Proceedings of the Fifteenth Internet Engineering Task Force University of Hawaii

October 31 - November 3, 1989

Compiled and Edited
by
Phillip G. Gross
Gregory M. Vaudreuil

Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, Virginia 22091

ACKNOWLEDGEMENTS

The Fifteenth IETF session was held at the University of Hawaii and was hosted by Torben Neilson. The University of Hawaii proved to be not only a beautiful location, but conducive to a remarkably productive meeting.

On behalf of the IETF and the Internet community, I would like to thank Torben Neilson and his dedicated staff, for the support and excellent facilities at the University of Hawaii.

Thank you to Monica Hart (NRI) for her dedicated efforts to incorporate the exploding number of working groups and the new Internet Engineering Steering Group into these Proceedings. I offer a heartfelt thank you to Karen Bowers (NRI) who has given me considerable support during my transition into NRI and the Internet community.

Gregory M. Vaudreuil/NRI

Contents

Chairman's Message

This is an exciting time for the IETF. Since the announcement of the IAB reorganization at the Stanford University IETF meeting in July and the formation of the IETF Steering Group (IESG), the number of working groups in the IETF has risen from 22 to approximately 40. I believe this reflects both the increased attention to pressing concerns offered by the IESG and the newly revitalized IAB. This also demonstrates a change in emphasis towards smaller more focused groups. During this period we have also completed and published as RFCs two important protocol efforts - the OSPF Routing Protocol, RFC1131, and the Point-to-Point serial line Protocol (PPP), RFC1134.

Working Group Status Reporting

In the past, we had three separate documents for each WG - the charter, the status report, and the current meeting report. We are now eliminating the separate status report. In the future, information about the WG chair(s) and mailing lists will be included in the charter, and the progress-to-date will be folded into the Current Meeting Report.

To better track working group progress, we are also revising the WG "chartering" process to include more specific milestones and deadlines. We can look again to the leadership of Dave Crocker (Network Management Area Director) and Marshall Rose for being on the forward edge. Please see the new charter of the SNMP Working Group, chaired by Marshall Rose, as an example of the direction we are pursuing. We will be asking new working groups to supply this type of additional detail by following a new charter format.

Security and Applications Area Directors

I would like to welcome two new important members to the IESG. Steve Crocker of Trusted Information Systems (TIS) joined us in November as the Director of the Security Area, and Russ Hobby (UC-Davis) has joined us more recently as the Applications Area Director. Among his other projects at TIS, Steve is involved with developing a secure email system based on RFCs 1113-1115. Russ has already proposed some specific application projects. Both Steve and Russ will be at the next IETF meeting. This leaves only the Operations area unfilled. Until filled, I will continue to serve as the interim director.

IETF at Florida State University (February 6-9, 1990)

The next IETF meeting is at Florida State University at Tallahasse on February 6-9, 1990. The local host is Ken Hayes, and the meeting is partly sponsored by the Department of Energy.

There will be an open meeting of the IESG at the February IETF. At the suggestion of Mike Karels at the last IETF meeting, we have scheduled the IESG from 4-7pm on Thursday so it does not conflict with other WG sessions or the technical presentations. The primary topic of this open IESG meeting will be the important issue of intra-AD routing protocol (i.e., IGP) standardization.

Logistics information about upcoming IETF meetings is announced on the ietf@isi.edu mailing list. To be added to that mailing list, send a request to ietf-request@isi.edu. Information on hotel and travel, preliminary agenda, working groups, and draft documents relevant to upcoming IETF meetings are also available online at NIC.DDN.MIL and NNSC.NSF.NET. See Section 1.3 "Online IETF Information" in the IETF Overview in these Proceedings.

Welcome to Greg Vaudreuil

I'm sad to report that Karen Bowers has left the IETF effort. Counteracting that bad news is the good news that she has taken on a new important project at NRI. Karen was with us during a crucial time of rapid growth and change in the IETF. We wish her well in her new duties at NRI. Karen has been replaced by Greg Vaudreuil.

Greg is already making his presence felt. He is continuing Karen's diligent monitoring of working group activity, and, with the IESG, he has begun to reorganize how we track and report WG activity. Greg is adopting standard abbreviations for all IETF working groups, and setting up working group mail reflectors at NRI using the standard abbreviation. In most cases, these standard working group mail reflectors will simply forward to a separate list maintained by the working group chair. However, this provides an easy way for newcomers to join working group discussions. In these Proceedings, the abbreviation will appear next to the name of the working group. These abbreviations will also be used in constructing filenames in the on-line IETF and Internet-Drafts directories at NNSC and the NIC.

Phillip G. Gross
Chairman, Internet Engineering Task Force

Final Agenda of the Fifteenth IETF

(October 31 - November 3, 1989)

TUESDAY, OCTOBER 31

9:00 am - 12:00 pm	MORNING WORKING GROUP SESSIONS • OSI X.400 (Rob Hagens/UWisc) • Open Distance Vector Routing (Charles Hedrick/Rutgers) • Alert Management (Louis Steinberg/IBM) • Connection IP (Claudio Topolcic/BBN) • User Documentation (Karen Roubicek/BBN, Tracy LaQuey/UTexas)
12:00 pm	IESG and WG Chair Working Lunch
1:00 pm - 4:00 pm	AFTERNOON WORKING GROUP SESSIONS • OSI X.400, Domain Name System (Joint Meeting) • Connection IP (Claudio Topolcic/BBN) • Point-to-Point Protocol Extensions (Russ Hobby/UCDavis, Phill Gross/NRI) • User Services (Karen Bowers/NRI and Craig Partridge/BBN) • Network Management Services Interface (Oscar Newkerk/DEC)
4:15 pm - 5:30 pm	TECHNICAL PRESENTATIONS • "Hyper MIB Demonstration", Steve Hunter/LLNL (15 minutes) • "The CERT", Richard Pethia/CMU (1 hour)
5:30 pm	RECESS

WEDNESDAY, NOVEMBER 1

9:00 am - 9:15 am — TECHNICAL PRESENTATION

- "Internet Status Report", Zbigniew Opalka/BBN (15 minutes)

9:15 am - 12:00 pm — MORNING WORKING GROUP SESSIONS

- Internet User Populations (Craig Partridge/BBN)
- OSI General (Rob Hagens/UWisc, Ross Callon/DEC)
- Dynamic Host Configuration (Ralph Droms/NRI)
- Open Systems Routing (Marianne Lepp/BBN)
- Point-to-Point Protocol Extensions (Russ Hobby/UCDavis, Phill Gross/NRI)
- NOC Tools (Bob Enger/Contel, Bob Stine/Sparta)
- Interconnectivity (Guy Almes/Rice)

1:00 pm - 4:00 pm — AFTERNOON WORKING GROUP SESSIONS

- TCP Large Windows (Craig Partridge/BBN)
- OSI General (Ross Callon/DEC, Rob Hagens/UWisc)
- User Documentation, NOC Tools (joint meeting)
- Connection IP (Claudio Topolcic/BBN)
- Dynamic Host Configuration (Ralph Droms/NRI)
- Open Systems Routing (Marianne Lepp/BBN)
- Interconnectivity (Guy Almes/ Rice)
- IP over FDDI (Dave Katz/ Merit)

4:15 pm - 5:30 pm — TECHNICAL PRESENTATIONS

- "Selective Binary Scheme for Congestion Avoidance", K.K. Ramakrishnan/DEC (1 hour)
- "ESnet Status Report", Tony Hain/LLNL (15 Minutes)

5:30 pm — RECESS

THURSDAY, NOVEMBER 2

9:00 am - 9:15 am — TECHNICAL PRESENTATION

- "NIC Update", Mark Lottor/SRI (15 minutes)

9:15 am - 12:00 pm — IETF STEERING GROUP MEETING

9:15 am - 12:00 pm — MORNING WORKING GROUP SESSIONS

- PDN Routing (Open session)
  (CH Rokitansky/Fern University of Hagen)
- Telnet (Dave Borman/Cray)
- Connection IP (Claudio Topolcic/BBN)
- Open Systems Routing (Marianne Lepp/BBN)
- Interconnectivity (Guy Almes/Rice)
- Domain Name System (Paul Mockapetris/USC-ISI)
- OSI Internet Management (Lee LaBarre/MITRE)
  (formerly CMIP over TCP)

12:00 pm — IESG and WG Chair Working Lunch

1:00 pm - 5:30 pm — TECHNICAL PRESENTATIONS

- "Talking Roads and Networked Cars",
  CH Rokitansky/Fern University of Hagen
- Pacific Rim Interconnectivity
  - "PACOM and Hawaii: Present and Future Plans", Torben Nielsen/University of Hawaii
  - "Agency Requirements in the Pacific Rim",
    Milo Medin/NASA and Tony Hain/LLNL
  - "The Australian Academic and Research Network (AARN)" Geoff Huston/AARN
  - "Internetworking in the South Pacific" ,
    Robert Elz/University of Melbourne
  - "Internetworking in Japan and the North Pacific",
    Jun Murai/University of Tokyo
- "White Pages Pilot Program",
  Marshall Rose/NYSERNet
- "NSFnet Status Report", Bilal Chinoy/Merit
- "Routing and Fair Pricing in Internets
  with Packet Loss", Vlad Rutenburg/SRI

5:30 pm — RECESS

FRIDAY, NOVEMBER 3

9:00 am - 11:30 am	WORKING GROUP AREA AND SELECTED WORKING GROUP PRESENTATIONS
	• Network Management (Dave Crocker/DEC) – SNMP (Marshal Rose/Nysernet) – OSI Internet Management (Lee LaBarre/MITRE) • Host and User Services (Craig Partridge/BBN) • Internet Services (Noel Chiappa/Proteon) – Point-to-Point Protocol (Russ Hobby/UCDavis) • OSI Interoperability (Rob Hagens/UWisc and Ross Callon/DEC) • Routing (Bob Hinden/BBN) • Applications (Phill Gross/NRI) • Operations (Phill Gross/NRI) • Security (Phill Gross/NRI)
11:30 am - 12:00 pm	CONCLUDING REMARKS (Phill Gross/NRI)
12:00 pm	ADJOURN

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) is a large open community of network designers, operators, vendors and researchers concerned with the smooth operation and evolution of the Internet. There is no formal membership in the IETF. The work is done by individuals who share an interest in the resolution of particular problems.

The IETF mission includes:

- Responsibility for specifying the short and mid term Internet protocols and architecture for the Internet Activities Board.
- Identification of pressing and relevant short to mid range operational and technical problem areas and convening of Working Groups to explore solutions.
- Provision of a forum for the exchange of information within the Internet community.

The IETF is organized around eight technical areas. Within each area, technical issues are addressed by working groups. Each area is led by an area director. Overall guidance of the IETF is provided by the IETF Steering Group (IESG). The IESG is composed of the area directors and the Chair of the IETF. The IESG has the general responsibility for making the Internet operate smoothly by identifying and resolving the short and mid term issues and problems. Each area director has primary responsibility for one area of IETF activity.

The current areas and directors are:

IETF and IESG Chair:	Phill Gross/ NRI
Applications:	Russ Hobby/ UC-Davis
Host and User Services:	Craig Partridge/ BBN
Internet Services:	Noel Chiappa/ Consultant to Proteon
Routing:	Robert Hinden/ BBN
Network Management:	Dave Crocker/ DEC
OSI Interoperability:	Rob Hagens/ U-Wisc and Ross Callon/ DEC
Operations:	Phill Gross/ NRI (interim)
Security:	Steve Crocker/ TIS

The work of the IETF is conducted in Working Groups, each of which is convened to solve a particular problem, work on an enhancement or exchange information vital to the operation of the Internet. There are currently over thirty working groups. The working groups conduct business via electronic mail on mailing lists established for each group and during plenary meetings of IETF and other meetings. Information about current activities is distributed by the IETF mailing list. Send a request to join the list to ietf-request@isi.edu.

The IETF holds quarterly plenary sessions composed of working group sessions, technical presentations and network status briefings. Proceedings are produced for each quarterly IETF meeting. This document is the Proceedings of the Fifteenth IETF.

1.1 IETF Working Group Summary (by Area)

HOST AND USER SERVICES AREA
Craig Partridge/BBN
craig@bbn.com

Dynamic Host Configuration (dhc) — Ralph Droms/NRI, rdroms@nri.reston.va.us

WG mail: host-conf@rutgers.edu
Status: Continuing, met Nov. 1, 1989
Drafts or Recent RFCs:

"Dynamic Configuration of Internet Hosts", Ralph Droms/Bucknell University
Internet-Draft: <draft-ietf-dhc-problem-stmt-00.txt> August 1989

Internet User Population (iup) — Craig Partridge/BBN, craig@bbn.com

WG mail: ietf@venera.isi.edu
Status: First meeting Nov. 1, 1989

TCP Large Windows (tcplw) — Craig Partridge/BBN, craig@bbn.com

WG mail: ietf@venera.isi.edu
Status: First meeting Nov. 1, 1989

User Documents (userdoc) — Karen Roubicek/BBN (NNSC), roubicek@nnsc.nsf.net; Tracy LaQuey/U of Texas, Tracy@emx.utexas.edu

WG mail: user-doc@nnsc.nsf.net
Status: Continuing, met Nov. 1, 1989

User Services (uswg) — Joyce Reynolds/ISI, jkrey@venera.isi.edu

WG mail: us-wg@nnsc.nsf.net
Status: Continuing, met Oct. 31, 1989

INTERNET SERVICES AREA

Noel Chiappa/Consultant/Proteon
jnc@ptt.lcs.mit.edu

Connection IP (cip) — Claudio Topolcic/BBN, topolcic@bbn.com
WG mail: cip@bbn.com
Status: Continuing, met Nov. 1, 1989

Router Discovery (rdisc) — Steve Deering/3com, deering@pescadero.stanford.edu
WG mail: gw-discovery@gregorio.stanford.edu
Status: New Group

IP MTU Discovery (mtudisc) — Jeff Mogul/DEC, mogul@decwrl.dec.com
WG mail: mtudwg@decwrl.dec.com
Status: New Group

IP over FDDI (fddi) — Dave Katz/Merit, dkatz@merit.edu
WG mail: FDDI@merit.edu
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

RFC1103: "Proposed Standard for the Transmission of IP Datagrams..", D. Katz/Merit, June 1989

"The Transmission of IP Datagrams over FDDI Networks" edited by Dave Katz for the IP over FDDI Working Group, January 1990
Internet-Draft: <draft-ietf-fddi-ipdatagrams-00.txt>

Performance and Congestion Control (pcc) — Allison Mankin/MITRE, mankin@gateway.mitre.org
WG mail: ietf-perf@gateway.mitre.org
Status: Concluding
Draft or Recent RFC:

"Gateway Congestion Control Policies", A.J. Mankin/MITRE and K.K. Ramakrishnan/DEC, July 1989
Internet-Draft: <draft-ietf-perfcc-gwcc-00.txt>

Point-to-Point Protocol (ppp)
WG mail: ietf-ppp@ucdavis.edu
Status: Concluded

Drew Perkins/CMU
ddp@andrew.cmu.edu
Russ Hobby/UC Davis
rdhobby@ucdavis.edu

Draft or Recent RFC:

"Requirements for an Internet Standard Point-to-Point Protocol (PPP)",
Drew Perkins/CMU, June 1989,
Internet-Draft: <draft-ietf-ppp-req-00.txt>

RFC1134: "The Point-to-Point Protocol (PPP), A Proposed Standard ..",
Drew Perkins/CMU, December 1989

Point-to-Point Protocol Extention (pppext)
WG mail: ietf-ppp@ucdavis.edu
Status: First meeting Nov. 1, 1989
Draft or Recent RFC:

Russ Hobby/UC Davis
rdhobby@ucdavis.edu

"The Point-to-Point Protocol (PPP) Initial Configuration Options",
Drew Perkins/CMU, November 1989, Internet-Draft: <draft-ietf-ppp-options.txt>

IP over Switched Megabit Data Service (ipsmds)
WG mail: ip-smds@nri.reston.va.us
Status: New Group

George Clapp/Ameritech
meritec!clapp@bellcore.bellcore.com
Mike Fidler/Ohio State
ts0026@ohstuma.ircc.ohio-state.edu

Router Requirements (rreq)
WG mail: r-req@nri.reston.va.us
Status: New Group

Jim Forster/cisco
forster@cisco.com
Philip Almquist
almquist@jessica.stanford.edu

ROUTING
Bob Hinden/BBN
hinden@bbn.com

Interconnectivity (iwg) — Guy Almes/Rice, almes@rice.edu

WG mail: iwg@rice.edu
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

RFC1105: "Border Gateway Protocol (BGP)",
K. Lougheed/cisco, Y. Rekhter/IBM, June 1989

Open Shortest Path First IGP (ospf) — Mike Petry/UMD, petry@trantor.umd.edu; John Moy/Proteon, jmoy@proteon.com

WG mail: ospfigp@trantor.umd.edu
Status: Concluded
Draft or Recent RFC:

RFC1131: "OSPF Specification", John Moy/Proteon
Oct. 1989

Open Systems Routing (orwg) — Marianne Lepp/BBN, mlepp@bbn.com

WG mail: open-rout-interest@bbn.com
Status: Continuing, met Nov. 1, 1989

Open Distance Vector IGP (odv) — Charles Hedrick/ Rutgers, hedrick@aramis.rutgers.edu

WG mail: odvigp@rutgers.edu
Status: First meeting Nov. 1, 1989

ISIS for IP Internets (isis) — Ross Callon/DEC, callon@erlang.dec.com

WG mail: isis@merit.edu
Status: New Group

"Use of OSI IS-IS for Routing in TCP/IP and Dual Environments"
edited by Ross W. Callon for the ISIS for IP Internets WG,
January 1990 <draft-ietf-isis-spec-00.ps>

Private Data Network Routing (pdnrout)
WG mail: pdn-wg@bbn.com
Status: Continuing Work
Draft or Recent RFC:

CH Rokitansky/Fern Uni-Hagen
roki@isi.edu

"Internet Cluster Addressing Scheme", August 1989,
Internet-Draft: <draft-ietf-pdn-clusterscheme-00.txt>

"Application of the Cluster Addressing Scheme to X.25 Public Data Networks...", August 1989,
Internet-Draft: <draft-ietf-pdn-pdncluster-00.txt>

"Assignment/ Reservation of Internet Network Numbers for the PDN-Cluster", August 1989,
Internet-Draft: <draft-ietf-pdn-pdnclusternetassignm-00.txt>

NETWORK MANAGEMENT AREA
Dave Crocker/DEC
dcrocker@nsl.dec.com

Alert Management (alertman)
WG mail: alert-man@merit.edu
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

Louis Steinberg/IBM
louiss@ibm.com

"Managing Asynchronously Generated Alerts", September 1989,
Internet-Draft: <draft-ietf-alertman-asyncalertman-01.txt>

OSI Internet Management (oim)

Lee LaBarre/MITRE
cel@mbunix.mitre.org
Brian Handspicker
bd@vines.dec.com

WG mail: netman@gateway.mitre.org
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

RFC1095: "Common Management Information Services and Protocol.. (CMOT)".
April 1989

Management Services Interface (msi)
WG mail: MSI@nri.reston.va.us
Status: New Group

Oscar Newkerk/DEC
newkerk@decwet.enet.dec.com

LAN Manager (lanman)
WG mail: lanmanwg@spam.istc.sri.com
Status: Continuing Work

Jim Gruel/HP
jimg@hpcndpc.cnd.hp.com

NOC-Tools (noctools)
WG mail: noctools@merit.edu
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

Bob Enger/Contel
enger@sccgate.scc.com
Bob Stine/Sparta
stine@sparta.com

"A Network Management Tools Catalog: Tools For ..",
December 1989, Internet-Draft: <draft-ietf-noctools-debugging-02.txt>

SNMP (snmp)

WG mail: snmp-wg@nisc.nyser.net

Status: Continuing, met Nov. 1, 1989

Draft or Recent RFC:

Marshall T. Rose/Nysernet
mrose@cheetah.nyser.net

"Management Information Base for Network Management of TCP/IP...", September 1989, Internet-Draft: <draft-ietf-snmp-mib2-01.txt>

OSI INTEROPERABILITY AREA

Ross Callon/DEC
callon@erlang.dec.com
Rob Hagens/UWisc
hagens@cs.wisc.edu

OSI General (osigen)

Rob Hagens/UWisc
hagens@cs.wisc.edu
Ross Callon/DEC
callon@erlang.dec.com

WG mail: ietf-osi@cs.wisc.edu
Status: Continuing, met Nov. 1, 1989
Draft or Recent RFC:

"An Echo Function for ISO 8473", October 1989,
Internet-Draft: <draft-ietf-osi-iso8473-00.txt>

OSI-X.400 (osix400)

Rob Hagens/UWisc
hagens@cs.wisc.edu

WG mail: ietf-osi@cs.wisc.edu
Status: Continuing, met Nov. 1, 1989

SECURITY AREA

Steve Crocker/TIS
crocker@tis.com

IP Authentication (ipauth)

Jeff Schiller/MIT
jis@athena.mit.edu

WG mail: awg@bitsy.mit.edu
Status: Continuing Work
Draft or Recent RFC:

"The Authentication of Internet Datagrams", August 1989,
Internet-Draft: <draft-ietf-auth-ipauthoption-00.txt>

SNMP Authentication (snmpauth)

Jeff Schiller/MIT
jis@athena.mit.edu

WG mail: awg@bitsy.mit.edu
Status: New Group

OPERATION
Phill Gross/NRI (Interim)
pgross@nri.reston.va.us

Joint Monitoring of Adjacent NSFnet Networks (jomann) — Susan Hares/Merit, skh@merit.edu
WG mail: njm@merit.edu
Status: Continuing Work

Benchmarking Methodology (bmwg) — Scott Bradner/Harvard, sob@harvard.edu; Mick Scully, mcs@ub.com
WG mail: bmwg@harvisr.harvard.edu
Status: New group

APPLICATIONS
Russ Hobby/UCDavis
rdhobby@ucdavis.edu

Domain Name System (dns) — Paul Mockapetris/USC-ISI, pvm@isi.edu
WG mail: namedroppers@nic.ddn.mil
Status: Continuing Work
Draft or Recent RFC:

RFC1101: "DNS Encoding of Network Names and Other Types", April 1989

TELNET (telnet) — Dave Borman/Cray, dab@cray.com
WG mail: telnet-ietf@cray.com
Status: Continuing Work
Draft or Recent RFC:

RFC1116: "Telnet Linemode Option", August 1989

1.2 Future IETF Meeting Sites

Winter 1990

Florida State University
Host: Ken Hays
February 6-9, 1990

Spring 1990

Pittsburgh Supercomputer Center
Host: Gene Hastings
May 1-4, 1990

Summer 1990

University of British Columbia
Host: John Demco
July 31- August 3, 1990

1.3 On Line IETF Information

The Interent Engineering Task Force maintains up-to-date on-line information on all its activities at NIC.DDN.MIL and NNSC.NSF.NET. On each of these hosts, there are two directories containing Internet-Draft documents and IETF working group information. All this information is available for public access.

The "IETF" directory has been created as an aid to both veteran IETF members and newcomers. It contains a general description of the IETF, summaries of ongoing working group activities and provides information on past and upcoming meetings. The directory generally reflects information contained in the most recent IETF Proceedings and Working Group Reports.

The "Internet-Drafts" directory has been installed to make available, for review and comment, draft documents that will be submitted ultimately to the RFC Editor to be considered for publishing as an RFC. Comments are welcome and should be addressed to the responsible person whose name and email addresses are listed on the first page of the respective draft.

In each directory there is a 00README file.

To access these directories, use FTP to NIC.DDN.MIL or NNSC.NSF.NET. After establishing a connection, Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

At NIC.DDN.MIL

cd internet-drafts:
cd ietf:

At NNSC.NSF.NET

cd internet-drafts
cd ietf
Note: The only difference is the colon required by the NIC Tops 20 machine.

Individual files can then be retrieved using the GET command:

get <remote filename>   <local filename>
e.g., get 00README   readme.my.copy

1.4 Guidelines to Authors of Internet Drafts

The Internet Drafts Directory is available to provide authors with the ability to distribute and solicit comments on documents they plan to submit as RFC's. Submissions to the Internet Drafts Directory should be sent to internet-drafts@nri.reston.va.us. Unrevised documents placed in the Internet Drafts Directory have a maximum life of six months. After that time, they will either be submitted to the RFC editor or will be deleted. After a document becomes an RFC, it will be replaced in the Internet Drafts Directory with an announcement to that effect for an additional 6 months.

Internet Drafts (I-D's) are generally in the format of an RFC with some key differences. The Internet Drafts are not RFC's and are not a numbered document series. The words INTERNET-DRAFT should appear in place of RFC XXXX in the upper left hand corner. The document should not refer to itself as an RFC or as a Draft RFC. The Internet Draft should not state nor imply that it is a proposed standard. To do so conflicts with the role of the IAB, the RFC editor and the IESG.

The document should have an abstract section, with a one to two paragraph abstract to follow the Status of this Memo section. If the draft becomes an RFC, the Status of the Memo section will be filled in by the RFC editor with a status assigned by the IAB. As an Internet Draft, that section should contain one of the following statements.

a) This draft document will be submitted to the RFC editor as a protocol specification. Distribution of this memo is unlimited. Please send comments to ....

b) This draft document will be submitted to the RFC editor as an informational document. Distribution of this memo is unlimited. Please send comments to ....

1.5 Current Internet Drafts

This summary sheet provides a short synopsis of each Internet Draft available within the "Internet-Drafts" Directory at the NIC and NNSC.

"Managing Asynchronously Generated Alerts," edited by Louis Steinberg/IBM for the Alert Management Working Group, September 1989 <draft-ietf-alertman-asyncalertman-01.txt>

> This draft defines mechanisms to prevent a remotely managed entity from burdening a manager or network with an unexpected amount of network management information, and to ensure delivery of "important" information. The focus is on controlling the flow of asynchronously generated information, and not how the information is generated. Mechanisms for generating and controlling the generation of asynchronous information may involve protocol specific issues.
>
> There are two understood mechanisms for transferring network management information from a managed entity to a manager; request-response driven polling, and the unsolicited sending of "alerts". Alerts are defined as any management information delivered to a manager that is not the result of a specific query. Advantages and disadvantages exist within each method. This draft discusses these in detail.

"The Authentication of Internet Datagrams", edited by Jeff Schiller/MIT for the Authentication Working Group, August 1989 <draft-ietf-auth-ipauthoption-00.txt>

> This draft RFC describes a protocol and IP option to allow two communicating Internet hosts to authenticate datagrams that travel from one to the other. This authentication is limited to source, destination IP address pair. It is up to host-based mechanisms to provide authentication between separate processes running on the same IP host. The protocol will provide for "authentication" of the datagram, not concealment from third party observers.

"The Transmission of IP Datagrams over FDDI Networks" edited by Dave Katz for the IP over FDDI Working Group, January 1990 <draft-ietf-fddi-ipdatagrams-00.txt>

> The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams and ARP requests and replies over FDDI networks.

"Use of OSI IS-IS for Routing in TCP/IP and Dual Environments" edited by Ross W. Callon for the IS-IS for IP Internets Working Group, January 1990 <draft-ietf-isis-spec-00.ps>

This internet draft specifies an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments and dual environments. This specification was developed by the IS-IS working group of the Internet Engineering Task Force. Comments should be sent to "is-is@merit.edu".

"A Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices" edited by Robert Stine/Sparta for the NOC-Toools Working Group, November 1989 <draft-ietf-noctools-debugging-01.txt>

This draft contains a catalog with descriptions of several tools available to assist network managers in debugging and maintaining TCP/IP internets and interconnected communications resources. Entries in the catalog tell what a tool does, how it works, and how it can be obtained.

"An Echo Function for ISO 8473", edited by Robert Hagens/U-Wisconsin for the OSI Working Group, October 1989 <draft-ietf-osi-iso8473-00.txt>

This draft defines an echo function for the connectionless network layer protocol. Two mechanisms are introduced that may be used to implement the echo function. The first mechanism is recommended as an interim solution for the Internet community. The second mechanism will be progressed to the ANSI X353.3 working group for consideration as a work item.

This draft is not intended to compete with an ISO standard. When an ISO standard is adopted that provides functionality similar to that described by this memo, then this memo will become obsolete and superceded by the ISO standard.

"Internet Cluster Addressing Scheme", by Carl-Herbert Rokitansky/Fern Uni-Hagen, August 1989 <draft-ietf-pdn-clusterscheme-00.txt>

In this document, the new concept of an addressing scheme, similar, but inverse to the subnetting scheme, is proposed, in which a set of Internet networks is associated to an Internet cluster.

This "Cluster Addressing Scheme" is of interest especially for wide-area networks, whose structure should be visible to the outside world for (global) routing decisions. In addition, the use of an address-mask (called "Cluster-Mask") for routing decisions within the cluster is discussed.

"Application of the Cluster Addressing Scheme to X.25 Public Data Networks and Worldwide Internet Network Reachability Information Exchange", by Carl-Herbert Rokitansky/Fern Uni-Hagen, August 1989 <draft-ietf-pdn-pdncluster-00.txt>

In this document, the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks is discussed and a new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is proposed.

"Assignment/Reservation of Internet Network Numbers for the PDN-Cluster", by Carl-Herbert Rokitansky/Fern Uni-Hagen, July 1989 <draft-ietf-pdn-pdnclusternetassignm-00.txt>

This document contains a proposal for the reservation of Internet network numbers for the PDN-cluster and the assignment of these PDN-cluster networks to all national X.25 public data networks (DNICs), which are worldwide already in operation.

"Gateway Congestion Control Policies", edited by A.J. Mankin/Mitre and K.K. Ramakrishnan/DEC, July 1989, <draft-ietf-perfcc-gwcc-00.txt>

The task remains for Internet implementors to determine effective mechanisms for controlling gateway congestion. This paper describes the characteristics of one experimental gateway congestion policy, Random Drop, and several that are better-known: Source Quench, Congestion Indication, Selective Feedback Congestion Indication, and Fair Queueing. Random Drop needs further study and does not offer solutions to the resource allocation problems that are the generalization of the congestion control problem. However, a motivation for documenting it now is that it has as primary goals low overhead and suitability for scaling up. Both of these are important goals for future gateway implementations that will have fast links, fast processors, and will have to serve large numbers of interconnected hosts.

"The Point-to-Point Protocol (PPP) Initial Configuration Options" edited by Drew Perkins/CMU for the PPP Working Group, November 1989 <draft-ietf-ppp-options.txt>

The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of

1. a method for encapsulating datagrams over serial links,
2. an extensible Link Control Protocol (LCP), and
3. a family of Network Control Protocols (NCP) for establishing and configuring different network-layer protocols.

The PPP encapsulating scheme, the basic LCP, and an NCP for controlling and establishing the Internet Protocol (IP) (called the IP Control Protocol, IPCP) are defined in The Point-to-Point Protocol (PPP) [1].

This document defines the intial options used by the LCP and IPCP. It also defines a method of line quality monitoring and a simple authentication scheme.

"Requirements for an Internet Standard Point-to-Point Protocol", edited by Drew Perkins/CMU for the Point to Point Protocol Working Group, June 1989 <draft-ietf-ppp-req-00.txt>

This draft document discusses the requirements for an Internet standard data link layer protocol to be used with point-to-point links. Although many industry standard protocols and ad hoc protocols already exist for the data link layer, none are both complete and sufficiently versatile to be accepted as an Internet standard. In preparation to designing such a protocol, the features necessary to qualify a point-to-point protocol as an Internet standard are discussed in detail. An analysis of the strengths and weaknesses of several existing protocols on the basis of these requirements demonstrates the failure of each to address key issues.

"Management Information Base for Network Management of TCP/IP-based internets", edited by M. T. Rose/ NYSERNET for the SNMP Working Group, September 1989 <draft-ietf-snmp-mib2-01.txt>

This memo defines the second version of the Management Information Base (MIB-II) for use with network management protocols in TCP/IP-based internets. In particular, together with its companion memos which describe the structure of management information (RFC 1065) along with the network management protocol (RFC 1098) for TCP/IP-based internets, these documents provide a simple, workable architecture and system for managing TCP/IP-based internets and in particular the Internet community.

"OSI Connectionless Transport Services on top of the UDP: Version 1", edited by C. Shue/OSF, W. Haggerty/Wang and K. Dobbins/ Cabletron, November 1989 <draft-osf-shue-osiudp-00.txt>

This draft proposes a method for offering the OSI connectionless transport service (CLTS) in TCP/IP-based Internets by defining a mapping of the CLTS onto the User Datagram Protocol (UDP). If this draft becomes a standard, hosts on the Internet that choose to implement OSI connectionless transport services on top of the UDP would be expected to adopt and implement the methods specified in this draft. UDP port 102 is reserved for hosts which implement this draft. Distribution of this memo is unlimited.

This memo serves as a companion document to RFC 1006 "ISO Transport Service on top of the TCP, Version 3".

"The Knowbot Information Service" by Ralph Droms/Bucknell, December 1989 <draft-nri-droms-kis-00.txt> and <draft-nri-droms-kis-00.ps>

Within the metanetwork of networks that exchange electronic mail, there are many directory services that provide partial coverage of network users; that is, directories with information about some subset of a particular network's user population. Searching the collection of available directories is time-consuming and requires knowledge of each directory's user interface. Although X.500 is currently under study as a basis for an Internet-wide directory service, it is unlikely that a universal user registry will be in place in the near future. The Knowbot Information Service provides a uniform interface to heterogeneous directory services that simplifies the task of locating users in the combined network.

Chapter 2

Area and Working Group Reports

2.1 Applications Area

Interim Director: Phillip Gross/NRI

I would like to welcome Russ Hobby/UC-Davis to the IESG as our new Applications Area Director. Russ's responsibilities at UC-Davis include network administration and new services to end users. Russ is involved in the California Internet Federation, and has been a longstanding IETF member. Russ was instrumental in the final push which completed the Point-to-Point Protocol (PPP) basic protocol as RFC1134. He currently continues to chair the PPP extension Working Group. Russ will be reporting in this spot in the future.

There are currently only two working groups in this area:

- Telnet (Borman, Cray)
- Domain Name System (Mockapetris, ISI)

Each of these WGs has produced RFCs, but continue to work on related topics in their areas.

I would like to encourage activity in other application areas as well. The following topics have been proposed:

- Common protocol for remote printing that could be used by both TCP/IP and PC-based networks
- Remote back-up facility for both TCP/IP and PC-based networks
- Electronic mail (bitmaps for SMTP, standardizing addressing hacks)
- Usage of DECnet Naming Service in TCP/IP networks

Some of this work would need to be coordinated with prospective application services (e.g., RCP) under the Host and User Services Area.

Please contact Russ Hobby (rdhobby@ucdavis.edu) with any comments, suggestions, or proposals you may have about new end-user network applications.

2.1.1 Domain Name System Working Group (dns)

CHARTER

Chairperson: Paul Mockapetris/USC-ISI, pvm@isi.edu

Mailing Lists: namedroppers@sri-nic.arpa

Description of Working Group:

The goal of the Domain Working Group is to advise on the administration of the top levels of the DNS ("the root servers"), consider proposed extensions and additions to the DNS structure and data types, and resolve operational problems as they occur.

Specific Objectives:

The specific short-term objectives are:

- Adding load balancing capability to the DNS.
- Adding DNS variables to the MIB.
- Implementation catalog for DNS software.
- Responsible Person Record.
- Adding network naming capability to the DNS.
- Evaluate short term measures to improve, or at least describe the security of the DNS.

Estimated Timeframe for Completion:

1. The preferred method for Load Balancing was decided upon at the April '89 IETF meeting at Cocoa Beach. A short RFC will be written in the near future.
2. Questionaire sent, responses data being organized, summary and detail to appear.
3. RFC issued April 89, implementations to follow.

CURRENT MEETING REPORT

Reported by Paul Mockapetris/USC-ISI

AGENDA

1. Joint meeting with X.400/X.500 WG
2. Discussion re DNS problems and solutions document
3. Future plans

MINUTES

The Domain WG met twice at the Oct-Nov IETF meeting in Hawaii.

The first meeting was a joint meeting with the X.400 WG to try to create coexistence plans for migration to X.400/X.500. The bulk of the discussion was OSI related. However, a plan to build a X.500 style (if not substance) name space was presented. The talks were frank and constructive, but the participants needed some OSI expertise which was not represented.

The second meeting discussed a draft document about DNS problems and solutions. Some changes were indicated, but the draft met with general approval, and a final version is to be prepared by Paul Mockapetris.

Additional proposals were to tackle the problem of dynamic update, and/or private domains. Since no concrete proposals were available, these issues were deferred till the next IETF, where they will be addressed or the WG will, once again, disband.

ATTENDEES

Arnold, Susan
Buroan, Jeffrey
Enger, Robert
Fuller, Vince
Huston, Geoff
Knowles, Stev
Lazear, Walt
Lekashman, John
Long, Dan
Lottor, Mark
Mockipetris, Paul
Reilly, Michael
Reschly Jr., Roberr J.
Roseustein, Mark
Rust, Bill
Schoch, Steven

2.1.2 TELNET Working Group (telnet)

CHARTER

Chairperson: Dave Borman/Cray, dab@cray.com

Mailing List: telnet-ietf@cray.com

Description of Working Group:

The TELNET working group is to look at RFC 854, "Telnet Protocol Specification", in light of the last 6 years of technical advancements, and determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the numerous TELNET options, and decide which of them are still germane to current day implementations of the TELNET protocol.

Specific Objectives:

- Either re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol, or issue a companion RFC to update and expand on fuzzy areas of RFC 854.
- Create or update RFCs for TELNET options to clarify or fill in any missing voids in the current option set. (Most noteably, some method to allow automatic user authentication is needed).
- Act as a clearing house for all proposed RFCs that deal with the TELNET protocol.
- When the above objectives have been met the group will go dormant, and will be re-activated as needed to fullfill the objective of being a clearing house for future extensions to the TELNET protocol.

Estimated Timeframe for Completion:

Will be determined during the next meeting.

CURRENT MEETING REPORT

Reported by David Borman/Cray Research

AGENDA

The meeting was started by looking getting a list of agenda items.

1. Review initial drafts of proposed options that were listed at the last meeting.
2. Discuss which of the current options are good or bad, and which ones should, may, or should not be used.
3. Talk about the TELNET RFC
4. Discuss option negotiation, and option negotiation loop avoidance.
5. Discuss flushing of data streams.

Item No. 5 was addressed immediatly. At the previous meeting we had discussed how to properly flush the data stream. After that meeting, final wording for the Host Requirements document was hammered out to give guidence to implementors about how to flush the data stream. So, the group felt that this issue has been dealt with.

Item Nos. 2, 3, and 4 were not discussed, as the rest of the meeting focused on item 1, review of draft proposals.

There were six draft options brought forward for discussion:

- Environment
- User Name
- System Type
- Authentication
- Encryption
- Compression

The encryption and compression options were not discussed due to time constraints.

The System Type option was voted down. After some discussion, it was decided that no one had any need at this time to find out the system type of the remote machine, and that there was no clear method of specifying the system type. (The example discussed was all the different variants of UNIX.) If at some later point in time a reason for needing the System Type is found, this option could be resurrected.

The User Name option was thought to be a straight forward option. Then discussion brought up that some systems also have an account id, in addition to the user name. Louis Mamakos pointed out that there are also systems that require a user name, and account id, and a project id. With this in mind, rather than expanding the User Name option to contain all this information, it was decided to get rid of the User Name option, and fold the functionallity into the Environment option.

So, the next option discussed was the Environment option. Modification were made to the draft. Rather than specifiying environment options as ENVIRON IS ... [IS ... [...]] where "..." is "VARIABLE=VAR", it has been changed to: ENVIRON IS VAR ... [VALUE ...] [VAR ... [VALUE ...]] [...]

This will allow for more flexability, and is not OS dependent. We will pre-define a few well known environtment variable names:

- USER
- ACCT
- PROJ
- PRINTER
- DISPLAY

The authentication option was discussed at some length. It was decided that the framework provided by the option would probably be sufficient for passing authentication information back and forth.

However, in order to be useful, the authentication option has to be defined in terms of at least one specific form of authentication. Kerberos was discussed. Some people felt that we needed something besided Kerberos. After much more discussion, Louis Mamakos agreed to write up something for a simple user authentication, and Milt Roselinsky agreed to write up something for system authentication. (The idea behind system authentication is once you authenticat the remote system, you can use an rlogin style of .rhosts file to log in the user.) Mike Karels volunteered Kevein Fall to write up how to use Kerberos with this option.

It was also decided that as new forms of authentication/encryption were added to the authentication/encription options, an entire new RFC would be issued that would update the previous RFC with the new scheme. We decided this was preferable to having lots of RFCs for each type of authentication/encryption.

And that covers most of the main points of the meeting. We will hold our next meeting at the next IETF meeting.

ATTENDEES

Bagnall, Doug
Borman, Dave
Hedrick, Charles
Karels, Mike
LoVerso, John
Mamakos, Louis A.
Miller, Dave
Reynolds, Joyce K.
Roselinsky, Milt
Solensky, Frank
Westfield, Bill
Wilder, Rick

2.2 Host and User Services Area

Director: Craig Partridge/BBN

This area combines two distinct activities: work on improving the quality of host-based services from the transport layer up to, but not including, applications; and work on developing and improving the quality of user services available on the Internet.

In the area of Host-Based services, the plan for the next year is to actively encourage new work on standards for various support services such as remote procedure call, external data formats, distributed file systems, and network graphics.

Under User Services, the key focus is on improving the services that already developed, and encouraging and fostering new activities such as the publication of the User Directory, the FYI notes, and the SIGUCCS project, that hold promise for improving user services offerings in the Internet.

As for what happened in Honolulu:

The major activities in Host-Based Services were the meeting of the Dynamic Host Configuration Working Group (which is progressing faster than expected towards developing a configuratihon protocol) and the Ad-Hoc TCP Options WG (which got bogged down a bit at its meeting).

In User Services, the major news is that Karen Bowers, after doing a wonderful job getting the User Services WG started under IETF, has taken on new responsibilities at NRI which make it difficult for her to continue as USWG chair. Joyce Reynolds has agreed to be the new chair of the USWG, effective immediately.

2.2.1 Dynamic Host Configuration Working Group (dhc)

CHARTER

Chairperson: Ralph Droms/Bucknell, rdroms@nri.reston.va.us

Mailing List: host-conf@rutgers.edu

Description of Working Group:

The purpose of this working group is the investigation of network configuration and reconfiguration management. We will determine those configuration functions that can be automated, such as Internet address assignment, gateway discovery and resource location, and that which cannot (i.e., those that must be managed by network administrators).

Specific Objectives:

1. We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to:
   (a) Exchange Internet packets with other hosts (e.g., discover own Internet address).
   (b) Obtain packet routing information (e.g., discover local gateways).
   (c) Access the Domain Name System (e.g., discover a DNS server).
   (d) Access other local and remote services.
2. We will summarize those mechanisms already in place for managing the information identified by objective 1.
3. We will suggest new mechanisms to manage the information identified by objective 1.
4. Having established what information and mechanisms are required for host operation, we will examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.

Estimated Timeframe for Completion:

1. Problem statement will be submitted as an RFC.
2. New Protocol document in one year.

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell University

MINUTES

The Dynamic Host Configuration Working Group met in two half-day sessions in Hawaii. The primary purpose of the meetings was to discuss the Working Group's Internet Draft describing the dynamic host configuration problem. The Working Group agreed on a final draft, which has been added to the IETF Internet Draft series.

The Working Group has decided to concentrate first on the problem of initializing the network layer. This initialization step includes the allocation of an IP address to the host, and transmission of that address, along with other subnet parameters such as subnet mask, MTU and broadcast address to the host. Current protocols that address the network layer initialization problem include RARP, BOOTP, the Athena project's Network Information protocol (NIP), and Sun's diskless workstation initialization mechanism. Based on the design parameters laid out in the draft problem statement document, the Working Group expects to define a new initialization protocol based primarily on BOOTP and NIP.

The Working Group's problem statement has been made available for comment as an Internet Draft. This document will be submitted for publication as an RFC in after the February, 1990 IETF meeting. At the February meeting, work will begin on the definition of the new network layer initialization protocol. Upon completion, the protocol definition will be submitted as an RFC. An experimental version of the new protocol will be developed, based on the current version of NIP, after the February meeting.

ATTENDEES

Almquist, Philip	Melohn, Bill
Bagnall, Doug	Mockapetris, Paul
Borman, Dave	Pleasant, Mel
Brackenridge, Billy	Reschly, Robert J.
Catlett, Charlie	Rosenstein, Mark
Cook, John	Schiller, Jeff
Easterday, Tom	Solensky, Frank
Lear, Eliot	Vaudreuil, Greg
Lekashman, John	Wilder, Bruce
LoVerso, John	Yasaki, Brian
Mamakos, Louis A.	

2.2.2 Internet User Population Working Group (iup)

CHARTER

Chairperson: Craig Partridge/BBN, craig@nnsc.nsf.net

Mailing List: ietf@venera.isi.edu (interim address)

Description of Working Group:

To devise and carry out an experiment to estimate the size of the Internet user population.

Specific Objectives:

We expect to produce two documents: (1) a description of the experimental procedure and (2) an RFC that gives the results of the experiment. We may also produce a short paper for publication in a networking magazine.

Estimated Timeframe for Completion:

The firm hope is that this will only take two meetings: Hawaii to determine the experimental design and then the next meeting to report the results.

CURRENT MEETING REPORT

Reported by Craig Partridge/BBN

MINUTES

The group met and discussed methods for determining the user population. Two approaches were discussed:

- generate a list of hosts from the DNS and then take samples of users on different hosts. Use these samples to extrapolate to the total population.
- poll on a per-domain basis on the assumption that site admins know how many users they have.

Craig Partridge has agreed to write up plans for doing both types of surveys.

ATTENDEES

Borman, David
LaQuey, Tracy
Lottor, Mark
Moore, Berlin
Partridge, Craig
St.Johns, Mike
Wintringham, Dan
Yuan, Aileen

2.2.3 TCP Large Windows Working Group (tcplw)

CHARTER

Chairperson: Craig Partridge/BBN, craig@bbn.com

Mailing List: ietf@venera.isi.edu

Description of Working Group:

> This is a short term, ad hoc, single question working group chartered to make some progress on the various proposals for TCP in long fat pipes.

Specific Objectives:

> Choose a proposed standard for the TCP extended window size option.

Estimated Timeframe for Completion:

> No later than the February IETF Meeting in Talahassee Florida.

CURRENT MEETING REPORT

Reported by Craig Partridge/BBN

MINUTES

The TCP Large Windows WG met for half a day to discuss the two proposals (RFCs 1072 and 1106) for improving TCP for large delay-bandwidth paths. During the meeting two key issues were raised.

The group determined that a key problem was how large to permit the window to be. A larger window makes it easier to consume the 32-bit sequence quickly. An example may help here. If one permits a window of 2^30 bytes, then in each round-trip time, one quarter of the sequence space can be consumed, and in four RTTs, the sequence space will recycle. However, a TCP cannot cycle the sequence space until it is sure the TTL of prior segments has expired (the forbidden zone problem). So, we were faced with choosing window sizes, that at anticipated speeds, didn't cause the sequence space to roll over in less than the anticipated TTL. In the end, the group was uncomfortable with this problem and has asked Van Jacobson and Bob Braden (both of whom have looked at this issue in more detail) to attend the next meeting.

Another issue was whether we preferred to use options in every segment to expand the window, or preferred to find a way that didn't cause implementations to do expensive option handling. The consensus was to avoid option handling (which meant we preferred the rfc 1072 approach). Some discussion was given to generating a larger TCP header, but this conversation foundered when we checked the TCP header and found it lacked a version number.

The group did not have time to consider another interesting proposal (passed on from the IETF Hosts group) to allow text error messages in RST segments.

ATTENDEES

Borman, Dave	McCloghrie, Keith
Elz, Robert	McKenney, Paul
Fox, Richard	Miller, Dave
Galvin, James M.	Solensky, Frank
Hedrick, Charles	St.Johns, Mike
Karels, Mike	Yasaki, Brian
Love, Paul	

2.2.4 User Documentation Working Group (userdoc)

CHARTER

Chairpersons: Tracy LaQuey/University of Texas tracy@emx.utexas.edu
Karen Roubicek/BBN roubicek@nnsc.nsf.net

Mailing List: user-doc@nnsc.nsf.net

Description of Working Group:

The USER-DOC Working Group will prepare a bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "how to use the Internet". (Target audience: those individuals who provide services to end users and end users themselves.)

Specific Objectives:

1. Identify and categorize useful documents/reference materials/training tools.
2. Publish both an on-line and hard copy of this bibliography.
3. Develop and implement procedures to maintain and update the bibliography. Identify an organization or individuals to accept responsibility for this effort.
4. As a part of the update process, identify new materials for inclusion into the active bibliography.
5. Set up procedures for periodic review of the biblio by USWG.

Estimated Timeframe for Completion:

1. Format for the bibliography will be decided upon by the July IETF session, as well as identification of "sources of information" (e.g., individuals, mailing lists, bulletins, etc.)
2. Draft bibliography will be prepared by mid-December 89.

CURRENT MEETING REPORT

Reported by Tracy LaQuey/University of Texas

AGENDA

1. Review Charter, Objectives, Timeframe
2. Review last meeting's action items
3. Review draft and make comments
4. Discuss RFC situation
5. Discuss review process, researching for entries
6. Discuss format
7. Determine what's next, for example, distribution and updating
8. Determine action items for the next meeting

MINUTES

The meeting began with a review of the charter of the USER-DOC Working Group. To summarize, the purpose of the group is to prepare a bibliography of online and hardcopy documents, reference materials, and training tools addressing general networking information and "how to use the Internet". End users and people who help end users are the targeted audience. The group has been collecting documents since the first USWG meeting in Texas and have described some broad categories to cover. It was decided at the IETF meeting at Stanford to have a rough draft ready by this meeting.

Last meeting's action items were reviewed. They were:

- Write up and pass out questionnaire at Plenary Session: Enger, Roubicek, Bowers

    This questionnaire was written up and passed out at the Stanford meeting.

- Develop Template (LaQuey, Marine, Redfield, Roubicek)

    Karen Roubicek made up a template which was reviewed by the mailing list. She then sent it out to several mailing lists and USENET groups.

- Define set of key RFCs (Reynolds)

    Joyce Reynolds put together a "Basic Beige" list of RFC's.

- Research Andrew system and Federal databases (Breeden)
- Find information SOCRATES project (Breeden)

    Laura Breeden did not attend this meeting, so we heard no report on these two items.

- Choose set of mailing lists (Roubicek, Bowers)

  Mailing lists, groups, newsletters and other methods of distribution were discussed at the Stanford meeting by a full meeting of the User Services group.

- Liaison with library community (Roubicek)

  Since we don't really have a final draft, we are waiting before we do this.

- Schedule distribution meeting (Bowers)

  A joint meeting with the NOCTOOLS WG chaired by Martyne was held at Stanford.

- Write article about bibliography for ConneXions soliciting info (Perillo)

  We did not write an article between the Stanford and Hawaii meeting, but Ole has us scheduled for helping him with an article for the December ConneXions issue.

The rough draft in reference format was passed out to all members attending and the entries and categories were discussed. It was suggested that we include an entry that describes the format of online documents (ASCII, PostScript, etc.) The following categories were discussed and added to the scope of the bibliography: Glossaries, Marketing Materials, Pointers to NICs, Pointers to Maps (this was actually added at the last meeting but no entries were put into the rough draft), and pointers to bulletins (like DDN Security and CERT). The marketing materials will help those users who need to show the value of networks, and procedures for getting connected. Some specific suggestions were made for additional entries (the ACM article on security and ethics, FRICC documents - OSTP report and online Gore Bill, NREN document). Elizabeth Redfield had sent a suggestion that OSI documents be included, so we decided to include general documents on OSI protocols and tutorials. It was decided that we should also add an introduction and short explanations on networks.

Tracy presented a sample format for the bibliography - using keywords to group similar documents. Dave Crocker pointed out the annoyance of having to flip back and forth from an index to the body of the text. He suggested using more detailed key information that will give a better idea of what the documents are about.

Tracy passed out a letter sent by Elizabeth Redfield concerning some bibliographic issues and additional fields. We decided that the format board would review this message and report on it at the next meeting.

Inclusion of RFCs was discussed. We don't want to have any political problems with including certain RFCs and excluding others. Joyce Reynolds is going to get some feedback from Jon Postel and Vint Cerf about her "Basic Beige RFCs" list. It was decided to go ahead and include this list, though.

Since a lot of entries are incomplete and questionable, we decided that our editorial and research boards should go over all of the entries. The editorial board will decide if an entry should be included while the research board will work on incomplete entries. The chair of the editorial board is Karen Bowers. Members include Tracy LaQuey, Francine Perillo, Joyce Reynolds, April Marine and Jon Pugh. The Research committee now includes the NIC, Jon Pugh, Martyne Hallgren, Dana Sitzler, Karen Roubicek, Tracy LaQuey and Roxanne Streeter.Since the format of the final bibliography is still not decided, several members decided to discuss those issues. The "Formatting Board" is Karen B. and Mary Stahl, with help from Dana Sitzer, Karen R. and Laura Breeden.

We will update the bibliography annually. If the group is not actively in session with the IETF, then we will reconvene for some period of time. The draft is still located and available via anonymous ftp on emx.utexas.edu in the directory "user.wg", file "bibliography".

The following action items were assigned:

- Karen Bowers will schedule a videoteleconference for the Editorial Board the week of November 13. The board will look over the existing draft and make changes.
- Joyce Reynolds is proposing an F.Y.I. series of notes, allied to the RFC's, but providing information about who does what on the Internet. (For example, Joyce thought that the bibliography would be a good F.Y.I. document.)
- We will have a final rough draft sent out electronically by January 16. We will schedule a full day (preferably the first day of the next IETF meeting in Florida) to go over it.
- The Research Board will find more documents and complete existing entries.
- The Format Board will review Elizabeth Redfield's bibliographic suggestions.
- We will follow up existing connections in the Library community.
- Ole, Tracy and Joyce will work on an article for ConneXions.

ATTENDEES

Armstrong, Karen	Moore, Berlin
Bowers, Karen	Reynolds, Joyce K.
Choy, Joseph H.	Roubicek, Karen
Crocker, Dave	Sitzler, Dana
Enger, Robert M.	Stahl, Mary
Hallgren, Martyne M.	Streeter, Roxanne
Jacobsen, Ole	Wintringham, Dan
LaQuey, Tracy L.	Yuan, Aileen

2.2.5 User Services Working Group (uswg)

CHARTER

Chairperson: Joyce K. Reynolds, jkrey@venera.isi.edu

Mailing Lists:

us-wg@nnsc.nsf.net
us-wg-request@nnsc.nsf.net

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF WGs created to perform certain projects, or outside organizations such as SIGUCCS.

Specific Objectives:

1. Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should
   - clearly address user assistance needs;
   - produce an end-result (e.g. a document, a program plan, etc);
   - have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
   - and not duplicate existing or previous efforts.
2. Create WGs or other focus groups to carry out projects deemed worthy of pursuing.
3. Provide a forum in which user services providers can discuss and identify common concerns.

Estimated Timeframe for Completion:

This is an operational WG and, as such, has an indefinite lifetime.

CURRENT MEETING REPORT

Reported by Craig Partridge/BBN

MINUTES

The User Services WG meeting was devoted to discussion of the effects of various changes on the USWG.

The two key changes were:

1. Karen Bowers has new and expanded responsibilities at NRI that make it difficult for her to continue to devote as much time to the USWG, and so she has resigned as chair of the USWG as of the Honolulu meeting. Karen does plan to continue as a member of the WG.
   The group expressed its gratitude to Karen for the very considerable effort she has put into making the USWG a success.
2. The re-organization of IETF and the placement of USWG under Craig Partridge, the area director of Host Services. Craig talked about his goals for user services, in particular that the USWG continue to serve in its role as a focus for user services activities in the Internet community and that it would continue to meet regularly to discuss common concerns.
   Craig also made it a point to mention that User Services plays a vital role in IETF and the Internet as a whole, and that he views user services as an important activity under his directorship, and distinct from Host Services. (Based on suggestions from the WG, the IESG has agreed to rename the area to "Host-Based and User Services" to emphasize that Craig is directing two distinct activities).
3. Martyne Hallgren report that SIGUCCS had voted to fund her services proposal for a SIGUCCS committee to examine user services concerns from the university level (yeah!).
4. Joyce Reynolds presented to the group her concept of the FYI Series of Notes, a vehicle by which information can be provided to the user community. Joyce will pursue this further and query the user-svc mailing list for additional ideas: appropriate topics, volunteers for selected topics, etc.
5. Karen Bowers strongly encouraged the members of user-svc to take the opportunity to get involved in the other working group activities as well. This will provide the kind of "mesh" required to ensure those representing user services are aware of the technical issues and likewise make contributions in light of user requirements.

ATTENDEES

Almquist, Philip	Oattes, Lee
Armstrong, Karen	Partridge, Craig
Bagnall, Doug	Pleasant, Mel
Bowers, Karen	Reynolds, Joyce K.
Choy, Joseph H.	Roubicek, Karen
Enger, Robert M.	Sheridan, Jim
Garcia-Luna, Jose	Sitzler, Dana
Hallgren, Martyne M.	St. Johns, Mike
Huston, Geoff	Stahl, Mary
LaQuey, Tracy	Streeter, Roxanne
Malkin, Gary	Wintringham, Dan
Moore, Berlin	Yuan, Aileen

2.3 Internet Services Area

Director: Noel Chiappa/Consultant, Proteon

The Internet Services Area is a fairly active one, with many groups focusing on particular issues. A major recent accomplishment is the production of the standard Point-to-Point Protocol document , and an initial options document. Both of these documents are now available. The basic protocol document has been submitted to the RFC editor as a proposed standard. Further work on the Point-to-Point protocol will be accomplished in the Point to Point Extensions working group, which will add support for additional protocols to the basic specification.

Of the ongoing groups, the Connection IP group and the Performance and Congestion Control group, have documents well under way. A major long term issue in this area is congestion control. The initial round of mechanisms being discussed in the Congestion Control group may not be sufficient, and further work in the area (perhaps involving some research on the IRTF side) is needed. Router Discovery, MTU Discovery, and IP over FDDI have been formed and are quite active. Discovery of available routers is a pressing problem. Currently the community has to either configure hosts with router addresses, or use a variety of non- standard techniques to find them; something standard is clealy needed quickly.

Several new groups are in the process of being formed, and will be announced shortly. The most important is a Router Requirements group, to redo RFC-1009 and bring it up to the standards of the Host Requirements RFC's. A group is being put together to standardize use of IP with the Appletalk environment. Among other things, it will document the KIP protocol, which has come into wide use.

A working group is being set up to address IP over Multi-Media Bridges. There are a number of vendors who wish to offer multi-media bridges, but there are a number of technical issues to be solved before the IP protocol family will operate over such devices. Multi-cast has been in an interim state in the architecture for years. A number of RFC's on the issue have appeared, but they need to be brought forward to the 'recommended' state and mandated for use.

Other topics will be receiving attention soon. The Internet needs a clear standard for the use of variable length subnet masks. The original subnet RFC did not deal with the details of this issue, and it needs to be regularized. Detection of dead nodes, particularly dead routers is a pressing problem. Most hosts fail to recover gracefully from routers that crash, and although the Host Requirements RFC discussed the issue, more work is needed. A group is being created to consider the issues involved in integrating the Switched Megabit Data Service, SMDS, into the IP architecture. This new service to be offered by the phone companies will provide true packet service (i.e. no connections or connection setup) over a T3 rate interface. It looks like the world's largest LAN, and presents some scaling problems for the IP architecture; clearly, ARP cannot be used in its existing form!

2.3.1 Connection IP Working Group (cip)

CHARTER

Chairperson: Claudio Topolcic/BBN, topolcic@bbn.com

Mailing List: cip@bbn.com

Description of Working Group:

> Define the next version of the ST protocol, explore future connection oriented internet protocol, use the former as a testbed to perform experiments in support of the latter.

Specific Objectives:

- Produce a new specification of ST
- Produce a specification of a next generation connection oriented protocol

Estimated Timeframe for Completion:

1. Produce a new specification of ST. (2-3 months)
2. Produce a specification of a connection oriented protocol. (6-12 months)

CURRENT MEETING REPORT

Reported by Claudio Topolcic/BBN

MINUTES

Connection IP Meetings of October 31, 1989 and November 1, 1989

We discussed our objectives in the working group. In the short term we should participate in identifying an appropriate experimental platform for the research we want to perform. In the medium term, we are attempting to understand the issues and mechanisms involved in connection oriented protocols and how those mechanisms fit into a protocol. We should produce output in the form of research papers. In the longer term we should define a connection oriented protocol that incorporates what we have learned.

We discussed a draft "issues and requirements" paper written by Phil Park and Guru Parulkar. This prompted a number of useful discussions. We discussed the issue of what a guarantee is, and how a guarantee could be given in an Internet composed of numerous networks providing different services.

It was asked if an explicit setup phase is required or whether performance guarantees could be given based on passive observation of the traffic. We generally felt that it is not possible to make any form or performance guarantee unless there is an explicit setup phase. We agreed that the term "flow" would be a better term than "connection" because too may listeners assume reliability when they hear that term. We want to use "flow" and when WE say it we mean explicit setup. We also asked if a network can or should provide service before the setup phase is complete. The issue of how performance is specified was discussed. We decided that the best approach is to perform experiments to try out various parameters to check for completeness. Other issues that were touched upon were how to control offered load, what kind of information should be passed up from the internet protocol to higher level protocol, and whether a single protocol could properly support both high bandwidth applications across high performance networks as well as low bandwidth applications across low capacity networks. The consensus on the last issue was that we should be more interested in high bandwidth applications and high performance networks, rather than low bandwidths. At the end we decided that this paper should be cleaned up and put on line for access by the IETF community.

Guru Parulkar gave a presentation on ATM networks and a discussion ensued. Later, Danny Cohen said that he believes carriers WILL provide ATM networks in the future, so we should prepare now to take advantage of them by building experimental ATM networks ourselves and developing our protocols to work well with them. Connection oriented protocols may be a particularly good match for ATM networks.

We identified several research issues, including:

- The most significant research issue is resource management. It can be partitioned roughly into three parts; setup, monitoring and enforcement. Other issues include how resource management relates to routing, and particularly policy based routing; how internet level resource management can be performed across networks that do various levels of resource management; and the minimum requirements placed on the underlying networks in order to do resource management?
- What are the parameters needed to describe the service that an application requires?
- What will be the complexity of per packet forwarding and how does that interact with high performance networks?
- How can data transfer be supported before the setup phase is complete?

We looked at performing experiments in the near term to support research in these areas. We identified several possible experiments including:

- Experimenting with different types of resource management algorithms, such as Lixia Zhang's Flow Protocol and ATM network approaches.

- Exploring the co-existence of connection oriented and datagram applications and how they relate to resource management.
- Comparing connection oriented and datagram services.
- Evaluating the performance of different algorithms under stress.
- Exploring different sets of parameters to specify performance requirements.

We looked at potential experimental platforms. There is sentiment for using an easily available system such as BSD. Guru described the testbed that Washington University is building over the next two years. ST on the Butterfly is the only currently available implementation of a connection oriented internet protocol that operates across the Internet. It was proposed that this might be an appropriate testbed. We decided to continue the discussion of possible experiments at a later time by electronic mail and multimedia conferencing.

We agreed that there should be two presentations at the next meeting of this working group. Guru Parulkar will give an in-depth presentation on ATM network technology, and Claudio Topolcic will give a presentation on the ST-2 protocol.

ST Protocol Specification Meeting of November 1, 1989

The meeting concentrated on the draft ST-2 specification document written by Claudio Topolcic as a result of the discussions at the previous IETF meeting. Most of the discussion centered about the concept of "Groups" of streams. Groups provide a mechanism that supports MDHD (Multi-Destination Half Duplex) conferences and more elegantly supports pre-emption of established streams by higher priority streams. Steve Casner pointed out that the Group concept was incomplete because it does not support call blocking as elegantly as would be desired. Although this issue was not decided at the IETF meeting, at a later meeting, the two principals agreed to suggest removing Groups from the protocol specification.

The balance of the meeting was spent identifying other flaws with the draft ST-2 specification document. These included the following:

- The high level protocol description is OK.
- The term "connection" and "stream" should be better defined.
- The terms "high reliability" is inaccurate and should be replaced with the concept of "low probability of congestion loss".
- The concept of merging PTP and CONF connection types, and using three bits in the CONNECT message to select individual characteristics had not made it into the specification. The three bits defined are:
  1. the stream was to always remain two-way
  2. to construct the reverse path along with the forward path
  3. to use reverse HIDs

This concept and the description of these bits should be added back into the document.

- It was suggested we use the term "flow" rather than "connection" because many readers incorrectly infer reliability when they read "connection".
- The document should be formatted in Slate.

ATTENDEES

The attendee list is a combination of meetings on Tuesday, Wednesday and Thursday:

Boive, Rick
Casner, Stephen
Cohen, Danny
Fox, Richard
Guru, Parulkar
McKenney, Paul E.
Park, Phil
Ramakrishnan, K.K.
Solensky, Frank
Steenstrup, Martha
Su, Zaw-Sing
Topolcic, Claudio
Zhang, Lixia

2.3.2 Router Discovery Working Group (rdisc)

<u>CHARTER</u>

Chairperson: Steve Deering/Stanford, deering@pescadero.stanford.edu

Mailing Lists:

gw-discovery@gregorio.stanford.edu
gw-discovery-request@gregorio.stanford.edu

An archive of all mail to the list is available by anonymous FTP from host gregorio.stanford.edu, file gw-discovery/mail-log.

Description of Working Group:

The Gateway Discovery Working Group is chartered to adopt or develop a protocol that Internet hosts may use to dynamically discover the addresses of operational neighboring gateways. The group is expected to propose its chosen protocol as a standard for gateway discovery in the Internet.

The work of this group is distinguished from that of the Host Configuration Working Group in that this group is concerned with the dynamic tracking of gateway availability by hosts, as opposed to the initial configuration of hosts.

Specific Objectives:

1. Identify existing and proposed protocols, and if necessary develop a new protocol, for gateway discovery.
2. Evaluate the protocols identified in 1 for suitability as Internet standards, according to criteria to be agreed upon by members of the Working Group. For new protocols or extensions to existing protocols, the evaluation shall include prototype implementations before being proposed as a standard.
3. Produce an RFC recommending a standard protocol for gateway discovery.

Estimated Timeframe for Completion:

It is hoped that the Working Group can complete all of its objectives within 6 months of its initial meeting.

CURRENT MEETING REPORT

The first meeting is planned for the February IETF Meeting.

2.3.3 MTU Discovery Working Group (mtudisc)

CHARTER

Chairperson: Jeffrey Mogul/DEC, mogul@decwrl.dec.com

Mailing List:

mtudwg@decwrl.dec.com
mtudwg-request@decwrl.dec.com

Description of Working Group:

The MTU Discovery Working Group is chartered to produce an RFC defining an official standard for an IP MTU Discovery Option. "MTU Discovery" is a process whereby an end host discovers the smallest MTU along a path over which it is sending datagrams, with the aim of avoiding fragmentation.

