Summer 2000

Gridpoints

The Quarterly Publication of the Numerical Aerospace Simulation Systems Division

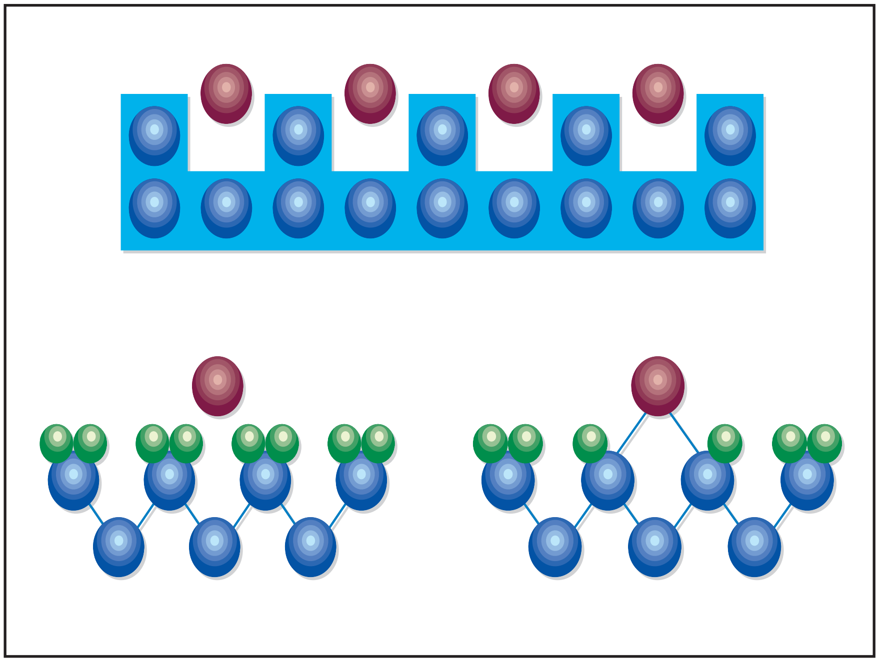

NAS researchers are developing atomic scale transistors to enable future microelectronics. Page 10

FEATURES

Features

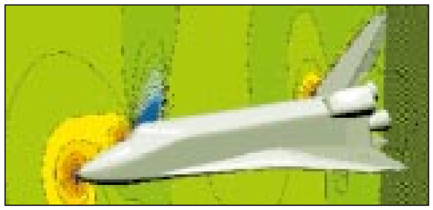

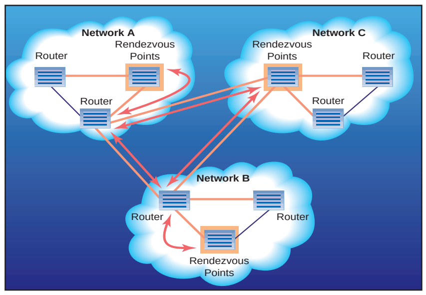

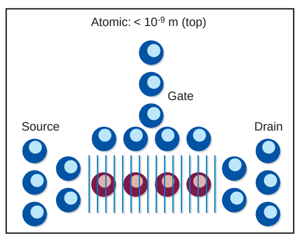

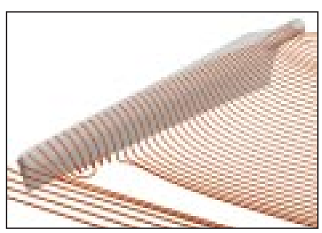

On The Cover:

NAS researcher Toshishige Yamada has conducted experiments with atomically precise structures using a substrate atomic lattice, floating of atoms above the substrate, and chemical bonding of atoms to the substrate. Yamada's research will enable future atomic scale devices. See page 10 for the full story.

Gridpoints

The Quarterly Publication of the Numerical Aerospace Simulation Systems Division

Vol. 1, No. 3 • Summer 2000

Executive Editors: Bill Feiereisen, John Ziebarth

Editor: Nicholas A. Veronico

Writers: Holly A. Amundson, Toshishige Yamada

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: Michael Aftosmis, Dick Anderson, Bill Arasin, Michael Boswell, Samson Cheung, Bob Ciotti, Jill Dunbar, Earl Duque, Chris Gong, Bryan Green, Jennifer Goudey, T.R. Govindan, Chris Henze, Tom Hinke, Vee Hirsch, Amara de Keczer, Tony Keefe, Chris Kleiber, Bob Kufeld, Dochan Kwak, Eric Langhirt, Tammy Leal, Bonita McPherson, George Makatura, Yvonne Malloy, Donovan Mathias, Meyya Meyappan, Gina Morello, George Myers, Todd Novak, Jeff Onufer, Marcia Redmond, Scott Schreck, Ryan Spaulding, Roger Strawn, Valerie Steele, Ron Sykora, Jim Taft, Leigh Ann Tanner, Bill Thigpen, Richard Thompson, Ray Turney, Judith Utley, Rob Van Der Wijngaart, Cliff Williams, Michael Wilson, Mindy Wilson, Toshishige Yamada, Jerry Yan, Maurice Yarrow

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 Research Inc., Advanced Management Technology Inc., Computer Sciences Corp., Eloret, Raytheon Company, and SGI.

Subscribe to *Gridpoints*

Subscriptions to *Gridpoints* are complementary. Please e-mail your name, company name, title, address, and zip code to:
gridpoints@nas.nasa.gov

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

From The Division Chief

Many years after winning the Nobel Prize for co-inventing the transistor, Walter Brattain said, "The only regret I have about the transistor is its use in rock and roll." On page 10, NAS researcher Toshishige Yamada details his patent-pending research into atomic chain electronics – a new, nanotechnology-based approach to device miniaturization that will revolutionize the future of electronics. Reflecting on Brattain's comments, I wonder what Yamada will say about his invention 30 years from now?

Development of NASA's prototype Information Power Grid (IPG) is progressing steadily. In April, the IPG's first important milestone to demonstrate data transfer rates was achieved using data mining software. To meet the milestone, the IPG team retrieved data from four remotely connected sites simultaneously. The success of our data mining demonstration will allow the NAS team to move forward in expanding the IPG's infrastructure, some of which will be demonstrated this September. The IPG project is directly related to a part of the NAS Division's mission statement: "implementing and integrating new high performance computing technologies into useful production systems."

The NAS Facility is home to *Lomax*, a 512-processor SGI Origin 2800, currently the world's largest single-image system computer. *Lomax* supports the IPG and a variety of NASA research efforts. NAS scientists are using the software program CART3D on *Lomax* to develop "virtual flight" through computational fluid dynamics. CART3D generates grids to calculate flow fields around an air or space vehicle. Recently, NAS researchers studied four Space Shuttle nose configurations, and the results will be "flown" in Ames Research Center's Vertical Motion Simulator this fall.

The NAS Systems Division supports numerous NASA enterprises with high performance computing resources, and conducts research that will have far reaching impacts on the NASA mission, the scientific community, and eventually, on the public. For example, scientists from the National Renewable Energy Laboratory are working to improve wind turbine designs using wind tunnel testing and NAS computational modeling resources. Within the next 20 years, this clean source of power could provide more than 10 percent of the world's electricity.

In addition, the NASA Research and Education Network (NREN) is improving the Internet with the global adoption of the Multicast Protocol. Multicasting will allow a single data packet to be replicated and delivered to multiple users simultaneously, rather than each packet being delivered serially. Through improved Internet speed, we will all benefit from multicasting's applications.

As always, I welcome your feedback.

Bill Feiereisen

wfeiereisen@mail.arc.nasa.gov

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.

News From NAS

NAS Researchers Lecture on Carbon Nanotubes

NAS nanotechnology researcher Deepak Srivastava chaired the conference "Carbon Nanotubes: Advances in Cutting Edge Applications and Scalable Production," in Miami, Fla., April 10-11. The meeting emphasized advances in commercial applications and the prospects for scalable production of carbon nanotubes. Srivastava also gave an invited talk on plasticity and hydrogen diffusion on carbon and boron nitrate nanotubes, and Ames' experimental nanotechnology researcher Allen Cassel spoke about nanotube applications in nanoelectronic circuitry and combinatorial chemical approaches for tuning the production of nanotubes on different substrates. On his return trip, Srivastava also gave an invited talk, "Nanotechnology of Molecular Materials, Electronics and Machines: Carbon Nanotubes," at the Quantum Theory Project (QTP) seminar, at his alma mater, University of Florida, Gainesville, on April 12.

Chimera Grid Tools Version 1.3 Now Available

The Chimera Grid Tools (CGT) version 1.3, a package of programs and scripts used to generate overset grid systems for complex viscous computational fluid dynamics (CFD) problems, has been released. CGT's new tools include: surface and volume grid generation programs; grid manipulation, smoothing, and projection programs; a modular, general-purpose scripting system which can help build an entire grid system and produce Overflow flow solver inputs; and a graphical user interface known as Overgrid – used for generating grid systems. This version will significantly reduce CFD cycle times for a wide range of applications.

Version 1.3 is the final release of the software that has been funded by the DoD Common High Performance Computing Software Initiative CFD-4 project. Work is continuing toward a fully automated grid generation process with the incorporation of an AI-based surface and volume grid generation program, as well as the new, automated version of the overset connectivity software, Pegasus 5. CGT's authors are William Chan (NASA Ames), Stuart Rogers (NAS), Steve Nash (NAS), and Pieter Buning (NASA Langley). gp

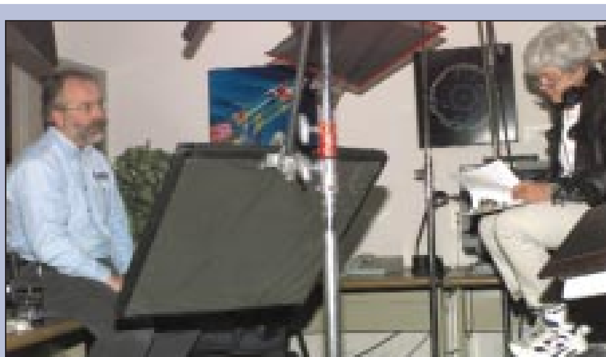

SGI Special: Lights, Camera, Action!

An SGI film crew producing a corporate documentary on the company's supercomputers and their uses employed the NAS Systems Division as a backdrop. Ames Center Director Henry McDonald and NAS Division Chief Bill Feiereisen were featured guests in the production. Feiereisen, above left in the visualization lab, discussed science conducted on SGI machines at NAS. Taping on June 13 lasted the entire day.
(Michael Boswell photo)

Ames Researchers Designing DNA Electronic Devices

IN the ongoing quest to solve challenges with creating nanoscale molecular electronic devices, NAS researchers are working to design nanoscale DNA diodes, transistors, and simple logic gates. The team is using a unique approach that provides a mechanism for self-assembly in future molecular-scale electronics, based on state-of-the-art nanotechnology and molecular electronics principles. The research, conducted by Jie Han in the NAS Systems Division, is funded by a two year, $80,000 grant from the Ames Director's Discretionary Fund.

Much of the recent progress in computer technology is based on creating smaller and smaller transistors for logic devices and integrated circuits. Rapid advances have been made in fabricating and demonstrating individual electronic wires and devices. However, these new technologies still face problems, including the requirement for devices to self-assemble.

Recent studies have shown that DNA has all the components needed to build an electronic device as well as the self-assembly characteristics needed to form complex electronic circuits. Even more exciting to the research team, DNA has the unique feature of self-replication. The work will focus on theoretical and computational proofs-of-concept, and experimental demonstration will be considered using nanotechnology and biology facilities at Ames.

Han is collaborating with the nanotube biosensor experimental team at Ames and Nanogen, Inc., based in San Diego. gp

Protein-folding Code Hums With NAS Optimization

A scientific code used to simulate protein folding sequences now performs 36 times faster than was previously possible, thanks to code optimizations made by NAS Systems Division researcher James Taft. The code, called

Corrections: In the Spring 2000 issue on page 10, the caption for Figure 3 should have read: "This image shows the spontaneous attachment of a protein to the modified messenger RNA that encodes it. This allows an evolutionary approach to the discovery of new functional proteins in the laboratory. Researchers at Harvard Medical School start by generating trillions of random proteins in a single tube. Then, they select rare proteins exhibiting a desired function. This process is a massively parallel analog computation of a problem not currently solvable by serial digital computational approaches. Current studies are directed toward the discovery of functional proteins in synthetically generated libraries of random-sequence proteins, rather than libraries isolated from biological systems. It is presumed that the proteins first utilized by early living organisms were recruited from such random-sequence libraries, and that this study will indicate what proportion of such proteins are functional. (Credit: Glen Cho, Anthony D. Keefe, Rihe Liu, David S. Wilson, and Jack W. Szostak) Additionally, Ling-Jen Chiang's credit was omitted from the graphic on page 21.

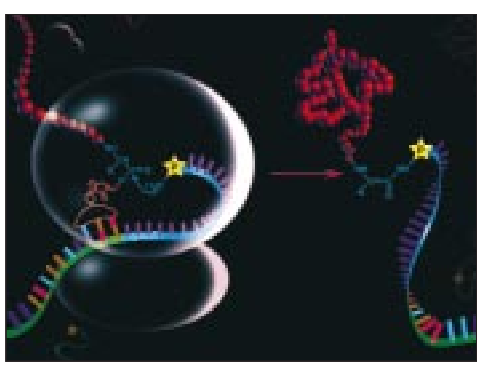

COSMOS, is a molecular dynamics model used in NASA's astrobiology research.

