Summer 2002

Gridpoints

The Quarterly Publication of the NASA Advanced Supercomputing Division

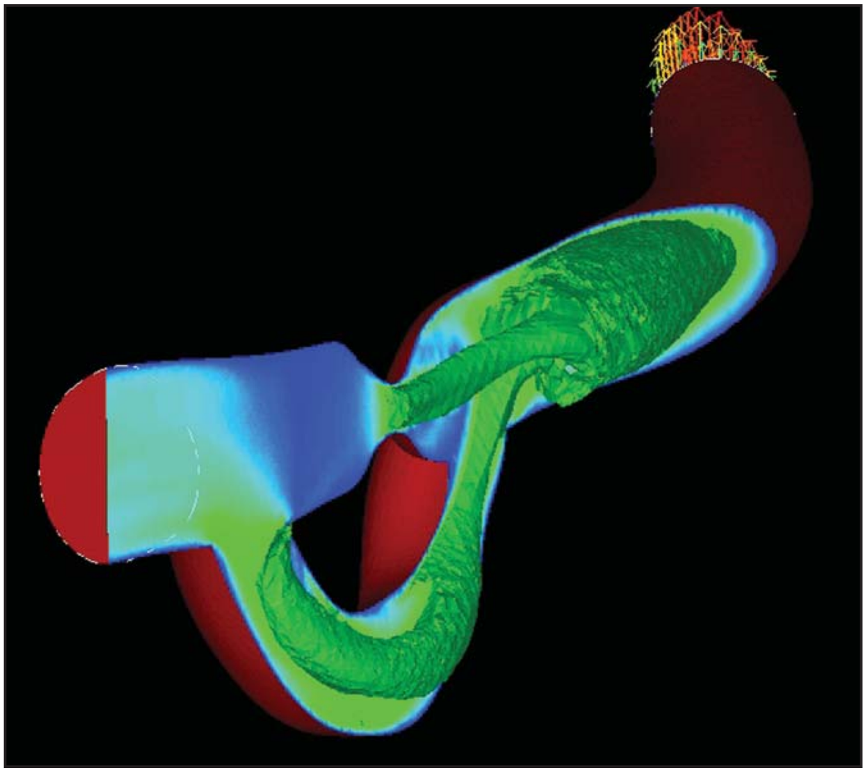

Computational methods are now being used to create preoperative surgical models for heart bypass patients. **See page 10**

FEATURES

Features

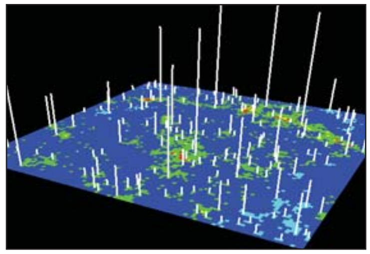

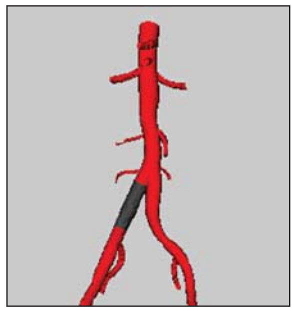

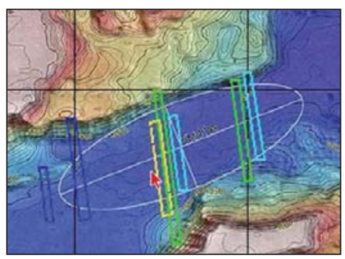

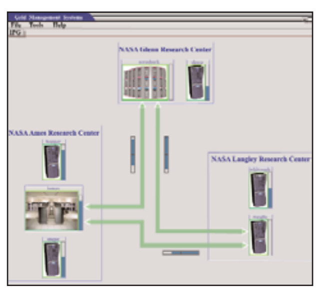

Also in this issue:

On The Cover:

Computational simulation of a pig's aorta with an implanted bypass graft, during the acceleration phase (systole) of the cardiac cycle. The curved model represents a 2-D slice of the velocity magnitude, plus an iso-surface for a given value of the velocity magnitude. Red defines the outside of the artery wall, while green denotes high blood flow velocity. (Visualization by Valerie Favier)

Gridpoints
The Quarterly Publication of the NASA Advanced Supercomputing Division
Vol. 3, No. 2 • Summer 2002

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: John Adams, Dick Anderson, Bryan Biegel, Kathleen Cain, Glenn Deardorff, Jennifer Dungan, Valerie Favier, Michael Frumkin, Chris Gong, T.R. Govindan, Virginia Gulick, Chris Henze, Esther Johnson, Randy Kaemmerer, David Kao, Stephanie Langhoff, Jeff Leibow, Tony Lisotta, George Myers, Marcia Redmond, Cynthia Rodriguez, Warren Smith, Allan Snavely, Deepak Srivastava, Charles A. Taylor, Rob Van Der Wijngaart, David Walford, Mary Walsh, Steven Zornetzer

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: Subscribers outside the continental United States can access *Gridpoints* on the web in PDF format at: http://www.nas.nasa.gov/gridpoints

News From NAS

Ziebarth Appointed Acting Division Chief

The NAS Division welcomed John Ziebarth to its helm as Acting Division Chief in April, following the move of Bill Feiereisen to director of Computer and Computational Sciences at Los Alamos National Laboratories. Ziebarth has been NAS's deputy division chief since 1999 and promises to move the division forward in its mission to serve as the flagship of supercomputing innovation.

"John Ziebarth comes to the acting division chief position with a wealth of experience in high-end computing (HEC) and networking," says Steven F. Zornetzer, director of Information Sciences and Technology at Ames. "John's leadership is important during this dynamic period of changing paradigms from vector to parallel machines and architectures and the rapid evolution in the Information Power Grid community. The NAS Division has a rich heritage of national and international leadership in HEC, which needs to be reasserted. We look to John to provide this leadership."

Ziebarth obtained his doctorate in Aerospace Engineering from Mississippi State University in 1983 and formerly served as director of the NASA Consolidated Supercomputing Management Office (CoSMO) and an associate

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

I am honored to serve NAS and Ames Research Center in this new capacity. For the past three years, I have served as deputy division chief. Two years prior to that, I was director of the Consolidated Supercomputing Management Office. During this time, the capabilities of our high-end computers has increased tremendously, enabling our researchers to make significant progress in computational sciences. This progress is reflected in a number of the division's highlights, which are listed below and detailed within the pages of this issue:

In April, the DeBakey Ventricular Assist Device (VAD) was named NASA's Commercial Invention of Year for 2001 (see page 3). NAS Division researchers developed and implemented a number of modifications for the VAD using computational fluid dynamics, enabling the device to run without failure for more than 120 days.

Additionally, the HiMAP (High Fidelity Multidisciplinary Analysis Process) software tool, developed by division researchers, has received the NASA Space Act Software Release award (see page 2). The American Institute of Aeronautics and Astronautics' Fluid Dynamics Technical Committee named a NAS collaborative paper, "Multilevel Error Estimation and Adaptive h-Refinement for Cartesian Meshes with Embedded Boundaries," best paper for 2002 (see page 4).

In the hardware arena, *Chapman*, our 1,024-processor machine, has been migrated from a prototype to a production machine. This is the world's largest single-system image supercomputer – it will enable the division to more than quadruple the division's computing cycles during the next year. In 2002, we will run more than 3 million production cycles, and we expect to grant nearly 14 million in 2003.

Recently, cracks in the Space Shuttle's fuel delivery system liners were discovered, grounding the fleet until a solution is approved. Since the end of June, the division has been supporting the effort of returning the Space Shuttle to flight status. Researchers from Marshall Space Flight Center have used more than 240,000 processor hours on *Chapman* to determine the cause of the cracks, their impact on safety, as well as the pros and cons of each repair choice. Selection of the final repair plans is being made as this publication goes to press.

Chapman's computational power and multi-level parallelism programming techniques are now being applied to POP, the Parllel Ocean Program developed by the Department of Energy's Los Alamos National Laboratory. POP is part of the Community Climate System Model that encompasses water circulation, sea ice, land/surface heat, atmospheric air circulation, the sun's radiation, and Earth's radiation into space, to simulate how the world's oceans affect the Earth's climate. NAS Division researchers have scaled the code to run on all 1,024 of *Chapman's* processors, and are working to further optimize the code. Today, it runs on 900-processors, and can simulate data for three-quarters of a year per day – an order of magnitude improvement over previous runs.

I hope you enjoy this issue, and always, I welcome your feedback.

John Ziebarth
jziebarth@mail.arc.nasa.gov

Continued from page 1
director at the National Center for Supercomputing Applications (NCSA).

"My background in high-end computing and computational engineering and science will allow me to work well with the top-notch scientists in this division," says Ziebarth. "Many firsts have come out of the division, and I think part of our goal is to continue to be innovators and leaders in the HEC field."

Ziebarth says he's confident that NAS will continue to achieve significant milestones while supporting research projects throughout the agency, despite the challenges of declining budgets and changes in the way NASA does business. To support Ziebarth in his leadership role, John Parks has agreed to take on the post of acting deputy, bringing with him more than 20 years of experience in HEC systems at Ames. Drawing on Parks' experience for daily operations issues, Ziebarth says, will allow him to spend more time on the research side of the division.

Barth Appointed to Springer-Verlag Editorial Board

Timothy Barth, a computer scientist in the NAS Division's Physics and Simulation Modeling Office, has been appointed to the editorial board of a book series published by prestigious international scientific publisher Springer-Verlag.

Established in 1997, the series, "Lecture Notes in Computational Science and Engineering," is a part of the publisher's mathematics division. Barth has been a contributing author to two editions of the lecture series. The most recent, released in January, is devoted to the topic of multi-scale and multi-resolution methods. All editions are in some form related to computational science and engineering (CSE), which is the application of applied mathematics and computations.

"Springer is really into CSE – they are embracing this as a new subject area, even though it's not a very focused subject. It's all applied mathematics with an emphasis on computations," Barth explaines.

His 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.

Growing the Next Generation of Computers

NAS researchers Deepak Srivastava and Chris Henze have been awarded the Ames Research Center Director's Discretionary grant for a new proposal encompassing nano-, bio-, and information technology. Their proposal for the "biomimetic simulation of signal transmission in nanotube-based dendritic networks" represents an innovative step forward in the division's work on nanotechnology. Inspired by the structure and operations of biological neural systems, the scientists plan to investigate the properties of a 3-D, tree-like architecture of carbon nanotubes, and prove its viability.

The field of molecular electronics using carbon nanotubes has progressed rapidly in recent years, resulting in several newly patented technologies, which experts believe will ultimately lead to the next generation of electronic computing devices. The National Science Foundation estimates that by 2015, the market for products manufactured using nanotechnology will reach $1 trillion, and in his latest budget proposal, President Bush included a 17 percent increase in spending on nanotechnology research.

Henze and Srivastava's proposal of a 3-D, tree-like circuitry takes a unique approach to current research by private industry and academia, which has so far focused on 2-D models. "We recently showed that three-terminal nanotubes grown as a y-junction serve as rectifier switches and have analog logic capacity," says Srivastava, senior scientist and lead of computational nanotechnology investigations at NAS. "These properties inspired us to consider the concept of a system architecture similar to the biological neural system, but made up of synthetic material."

Taking data derived from activating the nerve cells of crickets, Henze, a senior research scientist, can simulate their electronic behavior. Such simulations are a major tool for determining exactly how signals are transformed as they pass through the branch structures of neurons. "What's becoming clear is that there's a lot of non-linear processing going on in the complex branching structures," says Henze, "and it has a lot to do with the geometry of the branching." This observation, combined with Srivastava's simulations in building branched nanotube systems, led to a confluence of ideas for the proposal that the scientists hope will ultimately demonstrate how signals could be transmitted and processed as they pass through a branched-tree system made of carbon nanotubes.

The Director's Discretionary Fund is a well-established grant that provides seed money for innovative and high-risk research that, if successful, could lead to fundamental scientific breakthroughs. The NAS scientists hope that showing a high level of functionality will inspire other groups to try growing the tree and eventually lead to very complex sensing and computing applications.

HiMAP Software Wins Space Act Award

A team led by Guru Guruswamy of the NAS Division at Ames Research Center recently received the 2002 NASA Space Act Software Release award for developing HiMAP, the High Fidelity Multidisciplinary Analysis Process. HiMAP software efficiently integrates software analysis tools

Life Saving Device Named Invention of the Year

Resulting from the collaborative efforts of researchers at NASA Ames, NASA Johnson Space Centers, and MicroMed Technology, Inc. of Houston, a life-saving heart pump known as the DeBakey Ventricular Assist Device (VAD) was named NASA's 2001 Commercial Invention of the Year. The Invention of the Year Award program was developed to help promote part of NASA's mission, which is to transfer space-based technology to the private sector. Exclusive rights to the heart assist device were granted to MicroMed in 1996.

To date, the VAD has been successfully implanted in 115 individuals in Europe with no reported device failures. Trials in the United States are currently underway – 180 implantations of the device are planned. The overwhelming success of the heart assist device is largely due to the integration of several design modifications recommended by computer scientists Cetin Kiris and Dochan Kwak of the NAS Division at Ames.

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

Drawing on previous experience with modeling fluid flow through rocket engines, the team carefully examined seven different designs of the assist device by simulating blood flow patterns through the pump using NASA-developed computational fluid dynamics (CFD) technology. Specifically, the team used unsteady turbopump simulation technology originally developed for the Shuttle's main engine. "Although the technology applied here comes from technology developed for space transportation systems analysis, biofluid simulations require quite a bit of additional research," explains Kwak, chief of the NAS Applications Branch.