Specific Objectives:

1. Decide if the proposal in RFC 1063 is sufficient, or if there are flaws to be corrected, or possible improvements to be made. Or, decide that it is unwise to create an official standard.
2. Unless the proposal in RFC 1063 is acceptable, write a new RFC describing a different approach.
3. Encourage the participation of gateway implementors, since the MTU discovery process affects the design and performance of IP gateways.
4. Encourage sample implementations of end-host and gateway portions of MTU Discovery for popular software (BSD-derived kernels, primarily). (b) Encourage rapid implementation by major gateway vendors, since this option is relatively useless without widespread support.

Estimated Timeframe for Completion:

The first two objectives should be completed by April 1990. Objective 4a (sample implementations) should be attempted before the final RFC is released, to alert us to any pitfalls. Objective 4b (implementation by gateway vendors) may take longer.

CURRENT MEETING REPORT

The first meeting is planned for the February IETF Meeting.

2.3.4 IP Over FDDI Working Group (fddi)

CHARTER

Chairperson: Dave Katz/Merit, dkatz@merit.edu

Mailing Lists:

fddi@merit.edu
fddi-request@merit.edu

Description of Working Group:

The IP Over FDDI Working Group is chartered to create Internet Standards for the use of the Internet Protocol and related protocols on the Fiber Distributed Data Interface (FDDI) medium. This group is specifically not chartered to provide solutions to mixed media bridging problems.

Specific Objectives:

To create Internet Standards for the use of IP, ARP, and related protocols on the FDDI medium.

To provide support for the wide variety of FDDI configurations (e.g., dual MAC stations) in such a way as to not constrain their application, while maintaining the architectural philosophy of the Internet protocol suite.

To maintain liason with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards.

This working group is not chartered to provide solutions to mixed-media bridging problems, although results produced by this working group should not preclude such solutions.

Estimated Timeframe for Completion:

An Internet Standard or Standards should be produced within six months, with an estimated completion date of May, 1990.

CURRENT MEETING REPORT

The first meeting is planned for the February IETF Meeting.

2.3.5 Performance and Congestion Control Working Group (pcc)

CHARTER

Chairperson: Allison Mankin/MITRE, mankin@gateway.mitre.org

Mailing Lists:

ietf-perf@gateway.mitre.org
ietf-perf-request@gateway.mitre.org

Description of Working Group:

The IETF Performance and Congestion Control Working Group is chartered to collect and develop short-term techniques for improving Internet performance, methods like TCP Slow-start, which are retrofittable and inexpensive to implement. After a preliminary draft of a white paper documenting such performance enhancements for hosts and gateways, it was decided to sharpen the focus and divide the material into two papers.

Specific Objectives:

The first paper is the Internet-Draft on gateway congestion control policies and algorithms. The intent of this paper is to present what is now known about the difficult problem of avoiding congestion in Internet gateways. It describes proposed policies such as Random Drop, Congestion Indication, and Fair Queuing, and sketches ground-rules for their adoption. An additional goal of the paper (achieved during the writing) is to generate dialogue on longer-term Internet gateway performance problems.

The other paper is an RFC on TCP performance. This describes TCP algorithms such as Retransmit Backoff, Slow-start, Nagle (Small-Packet Avoidance), and Delayed Ack, as well as their correct interaction. The scope is to expand the treatment of TCP performance found in the Host Requirements RFC.

Estimated Timeframe for Completion:

To be determined.

CURRENT MEETING REPORT

Did not meet

2.3.6 Point-to-Point Protocol Working Group (ppp)

<u>CHARTER</u>

Chairpersons:

Drew Perkins/CMU ddp@andrew.cmu.edu
Russ Hobby/UC Davis rdhobby@ucdavis.edu

Mailing Lists:

ietf-ppp@ucdavis.edu
ietf-ppp-request@ucdavis.edu

Description of Working Group:

The working group is defining the use of serial lines in data networks. While the main intent is to standardize the connection of IP networks over point-to-point links, the protocol is being designed to be extensible to other network protocols as well. The protocol will provide the capability of establishing the link parameters, authentication, link encryption, link testing, as well as control of the link while it is up. The protocol will also allow configuration and control of the higher level protocols such as IP, OSI, 802.3 bridging, and others.

Specific Objectives:

The main objective of the workgroup is to produce an RFC defining the protocol for the link and IP levels.

Estimated Timeframe for Completion:

The final draft of the RFC will be completed for the Fall 89 IETF Meeting.

<u>CURRENT MEETING REPORT</u>

Reported by Russ Hobby/UC Davis

MINUTES

The PPP WG met on October 31 and November 1 at the IETF meeting at the University of Hawaii. This was the last meeting of the working group, however, further work on PPP will be done by a new working group to define extensions to PPP such as new options and the use of other protocols on PPP.

During video conferences following the Stanford meeting the PPP document was split into two separate documents. The first is a base PPP document describing the packet encapsulation and the process for control and option negotiation for the link and IP, but does not describe the options themselves.

The second document describes the options for the Link Control Protocol (LCP) and the IP Control Protocol (IPCP). This document also describes a simple authentication method that can be used when bringing up the link and a means of doing link quality monitoring (formally keep-alives) of the link during operation.

Final editing of the first document was done and it was submitted for review as an RFC. The working group reviewed the second document to clear up any technical details but put most the emphasis on the method of Link Quality Monitoring (LQM). The details of (LQM) where discussed and agreed upon and would be written into the second document after the meeting. Both the base document as it was submitted to be an RFC and the second document are being submitted as Internet Drafts for public review.

Areas in need of further work are:

1. Stronger Authentication Protocols
2. Definition of encryption methods
3. Definition of the use of other high level protocols

As a beginning of the new extensions WG, members volunteered to write definitions for the use of bridging, DECNET (Phase IV and V), and XNS over PPP.

ATTENDEES

This is a consolidated attendee list for the meetings of Tuesday and Wednesday.

Arnold, Susan	Knowles, Stev
Baker, Fred	LoVerso, John R.
Berggreen, Art	Mamakos, Louis A.
Boivie, Rick	McKenney, Paul E.
Broersma, Ron	Melohn, Bill
Carvalho, Charles	Opalka, Zbigniew
Catlett, Charlie	Reilly, Michael
Chiappa, Noel	Schoch, Steven
Farinacci, Dino	Westfield, Bill

2.3.7 Point-to-Point Protocol Extensions Working Group (pppext)

CHARTER

Chairperson: Russ Hobby/UC Davis, rdhobby@ucdavis.edu

Mailing lists:

ietf-ppp@ucdavis.edu
ietf-ppp-request@ucdavis.edu

Description of Working Group:

The Point-to-Point Protocol (PPP) was design to encapsulate multiple protocols. IP was the only network layer protocol defined in the original documents. The working group is defining the use of other network level protocols and options for PPP. The group will define the use of protocols including: bridging, ISO, DECNET (Phase IV and V), XNS, and others. The group will also define new PPP options for the existing protocol definitions, such as stronger authentication and encryption methods.

Specific Objectives:

The main objective of the working group is to produce an RFC or series of RFCs to define the use of other protocols on PPP.

Estimated Timeframe for Completion:

The RFC(s) should be complete during the year.

CURRENT MEETING REPORT

First meeting is planned for the February IETF.

2.3.8 Router Requirements Working Group (rreq)

CHARTER

Chairpersons: Jim Forster/Cisco forster@cisco.com
Philip Almquist/Stanford almquist@jessica.stanford.edu

WG Mailing List: TBD

Description of Working Group:

The Router Requirements Working Group has the goal of rewriting the existing Router Requirements RFC, RFC-1009, and a) bringing it up to the organizational and requirement explictness levels of the Host Requirements RFC's, as well as b) including references to more recent work, such as the RIP RFC and others.

Specific Objectives:

- Produce a draft document for initial comment by the community by the summer of 1990.

Estimated TimeFrame for Completion:

The objective is to have a completed document ready to be made into an RFC by early in 1991.

CURRENT MEETING REPORT

First meeting is planned for the February IETF.

2.4 Routing Area

Director: Bob Hinden/BBN

The major issue in this area is the topic of a standard Internal Gateway Routing Protocol (IGP). The IESG discussed this in detail at the open meeting in Hawaii. We plan to make this topic the focus of a special meeting at the next IETF meeting at Florida State University (February 6-9, 1990).

Because of its importance and its early promise, we have also decided to form a WG to specifically examine at the experimental Border Gateway Protocol (BGP). One possible outcome would be for BGP to eventually replace EGP as the Exterior Gateway Routing Protocol. Another possible outcome might be that the better parts of BGP could become a basis for a new or better EGP.

ROUTING

Bob Hinden

Nov 1989 IETF

"CORE" MODEL

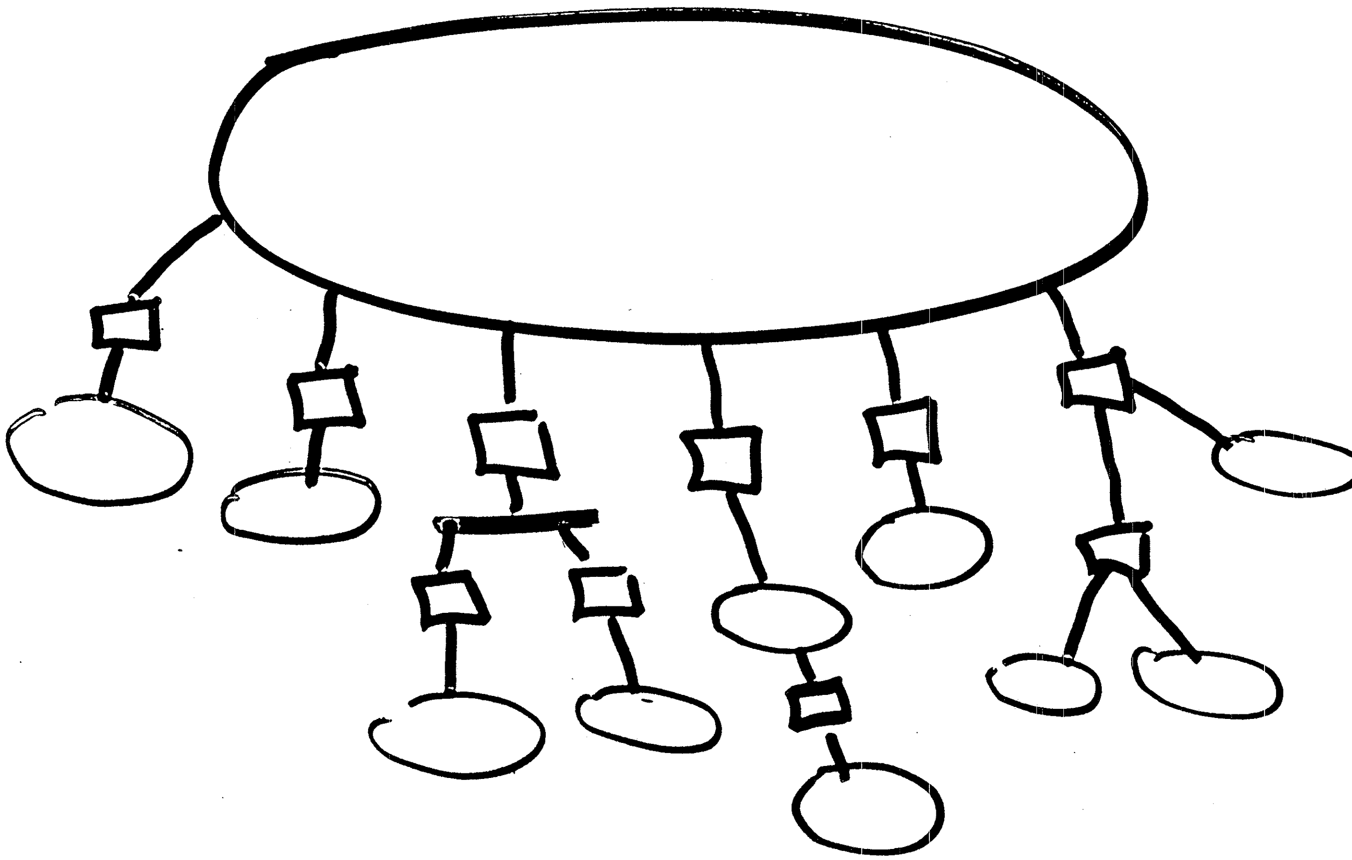

CURRENT INTERNET

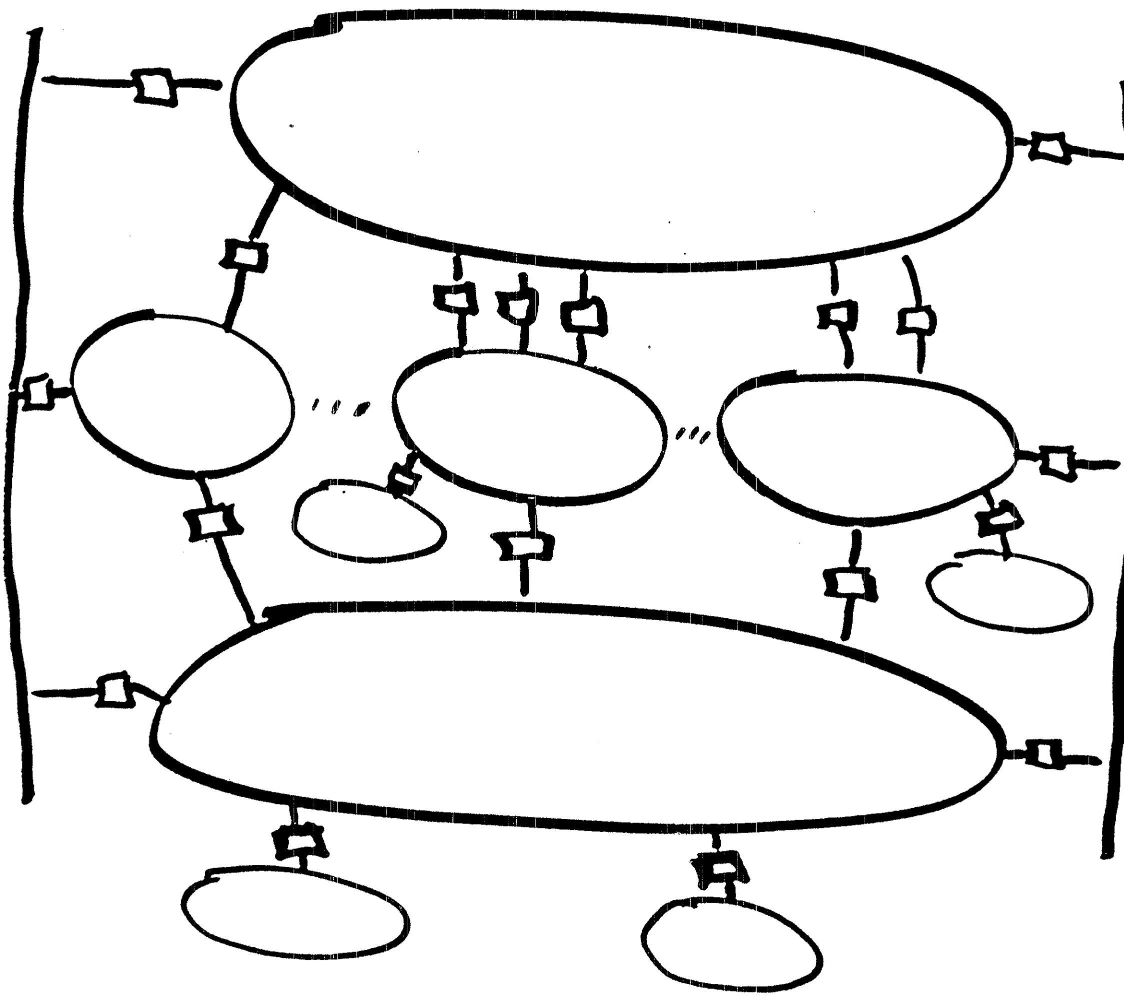

FUTURE INTERNET

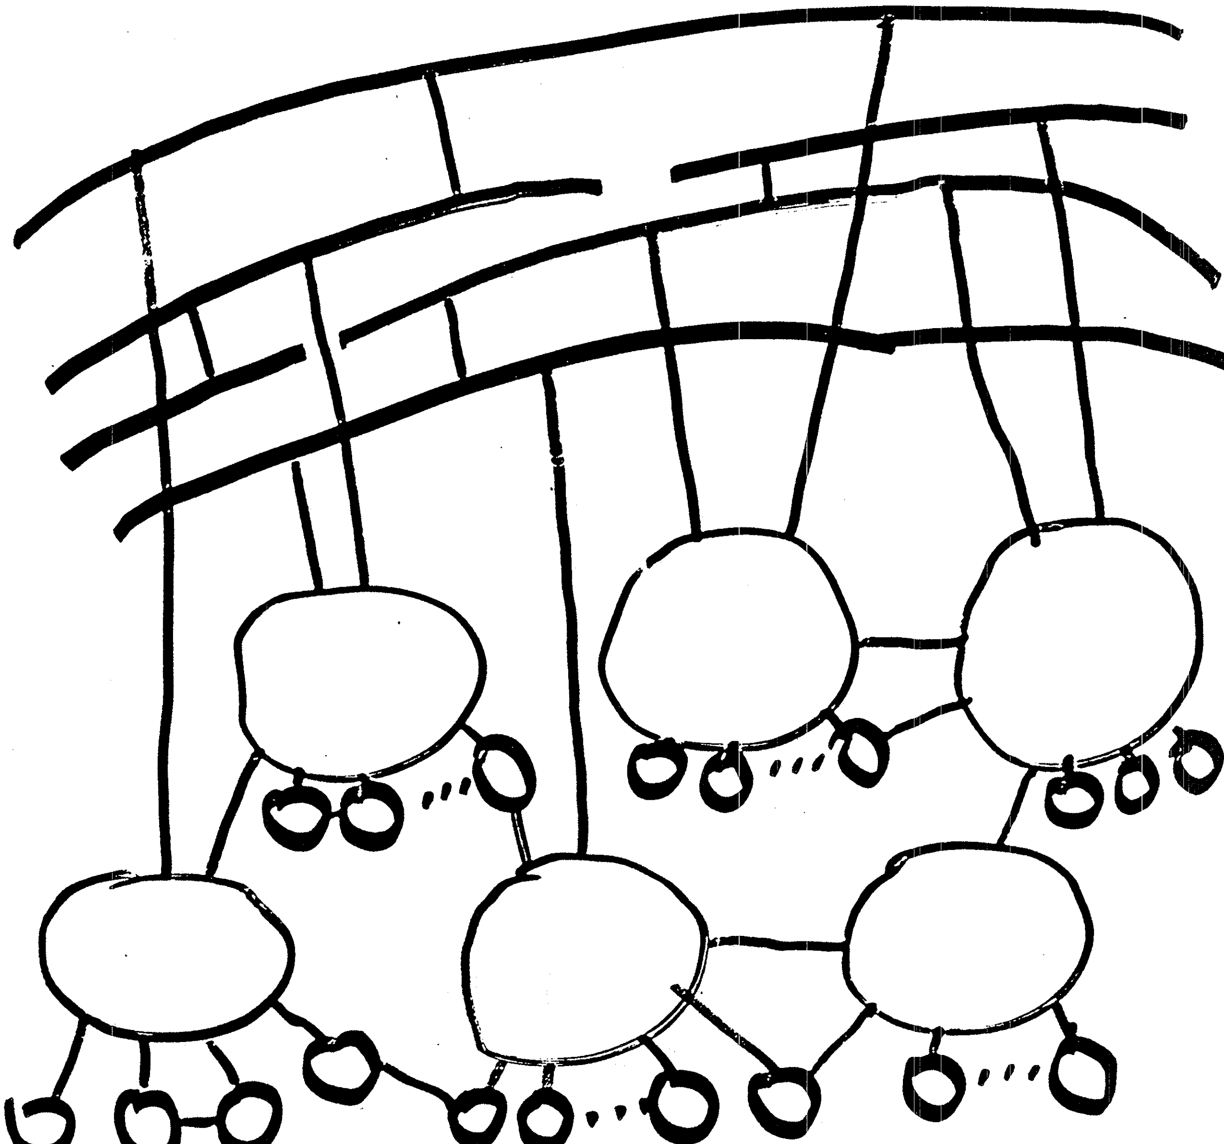

ROUTING ISSUES

- Standard IGP
- Inter-Domain (a.k.a Autonomous System) Routing
- Multiprotocol Routing
- Internet Growth
  - Size
  - Topology
  - Limits ?
- Routing Authentication

ROUTING WORKING GROUPS

- Open SPF (Petry, Moy)
  - OSPF RFC Released as "Proposed Standard"
  - 6 month Comment Period

- EGP3 (Lepp)
  - Work completed
  - RFC ?

- Intercomectivity (Almes)
  - Current Connectivity Issues Backbones and Regionals

2.4.1 Interconnectivity Working Group (iwg)

CHARTER

Chairperson: Guy Almes, almes@rice.edu

Mailing List: iwg@rice.edu

Description of Working Group:

We aim to improve practical inter-autonomous routing in the internet.

Specific Objectives:

1. Monitor the *actual* state of interconnectivity, both among national backbones and among mid-level networks.
2. Monitor policy, as articulated by those responsible for the various national backbone and mid-level networks, with a view of moving toward a common consistent architecture for interconnectivity.
3. Monitor implementation of inter-AS routing, both using current tools and using the coming BGP tools, for engineering sanity.
4. Beginning with the February 1990 IETF, begin to include people outside the United States in our discussions of interconnectivity state, policy, and implementation.

Estimated Timeframe for Completion:

1. The nature of this work is ongoing and periodic. Thus no 'completion date' is given.
2. Beginning with the Winter 1990 IETF, however, specific reports to the Internet community will be given reflecting what we learn each quarter. This periodic report will be of use to the IETF, to FARnet, and to the CCIRN members.

CURRENT MEETING REPORT

Reported by Guy Almes/Rice

Session 1: Discussion of Current Interconnectivity and Filtering

The first session dealt with Interconnectivity in the current Internet, and focussed on:

1. What routing policies were taken by the national backbones.
2. What uses of in- and out-filtering were used to implement them.
3. What uses of static preferences were used to implement them.
4. Whether the ASes of the current Internet form an hierarchy.
5. Whether the ASes of the current Internet support symmetric routes.

Session 2: Discussion of the BGP Usage Draft

The second session dealt with three topics:

- Comments on the draft specification of BGP (RFC 1105),
- The status of current BGP tests, and
- A discussion of the BGP usage draft.

Comments on RFC 1105:

- Remove the notion of one AS being higher or lower in relationship to a neighbor AS. This was seen as having no technical content, and removing it clarifies BGP as supporting a general graph topology.
- Remove mention of explicit ACK to connect.
- Fix the Finite Automaton to allow an implementor to use it to guide a BGP implementation based on the common 4.3 TCP.
- We are open to other changes; send mail to bgp@rice.edu

There followed a brief discussion of current BGP test implementations:

- cisco supports test BGP in versions of its gateway software
- the NSFnet Backbone gateways support it
- the current Cornell gated supports it
- a variant of gated developed by IBM also supports it
- NSF BB is the only implementation that allows connection not over a shared network.
- NSS-to-cisco tested at MERIT, and will be used to carry net-35 in test mode.
- NSS-to-gated-variant at Yorktown, used to pass pseudo-production IBM traffic.

Discussion of BGP usage led by Matt Mathis. For each AS path, a BGP speaker will know the following:

- The next Gateway IP Address,
- The AS path: AS#a AS#b .. AS#z (The next Gateway is a gateway of AS#a),
- Source Protocol: how AS#z derived the route (e.g., IGP, EGP2, or incomplete),
- Metric to the next gateway, and
- List of networks for which the above is valid.

There are two important variations of BGP usage:

- Exterior BGP: Between BG1 of AS#a and BG1 of AS#b.
- Interior BGP: Between BG1 of AS#a and BG2 of AS#a.

Exterior BGP is the simpler case by far. Interior BGP is quite hard, particularly with a RIP IGP implemented without flash updates.

We discussed a pathological case, in which an AS has three BGs: BG1, BG2, and BG3.
We are concerned with what BG3 sees. BG1 has best route, but the best route is currently down. BG2 has fair route, and it is currently up. BG1, BG2, and BG3 are connected via Internal BGP (IBGP).

Now, BG1's exterior route comes up! BG1 has a timer that tells it when to turn on. This situation is difficult because we need to bring the new attractive route into use without creating routing loops either at the Intra- or Inter-AS level.

NOTE: The interconnectivity Working Group has been split into two separate groups: (1) The Interconnectivity Working Group and (2) The BGP Working Group.

ATTENDEES

Almes, Guy
Bagnall, Doug
Brim, Scott
Burgan, Jeffrey
Choy, Joseph H.
Collins, Mike
Elz, Robert
Farinacci, Dino
Fidler, Mike
Froyd, Stan
Fuller, Vince
Garcia-Luna, Jose
Hedrick, Charles
Hinden, Bob
Honig, Jeffrey C.
Karels, Mike
Katz, Dave
Long, Dan
Love, Paul
Mathis, Matt
Medin, Milo
Merritt, Don
Parulkar, Guru
Pomes, Paul
Rekhter, Yakov
Solensky, Frank
St. Johns, Mike
Streeter, Roxanne
Veach, Ross
Willis, Steve

2.4.2 Open SPF-Shortest Path First IGP Working Group (ospf)

CHARTER

Chairpersons: Mike Petry/UMD petry@trantor.umd.edu
John Moy/Proteon jmoy@proteon.com

Mailing List: ospfigp@trantor.umd.edu

Description of Working Group:

The OSPF working group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage muliple vendor implementations.

Specific Objectives:

- Design the routing protocol, and write its specification.
- Develop multiple implementations, and test against each other.
- Obtain performance data for the protocol.
- Make changes to the specification (if necessary) and publish the protocol as an RFC.

Estimated Timeframe for Completion:

We have a complete protocol specification. Implementation experience and performance data should be obtained during the summer of 1989. The specification should be ready for final review by the October-November IETF.

CURRENT MEETING REPORT

Did not meet

2.4.3 Open Systems Routing Working Group (orwg)

CHARTER

Chairperson: Marianne Lepp/BBN, mlepp@bbn.com

Mailing List: open-rout-interest@bbn.com

Description of Working Group:

The Open Systems Routing Working Group is chartered to develop a policy-based AS-AS routing protocol that will accommodate large size and general topology.

Specific Objectives and Milestones:

- Architecture
- Draft Protocol Specification of key elements of the protocol

Estimated Timeframe for Completion:

February 1990

CURRENT MEETING REPORT

Reported by Marianne Lepp/BBN

MINUTES

The working group met for 3 half-day sessions. The first two were editing sessions for the architecture draft paper that will be out this month. The third was planned as an open informational session, but the open steering group meeting was taking place at the same time and took away our audience.

The agenda of the editing session was the last section of the architecture paper: defining the first step in an evolution path to implementing the entire protocol.

For the first phase, we chose a simplified virtual gateway protocol in which only two border gateways participate. We discussed the policy gateway to policy gateway protocol which runs among policy gateways inside an Administrative Domain. It was felt that this is an essential protocol for a first phase implementation. The mechanism for disseminating data, including electing a policy gateway spokesman, determining when data should be sent, and the mechanisms for limiting the extent of its flow, were also determined to be essential. A route query would allow experimentation with policy routes and validation without actually implementing the automated part of policy validation, etc. Finally, we discussed a mixed environment, when source routes could be installed at 'tack' points, while the gateways in-between are doing routing as usual.

ATTENDEES

Marianne Lepp
Michael Little
Lixia Zhang
Noel Chiappa
Martha Steenstrup

2.4.4 Open Distance Vector IGP Working Group (odv)

CHARTER

Chairperson: Charles Hedrick/Rutgers University, hedrick@cs.rutgers.edu
Mailing Lists:

odv@rutgers.edu
odv-request@rutgers.edu

Description of Working Group:

The Open Distance Vector Working Group is chartered to sponsor working on distance vector based routing protocols, and related work.

Specific Objectives:

1. Produce RFC describing IGRP. Should be ready by spring 89.
2. Sponsor and review work comparing distance vector and SPF algorithms. Timing depends upon actions of funding agencies. This is probably at least a one-year task.
3. Design a new distance vector protocol. This is a long-term goal.

CURRENT MEETING REPORT

Reported by Charles Hedrick/Rutgers University

MINUTES

This was an organizational meeting for the ODV group. The first meeting was a large one. (The attendance list is given at the end of this message.) It discussed primarily general issues. There was a brief meeting of a smaller group of people in the evening, to explore doing some actual implementation work.

The first meeting discussed primarily the question of whether there should be an ODV protocol at all. In addition, issues raised by the cisco patent application were discussed. A major part of the meeting was taken for a presentation by Jose Garcia-Luna of some research of his.

Many people would like there to be only one routing protocol. This has obvious advantages in terms of interoperability. Since OSPF is now at the

RFC stage, it has a head start in terms of IETF politics. The question is whether it makes sense to work on another protocol. Raising this issue is about as far as one can go. The IETF charter does not make it possible to prevent a group of people from working on a protocol. So we didn't vote on the question of whether work should proceed. But I will note here that many people were very sceptical.

Part of the problem is that it is difficult to prove in any unambiguous way what protocol is the best way in the long run. Jose Garcia-Luna's simulations attempted to compare SPF and distance vector approaches, but the routing algorithms simulated were not based on the best implementations of either approach. As part of the work of this group, we are going to try to get the resources to carry this work further. (This may actually be a more important activity than designing another protocol.) My feeling is that routing is still an unsolved problem. It is unrealistic to expect progress in this area to stop, leaving some current protocol as "the answer" for all time.

In response to the concern about extra protocols, I believe we are going to proceed as follows:

- Some subset of us will attempt to bring a description of IGRP to the stage of an RFC. The whole issue of whether it should be considered an alternative to OSPF is one for those who care about such issues to negotiate with the IAB. I do not plan to involve myself in that. My feeling is that enough people in the community are using IGRP that it at least makes sense to have a generally available document that describes it. If network politics make it impossible to issue it as an RFC, it will be available as a Rutgers University technical report.
- We will pursue Jose's work. This is more of an attempt to advance the state of the art than to produce an immediate competitor to OSPF. I believe it will be one to two years before anything concrete comes out of this. This work will include analysis as well as protocol design. We will try to avoid producing a protocol unless it worth doing.

There was a discussion about the implications of the IGRP patent application. There was a very strong feeling against an IETF-sponsored protocol that is tied up in patent rights. Some caveats:

- There is precedent for a protocol that involves a patent. The privacy taskforce is advocating an approach to Email that requires a license from RSA, Inc.
- The concern was primarily that it should be possible to distribute public-domain implementation through mechanisms such as the BSD tape, for use by recipients. This does not necessarily rule out all licensing. This request would be consistent with allowing internal use by recipients of the BSD tape, but licensing any products based on it.

We took a straw poll about licensing. 27 people objected to a protocol that involved a license. 3 saw no problem with it. 12 abstained. However it is not entirely clear what this vote meant. My best guess, based on a small number of conversations with individuals, is that the 27 people might be satisfied with a public-domain implementation that allowed free use, but required a license for incorporation into a product. At any rate, I believe that the committee will do everything possible to make any new protocol it designs unencumbered. This means that it will not be based directly upon IGRP. To the extent that it shares the same roots as IGRP, there may still be similarities. However we will try to make sure that we have sources in the literature predating IGRP for any mechanisms that we share with IGRP. Obviously the attempt to produce an RFC for IGRP will not adhere to these guidelines.

Jose Garcia-Luna's presentation was based on a published paper, so I don't intend to describe it here. (I have managed to lose my copy of the paper. Hopefully Jose will send a citation to the list.)

ATTENDEES

Almquist, Philip	Hinden, Bob
Arnold, Susan	Honig, Jeffrey C.
Bagnall, Doug	Huston, Geoff
Baker, Fred	Karels, Mike
Berggreen, Art	Knowles, Steve
Borman, David	Lear, Eliot
Burgan, Jeffrey	Little, Mike
Catlett, Charlie	Long, Dan
Chiappa, Noel	Merritt, Don
Chinoy, Bilal	Miller, David
Choy, Joseph H.	Opalka, Zbigniew
Collins, Mike	Pleasant, Mel
Coltun, Rob	Rosenstein, Mark
Elz, Robert	Rutenberg, Vald
Farinacci, Dino	Schiller, Jeff
Fidler, Mike	Sheridan, Jim
Forster, Jim	Vaudreuil, Greg
Fuller, Vince	Veach, Ross
Garcia-Luna, Jose	Willis, Steven
Gross, Phill	Yasaki, Brian
Hays, Ken	Youssef, Mary
Hedrick, Charles	

2.4.5 PDN Routing Working Group (pdnrout)

CHARTER

Chairperson: Carl-Herbert Rokitansky/Fern University of Hagen
roki@DHAFEU52.BITNET or roki@ISI.EDU

Mailing Lists:

- pdn-wg@BBN.COM: For internal discussions and information exchange between members of the PDN Routing working group.
- pdn-interest@BBN.COM: For information about:
  - Status report and proceedings of the PDN Routing WG
  - Draft proposals of documents and papers
  - Documents and papers published by PDN WG members
  - Important discussion on PDN Routing issues.
- pdn-request@BBN.COM: For people interested in being put on the "pdn-interest" mailing list.

Description of Working Group:

The DoD INTERNET TCP/IP protocol suite has developed into de facto industry standard for heterogenous packet switching computer networks. In the US, several hundreds of INTERNET networks are connected together; however the situation is completely different in Europe: The only network which could be used as a backbone to allow interoperation between the many local area networks in Europe, now subscribing to the DoD INTERNET TCP/IP protocol suite, would be the system of Public Data Networks (PDN). However, so far, no algorithms have been provided to dynamically route INTERNET datagrams through X.25 public data networks. Therefore, the goals of the Public Data Network Routing working group are the development, definition and specification of required routing and gateway algorithms for an improved routing of INTERNET datagrams through the system of X.25 Public Data Networks (PDN) to allow worldwide interoperation between TCP/IP networks in various countries. In addition, the application and/or modification of the developed algorithms to interconnect local TCP/IP networks via ISDN (Integrated Services Digital Network) will be considered.

Specific Objectives and Estimated Timeframe for Completion:

1. Application of the INTERNET Cluster Addressing Scheme to Public Data Networks. (Already done, see produced documents)

2. Development of hierarchical VAN-gateway algorithms for worldwide INTERNET network reachability information exchange between VAN-gateways (Already done, see produced documents)
3. Assignment of INTERNET/PDN-cluster network numbers to national public data networks. (Mapping between INTERNET network numbers and X.121 Data Network Identification Codes (DNICs) (Already done, see produced documents)
4. Assignment of INTERNET/PDN-cluster addresses to PDN-hosts and VAN-gateways according to the developed hierarchical VAN-gateway algorithms (Almost done, see produced documents)
5. Definition of the PDN-cluster addressing scheme as an Internet standard (Already done, [earlier than expected - a case that happens very seldom!] see produced documents)
6. Specification of an X.121 Address resolution protocol (RFC-Draft, expected to be completed by October '89)
7. Specification of an X.25 Call Setup and Charging Determination Protocol (RFC-Draft, expected to be completed by Fall '89)
8. Specification of an X.25 Access and Forwarding Control Scheme (to be written up as an RFC-Draft by Fall '89 or later)
9. Specification of routing metrics taking X.25 charges into account (to be written up as an RFC-Draft by Fall '89 or later)
10. Delayed TCP/IP header compression by VAN-gateways and PDN-hosts (new objective, will be considered Fall '89 or later)
11. Provide a testbed for worldwide interoperability between local TCP/IP networks via the system of X.25 public data networks (PDN) (starting June '89)
12. Implementation of the required algorithms and protocols in a VAN-BoX (Test version towards End '89)
13. Interoperability between ISO/OSI hosts on TCP/IP networks through PDN (1989/90)
14. Consideration of INTERNET Route Servers (1990)
15. Interoperability between local TCP/IP networks via ISDN (1990)
16. Development of Internetwork Management Protocols for worldwide cooperation and coordination of network control and network information centers (starting 1990).

CURRENT MEETING REPORT

Reported by Carl-Herbert Rokitansky/Fern University of Hagen

AGENDA

"A WORLDWIDE INTERNET - What's missing ? - What do we need to do ?"

- Introduction
- Background information (European network situation, current status of X.25 Research Network, future plans, etc.)
- Network situation in some oversea countries (Argentina, Australia, Brazil, Indonesia, Israel, Japan, etc.)
- Status report on BBN-VAN-GATEWAY (butterfly replacement, EGP, etc.) (Chet Birger, BBN) - Discussion
- Discussion on submitted RFC-Drafts:
  1. Internet Cluster Addressing Scheme
  2. PDN Cluster and Hierarchical VAN-Gateway Algorithms
  3. Assignment / Reservation of Internet Network Numbers to National X.25 Public Data Networks (DNICs)
- Hierarchical Scheme for the Assignment of PDN-Cluster Addresses, Draft RFC (Roki)
- X.121 address resolution protocol, Draft RFC (Roki); detailed technical discussion
- Access control and reverse charging on international X.25 connections, draft proposal; detailed technical discussion
- VAN-BoX (Specification of required protocols using formal description techniques)
- Coordination of international PDN Routing performance tests with partners in: Germany, Sweden, United Kingdom, Austria, Italy, Norway (Zone 2), USA (Zone 3), Japan (Zone 4), Australia, Indonesia (Zone 5)
- Detailed technical discussion and definition of a test plan
- Discussion on documents to be published by members of the PDN Routing WG
- Assignment of action items
- Miscellaneous (mailing lists, etc.)

MINUTES

Report of the Open PDN Routing WG Meeting, IETF, November 2, 1989

Network Situation in Europe and Some Overseas Countries (reported by Roki):

- German X.25 Research Network:

According to the plans, a German X.25 Research Network (X.25 WIN) will be installed and operated by the German PTT, starting January 1990. A large number of German universities and research institutes will be connected to this X.25 Research Network at fixed costs. A gateway to the German DATEX-P network will allow interoperation with the worldwide system of X.25 Public Data Networks (PDN).

Due to the charging policy for the X.25 Research Network (fixed costs), most universities, having local TCP/IP networks, are especially interested in exchanging TCP/IP datagrams with each other through this X.25 research network.

The PDN Routing and VAN gateway algorithms, which have already been published or are currently specified, are expected to improve the interoperability between these local TCP/IP network and to reduce the amount of network management significally.

- International X.25 Interconnect (IXI):
  The Commission of the European Communities (CEC) and the Dutch PTT on behalf of all European PTTs have signed an agreement for an international X.25 infrastructure in Brussels, in September 1989.
- NORDUNET:
  This large network interconnects hosts in the Scandinavian countries (Denmark, Finland, Iceland, Norway and Sweden). A satellite connection from the NORDUNET (Sweden) to the US TCP/IP Internet exists. The supported networks are: EARN, NSFnet/Internet, SPAN/HEPNET, OSI Pilot Services, EUnet, etc. The supported protocols are: TCP/IP, DECnet, ISO IP, X.25, RSCS, etc. The Swedish Institute of Computer Science (SICS) will participate in international PDN tests performed by the PDN Routing WG.
- TCP/IP Networks in Oversea Countries:
  Several TCP/IP networks exist already in Australia, Japan, Argentinia, Brazil, etc.

  With the support of DLR/FernUni a TCP/IP network is currently installed in Indonesia at LAPAN (Air and Space Research Establishment) and other national agencies (BPPT, etc.). Satellite communications will be provided by VSAT using X.25 protocols. It is intended to integrate these networks in international PDN tests.

Status Report on BBN-VAN-GATEWAY (Zbigniew Opalka, BBN):

The LSI-11/23 has been replaced by a butterfly gateway, which runs EGP for network reachability information exchange (with CNUCE, Italy, etc.). The BBN-VAN-GATEWAY will participate in international PDN tests (contact: Zbigniew Opalka and Chet Birger).

X.121 Address Resolution Protocol, Draft RFC (Roki), technical discussion:

The proposal for an X.121 address resolution protocol (developed at the FernUni), which is currently being written up as an Internet Draft, has been discussed in detail.

For a dynamic routing of Internet datagrams through X.25 Public Data Networks (PDN) an X.121 Address Resolution Protocol (X.121 ARP) is required to determine the mapping between the 32-bit Internet address of a PDN-host/ VAN-gateway and its X.121 address on the X.25 network. This X.121 address resolution can be performed by:

- a table lookup on the local host/gateway
- a fast X.121 address resolution using the user data field in X.25 calls
- an X.121 address resolution retrieved from a remote X.25 host/gateway
- by information exchange with an X.121 address resolution server

X.121 Address Resolution by Table Lookup:

The mapping between the Internet address and the corresponding X.121 address is contained in a data file ("XARP.PDN"). X.121 address resolution is simply performed by a local table lookup. A standard for the format of this file ("XARP.PDN") will be specified, so that it can be distributed to other PDN-hosts and VAN-gateways (by FTP) for X.121 address resolution. Whenever a PDN-cluster address is assigned to some PDN-host or VAN-gateway, the corresponding X.121 address will be updated in the (original) XARP.PDN file.

The following format has been discussed for the XARP.PDN file:

<IP-ADDRESS>:<X.121-ADDRESS>[;<FACILITIES>[;<ACCESS CONTROL>]][;<COMME

It has been agreed, that the fields <FACILITIES>, <ACCESS CONTROL> and <COMMENT> (e.g., containing the host/gateway name) should be optional.

Packet Identifier in X.25 Call Setup Packets:

According to RFC-877, IP datagrams are identified by a value of CC (hex) in the first octet of the user data field in an X.25 call setup request packet.

A (new) interpretation of the bits of the first octet in the user data field has been discussed:

```
1 1 0 0 x x x x
        | | | |
        | | | +---- 0: NO Reverse Charging
        | | |       1: Reverse Charging requested
        | | +------ 0: NO ARP
        | |         1: ARP
        | +-------- 0: Extended User Data Field (20 octets)
        |           1: Regular User Data Field (packet identifier only)
        +---------- 0: NO IP datagrams
                    1: IP datagrams will be transmitted
```

According to the new interpretation of the bits of the first octet of the user data field, the following packet identifiers would be used for:

Type		hexadec.		binary	Comment
Fast X.121 Address Resolution	(FXARP)	C2		11000010	ARP
X.121 Address Resolution	(XARP)	C6		11000110	regular, ARP
Piggy-Backed Fast X.121 ARP	(PFXARB)	CA	(CB)	11001010	IP, ARP
IP Datagram (RFC 877)	(IP)	CC	(CD)	11001100	IP, regular
(X.121 Address Res. using UDP	(IP)	CC	(CD)	11001100	IP, regular)

NOTE: A possible conflict with already defined values of the first octet of the user data field (C0-CF) for other applications has been mentioned. In this case the following alternative has been discussed:
The first octet of the user data field would always contain CC (hex), according to RFC 877. Then, the bits of the second octet would indicate the packet type as follows (similar to the specification above):

```
0 0 x x x x x x
| | | | | | | |
| | | | | | | +---- 0: NO Reverse Charging
| | | | | | |       1: Reverse Charging requested
| | | | | | +------ 0: NO ARP
| | | | | |         1: ARP
| | | | | +-------- 0: Extended User Data Field (20 octets)
| | | | |           1: Regular User Data Field (packet identifier only)
| | | | +---------- 0: NO IP datagrams
| | | |             1: IP datagrams will be transmitted
| | | +------------ 0: Request
| | |               1: Reply
| | +-------------- 0: NOT unsolicited (request or reply)
| |                 1: Unsolicited
| +--------<reserved>: for future use
+----------<reserved>: for future use
```

Fast X.121 Address Resolution Using the User Data Field in X.25 Calls:

A fast X.121 address resolution can already be performed during the call setup time by specifying the X.121 address resolution request/reply information in the user data field of the X.25 call setup request/accept packets, in the following format:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Packet Identifier     | 0  0  x  x| (Pointer) |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Max. X.121 length = 8 | Length of IP address=4|
+-----------------------+-----------------------+
| IP address of sender                          |
+-----------------------+-----------------------+
| IP address of target                          |
+-----------------------+-----------------------+
|Len. X.121 |           X.121 address of target |  in X.121 ARP reply or
|         (encoded in quartets, padded with 0's) |     unsolicited X.121 ARP
+-----------------------------------------------+
```

Note: that the X.121 address of the sender is contained in the X.25 packet header (calling DTE address).

Example:

```
194  22  4  14 188  1 255  1 189  42 128  1  227  17  6  23  0  2  80  0
```

Fast Reply

X.121 ARP	IP X121 Ptr 6	IP address Len.	IP address of sender	X.121 of target	address (311061700025) of target

X.121 Address Resolution Protocol for X.25 Hosts/Gateways:

The following X.121 Address Resolution Protocol (similar to the Ethernet Address Resolution Protocol (RFC-826)) for Internet hosts and gateways, which are directly connected to an X.25 Public Data Network, has been discussed:

The first (second) octet of the user data field in an X.25 call setup request packet contains C6 (06). When the X.25 connection is established, then, the following X.121 address resolution request/reply, is transmitted in the data field of subsequent X.25 packets:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Hardware Address Space (PDN)                  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Protocol Address Space (DoD_Internet)         |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Max. X.121 length = 8 | Length of IP address=4|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode                                        |
+-----------------------------------------------+
|Len. X.121 |           X.121 address of sender |
|        (encoded in quartets, padded with 0's) |
+-----------------------------------------------+
| IP address of sender                          |
+-----------------------------------------------+
|Len. X.121 |           X.121 address of target |
|        (encoded in quartets, padded with 0's) |
+-----------------------------------------------+
| IP address of target                          |
+-----------------------------------------------+
```

X.121 Address Resolution using UDP:

In case, that the X.121 address resolution cannot be obtained from a host or gateway directly connected to the X.25 network, the X.121 address resolution request/reply messages are sent to an X.121 address resolution server in IP datagrams using UDP. The data field contains the X.121 address resolution request/reply as specified above.

Access Control and Reverse Charging on International X.25 Calls:

An "X.25 Call Setup and Charging Determination Protocol" has been developed at the FernUni and is currently being specified as an Internet draft. This protocol will allow reverse charging on international X.25 calls, which is of special importance for a worldwide interoperability of TCP/IP networks.

Also, the specification of an access control scheme has been discussed, but most people suggested that access control should be done by higher layers.

Implementation of the Proposed Algorithms in a VAN-BoX (or on a Workstation):

The IETF-PDN Routing working group has already developed and specified most of the required PDN addressing schemes and gateway algorithms to allow a dynamic routing of TCP/IP datagrams through the worldwide system of X.25 Public Data Networks (PDN). The required algorithms and protocols include:

- PDN-cluster addressing scheme: published ICCC'88 and RFC Draft
- Hierarchical VAN-gateway algorithms: published in ITG/GI'89 and RFC Draft Assign. and Res. of PDN-cluster net no.: Internet Draft to be published as RFC
- Assign. and Res. of PDN-cluster addr.: being finished as an Internet draft
- X.121 Address Resolution Protocol: being finished as an Internet draft
- X.25 Call Setup and Charging Determ: being written up as an Internet draft
- Modified EGP2 or EGP3 between VANs: currently in progress to be defined
- Delayed TCP/IP header compression: will be considered (new objective)

By putting all these pieces together, it is intended to implement these algorithms, with support of the gateway companies (BBN, Proteon, SUN, 3COM, ACC, cisco) and eventually the University of Salzburg and the University of Tokyo, in a small "VAN-BoX" (and on a workstation) with an Ethernet and an X.25 interface. By placing this "VAN-BoX" between a local TCP/IP network and an X.25 public data network, the implemented gateway algorithms will automagically exchange network reachability information to provide worldwide INTERNET interoperability between local TCP/IP networks through X.25 Public Data Networks.

Coordination of International PDN Routing Performance Tests:

The developed PDN addressing schemes and VAN-gateway algorithms will be tested with participating sites in the following countries:
Zone 2 (Europe):

Germany:	Fern University of Hagen (all VAN-gateway levels)
	GMD, St. Augustin (DFN-Gateway)
	University of Dortmund (UUCP-Gateway)
	University of Karlsruhe (BELWUE)
	University of Stuttgart (BELWUE)
Austria:	University of Salzburg
Finland:	University of Helsinki (NORDUNET)
Italy:	CNUCE, Pisa *
Norway:	NTARE, Oslo, (NORDUNET) *
Sweden:	SICS, Stockholm (NORDUNET)
UK:	Portsmouth Polytechnic
	University College London (INTERNET Gateway) *

Zone 3 (North America):

USA:	ACC *
	BBN, Cambridge, MA
	CISCO, Menlo Park, CA *
	PROTEON *
	SRI, Menlo Park, CA *
	SUN, Mountain View, CA *
	3COM *

Zone 4 (Asia):

Japan: University of Tokyo

Zone 5 (Pacific):

Australia: CSIRO
Indonesia: LAPAN

Zone 6 (Africa): Egypt ?

Zone 7 (South America): Argentina ?, Brazil ?

(* ... intended, but not yet agreed) (? ... these countries will be contacted for participation, to have at least one representative site for each zone).

First tests have already been started within Germany. International PDN-tests are expected to start in January '90 between BBN and sites in Europe, Australia, Japan and Indonesia.

PDN-cluster addresses, according to the developed hierarchical scheme, will be assigned to all participating sites, in Jan '90. An appropriate application form has been prepared recently. SRI-NIC will be informed about the assigned PDN-cluster addresses. The XARP.PDN file containing the mapping between the Internet PDN-cluster address and the corresponding X.121 address will be updated after each assignment.

Assignment of action items:

Stahl: Check assignment and specification of INTERNET/PDN-cluster network numbers for national public data networks in the North America cluster for correctness (O3, Jan '90)

Roki: Finish Internet Draft "Addressing Scheme for the Assignment of INTERNET/PDN-Cluster Addresses to VAN-Gateways and PDN-Hosts" for submission to the IETF Chair and Reviewers (O4, Jan '90).

Roki: Check the coding of the first octet of the user data field in X.25 call setup request packets (related to o6, fast X.121 address resolution protocol, Jan '90).

Roki: Finish Internet Draft "X.121 Address Resolution Protocol", for submission to the IETF Chair and Reviewers (O6, Jan '89).

Roki: Continue Internet Draft "X.25 Call Setup and Charging Determination Protocol" (O7, expected to be completed by Feb '90)

Roki: Perform international PDN-tests according to the developed PDN-cluster addressing scheme and hierarchical VAN-gateway algorithms between USA (BBN) and sites in Europe (Fern University of Hagen, University of Dortmund, University of Salzburg, Portsmouth Polytech (UK), SICS (Sweden), etc.), Australia, Japan and Indonesia, starting January '90 (O11).

Comments, suggestions and contributions to the work being done in the PDN Routing working group are highly appreciated.

ATTENDEES

Berggreen, Art	Opalka, Zbigniew
Carvalho, Charles	Rokitansky, Carl-Herbert
Cook, John	Stahl, Mary
Malkin, Gary	Youssef, Mary N.

2.4.6 IS-IS for IP Internets Working Group (isis)

CHARTER

Chairperson: Ross Callon, callon@erlang.dec.com

Mailing Lists:

isis@merit.edu
isis-request@merit.edu

Description of Working Group:

The IETF IS-IS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Specific Objectives:

1. Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.
2. Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.
3. Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

Estimated Timeframe for Completion:

We intend to have completed objectives 1 and 2 by February, 1990.

CURRENT MEETING REPORT

The February IETF is the first and the last meeting.

2.5 Network Management Area

Reported by David Crocker/DEC

The Network Management IETF Area has recently seen a flurry of activity and coalescence. Each of its three major areas has had development progress. In addition to the technical work, there appears to be a degree of stabilization to the specification process for network management.

The area is broadly divided into three technical domains: SNMP-related protocol issues, CMIP-related protocol issues, and MIB-related data structure definitions. The MIB-related work further sub-divides into Transmission media – broadly defined as anything below the IP layer – and the rest of the MIB.

The SNMP Working Group has tried to keep the SNMP protocol and the SMI framework for data structures completely stable, in order to minimize operational impact, so that the focus of their work has been to upgrade the core MIB, with 100 variables, up to about 170 variables. The core MIB was restricted to pure TCP/IP issues, except for very minor host-specific information. This emphasis has been retained in the upgrade, which is called MIB II.

The NetMan Working Group, sometimes referred to as the CMOT Working Group, has renamed itself to OSI Internet Management Working Group. The group is continuing to pursue long-term use of the OSI CMIP protocol. A current debate is between the currently-published CMOT specification, versus a revision which would use the full CMOT, full SMI, and the full OSI upper stack – as opposed to operating over the lightweight Presentation Layer, as currently defined. This would rely upon ISO advancing CMIP to full International Standard, which is expected to happen shortly.

The debate between SNMP and CMIP/CMOT is being conducted by demonstration and use, rather than by direct discussion. Each group, in fact, is attempting to take into consideration the possible use of work by the other group. The best example of this is that the SNMP group's enhancement to the core MIB appears to be reasonable for adoption by the OIM group, which expects to adopt it, after technical review.

The "debate" will, however, stay interesting, given that the SNMP WG also has defined OSI management variables and has demonstrated management of an OSI stack using SNMP...

An open issue is the need to reconcile any ISO OSI MIB variables with the relevant SNMP WG OSI MIB variables.

In addition to the core MIB enhancement, there has been a flurry of specification efforts for media-related modules, such as FDDI and X.25. These fall under the broad category of "Transmission" MIB and an oversight group has been formed to coordinate Transmission MIB development. (It should be noted that some of this is far from straightforward. Besides the predictable issue of looking for variables that are common to more than one medium, and therefore are candidates for a 'generic' transmission MIB sub-tree, some media can be structured into "mesh" functionality, so that a pure hierarchy of information is not possible. FDDI is particularly ripe with this unfortunate opportunity.)

Other MIB efforts have begun for Remote management and Bridge management. (It should be noted that the MIB II enhancement to the core also created the first Application MIB variables. In this case, the application is - either reflexively or recursively - SNMP, itself.) The Remote WG effort will focus upon remote control of network monitors - e.g., devices which promiscuously capture LAN packets. The Bridge WG will focus upon upon the Filtering Repeater technology which often is used as an alternative to IP routers.

The AlertMan WG is discussing asynchronously-generated information, also known as traps or events. There is a trade-off between complexity in the remote agent, versus timeliness of information and polling overhead. If the remote agent is kept minimal, then it can have no intelligence to know when to send information, except when explicitly queried by the management station. In the extreme, the constant polling for information can keep a station too busy and can create excessive network overhead. On the other hand, requiring complex rule-processing by the agent will make the resulting agent software substantially more complex, limiting the range of platforms that can provide it. The AlertMan WG is trying to walk the resulting narrow line.

Though not yet formed, another group is developing. The range of specification efforts has the potential for unexpected and undesirable interactions, such as between MIB variables. Consequently, I am creating an advisory group to assist with coordination of the MIB specification(s) and to resolve any technical conflict. I hope to have it in place by the next IETF.

Network Management

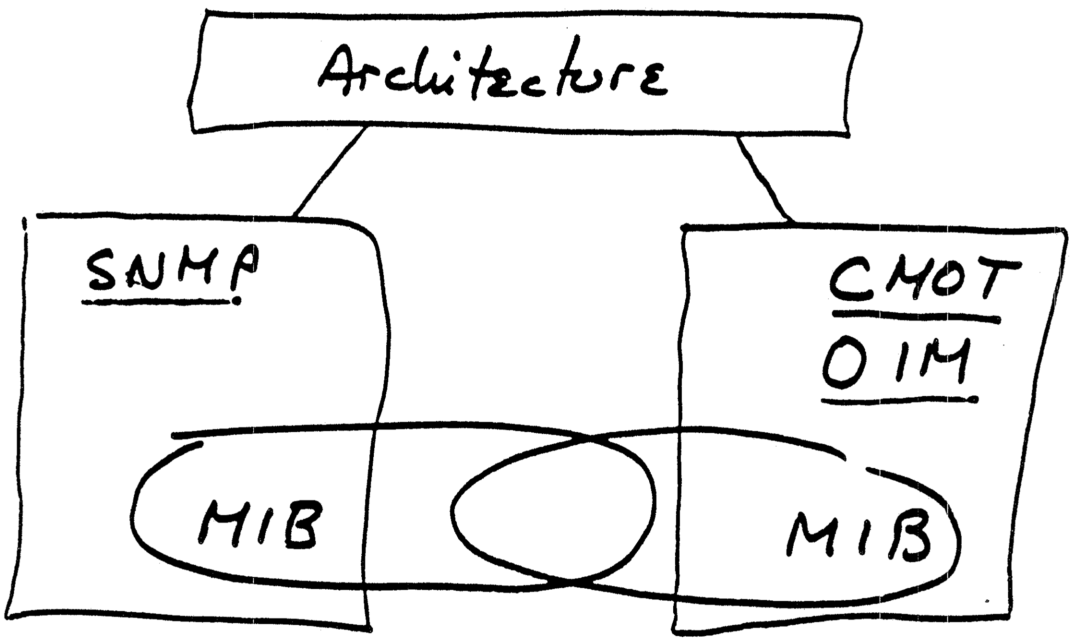

Architecture

- Monitor MIB collisions
- Monitor MIB "cleanliness"
- Resolve MIB conflict
- Review protocol enhancements

Protocols

o SNMP + Core MIB

- M Rose, et al

o OSI Internet Mgmt

- NetMan → OIM
- L. Labarre (+ B. Handspicker)

MIB

- Core (S. Nimp)
- Alert (L. Steinberg)
- Transmission (Fizz Mib) (J. Cook)
  - Ethernet ✓
  - X.25 ✓
  - TR < R/B ✓
  - T1
  - BitSync
  - FDDI (C. Feil? J. Case?)
- Spanning Tree (Satish)

Other

- NOC Tools (B. Stine, B. Enger)
- Mgmt Svc I/f (O. Newkerk, J. Case?)
- SNMP Authentication (J. Schiller)

2.5.1 Alert Management Working Group (alertman)

<u>CHARTER</u>

Chairperson: Louis Steinberg/IBM, louiss@ibm.com

Mailing Lists:

alert-man@merit.edu
alert-man-request@merit.edu

Description of Working Group:

The Alert Management Working Group is chartered with defining and developing techniques to manage the flow of asynchronously generated information between a manager (NOC) and its remote managed entities. The output of this group should be fully compatible with the letter and spirit of SNMP (RFC 1067) and CMOT (RFC 1095).

Specific Objectives:

1. Develop, implement, and test protocols and mechanisms to prevent a managed entity from burdening a manager with an unreasonable amount of unexpected network management information. This will focus on controlling mechanisms once the information has been generated by a remote device.
2. Write an RFC detailing the above, including examples of its conforment use with both SNMP traps and CMOT events.
3. Develop, implement, and test mechanisms to prevent a managed entity from generating locally an excess of alerts to be controlled. This system will focus on how a protocol or MIB object might internally prevent itself from generating an unreasonable amount of information; examples of such techniques might include limiting number of alerts per time period, delayed reporting of "good news" (as in the link up sgmp trap on NSFNET), or the use of thresholds.
4. Write an RFC detailing the above. Since the implementation of these mechanisms is protocol dependent, the goal of this RFC would be to offer guidance only. It would request a status of "optional".

Estimated Timeframe for Completion:

A draft of the first RFC (alert flow control) will be written and reviewed by the July IETF meeting, with final review expected at the October IETF meeting. The second RFC draft will be submitted for initial review at the October IETF meeting. A date for final review of this document has not yet been determined.

CURRENT MEETING REPORT

Reported by Louis Steinberg/IBM

AGENDA

1. Introduction
   (a) Who's who
   (b) General administrivia
      - Attendance
      - Someone to write up minutes
2. Flow Control Draft
   (a) Final questions, comments, protests
   (b) Did I leave anyone off the credits (algorithm was...)
      - attend 2 meetings OR
      - contribute ideas to e-mail
   (c) time to follow up on status...looking for RFC, status= recommended
3. Alert Generation Draft
   (a) John Cook to be primary author...others interested?
   (b) format discussion
   (c) questions (and answers) raised about format
   (d) Specific implementations– volunteers to write up one each
      i. thresholds on counters (uni directional)
      ii. time based hysteresis
         - delayed reporting good as a better refined version
      iii. value based hysteresis (stored with threshold)...esp on gauges
      iv. "pin-per-threshold"
      v. adaptive thresholds
      vi. others?

MINUTES

1. Call to order, introductions, status, questions, complaints
2. Asynch Generation Control Doc.
   (a) short overview of current DRAFT
   (b) Discussion of local vs. remote log, addendum to doc.
      - decision to keep local required log as necessary for reliability.
   (c) Reaffirmed decision to have alert log optionally writeable/ deleteable.
   (d) Decision (majority vote) to standardize result of log entry to a "full" log; alert log will wrap on "overflow", deleting oldest logged alert prior to adding new alert.
   (e) Discussion to allow agent to automatically reset feedback "EnableAsynchAlerts" after a timeout period. Felt by many to be unadviseable, but that document would not prevent it.

unadviseable, but that document would not prevent it.

(f) Discussion of Log Table format. Since Logged alert is a construct type, the encapsulating type must be a construct. Constructed Octed String is not allowed in the SMI (RFC 1065?), so the encoding must be of type OPAQUE.

3. Alert Generation Document
   (a) General discussion of objectives and format for use in intro
      - implementor's guide, kinds of alert controls, how they are sent and to whom
   (b) discussion of alert based environments for use in intro.
      i. are alerts directed by type
      ii. severity imbedded in alert or determined by manager
         - opinion that alerts often should not attempt to convey severity
      iii. Do collecting agents issue their own alerts?
   (c) Specific techniques for controlling alert generation (for use in techniques section)
      i. thresholds on counters (uni directional)
      ii. time based hysteresis
         - delayed reporting good as a better refined version
      iii. value based hysteresis (stored with threshold)...esp on gauges
      iv. "pin-per-threshold"
      v. adaptive thresholds
      vi. others???
   (d) Authorship
      - DRAFT will be prepared by John Cook
   (e) Agreement to include specific techniques and experiences of each member/vendor (with advice?)
   (f) vendors will be encouraged to advertise their alert-related mib objects, their encodings, and details of use in a standard format

ATTENDEES

Carvalho, Charles	Kerby, Kathy
Cook, John	Malkin, Gary
Crocker, Dave	McCloghrie, Keith
Easterday, Tom	Newkerk, Oscar
Froyd, Stan	Norton, Bill
Gerlach, Chuck	Oattes, Lee
Handspicker, Brian D.	Westfield, Bill
Hunter, Steven	Wilder, Bruce
Joshi, Satish	

2.5.2 OSI Internet Management Working Group (oim)

CHARTER

Chairpersons: Lee LeBarre/Mitre cel@mbunix.mitre.org
Brian Handspicker/DEC bd@vines.dec.com

Mailing Lists:

oim-request@mbunix.mitre.org
oim@mbunix.mitre.org

Description of Working Group:

- Specify management information and protocols necessary to manage IP-based and OSI-based LANs and WANs in the Internet based on OSI Management standards and drafts, NIST Implementors Agreements and NMF Recommendations.
- Provide input to ANSI, ISO, NIST and NMF based on experience in the Internet, and thereby influence the final form of OSI International Standards on management.

Specific Objectives:

1. Develop implementors agreements for implementation of CMIP over TCP and CMIP over OSI.
2. Develop extensions to common IETF SMI to satisfy requirements for management of the Internet using OSI management models and protocols.
3. Develop extensions to common IETF MIB-II to satisfy requirements for management of the Internet using OSI management models and protocols.
4. Develop prototype implementations based on protocol implementors agreements, IETF OIM Extended SMI and Extended MIB.
5. Promote development of products based on OIM agreements.
6. Provide input to the ANSI, ISO, NIST and NMF to influence development of OSI standards and implementors agreements.
7. Completion of the following drafts:
    - Implementors Agreements
    - Event Management
    - SMI Extensions
    - MIB Extensions
    - OSI Management Overview
    - Guidelines for the Definition of Internet Managed Objects

Estimated Timeframe for Completion:

Current specific objectives should be completed by December 1990.

CURRENT MEETING REPORT

Reported by Brian D. Handspicker/DEC

MINUTES

The November 2, 1989 meeting of Netman/CMOT WG was opened by co-chair Brian Handspicker. Lou Steinberg was appointed recording secretary for this meeting.

1. NEW GROUP NAME

   The group has changed its name to better reflect its charter to "OSI INTERNET MANAGEMENT". The charter will be clarified to reflect our goals to specify management of IP-based and OSI-based local area and wide area networks in the Internet. The management recommendations specified by this group will be based on OSI management standards and working drafts, NIST implementors agreements and Network Management Forum recommendations.

   For the most part, this group is not defining new standards, but rather is recommending how existing OSI specifications and implementors agreements can be used for the management of the Internet.

2. NEW DOCUMENTS

   To follow through on this charter, five documents will be generated and circulated by the end of 1989:

   - Implementors Agreements
   - Event Management
   - SMI Extensions
   - MIB Extensions
   - OSI Management Overview
   - Guidelines for the Definition of Internet Managed Objects

   These documents will all have the super-title: OSI Management for the Internet. These documents will be circulated as Internet Drafts with the intention that prototypes of each these agreements will be completed and tested before each is proposed as an Internet RFC. It is hoped that this testing can be completed by March/April 1990.

Throughout the creation and testing of these drafts the OIM WG will attempt to maintain close alignment with the Alert-Man and Management Services Interface WGs.

(a) Implementors Agreements

The implementors agreements will specify protocol, SMI and MIB agreements. The protocol agreements will reference the new IS version of ISO CMIP. The IS CMIP is expected to be registered in early 1990. Experts have estimated that it will take about 3 man-weeks to align a DIS-based CMIP implementation with the IS draft. This was considered to be insignificant compared to the value of providing initial CMOT products based on the IS. In addition, the protocol agreements will be drafted to specify both CMIP Over lpp over Tcp (CMOT) and CMIP over full OSI stack (CMIP). In either case, the Application Layer protocol is identical. The SMI agreements will reference the Internet extended SMI. The MIB agreements will reference the Internet extended MIB-II.

There is the potential for future work on a version of CMIP that runs on top of full ISO Session and Presentation on top of TCP instead of LPP. This may provide improved interoperability between CMOT and CMIP implementations. This may not be necessary if dual stack systems become popular. This issue will not be addressed in the current documents.

The implementation examples in the appendix of the current CMOT document will be retained in the new Implementors Agreements.

(b) Event Management

An Event Management Model has been proposed which aligns with current OSI Event Management and Reporting. Some concern was expressed that the OIM Event Management Model align with the work being done within the Alert-Man WG. In addition there is an opportunity to align SNMP traps and OIM events codes and semantics.

(c) SMI

The SMI Extensions document will reference the current Internet SMI and then specify extensions as necessary to support OSI Management of the Internet. In addition, the SMI document will reference the current ISO version of SMI in an attempt to align with ISO.