Taft modified the code to utilize 343 processors on *Lomax*, a parallel processor SGI Origin 2800 located at the NAS Facility. With 512 CPUs, *Lomax* is currently the largest single-image system in the world. Using a shared-memory parallel programming approach, Taft and COSMOS developer Andrew Pohorille, acting director of the NASA Computational Astrobiology Institute, have been rebuilding the code to run 30,000-atom problems on RISC microprocessor-based systems. Historically, many important molecular dynamics problems of this size have not scaled well on "clustered" parallel systems. *Lomax*, with very fast shared memory, is ideal for such computations.

"The results of this work can benefit the molecular modeling industry in general because the newly developed numerical techniques are applicable to many popular industry standard models used by pharmaceutical research companies in the U.S.," says Taft.

The two-month optimization effort focused on inserting the Ames-developed shared memory Multilevel Protocol (MLP) parallelization technique into the code. The MLP technique allowed dramatic improvements on the code's two most time-consuming routines. The overall speedup arises from scaling efficiencies in the MLP-based parallel algorithm, coupled with greater reuse of encached data. gp

NAS's Nanotechnology Team Members Recognized

Members of the Ames Integrated Product Team on Devices and Nanotechnology – representing the largest nanotechnology research efforts within the federal government – received a group achievement award for "excellence in research and outstanding contributions to the emerging field of nanotechnology" at the 2000 NASA Honor Awards ceremony on June 7.

Project manager Meyya Meyyappan, credits the NAS Systems Division as "the birthplace" of this work, originating with a small theory group doing computational nanotechnology at the NAS Facility. The team has expanded to include experimental research in the laboratory. Experiments on growing nanotubes for applications such as cancer diagnostics attract a steady stream of visitors from the high-tech arena, including former House science and technology committee chair Bob Walker, shown at right, who toured the "nanotube garden" on June 6.

Among those receiving the group achievement award were NAS researchers M.P. Anantrim, Bryan Biegel, T.R. Govindan, Jie Han, Cun-Zheng Ning, Deepak Srivastava, Alexei Svizhenko, and Toshishige Yamada. gp

Information & Technology Chairman Stephen Horn Tours NAS Facility

Rep. Stephen Horn (R-CA, Long Beach), chairman of the 106th Congress' Subcommittee on Government Management, Information, and Technology toured the NAS Facility on April 24 before hosting a hearing at Ames on "Emerging Technologies: Where is the Federal Government in the High Tech Curve?"

Senior research scientist Chris Henze (right) briefs Rep. Stephen Horn on the use of NAS's Responsive Workbench in the NAS Visualization Lab. (Nicholas A. Veronico)

Division Chief Bill Feiereisen briefed the congressman on NAS's advances in high performance computing, nanotechnology, and the Information Power Grid. Horn then toured the Visualization Lab and was given a Virtual Mechanosysthesis demonstration by Chris Henze and Bryan Green of the NAS data analysis group. A carbon nanotube and a graphite T-junction were projected on the Responsive Workbench and Horn was able to interact with individual carbon atoms. At the conclusion of his tour, Rep. Horn said, "This amazing technology presents exciting possibilities for science and technology in the 21st century."

— *Nicholas A. Veronico*

Fact Finding Tour of NAS And Ames

Former Pennsylvania Congressman Robert Walker, center, and Ames Deputy Director of Information Systems Paul Kutler, left, toured the Visualization Lab and were briefed on the capabilities of the 512 processor SGI Origin 2800, *Lomax*, by Bill Thigpen, NAS's Engineering Branch chief. Walker, former chair of the House Science Committee, was on a fact finding mission at Ames on June 6. The trio stand inside *Lomax's* processor banks in the computer room at NAS. — *Nicholas A. Veronico*

CART3D: A New Generation of Grids

3-D grid software automatically generates vehicle simulations from complex geometries.

Researchers Donovan Mathias and Jeff Onufer of NAS's computational physics and simulation group, tested four different nose designs for the Space Shuttle on the 512-processor SGI Origin 2800, *Lomax*, during the weekend of March 25-26. "The goal of this project was not to find the best-suited nose design for the orbiter, but to use computational fluid dynamics to support virtual flight. We are developing a process to bring virtual flight into the early design phase, so the vehicle can be 'flown' before a single piece of metal is cut," says Mathias who believes testing a design early in its development will save both time and money.

For the team's work on *Lomax*, they employed a beta version of the CART3D (Three-Dimensional Cartesian Simulation System for Complex Geometries) software package to generate the grids necessary for calculating flow fields surrounding the orbiter. Results from the *Lomax* tests will be used in the Vertical Motion Simulator (VMS, see sidebar, *NASA's Full Motion Simulator*, page 6) at NASA Ames in mid-September. The VMS will "virtually" fly the Space Shuttle, evaluating flight performance of the new orbiter designs.

During 21 hours of computations on *Lomax*, the team tested four configurations – three new re-designs (created by Dave Kinney from the Systems Analysis Branch at NASA Ames) plus the Shuttle's original nose configuration. "We were able to study quite a number of the orbiter's components during our weekend on *Lomax*," says Onufer. Each orbiter redesign took 60 simulations, requiring eight processors per simulation to study the control surfaces, including the elevons (airfoils used to control the movement of the shuttle), rudder, and speed brakes. The team also examined flight conditions of the vehicle: Mach number, angle of attack, and angle of sideslip. This research was funded by the Information Technology Program.

Once the orbiter grids were generated, they were uploaded to *Lomax* and used to calculate flow solutions using "Tiger,"

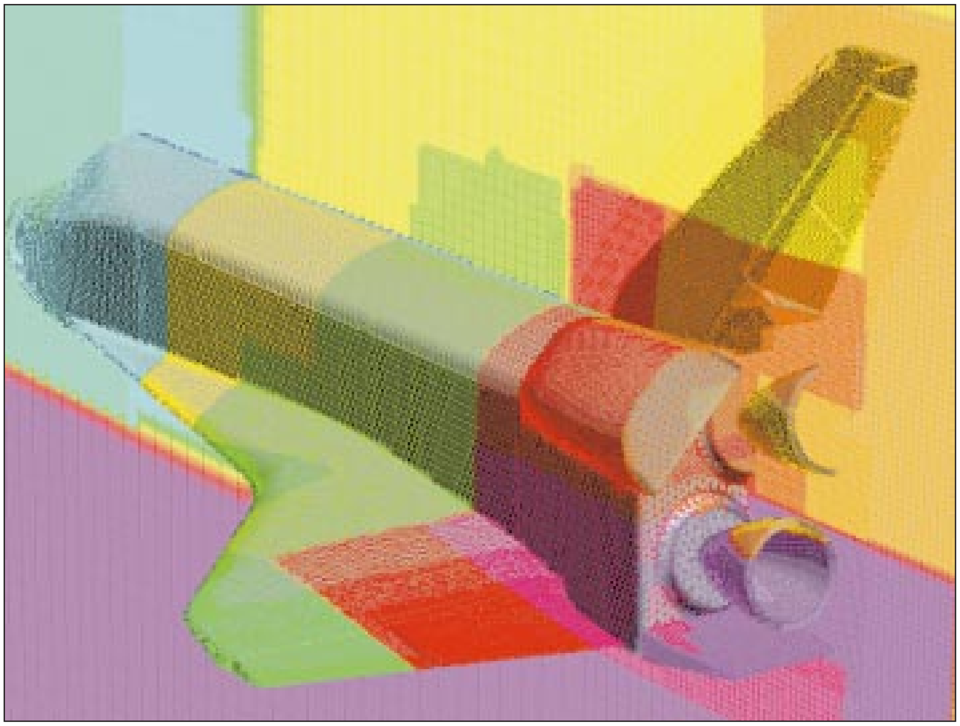

Mesh partitions flow over a Space Shuttle orbiter, with each color representing a set of equations assigned to different CPUs. Sixteen CPUs were used for this calculation; thus there are 16 divisions. Partitioning is done automatically on any number of processors. (Aftosmis) the flow solver module for CART3D. The Tiger code, developed by John Melton of Ames' Aeronautical Projects and Programs Office, models the inviscid flow field surrounding a vehicle such as the shuttle orbiter.

Enabling Virtual Flight

Aiming to test the shuttle designs in the VMS, researchers will merge their flow solutions calculated on *Lomax* with VMS's existing mathematical model, which uses a set of equations to describe the vehicle's performance. Merging the flow data with the simulation model, a pilot can evaluate the performance of the shuttle orbiter redesigns with high-fidelity simulations. "Ideally, we would like real astronauts to fly the simulation to provide the most useful feedback on the flight performance of the vehicle," says Mathias. Currently, NAS research scientist Michael Aftosmis is working on a scalable flow solver to replace the Tiger code. Aftosmis' anticipated flow solver will drastically reduce the time necessary to compute flow fields by distributing the mesh generated onto different CPUs. Mesh decomposition

Grid Generation Using CART3D

NAS research scientist Michael Aftosmis, John Melton from Ames' aeronautical projects and programs office, and Marsha Berger, co-director of the Courant Institute at New York University, released CART3D in 1997, after five years of development. CART3D is a preliminary design analysis package created to solve the Euler Equations of Motion with Cartesian grids around complex vehicles. Aftosmis' work is funded by NASA's High Performance Computing and Communications (HPCC) Program.

CART3D can import geometry from many different sources, including most major computer-aided design (CAD) packages, and computes a complete flow simulation around any complex vehicle. The software takes the surface definition of a vehicle and generates a mesh, or set of grids, necessary to compute the air flow around it using a flow solver. "The way CART3D sees geometry is unique – parts of a whole can be moved, and the entire configuration does not have to be regenerated. The software uses a component-based approach enabling the automatic readjustment of the grids, saving a ton of time," Aftosmis says.

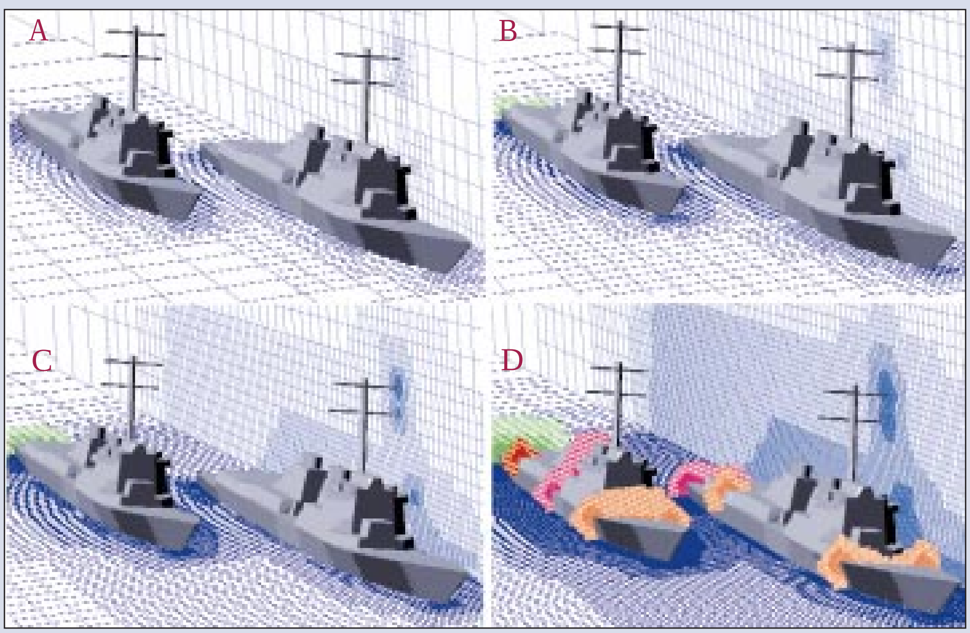

The efficiency of CART3D has increased since it was first released. "A majority of a researcher's time was being consumed in converting CAD geometry into something that could be used for mesh generation. CART3D has reduced the execution time of this step, and the ultimate goal is to put the whole thing in a box, eliminating the human element," says Aftosmis. "CART3D is a very robust software package, and can be quickly applied to almost any air or space vehicle," adds researcher Donovan Mathias. "Grids are essentially generated in a matter of minutes on just about any UNIX or PC platform using CART3D."

The sequence of grids (A through D, below), 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. This method utilizes a sequence of coarse meshes to reduce the calculation time of the final solution of a detailed, highly refined mesh. The amount of computational work performed is proportional to the complexity of the mesh – less complex grids enable rapid computations using solution estimates. This multigrid technique uses a hierarchical approach – the coarse solution estimates, (figures A through C), are used to "drive" the complete solution on the fine mesh (Figure D). The grids are automatically produced and distributed over *Lomax's* processors. In this example, the finest mesh has 4.5 million grid points while the coarsest has only 4,500. (Aftosmis and Mathias)

NASA's Full Motion Simulator

Space Shuttle configurations calculated on the NAS Facility's *Lomax* high performance computer will take flight in Ames Research Center's Vertical Motion Simulator in September. Each of the four newly computed Space Shuttle designs will be flown in the simulator, which is capable of replicating vertical, lateral, and longitudinal movements as well as roll, pitch, and yaw.