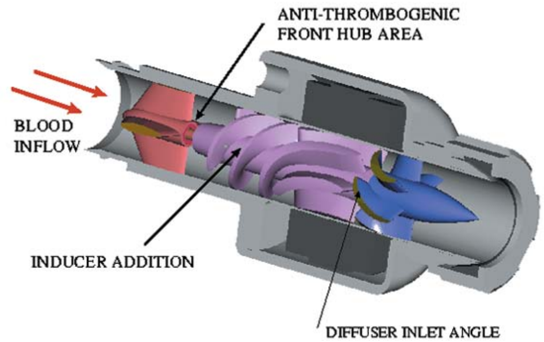

Using CFD analysis, NAS researchers Dochan Kwak and Cetin Kiris found that major design modifications to the DeBakey Ventricular Assist Device were necessary. The result of these changes: overall efficiency of the device was increased by 22 percent. (Courtesy MicroMed Technology)

Altering several components of the VAD, including cavity shape and impeller tip clearance size, led the researchers to propose three design modifications. One design modification was the addition of an inducer that spins with the impeller to eliminate irregular blood flow patterns, which led to blood clotting. The second modification was the alteration of the cavity shape between the flow straighteners and the inducer. The third change was the modification of the diffuser angle to adjust the flow angle of blood exiting the device. The combination of these three design modifications eliminated cell damage, while improving blood flow through the device.

Currently, the device is used as a "bridge to transplant" for heart failure patients awaiting a transplant. It is also used as a "bridge to recovery" for those patients whose hearts are healthy enough to heal naturally, and do not require a transplant. In the future, the team hopes to further improve the design so that it can be permanently implanted in humans, circumventing the need for any heart transplant.

to solve large-scale multidisciplinary problems on massively parallel supercomputers.

HiMAP integrates disciplines with diverse physical characteristics by retaining the efficiency of individual disciplines. Results are demonstrated for large-scale aerospace problems on several supercomputers. HiMAP has been successfully used by NASA's High Speed Civil Transport (HSCT), the Defense Advanced Research Projects Agency's Unmanned Combat Air Vehicle project, the Navy's Abrupt Wing Stall project, as well as Boeing's and Lockheed Martin's Internal Research and Development projects. "Anticipated savings over the next fiscal year to a number of aircraft programs is expected to be in the $100,000s. The added analysis capability due to the HiMAP software provides a capability that previously did not exist within NAVAIR," said Dr David Findlay, Manager, Naval Air Systems Command. Engineers at Sun Microsystems have ported HiMAP to the latest parallel systems running Sun HPC ClusterTools 4 Software.

HiMAP technology could prove useful in other fields such as automotive, mechanical, civil- and bio-engineering, where analysis of fluids and structures interactions plays an important role. Potential future applications include analysis of floating runways (similar to that under consideration at San Francisco airport), modern cable-stayed girder bridges, and human hearts.

Researchers Awarded for Development of New Mesh Refinement Strategy

On June 25, the American Institute of Aeronautics and Astronautics (AIAA) recognized the work of NAS Division researcher Michael Aftosmis and Marsha Berger, professor and deputy director of Courant Institute at New York University, in the area of error estimation and adaptive meshing. The team's paper, "Multilevel Error Estimation and Adaptive h-Refinement for Cartesian Meshes with Embedded Boundaries," has been named best paper by the AIAA Fluid Dynamics Technical Committee for 2002. The award was presented at a special luncheon held in conjunction with the 32nd AIAA Fluid Dynamics Conference in St. Louis, Missouri.

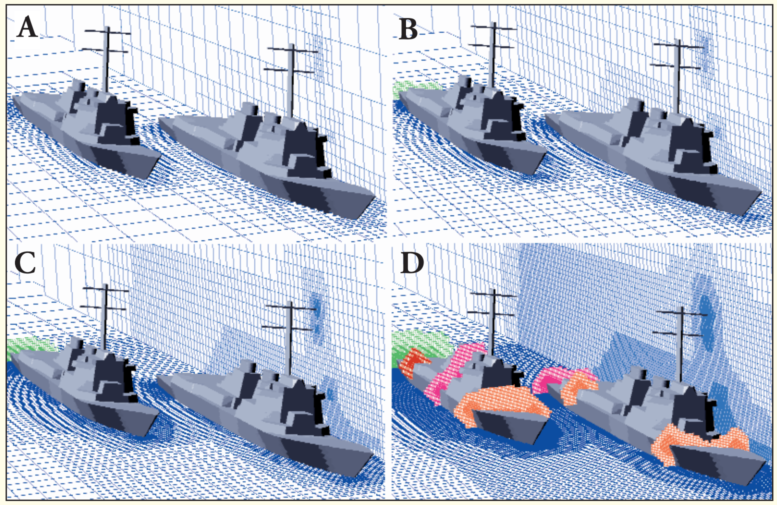

The sequence of grids (A through D, above) surrounds two advanced destroyer-type surface ships. These simulations were generated using Cart3D's flow solver module, which employs a multigrid technique to improve the efficiency of the simulation.
(Michael Aftosmis and Donovan Mathias)

The winning technical paper describes Aftosmis' and Berger's development of new techniques for error estimation and adaptive refinement for computational fluid dynamics (CFD) solutions. The new adaptive mesh refinement, or h-refinement technique, provides an optimal mesh refinement strategy for flow field adaptive CFD solutions. Meshes, or grids, are used to calculate flow fields surrounding vehicles, such as the Space Shuttle. "The novelty in our work is that it removes the ambiguity of mesh generation and identifies the single best solution," says Aftosmis.

Several features put this new h-refinement strategy a cut above previous adaptive meshing techniques. "It is general enough that it can be applied to any type of unstructured or hierarchical mesh and, since it is parameter free, it can be fully automated," explains Aftosmis. The new meshing strategy was also designated to remove the dependence on user skill, making it very quick and easy to generate meshes. CFD groups at Department of Energy laboratories, Michigan State University, and Massachusetts Institute of Technology have been experimenting with the new error estimator and refinement criteria. Aftosmis and Berger plan to integrate these error estimation and refinement techniques into their grid generation and solution software package, Cart3D, this fall.

Nanotechnology Research Published

A new take on a well-known nanotube experiment was recently published by NAS researcher Toshishige Yamada in *Applied Physics Letters* (May 2002). Yamada's paper, "Modeling of Kink-Shaped Carbon-Nanotube Schottky Diode with Gate Bias Modulation," challenges findings of a gate voltage modulation experiment by a research group at Delft University of Technology, The Netherlands.

In 1999, the Delft researchers published the first experiment clearly demonstrating that a transistor action principle called the "gate modulation effect" survives at the atomic scale as an intermolecular device made of up only of nanotubes. However, the team could not conclude whether the nanotube was p-type (positive-carriers rich) and where the rectification took place. Yamada analyzed their data and clarified that the nanotube must be n-type (negative-carriers rich), rather than p-type, as was previously believed. He also showed that rectification occurred at the kink part of the nanotube, rather than at the electrode contact.

Yamada's work is meaningful from a pure physics standpoint, and helps solve a key issue in semiconductor electronics. All semiconductors are either p-type or n-type. Today's semiconductor electronics require a sophisticated combination of both types to create transistors. "Knowing whether a semiconductor is either p-type or n-type is the first step towards nano-transistors and nano-electronics," explained Yamada. "The next issue to tackle is why the Delft group had an n-type nanotube, while others are seeing p-type nanotubes routinely in their laboratories."

Other experimental nanotube devices fabricated in research labs have all been hybrids of a semiconducting nanotube and macroscopic metallic electrodes, and will not guarantee the smallest miniaturization, according to Yamada. His paper shows that the Delft device made solely of nanotubes operated with the same transistor action principle used in today's devices, and strongly suggests the feasibility of an "all nanomaterial transistor" in the future.

Setting the Standard for Grid Benchmarking

Growing enthusiasm for grid computing among the business and research communities has prompted participants of the Global Grid Forum (GGF) to turn to NAS scientists for guidance on how to best measure grid efficiency and user-friendliness.

With a background rich in experience following the success of the NAS Parallel Benchmarks (NPB) for high-performance computers, NAS scientists Michael Frumkin and Rob Van Der Wijngaart have just received GGF approval of their charter to establish a formal research group on grid benchmarks.

Official approval means the scientists can hold regular sessions at GGF workshops and continue their work developing grid benchmarks in collaboration with other organizations interested in developing global grid standards.

"Benchmarking is a long-standing trade at NASA," explains Van Der Wijngaart. "Computational grids are really at the same stage as parallel computers were about 10 or 15 years ago. So we thought, 'let's take these NAS Parallel Benchmarks, which are highly respected, and use them as building blocks to construct more complicated distributed tasks that need to share information with each other on computational grids.'"

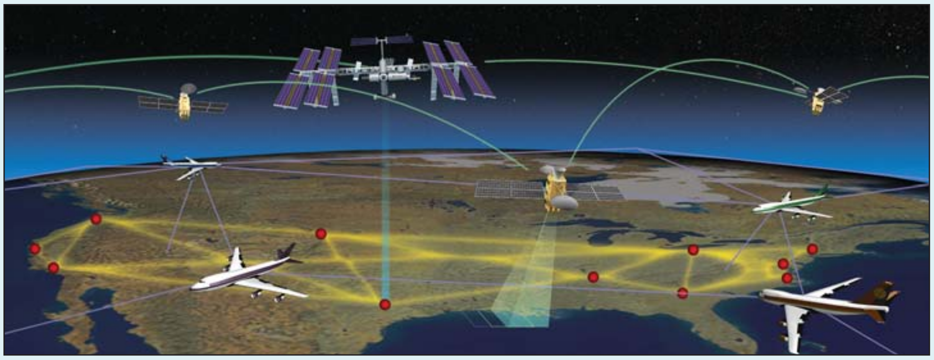

Artist's conception of the Information Power Grid in action. (IPG/NAS Division)

Although many grids already exist, including NASA's Information Power Grid, the European DataGrid, UNICORE, Seti@home, and Condor, there are currently no generally accepted tests in place to determine the efficiency of the varying standards and toolkits, like Globus and Legion.

"We need to find a common denominator to test the best solution," says Frumkin. "All these grid communities embrace the idea of benchmarking because it gives them the opportunity to show that their products work well."

So far, NAS scientists have defined four different model grid applications, derived from common NASA operations, such as flow visualizations and parameter studies. For each of these applications, they have provided exact specifications, resulting in the NAS Grid Benchmarks (NGB). For ease and convenience, they have also developed some reference implementations of the NGB, using grid tools developed using Java and Globus.

Frumkin and Van Der Wijngaart first presented an implementation of their suite of benchmarks at GGF3 in Rome last year, and although the presentation was late in the afternoon and unannounced, the meeting room was packed. This large show of interest led to a "birds of a feather" session in Toronto this past February, and contact with other groups interested in collaboration.

Allan Snavely, the Performance Modeling and Characterization Laboratory leader at the San Diego Supercomputing Center, says that several months ago his group also recognized that the lack of ways to measure grid performance posed a real problem for grid developers. "We have a window of opportunity to develop grid metrics and influence the design of grid systems in a very positive way," says Snavely. "We can't leave it aside – it just won't be as interesting or compelling in a year's time."

Snavely's team is taking a different approach to develop benchmarks that can measure performance and compare configurations. By talking and working with groups of application developers, they hope to develop a skeleton suite of applications that will represent the ways applications exercise and stretch grid systems. Snavely sees this approach as complementary to NAS's, and will co-chair the GGF research group with Van Der Wijngaart.

The next step for NAS, says Van Der Wijngaart, is to generate new problem sizes for each family of benchmarks so that tasks can be scaled up. For the NAS research group as a whole, Van Der Wijngaart believes that some hard and disciplined thinking now needs to take place, especially focusing on areas not yet covered by the NGB, like data-intensive applications.

But Frumkin is optimistic that ultimately the grid benchmarks will generate useful results and guide tool developers to write more efficient code. "When the NAS parallel benchmarks were released, everyone tried to show that their machine was better because it took less time to successfully execute the benchmarks. We are expecting the same thing to happen with the grid benchmarks."

— *Julie Jervis*

Visualizing Uncertainty in Earth Observing System Satellite Data

NASA researchers are developing sophisticated visualization tools for displaying statistical summaries that will help Earth scientists interpret variations in satellite data.

While scientists continue to explore the mysteries of the Earth's biosphere, innovative technology from the NASA Advanced Supercomputing (NAS) Division is helping researchers unravel the complex web of interactions between natural phenomena and human activity on our planet.

Using visualization tools and techniques developed by the NAS Division, scientists in NASA Ames' Ecosystem Science and Technology Branch are working to gain a better picture of the Earth's surface. The goal is to gain a better understanding of the Earth's primary forces, its climates and ecosystems, and the effects of human activity on them. Ultimately, scientists hope to be able to predict change in the Earth's biosphere and its consequences for human civilization.

