TransPAC Annual Report

Indiana University

NSF ANI-9730210

1 August 1998 - 31 July 1999

Table of Contents

A. Introduction

The following report provides a summary of the primary accomplishments and expenditures in the first year of funded activity under NSF Cooperative Agreement ANI-9730210, TransPAC: A High Performance Network Connection for Research and Education between the vBNS and the Asia-Pacific Advanced Network (APAN). A Program Plan and Budget for the second year (1999-2000) are presented.

B. TransPAC Organization and Management Committee

TransPAC Management Committee (TMC) is co-chaired by the TransPAC Principal Investigator (Michael A. McRobbie) and an APAN representative (Shigeki Goto). This committee establishes policy and provides overall direction for the TransPAC Project. TransPAC Management Committee members include:

APAN

Australia
- Robin Stanton (The Australian National University)
- Markus Buchhorn (Advanced Computational Systems, ACSys)

Japan
- Shigeki Goto (Waseda University)
- Kazunori Konishi (Kokusai Denshin Denwa, Co. Ltd.)

Korea
- Yong-Jin Park (Hanyang University)
- Kyungpyo Hong (Korea Telecom)

Singapore
- Francis Lee (National Technological University)
- Tham Chen Khong (National University of Singapore)

APAN Chair
- Kilnam Chon (Korea Advanced Institute for Science and Technology)

Tokyo XP
- Akira Kato (University of Tokyo)

Management Liaison
- Hiromasa Kaneko (Japan Science and Technology Corporation)

United States

Indiana University
- Michael McRobbie
- Doug Pearson
- Steve Wallace
- Jim Williams
- Rick McMullen
- Dennis Gannon
- Karen Adams

AT&T
- To be determined

STAR TAP
Linda Winkler (Argonne National Laboratories)

Management Liaison
Steve Goldstein (National Science Foundation)

The TMC met twice during this year (26 September 1998 in San Francisco USA, 17 February 1999 in Osaka Japan).

C. Milestones

1999

April 28	APAN, UCAID and TransPAC representatives meet to discuss terms of APAN/ UCAID MoU and establish outline for APAN/Abilene peering via TransPAC agreement
April 6	First HPIIS Team Meeting in Chicago
March 22-26	TransPAC-based demonstrations at GOIN'99
March 14	JST funded TransPAC network upgrade to 73Mbps bandwidth completed

1998

November 7-13	TransPAC-based iGrid demonstrations at SC'98
November 25	APAN / ESnet peering
November 24	APAN / NREN peering
October 23	APAN / vBNS peering
October 18	24x7 IU TransPAC NOC operational
September 21	NSF TransPAC 5-year award to Indiana University formally announced at Event in Washington, D.C.
September 21	Representatives of Indiana University and Japan Science and Technology Corporation sign MoU to establish link
September 8	Cooperative Agreement for TransPAC signed
August 18	IP routing KDD <-> IU via vBNS Downers Grove
July 30	ATM service established APAN Tokyo XP to STAR TAP

D. Progress Report

1. Progress Report: Attainment of HPIIS Objectives as Defined in the Statement of Work

A. Furnish, operate, and maintain a direct connection for high-performance traffic between the vBNS and APAN networks via the STAR TAP (Chicago).

Beginning August 1, 1998, TransPAC operated a 35 million bit per second (Mbps) variable bit rate non-real-time (VBR-nrt) Asynchronous Transfer Mode (ATM) permanent virtual path (PVP) connecting the APAN network at the Tokyo Exchange Point (XP) and KDD networks to the Ameritech Lucent switch located at the STAR TAP in Chicago. In May 1999 this infrastructure was expanded to 73Mbps using funding from the Japan Science and Technology Corporation (JST). The physical TransPAC network and the networks with which it peers are illustrated in Figure 1.

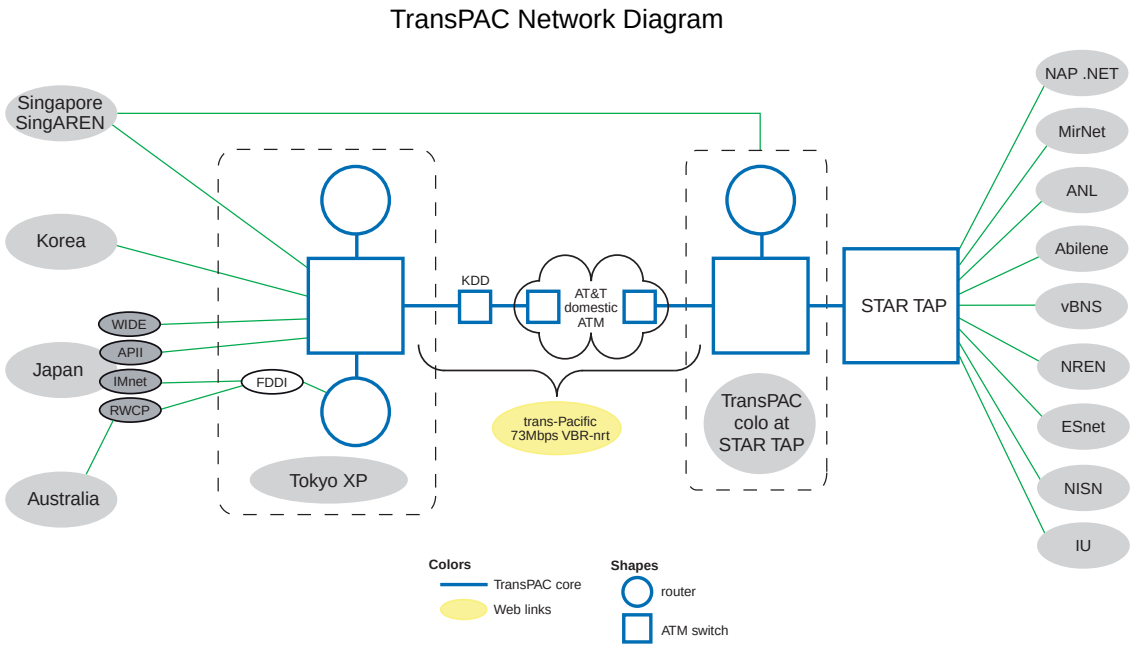

Figure 1

In addition to the trans-Pacific ATM service, the physical TransPAC network comprises equipment co-located with the STAR TAP at the Ameritech NAP facility in Chicago. This equipment includes:

Equipment rack:
- Western Telematic Netreach 16-port control port manager
- Western Telematic Netreach modem
- ServerTechnology Inc, Sentry 48 VDC remote power manager ServerTechnology Inc
- Sentry R-2000 110-VAC remote power manager
- HP Ethernet 8-port Ethernet repeater
- (2) Phone lines for remote access management

Network management and testing
- Fluke OC3port PLUS M/M ATM tester with remote software
- Accumatics rack mount PC for network management

ATM level connectivity
- Cisco Lightstream LS1010 ATM switch, with redundant power supplies
- ATM switch processor with Per Flow Queuing feature
- 4-port DS3 interface and a mixed OC3 interface (3 port multimode and 1 port single mode)

Layer 3 (IP) connectivity
- Cisco router (on loan STAR TAP)
- Cisco Ethernet 2E interface for 7500-class router
- Cisco ATM AIP interface for 7500-class router

TransPAC provides a Layer 3 Internet Protocol (IP) service via the TransPAC router co-located at STAR TAP. The router provides a peering point for the interconnection of APAN, the vBNS and other high performance research and education networks including Abilene, NREN and ESnet. In addition to facilitating peering, the TransPAC router provides a management point for the network, provides better behavior during network congestion (drop packets, not cells) and allows implementation of traffic shaping, differentiated services and policy.

As illustrated in the TransPAC Network Diagram, Figure 1 (page 9), a link to a commercial Internet service provider, NAP.NET is provided via a small portion of the increased bandwidth unilaterally funded by JST. TransPAC's original 35 Mbps bandwidth was augmented with an additional 38Mbps provided by JST. Of the 38Mbps increase, 3Mbps is used for the commercial connection. This service is provided by distinct ATM PVC and does not impact the HPIIS service.

TransPAC NOC service was officially begun in October 1998. The TransPAC NOC is an evolving partnership between Indiana University (IU), APAN and KDD (managers of the APAN Tokyo XP). The IU TransPAC NOC is physically located on the Indianapolis campus of Indiana University. Figure 2 illustrates the relationship between these NOCs and service providers.

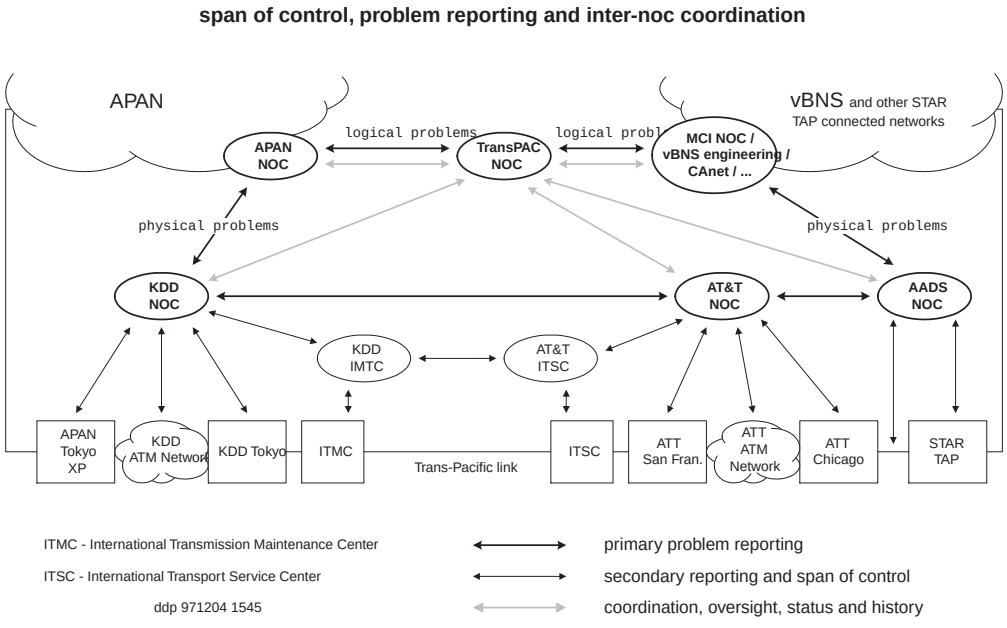

Figure 2

The IU TransPAC NOC is staffed 7 x 24 x 365. The NOC provides operations service (network monitoring, outage logging, trouble ticket generation and updating, maintenance coordination, etc.) and engineering service (circuit and router troubleshooting, configuration, etc.). The NOC web page is at http://www.transpac.org/noc/index.html.

