NASnews
July – August 1998 • Volume 3, Number 4

In this issue...

Where Air Meets Wing: New Software Tool `Draws the Line'

A new algorithm developed by NAS Systems Division researcher David Kenwright detects more flow separation and attachment lines and can dramatically reduce aero engineers' data analysis time.

Stare out the window the next time you're taking off in a passenger jet, and try to imagine the trillions of molecules flung into a ghostly dance as the wings cut through the air. These particles form whirling vortices that brush against the wings, creating drag that, in extreme cases, can lead to a sudden loss of lift. Because problems like this can't be seen, they must be anticipated by aircraft pilots and designers. But even when the aircraft is a simulation inside a computer -- where airflow along its skin can be depicted in bright-colored streamlines -- designers studying these vortices often have a hard time picking out the places where they strike or separate from the wings.

To The Article...

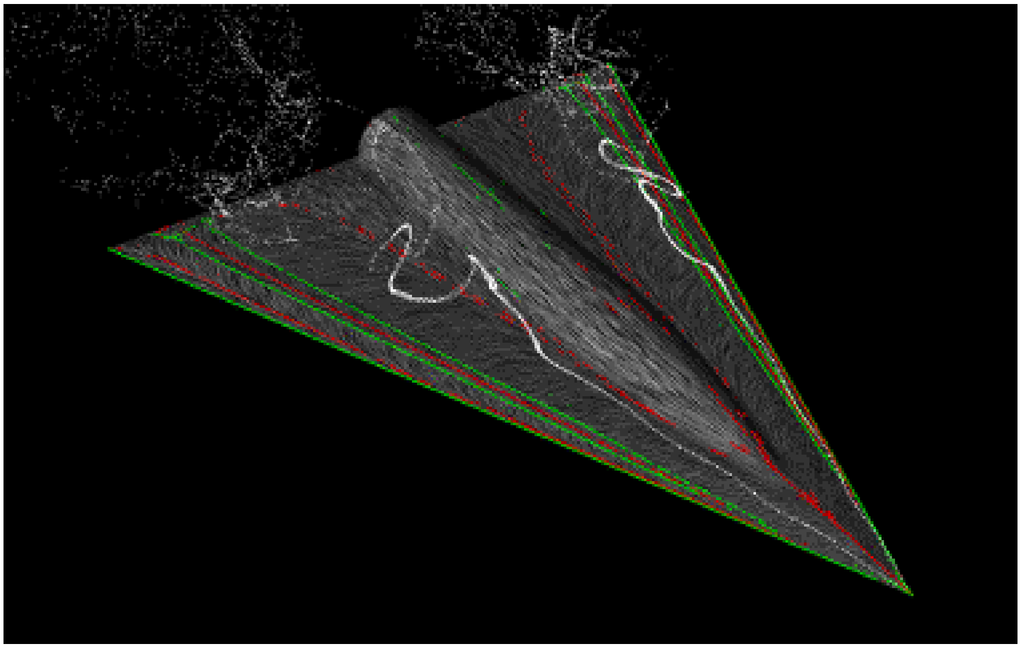

In this visualization of airflow near the surface of a delta wing aircraft, turbulent air separates from the wings along the white lines and reattaches along green lines. Locating these separation and attachment lines in a 13-gigabyte dataset took less than one second using new automatic feature-detection software developed by David Kenwright, a member of the NAS data analysis group. Delta wing dataset by Neal Chaderjin, NASA Ames Applied Computational Aerodynamics Branch. Graphic by David Kenwright.

To The Article...

Where Air Meets Wing: New Software Tool 'Draws the Line'

by Wade Roush

Stare out the window the next time you're taking off in a passenger jet, and try to imagine the trillions of molecules flung into a ghostly dance as the wings cut through the air. These particles form whirling vortices that brush against the wings, creating drag that, in extreme cases, can lead to a sudden loss of lift. Because problems like this can't be seen, they must be anticipated by aircraft pilots and designers. But even when the aircraft is a simulation inside a computer -- where airflow along its skin can be depicted in bright-colored streamlines -- designers studying these vortices often have a hard time picking out the places where they strike or separate from the wings.

Now, however, NAS researcher David Kenwright has devised software that causes these flow attachment and separation lines to stand out like the stripes on a zebra. The new algorithm, described by Kenwright in a paper to be presented at the IEEE Visualization '98 conference in October, automatically analyzes the surface flow on the aircraft, checking for mathematically significant features that indicate the arrival or departure of air particles. The result: a simulated aircraft showing separation and attachment lines that sometimes snake back and forth wildly as the vortices churn above them.

'Big Improvement' Over Current Methods

Aeronautics industry researchers say that Kenwright's new feature detection technique is a big improvement over today's most commonly used methods, which are spotty and time-consuming. "Some of our solutions have 35 million grid points, so for us to search manually for these important phenomena would take years," says Susan Ying, an aerodynamicist at The Boeing Company in Long Beach, Calif.

The total time saved by professionals in the computational fluid dynamics community could be considerable, Ying says, especially since the NAS data analysis group plans to include Kenwright's algorithm in a free library of visualization techniques, planned for release in September.

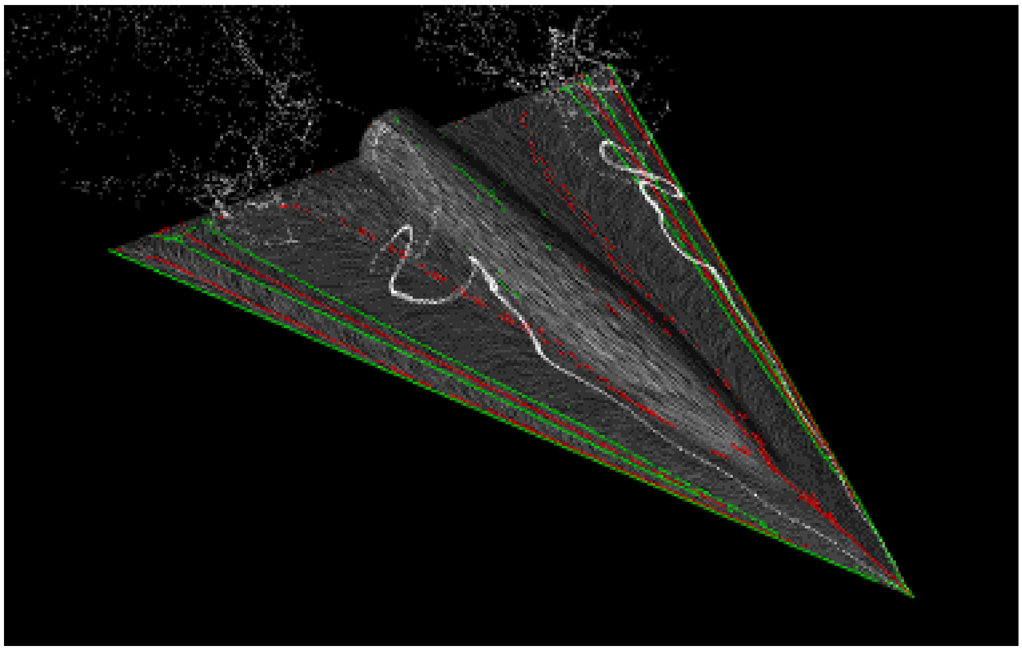

In this visualization of airflow near the surface of a delta wing aircraft, turbulent air separates from the wings along the white lines and reattaches along green lines. Locating these separation and attachment lines in a 13-gigabyte dataset took less than one second using new automatic feature-detection software developed by David Kenwright, a member of the NAS data analysis group. Delta wing dataset by Neal Chaderjin, NASA Ames Applied Computational Aerodynamics Branch. Graphic by David Kenwright.

