

AD-A236 542



December 1990

2

University
of Southern
California



Network Communications
Research

Final Report

DTIC
ELECTE
MAY 28 1991
S B D

91-00554



91 5 24 076

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

INFORMATION
SCIENCES
INSTITUTE



213/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

8c. ADDRESS (City, State, and ZIP Code) Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209			10. SOURCE OF FUNDING NUMBERS <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 25%;">PROGRAM ELEMENT NO. -----</td> <td style="width: 25%;">PROJECT NO. -----</td> <td style="width: 25%;">TASK NO. -----</td> <td style="width: 25%;">WORK UNIT ACCESSION NO. -----</td> </tr> </table>				PROGRAM ELEMENT NO. -----	PROJECT NO. -----	TASK NO. -----	WORK UNIT ACCESSION NO. -----		
PROGRAM ELEMENT NO. -----	PROJECT NO. -----	TASK NO. -----	WORK UNIT ACCESSION NO. -----									
11. TITLE (Include Security Classification) Network Communications Research Final Report (Unclassified)												
12. PERSONAL AUTHOR(S) Postel, J.; Anderson, C; Braden, R.; Casner, S.; Finn, G. G.; Mockapetris, P.												
13a. TYPE OF REPORT Final Report		13b. TIME COVERED FROM 11/87 TO 10/90		14. DATE OF REPORT (Year, Month, Day) 1990, December		15. PAGE COUNT 50						
16. SUPPLEMENTARY NOTATION												
17. COSATI CODES <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">FIELD</td> <td style="width: 33%;">GROUP</td> <td style="width: 33%;">SUB-GROUP</td> </tr> <tr> <td>09</td> <td>02</td> <td></td> </tr> </table>			FIELD	GROUP	SUB-GROUP	09	02		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) BFTPTool, congestion control, DARTnet, dig, doc, domain name, echo cancellers, groupware, high-capacity networks, high-speed networking, long-haul networks, Los Nettos, MBFTPTool, MMCC,			
FIELD	GROUP	SUB-GROUP										
09	02											
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This report summarizes work performed in the area of Network Communications Research. This effort was divided into four projects: 1) the Internet Concepts project, 2) the Multimedia Conferencing project, 3) the Protocol Accelerator project, and 4) the Supercomputer and Workstation Communication project. Under these projects, ISI research staff engaged in research into packet communication technologies, including, but not limited to, combinations of broadband satellite, internet protocols, and routing algorithms to assess their applicability to generic military command and control problems. Command and control system architectures for utilizing geographically distributed processing capabilities such as mixed voice, data and graphics conferencing, electronic message handling, and other applications were investigated, tested, and evaluated.												
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT												

18. MMConf, multimedia conferencing, network monitoring, *NNStat*, packet video, packet voice, personal conferencing, policy routing, PVP, routing, source quench, *statspy*, remote access protocols, remote window protocols, Request for Comments (RFC), US Domain



Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

UNIVERSITY OF SOUTHERN CALIFORNIA
INFORMATION SCIENCES INSTITUTE

**1990
FINAL
TECHNICAL
REPORT**

November 1987 - October 1990

**NETWORK COMMUNICATIONS
RESEARCH**

Sponsored by
Defense Advanced Research Projects Agency (DoD)
Information Science and Technology Office
Information Processing/Network Communications
Contract #MDA903-87-C-0719
ARPA Order No. 6180

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Table of Contents

I. Internet Concepts Project	2
I.A. TASK OBJECTIVES	2
I.B. TECHNICAL PROBLEM	3
I.C. GENERAL METHODOLOGY	3
I.D. TECHNICAL RESULTS	4
I.D.1. Protocol Studies	4
I.D.1.1. Congestion Control	4
I.D.1.2. NNStat and Statspy	5
I.D.1.3. Attack Resistant Protocols	6
I.D.1.4. Policy Routing	6
I.D.2. Intermail	7
I.D.3. The Domain Name System	7
I.D.3.1. DOC and DIG (Domain Tools)	8
I.D.3.2. The US Domain	8
I.D.4. Network Operations	9
I.D.4.1. Los Nettos	9
I.D.4.2. DARTnet	11
I.D.5. Internet Management	11
I.D.5.1. Internet Infrastructure and IAB Support	11
I.D.5.2. Host Requirements RFC	12
I.D.5.3. Requests for Comments	12
I.E. IMPORTANT FINDINGS AND CONCLUSIONS ..	12
I.F. SIGNIFICANT HARDWARE DEVELOPMENT ..	13
I.G. SPECIAL COMMENTS	13
I.H. IMPLICATIONS FOR FURTHER RESEARCH ...	13
I.H.1. Protocol Studies	13
I.H.1.1. Source Quench	13
I.H.1.2. NNStat and statspy	14
I.H.2. Domain Name System	14
I.H.3. Network Operations	14
I.H.3.1. Los Nettos	14
I.I. INC REPORTS AND PAPERS	14
 II. Multimedia Conferencing project	 17
II.A. TASK OBJECTIVES	17
II.A.1. Architecture and protocols	17
II.A.2. Prototypes	17
II.A.3. Media support	17
II.B. TECHNICAL PROBLEM	17
II.C. GENERAL METHODOLOGY	18
II.D. TECHNICAL RESULTS	18
II.D.1. Architecture and Protocols	18
II.D.1.1. Connection Management	18
II.D.1.2. ST-II	19
II.D.2. Prototypes	20
II.D.2.1. Packet Teleconference system	20
II.D.2.2. Installation of ISI video systems at DARPA and SRI	21
II.D.2.3. Extensions for multipoint teleconferencing ...	22

II.D.2.4. Improved Resolution	23
II.D.2.5. Audio Improvements	24
II.D.2.6. Port of packet voice and video processing to less expensive platform	24
II.D.3. Media Support	25
II.D.3.1. Refinements to video system implementation .	25
II.D.3.2. Echo Canceller	25
II.D.3.3. Image scanning and printing	26
II.E. IMPORTANT FINDINGS AND CONCLUSIONS .	26
II.F. SIGNIFICANT HARDWARE DEVELOPMENT ..	27
II.G. IMPLICATIONS FOR FURTHER RESEARCH ..	27
II.H. MMC REPORTS AND PAPERS	27
 III. Protocol Accelerator project	 29
III.A. TASK OBJECTIVES	29
III.B. TECHNICAL PROBLEM	29
III.C. GENERAL METHODOLOGY	29
III.D. IMPORTANT FINDINGS AND CONCLUSIONS	29
III.E. SIGNIFICANT HARDWARE DEVELOPMENT .	30
III.F. SPECIAL COMMENTS	30
III.G. IMPLICATIONS FOR FURTHER RESEARCH .	30
 IV. Supercomputer and Workstation Communication	 31
IV.A. TASK OBJECTIVES	31
IV.B. TECHNICAL PROBLEM	31
IV.C. GENERAL METHODOLOGY	31
IV.D. TECHNICAL RESULTS	31
IV.E. IMPORTANT FINDINGS AND CONCLUSIONS	32
IV.F. SIGNIFICANT HARDWARE DEVELOPMENT .	32
IV.G. IMPLICATIONS FOR FURTHER RESEARCH .	32
IV.H. SWC REPORTS AND PAPERS	33
 V. Appendix	 35

NETWORK COMMUNICATIONS RESEARCH

Final Technical Report

Overview

This report summarizes the work of the Internet Concepts Research project, the Multimedia Conferencing project, the Protocol Accelerator project, and the Supercomputer and Workstation Communication project, at USC/Information Sciences Institute under DARPA contract MDA903-87-C-0719 titled "Network Communications Research". The substantive results of this contract are reported in detail in the publications cited in this report.

Under this contract, the projects listed above engaged in research into new packet communication technologies, including, but not limited to, combinations of broadband satellite, internet protocols, and routing algorithms to assess their applicability to generic military command and control problems.

The projects investigated, implemented, tested, evaluated, and documented command and control system architectures for utilizing geographically distributed processing capabilities such as mixed voice, data and graphics conferencing, electronic message handling, and other applications.

I. Internet Concepts Project

Project Leader: Jon Postel

Research Staff: Research Assistants: Technical Staff: Support Staff:

Robert Braden

Steve Hotz

Joyce Reynolds

Celeste Anderson

Danny Cohen

Ann Westine

Kathleen McLaughlin

Gregory G. Finn

Paul Mockapetris

Walter W. Prue III

The Internet Concepts project performed studies and produced reports on various issues related to internetwork communications. Routing in very large-scale networks and protocol issues in very high-speed networks were addressed. Technical support was provided to the DARPA research community by maintaining the protocol specifications and by editing and publishing the Request for Comments (RFCs).

The work of the Internet Concepts project is discussed under five topic areas: (1) Protocol Studies, (2) Intermail, (3) Domain Name System, (4) Network Operations, and (5) Internet Management.

I.A. TASK OBJECTIVES

The overall objective of this project is to advance the state of the art in network protocols by focused studies of particular problems in the Internet System (e.g., congestion control) by developing prototype implementations of applications and services (e.g., Intermail and the Domain Name System), and by participating in Network Operations and Management activities.

The development of packet communication technology and the DARPA research program in networking have produced both physical systems (the packet switches) and logical systems (software architecture, protocols) that have been adopted for use by the Military (the DDN, for example).

This successful demonstration of a communications concept and its transfer into the operational service provides a point from which new long-term objectives, specifically in the area of very high-speed and large-scale networks, can be contemplated. The extension of powerful communications networks will eventually allow everyone in military service, or civilian service as well, to have access to these very high-speed networks, resulting in a very large network addressing and routing environment.

The Internet Concepts project explored new ways to utilize these packet-switching concepts with emerging technologies and in new environments. The areas of very high-speed networks and very large-scale networks require investigation into practical routing procedures and the incorporation of new technologies for providing enhanced performance in order for them to support the many specialized command and control requirements of the Military.

The nature of work in the Internet Concepts project has evolved over time as the active technical issues changed. The current operational Internet system has grown and is still growing rapidly, stressing the need for current procedures for routing and system management.

As the Internet has expanded, new types of applications have been added that necessitate evaluation as to whether the basic mechanisms continue to perform reliably, efficiently, and predictably. Several issues were investigated.

I.B. TECHNICAL PROBLEM

As the Internet has grown some problems that we previously ignored or deferred now require solutions. One of these problems is routing, another is congestion control. Work by Van Jacobson at LBL [RFC-1185] has improved the performance of TCP in the presence of congestion. However TCP is only one transport protocol, a solution is needed at the IP level.

The demonstration of prototype applications and services in the Internet shows how new capabilities can be provided and encourages others to develop other new applications or services. The Intermail application demonstrates the interoperation of mail systems that were designed separately and are incompatible. The Domain Name Service demonstrates a tree-structured, partially replicated, distributed, database service using datagram queries.

Advancing the state of the art in networking requires the operation of high performance networks such as the Los Nettos regional network and the DARTnet national testbed network.

Of course, the Internet requires some ongoing management activity, and ISI does its part by being involved in the standards process and the preparation of RFCs.

I.C. GENERAL METHODOLOGY

The general approach taken by ISI on protocols and networking issues is to work at two levels. First, we develop new protocols (or extensions to existing protocols), and second, we develop prototype applications or services to stress test the new protocols.

