

Visualizing Science: The OptIPuter Project

For many areas of science and engineering, visualization has become more than a convenient tool: it has become a necessity for interpreting the enormous amounts of data being produced by large-scale instrumentation, experiments and simulations. One of the most creative—and useful—ways of presenting scientific data is the large-scale tiled display. Ranging from tabletop versions to fully immersive 3D stereo virtual reality environments, these displays enable users to explore features that may be hard or impossible to identify with conventional, personal computers.

Concluding the first session of the SciDAC 2007 conference in Boston, SciDAC Program Director Dr. Michael Strayer commented that one of the major challenges facing the SciDAC multidisciplinary teams was “developing new visualization and data management systems to support moving, storing, analyzing, manipulating, and visualizing multi-petabytes of scientific data and objects.”

Scientific visualization indeed has become a critical tool in understanding large datasets and the behavior of complex systems. And a special form of scientific visualization—large-scale tiled displays—is enabling scientists and engineers to explore a wide range of applications, from structural engineering of large buildings to medical studies of extremely small cellular structures.

Tiled displays are not new. During the early 1990s a number of planar tiled displays were built and used for scientific studies. Several factors limited their wide acceptance, however. They were expensive, large, and hence unsuitable for the typical office environment. To address this situation, the California Institute for Telecommunications and Information Technology (Calit2) at the University of California—San Diego (UCSD), and the Electronic Visualization Laboratory (EVL) at the University of Illinois—Chicago (UIC) initiated the OptIPuter project.

The goal of the OptIPuter project is to develop innovative methods for scientists to explore very large remote data objects in an interactive and collaborative fashion. It is an ambitious goal, and the

developers have enjoyed remarkable success. Without a doubt the highlight is the “OptIPortal”—a commodity-priced, easily replicable display, with tens to hundreds of megapixels, which can be used by individual researchers or collaborative groups (figure 1).

“The OptIPortal is not a product,” emphasizes Calit2 research scientist Dr. Thomas DeFanti, UIC distinguished professor emeritus of computer science and co-principal investigator (co-PI) of the OptIPuter. (Larry Smarr of Calit2 is the PI, Jason Leigh of EVL is another co-PI, as are UCSD’s Phil Papadopoulos and Mark Ellisman). The OptIPuter website provides system documentation and hardware specifications, enabling researchers to build their own OptIPortals. The hardware is based on commodity PCs. The display tiles can be built from LCD displays, such as those used for desktop computers, plasma TVs, or video projectors. The software, based on EVL’s Scalable Adaptive Graphics Environment (SAGE) specialized middleware, is easily downloadable and installed, along with other visualization libraries, by UCSD’s open source Rocks configuration technology. The typical OptIPortal PC cluster node includes two gigabit Ethernet network interface cards (one for normal networking and control and the other for data transfer over the 10-gigabit OptIPuter private network), as well as national and international multi-gigabit networks, which connect to remote instruments, supercomputers, and data storage systems.

A special form of scientific visualization—large-scale tiled displays—is enabling scientists and engineers to explore a wide range of applications, from structural engineering of large buildings to medical studies of extremely small cellular structures.



Figure 1. Faculty and students conduct weekly research progress meetings in front of the 100-megapixel LambdaVision display at EVL/UIC. Using EVL's SAGE middleware, a variety of high-resolution information is juxtaposed on the walls to enable group discussion. Participants can be in the room or communicating over distance mediated via high-definition video conferencing.

OptIPortal Deployment Worldwide

OptIPortals are being deployed at academic institutions around the world:

- **Antarctica** McMurdo Station
- **Australia** University of Melbourne
- **Canada** Canadian Research Centre
- **China** Computer Network Information Center of the Chinese Academy of Sciences
- **Czech Republic** Masaryk University
- **Japan** Kyoto University; National Institute of Advanced Industrial Science and Technology (AIST); Osaka University
- **Korea** Korea Institute of Science and Technology Information (KISTI); Gwangju Institute of Science and Technology (GIST)
- **Mexico** Centro de Investigacion Cientifica y de Educacion Superior de Ensenada (CICESE)
- **Russia** Space Research Institute of the Russian Academy of Sciences in Moscow; Science and Innovation Center in Chernogolovka, Russia
- **The Netherlands** University of Amsterdam; SARA
- **Switzerland** University of Zurich
- **Taiwan** National Center for High-performance Computing (NCHC)
- **United Kingdom** Oxford University
- **United States** Adler Planetarium; Massachusetts Institute of Technology; NASA Goddard; Purdue University; Texas A&M University; U.S. Geological Survey (USGS) Earth Resources Observation and Science (EROS); UC-Davis; UC-Irvine/Calit2; UC-Santa Cruz; UCSD/Calit2; UCSD National Center for Microscopy and Imaging Research; UCSD Scripps Institution of Oceanography; University of Hawaii; University of Illinois-Chicago/EVL; University of Illinois-Urbana-Champaign/National Center for Supercomputing Applications (NCSA); University of Michigan; University of Southern California; University of Washington

OptIPortals are also being used by several corporations, including:

- Sharp Laboratories of America
- Rincon
- Nortel Networks

The goal of the OptIPuter project is to develop innovative methods for scientists to explore very large remote data objects in an interactive and collaborative fashion.



