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

NSF Cooperative Agreement No. ANI-9730202
December 2001 Quarterly Status Report

February 18, 2002

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

Linda Winkler, Bill Nickless
Argonne National Laboratory

Jim Williams, Stephen Peck
Indiana University

Table of Contents

A. Significant Results or Events in the Past Quarter

- EVL hosted the NSF CISE "Grand Challenges in e-Science" Workshop, sponsored by ANIR with participation from ACIR
- EVL and UIC National Center for Data Mining researchers demonstrated Terra Wide Data Mining Testbed Project at SC'01
- Jason Leigh and Oliver Yu received a 3-year, $540,000 award from NSF to develop QUANTA, a system to provide an intelligent API to translate high-level data distribution requirements into low-level optimized networking protocols and parameter settings
- SURFnet connected a dedicated 2.5Gb experimental network to StarLight, in addition to two 1Gb production links.
- Ireland's HEAnet connected to STAR TAP

- The testing and documentation phases of EVL-developed UCAN (Unified Collaboratory for Analyzing Networks) were initiated
- An IPv6 Router was installed at StarLight (another remains at STAR TAP)
- RENATER2 upgraded its DS-3 link to OC-3 at STAR TAP (Ameritech NAP)
- NORDUnet upgraded it DS-3 link to OC-3 and now connects to StarLight

B. Expected Results or Events in the Coming Quarter

- The first in a series of iGrid 2002 planning meetings will take place in Amsterdam
- A cluster-to-cluster data transfer (likely a stereo animation) will be sent over the NetherLight/StarLight link and displayed on LCD panel systems at both endpoints (EVL and SARA in Amsterdam)
- CERN will upgrade its OC-3 service to OC-12 and terminate its connection at StarLight
- UCAN network experiments will expand to include more STAR TAP/StarLight member network sites
- The Global NOC will issue an online newsletter and quarterly reports using the Footprints system

C. Summary of Technical Activities

C.1. Euro-Link Network Status and Institutions

C.1.a. CERN

CERN plans to convert its current NAP-based OC-3 service to a connection terminating at the StarLight facility in early 2002. An upgrade to OC-12 using KPNQwest is expected in Spring 2002. The European Union (EU)-funded "DataTAG" link[1], an OC-48 (2.5Gb) from CERN to StarLight, is planned for summer 2002. By late 2002, CERN hopes to upgrade DataTAG to a 10Gbps connection to Chicago, and to connect to the TeraGrid at 2.5Gbps. CERN hopes to have a 10Gbps TeraGrid connection in 2003.

C.1.b. IUCC

Due to budget constraints, Israel will disconnect its STAR TAP connection around April 1, 2002. Since the EU's Géant project pays 50% of the cost of a transatlantic circuit, Israel will connect to Géant via an STM-1 connection in order to connect to the US and Abilene.

C.1.c. NORDUnet

In late September, NORDUnet upgraded its Chicago connection from DS-3 to OC-3 and moved to the StarLight facility. NORDUnet peers with Abilene in New York and brings non-Abilene research traffic to StarLight.

C.1.d. Renater2

In November, RENATER2 upgraded its link to STAR TAP from a DS-3 to an OC-3. This link will soon be upgraded to an OC-12, and Dany Vandromme is considering connecting to StarLight. He is also considering bringing a 2.5Gbps circuit to StarLight.

C.1.e. SURFnet

SURFnet has two OC-12 production links from Amsterdam to StarLight; installation and peering were completed in November. SURFnet peers with Abilene over I-WIRE fiber to the Qwest PoP in Chicago, and uses the STAR TAP/StarLight OC-12 "transit network" to send all its other research traffic to STAR TAP.

In late December, SURFnet's 2.5Gbps link came up. SURFnet put a Cisco ONS 15454 (optical transport platform) at StarLight, which Kees Neggers invites others (e.g., CERN) to plug into, as needed. The UIC National Center for Data Mining (NCDM) 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.

[1] The EU is supporting a 2.5Gb trans-Atlantic link to StarLight with first year matching funds provided by the NSF HPIIS Euro-Link award. The two-year project begins January 1, 2002. DataTAG partners are CERN (prime contractor), Particle Physics And Astronomy Council (PPARC) in the United Kingdom, the Italian National Physics Network (INFN) and the University of Amsterdam in The Netherlands. <http://www.datatag.org>

C.1.f. Non-EuroLink European Connections

Ireland's HEAnet connected to STAR TAP and the STAR TAP Router on October 13.

C.1.g. Others

MREN (Chicago's Metropolitan Research & Education Network) sites are in the process of moving from STAR TAP to StarLight, connecting at 1Gbps speeds.

TransPAC, as of November 1, connects to StarLight at OC-12.