To achieve this monumental task, NASA scientists have been studying land, climate, and environmental changes using images from a variety of sensors in low-earth orbit. One of these sensors, the Moderate Resolution Imaging Spectro-

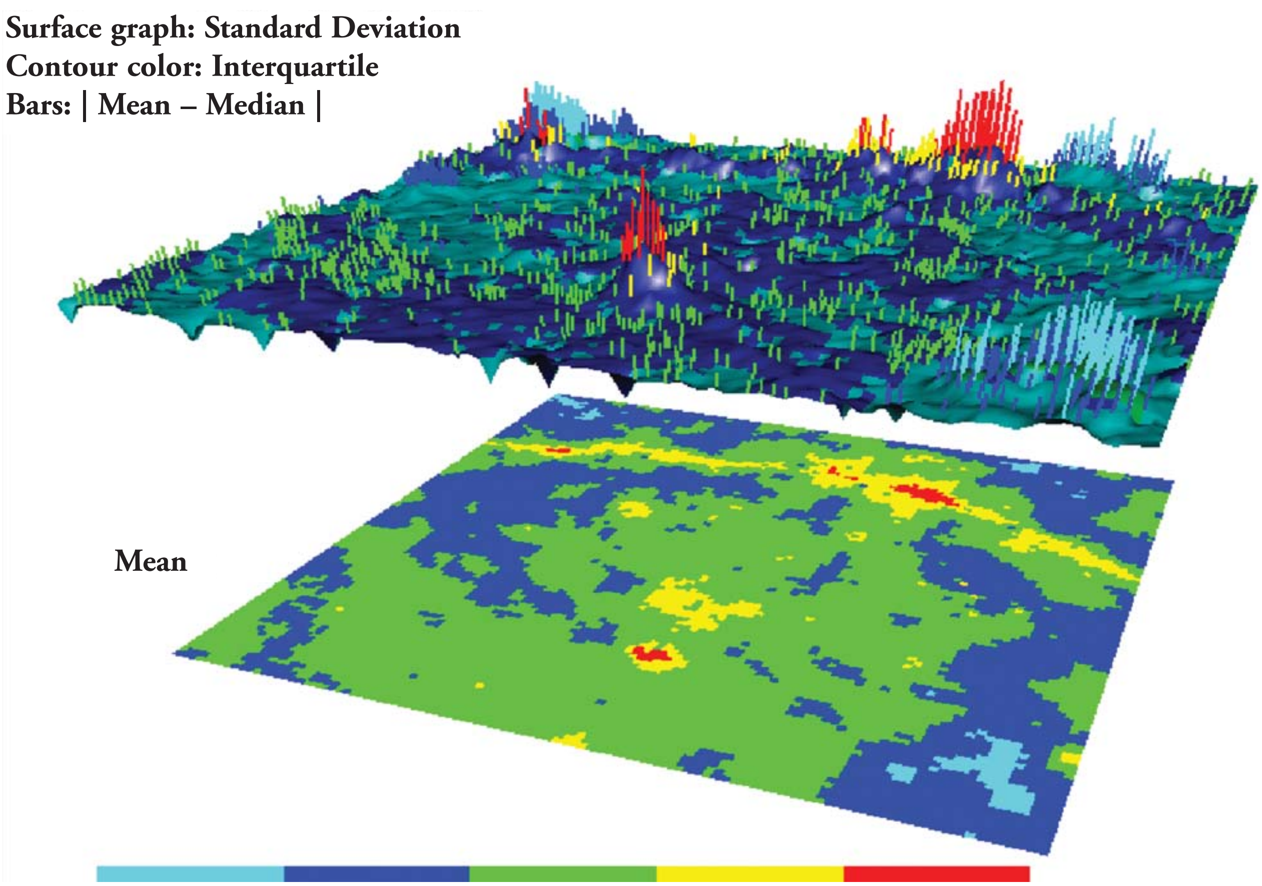

Figure 1: *Earth Observing System (EOS) satellite data interpreted using David Kao's new visualization techniques: The bottom plane shows the mean field colored from non-forest (blue) to closed forest (red). The upper plane is generated from three measurements: the surface is deformed by the standard deviation and colored by the interquartile range, and the heights of the vertical bars represent the difference between the mean and median measurements, similarly colored. Places where mean-median is largest are not parametric, so this is where the interquartile range shows a better picture of uncertainty than standard deviation. (Data courtesy Jennifer Dungan, Visualization by David Kao)*

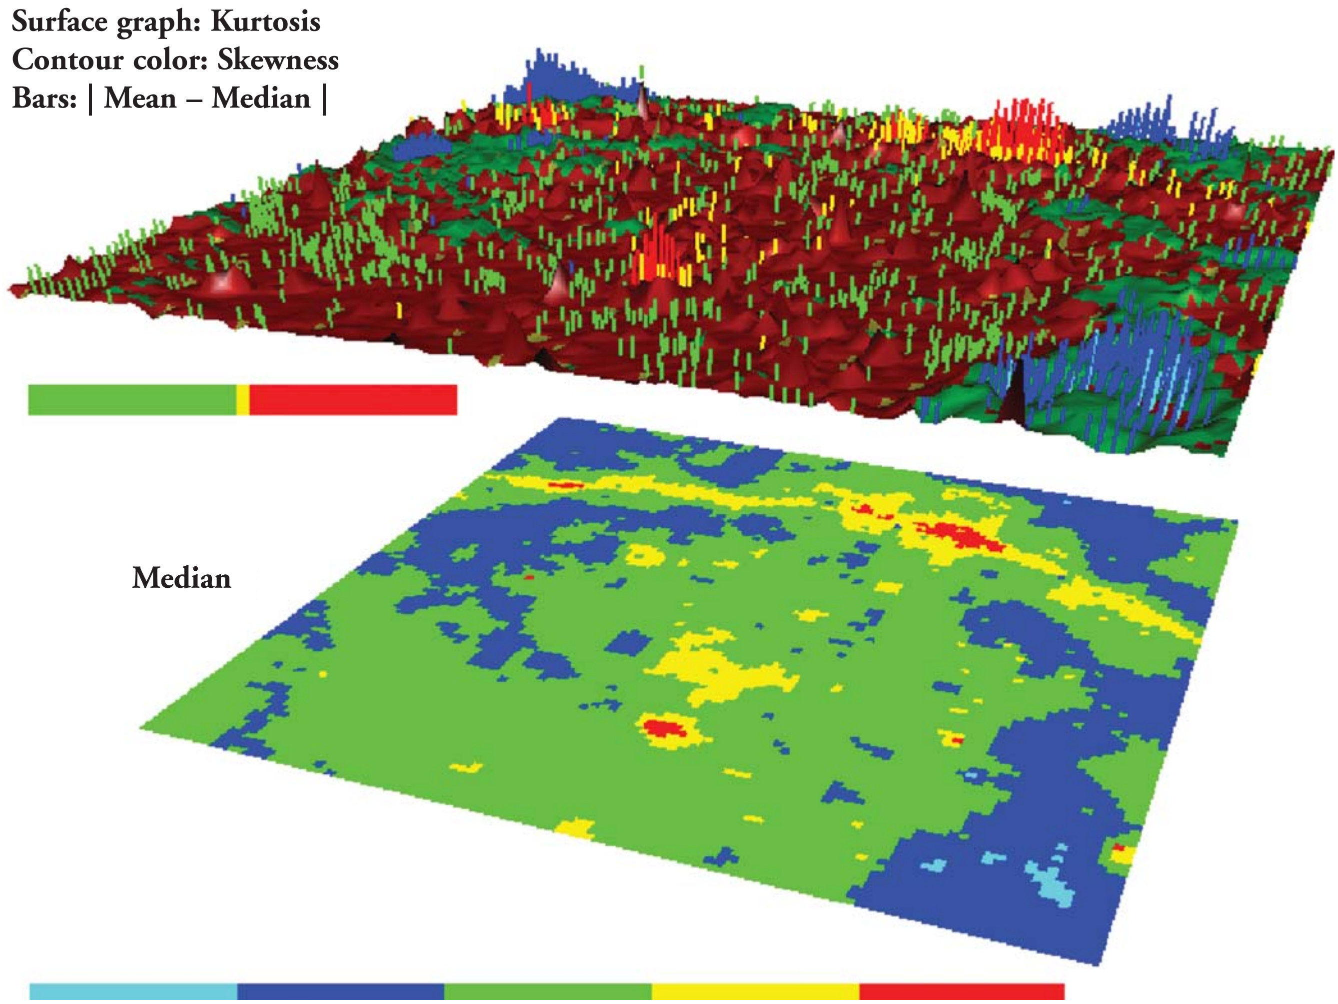

Figure 2: *The bottom plane shows the median field, and the upper plane is deformed by the degree of peakedness (kurtosis) and colored by the skewness measurement (green denotes negative and red denotes positive). Areas where skewness is large are also places where an understanding of uncertainty is not gained through standard deviation alone.*

(Data courtesy Jennifer Dungan, Visualization by David Kao)

radiometer (MODIS), was launched on the Terra satellite in 1999. Terra and its sister satellite Aqua, due to go live this fall, are part of the Earth Observing System (EOS), which provides scientists with daily images of the Earth's surface. With resolutions of 250 by 250 meters per pixel or smaller, the EOS creates huge challenges in charting the entire Earth's surface, especially monitoring changes over time. One of the challenges is how to detect and interpret errors or uncertainties in this immense quantity of data.

Jennifer Dungan, a research scientist in NASA Ames' Ecosystem Science and Technology Branch explains: "NASA's EOS satellites are producing large quantities of images that represent reflected light within a spectral region. Measurements of the amount of light reflected have error in them, which could be due to several factors – including sensor error or atmospheric contamination – and that adds uncertainty. We're interested in getting a gestalt view of this uncertainty, as well as being able to query specific variables." Identifying this uncertainty is not just important for mapping, adds Dungan, but could have a major impact on international policies, such as the Kyoto Protocol, a United Nations' initiative to limit the amount of greenhouse gases emitted into the atmosphere.

NAS researcher David Kao has been working with Dungan, Professor Alex Pang of the University of California at Santa Cruz, and Professor Han-Wei Shen at Ohio State University for more than a year on innovative visualization techniques for analyzing distribution data. Jointly, they responded to a NASA Intelligent Systems Program research announcement with a proposal to compare and understand distributions in EOS applications. Despite strong competition, the project team was awarded a three-year grant under the Intelligent Data Understanding topic led by Joseph Coughlan at NASA Ames, and started work in March 2001. To date, research has resulted in compelling new visualization tools that show the distribution of data for each pixel in each grid cell of a map in ways Earth scientists have not seen before.

"Scientists currently use common geographic information system packages that only allow them to look at one realization –

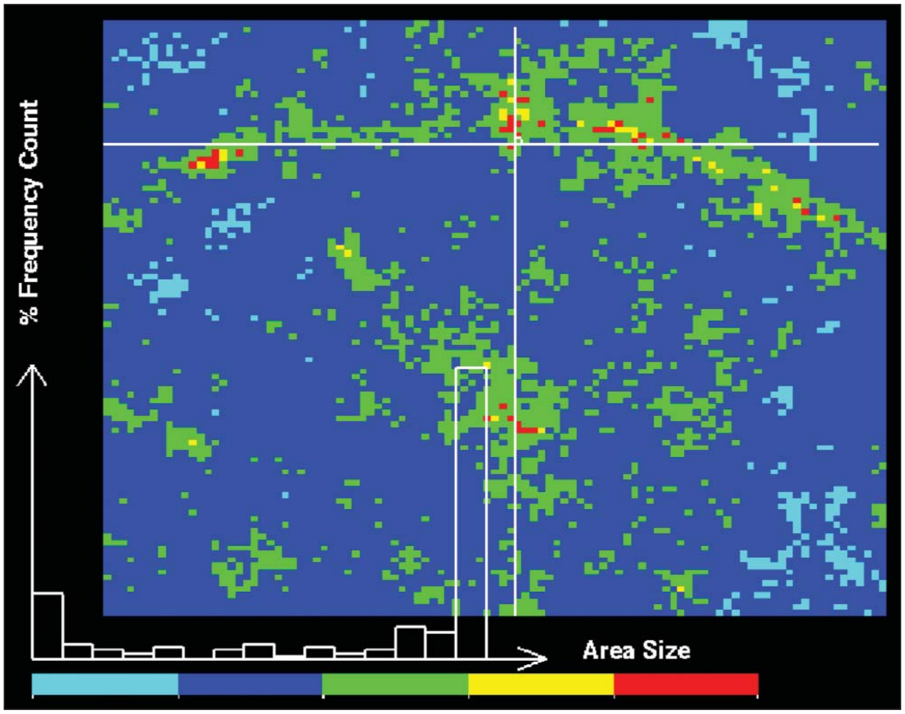

Figure 3: *Positioning the probe (represented by the white crosshair) in a chosen location provides a histogram of the surrounding area. The x-axis represents the area size and the y-axis represents the frequency as a percentage, that is, the number of realizations found with the same class at the probe location divided by the total number of maps. The histogram of clump area at current probe position shows that the green clump's area is likely to be high among all of the realizations.*
(Data courtesy Jennifer Dungan, Visualization by David Kao)

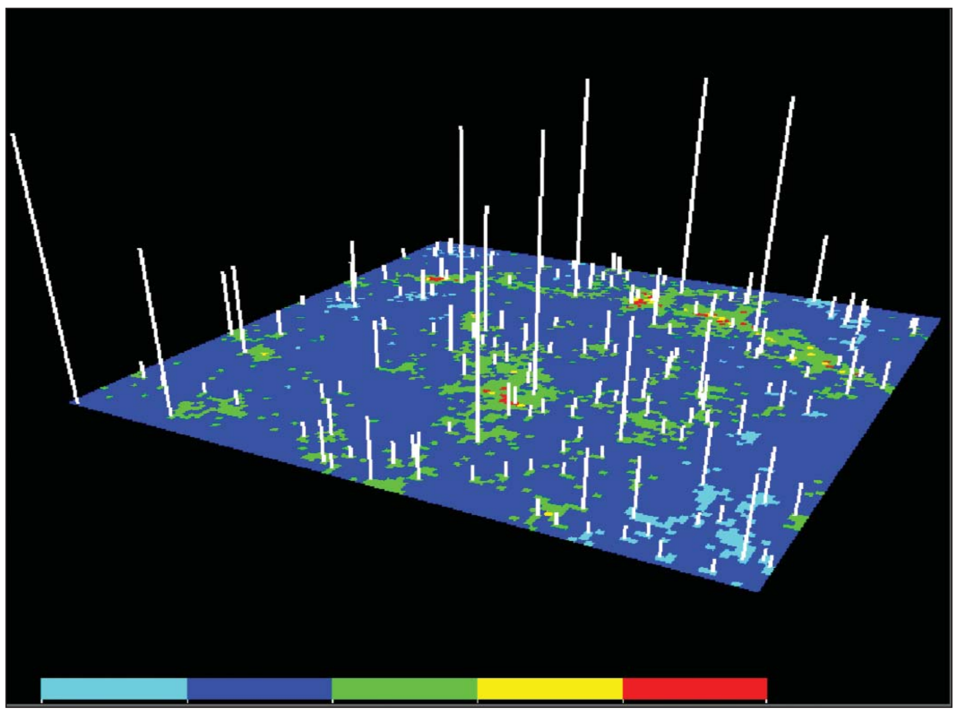

Figure 4: *Vertical line bars provide visual markers for each clump. The length of these lines indicates the size of the area, a visual guide that, unlike traditional color-coding, is less likely to be fooled by complex or circuitous shapes.*
(Data courtesy Jennifer Dungan, Visualization by David Kao)

which is a likely outcome of a map – at a time," says Kao. "Our tool not only allows them to associate information among hundreds of realizations, but can display four different statistics simultaneously."

