

The International Grid (iGrid): Empowering Global Research Community Networking Using High Performance International Internet Services

Maxine D. BROWN (maxine@uic.edu), University of Illinois at Chicago, USA
Thomas A. DEFANTI (tom@uic.edu), University of Illinois at Chicago, USA
Michael A. MCROBBIE (vpit@indiana.edu), Indiana University, USA
Alan VERLO (alan@eecs.uic.edu), University of Illinois at Chicago, USA
Dana PLEPYS (dana@eecs.uic.edu), University of Illinois at Chicago, USA
Donald F. MCMULLEN (mcmullen@indiana.edu), Indiana University, USA
Karen ADAMS (kadams@indiana.edu), Indiana University, USA
Jason LEIGH (jleigh@eecs.uic.edu), University of Illinois at Chicago, USA
Andrew E. JOHNSON (ajohnson@eecs.uic.edu), University of Illinois at Chicago, USA
Ian FOSTER (foster@mcs.anl.gov), Argonne National Laboratory, USA
Carl KESSELMAN (carl@isi.edu), University of Southern California, Information Sciences Institute, USA
Andrew SCHMIDT (andrew.g.schmidt@ameritech.com), Ameritech Advanced Data Services
Steven N. GOLDSTEIN (sgoldste@nsf.gov), National Science Foundation, USA

Abstract

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago and Indiana University collaborated on a major research demonstration at the IEEE/ACM Supercomputing '98 (SC'98) conference in Orlando, Florida, November 7-13, 1998, to showcase the evolution and importance of global research community networking. Collaborators worked together to solve complex computational problems using advanced high-speed networks to access geographically-distributed computing, storage, and display resources. It is this collection of computing and communication resources that we refer to as the International Grid (iGrid).

This paper presents an overview of the iGrid testbed, some of the underlying technologies used to enable distributed computing and collaborative problem solving, and descriptions of the applications. It concludes with recommendations for the future of global research community networking, based on the experiences of iGrid participants from the USA, Australia, Canada, Germany, Japan, The Netherlands, Russia, Switzerland, Singapore, and Taiwan.

Table of Contents

- What is iGrid?
- What is STAR TAP?
- iGrid Enabling Technologies
- iGrid Applications Classified by Type
- Recommendations for the Future

- Conclusions
- Acknowledgments
- References

What is iGrid?

In the research community, *computational grids* are emerging. They are the aggregate hardware and software resources that scientists require to solve extremely complex problems. On the hardware front, a computational grid is a collection of geographically-distributed resources: networks, computers, data stores, and visualization/virtual-reality displays. On the software front, it is the “middleware” necessary to integrate this ensemble so that its many and varied pieces operate as if they were one. We know this type of distributed computing can work on individual heroic projects—the challenge is to make it work seamlessly, efficiently, and routinely, independent of geographical boundaries, so that it becomes as ubiquitous and encompassing as the electrical power grid is today.[3]

The National Computational Science Alliance [<http://alliance.ncsa.uiuc.edu>], one of two efforts supported by the National Science Foundation (NSF) as part of its Partnerships for Advanced Computational Infrastructure (PACI) initiative, is a partnership of over 50 USA academic institutions, industrial research facilities, and government laboratories. The Alliance is building the National Technology Grid—a prototype of the 21st century’s computational and information infrastructure. EVL and Indiana University, Alliance partners, are extending this effort internationally and created the International Technology Grid (iGrid) testbed [<http://www.startap.net/igrid>] at SC’98 [<http://www.supercomp.org/sc98>] to demonstrate global collaborations.

Randy Bramley, a researcher from Indiana University who participated in iGrid, noted that: “The ‘Grid’ as a set of hardware (machines and networks) has existed for at least 10 years now. What distinguishes current Grid research, however, is viewing the Grid as a distributed set of services and providers. If we are to persuade people that iGrid really exists, it will require demonstrating unique functionality that was not available before and is not available otherwise.”

What is STAR TAP?

The centerpiece of iGrid is the NSF-sponsored initiative called STAR TAP, the Science, Technology And Research Transit Access Point [<http://www.startap.net>]. STAR TAP is a persistent infrastructure to facilitate the long-term interconnection and interoperability of advanced networking in support of applications, performance measuring, and technology evaluations. It is managed by EVL, Argonne National Laboratory [<http://www.mcs.anl.gov>], and Chicago’s Ameritech Advanced Data Services (AADS) [<http://nap.aads.net/main.html>].

