NASA / Fall 2000

InSights

High-Performance Computing and Communications

Issue 15, Fall 2000

Features

Events

NASA software, page 2

The Perfect Storm, page 10

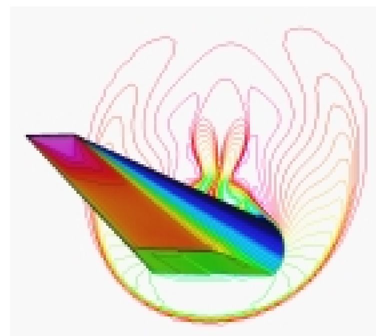

Reusable Launch Vehicle, page 18

Gigabit networking, page 22

Supercomputing Conference, page 29

Interactive software helps

By Lou Varricchio

When Sheri Zakarowsky received a letter from NASA's Learning Technologies (LT) Project team, she knew a unique opportunity was knocking—an opportunity to inspire high-school students to learn about aeronautics. The letter put out a call to science teachers interested in teaching aerodynamics to middle- and high-school students. Zakarowsky responded immediately. Here was a chance, she thought, to bring NASA technology and information down-to-earth and get it into the hands of students.

Zakarowsky, a science teacher at Richmond Heights High School in Ohio, learned from the letter that the LT team at the NASA Glenn Research Center (GRC) in Cleveland, Ohio, was developing FoilSim, a unique, interactive graphical software created in 1997 to help students perform simulated flight experiments.

"I was looking for exciting, new teaching aids to egage my students in learning about science," says Zakarowsky. "The LT group wanted teachers to help develop the accompanying lessons for the *Beginner's Guide to Aeronautics*, FoilSim's tutorial about flight dynamics which is packaged with the software. As the LT team developed the software and beginner's guide, several science teachers, in turn, would help create the lessons."

students learn about flight

According to Zakarowsky, FoilSim is a learning tool that captures the imagination of students and teachers alike. She knew the software had the potential to become an outstanding classroom aid so she joined seven fellow Ohio teachers—all selected by the LT group—to *test fly* FoilSim and work on lessons. The teachers met periodically at GRC for planning and working sessions.

"We worked together, developing the lessons at one of the Glenn computer labs," she says. "This was an opportunity to make students think about the world around them. Young people can enjoy FoilSim like a video game, but at the same time, they really learn something about science and engineering."

Under the guidance of the LT team, the lessons were soon ready for packaging with FoilSim and eventual classroom introduction. Beginning with the 1998-99 school year, the FoilSim software and lessons were available free via the World Wide Web. That same year Zakarowsky began using the software with her high-school physics students. Even as they use it in class, she comments, few high-school teachers or students know about the genesis of FoilSim.

Students at Marina Village Middle School in El Dorado Hills, Calif., build and fly kites to study the effects of lift and drag on airfoils.

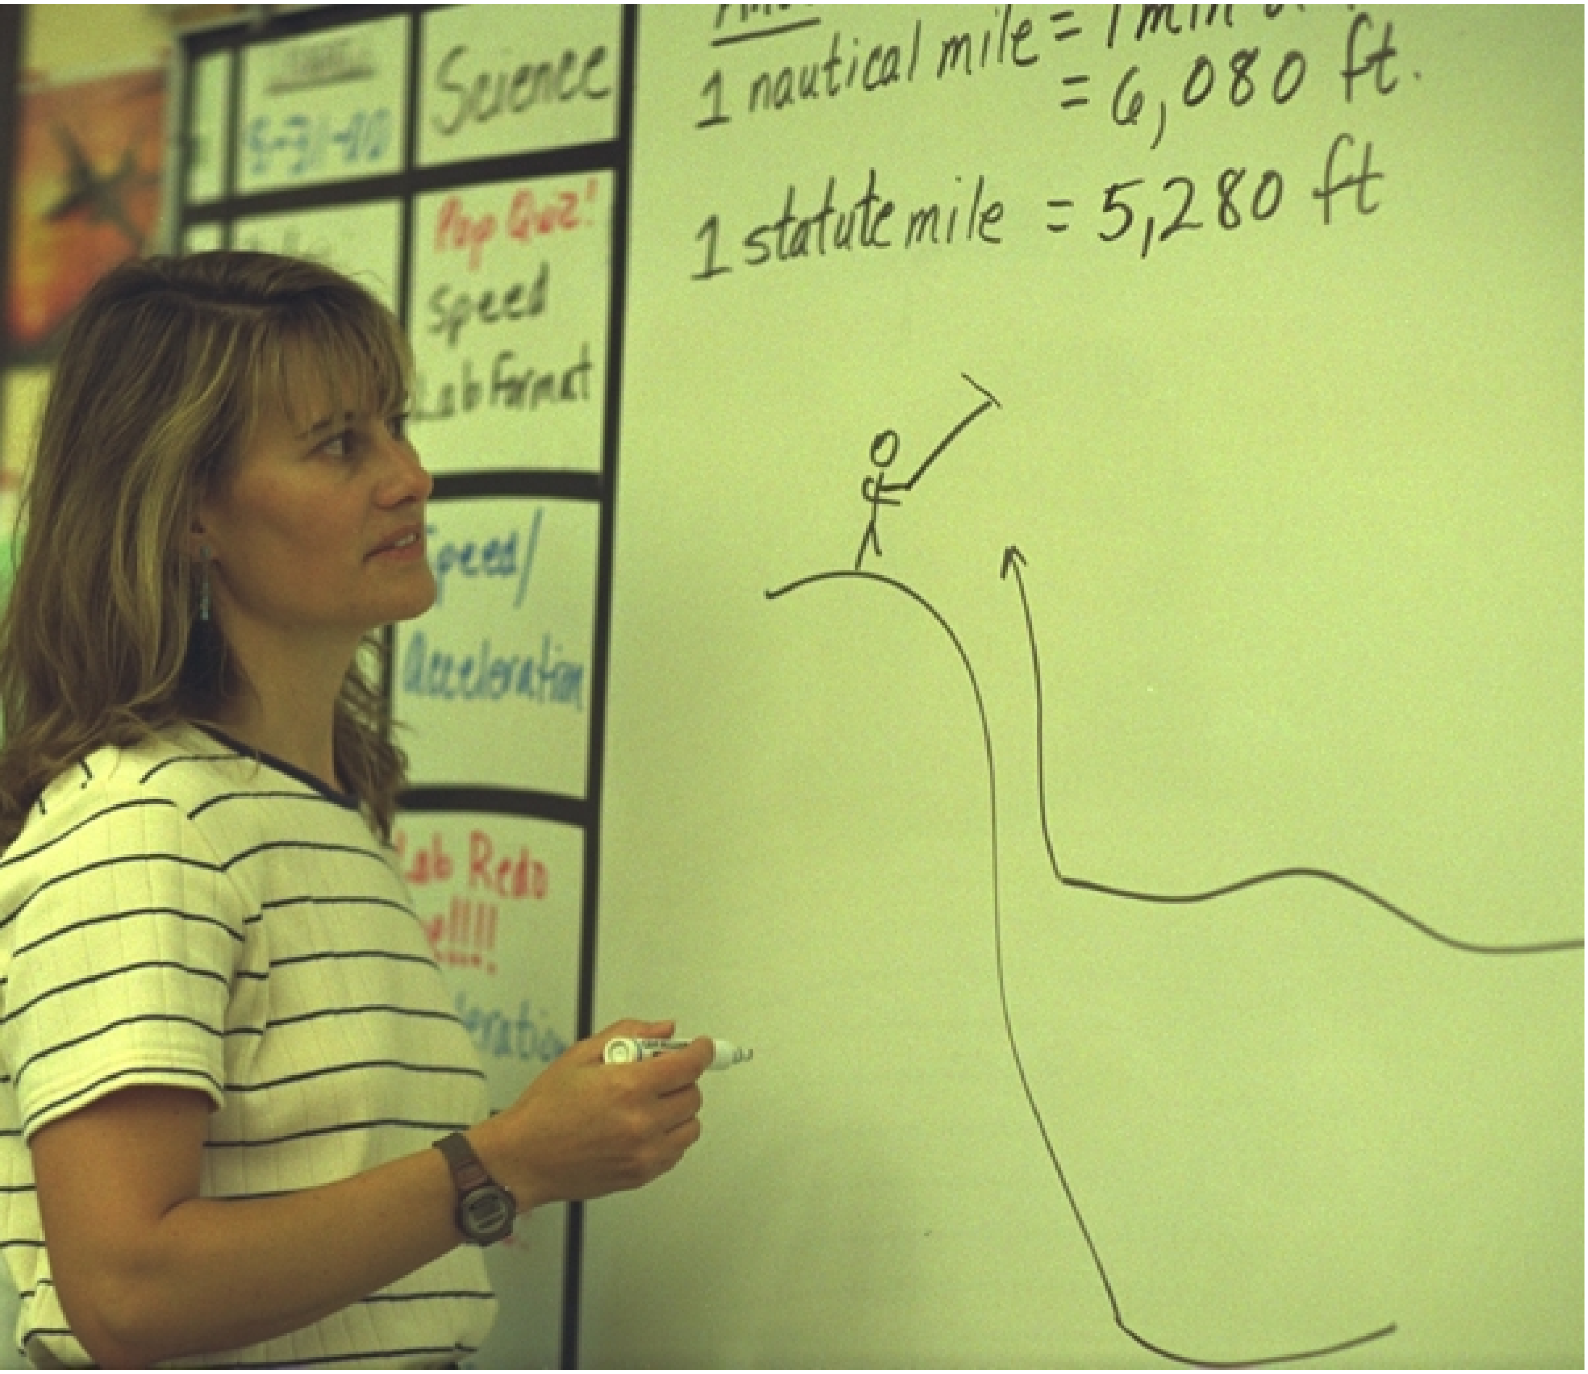

Diane Dunnigan, a science teacher at Marina Village Middle School, uses NASA FoilSim software to teach her students about aerodynamics.

Desktop wind tunnel

As with most creative software, developing FoilSim into a practical, easy-to-operate learning tool for students began as a problem-solving challenge for its creator, Tom Benson, an aerospace engineer, and member of the LT Project team at GRC. Headquartered at Ames Research Center in California, the LT Project is an integral part of NASA's High Performance Computing and Communications Program.

"The idea for FoilSim began in the mid-1990s," says Benson. "I'm a research scientist with 30 years' experience in aerodynamics. I worked on several projects, including NASA's National Aerospace Plane in the 1980s by developing computer models but it was time for me to change what I was doing.

"So I became very interested in the idea of simple, interactive software that would let the user experience what things affect an aircraft in flight. I wanted to be able to go to a computer with an on-screen slider or button and be able to change the wing shape easily, then see how the airflow changes when you alter it. I decided to create the software, use a high-speed computer, and keep it simple."

When Benson was in college in the 1960s, he says, research wind tunnels were accessible to undergraduate students. "But now," he says, "with limited funds, those research facilities are set aside for graduate and doctoral students."

With that fact in mind, Benson imagined an academic simulator that could be used on classroom workstations. He began programming.

He started by simulating the flow of air through a high-speed inlet. As he changed the flow and speed of air to supersonic velocities, he simulated the effects of shock waves on the inlet's performance. With his new software, Benson could simulate a wind tunnel in miniature.

"I could string shock waves together to simulate an F-15 fighter jet in supersonic

Matt (center) of Marina Village Middle School joins a classroom discussion about aerodynamics. Scott (right) listens attentively to his classmate.

flight that was pretty slick," Benson muses, "and I was able to get the whole computer program done in just six months. My career changed right there. From that time on, I went down the road of building educational software."

Benson's version of what would become FoilSim—called Vu-Foil—was targeted for college and university students studying aeronautics and related disciplines. The software was versatile and operated inexpensively on Unix workstations or PCs running the Linux operating systems.

Eventually, more than a dozen colleges and universities in the U.S. began using Vu-Foil and its derivatives in the classroom.

But it was only after Benson met high-school science teacher Roger Storm and computer technologist John Eigeneaur that he considered making the software useful for younger students.

"I liked the idea of making this technology available to teachers and students in the lower grades, letting them get comfortable with the concepts of flight and have fun at the same time. So the challenge was to change Vu-Foil into FoilSim," he notes.

Using a colorful display, FoilSim was conceived as a desktop wind tunnel so students could easily understand the physical forces affecting the lift of a wing using readily available PCs or Macs. And now that it's a part of a growing number of high-school science courses around the U.S., FoilSim has achieved a respectable status.

Classroom success

Students who might not otherwise find flying and aeronautics interesting become enthralled whenlearning about these topics in a hands-on way with FoilSim, says science teacher Zakarowsky.

Several teachers, she notes, report that FoilSim inspires them to rethink teaching science by doing more with non-conventional teaching aids. Having the right tools for the job makes a difference.

Diane Dunnigan (lower right) of the Marina Village Middle School instructs her science students about aerodynamics. Students (clockwise from lower left) Joe, Jennifer, Joe, Barrett and Luke discover the wonders of flight.

Joining Sheri Zakarowsky in introducing FoilSim to the classroom is Diane Dunnigan. A former aeronautics engineer, Dunnigan now teaches science to eighth grade students at Marina Village Middle School in El Dorado Hills, Calif.

"I teach flight science as an elective," Dunnigan says. "I grew up near Edwards Air Force Base where my father worked for Lockheed's Skunk Works, so I've been around aircraft all my life. While some of my students have been on commercial jets, few have actually been around them close up to touch them or to look at the different design characteristics. So FoilSim is ideal for learning about aircrafts and the concept of flight."

Dunnigan says FoilSim and her own enthusiasm for aviation have prompted several students to take flying lessons with the Civil Air Patrol at a nearby airport.

"FoilSim makes my students aware of what's happening to change the lift characteristics of the aircraft," she adds.

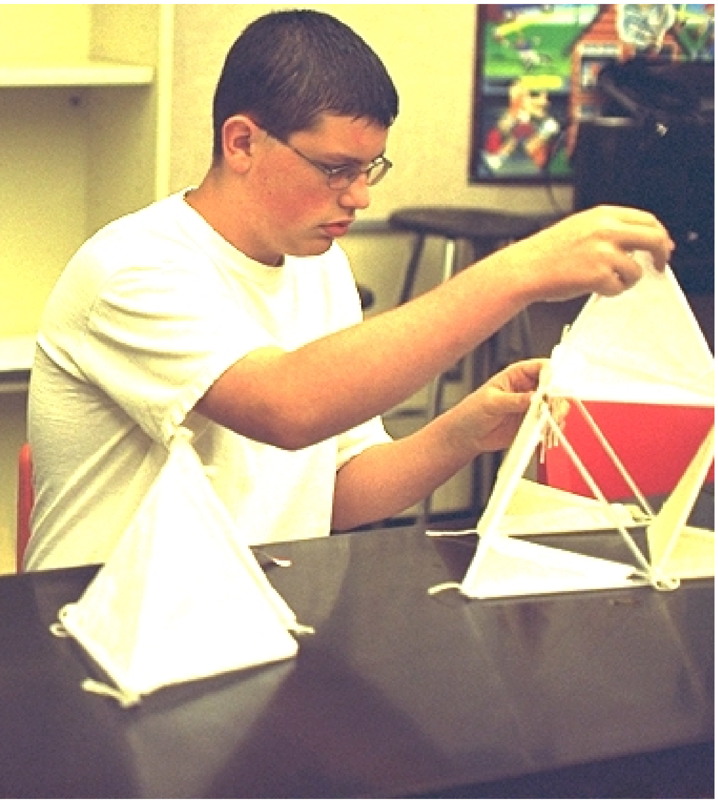

"What's useful, I believe, are the FoilSim graphics. Students see the actual lift changing. They find that really interesting because they see that the amount of force on the wing varies as they change other parameters. With that in mind," she adds, "I ask them to think about flight the next time they get on a plane. 'Sit by a window,' I say, 'and try to be over the wing. Look at the flight control surfaces; how do they move in takeoff versus descent?'"