Using distribution data derived from satellite images, Kao's team developed two approaches to visualize such data: pixel-wise summaries and feature-wise summaries. When presented with a set of possible numbers in one pixel (a distribution), the standard approach is to compute the mean and the standard deviation of the given data. Kao's software computes several statistical summaries, including the mean, median, standard deviation, interquartile range, kurtosis, and skewness at each pixel across the realizations (see Figure 1, page 6 and Figure 2, page 7). This provides the scientist with a pixel-wise summary of measurements that are color-coded onto the whole image and can be displayed separately or simultaneously in the same viewing space. This enables the scientist to study relationships among the measurements at each pixel.

In addition to examining uncertainty about individual pixels, scientists are interested in understanding clumps of pixels with similar values. These clumps may be categorized into classes, such as forest, grasses, cropland, and water, depending on the resolution of the image. Clump information, represented as a feature-wise summary, helps scientists identify the areas and shapes of individual clumps, as well as understand the uncertainty about the area. “When you navigate with a highly accurate road map, you can see when you’ll get to the end of the road,” says Dungan. “When you’re on the trail and using a map that’s colored according to what is on the Earth’s surface, you need to be able to understand why an area that’s mapped as grassland is actually forest. A map is just a model of the world and we need information about the quality of that model to determine the plus or minus percentage of error at each grid cell.” To help researchers analyze clump area statistics, Kao and his team have designed tools that allow scientists to both interactively interrogate the clump area’s statistics with a crosshair probe, and highlight the sizes and locations of clumps (see figures 3 and 4, page 8).

Dungan and Kao discussed their work at the July 2002 NASA-sponsored MODIS Vegetation Workshop. Hosted by the University of Montana in Missoula, the MODIS Land Science shared step-by-step details of the processing, distribution, analysis and interpretation of the MODIS vegetation variables to the Earth science community (see figure 5). In addition, Kao’s team continues to translate visualization ideas and techniques into methods and algorithms for understanding distribution data. The next stage, Kao says, is to investigate ways to describe both visually and numerically the differences in the distribution of data between pixels and features.

— *Julie Jervis*

For more information, see: David Kao, Jennifer Dungan, and Alex Pang, "Visualizing 2D Probability Distributions from EOS Satellite Image-Derived Data Sets: A Case Study," in Proceedings of *IEEE Visualization '01*, pages 457-460.

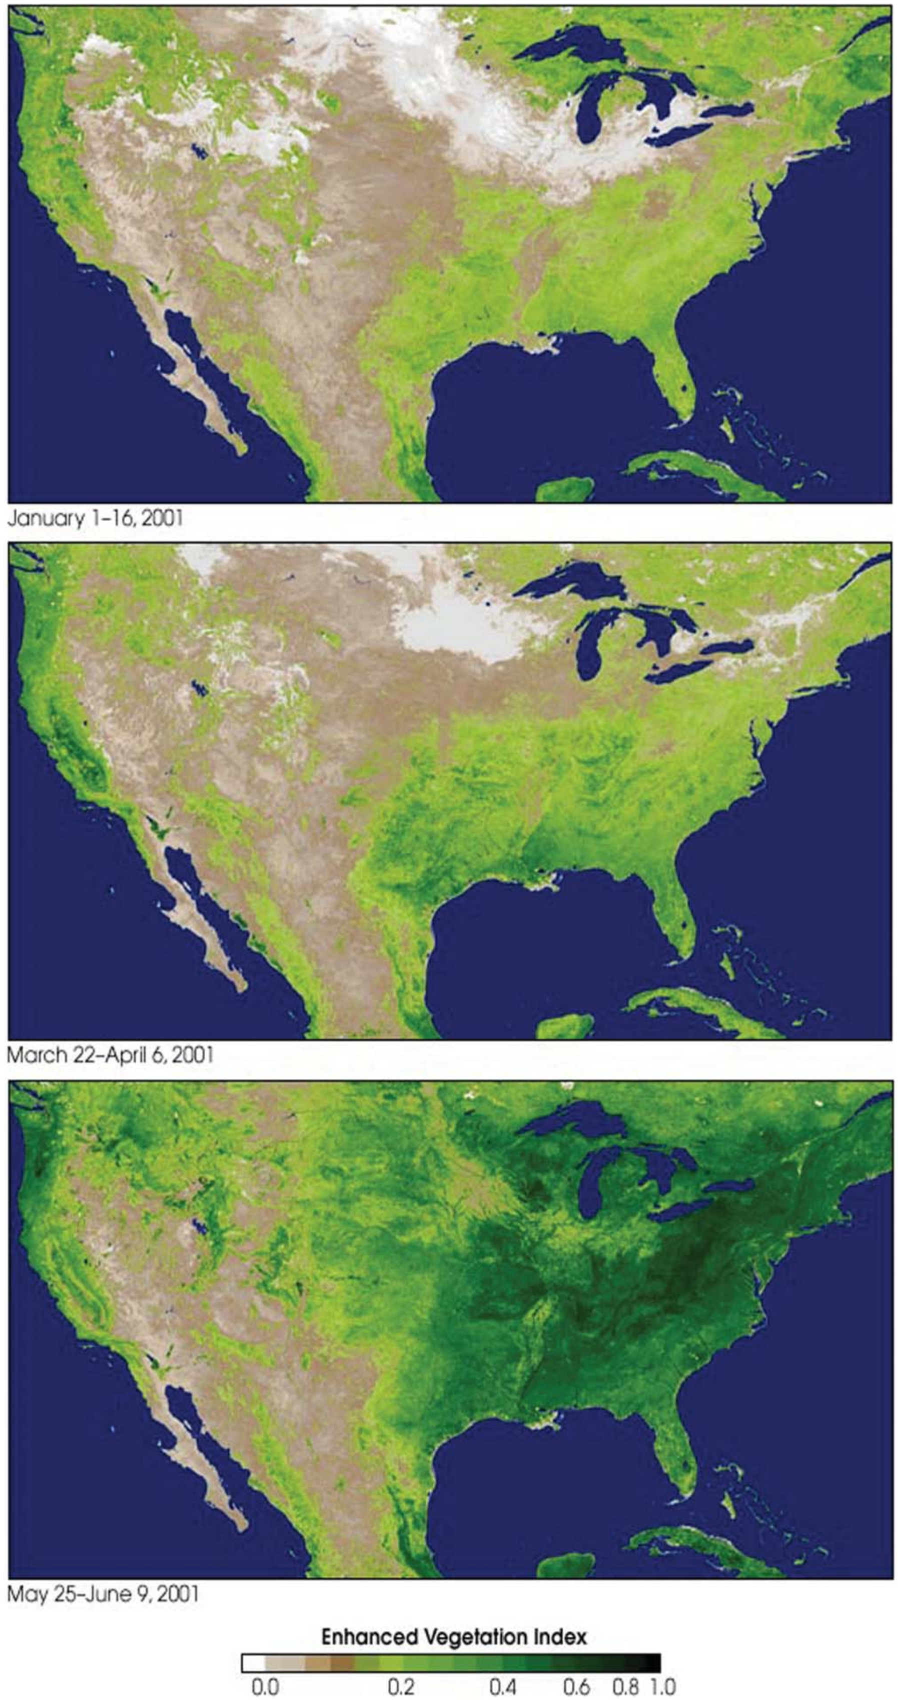

Figure 5: *MODIS Enhanced Vegetation Index observations from January to June 2001 for the United States show the cycle of vegetation from winter’s nadir to the new growth of spring. Vegetation ranges from 0, indicating no vegetation, to nearly 1, indicating densest vegetation. (Courtesy of NASA-Goddard/University of Arizona)*

Computational Models to Predict Outcome of Vascular Surgery

Scientists at NASA's Center for Turbulence Research and Stanford University are using magnetic resonance imaging data to create numerical flow simulations of blood flow through vessels of patients with cardiovascular disease.

Ranking number one on America's "top ten list" of killers, heart disease claims the lives of nearly one million Americans each year. At this very moment, 61.8 million Americans are battling with a cardiovascular disease – one of these people dies every 33 seconds. Surgeons perform vascular surgery on seven million patients every year, with no way to predict, with certainty, the effectiveness of a particular treatment.

Charles A. Taylor, assistant professor of surgery and mechanical engineering at Stanford University, Valerie Favier, a Stanford post-doctoral student funded by the Center for Turbulence Research (CTR), which is co-managed by the NASA Advanced Supercomputing (NAS) Division, and colleagues at Stanford University are working to cut the guesswork out of cardiovascular surgery. The group is creating computational models of the circulatory system in patients with cardiovascular diseases (outlining aspects such as vorticity, recirculation areas, and flow stagnation) to assist doctors with planning future surgical procedures. "This is something that will really impact vascular and cardiovascular surgery – it is a better way of planning treatments for patients," says Taylor.

When creating these computational models, one of the group's primary focuses is on the connection between the way blood flows through the circulatory system and how disease forms in the body. Using computational methods, Taylor and his group have been studying blood flow patterns in disease-prone areas of the bloodstream for the past eight years. "How blood flows through the circulatory system doesn't just affect where we get disease, but how well treatments work," says Taylor. One of these unfortunate locations is the carotid bifurcation, a branch point in the artery going up to the brain. Plaque deposits form more readily at this location, because its unique geometry causes backflow of the blood. When plaque ruptures, it blocks blood flow to the

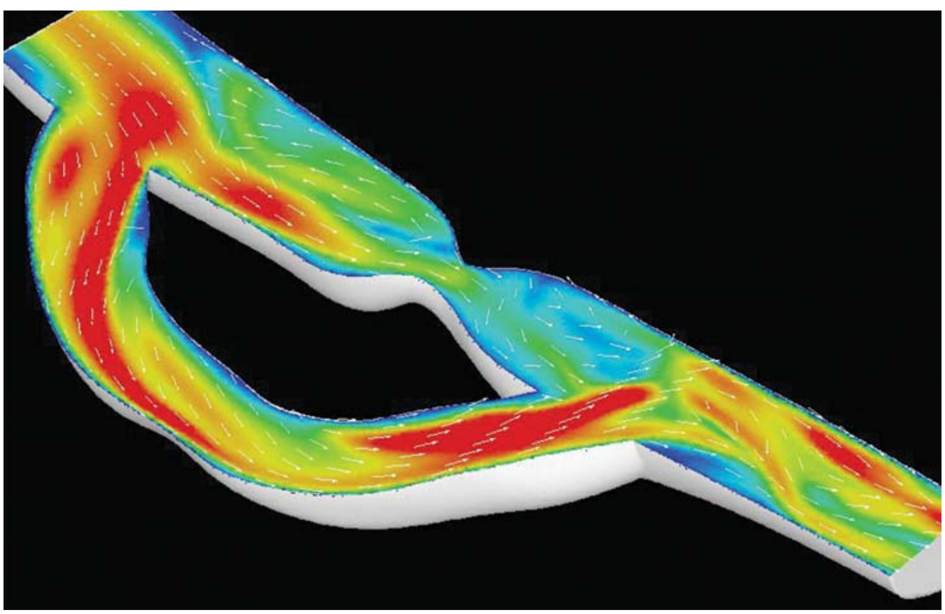

Figure 1: *Sectioned view of a pig's thoracic aorta bypass showing the internal fluid dynamics. Red depicts high blood flow velocity. Blue shades represent low pressure areas along the artery wall where plaque is most likely to form, leading to blockages.*
(Courtesy Charles Taylor, Stanford University)

brain, causing a stroke. Similar observations of blood flow patterns in the coronary arteries indicate a connection to heart attack: clots break off and travel downstream, closing off major vessels, leading to a heart attack.

While every patient's condition is unique, several generalizations can be made about blood flow patterns and how they are related to the formation of disease: Areas in the circulatory system with recirculation, low wall shear stress (low surface forces), and irregular blood flow are areas where plaque is most likely to form on artery walls.

Blood flow velocity is boosted significantly with an increase in physical activity. When sitting, there is a five-fold reduction in the demand for blood flow to the lower extremities. This interruption in blood flow causes irregular flow patterns and slows the movement of blood in some locations along the artery walls. Under this sedentary state, the chance of fatty deposits attaching to artery walls increases significantly. Once formed, these deposits have the potential to break away into the blood stream, causing blockages, which can lead to stroke or heart attack. Conversely, when a person maintains a physically active lifestyle, the blood flow velocity in all areas of the circulatory system increases, boosting shear stress on the artery walls, reducing the opportunity for lipids to attach themselves to the artery walls.

Converting Data to Computational Models

In the late 1960s, computers evolved to a point where the aerospace industry could use computational methods to model fluid flow over aircraft configurations to improve their design. The medical community reached this same conclusion in 1976, creating one of the first computational simulations modeling Newtonian behavior of blood. Computational methods enable Taylor and his group to visualize blood flow very effectively and efficiently.

Before the use of computational methods, rigid, hand-blown glass heart models were used to collect data about fluid flow. Making any changes to the configuration of the glass model was a time-consuming, costly, and laborious process. "With computer models, conditions can be altered very quickly, and there is an opportunity to do things that are both subject- and patient-specific," explains Taylor. "NASA/CTR has played an important role in the development of numerical algorithms for this project." Previous computational work was conducted on a NAS Division-supported IBM SP2, in 1995. Now, Stanford University has its own supercomputing facility that houses a 128-processsor SGI Origin 3800 supercomputer, dedicated to bio-computations.

