Network Research and Internet2

This document describes network research projects using Internet2 facilities, ether recently completed or currently in progress, and several projects being undertaken by Internet2 to enhance the potential for network research capabilities. The focus in this document is on the research projects themselves and not the facilities per se, but it also important to understand those facilities and how research projects utilize them.

The Abilene network and the data collection and storage project known as the Abilene Observatory provide a platform on which packet infrastructures can be studied in a production environment. The MAN LAN exchange point consists of lower layer switching facilities capable of supporting network research projects, and the new HOPI testbed will provide the capability to examine a hybrid of packet and circuit switched architectures. Finally, some equipment is available to research

Network research on these facilities falls naturally into three categories: projects that deploy dedicated equipment in the facilities, projects that use measurement equipment already deployed within the facilities, and those that use (special) features of the existing facilities. Other specialized projects are also possible, and support for research proposals also plays a role in the network research projects.

I. Research Projects

Three research projects now deploy specialized equipment in Abilene router nodes and provide a platform for a variety of different types of network research. These projects are collocation projects under the Abilene Observatory.

A. Collocation Projects

The three collocation projects are PlanetLab, the AMP project, and the Passive Measurement and Analysis Project.

- **PlanetLab**

  PlanetLab is a global overlay network for developing and accessing new network services. The goal is to grow to 1000 geographically distributed nodes, connected by a diverse collection of links. Toward this end, PlanetLab nodes have been located in edge sites, co-location centers, and routing centers (e.g., in each of the router nodes on the Abilene backbone). PlanetLab is designed to support both short-term experiments and long-running services. Currently running services include network weather maps,

network-embedded storage, peer-to-peer networks, routing and multicast overlays, and content distribution networks.

The research lead for PlanetLab is Larry Peterson of Princeton University.

- **The AMP Project**

AMP performs site-to-site active measurements (path, round-trip-time, packet loss and on demand throughput tests) and analyses that enable network researchers and engineers to track problems and changes in HPC performance. The monitoring mesh of more than 150 AMPs enables engineers to quickly identify the location, extent, and duration of network events. AMP is moving into new domains: other national meshes, deeper into campuses, more international sites, and a meta measurement system with redefinable meshes. NLANR/MNA is reimplementing the original AMP software into a package easily installed on a variety of hardware and software bases. The AMP package will be made freely available

Tony McGregor NLANR/MNA, Waikato University, is research lead for the project.

- **Passive Measurement and Analysis**

The objective of PMA is to deliver new insights into the operation, behavior, and health of the Internet for the benefit of network users, operators, and researchers. Passive header trace data provides the means to study workload profiles at a number of strategically located measurement points in high speed environments. PMA is collecting daily packet header trace samples from about two dozen sites within the United States. HPC networks are deployed primarily on the Abilene network or its connectors. PMA and Internet2 enjoy a strong collaboration, which is exemplified through the Indianapolis router instrumentations in 2002 and 2004.

Joerg Micheel, NLANR/MNA, San Diego Supercomputer Center, UCSD, is research lead on the project.

B. Projects Using Abilene Observatory Data

The research projects described in this section use data gathered by measurement devices in the Abilene backbone collected by Abilene engineers. The current datasets include flow data, usage data, routing data, throughput data, latency data, router data, and syslog data. The data sets are made available to all researchers.

The Abilene Observatory welcomes recommendations for additional data sets. Network researchers who need other types of data should make requests to abilene@internet2.edu.

- **WAIL: The Wisconsin Advanced Internet Laboratory**

  WAIL is studying flow sampling and anomaly detection using Abilene flow data.

  Research lead on the project is Paul Barford of the University of Wisconsin. See http://wail.cs.wisc.edu for more information.

