LDRD

LABORATORY DIRECTED RESEARCH AND DEVELOPMENT

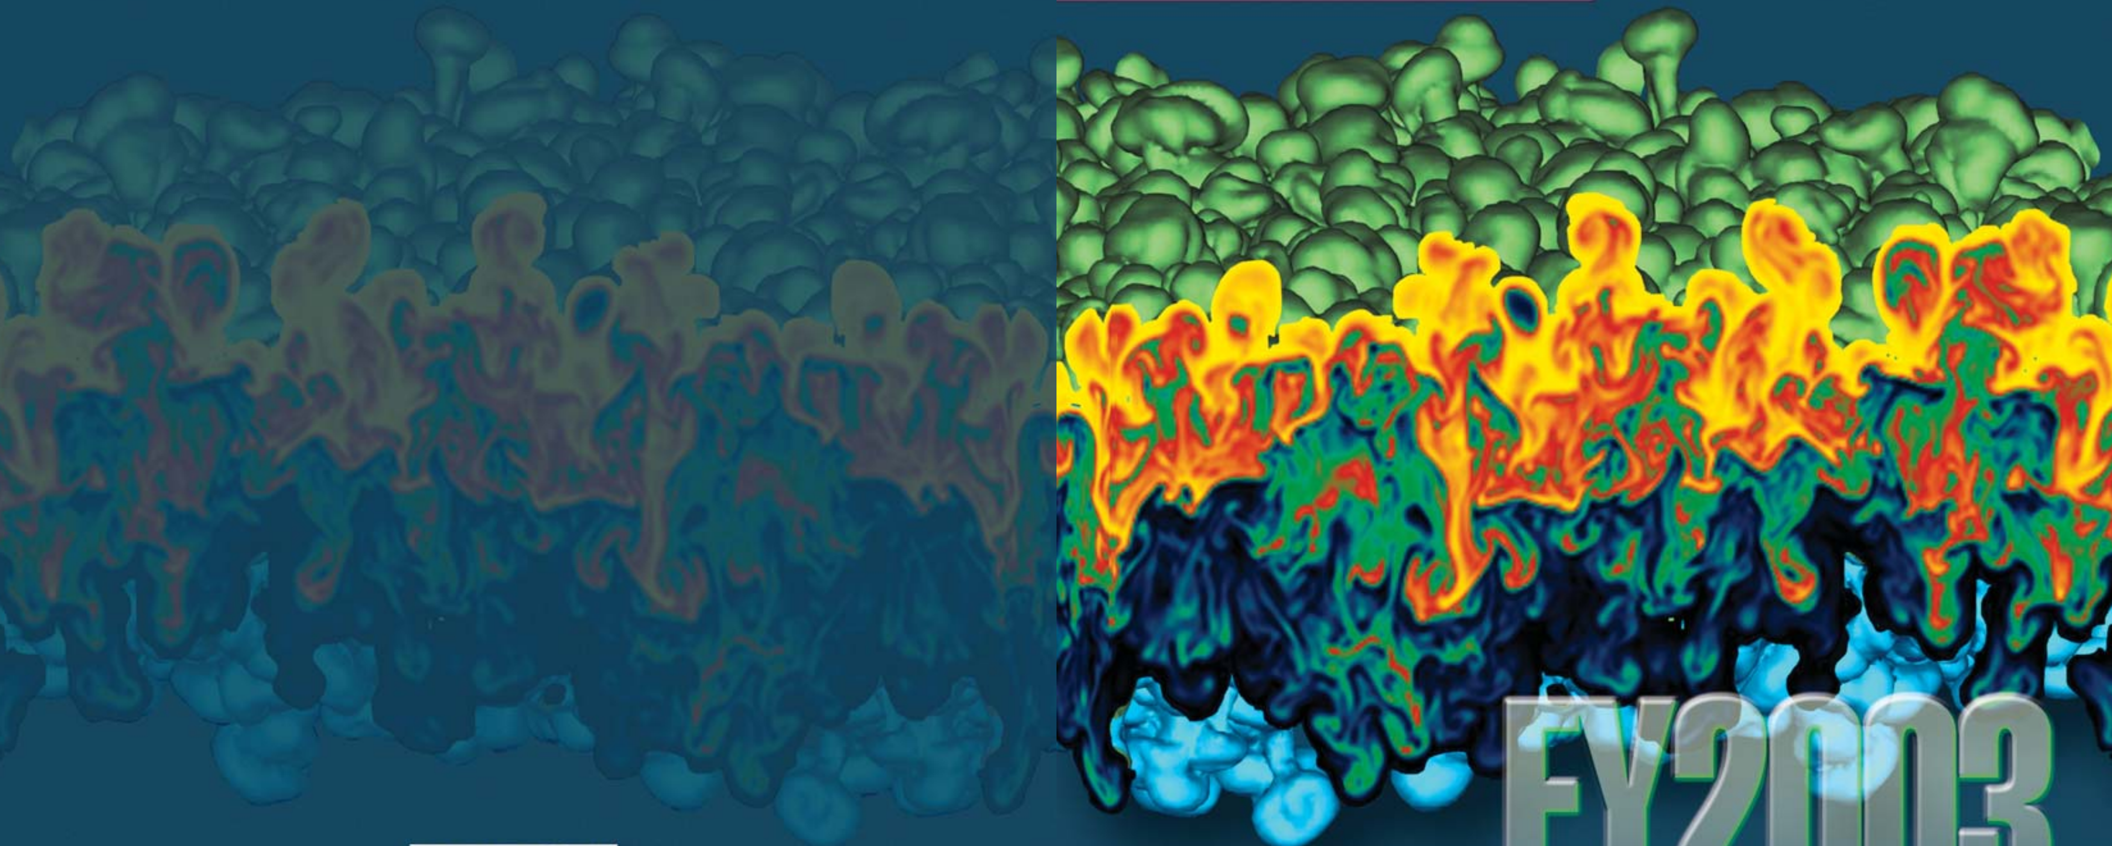

FY2003

ANNUAL REPORT

NNSA
National Nuclear Security Administration

LAWRENCE LIVERMORE NATIONAL LABORATORY

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
http://apollo.osti.gov/bridge/

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
http://www.ntis.gov/

or

Lawrence Livermore National Laboratory
Technical Information Department Digital Library
http://www.llnl.gov/tid/Library.html

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Acknowledgments

This Annual Report provides an overview of the FY2003 Laboratory Directed Research and Development (LDRD) Program at Lawrence Livermore National Laboratory (LLNL) and presents a summary of the results achieved by each LDRD project. At LLNL, Laboratory Director Michael Anastasio and Deputy Director for Science and Technology Harold Graboske are responsible for the LDRD Program and delegate responsibility for the operation of the Program to the Associate Deputy Director for Science and Technology and the Director of the Laboratory Science and Technology Office (LSTO), Rokaya Al-Ayat. The LDRD Program at LLNL is in compliance with Department of Energy (DOE) Order 413.2 and other relevant DOE orders and guidelines.

The LDRD Program extends its sincere appreciation to the principal investigators of the FY2003 projects for providing the content of the Annual Report and to the publications team. A special thanks goes to Adam Schwartz for his generous assistance. The Program also thanks the following members of the LSTO team for their many contributions to this publication: Mary Callesen, administrator; Nancy Campos, database manager; Andrew Hurst, computer specialist; and Cathleen Sayre, resource manager.

Scientific Editors
Kenneth Jackson, Karl van Bibber

Publication Editors
Karen Kline, Paul Kotta

Publication Designer
Daniel Moore

Compositor
Louisa Cardoza

Visualizing ASC-Scale Simulations on a Laptop

ViSUS, a scalable, streaming visualization technology, frees scientists from large-scale systems

Science-based stockpile stewardship, a core mission of the Laboratory, is based on computer simulations of the complex physics involved in nuclear weapons. The supercomputers on which such simulations are run process massive amounts of data, often many terabytes—trillions of bytes. Such simulations eliminate the need for actual weapons testing.

