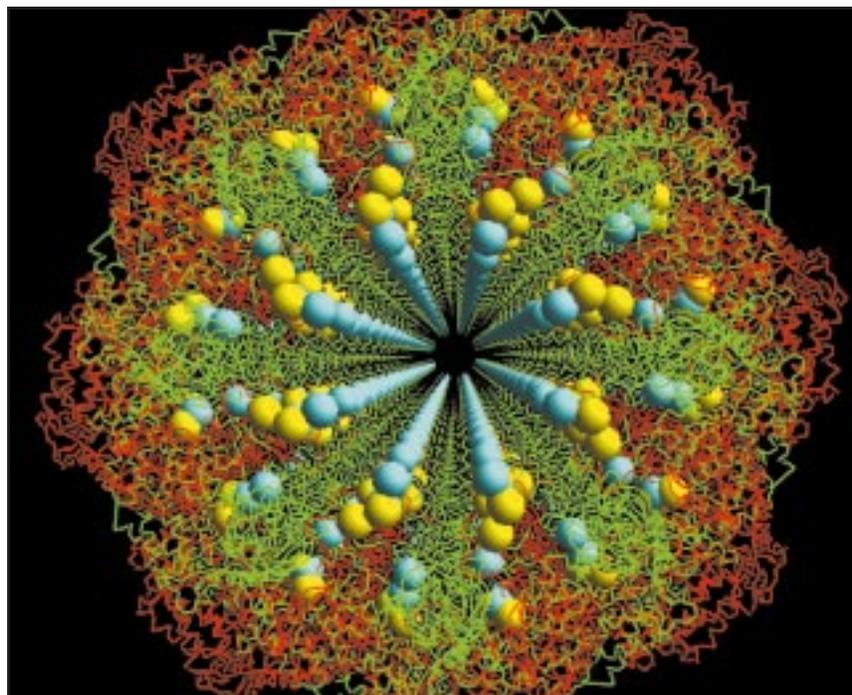


# Gridpoints

The Quarterly Publication of the Numerical  
Aerospace Simulation Systems Division



**Visual representation of protein nanotubes are the result of a collaborative effort between NAS researchers and the NASA Computational Astrobiology Institute. Page 11**

## FEATURES

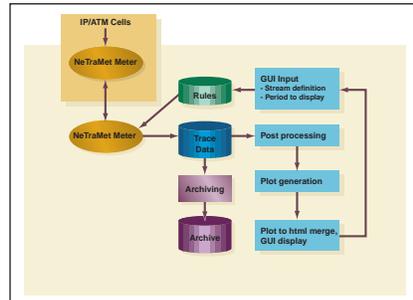
- **News from the NAS Systems Division — 1**
- **PC-based Network Traffic Monitoring — 4**
- **SC2000 Special Insert (follows Page 6)**
- **Calculating Semiconductor Etch Processes — 7**

# Features

4

## PC-based Network Traffic Monitoring Optimizes Existing Bandwidth

*NASA researchers are now developing and installing PCMon monitoring platforms at selected NREN network locations*  
Mark Foster



1A

## NASA at SC2000

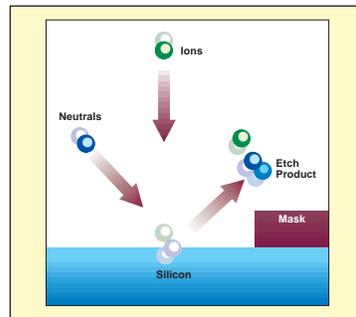
*Five NASA centers are showcasing recent advances in high-performance computing at SC2000 in Dallas. An eight-page special edition describing NASA's demonstrations follows page 6*



7

## Calculating Feature Evolution in Semiconductor Processing

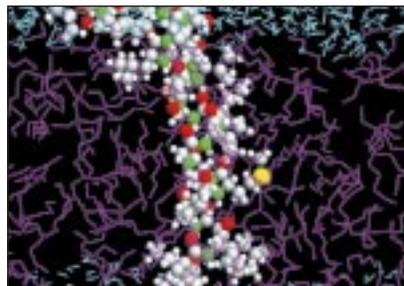
*Ames researchers have developed a computational code to understand the behavior of plasma on a silicon substrate used in integrated circuit manufacturing*  
Helen H. Hwang



11

## VMS Architecture Evolves As New Modules Are Developed

*The new tool, 'growler,' is capable of supporting multiple servers in an integrated collaborative environment*  
Holly A. Amundson



## On The Cover:

A protein nanotube is visualized using the Protein Data Bank (PDB) viewer module of 'growler.' Visualizations like this help stretch the limits of protein research for scientists at NASA. See page 11 for the full story. (Visualization by Chris Henze and Bryan Green)

# Gridpoints

The Quarterly Publication of the Numerical Aerospace Simulation Systems Division

Vol. 1, No. 4 • Fall 2000

**Executive Editors:** Bill Feiereisen  
John Ziebarth

**Editor:** Nicholas A. Veronico

**Writers:** Holly A. Amundson  
Mark A. Foster  
Helen H. Hwang

**Graphic Designer:** Sue Kim

**Illustrator:** Gail Felchle

**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: Dick Anderson, M. Jahed Djomehri, Jill Dunbar, Ian Foster, Mark A. Foster, Ken Freeman, Chris Gong, Bryan Green, T.R. Govindan, Chris Henze, Marjorie Johnson, Bill Johnston, Pat Kaspar, Chris Kleiber, Dochan Kwak, Meyya Meyappan, Sally Miller, Gina Morello, Todd Novak, Andrew Pohorille, Marcia Redmond, Frank Sabadin, Scott Schreck, Ryan Spaulding, Valerie Steele, Ron Sykora, Jim Taft, Leigh Ann Tanner, Bill Thigpen, Arsi Vaziri, Cliff Williams, Jerry Yan, Maurice Yarrow, John Ziebarth

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

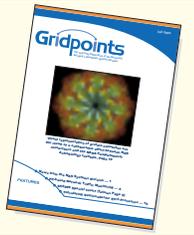
*Gridpoints* acknowledges that some words, model names and designations mentioned herein are the property of the trademark holder. *Gridpoints* uses them for identification purposes only.

### Subscribe to *Gridpoints*

Subscriptions to *Gridpoints* are complementary. Please send e-mail with your name, company name, title, address, and zip code to: [gridpoints@nas.nasa.gov](mailto:gridpoints@nas.nasa.gov) or subscribe online at: [www.nas.nasa.gov/gridpoints/subscribeform.html](http://www.nas.nasa.gov/gridpoints/subscribeform.html)

Visit us on the web at:  
<http://www.nas.nasa.gov/gridpoints>

## News From NAS



### IPG Team Meets Level One Milestone

**N**ASA's Information Power Grid (IPG) team demonstrated a complete version of a grid system, meeting another part of their level one milestone at the end of September. To demonstrate high-speed data access on the grid's infrastructure, the group employed the IPG Virtual Laboratory (ILab) program used for managing parameter studies. In addition, the program manager Condor provided access to idle workstations, thereby adding more computing power for executing jobs on the grid.

The IPG is designed to take a large collection of dispersed and heterogeneous resources – computing systems, storage systems, and instruments – and define a standard set of services for accessing those resources for scientific research. NASA Ames, Glenn, and Langley Research Centers are collaborating to develop infrastructure for the grid.

ILab is a problem-solving tool designed by NAS researchers Maurice Yarrow and Karen McCann to manage parameter studies in the IPG environment.

*Continued on Page 2*

### NAS Mission

To lead the country in the research and development of high-performance computing for NASA Programs and Missions by being the first to develop, implement, and integrate new high-performance computing technologies into useful production systems.

To provide NASA and its customers with the most powerful, reliable, and usable production high-performance computing systems in the country.

## From The Division Chief

**O**ver the past two years, the NAS Systems Division has instituted a very large collaborative research project to develop NASA's Information Power Grid (IPG). Researchers at NAS are partnering with experts from the National Science Foundation, the National Computational Science Alliance (NCSA), the Department of Energy, San Diego Supercomputer Center, Argonne National Laboratory, and the University of Indiana, to tie high-performance computing and data handling facilities across the country into one grid that crosses administrative boundaries. The IPG concept will be known by a variety of names under different organizations, but the idea is – in some respects – like a web for scientific and engineering computing. It will give researchers access to remotely located data and large computational facilities, and scientific instruments.

In September, the IPG team met a “level one” milestone by demonstrating distributed data access over the grid using the IPG Virtual Laboratory, or ILab software program and the Condor system manager. (See News From NAS). To



meet this milestone, our researchers and collaborators have worked to build the IPG's infrastructure to provide a uniform interface for launching, monitoring, and analyzing computational jobs. The enormous power of the grid will allow researchers to take their science to its next level, increasing our knowledge in areas as diverse as the bonding properties of atoms to remotely controlled vehicles on Mars.

The IPG will distribute computing power, and to this end NASA and SGI are working together to complete the next evolution of power generators in the form of a

1,024 processor machine. The system is expected to be up and running next year, and we will keep you informed on its progress in the Winter 2001 issue of *Gridpoints*.

Currently, the 512-processor SGI Origin 2800 housed at the NAS Facility can run at 60 gigaflops of sustained performance. The MLP (Multi-Level Parallelism) programming model developed at NAS is expected to linearly scale in calculation speed. This will enable new discoveries in computational fluid dynamics and computational astrophysics, as well as other disciplines.

Achievements in science by NAS researchers will be showcased at SC2000 November 4-10 in Dallas. This gathering of high-performance computing and network technologies experts is a forum for the exchange of ideas and techniques, and a tremendous learning arena for our researchers. At the conference, five NASA centers will offer a wide variety of demonstrations, videos, panel discussions, and paper presentations. If you are unable to attend the conference, an eight-page special SC2000 edition is included in this issue. I invite you to visit the NASA SC2000 website at: [www.nas.nasa.gov/SC2000](http://www.nas.nasa.gov/SC2000) for complete coverage of NASA's work at the conference.

As always, I welcome your feedback.

*Bill Feiereisen*

[wfeiereisen@mail.arc.nasa.gov](mailto:wfeiereisen@mail.arc.nasa.gov)

*Continued from Page 1*

Running the ILab program on the grid dramatically reduces turnaround time for complex computations. The system was recently used to manage a parameter sweep study on the X-38 Crew Return Vehicle across a collection of resources using the new IPG infrastructure. The ILab/X-38 test demonstrates uniform access to different IPG resources. The uniform interface for batch queuing systems, known as the Globus middleware toolkit, was recently integrated into

the Condor system manager. Condor utilizes Globus resources to provide more computing power by scavenging processor hours from 60 to 100 Sun and SGI workstations in the Condor pool. "I think the Condor project is part and parcel of getting the IPG infrastructure integrated into NAS," says IPG Project Manager Bill Johnston. NAS researcher Al Globus has been taking advantage of Condor to run a collection of genetic algorithms, exploring different molecular structures. Globus reports that in about a year,

## Information Power Grid Workshop Enhances Collaboration Efforts

**O**n September 19 and 20, NAS hosted the second Information Power Grid (IPG) Workshop at Ames Research Center, lead site for NASA's IPG project. The technical workshop marks two years of progress on the project, designed to provide uniform access to remotely located computational and storage resources, and instruments for collaborative problem solving. Addressing issues, reviewing progress, demonstrating accomplishments, and sharing ideas with colleagues on the future of the grid was part of the two-day event.

