

The global lambda visualization facility: An international ultra-high-definition wide-area visualization collaboratory

Jason Leigh^{a,*}, Luc Renambot^a, Andrew Johnson^a, Byungil Jeong^a, Ratko Jagodic^a, Nicholas Schwarz^a, Dmitry Svistula^a, Rajvikram Singh^a, Julieta Aguilera^a, Xi Wang^a, Venkatram Vishwanath^a, Brenda Lopez^a, Dan Sandin^a, Tom Peterka^a, Javier Girado^a, Robert Kooima^a, Jinghua Ge^a, Lance Long^a, Alan Verlo^a, Thomas A. DeFanti^a, Maxine Brown^a, Donna Cox^b, Robert Patterson^b, Patrick Dorn^b, Paul Wefel^b, Stuart Levy^b, Jonas Talandis^c, Joe Reitzer^c, Tom Prudhomme^c, Tom Coffin^d, Brian Davis^e, Paul Wielinga^f, Bram Stolk^f, Gee Bum Koo^g, Jaeyoun Kim^h, Sangwoo Han^h, JongWon Kim^h, Brian Corrieⁱ, Todd Zimmermanⁱ, Pierre Boulanger^j, Manuel Garcia^j

^a *Electronic Visualization Laboratory, University of Illinois at Chicago, United States*

^b *National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign, United States*

^c *Technology, Research, Education, and Commercialization Center/NCSA, United States*

^d *ACCESS Center/NCSA, United States*

^e *Earth Resources Observation Systems (EROS), US Geological Survey (USGS), United States*

^f *SARA Computing and Networking Services, Netherlands*

^g *Korean Institute of Science and Technology Information (KISTI), Republic of Korea*

^h *Gwang-ju Institute of Science and Technology, Republic of Korea*

ⁱ *Simon Fraser University, Canada*

^j *University of Alberta, Canada*

Available online 2 May 2006

Abstract

The research outlined in this paper marks an initial global cooperative effort between visualization and collaboration researchers to build a persistent virtual visualization facility linked by ultra-high-speed optical networks. The goal is to enable the comprehensive and synergistic research and development of the necessary hardware, software and interaction techniques to realize the next generation of end-user tools for scientists to collaborate on the global Lambda Grid. This paper outlines some of the visualization research projects that were demonstrated at the iGrid 2005 workshop in San Diego, California.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Graphics; Visualization; Collaboration; Remote networking; High-definition video

1. Introduction

In a decade's time, high-performance computing has proven its value in Science, Medicine, Engineering, Education, and

even in Art. These data-intensive domains rely on high-performance computing to process the terabytes of raw data (for example, from sensors) to produce meaningful insight, often enabled by high-quality visualizations. The need for the computing infrastructure to support collaborative work between distributed users has also grown, as research has increasingly become multi-disciplinary endeavors.

Up until the last five years however, the predominant model for supporting data-intensive collaborative science and

* Corresponding address: University of Illinois at Chicago, Department of Computer Science, 851 S Morgan St, MC 152, 1120 SEO, 60607 Chicago, IL, United States. Tel.: +1 312 9963002.

E-mail address: spiff@uic.edu (J. Leigh).

engineering involved replicating and co-locating the computing resources with the storage systems, and then transmitting the highly distilled results of the computation to expensive remote visualization systems at each collaborator's site. This was because bandwidth was a scarce commodity and so it was more efficient to co-locate the data with the computation than to move the data, on-demand, to the computation. This meant that computer centers not only had to maintain clusters of computers for number crunching, they also had to maintain backed-up storage facilities to keep copies of the raw data for their users. Recently, with rate of decline of the cost of bandwidth far exceeding that of computing and storage, it has become viable for scientists to connect to ultra-high-speed networks, more cheaply than they can build and maintain large computing, storage, and visualization systems¹ [4]. This is the fundamental premise behind the concept of shared cyber-infrastructure, and is being deeply researched amongst the international partners in the OptIPuter project [6]. By bridging high-performance computing systems with ultra-high-speed networks that can approach the capacity of computer system buses, one can potentially reduce the need to duplicate systems except where redundancy (for example, for backup) is important. In the same way that users today search for information on the Web with a browser and rely on Google to maintain the vast clusters of computers that organize information on the Internet, scientists who could not ordinarily afford high-performance computing, can collaborate and access the network of cyber-infrastructure (the Lambda Grid) using low-cost, scalable, high-resolution network-aware displays and high-definition video conferencing [5].