- **Network Research Lab at Case Western Reserve**

  The main objective of this project is to assess the presence and incidence of alpha flows in backbone links. A flow is said to be an alpha flow if it sends massive amount of data (an "elephant flow") and moreover it has high end-to-end available bandwidth compared to that of the link where the measurement is taken. The relevance of alpha flows is that they can saturate buffers, thereby leading to packet losses and poor levels of service even in over-provisioned best-effort network.

  Vincenzo Liberatore, Case Western Reserve University, is research lead on the project. http://vorlon.cwru.edu/~vxl11netlab/ for more information.

- **Spatio-Temporal Network Analysis, Boston University**

  Many studies have characterized traffic at the level of individual IP flows, and at the link level. In this project we are interested in the properties of traffic at an intermediate level -- the level of source-destination flows, ie, all traffic flowing from a given origin router to a given destination router. We are specifically interested in traffic properties at small time scales over a long period of time. Correlation with Abilene routing data is used in this project. (i.e. IS-IS weights used for routing).

  Mark Crovella and Eric Kolaczyk are research leads on the project, from Boston University, Department of Computer Science and Department of Mathematics and Statistics

- **MINDS Project, University of Minnesota**

  The overall objective of MINDS research is to develop high performance data mining algorithms and tools that will provide support required to analyze the massive data sets generated by various processes that monitor computing and information systems in order to respond to cyber threats in a timely manner. The main approach we take in analyzing such datasets is anomaly

detection. Anomaly detection has the unique property that it can discover the presence of previously unseen behavior such as new types of cyber attacks. It is also useful from a network management standpoint as it will bring the new network behavior to the administrator's attention. The input to MINDS system is NetFlow V5 data, and a version that uses tcpdump is under development.

Vipin Kumar, University of Minnesota, lead on the project. See http://www.cs.umn.edu/research/minds/ for more information.

- **Computer Network Research Group, University of Massachusetts at Amherst**

The project is studying the temporal-spatial correlations in network traffic on the Abilene network. Statistical techniques used include calculation of correlations, power spectral densities, wavelet analyses, and information theoretic techniques using calculation of entropy and conditional entropy.

Don Towsley, Department of Computer Science, University of Massachusetts at Amherst is project lead and Jing Weng is a project participant. See http://www-net.cs.umass.edu for more information.

- **Internet Tsunami Warning System Project, Carnegie Mellon University**

The Internet Tsunami Warning System project aims to develop a distributed high-speed network monitoring system to automatically detect and react to Internet attacks in their early stages. We use the NetFlow data to evaluate our system.

Dawn Song, Department of Computer Science, Carnegie Mellon University, is project lead. See http://www.ece.cmu.edu/~dawnsong for more information.

- **Niagara Project, University of Wisconsin-Madison and the Oregon Health & Science University.**

This portion of the Niagara project addresses query processing over data streams. A particular interest is accommodating streams that are slightly disordered on the attributes of interest. Flow records are processed to measure the disorder that naturally occurs. For example, the order in which flow records are collected depends on eviction of flows from the router cache and may not exactly coincide with the order in which the last packet of each flow was sent. The data distribution can then be used to generate synthetic flow records with varying degrees of disorder, for testing of alternative query evaluation techniques.

Project Leads are Professor David Maier (OHSU/OGI School of Science & Engineering), Professor David DeWitt (UW-M), and Professor Jeffrey Naughton (UW-M). Project Participants (OHSU) include Kristin Tufte, Jin Li, Peter Tucker,and Vassilis Papadimos

- **Department of Computer Science, University of Pennsylvania**

Impact of Aggregation of Traffic on Routing Performance: The principal focus of our project is to determine how aggregation of traffic affects routing performance. Specifically, we construct traffic matrices from actual traffic traces by aggregation at various levels of granularity (prefixes of length /0, /4, /6 and /8) and explore how they affect routing issues like load balancing, long term/short term stability etc.

Dr. Roch Guerin, University of Pennsylvania; Project Participants: Ashwin Sridharan.