To build a computational model of blood flow through the circulatory system, the group begins with MRI (magnetic resonance imaging) data gleaned from select cardiovascular disease patients at Stanford Hospital. From the MRI data, the team uses Stanford-developed software, known as Geodesic, to identify and "trace out" patterns that the blood

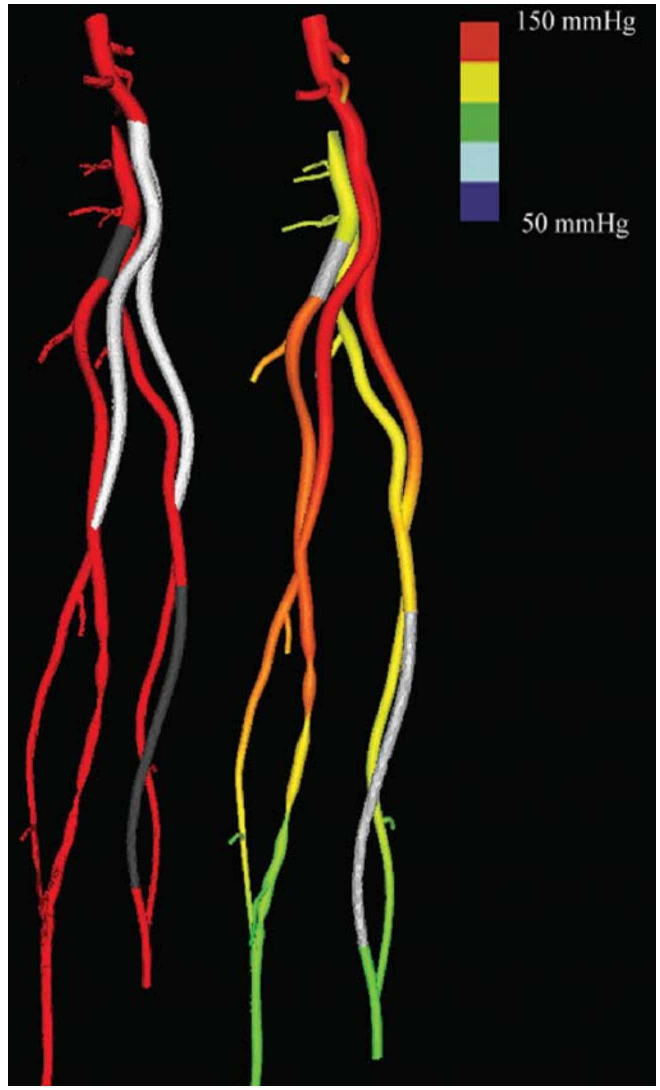

Figure 2: *Looking at the options presented by the ASPIRE System, doctors can see which treatments will produce acceptable pressure drops in patients, guiding their decision for treatment type. The pressure map offers a quick way to evaluate pressure losses associated with each treatment. Here, two treatment options are provided; the pressure distribution (generated using finite element methods) is captured during peak systole for a pre-operative case of ilio-occlusive disease, or complete closure of the air passage. The gray bands represent a complete blockage of the artery due to plaque formation. Graft implant: end-to-end aorto-iliac bypass.*

(Courtesy Charles Taylor, Stanford University)

vessels follow. Once the boundaries of the blood vessels are defined, a series of 2-D slices are taken at various locations along the vessel, and using the Level-Set method, the team is able to reconstruct an anatomic model of the blood vessel (see *Preoperative Model Creation Process*, page 12). "When we first started constructing models, it would take months to go from MRI data to a prototype. Now, we can acquire data and have a preoperative model constructed in about an hour's time," says Taylor.

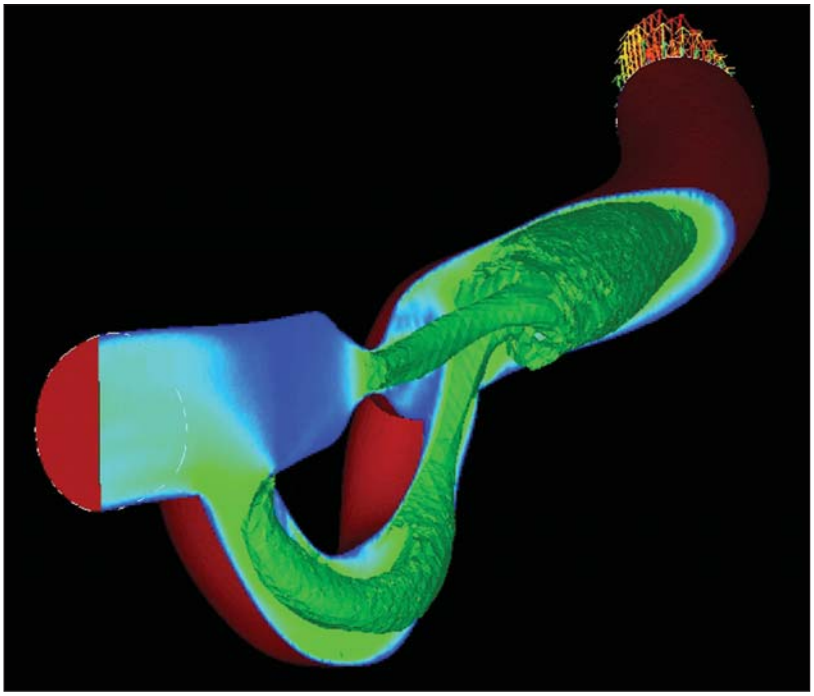

Figure 3: *Computational simulation of a pig aorta with an implanted bypass graft, during the acceleration phase (systole) of the cardiac cycle. The curved model represents a 2-D slice of the velocity magnitude, plus an iso-surface for a given value of the velocity magnitude. Red defines the outside of the artery wall, while green denotes high blood flow velocity. (Visualization by Valerie Favier)*

Once a simulation model is created, a doctor can implant a bypass graft in various locations along the blood vessel to see how it will affect blood flow (see Figure 3). Taylor and his group have worked closely with vascular surgeons on this process to ensure that the capabilities of the models will be something that will meet their needs.

Gathering Results

One very important phase in the engineering design cycle is the creation of a prototype. At this stage, researchers have the opportunity to make changes to their model to avoid design failure in the real world. In the medical field, however, doctors and scientists do not have the same freedom to test their theories on patients. In order to validate their simulation models, the group must have experimental results for comparison. Joy Ku, a doctorial candidate in Electrical Engineering at Stanford University and one of Dr. Taylor's students, is gathering experimental results from 15 different swine. “We decided to use pigs as our experimental animal, because the aorta of a pig is similar in size to that of a human. In addition, the blood clotting (thrombus) process in pigs is similar to humans,” explains Favier. The team implants grafts in each subject and takes velocity, flow rate, and pressure measurements before and after each procedure to gather data for the computational models (see Figure 1, page 10).

Selecting the Right Treatment

How well a surgical procedure works depends on blood flow conditions as well as a number of other factors. “Medicine is really a build and test approach in the sense that we acquire diagnostic data that tells us what the problem is, and then surgeons have to rely on pencil and paper sketches thinking, ‘Okay, this is what I’m going to do – this is what worked for me

Continued on page 14

Preoperative Model Creation Process

To create anatomic models of the vessels in cardiovascular patients, Charles A. Taylor and his team must work carefully through several steps. The process for preparing a simulation takes only two days – an acceptable timeframe for preoperative planning of bypass graft procedures. Beginning with a surgical procedure, the team gathers experimental results with which to compare their computational results. To construct the models, the team uses Magnetic Resonance Image (MRI) data gathered at Stanford. From the MRI data, the paths of vessels are traced out and their boundaries defined, in order to build the geometric models. To complete the model creation process, the team develops numerical flow simulations, which include both velocity and pressure components.

While the anatomic data for a patient is very important to understand when selecting the proper surgical treatment, it only tells a part of the story – physiological data is also important. “If you're going to do blood flow simulations, you have to know how much flow is coming in, and what all of the conditions are,” explains Taylor. The team also uses Magnetic Resonance Imaging to acquire blood flow velocity data in given planes. This image data is used to prescribe velocity conditions for the computational models, making them both anatomic and physiologic.

Step I. Surgery

A vascular bypass graft made with a polyester fabric is implanted during surgery. The top and end of the graft are connected to the thoracic aorta, creating two anastomoses. Then, the thoracic aorta is pinched with polyester umbilical tape midway between the two anatomoses. The last step simulates the presence of a stenosis with an occlusion of approximately 80 percent.

Step II. Magnetic Resonance Imaging

a. Measurements are taken using Magnetic Resonance Angiography (MRA) – about 100 slices, each with a thickness of 2.6 millimeters, are used to construct the 3-D anatomy. Even the smallest vessels are visible using this technique.

b. Cine-phase contrast Magnetic Resonance Imaging (MRI) techniques are used to measure the three components of blood velocity at select locations. Twenty-four frames in the cardiac cycle are obtained, and the 2-D images are segmented using the Level-Set method. The flow rate over the cardiac cycle is extracted from the velocity across the area of the vessel.

Step IIIa. Creation of Geometric Models

A path is created through points along the vessels of interest. Normal to this path, 2-D slices are used and segmented using the Level-Set method to extract the closed boundaries of the lumen. A surface is lofted through these curves, and a solid model is constructed using the Parasolid geometry kernel (Unigraphics Solutions, St. Louis, Missouri). The model is then meshed using an automatic mesh generator, developed by MeshSim, Simmetrix, Inc., New York. Typical mesh sizes are one million element meshes.

Step IIIb. Processing Flow Data for Inlet Boundary Conditions

For numerical simulation, a velocity profile is imposed as a boundary condition at the inlet. The physiologic pulsatile velocity is extracted during this step.

Step IV. Numerical Flow Simulations

A finite element code, PHASTA, developed by Kenneth Jansen, an associate professor at Rensselaer Polytechnic Institute (RPI), Troy, New York, and former NASA CTR fellow, is used to solve the incompressible Navier-Stokes equations. The velocity and pressure are computed in this realistic anatomic model, with a physiologic velocity profile. For the purpose of these simulations, the walls are rigid, and blood is behaving as a Newtonian fluid.

Note: In addition to working closely with Stanford's Departments of Surgery, and Radiology, the team is also collaborating with Rensselaer Polytechnic Institute on the numerical aspects of this research.

