
SC2003
Supercomputing Conference
NASA
Visit booth 411 to explore some of NASA's contributions to high-performance computing and networking.
www.nasa.gov

SUPERCOMPUTING CONVENTION WELCOME

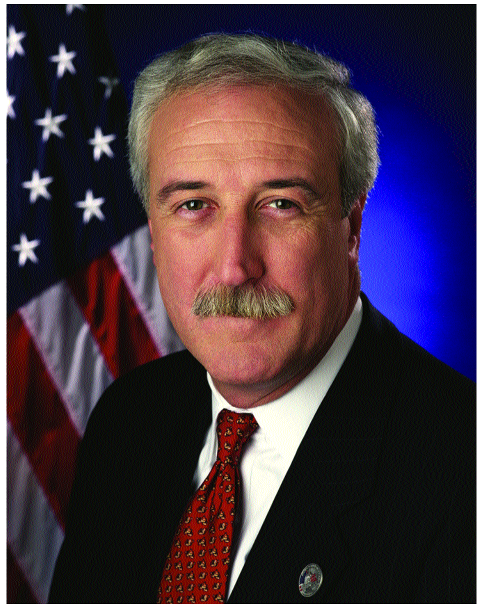

Sean O'Keefe, NASA Administrator

Welcome, SC2003 participants.

I'm pleased to welcome you to the 15th Supercomputing Conference in Phoenix, a city as forward looking as your high-performance computing and networking field.

This conference occurs shortly before our Nation and the world will commemorate the centennial of flight. Fittingly, many technological breakthroughs that have vastly improved the safety and efficiency of our air transportation system and enabled brave explorers to extend our horizons heavenward are derived from advancements in the supercomputing field.

We also know that progress in aviation and space flight requires a willingness to learn from and overcome major setbacks. Appropriately, at SC2003 NASA will show how we were able to utilize supercomputing capabilities in the critical investigation of the tragic Columbia accident.

In this, our 45th year of pioneering the future on behalf of the American public, NASA appreciates the opportunity to share with you how we are also using supercomputing resources to advance our ambitious mission goals. Those goals are to understand and protect our home planet, explore the Universe and search for life, and inspire the next generation of explorers.

At SC2003 you will learn how earlier this year, NASAdemonstrated a capability to extend powerful computing resources to remote locations. This capability could someday allow remote geological classification of rocks and minerals on distant sites, such as on the surface of Mars. We will also demonstrate the Earth System Modeling Framework, a national-scale collaboration to build a software infrastructure that allows weather and climate model components from different researchers to operate together on parallel supercomputers.

We hope you will have the opportunity to learn about these and other impressive supercomputing applications that are taking NASA into the second century of flight. Best wishes for a most successful SC2003!

— Sean O'Keefe

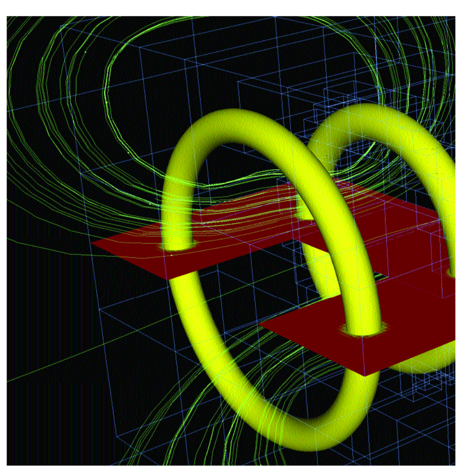

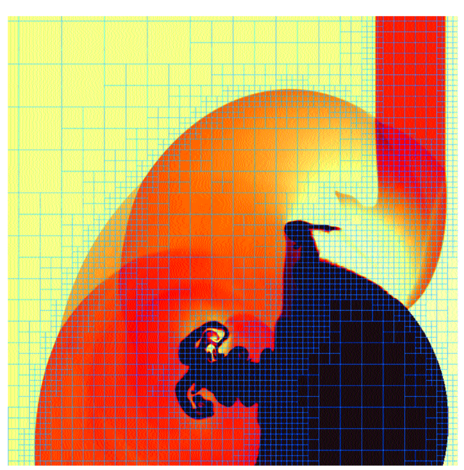

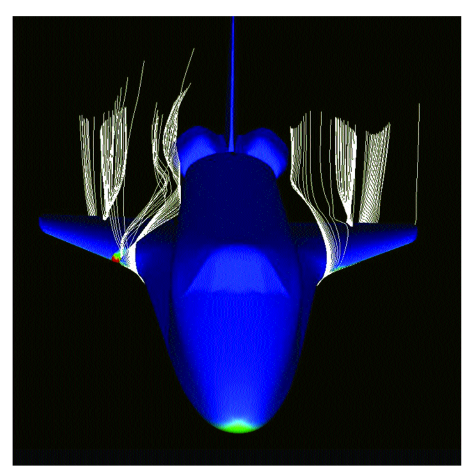

NASA AT SUPERCOMPUTING 2003

NASA is showcasing many of its exciting scientific research projects at SC2003. The following are descriptions of demonstrations listed by participating NASA centers and collaborators.

AMES RESEARCH CENTER

Mountain View, California

Ames Research Center (ARC) and its personnel work to develop technologies that enable the Information Age, expand the frontiers of knowledge for aeronautics and space, improve America's competitive position, and inspire future generations. Ames specializes in research geared toward creating new knowledge and new technologies that span the spectrum of NASA interests.

