Abstracts of INET '93

Editors: *B. Leiner*
L. Todd

International Networking Conference
San Francisco, California, USA
17-20 August 1993

The Annual Conference of the Internet Society

To order copies, contact the Internet Society at:
Internet Society
1895 Preston White Drive, Suite 100
Reston, VA 22091
USA
+1 703 620-8990
isoc@isoc.org

From the Conference Chair

The historical significance of INET□’93 lies in its focus on networking technology and how human beings use this technology to solve problems, work together and improve their lives, in a year when governments and citizens all over the world are suddenly gaining awareness, familiarity and hopes about the promise of communication infrastructures. The geographical significance of INET□’93 is that it takes place in San Francisco, just a few miles north of Silicon Valley, the home of individuals, universities, government and commercial organizations that all had a predominant role in the invention and the continued growth of data networking. For these contextual reasons, as well as for the intrinsic quality of the INET□’93 program, for the breadth of its participation and for the talent of its program committee, I hope you will find this event enriching and stimulating. These proceedings reflect the work of hundreds of individuals around the world and the leadership of Program Co-Chairs Barry Leiner and William Yundt. I am proud to acknowledge their work. INET□’93 itself reflects the dedication of an enthusiastic conference committee and the financial contributions of dozens of organizations, who espouse the ideals of the Internet community.

Eric A. Benhamou
INET□’93 General Conference Chair
3Com Corporation

From Internet Society Vice President for Conferences

INET□’93, the annual conference of the Internet Society, is the first global networking conference to take place since the existence and availability of networks and their services have become known to the general public. When the first forerunner to INET□’93 was held in London eleven years ago, there were only a handful of attendees from Europe and the United States and very little popular understanding of this new communications medium. All of this has now changed; throughout the world there is a new awareness of the potential societal benefits of network-based communication.

The goal of the INET conferences is to bring together network users, developers, planners, researchers, maintainers and sponsors from all corners of the world. This year’s “Training Workshop for Developing Countries,” led by George Sadowsky and held the week before INET□’93 , provided a first opportunity for many to learn about the practical aspects of network technology and services. The 120 INET□’93 attendees, from over 70 countries, who participated in the Workshop, provide a unique opportunity for all INET□’93 participants to engage in interpersonal networking on a global scale.

The richness and diversity of INET□’93 are due, in no small measure, to the vision and drive of its Conference Chair, Eric Benhamou. Working under his direction, the Program Co-Chairs, Barry Leiner and Bill Yundt have produced an exceptional program. Much of the credit is also due Liz Barnhart and Mike Roberts whose financial and organizational expertise provided coherence to the many conference-related activities.

We welcome you to INET□’93 and hope that you will enjoy the program, the people and the look into the future that you will encounter.

Lawrence H. Landweber
Vice President for Conferences
Internet Society

From the President of the Internet Society

It is a great pleasure to add my welcome to INET□’93, the annual conference of the Internet Society. As you have already noticed, the general chairman, Eric Benhamou, has organized and led an outstanding team this year. Program Co-Chairs Barry Leiner and William Yundt presided over the assembly of a timely and highly relevant collection of sessions. The Society is proud to host a special workshop, organized by George Sadowsky, for attendees from technologically emerging countries. We hope this will be a permanent part of all future INET conferences. The Society is most grateful to the generous sponsors of INET□’93 whose names are shown elsewhere in this proceedings. Without them, the conference would not have been able to offer the rich variety of activities available to you.

1993 has been a year filled with drama in the telecommunications sector, in the business world, on the technological side, and in government. In the U.S., certainly, the National Information Infrastructure initiative has captured the attention of the scientific and technical communities, as well as a wide range of government agencies, service, equipment and software providers, to say nothing of the multitude of new users in education, business, government and even in the residential community(!).

The Internet has expanded very rapidly outside the U.S. The most recent statistics suggest that 40% of the system is situated beyond U.S. borders. The applications show remarkable diversity, and an accelerating trend towards business use. Major efforts to organize backbone facilities in Europe have borne fruit, and similar work is underway in Latin America, Africa, and the countries of Eastern Europe and the former Soviet Union. Everywhere there is evidence of interest in provision of commercially useful services - especially in Europe where there appears to be a virtual revolution in Internet service provision taking place.

We are entering a period filled with creative ferment, with extraordinary potential for restructuring of the telecommunications industry and for unprecedented expansion of the Internet and related services into all sectors of the global economy.

In a microcosm of global networking interests, Silicon Valley is involved in a widespread and energetic exploration of the potential of information infrastructure in its Smart Valley project. John Young, former CEO of Hewlett-Packard, is a major force behind this initiative, as is Eric Benhamou. Along with hundreds, if not thousands of others, the Smart Valley participants are determined to create a sample of the 21st century economy well ahead of schedule. I expect we will hear a bit more about this in John Young’s keynote address.

Your participation in INET□’93 is only one of many ways for you to take an active part in the Internet Society’s programs. As you make use of the Internet on a daily basis, take a moment to consider how you use it, and share this knowledge with others. Send a summary to the Internet Society News editor; help a friend get onto the Internet; go to an IETF meeting if you have technical contributions to make; and volunteer to help with INET□’94 in Prague!

Vinton G. Cerf
President
Internet Society

From the Program Committee Co-Chairs

Continuing the precedent of INET '92, the program committee decided to set up six parallel tracks, allowing expansion of sessions exploring internet technologies while retaining the traditional international character of the INET conferences. The INET□'93 program thus has parallel tracks in Network Technology, Network Engineering, Application Technology, User Applications, Policy Issues, and Regional Issues.

Each of these tracks was in essence a mini-conference, ably organized by the track leaders (respectively Christian Huitema & Phill Gross, Elise Gerich & Francois Fluckiger, Eric Hood & John Hine, Laura Breeden & Rolf Nordhagen, Haruhisa Ishida & Steven Goldstein, and Richard Mandelbaum & Tomaz Kalin) with assistance from the program committees, shown in Appendix C. We are most grateful for their efforts in pulling together the program.

We set a goal to increase the number of refereed papers -- and succeeded. In addition to the many fine invited papers, refereed papers were submitted from around the globe to all of the tracks as follows (note: all numbers are for papers, not panelists, and are as of the date these proceedings went to press: June 28, 1993).

Track	*Submitted*	*Accepted*	*Invited*
Network Technology	7	4	5
Network Engineering	14	9	7
Application Tech.	18	6	14
User Applications	14	8	5
Policy Issues	8	4	2
Regional Issues	29	17	8

Another goal we had was that the conference maintain its international character (i.e. that the program not be dominated by USA speakers.) We also were successful at achieving this goal. Contributions were received from 48 countries around the globe, and the program (as of June 28, 1993) includes papers from 30 countries.

	All	*USA*
Submitted	90	28
Accepted	48	11
Invited	41	27

In addition to the main conference, a Training Workship for Developing Countries was organized by George Sadowsky with Steve Fram, Randy Bush and Art St. George as leaders of the Basic and Advanced Technical Tracks and the Network Services Track, respectively. Major sponsors of the workshop are the Internet Society and Novell, Inc., and it is being hosted by Stanford University on the Stanford campus. The workshop will be held the week before the conference and will feature hands-on network implementation, configuration/managment, and use instruction in four laboratories with equipment and software donated by IBM, Novell, Stanford and 3Com. An intensive curriculum has been planned and will be supported by the workshop track leaders, workshop committee and volunteers (see Appendix D).

To assure the Proceedings would be available for participants at the conference, it was essential that the speakers respond quickly once papers were selected. We used network means almost exclusively for the entire process, from original submittals of intents through paper reviews to the final submittal of camera-ready copy. Despite the anticipated problems and the difficulties with obtaining consistent formatting, we found the network process to be highly effective.

We would like to express our appreciation to all those involved in pulling together this program: the conference organizers, regional coordinators, track leaders, program committee, authors, panelists, keynote speakers, workshop organizers, and the many volunteers. We would particularly like to express our appreciation to Ms. Leslie Todd for her efforts in working with us and the authors in coordinating the program and producing a high-quality proceedings and book of abstracts.

Barry M. Leiner
INET□'93 Program Co-Chair
Universities Space Research Association

William Yundt
INET□'93 Program Co-Chair
Stanford University

Abstracts

Session BA: Security

chair: Al Hoover (ANS, USA, hoover@nis.ans.net)

BAB: Understanding the Internet Certification System

Stephen Kent (Bolt Beranek and Newman, Inc., USA, kent@bbn.com)

The Internet is establishing a public-key certification system. This system was established initially to support Internet Privacy Enhanced Mail (PEM), but it is intended to provide a certification infrastructure that can be exploited by a wide range of other, security-relevant protocols and applications. This paper begins with a brief overview of public-key certificates and reviews highlights of the Internet certification system. It then examines the rationale behind the decisions that underlie the system design, especially with regard to the use of X.500 names in certificates and constraints on the structure of the certification system.

BAC: Piloting Authentication and Security Services in the PASSWORD Project

Peter Kirstein (University College London, United Kingdom, p.kirstein@cs.ucl.ac.uk)
Peter Williams (University College London, United Kingdom, p.williams@cs.ucl.ac.uk)

In the beginning of 1992, under the auspices of the CEC VALUE Programme, a project was set up to pilot specific security services using OSI protocols. Three national consortia formed the PASSWORD project: Cambridge University, UCL and NeXor (UK); INRIA and E3X (France); GMD and DANET (Germany). The aim of the PASSWORD project is to pilot a number of OSI applications services with security enhancements. The secured services to be provided are the following: X.400 (88) Message services, X.500 Directory services, ODA Document services carried over X.400 (84) Message services, applications over Connection-oriented services, and interworking with the Internet Privacy Enhanced Mail.

The paper describes the current status of the project, and highlights some specific technical issues. The current status is that the Security Requirements and Security Policies for the applications have been defined. With regard to deployment, it is recognised that the first generation of certification services will deviate from the norms of good quality certification. This must be permitted, and promoted, in order that sufficient interest is generated in the benefits of the technology. To facilitate simple bootstrap of the security pilots, the ``push'' model of authentication is used. Any user initiates security services by supplying a complete forward certification path to the recipient, through to the National Root CA. The three National Root CAs have exchanged cross-certificates. A receiving Security Agent is expected to determine the cross-certificates required to bridge international interworking; these are appended to the forward certification path as necessary. The CAs are able to process correctly the certificates from the other consortia, and from US CAs developed for the Internet Privacy Enhanced Maiil (prm).

In the areas of X.400 implementations, UCL/NeXor uses PP; this includes a Message Store for which SOa end-to-end security service have been implemented. GMD/DANET bases their work on the DANET OSITEL/400 and INRIA/E3X on their UCOM.X 400-88 implementation. All three of the implementations are in regular service in their unsecured versions. The secured implementations have been shown to interwork for Content Integrity and Message Origin Authentication; two of the implementations interwork with content confidentiality, and the third should interwork there also soon. The British and Germans have also implemented the Internet PEM, and have been shown to interwork with two American implementations.

For X.500 Directory Services, the German and British consortia use QUIPU, which are now part of ISODE; the French Consortium uses UCOM.X 500, a commercialisation of PIZARRO. Here all three implementations provide secured DAP binds, and also signed DSA operations and results; only referral, not Chaining, is supported. Two other applications a secured ODA and a secured Connection Service have been implemented by the UK consortium. The pilot services will use the current research network infrastructure. This includes the various National research network services; internationally it uses both the EBONE and EUROPANET networks. The certificates for the authentication are being distributed using the X.500 Directory Service provided by the National research networks coordinated under the CEC VALUE PARADISE pilot. It is expected that serious piloting will start in July 1993.

Session BB: The Next IP

chair: Phillip Gross (Advanced Network Services, Inc., USA, pgross@nis.ans.net)

BBA: CIDR and the Evolution of the Internet Protocol

Peter Ford (Los Alamos National Laboratory, USA, peter@goshawk.lanl.gov)
Hans-Werner Braun (San Diego Supercomputer Center, USA, hwb@sdsc.edu)
Yakov Rekhter (T.J. Watson Research Center, USA, yakov@watson.ibm.com)

The tremendous growth of the Internet has resulted in several efforts addressing the need to be able to address and route the global public Internet. Many of these efforts call for large changes in the deployed Internet infrastructure. We believe a prudent course of action includes extending the usable life of IPv4 while evaluating the use of these new technologies. This requires incremental changes in the IPv4 routing technology base, but more importantly it will require additional coordination and management among the Internet service providers. We also believe that extending the life of IPv4 includes getting higher utilization of the IPv4 address space. Several strategies, mostly administrative in nature, are proposed to conserve the use of IP address space.

Session BC: Advances in Routing Protocols

chair: Christian Huitema (INRIA, France, christian.huitema@sophia.inria.fr)

BCA: Routing in Large Internets

Deborah Estrin (USC, USA, estrin@usc.edu)

This paper focuses on the scaling problems of routing in global scale networks, including the service requirements introduced by new multimedia (e.g., real time) applications.

To address the critical questions of scaling we begin first with an architecture for inter-domain routing that is currently under development and is a strong contender for use in the future Internet. Our architecture unifies hop by hop and source routing to support BOTH scalability and heterogeneity. The hop-by-hop component (Inter-Domain Routing Protocol(IDRP)) precomputes and installs routes that are used by many sources. The source routing component (Source Demand Routing Protocol (SDRP)) computes and installs specialized routes that are not shared by many sources, or that are needed sporadically. SDRP is demand-driven; routes are installed in response to traffic requirements.

SDRP, in combination with new information distribution and route selection mechanisms, can also be used to provide new services that are needed to support real-time traffic in future internets. We evaluate the applicability of adaptive alternate path routing in which the source selects from among several precomputed routes to carry the traffic generated by a particular session. This set of routes includes both shortest and alternate paths. To support load-based routing, dynamic load-based updates must be distributed to some subset of network nodes to be used by sources when making their routing decisions. The use of dynamic information enables the selection of routes capable of meeting the service requirements of applications. We evaluate mechanisms for distributing/obtaining load based information in a scalable manner.

