Proceedings of the Eleventh Internet Engineering Task Force

October 17-19, 1988 in Ann Arbor, MI

Compiled by

Phill Gross
Karen Bowers

January 1989

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

Acknowledgements

The 17-19 October 1988, IETF meeting was hosted at the Universtity of Michigan in Ann Arbor by Elise Gerich and Hans-Werner Braun of Merit, Inc. I wish to express very sincere appreciation to Elise Gerich (Merit), who handled all local arrangements. Her efficient planning beforehand and tireless help during the meeting made the meeting a pleasure to chair. She also held an informative and entertaining tour of the NSFnet Operations Center.

I'd also like to thank Gladys Reichlen and Allison Mankin of MITRE, who helped plan and run the meeting from the MITRE end.

Table of Contents

1. Chairman's Message

2. Overview and Status of IETF Working Groups

3. IETF Attendees

4. Final Agenda

5. Working Group Reports/Slides
   Ann Arbor, MI
   17-19 October 1988

6. Network Status Briefings and Technical Presentations

   - NSFnet Report - Part I (Hans-Werner Braun)
   - NSFnet Report - Part II (J.E. Drescher)
   - ARPANET/DDN Report (Marianne Lepp)
   - DDN Report: Transition of DDN Mailbridges from LSI-11 to Butterfly Gateways (Michael Brescia)
   - Interop 88 Network Report (Philip Almquist)
   - Internet Protocols for Amateur Radio (Phil Karn)
   - High Performance TCP Over an Ethernet (Van Jacobson)
   - Congestion Control Observations Using NETMAN (Allison Mankin)
   - What Is USENET? What Is NNTP? (Gene Spafford)
   - The NIC Domain Chart (Mark Lottor)
   - On Some T_1 Satellite Link Performance (John Lekashman)

7. Additional Working Group Updates

   - Interconnectivity (Guy Almes)
     NASA Ames Meeting
     27 September 1988
   - PDN Routing (Carl-H. Rokitansky)
     USNA, Annapolis Meeting
     15-17 June 1988

8. Papers Distributed at IETF

   - NSFnet Connectivity and Configuration (Susan Hares)
   - Management and Operation of the NSFnet Backbone Network Project (MERIT)
   - California Internet Federation

1. CHAIRMAN'S MESSAGE

PHILL GROSS
NRI

Chairman's Message

The IETF is growing. There are currently 16 active Working Groups in the IETF and the quarterly meetings are typically attended by 100-150 people. Chairing a group of this size, with this level of activity, is no longer a simple matter. The administrative details and logistics involved in planning meetings and producing the Proceedings had begun to detract from the more important mission of identifying key Internet problem areas and then organizing Working Groups to solve them.

I am pleased that, beginning with the January 18-20 IETF meeting, Karen Bowers (NRI, Senior Systems Analyst) will be working with me on many of these IETF matters. Karen will take almost complete responsibility for the Proceedings and many of the meeting planning activities. As a result, by next month we should be able to announce the dates and locations of the next 5 IETF meetings. She will also be working closely with me to facilitate the progress of the WGs. For example, we are considering a quarterly IETF electronic newsletter to announce WG meetings, documents, and status. This should help all those interested in IETF activities to be more aware of the activities of the various WGs. It may also help WGs maintain momentum between IETF plenary meetings.

A condensed status of the currently active IETF Working Groups is provided in the attached chart. Chapter 2 expands this information with an overview of each working group and a summary of progress to date.

For more detailed information (e.g., to obtain a description of the WGs; to obtain copies of the draft documents or WG reports; or to obtain information on meeting dates and locations), contact either the Chairs/Points-of-Contact directly (listed below) or send a request to bowers@sccgate.scc.com. We are now in the process of updating and reorganizing the IETF directory at SRI-NIC to make all this information more easily accessible online.

Phill Gross
(interim address: gross@sccgate.scc.com)
Corporation for National Research Initiatives (NRI)
1895 Preston White Drive, Suite 100
Reston, VA 22091
703-620-8990

Summary of IETF Working Group Status

(January 1989)

Working Groups	RFC or Draft?	Met Oct 88?	Current Report?	Meeting Jan 89?	Chair or POC (address)
Authentication	Yes	Yes	Yes	No	Jeff Schiller (MIT) jis@athena.mit.edu
CMIP-over-TCP (CMOT)	Yes	Yes	Yes	Yes	Lee LaBarre (MITRE) cel@mitre-bedford.arpa
Host Requirements	Yes	Yes	Yes	Yes	Bob Braden (ISI) braden@isi.edu
Interconnectivity	No	Yes	Yes	Yes	Guy Almes (Rice) almes@rice.edu
Internet MIB	Yes	Yes	Yes	Yes	Craig Partridge (BBN) craig@nnsc.nsf.net
NSFnet/Reg Monitoring	No	Yes	Yes	Yes	Susan Hares (Merit) skh@merit.edu
Open SPF-based IGP	Yes	Yes	Yes	No	Mike Petry (UMD) petry@trantor.umd.edu
Open Systems Routing	Yes	No	Yes	No	Marianne Lepp (BBN) mlepp@bbn.com
OSI Interoperability	Yes	NA	NA	Yes	Ross Callon (DEC) callon@erlang.dec.com
PDN Routing Group	No	Yes	Yes	No	C-H Rokitansky roki@isi.edu
Performance and CC	No	Yes	Yes	Yes	Allison Mankin (MITRE) mankin@gateway.mitre.org
Pt-Pt Protocol	Yes	Yes	Yes	Yes	Drew Perkins (CMU) ddp#@andrew.cmu.edu
ST and CO-IP	Yes	Yes	Yes	Yes	Claudio Topolcic (BBN) topolcic@bbn.com
TELNET Linemode	Yes	Yes	Yes	Yes	Dave Borman (Cray) dab@cray.com
User Services (New)	NA	NA	NA	Yes	Karen Bowers (NRI) bowers@sccgate.scc.com

2. OVERVIEW AND STATUS OF IETF WORKING GROUPS

OVERVIEW AND STATUS OF IETF WORKING GROUPS

This section provides the following basic information for all currently active IETF Working Groups (listed below):

1) Statement of charter and goals
2) Progress to date
3) Estimate of timeframe for completion
4) Dates of last and next meeting
5) Name of WG mailing lists
6) Names of key players

Working Groups	Chair or Reporter (address)
Authentication	Jeff Schiller (MIT) jis@athena.mit.edu
CMIP-over-TCP (CMOT)	Lee LaBarre (MITRE) cel@mitre-bedford.arpa
Host Requirements	Bob Braden (ISI) braden@isi.edu
Interconnectivity	Guy Almes (Rice) almes@rice.edu
Internet MIB	Craig Partridge (BBN) craig@nnsc.nsf.net
NSFnet/Reg Monitoring	Susan Hares (Merit) skh@merit.edu
Open SPF-based IGP	Mike Petry (UMD) petry@trantor.umd.edu
Open Systems Routing	Marianne Lepp (BBN) mlepp@bbn.com
OSI Interoperability	Ross Callon (DEC) callon@erlang.dec.com
PDN Routing Group	C-H Rokitansky roki@isi.edu
Performance and CC	Allison Mankin (MITRE) mankin@gateway.mitre.org
Pt-Pt Protocol	Drew Perkins (CMU) ddp#@andrew.cmu.edu
ST and CO-IP	Claudio Topolcic (BBN) topolcic@bbn.com
TELNET Linemode	Dave Borman (Cray) dab@cray.com
User Services (New)	Karen Bowers (NRI) bowers@sccgate.scc.com

Authentication Working Group

Jeff Schiller (MIT)
jis@athena.mit.edu

1) Brief statement of charter and goals

There are currently four main deliverables:

A) RFC specifying an authentication format which supports multiple authentication systems. [This document may wind up being specific to SNMP per discussions at the last working group meeting].

B) Document discussing the cost/benefit tradeoffs of various generic approaches to solving the authentication problem in the Internet context.

C) Document to act as a protocol designers guide to authentication.

D) RFC proposing A Key Distribution System (emphasis on "A" as opposed to "THE"). MIT's Kerberos seems the most likely candidate here.

2) Progress to date

As of this time there is an IDEA paper that is a description of the kerberos protocol. Jennifer Steiner at MIT is currently working on an RFC format document to submit that will describe the kerberos protocol in detail sufficient to code to.

3) Estimate of timeframe for completion

Hard to state clearly as the charter of the group (not to mention the membership) is still subject to change. However I would expect that the Kerberos RFC should be in draft format if not by January 17th, then before the IETF meeting following. We would like to also have a document defining authentication extensions to SNMP in draft format before the IETF meeting following the January meeting.

4) Dates of last and next meeting

Last Meeting: IETF meeting at Merit
Next Meeting: April 1989 IETF meeting (tentative)

5) Name of WG mailing lists

awg@bitsy.mit.edu

6) Names of key players

Jon Rochlis, Jeff Schiller and Jennifer Steiner

CMIP-over-TCP (CMOT) Working Group

Lee LaBarre (MITRE)

labarre@gateway.mitre.org

1) Charter: As described in RFC1052

   - Develop a long term approach to management of the Internet based on the OSI Network Management Framework and the Common Management Information Protocol (CMIP).

   - Provide input to the OSI standards process based on experience in the Internet, and thereby influence the final form of OSI International Standards on network management, in particular CMIS/P.

   - Approach

     a) Develop prototype implementors agreements on CMIP over TCP.

     b) Develop prototype implementations based on the CMOT agreements and IETF SMI and MIB agreements.

     c) Experiment with CMOT and extensions to the SMI and MIB.

     d) Develop final implementors agreements for CMOT.

     e) Promote development of products based on CMOT.

     f) provide input to the OSI Network Management standards process in time to effect the International Standards.

2) Expected duration of group:

   The groups work should be completed by June 1989.

3) List of Members:

   Member corporations are listed here.

Advanced Computing Environments	Communications Machinery Corp.
Convergent Technologies	Digital Equipment Corp.
Epilogue Technology Corp.	Excelan
Hewlett Packard Corp.	MITRE Corp.
SUN Microsystems	Sytek
3COM Corp.	Ungermann-Bass
Unisys Corp.	The Wollongong Group

4) Mailing List:

netman@gateway.mitre.org

5) Last meeting:

December 1988, Santa Clara, CA

7) Achieved goals: from (1)

a) Overview document (IDEA0012)
Thin Presentation layer (IDEA0017)
Prototype Implementors Agreements (IDEA0025)

b) Nine vendor prototype implementations demonstrated at INTEROP 88 in Santa Clara, CA.

c) Experimentation occurred during development of the INTEROP demo, and is continuing.

d) Draft implementors agreements are written for the DIS CMIP over TCP. Proposals for extending the SMI and MIB are in progress.

e) Thirteen corporations participated in the INTEROP 88 demo. Nearly all the vendors in that group have indicated that they expect to field products during 1989 based on CMOT implementors agreements.

f) Several Working Group members are participating in the OSI network management standards organizations and carrying the CMOT experience into that forum.

HOST REQUIREMENTS WORKING GROUP

Bob Braden (ISI)

braden@isi.edu

CHARTER AND GOALS:

The primary task of the Host Requirements Working Group (HRWG) is to prepare an RFC entitled "Requirements for Internet Hosts". This RFC will contain a comprehensive specification of the networking software requirements for an Internet host, to complement the Gateway Requirements RFC-1009.

As a secondary task, the WG has provided a forum for discussing particular solutions to pressing host problems, and has resulted in several RFC's by WG members.

The Host Requirements RFC covers the following topics:

- o Link Layer (only amendments to RFC-1009 discussion)
- o IP Layer (IP and ICMP)
- o Transport Layer (TCP and UDP)
- o Application Layer (SMTP, FTP, TFTP, and Telnet)
- o Support Programs (DNS, Booting, Network Management)

For each protocol, it amends and expands on the specification RFC(s). In those areas in which the referenced specifications contain ambiguous or incomplete information, the RFC contains further clarification, discussion, and guidance. The intent is to define the current architecture as completely and carefully as possible, not to invent new architecture.

PROGRESS TO DATE:

The draft document is nearly complete, after 5 meetings in 10 months. The 6th and last meeting is scheduled for the Austin IETF meeting in January 1989. The draft is now 175 pages.

ESTIMATE OF PUBLICATION DATE:

February 1, 1989.

MEETING DATES:

Last: Oct. 17-18, 1988 at Ann Arbor IETF meeting.

Next: Jan. 18-19, 1989 at Austin IETF meeting.

MAILING LIST:

To FTP the document, do anonymous FTP to host venera.isi.edu and fetch pathname:

pub/ietf-hosts.rfc.txt

This file is ~400KB. Change bars (and other symbols) mark all the content changes since the Ann Arbor meeting. Another file is available at the same host that contains only the text marked with change bars:

pub/ietf-hosts.rfc.chg

KEY PLAYERS:

Major contributions to the writing, revision, and editing have come from 25 people representing 20 organizations. At least 8 vendors have been represented.

SPINOFFS:

In writing this document, we came across a number of unresolved problems and undocumented areas. As a subsidiary task, the HRWG members have been inspired to prepare a number of RFC's on these topics. The RFC's for which we take credit are:

RFC-1063 ICMP MTU Discovery

RFC-1071 Internet Checksum Calculation

RFC-1073 Telnet Window Size Option

(draft) Telnet Terminal Type Extension RFC

(draft) Gateway Discovery RFC

(draft) TCP RST Extension RFC

IWG

Interconnectivity Working Group

Guy Almes (Rice)

almes@rice.edu

1) Statement of the charter and goal of the group

Within six quarters, specify, design, and demonstrate an initial production-quality implementation of inter-autonomous-system routing adequate to address the inadequate support for the NSFnet Model in current Inter-AS Routing.

Inadequate support for the NSFnet Model in current Inter-AS Routing:

Interconnectivity of the Internet no longer conforms to the stub model assumed by the designers of EGP. We currently suffer from (a) dangerous ad hoc interconnections due to the bold and (b) less interconnectivity due to the conservative. Further, we do not expect a new generation of inter-autonomous-system routing protocols to be designed, much less implemented, for several years. While the existing Open Inter-AS Working Group is needed to design a really new generation of protocols, and while the Short-Term Routing Working Group has made valuable contributions, we need a methodical approach to Inter-AS routing that can be applied in the context of the current Inter-agency Research Internet with its multiple national backbones, its evolving mid-level networks, and its exploding campus networks. This three-level NSFnet Model, while much more general than the older Stub Model, is much less general than the situation being addressed by the Open Inter-AS Working Group.

One possible technical approach is to appropriate and adapt the work of the EGP-3 Working Group. If no substantial improvement over the currently available tools can be produced within a short time frame, then it would be preferred to simply document what we've learned and await the product of the Open Inter-AS Working Group.

The costs of the current interconnectivity approach are large. They result in either having very labor intensive routing configurations or in less than adequate interconnectivity and the resulting long paths and lack of robustness.

2) Expected duration of the group: Six to eight quarters at the very most.

3) List of members: Initially, I am inviting:
   Guy Almes of Rice University/Sesquinet (almes@rice.edu), chair
   Mike Brescia of BBN/ARPAnet (brescia@alexander.bbn.com)
   Joe Choy of UCAR/USAN (choy@windom.ucar.edu)
   Phill Gross of MITRE/IETF (gross@gateway.mitre.org), ex officio
   Milo Medin of NASA/NSI (medin@nsipo.nasa.gov)
   Jacob Rekhter of IBM/NSFnet (yakov@ibm.com)

   Two of us (GA and JC) are active in NSFnet-related mid-level networks. Apart from PG, the others are active in different national backbones (ARPAnet, NSI, and NSFnet respectively). We need one more person from a mid-level and perhaps someone from ESnet. We are, as noted earlier, open to suggestions, but would like to keep the WG down to about eight members.

4) Give Mailing lists for the group:

   iwg@rice.edu

5) When was your last meeting?

   Last meeting at the Oct IETF,
   Next meeting will be at the Jan IETF.

MANAGEMENT INFORMATION BASE (MIB) WORKING GROUP

Craig Partridge (BBN)

craig@nnsc.nsf.net

1) Brief statement of charter and goals (ie, ~3-5 sentences with quantifiable goal like 'RFC specifying new Point-Point protocol')

   As defined in RFC 1052, the original purpose was to devise an Internet MIB and Structure of the Managment Information (SMI). When we finished, the WG stayed around as a forum where revisions of the MIB and SMI may be considered and approved.

2) Progress to date

   MIB (version 1) came out in the summer of 1988. RFCs 1065/1066.

   MIB (version 2) is planned for summer of 1989. Some proposals for changes in hand. First draft of new RFC expected in February.

   After MIB-2 the crystal ball gets hazy. The key unresolved questions are how long does the MIB have to work for both CMIP and SNMP (to forestall parties fighting for position, I've said very loudly that MIB-2 will but the question is open after MIB-2 is done).

3) Estimate of timeframe for completion

   As long as we need to keep tinkering with the MIB.

4) Dates of last and next meeting

   Last meeting: October IETF
   Next meeting: January 17th (IETF)

5) Name of WG mailing lists (if any; include address)

   mib-wg@nnsc.nsf.net is for the "core" members
   gwmon@sh.cs.net is for general discussion of network management issues

MIB

6) Names of key players

Karl Auerbach, Epilogue Technology
K. Ramesh Babu, Excelan
Lawrence Besaw, Hewlett-Packard
Terry Bradley, Wellfleet Communications
Jeffrey D. Case, University of Tennessee at Knoxville [OPEN-INOC WG]
James R. Davin, MIT (formerly Proteon)
Mark S. Fedor, NYSERNet
Phill Gross, NRI
Bent Torp Jensen, Convergent Technology
Lee Labarre, The MITRE Corporation [NETMAN WG]
Dan Lynch, Advanced Computing Environments
Keith McCloghrie, The Wollongong Group
Dave Mackie, 3Com/Bridge
Craig Partridge, BBN
Jim Robertson, 3Com/Bridge
Marshall T. Rose, The Wollongong Group
Greg Satz, cisco
Martin Lee Schoffstall, Rensselaer Polytechnic Institute
Lou Steinberg, IBM
Dean Throop, Data General
Unni Warrier, Unisys

JOMAAN

Joint Monitoring Access for Adjacent Networks focusing on the NSFNET Community Working Group. (A suggestion for an abbreviation is NSFNET Jo-MAAN, pronounced Joe - Man).

Sue Hares (Merit)
skh@merit.edu

Charter or Mission of NSFNET Jo-MAAN Working Group:

This Joint Monitoring Access for Adjacent Networks focusing the NSFNET Community Working group will:

- o discuss how to identify problems in the next hop network
- o create a list of existing tools which can solve these problems
- o Create a list of routing topology maps of regionals

We are focusing on the NSFNET community - the NSFNET backbone, the regional networks attached to the NSFNET backbone, campus networks, and peer networks for the NSFNET which includes the ARPANET and the MILNET.

Who should attend:

Technical representatives from mid-level or peer networks. In the future we may want to extend this to technical representatives from campus networks. However, in interest of getting a lot of work done quickly I would like to limit the initial working group.

Time duration for working group:

6-9 months (August 31, 1989)

Dates of Meetings:

1st - October 18th, 1988 at October IETF Meeting

2nd - January 18-20, 1989 at January IETF Meeting

3rd - March, at Routing Workshop help by NSFNET

JOMAAN

Mail group for working group:

njm@merit.edu

send requests to join to njm-request@merit.edu

Key players:

Susan Hares and Hans-Werner Braun. The idea came from David Wasley. However, all the regional technical representatives also play a key role.

Open SPF-based IGP (OSPFIGP) Working Group

Mike Petry (U. Maryland) and John Moy (Proteon)

Charter:

Design and developement of a multi-vendor SPF-based Internet Gateway Protocol. The protocol should draw on existing SPF routing technology, notably the work done by BBN and DEC.

Features of the protocol should include: stability in a large, heterogeneous AS; TOS support; the ability to pass external routing information transparently; explicit support for IP subnetting; authentication of participating routers.

The reasons for choosing an SPF base are 1) So the internet community can gain experience with a routing algorithm other than the current Ford-based algorithms and 2) To ease ISO transition, since the current ANSI proposal is SPF based.

The group should take the protocol through implementation and performance evaluation.

Goals:

meeting June 88 (1st draft of specification)
next meeting (trial implementations, spec revision)
next 2 meetings (performance evaluation, spec revision?)
Then we disband.

Membership: Open

Mailing list: oigp@trantor.umd.edu (open)

Last meeting: October IETF meeting
Next meeting: February (by video teleconference)

Progress:

March 88 - IDEA005 published (protocol requirements)
May 88 - IDEA020 published (comparison to DEC IS-IS)
June 88 - First third of spec released for public review

Open Routing Working Group

Marianne Lepp (BBN), co-chair mlepp@bbn.com
Robert Hinden (BBN), co-chair hinden@bbn.com

1) Charter and Goals of the Working Group

The charter of the working group is to design a policy-based routing protocol to run between autonomous systems to replace EGP and hand-configured tables. The protocol should deal gracefully with a large, heterogenous Internet with constraints determined administratively.

Document	Schedule
Requirements	completed -- IDEA 007
Draft archtiture	March, 1989
Draft specification	December, 1989

The group's final goal is an RFC draft specification.

2) Progress to Date

- o. Requirements paper complete.
- o. Several draft architectures are under consideration.
- o. There is a consensus on the basic points of the architecture -- hierarchical, source routing, route set-up, link-state, and other points.

3) Estimate of Timeframe for Completion

One year to complete charter of writing a draft specification.

4) Dates of Last and Next Meetings

Last meeting: Nov. 9,10 in Westboro, MA
Next meeting: Feb. on the West Coast.

5) Name of Working Group Mailing Lists

Private mailing list: open-rout-wg@bbn.com
Public mailing list: open-rout-interest@bbn.com

6) Names of Key Players and Liasons with Other Working Groups/Task Forces

Membership of the group (by invitation):

Robert Hinden	hinden@bbn.com
Ross Callon	rcallon%erlang.dec.com@decwrl.dec.com
Sergio Heker	heker@jvnca.csc.org
Noel Chiappa	jnc@xx.lcs.mit.edu
Mike Little	little@MACOM4.ARPA
Marianne Lepp	mlepp@bbn.com
Mike Petry	petry@TRANTOR.UMD.EDU
Zaw-Sing Su	zsu@tsca.istc.sri.com
Lixia Zhang	lixia@xx.lcs.mit.edu
Paul Tsuchiya	tsuchiya@gateway.mitre.org
Pat Clark	paclark%ford-cos1@ford-wdl1.arpa
Tassos Nakassis	nakassis@icst-ecf.arpa

Other relevant WG,TF

Auto-nets
Routing subcomittee of the FRICC

OSI Internet Interoperability Working Group

Ross Callon (DEC), Robert Hagens (U Wisc.)

This WG is reforming after a period of inactivity. The next meeting will be held at the January IETF meeting in Texas. A mailing list will be created at Wisconsin. It is the intent of the WG chairs to solicit participation from key OSI players, like NIST, ANSI X3S3.3, the Government OSI User's Group (ie, the originator's of GOSIP). The goal of the group are listed below.

1. Main Goal:

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 of the TCP/IP and OSI protocol suites.

2. Very Short Term Subgoals:

This section describes subgoals which are essential to initial deployment of OSI protocols in the Internet. We intend to work on these goals immediately, and finish initial action relatively quickly (hopefully within a couple of IETF meetings, and soon enough to influence initial OSI software releases).

2.1 Addressing

Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet.

2.2 EON

Provide documentation of the EON experimental effort.

2.3 Berkeley Release 4.4

Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley folks.

2.4 GOSIP

Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.

2.5 Getting Gateways Into the Internet

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.

Note, short term gateway sub-issues may include:

- Wonder whether funding is present to cause OSI gateways to happen
- Do we run dual gateways only and/or start with OSI over IP and/or vice versa. Does this depend on level of funding available?
- Determine what form of routing may be used in the short term (both within a domain, and between domains). Will this be fixed tables at first? (with migration to ANSI routing?)
- Recommend short term domain structure.
- Determine congestion control to be used in first release. Should this include use of the congestion experienced bit and related TP4 congestion algorithm?

3. Possible Short to Medium term sub-goals:

This section describes subgoals with are important to the success of OSI in the DoD Internet, but which are not essential to be completed before initial deployment of OSI protocols in the Internet. This is a tentative list, and is expected to be updated as we go along.

3.1 OSI Software Releases

Continue to work with Berkeley and CMU/MACH on future releases of OSI software.

3.2 Requirements for OSI Gateways / End Systems

Produce documentation on the requirements for OSI gateways and requirements for OSI end-systems, similar to the specs that have been produced for DoD protocol suites.

3.3 Dual Internets versus Encapsulation

Discuss relative strengths of dual internets versus encapsulation. Discuss possible problems with dual gateways (such as interaction between different congestion control schemes, and performance implications of running multiple routing schemes). Produce guidelines for a dual gateway.

3.4 Routing

Work on testing and deployment of the ANSI routing spec for OSI intra-domain routing in the Internet. We do not want to wait for a DP or a DIS. If we find bugs in the routing spec, then they become exponentially harder to fix as the standards process reaches further milestones.

Think about how a new inter-domain routing protocol may be used in the Internet.

3.5 Liaison

Continue liaison with GOSIP Users Group, as necessary.

Cooperate with ANSI and the NBS Implementors forums. Hopefully much or all of this can be done by phone, email, and overlap in corporate attendance, without the need for working group members to go out of their way to attend ANSI or NBS meetings.

3.6 Performance

Discuss performance of OSI. Determine which are implementation versus architectural factors in performance. Is the OSI releases in the kernel or user processes, what are layer interactions like, etc.

3.7 Directory Services

Outline the form of a possible Directory / Domain Naming service for the Internet. Should directory services for DoD and OSI suites be integrated? Are existing schemes suitable and available (e.g., current Internet directories, DEC DNA architecture).

4. NOT ISSUES FOR THIS GROUP:

IETF to ANSI liaison. There are a number of efforts in IETF that ANSI may be interested in for consideration in their future work. To a large extent, appropriate individuals in ANSI are already receiving IETF documentation and making use of this in their standards efforts. It is unlikely that we would be needed as a conduit for carrying documents, and we do not intend to be advocates to ANSI for IETF positions in general. Individual working group members may be advocates for some positions as they see fit.

PDNWG

Internet/Public Data Network Routing Group
("PDN Routing Group")

C-H Rokitanski, roki@isi.edu

1) Statement of the charter and goals of the group

The DoD INTERNET TCP/IP protocol suite has developed into de facto industry standard for heterogenous packet switching computer networks. In the US the ARPANET/MILNET connects several hundreds of INTERNET networks, 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 no algorithms are provided so far to dynamically route INTERNET datagrams through X.25 public data networks. Therefore the goal of the Internet/Public Data Network routing Group is to develop and to define the required routing and gateway algorithms for an improved worldwide routing of INTERNET datagrams through the System of Public Data Networks (PDN). Especially the following issues have been specified:

- Define the Cluster-Addressing Scheme and its application to public data networks as an INTERNET standard

- Specify gateway algorithms and protocols to be used by VAN-gateways

- Develop an X.121 Address Server/Resolution Protocol

- Develop (or support other working groups in developing) routing algorithms based on routing metrics other than hop-count: costs, delay, throughput, TOS, etc.

- Provide interoperability with ISO/OSI networks via the PDN

- Specification of protocols required for an European INTERNET/Public Data Network Information and Operation Center (cooperation with US-INTERNET NICs and NOCs)

- ISO-Migration of the INTERNET/PDN Cluster

2) Progress to date

See separate report of October meeting at the IETF. Mail to roki@isi.edu or gross@sccgate.scc.com for a copy.

3) Duration of the group

The PDN Routing Group should have a continuing nature, since

- Short-Term Issues (3 to 6 months)
- Medium-Term Issues (6 months to 2 years) and
- Long-Term   Issues (2 to 5 years)

were specified. (See last question below)

4) Dates of last and next meetings:

Last meeting - October IETF meeting

Next meeting - April IETF meeting

5) Mailing lists

No mailing lists are installed so far. Bill Melohn has offered to support such lists on SUN.COM.

6) Key players:

Mike Brescia, BBNCC, brescia@park-street.bbn.com
Thomas E. Brunner, SRI International, brunner@span.istc.sri.com
Ross Callon, BBNCC, rcallon@park-street.bbn.com
Noel Chiappa, MIT, jnc@xx.lcs.mit.edu
Bill Melohn, Sun Microsystems, melohn@sun.com
Carl-H. Rokitansky, DFVLR, roki@a.isi.edu

To keep the group a workable size it should probably not exceed 10 members.

Performance and Congestion Control Working Groups

Allison Mankin (MITRE)

mankin@gateway.mitre.org

1) Brief statement of charter and goals

Charter is to collect and develop short-term techniques of improving Internet performance, methods which like TCP ''Slow-start'' are retrofittable, inexpensive to implement, and contribute to globally better use of network resources. After a preliminary draft of a paper covering all Internet performance enhancement methods, it was decided to divide the material. Three RFCs are planned, whose tentative titles are:

Specification of Slow-start TCP
Gateway-Based Congestion Control
Proposal to Eliminate Source Quench

2) Progress to date

Produced a preliminary draft of guidelines for performance enhancement of IP, TCP, and a number of applications. Reviewed the draft at Annapolis meeting, decided at Ann Arbor meeting that the paper should be divided into the three listed above. This decision was encouraged by a suggestion from the Host Requirements WG that the documentation of TCP congestion control be separated and speeded up.

3) Estimate of timeframe for completion

Twelve months or more.

4) Dates of last and next meeting

Last meeting: Ann Arbor, October 16
Next meeting: Austin, January 18

5) Name of WG mailing lists

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

6) Names of key players (also include liaisons with other WGs or TFs)

Our attendance averages 20, though many attendees are observers. The following are the members who have contributed writing or editing so far (this list is as accurate as possible given the chair's case of the flu while listing it):

Art Berggreen	ACC
Dave Borman	Cray
Van Jacobson	LBL
John Lekashman	NASA/NAS
Allison Mankin	MITRE
Craig Partridge	BBN
K.K. Ramakrishnan	DEC
Bruce Schofield	DCEC

We are cooperating with the Host Requirements and Connection-Oriented IP WGs (the liaison people include John Lekashman for the former and Claudio Topolcic for the latter). We should have some liaison with the End to End Task Force, but we don't.

Point-to-Point Protocol Working Group

Drew Perkins (CMU) and Russ Hobby (UC Davis)

1) Statement of charter and goals:

The goal of the Point-to-Point Protocol Working Group is to publish an RFC defining a standard protocol for the encapsulation of IP Datagrams over point-to-point links including asynchronous and synchronous serial lines. The protocol will include encapsulation schemes as well as an extensible option negotiation protocol allowing negotiation of IP addresses, data compression, etc.

2) Progress to date:

The first step towards this goal was to document the requirements for such a protocol. A draft RFC discussing these requirements was sent to the IETF mailing list in October and will hopefully be published as an RFC soon. The purpose of this document is to make sure that everyone in the working group is aware of all the various issues. The second step is the definition of the standard protocol. A first draft defining the basic encapsulation scheme has been mailed to the PPP mailing list for review.

3) Estimate of timeframe for completion

The current estimate for completion is approximately April (or the next IETF meeting after the January meeting).

4) Dates of last and next meeting

The last meeting was at the October IETF, the next will be at the January IETF.

5) Name of WG mailing lists (if any)

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

6) Names of key players

Drew D. Perkins, ddp@andrew.cmu.edu
Philip Prindeville, philipp@oliver.cs.mcgill.ca
Russ Hobby, rdhobby@ucdavis.edu

ST and Connection-Oriented IP Working Group

Claudio Topolcic (BBN)

1) Statement of charter and goals

Produce a specification for the ST protocol that can be implemented by people outside the current small group of interested people and will support research in connection-oriented internet level protocols. Produce a gateway implementation of this protocol and at least one or two host implementations. Perform relevant experiments and gain experinece. Produce a specification for a next generation connection like protocol if the results of the preceeding experiments warrant it.

2) Progress to date

We have a preliminary draft of the ST specification, and we are talking it over and working toward a better draft. We have host implementations based on an older version of ST. We are almost done building a gateway implementation based on an older version of ST. We have a plan for how to look into producing a follow-on protocol. We have an outline of a "requirements document" which is the first step in this plan. We have not published any papers.

3) Estimate of timeframe for completion

The gateway implementation based on the older version of ST should be available in about 2 months. The ST specification should be available in 2 or 3 months. The host and gateway implementations based on the new ST specification should be available within six months of the specification, or about 8 or 9 months from now. The requirements document should be done within 3 months or so. The specification for the follow on protocol should be done in about a year.

4) Dates of last and next meeting

The last meeting was on Oct 17 1988 at Ann Arbor Michigan.
The next meeting will be on January 17 1989 at Austin Texas.

5) Name of WG mailing lists

Mailing list is "cip@bbn.com".

6) Names of key players

Claudio Topolcic, BBN, chairman, Ross Callon, DEC, Steve Casner, ISI, Phil Park, BBN, Guru Parulkar, Washington University, KK Ramakrishnan, DEC.

TELNET

TELNET Linemode Working Group

David Borman (Cray)

dab@cray.com

1) Statement of charter and goals

The TELNET Linemode working group is writing an RFC to describe a standard method of doing line mode TELNET (pushing the character processing to the front end when ever possible, and only sending completed lines across the network)

2) Progress to date

A draft RFC (IDEA16) has been produced. See below for timeframe to completion.

3) Estimate of timeframe for completion

The draft document (IDEA16) has been re-worked to be very close to what the final RFC will look like. The next meeting should be the last meeting needed to reach closure.

4) Dates of last and next meeting

Last meeting - Ann Arbor IETF
Next meeting - this next IETF

5) Name of WG mailing lists

linemode@uc.msc.umn.edu
linemode-request@uc.msc.umn.edu - To be added or deleted

6) Names of key players (also include liaisons with other WGs or TFs)

We have met twice. Below are all the people who have attended meetings, and which meetings they attended.

1	2	David Borman
1	2	Mike Karels
1		Bruce J Schofield
1	2	Louis A. Mamakos
1		Stuart Levy
1		Coleman Blake
1		David Wasley
1		Allan Fischer
1		Philip Prindeville
	2	Joyce Reynolds
	2	Bill Westfield
	2	Allen Cole

User Services Working Group

Karen Bowers (NRI)

bowers@sccgate.scc.com

This is a new working group. The first meeting will be held at the January IETF meeting in Texas. The draft charter and proposed goals are listed below. This will be finetuned at the initial meeting. A mailing list has not yet been established. For more information, send email to Karen Bowers (bowers@sccgate.scc.com).

The information below is organized as:

1) draft Charter with Key Objectives,
2) Selection Criteria for determining what issues/actions should be undertaken first
3) Issues/Actions for Consideration

CHARTER (draft): to provide a liaison among existing and newly forming network informations centers, network managers and the broad network user community.

Objectives: to consolidate and enhance the tools of existing user assistance and information services and make these pooled resources universally available to novice and experienced users alike.

to develop new and innovative network information/directory assistance techniques/methods in terms of general user support services (not technology-specific applications).

SELECTION CRITERIA (for projects/requirements to be addressed by the User Services Working Group):

1. Project/selected action must lend itself to accomplishment within a reasonable timeframe (say 1-3 years).

2. Must culminate in a measurable/quantifiable end result (production oriented; e.g. RFC, network users directory, etc.)

3. Must address user assistance needs and not technology specific requirements (e.g. routing)

4. Products/tools resulting from these efforts must not only address user information requirements but must be designed to be both maintainable and easily "updateable".

ISSUES/ACTIONS FOR CONSIDERATION (to be further expanded):

- A national directory (or directories) of existing networks and associated points of contact to include:

  1. short/concise description of each network, net #, and net maps,
  2. POCs for various actions: permission to connect, network engineering, network ops, 800#s, support services (such as assistance with routing/performance problems), etc.
  3. a standardized format describing how to connect: permission requirements, network specific procedures, guidance on physical (circuit/equipment) interface requirements and software (protocol) requirements, and Internet specific procedures (initial configuration requirements: net # assigned, name server, subnets, hand con configure routing tables...)

- How to set up and establish national procedures for net connections (=RFC).

- How to best answer new user problems

- A national (an international) network user directory (as a phone book)

- A guide to user training resources

3. IETF ATTENDEES

IETF ATTENDEES

The following is a list of people who attended all or part of the October 1988 IETF meeting. All organization affiliations are listed as submitted, and for brevity have not been expanded (Example: DCA vice Defense Communications Agency).

Name	Organization	Email Address
Almes, Guy	Rice University	almes@iapetus.rice.edu
Almquist, Philip	Stanford/self	almquist@Jessica.Stanford.EDU
Aronson, Cathy	Merit	CJA@merit.edu
Atlas, Stephen	BBN	Satlas@BBN.COM
Beeman, Roger	NWnet(Boeing)	beeman@boeing.com
Berggreen, Art	ACC	art@acc.arpa
Blunk, Larry	UMich	ljb@merit.edu
Boivie, Rick	IBM	rboivie@ibm.com
Borman, Dave	Cray Research	dab@cray.com
Bosack, Len	cisco Systems	Bosack@mathom.cisco.com
Braden, Bob	USC-ISI	braden@isi.edu
Bradley, Terry	Wellfleet Comm	(617) 275-2400
Bratton, Eric	UMich	ericb@caen.engin.umich.edu
Braun, Hans-Werner	UMich	hwb@mcr.umich.edu
Brescia, Mike	BBN	BRESCIA@BBN.COM
Brim, Scott	Cornell	SWB@DEVVAX.TN.CORNELL.EDU
Broersma, Ron	NOSC	ron@nosc.mil
Burruss, John	Wellfleet Comm	(617) 275-2400
Callon, Ross	DEC	callon%erlang.dec.com @decwrl.dec.com
Carpenter, Geoff	IBM Research	GCC%YKTVMX@CUNYVM.CUNY.EDU
Casner, Steve	USC/ISI	CASNER@ISI.EDU
Chiappa, J. Noel	MIT/Proteon	JNC@xx.lcs.mit.edu
Chinoy, Bilaz	Merit	bnc@merit.edu
Choy, Joe	NCAR	choy@ncar.ucar.edu
Chung, Anthony	Sytek	(415) 966-7430
Cohrs, Dave	U. of Wisconsin	dave@cs.wisc.edu
Cole, Allen	Univ. of Utah	cole@cc.UTAH.EDU
Collins, Mike	LLNL	collins@NMFECC.ARPA
Draughon, Phil	Northwestern	jpd@accuvax.nwu.edu
Drescher, J. E.	IBM Corporation	
Fedor, Mark	NYSERNET	Fedor@nisc.nyser.net
Finkelson, Dale	Midnet	dmf@westie.unl.edu
Frank, Randy	U Mich	frank@caen.engin.umich.edu
Geretz, Lionel	ACC	lionel@acc.arpa
Gerich, Elise	NSFNET/Merit	epg@merit.edu
Gerlach, Chuck	AT&T	cag@iwlcs.att.com
Gilligan, Bob	Sun	Gilligan@Sch.COM
Gross, Phill	NRI	gross@sccgate.scc.com
Gross, Martin	DCA	MARTIN@PROTOLABA.DCA.MIL
Hain, Tony	LLNL	Hain@NMFECC.ARPA
Hares, Susan	MERIT/NSFNET	skh@merit.edu

Hastings, Gene	PSC	hastings@morgul.psc.edu
Hobby, Russ	UC DAVIS	RDHOBBY@UCDAVIS.EDU
Hunter, Steven	LLNL	Hunter@NMFECC.Arpa
Jacobsen, Ole	ACE	ole@csli.Stanford.edu
Jacobson, Van	LBL	Van@LBL-CSAM.arpa
Jordt, Dan	U of Washington	danj@blake.acs.washington.edu
Karels, Mike	UCBerkeley	karels@acbarpa.Berkeley.edu
Karn, Phil	Bellcore	Karn@thumper.bellcore.com
Katz, Dave	Merit	Dave_Katz@um.cc.umich.edu
Knopper, Mark	Merit	MKnopper@Merit.edu
Krol, Ed	U. of Illinois	Krol@uxg.cso.uiuc.edu
LaBarre, Lee	MITRE	cel@mitre_bedford.arpa
Lakey, Jerry	Merit	JLL@merit.edu
LaQuey, Tracy	UTexas-Austin	tracy@emx.utexas.edu
Lazear, Walt	MITRE	lazear@gateway.mitre.org
LeKashman, John	NASA	lekash@orville.nas.nasa.gov
Lepp, Marianne	BBN	mlepp@bbn.com
Lottor, Mark	SRI	MKL@SRI-NIC.ARPA
Love, Paul	SDSC	LOVEEP@SDS.SDSC.EDU
Lowe, Ken	U of Washington	KEN@BLAKE.ACS.WASHINGTON.EDU
Lynn, Charles	BBN	CLYN@BBN.Com
Malkin, Gary	Proteon	GMALKIN@PROTEON.COM
Mamakos, Louis	Univ. of Md	lovie@trantor.umd.edu
Mankin, Allison	MITRE Corp.	mankin@gateway.mitre.org
Marshall,George	Ungerman-Bass	
Mathis, Matt	PSC	mathis@faraday.ece.CMU.edu
McCloghrie, Keith	Wollongong	kzm@twg.com
Medin, Milo	NASA/NSI	medin@nsipo.nasa.gov
Melohn, Bill	Sun Micro	Melohn@Sun.COM
Merritt, Don	BRL	merritt@BRL.MIL
Mockapetris, Paul	ISI	PUM@isi.edu
Morris, Don	NCAR	morris@ncar.ucar.edu
Moy, John	Proteon	jmoy@proteon.com
Mundy, Russ	DDN(DCS B602)	mundy@beast.ddn.mil
Natalie, Ron	Rutgers	Ron@Rutgers.Edu
Nguyen, Carolyn	AT&T	mhn@caelum.att.com
Nitzan, Rebecca	ESNET,DOE LLNL	NITZAN@NMFECC.ARPA
Norton, Bill	Merit	wbn@merit.edu
Opalka, Zbigniew	BBN	Zopalka@BBN.COM
Park, Phillipe	BBN	PPARK@BBN.COM
Parker, Paul	CMU	PAUL.PARKER@CS.CMU.EDU
Partridge, Craig	BBN STC	craig@nnsc.nsf.net
Parulkar, Guru	Washington Univ.	guru@flora.wastl.edu
Perkins, Drew	CMU	ddp@andrew.cmu.edu
Petry, Mike	Univ. of Md	petry@tranton????
Prindeville, Philip	McGill Univ.	philipp@cs.mcgill.ca
Ramakrishnan, K.	DEC	rama%erlang.dec@decwrl.dec.com
Reichlen, Gladys	MITRE	reichlen@gateway.mitre.org
Rekhter, Jacob	IBM	yakov@IBM.COM
Reschly, Robert	BRL	reschly@brl.mil
Reynolds, Joyce	USC/ISI	jKREY@ISI.EDU
Rochlis, Jon	MIT	jon@athena.mit.edu
Rokitansky, Carl	Fern U.Hagen	roki@DHAFEU52.Bitnet
Schiller, Jeff	MIT	jis@bitsy.mit.edu

Schofield, Bruce	DCEC	SCHOFIELD@EDN-UNIX.ARPA
Sheridan, Jim	IBM	jShERIDA@IBM.COM
Spafford, Gene	Purdue CS/SERC	spaf@cs.purdue.edu
St. Johns, Mike	DDN	StJohns@beast.ddn.mil
Stahl, Mary	SRI(NIC)	STAHL@SRI-NIC.ARPA
Stine, Bob	SPARTA	
Stone, Geof	Ntwk Sys Corp.	geof@nsco.network.com
Thixton, Cal	NeXT	Cal_Thixton@NeXT.COM
Ticknor, Paul	NASA	ticknor@prandtl.nas.nasa.gov
Topolcic, Claudio	BBN	topolcic@bbn.com
Veach, Ross	UIUC	RRV@UXC.CSO.UIUC.EDU
Vielmetti, Edward	UMich	emv@umix.cc.umich.edu
Waldbusser, Steve	CMU	waldbusser@andrew.cmu.edu
Ward, Carol	Westnet	cward@spot.colorado.edu
Warrier, Unni	Unisys	unni@cs.ucla.edu
Westfield, Bill	cisco Systems	BillW@cisco.com
Wilder, Rick	MITRE	rick@gateway.mitre.org
Wolff, Steve	NSF	steve@note.nsf.gov
Yu, Jessica	Merit	jyy@merit.edu

4. FINAL AGENDA

Final Agenda, 17-19 October 88 IETF

This was the final agenda for the October 17-19 IETF meeting at the University of Michigan in Ann Arbor. The meeting was hosted by Hans-Werner Braun and Elise Gerich of Merit.

MONDAY, OCTOBER 17

9:00 am Opening Plenary, Introductions and Local Arrangements

9:30 am Working Group Morning Sessions

- o Host Requirements, Members Only (Braden, ISI)
- o ST and Connection-Oriented IP (Topolcic, BBN)
- o CMIP-Over-TCP Net Management (Lee LaBarre, MITRE)
- o Interconnectivity and EGP3 (Almes, Rice)
- o Open SPF IGP (Petry, UMD and Moy, Proteon)

12:00 pm Lunch

1:30 pm Working Group Afternoon Sessions

- o Host Requirements, Open (Braden, ISI)
- o ST and Connection-Oriented IP (Topolcic, BBN)
- o CMIP-Over-TCP Net Management (Lee LaBarre, MITRE)
- o Interconnectivity and EGP3 (Almes, Rice)
- o Management Information Base (Partridge, BBN)

5:00 pm Recess

7:30 pm o Working Group for Joint Monitoring Access for Adjacent Networks focusing on the NSFNET Community (Hares, Merit)

TUESDAY, OCTOBER 18

9:00 am Opening Plenary
9:15 am Morning Working Group Sessions

- o Host Requirements, Members Only (Braden, ISI)
- o TELNET Linemode (Dave Borman, Cray)
- o Authentication (Schiller, MIT)
- o Performance and Congestion Control (Mankin, MITRE)
- o Point-Point Protocol (Perkins, Hobby, Prindeville)
- o PDN Routing (Rokitansky, FernUni Hagen)

11:30 am Lunch

1:00 pm Opening Plenary Statement (Gross, MITRE)

1:15 pm	Network Status Reports

- o Merit NSFnet Report (Braun, UMich)
- o IBM NSFnet Report (Drescher, IBM)
- o Arpanet/DDN Report (Lepp, BBN)
- o DDN Report (Brescia, BBN)
- o Interop 88 Network Report or 'How to build a complex internet in 2 days' (Almquist)

3:30 pm	Break
3:45 pm	Network Performance Presentations

- o Packets Over A Different Kind Of Ether, including Amateur Packet Radio Demonstration (Karn, Bellcore)
- o Keeping The Usual Ether Filled Up With High Performance TCP (Jacobson, LBL)

5:00 pm	Recess
7:00 pm	NSFNET NOC Tour

WEDNESDAY, OCTOBER 19

9:00 am	Congestion Control Observations Using NETMON (Mankin, MITRE)
9:30 am	Working Group Reports and Group Discussion

- o Authentication (Schiller, MIT)
- o CMIP-over-TCP (CMOT) (LaBarre, MITRE)
- o Interconnectivity (Brim, Cornell/Lepp, BBN)
- o Host Requirements (Braden, ISI)
- o Internet MIB (Partridge, BBN)
- o Joint NSFNET/Regional Monitoring (Hares, Merit)
- o Open SPF-based IGP (Petry, UMD)
- o Open Systems Routing (Lepp, BBN)
- o PDN Routing (Rokitansky, FernUni Hagen)
- o Performance and CC (Mankin, MITRE)

12:00 pm	Lunch
1:00 pm	Working Group Reports and Group Discussion (cont'd.)

- o Pt-Pt Protocol (Perkins, CMU)
- o ST and CO-IP (Topolcic, BBN)
- o TELNET Linemode (Borman, Cray)

1:45 pm	What Is Usenet?, What Is NNTP? (Spafford, Purdue)
2:30 pm	The NIC Domain Chart (Lottor, NIC)
2:45 pm	On Some T1 Satellite Link Performance (Lekashman, Ames)
3:15 pm	Concluding Plenary Remarks
3:30 pm	Adjourn (Rush to Airport)

5. WORKING GROUP REPORTS/SLIDES
ANN ARBOR, MI ACTIVITIES
17-19 OCTOBER 1988

Authentication

Jeff Schiller
MIT

Requirements document
SNMP authentication
Kerberos RFC

SNMP authentication:

using "communities" mechanism

define catagories of authentication protocols:

catagories	type
trivial	des
password	kerberos
safe	rsa
private	other

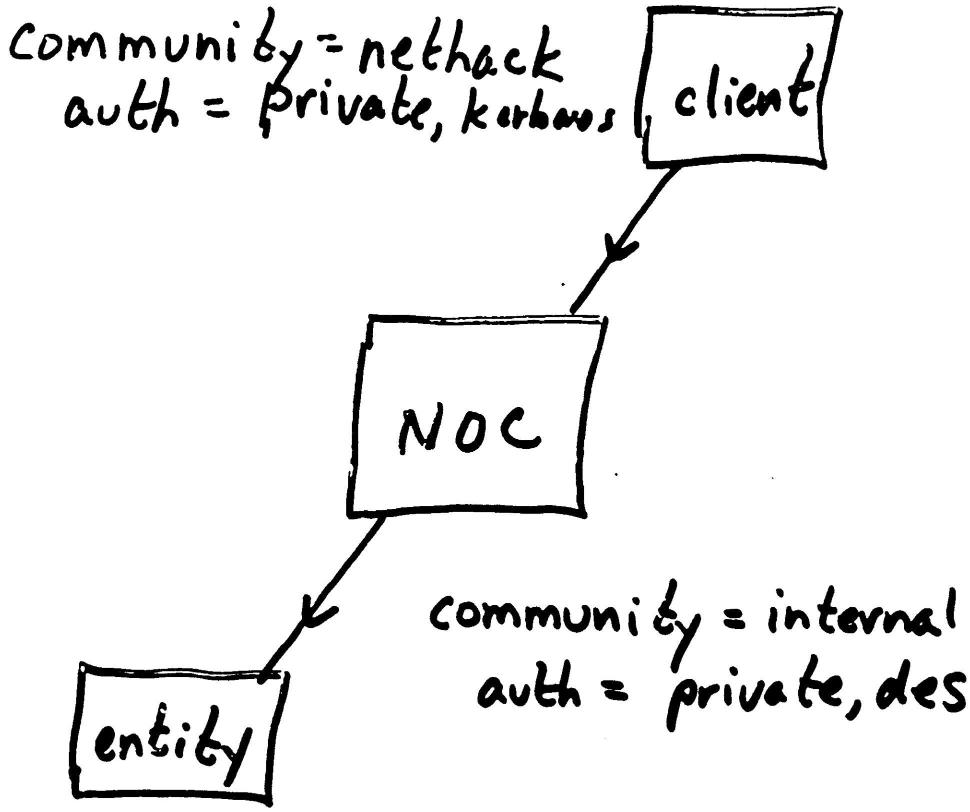

monitored entities are simple
complexity and access control
centralized at NOC

CMIP-over-TCP (CMOT)
(NETMAN)

Lee LaBarre
MITRE

Unni Warrier
UNISYS

CMIP-over-TCP (CMOT) WG Report — 17-19 Oct 1988
Reported by Lee LaBarre — Ann Arbor, MI

The NETMAN (CMOT) WG met Oct. 17 and 18 at the IETF meeting in Ann Arbor, MI. The meeting occurred in two separate morning sessions.

Morning of Oct. 17

- o Lee LaBarre provided a review of the groups charter, goals, and status as stated in IETF form 2.
- o The group defined a set of issues for consideration by the IETF

MIB Working Group, including:

- definition of the MIT (naming or object instance tree),
- the distinction between Object class and attribute,
- definition of distinguished attributes for objects,
- the specification of optional attributes in the MIB, and the impact on aggregate objects, e.g., table entries,
- the need for definition of procedures and objects for event and security management,
- the definition of thresholds.

These issues were raised in the MIB WG meeting, and all but thresholds received priority consideration for work this year. Work on thresholds will depend on contributions from the NETMAN WG, and is contingent on the existence of an event control mechanism.

Morning Oct. 18

- o We decided the NETMAN agreements would include the entire CMIS/P, ROSE, and ACSE protocol set, but stipulate a mandatory subset of services.
- o Recommendations were suggested for modifications to IDEA0017, the thin presentation layer, including:
  - investigate necessity for multiple PCIs, e.g., ROSE, ACSE,
    CMIP version, MIB version,
  - use of transports other than TCP and UDP, such as VMTP, etc.
  - negotiation of transport protocol for desired QOS,
  - investigate the multiplexing of associations across a single TCP connection,

- A decision was made to develop a proxy mechanism based on the use of object instance structure. This would minimize the number of associations and TCP connections used for proxy. It would also work in chaining a request through multiple managers.

- We reviewed the CMOT agreements document drafted by Unni Warrier and suggested revisions where appropriate. Lee LaBarre and Unni Warrier agreed to contribute new text to the document.

- The issue of specifying alternative QOS for management purposes was raised by Keith McCloghrie. He suggested that only low quality (UDP) service should be specified since manager applications might have to be prepared to deal with either QOS anyway, and UDP would place the lowest burden on agents. This issue will be addressed at the next meeting.

- The distribution list for the demo was opened up to a wider membership and the name changed to netman@gateway.mitre.org.

Slides Attached

NetMan Working Group Report

Internet Engineering Task Force (IETF)

Unni Warrier

(213) 829-7511 x5694

UNISYS Corporation
Distributed Architecture
Defense Systems
2400 Colorado Avenue
Santa Monica, CA 90406

CHARTER

Charter: As described in RFC1052.

- Develop a long term approach to management of the Internet based on the OSI Network Management Framework and the Common Management Information Protocol (CMIP).
- Provide input to the OSI standards process based on experience in the Internet, and thereby influence the final form of OSI International Standards on network management, in particular CMIS/P.

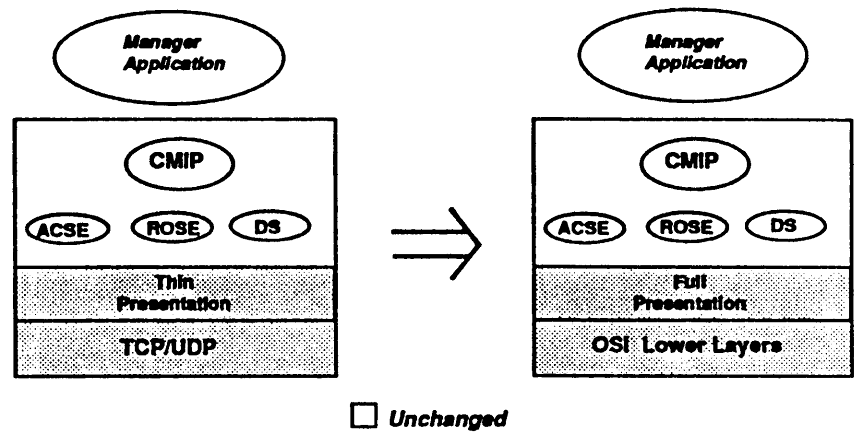

TCP to OSI Transition

Figure 2

APPROACH

a. Develop prototype implementors agreements on CMIP over TCP.

b. Develop prototype implementations based on the CMOT agreements and IETF SMI and MIB agreements.

c. Experiment with CMOT and extensions to the SMI and MIB.

d. Develop final implementors agreements for CMOT.

e. Promote development of products based on CMOT.

f. Provide input to the OSI Network Management standards process in time to effect the International Standards.

MEMBERSHIP

- Expected duration of group:
  The groups work should be completed by June 1989.
- Membership:
  Open

ACTIVITIES

Achieved goals:

a. Overview document (IDEA0012)
   Thin Presentation layer (IDEA0017)
   Prototype Implementors Agreements (IDEA0025)

b. Nine vendor prototype implementations demonstrated at INTEROP 88 in Santa Clara, CA.

c. Experimentation occurred during development of the INTEROP demo, and is continuing.

d. Draft implementors agreements are written for the DIS CMIP over TCP (CMOT). Proposals for extending the SMI and MIB are in progress.

ACTIVITIES (Cont.)

e. Thirteen corporations participated in the INTEROP 88 demo. Nearly all the vendors in that group have indicated that they expect to field products during 1989 based on CMOT implementors agreements.

f. Several Working Group members are participating in the OSI network management standards organizations and carrying the CMOT experience into that forum.

NETMAN WORKING GROUP REPORT

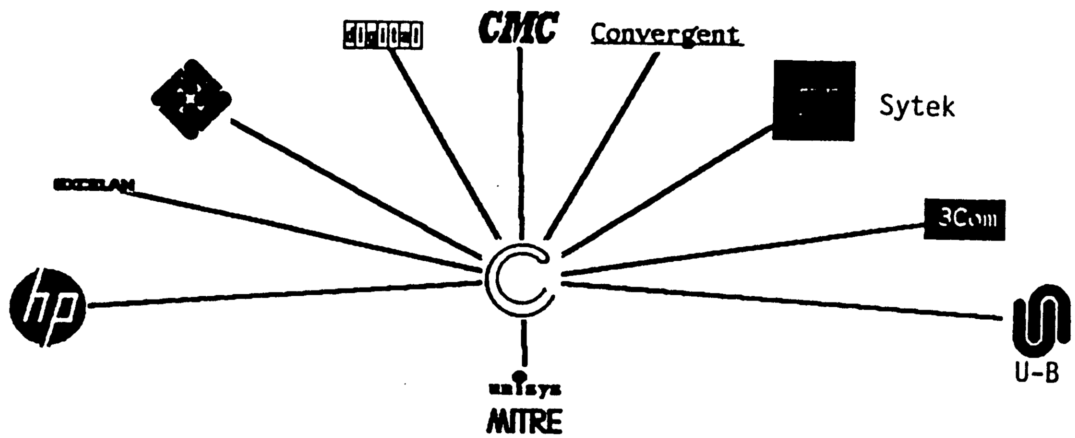

NetMan DEMO Participants

Advanced Computing Environments
Communications Machinery Corporation
Convergent Technologies
Digital Equipment Corporation
Epilogue Technology Corporation
Excelan
Hewlett Packard Corporation
MITRE Corporation
SUN Microsystems
Sytek
3COM Corporation
Ungermann-Bass
UNISYS Corporation

DEMO SCENARIO

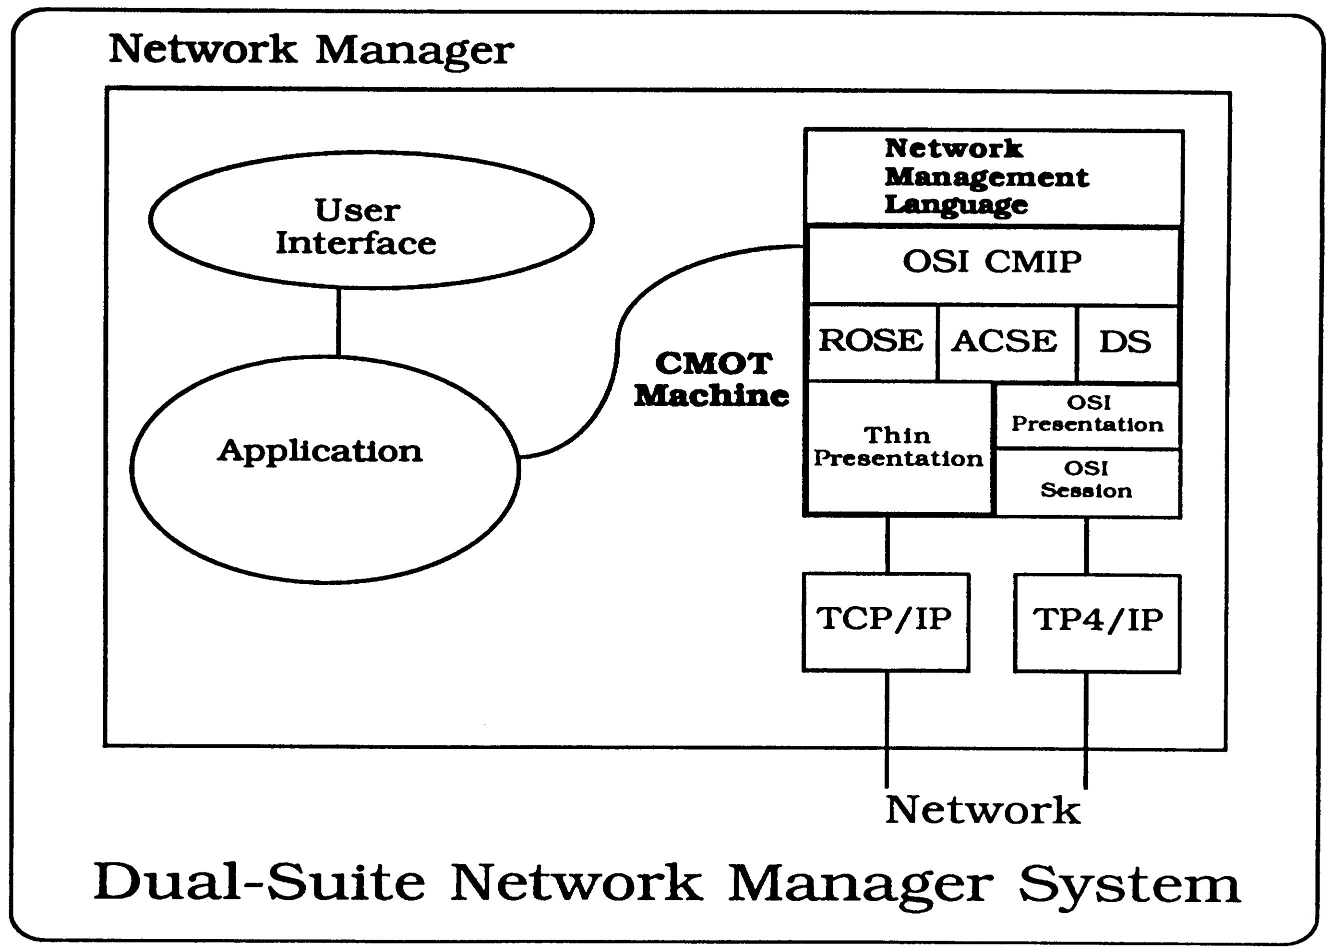

DEMO SCENARIO (Cont.)

- Variables aligned with MIB WG RFCs
- Experimental extensions for Events and Thresholds
- GETs
- SETs
- Event Reports
- Tables, rows of tables (aggregate objects)

LESSONS from DEMO

- Events worked.
- Thin Presentation worked.
- Agent code size as low as 40K on terminal server, PC.
- Modify agreements from this experience.

FUTURE WORK

- Align with DIS, new Implementer's agreement (RFC- CMOT) in progress.
- Extensions to MIB, SMI proposals in progress.
- IDEA17 (Thin presentation) wrap-up and make into RFC.

10/19/88

CMOT (NETMAN) WG

- Reviewed group charter/goals/status
- Defined issues for MIB WG to consider
  - object class/attribute/instance
  - distinguished attributes for objects
  - optional objects/attributes in MIB
  - Event management
  - thresholds
- Agreed develop RFC on total DIS CMIS/P, ROSE, ACSE and stipulate mandatory subset
- Recommend mods to IDEA017 (Thin Presentation)
  - multiple PCIs
  - map to other transports (TCP, UDP, VMTP?...)
  - negotiation for transport
  - multiplexing associations on one TCP conn.
  - use of Domain name service

10/19/88

CMOT (NETMAN) WG (cont)

- Develop proxy mechanism based on object instance
    - to reduce number of associations / TCP connections
- Reviewed CMOT Draft RFC
- Change name mail distribution list & open to new members

nmdemo88@gateway.mitre.org

⇓

netman@gateway.mitre.org

Host Requirements

Bob Braden
USC-ISI

IETF Host Requirements Working Group

REPORT FROM ANN ARBOR IETF MEETING
October 17-19, 1988
Bob Braden

I. INTRODUCTION

The Host Requirements Working Group met for 1.5 days at the IETF meeting in Ann Arbor, Michigan. This meeting was very important, since the Host Requirements RFC has reached a stage when it seems to be nearly finished, and because we are rapidly approaching our self-imposed deadline, the end of calendar 1988.

All discussions were based on the October 11, 1988 version of the spec.

II. SESSIONS AND ATTENDEES

* Monday, October 17, Morning Session

The Working Group met in closed session, with the following attendees:

Bob Braden (ISI), Dave Borman (Cray Research), Noel Chiappa (Proteon/MIT), Phil Karn (Bellcore), John Lekashman (NASA), Mark Lottor (SRI-NIC), Charlie Lynn (BBN), Paul Mockapetris (ISI), Allison Mankin (Mitre), Craig Partridge (BBN/NNSC), Drew Perkins (CMU), Bruce Schofield (DCEC), and Cal Thixton (NeXT).

Allison Mankin and Dave Borman both took minutes. A list of outstanding issues formed the agenda.

* Monday, October 17, Afternoon Session

The Working Group invited all interested people to an open session, in which the assembled group went through the entire document, section by section. There were 25 attendees, and most of the group kept picking the carcass clean until 6:30PM! Now, THAT is dedication. On the other hand, no one had been able to read the document all the way to the end, so that comments were quite sparse for the Application Layer and non-existant for the Support Services.

All those attending in the morning attended in the afternoon (except for Craig Partridge, who had to chair another meeting).

Additional people in the open session were:

Almquist (Stanford), Collins (MFENET-II), Gilligan (Sun), Jacobson (LBL), Karels (UCB), Katz (UMich), Melohn (Sun), Nitzan (MFENET-II), Opalka (BBN), Parker (CMU), Rochlis (MIT Athena), Schiller (MIT), and Westfield (Cisco).

Dave Borman again took minutes, for which I am immensely grateful.

* Tuesday, October 18, Morning Session

A final closed meeting was held, with Braden, Chiappa, Karn, Lekashman, Mockapetris, and Partridge in attendance. The group dealt with the remaining issues from the original list, and with some of the new issues raised at the open session. The major discussion item was Dead Gateway Detection.

III. QUESTIONS AND DECISIONS

We now summarize the important points that were raised in all these meetings, both those that were decided and those that are still undecided.

Introduction

o Section 1.1.4 Embedded Gateway Code

Suggested: there are advantages to embedded gateway functionality other than simple convenience [Melohn]. ACTION ITEM: Draft some text: Melohn.

Link Layer

o Section 2.3.1 Trailer Negotiation

Agreed: Need a definition of how trailer negotiation is done. ACTION ITEM: Draft some text: Karels.

o ARP

Agreed: ARP implementation MUST hold onto at least one packet [the most recent] destined for a given unresolved target address.

o Section 2.2.2 ARP Cache Validation

Suggested: the ARP cache timeout time of 60 seconds currently specified is much too short [Jacobson]; this is because ARP cache timeouts generate traffic that increases quadratically <<Ed: somewhat faster than linearly?>> with the number of hosts on the Ether; timeout should be at least 5 minutes.

<<Ed: The discussion of ARP cache validation in the current draft is based on experience at CMU with a particular timeout algorithm. Two specific ARP cache algorithms have been proposed, and one or both should be written up as RFC's. The argument for a 5 minute timeout is based on the idealistic assumption that Proxy ARP is broken and ought to be abolished; however, Proxy ARP has many dedicated supporters.

While the quadratic argument is somewhat theoretical, lots of experience shows that it would be a mistake to ignore it. Itis unclear how to resolve this issue.>>

o Section 2.4 Link/Internet Layer Interface

Agreed: RFC ought to define interface, including upcall for dead gateway discovery.

Internet Layer

o Section 3.2.1.6 Type of Service

The Host Requirements spec requires TOS at all levels (application, transport, Internet) in order to break the chicken-and-egg problem with gateway implementations of Type-of-Service. A future "Assigned Numbers" RFC will include recommended values for the TOS bits for use by the major application protocols.

It seems likely that gateways will implement TOS by granting one TOS attribute (low delay, high throughput, or high reliability) while diminishing the others to some extent. Because of this and for simplicity, the recommended values will set at most one attribute bit.

Suggested: the Host Requirements RFC should give the philosophy of the bits, even though the actual recommended values are in Assigned Numbers.

Agreed: An application SHOULD be able to change TOS during lifetime of TCP connection, to support single-connection applications like SMTP. This MAY take the form of setting TOS on every SEND call.

Agreed: The TOS values in applications must be configurable, because we can only guess at the actual service effects of particular TOS bit combinations, and because particular hosts will want to tune the TOS values for special situations.

Agreed: TCP segments in each direction will have TOS determined by application on sending side. If the applications at the two ends specify different TOS values, then ACK's will come back with different TOS than was used to send the data.

Agreed: A transport protocol MAY communicate to its application the TOS with which incoming datagrams arrived.

- o Section 3.2.1.7 Time to Live

TTL: is it a time, or a hop count? This has been debated at length by the Working Group, was debated in both the closed and open sessions at Ann Arbor, and is still unresolved.

There is considerable sentiment in favor of redefining the TTL field as a pure hop count. However, the editor believes this would be a fundamental change to the architecture, which precludes making this change in the present Host Requirments document. Those who support the hop-count-only position need to make a cogent argument, considering all facets of the problem, in a published paper or RFC.

- o Section 3.2.1.5 Identification Field

Agreed: Drop recommendation to base Ident field on the triplet: (src, dest, prot).

- o Section 3.2.1.8 Source Route Options

The Editor detected some willingness in the open meeting to take the Editor's side, against source routing by hosts.

- o Section 3.2.1.9 Mis-Addressed Datagrams

Agreed: An IP layer MAY check each incoming datagram for a bogus source IP address.

- o Section 3.2.2.4 Time Exceeded

Agreed: ICMP Time Exceeded (Reassembly) may be used to trigger an MTU discovery procedure (see e.g. RFC-1063) when one is standardized, but the present document should specify that these ICMP messages are to be ignored.

o Section 3.2.2.5 Parameter Problem

Agreed: do not need new code for missing option.

o Section 3.2.2.6 Echo Request/Reply

Agreed: Record Route and Timestamp options are to be returned in the Reply, with the present host entered (ie, as if the echoing host were a hop in the path); the options will not be truncated.

o Section 3.2.2.9 Address Mask Request/Reply

This area has gotten a lot of attention <<Ed: more than it deserves>> from the WG, and discussion continues.

The open meeting and the closed WG meeting differed on the importance of a host implementing a dynamic way to learn the Address Mask (open: MAY, WG: SHOULD). There has been difficulty figuring out how to limit replies to authoritative sources. We cannot decide whether a statically configured address mask should take precedence over a dynamically determined one; people have arguments for both.

Agreed: a host with a statically-configured mask MUST NOT automatically be authoritative for address masks; to control this, the configuration needs an "Address Mask authoritative" flag.

Agreed: authoritative source for address mask reply may be a gateway or designated host(s) (e.g., a file server for diskless workstations).

Suggested: Internet architecture should logically pair address mask and IP address, so address mask for an interface should be determined at boot time by the same mechanism that is used to determine the IP address of that interface; therefore, ICMP Address Mask messages solve the wrong problem [Braden].

o Section 3.3.1.3 Route Cache

Suggested: although the present draft recommends that the route cache be based upon destination hosts, the use of destination networks as the cache key is an important optimization [Karels].

Suggested: Timing out routing cache entries is a bad idea because of scaling arguments [Jacobson]. Pinging of gateways in use is acceptable when neither lower-level nor higher-level advice are available.

Agreed: A route cache entry should include a timestamp indicating when the gateway was last set or updated.

The Working Group is quite clearly and lamentably confused on the entire issue of dead gateway detection. In the Working Group, several different approaches have been proposed, discussed, drafted into RFC's, and later rejected.

The conclusion from the Ann Arbor meeting was that the best we can do currently is to list the alternatives and state the arguments.

- o Section 3.3.4 Multihomed Hosts

Suggested: the model of multihoming contained in the current draft should be replaced by a different model [Karels].

TCP

- o Section 4.2.2.2 Use of Push

An animated discussion of the section on Push was a mixture of confusion and religion. Some believe in Push, some do not. The people who don't believe in Push (falsely) accused those who do of using Push to improve performance. Agreed: Push has nothing to do with performance, only correctness. In fact, the opposite it true: it is NOT pushing that can improve performance in some systems.

- o Various Sections

Van Jacobson suggested changes to clarify or correct the text concerning the relationships between the Nagle and slow-start algorithms, between slow-start and the older "retransmit only the front of the queue" rule, and between Push and the Nagle algorithm. He also suggested improvements in the discussion of delayed ACK's.

<<Ed: The discussion of TCP performance requirements is included in the Host Requirements RFC because the Performance Working Group has not yet completed their task of creating a comprehensive RFC on the subject. The discussion in the Host Requirements RFC is necessarily fragmentary>>

o Section 4.2.2.12 Retransmission Timeout

Agreed: change text to avoid implication that there must be a "retransmission queue" (implying that segment boundaries are recorded in this queue) [Karels]. There may be implementation advantages in deferring packetization until a segment is sent.

o Section 4.2.2.13 Shrinking Window

Agreed: document a pitfall -- when window shrinks from the right and in fact goes to zero [Karels].

o Section 4.2.3.2 Delayed ACK's

Agreed: current text omits an important advantage of delayed ACK's -- letting application have a shot at the CPU before an ACK is sent [Jacobson].

o Section 4.2.3.3 SWS

Agreed: modify sender-side SWS algorithm to handle windows smaller than MSS [Karels].

o Section 4.2.3.4 Connection Liveness

Agreed: The current draft, which specifies that connection liveness ought to be based upon retransmission count rather than time, is correct.

o Section 4.2.3.4 Keepalives

Suggested: mechanism that is documented in current draft is not the latest spiffy idea [Karels]. ACTION: supply text: Karels.

<<Ed: In general, the WG seems quite firm against TCP keepalives, although the current text in the document is rather wimpy on the subject.>>

o Section 4.2.? Data with Control

Agreed: a TCP MUST support data with a FIN bit, and SHOULD support data with a SYN bit [Karn].

o Section 4.2.3.12 Invalid Address

Agreed: A TCP should ignore any datagram addressed to a broadcast or multicast address [Karn].

o Section 4.2.? SYN Overload

Agreed: it is OK for a TCP to indicate overload by sending a RST in response to a SYN. However, it would be worthwhile to follow Charlie Lynn's suggestion of a text error message in a RST segment. An RFC is needed.

SMTP Section 5.1.2.1 VRFY

Agreed: there needs to be a new 4xx response defined for VRFY when it cannot get an answer (e.g., because domain lookup fails) [Barns].

TELNET Section 5.4.2.* Status Option

Agreed: SHOULD implement the Telnet Status option.

TFTP Section 5.3.* Broadcast requests

Agreed: TFTP SHOULD ignore transfer requests sent to a broadcast address.

BOOTING

Suggested: RFC should contain separate discussions of dynamic configuration and of booting [Perkins].

