access

Alliance/NCSA
Vol. 14 No. 3
Fall 2001

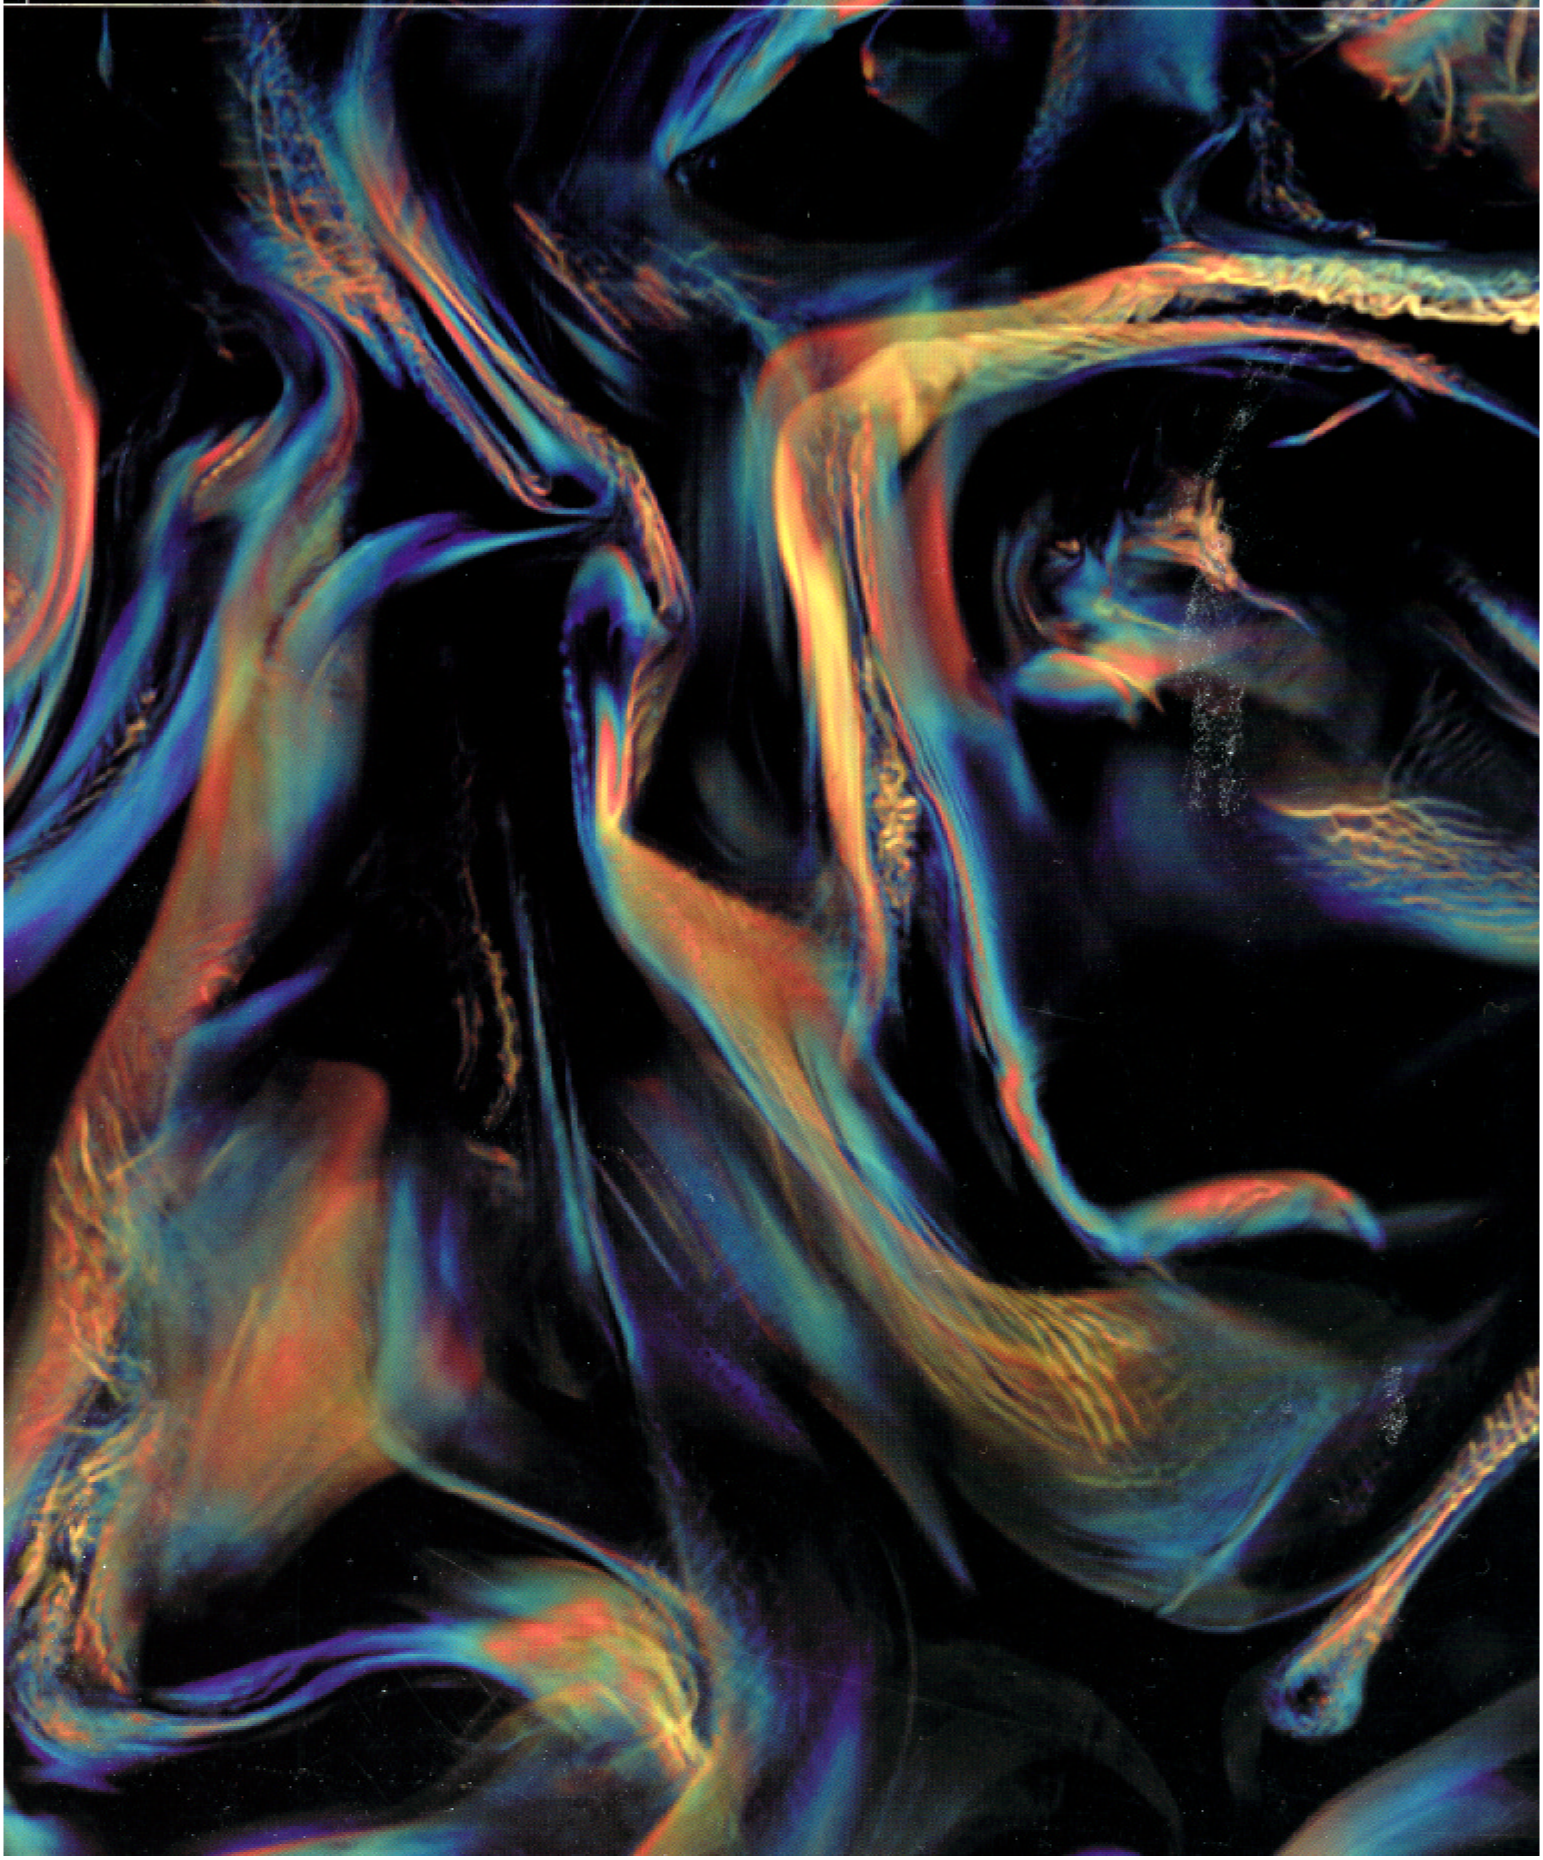

ALLIANCE
NCSA

access

Vol. 14 No. 3
Fall 2001

Executive Editor
John Melchi
jmelchi@ncsa.uiuc.edu

Managing Editor
J. William Bell
jbell@ncsa.uiuc.edu

Art Director
Carlton Bruett

Illustrators
Carlton Bruett
Luke Wroblewski

Copyeditors
Stephanie Drake
Karen Green

On the Web

http://access.ncsa.uiuc.edu
Access is published on the Web every two weeks, covering the latest developments from the Alliance and NCSA.

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
152 Computing Applications Building
605 E. Springfield Ave.
Champaign, IL 61820-5518
217-244-0072

ISSN 1064-9409

Printed with soybean ink on recycled paper.

Who we are

The National Computational Science Alliance (Alliance) is a partnership among more than 50 academic, government, and industrial organizations from across the United States to prototype an advanced computational infrastructure for the 21st century. This model infrastructure, called the Grid, will link together advanced supercomputers, visualization environments, and mass storage devices into a powerful, flexible problem-solving environment. This computing environment will be accessed via high-speed networks from anywhere in the country—eventually, the world.

The Alliance is led by the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign with major support from the National Science Foundation's Partnerships for Advanced Computational Infrastructure program. Additional funding for NCSA comes from the state of Illinois, the University of Illinois, industrial partners, and other federal agencies.

Cover

Supersonic turbulence in the outer envelope of a red giant star. Using NCSA's 64-processor prototype Itanium cluster, this 1-billion-cell turbulent flow simulation was created by Paul Woodward and his colleagues at the University of Minnesota's Laboratory for Computational Science and Engineering. The team will soon take advantage of Titan, the 1-teraflop Itanium Linux cluster under construction at NCSA, to develop a similar 8-billion-cell simulation.