The research outlined in this paper and demonstrated at the iGrid 2005 workshop² (www.igrid2005.org) in San Diego, California, marks an initial global cooperative effort between visualization and collaboration researchers to build a persistent virtual distributed facility (GLVF—the Global Lambda Visualization Facility) to enable the joint research and development of the hardware, software and interaction techniques, that are to become the primary interfaces between domain scientists and the Lambda Grid. The institutions that participated in the iGrid GLVF demonstration include:

- The Electronic Visualization Laboratory at the University of Illinois at Chicago;
- The California Institute for Telecommunications and Information Technology (CALIT2);
- The National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign, and the ACCESS Center in Washington DC;
- Earth Resources Observation Systems (EROS), US Geological Survey (USGS);
- The Technology, Research, Education, and Commercialization Center in Dupage County, Illinois;

- SARA Computing and Networking Services in Amsterdam, the Netherlands;
- Gwang-ju Institute of Science and Technology (GIST) in Gwang-ju, South Korea;
- Korean Institute of Science and Technology Information (KISTI) in Daejeon, South Korea;
- Simon Fraser University and the University of Alberta in Canada.

GLVF sites are all equipped with a wide range of visualization technologies including high-definition passive stereo displays, fully immersive CAVEs, auto-stereoscopic displays, ultra-high resolution tiled projection and LCD displays, and high-definition cameras (Fig. 1). In most cases these systems are already connected to high-speed multi-gigabit metropolitan, national and international networks (Fig. 2). Internationally, GLVF leverages the networking infrastructure that has been amassed as part of the Global Lambda Integrated Facility (GLIF—www.glif.is). This enables GLVF to become a one-of-a-kind resource for conducting long-term systematic research in network- and data-intensive visualization and collaboration. For domain scientists, GLVF can provide a means for them to learn how to apply Lambda Grid technology to transform their own areas of research.

2. GLVF at iGrid

At the iGrid workshop GLVF researchers streamed real-time ultra-high-resolution graphics, and high-definition video from their home sites for display on a variety of display systems including: a 100-megapixel LambdaVision tiled LCD display (www.evl.uic.edu/cavern/lambdavisio); a passive stereo XGA resolution GeoWall display (www.geowall.org); a passive stereo high-definition display; and the auto-stereoscopic Varrier display [10].

2.1. The Scalable Adaptive Graphics Environment

Luc Renambot (luc@mac.com), Byungil Jeong, Ratko Jagodic, Rajvikram Singh, Andrew Johnson, Jason Leigh—Electronic Visualization Laboratory (EVL), University of Illinois at Chicago.

The OptIPuter project envisions that in the future it will become routine for users to work and collaborate in rooms whose walls are made from seamless ultra-high-resolution displays that are fed by data streamed over ultra-high-speed networks from distantly located visualization, storage servers, and high-definition video cameras [6,7]. EVL has taken the first steps towards this vision by building LambdaVision—a 17-foot wide, tiled display built from an 11 × 5 array of LCD screens with a total resolution of 100 megapixels. LambdaVision is primarily enabled by the Scalable Adaptive Graphics Environment (SAGE) [9] a middleware system for managing visualization and high-definition video streams that are presented on the ultra-high-resolution display. Each visualization application (such as 3D rendering, remote desktop, video streams, large 2D maps) streams its rendered pixels (or graphics primitives) to SAGE, allowing for any given

¹ Gary Stix, The Triumph of the Light, Scientific American, January 2001—bandwidth is doubling every 9 months; computing at every 18 months; and storage at every 12 months.

² See editors' article in this issue of the Future Generation Computer Systems Journal for a description of iGrid 2005.



Fig. 1. Top left: SAGE showing several high resolution streams (high-definition thunderstorm simulation; high-definition live video; ultra-high-resolution aerial photography of New Orleans after Hurricane Katrina) simultaneously on the Lambda Vision display; Bottom left: TOPS showing a single high resolution montage of Amsterdam; top right: auto-stereoscopic video conferencing on the Varrier display; middle right: Solutions Server streaming graphics to a passive stereo XGA resolution display; bottom right: BPlayer streaming high-definition stereo movies to a high-definition passive stereo display.




