SC2002 High Performance Networking and Computing

A look at NASA exhibits,
events, and speakers at SC2002

A special section of
Gridpoints Magazine, Fall 2002

Welcome, SC2002 participants.

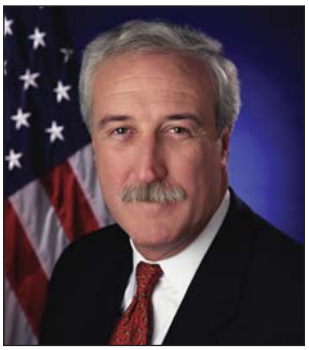
Bill Ingalls/NASA

Sean O'Keefe
NASA Administrator

I'm pleased to welcome you to a conference that represents everything NASA is all about: pioneering the future, pushing the envelope, and doing what has never been done before.

The city of Baltimore will provide a wonderful setting for productive professional interchange about the latest developments in information architectures, scalable computing, data analysis, applications, and collaborative technologies.

For our part, NASA will demonstrate at SC2002 how we are utilizing advanced supercomputing technologies to improve aviation safety and efficiency, probe more deeply into the mysteries of the universe, model terrain and pinpoint possible spacecraft landing sites on Mars, and work to better understand the dynamics of Earth's climatic system.

We're also delighted to showcase how the adept use of computational science was a key factor in the development of the NASA-(Dr. Michael) DeBakey Ventricular Assist Device, NASA's Commercial Invention of the Year. NASA computational techniques that normally model rocket fuel flow helped to perfect this miniature heart assist pump, which promises to improve the lives of thousands of heart patients.

When we think about touching lives, clearly the kind of work all of us are engaged in has the potential to help inspire the next generation of explorers, which not so coincidentally is one of NASA's core mission goals. Given the fact that fewer young people are choosing to enter technical fields, I'm very pleased that this conference offers excellent programs that will help teachers apply the exciting work of computational research into classroom activities. By assisting teachers in their efforts to connect students to the technology frontiers you explore on a daily basis, I believe that SC2002 will help in immeasurable ways to build a better America.

Sean O'Keefe

NASA is showcasing many of its exciting scientific research projects at SC2002. The following are descriptions of demonstrations, paper presentations, videos, and panel discussions 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.

Atomistic Simulations on USA and Japan Grid

Researchers will demonstrate results from multiscale quantum-mechanical and/or classical atomistic simulations on a grid of geographically distributed PC clusters, to study environmental effects of water molecules on fractures in silicon.

The multiscale simulation approach seamlessly combines atomistic simulations based on the molecular dynamics (MD) method and quantum mechanical (QM) calculations based on the density functional theory. The multiscale MD/QM simulation code has been grid-enabled using a modular additive hybridization scheme; multiple QM clustering; and computation/communication overlapping. A preliminary code run has achieved a parallel efficiency of 94 percent on 25 PCs distributed over three PC-clusters in the United States and Japan. A video trip through a silicon crack tip shows water-silicon reactions.

Cart3D: A Package For Automated Grid Generation and Aerodynamic Database Creation

Cart3D, NASA's Software of the Year for 2002, is a software package for fully automated geometry processing and computational fluid dynamics (CFD) simulation. The software enables users with access to networked computing clusters or supercomputers to synthetically generate the types of aerodynamic databases typically obtained only through extensive wind tunnel testing. SC2002 attendees can select and manipulate pre-built vehicle geometries, and watch while meshes suitable for CFD analysis are automatically generated on demand. An entire suite of CFD runs for the chosen configuration can then be submitted for processing on NASA's Information Power Grid. These simulation results constitute a synthetically generated aerodynamic database that gives analysts information needed for design. Animations, images, and results from previously computed simulation suites will also be available for playback, discussion, and demonstration. Additional information is available on the web at: www.nas.nasa.gov/~aftosmis/cart3d

The Hyperwall: A Multiple Flat-Panel Display Wall For High-Dimensional Visualization

The NASA Ames "hyperwall" consists of a seven-by-seven matrix of flat panel displays driven by 49 rack-mounted dual-CPU nodes, each with its own high-end graphics card. The hyperwall emphasizes displaying multiple independent but related images, and provides useful means for composing and controlling these image sets. In place of elaborate software or hardware crossbar switches, researchers rely on the human visual system for integration, synthesis, and pattern discrimination in complex and high-dimensional data spaces.

The tabular layout of the hyperwall supports many existing "multiview" visualization paradigms, including spreadsheet-style approaches, multidimensional/multivariate techniques, and brushing/linking.

DeBakey Ventricular Assist Device NASA's Invention of the Year – 2001

The use of computational fluid dynamics (CFD) technology has led to several major design improvements on the NASA-DeBakey Ventricular Assist Device (VAD). NASA Ames scientists Cetin Kiris and Dochan Kwak employed NASA Shuttle main engine technology and CFD modeling capabilities, coupled with the NASA Advanced Supercomputing (NAS) Division's high-performance computing

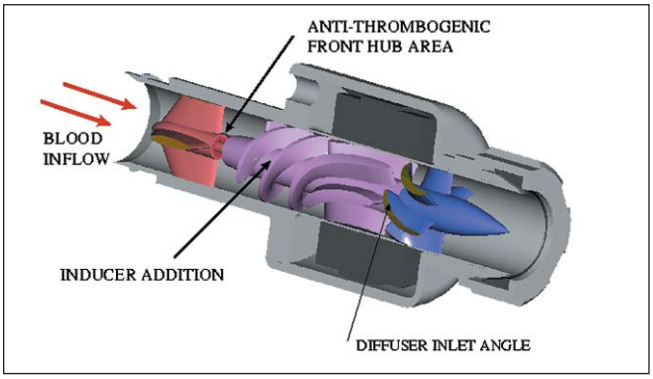

Using CFD analysis, NAS researchers found that major design modifications to the NASA-DeBakey Ventricular Assist Device were necessary. The result of these changes: overall efficiency of the device was increased by 22 percent.

(MicroMed Technology)

technology, to make several design modifications that vastly improved this miniature heart pump's performance. To date, it has been successfully implanted in 170 individuals in Europe, Asia, and the United States. In April 2002, the NASA-DeBakey VAD was named NASA's Commercial Invention of the Year for 2001.