An example of this is our congestion control study, which extends the IP protocol and needs to be stress tested by advanced applications that use UDP as well as other applications that use TCP. The Domain Name Service is a heavy user of UDP and may stress test the congestion control extensions we have developed.

I.D. TECHNICAL RESULTS

I.D.1. Protocol Studies

I.D.1.1. Congestion Control

One task area of the Internet Concepts Research project is to study protocol issues of current concern. We have looked at the problems that can arise when the system messages of the network (e.g., routing information updates) are corrupted. IP-level congestion control was also studied.

Currently, although there is a flow-control mechanism in the TCP protocol at the transport level, there is no congestion-control mechanism built into the Internet Protocol (IP), which is the only universal protocol of the Internet. There is an ICMP source quench (SQ) message designed for congestion control, but it has not been utilized in the IP layer yet.

A study of the behavior of Internet throughput when the SQ message is used for congestion control was undertaken. A network simulator sufficient to test these congestion control ideas was completed. The simulator allowed testing of several different approaches for IP congestion-control that depend upon the host IP module maintaining separate congestion control queues indexed by destination network. For that purpose, it is assumed that gateways generate Source Quench messages when their queues overflow. That information, when it returns to the source host, is used to clock out IP packets to the affected destination network much the way cars are clocked onto freeways when a freeway is congested. The simulator allowed testing of these approaches under differing topologies and loads, and with differing link characteristics. We developed a congestion-control algorithm based on the SQ message. The results of these simulations are reported in Greg Finn's paper, "A Connectionless Congestion Control Algorithm" [11].

The initial results indicated that, while the IP-level congestion-control algorithm is very effective in controlling congestion and is reasonably efficient in using available bandwidth, it unfairly distributes that bandwidth among competitive sources. This unfair distribution was not expected. The probability of receiving an SQ message from a congested gateway seemed out of proportion to that source's use of the congested path.

A modification was made that assumed that gateways return SQ messages for packets drawn at random from their overflowing queues, rather than always returning an SQ message for the packet that caused the overflow. This modification appears to improve fairness. The probability of a source's receiving an SQ message is related to the recent reception history of the overflowing queue. A large series of simulation runs demonstrated a marked improvement in distribution of resource. A comparison of the distribution of process-finishing times shows a much narrower spread when the random SQ approach is used. This new gateway policy is called "random drop".

Based on these successful simulation results, a plan for *in vivo* testing of the IP/SQ algorithm in a LAN setting was developed. Our earlier simulation result suggested that a

measure of fairness could be added to the IP/SQ algorithm if gateways 1) always generate SQ messages when discarding a message due to queue overflow, and 2) they randomly choose from that queue the message to discard rather than always choosing the one that caused the overflow to occur. To perform these tests required extensive modifications to the UNIX kernel IP driver and another set of kernel modifications to simulate a gateway.

The IP/SQ algorithm was implemented within the UNIX kernel for the Sun release 4.0 operating system. The pseudo-gateway was a dedicated Sun workstation with a modified UNIX kernel that implemented certain controls and features needed during testing. We required it to generate an SQ message whenever its queues overflowed and we required it to act as if it were configured with output channels of selected bandwidth and delay. We also implemented the "random drop" policy in the pseudo-gateway. This required moving a random number package into the UNIX kernel. We extracted the linear congruential random number package from the UNIX C library random package and installed it in the pseudo-gateway's kernel.

We conducted a formal series of tests, involving large scale monitored transfers from several Suns running the modified IP/SQ kernel through another Sun acting as a pseudo-gateway. These *in vivo* tests involved up to four hosts, sometimes on different local area networks, simultaneously transmitting equal-sized files through the pseudo-gateway whose behavior was controlled. It was also necessary to determine whether or not the operation of IP/SQ interfered significantly with Van Jacobson's modified TCP, which has its own congestion control algorithm. To avoid possible interference from restrictive transmission windows, TCP was not used for one series of the data transfers to create a believable "control" set of measurements. The same series of file transfers was then repeated with TCP and IP/SQ operating. Finally, the series was repeated again with TCP but without IP/SQ.

Hundreds of test runs were done, both with and without random drop operating in the pseudo-gateway, across a range of round-trip delays from 0.02 to 1.0 seconds and up to four simultaneous sources. As gateway capacity was changed between test runs, the amount of source data transferred was altered to keep the ideal fair completion time constant. We found that IP/SQ utilized 86% of available gateway capacity with a standard deviation of 6%. The algorithm also operated quite fairly. Mean completion time was 854 seconds with a standard deviation of 89 seconds. The interference effects between IP/SQ and TCP appear small.

1.D.1.2. NNStat and Statspy

Our work on network monitoring has produced software packages called *NNStat* and *statspy*. These programs allow one to monitor the traffic on ethernet and collect the results for analysis. The typical analysis consists of making histograms by packet type or by source or destination address.

NNStat is the network statistics gathering program that is used in the NSFNET backbone nodes (the NSSs) to measure the NSFNET traffic, as well as in a number of other regional and campus networks, and by a number of network experimenters.

Several releases of the *NNStat* package, incorporating a number of important extensions were made:

- Support for the Sun 4 hardware, support for a PC RT, access control, and subnet support.
- The much-requested ability to count bytes as well as packets.
- The extended language features for boolean expressions and case statements described in Bob Braden's SIGCOMM '88 paper [1].
- A major revision to the internal compilation algorithms, to improve performance and to support the entire extended language described in Bob Braden's SIGCOMM '88 paper [1].

An ISI technical report on *statspy* internal design was also completed; this is an expansion and revision of the SIGCOMM paper. It includes a more complete discussion of compilation algorithms and a brief overview of the program internals [3].

1.D.1.3. Attack Resistant Protocols

As long-haul computer networks become an accepted part of the day-to-day operations of business, the military, and the government, they become important and even vital to those operations. As use of networks grows, it is a wise precaution to assume that malicious attempts to sabotage the network will occur. A carefully constructed hardware and software system cannot protect itself indefinitely against internal failure or human intervention. Distributed network operating software should not make the network susceptible to widespread failure if one router, or even several, deviate from acceptable behavior. Network software should be resistant to this manner of attack while preserving the desirable network attributes of flexibility and efficiency.

A routing procedure that widely disseminates routing information among routers raises the possibility that a malicious attack or failure at any single router could severely disrupt the network. The shortest-path routing procedures examined by ISI showed the vulnerability to attack of some current network layer protocols. ISI developed an example of an attack-resistant protocol in an attempt to respond to the question: "How might one design network layer controlling protocols so that they are highly resistant to attack either accidentally or by deliberate intervention?" Our results are discussed in technical report ISI-RR/88-201 [10].

1.D.1.4. Policy Routing

We prepared guidelines and recommendations for DARPA on Policy-Based Routing [7]. There are many difficult problems with policy-based routing, and we focused on the knowledge that is needed to make the routing decisions, and where that knowledge must be located. We especially focused on the practical issues for the typical network topology of a regional network connected to multiple long-haul networks and to multiple

campuses, where two or more groups on each campus have different "rights" to use different long-haul networks.

I.D.2. Intermail

Under the charter to perform experiments in protocols and protocol interoperability, we had previously developed a system to forward electronic mail between the Internet (or ARPA-Mail) world and several commercial mail systems (e.g., Telemail, MCI-Mail). We continued to expand the capabilities of this experimental system and to operate it in order learn about such protocol interoperation.

The Intermail experiment in protocol interoperability links commercial mail systems with the Internet mail system [RFC-822]. It provides an electronic mail relay service between several commercial systems and the Internet. This service was developed and tested over several years. This has been a successful experiment and provides a useful service. One example is the communication of manuscripts for publications in IEEE Computer Society publications between authors on the Internet (primarily at universities) and the publication editors on COMPMAIL.

The Intermail mail-forwarding system was transferred to the new Commercial Mail Relay (CMR) project of the ISI Computer Center. In addition to Telemail, the Dialcom systems IEEE COMPMAIL, USDA, and CGNET are now currently operating under this system. The MCI-Mail relay is still operating on TOPS-20, but was moved from C.ISI.EDU to A.ISI.EDU. Some additions to the Intermail mail-forwarding system were made so that it would use a dial-out modem to exchange mail with two commercial mail systems, the IEEE Compmail system and the NSFMAIL system. In addition, some of the forwarding functions are being taken over by the new Commercial Mail system that is being developed by the ISI Computer Center. During these transitions, we provided consulting services to the Computer Center regarding CMR, gave talks at conferences, and wrote articles to inform the Internet community of this useful service [19][20][21][23][24][25][27][30] and RFC-1168].

I.D.3. The Domain Name System

In response to the growing difficulty of maintaining the data file on hosts in the Internet and their names and services, we had previously developed the Domain Name System (DNS). The DNS is a distributed database with partially replicated data, used to query for information indexed by a tree-structured name. Currently, the names are the names of host computers in the Internet and the related information comprises the network address and the services implemented.

The DNS is now fully integrated into the operation of the Internet. While the bulk of the tens of thousands of hosts involved in the Internet use the "Bind" software for the DNS distributed by Berkeley, the root servers use the "Jeeves" software developed by ISI. Understandably, some ongoing effort is required to support this key software. New domain software was installed at the TOPS-20 root servers to correct some minor bugs and allow a larger database. A paper on the DNS entitled "Development of the Domain

Name System," was published by Paul Mockapetris [15] as well as several RFCs [RFCs 1034, 1035, 1101].

1.D.3.1. DOC and DIG (Domain Tools)

To further our work in the design and efficient implementation of the DNS system, a series of questions were formulated to focus on various issues of system performance. A new DNS query tool was implemented to gather the necessary data. This tool, *dig* (domain information groper), is a command line tool which queries DNS servers in either an interactive or a batch mode. It was developed to be more convenient and flexible than "nslookup" for gathering performance data and testing DNS servers. Its features and options include most of those provided by nslookup, and several others. It is available via anonymous FTP from venera.isi.edu, file: pub/dig.1.0.tar.Z. Currently, an asynchronous version is being implemented so data can be collected in a more timely fashion.

The latest version of the DNS query tool (*dig* version 2.0) was also made available for anonymous FTP from venera.isi.edu, file: pub/dig.2.0.tar.Z. This version includes support for zone transfer, a more convenient way to make an address to domain name query, and various bug fixes.

An automated tool for testing the configuration of DNS nameservers was developed. The first implementation, "Doc", is a shell script that uses *dig* to query the nameservers for a specified domain. Doc (version 1.0) primarily tests that delegation information is consistent between the authoritative and delegating nameservers for a given domain. Doc is available for anonymous FTP from venera.isi.edu, file: pub/doc.1.0.tar.Z.

1.D.3.2. The US Domain

The US Domain is an official top-level domain in the Domain Name System (DNS) of the Internet community. It is registered with the Network Information Center at SRI International. The US Domain and all of its subdivisions are managed by the Domain Registrar at ISI.

The naming scheme of the US Domain is based on political geography, that is, the US Domain is subdivided into states, then cities, and so on. Any computer in the United States may be registered in the US domain.

Typical host names in the US domain are:

DOGWOOD.ATL.GA.US
GRIAN.CPS.ALTADENA.CA.US

Many of the names registered in the US Domain are DNS style names for computers in other systems. To make use of this feature, hosts on systems such as UUCP, and BITNET, must register their hosts with an Internet host. A mail exchanger (MX) record is then added to the US Domain zone file pointing to the Internet host for forwarding. The forwarding host must be directly on the Internet (that is, have an IP address). An example of a host entry in the zone file when using an MX record would be:

JOES-HOST.ACADEM.HOU.TX.US MX 10 GAZETTE.BCM.TMC.EDU

The US Domain is currently supported by four name servers: venera.isi.edu, vaxa.isi.edu, hercules.csl.sri.com, and nnsf.nsf.net.