Started in 1997, STAR TAP anchors the international component of the NSF very high-speed Backbone Network Service (vBNS) [<http://www.vbns.net>]. Canada’s CA*net II [<http://www.canarie.ca>], Singapore’s SingAREN [<http://www.singaren.net.sg>], and Taiwan’s TANet2 [<http://www.nsc.gov.tw>, <http://www.nhc.gov.tw>] are connected. Russia’s MirNET [<http://www.mirnet.org>] in cooperation with the University of Tennessee, is connected. The Asian Pacific Advanced Network (APAN) consortium [<http://www.apan.net>], which includes Korea,

Japan, Australia, and Singapore, in cooperation with Indiana University and the TransPAC initiative [<http://www.transpac.org>] is also connected. Connections from the Nordic countries' NORDUnet [<http://www.nordu.net>], France's Renater2 [<http://www.renater.fr>], The Netherlands' SURFnet [<http://www.surfnet.nl>], Israel's QMED networking initiative, and CERN [<http://www.cern.ch>] are imminent. Five USA Next Generation Internet (NGI) [<http://www.ngi.gov>] networks are connected: NSF's vBNS, US Department of Energy's ESnet [<http://www.es.net>], US Department of Defense's DREN [<http://www.hpcm.dren.net/Htdocs/DREN/>], and NASA's NISN and NREN [<http://www.nren.nasa.gov>]. And, the new Internet2/Abilene network [<http://www.internet2.edu/abilene/>] will also soon connect.

iGrid Enabling Technologies

iGrid applications depend on emerging services, such as resource control and reservation, new protocols, and high-bandwidth global grids. Particular emphasis was placed on distributed computing applications and the use of shared virtual spaces. [2] Two enabling technologies, used by a variety of the applications demonstrated, are described here.

CAVERNsoft [<http://www.evl.uic.edu/cavern>] is the systematic software architecture for the CAVE™ Research Network (CAVERN), an alliance of research and industrial institutions equipped with CAVEs, ImmersaDesks™, and high performance computing resources, interconnected by high-speed networks. [6][7] CAVERNsoft focuses on *tele-immersion*—the union of networked virtual reality and video in the context of significant computing and data mining—and supports collaborative virtual reality in design, training, education, scientific visualization, and computational steering. Developed at EVL, CAVERNsoft is designed to enable the rapid construction of tele-immersive applications; to equip previously single-user applications with tele-immersive capabilities; and, to provide a testbed for research in tele-immersion.

Globus [<http://www.globus.org>] is a toolkit of services to support innovative and high-performance distributed computing application development. [8] Globus services address key challenges that arise in wide-area, multi-institutional environments, such as communication, scheduling, security, information, data access, and fault detection. They make it possible, for example, for applications to locate suitable computers in a network and then apply them to a particular problem, or to organize communications effectively in tele-immersion systems. Globus services are used both to develop higher-level tools (e.g., CAVERNsoft) and directly in applications (e.g., Cactus, see application #1 below). The Globus toolkit is the work of a multi-institutional research and development team at Argonne National Laboratory and the University of Southern California, and also involves other institutions within the USA and around the world.

iGrid Applications Classified by Type

Foster and Kesselman, in their book *The Grid: A Blueprint for a New Computing Infrastructure* [3], propose a taxonomy of application types for grid-based computing. iGrid applications have been organized according to this taxonomy and are described below.

- **Distributed computing:** Parallel applications running on globally distributed computing elements, requiring Application Programmer Interfaces (APIs) and toolkits for resource allocation, authorization, authentication, and communications, such as Globus.

- **On-demand computing:** Remote access to instruments such as electron microscopes and accelerators.
- **Data-intensive computing:** Networked database search and visualization.
- **Collaborative computing:** Tele-immersion and shared virtual workspaces, typified by CAVERNsoft applications and collaborative medical diagnosis systems, which may also require access to computing or database resources in real time.

Distributed Computing

1. Industrial Mold Filling Simulation Using an Internationally Distributed Software Component Architecture [Canada, USA]

Indiana University, USA; Argonne National Laboratory, USA; Los Alamos National Laboratory, USA; Industrial Materials Institute, NRC, Quebec, Canada; Centre de Recherche en Calcul Appliqué (CERCA), Montréal, Canada

<http://www.extreme.indiana.edu/sc98/sc98.html>

This application (Figure 1) provides an integrated solution environment for a 3D parallel finite element code modeling industrial material processes, such as casting and injection molding. It won the SC'98 High Performance Computing (HPC) Challenge Award for Best Industrial Collaboration [<http://www.supercomp.org>].





Figure 1: This application ran on Indiana University's supercomputers (SGI Origin 200, IBM SP2 and SGI Power Challenge), NCSA's SGI Origin 2000s, machines at the Institute for Industrial Materials, the CAVE at the SC'98 SGI booth, and ImmersaDesks at the SC'98 Alliance booth and SC'98 iGrid booth. The first image shows the real-time demonstration. The second image is a mold for an engine piston. The filling "gate" is at the left, and two "pipes" provide reservoirs as the part cools and contracts. The gates and pipes are cut off later and the final part machined to tolerance. The mold itself is rendered translucent to allow viewing the piston and the cooling channels in the mold, and the fluid is color-coded according to pressure.

