

Available online at www.sciencedirect.com



Future Generation Computer Systems 22 (2006) 1032-1039



Real-time multi-scale brain data acquisition, assembly, and analysis using an end-to-end OptIPuter

Rajvikram Singh^{b,*}, Nicholas Schwarz^b, Nut Taesombut^c, David Lee^a, Byungil Jeong^b, Luc Renambot^b, Abel W. Lin^a, Ruth West^a, Hiromu Otsuka^e, Sei Naito^f, Steven T. Peltier^a, Maryann E. Martone^a, Kazunori Nozaki^d, Jason Leigh^b, Mark H. Ellisman^a

^a National Center for Microscopy and Imaging Research, University of California, San Diego, United States
^b Electronic Visualization Laboratory, University of Illinois at Chicago, United States
^c Department of Computer Science and Engineering, University of California, San Diego, United States
^d Cybermedia Center, Osaka University, Japan
^e KDDI Corporation, Garden Air Tower, 10-10, Iidabashi 3-chome, Chiyoda-ku, Tokyo 102-8460, Japan
^f KDDI Labs, 2-1-15 Ohara, Fujimino, Saitama, 356-8502, Japan

Available online 9 May 2006

Abstract

At iGrid 2005 we demonstrated the transparent operation of a biology experiment on a test-bed of globally distributed visualization, storage, computational, and network resources. These resources were bundled into a unified platform by utilizing dynamic lambda allocation, high bandwidth protocols for optical networks, a Distributed Virtual Computer (DVC) [N. Taesombut, A. Chien, Distributed Virtual Computer (DVC): Simplifying the development of high performance grid applications, in: Proceedings of the Workshop on Grids and Advanced Networks, GAN 04, Chicago, IL, April 2004 (held in conjunction with the IEEE Cluster Computing and the Grid (CCGrid2004) Conference)], and applications running over the Scalable Adaptive Graphics Environment (SAGE) [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, 23-24 September 2004, Nice, France, 2004]. Using these layered technologies we ran a multi-scale correlated microscopy experiment [M.E. Maryann, T.J. Deerinck, N. Yamada, E. Bushong, H. Ellisman Mark, Correlated 3D light and electron microscopy: Use of high voltage electron microscopy and electron tomography for imaging large biological structures, Journal of Histotechnology 23 (3) (2000) 261–270], where biologists imaged samples with scales ranging from 20X to 5000X in progressively increasing magnification. This allows the scientists to zoom in from entire complex systems such as a rat cerebellum to individual spiny dendrites. The images used spanned multiple modalities of imaging and specimen preparation, thus providing context at every level and allowing the scientists to better understand the biological structures. This demonstration attempts to define an infrastructure based on OptIPuter components which would aid the development and design of collaborative scientific experiments, applications and test-beds and allow the biologists to effectively use the high resolution real estate of tiled displays. © 2006 Published by Elsevier B.V.

Keywords: Graphics clusters; Multi-scale correlated microscopy; Montage images; HDTV streaming; Telescience; Tile displays; Tiled displays; Optical networks; Scientific visualization; Remote instrumentation

1. Introduction

1.1. The challenge

Researchers interested in assessing brain tissue at multiple resolutions are faced with a well-known problem when

E-mail address: rsingh@ncmir.ucsd.edu (R. Singh).

traversing scales: as investigations increase in resolution they typically decrease in scope. This gap between dimensional scales makes it difficult to understand how higher order structures are constructed from finer building blocks. A particular challenge for the nervous system is the need to bridge the dimensional range of 100s of microns to nanometers. This range is called "mesoscale" and encompasses cellular networks, dendritic and axonal architectures, synaptic connectivity and macromolecular constituents. These structures represent the

^{*} Corresponding author. Tel.: +1 312 404 4058.

⁰¹⁶⁷⁻⁷³⁹X/\$ - see front matter © 2006 Published by Elsevier B.V. doi:10.1016/j.future.2006.03.017

heart of information processing in the nervous system and are central to our understanding of the brain.