(d) MIB

Tentatively, the MIB extensions document will reference the current Internet MIB-II specification. We currently do not know of anything in MIB-II that causes problems to CMOT. This should be carefully reviewed by OSI experts. In addition, this MIB document will define extensions necessary to align with OSI Management. These extensions will include: DistinguishedAttributes for MIB-II "objects" and events. There is some concern that MIB-II should not include an in-line version no. in the variable codes. We were assured that the in-line version no. was not defined in MIB-II. The full MIB-II and all extensions defined in this document will be mapped into the ISO Template language.

After this MIB document, protocol groups are not expected to define new MIBs or MIB extensions. It is expected that as new objects are defined by other working groups (e.g. OSI) the management information associated with those objects will be specified by the WG that defines the new object.

(e) OSI Mangement Overview

The Overview sectionand the Examples appendix of the current CMOT document will be retained in a new Overview document.

(f) Guidelines for the Definition of Internet Managed Objects The IETF wrapup (closing plenary) participants recommended that the management groups write guidelines for defining managed objects. This will help the non-management groups (e.g. OSI) define the managed objects associated with their services. This document is not the same as the ISO GDMO draft. This document is specific to the IETF and may point to other document (such as GDMO) as additional reading.

3. INTEROPERABILITY TESTING

There is vendor interest in availability of an interoperability testing lab. DEC is willing to set up and run such a lab in the next few months. They are currently looking for facilities on the West Coast. Vendors interested in participating in such a lab should contact Dave Crocker.

HP will solicit comments on desired test cases and produce a document specifying test scripts.

HP has offered to host the next OIM meeting focusing on interoperability. This meeting will likely be in January in the Bay Area.

ATTENDEES

Halcin, Tom
Handspicker, Brian
Joshi, Satish
Kerby, Kathy
Nadler, Dennis
Newkerk, Oscar
Norton, Bill
Robertson, Jim
Steinberg, Louis
Wilder, Bruce

2.5.3 Management Services Interface Working Group (msi)

<u>CHARTER</u>

Chairperson: Oscar Newkerk/DEC, Newkerk@decwet.dec.com

Mailing List: TBD

Description of Working Group:

The objective of the Management Services Interface Working Group is to define a management services interface by which management applications may obtain access to a heterogenous, multi-vendor, multi-protocol set of manageable objects.

The service interface is intended to support management protocols and models defined by industry and international standards bodies. As this is an Internet Engineering Task Force Working Group, the natural focus is on current and future network management protocols and models used in the Internet. However, the interface being defined is expected to be sufficiently flexible and extensible to allow support for other protocols and other classes of manageable objects. The anticipated list of protocols includes Simple Network Management Protocol (SNMP), OSI Common Management Information Protocol (CMIP), CMIP Over Tcp (CMOT), Manufacturing Automation Protocol and Technical Office Protocol CMIP (MAP/TOP CMIP) and Remote Procedure Call (RPC).

Specific Objectives:

1. Determine the feasibility of a common interface across multiple management protocols.
2. Define the requirements for such an interface.
3. Define an architectural framework for such a service interface.
4. Define a specification that satisfies the architectural requirements.
5. Implement one or more prototypes of the interface.
6. Advance an RFC based on the specification and prototype experience.

Estimated Timeframe for Completion:

8 - 12 Months

CURRENT MEETING REPORT

First meeting is planned for the February IETF.

2.5.4 LAN Manager Working Group (lanman)

CHARTER

Chairperson: Jim Gruel/HP, jimg@hpcndpc.cnd.hp.com

Mailing List: lanmanwg@spam.istc.sri.com

Description of Working Group:

To define and maintain the MIB and relevant related mechanisms needed to allow management overlap between the workgroup environment (LAN Manager based) and the enterprise environment (based on TCP/IP management).

Specific Objectives:

This translates into three basic objectives:

- Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
- Propose extensions to the TCP/SMI when appropriate.
- Develop requirements for additional network management information, as needed, and work to extend the LAN Manager interfaces to support such information.

Estimated Timeframe for Completion:

Objective 1: Version 1 of the LANMAN MIB has been completed and is awaiting consideration by the RFC editor (two RFCs have been proposed: LANMAN-MIB for "conventional" objects, and LANMAN-MIB-EXPER for objects related to LAN Manager alert handling). Subsequent versions will be worked on as necessary after further experience is gained with version 1. There is no definite timeframe set for work on version 2.

Objective 2: No extensions to the SMI have been proposed, and there are no immediate plans for making such a proposal.

Objective 3: No modifications to the LAN Manager interfaces were required for version 1 of the LANMAN MIB. This issue will be reconsidered after further experience is gained with version 1.

CURRENT MEETING REPORT

Did not meet

2.5.5 NOC-Tools Working Group (noctools)

CHARTER

Chairpersons: Robert Enger/Contel enger@sccgate.scc.com
Robert Stine/Sparta stine@sparta.com

Mailing List: noctools@merit.edu

Description of Working Group:

The NOC-Tools Working Group will develop a catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

Specific Objectives:

1. Identify tools available to assist network managers in debugging and maintaining their networks.
2. Publish a reference document listing what tools are available, what they do, and where they can be obtained.
3. Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
4. Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.
5. Identify the need for new or improved tools as may become apparent during the compilation of the reference document.

Estimated Timeframe for Completion:

The first edition of the catalog will be submitted for final review at the October-November IETF meeting. Preliminary versions will be made available earlier.

CURRENT MEETING REPORT

Reported by Robert Enger/Contel

MINUTES

The NOCtools working group session was held on Wednesday morning, November 1st, during the recent Hawaii IETF meeting. In addition, a joint User-Doc/NOCtools meeting was held that afternoon. Attendance was larger than noted above; a complete attendee list was not available at the time of this writing. A notable absentee was co-chair Bob Stine of Sparta who was not able to attend this meeting.

Morning discussions included catalog re-organization and index design. It was suggested that the catalog be restructured, placing the tutorial as an appendix at the end, and adding a table indicating which keywords apply to which tools. Considerable controversy arose over catalog entries which contained multiple commands. The question was "how does one find a command by name, if the the entry containing multiple command names can appear only once in the (alphabetically organized) catalog". The group is reluctant to introduce a page number based index, because of the associated difficulty in making catalog updates. The members agreed to suggest to Bob Stine (the "book boss") that the two entries containing multiple command names be broken down into multiple entries, each containing one command (similar to the rest of the catalog entries). Bob has rejected this suggestion because he feels it will cause too much expansion and redundancy in the catalog.

Concurrent with the index design (table concept) was a discussion of the physical limitations of the table. Many of the problems result from the desire to make the document usable in a manual, off-line mode. To this end, it was decided that the table's ultimate horizontal size should be limited to two pages, so that an entire line can be viewed simultaneously on two adjoining pages (eg left and right side of a book). This in turn limits the number of keywords that can be listed in the table. One suggestion from the attendees was to anticipate future space problems by eliminating the "environment" (target hardware/software platform) keywords from the table to conserve space. Since there are currently no space problems it was decided to list the entire set of keywords across the top of the table.

During the morning meeting we also made numerous corrections to the text, as well as adding text to improve clarity and ease of use. We also acquired a new catalog entry: HyperMib, a HyperCard based tool allowing one to inspect the text of the MIB specification documentation, as well as a few new keywords.

Bob Stine reminds me to point out that, excepting the multiple-command entry suggestion, all other suggestions should now be reflected in the current draft. So, please review the current draft, and feel free to make additional suggestions or corrections.

After lunch, the joint UserDoc/NOCtools session was held. Discussion ranged widely from specific suggestions for product improvement to general questions of publicity, technical assistance, and distribution. We were joined by Dave Crocker, NOCtools' area director, and received cameo appearances from other luminaries too.

Specific suggestions for NOCtool catalog improvement included:

- Continuing the working group beyond publication of the first draft. This would allow the group to:
  - handle the expected volume of "me too" submissions
  - formulate written policy for document up-keep
  - locate an entity capable of assuming the update chores.
- It is expected that a second edition of the catalog will be published when the group disbands, probably around June.
- Tighten catalog entry format specifications, so that future entries will be submitted in near-perfect form, reducing the work load on the entity assuming update chores.
- Add some means of determining the "freshness" of a catalog entry, and consider whether entries should be removed. Suggestions:
  - Shelf life/expiration date
  - Date of initial insertion
  - Date of last update.
- The suggestion which received the most support was the one recommending the addition of a "last update" date.
- Before public announcement of the first edition, it was recommended that notification be sent to the IETF mailing list. This would provide entry-suppliers with a last opportunity to inspect the document before it goes public.

Gary Malkin volunteered to write the statement of work detailing the duties of the entity that assumes catalog update responsibility.

General discussions of import to NOCtools concerned the idea of trying to reduce the "administrative" or "procedural" load on working groups that produce documents. One suggestion was to elicit the assistance of professionals in the field, librarians. It was observed that library science is interested in learning how to utilize modern technology, and that perhaps the IETF (UserDoc?) could form a joint-research relationship with one of the schools. This would provide the IETF with valuable technical assistance in the area of document preparation, layout, etc, as well as professional assistance in the area of on-going document up-keep. It was agreed by most of the attendees that the IETF will be faced with more and more "living documents" (those requiring periodic update) as time goes on. Alternatively, it was suggested that the IAB/IETF approach the funding agencies with a request for money to pay for a full time document update and distribution service.

On the subject of publicity a number of attendees suggested other groups that should be contacted and notified of the existence of the documents. Conversation then went on to include suggestions for closer working relationships with parallel organizations in the Bitnet and uucp worlds. It is believed that each group has much to offer the others.

A mailing list, noctools@merit.edu, has been established for the working group. As usual, requests to join the list should be directed to noctools-request@merit.edu.

ATTENDEES

Karen Bowers
Robert Enger
Steven Hunter
Gary Malkin
Keith McCloghrie
Karen Roubicek
Mary Stahl

2.5.6 SNMP Working Group (snmp)

<u>CHARTER</u>

Chairperson: Marshall T. Rose/NYSERNet, mrose@nisc.nyser.net

Mailing List: snmp-wg@nisc.nyser.net
Description of Working Group:

> The SNMP Working Group has the goal of producing necessary SNMP centric RFCs especially in the area of the Management Information Base (MIB) and the Structure of Management Information (SMI) to provide for both critical operational management requirements and cooperative experimental work.

Specific Objectives:

> Provide a draft RFC for an enhanced backwardly compatible MIB in 4Q89 which can be implemented and interoperability tested by 1Q90 to address critical operational requirements. After multi-vendor testing, draft will be submitted to the RFC Editor for standardization.

Milestones

		SCHEDULED	ACTUAL
GOAL	Prepare MIB-II draft		
o TASK	- Initial meeting to assign actions	89-08-18	89-08-18
o TASK	- Actions due	89-09-01	89-09-08
o TASK	- Edit draft	89-09-15	89-09-22
o TASK	- QC draft and release	89-09-22	89-10-29
GOAL	Examine and tentatively agree		
o TASK	- Discussion meeting to review draft	89-10-16	89-10-16
o TASK	- Edit drafts and release		
	- MIB-II draft	89-10-20	
	- Ethernet-like draft	89-10-20	
	- T1-carrier draft	89-10-20	
	- Token-ring draft	89-10-31	
	- other drafts	TBD	
GOAL	Implement and report back		
o TASK	- Incremental editing of drafts	throughout	
o TASK	- 90 percent implimentation of relevant portions - along with interoperability testing	89-12-01	
GOAL	Evaluate and possibly iterate		
o TASK	- Determine if concensus is reached	89-12-01	
o TASK	- Final edit of drafts	89-12-08	
o TASK	- Submit drafts for standardization		
	- MIB-II draft	89-12-08	
	- Ethernet-like draft	89-12-08	
	- T1-carrier draft	89-12-08	
	- Token-ring and other drafts	N/A	

2.6 OSI Interoperability Area

Directors: Ross Callon/DEC and Robert Hagens/Univeristy of Wisconsin

The OSI Area has expanded rapidly. The initial OSI working group (OSI-IWG) has been converted into a general OSI WG (OSI-General). In order to meet the challange of operating OSI in a dual environment, the following new working groups have formed, or will be forming shortly.

List of Working Groups

- Name: OSI-General
- Chair: Callon and Hagens
- Scope:
  - Forum for OSI-related issues not covered by an existing WG
  - Initial starting point of any OSI issue

- Name: OSI-X.400
- Chair: Hagens
- Scope:
  - 822/X.400 gateway issues (including RFC 987 and successors)
  - Follow work of NIST X.400 groups

- Name: OSI-X.500
- Chair: Deutsch (tentative chair)
- Scope:
  - X.500 and DNS interactions
  - Evaluation of any missing pieces in X.500
  - Naming service requirements in a dual environment
  - Follow work of NIST X.500 groups

- Name: OSI-NSAP-ADMIN
- Chair: TBD
- Scope:
  - Produce NSAP administration guidelines

- Name: OSI-RA (Registration Authority)
- Chair: TBD
- Scope:
  - Produce X.400/X.500 name registration guidelines
  - Follow work of NIST and ANSI registration groups

- Name: OSI-MIB
- Chair: TBD
- Scope:

- Definition of MIB variables for dual protocol hosts

Readers interested in a summary of the work of the OSIIWG should consult the Current Meeting Report for the OSI-General working group.

OSI AREA

- OSI Interoperability Working Grp.
  - ECHO for 8473 (OSI IP)
  - NSAP ADDRESSING
  - CONGESTION AVOIDANCE
  - STATUS
    (ACCOMPLISHMENTS, HOLES, W.G.'s)
- OSI - X.400
  - X.400 operational issues (NIST liaison)
  - 822/X.400 GWY issues
  - MET WITH DOMAIN NAME GROUP
  - ADDRESS MAPPING
    - STATIC TABLES
    - USE DNS
    - (USE X.500 F.F.S.)

- OSI - X.500
  - NEW W.G.
  - "MISSING PIECES" IN X.500
  - INTERACTION WITH D.N.S.
  - SUPPORT OSI IN DUAL ENVIR. ("INTEROPERABILITY")

- OSI - RA
  - NSAP ADMIN. GUIDELINES
  - X.400/X.500 NAME REG. "

- OSI - MIB
  - ??

2.6.1 OSI General Working Group (osigen)

CHARTER

Chairpersons: Ross Callon/DEC callon@erlang.dec.com
Rob Hagens/UWisc. hagens@cs.wisc.edu

Mailing Lists:

ietf-osi@cs.wisc.edu
ietf-osi-request@cs.wisc.edu

Description of Working Group:

Help facilitate the incorporation of the OSI protocol suite into the Internet, to operate in parallel with the TCP/IP protocol suite. Facilitate the co-existence and interoperability of the TCP/IP and OSI protocol suites.

Specific Objectives:

The following are specific short-term goals and objectives for the OSI WG. Other mid-term objectives have also been identified and are available from the chairs.

- Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet. Coordinate addressing format with GOSIP version 2 and possibly other groups.
- Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley folks.
- Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.
- What routing should be used short term for (i) intra-domain routing; and (ii) inter-domain routing?
- For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Are there outstanding issues remaining here?
- Review short term issues involved in adding OSI gateways to the Internet. Preferably, this should allow OSI and/or dual gateways to be present by the time that Berkeley release 4.4 comes out.

Estimated Timeframe for Completion:

Indefinite

CURRENT MEETING REPORT

Reported by Robert Hagens/University of Wisconsin

AGENDA

- General Meeting
- Updates
  - BSD 4.4
  - New Revision RFC 1069
  - Echo RFC
  - GOSIP Comments
- OSI at Interop 89
- Results of the MITRE congestion avoidance experiments
- State of the OSIIWG – accomplishments and future work

MINUTES

The meeting was convened by co-chairmen Ross Callon and Rob Hagens. An attendance list will be published with the Proceedings of the IETF.

A series of brief status updates on the following topics were presented:

- BSD 4.4: An ISODE/BSD interface has been constructed and tested. Alpha copies have been distributed to a small number of sites. Work is still in progress fixing bugs, testing, etc.
- New revision of RFC 1069. The newest version of RFC 1069, compatible with the GOSIP V2 (if the OSIIWG comments are accepted) has been prepared. Its submission to the RFC editor will be delayed until GOSIP V2 is released.
- The ECHO RFC has been released as an Internet Draft. This RFC specifies how to implement an ECHO facility with ISO 8473. The WG reviewed the document and found (with 2 minor editing changes) it ready to be sent to the RFC editor.
- There is no official word from NIST regarding the OSIIWG GOSIP V2 comments. A representitive of the OSIIWG will attend the next GOSIP Advanced Requirements Committee meeting.
- GSA has a contract to administer ICD 0005 (although NIST still maintains authority). The DCA use of 0006 is unknown. NIST currently supports the use of 0005 by the entire Internet. Policies for the use of 0005 have not yet been established. Those with strong interests in future policy should contact:

  Mr. Gerard F. Mulvenna
  Technology Building, Room B-217
  National Institute of Standards and Technology
  Gaithersburg, MD 20899

Dave Katz presented his OSI experiences at Interop, 89.

Rick Wilder presented preliminary results of the MITRE congestion avoidance experiments.

Following this, the state of the OSIIWG was discussed. A list of new working groups that need to be formed was presented. This list includes the reorganization of the OSIIWG into the OSI-General WG.

Note: the OSI-RA group may be split into two separate groups, one to produce NSAP administration guidelines, and the other to follow upper layer registration policy.

Finally, the list of current and future work of the OSI Area was presented:

IETF OSIIWG STATUS/Callon and Hagens

Agreements and future work of the IETF OSIIWG

DRAFT

1. Physical Layer
   (a) Accomplishments and Agreements
       - None identified.

   (b) Future Work
       - None identified.
2. Link Layer
   (a) Accomplishments and Agreements
       - None identified.

   (b) Future Work
       - Distinguishing packets on the wire
       - HDLC
       - X.25
3. Network Layer
   (a) Accomplishments and Agreements
       i. Data transfer
          - ISO 8473/use as specified

       ii. Routing
          - ISO 9542/use as specified
          - Intra-domain routing/use ANSI IS-IS as presented as draft proposal
            use ANSI IS-IS as presented as draft proposal.
          - Inter-domain routing use static tables.

       iii. ISO 8473 Echo
            A draft RFC has been prepared. It describes an echo function that is realized by defining a new network selector that indicates an echo entity. This is backward compatable with existing 8473 packets.

iv. NSAP address format
   - RFC 1069 RFC 1069 has been updated to align with the GOSIP V2 NSAP address format.
   - NSAP Selectors OSIIWG comments on GOSIP V2 recommend that GOSIP V2 should not specify the format of the NSAP selector value.

(b) Future Work

i. ISO 8473 Echo

Initiate a new ANSI X3S3.3 work item to propose a CLNP echo function to ISO. This echo function is realized by defining a new protocol type field. This is not backward compatable with existing 8473 packets.

ii. NSAP address format
   - NSAP Administration Design and write procedures for administering NSAP address heirarchies.
   - ICD Usage Determine whether the Internet should register under ICD 0005 or ICD 0006 or both. Coordinate with any previous NIST/GSA agreements, or motivate new agreements.

iii. CO/CL

We should track the CO/CL interworking status in X3S3.3.

4. Transport Layer

(a) Accomplishments and Agreements

Recommend that GOSIP V2 mandate NIST agreements regarding congestion recovery algorithms and related retransmission timer algorithms.

(b) Future Work

None identified.

5. Session Layer

(a) Accomplishments and Agreements

None identified.

(b) Future Work

None identified.

6. Presentation Layer

(a) Accomplishments and Agreements

None identified.

(b) Future Work

None identified.

7. Application Layer

(a) X.400

i. Accomplishments and Agreements
   - PRMD name

The intended use of "NREN" as a PRMD name is to identify a management domain within which every registered Internet entity has a default X.400 Address. This address would be based upon the Internet domain name. We expect some or all currently registered entities to decide for them- selves whether they wish to use the default or register another name in another way. This default provides a useful and helpful option without constraining any individual entity to keep what the default provides for them.

ii. Future Work

A. GOSIP V2

Work with the GOSIP user's group to rewrite the X.400 ORAddress section.

B. 822 <-> X.400 gateway operation

- Table Maintenance
- Locating a Gateway
- ORAddress Structure

C. X.400 operation

- Default naming
- Taxonomy of issues Write a memo which describes the needs of X.400 addressing, X.400/RFC 822 address mapping, and utilization of an X.500 directory service. (In Progress).

(b) Registration and Naming

i. Accomplishments and Agreements:
See "NREN".

ii. Future Work

- NSAP administration See NSAP administration under Network Layer.
- NSAP and ORAddress relationships Explore the relationship between NSAP addresses and X.400 ORAddresses. Should the NSAP address field "oganization" under ICD 0005 be used in the X.400 ORAddress "organization" field to reduce administration complexity?
- Establishing Ownership Identify necessary steps we must take to assert that the name "NREN" belongs to the FRICC.

(c) Directory Services

i. Accomplishments and Agreements

None.

ii. Future Work

A. X.500 and Internet DNS

Explore coexistence/interactions between X.500 and the Internet DNS

B. Missing Pieces

Locate missing pieces required by a production system (format of objects, choice of dis- tinguished names, etc.)

C. Requirements of a dual protocol internet

- Application Gateways Identification of application gateways needed for communication between heterogenous, pure stack hosts. In addition, support for the deci- sion to gateway (i.e., forward as X.400 message or translate into RFC 822).
- Stack Choice Identification of optimal protocol stack choice for dual hosts (based upon the destination system).

(d) VTP

i. Accomplishments and Agreements

None

ii. Future Work

Look for problems with Telnet/VTP interaction.

(e) FTAM

i. Accomplishments and Agreements

None

ii. Future Work

Look for problems with FTAM/FTP interaction.

(f) Network Management

i. Accomplishments and Agreements

None

ii. Future Work

- CMIP
- OSI MIB

8. General Future Work

(a) Mixed Stack

GOSIP prohibits a mixed stack approach. Do mixed stacks have enought merit that they should be allowed?

(b) Mixed Technology

Can OSI problems be solved with internet technol- ogy? Will the Internet incorporate OSI technology? For example, can X.400 routing utilize the DNS, in the absense of X.500?

(c) Document Review
  - GOSIP
  - ANSI specifications
  - FRICC Multi-Protocol Implementation Plan

ATTENDEES

Almquist, Philip
Boivie, Rick
Callon, Ross
Cargille, Allan
Carter, Glen
Chinoy, Bilal
Colella, Richard
Coltun, Rob
Demar, Phil
Forster, Jim
Fox, Richard
Galvin, James M.
Gerlach, Chuck
Gross, Martin
Hagens, Rob
Joshi, Satish
Katz, Dave
Kerby, Kathy
Lazear, Walt
Miller, Dave
Nadler, Dennis
Norton, Bill
Oattes, Lee
Ramakrishnan, K.K.
Reilly, Michael
Roselinsky, Milt
Sheridan, Jim
Steinberg, Louis
Su, Zaw-Sing
Wilder, Rick
Wintringham, Dan
Youssef, Mary

2.6.2 OSI X.400 Working Group (osix400)

CHARTER

Chairperson: Robert Hagens/UWisc., hagens@cs.wisc.edu

Mailing Lists:

ietf-osi-x.400@cs.wisc.edu
ietf-osi-x.400-request@cs.wisc.edu

Description of Working Group:

The IETF OSI X.400 working group is chartered to identify and provide solutions for problems encountered when operating X.400 in a dual protocol internet. This charter includes pure X.400 operational issues as well as X.400 <-> RFC 822 gateway (ala RFC 987) issues.

Specific Objectives:

1. Develop a memo describing known issues and problems.
2. Develop a scheme to alleviate the need for static RFC 987 mapping tables.
3. Develop a scheme to support X.400 routing.
4. Consider ways in which directory services may be utilized in order to hide the details of RFC 822 and X.400 addressing.

Estimated Timeframe for Completion:

This memo is a working document. A first draft was discussed at the October 31, 1989 meeting.

CURRENT MEETING REPORT

Reported by Robert Hagens/Univeristy of Wisconsin

AGENDA

- Announcement of new name
- Status of the quest for "NREN"
- Review of Scope
  - 822 <-> X.400 gateway issues (RFC 987 and successors)
  - X.400 operational issues in a dual protocol internet
- Review of Issues

  - 822 <-> X.400 Gateways
    * RFC 987 gateway background
    * Table Maintenance
    * Locating a Gateway
    * ORAddress Structure
  - X.400 Operation
    * Routing to destination or 822 gateway
    * Use of Internet Technology
    * Mixed Stacks
    * MTA names
    * Use of "NREN"

  Presentation of a new, unified address structure
- Enumerating and discussion of major tasks

MINUTES

The meeting was convened by chairman Rob Hagens. An attendance list will be published with the Proceedings of the IETF. The Domain Name WG meet jointly with the OSI-X400 WG during the afternoon.

WORKING GROUP SCOPE

The scope of the WG was presented:

- RFC 822 <-> X.400 Gateway Issues
  - maintenance of RFC 987 mapping tables
  - routing toward a gateway
- X.400 Operational Issues
  - Structure of OR-Addresses in the Internet
  - X.400 Routing
  - Nameservers

The group determined that (with the exception of determining the structure of OR-Addresses in the Internet), they should not try to solve "pure-OSI" problems. These problems fall into the domain of other OSI groups. The WG should develop and maintain a close relationship with such groups:

- NIST X.400 SIG
- NIST X.500 SIG
- GOSIP X.400 committee

PRESENTATION OF ISSUES

Rob Hagens presented a list of issues facing the WG. That list is included here:

Issues, Problems, and Proposed Solutions to
X.400 and 987 Gatewaying in the Internet

1. X.400-RFC 822 Gateway Issues
   (a) Background

This background information serves as a very brief tutorial on RFC 987. The information presented below is far from complete. It is strongly recom- mended that anyone interested in the issues dis- cussed below should obtain and read RFC 987.
RFC 987 specifies how messages should be gatewayed between RFC 822 based systems and X.400 (1984) based systems. Although the RFC describes the translation of various protocol elements from one system to the other, the following discussion is limited only to the translation of addresses.
RFC 987 specifies that translation from one address space to another may occur in 2 ways. The normal method of translation (table lookup) is used when sub-trees of the different name spaces are associ- ated via mapping tables. The fall back method of translation (encoding in the other address space's format) is used when table lookup fails.
Table lookup is accomplished through the use of 2 separate tables: an RFC 822 -> X.400 table, and an X.400 -> RFC 822 table. Each entry in the tables is indexed by a key. The address to be mapped is compared against each key in the table. The com- parison that matches the most components is selected (i.e., the "longest" match). The value associated with the key is a template that is used to construct the translated address.

i. Table Driven Mapping

For example, the 822 domain "merit.edu" could be associated with the OR Address space "C=US, PRMD=NREN, O=MERIT.EDU" in a mapping table. Thus, when translating the 822 address "hwb@merit.edu", the domain specification "merit.edu" would be compared against the various keys in the table. Assuming that the table con- tains two keys "edu" and "merit.edu", the longest match "merit.edu" would be selected. The template associated with the key "C=US, PRMD=NREN, O=MERIT.EDU", would be used to produce the address "C=US, PRMD=NREN, O=MERIT.EDU, OU=CS, PN=HWB". In this example, the translation of the last domain component "cs" is performed systemat- ically. The translation of the right-hand side of the 822 address "hwb" is specified by RFC 987.
This example shows that a single entry can specify the translation for all addresses in the "merit.edu" domain. This entry associates the 822 domain "merit.edu" with the X.400 namespace under "C=US, PRMD=NREN, O=MERIT.EDU".
An analogous scheme is used for the opposite direction.

ii. Mapping Without Tables

If a mapping table entry is not present, transla- tion may still occur. However, in this case, the translation is less sophisticated. Translation, in this case amounts to encoding the address in the other system's format. RFC 987 specifies default rules that may be used to perform this encoding. These rules specify the manner in which an RFC 822 address may be encoded in X.400, and vice versa. The following examples consider each direction separately:

iii. RFC 822->X.400

In this direction, the domain-defined attribute "RFC-822" may be used to encode an RFC 822 address. For example, if an 822 address "hagens@janeb.cs.wisc.edu" was translated by a gateway that had an X.400 address "C=US, PRMD=NREN, O=MERIT.EDU", then that gateway (in the absence of a mapping table entry) would produce the address 'C=US, PRMD=NREN, O=MERIT.EDU, DD.RFC- 822="hagens@janeb.cs.wisc.edu"'.

iv. X.400->RFC 822

In this direction, left-hand side encoding may be used to encode an X.400 address within 822. For example, the X.400 address "C=FR, ADMD=FRENCH-PTT, O=INRIA, PN=HUITEMA", when considered by a gateway with the 822 address "merit.edu", would be translated to '"C=FR, ADMD=FRENCH-PTT, O=INRIA, PN=HUITEMA"@merit.edu'.

(b) Issues

i. Table Maintenance

The mapping table entries must be kept consistent among all the 987 gateways in the world. This is very difficult to accomplish by hand. How can the table maintenance task be automated?

ii. Finding the Gateway

How does a mail router find a 987 gateway? In the X.400-¿RFC 822 direction, it is the responsi- bility of X.400 routing. Note: X.400 routing is not defined by any standard. In the RFC 822- ¿X.400 direction, it is the responsibility of 822 routing. Conventional MX records could be util- ized to solve the problem.

iii. Structure of X.400 addresses

It is desirable to provide a default X.400 address for hosts within the Internet. This address will be structured so that the X.400 address space corresponds with the domain namespace. What is the best structure to use for this purpose?

The choice of format of X.400 addresses, and the corresponence of these addresses to 822 domains will determine the contents of the of 987 mapping table entries.

- Proposed Solution

  The currently proposed solution is to map the top and second level domains to the ORAddress "organization" attribute. Subsequent lower level domains will be mapped to a sequence of "organization unit" attributes. For example, "venera.isi.edu" would map to "O=isi.edu, OU=venera".
- Use of 'NREN' as a PRMD name

  The intended use of "NREN" as a PRMD name is to identify a management domain within which every registered Internet entity has a default X.400 Address. This address would be based upon the Internet domain name. We expect some or all currently registered entities to decide for themselves whether they wish to use the default or register another name in another way. This default provides a useful and helpful option without constraining any individual entity to keep what the default provides for them. Is it necessary to define a second PRMD name which would identify a management domain within the NREN that utilizes X.400 addresses that are not based upon Internet domain names? If this is true, is the original use of "NREN" incorrect?
  We need to show "ownership" of the name "NREN" so that other groups do not have the right to register it. Trademarking is the first step. Other uses of "NREN" should be looked into. Any way that we can show "use" of the name will help establish our "ownership".

2. X.500 Operation Issues
   (a) Issues
       i. Distinguished Names

          Who will determine the structure of X.500 dis- tinguished names (and the objects they locate) for use within the Internet community?
       ii. DNS coexistence

          How should the DNS and X.500 coexist?
       iii. Domain Distinguished Names

          Is it acceptable, for transition purposes only, to suggest that Domain names be used as Dis- tinguished names?

DISCUSSION OF ISSUES

Non-USA Internet Sites

The default OR Address may not be acceptable for Internet sites that are not within the USA. 1) The WG cannot mandate the format of addresses within a foreign country. 2) the NREN is a national object. Are these reasons sufficient to prevent the definition of a default name using NREN? At least, it should be made clear that the default name is valid for USA-Internet sites only. This may not be inappropriate if many foreign countries have already defined the X.400 registration policy that would affect the foreign Internet sites.

"NREN"

The name "NREN" was originally choosen to be a PRMD name. The purpose of this PRMD was to contain OR Addresses based upon Domain Names. It was suggested that perhaps "NREN" is not appropriate for this use. No other name was decided upon. Possible candidates are names that convey some concept of Domain Names, such as "DN". This change would allow the name "NREN" to be used by a FRICC-run PRMD.

Another option for a PRMD name would be to use the numeric form.

The effort to pre-register "NREN" as an ANSI OSI Organization name failed. It is not clear that the OSI X.400 WG should attempt to register the name until its exact use has been determined.

It was suggested that the WG should consider producing a specification for written OR Addresses.

PRESENTATION OF A NEW, UNIFIED ADDRESS FORMAT

Paul Mockapetris presented his ideas regarding a new style of address. He would like to see the world move forward with the development of a unified, simple address structure. His proposal is a format that has RFC 822 compatible syntax, whose semantic value is that of an X.500 distinguished name. These new addresses would be very short and user-friendly. The new addresses could be used to look up both X.400 ORAddresses as well as conventional 822 addresses. The look up mechanism could utilize the DNS as well as X.500.

GATEWAY SCENERIOS

A discussion of RFC 822 - X.400 gateway (987) scenarios produced the following questions:

- Will any 987 gateway provide connectivity to every X.400 MTA?
  The answer to this question will determine whether an 822 transfer agent must choose a specific 987 gateway based upon the destination

address, or if the closest, default 987 gateway will always suffice.
- Is there really benefit to table driven mappings or is it sufficient to simply use default encodings?

A scheme that utilizes the DNS to aid a 987 gateway was discussed. The scheme requires the following components:

- An ASCII (canonical) representation of ORAddresses.
- A new tree of the DNS that is based upon canonical ORAddresses strings (called ORADDR). This tree is populated with MX records (that store the SMTP 822 address of 987 gateways), and TO-SMTP RRs.
- Two new DNS resource records. TO-SMTP RRs are stored in the ORADDR tree. They contain the information necessary to translate an X.400 address into an 822 address. TO-X400 RRs are stored in the existing DN tree. They contain information necessary to translate SMTP 822 addresses into X.400 addresses. A distributed collection of TO-SMTP and TO-X400 records correspond to the 987 mapping tables X.400 to RFC 822 (mapping 1) and RFC 822 to X.400 (mapping 2), respectively.

A sample scenario would be:

822->X.400

Case A — The destination address is an SMTP address which has been previously associated with an ORAddress. This means that there is a TO-X400 RR that describes how to translate the SMTP 822 address into an ORAddress. The originating transfer agent will look up the destination address and receive an MX record and a TO-X400 RR. The MX record identifies a 987 gateway and is used to transfer the message to that gateway. The TO-X400 record is ignored by the originator.

When the 987 gateway receives the message, it will lookup the destination address and receive an MX and TO-X400 RR. The MX record is ignored, but the TO-X400 RR is used to translate the destination address into an ORAddress.

Case B — The destination address is an ORAddress. The originating transfer agent will look up the destination ORAddress in the ORADDR tree and receive an MX record. The MX record identifies a 987 gateway and is used to transfer the message to that gateway. The destination address sent in the SMTP envelope will contain "ORAddress"@gateway.

X.400->822

Routing to the 987 gateway is not within the scope of the WG; it is assumed that the message has already reached the 987 gateway.

Case A — The destination address is an ORAddress which has been previously associated with an SMTP 822 address (sub)tree. This means that there is a TO-SMTP RR that describes how to translate the ORAddress into an SMTP 822 address.

When the 987 gateway receives the message, it will lookup the destination address in the ORADDR tree and receive a TO-SMTP RR. The TO-SMTP RR is used to translate the destination address into an SMTP 822 address.

Case B — The destination address is an 822 address which has been encoded in an ORAddress.

When the 987 gateway receives the message, it will translate the destination address into an 822 address using the default encoding rules.

ATTENDEES

Brackenridge, Billy	Lottor, Mark
Callon, Ross	Love, Paul
Cargille, Allan	Mokapetris, Paul
Carter, Glen	Nadler, Dennis
Colella, Richard	Nitzan, Rebecca
Demar, Phil	Pomes, Paul
Galvin, James M.	Reschly, Robert J.
Gross, Martin	Roselinsky, Milt
Hagens, Rob	Rust, Bill
Hain, Tony	Schoch, Steven
Lazear, Walter D.	Sheridan, Jim
Long, Dan	St. Johns, Mike
	Wilder, Rick

2.7 Security Area

Interim Director: Phillip Gross/NRI

This is an incredibly important area that demands immediate attention. THerefore, I am very pleased to announce that Steve Crocker (TIS) has joined the IESG as the new Security Area Director. Among his other projects at TIS, Steve is involved with developing a secure email system based on RFCs 1113-1115.
There is currently only one working group in this area:

- IP Authentication (Schiller, MIT)

However, this WG has essentially completed its objective of developing an IP Authentication option, and has moved on to developing a method for SNMP Authentication. Therefore, in the interest of keeping those objectives distinct, this WG may be split into two - IP Authentication and SNMP Authentication. The goal would be to conclude the IP Authentication portion expeditiously, so full attention can be given to other matters.

We have identified the need for at least one additional near-term WG - the Secure Configuration WG. The goal will be to draft a short RFC documenting the proper ways to configure a new system to minimize the known windows for attack (eg, turn off STMP debug, etc). We have tentative agreement from the CERT to join us in this WG.
Steve Crocker will be reporting in this spot in the future.

2.7.1 IP Authentication Working Group (ipauth)

CHARTER

Chairperson: Jeffrey Schiller/MIT, jis@bitsy.mit.edu

Mailing List: awg@bitsy.mit.edu

Description of Working Group:

To brainstorm issues relating to providing for the security and integrity of information on the Internet, with emphasis on those protocols used to operate and control the network. To propose open standard solutions to problems in network authentication.

Specific Objectives:

1. RFC specifying an authentication format which supports multiple authentication systems.
2. Document discussing the cost/benefit tradeoffs of various generic approaches to solving the authentication problem in the Internet context.
3. Document to act as a protocol designers guide to authentication.
4. RFC proposing A Key Distribution System (emphasis on "A" as opposed to "THE"). MIT's Kerberos seems the most likely candidate here.

Estimated Timeframe for Completion:

This working group will hopefully complete its current objectives within one year. At this point the group will either disband or will move on to other related problems/issues.

CURRENT MEETING REPORT

Did not meet.

2.7.2 SNMP Authentication Working Group (snmpauth)

CHARTER

Chairperson: Jeffrey Schiller/MIT, jis@bitsy.mit.edu

Mailing List:

awg@bitsy.mit.edu
awg-request@bitsey.mit.edu

Description of Working Group:

To define a standard mechanism for authentication within the SNMP.

Specific Objective:

To write an RFC specifying procedures and formats for providing standardized authentication within the SNMP.

Estimated Timeframe for Completion:

By January 1, 1990.

CURRENT MEETING REPORT

Did not meet

2.8 Operations Area

Interim Director: Phillip Gross/NRI

I had not included an Operations Area in my original plans for an IETF steering group. Instead, I had included the important topic of User Services. In our early IESG discussions, we realized that there were several broad topics that generated what amounted to long-term standing WGs, and these often were operations-oriented topics. Therefore, our original thinking was to broaden out the User Services Area to be an Operations Area, which would include network operations (e.g., JOMANN), network information services (e.g., user services working group), and network connectivity planning (e.g., Topology planning and routing coordination).
/vspace.1in However, this plan did not come to quick fruition because we did not have a director for the Operations area. The IETF has a strong committment to user services. It was because of this strong personal committment that I asked Karen Bowers to form the User Services Working Group (USWG). The USWG had become a very active group, so rather than allow those efforts to languish in area without a director, we decided to move those efforts under Craig Partridge's Host Services Area. After a bit of a rocky handoff, this is now complete. Craig's newly renamed area is the Host and User Services Area.
To avoid this type of confusion in the near future, I will served as the interim director until a permanent director is identified. At this time, however, we feel it is better to leave most user services activities under Craig, rather than attempting to eventually move them back under the Operations area. Craig and I will simply need to carefully coordinate any activities that appear to overlap.

Currently, there are 2 active WGs in the Operations area:

- JoMANN (Hares, Merit)
- Benchmarking Methodology (Bradner, Harvard)

JoMANN is a long standing and productive WG. There have been some suggestions for a minor revision in its scope and format. For example, we may want to give this a more explicit FARNET spin. However, in any modification to its format, we will seek to keep it as active and useful as it has been in the past.

Benchmarking methodology is a new WG with the goal of developing standard methods for measuring performance in, for example, routers and bridges. It has already meet at least once, and will meet at the upcoming IETF.

We have identified the need for at least one more near-term WG - a TCP/IP installation guide. A prospective chair has been identified and we hope to hold the initial meeting shortly after the February IETF at FSU.

2.8.1 JOMANN Working Group (jomann)

CHARTER

Chairperson: Susan Hares/Merit, skh@merit.edu

Mailing Lists:

njm@merit.edu (Regional or National Net NOC people)
njm-interest@merit.edu (anyone interested)
njm-request@merit.edu

Description of Working Group:

The "Joint Monitoring Access for Adjacent Networks focusing on the NSFNET Community" working group is a continuing forum for the facilitation of common solution to operational problems in the NSFNET regional networks.

Specific Objectives:

The JOMANN Working Group will:

- discuss how to identify problems in the next hop network
- create a list of existing tools which can solve these problems (We will discuss to see if NOC-Tools Working Group can take over this. NSFNET will archive a list of these tools.)
- create a list of routing topology maps of regionals (possibly prepare a MAP Internet-Draft)

Estimated Timeframe for Completion:

Indefinite

CURRENT MEETING REPORT

Did not meet

2.8.2 Benchmarking Methodology Working Group (bmwg)

CHARTER

Chairpersons: Scott Bradner/Harvard, sob@harvard.harvard.edu
Mick Scully, mcs@ub.com

Mailing List: bmwg@harvisr.harvard.edu

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with standalone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for standalone devices has matured sufficiently, the group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Specific Objectives:

1. Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.
2. The document will also define various classes of standalone network devices, such as repeaters, bridges, routers, and LAN wiring concentrators, as well as detail the relative importance of various performance criteria within each class.
3. Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

In addition, this document will make specific recommendations on a common reporting structure for benchmark results.

The document will be organized such that each section::
- (a) Defines a device class.
- (b) Defines the performance characteristics important to this class of device.
- (c) Recommend a specific benchmark suite (FLINTSTONES) for this class of device.
- (d) Define a common reporting format for the results of the benchmark suite.

Estimated Timeframe for Completion:

We plan to issue a draft document for Objective No. 1 by late December 1989. A document for Objective No. 2 is planned for the end of February 1990 concentrating on a selected set of device classes. The effort will continue on Objective No. 2 and No. 3 with final reports available in the late 1990 time frame.

CURRENT MEETING REPORT

First meeting is planned for the February IETF Meeting.

Chapter 3

Network Status Briefings and Technical Presentations

3.1 "Hyper MIB Demonstration"

Presentation by Steve Hunter/ LLNL

The best way to summarize my presentation of my HyperMIB program is to give the entry that I wrote up for the NOCtools Catalogue (See the charter of the NOCtools working group).

NAME: HyperMIB

KEYWORDS: Macintosh, sourcelib, free.

ABSTRACT: HyperMIB is a hypertext presentation of the MIB (RFC1066). The tree structure of the MIB is presented graphically, and the user traverses the tree by selecting branches of the tree. When the MIB variables are displayed, selecting them causes a text window to appear and show the definition of that variable (using the actual text of the MIB document).

MECHANISM: The Apple Macintosh HyperCard utility is used. The actual text of the MIB document was read into scrollable text windows, and a string search is done on the variable selected. A person familiar with HyperCard programming could modify the program to suit their needs (such as to add the definitions for their company's private space).

LIMITATIONS: This program only gives the definition of the MIB variables, it cannot poll a node to find the value of the variables.
HARDWARE REQUIRED: Apple Macintosh computer with at least 1MByte of RAM.

SOFTWARE REQUIRED: Apple Macintosh operating system and HyperCard.

AVAILABILITY: This software may be copied and given away without charge. The files are available by anonymous FTP on CCC.NMFECC.GOV. The files are:

[Anonymous.programs.HyperMIB] Hyper_MIB.help	(ASCII text)
[Anonymous.programs.HyperMIB] Hyper.MIB	(binary)
[Anonymous.programs.HyperMIB] MIB.tree	(binary)'

The software is also available for a nominal fee from:

National Energy Software Center
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439
(312) 972-7250

Introduction Chapter

Objects Chapter

iso.org.dod.internet

iso.org.dod.internet.mgmt

1.3.6.1.2.1

System	1.3.6.1.2.1.1
Interfaces	1.3.6.1.2.1.2
Address Translation	1.3.6.1.2.1.3
IP	1.3.6.1.2.1.4
icmp	1.3.6.1.2.1.5
tcp	1.3.6.1.2.1.6
udp	1.3.6.1.2.1.7
egp	1.3.6.1.2.1.8

ESnet

written by:

Steven Hunter

This software is available from:

National Energy Software Center
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439

system (1)

- sysDescr (1) 1.3.6.1.2.1.1.1
- sysObjectID (2) 1.3.6.1.2.1.1.2
- sysUpTime (3) 1.3.6.1.2.1.1.3

Interfaces (2)

- ifNumber (1) 1.3.6.1.2.1.2.1
- ifTable (2) 1.3.6.1.2.1.2.2
  - ifEntry (1)
    - ifIndex (1) 1.3.6.1.2.1.2.2.1.1
    - ifDescr (2) 1.3.6.1.2.1.2.2.1.2
    - ifType (3) 1.3.6.1.2.1.2.2.1.3
    - ifMtu (4) 1.3.6.1.2.1.2.2.1.4
    - ifSpeed (5) 1.3.6.1.2.1.2.2.1.5
    - ifPhysAddress (6) 1.3.6.1.2.1.2.2.1.6
    - ifAdminStatus (7) 1.3.6.1.2.1.2.2.1.7
    - ifOperStatus (8) 1.3.6.1.2.1.2.2.1.8
    - ifLastChange (9) 1.3.6.1.2.1.2.2.1.9
    - ifInOctets (10) 1.3.6.1.2.1.2.2.1.10
    - ifInUcastPkts (11) 1.3.6.1.2.1.2.2.1.11
    - ifInNUcastPkts (12) 1.3.6.1.2.1.2.2.1.12

more

```
│
├── ifInDiscards (13)          1.3.6.1.2.1.2.2.1.13
├── ifInErrors    (14)         1.3.6.1.2.1.2.2.1.14
├── ifInUnknownProtos (15)     1.3.6.1.2.1.2.2.1.15
├── ifOutOctets        (16)    1.3.6.1.2.1.2.2.1.16
├── ifOutUcastPkts     (17)    1.3.6.1.2.1.2.2.1.17
├── ifOutNUcastPkts  (18)      1.3.6.1.2.1.2.2.1.18
├── ifOutDiscards      (19)    1.3.6.1.2.1.2.2.1.19
├── ifOutErrors        (20)    1.3.6.1.2.1.2.2.1.20
└── ifOutQLen          (21)    1.3.6.1.2.1.2.2.1.21
```

```
at (3)                                     1.3.6.1.2.1.3
│
└── atTable (1)                            1.3.6.1.2.1.3.1
    │
    └── atEntry (1)
        │
        ├── atIfIndex      (1)             1.3.6.1.2.1.3.1.1.1
        ├── atPhysAddress  (2)             1.3.6.1.2.1.3.1.1.2
        └── atNetAddress   (3)             1.3.6.1.2.1.3.1.1.3
```

ip (4)

- ipForwarding (1) — 1.3.6.1.2.1.4.1
- ipDefaultTTL (2) — 1.3.6.1.2.1.4.2
- ipInReceives (3) — 1.3.6.1.2.1.4.3
- ipInHdrErrors (4) — 1.3.6.1.2.1.4.4
- ipInAddrErrors (5) — 1.3.6.1.2.1.4.5
- ipForwDatagrams (6) — 1.3.6.1.2.1.4.6
- ipInUnknownProtos (7) — 1.3.6.1.2.1.4.7
- ipInDiscards (8) — 1.3.6.1.2.1.4.8
- ipInDelivers (9) — 1.3.6.1.2.1.4.9
- ipOutRequests (10) — 1.3.6.1.2.1.4.10
- ipOutDiscards (11) — 1.3.6.1.2.1.4.11
- ipOutNoRoutes (12) — 1.3.6.1.2.1.4.12
- ipReasmTimeout (13) — 1.3.6.1.2.1.4.13
- ipReasmReqds (14) — 1.3.6.1.2.1.4.14
- ipReasmOKs (15) — 1.3.6.1.2.1.4.15
- ipReasmFails (16) — 1.3.6.1.2.1.4.16

more

- ipFragOKs (17) — 1.3.6.1.2.1.4.17
- ipFragFails (18) — 1.3.6.1.2.1.4.18
- ipFragCreates (19) — 1.3.6.1.2.1.4.19
- ipAddrTable (20) — 1.3.6.1.2.1.4.20
  - ipAddrEntry (1)
    - ipAdEntAddr (1) — 1.3.6.1.2.1.4.20.1.1
    - ipAdEntIfIndex (2) — 1.3.6.1.2.1.4.20.1.2
    - ipAdEntNetMask (3) — 1.3.6.1.2.1.4.20.1.3
    - ipAdEntBcastAddr (4) — 1.3.6.1.2.1.4.20.1.4
- ipRoutingTable (21) — 1.3.6.1.2.1.4.21
  - ipRouteEntry (1)
    - ipRouteDest (1) — 1.3.6.1.2.1.4.21.1.1
    - ipRouteIfIndex (2) — 1.3.6.1.2.1.4.21.1.2
    - ipRouteMetric1 (3) — 1.3.6.1.2.1.4.21.1.3
    - ipRouteMetric2 (4) — 1.3.6.1.2.1.4.21.1.4
    - ipRouteMetric3 (5) — 1.3.6.1.2.1.4.21.1.5
    - ipRouteMetric4 (6) — 1.3.6.1.2.1.4.21.1.6
    - ipRouteNextHop (7) — 1.3.6.1.2.1.4.21.1.7
    - ipRouteType (8) — 1.3.6.1.2.1.4.21.1.8
    - ipRouteProto (9) — 1.3.6.1.2.1.4.21.1.9
    - ipRouteAge (10) — 1.3.6.1.2.1.4.21.1.10

icmp (5)	1.3.6.1.2.1.5
icmpInMsgs (1)	1.3.6.1.2.1.5.1
icmpInErrors (2)	1.3.6.1.2.1.5.2
icmpInDestUnreachs (3)	1.3.6.1.2.1.5.3
icmpInTimeExcds (4)	1.3.6.1.2.1.5.4
icmpInParmProbs (5)	1.3.6.1.2.1.5.5
icmpInSrcQuenchs (6)	1.3.6.1.2.1.5.6
icmpInRedirects (7)	1.3.6.1.2.1.5.7
icmpInEchos (8)	1.3.6.1.2.1.5.8
icmpInEchoReps (9)	1.3.6.1.2.1.5.9
icmpInTimestamps (10)	1.3.6.1.2.1.5.10
icmpInTimestampReps (11)	1.3.6.1.2.1.5.11
icmpInAddrMasks (12)	1.3.6.1.2.1.5.12
icmpInAddrMaskReps (13)	1.3.6.1.2.1.5.13

more

icmpOutMsgs (14)	1.3.6.1.2.1.5.14
icmpOutErrors (15)	1.3.6.1.2.1.5.15
icmpOutDestUnreachs (16)	1.3.6.1.2.1.5.16
icmpOutTimeExcds (17)	1.3.6.1.2.1.5.17
icmpOutParmProbs (18)	1.3.6.1.2.1.5.18
icmpOutSrcQuenchs (19)	1.3.6.1.2.1.5.19
icmpOutRedirects (20)	1.3.6.1.2.1.5.20
icmpOutEchos (21)	1.3.6.1.2.1.5.21
icmpOutEchoReps (22)	1.3.6.1.2.1.5.22
icmpOutTimestamps (23)	1.3.6.1.2.1.5.23
icmpOutTimestampReps (24)	1.3.6.1.2.1.5.24
icmpOutAddrMasks (25)	1.3.6.1.2.1.5.25
icmpOutAddrMaskReps (26)	1.3.6.1.2.1.5.26

```
tcp (6)
├── tcpRtoAlgorithm (1)                  1.3.6.1.2.1.6.1
├── tcpRtoMin (2)                        1.3.6.1.2.1.6.2
├── tcpRtoMax (3)                        1.3.6.1.2.1.6.3
├── tcpMaxConn (4)                       1.3.6.1.2.1.6.4
├── tcpActiveOpens (5)                   1.3.6.1.2.1.6.5
├── tcpPassiveOpens (6)                  1.3.6.1.2.1.6.6
├── tcpAttemptFails (7)                  1.3.6.1.2.1.6.7
├── tcpEstabResets (8)                   1.3.6.1.2.1.6.8
├── tcpCurrEstab (9)                     1.3.6.1.2.1.6.9
├── tcpInSegs (10)                       1.3.6.1.2.1.6.10
├── tcpOutSegs (11)                      1.3.6.1.2.1.6.11
├── tcpRetransSegs (12)                  1.3.6.1.2.1.6.12
└── tcpConnTable (13)                    1.3.6.1.2.1.6.13
    └── tcpConnEntry
        ├── tcpConnState (1)             1.3.6.1.2.1.6.13.1.1
        ├── tcpConnLocalAddress (2)      1.3.6.1.2.1.6.13.1.2
        ├── tcpConnLocalPort (3)         1.3.6.1.2.1.6.13.1.3
        ├── tcpConnRemAddress (4)        1.3.6.1.2.1.6.13.1.4
        └── tcpConnRemPort (5)           1.3.6.1.2.1.6.13.1.5
```

```
udp (7)
├── udpInDatagrams (1)     1.3.6.1.2.1.7.1
├── udpNoPorts (2)         1.3.6.1.2.1.7.2
├── udpInErrors (3)        1.3.6.1.2.1.7.3
└── udpOutDatagrams (4)    1.3.6.1.2.1.7.4
```

```
egp (8)
|
|--- egpInMsgs (1)                    1.3.6.1.2.1.8.1
|--- egpInErrors (2)                  1.3.6.1.2.1.8.2
|--- egpOutMsgs (3)                   1.3.6.1.2.1.8.3
|--- egpOutErrors (4)                 1.3.6.1.2.1.8.4
`--- egpNeighTable (5)                1.3.6.1.2.1.8.5
      |
      `--- egpNeighEntry (1)
            |
            |--- egpNeighState (1)    1.3.6.1.2.1.8.5.1.1
            `--- egpNeighAddr (2)     1.3.6.1.2.1.8.5.1.2
```

3. Introduction

As reported in RFC 1052, IAB Recommendations for the Development of Internet Network Management Standards [1], the Internet Activities Board has directed the Internet Engineering Task Force (IETF) to create two new working groups in the area of network management. One group is charged with the further specification and definition of elements to be included in the Management Information Base. The other is charged with defining the modifications to the Simple Network Management Protocol (SNMP) to accommodate the short-term needs of the network vendor and operator communities. The long-term needs of the Internet community are to be met using the ISO CMIS/CMIP [2,3] framework as a basis. An existing IETF working group, the "NETMAN" group, is already engaged in defining the use of CMIS/CMIP in a TCP/IP network, and will continue with responsibility for addressing the longer-term requirements.

4. Objects

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. Objects in the MIB are defined using Abstract Syntax Notation One (ASN.1) [9].

The mechanisms used for describing these objects are specified in the companion memo. In particular, each object has a name, a syntax, and an encoding. The name is an object identifier, an administratively assigned name, which specifies an object type. The object type together with an object instance serves to uniquely identify a specific instantiation of the object. For human convenience, we often use a textual string, termed the OBJECT DESCRIPTOR, to also refer to the object type.

The syntax of an object type defines the abstract data structure

3.2 "The CERT"

Presentation by Richard Pethia/ CMU

The following is a summary of the audience questions and discussion about the CERT.

Q: Noel Chiappa: Will you be using Tigre Teams? A: We may use them with prior authorization. We prefer to call it a Security Audit.

Q: Phill Gross: Is the primary focus of the CERT UNIX and the TCP/IP internet? A: Yes

Q: Phill: But you were involved in DECnet worm? A: Yes, because other emergency coordination groups do not yet exist, and the CERT is well known, so we took the lead. The lines of responsibility are not clearly defined.

Q: Phill: Those lines will become increasingly blurred as the internet moves to multiprotocol, including DECnet, and OSI. A: They will continue to be fuzzy. That seems to be OK with the people we work with.

Q: Phill: Did you say there were 68 incidences in the last few months? A: Yes. Most of those attempts were trying to get password files, or stealing cycles to run password guessing games. Many of these breakins were unknown to the users until we reported them and gave them back their own password files.

Q: Noel: It looks like you are having success finding the holes. Are you having any success finding the users of the holes? A: No. These folks are clever. We need FBI cooperation for phone line traces, but there is a reluctance to come forward to the FBI for three main reasons. 1) There is a concern about negative publicity, and a lack of understanding about investigative body. 2) There is a fear that the FBI will take over their operations, offering lots of hassle to have the monitoring and tracing necessary to stop an intruder. 3) In the past there has been little response when they do call. There is now an intense training program in the FBI, and they have been more successful prosecutions.

Q: A: There is often a cross purpose between the CERT and the FBI. We are in the business of helping others protect themselves. We tell affected people of the security breach immediately so they can close the hole. The FBI tries to catch the intruder, and likes to keep holes open to trap intruders in their next invasion. There is a new concern for the liabilities of leaving a system open if the open system may be used to attack other systems.

Q: Are you still liable if the Justice Department tells you to leave the hole open? A: An administrator is always subject to civil action... I'm not a lawyer and do not know the history, but liability is a factor in cooperation with investigations.

Q: When you find and publicize a problem, aren't you liable if other people use the holes you publicize to attack another system? A: We are aware of the problems, and are very careful to notify the users about the problem without specifying the problem. We then go to the vendors and tell them about their problem. If there is a security list, you can be sure that hackers are reading the list also, so we are very careful about what we send on the list. We try to notify affected users on a personal level, but where there is a widespread problem, we are forced to use public means of dissemination

We have received mixed reaction from vendors. We all talk a good game in security, but when we purchase equipment and software, we send a different message. We buy on performance, cost, features, and security is not high on the list.

Q: Has there been any discussion on vendor liability A: I have not heard much about it. Vendors don't like to talk about their liability. There is talk of using civil suits against vendors, but again, I am not a legal expert.

Q: There was a suggestion that the November 2nd worm exploited a hole known to the software provider and that there were possible lawsuits resulting. A: That is a real possibility. I just do not know the history.

Q: Concerning the November 16th DECnet worm. My impression is that the mailbridges were shut down as a defense. A: That is my impression also. Q: Were you aware that they were going to do that before they did it? A: No, we learned 1 hour after it occurred. Q: So I guess it would be unacceptable to criticize you. In the future, please notify the user community when major portions of the net are being isolated. It is a good policy to notify of any major outages. A: I understand some of the reasons for the mailbridges shutdown, but cannot comment about the timing or duration of the shutdown. Q: There was concern about the Mailbridges being down for 18 hours after the viruses were known to be inert across the bridges. If DCA won't trust you, who can be trusted. Why did they not believe you. A: Trust has to be built over time.

Q: Craig Partridge: There are several considerations that can be answered by analogy to the Post Office. While there is a continuous debate about proper use of the network, one principle should hold. As the Post Office makes sending cash through the mail a crime, and discourages the sending of blank checks, so to should we discourage the sending of anything valuable to an outsider over the net. A: It is hard for me to see how to implement this. Q: Craig: For example, it would not be wise for Nysernet to offer to pass in the clear ATM traffic, because it would expose them to more cracking. A: Concerns a user should have on the internet 1) do you want others to publish your papers before you do, 2) do you want people to change your programs, i.e., implement trojan horses. Because your work is in the public domain, you should not assume you are safe. I don't know what it is but there is value on the network. People are making a living mapping the network and grabbing passwords.

Q: Phill: Would it be reasonable to write a document on secure system configuration? If so, I encourage you to publish it as an RFC A: Sounds good.

Q: Do you collect data on security practices and make recommendations on avoiding those practices? A: No, but we are trying to gather information from other sources, like the rainbow collection. The information is useful, but has to be tailored into the internet community.

Q: Phill: Vendors sometime leave bugs in and administrators leave holes even when well known remedies would help. We need to write a document in a working group analogous to the Host Requirements document. A: Let's Talk.

Carnegie Mellon University
Software Engineering Institute

CERT

Computer Emergency Response Team

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Sponsored by the
U.S. Department of Defense

Carnegie Mellon University
Software Engineering Institute

CERT/CC Mission:

To supplement existing mechanisms by which informally organized experts deal with computer emergencies and their prevention

- reliable, trusted, 24-hour, single point of contact
- maintain accessible, secure repository of information
- facilitate communications
- conduct research targeted at improving security of existing systems without compromising functionality, performance, openness
- take proactive measures to raise awareness of security issues

Key Points (1)

Role in both response and prevention

Security awareness role from perspectives of user and technology producer

Five groups in each CERT

- executive
- action
- associate
- industry
- system administrators

INTERNET CERT (SEI/CERT/CC) is a prototype for others

CERT/CC has no authority

CERT System

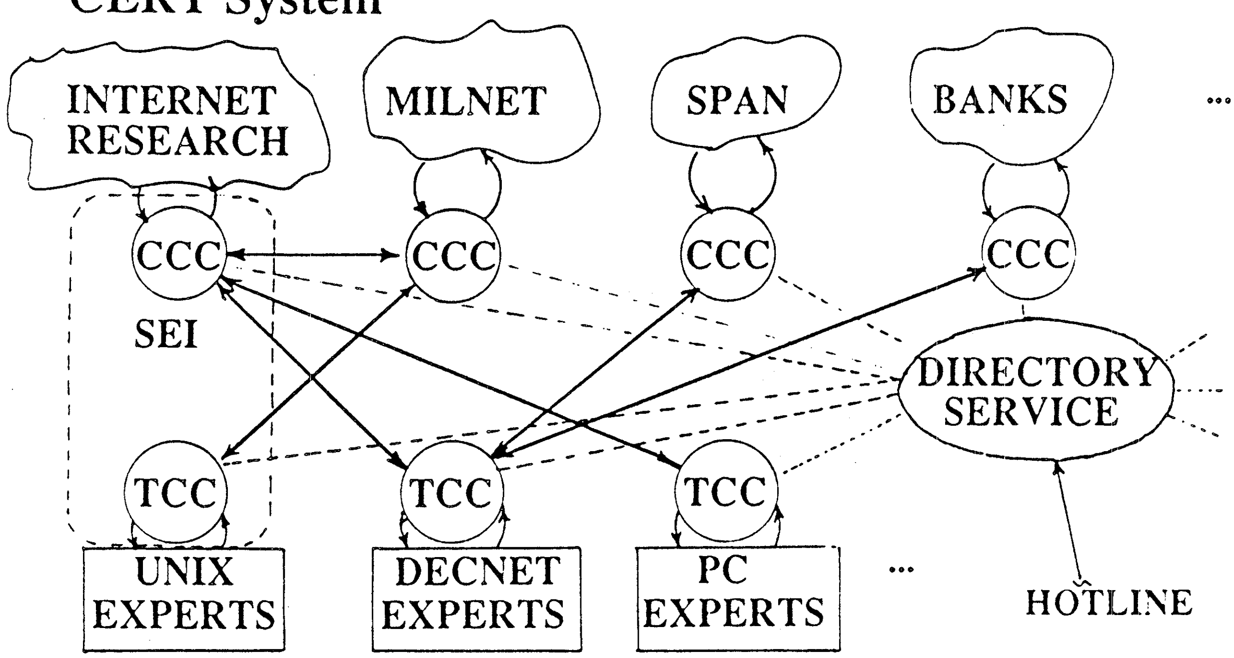

CCC - Constituency Coordinating Center
TCC - Technology Coordinating Center

Key Points (2)

DARPA because

- major sponsor of networking (ARPANET), operating systems (early UNIX and derivatives), and trust/security research
- sponsor of ongoing research facilitated by network

SEI because

- uniquely positioned among government, industry, academia
- chartered, organized, effective in technology transition: catalyst for change

Improved security *should not* hamper innovation, interoperability, performance, functionality, flexibility

Research community can act swiftly

Carnegie Mellon University
Software Engineering Institute

Organizing and Administering (1)

Establish contacts, communications mechanisms

Detail working relationships

Define, design, build tools and systems

Event-handling procedures for identification, classification, resolution

- determine nature and magnitude of threat
- assess vulnerability
- gauge response

Handle "sensitive" information

Produce guidelines and lessons learned

Organizing and Administering (2)

Run "system tests"

Maintain interaction histories

Develop response packages

Reactive Activities (1)

Determine nature of problem, magnitude of threat, vulnerability

- solicit help from associates and facilitate communication
- provide information to constituents
  - problem
  - counter measures
- assist constituents' efforts to assess vulnerability

Assist associates in problem resolution

Notify appropriate agencies and CERT system of problem/progress

Reactive Activities (2)

Maintain activity logs

Coordinate press releases

Facilitate postmortems to capture lessons learned

Proactive Activities (1)

Produce information packages: security issues

Develop effective distribution mechanisms: seminars, workshops, documents, video

Maintain registry of information on software packages: checksums, signatures, registered fixes

Support system administrators' efforts to verify the state of their software

Develop and document procedures for "security audits"

Assist the self-audit process

Proactive Activities (2)

Develop working relationships with vendors: inform them of problems and track progress

Learn the law and provide pointers to constituents

Learn the policies and provide technical guidance

Carnegie Mellon University
Software Engineering Institute

Help Organize the CERT System

Build working relationships with agencies and their contractors: NCSC, NIST, DCA, DOE, NASA, NSF, SRI, FBI, Treasury, BBN, MERIT

Build working relationships with industry and user groups: SUN, DEC, IBM, AT&T, ..., USENIX, /usr/group

Host and facilitate workshops

Develop mechanisms to spread information across CERTs

Develop information packages and dissemination vehicles to communicate CERT concepts, working relationships, status of CERT system

Develop mechanisms to gauge progress

Status

Phone system established

- 24-hour, 7-day/week coverage

Computer systems in place

- primary system for communications
- secondary (stand-alone) system for sensitive data

Event-handling procedures developing with experience

- active in several types of events

Over 500 contacts with industry, government, research community

Databases on vulnerabilities, fixes, configurations, and events being built

Carnegie Mellon University
Software Engineering Institute

Contact Information

For Emergencies:	(412) 268-7090
For Information:	(412) 268-7080
FAX:	(412) 268-5758
Electronic Mail:	CERT@SEI.CMU.EDU
U.S. Mail:	CERT/CC Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213-3890

3.3 "Internet Status Report"

Presentation by Zbigniew Oplaka/BBN

STATE OF THE INTERNET

Zbigniew Opalka
November 2, 1989

BBN Communications Corporation

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH

INTERNET GROWTH SUMMARY

- 997 Networks advertised
- 1710 Networks registered

NUMBER OF NETWORKS
December 1983-October 1989

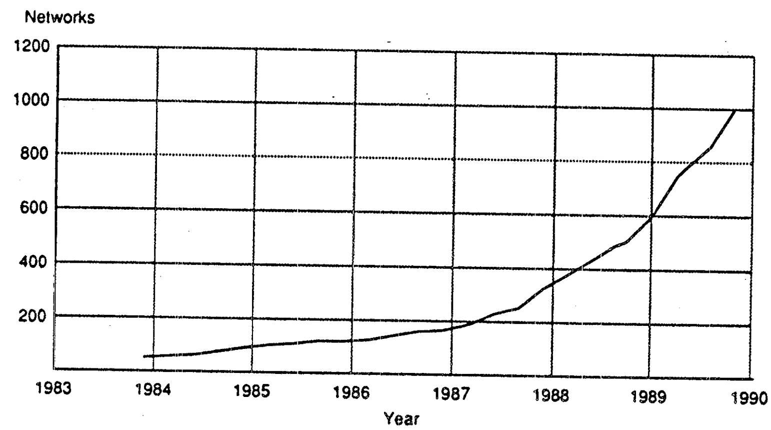

NUMBER OF NETWORKS
January 1989-October 1989

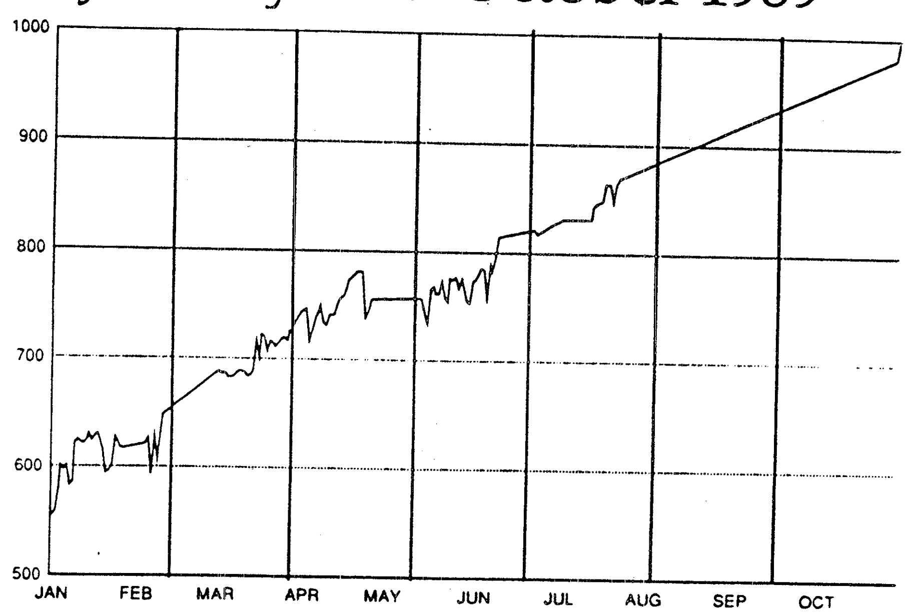

DDN MAILBRIDGES

CURRENT STATUS

- Six DDN Butterfly Mailbridges operational
- 214 EGP neighbors
- Ethernet interfaces added to Mitre and Ames mailbridges
  - 192.52.194-NSFTRANSIT 5
  - 192.52.195-NSFTRANSIT 6
- Access Control turned on sporadically
- BMILLBL has only one interface
  - ARPANET interface eliminated
  - Provides EGP server function on MILNET

TRAFFIC SUMMARY

- ~ 13,000,000 packets/day forwarded
- .3 - .7% packets dropped
- Average Bytes per packet
  Low - 68
  High - 175

BMILAMES DAILY THROUGHPUT

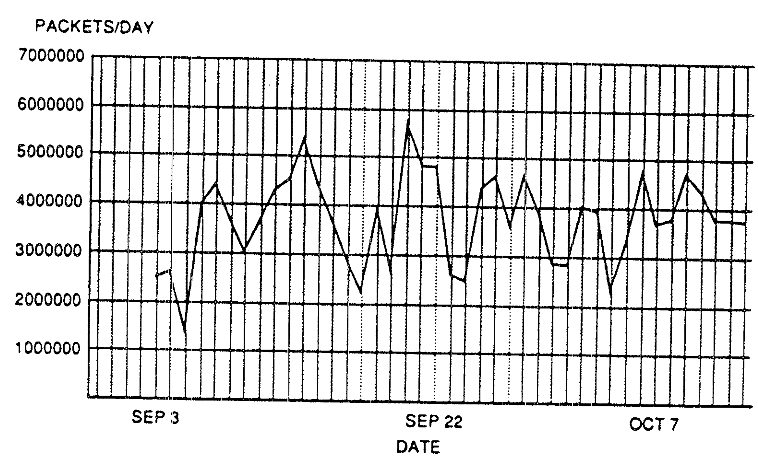

3.4 "A Selective Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer"

Presentation by K. K. Ramakrishnan/DEC

In this talk we discuss the problem of congestion in computer networks and introduce the concept of congestion avoidance. We then describe a scheme for congestion avoidance in a network with a connectionless network layer. This distributed scheme attempts to operate the network at the optimal point by having explicit feedback of information from congested nodes in the network.