In addition to providing TransPAC NOC service, the IU NOC provides NOC service (both engineering and operations) for Indiana University and for the Internet2 Abilene network. Further plans call for the expansion of IU NOC services to include STAR TAP and Eurolink (operational services only).

The current set of NOC tools ranges from standard, basic tools such as ping, pathchar, traceroute, MRTG and securing accounts on remote end-systems to more complex measurement tools such as OC3mon and Surveyor. In addition, the NOC and the TransPAC engineering staff are working to install more advanced troubleshooting tools.

The NOC trouble ticket system is based around a set of highly defined procedures and problem resolution protocols. The trouble ticket system resides on a server housed at the IU NOC, and related TransPAC trouble tickets are grouped together in an appropriate queue within the system.

When there is a problem on the TransPAC network, the NOC is alerted by either its network-monitoring software, or by other associated NOCs within the network. A trouble ticket is created immediately, and troubleshooting begins. An on-call NOC engineer is contacted, and works with the operations staff towards problem resolution. Appropriate support vendors and associated NOCs are contacted, if necessary. Problem reports and status updates are distributed to network clients via an email listserv.

The trouble ticket system is devised so varying levels of criticality and status can be assigned to a problem, and a built-in problem escalation method is employed. Closed trouble tickets are kept and studied to report network trends, and to build an online knowledge base of TransPAC network problems and resolutions.

B. Cooperate with STAR TAP and vBNS officials to ensure, to the extent supportable by prudent application of networking technology, that only approved institutions' traffic is permitted to use the high-performance connection:

The primary goal of the TransPAC network is to support high performance computing and communications applications and research between APAN institutions, and vBNS Authorized Institutions, and other global, high performance research and education networks meeting at the STAR TAP. To insure that uses of the TransPAC network are in support of this goal, all users must meet the following criteria:

1. Use is permitted for authorized educational institutions and national research laboratories conducting meritorious high performance applications. Traffic is permitted only between authorized institutions. Use of TransPAC by commercial organizations is not permitted without explicit approval of both the NSF and STA.

2. An institution's use of the TransPAC network must be in support of the TransPAC goal to support high performance computing and communications applications and research.

3. The TransPAC network is not to be used for general Internet connectivity. As such, TransPAC institutions should maintain separate Internet connections for commodity traffic or insure by other means that TransPAC facilities are not used for commodity Internet traffic.

4. An institution's use of TransPAC must not run counter to the AUP of any transited networks.

The TransPAC Management Committee evaluates all organizations seeking to use TransPAC resources to insure these criteria are met. The authorization process is defined as follows:

1. An institution submits an application to the TransPAC Management Committee for access to TransPAC resources.

2. The Committee, in consultation with other TransPAC committees such as the User Services Group, the Technical Coordinating Committee, and other partner organizations determines whether the application is or is not meritorious.

The TransPAC Acceptable Use and Authorization Policies are documented on the TransPAC web site at http://www.transpac.org/aup.html.

APAN institutions that are currently routed on TransPAC include:

Australia

Cooperative Research Center, Advanced Computational Systems (ACSys CRC)

Japan

Agency of Industrial Science and Technology (AIST)
Agriculture, Forestry and Fisheries Research Council (AFFRC)
Communications Research Laboratory (CRL)
Electrotechnical Laboratory (ETL)
Japan Advanced Institute of Science and Technology (JAIST)
Japan Science and Technology Corporation (JST)
Keio University Kokusai Denshin Denwa Laboratory (network management)
Kyoto University
Nara Advanced Institute of Science and Technology (NAIST)
National Cancer Center
National Cardiovascular Center
National Space Development Agency of Japan (NASDA)
Osaka University
Real World Computing Partnership (RWCP)
The Institute of Physical and Chemical Research (RIKEN)
Tokyo Institute of Technology (TIT)
University of Tokyo
Waseda University

Korea

Chungnam National University
Ewha University
Hanyang University
Information and Communications University
Inha University
Korea Advanced Institute of Science and Technology (KAIST)
Korean Meteorological Administration
Kyungpook University
Postech
Seoul National University
Soongsil University
Yonsei University

Singapore

Kent Ridge Digital Labs
Nanyang Technological University (NTU)
National University of Singapore

Thailand

Asian Institute of Technology (AIT)

Networks that peer with APAN at STAR TAP include the vBNS, Abilene, NREN, NISN and ESnet. An up-to-date list of the APAN institutions and peering networks is maintained on the TransPAC web pages at http://www.transpac.org/applications/routing.html.

To facilitate the use of TransPAC by APAN institutions that are on networks that contain a mix of authorized and unauthorized institutions, a source/destination policy routing infrastructure has been developed at the APAN Tokyo network exchange point.

The Indiana University Network Laboratory ran performance tests of Cisco source/destination policy routing as implemented in 7500-class RSP4 routers and found that router CPU utilization increased linearly and rapidly with the position matched in the access list. Standard source/destination policy routing would not work for a high performance network. A unique alternative strategy, a two-router solution, was developed. The process is as follows:

- The policy router (Router A) has a route map list that contains only the source network prefixes of authorized institutions. Destinations are wildcarded.
- The route map instructs that packets of any authorized institution are sent out an interface connected to Router B; all other packets are set out normal commodity paths.
- Router B performs normal routing functions. If the packet is bound for a STAR TAP connected network the packet is forwarded out TransPAC, otherwise the packet takes the default path back to Router A and is subsequently routed out the commodity path.

TransPAC Policy Routing - Conceptual Diagram

This highly simplified diagam shows only traffic flow from APAN institutions directed *toward* TransPAC/Internet

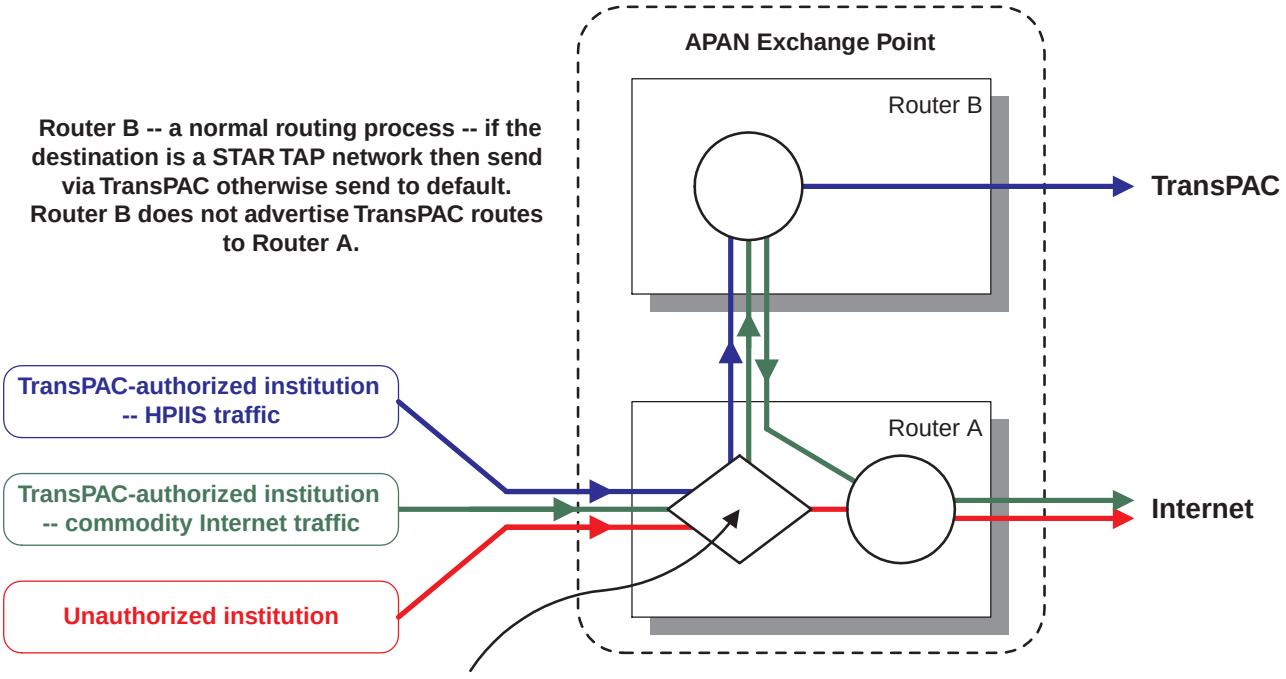

The policy routing architecture is documented on the TransPAC web pages at http://www.transpac.org/engineering.html.

C. Monitor the performance and use of the TransPAC connection:

Network utilization is monitored with Multi Router Traffic Grapher (MRTG). The five-minute moving average bits-per-second statistic for input and output is collected on the TransPAC Cisco LS1010 ATM switch. Daily statistics and weekly, monthly and yearly summaries are displayed on the TransPAC web page at http://www.transpac.org/noc/utilization.html.

Actual daily and yearly graphs of the interface to the trans-Pacific ATM service appear below. The time scale advances to the left as indicated by the red arrow. Green represents input bits-per-second (traffic coming from APAN). Blue represents output bits-per-second (traffic headed toward APAN). On the yearly graph, the additional colors of dark green and magenta represent the maximal 5-minute input and output rates that occurred during the summarization period; green and blue represent the average of the summarization period.

