Summer 2001

Gridpoints

The Quarterly Publication of the NASA Advanced Supercomputing Division

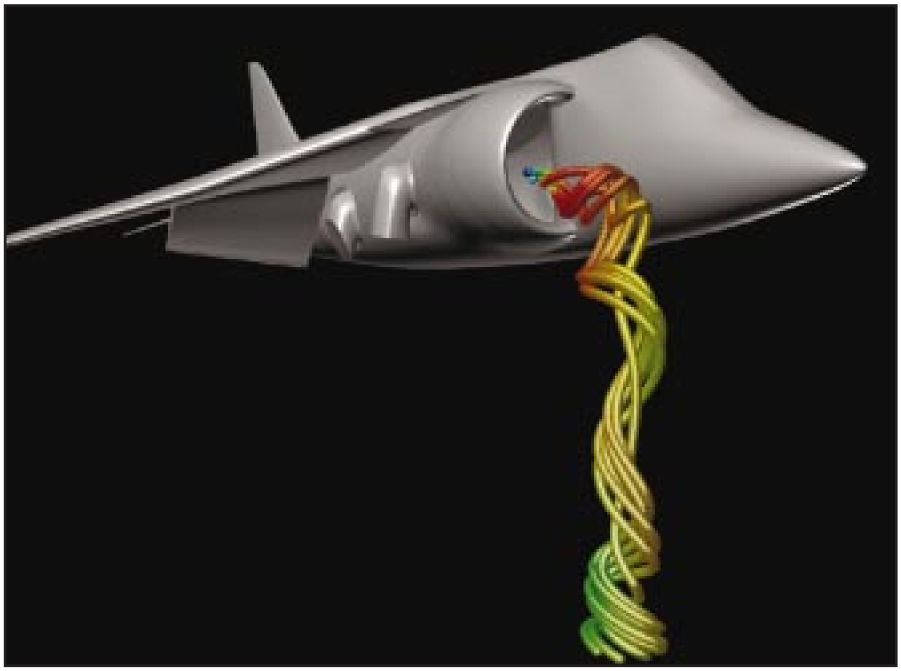

New algorithms have enabled NAS researchers to visualize extremely large datasets using personal workstations, cutting computation time and costs, while increasing resolution. See page 16

FEATURES

Features

On The Cover:

Newly developed algorithms have enabled NAS data analysis group researchers to visualize the ground vortex and jet exhaust from the Harrier aircraft from different angles. For more information, please turn to page 16.

Gridpoints
The Quarterly Publication of the Numerical Aerospace Simulation Systems Division
Vol. 2, No. 2 • Summer 2001

Executive Editors: Bill Feiereisen, John Ziebarth

Editor: Nicholas A. Veronico

Writers: Holly A. Amundson, Neal Chaderjian, Jill A. Dunbar, David Ellsworth, Timothy A. Sandstrom

Graphic Designer: Sue Kim
Illustrator: Gail Felchle

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: David Alfano, Dick Anderson, Charles Bauschlicher, Michael Boswell, Johnny Chang, Jill Dunbar, Chris Gong, T.R. Govindan, Daniel Herr, Mary Hultquist, Bill Johnston, Randy Kaemmerer, Chris Kleiber, MiYoung Koo, Marco Librero, Y.K. Liu, Bill McDermott, Nagi Mansour, Meyya Meyappan, Pat Moran, George Myers, Terry Nelson, Jana Nguyen, Marcia Redmond, Ryan Spaulding, Deepak Srivastava, Sumit Talwar, Leigh Ann Tanner, Bill Thigpen, Judith Utley, Arsi Vaziri, Cliff Williams

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

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

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

News From NAS

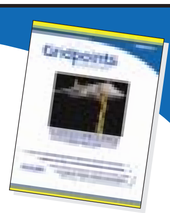

NAS Name Change Reflects High-End Computing Focus

After nearly 20 years of sustained growth, the Numerical Aerospace Simulation (NAS) Systems Division, part of the Information Sciences and Technology Directorate at Ames, is changing its name to the NASA Advanced Supercomputing (NAS) Division. The change more accurately reflects the division's leadership role within the high-performance computing and information technology communities.

The NAS Division was formed in the early 1980s to manage NASA's computational fluid dynamics efforts and the machines used in CFD. The division has since evolved as a world leader in the development of NASA's Information Power Grid (IPG), which allows remotely located researchers to access distant computational resources; and through its NASA Research and Engineering Network (NREN) affiliate, engineers within the NAS Division are working to build the Next Generation Internet.

The NAS Division will continue to focus in areas of research that are key to developing new, integrated high-performance computing environments such as the the IPG: applications, networks, problem-solving environments, high-performance computing, mass storage,

Continued on Page 2

NAS Mission

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

From The Division Chief

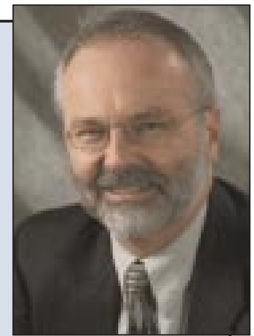

It's official – the NAS division has modified its name to more accurately reflect our expanded role in the supercomputing community within NASA and externally in the high-end computing community. Now known as the NASA Advanced Supercomputing Division, we were formed in the early 1980s to simultaneously develop supercomputers and advance the science of computational fluid dynamics. In the last five years the division has broadened its focus to support other areas of NASA computational sciences by bringing new developments in high-performance computing to bear in Earth-, space-, life-sciences, and nanotechnology, as well as our traditional area of aeronautics.

Distributed high-performance computing is the future model by which we will deliver computational horsepower to our scientists, allowing remotely located researchers to access distant computational resources through a common interface. The Information Power Grid is NASA's contribution to the worldwide community developing this technology.

We continue to co-develop the "power generators" on this grid and apply them to NASA's scientific research efforts. As an example, the computer scientists in our terascale applications group have worked during the last year with Earth scientists at NASA Goddard's Data Assimilation Office to improve the speed at which we are able to simulate the Earth's climate. The computational tools of global climate modeling are indispensable for answering questions related to our changing environment. But these questions demand computers and software tools that are among the largest and most intricate that have ever been produced. Working with Goddard and our partners in industry, we have been able to apply new programming techniques to the Data Assimilation System (DAS) that allow us to speed up our calculations by a factor of ten.

Running on the new 512-processor SGI Origin 3800 supercomputer, *Chapman*, using 384 processors, the DAS system can now process more than 200 days of weather information per day. Our terascale applications group researchers believe they can further increase the code's performance by a factor of two – delivering more than 400 days per day.

Using multi-level parallelism (MLP) coding techniques developed at NAS, researchers anticipate increasing the size of the Data Assimilation System's computations by adding more weather observations into the calculation while reducing the memory requirements needed to complete those calculations. Better utilization of our computational resources requires less capital investment in computer hardware, which translates into the NASA mantra of "faster, better, cheaper." The combination of MLP coding techniques and the single-system image architecture of the Origin 3800 supercomputers promises a new level of performance in global climate modeling and, therefore, in NASA's ability to address the important questions of global climate change.

As always, I welcome your feedback.

Bill Feiereisen
wfeiereisen@mail.arc.nasa.gov

grid common services, and other research areas, such as nanotechnology and device modeling.

NAS Engineer Honored at Annual Flemming Awards

NAS Division aerospace engineer Stuart Rogers was awarded the prestigious Arthur S. Flemming Award for his contributions to the field of computational fluid dynamics (CFD) during a ceremony at George Washington University, Washington, D.C., on June 5. Rogers' development and application of CFD tools has resulted in significant strides in solving real-world fluid dynamics problems, including analysis and design of complex aircraft, spacecraft, and submarines.

Stephen J. Trachtenberg, right, George Washington University president, presents NAS researcher Stuart Rogers with the Arthur S. Flemming Award during the June 5 ceremony in Washington, D.C. (Washington Photography Productions Inc.)

Rogers' work has had far-reaching benefits to NASA and the aerospace industry by providing tools that reduce the cost and design-cycle time of aerospace vehicles and components.

"Stuart has expertise not only in algorithms for computational flow simulation and in computer science for developing software, but he also has a deep understanding of flow physics," says Dochan Kwak, chief of NAS's Applications Branch. "On top of his multifaceted talents, Stuart is an extremely hard-working researcher. It's not surprising that he's made many critical contributions to Ames, NASA, and the nation's technical well-being."

Among Rogers' many important accomplishments in his 12 years at NASA Ames Research Center are the development of a software tools package that allows first-of-a-kind flow computation about an entire subsonic transport aircraft, including simulation of a complete Boeing 777 aircraft in a high-lift landing configuration. The Chimera Grid Tools software package, together with the OVERFLOW flow solver, received honorable mention in the 1998 NASA Software of the Year Award.

Rogers also co-developed the INS3D code resulting in world-class production software for the solution of incompressible Navier-Stokes equations. The code has had a broad impact on a wide range of applications, including advanced rocket pump development and the DeBakey Ventricular Assist Device, a miniaturized heart pump that prolongs the life of patients awaiting a transplant.

Rogers holds a bachelor's degree and a master's degree in Aerospace Engineering from the University of Colorado, and a doctorate in Aeronautics and Astronautics Engineering from Stanford University. His recent research and development work has been funded by NASA's High Performance Computing and Communications Program. The Arthur S. Flemming Awards, established in 1948, honors outstanding federal employees.

NASA Ames Gets New Cray SV1 Vector Supercomputer

NASA's Consolidated Supercomputing Management Office (CoSMO) has leased a Cray SV1e vector supercomputer from Cray Inc., Seattle, through Government Micro Resources Inc. The system, installed at the NAS Facility on May 25, will be used for scientific applications in aerospace, earth science, and space science.

The new SV1e, which will replace CoSMO's current CRAY C90, has 32 processors (CPUs) and four gigawords (GW) of main memory, compared to the C90's 16 CPUs and one GW of memory. Another four GW of solid-state memory will be installed later this year to create a Cray SV1ex. The air-cooled SV1e is a fourth-generation complementary metal-oxide semiconductor vector system designed to handle a broad range of application workloads.

"After installation and acceptance of the 'x' solid state memory, we hope to transition all users from the C90 to the Cray SV1ex," says CoSMO Deputy Director Ken Stevens. The SV1ex enhancement will include improved clock speed to 450 megahertz and improved cache over the Cray SV1e.

NAS staff will install additional applications (such as OVERFLOW and Laura) and math libraries on the SV1e to match the current C90 environment. After benchmarks are run, several experienced users will test-run their own applications. After running both systems in parallel for about a month, the C90 will be retired. Once testing is complete, users from government organizations – as well as industry and university users with government contracts, grants, or cooperative agreements – will be able to use the SV1e for their large-scale, scientific applications.

NAS Researcher Tackles Puzzling Nanotechnology Problem

NAS researcher Toshishige Yamada recently presented a new interpretation of a problem that has puzzled scientists for several years. In the past, researchers have observed strange current-voltage (I-V) patterns in a scanning tunneling microscope tip-carbon nanotube system, and attributed them to the intrinsic nanotube properties.

Now, Yamada has proposed a new model for the observed patterns. His work shows that a measurement electrode-nanotube junction is responsible for the observed behavior, and not the nanotube itself. Yamada's paper, "Modeling of Electronic Transport in Scanning Tunneling Microscope Tip – Carbon Nanotube System," was presented at the American Physical Society Meeting in Seattle, March 12-16, and published in *Applied Physical Letters*, Vol. 78 (12) pp. 1739-1741 (2001). *Applied Physical Letters* is on the web at:
http://ojps.aip.org/aplo

Report Compares Parameter Study Tools

A recent NAS technical report "A Comparison of Parameter Study Creation and Job Submission Tools," compares the differences among available general-purpose parameter study and job submission tools. NAS authors Adrian DeVivo, Maurice Yarrow, and Karen McCann focus on comparing a new NAS-developed software package, called ILab, with other tools. Results show that ILab is easier to use, completes jobs significantly faster, and is better suited to research and engineering environments.