Once a simulation has been finished—which can take weeks or months even on the fastest supercomputers—scientists need to interpret the results. One way to do this is by visualization, which converts the raw data produced into a three-dimensional (3-D) visual format. Because the massive amounts of information involved would overwhelm the capabilities of conventional office computers, such simulations are usually visualized either on large display systems with massive processing power, or by using high-end systems to generate off-line animations. However, the restrictions and limitations of the typical visualization system are many: they support only certain types of simulations or certain simulation systems, run only on high-end systems, and allow only limited modification of the parameters in the interactions. In many cases, the output is a static movie that cannot be changed at all. To test other combinations of variables, or view a different cross section of the simulation, the researcher must have additional movies made by trained personnel.

When a large display system is used, the dataset to visualize must be downloaded in its entirety to the visualization system. This can take hours or even days depending on the size of the dataset, which is a function of the simulation's complexity. Multiply such wait time and heavy use of network bandwidth by thousands of scientists and thousands of simulations across the DOE complex, and the need becomes clear for the technology to use existing resources more efficiently.

Limited access to these conventional visualization systems creates a frustrating slowdown in the overall process of scientific discovery. Scientists must wait as animation files are generated. In the case of a high-performance display, they must "wait in line" for access and must travel to wherever the system is located to use it. This approach precludes testing a sudden flash of inspiration on an ordinary office computer or with a laptop from an offsite location.

TOC

Enter ViSUS

An LDRD project is putting sophisticated visualization capabilities in the hands of the researchers themselves. "ViSUS: Visualization Streams for Ultimate Scalability" (02-ERI-003), led by computer scientist Valerio Pascucci, has developed a tool that lets researchers visualize massive simulation datasets remotely and adjust simulation variables on the fly to explore 3-D simulations interactively. This powerful, flexible tool can run on an ordinary desktop or laptop computer.

To change a variable or the perspective during visualization, the user simply moves a lever on the program's graphic user interface, as shown in Figure 1. In this visualization of trace gases in the atmosphere made with the IMPACT global atmospheric chemistry model, users can visualize different concentrations of ozone (O_3), hydrogen peroxide (H_2O_2), and bromine monoxide (BrO). To test the effect of different combinations of variables, the user makes changes and sees those changes instantly reflected in the images. Testing different scenarios—or the iterative process of finding the right set of variables to achieve a certain result—becomes much faster than it would be by generating large numbers of static movies or having to run new simulations on a supercomputer.

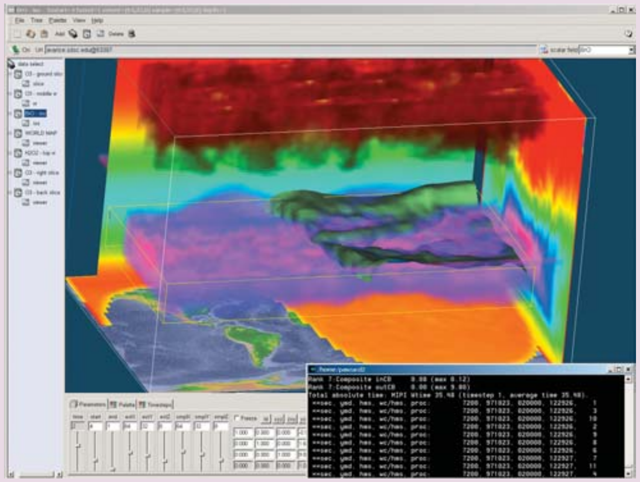

Figure 1. A screen capture of the ViSUS system being used to visualize an atmospheric simulation done with the IMPACT atmospheric chemistry model. Unlike other visualization systems, ViSUS allows the user to change a variable on the fly by simply moving a lever on the program's graphic user interface. In this case, ViSUS can test the effects of different concentrations of ozone (O_3), hydrogen peroxide (H_2O_2), and bromine monoxide (BrO). With ViSUS, the process of testing different scenarios becomes much faster than it would by generating large numbers of static movies.