2. Metacomputing the Einstein Theory of Spacetime: Colliding Black Holes and Neutron Stars Across the Atlantic Ocean [Germany, USA]

Max-Planck-Institut-fuer-Gravitationsphysik, Albert-Einstein-Institut, Germany; National Center for Supercomputing Applications, USA; Argonne National Laboratory, USA; Washington University, USA

<http://jean-luc.ncsa.uiuc.edu/SC98>, <http://bach.ncsa.uiuc.edu/SC98>

This simulation (Figure 2) of the complete set of 3D Einstein equations of general relativity, performed using a new parallel computer code called Cactus, enables astrophysicists to study colliding black holes, neutron stars, the formation of singularities, and other aspects of Einstein's theory that cannot be handled by analytic means. It won the SC'98 HPC Challenge Award for Most Stellar [<http://www.supercomp.org>], and the Heinz-Billing-Preis award given each year by the Max Planck Society in Germany for outstanding work in computational science [<http://www.mpg.de/billing/billing.html>].




Figure 2: This application ran on the RZG (MPG in Garching, Germany) T3E/600 (728 nodes), ZIB (Konrad Zuse Institut in Berlin Germany) T3E/900 (128 nodes), SDSC T3E/900 (256 nodes), NCSA NT Cluster (128 nodes), NCSA SGI Origin 2000, Argonne National Laboratory SGI Origin 2000, and the ImmersaDesk in SC'98 iGrid booth. The image on the left shows the real-time demonstration. The image on the right is an isometric top view of double precision neutron star collision.

3. Metacomputing and Collaborative Visualization [Germany, USA]

Sandia National Laboratories, USA; Pittsburgh Supercomputing Center, USA; High Performance Computing Center Stuttgart (HLRS), a division of RUS, the Computing Center of Stuttgart University, Germany

<http://www.hlrs.de/news/events/1998/sc98>

This project (**Figure 3**) demonstrates research and industrial projects, including advanced visualization of 3D resistive magneto-hydro-dynamic (MHD) equations, using Alegra software, and molecular dynamics simulations of the mechanical stability of quasi-crystals using IMD software.





Figure 3: The distributed computing environment used for this application is diagrammed in Figure 4. The image on the left shows the real-time demonstration. The image on the right depicts a tiny sphere bumping into a thin plate of only two atomic layers. This “Hopping Ball” simulation was developed to test that IMD code, a software package to perform classical molecular dynamics simulations, works correctly on parallel machines.

4. Maximum Likelihood Analysis Of Phylogenetic Data [Singapore, Australia, USA]

Indiana University, USA; National University of Singapore, Singapore; Cooperative Research Center, Advanced Computational Systems (ACSys CRC), Australia

<http://www.indiana.edu/~rac/hpc/index.html>, <http://bic.nus.edu.sg/sc98.html>

DNA data has accumulated more rapidly in recent years than compute power, so researchers must often exclude potentially informative data to make statistical analysis practical. This application demonstrates the use of the computationally intensive maximum-likelihood method of phylogenetic inference on three medium-to-large datasets: cytoplasmic coat proteins, mycosporidia, and cyanobacteria.

5. Construction of Numerical Wind Tunnel Based on Design Procedure: From Aircraft Geometric Definition To Aircraft Flow Solutions [Taiwan, USA]

National Center for High-Performance Computing, Taiwan; National Cheng Kung University, Taiwan; National Chiao-Tung University, Taiwan

<http://www.nchc.gov.tw/RESEARCH/NWT>,
<http://www.nchc.gov.tw/RESEARCH/CFDEE/Publications/index.html>

The Numerical Wind Tunnel (NWT) is dedicated to CFD solutions of an industrial scale. It enables users to easily and efficiently compute, manage and visualize data in real time, and devise engineering designs. For SC'98, NWT is applied to a conceptually designed inter-island small business aircraft.

6. Parallel Computation of High-Speed Train Aerodynamics [Taiwan, USA]

National Center for High-Performance Computing, Taiwan; University of Minnesota, USA

<http://www.nchc.gov.tw/RESEARCH/CFDEE/Publications/index.html>

High-speed trains will have a possible maximum speed of 500km/h, ~Mach 0.4. Side branches in a tunnel are commonly used to reduce the pressure gradient of a compression wave generated at the inlet as a train rushes into the tunnel. This project demonstrates a parallelized 3D compressible Euler solver for high-speed train aerodynamic simulations.

On-Demand Computing

7. Remote Visualization of Electron Microscopy Data [Japan, Singapore, USA]

University of Southern California, Information Sciences Institute (USC/ISI), USA; National Center for Electron Microscopy and Imaging Research, National Biomedical Computing Resource, and San Diego Supercomputer Center/NPACI, UCSD, USA; Argonne National Laboratory, USA; Osaka University, Japan; Tokyo Institute of Technology, Japan; Waseda University, Japan; National University of Singapore, Singapore

<http://www.mcs.anl.gov/xray-cmt>, <http://www.ncmir.ucsd.edu/>

This application remotely processes and visualizes electron microscope data. Users access remote datasets and perform computationally intensive tomography, a 3D image reconstruction technique, for immediate viewing on an ImmersaDesk. The goal of this project is remote control of scientific instruments.