Ames, SGI 'Grow' a New Supercomputer – The Altix 256 CPU SSI is Born

NASAAmes has a long-standing position as the leader in pushing the development of large single-system image (SSI) machines. Over the past five years, Ames and SGI in partnership have developed the world's first 256-, 512-, and 1,024- CPU global shared memory systems, based on SGI's Origin line of supercomputers. Recently, Ames and SGI partnered to expand the size of the SGI Altix line of supercomputers, which enable much simpler programming models, and dramatically simplify overall system management. As a result, they have many attractions for large research/production facilities such as Ames.

For several months, NASAAmes has been working closely with SGI engineering to rapidly build, benchmark, and production harden a 256-CPU Altix system. At each stage of development (64-, 128-, and 256-CPUs), the system achieved excellent sustained performance and near perfect scalability. Overall, performance was typically several times that of previous Ames platforms.

Currently, over 40 applications have been ported to the new system, including many large production codes, such as OVERFLOW, TLNS3D, and CART3D (aerosciences CFD), and ECCO, POP, and CCSM (Earth sciences climate/ocean models). While research continues, the Ames 256 CPU Altix system is already in limited production mode, with codes running 24/7 and scheduling managed by the Ames PBS batch processing system.

At press time, Ames expects to expand the current system to a 512- CPU SSI system. This effort is part of a joint project between NASA's Office of Aerospace Technology and Office of Earth Science, and the NASA Advanced Supercomputing (NAS) Division at Ames.

Automating Performance Analysis and Debugging of Parallel Programs

New components of an integrated parallelization environment that reduce the need for user interaction through automation have been developed at ARC. One is an automated process that analyzes performance inhibitors in OpenMP programs and makes suggestions for performance improvements. This component acts as an interface between the CAPO parallelization tool and the Paraver performance analysis system, to provide support for an optimal placement of OpenMP directives into sequential Fortran code.

A second component automates the use of relative debugging techniques, by running a multi-threaded OpenMP code side-by-side with its serial counterpart and finding where the values they compute begin to differ. These two new tools automate many of the time-consuming steps in the program development cycle, greatly decreasing the time required for parallelization. For more information, visit http://www.nas.nasa.gov/Tools/SC03/.

Dynamic Authorization in a Distributed Environment

Cardea, developed for the NASAInformation Power Grid, addresses the challenges unique to distributed authorization. The system decouples authorization and identity, distributes responsibility for authorization data, and separates access control decision from enforcement. Cardea directly evaluates authorization requests according to relevant request and resource characteristics rather than preconfigured local identities. Trusted authorities that issue attribute assertions, decision points that evaluate authorization requests, and policy points that store access control policies work cooperatively on each authorization request.

This demonstration highlights the potential of such an approach and showcases the key components working together to evaluate authorization decision requests. The demonstration also shows that each administrative domain can retain complete control over its access control policies and its internal authentication and enforcement mechanisms.

Earth Sciences Web Services Applications and Grid Data Services

NASA is building a unified Grid Data Services environment, and using an early prototype to demonstrate the combined use of web services applications and Grid Data Services as a viable front-end for NASA Earth Observing system (EOS) data. The goal is to provide a standard, flexible, and scalable environment, in order to make EOS data easily accessible to the Earth Science modeling and analysis community. This integration will make Grid technology geospatial-enabled, and make OGC technology Grid-enabled. The initial demonstration involves EOS data in archives at NASA Ames, NASA Goddard, and George Mason University in Virginia. The initial Grid Data service is GridFTP from the Web Services server to the archives. Find out more about NASA's grid effort at http://www.ipg.nasa.gov

Grid Miner

This demo shows the Grid supporting a data mining application, using NASA-developed tools. The Grid Miner uses a Resource Broker Service to select an appropriate computer on the grid on which to perform the mining. The Execution Management Service is then used to control the movement of the Grid Miner to the selected computer. The Grid Miner uses the data transport capabilities of the Grid to acquire the necessary mining operations from a repository, and then moves the data from the archive to the Grid computer where the Grid Miner mines the data. Upon completion of the mining, the Execution Manager transports the results to a location on the Grid where the results can be visualized.

The hyperwall

The hyperwall, a powerful new visualization and exploratory computing system was developed in March 2002 at NASAAmes Research Center. The hyperwall uses a 7x7 array of coordinated displays for interactive exploration of multidimensional or multivariate data and computations. The system leverages the human visual system for integration, synthesis, and pattern discrimination in the immense, high-dimensional data spaces increasingly produced by computational science and engineering. The hyperwall features more than 64 million pixels distributed over 55 square feet of viewing surface, with 100 gigabits per second of visual output. Applications for the hyperwall have been developed in aerospace engineering, computational chemistry, climate and weather modeling, (bio)nanotechnology, computational cosmology and astrophysics, neuroscience, bioinformatics, and others. The hyperwall has already enabled new scientific discoveries in several of these fields. Learn more about the hyperwall at http://www.nas.nasa.gov/Groups/VisTech/hyperwall/

NAS Parallel Benchmarks: New Territory, New Releases

The NAS Parallel Benchmarks (NPB) were among the first widely used application-level HPC benchmark suites. Several major extensions of the NPB suite have been created over the last year: Release of source code implementations of the NPB using new parallelization techniques; increase in memory footprint of the previous largest NPB class; addition of an I/O-intensive benchmark; and the addition of a suite of multizone benchmarks that feature multilevel exploitable parallelism. In addition, a completely new benchmark has been created, featuring an irregular, adaptive mesh and the concomitant unstructured irregular, dynamically changing memory accesses. Find out more about NPB at: http://www.nas.nasa.gov/Software/NPB

