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

http://www.euro-link.org

High Performance International Internet Services between Research and Education Institutions in the United States and Europe/Israel

NSF Cooperative Agreement No. ANI-9730202
Annual Report (April 1, 2002–March 31, 2003)

SUBMITTED May 19, 2003

Table of Contents

1. Participants

1.A. Primary Personnel

Participant's Name(s)	Project Role(s)	>160 Hours/Yr
Thomas A. DeFanti	Principal Investigator	Yes
Maxine Brown	Co-Principal Investigator	Yes
Andrew E. Johnson	Co-Principal Investigator	Yes
Daniel J. Sandin	Co-Principal Investigator	Yes

1.B. Other Senior Personnel (Excluding PI and Co-PI)

Additional people who contributed greatly to the project, and received a salary, wage, stipend or other grant support:

Participant's Name(s)	Project Role(s)	>160 Hours/Yr
Jason Leigh	Senior Personnel/Professional Staff	Yes
Michael McRobbie	Senior Personnel/Professional Staff	Yes
Doug Pearson*	Senior Personnel/Professional Staff	No
Jim Williams	Senior Personnel/Professional Staff	Yes
Linda Winkler+	Professional staff	Yes
Bill Nickless+	Professional staff	Yes
Alan Verlo	Professional staff	Yes
Laura Wolf	Professional staff	Yes
Patrick Hallihan	Professional staff	Yes

* Doug Pearson was listed as Senior Personnel on the proposal as he was in charge of the Euro-Link NOC at Indiana University. He has been replaced by Jim Williams at Indiana University

\+ Linda Winkler and Bill Nickless, while not compensated by the University of Illinois at Chicago, serve as part-time STAR TAP/StarLight/Euro-Link engineers.

1.C. Other Organizations That Have Been Involved as Partners

1.C.1. European National Research Networks (NRNs)

CERN (EUROPEAN LABORATORY FOR PARTICLE PHYSICS)

The European Laboratory for Particle Physics (CERN) <http://cern.ch> provides experimental facilities for particle physics experiments, mainly in the domain of high-energy physics (HEP).

CERN's current facility is the Large Electron Positron (LEP) collider in a 27km tunnel, the largest machine of this type in the world. Experiments are carried out by teams of hundreds of physicists from over 50 institutes spread across five continents. The next particle accelerator, to be completed in 2005, is the 14 TeV (1 Tera electron volt = 1 billion electron volts) Large Hadron Collider (LHC); it is being built using high-powered 14-meter superconductor magnets and will be installed in the existing LEP tunnel.

LEP experiments generate 25 TBs of data each year, which is stored on magnetic tapes. The LHC experiments are expected to produce several-orders-of-magnitude more data, namely tens of petabytes per experiment each year. The sheer volume of data combined with the complexity of the analysis to be performed, and the requirement that the data processing may also be done remotely, places heavy demands on the High Energy and Nuclear Physics (HENP) computing and networking infrastructure, which can only be met by using leading-edge technology and services.

As of March 2002, CERN has a 622Mb connection to StarLight. In addition, CERN manages the **European Union (EU) DataTAG** 2.5Gb circuit from Geneva to StarLight. The two-year EU project began January 2002; NSF matching funds are provided by the NSF HPIIS Euro-Link award. 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>. The EU DataTAG link was

successfully provisioned by T-Systems in August 2002. CERN plans to upgrade the EU DataTAG link to 10Gb, and to connect to the TeraGrid at 2.5Gb (and then upgrade to 10Gb), in 2003.

ISRAEL INTER-UNIVERSITY COMPUTATION CENTER (IUCC) (Internet-2/Ilan-2)

Israel's Inter-University Computer Center (IUCC) operates the Israel Academic Network (ILAN) <http://www.internet-2.org.il>. IUCC was created by eight Israeli universities, and is funded by the Committee for Planning and Budget of the Council for Higher Education.

In May 2002, Israel disconnected its STAR TAP connection and now no longer receives Euro-Link HPIIS funding. Because the European Union's GÉANT project pays 50% of the cost of a transatlantic circuit, Israel took advantage of its GÉANT connection to connect to Abilene in New York. ILAN-2 peers with GÉANT via two STM-1 (155Mb) lines that became operational in May and June 2002. One terminates in Milan, the other in Amsterdam. An additional 45Mb circuit is used as a backup when connectivity to GÉANT is down. IUCC connects with Abilene, ESnet and CANARIE via GÉANT's 2*2.5Gb connection to the USA.

NORDUNET

NORDUnet <www.nordu.net> serves the universities and publicly funded research institutions in Denmark, Finland, Iceland, Norway and Sweden. The national and international bandwidths are in the Gigabit range. The five national networks collaborate in the commonly owned NORDUnet to get international connectivity among themselves and the rest of the world, including connections to the general purpose commercial Internet via gateways and peerings.

NORDUnet's backbone is 2.5Gb. Core nodes in Stockholm, Oslo and Copenhagen form a network ring, giving the necessary resiliency in case of link failure. Denmark, Finland and Norway are all doubly connected to the NORDUnet backbone at 2.5Gb. Iceland's connection to NORDUnet was upgraded to 155 Mb in January 2002, replacing an overloaded 45Mb connection set up in September 2000. The total international bandwidth connected to NORDUnet's central node at KTH in Stockholm is now approximately 10Gb. NORDUnet also serves networks for research and education in Poland (NASK/12Mb), Estonia (EEnet/16Mb), Ukraine (UARnet/1Mb), and Russia (RUNNet/34Mb).

Currently NORDUnet brings a 622Mb link to New York, where it peers with Abilene. From New York, it brings a 155Mb link to StarLight, where it peers with other international networks and ESnet. NORDUnet wants to maintain this connection until August 2003, during which time it will intensify and reinvigorate Nordic/American research collaboration to justify a 2.5Gb wavelength to either StarLight and/or NetherLight in Summer 2003.

RENATER2

GIP RENATER <www.renater.fr> is a non-profit organization owned by five large research and education groups in France. Users of RENATER2, the national backbone of GIP RENATER, are advanced scientific laboratories and higher education institutions. RENATER2 is a nation-wide infrastructure connecting campuses, metropolitan and regional networks with IP and ATM services, and has international links to other European countries and a 155Mb connection to STAR TAP.

RENATER2 has a mesh topology, with 2.5Gb, 622Mb, 155Mb or 34Mb links between its central switching node and the regional PoPs. The central loop in the Paris area has a throughput capacity of 80Gb.

SURFNET

SURFnet is the national network for research and education in the Netherlands <www.surfnet.nl>. It receives funding from GigaPort, a joint project of the Dutch government, trade and industry, educational institutions and research institutes <www.gigaport.nl>.

The SURFnet5 network backbone is situated at two locations in Amsterdam for redundancy. Fifteen PoPs are connected to both backbone locations over two separate routes, ensuring one connection is always maintained in

case of a single line disruption. Nineteen routers exist within the network, four inside the collapsed backbone (core routers) and fifteen at the concentrator locations (connection routers).

In early 2002, SURFnet closed its New York PoP and brought up an OC-12 production link and a 2.5Gb research link between Amsterdam and StarLight. In September 2002, Level 3, which provisioned the 2.5Gb link, donated an additional 10Gb link from Amsterdam to Chicago for iGrid 2002. In October 2002, SURFnet negotiated to keep the 10Gb link and turned off its 2.5Gb link. In January 2003, SURFnet discontinued its OC-12 and selected T-Systems to replace Level 3 as the provider of the 10Gb link. SURFnet also has access to a 10Gb link from Amsterdam to New York donated by IEEAF/Tyco.

Starting in 2000, SURFnet began to construct **NetherLight**, a sister facility to the NSF-funded StarLight in Chicago. NetherLight, located at the Amsterdam Internet Exchange facility on the campus of the Amsterdam Science & Technology Centre, is an advanced optical infrastructure and proving ground for network services optimized for high-performance applications. Operational since summer 2001, NetherLight is a multiple Gigabit Ethernet (GigE) switching facility for high-performance access to participating networks and will ultimately become a pure lambda switching facility for wavelength circuits, as optical technologies and their control planes mature. NetherLight's international connectivity includes dedicated lambdas to the StarLight facility in Chicago and to CERN in Switzerland. On a national scale, SURFnet connects ASTRON/ JIVE in the region of Dwingeloo in northern Holland (ASTRON is the Netherlands' Foundation for research in astronomy and JIVE is the Joint Institute for VLBI [Very Long Baseline Interferometry] in Europe) to NetherLight by means of a 32-wavelength Dense Wave Division Multiplexing (DWDM) transport network.

Researchers use the NetherLight facility to investigate novel concepts of optical bandwidth provisioning and to gain experience with these techniques. In particular, researchers are investigating different scenarios on how lambdas can be used to provide tailored network performance for demanding grid applications. Important issues are: how to get traffic onto and out of lambdas; how to map load on the network to a map of lambdas; how to deal with lambdas at peering points; how to deal with provisioning when more administrative domains are involved; and, how to do fine-grain, near-real-time grid application-level lambda provisioning <www.science.uva.nl/~delaat/optical>.

1.C.2. US Management Team

ELECTRONIC VISUALIZATION LABORATORY (EVL), UNIVERSITY OF ILLINOIS AT CHICAGO (UIC)

EVL <www.evl.uic.edu>, over the past decade, has teamed with network engineers, computer scientists and computational scientists to collect, maintain, develop, distribute, and evaluate "tele-immersion" (virtual-reality over high-speed networks) tools and techniques for scientific computing. For Euro-Link, this expertise is being used to design, build, implement and maintain a new generation of sophisticated real-time network performance monitoring and persistent database tools tuned to application uses of high-performance networks.

INDIANA UNIVERSITY

Indiana University <www.indiana.edu> has a NSF HPIIS award for the design and deployment of the TransPAC consortium <www.transpac.org>, the high-speed backbone that interconnects research institutions in Japan, Korea, and other Pacific Rim countries that are part of the Asian Pacific Advanced Network (APAN) consortium with US institutions connected to the vBNS+, Abilene and the Fednets. Indiana also provides Global NOC services for Abilene, TransPAC, NaukaNet (formerly MIRnet), Euro-Link, STAR TAP/Star Light and AMPATH.

MATHEMATICS AND COMPUTER SCIENCE DIVISION (MCS), ARGONNE NATIONAL LABORATORY (ANL)

ANL <www.mcs.anl.gov> has been, and continues to be, a strong supporter of STAR TAP and Euro-Link activities. Linda Winkler has facilitated STAR TAP/StarLight engineering since its inception, and is the lead engineer today; her salary comes from ANL. Rick Stevens, director of the ANL Math and Computer Science Division, is the chair of STAR TAP's International Advisory Committee.

AMERITECH ADVANCED DATA SERVICES (AADS)

The AADS Network Access Point (NAP) <www.aads.net> is a major Internet Exchange Point in Chicago where ISPs meet to exchange traffic with other attached ISPs. The NAP is a large ATM switch providing both high speed and a high degree of scalability. Ameritech and STAR TAP have been models for next-generation NGI eXchanges (NGIXs).

METROPOLITAN RESEARCH AND EDUCATION NETWORK (MREN)

MREN <www.mren.org> is a regional network connecting Midwest-area research institutions. It is a model for GigaPoPs, or other regional networks. MREN is one of the world's most advanced high-performance broadband networks, developed to support a wide range of advanced research applications requiring high performance and high bandwidth. UIC is one of the founding members of MREN. Joe Mambretti, director of the iCAIR center at Northwestern University, is the director of MREN and Linda Winkler of ANL is MREN's technical director. In March 2002, MREN connected to the StarLight switch via Gigabit Ethernet.

INTERNATIONAL CENTER FOR ADVANCED INTERNET RESEARCH (iCAIR), NORTHWESTERN UNIVERSITY

The goal of iCAIR <www.icair.org>, under the leadership of Joe Mambretti, is to accelerate leading-edge innovation and enhanced global communications through advanced Internet technologies, in partnership with the international community. iCAIR was established to provide a focal point for leading-edge Internet research, innovation, and pre-production deployment. iCAIR is an international research and development center that creates large-scale, advanced digital communication systems based on Internet protocols, in part, by utilizing regional, national and international advanced research networks. iCAIR is very involved in the Digital Video Working Group of the Coordinating Committee for International Research Networks (DVWG, CCIRN) as well as the Internet2 Digital Video (I2-DV) Initiative, and works closely with several Euro-Link NRNs.

STAR TAP

STAR TAP <www.startap.net>, managed by UIC/EVL in collaboration with the Mathematics and Computer Science Division at Argonne National Laboratory (ANL) and operated by SBC/Ameritech Advanced Data Services, was initially created in 1997 to provide a persistent infrastructure for the long-term interconnection and interoperability of advanced international networking in support of applications, performance measuring, and technology evaluations. By the year 2000, STAR TAP successfully became a model for Next-Generation Internet eXchanges (NGIXs), so attention focused on ways to maintain and improve functionality. StarLight, the optical STAR TAP, is a direct outcome of working with the international research community to architect the networking component of a 21st century advanced scientific cyberinfrastructure. STAR TAP receives major funding from NSF awards ANI-9980480 and ANI-9712283 to UIC, and DOE funding to ANL.

STARLIGHT

Star Light <www.startap.net/starlight>, the optical STAR TAP, is managed by UIC/EVL, NU/iCAIR and ANL/MCS. It is located on the Northwestern University campus in downtown Chicago. StarLight is an advanced optical infrastructure and proving ground for network services optimized for high-performance applications. StarLight is being developed in partnership with Canada's CANARIE and Holland's SURFnet. StarLight connects to STAR TAP at the AADS NAP in Chicago via an OC-12 link. The major advantages of StarLight are space and collaboration opportunities for Gigabit-enabled applications, optical switching and, eventually, optical routing. StarLight is implementing a number of advanced layer 2 and layer 3 services of benefit to the Euro-Link community. StarLight receives major funding from NSF ANI-0229642 to UIC and NU, and DOE funding to ANL.

1.C.3. Consortia

Consortia of American universities and National Research Networks, which serve to facilitate connectivity to U.S. high-performance network service providers, such as the vBNS+, Abilene or the Fednets. Euro-Link, NaukaNet and

TransPAC are funded in part by the NSF High Performance International Internet Services (HPIIS) program. AMPATH also receives NSF funding.

AMPATH

AmericasPATH, or AMPATH,<www.ampath.fiu.edu> is a Florida International University (FIU) and Global Crossing (GC) collaborative project to interconnect the Research and Education networks of South and Central America, the Caribbean, and Mexico, to networks in the US and other countries.

NAUKANET (formerly FASTnet/MIRnet)

NaukaNet <http://www.naukanet.org/friends/naukanet/home.html> is an NSF HPIIS program with matching support from the Ministry for Science and Technology of the Russian Federation. The consortium includes National Center for Supercomputing Applications (NCSA) at the University of Illinois, the Russian Institute for Public Networks (RIPN), Moscow State University (MSU), Friends and Partners-Russia (F&P/R), the VUZTelecom Center of St. Petersburg, and other key Russian and US organizations. In September 2002, a 155Mb Telia-provisioned circuit between Moscow and StarLight became operational.

TRANSPAC

TransPAC <www.transpac.org> is a HPIIS-funded consortium of Indiana University and the Asian-Pacific Advanced Network Consortium, or APAN, which includes Australia, Japan, Korea and Singapore.

1.D. Other Collaborators or Contacts

ALLIANCE AND NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS (NCSA), UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

The NSF Partnerships for Advanced Computational Infrastructure (PACI) Cooperative Agreement to the National Computational Science Alliance (the "Alliance") <http://alliance.ncsa.uiuc.edu>, whose lead institution is the National Center for Supercomputing Applications (NCSA), funds, in part, UIC/EVL to deploy research results in virtual reality, networking, visual supercomputing, distributed computing and networked collaboration. (PACI does not fund basic research; it assumes partners already *have* research results and are funded by peer review in their disciplines.) NCSA/Alliance director Dan Reed is very supportive of STAR TAP and StarLight.

CALIFORNIA INSTITUTE FOR TELECOMMUNICATIONS AND INFORMATION TECHNOLOGY [Cal-(IT)2], UNIVERSITY OF CALIFORNIA SAN DIEGO (UCSD)

Cal-(IT)2 <www.calit2.net>, founded by Larry Smarr, is a distributed center, conducting research at both the UCSD and UC-Irvine (UCI) campuses. UCSD and UCI conduct research in core technologies needed to expand the reach and capacity of the global wireless Internet and its emerging all-optical core. Initially, it will use the new telecommunications infrastructure to advance applications important to California's economy, including education, environmental monitoring, health care delivery, transportation, and new media arts, but has ambitions to expand collaborations internationally. EVL currently partners with Cal-(IT)2 on the NSF-funded OptIPuter project. Cal-(IT)2 fully encourages the goals of STAR TAP/StarLight.

CAVERNUS

The CAVE Research Network User's Society (CAVERNUS) <www.cavernus.org> gives the worldwide community of virtual-reality device users a place to share ideas, solutions and discoveries as they interconnect over high-speed networks. The universities, research laboratories and commercial R&D facilities supporting this web site are also the primary users of EVL's CAVERNsoft and Quanta software toolkits. CAVERNUS hosts virtual-reality workshops and Birds of a Feather (BoF) or Special Interest Group (SIG) meetings at major conferences. As EVL continues to enhance CAVERNsoft and Quanta, we conduct network and visualization experiments with international collaborators via STAR TAP and StarLight.

GLOBUS AND THE GLOBAL GRID FORUM

Globus <www.globus.org> is a joint project of University of Chicago/ANL and the University of Southern California's Information Sciences Institute, with significant contributions from NCSA, NASA Ames, the Aerospace Corporation, and other partners. The Globus project is developing basic software infrastructure for computations that integrate geographically distributed computational and information resources. The Globus leadership recently founded the Global Grid Forum (GGF) <www.globalgridforum.org>, a community-initiated forum of individual researchers and practitioners working on distributed computing, or Grid, technologies; GGF participants come from over 150 participating organizations, with financial and in-kind support coming from sponsoring companies and institutions.

GRID PHYSICS NETWORK (GriPhyN)

GriPhyN <www.griphyn.org> is a team of experimental physicists and information technology (IT) researchers who plan to implement the first Petabyte-scale computational environments for data-intensive science in the 21st century. Driving the project are unprecedented requirements for geographically dispersed extraction of complex scientific information from very large collections of measured data. To meet these requirements, GriPhyN will deploy computational environments called Petascale Virtual Data Grids (PVDGs) that meet the data-intensive computational needs of a diverse community of thousands of scientists worldwide. While GriPhyN is an NSF-funded program under the leadership of Paul Avery of the University of Florida and Ian Foster of University of Chciago/ANL, the focus is on accessing data generated from CERN experiments. Hence, STAR TAP, and now StarLight, is of utmost importance to GriPhyN, and the European Union (EU) funded counterpart, the European DataGrid <www.eu-datagrid.org> project.

DataTAG

The DataGrid <www.eu-datagrid.org> is a European Union (EU) funded project to enable next-generation scientific exploration that requires intensive computation and analysis of shared large-scale databases across distributed scientific communities. The EU-funded DataTAG (Research and Technological Development of a Trans-Atlantic Grid) <www.datatag.org> initiative supports a 2.5Gb transatlantic link to StarLight to enable European and US grid projects to collaborate. (See CERN, Section 1.C.1.)

NETWORK FOR EARTHQUAKE ENGINEERING SIMULATION

The NEESgrid integration project is being organized by NCSA and funded by the NEES program at NSF <www.neesgrid.org, www.eng.nsf.gov/nees>. The goal of NEESgrid is to design and integrate experimental and computing and communications facilities for use by the earthquake engineering community. EVL collaborates with NEESgrid to deploy visualization and collaboration technologies. DeFanti is a member of the NEESgrid Executive Advisory Committee.

2. Activities and Findings

2.A. Research Activities

2.A.1. Euro-Link Goals and Statement of Work

Euro-Link is a consortium of the University of Illinois at Chicago (UIC) and several National Research Networks (NRNs) in Europe and Israel, formed to assist the US and European research and education communities with High-Performance International Internet Services (HPIIS). The NSF, through its HPIIS program, encourages NRNs to connect to high-performance internet service providers, such as the vBNS+ and the University Corporation for Advanced Internet Development (UCAID) Internet2's Abilene, via STAR TAP[1] and now StarLight[2].

European NRNs received NSF approval to connect to STAR TAP before they became members of the Euro-Link consortium. When the Euro-Link proposal was submitted to NSF in January 1999, there were four Euro-Link *charter NRNs*, NORDUnet, SURFnet, RENATER2, and Israel's IUCC. In late 1999, CERN received NSF permission to connect to STAR TAP and to join the Euro-Link consortium. In 2002, Israel/IUCC withdrew from Euro-Link due to budgetary constraints; rather than pay for their own link to the US, they now use GÉANT.

The NSF "Euro-Link" Cooperative Agreement has a clear Statement of Work that defines the goals and objectives of this consortium. The activities that support this Statement of Work, as well as additional activities, are documented in this section; they are:

- ***Furnish, operate, and maintain a direct connection for high-performance traffic between the vBNS+ and Abilene and European NRNs via the STAR TAP, and now StarLight (Section 2.A.2)***
- ***Ensure, to the extent supportable by prudent application of networking technology, that only approved institutions' traffic is permitted to use the high-performance connection (Section 2.A.3)***
- ***Monitor the network performance and use of the Euro-Link connections (Section 2.A.4)***
- ***Develop a new generation of performance analysis tools (Section 2.A.5)***
- ***Provide engineering services, both testbeds and production implementations, of new networking protocols (such as IPv6, Multicast, cache services, QoS) (Section 2.A.6)***
- ***Establish a Euro-Link Network Operations Center (Section 2.A.7)***
- ***Maintain a publicly-accessible Euro-Link-HPIIS web site (Section 2.A.8)***
- ***Support the use of Euro-Link-HPIIS for high-performance applications (Sections 2.A.9)***

[1] The Science, Technology And Research Transit Access Point (STAR TAP) <www.startap.net> was initially created in 1997 to provide a persistent infrastructure for the long-term interconnection and interoperability of advanced international networking in support of applications, performance measuring, and technology evaluations. By the year 2000, STAR TAP successfully became a model for Next-Generation Internet eXchanges (NGIXs), so attention focused on ways to maintain and improve functionality. STAR TAP was developed by the Electronic Visualization Laboratory at the University of Illinois at Chicago (UIC) and the Mathematics and Computer Science Division at Argonne National Laboratory (ANL), in partnership with SBC/Ameritech Advanced Data Services. STAR TAP receives major funding from NSF awards ANI-9980480 and ANI-9712283 to UIC, and DOE funding to ANL.

[2] StarLight <www.startap.net/starlight>, the optical STAR TAP, is a persistent infrastructure that supports advanced applications and middleware research, and aggressive advanced networking services. StarLight is a multi-vendor 1 Gb, 2.5 Gb, and 10 Gb experimental switching facility, serving as a nodal point for the "other end" or switching hub for national and international experiments. StarLight will ultimately become an anchor for wavelength-rich LambdaGrids, with switching and routing at the highest experimental levels, laying the foundation for fully-optical switching. It is being developed UIC, ANL and the International Center for Advanced Internet Research at Northwestern University, in partnership with Canada's CANARIE and The Netherlands' SURFnet. StarLight receives major funding from NSF award ANI-0229642 to UIC and NU, and DOE funding to ANL.

2.A.2. European NRN Connectivity to STAR TAP/StarLight

Each NRN procures and operates its own transatlantic services as an extension of its national services. Euro-Link provides engineering support and helps defray the costs of connecting to STAR TAP/StarLight.

CERN/DataTAG (one 622Mb; one 2.5Gb)

CERN's transatlantic connections to STAR TAP went from 12Mb in April 1999 to 21Mb in September 1999, 45Mb in April 2000 to 155Mb in January 2001, then 2*155Mb in December 2001. In late March 2002, CERN connected to StarLight via a 622Mb link, but was forced to temporarily transit its traffic via GÉANT when provider KPNQwest went out of business. In September 2002, CERN connected a new 622Mb link between Chicago and Geneva/CERN provisioned by Deutsche Telekom/T-Systems, and established connectivity with StarLight's IPv6 infrastructure.

In August 2002, a 2.5Gb circuit between CERN and StarLight was delivered on behalf of the European Union (EU) DataTAG project. DataTAG supports an intercontinental Grid testbed involving the European DataGrid project, several national projects in Europe, and related Grid projects in the US. 10GigE tests using beta-version Intel NICs in clusters between CERN, StarLight, Caltech and Sunnyvale (SLAC) recently took place.

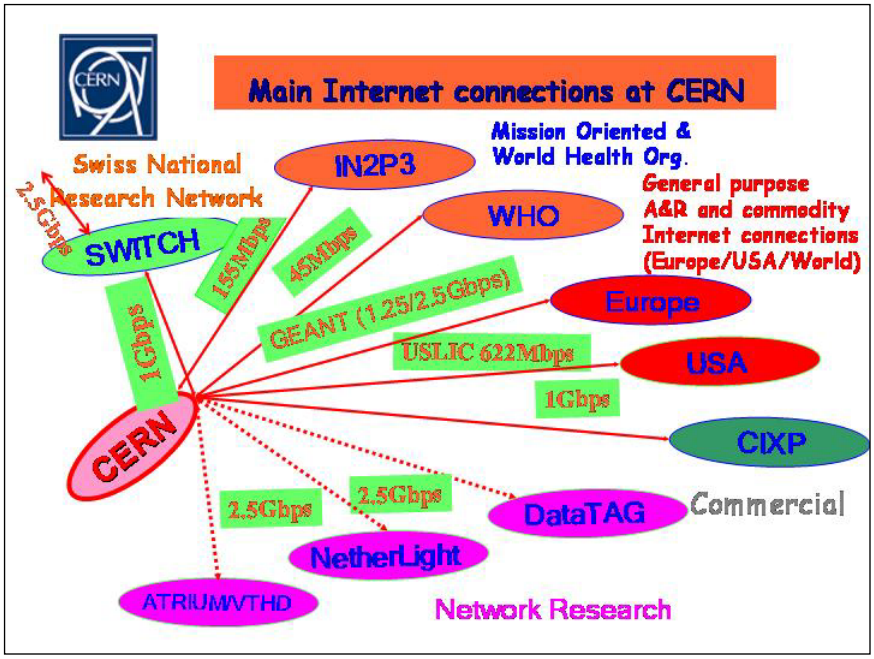

CERN plans to consolidate its two circuits to StarLight (622Mb for production traffic and 2.5Gb for DataTAG traffic) onto a single 10Gb transparent optical lambda on September 1, 2003, thus creating a 10Gb triangle between Amsterdam, Chicago and Geneva. CERN also hopes to connect to the TeraGrid in 2003.

