Winter 2003

Gridpoints

The Quarterly Publication of the NASA
Advanced Supercomputing Division

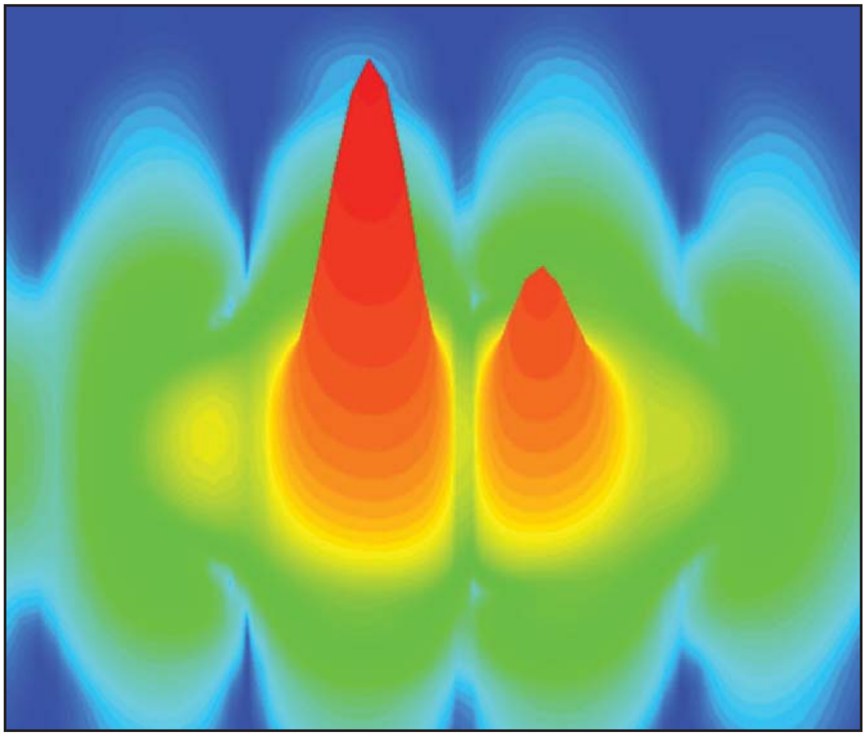

NAS researchers have filed a patent for using coupled lasers to achieve ultrafast modulation and switching in optical networks, one of the many accomplishments highlighted in the division's year-end review. **See page 8**

FEATURES

Features

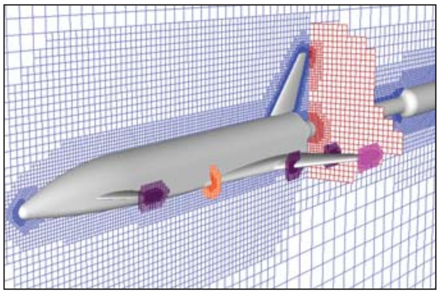

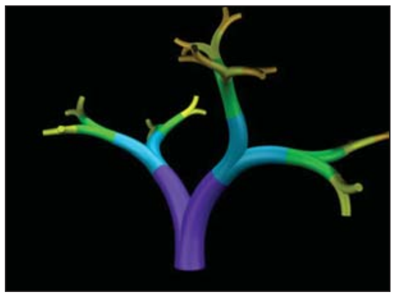

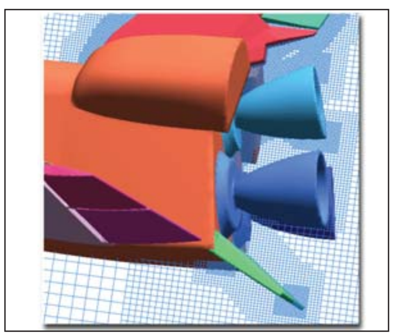

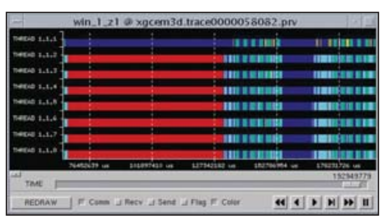

Also in this issue:

On The Cover:

NAS scientists have proven the effectiveness of using coupled lasers to achieve ultrafast modulation and switching in optical networks. Their method uses two coupled vertical-cavity surface-emitting lasers (VCSEL), each with a diameter of about ten microns, one-tenth the thickness of a human hair. (NASA/Cun-Zheng Ning)

Gridpoints

The Quarterly Publication of the NASA Advanced Supercomputing Division

Vol. 4, No. 1 • Winter 2003

Executive Editor: John Ziebarth

Editor: Nicholas A. Veronico

Writers: Holly A. Amundson, Jill A. Dunbar, Julie Jervis

Graphic Designer: Sue Kim

Illustrator: Cliff Williams

Photographer: Michael Boswell

Contact the Editor:

Nicholas A. Veronico
Gridpoints
M/S 258-6
Moffett Field, CA 94035-1000
Tel. (650) 604-3046 • Fax (650) 604-4377
nveronico@mail.arc.nasa.gov

The staff of *Gridpoints* would like to thank the following people for their assistance in preparing this publication:
Jan Aikins, Michael Aftosmis, Dick Anderson, Robin Amundson, M.P. Anantram, Bryan Biegel, Chris Buchanan, Galina Chaban, Christopher Dateo, Suhrit Dey, Heinz Erzberger, Graham Fletcher, Robert Frampton, Chris Gong, Peter Goorjian, Chris Henze, Tom Hinke, Winifred Huo, Dochan Kwak, Henry Jin, Gabriele Jost, Randy Kaemmerer, John Kitowski, Timothy Lee, Tony Lisotta, Nateri Madavan, Nagi Mansour, Cun-Zheng Ning, James Ramage, Marcia Redmond, Michael Rogers, Stuart Rogers, Subash Saini, Tim Sandstrom, David Schwenke, Chung-Lin Shie, Deepak Srivastava, Astrid Terlep, William Thigpen, Joseph Totah, Dunyou Wang, Charles Wiseman, Alan Wray, Jerry Yan, Maurice Yarrow

Gridpoints is a quarterly publication of the NASA Advanced Supercomputing Division at NASA Ames Research Center. The division's staff includes government employees and contractors from Advanced Management Technology Inc., Computer Sciences Corp., Cray Inc., Eloret, Raytheon Co., and SGI.

Subscribe to *Gridpoints*

Subscriptions to *Gridpoints* are complimentary. Please send e-mail with your name, company name, title, address, and zip code to:
gridpoints@nas.nasa.gov
or subscribe online at:
www.nas.nasa.gov/gridpoints

Note: We do not mail outside the United States. Foreign subscribers can access *Gridpoints* on the web in PDF format at:
http://www.nas.nasa.gov/gridpoints

News From NAS

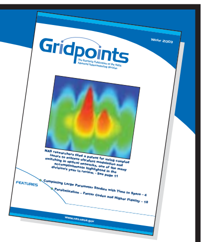

NATO Lecture Series Provides Guidance for Future UAV Design

Nearly 100 researchers from around the world gathered in Palo Alto, Calif., for the NATO (North Atlantic Treaty Organization) lecture series, hosted by NASA Ames Research Center on Oct. 29–30. Sponsored by NASA's Computing, Information, and Communication Technologies Program, the event was a unique opportunity for researchers to discuss their current work in the area of uninhabited air vehicles (UAVs).

UAVs are remotely or self-piloted aircraft that can carry cameras, sensors, communications equipment, or other payloads. Both NATO and NASA have applications in UAVs and have a strong interest in their design and development to meet national security requirements and NASA mission needs.

This year's series, "Applications, Concepts and Technologies for Future Tactical UAVs," highlighted topics including integrated mission systems design concepts and alternatives, key enabling technologies (sensors, situation awareness, communication), and autonomous operations. "Many of the issues that we have addressed in the past and that we continue to explore for the future here at NASA Ames Research Center are the same issues

Continued on page 2

NAS Mission

To lead the country in the research, development, and delivery of revolutionary, high-end computing services and technologies, such as applications and algorithms, tools, system software, and hardware to facilitate NASA mission success.

From The Division Chief

This past year has been an exciting time in the NAS Division. Our team made tremendous strides in computational science and computer technology in support of NASA missions, science, and engineering milestones. Researchers within our division received both the NASA Software of the Year award and the NASA Commercial Invention of the Year Award, and a number of patents were issued (see page 8). In addition, researchers in our Applications Branch received the first annual Ames Research Center Information Sciences and Technology Directorate "Technology Infusion Award."

Our computational scientists have demonstrated the highest performance on any U.S. computer architecture – by any vendor – on the 1,024 processor SGI Origin 3800 system, *Chapman*, for Earth science and aerospace codes. This illustrates that when the goal is high-performance, or what is referred to as "capability computing," the common shared-memory architecture is a tremendous tool for engineers and and the scientists using those architectures.

In November, NAS scientists showcased the division's latest supercomputing research and technologies at the SC2002 Conference of High Performance Networking and Computing in Baltimore. The division's new "hyperwall," a scientific visualization system, attracted large crowds (see *Gridpoints*, Fall 2002, page 11). Also in the NASA exhibit were researchers from the University of Utah who teamed with NAS scientists to demonstrate state-of-the-art remote, interactive ray-tracing. Computations were performed using more than 1,000 of *Chapman's* processors, and the images were displayed on the show floor in Baltimore. Ray-tracing is a powerful but extremely compute-intensive visualization technique. Due to excellent performance scaling of the ray-tracing code on *Chapman*, the Utah researchers were able to produce and display interactive (that is, multiple frames per second) visualizations of multi-gigabyte datasets, including an eight-gigabyte high-resolution scan of a human male thorax (data from the Visible Human Project, National Library of Medicine), and a 64-gigabyte terrain dataset, with a resolution of more than 17 billion polygons.

One of the biggest challenges of this demonstration was overcoming the inherent problems associated with TCP/IP when moving large amounts of real-time data across the country. Working closely with the Utah scientists, NAS network engineers tuned a variety of parameters at the kernel, system interface, and network layers to maximize throughput and minimize delay and delay variance (jitter). Once tuned, throughput was increased six-fold, delay was reduced by 20 percent, and jitter was reduced by 50 percent, enabling scientists to manipulate extremely large datasets in real time from 2,800 miles away.

On a final note, I have enjoyed more than five years with NASA Ames, both as Director of the Consolidated Supercomputing Management Office and as Deputy Division Chief and acting Chief of the NAS Division. I will soon be joining Los Alamos National Laboratory as head of the Advanced Computing Laboratory.

I would like to take this opportunity to thank the NAS Division and NASA for the opportunity to meet and work with many of the finest scientists and staff, as well as collaborators both in industry and academia, I have ever had the privilege to know and serve. Thank you.

John Ziebarth

Continued from page 1

being discussed by those of you attending this NATO event," said Jan Aikins, acting deputy director of Ames' Information Sciences and Technology Directorate. "Our researchers in the Information Sciences and Technology, Astrobiology and Space Research, and Aerospace directorates are dealing with both development and applications of autonomous control technologies for science, with human-machine interfaces for ground control systems, and with policies and regulations necessary for operating in the national airspace, among

Computational Modeling to Understand Breast Cancer's Behavior

On November 21, 2002, Eastern Illinois University professor of mathematics Suhrit Dey returned to NAS to present his latest breast cancer research findings (see *Gridpoints*, Fall 2001, page 4). Dey is developing a system of mathematical equations to model the effects of chemotherapy, radiation therapy, and immunotherapy on the growth and spread of breast cancer cells in the human body. Breast cancer currently affects more than two million women in the United States. Dey believes the effects of breast cancer can be managed by understanding how the disease behaves. His predictions on the disease's interactions within the body are based on mathematics and computer simulations.

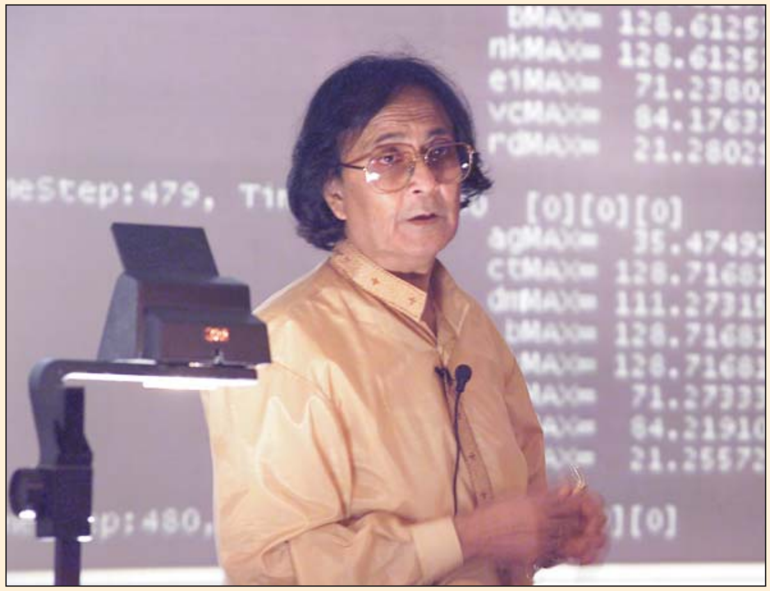

During the November 21, NAS New Technology seminar, Suhrit Dey, mathematics professor at Eastern Illinois University, presented several new equations for his 3-D dynamic animation used to model breast cancer in the human body. (NASA/Tom Trower)

Among Dey's recent work is a series of equations accounting for the growth and spread of cancer, which are used to model how these factors affect non-cancerous cells in the body. The new equations were integrated into Dey's 3-D dynamic animation, developed in part by his son Charlie Dey, also at the University of Illinois. Dey's animation enables him to rotate the model in any direction and observe movement of individual cancer cells in the body from all angles, especially in the breast area. Working with John Koontz of SkyNetworks in Champaign, Illinois, Dey has improved the code for generating his animations. The code requires massively parallel computations in the message-passing interface (MPI) environment.