Information Power Grid

Grids are an emerging technology that provide seamless and uniform access to geographically dispersed computational, data storage, networking, instruments, and software resources needed for solving large-scale scientific and engineering problems.

NASA's Information Power Grid (IPG) project is developing and deploying such a computing and data grid, and its goal is to use NASA's remotely located computing and data system resources to build distributed systems that can address problems too large or complex for a single site.

IPG Job Manager

The IPG Job Manager allows IPG users to reliably execute applications on remote computer systems specified by users. The Job Manager provides the following functionalities: moving files between file servers before an application begins executing (pre-staging) – a typical use of this is to move input data files to a computer system where an application will be run; moving files between file servers after an application completes executing (post-staging) – a typical use of this is to move output data files from a computer system where an application was run to a repository of project data; and reliably executes an application. The IPG Job Manager handles a variety of software and hardware failures; monitors the progress of the execution of a job; and maintains a history of the execution of an application that is available to users after the application finishes executing.

IPG Resource Broker

A computational grid such as the IPG consists of a large number and variety of computer systems. Users may not be familiar with all the available resources and their current state, so it can be difficult for users to select where to run an application. The purpose of the IPG Resource Broker is to provide suggestions for which computer system should be used to execute an application. Users provide: the number of CPUs and the amount of memory required; the executables to use for different operating systems; job execution parameters such as command line arguments, environment variables, scheduling queue, project, etc.; and the files to pre- and post-stage (optionally qualified by host name). From this input, the Resource Broker suggests where to run an application in the form of Job Manager Jobs. Each Job Manager Job contains all the information needed to execute the application, including the suggested computer system. The broker suggests systems by considering: the candidate hosts and operating systems (if any) that the user specified; the number of CPUs and amount of memory specified by the user; and the current status and load on the systems; the predicted completion times of the application on different systems.

Performance Prediction

Performance grids allow users to access a large number and variety of computer systems. This leads to the question of where users should run their applications. The NASA IPG team provides predictions of application execution and completion times to assist users in making this decision.

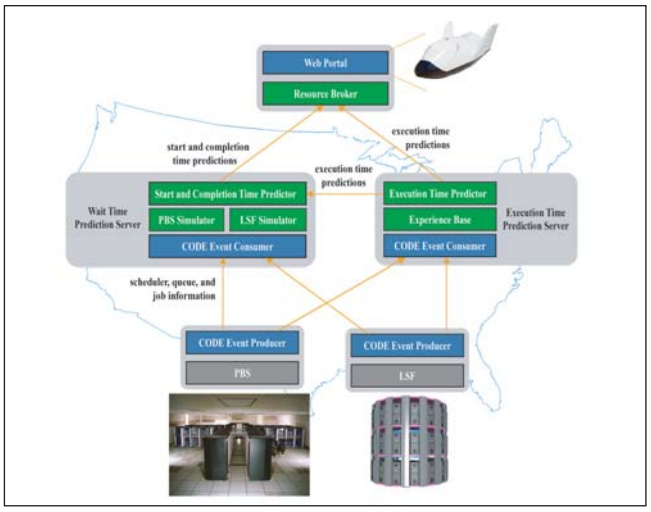

IPG's performance prediction architecture enables users to predict how an application will run on the grid, and how long the task will take. (NASA)

The team predicts execution times of applications on space-shared parallel computers using instance-based learning techniques. This tool maintains a database of experiences, called an experience base. The prediction for a query is formed using the outcomes of experiences in the database that are similar to the query. Experiences are job descriptions consisting of characteristics such as user name, number of CPUs, and executable name.

Access to most large computer systems is scheduled using queues and a scheduling algorithm that decides when to start jobs. Completion times of applications are predicted by simulating the operation of scheduling algorithms, given the jobs currently in the queues, and predictions of how long those jobs will execute.

Multi-level Parallelization and Optimization Environment

NASA researchers show three prototype tools working together to transform serial codes into multi-level parallel codes. Depicting a typical parallelization scenario, it starts with a serial cloud-modeling code and uses CAPO, a tool for assisting in parallelization, to produce an OpenMP version of the code. User inputs during parallelization are incorrect, and the resulting program does not produce the same answers as the original serial code. An automatic debugging tool is then applied, which uses information from CAPO to isolate the cause of the differences. Once the problem is fixed, the Paraver tool is used to demonstrate how to analyze performance of the parallel program. The scenario concludes with a visualization of the data computed by the program. Visit: www.nas.nasa.gov/Tools

Programmable Visualization of EOS DAAC Data

This demonstration shows programmable visualization components for the Earth Observing System (EOS) Distributed Active Archive Center (DAAC) data. The components include a central data model ("Field Model"), a metadata model ("Active Metadata"), and modules capable of interfacing to various data formats, including HDF-EOS files. The components can easily be assembled within a framework based on the Python programming language. Python includes numerous modules for web data access,

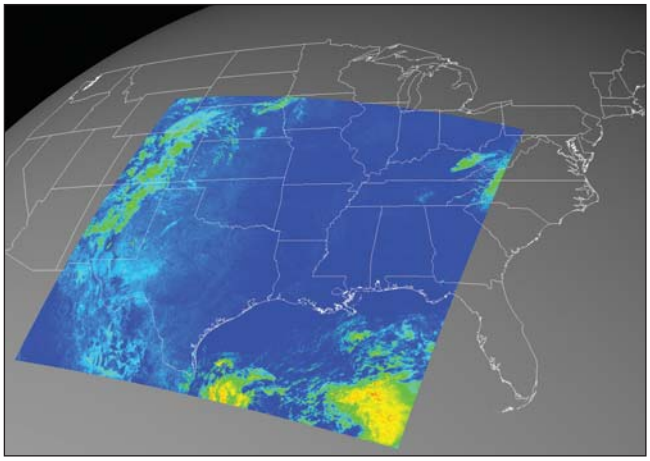

The visualization of the U.S. Gulf Coast region is based on Level One-B swath data acquired September 29, 2000, by the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument aboard the Terra satellite. Terra is the flagship in NASA's Earth Observing System. The datasets are made available by the Distributed Active Archive Center based at Goddard Space Flight Center, Greenbelt, Maryland.

offering the potential for automated access of DAAC data. The interactive visualization modules, coupled with the programmability of Python, enables users to complement the insights gained by interactive techniques with the efficiencies of programmable analysis. More information is available at: http://field-model.sourceforge.net

*-Ray: Large-Scale Data Rendering Software

Pronounced *star-ray*, *-Ray is a multi-threaded software package developed by researchers at the University of Utah and is capable of rendering large-scale scientific data sets at interactive rates. Star-ray rendering is entirely software based, taking advantage of the ccNUMA architecture of the SGI Origin series. Star-ray is able to achieve near linear scaling in performance, based on the number of pixels rendered and number of processors used. Ray tracing also provides an ideal framework for rendering large-scale data, because the rendering time scales sublinearly with scene complexity. The primary limitation on data size is main memory. With NASA's SGI Origin 3000, researchers will demonstrate the ability to interactively render datasets ranging from gigabytes to hundreds of gigabytes.

