Available online at www.sciencedirect.com

SCIENCE DIRECT®

Future Generation Computer Systems 22 (2006) 876–883

FGCS
FUTURE GENERATION COMPUTER SYSTEMS

www.elsevier.com/locate/fgcs

The first functional demonstration of optical virtual concatenation as a technique for achieving Terabit networking

Akira Hirano[a,*,1], Luc Renambot[a], Byungil Jeong[a], Jason Leigh[a], Alan Verlo[a], Venkatram Vishwanath[a], Rajvikram Singh[a], Julieta Aguilera[a], Andrew Johnson[a], Thomas A. DeFanti[a], Lance Long[a], Nicholas Schwarz[a], Maxine Brown[a], Naohide Nagatsu[b], Yukio Tsukishima[b], Masahito Tomizawa[b], Yutaka Miyamoto[b], Masahiko Jinno[b], Yoshihiro Takigawa[b], Osamu Ishida[b]

[a] *Electronic Visualization Laboratory, University of Illinois at Chicago, United States*
[b] *NTT Network Innovation Laboratories, NTT Corporation, Japan*

Abstract

The optical virtual concatenation (OVC) function of The Terabit LAN was demonstrated for the first time at the iGrid 2005 workshop in San Diego, California. The TERAbit-LAN establishes a lambda group path (LGP) for an application where the number of lambdas/L2 connections in a LGP can be specified by the application. Each LGP is logically treated as one end-to-end optical path, so during parallel transport, the LGP channels have no relative latency deviation. However, optical path diversity (e.g. restoration) can cause LGP relative latency deviations and negatively affect quality of service. OVC hardware developed by NTT compensates for relative latency deviations to achieve a virtual bulk transport for the Electronic Visualization Laboratory's (EVL) Scalable Adaptive Graphics Environment application.

---

1. Introduction

Parallelism is now penetrating into cutting edge applications, computers, and photonic networks to overcome the performance limitations resulting from using a single machine, single chip, or single lambda. The Scalable Adaptive Graphics Environment (SAGE) application developed at EVL accommodates 55 LCD displays driven by a 30 node cluster of PCs with the graphics rendering capacity approaching nearly a Tb/s. The high-end Linux cluster BlueGene, installed at Lawrence Livermore National Laboratories (LLNL) and developed by IBM has 65 536 PowerPC processors and 1152 GbE ports to communicate to other clusters. It is clear therefore that high-end applications and clusters need Tb/s capacity for their interconnections. We propose that an OXC-enabled photonic network will be the most promising way to realize a TERAbit-LAN [1–3]. There is a large number of on-going activities [4–6] evaluating and promoting photonic networking. The OMNInet testbeds cover various kinds of applications including actual production services [5]. We can find a sophisticated classification of such photonic network solutions in [6]. Among various kinds of users, the TERAbit-LAN project aims to meet high-end interconnection requirements with high reliability and reasonable cost.

To meet high-end Tb/s class demand, parallelism as mentioned above will play an important role. Currently, the I/O capacity of a PC is limited to 10 GbE. Thus multiple Network Interface Cards (NICs) will inevitably be required to support such demand. Even in a single PC, the PCI Express x32 interface which has 32 parallel lanes (each lane supports about 2 Gb/s) supports 64 Gb/s of capacity that exceeds the capacity of a conventional 10 GbE NIC. On the other hand, the number of parallel wavelengths in a DWDM transport system is approaching and will exceed 100 channels. This suggests that new schemes or mechanisms to manage multiple parallel NICs and transport through photonic networks will be essential to ensuring high throughput over LAN or WAN.

---

[*] Corresponding author.
E-mail address: ahirano@uic.edu (A. Hirano).

[1] On leave from NTT Network Innovation Labs. Japan.

0167-739X/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2006.03.027

The TERAbit-LAN project aims to develop such schemes or mechanisms to realize photonic networks that are tuned to accommodate terabit capacity communications which consist of a large number of parallel channels.

2. The TERAbit-LAN project

Fig. 1 shows a conceptual diagram of TERAbit-LAN. Three key components are needed: core optical cross connect switches (OXCs), transmission links, and end-interfaces. In this figure, we have shown a star-topology as the simplest example; this is by no means a constraint on the overall framework. In this configuration, parallelism plays an important role in photonic switching and transport. The end-systems generate flows of traffic in parallel and request parallel connections to target remote interfaces. The TERAbit-LAN accepts parallel connection requests and allocates parallel lambdas or parallel L2 connections and configures the OXC to establish dedicated parallel connections between these two end-interfaces.