The daily graph for Thursday, July 8, 1999:

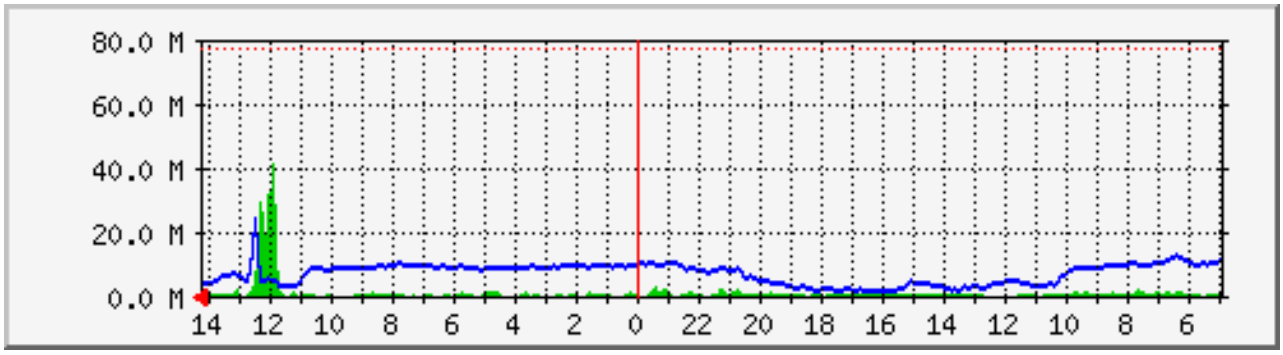

Figure 3

The yearly graph on Thursday, July 8, 1999:

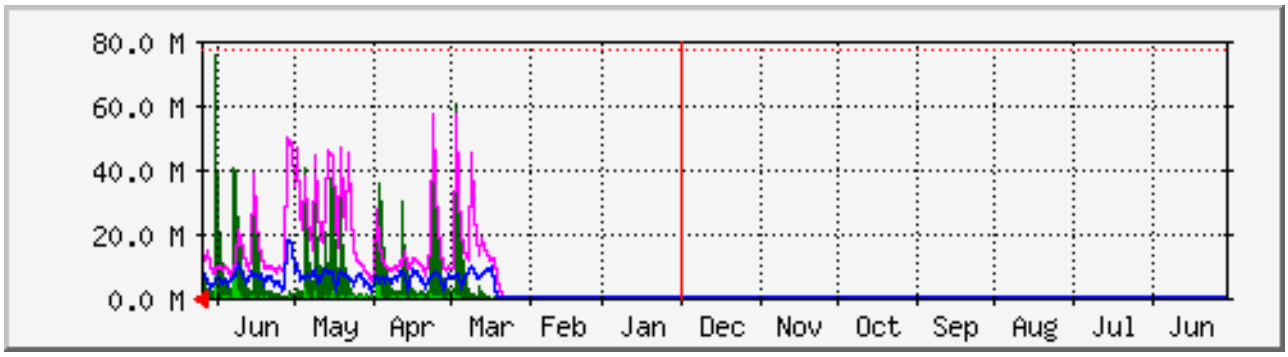

Figure 4

ANS Surveyor is employed to measure one-way delay and packet loss across the TransPAC/APAN networks. Surveyors have been installed at the Indiana GigaPoP, Korea Advanced Institute of Science and Technology and the APAN Tokyo network exchange point. Additional Surveyors will be installed

in Australia and Singapore. Surveyor statistics can be viewed at http://www.advanced.org/csg-ippm/. Figure 5 shows delay statistics and Figure 6 shows loss for Thursday May 22nd. The graphs shown for these days are not representative, but are included for illustration purposes.

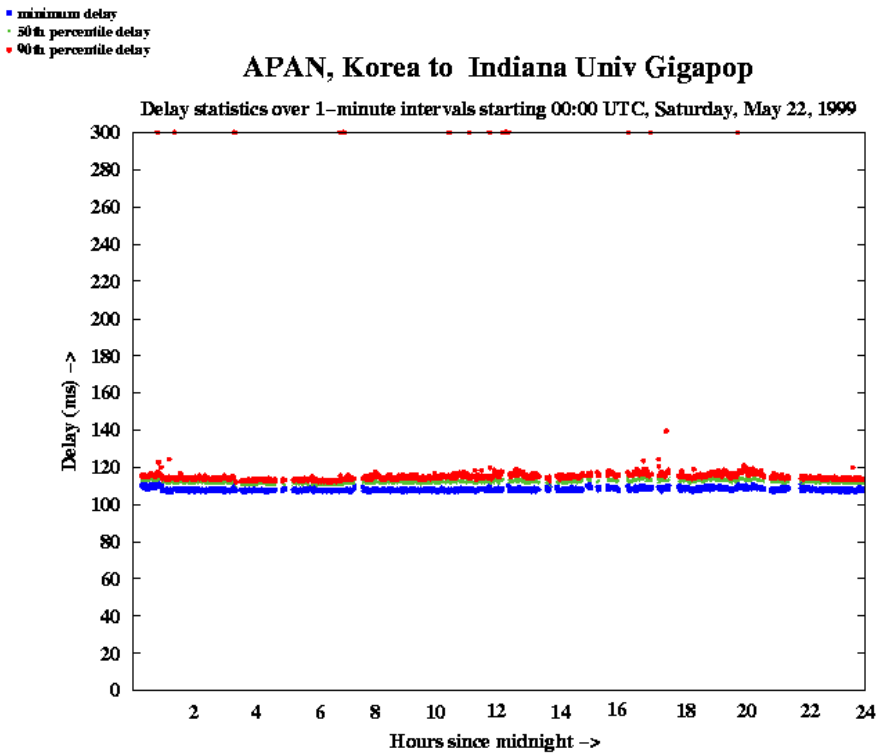

Figure 5

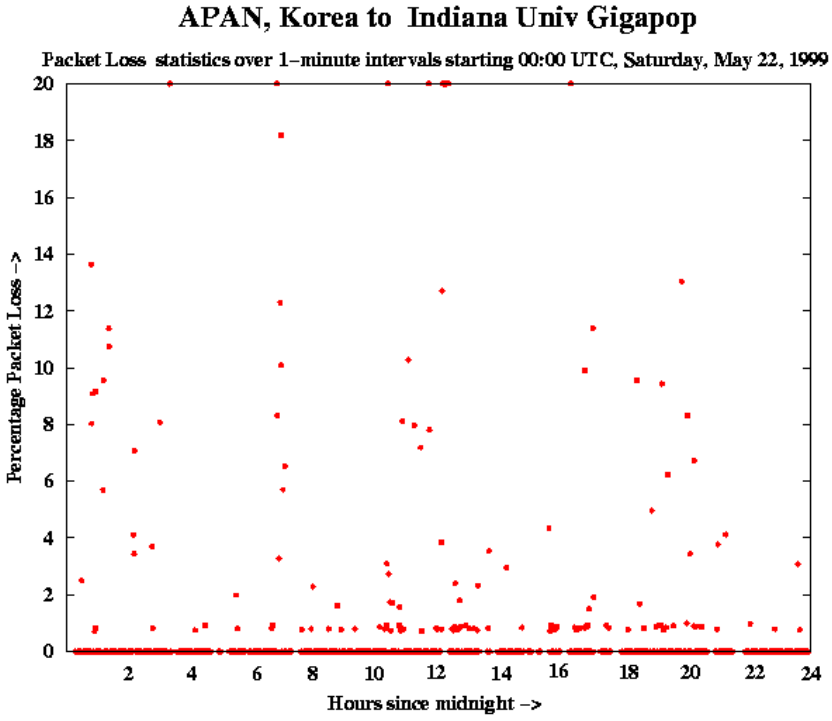

Figure 6

NLANR maintains an OC3mon for TransPAC, installed at the STAR TAP facility. Statistics are available in the NLANR Datacube, http://moat.nlanr.net/Datacube.

Availability of the TransPAC network is monitored with Ipswitch Whatsup. This system regularly pings critical network components and informs network operators and engineers of outages by graphical display and automated paging.

Over the period of time covered by this report the TransPAC network was largely stable. Specific network outages are detailed in the monthly reports submitted to the NSF. These reports can be found at http://www.transpac.org/documents.html. The causes of TransPAC outages show variation from month to month. No specific pattern is apparent. For example, in May there were significant problems with the AT&T circuit. In April there were a number of outages due to problems with the vBNS Downers Grove router.

TransPAC engineer, Linda Winkler, has performed significant analysis and characterization of the network IP throughput performance. This report is included in Section 4. Progress Report: Performance Analysis Report.

TransPAC maintains a remotely accessible OC3port ATM tester connected to the TransPAC ATM switch at STAR TAP. The OC3port is used to test quality of the underlying ATM network between STAR TAP at Chicago and the APAN Tokyo XP. Tests are performed during such ATM service changes as bandwidth upgrades.

D. Cooperate with the vBNS and the National Laboratory for Advanced Network Research (NLANR) Team (University of Illinois at Urbana (DAST), UCSD (MOAT) and Carnegie-Mellon University (NCNE) to develop testbed implementations and, as appropriate, production implementations of new versions of Internet networking protocols:

IPv6

A goal of TransPAC is to promote IPv6 development and gain operational experience. During this year, an IPv6 tunnel was established between JP and the 6bone via TransPAC. A 6bone link between WIDE and Singapore has been established.

The 6bone is a virtual network layered on top of portions of the physical IPv4-based Internet to support routing of IPv6 packets, as that function has not yet been integrated into many production routers. The network is composed of islands that can directly support IPv6 packets, linked by virtual point-to-point links called “tunnels.” The tunnel endpoints are typically workstation-class machines having operating system support for IPv6. The 6bone is thus focused on providing the early policy and procedures necessary to provide IPv6 transport in a reasonable fashion so testing and experience can be carried out.

Beginning in October 1998, production native IPv6 interconnections were established between ESnet, Internet2/vBNS, CANARIE, Cairn and WIDE. In addition, ESnet provides transit for all 6bone connected networks to all 6REN networks to guarantee continuity for early application and operational testing.

Multicast

A goal of TransPAC is to deploy an effective, ubiquitous, high-speed, persistent native IP multicast service with robust inter-domain connectivity to various high-performance backbone networks including Abilene, vBNS, and NREN. TransPAC has completed the first stage deployment of inter-domain multicast protocols MBGP/PIM-DM. Within APAN, inter-domain native IP multicast service is provided from Tokyo-XP to CRL, Seoul-XP, and SingAREN. In addition Tokyo-XP provides DVMRP

tunnels (non native-style multicast service) to Australia and IMnet. TransPAC has completed inter-domain multicast connectivity to the vBNS and NREN. Additional multicast services via access to the MBONE (traditional lower-speed non native-style Multicast Backbone services) are provided by a transit service across the vBNS.

It is worth mentioning a few words about the need for all of these multicast protocols. Multicast Border Gateway Protocol (MBGP) is standardized as RFC 2283, "Multiprotocol Extensions for BGP-4." MBGP uses these extensions to exchange inter-domain multicast routes.

Protocol Independent Multicast — Sparse Mode (PIM-SM), standardized as RFC 2362, was chosen primarily because it offers a relatively well-deployed implementation of a shared tree multicast protocol. In order to support the multi-megabit multicast streams between sites, a flood-and-prune type of distribution (used in dense mode protocols) is not feasible. The advantage of sparse mode on the backbone is that the distribution is join-oriented, and provides for a much more efficient distribution of data. Experience in the vBNS has proven PIM-SM can interconnect to other domains, some of which might use other multicast protocol types, e.g., DVMRP, other dense mode protocols, or other sparse mode protocols.

Multicast Source Discovery Protocol (MSDP) is a near-term solution for connecting intra-domain shared trees together in the inter-domain. This is accomplished by sources communicating information about sources among the domains. MSDP is applicable to shared tree protocols such as PIM-SM and CBT, as well as other protocols that keep active source information at the borders.