Agreed: BOOTP should be recommended, since it provides the most general solution to dynamic configuration, and since it works through gateways. However, the RFC ought also review the various partial solutions to dynamic configuration:

ICMP Information Request (=> Network number),
RARP (=> entire IP address),
ICMP Address Mask (=> Address mask).

BOOTP encompasses all of these and can also provide a list of default gateways.

ACTION: Write text about RARP: Melohn.

However, BOOTP is not sufficiently general to specify the configuration of all interfaces on a multihomed host. In this case, a host must either use BOOTP separately on each interface, or configure one interface using BOOTP and then access a file to configure the other interfaces.

Agreed: application layer configuration information will be taken from file(s), not obtained dynamically.

Agreed: there is a need for an IETF working group to create a general solution to the problems of dynamic configuration and booting, including the dynamic assignment of IP addresses.

Slides Attached

HOST REQUIREMENTS RFC

- Group Effort
  - 19 active contributors from 15 org's
    - 7 from vendors,
    - 12 univ, govt agency, res labs
- Intensive Effort
  - 5 meetings in 8 months
- Comprehensive -- all major protocols
- Explicit and detailed

- **OBJECTIVES:**

  - Define requirements
  - Point to essential documentation
  - Correct/update original documents
  - Fill gaps in specifications
  - Limit choices
  - Document known solutions to recurring issues

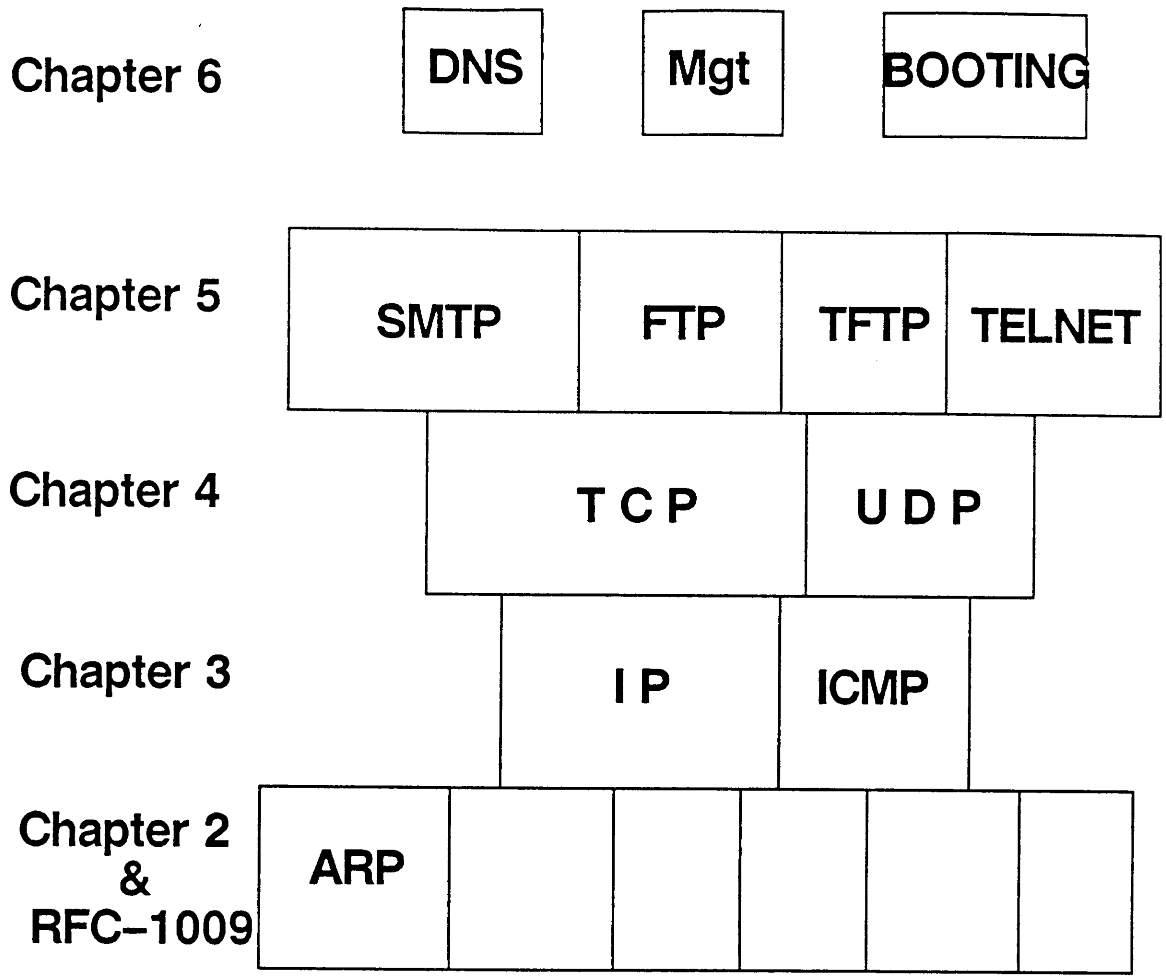
Chapter 6
DNS
Mgt
BOOTING
Chapter 5
SMTP
FTP
TFTP
TELNET
Chapter 4
T C P
U D P
Chapter 3
I P
ICMP
Chapter 2
&
RFC–1009
ARP

HOST REQUIREMENTS RFC

- Requirements ...
  - **"MUST" , "REQUIRED"**
- Recommendations ...
  - **"SHOULD" , "RECOMMENDED"**
- Options --
  - **"MAY" , "OPTIONAL"**

HOST CONFIGURATION

- ☞ IDEAL: Automagic
- ☞ REALITY: Not even close !!

- Require extensive configurability, but defaults will ease the pain.
- Many parameters must be adjustable --
  - ▶ Depend upon environment
  - ▶ Administrative requirements
  - ▶ Wizards are unsure
  - ▶ Interoperate with past mistakes

SOME HOST ISSUES

- Type-of-Service
  - To break deadlock, make it real
  - Host **must** implement mechanism
  - Recommended values in "Assigned Numbers" RFC
- Maximum Segment Size / Eff. MTU
  - Document current heuristic
  - Initiate experimentation with Mogul-Kent MTU-Discovery Option

    [IP-level reliable-delivery mechanism!]

    RFC-1063

- GATEWAY DISCOVERY
  - STATIC CONFIGURATION
  - DYNAMIC CONFIGURATION (BOOTP)
  - ICMP "GATEWAY DISCOVERY" MESSAGES

    ["Blew" the RFC]

- DEAD-GATEWAY DETECTION
  - LINK-LEVEL ADVICE (-) FROM:
    - ARPANET/MILNET
    - TOKEN RING
    - ETHERNET ???

- HIGHER-LEVEL ADVICE (±) FROM:
    - TCP
    - NFS (?)
    - ~~DNS~~
- GATEWAYS BROADCAST ICMP "UP-GATEWAYS" MSG
    - DRAFTED RFC
    - CURRENTLY DEAD

- "PING" ACTIVE GATEWAYS
  - SCALES BADLY
- TIMEOUT ROUTE-CACHE ENTRIES
  - SCALES BADLY (or worse)
- "WIRETAP" GATEWAY IGP
  - ARCHITECTURALLY UNSOUND

HOST: venera.isi.edu

Path:

pub/ietf-hosts.rfc.txt

Mailing list:

ietf-hosts-request@nnsc.nsf.net

Interconnectivity

Guy Almes
Rice University

INTERCONNECTIVITY WG

WHY? CURRENT ISSUES:

- EGP CAN'T SUPPORT WHAT MUST HAPPEN WITHOUT "TRANSIENT" LOOPS.

  MULTIPLE BACKBONES
  PEER<->PEER CONNECTIONS

- ORWG WON'T BE READY FOR "A WHILE"

- SOON OVER 1000 NETS

ELEGANCE IS NOT THE MAJOR ISSUE

CAN WE USE EGP3?

OTHER NEAR-TERM SMALL DESIRES

- MORE INFORMATION FOR MORE DECISION-MAKING BY MID-LEVELS.
- MAKE WHAT IS NOW IMPLICIT EXPLICIT, IN A PROTOCOL
- MORE INFORMATION FOR BOGON-HUNTING

GENERALLY, INCREASE KNOWLEDGE OF WHERE ROUTING INFORMATION CAME FROM; DON'T TRY TO BUILD LOOP-FREE PROTOCOL.

YES, WE CAN USE EGP3.

GOALS: (1) FINISH EGP3 CHANGES

(2) START ANOTHER COMPANION RFC; RECOMMENDATIONS AND WARNINGS, NORMATIVE USE IN INTER-AS ROUTING.

TRY TO MAKE IT GENERAL.

① ADD "PREVIOUS AS" TO POLL/UPDATE

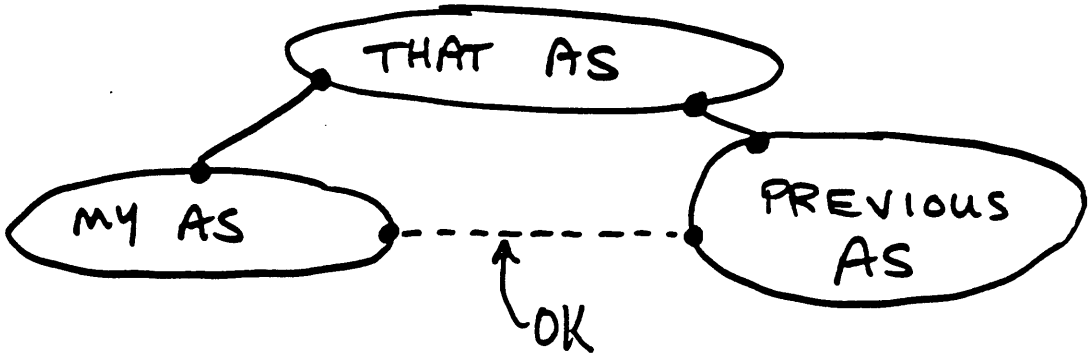

TABLE ENTRIES BY NET, NEXT GW, <u>PREV AS</u>

ALLOWS FORWARDING OF EGP INFORMATION ONE MORE HOP SAFELY.

WELL, WHY NOT ADD SOME MORE?

VARIABLE-LENGTH LIST OF ASs
ORIGINATING AS
PREV-PREV-AS

WOULD YOU REALLY USE IT?

WHAT TO DO WHEN CAN'T RESTRICT TO THREE AS HOPS?

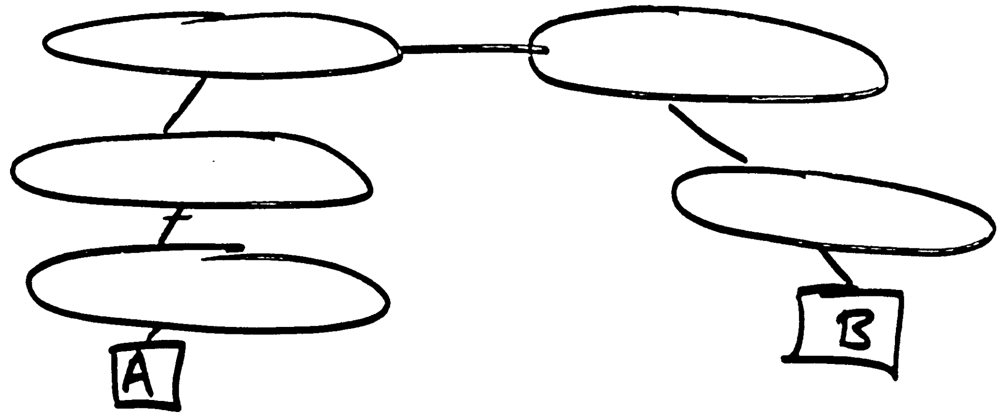

WHEN THERE IS NO OTHER CHOICE?
WITH STRONG CAVEATS,

② ADD "AS COUNT" = TOTAL # OF ASs ROUTING INFO HAS PASSED THROUGH.

MUST AVOID LOOPS!

RESTRICTION: IF INFO GOES DOWN, CAN NEVER GO UP AGAIN.

③ "THE REKHTER BIT": DIRECTION.

④ LATERAL AS COUNT = # OF SUCCESSIVE PEER BOUNDARIES CROSSED. $\boxed{-?}$

⑤ ENCODE ASSIGNED POSITION IN HIERARCHY !?

..

WHAT ABOUT METRICS?

- CURRENTLY EGP METRIC NOT USED (MUCH)

- ? USEFULNESS OF

  - OLD EGP METRIC - ASSIGN MEANINGS?
  - AS COUNT AS METRIC?
  - PRIMARY/SECONDARY INDICATOR? REAL USE NOT CLEAR.

WILL PROBABLY LEAVE SPACE FOR MULTIPLE METRIC TYPES/VALUES.

IGP (BBP) REQUIREMENTS

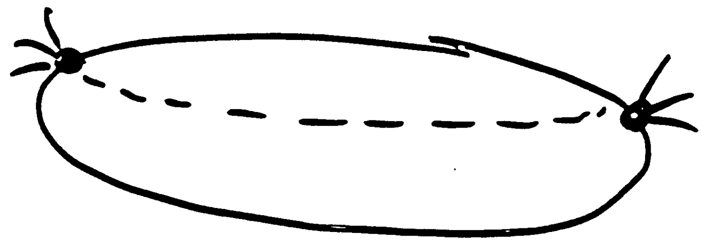

WILL HAVE

NEIGHBOR'S AS#
IP ADDRESS
YOUR OWN — " —
NEXT GW
PREVIOUS AS
AS COUNT
LATERAL COUNT
DIRECTION BIT

CAN PACK THESE INTO 24 BITS,
IF NECESSARY. WE'LL TALK.

another word about EGP3

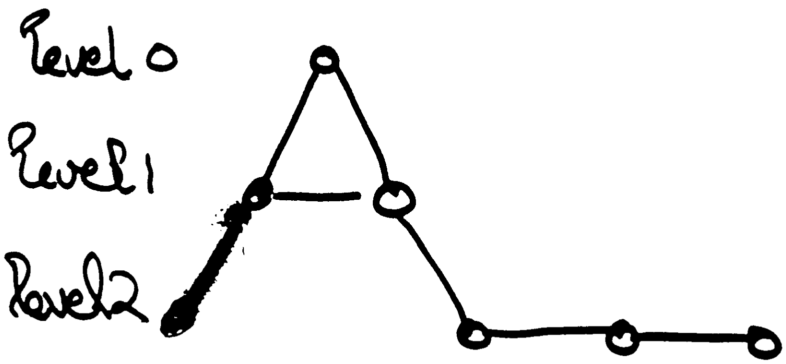

down bit — once down, set
Lateral bit

$\overset{0}{A} \longrightarrow \overset{1}{B} \longrightarrow x$ Last move

metrics — Probably not
Prim/secondary/tertiary
Do not propagate
information fields?

type	D	L	R	??	prev AS
4	1	1	2	8	

per net

need an explicit down

Internet MIB

Craig Partridge
BBN

[no report/slides provided]

NSFnet/Reg Monitoring
JOMAAN

Susan Hares
Merit, Inc.

NSFNET-Regional Meeting Report — 17 October 1988

Sue Hares, Chair, skh@merit.edu — NSFNET NOC @Ann Arbor

I. A collection of maps was distributed to all attendees. An effort to collect all kinds of maps will be made by Sue Hares.

   A. Maps of campuses, regionals, consortia, backbones should be sent in Postscript format to Sue Hares.

II. It was pointed out that On-line databases are kept at nis.nsf.net. Information such as Routing configurations are available.

III. Major discussion took place on the backup announcement of networks behind the regionals.

   A. Some major points:

      1. routing metric is interpreted *locally* by the NSS.
      2. multiple EGP peers can talk to one NSS with the same AS#.
      3. although the previous point is true, it was stressed that it is easier to manage the NSS when every peer has a unique AS#.
      4. every peer of an NSS should announce the shared net.
      5. NSFNET NOC needs one contact point within an AS#. Makes dealing with problems easier.

   B. Sue Hares discussed a "Cold Backup" strategy:

      1. configure two EGP neighbors.
      2. set egpmaxacquire to one.
      3. you would then peer with one at a time, trying the other only if you lost the first neighbor.
      4. must be careful with this because once your first neighbor came back up, you would not switch back to it until your second neighbor goes down. Sue Hares can configure this for your site if you wish. Contact her directly.

   C. Notification and confirmation of backup sites.

1. When adding new nets or when changing an additional network configuration, the NSFNET routing coordinator will make an effort to confirm the change with all parties involved. For example, checking with the primary announcer of a network before adding a secondary announcer for the same network.

2. The NSFNET routing coordinator will send out a mail message to NSFNET-SITE-PEOPLE notifying them of recent changes to the routing configurations. This message may be daily or as needed.

IV. SGMP/SNMP/CMOT based tools.

[ed. note: heavily involved in this discussion, notes are a bit scarce]

A. There are no CMOT based tools.

B. Shall we share SGMP sessions?

1. Some groups expressed concerns about the security of having one global SGMP session.

2. Concerns were expressed by many people regarding the changing of a global SGMP session every two weeks.

3. Agreed that it would be beneficial to all regionals and the NSFNET backbone to share SGMP information.

4. A read-only session called "monitor" should be added to all regional and NSFNET gateways by Friday, October 28, 1988. Progress on getting the sessions configured should be sent to Sue Hares and she will post a status report to NSFNET-SITE-PEOPLE.

5. It was suggested that people read Guy Almes's paper. It is available on the NIS machine. The NOC will post where it is located.

6. The common SGMP session must be considered private. Only the regional NOCs should be made aware of it. The session name will remain the same until there is some pressing need to change it.

C. There was loose consensus that problems between regionals should be hashed out by the concerned parties. The NSFNet backbone people would get involved in the event of a stand-still at solving the problem or in the case of an actual NSFNET backbone problem. Sharing SGMP information would make it a bit easier to pin-point the problem without NSFNET being the middle man.

D. Concern was expressed by certain regionals about lack of manpower in tracking down certain problems mentioned in the previous section without the help of NSFNET. At this junction, it was said that NSFNET would try and help.

E. Sue Hares of NSFNET will compile a list of available SGMP/SNMP tools. This will include vendors supporting SGMP/SNMP as well as public domain stuff. If you know of anything out there, please let her know.

V. Other Trouble-shooting tools.

A. Some other tools in detecting network problems were brought up. They were as follows:

1. Ping with record route.

a. doesn't show TTL exceeded.
b. will crash Ultrix.

2. Matt Mathis tool.

a. using TTL exceeded messages to trace the source of a route.

3. Ken Loewe's PC monitor program.

VI. Summary of Action Items.

A. Get your Postscript Maps to Sue Hares.

B. Make sure you are announcing the shared network to your NSS.

C. NSFNET routing coordinator should mail out messages to NSFNET-SITE-PEOPLE regarding routing changes.

D. Add an SGMP session called "monitor" to your regional gateways by October 28, 1988 and notify Sue Hares. She will then send out a status report to NSFNET-SITE-PEOPLE.

E. NSFNET NOC should post where Guy Almes's paper is located on nis.nsf.net.

F. Sue Hares will compile a list of available SGMP/SNMP tools.

(Notes by Mark Fedor of Nysernet. A big thank-you to Mark for a fine job... Sue Hares)

Slides Attached

Purpose

- Discuss how to find problems in the next hop network
- Create list of tools which can solve these problems
- Create a list of routing topology maps of regional networks

Agenda

1.) Introduction

2.) Routing Topology

Maps and Agreements

3.) Tools from Standards

SGMP/SNMP/CMOT/MIB

4.) Other Tools

5.) Methods

Methods

finding problems
in the next hop network
can use two methods:

- Verify not your end, and then call next hop network (NSFNET) who calls 2nd hop
- Debug via common tools whole path

Methods

- Complete list of contacts for campus, mid-level networks put on-line at NSFNET

  ongoing process

Tools

- Common SGMP sessions between NSFNET and regional networks

- SNMP once NSFNET supports SNMP

- Document on support of SGMP/SNMP in gateways and NSFNET

- Document on viewing tools for SGMP/SNMP

Other Tools

- Repository for tools at NSFNET IS machine (shareware status)

- List of Tools on NSFNET IS machine

MAPS

- On-Line Maps in simple postscript form on IS machine for campus, mid-level, and national networks
- Hard copy Maps collected too

Open SPF-based IGP

Mike Petry
University of Maryland

John Moy
Proteon

Open SPF-Based IGP WG Report
Reported by Mike Petry

17-19 Oct 1988
Ann Arbor, MI

A meeting was held to review and make comments on the draft specification of the OSPFIGP protocol written by John Moy. The most outstanding changes that were made were:

1) An encryption type field and fixing the size of the encryption field. I was decided that this field woul be used to validate the message using an out of band encryption method that was determined by the type field. This relieved the requirement to have a large or variable field set aside for things like large keys. Something like a cryptographic checksum of the packet was deemed more inline with the needs.

2) 32bit network mask. A full 32bit mask was allocated as a network mask. This allowed a more consistent determination of host routes vs. subnet routes vs network routes.

3) TOS - Some bit field adjustment were made so the TOS bits were easier to deal with. Including the precedence bits in this field is being considered.

4) The inclusion of a backup designated router, which was include in this draft, was explained.

A discussion of routing table representation was led by Van Jacobson. Van gave some insite on the merits of using Patricia Trees for compact routing table lookups.

Group Status

The OSPFIGP Requirments document remains completed. There has been little no changes to it since early spring 88.

The protocol specifications document has gone through what is hoped the last set of cosmetic changes. A few bits slid around, but no changes in philosophy were made.

The latest revision, in PostScript form, were made availabe via anonymous ftp from mank.proteon.com late in Dec. The packet formats should now be chiseled in stone.

There are three implementation of this protocol that are being work on.

1) MIT - for the MIT C gateway
2) Proteon - for the Proteon router
3) UMD - for 4BSD based systems

There is considerable collaboration between UMD and Proteon at this time. In fact, UMD has decicated a person to this task full time for the last five months. (Rob Coltun) The resultant UMD code will become public domain.

A common set of C header definition has been created that should aid in future implementations.

Here is a rough update of the UMD implementation:

> Code design and approximately a third of the OSPFIGP implementation
> has been completed. Currently finishing the SPF algorithm (which will include
> the new updates for the AS external and summary link updates) and the
> the receive packet routines. We expect to have a version by the April IETF
> that has been tested on a few local UMD machines and with the NeST
> simulation tool.

Because of timing problems, the OSPFIGP group has not planned to meet at the Jan IETF. Instead we are trying to get the NASA video conf system for the end of Feb.

Slides Attached

Open SPF-based IGP WG Report
Reported by Mike Petry

17-19 Oct 1988
Ann Arbor, MI

At the Oct IETF Meeting:

The draft specification of the OSPFIGP protocol written by John Moy was reviewed. The following modifications were made to his specificiation:

1) An encryption type field and fixing the size of the encryption field. I was decided that this field woul be used to validate the message using an out of band encryption method that was determined by the type field. This relieved the requirement to have a large or variable field set aside for things like large keys. Something like a cryptographic checksum of the packet was deemed more inline with the needs.

2) 32bit network mask. A full 32bit mask was allocated as a network mask. This allowed a more consistent determination of host routes vs subnet routes vs network routes.

3) TOS - Some bit field adjustment were made so the TOS bits were easier to deal with. Including the precedence bits in this field is being considered.

4) The inclusion of a backup designated router, which was include in this draft, was explained.

The protocol specifications document has gone through what is hoped the last set of cosmetic changes. A few bits slid around, but no changes in philosophy were made.

The latest revision, in PostScript form, were made availabe via anonymous ftp from mank.proteon.com late in Dec. The packet formats should now be chiseled in stone.

There are three implementations of this protocol being worked on.

1) MIT - for the MIT C gateway
2) Proteon - for the Proteon router
3) UMD - for 4BSD based systems

There is considerable collaboration between UMD and Proteon at this time. In fact, UMD has dedicated a person to this task full time for the last five months. The resultant UMD code will become public domain.

A common set of C header definition has been created that should aid in future implementations.

Here is a rough update of the UMD implementation:

Code design and approximately a third of the OSPFIGP implementation has been completed. Currently finishing the SPF algorithm (which will include the new updates for the AS external and summary link updates) and the the receive packet routines. We expect to have a version by the April IETF that has been tested on a few local UMD machines and with the NeST simulation tool.

OSPF IGP

- Subnet mask 32 bits
- Backup designated router
- TOS + Prec
- Multicast routes
- Multicast vs Broadcast
- Broadcast medium req. bidirectional
- Host routes supported by mask not bits
- Cryptographic type support
- ext info at untyped 32 bits

ospfigp@trantor.umd.edu
ospfigp-request

Open INOC

Jeff Case
UTK

[did not meet at Ann Arbor]

Open Systems Routing

Marianne Lepp
BBN

Open Routing Working Group

Charter: medium-term replacement of EGP

Requirements: IDEA 007

now consider policy crucial

Current work

Strawman Architecture

- Describe Policies
- Distribute Database Info
- Compute Routes
- Forward Packets

Policy

Links can carry information
Attribute list
Limited flow of info /data
nakossis 'Line up/down

Database

hierarchical
Will not flood everything

Compute Routes
Have not addressed in detail
Link state
Local break/ local fix

Forward Packets

source routing
Route Set-up

OSI Interoperability

[did not meet at Ann Arbor]

PDN Routing

C-H Rokitansky
FernUni Hagen

AGENDA OF THE 2nd PDN ROUTING WORKING GROUP MEETING AT IETF, Oct 18th '88

roki@DHAFEU52.BITNET AND roki@A.ISI.EDU

- INTRODUCTION
- BACKGROUND INFORMATION
- DETAILED TECHNICAL DISCUSSION AND SPECIFICATION OF SHORT TERM GOALS (VAN-GATEWAYS, EGP3, ROUTING METRICS, ETC.)
- DISCUSSION OF MEDIUM TERM GOALS
- LONG TERM GOALS
- STATUS REPORT ON BBN-VAN-GATEWAY (BUTTERFLY REPLACEMENT, EGP, ETC.) by (Mike Brescia, BBN)
- PROPOSAL FOR A MAPPING BETWEEN DNICs AND INTERNET/PDN-CLUSTER NETS - Discussion (Carl-H. Rokitansky, FernUni Hagen)
- DISCUSSION OF HIERARCHICAL GATEWAY ALGORITHM FOR ROUTING AND NETWORK REACHABILITY INFORMATION EXCHANGE BETWEEN LEVEL 1 TO 4 GATEWAYS (Carl-H. Rokitansky)
- PROPOSAL OF AN X.121 ADDRESS RESOLUTION PROTOCOL (Mike Brescia)
- PROPOSAL OF AN CALL SETUP & REVERSE CHARGING
- PROTOCOL (CRCP) FOR X.25 CONNECTIONS (Carl-H. Rokitansky)
- TECHNICAL DISCUSSION
- PDN ROUTING PERFORMANCE TESTS
- ASSIGNMENT OF ACTION ITEMS
- MISCELLANEOUS

PDN ROUTING WORKING GROUP STATUS

PUBLICATION

- INTERNET CLUSTER ADDRESSING PAPER IN PROCEEDINGS OF 9th INTERNATIONAL COMPUTER COMMUNICATIONS CONFERENCE (ICCC '88)

PROPOSALS:

- X.121 ADDRESS RESOLUTION PROTOCOL
- DNIC ↔ INTERNET PDN-CLUSTER NETWORK MAPPING
- HIERARCHICAL GATEWAT ALGORITHMS FOR PDN - CLUSTER
- CALL SETUP & CHARGING DETERMINATION PROTOCOL (SCDP)

DNIC Assignment

Zone	Area	DNICs assigned
2	EUROPE	49
3	N. AMERICA	39
4	ASIA	40
5	PACIFIC	20
6	AFRICA	8
7	S. AMERICA	44
		200

& INTERNET/PDN-CLUSTER Mapping

reserve	Hemisphere	PDN-Cluster Networks	Net MSB.
127	EAST	190.001 - 190.127	0000-0111
127	WEST	190.128 - 190.254	1000-1111
95	EAST	191.001 - 191.095	0000-0101
32	EAST	191.096 - 191.127	0110-0111
31	WEST	191.224 - 191.254	1110-1111
96	WEST	191.128 - 191.223	1000-1101
508			

PDN-CLUSTER Mask:

1111 1110 0000 0000 0....0 <254.0.0.0>

PDN-West Mask:

1111 1110 1000 0000 0....0 <254.128.0.0>

WEST			EAST		
6	7	3	2	4	5

HIERARCHICAL GATEWAY ALGORITHMS FOR PDN-CLUSTER

Level	Gateway
4	ZONE - GW
3	COUNTRY - GW
2	DATA NETWORK - GW
1	VAN - GW

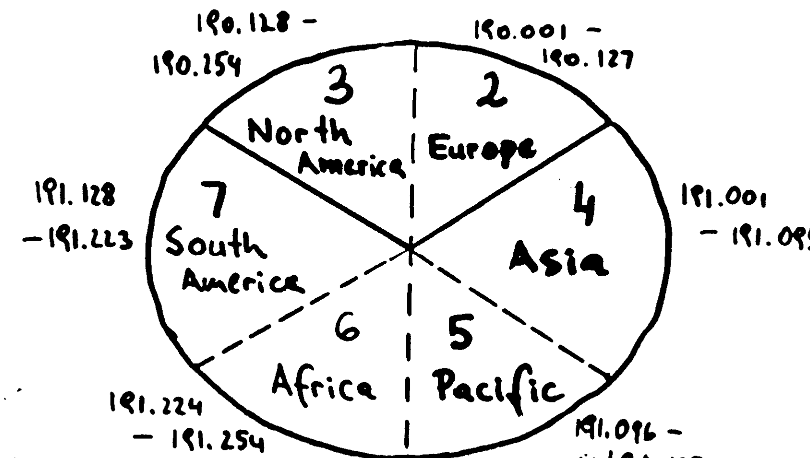

PDN-Cluster

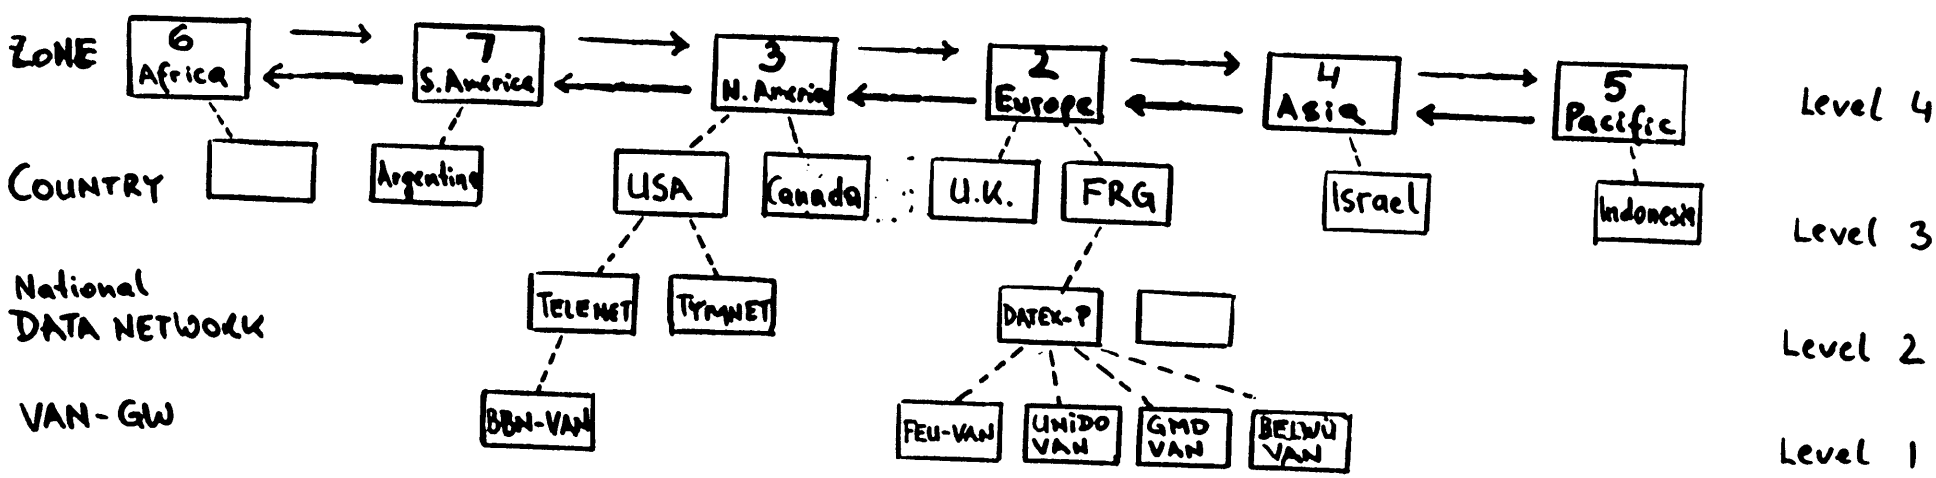

CALL SETUP & REVERSE CHARGING PROTOCOL (CRCP)

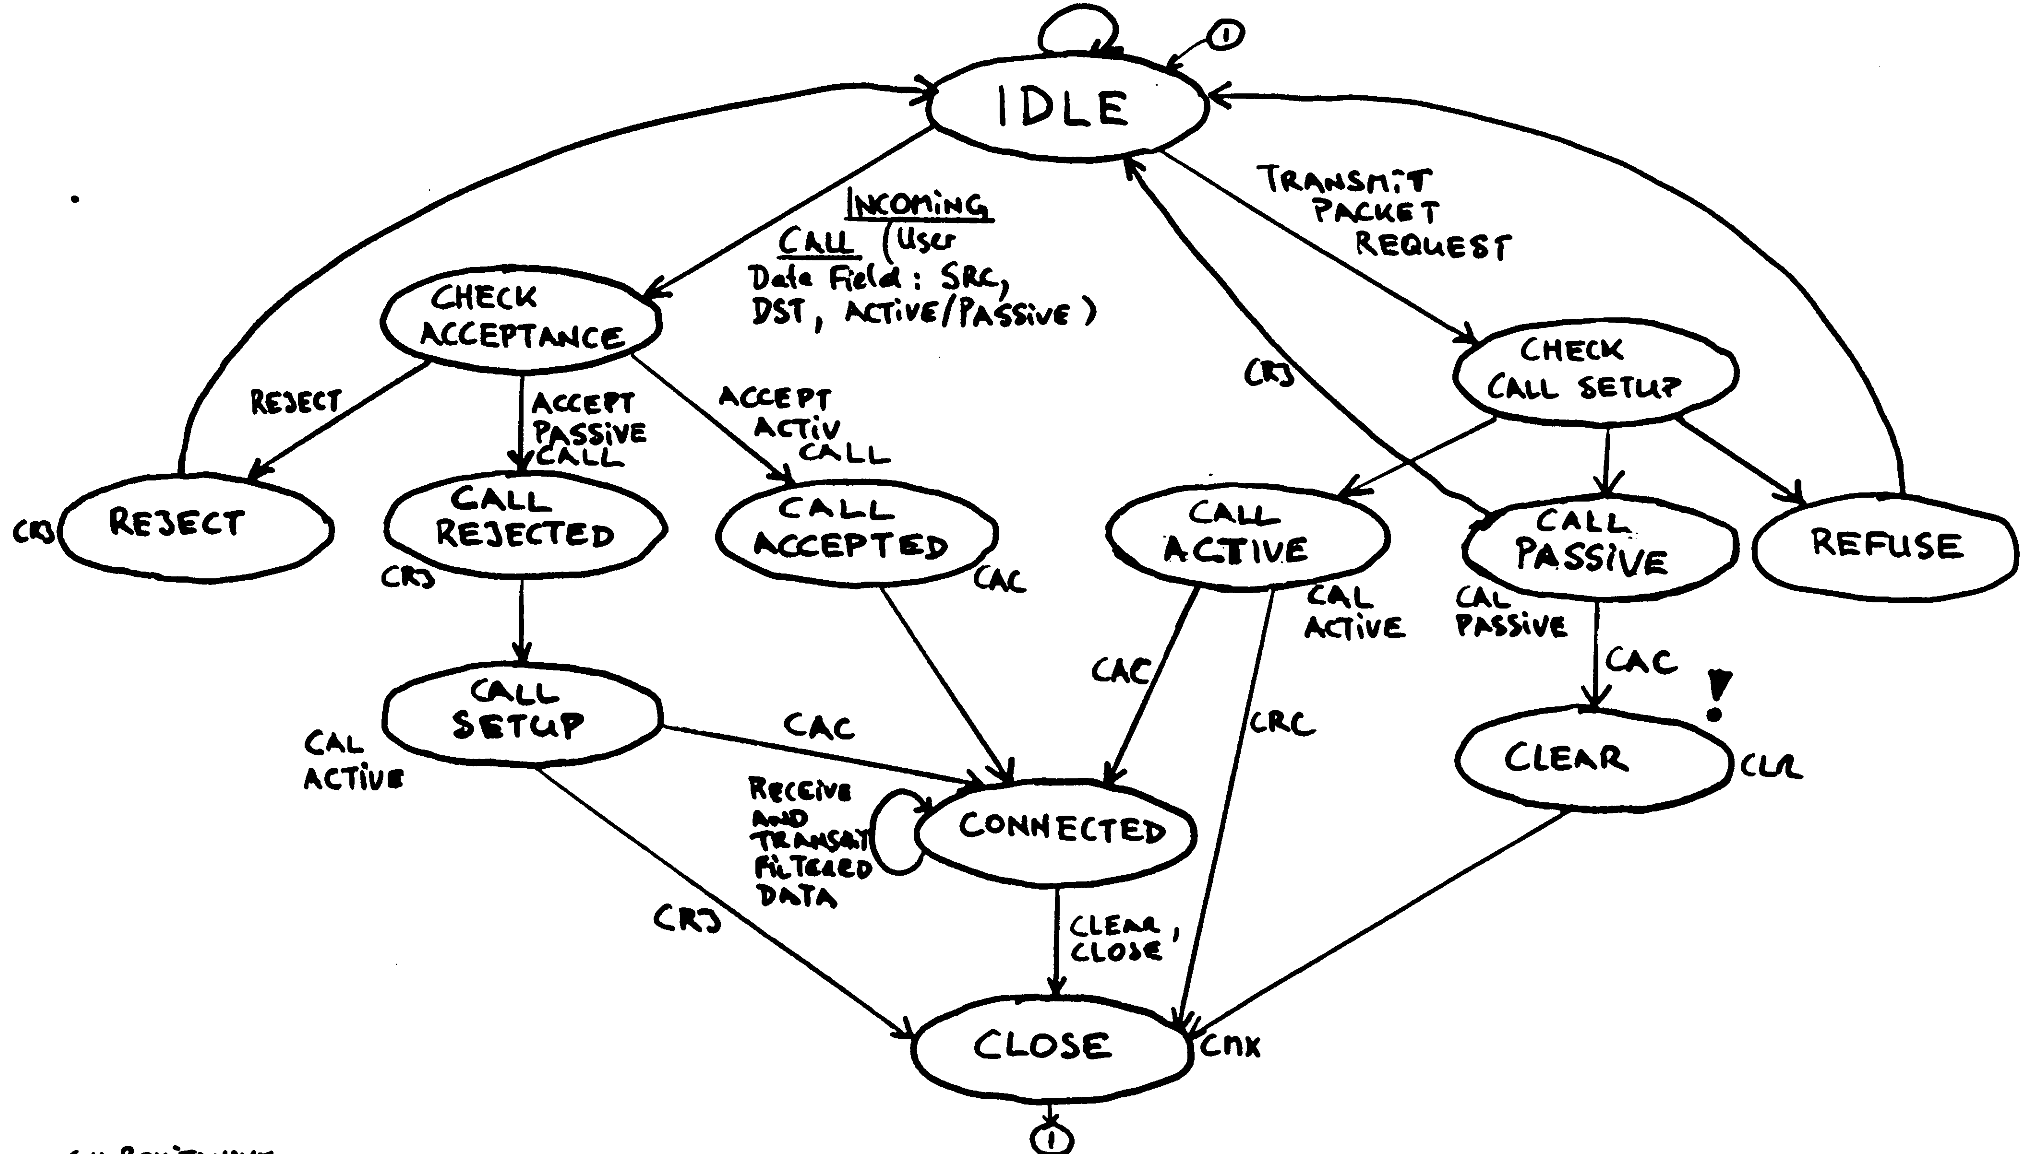

C.H. ROKITANSKY

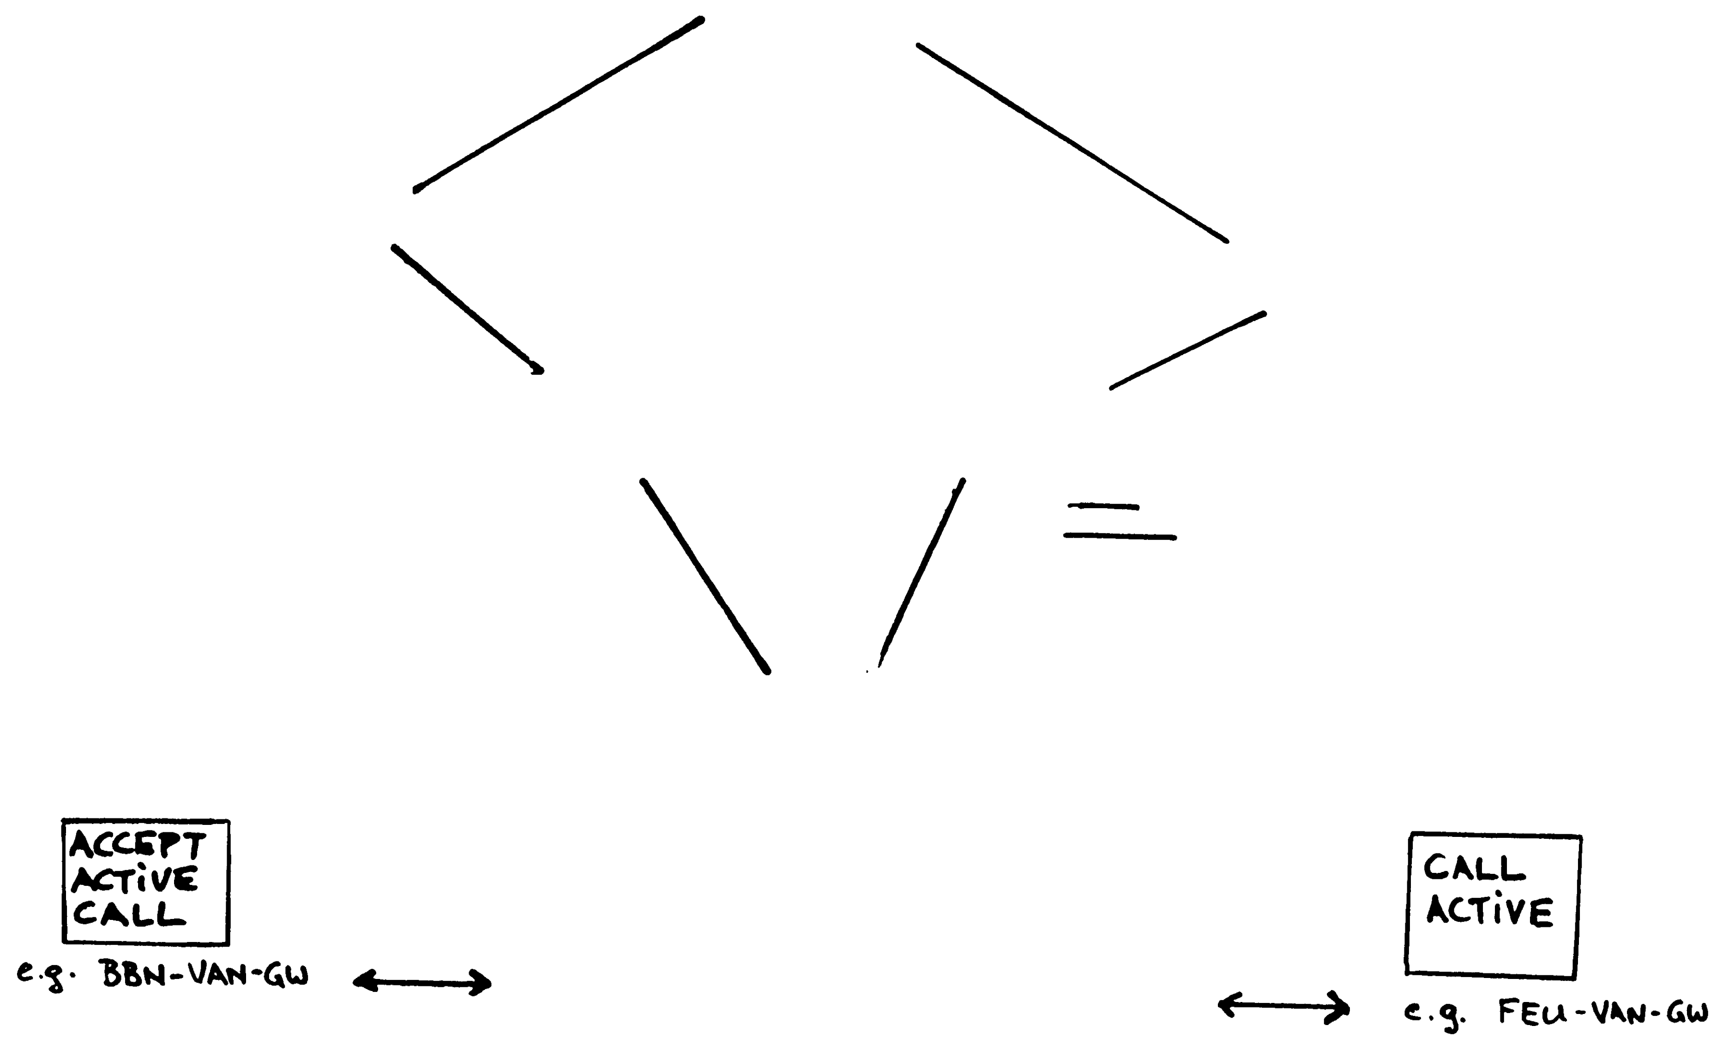

ACCEPT ACTIVE CALL
e.g. BBN-VAN-GW
CALL ACTIVE
e.g. FEU-VAN-GW

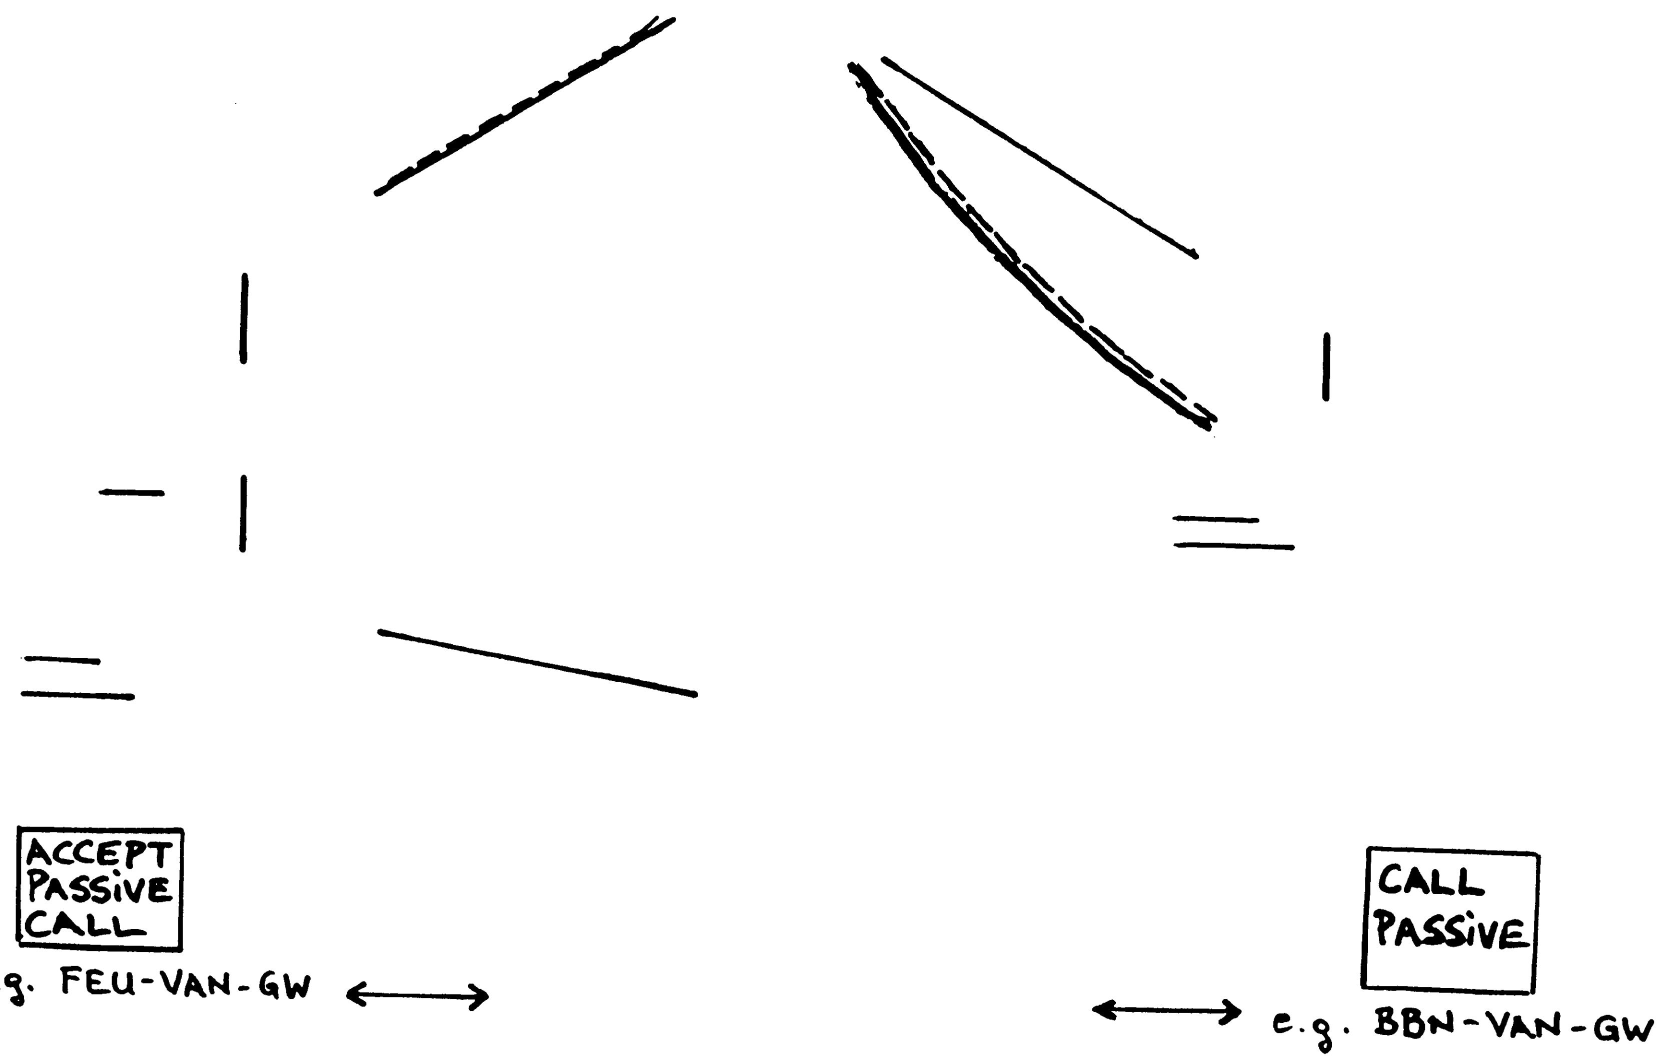
ACCEPT PASSIVE CALL
e.g. FEU-VAN-GW
CALL PASSIVE
e.g. BBN-VAN-GW

CALL SETUP & REVERSE CHARGING PROTOCOL (CRCP)

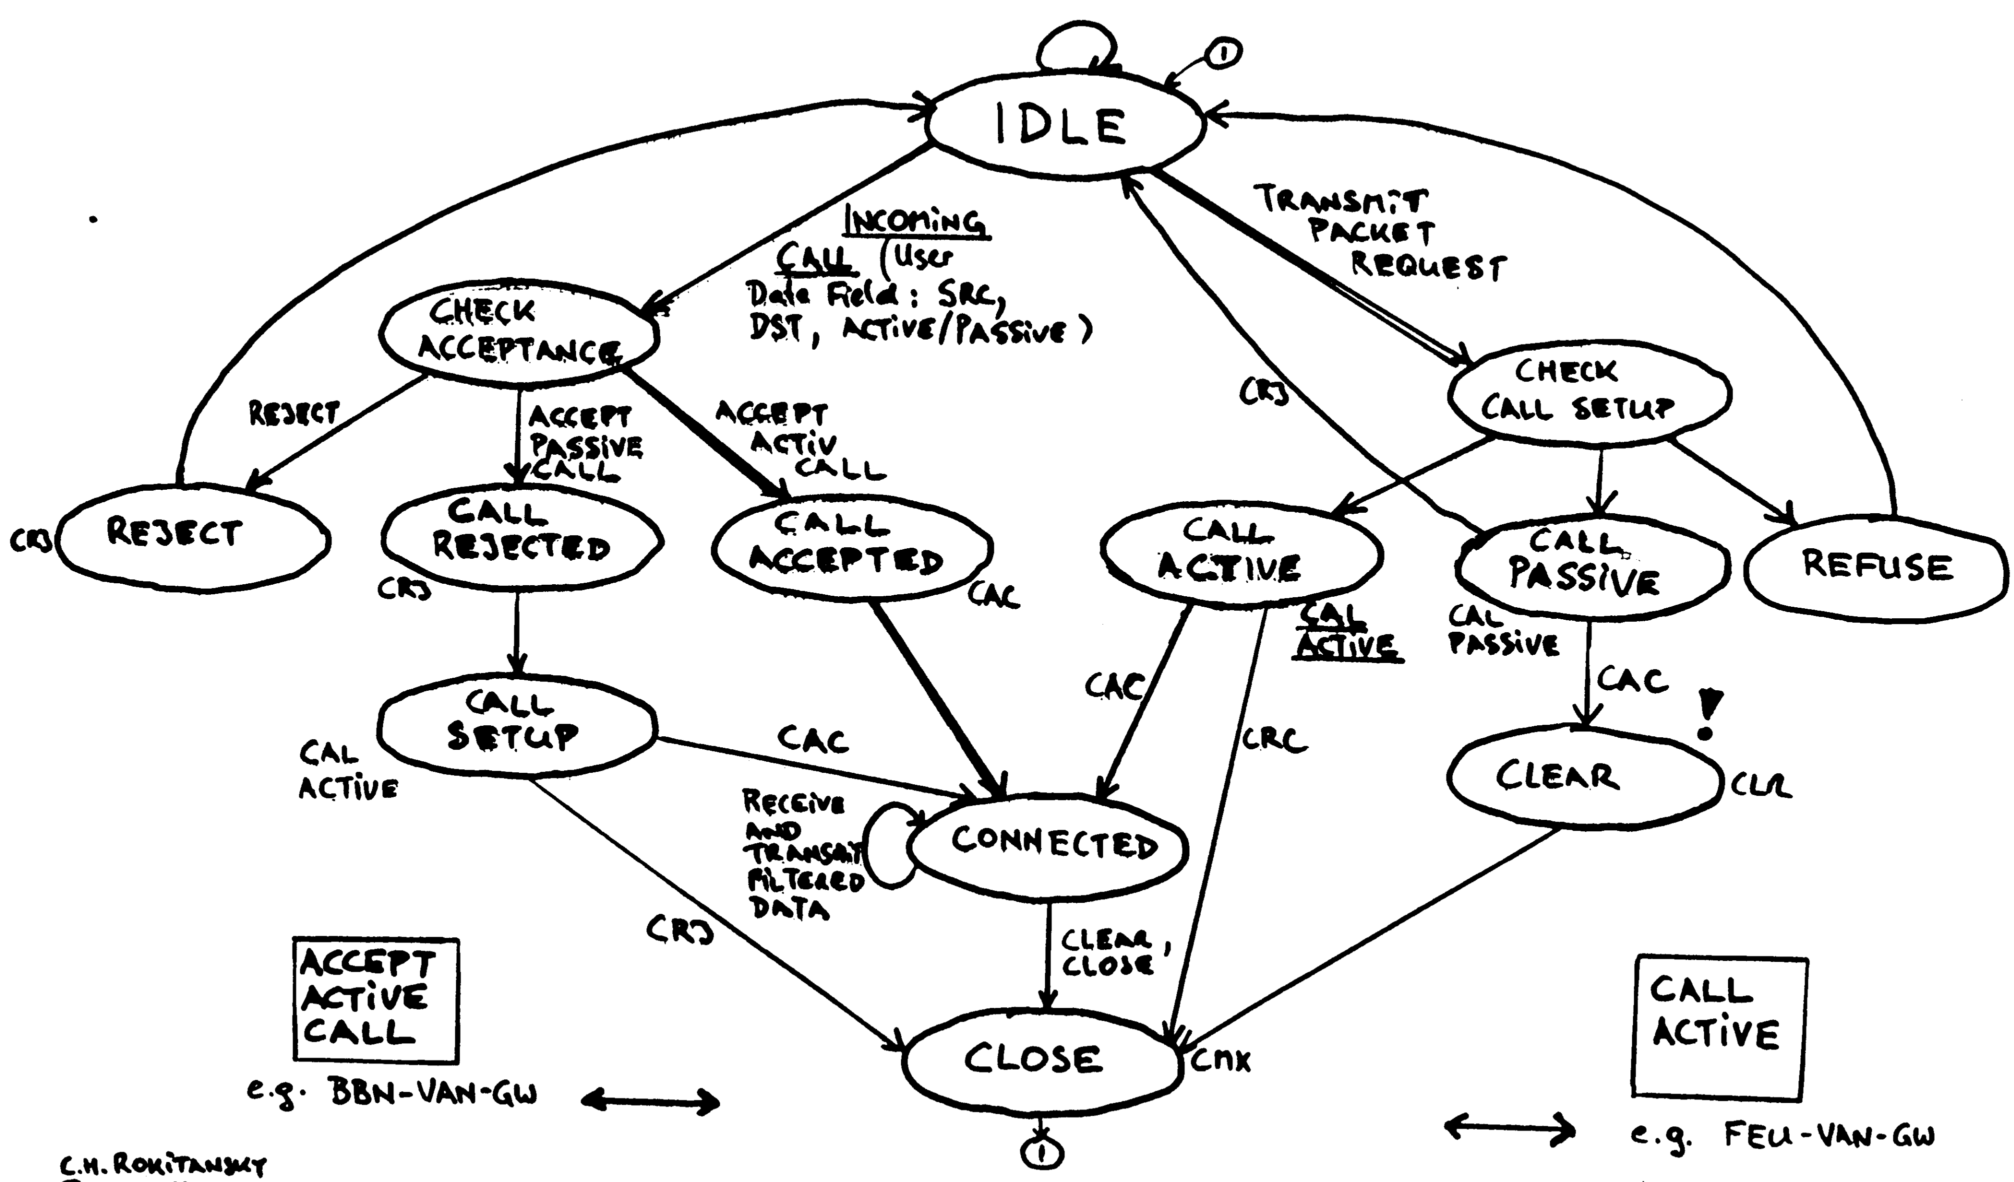

C.H. Rokitansky
FernUni Hagen

SIMPLIFIED STATE DIAGRAM

CALL SETUP & REVERSE CHARGING PROTOCOL (CRCP)

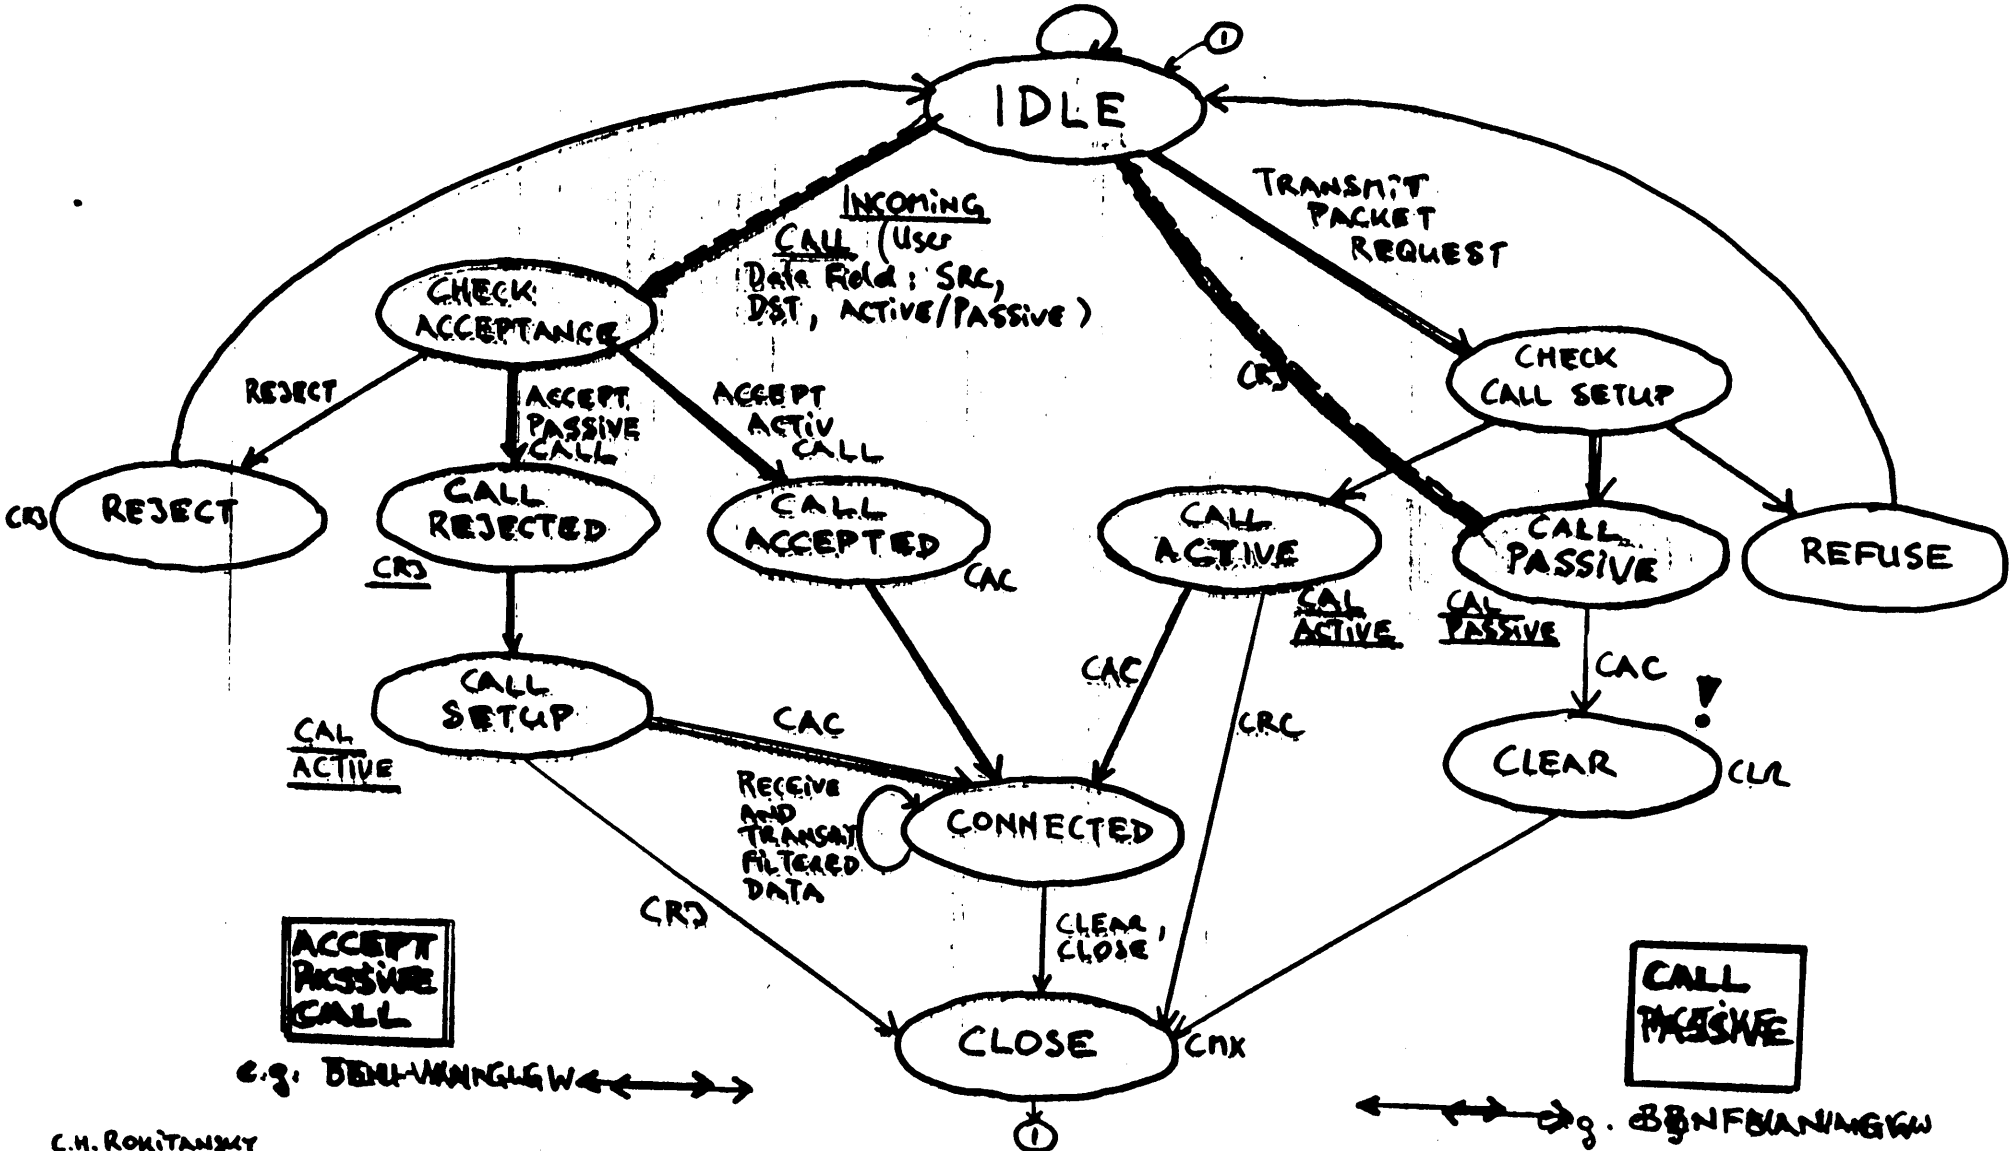

C.H. Rokitansky
Fernuni Hagen

PDN ROUTING WORKING GROUP PAPERS

- X.121 ADDRESS RESOLUTION PROTOCOL IDEA (Jan '89)
- MAPPING BETWEEN DNICs AND PDN-CLUSTER NETWORKS - IDEA (Jan '89)
- CALL SETUP & CHARGING DETERMINATION PROTOCOL (SCDP) - IDEA (Jan/Feb '89)
- HIERARCHICAL GATEWAY ALGORITHMS AND NETWORK REACHABILITY INFORMATION EXCHANGE FOR PDN-CLUSTER - IDEA (Jan/Feb '89 ?)

PDN - CLUSTER FUNCTIONALITY TESTS

BETWEEN USA AND EUROPE, EXPECTED TO START DEC '88

Performance and CC

Allison Mankin
MITRE

IETF Performance/Congestion Control WG — Oct 1988
Reported by A. Mankin & G. Reichlen (MITRE) — Ann Arbor

Attendees: Roger Beeman (Boeing), Art Berggreen (ACC), Scott Brim (Cornell), Steve Casner (USC/ISI), Bilal Chinoy (MERIT), Mike Collins (LLNL/ESnet), Gene Hastings (PSC), Van Jacobson (LBL), Paul Love (SDSC), Ken Lowe (Univ. of Washington), Allison Mankin (MITRE), Matt Mathis (PSC), Philippe Park (BBN), Paul Parker (CMU), Guru Parulkar (Washington Univ.), K.K. Ramakrishnan (DEC), Gladys Reichlen (MITRE), Robert Reschly (BRL), Bruce Schofield (DCEC), Geof Stone (Network Systems Corp.), Paul Ticknor (NASA/NAS), Claudio Topolcic (BBN), Steve Wolff (NSF), Jessica Yu (MERIT).

The Performance Working Group met on Tuesday morning. Currently this group has a paper in progress which addresses Internet performance for TCP and gateways. During an off-line discussion with Bob Braden, from the Host Requirements WG, it was recommended that the Performance WG produce a separate document (as a Draft RFC) specifying TCP congestion control methods, in particular Slow-start. This document would be an adjunct to the Host Requirements RFC. The WG agreed that this was a good approach. Therefore a draft of this new paper will be put together, and distributed to the WG via email for comment, before the January meeting.

Claudio Topolcic from the ST/Connection-Oriented IP WG briefed us on their group's direction. They are working on two documents: a modification to the current ST specification and a connection-oriented Internet protocol requirements document. In the requirements document they will be defining performance guarantees needed from the network for successful use of applications such as video-conferencing, in addition to the common ones (FTP etc.). Our two groups will cooperate.

In response to the Host Requirements RFC reiterating the definition of IP TTL as a time, not a hop count, the Performance WG discussed several issues: TTL as a time does not give enough range (that is, usual values of TTL, such as 30 seconds, could be quite unsafe with the current range of Internet transit times, if most gateways suddenly treated the TTL as a time. However, TCP not wanting to wrap sequence numbers while a segment is in the network requires the bounded lifetime implied by TTL as a time. A suggestion for an alternative that met some favor was to have gateway IP bound packet lifetime on the queue. The TTL maximum times the queue stay bound would have to be within the TCP Maximum Segment Lifetime.

Van Jacobson talked to us about his recent activities. The report in these minutes will be sketchy, since we hope to hear about these projects in detail in future IETF plenary presentations:

Gateway congestion control experiments: reserving bandwidth for packet video through gateways, in conjunction with an ARMA congestion predictor. The set-up is done with options and a special TOS is used during the lifetime of the video. Good success so far with reserving 250Kb/sec bandwidth for each video flow and still running TCP connections fairly.

Analytical modelling: he has a tractable model of transport and gateway with one gateway. It's not tractable with a larger Internet. The results so far support TCP window flow control (versus rate-based control) -- another reference on this was Aurel Lazar (Columbia Univ. Telecommunications).

Stimulated by a request from Phill Gross (visiting), we had a discussion of source quench, with the following as a brief summary of the various arguments about its effect:

1. SQ is not good, but not really bad except it takes bandwidth at a time when you really don't want to do that.

2. SQ has an underlying model-that congestion problems are being by a small number of hosts. This underlying model is not the same as the reality of transit gateways: confluences from a range of hosts at varying distances from the gateway.

3. Why not assume source quench is an early indication of packet drop? Because gateways are not held to using source quench to mean this. A connection can safely interpret source quench as meaning there is some congestion, but not what degree and not whether caused by itself; in LAN experiments, the SQ went consistently to the wrong host,, i.e. the host with the smallest windows and the most random sends. Slow-start therefore does a restart in response to SQ, but does not change the ssthresh, the size of the window above which further opening is done slowly.

4. SQ is essentially broken - even if you can guarantee you quench the right source, it is still not the right mechanism. Slow-start has a conservative handling of SQ, but it still has to be considered what harm it may do: synchronization effects and effects on control loops by taking action at a rate less than the round-trip time are two possibilities.

The group discussed whether it would make sense to produce a short ''kill SQ'' RFC--consensus was yes. Approach: a collection of existing data to support the con arguments of the discussion. Van and Allison have experimental evidence to contribute to the paper.

The remaining hour of the meeting was taken up with a discussion of gateway performance and the extent to which gateways, as they exist now, can support performance guarantees. The unsolved problem of how gateways can accurately signal bandwidth changes to TCP (and similarly behaved traffic) is a big obstacle.

Next Meeting:

At the next IETF meeting, the group will discuss a draft of the TCP document (coming). The agenda will also include new information gathering for the second document on gateway performance.

Pt-Pt Protocol

Drew Perkins
CMU

Point-to-Point Protocol WP Report — **17-19 Oct 1988**
Reported by Drew Perkins — **Ann Arbor, MI**

The PPP WG met at the IETF meeting in Ann Arbor, Michigan during the morning of Tuesday, October 18, 1988 and also briefly during the afternoon of Wednesday, October 19, 1988.

Attendees on Tuesday were Drew Perkins (CMU), Ron Broersma (NOSC), Michael Petry (UMD), Bob Gilligan (SUN), Mark Lottor (SRI-NIC), Terry Bradley (Wellfleet), Becca Nitzan (NMFECC), Milo Medin (NASA), John Moy (Proteon), Russ Hobby (UC Davis), Philip Prindeville (McGill), RIck Boivie (IBM), Jessica Yu (Merit), Cal Thixton(NeXT) and Phil Karn (Bellcore). Russ Hobby took minutes.

The first item on the agenda was a discussion of the "Point-to-Point Protocol Requirements" document. Three changes were suggested:

1. A section should be added discussing hardware vs. software requirements.

2. A definition of "fragment" should be added.

3. The section on "Sequencing" should mention that it should not be ruled out. There may be times when it is desired, such as for other protocols and possibly when the reliability bit is set in IP datagrams.

The rest of the meeting was spent discussing proposals for the protocol.

1. Discussion on liveness/up-down capabilities

    a. The protocol should make sure to use hardware status (carrier detect, etc) when possible.

    b. The use of the liveness protocol should be negotiated before line is brought up.

    c. Liveness protocol should compare frame counts sent to frame count received at other end for line line quality. Negotiate line quality (error rate) at which to take down and bring up the line.

2. Discussion on error detection/correction

    a. The protocol should send CRCs in ALL cases, other end does not necessarily need to check them if it does not want error detection (i.e. you want to pass through data even if it is know to be bad, may be the case in voice or video).

b. There was much discussion concerning error correction. Conclusion: error correction not used by default but may be enabled when it is necessary. Suggest using LAPB.

3. Discussion on async protocol

   a. We discussed two framing protocols for async links: the framing protocol used by Rick Adams' SLIP, and the Proposed Draft International Standard ISO 3309 Revised (E). The DIS ISO 3309 defines how to do HDLC framing for "Start/stop transmission", aka async links. Since backward-compatibility with SLIP is not one of our goals (SLIP provides so little that it doesn't make sense), we decided that we may as well abandon SLIP and standardize on ISO 3309 HDLC. This should work out well since HDLC is more likely to be supported in the future by modem and IC manufactures. It also clears up the confusion about back compatibility quite nicely (it won't work).