The first of the three key components is the OXC. We have developed an OXC prototype to serve as a core switch for TERAbit-LAN as shown in Fig. 2 [7], in which an 8 × 8 Planar Lightwave Circuit (PLC) optical switch is equipped [7]. It supports various kinds of interface for both Network Node Interface (NNI) and User-Network Interface (UNI), and has a supervisory and control unit for management and signaling communication with Generalized Multi-Protocol Label Switching (GMPLS) capability within the Control-plane (C-plane). The interface cards can accommodate both 10 GbE LAN-PHY, OC-192 (10 GbE WAN-PHY). These incoming signals will be converted into 10 G OTN signal (OTU2) format in the cards and launched into the PLC optical switch. So all the optical switching will be done in OTN format.

In order to realize 1 Tb/s capacity switching on an OXC based on 10 GbE channels, we have to handle 100 channels or lambdas at the same time, and switch these lambdas individually. However, it appears impractical to extend the switching matrix to such high dimensions as it would be expensive, difficult to control, and unreliable.

Fig. 3 shows a hypothetical TERAbit-LAN application represented on a simplified OSI layer diagram from the viewpoint of future high-end applications. As we have described in the introduction, some high-end Grid applications (large-scale cluster computing, high-end visualization, etc.) require bandwidth which exceeds the maximum capacity of a single NIC. Inevitably, multiple NICs have to work together to meet such requirements. Since these NICs serve from a single application, we believe they should be concatenated to ensure the best performance for the application. These NICs will generate multiple lambdas for a single application. Therefore, these lambdas will be expected to be used as a group of lambdas in a limited number of optical paths. These grouped lambdas should hence be switched and transported as a single path to simplify the switching complexity and to reduce cost.

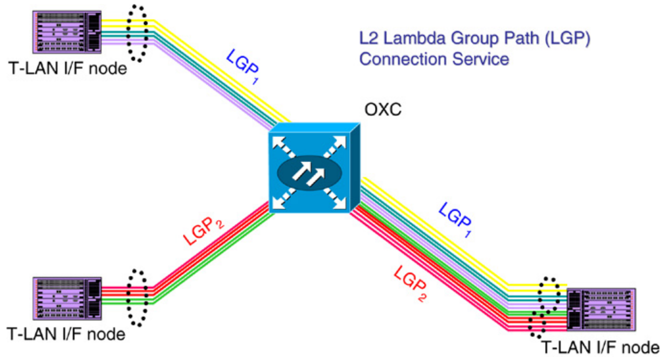

Fig. 1. Conceptual diagram of TERAbit-LAN.

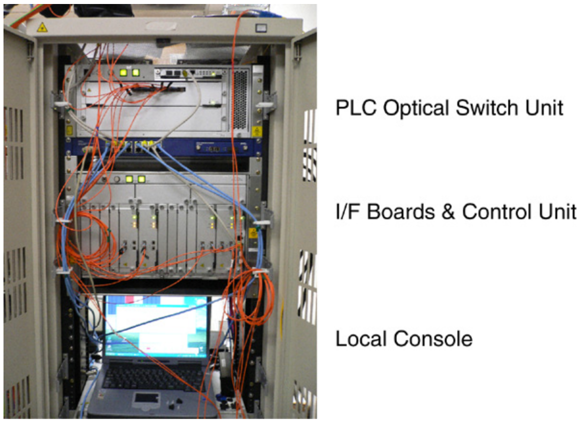

Fig. 2. OXC prototype.

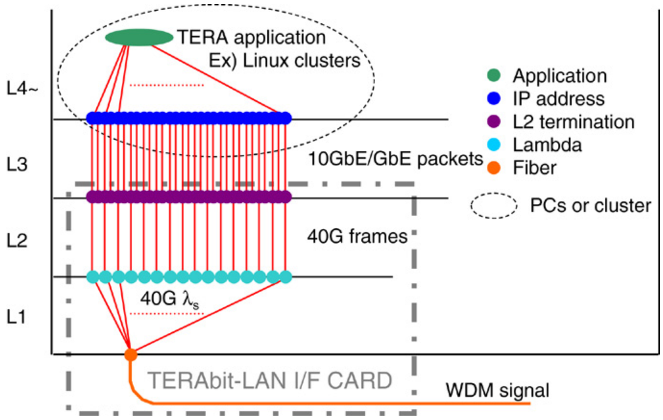

Fig. 3. An example of a Terabit application.

