Success Story

University of Utah's SCI Institute: Interactive Visualization of the World's Largest Scientific Data Sets

The Challenge

- Provide engineers and scientists with the ability to visualize multigigabyte data sets interactively
- Visualize very large high-resolution simulations with accurate, realistic movement of light, including refractions and reflections from glass and other transparent surfaces

The Solution

- An SGI® Reality Center™ facility, including a highly scalable, high performance, shared memory SGI® Onyx® family server with the CPU power to process massive industrial and scientific data sets and an SGI® storage array
- Star-Ray Interactive Ray Tracer, a ray-tracing software technology capable of high-resolution visualizations that accurately capture the physics of light

The Result

- Star-Ray, running on shared-memory SGI Onyx systems, enables researchers to visualize the largest data sets interactively
- Star-Ray running on SGI also delivers a uniquely accurate, interactive representation of light, a significant advantage in manufacturing applications such as headlight design

In 1992, the University of Utah formed an embryonic computer research group consisting of one professor and a handful of graduate students. Dr. Chris Johnson, the original group leader, still runs the program, but there have been some significant changes in the group: it has developed into a world-renowned research institute with a staff of 65 and is known as the Scientific Computing and Imaging Institute [SCI]. The National Institutes of Health [NIH] has named SCI the National Center for Research Resources for Bioelectric Field Modeling Simulation and Visualization. The SCI Institute has also been designated an Advanced Visualization Technology Center by the U.S. Department of Energy [DOE]. Significant projects for these institutions include the following:

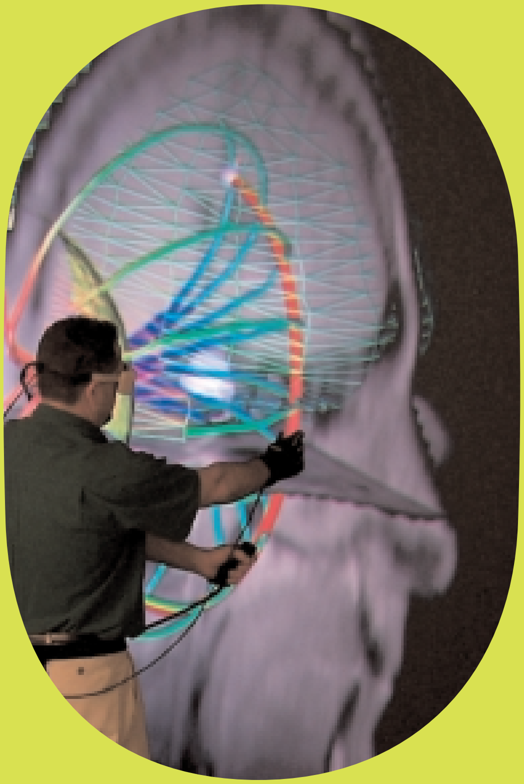

• For the NIH, the SCI Institute focuses on new techniques for modeling, simulating, and visualizing ECG and EEG data, including the localization of epileptic seizure sources and the design of defibrillation implants. The SCI Institute is also a part of a new NIH program aimed at integrating and visualizing data from multiple medical imaging and measurement technologies such as PET, MRI, and ECG.

• For the DOE, the SCI Institute is creating a problem-solving environment that will function as part of the SGI Accelerated Strategic Computational Initiative [ASCI] Blue Mountain installation at Los Alamos National Laboratory. Other SCI Institute projects for the DOE include development of a common-component architecture for software and data exchange and simulation methods for large fusion experiments as well as developing large-scale visualization techniques.

The SCI Institute has always worked closely with SGI. "We used SGI equipment from the beginning," says Johnson. "Then SGI combined high-performance computing with graphics power and ended up creating a computer that was perfect for the integrated environment we had in mind." The environment in

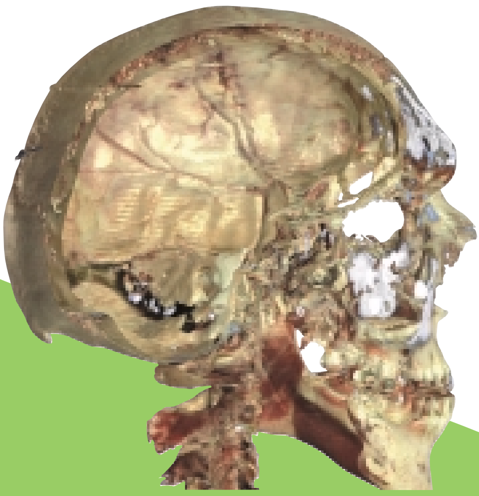

question became known as SCIRun (pronounced SKI run), and it enables the creation of scientific computing frameworks that integrate modeling, simulation, and visualization. The SCI Institute's Steven Parker, a graduate student turned research assistant professor, led the development of that project.

The SCI Institute currently operates a 24-processor SGI® Onyx2® with seven InfiniteReality® graphics pipes, as well as an SGI® Onyx® 3800 series system recently upgraded from 32 to 96 processors, with an InfiniteReality pipe that drives an immersive PowerWall display and a Fakespace® responsive workbench with a haptic device. An SGI® TP9400 storage array provides 6.5TB of storage. Scientific visualization remains the focus of the SCI Institute's work, and medical imaging is still a major interest.

Ray-Tracing: A Breakthrough Visualization Technology

Much of the SCI Institute's work has been aimed at developing technology that enables engineers and scientists to visualize extremely large data sets. "Scientists in the oil and gas, automotive, aerospace, and medical fields have data sets so big they've never been able to interact with complete pictures of them," says Dr. Greg Jones, associate director of the SCI Institute.