BCB: Multicast Reservation Protocols

Lixia Zhang (Xerox PARC, USA, lixia@parc.xerox.com)

This talk will present the design of a resource reservation protocol, RSVP, that supports both unicast and multicast applications, even when the number of sources and/or receivers is very large. Several novel features in the RSVP design lead to the unique flexibility and scalability of the protocol. RSVP is receiver-oriented, in that the receiver of the data flow is responsible for the initiation of the resource reservation; this allows RSVP to accommodate heterogeneous receivers in a multicast group. RSVP provides several reservation styles that allow applications to specify how reservations for the same multicast group should be aggregated at the intermediate switches; this results in more efficient utilization of network resources. Finally, RSVP uses "soft-state" in the switches, enabling it to gracefully support dynamic membership changes and automatically adapt to routing changes.

BCC: Multicast Routing Extensions for OSPF

John Moy (Proteon, Inc., USA, jmoy@proteon.com)

The OSPF protocol has been extended to handle the routing of IP multicast datagrams, resulting in a protocol called MOSPF. This potentially has applications in many areas, such as voice and video conferencing, distributed databases, games and simulations. This paper describes MOSPF features, including a characterization of its computed multicast datagram paths, the ability to mix multicast and non-multicast routers in an Autonomous System, and the support provided for forwarding IP multicasts between OSPF areas and Autonomous Systems.

A description of MOSPF protocol mechanisms is provided, and experiences to date with the MOSPF protocol are documented. An analysis of expected protocol performance is also presented. Possible applications within the existing MBONE are explored. Finally, a list of ongoing work relating to MOSPF is given.

Session BE: Mobility

chair: John Ioannidis (Columbia University, USA, ji@cs.columbia.edu)

BEA: Enhancement of VIP and Its Evaluation

Keisuke Uehara (The University of Electro-Communications, Japan, kei@cs.uec.ac.jp)
Jun Murai (KEIO University, Japan, jun@wide.sfc.keio.ac.jp)
Hideki Sunahara (The University of Electro-Communications, Japan, suna@cs.uec.ac.jp)
Fumio Teraoka (Sony Computer Science Laboratory; Japan; tera@csl.sony.co.jp)

``Mobile Computing'' is a buzz word in the area of computer communication nowadays. An important technology in mobile computing is how to communicate with mobile hosts. We introduced Virtual Internet Protocol (VIP) to achieve host migration transparency. Although VIP works correctly in a pure VIP environment, there are two major problems when VIP is used in existing IP networks: 1) backward compatibility with IP and 2) remaining obsolete cache entries. This paper enhances VIP to solve these problems. In enhanced VIP, the VIP header format is redesigned so that enhanced VIP is fully compatible with IP. To solve the problem of obsolete cache entries, ``temporary proxies'' and ``two-phase cache deletion'' are employed. Enhanced VIP is running on SONY NEWS workstations and PC/AT clones. Its overheads have been slightly improved though several functions such as two-phase cache deletion have been added.

Session BF: Network Support of Multimedia

chair: Bernhard Plattner (ETH, Switzerland, plattner@komsys.tik.ethz.ch)

BFA: Partial Order Transport Service for Multimedia Applications: Unreliable Service

Paul Amer (University of Delaware, USA, amer@droopy.laas.fr)
Christophe Chassot (Laboratoire d'Automatique et d'Analyse des Systemes, France,)
Thomas Connolly (University of Delaware, USA,)
Michel Diaz (Laboratoire d'Automatique et d'Analyse des Systemes; France; diaz@droopy.laas.fr)

This paper proposes a Partial Order Connection (POC) protocol. Motivated in particular by multimedia applications, POC is an end-to-end connection that provides a partial order service, that is, a service that requires some, but not all objects to be received in the order transmitted. [ACCD93] discusses a reliable version of POC which requires all transmitted objects to be eventually delivered. This paper presents an unreliable version which permits the service to lose a predefined subset of the objects. A metric of a partial order's complexity, $e_i(P)$, the number of linear extensions of partial order P in the presence of i lost objects, is proposed to quantify different combinations of partial order and reliability. Means for its calculation is presented when P can be modeled as a combination of sequential and/or parallel compositions of Petri-nets. This metric allows one to

compare and evaluate the complexity of different partial order services. A formal specification of POC, written in Estelle, indicates how a destination dynamically determines if arriving objects satisfy a given partial order (without or with losses).

BFB: Multimedia Application Requirements for Multicast Communications Services

Jon Crowcroft (University College London, United Kingdom, j.crowcroft@cs.ucl.ac.uk)
S Chuang (University College London, United Kingdom,)
S Hailes (University College London, United Kingdom,)
M Handley (University College London, United Kingdom, m.handley@cs.ucl.ac.uk)

Multimedia Conferencing can be seen as a new way of distributing communication between several applications (rather than a single monolithic application) that often involves more than two end-systems. This can take advantage of multidestination communications services provided by multicast system based on host groups and the enhancement of routers to form special purpose forwarding paths to these groups. This is mainly so that the number of copies of packets sent over mesh topology Wide Area Networks (X.25, IP, or CLNP, as well as SMDS and ATM based WANs) to group destinations may be reduced significantly over a simple replication at source, and so that the service interface is similar to that provide by multicast on networks with a physical broadcast technology (e.g. Satellite and LANs)

In this paper, we examine the snares and delusions that must be avoided when 4 different kinds of multicast application uses a multicast network service, namely: Shared Windows, Packet Digital Audio, Packet Digital Video and Shared Applications.

This is based on the experience of a real (non-toy system operating over international ISDN links in Europe, and the participation in the (configuration and us of) IETF Internet Audio and Video Casts, to over 200 sites worldwide.

BFC: A Run-time Environment for Da CaPo

M. Vogt (Swiss Federal Institute of Technology Zuerich, Switzerland, vogt@komsys.tik.ethz.ch)
Thomas Plagemann (Swiss Federal Institute of Technology Zuerich, Switzerland, plagemann@komsys.tik.ethz.ch)
Bernhard Plattner (Swiss Federal Institute of Technology Zuerich; Switzerland; plattner@komsys.tik.ethz.ch)
Thomas Walter (Swiss Federal Institute of Technology Zuerich; Switzerland; walter@komsys.tik.ethz.ch)

It is a well known fact that the performance bottleneck of modern high speed networks is located in the end systems. Reasons are the complexity of the applied protocols and the embedding of protocols in the end systems. In addition, a variety of new application with requirements different from the ones traditionally offered by networks emerges. It can be claimed that a single fixed protocol stack is not an optimal solution to equally satisfy all applications with their varying requirements.

Da CaPo (Dynamic Configuration of Protocols) provides an environment that allows the dynamic configuration of protocols with respect to these application requirements and the properties of the offered network services. The goal of the configuration is to reduce the protocol complexity and to increase the performance. The main idea is to subdivide the communication subsystem into three layers. The lowest layer, the transport infrastructure, connects end-systems. The offered service might be the service of a MAC layer, an ATM adaptation layer or the service of a transport protocol like TCP or XTP. This T service generally has to be enhanced within the C layer to allow applications (located in the A layer) to cooperate.

Every C service, offered at the AC interface, is decomposed into a set of inter-related protocol functions. These protocol functions are connected by a protocol graph to account for the dependencies between them. Protocol functions might use different protocol mechanisms, which in turn are implemented in either hardware or software. An implementation of a protocol mechanism is termed module. Each module has a set of characterising properties, e.g. throughput, cost, etc. These properties are used to calculate an optimal protocol.

A unified module interface allows the unrestricted configuration of modules to protocols. The run-time environment of Da CaPo links modules to protocols in one UNIX process. It realizes an efficient data transport

inside the end system because operations reducing performance like data copying or UNIX context switches are minimized. The run-time environment awaits data either from the application or the T service. It then guides the data packet from one module to the next until the data packet reaches its local destination (T service or application) or is blocked by a module, e.g. to enforce a rate control. In both cases the process continues to fetch a new data packet from the application or the T service. The application is also integrated into this process as it uses the same interface as the modules. This way the application can directly benefit from the efficient buffer management provided by the run-time environment.

An additional advantage of Da CaPo is the possibility to compare the performance of single protocol mechanisms in the same environment.

Session CA: Practical Considerations for Routing and Addressing

chair: Elise Gerich (Merit Network, Inc., USA, epg@merit.edu)

CAA: Technical Implementation of the RIPE Route Server

Tony Bates (RIPE Network Coordination Centre (NCC), Netherlands, t.bates@ripe.net)

This paper outlines both the model and technical aspects in the implementation and deployment of the European Route Server. Some history and background is given as to why the Route Server is needed and in what context the European Route Server project came about. As an integral part of the Route Server a brief introduction is given of the general routing policy representation language used not only within the Route Server project but also within the European TCP/IP networking community.

This work was done as part of the joint project "Implementation of a Route Server for policy based routing across the GIX" done in conjunction with the RARE technical programme and RIPE. The main funding for the project has been provided by SURFnet through RARE.

CAB: Practical Considerations for Network Addressing using a CIDR Block Allocation

Haavard Eidnes (University of Trondheim, Norway, havard.eidnes@runit.sintef.no)

The strategies for IP address allocation are about to change. The reason for this is the increasing growth the Internet has seen lately. Some short-term solutions to the strains this growth is causing have been proposed, such as CIDR and C-block address allocations. This paper will try to explain these changes, and how local network managers needing new IP addresses should plan address assignment for the benefit of all.

Author information:

Havard Eidnes got his M.Sc. in Computer Science from the Norwegian Institute of Technology at the University of Trondheim in 1986. Since 1988 he has worked part time for the Norwegian academic and research network Uninett as responsible for the IP network service.

CAC: Autonomous System Partitioning and Policy Routing in the Japanese Internet

Kazuhiko Yamamoto (Kyushu University, Japan, kazu@csce.kyushu-u.ac.jp)
Noritoshi Demizu (OMRON /Corporation, Japan, demizu@nff.ncl.omron.co.jp)
Masaki Hirabaru (University of Tokyo, Japan, hi@nc.u-tokyo.ac.jp)
Akira Kato (Keio University, Japan, kato@wide.ad.jp)

As the number of connected networks increase and multiple service providers exist in a single AS (Autonomous System), the networks begin to face routing problems. This paper presents a case study on AS partitioning to solve the routing problems with the Japanese Internet.

The Japanese Internet was established in 1987, and has steadily grown to where it now includes more than 200 connected networks. The number of service providers has increased and their policies have been diversified. The growth of the Japanese Internet and the traditional routing technology, RIP(Routing Information Protocol),

produce routing problems including dependency between service providers, limitation of RIP metric, overhead in terms of bandwidth, difficulty of path design, and inadequate metric.

To solve these problems, JEPG(Japan Engineering and Planning Group) and the authors are planning to divide the Japanese Internet into multiple ASes and combine current reliable routing technology, namely BGP(Border Gateway Protocol), and appropriate intra-domain routing protocols, such as OSPF(Open Shortest Path First).

Moreover, this paper describes policy routing in the Japanese Internet. It is shown that policy that gives preference to links can be express by dividing the ASes which hold less preferred links into multiple ASes. Some policies can't be expressed due to the limitations of packet forwarding technology and AS technology.

Session CB: Capacity Forecasting and Planning

chair: Toshiya Asaba (Internet Initiative Japan, Inc. Japan, asaba@iij.ad.jp)

CBA: Long Term Traffic Aspects of the NSFNET

Kim Claffy (San Diego Supercomputer Center, USA, kc@cs.ucsd.edu)
Hans-Werner Braun (San Diego Supercomputer Center, USA, hwb@sdsc.edu)
George Polyzos (University of California, San Diego, USA, polyzos@cs.ucsd.edu)

This paper presents the architecture of data collection for the T3 NSFNET backbone service, and difficulties with using the collected statistics for long-term network forecasting of certain traffic aspects. We describe relevant aspects of the T3 backbone architecture, describe the instrumentation for the statistics collection process, and how it differs from that on the T1 backbone. We then present long-term NSFNET data to elucidate long term trends in both the reachability of Internet components via the NSFNET as well as the growing cross-section of traffic. We focus on the difficulties of forecasting and planning for these two traffic aspects in an infrastructure whose protocol architecture and instrumentation for data collection was not designed to support such objectives.

CBB: A Measurement Study of Changes in Service-Level Reachability in the Global Internet

Michael Schwartz (University of Colorado, Boulder, USA, schwartz@cs.colorado.edu)
John Quarterman (Matrix Information & Directory Services, Inc. (MIDS), USA, jsq@tic.com)

The Internet is currently in a period of exponential growth, as measured by domain registrations and packet counts. Increasingly often, people want to know how fast a particular part of the Internet is growing - to do capacity planning, gauge commercial promise, or simply to understand this important change in our society. Yet, looking only at registrations and packet counts does not uncover the full complexity of the situation. There are a variety of ways that sites can connect to the Internet, each offering different capabilities, costs, and technical problems. Moreover, growing awareness of network security problems is changing the way people think about connecting to the Internet, based on mechanisms such as firewall gateways. In this paper we analyze Internet growth based on measurements of which of a dozen common TCP services could be reached at each of over 13,000 domains worldwide, tested four times over the course of 1992. We analyze this data as a function of country, type of institution, and type of service. We also derive mathematical models that can be used to project growth and disconnection rates for individual countries and the global Internet.