4. Discussion on sync protocol

   a. No one questioned that the "obvious" thing to do is use HDLC framing, with addresses 1 and 3 and UI in the control field. This is very nice because full LAPB can be run in parallel simultaneously if desired.

5. Discussion on packet format

   a. We decided to use our own numbering system for the type field with standard values independent of MAC layer (async/sync/etc). Conflicting goals of even packet boundaries for high-speed links and high link efficiency for low-speed links led to agreement on an ISO'ish protocol (reminiscent of HDLC addresses). All protocol types values can be represented in 15 bits or less. For the foreseeable future it is likely that there will be very few protocols, probably less than 32. Therefore, the type field will normally be a single octet for async links, but will be extended to two octets when necessary (protocol type exceeds 1 octet). On sync links, the two octet representation will be used at all times. This is accomplished by using the MSB of the first octet transmitted/received to indicate a one/two octet type field. When the MSB is one, the field is 16 bits and remaining 15 bits are the type value. When the first bit is zero, the field is 8 bits and the remaining 7 bits are the type value.

The initial values will be as follows:

zero - reserved
1 - link control
2 - IP
3 - ISO
4 - XNS
5 - MAC bridge
6 - DECNET
32767 - reserved (all 1's)

6. Discussion on Link Control Packets
   a. Line Reset
   b. Line going down
   c. Others

7. Discussion on Option Negotiation Packets
   a. One item will be negotiated per packet, but packet may have multiple parts (ie: a list of addresses)

   b. Option Packet Fields
      Option type - 16 bits
      Length - 16 bits
      Data

   c. Items considered for negotiation

ITEM	DEFAULT	
MTU	576	
Compression	Off	
Liveness (Up/Down)	Off	
LAPB (error correction)	Off	
Addresses	None	(use ARP format?)
Authentication		
Encryption	Off	
Character mapping	Off	

   d. General strategy for bringing up line.

      Start dumb, learn smarts. Start with basic communications and negotiate other capabilities. This ensures compatibility at start.

8. Discussion on problem of loopback detection and Master/Slave establishment. Protocol: Send random number (64 bit) challenge. Get response. Compare. If response is the same number, may be loopback, try new random number. If get back same number after N tries, assume loopback. Possible sources for random number:

   MAC address
   Machine serial number
   Non-volatile memory configuration
   Low bits of clock

Result of comparison determines master and slave. Higher number is master. For HDLC, higher number is DCE (address 1), lower number is DTE (address 3).

ST and CO-IP

Claudio Topolcic
BBN

ST and connection oriented internet protocol

- Two parallel tracks
  - ST specification
  - Longer term connection oriented internet protocol issues
- ST specification
  - Identified a number of issues
  - Did not resolve any
  - Will meet by multimedia conferences
  - Will exchange mail
  - Intend to have a good draft by next IETF
  - Implement after
- Connection oriented internet protocol
  - Progress understanding what we mean
  - Plan
    - Identify requirements
    - Specify options
    - Possibly test some options on ST
    - Incorporate results from Inter Domain Routing
    - Write a specification
  - Requirements document
    - Driven by applications
    - Req's of protocol
    - Req's of networks

    - Have an outline
    - Have writing assignments
    - Will continue by E-mail
    - Plan to have a draft by next IETF

TELNET Linemode

Dave Borman
Cray

TELNET Line Mode

- Reviewed & modified draft #2 of IDEA 16
- Draft #3, this meetings changes, will be very close to RFC
- Get new draft out for 1-2 months review, then submit as an RFC

User Services

Karen L. Bowers
NRI

[newly formed; will meet 18-20 Jan 89]

6. NETWORK STATUS BRIEFINGS AND TECHNICAL PRESENTATIONS

NSFnet Report

Part I

Hans-Werner Braun
Merit, Inc.
(University of Michigan)

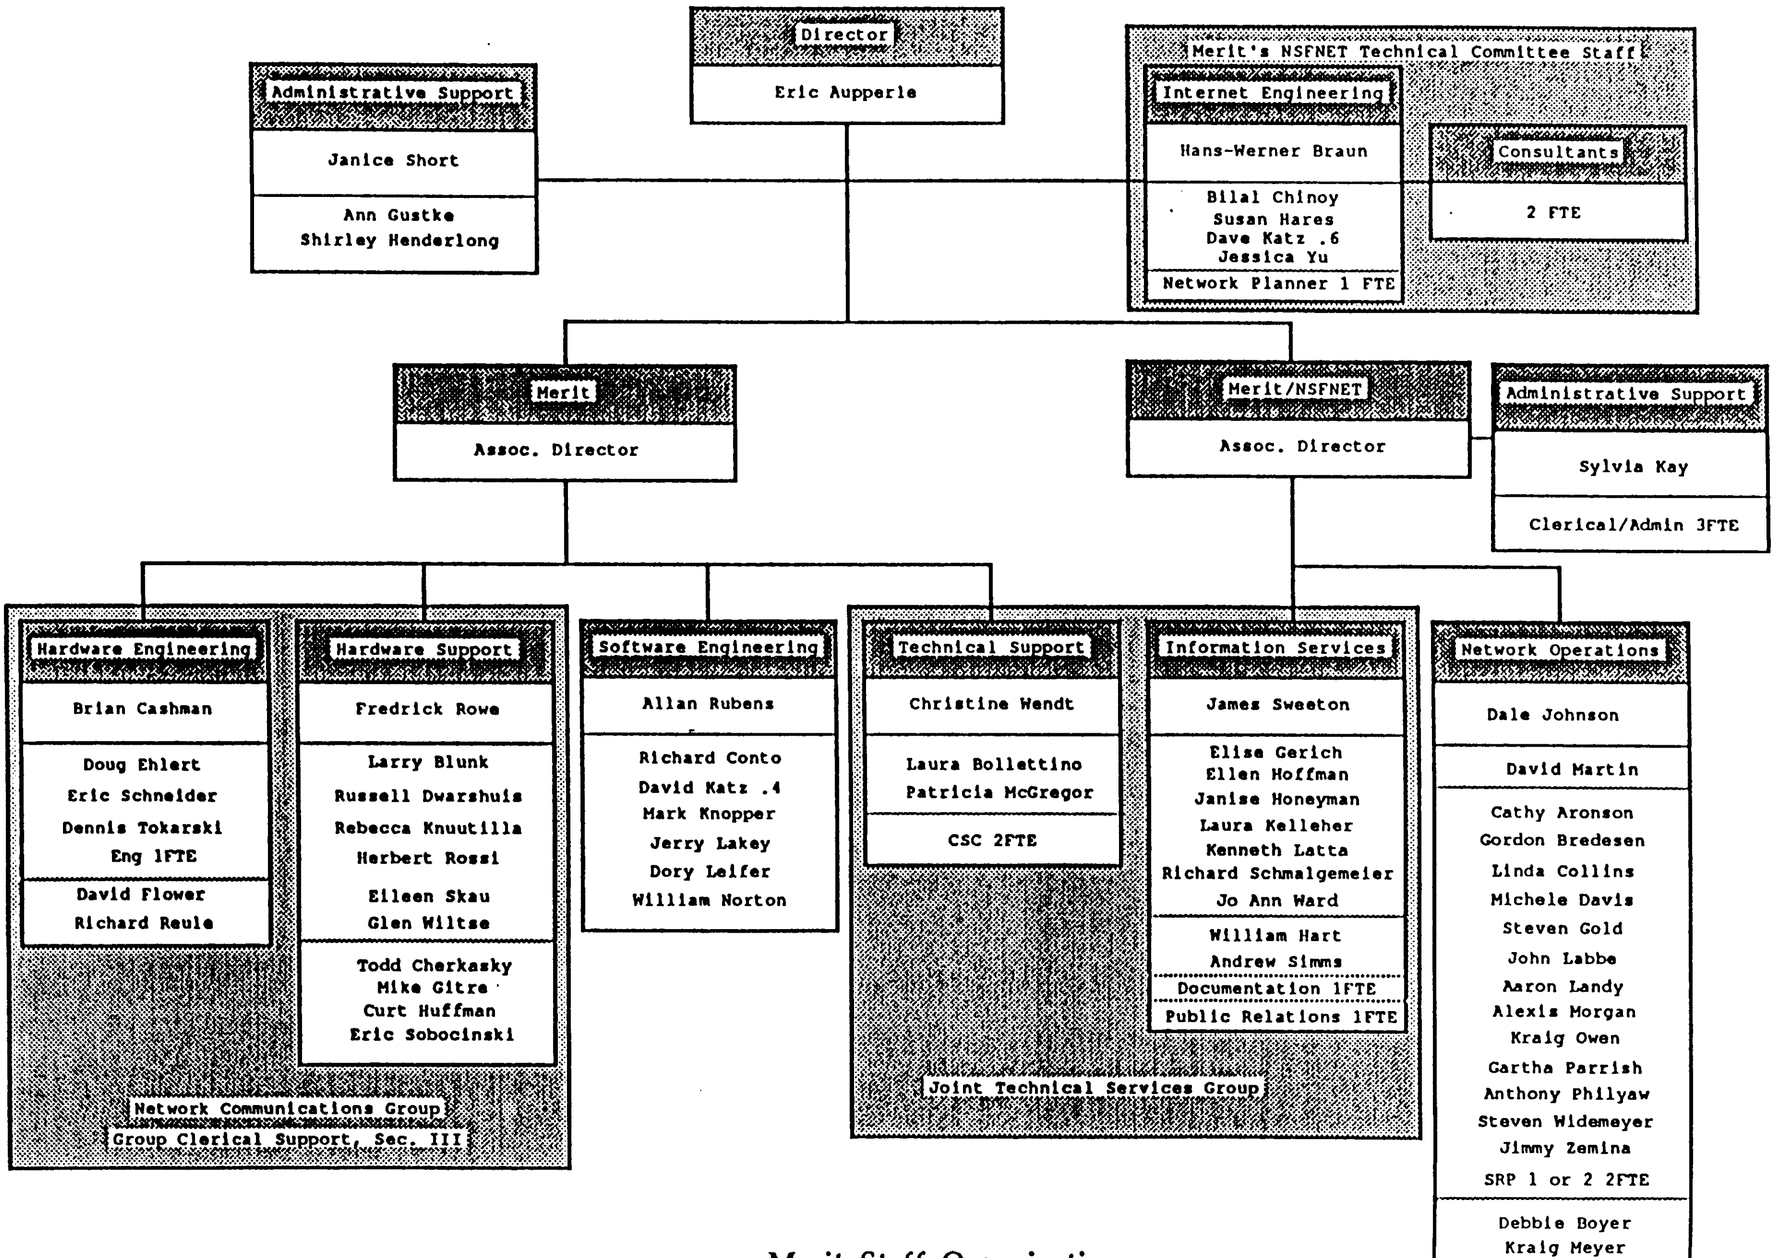

Merit Staff Organization
including NSFNET

8-19-88

NSFNET
T-1 DATA NETWORK

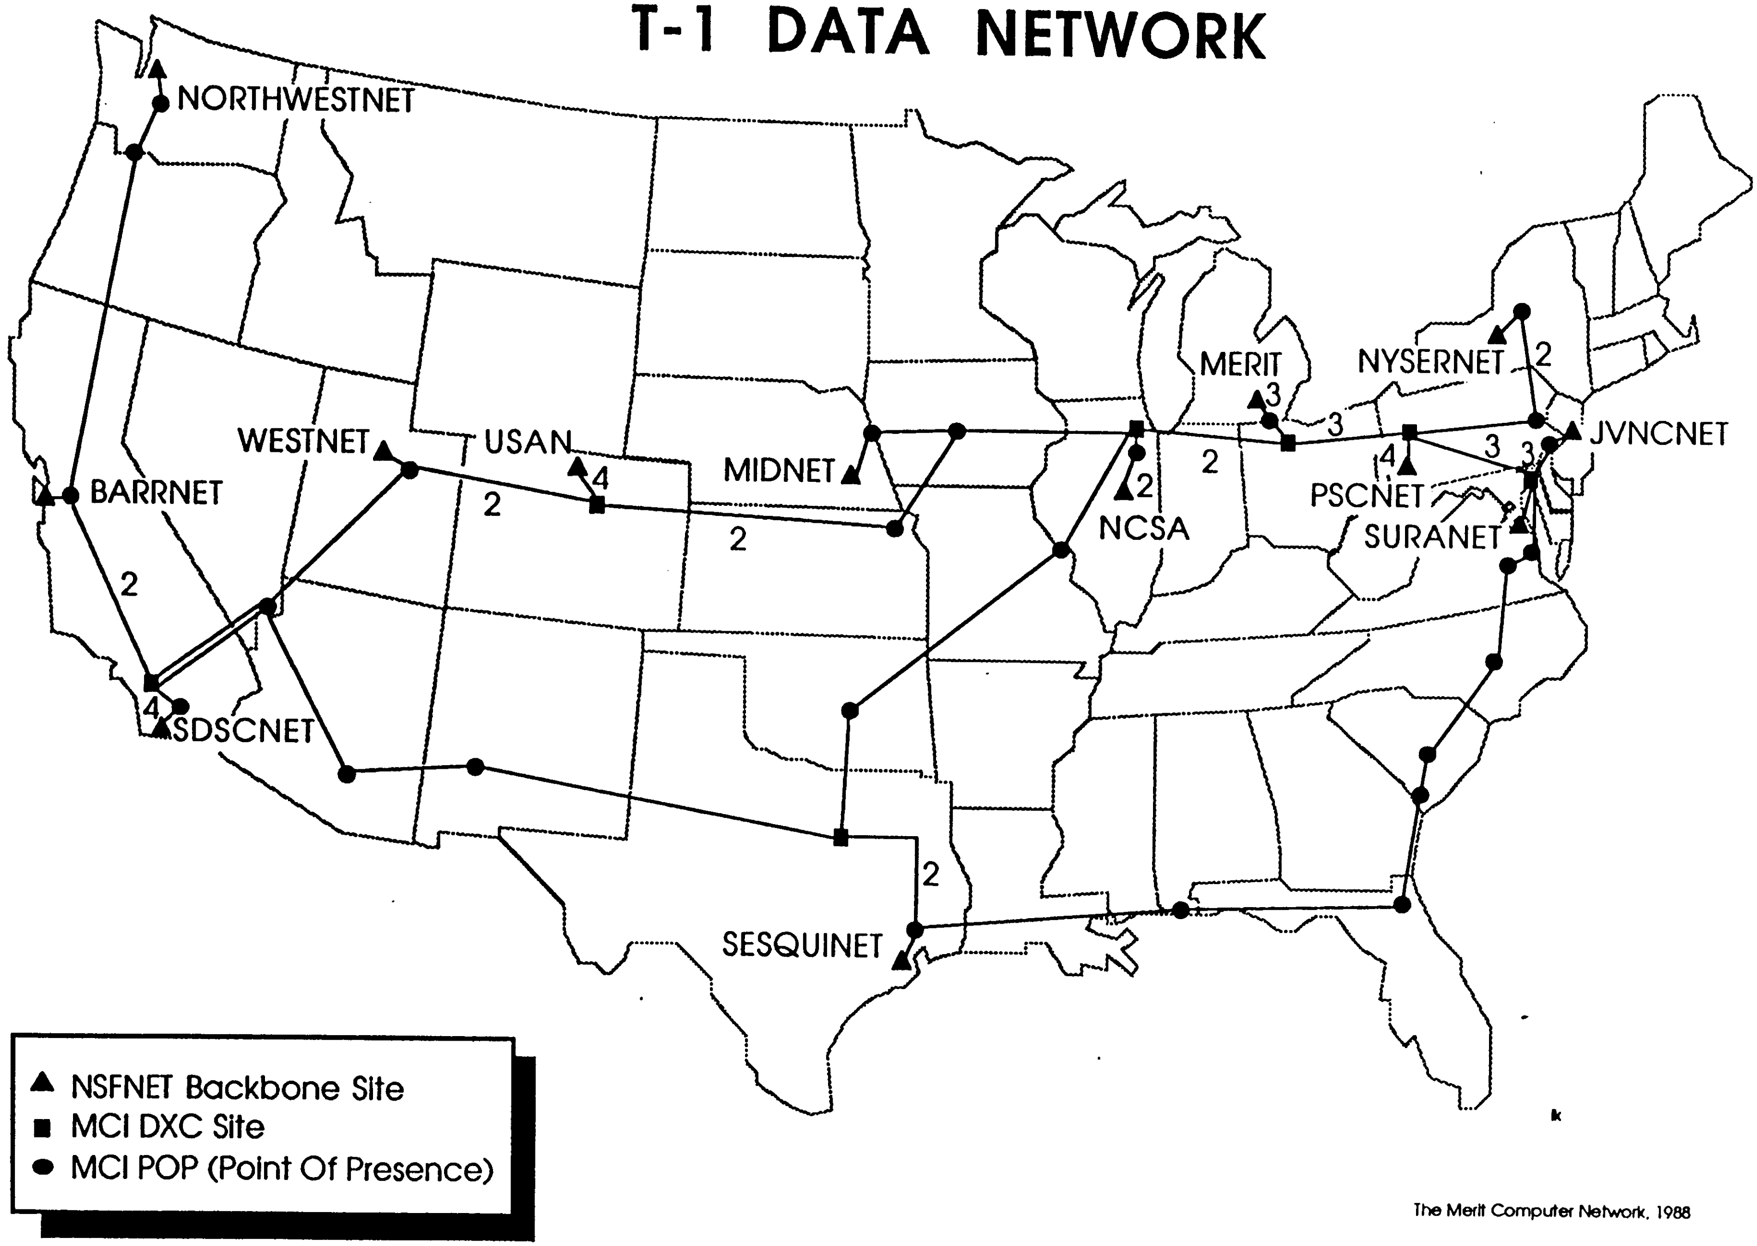

The Merit Computer Network, 1988

Physical Initial NSFNET Topology

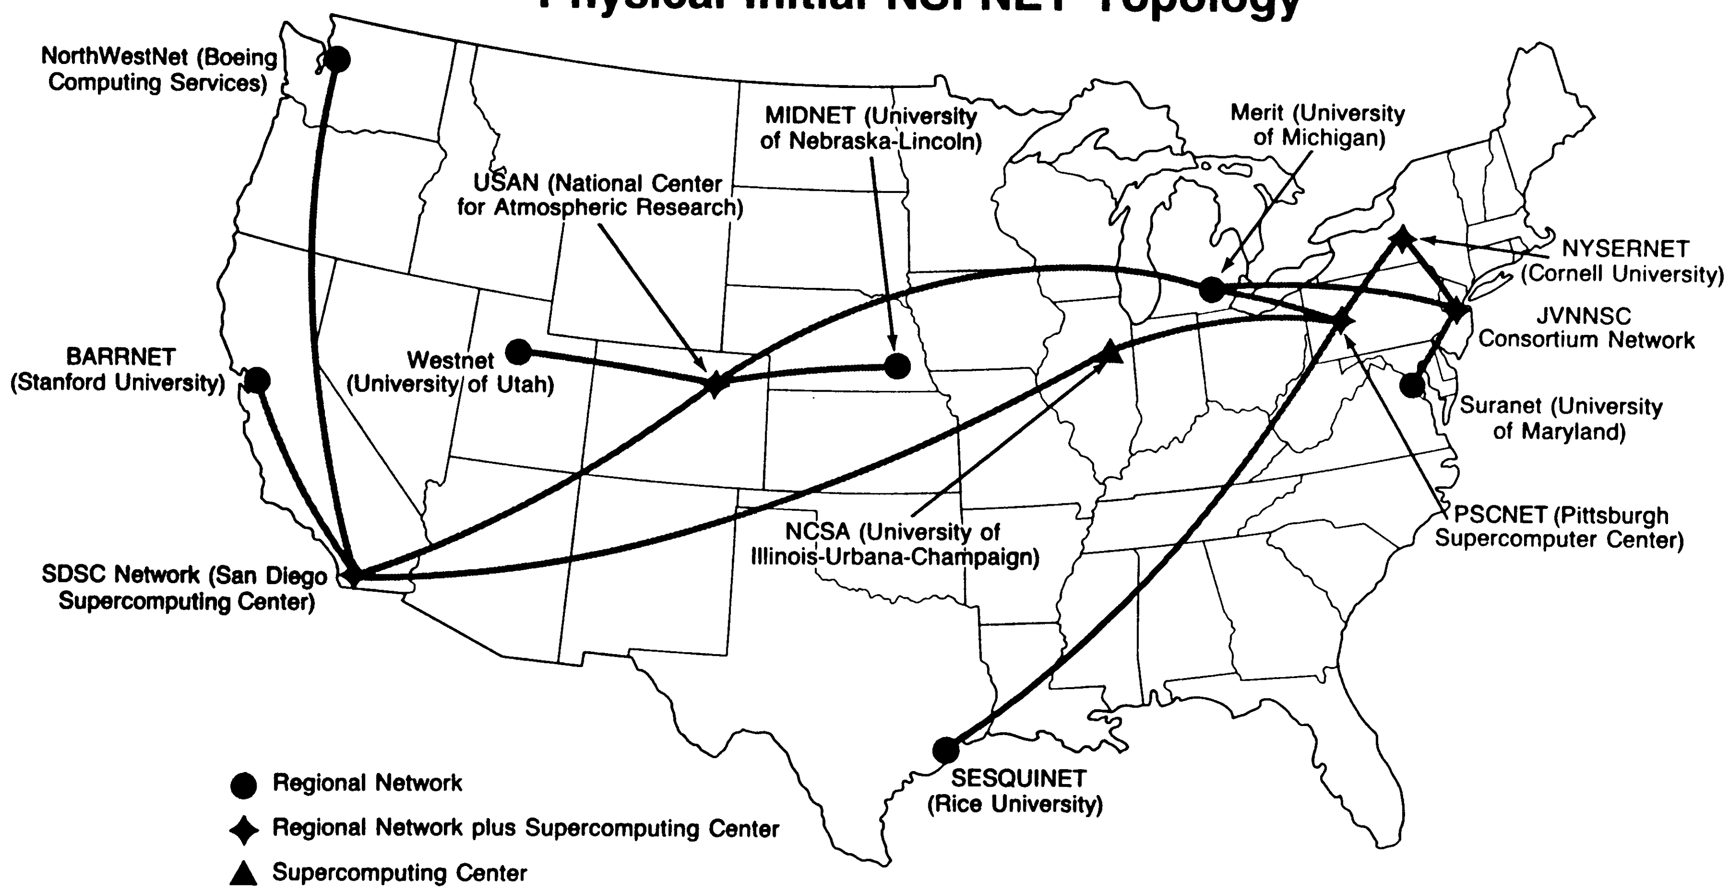

Center for Cartographic Research and Spatial Analysis, Michigan State University, 2/88

NSFNET Backbone
Logical Topology

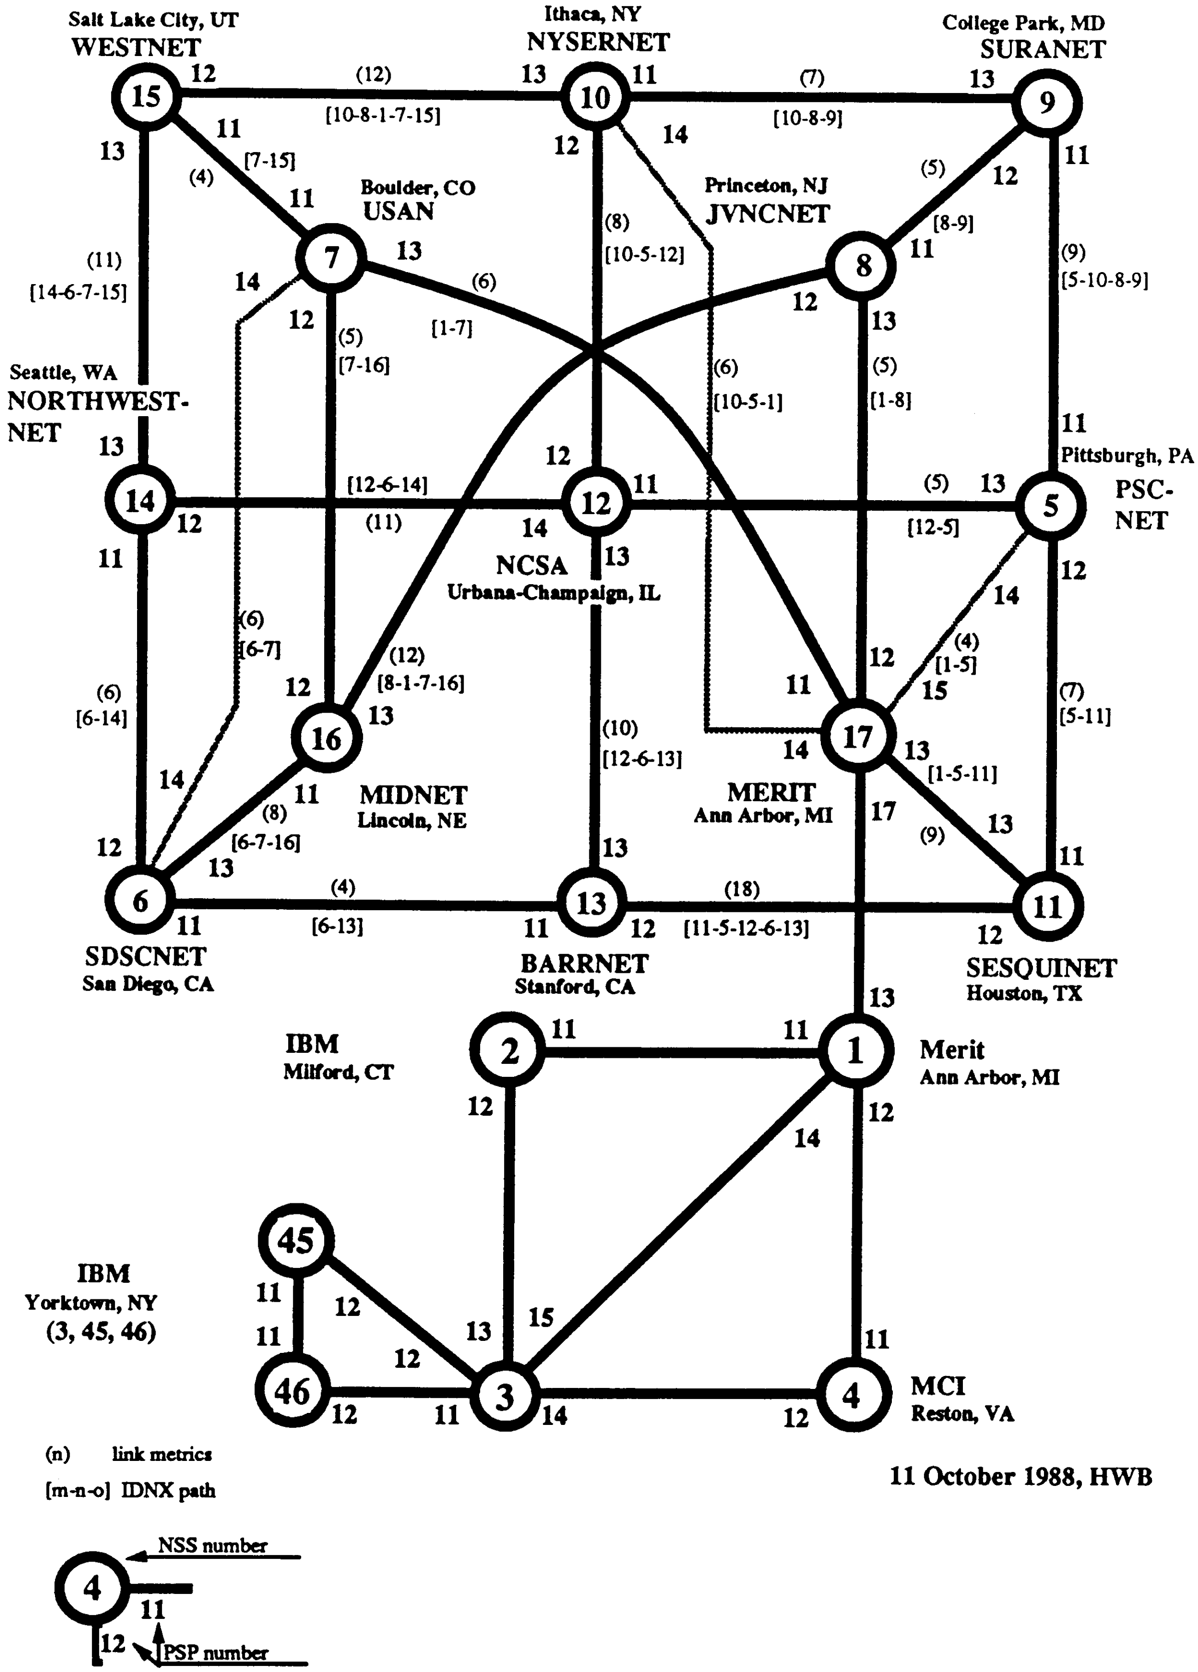

11 October 1988, HWB

Logical Initial NSFNET Topology

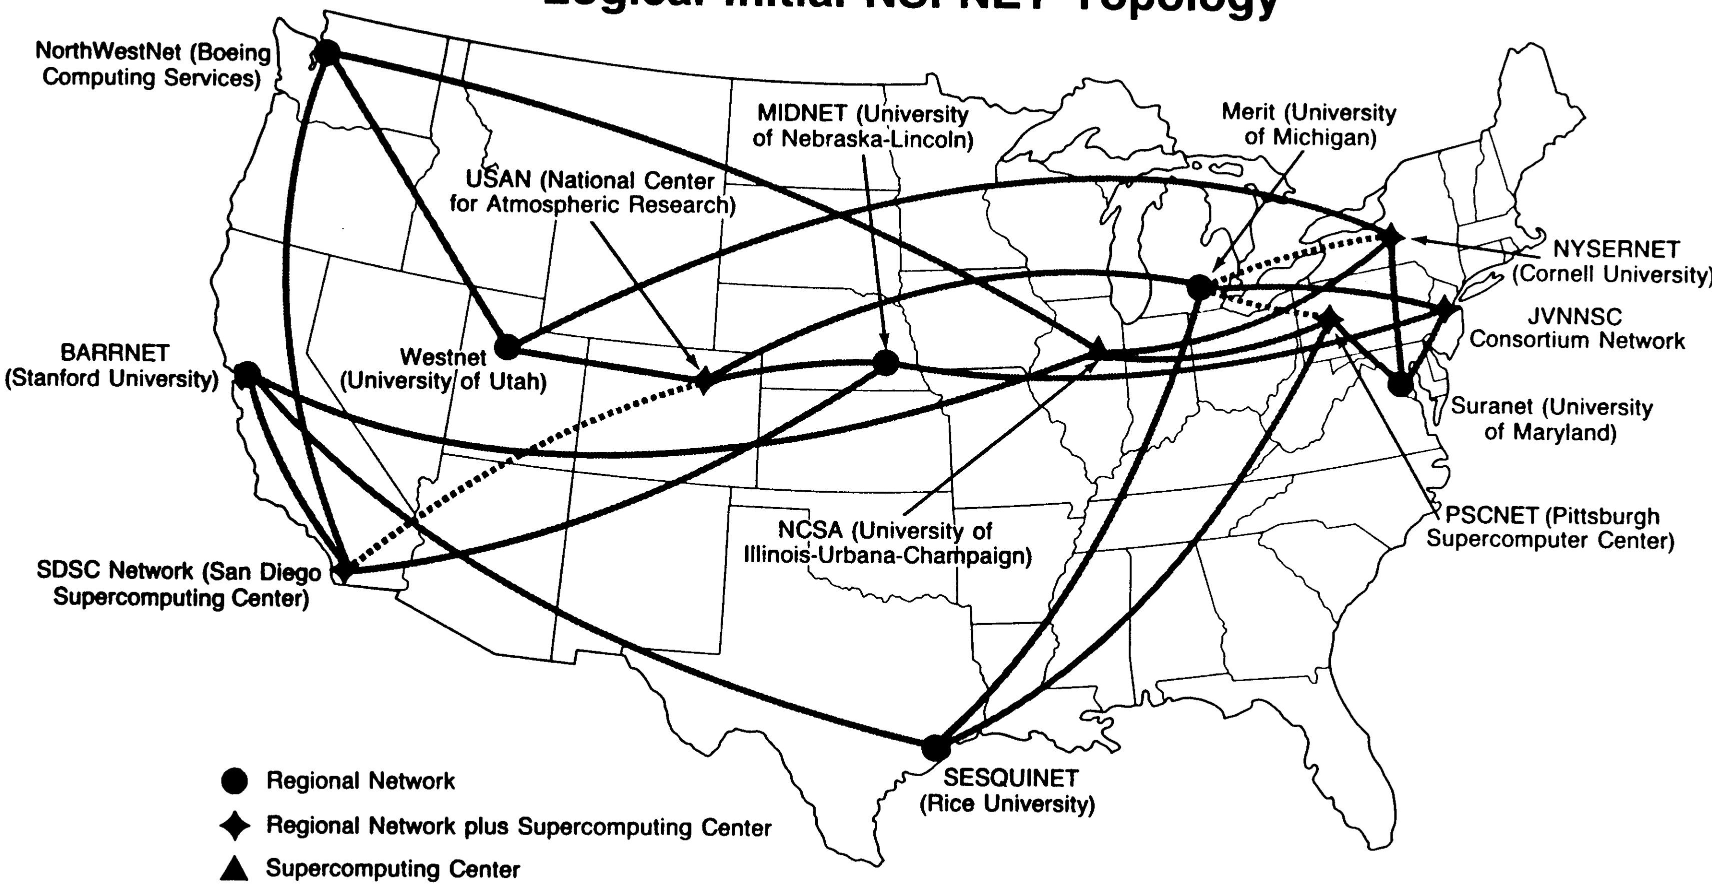

Center for Cartographic Research and Spatial Analysis, Michigan State University, 2/88

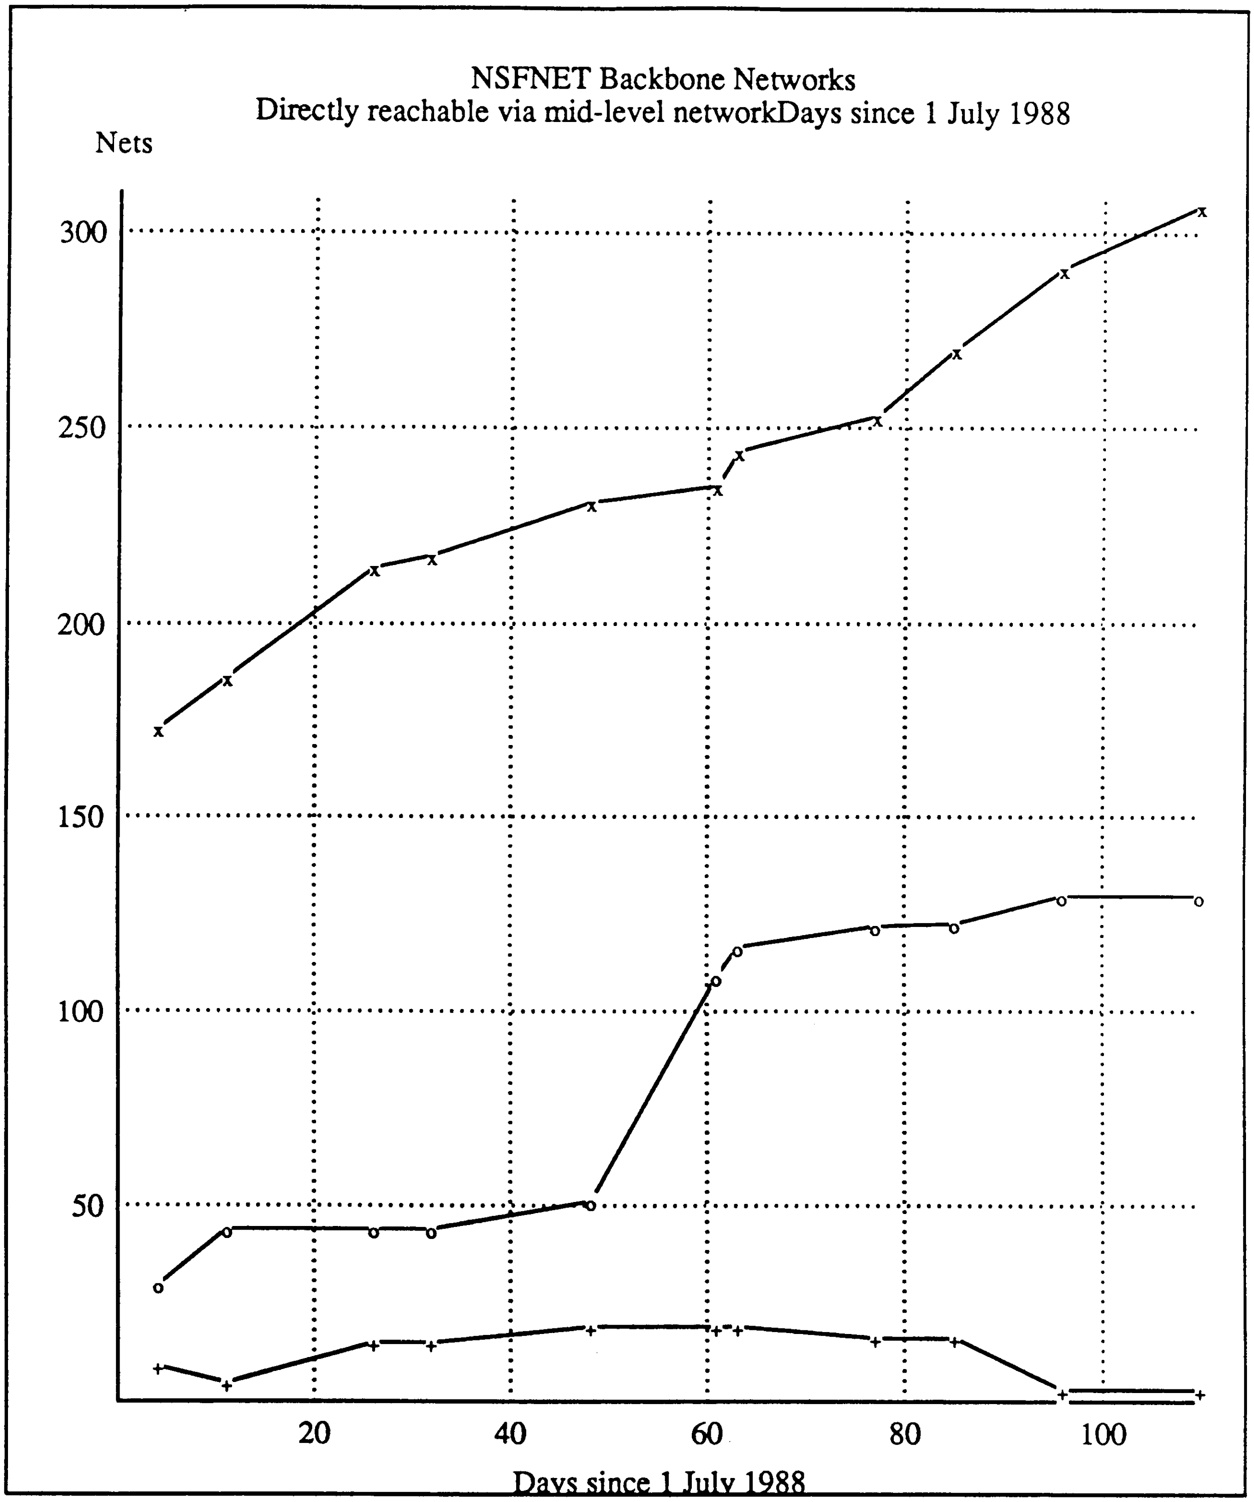
NSFNET Backbone Networks
Directly reachable via mid-level networkDays since 1 July 1988
Nets
300
250
200
150
100
50
20
40
60
80
100
Days since 1 July 1988

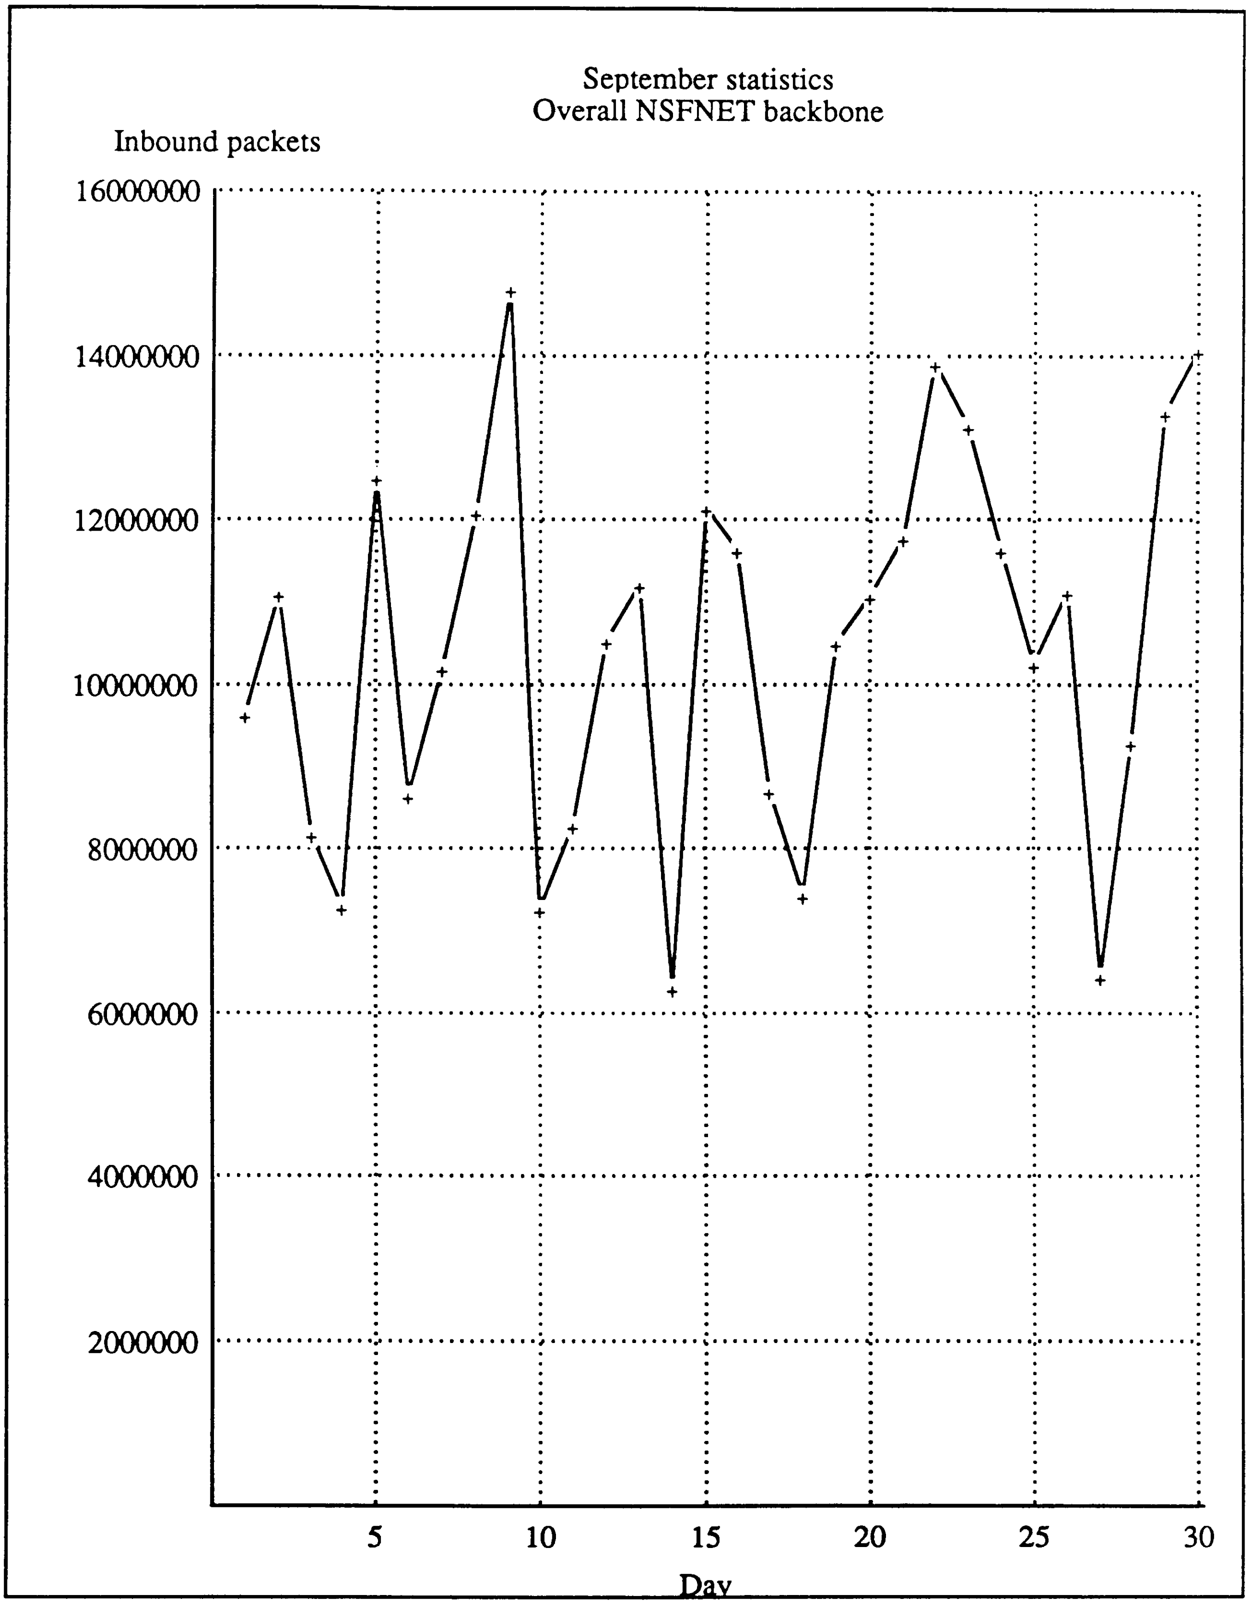
September statistics
Overall NSFNET backbone
Inbound packets
16000000
14000000
12000000
10000000
8000000
6000000
4000000
2000000
5
10
15
20
25
30
Day

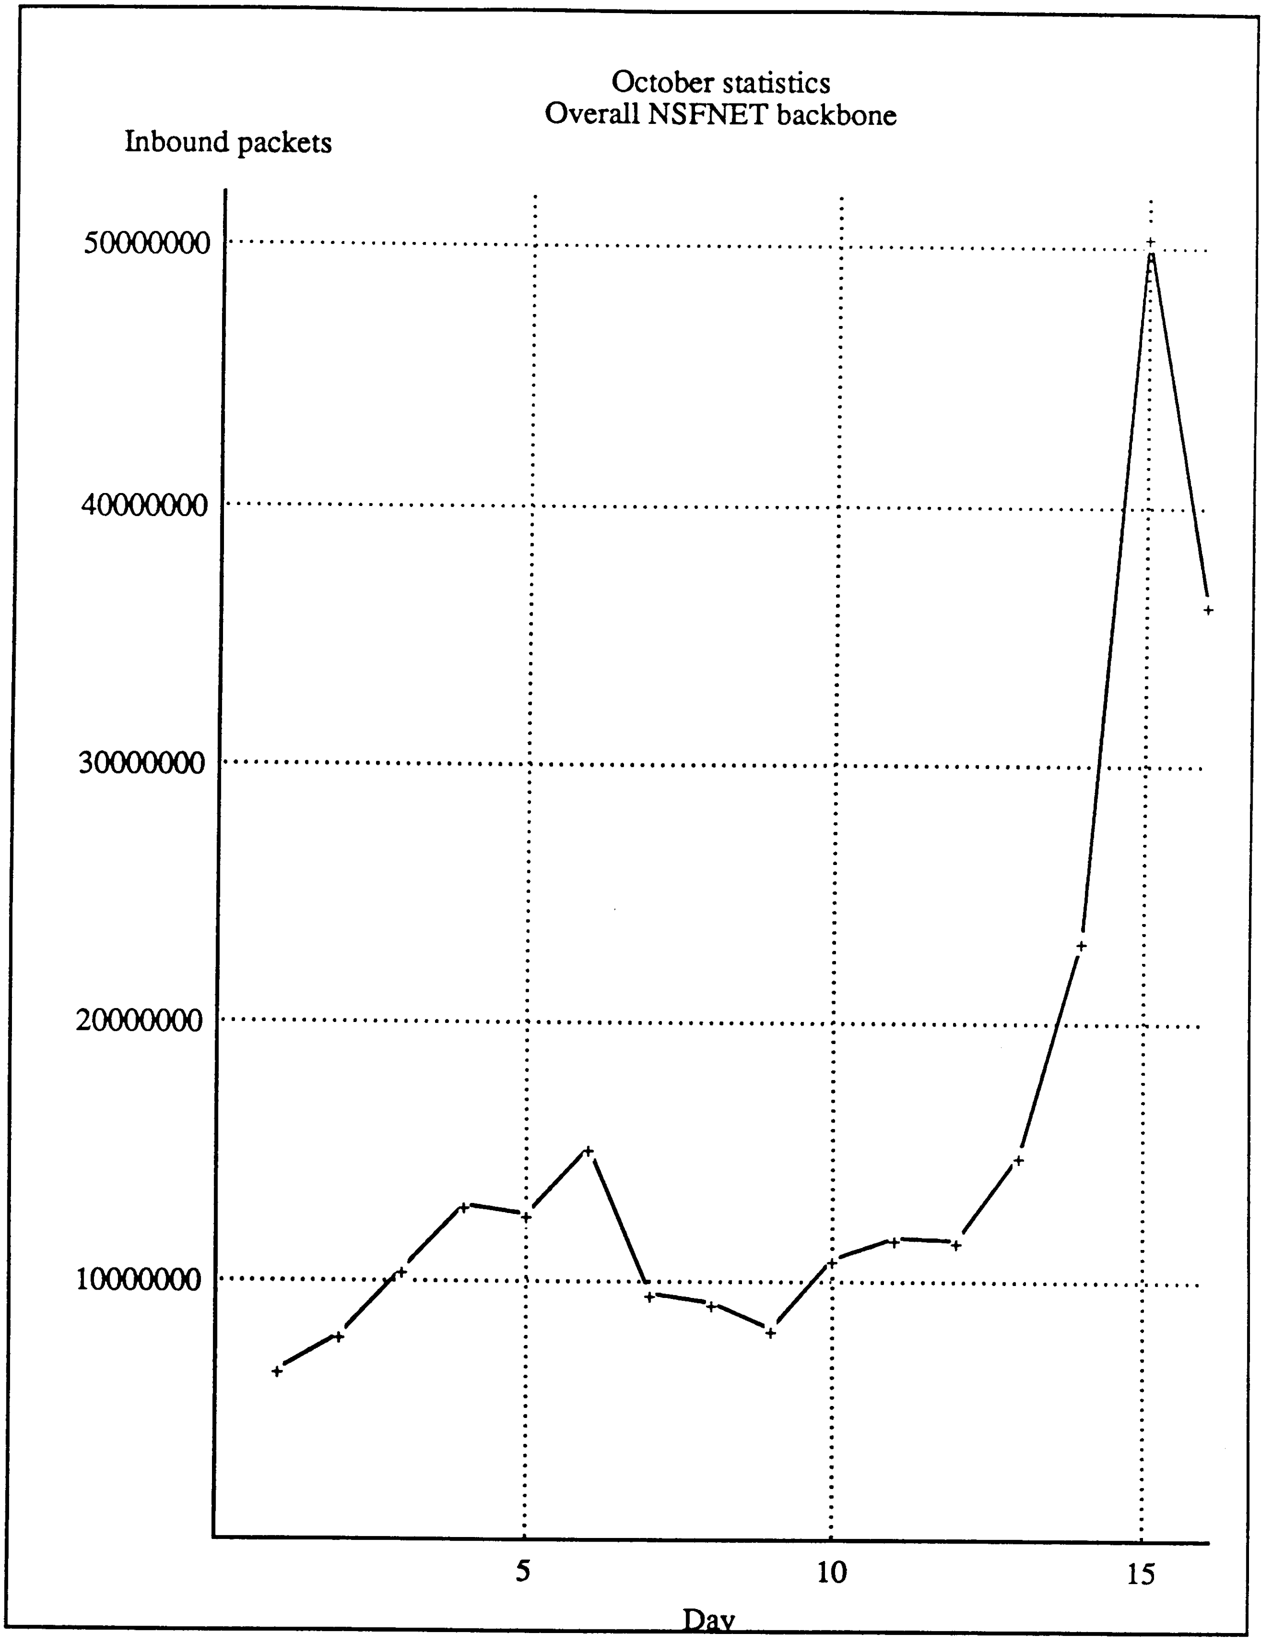
October statistics
Overall NSFNET backbone
Inbound packets
50000000
40000000
30000000
20000000
10000000
5
10
15
Day

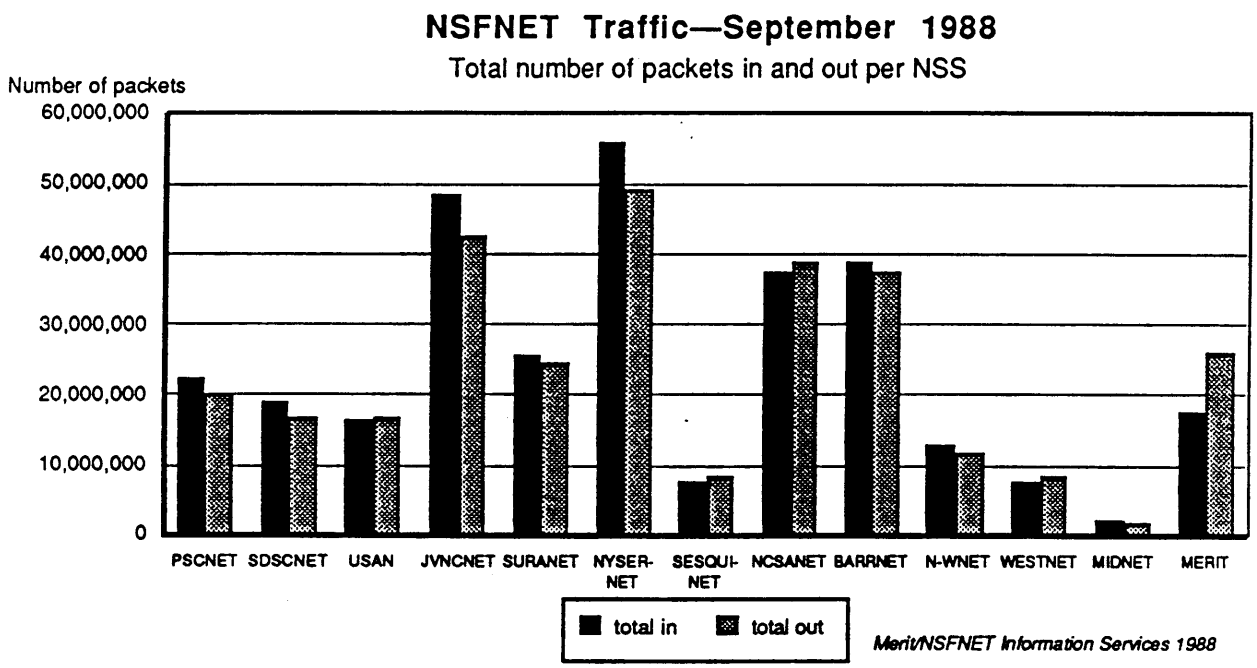
NSFNET Traffic—September 1988
Total number of packets in and out per NSS
Number of packets
60,000,000
50,000,000
40,000,000
30,000,000
20,000,000
10,000,000
0
PSCNET
SDSCNET
USAN
JVNCNET
SURANET
NYSER-NET
SESQUI-NET
NCSANET
BARRNET
N-WNET
WESTNET
MIDNET
MERIT
total in
total out
Merit/NSFNET Information Services 1988

NSFNET Traffic—Weekly packet counts 1988

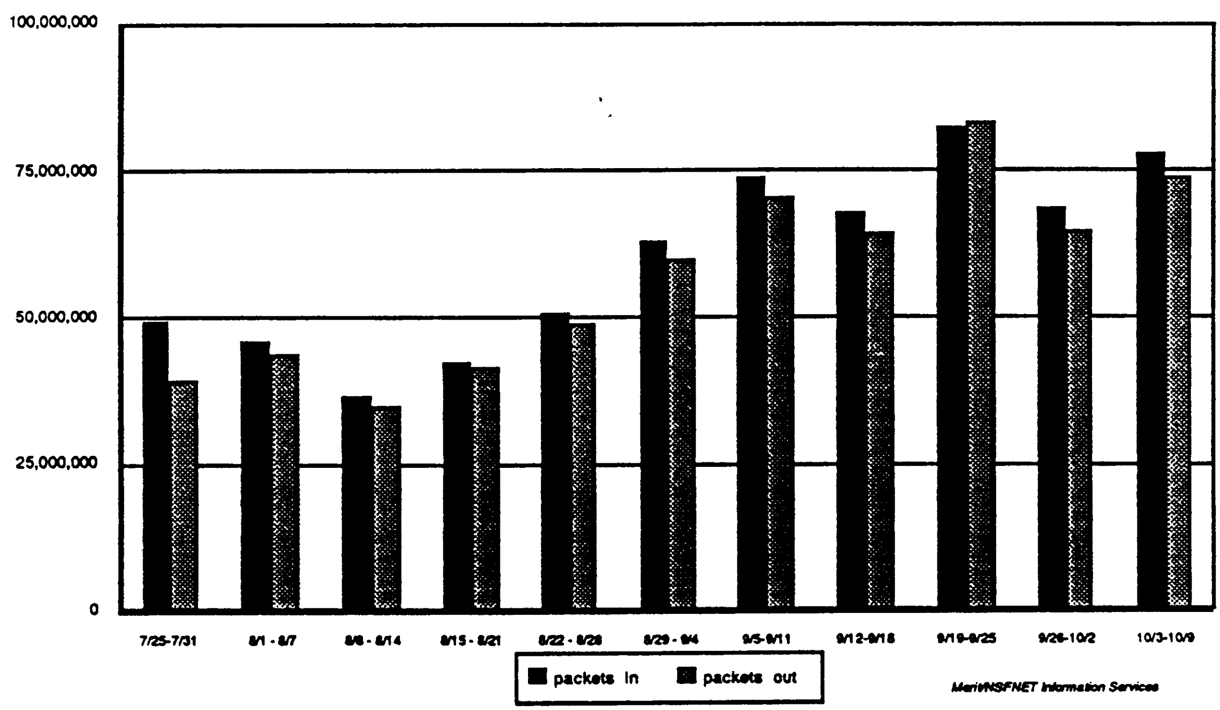

NSFNET Traffic—September 1988

Daily traffic in and out
for all NSSs

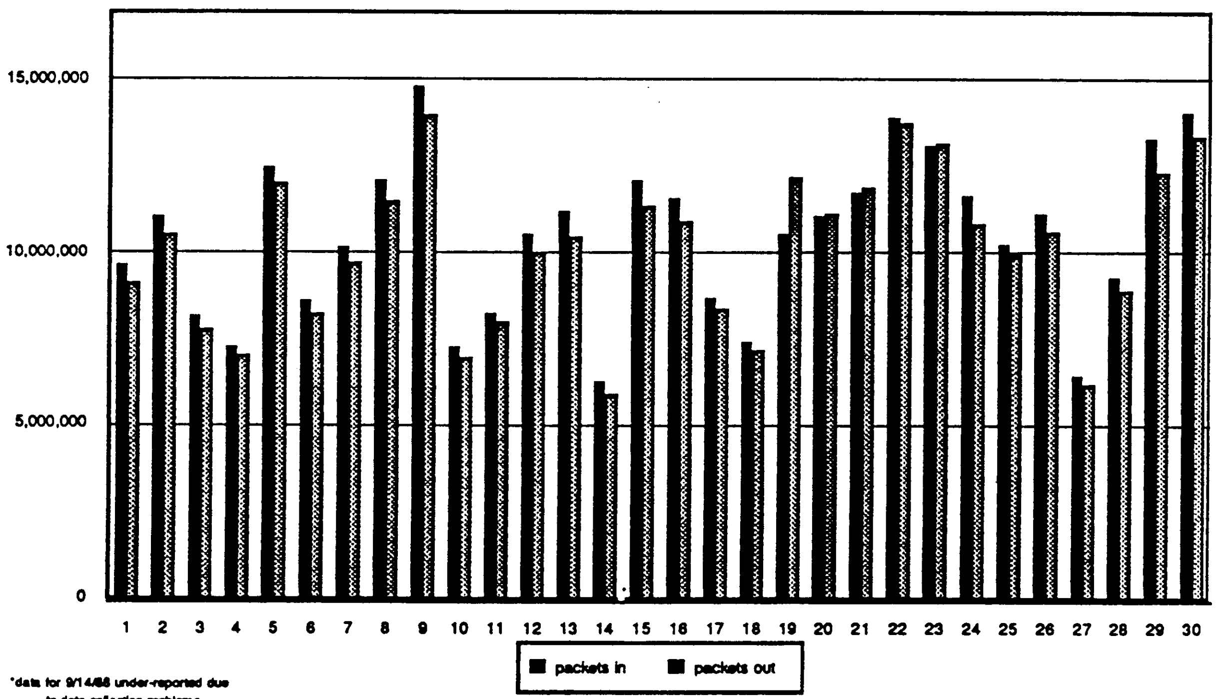

*data for 9/14/88 under-reported due
to data collection problems

NSFNET Traffic—1987-1988

Old Backbone (Aug. 87-Jun. 88)
New Backbone (Aug.-Sept. 88)

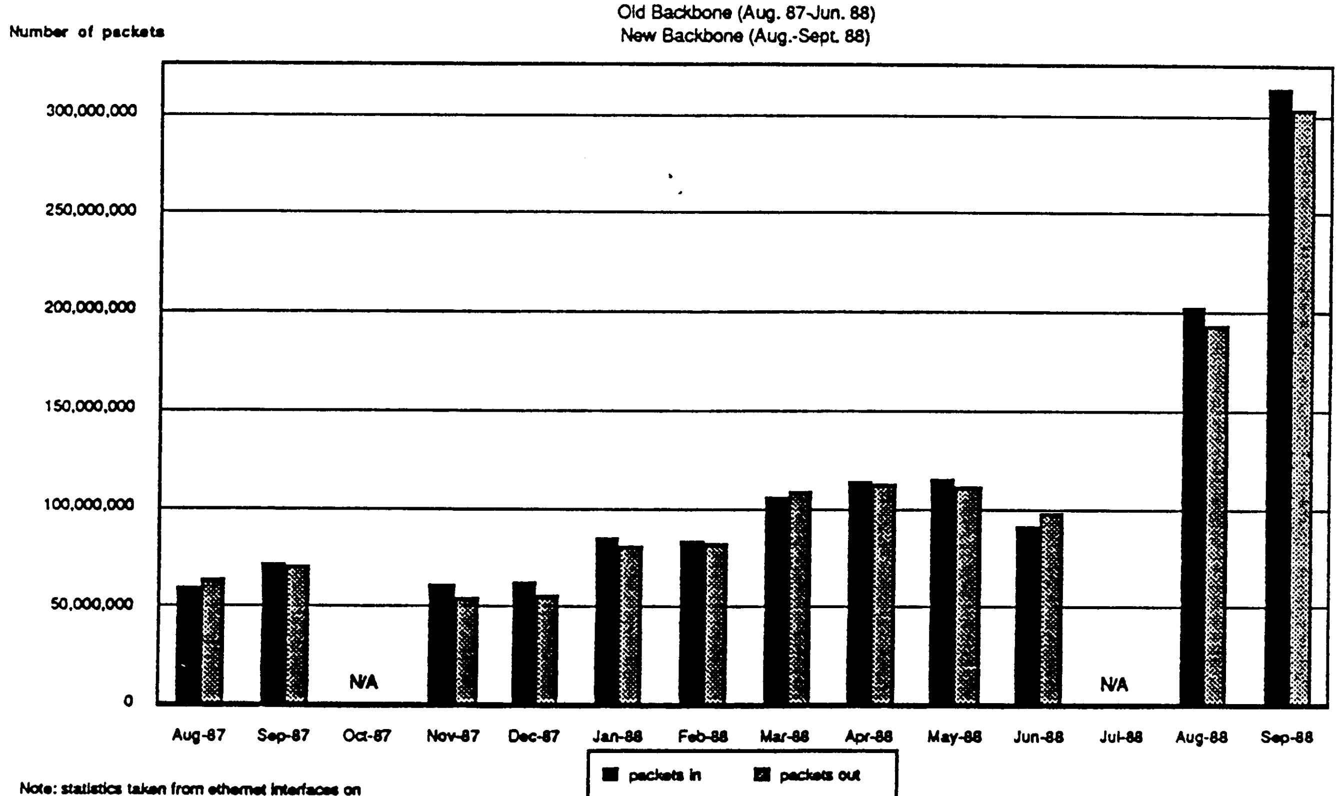

Note: statistics taken from ethernet interfaces on
old backbone, from token ring interface for new backbone

Merit/NSFNET Information Services

NSFnet Report

Part II

J.E. Drescher
IBM Corporation

National Science Foundation Network

STATUS REPORT

J.E. (Jack) Drescher
IBM Corporation
Technical Computing Systems
Milford, Connecticut

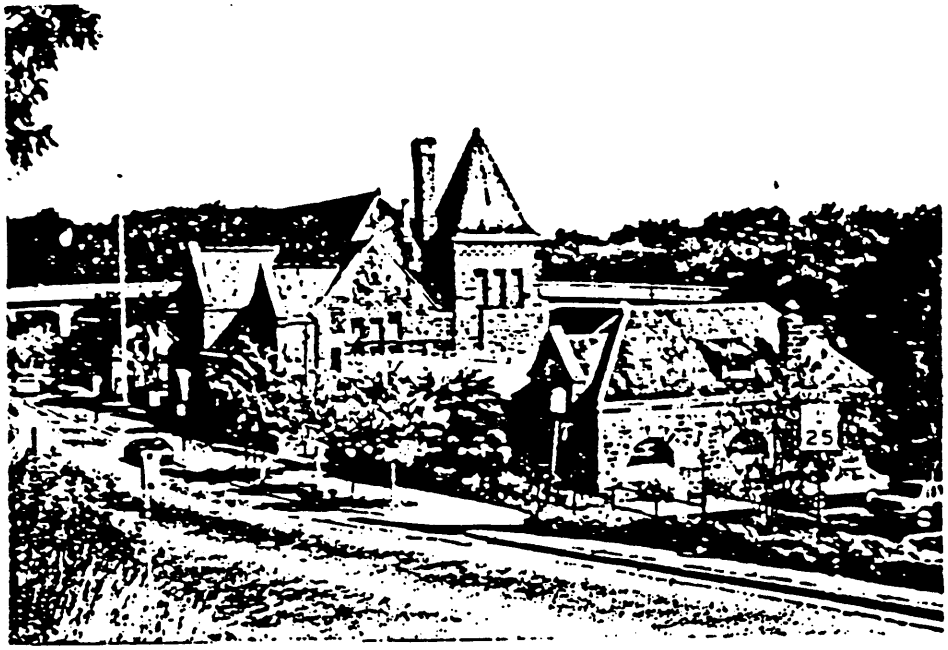
25

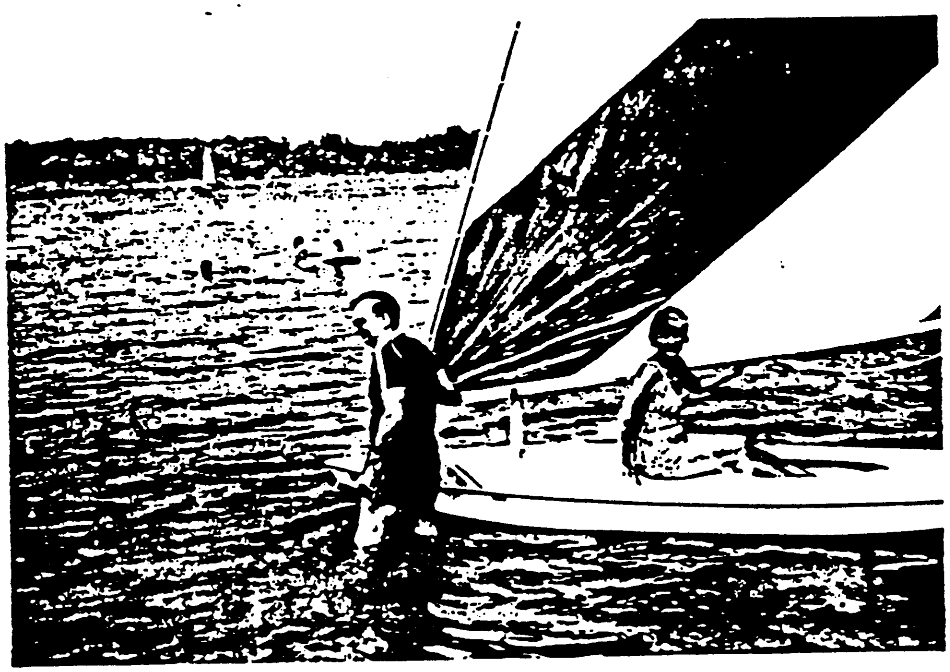

SOME IBM PEOPLE

RESEARCH

YORKTOWN

BARRY APPELMAN | **JORDAN BECKER**

MATT KORN

- JACOB REKHTER
- GEOFF CARPENTER
- FRED ROBBINS
- BILL RUBIN
- JED KAPLAN

NGUYEN HIEN

- WOLFGANG SEGMULLER
- BERT WIJNEN
- MILTON LILIE

TECHNICAL COMPUTING SYSTEMS

ANN ARBOR

JACK DRESCHER

- JIM SHERIDAN
- RICK UEBERROTH
- MYRON HEPNER

MILFORD

WALTER WIEBE

PAUL BOSCO

- RICK BOIVIE
- LOU STEINBERG
- STEVE CAPORALE
- FRANK BARTUCCA
- MATHI PACKIAM
- TIM ROLFES
- SUE WANG

GAITHERSBURG

TOM STIX

- AL WATSON
- BILL CROSTHWAIT
- MIKE SABOL

SKK 10/13/88

NETWORK MANAGEMENT OPPORTUNITY

WE CAN BECOME THE NATIONAL SHOWCASE

-- PUBLICITY
-- FACILITY
-- ATTITUDE
-- SKILLS
-- BASE PRODUCTS

WE CAN LEAD WAY IN AUTOMATING/SYNERGIZING PRODUCTS TO PROCESS

DIRECTION

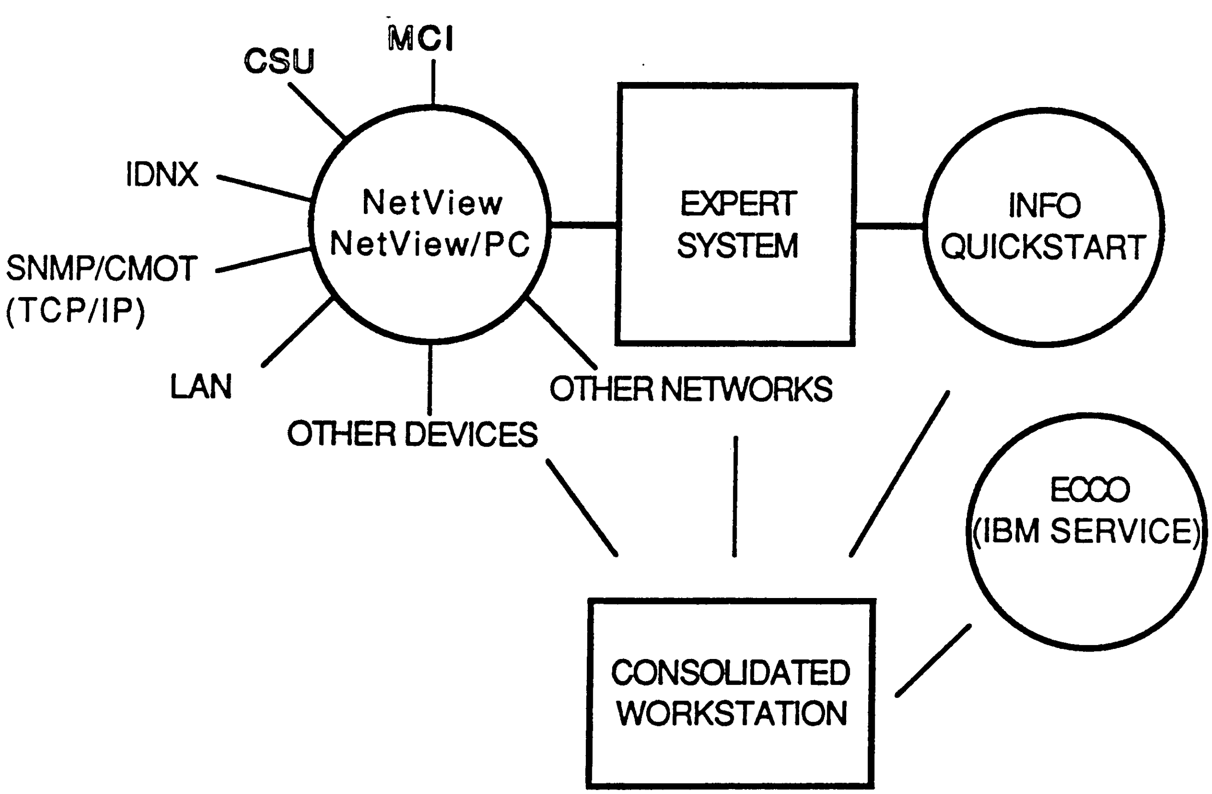

NSFNET SHORT TERM FOLLOW-ON FUNCTION

• STIMULATED BY ACCEPTANCE/TRAFFIC GROWTH
  - AND POSITIVE USER RESPONSE

• INCREMENTAL DELIVERY: 4Q88 THRU 4Q89

• APPROXIMATELY 50 LINE ITEM CANDIDATES DEFINED

- IBM HAS ASSIGNED INITIAL PRIORITIES, TARGET DATES
- PARTNERS REVIEW/INPUT 10/20/88
- TARGET TO CLOSE PLAN 11/01/88 (TIGHT)

NOTE: DOES NOT INCLUDE UPGRADE TO T_3

• FEATURES

- PERFORMANCE/CAPACITY ENHANCEMENTS
- MORE AUTOMATED CONTROL/ADJUSTMENT
- ENHANCED NETWORK MANAGEMENT
- ADVANCED PROTOCOLS e.g., SNMP. ➡ CMOT, EGP_3
- STREAMLINED CONFIGURATION ITEMS

• SOME EARLY SPECIFICS

- MIB INTERFACE TO NSS
- X.25 ARPANET ADAPTOR
- 3 COM ETHERNET ADAPTOR
- RTIC IDNX DRIVER (T_1)

ARPANET/DDN Report

Marianne Lepp
BBN

ARPANET STATUS

Oct 18, 1988

Marianne Lepp

ARPANET 6-23-88 TOPOLOGY

— VSAT Trunk

— BSAT/Wideband Net Trunk

--- Terrestrial Trunk To Be Removed Soon*

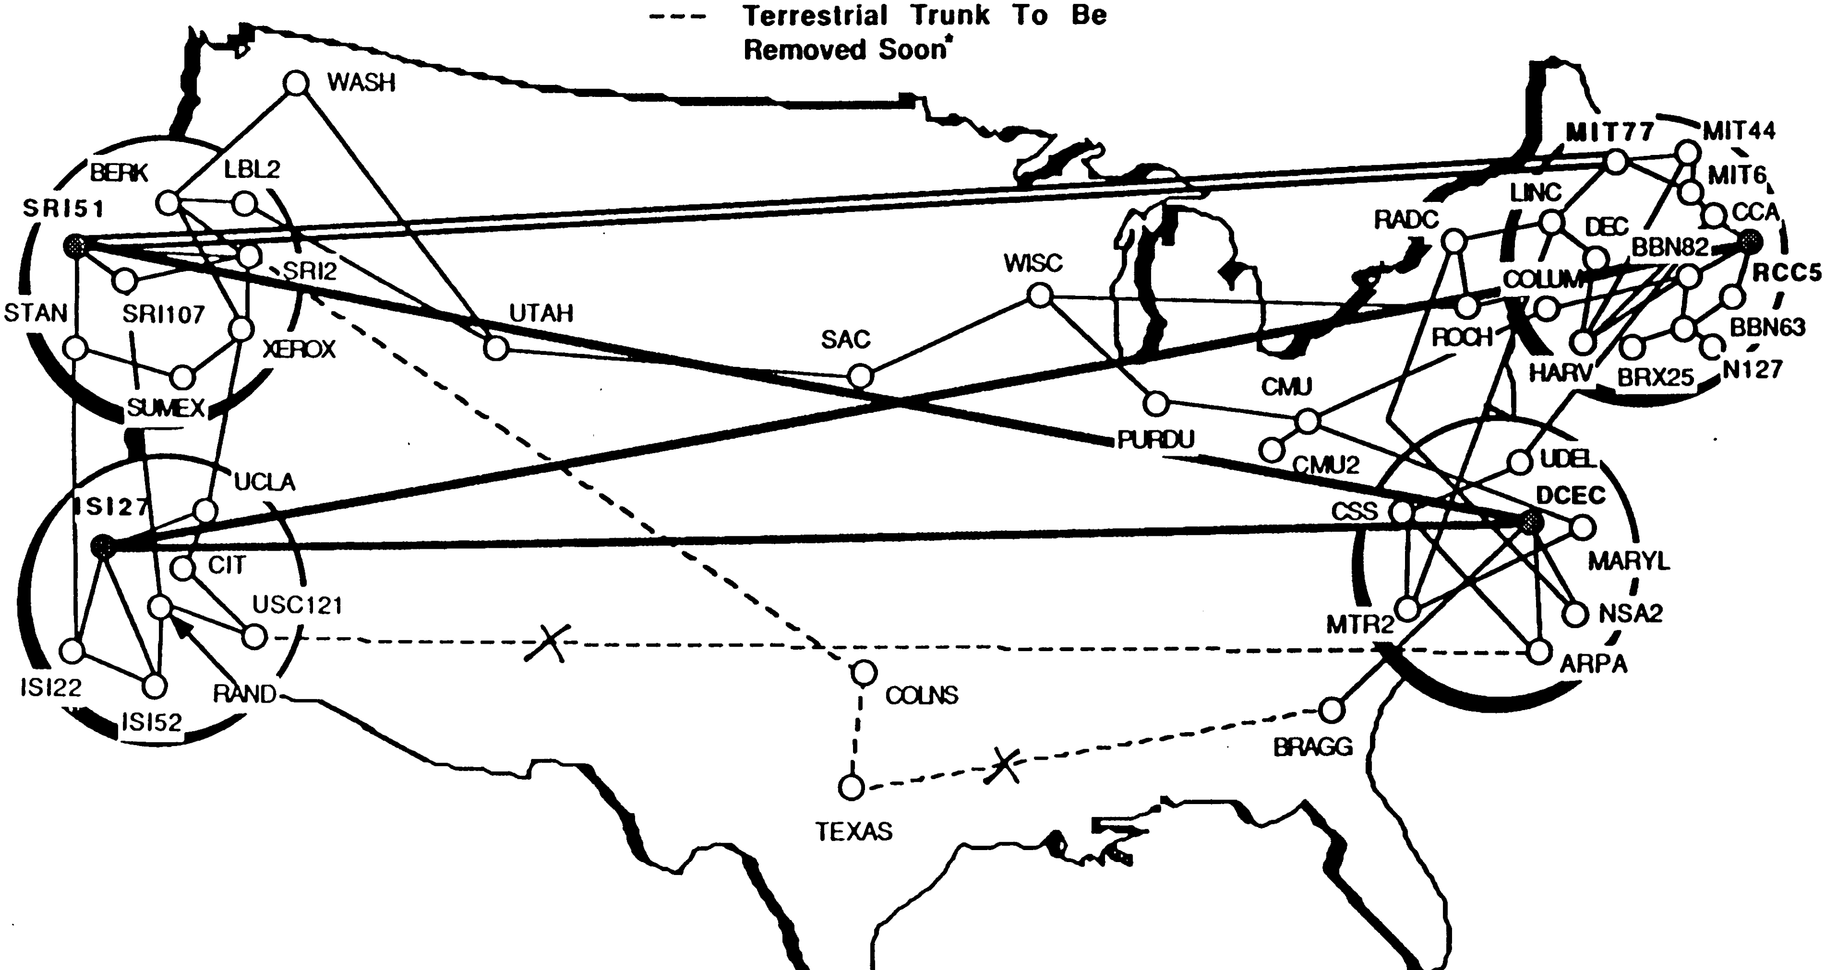

* Path Through SAC To Be Disc. 8/1/88(??)