CERN's OC-12c circuit supports 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 StarLight.

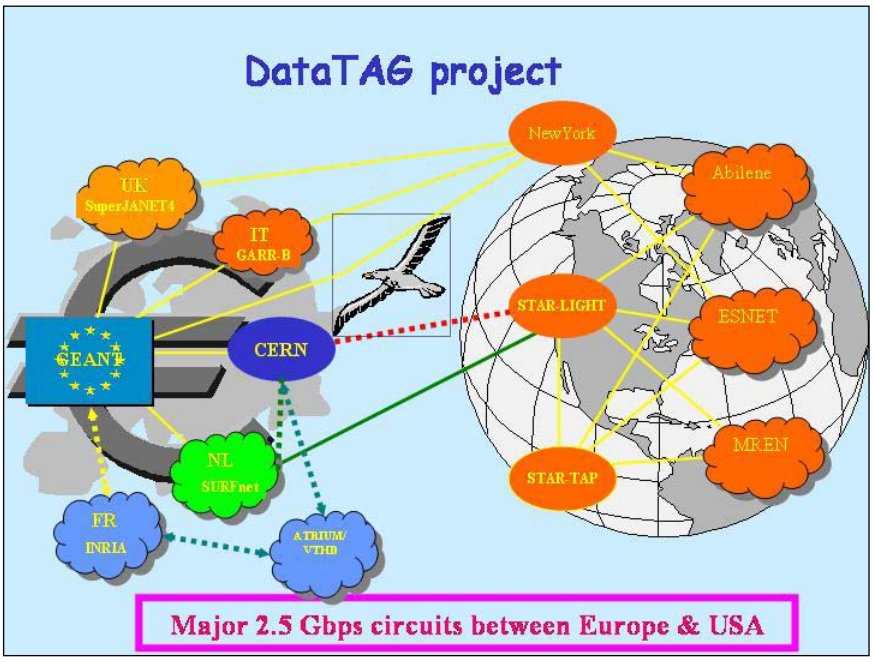

ISRAEL IUCC

Israel disconnected its STAR TAP connection in May 2002, which also concluded its participation in Euro-Link funding. Since the European Union's GÉANT project pays 50% of the cost of a transatlantic circuit, Israel now connects to Abilene in New York via its GÉANT connection.

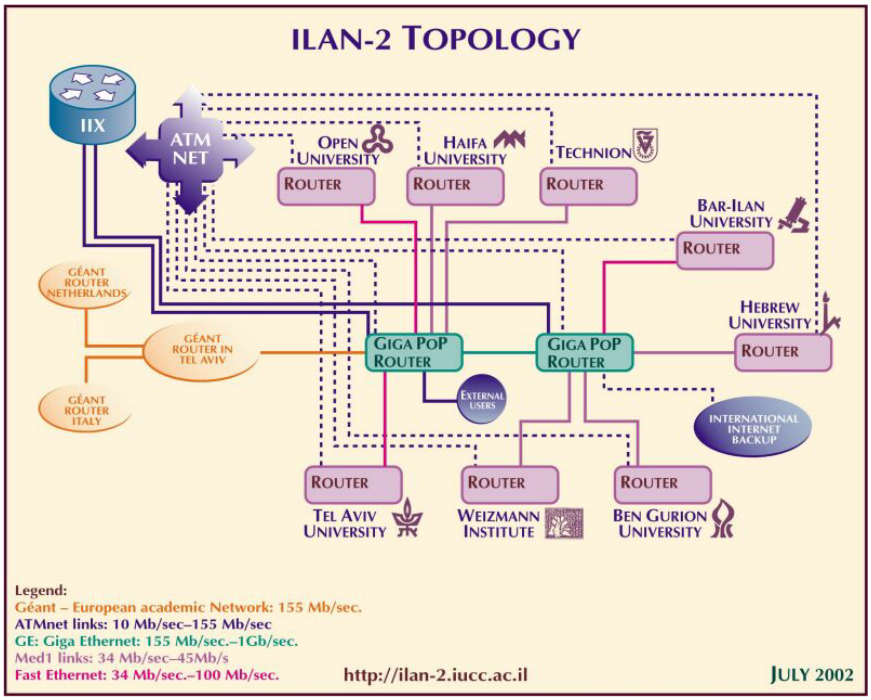

NORDUNET (155Mb)

In spring 2001, NORDUnet contracted for an STM-4c (622Mb) link between Stockholm and the Abilene PoP in New York, and an OC-3 (155Mb) link from New York to STAR TAP in Chicago to triple its DS-3 STAR TAP connection. Deutsche Telekom/T-Systems provides NORDUnet's 622Mb transatlantic link to Abilene's New York PoP, and Qwest provides its 155Mb link from New York to StarLight. NORDUnet intends to maintain this connection until August 2003, when it plans to bring a 2.5Gb wavelength to NetherLight.

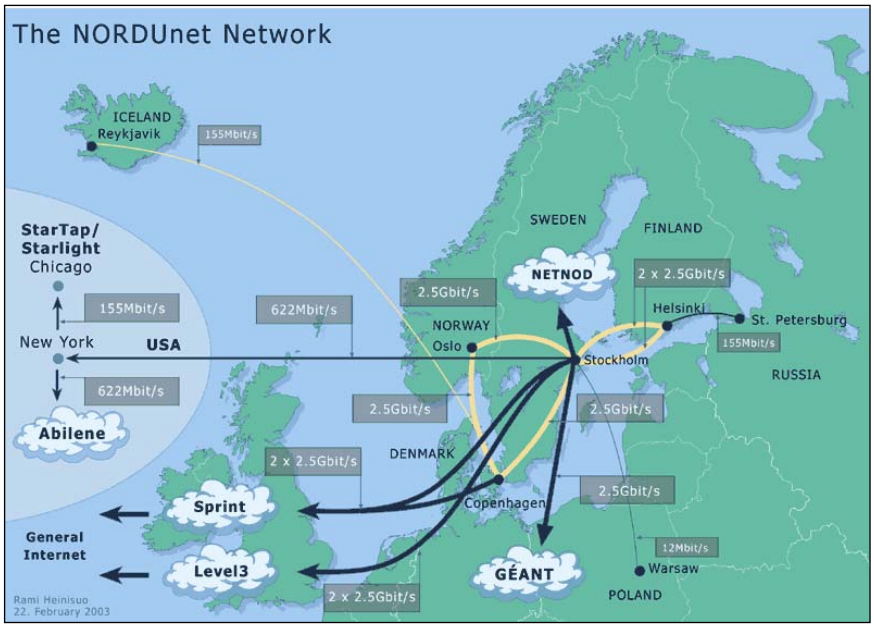

RENATER2 (155Mb)

RENATER2 maintains an OC-3 link to STAR TAP, upgraded from DS-3 in November 2001. No information is available on possible upgrades to OC-12 or 2.5Gb to StarLight, both of which are being considered by RENATER2.

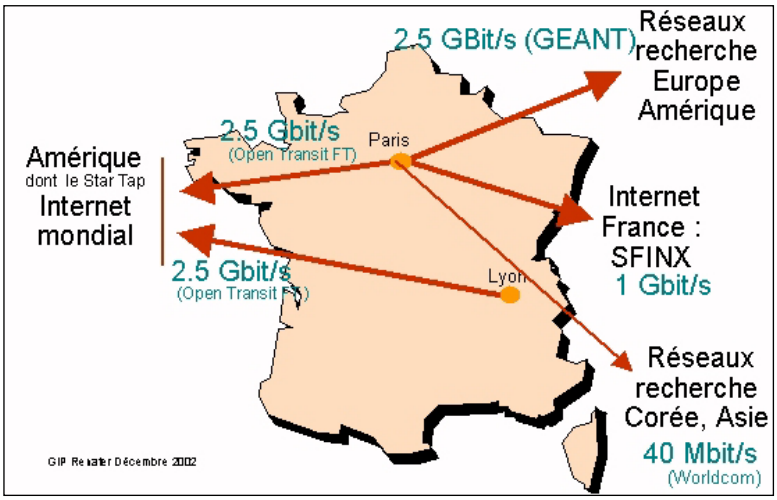

SURFNET (10Gb)

Between October 1 and December 31, 2003, Level 3 maintained a 10Gb Chicago-to-Amsterdam transoceanic link, donated for the month of September in support of iGrid 2002 and then retained by SURFnet through the end of the year as a replacement for its 2.5Gb link. As of January 2003, T-Systems provides the SURFnet 10Gb lambda. SURFnet also has access to a 10Gb link from Amsterdam to New York donated by IEEAF/Tyco.

SURFnet uses 1Gb of its 10Gb lambda for border router connectivity. SURFnet has a Cisco ONS 15454 (optical transport platform) at StarLight, which others (e.g., CERN, MREN) are welcomed to use.

SURFnet's four major peers in the US are Abilene, CA*net, Esnet and vBNS+. Abilene, ESnet and CA*net4 are currently connected to the StarLight router, with the vBNS+ planning to connect soon. Until then, SURFnet continues to use the shared OC-12 between StarLight and STAR TAP (at Ameritech) in order to reach the vBNS+.

The UIC National Center for Data Mining (NCDM) has clusters at both StarLight and NetherLight in Amsterdam. EVL and NCDM are testing applications between Chicago and Amsterdam at 2xGigE rates.

2.A.3. Compliance with NRN and NSF [vBNS] Acceptable Use Policies

To comply with the Acceptable Use Policies (AUPs) of NSF-approved High-Performance Internet Service Providers (e.g., vBNS+, Abilene), Euro-Link NRNs segregate their research and commodity traffic either at home or on the US East Coast prior to connecting to STAR TAP/StarLight.

2.A.4. Network Performance and Usage

The StarLight NOC maintains network statistics. From the StarLight Engineering web page <www.startap.net/starlight/ENGINEERING>, click on "StarLight Network Operations Center" and then "Network Monitoring". Alternatively, one can go directly to <http://loadrunner.uits.iu.edu/mrtg-monitors/starlight>. In addition, the NRNs maintain their own traffic statistics, some of which are publicly accessible and others that are password protected and/or available upon request.

2.A.5. Application-Level Network Performance Analysis Tools

EVL is developing a new generation of performance analysis tools. We are adopting and adapting emerging network performance monitoring technologies and developing advanced visualization/virtual-reality and database techniques to create new applications-oriented performance measuring tools for HPIIS.

As they are developed, these tools are incorporated into EVL's Quanta, applications-centric communications middleware, described below. UIC receives major support for Quanta from NSF award ANI-0129527, with additional support from the NSF HPIIS Euro-Link cooperative agreement ANI-9730202 as well as NSF Geowall award EAR-0218918, NSF PACI Alliance cooperative agreement ACI-9619019 and NSF ITR OptIPuter cooperative agreement ANI-0225642. This research is an outgrowth of EVL's CAVERNsoft research, which was supported by the above-mentioned awards as well as NSF Research Infrastructure EIA-9802090.
<www.evl.uic.edu/cavern/teranode/quanta.html>

2.A.5.a. Cluster Resources

Two new 16-node Linux clusters, one for the StarLight facility and one for EVL, were purchased. The EVL cluster, connected to a 3 x 5 tiled display, became operational in November 2002. A partially assembled system was installed at StarLight in September to support iGrid 2002 (6 nodes of 16 are currently operational); facility power issues will delay the full StarLight installation until the end of summer 2003.

UIC's National Center for Data Mining expanded its 4-node cluster to 7 nodes at StarLight to support its iGrid 2002 application. NCDM's Terra Wide Data Mining (TWDM) testbed consists of a 3-node cluster (each node has a 1GigE connection, for a total of 3GigE) in Amsterdam, a 4-node cluster with 4 OC-12s in Ottawa, and a 7-node cluster with 7 GigE at StarLight. NCDM added a temporary 4-node cluster with 4 GigE at Baltimore for SC'02.

2.A.5.b. Photonic Network Monitoring Tools

BANDWIDTH UTILIZATION RADAR MAP

EVL is developing Multi-modal Interactive Graphical display (Migraph) software that will be capable of graphically depicting StarLight network traffic.

UCAN: UNIFIED COLLABORATORY FOR ANALYZING NETWORKS

UCAN is a tool for collaborative network performance monitoring, testing and management. The first prototype of UCAN underwent miscellaneous bug fixes. The new version is incorporated into Quanta and includes automated installation scripts for easy configuration of the libraries. Microsoft Research now uses UCAN for internal application performance monitoring. Future plans for UCAN include using it in further DiffServ tests on the EMERGE-2 testbed between EVL and Korea.

QoSIMoTo (QoS INTERNET MONITORING TOOL)

QoSIMoto <www.evl.uic.edu/cavern/qosimoto> is available for IRIX, Linux and Windows 98/2000/XP. It has been adapted to use the Quanta networking library (see Section 2.A.5.c) and is now being released along with UCAN.

2.A.5.c. Photonic Network Transmission Protocols

QUANTA: APPLICATIONS-CENTRIC COMMUNICATIONS MIDDLEWARE

Quanta is a cross-platform adaptive networking toolkit that supports the data delivery requirements of interactive and bandwidth-intensive applications. Quanta's goal is to make it easy for programmers to specify the data transfer characteristics of their applications at a high level, which Quanta then transparently translates into appropriate networking decisions. Decisions include making dedicated lightwave reservations on optically switched networks, Quality of Service (QoS) reservations on DiffServ enabled networks, and adaptive utilization of transport protocols.

Quanta is based on, and is backward compatible with, CAVERNsoft, widely used by the CAVE community to develop advanced tele-immersive applications. Consequently, Quanta inherited all the data sharing abstractions from CAVERNsoft including: message passing, cross-platform data conversions, 64-bit support, databases using distributed shared memory, remote procedure calls, remote file I/O, Forward Error Corrected UDP, Parallel TCP for bulk data transfer, Reliable Blast UDP (RBUDP), and collaborative performance monitoring.

Quanta v0.2, planned for release at the end of March 2003, will include new classes that provide basic optical switching capabilities for establishing light-path connections, message passing using a scatter-gather mechanism, utility classes for object serialization, synchronization via condition variables and barriers, and portable Abstract Data Type (ADT)-like list, stack, queue and vector. The distribution will also include both static and dynamic libraries, providing users the flexibility of dynamically loading objects from Quanta.

The new photonic switching classes in Quanta will enable control of DiamondWave optical switches (from Calient Networks Inc.), via TL1 commands using TCP. The capabilities will include dynamic creation of new cross-connections, deletion of an existing connection and querying the list of connections on a switch. These basic classes provide low-level control of these optical switches, enabling users to create their own topologies. One can envision interconnecting machines belonging to different IP domains, as part of a single domain over an optical network. Quanta will provide a generic interface for extending the implementation to new switches in the future. Work is underway to provide an API to control photonic switches from GlimmerGlass Networks.

At iGrid 2002, Quanta's current capabilities were demonstrated to illustrate the enabling power of supporting high-performance global applications with next-generation wave-length-based networking. This includes integrating applications with the optical network control plane and high-performance transmission protocols.

PHOTONIC INTER-DOMAIN NEGOTIATOR (PIN)

EVL and Northwestern University are working on designs to support user-centric inter-domain light-path provisioning over multiple domains. Collaborative optical networking enables researchers from different institutions to dynamically setup direct optical cross-connect paths in support of collaborative applications. Dynamic configuration of multiple optical-VPNs (virtual private networks) provides isolation for simultaneous running applications. Path rerouting and network load balancing is accomplished via dynamic modification of light-path connectivity within the optical-VPN.

PIN is an applications-level API between applications and photonic networks. PIN enables applications to request and release lambdas between a pair of cluster nodes or stand-alone hosts. PIN maintains a simple database of the existing topology and manages connection requests from multiple users. PIN and its associated photonic switching classes will be included in the forthcoming release of Quanta.

To allocate a light path between two end points, an application contacts its local PIN, which dispatches generic light-path signaling messages to neighboring PINs until the final destination is reached. Each PIN translates the generic light-path signaling message into the native photonic signaling message that is understood by the local intra-domain light-path signaling facility. This facility then signals the photonic switch to make adjustments to its internal MEMS switches to establish the connection. At the present time a prototype of PIN is being developed, and TL1 command sets and APIs from multiple vendors, such as Nortel, Glimmerglass, and Calient, are being examined to identify common commands that PIN needs to support.

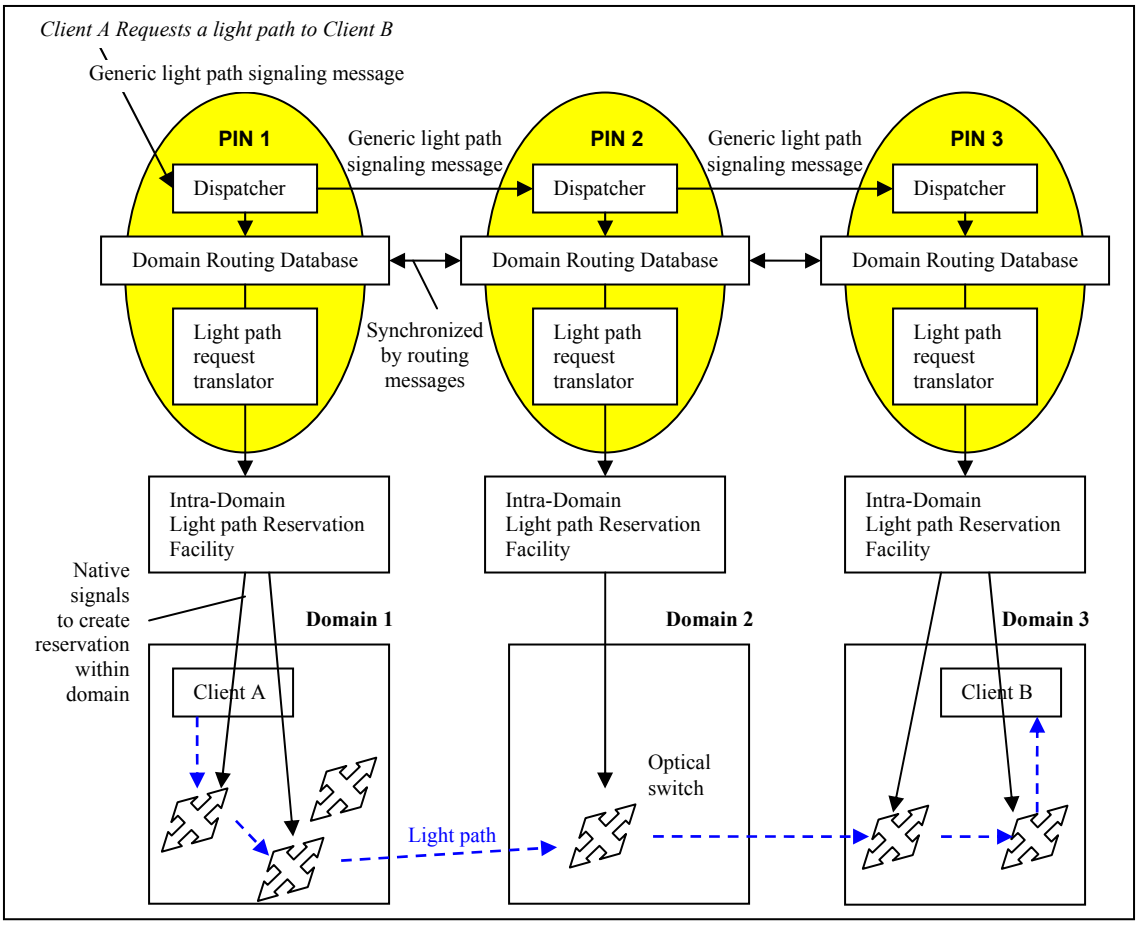

RELIABLE BLAST UDP (RBUDP)

RBUDP, integrated into Quanta, is a technique to accelerate reliable data transmission over fat networks. In RBUDP, the Sender blasts all the data (each packet is identified by a sequence number) to the Receiver. Upon receipt, sequence numbers are checked and 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. This procedure continues until the Receiver receives all packets. This technique is believed to be most effective when used with QoS, since guaranteed bandwidth may minimize transmission errors. The RBUDP scheme exploits low transmission errors to maximize throughput. The current RBUDP is used primarily for bulk data transfer, although it could be used to stream data with some tradeoff in latency.

This chart shows how well RBUDP performs on a single RBUDP stream transmitted between a pair of computers using GigE Network Interface Cards (NICs) over a 2.5Gb Chicago/Amsterdam link. The red line shows the maximum UDP throughput the NIC is able to achieve. The blue line shows that RBUDP is able to utilize almost all this bandwidth for data transmission. An *analytical (mathematical) prediction*

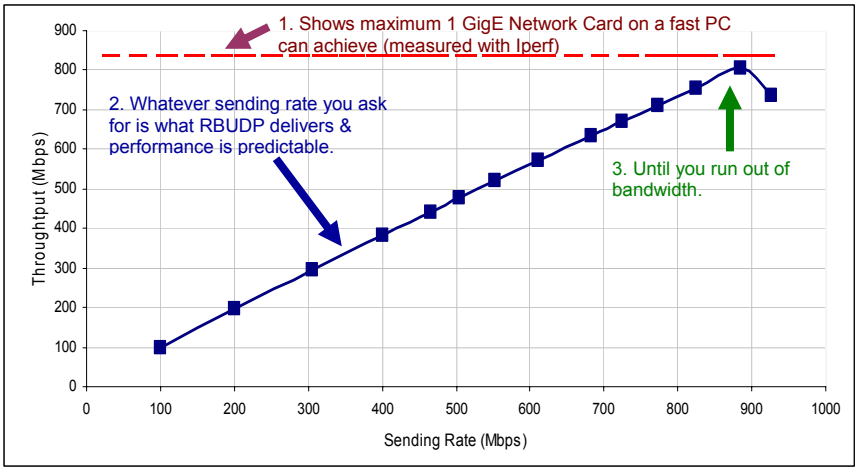

Performance of Quanta's Reliable Blast UDP Bulk Data Transfer Protocol.

function allows users to predict, given the sending rate and round-trip time, the expected throughput of RBUDP; a paper on this work was presented at the IEEE Cluster 2002 conference.

The following illustration shows the importance of this *mathematical prediction function*. In this example, a graphics application would like to reliably stream a sequence of animations from Chicago to Amsterdam (with 140ms round-trip delay). One of the caveats of RBUDP is that throughput is high only for large payloads (because a single acknowledge takes at least half the round-trip time). Several 1024x768 24-bit color animation frames must therefore be packed together to form a large payload. This means that the viewer at the endpoint will experience a certain amount of latency depending on the number of frames that need to be packed. High degrees of latency may be tolerable for passive viewing, but if the goal is to stream *interactive* graphics, achieving low latency overhead is important. Using the RBUDP prediction model, we can create a graph that allows the application to ask "What If?" questions, such as: given the desired frame rate, and the desired send rate for RBUDP, how much latency will be incurred to achieve the desired frame rate? These kinds of prediction models, coupled with information about the physical links, is what will allow Quanta's Adaptive Network Decision Manager to help the application select the best transmission services to meet its performance requirements.

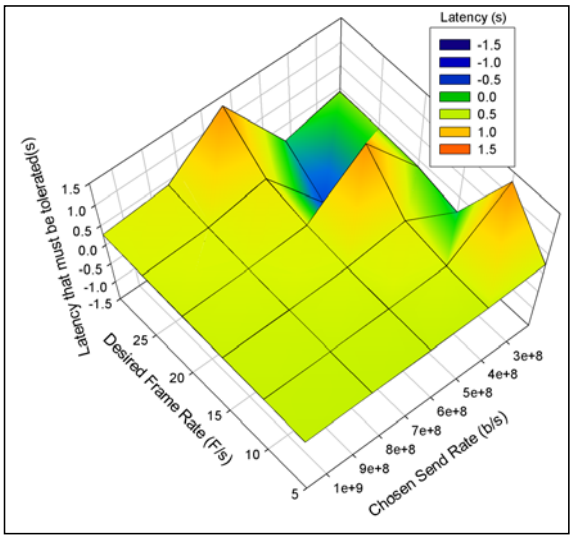

Quanta's RBUDP prediction function helps determine how much latency must be tolerated by a graphics streaming application to pack a large enough payload to achieve the desired frame rate using a chosen sending rate. (Assumptions: 1024x768 24-bit Round-Trip Time= 140ms.) Negative values of latency mean it is impossible to achieve the desired frame rate.

During iGrid 2002 in September 2002, a researcher at Vrije University in Amsterdam applied RBUDP to a parallel graphics streaming application called "Griz." Using the prediction model, EVL researchers were able to determine the number of animation frames that Griz had to package into a single payload to achieve full utilization of the StarLight/NetherLight link. After iGrid 2002, the EVL team recreated the RBUDP experiments and achieved a 725Mb data transfer rate between a pair of Linux PCs located in Chicago and Amsterdam.

New RBUDP prediction models, optimized for real-time streaming applications, have been completed. Target release is summer 2003.

FORWARD ERROR CORRECTION (FEC)

In summer 2003, FEC will be applied to RBUDP to reduce end-to-end latency of real-time streaming applications, and new protocols will be incorporated into the Quanta framework.

SIMPLE AVAILABLE BANDWIDTH UTILIZATION LIBRARY (SABUL)

SABUL, developed by the UIC National Center for Data Mining, enables very-high-throughput over long-haul networks. As an application layer protocol, SABUL merges features of UDP and TCP to produce a lightweight protocol with flow control, rate control and reliable transmission mechanisms. Specifically, SABUL transfers data and control information via UDP and TCP, respectively.

SABUL is a rate-based protocol similar to NETBLT, with two fundamental improvements. Both simulations and real-system experiments have shown that SABUL is TCP-friendly.

- Continuous data sending and rate adjustment without the concept of data blocking
- RTT-independent rate control.

As illustrated in the diagram below, SABUL uses a UDP channel to transfer large amounts of data from Sender to Receiver. The Receiver sends the state information of packets received (positive acknowledgement) and packets lost (negative acknowledgement) to the Sender using the TCP channel. The Sender dynamically updates its sending rate based on packet loss rate and loss frequency, and thereby reaches a data rate that guarantees efficiency and fairness.