NASA Open Source Software

This demo highlights progress in the development of an Open Source option for NASA software. {<italics> Open Source} means that source code and specific rights with respect to use and reuse of the code are provided as part of a software package distribution. For NASA, the adoption of an Open Source option for software distribution would be beneficial in three ways: Enhanced software through community review and development; enhanced collaboration through sharing of NASA-originated software; and more efficient and effective dissemination of research products (such as software) to the public. NASA has begun several pilot projects in Open Source distribution. Assuming these projects go as expected, Open Source distribution of NASA-originated software could extend to codes in scientific visualization, software development tools, software frameworks, Grid computing, and many other activities at NASA. More information about Open Source is located at: http://www.nas.nasa.gov/Research/Software/Open-Source/index.html

Navigation in Grid Space

Navigation or dynamic scheduling of applications on computational grids can be improved through a dynamic characterization of grid resources. The NAS Grid Benchmarks are used to map performance of grid resources into a GridScape representing the dynamic state of the grid. The GridScape is used for automatic assignment of the tasks to grid resources. Based on this approach, a Grid Resources Allocation Tool (GRAT) reduces the time-to-solution of a data mining application and of a flow simulation application by 25-35 percent.

Scientific Visualization Inside a Crystal Ball

Researchers demonstrate how several datasets from current biology and nanotechnology research can be analyzed more effectively using state-of-the-art display technology in 3-D graphics. With a 360-degree viewable volumetric 3-D display inside a 20-inch glass dome, scientists are able to view their data in 3-D more effectively and with more ease than conventional 2-D displays. The 3-D display system called "Perspecta" is made by Actuality Systems, Inc. and produces new ways of exploring scientific data not possible before.

Several scientific visualization applications are demonstrated using this new 3-D display system. This demonstration showcases the latest results from NASA Ames' research on the next-generation display technology for scientific visualization.

NASA at Supercomputing 2003

Supercomputing Support for STS-107 Columbia Accident Investigation Board (CAIB)

Following the Columbia tragedy of Feb. 1, 2003, all of NASA worked together to understand what happened to STS-107. Ames contributed key assets, tools, and expertise to the effort. State-of-the-art computational fluid dynamics (CFD) codes were used to simulate steady and unsteady flow fields around Columbia during ascent. Simulation results prompted the use of a higher velocity and density in foam impact testing done under the CAIB (June 10, 2003), which showed massive damage to the orbiter wing RCC panels and damaged T-seals due to foam impact. Simulations also provided insight into the mechanism of debris shedding from the bipod-ramp region. Each moving-body simulation required 1,000 to 5,000 CPU hours on a 1,024-CPU SGI Origin supercomputer, for a total of ~600,000 CPU hours in over 450 full simulations over a very short time period.

John H. Glenn Research Center

Cleveland, Ohio

As a diverse team working in partnership with government, industry, and academia to increase national wealth, safety, and security, protect the environment, and explore the universe, the John H. Glenn Research Center develops and transfers critical technologies that address national priorities through research, technology development, and systems development for safe and reliable aeronautics, aerospace, and space applications.

Extending Grid Computing to Remote Locations Using the Information Power Grid

NASA Glenn Research Center, in cooperation with NASA Ames and geologists from the University of Cincinnati and Bowling Green State University, has extended the computational capabilities of the Information Power Grid (IPG) to remote research sites. The combination of satellite (EOS) data acquisition and the IPG processing capabilities provides geologists with the ability to identify the key mineralogical features at the research site. The underlying connectivity for this research environment is provided by the NASA Research and Education Network (NREN) using a combination of terrestrial and mobile satellite-based networking solutions. This approach not only speeds the process of scientific discovery, but also serves as a simple demonstration of NASA's capacity for geological classification and exploration of remotes sites such as the Martian surface.

Interactive Large Screen Visualizations

NASA Glenn researchers showcase recent visualization activities ranging from an Information Power Grid geological ground truthing experiment to a full turbine engine simulation.

Securing Sensitive Flight and Engine Simulation Data using Smart Card Technology

NASA Glenn Research Center has developed a smart card simulation prototype that secures a hi-bypass jet engine simulation running the Numerical Propulsion System Simulation (NPSS). Upon initially starting the simulation, the smart card authenticates the user on a Solaris workstation. After successfully authenticating, the smart card simulation prototype software protects the flight and jet engine simulation files from unauthorized disclosure or alteration by encrypting and decrypting all files needed during the run.

Goddard Space Flight Center

Greenbelt, Maryland

Goddard Space Flight Center (GSFC) seeks to expand knowledge of the Earth and its environment, the solar system, and the universe through observations from space. High-performance computing interprets observational data both by processing it into understandable forms and by simulating observed and unobserved phenomena.

Chombo Framework for Block-Structured AMR Applications

GSFC is developing a software framework for implementing block-structured adaptive mesh refinement (AMR) algorithms to solve targeted problems in star formation, microgravity research, and space physics. The Chombo framework is a basis for this effort. Chombo is a set of object-oriented C++ libraries for implementing AMR applications that uses MPI to support distributed memory computing.

