



NSF Cooperative Agreement No. ANI-9730202 October 2001 Monthly Status Report

January 7, 2002

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

Linda Winkler
Argonne National Laboratory

Jim Williams, Stephen Peck
Indiana University

Table of Contents

| | | |
|------|---|---|
| A. | Summary of Technical Activities | 1 |
| A.1. | Euro-Link Network Status and Institutions | 1 |
| A.2. | Engineering Services | 2 |
| A.3. | NOC Services | 3 |
| B. | Euro-Link Performance Analysis Tools | 3 |
| C. | Accomplishments | 4 |
| C.1. | Meetings | 4 |
| C.2. | Publications | 5 |
| C.3. | Software Releases | 5 |
| D. | Collaboration Activities | 5 |

A. Summary of Technical Activities

A.1. Euro-Link Network Status and Institutions

A.1.a. CERN

Last month, the DOE/NSF High Energy Physics Transatlantic Network Working Group approved the US/CERN production link timeline outlined in the September 2001 Quarterly Status Report. Preparations are underway to meet the November 2001 deadline to upgrade CERN's one OC-3 to two unprotected OC-3s.

The EU-funded "DataTAG" link is planned for summer 2002; CERN will terminate an OC-48 (2.5Gb) at StarLight. This link will be purely for research, and will be used to carry out various network measurements, demonstrations and experiments. It will support a 2.5Gb trans-Atlantic link to StarLight with first year matching funds provided by the NSF HPIIS Euro-Link award. The project will last two years, with a starting date of 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.

A.1.b. IUCC

Hank Nussbacher notified us that due to budgetary considerations, Israel will most likely disconnect its STAR TAP connection around March 1, 2002 (to be confirmed within a month). Since the EU's Géant project pays 50% of the cost of a transatlantic circuit, Israel will probably order an STM-1 to Géant and maintain its connectivity to STAR TAP and Abilene the same way all of Géant/EU does (backup will be via a commodity ISP in Israel used only when the STM-1 goes down).

A.1.c. NORDUnet

In late September, NORDUnet tripled its research & education network connectivity from DS-3 to OC-3, and now connects to STAR TAP via StarLight.

A.1.d. Renater2

RENATER2's Dany Vandromme issued a tender to upgrade the capacity of both US links (commodity and STAR TAP), and expects an answer to his bid by mid-November (with possible negotiation with the short-listed ones.) He expects to get between 622Mbps and 2.5Gb for 2002, but won't be able to fill it from the outset.

A.1.e. SURFnet

SURFnet's new peering at the StarLight facility in Chicago came up October 11. Its old peering in New York City is still in place; STAR TAP engineer Linda Winkler advised the NOC to monitor both until SURFnet terminates the NYC connection. 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 other research traffic to STAR TAP.

SURFnet currently has one OC-12 from Amsterdam to StarLight in production, provided by Teleglobe, and a second OC-12 from Global Crossing, due to be operational in November. Their 2.5Gb research link, provided by Teleglobe, is imminent.

A.1.f. Non-EuroLink European Connections

Belgium's BELNET and Ireland's HEAnet are now connected to STAR TAP.

A.2. Engineering Services

A.2.a. StarLight/Abilene Connectivity

The Abilene OC-48 core router at the Chicago Qwest PoP is up and passing data in anticipation of SC'01. Abilene currently uses two pairs of I-WIRE fiber to connect to StarLight, one for StarLight participants (e.g., SURFnet) and one for MREN participants. NORDUnet peers with Abilene in New York and TransPAC peers with Abilene in Seattle.

A.2.b. StarLight/STAR TAP Connectivity

STAR TAP pays for two OC-12 connections from StarLight facility to STAR TAP (at the Ameritech NAP). One is for SURFnet's use; the other carries non-Abilene traffic from other research networks located at StarLight.

A.2.c. StarLight/STAR TAP Documentation

The StarLight web site was launched October 17. <<http://www.startap.net/starlight>>

The Networks web page was revised shortly thereafter to distinguish International Advanced Networks (connected and pending) from US Advanced Networks. <<http://www.startap.net/starlight/NETWORKS/>>

A.2.d. STAR TAP Router

No news to report.

A.2.e. 6TAP

ESnet staff has provided a second IPv6 router at the StarLight facility. Once installed, the StarLight community will be able to participate in the national IPv6 testbed operated from Lawrence Berkeley National Laboratory.

A.2.f. DiffServ

No updates to report.

A.3. NOC Services

Publication of the first Euro-Link quarterly network report using the new Footprints trouble ticket system has been postponed to next quarter. 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.

B. Euro-Link Performance Analysis Tools

B.1.a. Network Monitoring Tools

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. This Map has been suspended, and will be reinstated once our 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, also studying the effect of introducing background traffic.

Additions to UCAN are currently being added to CAVERNsoft, for its next release. Next steps include completion of API documentation of UCAN and extending network experiments to a test bed between EVL and other STAR TAP/StarLight member networks.

Vital Statistics Monitor (VitaMon)

EVL student Brenda Lopez is also 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 will employ RTPL to measure end-to-end bandwidth and delay between multiple points connected to STAR TAP.

B.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)

On October 31, the NSF awarded a three-year, \$540,000 grant to Jason Leigh and Oliver Yu to develop QUANTA. The system will provide an intelligent API for application programmers to translate high level data distribution requirements into low level optimized networking protocols and parameter settings. This work is being especially targeted for future high-speed 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 show that the throughput outperforms parallel TCP 60Mbps throughput from the testbed, as compared to 41Mbps using PTCP.

Forward Error Correction (FEC)

Eric also 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 commence testing. The next steps are to conduct more FEC experiments based on the new code, make RUDP responsive to network congestion, and incorporate the protocols into the QUANTA framework.

B.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. Other applications that EVL would like to test over SURFnet include VNC for streaming of clustered desktops, WireGL for streaming OpenGL visualizations and streaming stereoscopic visualizations.

TeraVision: CAVE Video Streaming

Brenda Lopez finished a stereo animation for display on several remote LCD panels as part of Jason Leigh's AGAVE-based Teravision project. Streaming tests to Northwestern University through the StarLight cluster will likely commence in January 2002.

EVL obtained a GigE card for its Onyx in anticipation of cluster-to-cluster throughput testing and graphics streaming to SURFnet. Once SURFnet connects its 2.5Gb link to StarLight, EVL will attempt to stream compressed CAVE video to SARA and University of Amsterdam. Vrije University is also interested in testing parallel cluster-to-cluster simulation and visualization codes. EVL has shared information about its cluster with Vrije to maximize compatibility between systems.

Miscellaneous: Thoughts on Parallel Rendering and Load-Balancing Strategies

After attending 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.”

C. Accomplishments

C.1. Meetings

October 30, 2001. Tom DeFanti met with Renater's Dany Vandromme on the Access Grid. 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.

October 18-19, 2001. Linda Winkler attended the 10-gigabit Ethernet Workshop, supported by the NSF and hosted by the San Diego Supercomputer Center.

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.

C.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.

C.3. Software Releases

No new software upgrades or releases.

D. Collaboration Activities

Tom DeFanti and Maxine Brown are working with people in the Netherlands's GigaPort Project and SURFnet to organize an iGrid event in Amsterdam next September 2002, to showcase 10Gigabit applications. In September, 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.