To date, we have approximately 300 registered Hosts in the US Domain.

I.D.4. Network Operations

I.D.4.1. Los Nettos

In response to the changing technology possibilities for applications and service, one of our goals was to explore higher speed regional and long-haul networking as a basis for new applications (e.g., packet video, shared screens).

A user-supported regional network was formed in the Los Angeles area in 1988 to provide connectivity between sites such as individual campuses and research centers in the area, and to provide connectivity to long-haul networks for all the campuses and centers.

This regional network for the Los Angeles area is *Los Nettos* [[18][22][26][28][29]]. *Los Nettos* connects to several long-haul and national networks, such as the NSFNET, the WBNET, and the DRI. *Los Nettos* is robustly linked to the NSFNET through CERFnet. The goal and motivation of *Los Nettos* is to provide a very high-speed, very low-delay network, that provides high-quality TCP/IP-based packet communication to its members. All connections and links are (and will be) at least T1 (1.5 Mbps) rate. *Los Nettos* is operated by the member organizations. Its creation has permitted a substantial reduction in the cost of network access in the region. In particular, the number of ARPANET IMPs in the region was reduced from eight to one.

During the first phase, five organizations joined *Los Nettos*: Caltech, ISI, TIS, UCLA, and USC. During the second phase JPL, TRW, RAND and IBM joined the network. During the last year, UNISYS and NOSC were added, but IBM Scientific Center closed its connection due to relocation of the staff. Currently, there are ten *Los Nettos* members.

I.D.4.1.1. Network Evolution

The first lines were installed at the end of October 1988 and by early December the network was operational with five sites and five lines up. For the first four sites the topology was a closed loop that provided a redundant path if any one of four links failed. This proved valuable when one site had to have power work done, and during a T1 link failure. It also provided additional bandwidth across the diagonal path using the load sharing capability of the cisco gateways.

We experienced minor delays in the installations of the lines, and significant delays in the delivery of the Datatel CSU/DSUs. We had serious support problems with Datatel early in this project, but these problems were resolved. There were compatibility problems between the Datatels and cisco routers. Fortunately, cisco was very supportive in helping us resolve such problems. These problems were fixed with version 1.2 Datatel firmware.

Los Nettos started with a connection to the ARPANET, and later through cooperation with CERFnet a connection to the NSFNET. The routing protocol configuration has changed several times in significant ways. In the course of these routing procedure developments, we detected a bug in the EGP updates from the core that was corrected by BBN, and some limitations in the cisco routing filter list that were subsequently extended by cisco.

As the ARPANET was scaled down and IMPs were removed from Los Nettos sites, we experienced a few temporary routing problems as site routing was changed to use Los Nettos. We worked with CERFnet to optimize the routes and provide backup and alternate routing. Cooperation between CERFnet and Los Nettos continues to be good.

The last ARPANET node in the Los Angeles area (at ISI) was turned off on 22 March 1990. Nationwide Internet access for Los Nettos is provided via CERFnet and NSFNET.

The DOE ESnet was installed through the Los Angeles area with nodes at Caltech and UCLA. Los Nettos is exchanging routing information with ESnet at Caltech. Los Nettos is getting only a handful of routes from the ESnet connection at Caltech and none through UCLA's connection. Los Nettos, therefore, does not use ESnet for any significant load.

The California Department of Water Resources (DWR) provided a north-south link interconnecting BARRnet and CERFnet that will be available for testing at the end of this year. BARRnet has installed OSPF within their network. This will allow them to install the appropriate routing controls to make this alternate path safe from long term routing loops.

I.D.4.1.2. Network Monitoring

Early in this project, we developed a procedure that periodically pings each interface in the network and can display a map of the network with non-communicative elements highlighted. This procedure is called *mon*.

The MIT SNMP development kit was brought up on a Sun workstation. It has the ability to show routes in both directions and to show default routes. We have found SNMP useful for debugging routing problems and for obtaining traffic statistics programmatically. An SNMP trap daemon was started for monitoring trap events. An SNMP-based tool for character graphic display of network routing status was created. It tells us if Los Nettos member networks are known by the NSFNET, if default routes are known by each node and, if optimal routes are being taken. Our SNMP-based route monitoring tool proved valuable for quickly showing problems that would affect performance and reliability but not connectivity [8].

One such example is when we experienced a routing problem with two of our networks. Although there were no links down to indicate that a problem existed, the SNMP-based route-monitoring tools immediately alerted us to the existence of a problem. SNMP-based diagnostic tools helped us to quickly isolate the problem and resolve it. The *ping* program was not enough to verify proper operation of a network. The CMU SNMP package was also adapted for monitoring Los Nettos.

For remote access to monitoring and control of the Los Nettos equipment we have developed a method for accessing the console ports of the cisco router and CSU/DSU's at remote member sites. A single dial up line is attached to a low cost any-to-any port selector. We can power cycle the routers and CSU/DSU's via dialup access to initiate a power reset by selecting a specific port on the port selector. Equipment for providing this remote console access to member ciscos and CSU/DSUs was packaged and configured. We installed these remote console access kits at Rand, Caltech, UNISYS in Camarillo, TIS, USC, TRW, and NOSC.

1.D.4.2. DARTnet

DARPA has established a new research network testbed for studying new protocols including policy-based routing, congestion and flow control, and other advanced protocol features. As planning for this testbed proceeded, several names were used: 1) Internet-Testbed, 2) T1GARNET, or "tiger-net", and finally, 3) DARTnet.

ISI assisted DARPA with planning experiments for this network including new gateways for routing experiments. Many teleconferences were organized to gather information about candidate hardware/software platforms for the open gateways and then decide what platform to use. Topics in these meetings included the schedule for hardware installation, the development of the necessary gateway and experimental software, and possible ways to facilitate collaboration among groups.

During this period, the DARTnet concept was developed, and the planning, equipment selection, and installation arrangements were made. Actual installation will not begin until the end of 1990. ISI performed system engineering and planned network operations center functions. We also worked with others to develop a plan of networking experiments to be performed using DARTnet once it is operational [[4][5][6]].

1.D.5. Internet Management

1.D.5.1. Internet Infrastructure and IAB Support

ISI provided support for Internet management. In particular, a portion of Bob Braden's effort during this period was concerned with serving as Executive Director of the IAB, as head of the End-to-End Research Group, and as a member of the IRSG. For the IAB, Bob arranged meetings and produced the meeting notes, action item lists, and conducted e-mail votes on issues between meetings. Pursuant to the interests of the End-to-End Research Group, Bob began to organize an internet experiment that could make use of a T1 testbed built on the DRI lines [2]. Also, Jon Postel serves as RFC Editor and a member of the IAB. He helped produce the IAB Official Standards Summaries represented by RFC-1130 and RFC-1140.

In addition, ISI assembled and distributed (via e-mail) the Internet Monthly Report, maintained mailing lists for the IAB and its Task Forces, and coordinated the assignment of protocol parameters (e.g., type codes and port addresses) [RFC-1060].

1.D.5.2. Host Requirements RFC

To clarify and tighten the specifications of the protocols used in the Internet, and to improve the interoperability of the products made by the numerous vendors now using these specifications, a document called the "Host Requirements" was developed. The Host Requirements document was split into two RFC's [RFCs 1122, 1123], one covering the link layer through the transport layer, and the other covering the application and support protocols. Completion of the two RFC's involved resolution of a number of difficult issues.

A commentary document that collects together the loose ends from the Host Requirements RFC was published itself as an RFC entitled, *A Perspective on the Host Requirements RFC* [RFC-1127].

1.D.5.3. Requests for Comments

ISI serves as the technical editor and "publisher" of the Internet document series called "Requests for Comments" (RFCs). During this contract, 151 RFCs were published. A complete listing of these RFCs is in Appendix I.

ISI research staff authored 21 RFCs. These RFCs were 1034, 1035, 1042, 1046, 1060, 1068, 1071, 1072, 1084, 1101, 1111, 1121, 1122, 1123, 1127, 1130, 1135, 1140, 1168, 1175, and 1177.

1.E. IMPORTANT FINDINGS AND CONCLUSIONS

Our experiments in IP-level congestions control have had promising results. We have developed and demonstrated a method of deriving stable congestion feedback from IP Source Quench messages, and we contributed to the understanding of fairness issues in gateways. We believe these experiments ought to be carried forward to a field trial over paths with significant delays and realistic traffic flows, i.e., using DARTnet.

Our NNStat network statistics package is in use at many sites, including the NSFNET backbone, for gathering traffic statistics, for traffic analysis, and for problem diagnosis. This variety of applications is due to the flexibility and ease of use of NNStat's *statspy* program. At a very small time cost, we have integrated into the NNStat distribution a number of useful extensions developed at half a dozen sites. Finally, we have found NNStat to be a flexible tool for network research, making measurements on traffic characteristics and models. The most significant problem with *statspy* has been the Sun promiscuous Ethernet interface software (NIT), whose quality has gotten steadily worse with later releases of Sun OS. To continue to use NNStat as a network research tool, it is necessary to modify *statspy* to use the new Berkeley packet filter instead of NIT.

Our work on organizing and installing DARTnet has been central to the project. Under a succeeding contract, ISI will set up and operate a Network Operations Center (NOC) for DARTnet, and we will continue to coordinate the research program using this valuable experimental facility. We also expect to make substantial contribution to the DARTnet

research program, both through our own experiments and also by providing our SPARC-based video teleconferencing software to other sites.

The INTERMAIL service has demonstrated the communications operability of electronic mail between very different systems and protocols. Our *dig* program has proved to be a useful tool in debugging the datastructures of the Domain Name System. Our work on many other areas (RFCs, US Domain, Los Nettos, etc.) has provided important contributions to the development of the Internet and the extension of the TCP/IP protocol suite.

I.E. SIGNIFICANT HARDWARE DEVELOPMENT

No significant hardware was developed in this effort.

I.G. SPECIAL COMMENTS

We believe that the success of the Internet and the TCP/IP protocol suite is a great success story for DARPA-sponsored research. A key factor in this is the open publication of the standards and other technical literature of the system as RFCs.

I.H. IMPLICATIONS FOR FURTHER RESEARCH

I.H.1. Protocol Studies

I.H.1.1. Source Quench

The testing of the IP/SQ algorithm in a small network setting is now finished. These tests have confirmed our expectations. Congestion seems well controlled and random drop improves fairness to what we believe is an acceptable level. In addition, little interference with the Van Jacobson TCP congestion control seems to take place. Nevertheless, questions about the validity of the recently completed tests remain.

Since DARTnet will not be available before the termination of this contract, a pseudo-gateway was used for the *in vivo* IP/SQ tests so as not to harm existing networks with congestion testing nor require modification of actual ISI gateways, which was considered impractical. Although the pseudo-gateway was carefully programmed to simulate an actual gateway, with variable load, capacity and channel delay, because it was not in fact a gateway there may be some hidden issues concerning the behavior of the IP/SQ algorithm not yet discovered.