The team compared the functionality of ILab, Nimrod/G, Nimrod/Clustor, Condor, and AppLeS/APST. For background on ILab, see the NAS feature story, "ILab: New IPG Tool Improves Parameter Study Implementation" (*Gridpoints*, Spring 2001, page 12). A postscript or PDF version of the technical report can be downloaded at:
www.nas.nasa.gov/Research/Reports/Techreports/2001/nas-01-002-abstract.html

Mansour Appointed Deputy Director for CTR

Nagi Nicolas Mansour, lead scientist for the NAS Division's Physics Simulation and Modeling Office, has been appointed Deputy Director for the Center for Turbulence Research (CTR) at Ames Research Center. Mansour has extensive experience in turbulence research, and has served as Ames coordinator to the CTR for the past six years. He will continue his duties in the NAS Division. Mansour holds a bachelor's degree in mechanical engineering, master's degrees in mechanical engineering and mathematics, and a doctorate in mechanical engineering. In recognition of his work, Mansour was elected as an Associate Fellow of the American Institute of Aeronautics and Astronautics, and Fellow of the American Physical Society for his work on turbulence, and drops and bubbles.

The CTR is a research consortium for fundamental study of turbulent flows, jointly operated by NASA Ames Research Center and Stanford University. The CTR's principal objective is to stimulate significant advances in physical understanding of turbulence, leading to improved capabilities for control of turbulence and turbulence modeling for engineering analysis. Particular emphasis is placed on probing turbulent flow fields developed by direct numerical simulations and/or laboratory experiments using new diagnostic techniques and mathematical methods, and on concepts for turbulence control and modeling. The CTR is directing its attention to the application of fluid dynamics to biology and the origin of life, protoplanetary disks, and atmospheric and geophysical flow phenomena.

Understanding How Electrons Spin

A study to determine how electrons spin has been published in the journal, *Physica Status Solidi-B*, Vol. 222, p. 523 (2000). The study, "Dyakonov-Perel Effect on Spin Dephasing in n-Type GaAs," conducted by Cun-Zheng Ning, NAS Research Branch, and M. Wu, University of California at Santa Barbara, explores the possibility of using electron spin coherence for application to faster, more energy-efficient information processing in the future.

Using a new quantum kinetic theory, Ning and Wu have conducted a large-scale computer simulation to determine the electron spin coherence lifetime in Gallium Arsenide (GaAs), a compound semiconductor material typically used for optoelectronics applications. The quantum kinetic approach and the results obtained will help in the understanding of spin coherence for future information technologies, and for quantum information processing and transport. A copy of the study is available online at:
www.wiley-vch.de/contents/jc_2232/2222.html

Gridpoints Readers Survey

The staff of *Gridpoints* would like your input about the publication. A short survey has been posted on the publication's home page at:

http://www.nas.nasa.gov/Gridpoints

Your opinion matters and is greatly appreciated. —*Eds.*

NASA Ames Center for Nanotechnology Earns High Marks from Reviewers

A biannual peer-review of the Center for Nanotechnology was held last spring to provide guidance to researchers while showcasing the work of Ames scientists.

A tremendous amount of progress has been made in the area of nanotechnology research at NASA Ames' Center for Nanotechnology since its inception in early 1996. The nanotechnology effort began with a handful of researchers from the NAS Division running computations on nanotubes, and has since grown into the robust Center for Nanotechnology. Encompassing both computational and experimental aspects of nanotechnology, the center is directed by Meyya Meyyappan.

Progress on research at the Center for Nanotechnology is peer-reviewed every two years – most recently on April 24 and 25. Ames Center Director Henry McDonald kicked off the April review session with a few words of encouragement, "I think it is important to have peer reviews. This group has made a significant impact, and I am pleased with the development in the area of nanotechnology."

Speaking before a panel of reviewers from industry and academia were members from Ames' nanotechnology research and development, computational nanotechnology, and computational electronics and optoelectronics groups, which make up the Center for Nanotechnology. Presentation topics ranged from controlled growth of carbon nanotubes to computer codes for modeling electron transport through carbon nanotubes, and quantum computing.

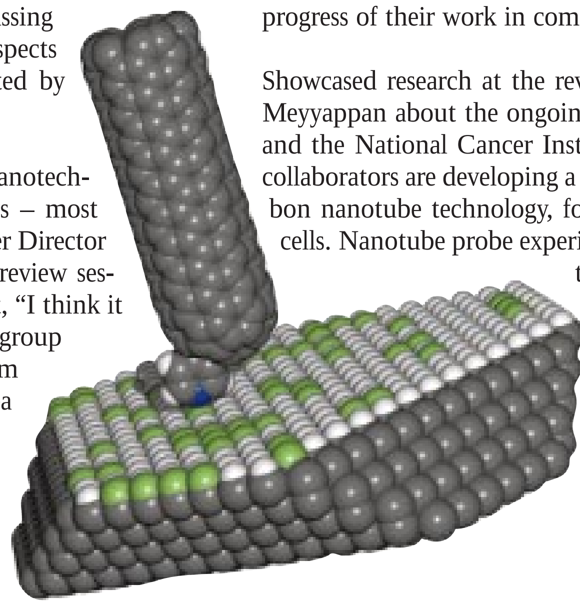

Figure 1: *A Buckytube probe (comprised of a pyridine molecule attached to, or functionalized with a carbon nanotube) serves as the tip of an atomic force microscope. The pyridine molecule was selected as the probe molecule because of its ability to identify atoms, or data on the surface of a sample (in this case, diamond with some of its hydrogen surface atoms replaced with fluorene). Developed by Charles Bauschlicher and Ralph Merkle; visualization by Al Globus.*

NAS researchers Deepak Srivastava, Charles Bauschlicher, Toshishige Yamada (see "The Smallest Nanoelectronics: Atomic Devices with Precise Structures," *Gridpoints*, Summer 2000, page 10), and Cun-Zheng Ning (see "VCSEL Lasers: Tiny Lasers, Huge Potential," *Gridpoints*, Spring 2000, page 12) were among the scientists to report on the progress of their work in computational nanotechnology.

Showcased research at the review included an update from Meyyappan about the ongoing collaboration between Ames and the National Cancer Institute (NCI). The NASA/NCI collaborators are developing a miniature probe, based on carbon nanotube technology, for early detection of cancerous cells. Nanotube probe experiments have reached an *in vitro* testing stage using biological samples (see "NAS's Nanotube Technology Used to Develop Biomedical Devices," *Gridpoints*, Spring 2000, page 3).

At the conclusion of the review session, the panel of evaluators reported to Ames Deputy Director Bill Berry: "The quality of the work is excellent. The energy and enthusiasm of the group was very visible which was an indicator of how excited they are about their work." In the past four years, researchers from the Center for Nanotechnology have published more than 170 papers and contributed more than 350 seminar presentations and invited talks to publicize the work accomplished in the area of nanotechnology at Ames.

For additional information on the Center for Nanotechnology, visit: www.ipt.arc.nasa.gov

— Holly A. Amundson

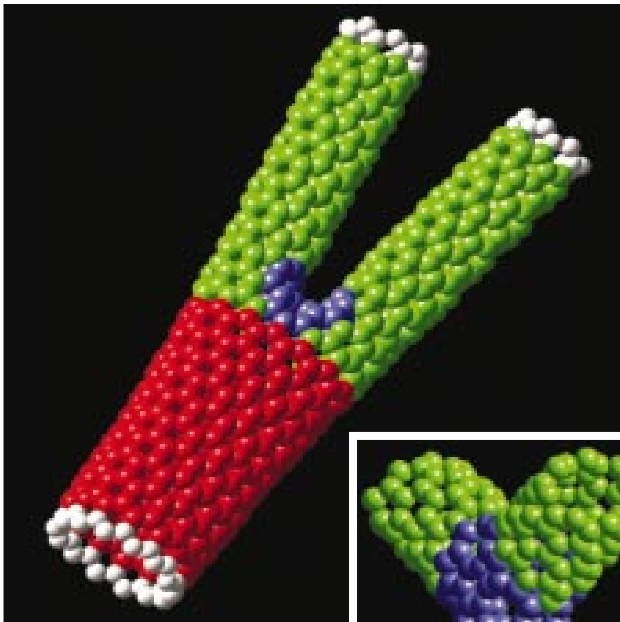

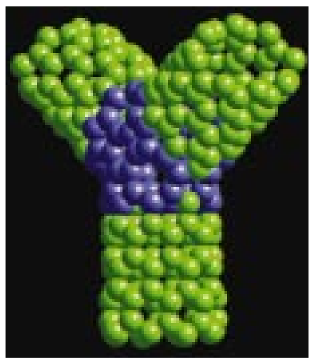

Figure 2 (above): *These images represent carbon nanotube junctions – the main figure is an acute angle y-junction, while the inset image is an obtuse angle y-junction. NAS researcher Deepak Srivastava and colleagues were the first to propose the use of three-terminal carbon nanotube junctions for nanoscale molecular electronic device applications. The switching behavior of the junctions is affected by structural symmetry – there is much less destructive interference between the electronic wave function in the two branches of a symmetrical y-junction. The configurations of these model junctions were tested using the NAS facility's Origin 2000 computers* Steger *and* Turing. *These figures were generated using a visualization program created by Srivastava and collaborator Madhu Menon from the University of Kentucky.*

Figure 3a–e (right): *In theory, nanotubes are capable of storing hydrogen in a very dense form and could potentially be used to fuel rocket engines to Mars. The figures at right demonstrate hydrogen storage in carbon nanotubes with different configurations of hydrogen molecule placement. Green molecules represent carbon nanotubes, white molecules represent hydrogen coverage on the nanotubes. There is an increase in bond strength with the addition of hydrogen molecules, as well as greater stability in structures with ordered hydrogen placement. The configuration in Figure 3a demonstrates the most stable structure and the highest binding energy, due to the careful placement, or symmetry of hydrogen atoms in the carbon nanotube. Research conducted by Charles Bauschlicher.*

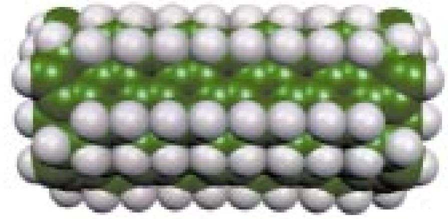

Figure 3a: 50 percent hydrogen coverage - pairs of lines configuration.

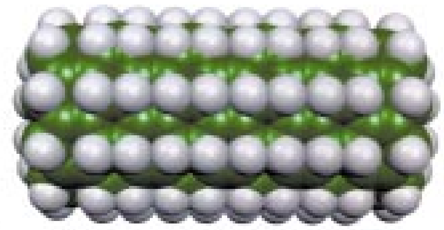

Figure 3b: 50 percent hydrogen coverage - lines configuration.

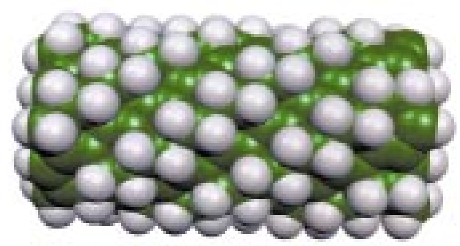

Figure 3c: 50 percent hydrogen coverage - spiral configuration.