Fig. 2. Network paths and bandwidth used during the iGrid GLVF demonstration. [NLR—National Lambda Rail; PNG—Pacific Northwest Gigapop; CAVEWave—UIC’s private 10G link over NLR; additions 10G link from Chicago to San Diego donated by Qwest network.]

layout onto the displays (e.g. the output of arbitrary M by N pixel rendering cluster nodes can be streamed to X by Y pixel display screens).

The graphics streaming architecture of SAGE addresses two non-trivial problems in scientific visualization. One is heterogeneity: since most visualization applications are closely tied to their graphics environment, it is difficult to integrate

various visualization applications into a unified graphics environment. The other is scalability: the ability of visualization software and systems to scale in terms of the amount of data they can visualize and the resolution of the desired visualization [8]. SAGE addresses the heterogeneity problem by decoupling graphics rendering from graphics display so that visualization applications developed on various environments

can easily migrate into SAGE by streaming their pixels into the virtual frame buffer. Also, SAGE provides scalability by supporting any number of rendering and displaying nodes, number of tiles, and screen resolution, and the SAGE visualization applications have extremely fast access to huge databases at remote or local sites taking advantage of affordable ultra-high-bandwidth networks.

At the iGrid 2005 workshop, EVL succeeded in simultaneously streaming a wide variety of content including: real-time parallel volume rendered images of the Visible Korean Human from KISTI (over KREONet2); NCSA's pre-rendered high-definition movies of tornado simulations from SARA (over SURFNet); and 1-foot resolution aerial photographs of New Orleans in the aftermath of Hurricane Katrina, and live high-definition video from EVL (over CAVEWave/National Lambda Rail). 10 Gb/s was allocated for the experiment, which ran for a total of 6 hours over 3 sessions. 9 Gb/s of sustained throughput was achieved; therefore approximately 4 terabytes of dynamic image data were transferred over an hour, and a total of 24 terabytes for the period of iGrid (equivalent to 6075 single-layer DVD disks).

2.2. TOPS—Technology for Optical Pixel Streaming

Paul Wielinga (wielinga@sara.nl), Bram Stolk—SARA Computing and Networking Services.

SARA has been studying the separation of the (virtual) graphics card and the display in the context of the OptIPuter principle. Our challenge is to cope with latency on long-distance connections. SARA's approach to driving tiled displays from remote render clusters is a lightweight approach. Our key insight into this problem is based on the short lifespan of a picture element in a video stream. At 50 frames per second, a picture element, or pixel, is only visible for 20 ms. If such a pixel is lost during communication, chances are it will remain unnoticed, because apart from spatial coherence, video streams also exhibit temporal coherence. Our approach therefore uses unreliable network transport (UDP) as opposed to a reliable (TCP) transport. This will give a valuable benefit, as high throughput is easily realized when acknowledgements from remote sites with large latencies are no longer required.

The viability of our approach was successfully demonstrated during iGrid 2005, where ultra-high-resolution images were streamed live from Amsterdam to San Diego over two 10 Gb/s connections, whereby the application consumed 18 Gb/s bandwidth, with peaks of 19.5 Gb/s. The application was controlled from the display side, and the low bandwidth control signals were sent reliably over TCP to Amsterdam. The high-bandwidth stream, representing ultra-high-resolution imagery, was displayed on the LambdaVision.

Bandwidth usage was controlled by the observer, by varying frame rate on the render cluster. Initially, the frame rate was limited by the receiving side, as the large amount of small 1500 byte packages locked up the CPU's of the receiving display nodes. The total connection between the render cluster and the display cluster was later jumbo-frame enabled. When using an MTU size of 9000 bytes, we could fill all available

bandwidth. The frame rate on the panel was increased to 9 frames/s, creating a data stream of 2.3 GB/s. This is the equivalent of transferring a complete DVD in 2 seconds over a 5000 miles distance. The goal is to be able to reach at least 15 frames/s, however this was beyond the 20 Gb/s bandwidth that was allocated during the demonstration. The opportunity will present itself again once a stable GLVF has been established.