The Vertical Motion Simulator's size and range of motion are key to producing high fidelity flight simulations. Housed at the Vertical Motion Flight Simulation Laboratory in a ten-story building, the simulator's vertical platform is 70 feet wide and can be elevated nearly 75 feet. On top of the vertical platform sits the interchangeable simulator cab where pilots can explore a craft's flight envelope.

The lab houses five different cabs that can be built up into a wide variety of helicopter, single-seat fighter jet, tilt rotor and transport aircraft, and the Space Shuttle Orbiter. Each cab can be configured to the research needs of various test projects by changing the seats, flight controls, and cockpit instrumentation. An image presentation system gives the pilot a virtual out-the-window view of the world.

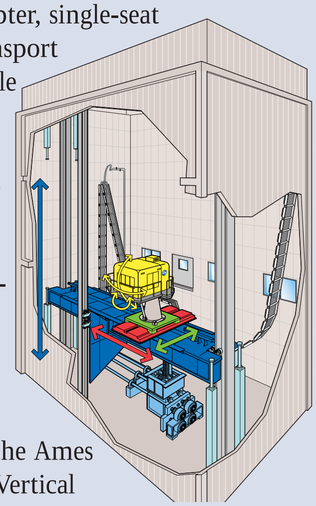
VMS/Neesham Wallace, Jennifer Goudey

For further information on the Ames Aviation Systems Division's Vertical Motion Flight Simulation Laboratory, visit: http://www.simlabs.arc.nasa.gov/vms/vms.html.

enhances the solver's scalability on parallel computers like *Lomax* by explicitly assigning different parts of the problem to the machine's processors. An application using the new solver with the same number of CPUs will work five to 30 times faster than the "Tiger" code. "We are really looking forward to putting this new solver to work – it will be a great deal faster than Tiger," says Onufer. Aftosmis plans to release the new code to a select group of government users in mid-June and apply it to high-speed applications: "We want to stretch the limits of the new solver by applying it to complex problems in all three flight regimes; subsonic, transonic, and supersonic."

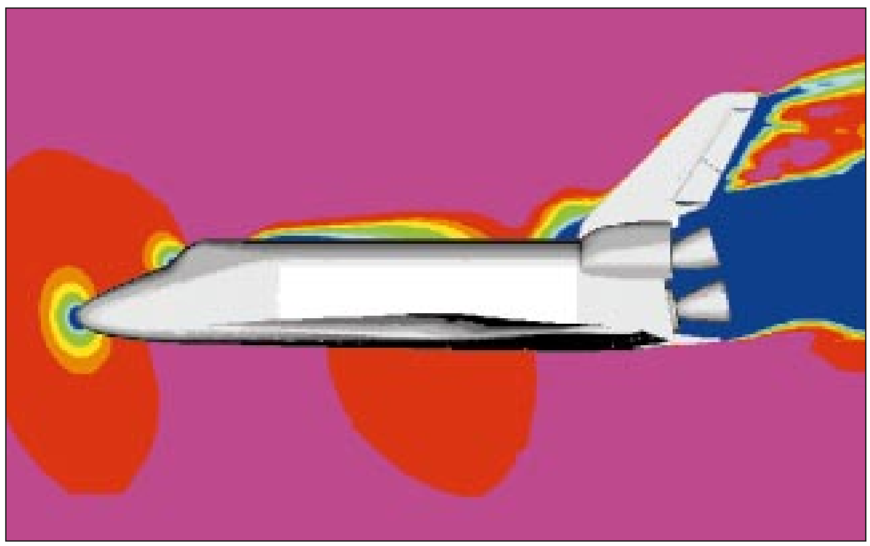

Above: The velocity component of the baseline shuttle orbiter as computed on *Lomax*. The angle of side-slip is zero degrees, and angle of attack is 12 degrees. Pink areas represent free stream air at MACH 0.9, while the darker colors correspond to a lower air-stream velocity. The blue area behind the shuttle represents low velocities. (Onufer, Mathias)

CART3D's Future

More than 80 institutions, including the U.S. government, industry, and academia currently use CART3D. Mathias, Onufer, and Aftosmis plan to continue their work on improving flow calculations for flight simulations in the future. "I think virtual flight is a great application of simulation software. The application stresses CART3D because it requires a lot of throughput, but is very doable with machines like *Lomax*," says Aftosmis. gp

—Holly A. Amundson

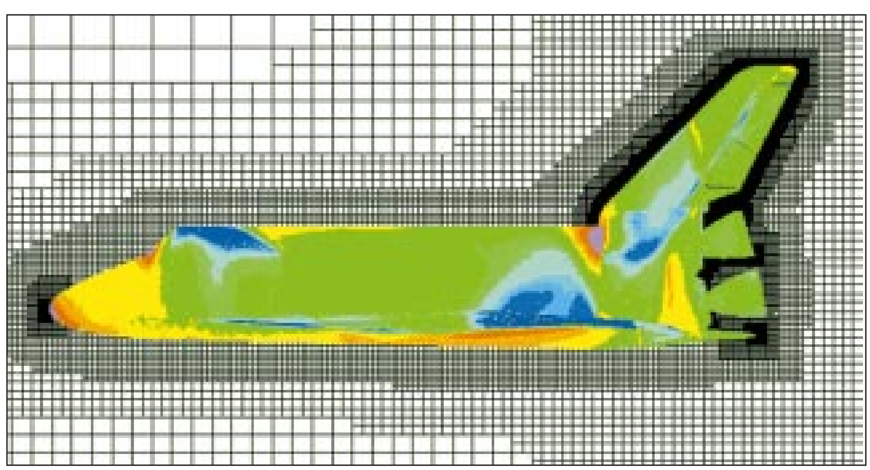

Right: The pressure coefficient used to compute lift and drag is depicted on the baseline shuttle orbiter. Red represents high pressure areas, while the darker colors denote low pressure. (Onufer, Mathias)

NASA Research and Education Network Drives Global Deployment of Multicast Protocol

Multicast will speed up the Internet by enabling a single data packet to be distributed to multiple Internet users simultaneously.

By Hugh LaMaster and Kenneth Freeman

Since 1998, the NASA Research and Education Network (NREN) – a component of the NAS Systems Division, has been a leader in the global deployment of new multicast routing and forwarding protocols used for sending information across the Internet. Protocol Independent Multicast–Sparse Mode (PIM-SM), Multicast Source Discovery Protocol (MSDP), and Multicast Border Gateway Protocol (MBGP) are now being adopted worldwide as the standard protocols for multicast data transfer. Multicast will improve the Internet's efficiency, and expand networking technology to allow scientists and researchers to collaborate while conducting experiments from remote locations. In the future, commercialization of multicast technology will benefit society as a whole, enabling advances in aerospace research, astrobiology, and other Earth sciences, as well as telemedicine.

How Does Multicasting Work?

Within the Internet, packets (pieces of data) are usually sent one at a time. This method, called "unicasting," is a simple mechanism for one-to-one communication. However, for one-to-many communication, unicasting can create extreme network congestion. With multicast, packets are simultaneously transmitted to multiple sites at the same time. As a result, multicasting conserves bandwidth, reducing the inherent bottlenecks created by one-to-many transmissions. During multicast, only one copy of the data is transmitted while network routers replicate the packets when needed. As a result, on both the sending and receiving side of the transmission, network bandwidth is conserved.

Multicasting is similar to broadcasting radio signals; like a radio signal, multicast packets originate from a single site. Anyone with the correct equipment can receive the signal,

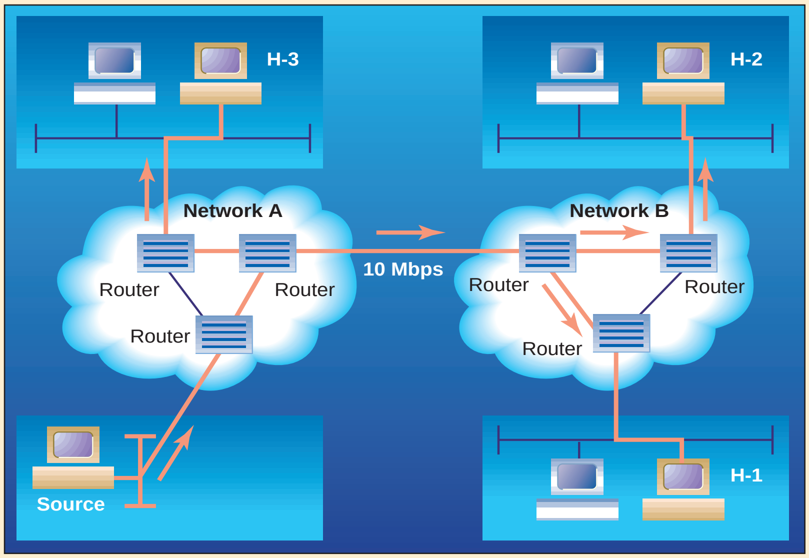

Figure 1. If the source computer is multicasting data at 10 Megabits per second (Mbps), 10 Mbps will be transmitted between Network A and Network B, reaching Host H-1 and Host H-2 simultaneously. Without multicast, 20 Mbps of data would be required to transmit data to Host H-1 and Host H-2, doubling the transmitted data time. (LaMaster and Freeman)

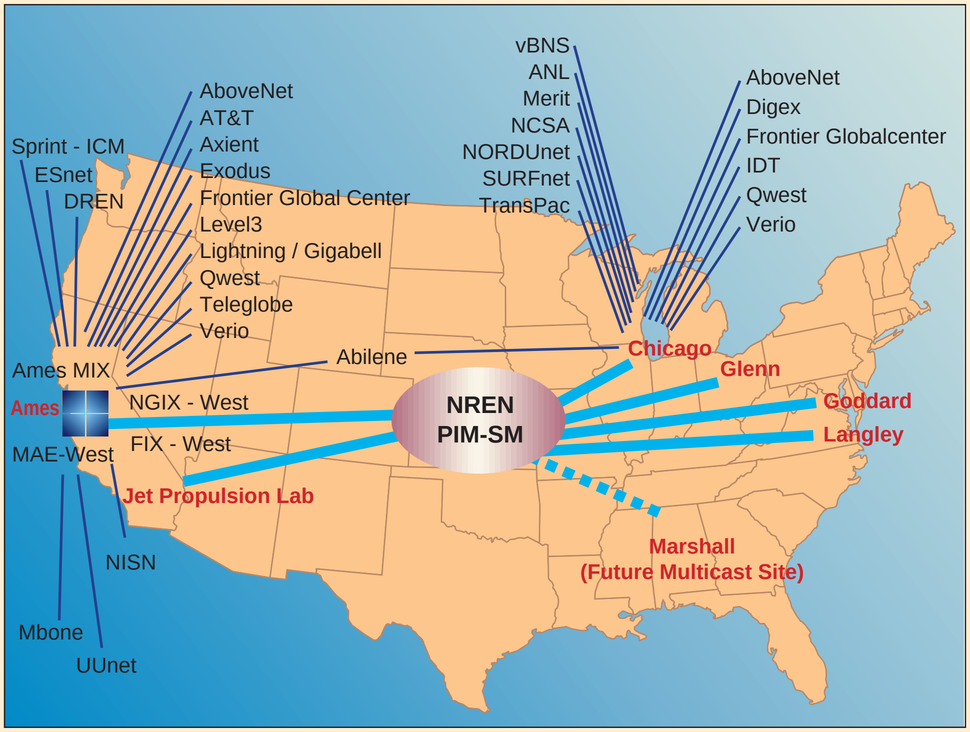

Figure 2. As of June 2000, NREN's multicast architecture has been adopted by numerous networks across the nation. Marshall Space Flight Center has been designated as a future multicast site.

while radio signals pass by anyone who does not want to receive the signal, or does not have the right equipment. Similarly, multicast is "broadcast" to the network but is only received by those interested in the data.

Evolution of Global IP Multicast

With the explosive growth of the Internet, networks are struggling to keep up with the increasing demand of millions of new users. In fact, the World Wide Web is often referred to as the "World Wide Wait." To reduce the amount of repetitive data transmissions, global IP (Internet Protocol) multicast technology allows a single packet stream to be distributed across the Internet simultaneously to any number of recipients. Multicast is often the only way to simultaneously distribute real-time data to numerous recipients efficiently. It also allows a data stream to be replicated at intermediate routers so that any link only receives a single copy, regardless of how many downstream recipients have requested the information.

Ames Research Center has been a focal point of IP multicast deployment since the early days of MBone, the multicast backbone. Global IP multicast routing, originally done via the distance vector multicast routing protocol (DVMRP), was the basis for the MBone. One of the first successes in using multicast was NASA's Space Bridge to Russia in 1996, where multiple remote sites were fed information simultaneously.

More recently, NREN has deployed a number of new networking protocols. Collaborating with the Ames Gateway Facility (Ames Research Center's data, video, and telecommunications relay point), Cisco Systems, and Sprint's research network (Sprint-ICM), the first Multicast Internet eXchange (MIX) was created; the first multi-network deployment of the Multicast Border Gateway Protocol (MBGP – a routing protocol that replaces DVMRP for inter-domain routing) was achieved; and the first multi-domain global deployment of the Protocol Independent Multicast (PIM – a forwarding protocol) was accomplished.