Strength in Numbers

An important part of Dey's work has been traveling the globe – from Scotland, to Italy, to India, to form collaborations with doctors and deliver invited talks. Discussing his research findings with doctors provides Dey with feedback about his mathematical models. Most recently, Dey has been working closely with Charles Wiseman, chief oncologist at Los Angeles' St. Vincent Hospital. Wiseman has supplied the Eastern Illinois University professor with data that supports the predictions of his three-dimensional model.

Dey has focused his research on the cause of recurring breast cancer cases. "The mathematical model reveals that when the main tumor is gone, that does not mean that the cancer is gone. The cancer cells could still be alive and gaining strength, especially if the immune system gets weaker," explains Dey. To determine what is needed to ward off recurring cases of cancer, and to increase the number of fighter T-cells, B-cells, and Natural Killer cells in the body, Dey is experimenting with variables in his mathematical models such as stress, depression, and side effects of medications that weaken the immune system.

To reach his long-term goal of putting most, if not all, of the human body's cells into a single configuration, Dey will require more computer power. Using a Pentium 4 machine, it currently takes two days to complete his calculations. Dey is confident that he could complete them in under 20 minutes on a multi-processor, shared-memory supercomputer such as those housed at the NAS facility.

Editor's Note: *Suhrit Dey's breast cancer research was initiated after receiving grants from Bob Augustine, dean of the Eastern Illinois University graduate school, Mary Anne Hanner, dean of the university's College of Sciences, and Bud May, director of the university's Grants and Research Department.*

others. We are pleased to be hosting this event, and are delighted to be working with you to further our common research goals for UAVs."

Among the panel of speakers were three members of the NATO lecture series team: John Kitowski, Robert Frampton, and James Ramage. Kitowski spoke on the economical and technological challenges associated with UAVs, while Frampton discussed their increasing role in support of combat missions. During his talk on the considerations for future UAV research and development, Ramage explained that further automation research is required to reach the design and performance goals of the UAV. In addition, Ramage believes that stringent safety regulations and technology maturity of uninhabited aircraft have, and will, present a huge challenge for UAV designers and developers. Concluding his presentation, Ramage said, "Future R&D directions should emphasize demonstration and validation of key enabling technologies in realistic integrated system operation environments, using mission relevant metrics."

Researchers have many challenges ahead before UAVs can function as fully integrated systems that can execute tasks such as automatically adjusting to the changing tactical requirements of combat. Essentially, UAVs must be able to "think" on their own, be self-directing, or autonomous. Autonomy, a system's ability to perform a task without direction from an external source, reduces and automates the workload, enabling the execution of tasks ranging from vehicle management, based on terrain and weather conditions, to threat avoidance.

In addition to advancing our nation's security, UAVs play a role in meeting NASA missions. Acting Chief of the Ames' Computational Sciences Division Daniel Clancy delivered a talk, "Autonomy for Deep Space Exploration as Relevant to UAVs," in which he described clear autonomy needs of the upcoming Europa mission, scheduled for launch in March 2008. This mission entails the robotic exploration of Europa to seek out evidence of life. Clancy emphasized: "Autonomy is a capability, not some flashy technology – it increases safety and capability."

Heinz Erzberger, senior scientist in Ames' Aviation Systems Division, also delivered a talk highlighting some of Ames' newest air traffic management technologies and decision support tools, which have been fielded with great success. Said Erzberger: "I saw several parallels in UAV automation and in air traffic management automation, especially with respect to human interface and situational awareness issues."

Next year's NATO event on this topic will take place in June 2003. For more information on the lecture series, visit NATO's Research and technology website at:
www.rta.nato.int/Meetings.asp.

NAS Division Tech Reports

Technical reports are one of the ways NAS Division researchers share their work with others in the scientific community. A large number of reports are available electronically in PDF format on the division's website, including an archive dating back to 1989. The most recent reports can be found at:
www.nas.nasa.gov/Research/Reports/techreports.html

NAS Technical Reports for 2003

"Performance Evaluation of Remote Memory Access (RMA) Programming on Shared Memory Parallel Computers," by Haoqiang Jin and Gabriele Jost • NAS-03-001

"NAS Parallel Benchmarks I/O Version 2.4," by Rob F. Van Der Wijngaart and Parkson Wong • NAS-03-002

NAS Technical Reports for 2002

"A Posteriori Error Estimates for Higher Order Godunov Finite Volume Methods on Unstructured Meshes," by Timothy Barth and Mats G. Larson • NAS-02-001

"Hybrid MPI+OpenMP Programming of an Overset CFD Solver and Performance Investigations," by Jahed Djomehri and Haoqiang Jin • NAS-02-002

"Resource Selection Using Execution and Queue Wait Time Predictions," by Warren Smith and Parkson Wong • NAS-02-003

"High-Order Central WENO Schemes for Multi-Dimensional Hamilton-Jacobi Equations," by Steve Bryson (NASA Ames) and Doron Levy (Stanford University) • NAS-02-004

"NAS Grid Benchmarks Version 1.0," by Rob F. Van Der Wijngaart and Michael Frumkin • NAS-02-005

"High-Order Semi-Discrete Central-Upwind Schemes for Multi-Dimensional Hamilton-Jacobi Equations," by Steve Bryson and Doron Levy • NAS-02-006

"NAS Parallel Benchmarks, Version 2.4," by Rob F. Van Der Wijngaart NAS-02-007

"The Efficiency and the Scalability of an Explicit Operator on IBM POWER4 System," by Michael Frumkin • NAS-02-008

"Implementation of the NAS Parallel Benchmarks in Java," by Michael Frumkin, Matthew Schultz, Haoqiang Jin, and Jerry Yan • NAS-02-009

Completing Large Parameter Studies — with Time to Spare

NAS Division researchers create an automated CFD system to run large parameter studies on IPG resources.

When designing an aircraft, every possible flight condition is scrutinized before it leaves the drawing board. Engineers spend countless hours testing thousands of parameters and vehicle configurations. With the evolution of computational fluid dynamics (CFD), a sophisticated computational analysis technique used to predict fluid flows, the time and expense involved in designing air- and spacecraft has been reduced substantially. Still, the CFD method is a labor-intensive activity, using hundreds of hours to complete the calculation of just one geometry under a single flight condition.

Researchers in the NASA Advanced Supercomputing (NAS) Division have automated the CFD process by creating the AeroDB system. AeroDB boils down the testing process of aerospace vehicle design, saving both time and money. The NAS AeroDB team of Stuart Rogers, Edward Tejnil, Michael Aftosmis, Shishir Pandya, Jasim Ahmad, and Neal

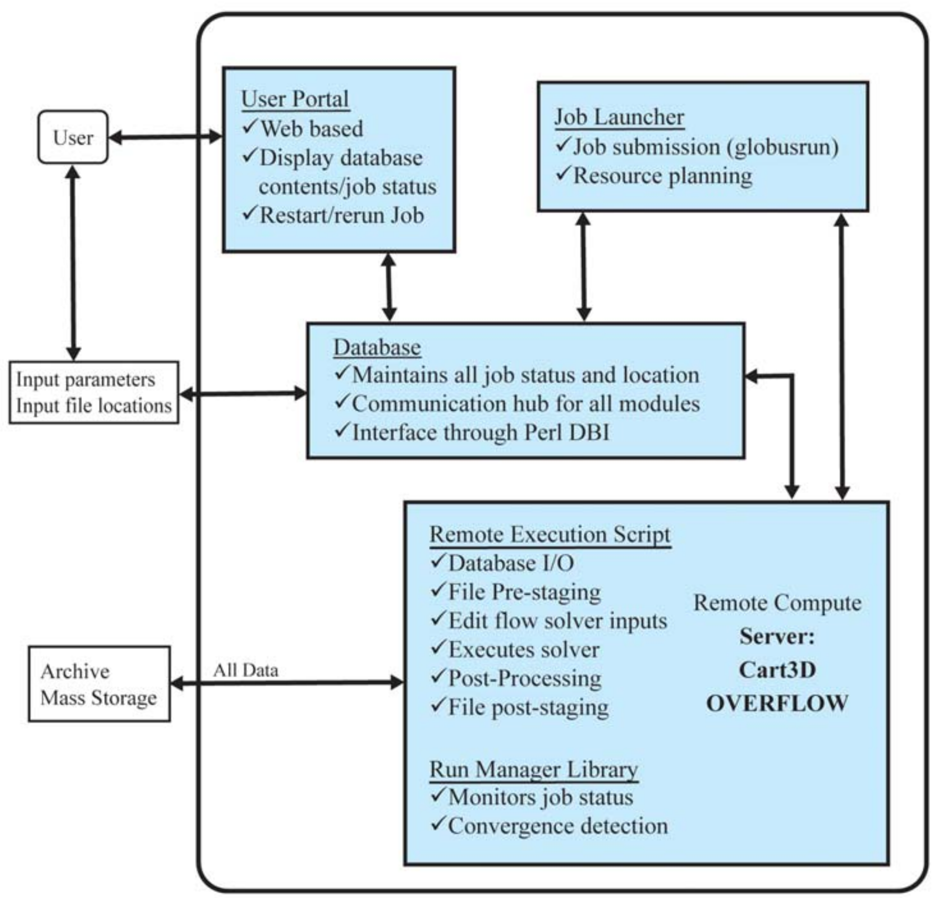

Figure 1: *This diagram represents several main components, or modules (written in the Perl programming language), that make up the AeroDB system. The database is the heart of the system — all communication between components takes place here. The other components include: the job submission script, which automatically creates a unique identification number for each job and provides the user with the interface to specify the case's input files and parameter values; the job launcher module, which is responsible for monitoring the database and launching each job submitted; a remote execution module, which draws on information from the database to run the flow solver; a run manager library, which monitors solver output to determine when the job converges or runs out of time; and a user portal, which serves as the interface for users to view the status of their jobs.*

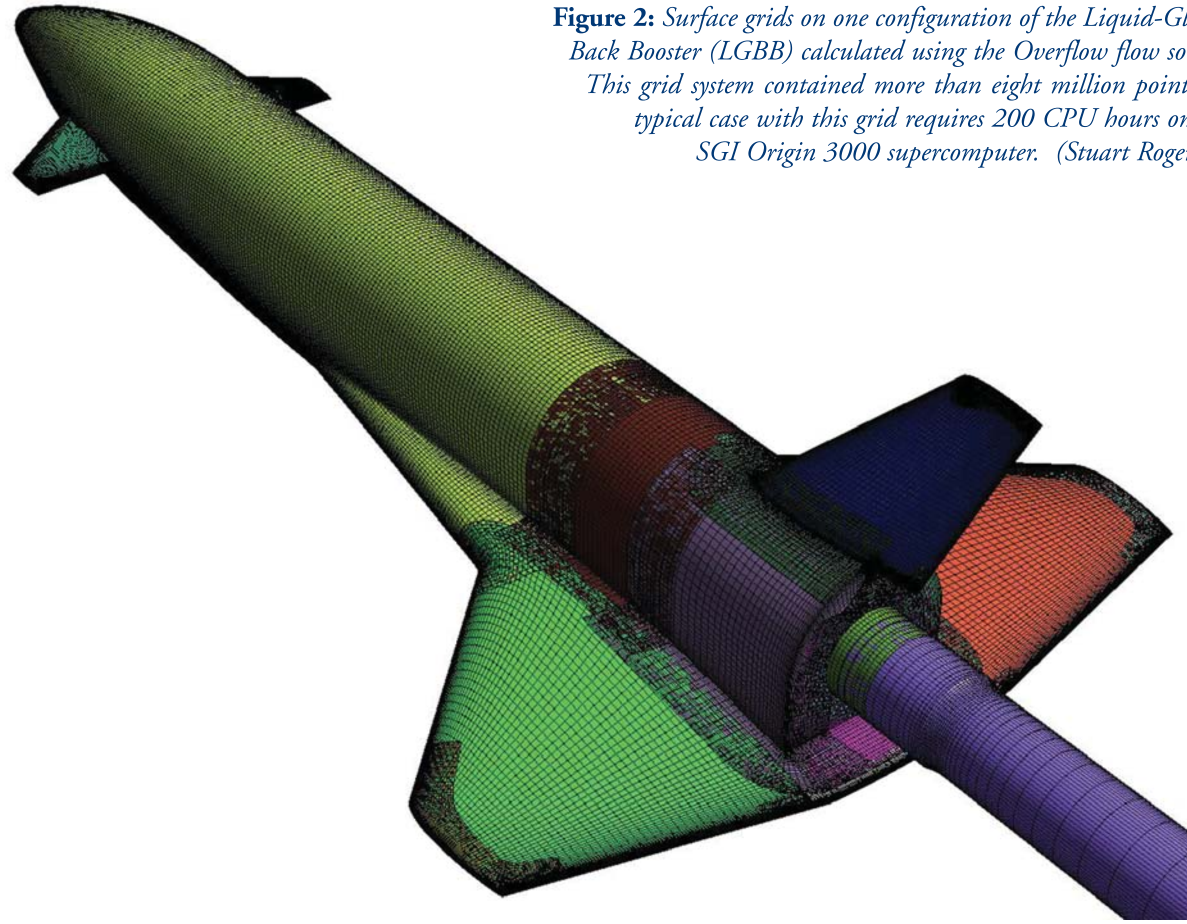

Figure 2: *Surface grids on one configuration of the Liquid-Glide-Back Booster (LGBB) calculated using the Overflow flow solver. This grid system contained more than eight million points. A typical case with this grid requires 200 CPU hours on an SGI Origin 3000 supercomputer. (Stuart Rogers)*

Chaderjian combined their knowledge and experience with CFD programming, the Perl scripting language, and the Overflow and Cart3D flow solvers to create this new software system. The project is supported by the NASA Computing, Information, and Communication Technologies (CICT) Program's Computing, Networking, and Information Systems (CNIS) Project.