- **Modular Strategies for Internetwork Monitoring, University of Michigan, University of Wisconsin, Boston University**

This project addresses the longstanding and difficult problem of detecting and classifying spatially distributed network anomalies from multiple monitoring sites. To characterize baseline vs. anomalous behavior of the Internet requires deployment of collaborative data collection, anomaly detection and pattern recognition for complex large scale systems. The project combines the forces of leading researchers in three complementary disciplines: (i) networking and data collection; (ii) statistical data analysis and signal processing; (iii) decentralized decision-making. The research goes well beyond the state-of-the art anomaly detection for centrally administered networks. In particular tools and practical data sharing algorithms are being developed for detecting coordinated intrusions, distributed denial of service attacks, and quality-of-service degradations in decentralized networks such as the Internet.

Al Hero, University of Michigan. **Project Participants:** Stephane Lafortune, Demos Teneketzis, and George Michailidis, University of Michigan; Paul Barford and Rob Nowak, University of Wisconsin; Eric Kolaczyk and Mark Crovela, Boston University

- **SWORD project, University Wisconsin Madison and University of Saskatchewan**

The main objective of this project is to characterize a variety of flow workloads using new techniques that, for each workload of interest, accurately capture and give insight into: (1) the variations in data flow arrival rate, and (2) the key

correlations in the workload statistics. The new analysis methods also include anomaly detection.

Derek Eager, Mary Vernon, and Su Zhang

- **Distributive Collaborative Data Streaming for Monitoring High-Speed Networks.**

In recent years, the problem of monitoring and analyzing the aggregate traffic flowing through many high-speed links has emerged as an important and challenging problem in network measurement and management. Monitoring the characteristics of this aggregate traffic is essential for detecting “global” events that are intrinsically distributed through the network. Examples of such events range from global top-traffic sources (global elephants) to incipient worm infections. It is hard to detect such events using traditional per-link monitoring mechanisms since the signal is usually too feeble to be observed locally. Such events may leave indelible signatures in the aggregate traffic, but only through the correlation of traffic among many links can this signature be revealed.

Professor Jun (Jim) Xu, Georgia Tech. **Project Participants:** Abhishek Kumar, Min-Ho Sung, and Qi Zhao, Georgia Tech.

- **Reconfiguration of the logical topology in WDM networks using simulation-based Morkov Decision Processes**

In this study we use average day to day behaviour of a network to predict the future trend of traffic and we use that as part of the process to determine how to reconfigure the virtual topology of the network. We would use the data from Abilene network both for obtaining this average day to day behaviour and as sample paths for our simulations.

Professor Mark Shayman, University of Maryland; **Participants:** Professor Steve Marcus, Professor Richard La, Pedram Fard, Kwangil Lee, Yuneng Xie, University of Maryland

- **Algorithms for Network Capacity Planning and Optimal Routing Based on Time-Varying Traffic Matrices**

We are designing algorithms for network capacity planning and optimal routing based on time-varying traffic matrices. We intend to extend the algorithms and techniques to other areas of network design too.

Professor Vishal Misra, Department of Computer Science, Columbia University; **Participants:** Abhinav Kamra, Department of Computer Science, Columbia University

- **Kent State University Computer Science Department**

Traffic Management and QoS Provisioning in IP Networks: The objective of this work is to investigate the impact of self-similar traffic on the performance of output buffers in switches and routers. It is a known fact that the superposition of independent alternating renewal processes (flows) can show self-similar characteristics. Since analytical and empirical studies have shown that self-similar traffic can have a detrimental impact on the QoS, finding an effective buffer management algorithm that can manage self-similar traffic has become an important problem in traffic engineering. Optimal resource allocation is directly affected by optimal buffer size and buffer management policy, bandwidth assignment and traffic management. In this project we study the effect of self-similar and bursty traffic on the triggered threshold buffer management algorithms. Besides the second-order self-similar traffic, we are investigating the effects of fractional Brownian motion on active queue management schemes.

Hassan Peyravi, Kent State University, is project lead.