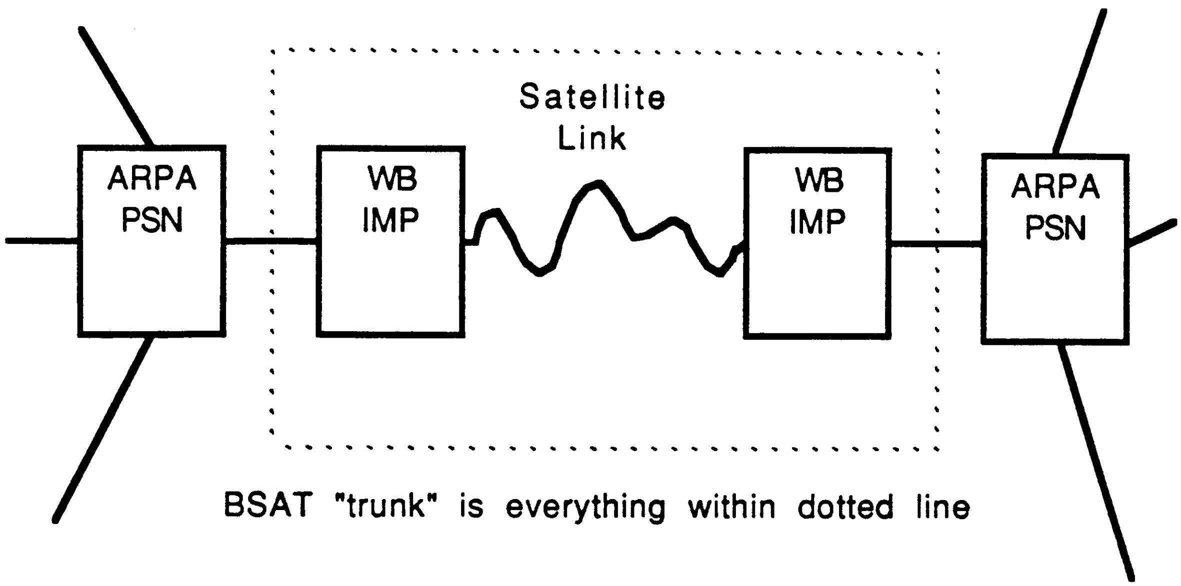

BSAT "trunk" is everything within dotted line

BSAT Lines

- Installed between RCC5 and ISI27
  DCEC and ISI27
  DCEC and SRI51

- Intended to replace transcontinental terrestrial trunks

- Network diameter reduced from 9 to 7 hops

- "Piggybacked" on Wideband Net

- Frequent outages caused by high and variable delay

- Improved by lengthening retransmission timer, increasing number of logical channels, and relaxing "line down" criterion

- Remaining outages caused by Wideband Net resets

Arpanet Topology

Summary Statistics

	Aug 87	Feb 88	Jun 88
Nodes	45	43	50
Trunks	67	68	82
Trunks per Node	3.0	3.2	3.3
Active Hosts	170	155	202
Hosts per Node	3.8	3.6	4.0

Arpanet Performance

Week Long Summary Statistics

	Aug 87	Feb 88	Jun 88
Host Traffic Msg / Sec	229	325	320
Total Internode Throughput (KB/S)	300	332	336

Arpanet Performance

Peak Hour Summary Statistics

	Aug 87	Feb 88	Jun 88
Host Traffic Msg / Sec	296	447	470
Total Internode Throughput (KB/S)	397	494	449

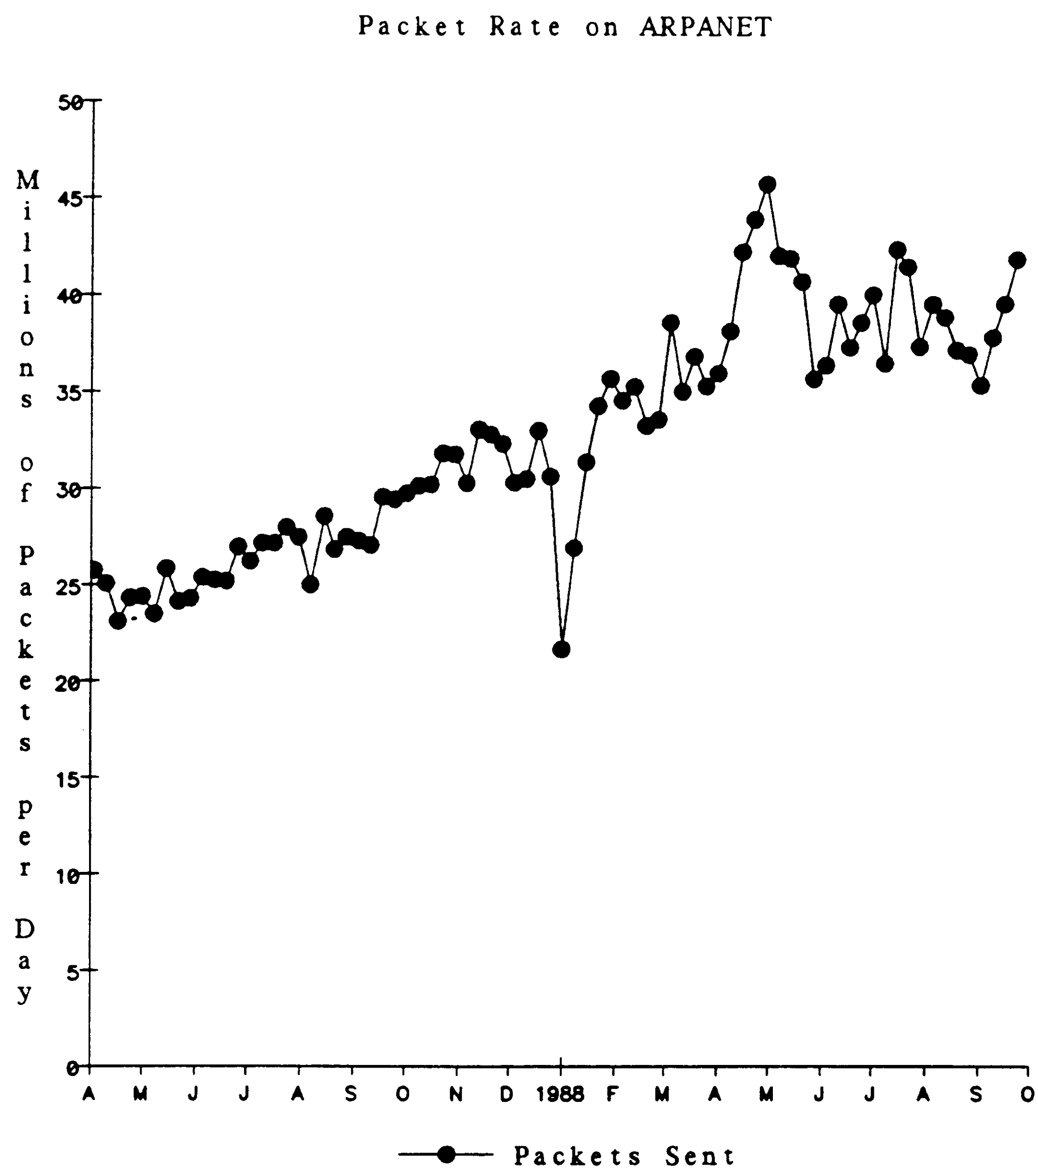
Packet Rate on ARPANET
Millions of Packets per Day
50
45
40
35
30
25
20
15
10
5
0
A M J J A S O N D 1988 F M A M J J A S O
Packets Sent

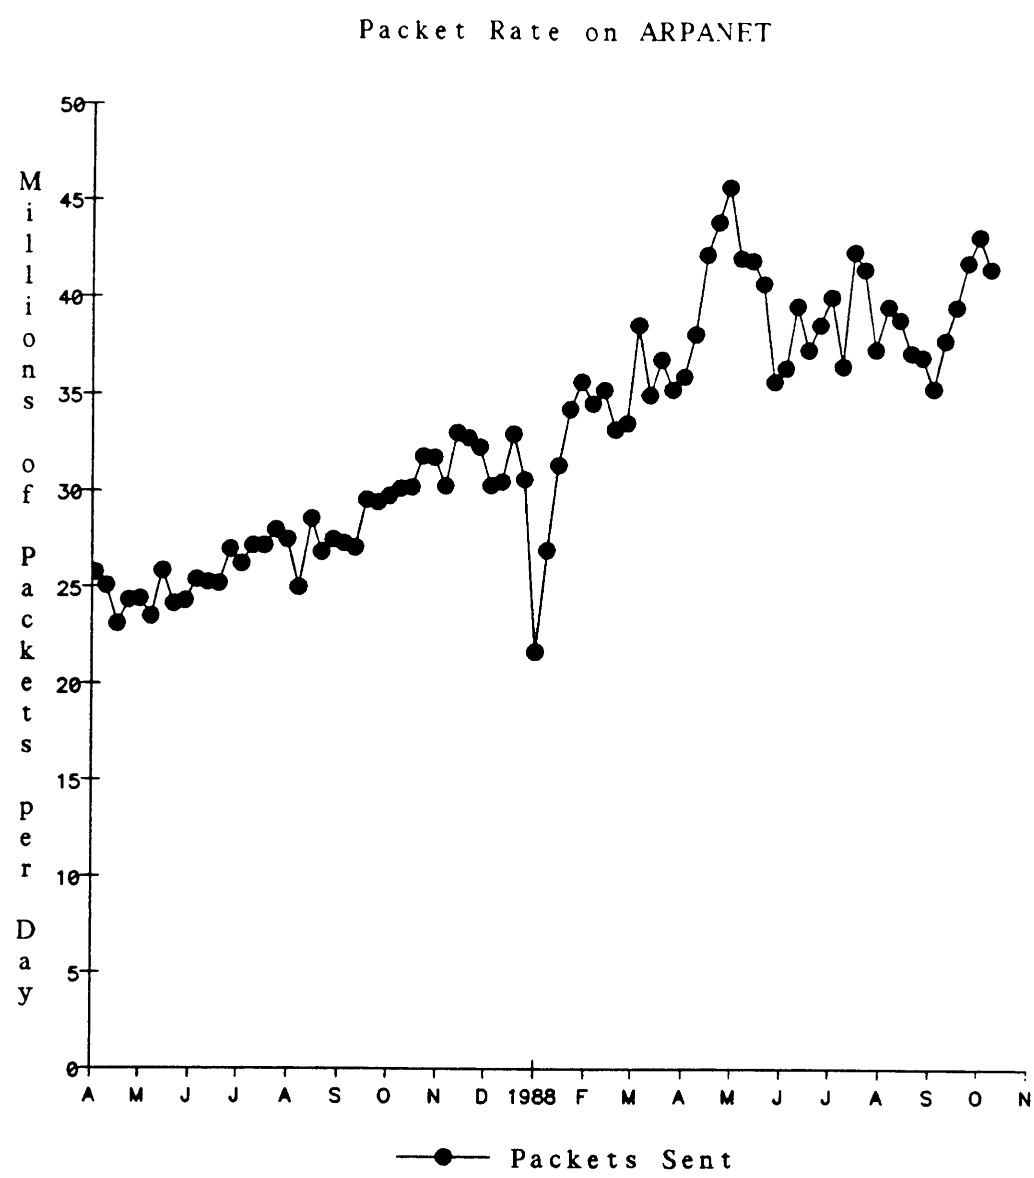
Packet Rate on ARPANET
Millions of Packets per Day
50
45
40
35
30
25
20
15
10
5
0
A M J J A S O N D 1988 F M A M J J A S O N
Packets Sent

CPU Utilization of ARPANET Node 14 (CMU)
6-10 June 1988

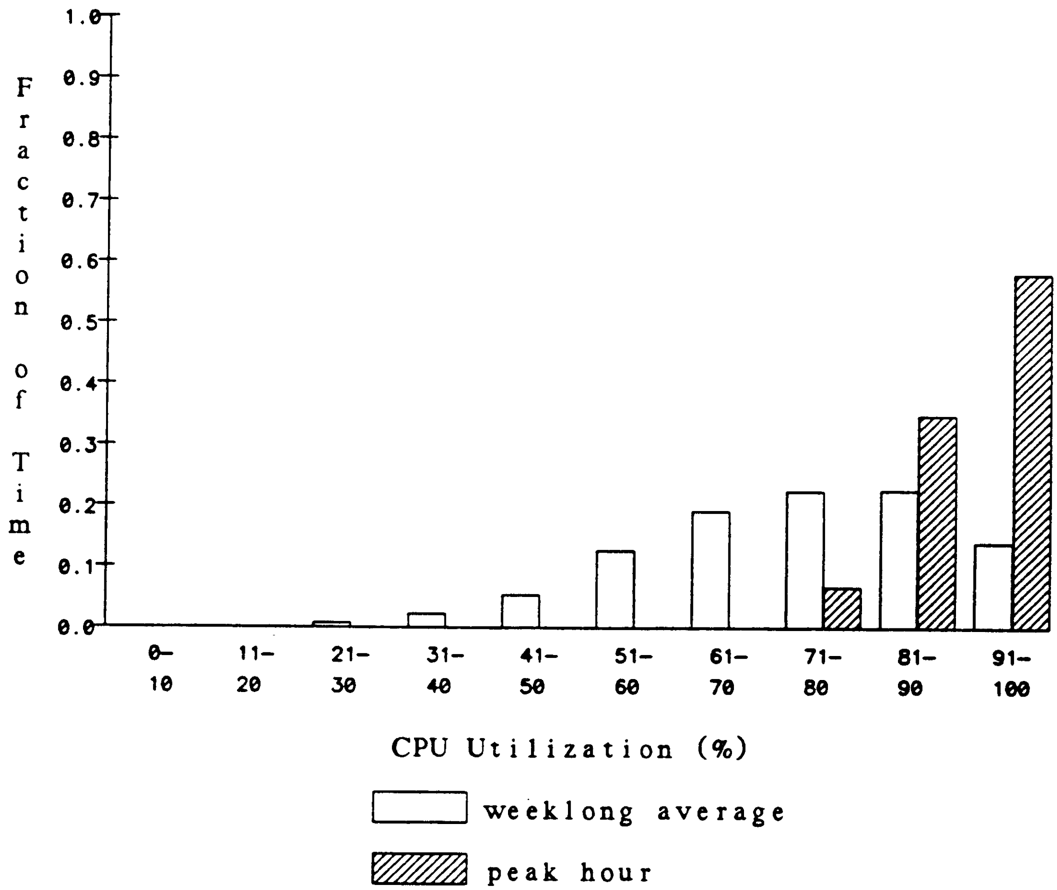

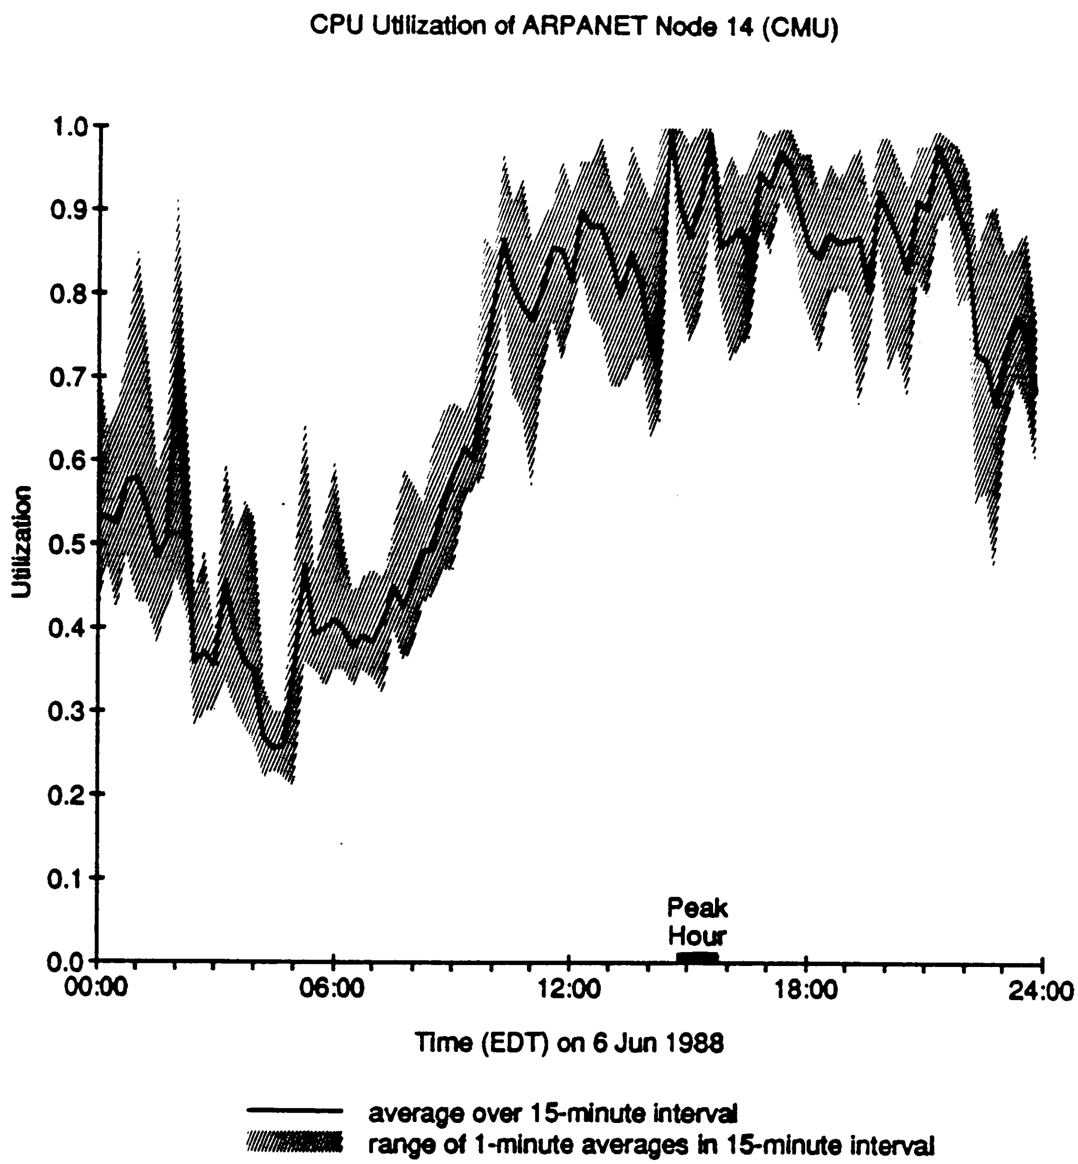
CPU Utilization of ARPANET Node 14 (CMU)
Utilization
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
Peak
Hour
00:00
06:00
12:00
18:00
24:00
Time (EDT) on 6 Jun 1988
average over 15-minute interval
range of 1-minute averages in 15-minute interval

CPU Utilization of ARPANET Node 27 (ISI27)

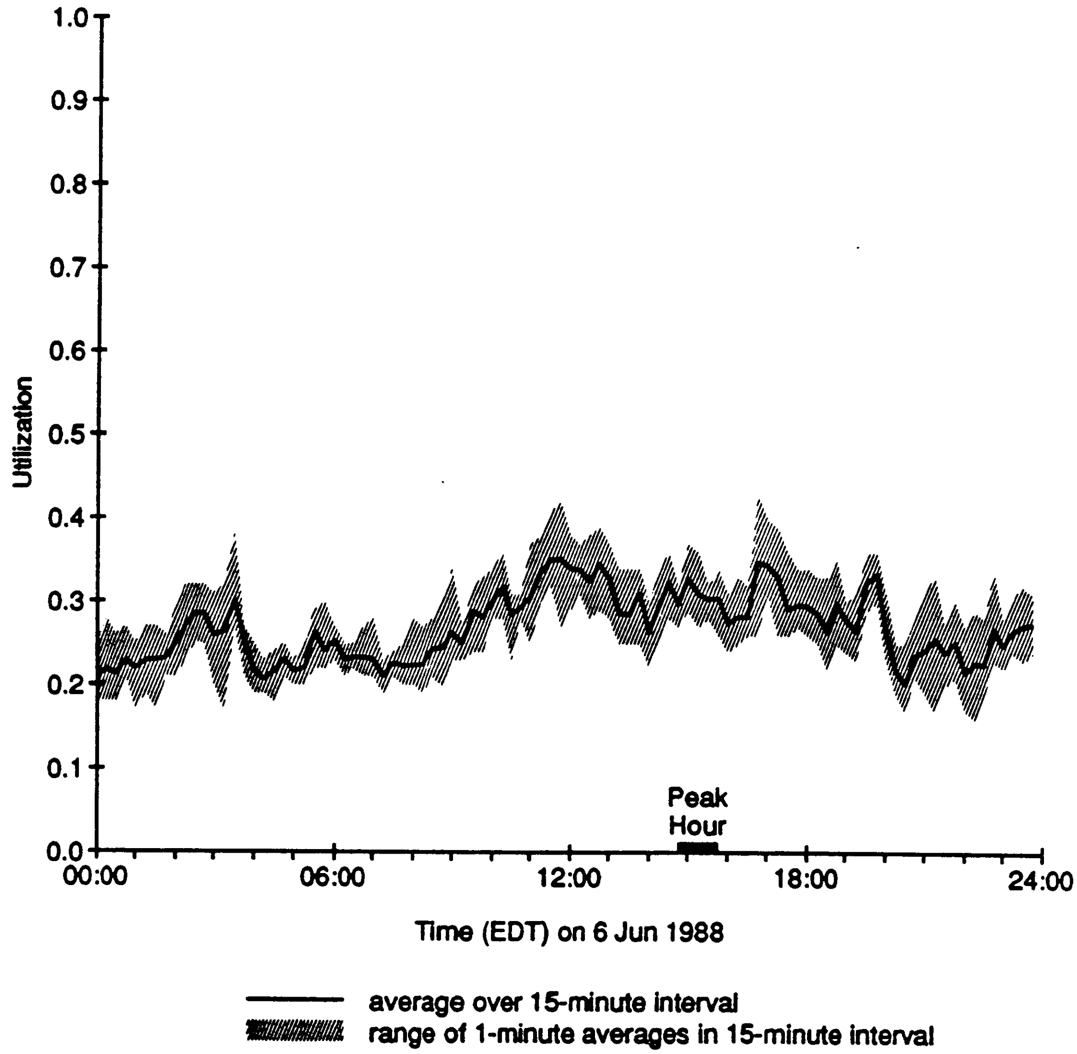

DDN Report: Transition of DDN Mailbridges from LSI-11 to Butterfly Gateways

Michael Brescia
BBN Communications Corporation

TRANSITION OF DDN MAILBRIDGES FROM LSI-11 TO BUTTERFLY GATEWAYS

Michael Brescia

BBN Communications Corporation

SEQUENCE OF EVENTS

INSTALLATION AND TESTING

- Assign Sites and PSN Connections
- Deliver and Install Butterfly Mailbridge Systems
- Test Hardware Installation and Software Integration
- BETA Testing with Selected Sites

SEQUENCE OF EVENTS

TRANSITION PLAN

- Prepare Mailbridge Transition Plan for Host and Gateways

- Design and Announce via "DDN Management Bulletin"
  - New Load Sharing Assignments
  - Schedule for Dropping Old Mailbridges

- Announce via EGP-PEOPLE Mailing List
  - New EGP Servers
  - Schedule for Dropping Old EGP Servers

SEQUENCE OF EVENTS
RETIREMENT OF LSI-11 MAILBRIDGES

- Remove Old Mailbridges from Service
- Remove Old EGP Servers

STATUS OF MAILBRIDGE TRANSITION

- Sites and Connections on Arpanet and Milnet Assigned
  - 3 on East Coast (DCEC, MITRE, BBN)
  - 3 on West Coast (AMES, ISI, LBL)

- July
  - Butterfly Mailbridge Hardware Installed

- August - Present
  - Software and Hardware Testing in Progress
  - Establishing Procedures with the Site Coordinators
  - Discovered EGP Problems
  - Examining Solutions

SCHEDULE FOR REMAINING MAILBRIDGE TRANSITION

- November/December
  - Announce BETA Test

- And beyond

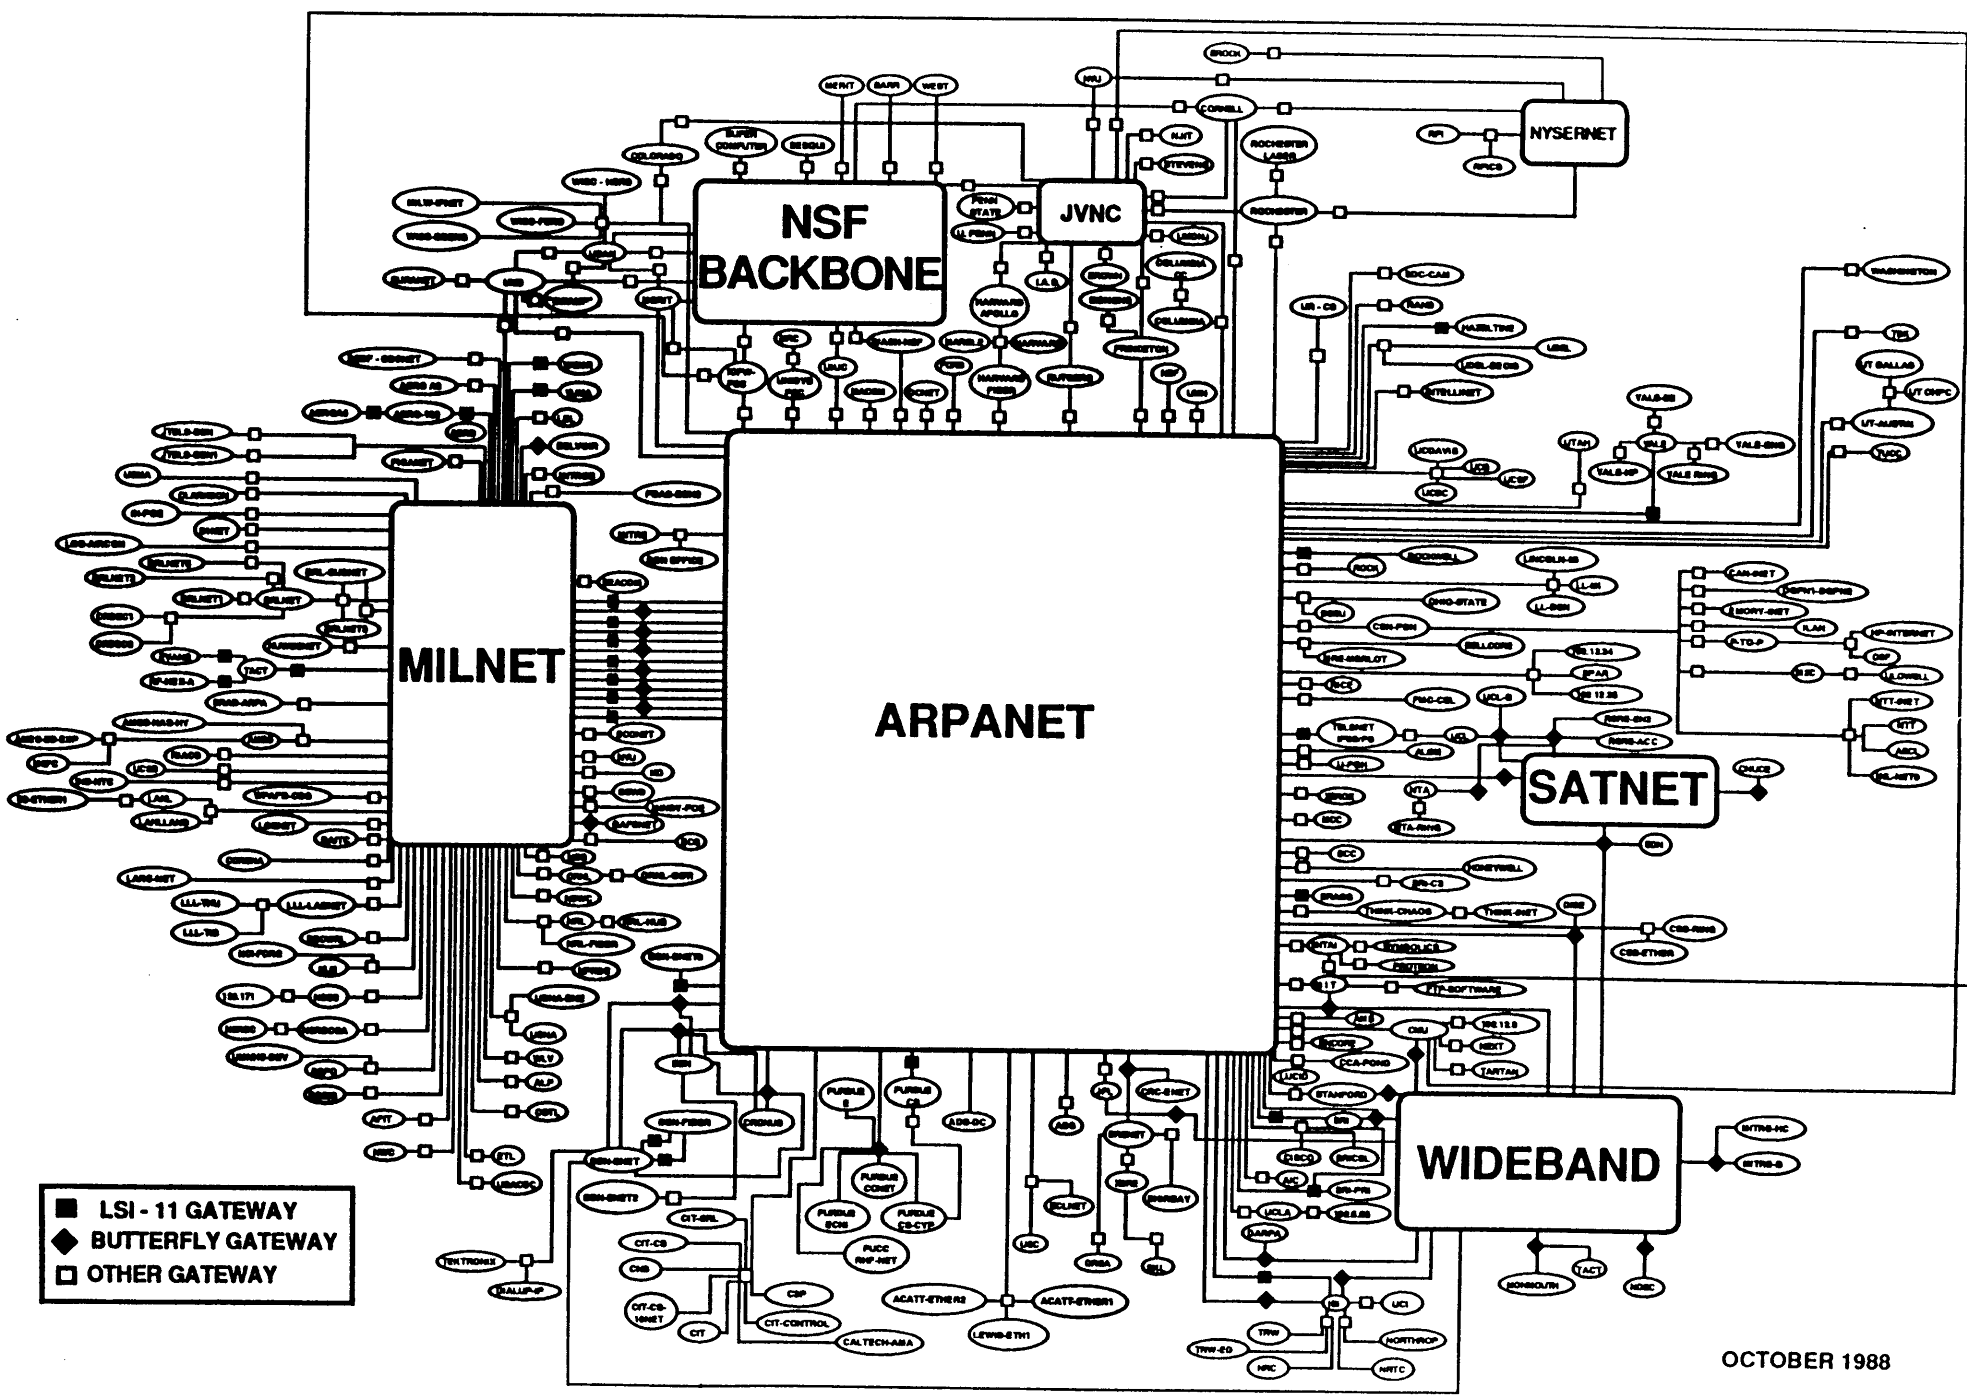

OCTOBER 1988

BBN Communications Corporation

Internet Growth in Networks

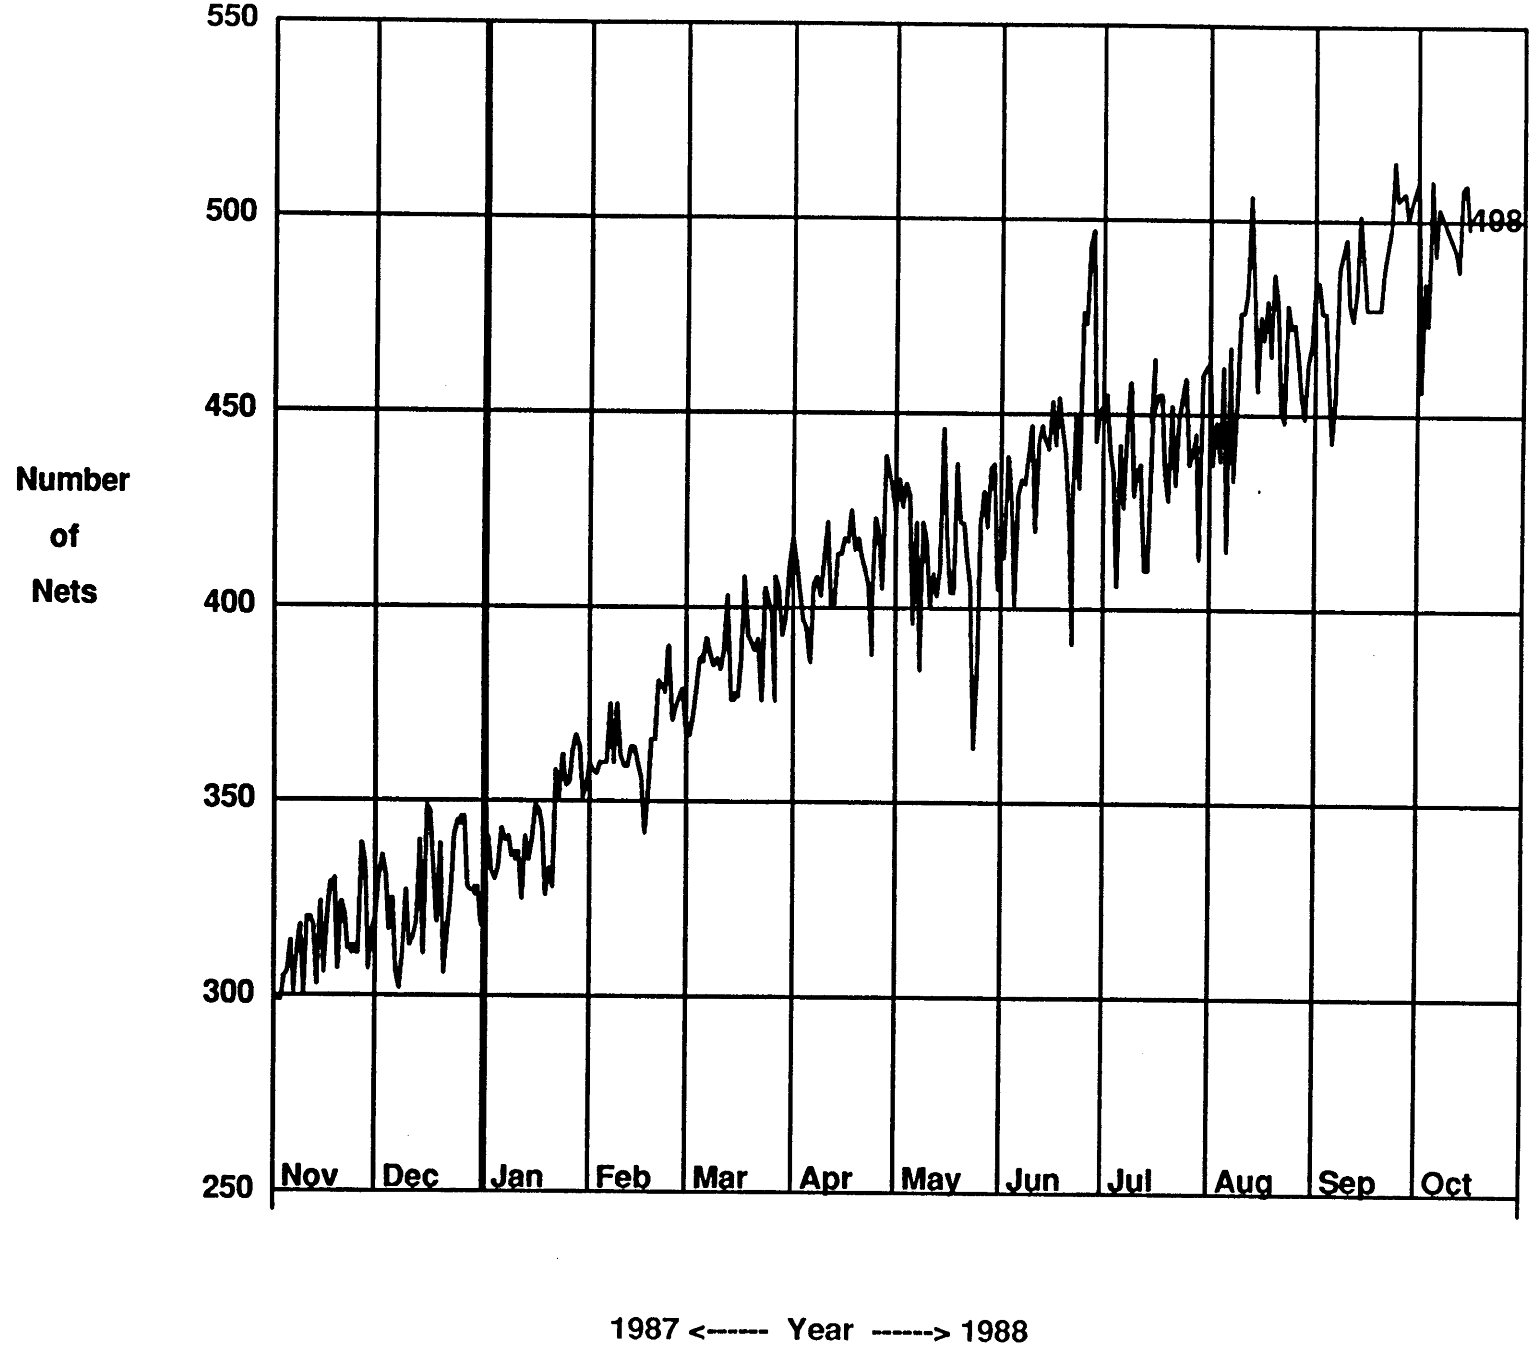

INTEROP 88 Network Report

Philip Almquist
Stanford University

The INTEROP 88 Network: Design, Problems, and Lessons Learned*

Philip Almquist

* WARNING: do not try this at home. Professional stunt driver required.

Introduction

- Large scale demonstration of TCP/IP Interoperability
    - 49 vendors
    - Approximately 250 hosts and gateways
    - Almost 2 miles of cabling
    - High-speed connections to ARPANet, MILNet, NSFNet, ...
- Standalone network for CMOT (NETMAN) demonstration
- Very successful
- Purposes of this talk
    - Inform
    - Stimulate IETF action

Description of the network

- Designed by Peter DeVries and myself
- Subnetted class B net
- Multiple media
    - Ethernet
    - Thin Ethernet
    - Ethernet over twisted pair
    - Ethernet over fiber
    - PRONet-80
    - IBM/802.5 token ring
    - SLIP
    - Packet radio
    - (also Hyperchannel, PRONet-10, T-1, and Ethernet over broadband in individual booths)
- Tree topology - no alternate routes
- Small subnets
- All backbone routers in NOC
- Built in 5 1/2 days by Peter, myself, 3 part-time technicians, and a horde of volunteers

INTEROP 88

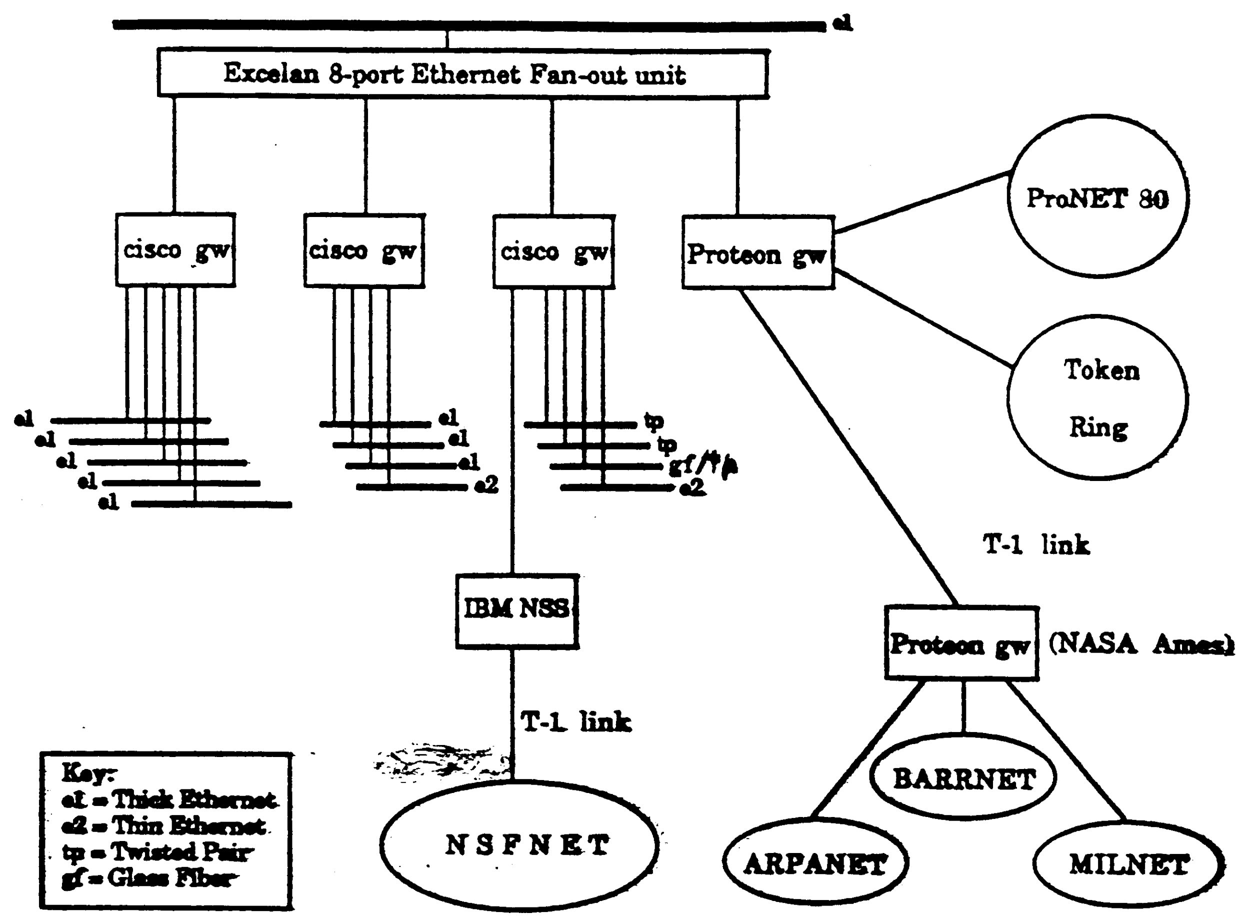

Show and Tel-Net Topology

Participating Vendors:

3Com
ACC
Apple Computer
Banyan Systems
BBN Communications
COMPUTERWORLD
CMC
Computer Network Technology
Concurrent Computer
Convergent Technologies
cisco Systems
DCA/SRI International
DEC
Encore
Eon Systems
Excelan/TGV/Kinetics
FTP Software
Halley Systems
Hewlett-Packard
Highland Software
IBM/MCI/Merit/CMU
Interactive Systems
InterCon
Interphase
Lachman Associates
Mitre/Unisys (NetMan)
Network General
Network Research
Network Solutions
Network Systems
Prentice-Hall
Prime Computer
Process Software
Proteon
Sirius Systems
Spider Systems
Sun Microsystems
SynOptics Communications
Syntax Systems/10Net
Sytek
Tandem Computers
TCL
TRW
Ungermann-Bass
UNIX World
Vitalink Communications
VXM Technologies/MIPS
Wellfleet Communications
Western Digital
The Wollongong Group
Xyplex

INTEROP '88
Show and Tel-Net
September 28-30, 1988
Santa Clara Convention Center

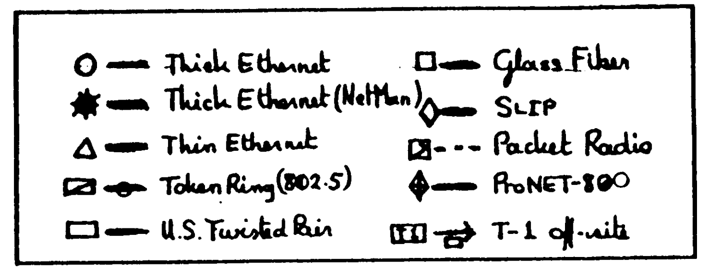

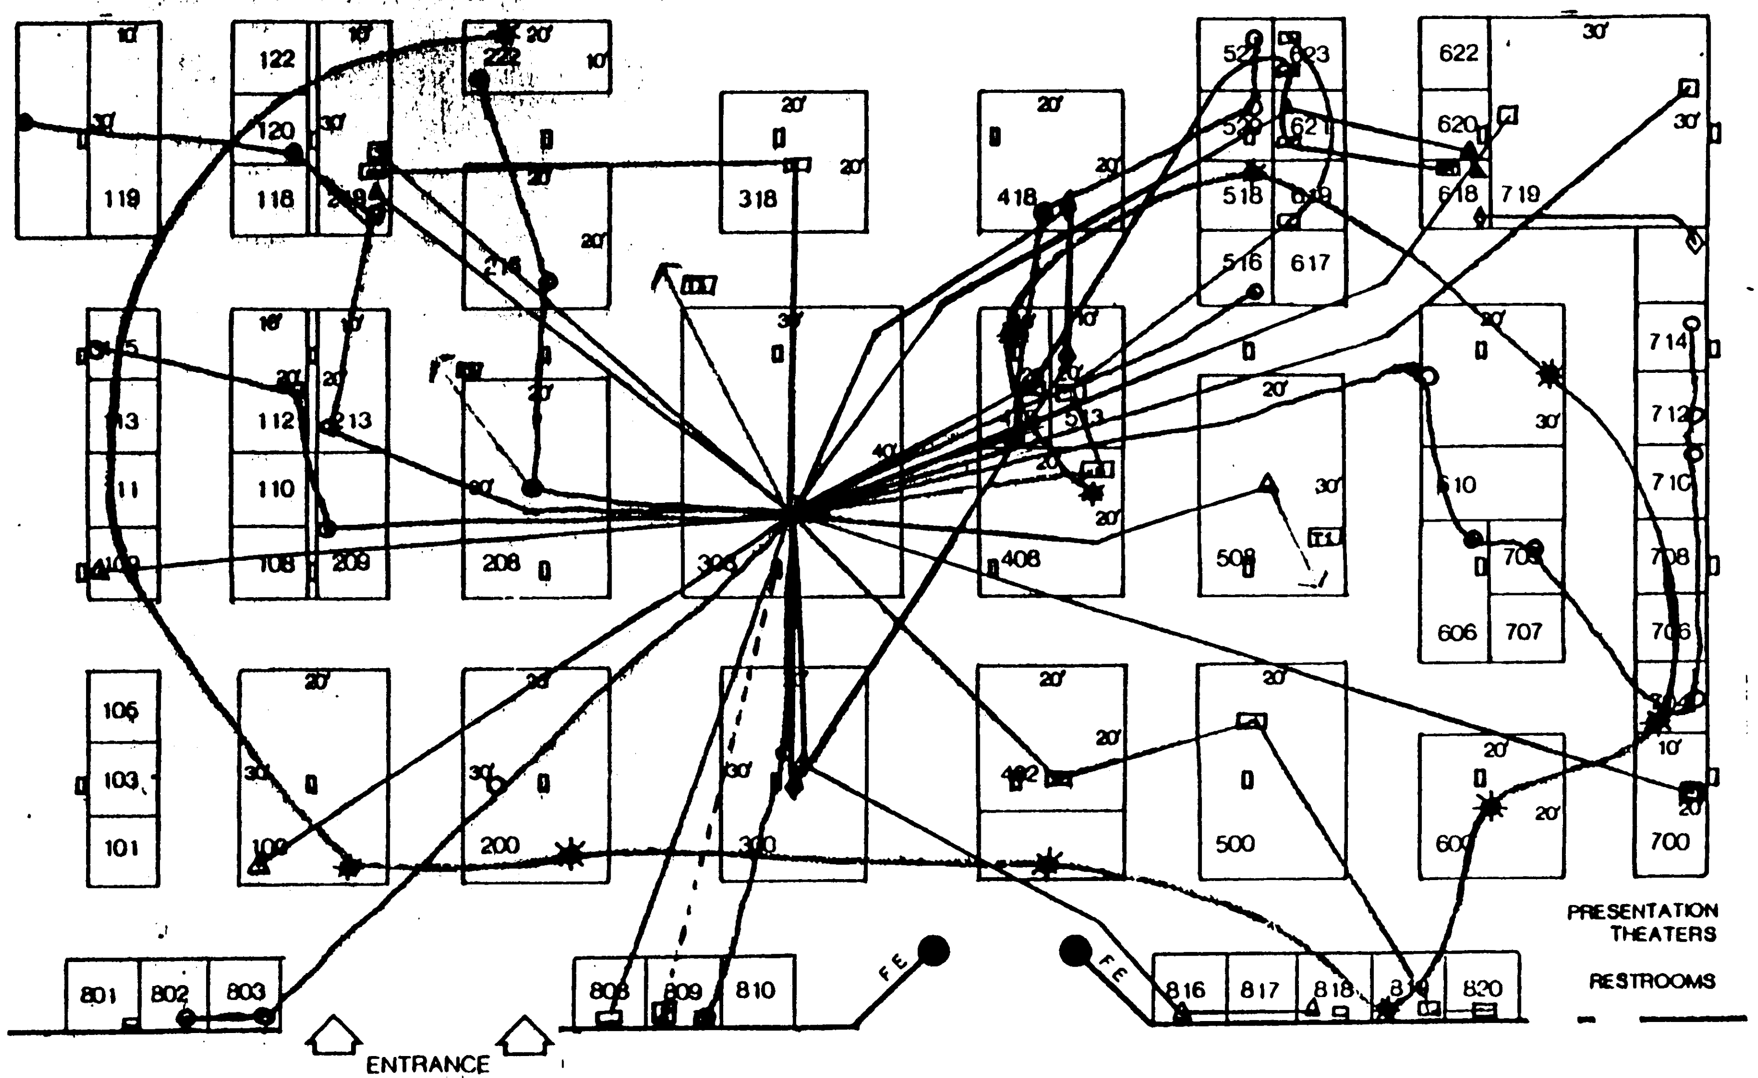

Sponsored by Advanced Computing Environments, Inc.

Designed, installed and managed by The Wollongong Group, Inc.

Cabling

- What we did
  - Cabling hung from ceiling
  - Intentionally very visible
  - Tranceivers reachable with a ladder
- Problems
  - Ran out of cable
  - T-1 didn't want to work (of course!)
  - Too many people inside the wiring center
  - One booth on wrong subnet because vendor rewired it!
  - Mysterious temporary failure of one Ethernet segment on second day of show
  - The usual minor problems...

IP address assignment/host table creation

- What we did
    - We obtained a domain: ShowNet.COM
    - Vendors filled out host questionnaires
    - We assigned IP addresses and created a zone file
    - A program read the zone file to generate the IN-ADDR.ARPA zone files and a HOSTS.TXT
- Problems
    - Questionnaires were returned late and filled out incorrectly
    - No host table czar
    - Zone file inaccessible until T-1 came up
    - Some vendors required /etc/hosts format

Domain service

- What we did
  - 3 authoritative servers (two off-site)
  - Off-site servers set up as secondaries
  - Small TTL's and refresh times
- Problems
  - Syntax errors in the master files
  - Little familiarity with domain software on primary
  - Miscommunication between the NIC and Wollongong
  - Root server update procedure failed
  - Primary not installed until the day before the show

Lessons
  - Make sure domain requests get honored well before you need them
  - Root server updates are probably not as robust as they should be
  - Hand-typed zone files require a syntax checker program

Network Management

- What we did
  - SUN running Wollongong/NYSERNet SNMP tools
  - Protocol analyzer
  - Smart Ethernet terminator
- Problems
  - pre-SNMP code on cisco routers the first day
  - bug in Proteon SNMP
  - SUN had incomplete/incorrect SNMP configuration files
  - Most segments didn't have extra tranceivers for monitoring
  - NOC personnel unfamiliar with the particular management tools available
- Lessons
  - Network management tools are useless if they can't be used quickly and easily when problems occur

Internet Protocol Police

Notice of Protocol Violation

IP Address of Offender: ____________

Domain Name of Offender: ____________________

Improper Configuration

- ☐ Wrong IP Address
- ☐ Wrong IP broadcast address
- ☐ Wrong Subnet Mask (or subnets not supported)
- ☐ Excessive Broadcasting
- ☐ ARPing for Broadcast Address
- ☐ Invalid Ethernet/Subnet address
- ☐ ______________________

Warnings

- ☐ Disabling UDP checksums
- ☐ Dropping packets while resolving addresses
- ☐ Tinygram generation
- ☐ Improper round-trip-timing
- ☐ Lack of congestion avoidance
- ☐ ______________________

Protocol Violations

- ☐ Forwarding broadcast packets
- ☐ TCP response to broadcast
- ☐ ICMP response to broadcast
- ☐ Ignoring ICMP redirects
- ☐ Ignoring ICMP source quench
- ☐ Broadcast TCP packets
- ☐ TCP Keepalives
- ☐ TCP aborts on ICMP message while connected
- ☐ Misc. protocol error
  - ☐ TCP ☐ IP
  - ☐ UDP ☐ ICMP
  - ☐ ARP ☐ _____
- ☐ ______________________

Inspector: ______________________ **Date:** ___/___/___

Internal routing

- What we planned
  - Use RIP throughout
  - Back doors were allowed only if not advertised
- What we actually did
  - Core routers sent all routes via RIP
  - Core routers believed RIP only from other core routers
  - Core routers had static routes to subnets behind non-core routers
  - Hosts and non-core routers to avoid RIP and use a static default route
  - Reasoning: possible bogus routes from misconfigured RIP-speakers
- Problems
  - Large and unnecessary RIP broadcasts (from NSFNet routes) caused prolems for PC's
- Lessons
  - Static routing is a b*tch

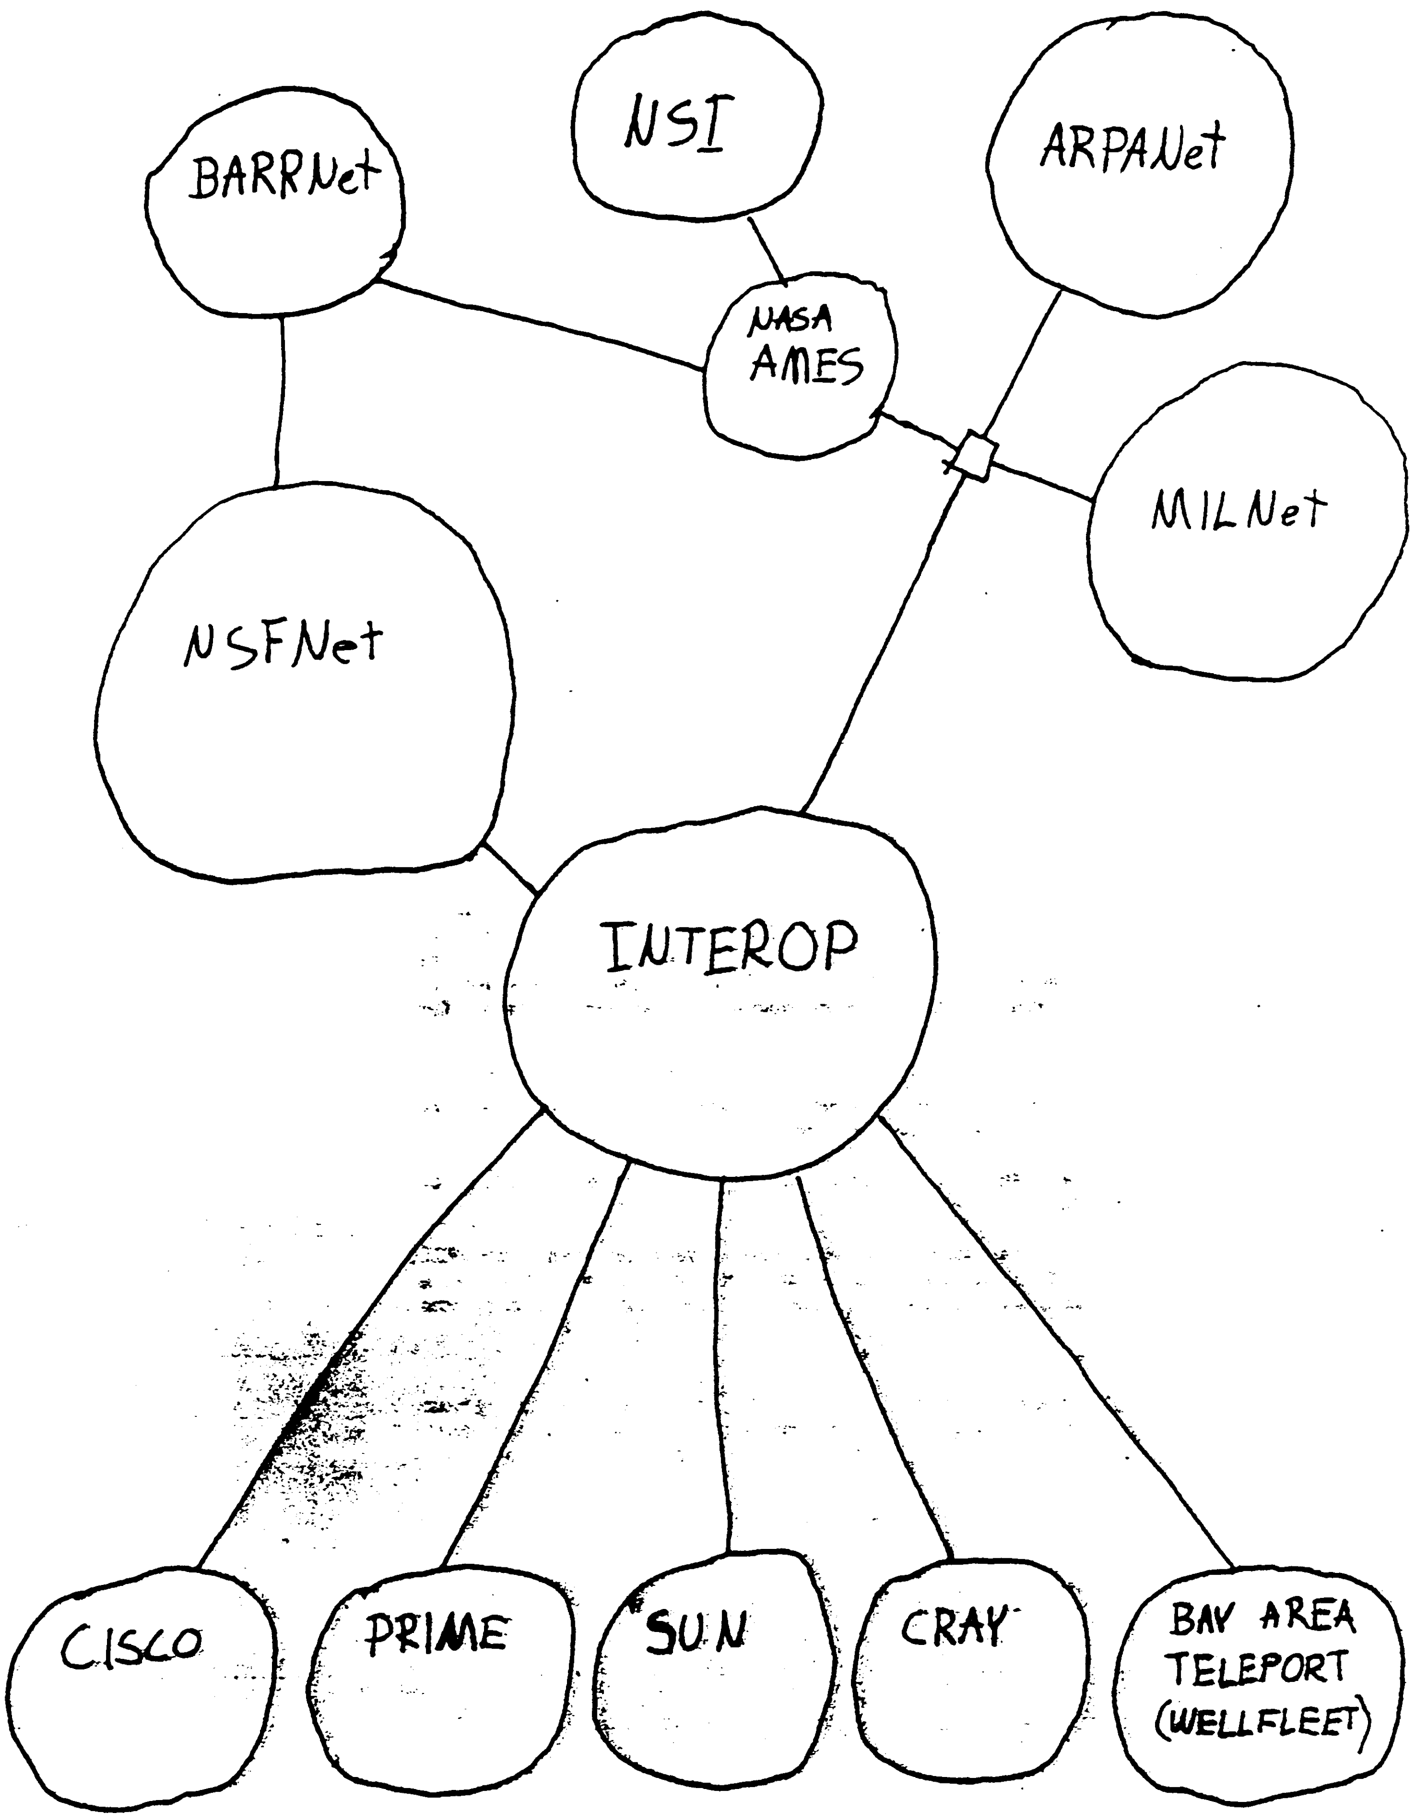
NSI
BARRNet
ARPANet
NASA AMES
MILNet
NSFNet
INTEROP
CISCO
PRIME
SUN
CRAY
BAY AREA TELEPORT (WELLFLEET)

External routing

- What we did
  - T-1 between core Proteon and AMES ARPANet/MILNet gateway
  - static routing over T-1
  - Proteon advertised RIP default
  - static routes to cisco, Prime, SUN Cray, Bay Area Teleport
  - Explicit RIP routes for NSFNet routes through IBM's NSS
- Lessons
  - cisco routers ignore RIP default

External routing - NSFNet

- What we did
    - NSFNet NSS in IBM booth
    - Secondary NSFNet path through BARRNet
    - IBM "subnet" was a class C net so EGP could treat it differently
    - PC/RT in IBM booth EGP peered with NSS and distributed RIP routes on the class C net
    - cisco core gateway also EGP peered with the NSS and distributed RIP routes on the class B net
    - Result: routing policy decisions by IBM and the NOC were independent of each other
    - NOC policy decision: always believe NSF routes (except for one afternoon when the NSFNet T-1 was flapping)

- Problems
  - We started out the show running old cisco code without NSFNet fixes to EGP
  - The NOC policy decision somewhat controversial...
  - Black holes occurred due to bad mixtures of static routes and firewalls in some of the regionals
- Lessons
  - Because of firewalls, it is dangerous to add a network to NSFNet without informing the regional networks.