CBC: IP Traffic Measurements and Analyze at CERN

Jessica Yu (MERIT Network Inc., USA, jyy@merit.edu)
Jean-Michel Jouanigot (CERN, Switzerland, jimi@dxcoms.cern.ch)
Olivier Martin (CERN, Switzerland, martin@cearn.cern.ch)

Because of the heavy concentration of lines into CERN a significant fraction of the total European IP traffic is crossing routers located on our site. CERN is thus ideally placed to collect information about the traffic flow, both at the macro and micro levels. In particular, an analysis of the trans-European and transatlantic (over the CERN/ANS line) flows has been made, and the results have been compared with those available from MERIT, the only other source of global Internet traffic statistics known to us.

Session CC: Mastering the Topology Stack

chair: Francois Fluckiger (CERN, Switzerland, fluckiger@vxcrna.cern.ch)

CCA: Domain Name System Forwarding Strategies

Alan Barrett (University of Natal, South Africa, barrett@ee.und.ac.za)

The Domain Name System (DNS) is a distributed database implemented by name servers on many hosts throughout the Internet. The DNS is used for several purposes, including mail routing, and mapping in both directions between system names and numeric addresses. In order to satisfy requests for information that is not available locally, name servers often have to forward requests to other name servers.

Configuration errors, network or host failures, and poor choice of forwarding strategy can have a significant adverse effect on network traffic volume and on name server load. These adverse effects are exacerbated by high delay network links. This paper will consider the traffic volume and name server load implications of some common forwarding strategies in the presence of some common problems.

CCB: How to Maximize Mail Address Flexibility and Reliability Using Standard Sendmail and BIND

Imazu Hideyo (Matsushita Electric Industrial Co., Ltd., Japan, himazu@isl.mei.co.jp)

I have developed an extremely flexible Sendmail/DNS scheme using the standard Sendmail and BIND software, neither of which require modification. In my implementation, both the destination user and host for a given mail address are obtained from a DNS query using the entire address, in contrast to obtaining only the destination host based only on the domain portion in the standard scheme. Likewise, the sender address is determined by a DNS query using both the originating user and host. Some of the key advantages are: two users of the same machine can use entirely different domain portions for their individual mail addresses, independent of the user's front end software, since this resolution is implemented in the delivery system. Multiple users of the same domain portion can be distributed among multiple hosts without any dependency on a central mail hub for mail distribution. Because this implementation is upward compatible with the standard scheme, it is possible to migrate individual divisions, without requiring sitewide modifications. We have been using this system for over 1000 users for the past 2 years, with significantly enhanced reliability as well as flexibility.

CCC: MBONE - the Multicast Backbone

Hans Eriksson (SICS, Sweden, hans@sics.se)

MBone the Multicast Backbone is technically a network of hosts running a multicast daemon. Functionally, MBone provides a multicasting facility to the Internet. This facility has been used to distribute audio and video from the IETF meetings, position data from a robot at the bottom of the Sea of Cortez, and a late night feature movie, just to mention a few items. The facility has also been the cause of severe Internet network meltdowns due to different reasons. The network providers have viewed various opinions about what to do with multicasting as it is in MBone, varying from forbidding to encourage pilot experiments. This paper will discuss the operational aspects of MBone, the symptoms that has surfaced and how it has evolved during the last year into a kind of semi- maturity.

Session CD: Dealing with Heterogeneous Networks

chair: Hans Eriksson (SICS, Sweden, hans@sics.se)

CDA: Mapping Networks and Services

John Quarterman (Matrix Information & Directory Services, Inc. (MIDS), USA, jsq@tic.com)
Smoot Carl-Mitchell (Texas Internet Consulting, USA, smoot@tic.com)
Gretchen Phillips (University of Buffalo, USA, phillips@acsu.buffalo.edu)

Computer networking is widespread, supported by the global Matrix of interconnected computer networks such as the Internet, BITNET, UUCP, and FidoNet. Perhaps 20,000,000 people use this global Matrix of computer networks, and several major global computer networks are growing exponentially, rapidly reaching new countries and new classes of users. This growth prompts network users and service providers to ask questions like:

o+ Where are these users? o+ What is the magnitude of the various network communities? o+ What are the available services? o+ How are services distributed? o+ And, finally, how can this data be represented in a meaningful graphical way?

We have chosen Latin America and the Caribbean (ALyC) as the geographic region for discussing these issues and demonstrating the mapping of currently available data. These are not traditional shaded geographic outlines nor node and link connectivity maps, but rather, maps that show locations and numbers of networked hosts and services. This representation of the network gives a presentation of where the networks go, rather than how they get there, and shows what areas are reachable by each network type. We do not show exactly where the users are, but we show where the hosts they use are. These maps can give network administrators and users alike a better understanding of the areas connected to the global Matrix, and the services available in each area.

CDB: The TipTop Gateway: Transporting K12Net Between the U. S. and Europe by Tunneling the Internet

Louis Van Geel (B- Net & K12Net, Belgium, lvg@k12.be)

K12Net is a grass-root network of school based/oriented BBS's throughout the world, using FidoNet technology and sharing curriculum-related conferences. The transport method used by FidoNet is mainly modem-based, and thus involves the normal charges for telephone calls. Although very cheap in operation on a local basis, moving data over long distances has restrained it's growth for many years.

With the gracious cooperation of IP networks, experiments have been carried out to tunnel FidoNet to FidoNet traffic over the Internet's intercontinental TCP/IP links.

The paper explains the technique used in both networks to tunnel the Internet. It also explains the procedures put in place on the FidoNet side to keep the mailflow under control.

CDC: Extending the LAN via ISDN

Graham Knight (University College London, United Kingdom, knight@cs.ucl.ac.uk)

This paper describes work carried out in the Department of Computer Science, University College London to make use of Basic Rate ISDN for teleworking. The hardware, software and network architecture used are described, and the costs of using the systems and the performance they offer are discussed.

CDD: Key Issues in Transporting TCP/IP over Public Data Networks

Michael Conn (MCI, USA, 0004387451@mcimail.com)
Dave Mcdysan (MCI, USA, dmcdysan@mcimail.com)

This paper summarizes and contrasts key issues associated with the use of the Frame Relay (FR), Switched Multi-Megabit Data Service (SMDS), Asynchronous Transfer Mode (ATM) and Broadband Integrated Services Digital Network (B-ISDN) public network protocols for transporting the Transmission Control Protocol/Internet Protocol (TCP/IP). This paper also analyzes the interaction of the TCP/IP protocol stack with these emerging public data network services and identifies how data networking performance can be tuned or improved. Key issues covered include IP encapsulation methods, traffic controls, congestion controls, routing methods, security implications, and the effects of the public network protocol on end-to-end performance. The current benchmark of IP routers connected by private lines in the Internet is used as the starting point. The implications of supporting such an IP network is shown to be sensitive to delay, which will be largest in initial frame relay networks, but decrease markedly with the deployment of SMDS and then ATM. It is shown that within a few years that public data networks will offer performance at levels similar to high speed private line based data networks in a cost effective manner by sharing the bandwidth on the backbone and the buffer resources within the network. The public data network will be able to reduce delay for multi-hop IP routes that take a long indirect path when route discovery mechanisms allow a more direct public network path to be found and used.

Furthermore, broader connectivity for both the access line to the IP routers, and the trunks between IP routers, can be achieved by public data networks.

Session CE: Network Operation and Management

chair: Peter Elford (Cisco Australia, Australia, pelford@cisco.com)

CEA: Accomplishing Performance Management with SNMP

Allan Leinwand (Cisco Systems, USA, leinwand@cisco.com)

As internetworks grow and expand, the need to provide a way to manage them is growing as well. Management requires information and a process to interpret and analyze it. The proliferation of the network management protocol SNMP has provided means for a wealth of information to be retrieved from network devices. These devices can include: personal computers, internetworking systems, minicomputers, mainframes, and so forth. Each of these devices can provide information which is useful to accomplishing network management. The problem is determining how to use this data in applications to help accomplish network management.

RFC1213 (MIB-2) and RFC1271 (RMON MIB) are documents which specify the data available from network devices. The information is found in individual pieces called objects. An object may be nearly any piece of information, ranging from a text description of the device to a number denoting the total number of packets sent out a single interface. MIB-2 describes the objects available from any device which is SNMP compatible while the RMON MIB defines objects available from remote network monitoring devices. These two documents will together form the basis of standard information available from many internetworks.

The objects in MIB-2 and the RMON MIB give network managers the ability to gather information applicable to the five critical areas of network management as specified by the Network Management Forum. These areas are: fault management, configuration management, security management, accounting management, and performance management. However, merely gathering the information does not automatically help in the management of a data network, exploration of the application of this information in each area of network management is necessary.

While exploring how to use this information in all the areas of network management is important, this paper will describe how to apply the information in the MIB-2 and the RMON MIB toward accomplishing performance management.

This paper will first define performance management as the process of measuring the performance of all elements of the network. This involves procedures to measure device throughput, link utilization, examine error rates. The performance management process can allow for the avoidance of congestion and can provide information to ensure that the network meets the users' needs.

Next, this paper will describe in detail how to use the objects in MIB-2 and the RMON MIB to produce applications which are useful for performance management. The description of the applications for performance management will be divided into two areas: monitoring current network performance and long term performance evaluation.

Monitoring current network performance helps managers examine devices and ensure that the network is operating with acceptable performance. This monitoring can help isolate potential faults and congestion before a disruption of network activity.

Long term performance evaluation can help network managers see trends in the usage of the network which can aid in capacity planning. Capacity planning can ensure that the growth of the network can always stay ahead of the demands of the users.

The paper will conclude by emphasizing that performance management can be accomplished by using the objects in MIB-2 and the RMON MIB.

CEB: Improved Network Management Using WIDE/PhoneShell

Hiroyuki Ohno (Tokyo Institute of Technology, Japan, hohno@is.titech.ac.jp)

Wide area networks today contain many types of equipment such as gateways, routers and modems. These are controlled and maintained by network managers from their terminals.

Often, there is no way to detect and repair problems on the network without manager intervention. This means all service could be stopped when problems occur. If no manager is actively working on the network, then several hours or days could pass before the network is restored. A solution to this common problem is to develop a mechanism for network maintenance without the necessity of computer terminal access. The WIDE/PhoneShell was developed specifically for this purpose.

The WIDE/PhoneShell was designed by the "network managers model" which is proposed in this paper. The model shows that network managers and computers exchange three basic types of messages to keep the networks stable: "request", "response", and "interrupt". The problem is one of how network managers can input such messages without terminal access.

The WIDE/PhoneShell solution consists of the three subsystems which are the DTMF I/O subsystem, Voice I/O subsystem and FAX I/O subsystem. These enable network managers to communicate with computers via telephone line. Thus, under the WIDE/PhoneShell environment, network managers can send commands and data using DTMF signals (telephone touch tones) and get feedback by computer generated voice. Network managers are thus able to monitor and repair computer networks without using modem or terminals.

Under the WIDE/PhoneShell environment, a pager control system named "WIDE/PCS" was also developed, which includes a "mail to pager" service. Recently some Japanese pagers with expanded memory capacity can display not only ASCII characters but also up to 36 kanji characters and pre-programmed illustrations on the pager's display. This greatly increases the complexity of alert messages possible to a network manager. These pagers are controlled by the DTMF I/O subsystem.

Another useful application under the environment is telephone-ping service. This is the oldest service in the WIDE/PhoneShell utilities. Managers can execute a ping command with DTMF and get an answer by computer generated voice. A simulator to learn the two stroke DTMF input method was created for the DTMF I/O subsystem.

"pingscript" is an important tool. It watches network status by sending ping packets several times an hour. When pingscript fails to send ping packets to a host, it assumes the host is unreachable. It will then ring the manager's pager with network status information. Usually pagers ring within a maximum of 6 minutes. They will also ring when the network or host recovers from a problem. Managers can now know the constant state of the network without begin tied to their terminals.

The WIDE/PhoneShell has been under development since 1989 and been used for daily network management in the WIDE project since 1990. The pager control system WIDE/PCS has been expanded for use on many sites to report problems. Recently, at the iNET'92 in Kobe Japan, it was introduced for network management.

CEC: Minerva: An Event-Based Model for Extensible Network Management

David Hughes (Bond University, Australia, bambi@bond.edu.au)
Zheng Da Wu (Bond University, Australia, wz@bond.edu.au)

The goal of the work outlined within the paper is to provide a Network Management System (NMS) capable of supporting "non-typical" environments. If is viewed that no NMS developer can forsee every requirement and ad hoc function that a network manager may desire. As such, the NMS must be flexible and extensible in terms of the acquisition, interpretation, storage and display of information. This should not be limited to information pertaining to network interface byte counts or other "typical" network related data, but also cater for any information the network manager, or systems administrator, may wish to record.

A model developed for use with the Minerva management system utilises an event based structure to accomodate these requirements. Rather than passing pieces of raw information to the NMS, notification of events are forwarded. An event may be defined to cater for virtually any situation, be it an interface on a router failing, a user logging into a host, or the central database failing to execute a batch update correctly. By definition, the receipt of such a notification may convey vast amounts of implicit information.

The user may define events to cater for the local environment. Each definition includes such information as the interaction with the database caused by the event and whether an alert should be sent to the operations staff.

To allow incorporation of this model into the working environment, tools have been designed to allow an event report to be generated from within C applications and shell scripts. This allows for simple integration of system administration programs and scripts that may currently be in use. Tools to provide the functionality common to todays suite of network management stations, including SNMP integration, have also been designed and are discussed.

Session DA: Tools for Networked Information Discovery and Retrieval

chair: Jim Fullton (CNIDR, USA, jim.fullton@cnidr.org)

DAA: UWIN - The University of Washington Information Navigator

Edward Lightfoot (University of Washington, USA, elight@u.washington.edu)

The University of Washington Information Navigator is a character-based program designed to provide easy access to campus, regional and Internet information resources and services. It has three aspects that may be of potential interest:

- An interface targeted at the non-computer users, using a multi-column browser designed to alleviate the potential confusion of deep menu hierarchies found in many other character interfaces.