Like Dunnigan, Zakarowsky's enthusiasm in the classroom is infectious to her students. She is especially proud of two students who have really taken their FoilSim lessons to heart.

"This year a couple of my flight science students are applying for NASA summer internships. This is the first time that ever happened at our small school," she says. "It's very satisfying when something like this results from everyone's dedication to developing FoilSim."

Zakarowsky will soon use the newest

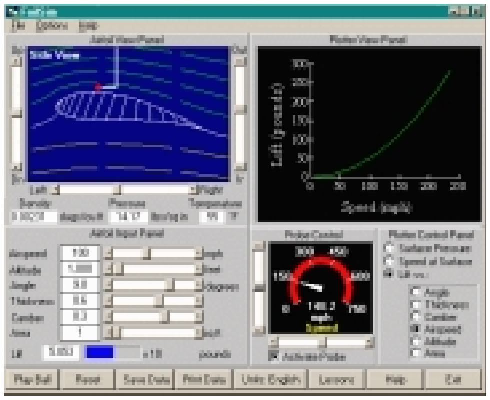

A FoilSim view panel shows a simulated wing and airflow. This NASA Learning Technologies software enables students to create airfoil shapes by moving sliders.

"FoilSim is ideal for learning about aircrafts and the concept of flight."

■ **Diane Dunnigan**
Marina Village Middle School

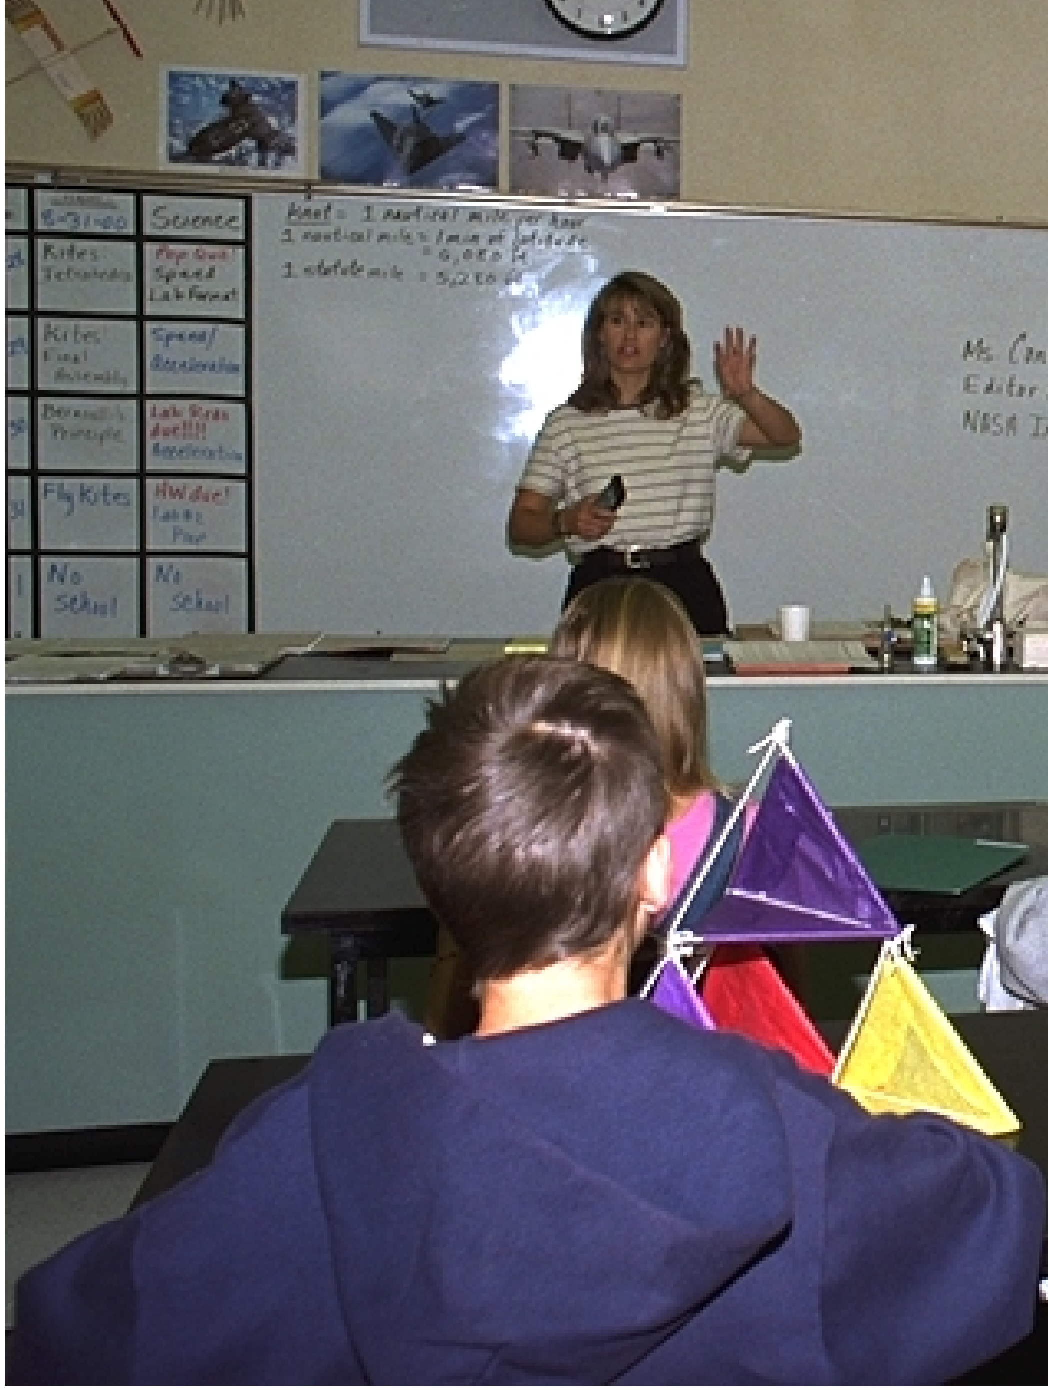

version of FoilSim in her classroom. This version includes a "stall mode" for the airfoil and also includes a model of the Martian atmosphere that helps students fly a Mars airfoil and conceptualize flight dynamics in alien atmospheres. Related software called EngineSim, which is currently available in its beta version and included with FoilSim, lets users test and modify four aircraft engines: a turbo jet, a jet with afterburners, a turbofan and a ramjet. More engine designs are planned.

"Like FoilSim, EngineSim is easy to download, interactive and a logical extension of the software," says LT's deputy project manager, Tom Dyson. "It makes sense to start with airfoils and then move on to study propulsion. We've received very high marks from the National Science Teachers Association for making FoilSim and EngineSim engaging—much so that students really do pause and think about the physical science behind flight," he says.

"The success of FoilSim and EngineSim," Dyson adds, "leaves everyone involved with a feeling of accomplishment. Plus the success of this educational product can be measured.

"We have 35,000 people each month visiting our website and downloading software," he says, "and the number of unique IP addresses" (the number of individual computer visits to the FoilSim Web site) "has been nearly three-quarters of

a million. We're pleased with the level of genuine interest the public is showing."

For creator Benson, there's more good news: NASA LT educational tools are helping teachers influence young people's career paths.

"I attended an airshow in Indiana a few years ago. A father and teenage son stopped by the NASA booth," says Benson with excitement in his voice. "The young man knew about FoilSim and was interested in learning about the aerospace field. I recommended a few colleges to him. Then someone snapped a photograph of us at the booth. Later I tacked up the photo on my office wall as a reminder to always focus on the kids."

However, Benson's story doesn't end there.

"Just a few weeks ago," he continues, "I was back at the airshow and a college student walks up to the booth, introduces himself, and says 'you don't remember me but I talked with you two years ago. I want you to know that you inspired me. I'm studying aeronautics in college.' Then it dawned on me—'hey, you're the student who inspires *me* everyday!' I flagged down our photographer. I now have a new photo of us, standing together, on my office wall."

Thanks to the tireless efforts of Tom Benson, Sheri Zakarowsky, Diane Dunnigan and many other men and women dedicated to inspiring young minds, a new generation of aerospace scientists and engineers is waiting in the wings.

Note: To download a copy of FoilSim and EngineSim software and the *Beginner's Guide to Aeronautics*, please visit NASA LT's website at: **http://www.grc.nasa.gov/WWW/K-12/aerosim/**

Science teacher Diane Dunnigan of the Marina Village Middle School exudes enthusiasm for aeronautics in her classroom. NASA FoilSim software, coupled with her own instruction, prompted several students to take flying lessons.

Virtual space weather from

By Adam Frank

The perfect storm packs the energy of 100 million atomic bombs

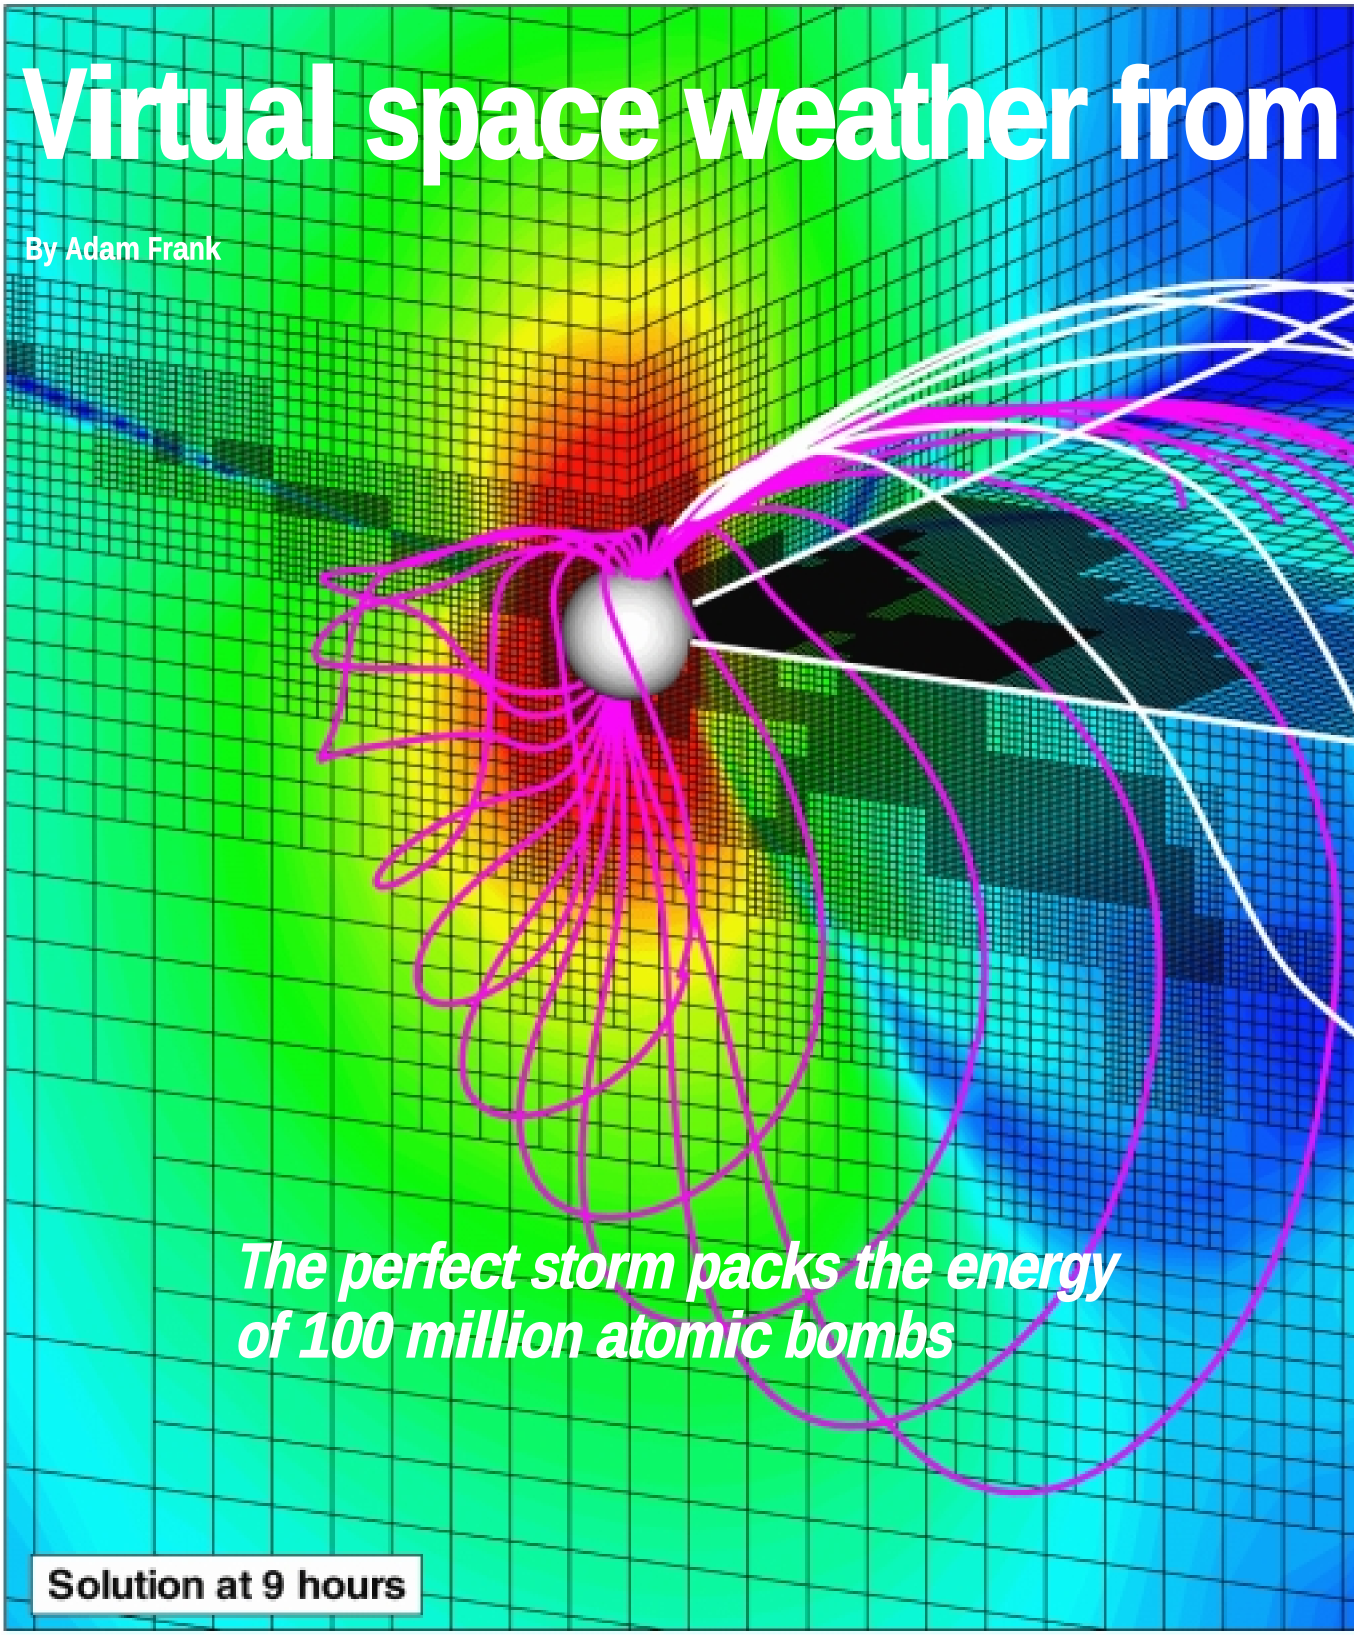

Simulation by D.L. De Zeeuw, T.I. Gombosi, C.P.T. Groth, K.G. Powell, and Q.F. Stout