Another important advantage of ViSUS is that it is a streaming solution that visualizes data in real time—as soon as retrieval of the data begins. This streaming concept is similar to the streaming audio and video now widely used on the Internet: instead of downloading a huge file, a data-streaming application plays the file as the data are received. Rather than requiring massive bandwidth between sender and receiver to download huge files quickly, all that is required is enough bandwidth for transmission to keep up with the rate at which the data are played.

Of course, ViSUS is much more complex than the Internet analogy suggests. The breakthrough technology at the heart of ViSUS is a set of algorithms that selectively choose the data that must be uploaded first in order for the images to be immediately useful. The algorithms first extract the data needed to produce coarse images that can later be "fleshed out" into higher-resolution images. Depending on the memory and processor of the computer on which ViSUS is running, the higher-resolution images may be viewable almost immediately. But even if memory, processing power, or bandwidth is limited, ViSUS can enable the real-time imaging of massive, complex scientific datasets.

"Ultimate Scalability" Means Ultimate Flexibility

This is the "ultimate scalability" in the name ViSUS: the same version of the visualization engine can run on anything from a low-end laptop to a high-end workstation. In the latter case, a high-resolution version can be viewed without sacrificing time. But the flexibility to trade speed for resolution allows a scientist to use an ordinary desktop or laptop computer to visualize data remotely and in real time. A scientist can run the visualization first at low resolution to get the big picture and then, if something intriguing appears, view it again at high resolution. Because the local algorithms scale down to low-end computers, the ViSUS program is installed on the user's own computer, rather than on a server—a requirement for fast visualization.

For comparison, Figure 2 shows a simulation of Richtmyer–Meshkov fluid mixing—a phenomenon that arises when a shock passes through the boundary of a heavy fluid sitting atop a light fluid—being visualized with ViSUS both on a large, multi-projector, very-high-resolution display (a "PowerWall") and on a laptop.

With its data-streaming and interactive capabilities, ViSUS could even be used to adjust the parameters of a supercomputer simulation in progress, provided that that the data source is set up for such interaction. Unsuccessful simulations could therefore be terminated midway, a boon for conserving supercomputer resources. If a simulation would otherwise take hours or even days to finish—during which time the computational resources would be unavailable for other work—then in the long run, ViSUS could save both supercomputer time and the time of the people running the simulation and waiting for its results.

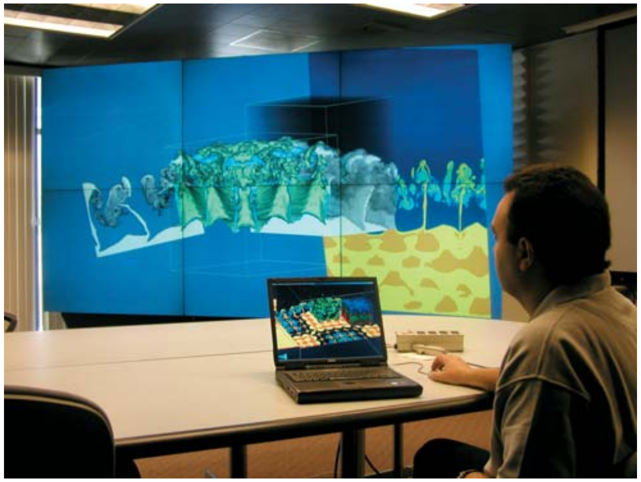

Figure 2. The same simulation—Richtmyer–Meshkov fluid mixing—being visualized on a large, multi-projector, very-high-resolution display and with ViSUS running on an ordinary laptop computer. Richtmyer–Meshkov instability is a phenomenon that arises when a shock passes through the boundary of a heavy fluid sitting atop a light fluid. This particular simulation won the 1999 Gordon Bell Award for best supercomputing performance.

