National Science Foundation, Directorate for Computer Information Science and Engineering
Division of Advanced Networking Infrastructure & Research (ANIR)

NSF Cooperative Agreement No. ANI-9730202
June 2002 Quarterly Status Report

July 22, 2002

Tom DeFanti, Maxine Brown, Andy Johnson, Dan Sandin, Jason Leigh, Alan Verlo, Laura Wolf
Electronic Visualization Laboratory
University of Illinois at Chicago

Linda Winkler, Bill Nickless, Caren Litvanyi
Argonne National Laboratory

Jim Williams, Stephen Peck
Indiana University

Table of Contents

A. Summary of Technical Activities

A.1. Euro-Link Network Status and Institutions

A.1.a. CERN

CERN connected to the StarLight facility in late March with an OC-12 (622Mb) and planned to bring a second, OC-48 (2.5Gb) link, on behalf of the European Union (EU)-funded "DataTAG" project, to StarLight this summer.

Unfortunately, with its provider KPNQwest announcing in late May that it might cease operation, CERN is temporarily using Géant to reach North American research networks. CERN is in the process of signing with new providers for its two circuits to Starlight. Its goal is to have the 622Mb circuit delivered by the end of July and the 2.5Gb circuit before September 1.

CERN's transatlantic connections went from 12Mb in April 1999 to 21Mb in September 1999, 45Mb in April 2000, 155Mb in January 2001, and then 2*155Mb in December 2001. By September 2002, CERN will have ~3Gb connectivity with its two new circuits. According to Olivier Martin, the interesting news is that the new 622Mb circuit costs about half the 12Mb ATM circuit provided by Cable & Wireless in April 1999.

A.1.b. IUCC

In May, Israel disconnected its STAR TAP connection and now no longer receives Euro-Link funding. Because the European Union's Géant project pays 50% of the cost of a transatlantic circuit, Israel is taking advantage of its Géant connection to connect to Abilene in New York. The Euro-Link and STAR TAP websites have been updated to reflect the IUCC disconnection.

A.1.c. NORDUnet

NORDUnet's 622Mb transoceanic link, provided by KPNQwest, failed on June 22 as the company ceased operation due to bankruptcy. NORDUnet is temporarily using Géant to reach New York. Previously, KPNQwest carried the 622Mb to New York City, where NORDUnet peered with Abilene, and then had a 155Mb link to StarLight in Chicago, where NORDUnet peered with other research networks, notably ESnet and CA*net.

NORDUnet is working hard to quickly reestablish direct connections to Abilene and STAR TAP/StarLight. The 622Mb from Stockholm to NORDUnet's POP in New York will now be delivered by T-Systems, who subcontracted to Qwest to connect from the T-Systems' New York location to the Abilene's NY POP and to StarLight in Chicago. Qwest is expected to establish these links by July 22; the transoceanic part is ready.

NORDUnet wants to maintain this connection until August 2003. Over the coming year, NORDUnet will intensify and reinvigorate Nordic/American research collaboration so there is justification for a NORDUnet 2.5Gb wavelength to the USA in Summer 2003.

Note: NaukaNet, the Russian research and education network, used Teleglobe as its transoceanic provider. Teleglobe went out of business in May 2002. NORDUnet is temporarily passing NaukaNet routes (as there is a link from the Nordic countries to St. Petersburg in Russia).

A.1.d. Renater2

RENATER2 maintains an OC-3 link to STAR TAP. We have received no information on possible future upgrades.

A.1.e. SURFnet

Teleglobe and Global Crossing were SURFnet telecom providers; as of May 31, Level (3) replaced Teleglobe, which went out of business.

SURFnet maintains two OC-12s (2*622Mb) and one OC-48 (2.5Gb) connection from Amsterdam to the StarLight facility in Chicago, which also serves as its POP. SURFnet has a Cisco ONS 15454 (optical transport platform) at StarLight, which Kees Neggers invites others (e.g., CERN, MREN) to plug into, as needed.

SURFnet's four major peers in the USA are Abilene, CA*net, ESnet and vBNS. Only Abilene and ESnet are currently connected to the StarLight router, although CA*net and vBNS plan to connect soon. Until then, SURFnet maintained a dedicated OC-12 between StarLight and STAR TAP in order to reach them. On May 1, they terminated their dedicated OC-12 link (due to Ameritech's increased charges) and now use an OC-12 shared link that STAR TAP maintains.