In July 1998, NREN began deploying the Multicast Border Gateway Protocol (MBGP). Initial sites included NREN, Sprint-ICM, and a transit-AS (Autonomous System). By fall 1998, connections to the Department of Energy's Energy Sciences Network (ESNET), the Defense Research and Engineering Network (DREN), as well as Internet service providers Verio, Exodus, and AboveNet, had been added. During this time, NREN began "peering" with vBNS (National Science Foundation and MCI's very highspeed Backbone Network Service) in Chicago. When Internet2's Abilene Research Network joined, all of the federal government's Next Generation Internet networks had adopted MBGP.

Multicast Source Discovery Protocol Deployment

Throughout the Mbone and early MBGP deployment, dense-mode protocols were used that can sometimes flood unnecessary data across wide area network (WAN) links. Using dense-mode, there was no way that very high bandwidth sources exceeding 10 Megabits per second (Mbps) could be supported across typical WAN links.

Beginning in the fall of 1998, NREN initiated testing of a new protocol, Multicast Source Discovery Protocol (MSDP), allowing the use of sparse-mode protocols globally. Sparse-mode enables multicast to be used without specific bandwidth limitations, and supports very high bandwidth data streams.

In spring 1999, driven by the requirement to support the Virtual Collaborative Clinic – an Internet-based telemedicine demonstration – NREN coordinated the deployment of MSDP and PIM-SM with the Internet2/Abilene network, the Corporation for Education Network Initiatives in California's CalREN-2 network, as well as Stanford University and the University of California at Santa Cruz. The deployment of PIM-SM and MSDP enabled data transfer rates of at least 24 Mbps to all sites (50 Mbps to some sites).

During summer 1999, the deployment of MSDP and PIM-SM continued throughout Internet2's customer networks, and among commercial Internet service providers. NREN provided the catalyst for this deployment by using PIM-SM and MSDP at the Ames MIX, as well as its peerings at Chicago. By January 2000, the conversion from dense-mode to MSDP/ PIM-SM was almost complete throughout the United States (see Figure 2, opposite).

Multicast Routing Protocols

Global multicast routing requires routers to replicate multicast packets and forward copies to each interface where a downstream receiver has requested a copy. This process is known as multicast forwarding, and it requires the creation of a multicast forwarding tree for each multicast group.

In addition to the creation of a forwarding tree, a routing protocol is required. In the usual configuration, the AS (Autonomous Systems) Interior Gateway Protocol (IGP) is the source of unicast routes that can also be used for multicast. A common IGP is OSPF (Open Shortest Path First). Among autonomous systems, an Exterior Gateway Protocol (EGP) is used to exchange routes. Today, the universal EGP is Border Gateway Protocol-4 (BGP-4). Multicast Border Gateway Protocol (MBGP) is an extension to BGP-4 that allows for an exchange of additional routes specifically for multicast. This is required because the data paths for unicast and multicast can differ.

In order for multiple autonomous systems to know about each other's multicast sources, a helper protocol is required. Today, this protocol is Multicast Source Discovery Protocol (MSDP). MSDP floods information about multicast sources among all the Rendezvous Points (RPs) in multiple domains, allowing receivers in one autonomous system to receive data from a different autonomous system (see Figure 3).

Protocol Independent Multicast-Sparse Mode

Sparse protocols use the explicit join of a multicast group by a host to trigger the creation of the multicast forwarding tree. Initially, the first-hop router joins the shared-tree from the rendezvous point. For high-bandwidth sources, it will then switch to a shortest-path tree from the first-hop router back to the source. PIM-SM does this through the use of Rendezvous Points that coordinate the creation of the multicast distribution trees throughout the PIM-SM domain (usually an AS). Within a defined network area, the RP creates a shared tree for every multicast group, and it coordinates the creation of a shortest path tree from a source to all recipients for high-bandwidth sources.

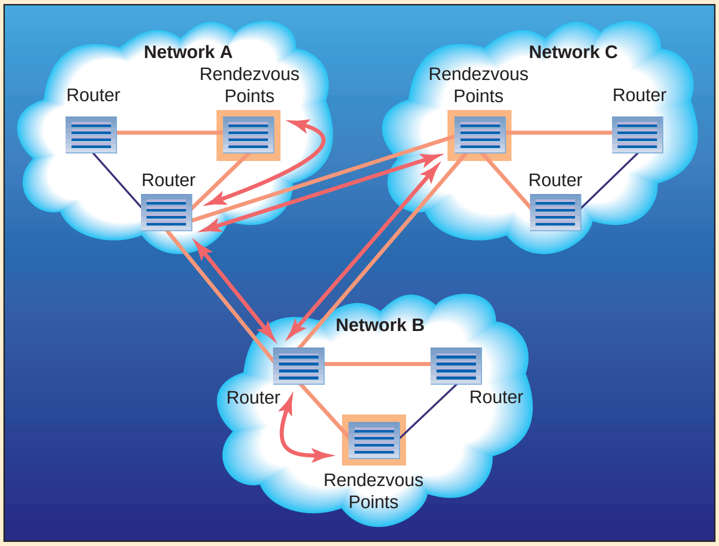

Figure 3. With inter-domain multicast routing, all routers with external MBGP connectivity have parallel MSDP connectivity. All routers with external MBGP and MSDP connectivity have meshed internal MBGP and MSDP connectivity. In addition, the RP must be part of the MSDP mesh. Each network represents an Autonomous System (AS). Moreover, the red arrows depict links with MBGP and MSDP connectivity.

The Future of Multicast

MSDP has always been regarded as an interim protocol for several reasons. It incorporates a feature for backwards-compatibility with the old Mbone, and it floods source information. This flooding, while vastly superior to the dense-mode data flooding it has replaced, still will not scale to an Internet with very large numbers of active sources. For some time now, development of global shared-tree protocols has been undertaken. Today this has to be considered a difficult and as yet unsolved problem.

In addition, with the commercialization of the Internet, a new constituency of commercial multicast Internet broadcasters is appearing. These broadcasters would like to support broadcast-only functionality, rather than the many-to-many multicast model. They also would like a single broadcaster to be able to deliver very large numbers of channels, which would create a multicast address allocation problem. A segment of the multicast address space has been allocated for this so-called source-specific multicast, which will allow both of these problems to be solved while maintaining most of the multicast address space for many-to-many communication. Source-specific multicast will require a simple modification of PIM-SM behavior to accommodate the different functionality. MSDP will not be required for source-specific multicast because the sources will typically be announced via the Web.

Today, IP Multicast is undergoing rapid development and deployment in the commercial world, and will be a key technology in providing efficient sharing of real-time data among large numbers of systems. gp

The Smallest Nanoelectronics: Atomic Devices with Precise Structures

Transistor miniaturization is reaching its limit. Atomic scale transistors are being developed to enable future microelectronics.

By Toshishige Yamada

Since its invention in 1948, the transistor has revolutionized everyday life – transistor radios and TVs appeared in the early 1960s, personal computers came into widespread use in the mid-1980s, and cellular phones, laptops, and palm-sized organizers dominated the 1990s. The electronics revolution is based upon transistor miniaturization; smaller transistors are faster, and denser circuitry has more functionality. Transistors in current generation chips are 0.18 µm (micron) or 180 nanometers in size, and the electronics industry has completed development of 0.13 µm transistors that will enter production within the next few years. Industry researchers are now working to reduce transistor size to 0.1 µm – a thousandth of the width of a human hair. However, studies indicate that the miniaturization of silicon transistors will soon reach its limit.

For further progress in microelectronics, scientists have turned to nanotechnology to advance the science. Rather than continuing to miniaturize transistors to a point where they become unreliable, nanotechnology offers the new approach of building devices on the atomic scale (see *Transistors: Electronic Workhorses*, opposite). One vision for the next generation of miniature electrical devices is atomic chain electronics, where devices are composed of atoms aligned on top of a substrate surface in a regular pattern. The Atomic Chain Electronics Project (ACEP), part of the semiconductor device modeling and nanotechnology group

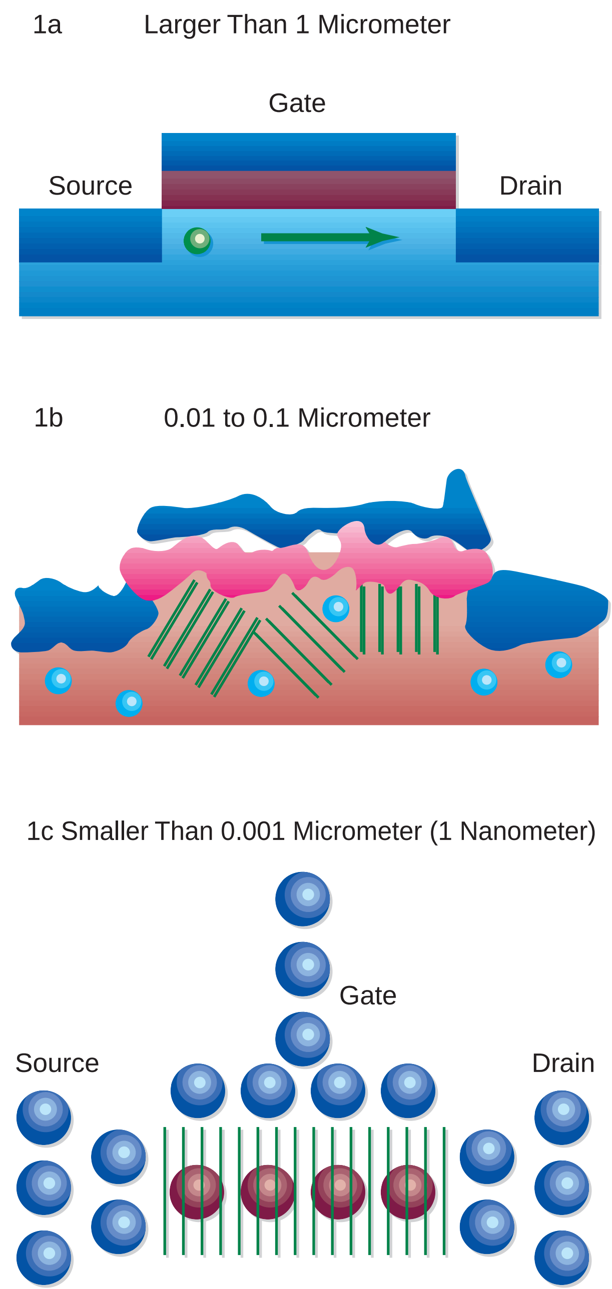

Figure 1: The evolution of transistor miniaturization has taken electronic devices from the macroscopic scale (1a), greater than 1 µm (micron). An electron particle (green) sees a jelly area (blue) formed by 10,000 dopant atoms, which is averaged out to be the same in devices, resulting in the same switch-on voltage. In transistors miniaturized to the mesoscopic scale (1b), 0.1 to 0.01 µm, an electron wave (green) can tell precisely where the dopant atoms are located, and this leads to different switch-on voltages. To solve this problem, researchers at the NAS Facility have patented an atomic device concept using atomic chains (1c) rather than further the miniaturization of conventional transistors.

integrated product team within the NAS Systems Division, has been developing the theory for understanding atomic chain devices. In this type of research, collaboration with

Transistors: Electronic Workhorses

Millions of tiny transistors are at the heart of today's electronics world. The main function of a transistor is to amplify a small input signal into a large output signal as shown in the diagram (a). The most common transistor used in today's electronics is a field-effect transistor (FET). To understand the operation of an FET, the popular analogy is that of a flexible hose system with weight on it. When the weight is small, water flow is not hindered, as shown at right. As the weight is increased, the hose becomes pinched, resulting in less water flow. In this analogy, the weight corresponds to the input signal and the water passing through the hose is the output.

In a field-effect transistor, the input is a voltage applied to the gate electrode, and the output is a current flowing between the source and the drain. When the gate voltage is low, electrons flow freely from the source to the drain. When the gate voltage is high, electrons are repelled under the gate and fewer electrons flow. This is known as the field effect.

In the analogy of water flow through a hose (b), the constriction caused by the weight depends on the material of the hose. For example, if the hose is a steel pipe, then the weight (input) has little effect on flow (output). In the same way, if the substrate of an FET is a metal, the gate voltage has little effect on current flow. Therefore, we need a material corresponding to a flexible hose – a semiconductor.

What FET property corresponds to the flexibility of the hose? That is doping, or the intentional inclusion of impurity atoms in the semiconductor substrate. By manipulating doping, a semiconductor can be changed from something almost an insulator (softest) to something almost a metal (hardest).