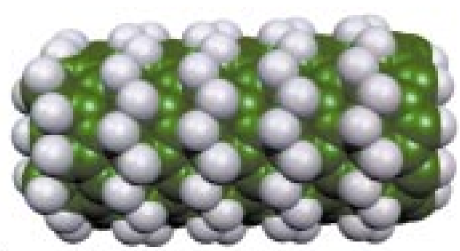

Figure 3d: 50 percent hydrogen coverage - rings configuration.

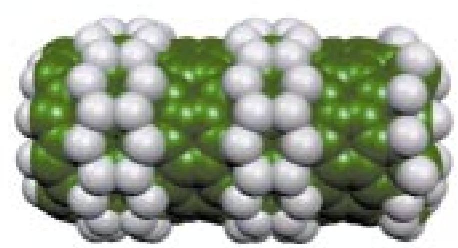

Figure 3e: 50 percent hydrogen coverage - pairs of rings configuration.

Reducing Jet Noise Through Computational Aeroacoustics

Scientists at the joint NASA/Stanford University Center for Turbulence Research are using NAS computational resources to predict the radiated jet noise of aircraft engines in an effort to reduce noise emissions.

In 1975, it was estimated that 7.5 million Americans lived in areas where aircraft noise exceeded 65 decibels – equal to having a TV set playing loudly in the background 24-hours a day. The Airport Noise and Capacity Act of 1990 set down what became known as "Stage 3" regulations that required more than 7,500 jet airliners to be modified with quieter engines or to be retired from the fleet by December 31, 1999. The direct result of Stage 3 regulations saw the number of people affected by jet noise reduced to about 600,000 according to the Federal Aviation Administration.

Currently, the International Civil Aviation Organization, in conjunction with aviation regulatory agencies in the United States and the European Union, are working to address noise and environmental issues from jet aircraft. The results of these meetings will establish noise reduction and emission guidelines for future Stage 4 regulations.

Two sources of noise from jet engines can be identified – one mechanical due to rotating components moving at high speeds, the other due to aerodynamical effects. The latter is also called aeroacoustic noise, which is the whistle that one hears when a high-pressure valve is opened, or the noise that the air makes when a car window is opened while moving at high speed on the highway. While the aviation community successfully reduced noise emissions emanating from a jet engine's fan and turbine section, reduction of noise caused by the mechanical components has nearly reached its limit. New methods to quiet an engine are being pursued actively by the jet engine manufacturers, NASA, and the academic community.

Scientists George S. Constantinescu and Sanjiva K. Lele of the Center for Turbulence Research (CTR) are continuing work on a computational model for predicting jet exhaust behavior using the NAS Facility's high-end computing resources. Their goal is to reduce the intensity of acoustic noise sources in a jet's exhaust that in turn will reduce engine noise.

Investigating Sound Attenuation Through Numerical Computations

Before researchers can offer solutions for reducing the sound emanating from a jet's exhaust, they must first understand the mechanisms that are responsible for noise generation. To accomplish this, researchers calculated the spatial and temporal distribution of the acoustic sources, a task that is "practically impossible to perform using experimental measurements," says Lele. These acoustic sources emanate from the interactions of the turbulent eddies in the flow of the jet.

"The LES method allows Constantinescu and Lele to accurately capture the distribution of acoustic sources and radiated sound waves in the jet near field. These waves are a direct result of the sound emitted by the acoustic sources in the jet."

The flow of fluids is characterized by a non-dimensional number, called the Reynolds number, which is the ratio of the inertia forces over the viscous forces. When the velocity of a fluid flow is small, the viscous forces are large compared to the inertia forces, thus the Reynolds number is small. A fluid flow in this regime can be observed to be regular and is characterized as laminar (or smooth). At high speeds, the Reynolds number is large and the observed fluid flow consists of eddying motions that are chaotic. The fluid flow in this regime is called turbulent. As the Reynolds number increases, the range of the size of the eddies also increases. Two simulation methods that are actively being pursued by NAS and

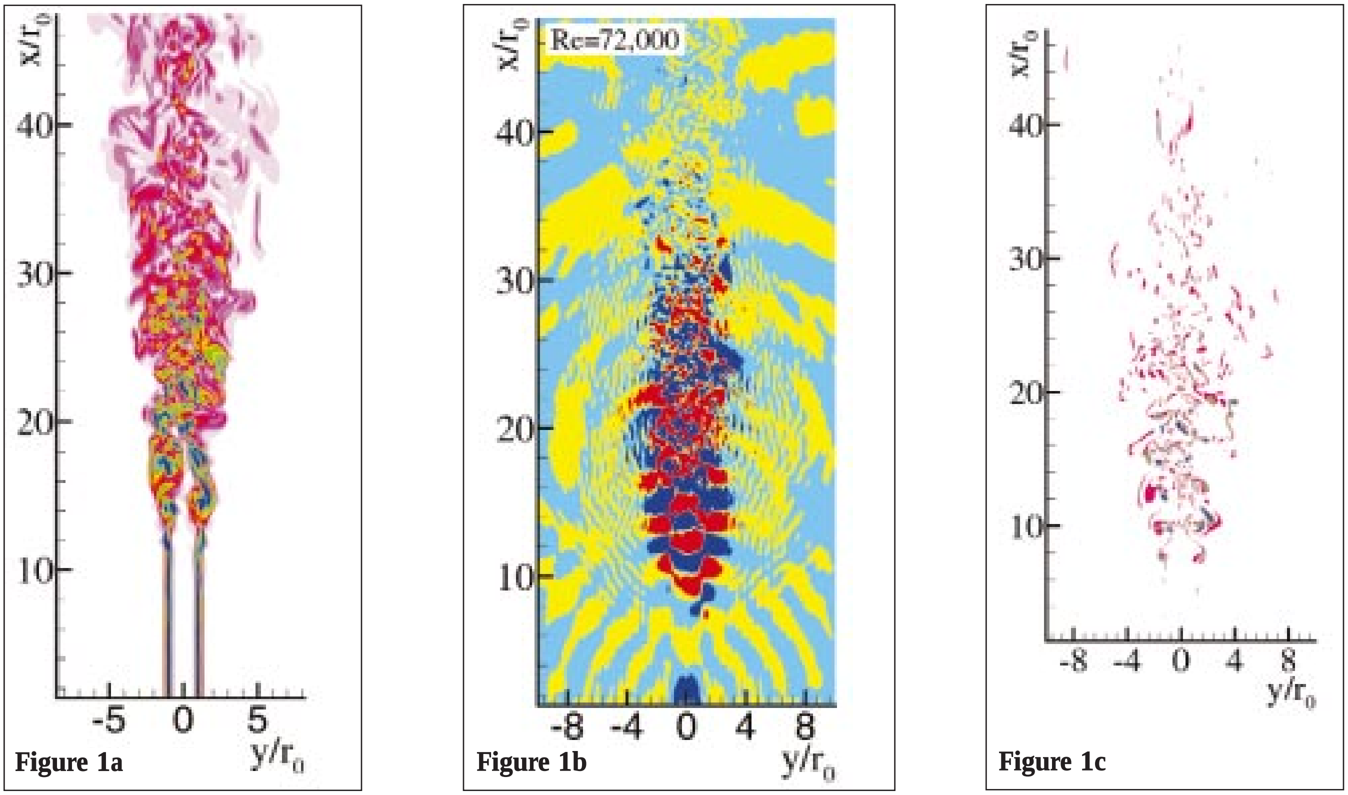

Figure 1: *Visualization of the turbulent jet at a Reynolds number of 72,000, and its radiated noise in the computational domain. Figure 1a: Instantaneous contours of vorticity magnitude, which are used to visualize the turbulent jet; Figure 1b: Instantaneous contours of dilatation contours show how the sound radiates from the centerline away from the jet; and Figure 1c: Instantaneous distribution of the Lighthill sound sources, which visualize the intensity of the sound sources. (Constantinescu)*

CTR scientists are the Direct Numerical Simulation (DNS) technique and the Large Eddy Simulation (LES) technique. In DNS all the scales of motion are resolved, which limits the simulations to moderate/low Reynolds numbers. Using LES, the energy containing large eddies are resolved and the effects of small eddies on the large eddies are modeled, which enable high Reynolds number simulations.

To determine where and how exhaust noise is generated, the acoustic group at CTR, led by professors Lele and Parviz Moin of Stanford University, have used both LES and DNS techniques to simulate noise generation from turbulent flows. Constantinescu and Lele have been employing the LES technique to accurately simulate turbulent eddies in jet exhaust. Future use of DNS for high-speed jets will be ideal, but, in the foreseeable future, will be limited to moderate Reynolds numbers.

Working with Moin and Lele in 1999, Jonathan B. Freund, currently a professor at the University of Illinois Urbana-Champaign, used DNS calculations to develop an understanding of the sound generation in exhaust jets. In his simulations, Freund was only able to simulate jet exhaust turbulence at a Reynolds number of 3,600 by using 25 million gridpoints – a very expensive simulation. DNS simulations at Reynolds numbers close to one million, which correspond to the actual operating conditions of jet engines, will be out of reach for supercomputers in the foreseeable future. Constantinescu and Lele's computational modeling may be seen as an extension of Freund's work for high-speed (high Reynolds numbers) jets.

In Constantinescu and Lele's computational model, they omitted the jet engine's exhaust nozzle. This was done primarily because modeling the nozzle would be too expensive. To compensate, the researchers imposed a velocity profile into the calculations similar to that measured at the jet's nozzle.

Center for Turbulence Research

The Center for Turbulence Research (CTR) is a joint NASA/Stanford University center that sponsors the work of post-doctoral fellows studying turbulence research. This includes studies of turbulence modeling, combustion modeling, turbulence control, aero-acoustics and, in general, research on flows where turbulence plays a dominant role. The center has a facility at Stanford and one at Ames Research Center (ARC). Co-located in the same building at ARC are civil service scientists from the Physics Simulation and Modeling Office of the NAS Division who collaborate with CTR post-doctoral fellows and visitors.

Figure 2a

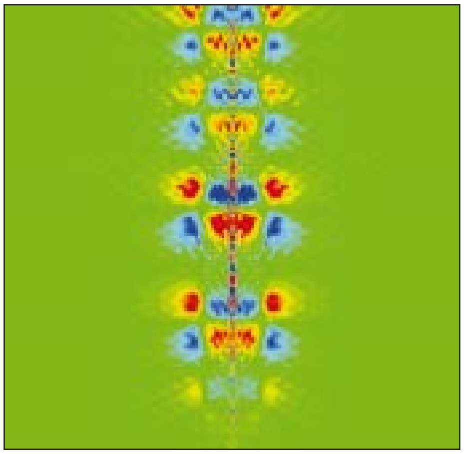

Figure 2b

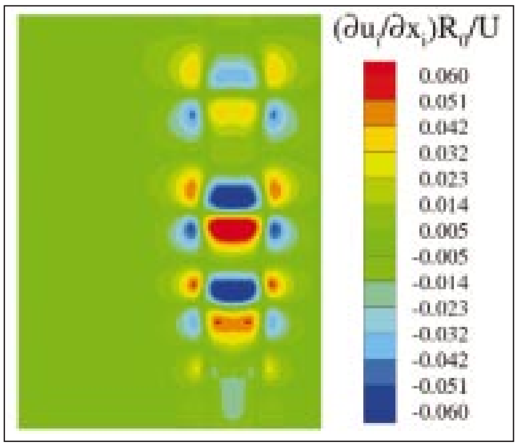

Figure 2: *Dilatation contours in a forced laminar jet. In the original method, Figure 2a, the equations at the centerline were solved in a Cartesian Coordinate system. In the present method, Figure 2b, a new set of equations at the centerline was obtained using series expansions. Calculations in Figure 2b maintain the same order of accuracy of the solution at the centerline as in the rest of the computational domain.*

(Constantinescu)

In addition, most of the noise is produced at the end of the exhaust flow – away from the nozzle, where the exhaust transitions from laminar, or smooth flow, into a turbulent flow.