The mesoscale gap arises in part from the requirement to use multiple imaging technologies to examine a specimen across scales. Each technology requires different expertise, specimen preparation techniques and contrast mechanisms, and also requires a severe reduction in the amount of tissue. For example, if the pipeline begins with an entire brain, the end results in one small block of tissue, <0.5 mm³. These requirements make it difficult for individual researchers to bridge scales, both because single researchers may not be familiar with a given technology and because there is significant loss of context as the scope decreases with increasing resolution of imaging technologies. Bridging techniques such as multiphoton microscopy and electron tomography, correlated microscopy is a key methodology for acquiring the necessary multi-scale data in order to fill in the resolution gaps between gross structural imaging and protein structure: data which is central to bridging the mesoscale gap and to the elucidation of the nervous system [20].

At iGrid 2005 we demonstrated the integration of highresolution tiled displays, Telescience [8,15], HDTV video streams, OptIPuter [7] system software and computational resources such as graphics clusters to extend the capabilities of bioscience instrumentation and informatics. Our purpose was to demonstrate how this integration enables researchers to bring multiple views of a specimen together at the same time without sacrificing resolution or context. This was done along with HDTV video streams from remote instruments and from collaborators at other sites. The experiment brought together multiple views of very large biological datasets acquired using multiple microscopes and multiple visualization modalities. Researchers could collaboratively view 2D scenes and 3D and 4D subsections of a scene while comparing them to dozens of possible contexts and matching these to live HDTV video output of an UHVEM (Ultra-High Voltage Electron Microscope). We are hopeful that this will ultimately result in the ability to actively guide the examination of a specimen collaboratively while remotely controlling rare instruments such as the 3 MeV UHVEM in Osaka Japan. It will provide contextual reference for the specimen under investigation, which is otherwise lost due to the use of multiple microscopies and differences in scale and resolution.

1.2. OptIPuter

The OptIPuter [5] is a National Science Foundation project that aims at developing an advanced distributed computing infrastructure for collaborative data exploration in the fields of Neuroscience and Geosciences. Computer scientists at the University of Illinois at Chicago and the University of California, San Diego are jointly working with the Biomedical Informatics Research Network, US Geological Survey's Earth Resource Observation System and the Scripps Institute of Oceanography for this project. For visualizing large geological or biomedical datasets one typically requires cluster driven tiled displays such as the one shown in Fig. 1. Remote datasets spanning over multiple terabytes are accessed and visualized on the high-resolution displays. The network bandwidth requirements for browsing these data stores or pushing the rendered pixels to remote tiled displays are in the range of several tens of gigabits per second.

In the past with traditional display and networking technologies, it has been virtually impossible to conduct experiments allowing correlation of multi-modality data from various sources and live instruments at a site. A typical HDTV video stream from a microscope requires 1 Gbps of bandwidth for streaming at 30 frames per second to a remote site. A montage dataset created by a light microscope spans multiple gigabytes and has typical pixel dimensions of several tens of thousands of pixels. Real-time control of remote instruments requires dedicated networks with guarantees on the quality of service to allow accurate control with minimal latency. The OptIPuter provides an effective framework for bringing together remote instrumentation, storage and computational resources over high-bandwidth optical networks to improve the speed and accuracy of scientific collaboration and experimentation.

1.3. Layers in the OptIPuter

The OptIPuter framework is classified into layers or software abstractions. The actual applications used by scientists are supported by several underlying layers which manage the distributed resources. These components are described in the following paragraphs.

1.3.1. Lightpath provisioning using Photonic Inter-domain Negotiator (PIN)

A key feature of the OptIPuter architecture is the provisioning of dedicated lambdas between remote resources using optical switches. An application wishing to allocate a lightpath between two end points contacts its local PIN [16] service, which dispatches generic lightpath signaling messages to neighboring PINs until the final destination is reached. Each PIN translates the generic lightpath signaling message into a native photonic signaling message that is understood by the local intra-domain lightpath signaling facility. This facility then signals the photonic switch to make adjustments to its internal MEMS (Micro-Electro-Mechanical Systems) switches to establish a connection.

1.3.2. Distributed Virtual Computer (DVC)

The DVC [3] enables simple application construction and high-performance execution by dynamically allocating globally distributed resources into a single virtual computer. It provides a set of abstractions that shield applications from the complexities of underlying software and hardware infrastructures, integrating them in a way to enable a simple resource use and performance models. Specifically, applications describe and acquire a set of distributed resources and dynamically configured optical networks through a scheduling agent, which transparently realizes their needs. The allocated resources are bound into a single domain and transparently managed by the middleware for security, highperformance, reliable communication and other forms of agreed service quality. The applications make use of these resources as a private resource context, using them to achieve good performance and reliable execution.