Automating the CFD process leaves designers and engineers free to focus on important aspects of vehicle analysis, such as design optimization. Part of this analysis includes the study of flow parameters, including angle of attack, sideslip angle, and Mach number. The process of setting up parameter studies (a collection of computer jobs with a slightly different input for each job) is very labor intensive and error prone. By reducing dependency on the user, AeroDB decreases the amount of error in the analysis. "The AeroDB system significantly simplifies the process of executing many CFD jobs. It is a big step toward automating the process to take the user out of the loop of having to monitor every single job," explains Rogers, senior scientist on the AeroDB project.

Automating the CFD Process

AeroDB is a collection of Perl scripts, a database, and a Web portal (see Figure 1, page 4). The system utilizes resources and tools provided by the Information Power Grid (IPG), NASA's geographically distributed network of computing and data resources, to support its automated system (see *Gridpoints*, Fall 2002, page 4).

The Perl scripts, developed largely by Tejnil, interact with the software tools provided by the IPG to launch, run, and monitor jobs on the grid's network of computers. A job submission script is used to enter basic information about new jobs (such as location of the input files, job size, type of flow solver, and flow parameters), into the database. A job launcher script sends all jobs for execution to the most suitable computer resources. The job launcher uses information provided by the IPG resource broker, which, in turn, obtains data from an information service called the Monitoring and Discovery Service (MDS). The MDS contains information about each IPG resource, including the number of jobs currently running, in order to suggest the optimal place to run a job.

To store information about CFD jobs, AeroDB has its own database. "The database is the heart of the system – the communication hub and repository for all information about the job," explains Rogers. The database contains both static and dynamic job information – unchanging computer host and user account information, and data that is constantly updated as a job is running. "AeroDB's database-centric approach provides the user with a convenient central warehouse for investigating the results of simulations," adds AeroDB team member Michael Aftosmis.

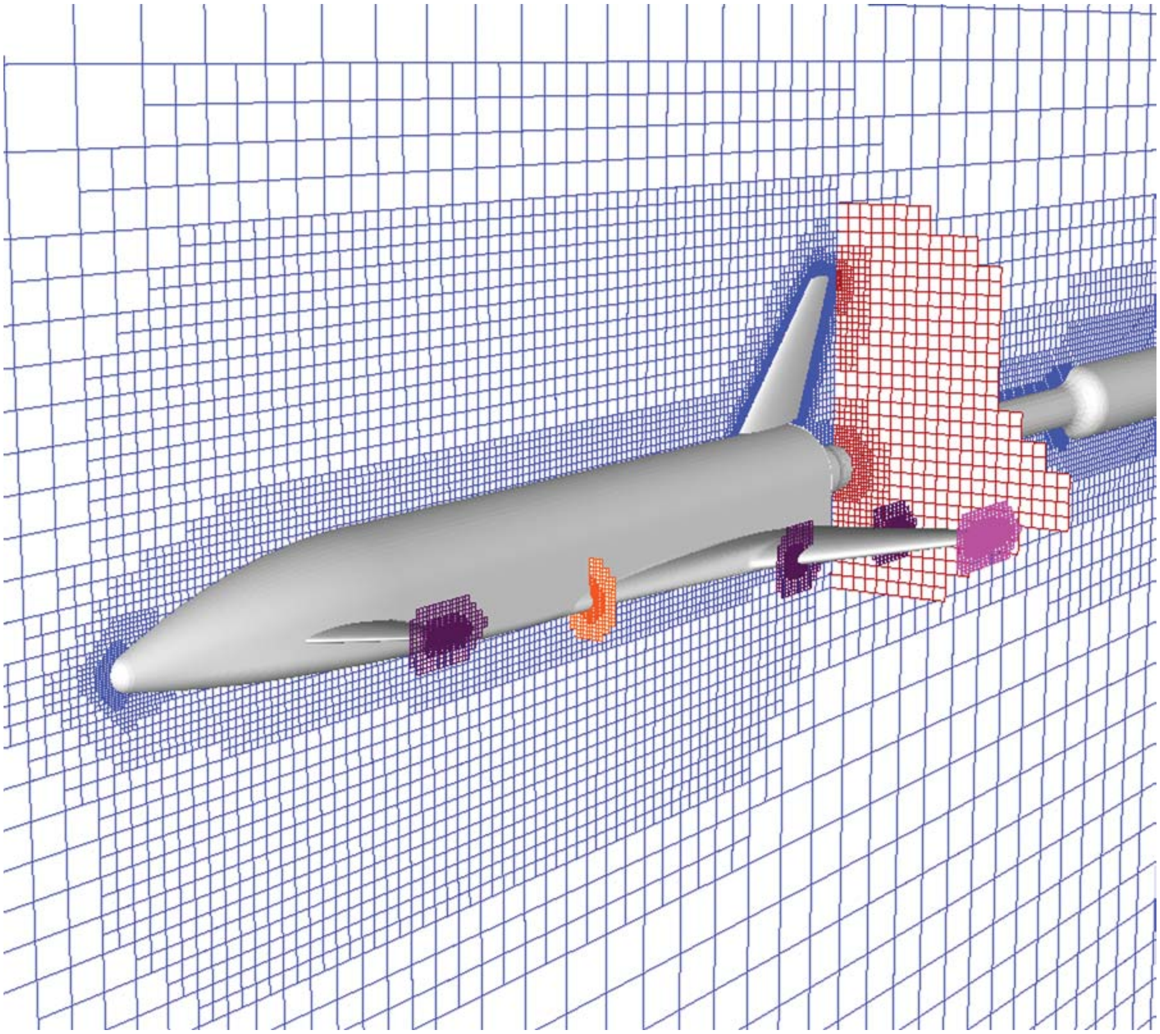

Figure 3: *This figure represents a Cartesian unstructured mesh generated with the Cart3D flow solver, for the same configuration of the Liquid-Glide-Back Booster as the Overflow grids in Figure 2 (see page 5). A Cart3D generated flow solution requires approximately one-tenth of the time; however, accuracy suffers somewhat due to the nature of the equations used in the flow solver, which overlook viscosity.* *(Michael Aftosmis)*

Controlling jobs and accessing information about them is made easy through a Web portal, the interface for the AeroDB system. Through this password-protected interface, users can view the status of each job in the database, and easily rerun and restart jobs as necessary using globusrun (a command used to manage job submission to grid resources). “AeroDB’s Web portal really adds to the convenience of the tool,” explains Aftosmis. It taps the system’s database, allowing the user to check the progress or manipulate jobs in a parameter study, from any Web browser.” To view detailed information about a particular job, the user just clicks on the automatically generated job ID number.

Working with the Flow

To demonstrate AeroDB’s capabilities, two flow solvers were used: Overflow and Cart3D. Overflow is used to solve Navier-Stokes equations, which mathematically describe how air or water flows around an object. When compiling, Overflow expends many CPU hours to complete calculations for even a simple geometry (see Figure 2, page 5). Cart3D solves Euler equations, which ignore viscosity (the resistant forces of a fluid’s flow), so flow can be represented with a lot fewer points; however, accuracy is somewhat lost when analyzing more extreme angles of attack (see Figure 3).

The sample dataset used to demonstrate AeroDB’s ability to automate the CFD process was the Liquid-Glide-Back Booster (LGBB), a conceptual design for a new, reusable launch vehicle. This vehicle is currently being studied under NASA’s Space Launch Initiative (SLI) Program, established to develop technologies to significantly reduce the cost of space transportation and enable future space missions. Some experimental wind tunnel data for the LGBB has already been generated at NASA’s Langley Research Center, Virgina, and has been used to validate CFD results generated using the AeroDB system. "Preliminary comparison between experimental and computed results look encouraging," says Rogers. One configuration of the LGBB geometry included the wing and fuselage, vertical tail, and canards. The team analyzed this flight vehicle over a wide range of flight conditions (see Figure 4).

Testing the Tool

To measure the success of the automated CFD system, several metrics were established: Within seven days, the system must successfully execute 1,000 Cart3D cases and 100 Overflow cases. This metric was surpassed in just three days, and at the end of the seventh consecutive day, 2,863 Cart3D cases and 211 Overflow cases had been completed successfully. These cases were run on 13 different IPG computers distributed over four different geographical locations: NASA's Ames and Glenn Research Centers, the Information Sciences Institute (ISI) at the University of Southern California in Marina Del Rey, and the National Center for Supercomputing Applications (NCSA) in Illinois.

At the same time the AeroDB system was utilizing the IPG's distributed resources, the application helped point out some areas for improvement within the grid environment. "We were helping improve the IPG environment by running our application. I think we've really tested the IPG to the extent where it has never been tested before – it was something that really needed to be done," says Rogers. "Working with the AeroDB team and their application was a valuable exercise. Many lessons were learned that will ultimately improve the IPG software. AeroDB was the first application we [the IPG team] have seen that stressed every aspect of the grid – from its technology to human factor issues such as communication and coordination. The AeroDB team has made a significant and welcome contribution to the NASA grid effort," adds NAS IPG lead Tony Lisotta.

Pushing the Limits

Even though the AeroDB system exceeded the metrics of the demonstration, the group plans to make several improvements, such as making the system accessible through the Web to any NASA user with an IPG certificate. The team also plans to add an automatic error recovery capability for instances when a job has to be restarted. This restart function will prevent the loss of work and valuable computer time. Since AeroDB is highly specialized, it can only accommodate Overflow and Cart3D, but with some additions to the Perl scripts, the system will be able to accommodate other flow solvers, enabling access to other CFD software packages.

In the future, the AeroDB team plans to integrate more of the grid's tools into their system, making it more versatile within the grid environment and a more seamless process for users. Specifically, the IPG job manager will be incorporated into the AeroDB system, shifting the responsibility for file transfer, staging, and interfacing with globusrun to the IPG team. This modification will provide a higher-level interface with these tools. "Incorporating the IPG job manager into AeroDB will not require a lot of changes," explains Rogers. "It should be fairly straightforward to do – our system was created in a modular fashion."

—Holly A. Amundson

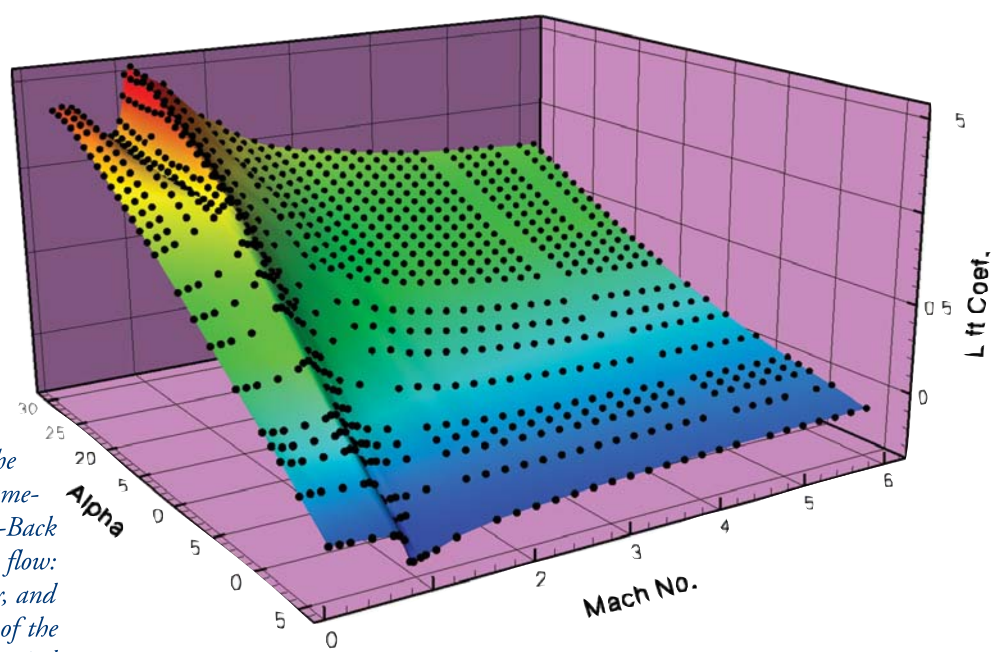

Figure 4: *This plot represents the relationship between several parameters for the Cart3D Liquid-Glide-Back Booster cases run with no sideslip flow: lift coefficient (CL), Mach number, and angle of attack. About 20 percent of the cases run within a seven-day time-period are represented by this plot, each dot representing the CL from one case.*

(Aftosmis/Rogers)

NAS Division in 2002: A Trip Down Shared-Memory Lane

The NAS Division looks back on a successful year of research in diverse areas, and recognition for outstanding scientific achievement.

The year 2002 was marked by many innovative developments from the NASA Advanced Supercomputing (NAS) Division, and major accomplishments in the areas of computational and computer science contributed to the success of NASA missions and science and engineering milestones throughout the agency.

In addition to receiving the NASA Software of the Year Award and the NASA Commercial Invention of the Year Award, the expertise and achievements of NAS scientists were recognized by a number of NASA Space Act Awards, patents, and invitations to sit on national and international oversight panels.

In the area of capability computing, NAS demonstrated the highest performance on any U.S. architecture, by any vendor, on the 1,024 processor SGI system for Earth science and aerospace codes. This accomplishment means that the common shared-memory architecture is proving increasingly valuable for scientists and engineers, and represents a major step forward in innovative and high-performance computing for science and missions throughout the agency.