Coronal mass ejections (CMEs) consist of ionized hydrogen and electrons that burst out from the sun at irregular intervals. Shown are 3-D magnetic field lines (white and magenta) and magnetic field strength (low=blue to high=red) nine hours after the initiation of a CME.

the sun to Earth

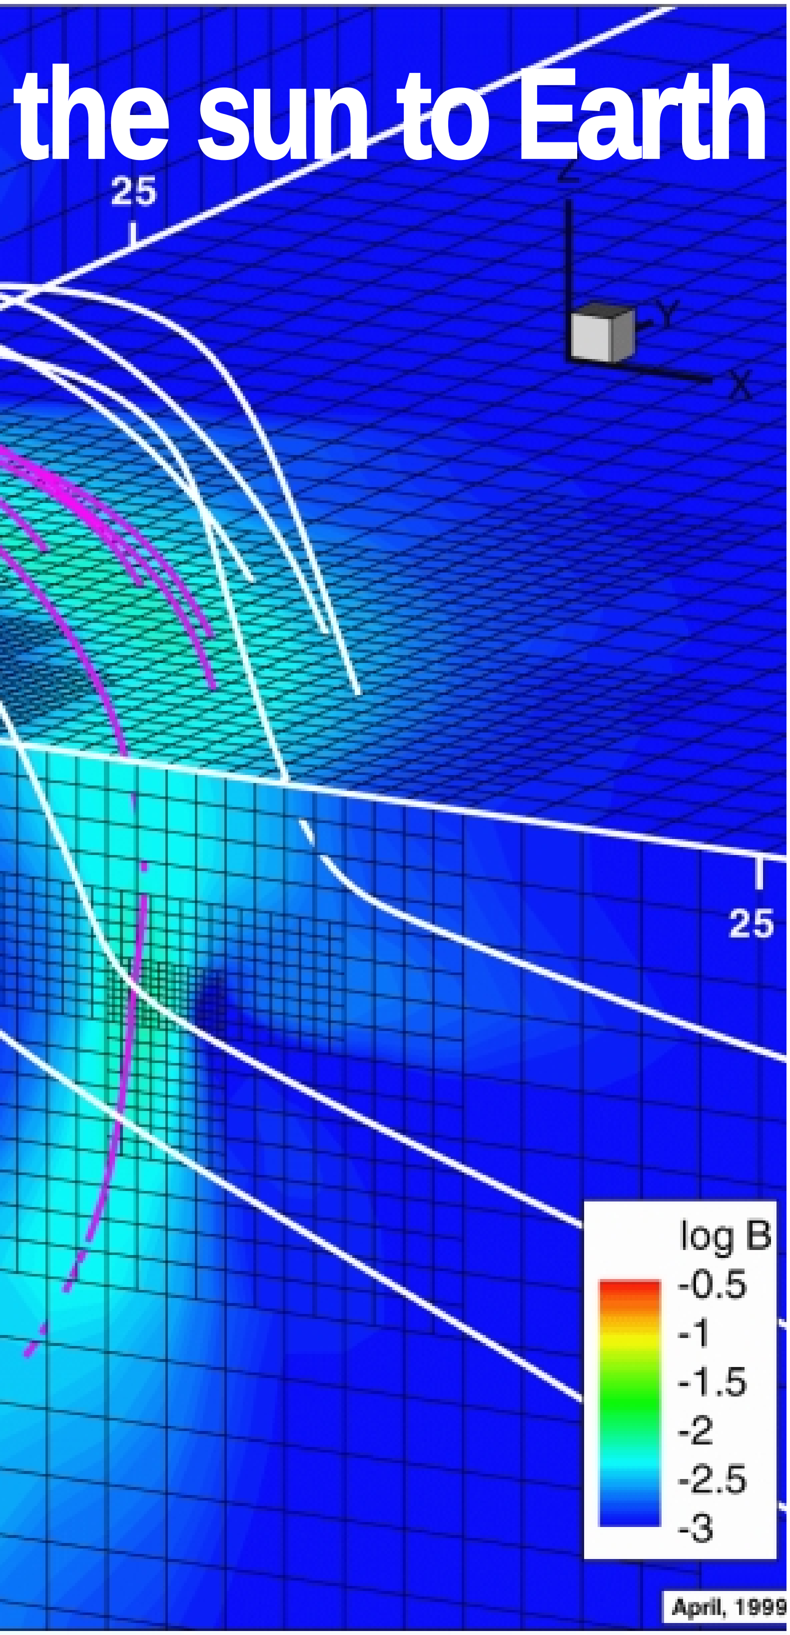

Photos by Tom Treuter

Magnetic field lines from a CME appear to envelop users of the Cave Automatic Virtual Environment virtual reality tool. University of Michigan's Scott Feldman and Aaron Jacobovits developed the visualization application, which allows researchers to fly through the simulation results and find features.

Bad weather—it came on fast and without warning. Its power was devastating. In its wake, more than six million people were left without power and in darkness.

It was a spring evening in 1989 when the severe space storm smashed into the Earth. More than a billion tons of electrically charged gas, which was blown off the sun three days earlier, plowed into the Earth's protective magnetic blanket, sending a torrent of current running toward the poles. In the Canadian province of Quebec, the flood of energy surged into electric lines, causing a system-wide power failure. The time from the onset of the storm to total blackout was just 90 seconds. When the lights finally came back on, millions of dollars in damage had been attributed to the storm.

Humanity is maturing into a space-faring race. As we move from exploring space to using it to benefit mankind, the danger of some aspects of *space weather* has slowly been recognized. Space weather is a new term in the lexicon of science. It refers to surges of matter and magnetic energy blown off the sun that periodically sweep across interplanetary space. Twenty years ago the only consequence of space weather was a good aurora. Now we live in the age of the global village, where beepers, cell phones and the Internet keep many of us connected 24 hours a day. It is an age in which satellites are vital tools in communication and commerce, an age when the risk posed by a space storm reaching the outer atmosphere has become all too real.

Wherever there is weather, there is also the need for prediction. While Earth-bound storms travel a few thousand miles at most, space storms must cover the 100 million kilometer distance from the sun to the Earth. There is a lot of room for error in 100 million kilometers, as well as a lot of unknown physics to tame before a space weather prediction system can be put in place. This is the challenge Tamas Gombosi and his University of Michigan team of High Performance Computing and

Simulation by D.L. De Zeeuw, T.I. Gombosi, C.P.T. Groth, K.G. Powell, and Q.F. Stout

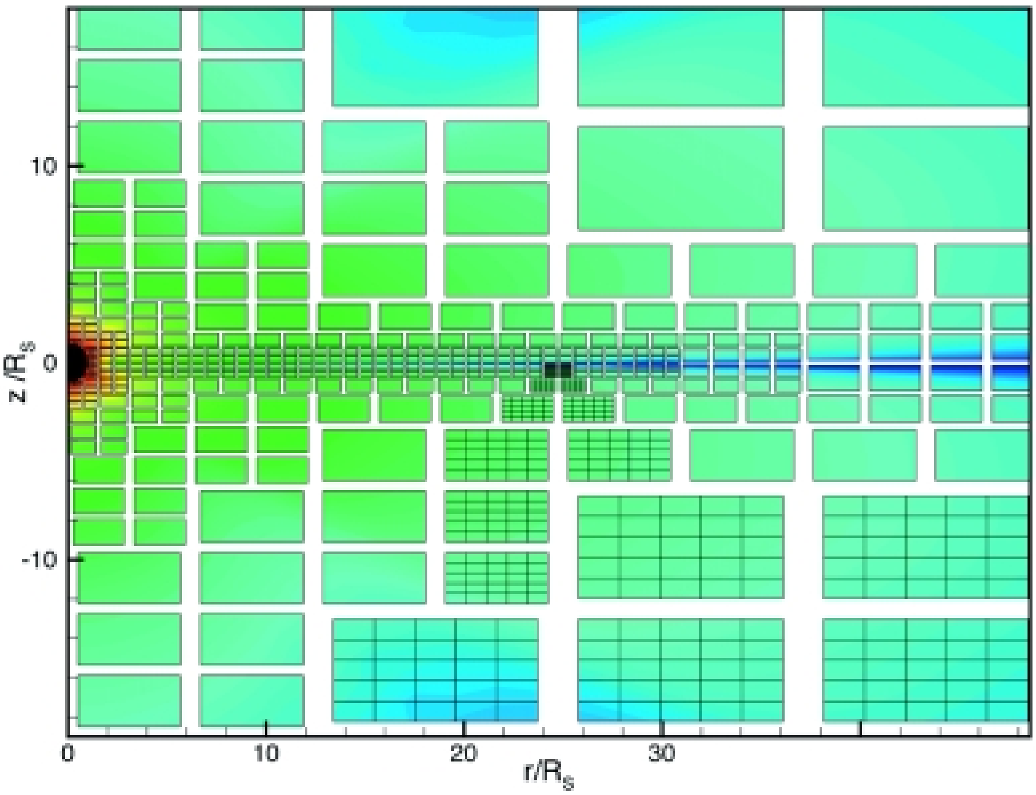

The University of Michigan's models slice simulated space into a grid of cells. Using adaptive-mesh refinement, cells automatically divide into smaller cells as finer-scale activity occurs in them. Cells are grouped into blocks for greater efficiency on parallel computers.

Tamas Gombosi

Using computational systems that push the boundaries of state-of-the-art techniques, Gombosi and his collaborators hope to chart the pathways of space weather from the sun all the way to the Earth.

Communications (HPCC) Earth and Space Science (ESS) researchers believe they can tackle. Using computational systems that push the boundaries of state-of-the-art techniques, Gombosi and his collaborators hope to chart the pathways of space weather from the sun all the way to the Earth. Their recent success with the first sun-to-Earth simulation of a single space storm brings that goal and NASA's hopes for a working space weather prediction system one step closer to reality.

On being the right size

"The scales involved are enormous and enormously different," says Gombosi, professor of space science and aerospace engineering at the University of Michigan and principal investigator (PI) on the ESS grant. He explains the key feature that has limited space weather prediction: the sun has a radius of a hundred thousand kilometers, but the radius of the Earth is 100 times smaller. The distance between the two bodies, however, is 200 times larger. "Given the disparities in the size scales," says Gombosi, "tracking space weather is a very challenging computational problem."

The most powerful form of space storm is what researchers call a coronal mass ejection (CME). In a CME, a piece of the sun's outer atmosphere is blown into space at more than 500 kilometers per second. The power unleashed in a CME is staggering. Each CME launches billions of tons of matter, with an energy equivalent to 100 million atomic bombs, into space. Tracking the details of a CME's evolution as it crosses interplanetary space is the first step in a space weather prediction system. The next requirement is following the details of an evolved CME's collision with the Earth and its subsequent effect on the near-Earth space environment. The Michigan project is ambitious, aiming to simulate details of both phases of space weather.

The principal tool of the project is "BATS-R-US," the imaginatively named computer code designed by the Michigan group to rein in the problem of space weather and its disparate scales. BATS-R-US stands for Block-Adaptive-Tree Solar-wind Roe Upwind Scheme. Embodied in that mouthful of a title is a series of innovative features that combine to make the Michigan team's effort unique in the world of computational physics.

The adjective *Adaptive* in BATS-R-US is one of the most important. It represents a significant advance in computational fluid dynamics in the last decade. Adaptive mesh refinement (AMR) is the capacity of a numerical code to put its resources where the action is. Resolution, which is the ability to capture the necessary level of detail, is among the most critical aspects of a simulation and often the best measure of its accuracy. If you can imagine creating a mosaic of the Mona Lisa with just four tiles, you have some idea of why low-resolution simulations are not worth much in terms of science. There is, however, an unfortunate trade-off in computational science between resolution and computer processing time. High-resolution simulations are also highly expensive.

"AMR codes get around this problem by dynamically changing the resolution," explains Ken Powell, co-principle investigator on the Michigan HPCC grant. Powell emphasizes that the BATS-R-US code can keep the resolution of the simulation high just where it is needed as the CME drives across interplanetary space toward the Earth. "The smallest zone in the simulations is 1/16 the size of the sun," he says, "while the largest is more than eight times the size of the sun. The resolution at any particular location in the grid changes, depending on what is happening there." The capacity to intelligently increase, and then relax, resolution as the simulation proceeds saves the Michigan group tens of thousands of computer processing hours. Without it the simulations and space weather prediction would become impossible.

"Our calculation was over in less time than a real storm would have taken to reach the Earth. We were faster than a speeding CME."

■ **Darren De Zeeuw**
University of Michigan

Darren De Zeeuw, a research scientist in the Michigan group, explains how their code called BATS-R-US was written from scratch for optimal speed.

Tamas Gombosi, University of Michigan professor of space science and aerospace engineering, leads this HPCC Grand Challenge investigation.

"We started out by simulating the normal solar wind in a domain that stretched from one solar radius all the way out to 215 solar radii (just beyond the Earth)."

■ **Tamas Gombosi**
University of Michigan

A journey of a thousand steps

Before they could simulate a space storm, the Michigan group first had to get the interplanetary equivalent of a beautiful day correct. A CME's propagation from the sun to the Earth could never be modeled correctly unless the space through which it propagated had the right properties. "We started out by simulating the normal solar wind in a domain that stretched from one solar radius all the way out to 215 solar radii (just beyond the Earth)," says Gombosi. The solar wind that fills interplanetary space has a considerable structure, even when the sun is quiet. Including a model of heating within the solar atmosphere helped the team get the background right. "We had to tweak our parameters a bit," says Gombosi, "but in the end we came up with an excellent fit to the solar wind data for the quiet sun." With the smooth seas worked out, Gombosi and his collaborators were ready to brew up a storm.