1.3.3. LambdaRAM

LambdaRAM [14] (Optically Connected Memory Cache) is a middleware designed to address long-haul latency in optical networks. LambdaRAM is a tool that aggregates pools of memory in clusters of gigabit-connected computers to provide a massive data cache. This minimizes access latency when data-intensive applications need to fetch data from distantly located data stores. It utilizes available network bandwidth to aggressively pre-fetch data before an application needs it. Applications accessing LambdaRAM see the distributed cache as a contiguous memory image. LambdaRAM is currently used in JuxtaView and TeraScope [6] (EVL's visual data mining software) to provide data correlation algorithms with fast access to distributed database tables.

1.3.4. LambdaStream

LambdaStream [12] is a transport protocol designed specifically to support gigabit-level streaming, which is required by streaming applications over OptIPuter. The protocol takes advantage of characteristics in photonic networks. It adapts the sending rate to dynamic network conditions while maintaining a constant sending rate whenever possible. One advantage of this scheme is that the protocol avoids deliberately provoking packet loss when probing for available bandwidth, a common strategy used by other congestion control schemes. Another advantage is that it significantly decreases fluctuations in the sending rate. As a result, streaming applications experience small jitter and react smoothly to congestion.

1.3.5. Group Transport Protocol (GTP)

The Group Transport Protocol (GTP) [4] is a highperformance data transport protocol, targeting efficient and fair sharing of source and sink capacities among active connections in high-speed long-distance network environments. GTP features distributed end-node based max-min fair rate allocation across multiple flows to avoid network resource contention at end systems and to support multipoint-to-point and multipoint-to-multipoint data movement. Experiments show that GTP performs as well as other UDP based aggressive transport protocols for single flows, and when converging flows (from multiple senders to one or multiple receivers) are introduced, GTP achieves both high throughput and much lower loss rates than others [19].

1.3.6. Scalable Adaptive Graphics Environment (SAGE)

The majority of the data visualization and collaboration problems in OptIPuter are being solved with SAGE [1]. SAGE, shown running on a tiled display in Fig. 1, provides a "window manager" and a dynamic pixel routing infrastructure for cluster driven tiled displays to allow several remote sources



Fig. 1. Biologists use a SAGE display for visualizing multiple datasets while collaborating with remote sites using HDTV video streams.



Fig. 2. A typical scientific visualization pipeline is made up of Storage \rightarrow Computation \rightarrow Visualization \rightarrow Display.

to simultaneously display their frame buffers on the tiled display. Thus it treats the entire tiled display as one large seamless desktop spreading over potentially unlimited pixels. Users can interact with the "windows" of a remote rendering machine and move or resize them on this desktop using intuitive user interfaces.

1.3.7. Application layer

The applications make use of all the underlying layers for allocating resources for achieving high-performance. The applications can run on individual machines or clusters, remotely or locally. JuxtaView [9], TeraVision [2], Magic Carpet [11], Vol-a-Tile [10] and HD movie streaming are some of the applications designed to run over the OptIPuter. As shown in Fig. 2, the typical scientific visualization pipeline consists of four stages. SAGE takes care of the display stage. LambdaRAM helps with computation and storage. GTP and LambdaStream bridge the various stages efficiently.

2. Microscopy experiments over OptIPuter

The OptIPuter model enables biologists to improve the speed and throughput of data production in correlated microscopy experiments, by providing infrastructure to allow existing data to be used for steering and/or refining real-time data acquisition. Tile displays running SAGE provide the screen real-estate necessary to simultaneously view very large datasets, from multiple imaging modalities (2D mosaics, 3D volumes, 4D time-series volumes, streaming media) adjacent to each other, while preserving resolution. The ability of the integrated OptIPuter model to manage dedicated links between remote resources allows these datasets to be sourced from remote resources including live instrument feeds. It is now possible to monitor and soon even guide a remote microscope experiment scanning cell structures while simultaneously referencing



Fig. 3. The various OptIPuter components required for running a distributed multi-scale correlated microscopy experiment.

multiple pre-collected datasets of varying modalities and magnification.