Here are some of the highlights gathered from the division in 2002:

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

To improve the performance of the DeBakey Ventricular Assist Device (VAD), NAS scientists Cetin Kiris and Dochan Kwak used computational fluid dynamics (CFD) technology developed for the Space Shuttle's main engine fuel turbopump to recommend several design modifications for solving blood flow problems.

The DeBakey VAD is a life-saving miniature heart assist device for human implantation, and the result of a collaborative project between researchers at NASA Ames and Johnson Space Centers and MicroMed Technology Inc., of Houston. Thanks to the improvements suggested by the NAS scientists, the device has now been implanted successfully in more than 170 people, and in April the DeBakey Ventricular Assist Device was named NASA's Invention of the Year for 2001.

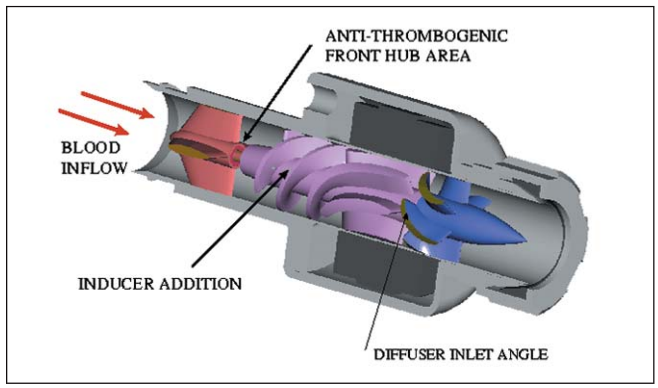

Using computational fluid dynamics analysis, NAS researchers Cetin Kiris and Dochan Kwak found that major design modifications to the DeBakey Ventricular Assist Device were necessary. The result of these changes increased overall efficiency of the device by 22 percent. (MicroMed Technology)

"Our involvement with the development and improvement of the DeBakey VAD makes me extremely proud of what we do at NASA," says Kiris. "This work will improve many lives."

NASA Software of the Year Award: Cart3D

Cart3D, an aerodynamic simulation tool developed by NASA scientists Michael Aftosmis and John Melton, in collaboration with Marsha Berger of the Courant Institute in New York, won NASA's prestigious Software of the Year award for 2002. The Cart3D software package provides

Cart3D is used to automate the grid generation process, shown at right, on the Space Shuttle orbiter. The software suite has significantly contributed to the NASA mission, and was awarded the agency's Software of the Year award in June. (Michael Aftosmis)

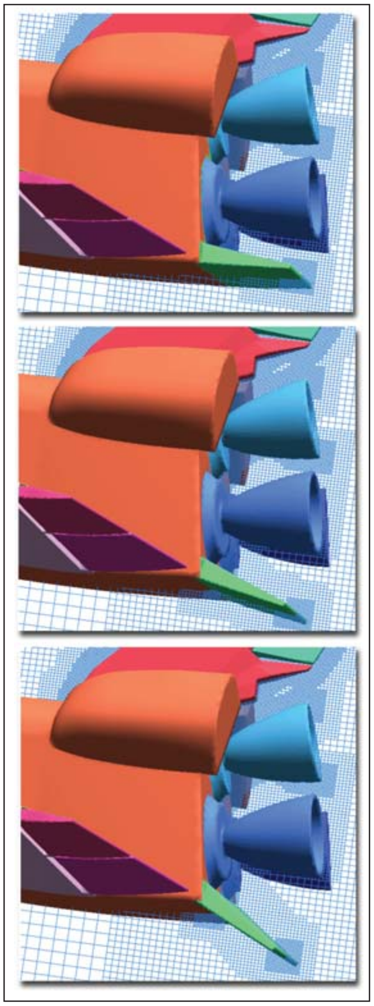

aerospace designers and engineers with a fully automated geometry processing and computational fluid dynamics simulation suite. This powerful tool streamlines the conceptual and preliminary analysis of new and existing aerospace vehicles.

"Cart3D reduces simulation time by a factor of at least 250," says Aftosmis. "This award recognizes the utility and impact that Cart3D is having throughout the engineering simulation community worldwide."

The software is already being used by a number of leading universities, and by an expanding list of more than 100 commercial users, including the Boeing Company and Northrop Grumman. The Ames Commercial Technology Office has granted a license for the software to ICEM CFD Engineering of Berkeley, California, a subsidiary of ANSYS, Inc., which will extend the use of Cart3D into other industries, including automotive and electronics.

Chapman Begins Operations

Chapman – the first and only 1,024 processor cache-coherent, shared-memory computer in the world – began operation in 2002, the result of collaboration between the NAS Division and SGI. The NASA/industry collaboration continues to build successively larger single-system-image computers and to advance the research and understanding in parallel supercomputing.

"Parallel shared-memory systems have many properties that provide significant advantages over clustered parallel systems," says NAS scientist Bob Ciotti. "For example, using the Multi-Level Parallel (MLP) library with other features and enhancements to the operating system, a simple yet powerful parallel programming environment allows users to more easily adapt engineering problems to parallel processing systems. This results in better utilization of scarce human programming resources, enabling the solution of NASA's most complex computational problems."

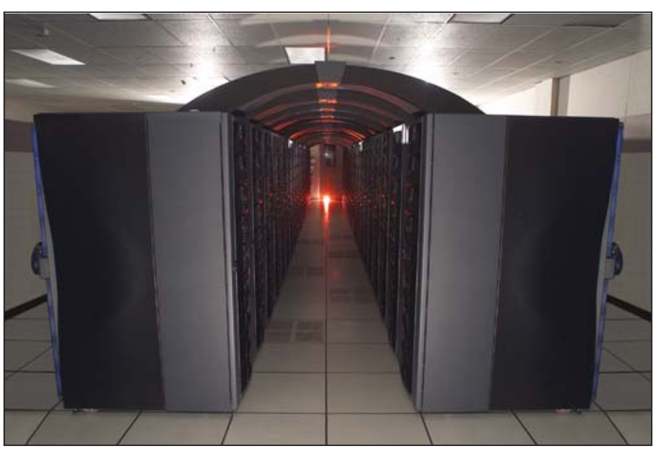

Chapman*, an SGI Origin 3800 parallel high-performance computer is now available for use by NASA researchers and their collaborators. (NASA/Tom Trower)*

UFAT Software Receives Space Act Award

The Unsteady Flow Analysis Toolkit (UFAT) software program, a pioneering tool for visualizing very large time-dependent or "unsteady" flow datasets from CFD simulations, won a NASA Space Act award in 2002. NAS scientist David Kao was recognized for his contribution to UFAT, which has been used to process and analyze high-fidelity simulation results for the Space Shuttle, military and commercial aircraft, artificial heart devices, and many other projects.

The software effectively reduces the analysis time of multi-gigabyte datasets from weeks to hours using state-of-the-art particle tracing and feature extraction algorithms developed by NAS scientists. For more information on UFAT, visit: www.nas.nasa.gov/Software/UFAT/

Other NAS recipients of the NASA Space Act Award include Nateri Madavan for his Aerodynamic Design using Neural Networks technology, and Guru Guruswamy for the High Fidelity Multidisciplinary Analysis Process (HiMAP) software system.

IPG Power Tools

Last summer, a team of developers working on NASA's Information Power Grid (IPG) deployed several new grid tools to make it easier for scientists throughout the United States to use the IPG's distributed resources. The new tools all use the Globus Toolkit, which provides a uniform interface to IPG resources, and were created to simplify the process of submitting jobs to the grid's distributed resources. New tools include a job manager to run jobs and track their progress, an abstraction tool for handling job scheduler information on multiple machines, and a resource broker for selecting grid resources (see *Gridpoints*, Fall 2002, page 4).

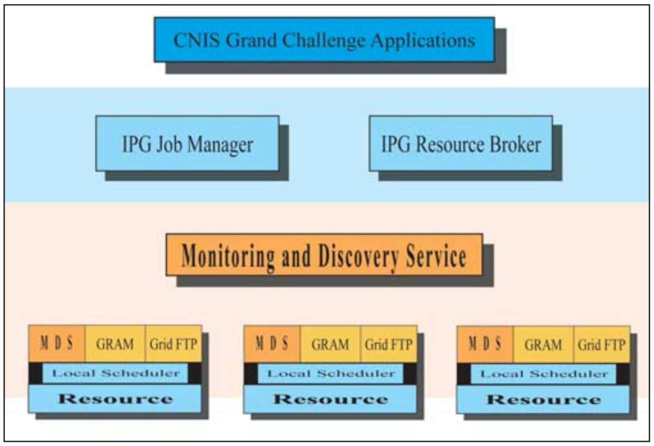

This diagram represents a conceptual view of the IPG software environment, including software components developed by the IPG team used to meet the requirements of a Computing, Information, and Communications Technology (CICT) Program milestone. The tools include the IPG resource broker and the Monitoring and Discovery Service (MDS).

The collaborative project drew on the expertise of IPG team members at NASA Ames, Glenn, and Langley Research Centers, together with researchers at the National Center for Supercomputing Applications (NCSA) in Illinois, the Information Sciences Institute at the University of Southern California, and Argonne National Laboratory in Illinois.

Delivery of the grid programming environment helped NASA's Computing, Information, and Communications Technology (CICT) Program's Computing, Networking, and Information Systems (CNIS) Project meet one if its major milestones. "This milestone provided us with a great opportunity to study the interoperability of grids," said NAS IPG task Tony Lisotta. "This will lead to future work in making grids more compatible with one another."

ILab meets IPG challenge

Another step forward for NAS's Information Power Grid (IPG) was accomplished in March, when ILab successfully completed 1,024 jobs using ten systems at four NASA sites (Ames, Glenn, and Langley Research Centers, as well as the Jet Propulsion Laboratory). The software package was developed by NAS scientists Maurice Yarrow and Karen McCann for automating the process of creating, submitting, and monitoring parameter studies on the IPG.

Following that accomplishment, NAS scientists again used ILab to run the Reusable Launch Vehicle (RLV) Grand Challenge computation as a five-by-five parameter study, using seven systems at three sites (Ames, Glenn, and Langley). Each computation ran through three separate restart phases, required 16 processors and 3.2 gigabytes of system memory, and generated output flow-data files of nearly 500 megabytes.

Largest Dataset Ever Visualized at NAS

Arguably the largest dataset visualized by NAS researchers was a simulation of the next generation of the Space Shuttle's main engine fuel pump, force-feeding liquid hydrogen into the main engine at 10,000 gallons per minute. The visualization focused on showing temperatures, pressure, and particle tracing, using 1,600 time steps – each with a gigabyte of data.

"We convert the dataset for each time step into a proprietary format that allows us to pull in paged chunks of data," explains NAS visualization specialist Timothy Sandstrom. "This

Patent Filed on Ultrafast Optoelectronics Device

Using semiconductor laser simulations, NAS scientists Cun-Zheng Ning and Peter Goorjian proved the effectiveness of using coupled lasers to achieve ultrafast modulation and switching in optical networks, resulting in a patent filed in June. Their method couples two vertical-cavity surface-emitting lasers (VCSEL), each with a diameter of about ten microns – one-tenth the thickness of a human hair, and carrying light at speeds of up to 40 gigahertz. (For more information on VCSELs, see *Gridpoints*, Spring 2000, page 12.)

"Coupling introduces a new frequency with a much greater limit than the intrinsic frequency of a single device," explains Ning. "Our method drives the higher frequency of the joint devices." Results from the project were published in the *Journal of Modern Optics, Optics Letters and Proceedings* of The International Society for Optical Engineering.

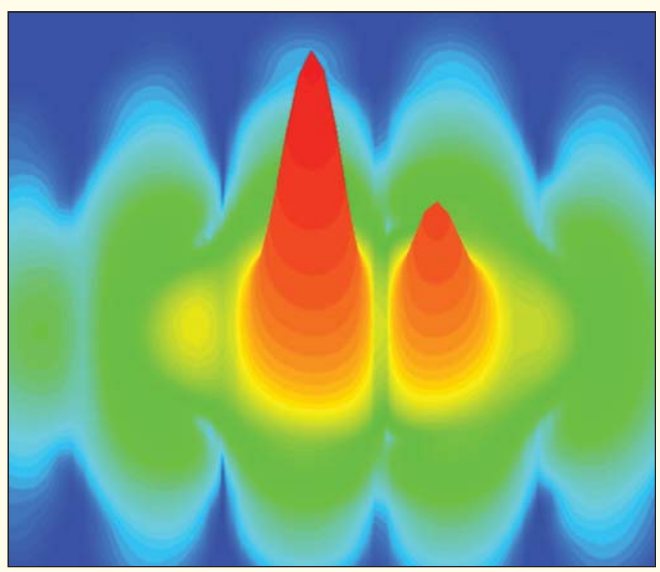

The image shows light intensity varying with angle of laser output from two coupled VCSELs. (Cun-Zheng Ning)

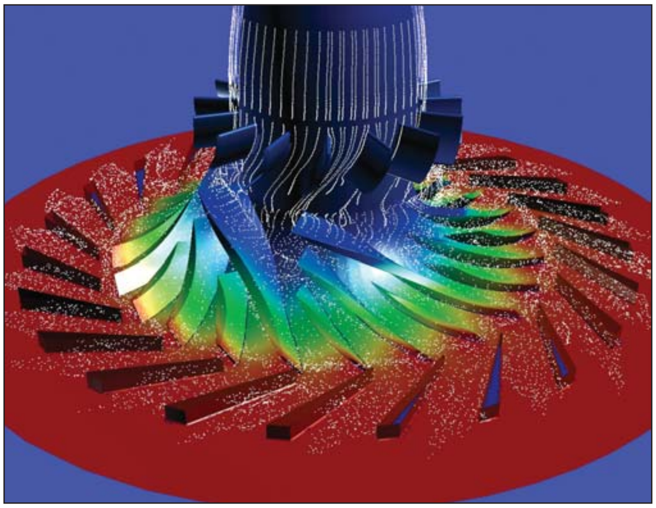

This Space Shuttle fuel pump simulation was the largest dataset computed at the NAS Facility during 2002. (Dataset: Cetin Kiris, Visualization: Timothy Sandstrom)

enables us to pull the exact data we need and focus on visualizing that one part."