Though the flow separation and attachment lines in Kenwright's demonstrations are simulated, aircraft engineers have a very real interest in them. Wings create the greatest lift at a high angle of attack -- that is, when they are pitched upward relative to the direction of the aircraft's forward motion. This is most useful during low-speed maneuvers such as takeoff and landing. Unfortunately, more vortices are generated at lower airspeeds and higher angles of attack; and this problem grows worse as the weight of an aircraft increases.

"One of the first lessons you learn as a pilot in training is that you don't follow immediately behind a big airplane," says Ying, "because they generate wake vortices, and even though you don't see them, [they're] really deadly if you get too close."

Separated airflow -- air moving abruptly away from the wings -- also increases drag. This in turn changes the stall speed, the speed at which, for a given angle of attack, turbulence above the wings drastically reduces lift. By modifying a wing's design in order to control where separation lines appear, and at what speeds and angles of attack, engineers can predict stalls and find ways to avoid them.

Reduces Analysis Time Significantly

Kenwright, a senior research scientist in the NAS data analysis group, was also involved in developing a method for detecting the cores of the wake vortices behind helicopter rotors (see NAS News, March-April 1997) He says the frustrating limitations of the two existing ways of finding attachment and separation lines first drew him to this new problem. In the most common approach, investigators simply "eyeball" the regions on a simulated aircraft surface where streamlines converge, relying on experience to judge whether these regions contain attachment lines or separation lines (which can have nearly identical appearances).

But as supercomputers have grown in speed and memory capacity, the amount of detail that can be depicted in whole-aircraft simulations has also increased, taxing the eyes of even the most ardent aerodynamicist. "This technique can reduce the time spent analyzing large datasets from days to minutes," says Kenwright.

No Dependence on Critical Points

In addition to saving time, the new algorithm detects more lines than the older methods. The other established technique for detecting separation and attachment lines overlooks a whole class of lines, Kenwright explains. That technique relies on a topological law dictating that streamlines connecting "critical points," where flow is momentarily stationary, must be separation or attachment lines. But "a lot of separation lines don't start or end at critical points," Kenwright notes. As far as the topological approach is concerned, these lines don't exist.

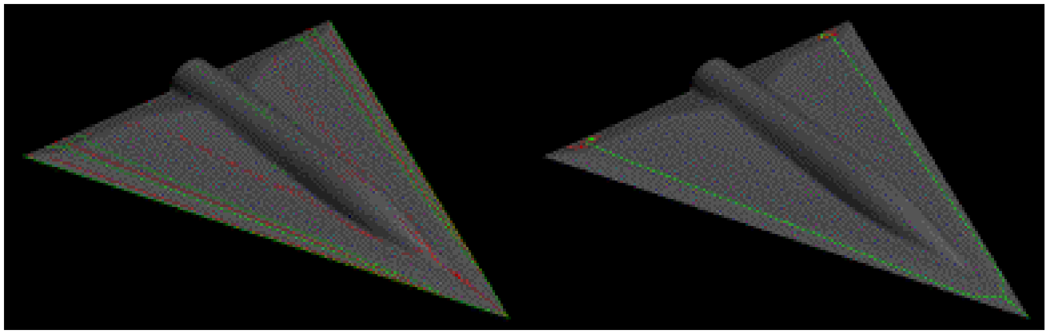

Previous techniques used to locate separation and attachment lines include visual inspecion of the surface streamlines (left) and vector field topology (right). Visual inspection of very large datasets is impractical, and vector field topology detects only the lines connecting points on the surface where flow velocity is zero. Kenwright's faster, more comprehensive method uses a new approach based on phase plane analysis. Datasets by Neal Chaderjian. Graphics by David Kenwright.

To correct that oversight, Kenwright devised a way of flushing out separation lines that doesn't depend on critical points. Instead, the new technique constructs the lines by analyzing airflow across the thousands of quadrilaterals that form an aircraft's skin in a computer model.

The algorithm first divides each quadrilateral into two triangles. For every triangle, the computer already knows the speed and direction of the velocity vectors originating at each vertex. The computer interpolates a continuous vector field between the vertices, and if certain topological features of the vector field happen to cross two of the triangles' edges, the line connecting the intersections must form part of a separation or attachment line. The computer highlights it, and the line segments in individual triangles merge into continuous separation or attachment lines.

Kenwright says the technique has detected all of the separation and attachment lines on several different simulated aircraft. Running in the C++ language on Kenwright's Silicon Graphics Inc. Onyx II workstation, the program examined approximately 100,000 triangles per second. At that rate, it took less than one second to locate all the separation and attachment lines on the surface of the delta-wing model shown on page 1. Kenwright has also applied the technique to time-varying aircraft simulations to reveal the dynamic behavior of separation lines, and has produced a five-minute, 9,000-frame video demonstrating the method and its applications.

Getting to the Root of the Matter

Ying, who intends to try Kenwright's technique on her own datasets at Boeing as soon as it is generally available, predicts that easier access to separation and attachment line locations will help aeronautical engineers tackle vortices at their source. "If you trace these vortices back, they come from the surfaces. This program allows you to go to that root cause," she says. "And some people claim that if you really do understand the root cause of the vortices, you can go back and destroy them" -- making this invisible hazard far more manageable.

For more information on this feature detection technique contact David Kenwright.

NAS Consultant Solves Hyper-X Simulation Problems

by Wade Roush

In February, while using a Cray supercomputer at the NAS Facility to simulate the crucial separation of the scramjet-powered Hyper-X from its Pegasus booster rocket -- a risky Mach-7 maneuver never before attempted -- the Hyper-X group at NASA Langley Research Center discovered a disturbing problem: The results of certain calculations turned out differently from run to run. "Id never encountered a case like that beforesays Doug Dilleyan assistant tech-nology manager and CFD lead for stage separation on the Hyper-X project. We needed to find out what was going how significant it was."

NAS scientific consultants Johnny Chang and Chuck Niggley flew to Virginia to meet with the Langley team. They soon traced the problem back to a Cray Fortran 90 compiler at the NAS Facility. The compiler, software that translates instructions written in the Fortran programming language into object code, was allowing the machine's multiple processors to share variables that were meant to be distributed among them privately.

Fortunately, Chang found a way to work around the bug, helping to put the Hyper-X project back in business in anticipation of a test flight over the Pacific in January 2000. "Johnny has done an amazing job," says Dilley. "Three seconds after separation, the vehicle has to be well clear [of the Pegasus booster] and has to have the right angle of attack," he explains. "It's very hard to get ground test data to simulate this. So we have to get results, and we don't have time to fool around with the Fortran 90 compiler."

'High-risk, Low-cost Program'

The Hyper-X (for Hypersonic Experimental Flight Vehicle) is a successor to the X-30 National Aerospace Plane, which was to fly eight times higher than current jet aircraft and at speeds up to Mach 25. Congress cancelled development of the X-30 in 1994 after NASA and the Department of Defense concluded that the necessary technologies were so immature that the plane could not be built within a reasonable budget. But that didn't kill all hopes for a future aerospace plane, and any such craft the U.S. develops (such as the envisioned High Speed Civil Transport) will need air-breathing engines such as the scramjet. Unlike rockets, scramjets do not have to carry their own oxygen supply, making them one of the most affordable candidate propulsion techniques for hypersonic passenger flight.

In keeping with the need for affordability, the Hyper-X project was designed as "a very high-risk, low-cost program," says Langley aeronautical engineer and Hyper-X assistant technology manager Sharon Stack. Indeed, the Langley team has pared down its mission objectives so much that the unpiloted vehicle will not be recovered, but will instead crash harmlessly in the ocean after its test. "We just want to demonstrate the scramjet," she says.

A scramjet is one evolutionary step above a ramjet, which is itself an advanced alternative to the conventional turbojet engine. In a turbojet, a rotating compressor driven by a turbine draws in outside air, compressing it before fuel is injected and the mixture is ignited. A ramjet is much simpler, consisting of little more than an inlet duct, a combustion chamber, and an exhaust nozzle. At the high speeds for which a ramjet is designed, air is compressed as it enters the conical inlet, making a compressor and turbines unnecessary. Without air constantly rushing through the engine at high Mach numbers, however, the exploding fuel in the combustion chamber would escape through both the exhaust and the inlet, so ramjets must be boosted to high speed using other propulsion systems before they can be turned on.