The UIC Laboratory for Advanced Computing (LAC) has clusters both at StarLight and at NetherLight (at SARA in Amsterdam). EVL and NCDM will begin testing applications from Chicago to Amsterdam at 2xGigE rates next quarter.

A.1.f. Non-EuroLink Connections

As we transition from STAR TAP to StarLight, below is a snapshot of who is currently connected to StarLight. We list it here, as it represents a major benefit to our Euro-Link members.

StarLight, the optical STAR TAP initiative, is an advanced optical infrastructure and proving ground for network services optimized for high-performance applications. Operational since summer 2001, StarLight is a 1 Gigabit Ethernet (GigE) and 10GigE switch/router facility for high-performance access to participating networks and will ultimately become a true optical switching facility for wavelengths. It is located in downtown Chicago on the Northwestern University campus. This telecommunications facility is designed by researchers, for researchers, for the benefit of global e-Science collaboration.

a. ***Four telecom carriers*** have fiber into StarLight: SBC/Ameritech, Qwest, Global Crossing and AT&T. Others are doing fiber builds or lease lines from the above carriers.
b. ***Three Federal Research & Education*** networks either connect (DoE's ESnet and NASA's NREN) or will soon connect (NSF's vBNS+) to StarLight.
c. ***Internet2's Abilene*** connects to StarLight at 1GigE, with plans to soon upgrade to 10GigE.
d. ***The Optical Metro Network Initiative (OMNInet),*** a 10GigE testbed developed by Nortel Networks and SBC/Ameritech with partners Northwestern University, University of Illinois at Chicago and CANARIE/Canada, has one of its collocation points at StarLight.
e. ***I-WIRE,*** the $7M State of Illinois dark fiber initiative that connects major universities and government laboratories in Illinois, has its hub located at StarLight.
f. ***Member institutions of the Midwest GigaPoP MREN*** are upgrading to 1GigE and moving to StarLight.
g. ***TeraGrid,*** an NSF-funded Major Research Equipment initiative, has its Illinois hub located at StarLight.
h. ***Nine international Research & Education networks/consortiums*** connect or will soon connect to StarLight, and many have equipment collocated at StarLight as well. Those currently at StarLight include: CERN (Laboratory for Particle Physics in Geneva), SURFnet (Holland), NORDUnet (Nordic countries), TransPAC (Asian-Pacific countries), NaukaNet (Russia) and AMPATH (South American countries). Those who will connect Summer 2002 are: CA*net4[1] (Canada) and DataTAG (the European Union-funded Trans-Atlantic Grid network). These activities are fostering interest from other Research & Education networks; the United Kingdom is in the process of developing UK-Light, with interest in creating linkages to NetherLight and CERN in Europe and StarLight in the US.
i. ***Developers of 6TAP, the IPv6 global testbed,*** notably ESnet and Viagenie (Canadian), have an IPv6 router installed at StarLight.
j. ***NLANR*** works with STAR TAP on network measurement and web caching; the NLANR AMP (Active Measurement Platform) is located at STAR TAP and the web cache is located at StarLight.

A.2. Engineering Services

A.2.a. StarLight/Abilene Connectivity

In April, Internet2's Heather Boyles said Abilene would likely maintain its direct connection to the STAR TAP NAP (currently OC-12, but may be reduced to OC-3 based on decreasing demand) until at least the fourth quarter of 2002 (April 2003 at the latest). After that, Abilene will decide whether to continue its direct connection based on the number of peer networks still connecting there.

Abilene's OC-48 core router, located at the Chicago Qwest PoP, connects to StarLight via two 1Gbps connections using I-WIRE fiber; one connection is for StarLight National Research Network traffic (e.g., SURFnet) and one connection is for MREN traffic.

A.2.b. StarLight/STAR TAP Connectivity

An OC-12 connection exists between the StarLight facility and STAR TAP (at the Ameritech NAP) that is available for shared use by StarLight customers.

A.2.c. STAR TAP Router

In our distributed STAR TAP/StarLight environment, there is a STAR TAP Router at the Ameritech NAP (Cisco 7505), a STAR TAP Router at 710 N. Lake Shore Drive (Juniper M10), and an MREN Router at 710 (Juniper M5). A Cisco 6509 serves as an additional MREN Router as well as the StarLight GigEthernet Exchange Point switch.