Simulation-Based Medical Planning for Vascular Surgery

The current paradigm for cardiovascular surgery planning relies exclusively on diagnostic imaging data to define the present state of the patient, empirical data to evaluate the

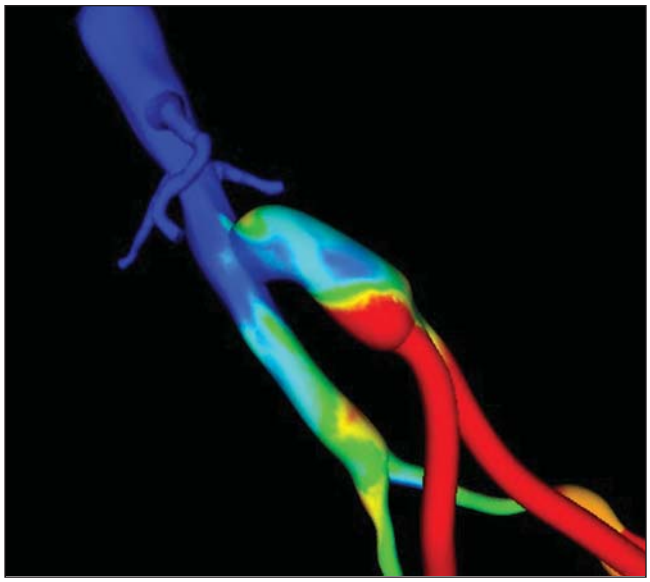

Flow simulations are performed for each preoperative plan and quantities of interest (such as wall shear stress, particle residence time) are calculated. (Stanford University)

effectiveness of prior treatments for similar patients, and the judgment of the surgeon to decide on a preferred treatment. Researchers at Stanford University and NASA are working to create a simulation-based medical planning system for cardiovascular disease that uses computational methods to evaluate alternative surgical options prior to treatment, using patient-specific models of the vascular system. The blood flow simulations enable a surgeon to see the flow features resulting from a proposed operation and to determine if they pose potential adverse effects such as increased risk of atherosclerosis and blood clot formation. Further details on this research can be found at:
www.med.stanford.edu/school/vascular

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.

Coupling Engineering Applications on NASA's IPG: A CORBA-Based Approach

Researchers will demonstrate the results of an engineering parameter study that pairs NASA Langley Research Center's VULCAN (Viscous Upwind aLgorithm for Complex flow ANalysis) code with CORBA-based services and applications developed at Glenn. The demo will feature this CORBA-based environment transferring data and executing the VULCAN code on IPG resources.

A secondary benefit of this research is the ability to couple CORBA-IPG middleware to Glenn's Numerical Propulsion System Simulation (NPSS) CORBA Component Developer's Kit (CCDK). This capability enables VULCAN to function as a service to both NPSS and stand-alone IPG client applications.

Interactive Large-Screen Visualizations

NASA Glenn researchers will showcase recent activities in adapting interactive large-screen visualizations of scientific and engineering research. Demonstrations will include: examples of research in space communications and distrib-

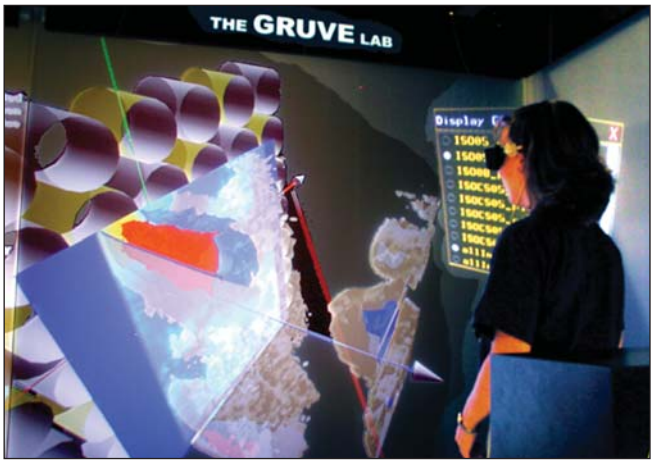

Glenn Research Center's interactive visualization theater will feature demonstrations of interactive computational research.

uted supercomputing solutions using advanced propulsion models; fuel cell design and research; biotechnology; and microgravity science, all supported by NASA's Computing, Information, and Communication Technologies Program.

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.

Modeling the Earth's Atmosphere at the Data Assimilation Office (DAO)

The NASA finite-volume General Circulation Model (fvGCM) is the DAO's next-generation modeling system. It is based on a unique, state-of-the-art finite-volume dynamical core and community-built physical parameterizations and land surface model. A highly efficient application of the MPI-2 remote memory access message-passing paradigm has been implemented within the NASA fvGCM. This implementation improves the throughput by as much as 39 percent, compared to the MPI-1 non-blocking communication paradigm. The NASA fvGCM produces high-resolution global forecasts capable of resolving atmospheric

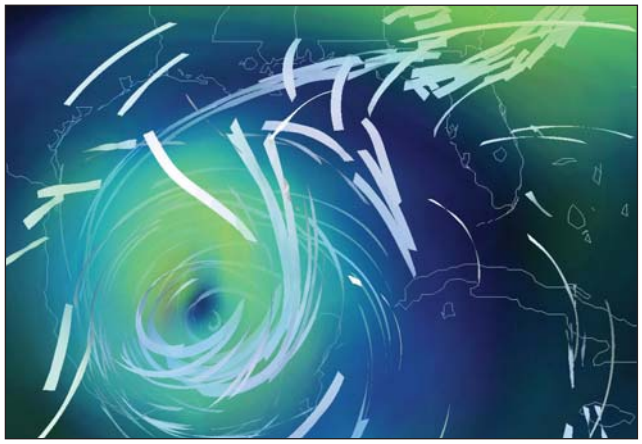

Snapshot of a three-dimensional animation of a hurricane developing in the Gulf of Mexico created from a 0.5- x 0.625-degree, 32-level run of the NASA fvGCM. (NASA)

motions from meso- to planetary-scale with a high throughput on multiple-processor, distributed memory computing platforms. For more information about the Data Assimilation Office, visit: http://dao.gsfc.nasa.gov