The availability of DARTnet will allow testing of the IP/SQ algorithm under actual network conditions. Under a new contract, we shall conduct a final series of tests to provide the data and confidence needed to determine whether or not to proceed with widespread adoption of IP/SQ as a part of the IP standard. We shall port the Sun 3 version of the IP/SQ kernel modifications to Sun SPARCstation and modify the DARTnet gateways to implement random drop. Once this is done we shall gather statistics on a series of transfers under congested conditions. The tools developed for DARTnet testing

by other DARPA contractors for generating traffic loads and for gathering statistics will aid us in this effort. At the conclusion of the testing a paper that outlines the results will be submitted to the Journal of Internetworking Research and Experience.

1.H.1.2. NNStat and statspy

To continue to use NNStat as a network research tool, it is necessary to modify *statspy* to use the new Berkeley packet filter instead of NIT.

1.H.2. Domain Name System

Although the DNS is in production use and is difficult to change, other naming systems may provide the impetus for additions. There is a need for a system design that allows for robustness and survivability, especially in terms of protection against bad information.

Support for X.500 style addresses for mail, etc., could be constructed as a layer on top of DNS without the sophisticated protection, update, and structuring rules of X.500.

1.H.3. Network Operations

1.H.3.1. Los Nettos

Los Nettos is currently operating at T1 line speeds. In the future, Los Nettos is looking to take the network to higher speeds such as T3, once tariffs and hardware make it practical.

1.I. INC REPORTS AND PAPERS

- [1] Braden, R., "A Pseudo-Machine for Packet Monitoring and Statistics", ACM SIGCOMM '88 Conference, Stanford, CA, August 1988.
- [2] Braden, R., "Internet End-to-End Task Force", Internet Activities Board (IAB), January 1989.
- [3] Braden, R., *A Packet Monitor Program*, USC/Information Sciences Institute, ISI/RR-88-207, March 1990.
- [4] Braden, R., "Research Testbeds", *National Net '90 Conference*, Washington, D.C., March 15, 1990.
- [5] Braden, R., "TCP/IP: Current Implementation and Future Evolution", Presented at the TCP/IP & OSI : The Interoperability Challenge Symposium, Washington, D.C., May 1990.
- [6] Braden, R., "Future Development of the Internet", *Interop '90 Conference*, San Jose, CA, October 10, 1990.
- [7] Cohen, D., R. T. Braden, P. Mockapetris, and J. Postel, "Policy Based Routing (PBR) for the NNT", Report to Ira Richer, DARPA/ISTO, May 1989.

- [8] Cohen, D., W. Prue, "Network Monitoring in Los Nettos", Fifth Annual Workshop on Networks and Supercomputers, San Diego, CA, July 1989.
- [9] DeSchon, A., and R. T. Braden, "Background File Transfer Program (BFTP)", *ConneXions — The Interoperability Report*, Vol. 2, No. 12, December, 1988.
- [10] Finn, G., *Reducing the Vulnerability of Dynamic Computer Networks*, USC/Information Sciences Institute, RR-88-201, June 1988.
- [11] Finn, G., "A Connectionless Congestion Control Algorithm", *ACM Computer Communications Review*, October 1989.
- [12] Mockapetris, P., "Domain Practices", 2nd TCP/IP Interoperability Conference, Arlington, VA, December 1987.
- [13] Mockapetris, P., "Simple Mail Transfer Protocol", *Computer Communications Standards*, William Stallings, ed., Macmillan, Chapter 5, Vol. 3, 1988.
- [14] Mockapetris, P., "The Domain Name System", NNSC Domain Tutorial, Cambridge, MA, August 1988.
- [15] Mockapetris, P., "Development of the Domain Name System", ACM SIGCOMM '88, Stanford, CA, August 1988.
- [16] Postel, J., G. Finn, A. R. Katz, and J. K. Reynolds, "An Experimental Multimedia Mail System, ACM Transactions on Office Information Systems, Vol. 6, No. 3, pp. 62-67, March 1988.
- [17] Postel, J., D. Cohen, and W. Prue, "Los Nettos", Third IEEE Workshop on Metropolitan Area Networks, Dana Point, California, March 1989.
- [18] Westine, A., "Los Nettos", *Users' Directory of Computer Networks*, University of Texas at Austin, Section 3, p. 20, July 1989.
- [19] Westine, A., "CMR - Commercial Mail Relay" and "Los Nettos - Los Angeles Research Network", *%@: A Directory of Electronic Mail Addressing and Networks*, pp. 58-61, O'Reilly and Associates, August 1989.
- [20] Westine, A., "Making Traffic Move in Los Angeles", *CERFnet News*, p. 5, August-September 1989.
- [21] Westine, A., "Making Traffic Move in Los Angeles", *Merit/NSFNET Link Letter*, Vol. 2, No. 4, September 15, 1989.
- [22] Westine, A., "Los Nettos" *Internet Resources Guide*, NSF NNSC, Section 5.4, p. 1, September 1989.

- [23] Westine, A., "Commercial Mail Relay (CMR)", *Internet Resources Guide*, NNSC, Section 5.20, p. 1, September 1989.
- [24] Westine, A., "TELEMAIL, DIALCOM, and MCI MAIL Instructions", March 1990.
- [25] Westine, A., A. DeSchon, J. Postel, and C. Ward, *Intermail and Commercial Mail Relay Services*, USC/ISI, RR-90-254, March 1990.
- [26] Westine, A., "Los Nettos: The High Performance Regional Network for Los Angeles", *NSF Network News*, April 1990.
- [27] Westine, A., "Commercial Mail Relay (CMR)", *!%@:: A Directory of Electronic Mail Addressing and Networks*, 2nd Edition, Pg. 70-73, O'Reilly and Associates, Inc., May 1990.
- [28] Westine, A., "Los Nettos - Los Angeles Research Network", *!%@:: A Directory of Electronic Mail Addressing and Networks*, 2nd Edition, Pg. 66-67, O'Reilly and Associates, Inc., May 1990.
- [29] Westine, A., "Los Nettos", *NSF Network News*, The NSF Network Service Center, No. 8, September 1990.
- [30] Westine, A., "Intermail and Commercial Mail Relay Service," *XVIII ACM SIGUCCS User Services Conference*, Cincinnati, Ohio, Sep 29-Oct 3, 1990.

II. Multimedia Conferencing project

Project Leader: Stephen Casner

Research Staff:

*Annette DeSchon
Eve Schooler
David J. Walden*

Research Assistants:

Brian Hung

Support Staff:

*Celeste Anderson
Kathleen McLaughlin*

II.A. TASK OBJECTIVES

The Multimedia Conferencing project had task objectives in three areas: architecture and protocols, prototypes, and media support. To this end, the project developed an architecture for multimedia conferencing and protocols, servers, and other facilities and for implementing systems using this architecture, developed prototype systems using this architecture, including the exchange of text, pictures, and voice data over digital packet-switched networks, and built and installed video systems at the DARPA offices in Washington, D.C. and at SRI in Menlo Park, CA. The specific objectives for each area are listed below:

II.A.1. Architecture and protocols

Connection Management
ST-II

II.A.2. Prototypes

Packet Teleconference System
Installation of teleconference sites at DARPA and SRI
Extensions for multipoint teleconferencing
Improved Resolution
Improved Audio
Port of packet voice and video processing to less expensive platform

II.A.3. Media support

Refinements to video system implementation
Echo canceller
Image scanning and printing

II.B. TECHNICAL PROBLEM

The purpose of multimedia conferencing is to enhance the productivity of cooperating users through the use of computers and networking. Multimedia conferencing is the next logical step in computer-assisted cooperation following previous steps taken in computer text mail and multimedia mail. The real-time nature of multimedia conferencing and the need to support multiple conferees are two significant added dimensions that provide new research challenges. This is more than a quantitative change in speed since it implies the

need for concurrent activity among all participants, rather than one-at-a-time delivery used for mail. In addition, it uses media, such as video and speech, which require more support outside of the traditional computer environment.

A typical multimedia conference is a combination of linked workstations, video, and voice. Each channel serves a distinct and useful purpose, and a system that can integrate all three not only reduces the amount of time required for a teleconference setup, but is also more easily replicated outside of the research environment.

Previously, each medium was supported by a separate conferencing system with separate facilities that required separate setup and used separate system components. This project proposed to tackle this problem by creating standards and models for a new conferencing infrastructure that would promote experimentation. Also, essential to this effort, was the creation of prototype systems with simple applications that could be useful services in themselves.

II.C. GENERAL METHODOLOGY

The work was divided into three areas of research: 1) architecture and protocols, 2) prototypes, and 3) media support.

Our exploration into architecture and protocols was accomplished by creating standard building blocks and models for the conferencing infrastructure.

The prototyping effort was aimed at trying out a variety of choices, and providing early versions of simple applications that could be useful services in themselves. Such early use was essential for services to evolve from experimental to production status.

In the area of media support, the expansion of the teleconferencing video system by two additional sites was accomplished by providing system integration of the new video systems, installing the system at the new sites, and then developing modifications to the video protocols to permit the image seen by one camera to be distributed via a satellite channel to more than one destination at a time.

Low-cost image scanning and printing options for image processing as an adjunct to the video channel for less capable hosts was developed and software was expanded to allow compression, transmission, and printing of facsimile-like images of a variety of commonly available laser printers. The system was then integrated into existing multimedia mail and conferencing systems.

II.D. TECHNICAL RESULTS

II.D.1. Architecture and Protocols

II.D.1.1. Connection Management

One objective of the project was to develop an integrated multimedia architecture to support a variety of distributed multimedia applications, including conferencing. To this end,

we collaborated with BBN on the integration of ISI's MMCC and BBN's Diamond/MMConf system.

The MMCC user interface was reconstructed using the Diamond Toolkit software as the first step in that integration. A design for a merged MMCC/MMConf conference manager was prepared and discussed with BBN in a teleconference.

During the course of this contract, the conference control program (MMCC), which runs in a small window on the Sun workstation screen, has been expanded substantially. Its control of local video camera and display was supplemented with control of remote camera selection. To control camera switching and full-screen/quadrant selection, we implemented a video control panel using software buttons in a small window sharing the Sun screen with MMConf/Diamond. In addition, it now performs an entirely new function of multisite teleconference connection management. Voice and video connections are established through MMCC using a mouse-and-button user interface. We expect this program will make the conference system much easier for novice users to manage without assistance from the implementors. This is particularly important as additional sites are installed.

The eventual goal is for MMCC to merge with the already existing MMConf program from BBN. MMConf also runs on the Sun workstation and coordinates distributed execution of the Diamond multimedia document editor for shared editing of documents in a conference mode. The MMCC functions were redesigned for wider application as part of the integrated program. The newly released Versatile Message Transaction Protocol (VMTP) with IP multicasting was considered as a replacement for MMCC's experimental multisite conferencing protocol, but we deferred action due to system requirements.

A report describing the multimedia conferencing project was published in a special issue of the *ACM SIGOIS Bulletin* on Computer-Supported Cooperative Work to serve as a guide for additional sites considering installation of the system [7]. Another paper on the effect of long distances on groupware design was published in the *Proceedings of the Groupware Technology Workshop* [8].

II.D.1.2. ST-II

We participated in the ST and Connection-Oriented IP Working Group of the Internet Engineering Task Force. This group worked on the definition the next version of the ST protocol used to carry packet video and packet voice. The goal was to get ST implemented more widely in the Internet so that teleconferencing and other applications can be used at more locations. The ST protocol specification was put into RFC-ready form through teleconferences and e-mail between ISI and BBN. After a comment period in the IETF Internet-Drafts directory, the draft was published as RFC-1190, *Experimental Internet Stream Protocol, Version 2 (ST-II)*.