The ability to terminate simulations before completion is particularly advantageous in climate simulation. Climate simulations are incredibly complex in that they involve numerous variables—such as wind speed, carbon dioxide content, temperature, and humidity—all interacting with each other; a change in any variable causes a chain reaction that affects all the others. Modeling climate events therefore entails some of the largest datasets in science. (For instance, the Earth Simulator, the Japanese supercomputer currently ranked as the world's fastest, was developed specifically for climate simulations.)

Figure 3 is a frame from a ViSUS visualization of Rayleigh–Taylor instability—another fluid-mixing phenomenon relevant to stockpile stewardship. The original dataset, over 3 terabytes, took over 2 months to generate using the Miranda code on a nearly 2000-processor supercomputer at LLNL that ranks among the world's fastest. With conventional visualization technology, scientists did not know the results until after the simulation was completed. ViSUS, however, would allow scientists to visualize such simulations while still in progress; if the simulation was not progressing as desired, it could be stopped or the variables could be adjusted.

This on-the-fly flexibility also permits researchers to explore large datasets to a groundbreaking degree. An ordinary visualization approach allows some changing of variables but is essentially a static movie. This approach may utilize only one-tenth of a percent of the data generated in the simulation. ViSUS, on the other hand, allows researchers more flexibility to change the variables of the experiment so that a much larger part of the dataset can be utilized.

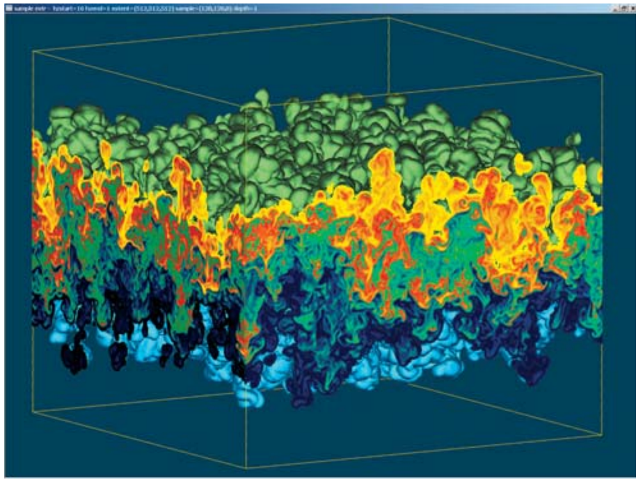

Figure 3. A frame from a ViSUS visualization of Rayleigh–Taylor instability. The original data set is over 3 terabytes in size and took over 2 months to generate using the Miranda code on a nearly 2000-processor supercomputer at LLNL. With ViSUS, scientists could visualize such simulations while still in progress and either stop the simulation or adjust the variables if it did not progress as desired.

The ViSUS Algorithms

The algorithms that selectively choose the data to transmit first—one of the major challenges in designing ViSUS—are known as progressive algorithms. Traditional visualization algorithms proceed from fine to coarse—they first read the complete high-resolution information, from which users can generate lower-resolution approximations according to the capabilities of their computer. Progressive algorithms, in contrast, are algorithms that transition through multiple resolutions automatically. Input data are read from course to fine, and output is produced in the same order—coarse imagery is produced first, and from that, high-resolution output can be produced. However, ordinary progressive algorithms have less accuracy at higher resolution because of loss due to compression. ViSUS ensures that no accuracy is lost due to compression.