The constraint and the resulting feature of the scheme is that there is only a single bit of congestion information fed back to the sources generating traffic. Sources react to this feedback information by adjusting their flow control window dynamically. Fairness considerations as well as the dynamic response to transients of the scheme is discussed.

In this talk we relax the assumption of having the same set of resources being shared by all the users of the network. We define a more general fairness goal to achieve in the light of the relaxation of the assumption. We present a solution alternative to achieve this goal, by having the routers in the network selectively feedback the congestion information only to those users that are using more than their fair share of each individual router's resources.

This work was performed jointly with Raj Jain and Dah-Ming Chiu.

A Selective Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer

K.K. Ramakrishnan, Dah-Ming Chiu and Raj Jain

Digital Equipment Corp.

550 King St. (LKG 1-2/A19)

Littleton, MA 01460-1289

ARPAnet: Rama%Erlang.dec@DECWRL.DEC.COM,
Chiu%Erlang.dec@DECWRL.DEC.COM,
Jain%Erlang.dec@DECWRL.DEC.COM

digital

kkr

Congestion Avoidance

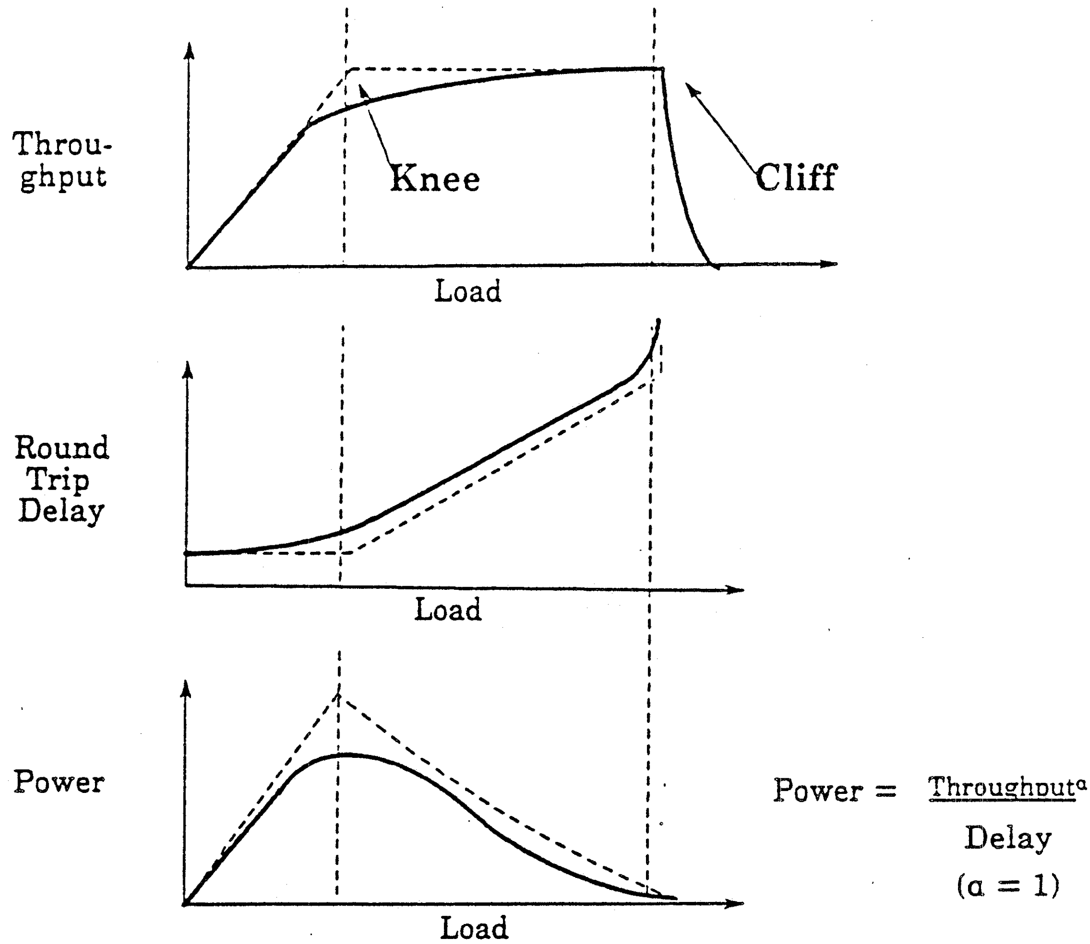

Congestion Control:

Recover from *zero* throughput and *infinite* delay state

"Left of cliff policies"(depend on number of buffers)

Congestion Avoidance:

Maintain *high* throughput and *low* delay.

"Operate at knee" policies (independent of buffers)

Need a congestion control mechanism to recover from impulse load

Design Goals and Requirements

1. Efficient.

Resource Efficiency = $\frac{\text{Throughput/Knee Throughput}}{\text{Response Time/Knee Response Time}}$

Network Efficiency = Efficiency of Bottleneck

2. Fairness.

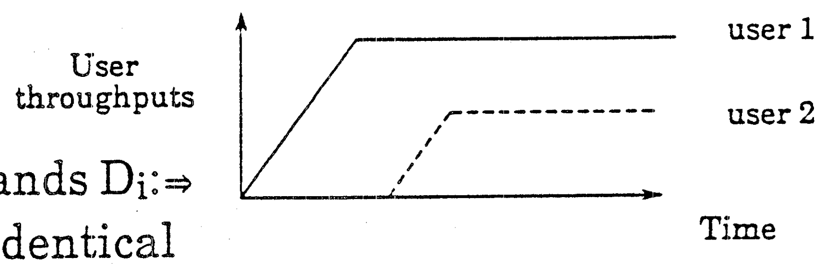

Identical user demands D_i: ⇒ allocations A_i also identical

3. Distributed Control

4. No new packets: during overload or underload.

5. All parameters should be dimensionless: no time values; suitable for all link speeds and network sizes.

6. Low parameter sensitivity; robust to noise.

7. Responsive

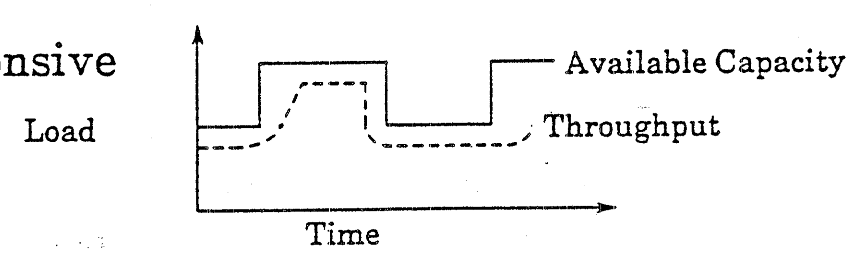

The Binary Feedback Scheme

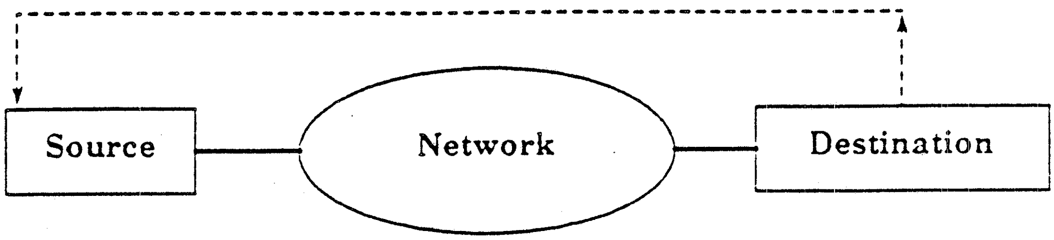

Network Layer Policies:

1. Routers average # packets in queue (incl. in service).

If congested, set congestion avoidance bit on **all** packets :

Ave. Queue Length ≥ 1, $\Rightarrow$ congestion.

User (Transport Layer) **Policies** :

Identical Policies at all users for window adjustment.

1. Decision Frequency ; Signal filter;

2. Increase/Decrease: increase : $+1$; decrease : aW $(0 < a < 1)$

Assumption: all sources share the same path.

Result: Identical window sizes for all users.

If paths are different, e.g.,with same bottleneck for 2 users:

$R_1 \neq R_2$ (R_i are round trip times).

Throughputs, $T_i = W/R_i$; $\therefore T_1 \neq T_2$.

Fairness goal: provide equal service at bottleneck resource

i.e., $T_1 = T_2$ is desired goal

Behavior of Window size with Non-identical demands

Configuration 1:

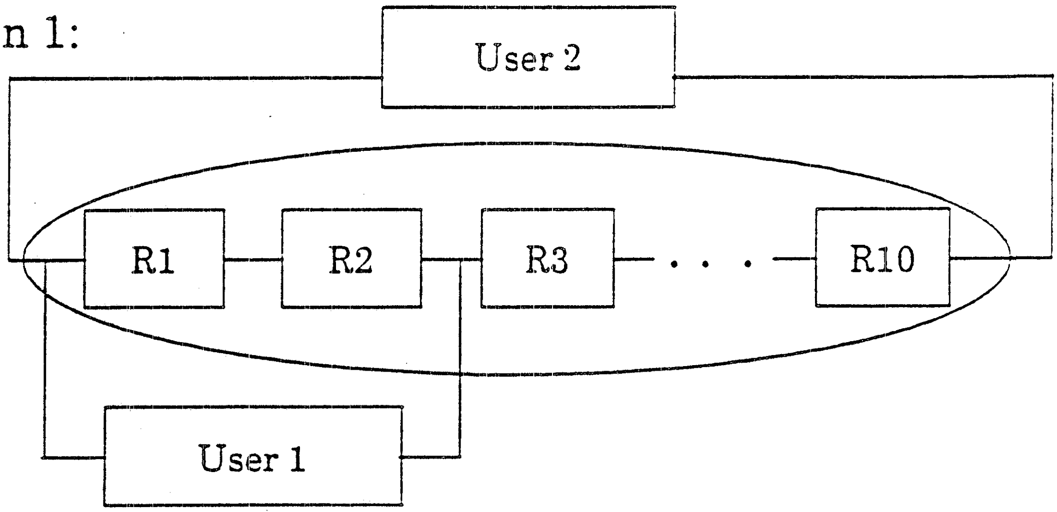

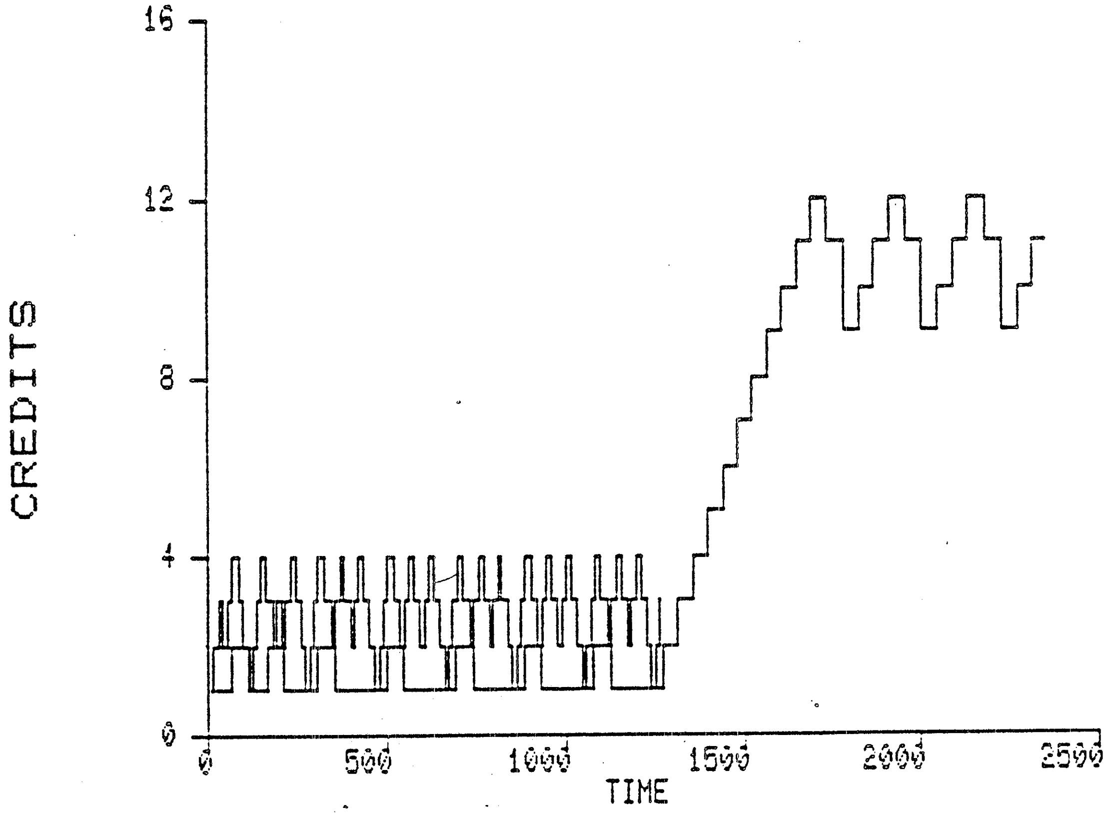

Behavior of Throughput with Non-identical demands

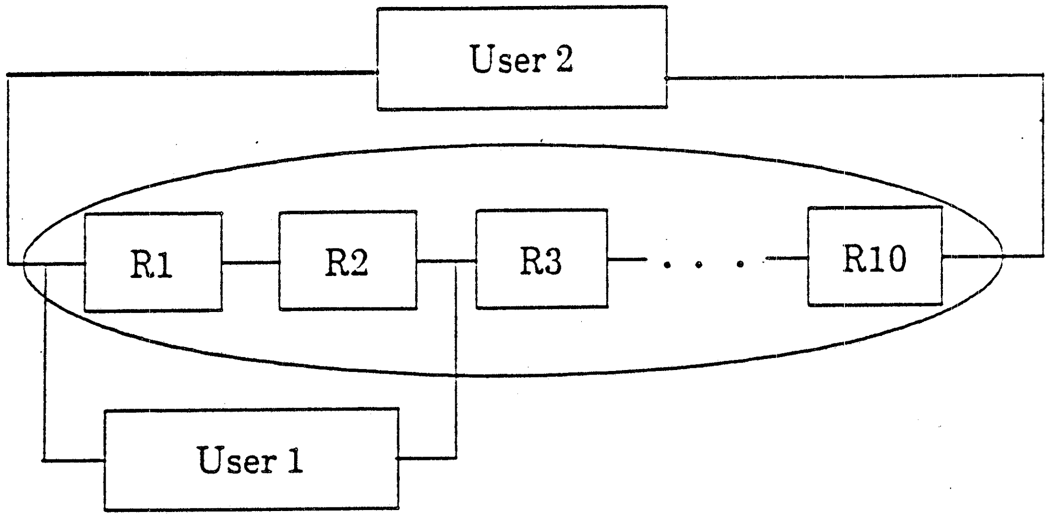

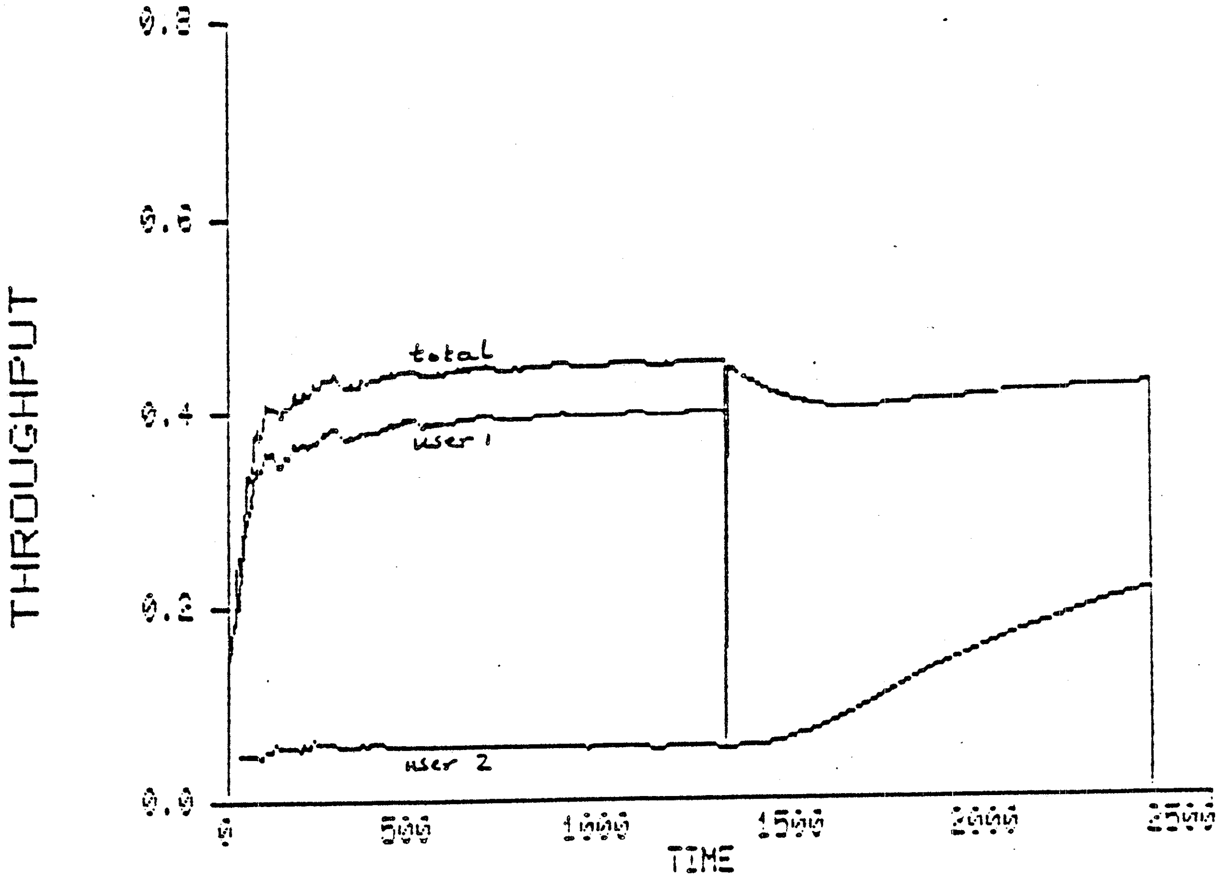

digital

kkr

Fairness : Users with unequal demands

Single Resource

Maximally efficient operating point: Knee

Fair share of resource with identical demands = D:

$$A_i = A \ (= C^{knee}/n)$$

When demands D_i are not equal:

If $A_i = C^{knee}/n$, part of resource may be wasted - inefficient.

Intuitively, want: $A_i = D_i$, for all i, when $D_i <$ fair__share
remaining capacity provided to others
i.e., A_j = fair__share for all $j \neq i$.

Maximally fair Allocation: $A_i^* = \min(D_i, A_{fair})$.

Finding A_{fair} and A_i^* is iterative, given C_{knee} & D_i.

Example: 3 users demand {60, 40, 10}
Knee capacity, $C_{knee} = 100$
Fair allocation = {50, 40, 10}

Multiple Resources/Multiple Users

Multiple users share different sets of resources

Consider Omniscient observer:

Goal: Achieve global optimal - fair and efficient allocation
$A^* = \{A_1^*, A_2^*, \ldots, A_n^*\}$

Let $R = \{1,2,\ldots,m\}$ = set of resources.
Let $C_j^{knee}, j = 1,2,\ldots,m$ = knee capacity of resource j.
Let $U = \{1,2,\ldots n\}$ = set of users.
Each User has a **Path** $P_i \subseteq \{R_1, R_2, \ldots, R_m\}$

Algorithm

1) Initally, $S = U$; $M = R$ and $C_j = C_j^{knee}$ for each j.

2) For each j, N_j = #users in S contending for R_j.

3) For each resource j, $B_j = C_j / N_j$ (computing per-user cap.)

4) Compute $B_k = \min(B_j)$ for all $j \in M$; let the resource be k.
(find bottleneck)

5) Remove resource k from M.

6) If resource $k \in P_i$, for user i, $A_i^* = B_k$

7) For each resource $j \in P_i$, and still in M, $C_j = C_j - B_k$
remove i from S.

8) If S is empty, STOP. Else, repeat steps 2 through 7.

Converge to maximally fair and efficient allocation.

Solution alternatives for Multiple Resources/Multiple Users

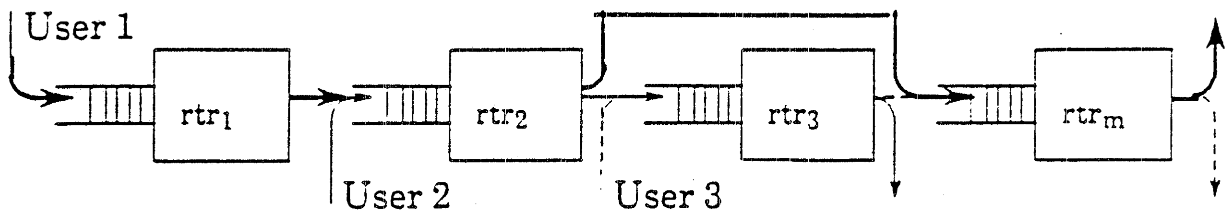

Alternatives:

(1) Keep Network policy the same:
Congested router sets bit on all packets.
Transport policies (user) are different.

(2) Keep Transport policies same:
Congested routers treat users differently
selective setting of congestion avoidance bit by router.

Goal:

(1) Achieve efficiency.
Resource efficiency = Resource Power/Resource Power at knee

Network Efficiency = Efficiency of the Bottleneck Resource

(2) Achieve fair allocation.

System Fairness = $(\Sigma x_i)^2/(n\Sigma x^2_i)$ where $x_i = A_i/A^*_i$

Selective Binary Feedback

Identify distinct 'customers' at each router.

Determine demands of *each* customer at *each* router.

When router is congested:
identify customers contributing to congestion
and receiving unfairly large share of capacity

Set 'congestion avoidance bit' on packets for these selected customers only.

For example: Demands={60,40,10} and C_{knee}=100,
only user 1 has bit set by the router.

Customer demand & allocation: Throughput (#pkts/sec).

Customers?

Destinations? sufficient?

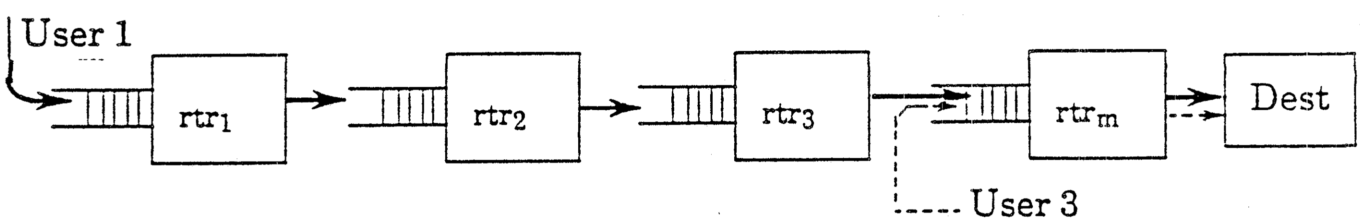

Sources? Similar difficulties..

Source-Destination pairs?

Decision based on granularity of fairness desired/overheads.

Selective Feedback Algorithm

Let C^{knee} = knee capacity of router.
Let D_i = demand of user i, $i = 1,\ldots,n$.
Let S = set of users whose allocations have to be found.
Let C = remaining capacity of router.

Problem: find optimally fair allocation at router:

$$A = \{A_1, A_2, \ldots, A_n\}$$

Invoke algorithm *only* if $\Sigma D_i > C^{knee}$

1) Initially, card[S] = n, $C = C^{knee}$, and $A_i = 0$, $i = 1,\ldots,n$.

2) $C = C^{knee} - \sum_{i \in S} A_i$

3) $A_{fair} = C/\,\mathrm{card}[S]$

4) for each user, $j \in S$, whose $D_j < A_{fair}$,
   $A_j = D_j$
   $C = C - A_j$
   remove j from the set S.

5) If no new allocations were made in Step 4, STOP.
   Else, repeat steps 3 through 4.

Set the bit on users, i, whose $D_i > A_i$.

Estimation of Knee Capacity

Knee capacity: Throughput of router operating at knee.

Router is beyond knee when: Average # at resource > 1.

Deterministic packet sizes:
knee capacity = # packets processed by 'congested router'.

Exponential packet sizes:
knee capacity $\approx\leq$ 0.5* packets processed by 'congested rtr'.

Knee capacity $\approx$ *cf* * # packets processed by 'congested rtr'

cf too small : more 'users' signalled than optimal.
a congested router's capacity estimated to be < C_{knee}

cf too large : fewer 'users' signalled than optimal.
a congested router's capacity estimated to be > C_{knee}

However, the mechanism is relatively insenstive to *cf*, in the range of $0.6 \leq cf \leq 0.9$.

Robust to configuration changes.

Interval for Demand Estimation

Average number at the resource estimated over a regeneration cycle : (busy + idle) interval.

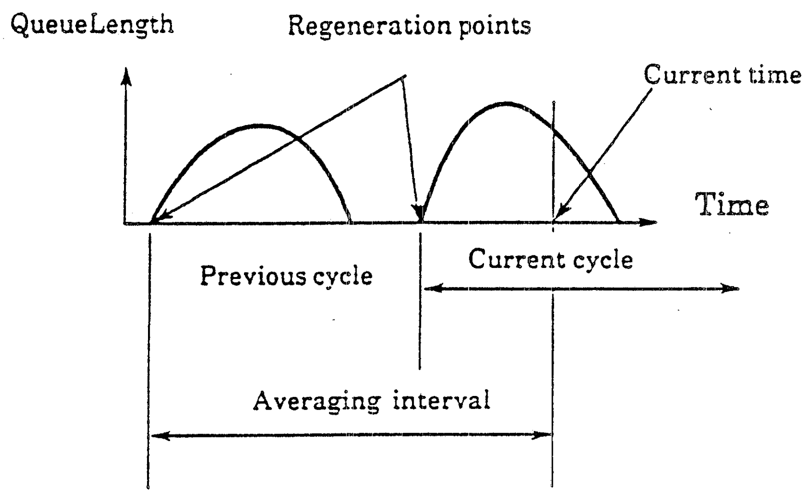

Average demands are also based on the same averaging interval.

Allocation action: setting the congestion avoidance bit.
Synchronized with the estimation of demand.

Behavior of Throughput with Selective Feedback

Configuration 1: user1: path - r1 - r2;
user2: path - r1 - r10.

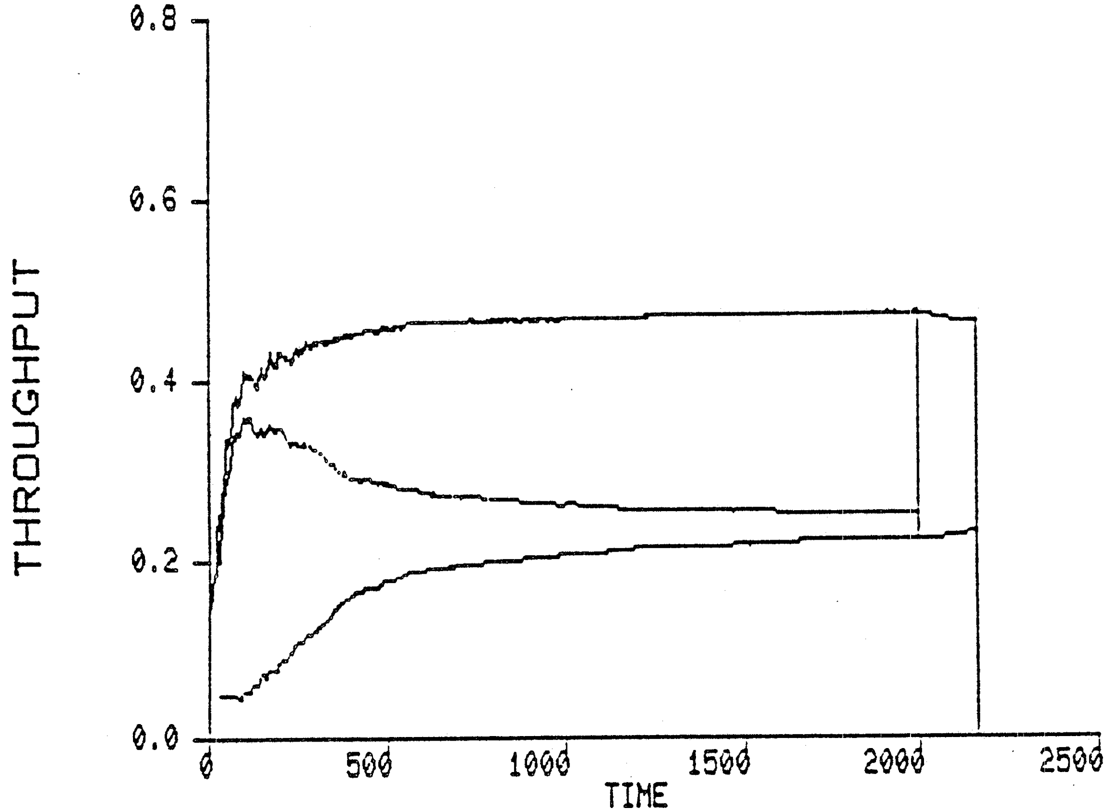

Observe: considerable improvement in fairness

Behavior under Overload

Configuration: 9 users sharing 4 routers.
Window size at knee = 3 (1/3 per user).

During a long busy period:

Accumulated demands of active customers still presenting load on router dominate.
Selective feedback algorithm has a delay in identifying new users obtaining greater than fair share of congested router.

Turn off "feedback selector" when severely overloaded.

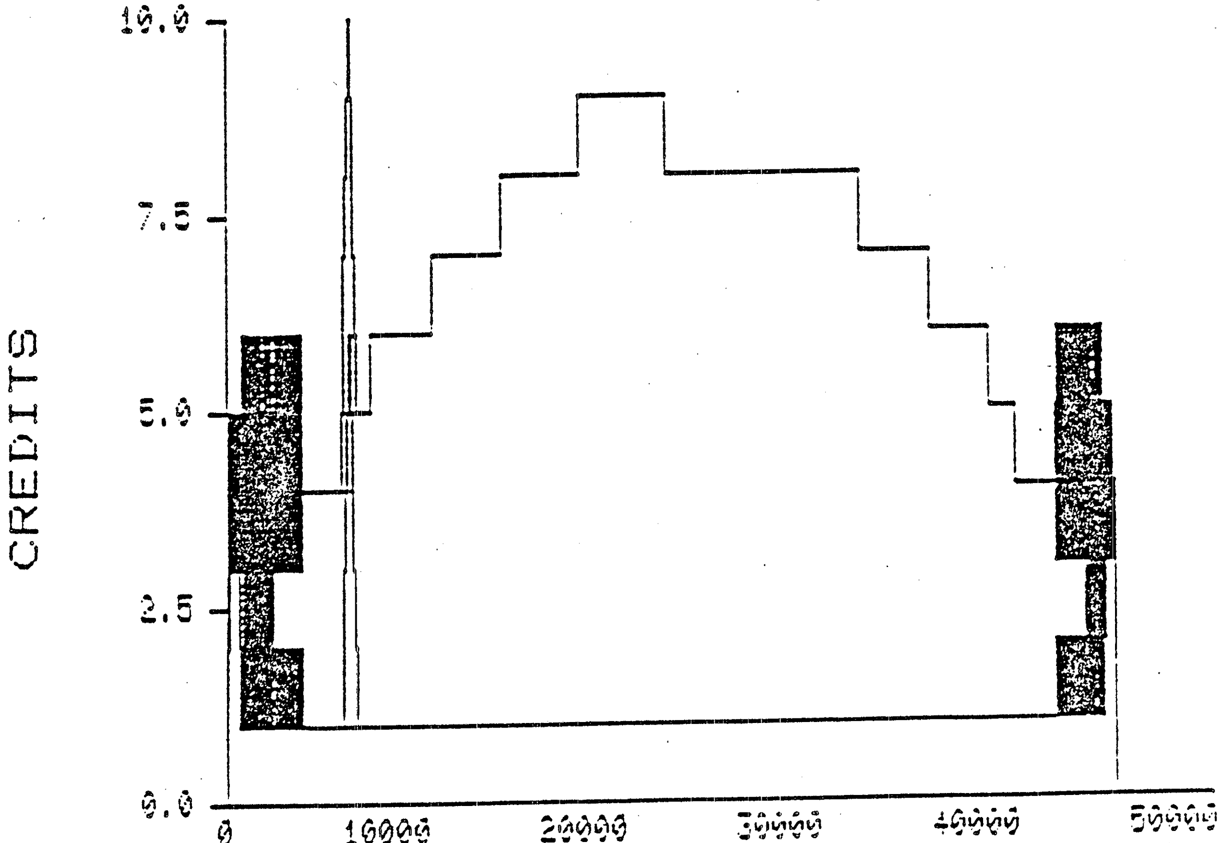

digital

kkr

Behavior with Random Service Time Distribution

- Packet Size: Uniformly distributed (0.5, 1.5)

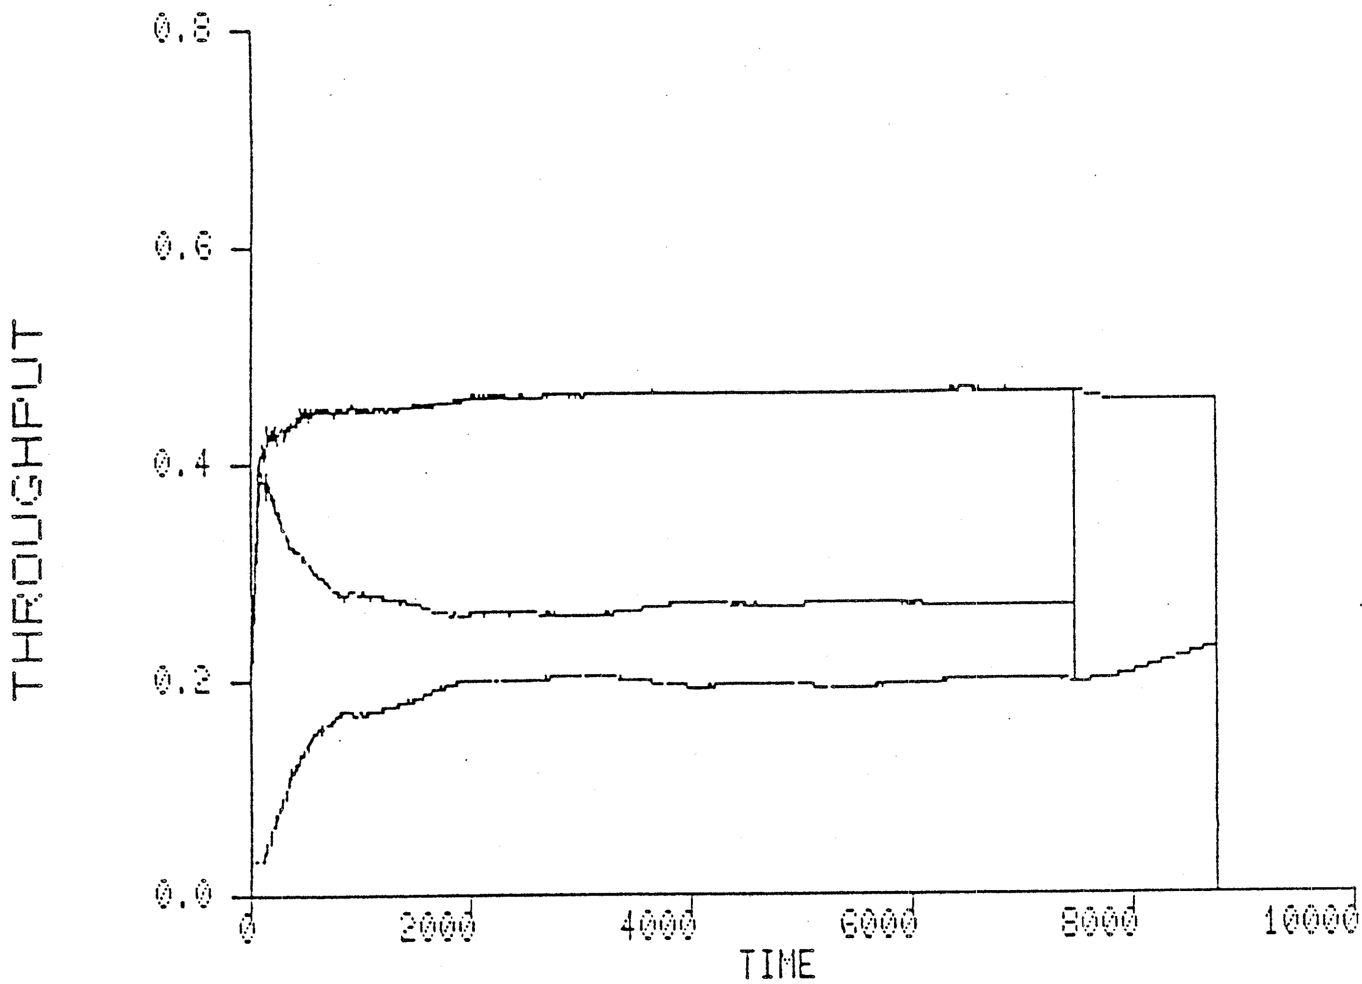

- Achieves Efficiency; convergence to Fair allocation takes longer.

Behavior with Users starting at Arbitrary Windows

- Behavior of Window Size

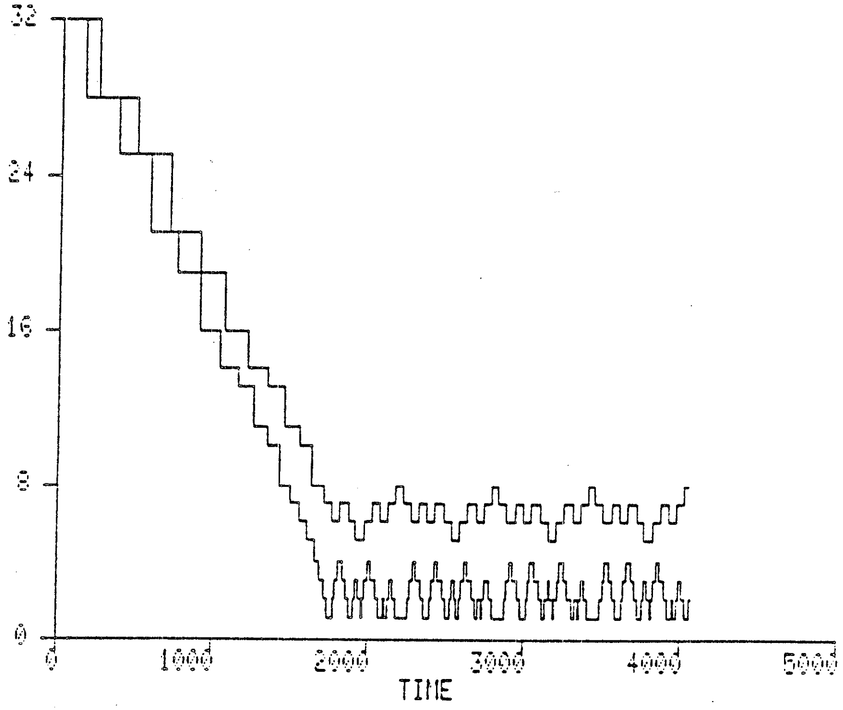

- Behavior of Throughput

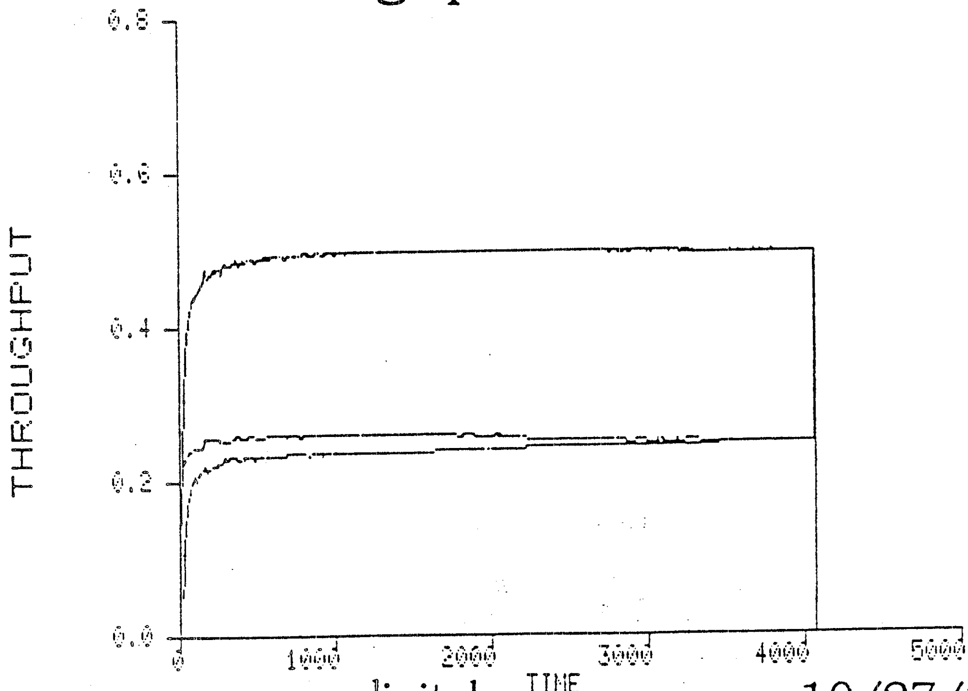

Behavior with Transient User Demand

- Configuration

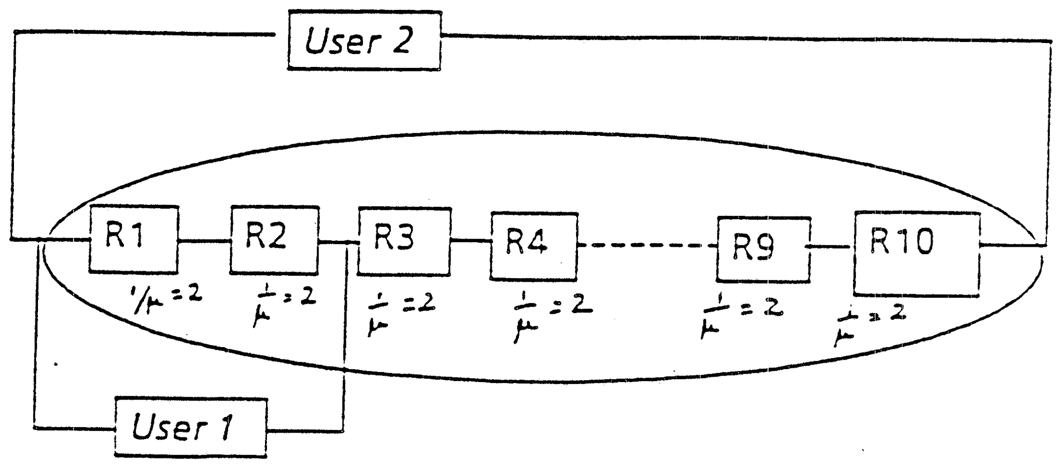

- Behavior of Throughput

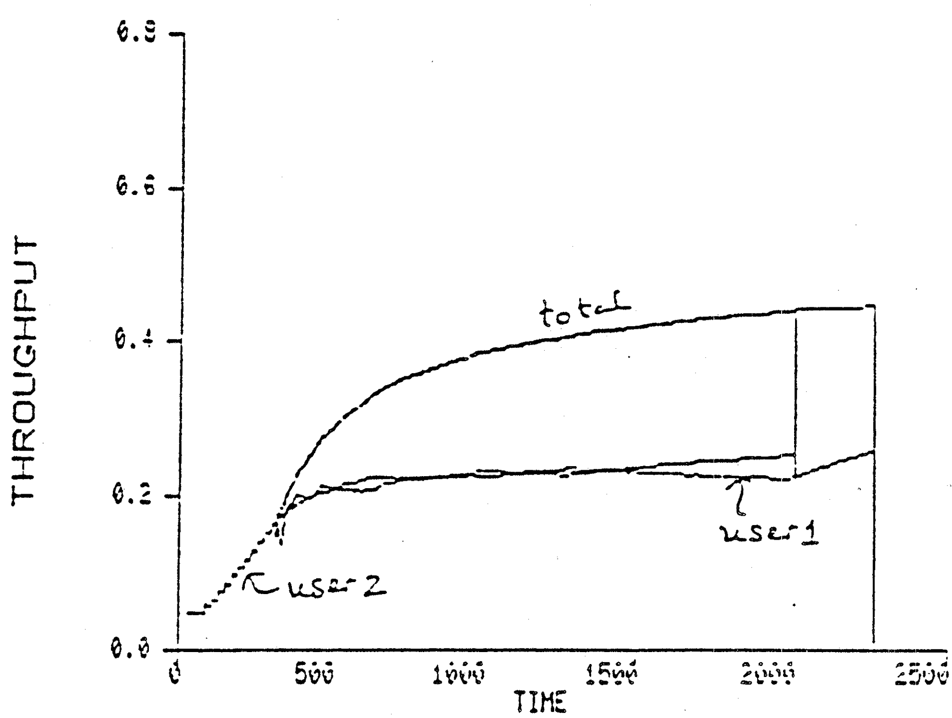

Summary

1. Congestion is not a static problem.

2. Congestion Avoidance:
   Operation with low delay and high throughput
   Independent of number of buffers.

3. Congestion can be avoided in connectionless networks.

4. Binary feedback Scheme:

Network Layer: Congestion Detection ($Q_{avg} = 1$)
Feedback filter (Avg since last cycle)
Feedback Selector (Only near knee)

Transport Layer: Decision fn (Collect $W_{old} + W$ bits, Examine the last W bits)
Signal filter (up if < 50 bits set)
Increase/Decrease ($W + 1$, $0.875W$)

5. The proposed scheme is efficient, fair, responsive, convergent, and robust.

3.5 "ESnet Status Report"

Presentation by Tony Hain

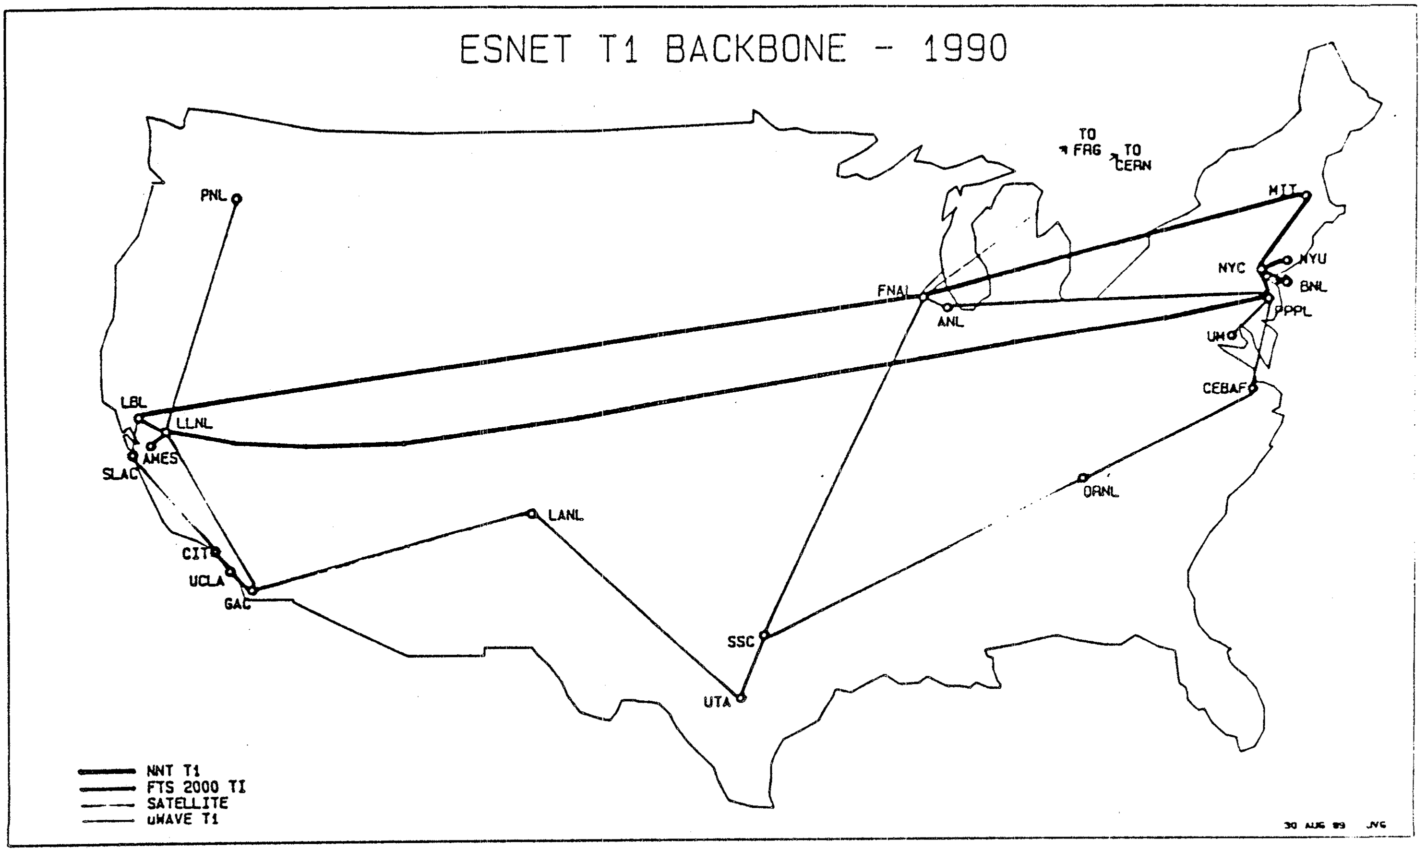

ESnet Nov. '89

ADMINISTRATIVE ISSUES:

NIS GROUP FORMED - BOB AIKEN G.L.

IP ROUTING GUIDLINES REVIEWED AUG. 89
FTP: CCC.NMFECC.GOV
[ANONYMOUS.SPECS]ESNET_IP_ROUTING.

DECNET ROUTING GUIDLINES EXPECTED JAN. 90
COPY IGW::CCC::
SYS$USER3:[ANONYMOUS.SPECS] TBD

ESnet

Nov. '89

PAST ACTIVITIES:

cisco ROUTERS RECEIVED SEPT . 89

NNT T1 CIRCUITS INSTALLED OCT . 89

INSTALLED SITES: LLNL, LBL, FNAL, ANL, PPPL,
BNL, NYU, MIT

ROUTING IP & DECNET 4

SNMP MANAGER RUNNING

ESnet

Nov. '89

PLANED ACTIVITIES:

DISCUSS ROUTING WITH SITES AND REGIONALS

INTERIM 56K SITES: ORNL, FSU, UT, GA NOV 89

MOVE CISCO TO GARCHING FRG. DEC 89

SWITCHED X25 SERVICE OVER BACKBONE JAN 90

FTS2000 T1 LINES FEB 90

3.6 "NIC Status Report"

Presentation by Mark K. Lottor/SRI International

Number of Networks: This graph shows the number of active networks (reported by BBN) plotted against the number of networks registered by the NIC (to be connected to the Internet).

Host Counts: The slide shows the number of hosts and domains on the Internet that were found by the domain survey program. This program recursively searches the domain tree counting everything it finds. Also listed is a breakdown of how many hosts had how many interfaces.

Hosts and Domains Graph: This graph plots the number of hosts and domain for the past two years using results from the domain survey program described above.

NIC Changes: The NIC will be dropping its direct ARPANET and MILNET IMP connections in the next few months and will be switching to a gatewayed configuration to allow better access via NSFNET. The changes will be announced in a future DDN Management Bulletin.

Domain Survey Statistics

Hosts	160,000
Domains	4800

Host Address counts:

0	11639
1	155523
2	2923
3	369
4	150
5	71
6	33
7	16
8	16
9	3
10	1
11	2
12	2
13	1
14	1
15	1
16	3
21	1
24	1
30	5
31	1

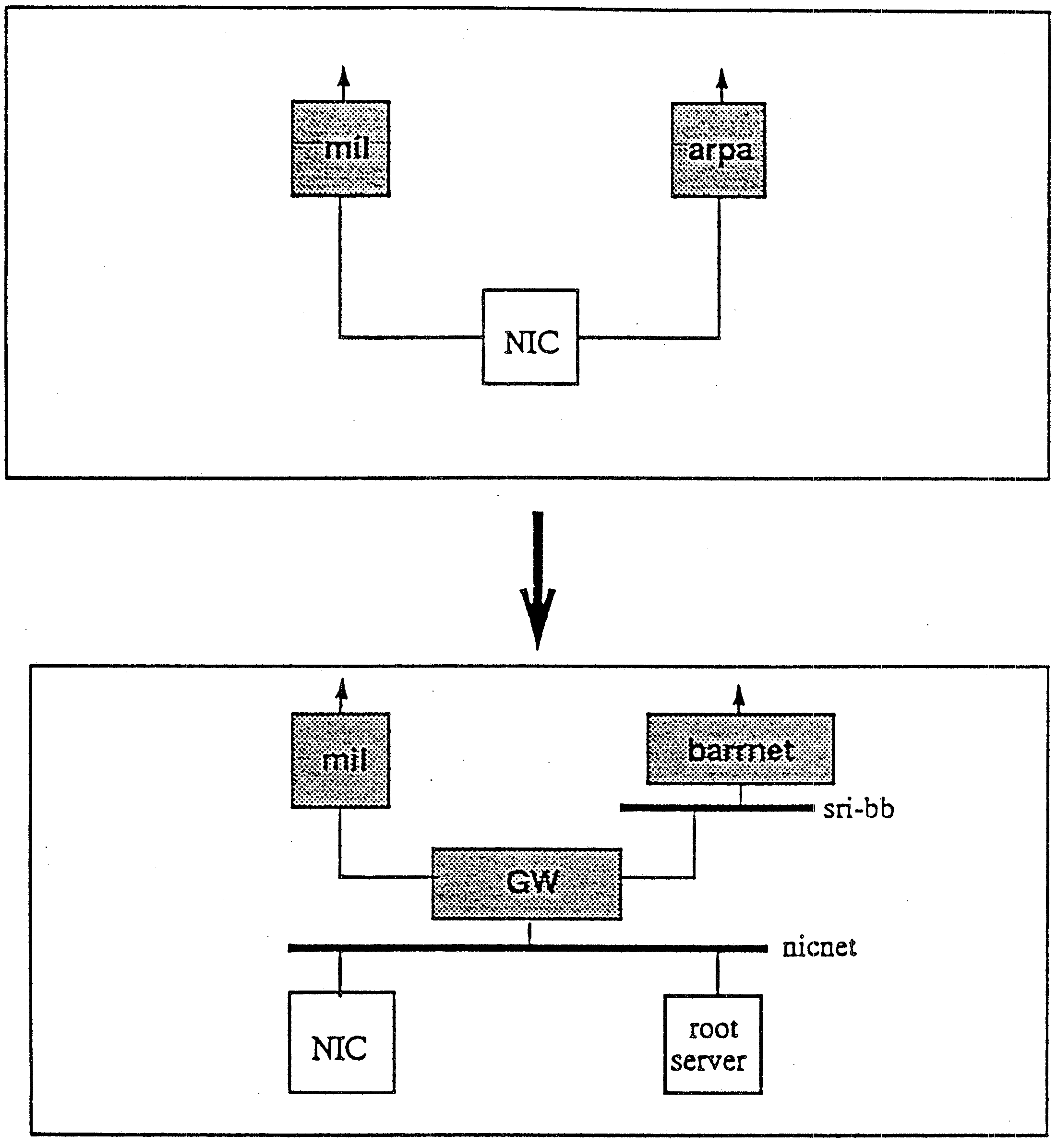
mil
arpa
NIC
mil
barrnet
sri-bb
GW
nicnet
NIC
root
server

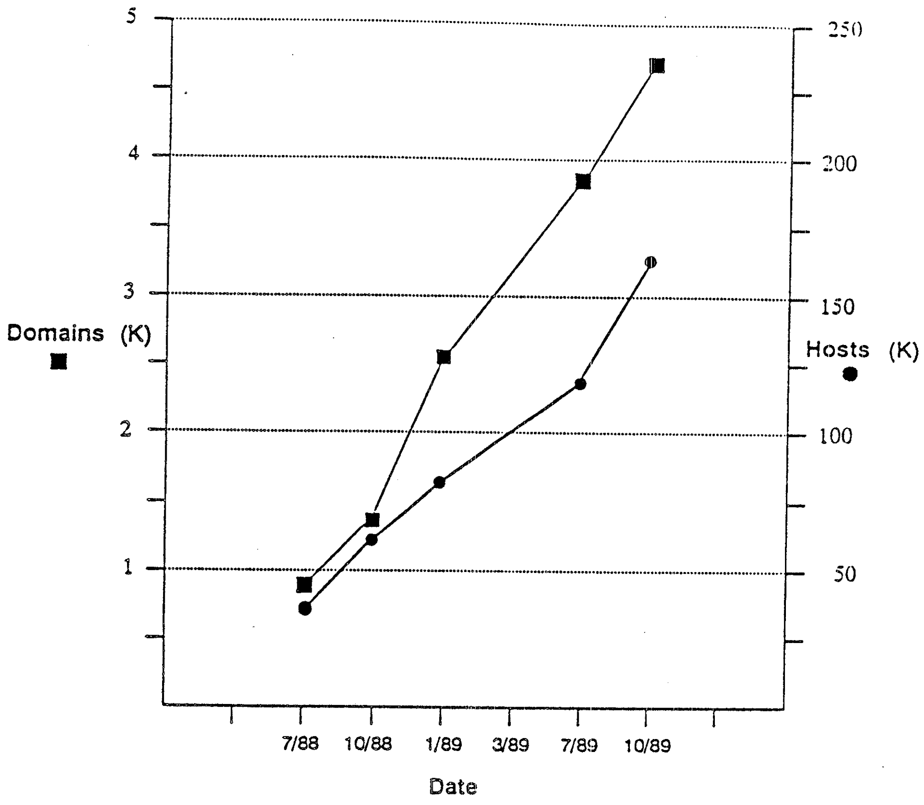

5
4
3
2
1
Domains (K)
250
200
150
100
50
Hosts (K)
7/88
10/88
1/89
3/89
7/89
10/89
Date

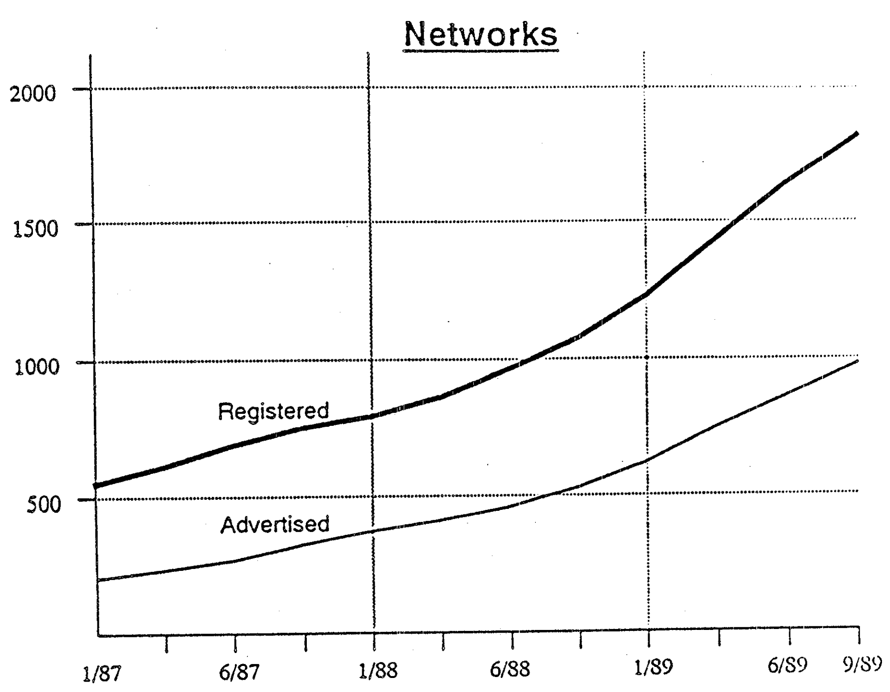

Networks
2000
1500
1000
500
Registered
Advertised
1/87
6/87
1/88
6/88
1/89
6/89
9/89

3.7 "Talking Roads and Networked Cars"

Presentation by Carl-Herbert Rokitansky/ Fern University of Hagen

Currently, there are two European-wide joint research projects to increase road-safety and road transport efficiency using advanced technologies in the fields of microelectronics, sensor engineering as well as telecommunications between vehicles and between vehicles and roadside infrastructure equipment:

PROMETHEUS is an acronym for "PROgraM for a European Traffic with Highest Efficency and Unprecedent Safety". This project has been launched by several European automobile companies. Contributions are focused on the development of computer support to the driver for his driving tasks, using sensoring systems, vehicle-to-vehicle and vehicle-to-roadside infrastructure communications, based on analysis of road traffic scenarios, artificial intelligence, a suitable network architecture and advanced communication protocols.

DRIVE is an acronym for "Dedicated Road Infrastructure for Vehicle Safety in Europe". It is a program of the Commission of the European Communities, in which the application of information technology and telecommunications to the development of Road Transport Informatics (RTI) is supported. DRIVE will contribute to the creation of an Integrated Road Transport Environment (IRTE).

In this presentation we try to give a brief overview of the system approach and of communication requirements for PROMETHEUS/DRIVE applications, and to discuss suitable routing strategies for large mobile networks, with a rapidly changing topology, as well as their performance evaluation by simulation of the developed communication protocols based on realistic dynamic networks and road traffic environments.