[1] On June 24, CA*net4 engineers began installing two (2) 1Gb connections in a Cisco ONS 15454 they co-located at Starlight, which will connect to CA*net4 Juniper routers in Toronto and Winnipeg. CA*net 3 STAR TAP turn down is July 31. There will be approximately one month overlap between CA*net 4 and CA*net 3 in Chicago.

A.2.d. STAR TAP/StarLight NLANR Web Cache

An NLANR web cache has been operational at StarLight since December 2001.

In March, STAR TAP/StarLight engineers moved Genuity peering that supports the web cache service from the STAR TAP Router to the StarLight Router.

A.2.e. 6TAP

Both STAR TAP and StarLight host IPv6 routers, provided by ESnet.

A.2.f. DiffServ

Oliver Yu and Jason Leigh are setting up the EMERGE 2 DiffServ testbed between EVL, Northwestern University, and KISTI (over KREONET/STARTAP). In EMERGE 1, priority queuing was tested over a testbed between EVL and Argonne National Lab. EMERGE 2 will attempt to employ Weighted Fair Queuing. CERN has expressed interest in participating in DiffServ experiments.

Students working under Oliver Yu are designing schemes for the following research objectives:

- Comparing QoS performances for different DiffServ types
- Investigating TCP tuning over DiffServ network
- Compare per-flow policing versus per-class policing

A.2.g. StarLight/STAR TAP Documentation

Soon, one will be able to get to the StarLight home page by using either <http://www.startap.net> or <http://www.startap.net/starlight>. The archived STAR TAP directories will be accessible via the StarLight site.

A.2.h. International Transit Network (ITN)

Rather than bring circuits to STAR TAP or StarLight, some international research networks peer with Abilene at one of the US coasts, where they can pass traffic to US universities (via Abilene) and to other international research networks (via Abilene and CA*net ITNs). Further information is available on the Abilene <www.ucaid.edu/abilene/html/itnservice.html> and CA*net web sites <www.canet3.net/optical/documentation.html>, as well as STAR TAP's <www.startap.net/CONNECT> page.

A.3. NOC Services

The Global NOC will move to the NetSaint network monitoring system in Summer 2002. <http://www.netsaint.org>

As of April 2002, the Global NOC and STAR TAP engineers have biweekly conference calls to discuss engineering and NOC issues.

MRTG graphs depicting StarLight traffic were added to the STAR TAP NOC web page under Network Monitoring <http://noc.startap.net/noc.html>. The NOC has traffic monitors for the StarLight M10, 6509 and LS1010. <http://loadrunner.uits.iu.edu/mrtg-monitors/starlight/>

B. Euro-Link Performance Analysis Tools

B.1. Network Monitoring Tools

Bandwidth Utilization Radar Map

Under Jason Leigh's supervision, EVL student Brenda Lopez is working with Linda Winkler to develop graphical radar maps for the STAR TAP and StarLight routers.

UCAN: Unified Collaboratory for Analyzing Networks

UCAN is an integrated, collaborative, extendable workbench tool for network engineers, researchers and application programmers using networks. UCAN can be used for running collaborative network tests between multiple sites, thus helping analyze network performance in a group environment. The first prototype, developed by EVL doctoral student Naveen Krishnaprasad, was tested between EVL and KISTI (Korea Institute of Science and Technology Information). The results can be seen at <http://www.evl.uic.edu/cavern/ucan/ucan_kisti.html>.

Next steps include miscellaneous bug fixing, and use of UCAN's collaborative modules to write a separate shared widget library over the CAVERNsoft networking library. This library can be used to develop collaborative GUIs for applications other than UCAN.

B.2. High-Bandwidth Transmission Over Long Distance Networks

Quality of service Adaptive Networking Toolkit (QUANTA)

EVL's CAVERN group will release Quanta in mid-July as well as a version of CAVERNsoft that uses Quanta. Quanta is an applications-level API that translates high-level data distribution requirements into low-level optimized networking protocols and parameter settings. This work is specifically targeted for optical networks.

CAVERNsoft provides a set of APIs that bridge graphics and networking, alleviating the need for visualization developers to become networking experts. CAVERNsoft focuses on problems related to networking and data distribution by providing a broad range of data distribution mechanisms including: message passing, distributed shared memory, remote procedure calls, forward error corrected UDP, and parallel TCP for bulk data transfer. Since 1997, we have observed the growing use of CAVERNsoft beyond VR to other networked applications. As a result, we recently separated the networking components of CAVERNsoft into its own toolkit called Quanta.