RSVP

The TransPAC NSF Proposal states in Advanced Services C.3.6.4 that "RSVP, among other methods, will be investigated as means to reserve bandwidth for IP applications." Resource Reservation Protocol (RSVP) is losing widespread favor as a method for providing end-to-end quality of service for IP applications. RSVP is more complex and demanding than Differentiated Services (DiffServ) in terms of router requirements and can negatively impact backbone router performance. TransPAC has suspended efforts to investigate RSVP as a method to provide backbone QoS.

QoS

TransPAC and APAN are participants in the Internet2 Quality of Service Backbone (QBone) project, within the context of a Joint Proposal organized by the International Center for Advanced Internet Research (iCAIR). Details regarding the Qbone project can be found at http://www.internet2.edu/qos/qbone/. The Joint Proposal and other details of participation are located on the TransPAC web pages at http://www.transpac.org/projects/qbone/index.html.

Participants in the Joint Proposal include:

APAN
Centre for Telematics and Information Tech., U of Twente
Electronic Visualization Laboratory, University of Illinois, Chicago
IBM Thomas Watson Research Center
iCAIR/Northwestern University
Indiana University
KDD - Kokusai Denshin Denwa Co. Ltd.
Korea Telecom
MREN - Metropolitan Research and Education Network
Nanyang Technological University
National University of Singapore
SingAREN

STARTAP
SURFnet
TransPAC

The objective of the proposal is to develop an optimal design and early implementations of state-of-the-art QoS concepts within an international testbed linked to the I2 networks. The international effort will address issues of common understanding of QoS design, definition of types and priorities of applications and required resource guarantees, management policy options for differentiation of service categories, distributed governance and resource allocation mechanisms across multiple domains, baseline standards, reference configurations and dependability parameters for auditing performance to contract, common response to non-compliance and priority conflicts, bandwidth brokers, policy servers, measurement and network management.

Progress in the effort has been slow. DiffServ standards are still in draft and router vendor implementations of pre-standard QoS features are incomplete and vary. The overall QBone architecture is just now beginning to take shape. TransPAC and APAN continue to be participants in this effort and will be early adopters.

Cache
The APAN Cache Working Group maintains a root Squid cache server at the APAN Tokyo XP. Reference to the service is http://cache.jp.apan.net/.

E. Maintain a publicly-accessible TransPAC-HPIIS web site containing information about the APAN-vBNS high-performance connection and research and education collaborations that it enables:

The TransPAC website (www.transpac.org) went online in fall 1998 and includes up-to-date international collaboration information, points of contact, engineering information links, documentation and mailing lists to facilitate additional collaboration and experience-sharing. Current statistical information and summaries about the performance of the TransPAC international connection are available online as well.

F. In concert with the methodologies developed by the NLANR Team, provide consultative user services supporting the use of TransPAC-HPIIS for high-performance computing and communications applications. Provide web-based front-end tools and direct user support to enable direct access to differentiated network services:

For most users of TransPAC the main resource to be allocated is a lower bound to bandwidth across the TransPAC link. We are approaching this problem in several ways, through participation in US and APAN QoS projects and testing of vendor implementations of differentiated services code for routers, and through manual scheduling and notification.

As mentioned above TransPAC engineering staff are participating in the I2 QBone QoS testbed. Digital video has been identified as a primary test application for diffserv-based QoS and digital video servers have been located in Japan and Singapore for testing router implementations of diffserv across TransPAC.

At present, without production quality QoS support for IP traffic, we are depending on manual coordination of individual high bandwidth applications to assure adequate resource availability. A web request and notification system for users requiring exceptional bandwidth is under development and will be placed in service next fiscal year.

Efforts in this fiscal year have been focused on improving overall layer 3 performance for IP applications. Specific problems existed with IP performance over the link early on that required extensive interaction between TransPAC, STAR TAP and AT&T engineering staffs. These problems have been resolved but this exercise pointed out the need for close cooperation between end-users and network engineers.

Because of the large bandwidth-delay product across the link (~1.3MB) TCP window size is a limiting factor for throughput. Work is currently underway to explore the TCP parameter space for optimum kernel and applications settings for applications with high throughput requirements. Tuning for specific applications has been accomplished using information from network test equipment and tcpdump. A set of instrumentation and tuning routines is under development to assist end-users in determining and improving the performance of their applications.

Performance problems diagnosed and corrected thus far are typified by the two applications described below. The first, mirroring of AP-Bionet data to and from Indiana University via ftp, suffered extremely low transfer rates. The second, transfer of large model data files to a prototyping machine using a custom Java application, also showed inexplicably poor performance. The ftp mirroring problem was diagnosed using network monitoring tools running on a workstation and resolved through tuning the TCP stacks at both ends of the transfer to allow larger buffer and TCP window sizes. The performance of the telemanufacturing application was improved simply by increasing the socket buffer sizes (TCP window scaling and large windows were already supported by the end stations).

TransPAC engineering staff are actively engaged with NLANR activities and participated in the fall ’98 and spring ’99 NLANR Joint Techs meetings.

2. Progress Report: Collaboration and Application Testbeds

TransPAC was deeply involved in two major networking projects during the past reporting year that were very successful at promoting international collaborations using advanced high-speed networks. The iGrid was sponsored by Indiana University and the University of Illinois at Chicago.

iGrid - International Grid Research Demonstrations at SC’98

Quoting from the iGrid web site www.startap.net/igrid, “The SC’98 International Grid (iGrid) research demonstrations showcase international collaborations using advanced high-speed networks, enabling researchers to work together, whether their colleagues live across the country or across the ocean, and to access geographically-distributed computing, storage, and display resources.”

“Applications in education, environmental hydrology, cosmology, medical imaging, molecular biology, and manufacturing use technologies such as remote instrumentation control, tele-immersion, real-time client server systems, multimedia, tele-teaching, digital video, distributed computing and high-throughput, high-priority data transfers. These applications and technologies depend on end-to-end delivery of multi-tens-of-megabits bandwidth with quality of service control, and need the capabilities of emerging Internet protocols for resource control and reservation.”

Specific projects involving TransPAC connectivity and assistance included:

1. Digital Video (DV) Stream on IEEE 1394 Encapsulated into IP Over the Long-Distance, High-Speed Internet Link — the first video demonstration between the USA and Japan over an IP network with DV 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 DV stream is directly input into a DV television or recorder.

2. Remote Visualization of Electron Microscopy Data — 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.

3. Tele-Manufacturing via International High-Speed Network — The Global Design and Manufacturing (GDM) project uses advanced networks to control rapid prototyping devices for manufacturing medical prostheses at Temasek Polytechnic. The devices are controlled by a Java application developed by GDM researchers.

4. Exploring CAVERNsoft Tele-Immersive Collaboratories through the iGrid Portal — This demonstration started with a virtual atrium that acts as a central tele-portation 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-Mogoa Grottoes of Dunhuang in China (NU/EVL); The CAVE Collaborative Console (Virginia Tech); and, FutureCamp '98 (IUPUI).

5. Maximum Likelihood Analysis of Phylogenetic Data — Two data sets were analyzed. In the first data set, contributed by the collaborators at the BioInformatics Center at the National University of Singapore, cytoplasmic coat proteins (involved in intracellular membrane transport) were sequenced from human, rat, bovine and yeast organisms. The second data set addressed the controversial phylogenetic placement of microsporidia (a parasite group including important human pathogens), with a dataset including representatives of most eukaryotic lineages (> 100 taxa). Some genetic studies find these to be highly degenerate fungi, while others, based upon small subunit rRNA, suggest an ancient eukaryotic lineage; resolving this question bears upon the reliability of ssurRNA-based phylogenetic analysis.

The systems, linked by the vBNS, TransPAC, and APAN networks, were at Indiana University, the Institute of High Performance Computing at the National University of Singapore, and the Cooperative Research Centre for Advanced Computational Systems (ACSys CRC) in Australia. Demonstrating computationally intensive analyses using a globally distributed collection of computational nodes paves the way for scientists connected by advanced networks to access remote servers, contribute key data sets, and collaborate with distant researchers. Our initial focus is on molecular biology in Indiana, Singapore, and Australia. Plans are underway to extend this partnership, addressing questions of performance analysis, virtual accounting schemes, and the development and expansion of the user community.

Global Observation Information Network (GOIN) 99

Quoting from the GOIN99/TransPAC web site www.transpac.org/applications/goin99/application.html "In 1993 President Clinton of the United States and Prime Minster Miazawa of Japan implemented GOIN under the Common Agenda for Cooperation in Global Perspective. The purpose of GOIN is to exchange global environmental data via information networks within and between the U.S. and Japan, to identify existing capabilities and weaknesses, to strengthen bilateral cooperation in Earth observation programs, and to serve as a prototype for future global observation network collaboration. GOIN workshops have been held yearly, 1996 through 1999 to bring together science and technology in support of the Common Agenda goal of protecting the global environment. The GOIN99 Workshop was held 22-25 March at the Hawaii IMIN International Conference Center, the East-West Center, University of Hawaii."

"As a first step toward establishing an earth observation (EO) prototyping plan over APAN-TransPAC, NASA and NASDA are cooperating with other agencies to provide prototyping demonstrations of EO applications for the Global Observation Information Network (GOIN) 1999 Plenary and Workshops.

Results from this first phase will then be used to develop a second phase of this proposal, with 'next step' activities to be defined and a management and prototyping plan to be submitted to APAN and TransPAC."

Projects that directly involved TransPAC facilities were:

1. IGBP Global Observation of Land Use and Land Cover — The IGBP-DIS Global 1km Land Cover Data Set (version 1), or more concisely, the IGBP DISCover, consists of the classification scheme, input data used to derive the classification, plus the validation database. The IGBP-DIS Global 1km Land Cover Data Set is archived and distributed from the EDC DAAC. Regional validation centers are located in North America, Eurasia, South America, Africa and Australia/Pacific. This project demonstrated distribution of IGBP-DIS Global 1km Land Cover data products to regional analysis centers and return of analyzed products to the product archive and distribution center in Sioux Falls.

2. Earth Monitoring and Disaster Warning — This project demonstrated the distribution of remote sensing data and derived information products to new sets of data users on APAN nodes. The demo democratized the decision making impact of the distributed data through local media outlets, universities, government laboratories, and non-governmental organizations (NGOs). The project monitors the feedback from data users and attempts to respond through changes and adjustments in the products and delivery patterns to better serve the objectives of the end users. The project addresses remote sensing of three phenomena which have wide impact on the region: 1) wildfires (biomass burning), 2) the intensity of fishing efforts, and 3) torrential rains.