Disappointments

- Network took one day too long to build
  - No time for interoperability testing
  - Network management not set up
  - No time for packet watching
  - Vendors pretty much left to sink or swim on their own
  - Network would have been more solid if it had run for a day before the show

Things I was particularly happy about

- It worked well enough...
- We got a tremendous amount of help from the Internet community

The reasons it all worked

Rick Boivie
Len Bosack
David Bridgham
Eric Brunner
Jeff Burgan
Myu Campbell
Mario Castro
Shelly DeVries
Steve Knowles
Susan Hares
Alex Latzko
Sandy Lerner
Milo Medin
Robert Michaels
Paul Mockapetris
Mike Moesler
Vince Raya
Sue Romano
Greg Satz
Mick Scully
Jim Shimoto
Mike St. Johns
James VanBokken
John Veizades

People who contributed to this talk

Peter DeVries
Milo Medin

Internet Protocols ("TCP/IP") for Amateur Radio

Phil Karn
Bell Communications Research

Internet Protocols ("TCP/IP") for Amateur Radio

Phil Karn, KA9Q

Amateur Packet Radio Station

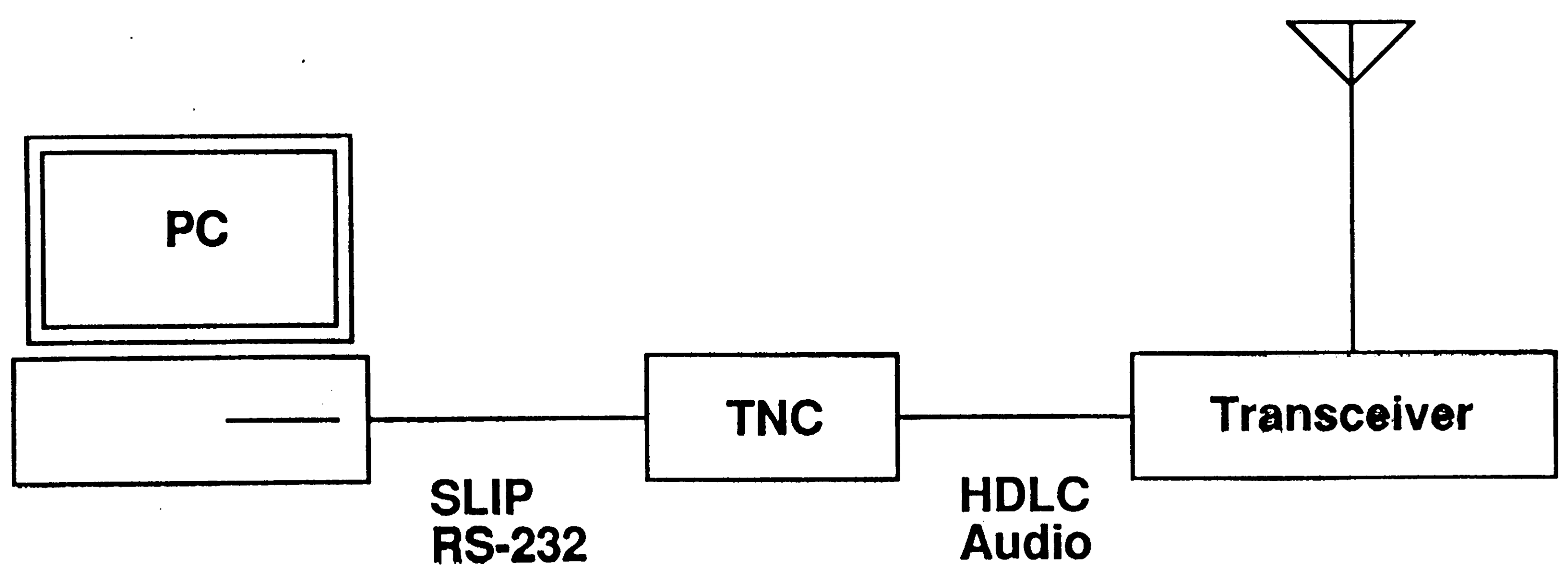

TITLE ____________ SEQUENCE NO. ____________

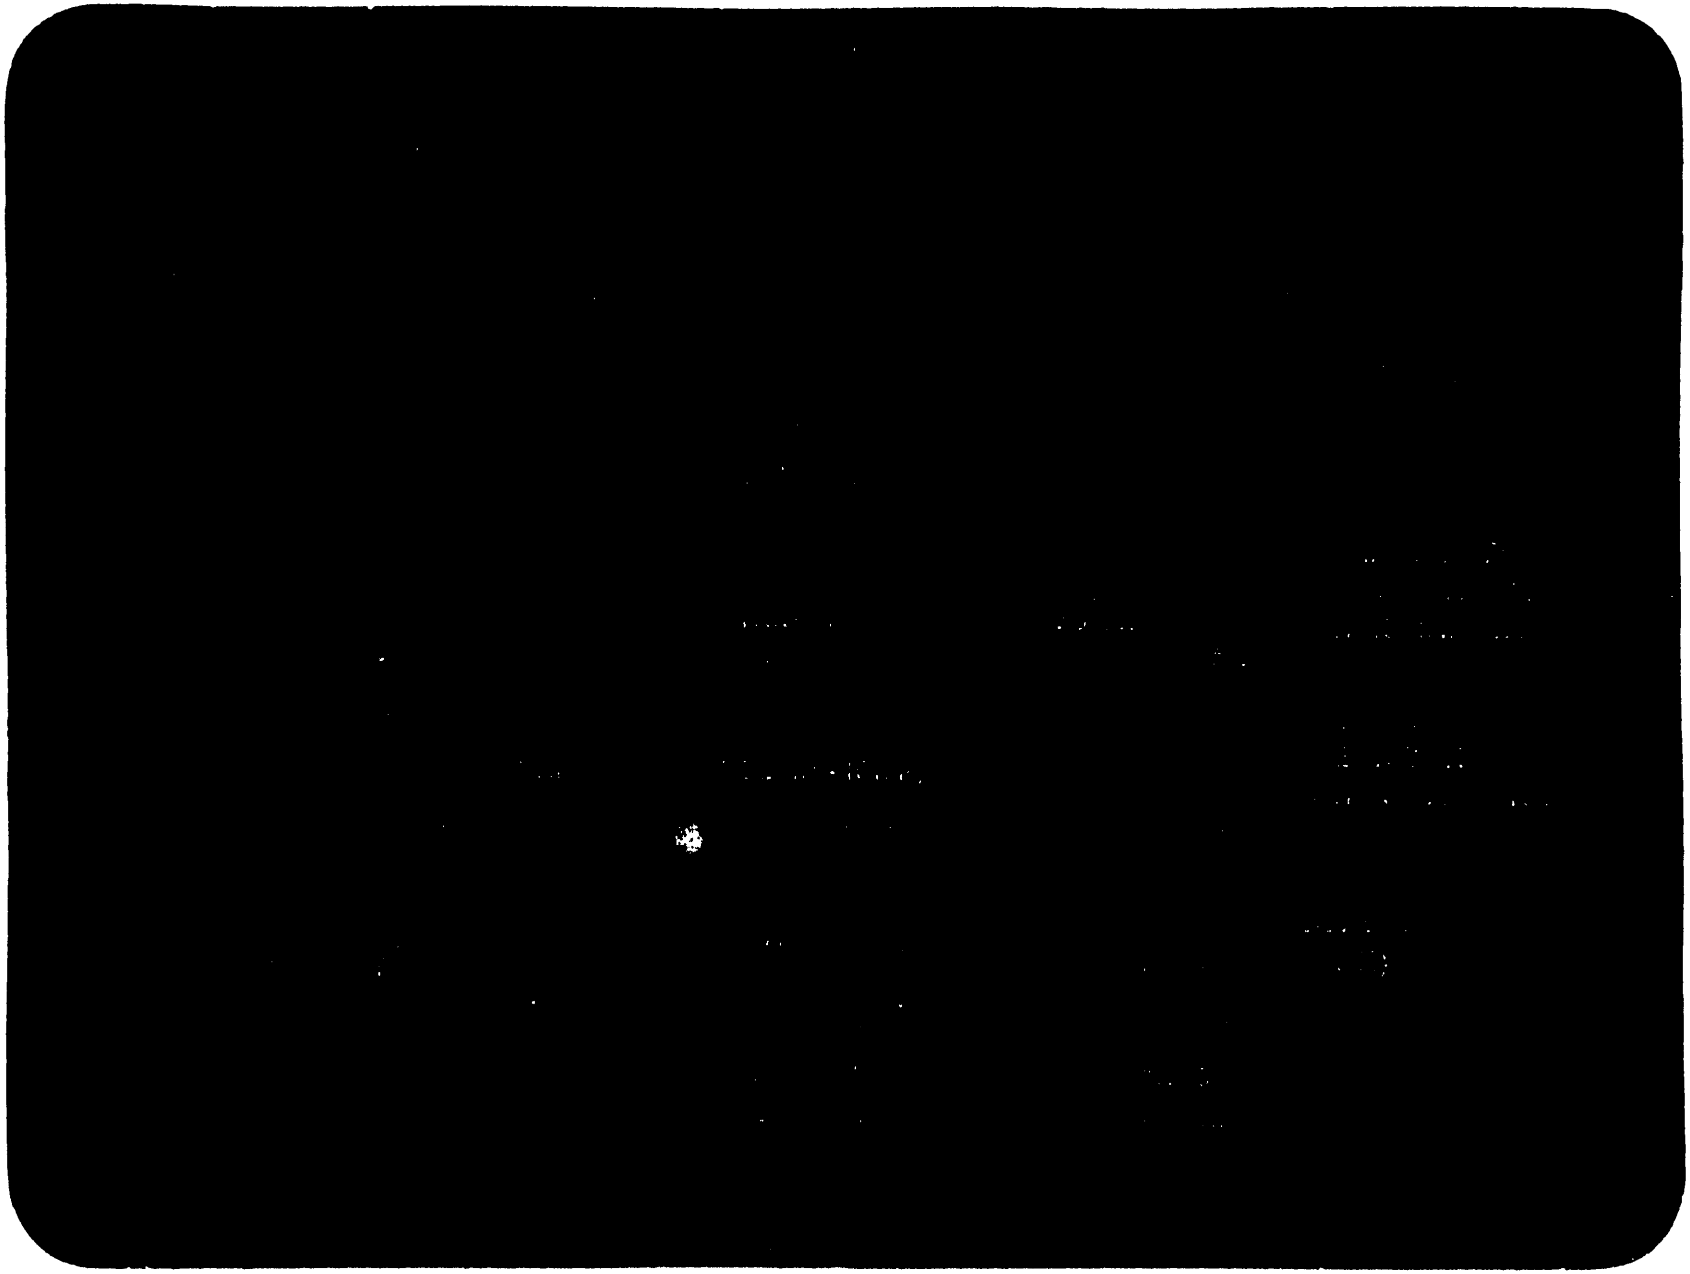

(6¾" x 9" APERTURE)

Bell
Communications
Research

SCO409

P. Karn

Packet Radio Frame

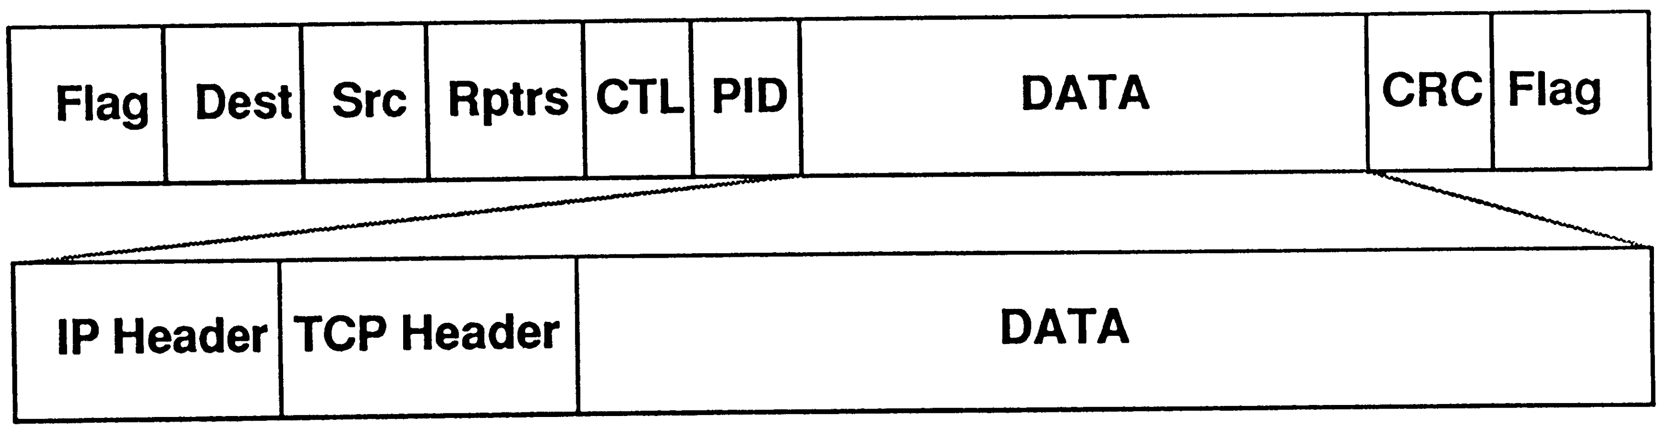

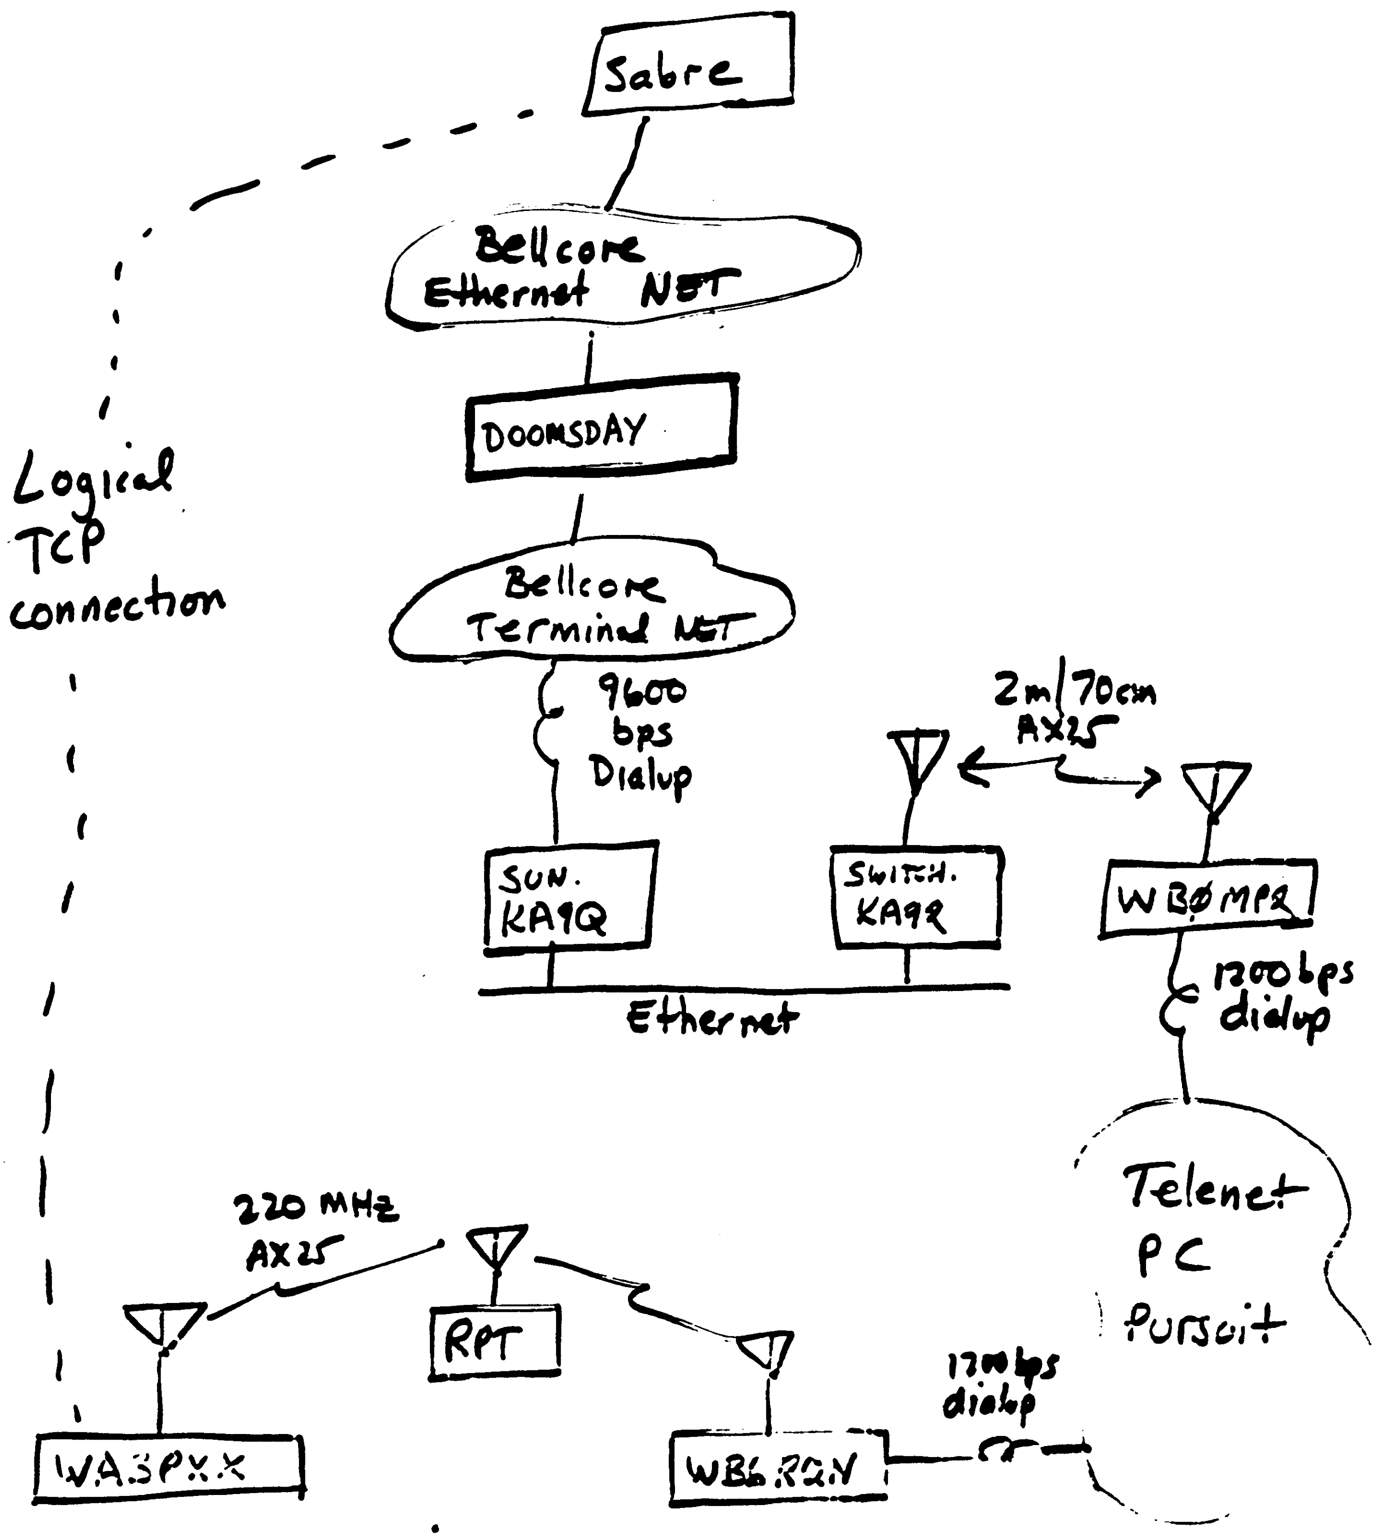
Sabre
Bellcore Ethernet NET
DOOMSDAY
Bellcore Terminal NET
Logical TCP connection
9600 bps Dialup
2m/70cm AX25
SUN. KA9Q
SWITCH. KA9Q
WB0MPQ
Ethernet
1200 bps dialup
Telenet PC Pursuit
220 MHz AX25
RPT
1200 bps dialup
WA3PXX
WB6RQN

Packet Modem Developments

- 56 Kbps MSK (WA4DSY)
  3 board kit from GRAPES (Atlanta)
  28 MHz IF to linear transverter

- 9600 bps FSK (G3RUH)
- 9600 bps FSK (K9NG)
  1 board kit
  Connects to FM voice radio
  Internal connections required

- 4800 bps (HAPN)

- 1200 bps PSK (TAPR/JAMSAT)
- 1200 bps PSK (G3RUH)
  1 board kit
  Connects to SSB/FM radio (FO-12)

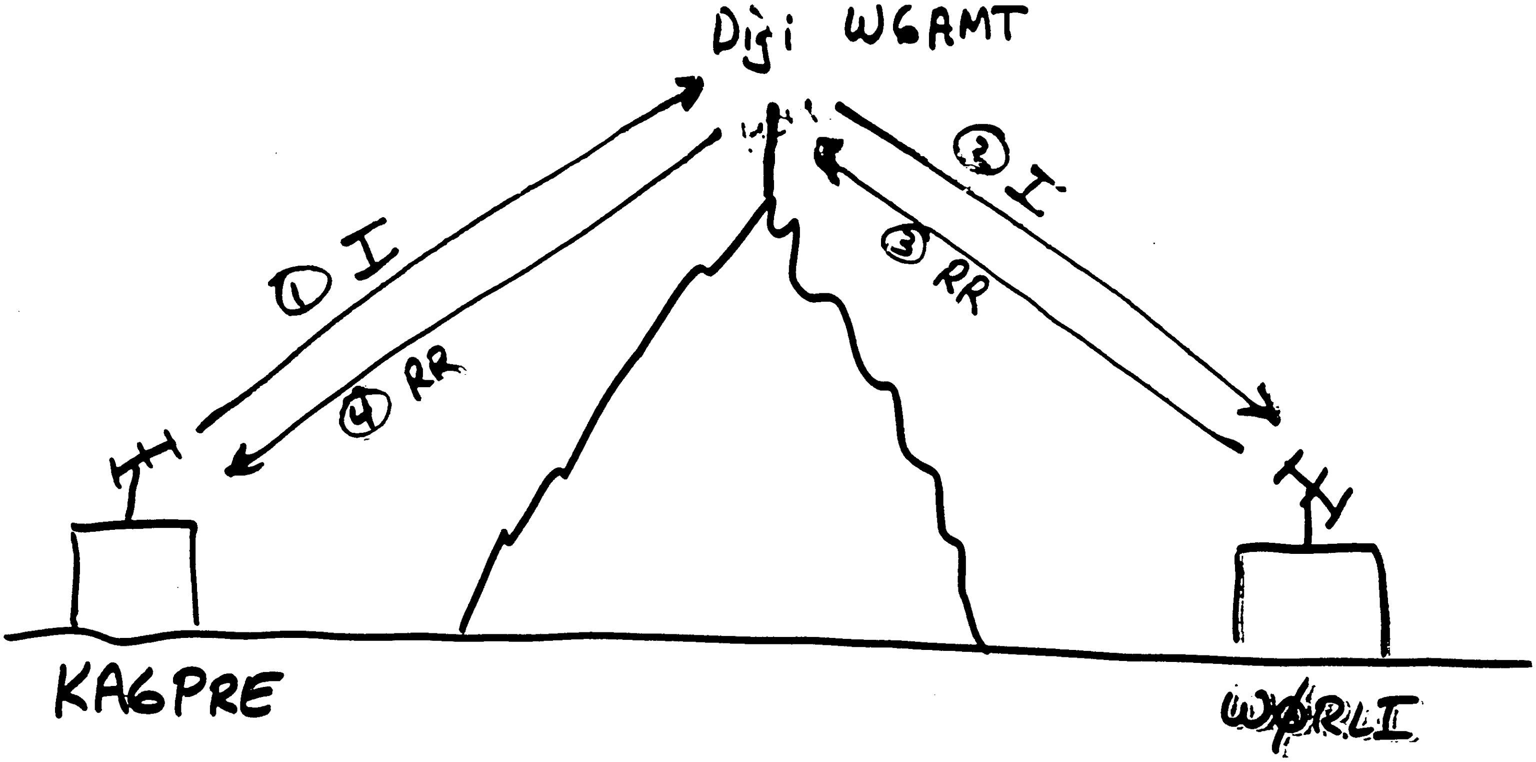
Digi W6AMT
1 I
4 RR
2 I
3 RR
KA6PRE
WØRLI

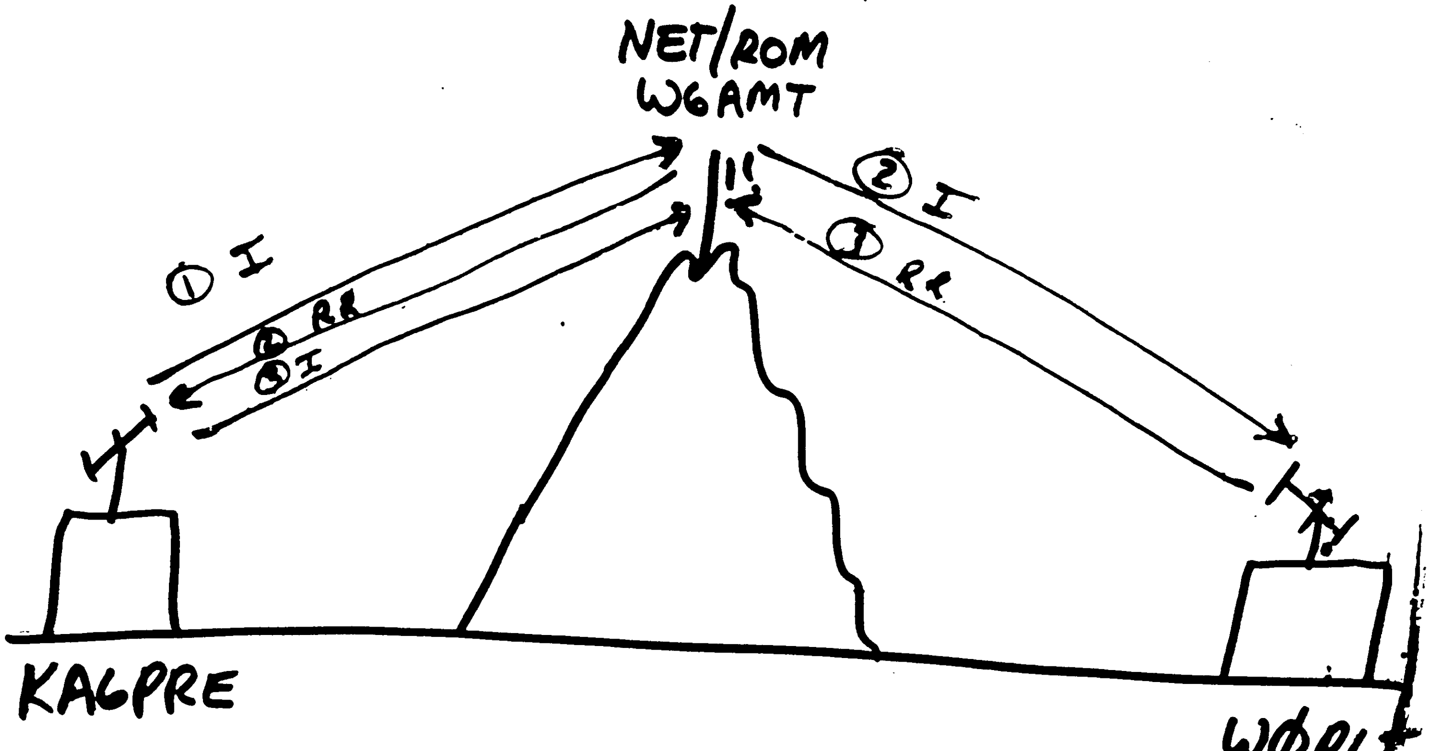
NET/ROM
W6AMT
1 I
RR
I
2 I
3 RR
KA6PRE
WØRLI

Problems

- Tinygrams
- Poor round trip timing algorithms
- Protocol violations, esp. on retransmit
- Short giveup timers
- Telnet echo
- Keepalives
- Fragmentation

Fixes

- **Nagle tinygram avoidance algorithm**
- **Jacobson, Karn RTT algorithms**
  Believe the numbers you get - no arbitrary limits!
- **Eliminate TCP giveup timers**
  Application, *not* TCP, should abort
- **Eliminate TCP keepalives**

Solving the "Retransmission Ambiguity" Problem

- Ignore RTTs of retransmitted packets
- Clamp backoff after retransmission and increase until valid RTT obtained

Partial Fixes

- Telnet local echo
- Link level retransmission
  - Not a replacement for "doing it right" at the physical layer
- Link level (transparent) fragmentation

High Performance TCP Over An Ethernet

Van Jacobson
Lawrence Berkeley Labs

Costs (in time) to Send a Packet

"**Fixed**" (per-packet):

Examples:

- Media acquisition time
- Packet headers & trailers
- Protocol processing
- Device & interrupt service

"**Variable**" (per-byte):

Examples:

- Bit time on wire
- Copy to/from user space
- Checksum data

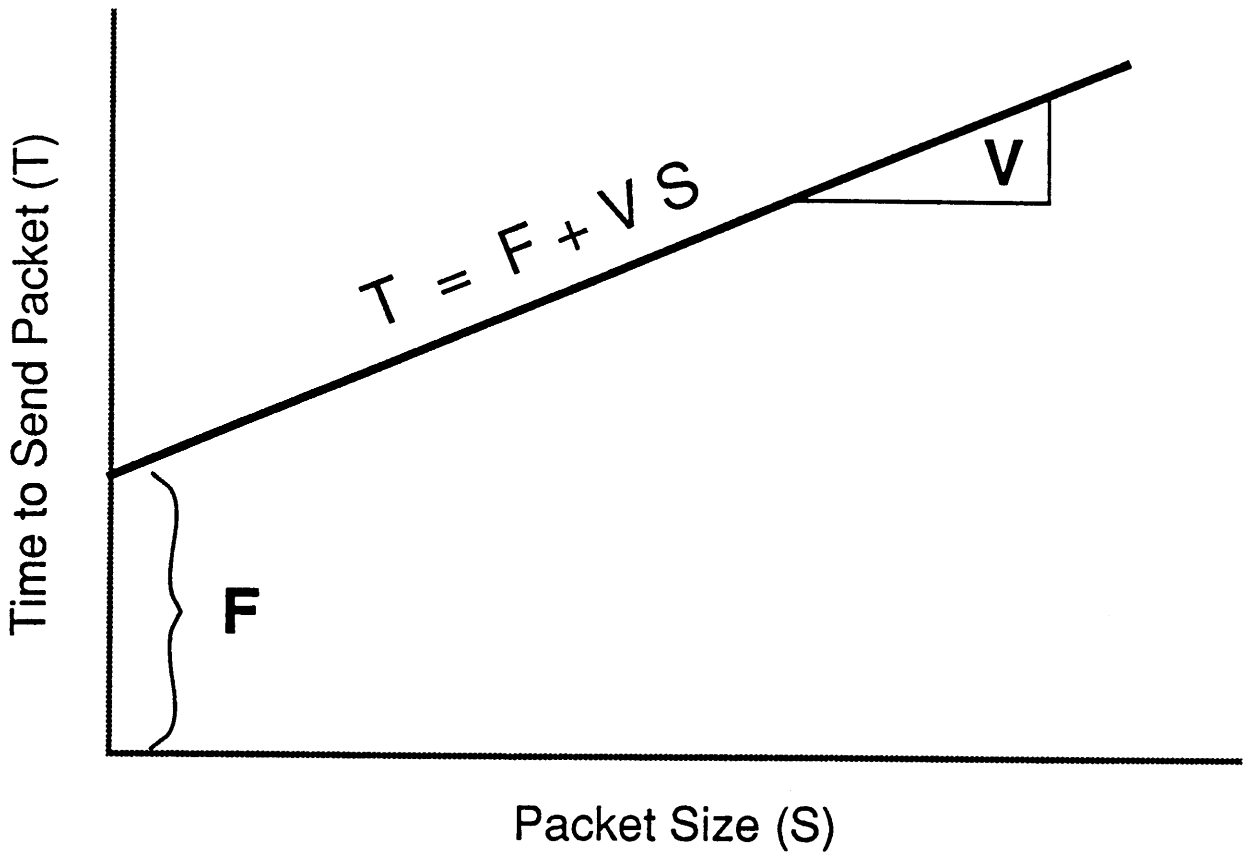

$$\text{Bandwidth} \equiv \frac{\text{size}}{\text{time}} \Longrightarrow B = \frac{S}{F + VS}$$

$$\boxed{B = \frac{1}{V + F/S}}$$

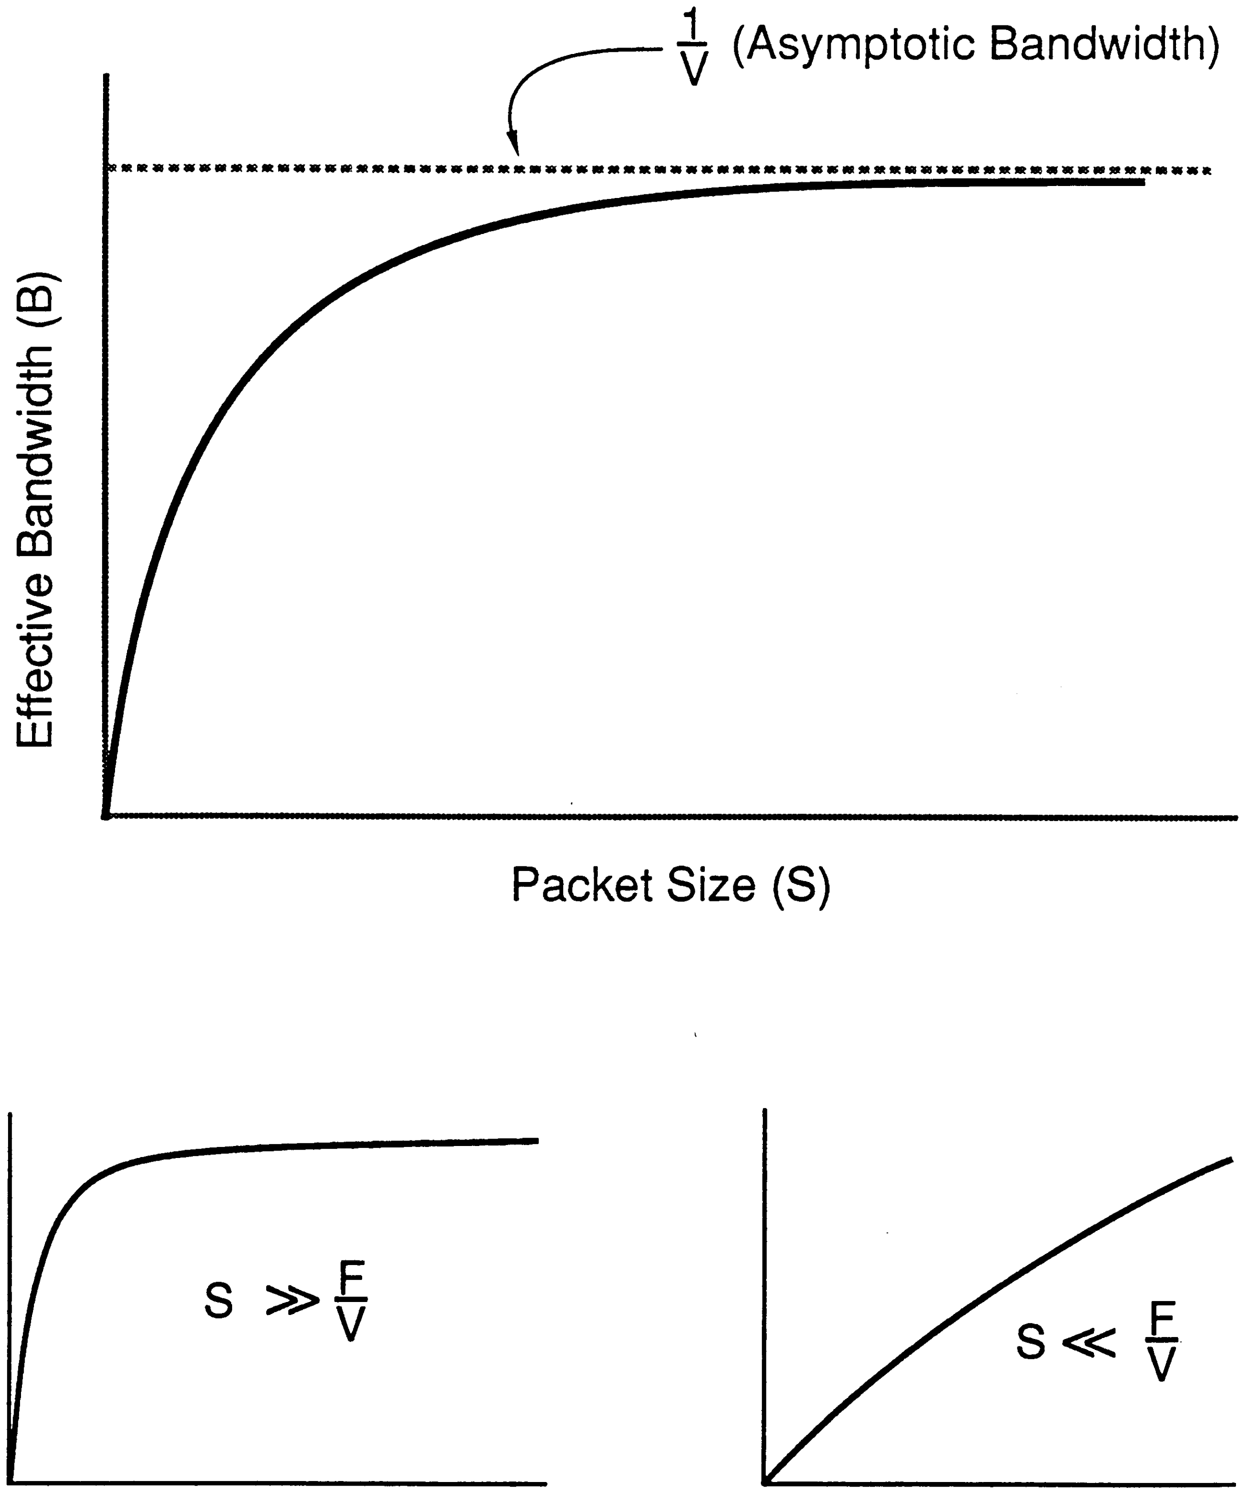

$\frac{1}{V}$ (Asymptotic Bandwidth)
Effective Bandwidth (B)
Packet Size (S)
$S \gg \frac{F}{V}$
$S \ll \frac{F}{V}$

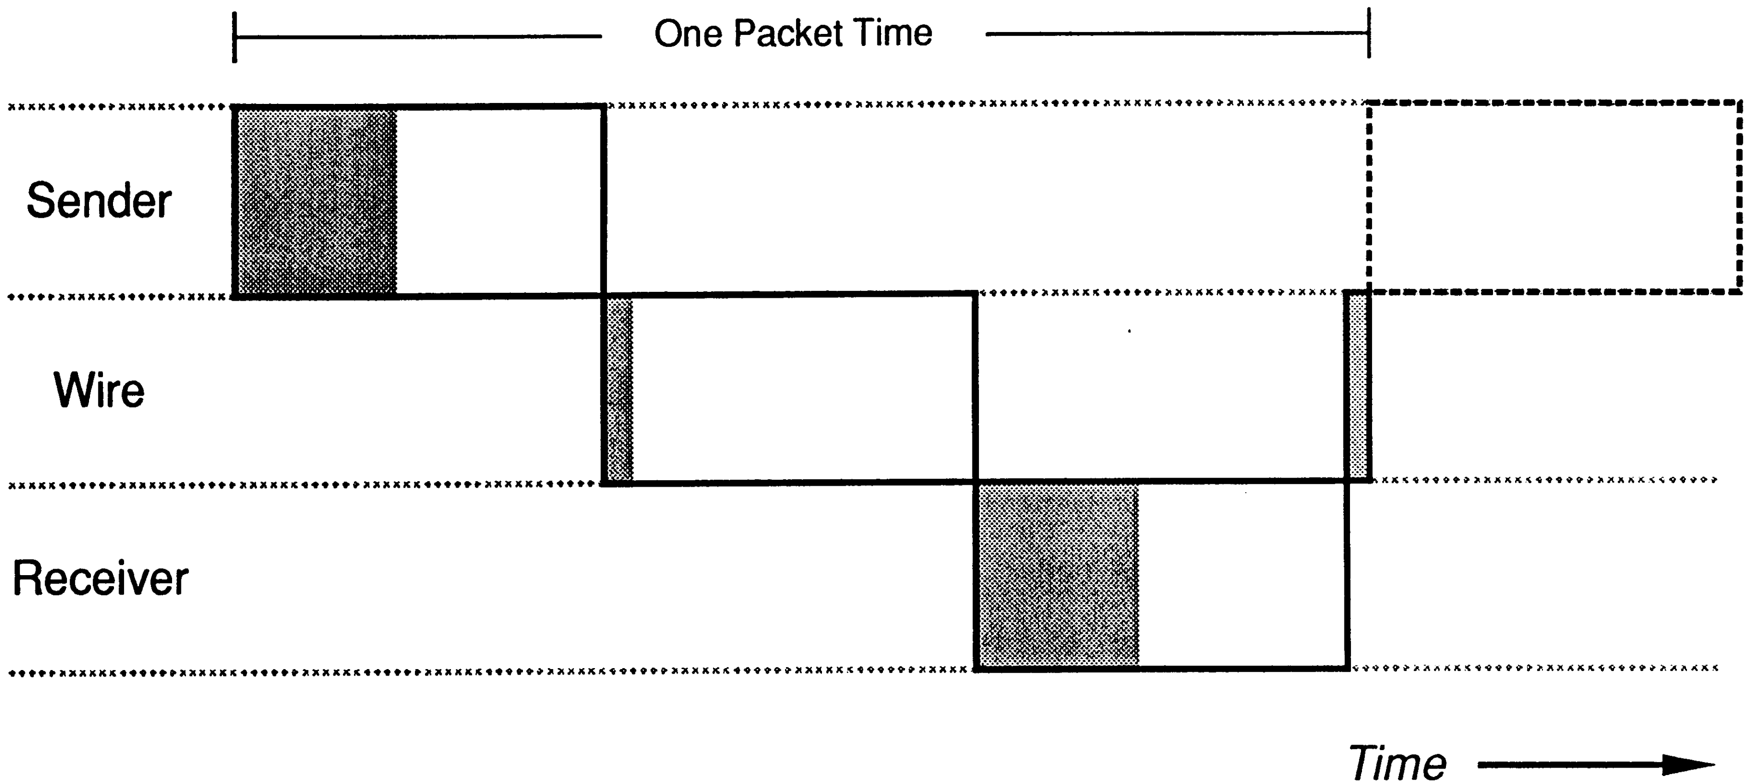

Sequence of Events: Sending One Packet

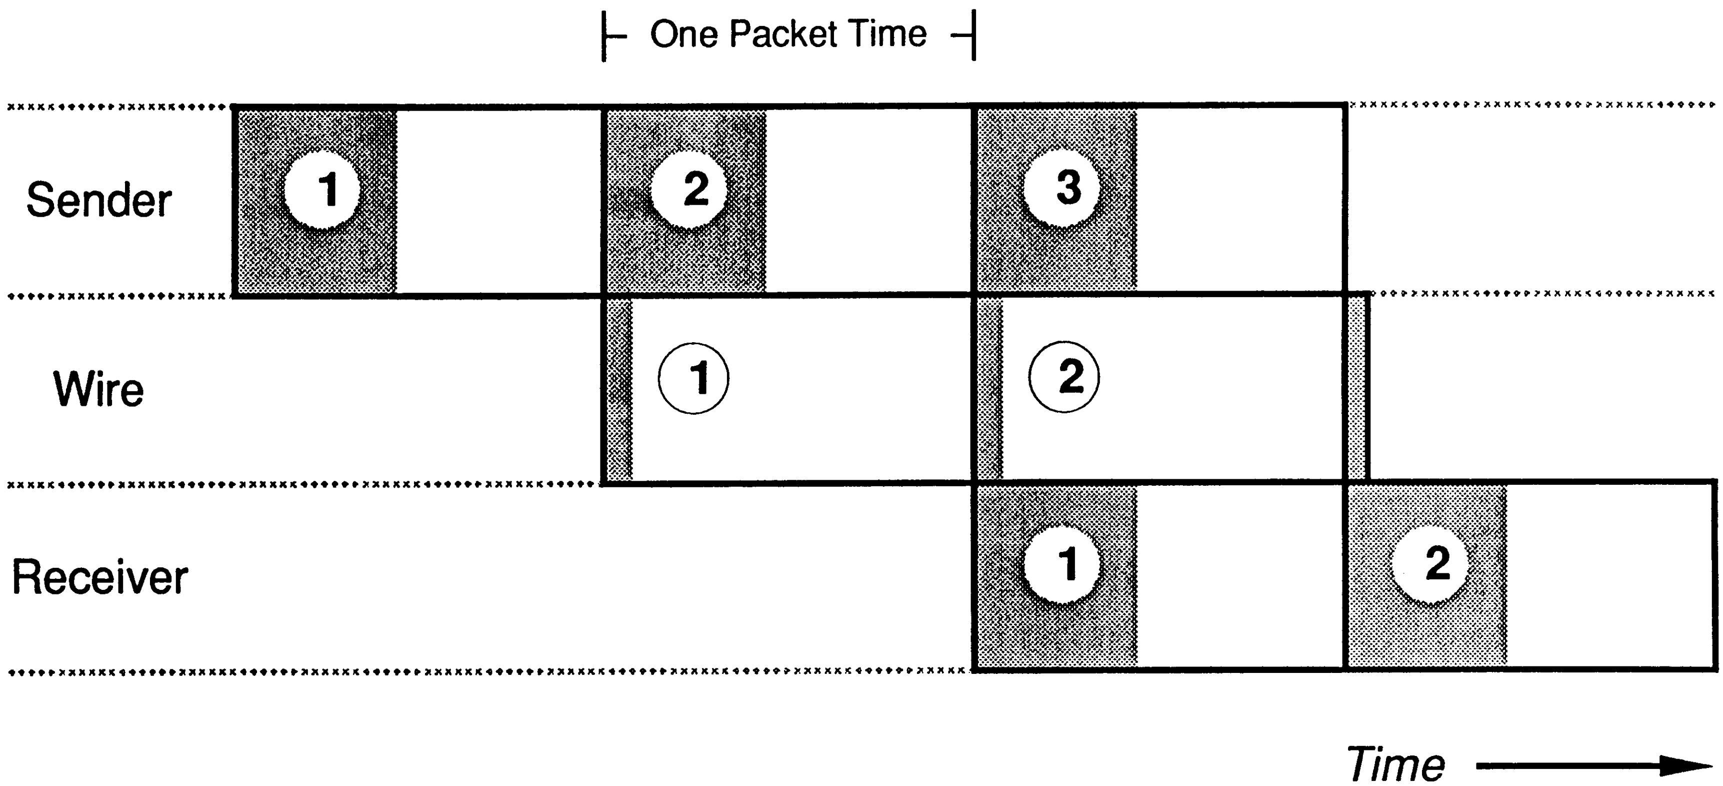

Sequence of Events: Pipelining

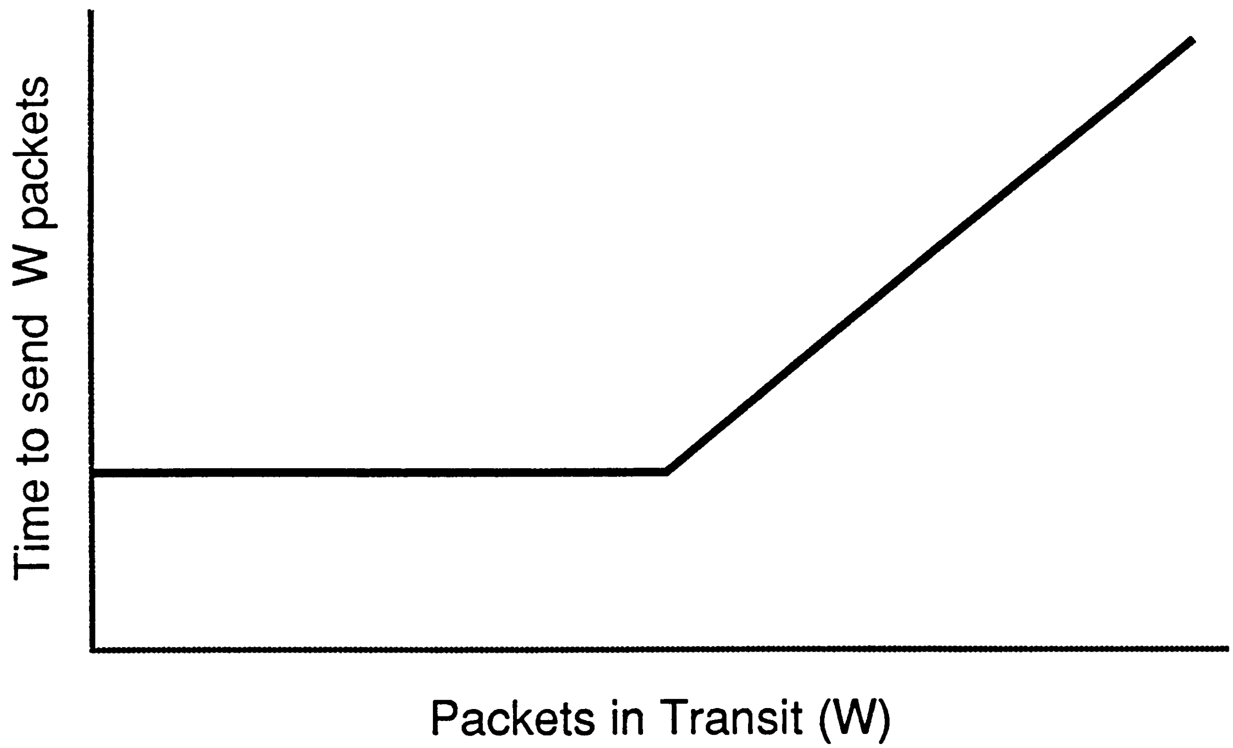
Time to send W packets
Packets in Transit (W)

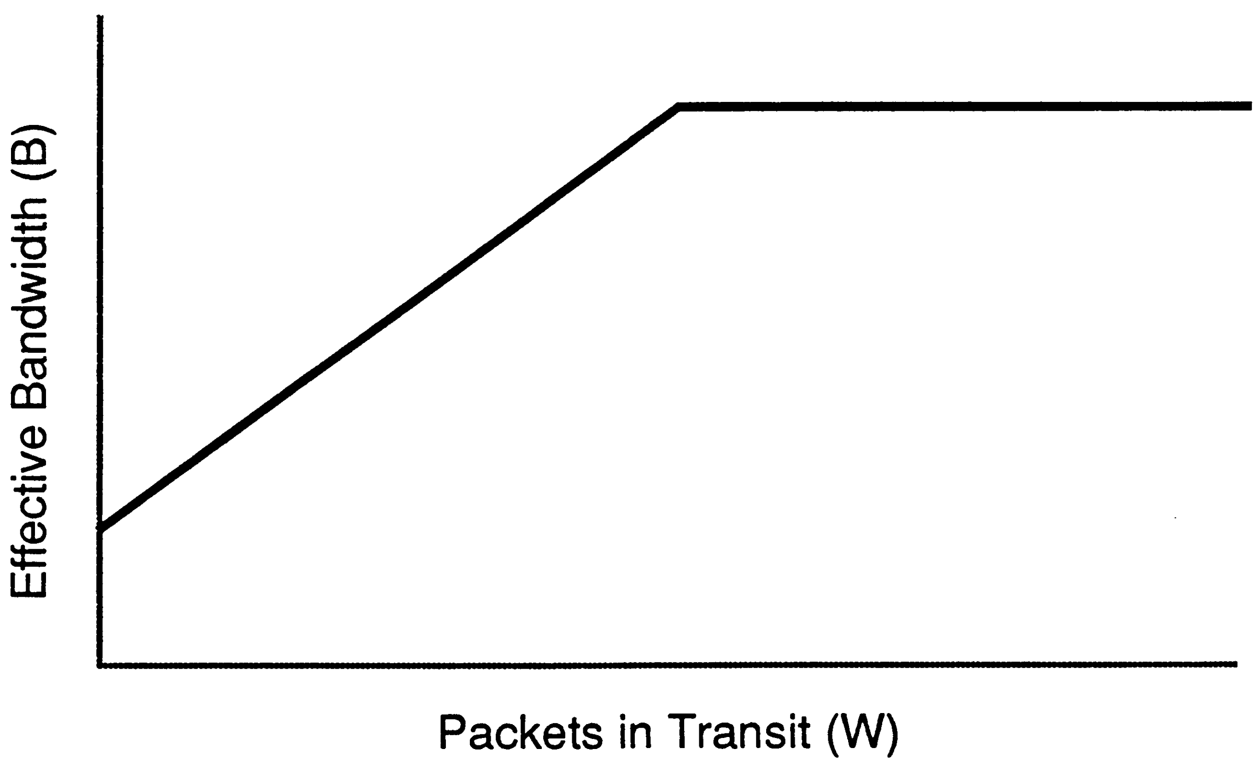
Effective Bandwidth (B)
Packets in Transit (W)

Ethernet Costs

- *Fixed costs:*

  24 byte IPG / Sync / CRC

  14 byte Ether header

  20 byte IP header

  20 byte TCP header

  ---

  78 bytes (= 62 us)

  ×1.5 (one ack per 2 data) **93 us / packet**

- *Variable Cost:*

  10 Mbps **0.8 us / byte**

F / V = 116 ⇒ want at least 1160 byte packets.

Max packet length is 1538 bytes.
1538 – 78 = 1460 bytes user data

max variable cost = 1460 × 0.8 = 1168 us

total cost for max length packet = **1261 us**

max efficiency = 1460 / 1577 = 93%

Max User Data Throughput of Ethernet Running TCP/IP

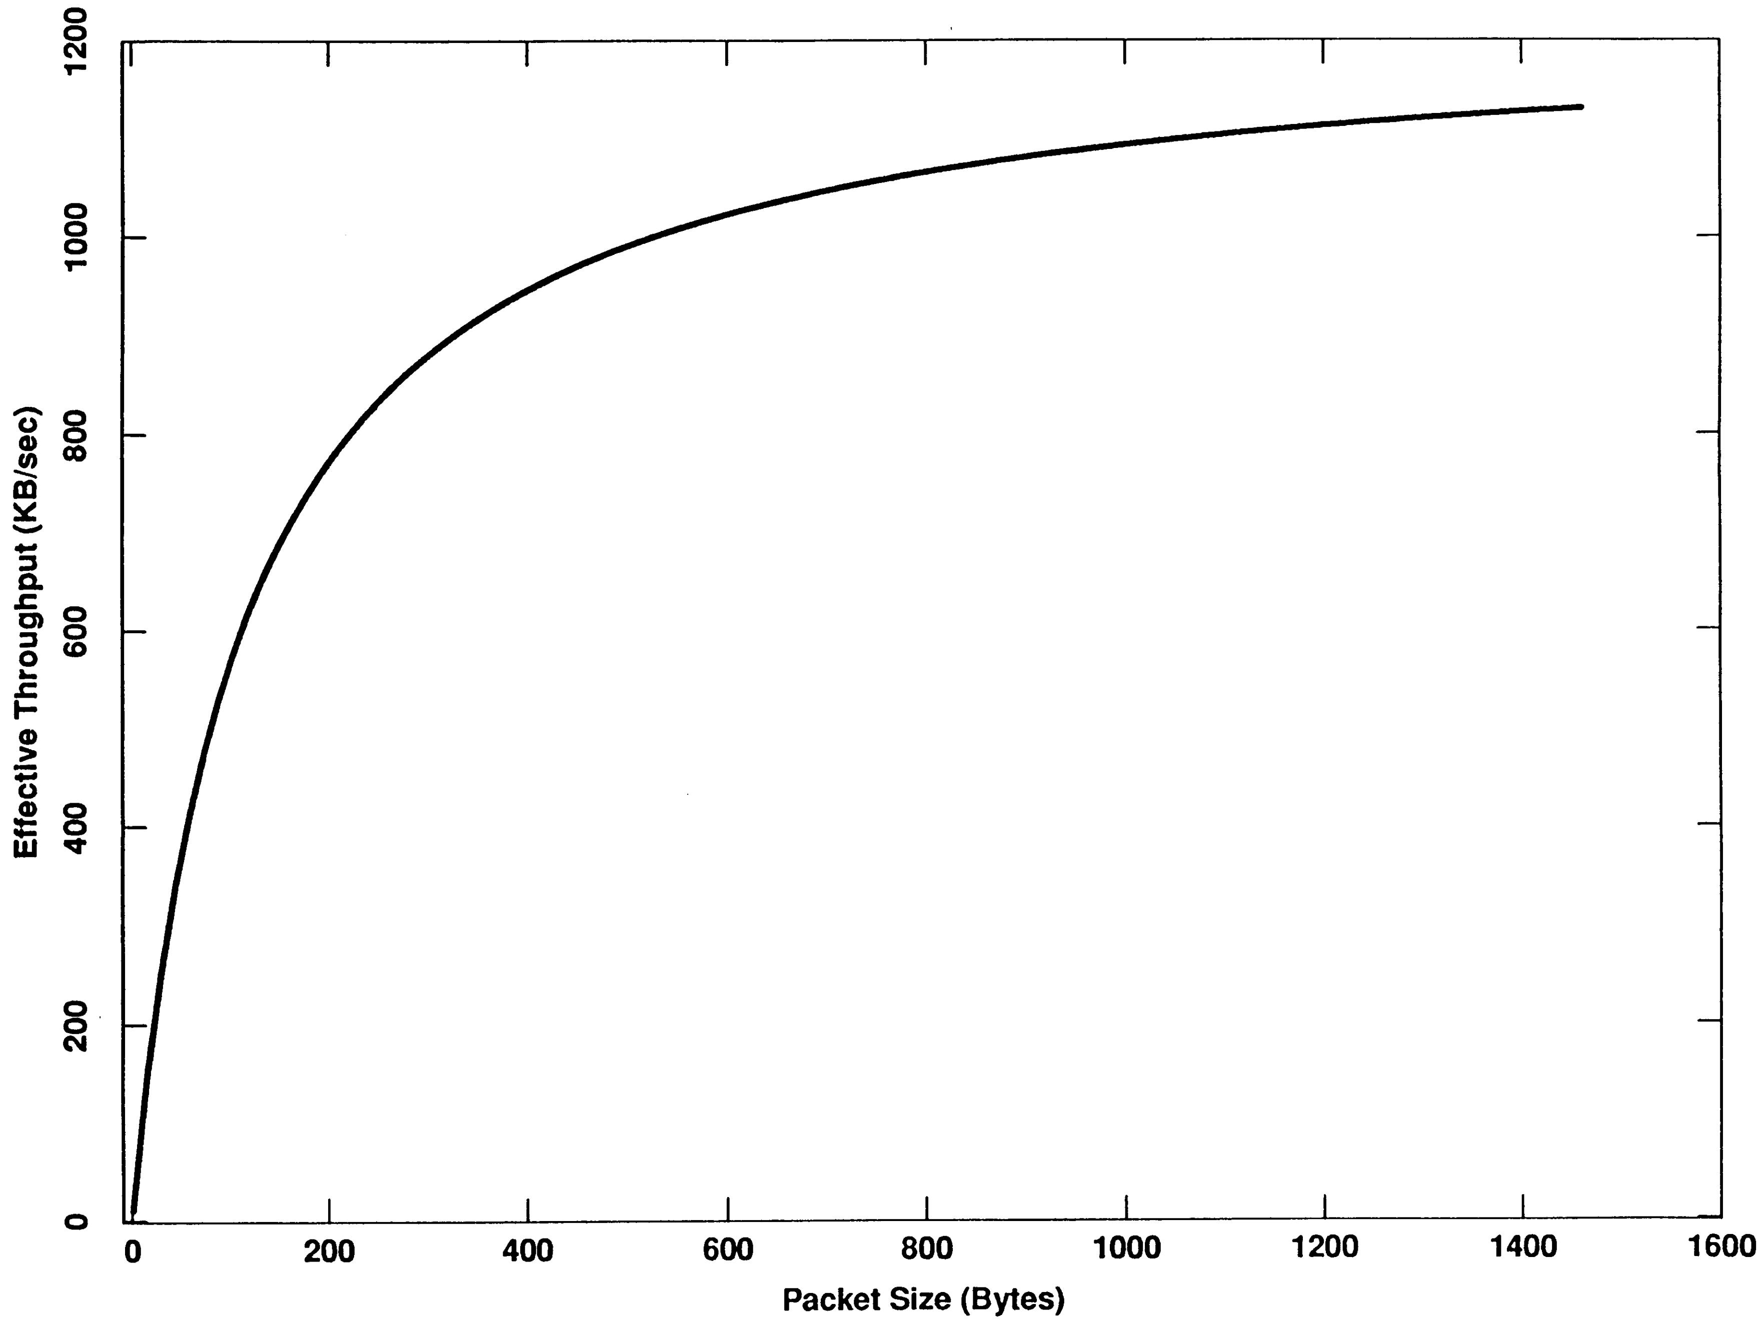

CPU / System Costs

(for 20MHz 68020 running 4BSD Unix)

- *Variable Costs (for 1460 byte packet):*

  (limitting bandwidth is memory @ 130 ns/byte)

User — System copy	200 us	
TCP Checksum	185 us	771 us
LANCE bus use	386 us	

- *Fixed Costs:*

LANCE (Ethernet) driver	100 us	
TCP / IP / ARP protocols	100 us	440 us
other OS functions (syscall, sleep, wakeup, 3 interrupts)	240 us	
Idle		200 us
		1411 us

Sept. 88 TCP Throughput Tests

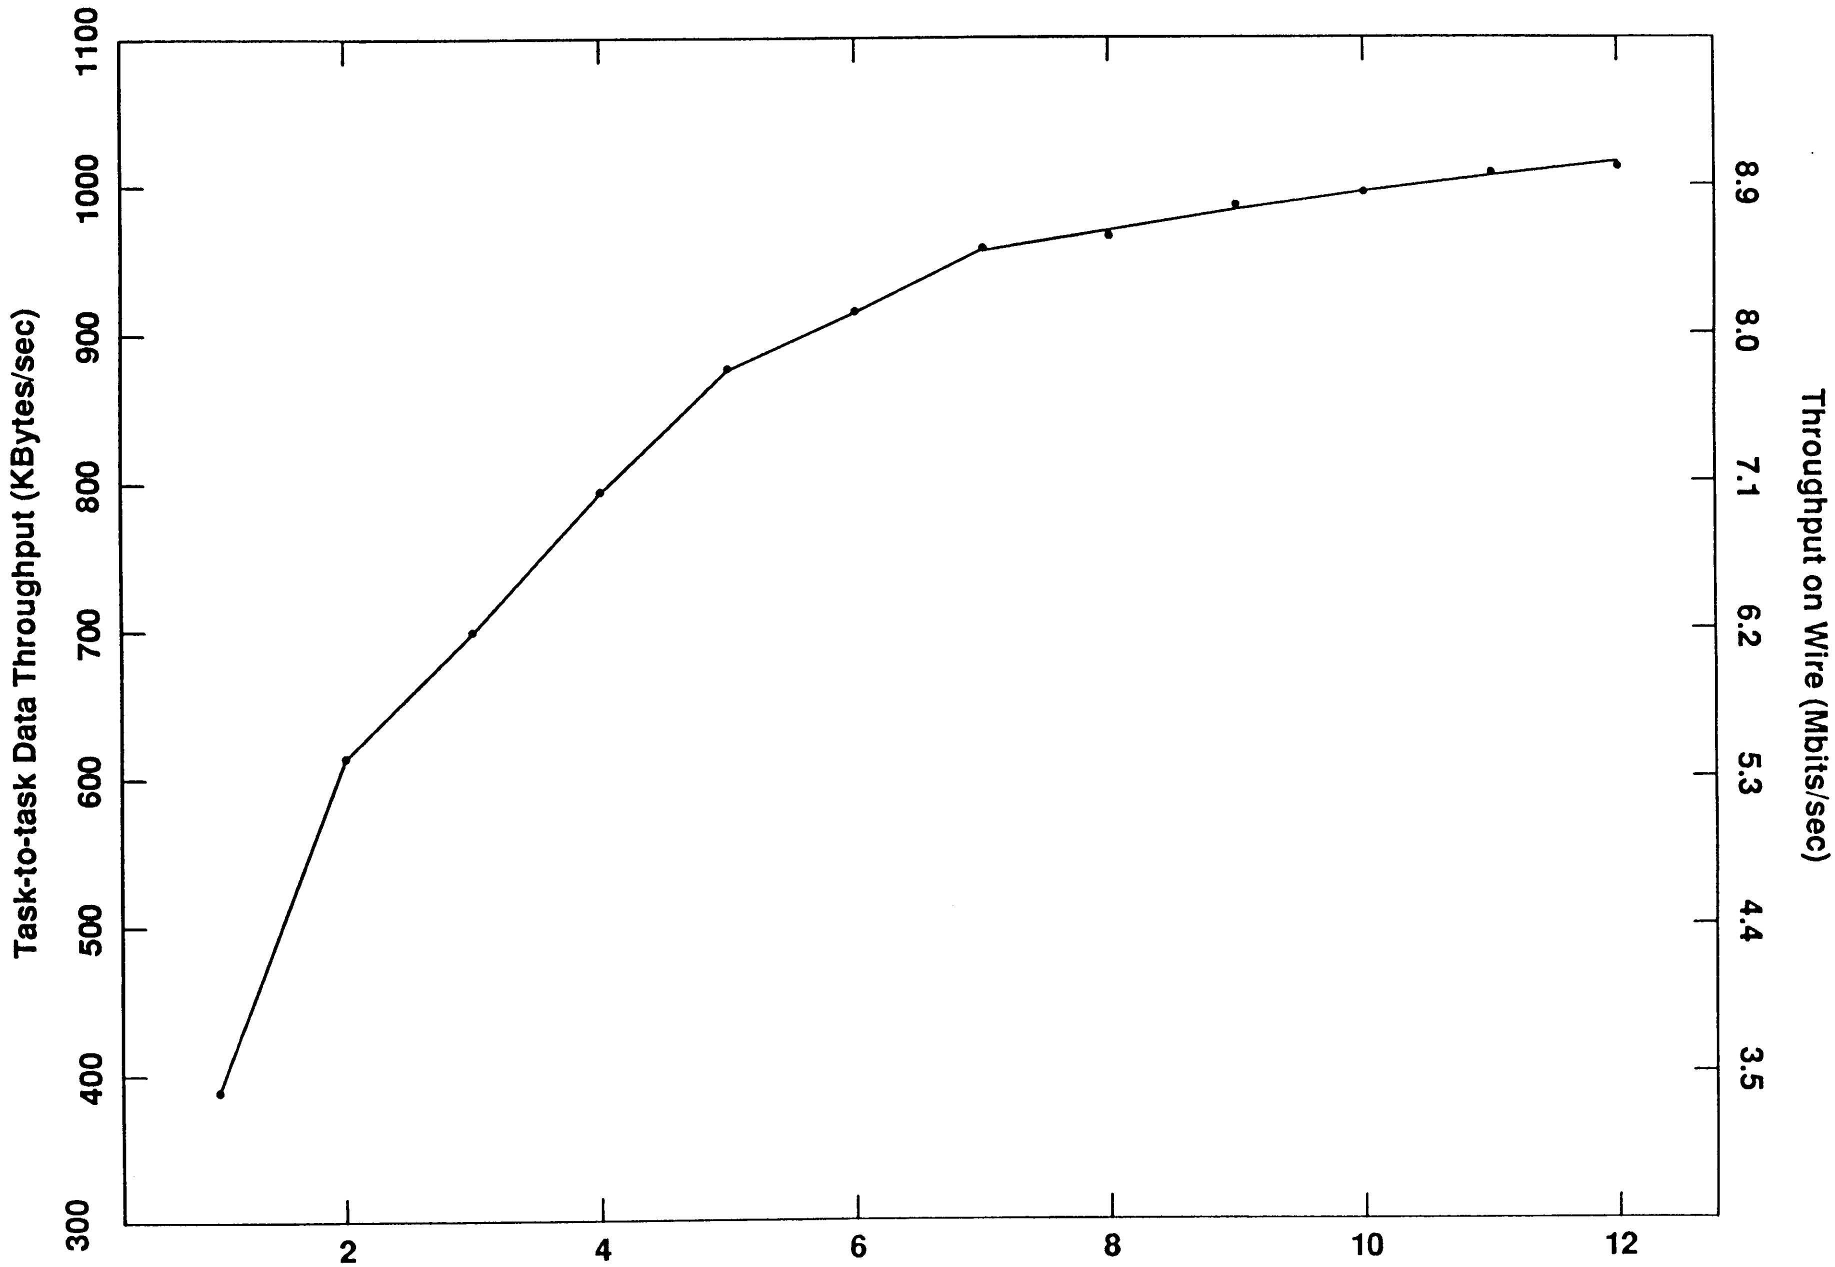

Receiver Buffer Size (in packets)
20MHz 68020 w/AMD LANCE

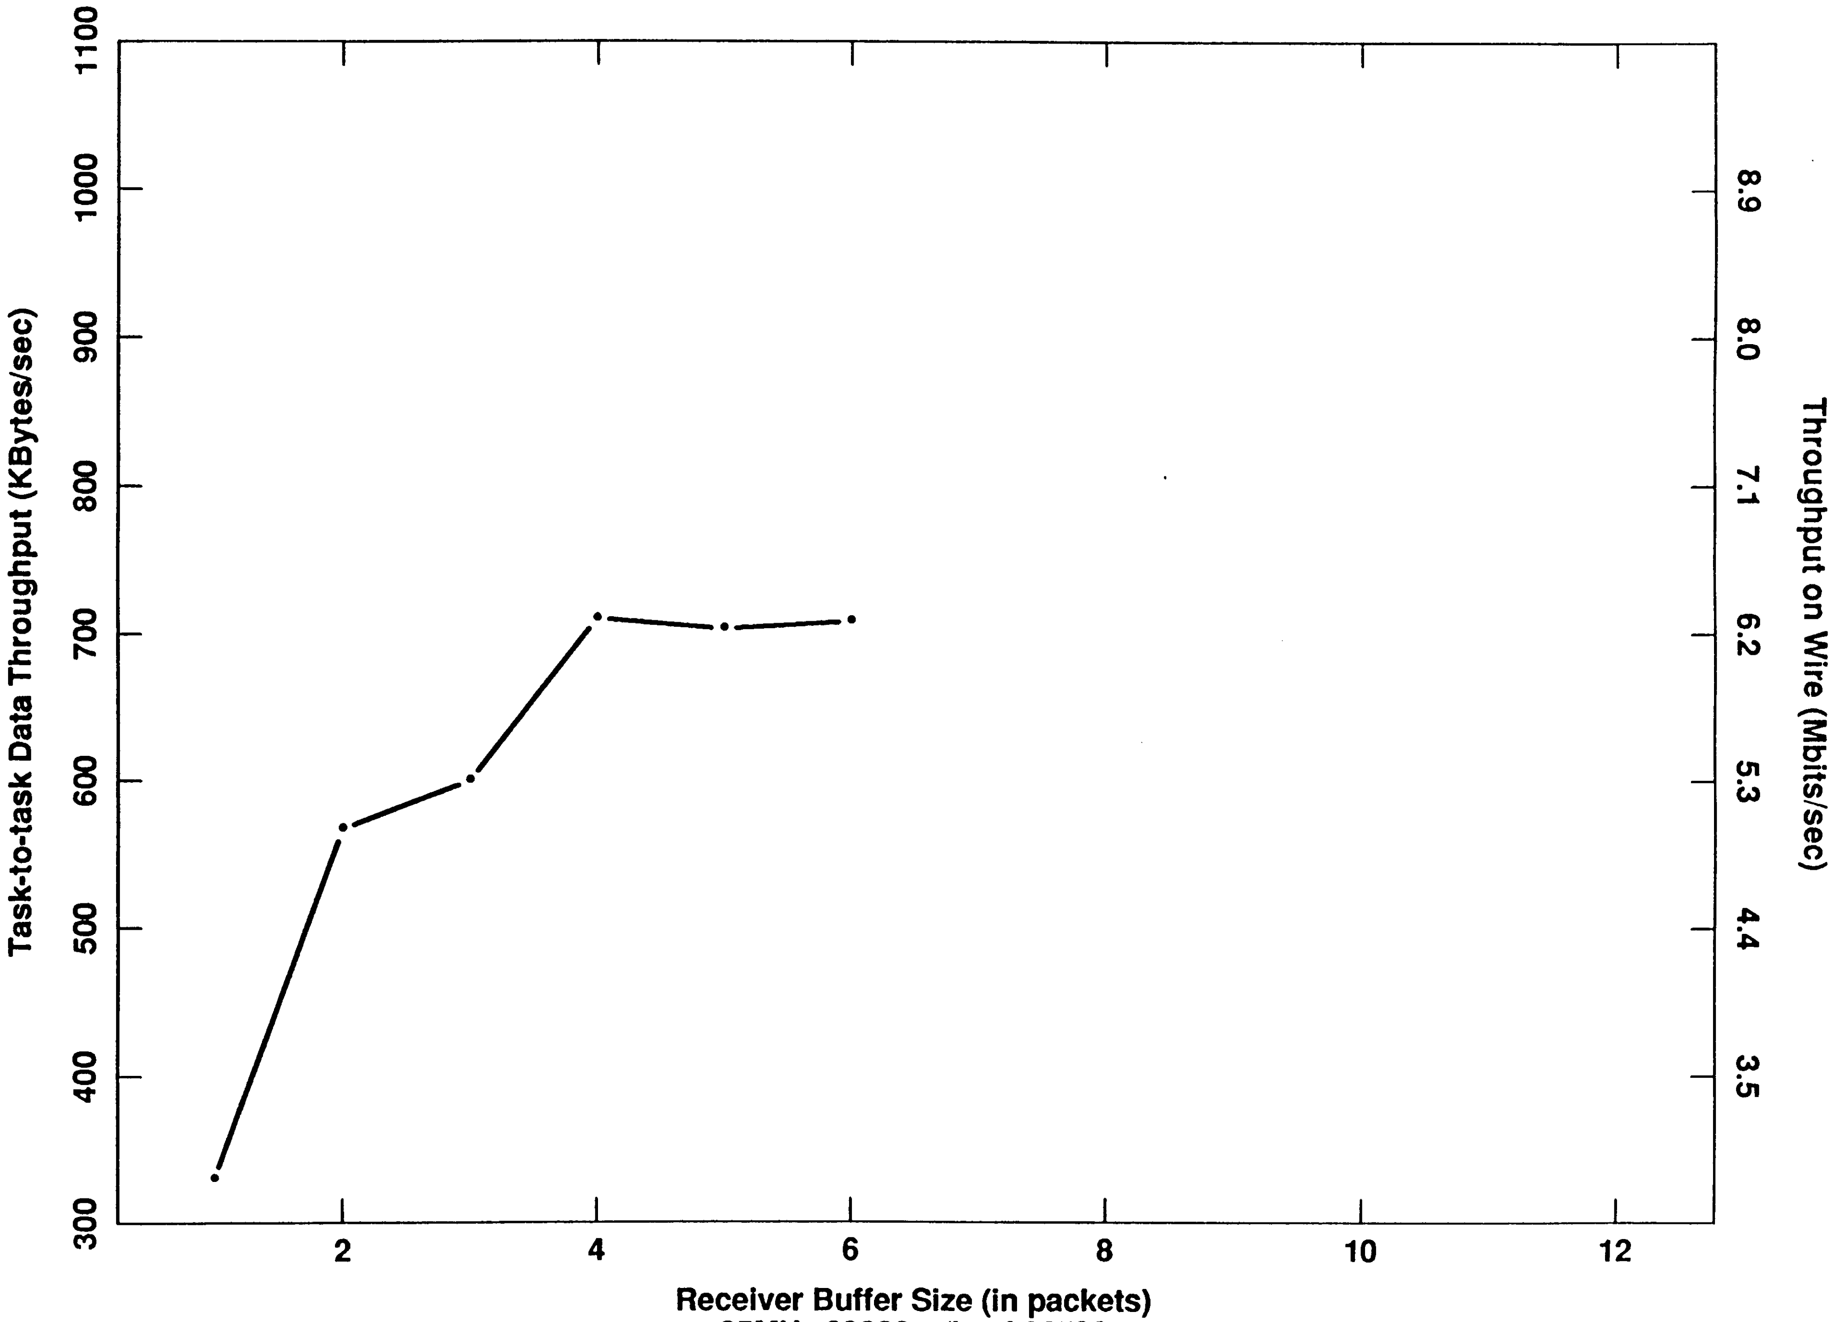

25MHz 68020 w/Intel 82586

Aug 88: New Network Code Xput Tests

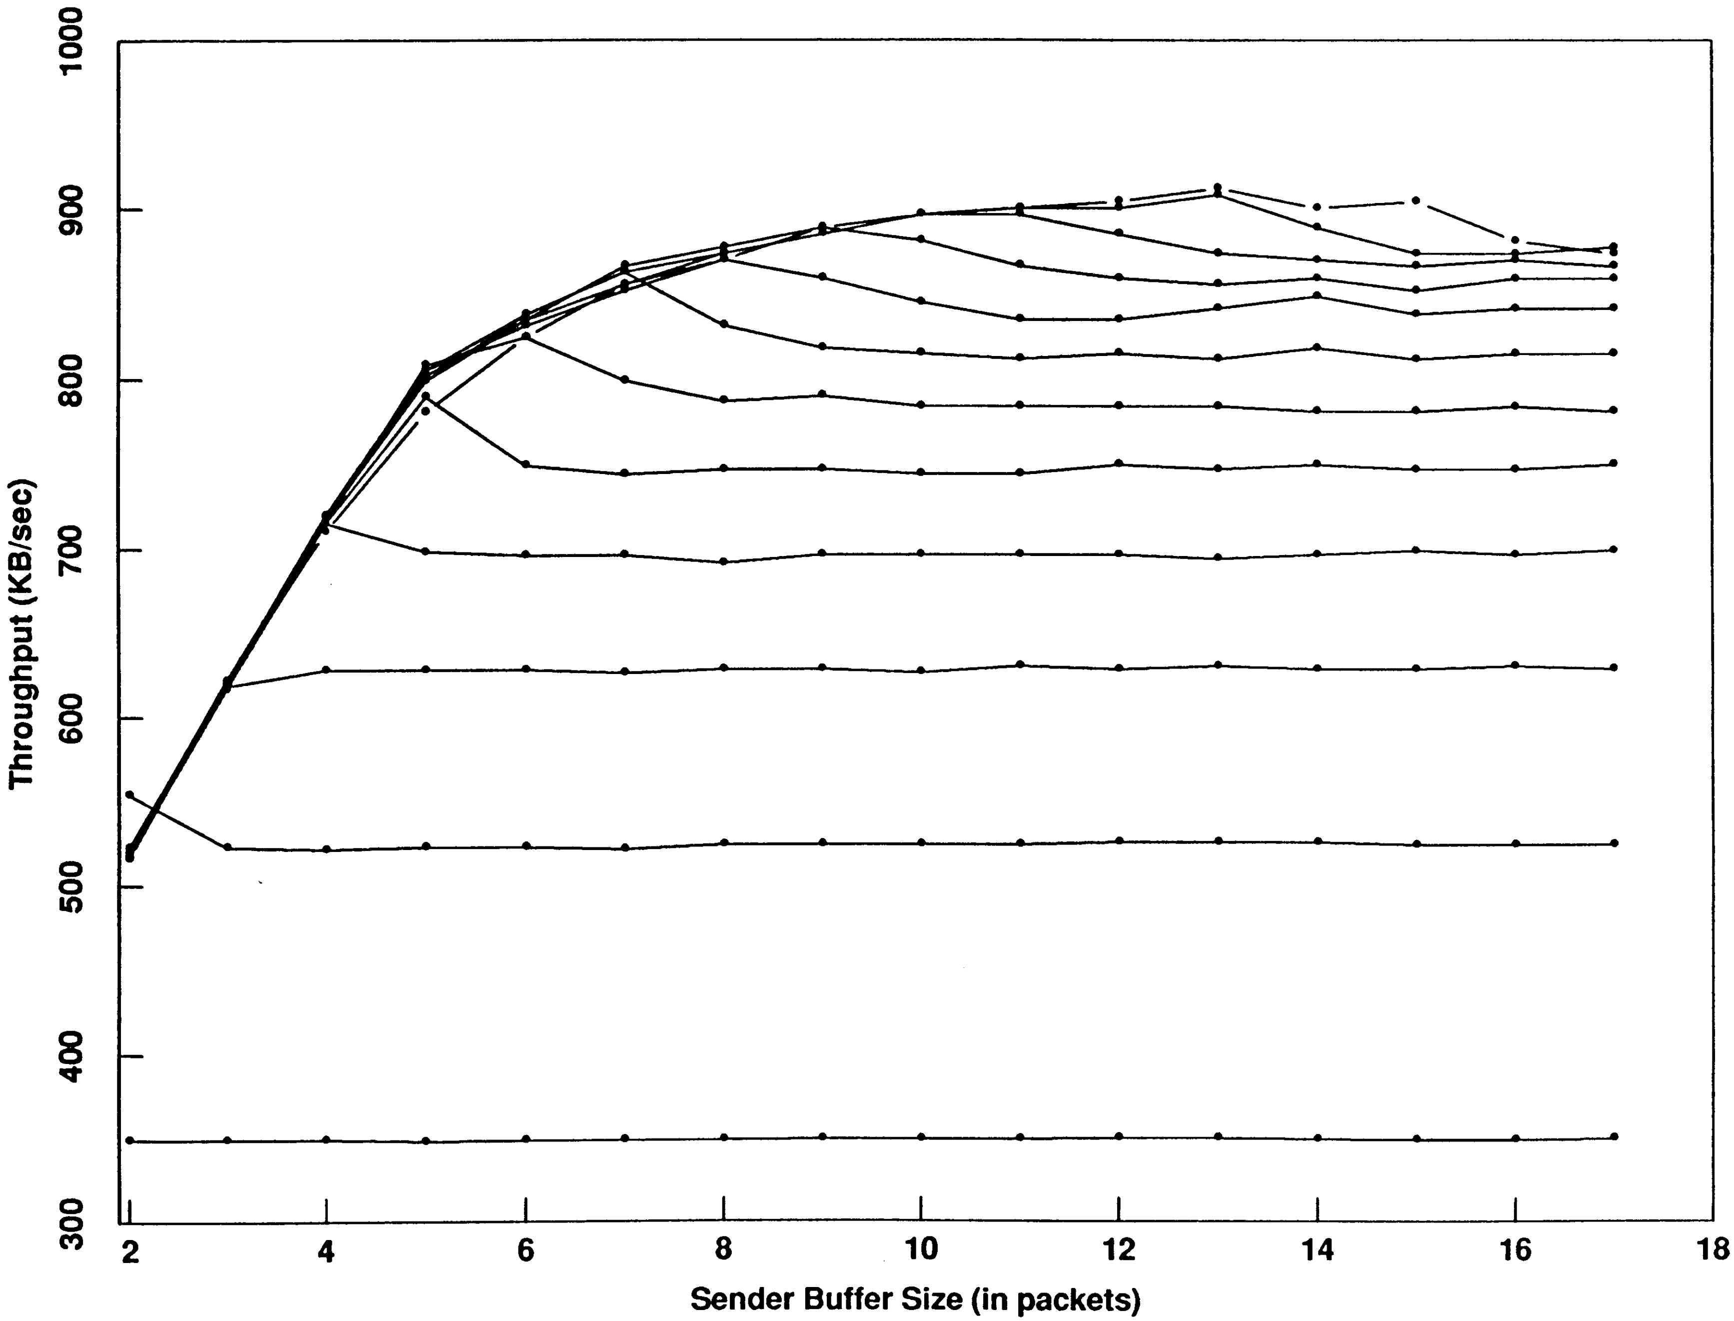

Aug 88: New Network Code Throughput Tests

(receiver buffer size fixed at 7 packets)