- A license management system that controls the number of concurrent sessions available for a given resource, thus allowing implementation of database and software licenses based on the number of concurrent users.

- A menu manager that tailors the available resources and services displayed for a user, based on user groups. It is currently implemented for three user groups.

-The license manager and menu manager, in combination, facilitate a consortium of users such as NorthWestNet in negotiating group licenses for services.

UWIN is currently in production at the University of Washington and is available for a limited set of resources and services on the Internet at uwin@washington.edu.

DAC: Whois++ and Centroids: A Next Generation White Pages

Joan Gargano (UC Davis, USA, jcgargano@ucdavis.edu)

The original white pages service, Whois, offerred by the first Network Information Center was intended to offer simple, referrential information on people and machines on the ARPANet. The X.500 protocol expanded the scope of directory services to a global scope, critical to the operation of the network, while introducing a high level of complexity to the protocol. The paper describe a new generation of white pages protocol, Whois++, which attempts to provide a global directory service with the functional capabilities of X.500 and the simplicity of Whois.

Session DB: Building Applications

chair: Eric Hood (North WestNet and FARNET, USA, ehood@nwnet.net)

DBA: Proposal for a Subject Oriented User Interface to the Internet

Judith Ahrens (Drexel University, USA, ahrensjd@duvm.ocs.drexel.edu)
Gerardo Esquer (Drexel University, USA, esquerga@duvm.ocs.drexel.edu)

Although Internet contains valuable information for many constituencies worldwide, its unfriendly user interface limits its long-term potential as an accessible information resource. To open the Internet resources, we propose a user interface architecture that matches the way people search for information -- by subject -- and that features automatic connection. These capabilities will free users from having to learn abstruse format access protocols. Specifically, we propose a subject-oriented user interface involving three search strategies as a front-end to Internet's current format- oriented interface. The three search strategies are: (1) search by specific subject (2) search from a list of subjects (3) search for a subject within a format. The result of this fundamental reorientation will be freer and easier access to and dissemination of information globally. Although the reorientation will require considerable efforts in standardization, coordination, and global cooperation, we believe this is a goal well worth pursuing.

DBC: WorldWide Web Initiative

Tim Berners-Lee (WorldWide Web Project, Switzerland, tbl@cernvax.cern.ch)
Robert Cailliau (CERN, Switzerland, cailliau@cernnext.cern.ch)
Nicola Pellow (CERN, Switzerland,)
Arthur Secret (CERN, Switzerland, secret@dxcern.cern.ch)

The World-Wide Web (W3) is a way of viewing all the online information available on the Internet as a seamless browsable continuum. Using hypertext jumps and searches, the user navigates through an information world partly hand-authored and partly computer-generated from existing databases and information systems. The web today incorporates all information from more basic information systems such as Gopher and WAIS, as well as sophisticated multimedia and hypertext information from many organizations. As a user interface tool, W3 clients provides a comprehensive point-and-click network access tools, while W3 servers provide an efficient, friendly method of providing data to real users. This paper answers some commonly asked questions about W3 such as those comparing it with other systems, and about recent developments and future directions global hypermedia will take.

Session DC: Multimedia Technologies

chair: John Hine (Victoria University of Wellington, New Zealand, john.hine@comp.vuw.ac.nz)

DCA: Piloting of Multimedia Integrated Communications for European Researchers (MICE)

Peter Kirstein (University College London, United Kingdom, p.kirstein@cs.ucl.ac.uk)
M Handley (University College London, United Kingdom, m.handley@cs.uc.lac.uk)
A Sasse (University College London, United Kingdom,)

The paper defines multimedia conferencing and describes the facilities currently available. It discusses briefly the activities which require standardisation, and the progress that has been in this direction to date. It gives an overview of the MICE pilot project, which utilises existing conferencing rooms workstations, codecs and software, and the existing network infrastructure, to offer researchers conferencing facilities within Europe, as well as a link to the US. The goals of the project, its achievements to date, and problems encountered are discussed in detail. Finally, we outline forthcoming activities.

DCB: The Packet Video Project at MCNC

Jack Drescher (Center for Communications at MCNC, USA, drescher@concert.net)

MCNC, a nonprofit corporation located in Research Triangle Park, NC, operates separate video and data (TCP/IP) networks. Its goal is to combine these networks into a single data network that offers packet-based conferencing as a featured service, providing users with videoconferencing capability among themselves and access to worldwide videoconferencing with the Internet community. The recent initiation of a statewide broadband network project in North Carolina increases the probability of a successful endeavor and broadens the potential user base. This paper describes some of the approaches that MCNC's Packet Video Group is using to achieve this goal.

DCC: Hyper-G: A Distributed Hypermedia System

Frank Kappe (Graz University of Technology, Austria, fkappe@iicm.tu-graz.ac.at)

Hyper-G is a general-purpose, large-scale, distributed hypermedia system currently developed at the Graz University at Technology. It was designed for handling large amounts of multimedia data, and care has been taken to provide mechanisms for automatic maintenance of a dynamically changing body of information, advanced user interface features, and efficient use of network and computing resources.

While Hyper-G was originally designed to run in fast, local-area networks, this paper shows how Hyper-G is currently being transformed into a global hypermedia information system distributed over Internet that retains Hyper-G's advanced functionality such as automatic link generation and maintenance, navigation facilities, access rights, and distributed searching, and makes efficient use of network bandwidth by extensive caching.

Session DD: Multimedia Tools

chair: Peter Kirstein (University College London, United Kingdom, p.kirstein@cs.ucl.ac.uk)

DDA: A Protype for a Multimedia Message Handling System Based on a Model of Semantic Base Info Representation With Multiple Media

Shingo Fujimoto (University of Electro Communications, Japan, shingo@cs.uec.ac.jp)
Satoshi Irioka (University of Electo-Communications, Japan, irioka@cs.uec.ac.jp)
Hideki Sunahara (The University of Electro-Communications, Japan, suna@cs.uec.ac.jp)

Current multimedia message handling systems have several problems in managing multimedia data efficiently since these systems are designed for text messages. In this paper, we propose a model of information representation with muttimedia data. A key idea of the model is the separation of semantics and presentation of information. Our model will provide the environment which can handle each media efficiently. The message can be selected by the recipients depending on their presentation type. It will give recipients some comprehensible information even if he doesn't have the player system for the message with the specific data type. This paper also describes the implementations of a prototype system of multimedia message handling based on the proposed model using the "X window system" as its user interface. A message and data format in the system is based on "MIME" format.

DDB: Architectural Foundations of PINE: a Multimedia Email and News Client for DOS and Unix

Terence Gray (University of Washington, USA, gray@cac.washington.edu)

Pine is a character-based mail and news program emphasizing ease-of-use. Other key design goals were: full Internet mail interoperability including support for multipart, multimedia MIME messages, support for a robust client-server email architecture using the Interactive Mail Access Protocol (IMAP), and integration of mail and network news. A PC version of Pine was developed to improve integration of the mail system and the desktop file system for DOS users. Pine(tm) is copyrighted but freely available, and is in use at more than 2000 sites in over 30 countries. The paper provides background on Pine, but focuses primarily on the requirements for a global email system and their design implications.

Session DE: Technologies for Cooperative Work

chair: Barry Kort (BBN Labs, USA, barry@michael.ai.mit.edu)

DEA: Toward Seamless Collaboration Media: From TeamWorkStation to ClearBoard

Hiroshi Ishii (The University of Toronto, Canada, ishii.chi@xerox.com)

I present an evolution of our collaboration media design from TeamWorkStation to ClearBoard. The key concept of our media design is "seamlessness". TeamWorkStation provides distributed users with a seamless shared workspace. Each user can continues to use favorite application programs or desktop tools, so there is only a minor cognitive seam between the individual and shared workspaces. To integrate the shared workspace and the partner's image seamlessly, we designed ClearBoard. Based on a metaphor of transparent glass board, ClearBoard permits co-workers in two different locations to draw with color markers or electronic pen while maintaining direct eye contact and the use of natural gestures. I believe these collaboration media will be promising applications for the future multimedia internets.

Session DF: NIDR Research and Development Initiatives

chair: Michael Schwartz (University of Colorado, Boulder, USA, schwartz@cs.colorado.edu)

DFA: Collaborative Information Retrieval: Gopher from MOO

Larry Masinter (Xerox Palo Alto Research Center, USA, masinter@parc.xerox.com)
Erik Ostrom (Gustavus Adolphus College, USA, eostrom@nic.gac.edu)

There are two visions of how use of the global network will evolve in the future. First, individuals will use the network as a resource, providing access to material from libraries and other suppliers of information and entertainment. Second, in addition to communicating with these data sources, people will communicate with each other, using a variety of interactive text, audio, and video conferencing methods.

This paper is about a system that combines the two uses: adding an information retrieval tool (Gopher) to a `text based virtual reality' environment (MOO). The combination allows informal collaboration using information retrieval to happen across the network.

DFB: Research Problems for Scalable Internet Resource Discovery

Michael Schwartz (University of Colorado, Boulder, USA, schwartz@cs.colorado.edu)
C. Mic Bowman (Pennsylvania State University, USA, bowman@cs.psu.edu)
Peter Danzig (University of Southern California, USA, danzig@usc.edu)

Over the past several years, a number of information discovery and access tools have been introduced in the Internet, including Archie, Gopher, Netfind, and WAIS. These tools have become quite popular, and are helping to redefine how people think about wide area network applications. Yet, they are not well suited to supporting the future information infrastructure, which will be characterized by enormous data volume, rapid growth in the user base, and burgeoning data diversity. In this paper we indicate trends in these three dimensions, and survey problems these trends will create for current approaches. We then suggest several promising directions of future resource discovery research, along with some initial results from projects carried out by members of the Internet Research Task Force Research Group on Resource Discovery and Directory Service.

Session DG: Distributed Applications Management

chair: Eric Hood (North WestNet and FARNET, USA, ehood@nwnet.net)

DGB: WWFS: An Evolutionary Framework for File Sharing in the Internet

Youki Kadobayashi (Osaka University, Japan, youki@ics.osaka-u.ac.jp)
Hideo Miyahara (Advanced Institute of Science and Technology, Japan, miyahara@ics.osaka-u.ac.jp)
Suguru Yamaguchi (Advanced Institute of Science and Technology, Japan, suguru@is.aist-nara.ac.jp)

This paper presents the design and implementation of the Cluster Server (CS), which is a distributed file system that works over the Internet. CS is designed specifically as the resource access mechanism of WWFS (WorldWide File System). CS interacts with electronic archives via file transfer protocols. In the meantime, CS handles client requests via file sharing protocols. CS currently implements FTP to access file servers and NFS to handle requests from clients, thereby enabling simple and transparent access to the Internet FTP archives. CS has been implemented as a user-level NFS server, that obviates the need to modify client BSD kernel and achieves transparency. CS caches transferred files for the serving cluster, so that we can significantly eliminate duplicate file transfer over long-haul links. CS makes it easy to implement multiple protocols within one process by providing basic framework for multi-threaded network daemon. With the multi-protocol framework, we can smoothly introduce new file transfer protocols and new file sharing protocols into the Internet.

DGC: Prospero: A Base for Building Information Infrastructure

Clifford Neuman (USC Information Sciences Institute, USA, bcn@isi.edu)
Steven Augart (University of Southern California, USA, swa@isi.edu)

The recent introduction of new network information services has brought with it the need for an information architecture to integrate information from diverse sources. This paper describes how Prospero provides a framework within which such services can be interconnected. The functions of several existing information storage and retrieval tools are described and we show how they fit the framework. Prospero has been used since 1991 by the archie service and work is underway to develop application interfaces similar to those provided by other popular information tools.

Session EA: Training and Support for Virtual Communities

chair: Paul Peters (Coalition for Networked Information, USA, paul@cni.org)

EAA: Educating Network Users in Distance Education Through Computer Networks - Experiences from Running Distance Education Courses Through UNINETT

Astrid Jenssen (University of Oslo, Norway, astrid.jenssen@usit.uio.no)

Only a small numbers of all academic and research staffs are using the applications available through computer networks. Linked to UNINETT (the Norwegian academic network connected to Internet), users have access to an international communication infrastructure through interconnected computers. The network provides computer-mediated communication (CMC) and resource sharing services.

At the Center for Information Technology Services (USIT), University of Oslo, we are making information technology available for use in distance education. One particular project, FIK - Distance Education through Communication Services, are providing distance education courses for educating network users through UNINETT. FIK is aimed at increasing the number of users and improving the knowledge of available network services. Through a distance education course the participants learn to use the services from their own desktop computers in familiar surroundings. The services are used for interaction, for educational purposes, as well as being a part of the curriculum. FIK has been offered since the autumn of '91, and different text-based services have been tried out as the main channel of communication.

There has been a continuing development process, improving the course offer, expanding the target group and tailoring the course offer for specific groups of participants. The target group has grown from participants mainly from the University of Oslo, to a national target group of users and potential users of UNINETT's services. This paper presents the FIK-project; the distance education model, experiences from using the network services in distance education and some of the challenges met in administrating and developing the course offer.

EAB: Supporting End-Users of Networked Bibliographic Databases

Terry Morrow (University of Bath, United Kingdom, t.m.morrow@bath.ac.uk)

In February 1991 a new UK bibliographic data service was launched. It is believed that it was a world first - a national service providing access to commercially supplied bibliographic databases, potentially available to tens of thousands of academics and hundreds of thousands of students, and free at the point of delivery. In all, around a million UK higher education students, research workers and staff are eligible to use this service.

BIDS (Bath Information & Data Services) currently provides access to bibliographic data supplied by two major commercial suppliers - the Institute for Scientific Information, and Excerpta Medica, a division of Elsevier Science Publishers. This data has previously been accessed by academics using printed publications, or by using CD-ROM, or through one of the commercial data services such as Dialog, Data-Star or DIMDI. Because of the pay-as-you-use charging mechanisms, the commercial search services are usually accessed by librarians acting as intermediaries.