L. RENARDI - EVL/UIC

Figure 2. Tellurion, a real-time planetary-scale data visualization being developed by EVL/UIC Ph.D. candidate Robert Kooima, is shown here on the Calit2 OptiPortable (the portable OptiPortal) at SC07. Tellurion uses NASA's Blue Marble Next Generation and Shuttle Radar Topography Mission data.

The cost depends in part on the application. OptiPortal applications fall into two broad categories: those that optimize for streaming video, striving for a balance between processor bandwidth and communications speed, and those that maximize pixel count per dollar. The former usually have LCD displays with 2.0–2.5 megapixels, each driven by one PC; the latter are generally built with two or four 4-megapixel panels per PC.

User Communities

This flexibility has resulted in the construction of many different OptiPortals. The OptiPuter team has developed a number of devices that range from a portable 15-tile, 2-megapixel OptiPortal called the OptiPortable (figure 2), which can be easily shipped to conferences, to a 70-tile, 286-megapixel display wall (figure 3). Already, a wide community of scientific users pursuing diverse applications of the OptiPortal has been established (sidebar “OptiPortal Deployment Worldwide” p33).

- **Biomedicine** The National Center for Microscopy and Imaging Research (NCMIR), located at UCSD, is home to the coordinating center of the Biomedical Informatics Research Network, which provides partner sites with information technology infrastructure. With funding

from Microsoft, NCMIR has created a Windows-based OptiPortal. The device was successfully demonstrated in Microsoft's booth at SC07—the international conference for high-performance computing, networking, storage, and analysis—as part of the NCMIR's research collaborations with colleagues at Oxford University.

- **Geosciences** EVL has developed and is exploring the use of CoreWall, an integrated environment for interpreting geoscientific data from sediment and crystalline cores (figure 4, p36). The device, which relies on OptiPuter technologies, is being used by the international Antarctic geological drilling program, a research partnership comprising the United States, Germany, Italy, and New Zealand.

- **Geoinformatics** The U.S. Geological Survey (USGS) National Center for Earth Resources Observation and Science (EROS) deployed a 15-panel OptiPortal, which was demonstrated to more than 200 researchers and federal officials visiting the USGS/EROS national center in 2007. USGS geophysicists have also worked with EVL and Calit2 staff to test the connectivity of their OptiPortal to the National LambdaRail, a nationwide infrastructure providing the U.S. research community with access to many different



Figure 3. The 70-tile, 286-megapixel Highly Interactive Parallelized Display Space (HIPerSpace) OptIPortal, developed by the Calit2 VisGroup, displays ultra-high-resolution imagery from the Mars Exploration Rover Mission in different wavelengths, notably, McMurdo, Cape Verde, and Duck Bay.

regional networks. Tests of the OptIPortal conducted with and without federal firewall constraints demonstrated that it could be used in times of emergencies, such as hurricanes, to display and distribute large data volumes using advanced optical network connectivity. These results were presented at the Geoinformatics Conference in 2007 in an invited presentation titled “From CAVEs to OptIPortals.”

● **Marine Microbial Analysis** As part of the Community CyberInfrastructure for Advanced Marine Microbial Ecology Research and Analysis (CAMERA) project, the Gordon and Betty Moore Foundation is funding Calit2 to assist the placement of OptIPortals at several leading metagenomic laboratories nationwide, including UCSD, UC–Davis, the University of Washington, Massachusetts Institute of Technology (MIT), and the J. Craig Venter Institute in Rockville, Maryland. The OptIPortals are being used for interactive study of microbial genomes and emerging applications, such as genetically engineered biofuel generation. At SC07, Dr. Ginger Armbrust, a researcher at the University of Washington, demonstrated how the OptIPortal is enabling new insights in the study of Puget Sound diatoms.