3. Distributed image spreadsheet — The project displays 3-dimensional time series animation data from various Earth observation satellites (including NOAA, GMS, GOES, TRMM, etc.) simultaneously to visualize climate and weather change globally or regionally by using remote data servers over high-end networks. The distributed image spread sheet system has only operated in a local area network environment due to high-end network requirements (about higher than 6 Mbps). At GOIN'99, NASDA installed a image spreadsheet server to provide TRMM data at EORC with a 6 Mbps ATM link to access APAN. NASA operated another server to provide NOAA, GMS, TRMM and GOES at GSFC with NREN connection. Then, both servers were accessed by an ISS specific client at GOIN99. See http://rsd.gsfc.nasa.gov/rsd/

3. Progress Report: Papers and Presentations, Meetings, and Conferences

Note: Many of these papers and presentations are available at www.transpac.org

1997, December 13-14
Indiana University; Bloomington, IN USA
TransPAC All-Hands International Meeting

1998, March 1
Tokyo, Japan
2nd Meeting of the TransPAC Consortium

1998, March 2
Tsukuba, Japan
APAN Committee and Working Group Meetings

1998, March 3-5
Tsukuba, Japan
Internet Workshop 1998 (IWS'98) and Worldwide Computing and Its Applications (WWCA'98) Conferences
Presentation: (IWS'98) TransPAC – A High Performance Network Connection for Research and Education between the vBNS and the Asia-Pacific Advanced Network (APAN); Doug Pearson

1998, April 9
Indiana University; Indianapolis, IN USA
Internet2 and IU
Presentation: TransPAC – A High Performance Network Connection for Research and Education between the vBNS and the Asia-Pacific Advanced Network (APAN); Michael McRobbie

1998, June 23-24
Tokyo, Japan
Japan Science and Technology Corporation and Indiana University Planning Meeting

1998, June 24-26
Tokyo, Japan
TransPAC-JP NOC meeting

1998, June 30-July 2
Seoul, Korea
KRNET'98
Presentation: TransPAC – Policy-based Routing and Differentiated Services in TransPAC; Doug Pearson and Allen Robel

1998, July 1-3
Seoul, Korea
APAN Committee and Working Group Meetings

1998, July 21-24
Geneva, Switzerland
INET'98
Paper: A High Performance Network Connection for Research and Education between the vBNS and the Asia-Pacific Advanced Network (APAN); Michael McRobbie, et al
Presentation: A High Performance Network Connection for Research and Education between the vBNS and the Asia-Pacific Advanced Network (APAN); Michael McRobbie

1998, July 28-31
Chicago, IL
7th International Symposium on High Performance Distributed Computing (HPDC-7 '98)
Paper: A High Performance Network for Research and Education Applications Between the vBNS and the Asia-Pacific Advanced Network (APAN); Michael McRobbie, et al

1998, August 29-30
Indiana University; Indianapolis, IN USA
TransPAC Engineering Meeting

1998, September 21
Washington DC USA
NASA Policy Meeting
Presentation: TransPAC - Network Objectives and Policies; Doug Pearson

1998, September 24
Boulder, CO USA
NASA Technical Meeting
Presentation: TransPAC; Doug Pearson

1998 September 26-29
San Francisco, CA USA
Internet2 Project Meeting and Applications Demonstrations

1998, September 26
San Francisco, CA USA
TransPAC Status Meeting

1998, November 7-13
Orlando, FL USA
Supercomputing'98 and iGrid Demonstrations

1998, December 7-11
Orlando, FL USA
IETF 43

1998, December 8
Orlando, FL USA
CCIRN QoS Working Group
Presentation: Proposed APAN I2 QBone Activities; Doug Pearson

1999, January 10-14
University of Maryland; College Park, MD USA
Chinese-American Networking Symposium
Presentation: Indiana University: From Intranet to TransPAC; Chris Peebles

1999, February 17
Osaka University; Osaka, Japan
TransPAC Management Committee Meeting

1999, February 18-21
Osaka University; Osaka, Japan
Internet Workshop 1999 (IWS'99);
Presentation: Internet2 and Other Research and Education Networks; Michael McRobbie

1999, March 22-25
Honolulu, HI
Global Observation Information Network'99 (GOIN99) Conference

1999, April 6
Chicago, IL
HPIIS Team Meeting

1999, April 28
Washington, DC
Internet2 Member Meeting and Applications Demonstration

1999, May 17
Indiana University; Indianapolis, IN USA
Meeting of Indiana University and members of Computer Network Information Center, Chinese Academy of Sciences
Presentation: TransPAC; Steve Wallace and Doug Pearson

1999, June 22
San Jose, CA
STAR TAP International Advisory Committee Meeting

1999, June 22
San Jose, CA
STAR TAP Technical Advisory Committee Meeting
Presentation: IU NOC; Doug Pearson

1999, June 22-25
San Jose, CA
INET'99 Conference
Joint Communications Research Laboratory, Japan and Indiana University demonstration of IEEE1394 Digital Video over IP

4. Progress Report : Application Status Reports

The TransPAC team recently surveyed all investigators and projects listed in the initial TransPAC grant submission. All project contacts were asked to respond to a short e-mail message by outlining the progress of their investigation, their current use of the TransPAC network and their future plans and needs for international connectivity. These responses are summarized below.

Original Projects Currently Using TransPAC

The following projects are included in the original TransPAC grant response and are actively using the TransPAC network. (Appendix B provides further detail on the projects and the scientific content.)

- Asia-US-Australia Collaboration in the Silicon Vertex Detector Project for the BELLE High Energy Physics Experiment at KEK
- APBionet: Asia-Pacific Bioinformatics Network
- RHIC Pioneering High-Energy Nuclear Interaction Experiment (PHENIX) High-energy Heavy-ion experiment
- Web Cache Meta Network
- Japan-US Collaboration on ICRF Heating and Current Drive Experiments and Modeling
- Nucleic Acid Database
- Development of Mirror Server by Using High Speed Data Transfer in Genome Science
- Applications of Networked Virtual Reality Systems
- Japan-US Collaboration in Sloan Digital Sky Survey over the Network

New Projects Using TransPAC

There are also new projects, not originally included in the TransPAC proposal, that are actively using the TransPAC network. (Appendix C includes further detail on these projects and the scientific content.)

- Bio-Mirror public service for high-speed access to biosequence data
- Trans-Pacific Telemicroscopy

Original Projects Not Yet Using TransPAC

Not all projects listed in the TransPAC grant proposal used TransPAC resources in the past year. Reasons included conflicting funding and political concerns, lack of current need for high performance bandwidth, and delays in beginning the research. More specific responses to inquiries are provided here.

The following ongoing projects are collaborations between the National Aeronautics and Space Administration (NASA) in the USA and the Institute of Space and Astronautical Science in Japan. ISAS is connected to the US via a SINET connection in both STARTAP and Palo Alto. The question of an additional ISAS connection to APAN and TransPAC is both financial and political.

- Data processing and distribution of ASCA data over the network (ISAS/NASA)
- Japan-US collaborations in developing ASTRO-E science instruments using the network (ISAS/NASA)
- Data processing and distribution of ASTRO-E ground test and flight data over the network (ISAS/NASA)
- The Geotail Data Analysis in Correlative Solar-Terrestrial Study (ISAS/NASA)
- Data processing of IRTS data over the network (ISAS/NASA)
- Exchange of data between Infrared Astronomical Satellites (ISAS/NASA)

The following project is currently underway using a NASA specific link. The project is planning to request additional budget to connect to APAN and TransPAC in their next funding cycle.

- Tropical Rainfall Measuring Mission (TRMM) Science Projects

The principal investigators are not actively pursuing the following projects.

- Turbulence Closure Models for the Analysis of Three Dimensional Flows Volume Image Visualization and Simulation for Telemedicine

The following project has no need for high bandwidth connectivity at present. As its needs change, the investigators will contact TransPAC.

- Medical Internet Exchange Cooperation

The following project is active on a local basis in Hawaii. The investigators plan to expand the project internationally, using TransPAC, in the next year.

- High Speed International Telemedicine between NCC and Hawaii Medical Network

Incomplete Responses

Several project contacts either responded to inquires with incomplete information or did not respond at all. The TransPAC team will continue to attempt to contact these investigators with the expectation that these projects will become active and operational in Year 2 of TransPAC.

- Performance Measurement of High Performance Internet
- 6Bone, Mbone, telepresence
- Global resource scheduling
- Telepresence/CAVE-CAVE
- Tele-education, Distance teaching
- Performance measurement of high speed network
- Wide Area Parallel Computing
- Numerical Weather Prediction model
- International research in High Energy Physics
- Remote Research Center for Rheology
- Code Development of Remote Operations for Accelerator Test facility at KEK, Japan
- International Research Forum on Flow Information Technology
- Modeling of plasmas for Tokamak fusion and industrial processing
- Alcohol consumption and CVD, other cause specific and total mortality in the multiple risk factor intervention trial (MRFIT).

5. Progress Report: Performance Analysis Report

Beginning with the initial installation of the TransPAC circuit, TransPAC engineers have been working to verify service and measure applications level performance levels. While this activity has involved somewhat of a moving target, significant results have been attained. From the start, the 35 Mbps service level was verified with a Fluke OC3PORT ATM cell generator. By filling a loopback VC, the engineer proved that the service level contract supplied by AT&T/KDD across the trans-Pacific link were indeed obtained. The relevant traffic characteristics of the initial service include: non-real -time variable bit rate service, sustained cell rate of 91,146 cells per second (~35 Mbps), peak cell rate of 96,000 cells per second (~36.8 Mbps), and maximum burst size of 1,000 cells.

Moving up the protocol stack to determine bandwidth attainable by applications, engineers attempted to measure throughput between Tokyo-XP and Chicago using a software package (window ping) written at Pittsburgh Supercomputer Center to stress a link and measure where queues may occur by filling a pipe until it overflows. The use of this tool was problematic for two reasons. The application is disruptive to other activity on the link because it does not use TCP and is, therefore, an inconsiderate network citizen. Also, the code was not developed to run on various flavors of Unix and problems arose running it under FreeBSD (the only operating system available on the Tokyo end of the link). TransPAC engineers worked with NCNE engineers to overcome the second problem and the code now runs fine under FreeBSD. The next problem occurred within the version of FreeBSD installed at

Tokyo. Despite the fact that the test programs were setting buffer sizes appropriate for large bandwidth delay product networks, FreeBSD reset the send and receive buffer size. The result was that we were not testing what we thought we were. UDP performance results were measured Tokyo to Chicago, of 31.5 Mbps. This is encouraging as the theoretical available bandwidth after accounting for ATM framing and IP overhead is 31.4 Mbps (assumes a 1500 byte packet length). However, in the Chicago to Tokyo direction, results were only 21.6 Mbps due to "the receive buffer reset bug" on the receiving host in Tokyo. Because of problems with the FreeBSD kernel on the machine in Tokyo, alternative receivers were investigated. An additional PC was contributed to the project, which had ATM connectivity but still had a FreeBSD operating system.