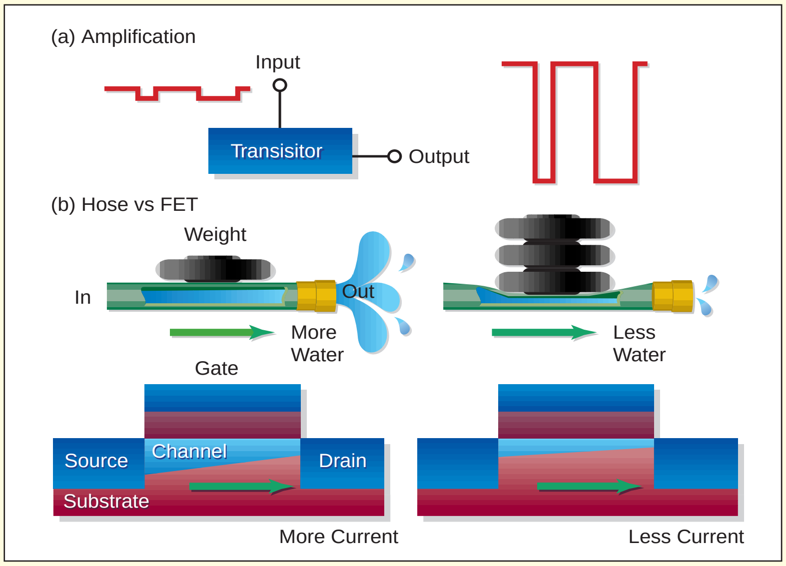

the experimentalists is also essential for success. ACEP members are collaborating with Dongmin Chen of the Nanophysics Group at Boston's Rowland Institute, who characterizes electronic properties of atomic structures using a Scanning Tunneling Microscope (STM). ACEP researchers are also collaborating with Richard A. Kiehl at the University of Minnesota's electrical engineering department, who studies the autoformation of atomic structures. Autoformation may lead to a breakthrough for mass production of atomic-scale devices.

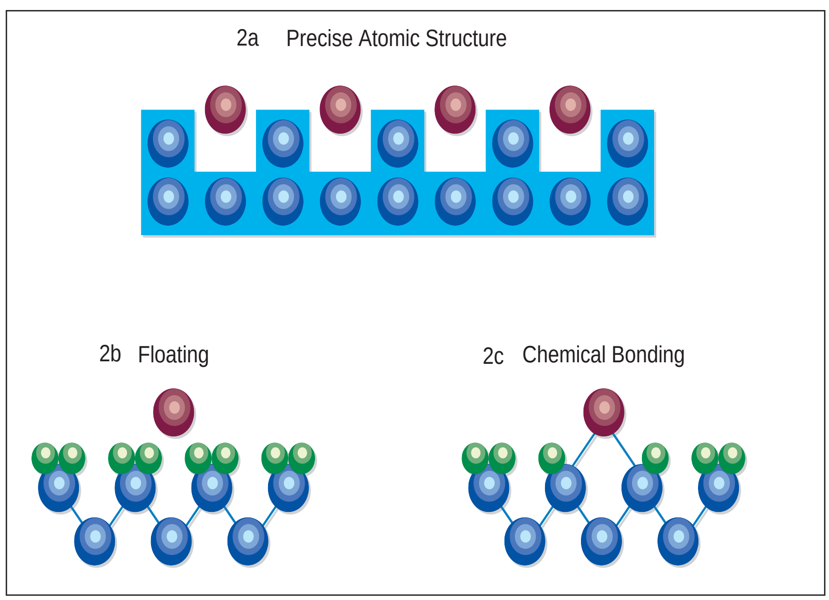

Figure 2: Rendering (2a) represents an atomically precise structure using the substrate atomic lattice (blue atoms as a template for atomic chain electronics (red atoms). Two confinement schemes for topmost atoms on the substrate are (2b), where atoms are floating, and (2c), where the atoms are firmly bonded to the substrate. Green dots represent hydrogen atoms, and they work as an electronically insulating layer with an atomic thickness.

Doping and Gain in Transistors

A critical element in transistor functioning is the use of dopants – impurities added to silicon to manipulate and control its electronic properties. Typically, there are several

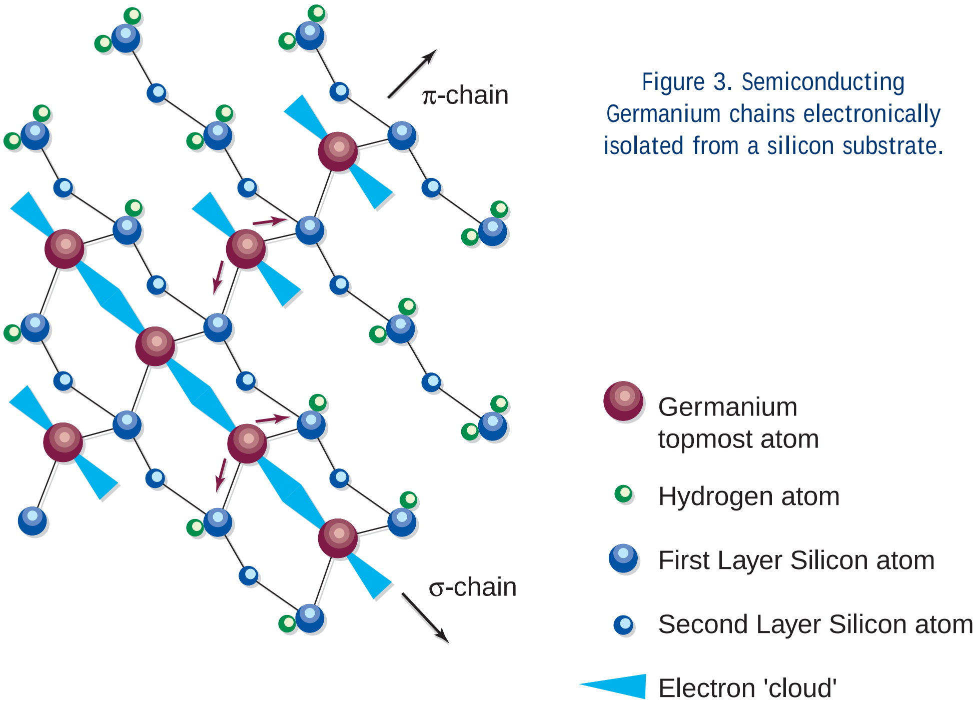

Figure 3. Semiconducting Germanium chains electronically isolated from a silicon substrate.

thousand dopant atoms in a transistor and these appear as a dopant "jelly" as in Figure 1(a) – on page 10. When using large numbers of dopant atoms in the construction of an electronic device, the precise location of these atoms in the transistor is not very important to the functioning of the transistor. However, when the transistor size is reduced to less than 0.1 µm as in Figure 1(b), the number of dopant atoms is small – less than 100. At 0.1 µm or less, each dopant's position affects how electrical current flows through the transistor. Present-day manufacturing technology cannot precisely control the location of so few dopant atoms, and this causes small variations in how the transistor functions. This is fatal when millions or billions of transistors are integrated on a computer chip, because these slight variations can cascade from one device to the next and eventually lead to malfunction. The solution to this variation problem devised by ACEP calls for all device structures to be built using atomic chains laid out in a regular, precise pattern by anchoring atoms to a substrate as in Figure 1(c).

A second critical element in transistor function is that the output of the device should be a magnification of its input – this is called *gain*. To create atomic chain transistors that produce gain, ACEP researchers have found that the phenomenon known as field effect (see *Transistors: Electronic Workhorses*, page 11) can be adopted as an operating principle. To build a field-effect transistor at the atomic scale (similar in construction to those of larger scales), scientists need to devise chains with both semiconductor and metal properties. First, the ACEP team tackled a simple problem theoretically: If researchers can arrange silicon atoms along a line floating in air, is this chain semiconducting? The answer is surprising: although bulk or thin film silicon is semiconducting, the silicon chains are always metallic. Fortunately, ACEP's research shows that a magnesium chain is semiconducting, even though ironically, bulk magnesium is metallic.

Substrate Determines Chain Properties

Of course, atoms cannot be floated in air, but must be placed on top of a substrate. Surface atoms of the substrate attract atoms in the chain and hold them at fixed positions. The substrate has atomic scale corrugations on the surface and these may be used to create a precise pattern as shown in Figure 2(a) – on page 11. However, ACEP members have found that this force is too weak to secure atoms reliably, see Figure 2(b). ACEP scientists concluded that a chemical bonding scheme to secure atoms to the substrate as in Figure 2(c) was the reliable approach to follow. Research soon showed that the chain properties were strongly influenced by the substrate material and surface orientation

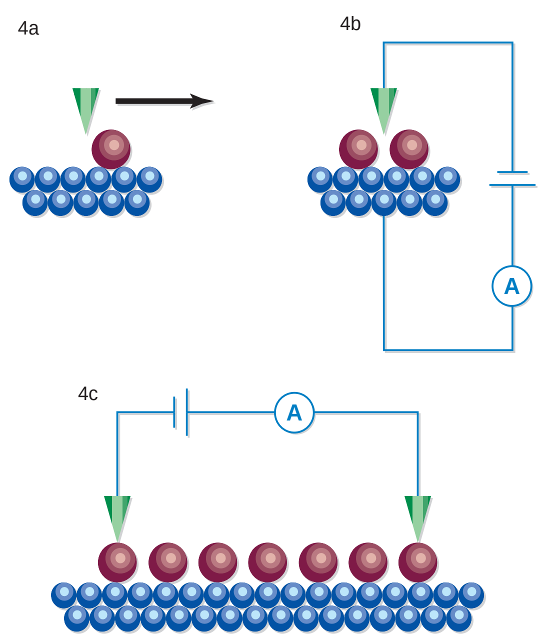

Figure 4: Atom manipulation and electronic characterization of small atomic structures are accomplished using a Scanning Tunneling Microscope (STM) tip (green). Researchers can push or pick up an atom (red) on the substrate (blue) using an STM tip as in (a). A created small atomic structure (red) can be electronically characterized by measuring current for applied voltage either in (b) a vertical geometry, or (c) a horizontal geometry. The latter needs a dual-tip STM and is technically quite challenging.

Figure 5: Carbon nanotube field effect transistors (FETs) experimentally built by Delft and IBM, with function similar to that predicted for an atomic chain FET in Fig. 1(c), have demonstrated the transistor action in the molecular/atomic dimension successfully.

Figure 5

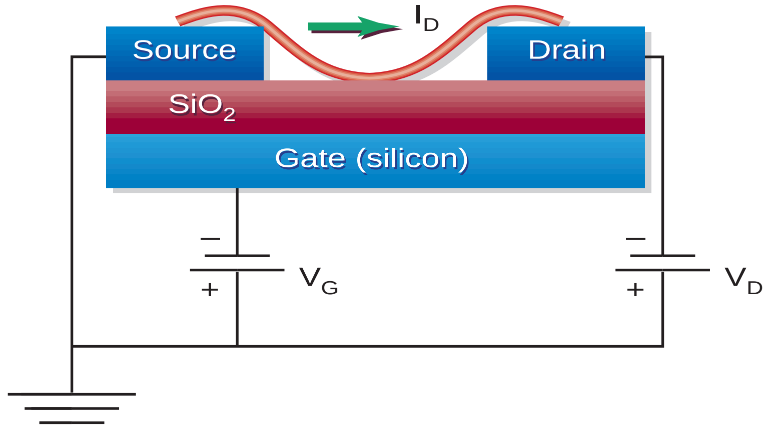

when chemical bonding is relied upon. In fact, the same atomic chain can be either metallic or semiconducting depending on the configuration of the substrate. For example, silicon chains with two chemical bonds to the substrate atoms are semiconducting – otherwise silicon chains are metallic. Aluminum chains with one chemical bond are semiconducting – otherwise, aluminum chains are metallic. A further concern was that electrons travelling along an atomic chain may detour into the substrate through chemical bonds and possibly exit via a neighboring chain – resulting in the short-circuit of two atomic chains. This short circuit is fatal for electronic applications and must be avoided. ACEP scientists have clarified the conditions for which the substrate surface acts as an insulator in the chemical bonding scheme. According to this, silicon and germanium crystals offer good insulating substrates.

Putting It All Together

The ACEP team has achieved the first step towards the development of atomic chain electronic devices. On top of a silicon substrate, germanium atoms can be attached with two chemical bonds as shown in Figure 3 – on page 12. This structure is a semiconducting chain on an insulating substrate. Nearby chemical bonds on the substrate are saturated with hydrogen atoms so that they are electronically inactive. This has resulted in the development of an essential component for an FET on the atomic scale.

These structures could be experimentally fabricated and characterized using an atom manipulation technology with a Scanning Tunneling Microscope (STM) as in Figure 4 – on page 12, or some other related techniques. Quite recently, the field effect in the molecular/atomic scale has been demonstrated in carbon nanotube FET experiments in Figure 5 by Cees Dekker's group at Delft, and Phaedon Avouris' group at IBM. This supports the ACEP's predictions for atomic devices, and the basis for atomic chain electronics is now being established. gp

Visit these sites for additional sources of information about nanotechnology:

Integrated Product Team on Devices, Nanotechnology
www.ipt.arc.nasa.gov
The Integrated Product Team on Devices and Nanotechnology's home page offers a gallery of nanotechnology experiments, NASA Ames' nanotechnology capabilities, and a list of projects and publications.

National Nanotechnology Initiative Home Page
www.nano.gov
Read the full text *National Nanotechnology Initiative: Leading to the Next Industrial Revolution*, supplement to President Clinton's FY 2001 budget.