Talking Roads and Networked Cars:

- An object-oriented approach
- Applications and Services
- Communication Characteristics
- Multi-hop Routing Strategies
- Simulation of Communication Protocols
  Based on Realistic Mobility Models

Prof. Dr. Bernhard WALKE
Paris, September 18, 1989

RESEARCH ACTIVITIES AT THE FERN UNIVERSITY OF HAGEN:

1) Capacity Management and Supply of Physical Layer (ISMA Protocol)
2) Frame/Slot Synchronization
3) Development and Accomodation of Layer-2 Protocols (CSAP2/DCAP etc.)
4) Analysis and Simulation of Mobility Dependent Connectivity Changes
5) Performance of Multi-Hop Protocols
6) Transmit Power Control
7) Acknowledged Broadcast for Network Management
8) Decentral Topology Update
9) Direction-Oriented and Knowledge-Based Routing
10) Simulation of Communications between Vehicles Based on Realistic Mobility Models
11) Communication Characteristics
12) Specification of IRTE Services
13) Addressing (broadcast, multicast, point-to-point)
14) Gateway Algorithms and Internetting
15) Strategy for Migration to ISO/OSI Protocols
16) Experimental Packet Radio Network and Field Trials
17) Test Scenario Descriptions
18) Implementation of (Layer 2 and 3) Communication Protocols for Demonstrators
19) Methods and Tools (Simulation/SDL/ESTELLE/EFSM)
20) Specification, Verification and Validation of the Developed Protocols

CURRENT AND FUTURE WORK

ISO Layer	Activities
1	* Synchronization * Spread-Spectrum Techniques * ISMA: Capacity Management
2	* CSAP/DCAP Integration * Power Control * Handover (Channel-Switch), Dezentralized Control * Switching without Bitmap * High Capacity Channels * Spot Beam Antenna Protocols * Busy Tone Techniques * ISMA: Management
3	* Multi-Hop Routing - Direction-Oriented - Knowledge-Based - Local/Global - Type of Service Specifics - Multi-Path Routing - Source Routing - Backward Learning * Addressing Techniques - Global - Lokal (Dynamic) * Internetting/Gatewayalg. * Flow Control * Redundanz (Multi Channel Connections) * Network Management
5-7	Service Specification * Communication Characteristiques. * Basic Functions

PROMETHEUS SYSTEM

MOBILE OBJECTS

VEHICLE

- PRO-COMPUTER (PC)
- SENSORS (SE)
- DISPLAY (DI)
- ACTUATORS (AC)
- DATABASE (DB)
- COMMUNICATION (CO)
  - INTERNAL
  - EXTERNAL

PEDESTRIAN

OBSTACLE

INFRASTRUCTURE

ROADSIDE INFRASTRUCTURE

- TRAFFIC SIGN
- TRAFFIC LIGHT
- TRAFFIC INFORMATION
- COMMUNICATION EQUIPMENT
  - MAILBOX, FRAN

MANAGEMENT CENTER

- TRAFFIC MANAGEMENT
  - SURVEILLANCE CONTROL
  - GUIDANCE
- NETWORK MANAGEMENT

GATEWAY to

- EMERGENCY CENTER
  - RED CROSS, POLICE, FIRE DEP.
- INFORMATION SYSTEMS
  - WEATHER, TIME TABELS
- RESERVATION SYSTEMS
  - HOTEL, FLIGHT, TRAIN
- PUBLIC SERVICES
  - BUSSES TAXI
- OTHER NETWORKS

PRO-NET/PRO-ROAD Scenario

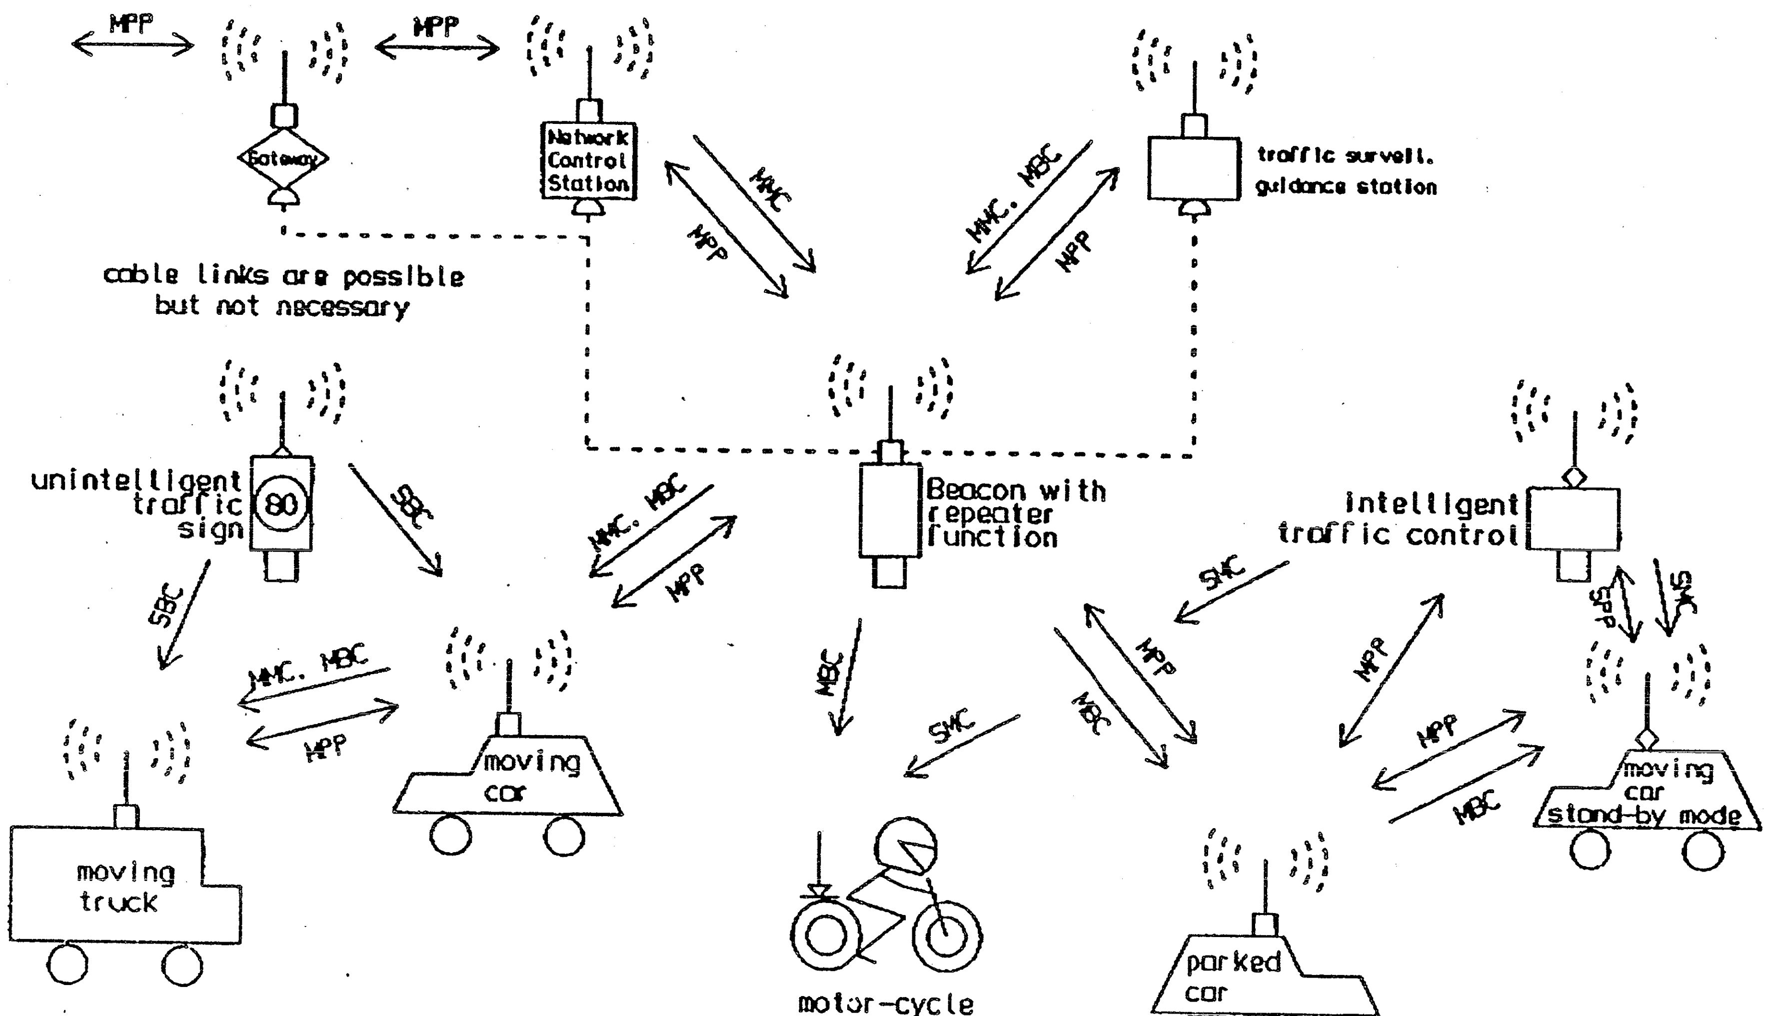

Symbols:

Mode:

NODE (store and forward)

TERMINAL (no forwarding)

GENERATOR (transmitting only)

SINK (receiving only)

Devices:

BEACON

HOST

STATION

GATEWAY

Mobility:

FIXED

QUASI-FIXED

QUASI-MOBILE

MOBILE

Messages:

SPP SINGLE-HOP POINT-TO-POINT

MPP MULTI-HOP POINT-TO-POINT

SMC SINGLE-HOP MULTICAST

MMC MULTI-HOP MULTICAST

SBC SINGLE-HOP BROADCAST

MBC MULTI-HOP BROADCAST

2. The PROMETHEUS Applications

Application Class	No.	Applications	Comments
PRO-NET	N1	Emergency Warning System (EWS)	car-to-car warning
	N2	Convoy driving	
	N3	Overtaking	
	N4	Cooperative manoeuvres	
	N5	Information exchange	
	...		
Combined PRO-NET/ PRO-ROAD	C1	Traffic coordination	
	C2	Mailbox services	
	...		
PRO-ROAD	R1	Traffic information	from traffic lights or signs,...
	R2	Traffic flow control	koordination between traffic lights,...
	R3	Traffic data aquisition	
	R4	Global route guidance	final driver destination
	R5	Local route guidance	parking places,...
	R6	Rescue request for EWS	
	R7	Automatic payment	toll roads, bridges, tunnels, parking places,...
	R8	Remote database access	tourist information, hotel booking,...
	R9	Network services	Email, File transfer, Virtual terminal,...
	...		

3.1 Basic Functions

Each Application contains a set of Basic Functions which provide essential functions like speed, position, etc. Basic Functions which are only of interest for PRO-CAR Applications are not considered here. The Functions are grouped according to their typical usage:

Service Class	No.	Basic Functions	Comments
Vehicle Characteristics	BV1	Speed	in m/sec
	BV2	Acceleration	in ± m/sec^2
	BV3	Direction	Current heading
	BV4	Position	relative, absolute, Lane No., ...
	BV5	Technical limits	max. speed, max. acceleration, kW, vehicle length, weight, ...
	BV6	Safety equipment	ABS, hazard lights, ...
	BV7	Communication equipment	PROMETHEUS, GSM, RDS, ...
	BV8	Intervehicle distance determination	sectors (front, rear , right, left, ...)
	BV9	Special status	overtaking, parked, accident
	...		
Road Characteristics	BR1	Road type determination	highway, urban road, ...
	BR2	One/two-way traffic	true, false
	BR3	Current Number of lanes	for both (all) directions
	BR4	Number of lanes ahead	for both (all) directions
	BR5	Slope	in %
	BR6	Junctions	roundabout,railway crossing,...
	BR7	Toll requirements	in ECU
	...		
Current Conditions	BC1	Weather conditions	rain, fog, visibility, ...
	BC2	Road conditions	friction coefficient (wet, icy,...), bumps, ...
	BC3	Visibility determination	ranges
	BC4	Restrictions	blocked lanes, road work, ...
	BC5	Traffic density measurement	autmatic measurement in cars/h
	BC6	Dangerous emission	leaking oil, flammable/explosive materials, dangerous gas, ...
	...		
Traffic Rules	BT1	fixed driving restrictions	no overtaking,... via traffic signs
	BT2	dynamic driving restrictions	red/green, direction, ... via traffic lights
	BT3	Common rules	right of way, ...
	...		
Intended Activities	BA1	Direction changes	turns, lane change, ...
	BA2	Intermediate destination	petrol station, restaurants, junctions, accommodation, ...
	BA3	Final travel destination	city, street name, ...
	...		

2.2.2 Characteristics and Parameters

For a unify detailed description typical characteristics will be defined for each Application and Basic Function. The following list shows the most important characteristics and their several values:

No.	Characteristics	Values
Ch1	*Zone of Relevance*	local station direct neighbors short range local area wide area (within PROMETHEUS) global (outside PROMETHEUS)
Ch2	*Information Flow*	vehicle - one/some/all nonvehicle - one/some/all
Ch3	*Lifetime*	indefinite (until explicitly canceled) relative (seconds) absolute (time/date)
Ch4	*Response Timeout*	nil relative (seconds) absolute (time/date)
Ch5	*Repetition Period (fixed/dynamic)*	periodical (seconds) event driven (depends on certain condition, e.g. passing beacon)
Ch6	*User Data Size (octetts)*	fixed (# bytes) variable (# bytes, depends on conditions)
Ch7	*Transfer Mode*	datagram stream
Ch8	*Priority*	emergency traffic/intracar control traffic/intracar information common user services
Ch9	*Acknowledgement*	required not required
Ch10	*Error Protection*	highly reliable normal simple

PRO-NET CHARACTERISTICS:

- BROADCAST and POINT-TO-POINT MESSAGES
- HIGH REPETITION RATE (→ CHANNEL ACCESS)
- SHORT LIFETIME (→ FEC)
- LIMITED ZONE OF RELEVANCE (→ SINGLE-HOP
  DOUBLE-HOP)

PRO-ROAD CHARACTERISTICS:

- POINT-TO-POINT and BROADCAST MESSAGES
- LOW REPETITION RATE
- LONG LIFETIME (→ ARQ)
- WIDE ZONE OF RELEVANCE (→ MULTI-HOP
  ALTERNATE ROUTES)

TOPOLOGY CHANGES

- ON BIDIRECTIONAL ROADS :

  TOPOLOGY CHANGES OCCUR FREQUENTLY

- ON UNIDIRECTIONAL ROADS :

  TOPOLOGY CHANGES OCCUR INFREQUENTLY

→ TO SIMPLIFY ROUTING ALGORITHMS AND FOR STABLE ROUTING

USE DIRECTION BASED ROUTING STRATEGIES

STATUS REPORT MESSAGE

TRANSMITTED PERIODICALLY BY EACH PROMETHEUS UNIT (PR), CONTAINS :

- TYPE OF MESSAGE (CONTROL MESSAGE)
- MESSAGE CLASS (= SINGLE-HOP & BROADCAST)
- TYPE OF ROUTING STRATEGY USED
- TRANSMITTING PR-ADDRESS
- STATUS (COMMUNICATION MODE, DEVICE, MOBILITY)
- LOGICAL DIRECTION
- SEQUENCE NUMBER
- DIRECT NEIGHBORS (PR-ADDRESSES)

- FIXED ROUTING ACCESS NODE (FRAN) ADDRESS
- HOPCOUNT
- NEXT PR ON ROUTE TO FRAN
- PRs REPORTING SAME FRAN AT HIGHER (OR EQUAL) HOPCOUNT

- QUASI-MOBILE ROUTING ACCESS NODE (QMRAN)
- HOPCOUNT
- NEXT PR ON ROUTE TO QMRAN
- PRs REPORTING SAME QMRAN

PNP Header (Draft)

HEADER LENGTH	PACKET LENGTH
SOURCE PRO-ID	
DESTINATION PRO-ID	
PREVIOUS PRO-ID	
TRANSMITTING PRO-ID	
RECEIVING PRO-ID	
HOP COUNT	SEQU. NUMBER
CONTROL	FLAGS

PRO-NET / PRO-ROAD SINGLE-HOP COMMUNICATION

SCENARIO: N1 (Vehicle Status Report) transmitted by 8

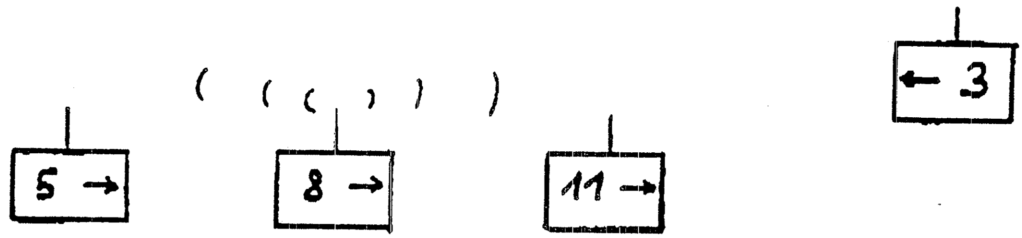

N1 Vehicle Status Report:

- Data: Speed, Acceleration, Direction, Position, Comms Equ, Status, Dir changes
- Zone of Relevance: DIRECT NEIGHBORS (→ single-hop)
- Information Flow: VEHICLE → VEHICLE: ALL
- Lifetime: 1 sec
- User Data Size: 10 bytes (assumed)
- Transfer Mode: DATAGRAM (→ UDP)
- Acknowledgement: NOT REQUIRED (→ LLC Type 1)

N1 Data

UDP: Sequence No, Checksum, etc.

UDP	N1 Data

IP: Source: 192.54.224.8
Destination: 192.54.224.255

IP	UDP	N1 Data

PNP:
Source PRO-ID: 8
Destination PRO-ID: 255
Previous PRO-ID: 8
Transmitting PRO-ID: 8
Receiving PRO-ID: 255

PNP	IP	UDP	Data

HDLC: Address: 255

HDLC	PNP	IP	UDP	Data	HDLC

PRO-NET / PRO-ROAD MULTI-HOP COMMUNICATION

SCENARIO: N4 (Overtaking) transmitted by 8

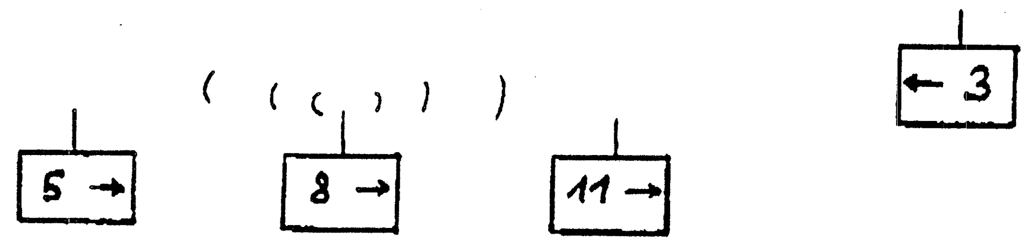

N4 Overtaking: Overtaking relevant data and Vehicle Status Report:

- Data: Speed, Acceleration, Direction, Position, Comms Equ, Status, Dir. changes
- Zone of Relevance: SHORT RANGE (→ multi-hop)
- Information Flow: VEHICLE → VEHICLE. ONE
- Lifetime: 1 sec
- User Data Size: 16 bytes (assumed)
- Transfer Mode: STREAM (→ TCP)
- Acknowledgement: REQUIRED (→ LLC Type 2)

N4 Data

TCP: Ports, Sequence No, SYN, Window, Checksum, etc.

TCP	N4 Data

IP: Source: 192.54.221.8
Destination: 192.54.221.3

IP	TCP	N4 Data

PNP:
Source PRO-ID: 8
Destination PRO-ID: 3
Previous PRO-ID: 8
Transmitting PRO-ID: 8
Receiving PRO-ID: 11

PNP	IP	TCP	N4 Data

HDLC: Address: 11

HDLC	PNP	IP	TCP	N4 Data	HDLC

MULTI-HOP ROUTING TO NEXT FIXED ROUTING ACCESS NODE

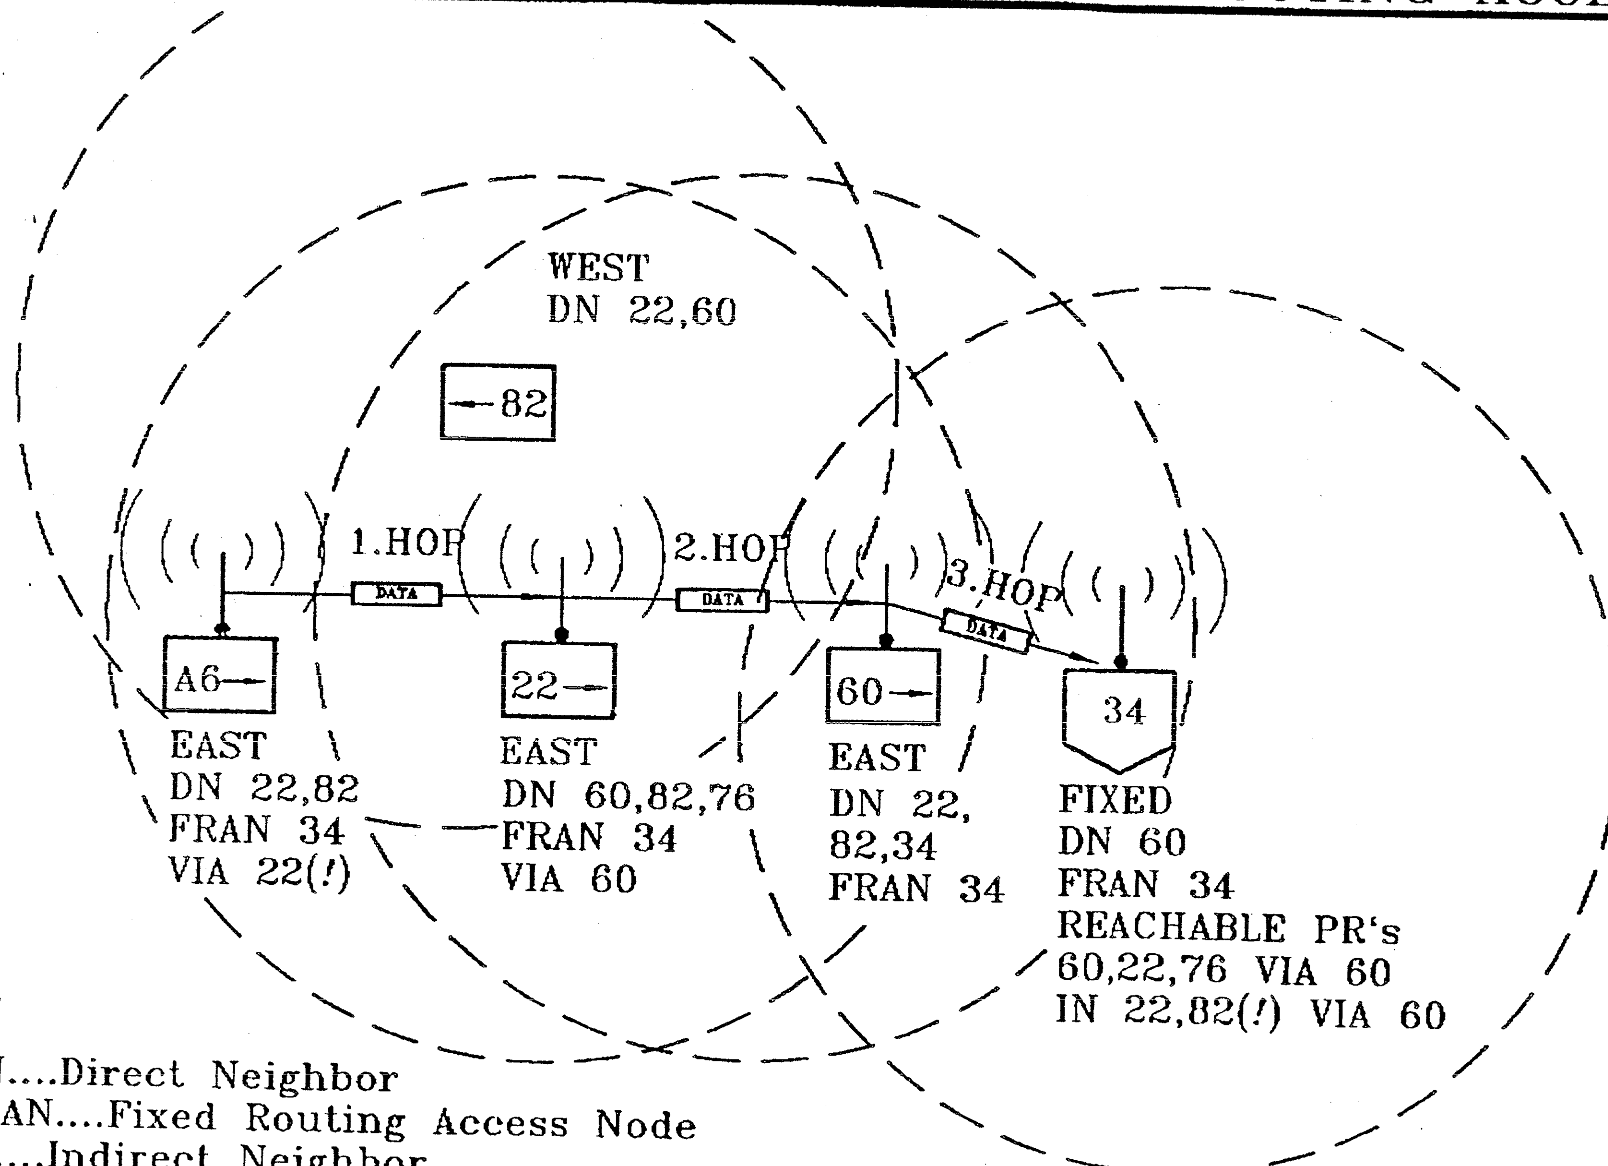

DN....Direct Neighbor
FRAN....Fixed Routing Access Node
IN....Indirect Neighbor

S imulation of
I ntelligent
M anoeuvering and
C
O mmunications

- Motivation

- Basic Simulation Model

- Simulation of Vehicle Mobility and Communications

- Simulation of PROMETHEUS Functions

- Simulation Architecture and Modules

- Modelling and Graphical Display of Selected Scenarios

- Conclusions

BASIC SIMULATION MODEL

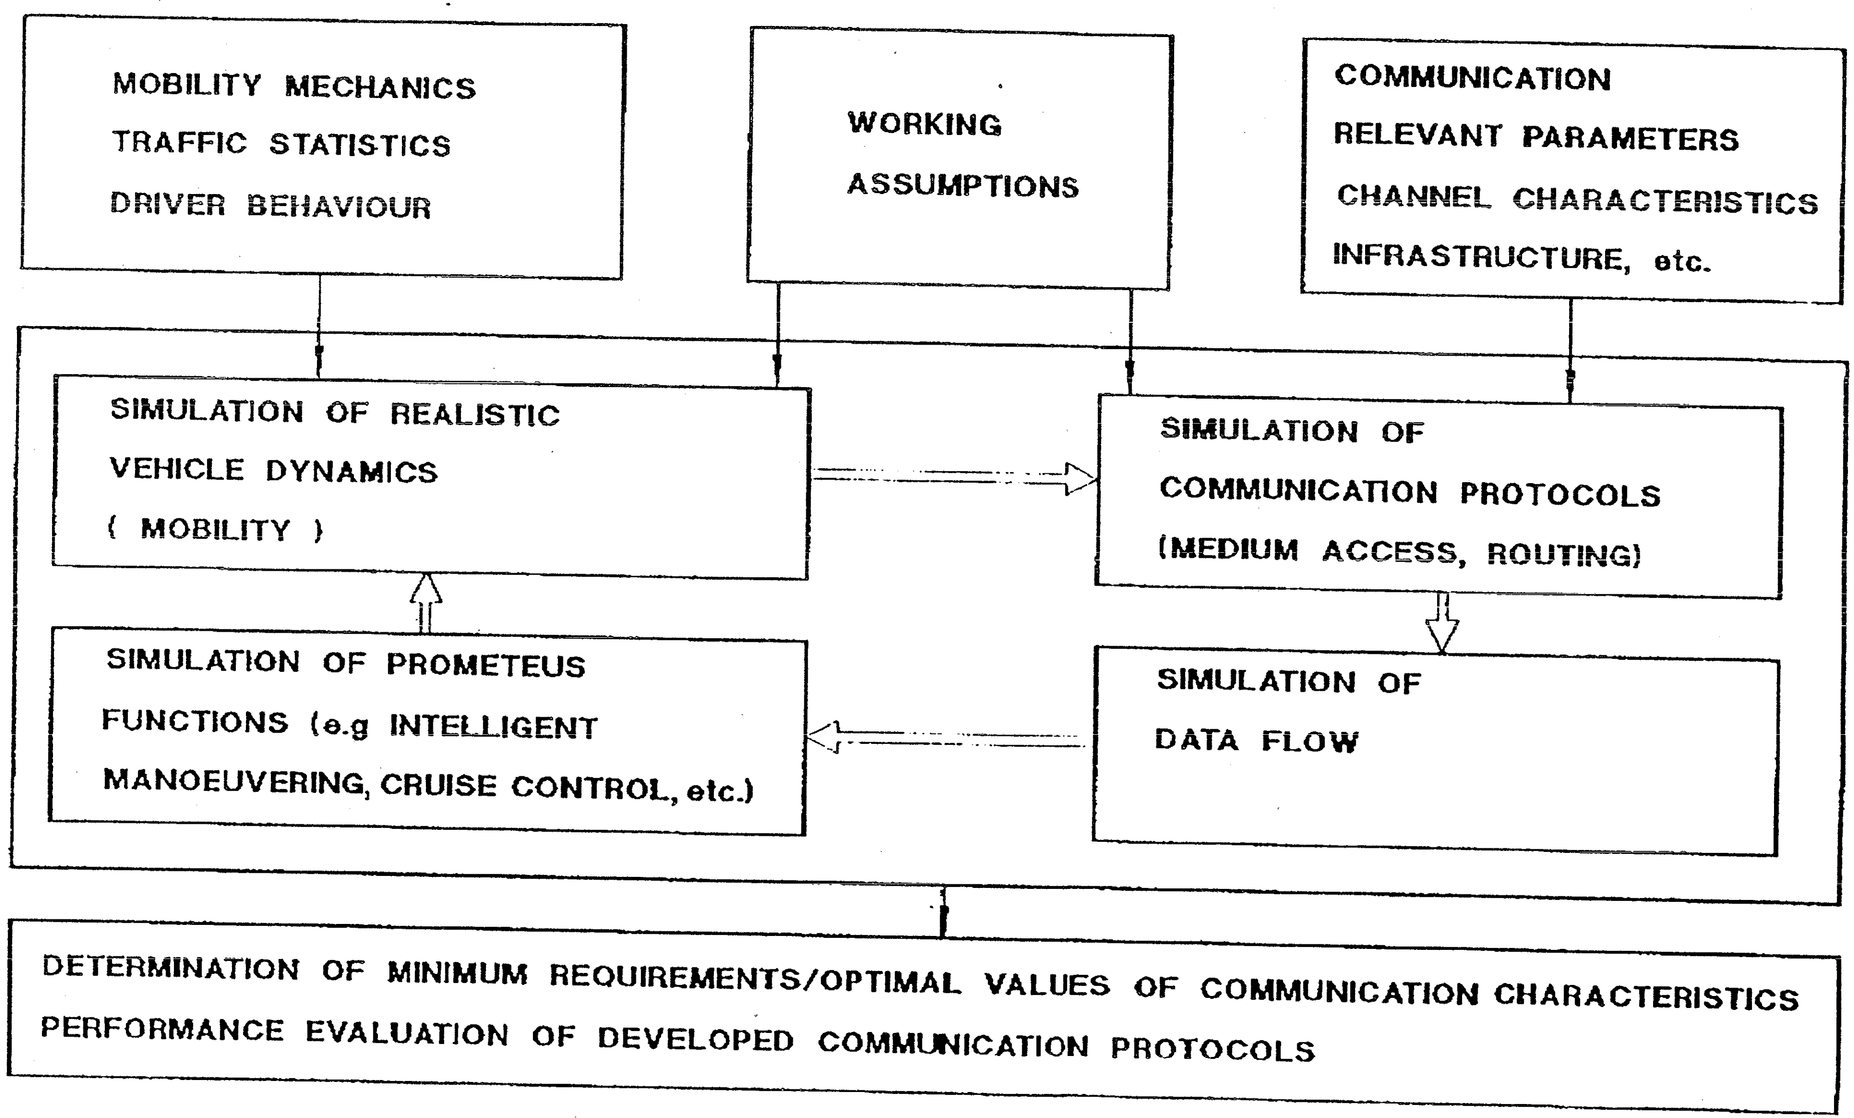

SIMCO - Simulation Scenarios

- Highways
- Rural Roads
- Intersections (controlled/not cotrolled by traffic lights)
- Road narrowing scenarios
- Access ramps (entrance/exit)

Lanes:
Up to 6 lanes (extendable); each direction.

Traffic density:
Up to 7200 vehicles per hour;
different values for each direction.

Minimum Speed: for the whole simulation range

Overtaking manoeuvers:
By specifying the leftmost lane, up to which overtaking manoeuvers are allowed, two-way road traffic scenarios with solid or dotted lines can be simulated.

Road Orientation: East/West or a North/South

Simulation Range: Up to 20.000 meters; different values for each road.

Individual Vehicle Characteristics

SIMCO maintains and updates for each vehicle:

Mobility Characteristics:
- Intended cruise speed
- Current speed
- Initial acceleration
- Current acceleration/deceleration
- Reaction time
- Vehicle length
- Position (x and y coordinates)
- Vehicle breakdown time
- PROMETHEUS equipped car (yes/no)

And for each PROMETHEUS equipped vehicle:

Communication Characteristics:
- Maximum and current transmission range
- Number and list of Direct Neighbors
- Short Range routing table
- SAME-/OPPOSITE-Direction routing table
- NON-MOBILE routing table
- Fixed Routing Access Node (FRAN) table
- Data queue (packets for transmission)

Modelling of Selected Scenarios

SIMCO uses only the following two road section characteristics to simulate the various traffic scenarios in a realistic and efficient way:

- High/low Maximum Speed Limits (for road sections with speed limits and intersections controlled by traffic lights)

- Blocked lanes (to simulate breakdown lanes, road narrowing scenarios, and access ramps)

The following traffic scenarios are discussed:
- Road sections with speed limits
- Intersections controlled by traffic lights
- Narrowing road section
- Breakdown lane / access ramp

SIMCO

PRO-COM SIMULATION BASED ON REALISTIC ENVIRONMENTS

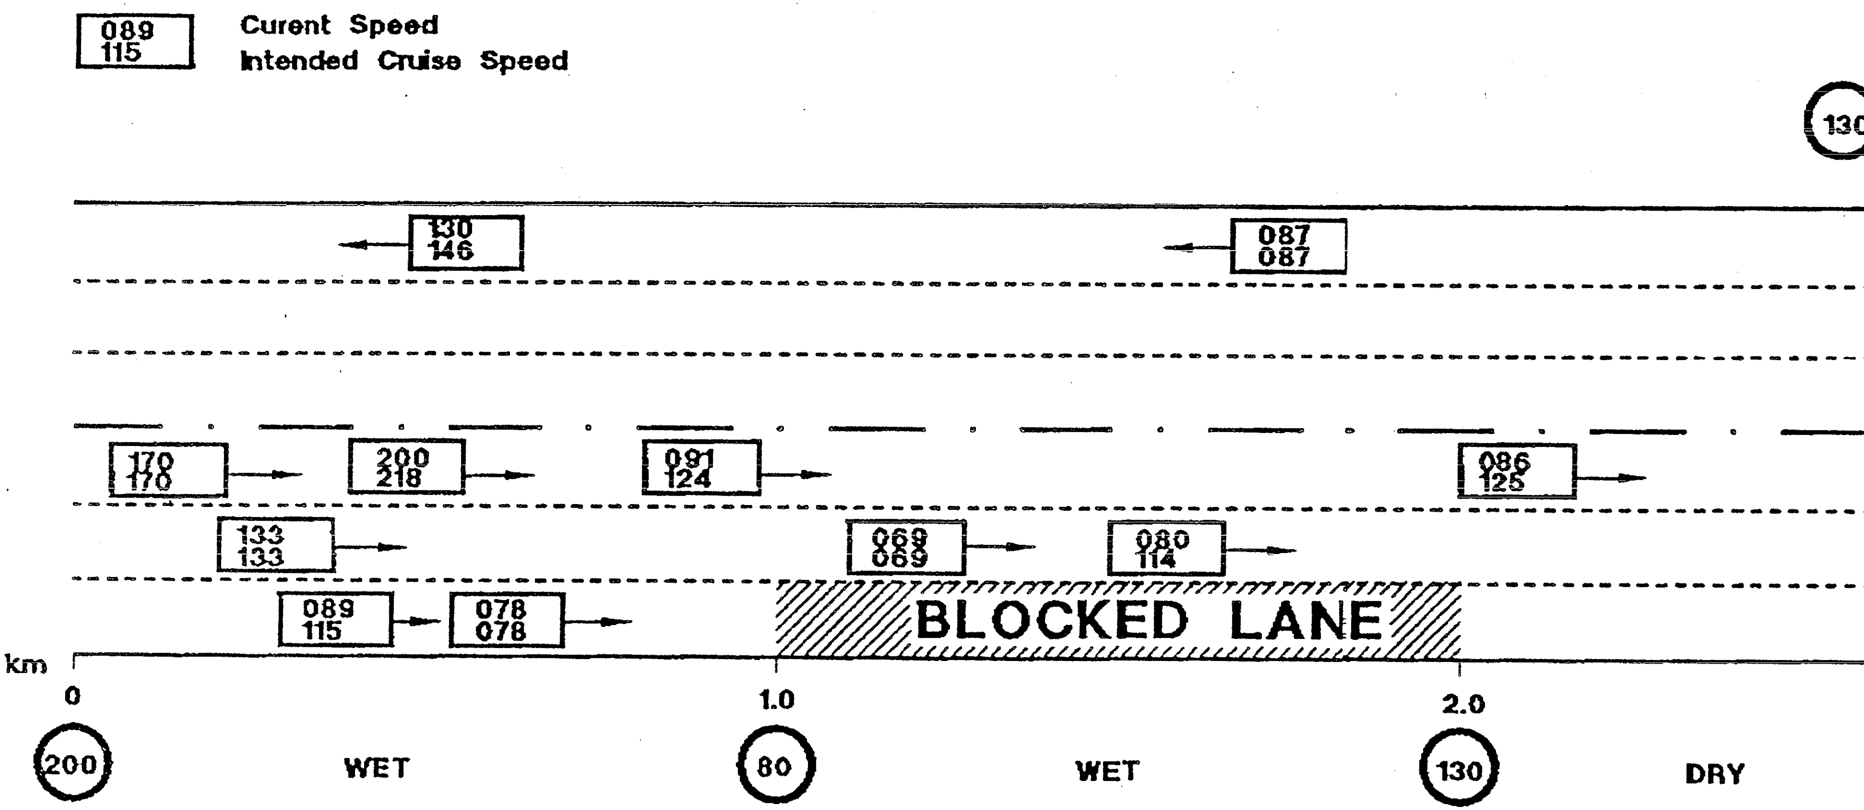

SIMCO

ROAD NARROWING SCENARIO

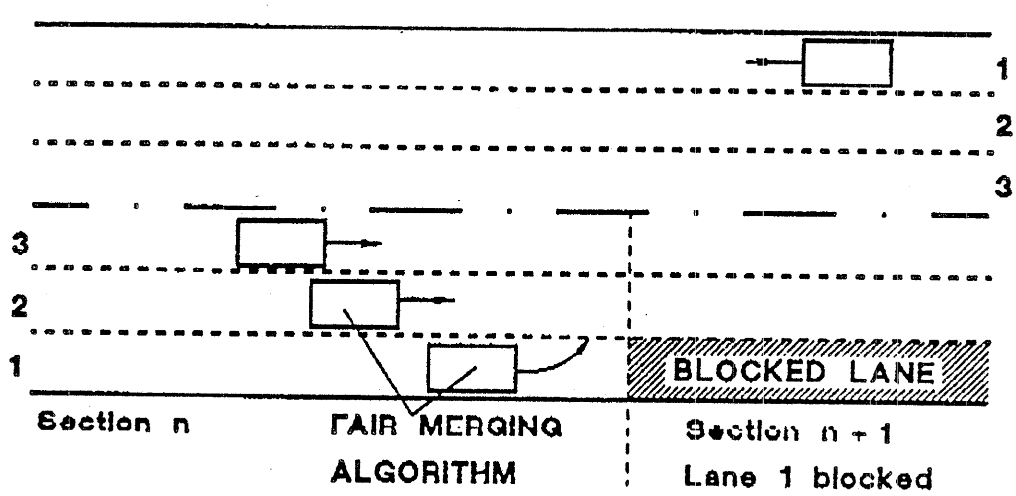

MERGING LANGE SCENARIO

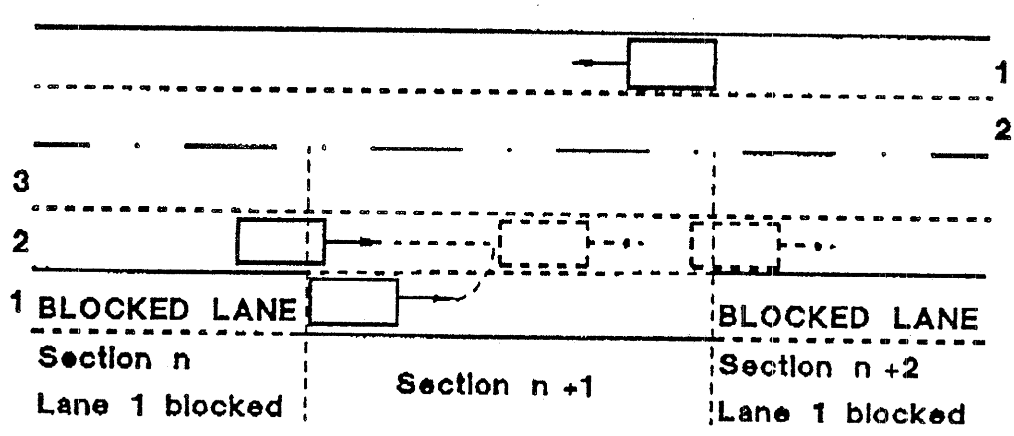

SIMCO

INTERSECTION SCENARIO

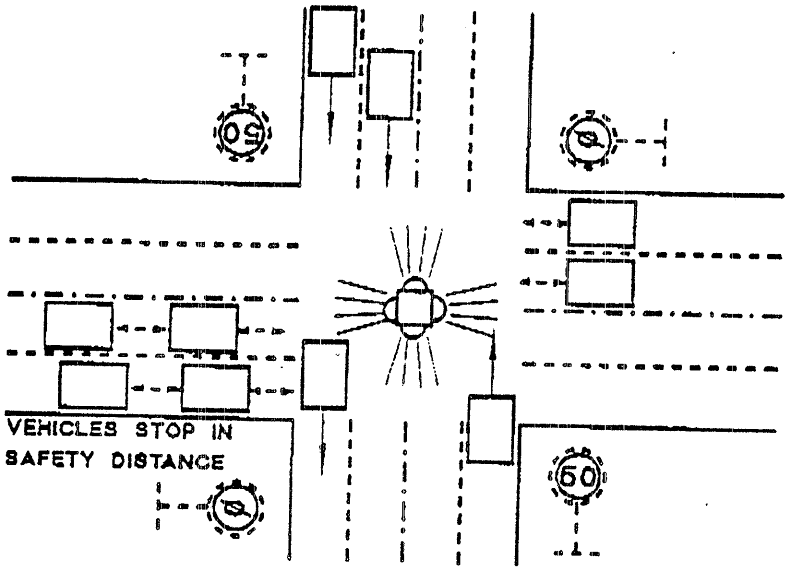

DYNAMIC
(SECTIONWISE)
CHANGES
OF SPEED
LIMITS

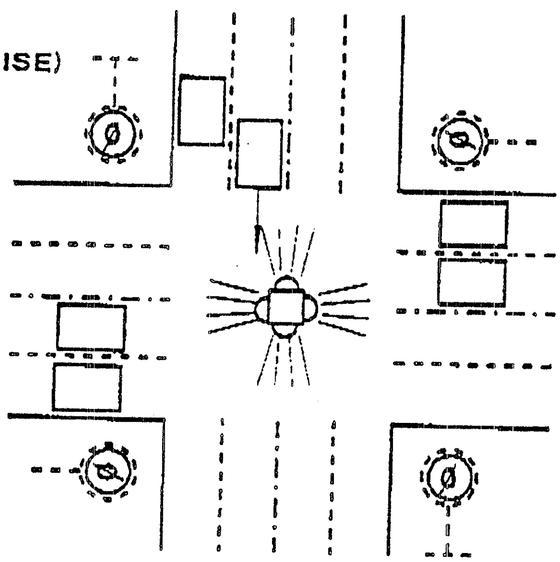

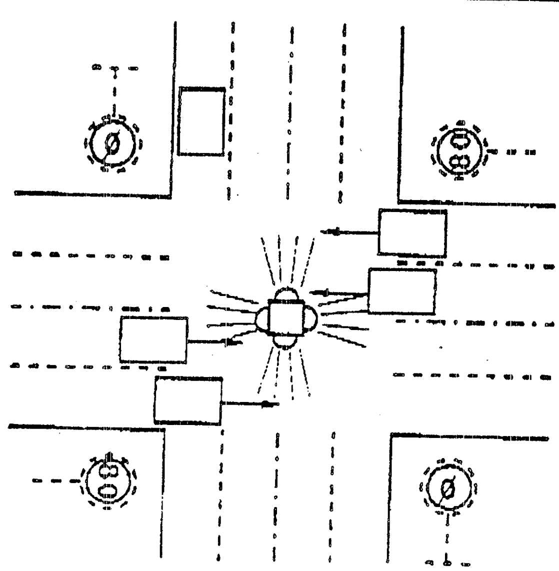

PROMETHEUS APPLICATIONS USING TCP/IP PROTOCOLS

ISO/OSI LAYERS
APPLICATION
PRESENTATION
SESSION
TRANSPORT
NETWORK
DATA LINK
PHYSICAL

<table>
<tr><th colspan="3">PROMETHEUS</th><th colspan="3">INTERNET</th><th colspan="3">ISO/OSI</th></tr>
<tr><td rowspan="2">R2
Traffic
Rules
(signs)</td><td rowspan="2">N1
Vehicle
Status
Report</td><td rowspan="2">N4
Overtaking</td><td rowspan="2">SMTP
Simple
Mail
Transfer
Protocol</td><td rowspan="2">FTP
File
Transfer
Protocol</td><td rowspan="2">TELNET
Virtual
Terminal</td><td>X.400
MHS
Message
System</td><td>FTAM
File
Transfer</td><td>VT
Virtual
Terminal</td></tr>
<tr><td colspan="3">TP0 Transport Protocol Class 0
RFC 1006</td></tr>
<tr><td colspan="2">UDP
User
Data
Protocol</td><td colspan="7">TCP
Transmission Control Protocol</td></tr>
<tr><td colspan="9">IP
Internet Protocol / Internet Control Message Protocol (KMP)</td></tr>
<tr><td colspan="9">PNP
Prometheus Network (PRO - NET / PRO - ROAD) Protocol</td></tr>
<tr><td colspan="3">IEEE 802.2
CSAP + DCAP
= CSAP II</td><td colspan="3">X.25 LAPB
S-ALOHA CSMA
BTMA single/Multislot</td><td colspan="3">I.440/441 LAPD
IEEE 802.X</td></tr>
<tr><td colspan="3">ISMA</td><td colspan="3">X.21</td><td colspan="3">IEEE 802.X</td></tr>
</table>

PROMETHEUS PROTOCOL STACK (PHASE 2)

RøKITANSKY / FERNUNIVERSITÄT DVT HAGEN

JUNI 1989

3.8 "PACIFIC RIM CONNECTIVITY"

3.8.1 "PACOM and Hawaii: Present and Future Plans"

Presented by Torben Nielson/University of Hawaii

Questions and Key Points follow:

Torben: The goal is to create an infrastructure. The reality is that funding comes from mission requirements. In so far as it is more economical to share links, infrastructure can be created.

The goal of PACOM is to encourage and use national networks to connect particular mission sites and then connect the national networks to form an internetwork.

Q: Milo Medin: Isn't it better to route to New Zealand if tariffs to Australia and New Zealand are the same? A: Yes, but there are some capacity concerns between New Zealand and Australia. If tariff costs are the same, I prefer to do rational engineering.

Q: Phill: Are you going to use OSPF? A: Yes we will cut over in January when the software is available. Q: What will the routers use to talk to each other. A: To Japan they will use PPP, the cutover will be in January or February.

Korea will come on line in January or February. The links will be direct into Hawaii, although that is not ideal. Korea should go into Japan.

Q: Who pays for the line from Japan to Australia? A: Talk to the Japanese and Australians. Q: So it is they who want it? A: No, I want it, but if they think of the connectivity from an engineering point of view they should also find it desirable. It is desirable from an engineering point of view but that does not mean that I have found someone to write the check.

There is a need for networking into Antarctica. The methodology needs to be worked out, but one method is to inscribe packets on the backs of penguins. It should be easy, they are black and white already.

There is thought of putting a satallite broadcast network for Pacific connectivity. There are a lot of ships that need cheaper connectivity than Merisat calls. We are looking at Comsat systems. There are islands that will need connectivity.

PACCOM

A PACIFIC NETWORK ARCHITECTURE

PACCOM originated in late 1987. The goals of the project were to

- Develop a sound network infrastructure in the Pacific.
- Meet agency connectivity requirements in the region.
- Encourage the development of national academic and research networks in the Pacific countries.

PACCOM currently links Australia, Japan and New Zealand. Korea is expected to be added in early 1990. Links were all put in during the Summer of 1989.

- Hawaii–NASA Ames Research Center: 512Kbps on HAW–4
- Hawaii–Australia: 56Kbps via satellite; slated for cutover to a 64Kbps ANZCAN circuit in November/December of 1989
- Hawaii–Japan: Three 64Kbps segments on TPC–3; one to Tokyo University, one to Keio University and one to the Institute for Supercomputing Research.
- Hawaii–New Zealand: 9.6Kbps on ANZCAN.

All PACCOM links are made using Proteon routers. Support is provided for:

- IP
- DECNET Phase IV
- CLNP

PACIFIC CONNECTIONS

PRESENT TOPOLOGY

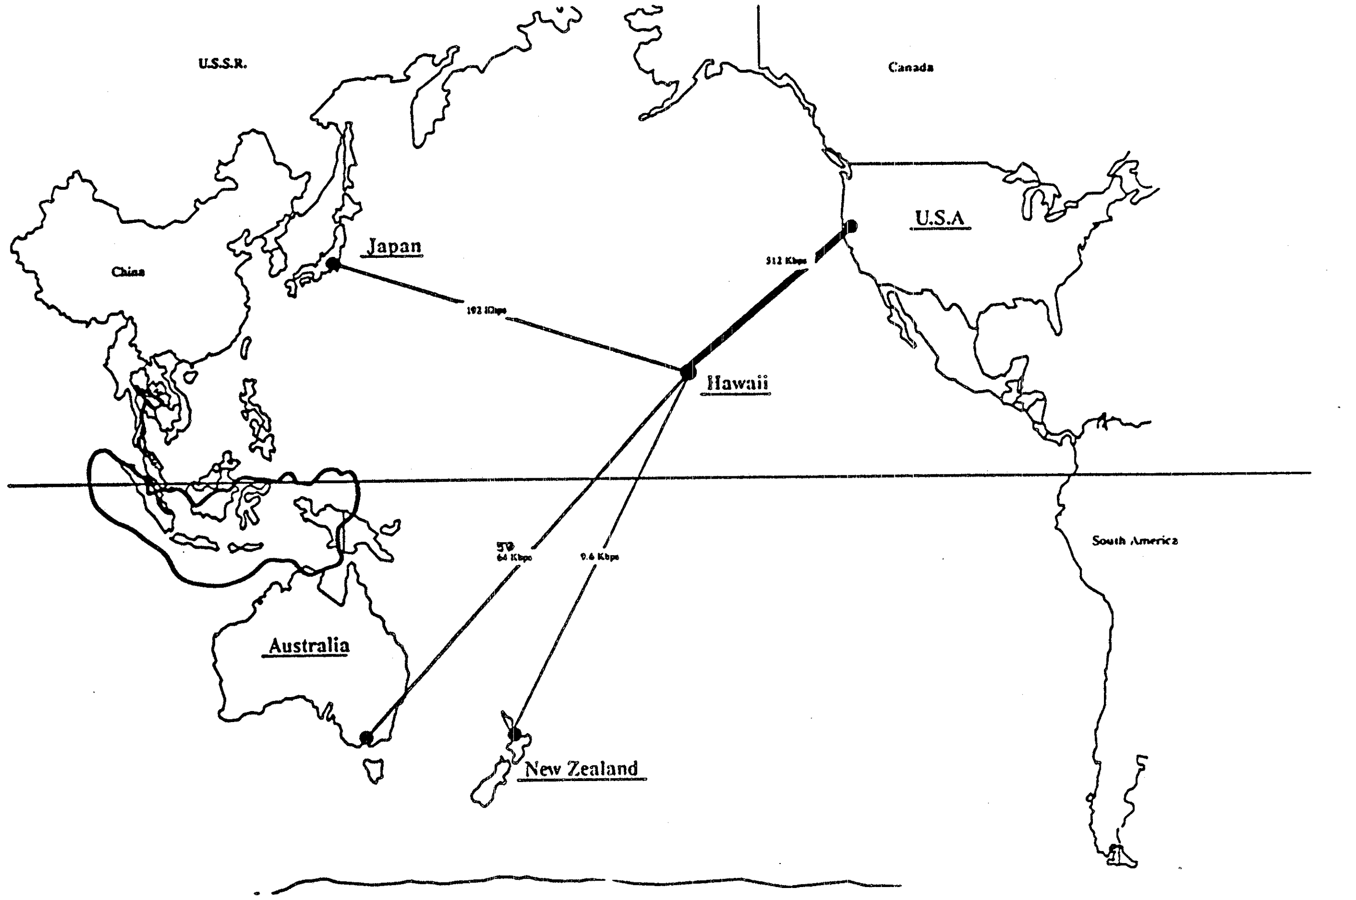

FUTURE TOPOLOGY

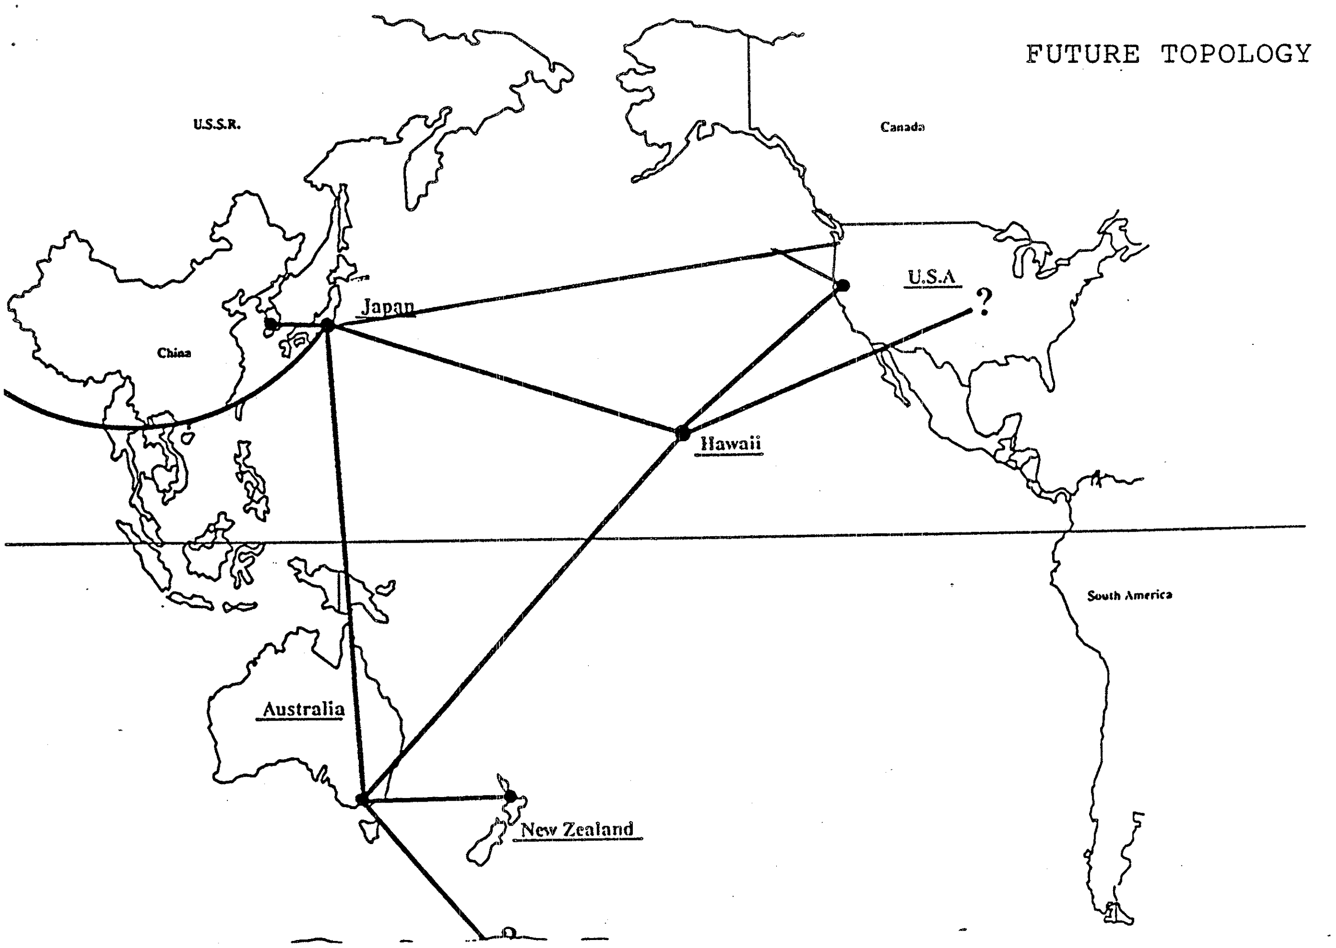

3.8.2 "Agency Requirements in the Pacific Rim"

Presentation by Milo Medin/NASA and Tony Hain/LLNL

NASA Requirements in the Pacific Rim

Milo S. Medin

Sterling Software Corporation
NASA Science Internet Project Office
NASA Ames Research Center

NASA Requirements

- Japan
- Australia
- New Zealand
- Hawaii

- 3 -

Japan

- ISAS - Tokyo
- Kyoto Univ - ~~Tokyo~~ Kyoto
- Univ. of Tokyo - Tokyo

- 4 -

Australia

- Bureau of Meteorology - Melbourne
- Australian Oceanographic Center - Sydney
- Anglo-Australian Obs. - Coonabarrabran
- Mt. Stromlo Obs. - Canberra
- CSIRO Radio Physics - Sydney, Parkes

- 5 -

New Zealand

- KAO Base - Christchurch

- 6 -

Hawaii

- Univ. of Hawaii - Manoa
- VLBI tracking station - Kauai
- Mauna Kea Obs. - Hawaii
- Meese Solar Obs. - Maui

ESNET

JAPANESE SITES UNDER CONSIDERATION

Kyoto Univ, Uji, Kyoto

IPP, Nagoya

JAERI, Naka

JAERI, Tokai

KEK National Laboratory for High Energy Physics,Tsukuba

National Institute of Genetics, Misima, Sizuoka-Ken

3.8.3 "The Australian Academic and Research Network"

Presentation by Geoff Huston/ AARN

The following pages reproduce an AARNet brochure which summarizes Geoff Huston's presentation. His presentation slides follow the brochure.

AARNet

The Australian Academic and Research Network

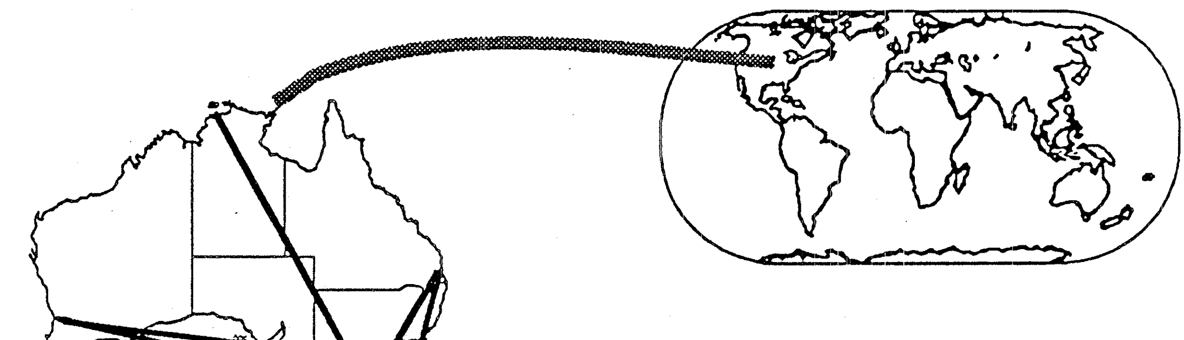

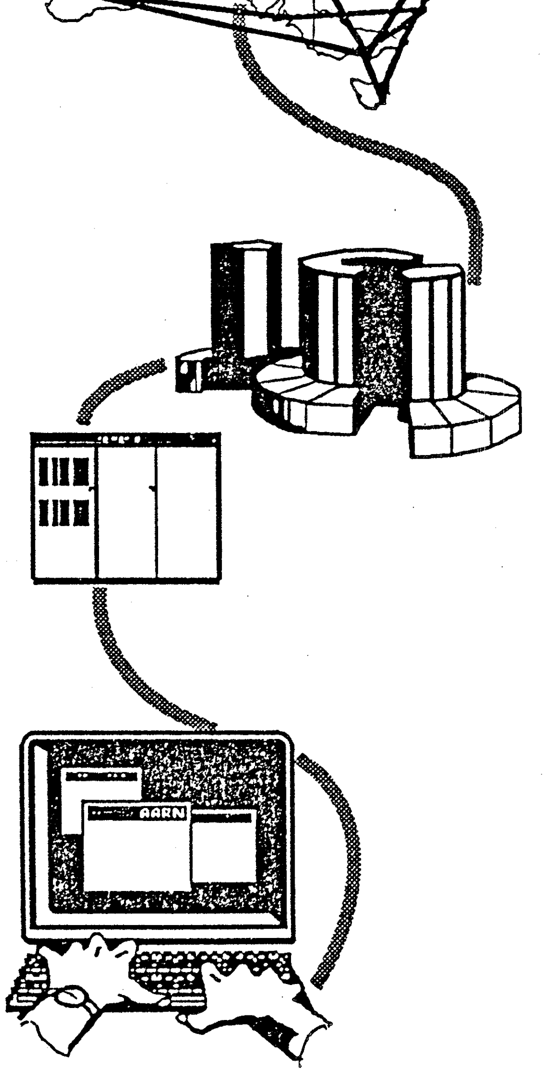

The implementation of a National research and academic computer network has been a matter of growing importance to Australian higher educational institutions and research bodies for some years. Over this period Australia has seen similar projects established and intensively used in peer nations, and the present lack of an Australian national network infrastructure for academic and research use is very much an anomalous situation.

The Australian Vice-Chancellors Committee (AVCC) and the Australian Committee of Directors and Principals Limited (ACDP) have moved to establish a high performance computer communications network within Australia. This national facility is to be constructed through the interconnection of Institutional Local Area Networks to form a set of Regional networks, and the interfacing of these Regional networks to a National network. This National network will also include interfaces to peer international research networks.

The Commonwealth Government has recently given a commitment to fund the establishment of the National and International components of this network for 1990 and 1991. In view of the importance to Australia's future research development the Commonwealth Scientific and Industrial Research Organisation (CSIRO) has also agreed to contribute to the development costs of the network.

The benefits of this program to Australia's research efforts lie in the fostering of national and international collaborative research efforts through a common communications service and in the direct impact on the productivity and effectiveness of research and academic activities by providing access to national and international communications facilities on the desktop. In doing so the program can provide a research communications infrastructure within Australia comparable to peer nations, with the consequent strategic benefits to the nation that a competitive and productive research environment can provide.

AARNet Objectives

The objective of the Australian Academic and Research Network is the provision of a high performance communications network to the members of the Australian academic and research community.

The activities that are proposed to be supported on this network include:

- Fostering collaborative activity through a common and effective communications medium. This encompasses the ability to exchange information, software and computer data between users of the network, enabling the support of geographically dispersed research groups with a common focus of activity.

- Support for a fast and reliable electronic mail delivery system for effective peer communication.

- The ability to access information sources, through either direct remote interactive access or through distributed database applications. This includes access to discipline-specific information sources, access to library catalogues, and similar. This activity can also facilitate the rapid dissemination of research methods and results throughout the research community.

- The ability to use local workstations to access remote high performance computing facilities in a productive manner across the network.

- Support of International collaboration, which is playing an increasingly important role within all areas of research activity. The network design includes the objective of enabling access to overseas networks using protocols and tools which provide direct access to overseas facilities.

The result of such a network is the creation of a distributed computing environment where each computing system or workstation can be used within a local, national and global networking environment to access other users or remote resources, and to provide the ability to publish local resources, information, software or data for access by other network members.

From the academic perspective the network's objectives are to construct the basic infrastructure for services and applications which can address many of the current and anticipated communications requirements of higher education institutions. This will include aspects of networking support for distance education programs and support for tertiary administration activities. Future areas of service provision may include the addition of audio and video services in the next generation of the data network.

Network Design

The AARNet network design methodology has been to nominate an architectural approach to address the issues involved in the construction and maintenance of a national network. The goals of this architecture are:

- To use existing networking technology;
- To use the expertise existing within the network member sites;
- Be readily implementable;
- Have architectural simplicity and uniformity;
- Be compatible with existing Australian and International networks;
- Allow for evolution of technology.

These goals can be achieved by the approach of using each institution's Local Area Network (LAN) as the basic connection unit of the network, and constructing a National Wide Area Network (WAN) to interconnect these LANs. Data packets originating within one LAN are taken by the WAN and delivered to the destination LAN.

This architecture effectively places a compatible superstructure of a national network above each local network, allowing networking services already used in each site to be extended to the national domain without significant alteration. The approach effectively utilizes one of the major assets which have been developed within each institution: the expertise in constructing uniform network services interconnecting a diverse computing environment.

Network Connection Policy

All higher education institutions and CSIRO have accepted invitations to participate in AARNet, and wil form a port of the network within the first phase of the implementation program.

In keeping with the aim of establishing a national research facility it is intended that Government research laboratories and similar instrumentalities will be able to connect to the network for the specific purposes of support of their research activity.

It is also intended that other bodies, including commercial and industrial research bodies, will be able to connect to the network on the terms that such a connection explicitly excludes all activity of a direct commercial nature from the network, and that the connection is in direct support of research activity of common interest to a higher education institution or CSIRO.

In the international domain AARNet will play an active role in establishing connections to overseas research and academic networks in order to provide Australian researchers and scholars a productive gateway to global computing and communications resources.

AARNet

The Australian Academic and Research Network

AARNet Implementation Program

The objective of the implementation program is to rapidly establish a national network which provides sufficent capacity to meet the requirements of the user community.

The implementation strategy is one of immediate implementation of a medium capacity network, and the phased introduction of higher capacities into the network over the next three years. It is anticipated that it will take some months following the initial installation of network connectivity for usage levels to generate intense network loads, and this staged introduction of additional network capacity is intended to pace this increasing load. This approach will also allow the network to take advantage of the high performance communications services being introduced by the national telecommunications carriers over the next few years.

The network will comprise a National Backbone network and eight State Regional networks. The National Backbone network will interconnect the Regional networks and also include interfaces to peer international networks. The Regional networks will provide connection services to each site.

The first implementation phase of the National Backbone network uses high capacity trunk links between the Regional networks of New South Wales, the Australian Capital Territory and Victoria. These high capacity links are capable of data transfers at a rate of 2 million bits per second. Other Regional networks will be served in the first instance by mid-speed links in a radial configuration from a National hub. These mid-speed links are capable of data transfers at a rate of 50 thousand bits per second. This phase will also fund a link to the research networks of the United States (which in turn have high capacity connectivity to European networks).