At about Mach 5, friction and shock waves raise the temperature of the air at a ramjet inlet to the engine's melting point. The scramjet, developed by aeronautical engineer and Hyper-X project leader Charles McClinton at Langley, gets around this problem using swept inlet walls that reduce the thermal load. By burning a combination of oxygen from the air and hydrogen fuel carried aboard the vehicle, the scramjet can reach hypersonic speeds -- that is, above Mach 5.

Getting to Supersonic Speeds a Challenge

The challenge for researchers who want to demonstrate scramjets in flight is that, like ramjets, they must be traveling at supersonic speeds before they can be ignited. A vehicle with dual turbojet-scramjet propulsion and maneuvering systems to allow runway takeoffs and landings was considered, but the Hyper-X team wanted to keep things simple and cheap. "Why pour all that money into getting it off the ground and worrying about recovery?" says Stack. "If we had to consider low-speed propulsion or flight controls for maneuvering and landing, we couldn't afford it."

Project engineers decided to get Hyper-X up to ignition speed in stages. In January 2000, a Boeing B-52 will take off from NASA's Dryden Flight Research Center near Edwards Air Force Base in southern California, carrying under one wing a 26-ton, winged Pegasus rocket with the diminutive Hyper-X -- measuring 12 feet in length -- attached to a special adapter on its tip. At 20,000 feet the B-52 will release the Pegasus, which will boost Hyper-X to an altitude of 100,000 feet and a speed of Mach 7.

Like experimental aircraft before it, the demonstration Hyper-X will get a "free ride" to flight altitude. In this artist's rendition, a Boeing B-52 will climb to 20,000 feet carrying a 26-ton Pegasus rocket under its wing, with the Hyper-X attached to the rocket's tip (above right). Then the rocket will detach from the Pegasus in an unprecedented "horizontal separation." Seconds later, the supersonic combustion ramjet (scramjet) engine inside the Hyper-X will fire, pushing the craft to Mach 11. To keep it from recontacting the Pegasus after separation, project engineers have prepared elaborate computer simulations of the crucial moment. Graphic courtesy of NASA's Langley and Dryden Research Centers.

The critical moment will come when mini-rockets inside the adapter fire, pushing Hyper-X slightly ahead of the booster while both are still speeding forward along the same line. A few seconds later, the scramjet will ignite. "Not only will this be the first time that a scramjet-powered vehicle has been demonstrated in flight, but we are not aware of anyone else who has done this type of horizontal separation," says Stack. "Of course, the concern is that the vehicles will recontact at some point. There are all kinds of potential aerodynamic nightmares, so we are trying to understand the air flow between the two vehicles."

Team Turns to Numerical Simulations
Operations at Mach 7 can't be fully modeled in a wind tunnel, and wind-tunnel tests at lower airspeeds can generate only static pictures of different moments in the separation process, Dilley explains. These snapshots reveal nothing about transient effects such as shock waves that might amplify the forces at work on the two vehicles. For that, the

Hyper-X team turned to numerical simulation.

Using a computational fluid dynamics package called SAMcfd (developed by ResearchSouth Inc.), Dilley and co-workers subjected computational grids representing the Hyper-X, its booster, and the air around them to millions of computations. They transformed the resulting data on changing forces, moments, and other flowfield quantities over time into a visualization of air flow around the two vehicles as they draw apart.

There was only one problem: Langley's production version of the SAMcfd code ran and compiled under the Fortran 77 compiler, designed for the older version of the Fortran language. Because the Fortran 77 code still worked well, the group had been reluctant to switch to Fortran 90. But when the NAS Facility finally decommissioned its Fortran 77 compiler from the Cray vector processors (Von Neumann, Eagle, and Newton) last March -- two years after Cray Research stopped support in favor of its Fortran 90 compiler -- the Hyper-X team was forced to switch to the new system.

"We tried to run the CFD code compiled under Fortran 90 and found that our results were not repeatable," recounts Dilley. "We would get some results, then run the program using the same input later the same day, and the solution would be different." The differences were insignificant when they first cropped up -- less than one-tenth of a percent. But because the results from each timestep in the simulation were used as input for the next step, even such tiny errors were quickly multiplying out of control, rendering the data useless. "We wanted to know what it was doing," says Dilley. "Was it a problem on the part of the new compiler or on the part of the simulation itself?"

Compiler Problem Identified -- and Solved

Enter Chang and Niggley, members of the NAS scientific consulting group who work with local and remote users of the Cray and parallel supercomputers to solve software-related problems. Niggley coordinated the effort and gathered the resources needed. Chang discovered a bug residing somewhere in the Fortran 90 compiler's "optimizer," which looks for opportunities to distribute work to several different processors. When he ran both the Fortran 77- and Fortran 90-compiled versions of the Langley team's code on a test grid with no optimization, they gave identical results. In practice, though, the full Hyper-X grid could never be processed without some optimization. "It would take years to run," according to Chang.

Chang wanted to find the exact spot in the Hyper-X simulation where optimization caused things to go awry. After returning to Ames, he broke the Langley team's SAMcfd code down into individual subroutines and compiled each subroutine twice -- once with and once without optimization. Then, by running object codes containing different mixtures of optimized and unoptimized subroutines and seeing which mixtures caused aberrations, "I was able to narrow the problem down to two subroutines out of more than 100," Chang explains. One of the two ceased to cause trouble after Cray Research released an updated version of its Fortran 90 compiler.

The problem with the other subroutine, which contained many different variables, was more stubborn. In order to optimize the subroutine to run on several processors, the compiler had to analyze which variables could be handled privately by individual processors and which needed to be shared among all the processors. In three cases, Chang found, it decided incorrectly. The compiler defined three variables that should have been private as being shared. When Chang explicitly specified which variables should be shared and which private, the problem vanished.

Return to Hypersonic Speed

The Langley engineers could finally get on with their simulation. "Johnny has been able to get consistent results from run to run," says Dilley. And, he says, that's good news because "the NAS Cray is our big workhorse for large-memory jobs." Now the team has returned to its real engineering problem: making the multistage approach work by ensuring a clean break when Hyper-X and the Pegasus booster part ways.

Chang, who alerted Cray Research about the optimization problem in the Fortran 90 compiler, enjoyed helping the Hyper-X team hunt down the bug. "It was a challenging problem," he says. "It wasn't something I immediately knew the answer to. It takes a lot of work and trying out new ideas to resolve some users' problems. And sometimes we find surprises."

If you've encountered problems with transferring your code to NAS's Fortran 90 compiler and would like assistance, contact the scientific consulting group or see the group's web site.

NASnews

▶ Main Menu July – August 1998 • Volume 3, Number 4

NAS Researchers' Simulations Give Shape to Vision of Nanotube Electronics

by Wade Roush

Most of the scientific visualizations displayed on the walls and computer screens of the NAS Facility are a fraction of the size of the aircraft they represent. But along one corridor here, the scales are reversed: pictures on monitors are tens of millions of times larger than the real thing. Physicist Deepak Srivastava and co-workers in the information technology modeling and simulation (ITMS) group are combining the NAS Facility's supercomputing power with state-of-the-art visualization techniques to simulate the behavior of carbon nanotubes, tiny cylinders only a few billionths of a meter in diameter (see NAS News, May-June 1997).

In recent conference papers and journal articles based on these simulations, Srivastava has reported new findings about the geometry and behavior of nanotubes that confirm their potential as tiny electromechanical sensors or semiconductor devices. He has shown, for example, that the laws of molecular dynamics allow the construction of nanotube T-junctions, key components in envisioned molecular-scale transistors.