The TERAbit-LAN/WAN project focuses on this parallelism in switching and transport of multiple lambdas. To realize parallel switching of these multiple lambda, TERAbit-LAN/WAN adopts lambda group switching and establishes the notion of a lambda group path (LGP). High-end application users would require multiple lambdas as a whole for their clusters, so a TERAbit-LAN/WAN switch would provision and establish an LGP upon request of these applications. The number of lambdas in each LGP will be determined and provisioned by such applications. The TERAbit-LAN/WAN switches these LGPs as a single end-to-end optical path. In comparison with the case where one would switch all the lambdas individually, LGP-based optical switching only requires a switching matrix dimension no larger than the number of LGPs. For example, when we switch 100 lambdas individually, we need roughly a 100 × 100-dimension optical

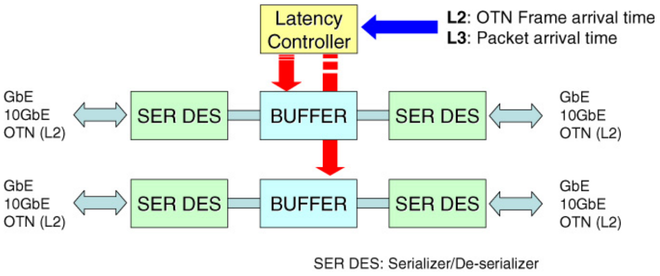

Fig. 4. An OVC circuit configuration.

switch. However, when we assume 10 lambdas can be grouped as an LGP, we need only a 10 × 10-dimension optical switch. The LGP scheme could therefore reduce switching complexity, the load of the operating system, and consequently Capex/Opex of the photonic network system.

The second key component is a transport or link system to support parallel lambdas. From the viewpoint of parallel transport of 10 GbE/GbE lambdas over TERAbit-LAN/WAN, relative latency deviations in each LGP could be a serious hindrance towards achieving high quality of service. This relative latency can arise for numerous reasons such as optical path diversity or group velocity dispersion in optical fibers. To overcome such issues, we have developed OVC (Optical Virtual Concatenation) [1]. The schematic configuration of an OVC circuit is shown in Fig. 4. It consists of serializer/de-serializer circuits, memory buffers, and a latency controller. The latency controller detects relative latency deviation among OTN frame arrival times for L2 implementation or Ethernet packet arrival times for L3 implementation. It is analogous to virtual concatenation in SDH/Sonet. We have extended it to multiple lambdas with the aid of ITU-T G.709 OTN [8]. OVC can be realized in L2 by an OTN (Optical Transport Network) function as defined in ITU-T Recommendation G.709. OVC can completely compensate for the aforementioned relative latency deviations to achieve virtual terabit bulk transport. We can achieve OVC accuracy of a single bit or timeslot (e.g. 0.1 ns for 10 Gb/s). OVC functionality can be realized also in L3. For this implementation, relative deviation of packet arrival time could be used to de-skew each channel to provide virtual bulk transport. However, the accuracy of OVC in L3 will degrade to milliseconds. With the help of OVC, visualization applications, clusters, and high-end PCs will not suffer any impairment that could result from the relative latency deviation in a single LGP.

The third key component of TERAbit-LAN consists of the end-interfaces. Currently, GbE/10 GbE interface cards are a popular interconnect for PCs or clusters, and of course these cards can be used to connect to TERAbit-LAN. OVC functionality can be achieved in L3 by monitoring the relative latency deviation of packet arrival or Round-Trip Time. However we also hope to develop a TERAbit-LAN card solution (Fig. 3). The solution accepts L2 connections or lambdas into one TERAbit-LAN card and can generate a WDM signal intended for an OXC node.

The capacity of a single channel is approaching 40 Gb/s, and we have developed a 40 Gb/s link system which is fully compliant with the ITU-T G.709 standard [8]. The system can accept both GbE and 10 GbE clients and transport by using advanced modulation schemes [9] with automatic group velocity dispersion (GVD) compensation [10,11] which ensures highly reliable operation of the system. In the near future, TERAbit-LAN will employ LGP-enabled OXC, 40 Gb/s/ch links with OVC, and will accept multiple 10 GbE channels as a single link.