Global Crossing provided AMPATH with 60-day trial 1Gbps (actually 155Mbps with a GigE interface) connection to StarLight, starting December 20, 2001. AMPATH organizers Julio Ibarra and Heidi Alvarez are exploring ways to keep the circuit in place after the trial. The significance of the trial is to showcase scientific research applications of importance to the US and, in particular, NSF.

C.2. Engineering Services

C.2.a. StarLight/Abilene Connectivity

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. (Note: TransPAC peers with Abilene both in Chicago and in Seattle. NORDUnet peers with Abilene in New York. AMPATH's Miami GigaPoP has a 1Gbps connection to StarLight.)

C.2.b. StarLight/STAR TAP Connectivity

There are two OC-12 connections from the StarLight facility to STAR TAP (at the Ameritech NAP) for non-Abilene research traffic. One OC-12 is for SURFnet's use. The second OC-12 link is shared among other StarLight customers.

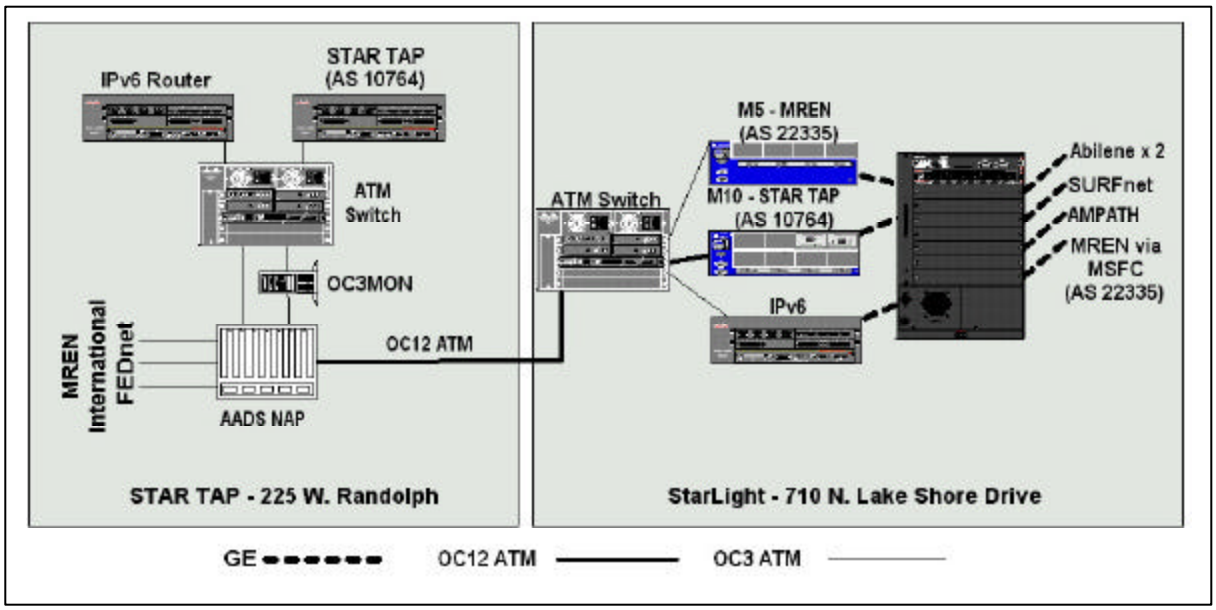

C.2.c. STAR TAP Router

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.

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

In December, STAR TAP engineers moved the NLANR web cache from STAR TAP to StarLight.

C.2.e. 6TAP

ESnet provided a second IPv6 router for the StarLight facility, which was installed in October. The original IPv6 router remains at STAR TAP (Ameritech NAP), enabling non-research networks to be part of the IPv6 testbed.

C.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.

C.2.g. StarLight/STAR TAP Documentation

The StarLight web site was launched October 17. <http://www.startap.net/starlight> Work is ongoing to reflect the growth of StarLight over time.

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

Several STAR TAP members, new and old, are opting to 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). Some, however, continue to bring smaller circuits to Chicago to peer at STAR TAP.

- SingAREN peers with Abilene in California, but is installing a new circuit to STAR TAP.
- Taiwan's TANet2 peers with Abilene in Seattle at the Pacific Northwest GigaPoP ("Pacific Wave") while maintaining a DS-3 link to STAR TAP (completed November 22). Asia Global Crossing-Taiwan Fixed Network Alliance is the new contractor.

C.3. NOC Services

Publication of the first Euro-Link quarterly network report using the new Footprints trouble ticket system has been postponed again. Although data has been collected, the Global NOC reported its software engineering group is in the process of automating and improving the NOC's reporting mechanisms.