"Because we are using a subgrid scale model we want to show that the model's contribution to the noise generation is very small and doesn't contaminate the sound field," Constantinescu says. "The purpose of doing these calculations is not to calculate the aerodynamic field. What we really want to do is capture the sound that is emitted by a propulsive jet. In order to do that, we must employ a very accurate numerical method.

"Sound is emitted through very small pressure fluctuations, which are two orders of magnitude less than the pressure fluctuations associated with the jet turbulence. If the goal is only to compute the average velocity field and the mean quantities that characterize turbulence in the jet, it can be accomplished using a second-order method. But for aero-acoustic applications, we really need a very accurate method. In our case, we chose to employ a six-order accurate method based on Padé schemes in the streamwise and radial directions, and Fast Fourier Transform methods in the azimuthal direction," Constantinescu explains. The higher fidelity of the six-order accurate method more efficiently represents the jet's turbulent eddies with a minimum of artificial or numerical viscosity.

More Information About LES

Research papers describing the Large Eddy Simulation techniques employed by Constantinescu and Lele can be viewed at the Center for Turbulence Research website: http://ctr.stanford.edu/publications.html

Solving The Singularity Problem

The LES method allows Constantinescu and Lele to accurately capture the distribution of acoustic sources and radiated sound waves in the jet near field. These waves are a direct result of the sound emitted by the acoustic sources in the jet. However, before successfully performing these simulations, an accurate treatment of the jet centerline had to be developed. When the researchers attempted to perform LES at high Reynolds numbers on relatively coarse meshes compared to DNS, they ran into convergence problems, which caused the code to diverge. The origin of the problem is a difficulty encountered by almost all numerical methods used to solve the governing equations in cylindrical coordinates – the flow equations are discretized in cylindrical coordinates, which is the natural coordinate system chosen to simulate flows that are nearly axisymmetric (the jet axis, or the centerline, plays this role in the jet simulations). However, in the cylindrical coordinate system, the governing flow equations are multi-valued at the centerline (known as the singularity problem). A special mathematical treatment had to be used on the centerline and in its vicinity.

Although Constantinescu and Lele tried to employ several methods previously proposed by other researchers to solve the singularity problem, they were not entirely satisfied with the results. "Many people have tried different approaches to deal with coordinate singularity, and what George and I have done is develop a rather nice way of dealing with the coordinate singularity in the context of high-order numerical methods," Lele says. "We are dealing with a flow that has a nearly circular symmetry, at least in an average sense. When a jet comes out of a circular nozzle, it has that statistical symmetry, but the eddies in the instantaneous flow are three-dimensional."

To resolve the singularity problem, the researchers developed a more accurate algorithm to solve the governing flow equations at the centerline. In this method, a new set of equations, valid at the centerline, was derived using the most general

Figure 3a

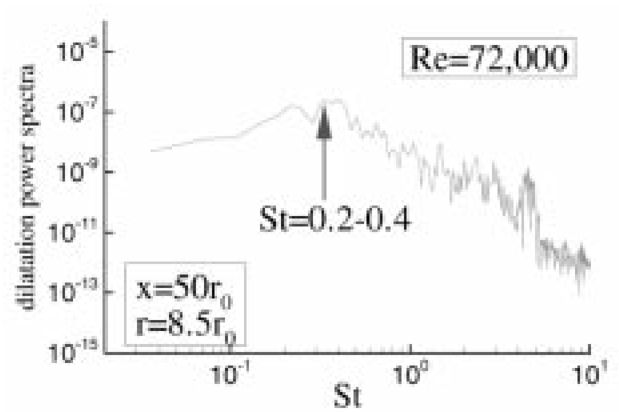

Figure 3b

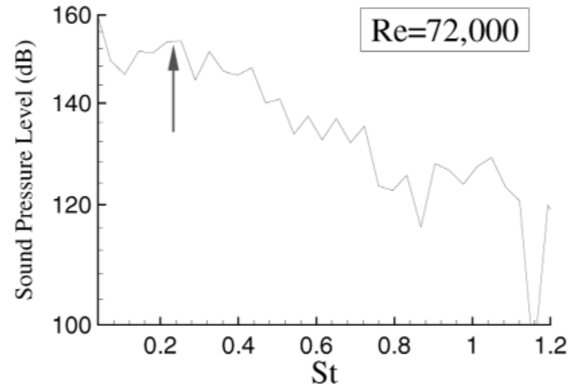

Figure 3: *Sound power spectrum in the acoustic near field of the jet at a Reynolds number of 72,000. Figure 3a presents the dilatation power spectrum with the arrow indicating the frequency at which the spectrum peaks. Experimental measurements show that the peak in the power spectrum is obtained for a non-dimensional frequency (Strouhal number) of 0.2 to 0.4. Figure 3b shows the same information represented in decibels. (Constantinescu)*

series expansions of the flow variables near the centerline. This maintains the same order of accuracy of the solution at the centerline as in the rest of the computational domain. Running a parallel code on the NAS Facility's Origin 2800 computers, employing 32 processors, Constantinescu and Lele can now run LES simulations at much higher Reynolds numbers (100,000) on coarser grids using only four million gridpoints. This has enabled the scientists to more accurately estimate the amount of sound generated by high-speed engine exhaust jets.

Continuing Research

Using the data provided by Constantinescu and Lele's simulations, engineers will be able to introduce disturbances with a specific frequency spectrum into the exhaust, thereby attenuating the intensity of the noise. This will result in a significant reduction of the sound radiated by the jets. Finding ways to force the jet to attenuate noise will be a challenging task. Constantinescu predicts it will take at least two more years of research to arrive at a point where scientists can physically control noise at Reynolds numbers corresponding to actual operating conditions. A parallel effort will be directed toward developing LES subgrid models that are suitable for noise calculations in high speed flows. These models should not introduce spurious noise that can destroy the directly simulated noise field from LES.

As computer modeling capabilities expand during the next few years, an additional goal of the researchers will be to simulate the presence of the jet nozzle . This will allow a more complete understanding of the flow physics and enable researchers to propose more efficient ways of controlling the noise.

—Nicholas A. Veronico

About The Researchers

George S. Constantinescu received his master of science degree in environmental engineering at the Civil Engineering Institute, Bucharest, Romania, and his doctorate from the Iowa Institute of Hydraulic Research (University of Iowa). He served as a postdoctoral associate in both the Department of Mechanical and Aerospace Engineering at Arizona State University and at the Center of Turbulence Research, Stanford University. Constantinescu has recently moved to Stanford's Center for Integrated Turbulence Simulations where he works on numerical simulations of gas turbine engines.

Sanjiva K. Lele is an associate professor with joint appointments in the Departments of Aeronautics and Astronautics and Mechanical Engineering at Stanford University. He received his bachelor of science degree from the Indian Institute of Technology, Kanpur, India, and his doctorate from Cornell University, both in Mechanical Engineering. Lele worked at NASA Ames from 1986 to 1990, then he joined the staff at Stanford. He has received the Francois N. Frenkiel award from the American Physical Society, Division of Fluid Dynamics, and the National Science Foundation's Presidential Young Investigator award. Last year he received best paper awards from both the American Society of Mechanical Engineers and the American Institute of Aeronautics and Astronautics.

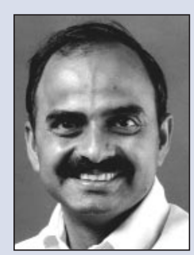

Public Key Infrastructure: Get a Passport to Grid Country

Members of the Information Power Grid team recently established a system to issue certificates (passports) giving users access to grid computer resources.

Every time you fly on a plane, drive a car, or cross the border into another country, you are required, by law, to carry a valid form of identification. Whether it's a driver's license to drive, or a passport to cross the border, these forms of identification can only be obtained by following a predefined set of instructions. Like the traveler who must fill out a passport application, users in the Information Power Grid (IPG) community must obtain a valid certificate to use the grid's computer resources. Members of the IPG team in the NAS Division have recently established their own trusted Certificate Authority (CA), enabling users to obtain certificates through a web browser or the "command line process," a new and improved procedure for requesting, processing, and issuing IPG certificates.

> ***"...Members of the grid community are currently developing a standard policy for the certificate itself that will apply to all grids."***
>
> *— Sumit Talwar*
> *IPG security team lead*

"Getting an IPG certificate is now very easy, fast, and efficient. An IPG certificate script has been modified to make the certificate process user-friendly and reliable," explains IPG certificate authority developer Jana Nguyen of the NAS Division. The new certificate request process takes less than 24 hours and requires users to follow only a few easy steps, thanks, in part, to an automated e-mail notification system and a customizable certificate authority management system (see *Getting a Certificate Using the Command Line Process*, page 11). Once users have an IPG account on the systems they plan to use for their job, a public and private key (a pair of keys that are used to encrypt and decrypt a message), and a valid certificate, they can then access the IPG.

Managing the Process

To help run an in-house certificate authority, the IPG team chose to use Netscape's Certificate Management System (CMS). They adapted the CMS with additional scripts from the National Computational Science Alliance (NCSA), which enabled the team to customize the certificate request and issuing process, and integrate it into the IPG environment. The Netscape software, built to assist with the certificate authority duties, is invisible to the user. "There is no real standard for designing a certificate authority, but members of the grid community are currently developing a standard policy for the certificate itself that will apply to all grids," explains Sumit Talwar, IPG security team lead.

Building Blocks

The first authority to issue certificates to grid users was the Globus CA at Argonne National Laboratory in Illinois. Because Argonne is a research organization, not a production facility for issuing grid certificates, other grid organizations have been encouraged to design their own certificate authority processes. "Developing our own certificate authority helps us build up our own in-house expertise in the security area of public key infrastructure – which includes a certificate management system to support the public key algorithm. This is good for Argonne and for us – it is a win-win situation," says NAS computer scientist, Y.K. Liu. NCSA and San Diego Supercomputing Center have also set up their own certificate authorities using the Netscape CMS.

Initially, the IPG team created a web browser interface, which would enable IPG users on all platform types – Unix, Mac, or PC, to request and receive certificates. The Globus interface developed at Argonne did not offer a web browser method, making it more difficult for PC and Mac users. "With the IPG certificate authority, you can use any platform – that was the team's goal," says Talwar. "Granted, it is

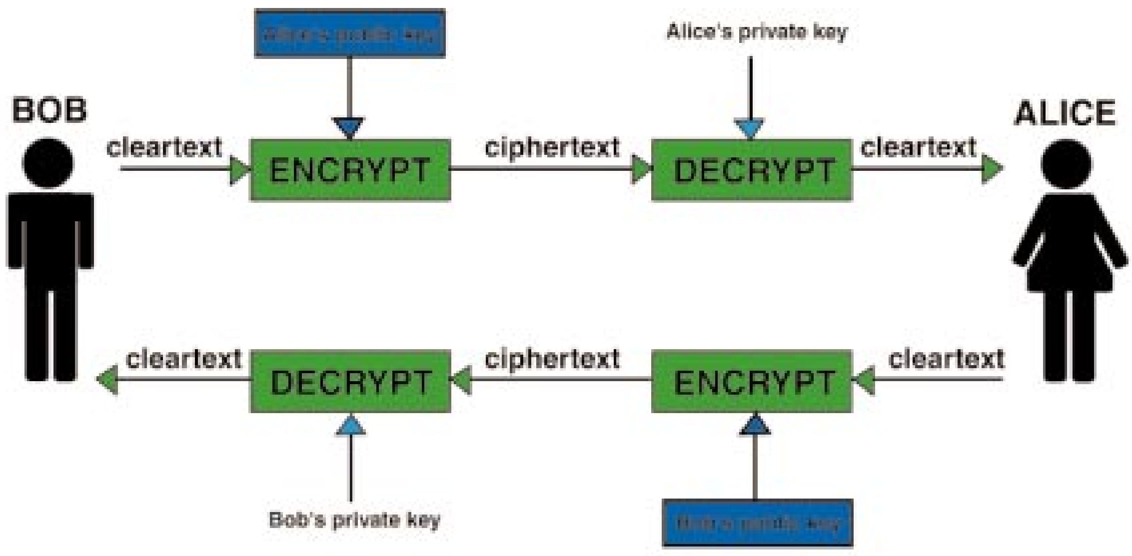

A simplified illustration of the public-key algorithm: When Bob wants to send an encrypted message to Alice, he will use Alice's public key to encrypt the message. After Alice receives the encrypted message from Bob, she will use her private key to decrypt his message. Conversely, if Alice wants to send an encrypted message to Bob she will use his public key.

a little more difficult to go through the web browser than it is to go through a command line, but this way users running on any platform type can request a certificate." The IPG certificate authority requires no additional software to request a certificate using the web method (only a web browser such as Internet Explorer version 5.5 or higher, or Netscape version 4.7 or later).