- **Intelligent Methods for Computer Network Management, Wuhan University, China**

The main objective of this project is to consider new intelligent methods for computer network management. We emphasize on two aspects: one is the analysis of the network measurement information, the network performance parameters are measured and analysized by different knowledge discovery methods so that the potential network performance can be diagnosed and predicted; the other is the fault resilient routing configuration, new applications (esp. e-commerce) need high reliable network and we try to make the network more reliable through dynamic routing configuration according to the traffic matrix.

Project Lead is Professor Yan Pu-liu and participants include Zhou jian-guo, Wu Jing, Chen Xiao, Chen li-jia of the Electronic Information College, Wuhan University, China.

- **The interaction between intradomain and interdomain routing, University of California at San Diego, University of Cambridge, and AT&T**

This project studies the interaction between intradomain and interdomain routing. In particular, it analyzes the impact of using hot-potato routing for selecting egress points. Data from an operational network is used to evaluate alternative egress selection policies under realistic scenarios.

Project leads include Renata Teixeira, University of California at San Diego; Jennifer Rexford, AT&T; Tim Griffin, University of Cambridge; and Geoffrey

M. Voelker, University of California at San Diego.

- **MIND: Multidimensional Indices for Network Diagnosis**

Detecting and unraveling incipient coordinated attacks on Internet resources requires a distributed network monitoring infrastructure. Such an infrastructure will have two logically distinct elements: distributed monitors that continuously collect packet and flow-level information, and a distributed query system that allows network operators to efficiently and rapidly access this information. The design of MIND encorporates a distributed indexing system that supports the creation of multiple distributed indices using proximal hashing to scalably respond to range queries.

Project leads and participants include Dr. Ramesh Govindan, University of Southern California, Dr. Christophe Diot, Intel Research at Cambridge, Dr. Wei Hong, Intel Research at Berkeley, Dr. Gianluca Iannaccone, Intel Research at Berkeley, Xin Li, Ph.D. candidate at University of Southern California, Fang Bian, Ph.D. student at University of Southern California, Hui Zhang, Ph.D. candidate at University of Southern California. See http://enl.usc.edu/projects/p2pMonitoring/index.html for more information.

C. Special Projects

In addition to providing data, Internet2 often works with researchers on special requests that are evaluated on a case-by-case basis. Examples are included below. Researchers requiring configuration changes to the network or other special facilities should contact Internet2 for further information. Not all requests can be accepted, but consideration will be given to strong proposals

Three recent projects are described below.

1. Sizing Router Buffers

In support of Sizing Router Buffers project of Stanford University's High Performance Networking Group, Internet2 staff have reduced the buffer size available to the most utilized backbone link, and used existing latency measurements to verify if this causes any contention on that link. This provides some empirical validation of the theories being developed, and some preliminary results were presented at SIGCOMM 2004 by Guido Appenzeller.

Project leads include Guido Appenzeller, Isaac Keslassy and Professor Nick McKeown. See http://yuba.stanford.edu/~appenz/pubs/sigcomm-extended.pdf and http://yuba.stanford.edu/~appenz/pubs/SIGCOMM04.ppt for more information.

2. IGP Routing Data

A Japanese researcher requested IGP routing updates from the Abilene backbone as part of his research project. Working cooperatively with Internet2 staff, the researcher developed a modified version of Zebra to collect IGP updates from each router. That development lead to the current IGP and EGP routing data available from the Abilene Observatory. See http://abilene.internet2.edu/observatory/data-collections.html#routing for more information.

3. Bandwidth Estimation

In support of CAIDA's Bandwidth Estimation project, Internet2 staff provided direct SNMP access to Abilene core routers to allow utilization to be measured at fine grain. In addition, staff ran a set of programs at CAIDA's request betwee two servers collocated at the router nodes to test some of these programs on a real network where the end hosts were 1 gigabit Ethernet connected with 9000-byte MTU ("jumbo frame") clean paths. See http://www.caida.org/projects/bwest/ for further information.