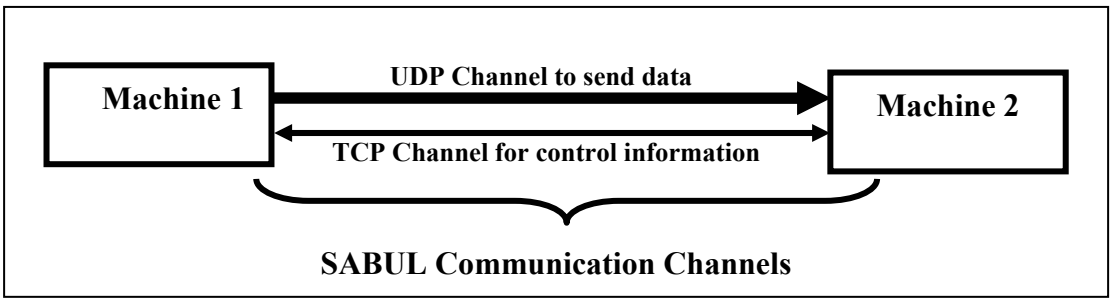

SABUL is a C++ library, implemented under Linux and ported to BSD platforms. During iGrid 2002, applications using SABUL achieved a record transfer speed of 2.8Gb over three 1Gb connections between Chicago and Amsterdam. SABUL has been integrated into several applications, including the DataSpace Transfer Protocol (DSTP), Lambda-FTP and Lambda-Mirror.

BEYONE STANDARD TCP: CERN/CALTECH EXPERIMENTS

CERN reported interesting results with respect to high-performance single stream TCP/IP streams using a variety of conventional and innovative new stacks (e.g., HSTCP, FAST, Scalable TCP, Grid DT). The results were achieved over the DataTAG circuit and its extensions to Amsterdam, Sunnyvale and Baltimore (site of the SC'02 conference) with DataTAG project partners and associated teams at Caltech and SLAC.

INTERRUPT COALESCING AND JUMBO FRAMES

The new 16-node cluster at EVL augments a four-node cluster currently used by EVL to run tests to Northwestern University. The new cluster has three times the bus bandwidth (400MHz x 64bits=2.98GBytes/s) of the current four-node cluster. In the past, EVL was able to achieve ~500Mb (local area) PC transmission rate out of the box. With appropriate tuning using *Interrupt Coalescing* and *Jumbo Frames*, EVL can achieve 850Mb.

EVL is working with Cees de Laat at University of Amsterdam and Paul Wielinga of SARA to perform tests over the 10Gb NetherLight/StarLight link. Tests include using AT&T's Virtual Network Collaboration (VNC) software to stream desktop graphics, as well as EVL's TeraVision hardware/software to stream mono/stereo graphics.

2.A.5.d. Photonic Network Application Toolkits

TERAVISION: ULTRA-RESOLUTION VISUALIZATION STREAMING

TeraVision is a hardware- and software-independent solution for real-time image distribution in advanced collaborative environments. The goal of the TeraVision project is to send high-resolution video streams between clusters using distributed servers and clients. The software enables multiple streams of synchronized video to be streamed between clusters, thus making it possible to send information to tiled or stereo displays. TeraVision multicasting and collaboration controls will enable disperse groups to make collaborative presentations, particularly when used in conjunction with the Access Grid. See: <http://www.evl.uic.edu/cavern/optiputer/teravision>.

The first version of the TeraVision system was demonstrated at iGrid 2002. In the demo, the left and right component video streams of a virtual-reality application were rendered by an SGI Onyx in Chicago, synchronized and streamed to Amsterdam, where they were resynchronized and displayed on a GeoWall. Performance statistics were monitored throughout the event to isolate bottlenecks in the system.

In December 2002, plans were made to develop TeraVision v2.0 with more flexibility. It will have the following advantages over version 1:

- New networking protocols, compression modules, frame grabbers and other video manipulation modules using a quick plug-in approach. The new design loads the appropriate module at run-time, making the design cleaner and easier to maintain.
- An RLE compression module. Newer compression algorithms will be integrated later.
- Multicasting capabilities (with scatter-gather calls) using EVL's RBUDP transmission scheme.
- A protocol for an intelligent server-client communication mechanism that will make it easier to control multiple clients and servers on different machines.
- An intuitive User Interface (UI) to allow the user to easily control more parameters in each of the modules

(i.e., compression quality or network packet sizes). The UI will contain visualization tools for analyzing performance of different system components.
- A console application for remote login and control, and a solution for the *X* x *Y* to *M* x *N* display problem.
- File streaming using different file formats (.gif, .bmp, .jpeg, .ppm, .xpm, etc.)

The modules have been implemented and are currently undergoing testing and debugging. Linux and Windows versions were developed concurrently and will be interoperable. A TeraVision system was installed at NCSA (UIUC) in March. EVL will demonstrate system version 2.0 at the 2003 Alliance All Hands Meeting April 30-May 2, 2003. <http://www.ncsa.uiuc.edu/Conferences/2003Meeting/aah-index.html>.

TERASCOPE

TeraScope is a set of visualization tools for massively parallelized Visual Data Mining applications running on high-speed networks. The goal is to interactively navigate and visualize more than a terabyte of data on either a high-resolution tiled display or a desktop workstation. For TeraScope, EVL is developing new algorithms and tools to interactively query and mine terabyte datasets <http://www.evl.uic.edu/cavern/optiputer/terascope.html>. An early prototype was demonstrated at iGrid 2002.

Message Passing Interface (MPI) is used to parallelize TeraScope code to run on clusters, and to incorporate visualization capabilities into its 3D graphical user interface (GUI) design. TeraScope provides parallel pipelines between the GUI and large datasets on remote DSTP servers, and parallelizes data querying, computing and display.

Latency is a key factor affecting user interaction in terabyte data computation and visualization. TeraScope includes a utility called *LamdaRAM,* which is a distributed memory cache that aggressively uses network bandwidth to prefetch application data from remote servers, thereby reducing latency. EVL is working on a benchmark for LamdaRAM in order to optimize it for other applications.

2.A.5.e. Advanced Collaborative Environments (ACE) Grid Working Group

Jason Leigh, EVL associate professor, and Rick Stevens of Argonne National Laboratory formed the Global Grid Forum (GGF) ACE Working Group <http://calder.ncsa.uiuc.edu/ACE-grid/> to complement other GGF Working Groups <http://www.gridforum.org/>. ACE aims to provide human-centered techniques and technologies for facilitating interactive, collaborative, and immersive access of Grid resources from anywhere, at any time.

Leigh attended the GGF7 meeting March 4-7, 2003 in Tokyo, Japan, via Access Grid, at which the penultimate draft of the ACE security document was discussed. This document details the security requirements of Advanced Collaborative Environments, and is intended to provide input into the GGF's overall grid security efforts. This document will be submitted to the GGF for review as a Grid Working Draft.

Additional topics discussed included the creation of future working groups in visualization services, collaboration services and Access Grid scheduling services.

Argonne's Bill Allcock presented EVL's survey of alternative high-speed transport protocols to the Data Transport Research Group at GGF7. EVL will consolidate all the information provided by members of that Group.

2.A.5.f. Transatlantic TCP Network Performance Studies

Yggdrasil (YG) is a script-based, authoring toolkit for networked virtual-reality applications that is built on top of Quanta; it was developed by Dr. Dave Pape, 2001 graduate of EVL. This software allows non-programmers to create effective, behavior-rich scientific and/or artistic tele-immersive environments. EVL co-director Dan Sandin extended the YG library for behaviors and network performance tests.

YG was used to develop tele-immersive EVL art installations exhibited in three separate international conferences/shows during the past year, all of which included networked participation from EVL, SUNY Buffalo's Department of Media Study <http://cas.buffalo.edu/media_study>, and the Tools for Creativity Studio at the Interactive Institute of Sweden in Umeå <http://www.interactiveinstitute.se/>.

- In conjunction with the Stockholm Art Fair, March 5-8, 2003, EVL's Dan Sandin and students remotely participated in a networked tele-collaborative VR art event which included Sandin's "Looking for Water," Josephine Anstey and Dave Pape's "PAAPAB," Umeå's theater, and Bino & Cool's Yggdrasil (collaborating via Stockholm's Center for Parallel Computers (KTH)).
- EVL's Dan Sandin attended the III Jornada en.red.ando Conference in Barcelona, Spain, January 31, 2003,

and ran a real-time, networked demonstration of his CAVE art application "Looking for Water."

- EVL student Todd Margolis held his MFA Thesis show "The Perfect Parlor," January 10, 2003, at the UIC Electronic Visualization Laboratory in Chicago featuring a networked tele-collaborative session between Chicago, New York, Indiana, Sweden and the Netherlands.

Sandin, Pape and EVL student Alex Hill are addressing issues related to TCP communication over transatlantic links. Previous demos (e.g. iGrid 2000) required several minutes of startup time when downloading a full virtual-reality world description. This is now reduced to a few seconds (3 seconds in the case of the Stockholm Art Fair).

When testing for the Barcelona show, Pape found that a YG repeater running on a Linux machine could not handle lots of small TCP-based updates quickly over a high-latency connection (~ 150 ms ping time). Modified YG code handles this better, but applications are still required to reduce the frequency of world-data updates, and to redirect as much of them as possible to UDP.

At the Stockholm Art Fair show, the master system was a PC. Pape and Sandin ran the repeater at KTH/PDC (Royal Institute of Technology, Center for Parallel Computers), to take advantage of its excellent connectivity to NORDUnet. The route from Buffalo went via Abilene to NORDUnet; from EVL it went through StarLight to NORDUnet. The typical round-trip ping time was 130 ms from Buffalo or Chicago to KTH. The potential bandwidth appeared to be on the order of 20 Mbps.

2.A.6. Euro-Link/STAR TAP Engineering Services

2.A.6.a. IPv6 Service at the 6TAP

Two 6TAP <www.6tap.net> IPv6 services are run by ESnet and CANARIE — one at STAR TAP and one at StarLight. 6TAP supports IPv6 over IPv4 tunnels and IPv6 performance measurement and statistics.

2.A.6.b. NLANR Web Cache

Due to greatly improved connectivity in Europe, STAR TAP engineers elected not to procure a contract with ISP provider Genuity (subsequently purchased by Level 3) to continue to support an NLANR web cache for international STAR TAP participants. Genuity had donated 1Mb Virtual Transit Service for this purpose since April 1998. Service formally ended October 25, 2002. The web cache was returned to NLANR.

2.A.6.c. NLANR Performance Measurement

An NLANR AMP (Active Measurement Platform) box is located at STAR TAP and information is accessible from the STAR TAP web pages <http://www.startap.net/ENGINEERING/PERFORM.html>.

2.A.6.d. EMERGE Differentiated Services (DiffServ) Testbed

Oliver Yu and Jason Leigh set up the EMERGE-2 DiffServ testbed between EVL, Northwestern University, and KISTI (Korea Institute of Science and Technology Information) over KREONET/STARTAP. In EMERGE-1, premium service provisioning for tele-immersive applications across a single domain was tested between EVL and Argonne National Lab. In EMERGE-2, assured service provisioning for multimedia applications across multiple domains is being deployed. The results of these experiments will be compared to the results of using the over-provisioned (uncongested) 10Gb SURFnet and 2.5Gb DataTAG links.

NLANR is providing EVL with an OC3-MON monitoring tool for the EMERGE-2 testbed. Students working under Oliver Yu are designing schemes for GARA/Globus-based resource reservation, TCP and RBUDP (Reliable Blast UDP) transmissions, differentiated transmissions of MPEG-2 video applications, and adaptive QoS control over this multi-domain DiffServ network.

Note: EMERGE was initially a DoE-funded effort to demonstrate DiffServ over MREN, ESnet and Abilene <http://www.evl.uic.edu/cavern/EMERGE/>. We extended these experiments to Europe and Asia via STAR TAP.

2.A.6.e. Multicast

NORDUnet and SURFnet have Native Multicast enabled. We are interested in creating an AUP-free IP multicast exchange point at StarLight, primarily due to ANL's Access Grid effort, which requires a solid IP multicast service.

2.A.6.f. Peering: Bi-Lateral Agreements

STAR TAP and StarLight run in an Acceptable Use Policy (AUP)-free mode; that is, connecting networks must agree pair-wise regarding acceptable use. For STAR TAP, we rely on mature ATM switching as provided by AADS; for StarLight, we use a Cisco 6509 switch/router. Once connected, the NRNs can peer with US Fednets, vBNS+, Abilene and one another. They either peer by bilateral agreement, which is functionally implemented at AADS with a full-mesh ATM Private Virtual Circuit (PVC) layer-2 service at the switch. Or, they can do level-3 peering using the STAR TAP/StarLight Routers. They may also connect to one or more ISPs at the AADS facility, which is outside the scope of STAR TAP and StarLight, but a useful capability nonetheless.

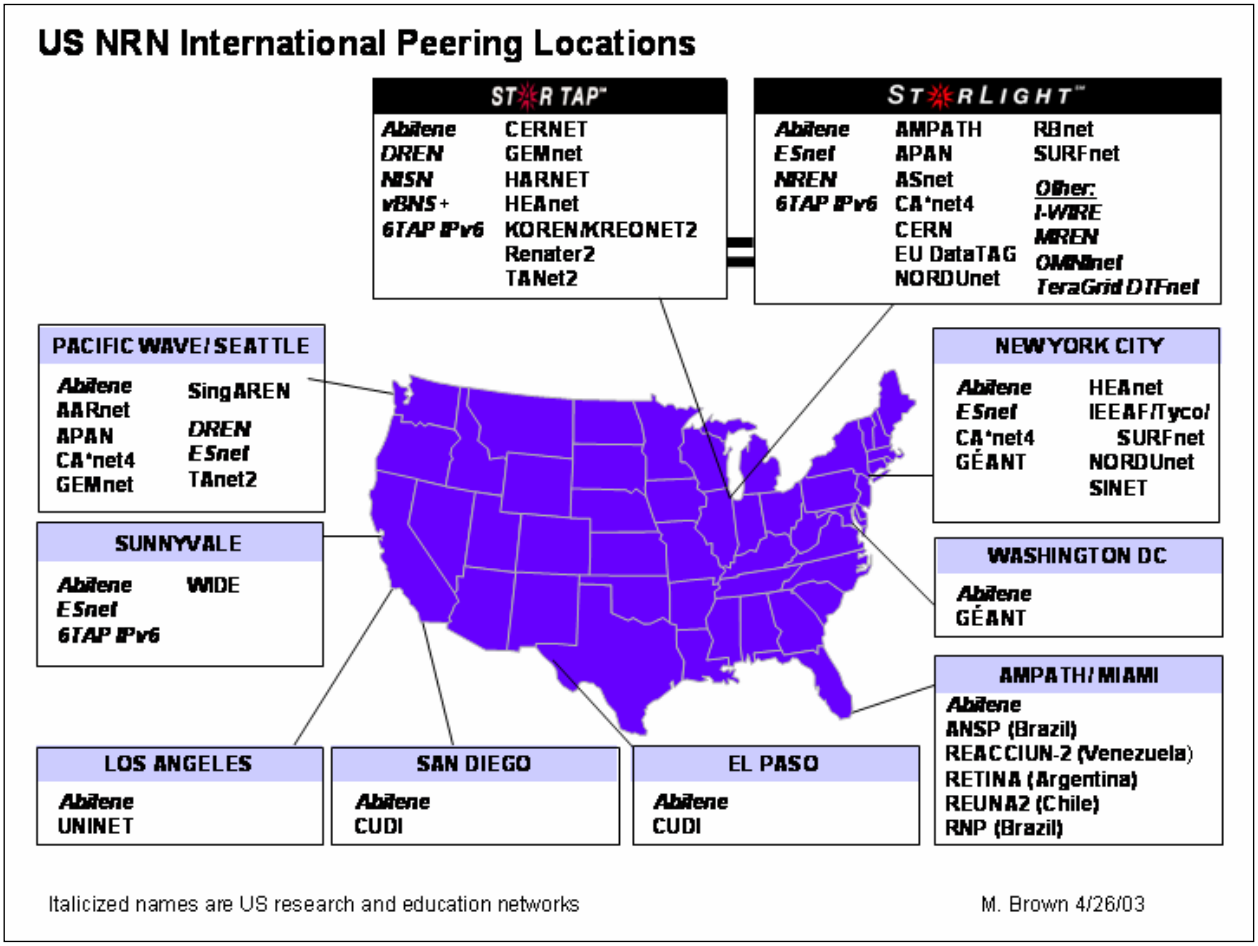

2.A.6.g. International Transit Network (ITN) Services

ITN service is offered by CANARIE and Internet2/Abilene to facilitate connectivity among international NRNs connecting to North American coasts. It became operational in October 2000. 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. [Note: Abilene is updating its web site and the link may temporarily be inoperative.]

2.A.6.h. StarLight/Abilene Connectivity

Abilene's Juniper T640 core router, located at the Qwest PoP, along with two OC-192 circuits to other Abilene backbone sites, connects to StarLight via two I-WIRE dark fiber connections. One 10GigE connection is for StarLight National Research Network traffic and a 1GigE connection is for MREN traffic.

Abilene will maintain its direct connection to the STAR TAP NAP (currently OC-12, but may be reduced to OC-3) until approximately August 2003. The number of peer networks still connected there will determine whether

Abilene continues its direct connection.

2.A.6.i. StarLight/STAR TAP Connectivity

An OC-12 connection exists between StarLight and STAR TAP (at the Ameritech NAP) for shared use by StarLight customers. This NSF-funded OC-12 StarLight/STAR TAP link will be maintained until April 30, 2004 after which StarLight may procure other funds if it deems it important to keep the link operational.

2.A.6.j. StarLight/US R&E Network Connectivity

ESnet and NREN connect to StarLight via 1Gb connections using I-WIRE fiber. DREN and vBNS+ expected soon.

2.A.6.k. STAR TAP/StarLight Routers

In our distributed STAR TAP/StarLight environment, there is a STAR TAP Router at the AADS NAP (Cisco 7505), a StarLight Router at the StarLight facility (Juniper M10), and an MREN Router at StarLight (Juniper M5). A Cisco 6509 serves as an additional MREN Router as well as the StarLight Gigabit Ethernet Exchange Point switch.

2.A.7. Euro-Link Network Operations Center (NOC)

The Euro-Link NOC can be accessed from <www.euro-link.org/ENGINEERING> or <http://noc.euro-link.org>. Indiana University's Jim Williams and Steve Peck operate the NOC. In addition to Euro-Link, Indiana's Global NOC operates the STAR TAP/StarLight, TransPAC, NaukaNet (formerly MIRnet), AMPATH and Abilene networks as five logical NOCs inside one physical NOC. Each European NRN operates its own NOC as well; this information is maintained for all STAR TAP-connected networks at <www.startap.net/ABOUT/points.html>. In addition to BGP session and syslog monitoring tools, the NOC has a router proxy that enables users to submit "show commands" to a STAR TAP/StarLight core node router. <http://loadrunner.uits.iu.edu/~routerproxy/startap/>.

The Euro-Link NOC monitors NRN connections to STAR TAP/StarLight, and when a particular network loses connectivity, troubleshooting begins. It contacts the European NRNs to inform them of problems through NOC-to-NOC email as well as formal email notices, and offers assistance as necessary. The Euro-Link NOC provides weekly reports of network availability (e.g., downtime, scheduled maintenance, etc.) to the peering networks.

2.A.8. Euro-Link Documentation and Dissemination of Information

The Euro-Link web site <www.euro-link.org> is continually updated with Euro-Link information, including general information, network performance studies, engineering data, applications and publications. This report will be posted on the web site shortly after its completion.

2.A.9. Euro-Link Applications Documented

Active US/Euro-Link collaborations are documented on the Euro-Link web site <http://www.euro-link.org/APPLICATIONS/>. An up-to-date list also appears in the Appendix of this document.

2.B. Research Findings

2.B.1. StarLight

StarLight has been operational since September 2001, with proposed expansion plans through 2006. Euro-Link members CERN, SURFnet and NORDUnet are currently connected.

2.B.2. Application-Level Network Performance Analysis Tools

Preliminary results are reported in Section 2.A.5 along with descriptions of the tools. Research papers and MS/PhD degrees are currently in progress.

2.B.2.a. Application-Level Network Performance Studies

Jason Leigh is working with SARA in Amsterdam to do network performance studies over long, fat networks using various transmission techniques (TCP, UDP, FEC, RBUDP). Leigh has also worked with CERN on RBUDP tests and is talking to them about DiffServ experiments. Dan Sandin also did some TCP-based network performance studies to Europe. These collaborations are documented in Section 2.A.5.

2.B.3. Network Performance Analysis Software Releases

Quanta <http://www.evl.uic.edu/cavern/quanta/> is application-centric communications middleware that supports the data delivery requirements of interactive and bandwidth intensive applications. Quanta v 0.1 was released in Spring 2002. Version 0.2 is expected in March 2003. CAVERNsoft G2, EVL's virtual-reality collaboration software, has been to use the Quanta networking library, which can be downloaded from <www.cavernsoft.org>. The Pre-Quanta version of CAVERNsoft G2 can be downloaded from<www.openchannelsoftware.com>.

QoSIMoTo (QoS Internet Monitoring Tool) <www.evl.uic.edu/cavern/qosimoto> is available for IRIX, Linux and Windows 98/2000/XP. It has been adapted to use the Quanta networking library and is now being released along with EVL's **UCAN** (Unified Collaboratory for Analyzing Networks), a collaborative network management toolkit.

2.B.4. Collaborations (April 2002–March 2003)

2.B.4.a. SC 2002

Tom DeFanti, Maxine Brown, Jason Leigh, Greg Dawe and several students attended SC'02. Jason Leigh and his students demonstrated TeraScope in the "Project DataSpace" research booth organized by Bob Grossman of the UIC National Center for Data Mining (NCDM).

Grossman demonstrated a collaborative project with researchers from Chicago, Ottawa and Amsterdam, and was awarded the SC'02 High Performance Bandwidth Challenge Award for Best Use of Emerging Infrastructure. Researchers from NCDM, CANARIE (Canada) and SARA (Amsterdam) have been working together over the past year to perform real-time data correlation over lambda networks.

At SC02, their demonstration had impressive results. A stream of data was moved from a cluster at SARA to a cluster at StarLight at over 2.8Gb. At the same time, a stream of data was moved between clusters in Ottawa and StarLight at over 2Gb. Both streams used the SABUL protocol, designed for high-performance data transport. The two data streams were merged at over 500Mb per node on StarLight's three-node cluster. These "lambda joins" are an important component of distributed data-mining applications.

2.B.4.b. iGrid 2002, Amsterdam, September 23-26, 2002

EVL, along with University of Amsterdam and Gigaport, co-organized iGrid 2002, an applications-driven testbed event that showcased 28 scientific and artistic applications from 16 countries. StarLight served as the optical exchange point for many of the applications. A donated 10Gb link from Chicago to Amsterdam temporarily donated by Level 3, and a donated 10Gb link from New York to Amsterdam donated by IEEAF/Tyco, augmented the SURFnet 2.5Gb link between Chicago and Amsterdam in support of the applications. The event was tremendously successful, providing the world's research community with the opportunity to work together to advance the state of the art in LambdaGrid-intensive computing. There were 250 attendees. Notable achievements include:

- A Terabyte of research data was transferred in less than a DVD/minute over a newly established "light path" extending 12,000 kilometers from TRIUMF in Vancouver to CERN. The demonstration required

dedicated portions of fiber-optic networks, spanning one provincial (British Columbia's BCnet) and two national research and education networks (CA*net4 and SURFnet) to establish the on-demand private network. Experiments were also conducted on the Chicago to CERN DataTAG link. The project culminated in establishing the first large-scale end-to-end "light path". ***Researchers were able to transfer a Terabyte of research data from disk-to-disk in under 3 hours – this is equivalent to transferring a full CD in less than 8 seconds (or a full-length DVD movie in less than 1 minute).*** Using *bbftp,* they could transfer 60GB files in 10 minutes; 1TB is about 17 of these files, or about 170 minutes transfer time. Using *Tsunami,* they achieved slightly better rates. Peak transfer rates in excess of 1Gb were achieved, twice the previous known record for this distance. This is the first establishment of an inter-domain end-to-end "light path" dedicated for a research application. The "light path" directly connecting TRIUMF and CERN is the longest known single hop network spanning the distance from Vancouver to Geneva via StarLight, NetherLight and DataTAG. (Note: Performance was affected by using 10GigE alpha-test cards from Intel. Subsequent testing with beta-test cards has shown that they perform roughly 4 times better.)

- UIC's Laboratory for Advanced Computing/National Center for Data Mining (LAC/NCDM) and the International Center for Advanced Internet Research (iCAIR) at Northwestern University set a transatlantic data-transfer record using Photonic Data Services (PDS), a technique they jointly developed. Using PDS, data was transmitted at 2.8Gb as part of a data-mining application. The demonstration was the first for PDS and shows the potential for data-mining applications to drive the use of available bandwidth.
- Artist Jackie Matisse's virtual-reality application "Kites Flying In and Out of Space" used a Grid model for real-time steering of calculations on computers distributed worldwide over high-speed networks. Each of the 12 kites appearing in the piece utilized up to 15Mb. The application used distributed servers in Chicago, Canada, Japan, Singapore and Virginia to compute each kite's physical dynamic properties. Each server sent a single kite's motions to iGrid where it was visualized in a CAVE. The application was scalable both computationally and geographically. It proved a good test of high-speed networking because the application required a multicast-enabled network to accomplish communications. As a side benefit, the kites become a visual metaphor for network performance as the kite motions (e.g., fast, slow) responded to latency of the network data.

On the last day of the demonstrations, SURFnet/SARA stress-tested the LAN networks and routers in the facility by generating and sustaining 8Gb of data in and out of the facility for one hour. For WAN and LAN bandwidth usage statistics, see the iGrid website <www.igrid2002.org> (scroll to "A page with bandwidth usage statistics is available") or see <http://www.igrid2002.org/bandwidth/index.html>.

2.B.5. Meetings Attended (April 2002–March 2003)

Because we also manage the NSF-funded STAR TAP initiative, many of the activities documented below overlap with those listed in our STAR TAP reports.

2.B.5.a. Euro-Link Annual Meetings

In lieu of a formal annual STAR TAP International Advisory Committee meeting in 2002, the Euro-Link management team met with individual member network architects at the iGrid 2002 event in September. On the final day of iGrid 2002, a LambdaGrid meeting was held where TransLight issues were discussed. (As of April 2003, TransLight is an NSF-funded one-year experimental networking trial between North America and Europe.)

Past STAR TAP/Euro-Link meetings held in conjunction with annual INET conferences are documented on the web <http://www.startap.net/ABOUT/meetingsIndex.html>.

2.B.5.b. HPIIS Team Meetings

While no formal HPIIS Team Meetings were held over the past twelve months (however, one was held in Washington DC in April 2003 and will be reported in next-year's annual report.), many informal discussions have taken place in the past year as HPIIS teams relocated their connections to StarLight in Chicago. Past HPIIS meetings are documented on the web site at: <http://www.euro-link.org/ABOUT/meetings.html#HPIIS00> and <http://www.startap.net/ABOUT/meetingHpiis99.html>

2.B.5.c. Euro-Link Participation in International Conference Events

In November, we attended SC'02 and participated in TeraScope and Tera Wide Data Mining project demos to

promote the goals of StarLight. (See Section 2.B.4.a.)