In the meantime, NAS scientists continue to develop tools that will enable them to scale up flow-solver simulations for use on the 1,024-processor supercomputer, *Chapman*, significantly shortening design time.

NAS Nanotechnology Leading the Way

Experiments conducted at the Rensselaer Polytechnic Institute in New York, the Science and Nanotechnology Center in Sussex, England, and Ulm University in Germany, produced actual carbon nanotube t-junctions for the first time, work originally proposed possible in 1997 by NAS division scientists based on simulations.

"Now we are seeing the theory we predicted in action," says NAS scientist Deepak Srivastava. Since then, Srivastava has been investigating more complex structures, and is currently testing thermal signals, the instantaneous heat flow on a branched nanotube structure, to verify the amount of information processing contained in this structure.

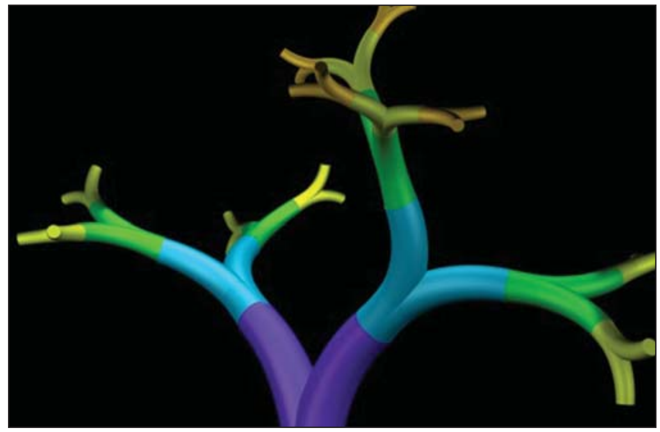

Biological neural systems were the inspiration for this simulation of a 3-D treelike structure of carbon nanotubes. NAS scientists are currently investigating the transmission properties of this type of architecture as a biomimetic approach to future computing paradigms. (Deepak Srivastava)

Shocking Chemistry

By simulating comet and meteor impacts on primitive atmospheres, NAS scientists are hoping to help researchers understand how life may have been formed on Earth, and whether the same processes could happen elsewhere.

At the American Chemical Society's National Meeting in Boston last August, NAS scientist Christopher Dateo gave a presentation of his group's findings, "Prebiotic Processing Induced by Comet and Meteor Impact." The study, funded by the Office of Space Science's exobiology program, simulated high energy shocks on four different atmospheres to determine the effect on the formation of molecules, and whether this might be a means of supplying organic material to a planetary atmosphere.

Dateo is developing the model and benchmarking it to experiments conducted by scientists in the Space Science Division at NASA Ames. He is optimistic that these studies could lead to the identification of bio-signatures or organic molecules.

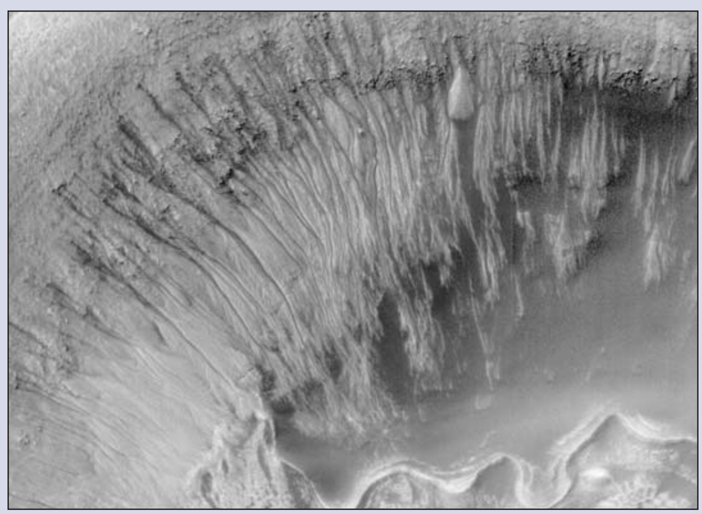

In an effort to determine how life began on Earth, scientists are researching the possibility that comet and meteor impacts brought organic material that could have eventually formed a planet's atmosphere. Mars' Newton Crater, formed by an asteroid impact that probably occurred more than three billion years ago, shows many narrow gullies, believed to have been formed by flowing water. (NASA/JPL/GRIN GPN-2000-001430)

Innovation in nanotechnology earned Srivastava and scientist Chris Henze a grant from the Ames Research Center Director's Discretionary Fund in 2002. Their proposal for the "biomimetic simulation of signal transmission in nanotube-based dendritic networks" was inspired by the structure and operations of biological neural systems. The scientists are investigating the properties of a 3-D, tree-like architecture of carbon nanotubes, and hope to prove its viability for applications involving the transmission and processing of signals.

Modeling Neural Systems

Researchers studying neural encoding and information theory analysis in the sensory system of a cricket's tail are using visualization and modeling techniques developed by NAS scientist Chris Henze to gain a better understanding of neural structures and how they operate.

Using a specially designed wind tunnel at the Center for Computational Biology at Montana State University, researchers Gwen Jacobs and John Miller record the crickets' responses to various wind stimuli, analyze the data, and figure out how the cricket's neural system encodes the input and output. In addition to visualizing a 3-D neural map for each stimulus, Henze has written software that allows the researchers to model what the cricket is doing, and has created a computational analogue of the laboratory set-up that harnesses the power of NAS' supercomputers. Researchers hope this work will provide a better understanding of the neural basis of perception and cognition.

Results from the project were presented at the international conference on Cooperative Dynamics of Neurosystems in Chile, the Computational Neuroscience Meetings in Chicago, and Neuroscience 2002, the annual meeting of the Society for Neuroscience.

Cataloging the Building Blocks of Life

Computational chemists in the NAS Division have developed a theoretical model that is already generating highly accurate data for a catalog of amino acids (see *Gridpoints*, Spring 2002, page 6). This catalog will help astronomers identify the molecules they observe in space, potentially providing more clues to solving the mystery of life's origin.

The model combines quantum and molecular mechanics in a unique vibrational infrared method. NAS scientist Galina

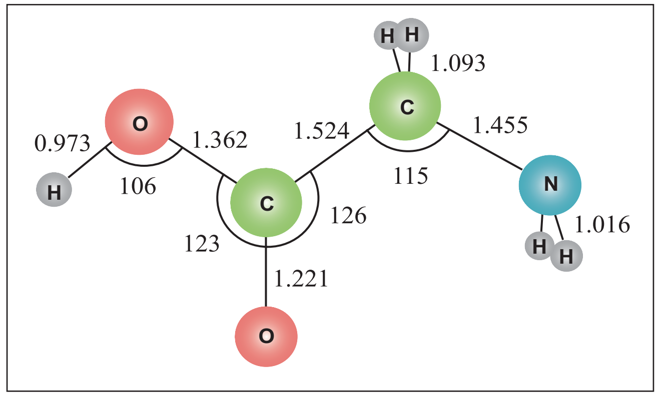

Illustration showing the complex of glycine, the simplest amino acid that occurs in many proteins and has already been observed in space. (Galina Chaban)

Chaban presented her findings to the second biennial Astrobiology Science Conference at NASA Ames Research Center in April, and since then has added glycine, alanine, and valine to the database.

NAS Scientist Receives International Award

For his distinguished work in computational nanosciences and carbon nanotubes, senior scientist Deepak Srivastava received the Eric Reissner Medal in August 2002 (see *Grid-*

The 'hyperwall' – From Concept to Debut

A team of NAS researchers has developed an innovative visualization system that allows scientists to view complex datasets and multiple parameters. The hyperwall, a seven-by-seven cluster of flat panel screens, (right), each driven by its own dual-processor computer with high-end graphics card, made its debut at the annual supercomputing conference, SC 2002, in Baltimore, and has proved popular among scientists throughout NASA (see *Gridpoints*, Fall 2002, page 11).

"The hyperwall allows scientists to look at data from different viewpoints," explains the system's chief architect, Chris Henze. "It can tile a single image like a powerwall can, but its main strength lies in its ability to support spreadsheet-style approaches to visualizing data."

Marsoweb Usage Doubles… Then Triples

Following an article in the Summer 2002 edition of *Gridpoints* and a subsequent press release, Marsoweb, a NAS-developed website for the Mars Exploration Program Landing Sites data, saw its user base soar to 70,000 – more than double the number of visitors in two months, compared with the past three years.

Marsoweb, developed by NAS visualization specialist Glenn Deardorff, is designed to provide a one-stop-shop for scientists involved in landing site selection for upcoming Mars Rover missions. In addition, it hosts a variety of images and an interactive data map that has also proven popular with Mars enthusiasts and the educational community.

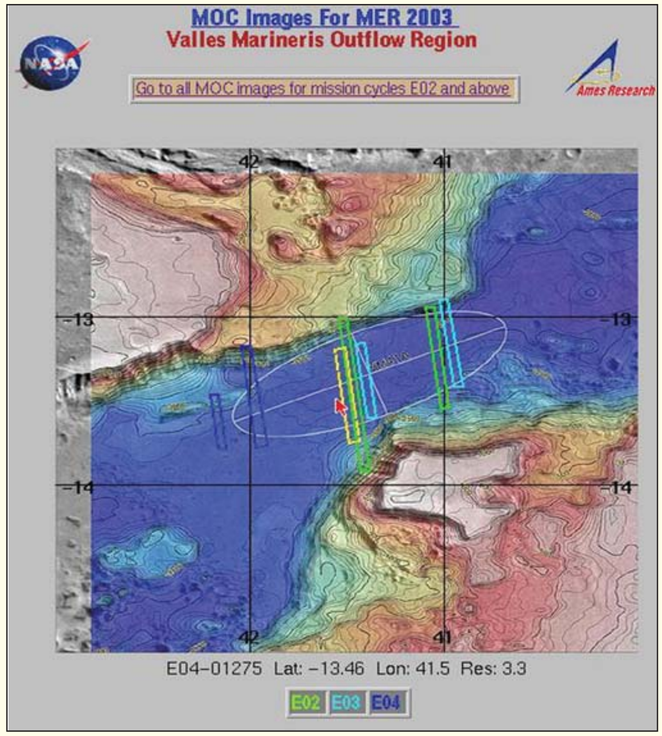

Mars Orbiter Camera (MOC) images for MER 2003 landing site studies can be displayed on regional navigation pages for closer scrutiny of the planet's surface. This user interface is used to select and view images in the MOC image archive. (Glenn Deardorff)

points, Fall 2002, page 1). The prestigious medal, named after the noted researcher and academic, is awarded every two years by the International Conference on Computational and Experimental Engineering and Sciences.

NAS Scalable Algorithms used in Materials Science

In a collaborative effort that joined two universities in the United States, four universities in Japan, and NASA Ames Research Center, a multiscale simulation was successfully conducted to study the environmental effects of water molecules on fractures in silicon. The simulation was performed on a grid of 25 PCs distributed over three clusters in the United States and Japan. NAS senior scientist Subhash Saini designs algorithms for large-scale applications and participated in the project. Saini co-authored several papers on scalable algorithms in 2002, including "ARMS: an agent-based resource management system for grid computing," which appeared in *Scientific Programming 10,* and "Collaborative Simulation Grid: Multiscale Quantum-Mechanical/Classical Atomistic Simulations on Distributed PC Clusters in the U.S. and Japan," published by the Institute of Electrical and Electronics Engineers.

NAS Research Advances Theory in DNA Electronics

Research in the conductive properties of DNA continues to challenge researchers around the world with contradictory experimental results. In 2002, NAS scientists M.P. Anantram, T.R. Govindan, and Alexei Svizhenko produced two papers from their theoretical and modeling work that they hope will help explain some of the experiments. The papers, "Environment and structure influence on DNA conduction" and "Influence of disorder on DNA conductance" are due to appear this spring. The scientists are also interested in learning how arbitrary sequences of DNA might be exploited in nanotechnology.

Parallelizing Quantum Chemistry

As large supercomputers and parallelization offer researchers greater amounts of processing power, the theoretical approach to quantum chemistry is becoming increasing relevant. Simulations using the laws of quantum mechanics can provide researchers with detailed information on the subtle interactions between electrons, without the high cost of experiments.

NAS scientist Graham Fletcher has been designing scalable algorithms to distribute data-intensive processing over a group of processors. In 2002, he developed a scalable, parallel, distributed-data Multi-Configurational Self-Consistent Field (MCSCF) program that allows researchers to study theoretical situations in which electrons are being separated. Fletcher believes this is particularly important for materials scientists designing lighter and stronger materials for space hardware. Fletcher's program is available on the General Atomic and Molecular Electronic Structure System, and maintained by members of the Gordon Research Group at Iowa State University.

NAS Releases Grid Benchmarks for Global Community

The Global Grid Forum (GGF) called on NAS scientists Rob Van Der Wijngaart and Michael Frumkin for guidance in measuring grid efficiency and user-friendliness. Following approval for the formation of an official research group on grid benchmarks in July, NAS promptly responded with the release of GridNPB3.0, a pencil-and-paper specification of four families of problems representing important applications on computational grids (see *Gridpoints*, Summer 2002, page 5, and *Gridpoints*, Fall 2002, page 15).

The release of GridNPB3.0 coincided with the release of a major new version of the NAS Parallel Benchmarks, designed to incorporate several important aspects of today's computing needs.