Algorithm components under development include finite-volume techniques for representing irregular geometries and multi-fluid interfaces, along with techniques for showing particles in incompressible fluids. The ChomboVis tool provides an interpretive environment for manipulating and visualizing AMR data in the Chombo framework. ChomboVis is built using python and VTK, enhancing portability to a variety of architectures. Both Chombo and ChomboVis use HDF5 for I/O. Visit http://seesar.lbl.gov/ANAG/NASA for additional information.

Coupling Earth System Models: An ESMF-CCA Prototype

An Earth system model typically consists of several complex model components coupled through data exchange. Because a software framework to facilitate coupling and to make models interoperable across organizations would be highly beneficial, NASA's ESTO/Computational Technologies Project has funded the development of the Earth System Modeling Framework (ESMF). The ESMF software consists of a superstructure (coupling) and an infrastructure (data structures and utilities). DOE supports the Common Component Architecture (CCA) project defining a minimal set of standard interfaces for a high-performance component framework. To ensure interoperability between ESMF and CCA, NASA has developed a prototype that adopts CCA's component registration and GUI for component connection. Component functionality and data exchange are designed to meet the ESMF requirements. More information is located at: http://ct.gsfc.nasa.gov/esmf_tasc.

The Coven Problem-Solving Environment

Coven is a problem-solving environment for development, execution, and maintenance of parallel applications. It provides a component-based model where user modules are plugged together to form programs. These modules use a shared data structure to pass data between each other, and use MPI to communicate with parallel instances of themselves. A graphical front-end for program assembly and OpenGL visualization tools is provided. New features include transparent multithreaded user modules, shared memory support, automatic logging with MPE, and a code performance profiler. The profiler allows for analysis of CPU load, memory utilization, and a history view of module execution. Additional information can be found at: http://www.parl.clemson.edu/Coven

Earth System Modeling Framework

The Earth System Modeling Framework (ESMF) is a national-scale collaboration to build a software infrastructure that allows weather and climate model components from different researchers to operate together on parallel supercomputers. It involves NASA, the National Oceanic and Atmospheric Administration, the National Science Foundation, the U.S. Department of Energy, and academia.

The ESMF strengthens communication among diverse groups, connecting scientists with software developers and weather forecasters with climate modelers. This cross-disciplinary venture reduces redundant programming through software that can run models at high performance on a variety of supercomputers. The ESMF also simplifies model construction and the exchange and incorporation of new sub-models. By enabling comparison of models based on alternative scientific approaches, the ESMF ultimately improves predictive capabilities. Find out more at http://www.esmf.ucar.edu

IBEAM: Interoperability Based Environment for Adaptive Meshes

IBEAM is an object-oriented component-based framework for adaptive mesh simulations of astrophysical phenomena on massively parallel architectures. The IBEAM project is a cooperative venture involving the Computational Technologies team at GSFC, the University of Illinois, the University of Chicago, the State University of New York at Stony Brook, and the Universities Space Research Association.

The main design goal is to provide a framework and a set of interoperable simulation components that will allow users to rapidly construct adaptive mesh refinement-based simulations. IBEAM components will provide the capability to simulate hydrodynamics, radiation transport, self-gravity, and a variety of other physical processes. The testbed problem is the modeling of light curves from gamma-ray bursts observed by NASA's Compton Gamma-Ray Observatory. Additional information located at: http://www.ibeam.org/sc2003

The Parallel Virtual File System Version 2

PVFS2 is a next-generation parallel file system for Linux clusters. It harnesses commodity storage and network technology to provide concurrent access to data that is distributed across a (possibly large) collection of servers. Parallel file systems have become an important aspect of large-scale scientific computing as a means to alleviate the I/O bottleneck inherent in processing large datasets. PVFS2 serves as a testbed for both parallel I/O research and Zas, a freely available production-level tool for the cluster community. New features of PVFS2 include distributed metadata servers, native support for several network protocols, and file system support for complex data types. Development and support is carried out in collaboration with Clemson University, GSFC, and Argonne National Laboratory. Visit http://www.parl.clemson.edu/pvfs to learn more.

Sci-Interactives: Innovative Science Outreach from Truth-N-Beauty Software and NASA CT

The advent of broadband Internet access together with high-speed processors for commercial PCs has opened significant new avenues for communicating science to the general public. Truth-N-Beauty Software, a science "e-" outreach and education company, has partnered with NASA's Computational Technologies (CT) Project to develop digital tools for teaching the public about the latest NASA Earth and space science research.

Sci-Interactives are Web-based applications that give people access to the real equations and data behind NASAscience with intuitive, graphically sophisticated tools and games. The Sci-Interactives explore phenomena ranging from the influence of sunlight on Earth's climate to the collapse of rotating interstellar clouds. More information can be found at: http://www.truth-n-beauty.com/nasact

JET PROPULSION LABORATORY (JPL)

Pasadena, California

The Jet Propulsion Laboratory's mission is to enable the nation to explore space for the benefit of humanity. The lab's mission is: to explore our own and neighboring planetary systems; to search for life outside the Earth's confine; to further our understanding of the origins and evolution of the Universe and the laws that govern it; to make critical measurements to understand our home planet and help protect its environment; to apply JPL's unique skills to solve problems of national security and national significance; and to inspire the next generation of explorers.