Work is now underway to provide Quanta with mechanisms to signal DiffServ-capable routers and make dedicated wavelength reservations on optically switched networks. Quanta's DiffServ testbed consists of a set of DiffServ routers at EVL, KISTI (Korea Institute of Science and Technology Information), and Northwestern University. Quanta's optically switch testbed consists of StarLight and OMNInet.

Reliable Blast UDP (RUDP)

EVL doctoral student Eric He and Argonne's Mike Papka are doing comparison tests between RUDP and NTT's TCP-BDTE and UDP-CMTE with token buckets, and parallel TCP, between StarLight and SARA.

Jason Leigh is supervising Eric to implement an RUDP transmission scheme to accelerate reliable data transmission over fat networks. In RUDP, the sender blasts all the data (each packet is identified by a sequence number) to the receiver. Upon receipt, the sequence numbers are checked and any lost packets are identified. The receiver then sends a lost packet report back to the sender through TCP. Upon receipt of the report, the sender retransmits the lost packets. The procedure continues until the receiver receives all packets. This technique is believed to be most effective when used in conjunction with QoS, since the guaranteed bandwidth will minimize transmission errors. The RUDP scheme exploits low transmission errors to maximize throughput.

Forward Error Correction (FEC)

Eric rewrote the FEC code for future integration into CAVERNsoft. Since the FEC protocol is mainly used to stream data, he created an interface specifically for streaming data between machines, and is ready to begin testing. The next steps are to conduct more FEC experiments based on the new code, to make RUDP responsive to network congestion, and to incorporate the protocols into the Quanta framework.

B.3. Ultra-High-Bandwidth Transmission Over Long Distance Networks (StarLight)

Interrupt Coalescing and Jumbo Frames

The two 16-node Linux clusters that EVL ordered last quarter are expected to arrive in late July; one for the StarLight facility and one for UIC. Engineers expect both to assembled and installed in August. The StarLight cluster will augment the four-node PC cluster currently used by EVL to run tests to Northwestern University. The new cluster will have three times the bus bandwidth (400MHz x 64bits=2.98GBytes/s) of the current four-node cluster. In the past, EVL was able to only achieve ~500Mbps (local area) PC transmission rate out of the box. With appropriate tuning using *Interrupt Coalescing* and *Jumbo frames*, EVL has achieved bandwidth of 850Mbps.

EVL is working with Cees de Laat at University of Amsterdam and Paul Wielinga of SARA (Amsterdam) to perform tests over the 2.5Gb NetherLight/StarLight link. Tests include VNC for streaming of clustered desktops, WireGL for streaming OpenGL visualizations, and streaming stereoscopic visualizations.

TeraVision: Ultra-Resolution Visualization Streaming

EVL student Raj Singh is doing weekly Gigabit LAN streaming tests to perfect the TeraVision software. He has integrated a video capture card and finished porting the code to Windows. He will soon begin streaming experiments between StarLight and SARA clusters. The goal is to stream a high-resolution animation to University of Amsterdam by August, then test a version of the TeraVision code that can be used to stream video to EVL's Access Grid Augmented Virtual Environment (AGAVE), in time for the iGrid 2002 event in September.

TeraVision's goal is to send high-resolution video streams between clusters using distributed servers and clients. The software allows multiple streams of synchronized video to be streamed between clusters, thus making it possible to stream tiled or stereo displays. Future versions of TeraVision will incorporate multicasting and controls for collaboration, which would enable different parties across the Access Grid to make presentations collaboratively.

TeraVision is a platform and software independent solution for real-time display distribution in advanced collaborative environments. TeraVision is a real-time method to distribute visual imagery from any graphics platform over the Access Grid that requires no setup, software, or hardware changes to a user's computer. The goal is to provide one solution for what is commonly referred to as the "Docking Problem/Display Pushing Problem" on the Access Grid; i.e., to provide a means for anyone on the Access Grid to plug-in their laptop and deliver a presentation without having to install or configure any software, or distribute any of the data files in advance.

GridRAM: Optically Connected Wide Area Network Memory

GridRAM is an application being developed to address long-haul latency in optical networks by collecting memory in a compute cluster and allocating it as a cache to minimize the effects of latency over long-distance, high-speed networks. GridRAM takes advantage of the gigabits of available network on the StarLight and OMNInet testbeds to pre-fetch information before an application is likely to need it (similar to how RAM caches work in a computer today). GridRAM extends this concept over high-speed networks.
<http://www.evl.uic.edu/cavern/teranode/gridram>

EVL student Charles Zhang has started GridRAM testing (UDP blasting) between the UIC Laboratory for Advanced Computing (LAC) cluster and the EVL cluster.