Interagency Working Group on Nanoscience, Engineering and Technology
www.whitehouse.gov/WH/EOP/OSTP/NSTC/html/iwgn/IWGN.Research.Directions/toc.html
This document lays the foundation for the future of nanotechnology research and development, and how research efforts should be coordinated within the scientific community.

Acknowledgements

The author acknowledges fruitful discussions with and/or continuous support for this project by Meyya Meyyappan, T. R. Govindan, C. W. Bauschlicher, Jr., H. Partridge, and B. A. Biegel. Careful reading by H. Yamada is gratefully acknowledged.

Toshishige Yamada is a Senior Research Scientist with Computer Sciences Corp. at NASA Ames Research Center. He obtained a Ph.D. in electrical engineering from Arizona State University. He held positions at NEC Microelectronics Research Labs., Japan, and Stanford University. NASA filed a patent application on his atomic chain electronics innovation, and the work was presented at an invited talk in the 196th Electrochemical Society Meeting last fall.

CFD Supports Wind Turbine Research

NAS provides the National Renewable Energy Laboratory with computational resources to help improve next-generation wind turbine designs.

Wind-generated electricity is the fastest growing, most cost-competitive renewable energy resource in the world. According to *Wind Force 10*, an international report by the European Wind Energy Association, by the year 2020, wind-generated electrical power will provide 10 percent of the world's electricity, reducing carbon dioxide emissions by more than 10 billion tons a year. Unlike fossil fuels, renewable energy does not exhaust the Earth's supply of non-renewable mineral resources such as coal and natural gas. A team of researchers from the Department of Energy's National Renewable Energy Laboratory (NREL), NASA Ames Research Center, and the University of California at Davis have launched a research program using both experimental and computational approaches to improve our understanding of wind energy systems. Researchers from Ames' Army/NASA Rotorcraft Division are using the high-performance computers at the NAS Facility to support the project.

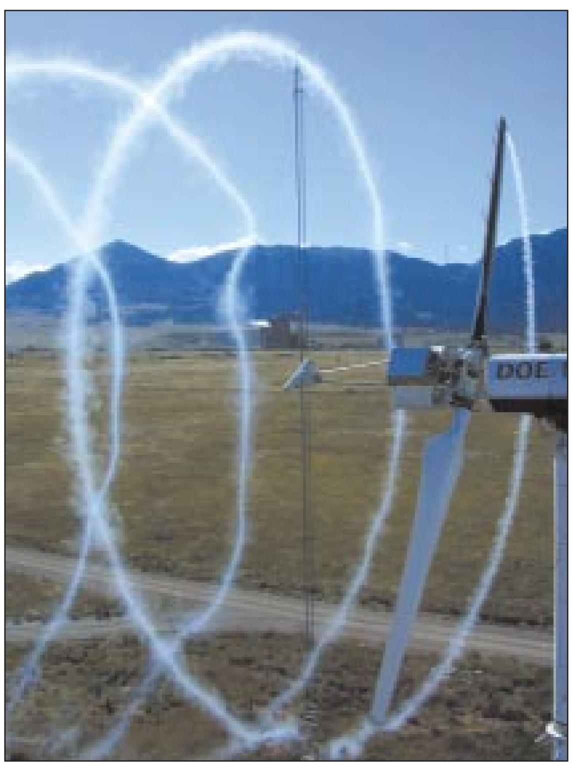

Before any testing took place in the wind tunnel at NASA Ames, experiments were conducted at NREL's field labs in Boulder, Colo. Researchers installed smoke generators in the tip of the wind turbine rotor to experimentally visualize the airflow and turbulence surrounding the blade. (NREL)

The experimental part of this program, known as the Unsteady Aerodynamics Experiment (UAE), tested NREL's wind turbine in the NASA Ames 80- by 120-foot wind tunnel from mid-April to early May. Computational fluid dynamics (CFD) predictions for the same wind conditions gathered during the wind tunnel experiment were computed on the CRAY C90 operated within the NAS Division and managed by the Consolidated Supercomputing Management Office (CoSMO). These computations require large amounts of data storage available through NAS's SGI Origin-based storage system, NAStore.

Computational Modeling

NREL's wind turbine experiment is the first in the United States to employ high-resolution Navier-Stokes calculations to investigate the aerodynamics of a wind turbine. Using CFD predictions, NREL and its collaborators plan to develop blade stall delay models for creating better turbine designs applicable to both the wind power and rotorcraft industries.

"Some problems encountered when designing an efficient wind turbine are related to those experienced with rotorcraft. Both wind turbines and rotorcraft have rotating blades with similar stall behavior," explains Bob Kufeld, NASA project lead for the UAE's wind tunnel experiments. Stall occurs when the rotor lift no longer increases with the angle of attack. Using results obtained from the well-known CFD code OVERFLOW, along with wind tunnel data, researchers can gain a better understanding of the rotor aerodynamics.

To improve wind turbine design, rotor design codes used by industry need a high degree of accuracy at low computational cost. "To date, there hasn't been a thorough set of experimental data to verify the accuracy of the computer codes because wind turbines operate in the inherently turbulent wind of the Earth's atmosphere. Current engineering models used are not very effective at higher wind speeds because our present theories do not accurately predict blade stall," explains Earl Duque, computational lead for the UAE in the Army/NASA Rotorcraft Division. Results from the UAE were obtained in the precisely controlled wind tunnel environment, and will provide industry with a valuable database to validate its computer models and CFD codes. This new data will enable designers to develop more reliable and cost-effective wind turbines.

Duque's extensive background in rotorcraft computations is beneficial for this project. "My experience working with others on rotorcraft computations in the past has made this project easier in the areas of visualization and grids." With the help of MIT researcher David Kenwright, formerly of NAS's data analysis group, Duque modeled vortex patterns over a rotorcraft blade, using Kenwright's flow visualization techniques. This work helped him visualize what is happening with the wind turbine blades. Working with NAS senior scientist Rupak Biswas, Duque also investigated grid adaptations, which enables accurate prediction of blade stall. Duque's earlier studies allowed him to understand how to build the grid systems for the NREL experiment.

Reducing the Cost of Energy

Charged by the Department of Energy to reduce the cost of electrical power produced from renewable sources, NREL began its wind experiment in 1989. The UAE research wind turbine has been field-tested in various configurations since the beginning of the project. "The current phase of the experiment (phase six) is the most realistic wind turbine blade – it's a two-blade rotor system, which is more modern and cost-effective than the three-blade configuration of earlier experiments. Two blades require less building materials but generate the same amount of power," says Kufeld.

When NREL began its studies, wind energy cost more than 20 cents per kilowatt-hour – which includes the system's initial capital investment, operation, and maintenance costs. Today, the price has dropped substantially to about four cents per kilowatt-hour. NREL's goal is to reduce this cost further, to two and a half cents per kilowatt-hour.

Stalled and Non-Stalled Turbine Blades (stream lines)

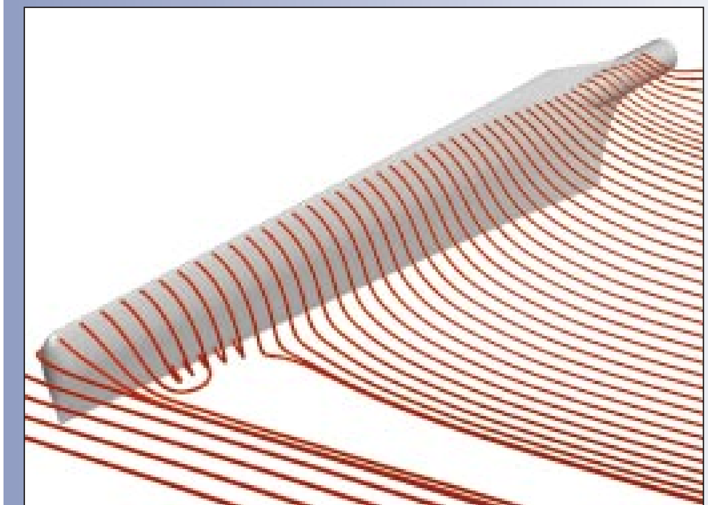

Representing the air surrounding the wind turbine blade, these visualizations of numerically simulated flow fields are based on wind tunnel test conditions of NREL's phase-six wind turbine. The visualizations were constructed using Fieldview Software, by Intelligent Light Inc., on a desktop computer. The red bands represent streamlines flowing over the blade at a single instant in time. Before the blade begins to stall (above), the air streams smoothly over the blade's surface. When the blade begins to stall, the airflow becomes unsteady, limiting power output (below). (courtesy Earl Duque)

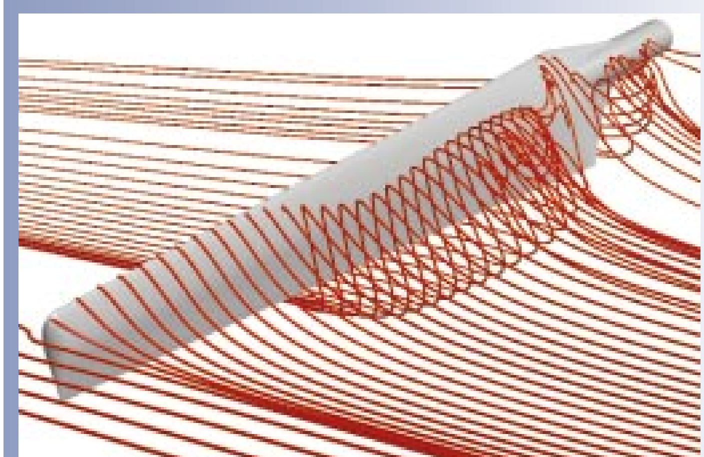

Comparisons: Actual vs. Computational

The wind tunnel test turbine is equipped with cameras, surface pressure taps, and strain gauges to gather data for validating numerical prediction models of aerodynamic blade responses. Testing in a controlled wind tunnel environment provides data that is free of interference from atmospheric turbulence, which substantially complicates turbine blade aerodynamic response. This data is ideal for validating computational blade stall models and CFD codes.

Continued on page 16

Wind Turbine Airfoil
(pre- and post-stall conditions)

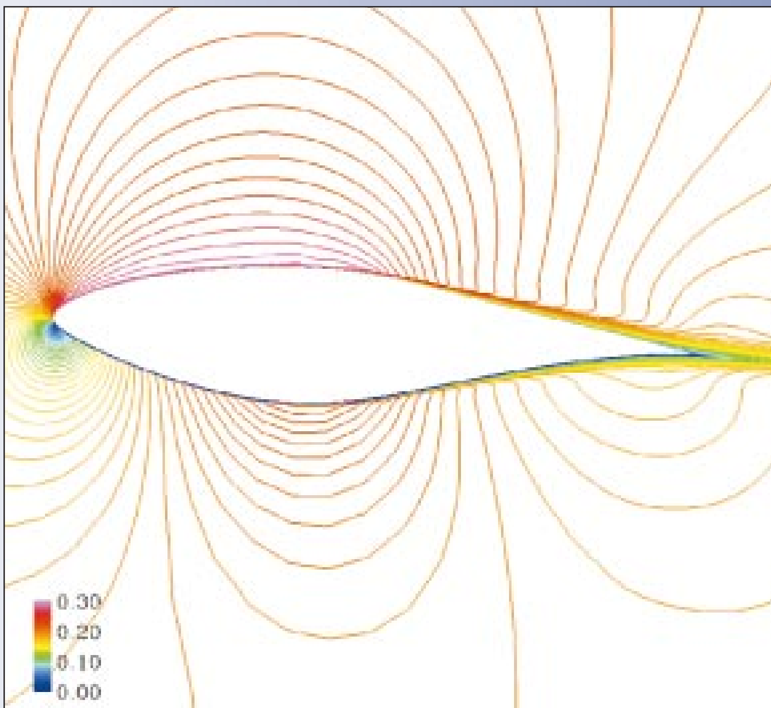

The graphic above represents the velocity contours of an S809 airfoil under pre-stall conditions. This airfoil is the same one used on NREL's Unsteady Aerodynamics Experiment wind turbine. Due to its design, the S809 airfoil was calculated to begin stalling at approximately 10-degrees angle of attack and to hold lift over a large range of angles. This lift characteristic helps maximize and control the power output of the turbine. Below, the airfoil is positioned at a 14-degree angle of attack, a post-stall condition. Wind turbines use stall to control power output of a system. Unit of measurement is Mach number. (Earl Duque)

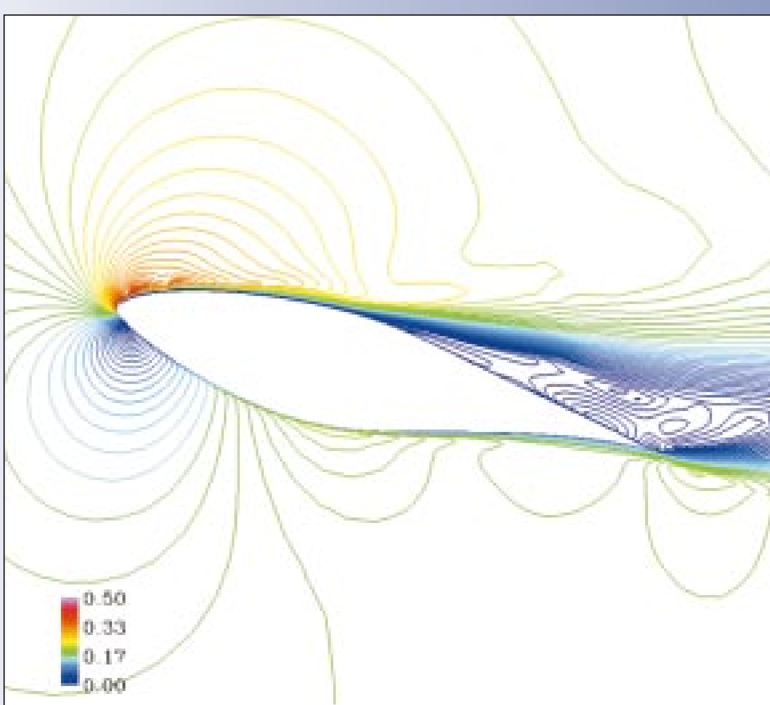

Continued from page 15

Working closely with NREL, Duque brought NAS computational resources to the project in 1997. "Computations extend the boundaries of what we can test," says Duque. "In the field or the wind tunnel, there are only so many things you can measure. With CFD you can see what's happening to the airflow around the blade that cannot be seen in field experiments." Only a few selected points on the rotor have sensors to collect data in an experiment. If the data from those points match well to computer calculations, every point of the flow field can be re-created using CFD.