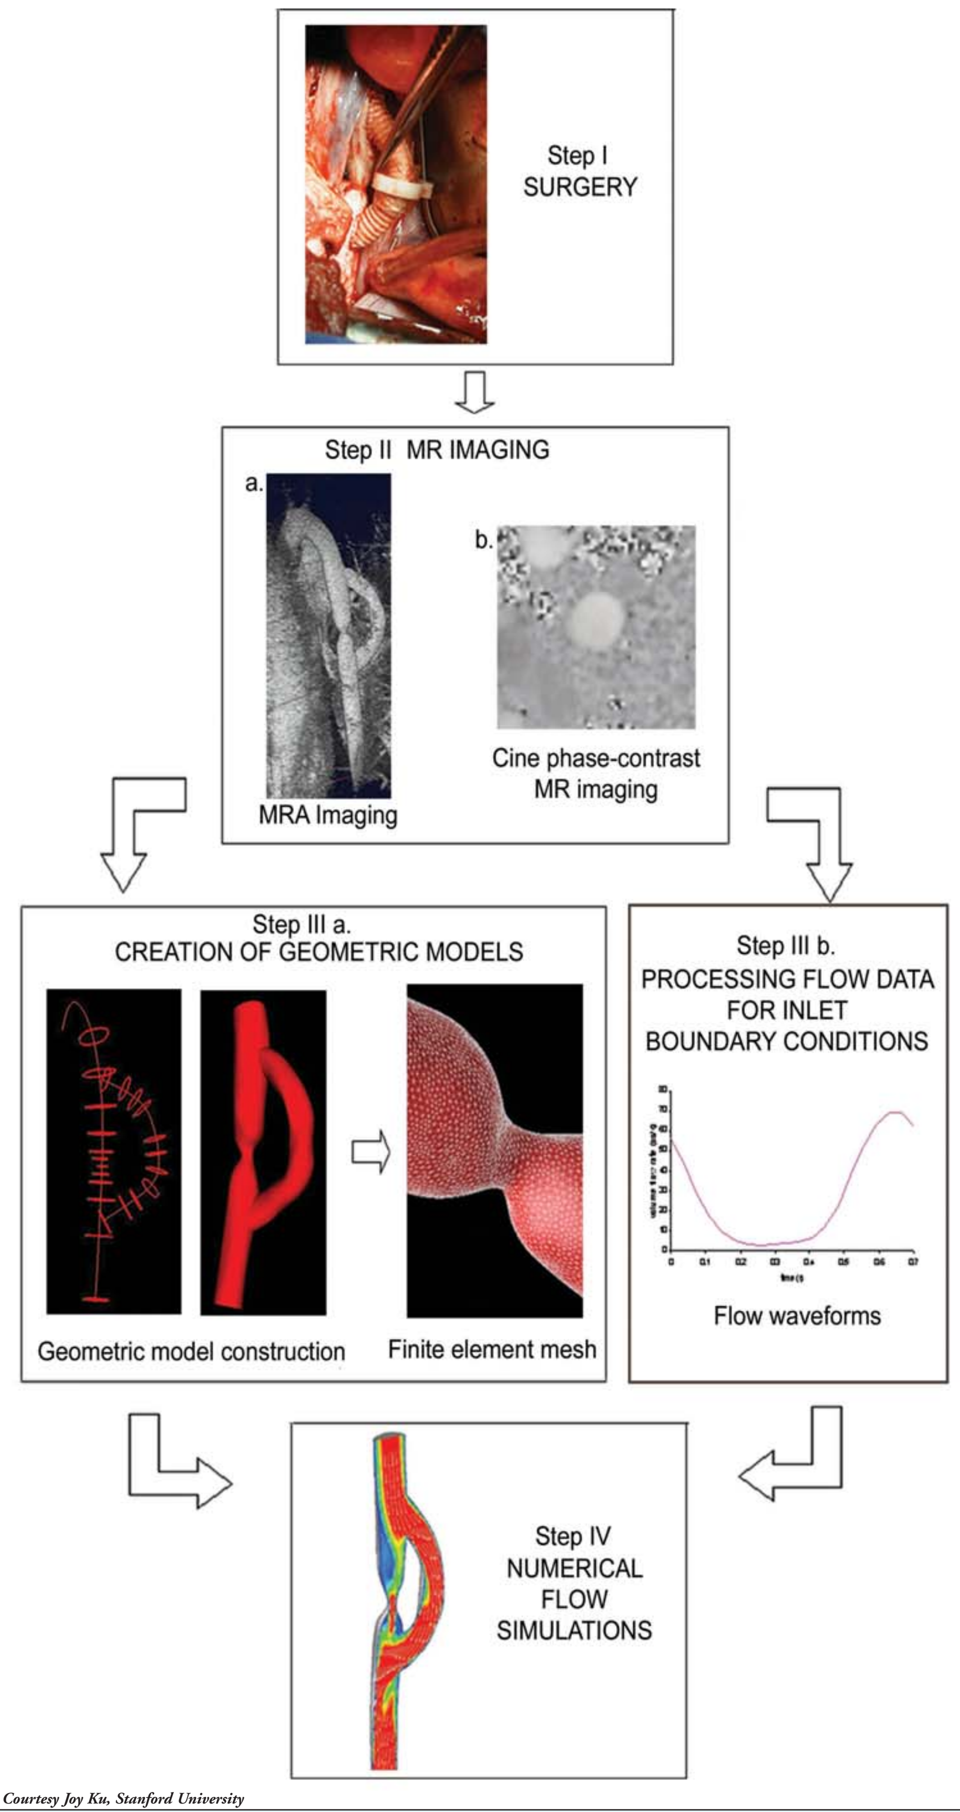

Courtesy Joy Ku, Stanford University

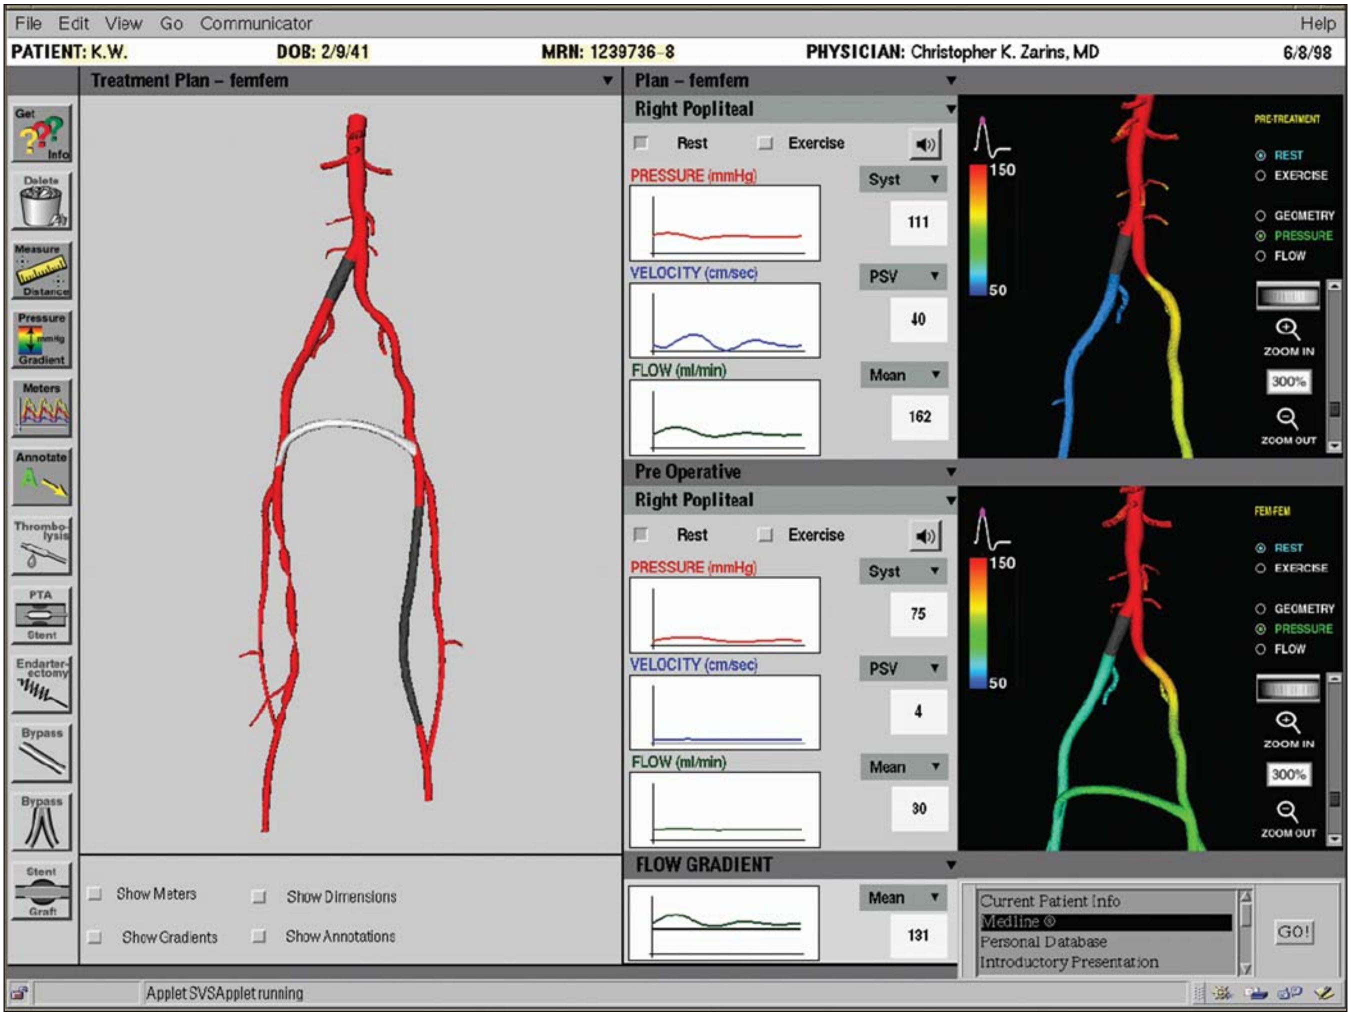

Figure 4: *The ASPIRE System's graphical user interface makes the tool easy for doctors to use, enabling quick feedback and preoperative planning. The visualization captured here represents one bypass possibility – a femoral to femoral artery bypass.*
(Courtesy Charles Taylor, Stanford University)

Continued from page 12
before in similar patients.' There is a certain amount of guesswork involved," says Taylor.

Although there are still many uncertainties, some things are known with confidence – the effects of maximizing volume flow, for example. If a surgeon implants a bypass graft, he or she wants to make sure there is enough blood flowing through it, in order to avoid blood flow re-circulation, which leads to clot formation or bypass failure (see Figure 2, page 11). "The purpose of implementing a bypass graft is to improve blood flow – modeling will help determine before a procedure is done, if it will do just that." To determine if a treatment will work, Taylor and his team look at factors such as a drop in pressure, which helps identify any restrictions in the artery.

Currently, medicine is largely diagnostic, and empirically based. "Diagnosis tells you what's there, but what we're trying to do is develop tools that will help surgeons answer the 'what ifs,'" explains Taylor. The team is working on a tool called the ASPIRE (Advanced Surgical Planning Interactive Research Environment) System that they hope will help surgeons answer those "what if" questions.

The ASPIRE System is a web and Java-based interface, best described as a "surgical sketch pad" (see Figure 4). Taking a preoperative model, surgeons can sketch out their proposed procedure, and the program then pulls in results from a simulation. The team has recently begun to apply this tool, retrospectively, to actual patient data. "The ASPIRE System will enable physicians to do design, using diagnostic data to create operative plans to predict outcomes, and actually figure out what's going to happen before they do it," says Taylor. "The dream is to be able to simulate the whole body and to enable the surgeon to do whatever he or she wants to do with the model – add a graft or stent, for example," says Favier. "This project is really difficult, because it's also research-in-progress in the experimental field – when we try to compare the computational results with experimental results, the experimentalists want to compare their results with the simulation results."

—*Holly A. Amundson*

Marsoweb: Analyzing the Mysteries of Mars

The Marsoweb website offers researchers a one-stop shop for Mars images and data, including tools for evaluating the best landing sites for the next Mars mission.

When the ancient Babylonians first looked to the night sky and named the portentous blood-red dot Nergal, God of the Underworld, they never dreamed that mankind's obsession with this celestial body would inspire centuries of religious and scientific debate. Nor could they have dreamed that by the end of the 20th century, humankind would have sent more than 30 spacecraft to the planet we now call Mars, hoping to capture elusive clues to the history of our solar system and possibly to the origins of life itself.

Although several Mars-bound spacecraft failed to arrive at the planet, the United States' success with Mariner 9, Vikings 1 and 2, the Mars Global Surveyor, and the Mars Odyssey spacecraft has provided the space research community with unprecedented images and data. Scientists are now using these data to guide NASA's next Mars exploration project, the twin Mars Exploration Rovers (MER) planned for launch June 2003.

The MER 2003 mission's aim is to search for evidence of liquid water in the planet's past that will provide more information about the geologic and climatic history of Mars. The twin rovers will land in areas where conditions may once have been favorable to life. By careful study of orbital data, NASA scientists are currently selecting landing

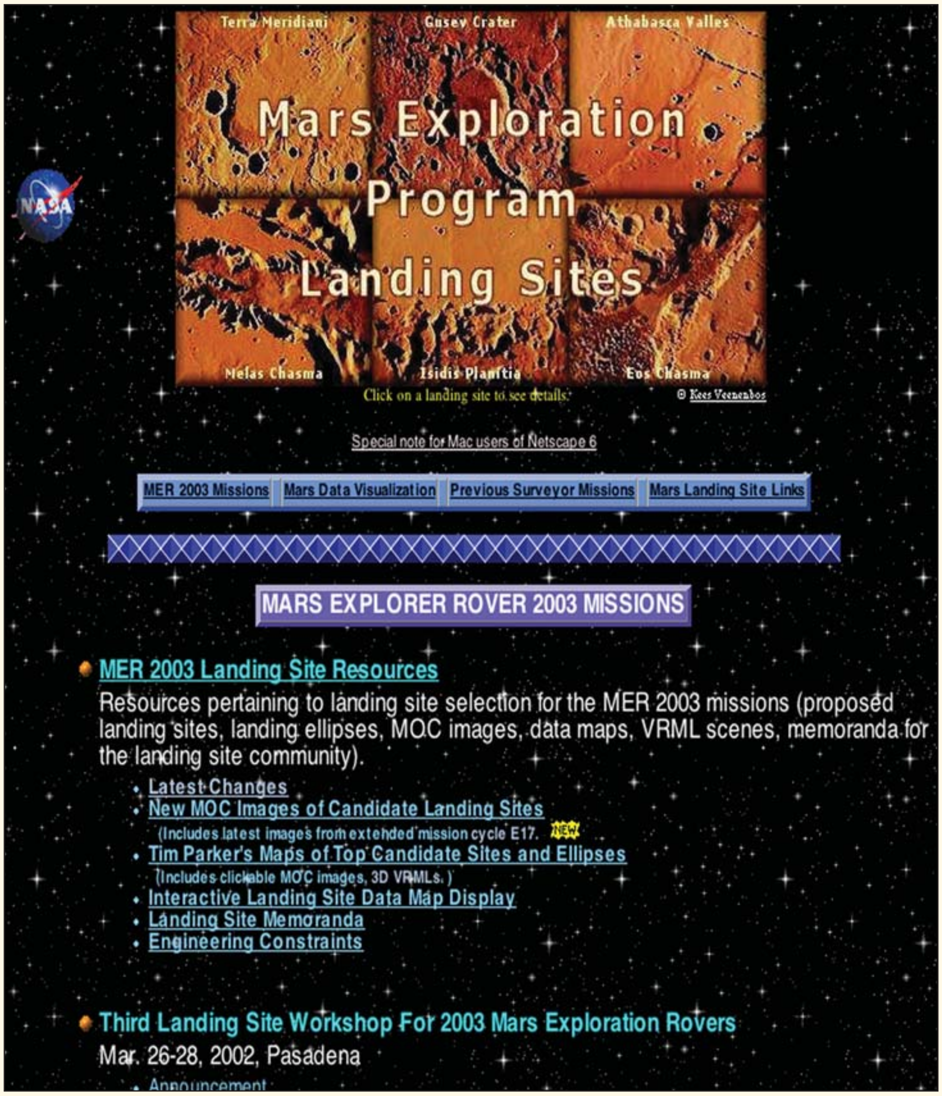

The Marsoweb site features resources related to landing site selection, as well as interactive global data archives. Visit the website at: http://marsoweb.nas.nasa.gov

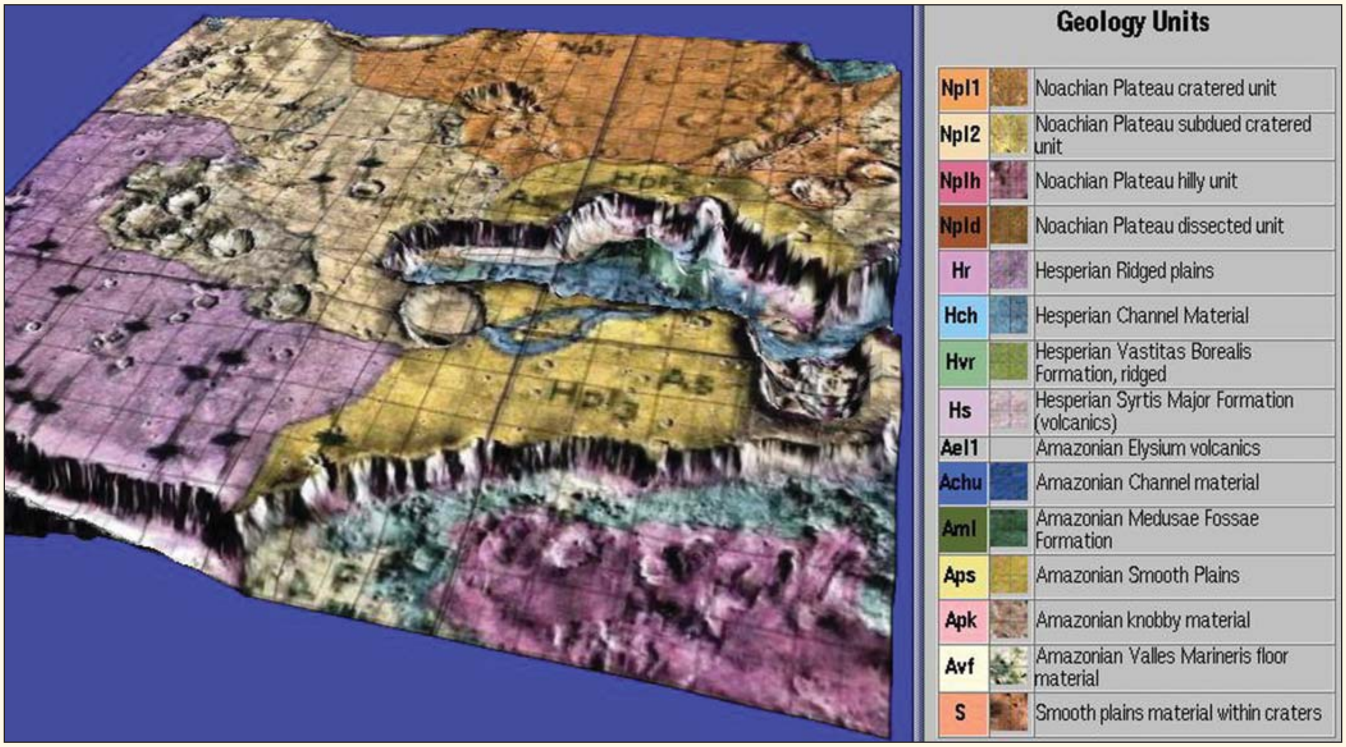

A 3-D Virtual Reality Modeling Language image of Viking terrain data is texture-mapped with a geology map in the Melas Chasma region. This tool is a component of the Mars Explorer Rover 2003 landing site selection resources. (Dataset courtesy Tim Parker and Matt Golombek, NASA/JPL, Visualization by Glenn Deardorff)

sites, specifically areas that contain minerals known to form under wet conditions or that appear to have been formed by water flow. To facilitate that selection process, Glenn Deardorff, a visualization technologist in the NASA Advanced Supercomputing (NAS) Division, and Virginia Gulick, a planetary scientist in the Space Science Division affiliated with the Center for Mars Exploration at Ames, have developed Marsoweb. Taken from the notion of "Mars-on-the-Web," Marsoweb is a NAS-hosted website dedicated to the online analysis of Mars orbiter data and a general repository for geological information about the Red Planet.