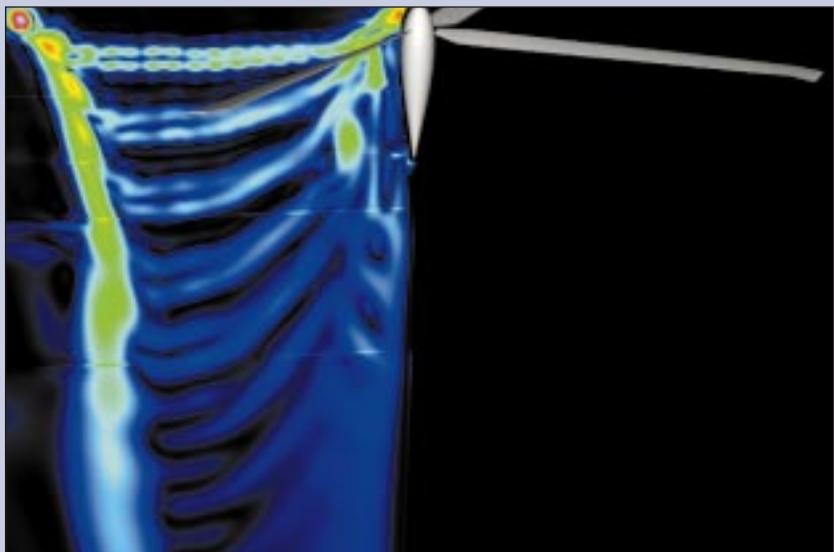
NAS Division Chief Bill Feereisen welcomed workshop participants, emphasizing the importance of collaboration to advance grid technology: "We are a partnership, we are trying to build a grid together. This is, by definition, a collaborative effort that crosses organizational boundaries. This partnership will build the infrastructure upon which science and engineering computing of the future will be based."

A number of organizations were represented at the workshop, and Argonne National Laboratory's Ian Foster was especially pleased with the meeting: "NASA is to be congratulated for putting together a meeting that not only showcased the progress made on IPG to date, but also for providing a wonderful forum for the exchange of ideas within the burgeoning grid community."

Workshop presentation subjects ranged from scheduling jobs on the grid, to debuggers for computational grid applications, to solving large problems using the grid. One of the featured speakers was M. Jahed Djomehri, a member of the algorithms, architectures, and applications group at NAS. Djomehri's presentation, "Large-Scale OVERFLOW Applications and Performance Assessment," outlined his attempts to find a technique that will

enable scientific computing of large-scale problems in a wide-area computational grid environment.

IPG Project Manager William Johnston is excited about the progress of grid technology for problem solving. "We are getting to the point where IPG is demonstrating that grids can indeed provide significant new computing and data handling capabilities to the NASA community for



This image is representative of a large-scale parallel computation executed on a grid system. Computed with 61 million grid points, the vorticity contours (blue) on a cutting plane located 45 degrees behind the rotor blade of an UH-60 helicopter model. (M. Jahed Djomehri, Roger Strawn, and Rupak Biswas)

dynamically building large-scale problem solving systems," Johnston says. "The breadth and depth of the presentations indicates that grids in general, and the IPG in particular, are rapidly becoming a reality. The innovative work of NAS and its National Science Foundation partners are advancing the state-of-the-art in grid computing, and the IPG engineering team is turning this work into real systems."

For abstracts of speakers' presentations and information about the workshop, visit:

<http://www.ipg.nasa.gov/workshop.html> 

he has accumulated nearly half a million CPU hours on this single problem using Condor.

"There are lots of pieces to the IPG and until you take a closer look, it seems like a trivial process to assemble the infrastructure. It takes a lot of work and effort. I would like to thank not only the IPG testbed group, but everyone in the NAS division as well as our outside collaborators for meeting this milestone," says Leigh Ann Tanner, deployment and integration project manager for the IPG. 

## NASA, SGI to Collaborate on New 1,024 Processor Origin 3000

**O**n August 23, NASA Ames Research Center and SGI announced plans to build a 1,024-processor Origin 3000 for the NAS Systems Division at Ames. Used by NASA researchers for aeronautics, Earth, and life science studies, it will be the largest single-system image (SSI) computer in the world. (An SSI system consists of a number of processors attached to globally shared memory and an input/output subsystem, all of which are controlled by a single operating system).

The Ames research community is anxious to run a number of important applications on the new system. "We want to run computational fluid dynamics codes, like OVERFLOW and NASA climate models on the 1,024 and see if they scale – the faster we can run code, the more rapidly we can solve problems vital to NASA missions," says Jim Taft, a senior scientific consultant at NAS.

Based on SGI's NUMAflex modular technology, the new 1,024-processor system, scheduled to be up and running early in 2001, will complete jobs substantially faster than NAS's previous 300-megahertz, 512-processor SGI Origin system, *Lomax*. "According to our projections, the new SGI Origin 3000 architecture is going to deliver about six times the performance at 1,024 processors as the 512-processor system," says NAS Division Chief Bill Feiereisen. SGI NUMAflex enables reconfiguration of the machine's architecture with the addition or removal of CPUs, depending on the desired scalability for an application.

The system utilizes the IRIX 6.5 operating system, and is built on SGI's "third-generation" NUMA architecture. It is designed to work with existing application software, and is fully compatible with other IRIX operating system-based servers and workstations. "Like all other large-CPU-count SGI Origin systems, the new 1,024 will continue to use the same IRIX operating system and compilers as other machines at NAS. This will be a distinct advantage to NASA researchers, who will be able to transition codes with a minimum of effort to the new system," says Taft. 

## Experts Meet to Solve Gigabit Networking Challenges

**N**etworking and applications development experts convened at NASA Ames Research Center August 14–16 to remove some of the roadblocks causing gigabit networking to screech to a halt just short of users' desktops. The workshop, "Gigabit Networking, the End-To-End View," was hosted by the NASA Research and Education Network (NREN) Project – an element of the agency's High Performance Computing and Communications (HPCC) Program, and was co-sponsored by the Next Generation Internet (NGI) Large Scale Networking Working Group.

HPCC Deputy Program Manager William Van Dalsem addressed the importance of gigabit networking to NASA enterprises. "The NASA HPCC Program goal," he said, "is to accelerate the development, application, and transfer of high-performance computing and computer communications technologies." Van Dalsem emphasized that these technologies are critical for meeting the engineering and science needs of the U.S. aerospace, Earth and space science, spaceborne research, and education communities. "The ultimate goal is to transfer these technologies to the American public," he said.

Applications demonstrated at the workshop outlined some of the difficulties in achieving true end-to-end gigabit networking. Demos included: the Digital Sky Virtual Observatory, Jet Propulsion Laboratory; Project Data Space, University of Illinois at Chicago; Visible Human, University of Michigan; Combustion Corridor – Visapult, Lawrence Livermore National Laboratory; Land Speed Record, Internet2 and Information Sciences Institute; High Definition TV, University of Washington; and Virtual Mechano-Synthesis (VMS), Ames Research Center. The VMS application, developed in the NAS division, was presented by researchers Bryan Green and Chris Henze, and NASA Center for Computational Astrobiology acting director Andrew Pohorille.

Panels of representatives from government agencies, industry, and academia set the stage for subsequent workshop breakout sessions by discussing the current status of gigabit networking, development of infrastructure, measurement, and gigabit end-to-end issues. The objective of the breakout sessions was to develop roadmaps for gigabit networking. Topics included: gigabit testbeds, platforms, measurement, and middleware/integration; teleseminars and telemeetings; models in real time, large databases, and remote instrumentation. Workshop proceedings and a final report are available on the NREN website:

[www.nren.nasa.gov/workshop5.html](http://www.nren.nasa.gov/workshop5.html) 

# PC-based Network Traffic Monitoring Optimizes Existing Bandwidth

**NASA researchers are now developing and installing PCMon monitoring platforms at selected NREN network locations**

**By Mark Foster**

**N**ASA's Research and Education Network (NREN) Project conducts research to enable the fusing of emerging network technologies into NASA mission applications. Many of these emerging technologies focus on building new network services. The NREN Project, which has shown significant success with creating a Multicast service for NASA networks, is also working on Quality of Service (QoS) technologies. QoS addresses the commitment of resources to specific applications, and ensures that the values of such performance parameters as bandwidth, latency, and packet loss stay within an acceptable range.

In June, NREN researchers demonstrated a new network traffic monitoring and measuring platform (PCMon) which has the ability to distinguish between data packets given "normal" treatment and data packets that are assigned "preferential" treatment. Transferring data with preferential treatment is used to provide a specific amount of bandwidth, or to

optimize data packet delay among the network's sending and receiving points. "Once we began working on QoS technologies, we realized that there were no commercial products that could verify that we were providing 'preferential treatment' to the end users," says NREN project manager Ken Freeman. The PCMon system was developed to meet this need, and has been installed at selected NREN connection points – Ames, Glenn, and Langley Research Centers, as well as the Next Generation Internet Exchange Point in Chicago (see Figure 1).

The ability to provide QoS to end-user applications will enable efficient sharing of network resources among multiple users, while providing preferential treatment to selected applications when network resources become scarce. "The digital video groups around NASA have long been frustrated

Figure 1 is an example deployment of PCMon monitoring platforms. An optical splitter has been located at a point where the target traffic can be monitored (Ames, Glenn, and Langley Research Centers, and the Next Generation Internet Exchange Point in Chicago). The splitter provides access to signals while also transmitting the signal on the primary path (in-band). Signals that have been split off (out-of-band) are sent to a measurement computer that runs a NeTraMet meter to collect and process Internet Protocol packets according to rules it has been provided. This meter process can send and receive data and configuration information via SNMP (Simple Network Management Protocol).

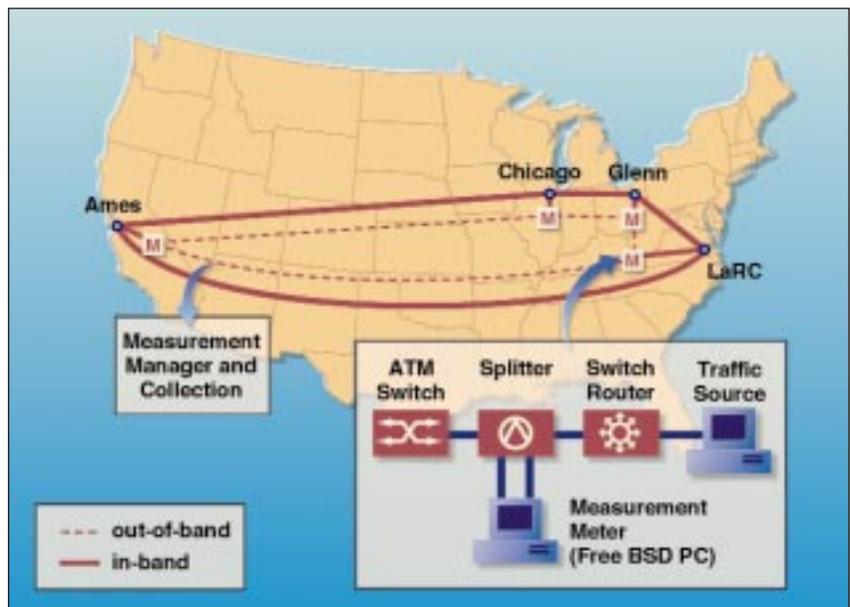


Figure 2: The NeTraMet manager process runs on a separate Unix platform and is used to control the actions of each meter process, downloading rules into the meters and retrieving information from the meters. In addition to incorporating the Realtime Traffic Flow Measurement features, This figure shows elements included in the NREN deployment.

by the data network's inability to consistently carry video traffic. With QoS, these video traffic flows can be given preferential treatment, thus providing performance guarantees during times of network congestion," says Marjory Johnson, NREN's associate project manager.

### PCMon Network Monitoring

The PCMon project's goal is to confirm the availability of network resources using monitoring devices that are independent of the applications being measured. Data transfer rates are measured by passively monitoring all traffic traveling on a network segment, then summarizing the amount of traffic for specific classes of data flows.

Data is collected using passive optical splitters combined with low-cost, commercial, off-the-shelf personal computer-based packet capture systems. This approach provides significant measurement flexibility across a wide range of applications and networking equipment, and can be done without modification to either the applications or the network core.