On March 14, the bandwidth on the TransPAC service was upgraded to 70 Mbps. Before the service was switched, TransPAC engineers once again verified the service level provided by AT&T/KDD utilizing the OC3PORT ATM cell generator for a 24-hour period. The relevant traffic characteristics of the upgraded service include: non-real-time variable bit rate service, sustained cell rate of 200,252 cells per second (~76.9 Mbps) and a peak cell rate of 182,292 cells per second (~70 Mbps). Due to difficulties with the ATM implementation under FreeBSD, alternative receivers were once again investigated. Once again, the window ping utility from PSC was utilized to achieve a UDP data stream of 71.7 Mbps across the trans-Pacific link from a PC in Tokyo to host in Chicago. Investigation is ongoing to demonstrate similar throughput in the other direction, and to find methods for achieving similar performance for TCP streams.

E. Program Plan

The following information details the plans for theransPAC project during the second award year year(99-00) in some key areas. The program plan will be further defined as the year progresses and the project continues to evolve.

Core Network Service

Prompted by increased competition and an increase in supply, we expect bandwidth prices to decrease over the course of the TransPAC award. We also expect the demand for bandwidth to increase as collaborations increase. Built into our contract with AT&T (and into any subsequent contracts) is a clause that allows IU to "test the market" each year to ensure TransPAC is receiving the benefits of changes in the market. Our intention is to leverage the decreasing price of trans-Pacific bandwidth to secure increasing TransPAC bandwidth at a constant price. (Annually)

The value of the TransPAC network to researchers is directly proportional to the number of research sites that can be reached via high performance network connections. TransPAC will pursue peering with other high performance research networks via the STAR TAP, guided by the needs of collaboration and research. (Ongoing)

The STAR TAP proposal to the NSF for renewed funding describes reasoning to establish a layer 3 service (IPv4 routing) as an optional service to be provided by STAR TAP. TransPAC currently provides a layer 3 peering service via the TransPAC equipment co-located at the STAR TAP. TransPAC and STAR TAP management and engineers will work to transition the TransPAC-specific service to a generalized STAR TAP service. (Q4?)

NOC

The Indiana University NOC, which provides operations and engineering services for the TransPAC, the Internet2 Abilene and IU networks, will provide operations services for STAR TAP and the Eurolink HPIIS networks, based upon sub-awards in these NSF proposals.

STAR TAP did not require a NOC during its first phase since it provided solely a layer 2 interconnect based upon the Ameritech NAP. The current program plan for STAR TAP includes offerings of layer 3 IPv4 routing, QoS/QBone, multicast, traffic measurement, performance measurement, web cache and IPv6 services. These network services will require the attention of a 24x7 NOC operations center.

The STAR TAP and Eurolink NOC operations will be fully integrated with the IU NOC that provides 24x7 support for TransPAC and the Internet2 Abilene network. Separate external identities will be maintained. (Q2, Q3)

TransPAC will install a Unix-based network performance and measurements workstation at the STAR TAP in July 1999. This tool will be integrated into NOC operations in Q1 and will serve as a platform for advanced performance and measurement analysis activities.

A collection of NOC tools will be implemented as follows:

Security – TACACS-based access control to network components will provide varied levels of authorities for IU TransPAC NOC and engineering and APAN NOC personnel. (Q1)

Configuration management – a scheduled process will automatically archive network component configuration files. (Q1)

Configuration management – a scheduled process will automatically log into network components, download the configurations, compare the existing and prior configurations and email differences to engineering personnel. (Q1)

Traffic map – a dynamic, web-based graphical map of network utilization and errors will be implemented. The map will have the ability to replay historical data. (Q1).

BGP reporting tool – a scheduled process will collect BGP route information from the TransPAC router and compare against the prior collected data. Differences will be e-mailed to the TransPAC NOC operations and engineering staff. (Q1)

Multicast monitoring – multicast session directory reporting will be hosted on the TransPAC management station to be installed at STAR TAP. (Q1)

Log watcher – system logs of the network components will be automatically watched and reports of significant events will be emailed to NOC operations and engineering. (Q1)

The current trouble ticket system used by the IU NOC for TransPAC and Abilene is based on a commercial system, Apriori. Web-based access to the trouble tickets is not provided outside of the NOC. A trouble ticket system will be developed that provides open (authenticated) access to current and archived tickets via a web interface. (Q2)

User Services

The first year's efforts in user services for TransPAC were focused on identifying key startup users and understanding their applications. To foster this approach IU and UIC sponsored the iGrid demonstration at SC98 and supported the implementation of a number of applications over TransPAC. The iGrid Asia-Pacific community has continued as the core TransPAC user group and it is expected that iGrid will be repeated at INET2000 in Yokohama, Japan.

TransPAC User Services will conduct an annual census of the research and education users of the TransPAC network. We will obtain detailed project information, classify projects according to type and area, offer assistance and solicit project status and feedback regarding the TransPAC network service. The current solicitation letter can be viewed at http://www.transpac.org/applications/9904_followup.html. (Annually)

A database containing this project information has been created and will be updated regularly. (Ongoing)

In the coming year user services will continue to emphasize the implementation and tuning of LFN applications. TransPAC user services will organize a hierarchical support group in cooperation with APAN working groups and relevant personnel in APAN countries. This, along with existing mailing lists and the TransPAC web site, will be the primary sources of information for TransPAC users. Specific application development assistance will also be available through the engineering positions currently being filled at IU. (Q1, Q2)

The most successful applications seem to be those that have good local support at both ends of the network and help from the TransPAC team during startup and tuning. We will continue this collaborative mode of operation for fielding new applications. (Ongoing)

Excellent progress in promoting TransPAC was made this year through conference demonstrations and papers. Efforts aimed at contacting and assessing potential users' needs were started this year and will be continued vigorously in the coming year. Supporting arrangements with other networks connecting to APAN to allow routing into TransPAC will also continue to be developed to facilitate the process of getting end-users online. (Ongoing)

Other support plans include developing an improved, more functional web site with links to development resources, and establishing a regular newsletter-style communications channel to which the user community can subscribe. (Q2)

NLANR training activities have made a valuable contribution to the expertise of the high performance network user community. It is our intention to hold a network application development and tuning workshop for TransPAC users this fiscal year, in cooperation with NLANR and associated with a major international networking or computing conference such as INET2000 or HPDC-9.

The current state of diffserv implementations is far from production quality, and the QoS picture is further confounded by a lack of inter-domain bandwidth brokers and mechanisms. Although TransPAC will continue to participate in QBone diffserv and bandwidth broker testing, for the coming year we expect that manual bandwidth scheduling will be required. To facilitate this we will establish a web-based allocation request form and scheduling system through which big bandwidth users can ask for exceptional resources. (Q1)

Application Collaborations and Demonstrations

As mentioned above, collaborations of many types are needed to develop and field international high performance network applications. One of our continuing goals is to foster these collaborations through the following mechanisms:

- Web-based information about existing collaborations, their projects and results (Q1)
- Direct end-user support for application development (Ongoing)
- Fostering new collaborations through participation in APAN and US networking community activities (Ongoing)
- Sponsoring and supporting demonstration testbeds such as iGrid and GOIN (Q2, Q3, Q4)

International conference demonstrations offer an ideal environment to stress test new applications, both for the application developers and network engineers. We hope to encourage further group demonstration projects as a way to bring developers together to exchange ideas and to put developers and network engineers in direct contact.

Following on the enormous success of the Supercomputing'98 iGrid demonstrations hosted by Indiana University and the Electronic Visualization Laboratory of the University of Illinois at Chicago, an iGrid event will be held at INET'2000 in Yokohama, Japan, July 18-21 2000.

The SC'98 iGrid research demonstrations showcased international collaborations using advanced high-speed networks, enabling researchers to work together and to access geographically-distributed computing, storage, and display resources. The iGrid booth provided global connectivity for collaborators from the United States, Australia, Canada, Germany, Japan, The Netherlands, Russia, Singapore, Switzerland, and Taiwan to demonstrate their use of advanced networks to solve complex computational problems.

The INET'2000 iGrid will be hosted by a consortium of Japanese universities. Initial discussions suggest that the iGrid, international research collaboration on advanced networks, will be one of the INET'2000 themes. TransPAC management, engineering, user services and operations personnel, along with the TransPAC network and STAR TAP will be major contributors to and participants in the event. The iGrid will be one of the major global events during the year 2000 to facilitate and showcase international collaborative advanced network-based applications. See http://www.transpac.org/applications.html for details on iGrid projects. (Q4)

Likewise, following on the success of the Global Observation Information Network '99 (GOIN99) conference and demonstrations held March 21-25 in Honolulu, HI, additional demonstrations will be held at the Committee of Earth Observation Satellites (CEOS) subgroup meetings, September '99 in Virginia and at the CEOS Plenary in November in Stockholm. (Q2,Q3)

One of the intents of the GOIN and CEOS demonstrations is to bring together researchers in earth observation sciences and experts representing advanced networking and to exploit this combined expertise for the benefit of global earth observation research. See http://www.transpac.org/applications.html for details on GOIN and CEOS projects.

The Internet2 Digital Video Network (I2 DVN) is an initiative that is developing enabling capabilities for research and education, primarily to support three modalities of advanced digital video — video conferencing, video-on-demand, and live transmission on national research and education high performance networks. This initiative also involves related activities such as digital video portals, digital media asset management (including digital libraries), infrastructure architecture, digital production studios, and content issues, such as channel design and development. In addition, the I2 DVN initiative is investigating network behavior, such as the effects of transmitting broadcast streams over multicast clouds, emerging media protocols, and implications of integrating various digital video modalities, which generally are implemented as separate architectures and separate technologies. This initiative is being led by the International Center for Advanced Internet research at Northwestern University (iCAIR), but includes many participating organizations. See http://i2dv.nwu.icair.org/ for details.