● **Real-Time Observatories** The ROADNet (Real-time Observatories, Applications and Data management Network) project at UCSD, funded by the National Science Foundation (NSF), is intended to enhance capacity to monitor and respond to changes in the environment by developing both the wireless networks and the integrated, seamless, and transparent information management system to deliver seismic, oceanographic, hydrological, and physical data in real time. ROADNet receives this information from various sensors throughout Southern California and then displays the images simultaneously on OptIPortals. The technology was demonstrated at the American Geophysical Union conference in December 2007.

● **Space Visualization** Little more than a year ago, the Adler Planetarium in Chicago opened its Space Visualization Laboratory, which features a variety of scientific visualization applications and display technologies, including an OptIPortal developed in collaboration with EVL. Visitors are encouraged to participate in user studies involving both application and pre-installation prototypes.

● **Technology Training** NSF funds Florida International University’s Global CyberBridges project




Figure 4. CoreWall is an integrated visualization tool for the study of lake and ocean sediment cores. It is a collaborative project among EVL/UIC, the U.S. National Lacustrine Core Repository (LacCore) at the University of Minnesota, the Integrated Ocean Drilling Program (IODP) through the Joint Oceanographic Institutions, the Lamont–Doherty Earth Observatory (LDEO) at Columbia University, and the ANDRILL Antarctic Drilling Program.

to use OptIPortals and SAGE in an international collaboration to improve the technology training for a new generation of scientists and to increase the rate of discovery for all domains. Project participants come from the Chinese Academy of Sciences, the City University of Hong Kong, and the University of Sao Paulo, Brazil.

StarCAVE

One of the most elaborate members of the OptIPortal family is the StarCAVE (figure 5). An immersive virtual reality system, it uses tiled projector displays and OptIPuter technologies to surround the viewers 360 degrees with audio and life-size images, enabling users to explore features that may be hard or impossible to identify with conventional, non-immersive techniques. The concept of the CAVE, a recursive acronym for Cave Automatic Virtual Environment, dates back to 1991 when it was conceived at EVL. When the CAVEs were first described, the idea of immersing oneself in a human cell or walking around the inside of a reactor seemed too much like science fiction. But after the initial showing to thousands of people at SIGGRAPH'92 and SC92, scientists began to recognize the possibilities of this innovative technology. At Argonne National Laboratory, for example, the

CAVE was used for real-time simulation and analysis of pollution control systems for commercial boilers and incinerators.

The original CAVEs suffered from several limitations, however. Smooth animation was difficult because the graphics engines were primitive; the lighting was dim because the projectors were not very bright; and the visual acuity was low, only about 20/140, because the projectors were relatively low resolution. The newest generation of the CAVE is brighter, bigger, and better than its ancestors. “The StarCAVE is fast, bright, and crisp,” says Dr. DeFanti. “Using 34 HDTV 1080p JVC projectors and Nvidia Quadro 5600 cards, we have been able to increase visual acuity in the StarCAVE to approximately 20/40. Increasing the acuity to 20/20 would cost five times as much in 2008.”

Developed at Calit2 by the team who developed the hardware for the original CAVEs (Dr. DeFanti, Dan Sandin, and Gregory Dawe), the StarCAVE operates at a combined resolution of more than 60 million pixels (30 million per eye) distributed over 15 rear-projected wall and two floor screens. It uses liquid crystal on silicon (LCOS) LCD projectors, offering smooth pixel edges and rich colors. Moreover, the heavy shutter glasses of the old CAVEs have been replaced by lightweight, circular polarized “sunglasses.” Sonification is achieved by using

When the CAVEs were first described, the idea of immersing oneself in a human cell or walking around the inside of a reactor seemed too much like science fiction, but after the initial showing to thousands of people, scientists began to recognize the possibilities of this innovative technology.




Figure 5. The Calit2 StarCAVE physical structure has five walls and a floor, with each wall consisting of three screens and the floor consisting of two screens. Shown here is the exterior of the StarCAVE with the entry wall almost closed.



Figure 6. StarCAVE displaying data from UCSD Protein Data Bank. RNA polymerase enzyme data from the UCSD Protein Data Bank is displayed in the StarCAVE. The RNA polymerase is from yeast cells and is composed of a dozen different proteins. Together, they form a machine that surrounds DNA strands, unwinds them, and builds an RNA strand based on the information held inside the DNA. Once the enzyme gets started, RNA polymerase marches confidently along the DNA copying RNA strands thousands of nucleotides long. StarCAVE virtual-reality development is by Philip Weber and Jurgen Schulze of Calit2.

wave-field synthesis speakers, and the users can interact with the data via an ART, Inc. wand/multi-camera wireless tracking system. Even the shape is nontraditional: the StarCAVE is a five-sided virtual reality room (plus floor), offering several advantages over the standard CAVE in the form of a cube, including minimization of glare and oblique views.