As demonstrated at iGrid06, this model is extremely useful for improving the remote electron microscopic (EM) tomography component of correlated microscopy experiments. In the latter, EM tomography data is acquired and registered within a wide-field scene previously acquired of the same sample via light microscopy (LM). The goal is to produce a multi-scale geometric model that provides detailed information within a wide-field context. This process is anything but straightforward. Samples that produce high contrast information in the light microscope are often lacking in contrast in the electron microscope, especially at the specimen thicknesses required for tomography. The structures of interest are also often difficult to distinguish at the EM level as detail will be visible that was otherwise obscured or impossible to resolve with the LM. The process to section and prepare the sample for EM can also further complicate the ability to keep track of single regions of interest (ROI) within fields of similarly labeled structures. As such, remote researchers often spend a tremendous amount of time scanning the sample, looking for these specific ROIs, in a process akin to searching for a needle in a haystack by looking through a soda straw. Under the OptIPuter model, many instances of data acquired in the life cycle of the sample can be presented to the user at the time of EM data acquisition to help with specimen navigation. Since the EM itself is a radiation source that is damaging the sample even as it is being imaged, there is a great advantage in accelerating this process. The ability of these data to inform this process is greatly amplified by the ability to present that data at or near native resolution without down-sampling. Also, the ability to stream HDTV video from the instrument and between participating collaborators with low latency further enhances the quality of the remote experiment and control.

A typical "small" microscopy dataset is approximately 20,000 by 20,000 pixels in size and spans multiple gigabytes and multiple files. Large datasets can span multiple terabytes of disk space. Applications such as JuxtaView [9] running over distributed shared memory middleware such as LambdaRAM allow intelligent access to these remote datasets with pre-fetching and caching capabilities. This most ostensibly enhances the user experience for browsing the datasets. LambdaStream and GTP provide efficient transport mechanisms at the lower layers of the middleware for streaming data between clusters. Future versions of these protocols aim to guarantee network performance parameters, viz. jitter and bandwidth, which are critical for remote-controlling instruments. DVC binds together all these layers by providing an efficient and secure resource allocation model for distributed OptIPuter resources.

Fig. 3 shows a diagrammatic representation of an OptIPuter setup for a distributed microscopy experiment.

3. iGrid demonstration

At iGrid we employed the various components of the OptIPuter model for successfully conducting a distributed multi-scale microscopy experiment. We used the 55-panel, dual-Opteron driven tiled display at the Calit2 [13] (California Institute for Telecommunications and Information Technology) building at the University of California, San Diego. The collective resolution of the display is 100 megapixels (17,600 \times 6000 pixels). The display can potentially sink 30 Gbps of video and data streams coming over the CAVEWave optical



Fig. 4. A LambdaCam [17] snapshot of the 100 megapixel tiled display at iGrid 2005 shows SAGE receiving video outputs from a remote microscope, an HDTV camera and a pre-recorded streamed video from Osaka, Japan. The picture also shows JuxtaView [9] showing datasets collected from a light microscope and an electron microscope. The tiled display was viewable over the web in real-time using the LambdaCam utility. The collective network bandwidth of all the applications, as reported by SAGE, was >4 Gbps.

network thus making it ideal for our correlated microscopy demonstration. We also used the 20-tile, 40 megapixel, dual-Opteron cluster at NCMIR for simulating a remote collaboration room. As seen in Fig. 4 multiple instances of JuxtaView [9] displayed montage and Cell EM datasets stored remotely on clusters at the Electronic Visualization Laboratory, Chicago and SARA, Amsterdam. Pre-recorded HD movies of cell division were also streamed from SARA. The datasets spanned several gigabytes and were accessed at interactive speeds as users zoomed/panned though them. JuxtaView employs LambdaRAM [14] as a middleware for distributing and intelligently fetching datasets from several remote machines. The largest dataset used was $17,000 \times$ 17,000 pixels in size and spanned across 2 GB of disk files. There was also an instance of Vol-a-Tile rendering a computed tomography dataset stored locally.