Ultra-High Resolution Astronomy: Phasing of Arrays of Formation Flying Spacecraft

Goddard Space Flight Center researchers are investigating arrays of multiple spacecraft, phased together and flying in formation. Systems of this type would enable ultra-high resolution science, including resolution of the stellar disks of nearby stars, as well as the planets around these stars. Systems with baselines (largest distance between spacecraft) of up to 500 meters are currently under study. GSFC scientists are researching and developing tools for systems modeling to determine the system design drivers, set science goals, and ascertain system performance. This is inherently a large-scale computing problem, both in terms of modeling and in terms of the algorithms to actively control the system. It is likely that such systems would also require autonomous on-board computing.
http://code935.gsfc.nasa.gov/cube%20Folder/OSCAR/index.html

IBEAM: Interoperability-Based Environment for Adaptive Meshes

IBEAM's technology goals are to develop a community framework for scalable astrophysical modeling simulations that promotes code interoperability. Design constraints require that the framework enable adaptive mesh refinement techniques; support interoperability of astrophysical modeling codes; be component-oriented; promote Object Oriented Design (OOD) concepts to enhance code reusability; and be coded in F90+MPI to promote performance and portability. The testbed problem is modeling gamma-

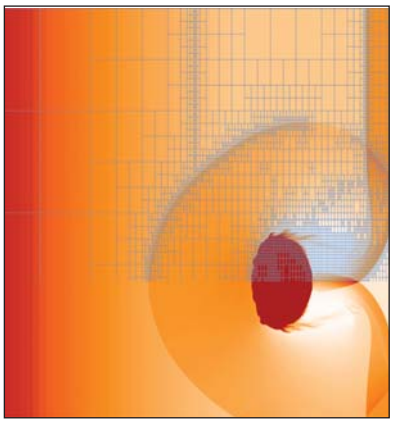

Adaptive mesh refinement simulation of a relativistic shock wave interacting with a dense gas cloud. (NASA)

Chombo Framework for Block-Structured AMR Applications

Researchers are 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 use MPI to support distributed memory computing.

Algorithm components in development include finite-volume techniques for representing irregular geometries and multi-fluid interfaces, along with techniques for 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 the Python scripting language and the Visualization Tool Kit, enhancing portability to a variety of architectures. Both Chombo and ChomboVis use HDF5 (Hierarchal Data Format) for input and output. For more details, visit: http://seesar.lbl.gov/ANAG/NASA

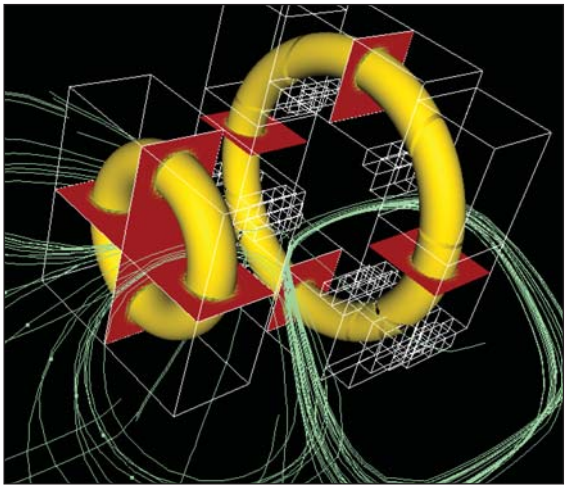

Adaptive mesh refinement (AMR) simulation of two co-rotating vortex rings (indicated by the yellow rings). Green lines show streamlines of the flow, while white boxes show where finer resolution is used to better capture details of the flow. Coherent vortices like those shown here are the principal means for mixing fluids in microgravity environments.

ray burst fireballs and afterglows. Project participants include NASA GSFC, University of Illinois, University of Chicago, SUNY Stony Brook, and the Universities Space Research Association. Details of this research can be found at: www.ibeam.org

The Parallel Virtual File System Version 2

PVFS is a 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. PVFS serves as both a testbed for parallel I/O research and as a freely available, production-level parallel file system for use in the cluster community. Development and support is carried out with collaboration between Clemson University, GSFC, and Argonne National Laboratory. For more information, visit: www.parl.clemson.edu/pvfs

Coven Software Development Framework

Coven is a framework for component-based collaborative software development. Its goals are to provide extensibility and usability, and to produce high-performance code. Extensibility is achieved through pluggable modules representing individual algorithms. Modules are connected to form a dataflow graph, created with assistance from a graphical user interface. Environments using this framework can build additional graphical components, which tailor to the specific needs of that problem domain (such as dataset visualizers, performance analyzers, debuggers, database query tools). Coven has been used to build environments for satellite remote sensing, n-body problems, molecular dynamics problems, and CFD fluid flow and heat transfer problems. For more information, visit: www.parl.clemson.edu/Coven

Beowulf/Mini-Grid System Software

A mini-grid is a campus-wide version of a computational grid, featuring a private switched network interconnecting grid resources where performance is not limited by Internet connections. All resources on the Clemson grid are Beowulf/Linux clusters. The Scyld Beowulf Linux OS has been modified to support the mini-grid with the addition of Balloc, the Beowulf node Allocator, which allows large jobs to transparently "borrow" nodes from adjoining clusters. Traffic for jobs on borrowed nodes runs through a private network, which means performance can be predicted. A single job can use all processors in the grid. The software project is funded by Goddard. Details are on the web at: www.parl.clemson.edu/minigrid

Cost-Effective Advanced Computing Technologies

Researchers at NASA are working on the use of reconfigurable logic to provide ultra-high-performance computing in both space-based and ground-based systems. Techniques include implementing applications directly in reconfigurable hardware, implementing specialized assist logic in reconfigurable hardware, and using reconfigurable logic blocks inside of an otherwise standard microprocessor. Information on this NASA/Clemson University collaboration is available at: www.parl.clemson.edu

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

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 results of NASA Earth and space science research.