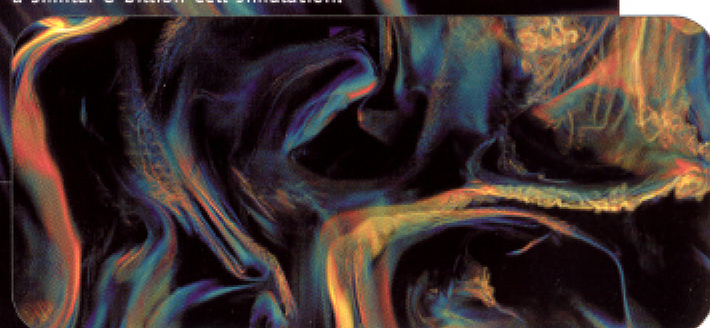

access

Contents

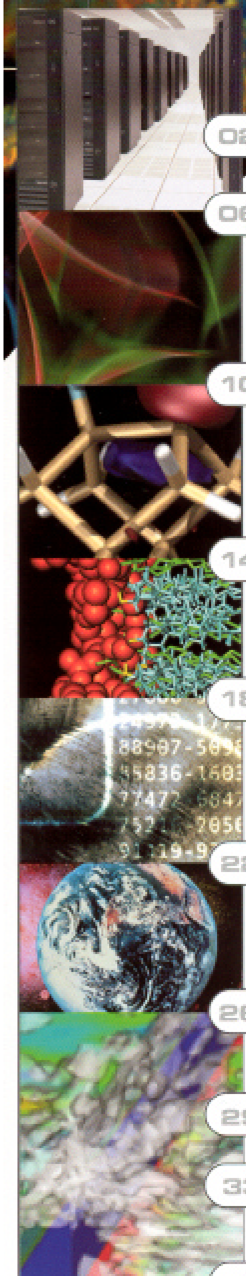