Srivastava got these results by exploiting the multiprocessor architecture of the Silicon Graphics Inc. (SGI) Origin2000 computers at both NAS and the National Center for Supercomputing Applications (NCSA) in Illinois. Software running on these machines can analyze model molecules containing far more carbon atoms than could previously be simulated, and can take into account -- for the first time -- some of the quantum-mechanical effects found at the atomic scale.

Since practical ways to build complex nanotube-based structures haven't yet been devised, Srivastava explains, such simulations are the only way to test the feasibility of theoretical designs -- designs that, ironically, may someday lead to the manufacture of nanotube-based computers that could dwarf the Origin2000s in processing power. "You can dream up new geometries for model junctions, but you need supercomputing to check if they're right," he says.

Where It All Started

Researchers first observed "buckminsterfullerenes," carbon-atom arrays of which nanotubes are one variety, in 1985. Nanotubes, discovered six years later by Japanese researcher Sumio Iijima, consist of flat, hexagonal arrays of carbon atoms rolled into cylinders. The tubes are grown in laboratories from hot carbon vapor and can have either a single wall or several concentric walls. Single-wall nanotubes are the strongest fibers yet discovered, and some researchers are studying how to combine them into super-strong ropes, among other materials.

The ITMS group's interest in nanotubes, however, arises from several other intriguing properties of the molecules, including their ability to act as semiconductors. In 1992, materials scientist Mildred Dresselhaus and others at the Massachusetts Institute of Technology theorized that single-wall nanotubes should vary in their electrical conductivity depending on their "chirality," or the angle at which the hexagonal sheets are rolled.

Dresselhaus predicted that when the hexagons were aligned perpendicular to the nanotube's central axis, the molecule would conduct electricity as efficiently as a piece of metal. But if the hexagon pattern was parallel to the tube axis or had a helical twist, the tube would act instead like a semiconductor, conducting an electrical current only when the voltage applied to the material was above a certain threshold. If this were true, it would imply that joining two nanotubes with differing chiralities ought to produce a tiny diode.

"That's where it all started," says Srivastava, referring to the excitement in scientific circles over the possibility that nanotubes could be joined into microelectronic devices. "Once theories showed that a nanotube could be both metallic as well as semiconducting, people started to speculate about heterojunctions and three-terminal devices." A heterojunction is a simple electronic device in which materials with differing conductivities are joined. For example, a diode has an excess of electrons on one side and a deficit on the other, allowing current to flow in one direction only.

Adding to the excitement, researchers in the U.S. and the Netherlands confirmed Dresselhaus's prediction earlier this year. At about the same time, Alex Zettl and co-workers at the University of California, Berkeley, observed the first nanotube diode in the laboratory. And NAS Systems Division researcher Jie Han has teamed with H.J. Dai at Stanford University to construct a variety of similar devices in the laboratory by inserting geometrical "defects," such as pentagons and heptagons, between nanotubes of differing chirality.

Next Milestone: The T-junction

Yet while two-terminal devices such as diodes have many uses, they lack the versatility of three-terminal semiconducting devices such as the transistor, which can switch on and off with input from a third terminal. "For transisting and amplifying devices such as those used in [microprocessor] logic gates, you need three-terminal structures," says Srivastava. "We were the first to propose this."

The construction of such structures, however, remains some time away. A three-terminal nanotube structure would involve a T-shaped junction, but the only technique researchers now have for joining nanotubes in the laboratory is to lay them end-to-end.

Researchers in chemistry and condensed-matter physics have debated whether nanotube T-junctions are even feasible. It's easy enough to fuse the ends of two nanotubes of differing chirality into angled heterojunctions by introducing a few pentagons on the inside of the bend and a few heptagons on the outside. But, Srivastava explains, there have been doubts about whether this approach would work for a T-junction -- until now.

Srivastava and collaborator Madhu Menon at the University of Kentucky, Lexington, decided to fly in the face of an accepted interpretation of Euler's theorem -- a topological rule that seemed to require that the number of pentagons and heptagons in a polyhedron such as a nanotube be equal -- by testing whether heptagon-heavy nanotube T-junctions might be physically plausible. Using commercial plotting software (including Cerius2, developed by Molecular Simulations Inc.) and some of their own code, the two researchers created 3-D computer models of two metal-semiconductor-metal T-junctions: one with six heptagons and no pentagons, the other with eight heptagons and two pentagons (as shown below).

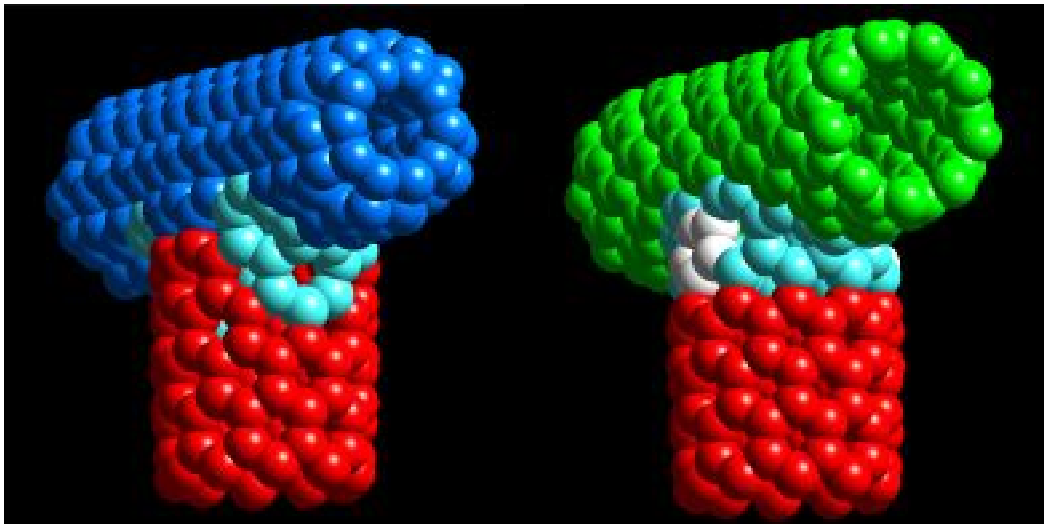

Given a plausible way -- both physical and geometric -- to join two nanotubes in a T shape, a tiny transistor could be built from nanotubes of differing conductivities. To form the simulated T-junctions above, Deepak Srivastava and co-workers in the information technology modeling and simulation group started with semiconducting (red) and metal-like (blue and green) nanotubes and inserted pentagons (white) and heptagons (light blue) at the inside bends. Supercomputer analysis proved that these geometries are stable under the laws of quantum molecular dynamics. Graphic by Deepak Srivastava.

But it's one thing to build an attractive computer model of a molecule and quite another to make it hold together in the real world. "It's almost like Tinkertoys," explains Srivastava. "You have these tubes. Give kids enough time, and they might figure out how to join them. But we have to check that it's done within the constraints of physics." The critical question for Srivastava and Menon was whether the angles and widths of the bonds between carbon atoms at the inside bends of the T were consistent with the known behavior of covalent bonds.

Better Results Through Visualization

They answered this question using the Generalized Tight Binding Molecular Dynamics (GTBMD) simulation software, introduced by Menon and another colleague in 1996. GTBMD analyzes simulated molecules designed using Cerius2 and other software, to determine whether the bonds between the atoms in the proposed molecules are allowable under the laws of quantum molecular dynamics.

Quantum GTBMD is more precise than the preceding tight-binding molecular-dynamics methods, Srivastava explains, because it takes into account a rule in quantum mechanics which states that electron orbitals in a covalent bond need not be orthogonal (at right angles to each other). With this rule in effect, the atoms in Srivastava and Menon's simulated T-junction could "relax" or squeeze together more closely than is possible with classical, more approximate methods.