II.D.2. Prototypes

II.D.2.1. Packet Teleconference system

One objective of the project was to develop an easy-to-use, robust, integrated multimedia teleconferencing system in order to encourage new users to participate in the system's use. To this end, improvements were made to the multimedia conference control program (MMCC) for more robust control of the conference through communication between MMCC, the Voice Terminal program (VT), and Packet Voice program (PVP). Additionally, an autopilot mode was incorporated into MMCC to allow experts to control connections remotely and in turn make it easier to get conferences going for new users at remote sites.

II.D.2.1.1. Conference control

The workstation conferencing environment consists of MMCC, the conference control program; MBFTPTool, a program for background distribution of documents among conference sites; and BBN's MMConf shared workspace. We improved coordination among the three programs so that the conference initiation takes the form of a single connection of the participants.

This coordination centers around a conference identifier assigned on a per-conference basis. It, in turn, is used as the name of a "conference directory" to which shared applications connect and in which shared files are placed. The conference identifier is distributed to participating sites by MMCC, the multimedia conference control program, which also makes sure the conference directory exists at each site. MMCC communicates the conference identifier and participant list to MBFTPTool. MMCC can also automatically establish and disconnect an MMConf session in parallel with voice and video connections, bringing up MMConf in the conference directory. If users create new shell windows during a conference for side activities, these shells will set the conference directory as their working directory.

The control protocol used by MMCC was enhanced to combat pathological behavior caused by network partitioning. Additional timeouts were put in place to avoid getting stuck waiting for another site to respond, and mechanisms were added for detecting state inconsistencies and restoring state information when sites reconnect after a partition. In MBFTPTool, a mechanism to coordinate with MMConf the distribution of files to the conference sites and a new user interface feature to provide "hints" to make it easy for the user to determine which parameters must be explicitly entered were added.

MMCC allows users to control whether or not other sites are allowed to join a conference, and to restrict remote control of cameras. MBFTPTool now has a conference mode, which sets special transfer modes and defaults for teleconferencing (as opposed to stand-alone) use. A mechanism for submitting file transfers to take place at a later time was also added.

A new version of the background FTP package, BFTP, was released, including the BFTPTool and the MBFTPTool. This version is now fully in the public domain since the

command parsing routines with restricted distribution were replaced with command parsing code from UC Berkeley and time parsing code from IBM/CMU, which may be freely distributed.

II.D.2.1.2. Network issues

We received the Image 30 video codec boards from Concept Communications in February 1988. We then installed packet video on these boards. We modified the on-board firmware of the Image 30 to implement the HDLC serial line protocol required to interface to the Butterfly. The Image 30 boards gave us several advantages compared to the ISI-built experimental codecs we used previously. First, they produced color video instead of black and white. Second, we were able to expand quickly from two to four teleconference sites. Third, the Image 30 takes two camera inputs, so we have at each site the room-view camera plus a copystand camera.

In order to accommodate the the packetization scheme for this new codec, substantial changes to the packet video program (PVP) that runs in the Butterfly were required. The major difference between processing for the Image 30 and the ISI experimental codec was that the frame rate is constant for the Image 30 and that the segments of each video frame must be delivered in order and uninterrupted. (Damaged frames may be discarded without loss of synchronization.) PVP was also modified to enable it to interface with the new ST Gateway being developed by BBN as part of the transition to the terrestrial Wideband Network. The new interface was successfully tested against the ST gateway in coordination with BBN. PVP was then modified to run as a background process that can communicate with a separate foreground program when operator control is required. This allowed the program to be rebooted automatically if a failure such as a power outage should occur, while making access easier when manual intervention is required to correct problems.

To allow establishment of the point-to-point, video-only calls (with in-band audio) that are required for the path to the UK, the Host Control Protocol (HCP) was incorporated into the multimedia conference control program, MMCC, and PVP during the last quarter of this contract. The new HCP control path will also allow implementation of multi-site conferencing with PictureTel codecs, but with the limitation that only one site is viewed at a time. HCP will be used to communicate with both PVP and VT in the SPARCstation.

II.D.2.2. Installation of ISI video systems at DARPA and SRI

One of MMC's efforts was to expand the two existing packet video sites to four. To accomplish this goal, we ordered commercial video codecs to replace the two experimental codecs built at ISI. We implemented modifications and extensions in several components of the system, and teleconference sites were installed at the DARPA offices in Washington, DC in mid-1988 and at SRI in late 1988, completing the planned system of four sites (BBN, DARPA, ISI, and SRI). Most of the Internet researchers who are likely to use the system are now reasonably close to at least one site.

The multimedia teleconference facility in Washington was eventually moved from the DARPA building at 1400 Wilson Blvd. to 1555 Wilson Blvd. in early 1989. The new room

was much larger and had a much better setup than the temporary installation at 1400. This move had been planned as a means to avoid contention for the room at 1400. It made teleconferences involving the Washington site much easier to arrange. However, the teleconference room at DARPA was moved back to the 12th floor at 1400 Wilson Blvd. in early 1990, because the space at 1555 Wilson was lost. A second room on the 1st floor, with lower noise and more space but also less availability, was set-up in parallel.

The SRI multimedia teleconferencing facility was also moved a short distance to RIACS in Mountain View, CA in September 1990. This was done so that the ST Gateway could be moved to NASA Ames where it connects to FIX-West. Also in September, a new teleconferencing site has been installed at University College, London, connected through the "UK Fat Pipe" and three ST gateways.

II.D.2.3. Extensions for multipoint teleconferencing

The on-board firmware of the Image 30 video codec from Concept Communications was modified to accept a merged stream of data packets from up to four sites and display each site in its own quadrant of the screen. The standard product Image 30, with its circuit-switched interface, can only handle a single site.

Modifications were also made to match the new codec scheme in the packet video program (PVP) that runs in the Butterfly. Incoming packet streams from the remote sites are queued separately for sorting and delivery to the codec. Packets are marked for the desired quadrant according to their source site.

The multimedia teleconferencing system made its official debut in a conference of more than two sites in 1988. The Autonomous Networks Task Force used it during a two-day meeting among members located at DARPA, ISI, and BBN. All three sites were displayed simultaneously in quadrants of the video screen. At each site, voice from the remote sites was mixed for playback, allowing all sites to talk at once if they wished. The Diamond/MMConf system was used for shared document display and editing among the three sites.

Test conferences were conducted with all four sites participating, but we determined that the Wideband Network was not robust enough at that traffic load for a real conference. This was due to performance limits in the Earth Station Interface (ESI). When the Wideband Network was converted to operation over terrestrial T1 lines in mid-1989, the ESI was no longer be part of the system. Delays were much lower with the Terrestrial Wideband Network (TWBNet) and this allowed us to conduct three-and four-site conferences regularly.

As multipoint conferencing progressed we modified PVP to automatically refresh the images of all conference participants upon the addition of each new participant to the conference. This "paints" a complete image of new participants rather than having images "build" slowly as motion occurs in parts of the scene. Similarly, PVP requests refreshes if it detects any lost packets. This removes any image anomalies that may be caused by the missing data. PVP has also been modified to use both point-to-point and conference modes of ST protocol connections to prepare for testing with the Butterfly ST Gateway.

II.D.2.4. Improved Resolution

One of the project's requirements was to demonstrate improved resolution in the conference video image. Resolution is a function of the commercial video codec and not of the components built by the project. The Image 30 codec has low delay and a fast frame rate (30/sec) compared to other codecs, but its resolution is lower and there are some unusual motion effects. We explored the prospects for improved resolution in the Concept Communications' Image30 codec, which was the codec in use at the beginning of this project.

Concept Communications was working on an improved version of the Image30 codec to give double the resolution in both dimensions. While this improvement was in progress, we augmented the packet video system to work with the Widcom codecs already owned by DARPA. These codecs provide better static resolution, but were limited to two-site conferences.

We adapted the packet video system to test the Widcom codecs. Since the Widcom is less tolerant of lost data than the Image30 codec, the packet order-restoration algorithm in the Packet Video Host (PVP) was enhanced to handle ST packets up to four packet positions late. This algorithm was also made more robust to tolerate the arrival of duplicate packets. Careful synchronization was required between PVP and the Widcom codec to insure the smooth flow of data with no underrun or overrun.

Using the Widcom codecs and the Terrestrial WBnet, we held a teleconference with DARPA in May of 1989 to demonstrate the video quality. While the static resolution of the Widcom was better than the Concept codec we use for multi-site video, motion quality was insufficient at a 56Kb/s data rate.

We then made arrangements with CLI, PictureTel, and VideoTelecom for loans of their codecs in order to test and demonstrate their use with packet video. We developed serial line protocol conversion software on a PC-based coprocessor board with four high-speed serial ports. This software interprets the various line-framing protocols of the codecs and encapsulates those frames in HDLC line framing to interface with PVP running in the Butterfly [1]. These improvements to PVP and the addition of a codec control panel to MMCC to select among the demonstration codecs and to set the data rate allowed us to test three other codecs in addition to the Widcoms.

We conducted a parallel demonstration of Compression Labs and PictureTel codecs between ISI and DARPA. The PictureTel codec was also used for the second of three DARPA/ISTO staff meetings held by teleconference between ISI and DARPA during 1989. The three-way comparison including the VideoTelecom codec was completed during the third DARPA/ISTO staff meeting.

As a result of these demonstrations, the PictureTel codec was selected as the best combination of picture quality and adaptability to the packet network. It provides twice the image resolution of the Image30, and allows filling between received frames so it can tolerate the asynchrony between transmit and receive clocks that results from packet

transmission. Its primary disadvantage is a higher internal delay, about 1/2 second, which impairs interactivity.

None of the three codecs tested allow multisite conferencing with all sites viewed simultaneously, as we can the Image30. Therefore, the new PictureTel codecs were installed in parallel with the Image30.

PictureTel video codecs were installed at all sites. These codecs provide better video quality, but can only be used for 2-site teleconferences. The existing Concept codecs are still available in parallel for multi-site, multi-quadrant teleconferences. Users may choose whichever codec best meets their needs. At conference initiation, the caller determines the codec type and video rate, then MMCC synchronizes these choices across all sites. Once a conference is set up, changes to these values are disallowed. RPC-based codec and crossbar servers were installed at all teleconferencing sites to switch video between the two codecs. Resolution of the Image30 may also be improved through the installation of a second-generation compression algorithm under development by Concept. In the future we plan to explore variable frame rate and/or data rate to try to minimize the motion effects. We may also be able to improve resolution by trading off some of the frame rate.

To allow better video motion fidelity when there are only two sites in conference, a new version of the packet video host, PVP, was released that allows codec types and codec data rates to be changed during an on-going ST connection. This facilitates quick comparisons of various codecs without requiring frequent loading and starting of different sets of PVP modules.

II.D.2.5. Audio Improvements

Shure acoustic echo cancellers on loan were used to allow conference participants to listen to a loudspeaker rather than using headphones. This was an advantage, but the echo cancellation was just barely adequate and sound quality was degraded somewhat. Background noise and room acoustics both affect the performance of the canceller.