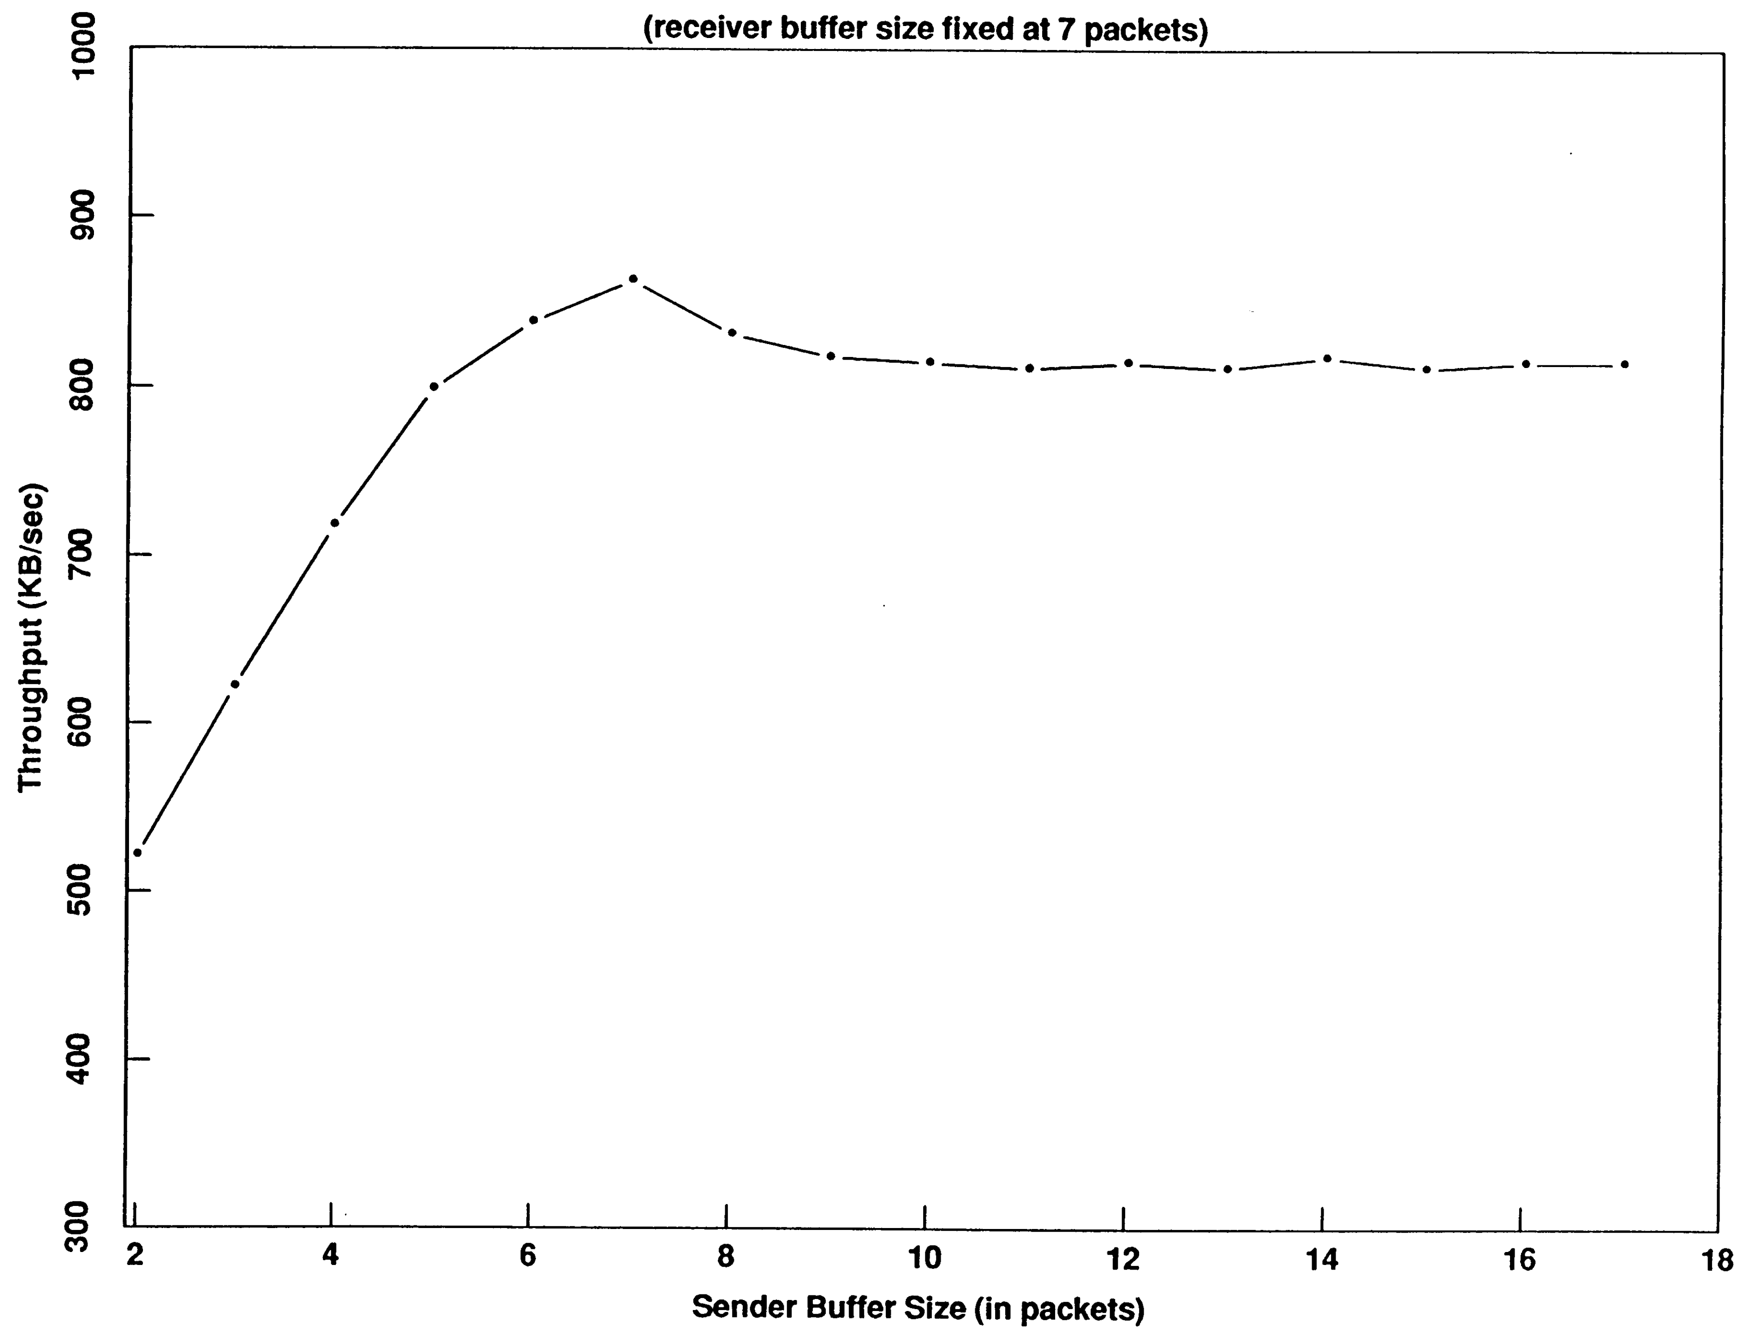

Aug 88: New Network Code Throughput Tests

(receiver buffer size fixed at 4 packets)

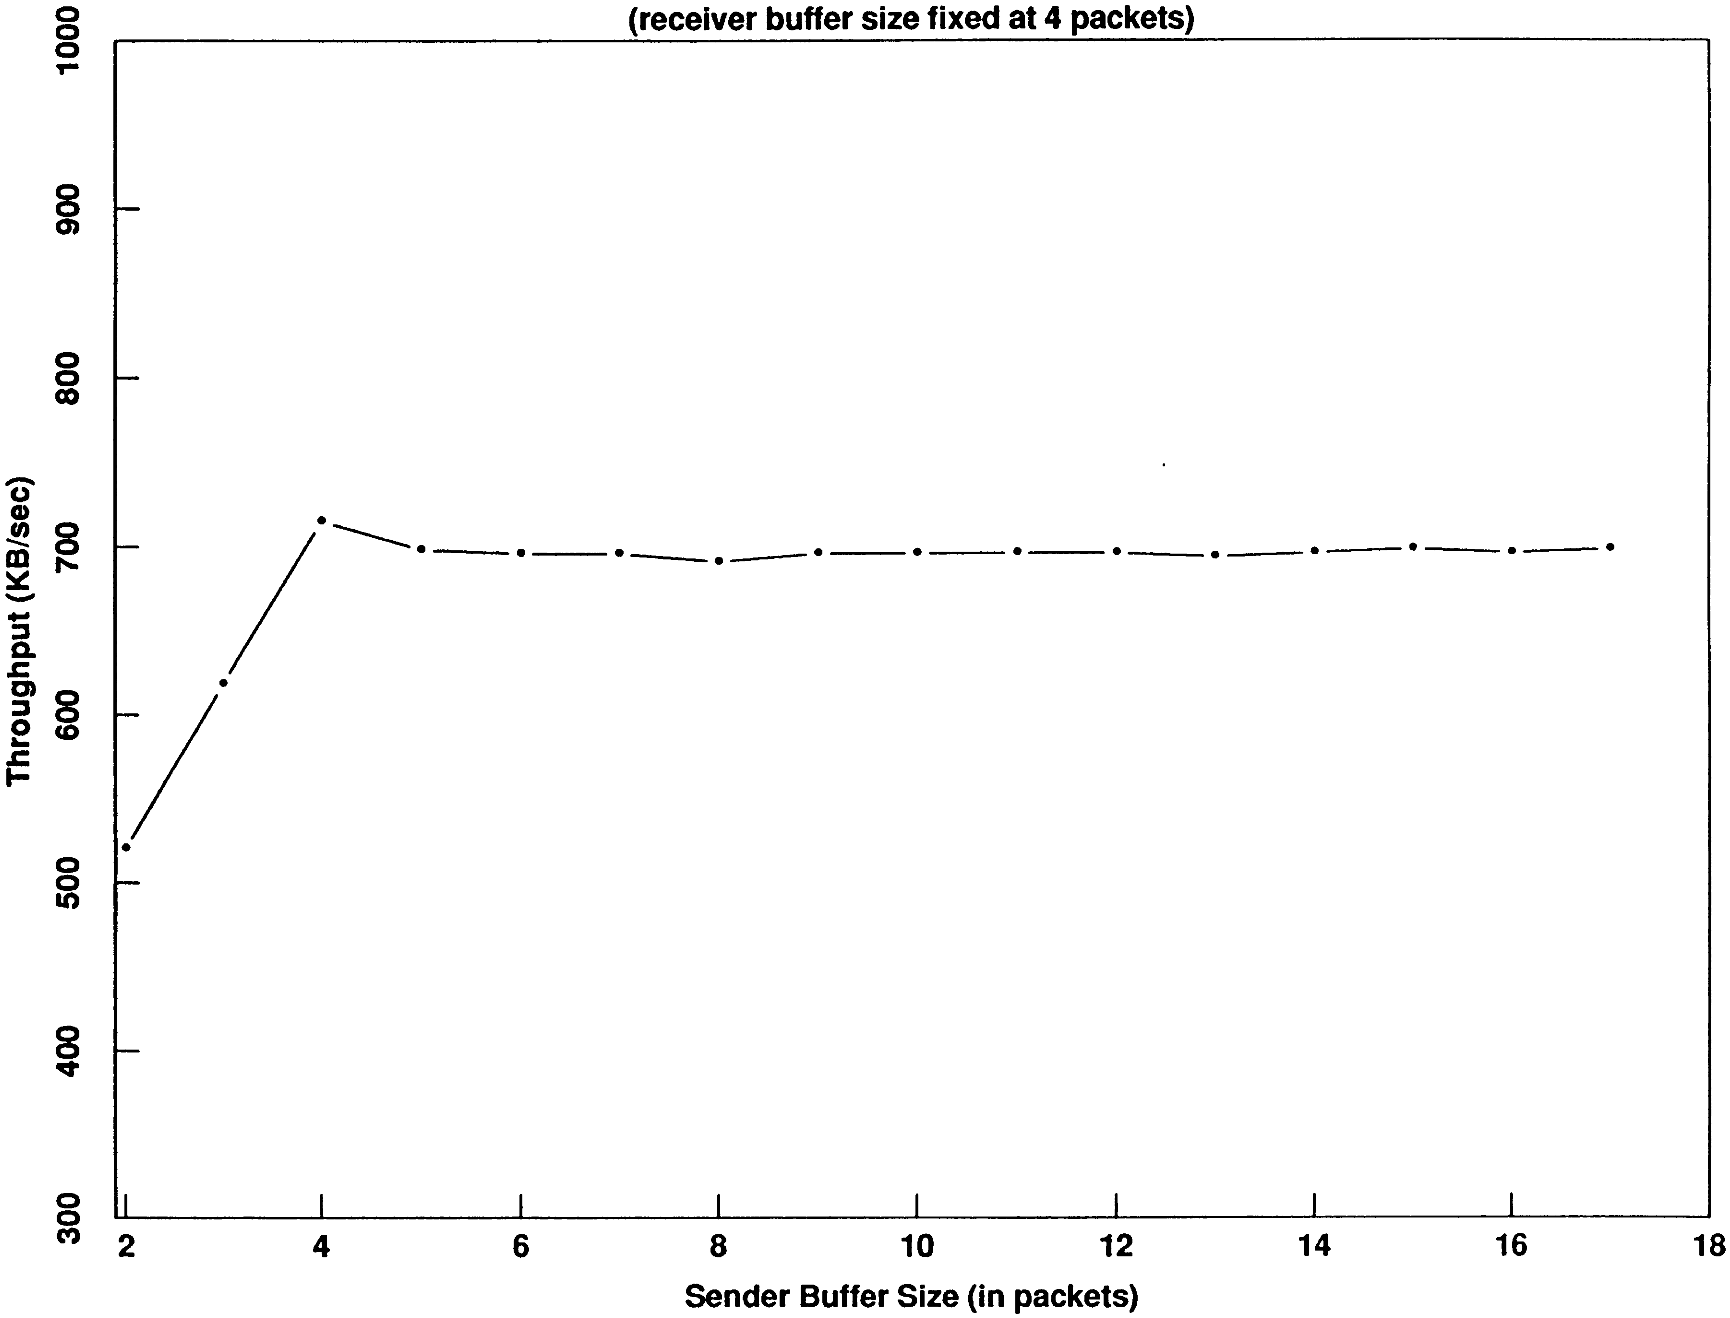

Aug 88: New Network Code Xput Tests

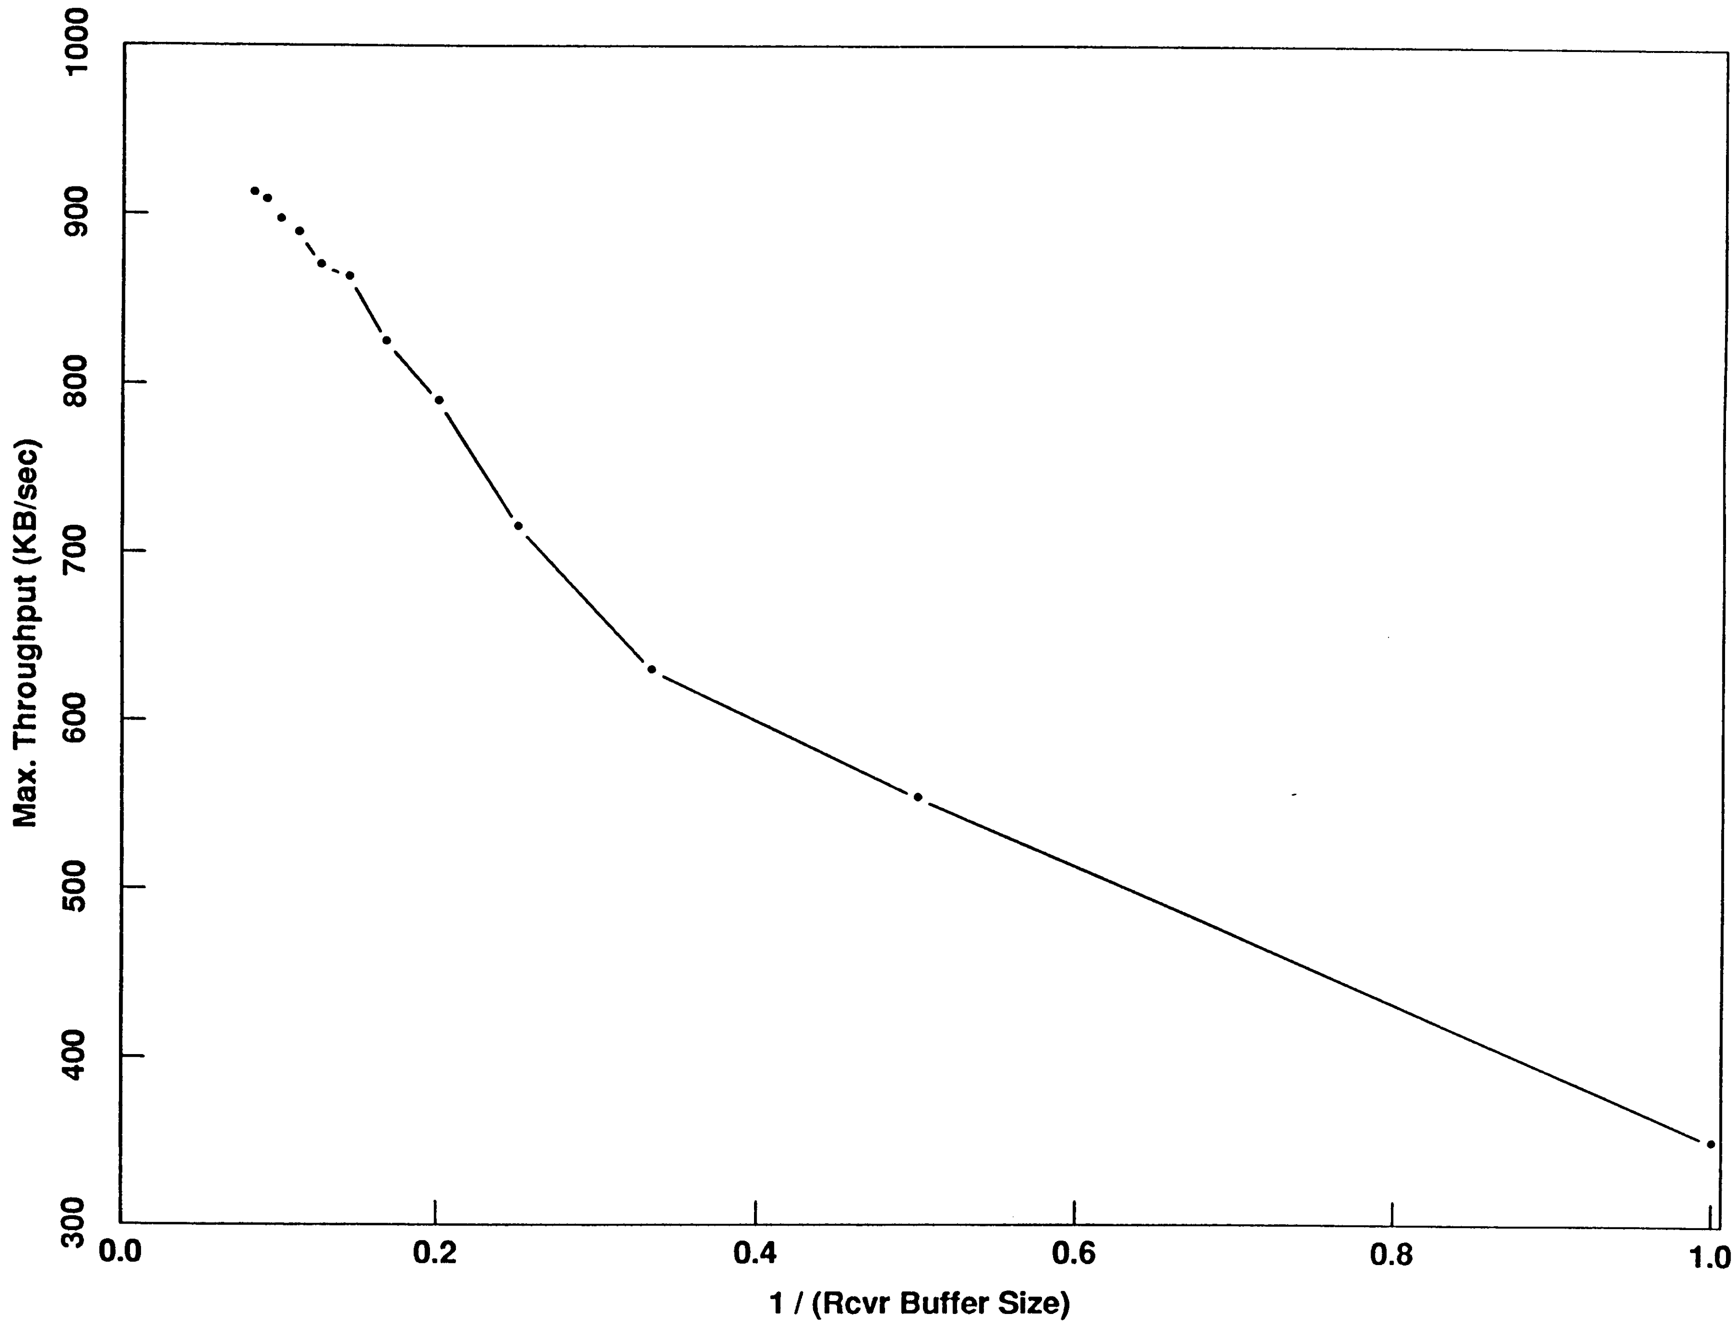

4BSD TCP "Header Prediction"

Identifying candidates (*ti* points to the incoming segment, *tp* points to the protocol control block).

Brute-force version:

```
#define FLAGS (SYN|FIN|RST|URG|ACK)

if (tp->state == ESTABLISHED &&
    (ti->flags & FLAGS) == ACK &&
    ti->seq == tp->rcv_nxt &&
    ti->win == tp->snd_wnd &&
    tp->snd_nxt == tp->snd_max) {
```

Minimalist version:

```
if ((ti->flags & FLAGS) == tp->pred_flags &&
    ti->seq == tp->rcv_nxt &&
    ti->win == tp->snd_wnd) {
```

"Receiver" processing: (I.e., data in the packet.) Make sure there's no piggy-backed ack, no packets on the reassembly queue, and enough buffer space to take the data.

```
if (ti->ack == tp->snd_una &&
    ti->len <= so->so_rcv.sb_cc)) {
        tp->rcv_nxt += ti->len;
        m->m_off += sizeof(struct tcpiphdr);
        m->m_len -= sizeof(struct tcpiphdr);
        sbappend(&so->so_rcv, m);
        sorwakeup(so);
        tp->t_flags |= TF_DELACK;
        return;
}
```

“Sender” processing: (i.e., no data in the packet.) Make sure something is acked, the ack is for data in-transit, and we’re not in the middle of slow-start or congestion avoidance.

If this segment was timed, update the round-trip timer. If all outstanding data is acked, stop the retransmit timer, otherwise restart it for the next segment. If there’s a process waiting to output, give the user a crack at the new space. Otherwise, if there’s data in the socket buffer, let the output routine decide whether to send it.

```
if (SEQ_GT(ti->ack, tp->snd_una) &&
    SEQ_LEQ(ti->ack, tp->snd_max)) {
        if (tp->t_rtt && SEQ_GT(ti->ack, tp->t_rtseq))
                tcp_xmit_timer (tp);
        sbdrop(&so->so_snd, ti->ack - tp->snd_una);
        tp->snd_una = ti->ack;
        tp->t_timer[REXMT] =
                tp->snd_una == tp->snd_max ?
                        0 : tp->t_rxtcur;
        m_freem(m);
        if ((so->so_snd.sb_flags & SB_WAIT) ||
            so->so_snd.sb_sel)
                sowwakeup(so);
        else if (so->so_snd.sb_cc)
                (void) tcp_output(tp);
        return;
}
```

Congestion Control Observations
Using NETMON

Allison Mankin
Mitre

Congestion Control Observations Using NETMON

Allison Mankin

MITRE

Internet Engineering Testbed

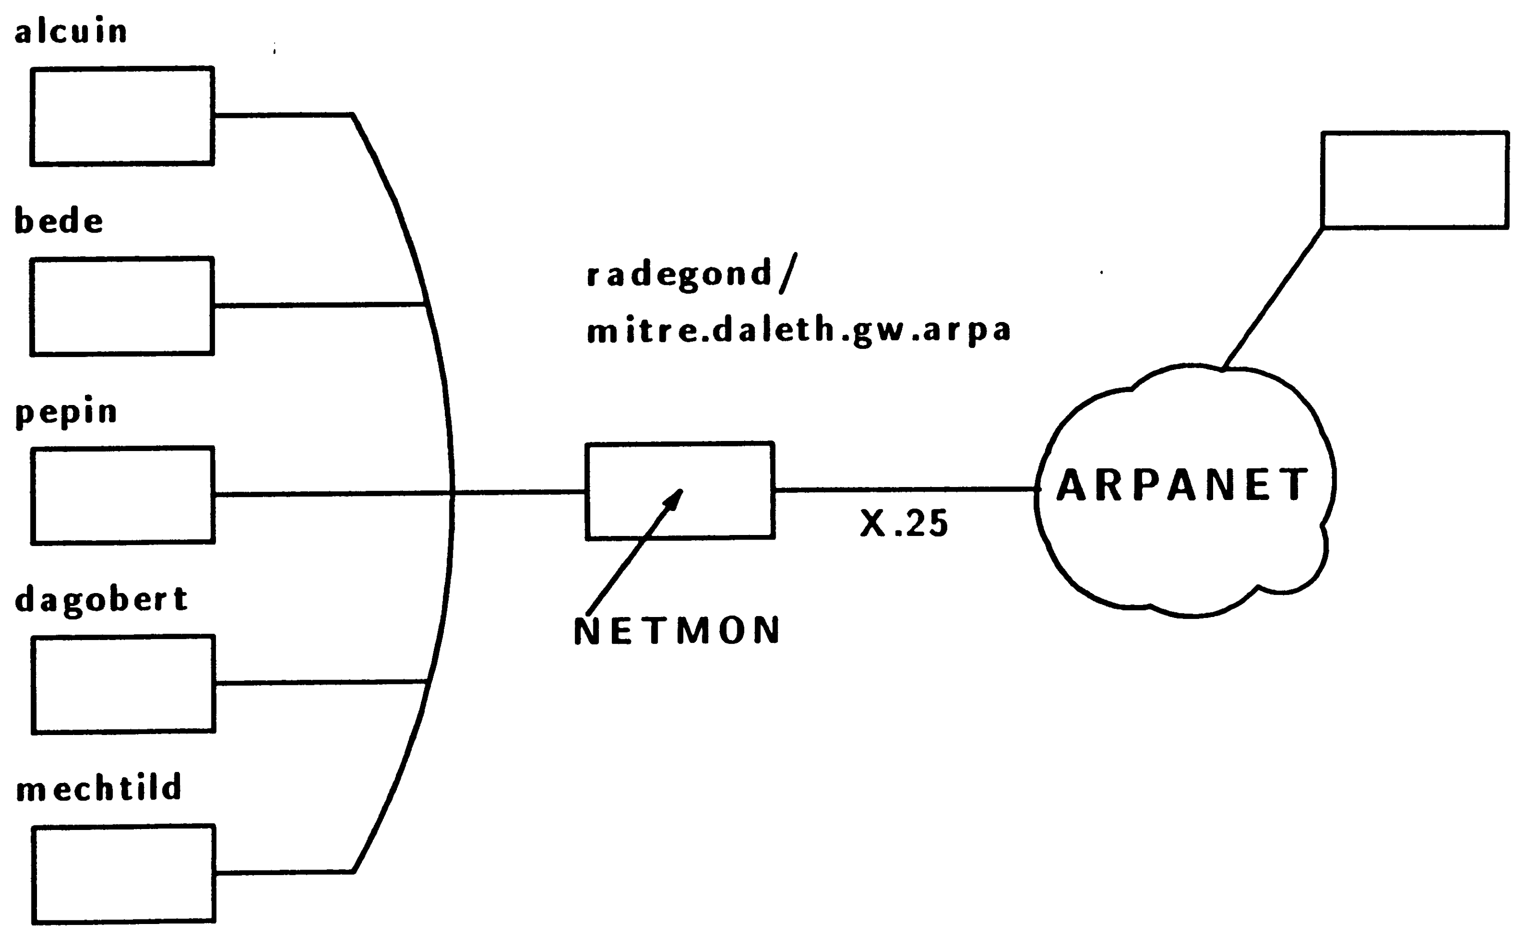

Expanded Internet Engineering Testbed

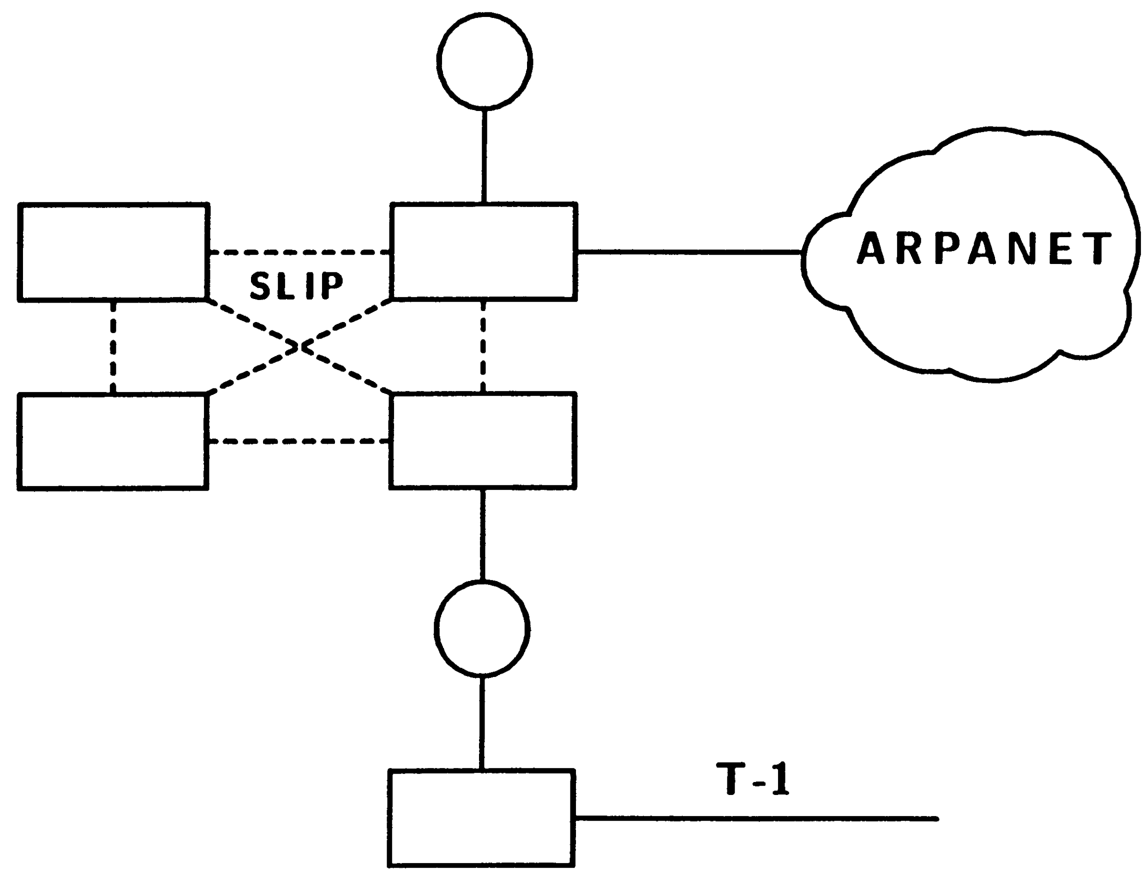

NETMON

- Internal packet monitoring for
  - BSD UNIX 4.3
  - SUNOS 3.X
  - ??
- Not restricted to Ethernet
- Would like some installation in other BSD clones

NETMON (continued)

- Double buffer design and fixed-size records to minimize amount of interrupt disabling
- Timestamp can be from alternate clock
- Access (reading the records, modifying parameters) through special file

NETMON Records

- MBUF ADDR
- TIMESTAMP
- LOCATION
- EVENT
- INFO WORD

NETMON Records (continued)

- QUEUE LEN
- PROTOCOL DATA
  - 12 BYTES
    (Plus next record whole, if header mode is selected instead of measurement mode)

- **Defined locations**

IP-IN	ETH_IN	X25_IN
IP-OUT	ETH_OUT	X25_OUT
	ETH_DEQ	X25_DEQ

- **Events**
  - DISCARD
  - FORW — DELIVER — REQUEST
  - AT_SEND_REQD
  - AT_SEND_FAIL
  - NEW_CIRC_REQD

(continued)

- INFO WORD
  (Varied Meaning)
  X25 LCN# (QUEUE ID)
  Interface ID

NETMON Probes in Gateway - 4BSD IP and Network Drivers

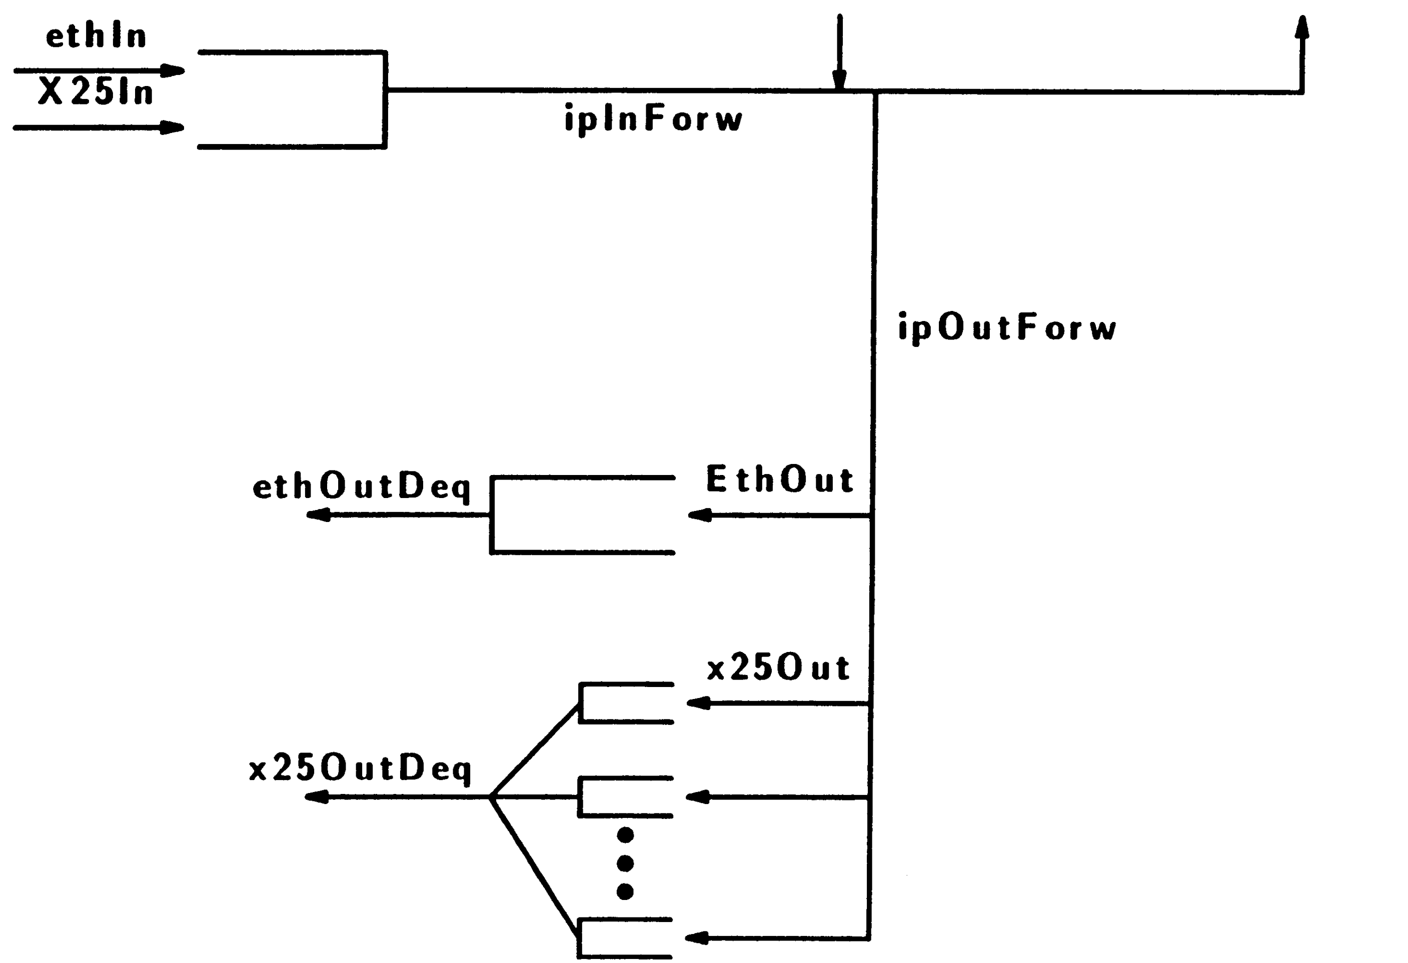

MITRE

Measurement Mode Example (as formatted by NETMONLOG)

4105	801E6700	16:45:11.84	ethIn>0	0	len -1:
4106	801E6700	16:45:11.84	ipInForw>0	qe0	len 40:
alcuin.mitre.org.1133>ucbarpa.berkeley.edu.9					
4107	801E6700	16:45:11.84	ipOutForw>-1	dda0	len 40:
alcuin.mitre.org.1133>ucbarpa.berkeley.edu.9					
4108	801E6700	16:45:11.84	x25Out>2	2	len -1:
•					
•					
•					
4135	801E6700	16:45:13.03	x25OutDeq>3	2	len -1:

Measurements of Slow-start TCP

- Competing 4.2 TCP connections do hog bottleneck queue
- Attempts to find a cliff with increasing numbers of connection sharing bottleneck queue
  - Couldn't breake RTT estimation - couldn't see any spurious retransmissions and no established connections gave up
  - From 18 on, some connections got to SYN giveup point (relic of 4.2 - 75 seconds counted by keepalive timer)
    - All of these - 4 SYNs in a row source quenched (SQ preference to small windows)
  - Repeated, frequent queue overflow starting with 3 connections to same destination

- Aside -
  - Need to be able to monitor for bottleneck queues like this 3-connection case
- ?? MIB Object -
  - First hop destination of last dropped packet
  - Poll in conjunction with a counter of drops

Preliminary Results on Random Dropping

- Algorithm selected with setsockopt
- First Cut:
  - ip_output gets Q empty/not empty indications from network drivers
  - hysteresis in starting and stopping random drop
  - when on, drop each jth (a random number) packet
  - source quench not sent
- Direct observation → fairness, queue dynamics

Random Dropping

- Queue max brought down to 5
- Many queue overflows in first 10 secs (out of 300) were totally eliminated by starting random drop at once instead of on next input
- Little decrease in total number of dropped packets - still have to "make mistakes to get information"
- 4.2 TCP gets large share of discards
- Innocent bystanders get very few discards (haven't yet seen more than 1 discard for any background connection)

What Is USENET? What Is NNTP?

Gene Spafford
Purdue University

What is USENET? What is NNTP?

Subtitle: Where did all my disk space go?

Gene Spafford

Dept. of Computer Sciences
Purdue University
W. Lafayette, IN

spaf@cs.purdue.edu

Some History

- A News

  Started as mailing lists in 1979 at Duke and UNC. Tom Truscott & Jim Ellis had the idea, based on UUCP.

  Steve Bellovin did first version of news, with Steve Daniel. Intended for less than 100 sites, less than few messages per group.

- B News

  B News at U. C. Berkeley by Mark Horton and Matt Glickman. 2.9 released in 1982.

  Notes written by Ray Essick and Rob Kolstad at same time. Based on Plato system, integrated with News in 1985-1987.

- Extensions

  2.10 was released in 1984 by Rick Adams @ seismo. Moderated groups were added at this time.

History (cont.)

- Directed Changes

  2.11 was released in 1986. Included batching, compression, sendme features, central consistency control.

- Next Generation

  Now in Beta Test — available 1989.

Structure

- Each article stored as a separate file
- Like articles are grouped in directories by topic
- Topics have hierarchies (comp, news, sci, soc, misc, rec, talk)
- Hierarchies differ by content and distribution. Examples; bionet, biz, world.
- Article structure defined in RFC 1036 — header and body. Simple files, simple text.
- Central control files contain pointers & authorizations
- Independent reader agents access files & display articles
- News posting and transfer agents interact with control files through well-defined functions.

Flow

- Articles copied to neighboring systems based on distribution

- Cycles rejected, too old articles rejected. Information in the article header used to determine validity.

- ''Flooding'' algorithm — redundancy built in

- Articles expired locally after set interval, or canceled

Transport

Primary transport for Usenet has always been UUCP.

- 1979 to 1982, 300 baud dial-up
- 1982 to 1985, 1200 baud dial-up
- 1985 to present, 2400 baud dial-up
    - * 1986 had LZ compression, UUCP-over-TCP
    - * 1986 saw NNTP arrive (RFC 977)
- 1987 to present, Telebit Trailblazers with MLZ and UUCP support
- 1986 to present, NNTP over TCP, UUCP over X-25
- 1985 to present, some sites get USENET via tape!

Traffic

Based on figures from R. Adams, H. Spencer, M. Horton, S. Bellovin and B. Reid:

- 1979: 3 sites, 2 articles per day
- 1980: 15 sites, 10 articles per day
- 1981: about 150 sites, 20 articles per day
- 1982: about 400 sites, 35 articles per day
- 1983: over 600 sites, 120 articles per day
- 1984: over 900 sites, 225 articles per day
- 1985: over 1300 sites, 375 articles per day, 1Mb+ per day
- 1986: over 2500 sites, 500 articles per day, 2Mb+ per day
- 1987: over 5000 sites, 1000 articles per day, 2.4Mb+ per day

Present Traffic (as of 10/1/88)

- Nearly 11000 sites.
- Sites in more than 17 countries, including Australia, New Zealand, Japan, Canada, England, Sweden, France, Italy and Germany.
- Over 1800 unique articles per day, average
- Over 4Mb of traffic per day, average
- Potential audience of 1,480,000 readers; actual readers in excess of 303,000.
- Most widely read group has over 40,000 regular subscribers.
- Over 450 active newsgroups
- over 80% of articles reach main sites in 1 day, over 97% in 3 days

Growth in Sites

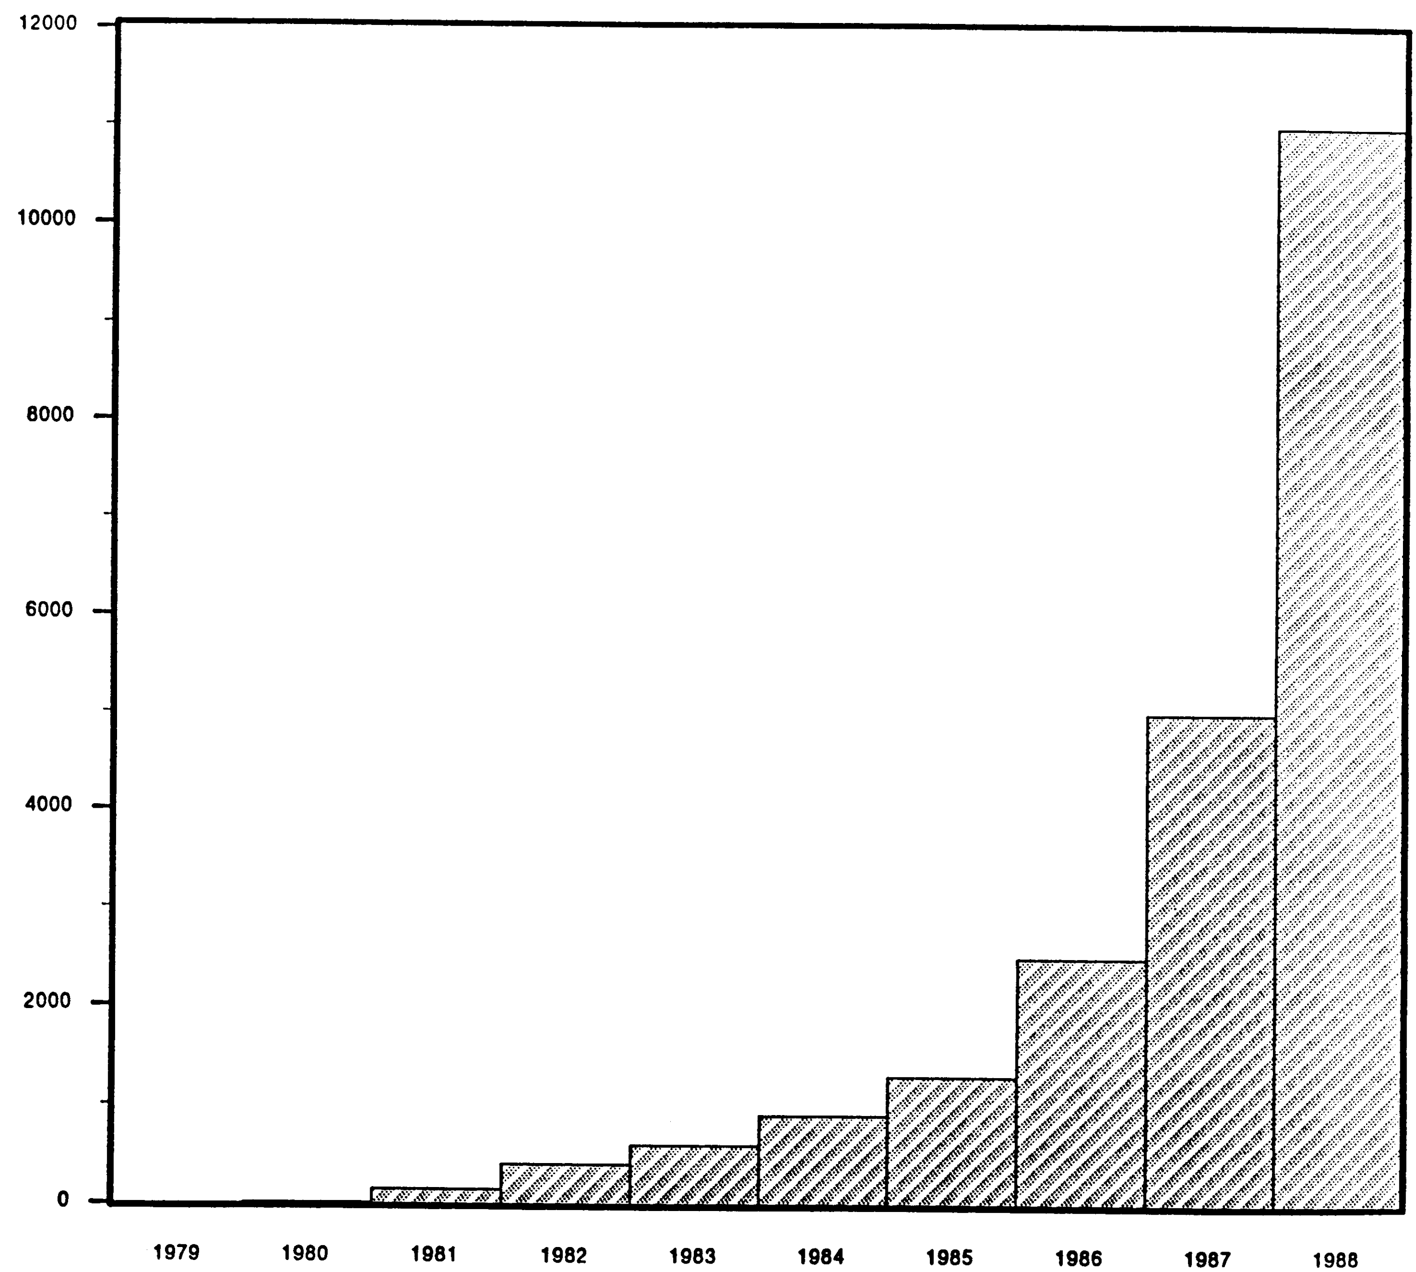

Growth in Traffic

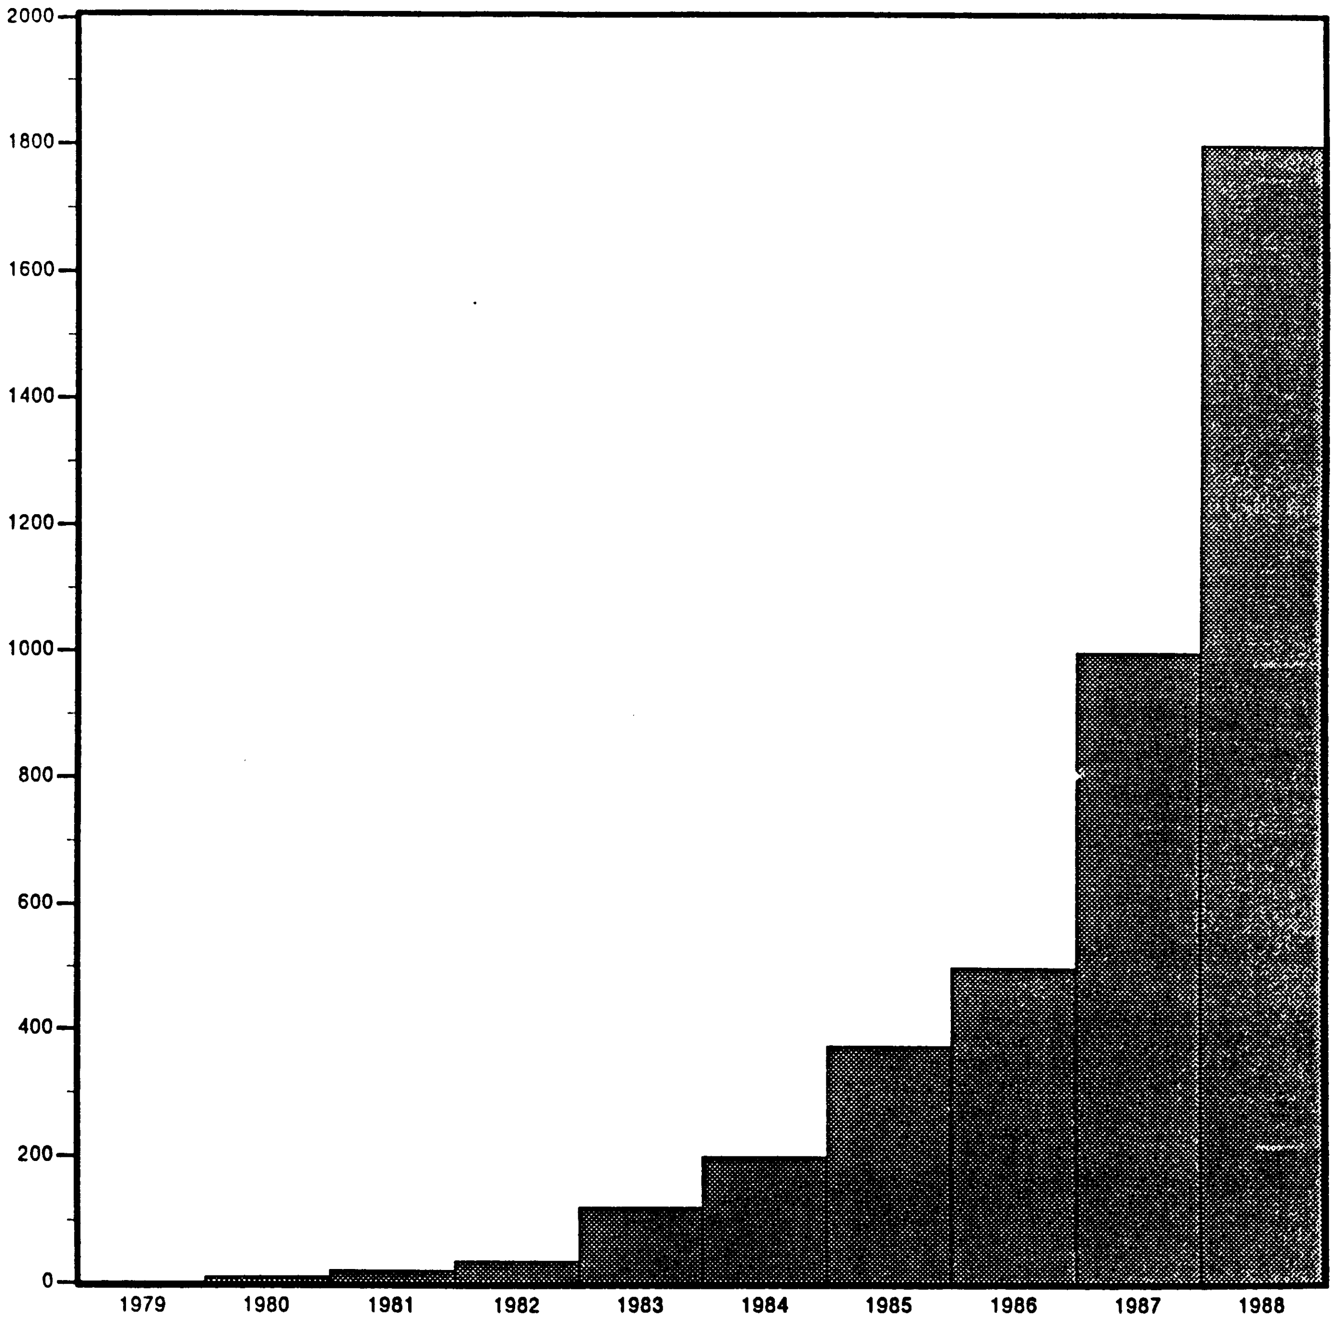

Control

...interesting feature — there is no authority!

Usenet operates on consensus and momentum.

- Any site can join
- No one controls flow, although some ''old-hands'' are listened to more carefully than others.
- Peer pressure is main control
- Abuses are surprisingly few and minor
- Voluntary education and upgrades — structure has some aids
- Increasing cost having more significance

NNTP

- Developed from independent work by Brian Kantor and Phil Lapsley, 1985.

- RFC 977, released in 1986

- Four major goals:
  * Reduce phone traffic for news transfer
  * Reduce "flooding" IP traffic
  * Allow diskless computers to access news
  * Reduce impact of mailing lists by integration with news

- Uses server daemon on TCP port.

- Supports posting, reading, transfers

- Reader agents for Unix, VMS, TOPS-20, MS-DOS, and Genera-7.

- Vastly increased connectivity; tremendous reduction in machine impact.

Concerns

- Increasing volume
- Educating users; maturity of users
- Comprehension of namespace
- Status of Usenet sites — not common carrier
- Costs — communications, CPU, disk, human
- Legal questions — copyright, trade secret, slander, over-zealous prosecutors
- Nutcases
- Continuity of software and guidance

Social Effects

- Citations to USENET
- Collaborative projects
- Conferences
- Software community
- Friendships, romances, marriages
- USENET as a condition of employment
- Image of schools and companies
- Growth of new services (uunet, for example)
- Source of research material
- Publications media — scholarly
- Publications media — fanzines (e.g., *OtherRealms*)

Some Futures

- Commercial USENET?
- Alternate networks?
- The ‘‘Balkanization’’ of USENET
- Reappearance of mailing lists
- Hypermedia

The NIC Domain Chart

Mark Lottor
SRI-NIC

ZONE

Zealot Of Name Edification

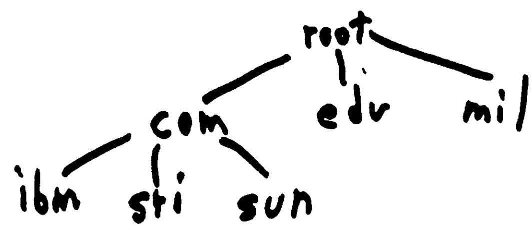

data collected

- domains, servers
- hostnames
- nicknames
- addresses
- hinfo
- wks
- MX

Domain Tree-Walker Statistics

Domains	1280
Domains (no data)	*140*
Internet Hosts	56000
Registered Hosts	*5700*
MX-only entries	3500
"" MX entries*	*550*
Gen'd host table	4340 kb
Official host table	*600 kb*
host table string searches	
"Sun"	17800
"Sun.Com"	7500
"Unix"	14200
"Vax"	5500
"IBM"	4700
"GW"	2200 – N
Registered GWs	260
"Tops-"	60

On Some T_1 Satellite Link Performance

John Lekashman
NASA-Ames

[slides not provided]

7. ADDITIONAL WORKING GROUP UPDATES

Interconnectivity

Guy Almes
Rice University

NASA Ames Meeting
27 September 1988

Interconnectivity WG Update
Reported by Guy Almes

27 September 1988
NASA Ames Rsch Ctr.
Moffett Field, CA

Attendees:

Guy Almes <almes@rice.edu>, chair
Hans-Werner Braun <hwb@mcr.umich.edu>
Michael Brescia <brescia@alexander.bbn.com>
Scott Brim <swb@tcgould.tn.cornell.edu>
Joe Choy <choy@windom.ucar.edu>
Phill Gross <gross@gateway.mitre.org>, ex officio as IETF chair

Milo Medin <medin@nsipo.nasa.gov>
Russ Mundy <mundy@beast.ddn.mil>
Mike St. Johns <stjohns@beast.ddn.mil>

Also in the Working Group, but unable to attend this meeting:

Marianne Lepp <marianne@alexander.bbn.com>
Jacob Rekhter <yakov@ibm.com>

The first meeting of the Interconnectivity Working Group was hosted by Milo Medin of NASA, and was called, more or less, to order at 9:00 a.m. Thanks to all those who could attend on such short notice and to Milo for serving as host under the twin disadvantages of recovering from a close encounter with a car and being torn away from InterOps preparations.

We discussed our short and long-term agenda. In the short-term, the IAB has asked Phill Gross for input on the status of EGP3, and he has asked us for recommendations since this matter falls squarely within our technical area. In the longer term we hope to improve inter-autonomous-system routing in practical ways that allow timely implementation. (Refer to the IWG Charter for a more detailed discusssion of this.) (NB: In hindsight, as the meeting progressed, it seemed to me as though these two agenda foci did not conflict as much as I had feared.)

Hans-Werner Braun reported on a meeting held at Ann Arbor on 15 August to discuss Inter-AS routing in the NSFnet context. There was considerable overlap both of participants and of technical focus, and we benefitted from their work and insight. (Refer to Hans-Werner's notes from this meeting.) There were two technical suggestions that arose at that meeting that proved important for our our meeting:

<1> Include in the entry for each destination network advertised the AS# (i.e., the 16-bit Autonomous-System Number) of the autonomous system from which the advertiser learned the route.

<2> Develop some EGP3 metrics that describe how the route was learned. Much of our meeting consisted of:

- <a> discussing how we thought Inter-AS routing should work and
- <b> discussing how EGP-3 with these two suggestions could allow this to take place.

In the following discussion we agreed that the hierarchical NSFnet Model of (a) Multiple national backbones (backbones for short), (b) Multiple mid-level networks (regionals for short), and (c) Many campus networks (campuses for short) was normative. Each regional connects to a generally large set of campuses, and to one or more backbones. It advertises these campuses to each of the backbones, and advertises all its known routes to its campuses (or else advertises default to its campuses). It will occasionally happen that a regional connects to another regional; great care must be taken in this case. Each backbone connects to a possibly large subset of regionals, and may also connect to one or more other backbones and possibly to some campuses. There is a so-called Two-Phase Rule that dictates that a packet travels across the internet in two phases. During the first phase, it travels 'up'the hierarchy; each Inter-AS hop in this phase either stays at the same level (e.g., backbone to backbone) or goes up a level (e.g., from regional to backbone). During the second phase, it travels 'down' the hierarchy; each Inter-AS hop in this phase either stays at the same level (e.g., backbone to backbone) or goes down a level (e.g., from regional to campus). Thus, once a packet goes 'down' the hierarchy once, it can never go 'up' again. In our consideration of EGP3, we tried to think about how it would enable smart gateways between AS's to make appropriate decisions without violating simple policy rules or creating routing loops. As usual, we want to determine strategies that improve the current situation while being deployable within the near-term future.

With specific regard to EGP3, we came up with the following:

<1> The EGP3 Idea paper should be revised and turned into an RFC as a Proposed Recommended Standard. We understand this will require work, and will help Marianne with the needed additions while keeping editorial leadership with her.

<2> Add to the current EGP3 design a Next-AS field in each route. This field denotes the immediate AS from which the advertising AS received the route. There needs to be some denotation of an empty value for this field.

<3> We will need a Metric Type that measures the number of AS's in the EGP chain from the originating AS. This metric is important in the case that a non-empty Next-AS value had to be "shifted out". The metric will have allowable values for other cases, but its presence will be manditory when this shifting out has occurred.

<4> We will need to describe recommendations for normative use. For example, we should describe how the protocol can be used in a fashion that avoids routing loops.

<5> We recommend that EGP3 be used within NSFnet, the NASA Science Network, the NSFnet-related mid-level networks, and other components of the national research internet. We understand that conversion of the DDN to EGP3 may take quite a long time for a variety of primarily non-technical reasons.

<6> In addition to the Metric Type for AS hop count, we also recommend a Metric Type that, for advertisements coming from the NSFnet Backbone, will mark the route as via the primary or secondary or tertiary Backbone exit point.

<7> We recommend that vendors and other implementers of external gateways (as distinct from intra-AS routers) try to exploit the possibilities presented by EGP3 in evolving toward greater sophistication. The trend we encourage is one in which the notion of Border Gateways that connect different AS's to each other grow in capability.

<8> We note that the route data communicated by the EGP3 packets can be split into two kinds: (1) information about the interconnection of various AS's and (2) information about which destination networks are reachable via these AS's. There is reason to think these two kinds of data will change in different patterns and that updates to them can be handled differently. Studying this distinction in practice and exploiting it are important for us to do.

<9> We stress that there is a great need for an active engineering effort in this area, and we urge both the refinement and implementation of EGP3 and its exploitation.

<10> This engineering effort will need to include the use of such measurement tools as Braden's statspy.

<11> This engineering effort will provide fruitful areas of interaction between the Interconnectivity Working Group and the FRICC's "Intersec" Workshop and the IETF's Open Inter-Autonomous System Routing Working Group. We look forward to this interaction.

With specific regard to the Core, we came up with the following:

<1> Part of our answer is implicit on our recommendations regarding EGP3.

<2> A certain amount of manually entered data, such as the so-called "Policy routing database" of the NSFnet backbone, will probably been needed for the foreseeable future.

<3> We discussed the following as normative patterns of routing exchange:

* There would be one Backbone that advertises to its regionals routes learned from other regionals. (That backbone can be thought of as serving as the core.)

* A regional may not advertise to one Backbone what it learned via another Backbone.

* A Backbone, on the other hand, may advertise to its regionals routes learned via another Backbone.

* We note that EGP3 allows more liberality than the current EGP without introducing dangerous exchanges of routes. Work and time will be needed to exploit this.

* There is a two-phase rule that we regard as (near) absolute: What goes up does come down, but what goes down never comes up. (This refers to packets going up and down the hierarchy of Backbone, Regional, and Campus.) (In this context, lateral motion is fine, but it makes the two-phase rule more difficult to enforce.)

We will meet at the October 1988 IETF meeting in Ann Arbor to work further on these issues.

PDN Routing

Carl-H. Rokitansky
Fern University of Hagan

USNA, Annapolis Meeting
15-17 June 1988

PDN Routing WG Update 3 October 1988
Provided by Carl-H. Rokitansky

Report of the Intial PDN Routing Group Meeting, June 16, IETF, USNA, Anapolis

(These summarizing notes of the initial PDN Routing group open meeting from the June 15 - 17 IETF were prepared by Carl-H. Rokitansky, Fern University of Hagen, FRG)

The PDN Routing group met on June 16, 1988 at IETF, USNA, Anapolis. The attendees were:

- Len Bosack, CISCO
* Mike Brescia, BBN
- Ed Cain, DCA
* J.J. Garcia-Luna, SRI
- Martin Gross, DCA
- Mike Little, M/A-COM
* Mark Lottor, SRI-NIC
* Bill Melohn, SUN
- John Moy, PROTEON
* Carl-H. Rokitansky (chair), Uni Hagen
* Greg Satz, CISCO
* Zaw-Sing Su, SRI

(* indicates membership of the PDN Routing group)

The meeting covered administrative items, background information and technical discussion:

1. Charter and Goal of the PDN Routing Group

   The DoD INTERNET TCP/IP protocol suite has developed into de facto industry standard for heterogenous packet switching computer networks. In the US the ARPANET/MILNET connects several hundreds of INTERNET networks, 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 no algorithms are provided so far to dynamically route INTERNET datagrams through X.25 public data networks. Therefore the goals of the Internet/Public Data Network Routing 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.

Main objectives of the PDN Routing group are:

- Define the cluster addressing scheme and its application to public data networks as an INTERNET standard

- Specify gateway algorithms and protocols to be used by VAN-gateways

- Develop an X.121 Address Server/Resolution Protocol

- Develop (or support other working groups in developing) routing algorithms based on routing metrics other than hop-count: costs, delay, throughput, TOS, etc.

- Provide interoperability with ISO/OSI networks via the PDN

- Specification of protocols required for an Eurpean INTERNET/Public Data Network information and operation center (cooperation with US-INTERNET NICs and NOCs)

- ISO-Migration of the INTERNET/PDN cluster

2. Mailing Lists

The intention was to install two mailing lists for the PDN Routing group. Members of the PDN Routing group will be put on an "IETF-PDN" list for internal discussion of proposals and group organization. People, interested in the ongoing work of the PDN Routing group will be put on an "IETF-PDN-INTEREST" list on request. First draft versions of proposals of the PDN Routing group will be sent to this list to encourage discussion and comments.

3. Meetings

The PDN Routing Group will meet periodically at the regular IETF meeting. These meetings will be open meetings. In addition, members might meet right before or after the IETF meeting. BBN has offered to host such (closed) PDN Routing Group meetings, if no other place is available.

4. PDN Routing Group - Short Term Goals (3 - 6 months)

4.1 PDN-Cluster

Reserve INTERNET network numbers for the PDN-cluster according to the cluster addressing scheme: check with Jon Postel and SRI-NIC

4.2 VAN-Gateways

Check which changes to the IP code would be required to support the cluster addressing scheme in existing VAN-gateways (BBN-VAN-GW).

4.3 INTERNET Gateways

Check if advertising a bunch of additional European INTERNET networks by means of EGP messages would cause a problem to the DoD INTERNET gateway system.

4.4 EGP3

Check for topological restrictions. Check if EGP3 satisfies the requirements for network reachability information exchange between VAN-gateways and if not develop a concept how a modified version of EGP3 could be used between VAN-gateways.

4.5 Routing Metrics

Develop a concept how PDN cost metrics can be taken into account in INTERNET routing decisions depending on hop count, etc.

4.6 Source Routing

Check which TCP/IP implementations (ULTRIX, TOPS-20, VMS, etc.) use the IP Source Route option, if specified in received datagrams, even in their reply packets; check with implementors if the IP Source Route option is neglected in reply packets.

4.7 Performance Tests

Provide a testbed for performance tests between PDN-hosts and INTERNET hosts via VAN-gateways subscribing to the cluster addressing scheme.

5. Medium-Term Goals (6 months to 2 years)

- Develop an X.121 Address Server/Resolution Protocol

- Develop (or support other working groups in developing) routing algorithms based on routing metrics other than hop-count: costs, delay, throughput, TOS, etc.

- Continue performance tests

- Specify the INTERNET/PDN-cluster as an INTERNET standard

- Interoperability with ISO/OSI networks in Europe and elsewhere

6. Long-Term Goals (2 - 5 years)

- Specification of protocols required for an European INTERNET/Public Data Network Information and Operation Center (cooperation with US-INTERNET NICs and NOCs)

- ISO-migration of the INTERNET/PDN cluster

7. ICCC '88 Presentation

The "Internet Cluster Addressing Scheme and its Application to Public Data Networks" will be presented at the 9th International Conference on Computer Communication (ICCC '88) in Tel Aviv, Israel, Oct 30 - Nov 4, 1988.

8. Report on the European situation (LANs using TCP/IP and ISO/OSI status)

- DFN: The German Research Network (DFN) favorizes the implementation and use of ISO/OSI protocols. However since these protocols are not fully specified and not generally available so far, most of the attached universities are now running LANs using TCP/IP protocols. Most sites would be very interested in an interoperation between these LANs through the national X.25 Public Data Network (DATEX-P) as well as to interconnect these LANs to the US INTERNET via international links (point-to-point links and SVC through X.25 by means of VAN-gateways). One disadvantage of using X.25 connections is the fact that the costs depend on the data volume transfered. However, fortunately, the DFN has agreed with the German PTT, that the PTT will probably offer an X.25 research network for universities and research establishments at fixed (reasonable) costs. Since similar projects are under consideration in other European countries (Netherlands, etc.), an European X.25 research network might be implemented within the next years. This would have a significant advantage for the interconnection of academic LANs now using TCP/IP, because the exchange of INTERNET network reachability information between attached LANs via X.25 research network links would not be cost sensitive at all.