Also at SC'02, CERN temporarily connected to the TeraGrid via the StarLight router to run high-performance single stream TCP/IP tests over the DataTAG circuit. The results were achieved over the circuit and its extensions to Amsterdam, Sunnyvale and Baltimore (SC'02 conference) with DataTAG project partners and associated teams at Caltech and SLAC.

The iGrid 2002 event held September 23-26 in Amsterdam challenged scientists and technologists to optimally utilize 2.5-10Gb experimental networks, with special emphasis on e-Science, Grid and Virtual Laboratory applications <http://www.startap.net/igrid2002/> (See Section 2.B.4.b.) Past events that we organized include iGrid '98 at SC'98 and iGrid 2000 at INET 2000.

2.B.5.d. Euro-Link Meeting, Workshop and Conference Participation (April 2002-March 2003)

March 28, 2003. University of Amsterdam's Bas van Oudenaarde and Freek Dijkstra arrived at EVL for two weeks to with Jason Leigh and student Eric He on learning to program the Calient optical switch (which will be shipped soon to Amsterdam), and to integrate their AAA software into EVL's Photonic Interdomain Negotiator (PIN).

March 21, 2003. Tom DeFanti and Maxine Brown attended the MREN Executive Committee meeting, where regional connectivity to StarLight, most notably FermiLab connectivity, was discussed. Several European sites are interested in high-bandwidth connectivity to FermiLab.

March 14, 2003. Tom DeFanti, Maxine Brown and Alan Verlo hosted a meeting at EVL to discuss TransLight's Internet2/GÉANT 10Gb connectivity requirements at StarLight. Attendees included Internet2's Rick Summerhill, Heather Boyles, Steve Corbato, Northwestern University's Joe Mambretti, and Argonne's Linda Winkler, Caren Litvanyi and Bill Nickless.

March 11, 2003. Greg Hidley, Cal-$(IT)^2$, UCSD, visited EVL to consult with directors and staff regarding visualization and networking technologies for OptIPuter partner Cal-$(IT)^2$.

March 5-8, 2003. In conjunction with a four-day art festival in Stockholm, EVL's Dan Sandin participated in networked tele-collaborative events between Chicago, Stockholm and New York.

February 25-27, 2003. Tom DeFanti and Maxine Brown attended an NSF-sponsored workshop on "The New Role of Science and Engineering in Risk Reduction" organized by Cal-$(IT)^2$, UCSD.

February 12, 2003. Calient Systems installed their 3D MEMS DiamondWave optical switches at EVL. EVL recently purchased two Calient switches, one for StarLight and one for NetherLight, which are being tested before being deployed. We plan to do optical switching tests between Chicago and Amsterdam, under the supervision of Jason Leigh at EVL and Cees de Laat at University of Amsterdam.

February 3-4, 2003. Tom DeFanti and Maxine Brown helped organize a ON*VECTOR Photonics Workshop, hosted by the Cal-$(IT)^2$, UCSD, and sponsored by NTT. ON*VECTOR (Optical Networked Virtual Environments for Collaborative Trans-Oceanic Research) is a joint project of NTT Network Innovation Laboratories, University of Tokyo and University of Illinois at Chicago, and managed by Pacific Interface Inc. (PII). This second annual workshop brought together optical networking researchers from Japan, the US, Canada and The Netherlands (Erik Radius of SURFnet and Cees de Laat of University of Amsterdam participated.)

January 29-31, 2003. FIU-AMPATH Workshop: Fostering Collaborations and Next Generation Infrastructure, was held in Miami. Tom DeFanti and Maxine Brown attended. DeFanti was on the panel "International Optical Networking" (chaired by Maxine Brown) and also gave a presentation entitled "From StarLight to TransLight: Building the LambdaGrid." Brown was on a panel "How to Work with Scientists" where she gave a brief talk on "Building Scientific Communities."

January 27, 2003. STAR TAP/StarLight engineers met with Northwestern University building engineers to discuss planned remodeling efforts, to ensure enough racks and power to handle demand.

January 9-10, 2003. The NSF ANIR Principal Investigator meeting was held at the Hyatt Regency in Reston, Virginia. This two-day workshop provided a unique opportunity to meet other NSF ANIR PIs and Program Directors. Attending were Tom DeFanti (representing Euro-Link, Oliver Yu (Intelligent Signaling), Maxine Brown (STAR TAP), Joe Mambretti (StarLight) and Jason Leigh (Quanta). There was interest in international connectivity from several researchers, notably Alan Whitney of Haystack Observatory at MIT who wants to get high-speed

connectivity to the JIVE vlbi center in Dwingeloo, The Netherlands.

December 17, 2002. Tom DeFanti, Alan Verlo and other STAR TAP engineers teleconferenced into the Joint Engineering Team (JET) meeting held at NSF in Arlington, Virginia. Tom gave a presentation about the future of STAR TAP and StarLight. Discussion took place about deploying IPv6 at various GigaPoPs, including StarLight.

November 18-22, 2003. Tom DeFanti, Maxine Brown, Jason Leigh, Greg Dawe and several students attended SC'2002. Jason Leigh and his students demonstrated TeraScope in the "Project DataSpace" research booth organized by Bob Grossman of the UIC National Center for Data Mining (NCDM).

November 20, 2003. Alan Verlo teleconferenced into the Joint Engineering Team (JET) meeting held at the SC'2002 conference in Baltimore. Linda Winkler and other STAR TAP engineers were in attendance.

October 18, 2002. A meeting of the Amsterdam/Chicago optical research collaboration team was held at EVL. Attendees included: Cees de Laat and Leon Gommans (UvA), Peter Clarke (UCL), Valerie Taylor, Joe Mambretti, Jim Chen, Elizabeth Bacon, Fei Yeh, David Lillethun (NU), Tom DeFanti, Maxine Brown, Jason Leigh, Oliver Yu, Mitch Theys, and Bob Grossman (UIC). Agenda and PowerPoints at <http://www.startap.net/starlight/ABOUT/meetOpResCollab02.html>.

October 15, 2003. Alan Verlo teleconferenced into the Joint Engineering Team (JET) meeting held at the ACCESS Center in Arlington, Virginia. He reported on the success of iGrid, and answered questions from Fednet representatives on StarLight/STAR TAP connectivity. JET is a standing committee of the Large Scale Networking (LSN) Coordinating Group (CG) of the Interagency Working Group (IWG) on Information Technology Research and Development (ITRD). The IWG reports to the National Science and Technology Council (NSTC) and its Committee on Technology (CT). The LSN and its teams serve as the forum for coordinating activities of the LSN Program Component Area.

October 14, 2002. In conjunction with Global Grid Forum 6 (GGF6) being held in Chicago, EVL hosted an open house on Monday, October 14, from 4-6 PM <http://www.gridforum.org/Meetings/ggf6/default.htm>. We showcased several iGrid demonstrations from UIC/EVL, UIC/NCDM, NU/ iCAIR and NCSA ACCESS Center at EVL as part of a self-guided tour of the lab to educate the global Grid community about StarLight as an enabling facility to large-scale global Grid applications. Featured applications included: TeraScope, Kites Flying In and Out of Space, High Performance Data Webs, and Photonic TeraStream. Approximately 250 visitors came, including Euro-Link partners from Amsterdam, CERN and the UK; photos from the event are documented at <ftp://ftp.evl.uic.edu/pub/INcoming/raj/ggf%20pics/>.

October 11, 2002. Tom DeFanti, Maxine Brown, Jason Leigh and Oliver Yu (EVL), and Joe Mambretti and Jeremy Weinberger (NU/iCAIR), met with Jack Waters (CTO), John Verduzco, Sarah Bleau and Jason Booma of Level 3 to discuss OMNInet and OptIPuter projects. Level 3 presented and demonstrated its onTAP product.

September 27, 2002. A small closed Lambda Workshop, a follow-up to the TERENA-organized workshop held in 2001, was held at University of Amsterdam. Representatives from NSF (Tom Greene), StarLight (Tom DeFanti, Maxine Brown, Joe Mambretti, Linda Winkler), I-Light/Indiana, TeraGrid, NORDUnet, SURFnet, Japan, TERENA, UK, Internet2, IEEAF, CERN, the Czech Republic and CANARIE attended. Minutes of the meeting are available upon request.

September 23-26, 2002. iGrid 2002 <www.igrid2002.org> took place. iGrid 2002 was organized by Dutch and USA organizations. Institutions in The Netherlands were: Amsterdam Science & Technology Centre, GigaPort Project, SARA Computing and Networking Services, SURFnet and Universiteit van Amsterdam/ Science Faculty. Institutions in the USA were: Argonne National Laboratory/ Mathematics and Computer Science Division, Indiana University/ Office of the Vice President for Information Technology, Northwestern University/ International Center for Advanced Internet Research, and University of Illinois at Chicago/ Electronic Visualization Laboratory. Major funding for iGrid 2002 was provided by the GigaPort Project, the Amsterdam Science & Technology Centre and the USA National Science Foundation, with in-kind support by SARA Computing and Networking Services (with funding from the NWO/NCF) and the Universiteit van Amsterdam.

September 17, 2002. Alan Verlo participated (via teleconference) in a JET meeting held at NSF in Arlington, VA.

September 15, 2002. Tom DeFanti and Maxine Brown met with Cees de Laat, University of Amsterdam, who was in Chicago for a brief time to attend a meeting of the Global Grid Forum Executive Committee.

August 26-28, 2002. Tom DeFanti attended the MAGIC (Middleware And Grid Infrastructure Coordination

Committee) Workshop sponsored by the Large Scale Networking (LSN) Coordinating Group of the Interagency Working Group for Information Technology Research and Development. The goal of the Workshop is to produce a report that describes a vision for middleware and Grid services that will enable scientific applications and cooperation (5 to 10 years out) and to identify needed Federal networking research to implement that vision. This *Blueprint for Future Science Middleware and Grid Research and Infrastructure* will also be used by LSN agency members to develop their research plans and programs. It will also be used to inform decision makers and researchers of the research needed in middleware and grids, and the need for a persistent infrastructure to support their use. <http://www.nsf-middleware.org/MAGIC/>

August 20, 2002. Alan Verlo participated (via teleconference) in a JET meeting held at NSF in Arlington, VA.

August 9, 2002. Level 3's Geoff Jordan, Kevin O'Hara, John Verduzco and Sara Bleau visited EVL to meet with Maxine Brown, Jason Leigh and Joe Membretti to learn about EVL and NU optical networking activities. Subsequent to the meeting, Level 3 donated a temporary 10Gb link from StarLight to SARA in support of iGrid 2002.

August 8, 2002. Level 3's John Verduzco and Sara Bleau visited EVL to meet with Maxine Brown, Jason Leigh and Laura Wolf.

August 7, 2002. Gert Svensson, KTH/PDC (Royal Institute of Technology, Center for Parallel Computers), visited EVL to meet with Jason Leigh to learn about current VR hardware and software research, including Linux clusters, and to get an overview of EVL's computing and networking environments and Grid activities.

July 28-August 1, 2002. Alan Verlo, Linda Winkler, Caren Litvanyi and Bill Nickless attended the NLANR/Internet 2 Joint Techs Workshop in Boulder, Colorado. On July 31, they participated in a JET (Joint Engineering Taskforce) meeting.

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 TERENA.

May 30-31, 2002. Tom DeFanti, Maxine Brown, Alan Verlo, Linda Winkler, Greg Dawe and Laura Wolf met with Cees DeLaat, Paul Wielange, Ed Mos and other members of the Dutch iGrid 2002 planning committee to discuss technical issues. 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 PACI Alliance All Hands Meeting at NCSA/UIUC. The EVL team participated in the poster session, demonstrated GeoWall, discussed StarLight, and participated in the Modern Graphics Hardware/Software tutorial.

April 25, 2002. Tom DeFanti, a member of the UCAID Board of Trustees' International Relations Committee, participated in a telephone conference call regarding 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. Ruzena Bajcsy, former head of the Computer Science Directorate at the National Science Foundation, toured EVL and met with Jason Leigh and Tom Moher.

April 15, 2002. Jason Leigh and EL student Shalini Venkataraman met with French artist Jackie Matisse 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 UIC computer science department faculty and students 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 talk on StarLight and iGrid 2002.

2.C. Research Training

There is clearly a collaborating team of professors, staff and engineers from UIC, ANL, NU, NCSA, MREN and Indiana University involved with Euro-Link (and STAR TAP/StarLight), facilitating greater advances in global networking than a single-investigator effort would afford. Moreover, if we count all the people involved in Euro-Link, not just the management team in the Chicago area, the involvement extends nationally (NLANR, NGI Fednets, Internet2) and internationally. All the people working on Euro-Link (and STAR TAP/StarLight) related projects are involved in furthering its goals, either within their respective disciplines, or by helping us better understand the benefits and future directions of long, high-bandwidth networks.

2.D. Education/Outreach

Our primary education and outreach activities include web documentation, journal articles, and conference presentations and demonstrations. We also provide PowerPoint presentations, and other teaching materials to collaborators to give presentations at many conferences, government briefings, etc.

Since 1986, EVL has partnered with NCSA, ANL, and more recently NU/iCAIR, in ongoing efforts to develop national/international collaborations at major professional conferences, notably ACM SIGGRAPH, ACM/IEEE Supercomputing (SC), IEEE High Performance Distributed Computing (HPDC) and INET, as well as Internet2 meetings. We have participated in European conferences (e.g., TERENA's LambdaGrid Workshop and INET 2001 in Stockholm), NORDUnet annual meetings and a UKERNA seminar on optical networking. Our success has been in the development of teams, tools, hardware, system software, and human interface models on an accelerated schedule to enable multi-site collaborations for complex problem solving. We have received a great deal of media attention for our work; news articles are posted at <www.euro-link.org/PUBLICATIONS/>.

Past events we have organized include iGrid '98 at SC'98, iGrid 2000 at INET 2000 in Yokohama, and iGrid 2002 in Amsterdam. We participated in the SC'02 conference <http://www.sc-conference.org/sc2002/> in Baltimore, November 16-22, 2002 and ran demos to promote the goals of StarLight. We also encourage our international collaborators to develop conference events (such as iGrid) to showcase meritorious applications among their own researchers.

EVL also collaborates with the NSF-funded GriPhyN, NEES and, most recently, Earthscope, initiatives, as well as the CAVE Research Network Users' Society (CAVERNUS) and the GeoWall consortium.

3. Publications and Products

3.A. Journals/Papers

In addition to overseeing the growth and development of Euro-Link, the faculty, staff and students of EVL/UIC are users of Euro-Link. Specifically, EVL/UIC develops tele-immersion tools and applications with collaborators worldwide. In addition, EVL/UIC studies the effects of long, fat networks on application performance. To this end, EVL is building tools into its CAVERNsoft communications library to facilitate optimal use, and is developing applications-level network performance analysis tools to help next-generation networks meet the high-bandwidth, quality-of-service (QoS) and connectivity needs of academic researchers running high-performance scientific applications. The publications listed below are relevant to all these activities.

Luc Renambot, Tom van der Schaaf, Henri E. Bal, Desmond Germans, Hans J.W. Spoelder, "Griz: Experience with Remote Visualization over an Optical Grid," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

Rajvikram Singh, Jason Leigh, Thomas A. DeFanti, "TeraVision: A High Resolution Graphics Streaming Device for Amplified Collaboration Environments," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

E. He, J. Alimohideen, J. Eliason, N. Krishnaprasad, J. Leigh, O. Yu, T. A. DeFanti, "QUANTA: A Tooklit for High Performance Data Delivery over Photonic Networks," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

Shalini Venkataraman, Jason Leigh, Tom Coffin, "Kites Flying In and Out of Space - Distributed Physically-based Art on the GRID," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

Charles Zhang, Jason Leigh, Thomas A. DeFanti, Marco Mazzucco, Robert Grossman, "TeraScope: Distributed Visual Data Mining of Terascale Data Sets Over Photonic Networks," Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

Thomas A. DeFanti, Jason Leigh, Maxine D. Brown, Daniel J. Sandin, Oliver Yu, Chong Zhang, Rajvikram Singh, Eric He, Javid Alimohideen, Naveen K. Krishnaprasad, Robert Grossman, Marco Mazzucco, Larry Smarr, Mark Ellisman, Phil Papadopoulos, Andrew Chien, John Orcutt, "Teleimmersion and Visualization with the OptIPuter," Proceedings of the 12th International Conference on Artificial Reality and Telexistence (ICAT 2002), The University of Tokyo, Japan, December 3-6, 2002, to be published by Ohmsha/IOS Press, <www.ic-at.org>.

Jason Leigh, Andrew Johnson, Kyoung Park, Atul Nayak, Rajvikram Singh, Vikas Chowdhry, Thomas A. DeFanti, "Amplified Collaboration Environments," VizGrid Symposium, Tokyo, November 2002.

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. <http://www.evl.uic.edu/cavern/papers/cluster2002.pdf>

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, March 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, March 2002.

3.B. Books/Publications

Cees de Laat, Maxine Brown and Tom DeFanti, (guest editors), Journal of Future Generation Computer Systems (FGCS), Elsevier Science Press, Volume 19, Issue 6, August 2003, (to appear).

3.C. Internet Dissemination

www.euro-link.org

3.D. Other Specific Products

Other than the information reported here, we have not developed any other specific product of significance.

4. Contributions

4.A. Contributions within Discipline

Euro-Link, by its very nature, is interdisciplinary. There is clearly a fine team of computer scientists, computational scientists and networking engineers involved with Euro-Link, facilitating greater advances in global networking than single-investigator efforts could produce. Euro-Link developed its management team in the Chicago area (EVL/UIC, ANL, MREN, iCAIR/NU), and leveraged the efforts of national networking groups (NLANR, NGI networks, Internet2) and international NRN technical and administrative contacts.

4.B. Contributions to Other Disciplines

Within the Computational Science and Engineering and the Computer Science communities, Euro-Link is a necessary and integral part of application advances and technological innovations. Euro-Link also impacts the network community by providing an infrastructure to study long-distant, high-bandwidth networks. NLANR is

working with Euro-Link/STAR TAP/StarLight on network measurement and web caching. ESnet and CANARIE are working with Euro-Link/STAR TAP/StarLight on 6TAP. Networking companies, such as Cisco and Juniper, have given Euro-Link/STAR TAP/StarLight router donations.

EVL not only manages the Euro-Link, STAR TAP and StarLight facilities, but is also one of Euro-Link's major users. EVL's networking interest is a natural outgrowth of its focus on visualization; i.e., EVL is not only interested in producing graphic images and display technologies, but is interested in moving visualizations over networks. In the early '90s EVL focused on distributed computing (connecting visualization/virtual-reality technologies to vector and parallel remote supercomputers), and by the mid-90s EVL focused on tele-immersion (collaborative virtual reality over networks, an extension of the "human/computer interaction" paradigm to "human/computer/human collaboration"). Now, in the '00 decade, EVL is focusing on latency issues in tele-immersion, particularly over ultra-high-speed networks. Transoceanic tele-immersion software and middleware must provide for latency-tolerant and time-shifted usage as well as archival storage and content-based retrieval of multi-flow, multi-participant virtual-reality sessions. EVL is examining the problems of managing these flows in real time and creating adaptive latency-tolerant solutions for international distances.

4.C. Contributions to Human Resource Development

We promote Euro-Link through web documentation, journal articles, demonstrations and presentations at major networking conferences (e.g., Supercomputing, INET, HPDC and Internet2), videotapes, PowerPoint presentations and other instructional material. We teach the infrastructure, the grid advancements, the technological innovations and the application advancements that global connectivity enables.

Euro-Link and NSF's companion HPIIS programs have helped change the way international science is done, by providing a persistent infrastructure for global collaboration. The HPIIS programs, through STAR TAP, and now StarLight, have enabled a worldwide community of application scientists, computer scientists, networking engineers and artists. STAR TAP has a mailing list of ~600 <stars@startap.net> individuals, from academia, government and industry, interested in information about international networking developments. The success of iGrid '98 and iGrid 2000 sparked the interest of the Europeans to sponsor iGrid 2002.

While we have no quantitative metrics to evaluate Euro-Link's role in education/human resources development, we have many testimonials from American and European users and NRN management. We have documented more than 80 applications on the Euro-Link web site, implying that hundreds of networking engineers, application programmers and discipline scientists are involved in Euro-Link-related international high-performance networking and applications development.

4.D. Contributions to Resources for Research and Education

In Section 4.B (Contributions to Other Disciplines), we note that Euro-Link is a necessary and integral part of application advances and technological innovations for the Computational Science and Engineering and the Computer Science communities, as well as of major interest to research network engineers. Euro-Link is a major -- and unique -- resource for Science and Technology. Euro-Link/STAR TAP/StarLight is an infrastructure and proving ground in which to implement new network engineering solutions to advance the state of the art.

4.E. Contributions Beyond Science and Engineering

Because of Euro-Link/STAR TAP/StarLight's interest in QoS, IPv6 and lambda switching, we have recently gotten inquiries from network equipment manufacturers and optical networking providers about partnering with us to create and showcase a marketplace for wavelength-based network services and products. We look forward to working with these companies and introducing them to the Nation's foremost university and Federal laboratory networking engineers, computer programmers and applications scientists, who are developing and using today's evolving grid technologies.

STAR TAP/StarLight (and Euro-Link) is evolving into a national/international optical-networking proving ground, to demonstrate an entirely new information architecture whereby bandwidth becomes the *enabling,* rather than gating, technology. Our research colleagues will showcase new optical networking capabilities to their collaborators

as well as Federal and corporate sponsors, thereby building new opportunities. Our users expect us to grow in capacity and sophistication, and we look forward to the engineering challenges ahead.

5. Special Requirements

5.A. Objectives and Scope

A brief summary of the work to be performed during the next year of support if changed from the original proposal.

NSF recently approved a one-year extension and supplement to Euro-Link, for the period April 1, 2003 – March 31, 2004. Funds will be used to continue Euro-Link, providing increased-bandwidth routed ("Layer 3") production HPIIS services between the US and GÉANT in Europe, as well as sustained support for TransLight, a global switched experimental ("Layer 2") networking trial already in progress. Supplement submitted to NSF in January.

5.B. Special Reporting Requirements

Do special terms and conditions of your award require you to report any information that you have not yet reported?

No.

5.C. Unobligated Funds

Do you anticipate that more than twenty percent of the funds under your NSF award will remain unobligated at the end of the period for which NSF currently is providing support?

No.

5.D. Animals, Biohazards, Human Subjects

Has there been any significant change in animal care and use, biohazards, or use of human subjects from what was originally approved (or approved later)?

No.

6. Summary of Award Expenditures (Year 4)

Available upon request.

7. Appendix: Euro-Link Applications

This appendix represents a major UIC/EVL effort to compile a list of meritorious applications involving US/Euro-Link researchers. This listing is cumulative, containing applications gathered over the past four years, since the inception of the Euro-Link award.

Data was gathered from documents submitted by NRNs, news lists and independent research. This information appears on the Euro-Link, STAR TAP and StarLight web sites. Applications are organized here (and on the Euro-Link web site) by NRN; they are organized by scientific discipline on the STAR TAP/StarLight sites.

7.A. CERN

WorldGRID

USA, CERN and participating EU DataGrid and DataTAG partners: *University of Wisconsin, Fermilab, University of Florida, Hampton University, Brookhaven National Laboratory, Caltech, University of Texas, University of Southern California/Information Sciences Institute, University of Chicago, Penn State University, University of California San Diego, Indiana University, Harvard University, Stanford University, Argonne National Laboratory, Northwestern University, Boston University, Johns Hopkins University, USA; EU DataTAG and EU Data Grid projects, Europe*

Ian Foster, University of Chicago and Argonne National Laboratory, foster@cs.uchicago.edu
Paul Avery, University of Florida, avery@phys.ufl.edu

This project is establishing an international Virtual Data Grid Laboratory (iVDGL) of unprecedented scale and scope, comprising heterogeneous computing and storage resources in the US, Europe and ultimately other regions linked by high-speed networks, and operated as a single system for the purposes of interdisciplinary experimentation in Grid-enabled, data-intensive scientific computing. This laboratory will drive the development, and transition to everyday production use of Petabyte-scale virtual data applications required by frontier computationally oriented science.

Project members are seizing the opportunities presented by the convergence of rapid advances in networking, information technology, Data Grid software tools, and application sciences, as well as substantial investments in data-intensive science now underway in the US, Europe, and Asia. Experiments conducted in this unique international laboratory will hopefully influence the future of scientific investigation by bringing into practice new modes of transparent access to information in a wide range of disciplines, including high-energy and nuclear physics, gravitational wave research, astronomy, astrophysics, earth observations, and bioinformatics.

iVDGL experiments will also provide computer scientists developing data grid technologies with invaluable experience and insight, therefore influencing the future of data grids themselves. A significant additional benefit of this project is that it will empower a set of universities who normally have little access to top-tier facilities and state-of-the-art software systems, hence bringing the methods and results of international scientific enterprises to a diverse, world-wide audience.

Note: The Grid Laboratory Uniform Environment (GLUE) collaboration was created in February 2002 to provide a focused effort to achieve interoperability between the US (iVDGL, GriPhyN and PPDG) and the European (EDG, DataTAG, CrossGrid, etc.) physics Grid projects. The scope includes the definition of a set of software, configuration recommendations, documentation and test suites that cover the basic interoperability requirements.

www.ivdgl.org

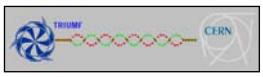

ATLAS Canada LightPath Data Transfer Trial (iGrid 2002)

TRIUMF, Carleton University, University of Victoria, University of Alberta, University of Toronto, Simon Fraser University, BCNet, CANARIE, Canada; CERN, Switzerland; Universiteit van Amsterdam, The Netherlands

Corrie Kost, TRIUMF, Canada, kost@triumf.ca

The Lightpath Trial is attempting to transmit 1 TeraByte (TB) of ATLAS data from TRIUMF, Canada's National Laboratory for Particle and Nuclear Physics, to CERN at high speed. ATLAS (A Toroidal LHC ApparatuS) is one of the experiments being constructed for the Large Hadron Collider (LHC) at CERN.

Using SURFnet's entire 2.5Gbps capacity between StarLight and NetherLight, as well as the planned 2.5Gbps links from Canada to StarLight, CERN to StarLight (EU DataTAG) and NetherLight to CERN, an end-to-end lightpath is being built between TRIUMF in Vancouver and CERN. The goal is to transfer a 1TB of Monte Carlo data between a cluster at TRIUMF and a cluster at CERN in under two hours.