8. Tele-Manufacturing via International High-Speed Network [Singapore, USA]

Temasek Polytechnic, Singapore; National University of Singapore, Singapore; Indiana University, USA

<http://www.cir.nus.edu.sg/teleman>

Advanced networks are used to control rapid prototyping devices for manufacturing medical prostheses at Temasek Polytechnic. The devices are controlled by a Java application.

Data-Intensive Computing

9. JavaCMS—A Java 3D Particle Collision Event Viewer Using a Distributed Object Database Management System [Switzerland, USA]

Caltech/CERN, USA/Switzerland

<http://pcbunn.cithep.caltech.edu/>

This application is part of the Globally Interconnected Object Databases (GIOD) project, a joint effort among Caltech, CERN, and Hewlett-Packard Corporation. This JavaCMS application enables remote viewing of individual events in the large (~1TB) datastore of fully simulated particle collision events from CERN's Large Hadron Collider (LHC), due to begin operation in 2005.

10. TerraVision on the Grid: Interactive Immersive Fly-Throughs Using Distributed Image Servers [USA]

Lawrence Berkeley National Laboratory, USA; SRI International, USA

<http://www.ai.sri.com/~magic>, <http://www.ai.sri.com/TerraVision>, <http://www.magic.net>

TerraVision, part of DARPA's MAGIC project, was developed to demonstrate a high-speed, wide-area IP/ATM network for real-time terrain visualization and high-speed distributed storage systems. Users roam in real time about a 3D landscape, created from elevation data and registered aerial images and comprised of a terabyte of data. TerraVision uses VRML 2.0, enabling users with VRML browsers to visualize large datasets from remote locations.

11. Reciprocal Net—A Global Shared Database for Crystallography [USA]

Indiana University, USA

<http://www.iumsc.indiana.edu/>

The Reciprocal Net project, a collaboration of 14 crystallography laboratories, attempts to create a new type of database structure for crystallographic data. The interface allows users to graphically examine and manipulate the data, using simple Web browsers, advanced workstations, stereographic equipped workstations, and immersive technologies.

Collaborative Computing

12. Distributed Virtual Reality Technology in Collaborative Product Design [Germany, USA]

National Center for Supercomputing Applications (in cooperation with Caterpillar, Inc.), USA; German National Research Center for Information Technology, (GMD), Germany

<http://www.ncsa.uiuc.edu/VEG/DVR/>

The Distributed Virtual Reality (DVR) system demonstrates the potential of high-performance computing and networking to enable engineers at geographically remote facilities to work together on product designs using virtual reality. The system supports collaborative product and process design reviews.

13. Architectural Walk-Through Coupled with a Parallel Lighting Simulation [Netherlands, USA]

Academic Computing Services Amsterdam (SARA), The Netherlands; Calibre bv, The Netherlands

<http://www.sara.nl/hec/CAVE>

This architectural walk-through accurately depicts room lighting conditions using a parallel radiosity simulation that runs on a supercomputer. Every time lighting conditions change, the simulation computes the new room shading and sends the resulting model over a high-speed network for display on an ImmersaDesk. Within seconds, the new room lighting can be evaluated.

14. Exploring CAVERNsoft Tele-Immersive Collaboratories through the iGrid Portal [Singapore, Australia, Japan, USA]

EVL, University of Illinois at Chicago (UIC), USA; Argonne National Laboratory, USA; Virtual Reality Medicine Lab (VRMedLab), UIC, USA; Northwestern University (NU), USA; Old Dominion University (ODU), USA; Institute of High Performance Computing, National University of Singapore (NUS) (in cooperation with Motorola), Singapore; Tokyo University, Japan; ACSys CRC, Australian National University, Australia; Indiana Univ. and Indiana University-Purdue University Indianapolis (IUPUI), USA; Virginia Tech, USA

<http://www.evl.uic.edu/cavern/events/sc98/>

A virtual atrium acts as a central tele-transportation point to various “tele-immersive” collaboratories around the world: Cave6D (ODU/EVL); Motorola Impact Visualization (NUS/EVL); Virtual Temporal Bone (VRMedLab); The Silk Road Cave Shrines—Mogao Grottoes of Dunhuang in China (NU/EVL); The CAVE Collaborative Console (Virginia Tech); and, FutureCamp ‘98 (IUPUI).