NEC echo cancellers were then installed at ISI and DARPA. These echo cancellers eliminate the need for conference participants to wear headphones at these sites. They also allow people who can't get to one of the teleconference sites to participate by telephone in the audio portion of a teleconference.

II.D.2.6. Port of packet voice and video processing to less expensive platform

Our next major effort in the project will be to port the real-time packet voice and video processing from the Butterfly multiprocessor to a less expensive and more widespread hardware platform. The NeXT machine was an appealing candidate for such a platform as was the Sun SPARCstation.

We developed a test program to investigate the sound I/O facilities on the SPARCstation and the NeXT machine. The program digitizes audio using the built-in codec, then either

loops it back internally to the local speaker or encapsulates it in UDP packets that are shipped over the Internet. Packets are either played back on another machine's built-in speaker or bounced off a remote host and played back locally. There was a noticeable delay of approximately one fourth of a second on the NeXT machine when the test program was run in internal loopback mode. We anticipate that this delay may be reduced in future releases of the NeXT operating system by additional DMA modes for the transfer of data between the DSP and the main processor. Delay in the SPARCstation was very low with small buffers. The addition of the network code seemed to add very little delay, compared to internal loopback. In tests run on our local Ethernet, the packet loss was between 0% and 1% and the speech was of satisfactory quality.

After running these tests, we selected the Sun SPARCstation as the first platform. We are also coordinating with MIT LCS, where packet video will be implemented on a 386 PC-AT hardware base, so our implementations can interoperate.

We then ported the Voice Terminal (VT) and Packet Video Host (PVP) programs to the SPARCstation. With this software and the built-in /dev/audio device on the SPARCstation, we were able to hold real-time audio conversations over ISI's local Ethernet. The programs form Stream Protocol (ST) Point-to-Point and Multi-Point connections using raw IP sockets to encapsulate ST packets in IP.

A new version of the Packet Video Host, PVP, was developed for testing. The packet-reordering algorithm that corrects for lost, out-of-order and late packets has been reimplemented to allow more flexible segmenting of video frame data. This allows sending larger packets at a slower rate, which will be important when sending video from SPARCstations across ethernets and other IP networks. A new debugging feature was also added to intentionally mis-order, delay or drop out-going packets in a specified pattern. This feature has been helpful in testing the new packet-reordering algorithm. For test purposes, PVP is also using the /dev/audio device as its source of real-time data. Once the High Speed Interface board and device driver are available, we will be able to connect video codecs to the SPARCstation so that PVP can send video data as intended. We are ready to begin testing the SPARCstation-VT with the BBN Butterfly-VT as soon as the new version of the ST gateway that supports IP encapsulation is deployed by BBN.

II.D.3. Media Support

II.D.3.1. Refinements to video system implementation

Video resolution was increased 25% in the horizontal dimension so the pixels now have a 1:1 aspect ratio. The packet video data rate was also increased, taking advantage of improved performance of the Wideband Net. The change to a commercial codec improved video quality.

II.D.3.2. Echo Cancellor

We designed and implemented an echo canceller algorithm ourselves to eliminate echo in a separate part of the packet voice system: hybrid echo in the Switched Telephone Net-

work Interface (STNI) cards. These cards allow calls to be placed over the packet voice system from standard telephones.

In this interface a "hybrid circuit" is required to match a two-wire telephone line to the separate receive and transmit paths of the digital packet network. Because the hybrid circuit cannot match perfectly, an echo is produced and must be suppressed. We designed and implemented an echo canceller algorithm in the NEC uPD7720 Signal Processing Interface signal processing chip for installation on the interface board. The echo canceller allows full-duplex communication instead of the half-duplex operation currently employed for echo suppression.

Simulations of the algorithm were done with different input signals, amplitudes, loop gain, and attenuation. The results were very useful in pointing out the value of loop gain necessary for a stable system and also for a particular convergence rate.

The canceller is now in operation after having been tested over the Wideband Net to tune algorithm constants for best performance. It is implemented on an NEC 77P20 digital signal processing chip mounted on the STNI card. The STNI, developed by ISI and BBN, sends and receives TouchTones and digitizes voice signals to interface between a telephone line and the packet voice terminal program. Echo occurs when the far-end signal is reflected from the 2-to-4 wire hybrid circuit required in each STNI to couple voice signals from the packet system into the two-wire telephone network. The echo canceller uses the transverse filter technique to estimate the echo and subtracts this estimate from the near-end signal. A document describing the theory and implementation of the echo canceller is available [[2][3][4]].

II.D.3.3. Image scanning and printing

Document scanning for teleconferencing at ISI is performed using a document scanner system based on an IBM-PC. We prepared a document describing the implementation of this system so it can be replicated elsewhere [5].

II.E. IMPORTANT FINDINGS AND CONCLUSIONS

One of our important findings was how to interface the commercial codecs to the packet system. We developed the technique of asynchronous interface to the commercial codecs and we achieved implementation of framing conversion in the ATCOMM board.

Over 100 teleconferences were held, 28 of them were multisite between 1 November 1987 and 30 September 1990. Various organizations including the Internet Activities Board (IAB), the Internet Engineering Task Force (IETF), the IRTF, and DARPA staff members have met using the teleconferencing system.

Several observations were made during these conferences that have allowed us to evaluate the effectiveness of different aspects of the system. It was felt that three-site conference were effective, but that participants needed to concentrate more, resulting in a more tiring session than in previous teleconferences with two sites. Some new mechanisms are being considered to help make such conferences easier to manage.

Although audio requires less bandwidth than video, audio is the more difficult medium to implement satisfactorily. It is important to provide full-duplex audio, either with headphones or an echo canceller. Because headphones can be annoying, echo cancellers are preferred but their cost is high and their performance may be marginal except in specially prepared rooms.

The resolution of the Image 30 codec is sufficient for room view, but not for viewgraphs on a copy stand. The best solution for this problem is to view the images on the MMConf shared workspace, either by scanning in the viewgraphs, or from material that originates on-line.

II.F. SIGNIFICANT HARDWARE DEVELOPMENT

ISI developed the TMS320-based codec for our initial experiments with packet video. It served well for that purpose, but was not practical to replicate; so we converted to a commercial codec made by Concept Communications.

While ISI did not build the hardware of the Concept codec, the project team did modify the firmware. We also integrated the ATCOMM board into the packet video system to connect the commercial codecs to the Butterfly.

II.G. IMPLICATIONS FOR FURTHER RESEARCH

Conferencing is useful, but still only works when the participants are close to the sites. We want to put conferencing on everyone's desktop by taking advantage of the high-speed packet interface on the workstation and the growing packet network infrastructure. Codec chips are being developed that will be integrated into the workstations to make the hardware cost reasonable.

Once personal conferencing is on a much larger scale, the conference traffic associated with it could cause severe congestion, if it is not managed properly. This is an area that needs to be addressed concurrently with the development of personal conferencing systems.

Additional research will be needed on the integration of ST-like services into IP-based networks. Packet voice and video need the guaranteed performance that ST provides, but ST is much different from IP. The IP networks need to support the growing demand for multimedia conferencing capability, but we need to determine whether it is best to install a protocol-like ST in parallel with IP, or to synthesize the two to make a new internet protocol that provides integrated services.

The groundwork for a coordinated plan for connection-oriented experiments on DARTNET has already been laid.

II.H. MMC REPORTS AND PAPERS

- [1] Casner, S., K. Seo, W. Edmond, and C. Topolcic, "N-Way Conferencing with Packet Video", Presented at RIACS, February 1990, Xerox PARC,

March 1990 and The Third International Workshop on Packet Video, Morristown, NJ, March 1990.

- [2] Hung, B., "ISI Internet Wideband Net Packet Voice Echo Canceler User Manual", Internal working paper, September 1988.
- [3] Hung, B., "ISI Internet Wideband Net Packet Voice Echo Canceler Software", Internal working paper, September 1988.
- [4] Hung, B., "A Voice Echo Canceler Implementation for A Packet Voice System", Internal working paper, September 1988.
- [5] Hung, B., "Multimedia Document Scanner System on a IBM-PC AT", Internal working paper, September 1988.
- [6] Schmandt, C., and S. Casner, "Phonetool: Integrating telephones and workstations", *GLOBECOM '89, Proceedings*, IEEE Communications Society, November 1989.
- [7] Schooler, E., and S. Casner, "Packet-Switched Multimedia Conferencing System", *ACM SIGOIS Bulletin*, pp. 12-22, January 1989.
- [8] Schooler, E., and S. Casner, "Multimedia Conferencing in the Internet: the Effect of Long Distances on Groupware Design", *Proceedings of the Groupware Technology Workshop*, IEIP Working Group 8.4, Palo Alto, CA, August 1989.

III. Protocol Accelerator project

Project Leader: Paul Mockapetris

Research Assistants:
Steve Hotz

Support Staff:
Kathleen McLaughlin

III.A. TASK OBJECTIVES

Typical networks have protocol processing software in hosts and switching nodes. The capabilities of protocol processing elements and links vary greatly, but the end-to-end performance of services seen by applications is limited by a series of delays: processing delays in the source host (scheduling, protocol processing); transmission delays on links; queueing and processing delays in switching nodes; and protocol processing and scheduling delays in the destination host.

The task objective was to define a method for accelerating the protocol processing task in hosts so that end-to-end performance improvements could keep track improvements in line signaling rates and new switching designs.

III.B. TECHNICAL PROBLEM

The problem was to define a general method for accelerating protocol processing that could be included or added to conventional workstation designs and deliver cost-effective improvements.

III.C. GENERAL METHODOLOGY

The work assumed compatibility with existing standards as a given. As a consequence, the effort had to span the entire system architecture, and work was divided into four areas:

1. *Architecture and methodology analysis*: evaluation of the costs and effectiveness of a variety of techniques for improving protocol processing performance. Selection of protocol tasks to be accelerated and simulation to demonstrate benefits.
2. *Hardware efforts*: design and implementation of a pair of test systems.
3. *Software efforts*: implement of protocol processing software for the hardware.
4. *Experiments*: evaluation of the demonstration system in several configurations.

III.D. IMPORTANT FINDINGS AND CONCLUSIONS

No fundamental difficulties were encountered in engineering a front-end which would offer significant performance improvements, however the system integration effort was much more difficult than anticipated, and protocol changes due to increasing emphasis on standardization as well as technical innovation obsoleted parts of the design before others could be begun.

We refocussed our efforts on optimizations which we felt could offer significant benefit but were seen as relatively insensitive to the evolution and standardization process. The problem we selected was the selection of the optimal server and address from a set of equivalent servers, some with multiple addresses. Our measurements and models revealed that this could be of considerable benefit for datagram services, but that choices could often be constrained by outside factors in more common services.

III.E. SIGNIFICANT HARDWARE DEVELOPMENT

No significant hardware was developed.

III.F. SPECIAL COMMENTS

After 15 months of research, on 13 February 1989, ISTO decided to discontinue work on this idea. The ideas formulated would work in a software environment, but the hardware development necessary to make the protocol accelerator ideas competitive with commodity networking chips was determined by ISI and DARPA staff to be impractical in this research environment.