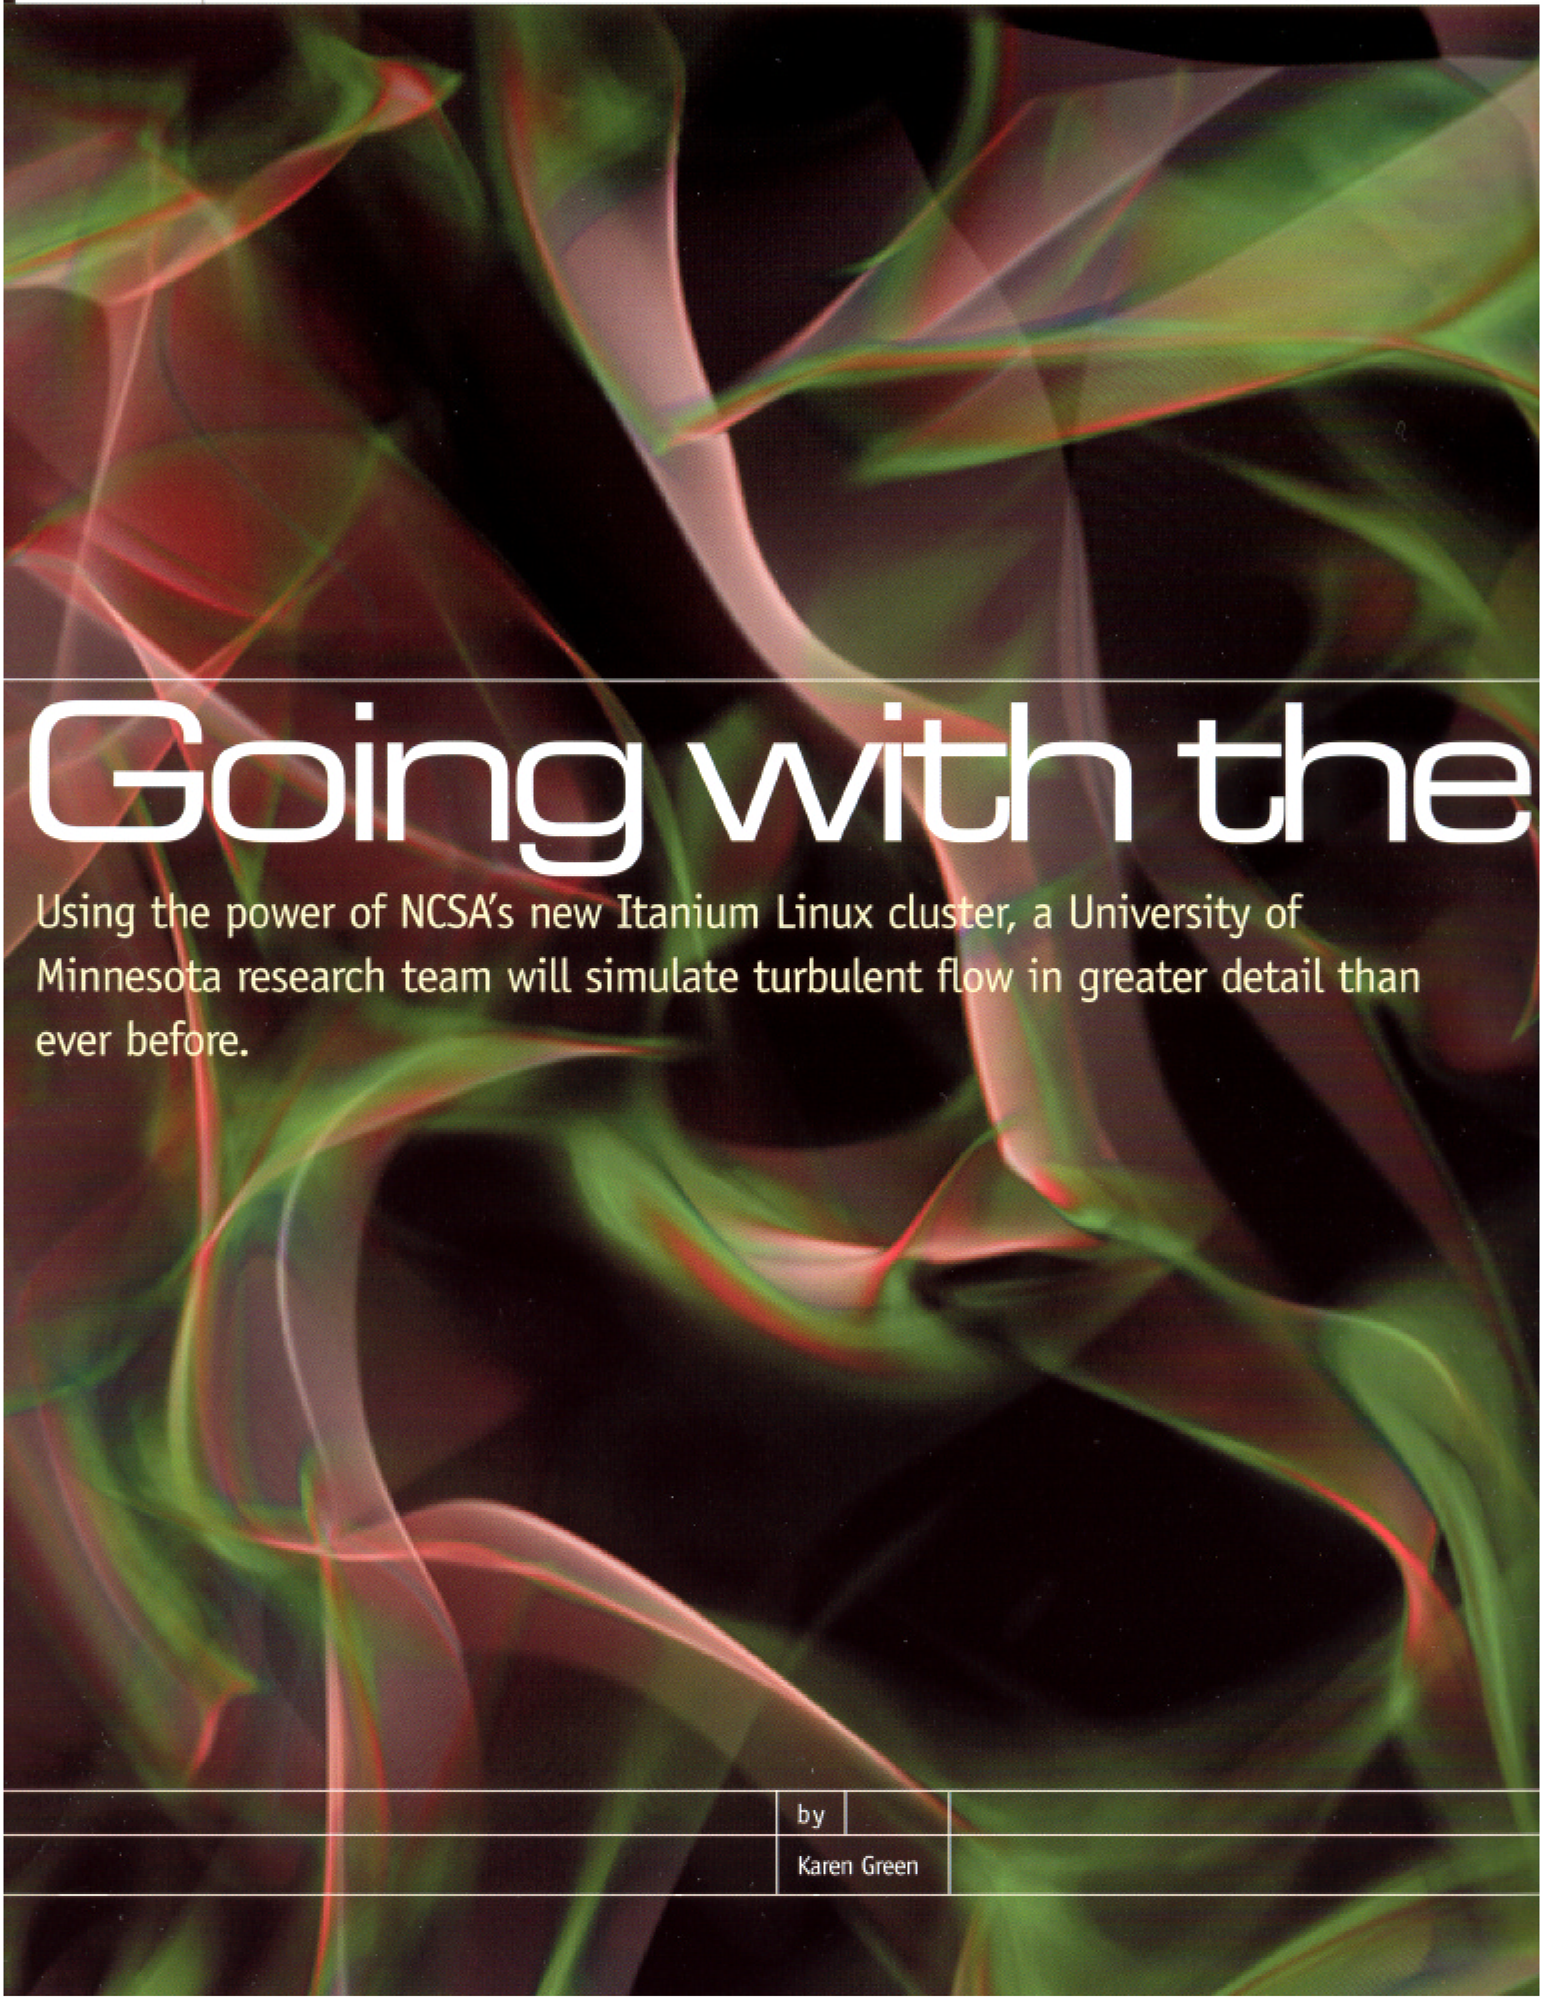

Going with the

Using the power of NCSA's new Itanium Linux cluster, a University of Minnesota research team will simulate turbulent flow in greater detail than ever before.