Sci-Interactives are web-based applications that let the science speak for itself. These tools give people access to the real equations and data behind NASA science with intuitive, graphically sophisticated tools and games.

Current projects include story development and Sci-Interactives for *Discover* Magazine's website. Tools explore the collapse of rotating interstellar clouds and the creation of aurora in the Earth's magnetosphere. Details of NASA's latest Earth and space science research can be found at www.truth-n-beauty.com

Simulation screen for a collapse and rotation Sci-Interactive.

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.

Shuttle Radar Topography Mission Science Processing on a Supercomputer

NASA researchers have developed software to process raw data acquired by the Shuttle Radar Topography Mission (SRTM), as part of an effort to generate a global digital elevation model database. The large amount of data required to form the mosaic was processed on *Alhena*, the SGI Origin 2000 supercomputer at JPL, using 32 processors in parallel. Digital elevation models are used for scientific applications in geophysics, geology and hydrology, as well as practical applications, such as highly detailed 3-D maps for civilian and military user communities. Information about this project, co-sponsored by NASA and the National Imagery and Mapping Agency, can be viewed at: http://photojournal.jpl.nasa.gov.

Dusty Rings: Signposts of Recent Planet Formation

Every planetary system forms in a thin disk of gas and dust around a young star. Web-based animations show the evolution of bright rings of dust in a disk where planets form. Planets grow from collisions and mergers of smaller bodies, called planetesimals, embedded in a disk of gas and dust in orbit around the star. Each visualization starts with the same initial conditions, roughly 10 billion billion planetesimals 1-100 meters across, embedded in a disk that extends from the orbit of Neptune to the outer edge of the Kuiper belt. The colors represent the intensity of the dust grains from low (dull reds) to high (blue/white). Individual ani-

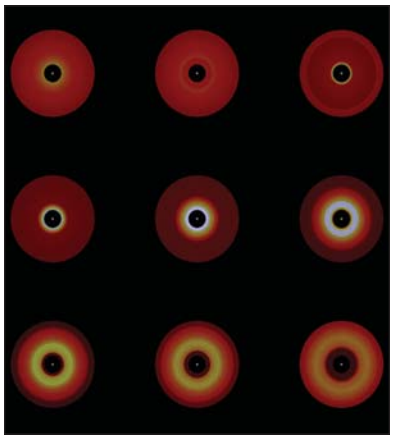

Estimating the Circulation and Climate of the Ocean: The Circulation Pathway of Subtropical-Tropical Exchange

JPL researchers are investigating mechanisms of subtropical-tropical water mass exchange in the Pacific Ocean, using an animation of a simulated passive tracer and its adjoint tracer. The forward passive tracer and adjoint passive tracer can be identified as describing where the tracer-tagged water mass goes to and where it comes from, respectively. The exchange plays a central role in the decadal change of *El Niño*, and is examined by explicitly tracing the water of the eastern Equatorial Pacific. In contrast to previous theoretical expectations, the exchange is found to be an "open-circuit," in which the origin is distinct from the destination. Temporal variability in ocean circulation is found to be fundamental in determining the pathway of the water mass, "short-circuiting" the otherwise circuitous route. Details of this research are available on the Internet at: http://ecco.jpl.nasa.gov/external

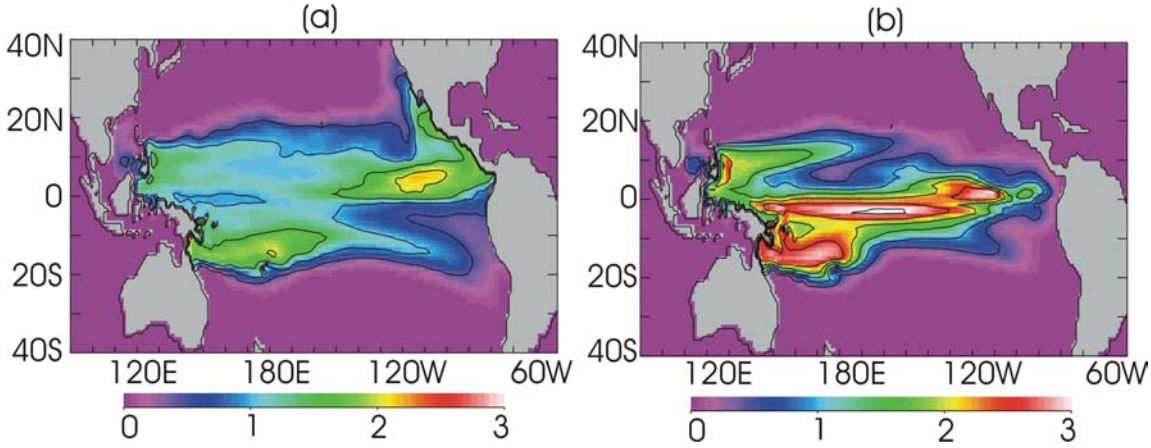

The image at left illustrates the average distribution of the water mass (depth integrated content) that reaches the eastern Equatorial Pacific five years prior to its terminal instant. The right panel shows the same data, except when the time-mean circulation is assumed. The difference between the two illustrates the significance of the circulation's temporal variability.

mations depict a passage of time of 2 billion years. In the "weaker bodies" animation, the bodies fragment easier than the stronger bodies, and thus produce more dust and brighter rings when they collide. The figure shown on page 8A depicts nine snapshots of the disk in one calculation. The bright point at the center is the star. The inner edge of the disk is at the orbit of Neptune. The outer edge is at the outer edge of the Kuiper Belt in our solar system. The top row of frames corresponds to a time, t = 0, 3, and 10 million years (Myrs); the middle row is t = 30, 100 and 300 Myrs; the bottom row is t = 600 Myrs, 1 billion years (Gyrs) and 2 Gyrs. Further details can be found at: http://cfa-www.harvard.edu/~kenyon/pf/sp/index.html

Using Hollywood for Science: Animation of a Proposed NASA Mission using Maya 4.0

The Satellite Test of the Equivalence Principle (STEP) was a proposed NASA small explorer mission which would test the principle that all masses, regardless of size or composition, will experience the same acceleration in a given