In addition, after Srivastava spent several weeks last year modifying GTBMD to run on Turing (one of the 64-processor Origin2000s at NAS), the software had sufficient power to analyze models containing more than 1000 carbon atoms each -- enough to minimize the unrealistic physical forces that crop up at chopped-off edges, often ruining simulations of smaller structures. "If you have to simulate only five, six, or ten atoms, a lot of people can do it," says Srivastava. "We've been able to extend that."

Simulating more atoms at once and accounting for quantum-mechanical effects turned out to be the cure for uncertainties over the T-junction. As Srivastava and Menon reported in Physical Review Letters last December, the modified GTBMD code confirmed that their model structures were workable. The structures also proved stable at room temperature in a classical molecular dynamics simulation. "The simulations prove that it is feasible to connect two nanotubes at right angles, and that the junction is fairly stable," sums up Srivastava, adding, "We have predicted that even more complex junctions will be possible."

Studying Nanotubes Under Stress

Using a similar set of technical advances, the ITMS group is also studying how regular, non-branching nanotubes behave over time -- especially under changing levels of mechanical stress. "Even though nanotubes are the strongest known material for compression and stretching, they are very flexible for twisting and bending and can even double over," says Srivastava. As strain concentrates at the pinched edges of a bent or twisted nanotube, he explains, the electrical charge and chemical reactivity of these regions may also change. If these effects can be understood and harnessed, the tubes might become useful as tiny electromechanical or electro-chemical sensors.

In past studies using a classical molecular dynamics simulation code employing a formulation called "Brenner's many-body atomic potential for C-C interactions," researchers found that a single-wall nanotube can absorb only a moderate amount of compressive force before it begins to buckle. Srivastava and Stephen Barnard, a computer scientist at NAS, guessed that a multi-wall nanotube might be stiffer than a single-wall tube and behave more predictably under stress. They knew that supercomputers would be required to test this hypothesis, since even a short section of a large-diameter, multi-wall nanotube can contain hundreds of thousands of carbon atoms; and no conventional computer has sufficient processing power and memory to apply Brenner's potential to a system this large.

Srivastava and Barnard used "classical techniques implemented in parallel" to simulate the large nanotubes, modifying the code applying Brenner's potential so that it would run on a multi-processor Origin2000 at NCSA. They also made an important change to Brenner's formulation, adding algorithms to calculate the van der Waals force, which has been shown in laboratory experiments to maintain a certain separation between concentric nanotube walls, increasing the stiffness of the whole assembly.

When Srivastava and Barnard applied the modified code to simulated nanotubes, the resulting computer animations confirmed that "multi-wall nanotubes behave rather differently from single-wall tubes," says Srivastava. Under low compressive force, a single-wall nanotube develops symmetric pinches distributed uniformly along its length. Under higher forces, the pinches become irregular, and strain begins to accumulate disproportionately in certain pinches (see below). Finally, the tube buckles, and deformations that were formerly elastic become plastic -- that is, the tube can't return to its original cylindrical shape even after the force is removed.

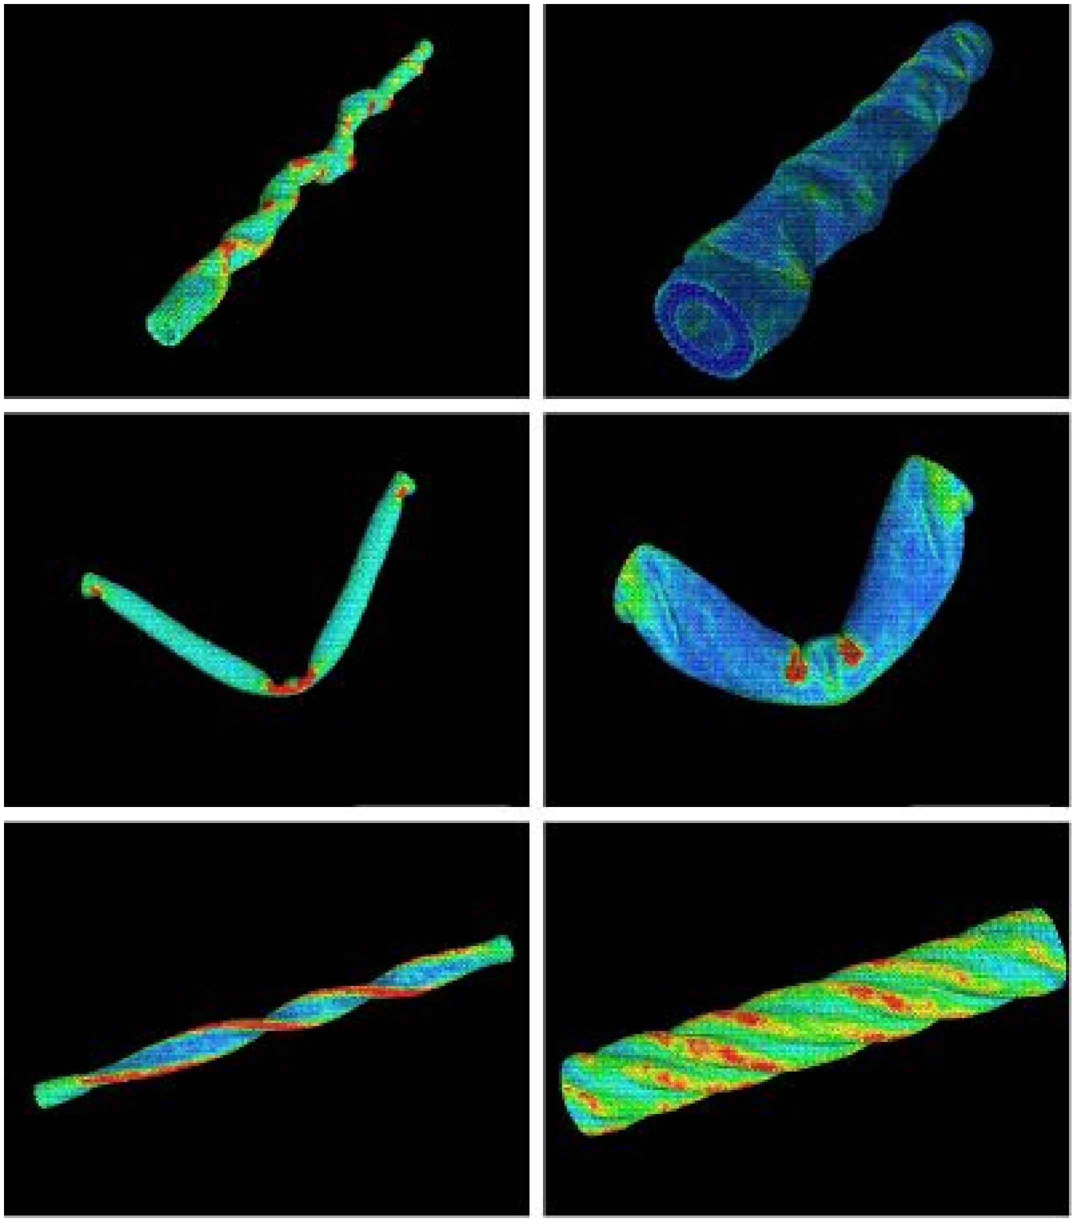

Changes in the shape and surface electrical charge of nanotubes under mechanical stress could make them useful as tiny electromechanical sensors. Srivastava and co-workers used computer simulations to model single-wall (left column) and multi-wall (right column) nanotubes under squeezing, bending, and twisting forces (top, middle, and bottom, respectively). Single-walled tubes developed kinks at unpredictable places, the researchers found, while the less flexible multi-walled tubes distorted symetrically. (Red = areas of high mechanical stress; green = medium; blue = low.) Graphics by Deepak Srivastava.