15. The Ganymede Telebot, an Enabling Technology for Teleconferencing [USA]

EVL/UIC, USA

<http://www.evl.uic.edu/EVL/RESEARCH/telebot.shtml>

Accessibility for the handicapped goes far beyond providing ramps and elevators in buildings, cars, airplanes, and buses. Some, such as those afflicted by Multiple Sclerosis, do not have the health to travel even when physical access is available. The Telebot integrates teleconferencing with life-sized display screens, robotics, and high-speed networking to ensure that handicapped participants have both a virtual presence and an *equal* presence.

16. IUPUI Future Camp [USA]

IUPUI, USA

http://www.science.iupui.edu/future_camp98, <http://www.vrve.iupui.edu>

IUPUI Future Camp is a one-week, multidisciplinary, virtual-reality, day camp. Demonstrated are 3 projects created by 18 campers, grades 9-11, in June 1998: Virtual Art Gallery, Virtual Ocean Colonization, and Virtual Indianapolis Zoo.

17. 3DIVE—3D Interactive Volume Explorer for Collaborative Investigation of Medical and Scientific Volumetric Data [USA]

IUPUI, USA

<http://www.vrve.iupui.edu/>

3DIVE is a tele-immersive application enabling multiple participants to display, manipulate, and discuss volumetric biological datasets—generated by MRI, CT, or sectioning.

18. Constrained Navigation Techniques for Collaborative Virtual Environments [USA]

Indiana University, USA

<http://www.cica.indiana.edu/~ewernert/projects/c2nav/>

The purpose of a collaborative virtual environment is two-fold: it facilitates information and viewpoint sharing (expert/student relationship) while it simultaneously promotes individual exploration and personal insight (peer/peer relationship). Several navigation methods are explored.

19. IMSRacer [USA]

Lawrence Technological University, USA; University of Michigan, USA

www.oakland.edu/~dong/IMSRacer

This research enables users from a variety of disciplines to navigate and interact with 3D graphics models on any virtual-reality display with little programming effort. New navigational tools are applied to IMSRacer, a 25-meter long sailing yacht.

20. Real-Time Digital Video Stream Over IP [Japan, USA]

WIDE Project, Keio University, Japan

<http://www.sfc.wide.ad.jp/DVTS>, <http://www.jain.ad.jp/workshop/IWS99/>,
<http://www.jp.apan.net/meetings/981022-SC98/SC98-Reports/DV/index.html>

This is the first video demonstration between the USA and Japan over an IP network with digital video quality. This system encapsulates a digital video stream from a normal digital video camera using IEEE 1394 into IP packets without encoding delays. On the receiving end, the IEEE 1394 digital video stream is directly input into a digital video television or recorder.

21. Education ICT Network Segment at Moscow State University [Russia, USA]

Moscow State University, Russia

<http://info.phys.msu.su/SC98>

This academic year Moscow State University began a distance-learning experiment between the physics and geology departments to remotely teach “General Physics,” “Ecology,” and “Foreign Language” courses. Their goal is to continue course development and create an archive of teaching materials, potentially accessible worldwide via high-speed networks (MirNET).

22. GiDVn: Globe internet Digital Video Network [Singapore, USA]

International Center for Advanced Internet Research (iCAIR), Northwestern University, USA; Materials Science and Engineering, Northwestern University, USA; National University of Singapore, Singapore; Institute of Materials Research and Engineering, Singapore

<http://www.icaair.org/gidvn/>

This is a demonstration of ongoing work to develop a common digital video platform to communicate with national and international colleagues in real time over international high-speed networks.

Recommendations for the Future

The purpose of iGrid is not to justify international collaboration; many of these collaborations existed prior to the SC'98 iGrid event and will continue long afterwards. As Breckenridge (#3), asserts, "Science and research need to operate in the global community... Each organization brings different strengths to the problem... The two organizations are solving the same problems at the metacomputing and collaborative visualization level. It makes sense to apply double the resources."

The purpose of iGrid, an applications-driven integration testbed, is to enable network engineers, application programmers, and computational scientists to discover what works, and what doesn't, with today's evolving computing and network infrastructure in order to help advance the state of the art.

When asked about their experiences, participants responded that iGrid-like events were beneficial and they overwhelmingly recommended that they continue, but with several caveats:

- Global network infrastructures have to be persistent to enable collaborative problem solving, but conference-like events advocate/justify/promote the benefits.
- Demonstration sites have to be persistent. Temporary conference sites are plagued with problems.
- Research can be better explained and promoted—among peers, students, and the general public—through interactive multimedia presentations.
- Research sometimes needs impetus for advancement (not only for the researchers, but the network providers, facilities managers, etc.)—and deadline-driven demonstration events are helpful.

Persistent Global Network Infrastructures

While Canadian and Asian-Pacific international participants had stable connections to STAR TAP, the European connections were either temporary or non-existent. And, a few of those using the vBNS complained that it too was slow. Bramley (#1) switched to the commodity Internet. "Most of the time," he noted, "the vBNS was slower and we were better off to avoid it—possibly because it was clogged with other SC'98 demos."