Although the web method is compatible with any platform type, it requires the user to navigate through nearly 20 steps,

Getting a Certificate Using the Command Line Process

All users of a Public Key Infrastructure must have a registered identity. These identities are stored in a digital format known as a public key certificate. Certificate Authorities (CAs) represent the people, processes, and tools to create digital certificates that securely bind the names of users to their public keys. Receiving a certificate from the Information Power Grid (IPG) CA is similar to being issued a driver's license from the Department of Motor Vehicles. To apply for a certificate, users must first be authenticated (the process of verifying a person's registered identification) and have an account on an IPG machine. Once users have met the criteria, they can request a certificate from the IPG CA using either the web browser method, or the new command line process. Below is an outline of the steps users must follow to receive a certificate using the command line:

1. The user requests a certificate by executing a script through the command line on one of the IPG systems.

2. This script sends the user's request to the CA system, which in turn generates an e-mail that is sent to the Certificate Authority administrator, notifying him or her of the request.

3. The CA administrator then decides whether or not to approve the request based on a positive identification (authentication) of the user.

4. If approved, the Certificate Authority software issues a certificate that includes both the encrypted digital signature of the CA and an assigned certificate identification number. "Signing" the certificate automatically generates and sends an e-mail notifying the user that their certificate is available for retrieval.

5. The user then executes a second command line script, "ipg-cert-retrieve" with the assigned identification number. This script automatically stores the "signed" certificate and sends e-mail to the IPG staff at four NASA Research Centers (Ames, Glenn, Jet Propulsion Laboratory, and Langley) requesting them to make an entry in the grid-mapfile, which identifies them as a grid user.

For security measures, if a certificate has not been retrieved within a certain amount of time, the CA administrator will check with the user to see if they received the notification, or why they have not retrieved their certificate.

Right: *The integration of IPG certificate scripts with three important systems (X.500 Directory Database, Netscape Certificate System, and Globus System) makes it easy for the user to request and retrieve a certificate. The scripts retrieve the user's information from the NASA X.500 directory database, connect to the Netscape Certificate System to request a certificate and, once the certificate request has been approved, sets up the certificate and the private key in the Globus system.*

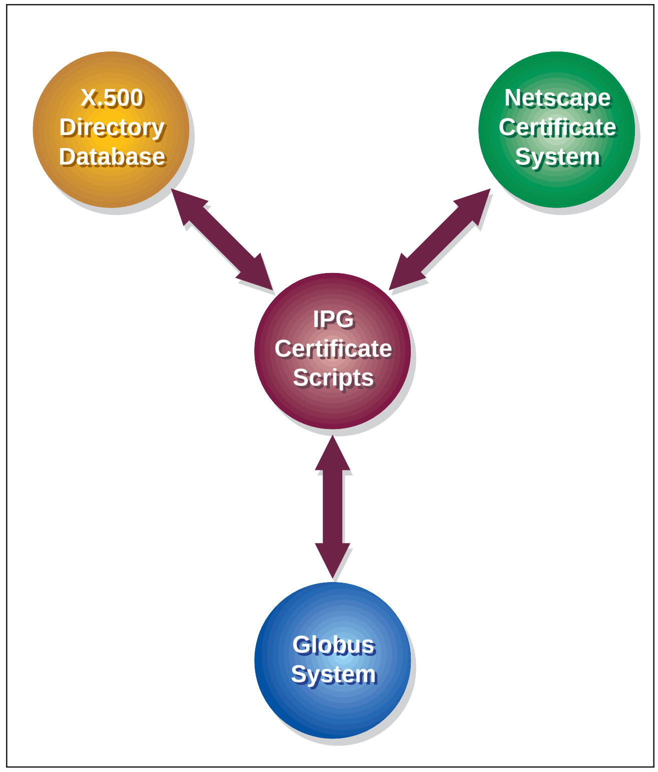

whereas the new command line method is completed in just five. "The original reason for going the web route was that our team thought it would be easier for the user," explains George Myers, NAS scientific consultant and portal development group lead. "Everyone has the idea that the web makes everything easier. In practice, we quickly realized that in this particular case, it was not true."

Command Line Features

Customized security features of the IPG CA include a script that searches the NASA X.500 directory database (an international standard that defines the format to store user's information) to confirm a user's identification. "We try to do as much security work in-house as possible – we want to limit the amount of security outsourcing," explains Liu. The team also developed a script to run an automated e-mail notification system between the IPG certificate authority, the Globus support staff, and the user. "The script streamlines the certificate request and retrieval process," says Liu.

> ***"It was very easy to get the certificate information into the files through the command line. This process was done automatically by script – much easier than through the web browser which requires you to manually manipulate files."***
>
> *— Scott D. Thomas*
> *NASA Ames programer*

For added security, certificates are only good for one year, and users must obtain a proxy certificate each time they use a grid resource (see "Infrastructure For NASA's Information Power Grid Nears Completion," *Gridpoints*, Summer 2000, page 18). The researcher has the freedom to set the number of hours the proxy certificate is valid for – the default being eight hours. Each user's IPG access session automatically closes at the end of the number of hours stated in the proxy certificate. "The proxy certificate is like a temporary pass or visa – if your paperwork or certificate is not on file, you will be denied access to the IPG resources," explains Liu.

Certificates for Everyone

In addition to users at Ames, NASA IPG partners at both NASA Langley (Virginia), and NASA Glenn (Ohio) Research Centers have successfully obtained certificates through the IPG Certificate Authority. Eventually, other NASA centers, grid organizations, and universities could be authorized to obtain certificates from the IPG Certificate Authority as well.

Since the end of February, about 50 users have followed the command line procedures to obtain their IPG certificate. According to one user, Scott D. Thomas, a scientific programmer at Ames, "It was very easy to get the certificate information into the files through the command line. This process was done automatically by script – much easier than through the web browser which requires you to manually manipulate files."

The team will continue to make software enhancements and add more features to increase robustness of the IPG CA. Says Talwar, "We encourage users to go through the IPG website or command line interfaces to get a passport (certificate) for accessing valuable IPG resources to facilitate their research."

— Holly A. Amundson

Certificate Authority Resources

For additional information on NASA's Information Power Grid Certificate Authority, visit the "New User's" section on the website at: www.ipg.nasa.gov/

Information on obtaining an account can be found at: www.ipg.nasa.gov/usersupport/newusers/accounts.htm

Netscape offers a website detailing its Certificate Management System: http://home.netscape.com/cms/v4.0

NAS Technical Training Seminars

A number of scientists have presented their research to audiences at the NAS Facility during the first half of 2001. Many of this year's presentations have been videotaped, and can be borrowed by sending e-mail to the NAS Documentation Center (doc-center@nas.nasa.gov). Five of these seminars, all of which were videotaped, are highlighted below:

• Terry Holst and Tom Pulliam of the NAS Division's Applications Branch presented "Genetic Algorithms Applied to Aerodynamic Shape Optimization" on May 15. They discussed a method for aerodynamic shape optimization using a genetic algorithm with real number encoding. This algorithm was demonstrated on three different problems: a simple hill-climbing problem; a two-dimensional airfoil problem using an Euler/ Navier-Stokes equation solver; and a three-dimensional trans-sonic wing problem using a nonlinear potential solver. Holst and Pulliam noted that although the emphasis is on aerodynamic shape optimization, the genetic algorithm presented is quite flexible and could be applied with minimal implementation effort to many single- or multi-objective problems in other fields, providing the means of evaluating the objective function exists.

• Stanford University's Charles Pierce discussed his research on the "Large Eddy Simulation of a Coaxial Jet Combustor" on March 29. Pierce reviewed the large eddy simulations that were conducted to test the performance of a new chemistry model in a methane-fueled coaxial jet combustor for which experimental data is available. Pierce's chemistry model utilizes transport equations for a mixture fraction variable, which tracks the mixing of fuel and oxidizer, and a progress variable that tracks the degree-of-reaction at each point in the flowfield.

• Frank Werblin, Department of Molecular and Cell Biology at the University of California, Berkeley, presented "A Visual Language of the Retina: A Dozen Simultaneous Movies" on March 27. Werblin discussed how the retina generates a parallel stack of at least a dozen different representations of the visual world, and how these representations, surprisingly sparse in both space and time – differ in their space-time filtering properties. They are formed through an intricate interaction between numerous excitatory and inhibitory patterns, and each selects for a different set of visual features.

• San Jose State University's James Wayman spoke on "Biometric Identification: Teaching Computers to Recognize People" on March 20. Biometric identification is the automatic identification or identity verification of living, human individuals based on behavioral and physiological characteristics. Technologies include fingerprint, hand, and voice recognition, as well as face and iris recognition. Customers for this technology have ranged from Disney World to the Super Bowl. Wayman discussed a scientific approach to analyzing applications and matching them to technologies, along with algorithms and mathematical models for predicting system performance based on experimental estimations of model parameters.

• On March 15, Suhrit Dey, from Eastern Illinois University and recently a visiting scientist at Ames, presented "Hemodynamics of T-cells with Applications to Breast Cancer." Dey discussed the role of white blood cells in fighting cancerous cells, noting that they are contained in less than one percent of blood's volume. T-cells originate from bone marrow, circulate through blood to the thymus (a lymphoid organ located inside the chest cavity above the breast and in the space between the lungs), and exit the thymus as "fighter" cells. Dey is studying the hemodynamics of breast cancer, and has found that a strong circulation of T-cells should help prevent its occurrence. In his talk, Dey discussed some things women can do to stimulate the thymus, increase the circulation of T-cells, and lessen the chance of contracting breast cancer. The next phase of Dey's research will involve studying the biomechanics of T-cells by applying computational models to explain surveillance, mobilization, and defense strategy of these lymphocytes.

