



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

Submitted September 17, 2001

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	7
C.1.	Meetings	7
C.2.	Publications	7
C.3.	Software Releases	8
D.	Collaboration Activities	8

### A. Summary of Technical Activities

#### A.1. Euro-Link Network Status and Institutions

##### A.1.a. CERN

The European Union agreed to fund the "DataTAG" (Research and Technological Development of a Trans-Atlantic Grid) proposal for an applied Grid network research; CERN and other groups involved in the proposal are now holding negotiations with the EU to finalize budget. This proposal 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 2 years. The starting date is 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.

CERN is preparing an RFP for two circuits between CERN and StarLight. The target date for the DataTAG 2.5Gb lambda, is July 2002. A SONET/SDH 622Mb production circuit, may be connected as soon as mid-March 2002, to support research production traffic between IN2P3 and CERN to vBNS, ESnet, Abilene, etc. (IN2P3 is the National Institute of Nuclear and Particle Physics, a CNRS (French National Center for Scientific Research) institute that currently sends traffic over the CERN link to STAR TAP.) The 2.5Gp network is purely for research, and will be used to carry out various network measurements, demonstrations and experiments.

EVL currently collaborates with CERN on various network transfer protocol tests on its existing 155Mb circuit. EVL student Eric He reports that direct BGP peering between UIC/EVL and CERN is up, and preliminary Reliable Blast UDP (RUDP) bandwidth tests over CERN's 100Mbps link are showing good results. See Section B.1.c.

#### **A.1.b. IUCC**

No activity to report.

#### **A.1.c. NORDUnet**

In Spring 2001, NORDUnet completed a tender for USA connectivity that was to take effect July 1, 2001. NORDUnet was to have an STM-4c (OC-12, 622Mbps) link between Stockholm and the Abilene PoP in New York, and 155Mbps link from New York to STAR TAP in Chicago. NORDUnet switched providers, from Teleglobe to KPNQwest. A series of problems, documented in previous reports, has unfortunately delayed the start date.

On August 2 the OC-3c circuit from Qwest's NY PoP to StarLight was finally working. However, logistical problems with NORDUnet's New York router continues to cause delays. Because NORDUnet had to formally terminate Teleglobe services on August 31, Abilene offered to provide transit between NORDUnet and the US Federal research networks, as well as Abilene international peers (including those connected through STAR TAP), *on a temporary basis*, until problems could be resolved.

#### **A.1.d. Renater2**

RENATER's peering with the STAR TAP Router continues to be disrupted due to France Telecom circuit problems. The STAR TAP technical information table has been amended to reflect this long-term disruption.

#### **A.1.e. SURFnet**

On August 20, Teleglobe performed maintenance on the transatlantic connection from SURFnet in the Netherlands to New York. SURFnet still hasn't closed its New York PoP yet, originally scheduled for July 1.

Teleglobe is in touch with QwestLink Chicago to use its I-WIRE fiber into the StarLight facility at 710 N. Lake Shore Drive for SURFnet's OC-12. STAR TAP engineers expect this will move up the original October due date to mid-to-late August. Teleglobe is also engineering a 2.5Gb lambda from Amsterdam, and expects the connection to take place September 1, as contracted. (Note: Subsequently, it was reported that the link would occur in mid-October.)

### **A.2. Engineering Services**

---

#### **A.2.a. StarLight/Abilene Connectivity**

Abilene engineer Steve Corbató reported Abilene's router is scheduled to arrive at the Qwest PoP August 31, and should be installed and functional by early September. The Abilene OC-48c backbone is expected to be in place by mid-to-late September, at which time it will be ready for the Gigabit Ethernet attachment to the StarLight switch.

#### **A.2.b. StarLight/STAR TAP Connectivity**

STAR TAP contracted for two OC-12 connections from the StarLight facility at 710 N. Lake Shore Drive to the Ameritech NAP. One will be dedicated to SURFnet/Abilene traffic; the second link is to be shared by SURFnet (for non-Abilene traffic) and other NRNs that co-locate at 710. [Note: If Abilene comes to the StarLight facility over I-WIRE fiber, then the second OC-12 may be used instead by MREN universities, who are discussing the possibility of upgrading their MREN connectivity to 1GigE and co-locating at 710.]

#### **A.2.c. StarLight/STAR TAP Documentation**

The StarLight web page is being redesigned to contain more extensive information on our GigE/10GigE/lambda switching exchange; the information will be organized using the same categories as STAR TAP information. Expected launch date is October 1. <<http://www.startap.net/starlight/>>

#### **A.2.d. International Transit Network (ITN)**

International Transit Network participants have been added to the STAR TAP networks web page and map. <<http://www.startap.net/NETWORKS/>>

### A.2.e. STAR TAP Router and Peering

On August 6, STAR TAP engineer Linda Winkler began running Caimis' RouteReporter software on the STAR TAP router <[www.caimis.com/products/te\\_manager/routereporter.shtml](http://www.caimis.com/products/te_manager/routereporter.shtml)>. The software monitors and reports on all BGP routing activity, and provides real-time and off-line analysis of routing performance. If it proves useful, a page displaying the data will be added to the STAR TAP and Euro-Link websites.