2.3. Improved high-quality media service architecture for multi-party collaboration environments

Jaeyoun Kim, Sangwoo Han, and JongWon Kim (jykim@nm.gist.ac.kr, swhan@nm.gist.ac.kr, jongwon@nm.gist.ac.kr)—Networked Media Lab (NML), Dept. of Information and Communications, Gwangju Institute of Science and Technology (GIST, formerly known as K-JIST).

With the advent of Gb/s-range Research and Engineering networks, it is now possible to accommodate bandwidth-demanding high-quality (DV/HDV/stereo HD/uncompressed HD) video services into computer-mediated collaboration environments such as the AccessGrid (AG) [1]. In such environments, the enhanced media service is expected to improve the Quality of Experience (QoE) of participants in a collaboration session. Thus, NML designed and implemented an architecture for high-quality video services for multi-party collaborative environments, such as the AG. This attempt to enable high-quality media services for the AG began in 2003. In 2004, the development of DV/HDV support for the AG was released for public use (<http://ace.nm.gist.ac.kr>). NML designed extensions for the AGTk (AG toolkit) video services, so called “ExtendedVideo{Producer, Consumer}Services” by considering such issues as flexible multicast address allocation, versatile video codec/application support, and negotiation of service capability. However, even though the AGTk augmentation enabled DV/HDV support, two key missing capabilities prompted NML to develop a new architecture, named AG Media/SVC (Scalable Visualization Consumer) [12, 13]. The first capability that the new architecture addresses is the automatic negotiation of common streaming media formats (DV/HDV/stereo HD/uncompressed HD) between remote sites that are participating in a collaborative session. The second much needed capability is a means to decouple CPU-intensive video decoding from image display so that the video can be easily manipulated when presented on displays of arbitrary resolution. During the summer of 2005, NML developed a version of SVC that used EVL's SAGE to distribute high-definition video to EVL's LambdaVision display. This was successfully demonstrated at iGrid 2005 by streaming Mpeg-2 (1440 × 1080i) video from Chicago to San Diego over CAVEWave at 25 Mb/s and decoding it into a 1 Gb/s stream for display on the LambdaVision.

2.4. BPlayer—high-definition stereoscopic graphics streaming

Donna Cox (cox@ncsa.uiuc.edu), Robert Patterson, Patrick Dorn, Paul Wefel, Stuart Levy—National Center for Supercomputing Applications.

Research projects at the University of Illinois in Urbana and Chicago have been brought together to accomplish the goal of streaming high-definition stereo movies long distances across optical networks. At iGrid 2005, visualizations of several scientific simulations were streamed from the National Center for Supercomputing Research in Urbana, Illinois to the iGrid facility in San Diego, California. Stereoscopic 3D, time-dependent, data were rendered to individual images at high-definition resolution. These images were then played back in sequence and streamed across a high-speed optical network to a stereoscopic HD projection system at iGrid. NCSA's BPlayer image playback software was augmented with EVL's SAGE to allow it to stream the playback over high-speed networks.

BPlayer is an application for displaying movies stored on disk as sequences of images. BPlayer software is used at NCSA to display HD content on a passive stereo display system driven by a single dual processor computer with a high performance hardware raid. Performance of this movie playing techniques depends on the ability to read image data from disk at a high rate. BPlayer uses a multi-threaded staggered sequential disk access scheme in order to realize the full bandwidth of the disk array. The integration of SAGE with BPlayer enabled it to stream images across the network to a remote display system. This combination of software was able to support the display of HD movies at various resolutions. The iGrid demonstration achieved approximately 14 frames/s display rate and used approximately 1.4 Gb/s. While not entirely satisfactory the frame rate was adequate given the distances involved. Further work is necessary to fully identify bottlenecks and address them. By inserting a microsecond sleep between SAGE's reception of multiple packets in the network receiving thread, we were able to increase the dynamic priority of the network receiving thread which is computed by the kernel based on average sleep time of a process [14]. As there were few competing processes running on the system, this priority increase led to an improved CPU share of the receiving thread to process the socket buffers and receive additional data which led to a lower packet loss that typically occurs due to socket buffers getting congested.