Exactly how a CME forms on the sun remains a mystery, so to create a storm in

Simulation by D.L. De Zeeuw, T.I. Gombosi, C.P.T. Groth, K.G. Powell, and Q.F. Stout

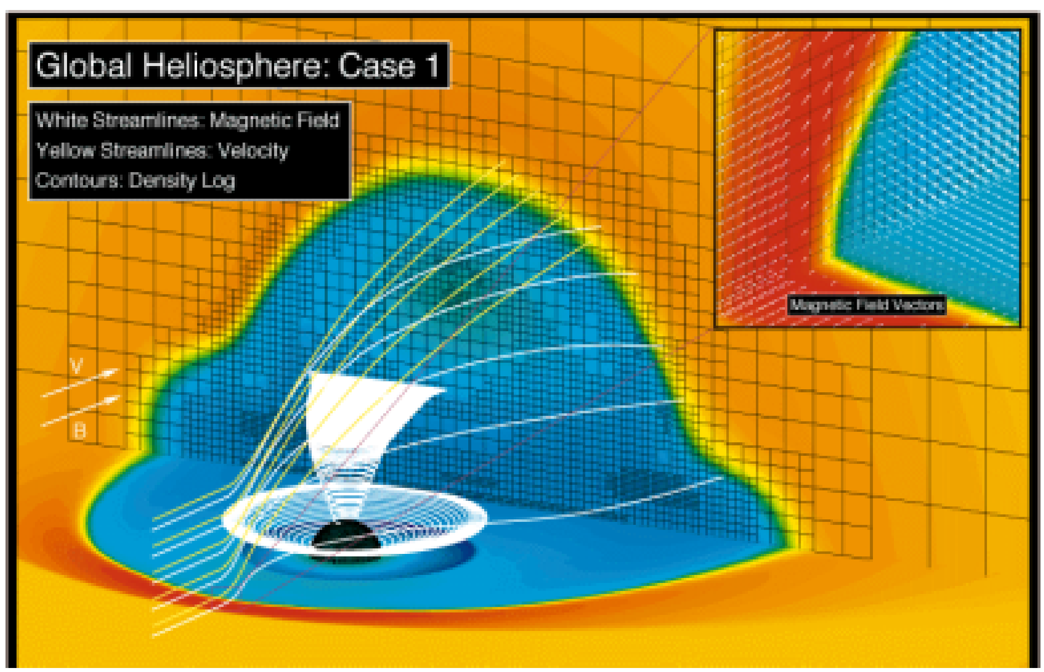

The outer heliosphere's slightly bent field lines trace the interstellar magnetic field while the spiral field lines convect out with the solar wind. The inset shows a detailed view of the bow shock.

their simulations, the Michigan group helped nature along a little. “At the beginning of the simulation, we injected a small region of high density and closed magnetic field loops into the corona,” says Gombosi. “It disturbed the delicate balance there, and we ended up with a nice CME.” As the virtual CME exploded into interplanetary space, the researchers watched carefully to see where it would cross the Earth’s orbit. Then, exercising a bit of God-like control, they reran the simulation, starting the Earth in just the right position to ensure a collision between it and the CME.

To catch the physics of the Earth-CME collision, the Michigan team faced a very tiny problem. The smallest cell in the team’s sun-to-Earth calculation was 440 times the size of our planet. This is the Mona Lisa mosaic problem all over again. “We needed to recreate the Earth’s grid so we could see what happened as the CME passed,” says Gombosi. The Michigan group solved its problem with two related simulations. First, team members performed the large-scale simulation that tracked the CME as it crossed the Earth’s orbit. Then they fed the values of the magnetic field, velocity and density in the passing storm into a second BATS-R-US simulation that focused on the near-Earth environment. “The smallest cell in that calculation was only a quarter of the Earth’s radius,” explains Gombosi. With the two linked calculations, the Michigan group could claim success in fielding the first fully 3D calculations for the evolution of space weather from the sun all the way to the Earth.

One of the more promising aspects of the Michigan group’s success is that the simulations were faster than the duration of a real space storm. “Our calculation was over in less time than a real storm would have taken to reach the Earth,” explains Darren De Zeeuw, a research scientist in the Michigan group. “We were faster than a speeding CME.”

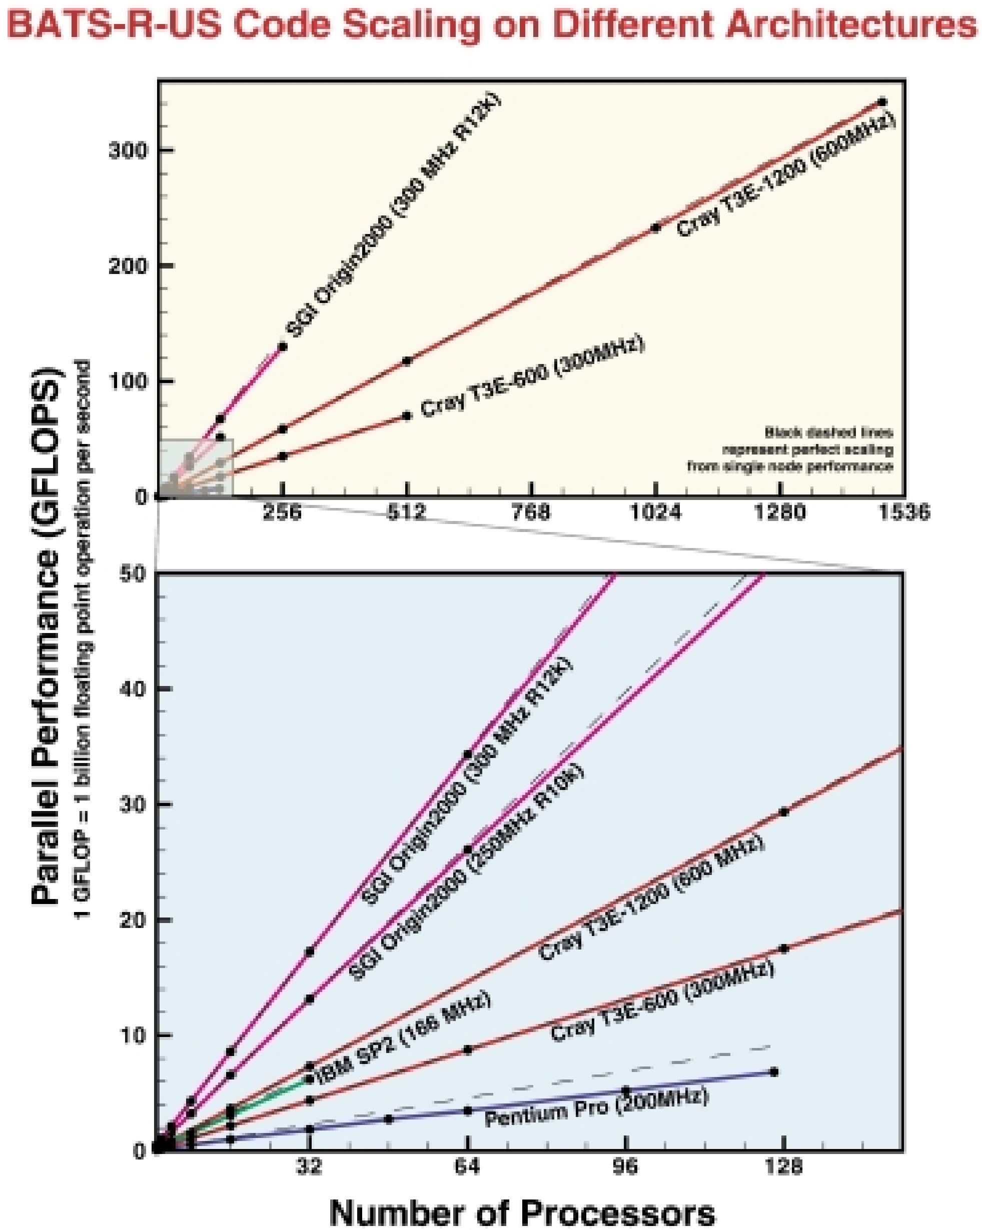

Code performance in billion floating point operations per second (gigaflops) on a variety of machines. Comparing black dashed line to colored line for a computer can show, for example, whether running on 32 processors will run 32 times faster than one processor.

The adaptive parallel universe

The extraordinary speed of the BATS-R-US code comes directly from the “Block” adjective in its name and represents one of the most innovative aspects of the Michigan team’s effort. “We knew we would be using parallel computers when we began the project,” says Powell, “so we tried to optimize the construction of the adaptive mesh refinement with the parallelization in mind.” A parallel computer is programmed to take a single problem and distribute it across many processors at once. If the researchers do their homework right, they can get the code to speed up by a factor nearly equal to the number of processors. The traditional approach to solving problems with parallel machines is *domain decomposition*. Each processor gets a chunk, or slab, of the physical domain that needs to be simulated. “Domain decomposition is fine when the grid is fixed,” explains Powell, “but it’s not so straight forward with an adaptive mesh code because the domain is changing as the resolution in a region is increased and then relaxed.”

Facing the problem head-on, the Michigan team dreamed up a novel solution. The most important thing in a parallel calculation is that no processor be left idle. To ensure this does not occur in the BATS-R-US code, the computational space is divided into blocks of grid cells. Each block was a fixed size, say 10X10X10 grid cells. The block becomes the *unit* of the calculation. Each processor is then tuned to work on so many blocks at once. If the resolution inside a block needs to increase, it is simply divided into new 10X10X10 blocks. The new blocks are redistributed across the processors to ensure each CPU has an equivalent workload. “We never have an unbalanced load on more than one processor,” explains De Zeeuw. “That means latency didn’t hurt us at all.” The Block strategy has proven enormously successful for the Michigan group. “We really have gotten AMR to work effectively on parallel machines,” continues De Zeeuw.

“The smallest zone in the simulations is 1/16 the size of the sun,while the largest is more than eight times the size of the sun. ”

■ **Ken Powell**
University of Michigan

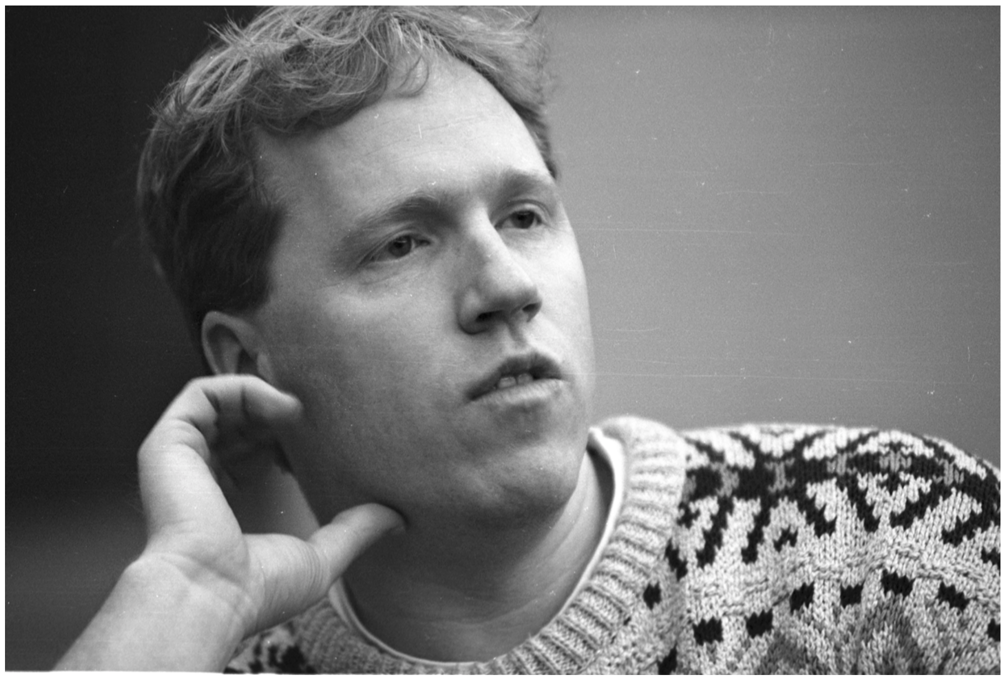

Ken Powell and the UM group optimized the performance of adaptive mesh refinement with parallel computers in mind.

“We see the speed of the code increase linearly with the number of processors, all the way up to massively parallel machines.” De Zeeuw explains that the group recently ran the code on a machine with 1490 processors and achieved more than 99 percent of the scaling speed up.

The real world

The group’s success marks only the beginning of its efforts to create a tool for real space weather prediction. “The model we used to describe the Earth is really too simple,” says Powell. The physics of the near-Earth environment is more complicated than that of interplanetary space. There are variations in the degree of ionization (how many free electrons are running around) and in the chemistry in the upper layers of Earth’s atmosphere. The coupling of these layers to the planet’s magnetic blanket (the magnetosphere) is also quite complicated. To produce models that can accurately follow the Earth’s response to a CME, the Michigan group has teamed up with a number of the nation’s best researchers in space physics.