Four new OC-3 circuits are being delivered by Teleglobe to STAR TAP

Korea's KISTI and Thailand's UniNet. Since KISTI already handles the KOREN/KREONet2 circuit to STAR TAP, we are currently trying to find out if this is an upgrade to their existing DS-3 or a separate link. HEAnet submitted a proposal to NSF requesting permission to connect to STAR TAP, and it was approved; see <[http://www.heanet.ie/Heanet/projects/nat\\_infrastruct/nsf-star-tap.html](http://www.heanet.ie/Heanet/projects/nat_infrastruct/nsf-star-tap.html)>. To our knowledge, BELnet and UniNet have not yet submitted proposals to NSF requesting permission to connect to STAR TAP, and we are in the process of confirming this.

Brazil's RNP network's DS-3 is ready for testing on the Ameritech end; however, RNP's New York router isn't yet available. The original July 12 connection date of Brazil's Sao Paulo research network FAPESP-ANSP has been pushed back due to FAPESP inter-office delays; 30 days is a conservative estimate. The São Paulo/Brazil Foundation for the Advancement of Research (FAPESP) 155 Mbps link to Miami has been operational since December 2000. NSF approved the FAPESP connection to STAR TAP on December 21, 2000. We are not aware of RNP's having submitted a proposal to NSF for approval, and are in the process of confirming this.

STAR TAP engineer Linda Winkler is continuing to work with Juniper to resolve intermittent M5 router stability problems.

### A.2.f. 6TAP

Bob Fink and Marc Blanchet will provide a second IPv6 Router for 710; eventually, Linda Winkler will remove the one at STAR TAP and backhaul IPv6 traffic from the NAP to 710.

### A.2.g. DiffServ

No updates to report.

## A.3. NOC Services

The NOC will initiate quarterly reports for all networks supported by the Global NOC. <<http://globalnoc.iu.edu>>  
The first Euro-Link quarterly network report using the new Footprints trouble ticket system will be issued at the end of September 2001. It will contain a breakdown of trouble tickets, outage causes, traffic statistics and trend analyses.

The NOC has added a 24x7 NOC staffing chart to its 'Contacts' section of the Euro-Link NOC web page, listing technician schedules. <<http://noc.euro-link.org>>

*Ongoing activities:* The Global Research NOC is preparing to issue a regular, online newsletter in the next few months. John Hicks is working on adapting MIRnet-type traffic graphs for the Euro-Link and TransPAC projects.

## B. Euro-Link Performance Analysis Tools

### B.1.a. Network Monitoring Tools

#### Bandwidth Utilization Radar Map

EVL student Brenda Lopez is designing a vital statistics monitor (VitaMon) for collaborators to use while running networked applications. She is also designing a graphical optical traffic map that will show all incoming/outgoing bandwidth among StarLight sites.

#### uCAN: unified Collaboratory for Analyzing Networks

EVL student Naveen Krishnaprasad continues to work on the unified Collaboratory for Analyzing Networks (uCAN). uCAN enables remote network researchers and application developers to collaboratively execute an application and monitor network utilization, as well as other application-specific parameters. uCAN also allows users to correlate, in real time, how the actions taken by an application directly impact the underlying networks, and

vice versa. A network researcher can also alter router configurations, such as a router's queuing algorithm, to determine how it might improve application throughput.

Development proved to be more difficult than anticipated, however, the resulting solution is general enough that it will be relatively simple to develop future collaborative desktop applications. This capability will be embedded in a future release of CAVERNsoft to allow tele-immersive applications to work with 2D desktop applications.

The first prototype of uCAN has been completed. Work is underway now to setup a testbed to prototype the capabilities of uCAN.

### B.1.b. Network Performance Studies for European/US Collaborative Art Project

EVL PhD graduate Dave Pape developed Yggdrasil (YG), a script-based, authoring environment for networked VR applications, which enables non-programmers to create effective, behavior-rich art and science virtual-reality environments. EVL co-director Dan Sandin is extending the library for behaviors, performing network performance tests and developing applications. Sandin supervised EVL student Joseph Tremonti in the development and execution of network performance tests to Austria, and later Sweden, Hungary and The Netherlands, throughout the Ars Electronica Center's Festival in Austria, September 1-6. [<http://www.aec.at>] A summary of the performance tests will appear in the September monthly report. (See Section D, Collaboration Activities)