Montage

Montage is a grid-capable astronomical mosaicking application. In this demonstration, Montage is used to re-project, background match, and finally co-add many 2MASS plates into a single-image mosaic. Montage is being developed under NASA's Earth Science Technology Office Computational Technologies project by a team that includes JPL, Caltech's Center for Advanced Computing Research, and its Infrared Processing Analysis Center. The grid-enabled prototype of Montage is being built in collaboration with the University of Southern California's Information Sciences Institute. The prototype uses the Pegasus planning system, which maps the complex Montage workflows onto the Grid resources. For additional information, visit http://montage.ipac.caltech.edu/ and http://pegasus.isi.edu

The QuakeSim Project

The QuakeSim project is developing a solid earth system science framework for creating an understanding of active tectonic and earthquake processes. New space-based techniques, coupled with high-performance computer models, now make it possible to measure the quiet motions associated with plate tectonics and the earthquake cycle, thus revolutionizing our understanding of earthquake processes and fault interactions. The QuakeSim project develops tools to simulate earthquake processes, and manage and model the increasing quantities of data now available. These tools include a research-oriented database system, a collaborative problem-solving portal environment, and three scalable parallel simulation codes (available for download). Without these applications, it is impossible to construct the more complex models and simulations necessary to develop hazard assessment systems critical for reducing future losses from major earthquakes. Additional information about the QuakeSim Project can be found at: http://www-aig.jpl.nasa.gov/public/dus/quakesim/, and the movies shown are located at http://pat.jpl.nasa.gov/public/RIVA/images.html.

LANGLEY RESEARCH CENTER

Hampton, Virginia

In alliance with industry, other agencies, academia, and the atmospheric research community, Langley Research Center (LaRC) undertakes innovative, high-payoff aerospace and scientific activities beyond the risk limit or capability of commercial enterprises and delivers validated technology, scientific knowledge, and understanding of the Earth's atmosphere. The center's success is measured by the extent to which its research results improve the quality of life.

Cluster-Based Intelligent Agents for Interactive Visualization in Immersive Environments

This work demonstrates a cluster-based architecture for interactively visualizing large-scale datasets in an immersive environment. The architecture supports a system of intelligent agents that move through the data-space searching for important features. The user interacts with the agents through a multi-modal interaction system that allows the user to specify important features and guide the agents in their search. The agents are based on the swarming concept that although each agent has limited intelligence, the overall behavior of the collection of agents exhibits advanced intelligence. This work is demonstrated on Computational Fluid Dynamics simulations for the Space Shuttle.

NASA Participation: Technical Papers, Tutorials, Panel Discussions
Several NASA scientists (both contractor and civil service personnel) and NASA collaborators are participating on SC2003 committees, panels, tutorials, and other conference events.

Paper Presentations

Job Scheduler Architecture and Performance in Computational Grid Environments
Tuesday, November 18, 4:30PM - 5:00PM
Hongzhang Shan and Leonid Oliker (both of Lawrence Berkeley National Laboratory); Rupak Biswas (NASA Ames Research Center)

Computational grids hold great promise for utilizing geographically separated, heterogeneous resources to solve large-scale scientific problems. However, several major technical hurdles, including distributed resource management and effective job scheduling, stand in the way of realizing these gains. The authors pose a novel grid superscheduler architecture and three distributed job migration algorithms, and also model the critical interaction between the superscheduler and autonomous local schedulers. Extensive performance comparisons with ideal, central, and local schemes using real workloads from leading computational centers are conducted in a simulation environment. Additionally, synthetic workloads are used to perform a detailed sensitivity analysis of the superscheduler. Several key metrics demonstrate that substantial performance gains can be achieved via smart superscheduling in distributed computational grids.

Evaluation of Cache-based Superscalar and Cacheless Vector Architectures for Scientific Computations
Wednesday, November 19, 2:30PM - 3:00PM
Leonid Oliker, Andrew Canning, Jonathan Carter, John Shalf, David Skinner (all Lawrence Berkeley National Laboratory); Stephane Ethier (Princeton University), Rupak Biswas (NASA Ames Research Center), Jahed Djomehri and Rob Van der Wijngaart (both of Computer Sciences Corporation)

The growing gap between sustained and peak performance for scientific applications is a well-known problem in high-end computing. The recent development of parallel vector systems offers the potential to bridge this gap for many computational science codes and to deliver a substantial increase in computing capabilities. This paper examines the intranode performance of the NEC SX-6 vector processor and the cache-based IBM Power 3/4 superscalar architectures across several scientific computing areas. Topics include: Performance of a microbenchmark suite that examines low-level machine characteristics; the behavior of the NAS Parallel Benchmarks; and the performance of several scientific computing codes. Results show that the SX-6 achieves high performance on a large fraction of the applications and often significantly outperforms the cache-based architectures. However, certain applications are not easily amenable to vectorization and would require extensive reengineering to utilize the SX-6 effectively.

Grid-Based Galaxy Morphology Analysis for the National Virtual Observatory
Thursday, November 20, 11:00AM - 11:30AM
Chair: Jay Larson (Argonne National Laboratory) and John Drake (Oak Ridge National Laboratory) Speaker(s)/Author(s): Ewa Deelman (ISI), Raymond Plante (NCSA), Carl Kesselman (USC/ISI), Gurmeet Singh (USC/ISI), Mei Su (USC/ISI), Gretchen Greene (Space Telescope Science Institute), Robert Hanisch (Space Telescope Science Institute), Niall Gaffney (Space Telescope Science Institute), Antonio Volpicelli (Space Telescope Science Institute), James Annis (Fermi Lab), Vijay Sekhri (Fermi Lab), Tamas Budavari (John Hopkins University), Maria Nieto-Santisteban (John Hopkins University), William O'Mullane (John Hopkins University), David Bohlender (Canadian Astrophysical Data Center), Tom McGlynn (NASA), Arnold Rots (Simthsonian Astrophysical Observatory), Olga Pevunova (NASA)

Part of the development of the National Virtual Observatory (NVO), a data grid for astronomy, is the creation of a prototype science application to explore the dynamical history of galaxy clusters by analyzing the galaxies' morphologies. The purpose of the prototype is to explore how grid-based technologies can be used to provide specialized computational services within the NVO environment. The paper describes the scientific goals of the application; the key technology components, particularly Chimera and Pegasus, which are used to create and manage the computational workflow; and how the components are connected and driven from the application's portal.