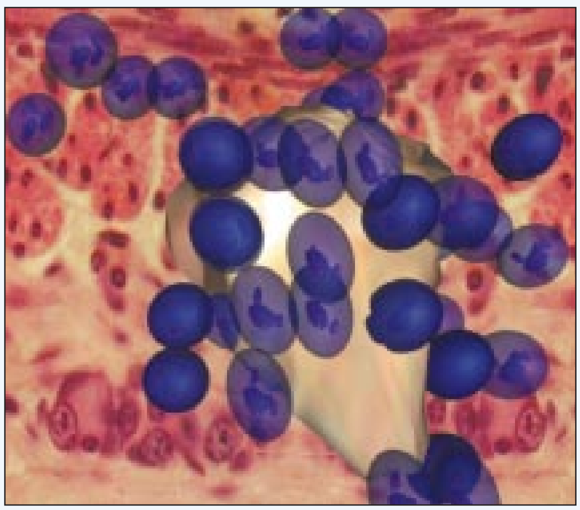

A visualization from Suhrit Dey's research (see March 15) shows a large, cancerous white blood cell being overpowered by relatively small (purple) T-cells.
(Visualization by Cliff Williams)

For information on these and other NAS technical training events, contact Marcia Redmond at: mredmond@mail.arc.nasa.gov. For information on future training sessions, see the *Gridpoints* Calendar of Events (page 21), or visit the NAS Technical Training Event Web site:
www.nas.nasa.gov/User/Training/training.html

IPG Supports National Airspace Simulation System Demonstrations

In a recent collaboration between the NASA Ames' Computational Sciences Division and members of the IPG team in the NAS Division, grid resources were employed to support simulations of commercial air traffic crossing the United States.

Since the outset of its development in fall 1998, NASA's Information Power Grid (IPG) has been designed to enable interdisciplinary collaborations. The IPG's distributed network of scientific instruments, data archives, and computational resources was recently applied to the area of aviation safety. With the help and collaboration of NAS Division IPG team members, the Aerospace Extranet group in the Computational Sciences Division at Ames Research Center employed grid resources to simulate commercial air traffic flow across the United States.

"The Aerospace Extranet Program demo was a major success for the IPG – it demonstrates the true capabilities of the grid," says Johnny Chang, NAS scientific consultant. By helping the Aerospace Extranet group, the IPG team gained valuable insights about the process of integrating new grid users. "We learned a good deal about the requirements IPG users are likely to need as a result of this demonstration," says Judith Utley, NAS Globus administrator. "We will now be able to respond faster and more smoothly to new IPG users and bring in new systems as well as more NASA centers."

> "Once all installations and modifications were complete, a seamless environment was provided between *Washington*, IPG systems at NAS, and an IPG system at Glenn."
>
> *— Judith Utley*
> *Globus Administrator*

On April 9 and 12, the Aerospace Extranet group (led by David Maluf, with William J. McDermott, principal investigator) used the IPG to demonstrate near-real-time, distributed modeling and simulation capabilities of the National Airspace Simulation System. The demonstrations were accomplished using a Sun Ultra Enterprise 450 computer, known as *Washington*, located within the Computational Sciences Division. *Washington* was connected to the IPG as a fully integrated client, allowing the Aerospace Extranet team to use *Washington*, as well as three IPG resources, *Evelyn* and *Turing* (both SGI Origin 2000s) at the NAS Facility, and *Sharp* (also an Origin 2000) at NASA Glenn Research Center, Cleveland, Ohio.

To enable the Aerospace Extranet group's application to send and receive data between *Washington* and the IPG, the Globus toolkit (software designed to enable different types of grid resources to interface with each other) was installed. Globus administrators Judith Utley and Mary Hultquist assisted Matt Linton (system administrator for *Washington*) with the installation, configuration, and testing of the Globus software package. Using Java applications, the Aerospace Extranet group created their own web interface on *Washington* to initiate Globus commands. Globus enables the team to interact with their simulation as it runs on the IPG resources.

To obtain user access to the grid, the IPG certificate authority software (see *Getting a Passport to Grid Country,* page 10) was also installed on *Washington*. Researchers relied heavily on the security of the IPG because of the sensitive nature of the simulation data. The group's massive datasets required resources that could simultaneously run multiple jobs using a high number of computational cycles – exactly what the IPG was built to provide.

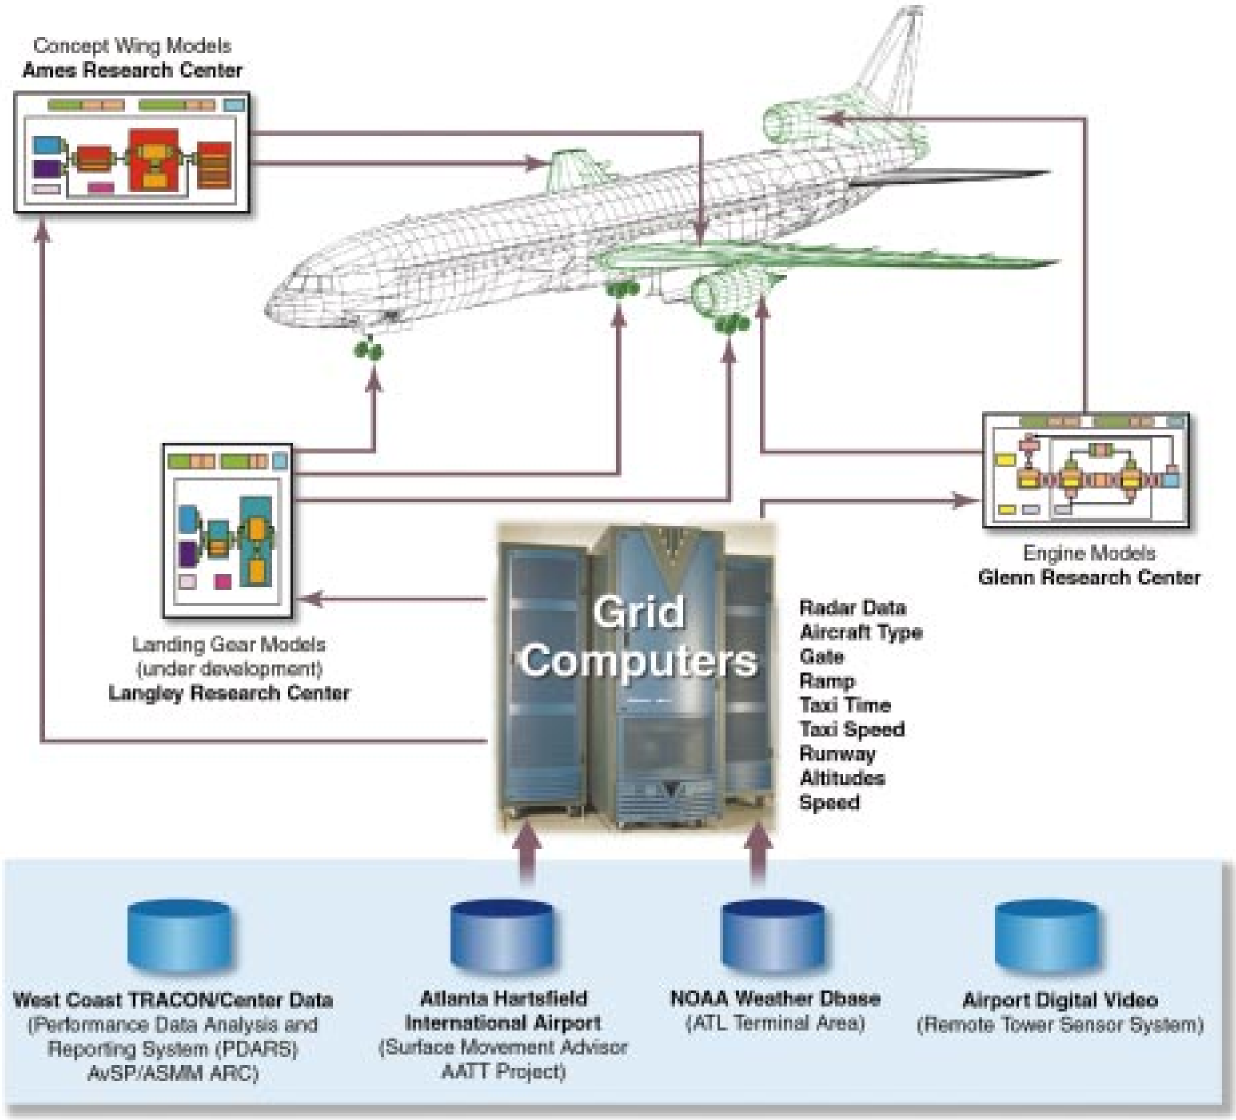

During the demonstration of the prototype National Airspace Simulation System, jobs were submitted from *Washington*, where the database files originated, and distributed to *Evelyn*, *Turing*, and *Sharp*. The output files, or engine simulation parameters, were generated using the Numerical Propulsion System Simulation (NPSS) software installed on *Sharp*. The engine parameters were then sent back to *Washington* where users had the ability to select, manipulate, and graph the output. The data used for the simulations was taken from radar tracks of flights arriving and departing from Atlanta, Georgia's Hartsfield International Airport.

Learn More About the IPG

Additional information about NASA's Information Power Grid (IPG) can be found at: www.ipg.nasa.gov

The Computational Sciences Division at NASA Ames Research Center is on the web at:
http://infotech.arc.nasa.gov:80/compsci.html

Details about the National Airspace System are available at: www.faa.gov/education/resource/National%20Airspace%20System%20(NAS).htm

"Once all installations and modifications were complete, a seamless environment was provided between *Washington*, IPG systems at NAS, and an IPG system at Glenn," explains Utley. "The demos went off without a hitch," adds Chang. Additional IPG team members involved with assisting the Aerospace Extranet group with the demonstrations included: IPG Project Manager Bill Johnston, Deployment and Integration Manager Leigh Ann Tanner, and NAS Scientific Consultants Terry Nelson and Chuck Niggley.

—*Holly A. Amundson*

Editor's note: *The Aerospace Extranet group's simulations are part of the NASA-funded Information Technologies Base Program, and are conducted on behalf of the National Airspace Simulation System to ensure safe and efficient movement of aircraft through the nation's airspace.*

The Power of Unsteady Flow Visualization

New algorithms have enabled NAS data analysis group researchers to visualize extremely large datasets using personal workstations, cutting computation time and costs while increasing resolution.

By Timothy A. Sandstrom and Neal M. Chaderjian

The enormous computational power of the Information Power Grid (IPG – a distributed network of scientific instruments, data archives, and computational resources) is enabling scientists in the NASA Advanced Supercomputing (NAS) Division to carry out flow simulations that were previously unattainable. The dataset size for a typical simulation consists of hundreds of gigabytes, which, if printed out, would create a stack of paper more than one and one-half miles high. The sheer scale of these datasets calls for a new breed of visualization tools that will allow the user to interact with the data in real-time. Commercial off-the-shelf software has difficulty performing real-time interaction with datasets of this size. This situation is further exacerbated by the need to visualize dozens, or even hundreds, of datasets. However, work being done by the NAS Research Branch's data analysis group specifically targets tool development that enables high levels of user interaction with very large datasets.

NAS Applications Branch scientists Neal Chaderjian, Jasim Ahmad, Shishir Pandya, and Scott Murman serve as a recent example of the need for visualization tools that enable users to interactively explore very large, unsteady datasets. The researchers computed 45 time-accurate Navier-Stokes flow simulations of a YAV-8B Harrier jet, which is capable of vertical take-off and landing, hovering above a tarmac with a 33-knot headwind (see figures 1a and 1b, page 17). This Advanced Technology Application (ATA) is part of the Computational Aero Sciences High Performance Computing and Communications Program.

Computational Visualization For Mission Safety

The purpose of the Harrier flow simulations is to demonstrate a technology to generate a stability and control database for an aircraft with complex flow physics, while reducing computational costs. Controlling an aircraft hovering near the ground can be a tricky proposition, as a fatal January 2001 crash of a Harrier demonstrates.