II. Research Grants

Internet2 currently participates in four research grants associated with network research. In three of them the primary role of Internet2 is to support the research grants and participants: The 100x100 project from the NSF, the Ultralight project organized by the high-energy physics community from the NSF, and the Distributed Security Data project under the Department of Homeland Security. In one we are working closely with the University of Virginia to investigate how Internet2 can better support research: the Network Research Facilities project supported by NSF. There are other projects and grants that Internet2 participates in, but these are the projects that involve network research.

A. The 100x100 project

The 100x100 project is driven by the network research community in the United States. The goal is to examine the requirements, both technological and economic, to bring 100 megabits per second connectivity to 100 million homes in the US. The project is funded by the National Science Foundation.

B. The Ultralight project

The Ultralight project is organized by the high-energy physics community to support the massive data needs that will be required when the Haydron Collider

comes on line at CERN. It is a project to examine new and innovative ideas for data transfer on a large scale.

C. The Distributed Security Data project

The Network Data project is a project to examine the security aspects of traffic on packet networks. The proposal will attempt to identify security issues and traffic patterns associated with threats on the network. The Abilene observatory will be a fundamental resource for the project.

D. The Network Research Facilities project

The Network Facilities project is an exploratory program to determine how the needs of the network research community can be enhanced by Internet2 facilities. Internet2 staff are working cooperatively with Jorg Liebeherr at the University of Virginia to reach out to broad sets of network researchers, inform them of Internet2 facilities, and solicit additional services that should be provided. The results will be published at the end of the project. The project is funded by the National Science Foundation under grant SCI-0441149.

III. Internet2 Facilities

This section describes current and planned Internet2 facilities that can be used by the network research community.

A. The HOPI Project

The Hybrid Optical and Packet Infrastructure (HOPI) project is dedicated to looking at new and innovative ideas for future network architectures. A design team consisting of network engineers and researchers from the Internet2 community has been assembled to create a testbed to examine new architectures for next generation networks. The design team has written a white paper describing a testbed that will examine circuit and packet switched infrastructures including elements of dynamic provisioning. The white paper is currently available for comment. A corporate advisory team has also been formed to support the project, a team of corporate engineering leaders having experience in optical switching and dynamic provisioning.

B. The MAN LAN Optical Facility

The MAN LAN Optical Facility is a layer 1 exchange point located in New York City that participates in the Global Lambda Integration Facility (GLIF). The GLIF is an international collaboration to bring layer 1 capabilities to the research and education community on a world wide basis. The MAN LAN facility provides both production and experimental capabilities including the ability to dynamically switch circuits. As part of the MAN LAN exchange point, it provides connectivity between international research and education networks, including federal networks supporting high performance scientific applications. It also provides a resource for the HOPI project.

C. The Abilene Observatory

The Abilene Observatory is a project to support network research on the Abilene network. It provides the capability for researchers to collocate equipment on the Abilene backbone, and it provides a large database of network attributes. Many different types of data are collected on the network and made available to the research community through easily accessible web interfaces. Over 25 research projects now use the observatory for experimental projects, including studies on the theory of flows, routing algorithms, and identification of traffic patterns.

D. Dedicated Abilene Connectivity

Internet2 is developing the ability to give researchers an interface to create limited-term MPLS tunnels to provide controlled dedicated connectivity between points within Abilene. This could be used, for example, to create an overlay among machines installed as part of the Abilene Observatory. [add more detail?]

E. The Network Research Facilitation Project

The Network Research Facilitation Project (NRFP) is a project to support the network research community by supplying equipment to researchers that is no longer used in Internet2 infrastructure. For example, the recent Abilene upgrade that included installation of new routers, made available eleven Cisco GSR routers for use by the research community. These routers are distributed on a yearly basis through a proposal process organized by Internet2. Additional equipment will be added to the project it becomes available.