Poster Presentation

Large-scale Simulations of Hybrid Inorganic-Organic Systems on Multi-Teraflop Linux Clusters and High-end Parallel Supercomputers
Tuesday, November 18, 5:01PM - 7:00PM
Satyavani Vemparala (Louisiana State University), Bijaya B. Karki (Louisiana State University), Hideaki Kikuchi (Louisiana State University), Rajiv K. Kalia (University of Southern California), Aiichiro Nakano (University of Southern California), Priya Vashishta (University of Southern California), Shuji Ogata (Nagoya Institute of Technology), Subhash Saini (NASA Ames Research Center)

A scalable parallel algorithm has been developed for organic macromolecules coupled with inorganic materials, based on a novel spatial-decomposition scheme and space-time multi-resolution algorithms. The algorithm used employs: 1) a novel parallel implementation of the fast multipole method; 2) penalty-function-based constrained, combined with a symplectic multiple time-scale method; 3) dynamic management of distributed data structures; and 4) local-topology-preserving computational-space decomposition for latency minimization and load balancing. Scalability has been tested on the 1,024-processor Intel Xeon-based Linux cluster, the 1,184-processor IBM SP4, and the 512-processor Compaq AlphaServer system. Efficiency is 0.87 for a 21.6 million-atom system on 1,024 SP4 processors. Video showing the design of electrically switching organic nano-films will be shown.

Bird-of-a-Feather

The Climate/Ocean/Weather and HPC Communities: The State of the Union
Wednesday, November 19, 5:00PM - 6:00PM
Chair: Jay Larson (Argonne National Laboratory) and John Drake (Oak Ridge National Laboratory)

Climate, ocean, and weather applications constitute a grand challenge in high-performance computing. During the 1990s, this community—like other computational science communities—worked hard to migrate software from vector platforms to commodity-based multiprocessors. This effort proved long and arduous, and some found it difficult to attain satisfactory levels of efficiency (defined as a fraction of peak performance) on these platforms. At SC01, one bird-of-a-feather (BOF) session asked the question: Is U.S. climate modeling in trouble?

This bird-of-a-feather session will be a forum in which we attempt to answer these questions such as: What do climate/ocean/weather scientists require, and how can the HPC community help? How well will current efforts meet the challenges posed by new processor/platform architectures? What is the current and future role of grid computing in climate?

Workshop

Petaflops Programming Models: Ameliorating Architectural Issues or Exacerbating Them?
Tuesday, November 18, 1:30PM - 3:00PM
Chair: John van Rosendale (Department of Energy), et al. Speaker(s)/Author(s): Moderator: Burton Smith (Cray), Panelists: Rusty Lusk (Argonne National Laboratory), Hans Zima (University of Vienna and NASA Jet Propulsion Laboratory), Kathy Yelick (UC Berkeley), Larry Snyder (University of Washington)

This panel examines evolving and expected programming models, and the way language and compiler developers hope to address the challenges posed by petaflops architectures. Given the range of issues such systems are raising, which issues should the programming model attempt to address, and which should be passed on to the end-user? "Hot-button" issues include: How to handle fault tolerance; very deep memory hierarchies; interoperability; coping with tens, or hundreds of thousands of threads; and the semantics of intelligent memory.

Panels

High Performance Computing System Performance Modeling
Thursday, November 20, 3:30PM - 5:00PM
Chair: Larry Davis (DoD High Performance Computing Modernization Program); Panelists: Allan Snavely (SDSC), Jack Dongarra (University of Tennessee), Walt Brooks (NASA Ames), Adolfy Hoisie (DOE), Henry Newman (Instrumental), John McCalpin (IBM)

This panel consists of government and government-sponsored researchers in the area of high performance computing system performance modeling. Benchmarks are structured to address a range of target configurations. They provide the High Performance Computing Modernization Program (HPCMP) with accurate performance information on available high performance computing (HPC) capabilities. The intent is to provide a set of program source code listings, makefiles, runtime scripts, input files, and validated results files which represents the type of computational work performed on HPC resources. Questions addressed by the panelists include:

1. Describe how and why benchmarks are used in your organization.
2. Describe a minimal spanning set of HPC system attributes that can accurately model the performance of HPC systems on a broad class of applications.
3. What is the potential for the use of synthetic benchmarks to accurately model the performance of a broad class of applications, both now and in the future?
4. For procurement related decisions, do you allow vendors to optimize the benchmarks for their particular architecture? Why or why not?