Arguably, the StarCAVE is expensive (about \$1 million). Nevertheless, the scientific benefits of this newest-generation CAVE are readily acknowledged by avid researchers, who can get right inside their simulations in real time, gaining new insights into complex phenomena. The StarCAVE is being used in a wide range of areas, from molecular biology to neuroscience. The following is a sampling of some applications written for the StarCAVE in 2007 and early 2008.

- Comparative Analysis** With the establishment of databases such as the Protein Data Bank (PDB), researchers can easily obtain access to more than 50,000 structures (figure 6). Not so easy, however, is analyzing those proteins, especially if one wishes to compare numerous protein structures at once. Several viewers and user interfaces have been developed, but these typically restrict the user to several proteins. To address this problem, researchers at Calit2, in collaboration with researchers at EVL, developed the Protein Viewer. This application connects the StarCAVE to the PDB data bank server to download one of approximately 50,000 protein structures and display it in 3D. PyMOL, a free molecular graphics package, is used to convert the PDB files to the Virtual Reality Modeling Language format, which the StarCAVE visualization software can read. The user can choose between different visualization modes (cartoon, surface, or sticks), load multiple proteins, align proteins, or display the corresponding amino acid sequence. The high resolution of the StarCAVE and its surround display capability provides a view of a large number of protein structures at once, which scientists like to find similarities between proteins.

- Neuroscience** Researchers at UCSD are currently working with neuroscientists to find out whether the human brain works similarly for way-finding tasks in the StarCAVE and in the real world. In this application, users are tracked as they walk through a virtual replica of the Calit2 building in the StarCAVE (figure 7).

- Geospatial Visualization** The open source software library ossimPlanet provides accurate, high-performance, double-precision 3D geospatial visualization. Used with OpenCOVER, it allows StarCAVE users to “fly” over the Earth with highly detailed satellite images and 3D terrain. This application programming interface (API) already has been used in several application domains, from the visualization of ocean currents to genomics.

- Social Science** Opened in 2003 in Florence, Italy, ClusterCity featured two virtual reality environments in the Palazzo Vecchio: one representing Florence in the Renaissance age and the other depicting Chicago in the gigabits age. Motivated by the success of this endeavor, UCSD researchers recently constructed a visualization of the Palazzo Vecchio itself. The application involves the visualization of a LIDAR scan of the Hall of the 500, with about 26 million points as a point cloud in the StarCAVE. A custom data distribution and rendering algorithm achieves a rendering rate of about 10 frames per second, giving a real feeling of “being there.”




Figure 7. StarCAVE running the Calit2 building architectural model. The Calit2 StarCAVE enables participants to walk through a 3D architectural model of the Calit2 building. The model was created by Ben Maggos of the UCSD Center for Research in Computing and the Arts and by Daniel Rohrlack of the UCSD Computer Science department and was ported to the StarCAVE by Jurgen Schulze of Calit2.

From Walls to the Desktop

From powerwalls to immersive virtual reality environments, the OptIPuter project has provided researchers with a range of devices. For instance, the EVL team has designed and developed a next-generation tabletop high-resolution display, called LambdaTable, and was recently funded to develop an immersive autostereo version, called OmegaTable.

The LambdaTable is designed for interactive group-based data visualization (figure 8, p40). The 24-megapixel digital table comprises six LCD 2560 × 1600 displays driven by six computers. Six infrared cameras are mounted overhead to track input from many users simultaneously interacting with various interface devices, or pucks. Each object is identified by a unique pattern of embedded infrared LEDs. A real-time computer vision library collects the camera video streams to determine the orientation and position of each input device, which is then passed to the application running on the LambdaTable display.

The Rain Table is one such application (figure 9, p40) and is being developed by EVL in collaboration with the Department of Geology at the University of Minnesota and the Science Museum of Minnesota to study rainfall. Using electronic puck input devices, museum visitors indicate the loca-

tion of rainfall on two-dimensional topological maps of locations on Earth. The visitors can then observe how the rain flows down mountains and develops channels of water that cut through slopes, merge, and flow into streams and oceans. The Rain Table exhibit will help the public understand how rain in one place can cause flash floods in another. By extension, the museum participants can also see how the distribution of pollutants, such as lawn chemicals, can impact water supplies.

What's on the Horizon?

With successful demonstration of the OptIPuter technologies in national-scale geoscience and bioscience projects, the OptIPuter team is ready to enhance the infrastructure and expand its user communities.