By contrast, a four-wall nanotube subjected to the same degree of strain undergoes symmetric pinching, then some sideways buckling, but no plastic deformation. In a paper presented at the SC97 conference, Srivastava and Barnard reported that severe deformation does not occur until the compressive force overwhelms the Van Der Waals force. Using their model, Srivastava says, researchers can now predict the exact amount of force required. Without in-house tools and commercial software such as Cerius2, and without supercomputers that can quickly evaluate the fitness of these molecular models -- particularly those that may contain hundreds or even thousands of atoms -- such predictions would be well beyond the reach of science. "You get much more understanding through a visualization," Srivastava says. "And in this case, understanding the exact pinching behavior will help us take advantage of it." For example, engineers designing nanotube sensors sensitive to differences in pressure, charge, or chemical concentration could capitalize on the fact that the distribution of electrical charge on a nanotube's outer surface varies predictably around pinching points.

Ongoing Nanotube Work

In his ongoing work, Srivastava is using simulations to determine whether networks of nanotubes joined at T-junctions and other kinds of junctions could form complex electronic devices. He is also investigating the properties of nanotubes under a wider array of conditions. Applying elements such as boron and nitrogen to pure graphene sheets, for example, may alter the nanotubes' electronic properties.

Nanotubes may also make ideal storage capsules for hydrogen fuel on future space missions. Storing hydrocarbon fuels in nanotubes would be "lighter and cheaper" than conventional methods, Srivastava says. The nanotube capsules could be burned along with the fuel, producing far fewer toxic by-products than conventional solid-fuel rockets such as the Space Shuttle's solid rocket boosters.

Using simulations, Srivastava wants to determine what diameter a tube should have to allow as many hydrogen atoms as possible to be packed inside, and how the tubes' end caps could be opened and closed for loading. These simulations will guide experiments that, Srivastava hopes, will someday make nanotube-based technologies as significant as CFD simulations have been for the aerospace community.

Researchers interested in collaborating with Srivastava can contact him at (650) 604-3486 or at deepak@nas.nasa.gov. Srivastava and Barnard's SC97 paper on nanotubes under stress and details on NAS Systems Division's carbon nanotube device modeling projects are available.

NASnews
▶ Main Menu
July – August 1998 • Volume 3, Number 4

Virtual Wind Tunnel Software Ready for Real World

by Wade Roush

Aerospace engineers need to understand the flow of air around their prototype vehicles in as much detail as possible. But getting close to a model in an operating wind tunnel is impossible -- even if they could withstand the gale-force winds, the engineers' presence would interfere with the flow itself. Numerical simulations are one way around this, and now a quick, intuitive way to view these simulations, previously open only to selected users, is available to the general aerospace community. The Virtual Wind Tunnel, software developed in the NAS Systems Division, allows users to immerse themselves in 3-D representations of their computational fluid dynamics (CFD) datasets using standard virtual-reality gear.

A new version of the program was released on July 1. Visitors to the NAS Facility and a few evaluators have had access to previous incarnations of the VWT, but this is the first version to be approved for release through negotiated agreements with interested parties, according to research scientist Steve Bryson, the prime mover behind the VWT and a member of the NAS data analysis group.

The VWT immerses users in a virtual space where colorful markings such as streamlines appear around computer-generated aerospace vehicles. As the system follows the movements of a user wearing head-tracked stereo glasses or a boom-mounted stereo display, the scene shifts accordingly in real time, allowing investigators to examine any part of the flow pattern from any point of view.

Researchers familiar with the software report that the VWT enables aerospace designers to interact with large CFD datasets more productively. "Being able to see a wider field of view and to feel present is a really powerful thing," says Val Watson, an Ames senior staff scientist who studies the visualization of scientific data.

Users at the Boeing Company and NASA's Langley and Goddard research centers have worked with earlier versions of the software, and interagency agreements allowing the Army and Navy to use the new version of the VWT are in the works, according to Bryson. Discussions on patenting the VWT and licensing it to interested users through a commercial software firm are nearing a conclusion. "Finally, it's going out into the world," he says with a note of satisfaction.

Challenge in Handling CFD Data

The VWT might have stepped into the real world sooner, Bryson explains, if not for the technical challenges involved in metamorphosing stereo representations of huge CFD datasets in near-real time. If a user moves an object in a virtual-reality environment, the object must be re-rendered in one-tenth of a second or less -- or all hand-eye coordination is lost. But when the scene has the complexity of a CFD simulation with millions of data points, that's no easy task. "The primary problems were not with the virtual-reality interface but with handling all that data," Bryson says. "It's taken a very long time for the techniques for handling these datasets to come together and for the technology to become powerful enough."

The solution first came into view when data analysis group leader Michael Gerald-Yamasaki suggested "breaking out" the different problems involved in visualizing complex datasets, then devising general solutions that could be used in the VWT and many other CFD applications. Bryson recalls that the group thinking was, "Why don't we make data handling work in the abstract -- that will be great for everyone."

The result was the Field Encapsulation Library (FEL), a collection of programs that access CFD data, track its location on disk or in a computer's memory, move it from one location to another, and transform it from one format to another. Bryson and research scientist David Kenwright, also in the data analysis group, completed the first version of FEL in 1996. (A new version, FEL II, was released on June 30.) Additions to Include Collaborative Engineering

Since 1996, Bryson and co-workers -- most recently, visualization software engineers Bryan Green and David Whitney -- have focused on giving the VWT added capabilities and a cleaner look. Even before the latest additions, VWT users could view flow information (which must be in PLOT3D format for regular or unstructured CFD grids) in many different forms, including streamlines, streaklines, cutting planes, particle paths, isosurfaces, tufts, or simple numerical displays. They could also manipulate this data using tools such as the "rake," a movable line that emits particles into the flow field.

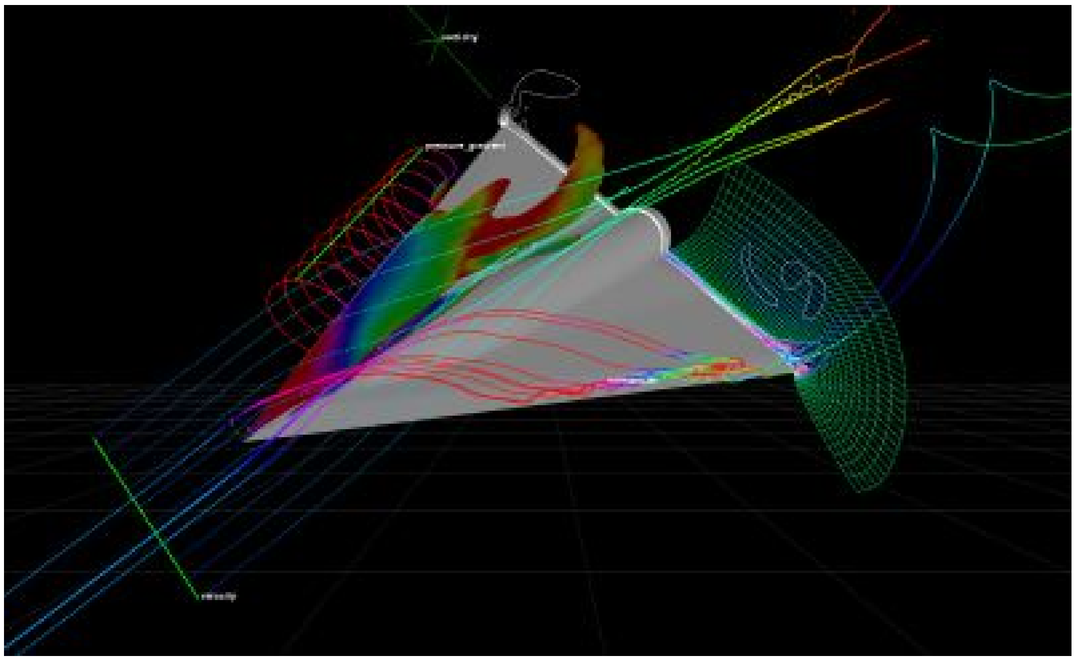

Some of the many visualization tools available in the Virtual Wind Tunnel. Streaklines (dotted lines) represent flow velocity colored by time. The lower set of streamlines shows momentum, colored by Mach number; the upper set shows pressure gradient (the direction of greatest pressure variation) colored by pressure value. A partial grid plane colored by vorticity appears along the trailing edge, with contours of vorticity magnitude (in white). An isosurface of density colored by velocity magnitude appears above the upper wing. Dataset by Neal Chaderjian. Graphic by Steve Bryson.