HPC Productivity
Friday, November 21, 8:30AM - 10:00AM
Chair: Jeremy Kepner (MIT Lincoln Laboratory); Panelists: Ken Kennedy (Rice University), David Kuck (Intel), Mark Snir (University of Illinois), Thomas Sterling (CalTech/NASA JPL), Bob Numrich (University of Minnesota), John Gustafson (Sun)

The value of a high-performance computer to a user involves many factors, including execution time on a par-

ticular problem, software development time, and direct and indirect costs. The DARPA High Productivity Computing Systems program is focused on providing a new generation of economically viable, high-productivity computing systems for national security and the industrial user community, in the 2007-2010 timeframe. The goal is to provide systems that double in productivity (or value) every 18 months.
This program has initiated a fundamental reassessment of how we define and measure performance, programmability, portability, robustness, and ultimately, productivity in the HPC domain. Panelists will present their views on two fundamental questions:

1. How should we define and measure productivity in HPC?
2. What are the implications for HPC designers and users?

Tutorial M05

Component Software for High-Performance Computing: Using the Common Component Architecture

Monday, November 17, 8:30AM - 5:00PM

Robert C Amstrong (Sandia National Laboratories), David E Bernholdt (Oak Ridge National Laboratory), Lori Freitag Diachin (Sandia National Laboratories), Wael R Elwasif (Oak Ridge National Laboratory), Daniel S Katz (Jet Propulsion Laboratory, California Institute of Technology), James A Kohl (Oak Ridge National Laboratory), Gary Kumfert (Lawrence Livermore National Laboratory), Lois Curfman McInnes (Argonne National Laboratory), Boyana Norris (Argonne National Laboratory), Craig E Rasmussen (Los Alamos National Laboratory), Jaideep Ray (Sandia National Laboratories), Sameer Shende (University of Oregon), Shujia Zhou (Northrop Grumman/TASC)

Content-Level: 25% Introductory 50% Intermediate 25% Advanced

This tutorial will introduce participants to the Common Component Architecture (CCA) at both conceptual and practical levels. Component-based approaches to software development increase developer productivity by helping manage the complexity of large-scale software applications and facilitating the reuse and interoperability of code. The CCAwas designed specifically with the needs of high-performance scientific computing in mind. It takes a minimalist approach to support language-neutral component-based application development for both parallel and distributed computing without penalizing the underlying performance, and with a minimal cost to incorporate existing code into the component environment. The CCA environment is also well suited to the creation of domain-specific application frameworks, whereas traditional domain-specific frameworks lack the generality and extensibility of the component approach.