During take-off and landing, the Harrier's thrust vectoring nozzles are directed toward the ground to provide lift needed to support the vehicle in low-speed flight. The vectored nozzle exhaust impacts the ground and interacts with the ambient flow to form a ground vortex. A jet fountain can also form when the jet flow impacting the ground turns up to strike the underside of the fuselage. The overall concern is that these flow features can cause vehicle safety and control problems. For example, hot gas ingestion by the inlet may cause the Harrier to crash due to a rapid loss of engine thrust, which in turn causes a loss of lift. Rapid flow accelerations on the underside of the vehicle can also cause a "suck-down effect" where the vehicle is pulled toward the ground. In addition, there is also a concern with ingestion of ground debris, and for the safety of ground personnel in the vicinity of the flying aircraft. Visualization of the flow field is crucial to understanding how the jet exhaust flow behaves, and for improving the safety of flight operations at low altitudes. Sixteen hundred files consisting of more than 100 Gigabytes of data were used to visualize the unsteady flow for a single Harrier simulation hovering at a height of 30 feet above the tarmac.

Using both Gel and Batchvis, software developed by the data analysis group, NAS researchers were able to discern flow features not readily attainable with traditional software packages such as FAST and Plot3D (see sidebar: *How a New Algorithm Powers Unsteady Flow Visualization*, page 18). The new software enabled the team to choose locations to release virtual smoke particles into the flow by grabbing "rakes" and dragging them around with the mouse. Gel then tracked the

Figure 1a: *The U.S. Marine Corps operates the Harrier fighter jet, capable of vertical take-off and landing. Computational models were used to visualize the air flow about the jet while hovering.* *(Boeing)*

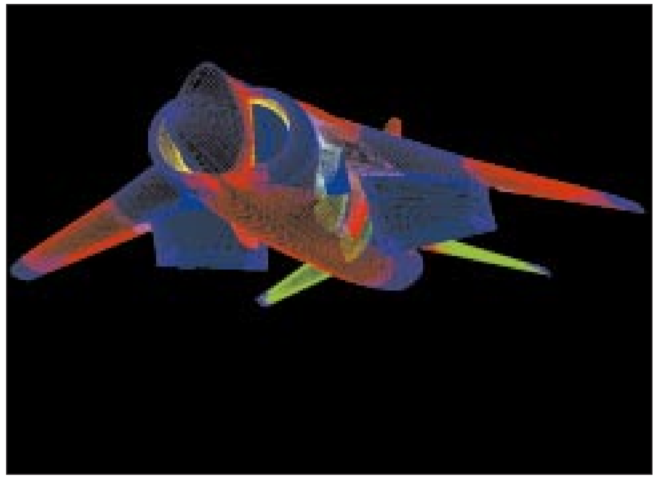

Figure 1b: *A Harrier grid system was used to compute 45 time-accurate Navier-Stokes flow simulations of a YAV-8B Harrier hovering above a tarmac with a 33-knot headwind.* *(Timothy A. Sandstrom)*

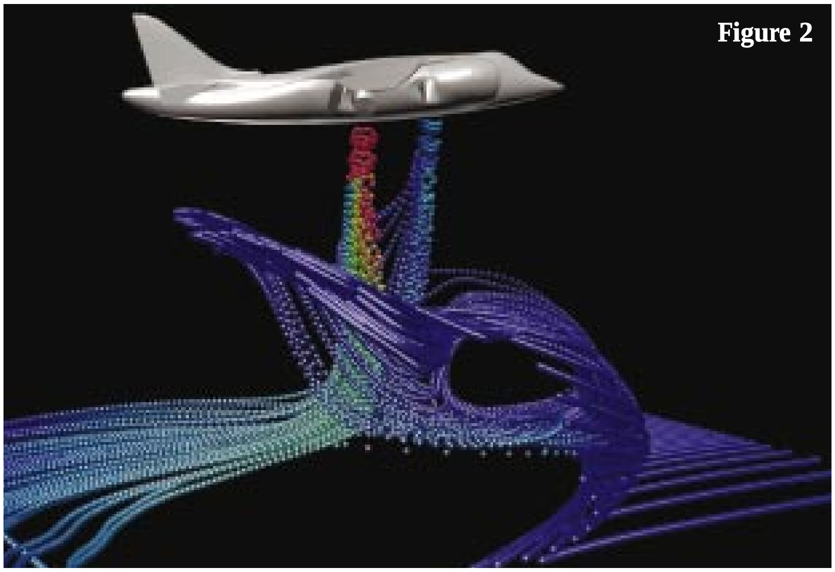

Figure 2 (left) and Figure 3 (bottom left): *Using Gel and Batchvis software, NAS data analysis group researchers are able to visualize the ground vortex and jet exhaust of the Harrier from different angles (Figure 2, side view, and Figure 3, oblique view). The jet's exhaust can form a fountain that strikes the underside of the aircraft. Hot gasses from the jet fountain, if ingested, can cause a rapid loss of engine thrust. The streaklines depicted in these figures are colored by temperature (blue depicts cool, red is hot).* *(Timothy A. Sandstrom)*

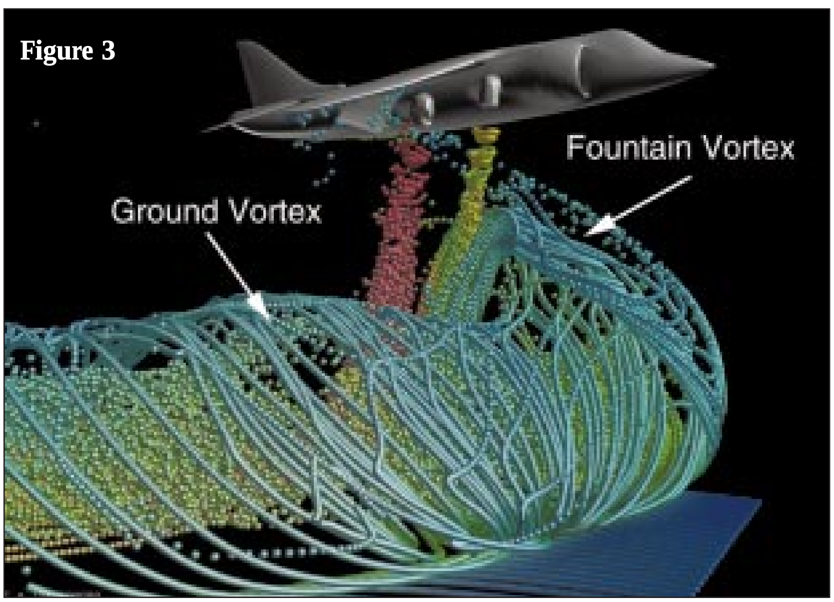

locations of thousands of particles as they moved about the aircraft. A vortex core detection algorithm was also used as a feature detection tool to automatically extract the location of vortical flows. A side view of the Harrier's jet exhausts interacting with the ambient flow to form a ground vortex is shown in Figure 2. The ground vortex is unsteady, periodically changing its size and position with time. The jet exhausts also oscillate back and forth as they interact with the time-varying ground vortex. The frequency of this oscillation was found to correspond precisely with the variation of lift with time. Thus, a direct correlation between the time-varying lift and unsteady ground vortex was established.

Figure 3 depicts an oblique view, where the virtual flow particles are rendered as small spheres rather than screen pixels to highlight the ground vortex. Notice that a fountain vortex

Continued on page 20

How a New Algorithm Powers Unsteady Flow Visualization

By David Ellsworth

The 100 gigabyte Harrier simulation is an example of the large datasets that can be produced by the 512- and 1,024-processor supercomputers at the NAS Facility. Once the datasets have been computed, they are analyzed using visualization software such as Gel and Batchvis running on a researcher's personal workstation. These datasets, however, are difficult to visualize on workstations with traditional visualization applications (such as FAST and Plot3D) because they assume that the entire dataset (or at least a pair of time steps for time-varying data) can be loaded into main memory. Unfortunately, many personal workstations do not have sufficient memory capacity to handle this workload. In addition, most workstations do not have sufficient disk space to hold large datasets, which means the information must be stored on a remote file server.

During the past four years, researchers in the NAS data analysis group have been developing new algorithms that enable the visualization of large datasets on personal workstations. These algorithms are called "out-of-core" visualization techniques because they leave most of the data on disk instead of loading all of it into main memory (previously called "core"). Since many visualizations only examine a small fraction of the dataset at any one time, the currently needed data is often small enough to fit into a workstation's main memory. Reducing the amount of data required is important because it allows the visualization to be computed at the workstation's maximum performance (see chart at right). If the data did not fit in main memory, it would have to be repeatedly retrieved from the disk drives, which are many thousands of times slower. Figures A and B (page 19) show how out-of-core algorithms reduce the data that must be read.

Animation on the web

Images and animation from the NAS data analysis group's research can be downloaded from the web at:

www.nas.nasa.gov/Groups/VisTech

Click on the Images and Movies button for a selection of more than 50 aerospace, biological, physical, and earth and space science visualizations.

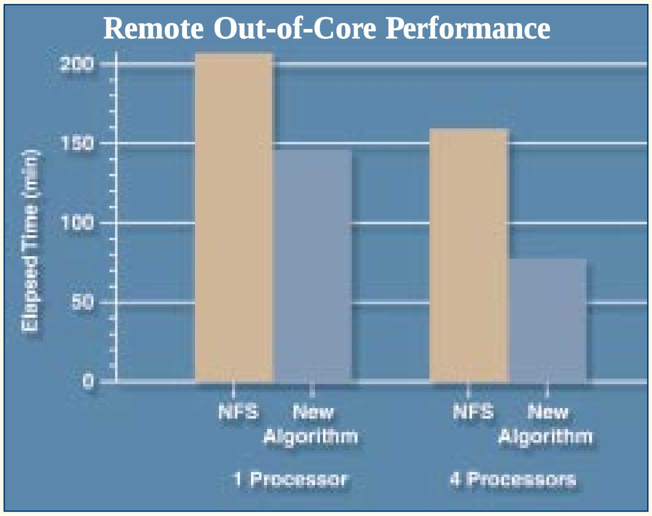

When compared to the existing Network File System (NFS), the new "out-of-core" algorithm dramatically improves the time needed to visualize the 100-gigabyte Harrier simulation.
(David Ellsworth)

Remote visualization, where data is visualized from a remote server, allows numerous researchers to share a dataset without making multiple copies of it. It also allows data to be visualized directly from the machine where it was computed, which makes it much easier to verify that the computation was completed correctly. Finally, remote visualization is practical because Gigabit Ethernet is fast enough to not be a bottleneck, and is now inexpensive enough to be deployed to the desktop.