Conference Hosts Biologists, Nanotechnologists, and Information Technologists

In the first workshop of its kind, nearly 100 experts gathered at NASA Ames Research Center last October to explore cross-disciplinary approaches to solving the science and technology challenges facing future NASA missions. NAS senior scientist T.R. Govindan co-chaired the Biology, Information

Meyya Meyyappan, director of the Center for Nanotechnology at Ames, presents an overview of the nanotechnology session of the workshop. (NASA/Tom Trower)

Science, Nanotechnology Fusion and NASA Missions Invitational Workshop, which was organized by NASA Ames and the University Space Research Association.

"Our goal was to start a conversation and gather new ideas," explains Govindan. "Based on the level of enthusiasm and interest in the fusion of these three fields, it's likely that this kind of conference might become a regular event." For more information, visit http://binfusion.arc.nasa.gov

Analyzing Spectral Congestion

Calculating the vibrational frequencies and infrared opacities of water and methane is important in studying planetary atmospheres, and in the astrobiological search for life. NAS scientists developed computational methods to help astronomers analyze their observations and understand the composition of planetary atmospheres, with an error rate of less than 0.00012 electron volt. This level of accuracy provides researchers with a computational method to ascertain methane presence, where spectral congestion has until now made this an impossible task.

NASA's Office of Space Science funded the project, and NAS scientist David Schwenke presented the research, "*Ab Initio* Theory of Water and Methane Frequencies and Opacities," to the American Geophysical Union Fall Meeting in December. Astronomers worldwide are now using this data.

Calculating Planetary Entry Reactions

When a spacecraft enters a planetary atmosphere at very high speed, it causes a reaction that excites, dissociates, and ionizes the ambient atoms and molecules, generating temperatures

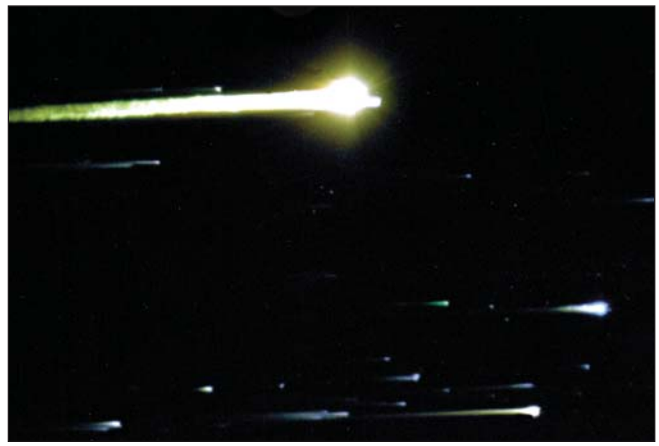

This photograph of Apollo 8 reentering the Earth's atmosphere was taken by a U.S Airforce Airborne Lightweight Optical Tracking System camera mounted on a KC-135A aircraft flying at 40,000 feet. NAS researchers are studying how atoms and molecules interact during vehicle reentry. Results from these studies will be used in future space vehicle designs.
(NASA/JSC/GRIN GPN-2000-001357)

in excess of 20,000 degrees Kelvin. To help heat-shield designers understand how molecules like oxygen and nitrogen react in such high temperatures, NAS scientist Dunyou Wang has completed the first accurate theoretical calculation of the rate reaction, providing important data for flow-field modeling and determining heat load. The paper, "Quantal Study of the Exchange Reaction for $N+N_2$ Using an *Ab Initio* Potential Energy Surface," will be published by the *Journal of Chemical Physics*.

Studying Diffuse Interstellar Bands

For more than 50 years, scientists have been trying to determine the origin of diffuse interstellar bands (DIBs) in the astronomical spectrum. More recently, speculation has focused on classes of polycyclic aromatic hydrocarbons (PAHs), and two collaborative studies between NAS scientists and the Astrochemistry Lab at NASA Ames Research Center have resulted in significant findings in the way PAHs behave. "We looked at the neutral, positive, and negative ions of a class of compounds called perylene, terrylene, and quaterrylene,"

Chimera Grid Tools and Pegasus5

NAS scientists William Chan and Stuart Rogers continued to improve grid generation processes in 2002, with the release of version 1.7 of the Chimera Grid Tools (CGT) suite, and Pegasus5 software.

The CGT package contains a variety of programs and scripts used in the overset CFD process for grid generation, diagnostics, and solution post-processing. The new release contains a module for automatic boundary conditions selection and a components module, allowing a fast and user-friendly way to create the inputs for simulations involving multiple components in relative motion.

Pegasus5 software is used to perform the preprocessing task of linking together a large number of randomly overset grids. During 2002, major enhancements were made to significantly reduce the CFD cycle-time.

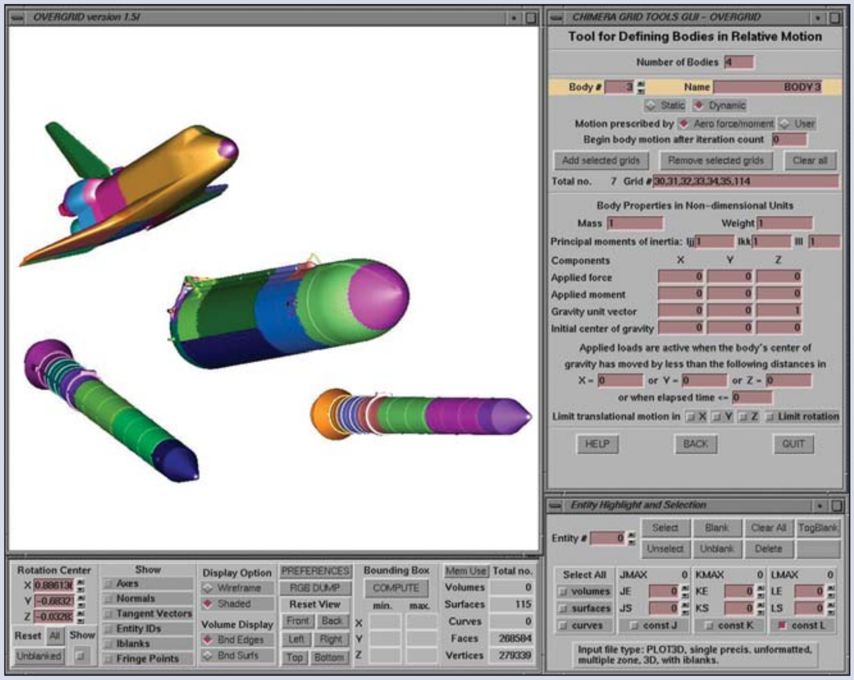

This screenshot from the latest version of the Chimera Grid Tools software package shows an example of a fictitious Space Shuttle Launch Vehicle separation procedure. Each body is shown in dynamic relative motion and contains multiple grids, depicted in different colors. *(NASA/William Chan, Stuart Rogers)*

For more information on CGT and Pegasus5, visit: http://www.nas.nasa.gov/~rogers/cgt/doc/man.html

explains NAS scientist Timothy Lee. "Our theoretical results explained all of the Astrochemistry Lab's observations."

Lee's collaborative paper, "Electronic Absorption Spectra of Neural Perylene, Terrylene and Quarterrylene and Their Positive and Negative Ions: Neon Matrix-Isolation Spectroscopy and Time Dependent Density Functional Theory Calculations," is due to appear in the *Journal of Physical Chemistry* this year. Results of a similar study, "Time Dependent Density Functional Theory Calculation of Large Compact PAH Catons: Implications for the Diffuse Interstellar Bands," is due to appear in the *Astrophysical Journal*.

Modeling the Probability of Ionization

Ionization – the process by which molecules kick out electrons – is important for the survival of mankind, as it provides a protective shield that deflects cosmic rays. Ionization also has a variety of other purposes, and in 2002 NAS scientists completed a project to develop an accurate model for calculating the probability of ionization. This model will help in the design of silicon chips and in studying how the human body reacts when exposed to radiation in space. The result, "Total Electron-Impact Ionization Cross-Sections of CFx and NFx (x = 1-3)" by Winifred Huo, Vladimir Tranovsky, and Kurt Becker, was published in *Chemical Physics Letters* by Elsevier, May 31, 2002. "The next step is to calculate the probability of ionization breaking up molecules," says NAS scientist Huo. "This is important in studying the health effects of ionization and how carcinogens damage DNA."

NAS Researcher Awarded for New Mesh Refinement Strategy

NAS researcher Michael Aftosmis and colleague Marsha Berger from the Courant Institute in New York received the American Institute of Aeronautics and Astronautics (AIAA) best paper award for 2002. The award was presented in June at a special luncheon during the 32nd AIAA Fluid Dynamics Conference for their joint paper, "Multilevel Error Estimation and Adaptive h-Refinement for Cartesian Meshes with Embedded Boundaries." This paper describes the team's development of new techniques for error estimation and adaptive refinement for CFD solutions, and formed the basis of a planned adaptation module for NASA's Cart3D analysis software.

NAS Scientist Appointed to Springer-Verlag Editorial Board

NAS scientist Timothy Barth was appointed to the editoral board for the book series *Lecture Notes in Computational Science and Engineering,* published by international scientific publisher Springer-Verlag. Previously, Barth was a contributing author to two editions of the lecture series, which was established in 1997 and is part of the publisher's mathematics

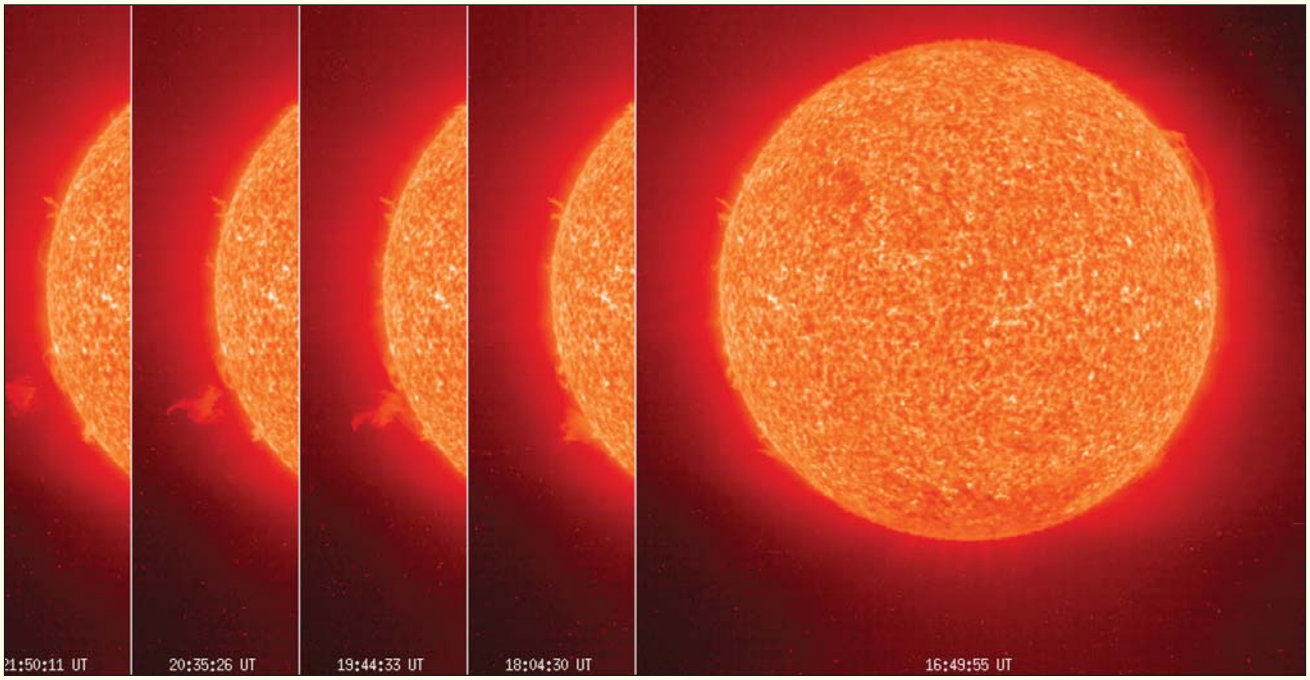

Modeling the Sun

The NAS Division's expertise in turbulence research, combined with the power of the division's supercomputers, is the driving force behind a collaborative project with Stanford University scientists interested in modeling the physical properties of the sun.

The main goal of the project is to understand the physics of a region inside the sun called the "convection zone," where energy produced in the core is transported to the surface in much the same way as a pot of boiling water bubbles and rolls. "The Earth's weather is strongly dependent on variations in solar radiation," explains NAS scientist Alan Wray. "We want to know how these variations occur, why, and in what timeframe." Results of this research may also help NASA protect its astronauts from spurious solar rays, and ensure communication satellites stay in optimum orbital positions. This work may also influence NASA's search for extra-solar life.

Above: *This sequence of images, captured by the NASA/ESA Solar and Heliospheric Observatory (SOHO), shows an eruptive prominence or gaseous emission from the sun measuring more than 80,000 miles long and traveling at more than 15,000 mph. Such events can disrupt power and communications in space. (NASA/GSFC/GPN-2002-000120)*

division. "Springer is embracing computational science and engineering as a new subject," says Barth, whose duties on the editorial board include approving new books for the series and providing input to the senior editor about potential authors and new topics.

Patent for Non-linear Optics Simulation

In addition to awarding NAS scientist Peter Goorjian a distinction in the 2002 NASA Space Act Awards for his patent on photonic switching devices using light bullets, NASA headquarters granted an exclusive license under the patent to Cyrospace, Inc., a technology mining, development, and commercialization company.

Goorjian's work studies light pulses, commonly known as light bullets because they measure approximately ten femtoseconds, or 1/100,000,000,000,000 of a second. Results from simulations of light bullets propagating through certain materials, like glass or fused silica, can help researchers determine how these short pulses may be used to make optical switches. The next step is for Cyrospace, Inc. to begin creating experiments with light bullets in the laboratory, using techniques from Goorjian's patent.