2.5. Auto-stereoscopic video conferencing

Dan Sandin (dan@uic.edu), Tom Peterka, Javier Girado, Robert Kooima, Jinhua Ge—Electronic Visualization Laboratory, University of Illinois at Chicago.

Prior teleconferencing experiences have largely been a 2D experience, but this changed at iGrid 2005, as conference attendees sat and discussed the benefits of *auto-stereoscopic* teleconferencing “face-to-face” in real-time 3D with scientists at EVL. This essentially meant that viewers could see stereoscopic 3D video of a remote participant without the need to wear special glasses. The display system consisted of an EVL-designed device called the Varrier auto-stereo display [10] driven by a small cluster of computers to track the viewer and render a proper view. Initial feedback from iGrid was positive and participants certainly appeared to enjoy the interaction.

The CAVEWave 10 Gb optical networking between UCSD and UIC facilitated audio and stereo video streaming in real-time, using approximately 220 Mb/s of bandwidth in each direction. End to end latency was approximately 0.3 s (one way), and auto-stereo images of participants were rendered at 25 Hz. Raw, unformatted YUV411 encoded images were transmitted, using one half the bandwidth as compared to a corresponding RGB format. The real-time display for the most part appeared flawless even though approximately 10% of packets were dropped. The robustness is due to the high transmission and rendering rates. The sending rate was 30 Hz, while the rendering rate was approximately 25 Hz, limited mainly by pixel fill rates and CPU loading. The net result is that the occasional lost packets were negligible. The excellent network quality of service can be attributed to the limited number of hops and routers between source and destination. There were exactly 4 hosts along the route, including source and destination; so only two intermediate routers were encountered. The success of the teleconferencing is directly attributable to the 10 Gb/s CAVEWave lightpath over the National LambdaRail, as the application simply sends raw stereo video and audio data and relies on the network to successfully deliver quality of service.

2.6. The Solutions Server

Brian Corrie (bcorrie@sfu.ca), Todd Zimmerman—Simon Fraser University; Pierre Boulanger, Curtis Badke, Jeff Ryan, Manuel Garcia—University of Alberta.

The Media LightPath (MLP) project is a CANARIE funded project that is exploring the use of user-controlled lightpaths (UCLP) [11] for advanced collaboration and visualization. One of the goals of the MLP project is to provide a UCLP-based collaboration and visualization capability to collaborating institutions that have lightpath capability. As a first step towards this end, the MLP project participated in the iGrid 2005 GLVF demonstration. Our demonstration consisted of a computational simulation (running at the University of Alberta in Edmonton) streaming data to a remote visualization resource (located at Simon Fraser University in Vancouver) over the gigabit WestGrid network. The simulation data was then rendered as a stereo pair of images, sent over a dedicated gigabit network to Vancouver to San Diego, and displayed on a passive stereo display at the iGrid venue. We utilized the University of Alberta's Solution Server software to stream computational data to the visualization machine and SGI's VizServer software to stream the rendered images to San Diego. The main result that MLP researchers took away from their iGrid experiences was the large impact that distance can have on software performance. Before iGrid, MLP had conducted some limited testing on a temporary lightpath between Vancouver and Edmonton as well as using the shared WestGrid lightpath. Application level throughput over the local campus backbone exceeded 600 Mb/s. Experiments over a test lightpath and the WestGrid network decreased application-level performance to approximately 200 Mb/s. When the same experiment was extended to iGrid in San Diego, that

performance further decreased to a mere 20 Mb/s. Although a very disappointing result, this provided a valuable data point for our research. It appeared that the fundamental problem lay in the inadequate tuning of the end systems and that this tuning was required at a variety of layers (application, client hardware, and networking). It is clear that more research is required to develop application level solutions that can adapt to these emerging advanced networking environments. It is also clear that such research is difficult to achieve in an environment where the networks and end-points exist only for short periods of time or for experimental events such as iGrid.