Access to the BIDS services is provided directly to end-users via JANET (the UK's Joint Academic NETwork), and JANET is also used to deliver search results back to the user via e-mail. Permission to use the service is granted by site license, and this license entitles any member of a subscribing institution (staff, postgraduate,

undergraduate, research worker etc) to unlimited access. The service continues to grow in popularity. At the time of writing there are around 14,000 accesses a week to the ISI data alone, and this number continues to rise. So far, 73 institutions have taken out licenses to BIDS-ISI, and 21 are licensed to use BIDS-EMBASE.

Because they are networked, these services are freely available from outside the library and this wide area access together with the scale of use has raised new problems of how to deliver effective support and instruction. One-to-one advice by library, computer centre, or BIDS service staff is simply not practical for the vast majority of users.

This paper will first attempt to define exactly what is meant by support. It will then describe the various strategies that were adopted to meet the problems of providing effective end-user support for these services. Topics covered include

- the involvement of supporters in the design and development of both the service and the support materials; the paper will describe how the contributions of a team of librarians were organised, and discuss the importance of this collaboration in designing appropriate and widely used materials

- design features which appear to have been critical to the success of the service such as simplicity of structure, ease of navigation, context-sensitive help, and short-cuts for the experienced

- the importance of well designed, professional printed support materials, both for end-users and for supporters and instructors

Finally, the paper will try to draw some lessons that might be applicable to other networked information services.

EAC: Community Computer Networks

Steve Cisler (Apple Computer, Inc., USA, sac@apple.com)
Dewayne Hendricks (Tetherless Access Ltd., USA, dewayne@netcom.com)

Since the late 1980's groups within a number of communities have joined together to start community computer networks that can bring affordable access to many people and organizations within the areas served. A community network consists of one or more computers providing services to people using computers and terminals to gain access to those services and to each other. This essay discusses several current models and examples of community networks, examines what impact they can have on the environment in which they grow, and makes a few wild guesses about the future of such networks given the changes in computing and telecommunications technologies. While many of these networks tie together geographically separated individuals linked by a common interest or profession, this essay will not cover scholarly or affinity group networks. Other models and systems exist but have been described elsewhere. Instead, community will be used in the sense of a municipal, county, regional area, or Indian nation. The use of wireless technologies used in a city experiment is discussed as well.

EAD: Building Electronic Communities

Jill Foster (Univ of Newcastle upon Tyne, United Kingdom, jill.foster@newcastle.ac.uk)

Building electronic communities is beset with many difficulties: lack of awareness of the possibilities, lack of a framework for collaboration, lack of training in using the network, complex user interfaces, and other hurdles.

The paper explores some of these problems and suggests some solutions based on the experience gained with various projects within Europe and in particular with the NISP/Mailbase Project (managed by the author).

The UK NISP Project aims to help groups of academics, and in particular non traditional computer users, build their electronic communities. The UK Mailbase Service (enhanced mailing list server) underpins this activity. The NISP Team provides help and encouragement to user groups on two levels:

- specific advice and encouragement to those groups who simply want to make use of the Mailbase mailing list facilities to improve communication amongst their members

- more general support for a small number of groups by training and helping them to build their electronic communities and to make good use of the global networked services.

The following is just a small selection of some of the groups using Mailbase:

- Aromatherapy Researchers - ART: the UK higher education Art community - Biology (teaching) - Chemists: computational chemistry and modelling - Childcare-quality: social science research group - Engineering Design - History - Latin American Studies group - Liturgy - Libraries - Medieval and Renaissance Music - Regional Research Laboratories (Social and Economic scientists)

More extensive support and training is given to groups which can provide a reasonable amount of volunteer or funded effort. The NISP Team works with a nominated subject specialist and trains them in the use of the network. NISP then helps this group "coordinator" set up a programme of activities and facilities for the group such as:

- Workshops to train the members of their community in the use of the various networked services.

- A coordinated set of mailing lists for discussions, sharing of reports, papers etc.

- Subject based resource guides and other documentation on using the network.

One of the main aims of the NISP Project is to promote the use of the network and the networked services across the whole spectrum of research areas. It is very rewarding to witness this happening and to see the enthusiasm of new groups as they realise the potential benefits of using the network.

Session EB: Panel: Virtual Culture and Community

chair: Jean Polly (NYSERNet, Inc., USA, jpolly@nysernet.org)

EBA: Virtual Culture and Institutes for Research

Ermel Stepp (Marshall University, USA, estepp@wvnvm.wvnet.edu)

Three institutes have been organized as global virtual communities of research fellows engaged in collaborative research. The three virtual institutes are: The Institute for Research on Virtual Culture, The Institute for Research on Visionary Leadership, and The Merton Research Institute. Each institute functions through optimal utilization of innovative networked information systems to sustain discourse and create and distribute documents. Each institute conducts a continuing electronic conference on substantive research, publishes an electronic journal, and maintains networked archives of the journal, research papers/reports, masters theses, doctoral dissertations, and related documents.

Session EC: Global Networks and K12 Education

chair: Bruce Nelson (Novell Corp., USA, bruce_nelson@novell.com)

ECA: The Boulder Valley Internet Project

Ken Klingenstein (University of Colorado, USA, kjk@spot.colorado.edu)

The Boulder Valley Project is a comprehensive effort by the Boulder Valley School District (BVSD) and the University of Colorado at Boulder to deploy Internet connectivity, develop curricula, and assess network impact on the education programs and operations of an entire school district. The project derives support from National Science Foundation, local corporations, district and university funds, and regional networking organizations. The goals of the project involve integrating networking into all aspects of the K-12 environment and, to a significant degree, the surrounding community.

The project is distinctive in its systemic and analytic approaches. It is systemic in that all aspects of the district environment - classrooms (teachers and kids), school libraries, administration, community involvement - are being provided with networking opportunities and are developing approaches to utilize its capabilities. It is analytic in both its carefully staged development and the formal assessment being incorporated. Moreover, while there are many instances nationally of good individual efforts in K-12 use of the Internet, there is little being done to design scalable approaches. Since widespread deployment of networking in K-12 will require structured methodologies, we intend to determine the key factors for sustainable networking.

This is the first year of the three-year project. The project is intended to pioneer a comprehensive approach to K-12 use of networking and the Internet, where planning and implementation involve district, teacher, university, community, and regional participation. A number of observations have already emerged, concerning appropriate connectivity models, the support of curriculum development, teacher training, political issues to be managed, assessment metrics, and outreach opportunities. After a brief overview of the goals and status of the project, this paper will focus on some of the early lessons learned and identify needs and issues that require further work. Given the remarkable impact that networking will ultimately have on learning environments (early education, i.e. K-12, higher education, and lifelong learning), it is important to establish effective models for others to follow.

ECB: Educational Networking: Meeting Educators' Needs

Robert Tinker (TERC, USA, bob_tinker@terc.edu)

Network professionals see the network as essential to their personal learning and assume that their style of network use should be extended to all of education. But the educational enterprise is more complex than such a model assumes. Many serious educators also see networking as a key ingredient in educational reform, but only as part of a larger whole. In this article examples of successful network based educational projects are presented in order to clarify the essential features of this larger picture. Using these principles, a large class of exciting, low-bandwidth network educational activities are described that could become self-supporting.

ECC: The Globe Project: A Place for Youth to Make a Difference by Working Together Globally Using Telecommunications

Peter Copen (Copen Family Fund, USA, pcopen@igc.apc.org)

There is clearly an exponential rise of serious problems in the world. This work - the networking of youth globally by telecommunications - provides the opportunity to link youth around the world in order for them to learn to solve these problems exponentially and to be more prepared to handle the enormous and critical responsibilities of the 21st Century.

The work of The Copen Family Fund has been to create a pilot project that demonstrates this opportunity. We now have done so!

This work is more than educational. We feel that it is a moral and ethical responsibility that we have to our children and our grandchildren for the future.

In the last 5 years, The Copen Family Fund has been a recognized leader in this field by setting up pilot projects that are at the cutting edge of global telecommunications for youth. Each of these pilot projects demonstrates how youth can work together exponentially and internationally, for the health and welfare of the planet and its people. Each of these projects builds and expands upon its predecessor.

These pilot projects have been based on the use of low-cost communications technologies such as electronic mail and video- speaker telephones which enable youth to work together on meaningful-joint projects.

The first of these projects started with 24 schools in Russia and the United States and then expanded to include over 400 schools in 20 countries. The second of these projects, a collaboration between 6 of the largest educational-telecommunications networks in the U.S., can impact over 35,000 schools in the U.S. and internationally. The third project is designed to include virtually any school, youth-service organization, network or person on the planet!

Session EE: Supporting Global Emergency Management

chair: Ann O'Beay (MCI, USA, 4005058@mcimail.com)

EEA: Toward a Global Disaster Management Information Network: The Role of the Internet

Peter Anderson (Simon Fraser University, Canada, anderson@sfu.ca)
David Butler (University of Colorado, USA, hazards@vaxf.colorado.edu)

This paper examines the potential role that new communication and information technologies might play in hazards research and disaster management. Included is a discussion of current national and international efforts to bridge the gap between scientists and disaster planners, policy makers and practitioners. The actual and potential integrative roles of the Internet and other international and national networks are also explored, including efforts to establish integrated domestic and international disaster management information networks.

EEB: Internetworking and Global Environmental Research

Herre Heersma (RIVM, Netherlands, herre@rivm.nl)
Otto Tissing (RIVM, Netherlands, tissing@rivm.nl)

After the Chernobyl accident we realized once more that the effect of environmental disasters does not stop at national borders but can extent to a continental or even global scale. However, in many cases timely information concerning the state of the environment on a large geographical scale has appeared difficult to obtain, mainly due to a lack of sufficient tools for data interchange. Our experience has shown that internetworking can play an important role in collecting and distributing environmental information as a basis for early warning systems and the development of potential abatement strategies. In the world of environmental research, ozone is interesting because it contribute substantially to smog formation and is involved in the greenhouse effect. This session will present the internetworking strategies of our institute related to a few research projects directed to the occurence and behaviour of ozone in troposphere and stratosphere. Ozone monitoring stations are operational in almost all European countries. From selected stations all data should be available to all project participants. A common system for interchanging these data has therefore been developed for this purpose. As a part of the global network for the detection of stratospheric change we have now established a remote controlled station in New Zealand to measure the thickness of the ozone layer. We intend to offer the control and data acquisition feature via Internet. Finally, continental-scale management of environmental data demands distributed data concepts. This implies that networks must go hand in hand with related services to organize interpersonal communication along with data interchange.

EEC: CERT Incidence Response and the Internet

Katherine Fithen (CERT Coordination Center, USA, kfithen@cert.org)
Barbara Fraser (CERT Coordination Center, USA, byf@cert.org)

When responding to computer security incidents throughout the Internet community, the CERT Coordination Center utilizes the Internet in many ways. CERT activities involve both individual Internet sites and the global Internet as a whole, and CERT uses resources from the Internet, as well as providing services to the Internet. Included in this paper are discussions of Internet resources employed by CERT, and information on how members of the Internet community can use CERT products and services.

Session EF: Cyber-Knowledge and Information-Space

chair: Eric Manning (Keio University, Japan, manning@chitose.doi.cs.keio.ac.jp)

EFA: Presence and Form in the Architecture of Cyberspace

Barre Ludvigsen (Ostfold Regional College, Norway, borrel@dhhalden.no)

The most important prerequisites to our understanding of architecture are time and presence. Le Corbusier's definition of architecture, which is as good as any, declares that "The business of architecture is to establish emotional relationships by means of raw materials". He goes on to describe the architect: "by the relationships which he creates he wakes profound echoes in us, he gives us the measure of an order which we feel to be in accordance with that of our world."

The state of architecture 70 years ago may with good reason be compared to that of the architecture of cyberspace today. Engineers have laid the foundations of beauty, form and order in their contributions of new raw materials to the dream of presence in time removed from "here and now". The natural and well known concept of consciously being somewhere else in the present.

Assuming the attraction of chaos and the vernacular of cyberspace, its existence is only assured through the juxtaposition with designed order. The various concepts, rules and fashions of graphic and visual design may readily be applied to those aspects of cyberspace that can be visually observed and that mimic the visual in physical architecture. Furthermore, manipulating form to follow function is probably what cyberspace has been all about until now. What seems to be lacking are those elements of physical architecture that give us most pleasure and to which we sometimes attribute beauty.

We need to study those elements of time and communication in cyberspace that have created the vernacular architecture of virtuality. We need to recognize the meaning of presence in cyberspace and the manners in which its raw materials can be used to form the orders from which we draw the very basic needs of security and confidence.

A primary aim in the definition of an architectural model for designing cyberspace must therefore be the recognition of form. That which is visible, tangible and audible in physical space as spacially bounding surfaces. What then are the boundaries of cyberspace. Might they be found in the definitions of the concepts of form and shape?

Although the concept of form in cyberspace may seem alien in nature, a glance at the various definitions of form reveals a surprising variety of connections to the previously discussed elements of presence, order, meaningful relationships and aesthetics.

As we attempt to shape and design the environments of cyberspace, the elements of architectural method become inseparable from the essentials embodied both in the animate and inanimate concepts of form. The aesthetics, meaning and presence in form as viewed collectively by athlete, philosopher and architect.