APAN and TransPAC are participants in the I2 DVN initiative. Specific immediate projects center on implementation of simulated channels from several I2 DVN stations at major universities. Channels can broadcast live events which can be archived, or stored content selected by the viewer. See http://pisces.ctd.anl.gov/IDVn/index.html. The I2 DVN application is the testbed application for TransPAC and APAN's QBone proposal. (Ongoing)

Advanced Application Services Testbeds

A key enabling component for widely distributed applications is a set of middleware services for resource discovery, user authentication and authorization, resource scheduling and execution control. The Globus toolkit provides these middleware services and was used in several iGrid demonstrations at SC98. We plan to expand support for Globus over TransPAC and to work with the Globus development group at ANL to provide support for the development of Globus programs over TransPAC. (Ongoing)

Advanced Network Services Testbeds

Multicast is becoming an important component of collaborative environments and digital video services. Activities to be undertaken in 99-00 to strengthen multicast support by TransPAC include: 1) Moving from PIM-DM to PIM-SM/MSDP across TransPAC and inter-domain connections. This change will require the deployment of more recent Cisco IOS images across TransPAC participants. 2) Establish multicast interconnect with Abilene. Currently, TransPAC and Abilene exchange only unicast traffic. 3) Deploy tools for debugging, monitoring, and measuring multicast traffic; and 4) Coordinate additional native IP multicast interconnectivity in conjunction with other STAR TAP participants. (Q1-Q3)

As described in detail in Section D.1 Progress Report: Attainment of HPIIS Objectives, TransPAC and APAN are participants in the Internet2 Quality of Service Backbone (QBone) project. TransPAC and APAN will continue to be participants in this effort and will be early adopters, but it is too early to

predict a timeline for significant deployment of QoS services. STAR TAP has proposed engineering activities including QBone and Multiprotocol Label Switching (MPLS) research. TransPAC engineers will actively participate with STAR TAP engineers in these activities. (Ongoing)

Although the timeline for implementation of an active QoS mechanism on TransPAC is uncertain at this time, TransPAC currently provides service for applications that require significant network resources, e.g. telemicroscopy, which requires approximately 35Mbps bandwidth for video display.

To provide reliable and equitable use of the network, a web-based resource request and reservation process will be developed. Owners of applications that require very significant network resources will be encouraged to schedule use of the network. The system will provide a reservation form and display of currently scheduled activities. (Q1)

Plans for the coming months include participation in the 6REN/6TAP. In order to facilitate the easy interconnection of 6REN participants, STAR TAP, CANARIE and ESnet are jointly sponsoring an IPv6 Exchange "6TAP" project to provide routing and route serving services at the STAR TAP in Chicago. The 6TAP will provide an IPv6 capable router co-located at STAR TAP to experiment with early route administration and peering services to assist in the development of IPv6 operational procedures. During June, a router supplied by ESnet was installed at STAR TAP to provide a native IPv6 interconnect point. The IPv6 efforts are a prime example of collaboration across projects, agencies and countries. ESnet personnel have been active in the 6bone deployment efforts and will provide routing services and operational management support. CANARIE will provide IPv6 route server and routing registry. (Ongoing)

STAR TAP has proposed working with NLANR to provide and run a web cache at the STAR TAP. TransPAC will actively participate in deployment of this service and will coordinate efforts between APAN and STAR TAP. (Ongoing)

Performance Analysis and Measurement

TransPAC engineer Linda Winkler performed a significant analysis of network performance at the network's original 35Mbps bandwidth. See http://www.transpac.org/engineering/performance/. Network bandwidth has now been increased to 73Mbps and engineering of the network has changed substantially. Another in-depth analysis will be performed. In addition to providing valuable insight into performance of the network and input for subsequent tuning, a goal of this analysis will be to deliver a significant published work that can serve to advance the state-of-the-art in managing performance of advanced research and education networks. (Q2)

ANS Surveyor systems (http://www.advanced.org/csg-ippm/), which provide measurement of one-way delay and packet loss in a community of surveyors, have been installed in Korea at the Korea Advanced Institute of Science and Technology and at the APAN Tokyo network exchange point. The Tokyo installation awaits completion of GPS antenna installation. Additional Surveyors will be installed in Australia and Singapore. (Q2)

The method for displaying reports of the Surveyor measurement is somewhat cumbersome owing to the large number of Surveyor sites and large amount of available data. TransPAC will work with ANS to develop a method to easily and prominently display current graphs from the TransPAC web pages. TransPAC will also work with ANS to implement methods for longitudinal analysis of the Surveyor data. (Q3)

NLANR operates an OC3mon for TransPAC installed at the STAR TAP facility. Data are collected, processed and stored in the NLANR Datacube (http://moat.nlanr.net/Datacube/). TransPAC will work

with NLANR to develop a process for regular analysis and reporting of the TransPAC OC3mon data. Graphical analysis will be reported on the TransPAC web pages. (Q3) Reports will include:

- Major tcp/udp applications by packet, byte and flow counts
- Average and maximum pps
- Average and maximum flows per second
- Cumulative histograms of packet sizes

Utilization statistics of the TransPAC international ATM PVP are currently reported by the PVP aggregate. Per-VC reports will be developed and displayed on the TransPAC web pages. (Q2)

Human Resources

Funded positions for a network engineer and high performance computing and communications engineer have not been filled. The positions have been advertised and candidates interviewed, but suitable candidates at the offered salary have been difficult to find. A promising offer has been made to a network engineer. As TransPAC moves into its second year these positions will become even more critical as participation in advanced services testbeds increases, the requirement for user support for high performance applications increases, the focus sharpens for facilitation of collaboration and applications, and daily operational issues increase.

We will seek to fill the existing position descriptions throughout Q1. If by early Q3 it becomes apparent that the positions cannot be filled, we will work with the NSF Program Official to redefine the positions, responsibilities, oversight, and management resource commitment to better match the positions to market availability.

Web-based Information

Descriptive information regarding TransPAC-based applications and collaborations will be enhanced. Information collected from regular applications follow-up efforts (Section E. Program Plan, User Services) will be posted to the web page. Detailed information regarding significant projects will be posted. (Q1)

Technical information regarding application performance and host tuning for LFN (long fat networks) will be posted. (Q2)

Additional display of performance data, including enhanced views of the TransPAC-specific ANS Surveyor one-way delay and loss statistics and display of analysis of OC3mon data will be posted. (Q2)

A web accessible trouble ticket system and archive will be developed. The system will be devised to incorporate IU NOC responsibilities for TransPAC, STAR TAP and Eurolink NOCs. (Q2)

HPIIS Team

IP-based, desktop video conferencing will serve to strengthen the communication and collaboration among the HPIIS Team. TransPAC will work with University of Tennessee Knoxville MIRnet to recommend appropriate video conferencing technology. (Q1)

Realizing that all HPIIS networks and STAR TAP will benefit from close cooperation, collaboration and sharing of expertise in administrative, policy and technical realms, TransPAC will work to increase the level of these interactions. (Ongoing)

Appendix A — Project Descriptions for Active Original TransPAC Projects

1. Asia-US-Australia Collaboration in the Silicon Vertex Detector Project for the BELLE High Energy Physics Experiment at KEK

Japan: KEK (High Energy Accelerator Research Organization)
Osaka University
Tokyo Metropolitan University
University of Tokyo
University of Tsukuba

Australia: University of Melbourne
University of Sydney

USA: University of Hawaii
Princeton University

Project contacts:

Professor H. Aihara, aihara@phys.s.u-tokyo.ac.jp
Professor D. Marlow, marlow@puphep.princeton.edu

Project web site: http://bsunsrv1.kek.jp/

Project description: The BELLE detector is the state-of-the-art detector to investigate CP violating phenomena with unprecedented precision at the KEK B meson factory.

The CP (C=Charge conjugation, P=Parity) violation is a key to explain why the universe is dominated by the matter, not by the anti-matter. The primary goal of the BELLE detector is to identify the origin of the CP violation. The BELLE collaboration consists of more than 40 institutions from Japan, Korea, China, Taiwan, India, Russia, USA, Australia and Europe.

At the heart of the BELLE detector a high precision particle trajectory detection system consisting of silicon microstrip sensors will be installed. This silicon system contains about 100K channels to be read out by a high-speed, online data-taking system. To achieve required precision all electronics channels must be constantly monitored and calibrated.

An international collaboration was formed to design and build the silicon vertex detector. In addition, the generated data will be jointly analyzed by the participating institutions to obtain physics results in a timely manner.

2. APBionet: Asia-Pacific Bioinformatics Network

Michigan State University (USA) and National Institute Genetics (Japan)

Project contacts:

Prof. Hideaki SUGAWARA
Center for Information Biology
National Institute of Genetics
E-mail: hsugawar@genes.nig.ac.jp

Prof. Herman D. Hughes
Computer Science & Dir. of HSNP
Michigan State University
E-mail: hughes@cps.msu.edu

Prof. J. Tiedje
Center for Microbial Ecology
540 Plant and Soil Sciences Building
Michigan State University
E-mail: tiedjej@pilot.msu.edu

Project web site: http://www.cme.msu.edu/RDP

Project Description: The APBionet was established to realize smooth information flow and sharing free from boundaries caused by physical distance, social system, technology and culture. Broadband networks are indispensable to APBionet that consists of:

- Fast and robust networks of databases and applications
- Facilities for visualization of a large scale data and interactive simulation of life phenomena for the advanced study on life sciences and biotechnology
- Distance learning systems based on multimedia servers to serve bioinformatics experts in the Asia-Pacific region

3. RHIC Pioneering High-Energy Nuclear Interaction Experiment (PHENIX) High-energy Heavy-ion experiment

University of Tokyo and Brookhaven National Laboratory

Project contacts:

Hideki Hamagaki <hamagaki@cns.s.u-tokyo.ac.jp>
Center for Nuclear Study, University of Tokyo (CNS)

Sam Aronson <aronsons@bnl.gov>
Brookhaven National Laboratory, Upton, NY

Project web site: http://www.phenix.bnl.gov/

Project Description: The PHENIX experiment is one of the two major experiments in the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. It aims to detect the signatures of Quark Gluon Plasma (QGP) deconfined phase, which is thought to exist in the early universe. More than 400 physicists, engineers and students from eleven countries are currently participating in the PHENIX experiment. The participating institutions from Japan are: University of Tokyo, CNS-Tokyo, KEK, Tsukuba, Waseda, Kyoto, Hiroshima, Nagasaki Institute for Applied Science, Tokyo Institute of Technology and RIKEN.

4. Web Cache Meta Network

NTT (Japan), KAIST (Korea), NLANR (USA), Keio University, Nihon University and another 30 institutions

Project contacts:

Masaaki Nabeshima, nabe@slab.ntt.co.jp (NTT Labs)
Kilnam Chon, chon@cosmos.kaist.ac.kr (KAIST)
Duane Wessels, wessels@nlanr.net (NLANR)

Project web site: http://cache.jp.apan.net/ and http://cache.kaist.kr.apan.net/

Project description: To ensure the effectiveness of the web cache meta network, cache servers and primary web servers should exchange their status information (e.g. object update information from primary servers, load status information from cache servers). The goal of this project is to make a network to exchange that meta information. Currently we are experimenting with a hierarchical web cache system between continents to get some operational data. NTT labs is developing a metadata exchange system (an alpha stage code is available).