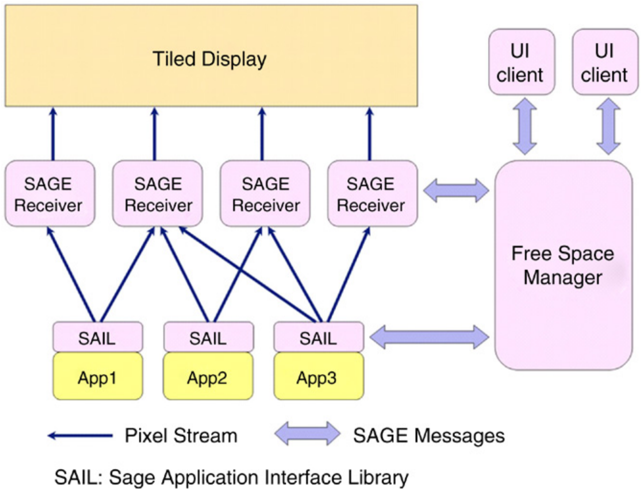

Fig. 5. SAGE framework.

3. The Scalable Adaptive Graphics Environment (SAGE)

The Scalable Adaptive Graphics Environment (SAGE) [12–14] is a middleware system for managing visualization and high-definition video streams for viewing on ultra-high-resolution displays such as EVL's 100 megapixel LambdaVision tiled LCD display wall. SAGE consists of Free Space Manager, SAGE Application Interface Library (SAIL), SAGE Receiver, and User Interface (UI client) as shown in Fig. 5. The Free Space Manager gets user commands from UI clients and controls pixel streams between SAIL and the SAGE Receivers. SAIL captures output pixels from applications, and streams them to appropriate SAGE Receivers. A SAGE Receiver can get multiple pixel streams from different applications, and displays streamed pixels on multiple tiles. Remote visualization applications (such as 3D rendering, remote desktop, video streams, and large 2D maps) stream their rendered pixels (or graphics primitives) to SAGE, allowing for any given 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).

4. The technology and iGrid experiences

4.1. Network configuration

Fig. 6 shows the network configuration for the demonstration. The SAGE sender PC is placed in EVL/Chicago and generates two parallel GbE video streams from its two GbE NICs. Each stream consists of half of a 3200 × 1200 graphics animation. GbE vlans were set up from EVL to iGrid so that each of

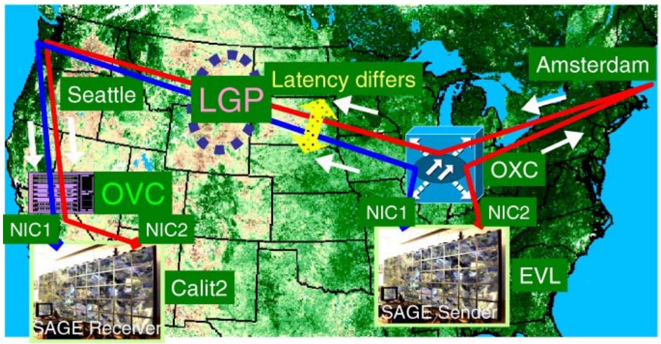

Fig. 6. Network configuration.

the streams followed decidedly different network paths to reach their destination. To demonstrate OVC functionality in an LGP, and to demonstrate its path restoration capability, we installed our OXC system at EVL. The L2 switch accepts GbE channels and maps them into 10 GbE LAN-PHY signals. Then the 10 GbE signals were converted to G.709 OTU2. Optical switching is performed on the OTN signal. The switched OTN signal was converted back into 10 GbE LAN-PHY and then into GbE signals.

Under normal conditions, these two GbEs would be transported to iGrid on two vlans (3707 and 3708) over CAVEWave. Thus there is no relative latency difference between the two links. The SAGE receiver PC would then stitch together the two halves of the image for display.

In the event of a fault in a GbE link, the OXC could switch to a restoration path to save the channel (the restoration path is provisioned in advance.) For this demonstration, we prepared entirely different paths (one through Amsterdam, but both originating from EVL) with significantly differing latency to show that OVC was able to transparently resynchronize the two flows. We also installed a network emulator which was kindly provided by Anue Systems Inc. for our demonstration to simulate expected latency from Amsterdam, should the real link fail during the workshop.

4.2. OXC system

The configuration of the OXC system used in the demonstration is shown in Fig. 7. It consists of an 8 × 8 switching unit, 10 GbE LAN-PHY interface cards, GMPLS-enabled control unit, and a L2 switching unit. The L2 switch accepted GbE client signals and mapped them into 10 GbE LAN-PHY signals. Then the 10 GbE signal was passed to the OXC I/F card which converted the incoming 10 GbE LAN-PHY signal into a 10 G OTU2 signal which is compliant with the ITU-T G.709 standard. All the optical switching was carried out as an OTN standard signal.

4.3. OVC implementation

There are two possible implementations for OVC function as discussed in The TERAbit-LAN section. These include L2 solution by OTN frame or L3 solution by Ethernet packet.