Now they can also control lighting effects and graphics quality; specify that visualization computations occur at high speed and low resolution or vice versa; use a virtual trackball to move tools through the virtual space; access a vastly increased number of "derived parameters" based on basic CFD values such as pressure, density, and temperature (see figure above); and get help from extensive online menus.

And within a few months, according to Whitney, multiple users will be able to work together within the VWT's virtual space from dispersed geographical locations, by connecting over the Internet to a central server responsible for rendering a scene. Collaborative virtual engineering programs such as Division Inc.'s dVise have been commercially available for some time, but not for CFD applications. Engineers and designers using such systems typically focus on a static model, not on how that model interacts with its environment.

Engineers at Boeing are eagerly awaiting the deployment of the VWT's collaborative feature, according to Whitney. „If we can prove the concept, they can convince their managers that they should be using this in a production sense."

Cheaper Than Wind Tunnel Tests

Using the VWT software to achieve full virtual-reality "telepresence" will, of course, require certain hardware. The basics are a Silicon Graphics Inc. (SGI) workstation and head-tracked stereo glasses, boom-mounted glasses, or a back-projection "virtual workbench," items that can reach $100,000 in total cost. But the space inside the VWT can also be explored using a standard two-dimensional interface (a flat SGI workstation screen); and according to Bryson, work in the next few months will allow the VWT to run on less-expensive platforms, including Hewlett-Packard workstations.

Such costs, while significant, are minor compared to the charge for wind-tunnel time. In fact, building a "numerical wind tunnel" to reduce the time and expense involved in real wind tunnel tests was part of the NAS Facility's original charter. NAS researchers set out in 1984 to prove that supercomputers could use the laws of aerodynamics to simulate the flow of air around an entire aircraft, and that the resulting data could be visualized in ways that were just as informative as conventional wind-tunnel testing.

But even before the computer hardware and software needed to carry out whole-aircraft simulations had matured, Bryson explains, researchers Scott Fisher and Mike McGreevey in the Ames Aerospace Human Factors Research Division and Creon Levit at NAS had begun to speculate that an immersive medium such as virtual reality would be needed to view these simulations effectively. "It was a natural marriage," says Bryson, who joined Fisher's group in 1987 and moved to NAS in 1990.

Displayed as columns of numbers or as flat images on a screen, he explains, large CFD datasets can be overwhelming and difficult to interpret. But with virtual reality, "You're exploring directly, moving things around, and there's an involvement that significantly enhances the meaningfulness of the experience."

3-D Adds 'Powerful Impact'

The main goal of the VWT project is to help engineers learn as much as they can from their simulations before beginning the costly work of building prototype aircraft and testing them in real wind tunnels or in flight. But Bryson admits it can't hurt if users have a little fun.

"The three-dimensional aspect of it has a very powerful impact on people," Bryson remarks. "In particular, it's my strong impression that when people explore on their own in a three-dimensional environment, they form a much stronger memory of what they did than when it's just on a screen. It's more fun because it's meaningful."

Some of the latest work on the VWT is described in a paper submitted to the IEEE Visualization 1998 conference to be held in Research Triangle Park, NC, in October. For more information on the VWT, including the user's manual, see http://www.nas.nasa.gov/Software/VWT, or contact Steve Bryson at bryson@nas.nasa.gov.

New Users: What You Need to Know When the 1999 Operational 'Clock' Starts

When the wall-clock in the NAS Facility's computer room strikes midnight on Thursday, October 1 -- marking the beginning of the 1999 operational period -- current users will turn into pumpkins, and a new set of users will get their turn at using the magic of supercomputing. Whether you end up among the pumpkins or the "royals," you'll want to know about the revised selection guidelines implemented earlier this summer, as well as the steps you'll need to take to set up your new account or find other supercomputing resources for your research.

The biggest change going into effect for 1999, according to Chris Kleiber, NAS control room manager, is that time on the NAS and Aeronautical Consolidated Supercomputing Facility (ACSF) CRAY C90 supercomputers, and one of the CRAY J90s, now comes with a price tag. Previously, time on these systems was given away through a competitive peer-review process.

The ACSF is handing over supervision of Von Neumann, Eagle, and the Newton cluster to the year-old Consolidated Supercomputing Management Office (CoSMO), which will sell time on Von Neumann and Eagle for $39 per hour and on the Newtons for $25 per hour. CoSMO will also sell time for production jobs on the five Silicon Graphics Inc. Origin2000 parallel systems at the NAS Facility. This won't be CoSMO's first experience in account administration. During the 1998 operational period, in an effort to promote CoSMO's growth, Ames Research Center placed half of the time available on Von Neumann under CoSMO's management.

'Self-sufficient' Supercomputing

CoSMO's charge is to manage and coordinate supercomputing throughout NASA, and one of the program's first goals is to make supercomputing at the agency more self-sufficient, according to CoSMO manager John Ziebarth. "As NASA moves toward the goal of full cost accounting, it's necessary to begin to recover some of the operations costs through a charge-back formula," Ziebarth explains. The agency will probably continue to pay for other costs, such as future capital costs and networking, through a centralized pool of resources, he says.

A large chunk of time on the Origin2000s will still be available on a competitive, zero-cost basis, says Eugene Tu of the Information Technology (IT) Base Program. As acting IT program manager, Tu, along with staff in the High Performance Computing and Communications (HPCC) Program, will run this peer-review selection process. Jointly, these programs are embarking on the new Information Power Grid (IPG) research and technology project to connect supercomputing resources into an agency-wide testbed for distributed computing, allowing jobs to be allocated wherever the most CPU time is available.

Given this focus, Tu notes, a joint IT-HPCC Announcement of Opportunity, released in late June, included new criteria for research proposals. This year, rather than soliciting projects generally related to aeronautics or computer science, the review committee will look for proposals with direct impact on the new IPG research objectives such as high-speed data transfer, data mining, multiplatform compilation, batch processing, and other key areas of heterogeneous distributed computing, according to Tu.

In addition, says Tu, proposals will require sponsorship from one of the base or focus programs within NASA's Aeronautics and Space Transportation Technology Enterprise. Notification of selections and allocations will be made by August 24.

'Royal' Red Tape

Investigators who are awarded time on the Origin2000 systems need to be aware of a few administrative requirements, says Kleiber. Before you can access the machines, you'll need an account request form, available online. To complete the form, you'll need the group identification number for the 1999 allocation, which will be disclosed to the project's principal investigator at the time of acceptance. The form should be printed out and signed by the principal investigator. Then, fax your completed forms to the NAS user interface coordinator at (650) 604-1777 before September 27 to ensure access to resources on October 1.

New accounts will be activated at about 6:00 a.m. Pacific time on October 1. The accounts of existing users who receive additional computer time for the 1999 operational period will simply continue, with their passwords unchanged. NAS help desk staff members will telephone new users with their passwords sometime before October 1; but only after the accounts are activated will users be able to log in and use the systems.

Meanwhile, as the wall-clock time for the 1998 operational period runs out, the accounts of existing users who didn't receive time in 1999 will be disabled no later than 6:00 a.m. on October 1. These users will have access to the NAS Facility's mass storage devices for another 90 days, during which NAS help desk staff will archive all remaining data on the Crays to the mass storage system for users to access.