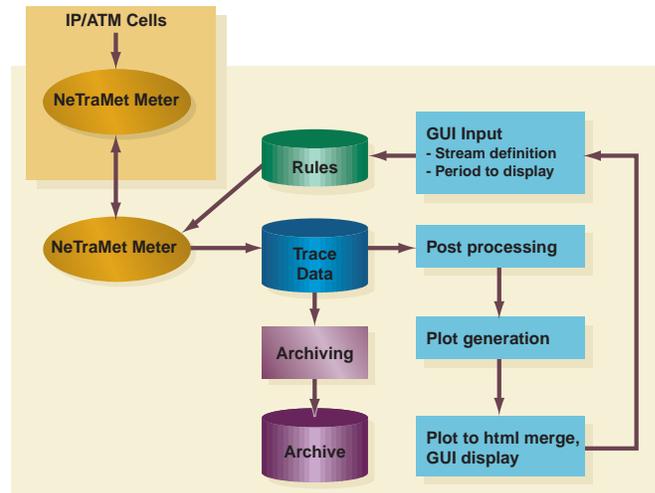
In addition to monitoring and summarizing traffic flows, the PCMon system includes a near-real-time graphical display that enables NREN researchers to correlate the behavior of a given application with the network's performance.

### PCMon Software

NREN's PCMon platforms incorporate NeTraMet software designed by Nevil Brownlee from the University of Auckland, New Zealand (see Figure 2). The package runs on a FreeBSD Unix operating system. The team chose NeTraMet for these key reasons:

- It supports a broad number of systems and interfaces which can be used collectively – Ethernet, ATM (Asynchronous Transfer Mode), and Packet-over-SONET (POS), with the collection and data representations transparent to these differences;
- NeTraMet is designed to be modular with a very lightweight collection "meter" that has limited CPU requirements while using powerful rule sets to filter and collect traffic of interest;
- It has the ability to incorporate Cisco Netflow values that can be useful in locations where installation of a passive probe is not feasible, and;
- NeTraMet provides ease of direct linkage to resource allocation services.

Because NeTraMet is SNMP (Simple Network Management Protocol)-based, the meters can easily provide meas-



urement values to any application that can make SNMP queries. To integrate NeTraMet, NREN's Jerry Young has been instrumental in adapting the software and building the rule set generator's graphical user interface – both fundamental to the creation of the PCMon platforms. Moreover, NREN's Celeste Banaag played an important role in the development of the data storage and archive elements of the PCMon system.

### Features of the PCMon System

NREN's PCMon system passively collects network traffic information and communicates that data with a high-speed (10 megabits per second or faster) out-of-band access mechanism. The system tracks at least four parameters: static maximum bandwidth of the link, preferred flow bandwidth reserved (maximum and minimum), best effort flow bandwidth used, and preferred flow bandwidth used. To track such parameters, the NREN team uses dynamically alterable criteria to classify and summarize traffic flows. In addition, summary data can be retrieved from multiple locations simultaneously.

Presenting an accurate display of the distributed measurements requires concurrent on-demand retrieval and near-real-time graphic display of the collected data. The measurements are then archived for possible post-processing or future analysis. Storage, processing, and data display enables users of the system to effectively control what flows are being measured and how they are viewed.

### Tools and Development

To get the platforms to interact with the network, the NREN team modified NeTraMet for ATM interfaces (Asynchronous Transfer Mode – manufactured by Efficient Networks) and defined NeTraMet rule sets for specifying what data the PCMon are to capture. In addition, NREN researchers created graphical representations for text-based input and measurement results (see Figure 3, page 6). Once the system counts the separate flows, the data is processed to produce GIF images. As with many network management and measurement systems, NREN researchers have made use of

Figure 3: A useful improvement to text-based rule set generation is a web form that can be used to define the variable details in the rule sets – for example, Differentiated Services (DS) field values, source and destination addresses. The output of this form is integrated into a new rule set that is subsequently transmitted to the NeTraMet meter.

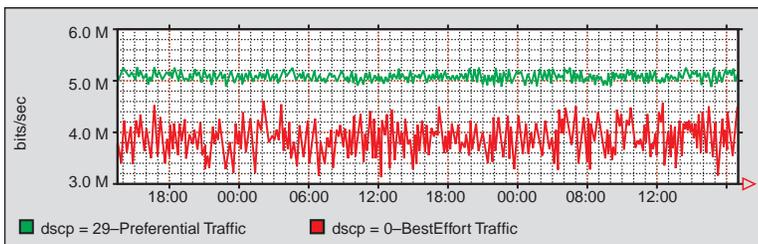
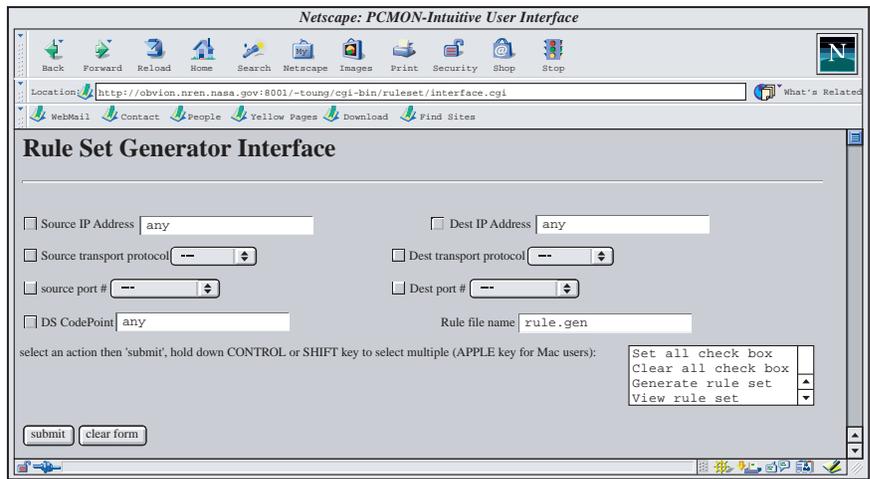


Figure 4: In this example, the traffic with a Differentiated Services codepoint of 29 used approximately 5 megabits per second (Mbps) while the other traffic (DSCP=0) about 4 Mbps. This data represents simulated traffic generation, thus the regularity of the patterns.

the RRDTool (Round Robin Database) system developed by Tobi Oetiker from the Swiss Federal Institute of Technology, Zurich; a prototype output sample is shown in Figure 4.

### Future PCMon Development

As demands upon the PCMon system increase, enhancements will be needed as its uses expand. These uses are likely to include higher speeds and Packet Over SONET interconnects. The University of Waikato, Hamilton, New Zealand, has been developing inexpensive PCI (Peripheral Component Interconnect) cards that are microprogrammable, and can support higher signaling speeds and alternative framing mechanisms. Advancements in PCI card development will enable the platforms to operate at gigabit speeds while adapting to future changes in the NREN network.

### Forthcoming Challenges

A large factor in the success of future NREN projects will be a network that provides specific bandwidth and performance capabilities to satisfy requirements for individual applications. Some of these applications will be able to request network services to achieve the required performance. While some network performance instrumentation can be incorporated into the services available to the applications, an independent, out-of-band mechanism is crucial to validate the effectiveness of the priority assigning mechanisms.

Measurement and characterization of Internet Protocol (IP) traffic has received considerable attention during the last few years; some projects, including this one, are now focusing on issues related to network QoS, such as delays in data

transfer. NREN researchers anticipate continued development in network QoS, particularly as more organizations deploy services that differentiate between types of IP traffic. NREN's major challenge for IP traffic measurement will be maintaining pace with the explosive growth in network transport speeds and interconnect technologies.

*Mark Foster is a chief network engineer with Recom Technologies at NASA Ames Research Center. He obtained an M.S. in computer science from Oregon Graduate Institute, Beaverton. He is one of the principle architects of the current NREN network, and leads the NREN applications and research group. He previously managed a systems and network support group in the University of Pennsylvania's Engineering School. Foster has participated in the architecture and deployment of NSFNet and vBNS network services to academic research organizations and reviews of IEEE Communications articles.*



### References to Related Projects

A number of projects focus on characterizing network traffic and network performance. NREN has incorporated some elements or ideas from a variety of projects. They include:  
 AMP – <http://amp.nlanr.net>  
 Surveyor – <http://www.advanced.org/surveyor>  
 WAND – <http://atm.cs.waikato.ac.nz/wand>  
 NeTraMet – <http://www.auckland.ac.nz/net/NeTraMet>  
 CoralReef /OCXMON  
<http://www.caida.org/tools/measurement/coralreef/>



NASA at SC 2000

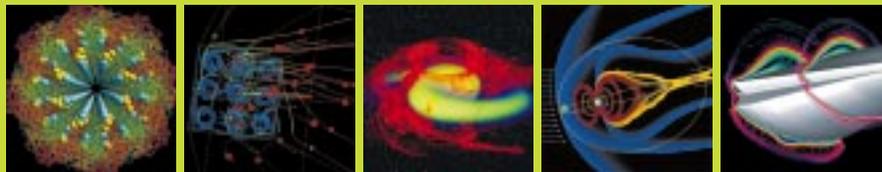


SPECIAL EDITION



# SC2000

National Aeronautics and Space Administration



A look at NASA exhibits, events,  
and speakers at SC2000:  
High Performance Networking  
and Computing Conference.

## Welcome To SC2000

I would like to personally welcome you to NASA at SC2000. The NASA High Performance Computing and Communications (HPCC) Program is pleased to be hosting NASA's presence with exhibits, events and speakers from all around the agency.

You will see that high-performance computing and networking are critical to NASA's quest to expand frontiers on the Earth, in the air and in space. From improving our understanding of observational Earth and space data, to incredible computational models of revolutionary aerospace vehicles, our planet, and even distant stars, high-performance computing is critical to improving our knowledge of ourselves, our transportation systems, our planet, and the universe. In order to pull all of this knowledge together, high-performance networking enables the best and brightest scientists and engineers to dream, design, and discover together, and then allows us to share these discoveries with the American public – especially our children.

**Dr. Eugene L. Tu**  
**HPCC Program Manager**



Judy Conlon/NASA

NASA is showcasing many of its exciting scientific research projects at SC2000. The following are descriptions of demonstrations, paper presentations, videos, and panel discussions listed by participating NASA researchers and collaborators.

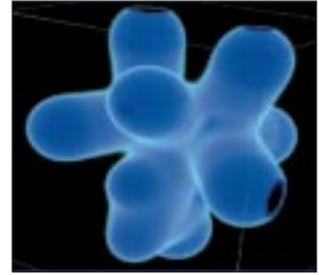
### Ames Research Center

*As NASA's Center of Excellence for Information Technology, Ames provides agency research leadership and world-class capability encompassing the fields of supercomputing and networking, high-assurance software development, verification and validation, automated reasoning, planning and scheduling, and human factors.*

### 'Growler' Visualization Tool

Growler is a modular scientific visualization system, whose development has been guided by NASA's desire to perform remote data visualization, remote computational steering, heterogeneous data analysis, and multi-user collaboration. Growler is based on a high-performance distributed component model, involving interdependent computational and graphical modules. Most modules are inherently collaborative, in that multiple remote users may take part in

and control the same visualization environment. Visualization modules include: VMS (Virtual Mechanosynthesis), the COSMOSViewer, and a computational fluid dynamics (CFD) visualization module which allows for not only the perusal of remote CFD datasets, but also the visualization and steering of an active CFD simulation run, executing on a remote supercomputer.



### Shared Memory Multi-Level Parallelism

The NASA Ames terascale application group has developed a new and highly efficient method for improving the parallel performance of NASA production codes. The new method is termed "shared memory multi-level parallelism," or MLP. The MLP technique has been used to convert a number of NASA codes in CFD (OVERFLOW), molecular dynamics (COSMOS), and climate modeling (FVCCM3). All of these conversions have demonstrated a five to 15 times performance improvement over the best previous parallelization efforts. All of the above examples demonstrate linear increases in parallel across all 512 CPUs when executing on the NASA Ames 512-CPU Origin 2800 system. For OVERFLOW, this amounts to 70 gigaflops of sustained performance. All indications are that these codes will scale to substantially more CPUs when the 1,024-CPU SGI Origin 3000 system is delivered in 2001.