3. Future challenges

iGrid 2005 proved to be a successful launching point for the Global Lambda Visualization Facility. The collaborating researchers found it highly informative to be able to evaluate their own techniques and to compare them against alternative approaches developed by their colleagues (these are explained in greater detail in separate articles in this journal issue). In the next few years the goal is to ensure the permanence of this one-of-a-kind testbed so that the level of cooperation that was initiated by iGrid can be further expanded. We envision international linked heterogeneous display-rich environments that will engage in both short-lived and long-lived, ad hoc and pre-scheduled, interactive collaborative scenarios that can support comprehensive and synergistic visualization/collaboration research, as well as applications in Science, Engineering, Medicine, Education and Art. An immediate challenge that the group will collectively address is a means to provide a common or compatible high-definition video and ultra-high-resolution visualization multicasting scheme between the sites. Some of this work can leverage existing AccessGrid [1] techniques, however the rendering of this information within heterogeneous room-filled display environments poses new challenges in terms of the innovative user-interaction methods that are needed to enable convenient use, and multicasting of the variety of data streams. Multicasting in all its forms (unreliable, reliable, and encrypted) at the magnitude of tens of gigabits is challenging both in terms of networking and the end-systems that must assimilate the multicast streams. Although a variety of techniques exist for supporting reliable multicast [2], *high-bandwidth (on the order of tens of gigabits/s) and low-latency, reliable* multicast is an unsolved problem and an active area of research within the Grid community [3]. It is a particularly challenging problem at the endpoints that must receive the multicast traffic; and must be solved for any applications that seek to use networked tiled displays. In the SAGE model, the rendering source is decoupled from the display and fed by a high-speed network. The bandwidth utilized is dependent on the image resolution of the rendering source. As windows on the tiled display are resized or repositioned on the walls, SAGE must re-route the multiple streams from the rendering source to the PC nodes that drive the displays. When the windowed area is smaller than the original source of the image, the source must be rescaled

so that the combined network traffic does not overwhelm the display node (limited to 1 Gb/s). This problem becomes much more complex when multicasting is needed to support independent window operations (such as position or scaling) at each display site with different configurations. A number of solutions will be examined with various trade-offs: 1. utilize a layered multicast approach to interleave each source stream over multiple multicast streams and coordinate the receivers so that each display node knows which layered stream to subscribe to in order to assemble the images for the desired screen resolution; 2. use high-performance PCs equipped with 10 Gb/s interfaces to intercept each multicast stream and perform resizing and re-routing operations on behalf of the thin-client displays.

Realising a persistent, robust GLVF provides us the opportunity to conduct thorough testing of these concepts on real network testbeds at a variety of distances, and enables us to guarantee that the developed solutions are compatible internationally.

Acknowledgments

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago specializes in the design and development of high-resolution visualization and virtual-reality display systems, collaboration software for use on multi-gigabit networks, and advanced networking infrastructure. These projects are made possible by major funding from the National Science Foundation (NSF), awards CNS-0115809, CNS-0224306, CNS-0420477, SCI-9980480, SCI-0229642, SCI-9730202, SCI-0123399, ANI-0129527 and EAR-0218918, as well as the NSF Information Technology Research (ITR) cooperative agreement (SCI-0225642) to the University of California San Diego (UCSD) for "The OptIPuter" and the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement (SCI-9619019) to the National Computational Science Alliance. EVL also receives funding from the State of Illinois, General Motors Research, the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC), and Pacific Interface Inc. on behalf of NTT Optical Network Systems Laboratory in Japan. The GeoWall, GeoWall2, Personal GeoWall2 (PG2), and LambdaVision are trademarks of the Board of Trustees of the University of Illinois.

SARA's TOPS research is funded in part by the GigaPort NG project. SARA is The Netherlands' High Performance Computing and Networking Center. It operates the SURFnet NOC and houses SURFnet5, SURFnet6 and NetherLight cores and The Netherlands' national supercomputers.

Simon Fraser University and the University of Alberta's MediaLight Path and Solution Servers projects are supported by CANARIE, SGI, CFI, ASRA, and BCKDF.

GIST (Gwangju Institute of Science and Technology) is supported through the BK-21 program, and EVL at UIC.

KISTI is supported through Korea's national e-Science program.