The video output of a light microscope at the National Center for Microscopy and Imaging Research (NCMIR) was fed to a TeraVision [2] box. The 1024×1024 video was then streamed to the SAGE display using TCP at 20 fps generating 500 Mbps of network traffic through the TeraVision box. The audience at the conference could witness the microscope scanning a biological sample in real-time. The montage datasets shown by JuxtaView were the final output of the scanned images.

Pre-recorded video of an experiment on the Hitachi H-3000 UHVEM microscope at Osaka University, Japan was streamed using KDDI Lab's JPEG2000 HDTV encoder [18]. The codec allows for real-time streaming of compressed HDTV video using specialized hardware and provides exceptionally low latency which is key for remote instrumentation control. The rate of the network streams can be controlled based on the quality of video desired. The output of the codec at the iGrid end was fed to a TeraVision box for the purpose of putting it up on the SAGE display.

An HDTV camera and a TeraVision box in a conference room at NCMIR transmitted uncompressed video at 1920 \times 1080 at 20 fps (1 Gbps). The audience could see scientists at NCMIR manipulating datasets on their 5 \times 4 tiled display. The goal is to allow scientific groups to be able to video-conference and share experimental results easily to enhance collaborative analysis. The collective bandwidth of all applications, as reported by SAGE, was greater than 4 Gbps.

Juxtaposition of collaborative tools such as HDTV video streaming alongside real-time instrument feeds and collaborative visualization of multi-scale, multi-modality datasets provides a dynamic context which directly helps active investigation of specimens and data analysis. We are certain that this approach will not only enhance collaboration amongst researchers throughout scientific workflows but also serve to increase the effectiveness and efficiency of correlated microscopy.

4. Conclusion and future work

In the past with traditional display and networking technologies, it has been virtually impossible to conduct collaborative microscopy experiments allowing correlation of multi-scale, multi-modal datasets and outputs of remote instruments. Such an experimental setup would help alleviate the mesoscale gap problem. We successfully conducted such an experiment at iGrid 2005 where video from remote microscopes, remote HDTV cameras, renderings of volumes (3D) and image (2D) datasets from local and remote data stores were all brought together on a SAGE display. The OptIPuter model presents an efficient and data friendly framework to allow biologists to conduct a variety of experiments in a shared virtual space. The high bandwidth networks and high resolution tiled displays meet the demanding requirements of remote instrumentation and the sheer number of pixels needed for experiments such as correlated microscopy.

For future work we would like to integrate the currently separate remote instrumentation control, within the OptIPuter environment. We also would like to integrate live instrument feeds more closely with cell structure databases to allow for automatic detection of biological structures during an experiment. Tools to help the scientists guide live experiments based on references extracted from pre-recorded datasets is also a priority for us. Work is currently underway to build userinterface devices to allow biologists to modify or annotate the datasets and save them back into data stores.

Acknowledgments

We would like to thank Cees de Laat at University of Amsterdam for providing the endpoints at SARA in Amsterdam to perform our demonstration. Alan Verlo, Lance Long and Tom Hutton provided us with constant networking and system support at iGrid. Chris Airriess and Ryan Biltz supported clusters at the NCMIR end. We are grateful for the demonstration time support by Hiroyuki Hakozaki and James Bouwer. They were responsible for running the instruments and tiled-display software at NCMIR during iGrid. Diana Price and Van Phung provided their pre-recorded light-microscope datasets used for the demo and this paper.

We would especially like to thank our collaborators Hirotaro Mori, Shinji Shimojo and Toyokazu Akiyama from Osaka University, Japan without whose help, the streaming of the recorded electron microscope experiment would not have been possible. Our collaborators from KDDI R&D Laboratory, Sei Naito and Atushi Koike, provided us with the equipment and valuable support in conducting our video streaming from Japan.

We are also grateful to Fang-Pang Lin, Sun-In Lin, and Shi-Wei Lo at the National Center for High Performance Computing, Taiwan for their continuous support in helping test the OptIPuter technologies.

These projects are made possible by major funding from NIH/NCRR for the National Center for Microscopy and Imaging Research (P41RR004050) NIH/NCRR for the National Biomedical Computational Resource (P41 RR008605) NIH/NINDS Center for Research in Computational Neuroscience (R01 NS046068) NSF for OptIPuter (NSF ANI0225642) NSF Pacific Rim Applications and Grid Middleware Assembly (NSF INT-0216895).