*Note: Subsequently, a Terabyte of research data was transferred in less than a DVD/minute over a "light path" extending 12,000 kilometers from TRIUMF in Vancouver to CERN. The demonstration required dedicated portions of fiber-optic networks, spanning one provincial (British Columbia's BCnet) and two national research and education networks (CA*net4 and SURFnet) to establish the on-demand private network. Experiments were also conducted on the Chicago to CERN DataTAG link. The project culminated in establishing the first large-scale end-to-end "light path".* ***Researchers were able to transfer a Terabyte of research data from disk-to-disk in under 3 hours – this is equivalent to transferring a full CD in less than 8 seconds (or a full-length DVD movie in less than 1 minute).*** *Using* bbftp*, they could transfer 60GB files in 10 minutes; 1TB is about 17 of these files, or about 170 minutes transfer time. Using Tsunami, they achieved slightly better rates. Peak transfer rates in excess of 1Gb were achieved, twice the previous known record for this distance. This is the first establishment of an inter-domain end-to-end "light path" dedicated for a research application. The "light path" directly connecting TRIUMF and CERN is the longest known single hop network spanning the distance from Vancouver to Geneva via StarLight, NetherLight and DataTAG. (Note: Performance was affected by using 10GigE alpha-test cards from Intel. Subsequent testing with beta-test cards has shown that they perform roughly 4 times better.)*

www.triumf.ca

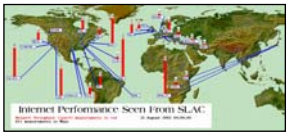

Bandwidth Challenge from the Low-Lands (iGrid 2002)

Europe: CERN, Switzerland; Daresbury Laboratory, Manchester University, Rutherford Appleton Laboratory-Oxford, University College London, UK; Institut National de Physique Nucléaire et de Physique des Particules (IN2P3), France; Istituto Nazionale di Fisica Nucleare (INFN), Milan, Italy; INFN, Rome, Italy; NIKHEF, The Netherlands

Japan: APAN, KEK High Energy Accelerator Research Organization, The Institute of Physical and Chemical Research (RIKEN)/ RIKEN Accelerator Research Facility

North America: TRIUMF, Canada; Argonne National Laboratory, Brookhaven National Laboratory, Caltech, ESnet, Fermi National Accelerator Laboratory, Internet2, Thomas Jefferson National Accelerator Facility (JLab), Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, LBNL National Energy Research Scientific Computing Center (NERSC), NASA Goddard Space Flight Center, Oak Ridge National Laboratory, Rice University, San Diego Supercomputer Center, University of California, San Diego, Stanford Linear Accelerator Center, Stanford University, Sun Microsystems University of Delaware, University of Florida, University of Michigan, University of Texas at Dallas, University of Wisconsin-Madison

Antony Antony, Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF), The Netherlands, antony@nikhef.nl
R. Les Cottrell, Stanford Linear Accelerator Center (SLAC), USA, cottrell@slac.stanford.edu

The avalanche of data already being generated by and for new and future High Energy and Nuclear Physics (HENP) experiments demands new strategies for how the data is collected, shared, analyzed and presented. For example, the SLAC BaBar experiment and JLab are each already collecting over a TB/day, and BaBar expects to increase by a factor of two in the coming year. SLAC and Fermilab's CDF (Collider Detector at Fermilab) and D0 experiments have already gathered well over a petabyte of data, and the CERN Large Hadron Collider (LHC) experiments expect to collect over 10-million terabytes. The strategies being adopted to analyze and store this unprecedented amount of data is the coordinated deployment of Grid technologies, such as those being developed for the Particle Physics Data Grid (PPDG) and the Grid Physics Network (GriPhyN). It is anticipated that these technologies will be deployed at hundreds of institutes that will be able to search out and analyze information from an interconnected worldwide grid of tens of thousands of computers and storage devices. This, in turn, will require the ability to sustain, over long periods, the transfer of large amounts of data among collaborating sites with relatively low latency.

This project demonstrates the current data-transfer capabilities to several sites worldwide that have high-performance links. The iGrid 2002 site acted as a HENP Tier 0 or Tier 1 site (an accelerator or major computation site) in distributing copies of raw data to multiple replica sites. The demonstration is over real live production networks with no efforts to manually limit other traffic. The results are displayed in real time. Researchers investigate/demonstrate issues regarding TCP implementations for high-bandwidth long-latency links, and create a repository of trace files of a few interesting flows. These traces, valuable to projects like EU DataTAG, help explain the behavior of transport protocols over various production networks.

www-iepm.slac.stanford.edu/monitoring/bulk/igrid2002

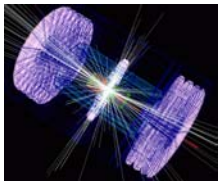

Bandwidth Gluttony: Distributed Grid-Enabled Particle Physics Event Analysis (iGrid 2002)

Argonne National Laboratory, Caltech, USA; CERN, Switzerland

Julian Bunn, Caltech, julian@cacr.caltech.edu
William E. Allcock, Argonne National Laboratory, allcock@mcs.anl.gov

Requests for remote virtual data collections are issued by Grid-based software that is itself triggered from a customized version of the High-Energy Physics (HEP) analysis tool called ROOT. These requests cause the data to be moved across a wide-area network using both striped and standard GridFTP servers.

For iGrid, distributed databases located at ANL, StarLight, Caltech, CERN and other HEP institutions are used. As the collections are instantiated on the client machine in Amsterdam, ROOT analyzes the data, rendering the results in real time. The virtual data collections are catalogued using the Globus Replica Catalogue. This scheme is a preview of a general Grid-Enabled Analysis Environment that is being developed for CERN's Large Hadron Collider (LHC) experiments.

In a closely related part of the iGrid 2002 demonstration, an attempt was made to saturate a 10Gbps (OC-192) link between Amsterdam, ANL and StarLight and a 2.5Gbps (OC-48) link between Amsterdam and CERN, by using striped GridFTP channels and specially tuned TCP/IP stacks. In this test, memory-cached data, in contrast to the file-based ROOT part of the demonstration, was used.

http://pcbunn.cacr.caltech.edu/iGrid2002/demo.htm

Fine Grained Authorization for GARA Automated Bandwidth Reservation (iGrid 2002)

University of Michigan, USA; CERN, Switzerland

William A. Adamson, University of Michigan, andros@umich.edu

This iGrid 2002 demonstration shows modifications to the Globus Toolkit General-purpose Architecture for Reservation and Allocation (GARA). Specifically, it shows (1) a fine-grained cross-domain authorization for GARA that leverages existing security and group services and (2) the elimination of the need for long-term Public Key credentials, currently required by the system. Also shown was a secure and convenient Web interface for making reservation requests based on Kerberos credentials.

GARA modifications are demonstrated by reserving bandwidth for a videoconference application running between sites with distinct security domains. Traffic generators overload the router interface servicing the video receiver, degrading the video quality when bandwidth is not reserved. Successful reservation occurs only when the reservation parameters are within policy bounds, and when the requestor is a member of the required groups. At reservation start time, the end-domain Cisco ingress routers are configured with the appropriate Committed Access Rate (CAR) limit, which marks the packets and polices the flow. The participating routers are statically configured with Weighted Random Early Dectection (WRED), Cisco's implementation of the Random Early Detection class of congestion avoidance algorithms. The router configurations are removed at the reservation end.

www.citi.umich.edu/projects/qos

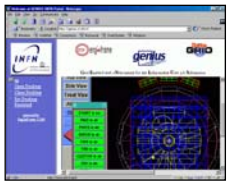

GENIUS (iGrid 2002)

Italy: Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Catania; NICE srl, Camerano Casasco; Università di Catania

CERN, Switzerland

Roberto Barbera, INFN, roberto.barbera@ct.infn.it
Peter Kunszt, CERN, peter.kunszt@cern.ch

The grid portal GENIUS (Grid Enabled web environment for site Independent User job Submission) is an interactive data management tool being developed on the EU DataGrid testbed. At iGrid 2002, researchers demonstrated GENIUS's data movement and discovery, security mechanisms and system monitoring techniques, as well as optimization and fail-safe mechanisms—for example, how to find network optimized files and how to detect system failure.

https://genius.ct.infn.it

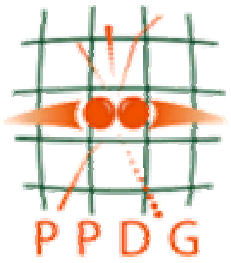

Particle Physics Data Grid (PPDG)

USA, CERN and participating EU DataGrid and DataTAG partners: *Argonne National Laboratory, Brookhaven National Laboratory, Caltech, Fermilab, Lawrence Berkeley National Laboratory, San Diego Supercomputer Center, SLAC, Thomas Jefferson National Accelerator Facility, UCSD, University of Florida, University of Wisconsin, Harvard University, USA; Global Grid Forum (worldwide); EU DataGrid project, Europe; HENP InterGrid Coordination Board (HICB) and HICB Joint Technical Board (worldwide) <www.hicb.org>.*

The PPDG collaboration was formed in 1999, with initial funding from the NGI initiative and later from the DOE MICS and HENP programs, to develop, acquire and deliver vitally needed Grid-enabled tools for data-intensive requirements of particle and nuclear

physics. Novel mechanisms and policies are to be vertically integrated with Grid middleware and experiment-specific applications and computing resources to form effective end-to-end capabilities.

PPDF is a collaboration of computer scientists with a strong record in distributed computing and Grid technology, and physicists with leading roles in the software and network infrastructures for major high-energy and nuclear experiments. Together they have the experience, knowledge and vision in the scientific disciplines and technologies required to bring Grid-enabled data manipulation and analysis capabilities to the desk of every physicist.

PPDG goals and plans are guided by the immediate, medium-term and longer-term needs and perspectives of the physics experiments, some of which will run for at least a decade from 2006 and by the R&D agenda of the computer science projects involved in PPDF and other Grid-oriented efforts. PPDG actively participates in iVDGL together with GriPhyN as a three-pronged approach to data grids for US physics experiments. It also works closely with complementary data grid initiatives in Europe and beyond: Global Grid Forum, EU DataGrid and as part of the HENP InterGrid Coordination Board (HICB) and HICB Joint Technical Board.

www.ppdg.net

European DataGrid

Led by CERN, European Organization for Nuclear Research, five main and fifteen associated partners. European research agencies include the European Space Agency (ESA), France's Centre National de la Recherche Scientifique (CNRS), Italy's Istituto Nazionale di Fisica Nucleare (INFN), the Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF) and UK's Particle Physics and Astronomy Research Council (PPARC). Associated partners include the Czech Republic, Finland, France, Germany, Hungary, Italy, the Netherlands, Spain, Sweden and the United Kingdom.

Fabrizio Gagliardi, CERN, Fabrizio.Gagliardi@cern.ch

The DataGrid project, funded by the European Union, aims to set up a computational and data-intensive grid of resources for the analysis of data coming from scientific exploration. Next-generation science will require coordinated resource sharing, collaborative processing and analysis of huge amounts of data produced and stored by many scientific laboratories in several institutions.

The project will devise and develop scalable software solutions and testbeds in order to handle many PetaBytes of distributed data, tens of thousands of computing resources (processors, disks, etc.), and thousands of simultaneous users from multiple research institutions. The project spans three years, from 2001 to 2003, with over 200 scientists and researchers involved. The first and main challenge facing the project is the sharing of huge amounts of distributed data over the current network infrastructure. The DataGrid project will rely upon emerging computational Grid technologies that are expected to make a giant computational environment out of a distributed collection of files, databases, computers, scientific instruments and devices.

http://www.eu-datagrid.org

GEANT 4

CERN (primary), Bulgaria, Canada, Finland, France, Germany, Hungary, India, Italy, Japan, Netherlands, Russia, Spain, Switzerland, United Kingdom, United States

See website for contact names (click Support, click Contacting Geant4)

Geant4 (GEometry ANd Tracking) is a toolkit for the simulation of the passage of particles through matter. It was designed for the next-generation of High Energy Physics (HEP) experiments, with primary requirements from the LHC, the CP violation and

heavy ion experiments. Its application areas include HEP and nuclear experiments, medical, accelerator and space physics studies. Geant4 exploits advanced software engineering techniques and object-oriented technology to achieve the transparency of the physics implementation, and hence provide the possibility of validating the physics results. It was developed by RD44, a worldwide collaboration of about 100 scientists from over 40 institutions and laboratories participating in more than 10 experiments in Europe, Russia, Japan, Canada and the United States.

http://geant4.web.cern.ch/geant4

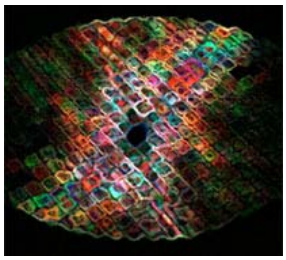

GiDVN: Global Internet Digital Video Network (iGrid 2000)

International Center for Advanced Internet Research (iCAIR), Northwestern, USA; Digital Video Working Group, Coordinating Committee for International Research Networks (DVWG,CCIRN), worldwide membership

Joe Mambretti, iCAIR, Northwestern University, j-mambretti@nwu.edu

GiDVN projects are enhancing media capabilities for the next-generation Internet, enabling new applications to interoperate throughout the world.

Peter Marshall, CANARIE Inc., Canada
Olivier Martin, Paolo Moroni, Philippe Galvez, Joop Joosten, CERN
Kazunori Konishi, APAN, Japan
Shuichi Matsumoto, Masahiro Wada, Shigeyuki Sakazawa, Yasuhiro Takishima, Tetsusi Yamashita, KDD, Japan
Jaehwa Lee, Seungmo Choe, APAN-KR, Korea
Yung Yi, Yanghee Choi, Wang Lijing, Seoul National University, Korea
Sungkwan Youm, Korea University, Korea
Egon M. Verharen, SURFnet, The Netherlands
Cesar Olvera, DGSCA-UNAM, Mexico
Manjeet Singh, Francis Lee, SingAREN, Singapore
Artur Serra, Sebastia Sallent, Joan Borràs, Universitat Politecnica de Catalunya, Spain
Björn Pehrson, Daniel Forsgren, Royal Institute of Technology, Sweden
Joe Mambretti, Jim Chen, Jeremy Weinberger, Tim Ward, Northwestern University, USA

http://www.icair.org/inet2000, http://iumrs.ms.nwu.edu

Authenticated Quality of Service (QoS)

CERN; University of Michigan, USA; Argonne National Laboratory, USA; Pittsburgh Supercomputing Center, USA

William A. (Andy) Adamson, University of Michigan, andros@umich.edu

This project provides the next step toward providing authenticated, authorized network QoS guarantees. The collaborators are mainly concerned with the signaling for the set-up of network QoS, specifically the design and transmission of authentication and authorization information. They will test with UDP-based applications for which current Cisco QoS configurations work, and will experiment with TCP-based applications.

The UDP traffic generator from University of Michigan's Center for Information Technology Integration (CITI) is being used to test QoS capabilities of network components. The basic form of these initial results will be followed when verifying network resource reservation at the end-point Bandwidth Broker networks.

This project aims to integrate the PKI-based Akenti authorization system into Argonne National Lab's Globus-based bandwidth broker and the Internet2 Middleware working group's LDAP directory schema, to provide a scalable, authenticated means to request network QoS.

http://www.citi.umich.edu/projects/qos/

GriPhyN: Grid Physics Network

CERN, Switzerland; 17 US research laboratories and universities, led by University of Chicago, USA and University of Florida, USA

Paul Avery, University of Florida, avery@phys.ufl.edu
Ian Foster, University of Chicago and Argonne National Lab, foster@cs.uchicago.edu
Harvey B. Newman, CERN and Caltech, newman@hep.caltech.edu

The GriPhyN collaboration is a team of experimental physicists and information technology (IT) researchers who plan to implement the first Petabyte-scale computational environments for data intensive science in the 21st Century. Driving the project are unprecedented requirements for geographically dispersed extraction of complex scientific information from very large collections of measured data: CMS (Compact Muon Solenoid), ATLAS (A Toroidal LHC Apparatus), LIGO (Laser Interferometer Gravitational-wave Observatory), and SDSS (Sloan Digital Sky Survey).

CMS and ATLAS experiments will use the Large Hadron Collider (LHC) at CERN to search for the origins of mass and probe matter at the smallest length scales; LIGO will detect the gravitational waves of pulsars, supernovae and in-spiraling binary stars; and, SDSS (Sloan Digital Sky Survey) will carry out an automated sky survey enabling systematic studies of stars, galaxies and large-scale structure.

To meet these requirements, which arise initially from the four physics experiments involved in this project but will also be fundamental to science and commerce in the 21st Century, GriPhyN will deploy computational environments called Petascale Virtual Data Grids (PVDSs) that meet the data-intensive computational needs of a diverse community of thousands of scientists spread across the globe. GriPhyN is a US National Science Foundation Information Technology Research (ITR) initiative led by the University of Chicago and the University of Florida.

http://www.griphyn.org

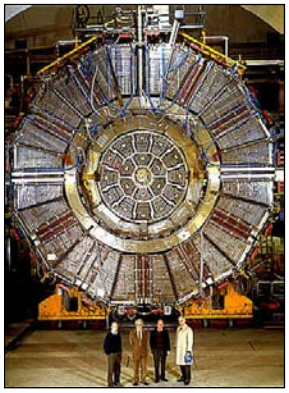

Networked Experiments of the European Laboratory for Particle Physics

CERN; Argonne National Laboratory, USA; California Institute of Technology, USA; Cornell University, USA; Fermilab, USA; Harvard University, USA; Lawrence Berkeley National Laboratory, USA; Massachusetts Institute of Technology, USA; Princeton University; USA; ESnet, USA.

Harvey Newman, Caltech, newman@hep.caltech.edu

CERN provides experimental facilities for particle physics experiments, mainly in high-energy physics (HEP). CERN's current major facility is the Large Electron Positron (LEP) collider in a 27-km tunnel, the largest machine of this type in the world. Four very large experiments in man-made caverns intersect the LEP tunnel, constituting half of CERN's total experimental program for the 1990s. Each of the experiments is carried out by teams of several hundred of physicists from over 50 worldwide institutes.

LEP experiments generate 25 terabytes of data each year, which is stored on magnetic tape cartridges, whereas the Large Hadron Collider (LHC) experiments, expected to commence in 2005, are expected to produce several order of magnitude more data. The sheer volume of the data combined with the complexity of the analysis to be performed, and the requirement that the processing of the data may also be done remotely, places heavy demands on the High Energy & Nuclear Physics (HENP) computing and networking infrastructure, which can only be met by using leading edge technology and services.

http://www.cern.ch

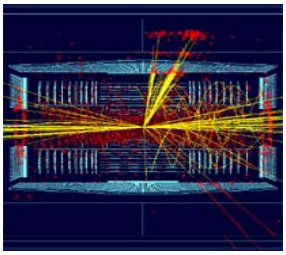

Distributed Particle Physics Research

CERN; Caltech, USA

Harvey Newman, Caltech and CERN, newman@hep.caltech.edu
Julian Bunn, Caltech, julian@cacr.caltech.edu

This application demonstrates remote viewing and analysis of particle physics events. The application is the front end to an engineered object-oriented global system that incorporates grid middleware for authentication and resource discovery, a distributed object database containing several terabytes of simulated events, and a component that enables queries issued by the front-end application to be matched to available computing resources in the system (the matchmaking service).

http://pcbunn.cacr.caltech.edu, http://cmsdoc.cern.ch/orca, http://iguana.web.cern.ch/iguana, http://vrvs.cern.ch

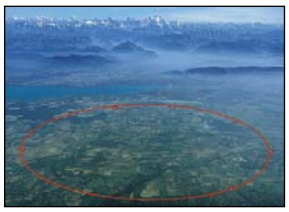

Large Hadron Collider (LHC) project

CERN; Caltech, USA; Others.

Harvey Newman, Caltech, newman@hep.caltech.edu

The LHC is an accelerator that brings protons and ions into head-on collisions at higher energies than ever before, enabling scientists to penetrate still further into the structure of matter, and recreate the prevailing conditions of the early post-"Big Bang" universe.

The LHC is a remarkably versatile accelerator. It can collide proton beams with energies around 7-on-7 TeV and beam crossing points of unsurpassed brightness, providing the experiments with high interaction rates. It can also collide beams of heavy ions such as lead with total collision energy in excess of 1,250 TeV–about 30 times higher than at the Relativistic Heavy Ion Collider (RHIC) under construction at the Brookhaven Laboratory in the US. Joint LHC/LEP operation can supply proton-electron collisions with 1.5 TeV energy, some five times higher than presently available at HERA at the DESY laboratory in Germany. The research, technical and educational potential of the LHC and its experiments is enormous.

Several LHC experiments are being developed; see http://lhc-new-homepage.web.cern.ch

Two LHC detectors, ATLAS and CMS, are described below.

ATLAS (A Toroidal LHC ApparatuS)

CERN; Caltech, USA; Others.

Harvey Newman, Caltech and CERN, newman@hep.caltech.edu

1850 collaborators in 150 institutions around the world are constructing the ATLAS experiment. It will study proton-proton interactions at the Large Hadron Collider (LHC) at CERN. The detector is due to begin operation in the year 2005. ATLAS is designed to improve our fundamental understanding of matter and forces. A prime physics goal of ATLAS is to understand the nature of mass.

http://atlasinfo.cern.ch/Atlas/Welcome.html

Compact Muon Solenoid (CMS)

CERN; Caltech, USA; Others.

Harvey Newman, Caltech and CERN, newman@hep.caltech.edu

The CMS detector is designed to cleanly detect the diverse signatures of new physics at the Large Hadron Collider (LHC). It will do so by identifying and precisely measuring muons, electrons and photons over a large energy range; by determining the signatures of quarks and gluons through the measurement of jets of charged and neutral particles

(hadrons) with moderate precision; and by measuring missing transverse energy flow, which will enable the signatures of non-interacting new particles as well as neutrinos to be identified.

The Globally Interconnected Object Databases (GIOD) Project

Caltech, USA; CERN; Hewlett Packard; Others.

Harvey Newman, Caltech, newman@hep.caltech.edu

Particle physicists are engaged in large international projects to address a massive data challenge, with special emphasis on distributed data access. The data is overwhelming. Even though data from the CMS detector will be reduced by a factor >107, over a petabyte (1015 bytes) of data per year will accumulate for scientific analysis.

The task of finding rare events resulting from the decays of massive new particles in a dominating background is even more formidable. Particle physicists have been at the vanguard of data-handling technology, beginning in the 1940s with eye scanning of bubble-chamber photographs and emulsions, through decades of electronic data acquisition systems employing real-time pattern recognition, filtering and formatting, and continuing on to the petabyte archives generated by modern experiments. In the future, CMS and other experiments now being built to run at CERN's Large Hadron Collider expect to accumulate of order of 100 petabytes within the next decade.

The scientific goals and discovery potential of the experiments will only be realized if efficient worldwide access to the data is made possible. Particle physicists are thus engaged in large national and international projects that address this massive data challenge, with special emphasis on distributed data access. There is an acute awareness that the ability to analyze data has not kept up with its increased flow. The traditional approach of extracting data subsets across the Internet, storing them locally, and processing them with home-brewed tools has reached its limits. Something drastically different is required. Indeed, without new modes of data access and of remote collaboration we will not be able to effectively "mine" the intellectual resources represented in our distributed collaborations.

http://pcbunn.cithep.caltech.edu

MONARC (Models of Networked Analysis at Regional Centers) for LHC Experiments

CERN; Caltech, USA; Others.

Harvey Newman, Caltech, newman@hep.caltech.edu

LHC experiments have envisaged computing models (CM) involving hundreds of physicists doing analysis at institutions around the world. Regional Centers are intended to facilitate access to the data with more efficient and cost-effective data delivery to the groups in each world region, using high-speed national networks. This project intends to study network-distributed computing architectures, data access and data management systems that are major components of the CM, and the ways in which the components interact across networks.

http://www.cern.ch/MONARC/

Virtual Room Videoconferencing System (VRVS)

CERN; Caltech, USA; Others.

Harvey Newman, Caltech, newman@hep.caltech.edu

The VRVS was introduced in early 1997, to provide a low cost, bandwidth-efficient, extensible tool for videoconferencing and collaborative work over networks within the High Energy and Nuclear Physics (HENP) communities, and to some extent, research and education at large.

Since it went into production, deployment of the Web-based system has expanded to include 100s of registered hosts running VRVS software in more than 28 countries. There are now 19 "reflectors" that create the interconnections and manage the traffic flow, at HENP labs and universities in the US and Europe. Virtual Room videoconferencing is regularly employed as part of ATLAS and CMS, and increasingly for other DOE-supported programs. The system is managed by the Caltech L3/CMS group working in collaboration with the CERN IT Division, under a joint project approved by the LHC Computing Board (LCB) in July 1997. Future plans for the system include deployment of additional reflectors to Asia and the spawning of other sets of "Virtual Rooms" in Russia.

http://VRVS.cern.ch/

7.B. Israel IUCC

TelePresence Microscopy (TPM)

Argonne National Laboratory, USA; Technion, Israel

Wayne D. Kaplan, Technion, kaplan@tx.technion.ac.il
Mike Lieberthal, Technion, mtmike@tx.technion.ac.il

TelePresence Microscopy (TPM) enables on-line cooperation between scientists, utilization of unique instrumentation by researchers who may not have access to such resources, and training of students. Video signals from the microscope (specimen view, detectors or peripheral instrumentation) are fed into a server, and can be directly accessed via a web browser capable of server push technology (such as Netscape).

Remote control is possible by sending instructions for magnification, movement of the specimen and focus using the web page interface. A control server sends this information into the microscope and the remote user can get instant feedback. Video conferencing enables discussion between local and remote parties. The local user determine whether the microscope session is open or closed to the general public for viewing, and if the remote user may have remote control. Sharing of data is also possible on-line.

http://www.technion.ac.il/technion/materials/TPM/index.html

The Israel Space Agency - Middle East Interactive Data Archive (ISA-MEIDA)

NASA/GSFC, USA; Tel Aviv University, Israel

Pinhas Alpert, Tel Aviv University, pinhas@cyclone.tau.ac.il

ISA-MEIDA is an Earth observing data center available to the research community and general public through the Internet. It is one of nine international data centers, or nodes, participating in NASA's Earth Observing System Data Information System (EOSDIS).

NASA EOSDIS nodes outside of the US interconnect Canada, Japan, England, Russia and Europe to eight data centers within the US. By 2002, the EOSDIS global archives will contain more than 260 data products, occupying over 3 petabytes (1 petabyte = 10^9 megabytes).