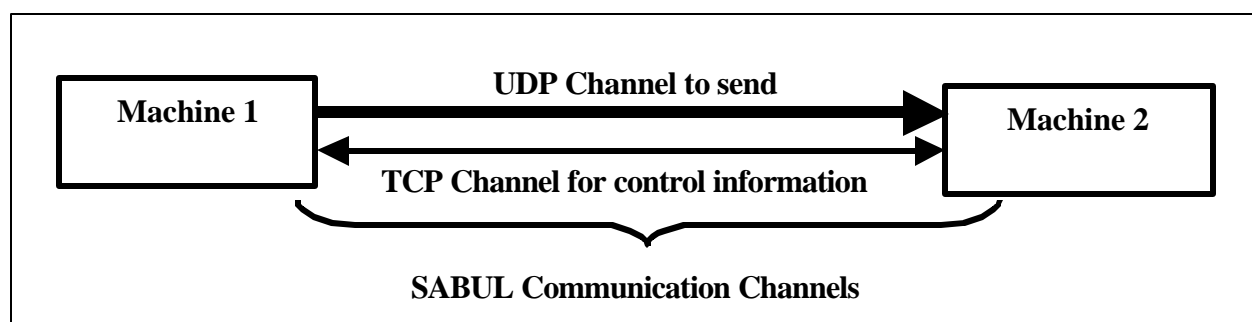
### B.1.c. High-Bandwidth Transmission Over Long Distance Networks

#### Parallel Socket Experiments

Eric He and Jason Leigh are combining their past work on Parallel TCP, Reliable Blast UDP, and Forward Error Correction into a system called QUANTA (Quality of service Adaptive Networking Toolkit) to 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)

UIC/EVL and UIC Laboratory for Advanced Computing's (LAC) Bob Grossman performed parallel TCP and RUDP tests between EVL and SARA in Amsterdam; a UIC Technical Report is available; see Section C.2. SABUL is a C++ library for large data transfers over high-speed wide-area networks. It is a rate-control-based protocol similar to NETBLT, but with two fundamental improvements. The first is the use of both TCP and UDP channels in the protocol. The second improvement over NETBLT and similar implementations of rate-controlled Reliable UDP protocols, is its continuous updating of communication state information, which helps control the data transfer rate to reduce packet loss.



**Figure 1: SABUL Data Transfer**


Figure 1, above, shows that the UDP channel is used to transfer large amounts of data from Sender to Receiver. The Receiver sends the state information of the packet loss and the list of lost packets to the Sender using the TCP channel. The Sender dynamically updates its sending rate based on the packet loss and thereby reaches a data rate which minimizes the packet loss. The SABUL library has been implemented in C++ on a Linux operating system. It has now been ported to various Unix platforms (Solaris, AIX, Irix, and Free BSD).

**Experimental Results.** The SABUL library was utilized for data transfer between two Linux nodes at Ann Arbor, MI and NCAR, Boulder, CO. Each node had 256 MB of RAM and fast Ethernet cards. The machines were interconnected through Internet2's Abilene network with an OC-3 uplink (155 Mb/s). Table 1 shows the results of our tests along with results obtained using Iperf, a network performance tool and PSocket, a library that stripes data across multiple TCP sockets. While using Iperf the network was tuned with an optimal TCP window size of 512.5 KB. SABUL was tuned to have a maximum packet loss of one percent.


Node Location	Library/Utility used	Throughput in Mb/s	Packet loss %
Boulder to Ann Arbor	Iperf with network tuning	83.3	0.0
	PSockets (19 sockets)	85.27	0.085
	SABUL	92.65	0.032
Ann Arbor to Boulder	Iperf with network tuning	10.1	0.11
	PSockets (19 sockets)	40.68	0.57
	SABUL	63.87	1.1

**Table 1: Performance of SABUL over Abilene network**

Table 1 shows that SABUL is able to achieve higher throughput than that of Iperf with TCP window tuning as well as PSocket. At the same time, SABUL was able to maintain the packet loss to be less than one percent. Figures 2 and 3 show the change in data rate as well as packet loss rate as SABUL tunes itself to have a packet loss of one percent. These graphs are results of runs between Ann Arbor to Boulder.



**Figure 2: SABUL date rate tuning**



**Figure 3: Packet loss percentage in SABUL**

Node Location	Library/Utility Used	Throughput in Mb/s	Packet Loss %
Amsterdam to Ann Arbor	Iperf with network tuning	39.4	0.27
	SABUL	90.47	0.01
Ann Arbor to Amsterdam	Iperf with network tuning	20.0	0.62
	SABUL	39.0	0.6
UIC (EVL) to Amsterdam	Iperf with network tuning	58.04	0.12
	SABUL	70.66	0.15
Amsterdam to UIC (EVL)	Iperf with network tuning	14.88	0.68
	SABUL	16.83	2.26

**Table 2: Performance of SABUL over an OC-3 link**

Table 2 compares the results of throughput and packet loss obtained while sending or receiving data from nodes at Ann Arbor, MI and Chicago, IL to a node in Amsterdam, Netherlands. It can be clearly seen that the throughput performance of SABUL is much superior to that of a single TCP socket tuned with appropriate TCP window size.