“We are working with NCAR (the National Center for Atmospheric Research) and Rice University,” says Gombosi. He explains that these groups have developed sophisticated codes to simulate different parts of the Earth’s upper reaches. The Rice team has a numerical model for the inner regions of the magnetosphere. The NCAR

Simulation by D.L. De Zeeuw, T.I. Gombosi, C.P.T. Groth, K.G. Powell, and Q.F. Stout

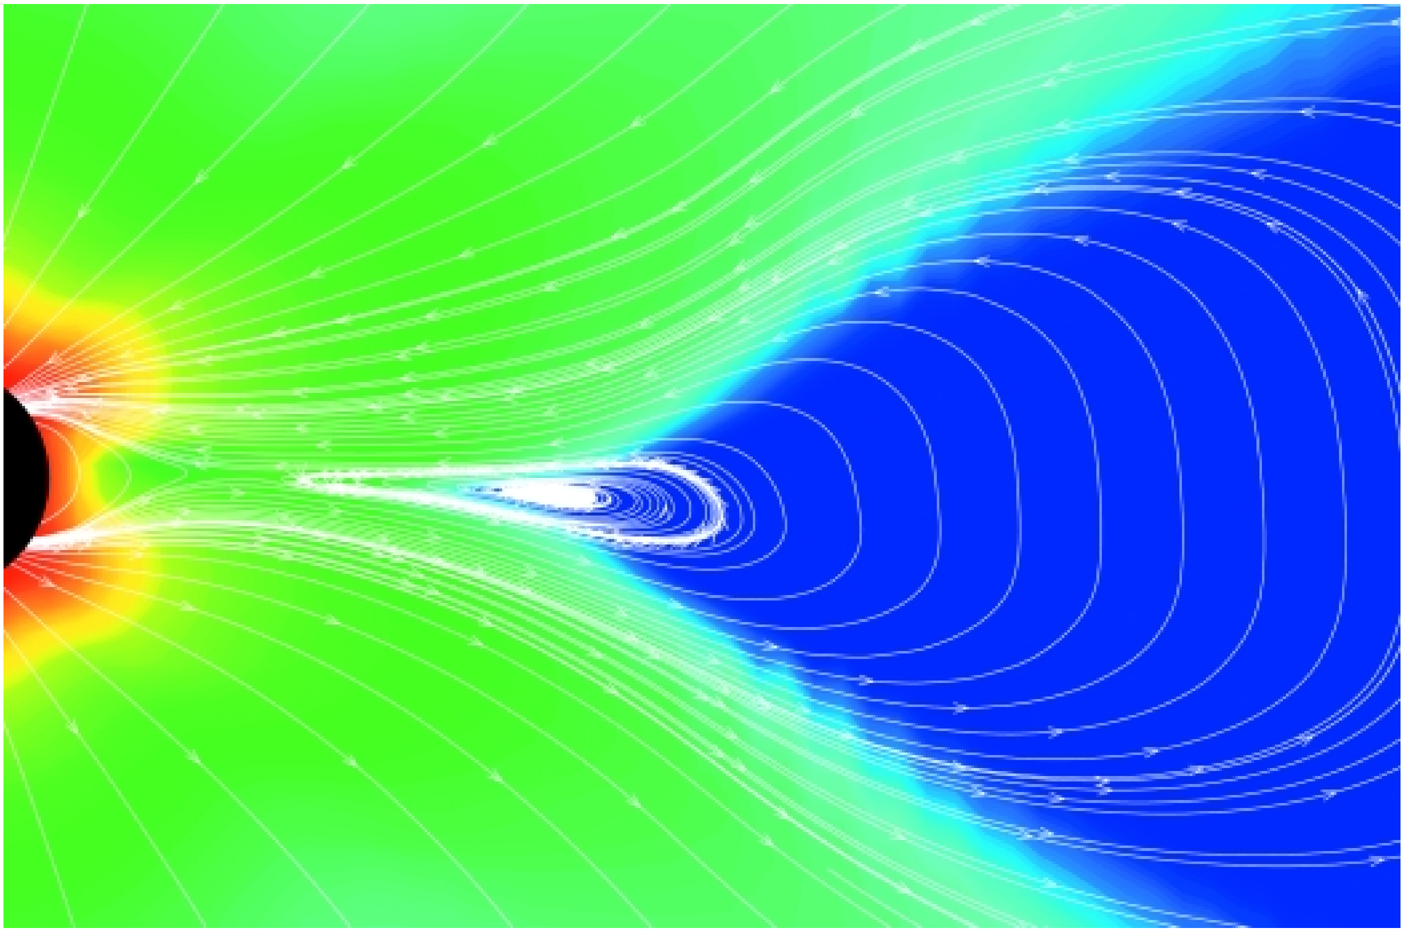

Coronal mass ejections (CMEs) burst out from the sun at irregular intervals. Shown are magnetic field lines (white) and magnetic field strength (low=blue to high=red) 20 hours into a CME simulation.

team has developed a numerical code called TIMEGCM or Thermosphere-Ionosphere-Mesosphere-Electro-Global-Circulation Model. As the name implies, it was built to handle a good portion of the outer regions of the Earth's atmosphere.

Linking the BATS-R-US codes to models built by other groups, but not optimized for parallel machines, poses a considerable challenge for the Michigan researchers. "We are basically taking their physics kernels and building new numerics around them," says Powell. "We will act as the glue that holds the different codes together." The payoff promises to make the effort more than worthwhile because the final, coupled codes will track current flow and chemistry in detail as a CME pours energy into the near-Earth environment. This is a feature that a true predictive system must include.

"When you start doing something as ambitious as this," says Gombosi, summarizing his team's approach to the problem, "you always have to make compromises in terms of the physics. Hopefully we are now moving to the point that we can make fewer of these compromises and get closer to the real systems." With society becoming ever-more-dependent on the environment hundreds to thousands of miles from the Earth, the Michigan team's progress holds the promise of a true sea change in our approach to space weather.

Reusable launch vehicle effort requires advanced computing resources

By Lou Varricchio

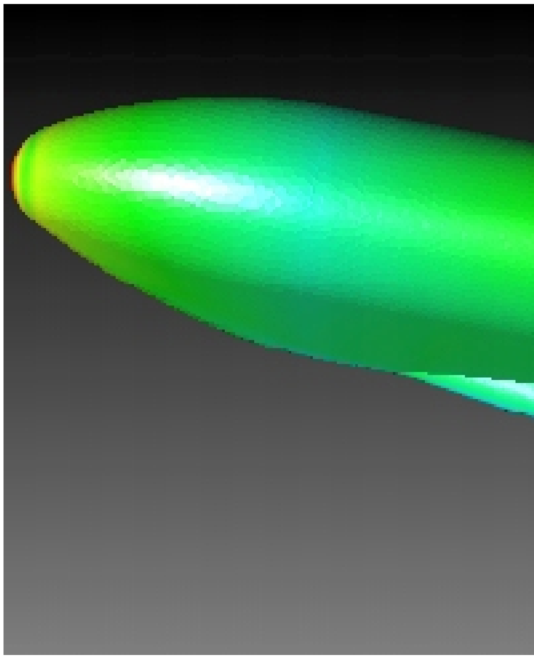

Imagine a space vehicle that behaves as both aircraft and spacecraft. It takes off like an airplane, delivers a satellite to orbit and then returns to Earth for a new payload within a few hours.

By the year 2040, according to many experts, a fourth generation reusable launch vehicle (RLV) will ultimately blur the distinction between aircraft and spacecraft. But before dreaming of a *fourth* generation reusable launch vehicle, how do we get from today's *first* generation, partially reusable launch vehicle — the space shuttle — to a *second* generation, truly reusable space vehicle? Designing a low-cost, reusable space vehicle will involve many important steps along the road toward gaining faster, cheaper, better access to space.

Making regular trips to orbit possible, RLVs will produce myriad benefits including new commercial and research opportunities plus an increased understanding of our environment via cheap access to space. But long before these vehicles are ever built, and before the dream of low-cost access to space is realized, their engines, structures and aerothermodynamic characteristics will be modeled and studied in 3-D detail using highly advanced computing tools.

Designing and building space vehicles has been an expensive proposition so far. However recent advances in material, propulsion and computing technologies will contribute to NASA's goal of lowering launch costs and help usher in a new age of space transportation. The Computational Aerospace Sciences (CAS) Project of NASA's High Performance Computing and Communications Program is involved, along with other agency teams, to help make routine access to space a reality.

NASA's second generation RLV effort requires sophisticated computational tools early in the design phase. Ascent-to-orbit and return-to-Earth are critical aspects to consider when designing space vehicles that combine features of both aircraft and spacecraft. Getting to and from orbit are the portions of a space vehicle's flight envelope — the bounds within which a flight system can operate — that impose a lot of physical stress. That's why researchers who can assist engineers, by developing efficient computing resources for various simulations and analyses, will make important contributions to second generation RLV development.

NASA's CAS Project has led the way in high-performance computational technologies for NASA's aviation projects. Recently, CAS was charged with helping the agency meet its goal in achieving low-cost access to space. As with CAS's computational accomplishments to date, these new space challenges will require a multidisciplinary approach to a variety of vehicle design and analysis problems.

CAS's three major areas of research and development are advanced aerospace applications, computing testbeds and system software that works as the interface between applications and end users. CAS's expertise in developing tools for 3-D simulations of propulsion systems and full-vehicle simulations grew out of its long-time, aviation-related work. This expertise is a valuable resource to NASA Aerospace Technology Enterprise efforts such as the second generation RLV. In the future, CAS computational technologies will enable the entire lifecycle of a new space vehicle to be simulated saving millions of dollars in development costs.

On the same path toward achieving regular access to space by a new generation RLV is the *first* generation RLV — the space shuttle. While the 1970s designed shuttle is the recognized forerunner, it is only partially reusable. Still, the shuttle remains a valuable, pioneering space vehicle with much to teach us. Even as each lesson learned informs us about what to anticipate when designing a second generation RLV, the shuttle fleet will continue to serve as NASA's link to the International Space Station.

Over the next few years, NASA plans to modify the shuttle to make it a safer, more efficient space vehicle. CAS will play a part in the effort by making computational studies of the shuttle's flight characteristics and structural integrity during various launch-abort scenarios. Upgrades to the shuttle, from avionics to airframe, will be incorporated into the fleet by 2005. These

Computations and image by Susan Cliff

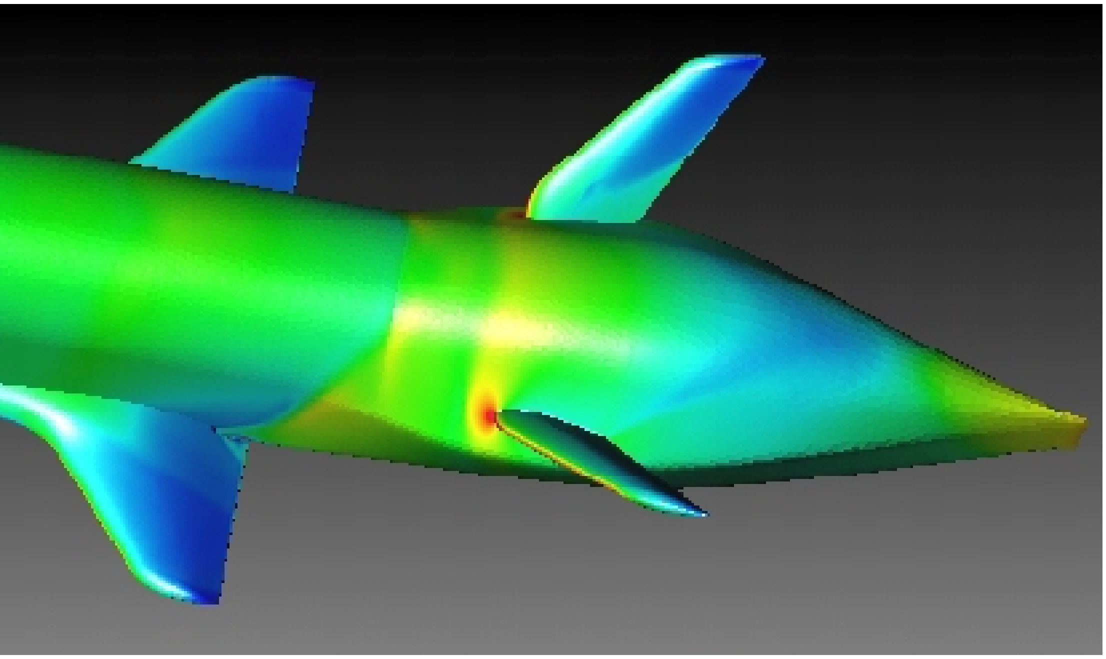

The X-37 aerospace vehicle will demonstrate new space technologies. In this flight simulation, the air pressure on the craft's surface appears in color.

upgrades will extend the operational life of the shuttle until its replacement is on the launch pad.

In the meantime, as the shuttle's role expands as the International Space Station's prime workhorse, researchers at NASA centers will continue work on follow-on, fully resusable space vehicles to service the new space station as well as future orbital outposts.

Early in 1999, CAS teams at NASA's Ames, Glenn and Langley centers began focusing their attention on important computing areas relating to advanced space vehicle development especially, for the near term, the second generation reusable launch vehicle.

One area demanding R&D attention is aerodynamic optimization — getting the most of a space vehicle's design so it performs well as it leaves and reenters the atmosphere. At Ames Research Center, researcher John Melton and his CAS team are supporting several NASA groups in meeting the computational challenges relating to space vehicle development.

Vehicle speed range is so great with proposed RLVs, he notes, that "you can't guarantee the vehicles will have good handling characteristics as you get ready to land." The challenge, he says, is in finding shapes that perform well during the entire flight envelope. CAS's computational tools will help find the way and enable engineers to reduce costs by optimizing the best designs and new materials.

Having concluded work on NASA's High-Speed Civil Transport, the CAS team at Langley Research Center in Virginia now focuses on computing requirements for solving intricate problems relating to future RLV designs.

"There are many unexplored problems associated with proposed RLVs," says researcher Jaroslaw Sobieski at Langley. "It's a type of vehicle that passes from subsonic and transonic to supersonic flight regimes rapidly. It travels through the atmosphere to outside the atmosphere, and back, with all the associated thermal effects. Plus an RLV's supporting structure weight has to be minimized to make way for payload needs. All this means designers must optimize the vehicle from the start. This will demand tremendous computing power to properly analyze all the engineering disciplines and design variants involved. It's exciting and we are just at the beginning."

"For the second generation RLV, we have to gear up to compute with a very large number of processors," Sobieski says. "For example, consider an advanced machine with tens of thousands of processors. We can't simply take existing codes, plant them on such a new machine and expect a proportional increase in speed. We have to redevelop codes from scratch. Parallel to that challenge, we're using existing multiprocessor codes that can engage a limited number of processors, in the range of hundreds, so that we can exploit this new capability to solve the problems currently at hand."

The SHARPer image

A recent experiment that could contribute to a breakthrough in aerodynamic optimization consists of a rocket nose cone called SHARP, short for Slender Hypersonic Aerothermodynamic Research Probe.

During the summer of 2000, NASA Future-X Project researchers from the Marshall Space Flight Center tested the unique probe with its ultra-high temperature ceramic jacket, a material that could radically improve thermal protection when a space vehicle reenters the atmosphere. Lofted to space by a modified Minuteman III rocket, SHARP yielded strong data for ongoing studies. The effort, which involves a coordinated effort between Marshall Space Flight Center, the Ames Research Center, Sandia National Laboratory and industry, may soon offer surprising payoffs.

According to John Melton, the CAS team is using high-performance computing tools to assist SHARP researchers in analyzing the experiment's aerodynamics, or forces acting on a vehicle whenever it leaves or reenters the atmosphere. "With the new ceramic material developed at NASA Ames, we can now think of new space-vehicle images — images that will include sharp, leading edges,” he notes. “If you look at the space shuttle’s leading edges, you’ll notice they’re all rather blunt. The dream of hypersonic aerodynamics has always been to develop *pointy* leading-edge vehicles; they’re the best shapes for the job. We couldn’t design these shapes until new ceramics became available. The new space vehicle shapes will help reduce the cost of getting into and back from space.”

As should be evident by now, the faster, cheaper, better road to space isn’t paved with just a single good idea or vehicle.

Also managed by the NASA Marshall Space Flight Center, the new X-37 vehicle is a testbed for a variety of space technologies. It’s hoped that the 28-foot-long pilotless craft will demonstrate the feasibility of making space travel affordable and reliable. CAS researchers at Ames are working with the X-37 team, employing advanced computational simulation tools to assist in vehicle design analysis and three-dimensional modeling. If all goes well, the X-37 concept should lead to space vehicles with fewer parts and more capabilities.

Even with fewer parts, a complex RLV design will present daunting challenges to engineers working on various segments of the vehicle. This kind of space vehicle will require a coordinated effort between humans and computers to solve a variety of engineering problems. In the RLV era, engineers must think in new ways as they integrate their many disciplines in tackling space-vehicle design problems.

A segment of CAS’s efforts to meet this challenge centers on a form of machine intelligence called a multi-agent system. Cooperative problem- solving — across disciplines—will be important during the design phase of future space vehicles. Multi-agent systems will act as a distributed computer network of semi-autonomous processing elements that work together to solve a single problem. Intelligent agents communicate with each other to perform collective tasks on behalf of the humans involved.