Continued on page 10A

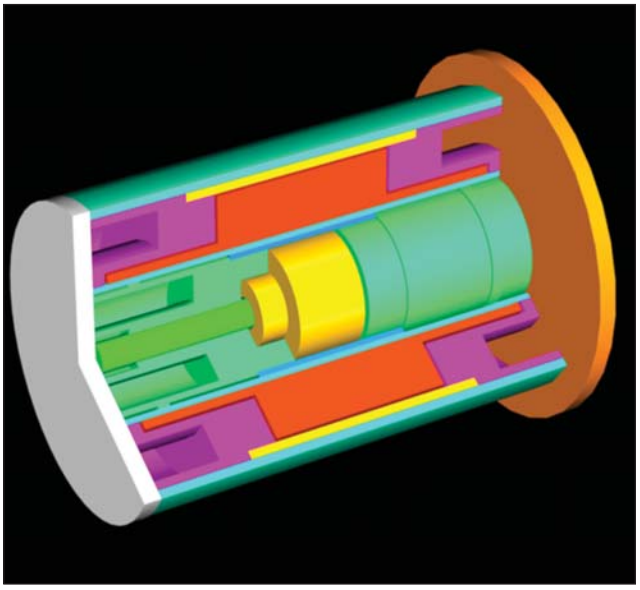

Using Hollywood for Science: Animation of a proposed NASA mission using Alias Wavefront's Maya 4.0 software.

3-D Adaptive Grid MHD Simulations of the Global Heliosphere

The visualizations show the interaction of the interstellar medium and the solar wind. The solar system is moving through the local interstellar medium. This complex system involves three different discontinuities: the termination shock, the heliopause, and the bow shock. In the visualization, the interstellar wind flows from the left and impinges on the solar wind. The sun is located at the center.

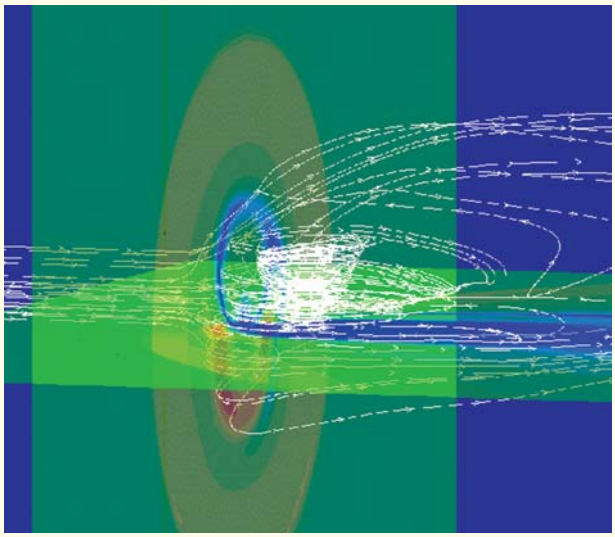

The code uses BATS-R-US, developed at the University of Michigan. It is a 3-D magnetohydrodynamics adaptive grid code. It uses an adaptive Cartesian grid made of a rectangular block of cells. The solar wind and the interstellar wind are modeled, including both magnetic fields and the plasma in a self-consistent way. The simulation used 4.5 million cells and 11 levels of refinement. The minimum size of the cells is 0.5 Astronomical Units (AU), and the maximum size is 36 AU.

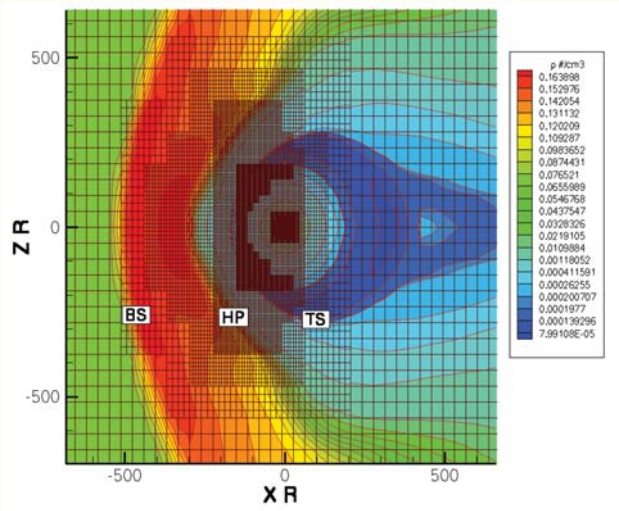

Above: *This graph shows the contours of the density. In this figure the interstellar medium flows from the left, with the Sun at the center of the figure. Black denotes different refined areas of the mesh. Also indicated: the location of the Bow shock (BS), Heliopause (HP), and the Termination shock (TS).*

Left: *Contours of the magnetic field are shown with blue representing the minimum intensity of the magnetic field and red depicting maximum intensity. The white lines indicate magnetic field lines. At the center, the solar magnetic field (the Parker spiral), can be seen with the Sun at the center. The interstellar medium flows from the left.*

Continued from page 9A

gravitational field. This principle is a fundamental assumption in Einstein's General Theory of Relativity. The movie contains a simulation of the relative motion of two test masses in the experiment in orbit about the Earth when the Equivalence Principle is violated, followed by an animation of the assembly of the various parts of the experiment. The modeling, animation, and rendering of the movie were performed using Alias Wavefront's Maya Complete software package, and took approximately three hours on eight processors of JPL's SGI Origin 2000 supercomputer. Further details on this research project can be found at:
http://sc.jpl.nasa.gov/success_stories/slides/slide44.html

Common Component Architecture: A Demonstration of the CCA Forum Technologies

The Common Component Architecture (CCA) Forum is a group of researchers from national labs and academic institutions committed to defining a standard component architecture for high-performance computing.

In the last year, researchers began testing the Common Component Architecture (CCA), designed for large, collaborative, high-performance computing applications. The testing is designed specifically for its suitability to NASA applications, such as climate simulations. Further details can be found at: http://pat.jpl.nasa.gov/public/projects

NASA Participation: Technical Papers, Tutorials, and Panel Discussions

Several NASA scientists (both contractor and civil service personnel) are participating on SC2002 committees, panels, tutorials, and other conference events.

Paper Presentation

Collaborative Simulation Grid: Multiscale Quantum-Mechanical/Classical Atomistic Simulations on Distributed PC Clusters in the USA and Japan

Hideaki Kikuchi (Louisiana State University), Rajiv K. Kalia, Aiichiro Nakano, Priya Vashista (University of Southern California), Hiroshi Iyetomi (Niigata University), Shuji Ogata, Takahisa Kouno (Yamaguchi University), Fuyuki Shimojo (Hiroshima University), Kenji Tsuruta (Okayama University), and Subhash Saini (NASA Ames).