by

Karen Green

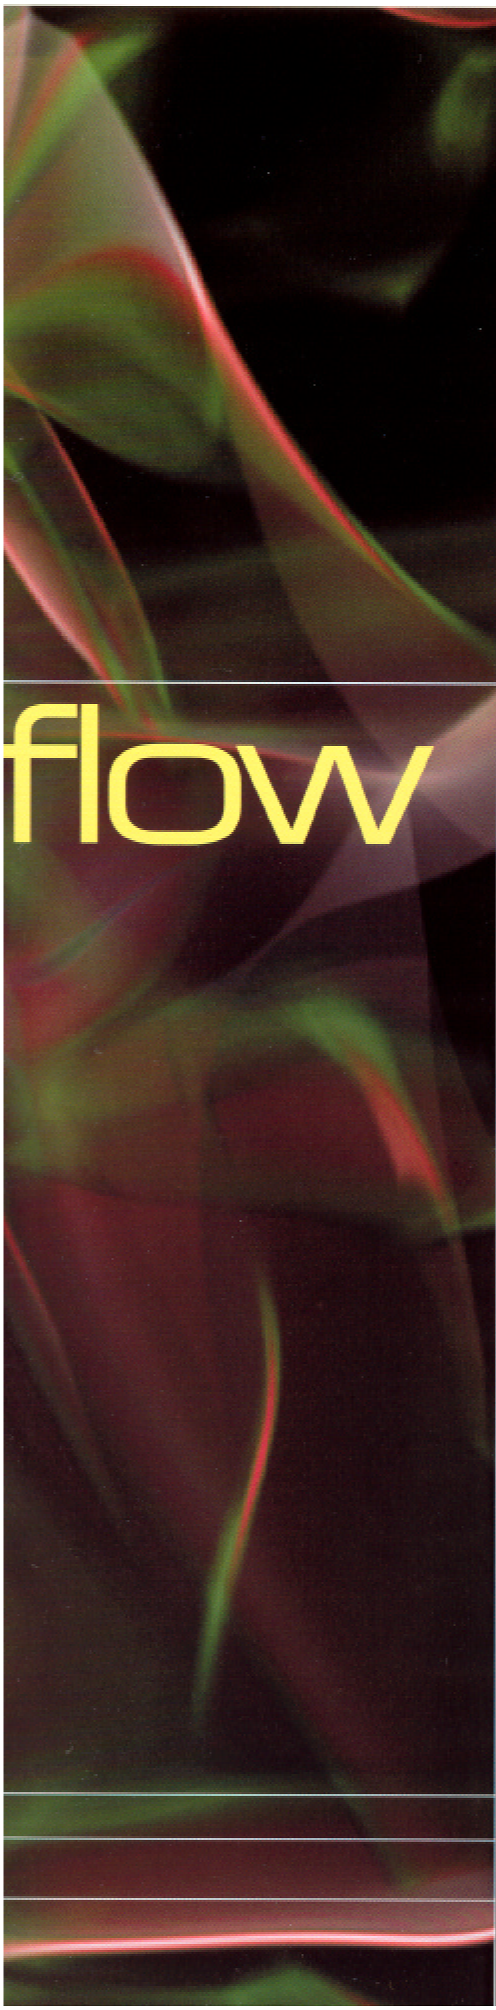

Anyone who has traveled on an airplane is aware of turbulence. This atmospheric phenomenon, which makes your tray table jiggle and leads the captain to turn on the "fasten seat belts" sign, is caused by instabilities of shear flow that result from atmospheric convection. This turbulence, although annoying and sometimes unnerving, pales in comparison to the supersonic turbulence found in the outer envelopes of red giant stars.

Simulations of convection in such stars by Paul Woodward's team at the University of Minnesota's Laboratory for Computational Science and Engineering (LCSE) show that heat rising from the interior of a star stirs up the outer envelope of the star and can result in gas motions that exceed the speed of sound near the stellar surface. These turbulent convective motions give rise to shocks—sudden compressions of the gas—of Mach 2 or more.

Turbulent motions in detail

To improve our understanding of such violent turbulent flows, Woodward's team is using the new Itanium Linux cluster at NCSA to simulate turbulent motions in detail. Violent turbulence that causes compression in gasses is less common than the turbulence airplanes encounter or the turbulence in the wake of a boat, which involve much slower relative velocities. But research so far suggests that the three are very similar and that understanding one will help in understanding the others. In the long run, a better understanding of turbulence will help in everything from analyzing weather patterns to designing airplanes and boats.

"Fluid turbulence is important in many areas of science and engineering," says Woodward. "It is a factor in the design of aircraft engines, boats, and even cars. It also influences the behavior of rivers, oceans, and the atmosphere, so this work should have an impact on a wide variety of disciplines."

Woodward's team concentrates on a phenomenon that is still not fully understood: what happens between the point at which stirring causes motion on a large scale (the big eddies right behind a boat, for example, or the large convection cells that cause cumulous clouds) and the point at which the resulting small-scale turbulent motion dissipates as heat. Figuring out what happens on the middle scales of turbulent flow is important because these turbulent motions are believed to strongly influence large-scale fluid flow. The steadily increasing power of supercomputing systems is just beginning to make possible these simulations of turbulence on smaller scales. Woodward and his colleagues were quick to jump at the opportunity to develop high-resolution simulations that could follow a turbulent flow from the macro level—where energy causes motion—down to the levels that lead to dissipation as heat.

An experimental approach

Woodward and his colleagues are no strangers to large, computationally intensive simulations. For years this research team has studied fluid dynamics in red giant stars in an effort to better understand stellar convection as well as the pulsation and ultimate mass ejection of red giants. The team's efforts to simulate the broad range of scales in turbulent flow date back a decade or more.

The largest of these earlier runs, a simulation of turbulence induced by a shock wave passing over an interface between two fluids of different density, won the 1999 Gordon Bell Award in the performance category. This simulation involved collaborators at Lawrence Livermore National Laboratory and IBM. Woodward's LCSE team most recently ran a 1-billion-cell simulation on a prototype 64-node Itanium Linux cluster at NCSA. Because of the configuration of the Itanium cluster and the fast network connection to the team's Minnesota lab, this newest run generated a rich dataset that documents the complete time history of the simulation. The team will use these data to validate much more conclusively ideas for turbulence modeling suggested by their earlier work, particularly the large simulation done at Livermore.

"Our approach to the study of turbulence is experimental," says Woodward. "We are trying to use the Itanium cluster to generate an extremely detailed set of data describing a turbulent flow, with the density, pressure, and three components of velocity sampled on a sequence of hundreds of times at each of a billion mesh points." Such mind-boggling kinds of data can be compared to existing turbulence theories or even used to develop new theories, he adds.