In the traditional raster-based visualization used by desktop PCs, a graphics card processes the entire data set to create a screen display. The system is graphics-card-limited; double the size of the data set, and you reduce performance by a factor of two. With the multigigabyte data sets of science and industry, interactive visualization becomes impossible. Star-Ray Interactive Ray Tracer, a ray-tracing technology developed by the SCI Institute, takes a totally different approach.

Instead of pushing the entire data set of polygons through a graphics card, Star-Ray uses CPU power to look back at the data from the viewpoint of the screen, casting a ray from each pixel to the data set to determine what will be displayed. It asks the question, What data will appear on the screen? The amount of data processed for display is determined by screen size only and does not vary with the size of the data set. Visualization performance is limited only by available CPU power.

Star-Ray and SGI: Taking Visualization to Another Level

Ray-tracing has been around since 1979, but has not come into wide use in the scientific community because its performance has been hardware-limited. Now Star-Ray, running on highly scalable, high-performance, shared-memory systems from SGI, has unlocked the technology's potential. "We can visualize a 25GB data set interactively," says Jones.

Star-Ray's scalability makes it a natural fit with the shared-memory SGI architecture, which allows easy scaling through the addition of bricks (CPU/memory modules). Unlike cluster configurations, it doesn't require modification of programs or data, and there is no memory-to-graphics or interprocessor performance loss. "The shared-memory architecture of SGI is a powerful environment for Star-Ray," says Jones.

"We can work with very large data sets and very high-resolution models, and we don't have to worry about writing applications to divide data across nodes and put it back together afterwards. We just put it into the SGI machine where all the processors can get to it. And when the models get bigger and bigger, SGI shared-memory really allows us to crank."

Star-Ray allows scientists to scale data sets to virtually any size with very little performance loss in interactive visualizations. In any case, visualization performance can be increased anytime by the addition of bricks. Scalability of the SGI architecture was recently demonstrated at NASA Ames, where one of the world's largest supercomputers, a 1,024-processor SGI® Origin® 3800 server, was used to create single system image climate and CFD visualizations. A single operating system controlled all CPUs, memory, and I/O functions to produce a single coherent image.

Star-Ray Shines Where the Data Sets Are Biggest

Star-Ray was demonstrated convincingly in the SGI booth in SIGGRAPH 2002 and is now being marketed by Visual Influence (*www.visual-influence.com*), an SCI spin-off of which Greg Jones is the president. Star-Ray is a high-end technology of great interest to engineers and scientists who need greater insight into the very large data sets being generated in science, medicine, manufacturing, and the oil and gas industry.

"The shared-memory architecture of SGI is a powerful environment for Star-Ray. We can work with very large data sets and very high-resolution models, and we don't have to worry about writing applications to divide data across nodes and put it back together afterwards. We just put it into the SGI machine where all the processors can get to it."

—Dr. Greg Jones,
President of Visual Influence,
SCI Associate Director

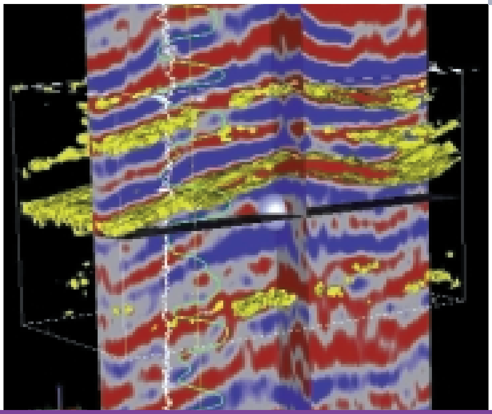

"Now we can volume-visualize multi-gigabyte, high-resolution data sets interactively on SGI systems and show researchers things they've never seen before."

—Dr. Chris Johnson,
Director,
SCI Institute

Above all, Star-Ray makes interactive visualization of very large data sets a reality. "If Star-Ray were being used significantly, people would see data sets they've never seen before," says Jones. "And they'd see them interactively." But ray-tracing has another major trick up its sleeve: it copies the physics of light in a way no other interactive graphics display technology does.

If, for example, you create a model of a car, truck, airplane, or even a building, ray-tracing will create shiny surfaces, including glass, with real-world reflections and refractions. Some automotive designers already use ray-tracing to study the way a curved windshield or a headlamp reflects or transmits light. Star-Ray makes their images more precise and realistic—and makes them more productive by providing fully interactive images. Aircraft designers can use Star-Ray to compare cockpit lighting scenarios with realistic, detailed, interactive graphics. Architects can visualize the way moving late-afternoon sunlight and shadows affect play on a football or baseball field, or conduct walk-throughs of building designs with the option of changing the camera path or material properties at any time. Oil and gas companies no longer need to downsample to visualize very large data sets, in which detail is discarded to achieve better performance.

The Future of Visualization: The View from the SCI Institute

"Back in the '80s we'd try to model a brain using 100 polygons," says Chris Johnson. "The neurosurgeons were not impressed. Now we can volume-visualize multigigabyte, high-resolution data sets interactively on SGI systems and show researchers things they've never seen before. When they first saw their data in an immersive environment that allowed them to look inside a 3D volume and decided to alter their surgery plans as a result—that's when we became value-added.

"Using that experience as an indicator, the visualization of these very large data sets is simply going to explode. More and more people are going to utilize very high-resolution models, simulations, and visualizations to do their work in science, engineering, and medicine. And SGI technology is going to play an important role."

Star-Ray, running on SGI shared-memory computing systems, has opened up a new world of understanding and insight to scientists and engineers who work with the world's largest simulation data sets. This power of interactive visualization, combined with realistic detail and an accurate portrayal of the movement of light, makes the combination of Star-Ray software and SGI shared-memory supercomputers a potent force at the highest levels of research and design.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 6771.0290