Subsequent phases of the National Backbone will extend additional capacity to the Regional networks of Queensland, South Australia, Western Australia and Tasmania, and will provide additional trunk circuits for the enhanced reliability of the backbone network. These phases also include additional capacity to the U.S. in response to anticipated usage levels, and a link to New Zealand, intended to rationalize the costs and capacity requirements of international links in the South Pacific region.

The first phase of the Regional networks use medium speed links radiating out from Regional hubs to each connected site. Subsequent phases of the Regional networks will include the installation of additional capacity. The exact program within each phase will be determined by the usage levels of the network links and the available funding.

As with local networks, the AARNet is a multi-protocol wide-area network. The network design includes the support for a number of different network protocols to coexist within the infrastructure of a single set of physical communications links. The initial phase of the network will support three protocol stacks; ISO OSI, TCP/IP and DECnet™ Phase IV. Attention will also be given to the appropriate mechanisms to support access into the international Packet Switched Networks using the X.25 interface protocol once the initial phase of the network has been set into production. Access to resources over other protocols, such as SNA, will be supported by the use of application-based protocol gateways. In the longer term it is intended that the network migrate to support of the standard ISO OSI network protocols.

AARNet Network Services

The services to be provided by AARNet are an extension of thos services already available on each institution's Local Area Net work (LAN). Such network applications allow users to exchang software and data, access remote systems, submit jobs for remot execution, and link network resources such as printers, disks, an processors to the local host system. The applications availabl over AARNet will be no different in many respects - the limita tions of the wide area network are related to capacity and traffi transit times, so that some network intensive applications, such a disk sharing, are not viable. However many of the networ applications in local use on a LAN will run unaltered acros AARNet.

The following is a brief list of some of the applications which wil initially be supported over AARNet, and the ways in which thes applications are commonly used:

- Electronic Mail; allowing the interchange of ideas, informa tion and resources between users, is one of the most commo network services. AARNet will enable efficient mail deliver by allowing the direct exchange of messages between mai systems, and, by using mail gateways, will interconnect th different mail systems in common use.
- Remote Access; allowing a network user to establish an inter active session on a remote host. This facility is used for a wid variety of purposes, including access to library catalogues databases and other information sources, as well as access t shared computing facilities.
- News; as well as mail delivery, the other major component o messaging networks is the exchange of public messages, cre ating a network-wide bulletin board. The USENET news net work in the U.S. currently delivers 3,000 messages per day t a global readership of well over 1,000,000 users. NEWS is use to distribute software, provide technical assistance on a pee basis, and to allow the interchange of ideas on a wide range o subjects of particular relevance to the research and academi community.
- File Transfer Facilities; providing a mechanism for the rapid distribution of documents, software, and data, and are also a integral component of the support for the productive use o remote computing facilities.

There are many other network applications, including network job submission, directory services, electronic document exchange, distributed databases, distributed filing systems and such. As with the trend within local sites towards a distributed computing model as a more productive and cost effective computing strategy than a single central computing resource, the academic and research community is now in a position to take advantage of the significant opportunities to productively utilize a national and global distributed computing environment.

Planned Activities for AARNet

- September - October 1989
  Information gathering process for supply of equipment and services for the network. Request for Proposals for the supply of equipment and services to AARNet released.

- November - December 1989
  Preparation for the implementation of phase 1 of both the National and Regional networks, including determination of final configurations, purchase and pre-delivery activities. Determination of National and Regional operational and management structures.

- January - April 1990
  Installation of the network, including the post-installation commissioning of equipment and connection services.

- April 1990
  Commissioning of AARNet into production.

- May - August 1990
  Integration of existing network services into the network, including ACSnet services, interfaces to Austpac-connected services and international network links.

- September - December 1990
  Pre-implementation activity for phase 2 of Network implementation strategy.

- January - June 1991
  Implementation of phase 2 of the network, and integration of these services within the overall network structure.

- July - September 1991
  Review of network progress and development of proposals for further development to be conducted. Planning of implementation of phase 3 of the network.

- October - December 1991
  Presentation of review and consideration of plans for the 1992 - 1995 period by AARNet parent bodies and their members.

AARNet Bodies

There are a number of committees and structures which have been set up to perform much of the planning activity to date. These are:

- The AARNet Steering Committee, chaired by Professor K. McKinnon, Vice-Chancellor, The University of Wollongong. The committee comprises representatives of higher education institutions and CSIRO. The brief of this committee is to provide overall direction and policy determination during this establishment phase of the project.

- The AARNet Technical Working Party, chaired by Dr R. Erskine, Director, Computing Services, The Australian National University. This working party includes network managers from higher educational institutions and CSIRO representation. The brief of this committee is to provide advice on the appropriate technologies to use within the design of the network.

- Regional Network Groups within each State. These groups include network managers drawn from all higher education institutions and CSIRO divisions within the State. These groups are to provide specific definition of the design of Regional networks, and also to provide the framework for the subsequent Regional operational and management structure.

For More Information...

Information on the Australian Acdemic and Research Network is available from a number of sources.

The Australian electronic news network is used to publish information regarding the AARNet, and is also an appropriate method for placing queries about any aspect of the network. The newsgroup for information and discussion about the AARNet is ***aus.comms***. Your Computing Services Section can provide you with assistance as to how to access this newsgroup. The electronic news network also regularly contains information regarding overseas networks - again your local Computer Services Section can be of assistance in accessing this information.

If you are a member of a higher education institution, then your local Network Manager will be able to provide answers to many initial queries, and will also know where to forward those queries which cannot be answered directly. The local Network Manager will also be able to provide details as to how national and international resources and facilities will be accessed from within your local network when the AARNet is commissioned into operation.

Within CSIRO, the contact point on technical matters is the CSIRO representative on the AARN Technical Working Party, Dr. Trevor Hales of the Division of Information Technology. On all other matters contact Mr Greg Bachelor, CSIRO Management Information Systems, or the AARNet Network Technical Manager.

All other enquires should be directed to the AARNet Network Technical Manager, Mr Geoff Huston. Queries may be sent via electronic mail to *gih900@csc.anu.oz.au*, phone (062) 493385, or fax (062) 473425. The postal address is: Network Technical Manager, AVCC, G.P.O. Box 1142, Canberra, ACT 2601.

Australian Academic and Research Network

Australian Vice-Chancellors Committee

Australian Committee of Directors and Principals

Commonwealth Scientific and Industrial Research Organisation

The objective of the network is to establish an Australian computer communications network to provide common set of networking services to member sites.

Network membership to include:

- Universities
- Colleges
- CSIRO Divisions
- Government Research Organisations
- Other research and commercial organisations with common interests to the core member bodies of AVCC, ACDP and CSIRO

AARNet Network Objectives

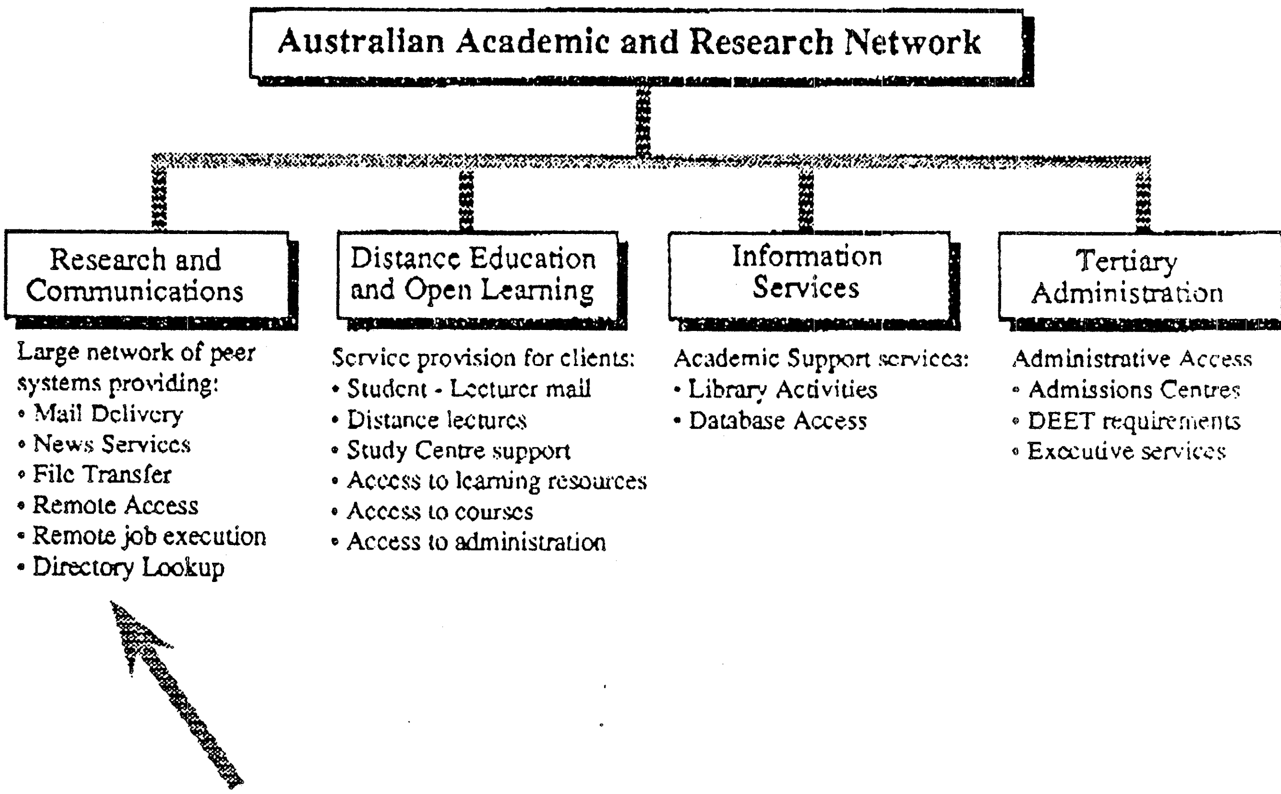

Phase 1 AARN network

Phase 1 network infrastructure to be implemented by April 1990.

Phase 1 services to be implementedthrough 1990.

The immediate objectives are concerned with the delivery of data communications services using readily available technology and available networking applications.

Longer term objectives concern the provision of an infrastructural facility in this country whichis intended to assist in the national academic and research endeavour by providing accessible and cost-effective tools for cooperative effort on a national scale.

AARNet Activities - 1989

- Refinement of 1988 Network Workshop objectives by the AARNet Technical Committee:

  - Use of multi-protocol routers as a cost-effective alternative to interconnected Remote Bridge and dedicated protocol routers

  - Private X.25 network deferred - X.25 gateway services to the public PSDN to be implemented as an alternative

  - Initial protocols to be supported to be TCP/IP, DECnet Phase IV and ISO OSI CLNS.

- Consultation with networking groups in each State to determine connectivity and equipment requirements for each site.

- Preparation of initial AARNet budget for the period to 1991.

- Presentation of AARNet proposal to AVCC, ACDP and CSIRO mid 1989 for endorsement to proceed and commitment to underwrite AARNet expenditure.

- Submission to Australian Research Counci for funding under the research infrastructure support program.

- Drafting of AARNet implementationschedules for first quarter 1990

AARNet - Phase 1

Phase 1 program encompasses:

- Installation of mid-speed (48K) and high speed (2Mbps) leased lines plus routers to form the trunk backbone.
- Installation of 48K leased lines for the tail-end links from each backbone interface to every member site.
- Support for national DECnet and national IPnet using multi-protocol routing units for both the backbone and the tail-end sites
- Support for ISO OSI (clns) anticipated mid 1990 - plus migration to DECnet Phase V routing
- Support for X.25 gateway access anticipated in late 1990

Service Issues

- Integration of ACSnet services with IP facilities
- IP management issues
- DECnet managment for Australian DECnet Phase IV
- Directory services
- Network management facilities
- IP / DECnet application gateways
- Introduction of ISO OSI support
- Gateways to other services

AARNet BACKBONE Network

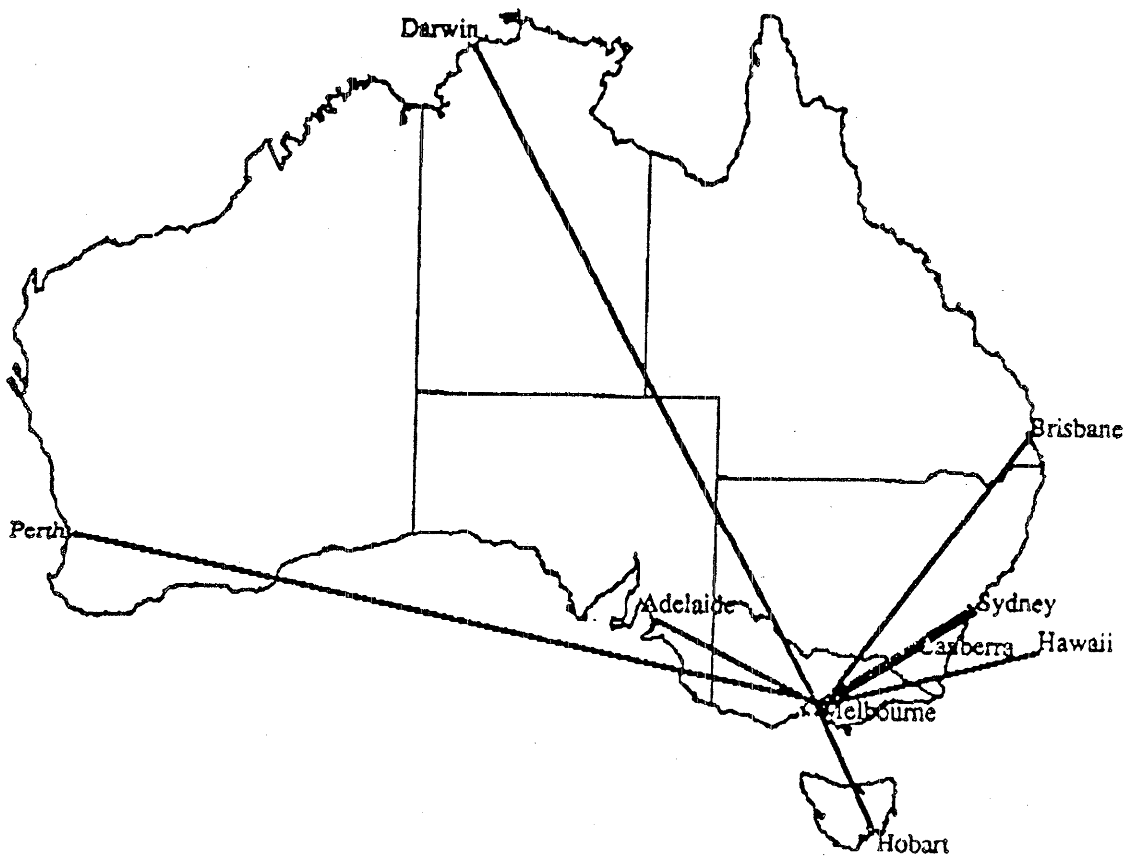

National Research Network Backbone

Phase 1 Topology

Backbone links will use midand high speed point-to-point digital lines for Phase 1: Telecom 48K Digital Data Service and Telecom 2Mbps Megalink for high speed lines)

Backbone links are terminated at Regional Network Centres (hosted by a University within each State Capital)

Major mid-speed trunk lines will be upgraded to 2Mb leased lines. Additional ISDN channels will also be configured as required to form a mesh topology for increased reliability as well as increased bandwidth and performance.

AARNet Multi-Protocol Routing Equipment

cisco routing equipment to be used:

- Supports routing functionality for TCP/IP, DECnet Phase IV, ISO OSI
- May be configured as an Ethernet Bridge on a per-protocol basis
- Support X.25 transit capabilities
- High routing performance (> 12,000 pps)
- Supports interfaces for:
  - Ethernet
  - Serial lines (up to 4Mbps)
  - Token Ring
- Extensive control parameters, including security filters
- Direct peer network access into the U.S. Internet and connected networks over TCP/IP.

The router decision effectively splits considerations of supported protocols and supported data transmissions technologies, allowing migration into new data technologies (such as Fastpac) without the requirement for extensive re-equipment of the network.

Futures for the AARN Backbone

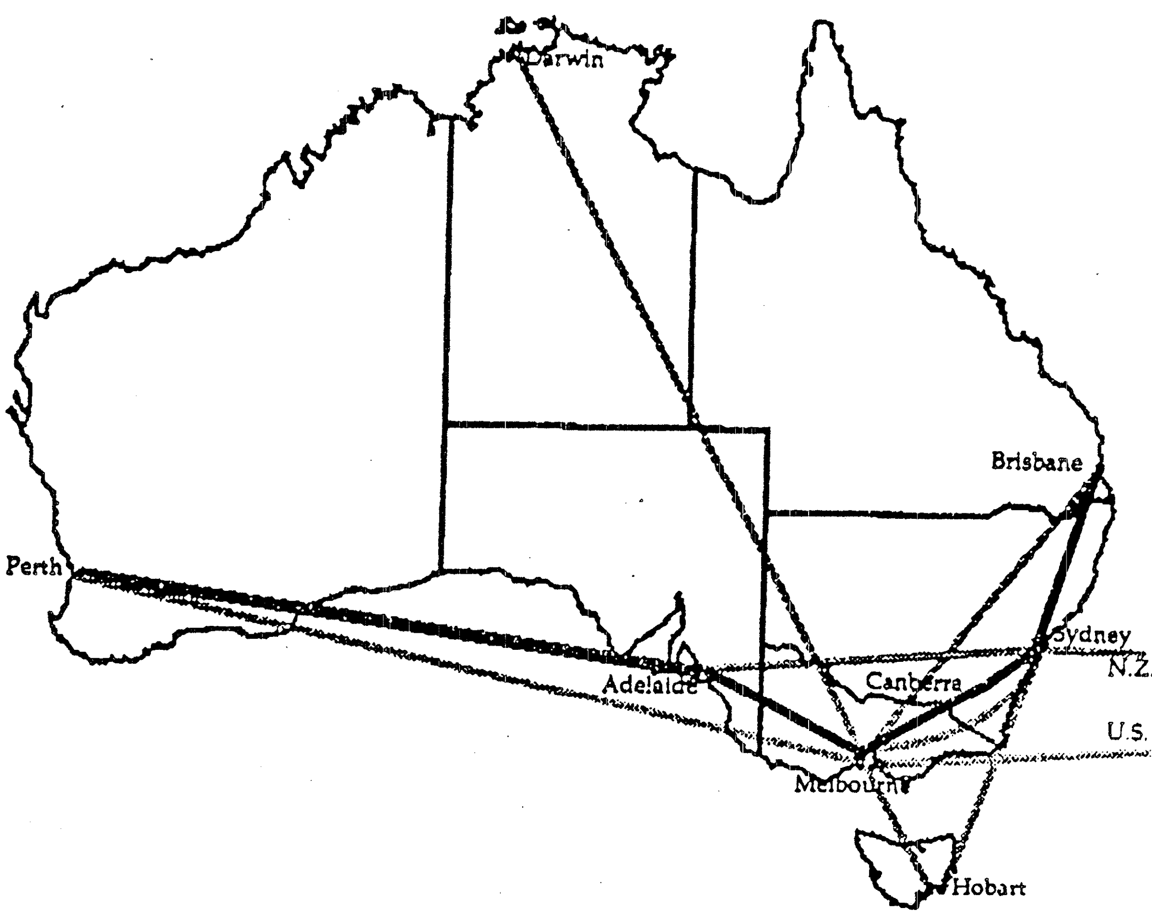

Increase bandwidth on trunk links - Megabit main trunk with additional 64K/144K ISDN links.

Major Regional Centres provided with dual paths, second hub located in Sydney for enhanced network throughput.

Link to Hawaii upgraded.

Possible international links to Japan and New Zealand.

Potential use of dynamic ISDN connections for peak load servicing

3.8.4 "Internetworking in the South Pacific"

Presented by Robert Elz/ University of Melbourne

New Zealand: There is primarily DECnet in the country. Some campuses are using IP, and there is IP over DECnet to get IP networks inter organizationally. New Zealand is connected directly into Hawaii.

Malisia Trying to get networking going. They have some internal networking on X.25 but international costs are prohibitive.

Singapore is on bitnet. The last I hear, they have 9600 baud dial up connection.

3.8.5 "Internetworking in Japan and the North Pacific"

Presentation by Jun Murai/ University of Tokyo

Questions from the presentation were as follows:

Q: There have been reports in the New York Times about price wars in Japan for new telecommunications services, including international services. Has this affected the usage of wide area networking yet? A: There is not a direct relationship between telephone price and network usage yet.

Q: Could you say a few words about ISDN in Japan? A: There is a lot of interest in ISDN. There is a lot of effort now to get it deployed. NTT has made a goal of providing ISDN service globally in Japan by the end of the century.

Q: Is there much interest in internationalizing the character sets? A: Yes. I have to clarify that a character is a byte. We have confronted those issues with the RFC 822 Kangi mail system, and are now looking at it in terms of X.400. It is not clear yet how it will work. Q: Does Kangi now use two bytes? A: Yes

Q: NTT has an IP link over X.25 to CSnet. Do you know of any other commercial links? A: Not at this time. We have held a lot of discussions with NTT in the past, and are not aware of any.

Internetworking in Japan

Jun Murai

WIDE Project
University of Tokyo

jun@u-tokyo.ac.jp

Networking in Japan

- Ministry of Education
  - N-1 Network for mainframes/super
- JUNET
  - CS Community
  - Email/RFC822+Japanese
- WIDE Internet
  - CS Community
  - IP Portion of JUNET
- TISN
  - Pysics/Astronomy Community
  - IP + DECnet
- HEPNET-J
  - Japan's HEPNET
- BITNET-J
- IP/NACSIS
  - IP/Private X.25

Ministry of Education

As a background..

- Seven Computer Centers
  - Recommendation of the Sci Council in 1963
  - Since 1965
  - Mainframes and Supercomputers to be shared
- N-1 Network
  - Development Started in 1973
  - 9.6Kbps and 48Kbps over X.25
  - Proprietary protocols
  - Resource Sharing: TSS Access, ftp and RJE
  - No Emails
  - Grouped in the 7 regions
  - Administration by the 7 Computer Center
  - Very Small Inter-Region Access
- NACSIS
  - Private X.25 Network for Universities
  - N-1 started to use it
  - IP over X.25

National Organization

Regions

- Region 1: Hokkaido
- Region 2: Aomori Pref., Iwate Pref., Miyagi Pref., Akita Pref., Yamagata Pref., Fukushima Pref.
- Region 3: Ibaraki Pref., Tochigi Pref., Gunma Pref., Saitama Pref., Chiba Pref., Tokyo, Kanagawa Pref., Niigata Pref., Yamanashi Pref., Nagano Pref., Shizuoka Pref.
- Region 4: Gifu Pref., Aichi Pref., Mie Pref.
- Region 5: Toyama Pref., Ishikawa Pref., Fukui Pref., Shiga Pref., Kyoto, Tottori Pref., Shimane Pref.
- Region 6: Osaka, Hyogo Pref., Nara Pref., Wakayama Pref., Okayama Pref., Tokushima Pref., Kagawa Pref., Ehime Pref., Kochi Pref.
- Region 7: Hiroshima Pref., Yamaguchi Pref., Fukuoka Pref., Saga Pref., Nagasaki Pref., Kumamoto Pref., Oita Pref., Miyazaki Pref., Kagoshima Pref., Okinawa Pref.

There are seven computer centers throughout Japan, each of which serves the universities in its region. Each center has a regional council and administration offices which researchers can contact when they wish to use the system. Researchers, in addition to using the center in their region, can also use any of the other centers in Japan. Thus, researchers are free to use the center with the characteristics most appropriate to their requirements. The system can be used through a nation-wide network, remote terminal, or the terminals at the centre.

Limitations on Usage

Each of the centers can only be used for pure research purposes, and any of the findings must be published.

Principal Equipment in Computer Centers

1. Hokkaido University Computing Center (1970)
   HITAC M-682H, S-810/10
2. Computer Center, Tohoku University (1969)
   NEC ACOS-2000, SX-1
3. Computer Centre, University of Tokyo (1965)
   HITAC M-682H×2, M-680H, S-820/80
4. Nagoya University Computation Center (1971)
   FACOM M780/20, VP-200
5. Data Processing Center, Kyoto University (1969)
   FACOM M780/30, VP-400E, VP-200
6. Computation Center, Osaka University (1969)
   NEC ACOS-2020, SX-2N
7. Computer Center, Kyushu University (1969)
   FACOM M780/20, VP-200

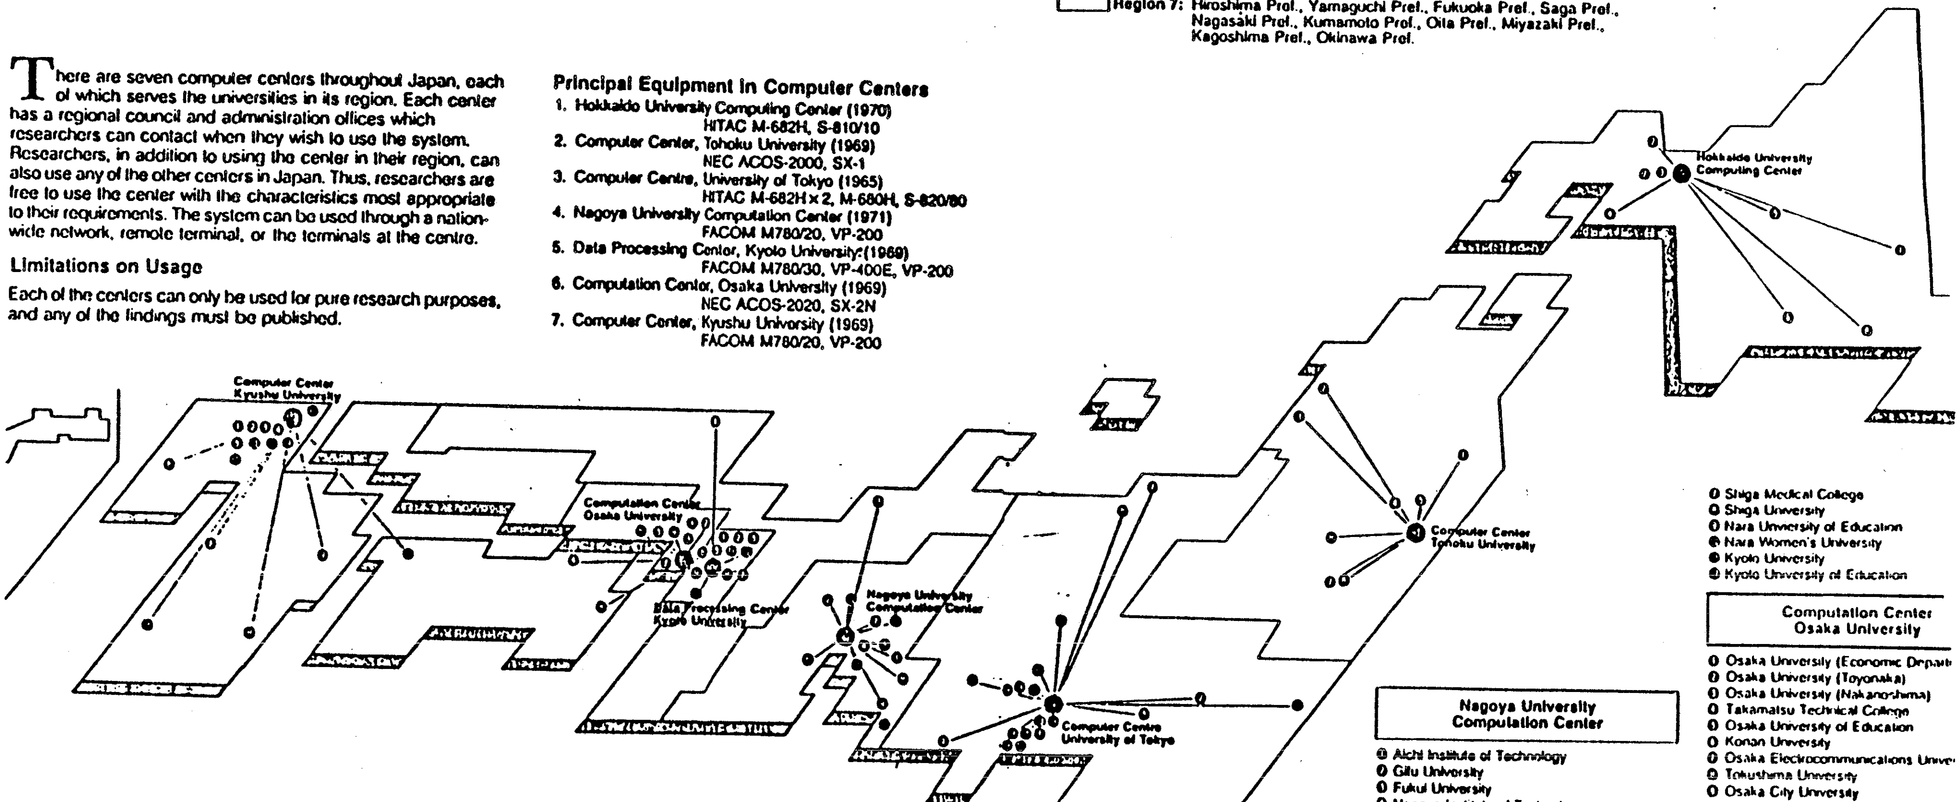

Remote Station Installations Connected to Computer Centers

Hokkaido University Computing Center

- Kitami Institute of Technology
- Otaru Commercial College
- Hokusei Gakuin Women's College
- Muroran Institute of Technology
- Asahikawa Medical College
- Obihiro Zootech College
- Hokkaido Institute of Technology
- Kushiro Technical College

Computer Center Tohoku University

- Iwate University
- Akita University
- Tsuruoka Technical College
- Miyagi Technical College
- Miyagi University of Education
- Fukushima University
- Fukushima Technical College
- Yamagata University (Yonezawa city)

Computer Centre University of Tokyo

- Tokyo Institute of Technology
- Tsukuba University
- Yamanashi University
- Shizuoka University
- Science University of Tokyo
- Joetsu University of Education
- Niigata University
- Saitama University
- Hitotsubashi University
- Ibaraki University
- Yokohama City University
- Gunma University
- Meiji Gakuin University
- Tokyo University of Agriculture and Technology
- Musashi Institute of Technology
- Azabu University
- Seijo University
- Tokai University
- Meiji University
- Risho University
- and twenty-nine other universities

Nagoya University Computation Center

- Aichi Institute of Technology
- Gifu University
- Fukui University
- Nagoya Institute of Technology
- Daido Institute of Technology
- Mie University
- Aichi University of Education
- Toyota Technical College
- Toyohashi Technical Science University
- Hamamatsu Medical College
- Nagoya University (Medical Department)
- Gifu Technical College
- Chubu University

Data Processing Center Kyoto University

- Kyoto Textile Industrial College
- Kyoto Medical College
- Doshisya University
- Tottori University
- Bukkyo University
- Shiga Medical College
- Shiga University
- Nara University of Education
- Nara Women's University
- Kyoto University
- Kyoto University of Education

Computation Center Osaka University

- Osaka University (Economic Depart…
- Osaka University (Toyonaka)
- Osaka University (Nakanoshima)
- Takamatsu Technical College
- Osaka University of Education
- Konan University
- Osaka Electrocommunications Unive…
- Tokushima University
- Osaka City University

Computer Center Kyushu University

- Fukuoka University of Education
- Kyushu University
- Fukuoka University
- Nakamura Gakuen University
- Kyushu Arts Engineering University
- Saga University
- Kumamoto University
- Miyazaki University
- Oita University
- Ehime University
- Kyushu Institute of Technology
- Kinki University (Kyushu Faculty of …
- Fukuoka Women's University
- Kyushu Industrial University
- Fukuoka Institute of Technology

	mainframe	supercomputer
Hokkaido Univ.	HITAC M-682H	S-810/10
Tohoku Univ.	ACOS2000	SX-1
Univ. of Tokyo	HITAC M-680H,M-682H	S-820
Nagoya Univ.	FACOM M-780/20	VP-200
Kyoto Univ	FACOM M-382,M-780/30	FACOM VP-200,FACOM VP-400E
Osaka Univ.	ACOS2000	SX-2N
Kyushu Univ.	FACOM M780/20	VP200

The Seven Computer Centers and their Computers

User's Region	Computer Centers							
	Hokkaido	Tohoku	Tokyo	Nagoya	Kyoto	Osaka	Kyushu	Total
Hokkaido	98.48	0.01	0.01	0.00	0.01	0.00	0.01	6.97
	94.89	0.00	0.03	0.00	0.17	0.00	0.00	11.82
Tohoku	0.05	99.25	0.01	0.37	0.02	0.00	0.00	6.36
	0.00	97.82	0.00	0.00	0.00	0.22	0.00	12.53
Tokyo	1.03	0.72	98.10	0.83	0.06	0.05	0.02	35.74
	5.11	2.18	98.70	0.12	4.89	0.36	0.00	25.49
Nagoya	0.40	0.02	0.08	96.56	0.17	0.00	0.01	12.26
	0.00	0.00	0.18	97.86	0.66	0.05	0.00	7.72
Kyoto	0.00	0.00	0.42	0.24	86.29	0.71	0.00	23.12
	0.00	0.00	1.00	1.97	83.48	0.35	0.00	21.66
Osaka	0.01	0.00	0.82	1.95	12.52	99.13	3.42	9.30
	0.00	0.00	0.08	0.05	7.32	99.02	18.36	13.60
Kyushu	0.03	0.00	0.56	0.05	0.93	0.11	96.54	6.25
	0.00	0.00	0.01	0.00	3.48	0.00	81.64	7.18
Total	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00

Regional Ratio in CPU
Time(Mainframes/Supercomputers)(1987)

4 Science Information Network

ıe NACSIS is conducting a project to construct a ıtionwide Science Information Network in order to omote the circulation of information among searchers. The Science Information Network is a ıvately operated packet switching network nploying high-speed digital circuits to terconnect packet switching nodes. The network ıs already expanded to cover major research eas, and is also interconnected with computing ıd information processing facilities located in universities and national research institutes under the Ministry of Education, Culture and Science. Most researchers belonging to universities and national research institutes are eligible to transfer their research-related information through the network and to exchange electronic mail between researchers. Also, the entire facilities provided by the Science Information System are accessible through the network.

The network will serve as the fundamental communication facility for the Inter-University Computer Network of Japan (N-1 Network) and the University Library Network, in which NACSIS plays a primary role. While NACSIS is promoting interconnections among Local Area Networks on university campuses, closed group of computers or researchers, and other data communication services in Japan and abroad in order to realize advanced use of applications for a wider range of researchers, the network will support network interconnection, closed group serviced, and vı network operation, as well as protocol conversı order to realize these network functions adequate efficiency, network protocols will adapted to an internationally standardized sch that is, OSI (Open System Interconnection). planned is the incorporation of an integı communication service into the network that include full text, graphics, voice, and ın communication.

Science Information Network

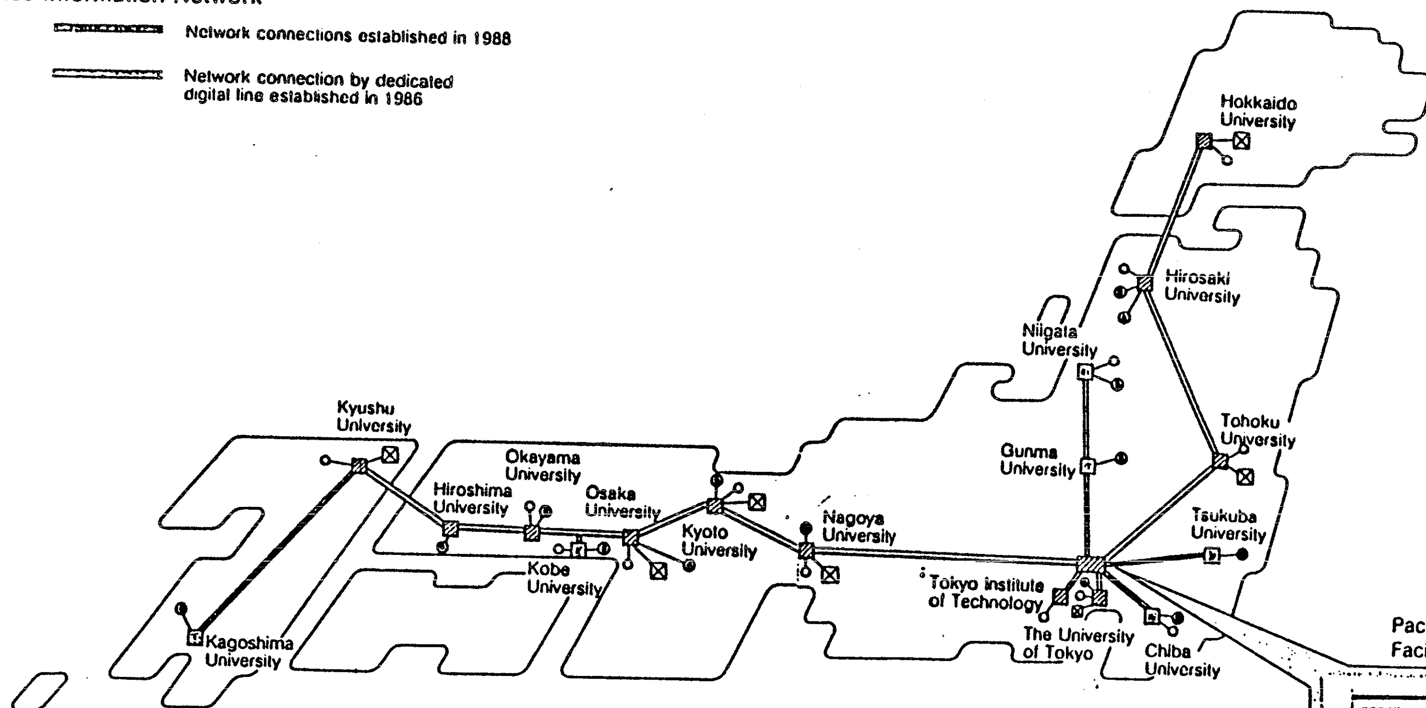

Packet-Switching Facility at NACSIS

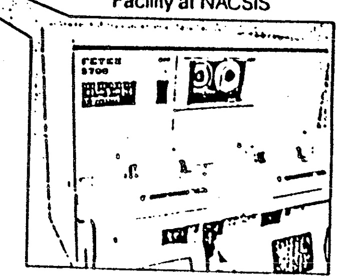

- Operating equipment of Science Information Network (packet switching modes, etc.)
- Facilities of Science Information Network established 1988 (as above)
- Inter-University computer centers
- Information processing centers
- University libraries

Message Exchange in JUNET

RFC822

- Japanese Extensions on Body
  - JIS X0208
    - Kanji Code Set
  - JIS X0202 (ISO 2022)
    - Introducing (Escape) sequences
- Japanized Version of X Windows client
- Kana-Kanji Conversion with inet IPC
- Gnuemacs/Bnews/etc..
- More than 90% of domestic traffic

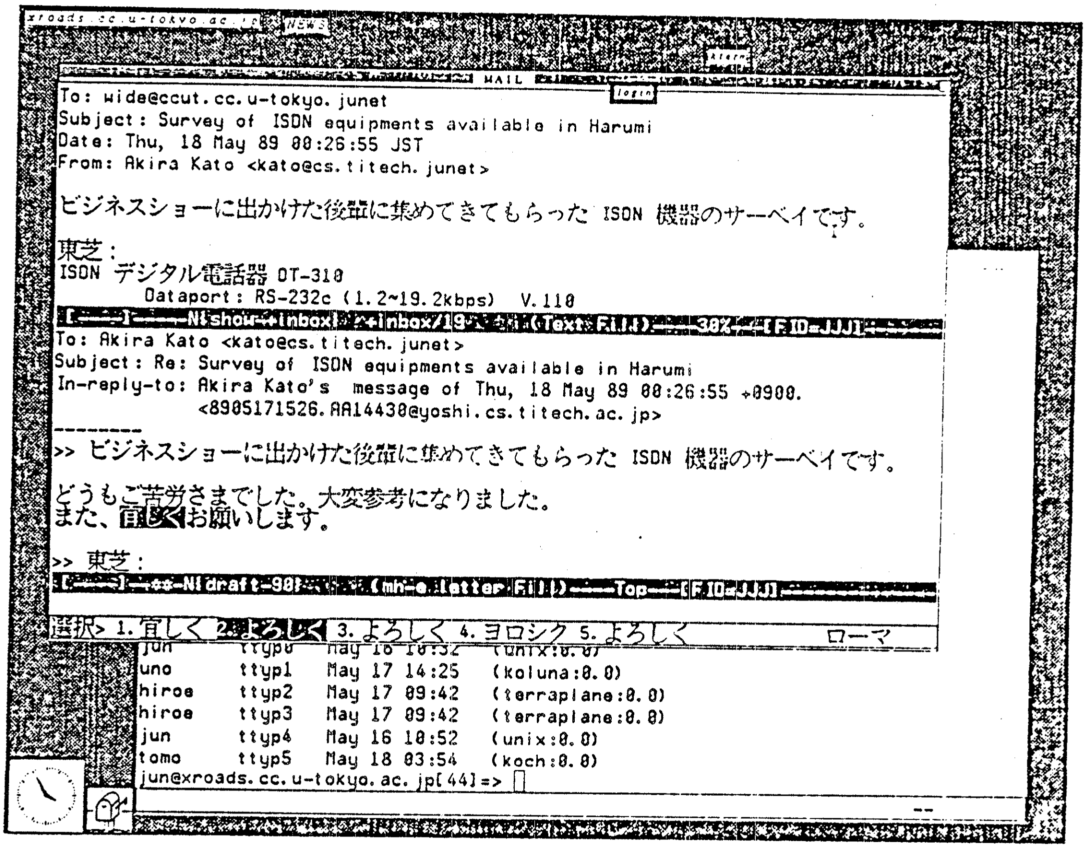
xroads.cc.u-tokyo.ac.jp
MAIL
To: wide@ccut.cc.u-tokyo.junet
Subject: Survey of ISDN equipments available in Harumi
Date: Thu, 18 May 89 00:26:55 JST
From: Akira Kato <kato@cs.titech.junet>
ビジネスショーに出かけた後輩に集めてきてもらった ISDN 機器のサーベイです。
東芝:
ISDN デジタル電話器 DT-310
Dataport: RS-232c (1.2~19.2kbps) V.110
To: Akira Kato <kato@cs.titech.junet>
Subject: Re: Survey of ISDN equipments available in Harumi
In-reply-to: Akira Kato's message of Thu, 18 May 89 00:26:55 +0900.
<8905171526.AA14430@yoshi.cs.titech.ac.jp>
>> ビジネスショーに出かけた後輩に集めてきてもらった ISDN 機器のサーベイです。
どうもご苦労さまでした。大変参考になりました。
また、宜しくお願いします。
>> 東芝:
選択> 1. 宜しく 2.よろしく 3. よろしく 4.ヨロシク 5.よろしく ローマ
uno ttyp1 May 17 14:25 (koluna:0.0)
hiroe ttyp2 May 17 09:42 (terraplane:0.0)
hiroe ttyp3 May 17 09:42 (terraplane:0.0)
jun ttyp4 May 16 10:52 (unix:0.0)
tomo ttyp5 May 18 03:54 (koch:0.0)
jun@xroads.cc.u-tokyo.ac.jp[44]=>

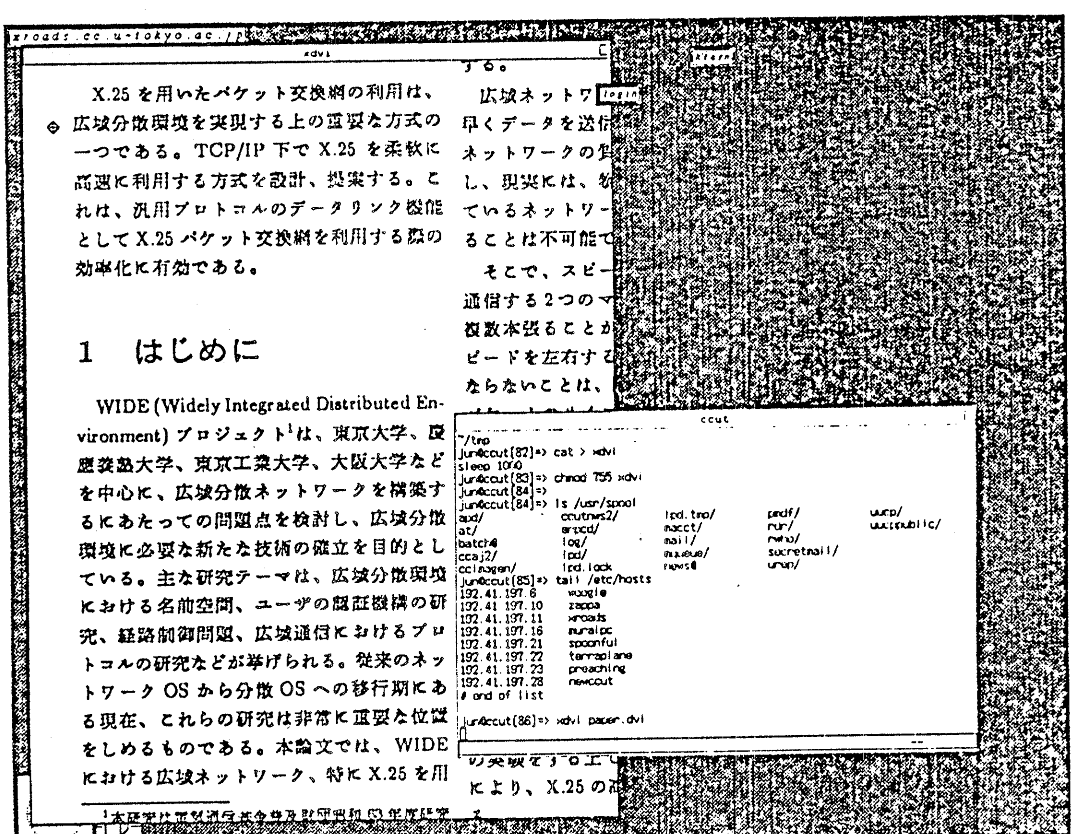
xroads.cc.u-tokyo.ac.jp
xdvi
X.25 を用いたパケット交換網の利用は、広域分散環境を実現する上の重要な方式の一つである。TCP/IP 下で X.25 を柔軟に高速に利用する方式を設計、提案する。これは、汎用プロトコルのデータリンク機能として X.25 パケット交換網を利用する際の効率化に有効である。
1 はじめに
WIDE (Widely Integrated Distributed Environment) プロジェクト[1]は、東京大学、慶應義塾大学、東京工業大学、大阪大学などを中心に、広域分散ネットワークを構築するにあたっての問題点を検討し、広域分散環境に必要な新たな技術の確立を目的としている。主な研究テーマは、広域分散環境における名前空間、ユーザの認証機構の研究、経路制御問題、広域通信におけるプロトコルの研究などが挙げられる。従来のネットワーク OS から分散 OS への移行期にある現在、これらの研究は非常に重要な位置をしめるものである。本論文では、WIDE における広域ネットワーク、特に X.25 を用
広域ネットワ
早くデータを送信
ネットワークの
し、現実には、
ているネットワー
ることは不可能で
そこで、スピー
通信する2つのマ
複数本張ることが
ピードを左右す
ならないことは、
により、X.25 の
ccut
jun@ccut[82]=> cat > xdvi
sleep 1000
jun@ccut[83]=> chmod 755 xdvi
jun@ccut[84]=> ls /usr/spool
jun@ccut[85]=> tail /etc/hosts
192.41.197.6 woogie
192.41.197.10 zappa
192.41.197.11 xroads
192.41.197.16 muraipc
192.41.197.21 spoonful
192.41.197.22 terraplane
192.41.197.23 preaching
192.41.197.28 newccut
end of list
jun@ccut[86]=> xdvi paper.dvi

JUNET

- UUCP + IP
- Since 1984
- Volunteer base
- 270 organizations
  - — Email/Bnews
  - — fj news groups
- Migrating into WIDE Internet

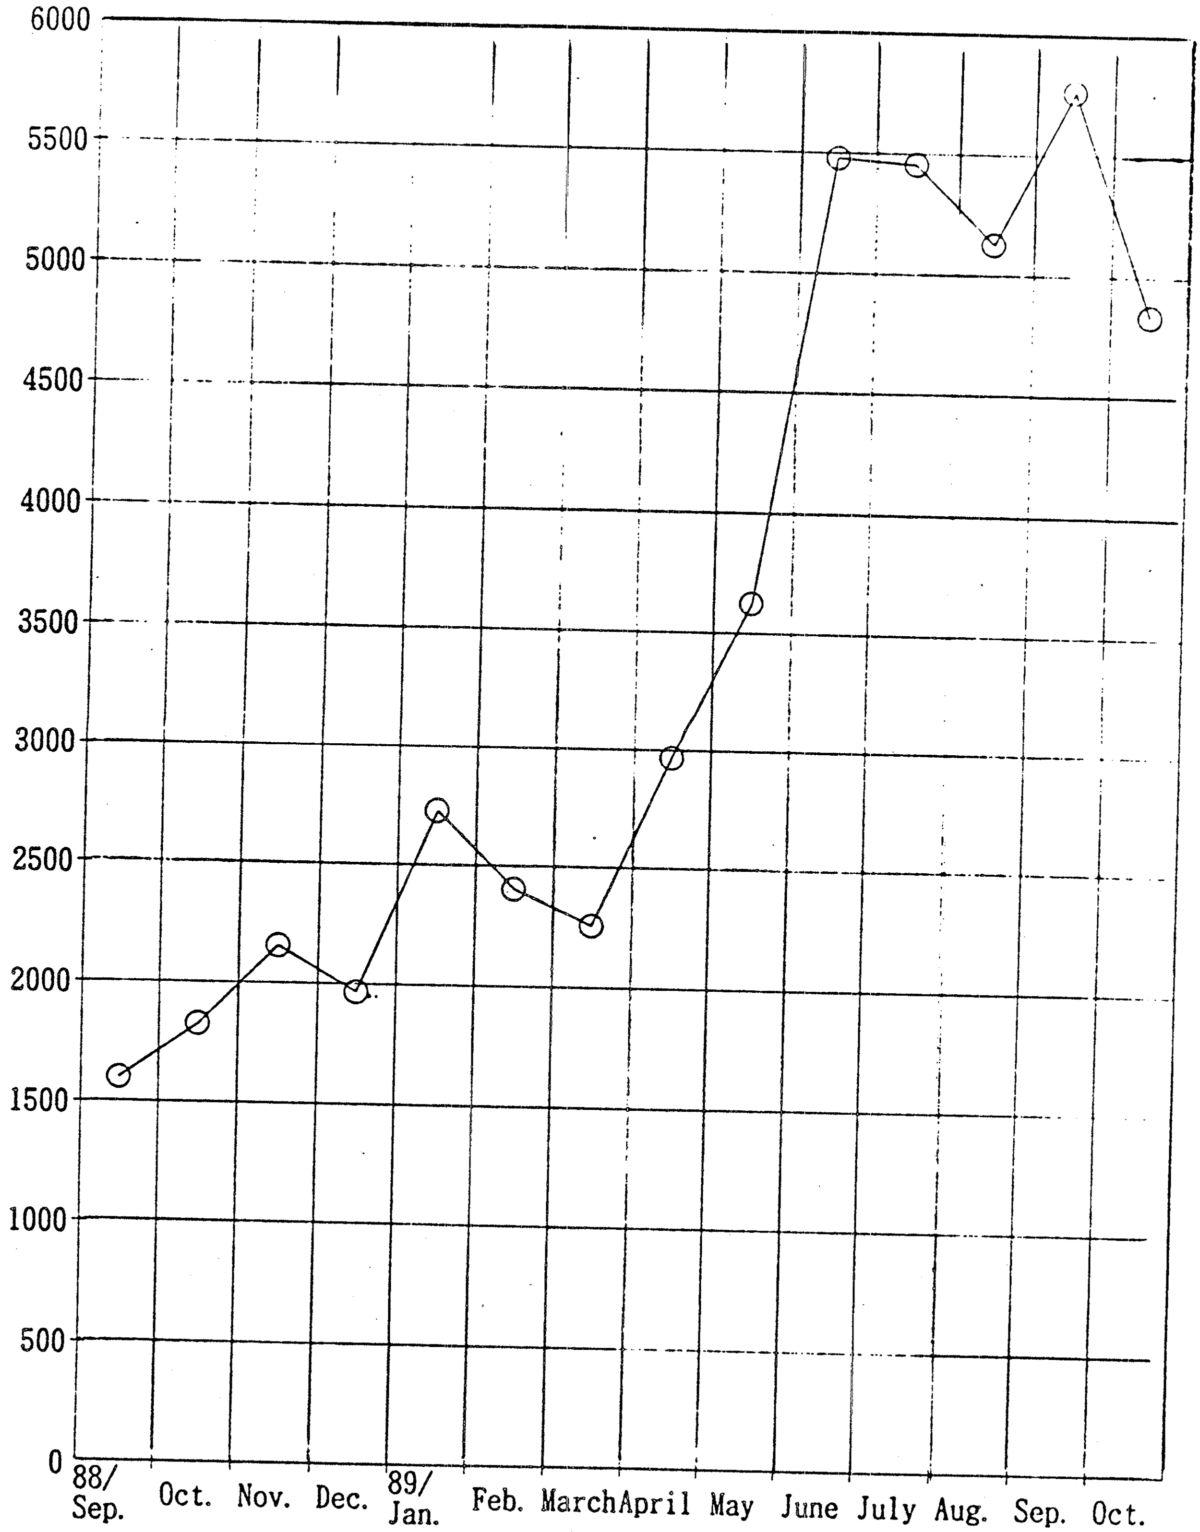

Number of Articles

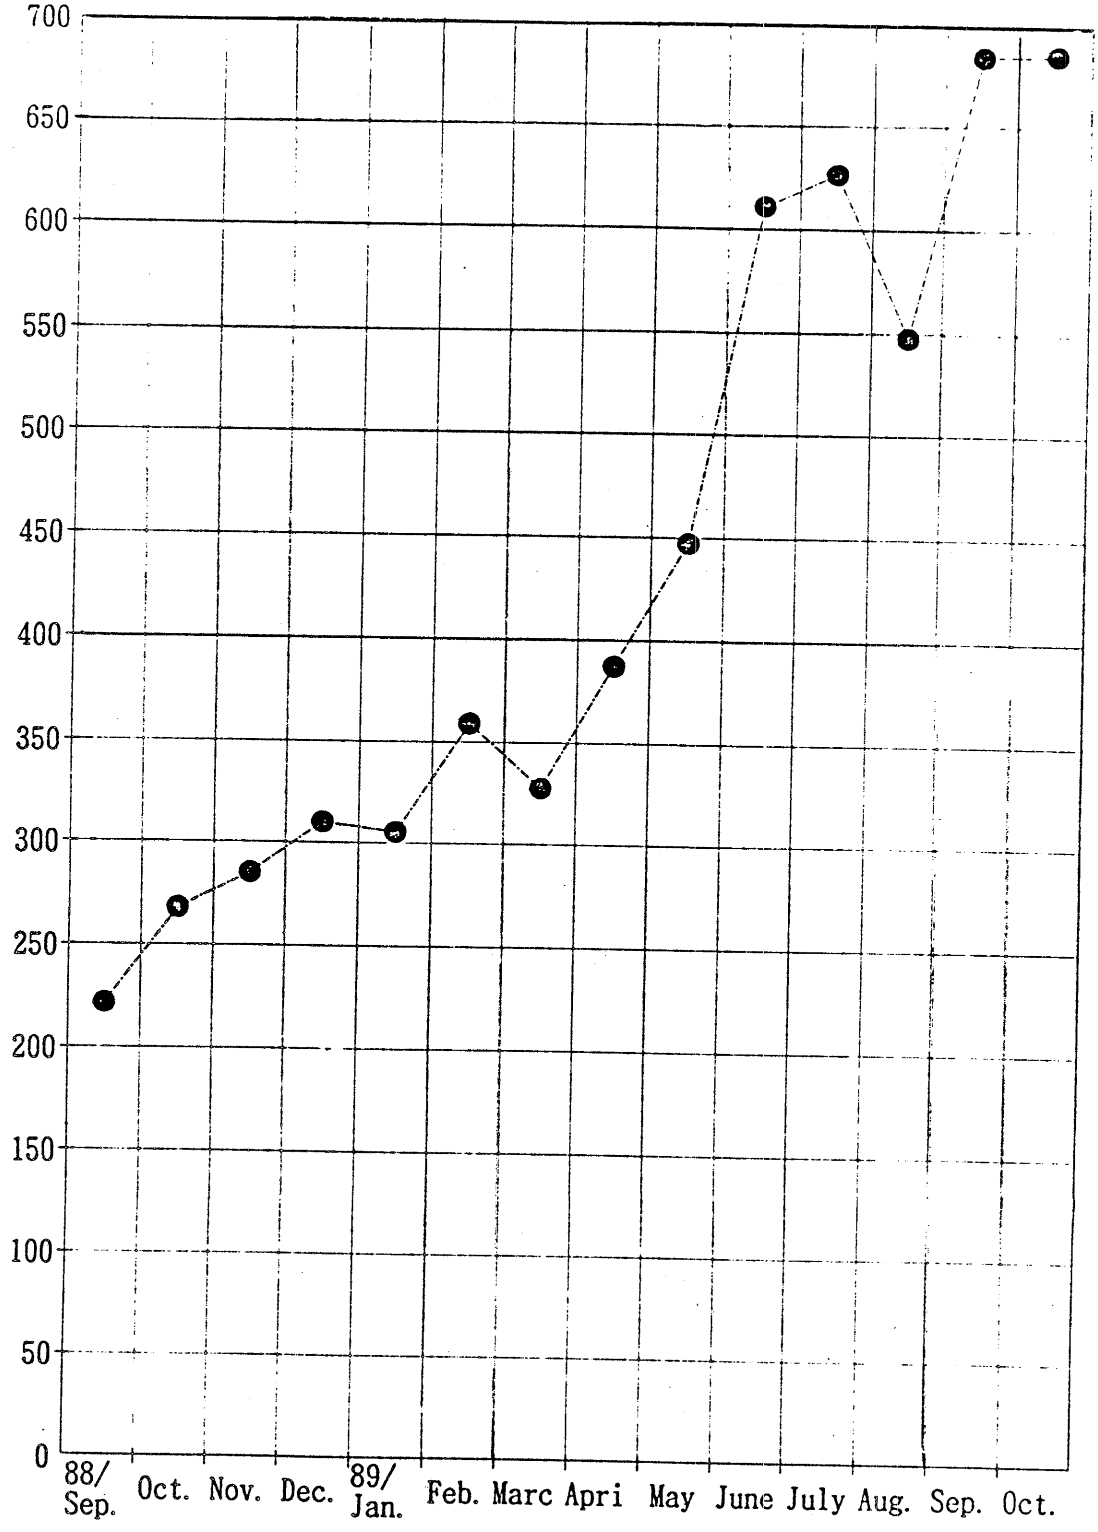

Number of Active Posters

DECNET in Japan

- KEK
  - National Laboratory for High Energy Physics
- Links
  - NACSIS Private X.25
  - NTT DDX-P X.25
  - FENIX - Fujitsu X.25
  - Leased lines: 9.6Kbps and 64Kbps

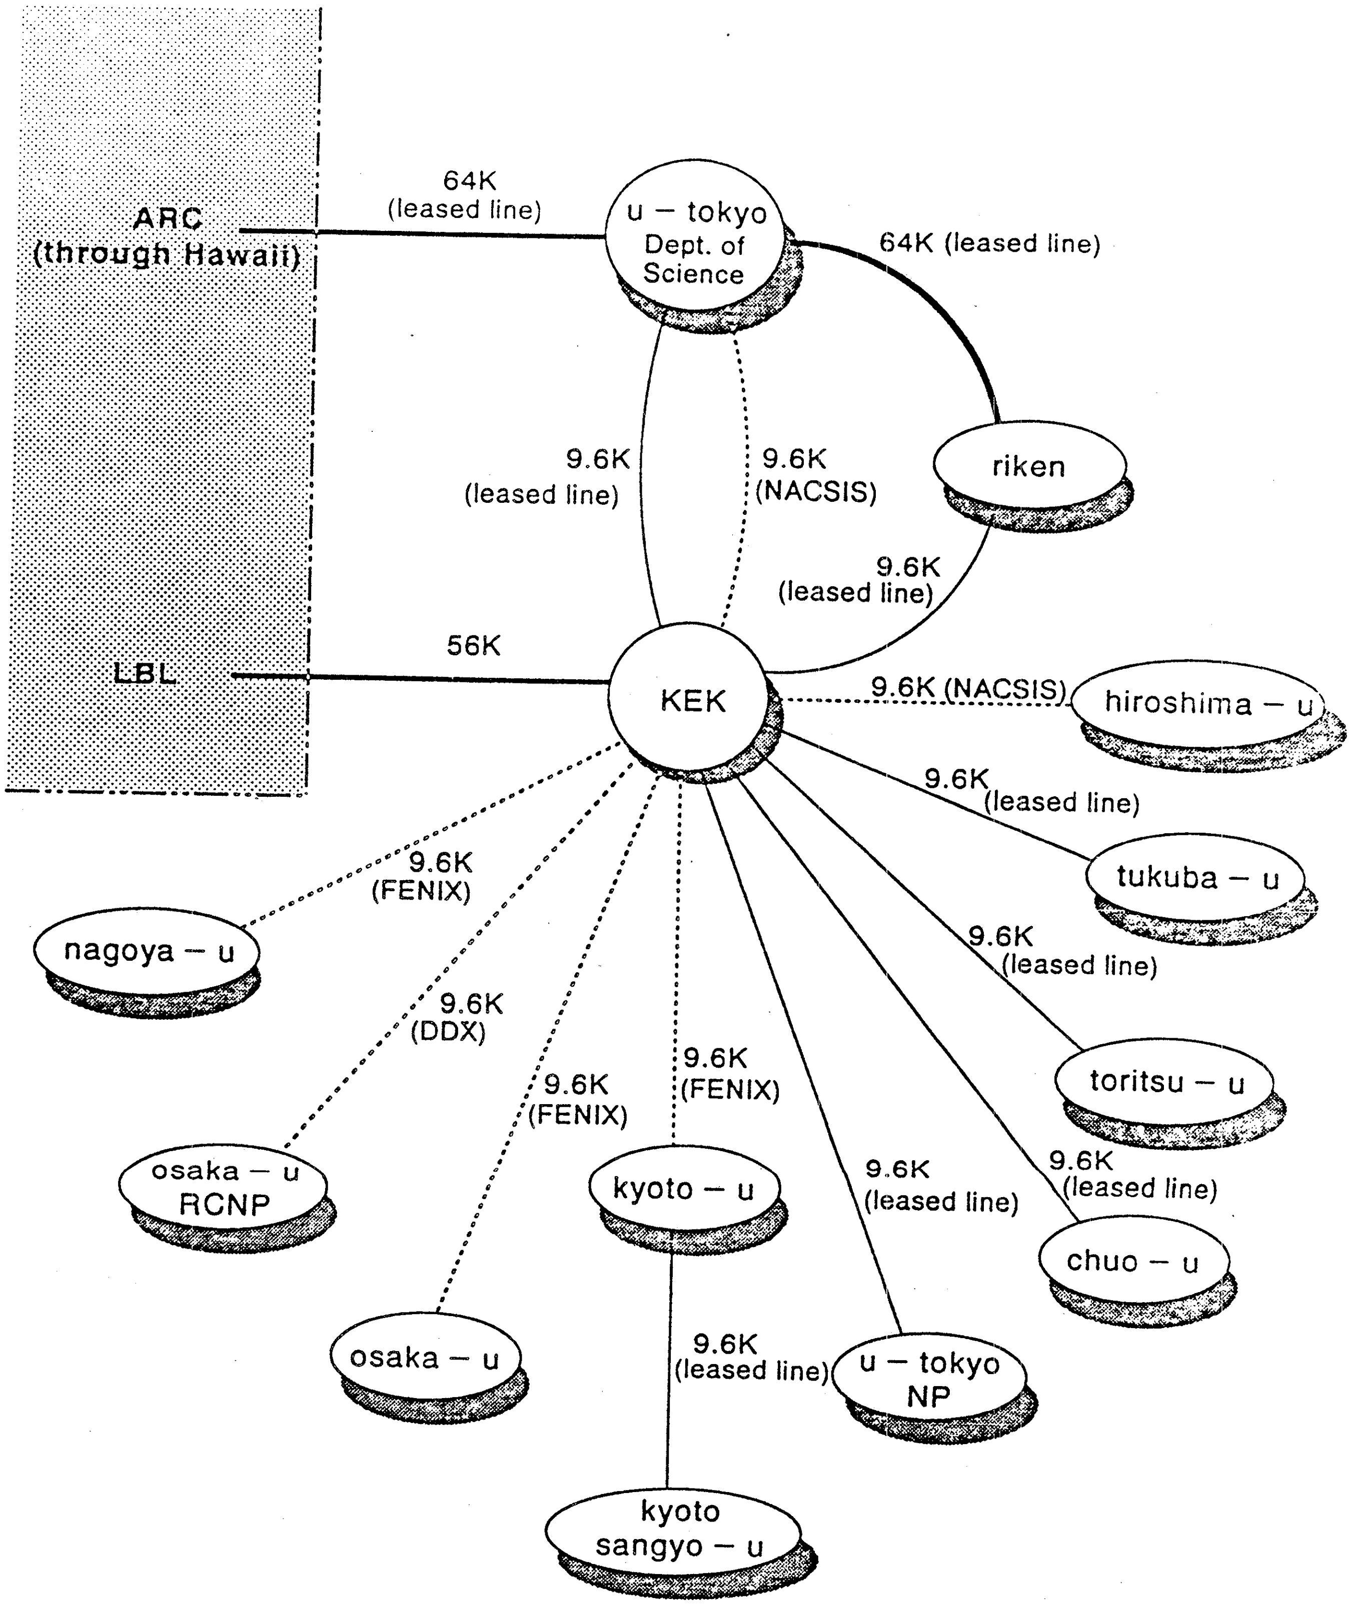

DECNET National Connections
HEPNET – J & TISN

IP in Japan

IP Activities

- WIDE Project
- IP/NACSIS Project
- TISN Project
- WIDE Project
  - US$ 0.4M: Various funding resources
  - Consortium of CS researchers
  - Office: U. Tokyo (KEIO Univ. from April 90)
- IP/NACSIS Project
  - M. of Education research
  - Actually done by WIDE Project
  - IP over X.25 (RFC877+)
- TISN Project
  - University of Tokyo
  - Astronomy and Physics communities
  - IP + DECNET