Good News for 'Pumpkins'
If you're among those who weren't awarded time on the supercomputers, but would still like to purchase cycles from CoSMO, contact your sponsoring NASA program office for more information. For information on CoSMO, contact John Ziebarth at (650) 604-2761, send email to ziebarth@nas.nasa.gov, or go to the CoSMO home page.

Users who want to purchase time from the CoSMO office must complete a Project Cover Sheet. This form must be received by the NAS user interface coordinator by September 1.

Go To Sidebar...

NASnews
Main Menu
July – August 1998 • Volume 3, Number 4

Help for New Users Available Online

The NAS web site has plenty of information to assist both new and returning users. Here are a few places to explore.

Information to get started:
- NAS home page
- NAS Core Technology page

General User Questions:
- These pages provide information on system status, as well as links to various forms, including the account request form.

Information about supercomputers supported by the NAS Systems Division:
- Provides links to home pages for all supercomputers supported by the division. In addition, links are provided to the Guide to Getting Started Quickly. For the Cray vector parallel systems, a link is provided to the System Customized User Guide.

The Scientific Consulting (SciCon) home page:
- Provides a variety of links to Math Library information, the SciCon online tutorials (see below), and NAS News articles related to the NAS scientific consultant¹s work. In addition, it provides links to a selection of other sites that may contain related and useful information, links to additional online resources.

The Scientific Consultants Online Tutorials:
- Contains links to a selection of tutorials presented by the NAS scientific consultants on topics pertinent to both the Cray parallel vector processors and parallel systems.

Return to main article

NASnews
▶ Main Menu
July – August 1998 • Volume 3, Number 4

Technical Seminars Available on Video

Many research seminars presented at the NAS Facility are accessible through a videotape loan program. Also available is information on past training events and procedures for obtaining training materials.

Here are summaries of the most widely attended technical seminars at the NAS Facility since the beginning of April.

Interaction of THz Transients and Optical Pulses in Quantum Wells: Anisotropic Electron-hole Wavepackets and Optical Switches. Steve Hughes, from the Department of Physics at Washington State University, has studied the creation of THz-driven, anisotropic electron-hole relative motion wavepackets in quantum wells. In the July 14 New Technology Seminar, Hughes analyzed a series of propagating wavepackets. He also discussed applications to the recently observed dynamical Franz-Keldysh effect and high-harmonic-generation in the THz regime. Hughes showed how an intense half-cycle THz pulse can be used in conjunction with a quantum-well optical amplifier to form the basis of an ultra-high-speed optical switch.

Within This Article...

Virtual Worlds on the Internet: A Bold New Medium for Interaction, Collaboration, Play, and Learning. The PC has evolved into a capable, real-time 3-D platform and is now serving as a portal into large-scale Internet-hosted 3-D virtual worlds inhabited by thousands of users. At the June 18 New Technology Seminar, Bruce Damer, president of DigitalSpace Corp., took the audience on a tour into several online virtual worlds, meeting and speaking with their inhabitants while visiting several ^{3}purpose-built^{2} areas. Damer also talked about the system and user interface, architectural issues of constructing these spaces, and the social craft of creating and sustaining the communities within them.

Cracking a Tough NUT* (*Nanocrystal Under Tension) with a Big Computer: Studying Materials Failure Using Millions of Atoms. At the June 2 Nanotechnology Seminar, Farid Abraham, of the IBM Research Division at Almaden Research Center, discussed cracking a nanocrystal under tension with large computers, and how conventional continuum theory is not capable of explaining the violent crack velocity oscillations that can occur. With the advent of scaleable parallel computers, Abraham said, atomistic simulations are providing immediate insight into the nature of fracture dynamics by allowing us to ^{3}see^{2} what is happening on the atomic scale. In a videotape of the fracture simulations, he showed how dislocations occurred and how this problem scales to the future teraflop regime in scientific computing.

Advances in Computing: Present to the Next Millennium. The past ten years have yielded tremendous advances in computing technologies. At a Special Presentation Seminar on May 22, N. ^{3}Radha^{2} Radhakrishnan, director of the Information Technology Laboratory (ITL) at the U.S. Army Waterways Experiment Station in Vicksburg, Miss., provided a synopsis of the evolution and present state of technology within the ITL and a glimpse of where the technology is headed in the next millennium. Radhakrishnan discussed the ITL's hardware environment, its interactions with the academic communities, computational highlights, metacomputing, and ITL's path to the future.

Supercomputing at the U.S. Environmental Protection Agency/The National Environmental Supercomputing Center. Arthur Cullati discussed the U.S. Environmental Protection Agency's (EPA), supercomputing capabilities at a Special Presentation Seminar on May 15. Cullati, director of the National Environmental Supercomputing

Center (NESC) -- the EPA's first and only supercomputing center -- traced the evolution of the NESC from its inception in 1992 to the present.

Tracking for Video-See-Through Augmented Reality. Augmented reality (AR) has demonstrated promise as a paradigm for providing users with real-time, high-quality visualization of a wide variety of information. However, some key technological problems have prevented AR from reaching its potential. At the May 12 Data Analysis Seminar, Mark Livingston, a graduate student at the University of North Carolina, discussed AR ultrasound visualization, describing the prototype system, along with the progression of approaches to improving tracking.

VisualEyes: Integrated Data Synthesis, Analysis, and Visualization of Parallel Adaptive Simulations. Chandrajit Bajaj, of the Texas Institute for Computational and Applied Mathematics at the University of Texas, discussed the VisualEyes project at the April 30 Data Analysis Seminar. The project focuses on research in some of the core technologies (compression, adaptive meshing, interrogative visualization, and error estimation) and integrates parallel simulation with data analysis and collaborative visualization of multivariate scalar/vector tensor fields. Bajaj discussed the details of a progressive mesh encoding scheme, as well as the use of vector field topology for several applications in interrogative visualization.

Representation and Image Comparison Metrics for Very Large Datasets. At the April 16 Data Analysis Seminar, Raghu Machiraju of the NSF Engineering Research Center for Computational Field Simulation at Mississippi State University, talked about representation and image comparison metrics for very large datasets. He addressed two important operations needed for the visualization of large (terascale) datasets: compressed domain representation and metrics for comparing datasets and images.

"Development and Parallelization of a Hybrid Particle/Continuum Method for Simulating Rarified Flow." A particle method such as the Direct Simulation Monte Carlo (DSMC) method simulates a gas flow by statistically modeling the behavior of a large number of virtual particles. Such methods are required for the simulation of rarefied flows for which the Navier-Stokes (NS) equations have become invalid due to failure of the constituent relations upon which they are based. Unfortunately, while more versatile than NS, DSMC also requires much more computational effort. Even on a parallel computer, simulation times can become long, and a less computationally intensive method is needed. At the April 14 New Technology Seminar, Craig Duttweiler, from Stanford University, discussed two approaches to the hybridization of NS and DSMC, which would allow the former to handle regions of higher density and the latter to handle regions of greater rarefaction.

To borrow a videotape of these or other NAS-sponsored seminars, send a request to the NAS Documentation Center.

NASnews
Main Menu
July – August 1998 • Volume 3, Number 4

In this issue...

Credits

Executive Editor: Thomas Lasinski

Editor: Jill Dunbar

Staff Writer: Wade Roush

Contributing Writers: Peter Adams, Marcia Redmond

Image Coordinator: Joel Antipuesto

Online Page Layout and Graphics: Peter Adams

Other Contributors: Peter Adams, Richard Anderson, Steve Barnard, Steve Bryson, Johnny Chang, Doug Dilley, James Donald, Chris Henze, Randy Kaemmerer, David Kenwright, Chris Kleiber, Art Lazanoff, Marcia Redmond, Sharon Stack, Deepak Srivastava, Eugene Tu, David Whitney, Susan Ying, John Ziebarth

Editorial Board:
Cristy Brickell, Jill Dunbar, Chris Gong, Thomas Lasinski, Nateri Madavan, Patrick Moran, George Myers, Wade Roush, Harry Waddell