Currently, the effectiveness of the OptIPortal implementations depends critically on the choice of the graphics processing unit (GPU) and PC, matching the slots and power supply capacity. Ascertaining that the most up-to-date configurations are offered requires ongoing analysis. One goal, therefore, is to design “standard” OptIPortal specifications, hardening both hardware and Rocks and SAGE core software specifications for ease of use by new user communities. Other plans include further integration with advanced communication

With successful demonstration of the OptIPuter technologies in national-scale geoscience and bioscience projects, the OptIPuter team is ready to enhance the infrastructure and expand its user communities.

From powerwalls to immersive virtual reality environments, the OptIPuter project has provided researchers with a range of devices.




Figure 8. LambdaTable visualizing a Purkinje cell. Collaborators sit at LambdaTable, a tiled LCD tabletop display that supports interactive group-based visualization of high-resolution data developed by EVL/UIC, examining a volume visualization of a Purkinje cell from a rat's brain that was imaged using a multi-photon microscope at the National Center for Microscopy and Imaging (NCMIR), UCSD. Each layer is a mosaic of smaller, high-resolution images that were stitched together, and then the layers were stacked to create a volumetric view of the neuron.



Figure 9. LambdaTable running the Rain Table museum exhibit. Rain Table is an interactive museum exhibit developed for the LambdaTable by EVL/UIC graduate student Dmitri Svistula in collaboration with the Department of Geology at the University of Minnesota and the Science Museum of Minnesota.

protocols; integrated teleconferencing support with various high-definition video and audio streams; and additional system and application monitoring, analysis, and performance.

EVL is also continuing to enhance its visualization tools. A new tool, called Tellurion by its author, EVL doctoral candidate Robert Kooima, was developed last year to enable the merging of

multi-resolution geographic and image data. Unlike traditional geographic rendering techniques, Tellurion's data composition is generated dynamically and in real time (figure 10). "Many terabytes of data exist describing the Earth, the Moon, and other planets. This continues to increase exponentially as more powerful sensors are put into service," said Mr. Kooima. "While




Figure 10. *Tellurion shown on Varrier autostereo tiled display. Tellurion, a real-time planetary-scale data visualization being developed by EVL/UIC doctoral candidate Robert Kooima, is shown here on the Calit2 Varrier autostereoscopic tiled display. Tellurion uses NASA's Blue Marble Next Generation and Shuttle Radar Topography Mission data.*

much of these data are related, little has been done to bring widely disparate data together into a common visualization. We're working on GPU-centric solutions to this problem that allows separate, but related, planetary-scale datasets to be flexibly composed in powerful ways."

The researchers also hope to integrate better audio technologies into the StarCAVE. "Better audio will provide an important dimension for scientists to better understand complex physical and biological phenomena," says Dr. DeFanti.

Contributors This article was written by the following individuals, with the assistance of Dr. Gail W. Pieper, senior coordinator of technical editing and writing in the Mathematics and Computer Science Division at Argonne National Laboratory: Dr. Thomas A. DeFanti, Calit2 at UCSD, and EVL at UIC; Qian Liu, Mason Katz, Dr. Phil Papadopoulos, Joseph Keefe, Dr. Greg Hidley, Greg Dawe, Ian Kaufman, Bryan Glogowski, Dr. Kai-Uwe Doerr, Dr. Jurgen P. Schulze, Dr. Falko Kuester, Prof. Peter Otto, Dr. Ramesh Rao, and Dr. Larry Smarr, Calit2 at UCSD; Dr. Jason Leigh, Dr. Luc Renambot, Alan Verlo, Lance Long, Maxine Brown, Prof. Dan Sandin, Venkatram Vishwanath, and Robert Kooima, EVL at UIC; Dr. Javier Girado, Qualcomm, Inc.; and Dr. Byungil Jeong, Sun Microsystems

Acknowledgments UCSD receives major funding from the National Science Foundation (NSF) for the OptiPuter, cooperative agreement OCI-0225642, Larry Smarr, principal investigator. UCSD also receives funding for OptiPortal and application development from the Gordon and Betty Moore Foundation for CAMERA, the Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis, and from the State of California for StarCAVE. UIC receives NSF support for OptiPuter via UCSD, as well as NSF award CNS-0420477 for LambdaVision/LambdaTable and NSF award OCE-0602117 for CoreWall. UIC also receives funding for OptiPortal and application development from the Science Museum of Minnesota, the Adler Planetarium & Astronomy Museum and the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC). Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies and companies

Further Reading

<http://www.optiputer.net>
<http://www.evl.uic.edu/cavern>
<http://www.calit2.net>
<http://www.scivee.tv/node/3648>
<http://www.youtube.com/evltube>