### Digital Sky Virtual Observatory (JPL/NREN)

The Digital Sky Virtual Observatory (DSVO) project will help create a "National Virtual Observatory" providing an intuitive, comprehensive way to remotely navigate immense celestial datasets. Researchers will be able to access and interrelate geographically distributed datasets from multiple sky surveys at arbitrary resolutions. DSVO routes large data streams from diverse locations over high-speed networks like NREN, CalREN2, Abilene, and NTON, combines them on the fly, and projects them for viewing using Virtual Observatory software. For additional information, visit:

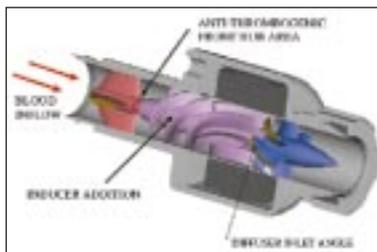
<http://alphabits.jpl.nasa.gov/DPAT/tasks.html#DSVOproj>

### DeBakey Heart Assist Device

Approximately 20 million people worldwide suffer annually from heart failure, a quarter of them in America alone. The use of computational fluid dynamics (CFD) technology has led to major design improvements on the heart assist device, enabling its human implantation. Ames scientists employed NASA Shuttle main engine technology and CFD modeling capabilities, coupled with the NAS Systems Division's high-performance computing technology, to make several design modifications that vastly improved the heart device's performance.

The research team investigated seven different designs, altering cavity shapes, blade curvature, inlet cannula shapes, and impeller tip clearance size.

They then suggested three major design modifications to solve the problems of cell damage resulting from exposure to high shear stress and interrupted regions of blood flow in the DeBakey Ventricular Assist Device (VAD).



### Interactive Multimedia

NASA is reaching into the classrooms of America to make a difference in math, science, and technology education.

The Internet is being used as the primary medium for live interactions with scientists, virtual electronic field trips, collaborative projects, and distance learning activities. The NASA Learning Technologies Project offers more than 50 interactive multimedia projects to provide standards-based learning environments for the use of NASA's scientific data.



### Foil Simulation and Engine Simulation

This technically accurate simulation software is intended to allow educators and students to explore key concepts of wing and engine design. Based on data derived from NASA and other research, it allows examination of cause and effects relationships by enabling the modification of key variables. The software is intended to help teachers and students to grasp related math and science concepts. The content is produced in conjunction with state and local education standards and is intended to draw additional students to aerospace careers.

### ILab: Parameter Study Creation and Submission on the Information Power Grid (IPG)

Parameter study creation and submission has been fully automated with this Graphical User Interface (GUI)-based tool. In addition, the greatest impediment to user acceptance of the IPG/Globus model has been overcome: the job control language, which is required for submitting into the IPG/Globus processing environment, has been eliminated. This GUI tool first automates the creation of parameter studies of arbitrary dimension and then automatically creates all Globus job control language (RSL decks and auxiliary shell scripts) required to actually set up and invoke



Globus jobs on the IPG. The Ames team has demonstrated the ILab capabilities by creating a 12-by-16 case parameter study in Mach number and Alpha (angle of attack) for the X38 Crew Return Vehicle. These 192 flow solution cases were automatically submitted, using Globus, onto two separate Origin 2000 parallel machines in the IPG testbed.

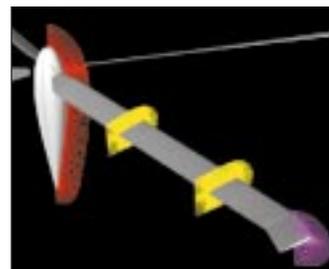
### Parallelization Toolkit in the Legacy Code Modernization Project

With the rapid evolution of high-performance computing systems, code parallelization for new architectures becomes more and more challenging. The process is time consuming and error prone. Automation with tools reduces porting cost and facilitates fast code development.

This parallelization toolkit performs accurate and in-depth program analyses, transforms codes as necessary, and generates parallel codes under nominal user interaction. The engine for the data dependence analysis is provided by CAPTools. The tool performs analysis at the loop nest-level and generates Open MP codes. A set of GUI browsers is implemented for users to interact with the parallelization process. APAPT is a tool that analyzes data affinity and optimizes data placement for the parallel codes. These tools have successfully been used to parallelize several large CFD applications.

### Large-Scale Parallel Distributed Computing

A virtual computing environment that integrates computer resources at different sites, collectively known as NASA's Information Power Grid (IPG), is employed to simulate a large-scale memory and CPU-intensive aerospace application. This illustrates how IPG technology can be used to obtain a rigorous wake flow solution around the hovering rotors of a helicopter, where flow features are dominated by complex vortex dynamics.



Researchers are simulating wake flows around a UH-60 Blackhawk model rotor on the NASA IPG testbed consisting of SGI Origin 2000 platforms. Distributed parallel Navier-Stokes flow computations are made using the OVERFLOW-D code, and adapted to the IPG execution environment by the Globus-MPICH-G toolkit. The flow domain is discretized by multiblock overset meshes.

### Support For Debugging Automatically Parallelized Programs

NASA researchers have developed a system to simplify the process of debugging programs produced by computer-aided parallelization tools. The system uses relative debugging techniques to compare serial and parallel executions to show where the computations begin to differ. If the original serial code is correct, errors due to parallelization will be isolated by the comparison.

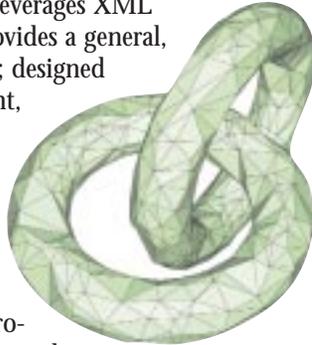


The debugging system uses information produced by the parallelization tool to drive the comparison process. In particular, the debugging system relies on the parallelization tool to provide information about where variables may have been modified and how arrays are distributed across multiple processes.

### AM/FM: Active Metadata/Field Model

NASA researchers present technologies from two complementary projects: Active Metadata (AM) and Field Model (FM). Active Meta-data is a flexible, extensible mechanism for data about data; metadata evolves over data lifetime; AM implementation leverages XML technology. Field Model provides a general, flexible model for field data; designed for very large data; consistent, demand-driven evaluation philosophy; and templated C++ implementation.

Together, the two projects will provide a more complete solution for the scientist. Field Model provides an efficient, extensible mechanism for handling large field data. Active Metadata provides the means to store and utilize the associated metadata, data that are typically much smaller in size but far more diverse in structure.



### IPG Launch Pad

The NASA IPG Launch Pad is a web-based user portal to the Information Power Grid. The Launch Pad provides IPG users with the ability to: view the status of grid resources; prepare, launch, and track jobs; manipulate files at remote locations; and customize the environment to create a personalized view of the portal.

The Launch Pad uses the Grid Portal Development Kit (GPDK) developed by the National Computational Science Alliance's Jason Novotny. GPDK is based on Java servlets and Java server pages. GPDK uses MyProxy, developed by Novotny and Von Welch, to provide security for the Launch Pad.

### Web-Based Interactive Mars Data Archive

This facility serves as a means for Web-based exploration of data returned by Mars missions. Users can select regions of Mars to navigate in 3-D as VRML scenes, which are created on-the-fly by a servlet. The VRML scenes use images and terrain data from the Viking missions, and can include current ultra-high-resolution Mars Orbiter Camera (MOC) images and Mars Orbiter Laser Altimeter (MOLA) elevation profiles as well. A MOC image archive displays MOC images within larger-scale Viking images for regional context, and includes online Java-based image processing. A MOLA data archive allows users to view the MOLA elevation profiles, and to query them for coordinates and elevations.

## Ames/Jet Propulsion Laboratory (JPL) Telescopes In Education

The Telescopes in Education (TIE) program provides the opportunity to use a remotely controlled telescope and charge-coupled device camera in a real-time, hands-on, interactive environment to students around the world. TIE enables students to increase their knowledge of astronomy, astrophysics, and mathematics. The TIE program currently utilizes a science-grade 24-inch reflecting telescope located at the Mount Wilson Observatory in Southern California. An additional telescope will be added this year in Chile to allow students to observe the Southern Hemisphere.

## Glenn Research Center

*Glenn Research Center is NASA's Lead Center for Aeropropulsion and Center of Excellence in Turbomachinery. The center is committed to developing innovative technology and leveraging its computational, analytical, and experimental expertise in turbomachinery to future aerospace programs. Glenn's aeropropulsion program plays a significant role in promoting safe and environmentally compatible U.S. aircraft propulsion systems.*

## NASA Glenn Aeroshark Cluster

An important objective of the High Performance Computing and Communications (HPCC) Program is to accelerate the development of affordable, high-performance computers by domestic vendors. This includes a wide range of computational abilities, from scaled parallel machines to cost-effective high-performance workstations – all networked together in a unified, heterogeneous computing environment.

The NASA HPCC Program is developing tools and techniques

that can be used to forecast the performance of Grand Challenge applications in various high-performance architectures. System evaluations consider the effect of the distributed system environments on performance and are used to analyze the requirements of Grand Challenge applications. For more information, visit: <http://hpcc.grc.nasa.gov/aeroshark/>



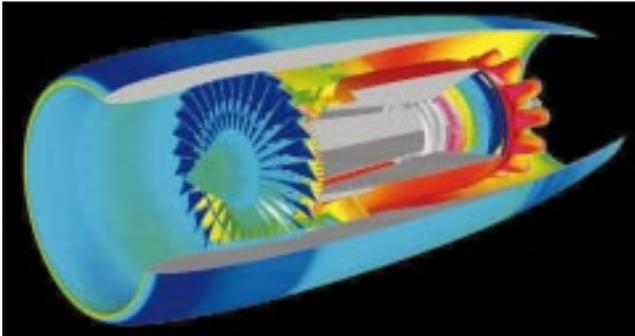
The Pentium III Aeroshark cluster is a network of 64 compute nodes (128 CPUs). The cluster is used to execute significant turbomachinery simulations in a production mode.

## NPSS Airline Seat Kiosk

Scientists at GRC are developing an advanced multi-disciplinary analysis environment for aerospace propulsion systems called the Numerical Propulsion System Simulation (NPSS). NPSS is currently being extended to support the Aviation Safety Program and space transportation.

NPSS is developing an object oriented framework for aerospace propulsion with the integration of multiple disciplines such as aerodynamics, structure, and heat transfer

with numerical zooming on component codes. Zooming is the coupling of analyses at various levels of fidelity. NPSS development includes using the Common Object Request Broker Architecture in the NPSS Developer's Kit to facilitate collaborative engineering. Visit NPSS at: <http://hpcc.lerc.nasa.gov/hpcc2/npssintro.shtml>



### Heterogeneous Cluster Computing

The Dynamic Load Balancing Tool improves the execution time of ADPAC (An Advanced Ducted Propfan Analysis Code) running in parallel on local area and wide area networks formed by combined Unix and Windows/NT environments. Visit the Dynamic Load Balancing web page at: <http://www.cfd.iupui.edu/~scomp99>

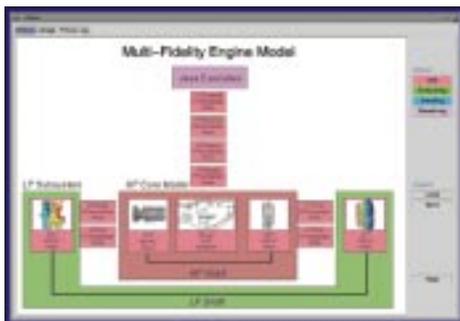
### Local and Wide Area Computing

NASA Glenn researchers are developing a distributed cooperative environment for dynamic load balancing for multiple parallel jobs on local and wide area networks with a Java-based multi-agent environment.

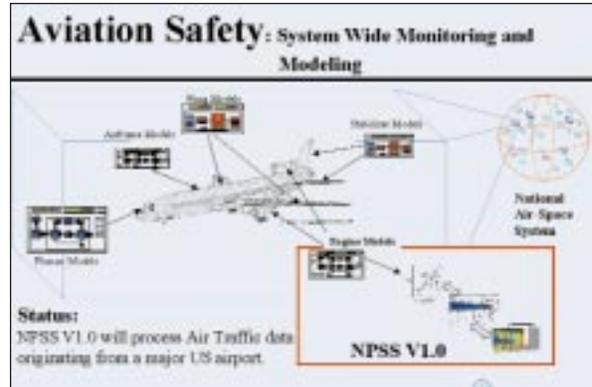
In several single-dynamic load balancing applications, without exchanging information of loads on computers, there will be no optimal parallel load distribution. The Dynamic Load Balancing Web Page can be accessed at: <http://www.cfd.iupui.edu/~scomp99>

### NASA Glenn Power Grid Challenge

The purpose of this simulation is to determine at what RPM the low- and high- pressure subsystems of the Energy Efficient Engine are balanced. During the simulation the Burner element computes a new high-pressure turbine inlet flow based on the exit flow and the combustor model. This high-pressure turbine inlet flow data is then copied to the high-pressure turbine element, constituting a "swap." Next, a selected number of coupled ADPAC (a parallel 3-D CFD code) iterations are performed to get the coupled simulations to settle to this new set of conditions.