ISA-MEIDA enables Israeli users to access Earth data from the US and other participating data centers. Similarly, international users have access to complementary datasets archived at the Israeli node. These datasets include detailed assessments and forecasts of atmospheric variables such as temperature, wind, humidity, clouds, dust, ozone, other aerosols and gases over the Middle East, as well as earth surface fields like land cover, soil moisture and vegetation. The full data from GEOS-1 and AVHRR is available through the ISA-MEIDA's connection to the global EOSDIS system. The Israel meteorological observations are the property of the Meteorological Service, which gave ISA-MEIDA permission to relay requests for this information.

http://www.nasa.proj.ac.il/

Israeli Mirror of the Los Alamos E-Print Server (ePrint arXiv)

Los Alamos National Laboratory, USA; Tel Aviv University, Israel

Marek Karliner, School of Physics and Astronomy, Tel Aviv University, marek@proton.tau.ac.il

Several theoretical problems in high-energy physics can only be solved using the most powerful computers available. One example is working out, in detail, the consequences of quarks – the force that holds together the most basic constituents of matter. Protons and neutrons are made out of quarks, and the basic theory describing their interactions is known as the Quantum Chromodynamics (QCD).

The most promising approach of solving QCD involves extremely large computer calculations, on the scale of many teraflops over a period of a year or more. Such calculations produce very large amounts of intermediate data, which then need to be processed to obtain the final results, i.e., scientific visualization.

Research results in theoretical high-energy physics are nowadays disseminated mostly through a system of electronic distribution of articles (e-prints), based in the Los Alamos National Laboratory, and run by Paul Ginsparg.

The Los Alamos server was initially set up in 1991 for distribution of articles in theoretical high-energy physics, but has grown over the last few years to encompass all fields of physics and mathematics. An Israeli mirror has been set up to ensure access to the e-print archive, with technical support provided by the Tel Aviv University computer center staff and Los Alamos server staff, working together via the Internet. The mirror server uploads a large amount of new articles from the Los Alamos server daily, and maintains a large archive (several Gbytes) of articles from the past seven years.

http://xxx.tau.ac.il/

Collaborative Learning Over Broadband Internet: "Dialog Through Music"

Israel; Canada

Martin Brooks, National Research Council of Canada, Canada, Martin.Brooks@nrc.ca
Peter Marshall, CANARIE Inc., Canada, marshall@canarie.ca

Maestro Pinchas Zukerman and the National Arts Center partnered with the National Research Council of Canada (NRC), Communications Research Centre Canada (CRC), CANARIE (Canada's Advanced Internet Development Organization) and the Canada-Israel Industrial Research and Development Foundation, in an exploratory application of tomorrow's Internet for international learning.

On October 4, 2000, while the National Arts Centre Orchestra was in the Middle East, Maestro Zukerman led "Dialog Through Music," where Israeli, Palestinian, and Canadian youth shared their emotional and creative responses to Beethoven's famous melody "Ode to Joy." Three groups of ten high school students connected by live video, with Israeli students and Maestro Zukerman in Tel Aviv, Palestinian students in Jerusalem, and Canadian students in Ottawa.

Maestro Zukerman's extensive experience with live video violin teaching, and the potential for the National Arts Center to use broadband learning technologies to build passion for the arts among all Canadians, makes this partnership between Arts and Sciences an exciting opportunity to explore our future.

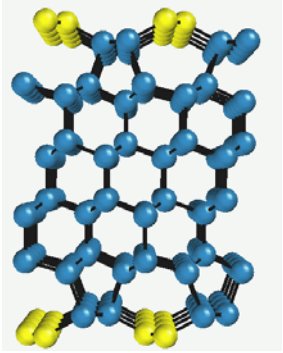

Computer Alchemy Using Virtual Reality

Israel; Others

Joan Adler, Technion, Israel, phr76ja@phjoan.technion.ac.il
Wayne D. Kaplan, Technion, kaplan@tx.technion.ac.il

Computational Condensed Matter Physics and Material Science is a rapidly growing field, fueled in part by the advent of computer alchemy. As computers have become more powerful and algorithms more robust, this practical research tool models molecularly dynamic systems in virtual reality–notably, atomistic material simulations up to several million particles.

Especially suitable for discussing simulation results with experimental collaborators, visualization is also essential for teaching quantum mechanics and condensed matter physics, since most effects on the atomic scale cannot be demonstrated in the usual way.

Current work at Technion involves atomic-level modeling of diamond and aluminium/alumina interfaces. Computational physics and experimental research faculty are collaborating to produce the 3D computer visualization models. High-bandwidth networks are used to transfer the visualizations, and developing protocols for interactive conferencing and discussion.

http://phycomp.technion.ac.il/

Israel One – A Broadband High School Network

North Carolina School of Science and Mathematics, USA; Tel Hai College, Israel

Iris Raviv, Israel's Ministry of Education, Culture and Sports, irisraviv@hotmail.com
Hamutal Hameiri, North Carolina Israel Partnership, hamutalm@netvision.net.il
Itzhak Yuli, Tel Hai College, tzachy@telhay.co.il
Peggy Manring, North Carolina School of Science and Math, manringp@academic.ncssm.edu

Israel One is a high school educational network. It uses distance learning technology and videoconferencing to reach remote schools, especially in the periphery where English and math teachers are scarce. Sixteen schools are expected to participate, representing the cultural diversity of Israel.

This project serves as a testbed for current broadband technology and ATM protocol, and utilizes voice, audio and data. The project provides teacher training to assist students in the use of the broadband technologies, and adopts a student empowerment program. The multi-cultural model of the project encourages cooperation among schools of different geographical locations (Galilee, Negev), old Israelis and new immigrants, Arab, Druze and Jews; secular and religious; development towns, kibbutzim and moshavim.

The program builds upon the experience of the North Carolina Information Highway (NCIH) and the North Carolina School of Science and Mathematics (NCSSM). It is carried out in cooperation with the Israel Internet II program and industrial partners Bezeq (Telecom), Motorola Israel, Bynet, TNN and Accord. International corporate partners include VTEL, PictureTel and GTE.

ProtoMap–A Map of Protein Space: Interactive Web Site for Biological and Biomedical Investigations

Hebrew University, Jerusalem, Israel; Mirror sites: Stanford University, USA; Cornell University, USA

Michal Lineal, Hebrew University, Jerusalem, michal@keonardo.is.huji.ac.il
Nathan Lineal, Hebrew University, Jerusalem, nati@cs.huji.ac.il

This site classifies and groups all of the proteins in the Swiss-Prot and TrEMBL

databases. Transitivity is used to identify homologous proteins, and within each group, every two members are either directly or transitively related. Transitivity is applied restrictively in order to prevent unrelated proteins from clustering together. The classification is done at different levels of confidence, and results in a hierarchical organization of all the proteins.

The resulting classification splits the protein space into well-defined groups of proteins, most of them closely correlated with natural biological families and superfamilies. The hierarchical organization may help to detect finer subfamilies that make up known families of proteins, as well as interesting relations between protein families.

http://www.protomap.cs.huji.ac.il/

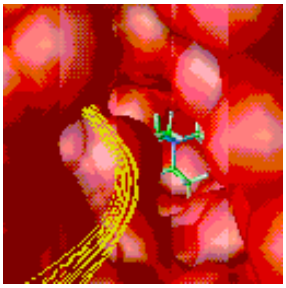

Visualization of Acetylcholinesterase: Nature's Vacuum Cleaner

Cornell University, USA; Weizmann Institute of Science, Israel

Richard Gillilan, Cornell University, richard@tc.cornell.edu
Joel L. Sussman, Weizmann Institute of Science, Joel.Sussman@weizmann.ac.il
Israel Silman, Weizmann Institute of Science, Israel.Silman@weizmann.ac.il

Acetylcholinesterase (AcChoEase) is an enzyme that plays a key role in the human nervous system. In vertebrates, nerve impulses travel from cell to cell by means of chemical messenger. When an electrical impulse reaches the end of one cell, messenger molecules–acetylcholine (AcCho), in this case–are released to diffuse though the fluid-filled, intercellular, synaptic gap. Upon reaching the destination cell, AcCho molecules dock into special receptors triggering a new electrical impulse. Like a vacuum cleaner, the enzyme AcChoEase is constantly sweeping up and hydrolyzing AcCho during this process, so that the whole cycle can begin again.

Chemicals that inhibit the action of AcChoEase are being used in the treatment of glaucoma, myasthenia gravis and, experimentally, Alzheimer's disease. In spite of the ability to exploit the enzyme, its precise mechanism of operation is still a mystery.

The recent solution of the X-ray structure for AcChoEase, places the active catalytic site deep within a gorge-like fold of the protein. Electrostatic computations reveal the enzyme to be a single massive dipole. Such a configuration of charge suggests an electrostatic mechanism for directing the positively charged AcCho into the gorge and towards the active site.

Interactive Simulation in the Field of Plant Nutrition

Penn State University, USA; Tel Aviv University, Israel

Jonathan P. Lynch, Penn State University, JPL4@psu.edu
Amram Eshel, Tel Aviv University, amram@post.tau.ac.il

This US-Israel Bi-national Agricultural Research and Development (BARD) Foundation project involves hypotheses testing and result evaluation, using an interactive graphic model. *SimRoot*, a 3D model developed at Penn State, which graphically describes the 3D deployment of plant root systems in soil. The model follows plant development and will predict the performance of plants under various environmental conditions, depending on their physiological characteristics.

Results are displayed graphically in 3D dynamic images that can be rotated by the viewer. This work was limited until now to users working at the console or on local-area networks. This application requires high-speed computer communication to allow real-time responses between Israel and the US. If successful, a whole new area of applications in cooperative ecological and physiological research and teaching will evolve.

SeaWiF Data for the Eastern Mediterranean and the Middle East Image

Goddard Space Flight Center, NASA, USA; Ben Gurion University, Israel

Gene Carl Feldman, NASA/GSFC, gene@seawifs.gsfc.nasa.gov
Arnon Karnieli, J. Blaustein Institute for Desert Research, Ben Gurion University, karnieli@spamergsfc.nasa.gov

SeaWiFS is a newly developed satellite, integrating advanced technology that acquires multi-channel data over land and sea. Daily SeaWiFS images over the Eastern Mediterranean and the Middle East are received at the J. Blaustein Institute in HRPT format (1.1 km resolution) using a PC-based receiving station. The volume of daily images will range from 30-130Mb, depending on the area of acquisition. These PCs receive raw HRPT-type data and convert it to level-0 to match the specifications of NASA documentation. Subsequently, each image is transferred from receiving stations to NASA/GSFC, to create a global image based on several HRPT stations across the globe.

http://seawifs.gsfc.nasa.gov/SEAWIFS.html

Large-Scale Atomistic Modeling of Semiconductors and Ceramics

University of Georgia, Athens, USA; Technion, Israel; Bar Ilan University, Israel

Joan Adler, Technion, Israel, phr76ja@phjoan.technion.ac.il
Dennis Rapaport, Bar-Ilan University, rapaport@mail.biu.ac.il
David Landau, University of Georgia, Athens, dlandau@uga.edu

The field of computational atomistic modeling is a research area of considerable importance and interest. In particular, the quantitatively accurate atomistic modeling of solids is a rapidly developing discipline with many practical applications. In this project we plan to use various kinds of molecular dynamics and Monte Carlo simulation techniques to explore a range of phenomena associated with semiconductors and ceramics.

The kinds of behavior that will be studied include the interfaces in mixtures of silicon, germanium and carbon, and the processes responsible for generating defects and the subsequent graphitization of diamond; the complexity of these phenomena is such that a detailed atomistic modeling approach is essential. The simulations will be closely coordinated with ongoing experimental studies. Because of the heavy computational requirements and the detailed data analysis involved, the work entails the development of algorithms to support parallel processing together with specialized visualization techniques.

Video of Demand (VoD)

Israel; Internet2,USA.

Hank Nussbacher, IUCC, hank@interall.co.il

VoD is a new technology that is enabled using advanced Internet capabilities. There are a select number of VoD servers located in Israel, Europe and the USA. This site documents all the various high-bit rate VoD servers located on advanced networks worldwide. A VoD system is one that allows a user to start the playback when he/she wishes, as well as pause, rewind and fast-forward that playback. High-bit rate is anything above 1Mbps. This site links to other worldwide VoD web sites, including: ViDe (Video Development Initiative), Digital Video for the next Millennium, Internet2's Digital Video Network and Internet2's Digital Video Initiative.

http://www.internet-2.org.il/vod.html

7.C. NORDUnet

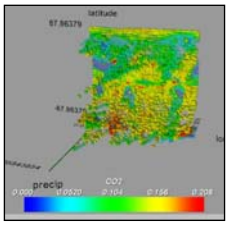

High Performance Data Webs (iGrid 2002)

University of Illinois at Chicago, USA; Dalhousie University, Canada; Imperial College of Science, Technology & Medicine, University of London, UK, Universiteit van Amsterdam, SARA, The Netherlands; Center for Parallel Computers, Royal Institute of Technology, Sweden

Robert Grossman, Laboratory for Advanced Computing (LAC), University of Illinois at Chicago, grossman@uic.edu

DataSpace is a high-performance data web for the remote analysis, mining, and real-time interaction of scientific, engineering, business and other complex data. DataSpace applications are designed to exploit the capabilities provided by emerging domestic and international high-performance networks so that gigabyte and terabyte datasets can be remotely explored in real time. It is an open, standards-based infrastructure that combines data web services, data grid services and semantic web services for remote data analysis and distributed data mining.

The Terra Wide Data Mining (TWDM) testbed consists of high-performance clusters worldwide linked by wide-area advanced networks, providing the data and computer services required.

www.ncdm.uic.edu
www.dataspaceweb.net

PAAPAB (iGrid 2002)

University at Buffalo, Res Umbrae, New York State Center for Engineering Design and Industrial Innovation, University at Buffalo, USA; Tools for Creativity Studio, Interactive Institute, Sweden

Dave Pape, Res Umbrae, dave.pape@acm.org

PAAPAB (Pick An Avatar, Pick A Beat) is a shared virtual-reality disco environment inhabited by life-size puppets, animated by users in CAVEs and ImmersaDesks around the globe. Users can tour the dance floor to see the puppets they animate, dance with the puppets, and dance with avatars of other users. This research focuses on creating interactive drama in virtual reality; that is, immersive stories. PAAPAB serves as a testbed for technology development as well as character and world design.

In addition to PAAPAB, iGrid 2002 attendees were treated to performances with the Interactive Institute's Incarnation of a Divine Being, an environment built using the same software infrastructure. Incarnation is a virtual-reality space based on ancient Greek theater, where people meet and together perform an improvised drama in cyberspace.

http://resumbrae.com/projects/paapab
www.ccr.buffalo.edu/anstey/VR/PAAPAB
www.nyscedii.buffalo.edu

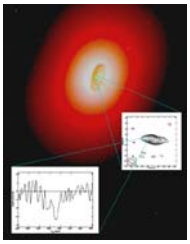

vlbiGrid (iGrid 2002)

Joint Institute for VLBI in Europe (JIVE), The Netherlands; Metsahovi Radio Observatory, Finland; Dept. of Physics and Astronomy at University College London, Jodrell Bank Observatory and Dept. of Physics and Astronomy at University of Manchester, UK; Haystack Observatory, Massachusetts Institute of Technology, USA; Universiteit van Amsterdam, The Netherlands

Steve M. Parsley, Joint Institute for VLBI in Europe (JIVE), The Netherlands, parsley@jive.nl

Very Long Baseline Interferometry (VLBI) is a technique in which an array of physically

independent radio telescopes observes simultaneously to yield high-resolution images of cosmic radio sources. Today, magnetic tape transports data from telescopes to data processors. The European VLBI Network (EVN) has access to multiple data sources that can deliver 1Gbps each and a dedicated supercomputer that can process 16 data streams simultaneously. High-speed networks are enabling the EVN to achieve many-fold improvements in bandwidth.

www.jive.nl
www.jb.man.ac.uk
www.haystack.edu

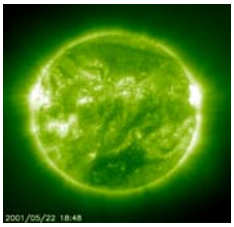

EVL: Alive on the Grid

University of Illinois at Chicago, Indiana University, State University of New York (SUNY), Buffalo, USA; Stichting Academisch Rekencentrum Amsterdam (SARA), Technische Universiteit Eindhoven (TU/e), V2 Lab/V2 Organization Institute for the Unstable Media, Rotterdam, The Netherlands; Interactive Institute - Tools for Creativity Studio, Umea, Sweden; C³ Center for Culture & Communication Foundation, Hungary

Dan Sandin, University of Illinois at Chicago, dan@evl.uic.edu

The Ars Electronica Festival, sponsored by the Ars Electronica Center (AEC), is an international, large-scale annual art festival held annually in Linz, Austria. In 2001, the AEC festival commissioned the University of Illinois at Chicago's Electronic Visualization Laboratory to produce an original virtual reality art application for the CAVE. The resulting work, "EVL: Alive on the Grid," is a collection of shared and tele-immersive virtual-reality environments linked by the Grid (a collection of networks, computers and virtual reality displays spanning the globe). During the five-day event, participants from three American and five European sites interacted in the virtual worlds.

http://www.aec.at/festival2001/

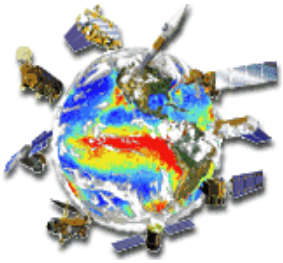

Earth Observing System (EOS)

NASA ,USA; worldwide collaborators. Euro-Link NRN-linked institutions in Denmark, Finland, France, Norway, Sweden

Michael King, Goddard Space Flight Center, king@climate.gsfc.nasa.gov

The Earth Observing System (EOS) is the centerpiece of NASA's Earth Science Enterprise (ESE). It consists of a science component and a data system supporting a coordinated series of polar-orbiting and low inclination satellites for long-term global observations of the land surface, biosphere, solid Earth, atmosphere and oceans. The EOS Project Science Office (EOSPSO) is committed to helping bring program information and resources to program scientists and the general public alike.

http://eospso.gsfc.nasa.gov

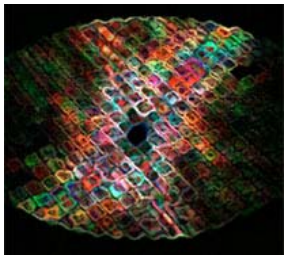

GiDVN: Global Internet Digital Video Network (iGrid 2000)

International Center for Advanced Internet Research (iCAIR), Northwestern, USA; Digital Video Working Group, Coordinating Committee for International Research Networks (DVWG, CCIRN), worldwide membership

Joe Mambretti, iCAIR, Northwestern University, j-mambretti@nwu.edu

GiDVN projects are enhancing media capabilities for the next-generation Internet, enabling new applications to interoperate throughout the world.

Peter Marshall, CANARIE Inc., Canada
Olivier Martin, Paolo Moroni, Philippe Galvez, Joop Joosten, CERN
Kazunori Konishi, APAN, Japan
Shuichi Matsumoto, Masahiro Wada, Shigeyuki Sakazawa, Yasuhiro Takishima, Tetsusi

Yamashita, KDD, Japan
Jaehwa Lee, Seungmo Choe, APAN-KR, Korea
Yung Yi, Yanghee Choi, Wang Lijing, Seoul National University, Korea
Sungkwan Youm, Korea University, Korea
Egon M. Verharen, SURFnet, The Netherlands
Cesar Olvera, DGSCA-UNAM, Mexico
Manjeet Singh, Francis Lee, SingAREN, Singapore
Artur Serra, Sebastia Sallent, Joan Borràs, Universitat Politecnica de Catalunya, Spain
Björn Pehrson, Daniel Forsgren, Royal Institute of Technology, Sweden
Joe Mambretti, Jim Chen, Jeremy Weinberger, Tim Ward, Northwestern Univ., USA

http://www.icair.org/inet2000, http://iumrs.ms.nwu.edu

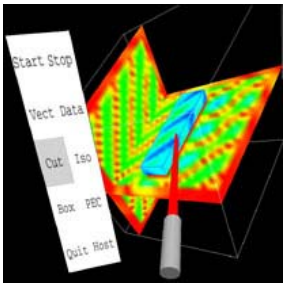

Steering and Visualization of Finite-Difference Code on a Computational Grid (iGrid 2000)

University of Houston, USA; Royal Institute of Technology, Sweden

Lennart Johnsson, University of Houston, johnsson@cs.uh.edu
Erik Engquist, Royal Institute of Technology, erike@pdc.kth.se
Per Öster, Royal Institute of Technology, per@pdc.kth.se

This application enables computational steering of electromagnetic simulations across distributed resources using interactive visualization in a virtual-reality environment. To handle the large computational requirements of both simulation and visualization, the system is distributed across multiple machines using Globus.

http://www.pdc.kth.se/projects/GEMSviz

WITAS Multi-Modal Conversational Interface

University of Stanford, USA; Linköping University, Sweden

Oliver von Klopp Lemon, Stanford University, lemon@csli.stanford.edu
Erik Sandewall, Linköping University, erisa@ida.liu.se

The WITAS Unmanned Aerial Vehicle (UAV) under development at Sweden's Linköping University is an autonomous mobile helicopter with onboard AI, adjustable with respect to the operating environment and operator decisions. A team at Stanford University's Center for the Study of Language and Information (CSLI) is collaborating to build a multi-modal communication interface to this robot, capable of complex dialogues about the UAV's tasks and state, and about situations as they unfold on the ground.

The interface supports dialogues between the operator and the UAV using natural conversational language. The multi-modal aspects of the interface derive from the ability to combine speech, text, graphics, gestures, live video, and sensor data in the same communication. The interface is also designed to be "conversational," in the sense that the multi-modal dialog is planned and managed, and that different display and communication strategies are activated in different dialog and resource/time-bounded contexts. Currently, the team is using the Open Agent Architecture to manage communicating processes, Nuance for speech recognition, and Gemini for NL parsing.

http://www-csli.stanford.edu/semlab/witas.html, http://www.ida.liu.se/ext/witas/

Learn 2: A Network of Incubator Spaces for Developing & Designing Environments for Lifelong Learning

Stanford University, USA; University of Oslo, Norway; Royal Institute of Technology (KTH), Sweden; Uppsala University, Sweden; Roskilde University, Denmark

Gunnar Liestol, University of Oslo, gunnar.liestol@media.uio.no

The LEARN2 project develops resources, procedures, conventions and production methods for high-bandwidth multimedia learning, and its related knowledge distribution

on the Web. An optimal format for lifelong learning of cross-disciplinary topics (based in the humanities) will be established and deployed in a series of actual courses.

The project takes into account that further progress in the field of multimedia learning environments must consider the formal shaping – the rhetoric and design – of the medium, and appropriate additional and closely-related pedagogical strategies to achieve optimal exploitation of the continued technological improvements.

The three goals of the project are: (1) Establishing an optimal network of "incubator sites," or high-bandwidth multimedia learning environments, (2) Focus on the development and design of high capacity multimedia courses, conventions and genres for learning within these spaces, and (3) Conduct research and evaluation of the technological, rhetorical and pedagogical solutions employed.

http://www.media.uio.no/learn2

Global Observation Information Network (GOIN) Demonstration in Stockholm

NORDUnet, Norway; NASA NREN, USA; APAN, Asia

Peter Villemoes, NORDUnet, Peter.Villemoes@adm.nordu.net

The Global Observation Information Network (GOIN) is a USA/Japan effort to strengthen bilateral cooperation in Earth observation information networks, involving both satellite and in-situ data.

In 1999, NORDUnet cooperated with NASA/NREN and APAN to support live demonstrations of the GOIN in Stockholm during the Committee on Earth Observation Satellites (CEOS) Plenary, hosted by the European Meteorological Satellite (EUMETSAT) organization, with the support of the Swedish Space Corporation (SSC). CEOS <www.ceos.org> is the worldwide technical coordination body for all agencies that develop and operate satellites that observe the earth from space.

http://www.nnic.noaa.gov/GOIN/GOIN.html

Tromsø And COrnell Moving Agents (TACOMA)

University of Tromsø, Norway; Cornell University, USA; University of California, San Diego, USA

Dag Johansen, Faculty of Science, University of Tromsø, Norway, dag@cs.uit.no

The TACOMA project focuses on operating system support for agents, and how agents can be used to solve problems traditionally addressed by other distributed computing paradigms; e.g., the client/server model. A series of TACOMA distributed systems have been completed where agents can be moved about in the Internet.

An agent in TACOMA is a piece of code that can be installed and executed on a remote computer. Such an agent may explicitly migrate to other hosts in the network during execution. We are currently focusing on fault-tolerance, security, applicability and management issues. The TACOMA platform has also been ported to new operating system architectures, in particular Windows NT, Windows CE and the PalmOS.

Several TACOMA applications are under construction. One example is a wide-area network weather monitoring system accessible over the Internet. This distributed application is StormCast. We are also investigating whether agents can be useful in extensible file system architectures.

http://www.tacoma.cs.uit.no/

Scalable High-performance Really Inexpensive Multi-Processor (SHRIMP)

University of Tromsø, Norway; Princeton University, USA

The SHRIMP project investigates how to construct high-performance servers with a

network of commodity PCs and commodity operating systems. The cost of a multi-computer server is substantially less than a commercial, custom-designed multi-computer. The goal is to study how to build a system that delivers performance competitive with or better than the commercial multi-computer servers. Research consists of several components: user-level, protected communication, efficient message-passing, shared virtual memory, distributed file system, performance measurement, scalable 3D graphics, and applications.

Princeton's Computer Science department is building a parallel computer using PCs running Linux as the processing elements. The first was a simple two-processor prototype that used a dual-ported RAM on a custom EISA card interface. A recent prototype will scale to larger configurations, using a custom interface card to connect to a "hub" that is essentially the same mesh routing network used in the Intel Paragon <http://www.ssd.intel.com/paragon.html>. Considerable effort has gone into developing low overhead "virtual memory mapped communication" hardware and support software.

http://www.CS.Princeton.EDU/shrimp/

European Incoherent SCATter (EISCAT)

University of Tromsø (Norway); US institutions involved in Upper Atmosphere Facilities and KDI.

C.M. Hall, University of Tromsø, Norway, chris.hall@phys.uit.no

The EISCAT Scientific Association operates radars and receivers in several Nordic cities. Several Incoherent Scatter facilities are distributed about the world, such as Millstone Hill Observatory (MHO), in Westford, Massachusetts. EISCAT studies the interaction between the Sun and the Earth as revealed by disturbances in the magnetosphere and the ionized parts of the atmosphere (these interactions also give rise to the spectacular aurora, or Northern Lights).