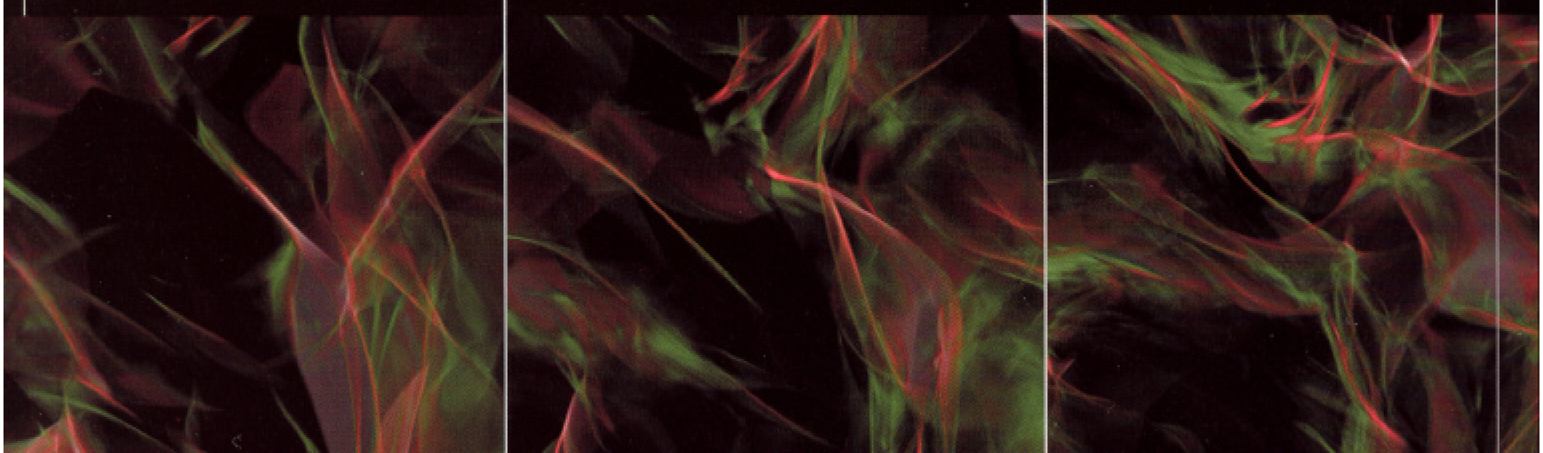

Three intervals of gas compression from a 1-billion-cell turbulent flow simulation by Paul Woodward and researchers at the University of Minnesota's Laboratory for Computational Science and Engineering. Red to white streaks indicate strong compression of the gas, while green to blue streaks show weaker compression.

The key to using data from computer simulations as if they were experimental data is to be very certain of the data's accuracy. Just as experiments involve errors in measuring data values, computer simulations involve numerical errors—failures of the computational model to accurately follow the behavior of a real gas. To simulate turbulent flow, Woodward and LCSE scientist David Porter use a method called piecewise parabolic method (PPM). In interpreting their data, Woodward and Porter carefully filter out the smallest-scale motions, where viscosity is an influence and numerical errors are most likely to occur. Since they are interested in the larger-scale motions where viscosity is not a factor, this filtering process removes errors without sacrificing simulation detail. To achieve the level of detail they require, they need a very fine grid that will result in a high-resolution simulation. For this reason, they are planning an 8-billion-cell turbulence simulation that will run on Titan, NCSA's brand-new 1-teraflop Itanium Linux cluster. With the data from the 8-billion-cell simulation, the researchers will have the data they need to make comparisons to their 1998 Livermore run.

The new simulation will probably require all of the cluster's 320 processors. A full Navier-Stokes simulation—a process that would follow the flow from its start to its dissipation as heat—would require a grid of 340 billion cells and 150 teraflops of computing power. Overall the LCSE team expects to generate 40 terabytes of data on turbulent flow in the most complete detail possible on today's computing systems.

"We view this data as much more than a bunch of numbers saved to a disk, because high-quality data like this has many uses," he adds. "If you have high-quality data, you can test theories and be confident of the results, and we believe these simulations will give us this special kind of data."

This research is supported by the National Science Foundation, the Department of Energy Office of Science, and the University of Minnesota's Minnesota Supercomputing Institute.

Access Online http://access.ncsa.uiuc.edu/CoverStories/itaniumflow/

For further information:
http://www.lcse.umn.edu

Team members
Benjamin Allen
Sarah Anderson
Kevin Edgar
Michael Jacobs
David Porter
Igor Sytine
Paul Woodward

Vorticity at three intervals in a turbulent flow. Blue and cyan indicate low vorticity, orange shows mid-level vorticity, and yellow indicates high vorticity. Over time the turbulent flow shows more vorticity.