NASA at Supercomputing 2003

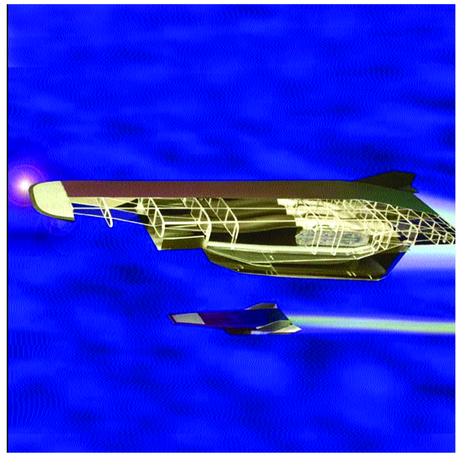

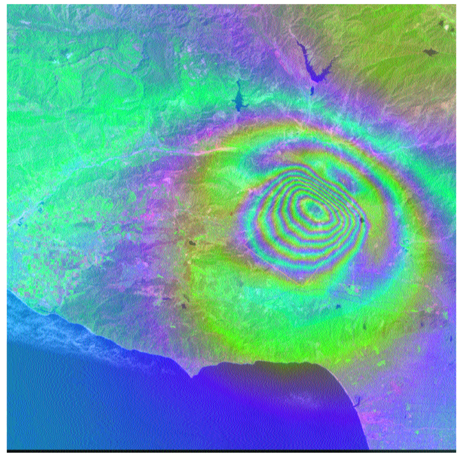

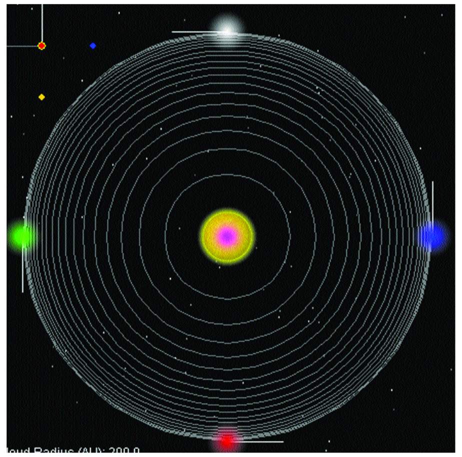

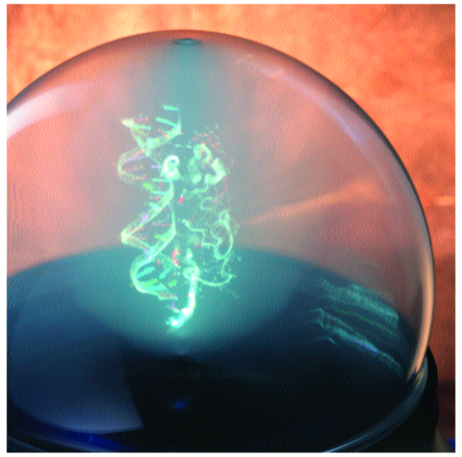

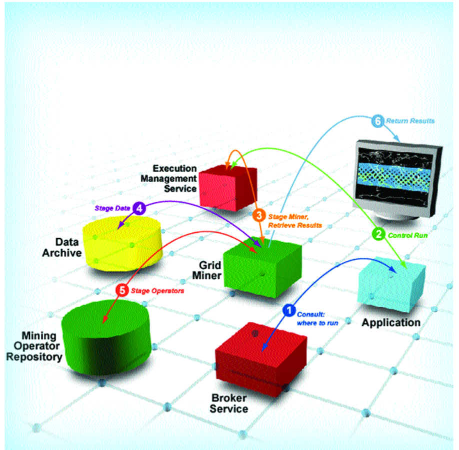

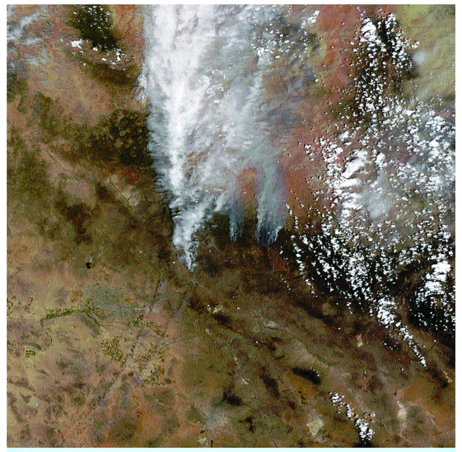

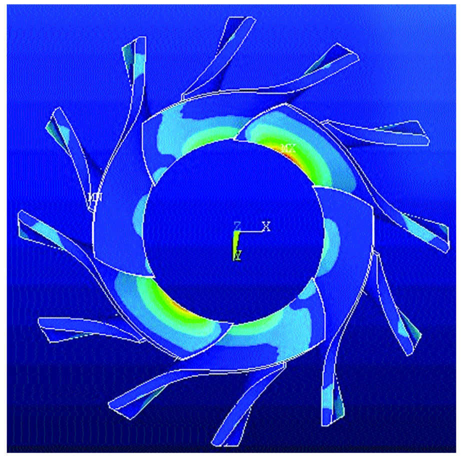

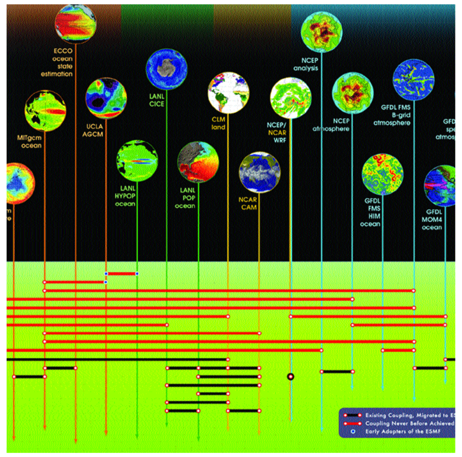

NASA Booth #411

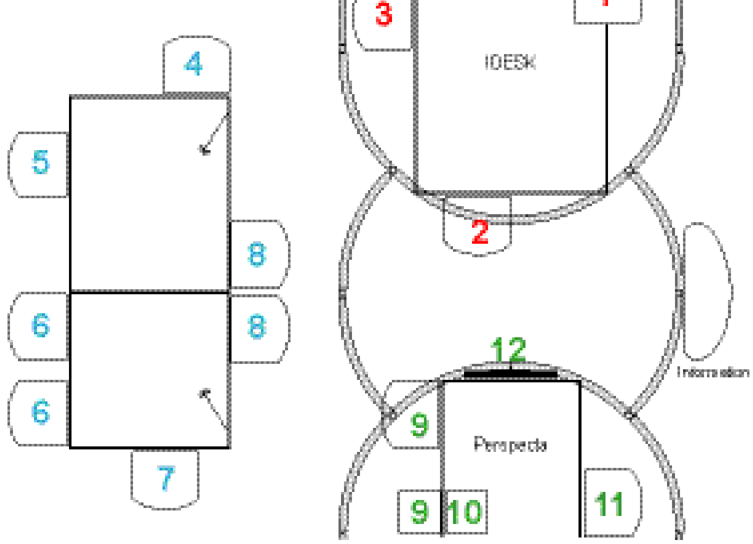

1
- Interactive Large Screen Visualizations (GRC)
- Cluster-Based Intelligent Agents for Interactive Visualization in Immersive Environments (LaRC)

2
- QuakeSim (JPL)
- Montage (JPL)

3
- Grid Miner (ARC)
- Dynamic Authorizaton in a Distributed Environment (ARC)
- Earth Sciences Web Services Applications and Grid Data Services (ARC)
- Navigation in Grid Space (ARC)

4
- Chombo Framework for Block-Structured AMR Applications (GSFC)
- IBEAM: Interoperability-Based Environment for Adaptive Meshes (GSFC)
- Sci-Interactives: Innovative Science Outreach from Truth-N-Beauty Software and NASA CT (GSFC)

5
- Extending Grid Computing to Remote Locations Using the IPG (GRC)

6
- The Parallel Virtual File System 2 (PVFS2) (GSFC)
- The Coven Problem-Solving Environment (GSFC)

7
- Automating Development of Efficient OpenMP Programs (ARC)
- NAS Parallel Benchmarks: New Territory, New Releases (ARC)

8
- Earth System Modeling Framework (GSFC)
- Coupling Earth System Models: An ESMF-CCA Prototype (GSFC)

9
- NASA Open Source Software (ARC)
- The hyperwall (ARC)

10
- Scientific Visualization Inside a Crystal Ball (ARC)

11
- Ames Performance on the Altix Supercomputer System (ARC)
- STS-107 Investigation Support (ARC)

12
- Lightbox: "War and Peace Nebula" (JPL)