“The more sophisticated an agent can be in problem-solving,” says Abhi Deshmukh, a CAS-funded researcher at the University of Massachusetts/Amherst’s Fundamental

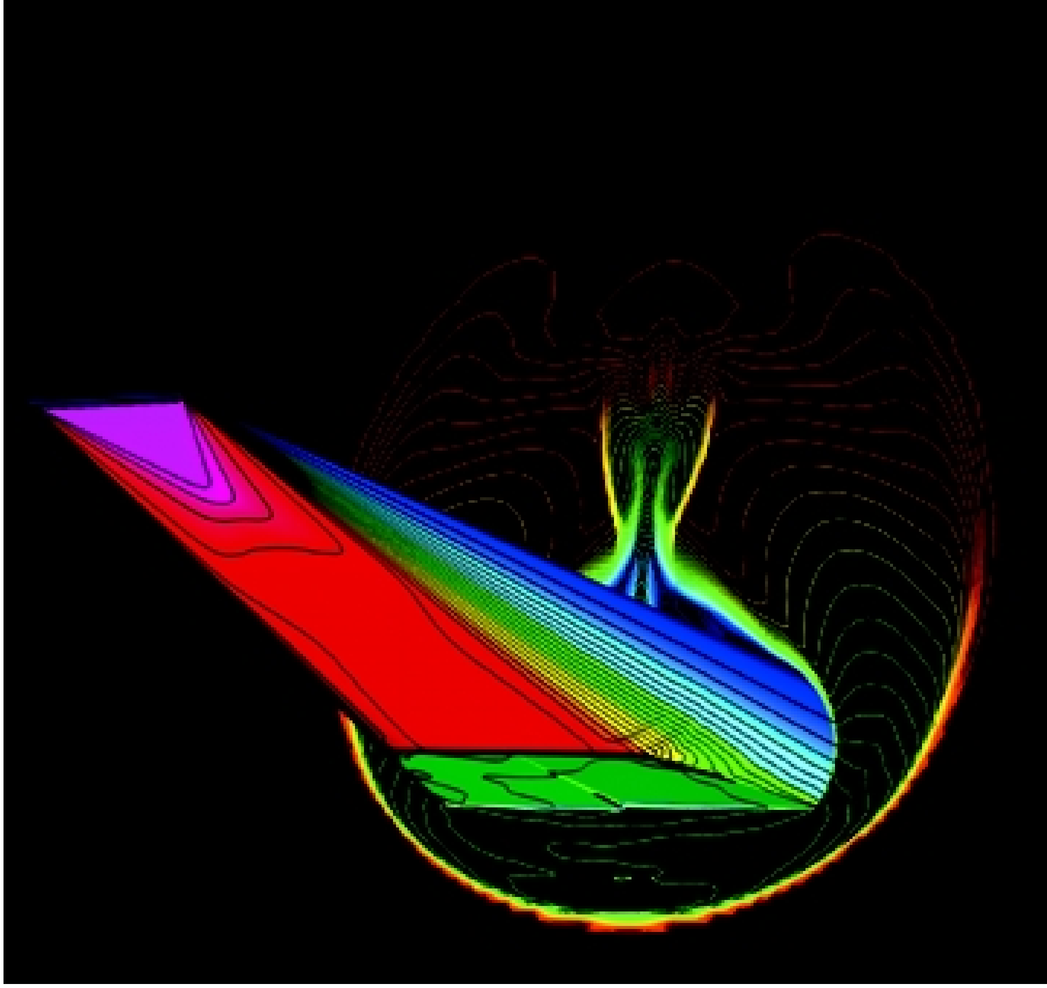

Surface pressures and Mach contours are shown surrounding the Slender Hypersonic Aerothermodynamic Research Probe vehicle at Mach 12. The data were computed on the NASA Ames Research Center’s Numerical Aerospace Simulation system in support of the vehicle’s aerodynamic and control system design.

Image by NASA

The X-37 vehicle demonstrates advanced space transportation technologies. The vehicle will support NASA’s goal of dramatically reducing the cost of getting to space.

Image by Scott Lawrence

and Applied Research in Multi-Agent Systems Laboratory, “the easier it is to achieve coordinated behavior among all the agents.” Deshmukh believes the key to success in designing future generations of RLVs will depend on ever more advanced multi-agent computing systems for processing an enormous amount of engineering data.

New engines for space

As the CAS team creates the computing tools required in designing tomorrow’s space vehicles, it also explores what will be needed to help engineers design low-cost, efficient engines to get them to and from space — yet another step along the way toward reliable space transportation.

NASA’s Numerical Propulsion System Simulation (NPSS) effort, which was begun by CAS researchers at the Glenn Research Center in the early 1990s as an engineering design and analysis tool to support aircraft engine development, is now playing a role in space vehicle engine design.

A concerted effort by the aerospace industry and academia to develop an advanced engineering environment, NPSS is actually an integrated collection of software programs for the analysis and design of engines and components. Its purpose is to dramatically reduce the time, effort and expense necessary to design and test a variety of engines.

According to researcher Karl Owen at the Glenn Research Center, NPSS generates sophisticated computer simulations of an aerospace object or system, thus permitting an engineer to examine various design options without having to conduct costly and time-consuming real-life tests. The ultimate goal of NPSS is to create a *numerical test cell* that enables engineers to design and analyze complete engine simulations on cost-effective computing platforms. Using NPSS, engine designers will analyze different parts of a rocket propulsion system simultaneously and perform different analysis simultaneously — for example, aerodynamic and structural — while performing the overall analyses faster and at lower cost.

Images by Mark Stewart

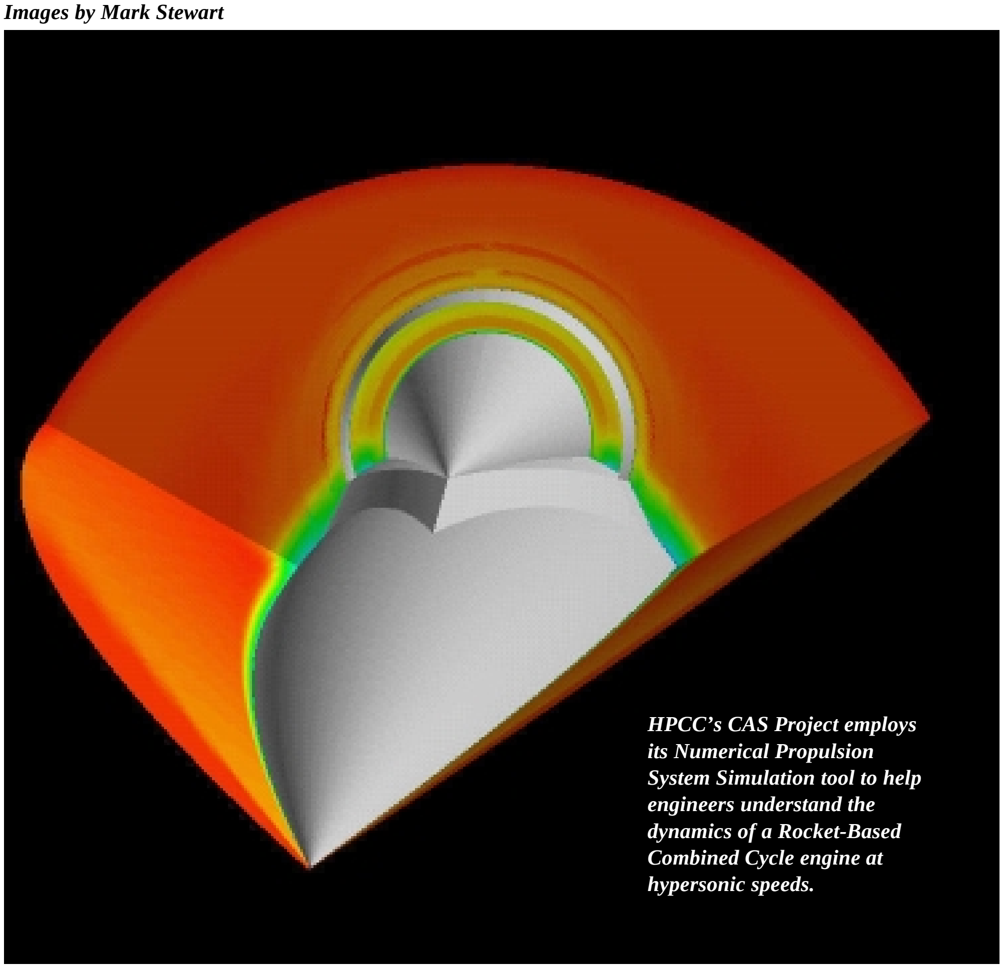

HPCC’s CAS Project employs its Numerical Propulsion System Simulation tool to help engineers understand the dynamics of a Rocket-Based Combined Cycle engine at hypersonic speeds.

Pratt & Whitney’s Mark auBuchhon recognizes the benefits of NPSS for the development of advanced jet engines.

NPSS, as Owen describes it, uses object-oriented software design as a way of organizing data and procedures into manageable packages called *objects*. One reason the object-oriented approach was chosen for NPSS is because it allows new codes to be introduced into the system quickly and easily. In other words, if an aerospace company develops a powerful new code for one rocket engine component, the object-oriented framework permits use of that code with all the other codes in NPSS even if that new code runs on a different type of computer.

According to Mark auBuchon, manager of Advanced System Simulations at Pratt & Whitney, “the NPSS engine modeling system offers great promise to industry and government by establishing a standard for the modeling methodology and execution environment for engines. This standardization," he continues, "will lower costs and ensure the successful development and transition of propulsion technology."

NPSS provides another important capability to RLV propulsion developers — zooming. As in photography, zooming or magnifying allows a rocket engine developer to analyze the performance of an engine component by zooming in on that component to evaluate its performance in great detail. This marks a major step forward because aerospace engineers will be able to perform detailed analyses of engine components within a system (the entire engine) rather than in isolation.

With a new generation of space vehicles on the horizon, CAS's efforts will reduce design and development time for new rocket engines and other key components to the national space transportation system. As with CAS's successful track record in high-performance computing applications for aviation during the 1990s, the goal is to produce similar efficiencies, capabilities and cost savings with future generations of RLVs — a vital step on the high road to space.

Photo by Judy Conlon

Mark auBuchon of Pratt and Whitney explains how the Numerical Propulsion System Simulation engine modeling system offers great promise by lowering costs and ensuring successful engine development.

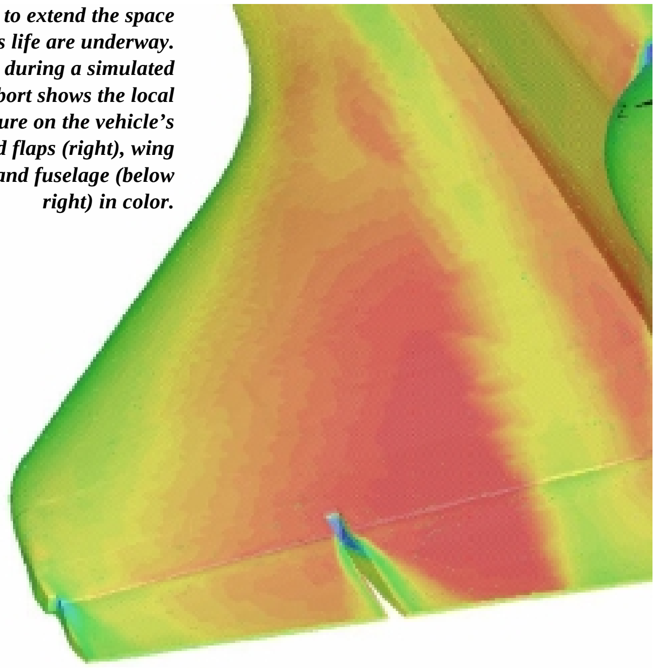

Studies to extend the space shuttle's life are underway. Analysis during a simulated launch abort shows the local pressure on the vehicle's elevons and flaps (right), wing (top), and fuselage (below right) in color.