The original out-of-core visualization application developed at NAS did not perform both the computation and disk access at the same time. Instead, when the program detected that data must be read, the calculation was stopped while the data was read from disk. If the dataset was read from disk on a remote server, the computation was further delayed while the data was transferred over the network. The data analysis group's new algorithm simultaneously overlaps the computation, disk access, and network transfer functions. This is accomplished by dividing the visualization into a number of tasks – when one task must wait for data to be retrieved from disk, a scheduling algorithm runs another task to keep the processor busy. Data retrieval from a local or remote disk

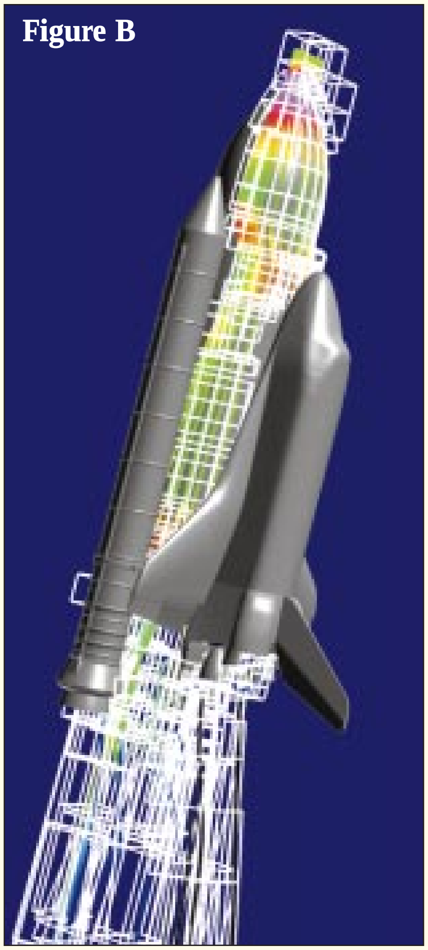

Figure A visualizes the airflow around the external tank of the Space Shuttle launch vehicle (data courtesy Ray Gomez, NASA Johnson Space Center). Figure B displays the outlines of the solution data blocks (in white) that were read by the out-of-core algorithm to compute the visualization. Since the out-of-core library only reads the data needed by the streamline calculation, only eight percent of the 573-megabyte dataset was read in this case. Note that the data blocks are only around and aft of the external tank. This demonstrates that the blocks around the orbiter and solid rocket boosters, as well as those blocks further away, were not read. (David Ellsworth)

proceeds independently while the requesting task waits. The new algorithm is general enough to support a variety of visualization techniques, which must be modified so that the work can be divided into a number of smaller tasks. However, such modifications are similar to those made to enable a visualization to run in parallel.

The chart on page 18 depicts the improved performance of the new out-of-core algorithm when computing a 1,600-frame animation of the Harrier dataset using parameters similar to the ones used in Figure 2 (see page 17). The dataset was read from a remote system over an 800-megabits-per second HIPPI (HIgh Performance Parallel Interface) network connection, and used SGI Onyx workstations for both the local and remote systems. The chart compares the performance results using the standard Network File System protocol to retrieve the data versus performance using the new algorithm. When using one processor, the time decreased from 207 minutes to 146 minutes, or 30 percent. The improvement in performance is attributed to the increased overlap of computation, disk access, and network data transfer. When four local processors were used, compute time decreased by more than half, from 159 to 77 minutes, demonstrating that the algorithm can make use of multiple processors. The computation time did not decrease by a factor of four when the number of processors was increased because the disk and network speed were the same for all visualization runs.

The new out-of-core algorithm will make it easier for researchers to visualize the output of simulations run on IPG systems. In the future, NAS researchers will test the performance of the algorithm using a wide area network to access data from remote IPG systems.

Editor's note: Michael Cox, formerly of the NAS Division, began work on the out-of-core algorithm in 1997. In the ensuing four years, the data analysis group has worked to develop the algorithm into the robust tool that it is today. The author and Tim Sandstrom wrote a number of the applications that use out-of-core technology, while Pat Moran and Chris Henze worked on the Field Encapsulation Library, which provided the framework for the current out-of-core implementation. For more details about the new algorithm, see: www.nas.nasa.gov/Research/Reports/Techreports/2001/nas-01-004-abstract.html

David Ellsworth is a research scientist for Advanced Management Technology, Inc. and is a member of the NAS data analysis group. He received a doctorate in Computer Science from the University of North Carolina at Chapel Hill, and has been researching methods for visualizing very large data sets since joining NAS in 1997.

Continued from page 17
forms in front of the primary ground vortex and is then swallowed up into the primary vortex.

When the Harrier descends to within 10 feet above the tarmac, the fountain vortex is ingested into the inlet (see Figure 4.) This results in the hot gas ingestion phenomenon that leads to a loss of thrust and subsequently lift, which in turn may cause the Harrier to crash.

In Figure 5, a vortical flow structure on the horizontal tail, as viewed from behind the aircraft, is displayed. This unusual vortex system is the result of very low-speed ambient flow being entrained by a complex interaction between the jet fountain coming up from the underside of the fuselage, and a wing flap vortex that convects back over the top of the horizontal tail. This vortex system would not occur under normal high-speed flight conditions.

Figure 4: *The fountain vortex has been isolated in this view to show how it travels up from the tarmac into the engine inlet (spheres colored by temperature). (Timothy A. Sandstrom)*

Expanding Visualization Tools For the Future

Interactive unsteady flow images help scientists and engineers better understand what is happening to the Harrier when close to the ground, and why. Researchers are tempted to use instantaneous streamline tools because this method only requires data at one time step (rather than 1,600 in the present case). However, the resulting static images do not properly capture the dynamic behavior of the unsteady flow, and can give misleading or even wrong impressions of what is actually happening.

NAS data analysis group researchers are working to develop emerging interactive tools, such as Gel, to provide the engineer and scientist with a powerful laboratory to explore flow features and improve the safety of flight operations as well as the working environment of ground personnel.

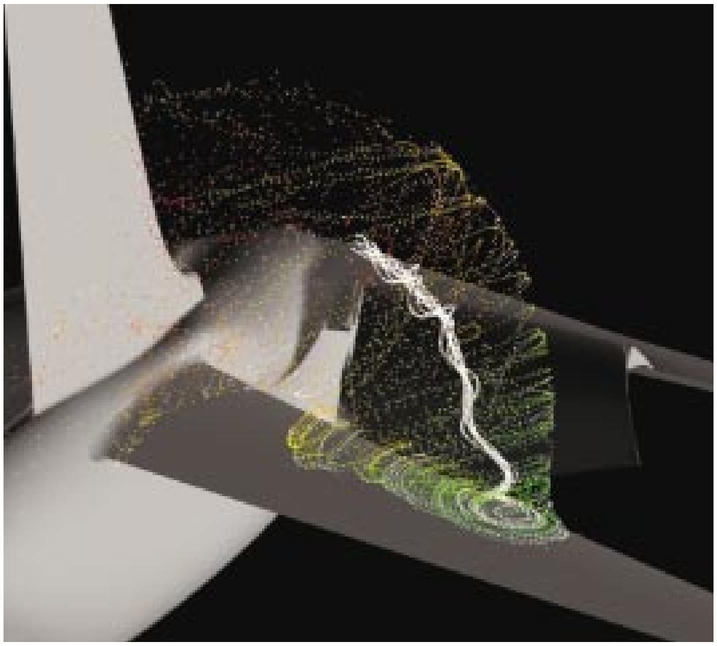

Figure 5: *Tail view of the Harrier showing a vortex (white lines) impacting the horizontal tail and causing an upward swirling flow (particles colored by temperature). This vortex is formed through a complex interaction between the jet fountain and wing flap vortex. (Timothy A. Sandstrom)*

Timothy A. Sandstrom is a senior systems analyst for Advanced Management Technology, Inc. and a member of the NAS data analysis group. He has been writing visualization software for 10 years and was a key member of the team that developed the Flow Analysis Software Toolkit (FAST), which won NASA's Software of the Year award in 1995.

Neal M. Chaderjian is a research scientist in the NAS Applications Branch, and group lead for the Computational Aerosciences Powered-Lift Advanced Technology Application project. He joined NASA after receiving his doctorate from Stanford University. He is currently interested in integrating CFD and IT technologies to enable first-of-a-kind unsteady flow simulations that impact NASA programs. He has received NASA and DOD awards, and is an Associate Fellow of the American Institute of Aeronautics and Astronautics.

Calendar of Events

10th IEEE International Symposium on High Performance Distributed Computing

San Francisco, California • August 7–9

The 10th International Symposium on High Performance Distributed Computing is a forum for presenting the latest research findings on the use of networked systems for high-performance computing. All aspects of high-performance distributed computing will be presented, including, visualization, collaboration, hardware technologies, network protocols, the middleware that ties distributed resources together into "computational, data, and collaboration grids," middleware that enables application use of grids, and tools and languages that support application development. Details are available at: www-itg.lbl.gov/HPDC-10

14th International Conference on Parallel and Distributed Computing Systems

Dallas, Texas • August 8–10

The 2001 International Conference on Parallel and Distributed Computing Systems is a forum for the exchange of knowledge and experience among researchers, engineers, and practitioners working with parallel architectures and systems, as well as distributed computing and information systems. On the web, visit: www.isca-hq.org/confr.htm

SIGGRAPH 2001

Los Angeles, California • August 12–17

SIGGRAPH 2001 showcases applications of computer graphics and interactive techniques developed by computer graphics scientists, artists, engineers, and educators. For additional information, visit: www.siggraph.org/s2001

2001 NATO Research and Technology Office – Short Course: Error Estimation and Solution Adaptive Discretization in CFD

NASA Ames Research Center • September 10–14

The NATO Research and Technology Office will host a four-day workshop on Error Estimation and Solution Adaptive Discretization in CFD at the NAS Facility. Discussions will cover: Delaunay triangulation, computational geometry; Introduction to A-posteriori error estimation, Giles/Pierce theory; Adaptive Finite Element Methods for fluid flow, model adaptivity; Implicit A-posteriori computation of bounds, "Energy" norms and outputs of interest; and A-posteriori error estimation, stability and error control. Visit www.nas.nasa.gov/User/Training/training.html for registration information.

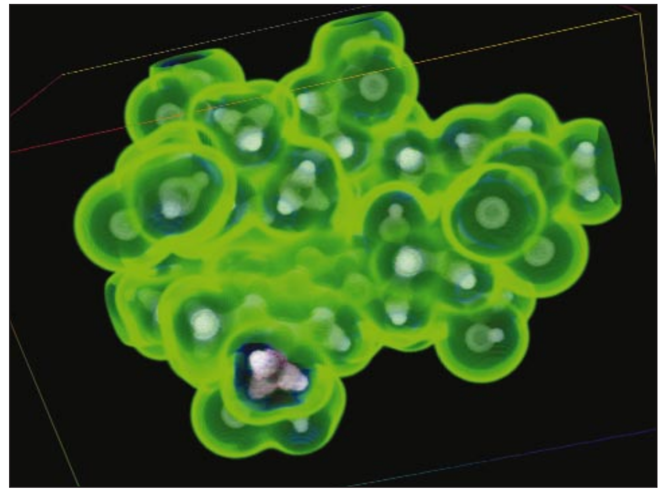

This visualization reveals the arrangement of electrons in a molecule of vitamin B12. The vitamin is used by many enzymes as a cofactor in reactions involving structural rearrangements of various molecules – in other words, the enzyme-vitamin complex serves as a kind of molecular reassembler or "nanomanipulator." This important functionality of vitamin B12 relies on the electronic behavior of a metallic cobalt atom in its core. To see a 3-D animation, visit the NAS website's video collection at: www.nas.nasa.gov/About/Media/videos.html

(visualization by Chris Henze)

Ninth Foresight Conference on Molecular Nanotechnology

Santa Clara, California • November 9–11

The Ninth Foresight Conference on Molecular Nanotechnology will review the advances in nanotechnology and cover topics such as molecular electronics and machines, scanning probes, self-assembly, nano-materials and structures, as well as nanoelectronics and nanodevices. Jie Han of the NAS Division's computational nanotechnology group is one of the featured speakers. Visit: www.foresight.org

SC2001

Denver, Colorado • November 10–16

SC2001 will bring together scientists, engineers, designers, and managers from all areas of high-performance networking and computing, and showcase the latest in systems, applications, and services. The conference website is: www.sc2001.org

2001 Information Power Grid Workshop

NASA Ames Research Center • December 5–6

The NASA Information Power Grid Workshop will be held at the NAS Facility on December 5 and 6. Planned topics include: IPG infrastructure and deployment; grid interoperability middleware; IPG security infrastructure; grid utilization tools; data access and grid computing; grid user interfaces; IPG applications and tools; as well as other related topics of interest to grid computing. For the latest information, visit: www.ipg.nasa.gov or NAS's training site: www.nas.nasa.gov/User/Training/training.html