The visualization process begins with data retrieval, in which ViSUS analyzes the data and uploads the information needed to begin visualization right away. The next step is to build the mesh—the 3-D grid appropriate for the type of simulation and resolution. Next is the generation of filed data, which are the quantities computed by the simulation to describe the progress of a physical phenomenon in a given frame, such as the concentrations of two fluids at each location in the 3-D space in a fluid-mixing simulation. From these coordinates, isosurfaces are created. Isosurfaces are the 3-D surfaces that represent, for instance, the complex outer shape of twisting plumes of fluid. The final step is view-dependent rendering—selecting only the data that are "visible" to the user from the chosen perspective; everything else is excluded from the frame to conserve memory and processor use, and thereby speed visualization.

In conventional visualization, data are accessed row by row. For example in a 3-D fluid-mixing simulation with a resolution of 2000 × 2000 × 2000 bytes, a single 3-D time step would contain 2000^3 bytes, or 8 gigabyes, which would be beyond the memory capacity of most desktop computers. A single 2-D cross section would be 4 megabytes. If the cross section is changed, the entire dataset would have to be loaded again, 8 gigabytes at a time, to pick out the necessary points from each row. The ViSUS system, however, uses an advanced computational technique known as streaming Z curves to change the visualization without having to read the data all over again.

The algorithms that allow the ViSUS visualization engine to run on any computer are known as out-of-core algorithms. Because the dataset is much larger than the computer's main ("core") memory, these algorithms move blocks of data in and out of the computer's memory as they are processed. Conventional out-of-core algorithms can only transfer blocks of data between the source disc and the computer's memory at a fixed ratio. ViSUS, however, uses a type of out-of-core algorithm called cache-oblivious algorithms, which are not limited to a predetermined ratio. This optimizes transfer directly from the storage disc to core on any computer platform.

Other Applications and Advantages

The potential uses of ViSUS go beyond stockpile stewardship and other advanced physics applications. In a medical-imaging application, a doctor could use ViSUS for real-time examination of magnetic resonance imaging (MRI) data, even from an office or other remote location. Although initial resolution would be low in such cases, the doctor could later zoom in on points of interest at higher resolution. And because ViSUS uses lossless compression, no accuracy is lost—potentially important details are preserved when moving from low to high resolution.

The ViSUS software package is also independent of any hardware configuration and has already been adapted for four operating systems—SUN, Irix, Linux, and Windows—which should make porting ViSUS to additional operating systems easy.

Using ViSUS can enhance data security, because the dataset remains in one location for centralized control. This also assures data integrity: ViSUS streams the data and does not actually download the dataset to the user's computer. Therefore, a team of researchers can simultaneously use a single copy of the dataset from different locations, with no potential problems arising due to different versions of the dataset.

Future Work

One of the new techniques that Pascucci and his team are investigating as a means of accelerating the visualization process is occlusion culling, which means avoiding the unnecessary repeated constructing and drawing of portions of isosurfaces not visible in a given frame. Isosurface creation will also be sped up by adapting the algorithms to utilize the full capabilities of commercially available graphics cards in workstations.

As these and other upgrades enhance the robustness of ViSUS, the project team plans to allow LLNL scientists to use ViSUS in actual research and to encourage use of ViSUS technology in other visualization projects. In addition to benefiting the research itself, this use will also provide valuable feedback for further enhancements.

This powerful, flexible visualization tool will enable researchers to study large, supercomputer-generated data sets quickly, flexibly, and thoroughly. ViSUS will free scientists from having to use application-specific software or compete for access to dedicated visualization systems. In short, ViSUS allows scientists to spend more time on the science and less time and money on the tools.

— Paul Kotta

For more information, see the project summary (on CD only) for the project covered in this article:

ViSUS: Visualization Steams for Ultimate Scalability
Valerio Pascucci, Principal Investigator, 02-ERI-003

Related projects on computational simulation:

Adaptive Mesh Refinement Algorithms for Parallel Unstructured Finite Element Codes
Dennis Parsons, Principal Investigator, 03-ERD-027
Long-Time-Scale Atomistic Simulations
Wei Cai, Principal Investigator, 03-LW-027
Strategic Initiative in Applied Biological Simulations
Michael E. Colvin, Principal Investigator, 01-SI-012