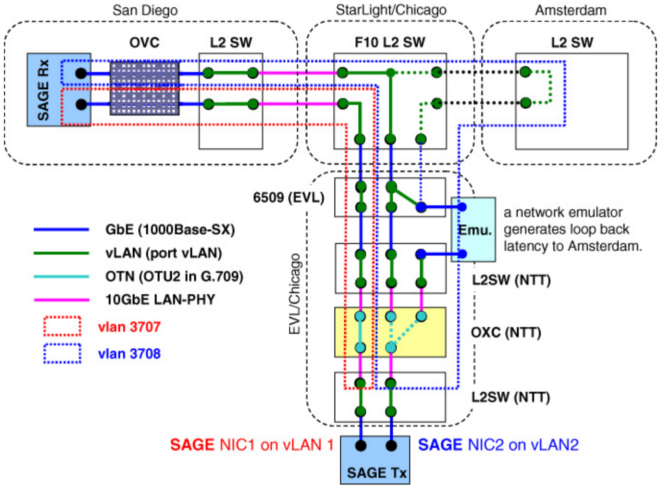

Fig. 7. Configuration of OXC system used in the demo.

As for the scale of latency deviations, we are supposing about 200 ms at maximum. It could be the case when a path restoration in global scale occurred. In our OVC hardware, we have implemented 1 GB of buffer that can sufficiently support around 500 ms of latency which is enough to meet with global scale path diversity. In the demonstration, OVC functionality was realized in L3 by integrating a latency detection module and a latency controller as shown in Fig. 4. Ideally OVC function should be implemented in L2 to realize highest accuracy. For such a purpose, all the networking equipment including L2 switches, routers, and repeaters must be compliant with the G.709 standard. But unfortunately, not all the equipment was ready to conform to the standard. Thus we adopted an L3 OVC configuration in the demonstration. The latency detection module detected RTT for each channel in a LGP. Then the differences in relative latency that were calculated by the obtained RTT times were used to control the latency control unit. The restored path had excessive latency of about 175 ms which has significantly impaired quality of streaming from EVL. To ensure keeping the quality of streaming, we have to compensate the latency deviation to less than about several tens of ms to avoid visible deterioration. The OVC circuit we have implemented has less than 1 ms of accuracy of compensation which is sufficient for this purpose. For L2 implementation, absolute latency determined by a change of geographical distance is a main component. In global scale systems, some small latency induced by electrical regeneration in link systems will also be needed to be considered. For L3 implementation, we need to consider store and forward latency to route packets. Usually, this latency component will be the second largest next to the absolute latency, and it will be statistically distributed. So we have averaged over 3 s to get a stable value for compensation. In other words, we have compensated at a fixed averaged value instead of packet by packet compensation. The averaging time could be optimized depending on network condition. We will be able to decrease it when we have no congestion, since we can expect stable, not scattered, latency deviation. However, when we faced severe congestion and large numbers

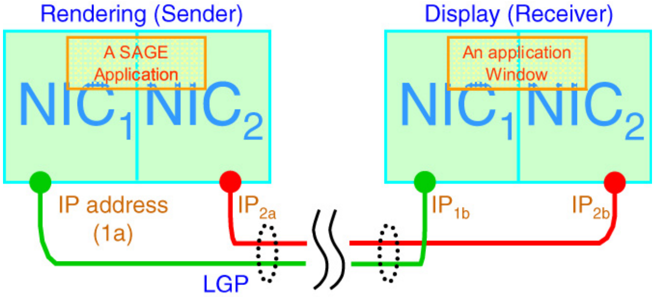

Fig. 8. SAGE with multi-network support.

of packet loss, we can hardly get stable averaged latency to compensate. When the congestion had cleared, we can get a stable latency deviation and OVC will re-start within 3 s as we mentioned above. In our demo, we developed TCL scripts on Windows OS to measure RTT and average, and to control the latency controller. In the demonstration problems occurred in the loop-back connection through Amsterdam, so the equivalent relative latency of 180 ms was provided by the network emulator placed between a port of the OXC and a port in Cisco 6509 which was connected to StarLight's E1200.