Part of the projects are funded by 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-02256402) 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.

References

- 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, 23–24 September 2004, Nice, France, 2004.
- [2] R. Singh, B. Jeong, L. Renambot, A. Johnson, J. Leigh, TeraVision: a distributed, scalable, high resolution graphics streaming system, in: Proceedings of Cluster 2004, San Diego, CA, 2004.
- [3] N. Taesombut, A. Chien, Distributed Virtual Computer (DVC): Simplifying the development of high performance grid applications, in: Proceedings of the Workshop on Grids and Advanced Networks, GAN '04, Chicago, IL, April 2004 (held in conjunction with the IEEE Cluster Computing and the Grid (CCGrid2004) Conference).
- [4] R. Wu, A.A. Chien, GTP: Group transport protocol for Lambda-Grids, in: Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid, CCGrid2004, Chicago, IL, April 2004.
- [5] J. Leigh, L. Renambot, T.A. DeFanti, M. Brown, E. He, N. Krishnaprasad, J.A. Meerasa, A. Nayak, K. Park, R. Singh, S. Venkataraman, C. Zhang, D. Livingston, M. McLaughlin, An Experimental OptIPuter architecture for data-intensive collaboration visualization, in: Third Annual Workshop on Advanced Collaborative Environments, Seattle, Washington, June 22, 2003 (co-located at the Twelfth IEEE International Symposium on High Performance Distributed Computing (HPDC 12) and Global Grid Forum 8).
- [6] C. Zhang, J. Leigh, T.A. DeFanti, TeraScope: Distributed visual data mining of terascale data sets over photonic networks, Future Generation Computer Systems 19 (6) (2003) 935–943.
- [7] Official OptIPuter website, http://www.optiputer.net/.
- [8] S.T. Peltier et al., The telescience portal for advanced tomography applications, Journal of Parallel and Distributed Computing: Computational Grid 63 (5) (2002) 539–550.
- [9] N. Krishnaprasad, V. Vishwanath, S. Venkataraman, A. Rao, L. Renambot, J. Leigh, A. Johnson, JuxtaView — a tool for interactive visualization of large imagery on scalable tiled displays, in: Proceedings of IEEE Cluster 2004, San Diego, CA.
- [10] N. Schwarz, S. Venkataraman, L. Renambot, N. Krishnaprasad, V. Vishwanath, J. Leigh, A. Johnson, G. Kent, A. Nayak, Vol-a-Tile – a tool for interactive exploration of large volumetric data on scalable tiled displays, in: IEEE Visualization 2004 Poster Compendium, 2004.
- [11] B. Jeong, R. Jagodi, L. Renambot, R. Singh, A. Johnson, J. Leigh, Scalable graphics architecture for high-resolution displays, in: Using Large, High-Resolution Displays for Information Visualization, Minneapolis, MN, USA, 24 October 2005.
- [12] C. Xiong, J. Leigh, E. He, V. Vishwanath, T. Murata, L. Renambot, T. DeFanti, LambdaStream — a data transport protocol for streaming network-intensive applications over photonic networks, in: Proceedings of the Third International Workshop on Protocols for Fast Long-Distance Networks, Lyon, France.
- [13] Official Calit2 website, http://www.calit2.net/.
- [14] Official LambdaRam website, http://www.evl.uic.edu/cavern/optiputer/ lambdaram.html.

- [15] D. Lee et al., Global telescience featuring IPv6 at iGrid2002, Future Generation Computer Systems 19 (6) (2003) 1031-1039.
- [16] Official PIN website, http://www.evl.uic.edu/core.php?mod=4&type=1& indi=216.
- [17] Official LambdaCam website, http://www.evl.uic.edu/cavern/ lambdacam/.
- [18] http://www.kddilabs.com/index2.html.
- [19] X. Wu, A.A. Chien, Evaluation of rate-based transfer protocols for Lambda-Grids, in: Proceedings of the 12th IEEE International Symposium on High-Performance Distributed Computer, HPDC, June 2004.
- [20] M.E. Martone, T.J. Deerinck, N. Yamada, E. Bushong, H. Ellisman Mark, Correlated 3D light and electron microscopy: Use of high voltage electron microscopy and electron tomography for imaging large biological structures, Journal of Histotechnology 23 (3) (2000) 261–270.