References

- [1] The AccessGrid—<http://www.accessgrid.org>.
- [2] M.P. Barcellos, M. Nekovec, M. Koyabe, M. Dawe, J. Brooke, High-throughput reliable multicasting for grid applications, in: Fifth IEEE/ACM International Workshop on Grid Computing, Grid'04, 2004, pp. 342–349.
- [3] M. den Burger, T. Kielmann, H.E. Bal, Balanced multicasting: high-throughput communication for grid applications, SC'05, Seattle, WA, 12–18 November 2005.
- [4] G. Stix, The Triumph of the Light, Scientific American, January 2001.
- [5] L. Renambot, A. Johnson, J. Leigh, LambdaVision, Building a 100 Megapixel Display, in: NSF CISE/CNS Infrastructure Experience Workshop, Champaign, IL, July 27, 2005.
- [6] L. Smarr, A.A. Chien, T. DeFanti, J. Leigh, P.M. Papadopoulos, The OptIPuter, Communications of the ACM 46 (11) (2003) 58–67.
- [7] J. Leigh, L. Renambot, T.A. DeFanti, et al., An experimental optiputer architecture for data-intensive collaborative visualization, in: 3rd Workshop on Advanced Collaborative Environments, Seattle, WA, June 2003.
- [8] W. Blanke, C. Bajaj, D. Fussell, X. Zhang, The metabuffer: a scalable multiresolution multidisplay 3-D graphics system using commodity rendering engines, Tr2000-16, University of Texas at Austin, February 2000.
- [9] L. Renambot, A. Rao, R. Singh, B. Jeong, N. Krishnaprasad, V. Vishwanath, V. Chandrasekhar, N. Schwarz, A. Spale, C. Zhang, G. Goldman, J. Leigh, A. Johnson, SAGE: the Scalable Adaptive Graphics Environment, in: Proceedings of WACE 2004, Nice, France, 09/23/2004–09/24/2004.
- [10] D. Sandin, T. Margolis, J. Ge, J. Girado, T. Peterka, T. DeFanti, The varrier autostereoscopic virtual reality display, in: Proceedings of ACM SIGGRAPH 2005, ACM Transactions on Graphics 24 (3) (2005) 894–903.
- [11] J. Wu, S. Campbell, J.M. Savoie, H. Zhang, G.v. Bochmann, B. St. Arnaud, User-managed end-to-end lightpath provisioning over CA*net 4, in: Proceedings of the National Fiber Optic Engineers Conference, NFOEC, Orlando, FL, USA, September 7–11, 2003, pp. 275–282.
- [12] S. Han, J.W. Kim, Adaptive high-quality video service for network-based multi-party collaboration, in: Proc. of SPIE ITCOM'2005: Multimedia Systems and Applications VIII, Boston, October 2005.
- [13] J. Kim, J.W. Kim, Decomposable decoding and display structure for scalable media visualization over advanced collaborative environments, in: Proc. of SPIE ITCOM'2005: Multimedia Systems and Applications VIII, Boston, October 2005.
- [14] Linux Scheduler 2.6, http://josh.trancesoftware.com/linux/linux_cpu_scheduler.pdf, 2005.



Jason Leigh is an Associate Professor of Computer Science and co-director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). His areas of interest include: developing techniques for interactive, remote visualization of massive data sets; and for supporting long-term collaborative work in amplified collaboration environments. Leigh has led EVL's Tele-Immersion research agenda since 1995 after developing the first networked CAVE application in 1992. He is co-founder of the GeoWall Consortium and visualization lead in the National Science Foundation's OptIPuter project.




Luc Renambot received a Ph.D. at the University of Rennes-1 (France) in 2000, conducting research on parallel rendering algorithms for illumination simulation. Then holding a Postdoctoral position at the Free University of Amsterdam, till 2002, he worked on bringing education and scientific visualization to virtual reality environments. Since 2003, he joined EVL/UIC as a PostDoc and now as Research Assistant Professor, where his research topics include high-resolution displays, computer graphics, parallel computing, and high-speed networking.



Andrew Johnson is an Associate Professor in the Department of Computer Science and member of the Electronic Visualization Laboratory at the University of Illinois at Chicago. His research interests focus on interaction and collaboration in advanced visualization environments.



Byungil Jeong received the B.S. and M.S. degrees in electrical engineering in 1997 and 1999 from the Seoul National University, South Korea. He worked as a researcher at Imaging and Media Research Laboratory, Korea Institute of Science and Technology until 2002. He is a Ph.D. candidate in computer science and working as a research assistant at Electronic Visualization Laboratory, University of Illinois at Chicago. His research interests include scalable graphics architecture, high performance graphics streaming and tiled high-resolution displays.