4.4. SAGE with multi-network support

SAGE was improved to accommodate multiple network connections to demonstrate OVC functionality in the TERAbit-LAN project. Fig. 8 shows a schematic of SAGE's new multi-network capability. A single SAGE application can use multiple IP addresses or multiple NICs and the bandwidth allocation or splitting on the SAGE window can be arbitrary configured. On the sender side, traffic of an application running on a SAGE window was split into two streams depending on the location of the application on the window. When the application ran on only the left half of the SAGE window, the traffic was pushed out from only one NIC which corresponds to the window. In the actual demonstration, we placed a real-time streaming application at the center of a SAGE window. Therefore, almost equal amounts of traffic were output from both NICs and received in the receiver side in Calit2.

4.5. LIVE streaming from EVL

Fig. 9 shows a snapshot of the demonstration station. We had placed a moving toy in front of the HD camera which was located in EVL/Chicago so that the audience could easily see tearing between the two video streams in real-time. In the control condition where there is no relative latency deviation in two constituent channels in a LGP, there was no tearing noticed on the display screen. Then, when an intentional fault occurred in a channel, audiences could clearly notice tearing in the video feed. The tearing resulted from additional latency induced by a change in optical route restored by the OXC system. When we turned the OVC circuit ON, the tearing was suppressed immediately. The recovery time needed was on the order of a few seconds, due to some time needed to detect and determine relative latency differences between the two GbE channels. We have encountered severe bandwidth fluctuation and sometimes intermittent connection to EVL/Chicago in the former part of our demo timeslot. But in the last 30 min, the network performance recovered well and showed good throughput. We achieved 320–400 Mb/s of peak bandwidth used for each vlan in the demo.

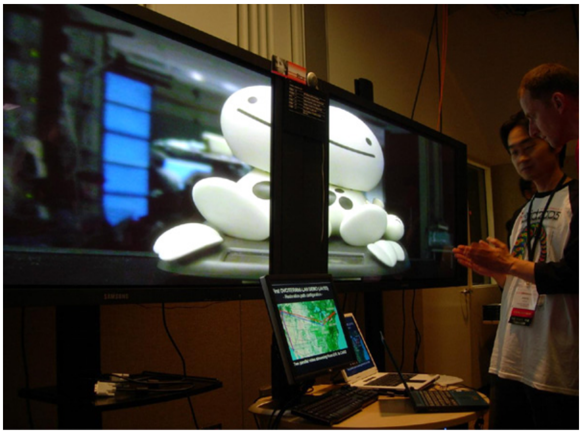

Fig. 9. LIVE double-width video streaming from EVL/Chicago.

5. Conclusions

We have successfully demonstrated OVC functionality in a TERAbit-LAN LGP with a high-performance graphics stream application. An OXC placed in EVL/Chicago established an LGP between EVL/Chicago and Calit2/San Diego. In the demonstration, OVC was realized in L3 by simple latency detection and control, since an OTN-enabled transport layer was not available in end-to-end. But we are anticipating that the G.709 OTN standard will penetrate into the transport network and fully realize the TERAbit-LAN concept. In our demo, we have investigated a live streaming application. To collect information concerning the accuracy of latency compensation needed for applications, we will investigate similar kinds of tests for a wider extent of applications such as tightly interactive ones. In addition, we have tested OVC in a LGP with two parallel vlans in the demo, and they were static, since we were focusing on OVC in this time. We plan to implement a novel *c*-plane network which can accept a required number of lambdas in a LGP from end-hosts and can configure end-hosts and optical switches to establish a LGP dynamically.

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). The LambdaVision is a trademark of the Board of Trustees of the University of Illinois. We are also grateful to Anue Systems Inc. for providing a network emulator for our demonstration.

References

[1] Masahito Tomizawa, Jun Yamawaku, Yoshihiro Takigawa, Masafumi Koga, Yutaka Miyamoto, Toshio Morioka, Kazuo Hagimoto, Terabit-LAN with optical virtual concatenation for grid applications with super computers, in: Proceedings of OFC2005, paper OthG6, 2005.

[2] TERAbit-LAN, http://www.evl.uic.edu/cavern/terabitlan/.

[3] Symposium Panels in iGrid2005, Advances toward economic and efficient Terabit LANs and WANs, http://www.igrid2005.org/program/symposium_panels.html.

[4] StarLight, http://www.startap.net/starlight/.

[5] Joe Mambretti, Recent progress on transLight, OptIPuter, OMNInet and future trends toward lambda grids, in: Proceedings, CPT2005 8th International Symposium on Contemporary Photonics Technology, 12–14 January 2005, Tokyo, Japan.

[6] Cees de Laat, Erik Radius, Steven Wallace, The rationale of the current optical networking initiatives, Future Generation Computer Systems 19 (2003) 999–1008.