**Conclusion.** We are not proposing that the days of TCP are numbered. Due to its widespread acceptance and many refinements over the years, it should still serve well as a general-purpose reliable transport protocol. We see SABUL as a more specialized protocol serving the needs of applications that require large data transfers on high-bandwidth networks over long distances.

### B.1.d. Ultra-High-Bandwidth Transmission Over Long Distance Networks

In anticipation of SURFnet's 2.5Gb connection to STAR TAP in September/October, Jason Leigh has been speaking with Cees de Laat (SURFnet and University of Amsterdam), Paul Weilinga (SARA) and Henri Bal (Vrije University, Amsterdam) to develop ideas for testing bandwidth-intensive applications over SURFnet. SARA and SURFnet are both interested in attempting to stream compressed CAVE video. EVL has ordered a GigE card for its Onyx in preparation for this. Vrije University is 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. Bob Grossman, director of UIC's National Center Data Mining (NCDM), wants to install a data-mining server in Amsterdam and performing tera-mining queries between Chicago, Amsterdam and Canada.

EVL is in the process of building and optimizing its Linux cluster to maximize throughput over GigE networks. 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 been able to achieve 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.

## C. Accomplishments

---

### C.1. Meetings

---

**September 4, 2001.** A StarLight Media Announcement brainstorming meeting was held at EVL. Joe Mambretti (NU) and Andy Schmidt (EVL) gave Tom Garritano (NSF) and Paul Francuch (UIC) a tour of the StarLight facility before the meeting. Also present were EVL's Tom DeFanti, Maxine Brown, Laura Wolf and Andy Schmidt, and

**August 27, 2001.** Jason Leigh, University of Amsterdam's Cees de Laat, Joe Mambretti, Linda Winkler and Andy Schmidt met at the StarLight facility to discuss STARLIGHT application requirements and future network research ideas.

**August 23-24, 2001.** Tom DeFanti, Maxine Brown and Linda Winkler attended the NLANR HPIIS Workshop <<http://moat.nlanr.net/Workshops/HPIIS-2001/>>.

**August 15-17, 2001.** Tom DeFanti and Maxine Brown attended the FIU AMPATH Workshop to Identify Areas of Scientific Collaboration between the US and the AMPATH Service Area, held at the Florida International University Kovens Conference Center at Biscayne Bay Campus <<http://www.ampath.fiu.edu/events.htm>>. Tom DeFanti was one of several interviewed by Discovery Channel Latin America for the television documentary *Vida @ Linea*.

**August 13, 2001.** Chip Cox of NSF ANIR visited EVL to learn more about current STAR TAP, Euro-Link and StarLight activities. Those attending this meeting included Tom DeFanti, Maxine Brown, Jason Leigh, Andy Schmidt, Linda Winkler, Bob Grossman and Joe Mambretti.

**August 10, 2001.** Tom DeFanti and Bob Grossman discussed StarLight equipment needs for an upcoming NSF RI proposal we wish to submit.

**August 2, 2001.** EVL's Jason Leigh hosted the Canadian Communications Research Centre's John Spence and Eric Tsang to discuss possible VR/tele-immersion and Access Grid collaborations. The Centre, a Canadian federal government laboratory, is one of the main testbeds in Canada supporting CANARIE and the CA\*Net 3.

### C.2. Publications

---

R.L. Grossman, M. Mazzucco, Y. Pan, H. Sivakumar, Q. Zhang, "Simple Available Bandwidth Utilization Library (SABUL) for High-Speed Wide Area Networks," University of Illinois at Chicago, Department of Mathematics, Statistics and Computer Science, Laboratory for Advanced Computing, Technical Report, 2001.

### C.3. Software Releases

---

No new software upgrades or releases.

### D. Collaboration Activities

---

Tom DeFanti and Maxine Brown are talking to people in the Netherlands's GigaPort Project and SURFnet5 about hosting an iGrid event in Amsterdam next September 2002, to showcase 10Gigabit applications. The iGrid 2002 invited call for participation Maxine sent to scientists whose sites will be connected to I-WIRE, DTF and StarLight, has so far resulted in four proposals and several expressions of interest.

Dan Sandin and several EVL students will present a collaborative, networked VR environment at the Ars Electronica Festival, September 1-6, in Linz, Austria. Participants include artists from Hungary's C3 laboratory, The Netherlands' V2 laboratory, The Interactive Institute of Sweden, and the United States (UIC, Chicago and SUNY, Buffalo).

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

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