis Duke, *Florida State University*
Betsy Schermerhorn, *Fermi National Accelerator Laboratory*
Larry Smarr, *University of California, San Diego/Entropia, Inc.*
Joe Thompson, *Mississippi State University*

SCXY INDUSTRY ADVISORY COMMITTEE
Tom Tabor, *HPCwire*
Larry Dunn, *CISCO*
Debra Goldfarb, *IDC*
Sam Adams, *NEC*
Ty Rabe, *Compaq/DEC*
Kathy Sowards, *Hewlett Packard*
Bill DeSimone, *Compaq/DEC*
Chris Willard, *IDC*
Ellen Roder, *Cray Inc.*
Mike Bernhardt, *The Bernhardt Agency*
George Funk, *Dell*
Jim Padian, *IBM*
Rich Brueckner, *SUN*
Steve Joachims, *Hewlett Packard*
Fabio Gallo, *SGI*

SCXY STEERING COMMITTEE
Term Expires 2000:
Beverly Clayton, *Pittsburgh Supercomputing Center*
Cherri Pancake, *Oregon State University*
Margaret Simmons, *UCSD/SDSC*
Term Expires 2001:
David Culler, *University of California, Berkeley*
Louis Turcotte, *Rose-Hulman Institute of Technology*
Steve Wallach, *CenterPoint Ventures/Chiaro Networks*
Term Expires 2002:
David Cooper, *Lawrence Livermore National Laboratory*
Sally Haerer, *Oregon State University*
Charles Slocomb, *Los Alamos National Laboratory*
Term Expires 2003:
Jeff Huskamp, *East Carolina University*
Member Emeritus:
George Michael, *Lawrence Livermore National Laboratory*
Ex-Officio Members:
Anne Marie Kelly, *IEEE Computer Society*
Mary-Kate Rada, *IEEE Computer Society*
Alisa Rivkin, *ACM*

www.acm.org/sigarch/ join today!

SIGARCH & acm

www.acm.org

The ACM Special Interest Group on Computer Architecture serves a unique community of computer professionals working on the forefront of computer design in both industry and academia. It is ACM's primary forum for interchange of ideas about tomorrow's hardware. The newsletter *Computer Architecture News* publishes 5 times yearly and includes the ISCA and ASPLOS proceedings. SIGARCH also offers Member Plus proceedings benefits to members.

The Association for Computing Machinery (ACM) is a not for profit educational and scientific computing society. Benefits include a subscription to *Communications of the ACM*, *Membernet*, discounts on books and conferences, and the option to subscribe to the ACM Digital Library (www.acm.org/dl).

		Expedited Air Rates for outside North America
o	SIGARCH (ACM Member) . . . . . $ 28*	
o	SIGARCH (ACM Student Member) . . . . . $ 14*	
o	SIGARCH (Non-ACM Member) . . . . . $ 28*	
o	ACM (Professional Member $95) & SIGARCH($28) $123	
o	ACM (Professional Member $95) & SIGARCH($28) + ACM Digital Library ($90) $213	
o	ACM (Student Member $38) & SIGARCH ($14) . . . . . $ 52	
o	Computer Architecture News only . . . . . $ 55	
o	Expedited Air (outside N. America) . . . . . $ 25	
o	Expedited Air for Communications (outside N. America) . . . . . $ 37	
	SIGARCH Member Plus Package (Must be a SIGARCH Member)**	
o	SPAA: ACM Symposium on Parallel Algorithms and Architectures (July) . . . . . $ 22	$ 20
o	International Conference on Supercomputing (July) . . . . . $ 23	$ 20
o	Supercomputing Conference (Nov.) . . . . . $ 39	$ 20

*Annual subscription cost of $19.95 is included in member dues of $28 (for students cost is included in $14).
**International members please contact ACM Member Services for expedited air prices.

Payment Information

Name _____

ACM Member # _____

Mailing Address _____

_____

City/State/Province _____

Country/ZIP/Postal Code _____

Email _____

Phone _____

Fax _____

Credit Card: o AMEX o VISA o MC

Credit Card # _____

Exp. Date _____

Signature _____

Make check payable to ACM, Inc.

In addition to U.S. Dollars, ACM accepts bank check and Eurocheque payments in several European currencies. For information on currencies accepted and conversion rates, see contact information below.)Prices include surface delivery charge. Expedited Air Service, which is a partial air freight delivery service, is available outside North America. Contact ACM for further information.

MAILING LIST RESTRICTION

ACM occasionally makes its mailing list available to computer related orgainzations, educational institutions and sister societies. All email addresses remain strictly confidential. Check one of the following if you wish to restrict the use of your name

- O ACM announcements only (1)
- O ACM and other sister society announcements (2)
- O ACM subscription and renewal notices only (3)

ACM Headquarters
1515 Broadway
New York NY 10036
voice: 212-626-0500
fax: 212-944-1318
email: acmhelp@acm.org

Remit to:
ACM
PO Box 11315
New York, NY 10286-1315

S002N

www.acm.org/sigarch/ join today!

Denver, Colorado • November 10 – 16, 2001

beyond

SC GOES GLOBAL High-performance computing—scientific computing, networking, and communications—are fundamental drivers for today's computing revolution in e-commerce, simulation-based design, and data mining and more. For 13 years the SC conference has been the place where industry and academia meet to discuss trends and developments in these important areas. With a first-class technical program that includes a spectrum of tutorials and workshops, and a unique show floor—combining both industry exhibits and research exhibits involving many of the world's leading laboratories and universities—SC2001 remains the place to taste what is new and glimpse the future.

boundaries

The conference structure breaks down the usual barriers between here and there. The SC2001 "Beyond Boundaries" conference will link the Denver Convention Complex with constellation sites across the U.S.—and worldwide—using advanced Access Grid collaboration technology. This becomes the first truly global technical conference: a multi-national and multi-cultural meeting place for communication and discussion of high-performance computing and communications ideas, and their impacts on science and society.

SC2001

scientific computing, advanced networking, and communications

Conference: Nov. 10 – 16 • Exhibition: Nov. 12 – 15

Plan to participate in SC2001, the conference for scientific computing, networking and communications. Featured activities at the conference include • MasterWorks Presentations • Industry Exhibit and Exhibitor Forum • Technical Papers • Tutorials and Workshops • Invited Speakers and Panels •• State-of-the-Field Speakers • Research Exhibits and Posters • Education Program • Birds-of-a-Feather Sessions • SCinet • HPC Games

beyond

The exciting MasterWorks Presentations highlight novel, innovative uses and practices of advanced computing and communications in solving challenging, real-world problems. Learn how the latest waves in technology are impacting the larger scientific community and even the general public.

Commercial Exhibits highlight the industry's cutting-edge research, products, services, and developments. Anticipating the future, far-reaching concepts and technologies are presented and demonstrated in the Research Exhibits.

SC2001 will take place in the Denver Convention Complex, which has a high-speed Internet and fiber optic infrastructure routed to more than 600,000 square feet of convention space including meeting rooms, ballrooms, and exhibition space.

boundaries

SC2001 INFORMATION
www.sc2001.org
info@sc01.org

SC2001 EXHIBITOR INFORMATION
www.sc01.org/exhibits
industry@sc01.org
dc@dcexpo.com

DENVER INFORMATION
www.denver.com
www.denver.org
www.downtowndenver.citysearch.com
www.denverconvention.com

SC2001 CONFERENCE CHAIR
Charles Slocomb
Los Alamos National Laboratory

SC2001

Sponsored by ACM SIGARCH and IEEE Computer Society