After a selected number swaps between the Burner and other elements, the calculated LP and HP net horsepower values are compared against the maximum allowable imbalance value. If necessary, new estimated shaft RPM values are calculated, and another trial begins. The simulation terminates when the shaft powers are balanced.



### Flight Safety

The aviation safety prototype (AvSP) is an airline risk assessing, and parametric performance driven, collaborative development effort between Glenn Research Center (GRC) and Ames Research Center (ARC). AvSP gives Internet users the ability to execute and interact with a secure system that is monitoring and modeling the air traffic over a major U.S. airport. The AvSP processes flight data was previously gathered for arrivals and departures through take-off and landing, as well as weather information, aircraft models, GRC engine models, ARC wing models, and combined systems, all executing within an integrated framework.

The AvSP combines various new technologies including GRC's Numerical Propulsion System Simulation (NPSS) V1.0 to process 0D engine data, Common Object Request Broker Architecture (CORBA), and Information Power Grid (IPG). The Aviation Safety Prototype Web Page is located at <http://as.grc.nasa.gov/AvSP>

### Goddard Space Flight Center

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

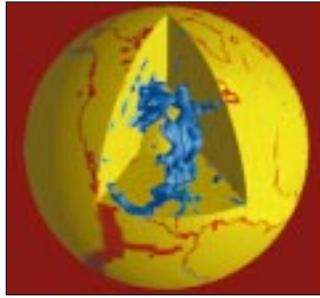
### Journeys Through Earth and Space

Why are the Rocky Mountains so far inland? How do we preserve the changing Amazon rain forest? When will the sun fling parts of itself towards Earth?

NASA is tackling questions like these inside supercomputers. Here, billions of calculations per second recreate the universe mathematically. Supercomputers can process observations into a motion picture. Or, they can solve equations that describe realities seen and unseen. To understand and predict nature through computation,



NASA started the Earth and Space Sciences (ESS) Project. It is a unit of the agency's High Performance Computing and Communications (HPCC) Program. This video magazine follows three ESS research teams on their journeys to discovery. For more information, visit: <http://esdcd.gsfc.nasa.gov/ESS/grand.st2.html>



### Beowulf System Software

The latest system software works to implement a Beowulf-aware environment that presents a single system image to the user. No longer simply a collection of computers, Beowulf is a computing system with a collection of processing resources. Up-to-date Beowulf news can be found at: <http://www.parl.clemson.edu>

### Jet Propulsion Laboratory (JPL)

*The mission of the Jet Propulsion Laboratory (managed for NASA by the California Institute of Technology) is to expand the frontiers of space by conducting challenging robotic space missions for NASA to: explore our solar system; expand our knowledge of the universe; further our understanding of Earth from the perspective of space; and pave the way for human exploration.*

### Distributed In-situ Mission Simulation

The Virtual Mission project at JPL employs modeling and simulation technology to create a software-based mission lifecycle in order to validate the mission system design in a comprehensive manner. The Virtual Mission Lifecycle develops subsystem behavior models based on performance and operation properties of each subsystem, integrates the subsystem models via an operation executive, and executes realistic observation scenarios by providing a synthetic mission environment. Team members include Meemong Lee, Wenwen Lu, Jeng Yen, and Roberto Mendoza. For more information, visit: <http://wwwmsim.jpl.nasa.gov/MS>

### PYRAMID – Parallel Unstructured Adaptive Library

PYRAMID is a software library for performing parallel adaptive mesh refinement on unstructured meshes. It uses triangular or tetrahedral elements for the adaptive mesh. A suite of well-designed and efficiently implemented modules that perform operations in a typical parallel AMR process are included.

The library is implemented in Fortran 90 with MPI, and it supports mesh quality control, adaptive refinement, load balancing, mesh migration, and visualization all in parallel. More information on PYRAMID can be found at: <http://hpc.jpl.nasa.gov/APPS/AMR>

### Langley Research Center

*The main focus of the High Performance Computing and Communications (HPCC) Program Computational Aerospace Sciences Project activity at Langley Research Center is the application of high-performance computing technology to the multidisciplinary design optimization (MDO) of aerospace vehicles.*

### Multidisciplinary Optimization

At Langley Research Center, research is underway to apply high-performance computing technologies to the multidisciplinary design optimization (MDO) of Reusable Launch Vehicles (RLV). Various analysis codes are being integrated using a commercially developed framework such as iSIGHT. The framework allows the use of networked heterogeneous workstations, vector supercomputers, and a scalable parallel testbed.

Initially, coarse-grained parallelism is being used to generate many vehicle designs simultaneously for trade studies. This capability to perform conceptual MDO using low-fidelity codes will be demonstrated on a full RLV configuration during the coming year. Later, high-fidelity analysis codes will be incorporated, giving RLV designers, for the first time, the capability to perform preliminary vehicle design with complex multidisciplinary interactions in a reasonable amount of time. The group's web address is: <http://hpccp-www.larc.nasa.gov>



### NASA Participation: Technical Papers

*A number of NASA scientists have been honored by having their research papers selected for presentation by the SC2000 Program Committee. From the 179 papers submitted to the committee, 62 will be presented at Dallas. Papers presented by NASA and its collaborators follow.*

#### Tuesday, November 7

Session: Scheduling • 1:30 p.m. to 3 p.m.

#### An Object-Oriented Job Execution Environment

*Lance Smith and Rod Fatoohi*

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

Session: MPI/OPENMP • 1:30 p.m. to 3 p.m.

### **A Comparison of Three Programming Models for Adaptive Applications on the Origin 2000**

*Hongzhang Shan, Jaswinder P. Singh, Leonid Oliker, and Rupak Biswas*

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

#### **Wednesday, November 8**

Session: Applications II • 1:30 p.m. to 3 p.m.

### **Parallel Unsteady Turbo-Pump Simulations For Liquid Rocket Engines**

*Cetin C. Kiris, Dochan Kwak, William Chan*

This paper reports the progress being made towards complete turbo-pump simulation capability for liquid rocket engines. The Space Shuttle Main Engine (SSME) turbo-pump impeller is used as a test case for the performance evaluation of the MPI, hybrid MPI/Open-MP, and MLP versions of the INS3D code. Then, a computational model of a turbo-pump has been developed for the shuttle upgrade program. Relative motion of the grid system for rotor-stator interaction was obtained by employing overset grid techniques. Unsteady computations for an SSME turbo-pump, which contains 101 zones with 31 million grid points, are run on Origin 2000 systems at NASA Ames. The approach taken for these simulations, and the performance of the parallel versions of the code are presented.

#### **Thursday, November 9**

Session: Data Grid • 10:30 a.m. to Noon

### **Computing and Data Grids for Science and Engineering**

*William E. Johnston, Dennis Gannon, Bill Nitzberg, Leigh Ann Tanner, Bill Thigpen, Alex Woo*

We use the term “grid” to refer to a software system that provides uniform and location-independent access to geographically and organizationally dispersed, heterogeneous resources which are persistent and supported. While, in general, grids will provide the infrastructure to support a wide range of services in the scientific environment (collaboration and remote instrument control) in this paper we focus on services for high-performance computing and data handling. We describe the services and architecture of NASA’s Information Power Grid – an early example of a large-scale Grid, and some of the issues that have come up in its implementation.

Session: Gordon Bell II • 10:30 a.m. to Noon

### **Efficient Simulation of Carbon Detonations**

*A.C. Calder, B.C. Curtis, L.J. Dursi, B. Fryxell, G. Henry, P. MacNeice, K. Olson, P. Ricker, R. Rosner, F.X. Timmes, H.M. Tufo, J.W. Turan, M. Zingale*

Simulations and performance results of nuclear burning fronts in supernovae on the largest domain and at the finest spatial resolution studied to date will be presented. These simulations were performed on the Intel ASCI-Red machine at Sandia National Laboratories using FLASH, a code developed at the Center for Astrophysical Thermo-nuclear Flashes at the University of Chicago. The panel will discuss the key algorithms and their implementation as well as the optimizations required to achieve sustained performance of 238 GFLOPS on 6,420 processors of ASCI-Red in 64-bit arithmetic.

### **NASA Participation: Panel Discussions**

Thursday, November 9 • 3:30 p.m. to 5 p.m.

### **Petaflops Around the Corner: When? How? Is it Meaningful?**

Moderator: *Neil Pundit*

Panelists: *Marc Snir, Bill Camp, Thomas Sterling, Paul Messina, Rick Stevens, Pete Beckman*

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

Friday, November 10 • 10:30 a.m. to Noon

### **Megacomputers**

Moderator: *Larry Smarr*

Panelists: *Andrew Chien, Ian Foster, Thomas Sterling, David Anderson, Andrew Grimshaw*

We are nearing a major discontinuity in the evolution of high-performance computing, creating a third era of supercomputing. The first era, which ended with the Cray 1, sought performance by building the fastest single processor possible. The second era, starting with the Illiac IV, sought performance through tens to thousands of identical processors. The third era couples very large numbers of heterogeneous computers on networks to create a virtual supercomputer. Questions to be addressed by the panel include: What is the software architecture model? What are the classes of applications most likely to move to this new computing fabric? And what new computer science research frontiers are important in this more “biological” effervescent style of architecture?



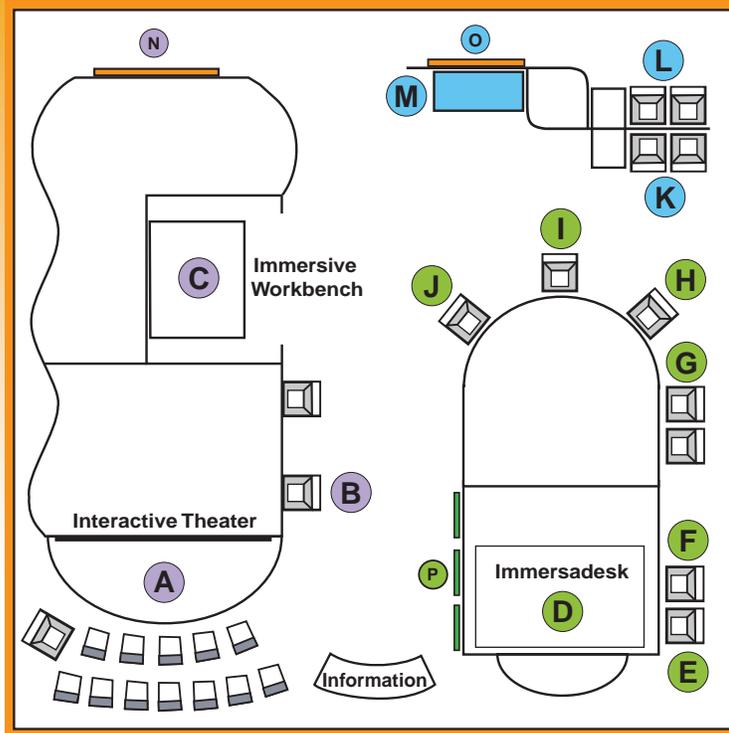
SC2000





National Aeronautics and Space Administration

# SC2000 Booth Map



- A** Interactive Theater:  
 Distributed In-situ Mission Simulation  
 Jim Taft Presentation - Origin Systems at ARC  
 Pyramid: Parallel Unstructured Adaptive Mesh Refinement  
 Remote Exploration and Experimentation (REE)  
 Videos:  
 Computational Aerospace Sciences  
 Journeys Through Earth and Space  
 Learning Technologies  
 NASA Research & Education Network  
 Numerical Propulsion System Simulation  
 Remote Exploration and Experimentation
- B** Multidisciplinary Optimization
- C** Visualization@Ames (Immersive Workbench)  
 AM/FM: Active Metadata / Field Model  
 Growler  
 Web-based Interactive Mars Data Archive
- D** Immersadesk:  
 Digital Sky Virtual Observatory  
 Immersive Physics
- E** Digital Sky Virtual Observatory
- F** NASA Employment Information
- G** FoilSim / EngineSim  
 Interactive Multimedia  
 Telescopes in Education
- H** Parallelization Toolkit in the Legacy Code  
 Modernization Project  
 Support for Debugging Automatically Parallelized  
 Code
- I** Job: Parameter Study Creation and Submission on  
 Large-Scale Parallel Distributed Computing
- J** DeBakey Heart Assist Device  
 IPG Launch Pad
- K** NASA Glenn Aeroshark Cluster  
 NASA Glenn Power Grid Challenge  
 Heterogeneous Cluster Computing  
 Local and Wide Area Computing  
 Multiclustering Using Load Sharing Facility (LSF)  
 Wrapping Legacy Scientific Applications
- L** Beowulf System Software  
 CERSe: Component-Based Environment for  
 Remote-Sensing  
 Problem Solving Environments  
 PVFS: Parallel Virtual File System
- M** Flight Safety
- N** **O** Landsat Mosaic Images
- P** Literature