[7] Toshio Watanabe, Shunichi Sohma, Takashi Goh, Tomohiro Shibata, Hiroshi Takahashi, Compact 8 × 8 Silica-based PLC switch with compressed arrangement, in: Tech. Digest of ECOC2004, Th3.6.3, 2005.

[8] ITU-T Recommendation G.709, Interfaces for the Optical Transport Network (OTN), 2003.

[9] Yutaka Miyamoto, Akira Hirano, Kazushige Yonenaga, Akihide Sano, Hiromu Toba, Koichi Murata, Osamu Mitomi, 320 Gbit/s (8 × 40 Gbit/s) WDM transmission over 367 km zero-dispersion-flattened line with 120 km repeater spacing using carrier-suppressed return-to-zero pulse format, in: Tech. Digest of OAA99, PDP4.

[10] Akira Hirano, Yutaka Miyamoto, Novel modulation formats in ultra-high-speed optical transmission systems, and their applications, in: Tech. Digest of OFC2004, invited paper ThM1, 2004.

[11] Shoichiro Kuwahara, Akira Hirano, Hiroji Masuda, Hiroto Kawakami, Yutaka Miyamoto, Mikio Yoneyama, Yoshiaki Kisaka, Kenji Sato, Masahito Tomizawa, Yasuhiko Tada, Tomoko Sawabe, Teturo Fujii, Kazushige Yonenaga, WDM field demonstration of 43 Gbit/s/channel path provisioning by automatic dispersion compensation using tone modulated CS-RZ signal, in: Tech. Digest of ECOC2003, paper Tu1.6.2, 2003.

[12] Byungil Jeong, Luc Renambot, Rajvikram Singh, Andrew Johnson, Jason Leigh, High-performance scalable graphics architecture for high-resolution displays, EVL Technical Document Technical publication 20050824_Jeong, 2005.

[13] Luc Renambot, Arun Rao, Rajvikram Singh, Byungil Jeong, Naveen Krishnaprasad, Venkatram Vishwanath, Vaidya Chandrasekhar, Nicholas Schwarz, Allan Spale, Charles Zhang, Gideon Goldman, Jason Leigh, Andrew Johnson, SAGE: The scalable adaptive graphics environment, in: Proceedings of WACE 2004, September 2004, Nice, France.

[14] SAGE, http://www.evl.uic.edu/cavern/sage/.

Akira Hirano is a senior research engineer at NTT Network Innovation Laboratories in Japan. He received the B.S. degree in physical chemistry and M.S. degree in astrophysics from Osaka University, Japan. In 1993, he joined Nippon Telegraph and Telephone Corporation working in the high-speed optical communications area. He was a visiting researcher at University of Illinois at Chicago, USA (2005–2006), where he engaged in the research of photonic networking. His main interests are in high-speed optical transport systems, optical signal processing, and photonic networking. He is a member of the Institute of Electrical and Electronics Engineers (IEEE) and Institute of Electronics, Information, and Communication Engineers (IEICE) of Japan.

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.

Byungil Jeong received the BS and MS degree 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.

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.

Alan Verlo is an Associate Director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago, responsible for EVL's high-performance computing and communications infrastructure. He is also a member of the StarLight network engineering team, which supports international, national, regional and local research and education network connections. For many years, Verlo has also been active with the Supercomputing (SC) SCinet networking team, with the Technical Engineering Working Group of the Global Lambda Integrated Facility (GLIF), and with the US National Coordinating Office's Large Scale Network (LSN) Joint Engineering Team (JET).

Venkatram Vishwanath is a doctoral candidate under Dr. Jason Leigh at the Electronic Visualization Laboratory in the Department of Computer Science at the University of Illinois at Chicago. He is currently involved with the OptIPuter Project. His current research interests include high-performance network based computing, end-system and network-aware synergistic high-speed transport protocols, wide-area network memory caches for extremely-large data-intensive applications. He is the author of LambdaRAM, LambdaStream and a coauthor of MAGNET.

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.

Julieta Cristina Aguilera is a graduate student and research assistant at the Electronic Visualization Laboratory. She studied design at the School of Architecture of the Universidad Católica de Valparaíso, Chile, and at the University of Notre Dame, Indiana, USA. Her interests are in grids and interfaces.

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.

Thomas A. DeFanti, Ph.D., at the University of Illinois at Chicago, is a distinguished professor of Computer Science and director of the Electronic Visualization Laboratory (EVL). At the University of California, San Diego, DeFanti is a research scientist at the California Institute for Telecommunications and Information Technology (Calit2). He shares recognition along with EVL director Daniel J. Sandin for conceiving the CAVE virtual reality theater in 1991. Striving for a more than a decade now to connect high-resolution visualization and virtual reality devices over long distances, DeFanti has collaborated with Maxine Brown of UIC to lead state, national and international teams to build the most advanced production-quality networks available to academics.