A multidisciplinary, collaborative simulation has been performed on a grid of geographically distributed PC clusters. The multiscale simulation approach seamlessly combines atomistic simulation based on the molecular dynamics (MD) method and quantum mechanical (QM) calculation based on the density functional theory, so that accurate but less scalable computations are performed only where they are needed. The multiscale MD/QM simulation code has been grid-enabled using a modular, additive hybridization scheme, multiple QM clustering, and computation-communication overlapping. The 'gridified' MD-QM simula-

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.

Biological Metaphors for Feature Extraction in Large Datasets

This demonstration describes a new visualization metaphor for quickly extracting and visualizing features in large datasets. The data swarming metaphor is a biologically based collection of agents that move through space interacting with local data. Although each agent follows a few simple rules, the collection of agents can find complex patterns in the dataset. In particular, this demonstration portrays a general agent architecture that can be used to create unique visualization tools as well as support common visualization techniques. The data swarming paradigm is applied to computational fluid dynamics visualization (below), and is implemented as a virtual environment that runs on immersive displays.

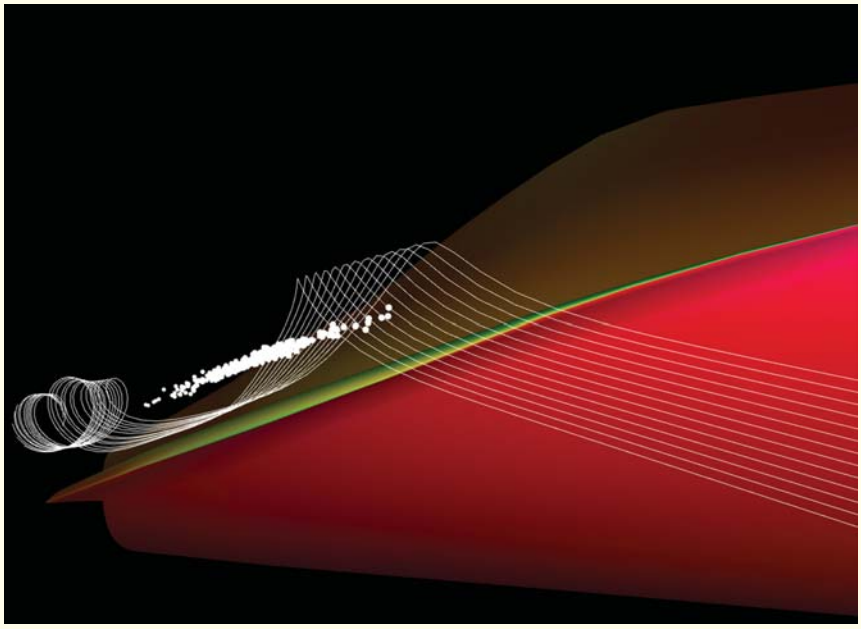

tion code has been used to study environmental effects of water molecules on fractures in silicon. A preliminary run of the code has achieved parallel efficiency of 94 percent on 25 PCs distributed over three PC clusters in the United States and Japan. A larger test involving 154 processors on five distributed PC clusters is in progress.

Tutorial S11

November 18, 8:30 a.m. – 12:00 p.m.

Requirements Analysis for Scientific Applications

Scientific applications are increasingly complex and collaborative. In order to coordinate, plan, and validate scientific software, it is useful to gather requirements before development begins. In this tutorial, developers of the Earth System Modeling Framework explore how requirements analysis can improve development efficiency and reduce misunderstandings and defects. They will look at how requirements collection can be implemented for distributed collaborations; review sample requirements documents and templates; and discuss how the requirements process impacts other aspects of software development, such as design and verification. Throughout this process they will emphasize sensitivity to the unique aspects of the research-oriented HPC computing environment. More information about the Earth System Modeling Framework is at: http://www.esmf.ucar.edu

Tutorial S3

November 16, 8:30 a.m. – 5:00 p.m.

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

Rob Armstrong, Jaideep Ray (Sandia National Laboratories), David Bernholdt, Wael Elwasif, Jim Kohl (Oak Ridge National Laboratory), Lori Freitag, Lois McInnes, Boyana Norris (Argonne National Laboratory), Dan Katz (NASA–Jet Propulsion Laboratory), Gary Kumfert (Lawrence Livermore National Laboratory), Craig Rasmussen (Los Alamos National Laboratory)

This tutorial will introduce the Common Component Architecture (CCA) at both conceptual and practical levels. Components are software objects that encapsulate useful functionality and interact with other components only through well-defined interfaces. Component-based approaches to software development, including the CCA, provide a means to manage the complexity of large-scale software applications and facilitate the reuse and interoperability of code.

The component concept is widely used in the commercial and business world (e.g., CORBA, COM, and Enterprise JavaBeans) and is popular in visualization as well (e.g., Data Explorer), but none of these environments is well suited to use in high-performance computing environments. The CCA was designed specifically with the needs of high-performance scientific computing in mind, including the existence of large bodies of existing code that cannot easily be rewritten. The tutorial will cover the concepts of components and the CCA in particular; the tools provided by the CCA environment – including the Babel system for language interoperability; the creation of CCA-compatible components; and their use in scientific applications. A combination of traditional presentation, and live demonstration will be used during the tutorial. After the session, participants will be able to download the tools and software to reproduce the live examples in their local environment. For more information, see:
www.cca-forum.org/tutorials/sc02.html.

Panels

November 22, 8:30 – 10:00 a.m.

Are Designer Supercomputers an Endangered Species?

Moderator: *Aruna Ramanan (IBM)*
Panelists: *Thomas Sterling (Center for Advanced Computing Research, California Institute of Technology and Jet Propulsion Laboratory); Gita Alaghband (University of Colorado, Denver); Jamshed Mirza (IBM Corporation); Tadaski Watanabe (NEC Corporation); and Candace Culhane (NSA)*

At the dawn of the high-performance computing era, all supercomputers were "designer machines." Every aspect of these machines was designed specifically to lend the highest performance level achievable, using the state of the art technology available at the time. While the attention to design still remained, the focus changed from exploiting fine-grained parallelism using vector processors to coarse-grained parallelism using "massive parallel processors" or MPPs. The emergence of MPPs changed the focus from the processor to the interprocessor communication network. The systems were still designer systems. Then came the shift to specially designed clusters.