# Calculating Feature Evolution in Semiconductor Processing

Ames researchers have developed a computational code to understand the behavior of plasma on a silicon substrate used in integrated circuit manufacturing

By Helen H. Hwang

While the quest for the ever-faster semiconductor transistor continues, the size of the devices has been steadily decreasing. The current generation of processors has dimensions of 0.18  $\mu\text{m}$  (microns), while the industry hopes to shrink them by a factor of 20 within a decade. The production of these transistors involves many complex steps of etching trenches and gates on a semiconductor wafer (See *How to Etch a Trench*, page 9). NASA scientists are currently developing methods to accurately predict the etched trench shape under a variety of processes.

In the Plasma Processing Modeling group, part of the Devices and Nanotechnology Integrated Product Team at Ames Research Center, an effort is underway to simulate trench etching in order to optimize conditions to create better etch profiles and rates. The modeling project involves developing a simulation which calculates etch rates and profiles based on material properties as well as properties of the plasma.

## SPELS: Computations Are the Key to Success

The team has developed the SPELS (Simulation of Profile Evolution using Level Sets) computational code to predict etch profile evolution. This code is based on an approach known as “Level Sets,”

which is a significant departure from traditional methods relying on so-called “string” algorithms. The latter assumes

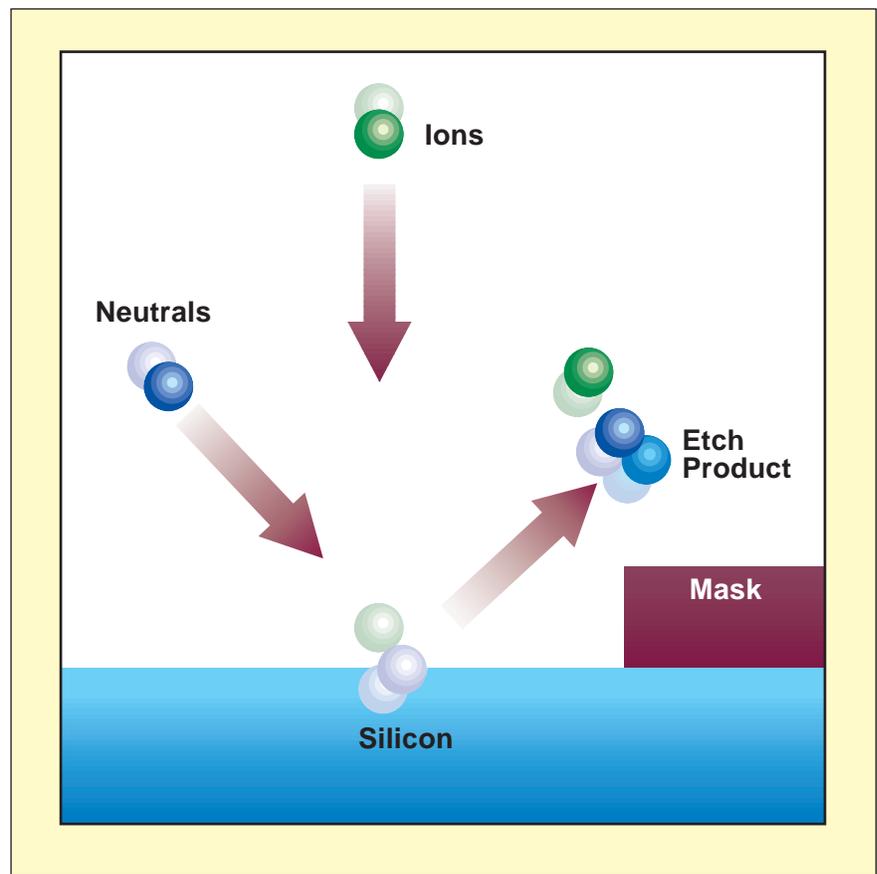
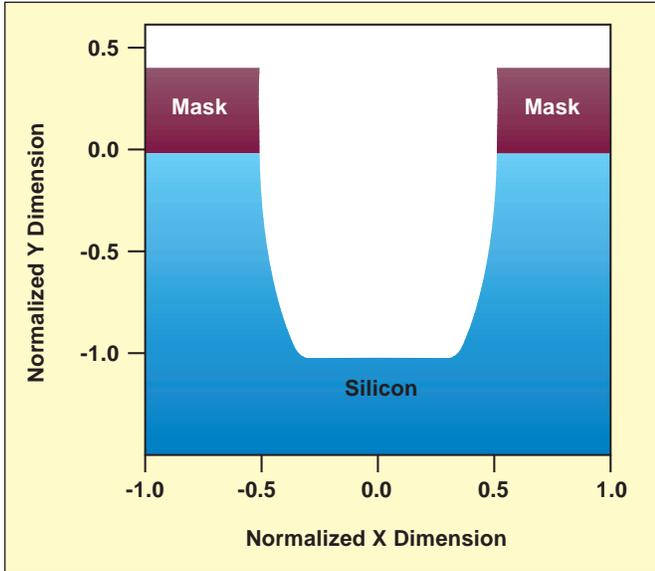


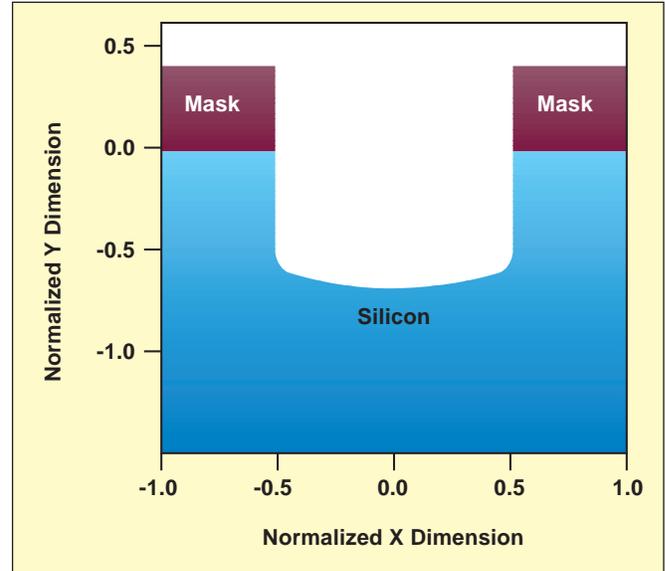
Figure 1: Plasma etching of silicon uses chemicals that are 'directionally blind.' Transistors require etched trenches with vertical walls, thus the etching process must be controlled with atomic precision. Masking prevents etching in unwanted areas.

that a moving surface can be described as a string with many nodes, and that the string is moved according to the local etch rate at each node. Due to numerical issues, the string can sometimes cross itself, or create a loop. Removing the loops, or de-looping, requires a significant amount of computer simulation time. In contrast to

string algorithms, the Level Set method accounts for the changes between the plasma and the silicon wafer's surface through a partial differential equation (PDE). SPELS not only solves this PDE, but more importantly, considers the physics and chemistry associated with the calculation of the etch rate itself.

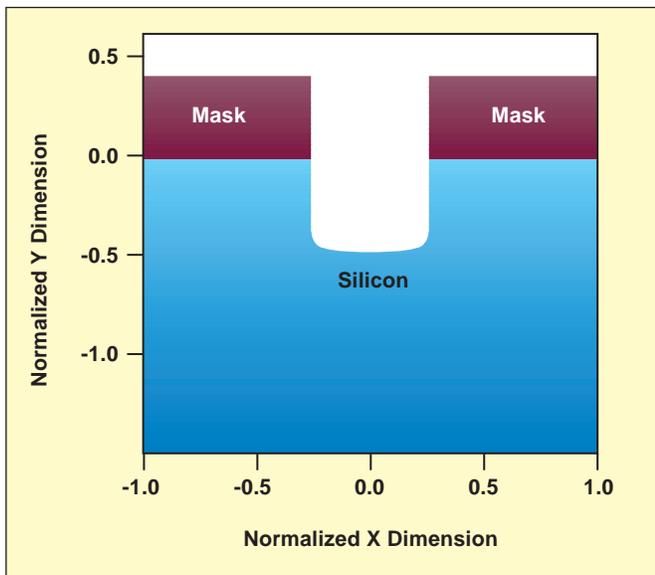


**Figure 2a**

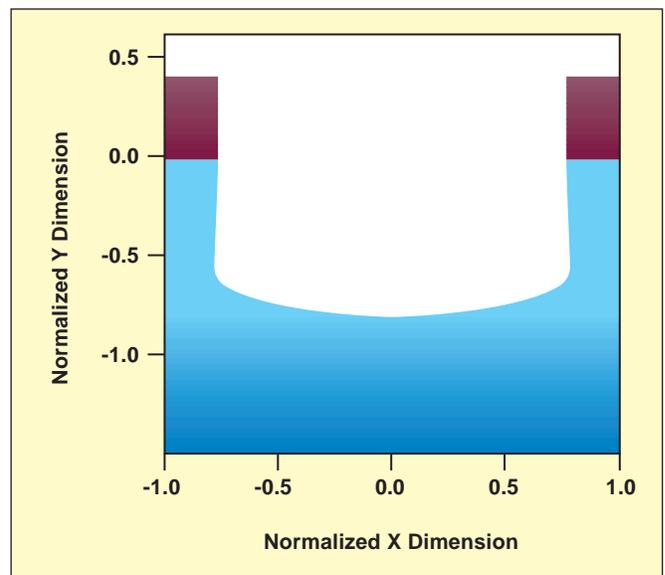


**Figure 2b**

The examples above were generated from SPELS: For plasma conditions with high neutral flux, the neutrals fully passivate the surface and etching occurs quickly. The profile has a tapered shape but the overall etched depth is large, as shown in Figure 2a. Changing the plasma conditions such that the ion flux is substantial, the overall profile has straighter sidewalls than for the high neutral flux case. However, the total etched depth is considerably less, as fewer neutrals are available to react with the silicon; thus less surface products can be removed. The center of the trench can 'collect' more neutrals and thus etches faster, which leads to the curved trench bottom, as seen in Figure 2b.



**Figure 3a**



**Figure 3b**

Smaller structures etch slower due to Reactive Ion Etching (RIE) lag. Neutrals arrive from the plasma in all directions due to collisions, and neutrals at large angles are screened by the mask and trench opening. Fewer neutrals passivate the surface of smaller trench openings (Figure 3a). In Figure 3b, increasing the width of the trench opening allows more neutrals to passivate the trench surface, which leads to a larger etched depth.