EFB: Text, Internet: InterText

Barton Thurber (University of San Diego, USA, thurber@usdcsv.acusd.edu)
Jack Pope (University of San Diego, USA, pope@usdcsv.acusd.edu)

Modern computer networks, offer new mechanisms for human interaction, communication and understanding. As change agents, networks force social, economic and legal questions upon us that have no easy answers. What are the protections afforded to electronic mail and how do they differ (or should they differ) from those protections afforded standard "hard copy" mail? How is information "owned" on a network? What is the nature of ownership in the electronic world? Who are the authors, who are the readers, and is that distinction still meaningful? Within the Internet community, new tools have emerged to help users on a worldwide network cope with the management of information; but typically these new tools do not help with these large questions or, more generally, with questions about the nature of the information being managed. Gopher, for example, "is a distributed document delivery service. It allows a neophyte user to access various types of data residing on multiple hosts in a seamless fashion." But if the user incorporates any of this data into a project, conversation or interactive document he/she is writing or is a member of, the data is ontologically homeless-- especially as it is used and re-used in multiple, unpredictable and equally "homeless" electronic conversations. We argue that the essence of distributed electronic communication therefore involves much more than "information", and that what is involved is correspondingly more than information "management." We suggest that what is happening at the network level is also what has been happening to individual documents, that exchanges on the Internet are in important ways the same as exchanges within interactive hypermedia, and that what is being developed is less a network architecture as such but a cyberworld, in which individual documents as well as the network that generates access to those documents share the same purpose and structure. On both macro and micro levels, traditional author/reader distinctions are blurred, texts change according to decisions the reader makes about his or her own relation to the narrative which, in turn, allows for and in some cases requires, the reader to become the writer, placing his or her own writing within the narrative "field," not simply the text, created by the original author(s). We propose, finally, that what is at stake, in designing networks no less than in designing documents, is not the delivery of information per se, but the nature and structure of knowledge itself.(1)

(1) We are both member s of the Center for Research In Interactive Technologies, (CRIT), which has developed the NewBook Editor (Chariot Software), an authoring tool for the creation of interactive multimedia. We are currently working on a version of the NewBook Editor incorporating Gopher links, which would render the discussion above more than simply metaphoric.

EFC: Gender Swapping on the Internet

Amy Bruckman (MIT Media Lab, USA, asb@media-lab.media.mit.edu)

In text-based virtual reality environments on the Internet called MUDs, it is possible to pretend to be the opposite gender. In these virtual worlds, the way gender structures basic human interaction is often noticed and reflected upon. This paper introduces MUDs, and then presents a community discussion about gender issues that MUDs inspired. Next, a case study of one gender swapper examines why he chose to gender swap and what he and others around him learned from that experience. Gender swapping is one example of ways in which network technology can impact not just work practice but also culture and values.

Session FA: Economic Policy

chair: Steven Goldstein (NSF, USA, goldstein@note.nsf.gov)

FAA: The Internet and the Market System: Externalities, marginal cost, and the Public Interest

Sandra Schickele (Sonoma State University, USA, schick@mvax.sonoma.edu)

This paper lists conditions which must be met if an efficient private market for Internet services is to exist. Special consideration is given to the presence of externalities and the possibility of marginal cost pricing of network services. It is argued that the Interment does not satisfy the conditions for efficient production through the private market and that government subsidy and regulation are necessary.

FAB: Research Networks in Developing Countries: Not Exactly the Same Story!

Daniel Pimienta (Union Latina, Redalc Office, Dominican Republic, pimienta!daniel@redid.org.do)

This paper presents a structured set of guidelines to help start and operate research networks in developing countries. The proposed methodology is the result of a combination of studies and field experiences in Latin America since 1989 (REDALC Project). The introduction identifies the implicit model and the key factors for the success of research networks creation in the industrial countries. A comparison is made with the developing country's environment calling for a different approach. A summary list of the activities linked to network building and operation is shown which demonstrates that the bulk of activities is in managerial tasks rather than in technical one's. A hierarchical approach for problem solving is described in four layers: institutional, organizational, financial and, last, technical. The activities belonging to each layer are described with a set of associated guidelines. Finally, some success prone ingredients are presented. References are made to practical examples derived from the cases of the Dominican or Peruvian networks.

FAC: Network Analysis in Support of Internet Policy Requirements

Hans-Werner Braun (San Diego Supercomputer Center, USA, hwb@sdsc.edu)
Kim Claffy (San Diego Supercomputer Center, USA, kc@cs.ucsd.edu)

In this paper we describe the importance of network analysis in support of network policy considerations and evaluate a number of examples. We propound and support our hypothesis that, in the face of today's critical era in the evolution of global information infrastructure, Internet policy considerations and network analysis must begin to interact in ways not previously recognized or implemented.

Session FD: Future Global Network Policy

chair: Haruhisa Ishida (University of Tokyo, Japan, ishida@u-tokyo.ac.jp)

FDA: Global Networking from Asia-Pacific Perspective

Kilnam Chon (KAIST, Korea, chon@cosmos.kaist.ac.kr)

Computer networking is growing very rapidly and is becoming truly global. This paper addresses issues on global networking from the Asia-Pacific perspective, and proposes solutions to some of the issues.

Coordination for Asia-Pacific including link and organizational coordination is considered with respect to Global Internet Exchange and network information centers. Local language support is a rather unique issue to the region due to complexity of some of their characters. Other noteworthy issues are also considered.

FDB: On Policy Issues for the Operation of European Data Networks for the Research Community

Klaus Ullmann (Direction Operational Unit Ltd., United Kingdom, ullmann@dfn.d400.de)

Any discussion on European data networking has to review the tariff situation for both cross-border and national data transmission. This situation has improved compared with the situation a few years ago but still Europe can only be described on both relative and absolute terms as an expensive world region for data communication. The difference in relation to US tariffs is roughly a factor of 10 to 20. As a consequence the commercial market for data communication is not very much developed which in turn means that the research community in Europe has fewer options to rely on commercial offerings than the US research community.

A second specific characteristic in Europe is the absence of a permanent funding body for networking at the European level. This results in a funding structure, where the main funding comes from national funding bodies.

Having both facts in mind a growing consensus amongst people responsible for research networking in Europe for the selection of an appropriate organisational structure is emerging and a structure is being prepared which enables funding for pan-European networking and economic use of available market offerings. Future organisational needs have been formulated in the "European Consultative Forum for Research Networking" (ECFRN) and basically require a three layer structure - a European layer and a national layer for service provision and a political layer for funding. The political layer is now being established and is known as HIgh-Level Officials Group (HILOG). At the level of service provision the Operational Unit (OU) complements already now existing national networks - it has been established recently as a company limited by shares and is developing its business case. The main part of the business case is the provision of a 2 Mbps backbone immediately and the implementation of a 34 Mbps backbone in the next year.

FDC: EBONE - Backbone and Neutral Interconnect

Frode Greisen (EARN, Denmark, frode.greisen@uni-c.dk)

Ebone is a consortium of 22 academic and commercial IP service providers, most of which are European. Started in 1992, Ebone quickly came to cover most of Europe with a T1 backbone including T1 links to NSFnet and other US networks. The basic principle that the backbone is policy free made a rapid growth and effecient routing possible.

In 1992, Ebone was built on voluntary contributions, and in 1993 the partners established a cost sharing formula based on accees speed and a lightweight decentralized organization was set up. By mid 1993, the partners are negotiating the future as a number of options are emerging: To continue the current organization, to establish a new organizational form, or to transfer responsibility for backbone provision for the research partners to the newly established Dante Ltd. while the commercial partners cooperate on setting up a European CIX.

At any rate, both economical and technical cooperation is needed in Europe in order to secure end user service and connectivity in an environment of expensive international lines and increasingly complex routing requests in the growing Internet.

Session GA: Asia Pacific

chair: Haruhisa Ishida (University of Tokyo, Japan, ishida@u-tokyo.ac.jp)

GAA: Sustainable Development Networks/Concepts & Practice in Australia/Pegasus Network Experience

Ian Peter (Pegasus Networks, Australia, ianp@peg.apc.org)

In the last year activity has substantially increased on establishment of "sustainable development networks". Although this began before the 1992 Earth Summit, the inclusion of the concept in Agenda 21 has given rise to more activity.

This paper tells the story of the establishment of Pegasus Networks in Australia. There is much in the establishment and growth of Sustainable Development Networks which can be learnt from the experiences here, which are now into their fourth year of operation.

The paper is deliberately non-technical, and addresses the management, marketing, and political questions surrounding network development. It is intended as a guide to the important non--technical aspects of devloping networks to serve communities.

GAB: Mule: A MULtilingual Enhancement to GNU Emacs

Ken'ichi Handa (ElectroTechnical Laboratory, Japan, handa@etl.go.jp)
Mikiko Nishikimi (ElectroTechnical Laboratory, Japan, nisikimi@etl.go.jp)
Satoru Tomura (ElectroTechnical Laboratory, Japan, tomura@etl.go.jp)

Mule: MULtilingual Enhancement to GNU Emacs handles not only ASCII characters (7 bits) and ISO Latin-1 (8 bits), but also Greek, Cyrillic, Hebrew, Thai, Vietnamese (8 bits), Japanese, Chinese, Korean (16 bits) coded in the ISO2022 standard and its variants (e.g. EUC, Compound Text).

A text buffer of Mule holds a mixture of characters of these languages. As for inputting, Mule provides various input methods and in case Mule is running under a terminal emulator, whatever the emulator provides can be used.

Mule has extended GNU Emacs's internal structure on many points including buffer, string, character and character-syntax so that EmacsLisp can freely manipulate multilingual text. In addition, Mule converts multilingual text into appropriate codes while handling many types of I/O (e.g. file, process, keyboard, display).

Since Mule is upper compatible with the original GNU Emacs, there's no regression for users of the original even if they move to Mule.

GAC: Building an International Network in Asian-Pacific Rim: the Taiwan Strategy

Chang-Sung Yu (National Taiwan University, Republic of China, yucs@ccms.ntu.edu.tw)

International networking is a global trend toward an information age. Since 1987, a series of networking projects were carried out in Taiwan, and finally an international network, Taiwan Academic Network (TANet), was established.

The author believes that TANet is the most successful research purpose internetwork in the Asia-Pacific Rim. This paper presents the internetworking strategy implemented in Taiwan and analyzes the critical success factors.

Session GB: Central and Eastern Europe

chair: Jan Gruntorad (Czech Technical University Prague, Czech Republik; tkjg@earn.cvut.cz

GBA: Research Networks in the Baltic States

Ants Work (Institute of Cybernetics, Estonia, ants@ioc.ee)

In all three Baltic Countries works have been started to develop the research and education networks, called ESTNET, LATNET and LITNET respectively. A joint project was launched to implement a backbone network on TCP/IP for five university cities (Tartu-Tallinn-Riga-Vilnius-Kaunas). This project and the backbone are called BALTBONE.

There is some communication equipment available and implemented already (Cisco AGS, two IGS-R and five CS-516) and some more will be implemented in the near future (supported by UNESCO). The backbone is partly ready and it is linked via FUNET/NORDUnet to the global Internet from March 1993 using a 64 Kbps link Tallinn-Helsinki. A leased line 9.6 Kbps between Tallinn and Riga is in operation from Oct. 1992 and a 64 Kbps link Tallinn-Tartu on 12 group of analog channels was introduced in March 1993 to substitute the satellite link Tartu- Stockholm, which is in operation from April 1992, but which is several times more expensive.

In the paper the progress of the BALTBONE project is described, a short survey of research and education networks in the Baltic States is given and some major problems are discussed.

GBB: Telecommunications and Networking in Baltic Countries

Algirdas Pakstas (Inst. Math & Informatics, Lithuania, algirdas.pakstas@idt.unit.no)

Last changes in Central and Eastern Europe (CEE) are stimulating factors for development of modern telecommunications and networking infrastructures. Authors intended to review concrete region - Baltic countries. Estonia, Latvia and Lithuania inherited information technologies from the ex-USSR. As consequence of that, at the moment these countries have a lot of similarities both with CEE and with so called "less industrialized countries".

Paper shows existing technological background and telecommunications infrastructures. The average Soviet computer centre is, even now, equipped with technology being ten years behind the rest of the world. In the Baltics heavily dependent on Soviet technology the situation was very similar.

Baltic nations didn't start the building of telecommunication industry from the scratch but definitely telecommunications have been "chronically in poor shape". However, the telephone network in the Baltic countries is the most modern of all ex-USSR republics with Lithuania being used as pilot for the Soviet telecommunication vendors.

In October, 1992 ITU have allocated the following codes: Lithuania=370, Latvia=371, and Estonia=372. Full operation of national numbering zones is expected in 1993.

Research and education institutes in the Baltics are: Academies of Sciences (research and PhD studies) and Universities (education and some research). Main institutes involved in Information Technology area are presented.

Networking activity in Soviet period is reviewed presenting remote terminal facilities via IASNET/VNIIPAS and experience of participation in ACADEMNET project.

Developments of the last years are described in more details. Mailing lists are presented. Current status and networking alternatives are discussed (FIDOnet, Internet-UUCP, EARN/BITNET, X.400), TCP/IP-oriented BALTBONE project is presented.

Current situation can be estimated as more or less acceptable because main research and education centers already have access to e-mail and some of them even Internet on-line services (FTP, TELNET, RLOGIN). But there networks represent only fast and reliable communication media for academic and some business organizations (commodities exchanges). Networks are not yet used directly for educational purposes.

Available statistics shows that traffic of one UUCP node can be even higher that of whole X.400 network and that foreign traffic is main part of academic UUCP-node traffic. So, we can agree with D.Pimenta (Project REDALC) that UUCP is more affordable entry solution for less industrialized countries. The restricting factor for dial-up connected UUCP nodes is high prices for long-distance calls indirectly pressing on the structure of traffic rather than availability of leased lines or its bandwith.

The main reason of current ``status quo'' is shortage of state-wide development programs. All developments are based mostly on enthusiasm of people.

Real network administration problems will arise when networks became as common as in industrial western countries. Also, it should be noted that special educational programs for personnel of newly established networks and popular publications for end-users about available on-line resources are still urgent (sometimes even more needed than donation of hardware and software).