D. Euro-Link Performance Analysis Tools

D.1.a. Network Monitoring Tools

Note: Graduate students assist in much of the research and experimentation described in this section. Work was suspended during December due to the semester break.

Bandwidth Utilization Radar Map

Because the STAR TAP engineers are currently moving equipment between the Ameritech NAP and StarLight facilities, we are unable to gather statistics usage at this time. The engineers will reinstate the map once the environments are again stable.

UCAN: Unified Collaboratory for Analyzing Networks

In October, EVL student Naveen Krishnaprasad started the testing and documentation phase of UCAN (formerly uCAN). UCAN is an extendable integrated framework that enables network researchers and application developers to monitor network utilization of applications and obtain network state information at the same time. UCAN is fully collaborative, allowing authorized users to join a collaborative session and remotely control network operations initiated by others in the session.

An internal experimental network was set up at EVL and experiments were conducted to collaboratively monitor network performance of applications, and to study the effect of background traffic. Additions to UCAN are currently being added to CAVERNsoft for its next release. Next steps include completion of API documentation and extending network experiments to a testbed between EVL and other STAR TAP/StarLight member sites.

Vital Statistics Monitor (VitaMon)

EVL student Brenda Lopez is designing a vital statistics monitor (VitaMon) for collaborators to use while running networked applications. It will be a graphical optical traffic map that shows all incoming/outgoing bandwidth among StarLight sites. VitaMon may either employ RTPL or Globus' Network Weather Service to measure end-to-end bandwidth and delay between multiple points connected to STAR TAP.

D.1.b. High-Bandwidth Transmission Over Long Distance Networks

Parallel Socket Experiments

No news to report at this time.

Quality of service Adaptive Networking Toolkit (QUANTA)

In February 2001, NSF awarded a three-year, $540,000 grant to Jason Leigh and Oliver Yu to develop 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.

Reliable Blast UDP (RUDP)

In October, Eric He's revised RUDP code was integrated into the new version of CAVERNsoft. The new classes allow application developers to transfer bulk data over high-speed networks. October RUDP experiments over the CERN-EVL OC-3 link showed that the throughput outperformed Parallel TCP (PTCP) ? 60Mbps throughput from the testbed, as compared to 41Mbps using PTCP.

Forward Error Correction (FEC)

Eric rewrote the FEC code for future integration into CAVERNsoft. Since 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.

D.1.c. Ultra-High-Bandwidth Transmission Over Long Distance Networks (StarLight)

Interrupt Coalescing and Jumbo Frames

Two 16-node Linux clusters are on order; one for the StarLight facility and one for UIC. Engineers expect both to be installed in the first quarter of 2002. 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 A msterdam and Paul Wielinga of SARA 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.

Tera-Vision: Ultra-Resolution Visualization Streaming

In October, EVL obtained a GigE card for its Onyx in anticipation of cluster-to-cluster throughput testing and compressed CAVE video streaming over the 2.5Gb NetherLight/StarLight link to SARA and University of Amsterdam. Vrije University in Amsterdam is also interested in testing parallel cluster-to-cluster simulation and visualization codes. EVL shared information about its cluster with Vrije to maximize compatibility among systems.

In October, Brenda Lopez finished a stereo animation for display on several remote LCD panels as part of Jason Leigh's AGAVE-based Tera-Vision project, proposed for iGrid 2002. Ultra-resolution visualization streaming tests to Northwestern University through the StarLight cluster will commence in January 2002.

Tera-Immersion

In late December, EVL researchers compared the performance of its Onyx2 and its PCs over GigE between EVL and StarLight in preparation of the Tera-Immersion project proposed for iGrid 2002.
<http://www.evl.uic.edu/cavern/teranode/ (click Tera-Immersion)>

Miscellaneous: Thoughts on Parallel Rendering and Load-Balancing Strategies

After attending the IEEE Visualization 2001 conference in San Diego, October 21-26, EVL computer science student Shalini Venkataraman reported the following as relevant to remote visualization over optical networks:

"Our intent to leverage widely available commodity graphics cards on cluster computers linked over high-speed networks for large data visualization poses challenges in two areas: parallel rendering and load-balancing strategies. Various strategies to address these challenges were propounded in the conference.

Parallel rendering strategies can be divided into sort-first and sort-last depending on how geometric primitives are sorted from object-space to screen-space. In contrast to sort-first, sort-last approaches are very scalable with input data size while requiring a high-bandwidth network for image composition. This makes it very amenable to high-speed networks. The disadvantage, however, is the latency caused by the image readback and composition. To address this, Sandia National Labs presented its research on parallel composition strategies ? the Virtual trees and the Tile split and delegate approaches.