As a step toward a unified view of caching and replication, we are developing RepliCache (Large Object Cache) system. RepliCache will be located at GigaPoPs and it will serve large-bandwidth data (e.g., multimedia data) to users or collaborating caches. In the context of replication, resolution and contents, we will cooperate with the Internet 2 Distributed Storage Infastructure project team.

5. Japan-US Collaboration on ICRF heating and Current Drive Experiments and Modeling

University of Tokyo and Massachusetts Institute of Technology

Project contacts:

Professor Y. Takase, takase@phys.s.u-tokyo.ac.jp
Dr. M. Greenwald, g@psfc.mit.edu

Project web site: <None listed>

Project description: The ICRF heating experiments on the Alcator C-Mod Tokamak at MIT has been highly successful. With the additional power and current drive capabilities currently being added in collaboration with PPPL (Princeton Plasma Physics Laboratory), the emphasis of research will shift toward steady-state advanced Tokamak experiments using current drive and profile control. MIT is already preparing the infrastructure necessary to support remote collaborators.

Participation in these experiments by the University of Tokyo group, including planning of experiments, real-time participation, and data analysis became possible with the increased bandwidth provided by APAN. The University of Tokyo group, in collaboration with MIT physicists, performed theoretical modeling of various heating and current drive scenarios in 1998. The results indicate potential usefulness of the high harmonic fast wave for current drive in both conventional and spherical Tokamaks.

6. Nucleic Acid Database

Japan Agency of Industrial Science and Technology and Rutgers University

Project contacts:

Yoshikuni Okada, Director of RIPS Center
Agency of Industrial Science and Technology
rdire@aist.go.jp

Prof. Helen M. Berman, Director
Department of Chemistry
Rutgers, The State University of New Jersey
berman@dnarna.rutgers.edu

Project web site: ndbserver.rutgers.edu

Project Description: The Nucleic Acid Database (NDB) Project was established to serve as a resource for researchers who study the structure of nucleic acids. It provides a repository for the coordinates of oligonucleotide crystal structures. In addition, the NDB provides information of general interest to researchers in the field, and develops and distributes standard geometric information for use in molecular refinement and modeling programs.

This project is supported financially by the National Science Foundation and the Department of Energy. In Asia, the NDB WWW site is mirrored at the Structural Biology Center at AIST, Japan. For mirroring, the Japan mirror frequently needs to receive updates from the NDB WWW site in the US.

7. Development of the Mirror Server by Using High Speed Data Transfer in Genome Science

Project contacts:

Yoshihiro Ugawa <ugawa@disc.dna.affrc.go.jp>
Akira Mizushima <goddila@maffin.ed.jp>
Don Gilbert <gilbertd@chipmunk.bio.indiana.edu>

Project web site: <None listed>

Project Description: DNA/protein biological sequence database is essential for advanced studies in genome research. These sequence data have been collaboratively collected by the US, Japan and Europe since 1984. Software search engines such as FASTA/BLAST are used to find homology sequences. Huge local disks are required to store these databases.

Some mirror servers have been developed to provide updated data for the DNA/protein biological sequence database. Currently, however, these servers often fail due to the lack of adequate network bandwidth. In this project, we will develop a reliable mirror server with high-speed data transfer in HPIIS.

In genome research the data from DNA databases have increased tremendously; about 10% are Japanese contributions. However existing servers are mainly mirroring US data to Japan. A bi-directional mirror server is needed to update advanced data that are collected by a number of other countries and deliver them to researchers. This mirror server cannot be established without a reliable, high-speed and connection.

8. Applications of Networked Virtual Reality Systems

Project Contacts:

Prof. Kwangyun Wohn (wohn@acm.org)
KAIST (Korea Advanced Institute of Science & Technology)

Tom DeFanti, UIC (tom@uic.edu)
University of Illinois at Chicago

Project web site: <None listed>

Project description: <Not Available>

9. Applications of Networked Virtual Reality Systems

Japan-US Collaboration in Sloan Digital Sky Survey over the Network

Project contacts:

Professor S. Okamura, okamura@astron.s.u-tokyo.ac.jp
Professor A. Szalay, szalay@tardis.pha.jhu.edu

Project web site: http://www.sdss.org

Project description: Sloan Digital Sky Survey (SDSS) is a project to carry out imaging and spectroscopic surveys of half the northern sky using a dedicated, wide-field, 2.5-m telescope. The imaging survey with a large mosaic CCD camera will produce digital photometric maps of the sky in five color bands. These maps will be used to extract the position and various photometric parameters of about 100 million galaxies and close to the same number of stars. Among the extracted objects, about 1 million galaxies and 100 thousand quasars are selected, for which medium resolution spectra will be obtained.

The SDSS is a collaborative project between the US and Japan involving seven US institutions and the Japan Promotion Group (JPG). The observation, i.e., data taking, will be carried out at the Apache Point Observatory, New Mexico, where some online data processing is performed. The bulk of the data reduction will be done at FNAL and the master database will be maintained there. The JPG will also maintain the whole dataset in Japan minus the raw data. The JPG is planning to produce the merged pixel map noted above from the flat-fielded data. The data processing required to construct the merged pixel map involves reference to the raw data and to the flat-fielded data.

Significant scientific analyses of these data will often produce results whose data storage requirements are close to the amount in the input catalog itself. Thus, a network capable of transferring some 10GB a day is of critical importance in promoting active, timely discussion between the JPG and US astronomers across several institutions. Similar imaging data taken with the Subaru Telescope will sometimes be useful for interpreting the results of the SDSS data analysis.

Appendix B — Project Descriptions for New TransPAC Projects

1. Bio-Mirror public service for high-speed access to biosequence data

Project contacts:

Tan Tin Wee <tinwee@pobox.org.sg>
Don Gilbert <gilbertd@chipmunk.bio.indiana.edu>

Project web site: www.bio-mirror.net

Project description: This is a worldwide bioinformatic public service for high-speed access to up-to-date DNA/protein biological sequence databanks. In genome research, these databanks have been being growing tremendously, so much so that distribution of them is hampered by existing Internet speeds. The Bio-Mirror project is devoted to facilitate timely access to important large data sets for this research. High-speed access is provided by Internet2 infrastructure of the Very High Speed Backbone Service(vBNS), Abilene, TransPAC, and the Asia-Pacific Advanced Network (APAN).

Currently available servers:

Japan
[1]ftp://bio-mirror.jp.apan.net/pub/biomirror/

Australia
[2]ftp://bio-mirror.au.apan.net/biomirrors/

Singapore
[3]http://www.bic.nus.edu.sg/bioresources.html
[4]ftp://dbase.bic.nus.edu.sg/biomirrors/

USA
[5]http://bio-mirror.us.apan.net
[6]ftp://iubio:iubio@bio-mirror.us.apan.net

Current data sets:

DNA biosequence data include GenBank, EMBL, DDBJ. Protein biosequence data include SWISS-PROT*, TrEMBL, PIR. Other data include BLOCKS, ENZYME, PROSITE*, REBASE.

Data currently total about 10 Gigabytes in compressed format, and are updated from the primary sources nightly. * Commercial restrictions on SWISS-PROT and PROSITE exist.

These servers are publicly available sites for high-speed access to up-to-date DNA/protein biological sequence databanks. High-speed access between the sites is provided by the network infrastructure developed by Very High Speed Backbone Service (vBNS), TransPAC (Trans-Pacific network), and Asia-Pacific Advanced Network (APAN), and these sites are well connected to national research and education networks within each country.

The DNA/protein biological sequence database is essential for advanced studies in genome research. These sequence data have been collected by the US, Japan, and Europe since 1984.

Software search engines such as FASTA/BLAST are used to find homology sequences. Huge local disks are required to store these databases.

Some mirror servers have been developed to provide updated data for the DNA/protein biological sequence database. Currently, however, these servers often fail due to the lack of adequate network bandwidth. In this project, we will develop a reliable mirror server with high-speed data transfer in HPIIS.

In genome research data from DNA databases have increased tremendously; about 10% are Japanese contributions. However existing servers are mainly mirroring US data to Japan. A bi-directional mirror server is needed to update advanced data collected by a number of countries and deliver them to researchers.

2. Trans-Pacific Telemicroscopy

Project contacts:

Martin Hadida-Hassan <marty@SDSC.EDU> - primary
Mark Ellisman <mark@alex.UCSD.EDU>
Youki Kadobayashi <youki@center.osaka-u.ac.jp>

Project web site: http://www-ncmir.ucsd.edu/CMDA/ and http://www.uhvem.osaka-u.ac.jp/official/news.html

Project description: The National Center for Microscopy and Imaging Research (NCMIR) - an NIH-funded Resource - has been leading the efforts in the field of "Telemicroscopy" as part of an on-going research project started in 1992, the Collaboratory for Microscopic Digital Anatomy (CMDA). The CMDA contains a rich set of software tools that provide for remote operation of an electron microscope, the JEOL4000 IVEM, located at NCMIR in San Diego. Using CMDA tools, remote researchers can interactively steer this specialized microscope to investigate their specimens and collect high-resolution digital images of selected areas of the specimen. Additional CMDA tools allow for processing, analysis and visualization of the acquired 2D and 3D datasets.

NCMIR has inspired similar efforts to be conducted by our collaborators at the Ultra-High Voltage Electron Microscopy (UHVEM) laboratory in Osaka University. The UHVEM laboratory has developed tools similar to the CMDA that provide for remote control of their unique electron microscope - the Hitachi H-3000, one of the most powerful electron microscopes in the world.

During the past 6 months, the NCMIR and UHVEM laboratories have been conducting "Trans-Pacific Telemicroscopy" experiments together. Scientists at the UHVEM lab have remotely operated the IVEM in San Diego, and scientists at NCMIR have remotely operated the UHVEM in Osaka. These experiments involve transmission of both microscope control information and video feedback of the specimen over the vBNS, APAN and TransPAC networks. As an example, we have been successful in transmitting up to 36.5 Mbps of Digital Video traffic from Osaka to San Diego. Results so far have been very encouraging and genuine research in the fields of neuroscience, materials science and high-speed networking has taken place using this system. We hope to continue these efforts in order to improve our tools and make a firm statement about the benefits of high-speed networking across the Pacific Ocean.