- BELWUE: The experimental Baden-WUErtemberg Extended Lan (BELWUE) is a high speed network at 140Mbit/sec (!), also subscribing to the TCP/IP protocol suite and interconnects computers and supercomputers (CRAY, etc.) at the University of Stuttgart and the Univerisity of Karlsruhe. Several universities and some companies in the Stuttgart area would be interested to be connected to this high speed network for online use of CRAY services via X.25 links.

- other: Several other networks in Europe are using (e.g., EUNET), or plan to use TCP/IP protocols, and are interested to be connected to the US INTERNET (point-to-point links or X.25 connections).

9. X.121 Address Server/Resolution

An important issue is the development of an X.121 Address Resolution Protocol. X.25 specific characteristics (no broadcast feasibility, cost sensitive, no reverse charging on international calls) must be taken into account.

10. Routing of INTERNET datagrams through X.25 networks

To allow worldwide interoperation between LANs now using TCP/IP protocols via VAN-gateways and X.25 links, network reachability information must be exchanged. The question is, whether this information should be spreaded worldwide, and maintained and updated in all INTERNET gateways or it should be gathered and updated in specific route servers, and provided on request.

11. Action items

- Development of an X.121 address resolution protocol (Mike Brescia)

- Discussion of methods and requirements involving route servers (Len Bosack/Greg Satz)

- Development of hierarchical gateway algorithms for PDN routing and network reachability information exchange between level-1 and level-2 VAN-gateways (Carl-H. Rokitansky)

- Submission of final version of the INTERNET cluster addressing scheme paper for publication in Proceedings of the ICCC'88 (Carl-H. Rokitansky)

- Proposal for a sophisticated mapping between DNICs and INTERNET/PDN-cluster network numbers (Carl-H. Rokitansky)

- Procedure of assigning and organizing PDN-cluster network numbers (Zaw-Sing Su/Mark Lottor)

12. Next meeting

The next (open) meeting of the PDN Routing group will be at the IETF meeting at Ann Arbor in October.

Carl-H. Rokitansky
----

8. PAPERS DISTRIBUTED AT IETF

NSFnet Connectivity and Configuration

Susan Hares
Merit, Inc.

Pittsburgh Supercomputer Center (PSC)

Pittsburgh, PA

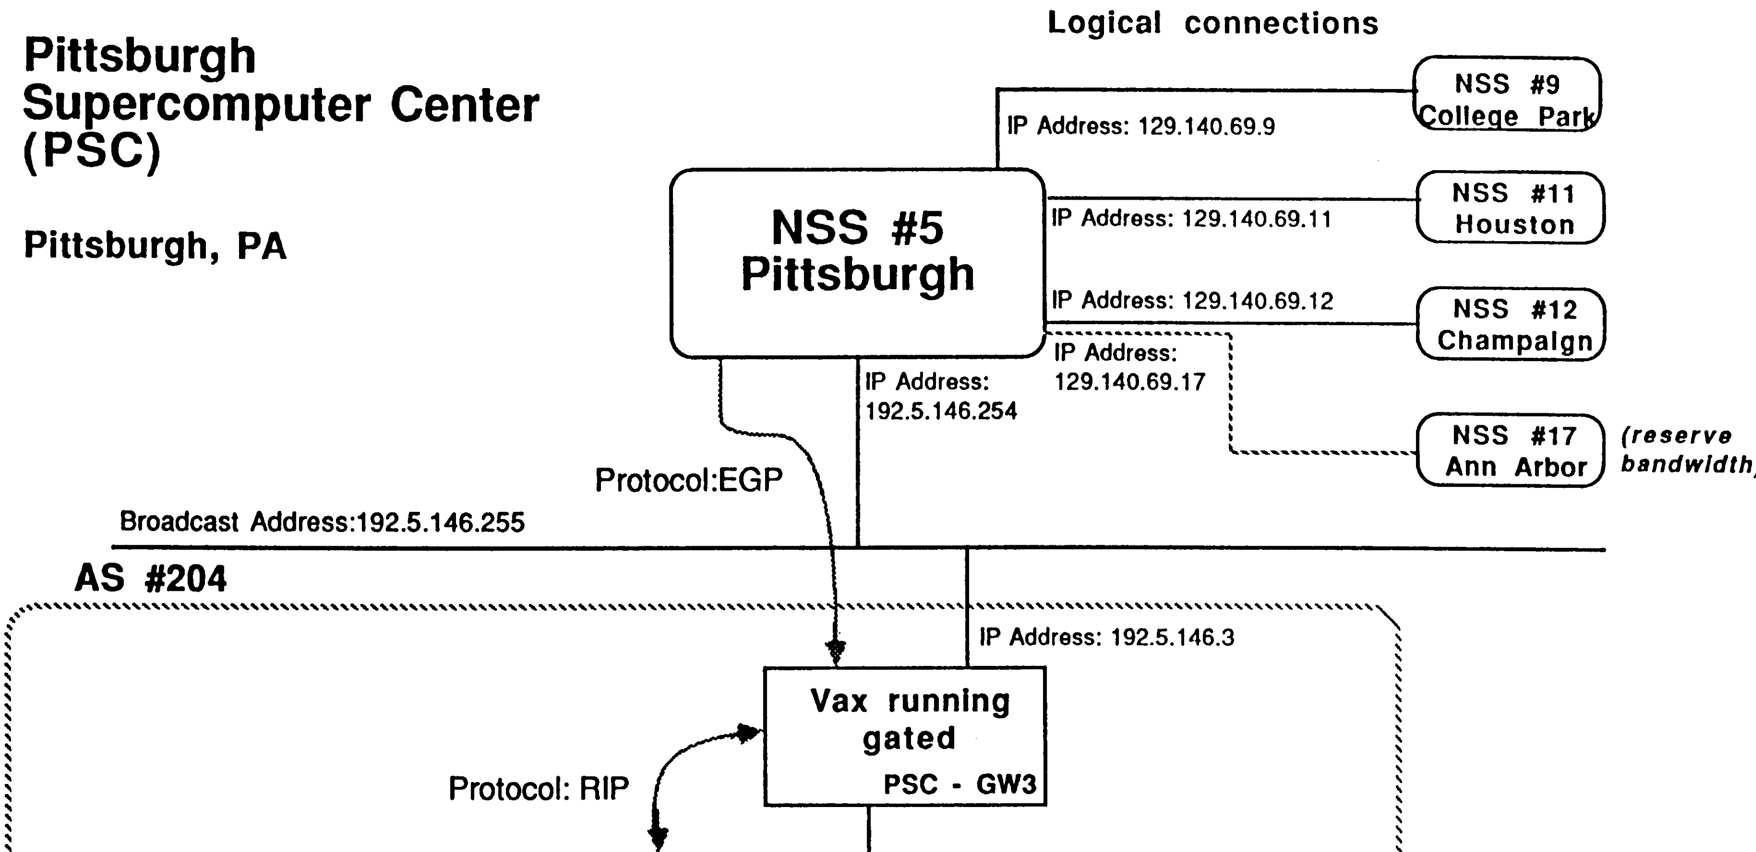

PSCnet

Net Number	Net Name	Network Location
128.2	CMU-NET	Carneige Mellon University, Pittsburgh, Pennsylvania
128.146	OHIO-STATE	Ohio State University, Columbus, Ohio
128.182	PSCNET	Pittsburgh Supercomputing Center, Pittsburgh, Pennsylvania
129.1	BGSU	Bowling Green State University, Bowling Green, Ohio
129.22	CWRUNET	Case Western Reserve University, Cleveland, Ohio
129.137	UN-OF-CINCI	University of Cincinatti, Cincinatti, Ohio
129.25	PRPNET	PRPNET, Pittsburgh, Pennsylvania
130.49	U-PITT	University of Pittsburgh, Pittsburgh, Pennsylvania
192.5.146	CPW-PSC	PSCNET NSS 5, Pittsburgh, Pennsylvania
192.31.3	ALCOA-NET	Aluminum Company of America, Alcoa Center, Pennsylvania
192.35.79	CCFNET	Cleveland Clinic Foundation, Cleveland, Ohio

Primary

10-15-88

San Diego Supercomputer Center (SDSC)

San Diego, CA

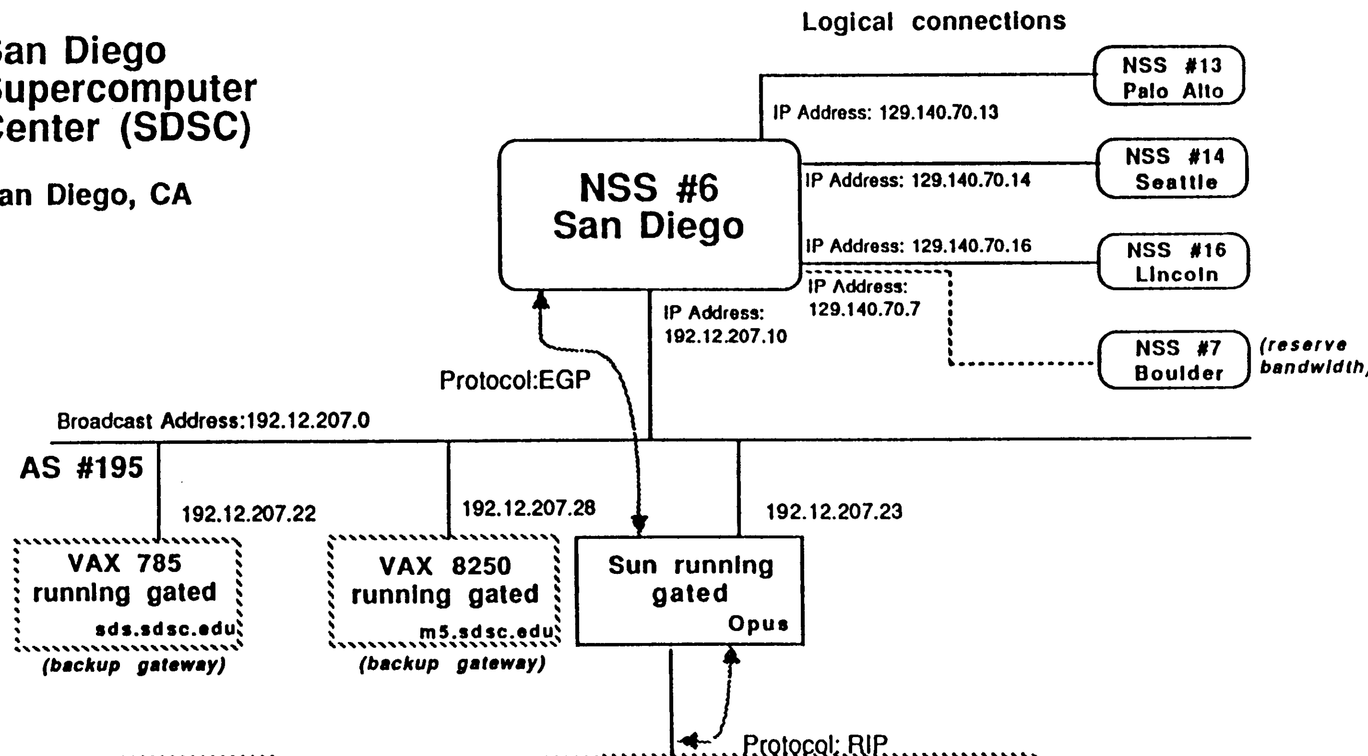

SDSC

Net Number	Net Name	Network Location
31	UCDLA-NET	UC Division of Library Automation
128.54	UCSD	University of California, San Diego, California
128.111	UCSB	University of California, Santa Barbara, California
128.171	HAWAII	University of Hawaii, Honolulu, Hawaii
128.2	UCI-NET	University of California, Irvine, California
129.8	CSUFRESNO	California State University at Fresno, Fresno, California
129.65	CALPOLY	California Polytechnic State University, San Luis Obispo, California
130.15	CSUNET-IP	California State University
192.12.207	SUPERCOMP	San Deigo Supercomputer Center, La Jolla, California
192.31.21	SDSC-APOLLO	San Deigo Supercomputer Center, La Jolla, California
192.31.146	UCR	University of California, Riverside, California
192.31.153	SALKNET	Salk Institute, La Jolla, California
192.35.201	HAWAII-HIG	University of Hawaii, Honolulu, Hawaii
192.35.202	HAWAII-RIT	University of Hawaii, Honolulu, Hawaii

Primary

10-15-88

National Center for Atmospheric Research (NCAR)

Boulder, CO

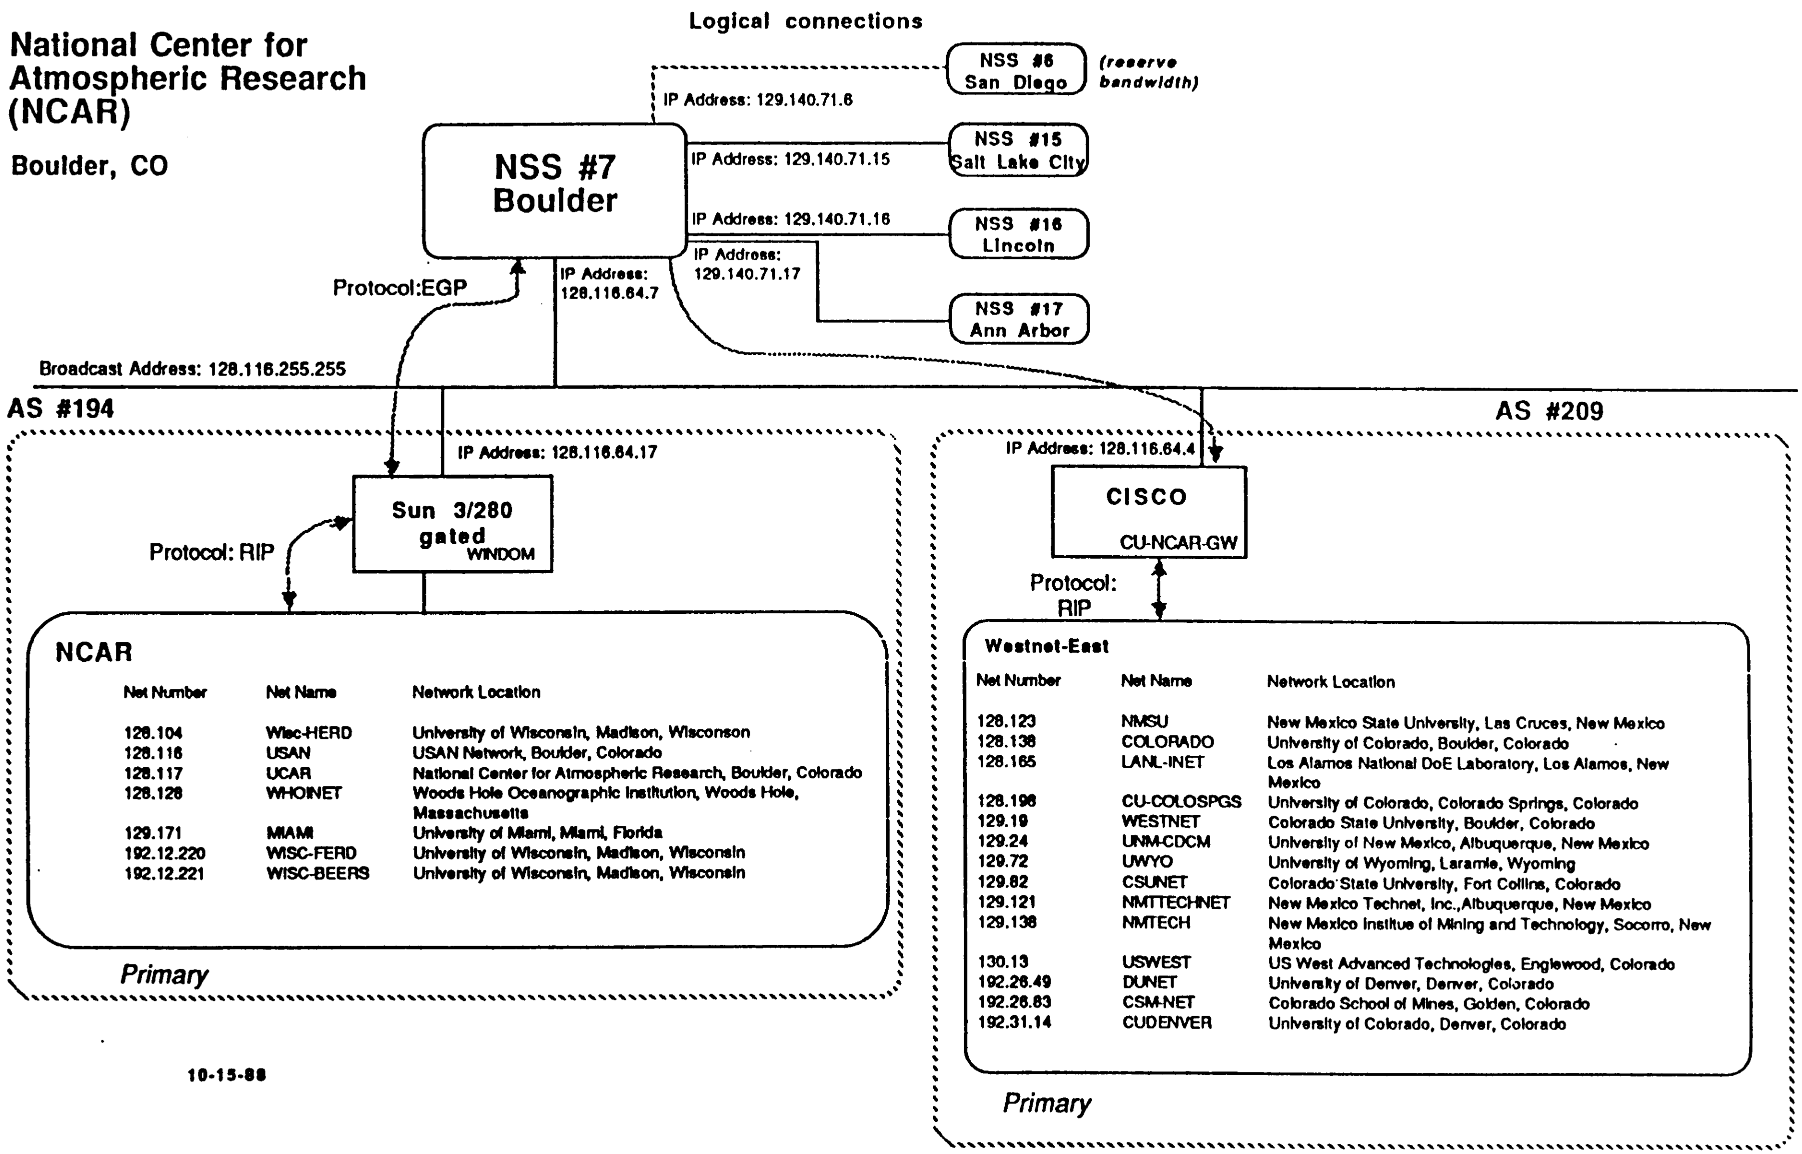

NCAR

Net Number	Net Name	Network Location
128.104	Wisc-HERD	University of Wisconsin, Madison, Wisconson
128.116	USAN	USAN Network, Boulder, Colorado
128.117	UCAR	National Center for Atmospheric Research, Boulder, Colorado
128.128	WHOINET	Woods Hole Oceanographic Institution, Woods Hole, Massachusetts
129.171	MIAMI	University of Miami, Miami, Florida
192.12.220	WISC-FERD	University of Wisconsin, Madison, Wisconsin
192.12.221	WISC-BEERS	University of Wisconsin, Madison, Wisconsin

Westnet-East

Net Number	Net Name	Network Location
128.123	NMSU	New Mexico State University, Las Cruces, New Mexico
128.138	COLORADO	University of Colorado, Boulder, Colorado
128.165	LANL-INET	Los Alamos National DoE Laboratory, Los Alamos, New Mexico
128.198	CU-COLOSPGS	University of Colorado, Colorado Springs, Colorado
129.19	WESTNET	Colorado State University, Boulder, Colorado
129.24	UNM-CDCM	University of New Mexico, Albuquerque, New Mexico
129.72	UWYO	University of Wyoming, Laramie, Wyoming
129.82	CSUNET	Colorado State University, Fort Collins, Colorado
129.121	NMTTECHNET	New Mexico Technet, Inc.,Albuquerque, New Mexico
129.138	NMTECH	New Mexico Institue of Mining and Technology, Socorro, New Mexico
130.13	USWEST	US West Advanced Technologies, Englewood, Colorado
192.26.49	DUNET	University of Denver, Denver, Colorado
192.26.83	CSM-NET	Colorado School of Mines, Golden, Colorado
192.31.14	CUDENVER	University of Colorado, Denver, Colorado

10-15-88

John Von Neumann
National Supercomputer Center

(JVNC)
Princeton, NJ

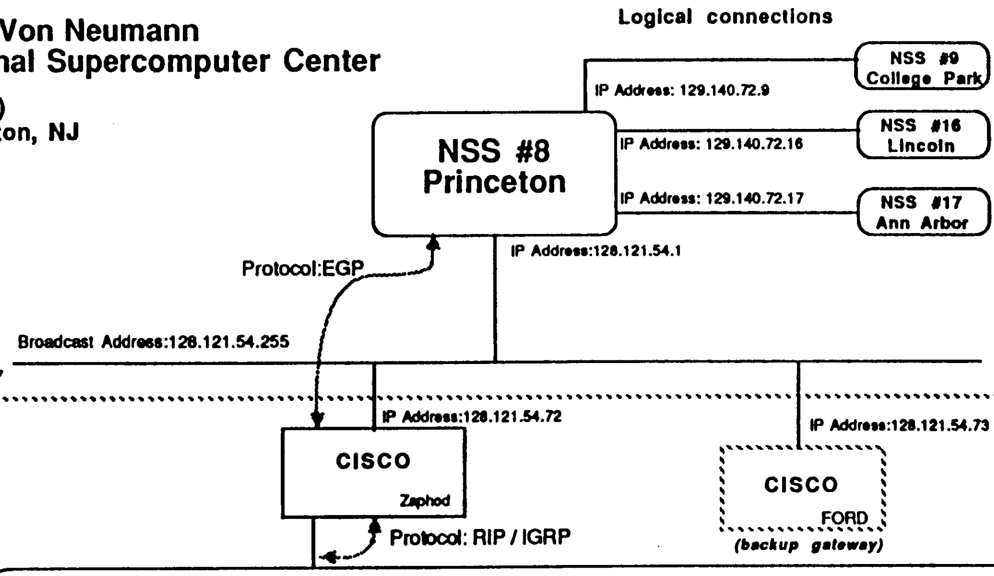

JVNCnet

Net Number	Net Name	Network Location
18	MIT-Temp	Massachusetts Institute of Technology, Boston, Massachusetts
128.6	Rutgers	Rutgers University
128.36	YALE-NET	Yale University, New Haven, Connecticut
128.91	UPENN	University of Pennsylvania, Philadelphia, Pennsylvania
128.93	INRIA-NET	Domaine de Volceau, Rocquencourt, France
128.103	HARV-FIBER	Harvard University, Cambridge, Massachusetts
128.112	PRINCETION	Princeton University, Princeton, New Jersey
128.118	PENN-STATE	Pennsylvania State University, University Park, Pennsylvania
128.119	UMASS-NET	University of Massachusetts, Amherst, Massachusetts (UMass COINS Dept LAN)
128.121	JVNC-NET	John von Neumann National Supercomputer Center
128.148	BROWN-UNIV	Brown University, Providence, Rhode Island
128.18	LEHIGH	Lehigh University, Bethlehem, Pennsylvania
128.197	BU-NET	Boston University, Boston, Massachusetts
128.235	NJIT	New Jersey Institute of Technology, Newark, New Jersey
129.1	NORTHEASTERN-NET	Northeastern University, Boston, Massachusetts
129.2	IRISA-NET	Universite de Rennes, Rennes, France
129.25	DREXEL	Drexel University, Philadelphia, Pennsylvania
129.32	TEMPLE	Temple University, Philadelphia, Pennsylvania
129.133	WESNET	Wesleyan University, Middleton, Connecticut

Net Number	Net Name	Network Location
129.17	DART-ETHER	Dartmouth College, Hanover, New York
129.199	ENS-NET	Ecole Normale Superieure, Paris, France
192.5.60	INRIA-SOPHIA	Domaine de Volceau, Rocquencourt, France
192.5.88	YALE-EE-NET	Yale University, New Haven, Connecticut
192.12.9	YALE-EE2-NET	Yale University, New Haven, Connecticut
192.12.216	STEVENS-TECH	Stevens Institute of Technology, Hoboken, New Jersey
192.16.204	IASNET	Institute for Advanced Study, Princeton, New Jersey
192.26.88	YALE-ENG-NET	Yale University, New Haven, Connecticut
192.26.148	UMDNJ	University of Medicine and Dentistry of New Jersey, Newark, New Jersey
192.31.27	ALTAIR-NET	Domaine de Volceau, Inria, France
192.33.144	CERISI-NET	Institut National de Recherche en Informatique et Automatique, Domaine de Volceau, Rocquencourt, France
192.33.149	CMA-NET	Institut National de Recherce en Informatique et Automatique, Domaine de Volceau, Rocquencourt, France
192.33.170	INRIA-FRANCE	Institut National de Recherce en Informatique et Automatique, Domaine de Volceau, Rocquencout, France

PRIMARY

10-15-88

Suranet

University of Maryland
College Park, Md

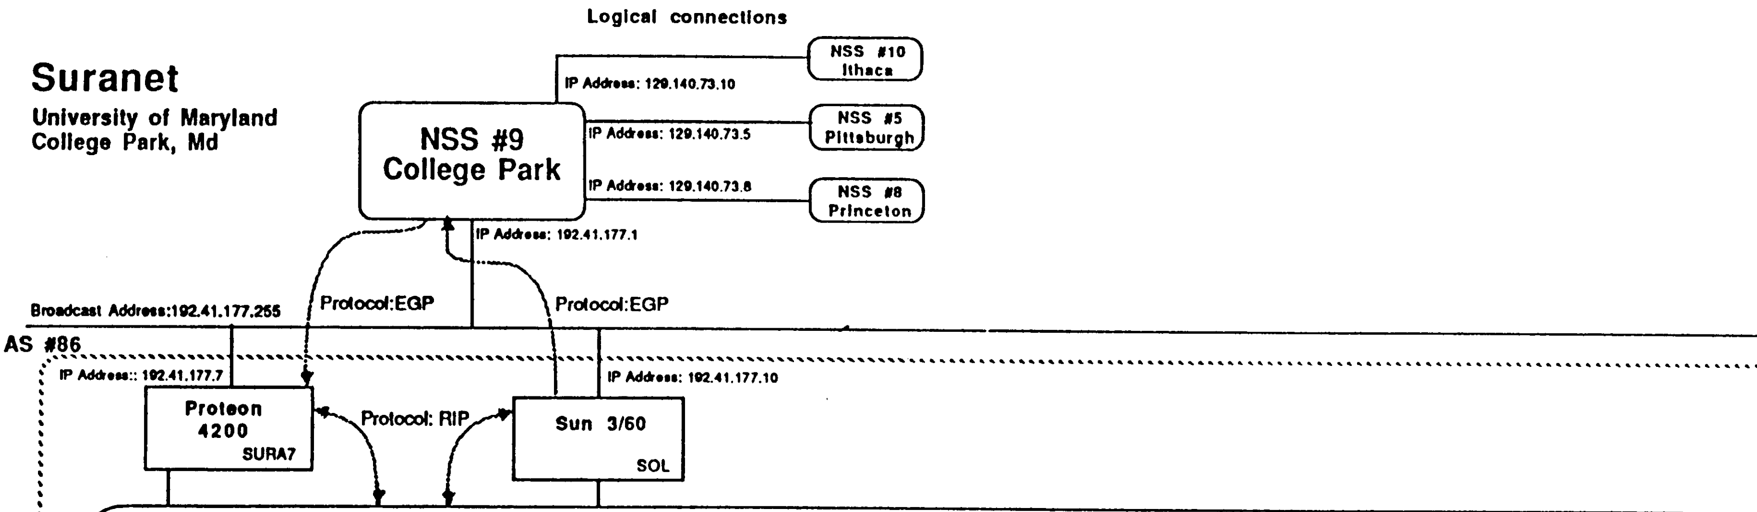

Suranet

Net Number	Net Name	Network Location
128.4	DCNET	DCN
128.8	UMDNET	University of Maryland, College Park, MD
128.8	NRL-LAN	Naval Research Laboratory, Washington D.C.
128.61	GATECH	Georgia Institute of Technology, Atlanta, Georgia
128.82	ODU	Old Dominion University, Norfolk, Virginia
128.109	TUCC-MCNC	Triangle Universities Computation Center (Duke University, North Carolina State University, University of North Carolina) Research Triangle Park, North Carolina
128.14	EMORY-INET	Emory University, Atlanta, Georgia
128.143	VIRGINIA	University of Virginia, Charlottesville, Virginia
128.15	NSF-LAN	National Science Foundation, Washington, D.C.
128.163	UKY	University of Kentucky, Lexington, Kentucky
128.164	GWU-GATE	George Washington University, Washington, D.C.
128.167	SURA	Suranet
128.169	UTK	University of Tennessee, Knoxville, Tennessee
128.172	VCU-LAN	Virginia Commonwealth University, Richmond, Virginia
128.173	VA-TECH	Virginia Polytechnic Institute, Blacksburg, Virginia
128.175	UDELNET	University of Delaware, Newark, Delaware
128.183	GFSC	NASA - Goddard Space Flight Center, Greenbelt, Maryland
128.186	FSU	Florida State University, Tallahassee, Florida
128.192	UGA	University of Georgia, Athens, Georgia
128.22	JHU	John Hopkins University, Baltimore, Maryland
128.227	UFNET	University of Florida, Gainesville, Florida
128.231	NIH	National Institutes of Health, Bethesda, Maryland
128.239	WM-NET	College of William & Mary, Williamsburgh, Virginia
129.6	NBS	National Bureau of Standards, Gaithersburg, Maryland
129.43	NCI-FCRF	National Cancer Institute, Frederick, Maryland
129.57	CEBAF	Continous Electron Beam Accelerator Facility, Newport News, Virginia

Net Number	Net Name	Network Location
129.59	VANDERBILT	Vanderbilt University, Nashville, Tennessee
129.66	ALABAMANET	University of Alabama, Birmingham, Alabama
129.71	WVNET	West Virginia Network for Educational Telecomputing, Morgantown, West Virginia
129.81	TULANE	Tulane University, New Orleans, LA (TCS/Computer lab)
130.11	GEORES	U.S. Geological Survey, Reston, Virginia
130.14	NLM-ETHER1	National Library of Medicine, Bethesda, Maryland
130.18	MSSTATE	Mississippi State University, Mississippi State, Mississippi
130.39	TIGERLAN	Louisiana State University, Baton Rouge, Louisiana
130.207	GIT	Georgia Institute of Technology, Atlanta, GA
192.5.39	UDEL-EECIS	University of Delaware, Newark, Delaware
192.5.45	FCCC	Fox Chase Cancer Center, Philadelphia, Pennsylvania
192.5.57	UDEL-CC	University of Delaware, Newark, Delaware
192.5.82	FSUSTAT	Florida State University, Tallahassee, Florida
192.5.214	DECUANET	Digital Equipment Corporation, Landover, Maryland
192.5.215	MASONNET	George Mason University, Fairfax, Virginia
192.5.219	CLEMSON	Clemson University, Clemson, South Carolina
192.16.175	GUACC	Georgetown University, Washington, D.C.
192.16.176	LSUNET	Louisiana State University, Baton Rouge, Louisiana
192.26.10	GALLAUDET	Gallaudet University, Washington, D.C.
192.26.11	NRL-HUBNET1	Naval Research Laboratory, Washington, D.C.
192.26.12	NRL-NUBNET2	Naval Research Laboratory, Washington, D.C.
192.26.13	NRL-HUBNET3	Naval Research Laboratory, Washington, D.C.
192.26.14	NRL-HUBNET4	Naval Research Laboratory, Washington, D.C.
192.26.17	NRL-HUBNET7	Naval Research Laboratory, Washington, D.C.
192.26.26	NRL-FIBER	Naval Research Laboratory, Washington, D.C.
192.31.192	SUPER	IDA/Supercomputing Research Center, Lanham, Maryland
192.33.115	NRAO-CV	National Radio Astronomy Observatory, Charlottesville, Virginia
192.41.177	SURA-NOC	University of Maryland, College Park, Maryland

Primary

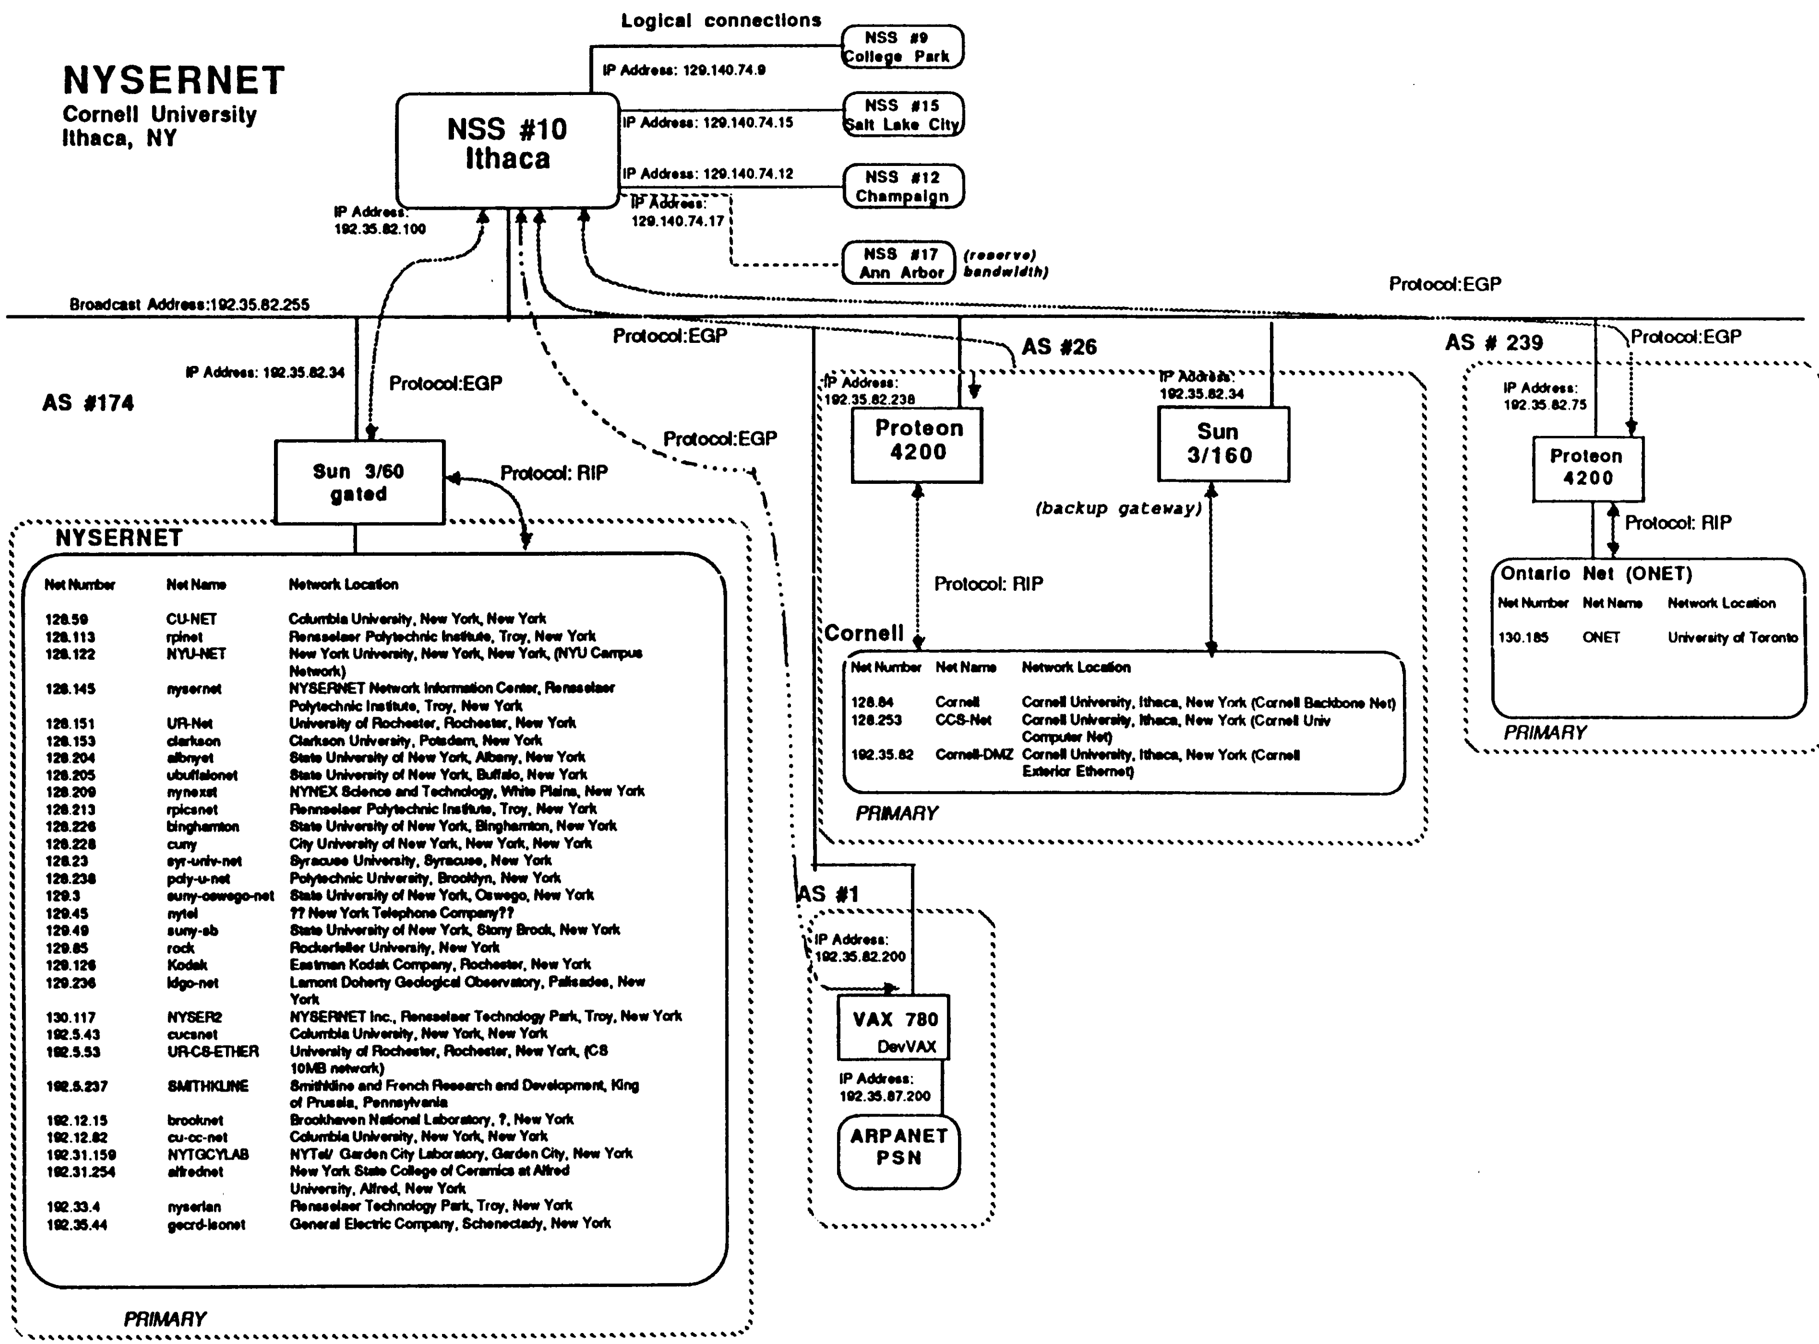
NYSERNET
Cornell University
Ithaca, NY
Logical connections
NSS #10
Ithaca
IP Address: 192.35.82.100
IP Address: 129.140.74.9
NSS #9
College Park
IP Address: 129.140.74.15
NSS #15
Salt Lake City
IP Address: 129.140.74.12
NSS #12
Champaign
IP Address: 129.140.74.17
NSS #17
Ann Arbor
(reserve) bandwidth)
Broadcast Address:192.35.82.255
Protocol:EGP
Protocol:EGP
Protocol:EGP
Protocol:EGP
Protocol:EGP
AS #174
IP Address: 192.35.82.34
Sun 3/60
gated
Protocol: RIP
NYSERNET
Net Number Net Name Network Location
128.59 CU-NET Columbia University, New York, New York
128.113 rpinet Rensselaer Polytechnic Institute, Troy, New York
128.122 NYU-NET New York University, New York, New York, (NYU Campus Network)
128.145 nysernet NYSERNET Network Information Center, Rensselaer Polytechnic Institute, Troy, New York
128.151 UR-Net University of Rochester, Rochester, New York
128.153 clarkson Clarkson University, Potsdam, New York
128.204 albnyet State University of New York, Albany, New York
128.205 ubuffalonet State University of New York, Buffalo, New York
128.209 nynexst NYNEX Science and Technology, White Plains, New York
128.213 rpicsnet Rensselaer Polytechnic Institute, Troy, New York
128.226 binghamton State University of New York, Binghamton, New York
128.228 cuny City University of New York, New York, New York
128.23 syr-univ-net Syracuse University, Syracuse, New York
128.238 poly-u-net Polytechnic University, Brooklyn, New York
129.3 suny-oswego-net State University of New York, Oswego, New York
129.45 nytel ?? New York Telephone Company??
129.49 suny-sb State University of New York, Stony Brook, New York
129.85 rock Rockefeller University, New York
129.126 Kodak Eastman Kodak Company, Rochester, New York
129.236 ldgo-net Lamont Doherty Geological Observatory, Palisades, New York
130.117 NYSER2 NYSERNET Inc., Rensselaer Technology Park, Troy, New York
192.5.43 cucsnet Columbia University, New York, New York
192.5.53 UR-CS-ETHER University of Rochester, Rochester, New York, (CS 10MB network)
192.5.237 SMITHKLINE Smithkline and French Research and Development, King of Prussia, Pennsylvania
192.12.15 brooknet Brookhaven National Laboratory, ?, New York
192.12.82 cu-cc-net Columbia University, New York, New York
192.31.159 NYTGCYLAB NYTel/ Garden City Laboratory, Garden City, New York
192.31.254 alfrednet New York State College of Ceramics at Alfred University, Alfred, New York
192.33.4 nyserlan Rensselaer Technology Park, Troy, New York
192.35.44 gecrd-lsonet General Electric Company, Schenectady, New York
PRIMARY
AS #26
IP Address: 192.35.82.238
Proteon
4200
IP Address: 192.35.82.34
Sun
3/160
(backup gateway)
Protocol: RIP
Cornell
Net Number Net Name Network Location
128.84 Cornell Cornell University, Ithaca, New York (Cornell Backbone Net)
128.253 CCS-Net Cornell University, Ithaca, New York (Cornell Univ Computer Net)
192.35.82 Cornell-DMZ Cornell University, Ithaca, New York (Cornell Exterior Ethernet)
PRIMARY
AS # 239
IP Address: 192.35.82.75
Proteon
4200
Protocol: RIP
Ontario Net (ONET)
Net Number Net Name Network Location
130.185 ONET University of Toronto
PRIMARY
AS #1
IP Address: 192.35.82.200
VAX 780
DevVAX
IP Address: 192.35.87.200
ARPANET
PSN

10-16-88

SESQUINET

Rice University
Houston, TX

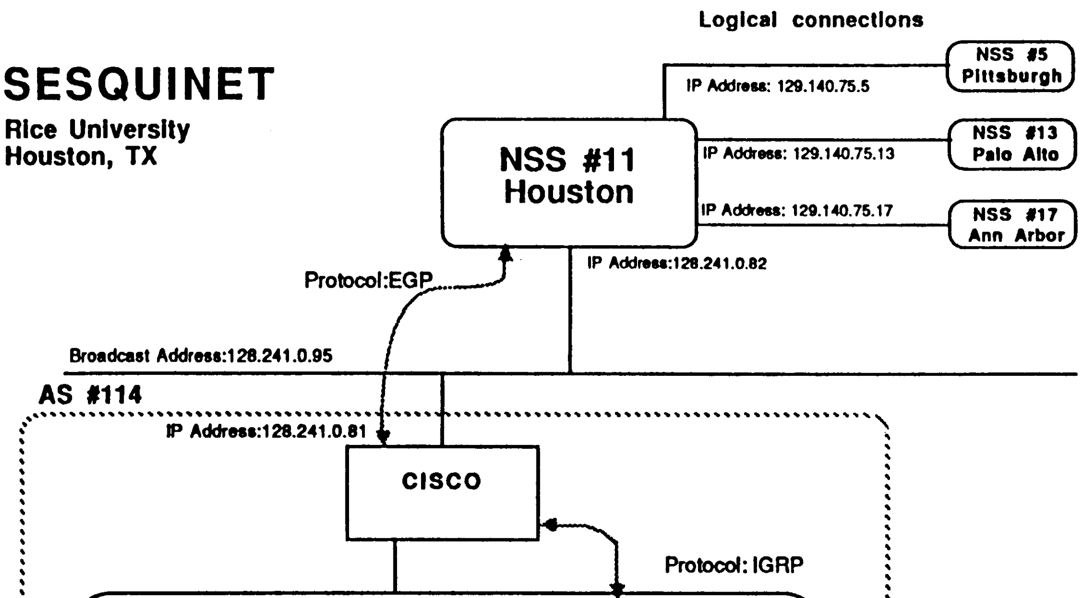

SESQUINET

Net Number	Net Name	Network Location
128.42	Rice-Net	Rice University, Houston, Texas
128.83	UTAustin	University of Texas, Austin, Texas
128.194	TAMU-Net	Texas A&M, College Station, Texas
128.241	Sesquinet	Rice University, Houston, Texas
128.249	TMC-Net	Texas Medical Center, Houston, Texas
129.7	UH-Net	University of Houston, Houston, Texas
129.106	UTHOUSTON	University of Texas, Houston, Texas
129.107	UTArlington	University of Texas at Arlington, Arlington, Texas
129.108	UTEIPaso	University of Texas at El Paso, El Paso, Texas
129.109	UTGalveston	University of Texas at Galveston, Galveston, Texas
129.11	UTDallas	University of Texas, Richardson, Texas
129.111	UTHCSA	University of Texas, Health Science Center, Dallas, Texas
129.112	UTSWMED	University of Texas, Southwestern Med Center, Dallas, Texas
129.113	UTPBasin	University of Texas of the Permian Basin, Odessa, Texas
129.114	UTCCSPRD	University of Texas, MD Anderson Cancer Center,
129.116	CHPCHYPERHOSE	University of Texas, Center for High Performance Computing
129.118	TXTech	Texas Tech University, Lubbock, Texas
129.12	UNTexas	North texas University, Denton, Texas
129.207	PVAMU-Net	Prairie View A&M University, Prairie View, Texas
192.6.201	UTSanAntonio	University of Texas at San Antonio, San Antonio, Texas
192.16.72	UTCHPC	University of Texas, Austin, Texas
192.31.87	HARC-Net	Houston Area Research Center, The Woodlands, Texas
192.31.88	BCMTech-Net	Baylor College of Medicine, The Woodlands, Texas
192.31.101	TSU-Net	Texas Southern University, Houston, Texas
192.31.152	rosenet	Roserra Inc., Houston, Texas

PRIMARY

10-15-88

SESQUINET

Rice University
Houston, TX

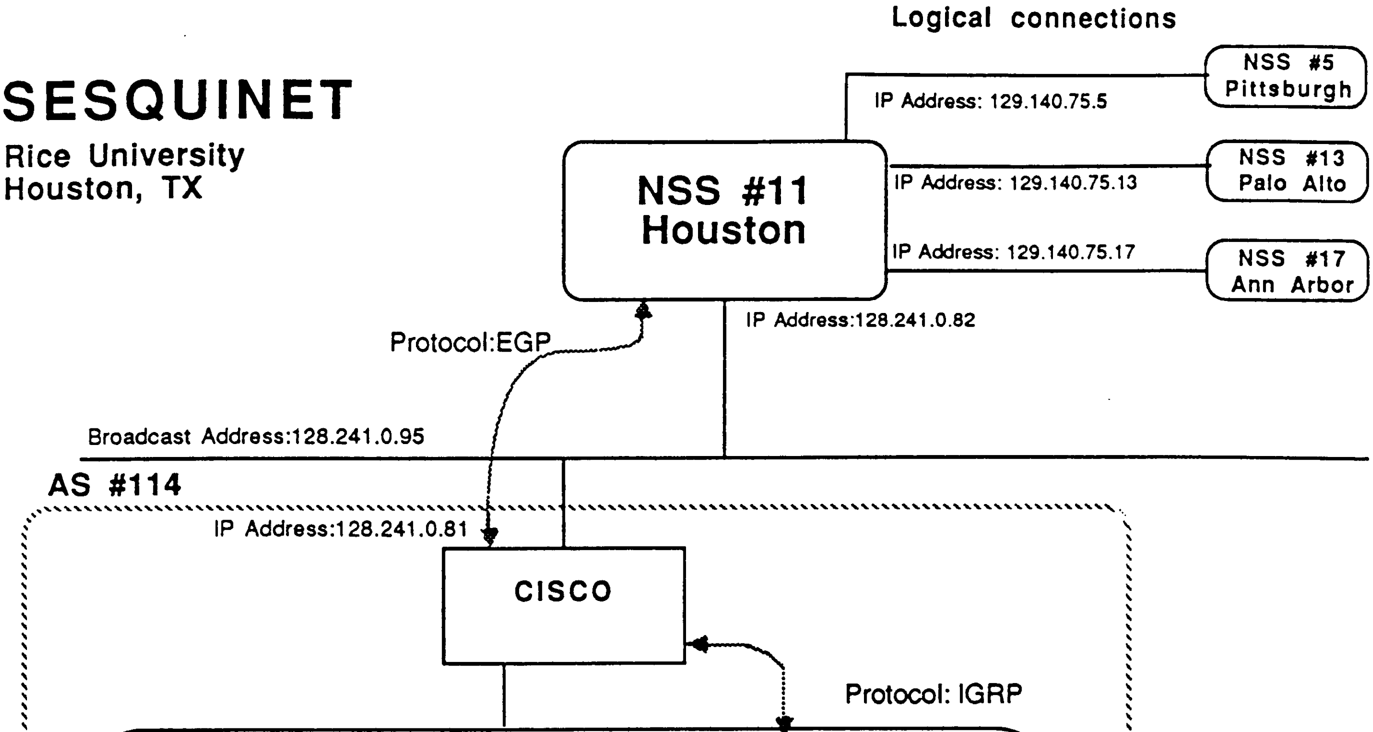

SESQUINET

Net Number	Net Name	Network Location
128.42	Rice-Net	Rice University, Houston, Texas
128.83	UTAustin	University of Texas, Austin, Texas
128.194	TAMU-Net	Texas A&M, College Station, Texas
128.241	Sesquinet	Rice University, Houston, Texas
128.249	TMC-Net	Texas Medical Center, Houston, Texas
129.7	UH-Net	University of Houston, Houston, Texas
129.106	UTHOUSTON	University of Texas, Houston, Texas
129.107	UTArlington	University of Texas at Arlington, Arlington, Texas
129.108	UTEIPaso	University of Texas at El Paso, El Paso, Texas
129.109	UTGalveston	University of Texas at Galveston, Galveston, Texas
129.11	UTDallas	University of Texas, Richardson, Texas
129.111	UTHCSA	University of Texas, Health Science Center, Dallas, Texas
129.112	UTSWMED	University of Texas, Southwestern Med Center, Dallas, Texas
129.113	UTPBasin	University of Texas of the Permian Basin, Odessa, Texas
129.114	UTCCSPRD	University of Texas, MD Anderson Cancer Center,
129.116	CHPCHYPERHOSE	University of Texas, Center for High Performance Computing
129.118	TXTech	Texas Tech University, Lubbock, Texas
129.12	UNTexas	North texas University, Denton, Texas
129.207	PVAMU-Net	Prairie View A&M University, Prairie View, Texas
192.6.201	UTSanAntonio	University of Texas at San Antonio, San Antonio, Texas
192.16.72	UTCHPC	University of Texas, Austin, Texas
192.31.87	HARC-Net	Houston Area Research Center, The Woodlands, Texas
192.31.88	BCMTech-Net	Baylor College of Medicine, The Woodlands, Texas
192.31.101	TSU-Net	Texas Southern University, Houston, Texas
192.31.152	rosenet	Roserra Inc., Houston, Texas

PRIMARY

10-15-88

National Center for Supercomputer Applications (NCSA)

University of Illinois-
Urbana/Champaign
Champaign, IL

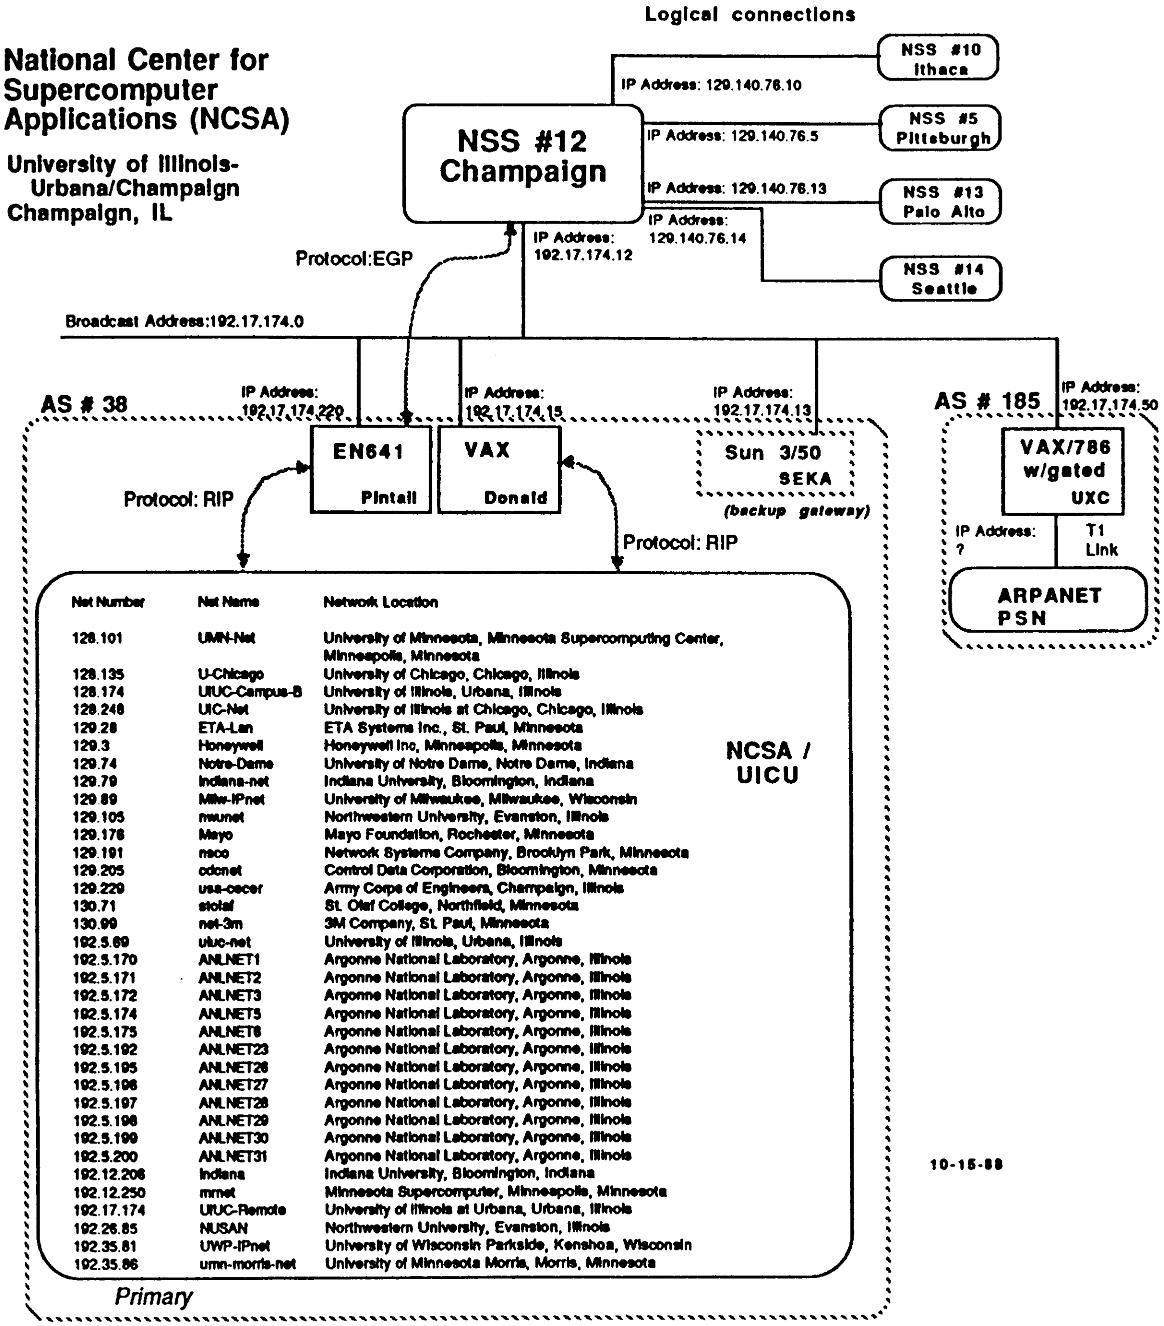

Net Number	Net Name	Network Location
128.101	UMN-Net	University of Minnesota, Minnesota Supercomputing Center, Minneapolis, Minnesota
128.135	U-Chicago	University of Chicago, Chicago, Illinois
128.174	UIUC-Campus-B	University of Illinois, Urbana, Illinois
128.248	UIC-Net	University of Illinois at Chicago, Chicago, Illinois
129.28	ETA-Lan	ETA Systems Inc., St. Paul, Minnesota
129.3	Honeywell	Honeywell Inc, Minneapolis, Minnesota
129.74	Notre-Dame	University of Notre Dame, Notre Dame, Indiana
129.79	Indiana-net	Indiana University, Bloomington, Indiana
129.89	Milw-IPnet	University of Milwaukee, Milwaukee, Wisconsin
129.105	nwunet	Northwestern University, Evanston, Illinois
129.176	Mayo	Mayo Foundation, Rochester, Minnesota
129.191	nsco	Network Systems Company, Brooklyn Park, Minnesota
129.205	cdcnet	Control Data Corporation, Bloomington, Minnesota
129.229	usa-cecer	Army Corps of Engineers, Champaign, Illinois
130.71	stolaf	St. Olaf College, Northfield, Minnesota
130.99	net-3m	3M Company, St. Paul, Minnesota
192.5.69	uiuc-net	University of Illinois, Urbana, Illinois
192.5.170	ANLNET1	Argonne National Laboratory, Argonne, Illinois
192.5.171	ANLNET2	Argonne National Laboratory, Argonne, Illinois
192.5.172	ANLNET3	Argonne National Laboratory, Argonne, Illinois
192.5.174	ANLNET5	Argonne National Laboratory, Argonne, Illinois
192.5.175	ANLNET6	Argonne National Laboratory, Argonne, Illinois
192.5.192	ANLNET23	Argonne National Laboratory, Argonne, Illinois
192.5.195	ANLNET26	Argonne National Laboratory, Argonne, Illinois
192.5.196	ANLNET27	Argonne National Laboratory, Argonne, Illinois
192.5.197	ANLNET28	Argonne National Laboratory, Argonne, Illinois
192.5.198	ANLNET29	Argonne National Laboratory, Argonne, Illinois
192.5.199	ANLNET30	Argonne National Laboratory, Argonne, Illinois
192.5.200	ANLNET31	Argonne National Laboratory, Argonne, Illinois
192.12.206	Indiana	Indiana University, Bloomington, Indiana
192.12.250	mmet	Minnesota Supercomputer, Minneapolis, Minnesota
192.17.174	UIUC-Remote	University of Illinois at Urbana, Urbana, Illinois
192.26.85	NUSAN	Northwestern University, Evanston, Illinois
192.35.81	UWP-IPnet	University of Wisconsin Parkside, Kenshoa, Wisconsin
192.35.86	umn-morris-net	University of Minnesota Morris, Morris, Minnesota

10-15-88

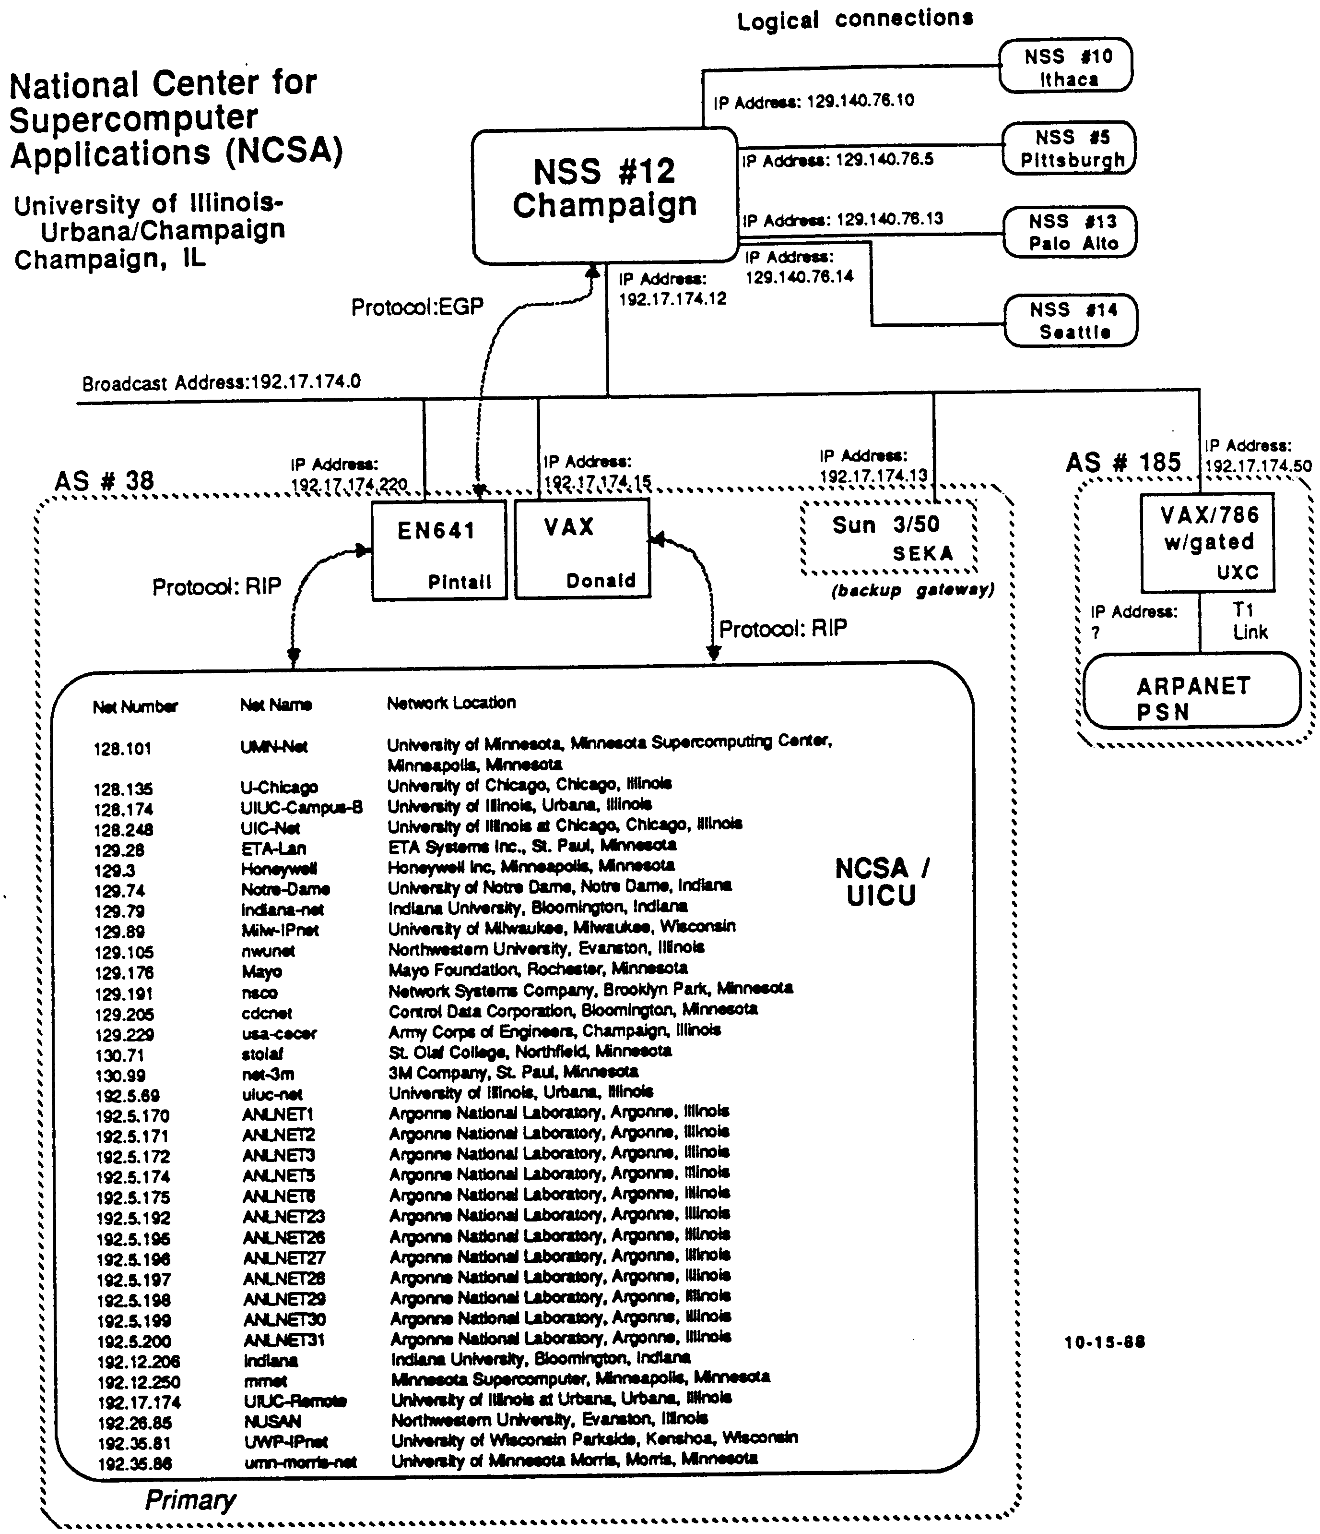

Net Number	Net Name	Network Location
128.101	UMN-Net	University of Minnesota, Minnesota Supercomputing Center, Minneapolis, Minnesota
128.135	U-Chicago	University of Chicago, Chicago, Illinois
128.174	UIUC-Campus-B	University of Illinois, Urbana, Illinois
128.248	UIC-Net	University of Illinois at Chicago, Chicago, Illinois
129.28	ETA-Lan	ETA Systems Inc., St. Paul, Minnesota
129.3	Honeywell	Honeywell Inc, Minneapolis, Minnesota
129.74	Notre-Dame	University of Notre Dame, Notre Dame, Indiana
129.79	Indiana-net	Indiana University, Bloomington, Indiana
129.89	Milw-IPnet	University of Milwaukee, Milwaukee, Wisconsin
129.105	nwunet	Northwestern University, Evanston, Illinois
129.176	Mayo	Mayo Foundation, Rochester, Minnesota
129.191	nsco	Network Systems Company, Brooklyn Park, Minnesota
129.205	cdcnet	Control Data Corporation, Bloomington, Minnesota
129.229	usa-cecer	Army Corps of Engineers, Champaign, Illinois
130.71	stolaf	St. Olaf College, Northfield, Minnesota
130.99	net-3m	3M Company, St. Paul, Minnesota
192.5.69	uiuc-net	University of Illinois, Urbana, Illinois
192.5.170	ANLNET1	Argonne National Laboratory, Argonne, Illinois
192.5.171	ANLNET2	Argonne National Laboratory, Argonne, Illinois
192.5.172	ANLNET3	Argonne National Laboratory, Argonne, Illinois
192.5.174	ANLNET5	Argonne National Laboratory, Argonne, Illinois
192.5.175	ANLNET6	Argonne National Laboratory, Argonne, Illinois
192.5.192	ANLNET23	Argonne National Laboratory, Argonne, Illinois
192.5.195	ANLNET26	Argonne National Laboratory, Argonne, Illinois
192.5.196	ANLNET27	Argonne National Laboratory, Argonne, Illinois
192.5.197	ANLNET28	Argonne National Laboratory, Argonne, Illinois
192.5.198	ANLNET29	Argonne National Laboratory, Argonne, Illinois
192.5.199	ANLNET30	Argonne National Laboratory, Argonne, Illinois
192.5.200	ANLNET31	Argonne National Laboratory, Argonne, Illinois
192.12.206	Indiana	Indiana University, Bloomington, Indiana
192.12.250	mmet	Minnesota Supercomputer, Minneapolis, Minnesota
192.17.174	UIUC-Remote	University of Illinois at Urbana, Urbana, Illinois
192.26.85	NUSAN	Northwestern University, Evanston, Illinois
192.35.81	UWP-IPnet	University of Wisconsin Parkside, Kenshoa, Wisconsin
192.35.86	umn-morris-net	University of Minnesota Morris, Morris, Minnesota