III.G. IMPLICATIONS FOR FURTHER RESEARCH

Various software and hardware techniques for optimizing or expediting protocol processing can work, as demonstrated in this effort and others. The key issue is identifying opportunities where the benefit is worth the effort. This is relatively simple if the task to be accelerated is viewed as static, but in the real world this is not so; protocols change, operating systems interfaces change, and the benefit of any special-purpose implementation must compete with the always greater effort expended on improvements on general purpose CPUs and portable protocol implementations. Obsolescence is a constant danger; an accelerator for a particular version of an operating system may take as long to develop as the next version of the operating system.

Future work should focus on efforts that incorporate acceleration techniques into the portable protocol implementations and harness the power of general purpose parallel machines rather than attempting to develop custom hardware. Specific niches will exist where this is not so, and silicon compiler advances may enlarge these niches, but adding performance to protocol processing through pipelining and parallelization techniques should leverage off of general purpose efforts in these areas.

IV. Supercomputer and Workstation Communication

Project Leader: Stephen Casner

Research Staff:

Alan Katz

Support Staff:

Kathleen McLaughlin

The Supercomputer and Workstation Communication project was a one-year effort that explored new techniques for using high-capacity networks to communicate between personal workstations and remote supercomputers. As a part of this effort, the SWC project evaluated and tested the use of standard IP/TCP-based protocols, as well as the newer remote window protocols such as the X Window System and Sun Microsystems' NeWS. The DARPA Wideband Satellite Network provided the high-capacity, long-distance communication path for many of these tests. User requirements of the scientific community that will be using these systems were also investigated.

IV.A. TASK OBJECTIVES

The main objective of the Supercomputer Workstation Communication project was to develop new techniques to allow personal workstations to communicate with remote supercomputers (such as a Cray 2) over high-capacity networks. It is common for supercomputer centers to provide high-bandwidth access for local users over Ethernets, Hyperchannels, or other high-speed networks, but remote users are often constrained by low-data-rate links. The DARPA Wideband Satellite Network, with its 3 Mb/s channel speed, provided the combination of high capacity and long distance we needed for tests of remote access.

IV.B. TECHNICAL PROBLEM

The project addressed the need for effective, high-bandwidth communication between powerful remote computers and their users. The need for continued research in this area is demonstrated by the rapid growth of the Defense Data Network (DDN) (and more recently of NSFNET), and especially by the advance of technology that allows higher speed long-haul networks.

IV.C. GENERAL METHODOLOGY

Our research was divided into four areas: 1) evaluation and testing of existing remote access protocols, 2) exploration of the possibility of an Intelligent Communication Facility that would allow users to connect the output of a program running on one machine to the input of a program running on a different machine, 3) research into supercomputer and workstation interaction, and 4) prediction of future user requirements of the scientific community that will be using workstations to access supercomputers.

IV.D. TECHNICAL RESULTS

The results of this project are described in detail in the final report submitted at the end of this project, *Supercomputer Workstation Communication*, USC/Information Sciences Institute, SR-89-235, June 1989.

IV.E. IMPORTANT FINDINGS AND CONCLUSIONS

Our testing of new and existing protocols will help to determine which protocols should be used in the supercomputer/workstation environment. The development of a system similar to our proposed Intelligent Communication Facility will allow users to run existing programs on different machines without having to rewrite them or to become knowledgeable about low-level network protocols.

The current trend away from the use of large, timeshared mainframes and toward the use of powerful workstations connected to high-speed networks indicates that more and more researchers will use the new remote window protocols such as X and NeWS. Our studies in this area have given examples of how to best use these protocols and have identified several major issues that must be resolved in order for them to work.

IV.F. SIGNIFICANT HARDWARE DEVELOPMENT

No significant hardware was developed.

IV.G. IMPLICATIONS FOR FURTHER RESEARCH

There is ample opportunity for follow-on work. A General Internet remote execution protocol must be developed. This protocol should allow for user authentication and for security. Such a protocol is needed for the remote window protocols to be used effectively, as well as for applications such as our split editor and the Intelligent Communication Facility.

An Internet-based batch protocol is also needed. The National Center for Atmospheric Research (NCAR) has developed such a system on top of existing protocols, but this system requires the user to register a password for his local machine with the batch processor. A separate protocol is really needed, perhaps a part of a general remote execution protocol.

Much more work should be done to test and tune both X and NeWS on supercomputers over high-capacity networks. At the time of our study, neither X nor NeWS was available on any of the NSF-sponsored supercomputers, but they should become available in the near future. It is unknown how much the widespread use of these protocols will load the supercomputers and the networks.

A standard must be defined to allow for the interchange of equations among electronic mail messages, various document preparation systems, and symbolic mathematics systems. As more and more scientists become connected to the Internet, they will demand the ability to send mathematical expressions via electronic mail. It would be useful to prototype the Intelligent Communication Facility described ISI/SR-89-225. Having such a system in addition to X or NeWS would allow more effective use of both remote and local resources.

IV.H. SWC REPORTS AND PAPERS

- [1] Katz, A., and S. Casner, *Supercomputer Workstation Communications*, USC/Information Sciences Institute, SR-89-235, June 1989.

V. Appendix

Request for Comments

- RFC 1028: Davin, J., J. Case, M. Fedor, and M. Schoffstall, *A Simple Gateway Monitoring Protocol*, November 1987.
- RFC 1029: Parr, G., *A More Fault Tolerant Approach to Address Resolution for A Multi-LAN System of Ethernets*, University of Ulster, May 1988.
- RFC 1030: Lambert, M., *On Testing the NETBLT Protocol over Diverse Networks*, MIT Laboratory for Computer Science, November 1987.
- RFC 1031: Lazear, W., *MILNET Name Domain Transition*, MITRE, November 1987.
- RFC 1032: Stahl, M., *Domain Administrators Guide*, SRI International, November 1987.
- RFC 1033: Lottor, M., *Domain Administrators Operations Guide*, SRI International, November 1987.
- RFC 1034: Mockapetris, P., *Domain Names – Concepts and Facilities*, USC/ISI, November 1987.
- RFC 1035: Mockapetris, P., *Domain Names – Implementation and Specification*, USC/ISI, November 1987.
- RFC 1036: Horton, M., and R. Adams, *Standard for Interchange of USENET Messages*, December 1987.
- RFC 1037: Greenburg, B., and S. Keene, *NFILE – A File Access Protocol*, December 1987.
- RFC 1038: St. Johns, M., *Draft Revised IP Security Option*, IETF, January 1988.
- RFC 1039: Latham, D., *A DoD Statement on Open Systems Interconnection Protocols*, DoD, January 1988.
- RFC 1040: Linn, J., *Privacy Enhancement for Internet Electronic Mail: Part I: Message Encipherment and Authentication Procedures*, IAB Privacy Task Force, January 1988.
- RFC 1041: Rekhter, J., *Telnet 3270 Regime Option*, T.J. Watson Research Center, IBM, January 1988.
- RFC 1042: Postel, J., and J.K. Reynolds, *A Standard for the Transmission of IP Datagrams over IEEE 802 Networks*, USC/ISI, February 1988.
- RFC 1043: Yasuda, A., and T. Thompson, *TELNET Data Entry Terminal Option DODIIS Implementation*, Defense Intelligence Agency, February 1988.

- RFC 1044: Hardwick, K., and J. Lehashman, *Internet Protocol on Network Systems HYPERchannel Protocol Specification*, February, 1988.
- RFC 1045: Cheriton, D., *VMTP: Versatile Message Transaction Protocol, Protocol Specification*, Stanford University, February 1988.
- RFC 1046: Prue, W., and J. Postel, *A Queuing Algorithm to Provide Type-of-Service for IP Links*, USC/ISI, February 1988.
- RFC 1047: Partridge, C., *Duplicate Messages and SMTP*, CIC at BBN Labs, February 1988.
- RFC 1048: Prindeville, P., *BOOTP Vendor Information Extensions*, McGill University, February 1988.
- RFC 1049: Sirbu, M., *A Content-Type Header Field for Internet Messages*, CMU, March 1988.
- RFC 1050: *RPC: Remote Procedure Call Protocol Specification*, Sun Microsystems, Inc., April 1988.
- RFC 1051: Prindeville, P., *A Standard for the Transmission of IP Datagrams and ARP Packets over ARCNET Networks*, McGill University, March 1988.
- RFC 1052: Cerf, V., *IAB Recommendations for the Development of Internet Network Management Standards*, NRI, April 1988.
- RFC 1053: Levy, S., and T. Jacobson, *Telnet X.3 PAD Option*, Minnesota Supercomputer Center, April 1988.
- RFC 1054: Deering, S., *Host Extensions for IP Multicasting*, Stanford University, May 1988.
- RFC 1055: Romkey, J., *A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP*, June 1988.
- RFC 1056: Lambert, M., *PCMAIL: A Distributed Mail System for Personal Computers*, MIT, June 1988.
- RFC 1057: *RPC: Remote Procedure Call Protocol Specification, Version 2*, Sun Microsystems, Inc., June 1988.
- RFC 1058: Hedrick, C., *Routing Information Protocol*, Rutgers University, June 1988.
- RFC 1059: Mills, D., *Network Time Protocol (Version 1) Specification and Implementation*, University of Delaware, July 1988.
- RFC 1060: Reynolds, J.K., and J. Postel, *Assigned Numbers*, USC/ISI, March 1990.
- RFC 1061: Not issued yet.

- RFC 1062: Romano, S., M. Stahl, and M. Recker, *Internet Numbers*, SRI International, August 1988.
- RFC 1063: Mogul, J., C. Kent, C. Partridge, and K. McCloghrie, *IP MTU Discovery Options*, July 1988.
- RFC 1064: Crispin, M., *Interactive Mail Access Protocol - Version 2*, SUMEX-AIM, July 1988.
- RFC 1065: Rose, M., and K. McCloghrie, *Structure and Identification of Management Information for TCP/IP-Based Internets*, TWG, August 1988.
- RFC 1066: McCloghrie, K., and M. Rose, *Management Information Base for Network Management of TCP/IP-Based Internets*, TWG, August 1988.
- RFC 1067: Case, J., M. Fedor, M. Schoffstall, and J. Davin, *A Simple Network Management Protocol*, August 1988.
- RFC 1068: DeSchon, A., and R. Braden, *Background File Transfer Program (BFTP)*, USC/ISI, August 1988.
- RFC 1069: Callon, R., and H.W. Braun, *Guidelines for the use of Internet-IP addresses in the ISO Connectionless-Mode Network Protocol*, UMich, February 1989.
- RFC 1070: Hagens, R., Hall, N., and M. Rose, *Use of the Internet as a Subnetwork for Experimentation with the OSI Network Layer*, U of Wisconsin - Madison, The Wollongong Group, February 1989.
- RFC 1071: Braden, R., D. Berman, and C. Partridge, *Computing the Internet Checksum*, September 1988.
- RFC 1072: Jacobson, V., and R. Braden, *TCP Extensions for Long-Delay Paths*, USC/ISI, October 1988.
- RFC 1073: Waitzman, D., *Telnet Window Size Option*, BBN STC, October 1988.
- RFC 1074: Rekhter, J., *The NSFNET Backbone SPF based Interior Gateway Protocol*, T.J. Watson Research Center, IBM, October 1988.
- RFC 1075: Waitzman, D., C. Partridge, (BBN STC), and S. Deering, (Stanford), *Distance Vector Multicast Routing*, November 1988.
- RFC 1076: Trewitt, G., and C. Partridge, *HEMS Monitoring and Control Language*, Stanford University, November 1988.
- RFC 1077: Leiner, B., *Critical Issues in High Bandwidth Networking*, Editor, Gigabit Working Group, November 1988.
- RFC 1078: Lottor, M., *TCP Port Service Multiplexer (TCPMUX)*, SRI-NIC, November 1988.