Domestic Links

- JUNET
  - UUCP/Dial-UP
  - UUCP/X.25
- WIDE
  - IP/64Kbps Leased
  - IP/192Kbps Leased
  - IP/2*64K ISDN
  - IP/SLIP 3.4KHz Voice Leased
  - IP/NACSIS-X.25 9.6Kbps
- TISN
  - IP/64Kbps Leased

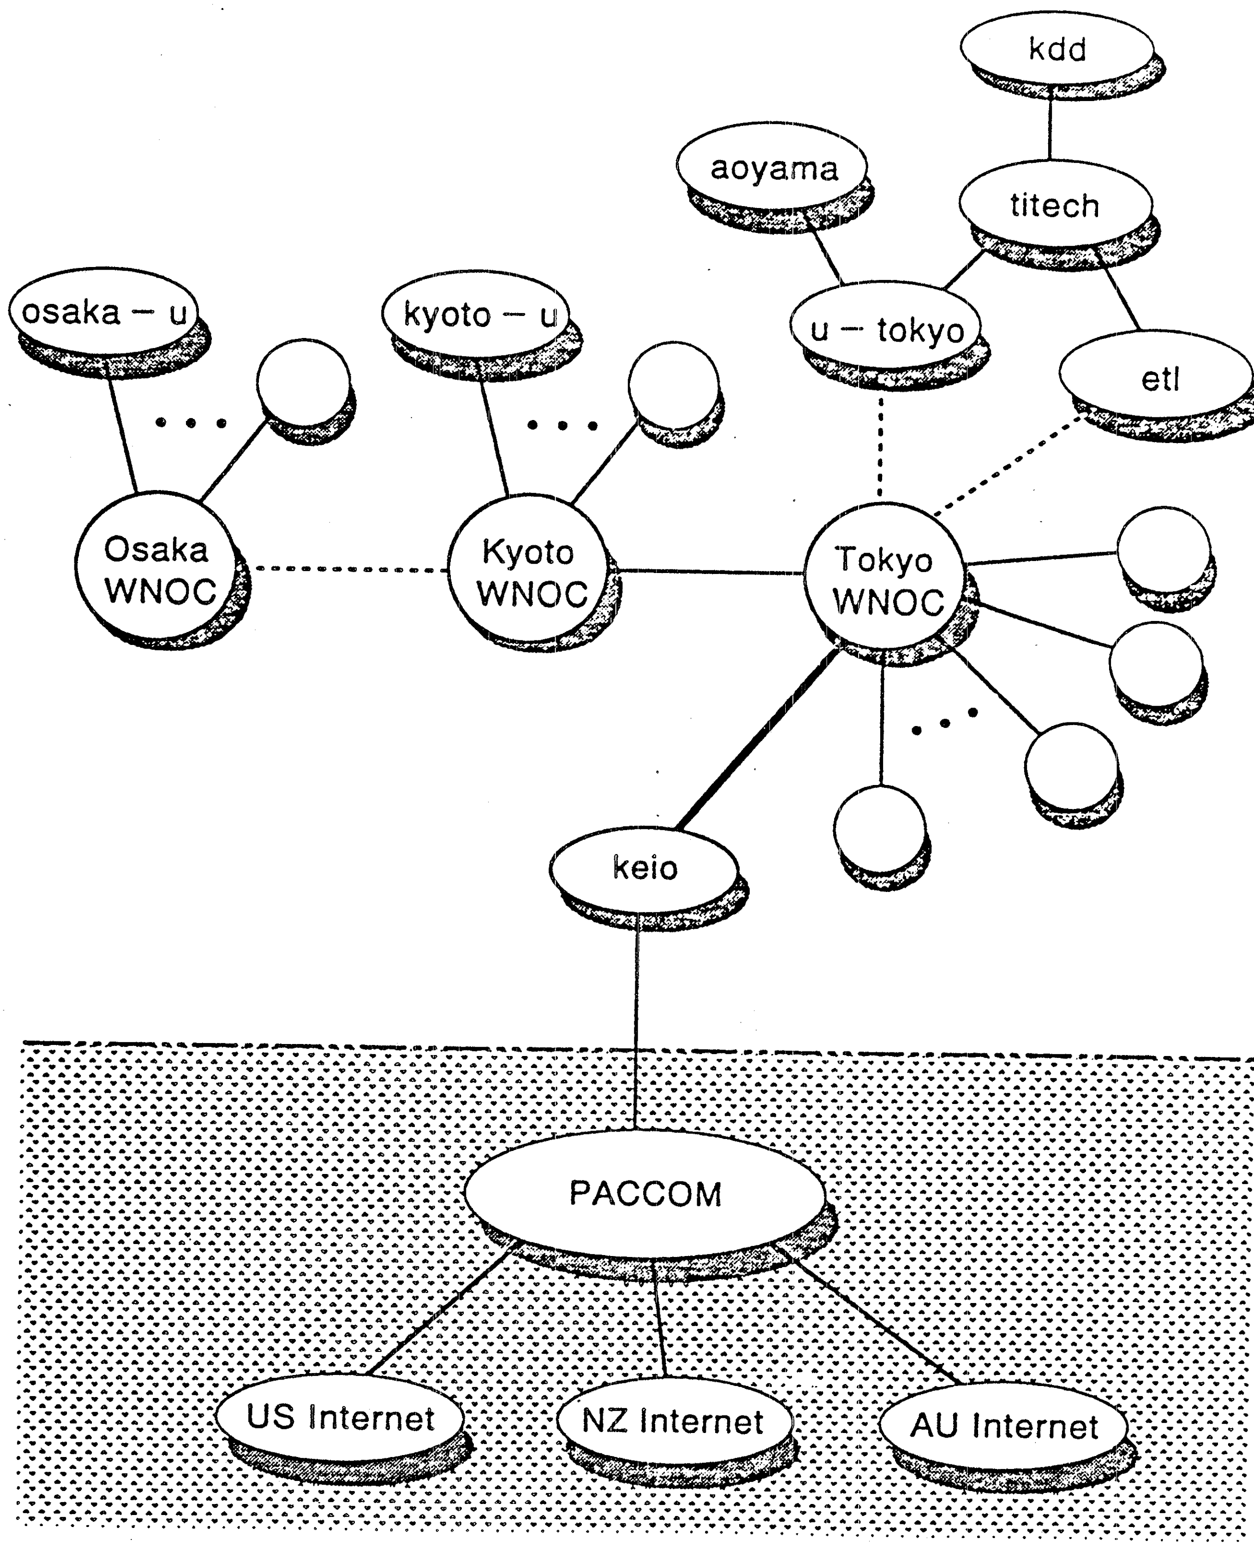

WIDE Logical Connections

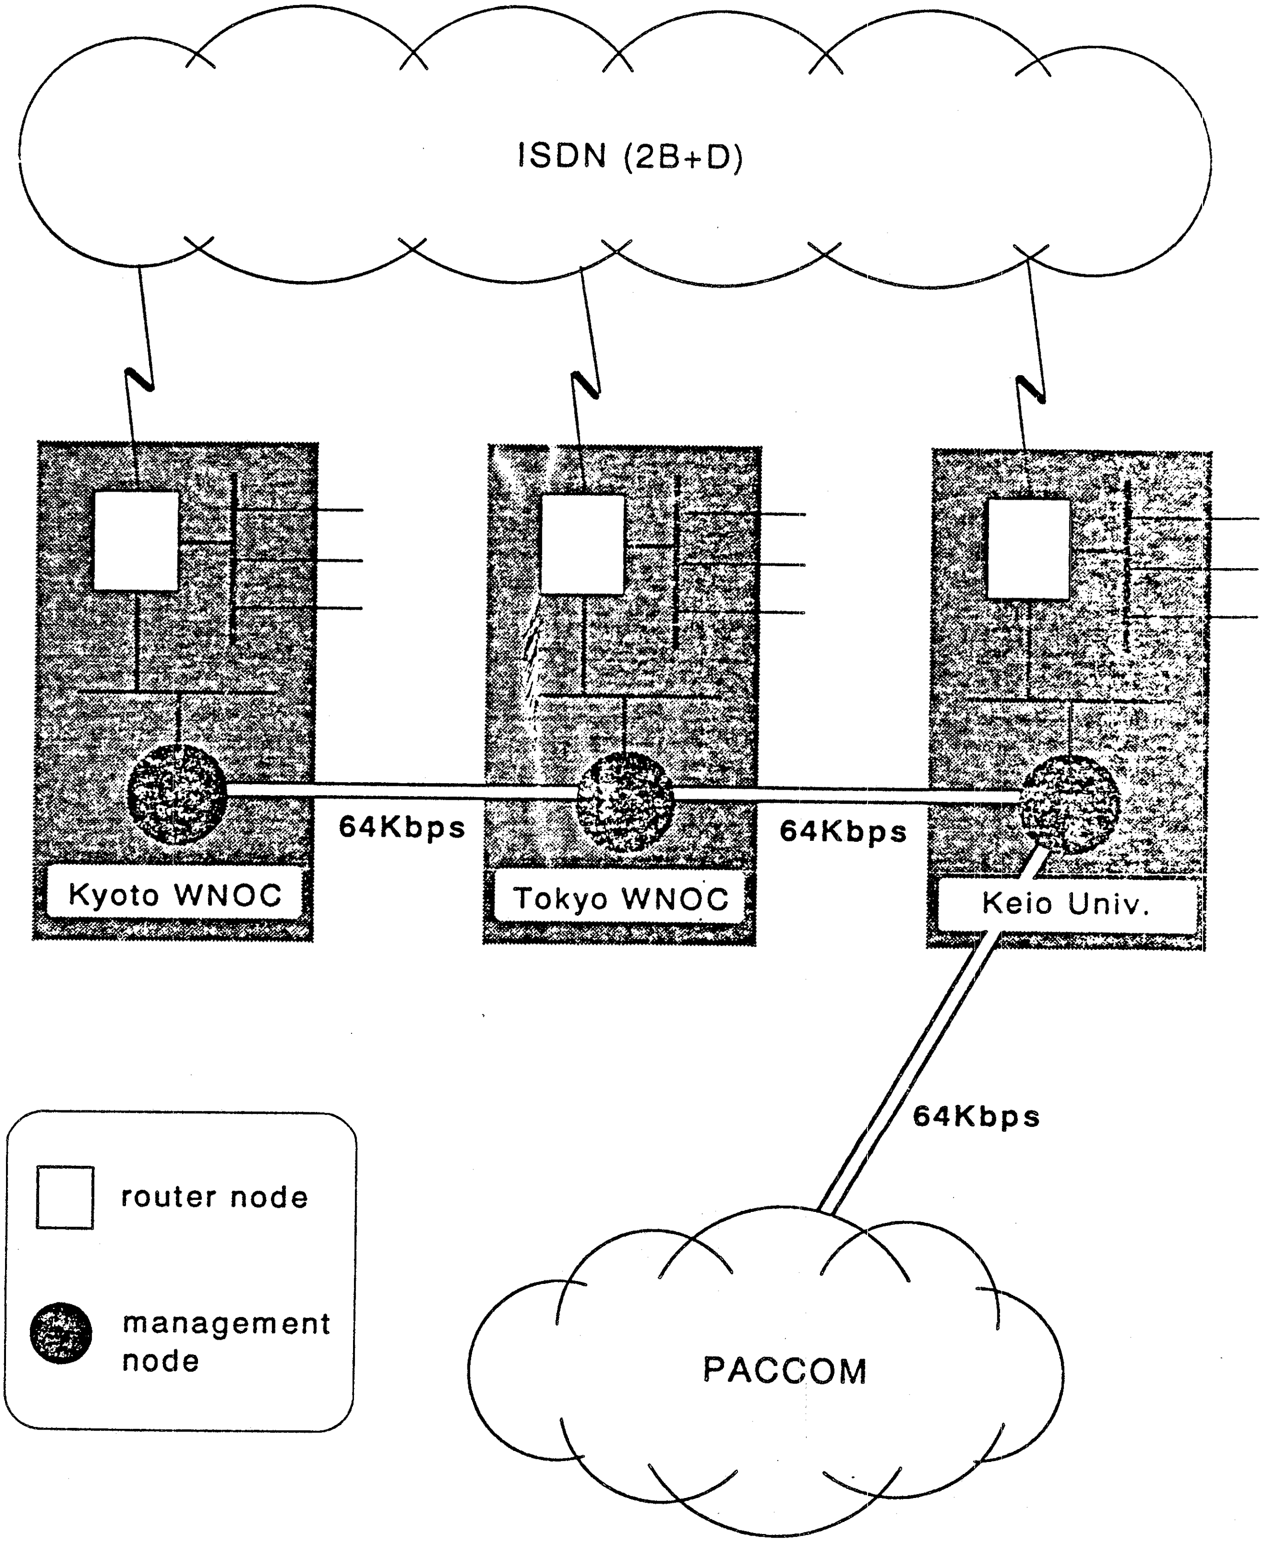

WIDE Backbone Structure

	AREA		WIDE	NACSIS	TISN
①	Sendai	Tohoku Univ.		✓	
②	Tsukuba	ETL	✓		
		KEK (National Lab. for High Energy Phisics)			✓
③	Tokyo	Tokyo WNOC	✓		
		Univ. of Tokyo	✓	✓	✓
		Keio Univ.	✓		
		Tokyo Institute of Technology	✓	✓	
		NACSIS		✓	
		Aoyama Univ.	✓		
		Instisute of Phisical & Chemical Research			✓
		KDD	✓		
④	Nagoya	Nagoya Univ.		✓	
⑤	Kyoto	Kyoto WNOC	✓		
		Kyoto Univ.	✓	✓	
⑥	Osaka	Osaka WNOC	✓	✓	
		Osaka Univ.	✓		
⑦	Kyushu	Kyushu Univ.	✓	✓	

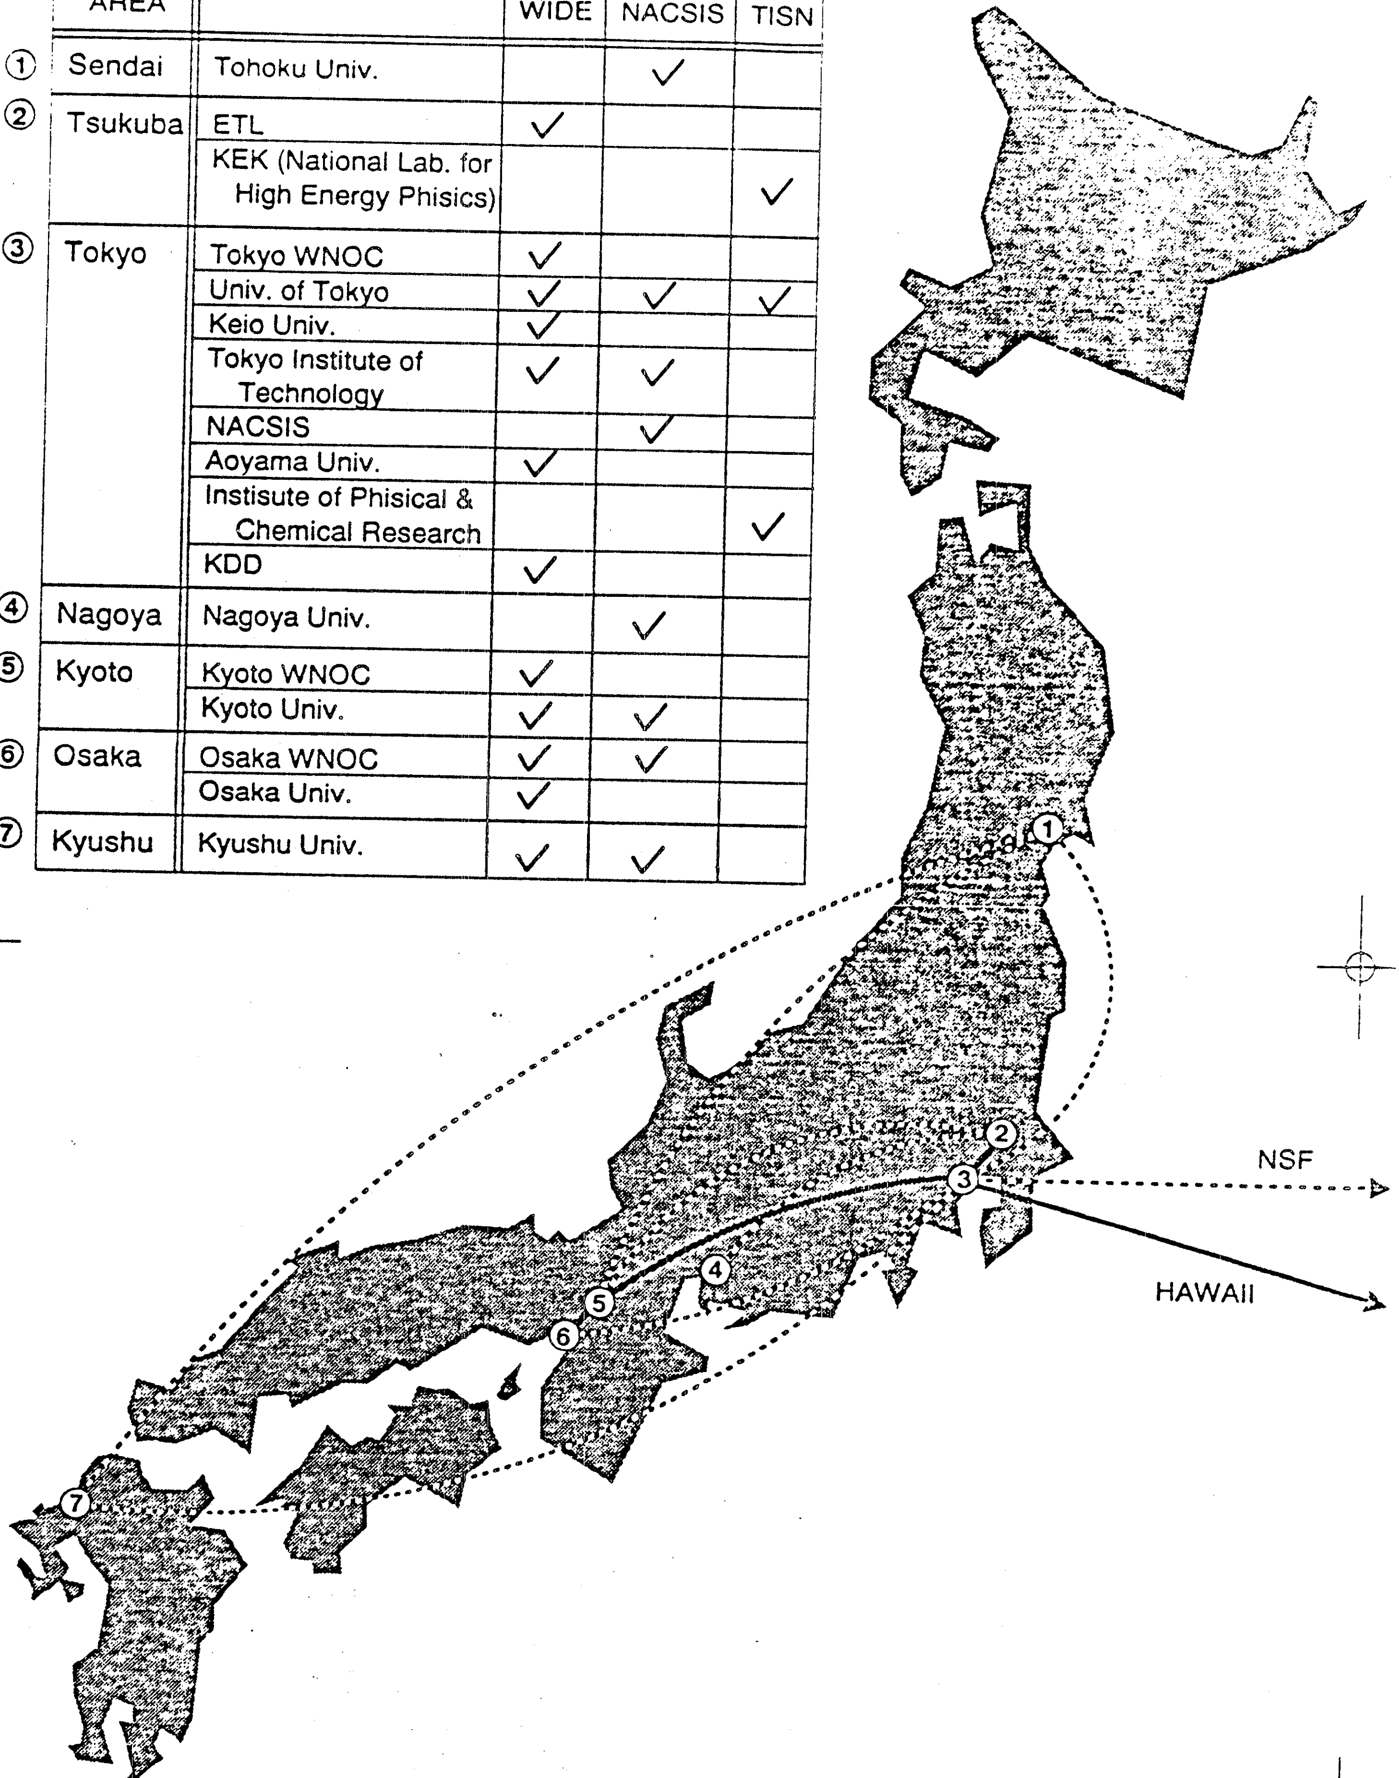

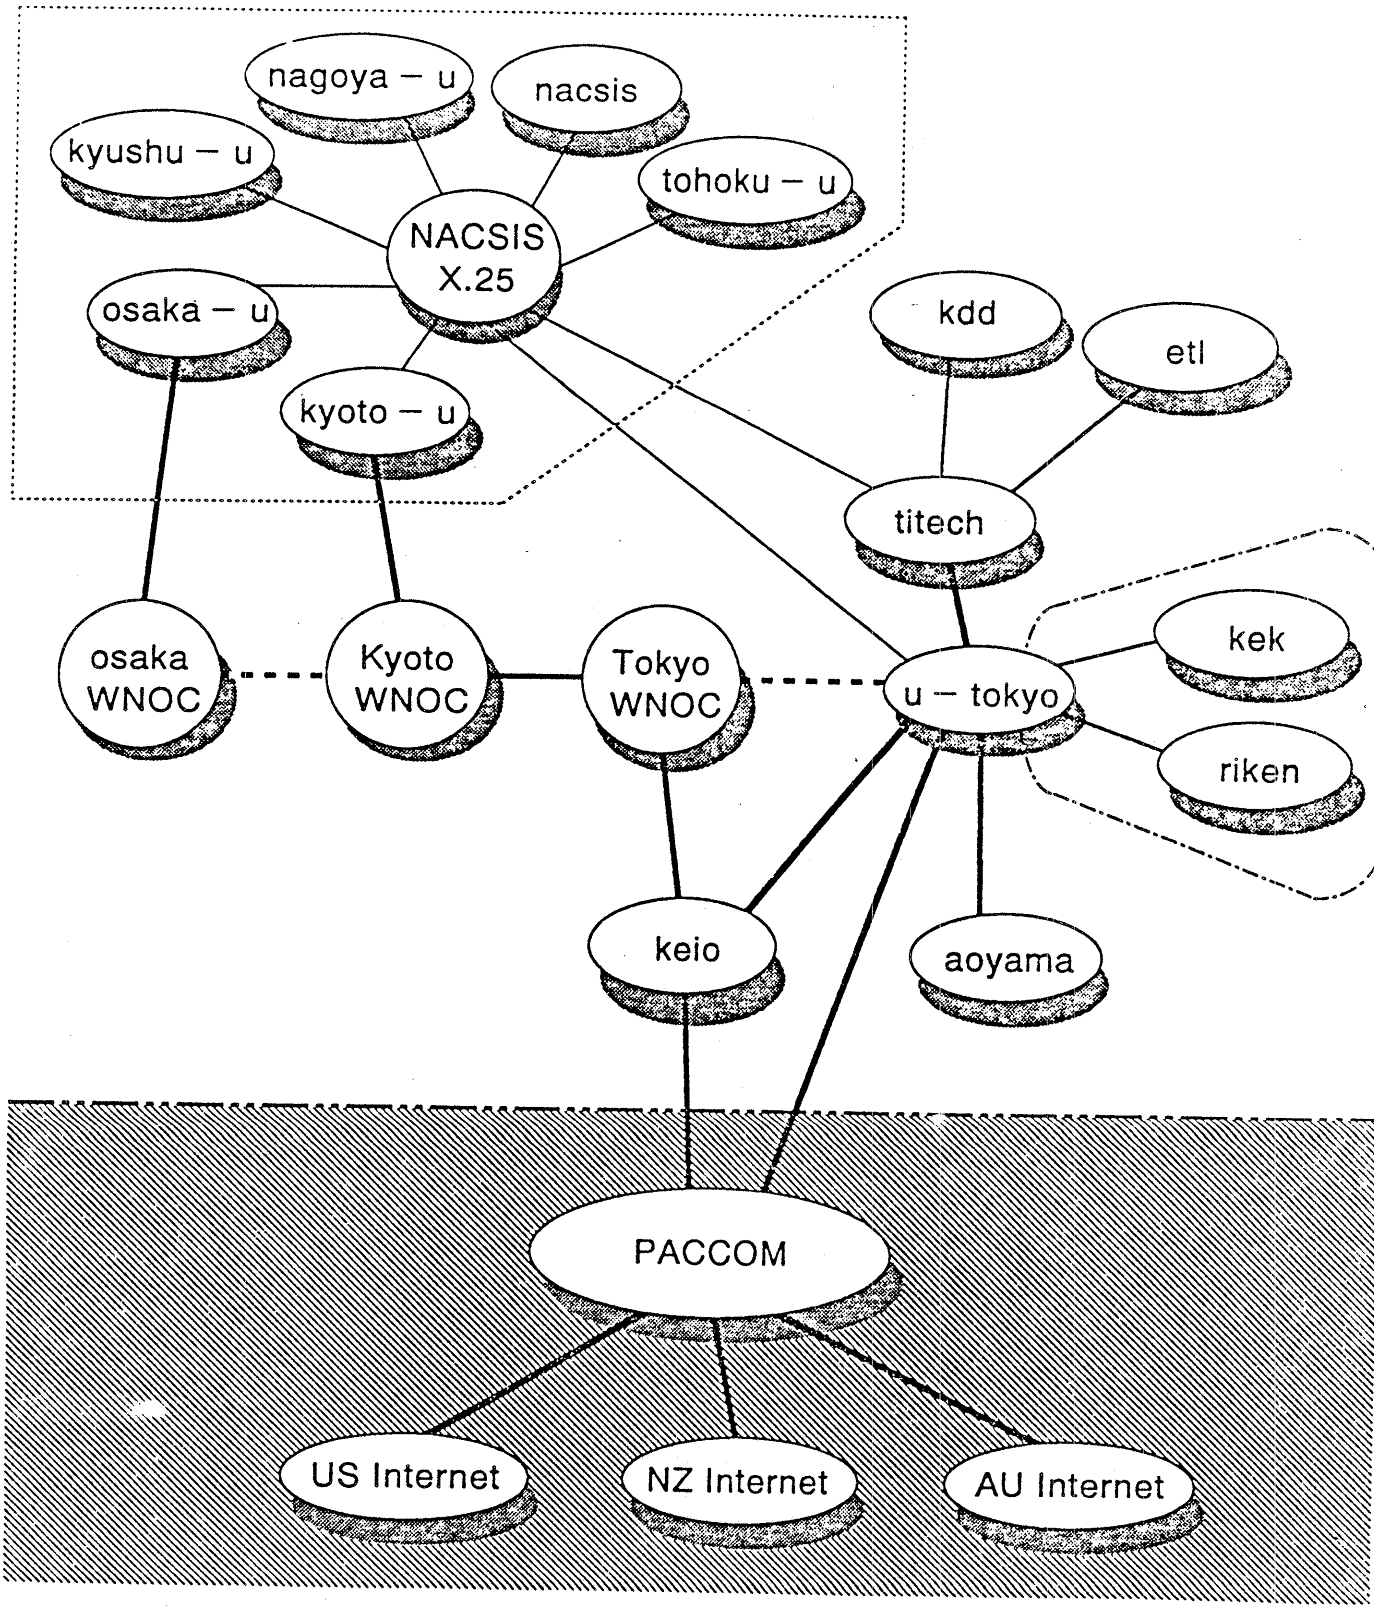

WIDE Logical Connections

IP Administration

- Address:
  - WIDE Project(jun)
    - Re-allocation of SRI-NIC's group allocation
- Domain:
  - JP-DOM (.jp) at jp-gate.wide.ad.jp
    - ac: universities
    - ad: network administrators
    - co: commercial
    - go: government
    - or: (non-profit) organization
- Domestic Connectivity:
  - WIDE, TISN and IP/NACSIS
- International Connectivity:
  - WIDE and TISN

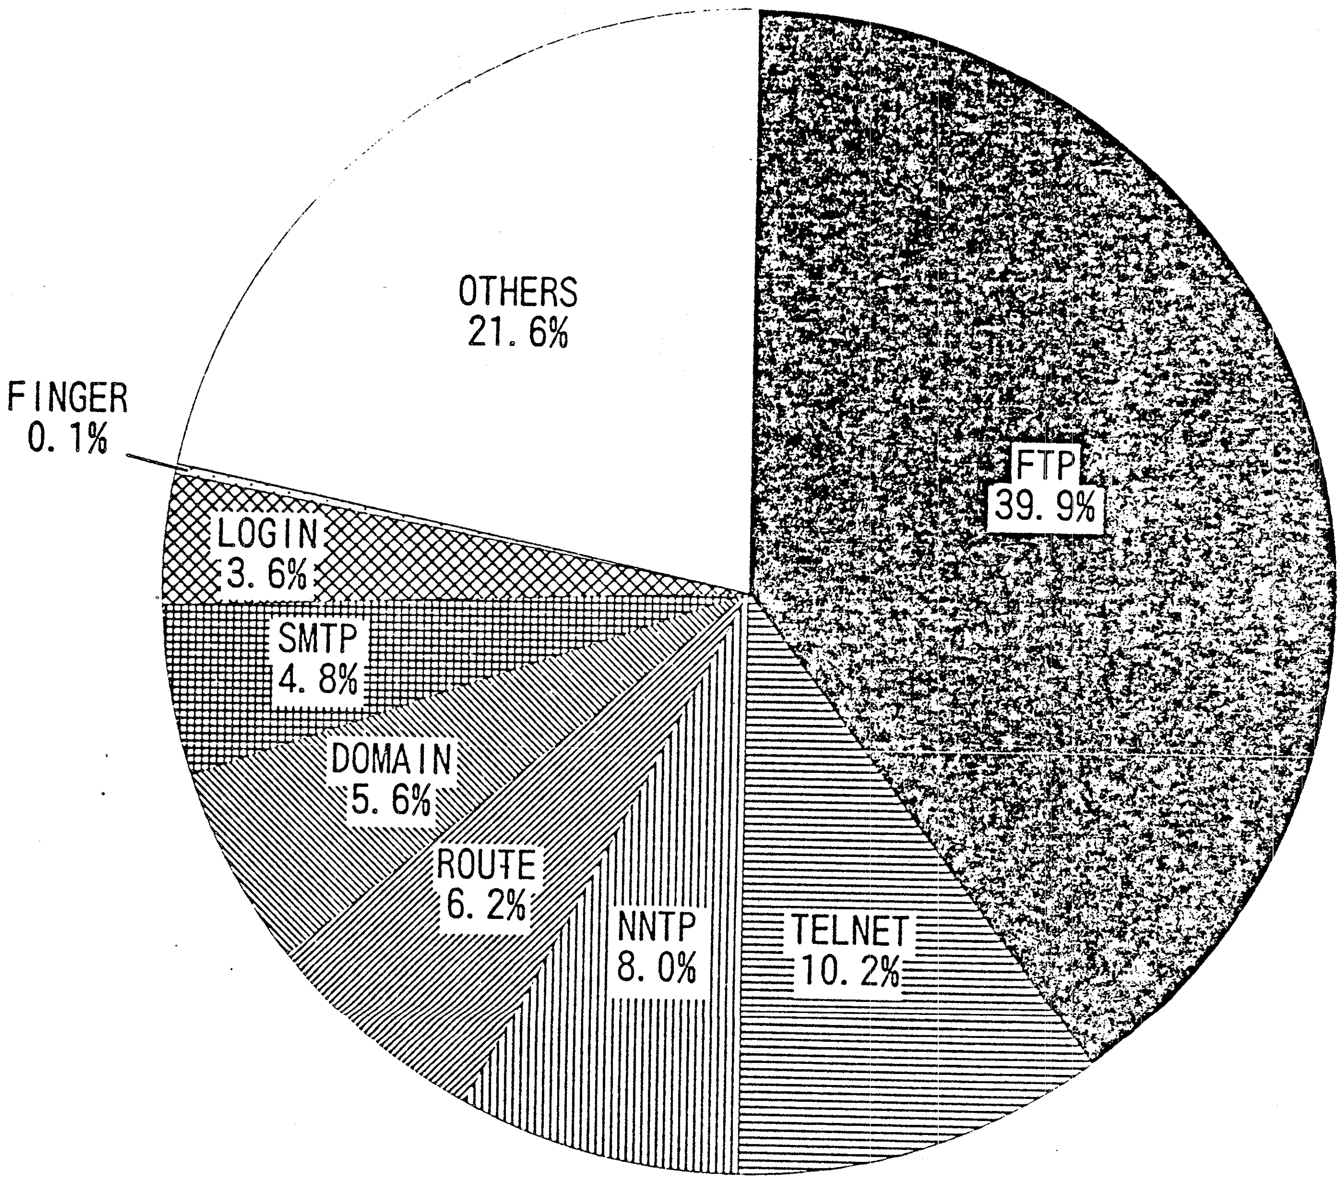

WIDE Backbone(keio<->u-tokyo) traffic

International Links

- IP
  - WIDE - 64Kbps/TPC3
    - PACCOM/U. of Hawaii
  - TISN - 64Kbps/TPC3
    - PACCOM/U. of Hawaii
  - NACSIS - 9.6Kbps
    - SURAnet/NSF
- BITNET
  - Sci. Univ of Tokyo - 9.6Kbps
    - CUNY
- DECNET
  - KEK
    - LBL - 56Kbps Satellite

NETWORK TRAFFIC INTO AND OUT OF JAPAN

IN
OUT

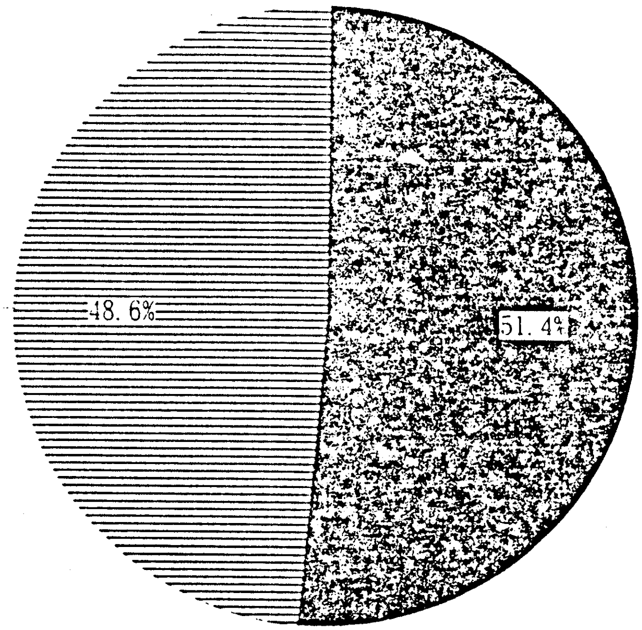

TISN

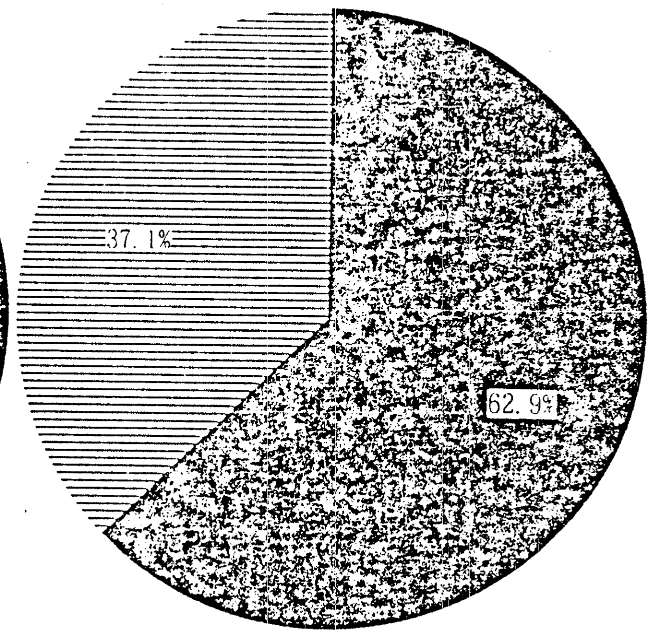

TOTAL

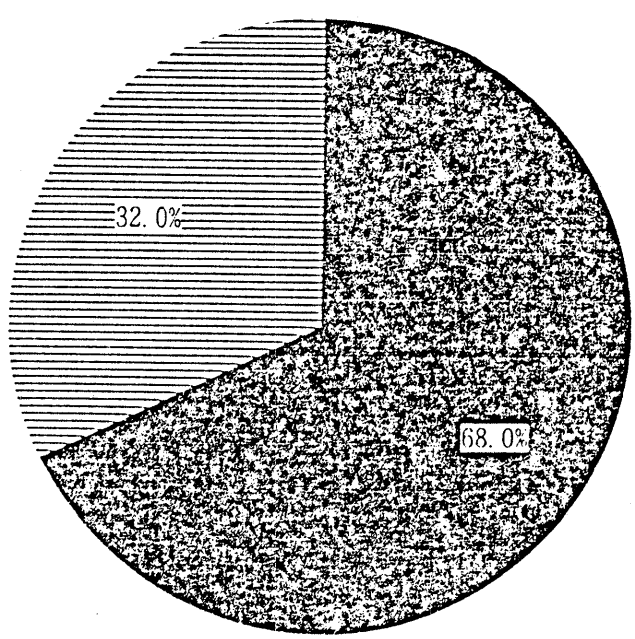

WIDE

North Pacific Networking

- SDN/PACNET(KR)
  - UUCP and IP
  - Phone, X28,9/X25, 9.6Kbps Leased
  - ISO-OSI Researches
  - ODP/X.400/VTAM/CASE/FTAM
  - will be on PACCOM soon..
- HARNET (HK)
- China, People's Republic of {CN}
  - CSNET: Karlsruhe -> Beijing
    - user%beijing%ira.uka.de@relay.cs.net
    - 1200bps Kermit link to Vienna in High Energy Physics Institute

3.9 "White Pages Pilot Program"

Presentation by Marshall Rose/ NYSERNet

From a posting to the TCP-IP mailing list on December 20, 1989

Back in July, NYSERNet started a White Pages Pilot Project using X.500 over TCP/IP as the underlying technology. At the three month mark last October, we hit nearly 100K entries at approximately 30 sites, about half of these were NYSERNet sites. During the last three months, we (NYSERNet and University College London) have spent a lot of effort making the software more robust, performant, and usable, based on our initial experiences. Well, as we enter the next three months, I'd like to extend an invitation to Internet sites in the US and CA to join our pilot. Here are the details:

1. You will need to run your own Directory Service Agent (DSA). This should run on just about any 4BSD-derived platform, although the recommended platform is a Sun-3 or Sun-4. You will need 30MB free disk for sources and executables. In addition, for each person you intend to have registered, the DSA will require approximately 1K of primary memory. (Yes, the DSA keeps entries resident in core, does its own memory management, etc., etc., there are obscure technical reasons for this.) I'm the first to admit that the memory requirement is "noteworthy", but just think of it as the price of admission.
2. The machine you run your DSA on will have to be on the Internet (direct IP access) and your organization must reside in the United States or Canada. The Canadian DMD (Directory Management Domain) is still being set-up at the University of Toronto, but should be operational before year's end. If there is an IP-connected site in Mexico, contact me: I'd like to get c=MX up and running sometime. It would be nice to offer White Pages over dial-up or something, but no dice. Think of the IP-connectivity requirement as another price of admission.
3. You will need to be able to devote time to installing the software and maintaining it. You will also need to check on your DSA regularly (i.e., once each morning) just to see that things are fine. In addition, if users at your site need help, you will be the first point of contact. This really isn't such a drain, considering that if you're the PostMaster at this site, you perform the exact same functions already.

So, after committing all this what do you get?

Well, if you want a "hype" answer:

- you get to join a large distributed information service which is administered by different organizations;
- you get to take part in the first production-quality field test of the OSI Directory (X.500);

- you get to take part in the first large scale production application of OSI technology on top of the TCP/IP suite of protocols; and,
- you get to add this experience to your resume, which will look quite good.

But, if you want the real answer:
You get to offer an exciting new service to your users. White Pages is just one of many applications you can host on top of the OSI Directory. By getting the Directory installed at your site, you are bootstrapping yourself to support the next generation of applications which need Directory Service, e.g., MHS (X.400).

Besides, it's fun to run the White Pages software to track people down, display their photos, find out their favorite drink, etc.

For more information, use anonymous FTP to host nisc.nyser.net, and retrieve the file: pilot/src/pilot-ps.tar.Z in BINARY mode. This contains a compressed tar image of several postscript files containing four documents: an introduction, an Admin Guide, a User Manual, and a presentation. Print these. The Admin Guide says how to get the software.

NYSERNet WHITE PAGES
PILOT PROJECT

Marshall T. Rose
NYSERNet, Inc.

October 26, 1989

OUTLINE

PART I: THE WHITE PAGES

PART II: ACCESSING THE SERVICE

PART III: THE FUTURE

SUMMARY

- A LARGE DISTRIBUTED INFORMATION SERVICE INVOLVING ADMINISTRATION BY DIFFERENT ORGANIZATIONS
- THE FIRST PRODUCTION-QUALITY FIELD TEST OF THE OSI DIRECTORY (X.500)
- THE FIRST LARGE SCALE PRODUCTION APPLICATION OF OSI TECHNOLOGY ON TOP OF THE TCP/IP SUITE OF PROTOCOLS

PART I

THE WHITE PAGES

- INTRODUCTION
- RELATION TO OSI DIRECTORY
- REALIZING THE WHITE PAGES SERVICE

INTRODUCTION

- NETWORKS PROVIDE THE INFRASTRUCTURE BETWEEN USERS
- NEED INFRASTRUCTURAL INFORMATION TO FACILITATE INTERACTIONS

  e.g., E-MAIL ADDRESSES

- WHITE PAGES CONTAIN

  INFRASTRUCTURAL INFORMATION

WHITE PAGES IN THE REAL WORLD

- THE TELEPHONE BOOK IS THE BEST EXAMPLE
- MANY PROVEN FEATURES:

  MULTIPLE TYPES OF INFORMATION
  (USEFUL IN FINDING THE "RIGHT" ENTRY)

  YELLOW PAGES KEYED BY BUSINESS SERVICE

  LOCALITY OF INFORMATION

  DIRECTORY ASSISTANCE
  (IMPRECISE MATCHING)

WHITE PAGES IN THE COMPUTER WORLD

- CONTAINS TELEPHONE BOOK INFORMATION
  - ALONG WITH LOCAL "PHONE" INFORMATION
  - SUGGESTS BOTH LOCALITY AND ACCESS CONTROL
- CONTAINS NETWORK-SPECIFIC INFORMATION
  - E-MAIL ADDRESSES
  - PRIVATE MAIL
  - NETWORK MANAGEMENT
- ULTIMATELY: "THE" REPOSITORY OF ALL SYSTEM AND NETWORK ADMINISTRATIVE INFORMATION

A SMALL DISTINCTION

- WHITE PAGES IMPLIES SEARCH BASED ON <u>NAME</u>
- YELLOW PAGES IMPLIES SEARCH BASED ON <u>ATTRIBUTES</u>
- NETWORK SERVICE HAS FEATURES OF BOTH
  - PERHAPS <u>RAINBOW PAGES</u>™ IS BETTER TERM

RELATION TO OSI DIRECTORY

- NEED THREE THINGS TO IMPLEMENT THE SERVICE
  - OSI INFRASTRUCTURE
  - OSI DIRECTORY
  - WHITE PAGES ABSTRACTION
- ISODE PROVIDES OSI INFRASTRUCTURE
- QUIPU PROVIDES OSI DIRECTORY
- NOW NEED TO SEE HOW DIRECTORY TECHNOLOGY INFLUENCES WHITE PAGES SERVICE

RAMIFICATION 1: UNIQUE IDENTIFICATION OF USERS

- EACH ENTRY IN THE WHITE PAGES IS IDENTIFIED BY A HANDLE
- ONE TO ONE MAPPING BETWEEN ENTRIES IN DIRECTORY AND WHITE PAGES
- USE DIRECTORY DISTINQUISHED NAME, e.g.,

```
c=US
    @o=NYSERNet Inc.
    @ou=Research and Development
    @ou=Western Development Office
    @cn=Marshall Rose
```

FOR WHITE PAGES HANDLE

RAMIFICATION 2:
SEARCHING THE WHITE PAGES

- SEARCHES OCCUR RELATIVE TO AN AREA

- INTERACTIVE MODEL

  FIRST, IDENTIFY AREAS LIKELY TO CONTAIN INFORMATION

  SECOND, SEARCH AREAS

- SINCE AREAS ARE JUST PARTS OF THE DIRECTORY, BOTH STEPS INVOLVE SEARCHING

- USER INTERFACE PROVIDES SIMPLE SYNTAX FOR DOING BOTH AUTOMATICALLY

RAMIFICATION 3:
STRUCTURE OF INFORMATION

- INFORMATION IS STRUCTURED USING ASN.1:

  FORMAL DEFINITION

  DIRECTORY ENFORCES SYNTAX

  USERS ENFORCE SEMANTICS

- IT'S ALL BINARY

  SO USER INTERFACE MUST SELECT PLEASING OUTPUT STRATEGY

  (UNLIKE MOST INTERNET-STYLE INTERFACES)

REALIZING THE WHITE PAGES SERVICE

- PILOT COMMUNITY IS NYSERNet MEMBERSHIP
  AND BROAD INTERNET COMMUNITY
- ISODE PROVIDES OSI INFRASTRUCTURE OVER TCP/IP USING RFC1006
- THE WHITE PAGES ABSTRACTION
  ADMINISTRATIVE DISCIPLINE
  USER INTERFACE

ADMINISTRATIVE DISCIPLINE

- BASED ON THE FOUR MODELS OF THE DIRECTORY
  INFORMATIONAL PERSPECTIVE
  FUNCTIONAL PERSPECTIVE
  ORGANIZATIONAL PERSPECTIVE
  SECURITY PERSPECTIVE

INFORMATIONAL PERSPECTIVE

- EACH ENTRY IN THE WHITE PAGES CORRESPONDS TO AN ENTRY IN THE OSI DIRECTORY
- SINCE DISTINGUISHED NAMES ARE HIERARCHICAL, SO ARE HANDLES IN THE WHITE PAGES
- ONLY LIMITED INFORMATION TYPES SUPPORTED
  - ORGANIZATIONS
  - ORGANIZATIONAL UNITS AND ROLES
  - LOCALITIES
  - PERSONS

ATTRIBUTE TYPES FOR PERSONS

commonName	physicalDeliveryOfficeName	stateOrProvinceName
description	photo	streetAddress
facsimileTelephoneNumber	postOfficeBox	surName
favouriteDrink	postalAddress	telephoneNumber
mobileTelephoneNumber	postalCode	title
otherMailbox	rfc822Mailbox	userPassword
pagerTelephoneNumber	roomNumber	userid

FUNCTIONAL PERSPECTIVE

- DUA CONTACTS DSA FOR INFORMATION
- IF DSA DOES NOT HAVE INFORMATION RESIDENT, IT EITHER

  CHAINS REQUEST TO A DSA CLOSER TO THE INFORMATION

  REFERS DUA TO A DSA CLOSER TO THE INFORMATION

- WHAT DOES RESIDENT MEAN?

ENTRY DATA BLOCK

- AN ENTRY DATA BLOCK (OR BLOCK) CONSISTS OF A SMALL PORTION OF THE TREE

  THE NAMES AND ATTRIBUTES OF THE IMMEDIATE CHILDREN OF OF A PARTICULAR NODE

- THREE KINDS OF BLOCKS

  SLAVE COPY: COMPLETE AND AUTHORITATIVE

  REGULARLY UPDATED FROM UPSTREAM DSA

  CACHE COPY: POSSIBLY PARTIAL INFORMATION

  INFORMATION DETERMINED FROM CHAINING

  INVALIDATED RELATIVE QUICKLY

  MASTER COPY: THE ORIGINAL

THE RESIDENCY REQUIREMENT

OPERATION REQUESTED	COPY REQUIRED FOR RESIDENCY
READ, COMPARE	MASTER, SLAVE, OR CACHE
LIST, SEARCH	MASTER, OR SLAVE
UPDATE	MASTER

- IN ORDER TO IMPROVE SEARCHING, SLAVE COPIES OF THE ROOT and c=US ARE KEPT AT EACH DSA

- UPDATES STILL RELY ON CENTRALIZED ENTITY

ORGANIZATIONAL PERSPECTIVE

- RESPONSIBILITY FOR INFORMATION DIVIDED INTO DIRECTORY MANAGEMENT DOMAINS (DMDs)

- LEVEL-0: HIGHLY-AVAILABLE AUTHORITATIVE SERVERS

  ROOT

  c=US

- LEVEL-1: AUTHORITATIVE SERVER FOR EACH ORGANIZATION

- LEVEL-2: OVERFLOW DSAs FOR AN ORGANIZATION

  USE NOT RECOMMENDED AT PRESENT

TOPOLOGY OF THE PILOT PROJECT

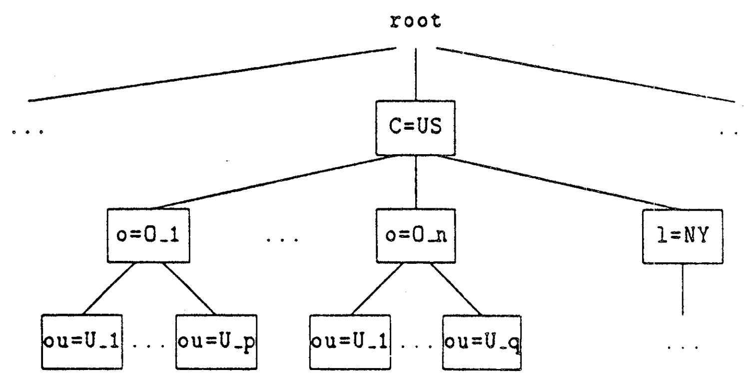

SECURITY PERSPECTIVE

- <u>SIMPLE</u> SECURITY MODEL (PASSWORD-BASED)

  WON'T TOUCH <u>STRONG</u> SECURITY MODEL WHICH USES PKS

- ACCESS CONTROL LISTS

  NONE, DETECT, COMPARE, READ, ADD, WRITE

  FOR ENTRIES, ATTRIBUTES, AND CHILDREN

- PASSWORDS MUST NOT BE REPLICATED OUTSIDE OF ORGANIZATION'S DMD

PART II

ACCESSING THE SERVICE

- FRED
- FACES
- MH

FRED

- BASED ON SRI-NIC WHOIS SERVICE
  - EXPERIENCE SHOWS SYNTAX IS WELL-LIKED
  - TRAINING PROBLEM REDUCED
- INTERACTIVE PROGRAM
  - ALSO AVAILABLE VIA NETWORK AND MAIL

WHOIS COMMAND

`whois input-field [record-type] [area-designator] [output-control]`

- PARTIAL NAME, e.g.,

    `rose`

- FULLY-QUALIFIED HANDLE, e.g.,

    `@c=US@cn=Manager` OR `!1`

- MAILBOX SPECIFICATION, e.g.,

    `mrose@nisc.nyser.net`

MATCHING RULES

- IF "*"-SIGN PRESENT, USE WILDCARDING, ELSE
- IF USER WANTS IMPRECISE MATCHING, USE SOUNDEX, ELSE
- IF USER WANTS SURNAME MATCHING, LOOK THERE, ELSE
- FORCE LIBERAL WILDCARDING, e.g., `*rose*`

ALTERNATE QUERY FORMS

- USE THE -title SWITCH WHEN LOOKING FOR PEOPLE
- FOR EXAMPLE,

  ```
  whois rose -title scientist
  ```

  LOOKS FOR SOMEONE NAMED rose WHO IS A SCIENTIST
- WHILE

  ```
  whois -title operator
  ```

  LOOKS FOR ANYONE WHO IS AN OPERATOR
- SEARCHES ARE RELATIVE TO THE APPROPRIATE AREA

AREA DESIGNATOR

- SAYS WHERE TO SEARCH, EITHER
- DIRECT REFERENCE, e.g.,

  "@c=US@o=NYSERNet Inc." OR !3
- INDIRECT REFERENCE, e.g.,

  -org nyser
- INDIRECT REFERENCE CAUSES IMPLICIT SEARCH TO DETERMINE LIST OF AREAS FOR SEARCH

AN EXAMPLE

```
fred> whois goodfellow -org anterior
Trying @c=US@o=Anterior Technology ...
Geoffrey Goodfellow (2)         Geoff@Fernwood.MPK.CA.US
     aka: Geoffrey S. Goodfellow

President
Anterior Technology
  POB 1206
  Menlo Park, CA  94026-1206

Telephone: +1 415 328 5615
FAX:       +1 415 328 5649
TELEX:     number: 650 103 7391, country: US, answerback: MCI UW

Mailbox information:
  MCI-Mail: Geoff
  Internet: Geoff@Fernwood.MPK.CA.US
  UUCP: fernwood!Geoff

Drinks:   chilled water
Picture:  /usr/etc/g3fax/Xphoto invoked

Handle:   @c=US@o=Anterior Technology@ou=Corporate@cn=Geoffrey Goodfellow (2)
Modified: Fri Jul 21 11:41:27 1989
```

FACES

- YOU CAN STORE ARBITRARY DATA IN THE WHITE PAGES
- ONE ATTRIBUTE IS A FACSIMILE IMAGE CALLED photo
- THERE ARE TWO X WINDOWS PROGRAMS WHICH DISPLAY THIS INFORMATION

  XFACE: WHEN READING A MESSAGE WITH MH, DISPLAYS PHOTO

  XWHO: LIKE RWHO, BUT WITH PHOTOs

MAPPING TO HANDLES

- MAPPING local@domain to A DISTINGUISHED NAME IS A PROBLEM
- COULD USE SEARCH ON rfc822Mailbox ATTRIBUTE, BUT HOW TO LIMIT SEARCH?
- ON LOCAL AREA NETWORK (XWHO), PROBLEM IS SIMPLER
  - SEARCH LOCAL PART OF TREE FOR userid ATTRIBUTE
- FOR E-MAIL (XFACE), USE DIRECTORY TO PROVIDE INVERSE MAPPING TO DOMAIN NAMES
  - COUNTRIES, ORGANIZATIONS HAVE domainRelatedObject IN OBJECT CLASS
  - OBJECTS OF THIS CLASS HAVE associatedDomain ATTRIBUTE
  - USE A RECURSIVE ALGORITHM TO DERIVE DI-SUBTREES LIKELY TO HAVE DESIRED INFORMATION

MH

- WHEN COMPOSING MAIL, IT WOULD BE NICE TO USE THE WHITE PAGES TO GET E-MAIL ADDRESSES
- MH IS MODIFIED TO USE FRED FOR THIS PURPOSE

MH USE OF FRED

- SPECIFY A NAME BY USING "{" AND "}" INSTEAD OF AN ADDRESS, e.g.,

  ```
  To: { rose -org nyser }
  ```

- SEND AND WHOM COMMANDS WILL EXPAND QUERY

  USER IS ASKED TO CONFIRM/REFINE

- NOT AVAILABLE FOR PUSH

  (MUST BE INTERACTIVE USE OF MH)

PART III

THE FUTURE

- LEVEL-1 SLAVE DSAs
- DEVELOPMENT
- NAME CHANGES

LEVEL-1 SLAVE DSAs

- LEVEL-1 DSAs ARE A SINGLE POINT OF FAILURE
- START LOOKING FOR PEER TO RUN LEVEL-1 DSA FOR YOUR DATA

DEVELOPMENT

- WHENEVER MAINTENANCE COMES UP, DEVELOPMENT GOES OUT THE DOOR!
- SO, MAJOR DEVELOPMENT IS OVER

SOME AREAS OF INTEREST

- PRIVATE MAIL (KEY RETRIEVAL)
- USER-DEFINED TEMPLATES FOR FRED OUTPUT
- DOCUMENT SEARCHING

NAME CHANGES

- PUTTING EVERYONE UNDER c=US WON'T SCALE
- SO, UNLESS YOUR ORGANIZATION IS INTER-STATE, PLAN ON

  c=US@l=NY@o=O_i

- ALSO PLAN ON HAVING ONLY INCOMPLETE INFORMATION ON c=US

CONCLUSIONS

- WHITE PAGES SERVICE

  USING THE OSI DIRECTORY

  RUNNING IN A TCP/IP-BASED INTERNET

  SOUNDS UNLIKELY!

- ACTUALLY, IT'S OPTIMAL FOR A PILOT

  SHEER SIZE OF THE INTERNET

  INTERNET SUITE OF PROTOCOLS FOR STABLE END-TO-END SERVICES

  OSI DIRECTORY SERVICE FOR RICH SERVICE ENVIRONMENT

3.10 "NSFnet Status Update"

Presentation by Bilal Chinoy/Merit

The NSFnet backbone reconfiguration was complete by July, 1989. The new topology has 19 T1 links, as opposed to 14 in the previous topology. This increased the network redundancy and performance characteristics. In addition, there are no single connected nodes.

Traffic through the backbone continues to increase, and the backbone switched approximately 1.59 billion packets in September, 1989.

Application distribution (in packets) shows a consistent trend. Interactive traffic and File Transfer Protocol continue to be the dominating applications.

We now have direct operational connections to the Arpa/Milnet, through T1 connections on both coasts. The east coast connection is between the NSS at College Park and the Mailbridge at Mitre. The west coast connection is between the NSS at Palo Alto and the Mailbridge at NASA Ames. This has dramatically improved the reliability and quality of the NSFnet - DDN connectivity.

The NSFnet research network was upgraded to 5 nodes and 7 T1 links, going from 4 nodes and 4 T1 links.

Cylink ACSU's are being tested on the research network, for early 1990 deployment on the operational network.

Border Gateway Protocol (BGP) testing continues on the research network.

The NSFnet demonstration at InterOp '89 was well received. We demonstrated OSI (CLNP) switching through the NSS's at San Jose, and on the research network.

In addition, the NSS at San Jose also was fully connected to the show FDDI ring. This demonstrated the NSS FDDI functionality.

RFC 1105 on the BGP protocol was announced. In addition, draft RFC's on BGP usage, environment, etc., and on routing models, also are being finalized.

All the responses from the NSFnet expansion solicitation have been received by NSF.

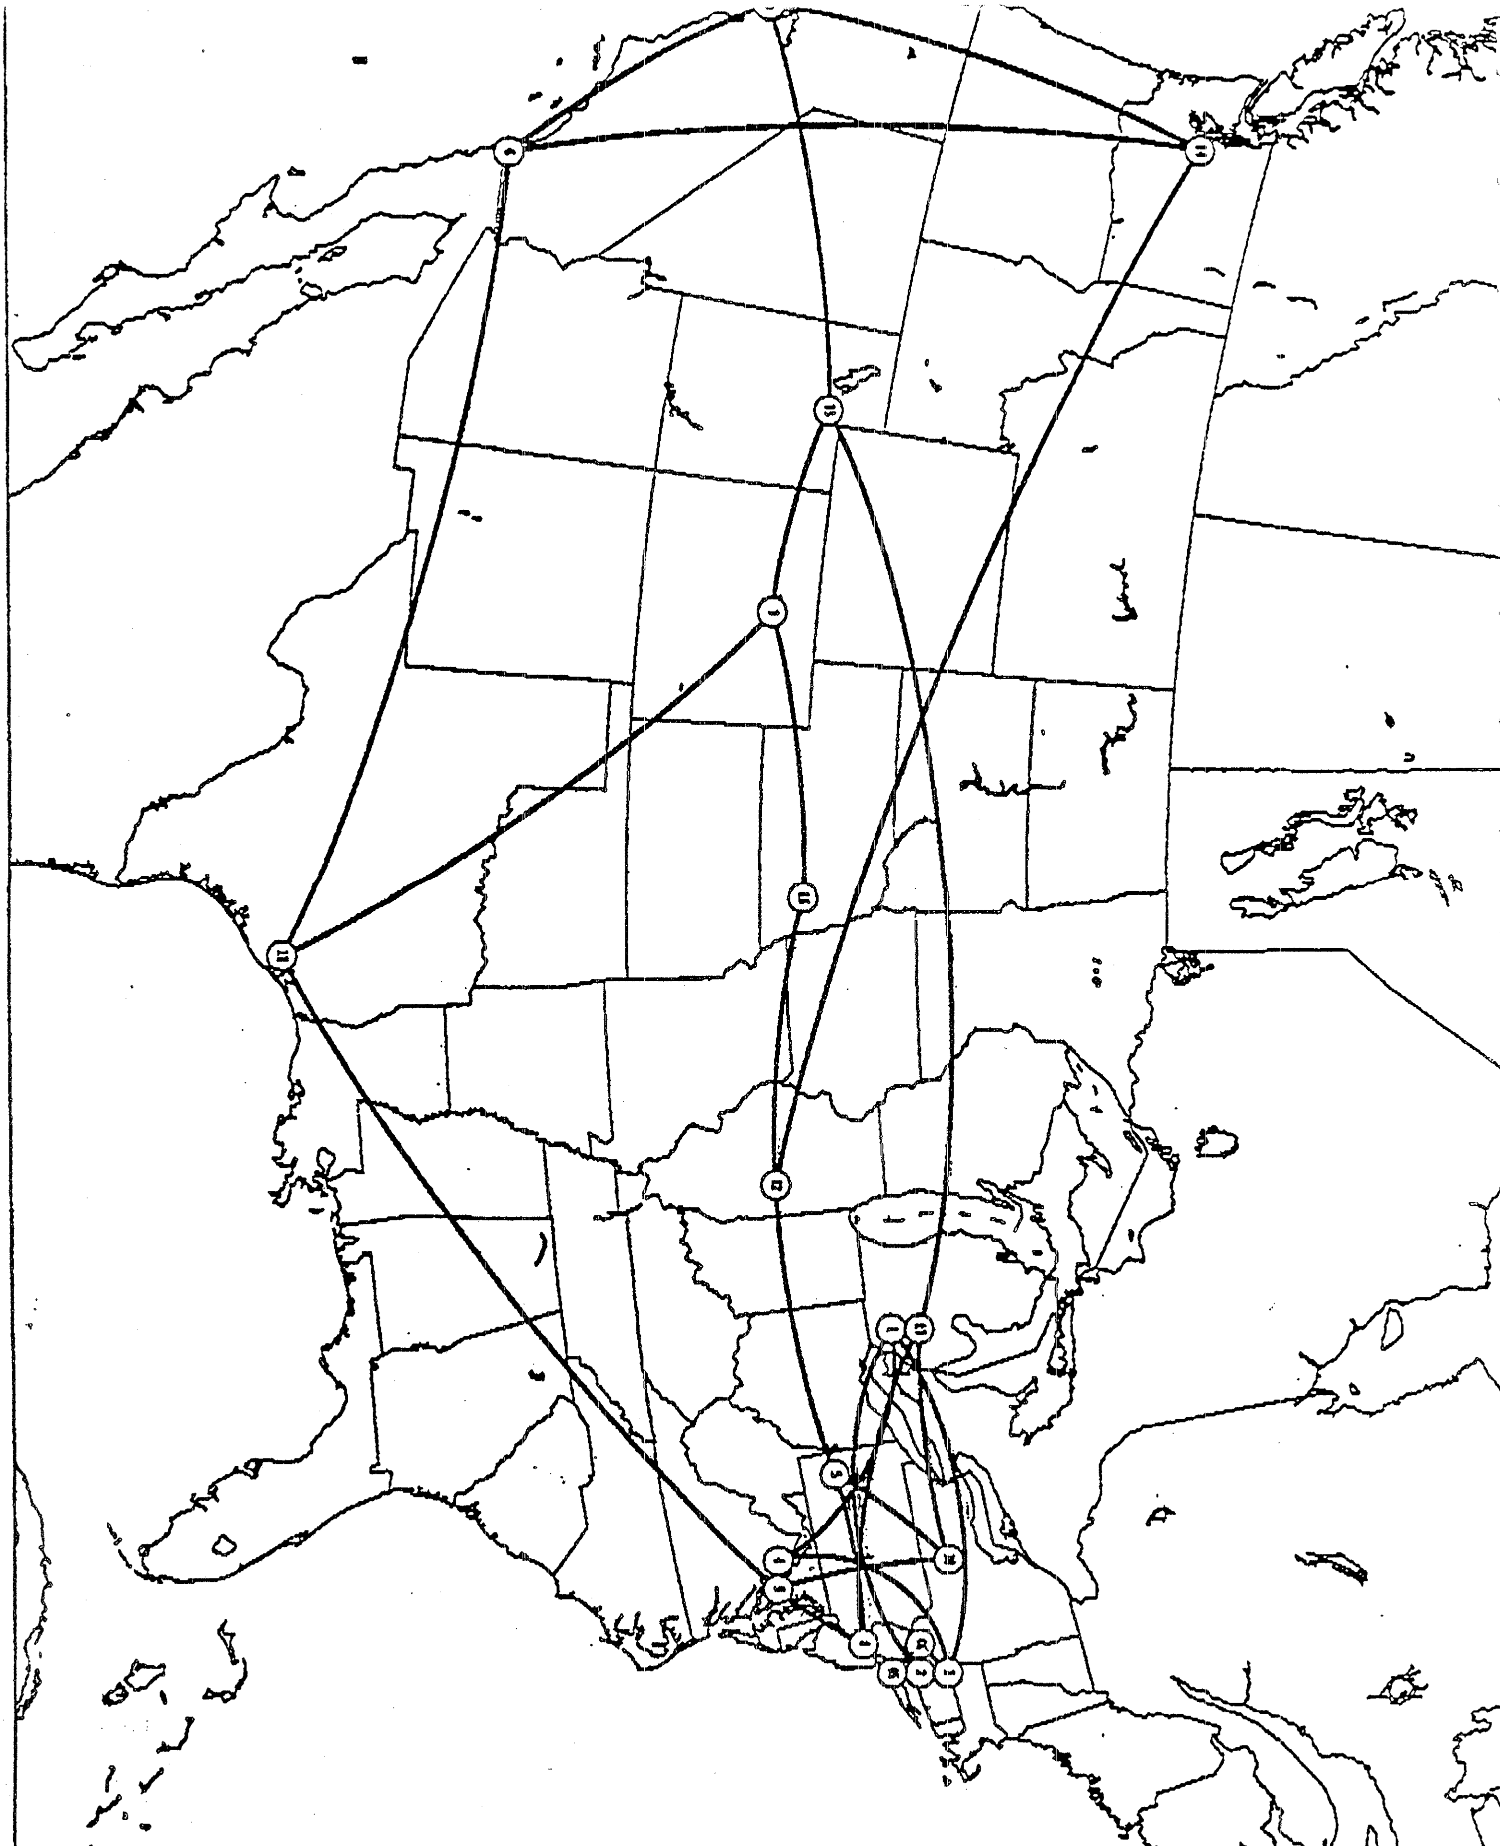

New NSFNET Backbone

Circles contain NSS number (1-4 and 45-48 are test nodes)

Prepared by NSFNET-info@merit.edu at Thu Apr 6 18:57:09 1989
netmap-1.5 program by Brian Reid, map data from World Data Bank II
Lambert Conformal Projection [44°N,33°N], Map center: [40°N, 96° 30'W]
Image resolution 300/in., stroke limit 1 pixels

National Science Foundation Network (NSFNET)
Monthly Network Traffic Statistics

NSFNET is a cooperative project of Michigan's Merit Computer Network, the National Science Foundation, IBM, and MCI Telecommunications.