For example, the German/USA applications (#2, 3, and 12) had temporary connections provided by Deutsche Telekom to Teleglobe to CA*net to STAR TAP to vBNS. By demonstrating what's possible, Shalf (#2) hoped to "apply more pressure on the organizations involved to make this capability permanent." Kindratenko (#12) had to learn "...how to establish reliable connections between remote sites involving quite a few network service providers. This is a challenging task,

knowing that they have different policies, run different equipment, support different and sometimes incompatible protocols, etc.” Breckenridge (#3) who provided a network diagram for her application (Figure 4), noted “the SC’98 network was a great testbed and now needs to be implemented for general usage.”

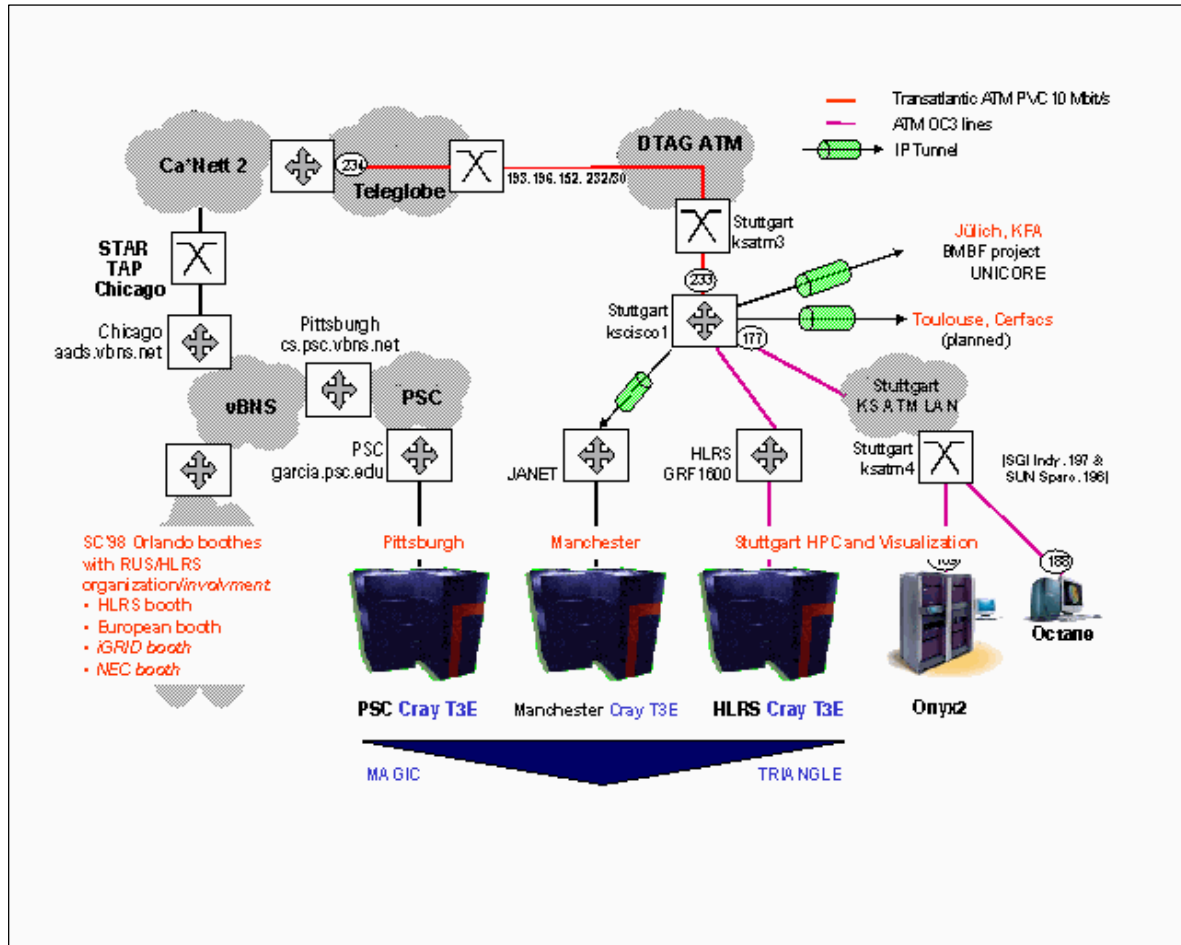


Figure 4: An example of the temporary connectivity established between Stuttgart University (Germany) to SC’98 in Orlando, Florida (USA) for application #3.

SURFnet is in the process of connecting The Netherlands to STAR TAP. While temporary connections could be made from some European countries to the USA for iGrid, it was not possible to connect SARA in Amsterdam. Researchers were given an account on NCSA’s SGI Origin2000 and relied on the vBNS for their demonstration. According to SARA researcher Janssen (#13), iGrid participation was important now, before permanent connectivity was in place (see Figure 5): “...to show the high performance community that SARA is committed to the STAR TAP initiative, and to show the SURFnet organization that SARA has the expertise and the applications to use such networks effectively.”