GBC: Networking in the Czech Republic - The CESNET Project

Jan Gruntorad (Czech Technical University, Prague, Czech Republik, tkjg@earn.cvut.cz)

Paper describes networking situation in the Czech Republic where the CESNET (Czech Educational and Scientific Network) was created. The main goal of the CESNET project was to create a multiprotocol network infrastructure in all over the country for providing services of all available academic networks - Internet, EUnet, EARN, HEPNET, etc. CESNET topology, technology (CISCO routers, 64/19,2 kbps lines) and services are described. Activities aiming to enable users from outside the academic community to use network services are also described. The main problem seems to be some legal restrictions still valid in our country.

GBD: SANET: Evolution and Services

Karol Fabian (Institute of Automation and Communication/UAKOM, Slovakia, fabian@uakom.sk)
Lubos Elias (Institute of Automation and Communication, Slovakia, elias@uakom.cs)
Dusan Hruby (Institute of Automation and Communication, Slovakia, hruby@uakom.cs)

The paper presents an overview of academic computer networking in Slovakia. A topology and history of the SANET backbone is presented. Organizational background and funding are mentioned together with the list of important international projects, which have helped to build the network. Services available in SANET are described with a detailed description of X.500 SANET directory structure The new extension of the EUROPAnet project to Slovakia will enhance SANET connectivity. In the frame of this project X.500 infrastructure for the SANET backbone will be provided. Statistics included in the paper were provided and evaluated at Institute of Automation and Communication which serves as the Network Operation Center. The main problems encountered in the building of the infrastructure for SANET are mentioned. Problem causes are discussed. The split of Czechoslovakia has introduced the creation of two new top level domains which will have to exist concurrently with the original .cs domain of the two former states.

GBE: Romanian Computer Network for Education and Research - the First Steps

Paul Cristea (Polytechnic Institute of Bucharest, Romania, pcristea@roearn.ici.ac.ro)
Trandafir Moisa (Politehnica University of Bucharest, Romania, moisa@roearn.ici.ac.ro)

The paper presents a project to build a nation-wide network aimed to link the main universities and research institutions of Romania to the international academic networks. The guiding principle of the project is to design and implement an open system based on OSI standards, without technical barriers in communications between users of different computing hardware. The main effort of the project is targeted to the development of the network infrastructure. Another objective is to support and encourage the development of LANs, offering frame solutions for the Romanian academic institutions. The network topology and the development steps of the Romanian Computer Network (RCN) for Research and Education are analyzed. The global RCN architecture is structured on four levels: the International GateWays, the National Communication Infrastructure for Data Transport and Switched Packets, the Regional WANs and the Local Communication Infrastructure of R&D Units and Universities. The National Data Communication Infrastructure has a ring topology and includes six nodes in universities, located in major towns. A fairly uniform distribution throughout the country has been observed. After completion, the network will link 200 institutions from 25 towns, offering the basic network services to most of the Romanian academic community.

GBF: New Developments and Services in the Hungarian Academic and Research Network

Peter Bakonyi (Hungaria Ltd. - HUNGARNET, Hungary, h25bak@ella.hu)
Laslo Csaba (Hungaria Ltd. - HUNGARNET, Hungary, h26csa@ella.hu)

At the previous two INET conferences overview were given about the first two major development phases of HUNGARNET. In the first phase, based on home made products and services, Hungary could join the networks of the Academic World on the level of e-mail and interactive services. This was done through EUnet and EARN/BITNET. In the second phase the well-known and accepted networking technologies were planned and partly introduced and that was the subject of presentation in Kobe. In this report the current full scale integration to international networks is outlined. Special emphasis is given to Internet related developments in Hungary.

The development of the Hungarian Academic and Research Network has been closely related to a national initiative called "Information Infrastructure Development Program" (IIF). The IIF program was launched in 1986 to build a computer network providing communication and information services to the research, development and higher education community. The first phase (1986-1990) of the IIF program has been successfully implemented. All the major universities, academic institutions and some large companies joined the IIF system, and have been connected to the network.

Now the program is in its second phase (1991-1994).

The objective of the program is the establishment of European level network and information services. In the second phase the network will have a distributed nature, in contrast to the previous centralised home made system. The current concept is the adaptation of the latest results of information technology, the acquisition of up to date HW, SW and network equipments and the full integration of HUNGARNET into the world-wide research networks.

At present there are 450 IIF institutions and about 300 out of them have joined HUNGARNET which is the umbrella organization of the academic community. The major new developments of the last year were the introduction of the Internet technology, the acquisition of the first UNIX based client/server configurations for universities and for research sites and the installation of a FDDI ring at three Universities in Budapest.

GBG: Croatian Academic and Research Network

Predrag Pale (Zagreb University, Croatia, predrag.pale@etf.hr)

The fundamental importance of computer networks for development of a young state has been recognized in Croatia. The Ministry of Science and Technology has initiated Croatian Academic and Research Project (CARNet). The first phase has established the national backbone and international links as well as provided necessary equipment (UNIX workstations, modems, lines, etc.) for new institutions. Thus, 100 institutions (40% of all those financed by the Ministry) and 7000 researchers (70%) are getting international connectivity. The network is built with leased lines carrying TCP/IP protocol and connecting UNIX workstations and LANs of UNIX and MS-DOS computers.

E-mail, FTP, TELNET and Gopher services are designated as basic network services everybody should be able to use. A professional organization, University Computing Center (SRCE) has been assigned with the tasks of running the network and providing user support and information system development.

Besides encouraging intensive use of worldwide available information, development of Croatian databases and information services is supported.

Although academic in its nature, this network should be the model and pilot for developing similar computer networks in Croatia. Until then, it will provide the basic connectivity and services to other segments of Croatian society (government, schools, industry, etc.).

GBH: Networking in Poland

Jerzy Zenkiewicz (Torun University, Poland, jezenk@cc.torun.edu.pl)

The Polish Research and Academic Computer Network called NASK,links nine regional networks and provides international access to european and world networks. Top level country Coordinating Commmitee and Technical Group were established to developed and maintain NASK. During the last two years the rapid development of NASK has provided more than 30 000 users of academic and research community the access to network services. Now the main goal is to create a full network servi- ces for all avaible protocols in NASK:IP,X.25,DECnet, and SNA, where the IP protocol was selected as the basic protocol. In this paper I will briefly describe present status of networking in Poland.Specially,some important is- sues related to architecture,maintaining,development problems and avaible services will be discussed.

GXX: System Management for Large Computer Network: Experience on Design and Creation of the Azerbaijan Republic Information Computer Network

A. Abbasov (Academy of Sciences, Azerbaijan, rkasumov@sovamsu.sovusa.com)
R. Alguliev (Academy of Sciences, Azerbaijan, rkasumov@sovamsu.sovusa.com)
E. Aliyev (Academy of Sciences, Azerbaijan, root@crack.azerbaijan.su)
V. Gasumov

Experience of design and creation of system management (SM) for Republic Information Computer Network (RICN) is described. The common user-technical and software structure of RICN is presented. The tasks on syntheses of SM and application processes interaction on distributed systems are settled. The software structure of SM and its components and technological schemes of distributed data processing on network are described.

The user interface (UI) problems for large distributed data processing system DDPS are investigated. The intellectual and distributed UI structure are presented. The experimental and maintenance results and their analysis for RICN on mentioned problems are presented.

Session GD: Empowering New Users

chair: Steve Ruth (George Mason University, USA, ruth@gmuvax.gmu.edu)

GDA: Personal Computer-Based Communications Training: The Future in African Networking

Gary Garriott (Volunteers in Technical Assistance, USA, gary.garriott@f165.n109.z1.fidonet.org)

Information technology (IT) to date has not lived up to its potential to assist social and economic development processes in Africa. One reason is that foreign consultants introduce computer systems with little long-term stake in what they are being asked to accomplish. They have little interest or incentive to learn about cultural norms and social mores that would influence not only how the machines will be used, but also how to successfully adapt and implement training programs to African audiences. Needed is an "Africanized approach to training and curricu- lum development in electronic networking which emphasizes a practical, "hands-on" experience continuing long after the formal training has ended. Follow-up technical assistance and "seed provisioning" of low-cost hardware and software is essential. Initial selection of student "trainer of trainers" would be done through careful analysis and profiling, Given the ubiquity of personal computers in Africa, a PC-computer literacy module followed by a specially-tailored module for providing low-cost message and file transfer capability using 1) digital radio in terrestrial environments, 2) low earth orbiting satellites, and 3) telephone-based analog modems is proposed. This approach could produce hundreds and perhaps thousands of individuals skilled in low-cost electronic networking tech- niques within five years. Compared with the few people currently available (perhaps a several dozen), such numbers could provide the "critical mass" required for a sustainable indigenous capa- bility, with the future capacity for local innovation and re- invention provided by the Africans themselves.

GDB: Computing under "Socialism": A Retrospective and Perspective

Seymour Goodman (University of Arizona, USA, goodman@misvax.mis.arizona.edu)

We consider the intersection of two global processes: (1) the diffusion of computing, most notably microcomputers and international networks; and (2) the traumatic changes taking place in current or former "socialist" economies. The latter are, or were previously, characterized by: (a) a single or dominant political party; (b) State ownership or very strong forms of control of the "commanding heights" of the economy; (c) strong forms of protectionism; (d) an official revolutionary ideology (e.g., Marxist-Leninist, nationalist, anti-colonial); and (e) systemic political controls on information.

Once considered at the forefront of developing nations, and spanning the alphabetical spectrum from Afghanistan to Zambia, these countries comprise a significant fraction of the world, both geographically and demographically. Many invested heavily in the development, import, and application of technology. A large and comprehensive, but seriously deficient, computing industry was built among an important subset of these countries centered around the former USSR.

Today, the computing communities in most of these countries are undergoing dramatic and often difficult transitions. These transitions are driven by a combination of two sets of factors: (i) responses to past internal economic and policy failures; and (ii) what are in effect externally imposed economic and technological pressures and opportunities, e.g., the existence of rapidly evolving and inexpensive computing resources and unprecedented opportunities for employment by Western companies. Although the diversity of problems and reactions varies greatly over the very large group of countries considered, there are some significant common features emerging across many of these countries with regard to state- supported hardware industries, embryonic software industries, and international technological relations.

GDC: Empowering Low-Bandwidth Users

Laurence Press (UC Dominguez Hills, USA, lpress@isi.edu)

To empower people who are not computer professionals, systems must be conceptually simple and easy to use. Unfortunately, non- professional users are likely to have low-bandwidth connectivity, particularly in less-industrialized nations and regions. One strategy in such cases is asynchronous operation, where a bi- directional transfer of messages occurs when the user logs onto the Net. Seamless integration with the local file system and other tools is also necessary. Three examples of such systems are mentioned, but none are in widespread use on the Internet today.

GDD: Predicting and Facilitating the Success of Developing Nations in the Use of International Networks: Insights and Proposals

Stephen Ruth (George Mason University, USA, ruth@gmuvax.gmu.edu)

This paper offers a brief insight into some of the evaluation problems inherent in the entry of developing nations into the use of networks. Since three fourths of the world's countries have poor connectivity or no connectivity it is helpful to establish some typical predictor variables as well as success metrics to make it easier to find indicators from one country's entry and progress to foretell the results of another's. A number of qualitative and quantitative variables are proposed. To operationalize the variables, data is presented from five developing nations, four with good and improving connectivity (Tunisia, CSFR, Romania and Chile) and one, Pakistan, not yet connected. These data can be a first step in establishing a dialogue concerning predictors for success and techniques for avoiding recurring problems. Pakistan, the least networked of the five nations examined, has the potential to make rapid strides at low unit cost. A proposed approach to more rapid entry of Pakistan into the networking community is offered.

Session GE: European Developments

chair: Tomaz Kalin (RARE, Netherlands, kalin@rare.nl)

GEA: NJE Services to the Users ... and More

Daniele Bovio (EARN Office, France, hi@frors12.bitnet)
Frode Greisen (EARN, Denmark, frode.greisen@uni-c.dk)

The main goal of the recently completed regionalization of EARN (EARN2) was to improve the performance of the whole network reducing the load on the key backbone sites in Europe and exploiting several existing IP lines, both in Europe and between European network organizations (e.g.like NORDUNET and JANET), and the USA. Thanks to the new full mesh structure, achieved in September 1992, and the increased connectivity with the United States the EARN2 plan has proved to be, in fact, extremely effective, in particular from the users' perspective, for several reasons. Round trip time for mail and files, number of average store and forward stations, number of files queued on the backbone links, and time spent by files waiting to be transmitted, are key aspects which largely influence the quality of the service perceived by the users. The statistics collected by the EARN Office show a clear reduction of all the above average values during and after the completion of the regionalization and, for the same period, a general increase of the traffic volume, confirming the effectiveness of the new structure. This improvement is experienced in all 31 connected countries, both in those who have an IP service and those who do not.

The EARN2 backbone is currently monitored and weekly reports covering all the above aspects are produced. An index that takes in account all these different parameters and gives "at a glance" the status of health of the network has also been studied and it is published along with the detailed results.

EARN is nowadays focused both on maintaining the satisfactory quality of service reached, on offering more and new services on top of the network layers, and facilitating the mail/file exchange with other networks. The global connectivity with the Internet is being improved. 9 INTERBIT gateways (for mail purposes) are already operational in Europe. Also a set of servers (BITFTP) that allow NJE and mail-only users to FTP files from Internet-only sites with simple commands will be made available during 1993. Brand new services in the users support area have also been started: a net-accessible help desk, run by EARN Staff and other network experts is available, and an impressive amount of documentation about the main network resource tools has been published for the users.