"The NAS Division was initially asked to provide an interface for landing site selection processes by NASA Ames' Center for Mars Exploration," explains Deardorff. "Now the site has more than 31,000 users, including scientists at NASA, the U.S. Geological Survey office, various government laboratories, academia, and the general public." Working closely with Gulick, who also serves on the MER landing site selection steering committee, Deardorff has spent the past three years developing the website into a collaborative facility for Mars data analysis.

"The main goal of our effort has been to provide online analysis and visualization tools so the science community can use images of the planet taken at different resolutions and other datasets," explains Gulick. "We wanted to be able to place Global Surveyor images onto Viking images, overlay other datasets, including various engineering, topographical, compositional and geological data, then run comparisons, and make the data available in a user-friendly format to support the workshops on landing site selection."

Using a unique interactive data map viewer also developed by Deardorff, scientists can view and query the candidate landing sites, (which are located in the equatorial area of the planet from 15 degrees north to 15 degrees south), or download the data maps into their own systems for further study. Deardorff's viewer is a Java-based graphical user interface, which incorporates various types of maps, most of which can be selected for further study with interactive querying tools. The background images are from Viking, overlaid by higher resolution images from the Mars Orbital Camera (MOC) on board the Mars Global Surveyor (MGS), some of which show detail at less than three meters per pixel. "In many cases, the coordinates were off a little," says Deardorff. "I had to hand-place most of the thousand or so initial MOC images, figuring out exactly how to line them up, skew, and orient them. I developed a real eye for geographic features and shading." Since then, more than 93,000 MOC images have been released, and are housed at Malin Space Systems in San Diego, Calif. All of the narrow-angle MOC images (more than 41,000 of them) can be accessed through Marsoweb.

Deardorff also developed a Virtual Reality Modeling Language (VRML) component in the map viewer, which provides an interactive 3-D image of Mars' surface on the user's screen (see image above). VRML images change as the user

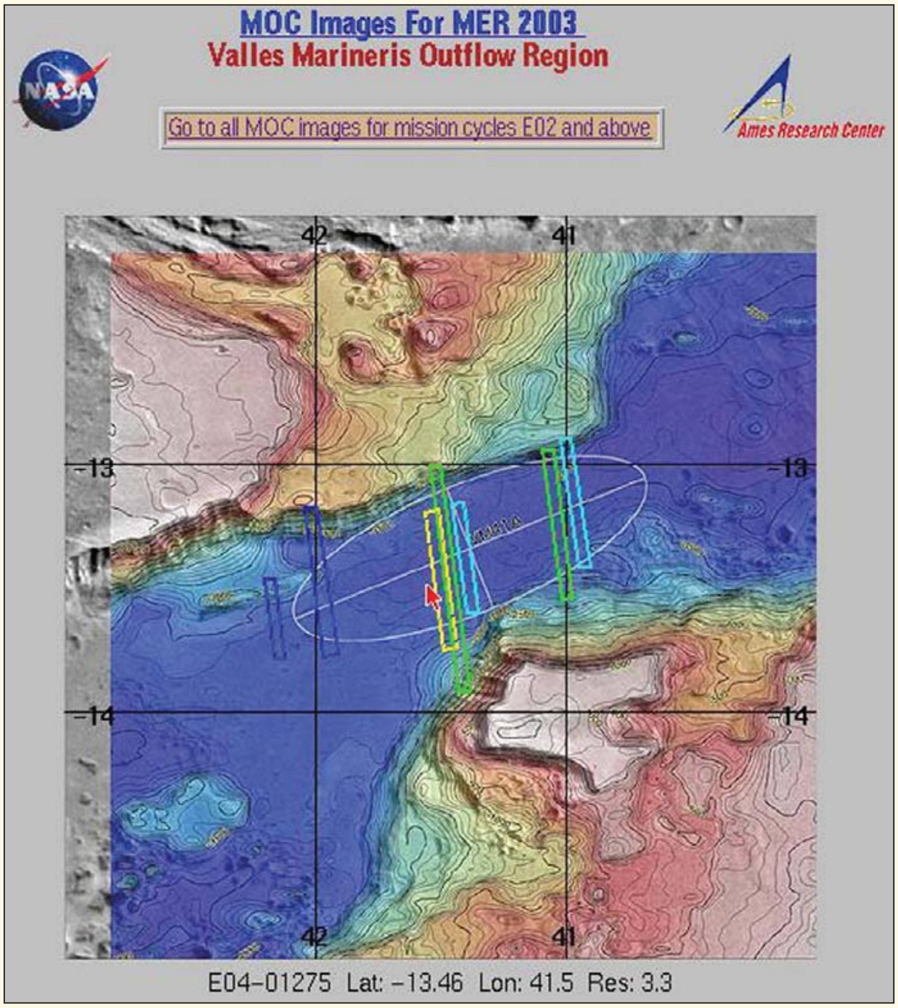

Mars Orbiter Camera (MOC) images for MER 2003 landing site studies can be displayed on regional navigation pages for closer scrutiny of the planet's surface. This user interface is used to select and view images in the MOC image archive. (Glenn Deardorff)

controls the mouse to simulate movement, enabling virtual 3-D exploration. Researchers can then zoom through Mars' canyons and valleys, or fly over its volcanoes and desert dune fields.

Deardorff says that support for this kind of interactive data exploration tool has been overwhelming. For the past three years, the project has been supported by NASA's Mars Data Analysis Program. Last year, Deardorff submitted a proposal for "A Web-based Collaboratory for Interactive Planetary Data Exploration" to NASA's Office of Space Sciences highly competitive Applied Information Systems Research Program. The proposal was awarded a one-year grant, which has enabled Deardorff to extend many of the website's interactive capabilities.

A relatively new feature allows users to interact with data from the Mars Orbiter Laser Altimeter (MOLA), an elevation measurement tool on board the Global Surveyor. Users are able to create and query elevations, cross-sections, and profiles, and can select a MOLA track to examine how the profile features correspond to the track's plan view. Other queriable features include data from the MGS Thermal Emission Spectrometer (TES), which shows measurements of thermal radiation, and digital maps of geology units.

Although the initial focus of Marsoweb has been to provide essential data accessing tools for landing site selection, Gulick acknowledges that the site continues to evolve as a "one-stop shop for Mars data and user-friendly analysis capabilities," offering an easy-to-access online library. Marsoweb currently hosts a huge repository of data, including almost 400 MOC images of candidate landing sites, global archives of over a thousand MOC images from the first orbital phases of the Global Surveyor, interactive archives of MOLA elevation tracks, and other instrument and geological data.

To keep the public up-to-date on the landing site selection process, Deardorff has posted memoranda, presentations, and notes from workshops. "There's a lot of public interest in the website," says Deardorff. "We've even used voluntary contributions from outside enthusiasts, who have spent a lot of time making sure that every pixel on the MOC images is enhanced as much as possible."

As the MER 2003 mission approaches, scientists continue to evaluate landing sites that balance safety with scientific merit. To aid in the search, Deardorff and Gulick are planning several enhancements to make Marsoweb even more compelling to the Mars community. The most pressing, Deardorff says, will be to implement an electronic notebook and a system for sharing screen viewing for remote collaboration among scientists, to integrate the MER 2003 landing site and global data archives, and to incorporate data from the latest orbiter, Mars Odyssey.

Astronomer Johannes Kepler worked out the intriguing motions of Mars in 1609, the year before Galileo first pointed his telescope toward the stars. Almost 400 years later, Marsoweb is bringing thousands of close-up views of the entire martian surface to desktops around the world.

— *Julie Jervis*

CODE Framework Brings Structure to Grid Computing

IPG Development team applies standards and protocols established in the Grid Forum to a tool for monitoring and managing grid resources, even from remote locations.

NASA's Information Power Grid (IPG) team has actively participated in the Grid Forum since its inception in June 1999. Established to address issues associated with grid technologies, the Grid Forum aims to promote and develop standards and protocols to help guide the growth and implementation of emerging grid environments. One aspect of grid computing addressed by the Grid Forum is of particular interest to the IPG team: Grid Monitoring Architecture (GMA), which encompasses the identification of system parameters and data formats required for monitoring.

Members of the Grid Forum have recently begun to agree on GMA standards. "The current challenge is dealing with the changing type of grid infrastructure – now everything is XML-based, but the grid environment will be shifting to web services, or grid services in the near future," explains Warren Smith, IPG development group lead. Smith has applied some of these new standards and protocols to a grid framework called Control and Observation in Distributed Environments, or CODE. The team's motivation for constructing CODE was very practical: The IPG environment needs a system for observing and controlling the resources, services, and applications that make up the large production grid. Adding resources and services to the grid environment makes it increasingly difficult to keep track of everything, Smith explains.

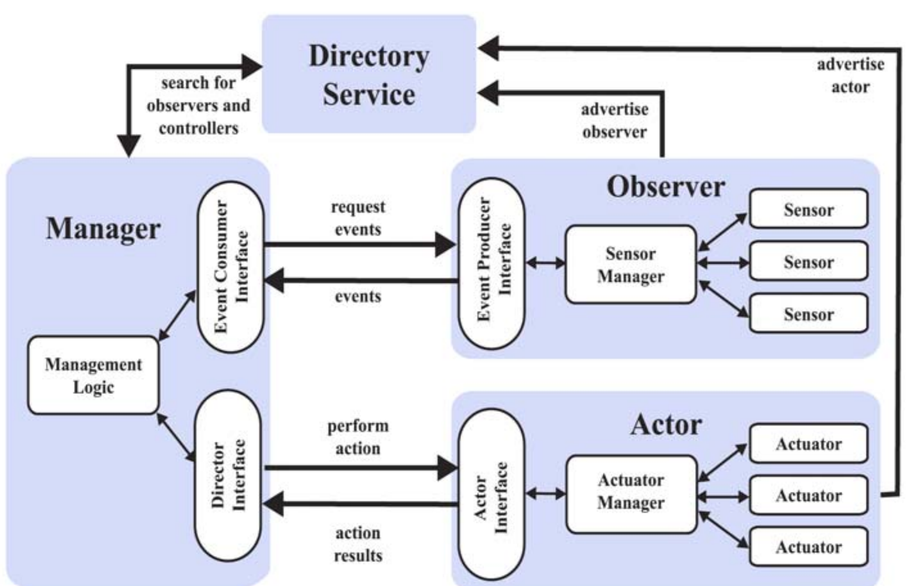

Figure 1: *The CODE framework is composed of four components that facilitate the control and observation of resources, services, and applications supporting large computational grids. The "actor" performs actions that affect local or remote resources, services, and applications. The "observer" provides information, such as scheduling queue information and local network usage, in the form of events. The "manager" asks the observers for information, examines the information and makes decisions about it, and requests actions from actors. The "directory service" is a searchable database used to store information about locations of observers and actors, which enables managers to search for specific observers and actors.*

"Our tool provides a secure framework for making observations and performing actions on remote computer systems, directing this data where it is needed, and analyzing data to determine what actions should be taken," explains Smith. The tool is composed of four main components: The "observers" of the framework perform and report observations, the "actors" perform actions, "managers" receive observations, and the "directory service" maintains a database for the location of observers and actors (see Figure 1).

CODE provides a set of commonly used sensors for observing a variety of properties – monitoring hosts and networks, and actuators for performing various actions, such as sending e-mail and restricting demons. Users can easily integrate their own components into the framework.

Structure of the Framework

At the lowest level in the CODE framework, the distributed event service provides the user with notifications when certain things happen, for example, when an application is complete or when a machine goes down. All of these notifications are transmitted through Internet connections, using a Grid Forum protocol based on XML. "Right now, these notifications are being done over a socket connection – it can be done using TCP with no security, or it can be done using the Grid Security Infrastructure," says Smith.

Integrating the CODE framework into the IPG environment has several notable benefits. "The new framework will provide notifications – a very important paradigm for distributed computing," explains Smith. CODE is also modular, allowing users to easily build a variety of "sensors" to monitor anything within the grid environment – an application or a piece of hardware, for example. In addition to its monitoring capabilities, the tool allows users to manage distributed systems or computations, and make decisions about their operations.