Figure 5: SARA researchers from The Netherlands demonstrating lighting simulations in an architectural walk-through (see application #13).

While the transoceanic connections were, for the most part, stable, much was learned about the way broadband networks operate. (Nakamura, #20) “iGrid let us experiment with Digital Video Streams over long distance networks.” (Lin, #5 and Chung, #6) “iGrid let us discover whether the current bandwidth would provide reasonable service for real-time tele-immersion.” (Ming, #8) “From Orlando, we controlled the Fusion Deposition Modeling machine in Singapore. The program didn’t run as fast as we would have liked. Perhaps this is because we are using TCP as the basic transmission protocol and TCP does not perform well in LFN network.” [Note: *High-speed networks between continents are commonly characterized by high bandwidths and high latencies. They are also referred to as Long Fat Networks (LFNs, pronounced “elephants”).*]

Persistent Demonstration Sites

Demonstrations at a conference site where a Nation’s networking, and now Global networking, has to be temporarily pulled into a convention center in less than a week’s time creates difficult conditions, not only for the applications people, but for the networking staff. Clearly, any criticism is not directed at the dedicated volunteers responsible for the conference infrastructure, but is a cry for persistent stable environments with high-end connectivity. For SC’98, multicast was not handled well, so many of the applications requiring multicast failed at their onsite attempts. Since iGrid demos were primarily multi-site, participants were relieved to know that multicast worked elsewhere. Leigh (#14) put a positive spin on otherwise frustrating conditions: “We learned that we needed to do additional testing of constantly collapsing conference networks to observe how our applications behaved under the worst situations.”

Promoting Research Among Peers, Students, and the Public

Among peers... (Nakamura, #20) “Many SC98 attendees were interested in our Real-Time Digital Video Stream (DVTS) project. Before SC’98, we thought that there were not many environments that could use DVTS, because DVTS needs so much bandwidth. But, we found out that many users have good environments, such as DoE, NASA, and university and company LANs. We have therefore accelerated our time schedule to complete our project and distribute the software. Current DVTS does not support the multicast. We should support multicast.”

(Lin, #5 and Chung, #6) “Through iGrid (see Figure 6), we hoped to share our numerical wind tunnel and high-speed train project experiences with other research groups and to seek opportunities for possible joint collaborations for application development. This goal was achieved... In December 1998, we collaborated with the ACR group at University of Iowa, headed by Judy Brown, using EVL’s CAVERNsoft, running from ImmersaDesk-to-ImmersaDesk. We provided our business jet and high-speed train geometric models and associated numerical solutions and collaboratively visualized the data. The quality of the connection was good for both audio and graphics. There was a 5~10 second delay in the audio, but the graphics was nearly real time. We also worked with the SCV group of Boston University, headed by Glenn Bresnahan, on a collaborative tele-immersive visualization. Our connection was midnight Friday (US time) and early morning Saturday (Taiwan time), which improved the quality of the networking. We got genuine real-time communication. Both collaborations were recorded and shown at Taiwan’s governmental opening of the T3 connection to STAR TAP/vBNS on December 22, 1998. We are currently following up with these two institutions for future collaborations.”

Among students... (Dong, #19) “I would like to utilize, at least part of, the wonderful collection of immersion and tele-immersion techniques in my classroom.” (see Figure 6)... (Sandalov, #21) We [Moscow State University] will be glad to collaborate with US universities in open distributed education on the networking level as well as on an education level. In Russia, we will set up server clusters for remote education proposed with Tver and Syktuvkar State Universities. In Germany, we will set up server clusters for remote education proposed with University at Duesseldorf.”

Among the public... (Kindratenko, #12) “Participation in iGrid allowed us to showcase the technology we have been developing and to see the reaction of audience, get feedback regarding the usability, practicability, possible applications, exchange ideas, etc.”... (Ming, #8) “It showed the potential of the Internet for manufacturing or prototyping.”




Figure 6: Numerical Wind Tunnel demonstration from Taiwan (#5) is shown on left. IMRacer demonstration from USA (#19) is illustrated on the right.

Advancing Research Through Deadline-Driven Events

(Kindratenko, #12) “Previous trials involved participants from two remote sites, but this time we collaborated with three.”... (Curry, #14) “We felt extremely validated in our conceptual efforts.”... (Lin, #5 and Chung, #6) “Our numerical wind/high-speed train demonstration...gave us confidence for the future of trans-Pacific collaborative scientific visualization work.”... (Shalf, #2) “Globus infrastructure is starting to gel (all of the low level stuff is in place where we need it and it mostly works).”... (Kesselman, #7) “It served as a catalyst for moving an existing APAN connection forward and motivated the deployment of Globus software to APAN sites.”... (Chen, #10) “This project initiated and demonstrated that digital video will be an important service in global digital networking.”... (Leigh, #14) “This event provided a forum to try out new CAVERNsoft technology; to collect performance data for off-line analysis; and, to meet other groups who want to collaborate with us. We’ve already started building improvements, based on our iGrid experiences, into CAVERNsoft.”... (Wong, #14) “The demonstration was a success. We were able to have simultaneous connections from Singapore to sites in Japan, Australia, Chicago, and Orlando.”