Jasim Ahmad and Roger Strawn in the Army/NASA Rotorcraft Division have improved the algorithm that Duque used to run CFD calculations of the wind turbine in phase two of the experiment. Phase two computations took 150 CPU hours to complete. Now, phase six calculations can be done in 40 CPU hours on CoSMO's CRAY C90.

Duque will use the experimental data collected from the wind tunnel test and compare it to the predictions he creates using OVERFLOW (version 1.8), a Reynolds-averaged thin-layer Navier-Stokes (RaNS) code. He plans to use OVERFLOW-D for the next set of wind turbine computations. The first version of OVERFLOW, developed by Pieter Buning, currently at NASA Langley Research Center, was later modified by Ahmad for unsteady and steady rotor applications. OVERFLOW predicts the power output and the flow field around a turbine. Strawn is consulting with NAS researchers Rupak Biswas, Rob Van der Wijngaart, and Maurice Yarrow on a parallel version of OVERFLOW-D for rotorcraft.

In the future, OVERFLOW-D may be used by the Information Power Grid (IPG) community (see the IPG story on page 18) for other fluid flow problems requiring highly refined grids, resulting in massive amounts of computational data such as helicopter and tiltrotor configurations.

"OVERFLOW really tests the IPG system – there is too much data to process on any single machine at NAS, but it can potentially be done on the IPG once it is fully operational," says Strawn. The team anticipates running OVERFLOW on an SGI Origin 2000 for the phase six calculations in winter 2001.

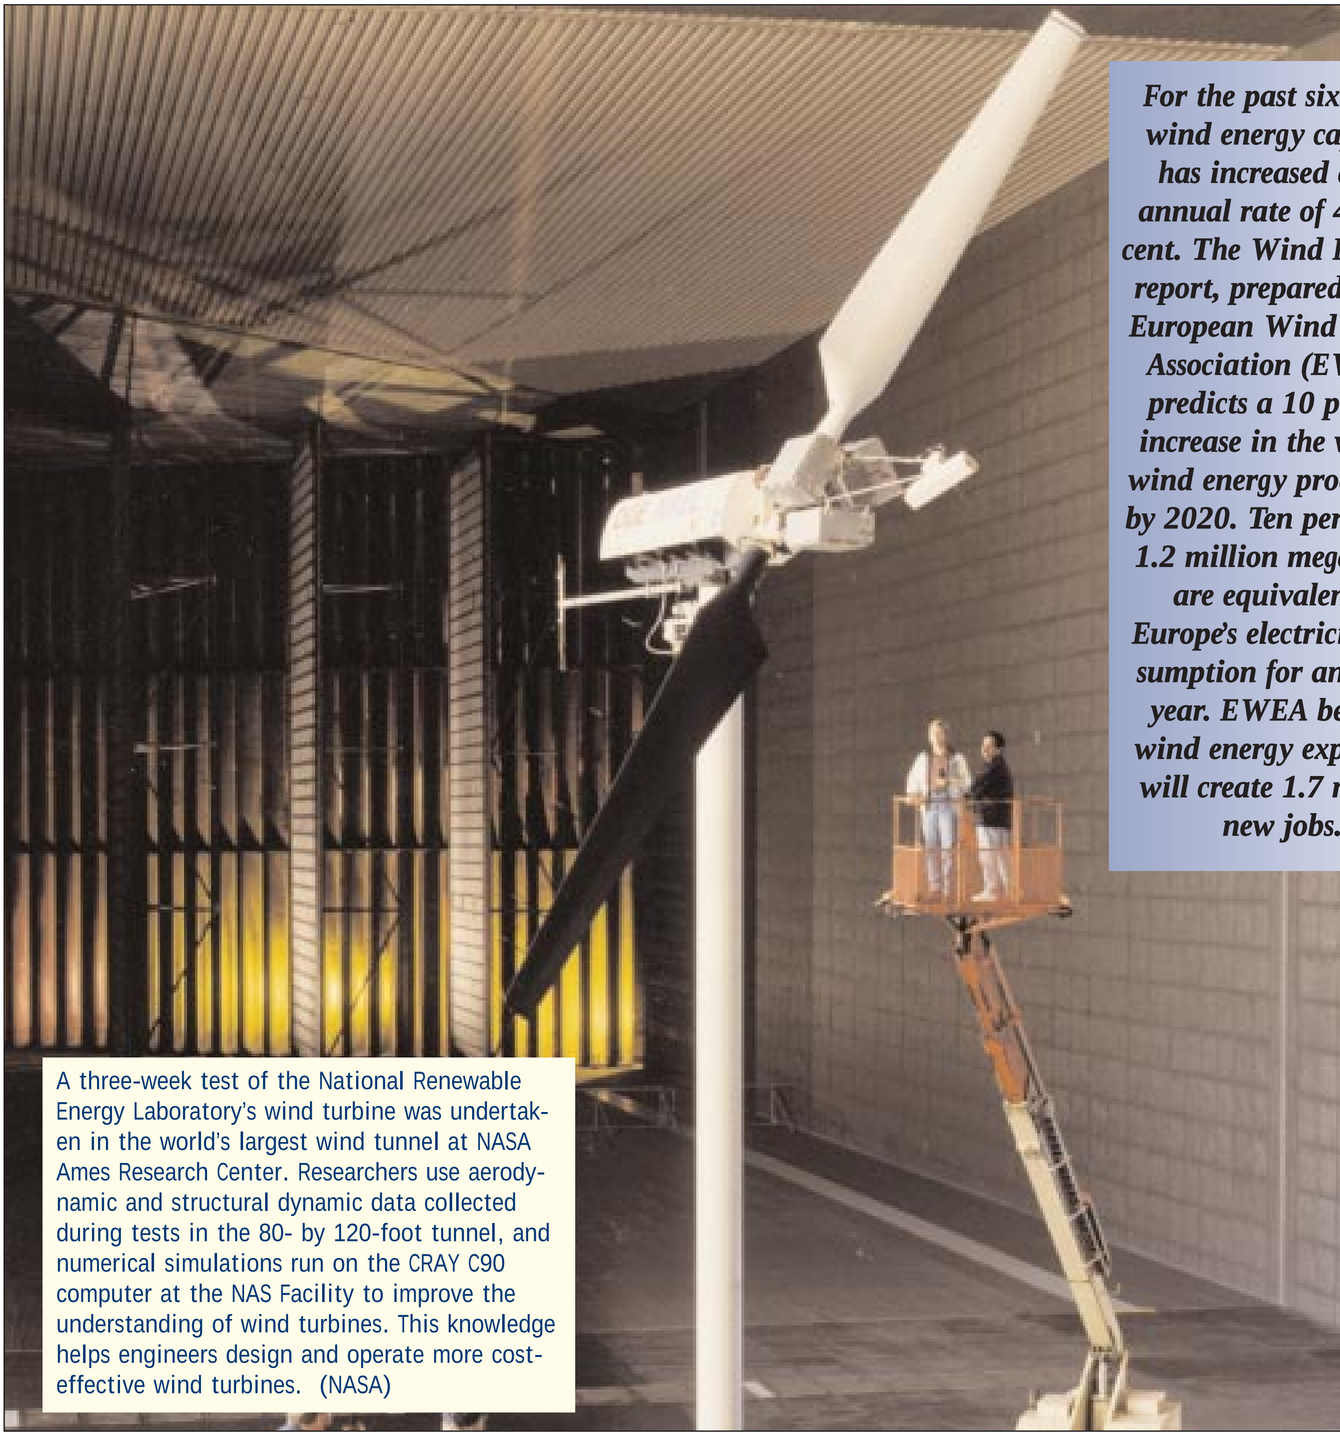

A three-week test of the National Renewable Energy Laboratory's wind turbine was undertaken in the world's largest wind tunnel at NASA Ames Research Center. Researchers use aerodynamic and structural dynamic data collected during tests in the 80- by 120-foot tunnel, and numerical simulations run on the CRAY C90 computer at the NAS Facility to improve the understanding of wind turbines. This knowledge helps engineers design and operate more cost-effective wind turbines. (NASA)

For the past six years, wind energy capacity has increased at an annual rate of 40 percent. The Wind Force 10 report, prepared by the European Wind Energy Association (EWEA), predicts a 10 percent increase in the world's wind energy production by 2020. Ten percent, or 1.2 million megawatts, are equivalent to Europe's electricity consumption for an entire year. EWEA believes wind energy expansion will create 1.7 million new jobs.

NREL, NASA Ames, and UC Davis have made measurable progress in the area of wind power production by providing the industry with detailed experimental data and valid computational methods. From gathering detailed pressure data on wind turbine blades to predicting the same pressure on CoSMO's CRAY C90, every accomplishment brings the Department of Energy closer to its goal – more affordable, clean, and sustainable electricity. gp

— *Holly A. Amundson*

Acknowledgments

Funding for the UAE wind turbine is provided by NREL's Chief of Applied Research, Mike Robinson, through the Army/NASA CFD Code Project. C.P. Van Dam and students at the University of California at Davis provide support for the computational component of NREL's UAE project.

Editor's Note: *As* Gridpoints *went to press, Earl Duque accepted a new challenge at the College of Engineering and Technology, Northern Arizona University in Flagstaff as an associate professor of mechanical engineering. Duque is continuing his research in wind turbine technology and applied computational fluid dynamics.*

Infrastructure For NASA's Information Power Grid Nears Completion

More than a dozen components of the IPG are being readied for roll-out by September 30.

During the 1970s, researchers were constructing a network of computers to share information in an effort to expand science. The resulting computer network of the '70s became the Internet of the 1980s and '90s. Today, researchers led by the NAS Systems Division are expanding the Internet model to build NASA's Information Power Grid (IPG) – a distributed system linking remotely located computer networks, storage devices, and instruments into a single computational grid.

NASA has selected Ames as the Center of Excellence for Information Technology and the NAS Systems Division as the focal point for high performance computing as well as the development of NASA's prototype IPG. In September, NASA researchers will demonstrate the IPG's capability and its ability to access geographically distributed systems. NASA's IPG team will deliver a "persistent testbed" for scientists to test their applications, and to provide the framework for systems software development. To meet the September 30 milestone, the team is developing more than a dozen tools to expand the grid's infrastructure.

IPG's Network Connectivity

The first measurable step in the grid's development was taken last April when NAS researchers demonstrated the IPG network's connectivity and its ability to access distributed archival data. The April milestone sought a data transfer rate between 10 and 50 megabits per second while retrieving files from two remote sites. Using NREN (the NASA Research and Education Network) for the connection, the tests delivered a connectivity rate of 58 megabits per second with data transfers from four separate locations.

"While the April demonstration is by no means high speed, network connectivity lays the foundation for the next step – a high-speed interconnect test," says Leigh Ann Tanner, manager, advanced systems development.

'Launch Pad' is Gateway to the Grid

The most visible piece of the grid's infrastructure is the web-based user interface, or portal, called the NASA IPG "Launch Pad." This interface is being developed in collaboration with researchers at San Diego Supercomputer Center, the National Computational Science Alliance (NCSA), Argonne National Laboratory, and the University of Indiana. The Launch Pad will allow researchers to create and submit a job, track its progress, and determine which machine is processing the work. This simple and easy-to-use web interface will allow a user to specify the resources required to run an experiment, such as the amount of memory required and the number of computer processors, and then "launch" the job to the grid.

"Scientists often perform the same basic task repeatedly," says George Myers, NAS scientific consultant and portal development group lead. "Being able to upload previously created job templates will only require changes to a line or two of code, or the addition of new file specifications, without completely rebuilding job submission information each time. We are working to make the job template modular so that researchers can customize their view of the portal."