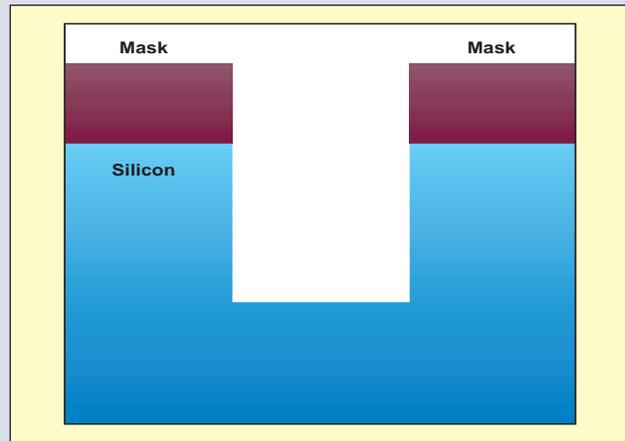
## 8 Gridpoints

# ***Chips 101: How to Etch a Trench***

**E**tching refers to removing materials selectively from specified locations, while leaving the remaining structure unaffected. Areas that are under a mask are “hidden” from the reactive chemicals, while the rest of the material can be etched (see Figure 1, page 15, and Figure 4 at right).

For a long time, etching has been accomplished by using chemicals. This process is “directionally blind” – etching proceeds isotropically in all three directions at the same rate. However, modern transistors require trenches and gates with vertical walls, which means that the etch action must stop with atomic precision when encountering the valuable, active semiconductor layer. Chemical etching, while quick, lacks this precision.

A new method, called “plasma etching,” has become the industry standard in the last decade because it is directional in nature. Plasma-based etching is done in plasma chemical reactors which consist of a vacuum chamber, power supply, and gas handling system. The gas-phase chemical compounds are separated into neutral fragments, positive and negative ions, and electrons. A reactor with an inductive power supply, known as an Inductively Coupled Plasma (ICP) reactor, is one such tool used in industry for plasma etching. Typically, chlorine-containing mixtures of gases are used to etch silicon; fluorocarbon gases, such as  $\text{CHF}_3$  and  $\text{C}_2\text{F}_6$ , are used to etch silicon dioxide ( $\text{SiO}_2$ ).



**Figure 4:** The ideal etched trench consists of straight walls and bottom, and square corners.

Plasma etching can be highly directional, unlike chemical etching. Some of the neutral fragments of the plasma react with the material in the trench to produce a protective film. The positive ions respond to the voltage difference between the center of the plasma and wafer surface and accelerate toward the trench. The combination of the electric field direction and lack of collisions en route makes the ions bombard the wafer surface vertically, thus removing the protective film on the horizontal surface but not the sidewall. A vertical feature, as shown in the figure above, would evolve as this action continues for a few minutes.

Etch rates will depend on reactor conditions, such as operating power and gas pressure, as well as material properties – such as the composition of the wafer and reacting gases used. In addition to knowing how fast the material can be etched, the shape of the structure that is etched must be determined (a rectangular trench is desired to manufacture transistors). In order to calculate the etch rates and overall etched form, researchers must understand the plasma’s interaction with the semiconductor wafer.

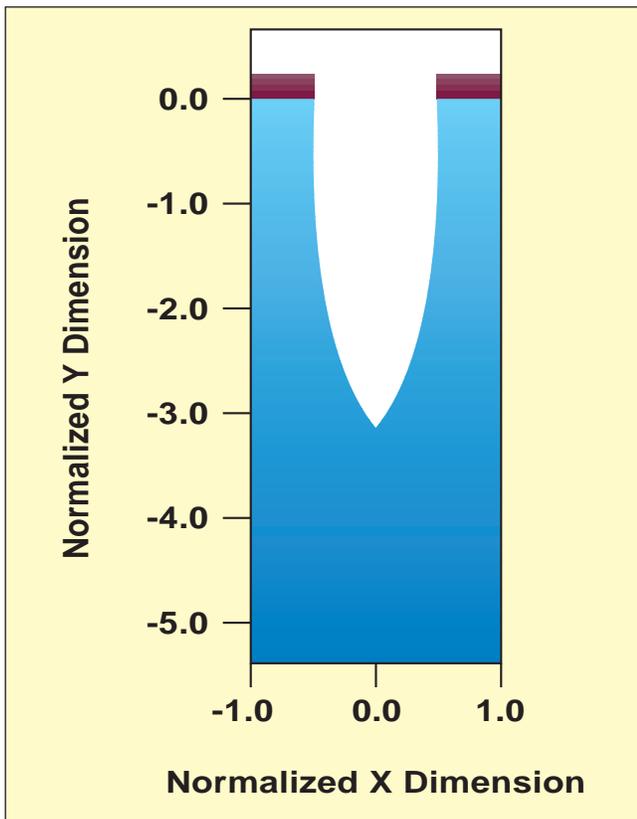
## **Predicting Etch Rate and Trench Shape**

Although energetic ions can directly sputter the wafer, generally the mechanism for etching in low-temperature plasmas is done by ion-enhanced etching. The first step is called passivation, where the neutral species – such as atomic and molecular chlorine, chemically react with the surface of the wafer. For silicon wafers, the products formed are silicon chlorides ( $\text{SiCl}_x$ ). These surface compounds are then detached by the energetic ions, and thus the semiconductor material is etched. Although ions can directly sputter off silicon atoms from the surface, this requires ions with energies of more than 1 kilo-electronvolts. Impact by such high energy ions can cause undesirable damage to the silicon wafer.

For the plasmas used in these applications, ions generally have energies of only 10 to 100 electronvolts. Ions also arrive perpendicular to the surface, due to acceleration from electric fields in the plasma. The normal incidence of the ions helps to etch away the surface products. Thus the ions control the etch rate, but etching cannot occur without passivation, making the flux of the neutrals to the surface equally important.

One determining factor of the etch rate is the ratio of the rate of neutrals to the rate of the ions arriving at the surface. In general, a higher ratio of neutral flux to the surface leads to a higher etch rate. This results from surface silicon transforming into silicon chlorides. If there is an insufficient amount of passivation, the semiconductor will not etch, regardless of the number of ions that bombard the silicon surface.

An ideal trench consists of straight walls and a flat bottom, with square corners (see Figure 4, above). Any non-ideal features, such as bowed walls or rounded corners, can impact subsequent processing steps, such as depositing material into the trench. The resulting device may not



**Figure 5:** A trench etched in a pure chlorine plasma has an aspect ratio of 3:1 (depth:width). Etch rates are lower in the corners due to neutral shadowing. Obtaining high aspect ratios, such as 50:1, is impossible under these conditions. Other gas mixtures, often containing hydrogen bromide, are required for etching high aspect ratio devices.

work properly, or may prevent packing devices closer together as desired.

Obtaining an ideal trench shape with a fast etch rate is difficult, however. When the neutral flux is high, the etch rate for the silicon is fast because the neutrals passivate the surface and the ions are able to remove the silicon chlorides. However, the trench has a tapered shape, as shown in Figure 2a. Etching trenches in plasmas with higher ion fluxes leads to straighter sidewalls (see Figure 2b, page 8). Although the trench has a more ideal shape, the etch rate is lower, due to fewer neutrals passivating the surface. The trench bottom is curved because the center collects more neutrals than the corners.

Another factor that can influence etch rate is the size of the trench opening. Both experiments and simulations have shown that smaller devices etch at a slower rate than larger devices (see Figures 3a and 3b, page 8). This scaling difference is called RIE lag, or Reactive Ion Etching lag. In the plasma, neutrals collide with other neutrals and the reactor walls, so by the time they arrive at the trench opening, they have widely varying impact angles. Neutrals that arrive with a severe angle to the opening are blocked by the

height of the mask, especially for small trench openings. In other words, large trench openings are more efficient at collecting the neutrals, because fewer of them are “screened” due to the mask. Because the larger trenches have more neutrals that arrive at the bottom of the trench, they etch faster than the smaller trenches.

This screening effect of the neutrals is often referred to as neutral shadowing. Shadowing can also affect the overall etched depth. Although chlorine plasmas lead to high etch rates with silicon, rarely is chlorine used alone. Instead, chlorine is usually used in combination with other gases in commercial usage.

One well-known problem is that chlorine is not able to etch features with high aspect ratios, or features that are much deeper than they are wide. To obtain a fast etch rate, the neutral flux should be large, which leads to tapered shapes where the bottom of the trench is narrower than the surface opening. In etching a higher aspect ratio feature, the bottom of the trench will eventually close off under these conditions. For typical conditions in an Inductively Coupled Plasma discharge (ICP) using pure chlorine, the amount of neutrals that reaches the bottom is restricted due to shadowing. As seen in Figure 5, the maximum obtainable aspect ratio is about 3:1, which falls short of the desirable ratios of around 50:1. Clearly, some other gas chemistry is necessary to obtain straight wall etching for higher aspect ratio features in silicon and assist etching at the bottom of the trench.

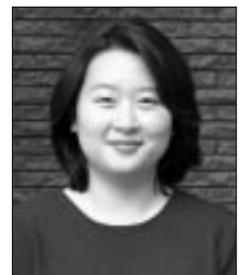
### Collaboration and Further Studies

One effort underway is a collaboration with Infineon Technologies, in which trenches etched in hydrogen bromide are being studied. Although hydrogen bromide etches in a manner similar to chlorine, higher aspect ratio features (around 9:1) have been achieved. Achieving such high aspect ratios is believed to be critical for the next generation of semiconductor chips. By matching simulation results to the experiments, we hope to discover what makes hydrogen bromide amenable to high aspect ratio etching and help optimize the process. 

### Acknowledgements

The author would like to thank T. Barth, D. Bose, T. R. Govindan, and M. Meyyappan for valuable discussions.

*Helen Hwang is a research scientist in the NAS Systems Division's Research Branch. She obtained a Ph.D. in electrical engineering from the University of Illinois at Urbana-Champaign. She has studied various aspects of plasma processing for the past 10 years, including modeling feature profile evolution, particulate contamination, and radiofrequency discharges.*



# VMS Architecture Evolves As New Modules Are Developed

**The new tool, 'growler,' is capable of supporting multiple servers in an integrated collaborative environment**

**A** NAS-developed application for scientific visualization has expanded into an interactive, distributed, and collaborative virtual environment for helping researchers visualize and interact with atoms in an integrated environment. The prototype application, called “growler” (formerly Virtual Mechanosynthesis), demonstrates the use of distributed modules and computational steering (interaction with running simulations). Developed by NAS researchers Chris Henze and Bryan Green, growler is a distributed framework that provides users with a tool to computationally and graphically explore the molecules and materials they are working with.

The most recent addition to the prototype growler environment is COSMOS (COmputer Simulations of MOlecular Systems), a scientific code used to simulate chemical and biomolecular processes such as protein folding, protein-DNA interactions, and functions of cell membranes.

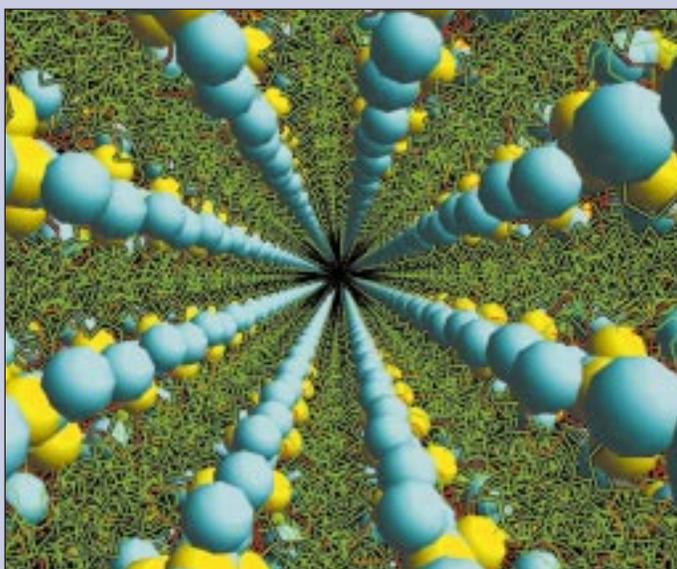
“Growler provides a distributed component framework, and there are currently two primary types of modules that plug into it – viewers and simulation modules,” explains Henze. “Growler mediates all communication – it sets up a communication link between the simulation and the viewer,” adds Green. “The viewer renders the data it gets from the server as it arrives.” For example, as soon as COSMOS completes a time step,

it sends the updated atom positions to the viewer, which is responsible for updating the view of the molecules for users.