GEB: Operational Network Services for the European Research Community

Howard Davies (University of Exeter, UK, United Kingdom, h.e.davies@exeter.ac.uk)

The history of network services for European researchers demonstrates a complex, but progressive rationalization which still leaves much scope for further progress. This paper provides a description the current situation and its evolution. It concentrates on the activities of RARE and on the role of the Operational Unit recently re-christened DANTE which has been established as a not-for-profit company by the national networks of 12 Western European countries to provide international services both within Europe and to North America. The planned service portfolio of DANTE and other current developments are also described.

GEC: SuperJANET Applications

John Dyer (Joint Network Team, United Kingdom, j.dyer@jnt.ac.uk)

SuperJANET is the new UK academic high speed network based on fibre optic technology. The installation of the network which started early in 1993 has given the UK academic community an opportunity to pilot new applications.

The application areas that have been identified for piloting are : Distributed Multimedia Teaching, Group Communications, Information Services, Remote Access, Remote Consultation and Supercomputer Support including Visualisation.

The paper reviews the areas in which application piloting is currently taking place and describes some of the pilot projects in detail.

The SuperJANET pilot applications should be regarded as new wave applications, NOT merely new versions of old applications with "go faster stripes". They are designed to take advantage of the increase in network bandwidth and reductions in latency and jitter.

The paper will review future directions and invite discussion on users expectations and desires.

GED: EBONE, A Technical Overview

Bernhard Stockman (SUNET/NORDUnet/EBONE Network Op Ctr, Sweden, boss@ebone.net)

The requirement for Europe-wide network connectivity at increasing bandwidth and for high-speed connectivity to non-European networks has long been recognized. In recent years, this demand has focused on IP services. Until late 1991 such connectivity was available only to individual national and international research networks. Ebone is an initiative to promote close collaboration among today European network service providers to provide a general purpose pan-European backbone.

The paper will focus on the technical details of the EBONE backbone. It will give a functional description including item such as services, topology and the routing implementation. The paper continues with a description of the EBONE intercontinental connectivity and how this relates to the Global Internet eXchange (GIX). A description of the management and operations of the backbone will also be given together with some traffic flow reports. The paper will finally describe how the EBONE architecture is planned to evolve in the near term with respect to changes in earlier EBONE Boundary System (EBS) definitions and why these changes where necessary, changes in the internal and external routing protocols, implementation of CIDR and full routing and a method for improved interconnectivity between European backbone providers at large.

Session GF: Latin America

chair: Xavier Baquero (EcuaNet, Ecuador, xbaquero@ecnet.ec)

GFA: Networking Latin America and the Caribbean: Creating Alternatives

Luis German Rodriguez Leal (Universidad Central De Venezuela, Venezuela, lgrodrig@conicit.ve)

There is an ongoing process of integration of academic networks in Latin American and Caribbean countries. This papers analyzes how different factors have shaped this process and explains the main conclusions of a recent meeting in Guadalajara, Mexico (November 1992). This process has gone through several stages of maturity for over four years and it can be said that it has not followed patterns of development observed in other regions of the world. The evolution of the effort for connecting and organizing the networks of the region is seen through the results of the different meetings where their actors and promoters have confronted their goals against the facts that condition them. There have been five regional meetings devoted to this goal (from San Jose, Costa Rica, in 1989 to Guadalajara, Mexico, in 1992). The stages can be associated with several factors, one of them is the development of the national initiatives; some countries have more than one network pretending to cover the academic community and in many cases these initiatives have conflicts among them. Another factor is the support of that these initiatives have from their national organizations for science and technology; some are officially backed by their governments while others are still fighting to be recognized or at least to be considered as a helpful tool for the academic sector. The partipation of organizations external to the region and interested on promoting the integration process (Organization of American States, FUNDESCO, National Science Foundation, UNESCO, PNUD, etc.) has also played a decisive role in it. The initiatives of the region led to the decision, at the last meeting of academic networks for Latin America and the Caribbean, to create an open forum devoted to monitoring theprocess. The evaluation will be done during the annual gathering. The achiement of the defined goals will be checked against the development of specific tasks associated with the established strategies.

GFB: Non-Commercial Networking in Brazil

Michael Stanton (Pontificia Universidade Catolica do Rio de Janeiro, Brazil, michael@inf.puc-rio.br)

As a country of continental dimensions, Brazil presents interesting challenges to the establishment of telecommunications and computer networking infrastructure. The article briefly describes the present telecommunications infrastructure, and explains how this was used since 1988 to set up a number of different non-commercial networks integrating Brazilian organizations and providing them with access to international networking.

The article goes on to describe the creation of a "second generation" network, based on Internet technology, which began experimental operation in 1991, and which currently provides the principal non-commercial service in the country. This is based on a number of distinct regional networks, usually supported by the respective state governments, and the National Research Network set up by the national (federal) government, which provides national backbone connectivity. Brief mention is made of the organizational and support infrastructure established for the national network, as well as some of the network services and the information and computing resources currently or soon to be made available.

The article concludes with a look at present and future international connectivity of the Brazilian network, and also at the perpectives for the introduction of broadband networking in the short and medium term, as a result of current and planned improvements in telecommunications infrastructure.

GFC: The Chilean Internet Connection or I Never Promised You a Rose Garden

Patricio Poblete (Universidad de Chile, Chile, pvpoblete@daisy.uwaterloo.ca)
Ricardo Baeza-Yates (Universidad de Chile, Chile, rbaeza@dcc.uchile.cl)
Jose Piquer (Universidad de Chile, Chile, jpiquer@dcc.uchile.cl)

Chile has had two international 56Kbps links to the Internet since January 1992. This online connection to the world is having a great impact on academic research, extending state-of-the-art communication technology to our country. The impact will be even greater considering the traditional isolation of the country, surrounded by mountains and sea at the end of the world, and a very particular geography: almost 4350 kilometers long from north to south with an average width of 190 kilometers.

However, some problems (technical and political) have also been generated by this connection. For instance, the two links are operated by independent organizations, and there is no local connection between their respective networks. Some of the problems we discuss are unique to Chile, but many may also arise in other countries connecting to the Internet.

This paper presents an analysis and a critique of the solutions implemented in the Chilean networks, focused mainly on the pricing, accounting and billing of the international link. This problem is quite similar to the NSF backbone pricing problem: a shared and expensive resource. We propose some new viewpoints and improved long-term solutions. In particular, we strongly believe that the current situation is not sustainable, and we advise against using a similar solution anywhere else, as we explore better alternatives.

We also discuss the trend towards widespread use of this technology in Chile, and the steps that are being taken to make it compatible with commercial use.

Session GG: Africa

chair: Ahmed Abdel-Bassit (Estinet, Egypt, estinet!ab@relay.eu.net)

GGA: Lessons in Network Building in Africa: The Need for Development of Network Supportive Infrastructure

Lishan Adam (UNECA/PADIS ADDIS ABABA, Ethiopia, lishan@padis.gn.apc.org)

Africa lacks both the infrastructure for and environment conducive to network building. Physical (telecommunications, institutional, and computing), information, marketing and human resource infrastructures are all too weak to support network set up. The current socio-economic, cultural and political situation of the region make networking difficult. Factors such as communication policies, tariff and traffic structures are the main stumbling blocks to building national and regional communication networks in the region.

Experiences in the Africa region indicate a need for a systematic approach to network building and a need for portable low cost networks with the ability to global connectivity. Sustainability and localization are other issues to be resolved in order for the current African network map to become much more denser. The major factors determining network building in the region include:

- the approach to be followed in introduction of networking (bottom-up, phased, modular, etc.) based on the socio-economic situation of the region;

- network technology (protocol, platform, etc.) depending on the available technology for development;

- the supportive infrastructure for network building.

This paper outlines these major issues in assessing the situation in the region, recommends solutions and discusses a practical project.

GGB: Academic Networking in Tunisia

Nejib Abida (Regional Institute of Science Info.& Tel, Tunisia, abida@tnearn.bitnet)

Computer communications networks is part of the new information revolution made possible by new developments in information technology. It is born out of the growing need for new capabilities for the generation, processing, transmission and display of data required for more effective management of modern organisations. Often the time available for making critical decisions has tended to diminisch while the consequences of errors have grown more severe,making on-line real-time data processing and communication and system reliability critical futures of computer communications networks.

These futures pose great demand for complex telecommunications infrastructure requisite technical and managerial manpower, new organisational structure and many other technology support systems, which, through largely taken for granted in the developed countries, are at different levels of development in various developping countries. The situation in Tunisia today is adequate for the computer communication networks development and the academic network is start-UP !

The development of the Tunisian Academic and Research network has been closely interconnected with the initiative of IRSIT effort to set up EARN and Internet nodes. Today the services of those nodes is at the disposal of more then 20 institutions and abaout 200 users. The politiacl system in Tunisia is today encourage to develop a national wide area network and the first phase of the RNRT (Reseau National pour la Recherche et la Technologie) is on going... The RNRT should be based on the public switched data network which is providing a volume charged service for the institutions. The number of the institutions connected to the packet network is over 50. More than 12 LAN will be supported by the packet network and we expect severel tens interconnected PC-s as well. The introduction of the INTERNET technology has been decided, therefore an IP router must be installed for each regional LAN, and a dedicated IP backbone network will be established to carry the high volume IP traffic while the low volume will be transmitted over X.25. Tunisia would like to be connected to the European backbone networks like IXI, EBONE 92) and run X.400, OSI CLNS.

Research centers on the RNRT:

8 research centers are to be built in the first phase of the RNRT. The basic configurations of the centers are built around one or several LANs each of which should consist of one or two servers, 5-10 worksations and 10-20 terminals. The centers will be Internet networks. The first installation will be financed by the research ministery). Optical cable based CAMPUS networks are being built. ETHERNET will run over the cables, and the first three FDDI's are being installed as well.

INET '93 Organizing Committee

Conference Committee Chairs

General Conference: Eric Benhamou (3COM Corp, USA, eric_benhamou@3mail.3com.com)
Program Committee: Barry Leiner (USRA, USA, leiner@nsipo.nasa.gov)
William Yundt (Stanford University, USA, gd.why@forsythe.stanford.edu)
Local Arrangements: David Ernst (UCSF, USA, dernst@ucsfvm.ucsf.edu)
Janet Perry (Novell, USA, janetp@novell.com)
Workshop: George Sadowsky (NYU, USA, sadowsky@nyu.edu)
Conference Manager: Elizabeth Barnhart (Educom, USA, barnhart@educom.edu)
ISOC Liaison: Michael Roberts (Educom, USA, roberts@educom.edu)
Larry Landweber (University of Wisconsin, USA, lhl@cs.wisc.edu)

Track Leaders

Network Technology: Christian Huitema (Christian Huitema INRIA, France, christian.huitema@sophia.inria.fr)
Phillip Gross (Advanced Network Services, Inc., USA, pgross@nis.ans.net)
Network Engineering: Elise Gerich (Merit Network, Inc., USA, epg@merit.edu)
Francois Fluckiger (CERN, Switzerland, fluckiger@vxcrna.cern.ch)
Application Technology: Eric Hood (North WestNet and FARNET, USA, ehood@nwnet.net)
John Hine (Victoria University of Wellington, New Zealand, john.hine@comp.vuw.ac.nz)
User Applications: Laura Breeden (FARNET, USA, breeden@farnet.org)
Rolf Nordhagen (University of Oslo, Norway, rolf.nordhagen@usit.uio.no)
Policy Issues: Haruhisa Ishida (University of Tokyo, Computer Centre, Japan, ishida@u-tokyo.ac.jp)
Steven Goldstein (National Science Foundation, USA, goldstein@note.nsf.gov)
Regional Issues: Richard Mandelbaum (Polytechnic University of New York, USA, rma@tsar.poly.edu)
Tomaz Kalin (RARE, Netherlands, kalin@rare.nl)

Regional Coordinators

Ahmed Abdel-Bassit (Estinet, Egypt, estinet!ab@relay.eu.net)
Xavier Baquero (Ecuanet - Corporacion Ecuatoriana de Informacion, Ecuador, xbaquero@ecnet.ec)
Klaus Birkenbihl (Center for Science Networks, Germany, klaus.birkenbihl@gmd.de)
John Demco (The University of British Columbia, Canada, demco@cs.ubc.ca)
Geoff Huston (AARNet, Australia, g.huston@aarnet.edu.au)
Haruhisa Ishida (University of Tokyo, Computer Centre, Japan, ishida@u-tokyo.ac.jp)
Valery Udalov (Latvian SSR Academy of Sciences, Latvia, ms@mvm.riga.lv)

Advisory Committee

David Brandin (Interop Company, USA, dbrandin@interop.com)
Lee Caldwell (Novell, USA, lcaldwel@novell.com)
Michael Carter (Apple, USA, mickey@applelink.apple.com)
Vinton Cerf (Internet Society/CNRI, USA, vcerf@cnri.reston.va.us)
David Ernst (UCSF, USA, dernst@ucsfvm.ucsf.edu)
Dan Lynch (Interop, USA, dlynch@interop.com)
Michael Majdalany (Sun, USA, michaelm@corp.sun.com)
Richard Mandelbaum (Polytechnic University of New York, USA, rma@tsar.poly.edu)
Janet Perry (Novell, USA, janetp@novell.com)
Vincenzo Puliatti (United Nations Developement Program, Peru, enzo@undp.org)
Robert Shahan (IBM, USA, shahan@rhqum21.ibm.com)
Katherine Webster (Sun, USA, katherine.webster@corp.sun.com)

Administrative Support

Paul Baer (Stanford University, USA, baer@jessica.stanford.edu)
Jo-Anne Scott (Stanford University, USA, scott@mordor.stanford.edu)
Sunnie Simmons (3COM Corp., USA, sunnie_simmons@3mail.3com.com)
Leslie Todd (USRA, USA, todd@nsipo.nasa.gov)