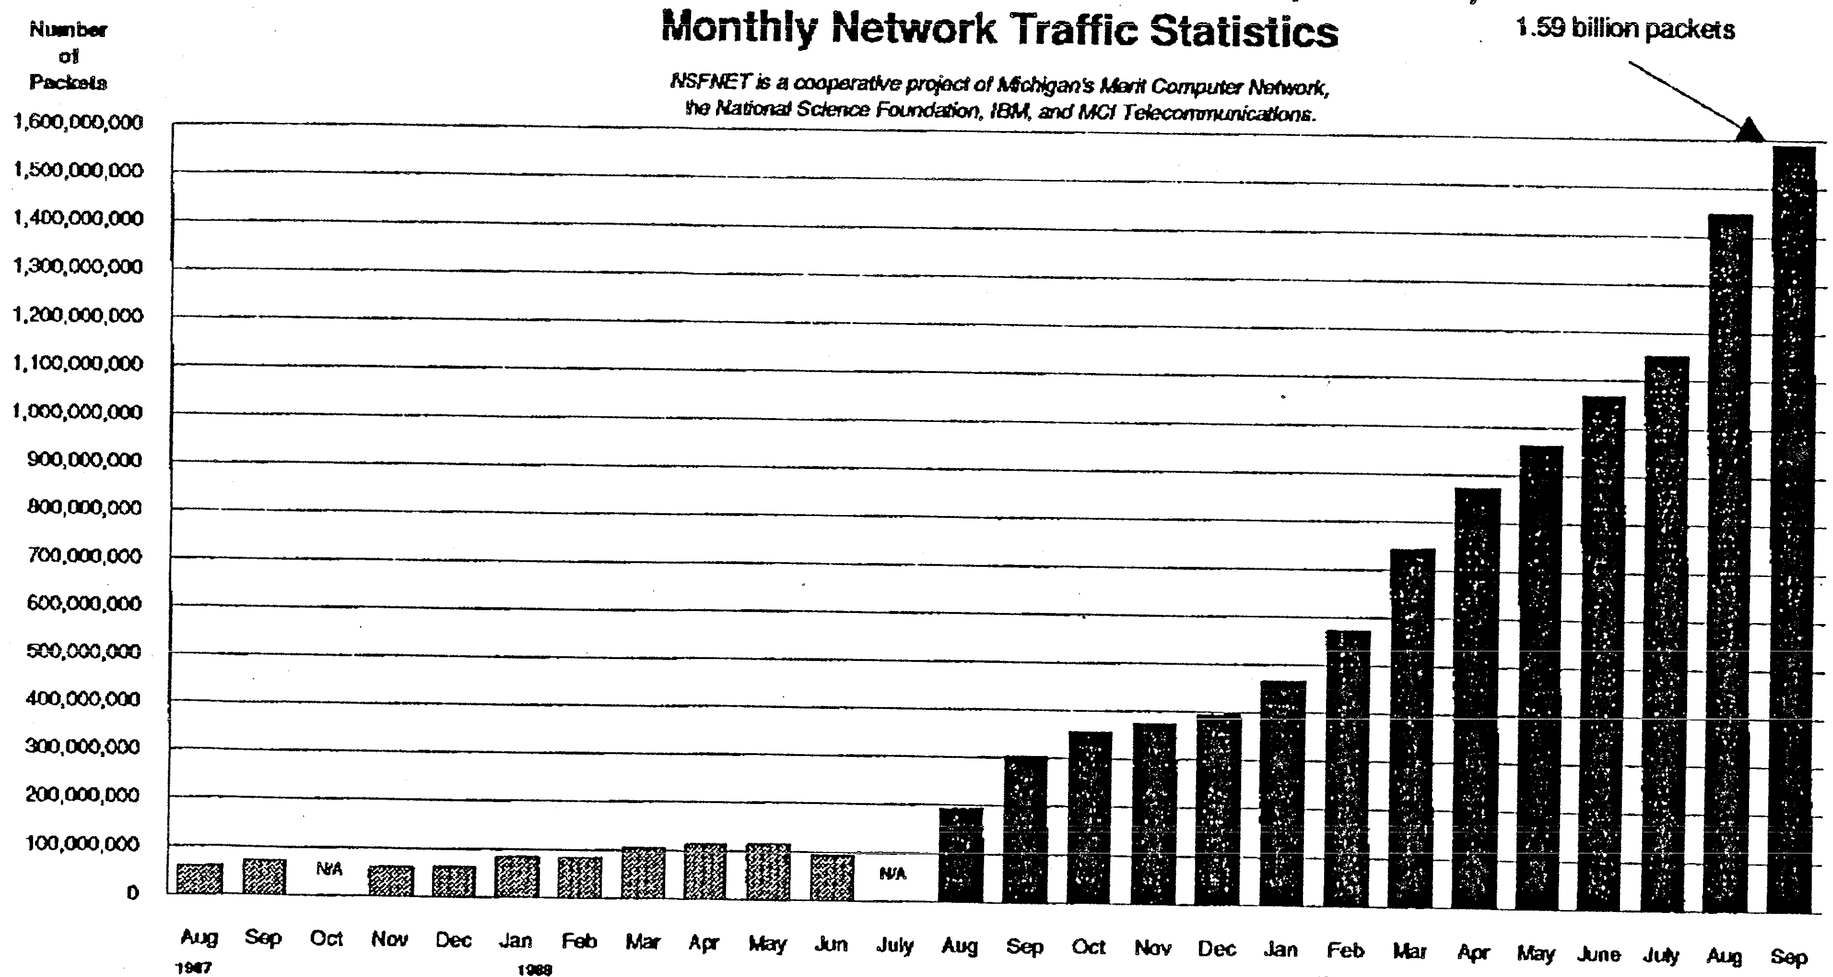

APPLICATIONS DISTRIBUTION (Oct 1 -18)

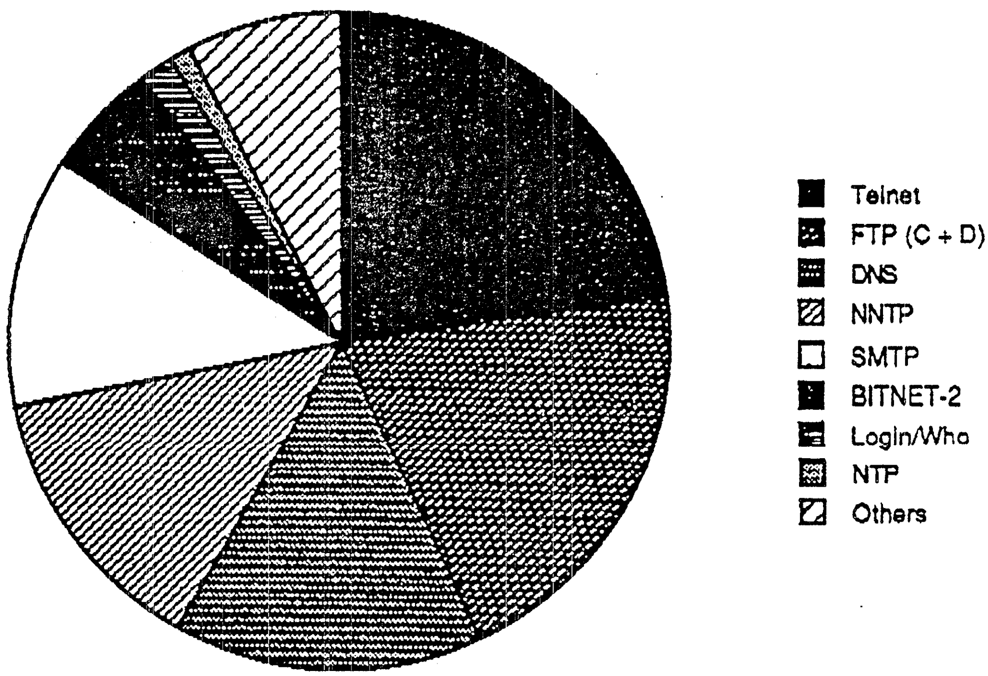

bac

NSFNET/NSN/DDN connection at the University of Maryland

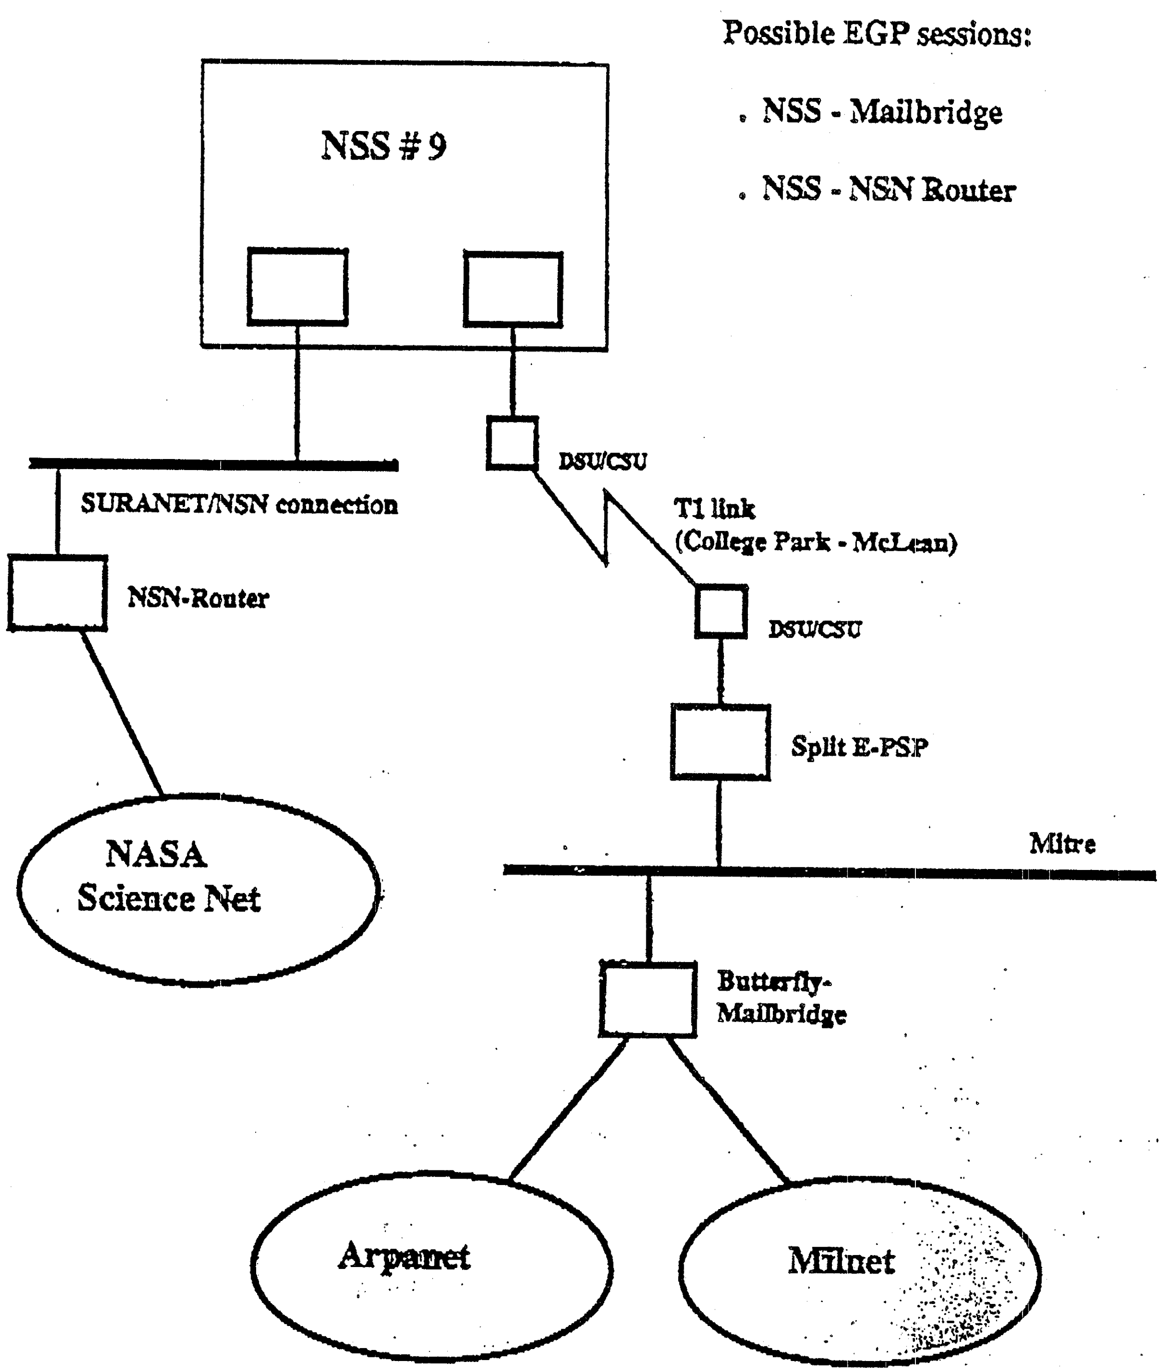

RESEARCH NETWORK

- 5 NODES, 7 LINKS
- ACSU PHASEOVER
- BGP TESTING

NSFNET Research Network

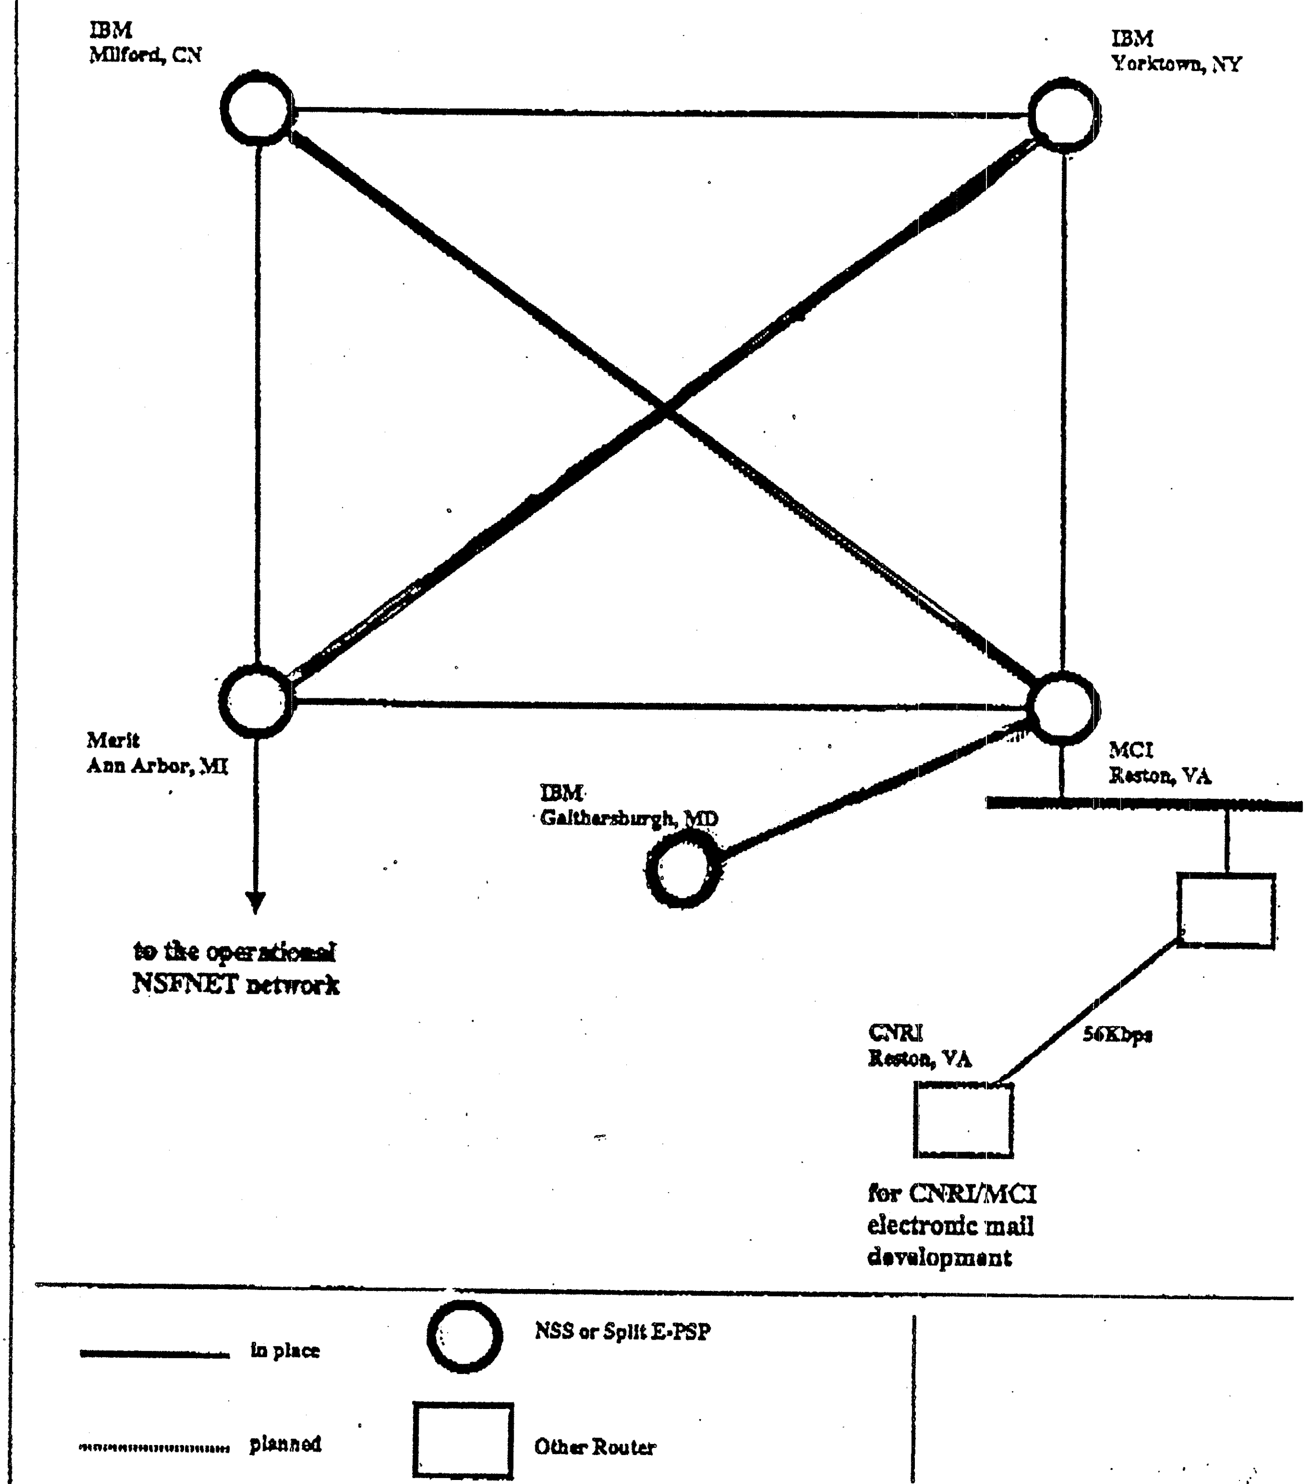

NSFNET
InterOp '89

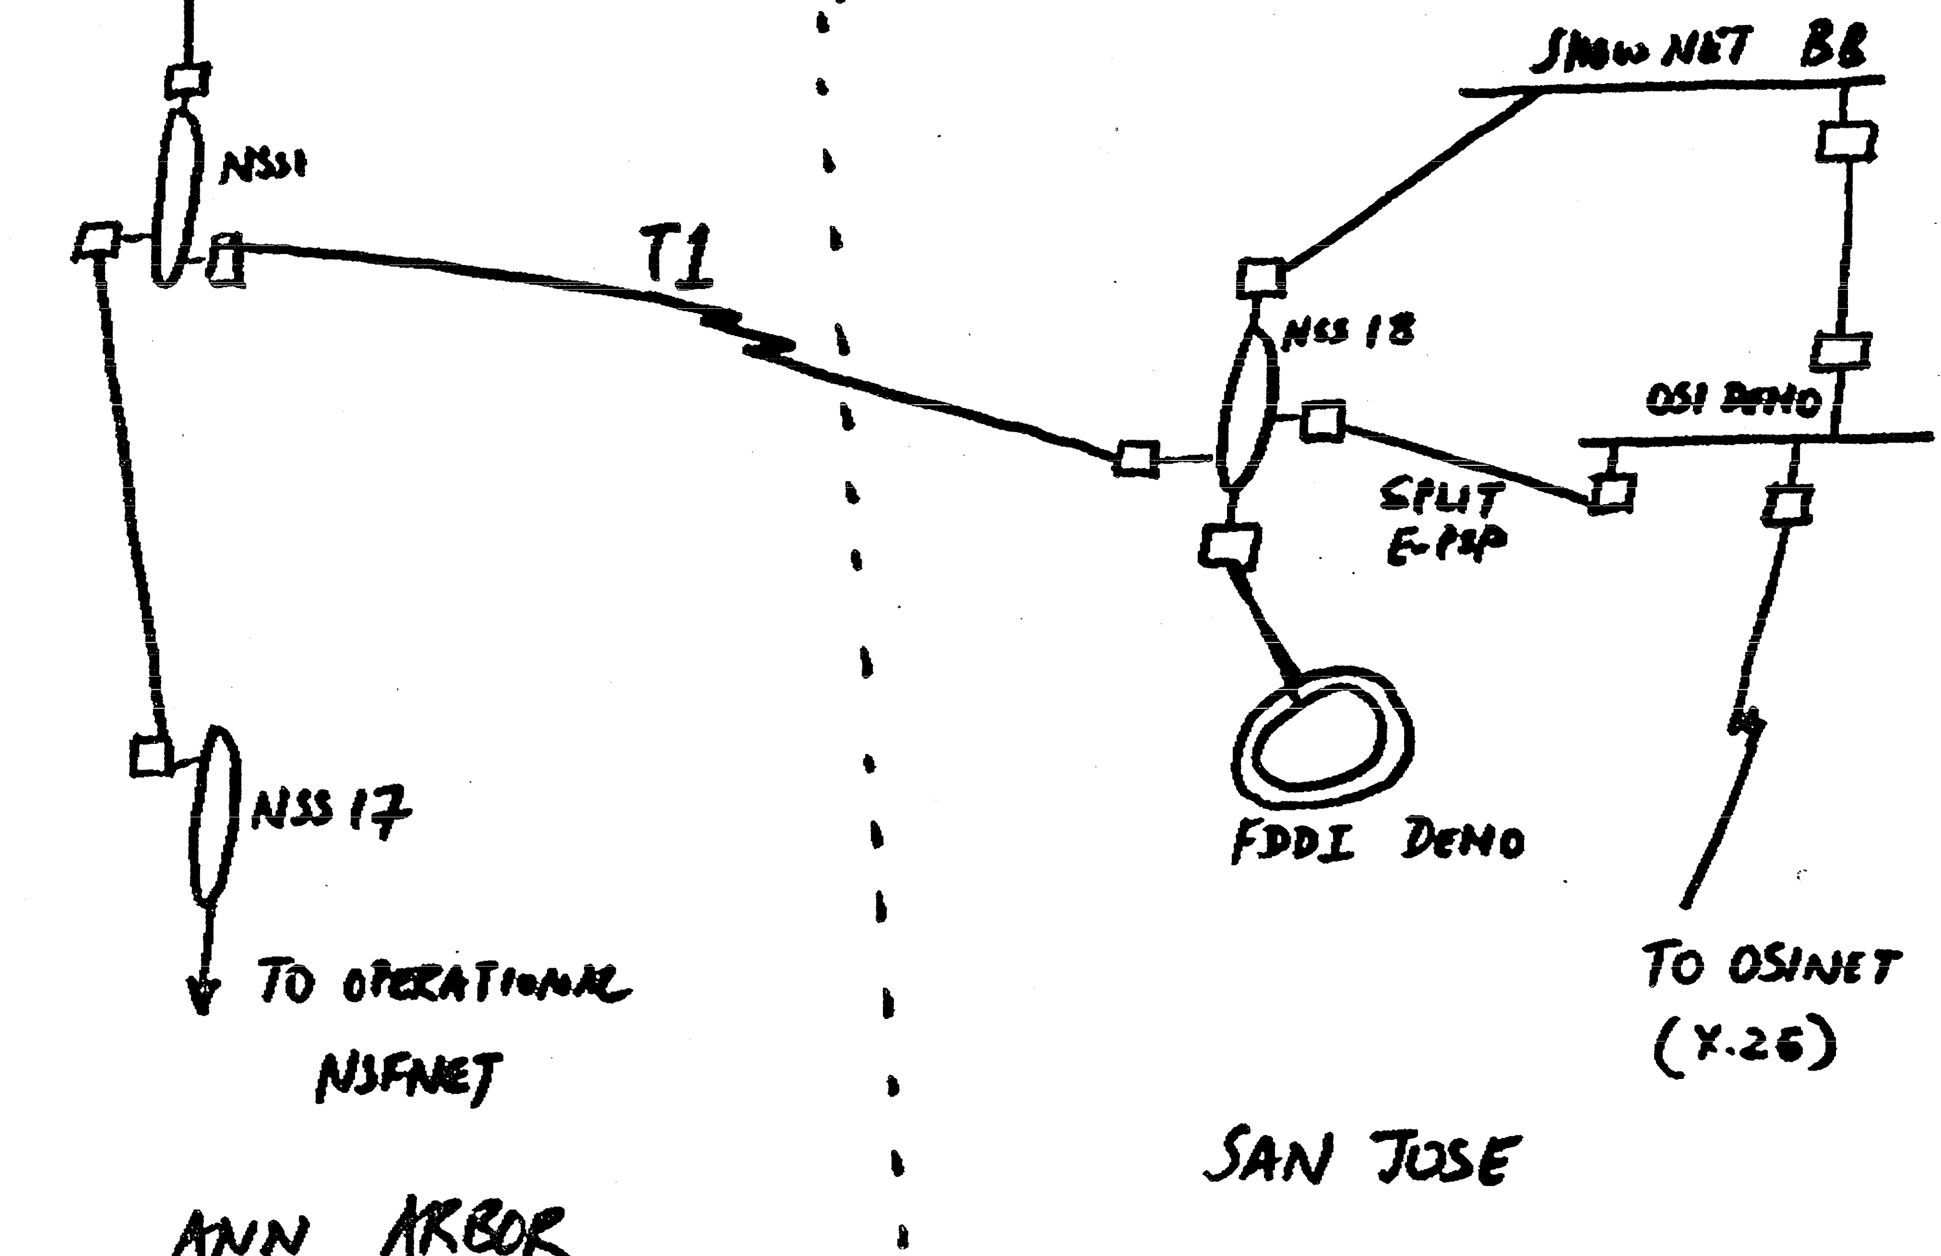

BORDER GATEWAY PROTOCOL

- RFC 1105 - PROTOCOL
- DRAFT ON USAGE, ENVIRONMENT, ETC.
- DRAFT ON AD/RD's, etc.
- 3 IMPLEMENTATIONS
  - NSS
  - CISCO
  - GATED
- NEW WORKING GROUP
- EXTENSIVE TESTING

NSFNET EXPANSION

- ALL RESPONSES IN

- TARGETED FOR FIRST HALF OF 1990

3.11 "Routing and Fair Pricing in Internets with Packet Loss"

Presentation by Vlad Rutenburg/ SRI International

The problem of fair charging mechanisms for interdomain communication with packet losses represents the an important issue in policy based routing. We present a mechanism that assigns to individual domains fair wages for successful delivery of packets and fair penalties for losing packets. This method makes sure that each domain breaks even in the long run.

Suppose a packet travels along a path P from source to destination through networks $N_1, N_2, ..., N_k$. Let c_i denote the cost to network N_i of handling a packet, and p_i denote the frequency with which network N_i loses messages. Let e_i denote the fair tariff that network N_i should charge to the sender for successfully transitioning a packet, and let u_i denote the penalty that network N_i should pay to the sender for losing a packet. Also let $E_{i-1} = e_1 + e_2 + \cdots + e_{i-1}$ denote the total cost accumulated by a packet entering network N_i. We have derived that in order to be fair, the pricing policy should be as follows: If region N_i loses the packet, it should reimburse the source in the amount u_i equal to E_{i-1}. If region N_i successfully transitions the packet, it should be paid the amount $e_i = (p_i E_{i-1} + c_i)/(1 - p_i)$. Notice that the network needs to charge more than the handling cost, because it has to cover the "cost of loss insurance". Thus, the above tariffs reflect both the delivery costs and the risks of losses in an economically optimal way.

If we add all the costs e_i for all the regions, the total cost for successful packet delivery from source to destination along the path P is $J_P = (c_1 + c_2 R_1 + \cdots + c_k R_{k-1})/R_k$, where $R_i = (1 - p_1)(1 - p_2)(1 - p_3) \cdots (1 - p_{i-1})(1 - p_i)$. Based on this new cost measure, one can now find the path that minimizes the value of J_P. Even though the commonly used efficient destination-based versions of algorithms (Dijkstra, Bellman-Ford) for computing optimal routing do not work with respect to the new measure, the equally efficient source-based versions of these algorithms, do work, providing an efficient way for computing optimal routing between different source-destination pairs.

An important aspect of the above results is that even if cost accounting is not part of the interdomain policy, the above measures and algorithms should still be used in order to minimize the average global resource cost of successful packet delivery, as has been proved by Rutenburg and Shacham. Namely, given a path P in G, the expected total system cost with retransmission until successful delivery is equal to $(c_1 + c_2 R_1 + \cdots + c_k R_{k-1})/R_k$, which is exactly equal to the path cost J_P derived in the previous paragraphs. Thus, the optimal routes with respect to J_P also minimize the global resource use.

OBJECTIVE

- To develop fair charging and accounting mechanisms for unreliable internetworks (i. e., with packet losses)
- To design efficient algorithms for computing the best packet routing in such internetworks

CHARGING AND ROUTING IN RELIABLE NETWORKS (REVIEW)

PRICING:

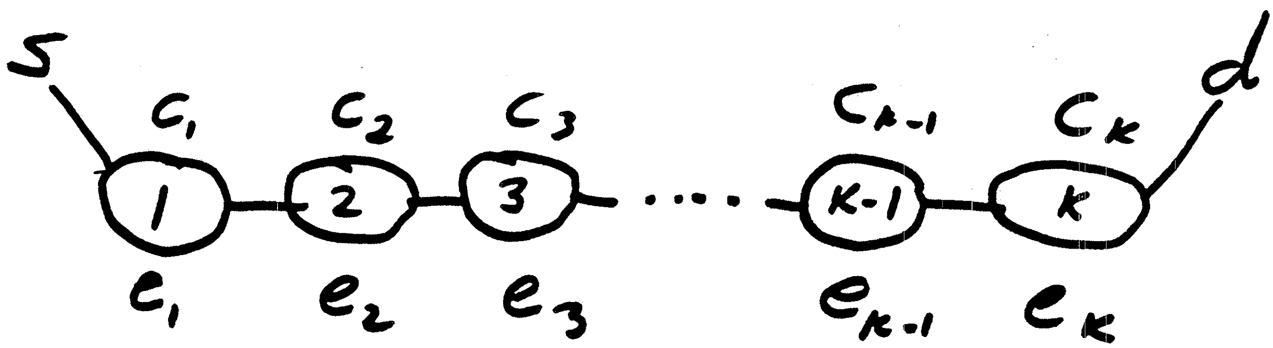

e_i - the price of using autonomous region (network) i

- If want to be fair - set the prices e_i to the costs c_i of delivery

ROUTING:

- Compute the least expensive path from every node (region) v to the fixed destination node (region) d

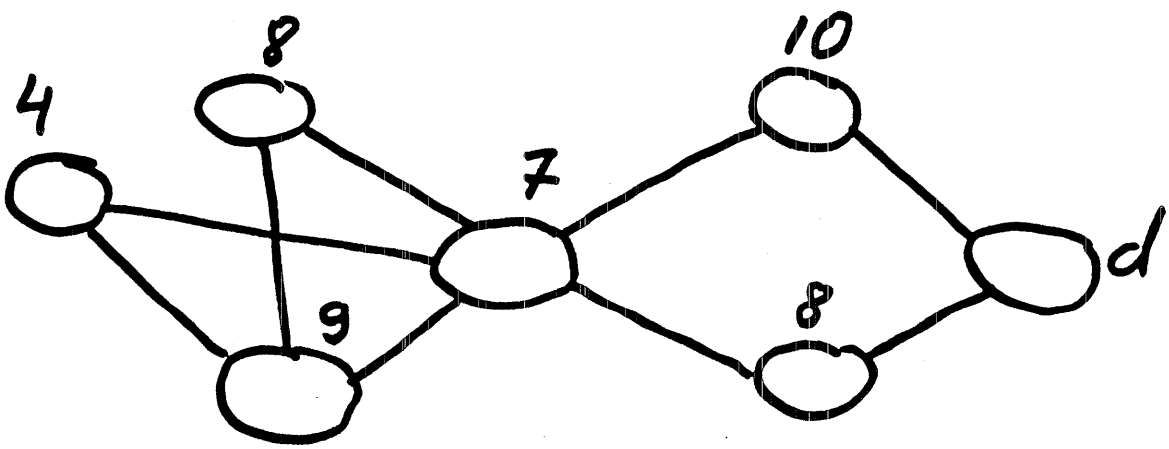

- Can be solved using Dynamic Programming (DP) algorithms, because of the memoriless property: the portion of the optimal path after some intermediate node w is independent of the path before w

- Can be solved fast and efficiently:
  - link state - SPF (Dijkstra) algorithm
  - distance vector - Bellman - Ford algorithm

ACCOUNTING IN UNRELIABLE NETWORKS

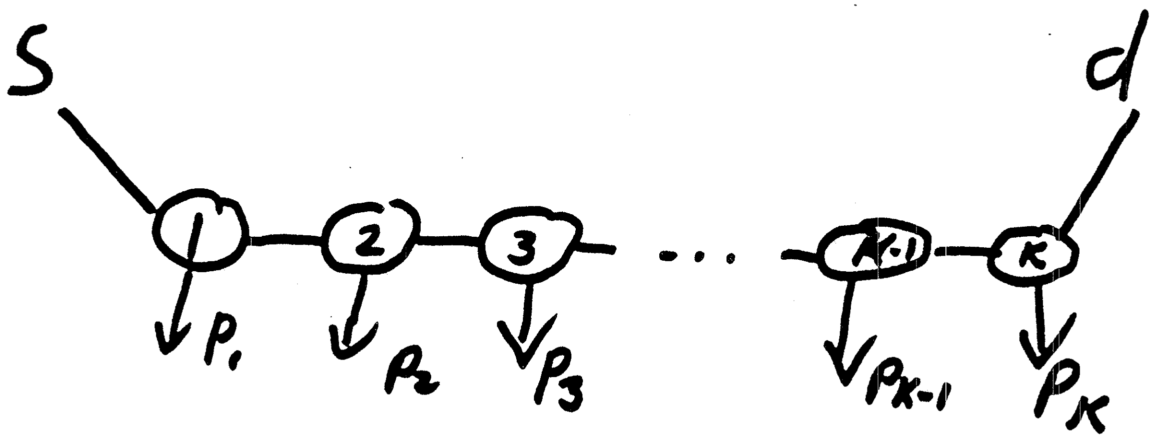

p_i - the probability that a message will be lost inside region i

Q: If message is dropped in region 3, who should pay and how much?

A: Region 3 should pay e_1 to region 1 and e_2 to region 2 and absorb its own expenses

- Fair Pricing: need to set prices that allow each region to break even in the long run

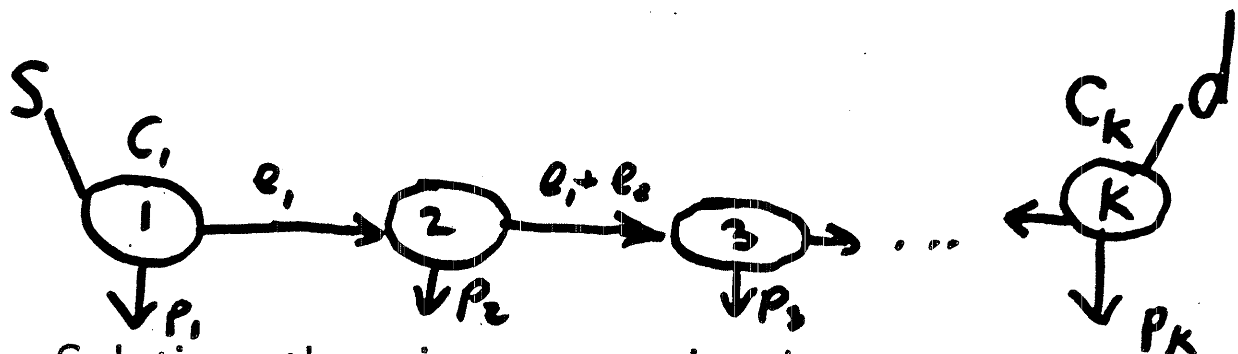

- Solution: the price e_i at region i

$$e_i = \frac{p_i E_{i-1} + c_i}{1 - p_i},$$

$E_{i-1} = e_1 + e_2 + \cdots + e_{i-1}$ - total accumulated cost

- Can be viewed as: a region needs to charge more than the handling cost in order to cover the "cost of disaster insurance"

- The total cost from source to dest is

$$J_P = \frac{c_1 + c_2 R_1 + \cdots + c_k R_{k-1}}{R_k},$$

$$R_i = (1 - p_1)(1 - p_2)(1 - p_3) \cdots (1 - p_{i-1})(1 - p_i)$$

ROUTING:

- No longer memorilessness: the choice of the future portion of the path depends on the total accumulated cost

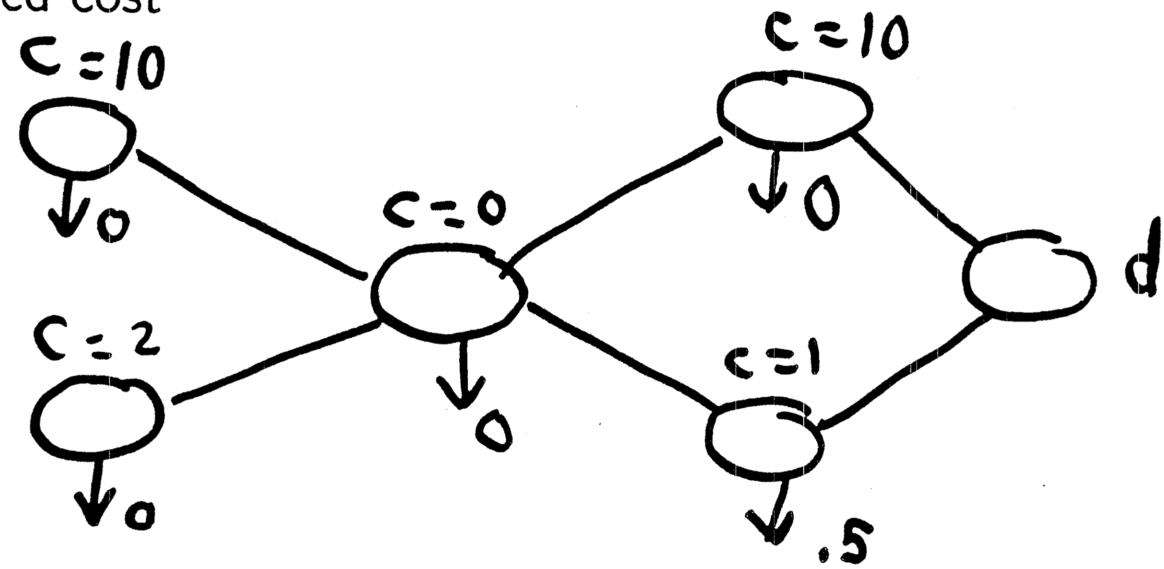

- Cannot run the old DP algorithms with respect to a fixed destination

- Solution: Solve for a fixed source, not destination:

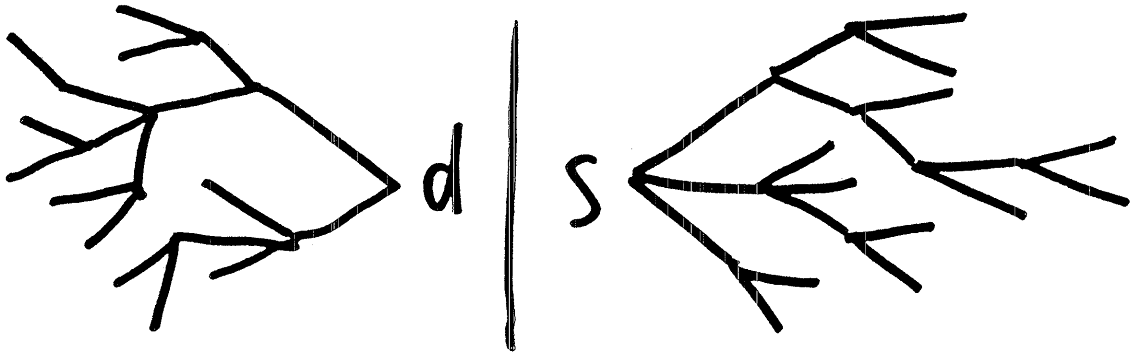

- Approach: Given internetwork G, and a fixed source node (region) s, for every node v in G, find the least expensive path from s to v
- Algorithm: Compute optimal value $E_{source,v} = \min_P E_P$ for all i using update formula

$$E_{source,v} = \frac{\min_{j \in N(v)} E_{source,j} + c_v}{1 - p_i}$$

- Can be solved using Dynamic Programming algorithms, because of the memoriless property: the portion of an optimal path before some intermediate node w is independent of the path after w

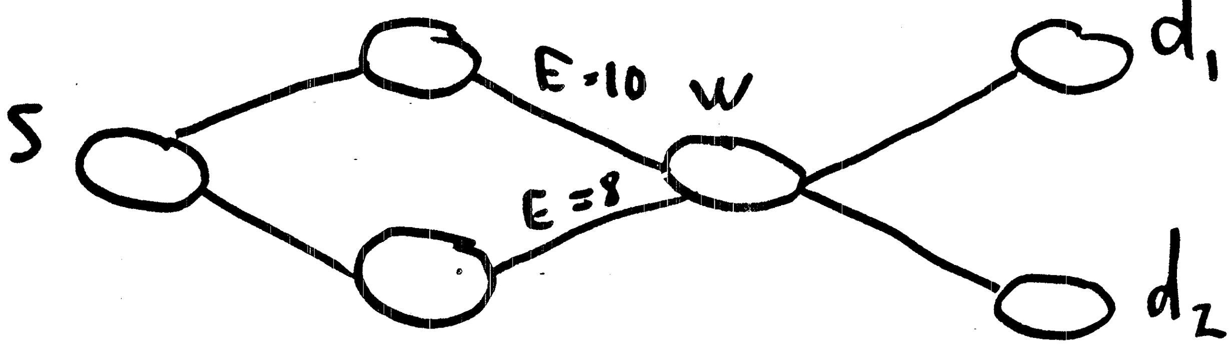

- Again, can be solved fast and efficiently using the "reversed" versions of:
  - link state - SPF (Dijkstra) algorithm
  - distance vector - Bellman - Ford algorithm
- If need to find optimal paths between all the pairs of nodes, can run the algorithm n times: for each source separately (or in parallel)

DISCUSSION

- An "economically" proper and simple pricing mechanism for unreliable internetworks presented
- This pricing mechanism properly balances the costs and the risks
- The optimal routing can be computed using source-based versions of the simple and efficient shortest path algorithms
- Even if money - not an object, still should use the above algorithms to compute the routes that deliver (with retransmission) messages at the lowest expected costs from the global point of view:

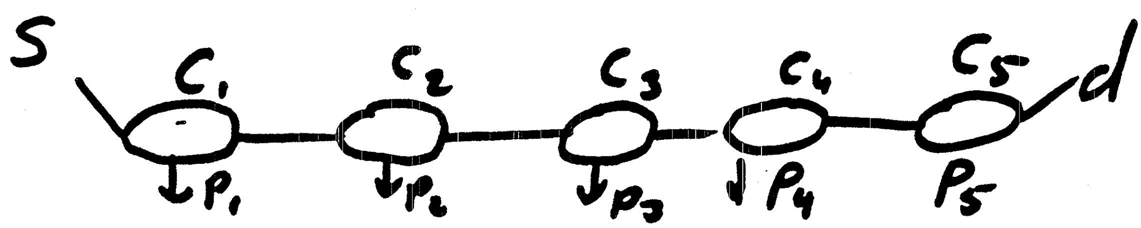

Theorem: Given a path P in G, the expected total system cost until successful delivery is

$$J_P = \frac{c_1 + c_2 R_1 + \cdots + c_k R_{k-1}}{R_k}$$

Appendix A

Attendees

Guy Almes

Rice University
PO Box 1892
Dept of Computer Science
Houston, TX 77251-1892
713-527-6038
almes@rice.edu

Philip Almquist

Consultant
214 Cole Street
San Francisco, CA 94117
415-752-2427
almquist@jessica.stanford.edu

Karen Armstrong

San Diego SCC
PO Box 85608
CERFNet
San Diego, CA 92138-5608
619-534-5077
armstrong@sds.sdsc.edu

Susan Arnold

San Diego Supercomputer Center
CERFnet
PO Box 85608
M/S SDSC
San Diego, CA 92138
619-534-5047
susie@sds.sdsc.edu

Douglas Bagnall

Hewlett Packard
Apollo Division
330 Billerica Road
CHR 03 DC
Chelmsford, MA 01824
508-256-6600
bagnall_d@apollo.hp.com
mit_eddie!apollo!bagnall_d

Fred Baker

Vitalink
6607 Kaiser Drive
Fremont, CA 94555
415-794-1100
baker@vitalink.com

Art Berggreen

Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
805-963-9431
art@sage.acc.com

Rick Boivie

IBM
472 Wheelers Farms Road
Milford, CT 06460
203-783-7284
rboivie@ibm.com

Dave Borman

Cray Research
1440 Northland Dr.
Mendota Heights, MN 55120
612-681-3398
dab@cray.com
dab%oliver.cray.com@uc.msc.umn.
edu

Karen Bowers

Corp. for Natl. Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091
703-620-8990
kbowers@nri.reston.va.us

Billy Brackenridge

29160 Upper Ramirez Motorway
Rural Route 4
Malibu, CA 90265
213-457-2026
billy@isi.edu

Scott Brim

Cornell University
265 Olin Hall
Theory Center
Ithaca, NY 14853
607-255-8686
swb@devvax.tn.cornell.edu

Ronald Broersma

Naval Ocean Systems Center
Code 914
San Diego, CA 92152-5000
619-553-2293
ron@nosc.mil

Jeffrey Burgan

NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035
415-694-6440
jeff@nsipo.nasa.gov

Ross Callon

Digital Equipment Corporation
550 King Street
LKG1-2/A19
Littleton, MA 01460-1289
508-486-5009
callon@bigfut.dec.com

C. Allan Cargille

University of Wisconsin-Madison
1204 W. Dayton St.
Madison, WI 53705
608-262-5084
cargille@cs.wisc.edu

Glen Carter

Defense Communications Agency
US Air Force
HQ DCA Code B615
Washington, DC 20305-2000
703-285-5481
gcarter@ddn1.dca.mil

Charles Carvalho

Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
805-963-9431
charles@sage.acc.com

Stephen Casner

University of Southern California
4676 Admiralty Way
ISI
Marina del Rey, CA 90292
213-822-1511
casner@isi.edu

Charlie Catlett

NCSA
605 E. Springfield Avenue
CAB 56
Champaign, IL 61820
217-333-1163
catlett@ncsa.uiuc.edu

J. Noel Chiappa

MIT/Proteon
Two Technology Drive
Westborough, MA 01581-5008
617-898-2800
jnc@lcs.mit.edu

Joseph Choy

NCAR
PO Box 3000
Boulder, CO 80307-3000
US
303-497-1222
choy@ncar.ucar.edu

Richard Colella

Nat. Inst. of Standards & Technology
Building 225
Room B217
Gaithersburg, MD 20899
301-975-3627
colella@osi3.ncsl.nist.gov

Rob Coltun

University of Maryland
Computer Science Dept.
College Park, MD 20742-2411
301-454-2946
rcoltun@trantor.umd.edu

Dave Crocker

Digital Equipment Corporation
Network Systems Lab
100 Hamilton Ave.
Palo Alto, CA 94301
415-688-6820
dcrocker@nsl.dec.com

Bilal Chinoy

Merit Computer Network
1075 Beal Avenue #116
Ann Arbor, MI 48109
313-769-6658
bac@merit.edu

Danny Cohen

University of Southern California
4676 Admiralty Way #1000
ISI
Marina del Rey, CA 90292
213-822-1511
cohen@isi.edu

Michael Collins

Lawrence Livermore National Lab
PO Box 5509
M/S 561
Livermore, CA 94551
415-422-4017
collins@ccc.mfecc.arpa

John Cook

Chipcom Corp.
195 Bear Hill Rd.
Waltham, MA 02154
617-890-6844
cook@chipcom.com

Philip DeMar

Fermi National Accelerator Lab
PO Box 500
M/S 234
Batavia, IL 60510
312-840-3678
demar@fnalb.fnal.gov

Ralph Droms

Corp. for National Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091
717-524-1145
rdroms@nri.reston.va.us
droms@sol.bucknell.edu

Tom Easterday

Ohio State University
1971 Neil Avenue
Room 406
Columbus, OH 43210
614-292-4843
tom@niscuiircc.ohio-state.edu

Robert Elz

University of Melbourne
Computer Science
Parkville, VIC 3161
AUSTRALIA
+61-3-344-5225
kre@munnari.oz.au

Robert Enger

Contel
1300 Quince Orchard Blvd.
Gaithersburg, MD 20878-4199
301-840-4040
enger@sccgate.scc.com

Dino Farinacci

3Com
2081 N. Shoreline Blvd.
Mountain View, CA 94043
415-940-7661
dino@bridge2.3com.com

Michael Fidler

Ohio State University
1971 Neil Avenue, Rm 406
IRCC
Columbus, OH 43210-1210
614-292-4843
ts0026@ohstvma.ircc.ohio-state.edu

James Forster

cisco Systems
1350 Willow Road
Menlo Park, CA 94025
415-326-1941
forster@cisco.com

Richard Fox

Hughes LAN
950 Linden Ave. #208
Sunnyvale, CA 94086
415-966-7924
sytek!rfox@sun.com

Stanley Froyd

Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
805-963-9431
sfroyd@salt.acc.com

Vincent Fuller

Stanford University
Networking Systems
115 Pine Hall
Stanford, CA 94305
415-723-6860
fuller@jessica.stanford.edu

Alison Gabriel-Reilly

Digital Equipment Corporation
BXB1-1/D03
85 Swanson Road
Boxboro, MA 01719
508-264-5130
g_reilly%ultra.enet.dec.com@decwrl.
dec.com

James Galvin

Trusted Information Systems, Inc.
3060 Washington Road
Glenwood, MD 21738
301-854-6889
galvin@tis.com

Jose Garcia-Luna

SRI International
Telecommunications Sciences Ctr
333 Ravenswood Avenue
Menlo Park, CA 94025
415-859-5647
garcia@sri.com

Chuck Gerlach

AT&T Bell Laboratories
1100 East Warrenville Rd
IW 1F305
Naperville, IL 60566-7044
312-979-7325
cag@iwlcs.att.com

Phill Gross

Corp. for Natl. Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091
703-620-8990
pgross@nri.reston.va.us

Martin Gross

DCA
1860 Wiehle Avenue
Code R640
Reston, VA 22090-5500
703-437-2165
martin@protolaba.dca.mil

Robert Hagens

University of Wisconsin-Madison
Computer Science Dept.
1210 West Dayton Street
Madison, WI 53706
608-262-1017
hagens@cs.wisc.edu

Tony Hain

Lawrence Livermore National Lab
PO Box 5509
Livermore, CA 94550
415-422-4200
hain@nmfecc.arpa

Tom Halcin

Hewlett-Packard
19420 Homestead Rd.
M/S 43lh
Cupertino, CA 95014
408-447-2480
halcin%hpinddm@hplabs.hp.com

Martyne Hallgren

Cornell University
265 Olin Hall
Theory Center
Ithaca, NY 14853-5210
607-255-8686
martyne@tcgould.tn.cornell.edu

Brian Handspicker

Digital Equipment Corporation
550 King Street
LKG1-2/E19
Littleton, MA 01460
508-486-7894
bd@vines.dec.com

Kenneth Hays

Florida State University
Super. Computations Res. Inst.
400 Science Center Library
Tallahassee, FL 32306-4052
904-644-7053
hays@scri1.scri.fsu.edu

Charles Hedrick

Rutgers University
Hill Center, POB 879
CCIS
Piscataway, NJ 08854
201-932-3088
hedrick@aramis.rutgers.edu

Robert Hinden

BBN Communications
50 Moulton Street
Cambridge, MA 02138
617-873-3757
hinden@bbn.com

Russell Hobby

University of California
Computing Services
Surge II - Room 1400
Davis, CA 95616
916-752-0236
rdhobby@ucdavis.edu

Jeffrey Honig

Cornell University
265 Olin Hall
Theory Center
Ithaca, NY 14853-5201
607-255-8686
jch@tcgould.tn.cornell.edu

Steven Hunter

Lawrence Livermore National Lab
PO Box 808
M/S L-561
Livermore, CA 94550
415-423-2219
hunter@ccc.mfecc.arpa

Geoff Huston

Australian Academic & Research Network
c/o CSC, Australian National Univ.
GPO Box 4
Canberra, A.C.T. 2601
AUSTRALIA
61-62-493385
gih900@csc.anu.oz.au

Ole-Jørgen Jacobsen

Advanced Computing Environments
806 Coleman Avenue
#9
Menlo Park, CA 94025
415-325-9542
ole@csli.stanford.edu

Dan Jordt

University of Washington
Academic Computing Services
3737 Brooklyn Avenue NE
Seattle, WA 98105
206-543-7352
danj@blake.acs.washington.edu

Satish Joshi

SynOptics Communications, Inc.
501 East Middlefield Road
Mountain View, CA 94043
415-960-1100
sjoshi@mvis1.synoptics.com

Michael Karels

University of California
CSRG Computer Science Div., EECS
457 Evans Hall
Berkeley, CA 94720
415-642-4948
karels@berkeley.edu

Dave Katz

Merit Computer Network
1075 Beal Avenue
Ann Arbor, MI 48103
313-763-4898
dkatz@merit.edu

Kathy Kerby

BBN Communications
50 Moulton St.
Cambridge, MA 02138
617-862-8784
kkerby@bbn.com

Stev Knowles

FTP Software
26 Princess Street
Wakefield, MA 01880-3004
617-246-0900 x270
stev@ftp.com

Tracy LaQuey

University of Texas
Computation Center
M/S COM 1
Austin, TX 78712
512-471-3241
tracy@emx.utexas.edu

Walter Lazear

Mitre Corporation
7525 Colshire Drive
M/S W422
McLean, VA 22102
703-883-6515
lazear@gateway.mitre.org

Eliot Lear

IntelliGenetics
700 East El Camino Real
Mountain View, CA 94040
415-962-7323
lear@presto.ig.com

John Lekashman

NASA Ames Research Center
MS 258-6
Moffett Field, CA 94035
415-694-4359
lekash@orville.nas.nasa.gov

Marianne Lepp

BBN Communications
150 Cambridge Park Drive
Cambridge, MA 02140
617-873-2458
mlepp@bbn.com

Mike Little

SAIC
8619 Westwood Center Dr.
Vienna, VA 22182
703-749-5360
little@saic.com

Dan Long

BBN Communications
NEARNet
10 Moulton St.
Cambridge, MA 02138
617-873-2766
long@bbn.com

Mark Lottor

SRI International
333 Ravenswood Ave
M/S EJ282
Menlo Park, CA 94025
415-859-2652
mkl@nic.ddn.mil

Susan LoVerso

Encore Computer Corp.
257 Cedar Hill St.
Marlboro, MA 01752
508-460-0500
sue@encore.com

John LoVerso

Xylogics, Inc.
53 Third Ave.
Burlington, MA 01803
617-272-8140
loverso@xylogics.com

E. Paul Love

San Diego Supercomputer Center
PO Box 85608
San Diego, CA 92138
619-534-5043
loveep@sds.sdsc.edu

Gary Malkin

Proteon
2 Technology Drive
Westborough, MA 01581
508-898-2800
gmalkin@proteon.com

Louis Mamakos

University of Maryland
Computer Science Center - Syst
College Park, MD 20742
301 454-2943
louie@trantor.umd.edu

Matt Mathis

PSC
4400 5th Ave.
Pittsburgh, PA 15213
412-268-3319
mathis@fornax.ece.cmu.edu

Keith McCloghrie

Hughes LAN Systems
1225 Charleston Road
Mountain View, CA 94043
415-966-7934
sytek!kzm@hplabs.hp.com

Paul McKenney

SRI International
333 Ravenswood Ave.
EK360
Menlo Park, CA 94025
415-859-4910
mckenney@sri.com

Bill Melohn

Sun Microsystems
2550 Garcia Avenue
Mountain View, CA 94043
415-336-2941
melohn@sun.com

David Miller

Mitre Corporation
Burlington Road
M/S EO66
Bedford, MA 01730
617-271-3993
dtm@mitre.org

Berlin Moore

PREPnet
530 N. Neville
Pittsburgh, PA 15213
412-268-7870
prepnet@andrew.cmu.edu

Dennis Nadler

US Air Force
1842 EEG/EETSS
Scott AFB, IL 62225
618-256-2656
1842eeg-eets@afcc-0a1.af.mil

Milo Medin

NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035
415-490-9157
medin@nsipo.nasa.gov

Donald Merritt

Ballistic Research Laboratory
Attn: AMXBR-SECAD
Aberdeen Proving Gro, MD 21005-5066
301-278-6808
don@brl.mil

Paul Mockapetris

University of Southern California
4676 Admiralty Way
ISI
Marina del Rey, CA 90292
213-822-1511
pvm@isi.edu

Jun Murai

University of Tokyo
2-11-16 Yakoi, Bunkyoku
Computer Center
Tokyo 113,
JAPAN
81-3-5684-3400
jun@u-tokyo.ac.jp

Oscar Newkerk

Digital Equipment Corporation
14475 West 24th St.
Bellevue, WA 98007
206-865-8913
newkerk@decwet.dec.com

Torben Nielsen

University of Hawaii
Dept of ICS
2565 The Mall, Keller 304A
Honolulu, HI 96822
808-949-6395
torben@foralie.ics.hawaii.edu

Rebecca Nitzan

NASA
600 Maryland Ave., SW
Suite 200 East
Washington, DC 20024
703-765-6115(H) 202-554-8677
nitzan@nsipo.nasa.gov

Bill Norton

Merit Computer Network
1075 Beal Avenue
Ann Arbor, MI 48104
313-936-2656
wbn@merit.edu

Lee Oattes

University of Toronto
Computer Services
4 Bancroft Avenue, Rm. 102
Toronto, ON M5S 1C1
CANADA
416-978-5448
oattes@utcs.utoronto.ca

Zbigniew Opalka

BBN Communications
50 Moulton St.
Cambridge, MA 02238
617-873-2888
zopalka@bbn.com

Philippe Park

BBN Systems & Technologies
10 Moulton Street
Cambridge, MA 02238
617-873-2892
ppark@wilma.bbn.com

Craig Partridge

BBN Systems & Technologies
10 Moulton Street
Cambridge, MA 02138
617-873-2459
craig@bbn.com

Gurudatta Parulkar

Washington University
Campus Box 1045
St. Louis, MO 63130
314-889-4621
guru@flora.wustl.edu
parulkar@udel.edu

Richard Pethia

Carnegie Mellon University
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-3890
412 268-7739
rdp@sei.cmu.edu

Mel Pleasant

Rutgers University
H019 Hill Center
Busch Campus
Piscataway, NJ 08855-0879
201-932-2023
pleasant@rutgers.edu

Paul Pomes

University of Illinois, CSO
1304 W. Springfield Avenue
Urbana, IL 61801-2987
217-333-6262
paul@uxc.cso.uiuc.edu

Michael Reilly

Digital Equipment Corporation
550 King Street
LKG1-3/L06
Littleton, MA 01460
508-486-5283
reilly%atari.nac.dec.com@decwrl.dec.
com

Robert Reschly

Ballistic Research Laboratory
ATTN: SLCBR-SF (Reschley)
Aberdeen Proving Gro, MD 21005-5066
301-278-6678
reschly@brl.mil

Jim Robertson

3Com
2081 N. Shoreline Blvd.
Mountain View, CA 94043
415-940-2683
jar@bridge2.3com.com

Marshall Rose

NYSERNet Inc
420 Whisman Court
Mountain View, CA 94043-2112
415-961-3380
mrose@nisc.nyser.net

K.K. Ramakrishnan

Digital Equipment Corporation
LKG 1-2/A19
550 King Street
Littleton, MA 01460-1289
508-486-7267
rama%erlang.dec.com@decwrl.dec.
com

Yakov Rekhter

IBM
TJ Watson Research Center
Route 134 PO Box 218
Yorktown Heights, NY 10598
914-945-3896
yakov@ibm.com

Joyce K. Reynolds

University of Southern California
4676 Admiralty Wy #1001
ISI
Marina del Rey, CA 90292-6695
213-822-1511
jkrey@venera.isi.edu

Carl-Herbert Rokitansky

Fern Univ. of Hagen, FB ET/DVT
Data Processing Techniques
Frauenstuhlweg 31
D-5860 ISERLOHN,
WEST GERMANY
49-2371-566-235
roki@dhafeu52.bitnet

Milt Roselinsky

Communication Machinery Corporation
125 Cremona Drive
Santa Barbara, CA 93117
805-562-3194
cmcvax!milt@hub.ucsb.edu

Mark Rosenstein

MIT
Amherst St.
Room E40-3311
Cambridge, MA 02139
617-253-1530
mar@athena.mit.edu

Karen Roubicek

BBN Systems & Technologies
10 Moulton St.
NSF Network Service Center
Cambridge, MA 02138
617-873-3361
roubicek@nnsc.nsf.net

Betsey Ruben

Prime Computer
500 Old Conn. Path
Framingham, MA 01701
508-879-2960
bnr@s62.prime.com

Bill Rust

FTP Software
26 Princess Street
Wakefield, MA 01880-3004
617-246-0900
wjr@ftp.com

Vlad Rutenburg

SRI International
333 Ravenswood Avenue
EJ 386
Menlo Park, CA 94025-3493
415-859-2683
vlad@sri.com

Jeffrey Schiller

MIT
1 Amherst Street
Room E40-311
Cambridge, MA 02139
617-253-4101
jis@bitsy.mit.edu

Steven Schoch

NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035
415-694-4121
schoch@ames.arc.nasa.gov

Jim Sheridan

IBM
166 East Shore Drive
PO Box 334
Whitmore Lake, MI 48189
313-393-6537
jsherida@ibm.com

Dana Sitzler

Merit Computer Network
1075 Beal Avenue
Ann Arbor, MI 48109-2112
313-936-3000
dds@merit.edu

Frank Solensky

Racal InterLAN
155 Swanson Rd
Boxborough, MA 01719
508-263-9929
solensky@interlan.interlan.com

Michael St. Johns

Department of Defense, Attn:T41
9800 Savage Rd.
Ft. Meade, MD 20755
301-688-6742
stjohns@umd5.umd.edu

Mary K. Stahl

SRI International
Net Info Sys Ctr
333 Ravenswood Ave. Rm EJ 296
Menlo Park, CA 94025
415- 859-4775
stahl@nisc.sri.com

Martha Steenstrup

BBN Communications
150 CambridgePark Dr.
Rm 20/665
Cambridge, MA 02140
617-873-3192
msteenst@bbn.com

Louis Steinberg

IBM
472 Wheelers Farms Rd
M/S 91
Milford, CT 06460
203-783-7175
louiss@ibm.com

Roxanne Streeter

NASA Ames Research Center
Sterling Software
M/S 233-18
Moffett Field, CA 94035
415-694-4845
streeter@nsipo.arc.nasa.gov

Zaw-Sing Su

SRI International
333 Ravenswood Ave.
EJ280
Menlo Park, CA 94025
415-859-4576
zsu@tsca.istc.sri.com

Claudio Topolcic

BBN Systems & Technologies
10 Moulton Street
Cambridge, MA 02138
617-873-3874
topolcic@bbn.com

Gregory Vaudreuil

Corp. for Natl. Research Initiatives
1895 Preston White Drive
M/S 100
Reston, VA 22091
703-620-8990
gvaudre@nri.reston.va.us

Ross Veach

University of Illinois
1304 West Springfield
Urbana, IL 61801
217-244-4274
rrv@seka.cso.uiuc.edu

Bill Westfield

cisco Systems
1350 Willow Rd.
Suite 201
Menlo Park, CA 94025
415-326-1941
billw@cisco.com

Bruce Wilder

IBM
PO Box 12195
Research Triangle Park
Dept. G20/B503
Raleigh, NC 27709
919-254-4087
bwilder@yktvmx.watson.ibm.com

Rick Wilder

Mitre Corporation
7525 Colshire Dr
M/S W425
McLean, VA 22102
703-883-6174
rick@gateway.mitre.org

Steve Willis

Wellfleet Communications
12 DeAngelo Drive
Bedford, MA 01730
617-275-2400
swillis@wellfleet.com

Dan Wintringham

Ohio Supercomputer Center
1224 Kinnear Road
Columbus, OH 43212
614-292-0901
danw@igloo.osc.edu

Brian Yasaki

Wollongong Group
7799 Leesburg Pike
Suite 1100, North Tower
Falls Church, VA 22043
703-847-6340
bky@twg.com

Mary Youssef

IBM
472 Wheelers Farms Rd.
M/S91
Milford, CT 06460
203-783-4338
mary@ibm.com

Aileen Yuan

Mitre Corporation
7525 Colshire Dr.
McLean, VA 22102
703-883-7023
aileen@gateway.mitre.org

Lixia Zhang

Xerox PARC
333 Coyote Hill Road
Palo Alto, CA 94304
415-494-4415
lixia@arisia.xerox.com