BARRNET

Stanford University
Palo Alto, CA

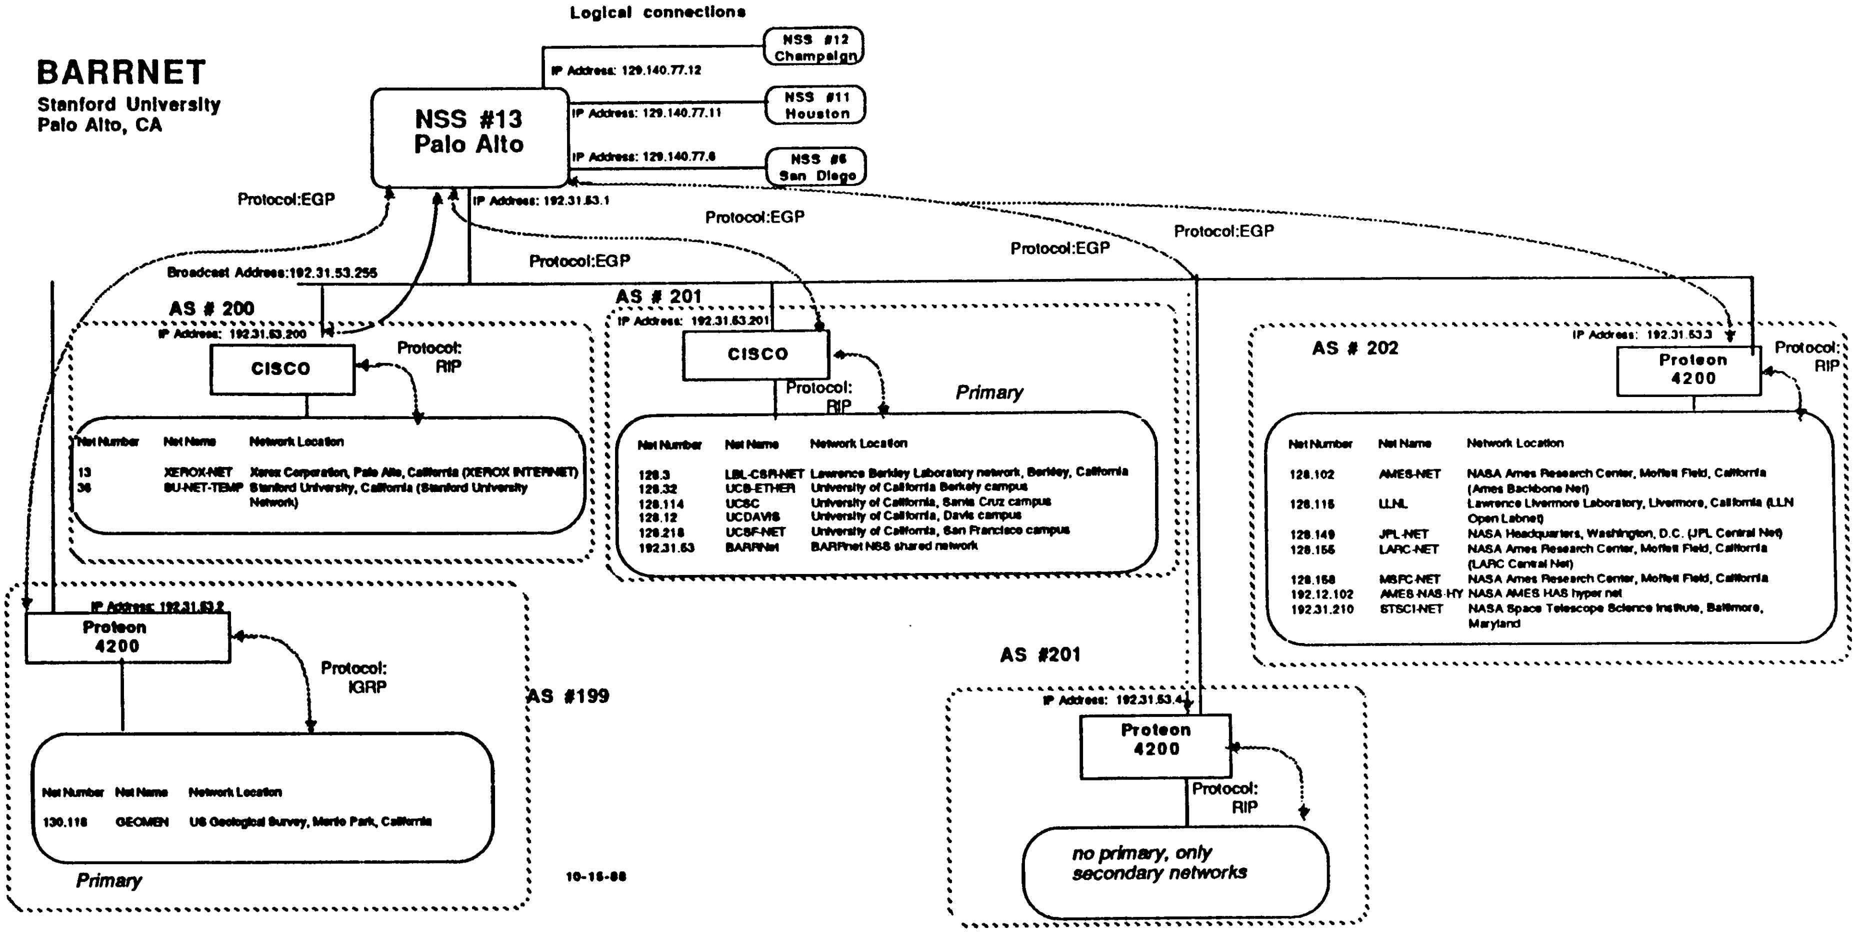

NorthWestNet

University of Washington
Seattle, WA

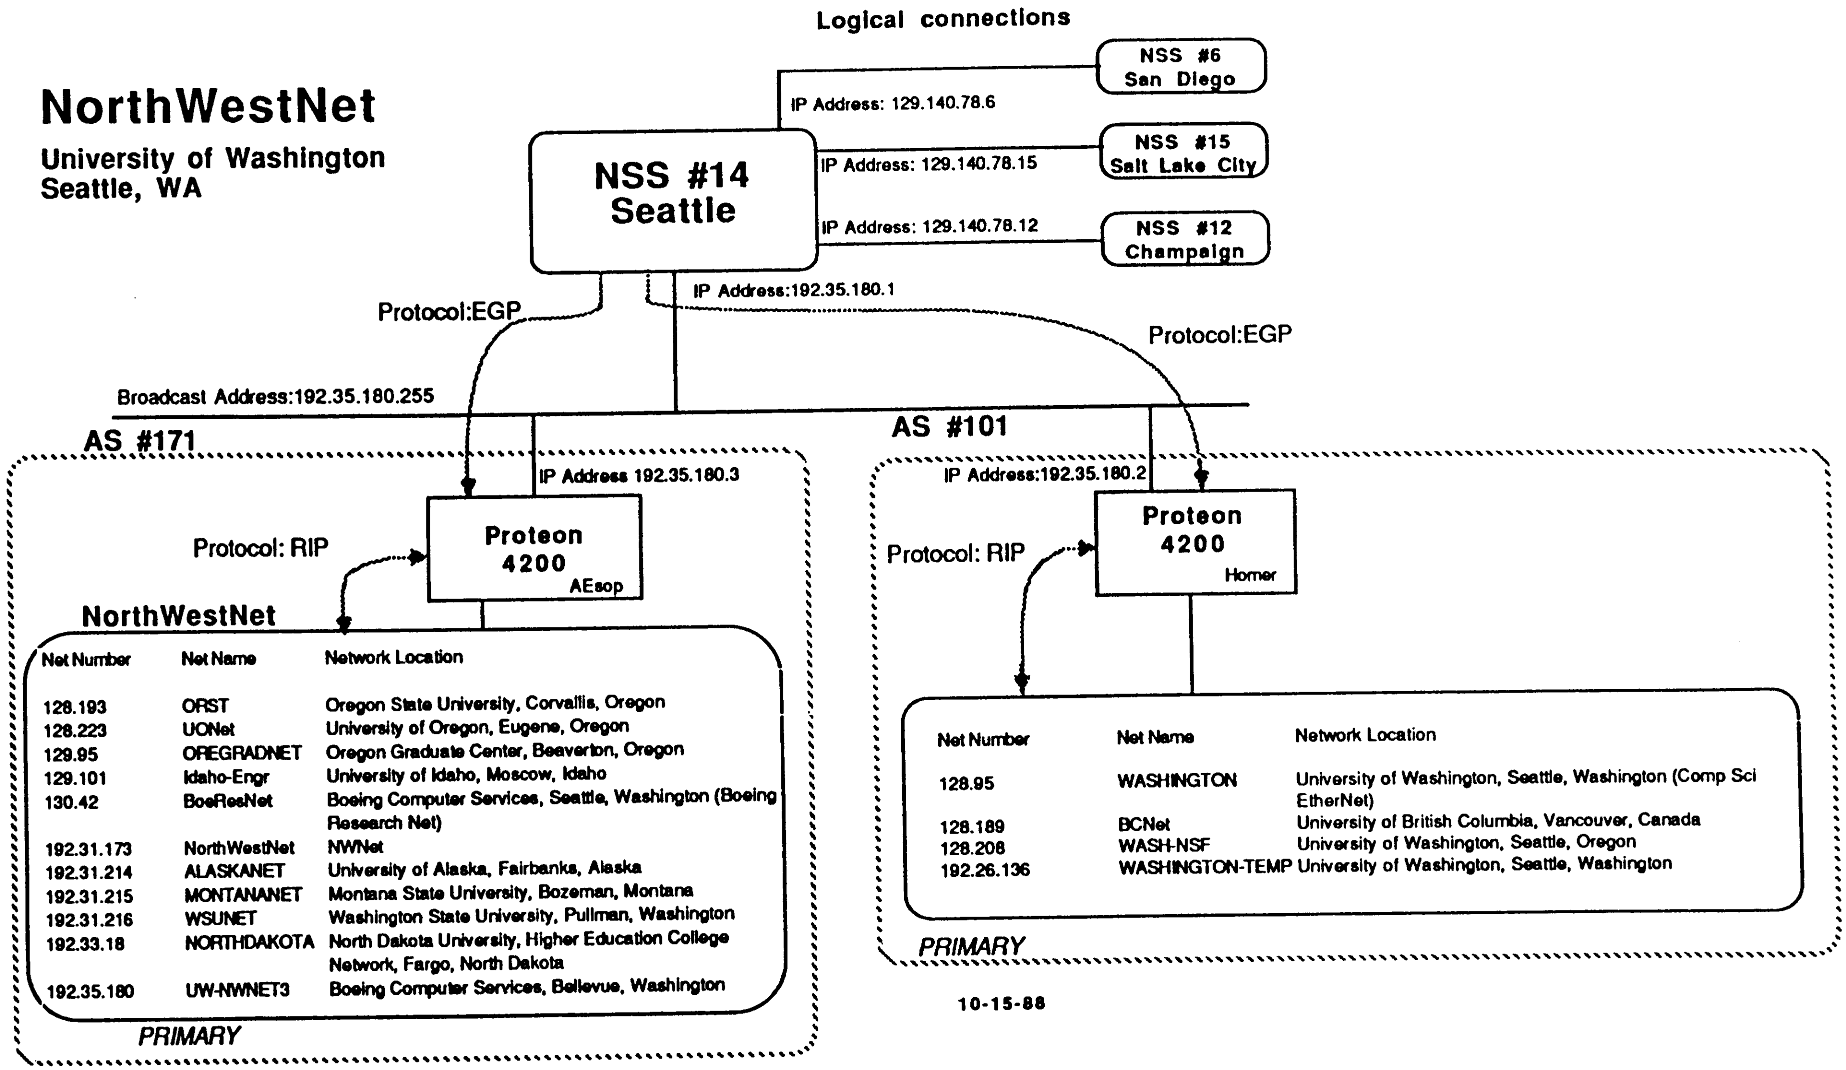

NorthWestNet (AS #171)

Net Number	Net Name	Network Location
128.193	ORST	Oregon State University, Corvallis, Oregon
128.223	UONet	University of Oregon, Eugene, Oregon
129.95	OREGRADNET	Oregon Graduate Center, Beaverton, Oregon
129.101	Idaho-Engr	University of Idaho, Moscow, Idaho
130.42	BoeResNet	Boeing Computer Services, Seattle, Washington (Boeing Research Net)
192.31.173	NorthWestNet	NWNet
192.31.214	ALASKANET	University of Alaska, Fairbanks, Alaska
192.31.215	MONTANANET	Montana State University, Bozeman, Montana
192.31.216	WSUNET	Washington State University, Pullman, Washington
192.33.18	NORTHDAKOTA	North Dakota University, Higher Education College Network, Fargo, North Dakota
192.35.180	UW-NWNET3	Boeing Computer Services, Bellevue, Washington

PRIMARY

(AS #101)

Net Number	Net Name	Network Location
128.95	WASHINGTON	University of Washington, Seattle, Washington (Comp Sci EtherNet)
128.189	BCNet	University of British Columbia, Vancouver, Canada
128.208	WASH-NSF	University of Washington, Seattle, Oregon
192.26.136	WASHINGTON-TEMP	University of Washington, Seattle, Washington

PRIMARY

10-15-88

Westnet

University of Utah
Salt Lake City, UT

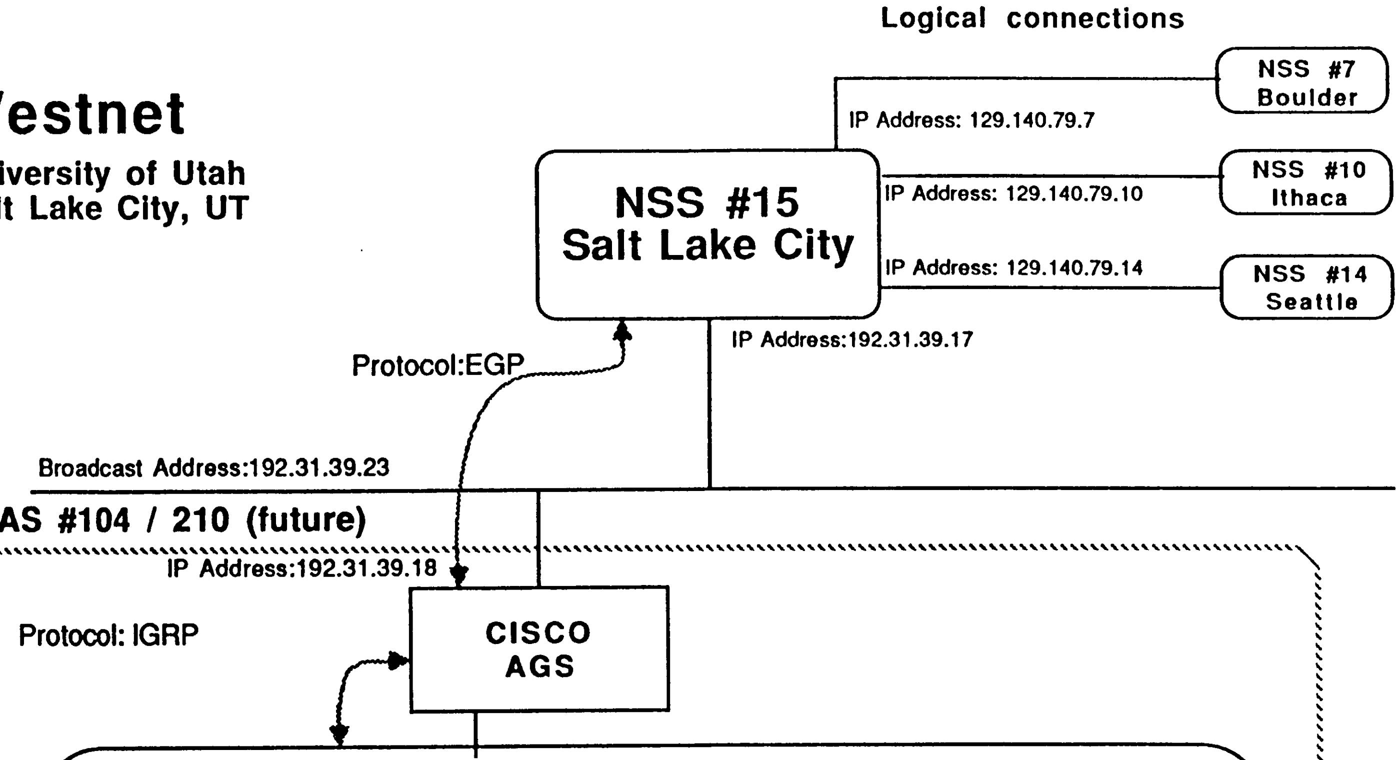

Westnet

Net Number	Net Name	Network Location
128.11	Utah-Net	University of Utah, Salt Lake City, Utah
128.119	UMASS-NET	University of Massachusetts, Amherst, Massachusetts (UMass COINS Dept LAN)
128.187	BYU-Net	Brigham Young University, Provo, Utah
128.196	UNIV-ARIZ	University of Arizona, Tuscon, Arizona
129.123	USU	Utah State University, Logan, Utah
129.219	ASU-NET	Arizona State University, Tempe, Arizona
192.12.56	Utah-Ap-Net	University of Utah, Salt Lake City, Utah (Utah Apollo Ring Net)
192.31.28	STEWARD-OBS	Unviversity of Arizona, Tuscon, Arizona
192.31.39		WestNet-West subnet
192.31.165	NOAO-TUCSON	University of Arizona, Tuscon, Arizona
192.33.140	SAO-NET	University of Arizona, Tuscon, Arizona
192.35.195	ARIZONA-CMI	University of Arizona, Tuscon, Arizona
192.35.203	ece-ariz	University of Arizona, Tuscon, Arizona
192.42.108	UMC-NET	Utah-Michigan-Chicago Cosmic Ray Experiment, Dugway, Utah (Fly's eye)

10-15-88

PRIMARY

MIDNET

University of Nebraska-Lincoln
Lincoln, NE

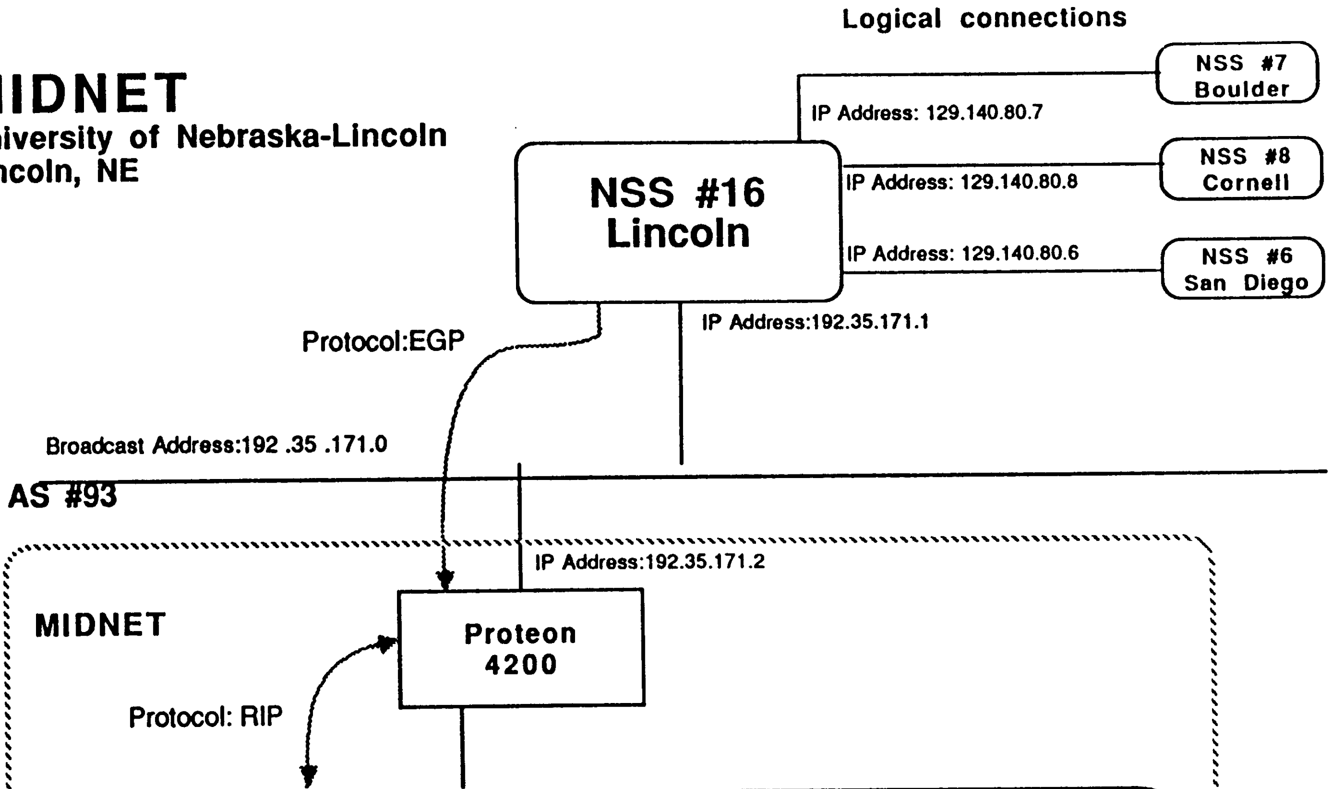

Net Number	Net Name	Network Location
128.206	MONET	University of Missouri, Columbia, Missouri
128.242	MIDNET	Midwest Regional Network, Champaign, Illinois
128.252	WASHINGTON-U	Washington University, St. Louis, Missouri
128.255	UIOWA	University of Iowa, Iowas City, Iowa
129.15	UOKNOR	University of Oklahoma, Norman, Oklahoma
129.93	HUSKERNET	University of Nebraska, Lincoln, Nebraska
129.13	KSUNET	Kansas State University, Manhattan, Kansas
129.186	CYCLONENT	Iowa State University, Ames, Iowa
129.237	JAYHAWKNET	University of Kansas, Lawrence, Kansas
129.244	KehNet	University of Tulsa, Tulsa, Oklahoma
192.31.83	OSUNET	Oklahoma State University, Stillwater, Oklahoma
192.31.86	RAZORNET	University of Arkansas, Fayetteville, Arkansas
192.35.96	UOKECNA	University of Oklahoma, Norman, Oklahoma
192.35.97	UOKECNB	University of Oklahoma, Norman, Oklahoma
192.35.98	UOKECNC	University of Oklahoma, Norman, Oklahoma
192.35.171	midnet-net2	University of Nebraska, Lincoln, Nebraska

PRIMARY

10-15-88

Merit Computer Network

(University of Michigan)

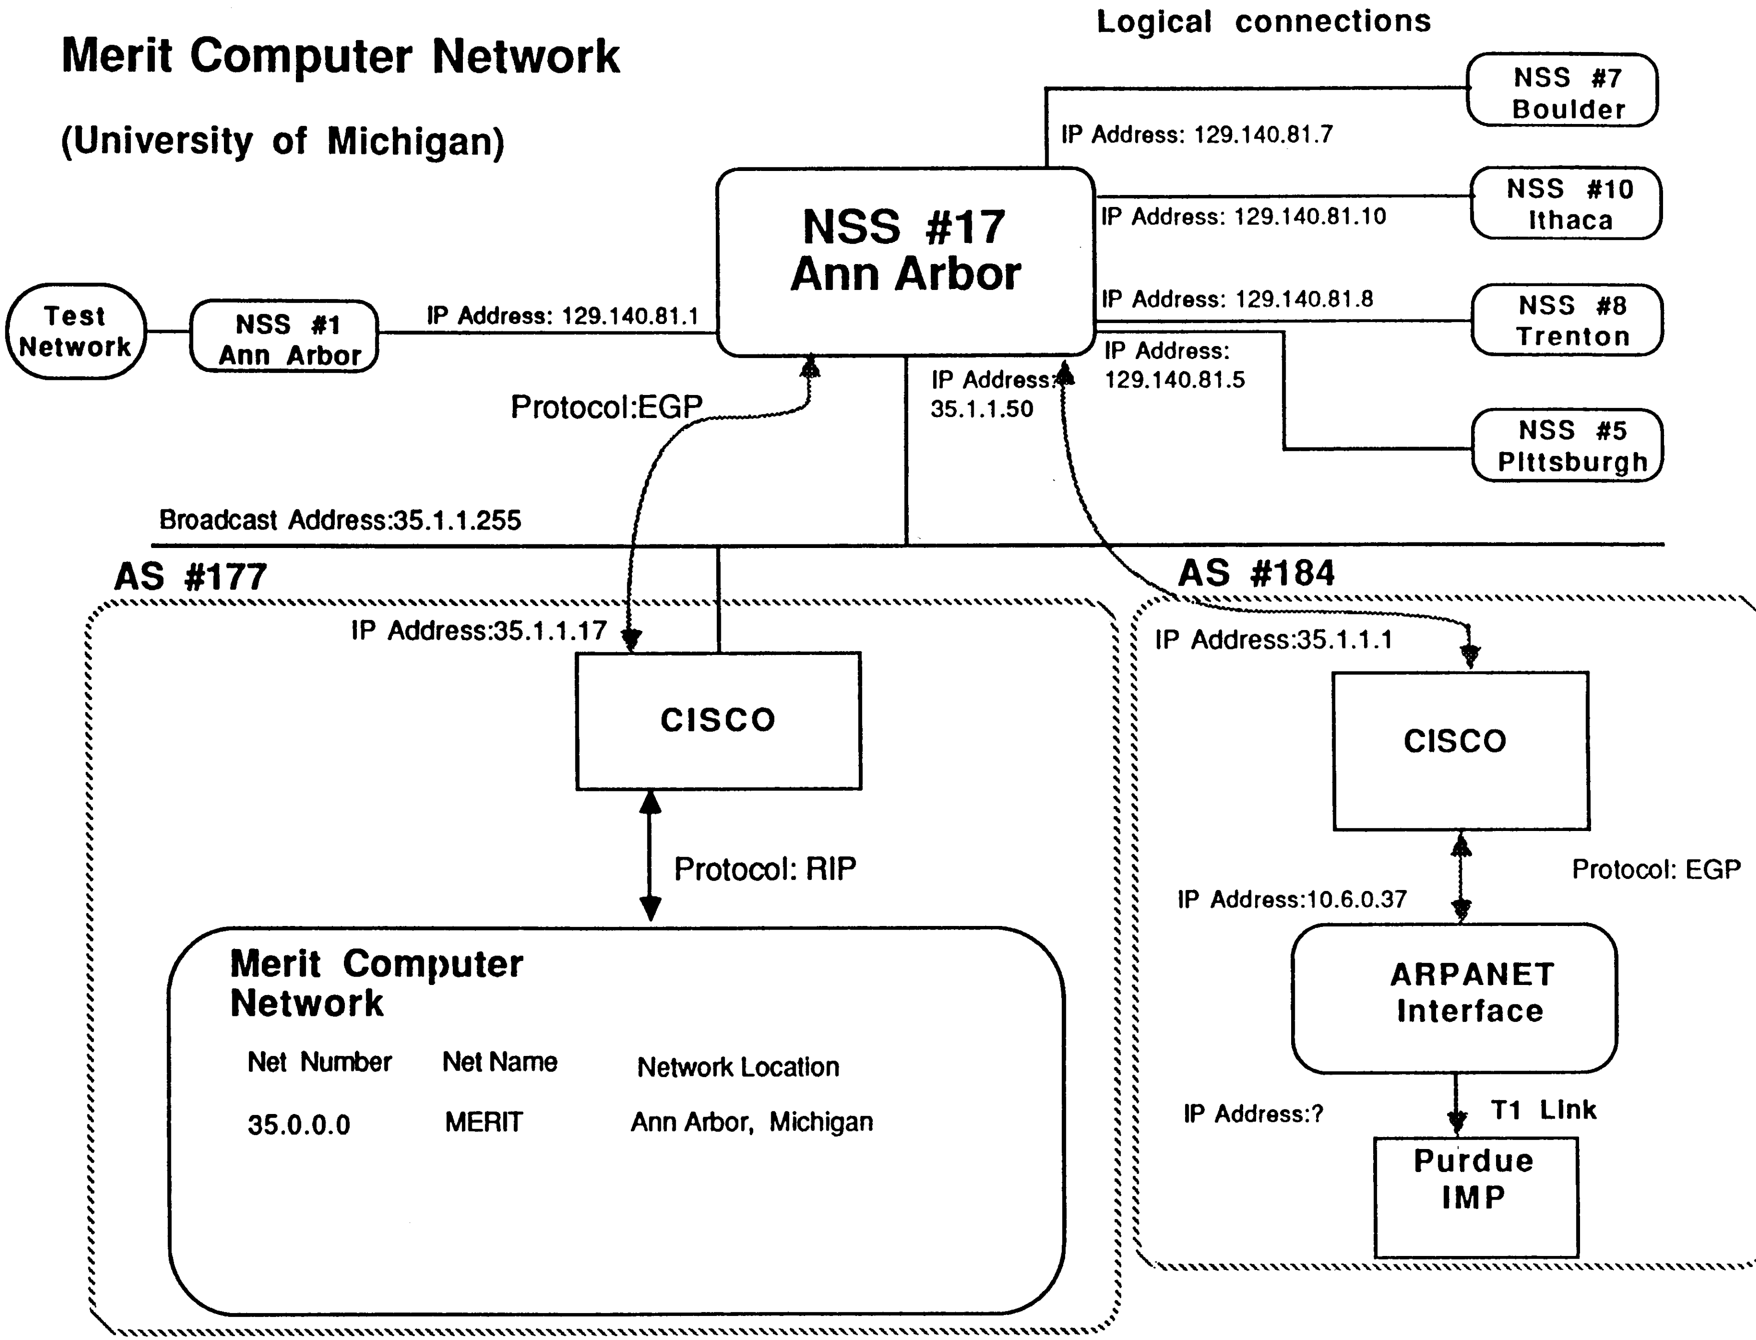

10-16-88

PRIMARY

Merit
Computer
Network

MANAGEMENT AND OPERATION OF THE NSFNET BACKBONE NETWORK PROJECT

Monthly Report
August 1988
Merit Computer Network

The NSFNET Backbone Network Project, is managed and coordinated by The Merit Computer Network under sponsorship of the National Science Foundation, Agreement No. NCR 8720904.

Management and Operation of the NSFNET Backbone Network

August Monthly Report

Merit Computer Network

NSFNET Traffic Statistics

With the advent of August, the production NSFNET backbone marks its first full month of sustained operation for all nodes. This is reflected in the traffic statistics, which are reported for all 31 days of the month. Again, the packet numbers are collected hourly and reflect packets into and out of each NSS as measured at the LAN-0 interface. These counts are collected via SGMP for each node and stored in a SPIRES database on the Information Services host machine. Figures A through F summarize the findings for August, with actual numbers shown in the spreadsheet in Table 1.

For August, the total packet counts are 202,641,056 packets in and 194,041,532 out. The daily number of packets peaked at over 20 million, more than 5 million higher than the maximum reported for July. Again in August, there is a consistent drop in traffic on weekends, with the lowest days recorded on Sundays. Highest counts are appearing on Wednesday (a trend which changes in September, judging by our preliminary analyses of September data). The lower rates on days around weekends may in part relate to the prevalence of vacations and long weekends many researchers have during August. In general, packet counts are steadily increasing, with the higher counts all occurring during the later part of August.

The packet numbers vary dramatically by NSS, with two sites reporting usage much greater than the other thirteen. Both JVNCnet (NSS 8) and CNSF/NYSERNet (NSS 10) have monthly totals greater than 60 million each. By contrast, Westnet (NSS 15) and Midnet (NSS 16) report less than 10 million packets each. This in part is related to the number of attached networks at each NSS, although other factors also are affecting these counts including the maturity of the regional networks and the co-location of a supercomputer center. The relationship between these factors will be further analyzed in future reports.

One-way delay data

As with packet data, August marks the first month for which a complete report of one-way delay data is available for the new backbone. Pings were recorded once each day between all pairs of NSSs. (However, technical problems resulted in the loss of data for NSS 8 during August.) The ping data has been divided by two to get the one-way delays in accordance with our agreement with NSF. Table 2 shows the minimum, maximum, and average times recorded for each pair. Perhaps the most notable trend is the fact that the average is at the lower part of each range, indicating that the higher delays are relatively rare. Variation in average times is partially related to real distance as well as the number of IDNX hops required between any given pair. These factors are continuing to be analyzed.

Significant Network Events

The data from the month of August shows overall stability of the network. Generally, the outages are short and infrequent. These data are presented in Table 3. Outages were divided into two categories "Class One" is full node outages and "Class Two" is partial node outage resulting in reduced performance relative to the backbone. As the tracking mechanisms develop and problem determination improves, it is our hope that certain outages will be avoidable.

All "Class One," full node outages, were limited to only a few hours. "Class Two," outages resulting in minor performance degradation, was limited to less than a day.

The longest "Class Two" outage was that of JVNC on August 8. Even in this case, full recovery was made in less than 24 hours. In other "Class One" outages, JVNC's link to SURANet was lost, and work is being done to determine exactly why this occurred and what can be done to prevent it from happening again.

This report includes the following information:

Table 1: **Raw packet counts in and out of the NSFNET backbone**
Shows the total number of packets per day for each NSS for the month of August.

Figure A: **Daily aggregate packet counts**
Shows the total packet count for all nodes for each day during the month of August.

Figure B: **Weekly aggregate packet counts**
Shows the total packet counts for all nodes by week during the month of August

Figure C: **Average packet counts by day of the week**

Shows the average number of packets in and out by day of the week for all NSSs.

Figure D: **Aggregate packet counts by node**
Shows the total packet count by node for the month of August

Figure E **Aggregate weekly packet counts in and out per NSS**
Shows the weekly number of packets *in* per NSS.
Shows the weekly number of packets *out* per NSS.

Figure F: **Daily range and average of packets in and out per NSS**
Shows minimum, maximum and average packets for the Month of August *in* of each node.
Shows minimum, maximum and average packets for the Month of August *out* of each node.

Table 2 **Average one-way delay times (in milliseconds)**
Shows minimum, maximum, and average between all NSS pairs for the month of August.

Table 3 **NSFNET Significant Network Events**
Shows outages, the resolution of the problem, and the classification for each outage.

The following figures are by NSS number. The key for these is:

5 PSCNET
6 SDSCNET
7 USAN
8 JVNCNET
9 SURANET
10 CNSF/NYSERNET
11 SESQUINET
12 NCSA
13 BARRNET
14 NORTHWESTNET
15 WESTNET
16 MIDNET
17 MERIT

Table 1
NSFNET Traffic–August 1988
Packets per day for each NSS

TRAFFIC PATTERNS—AUGUST

	nss5	nss6	nss7	nss8	nss9	nss10	nss11	nss12	nss13	nss14	nss15	nss16
Packets In												
8/1	252,006	461,233	445,926	1,115,228	1,067,611	1,100,730	407,431	643,968	779,087	200,153	161,755	116,446
8/2	377,221	558,071	365,924	1,137,717	1,202,341	1,138,780	257,687	807,463	844,104	281,690	135,530	104,117
8/3	315,926	783,008	466,948	2,261,856	1,063,481	1,295,505	269,134	711,149	642,378	277,560	142,459	117,139
8/4	383,848	509,499	414,330	1,552,535	706,463	1,414,877	294,781	890,074	766,099	308,944	106,468	82,811
8/5	429,381	377,635	310,058	1,279,183	527,114	1,504,711	256,890	652,023	450,304	282,963	189,025	123,601
8/6	142,913	222,446	162,418	849,269	303,854	624,067	171,206	529,821	515,195	191,575	72,283	77,589
8/7	114,017	132,440	182,558	791,479	290,982	664,936	188,639	611,219	388,568	249,778	72,078	11,485
8/8	254,437	387,417	344,784	1,141,206	693,973	1,374,886	435,885	719,426	771,907	316,464	123,848	90,629
8/9	439,643	587,400	335,664	1,345,367	719,363	1,271,011	372,149	890,305	765,194	330,859	146,439	69,934
8/10	378,989	657,832	450,916	1,752,468	713,086	1,057,245	319,884	849,993	888,120	368,337	170,838	87,858
8/11	333,410	29,949	8,385	101,488	653,999	325,920	183,939	647,702	88,969	225,044	200,757	2,864
8/12	86,146	190,494	76,251	278,902	152,042	294,172	56,673	198,254	328,776	112,647	24,652	9,344
8/13	271,103	312,219	215,182	704,243	410,276	780,576	148,385	447,187	579,297	186,329	66,811	42,187
8/14	230,661	212,353	343,979	818,572	302,548	665,550	208,657	445,776	464,636	168,599	99,976	61,284
8/15	263,351	550,760	353,069	1,093,713	497,531	920,099	272,702	578,685	829,057	230,222	150,657	68,911
8/16	318,142	823,225	450,581	1,175,316	832,203	1,387,200	337,287	733,642	693,632	414,385	206,515	102,324
8/17	312,022	591,051	426,027	765,038	745,773	1,248,115	319,356	713,423	381,638	415,438	233,641	112,440
8/18	299,425	589,821	488,084	1,161,939	554,068	1,181,335	309,208	723,168	305,215	263,568	134,924	82,792
8/19	349,924	662,229	380,306	1,359,742	n/a	1,097,525	235,928	655,297	330,186	490,349	260,898	73,139
8/20	266,120	229,489	231,947	1,063,313	n/a	753,110	133,081	545,455	288,743	297,191	90,181	53,713
8/21	229,928	336,063	292,817	909,003	287,463	589,733	176,783	408,132	203,351	259,795	66,744	21,796
8/22	156,989	639,606	399,634	1,212,499	282,970	1,457,340	338,346	706,375	423,598	466,575	206,973	84,341
8/23	164,648	425,370	615,635	1,792,461	667,284	1,541,932	239,610	720,979	235,993	239,993	98,050	58,082
8/24	425,210	551,698	798,924	1,980,403	717,151	1,975,259	273,095	1,053,892	492,855	306,494	181,204	103,889
8/25	262,662	465,465	494,622	1,079,866	616,965	1,137,529	215,150	808,373	870,079	269,638	190,172	90,525
8/26	342,510	569,742	734,744	2,214,211	1,032,153	2,041,251	241,931	954,131	1,283,757	403,337	189,773	85,148
8/27	217,319	271,796	343,392	1,601,178	733,296	1,351,982	273,748	761,462	823,163	154,761	139,688	36,790
8/28	n/a	n/a	153,332	599,813	505,747	305,076	61,725	475,339	389,641	227,121	48,568	51,179
8/29	277,714	331,926	360,768	1,174,004	688,002	1,036,581	231,271	738,694	567,187	195,095	98,400	64,163
8/30	327,087	598,087	612,920	1,564,672	1,237,850	1,662,409	413,018	1,185,526	1,064,846	336,065	334,203	131,031
8/31	596,390	662,572	679,445	1,841,653	1,188,315	1,678,332	375,922	1,032,600	1,214,623	442,188	258,182	81,441
Packets out												
8/1	221,102	427,640	400,017	896,310	879,874	971,815	393,733	892,112	1,036,189	267,145	71,619	51,673
8/2	317,144	462,059	399,505	897,879	1,091,448	919,670	271,384	813,698	1,171,476	297,867	68,221	66,878
8/3	324,303	464,234	1,311,271	1,293,423	922,626	880,467	298,834	754,585	1,074,846	310,541	54,955	74,269
8/4	467,414	413,961	415,062	1,164,507	866,415	1,094,185	290,935	996,286	755,637	360,982	64,016	75,692
8/5	331,624	290,941	336,352	1,078,955	587,996	839,392	276,988	741,227	608,336	384,062	208,412	63,487
8/6	248,933	214,882	189,039	290,563	423,633	508,871	165,193	604,215	524,368	209,206	57,358	29,695
8/7	124,120	116,293	202,829	511,174	366,978	630,751	195,166	659,408	397,076	280,209	60,134	10,250
8/8	247,003	351,249	372,371	732,546	710,694	913,302	487,584	864,448	607,492	367,027	76,436	48,229
8/9	421,439	455,021	422,854	1,264,977	757,190	894,316	379,915	945,905	660,630	321,701	83,454	53,924
8/10	359,124	610,869	493,550	1,233,449	786,808	1,146,270	346,106	962,175	762,315	425,464	111,989	79,393
8/11	294,367	426,015	26,788	51,539	747,251	1,000,571	10,457	33,495	67,794	20,857	104,310	74,419
8/12	48,525	178,596	85,961	266,152	176,881	240,613	65,392	246,711	272,567	122,133	16,966	9,473
8/13	174,799	262,458	242,240	660,891	444,468	641,724	143,762	585,112	491,149	245,081	60,293	41,541
8/14	305,196	197,313	374,862	556,115	316,961	665,666	230,712	624,645	388,767	211,876	77,620	53,319
8/15	262,592	451,744	346,761	620,376	697,400	869,986	316,425	667,034	761,086	279,410	130,069	61,454
8/16	292,083	534,868	445,177	878,380	1,000,150	974,351	373,154	937,414	1,021,944	487,873	166,948	68,667
8/17	359,152	310,630	409,616	635,394	687,407	733,855	336,924	720,989	776,091	427,981	194,470	67,917
8/18	275,265	340,697	459,449	767,900	656,462	1,123,243	384,466	852,344	842,838	292,583	116,445	61,140
8/19	366,829	401,750	390,120	864,046	692,127	1,235,259	308,460	803,226	748,844	335,861	162,729	64,054
8/20	235,576	162,285	249,544	792,411	376,950	742,237	176,701	603,885	569,796	164,703	84,327	62,398
8/21	287,162	216,369	350,551	650,644	347,407	400,484	232,814	520,435	447,862	284,837	39,125	27,538
8/22	292,562	399,476	432,361	764,918	334,008	1,121,880	330,890	719,997	629,080	611,853	174,897	90,157
8/23	162,617	300,502	608,689	1,266,423	702,948	1,447,096	267,925	580,016	519,157	283,741	117,382	43,345
8/24	357,004	356,419	819,687	1,499,027	886,979	1,597,386	316,932	885,868	1,098,676	346,163	240,633	84,281
8/25	278,252	418,485	487,758	848,203	638,202	1,095,924	266,874	755,152	775,085	314,355	224,328	81,709
8/26	414,377	547,147	711,399	1,697,862	788,590	1,998,438	280,153	1,051,484	1,144,816	392,146	222,559	84,182
8/27	142,438	255,202	316,749	1,065,841	481,479	1,772,877	274,600	645,694	845,228	184,127	160,694	46,751
8/28	n/a	n/a	152,982	473,559	320,337	850,776	85,165	348,144	320,507	148,914	25,697	52,308
8/29	309,270	290,698	338,223	820,867	555,052	1,058,251	281,426	513,097	540,802	182,185	101,715	51,900
8/30	319,806	453,561	519,886	1,176,616	881,012	1,776,255	488,779	1,254,062	1,012,637	400,730	259,994	90,196
8/31	450,896	508,267	653,264	1,405,300	956,204	1,852,981	420,389	1,059,445	960,848	541,212	255,533	73,954

Figure A

NSFNET Traffic—August 1988

Daily packets in and out
for all NSSs

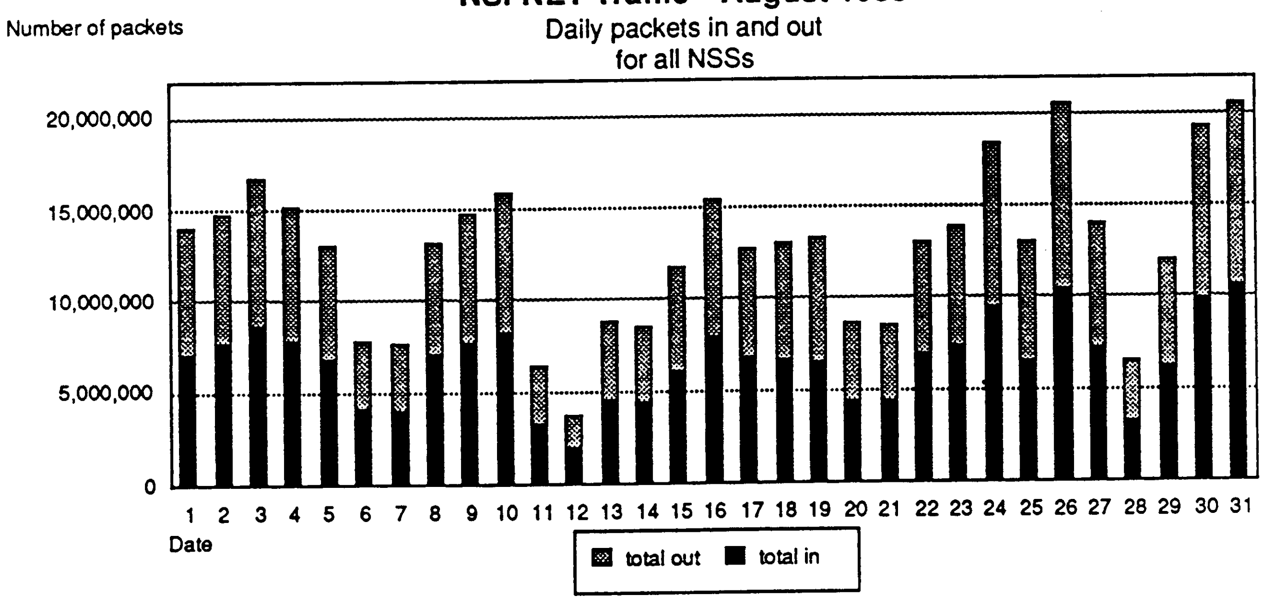

Figure B
NSFNET Traffic—August 1988

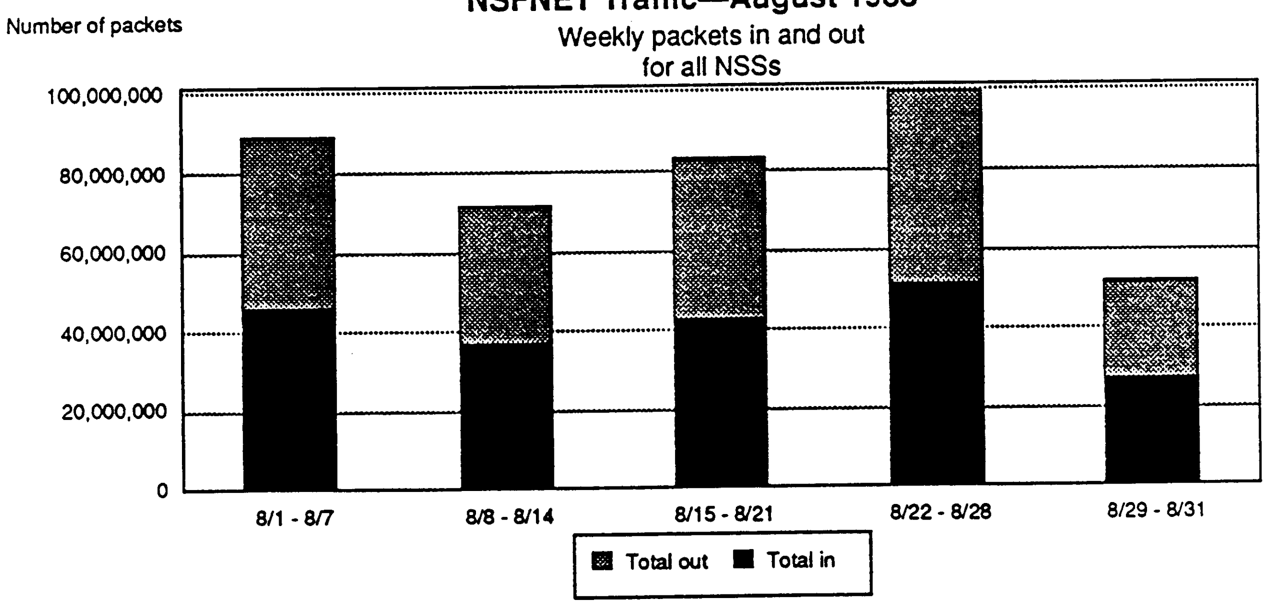

Figure C
NSFNET Traffic—August 1988
Average number of packets in and out
by day of the week for all NSSs

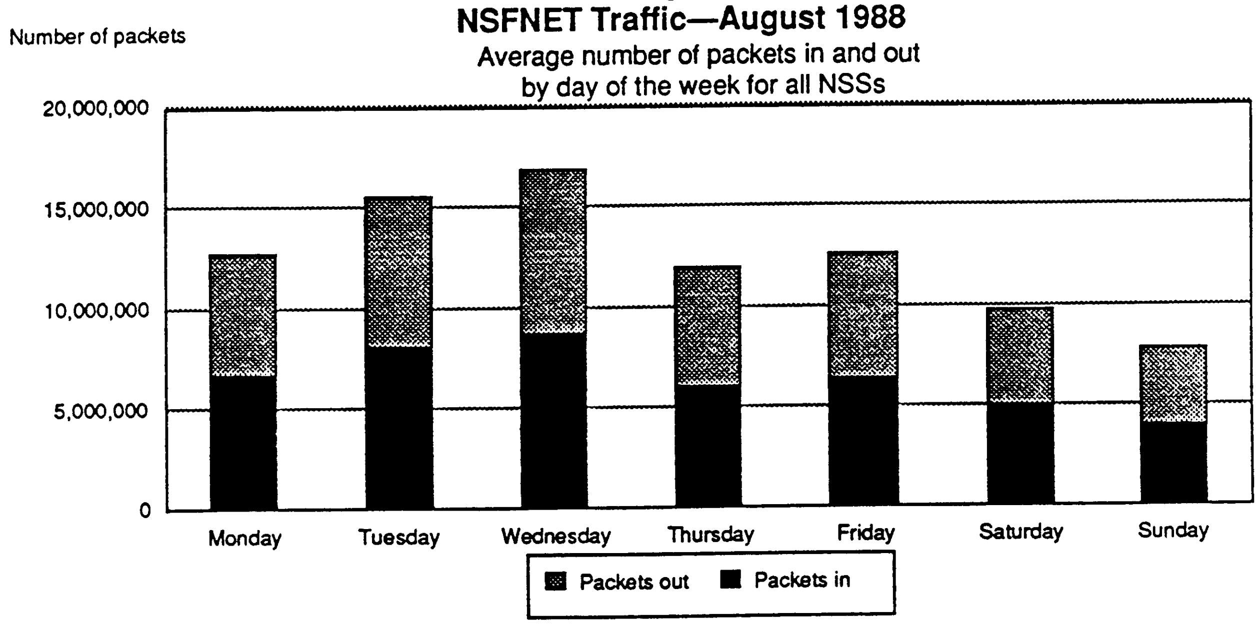

Figure D
NSFNET Traffic—August 1988

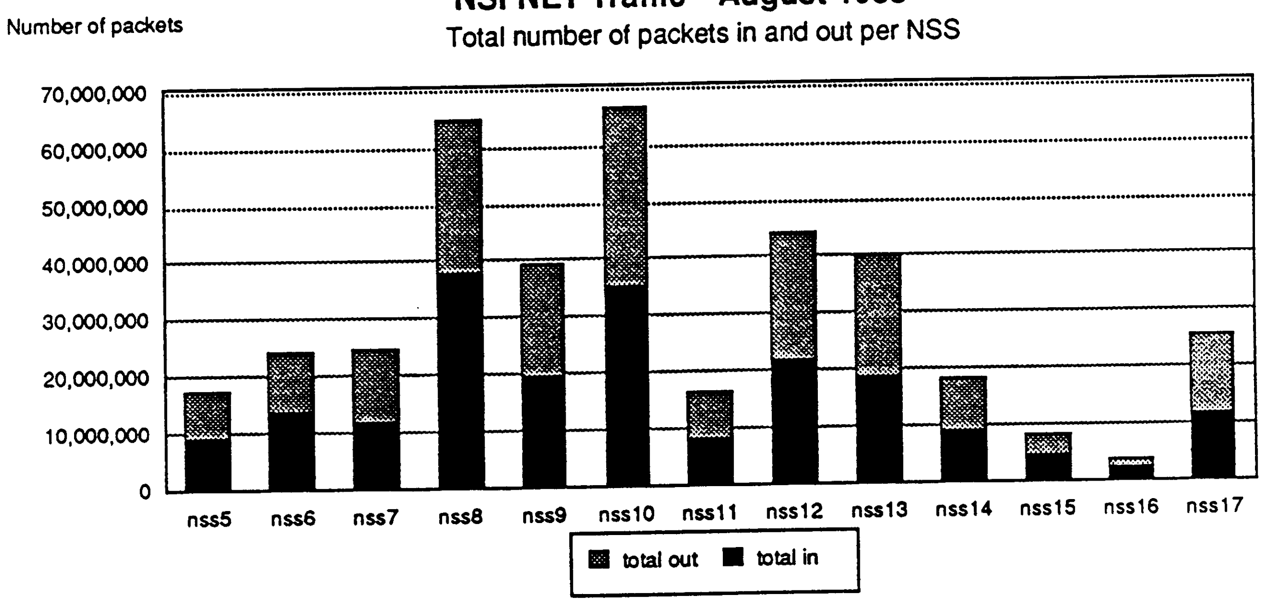

Figure E

NSFNET Traffic—August 1988

Weekly packet counts in per NSS

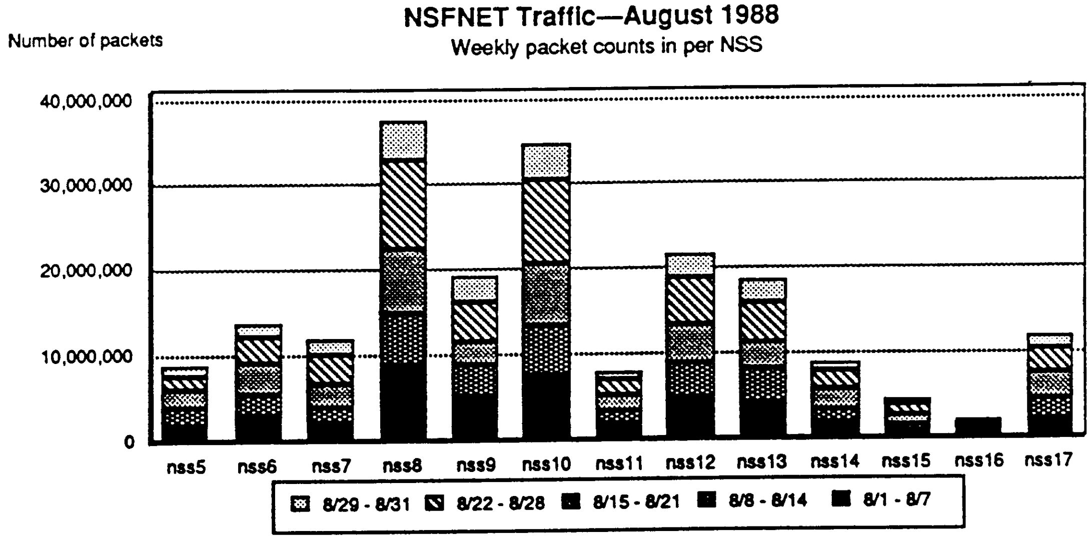

NSFNET Traffic—August 1988

Weekly packet counts out per NSS

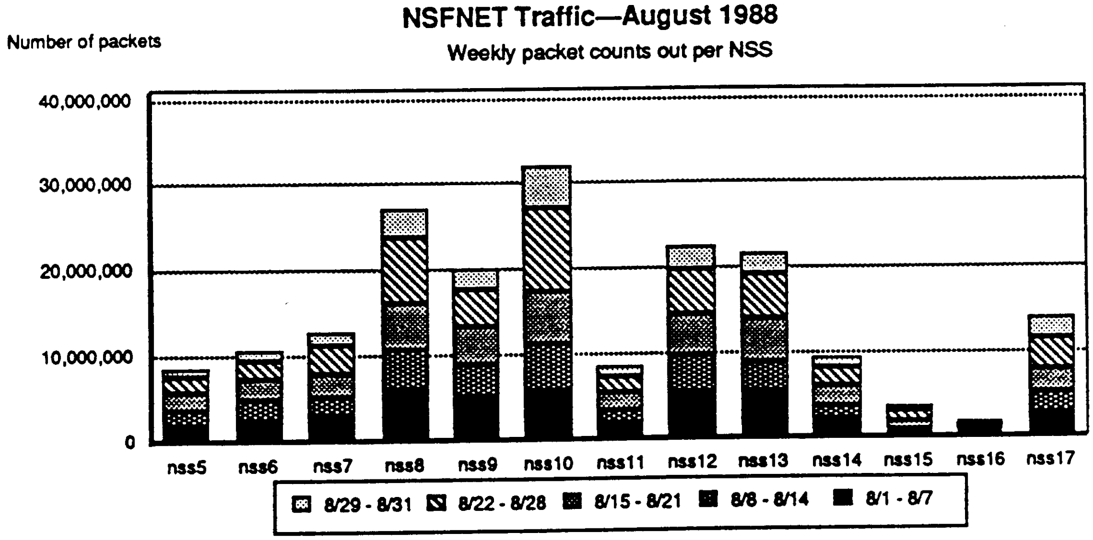

Figure F

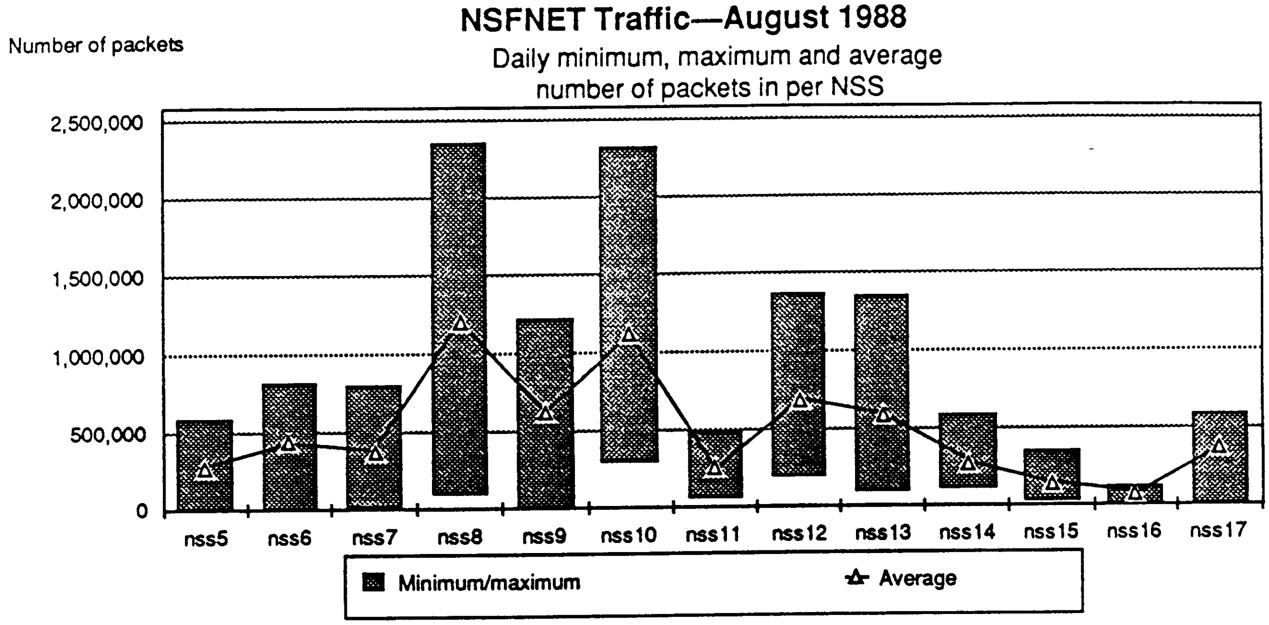

NSFNET Traffic—August 1988
Daily minimum, maximum and average
number of packets in per NSS
Number of packets
2,500,000
2,000,000
1,500,000
1,000,000
500,000
0
nss5
nss6
nss7
nss8
nss9
nss10
nss11
nss12
nss13
nss14
nss15
nss16
nss17
Minimum/maximum
Average

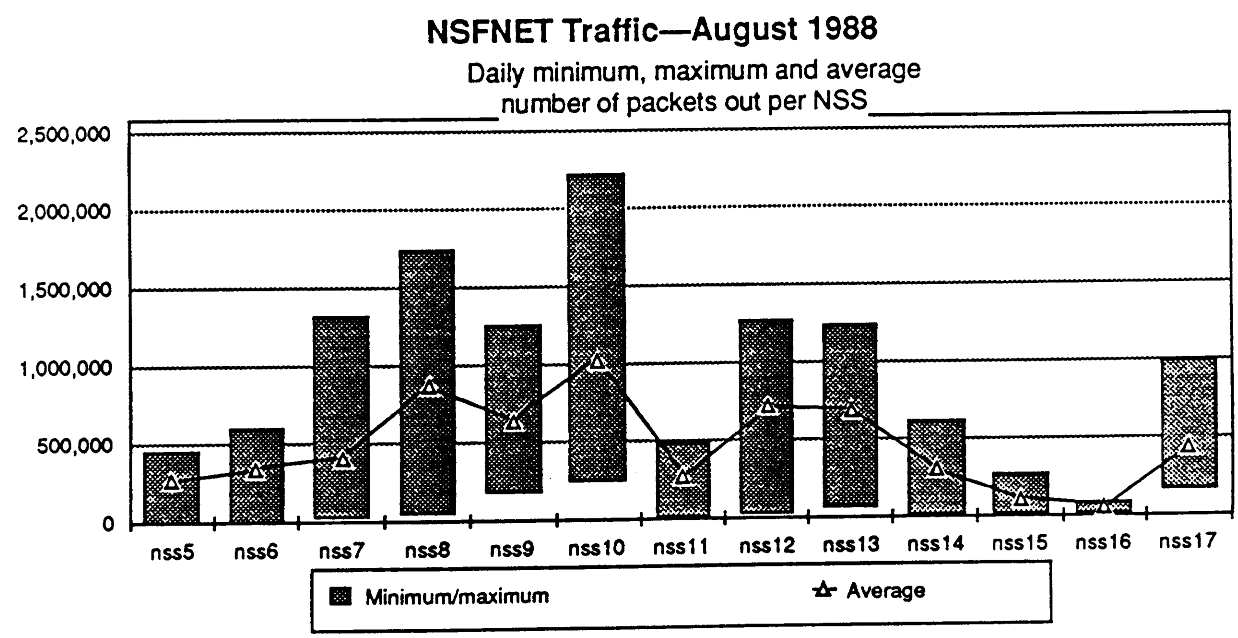

NSFNET Traffic—August 1988
Daily minimum, maximum and average
number of packets out per NSS
2,500,000
2,000,000
1,500,000
1,000,000
500,000
0
nss5
nss6
nss7
nss8
nss9
nss10
nss11
nss12
nss13
nss14
nss15
nss16
nss17
Minimum/maximum
Average

Table 2
NSFNET Traffic–August 1988
Average one-way delay times

		One-way delay time		
from	to	max	min	avg
NSS5	NSS5	4.4	1.4	1.7
NSS5	NSS6	86.9	74.8	78.1
NSS5	NSS7	69.1	50.5	56.7
NSS5	NSS8	63.6	42.9	50.0
NSS5	NSS9	52.3	38.6	44.0
NSS5	NSS10	63.8	50.2	53.8
NSS5	NSS11	57.1	40.2	47.8
NSS5	NSS12	48.0	30.7	36.8
NSS5	NSS13	86.7	72.4	77.8
NSS5	NSS14	83.7	78.9	79.8
NSS5	NSS15	70.4	64.2	65.2
NSS5	NSS16	88.1	74.6	78.4
NSS5	NSS17	34.2	20.4	24.7
NSS6	NSS5	77.7	71.9	72.9
NSS6	NSS6	3.4	1.5	1.6
NSS6	NSS7	30.8	27.6	27.9
NSS6	NSS8	87.6	76.6	78.2
NSS6	NSS9	101.1	93.4	94.7
NSS6	NSS10	100.3	87.7	89.3
NSS6	NSS11	105.7	97.1	98.0
NSS6	NSS12	69.1	67.0	67.2
NSS6	NSS13	23.3	21.5	21.7
NSS6	NSS14	28.8	28.1	28.2
NSS6	NSS15	46.6	44.4	44.7
NSS6	NSS16	44.1	42.9	43.0
NSS6	NSS17	57.9	55.2	55.6
NSS7	NSS5	65.9	47.4	49.8
NSS7	NSS6	29.3	27.6	27.8
NSS7	NSS7	3.4	1.4	1.6
NSS7	NSS8	59.8	52.4	53.6
NSS7	NSS9	77.9	68.9	70.3
NSS7	NSS10	67.1	63.6	64.1
NSS7	NSS11	76.4	71.8	72.3
NSS7	NSS12	91.1	71.1	74.4
NSS7	NSS13	48.9	46.0	46.4
NSS7	NSS14	53.9	52.4	52.4
NSS7	NSS15	21.9	20.4	20.6
NSS7	NSS16	29.2	25.5	25.8
NSS7	NSS17	38.9	29.9	31.7

		One-way delay time		
from	to	max	min	avg
NSS9	NSS5	39.9	38.3	38.6
NSS9	NSS6	105.3	93.4	95.2
NSS9	NSS7	78.1	68.9	70.4
NSS9	NSS8	30.6	21.5	22.8
NSS9	NSS9	1.8	1.4	1.5
NSS9	NSS10	35.5	29.6	30.4
NSS9	NSS11	76.8	70.4	71.1
NSS9	NSS12	70.4	62.0	63.3
NSS9	NSS13	117.8	106.9	108.9
NSS9	NSS14	122.2	113.5	114.6
NSS9	NSS15	93.9	88.4	89.2
NSS9	NSS16	81.4	77.9	78.4
NSS9	NSS17	44.6	41.6	42.1
NSS10	NSS5	58.8	50.8	51.9
NSS10	NSS6	92.9	87.9	88.7
NSS10	NSS7	67.8	63.7	64.4
NSS10	NSS8	61.2	55.3	56.9
NSS10	NSS9	31.8	29.6	29.8
NSS10	NSS10	2.1	1.4	1.6
NSS10	NSS11	84.9	76.1	77.4
NSS10	NSS12	40.5	35.9	36.3
NSS10	NSS13	84.3	81.1	81.4
NSS10	NSS14	94.4	87.2	88.3
NSS10	NSS15	63.7	62.2	62.3
NSS10	NSS16	93.8	85.9	87.2
NSS10	NSS17	48.7	43.5	44.3
NSS11	NSS5	37.0	33.9	34.1
NSS11	NSS6	102.1	98.4	98.9
NSS11	NSS7	79.4	72.6	73.6
NSS11	NSS8	70.6	68.3	68.6
NSS11	NSS9	70.0	68.3	68.5
NSS11	NSS10	94.3	84.4	85.5
NSS11	NSS11	2.1	1.5	1.5
NSS11	NSS12	61.8	57.5	58.0
NSS11	NSS13	89.9	89.0	89.0
NSS11	NSS14	114.0	108.8	109.8
NSS11	NSS15	93.1	89.8	90.1
NSS11	NSS16	98.6	94.8	95.5
NSS11	NSS17	47.1	45.3	45.4

Table 2
NSFNET Traffic–August 1988
Average one-way delay times

from	to	One-way delay time max	min	avg
NSS12	NSS5	38.5	30.6	31.8
NSS12	NSS6	73.1	66.9	67.7
NSS12	NSS7	82.9	72.3	75.2
NSS12	NSS8	74.9	66.5	68.0
NSS12	NSS9	69.3	62.1	63.1
NSS12	NSS10	43.3	35.9	36.7
NSS12	NSS11	67.7	60.6	61.9
NSS12	NSS12	1.9	1.4	1.5
NSS12	NSS13	50.2	48.6	48.8
NSS12	NSS14	56.8	55.2	55.4
NSS12	NSS15	95.4	89.2	90.3
NSS12	NSS16	102.0	94.9	96.3
NSS12	NSS17	51.1	45.6	46.7
NSS13	NSS5	77.9	72.5	73.0
NSS13	NSS6	23.7	21.4	21.7
NSS13	NSS7	51.3	46.0	46.5
NSS13	NSS8	105.5	95.0	96.9
NSS13	NSS9	129.7	107.0	110.8
NSS13	NSS10	92.9	81.1	82.4
NSS13	NSS11	90.0	89.2	89.2
NSS13	NSS12	50.6	48.6	48.8
NSS13	NSS13	1.9	1.5	1.5
NSS13	NSS14	48.8	46.4	46.6
NSS13	NSS15	65.1	62.8	63.2
NSS13	NSS16	67.8	61.5	62.1
NSS13	NSS17	78.0	73.0	73.8
NSS14	NSS5	83.8	78.9	79.8
NSS14	NSS6	29.7	28.2	28.3
NSS14	NSS7	55.2	52.4	52.9
NSS14	NSS8	111.5	101.4	102.9
NSS14	NSS9	118.4	113.6	114.4
NSS14	NSS10	94.2	87.3	87.9
NSS14	NSS11	121.3	110.3	111.9
NSS14	NSS12	56.5	55.3	55.3
NSS14	NSS13	53.7	46.2	47.6
NSS14	NSS14	2.2	1.5	1.5
NSS14	NSS15	57.3	56.0	56.2
NSS14	NSS16	72.8	67.8	68.3
NSS14	NSS17	83.5	79.8	80.5
NSS15	NSS5	74.0	64.3	65.8
NSS15	NSS6	46.9	44.6	44.9
NSS15	NSS7	25.3	20.5	21.0
NSS15	NSS8	78.2	69.2	70.5
NSS15	NSS9	92.9	88.3	88.8
NSS15	NSS10	73.6	62.2	63.4
NSS15	NSS11	96.6	88.8	89.8
NSS15	NSS12	95.7	87.9	89.8
NSS15	NSS13	69.8	62.8	63.5
NSS15	NSS14	57.0	56.1	56.1
NSS15	NSS15	2.0	1.5	1.6
NSS15	NSS16	46.8	42.3	42.6
NSS15	NSS17	49.6	47.7	47.9
NSS16	NSS5	84.1	72.5	74.8
NSS16	NSS6	48.9	43.0	43.5
NSS16	NSS7	27.9	25.5	25.6
NSS16	NSS8	70.9	62.7	63.6
NSS16	NSS9	80.8	77.9	78.3
NSS16	NSS10	92.3	86.0	86.9
NSS16	NSS11	97.9	94.0	94.7
NSS16	NSS12	103.5	93.4	95.3
NSS16	NSS13	63.2	61.5	61.7
NSS16	NSS14	70.3	67.7	68.1
NSS16	NSS15	44.8	42.3	42.4
NSS16	NSS16	2.4	1.4	1.5
NSS16	NSS17	55.0	53.0	53.1
NSS17	NSS5	23.3	20.4	20.8
NSS17	NSS6	57.1	55.1	55.4
NSS17	NSS7	32.2	30.7	30.9
NSS17	NSS8	32.9	24.9	25.9
NSS17	NSS9	44.4	41.6	42.1
NSS17	NSS10	54.1	43.4	44.7
NSS17	NSS11	49.2	44.4	44.8
NSS17	NSS12	54.4	43.7	45.8
NSS17	NSS13	79.4	73.1	73.9
NSS17	NSS14	83.6	79.8	80.3
NSS17	NSS15	49.7	47.6	47.8
NSS17	NSS16	55.8	52.9	53.4
NSS17	NSS17	1.7	1.5	1.5

Table 3

NSFNET Significant Network Events
August 1988

Outage classifications:
- Class 1: Full node outage
- Class 2: Partial node outage with impact to addtional nodes

DATE	PROBLEM	RESOLUTION	CLASS
NSS #5 Pittsburgh Supercomputer Center / PSCNET			
August 01	PSP-5-13 & PSP-5-17 down temporarily	Almaden cards changed	2
August 11	All machines down for 40min.	RCP rebooted	1
August 09	IDNX link down for 2hrs.20min.	Trunk card swapped	1
August 19	All machines down for 2hrs.10min.	MCI switching circuits	1
	All Machines down for 2hrs.5min.	MCI switching circuits	1
	All machines down for 45min.	MCI switching circuits	1
	Ann Arbor to Pittsburgh link down 5hrs.20min.	Bad repeater	2
NSS #6 San Diego Supercomputer Center / SDSCNET			
August 04	All machines down for 40min.	RCP rebooted	1
NSS #7 National Center for Atmospheric Research / NCAR			
August 01	PSP-7-12 down for 19hrs.	Hard drive replaced	2
August 20	All machines down for 4hrs.35min.	Fiber break	1
NSS #8 John Von Neumann National Supercomputer Center / JVNCNET			
August 06	All machines down 4hrs.30min.	Power outage-construction	1
August 08	All machines down temporarily	Power outage-power co.	1
August 08	PSP-8-11 down 23hrs.30min	PSP rebooted	2
August 17	All machines down 3hrs.45min.	Electrical storm	1
August 22	All machines down 1hr.	Generator problems	1
August 25	All machines down 3hrs.40min	Electrical storm	1
NSS #9 University of Maryland College Park, MD / SURANET			

Impacted by JVNC Power Outages

DATE	PROBLEM	RESOLUTION	CLASS
NSS #10 Cornell University Ithaca, NY / CNSF/NYSERNET			
August 30	PSP-10-12 down 5hrs.	PSP rebooted	2
NSS #11 Rice University Houston, TX / SESQUINET			
August 04	All machines down 1hr.	Sliding cable locks repaired	1
August 31	PSP-11-13 down 9hrs.	PSP rebooted	2
NSS #12 National Center for Supercomputer Applications / NCSA			
August 10	All machines down 1hr.15min.	No disk space	
NSS #13 Stanford University Palo Alto, CA / BARRNET			
No Major Problems in August			
NSS #14 University of Washington Seattle, WA / NORTHWESTNET			
August 18	All machines down 5hrs.20min.	MCI replaced a link part	1
NSS #15 University of Utah Salt Lake City, UT / WESTNET			
August 20	All machines down 4hrs.35min.	Fiber break	1
NSS #16 University of Nebraska-Lincoln Lincoln, NE / MIDNET			
August 04	PSP-16-10 was down 3hrs.	Disk controller replaced	1
August 06	All machines down 9hrs.	Scheduled power outage	1
August 13	All machines down 3hrs.	Fiber break	1
August 20	All machines down 4hrs.35min.	Fiber break	1
NSS #17 University of Michigan Ann Arbor, MI / MERIT			
August 15	PSP-17-14 down temporarily	PSP rebooted	2

California Internet Federation Participants

BARRNet

CERFNET

California State University

Los Nettos

NASA

San Diego Supercomputer Center

State of California - Department of Water Resources

University of California

California Internet Federation

The California Internet Federation of computer network organizations has the purpose of providing coordination and support of educational and research networking in California. California is recognized as a leader in high technology. To maintain this leadership, however, California's educational and research institutions require the communication tools to share information, resources and ideas. Isolated facilities can no longer compete in today's fast paced age of information. The California Internet Federation has been formed to insure that high quality communication tools are available for education and research to keep California in its position of leadership in these areas.

California Internet Federation Objectives

1) Coordinate interconnection of educational and research networks in California. Areas of coordination include:

> a) Design of cost-effective and reliable interconnection among these computer networks.
>
> b) Assist with agreements among network administrations in support of interconnections.
>
> c) Implementation of connections and routing strategies.
>
> d) Management schemes for the connection of interconnected networks.

2) Provide coordination for the connection of California networks with national and international networks.

3) Support of educational and research networking by promoting:

> a) Use of standards and compatibility of networks.
>
> b) The understanding of internetwork technologies.
>
> c) dissemination of information about resources available via the internet.
>
> d) Development of new resources available via the internet.
>
> e) Collaboration between private and public sectors.

4) Increase visibility of internetworking and demonstrate its importance to California.

Calfornia Internet Federation Meeting
August 23,1988

Name	Organization	Network connections.	Email Address
Clark, Ray	CSU Chancellor	CSUNET	rclark@calstate.bitnet
Cooling, Mike	CSU Sacto	CSUNET	cssexb!cooling@ucdavis.edu
Darling, Gary	State DWR	DWR nets	caldwr!gary@ucdavis.edu
DeJarnett, Steve	Cal Poly SLO	CSUNET	steve@polyslo.calpoly.edu
Estrada, Susan	SDSC	CERFNET	estradas@sds.sdsc.edu
Hobby, Russ	UC Davis	BARRNet UCNET	rdhobby@ucdavis.edu
Jones, Bill	NSI/NASA Ames	NSI/NASA	jones@nsipo.nasa.gov
Liu, Mei-ling	Cal Poly SLO	CSUNET	mliu@polyslo.calpoly.edu
Love, E. Paul	SDSC	CERFNET	loveep@sds.sdsc.edu
Lynch, Clifford	UCOP	UCNET	lynch@postgres.berkeley.edu
Madden, Jim	UC San Diego	CERFNET UCNET	madden@ucsd.edu
Neuman, Gerald K	SDSC	CERFNET	gkn@sds.sdsc.edu
Ollikainen, Ari	RIACS/NASA	BARRNet NASA Nets	ari@riacs.edu
Postel, Jon	USC/ISI	Los Nettos	postel@isi.edu
Prue, Walt	USC/ISI	Los Nettos	prue@isi.edu
Scott, Greg	UC Santa Cruz	BARRNet UCNET	greg@ucscm.ucsc.edu
Smith, Dick	CSU Sacto	CSUNET	lll-crg!csusac!dsmith
Taylor, Chris	CSU Chancellor	CSUNET	1gtlfct@calstate.bitnet
Tomcheck, Dave	UC Irvine	CERFNET	tomcheck@uci.bitnet
Walker, David	UC Irvine	CERFNET UCNET	dhwalker@uci.bitnet
Wasley, David	UC Berkeley	BARRNet UCNET	dlw@violet.berkeley.edu
Wills, Dave	UCOP	UCNET	
Yundt, Bill	Stanford	BARRNet	gd.why@forsythe.stanford.edu

Calfornia Internet Federation Meeting
September 28, 1988

Name	Organization	Network connections.	Email Address
Darling, Gary	State DWR	DWR nets	caldwr!gary@ucdavis.edu
Ferrin, Tom	UCSF	BARRNet	tef@cgl.ucsf.edu
Fink, Robert	LBL	BARRNet	rlfink@lbl.gov
Griffiths, Darren	LBL	BARRNet	dagg@lbl.gov
Harel, Elie	UCLA	UCNET	aceh0@uclaais.bitnet
Hobby, Russ	UC Davis	BARRNet UCNET	rdhobby@ucdavis.edu
Jones, Bill	NSI/NASA Ames	NSI/NASA	jones@nsipo.nasa.gov
Lynch, Clifford	UCOP	UCNET	lynch@postgres.berkeley.edu
Madden, Jim	UC San Diego	CERFNET UCNET	madden@ucsd.edu
Neuman, Gerald K	SDSC	CERFNET	gkn@sds.sdsc.edu
Ollikainen, Ari	RIACS/NASA	BARRNet NASA Nets	ari@riacs.edu
Prue, Walt	USC/ISI	Los Nettos	prue@isi.edu
Reese, David	CSU Chancellor	CSUNET	1gtlfct@calstate.bitnet
Scott, Greg	UC Santa Cruz	BARRNet UCNET	greg@ucscm.ucsc.edu
Stefferud, Einar	NMA-Northrop		stef@nrtc.northrop.com
Taylor, Chris	CSU Chancellor	CSUNET	1gtlfct@calstate.bitnet
Walker, David	UC Irvine	CERFNET UCNET	dhwalker@uci.bitnet
Wasley, David	UC Berkeley	BARRNet UCNET	dlw@violet.berkeley.edu
Wills, Dave	UCOP	UCNET	
Yundt, Bill	Stanford	BARRNet	gd.why@forsythe.stanford.edu