Computations and images by John Melton

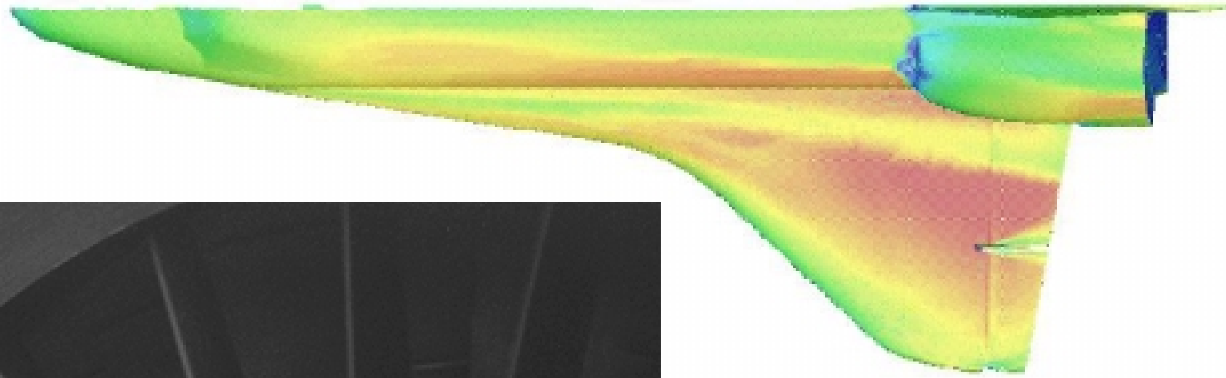

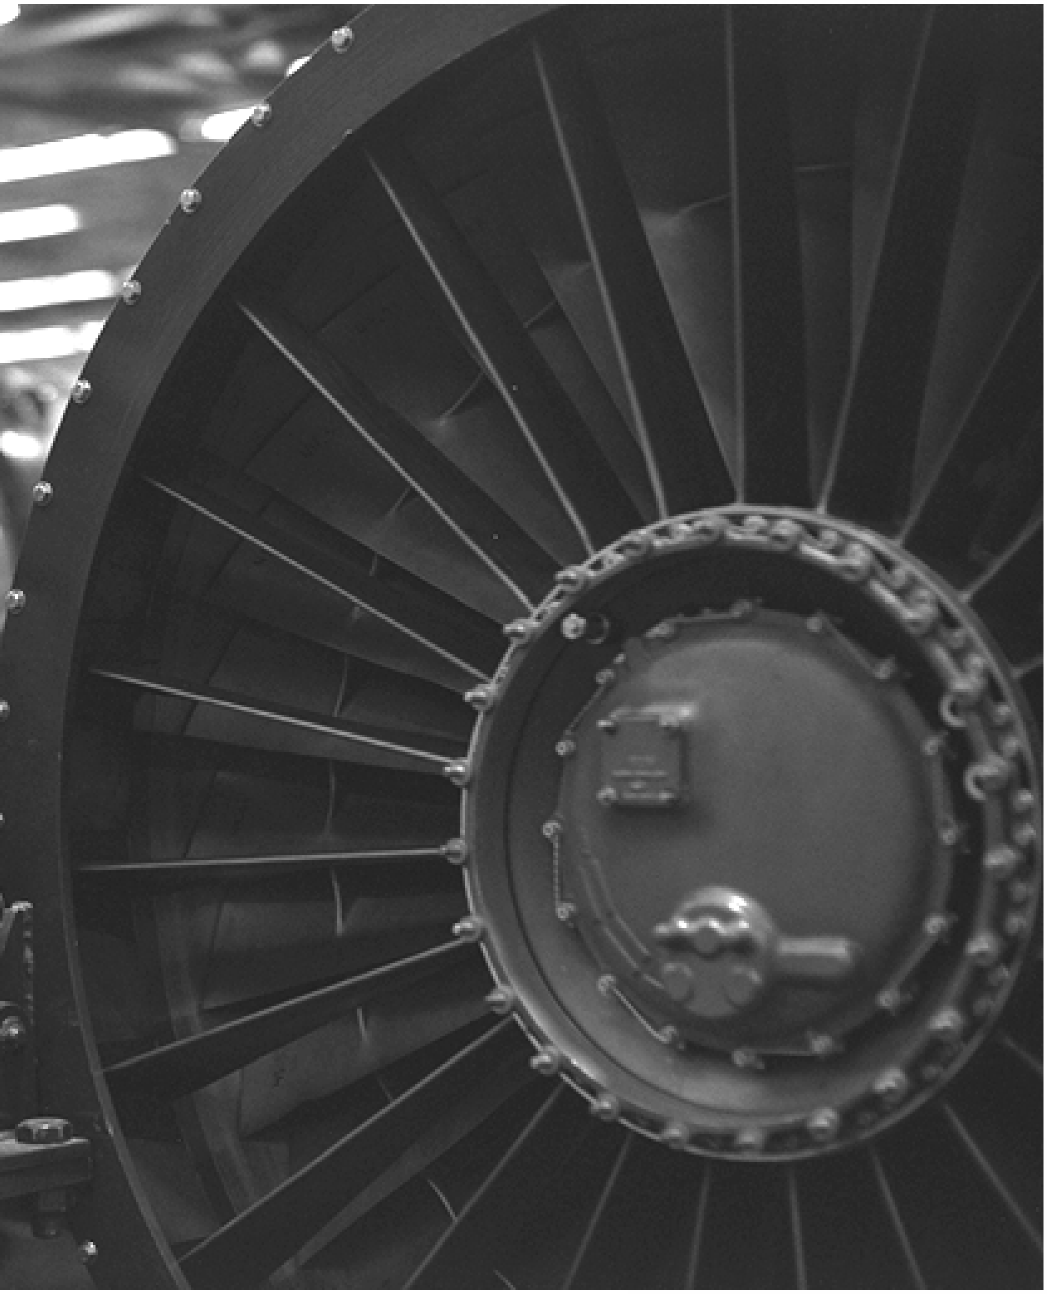

“The NPSS engine modeling system offers great promise to industry and government by establishing a standard for the modeling methodology and execution environment for engines.”

■ **Mark auBuchon**
Pratt & Whitney

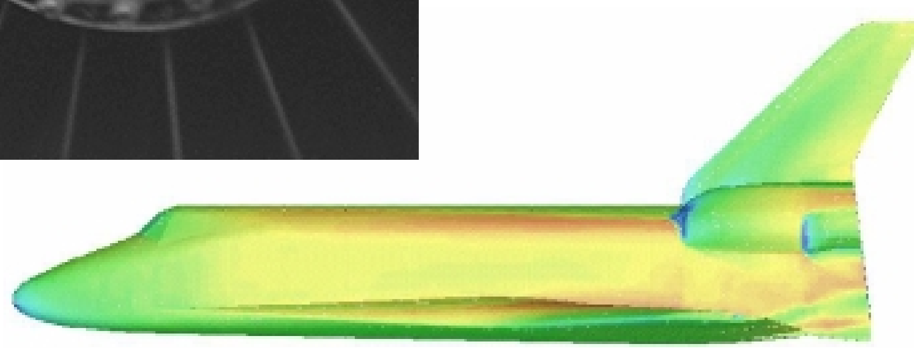

Ultra-fast gigabit networking technology will revolutionize data transmission

By Kenneth Freeman and Hugh LaMaster

Photos by Fritz Hasler

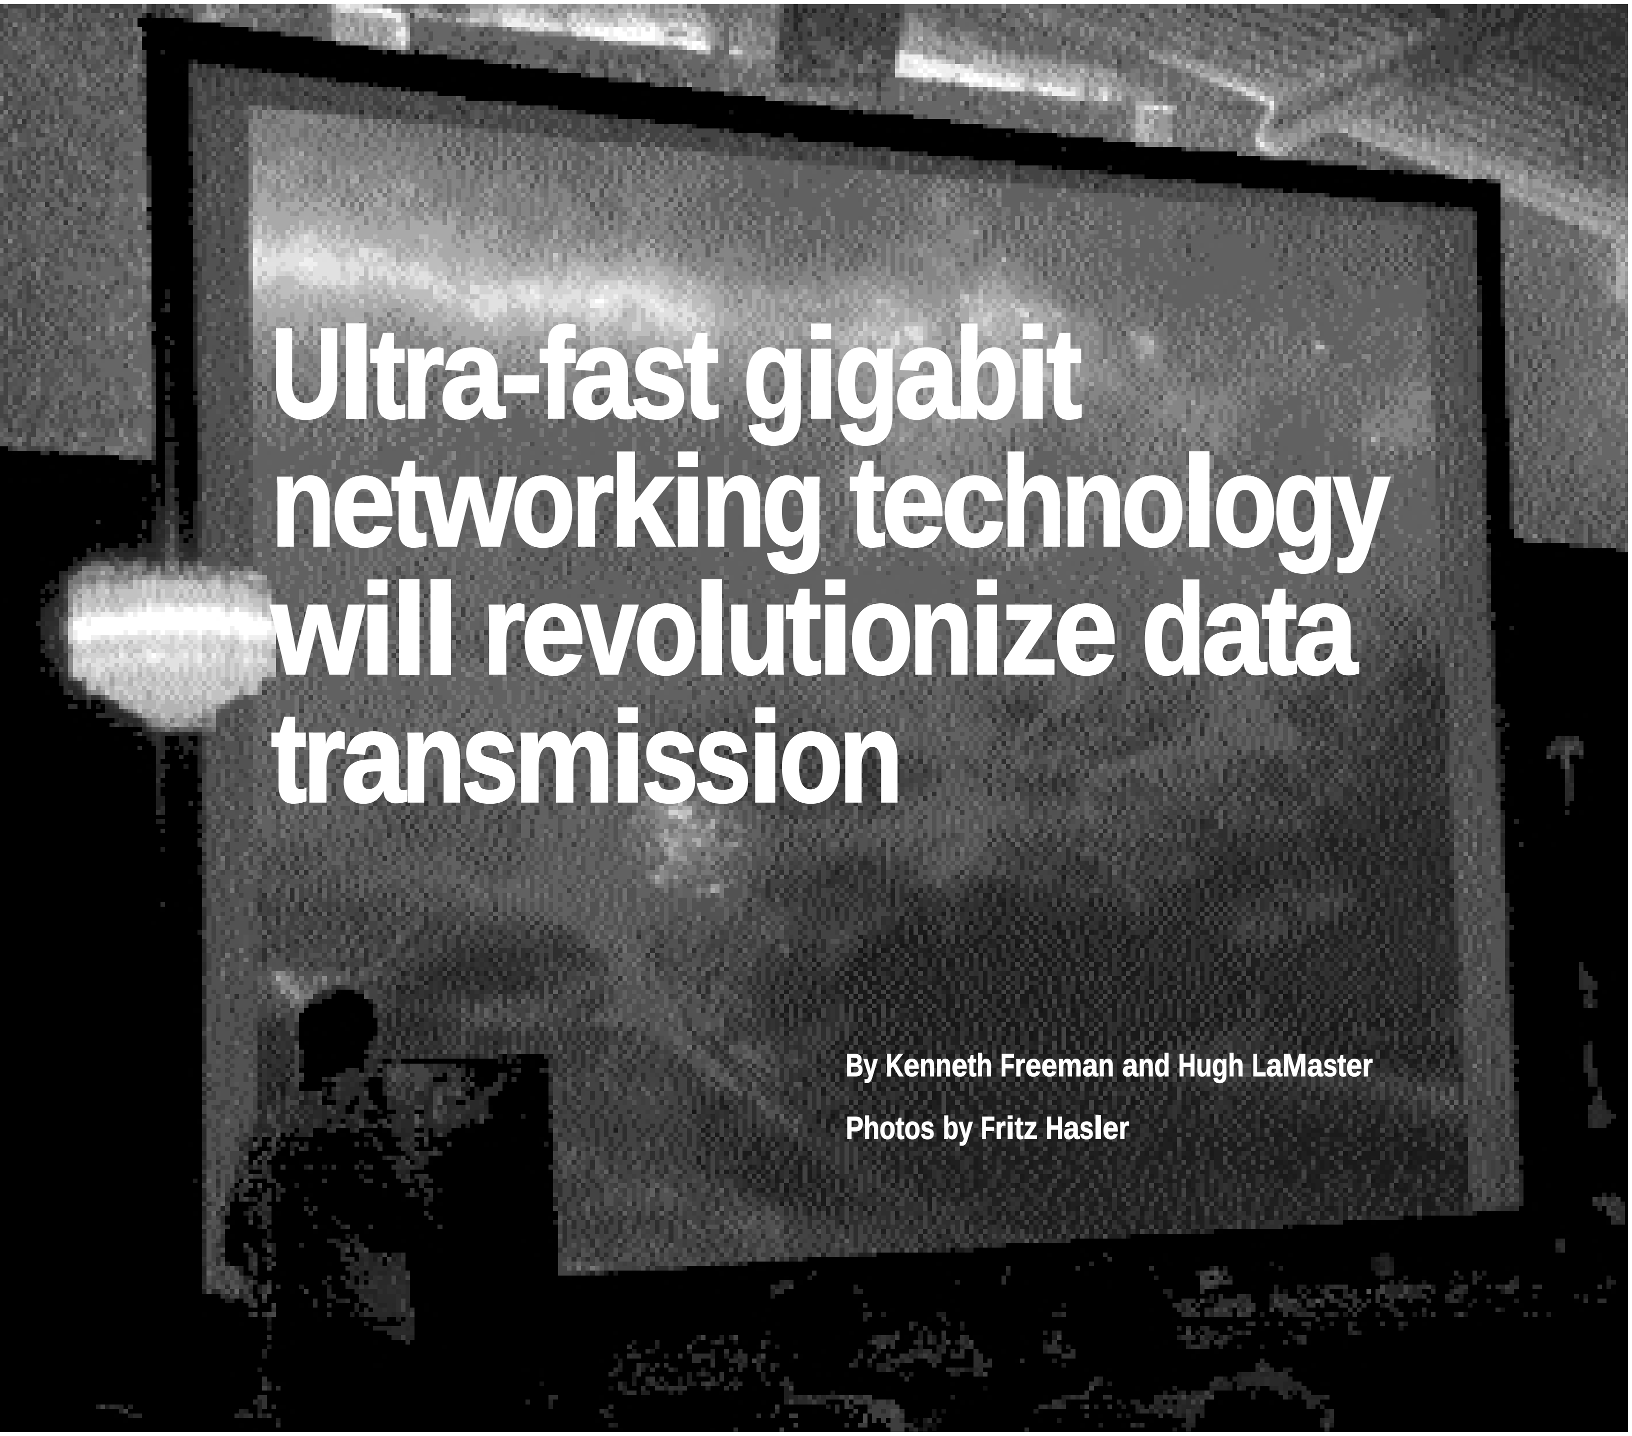

David Curkendall of the Jet Propulsion Laboratory (JPL) is silhouetted before a projected Infrared Astronomical Satellite image of the Milky Way at a recent NREN gigabit-networking workshop. Gigabit-networking technology will enable researchers to exchange astronomical data via JPL's Digital Sky Virtual Observatory.

As you read this sentence, multiple streams of data—totaling a billion bits per second—are being transmitted over strands of optical fiber, which is hair-thin plastic or glass *wire* used to shunt information using infrared or visible light.

While this kind of high-speed networking in the aggregate is achievable today—that is, the ability to transfer, at once, billions of bits of information through a single fiber-optic line across multiple data streams—the real technical challenge remains in delivering a single, high-speed data stream, from end-to-end, directly to the user's desktop.

The NASA Research and Education Network (NREN) Project is administered by the agency's High Performance Computing and Communications Program. Along with partners from other federal agencies, industry and academia, a dedicated NREN team is already demonstrating the future when ultra-fast, data-intensive, computer networking will be commonplace. An emerging technology of interest to NASA—known as *gigabit networking*— promises to deliver the high-speed goods on a number of fronts. A gigabit is one billion bits of information, and a gigabit network is the way to transmit enormous amounts of information. While maturation of this technology lagged for years, due primarily to computer network bandwidth limitations, it now has the proverbial *green light* to come of age.

Through the NREN Project, whose goal is to infuse emerging networking technologies such as gigabit networking into NASA applications, gigabit networking is fast becoming a vital part of the Next Generation Internet (NGI) initiative. NGI is a national effort among government, academia and the private sector to create the Internet of the future. As part of the initiative, the NREN Project seeks to infuse gigabit networking technologies throughout the agency. Eventually, as the technology becomes essential to NASA's operational networks and applications, it will provide NASA computer users the remarkable ability to work with each other at remote locations as if they were next door.

Gigabit networking applications within NASA will include rapid data transfer in everything from research and development and space missions to public outreach and educational activities, according to Marjory Johnson, NREN associate project manager.

"With our gigabit networking applications, NREN hopes to set the pace for future commercial use of the technology," Johnson says. As a multitude of uses for gigabit networking are being identified, the NREN team—as with other groups—is working to fine-tune the technology, as well as to demonstrate its capabilities, according to Johnson.

Tim Boyce, Celeste Banaag and David Guevara of NASA check out a high-speed networking demonstration at a recent NREN workshop that showcased gigabit-networking technology.

Moving billions of bits

What is gigabit networking? Simply put, a gigabit-per-second network speed is one billion bits of electronic information transmitted per second. Gigabit networking, then, offers the capability of moving a billion bits of data per second—many times faster than the computer networks that most people are currently familiar with. This technology opens up possibilities for exciting new applications. Distributed multimedia conferencing for graphics, voice, video, high-definition TV and virtual reality will rely on gigabit networking. More importantly, demanding NASA science and engineering applications can be deployed that use visualization, computation and storage resources at multiple network locations, achieving an unprecidented ability for researchers distributed throughout NASA to work collaboratively.

What kind of gigabit networks are deployed? One type is a local area network (LAN). The most familiar type of LAN is Ethernet, which originally ran at a speed of ten megabits per second. A megabit is one million bits. Ten-megabits-per-second Ethernet is widely deployed in offices, computer rooms and even some homes. It has been the most common method to connect PCs, printers and network servers during the past two decades. Today, gigabit networking such as Gigabit Ethernet can run one hundred times faster than the original Ethernet. Ten-Gigabit Ethernet is currently being standardized and will be 1,000 times faster than the original Ethernet.