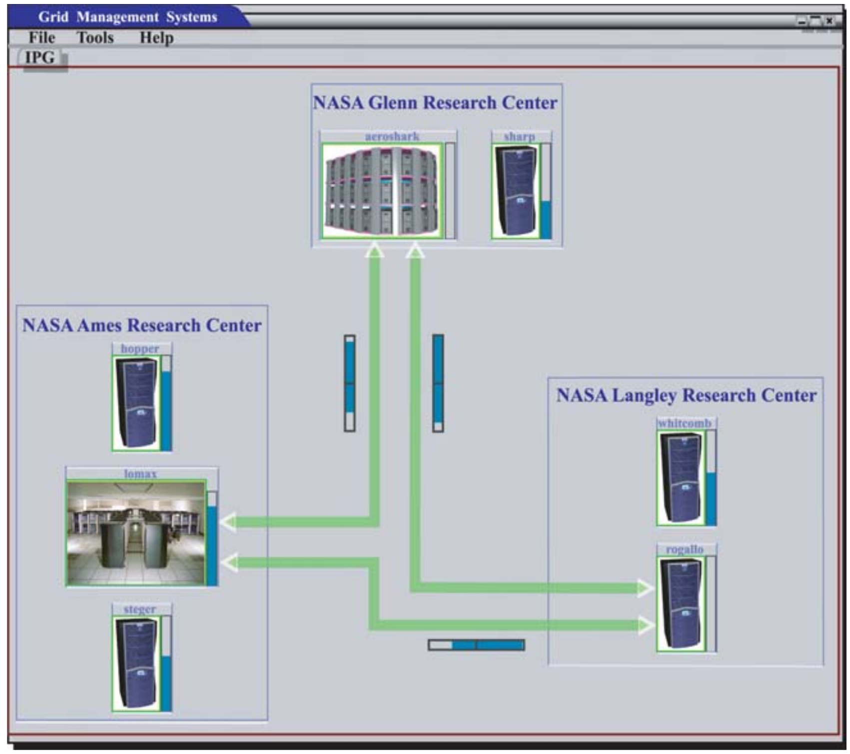

Figure 2: *This Java-developed graphical user interface allows grid administrators to view system status, including current load of machines at three NASA centers (Ames, Langley, and Glenn Research Centers). The color-coded system makes it easy to determine what is happening. For example, a green line represents a successful query of the system in the direction of the arrow, yellow means a query in that direction failed, and red denotes two consecutive query failures. The status bars next to the lines (light violet rectangles) represent current bandwidth usage for each resource.*

Although Smith would like to see other members of the Grid Forum community adopt CODE into their distributed computing environments, the framework was carefully designed to utilize the standards and protocols set forth in the Grid Forum, making it easy to communicate with other grids following the standards. "Even if people don't use the CODE framework, they can use their own tool and interface it to CODE. It will interoperate with other grid environments – that's the whole point of defining protocol standards – it's like FTP, it enables a user to talk to someone else's server," explains Smith.

Working Within the Framework

Using CODE, Smith has also developed a tool called the Grid Management System, designed for monitoring and managing a computational grid. The management tool takes advantage of the secure, scalable, and extensible architecture of CODE. "The idea behind the Grid Management System is to allow a grid administrator to keep an eye on a grid and manage it – this tool will be tailored to the needs of system administrators," says Smith.

The Grid Management System includes a clear graphical user interface, providing system administrators with an up-to-the-minute status of major grid resources, and their connecting networks (see Figure 2), all at a quick glance. Administrators can also perform administrative functions such as adding, removing, or modifying users in the grid mapfile, or database, on remote computer systems.

Several other systems for remote monitoring and resource management are available, however, many of them are costly, and lack standards and compatibility with other products. The Grid Management System, built within the CODE framework, adheres to Grid Forum standards and is portable, while providing the necessary functionality to use and manage grid resources. In addition, the Grid Management System supports the grid security infrastructure.

Smith and other members of the IPG development team will continue to integrate the latest Grid Forum standards and protocols into CODE. In addition, they will improve and expand the functionality of the tool based on user feedback. Says Smith: "We hope that the CODE framework will be one of the basic components used on the grid – it looks like it will be deployed into the IPG within the next year."

To learn more about the CODE framework, visit : http://www.nas.nasa.gov/Research/Reports/Techreports/2001/nas-01-006-abstract.html

NAS Technical Training Seminars

During the first half of 2002, several well-attended new technology seminars were presented at the NASA Advanced Supercomputing (NAS) Facility. Many of these seminars have been videotaped, and can be borrowed by e-mailing the NAS Documentation Center (doc-center@nas.nasa.gov). The seminars highlighted below are all available on videotape. Please include the date when ordering.

Computational Electromagnetics

Pavel Bochev from the Computational Mathematics and Algorithms Group at Sandia National Laboratories talked about computational electromagnetics and differential complexes. Bochev noted that nearly three decades after the first finite difference scheme for the time-domain Maxwell's equations became a standard in computational electromagnetics, the field abounds with discrete models of Maxwell's equations based on finite-difference, finite-volume, and finite element paradigms. Bochev showed that successful schemes all generate a discrete version of the De Rham complex, which guarantees a physically correct approximation of the curl-curl operator, and gave side-by-side comparisons of two time domain simulations of magnetic diffusion. The talk concluded with the examination of some examples of differential complexes related to the Laplace and Stokes equations. (May 20)

Magnetic Fields

Tatyana Sharpee from the University of California San Francisco's Sloan-Swartz Center for Theoretical Neurobiology in the Department of Physiology discussed a semi-classical solution to tunneling decay in a magnetic field. Sharpee's talk, "Many-electron Tunneling in a Magnetic Field," presented methods that have been applied to tunneling from a strongly correlated 2-D electron system in the presence of magnetic fields parallel to the electron layer, noting that the tunneling matrix element can be used as a sensitive probe of in-plane electron correlations. (April 18)

Numerical Modeling of Coating Flow

Len Schwartz from the Department of Mechanical Engineering and Mathematical Sciences at the University of Delaware discussed theoretical and numerical techniques for predicting the slow free-surface flow of liquids on solid substrates with a variety of driving mechanisms. He presented a simulation of the three-dimensional unsteady motion of Newtonian liquid layers and isolated droplets. Schwartz concluded with new results relevant to biological cell division. (April 3)

Development Tools

Prakash Narayan, Ferhat Hatay, and Chansup Byun from Sun Microsystems detailed three of Sun's new development tools: Forte Tools provides the user tools to create scalable, robust, high-performance computing (HPC) applications; Sun's HPC ClusterTools Software enables high-performance computing users to create, tune, and deploy scalable parallel computing applications in an open and stable environment; and HiMAP provides basic recommendations for compilation, debugging, and execution of message passing interface applications on a Sun HPC environment. (March 26)

Surgical Planning

Dr. Charles Taylor, Assistant Professor of Surgery and Mechanical Engineering at Stanford University, reviewed his application of computational fluid dynamics to cardiovascular surgery planning. Several aspects of simulation-based medical planning systems were described, including image-based modeling, operative planning, and three-dimensional and one-dimensional blood flow computational methods. (See *Computational Models to Predict Outcome of Vascular Surgery* on page 10 for additional information.) (March 21)

Dynamic Vision

UCLA Computer Science Professor Stefano Soatto presented an overview of "Vision as a Sensor for Control and Interaction with the Physical World." He noted that all of the projects fall under the umbrella of "dynamic vision," that is, the study of a vision as a sensor for computers to interact with physical space. Soatto discussed recent work on the estimation of three-dimensional shapes from collections of images in the estimation of three-dimensional motion and its application to real-time virtual insertion. In addition, Soatto also presented preliminary results in modeling dynamics of complex visual phenomena such as smoke, water, as well as modeling and recognizing human gaits. (March 19)

Breast Cancer Computational Studies

Professor Suhrit Dey from Eastern Illinois University returned to NAS to present his latest findings in breast cancer research. In this third in a series of seminars by Dey, the focus was on computational studies on the antigen-specific immune response to breast cancer. Dey described a new mathematical model to fight breast cancer. The model is three-dimensional and consists of five nonlinear reaction-diffusion type equations. They were solved computationally by perturbed functional iterations. The results obtained so far appear to be in agreement with those recorded in medical literature. (March 1)

Automated Systems Administration

Steve Traugott, founder of Infrastructures.Org and a former Ames employee, discussed the practices and terminology of automated systems administration such as concurrent versions systems; makefiles that run at boot; isconf; cfengine; lightweight directory access protocol; kickstart; jumpstart; and rsync. Traugott also reviewed many of the key components of the NAS infrastructure (which he helped create) including the division's CVS repository, the LAN engine toolset, and the Condor pool deployment and management tools. (Feb. 28)

Calendar of Events

Pressure contours behind the Shuttle Orbiter in high-supersonic flight computed using Cart3D — a grid generation tool suite collaboratively developed by NASA Ames researchers Michael Aftosmis and John Melton, and Marsha Berger, deputy director of the Courant Institute at New York University. (Computations and visualization by Michael Aftosmis)

Call for Papers: Graphics Applications for Grid Computing

Submissions due: August 31

The theme for a special March 2003 issue of IEEE Computer Graphics and Applications is graphics and visualization applications for grid computing. The ability to take advantage of the disparate resources available in grid-enabled applications is both exciting and difficult. The magazine is soliciting papers that describe innovative results in this area. For potential topics and details, visit: www.sci.utah.edu/grids.html

15th International Conference on Parallel and Distributed Computing Systems

Louisville, Kentucky • September 19–21

The 15th International Conference on Parallel and Distributed Computing Systems (PDCS) provides researchers, designers, educators, and other interested parties an opportunity to explore and exchange information about all aspects of parallel and distributed computing systems. Find out more at: www.isca-hq.org/confr.htm

The Grace Hopper Celebration of Women in Computing 2002

Vancouver, B.C., Canada • October 9–12

The Grace Hopper Celebration of Women in Computing 2002 is the fourth in a series of conferences designed to bring the research and career interests of women in computing to the forefront. Presenters are leaders in their respective fields, representing industrial, academic, and government communities. Leading researchers present their current work, while special sessions focus on the role of women in today's technology fields. For more information, please visit: www.gracehopper.org

10th Foresight Conference on Molecular Nanotechnology

Bethesda, Maryland • October 11–13

Rapid advances in our ability to image, manipulate, and probe the properties of matter at the atomic scale – together with emerging insights into structure, function and self-assembly in biological systems – is bringing to fruition the tremendous promise of nanotechnology. The 10th Foresight Conference on Molecular Nanotechnology covers a broad range of topics, including: nanodevices and nanostructures, biomolecular machinery, molecular machines, scanning probes, self-assembly, and computational chemistry. See: www.foresight.org/Conferences/MNT10/index.html

Next Generation Networks (NGN) Conference

Boston, Massachusetts • October 14–18

The NGN conference focuses exclusively on future trends in high-performance networking, looking at decisions that vendors, service-providers, and enterprise users will have to make two to five years down the road. Full-day and half-day tutorials will cover: advanced internet protocols, emerging security technologies, optical networking technologies, trends in broadband networking, and more. Go to: www.bcr.com/ngn

11th Conference on Current Trends in Computational Chemistry

Jackson, Mississippi • November 1–2

This symposium, organized by Jackson State University, covers all areas of computational chemistry as well as quantum chemistry. The format consists of a series of (invited) plenary lectures and poster presentations on Friday and Saturday, covering applications as well as theory. Get details at: http://ccmsi.jsums.edu/cctcc

Supercomputing 2002 (SC2002)

Baltimore, Maryland • November 16–22

"From Terabytes to Insights" is the theme for this year's supercomputing conference. SC2002 will bring together scientists, engineers, designers, and managers from all areas of high-performance networking and computing, showcasing the latest in systems, applications, and services. Check out the conference website at: www.sc2002.org

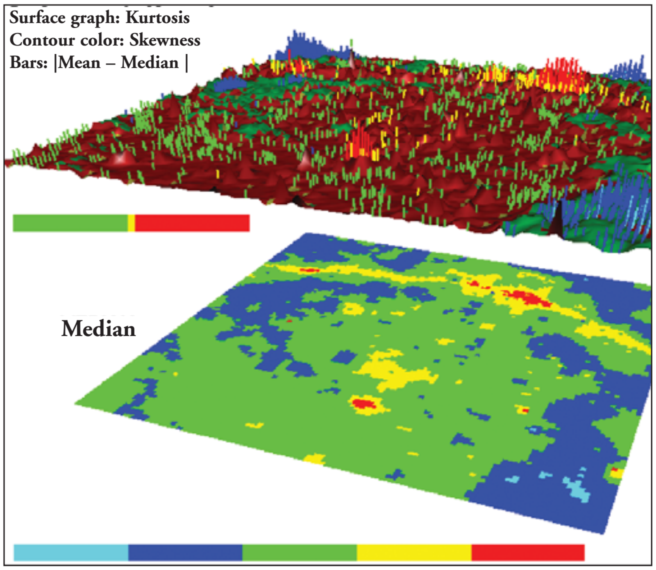

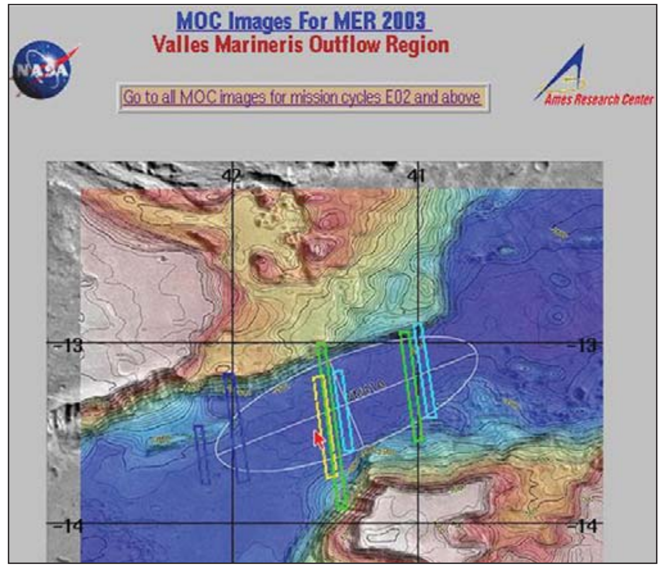

Visualizing Uncertainty in Earth Observing System Satellite Data

Sophisticated scientific visualization tools help Earth scientists interpret variations in satellite data. See page 6.

Marsoweb: Analyzing the Mysteries of Mars

The Marsoweb website offers researchers a one-stop shop for Mars images and data, including tools for evaluating the best landing sites for the next Mars mission. See page 15.

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 94035-1000

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