The portal group is finalizing the look and feel of the NASA Launch Pad, which will use a grid portal development kit adapted from NCSA's Jason Novotny and Vaughn Welch's

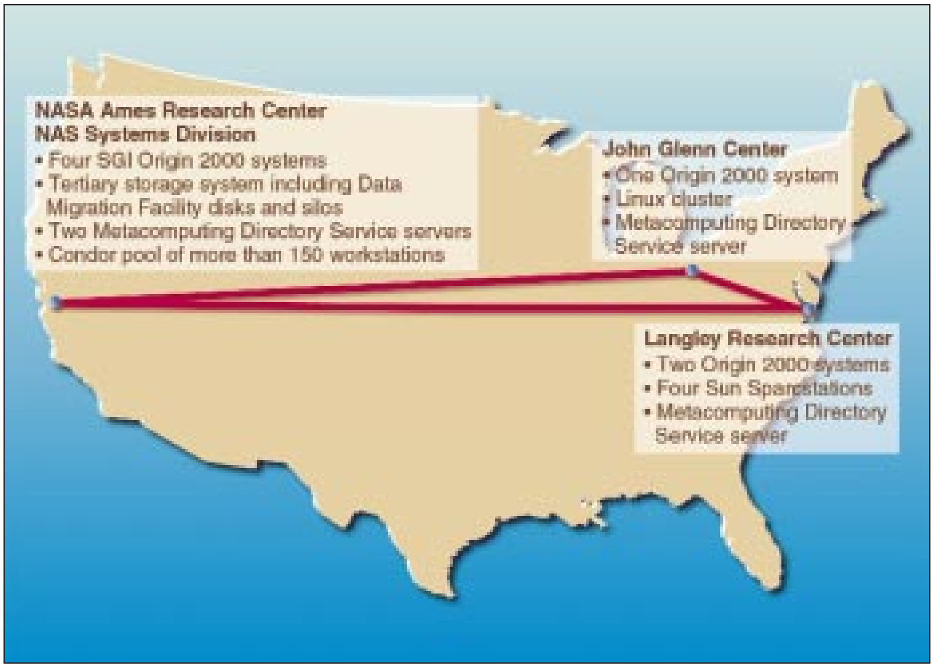

NASA's Information Power Grid links high performance computers at Ames, Glenn, and Langley Research Centers. Infrastructure enabling researchers to test applications and develop grid related systems software will be completed by September 30.

interface to the grid. Dennis Gannon of the University of Indiana has also been advising the portal development group.

User Support, Middleware Under Development

To answer grid users' questions and troubleshoot problems, NASA is developing a new grid support model that will provide 24 hour, seven day per week assistance to users, regardless of an end-user's location. The support model team is tasked with developing procedures for problem reporting and tracking, routing questions to the proper support group and how the situation will be remedied.

Varying levels of support exist at the three participating NASA centers, ranging from phone or e-mail support only to the NAS Systems Division's "24 x 7" assistance. For more extensive support, the NAS scientific consultants are available Monday through Friday from 7 a.m. to 5:30 p.m. (PST). Answers requiring other specific knowledge areas may reside with support staff located at other participating NASA research centers – currently Glenn and Langley. To address the intricacies of user assistance, the IPG team is developing one mechanism to standardize all of the support functions that will be in place by September 30.

NASA's prototype IPG will use the Globus middleware toolkit – developed by Ian Foster of Argonne National Laboratory and Carl Kesselman of the University of Southern California's Information Sciences Institute, to enable different types of applications and interfaces to communicate with the grid. Judith Utley, a NASA IPG team member, coordinates upgrades and new releases of Globus to each of the grid sites, and resolves software bugs. Utley also manages the deployment of the Grid Information Service (GIS) that allows researchers to poll the GIS to determine what jobs are running, as well as types of systems and processors are available. "Globus deployment and grid information services require close coordination between all NASA IPG centers," says Utley. "Our weekly teleconferences provide the avenue to complete this phase of the project." Utley collaborates on GIS deployment with Glenn Research Center's Kim Johnson and Allen Holtz as well as Langley Research Center's Mike Perry and Greg Cates.

Keeping the Grid Secure

"Security for both the user and the grid is of the utmost importance," says NAS systems engineer Yvonne Malloy, who is implementing the IPG's certificate infrastructure using Netscape CMS (Certificate Management System) to address this issue. To obtain authorization to run jobs on the IPG, a user will request a certificate through the online grid registration authority. Once a user has been authenticated and authorized, a certificate is issued and users can download it onto their web browser. Each job is "signed" with the certificate. Malloy is currently incorporating the certificate system into the Globus structure, and is planning to take user security to the next level using the 'MyProxy' software system also developed by Welch and Novotny.

"MyProxy is a way for users to limit the exposure of their private key password," says Malloy. "The private key is decrypted long enough to generate proxies and then it is re-encrypted. These proxies are valid for short amounts of time, but are treated like the user's actual certificate." MyProxy is expected to be installed by the end of August.

Managing the Results

After running a job on the IPG, results may be stored in a location capable of handling tremendous volumes of data. For data management, NAS engineers are using a mass storage system for NASA's IPG. Known as the Data Migration Facility (DMF), it consists of two SGI Origin-based systems, a large fiber channel RAID (Redundant Array of Inexpensive Disks) array capable of storing 1.3 terabytes of data, and four attached tape drives. The tape drives will migrate stored files from the RAID array to removable tape cartridges. "Off-loading a file's data blocks from the RAID array to tape will give us virtually unlimited storage capacity," says Bill Arasin, a systems analyst in the high speed processor group. "There will be a brief delay when users retrieve their files, but this offers the ability to store as much data as we have tape."

The DMF is a redundant system that maintains a copy of each file on two separate disks in the event of a disk failure. To add further redundancy, both systems will be configured to run Failsafe DMF which enables either system to take over all the services of the other if a system outage occurs. "Failsafe DMF provides a seamless shift between servers with no interruption of service detectable by users," says Arasin.

From left to right: Ames Center Director Henry McDonald, NAS Division Chief Bill Feiereisen, NASA Administrator Daniel Goldin, and NAS senior scientist Jim Taft discuss high performance computing and system scalability in the NAS Facility's computer room. *Lomax*, the 512-processor SGI Origin 2800, shown in the background, is a node on the IPG and was a key component in demonstrating the grid's connectivity and ability to access distributed archival data.
(Nicholas A. Veronico)

Proving the Grid's Reliability

To ensure that IPG users will have a reliable computing environment, NASA engineers have been testing the grid's performance. Validation encompasses four areas: Each piece of the grid's infrastructure is repeatedly tested to ensure that it works and will perform adequately; that it will interface with other systems and software; that it will perform uninterrupted while handling a large number of jobs; and that it remains stable as more computations are added into the system. "Our tests are designed to model the tasks that will most frequently be done by users on a day-to-day basis in the grid environment," says Ray Turney, IPG test engineer.

"All building blocks of the grid's infrastructure are interdependent upon the others," Turney continues. "For example, when new software versions are introduced, the system environment may change." Nightly tests run simulated jobs through the systems at Glenn, Langley, and the local machines at Ames. Turney is currently testing the Globus environment and the Metacomputing Directory Service.

Condor: Adding Flexibility to the IPG

In addition to high performance computing, where job speed is measured in floating point operations per second, researchers can also use high throughput computing resources where large numbers of processors work for long periods of time – measured in weeks or even months. "We're not only linking supercomputers," says Eric Langhirt, NAS Condor systems specialist, "but workstations that are idle will now become a resource for the grid." Condor, developed by Miron Livny and a team of researchers at the University of Wisconsin, is a distributed batch processing system that scavenges compute cycles from idle workstations on a network and runs jobs while the local user is away.

The software system works by monitoring the CPU load and keyboard activity of a pool of workstations to determine whether a terminal is idle and available to run jobs. The system matches a job's requirements for resources with workstations that have available computing cycles. Once a match is made, the job is sent to the workstation for processing. Condor continually monitors the workstation's CPU load and its keyboard; movement of the mouse or a keystroke will restore control of the computer to the local user. Condor then moves the job to the next idle workstation in the pool and continues processing.

Using the Condor system, the NAS Facility's more than 400 workstations will provide nearly 2 million CPU hours each year – equal in power to one of the NASA's 256-processor high performance computers. When all three NASA grid centers are linked, this number is expected to nearly double.

Looking Ahead

Additional infrastructure components slated for completion this fall include a distributed debugging and performance analysis system, a wide area network connection capable of transmitting at least 100 megabits per second, a computer hours allocation and accounting system, and the integration of the CORBA programming language to the IPG.

Each completed and integrated piece of the grid's infrastructure brings the IPG team closer to meeting its September goal of "demonstrating seamless operations within a heterogeneous distributed computing environment." gp

— *Nicholas A. Veronico*

Calendar of Events

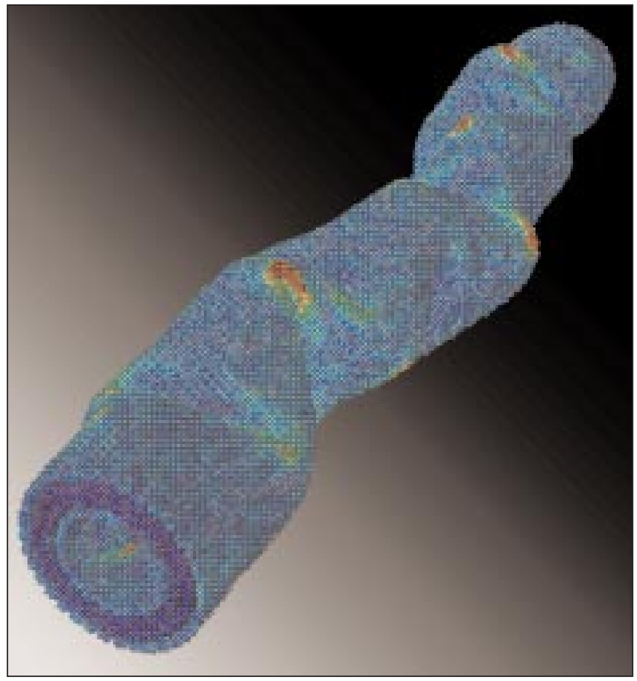

Multi-wall nanotube compression (buckling), simulation and visualization by Deepak Srivastava

September 11-15 – Albuquerque, New Mexico
The Linux Supercluster Users Conference is aimed at sharing information on scientific computing techniques. The objective of this meeting is to help computational scientists and engineers develop applications to achieve maximum performance and scalability on Linux Supercluster systems. Hosted by the National Computational Science Alliance (NCSA), the Advanced Computing Technology Center at IBM Research, and the High Performance Computing Education and Research Center. For additional information, visit:
http://clumber.tc.cornell.edu/actc/Supercluster/index.html

September 19-20 – Mountain View, California
NASA's Information Power Grid project will host a workshop to demonstrate and review the IPG's growth, examine current issues, and plan future directions. Workshop proceeding and selected papers will be published in a special issue of "Scientific Programming." Contact Marcia Redmond at (650) 604-4373 or visit the workshop's website at: http://www.ipg.nasa.gov/ workshop.html

September 18-20 – Davis, California
The Fifth Symposium on Overset Grids and Solution Technology will be held at UC Davis. This international symposium provides a joint forum where mathematicians, scientists, and engineers from academia, industry, and federal laboratories can exchange research ideas, new algorithmic and software developments, and applications pertinent to the state-of-the-art and state-of-the-practice of Overset/ Chimera Grid and Solution Technology. More information about the symposium can be found at:
http://ntserver.itd.ucdavis.edu/Chimera2000/

September 28 – Mountain View, California
SGI will roll out its new Origin and Onyx 3000 modular computing systems at its Mountain View headquarters from 2 p.m. to 7 p.m. Bill Feiereisen, NAS Systems Division chief, is a featured guest speaker. For directions, registration, and additional information, call (800) 318-6432, or visit: http://www.sgi.com/events/thinkahead

October 9-13 – San Jose, California
The Microprocessor Forum 2000 provides designers and users of high-performance microprocessors with insights into new processor designs – for applications ranging from PCs and workstations to information appliances and the Internet. For additional information, visit the conference's website at: http://www.MDRonline.com/MPF/conf.html

October 30-November 1 – Orlando, Florida
Storage Networking World Conference and Expo features panel discussions and case studies on the future of Storage Networking and topics such as storage outsourcing, managing storage and server assets, Storage Area Network interoperability, the role of storage systems within the Internet infrastructure and integrating diverse environments, along with a separate track covering in-depth technical issues and deployment examples associated with storage networks. More information is available at:
http://www.storagenetworkingworld.com

November 3-5 – Bethesda, Maryland
Eighth Foresight Conference on Molecular Nanotechnology. This conference is a meeting of scientists and technologists working in fields leading toward molecular nanotechnology: thorough three-dimensional structural control of materials and devices at the molecular level. An introductory tutorial on foundations of nanotechnology will be held November 2. For more information, phone (650) 917-1122, or fax (650) 917-1123. The conference's web address is:
http://www.foresight.org/Conferences/MNT8/index.html

November 4-10 – Dallas, Texas
SC2000, the conference of high performance networking and computing, will meet at the Dallas Convention Center for a series of technical programs, demonstrations, educational outreach, and visualizations of computational data. SC2000 is sponsored by the Institute of Electrical and Electronics Engineers Computer Society and the Association for Computing Machinery's Special Interest Group on Computer Architecture. Be sure to visit the NASA booth. Further details can be obtained at: http://www.sc2000.org