Nicholas Schwarz is a graduate student at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). His research interests include the visualization of very large multi-dimensional data for scalable high-resolution displays and virtual-reality environments. He works extensively with domain experts in the bio-science and geo-science communities.



Paul Wielinga is manager of High Performance Networking business unit at SARA Computing and Networking Services, the e-Science Support Center of The Netherlands. He has a Masters Degree in Physics from the University of Amsterdam and has over 30 years of IT experience, in the fields of Operating Systems, High Performance Networking and advanced Scientific Visualization. Since 1996 he has been involved in the developments in the field of Visualization and Virtual Reality for scientific and commercial applications. He was manager of the Visualization group at SARA from 1992 till beginning of 2005. He is main contact person at SARA for the international OptIPuter collaboration. Wielinga is also coordinator for Research and Development at SARA.



Bram Stolk is a Virtual Reality specialist with SARA. He received his M.Sc. in Computer Science from the University of Amsterdam. He is the principal author of the visual datamining tool for genomics, called SARAgene. Bram currently researches high resolution visualization technology over high speed networks. His interests include Self Organizing Feature Maps, Computer Graphics and Virtual Reality.



Gee Bum Koo is leader of the visualization group at KISTI Supercomputing Center. His research interests include volume visualization, visualization service in grid environment, and high-end visualization systems. He has a B.A. and M.S. in computer science from SoGang university, Korea. He is one of the steering committee members of the national e-Science program in Korea.



Dr. JongWon Kim received the B.S., M.S. and Ph.D. degrees from Seoul National University, Seoul, Korea, in 1987, 1989 and 1994, respectively, all in control and instrumentation engineering. After working at Kongju National University, Kongju, Korea, (94–99) and University of Southern California, Los Angeles, USA (97–01), he joined as an Associate Prof. at Gwangju Institute of Science & Technology (GIST), Gwangju, Korea. Under the slogan of “Reliable and flexible delivery system for integrated multimedia over wired/wireless IP networks”, he is focusing on networked media systems and protocols. He has co-authored around 200 technical publications in international/domestic journals and conferences. Dr. Kim is a senior member of IEEE, and a member of ACM, SPIE, KICS, IEEK, and KISS.



Brian Corrie is the Collaboration and Visualization Coordinator for the WestGrid Collaboration and Visualization research program and SFU’s Centre for Interdisciplinary Research in Math and Advanced Computational Science. Brian’s research interests are in advanced collaborative environments, computer graphics, scientific visualization, virtual environments (VE), and coupling computational simulation to visualization. Over the last ten years Brian has been involved in the establishment of three virtual environment facilities including the Virtual Environments lab at the Australian National University (ANU) in Canberra, Australia, the Virtual Environments Technology Centre (VETC)

at the National Research Council of Canada’s Integrated Manufacturing Technologies Institute in London, Ontario, and the Immersive Media Lab at the New Media Innovation Centre (NewMIC) in Vancouver, British Columbia. He was the Project Leader for the Collaborative VE project at ANU, the Technical Leader at the VETC, and the Focus Area Leader for the Immersive and Collaborative Environments research program at NewMIC. Brian joined WestGrid and IRMACS in 2004 to coordinate and manage their collaborative and visualization research programs and facilities.



Dr. Pierre Boulanger graduated from Laval University in Engineering Physics. He also received his Masters in Physics from the same university and his Ph.D. in Electrical Engineering from the University of Montreal. He worked for 18 years at the National Research Council of Canada as a senior research officer where his primary research interest were in 3D computer vision, rapid product development, and virtualized reality systems. Since July 1st 2001, he is working as a professor at the Department of Computing Science at the University of Alberta doing research and teaching on virtualized reality systems. He is also an adjunct scientist and principal investigator for new media at TRILabs and at the Banff Centre. In 2004, Dr. Boulanger was awarded an iCORE industrial chair in Collaborative Virtual Environment. He has published more than 150 scientific papers in various journals and conferences. Dr. Boulanger is also on many international committees and frequently gives lectures on rapid product development and virtualized reality.