More recently, "Components Off the Shelf" or COTS clusters, have emerged in a big way. Strides in the cyber communication infrastructure, together with the narrowing gap between the computing world and the networking world has opened up the possibility of grid computing and peer-to-peer computing. The Seti@home project has demonstrated the power of utilizing distributed resources. These developments have led many to believe less and less in designer supercomputers and more and more in constructing high-performance systems using widely available components used in mainstream computing. But the emergence of the Japanese Earth Simulator as the top performer has stirred up the high-performance computing community. This panel will explore the strategies needed to transform terabytes of information into insights, and examine the challenges that need to be addressed by the two different approaches to supercomputing.

SC2002

BOOTH MAP

1 Immersive Workbench
- Interactive Large Screen Visualizations
- IBEAM: Interoperability-Based Environment for Adaptive Meshes
- Biological Metaphors for Feature Extraction in Large Datasets

2
- Coupling Engineering Applications on the Information Power Grid: A CORBA-Based Approach
- 128-Node Information Power Grid Cluster

3
- Real-Time Ray Tracing with the University of Utah
- Programmable Visualization of EOS DAAC Data

4
- Estimating the Circulation and Climate of the Ocean: The Circulation Pathway of Subtropical-Tropical Exchange
- Dusty Rings: Signposts of Recent Planet Formation
- Shuttle Radar Topography Mission (SRTM): Science Processing on a Supercomputer
- 3D Adaptive Grid: MHD Simulations of the Global Heliosphere
- Using Hollywood for Science: Animation of a Proposed NASA Mission Using Maya 4.0
- Common Component Architecture (CCA): A Demonstration of the CCA Forum Technologies

5
- Information Power Grid (IPG)
- IPG Job Manager
- IPG Resource Broker
- Performance Prediction
- AeroDB: Automated CFD Parameter Studies on Distributed Parallel Computers
- Tool Agent Framework-Java (TAF-J) - Connecting Aerospace Analysis Tools with the Information Power Grid

6
- Multi-Level Parallelization and Optimization Environment

7
- Cart3D: A Package for Automated Grid Generation and Aerodynamic Database Creation
- Atomistic Simulations on USA and Japan Grid

8 & 9
- Beowulf/Mini-Grid System Software
- Coven Parallel Programming Framework
- The Parallel Virtual File System Version 2
- Reconfigurable Computing Technology

10
- Simulation-Based Medical Planning for Vascular Surgery

11
- Ultra-High-Resolution Astronomy: Phasing of Arrays of Formation Flying Spacecraft
- Modeling the Earth's Atmosphere at the Data Assimilation Office
- Sci-Interactives: Innovative Science Outreach from Truth-N-Beauty Software and NASA CT

12
- Chombo Framework for Block-Structured AMR Applications
- IBEAM: Interoperability-Based Environment for Adaptive Meshes

13
- The hyperwall: A Multiple Flat-Panel Display Wall for High-Dimensional Visualization

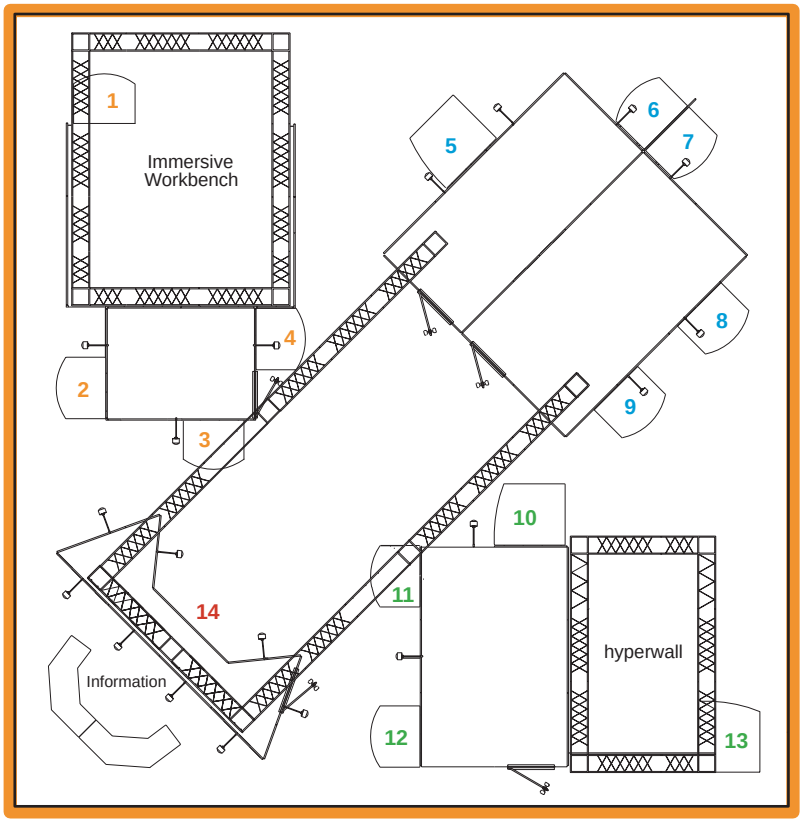

14 Video and Poster Wall
- SC2002 Conference Opening Video
- NASA Research Highlights Video
- Numerical Propulsion System Simulation (NPSS)
- Supercomputing at Goddard Space Flight Center
- Computational Science at Goddard Space Flight Center
- MU-SPIN: Extending Computing Opportunities to Minority Institutions
- NASA/SGI Develops 1,024-Processor Coherent Shared-Memory Origin 3800
- Capability Computing at the NASA Ames Research Center: Solving the Toughest Computational Problems
- NASA-DeBakey Ventricular Assist Device (VAD) - 2001 NASA Commercial Invention of the Year
- HiMAP: High-Fidelity Multidisciplinary Analysis Process
- Cart3D: 2002 NASA Software of the Year
- Computing, Information, and Communications Technology Program (CICT)
- Computing, Networking, and Information Systems Project (CNIS)
- Consolidated Supercomputing Management Office (CoSMO)
- Information Power Grid (IPG) Accomplishments - 2002