Another type of network in which gigabit networking is being deployed is the wide area network (WAN). A WAN is a geographically dispersed telecommunications network.

Today, the highest-speed WANs run at 2.5 gigabits per second or 250 times the speed of Ethernet. Even-faster links are being tested and are in operation. These fiber-optic links run at ten-gigabits-per-second speeds, the same as the Ten-Gigabit Ethernet and 1,000 times the speed of the common ten-megabits-per-second Ethernet speed.

Advances in optical fiber, lasers and semiconductors have enabled high-bandwidth, long-distance optical technology connection speeds to exceed that of early WAN links by a factor of a million. At the same time, between 16 and 100 separate wavelengths of light traveling over a single optical fiber permits many connections to travel over the WAN link simultaneously. This amazing capability—going from one wavelength to 100—allows hundreds of gigabits per second, perhaps even one terabit (a trillion bits per second) of total traffic to travel over a single optical communications line.

Beyond exploring the best options for deploying gigabit networking, the NREN team is demonstrating that many factors besides raw network bandwidth must be examined, especially when rapidly accessing and manipulating enormous amounts of data over networks is required. For example, one area NREN explores concerns packets. A packet is the basic unit of network data. Messages are assembled in packets and marked with an address and other information for routing purposes. The size of packets, the time it takes packets to travel over a WAN, the likelihood of a packet being dropped along the way, and the total amount of data sent but not acknowledged are all factors that must be managed consistently in order for the user's application to achieve the promised speed.

Gigabit networking applications

From the study of ultra-small, sub-atomic particles to immense galactic structures, gigabit-networking technology will help NASA researchers, who rely on high-performance computing and communications technology, to probe the mysteries of the universe.

Working closely with both NASA's Numerical Aerospace Simulation (NAS) Systems Division and the Astrobiology Institute at Ames Research Center in California, the NREN Project is providing gigabit networking technology to demonstrate Virtual MechanoSynthesis (VMS), a 3-D simulator that allows users to see, move and even feel simulated molecular structures. VMS is an important tool to help scientists better understand how to design nano-electronic components, chemical and bio sensors, and nanotubes.

Gigabit networking handles the data-intensive VMS communications at high speed and enables a VMS simulation environment to run simultaneously in multiple locations so groups of scientists can work in real time.

"With VMS a user can grab an individual atom with a wand, move it about, and build complex structures," says Jon Guice, an information technology researcher at NASA's Astrobiology Institute. "VMS uses computational steering—the ability to design and modify simulations interactively."

Guice says gigabit networking is important to VMS's viability because it is the best way to deliver the enormous data required by multiple, distant users. VMS, he notes, has many potential applications in biology and medicine, and NREN's involvement in testing VMS on a high-speed network is a critical step on the road toward demonstrating its value as a research tool.

In another example of demonstrating the value of gigabit technology for NASA

Image by Michael Boswell and Christopher E. Henze

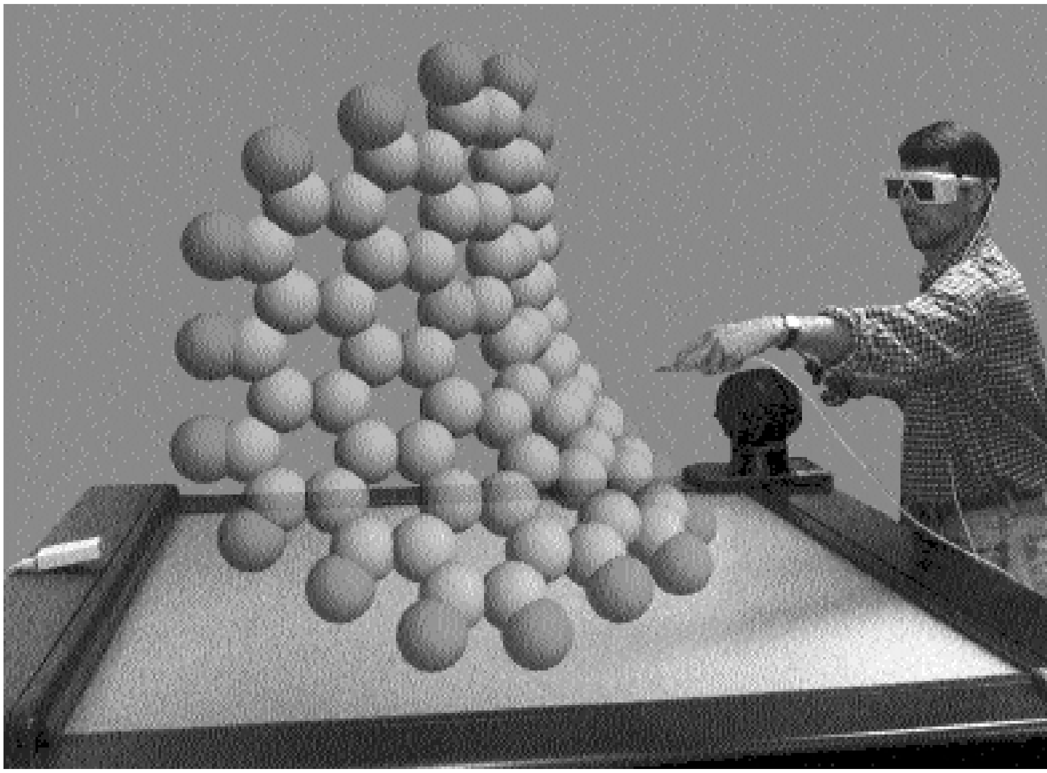

NASA's Christopher E. Henze uses a tool to manipulate a molecular model of graphite with the Virtual MechanoSynthesis (VMS) 3-D simulator. Gigabit networking, a means to deliver data to distant users, is important to VMS's viability.

operations, NREN engineers are working with scientists at the Jet Propulsion Laboratory (JPL) to deliver data and link users of the Digital Sky Virtual Observatory. Digital Sky is an astronomical data treasure trove that will enable viewing of tens of terabytes of archived space image and related information.

David Curkendall, manager of JPL's High Performance Information Technology Office, says, "Digital Sky and the future National Virtual Observatory are exciting new NASA/National Science Foundation programs knitting together geographically distributed collections of data of the entire sky at various wavelengths of the electromagnetic spectrum. If the goals of these programs are to be met, it must be possible to routinely transmit terabytes of data from the distributed archives to the points of virtual observation and analysis." That fact alone, Curkendall stresses, is what makes NREN's gigabit networking knowledge crucial to the success of Digital Sky.

Some day, astronomers combing through Digital Sky's vast holdings may discover new galaxies, stars and planets, showing that by linking basic research with the power of high-speed communication, the sky literally is the limit.

"The know-how to do gigabit networking is clearly out there, but it is still very difficult to do today," says NREN's Marjory Johnson. "Demonstrations such as Virtual MechanoSynthesis and Digital Sky are proof that gigabit networking is an effective tool for accessing data from large databases."

Challenges remain

As these new applications demonstrate the versatility of gigabit networking technology, the development of network systems continues and technical challenges remain. For example, workstation tuning and application enhancements will be necessary to take advantage of the potential increase in network bandwidth. Also, satellites such as the new Terra environmental satellite transmit an enormous amount of data to Earth on a steady basis. To distribute these large data sets across the U.S. and even the world NASA engineers must tackle the communications challenge this problem presents. Data streams from the Terra satellite are acting as a testbed for proving advanced technologies, including gigabit networking.

According to NASA's Jeff Smith, senior networks system engineer, "the Terra spacecraft generates up to one terabyte of data per day. A terabyte is a trillion bytes. This enormous flood of data can be distributed to more than 500 Terra researchers only over high-performance research and education networks such as NREN. With the scheduled launch of the Aqua spacecraft in 2001, even more data are on the way. Gigabit networking is the only possible way that NASA will be able to distribute these massive amounts of Earth-observation data."

What's ahead?

Today the NREN Project is sharing its knowledge and experience of what gigabit networking can accomplish through technology demonstrations and important engineering developments. At the same time, the project fosters a diverse community interested in utilizing this remarkable communications technology fully.

Image mosaic by Caltech's Gene Kopan

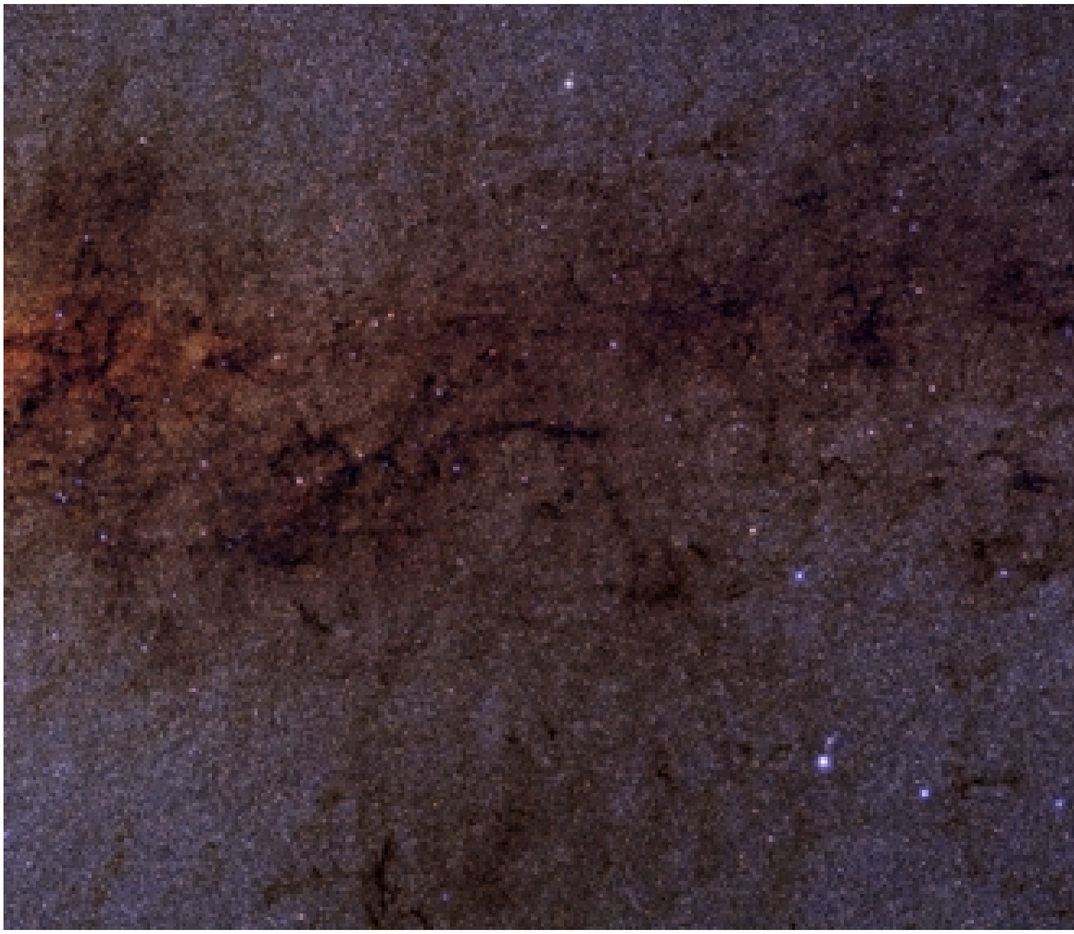

Infrared images of the Milky Way's center and the composite data of our galaxy will be archived at the Digital Sky Observatory. The astronomical database will allow viewing of tens of terabytes of space images and related information.

David Curkendall of the Jet Propulsion Laboratory sees the Digital Sky Virtual Observatory as a prime application for gigabit-networking technology.

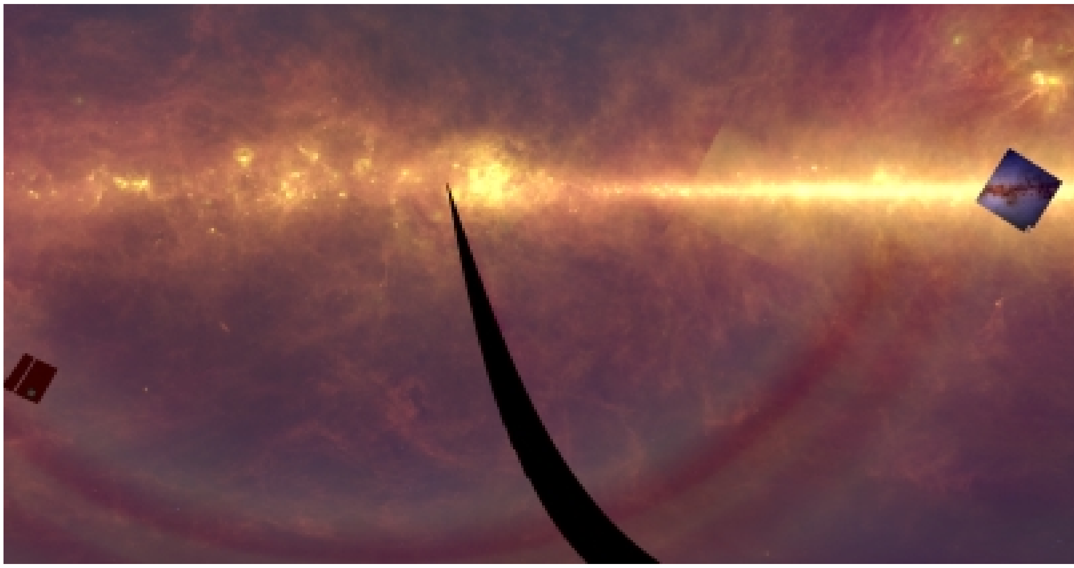

Image mosaic by Joe Jacob of JPL/Caltech

Recently, gigabit networking experts and others interested in the technology's potential from around the U.S. gathered at an NREN-sponsored workshop at NASA's Ames Research Center. Workshop participants witnessed a VMS demonstration and joined discussions of what milestones need to occur over the next two years to advance gigabit networking to its next level. Right now, issues such as connectivity, moving gigabits per second of data from a source to a user, workstation problems, and how to better use computer architectures, operating systems and network interface cards in a gigabit world commanded center stage during the discussion.

The experts agree that without gigabit networking technology to handle an enormous flood of information, the research community—the scientists and engineers sharing new ideas and data over these networks—will suffer.

While delivering gigabit capabilities to the desktop was once considered nearly impossible, today it can be accomplished when the user taps the right kind of networking expertise. Advancing gigabit technology in support of NASA missions, the NREN Project team is already demonstrating that the world of super-fast networking will soon be routine.

Image by MITI, ERSDAC, JAROS and U.S./Japan ASTER

A false-color image of San Francisco Bay from NASA's Terra spacecraft shows visible and near-infrared light to reveal sediment in the bay, vegetation, and urban environment details. Gigabit networking will assist scientists in accessing and exchanging large satellite datasets.

Networking experts gathered for exciting demonstrations of gigabit-networking technology at a recent NASA workshop.