The performance of the above mentioned strategies depend critically on balancing the load between the cluster processors. The taxonomy of *load-balancing algorithms* was expounded ? broadly classified as static, dynamic or adaptive depending on how and when the data is partitioned. The adaptive load-balancing algorithms provide the best solution currently but involve dynamic data redistribution at run time. This entails increased communication costs and so is very apt to use in high-speed networks."

E. Accomplishments

E.1. Meetings

December 10, 2001. Linda Winkler, Andy Schmidt and Tom DeFanti met with CERN's Sylvain Ravot to discuss CERN's plans to convert its current NAP based OC-3 service to OC-12 service terminating at StarLight. They also discussed plans for the 2.5Gbps "DataTAG" link to StarLight. <http://www.datatag.org> <http://cern.ch/datatag> <http://cern.ch/datatag/presentations/DataTAG-RN-brochure.ppt>

December 5-6, 2001. The NSF ANIR Workshop on "Grand Challenges in e-Science" was held at EVL to discuss the challenges, applications and funding needs of emerging experimental and commodity networks. For further information, see <http://www.evl.uic.edu/activity/NSF/final.html>

November 7, 2001. Tom DeFanti addressed Korea's NGI fall workshop via videoconferencing, and presented "STAR TAP/StarLight Status and Future Plan."

November 1, 2001. Tom DeFanti met with EVL scientists and affiliated faculty and Northwestern's Joe Mambretti to discuss the visualization cluster location, shipping, networking and software issues.

November 1, 2001. Tom DeFanti discussed changes to the Global NOC service model with IU's Jim Williams.

October 30, 2001. Tom DeFanti gave a presentation via the Access Grid about StarLight to a French delegation visiting NSF in Washington DC. Dany Vandromme, head of RENATER2, was a member of the delegation. Vandromme discussed his intention to increase Renater's bandwidth by bringing in a 622Mb to STAR TAP (perhaps via StarLight), primarily for ESnet and CA*net connectivity. <http://www.france-science.org/presentationppt/fichiers_powerpoint.htm>

October 18-19, 2001. Linda Winkler attended the 10 Gigabit Ethernet Workshop, supported by the NSF and hosted by the San Diego Supercomputer Center. <http://www.sdsc.edu/10GigE>

October 10, 2001. Tom DeFanti, an alternate CENIC Board of Directors (Larry Smarr is the member), attended CENIC board meeting in Los Angeles. Of particular interest was discussion of their Optical Network Initiative (ONI) for the State of California, with interest in connecting to the Pacific Northwest and StarLight.

October 2, 2001. EVL student Atul Nayak attended the CAVE Programming Workshop organized by the Center for Parallel Computing (PDC), Royal Institute of Technology (KTH), in Stockholm, Sweden. He presented a tutorial on CAVERNsoft titled, "Creating Collaborative Virtual Environments (CVE)," in which he described CAVERNsoft G2, a toolkit for high performance tele-immersive collaboration applications for creating CVEs. See: http://www.pdc.kth.se/projects/vr-cube/workshop.html

October 2, 2001. Tom DeFanti and Maxine Brown attended morning plenary sessions of Internet2's Austin-based Fall 2001 Member Meeting via the Access Grid.

E.2. Publications

Krishnaprasad, Naveen, "A Unified and Collaborative Approach for Analyzing Networks," Master's Thesis, Department of Computer Science, University of Illinois at Chicago, November 2001.

E.3. Software Releases

No new software upgrades or releases.

F. Collaboration Activities

Researchers from EVL and UIC's National Center for Data Mining demonstrated the Terra Wide Data Mining Testbed project at SC'01. UIC's four-node visualization cluster was introduced at the conference. GUIs and visualization tools were developed by and coded for the demo by the team, including:

- TeraGUI: the uniform GUI for driving the visualization tools on "TeraNodes" (i.e., small clusters that we have conceived as "on-ramps" to the TeraGrid)
- Tera2D: TeraOverview for the large multi-dimensional dataset and TeraParallel, which provides users with functions such as dimension selection, parallel coordinate rendering and brushing
- Tera3D: implemented with VTK, to provide users with a 3D visualization GUI to rotate, pan and fly into data

Tom DeFanti and Maxine Brown are working with people in the Netherlands' GigaPort Project and SURFnet to organize an iGrid event in Amsterdam in September 2002, to showcase 10Gigabit applications. In September 2001, Maxine sent an invited iGrid 2002 Call for Participation to computational scientists and engineers whose sites will be connected to I-WIRE, DTF and StarLight; this has so far resulted in four proposals and several expressions of interest. [Note: A revised Call for Participation was sent to the general computational science and engineering communities on January 1, 2002.]

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.

G. Problems

No significant problems were encountered this quarter.

H. Any Proposed Changes in Future Plans

No changes to date.

I. Summary of Award Expenditures (October-December)

Available upon request.