The Incoherent Scatter Radar technique requires sophisticated technology and EISCAT engineers are constantly involved in upgrading the systems.

http://www.eiscat.uit.no/

Distributed Virtual Reality

Center for Parallel Computers (PDC), Royal Institute of Technology, Stockholm; NCSA/University of Illinois, USA; University of Illinois at Chicago, USA; University of Utah, USA; University of Houston, USA.

Johan Ihren, Parallel Computing Center, RIT, Stockholm, johani@pdc.kth.se

PDC's Cube is a fully immersive visualization environment that displays images on all surrounding surfaces, including the floor and the ceiling. Virtual environments can be used for everything from science to art, and from industrial simulations to education. Examples of projects at PDC in this area include EnVis, a tool for CFD (Computational Fluid Dynamics) visualization.

PDC has been involved in a number of activities:

- PDC participates in the Globus/GUSTO testbed <www.globus.org>. At SC'97, PDC contributed with a computational electromagnetic application running across a number of sites.
- At the NCSA/Alliance'98 conference, PDC participated in a global VR demonstration. Together with 4 US sites (University of Utah, University of Houston, University of Illinois at Chicago, and University of Illinois at Urbana/Champaign) a collaborative application designed by Caterpillar was demonstrated.
- PDC is an international affiliate partner to the US National Partnership for Advanced Computing Infrastructure (NPACI). General areas of collaboration are grid computing, advanced scientific visualization, mass storage and computational chemistry.

- As an experienced IBM SP site, PDC has agreed with University of Houston to take care of IBM SP related systems management and training at UH. This is partly an experiment in remote (transatlantic) systems work. The IBM SP with mass storage system at UH is working as a data cache within NPACI. An activity similar activity will take place at PDC.

PDC will participate in the following, future activities:

- PDC works as data cache for the Swedish Space Corporation. Data from satellite-based experiments are transferred to PDC from the down-link point and made globally accessible. An example is the ODIN experiment with researchers from Sweden, Finland, France and Canada accessing the data.
- Within NPACI, Mark Ellisman is leading a project on federating brain data. PDC is a partner in similar Swedish and European projects. Planning is underway on how to integrate results of these projects, a task that will put very high demands on the transatlantic link.
- Within the NPACI project, the Telescience ALPHA project would like to do some experiments with digital video between the Karolinska Institute, Stockholm, and UCSD, and possibly some other US sites. The use of IPv6 is high on the list of features for this application.
- Plans are being made for a global data analysis grid for the forthcoming experiments at the Large Hadron Collider at CERN. Full-scale experiments will start in 2005-2006, producing petabytes of data. Testbeds are planned to start in 2000-2001 and PDC plans to participate on behalf of the Swedish particle physicists. Coordination with the US HEP project is also planned. Testbeds will put high demand on the transatlantic links.

7.D. RENATER

Atacama Large Millimeter Array (ALMA)

USA, Chile, France, Germany, The Netherlands, UK and Europe*: The US side of the project is run by the National Radio Astronomy Observatory (NRAO), operated by Associated Universities, Inc., under cooperative agreement with the National Science Foundation (NSF). The European side of the project is a collaboration among the European Southern Observatory (ESO), the Centre National de la Recherche Scientifique, the Max-Planck-Gesellschaft, The Netherlands Foundation for Research in Astronomy and Nederlandse Onderzoekschool Voor Astronomie, and the United Kingdom Particle Physics and Astronomy Research Council.*

The Atacama Large Millimeter Array (ALMA) is a millimeter wavelength telescope – the world's most sensitive, highest resolution, millimeter-wavelength telescope (64 12-meter antennas) in Llano de Chajnantor, Chile. ALMA will bring to millimeter and sub-millimeter astronomy the aperture synthesis techniques of radio astronomy that enable precision imaging to be done on sub-arcsecond angular scales. The richness of the celestial sky at millimeter wavelengths is provided by thermal emission from cool gas, dust, and solid bodies, the same material that shines brightly at far infrared wavelengths. Presently, such natural cosmic emission can be studied only from space with the coarse angular resolution and limited sensitivity that small orbiting telescopes provide.

ALMA will image at 1 mm wavelength with the same 0.01" resolution that will be achieved by the Next Generation Space Telescope. It will provide scientific insight at longer wavelengths that is complementary to that of the VLT and will do so with the same image detail and clarity. In addition, the reconfigurability of ALMA antennas gives ALMA a zoom-lens capability so that it can also make high-fidelity images of large regions of the sky. ALMA is astronomy's complete imaging instrument.

www.alma.nrao.edu

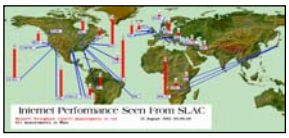

Bandwidth Challenge from the Low-Lands (iGrid 2002)

Europe: CERN, Switzerland; Daresbury Laboratory, Manchester University, Rutherford Appleton Laboratory-Oxford, University College London, UK; Institut National de Physique Nucléaire et de Physique des Particules (IN2P3), France; Istituto Nazionale di Fisica Nucleare (INFN), Milan, Italy; INFN, Rome, Italy; NIKHEF, The Netherlands

Japan: APAN, KEK High Energy Accelerator Research Organization, The Institute of Physical and Chemical Research (RIKEN)/ RIKEN Accelerator Research Facility

North America: TRIUMF, Canada; Argonne National Laboratory, Brookhaven National Laboratory, Caltech, ESnet, Fermi National Accelerator Laboratory, Internet2, Thomas Jefferson National Accelerator Facility (JLab), Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, LBNL National Energy Research Scientific Computing Center (NERSC), NASA Goddard Space Flight Center, Oak Ridge National Laboratory, Rice University, San Diego Supercomputer Center, University of California, San Diego, Stanford Linear Accelerator Center, Stanford University, Sun Microsystems University of Delaware, University of Florida, University of Michigan, University of Texas at Dallas, University of Wisconsin-Madison

Antony Antony, Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF), The Netherlands, antony@nikhef.nl
R. Les Cottrell, Stanford Linear Accelerator Center (SLAC), USA, cottrell@slac.stanford.edu

The avalanche of data already being generated by and for new and future High Energy and Nuclear Physics (HENP) experiments demands new strategies for how the data is collected, shared, analyzed and presented. For example, the SLAC BaBar experiment and JLab are each already collecting over a TB/day, and BaBar expects to increase by a factor of two in the coming year. SLAC and Fermilab's CDF (Collider Detector at Fermilab) and D0 experiments have already gathered well over a petabyte of data, and the CERN Large Hadron Collider (LHC) experiments expect to collect over 10-million terabytes. The strategies being adopted to analyze and store this unprecedented amount of data is the coordinated deployment of Grid technologies, such as those being developed for the Particle Physics Data Grid (PPDG) and the Grid Physics Network (GriPhyN). It is anticipated that these technologies will be deployed at hundreds of institutes that will be able to search out and analyze information from an interconnected worldwide grid of tens of thousands of computers and storage devices. This, in turn, will require the ability to sustain, over long periods, the transfer of large amounts of data among collaborating sites with relatively low latency.

This project demonstrates the current data-transfer capabilities to several sites worldwide that have high-performance links. The iGrid 2002 site acted as a HENP Tier 0 or Tier 1 site (an accelerator or major computation site) in distributing copies of raw data to multiple replica sites. The demonstration is over real live production networks with no efforts to manually limit other traffic. The results are displayed in real time. Researchers investigate/demonstrate issues regarding TCP implementations for high-bandwidth long-latency links, and create a repository of trace files of a few interesting flows. These traces, valuable to projects like EU DataTAG, help explain the behavior of transport protocols over various production networks.

http://www-iepm.slac.stanford.edu/monitoring/bulk/igrid2002

Kites Flying In and Out of Space (iGrid 2002)

University of Illinois at Chicago, Virginia Tech, Virginia Tech Foundation, Virginia Polytechnic Institute & State University, USA; SARA, The Netherlands; Sorbonne and La Cite Museum de Musique Paris, France; Virtual Reality Development and Research Laboratory, Tohwa University, Japan; Institute of High Performance Computing,

Singapore; New Media Innovation Center, Vancouver, British Columbia, Canada

Tom Coffin, National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign (UIUC), tcoffin@ncsa.uiuc.edu

This virtual-reality art piece is a replication and study of the physical properties of the flying kinetic artwork of Jacqueline Matisse-Monnier. The complexity involved with calculating and rendering data is facilitated by distributed computing over high-speed networks.

Because the calculations for these kinetic art pieces (kites) are so computationally intensive, a single PC can only support the simulation of one kite. To support the many kites flown at iGrid 2002, collaborators with computing resources around the world performed the physically-based kite simulations at their home institutions and then streamed the results of the calculations, in real time, to Amsterdam. In essence, this is grid computing for arts.

Note: This application used a Grid model for real-time steering of calculations on computers distributed worldwide over high-speed networks. Each of the 12 kites appearing in the piece utilized up to 15Mbps. The application used distributed servers in Chicago, Canada, Japan, Singapore and Virginia to compute each kite's physical dynamic properties. Each server sent a single kite's motions to iGrid where it was visualized in a CAVE. The application was scalable both computationally and geographically. It proved a good test of high-speed networking because the application required a multicast-enabled network to accomplish communications. As a side benefit, the kites become a visual metaphor for network performance as the kite motions (e.g., fast, slow) responded to latency of the network data.

http://calder.ncsa.uiuc.edu/ART/MATISSE

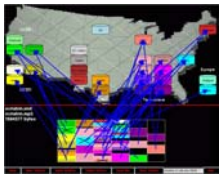

Video IBPster (iGrid 2002)

University of Tennessee, University of California Santa Barbara, University of California San Diego, USA; ENS, Lyon, France; Universita del Piemonte Orientale, Alessandria, Italy; High Performance Computing Center, Rechenzentrum Universitat Stuttgart, Germany

Alessandro Bassi, Logistical Computing and Internetworking (LoCI) Laboratory, University of Tennessee, abassi@cs.utk.edu

Logistical Networking is the global scheduling and optimization of data movement, storage and computation. At LoCI, scientists develop tools such as Data Mover for fast data transfer using as much bandwidth as is available. At iGrid 2002, a geographically distributed abstraction of a file was replicated, transported to depots that were closer according to network proximity values calculated in real time using the Network Weather Service (NWS), and downloaded from the nearest site in a completely transparent way for a high-level application.

http://loci.cs.utk.edu
http://nws.cs.ucsb.edu

BABAR

DAPNIA/SPP, Saclay, France; Stanford Linear Accelerator Center (SLAC), USA; CERN.

Roy Aleksan, DAPNIA/SPP, Saclay, France, aleksan@hep.saclay.cea.fr, aleksan@dapnia.cea.fr, roy@slac.stanford.edu

The BaBar detector was built at SLAC to study the millions of B mesons produced by the PEP-II storage ring. The BaBar collaboration consists of around 600 physicists and engineers from 85 institutions in 9 countries.

http://www.slac.Stanford.edu/BFROOT, http://www-dapnia.cea.fr,

http://www.SLAC.Stanford.edu

The DØ Experiment

Worldwide collaborations, including CERN, based at Fermilab, Illinois. French institutions include: DAPNIA/SPP, Saclay; Centre de Physique des Particules de Marseille; Institut des Sciences Nucleaires de Grenoble; LPNHE, Universités Paris VI and VII; and Laboratoire de L'Accelerateur Lineaire

Armand Zylberstejn,FNAL, USA, azylber@fnald0.fnal.gov, azylber@hep.saclay.cea.fr

The DØ Experiment is a worldwide collaboration of scientists conducting research on the fundamental nature of matter. The experiment is located at the world's premier high-energy accelerator, the Tevatron Collider, at Fermilab.

Worldwide collaborations include Fermilab, Brookhaven National Lab; CERN; Cornell University; DESY, Germany; KEK, Japan; Lawrence Berkeley Lab and Stanford Linear Accelerator Center; several French institutions (DAPNIA/SPP, Centre de Physique des Particules de Marseille, Institut des Sciences Nucleaires de Grenoble, LPNHE, Universites Paris VI and VII, and Laboratoire de L'Accelerateur Lineaire.

http://www-d0.fnal.gov, http://www-dapnia.cea.fr

iMAGIS: Models, Algorithms, Geometry for Graphics and Image Synthesis

INRIA, France; MIT, USA.

Claude Puech, INRIA, France, Claude.Puech@inria.fr

This project develops new visualization techniques to enable the interactive manipulation of urban data. To achieve this goal, efficient image caching and interpolation techniques are combined with traditional 3D techniques. This is important for applications such as project review, civil and military simulators, virtual tourism, education, and climate/environmental studies.

http://www.inria.fr/Equipes/IMAGIS-fra.html, http://www.inria.fr/Unites/RHONE-fra.html, http://graphics.lcs.mit.edu/

Image/Video Transmission, Storage and Manipulation of 3D Images

LIMSI-CNRS, Orsay, France; Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, USA.

3D interactive modeling, with real-time constraints, in a strongly reactive context, constitutes the main collaboration between these two laboratories.

http://www.renater.fr/International/STARTAP_Peerings/Projets/LIMSI_1.htm, http://www.LIMSI.fr

Network Computing

LIP, Ecole Normale Superieure de Lyon, France; Région Rhône-Alpes, France; INRIA, France; CNRS, France; Innovative Computing Laboratory, Univ. of Tennessee, USA.

This research involves setting up metacluster computing with SCILab and NetSolve, organizing videoconferencing for project meetings, developing and testing a distributed web cache, and installing IBP. The LHPC project is a common laboratory for studying parallel computers; it was created by LIP and Matra Systèmes & Information. The first parallel computer of LHPC was installed in March 1995. This machine has evolved and is now made up of around 100 processors.

http://www.ens-lyon.fr/LIP/

Accélération de Particules par Ondes de Choc

Institut Astrophysique de Paris, France; North Carolina State University, USA.

Development of hydrodynamic models for the acceleration of particles in Supernovae remains.

http://www.IAP.fr

Catalogue des Étoiles au Beryllium

Institut Astrophysique de Paris, France; Limber Observatory, Texas, USA.

Catalogue of Beryllium stars.

http://www.IAP.fr

Détection des Étoiles Doubles par Haute Résolution Angulaire

Institut Astrophysique de Paris, France; Center for High Angular Resolution Astronomy, Georgia State University, USA.

Detection of binary stars.

Etude d'Amas de Galaxies en Ultraviolet

Institut Astrophysique de Paris, France; University of Alabama, USA.

Analysis of observations of galaxy clusters with the EUVE satellite.

http://www.IAP.fr

Etude de l'Amas ZWICKY 3146

Institut Astrophysique de Paris, France; CFA, Cambridge, Massachusetts, USA.

Studies of properties of clusters observed with X telescopes.

http://www.IAP.fr

Etude des Disques de Gaz

Institut Astrophysique de Paris, France; Yerkes Observatory, Univ. of Chicago, USA.

Study of gaz disks around young planetary systems such as Beta Pictoris.

http://www.IAP.fr

Etude des Effets de la Poussière dans les Galaxies

Institut Astrophysique de Paris, France; Space Telescope Science Institute (STSCI), Maryland, USA.

Theoretical prediction of observations with new telescopes in the infrared.

http://www.IAP.fr

Etude Théorique et Expérimentale des Profils de Raies

Institut Astrophysique de Paris, France; University of Louisville, Kentucky, USA.

Theoretical computation of ray profiles and comparison with experimental spectra.

http://www.IAP.fr

Fluctuations de Brillance de Surface

Institut Astrophysique de Paris, France; University of California at Berkeley, USA.

Theoretical predictions and observations of signal fluctuations in infrared imaging of galaxies.

http://www.IAP.fr

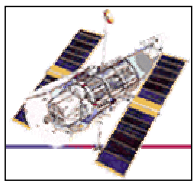

Hubble Space Telescope: Recherches de Raies d'Absorption

Institut Astrophysique de Paris, France; Space Telescope Science Institute (STSCI), Maryland, USA.

Search for absorption rays in spectra from the Hubble Space Telescope, in the framework of the "Absorption Line Systems in Quasars" program.

http://www.stsci.edu/

Large Scale Structure and Cluster Formation

Institut Astrophysique de Paris, France; CFA Cambridge, MA, USA; NASA, USA.

A long-term NASA project of combined analysis in optical, X-ray and radio.

Local Interstellar Cloud

Institut Astrophysique de Paris, France; LPL-WEST, Tucson, Arizona, USA.

The study of local interstellar clouds.

Modèles de Formation d'Etoiles avec Vent Galactique

Institut Astrophysique de Paris, France; University of Illinois at Urbana-Champaign, Dept. of Astronomy, USA.

Models for formation and evolution of galaxies with loss of mass. Consequences for the chemical evolution of elements D, He, O, Fe and others.

Nucléosynthèse du Big Bang dans le Modèle Cosmologique Standard

Institut Astrophysique de Paris, France; University of Minnesota Theoretical Physics Institute School of Physics and Astronomy, USA.

Development of a model for the primordial nucleosynthesis. Study of cosmologic parameters and comparison with observational constraints.

Origine et Evolution du Lithium, du Beryllium et du Bore dans la Galaxie

Institut Astrophysique de Paris, France; University of Chicago, USA.

The observation of Lithium, Beryllium and Bore elements in halo stars and the development of astrophysics models that detail the chemical evolution of these elements in the galaxy.

Photochimie des Enveloppes Circumstellaires

Institut Astrophysique de Paris, France; Physics Dept., New York University, USA.

Theoretical computation dealing with distribution of gaseous components of circumstellar shells.

FUSE (Far Ultraviolet Spectroscopic Explorer)

NASA, USA; The Johns Hopkins University, USA; The University of Colorado at Boulder, USA; The University of California at Berkeley, USA; Canadian Space Agency (CSA), Canada; French Space Agency (Centre National d'Etudes Spatiales, or CNES), France; Institut Astrophysique de Paris (IAP), France.

FUSE is a NASA-supported astronomy mission that was launched on June 24, 1999 to explore the universe using the technique of high-resolution spectroscopy in the far-ultraviolet spectral region.

http://www.iap.fr/ProgrammesCollaboration/Fuse/, http://fuse.pha.jhu.edu/

Serveur d'Éphémérides MPC/CBAT de l'U.A.I.

Institut Astrophysique de Paris, France; Smithsonian Astrophysical Observatory, Cambridge, Massachussetts, USA.

Using a server for computation of asteroids and comets ephemeredes.

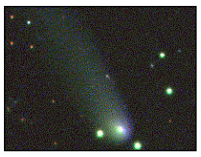

http://cfa-www.harvard.edu/newtop/saohome.html

SLOAN Digital Sky Survey (SDSS)

Institut Astrophysique de Paris, France; Johns Hopkins University, Dept. of Astronomy,

USA; CFHT Corp. (Hawaii); University of Hawaii Institute for Astronomy, Honolulu.

The SDSS enables the automatic, systematic study and exchange of data of stars, galaxies, nebula, and large-scale structure.

http://www.sdss.org/sdss.html

Surveys Radio

Institut Astrophysique de Paris, France; National Radio Astronomy Organization (NRAO), Virginia, USA.

Properties of deep radio surveys.

TERAPIX (Traitement Elementaire Reduction et Analyse des PIXels)

Institut Astrophysique de Paris, France; CFHT Corp., Kamelua, Hawaii, USA; University of Hawaii Institute for Astronomy, Honolulu, USA.

Data analysis of images form MEGAGAMMA/PRIME. TERAPIX is an astronomical data processing center at the Institut d' Astrophysique de Paris dedicated to very large CCD images and massive data flow provided by the MEGACAM camera. TERAPIX organizes the MEGACAM image processing and provides images and catalogues to the Canada-France-Hawaii-Telescope (CFHT) user community.

http://terapix.iap.fr

CASSINI-HUYGENS Cluster 2 WBD

CETP (IPSL,Vélizy), France; Dept. of Physics and Astronomy, University of Iowa, USA.

Transfer and real-time visualization of data from the Radio and Plasma Wave Science on board the spacecraft.

http://despa.obspm.fr/plasma/cluster/cluster.html, http://despa.obspm.fr/

CASSINI-HUYGENS Cluster 2 WEC

CETP (IPSL, Vélizy), France; University of California at Berkeley Space Science Lab, USA; Dept. of Physics and Astronomy, University of Iowa, USA.

Technical coordination of the WEC consortium of the CLUSTER-2 project.

http://despa.obspm.fr/plasma/cluster/cluster.html, http://despa.obspm.fr/, http://www.CETP.IPSL.fr

GALILEO

CETP (IPSL, Vélizy), France; Dept. of Physics and Astronomy, University of Iowa, USA.

Study of the environment of Jupiter.

http://www.CETP.IPSL.fr

GALILEO / NIMS

DESPA (Observatoire de Paris Meudon), France; NASA Jet Propulsion Laboratory, USA.

Infrared spectro-imagery of Jupiter with the GALILEO spacecraft.

http://despa.obspm.fr, http://www.jpl.nasa.gov/

CASSINI-HUYGENS DISR

DESPA (Observatoire de Paris Meudon), France; University of Arizona Lunar Planetary Lab, USA.

Analysis of cloud particles from TITAN: Electronic maps set up at DESPA

http://despa.obspm.fr/planeto/disr.html, http://www.jpl.nasa.gov/cassini/, http://despa.obspm.fr/planeto/cassini.html

High Energy Solar Spectroscopic Imager (HESSI)

DESPA (Observatoire de Paris Meudon), France; University of California at Berkeley/Space Science Laboratory, USA

Robert Lin, University of California, Berkeley, boblin@ssl.berkeley.edu
Nicole Vilmer, Observatoire de Paris-Meudon, vilmer@obspm.fr

The HESSI mission consists of a single spin-stabilized spacecraft in a low-altitude orbit inclined 38 degrees to the Earth's equator. The only instrument on board is an imaging spectrometer with the ability to obtain high fidelity color movies of solar flares in X rays and gamma rays. It uses two new complementary technologies: fine grids to modulate the solar radiation, and germanium detectors to measure the energy of each photon very precisely.

http://hesperia.gsfc.nasa.gov/hessi/, http://www.obspm.fr

IRON

DESPA (Observatoire de Paris Meudon), France; Dept. of Astronomy, Ohio State University, USA.

Computation of collision and radiative atomic data for astrophysics.

http://www.obspm.fr

VIMS

DESPA (Observatoire de Paris Meudon), France; Lunar and Planetary Laboratory, Dept. of Planetary Science and Astronomy, University of Arizona, USA.

Infrared spectro imagery on the Cassini spacecraft: observations of Jupiter, Saturn, Titan between 2000 and 2006; transmission of camera images.

http://www.obspm.fr

ASTRAIA

CETP (IPSL, Vélizy), France; NCAR, USA.

Doppler meteorological airborne radar with two beams.

http://www.CETP.IPSL.fr

Capteur Spatial SEAWIFS: Ameliorations

Laboratoire des Sciences du Climat et de l'Environnement (LSCE), CEA Saclay, France; Rosenstiel School of Marine and Atmospheric Science (RSMAS), Univ. of Miami, USA.

Enhance atmospheric corrections for the space measurement system SEAWIFS for sea color.

CLIMSERV-CDC

(Note: il ne s'agit pas d'un projet, mais d'une activité de service effectuée par l'IPSL)

LMD (IPSL, Ecole Polytechnqiue-Palaiseau), France; Climate Diagnostic Center (CIRES), U of Colorado, USA

Data transfers from the NOAA Climate Diagnostic Center database to the French database CLIMSERV.

http://perceval.polytechnique.fr, http://www.lmd.jussieu.fr

CLIMSERV-EOSDIS / LARC

(Note: il ne s'agit pas d'un projet, mais d'une activité de service effectuée par l'IPSL)

LMD (IPSL, Ecole Polytechnqiue-Palaiseau), France; EOSDIS/LARC, NASA Langley Atmospheric Sciences Data Center, Virginia, USA.

Data transfers from the EOSDIS database to the French database CLIMSERV.

http://perceval.polytechnique.fr, http://www.lmd.jussieu.fr

International Field Experiment in the Indian Ocean (INDOEX)

LMD (IPSL, Ecole Polytechnqiue-Palaiseau), France; Center for Clouds, Chemistry and Climate, Scripps Institute for Oceanography, University of California, San Diego, USA

V. Ramanathan, Scripps Institution of Oceanography of the University of California, San Diego, vramanathan@ucsd.edu

INDOEX addresses questions of climate change that are of high priority and great value to the US and the international community. The project's goal is to study natural and anthropogenic climate forcing by aerosols and feedbacks on regional and global climate.

http://www-indoex.ucsd.edu/, http://www.lmd.ens.fr/INDOEX/

7.E. SURFnet

Atacama Large Millimeter Array (ALMA)

USA, Chile, France, Germany, The Netherlands, UK and Europe*: The US side of the project is run by the National Radio Astronomy Observatory (NRAO), operated by Associated Universities, Inc., under cooperative agreement with the National Science Foundation (NSF). The European side of the project is a collaboration among the European Southern Observatory (ESO), the Centre National de la Recherche Scientifique, the Max-Planck-Gesellschaft, The Netherlands Foundation for Research in Astronomy and Nederlandse Onderzoekschool Voor Astronomie, and the United Kingdom Particle Physics and Astronomy Research Council.*

The Atacama Large Millimeter Array (ALMA) is a millimeter wavelength telescope – the world's most sensitive, highest resolution, millimeter-wavelength telescope (64 12-meter antennas) in Llano de Chajnantor, Chile. ALMA will bring to millimeter and sub-millimeter astronomy the aperture synthesis techniques of radio astronomy that enable precision imaging to be done on sub-arcsecond angular scales. The richness of the celestial sky at millimeter wavelengths is provided by thermal emission from cool gas, dust, and solid bodies, the same material that shines brightly at far infrared wavelengths. Presently, such natural cosmic emission can be studied only from space with the coarse angular resolution and limited sensitivity that small orbiting telescopes provide.

ALMA will image at 1 mm wavelength with the same 0.01" resolution that will be

achieved by the Next Generation Space Telescope. It will provide scientific insight at longer wavelengths that is complementary to that of the VLT and will do so with the same image detail and clarity. In addition, the reconfigurability of ALMA antennas gives ALMA a zoom-lens capability so that it can also make high-fidelity images of large regions of the sky. ALMA is astronomy's complete imaging instrument.

www.alma.nrao.edu

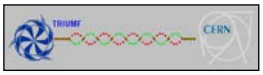

ATLAS Canada LightPath Data Transfer Trial (iGrid 2002)

TRIUMF, Carleton University, University of Victoria, University of Alberta, University of Toronto, Simon Fraser University, BCNet, CANARIE, Canada; CERN, Switzerland; Universiteit van Amsterdam, The Netherlands

Corrie Kost, TRIUMF, Canada, kost@triumf.ca

The Lightpath Trial is attempting to transmit 1 TeraByte (TB) of ATLAS data from TRIUMF, Canada's National Laboratory for Particle and Nuclear Physics, to CERN at high speed. ATLAS (A Toroidal LHC ApparatuS) is one of the experiments being constructed for the Large Hadron Collider (LHC) at CERN.