Conclusions

The majority of these collaborations existed prior to iGrid, and will continue in the months and years to come. iGrid did facilitate several new collaborations, which will hopefully also continue. If anything, iGrid helped scientists and programmers prioritize and focus their development activities for the coming year. One lesson learned by all involved is that as problems are identified and fixed, new problems emerge. Janssen (#13) noted, “Preparation is very important but you cannot prepare for everything.” In some cases there were hardware failures. In other, people failures. Lau (#10) actually never got his TerraVision application to run in the iGrid booth; when he experienced hardware failures at a remote site, he couldn’t locate someone to reboot the system! Says Lau, “The iGrid booth infrastructure worked great. Unfortunately, we were victims of software glitches, network failures, timing, and hardware failures beyond our control...Next time, I’ll have a videotape handy.”

A unique experience for everyone was the sociology of dealing with people from other cultures, and other time zones. (Kobayashi, #20) “The TransPAC link was a great; however, the time difference between Orlando and Tokyo is large. I could not take sleep during the demonstration.”... (Kesselman, #7) “Coordination of complex applications over multiple time zones is complex. Having a persistent infrastructure such as Globus and TransPAC will make it easier to perform these collaborations in the future. However, we need to advance asynchronous collaboration technologies beyond email in order to decrease turn-around time.”

Acknowledgments

Support by the National Science Foundation through grants CDA-9303433, with major support from DARPA (iGrid), ANI-9712283 (STAR TAP) and ANI-9730201 (TransPAC) is gratefully acknowledged.

References

1. Leigh, Jason, Andrew Johnson, Tom DeFanti, Maxine Brown, Mohammed Ali, Stuart Bailey, Andy Banerjee, Pat Banerjee, Jim Chen, Kevin Curry, Jim Curtis, Fred Dech, Brian Dodds, Ian Foster, Sara Fraser, Kartik Ganeshan, Dennis Glen, Robert Grossman, Randy Heiland, John Hicks, Alan Hudson, Tomoko Imai, Mohammed Khan, Abhinav Kapoor, Robert Kenyon, John Kelso, Ron Kriz, Cathy Lascara, Xiaoyan Liu, Yalu Lin, Theodore Mason, Alan Millman, Kukimoto Nobuyuki, Kyoung Park, Bill Parod, Paul Rajlich, Mary Rasmussen, Maggie Rawlings, Daniel Robertson, Samroeng Thongrong, Robert Stein, Kent Swartz, Steve Tuecke, Harlan Wallach, Hong Yee Wong, Glen Wheless. March 1999. A Review of Tele-Immersive Applications in the CAVE Research Network. IEEE Virtual Reality '99, Houston, TX (accepted for publication).
2. Donald F. McMullen, Michael A. McRobbie, Karen Adams, Douglas Pearson, Thomas A. DeFanti, Maxine D. Brown, Dana Plepys, Alan Verlo, Steven N. Goldstein. February 18-20, 1999 (accepted for publication). The iGrid Project: Enabling International Research and Education Collaborations through High Performance Networking. Internet Workshop '99 (IWS'99), Osaka, Japan.
3. I. Foster and C. Kesselman (eds.), *The Grid: A Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, 1999 [<http://www.mkp.com/grids>].
4. Holly Korab, "What Exactly is the Grid?," HPCWire, January 29, 1999, hpcwire@tgc.com
5. T.A. DeFanti, S.N. Goldstein, B. St. Arnaud, R. Wilder, R.C. Nicklas, P. Zawada, B. Davie, and M.D. Brown, "STAR TAP: International Exchange Point for High Performance Applications," INET '98, Internet Society's 8th Annual Networking Conference, July 21-24, 1998, Geneva, CD ROM.
6. J. Leigh, A. Johnson, and T. DeFanti, "CAVERN: Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments," *Journal of Virtual Reality Research, Development and Applications*, Virtual Reality Society, Vol. 2, No. 2, December 1997, pp. 217-237.
7. J. Leigh, *CAVERN and a Unified Approach to Support Real Time Networking and Persistence in Tele-Immersion*, PhD Thesis, University of Illinois at Chicago, 1997.
8. Foster and C. Kesselman, "Globus: A Metacomputing Infrastructure Toolkit," *International Journal for Supercomputer Applications*, Vol. 11, No. 2, 1997, pp. 115-128.