Center for Turbulence Research's Summer Program 2002

The biennial summer program at the Center for Turbulence Research (CTR) welcomed 46 participants from 11 countries, nine states, and 40 institutions worldwide in a collaborative effort that brought theorists and experimentalists together in a number of turbulence research projects. CTR is a cooperative program between Stanford University and NASA Ames Research Center, and the NAS Division supports the program with supercomputer access and project advice and management.

Each of the 32 projects conducted in the summer were based on researchers' own proposals that matched with NASA's interests. Research focused on nine areas: turbulence acoustics, Reynolds-Averaged Navier-Stokes (RANS) modeling, large eddy simulation (LES), numerical methods for LES, turbulence fundamentals, stratified shear flow turbulence, flow optimization, nanofluidics and biology, turbulent combustion and sprays.

A summary of the findings will appear in the journal *Physics of Fluids*, followed by a full volume of final papers due to be published by CTR this year.

NASA-operated Boeing 727 with wing tip smoke generators graphically illustrates the aircraft's wing tip trailing vorticies. NAS scientists are collaborating with partners in France to determine how trailing vorticies affect the cloud formation. *(NASA/Dryden ECN-3831)*

Strained Wakes

The ability to predict the flow over a high-lift airfoil using standard CFD codes depends on understanding how turbulence behaves in the wake of multicomponents (flaps, spoilers, and leading edge slats) that separate to increase lift for take-off and landing. NAS scientist Michael Rogers has been simulating strained wakes – regions of reduced velocity behind an obstacle – and mixing layers where two streams of different speeds come together, then analyzing their effect on the types of turbulence that are found in high-lift airfoils. Rogers' paper, "The Evolution of Strained Turbulent Plane Wakes," was published in 2002 in the *Journal of Fluid Mechanics*, vol. 463, pp. 53-120.

Stratified Turbulent Flows

Another area where Michael Rogers employs his turbulence modeling expertise is in helping researchers understand stratified turbulence for making global climate models. Stratified turbulence is found in the atmosphere around the Earth, particularly relating to cloud physics and formation. During CTR's Summer Program, Rogers participated in several projects analyzing data from numerical simulations, including one effort that attempted to distinguish between fluctuating motions caused by internal gravity waves and classical stratified turbulence.

"Scientists at CTR have been developing new statistics that will help with turbulence modeling," explains Rogers. "Researchers from all over the world are interested in becoming more familiar with CTR's modeling ideas."

Contrails and Cloud Formation

NAS scientists studying aircraft trailing vortices focused on aircraft contrails, the lines that trail across the sky like smoke behind an airplane, in one of CTR's summer program experiments that will ultimately be useful to Earth scientists modeling the global climate.

NAS scientist Karim Shariff collaborated with researcher Roberto Paoli, a post-doctorate scientist with the Computational Fluid Group at CERFACS, the European Centre for Research and Advanced Training in Scientific Computation in France. Paoli's research project simulated the interaction of jet exhaust with the vortex of the aircraft, and drew upon prior knowledge gained by researchers in the Earth Science Division at Ames to determine how contrails artificially increase cloudiness and trigger the formation of cirrus clouds.

Wireless Firewall Gateway Gains National Recognition

Development of a system that ensures secure interoperability of a wireless network, using the 802.11b standard, gained national recognition when the magazine *Government Computer News* featured an article on the project on the front cover of their October 8 issue. The Wireless Firewall Gateway (WFG) promotes easy and secure wireless networking using techniques that could be applied inexpensively anywhere. NAS network and security team members Nichole Boscia and Derek Shaw worked with Dave Tweten, computer security official at NAS, and were interviewed for the article. Their paper, "Wireless Firewall Gateway," was also published in a book titled *Wireless Security Essentials*.

The NAS Division's achievements in computational science and information technology cover a broad spectrum of research and development throughout NASA, and reflect the diversity, drive and determination of each member of the NAS team. By expanding its high-performance computing power and forming collaborative partnerships with industry, academia, and other NASA centers, NAS has affirmed its leadership position in the future of capability computing.

— Julie Jervis

Parallelization – The Key to Faster Codes, Higher Fidelity Simulations

NAS Division researchers create a computer-aided tool to help automate the tedious process of parallelizing serial code.

To create better cloud models, researchers are using a computer code recently parallelized by NASA Advanced Supercomputing (NAS) Division researchers Henry Jin and Gabriele Jost. Parallelizing the code enabled scientists at NASA Goddard Space Flight Center, Greenbelt, Md., to run larger cloud simulations in shorter periods of time, helping to shed light on phenomena such as air-sea interactions and global climate changes.

Jin and Jost worked with scientists Wei-Kuo Tao, Dan Johnson, and Chung-Lin Shie of NASA Goddard, to improve the three-dimensional Goddard Cumulus Ensemble (GCEM3D) Code. To accomplish this work, NAS researchers applied their tool, Computer-Aided Parallelization and Optimizer (CAPO), designed to automate the tedious and error-prone steps of parallelizing serial code, and to take advantage of shared-memory parallel computers. "The purpose of parallelizing code is to make it run faster. The idea is, if it runs faster on one processor, it should run 'n' times as fast on 'n' number of processors," explains Jost. Parallelizing code enables researchers to solve bigger problems – problems that may not have been previously solvable due to a lack of time or available resources.

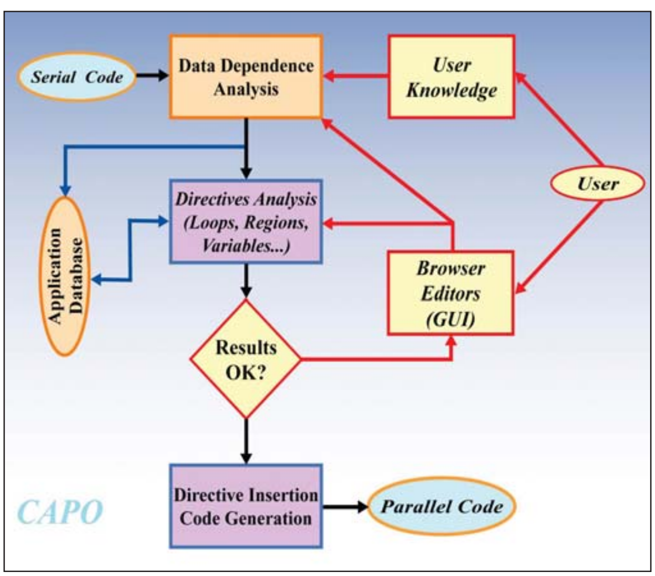

Building Tools with Tools

Development of the CAPO tool in 1998 was inspired by a collaboration with the developers of CAPTools at the University of Greenwich in the United Kingdom. CAPTools is software that generates code to run on distributed memory machines. Jin and Jost believed that they could extend CAPTools to produce a parallel program for shared-memory machines to take advantage of the unique computing environment at the NAS Facility (see *Gridpoints*, Spring 2002, "The Evolution of High-Performance Computers in the NASA Advanced Supercomputing Division," page 1A). "I think the CAPO tool is very useful, especially for model runs demanding lots of CPUs and massive amounts of memory," says Goddard's Shie.

The CAPO tool traverses through several steps in order to parallelize a serial code (see Figure 1). First, the tool performs a data dependency analysis (to determine how different variables depend on one another), then a loop level analysis (looking for repeated sequences of instructions in the code), followed by a series of graphical user interfaces that illustrate the parallelization process to the user. "The tool takes away the tedious and error-prone work of parallelizing code, allowing the user to focus on optimization of critical parts of the code, all through a single interface," explains Jost.

Dependency analysis, the core part of the CAPO software, determines the relationships between variables within a serial code. This process is time consuming in most cases because complex codes have many routines and subroutines (sets of programming instructions designed to perform specific tasks). "Dependency analysis is an important step in the process, because you have to decide whether operations can

Figure 1: *The process of parallelizing serial code using the CAPO tool begins by importing serial code (top left of figure). The tool then analyses data dependencies to determine the relationship among variables in the code. After the data dependency analysis step, a loop level analysis helps determine how to divide the code to run in parallel. Then, a series of graphical user interfaces guides the user through the process of addressing sections of code that cannot be parallelized in original form.*

be performed in parallel or not. If memory access is dependent on other operations, the operations have to be performed in a certain order [versus running operations in parallel]," explains Jin.

All data generated during the dependency analysis process is stored in a database. If the CAPO system runs into an obstacle while conducting the analysis, information about the nature of and reasons for the obstacle are stored in a database. A user can later examine the database and potentially remove obstacles based on his or her knowledge about the application and its input parameters. "The user guides the tool, but the tool also guides the user – it's an interaction," explains Jost.

Following the dependency analysis, the loop level analysis examines the code. "CAPO examines the loops for potential data dependencies that might prohibit parallelization. If you have a loop which iterates over some repeated sections of the code, you can actually break this loop into individual pieces, such that you can run them concurrently on the processors," explains Jin. "That's how you get the speed-up in code performance." Speed-up is the ratio of the rate at which work is done when a job is run on 'n' processors to the rate at which it is done by just one processor.

Once the loop level analysis is complete, users are guided through a series of the tool's built-in graphical user interfaces. CAPO's interface enables users to view all instances where the code did not parallelize. The more errors or obstacles a user is able to remove, the higher level of parallelization that can be achieved (See Figure 4, page 20).

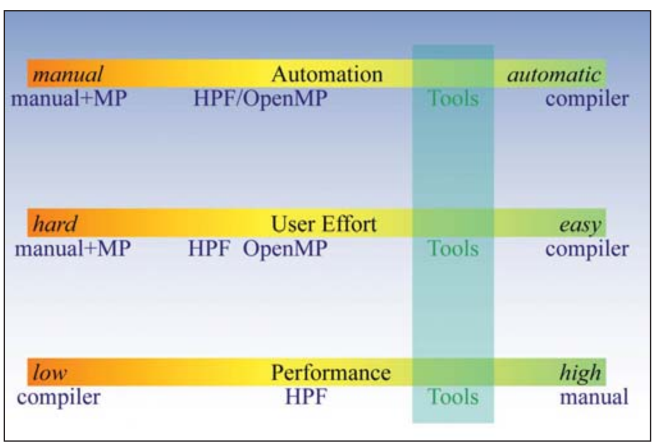

Figure 2: *The different levels of ease-of-use and performance for various parallelization approaches are illustrated in this diagram. Hand-generated parallel code usually shows very high performance, however, this method requires a lot of user time and effort. High-performance Fortran (HPF) or OpenMP are programming paradigms to bridge this gap. Compiler-based automatic parallelization is easy to use; however, performance of the generated code is often limited. Using a parallelization tool such as CAPO, considerably eases the workload for users, while maintaining acceptable performance of the generated code.*

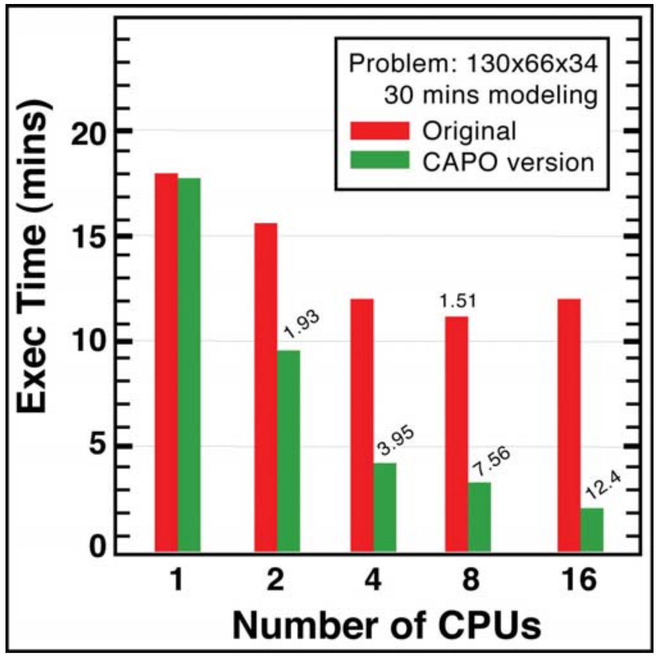

Figure 3: *This bar chart represents the contrast in performance between the original GCEM3D code (in red), and the parallel CAPO version (in green). This specific example is a 130-by-66-by-34 case run on an SGI Origin 3000. While the original did not scale beyond four processors (CPUs), the CAPO version achieved a speedup of 12.4 on 16 CPUs, improving performance of the code by a factor of 8.2.* *(NASA/IPG)*

Proof Is in the Numbers

Before applying CAPO to GCEM3D, the cloud modeling code was only able to run very small cases, scaling up to four processors on a PC. After applying the CAPO tool and making a few adjustments, Jin and Jost achieved a speed-up of a factor of 12 when running on 16 processors of an SGI Origin 3000 for a 130-by-66-by-34 test case (see Figure 3). Using larger cases, the code was able to scale up to 64 CPUs. Test cases were also run on an SGI Origin 2000, a Sun Enterprise, and a Dell PC, demonstrating the portability of the code CAPO generates. Wherever the OpenMP programming model is supported, CAPO-generated code can be run. OpenMP is a portable, scalable model that provides parallel programmers a simple and flexible interface for developing shared-memory parallel applications.