C. Accomplishments

C.1. Meetings

June 21-22, 2002. Tom DeFanti and Maxine Brown attended the Coordinating Committee for International Research Networks (CCIRN) meeting, held annually in conjunction with INET, which was in Washington DC this year. DeFanti gave a presentation on StarLight to the international group of attendees. Note: Kees Neggers is one of the CCIRN organizers, as is Karel Vitsch of DANTE.

May 30-31, 2002. Tom DeFanti, Maxine Brown, Alan Verlo, Linda Winkler, Greg Dawe, and Laura Wolf met with the Cees DeLaat, Paul Wielanga, Ed Mos and others from the Dutch iGrid 2002 planning team for two days of technical meetings. Discussed were application network and computing needs, funding efforts, content (symposia and demonstrations), advertising, registration and staffing needs.

May 8-10, 2002. Tom DeFanti, Maxine Brown, Dan Sandin, Jason Leigh, Vikas Chowdhry, Atul Nayak, Dan Sandin and Greg Dawe attended the Alliance All Hands Meeting at NCSA/UIUC; an event to present the Alliance partner technology development and deployment focus areas. Members of the EVL team participated in the poster session, where they demonstrated the Access Grid Augmented Virtual Environment (AGAVE) system and discussed the StarLight facility, and also participated the Modern Graphics Hardware/Software tutorial.

April 25, 2002. Tom DeFanti, a member of the UCAID Board of Trustees' International Relations Committee, and participated in a telephone conference call on Abilene and connectivity to international Research Networks.

April 23, 2002. Jason Leigh held a videoconference session with Maria Roussos of the Foundation of the Hellenic World in Greece to discuss network requirements for a collaborative virtual reality application currently under development for iGrid 2002.

April 18, 2002. Dr. Ruzena Bajcsy, former head of the Computer Science Directorate at the National Science Foundation, toured EVL and met with Jason Leigh and Tom Moher. Dr. Bajcsy was visiting the campus as an invited Distinguished Lecturer.

April 15, 2002. Jason Leigh and EVL student Shalini Venkataraman met with French artist Jackie Matisse's to discuss development of a physically based art application for iGrid 2002.

April 12, 2002. Jason Leigh presented "Network-Centric Techniques for Advanced Collaborative Environments," to the UIC computer science department as part of his faculty interview.

April 12, 2002. Andy Johnson of EVL attended the First AMPATH International Conference at the Universidad Austral de Chile in Valdivia, Chile. He gave a presentation on StarLight and on iGrid 2002.

March 21, 2002. CERN's Paolo Moroni and Caltech's Sylvain Ravot met with StarLight engineers and supervised the installation and configuration of CERN equipment moving from the Ameritech NAP to the StarLight facility, which now serves as its Chicago PoP.

March 20, 2002. SURFnet's Erik-Jan Bos visited the StarLight facility to supervise equipment installations and meet with Tim Ward, Linda Winkler and Tom DeFanti. Project updates and OC-12c's to the NAP were discussed.

C.2. Publications

E. He, J. Leigh, O. Yu, T.A. DeFanti, "Reliable Blast UDP: Predictable High Performance Bulk Data Transfer," 4th IEEE International Conference on Cluster Computing, Chicago, Illinois, September 2002.

M. Thorson, J. Leigh, G. Maajid, K. Park, A. Nayak, P. Salva, S. Berry, "AccessGrid-to-Go: Providing AccessGrid Access on Personal Digital Assistants," in Proceedings of Access Grid Retreat, La Jolla, California, 2002.

J. Leigh, J. Girado, R. Singh, A. Johnson, K. Park, T.A. DeFanti, "TeraVision: a Platform and Software Independent Solution for Real Time Display Distribution in Advanced Collaborative Environments," in Proceedings of Access Grid Retreat, La Jolla, California, 2002.

C.3. Software Releases

No new software upgrades or releases.

D. Collaboration Activities

Working with SARA to do network performance studies over long, fat networks using various transmission techniques (TCP, UDP, FEC, RUDP).

Working with CERN on RUDP tests. EVL is talking to CERN about DiffServ tests.

E. Problems

No significant problems were encountered this quarter.

F. Any Proposed Changes in Future Plans

No changes to date.