Using SURFnet's entire 2.5Gbps capacity between StarLight and NetherLight, as well as the planned 2.5Gbps links from Canada to StarLight, CERN to StarLight (EU DataTAG) and NetherLight to CERN, an end-to-end lightpath is being built between TRIUMF in Vancouver and CERN. The goal is to transfer a 1TB of Monte Carlo data between a cluster at TRIUMF and a cluster at CERN in under two hours.

*Note: Subsequently, a Terabyte of research data was transferred in less than a DVD/minute over a "light path" extending 12,000 kilometers from TRIUMF in Vancouver to CERN. The demonstration required dedicated portions of fiber-optic networks, spanning one provincial (British Columbia's BCnet) and two national research and education networks (CA*net4 and SURFnet) to establish the on-demand private network. Experiments were also conducted on the Chicago to CERN DataTAG link. The project culminated in establishing the first large-scale end-to-end "light path".* ***Researchers were able to transfer a Terabyte of research data from disk-to-disk in under 3 hours – this is equivalent to transferring a full CD in less than 8 seconds (or a full-length DVD movie in less than 1 minute).*** *Using* bbftp*, they could transfer 60GB files in 10 minutes; 1TB is about 17 of these files, or about 170 minutes transfer time. Using Tsunami, they achieved slightly better rates. Peak transfer rates in excess of 1Gb were achieved, twice the previous known record for this distance. This is the first establishment of an inter-domain end-to-end "light path" dedicated for a research application. The "light path" directly connecting TRIUMF and CERN is the longest known single hop network spanning the distance from Vancouver to Geneva via StarLight, NetherLight and DataTAG. (Note: Performance was affected by using 10GigE alpha-test cards from Intel. Subsequent testing with beta-test cards has shown that they perform roughly 4 times better.)*

www.triumf.ca

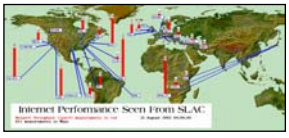

Bandwidth Challenge from the Low-Lands (iGrid 2002)

Europe: CERN, Switzerland; Daresbury Laboratory, Manchester University, Rutherford Appleton Laboratory-Oxford, University College London, UK; Institut National de Physique Nucléaire et de Physique des Particules (IN2P3), France; Istituto Nazionale di Fisica Nucleare (INFN), Milan, Italy; INFN, Rome, Italy; NIKHEF, The Netherlands

Japan: APAN, KEK High Energy Accelerator Research Organization, The Institute of Physical and Chemical Research (RIKEN)/ RIKEN Accelerator Research Facility

North America: TRIUMF, Canada; Argonne National Laboratory, Brookhaven National Laboratory, Caltech, ESnet, Fermi National Accelerator Laboratory, Internet2, Thomas

Jefferson National Accelerator Facility (JLab), Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, LBNL National Energy Research Scientific Computing Center (NERSC), NASA Goddard Space Flight Center, Oak Ridge National Laboratory, Rice University, San Diego Supercomputer Center, University of California, San Diego, Stanford Linear Accelerator Center, Stanford University, Sun Microsystems University of Delaware, University of Florida, University of Michigan, University of Texas at Dallas, University of Wisconsin-Madison

Antony Antony, Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF), The Netherlands, antony@nikhef.nl
R. Les Cottrell, Stanford Linear Accelerator Center (SLAC), USA, cottrell@slac.stanford.edu

The avalanche of data already being generated by and for new and future High Energy and Nuclear Physics (HENP) experiments demands new strategies for how the data is collected, shared, analyzed and presented. For example, the SLAC BaBar experiment and JLab are each already collecting over a TB/day, and BaBar expects to increase by a factor of two in the coming year. SLAC and Fermilab's CDF (Collider Detector at Fermilab) and D0 experiments have already gathered well over a petabyte of data, and the CERN Large Hadron Collider (LHC) experiments expect to collect over 10-million terabytes. The strategies being adopted to analyze and store this unprecedented amount of data is the coordinated deployment of Grid technologies, such as those being developed for the Particle Physics Data Grid (PPDG) and the Grid Physics Network (GriPhyN). It is anticipated that these technologies will be deployed at hundreds of institutes that will be able to search out and analyze information from an interconnected worldwide grid of tens of thousands of computers and storage devices. This, in turn, will require the ability to sustain, over long periods, the transfer of large amounts of data among collaborating sites with relatively low latency.

This project demonstrates the current data-transfer capabilities to several sites worldwide that have high-performance links. The iGrid 2002 site acted as a HENP Tier 0 or Tier 1 site (an accelerator or major computation site) in distributing copies of raw data to multiple replica sites. The demonstration is over real live production networks with no efforts to manually limit other traffic. The results are displayed in real time. Researchers investigate/demonstrate issues regarding TCP implementations for high-bandwidth long-latency links, and create a repository of trace files of a few interesting flows. These traces, valuable to projects like EU DataTAG, help explain the behavior of transport protocols over various production networks.

www-iepm.slac.stanford.edu/monitoring/bulk/igrid2002

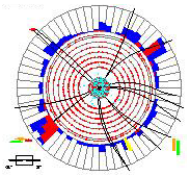

D0 Data Analysis (iGrid 2002)

Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF), The Netherlands; Fermi National Accelerator Laboratory (Fermilab), Michigan State University, USA

Wim Heubers, Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF), The Netherlands, wimh@nikhef.nl

The D0 Experiment, which relies on the Tevatron Collider at Fermilab, is a worldwide collaboration of scientists conducting research on the fundamental nature of matter. The research focuses on precise studies of interactions of protons and antiprotons at the highest available energies as part of an intense search for subatomic clues that reveal the character of the building blocks of the universe.

Currently, raw data from the D0 detector is processed Fermilab's computer farm and results are written to tape. At iGrid 2002, researchers showed that by using the transoceanic StarLight/NetherLight network, it is possible for Fermilab to send raw data

to NIKHEF for processing and then have NIKHEF send the results back to Fermilab.

www-d0.fnal.gov
www.nikhef.nl

Griz: Grid Visualization Over Optical Networks (iGrid 2002)

Vrije University, The Netherlands; University of Illinois at Chicago, USA

Henri Bal, Vrije Universiteit, bal@cs.vu.nl

Aura, a distributed parallel rendering toolkit, was used to remotely render data on available graphics resources in Chicago and in Amsterdam for local display at the iGrid 2002 conference. Aura is applied to real world scientific problems; notably, the visualization of high-resolution isosurfaces of the Visible Human dataset and an interactive molecular dynamics simulation.

Interactive and collaborative applications have a near-real-time requirement. For interaction over long distances, network delay is a key factor. Optical networks seem to have a predictable latency, making near-real-time interactive behavior easier, and the higher bandwidth allows faster access to large datasets and remote visualization machines.

www.cs.vu.nl/~renambot/vr/html/intro.htm

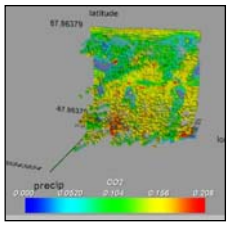

High Performance Data Webs (iGrid 2002)

University of Illinois at Chicago, USA; Dalhousie University, Canada; Imperial College of Science, Technology & Medicine, University of London, UK, Universiteit van Amsterdam, SARA, The Netherlands; Center for Parallel Computers, Royal Institute of Technology, Sweden

Robert Grossman, Laboratory for Advanced Computing (LAC), University of Illinois at Chicago, grossman@uic.edu

DataSpace is a high-performance data web for the remote analysis, mining, and real-time interaction of scientific, engineering, business and other complex data. DataSpace applications are designed to exploit the capabilities provided by emerging domestic and international high-performance networks so that gigabyte and terabyte datasets can be remotely explored in real time. It is an open, standards-based infrastructure that combines data web services, data grid services and semantic web services for remote data analysis and distributed data mining.

The Terra Wide Data Mining (TWDM) testbed consists of high-performance clusters worldwide linked by wide-area advanced networks, providing the data and computer services required.

www.ncdm.uic.edu
www.dataspaceweb.net

Kites Flying In and Out of Space (iGrid 2002)

University of Illinois at Chicago, Virginia Tech, Virginia Tech Foundation, Virginia Polytechnic Institute & State University, USA; SARA, The Netherlands; Sorbonne and La Cite Museum de Musique Paris, France; Virtual Reality Development and Research Laboratory, Tohwa University, Japan; Institute of High Performance Computing, Singapore; New Media Innovation Center, Vancouver, British Columbia, Canada

Tom Coffin, National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign (UIUC), tcoffin@ncsa.uiuc.edu

This virtual-reality art piece is a replication and study of the physical properties of the flying kinetic artwork of Jacqueline Matisse-Monnier. The complexity involved with

calculating and rendering data is facilitated by distributed computing over high-speed networks.

Because the calculations for these kinetic art pieces (kites) are so computationally intensive, a single PC can only support the simulation of one kite. To support the many kites flown at iGrid 2002, collaborators with computing resources around the world performed the physically-based kite simulations at their home institutions and then streamed the results of the calculations, in real time, to Amsterdam. In essence, this is grid computing for arts.

Note: This application used a Grid model for real-time steering of calculations on computers distributed worldwide over high-speed networks. Each of the 12 kites appearing in the piece utilized up to 15Mbps. The application used distributed servers in Chicago, Canada, Japan, Singapore and Virginia to compute each kite's physical dynamic properties. Each server sent a single kite's motions to iGrid where it was visualized in a CAVE. The application was scalable both computationally and geographically. It proved a good test of high-speed networking because the application required a multicast-enabled network to accomplish communications. As a side benefit, the kites become a visual metaphor for network performance as the kite motions (e.g., fast, slow) responded to latency of the network data.

http://calder.ncsa.uiuc.edu/ART/MATISSE

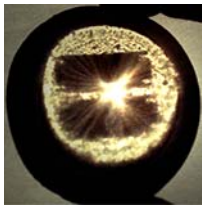

Photonic TeraStream (iGrid 2002)

Northwestern University, University of Illinois at Chicago, Argonne National Laboratory, USA; Universiteit van Amsterdam, The Netherlands

Joe Mambretti, International Center for Advanced Internet Research (iCAIR), Northwestern University, j-mambretti@nwu.edu

ICAIR, in partnership with the Materials Sciences Research Center at Northwestern University, is developing an International Virtual Institute (IVI) for Materials Science. The IVI wants to be able to instantaneously discover, gather, integrate, and present information–whether large-scale datasets, scientific visualizations, streaming digital media, results of computational processes–from resources worldwide. To accomplish this, iCAIR is developing "Global Services-on-Demand" technologies for optical networks.

The Photonic TeraStream is supported by OMNInet, the Chicago-area Optical Metro Network Initiative. OMNInet is designed and developed by SBC/Ameritech, Nortel Networks and iCAiR, in collaboration with EVL, CANARIE and ANL. It is an experimental networking testbed, enabling researchers to assess and validate next-generation optical technologies, architectures and applications in metropolitan networks. For iGrid 2002, however, the OMNInet testbed was extended to Amsterdam through StarLight and NetherLight to demonstrate that photonic-enabled applications are possible, not only at the metro level, but also on a global scale (the global LambdaGrid).

Researchers are using OMNInet to prototype tools for intelligent application signaling, dynamic lambda provisioning, and extensions to lightpaths through dynamically provisioned Layer2 and Layer3 configurations, in part, to allow for access to multiple types of edge resources. In turn, these network-control capabilities are being incorporated into next-generation large-scale global applications, which include high-performance data transfer (based on GridFTP), digital media streaming (270Mbps encoding), and high-performance remote data-access methods (based on iSCSI).

At iGrid 2002, iCAIR presented its innovative dynamic lambda provisioning capability–the Optical Dynamic Intelligent Network (ODIN) service layer. Applications use intelligent signaling to provision their own lightpaths with ODIN in order to optimize network-based resource discovery and performance; for example, to access and to

dynamically interact with very large amounts of distributed data. Applications supported by dynamic lambda switching provide for significantly more powerful capabilities than those based on today's communication infrastructure.

www.icair.org/igrid2002
www.uva.nl
www.icair.org/omninet

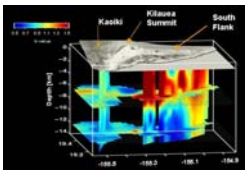

The OptIPuter

University of California San Diego, University of Illinois at Chicago, Northwestern University, San Diego State University, Information Sciences Institute/University of Southern California, University of California-Irvine, Texas A&M University, USGS, USA;University of Amsterdam, The Netherlands.

Larry Smarr, University of California San Diego, lsmarr@ucsd.edu
Tom DeFanti, University of Illinois at Chicago, tom@evl.uic.edu

The OptIPuter is a powerful distributed cyberinfrastructure to support data-intensive scientific research and collaboration. It has two application drivers – the NIH Biomedical Informatics Research Network and the NSF EarthScope – where scientists are generating multi-gigabytes of 3D volumetric data objects that reside on distributed archives that they want to correlate, analyze and visualize. The OptIPuter is being designed as a "virtual" parallel computer in which the individual "processors" are widely distributed clusters; the "memory" is in the form of large distributed data repositories; "peripherals" are very-large scientific instruments, visualization displays and/or sensor arrays; and the "motherboard" uses standard IP delivered over multiple dedicated lambdas.

www.calit2.net/news/2002/9-25-optiputer.html

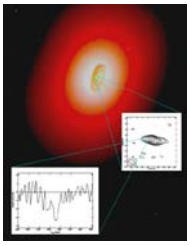

vlbiGrid (iGrid 2002)

Joint Institute for VLBI in Europe (JIVE), The Netherlands; Metsahovi Radio Observatory, Finland; Dept. of Physics and Astronomy at University College London, Jodrell Bank Observatory and Dept. of Physics and Astronomy at University of Manchester, UK; Haystack Observatory, Massachusetts Institute of Technology, USA; Universiteit van Amsterdam, The Netherlands

Steve M. Parsley, Joint Institute for VLBI in Europe (JIVE), The Netherlands, parsley@jive.nl

Very Long Baseline Interferometry (VLBI) is a technique in which an array of physically independent radio telescopes observes simultaneously to yield high-resolution images of cosmic radio sources. Today, magnetic tape transports data from telescopes to data processors. The European VLBI Network (EVN) has access to multiple data sources that can deliver 1Gbps each and a dedicated supercomputer that can process 16 data streams simultaneously. High-speed networks are enabling the EVN to achieve many-fold improvements in bandwidth.

www.jive.nl
www.jb.man.ac.uk
www.haystack.edu

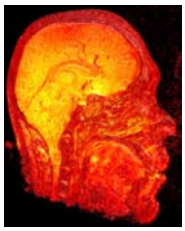

Quanta: Application-Centric Communications Middleware

University of Illinois at Chicago, USA; SARA, The Netherlands

Jason Leigh, University of Illinois at Chicago, spiff@evl.uic.edu

Quanta, targeted for optical networks, is an applications-level API that translates high-level data distribution requirements into low-level optimized networking protocols and parameter settings. Application developers specify their application's data transfer

characteristics, and then Quanta transparently translates these requirements into transmission protocols and network services. Quanta provides a rich set of networking tools and data distribution mechanisms, including: message passing, distributed shared memory, remote procedure calls, remote file I/O, Forward Error Corrected UDP, Parallel TCP for bulk data transfer, and collaborative performance monitoring. The data transfer scheme Reliable Blast UDP (RBUDP) was recently added to Quanta. Quanta is used in the OptIPuter project; illustrated here is a data from the NIH BIRN project and SDSC, and is being tested over long, fat networks between Chicago and Amsterdam.

www.evl.uic.edu/cavern/teranode/quanta.html

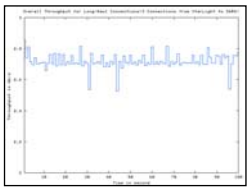

Lambda Joins

University of Illinois at Chicago, USA; CANARIE, Canada; SARA, The Netherlands

Robert Grossman, National Center for Data Mining (NCDM), University of Illinois at Chicago, grossman@uic.edu

Database joins are one of the key technologies that make data processing practical. As more data is distributed over the Internet, the ability to *join* data located in two different global locations is becoming critical. There are two fundamental problems: finding efficient protocols to move data over long distances, and finding efficient algorithms to merge two data streams.

For a demo at SC 2002, Project DataSpace, in collaboration with researchers from Chicago, Ottawa and Amsterdam, a stream of data was moved from a computer cluster at SARA in Amsterdam to a three-node computer cluster at StarLight in Chicago, at over 2.8 Gbps. At the same time, a stream of data was moved from a cluster at CANARIE in Ottawa to the same cluster at StarLight in Chicago, at over 2 Gbps. The two streams of data were then merged using the StarLight cluster at over 500 Mbps, per node.

Both the algorithm for joining the lambda streams and the high-performance data transport streaming protocol used (SABUL) were developed by NCDM/Laboratory for Advanced Computing at the University of Illinois at Chicago.

To many network engineers, lambda and lightpath are used interchangeably to describe a low layer end-to-end dedicated communications channel of effective guaranteed bandwidth. Using protocols such as SABUL, it is now possible to use lambdas to move large datasets over long distances as fast as the data can be pulled from disk. Using lambda joins, it is now possible to merge two such streams and look for patterns in data, even if the data is distributed worldwide.

Project DataSpace won the SC 2002 High Performance Bandwidth Challenge Award for Innovative, High Speed, Data Correlation – Best Use of Emerging Infrastructure.

www.dataspaceweb.net

Project DataSpace and the Terra Mining Testbed

Project DataSpace: University of Illinois at Chicago, USA; University of Pennsylvania, USA; Caltech, USA; *NCAR, USA; University of California Davis, USA; Magnify Research, Inc., USA; Imperial College, England; ACYys, Canberra, Australia*

Terra Mining Testbed: University of Illinois at Chicago, USA; SARA, The Netherlands; Dalhousie University, Canada, Imperial College of Science, Technology and Medicine, England; Virginia Tech and ACCESS DC, USA, Internet2, USA; University of California Davis, USA; University of Pennsylvania, USA.

Robert Grossman, University of Illinois at Chicago, grossman@uic.edu

The web today provides an infrastructure for working with distributed multimedia documents, but not for remotely exploring data. Project DataSpace is an attempt to provide such an infrastructure. It contains protocols for mining distributed data and is

effective for distributed workstation clusters connected with high performance networks (super-clusters) and commodity networks (meta-clusters). The Terabyte Challenge, the testbed for Project DataSpace, will link 12 sites on five continents and demonstrate a variety of applications which will publish, access, analyze, correlate and manipulate remote and distributed data.

The Terra Mining Testbed is an infrastructure built on top of DataSpace for remote analysis, distributed data mining and real-time interaction with large, complex data sets. In a demonstration at SC'02, the DataSpace team accessed, correlated and then visualized data from the National Center for Atmospheric Research (NCAR), National Oceanic and Atmospheric Administration (NOAA) and World Health Organization (WHO)—a CERN linked institution—to study the correlation between El Nino and cholera outbreaks. Terra Mining applications are designed to exploit the capabilities of emerging domestic and international optical networks so that gigabyte and terabyte datasets can be remotely explored in real time.

The demonstration also showcased PC-based clusters called TeraNodes, now being deployed throughout the world, which will be dedicated to massive computation, data mining or visualization over national and international high performance networks. In coming years, as optical technology transforms networking capabilities, TeraNodes will become the building blocks for an optically connected web of data.

Transatlantic Broadband Videostream Using Motion JPEG-over-IP

Univ. of California, Berkeley, USA; Univ. of Tennessee, USA; SURFnet, The Netherlands

Lawrence Rowe, BMRC, University of California at Berkeley, rowe@cs.berkeley.edu
Chris Hodges, SunSITE, University of Tennessee, hodge@ns.utk.edu
Bart Kerver, SURFnet, bart.kerver@surfnet.nl

UC-Berkeley, Univ. of Tennessee and SURFnet have implemented the first live transatlantic broadband audio and video streams using Motion JPEG-over-IP. The implementation of this technology demonstrates how full-screen, near-TV quality video can be delivered over the Internet using simple, inexpensive hardware and freely available, open source software. The technology provides significantly better video quality than those relying on the H.323 standard—the current standard governing interactive audio, video and data communications in a networked environment.

High-quality real-time video, video-on-demand and videoconferencing are critical for applications such as distance education, telemedicine, and remote scientific collaboration. Based on initial success of trials run in December 2001, collaborators will continue to explore ways it can be used to strengthen and expand academic instruction and research.

http://www.bmrc.berkeley.edu/

SC Global

Ian Foster, Argonne National Laboratory, foster@mcs.anl.gov

The annual SC conference is the premiere technical and industrial meeting for high-end networking and computing and computational science. The showcase event at SC'2001 held in Denver, Colorado was SC Global—the first truly global technical conference held on the Grid.

SC Global employed Access Grid technology to link the SC Core at the Denver Convention Center with dozens of SC Constellation Sites throughout the world. The event became a multi-national meeting place for demonstrations and discussions related to high-end computing and communications, and featured real time panels, workshops, and birds-of-a-feather sessions.

Disclaimer: Thirty-six sites participated in SC Global, 28 of which contributed content to the conference. These included sites all across the US and in Beijing, China; Juelich,

Germany; Tsukuba, Japan; Toronto, Canada; Manchester, England; Stuttgart, Germany; Bologna, Italy; Sydney, Australia; Porto Alegre, Brazil; and the National Science Foundation's South Pole Research Center. The SARA supercomputer center in Amsterdam subsequently purchased an Access Grid. While they did not participate in SC'2001, they are now able to participate in other global projects.

http://www.accessgrid.org
http://www-fp.mcs.anl.gov/scglobal/default.htm

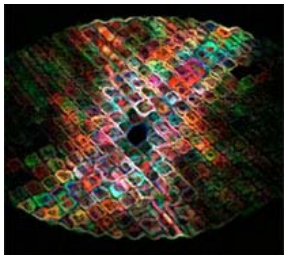

GiDVN: Global Internet Digital Video Network

International Center for Advanced Internet Research (iCAIR), Northwestern, USA; Digital Video Working Group, Coordinating Committee for International Research Networks (DVWG,CCIRN), worldwide membership

Joe Mambretti, iCAIR, Northwestern University, j-mambretti@nwu.edu

GiDVN projects are enhancing media capabilities for the next-generation Internet, enabling new applications to interoperate throughout the world.

Peter Marshall, CANARIE Inc., Canada
Olivier Martin, Paolo Moroni, Philippe Galvez, Joop Joosten, CERN
Kazunori Konishi, APAN, Japan
Shuichi Matsumoto, Masahiro Wada, Shigeyuki Sakazawa, Yasuhiro Takishima, Tetsusi Yamashita, KDD, Japan
Jaehwa Lee, Seungmo Choe, APAN-KR, Korea
Yung Yi, Yanghee Choi, Wang Lijing, Seoul National University, Korea
Sungkwan Youm, Korea University, Korea
Egon M. Verharen, SURFnet, The Netherlands
Cesar Olvera, DGSCA-UNAM, Mexico
Manjeet Singh, Francis Lee, SingAREN, Singapore
Artur Serra, Sebastia Sallent, Joan Borràs, Universitat Politecnica de Catalunya, Spain
Björn Pehrson, Daniel Forsgren, Royal Institute of Technology, Sweden
Joe Mambretti, Jim Chen, Jeremy Weinberger, Tim Ward, Northwestern Univ., USA

http://www.icair.org/inet2000, http://iumrs.ms.nwu.edu

Application-Level Network Performance Analysis Tools

Stichting Academisch Rekencentrum Amsterdam (SARA), The Netherlands; University of Illinois at Chicago, USA.

Jason Leigh, University of Illinois at Chicago, USA spiff@uic.edu
Paul Wielinga, SARA, The Netherlands, wielinga@sara.nl

The UIC Electronic Visualization Laboratory and SARA collaborate on a number of network performance projects. Specifically, in fulfillment of the goals of the Euro-Link project, the two sites have been working various networking QoS schemes to improve transmission rates, notably TCP, UDP, FEC and RUDP.

ALIVE: Architectural Linked Immersive Environment

Stichting Academisch Rekencentrum Amsterdam (SARA), The Netherlands; University of Illinois at Chicago, USA.

Jason Leigh, University of Illinois at Chicago, USA spiff@uic.edu
Ed Breeveld, SARA, The Netherlands, edward@sara.nl
Frans Blok, Office of Metropolitan Architecture, Rotterdam, The Netherlands

ALIVE is used to evaluate the usability of collaborative virtual reality for architectural design. The ALIVE project started February 1999 at SARA in cooperation with EVL and the Office for Metropolitan Architecture. In February 1998, architect Rem Koolhaas won

the Richard H. Driehaus Foundation International Design Competition for the new Campus Center at Illinois Institute of Technology's historic Mies van der Rohe campus. A walkthrough of the plan has been modeled for the CAVE®.

http://www.sara.nl, http://www.archfonds.nl

The MegaConference

Hosted at Ohio State University. Augmented by other organizations as needed. SURFnet is one provider of MCU capacity, and broadcasts The Megaconference via RealVideo and MPEG1

Malik Amer Khan, OARnet, mkhan@oar.net
Bob Dixon, Ohio State University, Bob_Dixon@osu.edu

The Megaconference is a permanent, ongoing multipoint H.323 Internet videoconference. It is free and open to anyone with adequate equipment to provide good quality Internet video at the speed of 384 Kbps. SURFnet has participated in H.323 videoconferences/demos between Ohio State University, University of South Carolina, NYSERnet, Buffalo University, Rochester University and Syracuse University.

http://www.mega-net.net/megaconference

6TAP

The Netherlands; CANARIE (Canada); ESnet (USA).

The 6TAP project provides native and tunneled IPv6 interconnections at STAR TAP to early IPv6 production networks to enable them to build and demonstrate IPv6-based applications. 6TAP will develop: IPv6 route server technology, network tools for network measurement, analysis and display, and experience in supporting, provisioning and operating IPv6 Internet exchange points.

6TAP is co-sponsored by ESnet and CANARIE, with participation from vBNS, Abilene/Internet2, SURFnet, APAN, CERN, SingAREN, four US national labs (ANL, LBNL, ORNL, SLAC), U. Wisconsin, UCSD and Sun Microsystems.

http://www.6tap.net