Lance Long is a visiting research programmer at the Electronic Visualization Laboratory. He holds a M.S. in Computer Science from the University of Illinois at Chicago. He develops and deploys clustering and networking solutions for advanced scientific visualization. He participates in related conferences, providing support and promoting collaboration for emerging research technologies.

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.

Maxine D. Brown is an Associate Director of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). Currently, she is co-principal investigator of the NSF international networking TransLight/StarLight award, project manager of the NSF-funded OptIPuter initiative, a member of the PRAGMA steering committee, and a member of the Florida International University's CHEPREO external advisory board. Brown is a founding member of GLIF and is co-chair of the GLIF Research & Applications (RAP) working group along with Larry Smarr.

Naohide Nagatsu received the B.E. and M.E. degrees in electrical and information engineering from Yokohama National University, Kanagawa, Japan, in 1990 and 1992, respectively. He joined NTT Transmission Systems Laboratories in 1992, and has engaged in the research of photonic networking. He was a visiting researcher at University College London, United Kingdom (2003–2004), where he engaged in the research of optical burst switching. He has been active in ITU-T SG 15 for the standardization on the optical transport network (OTN) architecture since 1997. Mr. Nagatsu is a member of IEEE and IEICE.

Yukio Tsukishima received his B.C and M.S from Tokyo Institute of Technology in 1999 and 2001. He became a staff at Network Innovation Laboratories of NTT in 2001. He is engaged in research and development of IP and Optical layer-interworking technologies using GMPLS functions. He is also researching and developing Grid Network Management Systems on GMPLS-NW.

Masahito Tomizawa received his B.E. in applied physics in 1990 and M.Sc. in physics in 1992 from Waseda University, Tokyo, Japan. He received the Dr. Eng. Degree in applied physics in 2000 from the same University. In 2003 to 2004, he was invited as a visiting researcher from Laboratory for Information and Decision Systems in Massachusetts Institute of Technology (MIT), Massachusetts, USA. He received the outstanding Paper Award from IEICE in 2003. He is a member of IEEE (Institute of Electronics, Electrical Engineers).

Yutaka Miyamoto received the B.E. and M.E. degrees in electrical engineering from Waseda University, Tokyo, Japan, in 1986 and 1988, respectively. He is now a senior research engineer, supervisor of photonic transmission systems research group in NTT network Innovation Laboratories. His current research interest includes high-capacity optical transport systems with advanced modulation formats. He is a member of IEEE and IEICE of Japan. He received the Best Paper award of the first optoelectronics and communication conference (OECC) in 1996 from the IEICE, and the Best Paper award from the IEICE Communication Society in 2003.

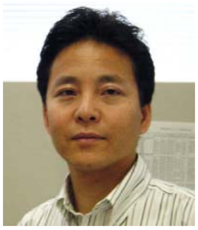

Masahiko Jinno is a senior research engineer, supervisor at NTT Network Innovation Laboratories. He received the B.E. and M.E. degrees in electronics engineering from Kanazawa University, Kanazawa, Japan, in 1984 and 1986, respectively, and the Ph.D. degree from Osaka University, Osaka, Japan, in 1995. He received the Young Engineer Award from the Institute of Electronics, Information, and Communication Engineers (IEICE) in 1993, the Best Paper Award in 1997 and 1998 from the Second and Third Optoelectronics and Communications Conference (OECC), and the NTT president award in 2000.

Yoshihiro Takigawa received his BS and MS degrees from Waseda University, Tokyo, Japan in 1980 and 1982 respectively. He has been with NTT from 1982. He was in charge of research and development of ISDN subscriber systems and ATM subscriber systems. He was engaged in the NTT Multimedia Joint Trial and the Asia Multimedia Forum. From 2000 he has been responsible for researches on Photonic networking, Optical Cross Connect systems, and GMPLS protocols.

Osamu Ishida joined NTT Transmission Systems Laboratories in 1988. His research includes coherent optical communication systems, optical cross-connect systems, and 10Gb/s Ethernet transport. He currently leads Media Networking Systems Research Group, NTT Network Innovation laboratories, where he investigates Terabit LAN. Mr. Ishida is a member of IEEE. He received the Young Engineers Award from IEICE in 1995. He is co-author of a Japanese text book on 10b/s Ethernet Technologies.