Rajvikram Singh is a Ph.D. student in computer science at the University of Illinois at Chicago. He received a Master of Science in computer science from UIC while working as a research assistant at the Electronic Visualization Laboratory with Jason Leigh's group. His current research interests include developing applications and protocols for high-speed networks, concurrent computing and high-definition video streaming. He is also currently working with

Mark Ellisman's group at the National Center for Microscopy and Imaging Research, University of California, San Diego as a computer scientist.



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 multidimensional data for scalable high-resolution displays and virtual-reality environments. He works extensively with domain experts in the bio-science and geo-science communities.



Nut Taesombut received the B.Eng. degree in Computer Engineering from Chulalongkorn University, Thailand, and the M.S. degree in Computer Science from the University of California, San Diego (UCSD), in 2000 and 2003, respectively. He is currently working towards the Ph.D. degree in Computer Science at the University of California, San Diego. His research interests are in the general area of highperformance computing and communication; specifically, Grid computing, wide-area federated systems, and advanced optical

network



David Lee is currently pursuing a Masters degree at the University of Michigan School of Information. He has previously worked as an applications engineer at the Microscopy and Imaging Research at the University of California San Diego School of Medicine. His research at the NCMIR mostly consisted of building advanced visualization solutions for high resolution biomedical applications. Previous work was also done in building computational web

portals for biologists, abstracting high-performance computing and data grids into a simple Web form.



Byungil Jeong received the BS and MS degrees in electrical engineering in 1997 and 1999 respectively 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.



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.

Abel W. Lin is a programmer/analyst at the National Center for Microscopy and Imaging Research.

Ruth G. West is an artist with background as a molecular genetics researcher. Working predominantly with computer-based media, West explores how artistic practice and aesthetic experience can nurture scientific discovery. West is Director, Visual Analytics and Interactive Technologies for the National Center for Microscopy and Imaging Research and a Research Associate at the UCSD Center for Research in Computing and the Arts, where she is the first Cal-(IT)2 New Media Artist crossing over to the Digitally Enabled Genomic Medicine Layer. She is the founder of "in silico v1.0", a collaborative of biologists, computer scientists and artists. Web: http://www.viewingspace.com and http:// www.insilicov1.org.



Hiromu Otsuka has graduated Kyoto University from the Faculty of Law in 1997. He is currently working at KDDI Corporation, Ubiquitous Networking Section Technology Development Department. His research interests are ubiquitous networking, grid computing, network security and cyber law.



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.



Mark H. Ellisman is a professor of Neuroscience and Bioengineering and the Director of the Center for Research in Biological Systems at UCSD. Prof. Ellisman directs the National Center for Microscopy and Imaging Research (NCMIR), an internationally acclaimed technology development center and research resource established by the National Institutes of Health (NIH). He has received numerous awards including a Jacob Javits award from

the NIH and the Creativity Award from the National Science Foundation, and he is a Founding Fellow of the American Institute of Biomedical Engineering. His scientific contributions include highly regarded work on basic molecular and cellular mechanisms of the nervous system and development of advanced technologies in microscopy and computational biology. He is a pioneer in the development of three-dimensional light and electron microscopy and combined application of these image acquisition tools and computational technologies to achieve greater understanding of the structure and function of the nervous system. His group was the first to introduce the idea of "Telemicroscopy" by demonstrating the network-enabled remote use and sharing of a high-energy electron microscope in 1992 and then developed practical systems now in use by researchers in the US and abroad. He led the successful Neurosciences Thrust for the National Partnership for Advanced Computational Infrastructure (NSF-NPACI) which resulted in the development of a stable collaborationenabling distributed system or "cyberinfrastructure" in the late 1990s. Most recently he has taken on the task of leading a large cyberinfrastructure initiative for the National Institutes of Health, creating the Biomedical Informatics Research Network (BIRN), linking researchers at nearly 40 universities in the US and Europe. BIRN provides cyberinfrastructure for collaboration and cooperative work related to neurodegenerative disorders and is regarded by many as the most advanced large-scale cyberinfrastructure project currently serving a large and geographically distributed scientific community.