## Designed for Distributed Communications

The growler framework is designed for distributed communication – the simulations being studied do not have to reside on the same computer from which they are being viewed and/or controlled. Scientists can access and manipulate different modules within a single graphical environment to compare data and view interrelated simulations. Growler works like a standard editor – it can have multiple files open (or viewers, in the case of growler), but only one

file is active at a time. With a simple function command, users can bring remaining files to the foreground at anytime. Instead of editing files, however, growler allows users to interact with a 3-D graphical environment defined by the viewer module. Users can work with the most up-to-date simulations – the viewer renders the data provided by the server in real time.



This image provides a view down the length of a protein nanotube, which was visualized using the Protein Data Bank (PDB) viewer module of growler. The crystal structure from the PDB (1A6E) actually contains information for only a single pair of red and green strands – the viewer imposes symmetry operations to arrange eight pairs around the circumference of the tube (which is also opened up to expose its interior). The green and yellow spheres represent particular amino acid residues. The image is representative of a collaborative effort between NAS researcher Chris Henze and Astrobiology researcher Jonathan Trent.

(Visualization by Chris Henze and Bryan Green)

Several simulations or viewer modules can be connected to growler simultaneously to view multiple computational runs. One module currently operating in the growler environment is COSMOS. The COSMOS molecular dynamics code allows researchers to model the dynamics of biomolecules. A COSMOS server ships data and atom positions to the COSMOS

At right is a snapshot from a 'live' simulation using the biomolecular dynamics code COSMOS within the growler environment. The image shows a Glycophorin-A protein molecule spanning a membrane-like layer of hexane (purple molecules) sandwiched between two layers of water (blue-green molecules). The variously colored spheres represent different atom types in the protein molecule. COSMOS is a molecular dynamics code developed by a group led by the NASA Center for Computational Astrobiology acting director Andrew Pohorille.

(Visualization by Chris Henze and Bryan Green)

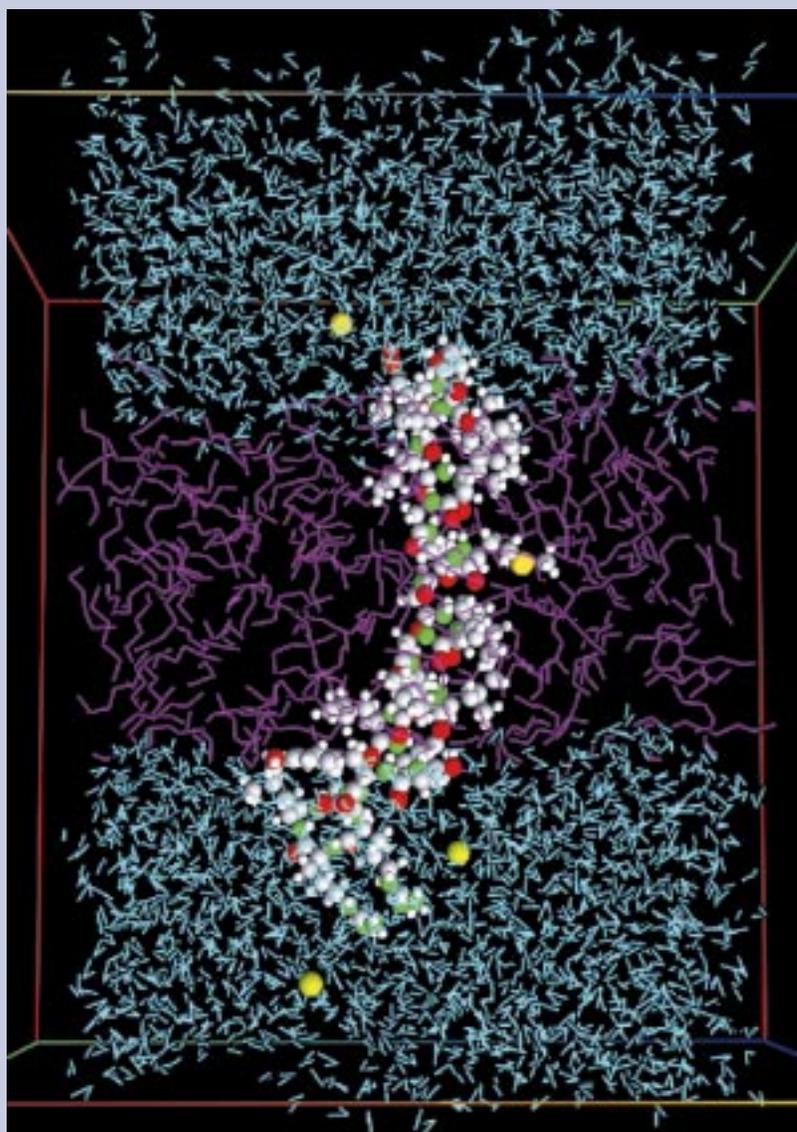
viewer, both of which are plugged directly into the growler framework. Running COSMOS in growler stretches the limits of protein research for scientists like Andrew Pohorille, acting director of the NASA Center for Computational Astrobiology. "Growler is a wonderful tool," says Pohorille, "because its graphical representation provides a better understanding of structure and function of proteins."

The nascent capabilities of COSMOS running in growler do not yet include manipulation of molecular structures. "Basically, we're just watching the simulation right now, we can't actually 'tweak' the data," explains Green. The team is currently evaluating what types of manipulations would be useful for COSMOS users. Henze and Green first demonstrated COSMOS running in growler for Pohorille on August 15, using a Glycophorin-A protein in a membrane.

Other modules operating in the growler environment include Brenner and a Protein Data Bank (PDB) file viewer. Developed by Don Brenner of North Carolina State University, Brenner is a molecular dynamics code implementing the Brenner Potential to study the behavior of hydrocarbons, and particularly fullerenes. The PDB is a collection of biomolecular structural information, including proteins. Running the PDB file reader and viewer in growler enables researchers to view static structures of proteins.

### Work Continues on Prototype

In the near future, Henze and Green will be running a parallelized version of COSMOS in the growler environment. NAS researcher Jim Taft, who specializes in scientific code optimization, has been working closely with Pohorille to parallelize the COSMOS code so it executes more time steps per second. Henze and Green have integrated OVERFLOW, a fluid dynamics code for simulating flow around aircraft and through jet engines, into the growler environment. They plan to add computational steering to the OVERFLOW package in the near future. OVERFLOW



will provide dynamic, run-time visualization, and steering of simulations. Currently, COSMOS is limited by a factor of time – the most interesting biological processes take milliseconds, three orders of magnitude greater than current capabilities. Extending simulation time will require faster computers, longer runs, and more efficient algorithms.

While growler is a "pre-alpha" prototype, Green and Henze believe it has the potential to be a highly useful problem-solving environment for NASA researchers in several disciplines. The Ames team will present growler's capabilities at the SC2000 conference in Dallas this November (see SC2000 insert page 2A). 

— Holly A. Amundson

**Editor's Note:** *The following people have contributed to the development of the COSMOS code; Eric Darve, Vitaly Galinsky, and Kevin Lin, all in the Ames Biomolecular and Cellular Modeling Program; Nagi Mansour, Karim Sharif, Jim Taft, and Alan Wray, in the NAS Systems Division, and Petros Koumoutsakos and Jens Walther from ETH, Zurich.*

# Calendar of Events

## **December 10 – Monterey, California**

The second workshop on media processors and digital signal processors (MP-DSP) brings together researchers, computer architects, and engineers working in fields related to multimedia and digital signal processing. The workshop will provide a forum for presenting and exchanging new ideas and experiences in this area. Topics include micro-processor architectures (superscalar, VLIW, multithread, vector machines), compiler techniques, and software optimizations for multimedia, and DSP applications. Visit the workshop homepage at: <http://www.ee.ucla.edu/mp-dsp/>

## **March 5-7 – Amsterdam**

The Sixth Grid Forum Workshop (Global Grid Forum 1) will be held March 5-7. Grid Forum is for individual researchers working on distributed computing. The forum focuses on the promotion and development of grid technologies. For more information, visit: <http://www.gridforum.org/>

## **March 19-21 – Chapel Hill, North Carolina**

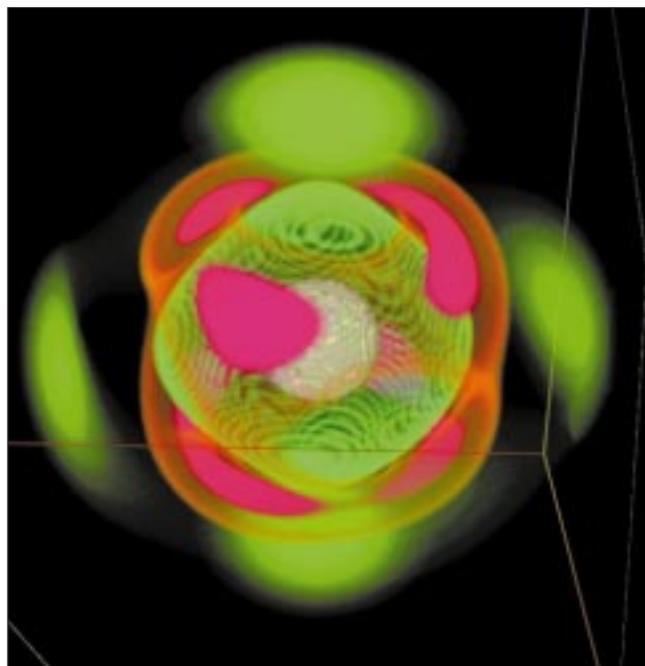
The 2001 Symposium on Interactive 3-D Graphics focus on the concept of real-time interactive 3-D computer graphics and multimedia. The symposium will include formal paper sessions and demonstrations by research groups and vendors. Workshop topics include: interaction techniques for 3-D systems, high-performance 3-D graphics architectures, hardware, and software for virtual world interaction. For more details see: <http://www.siggraph.org/conferences/i3d/>

## **March 26-29 – Oxford, United Kingdom**

The seventh international Conference on Numerical Methods for Fluid Dynamics is organized by the Institute for Computational Fluid Dynamics (ICFD) to bring together scientists and engineers in the field of CFD to review recent advances and exchange ideas on techniques for modeling fluid flows. For more information about this conference, visit: <http://web.comlab.ox.ac.uk/oucl/people/bette.byrne.html>

## **April 23 – San Francisco**

HCW 2001, the tenth Heterogeneous Computing Workshop, is associated with the International Parallel and Distributed Processing Symposium, sponsored by the IEEE Computer Society. Some of the topics covered will be: task mapping and scheduling, resource management, network-centric computing, and real-time distributed systems. Further details are available at: <http://myrtle.cs.umanitoba.ca/hcw/>



Electronic charge distribution around a cobalt atom. Simulation and visualization by Chris Henze

## **April 23-27 – San Francisco**

The International Parallel and Distributed Processing Symposium (IPPS/SPDP) is designed for scientists and engineers from around the world to present their research and latest discoveries in all facets of parallel computation. The symposium will include technical sessions, workshops, tutorials, and commercial exhibits. For conference proceedings and other details, visit: <http://www.ipdps.org>

## **April 27 – San Francisco**

The eighth Reconfigurable Architectures Workshop (RAW 2001) will be held in conjunction with the International Parallel and Distributed Processing Symposium (IPDPS 2001). RAW 2001 is a gathering designed for researchers to present ideas, results, and ongoing research on both theoretical and industrial advances in reconfigurable computing. The main focus of the workshop is on run-time reconfiguration. For more details see: <http://www.cse.unsw.edu.au/~raw/>

## **May 21-25 – Indian Wells, California**

Cray User Group Summit 2001: Hosted by NASA's Jet Propulsion Laboratory, the conference will provide an opportunity to share the latest information on high-performance computing solutions. The conference will include technical sessions, general sessions, and tutorials. Paper topics will include; communications and data management (mass storage, networking), high performance solutions (applications, performance, visualization), and programming environments (compilers and libraries, software tools). Submit abstracts by December 8, 2000. For additional information, see: <http://www.fpes.com/cugs01>