The newly parallelized GCEM3D code enabled researchers at Goddard to increase the resolution of their case studies. They successfully ran a large test case of 1,026-by-1,026-by-34, using more than seven gigabytes of memory – a new feat using this application. "Our goal of cloud modeling not only aims to better understand the microphysical and dynamical processes of the cloud system itself, but also to improve their representation for large-scale applications, such as studies on the precipitating convective system, air-sea interactions and cloud-aerosol interactions, as well as the global change in climate and hydrology," explains Shie. Having the ability to

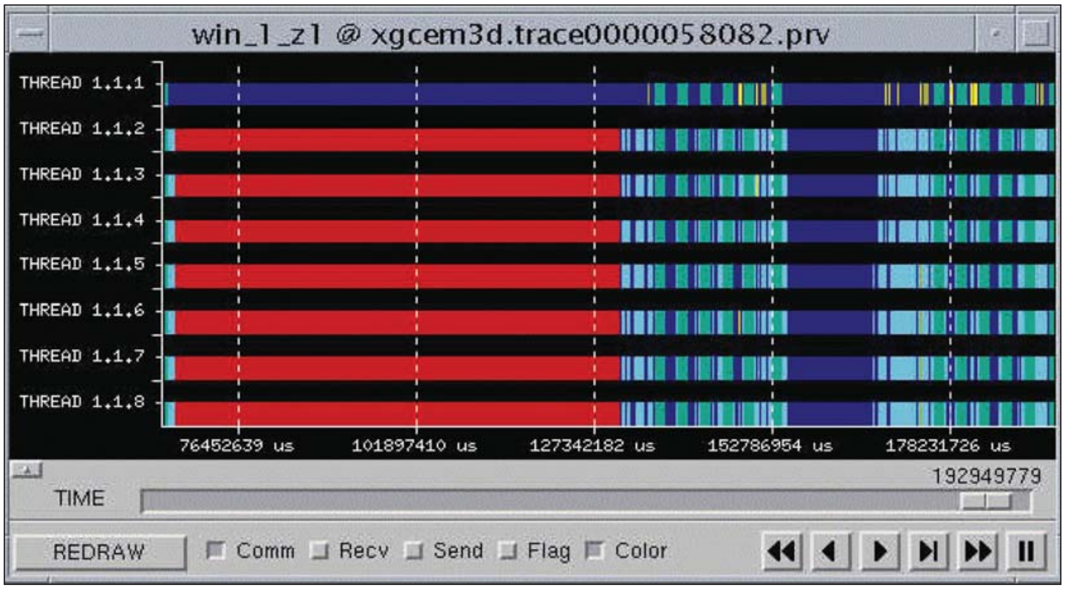

Figure 4: *These diagrams represent timelines of the first few iterations of a loop in a routine for eight threads. (Threads are portions of a program that can be run independently of, and concurrently with, other portions of the program.) Dark blue coloring indicates a running thread, light blue indicates an idle thread, and red represents inactive threads.*

Case No. 1, top: *In the original GCEM3D code, only one thread (the master thread) is active most of the time, while the remaining threads are idle.*

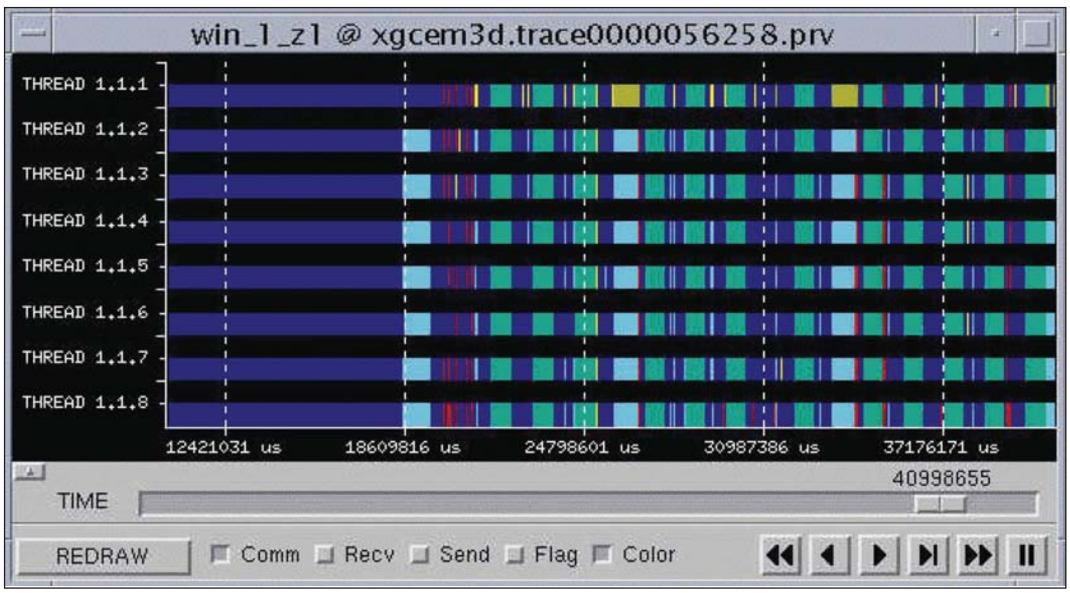

Case No. 2, right: *In this case, the CAPO-parallelized version of the cloud modeling code, all threads are active most of the time (represented by the dark blue coloring).*

run much larger cases enables NASA Goddard scientists to achieve these research goals faster.

Sticking to What Works

Although there are other methods for parallelizing code, CAPO is optimal for achieving performance, ease of use, and automation with a single tool. Other methods include generating the modified code by hand, using an automatic compiler, and employing other parallelization tools such as high-performance Fortran and OpenMP (see Figure 2, page 19). The CAPO tool is superior in its ability to pinpoint errors versus an automatic compiler. “Unlike a compiler, CAPO allows the user to provide the knowledge about code structures or input parameters in order to remove unnecessary data dependencies,” explains Jin. Using an automatic compiler to parallelize code will not result in the kind of speed-up the CAPO tool can achieve, he adds.

After successfully parallelizing their cloud modeling code, researchers at NASA Goddard are interested in applying the CAPO tool to some of their other serial codes. “I will apply CAPO to other codes in the future because of the substantial improvement in model performance due to computational efficiency and memory extension,” says Shie.

Tao and Shie visited NASA Ames in September 2002 to learn more about the parallelization tool. And recently, Jin and Jost visited NASA Goddard to demonstrate the tool to a group of researchers. The CAPO tool team’s eventual goal is to transfer the knowledge of the tool so that each user can apply the tool to codes individually. Decades of effort have been consumed generating serial code that now needs to be parallelized so it will run more efficiently on the shared-memory parallel systems at the NAS Facility.

—Holly A. Amundson

Calendar of Events

17th Annual HPCC Conference

Newport, Rhode Island • March 25–27, 2003

The National High Performance Computing and Communications Council holds an annual High Performance Computing and Communications Conference each spring. This is one of the few conferences that emphasizes communication between manufacturers and users, as well as academics and the government agencies which establish policy and regulate the use of advanced technologies. Topics covered include: wireless computing, e-government, Internet security, grid computing, mass storage, homeland cyber security, and bioterrorism. For more information, visit: www.hpcc-usa.org/genconf.html

Linux Clusters: The HPC Revolution 2003

Las Vegas, Nevada • June 17–20

The Linux Clusters: the HPC Revolution 2003 Conference is the premier international forum for Linux cluster users and system administrators to share information on applications and tools development, scientific computing techniques, and systems administration of Linux clusters. More information is available at:
www.linuxclustersinstitute.org/Linux-HPC-Revolution

High Performance Distributed Computing Conference (HPDC)

Seattle, Washington • June 22–24

The Twelfth IEEE International Symposium on High-Performance Distributed Computing will be a forum for presenting the latest research findings on the design and use of highly networked systems for computing, collaboration, data analysis, and other innovative tasks. HPDC-12 will be held in Seattle, Washington immediately preceding the 8th Global Grid Forum. Details are available at:
www-csag.ucsd.edu/HPDC-12

International Workshop on Active Middleware Services

Seattle, Washington • June 25

Held in conjunction with the 12th International Symposium on Grid Computing (HPDC 12), the fifth annual International Active Middleware Workshop will focus on Autonomic Computing and will bring together leading researchers and ideas in this emerging discipline. Submission deadline for abstracts is February 28. More information can be found at: www.caip.rutgers.edu/ams2003

SIGGRAPH 2003

San Diego, California • July 27–31

SIGGRAPH 2003 hosts the largest, most comprehensive exhibition of products and services for the computer graphics and interactive marketplace. Details at:
www.siggraph.org/s2003/conference

2003 NASA Information Power Grid (IPG) Tutorial and Workshop

Palo Alto, California • February 4-6

The IPG Tutorial and Workshop will be held February 4, providing an introduction to the NASA grid and its various tools. Tutorials will review the general capabilities of the IPG, the basic functions available to its users and those developing IPG applications. The tutorial requires no previous IPG experience.

In contrast to other grid workshops and events, the IPG workshop is focused specifically on NASA's Information Power Grid. The objectives of the IPG workshop are to provide a forum in which: IPG system developers and implementors can share their experiences; IPG application developers and users can share their experiences; potential IPG users can gain a broad understanding of its capabilities; and IPG research and development personnel and the NSF PACI partners can report on the progress of their grid research and development.

Presentations will include talks from research and development teams from each of the major IPG-related organizations including: NASA centers; Argonne National Laboratory; National Computational Science Alliance; San

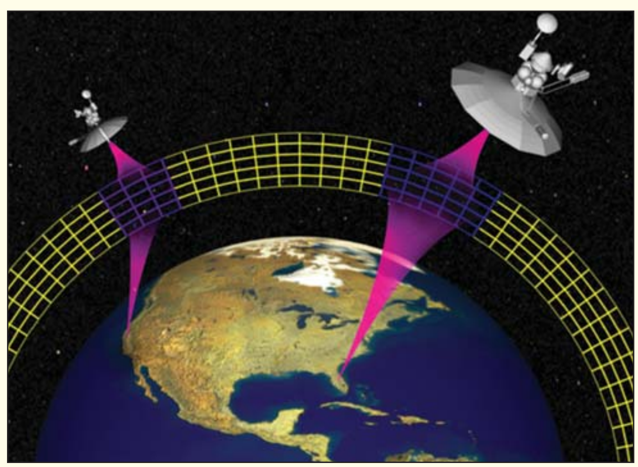

Diego Supercomputer Center; University of Southern California's Information Sciences Institute, and other participating universities.

For additional information on this event, visit the website: www.nas.nasa.gov/2003ipg or contact Marcia Redmond at: mredmond@mail.arc.nasa.gov.

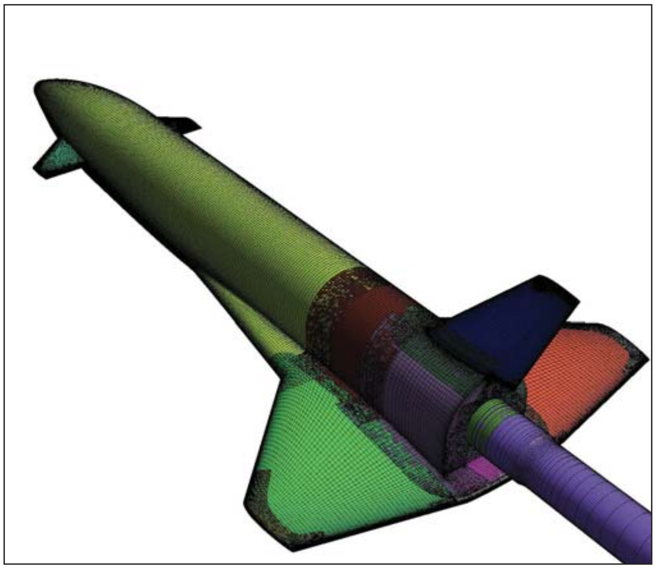

Completing Large Parameter Studies – with Time to Spare

NAS Division researchers create an automated CFD system to run large parameter studies on IPG resources. *See page 4.*

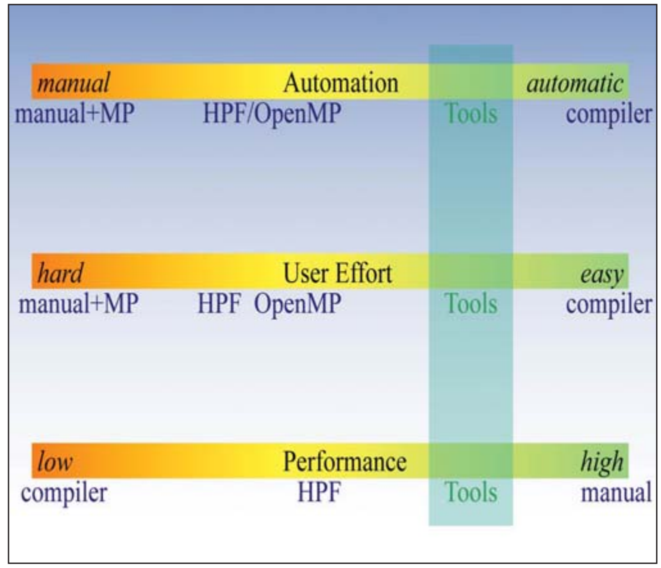

Parallelization — The Key to Faster Codes, Higher Fidelity Simulations

Scientists within the NAS Division have designed a computer-aided tool to help automate the tedious process of parallelizing serial code. *See page 18.*

www.nas.nasa.gov/gridpoints

National Aeronautics and Space Administration
NASA Advanced Supercomputing Division
Ames Research Center
MS IN258-6
Moffett Field, CA 94035-1000

Official Business
Penalty for Private Use $300

FIRST CLASS MAIL
POSTAGE & FEES PAID
NASA
Permit No. G-27

Address Correction
If the above address or mail stop is incorrect, please e-mail a current address to gridpoints@nas.nasa.gov, or write to:
Gridpoints Editor, MS 258-6, Moffett Field, CA 94035

PLEASE RECYCLE
Printed on recycled and recyclable paper with vegetable-based inks