- RFC 1079: Hedrick, C., *Telnet Terminal Speed Option*, Rutgers University, December 1988.
- RFC 1080: Hedrick, C., *Telnet Remote Flow Control Option*, Rutgers University, December 1988.
- RFC 1081: Rose, M., *Post Office Protocol – Version 3*, TWG, November 1988.
- RFC 1082: Rose, M., *Post Office Protocol – Version 3, Extended Offerings*, TWG, November 1988.
- RFC 1083: IAB, *IAB Official Protocol Standards*, USC/Information Sciences Institute, January 1989.
- RFC 1084: Reynolds, J.K., *BOOTP Vendor Information Extensions*, USC/ISI, December 1988.
- RFC 1085: *ISO Presentation Services on Top of TCP/IP-based internets*, M. Rose, TWG, December 1988.
- RFC 1086: Onions, J., Nottingham, and M. Rose, *ISO-TP0 Bridge Between TCP and X.25*, TWG, December 1988.
- RFC 1087: IAB, *Ethics and the Internet*, January 1989.
- RFC 1088: McLaughlin, L., *A Standard for the Transmission of IP Datagrams over Net-BIOS Networks*, TWG, February 1989.
- RFC 1089: Schoffstall, M. (Rensselaer Polytechnic Institute), C. Davin (MIT), M. Fedor (NYSERNET), and J. Case (University of Tennessee at Knoxville), *SNMP over Ethernet*, February 1989.
- RFC 1090: Ullmann, R., *SMTP on X.25*, Prime Computer, February 1989.
- RFC 1091: VanBokkelen, J., *Telnet Terminal-Type Option*, FTP Software, February 1989.
- RFC 1092: Rekhter, J., *EGP and Policy Based Routing in the New NSFNET Backbone*, IBM T. J. Watson Research Center, February 1989.
- RFC 1093: Braun, H. W., *The NSFNET Routing Architecture*, Merit, February 1989.
- RFC 1094: Sun Microsystems, *NFS: Network File System Protocol Specification*, March 1989.
- RFC 1095: Warriar, U. (Unisys), L. Besaw (Hewlett-Packard), *The Common Management Information Services and Protocol over TCP/IP (CMOT)*, April, 1989.
- RFC 1096: Marcy, G., *Telnet X Display Location Option*, Carnegie-Mellon University, March 1989.

- RFC 1097: Miller, B., *Telnet Subliminal-Message Option*, CMU-NetDev, April 1989.
- RFC 1098: Case, J., (UTENN), M. Fedor, (NYSERNET), M. Schoffstall, (Rensselaer), C. Davin, (MIT), *A Simple Network Management Protocol (SNMP)*, April 1989.
- RFC 1099: Not issued yet.
- RFC 1100: Internet Activities Board, *IAB Official Protocol Standards*, April 1989.
- RFC 1101: Mockapetris, P., *DNS Encoding of Network Names and Other Types*, ISI, April 1989.
- RFC 1102: Clark, D., *Policy Routing in Internet Protocols*, MIT, Laboratory for Computer Science, May 1989.
- RFC 1103: Katz, D., *A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks*, Merit/NSFNET, June 1989.
- RFC 1104: Braun, H-W., *Models of Policy Based Routing*, Merit/NSFNET, June 1989.
- RFC 1105: Lougheed, K., (CISCO), Y. Rekhter, (IBM), *A Border Gateway Protocol (BGP)*, June 1989.
- RFC 1106: Fox, R., *TCP Big Window and Nak Options*, Tandem, June 1989.
- RFC 1107: Sollins, K., *A Plan for Internet Directory Services*, MIT, July 1989.
- RFC 1108: Not issued yet.
- RFC 1109: Cerf, V., *Report of the Second Ad Hoc Network Management Review Group*, NRI, August 1989.
- RFC 1110: McKenzie, A., *A Problem with the TCP Big Window Option*, BBN STC, August 1989.
- RFC 1111: Postel, J., *Request for Comments on Request for Comments, Instructions to RFC Authors*, ISI, August 1989.
- RFC 1112: Deering, S., *Host Extensions for IP Multicasting*, Stanford University, August 1989.
- RFC 1113: Linn, J. (DEC), IAB Privacy Task Force, *Privacy Enhancement for Internet Electronic Mail Part I — Message Encipherment and Authentication Procedures*, August 1989.
- RFC 1114: Kent, S. (BBNCC), J. Linn (DEC), IAB Privacy Task Force, *Privacy Enhancement for Internet Electronic Mail: Part II — Certificate-Based Key Management*, August 1989.

- RFC 1115: Linn, J. (DEC), IAB Privacy Task Force, *Privacy Enhancement for Internet Electronic Mail: Part III -- Algorithms, Modes, and Identifiers*, August 1989.
- RFC 1116: Borman, D. (Cray), Internet Engineering Task Force, *Telnet Linemode Option*, August 1989.
- RFC 1117: Romano, S., M. Stahl, M. Recker, *Internet Numbers*, SRI-NIC, August 1989.
- RFC 1118: Krol, E., *The Hitchhikers Guide to the Internet*, University of Illinois Urbana, September 1989.
- RFC 1119: Mills, D. L., *Network Time Protocol (Version 2) Specification and Implementation*, UDEL September 1989.
- RFC 1120: Cerf, V., *The Internet Activities Board*, NRI, September 1989.
- RFC 1121: Postel, J., (ISI), L. Kleinrock (UCLA), V. Cerf, (NRI) B. Boehm, (UCLA), *Act One - The Poems*, September 1989.
- RFC 1122: Braden, B., *Requirements for Internet Hosts -- Communication Layers*, USC/ISI, October 1989.
- RFC 1123: Braden, B., *Requirements for Internet Hosts -- Application and Support*, USC/ISI, October 1989.
- RFC 1124: Leiner, B., *Policy Issues in Interconnecting Networks*, RIACS, October 1989.
- RFC 1125: Estrin, D., *Policy Requirements for Inter-Administrative Domain Routing*, USC/Computer Science Department, November 1989.
- RFC 1126: Little, M., *Goals and Functional Requirements for Inter-Autonomous System Routing*, SAIC, October 1989.
- RFC 1127: Braden, B., *Perspective on Host Requirements*, USC/ISI, October 1989.
- RFC 1128: Mills, D., *Measured Performance of the Network Time Protocol in the Internet System*, UDEL, October 1989.
- RFC 1129: Mills, D., *Internet Time Synchronization: The Network Time Protocol*, UDEL, October 1989.
- RFC 1130: Postel, J., *IAB Official Protocol Standards*, USC/ISI, October 1989.
- RFC 1131: Moy, J., *The OSPF Specification*, Proteon, October 1989.
- RFC 1132: McLaughlin, L., *A Standard for the Transmission of 802.2 Packets over IPX Networks*, The Wollongong Group, November 1989.

- RFC 1133: Yu, J., and H-W. Braun, *Routing between the NSFNET and the DDN*, Merit Computer Network, November 1989.
- RFC 1134: Perkins, D., *The Point-to-Point Protocol: A Proposal for Multi-Protocol Transmission of Datagrams Over Point-to-Point Links*, CMU, November, 1989.
- RFC 1135: Reynolds, J., *The Helminthiasis of the Internet*, USC/ISI, December 1989.
- RFC 1136: Hares, S., and D. Katz, *Administrative Domains and Routing Domains A Model for Routing in the Internet*, Merit/NSFNET, December 1989.
- RFC 1137: Kille, S., *Mapping Between Full RFC 822 and RFC 822 with Restricted Encoding*, University College London, December 1989.
- RFC 1138: Kille, S., *Mapping Between X.400 (1988) / ISO 10021 and RFC 822*, University College London, December 1989.
- RFC 1139: Hagens, R., *Echo function for ISO 8473*, IETF-OSI Working Group, January 1990.
- RFC 1140: Postel, J., *IAB Official Protocol Standards*, Internet Activities Board, USC/ Information Sciences Institute, May 1990.
- RFC 1141: Mallory, T., and A. Kullberg, *Incremental Updating of of the Internet Checksum*, BBN Communications, January 1990.
- RFC 1142: Not issued yet.
- RFC 1143: Bernstein, D., *The Q Method of Implementing TELNET Option Negotiation*, NYU, February 1990.
- RFC 1144: Jacobson, V., *Compressing TCP/IP Headers for Low-Speed Serial Links*, LBL, February 1990.
- RFC 1145: Zweig, J. (UTUC), and C. Partridge (BBN), *TCP Alternate Checksum Options*, February 1990.
- RFC 1146: Zweig, J., (UTUC) and C. Partridge (BBN), *TCP Alternate Checksum Options*, March 1990.
- RFC 1147: Stine, R., *FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices*, Sparta, Inc., April 1990.
- RFC 1148: Kille, S., *Mapping between X.400(1988) / ISO 10021 and RFC 822*, UCL, March 1990.
- RFC 1148: Kille, S., *Mapping between X.400(88) and RFC-822*, University College London, March 1990.

- RFC 1149: Waitzman, D., *A Standard for the Transmission of IP Datagrams on Avian Carriers*, BBN STC, April 1990.
- RFC 1150: Malkin, G., *F.Y.I. on F.Y.I. - Introduction to the F.Y.I. Notes*, March 1990.
- RFC 1150: Malkin, G. (Proteon), and J. Reynolds (ISI), *FYI on FYI: Introduction to FYI Notes*, March 1990.
- RFC 1151: Partridge, C., (BBN), and R. Hinden (BBN), *Version 2 of the Reliable Data Protocol (RDP)*, April 1990.
- RFC 1152: Partridge, C., *Workshop Report Internet Research Steering Group Workshop on Very-High-Speed Networks*, BBN, April 1990.
- RFC 1153: Wancho, F., *Digest Message Format*, WSMR, April 1990.
- RFC 1154: Robinson, D., and R. Ullman, *Encoding Header Field for Internet Messages*, Prime Computer, Inc., April 1990.
- RFC 1155: Rose, M., (PSI) and K. McCloghrie (HUGHES), *Structure and Identification of Management Information for TCP/IP-based Internets*, May 1990.
- RFC 1156: McCloghrie, K., (HUGHES) and M. Rose, (PSI) *Management Information Base for Network Management of TCP/IP-based Internets*, May 1990.
- RFC 1157: Case, J. (SNMP), M. Fedor (PSI), M. Schoffstall (PSI), and J. Davin (MIT), *A Simple Network Management Protocol (SNMP)*, May 1990.
- RFC 1158: Rose, M., *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*, Performance Systems International, May 1990.
- RFC 1159: Nelson, R., *Message Send Protocol*, Clarkson University, June 1990.
- RFC 1160: Cerf, V., *The Internet Activities Board*, CNRI, May 1990.
- RFC 1161: Rose, M., *SNMP over OSI*, Performance Systems International, Inc., June 1990.
- RFC 1162: Satz, G., *Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542) Management Information Base*, cisco Systems, Inc., June 1990.
- RFC 1163: Loughheed, K. (cisco), and Y. Rekhter (IBM), *A Border Gateway Protocol (BGP)*, June 1990.
- RFC 1164: Honig, J. (CORNELL), D. Katz (MERIT)/NSFNET), M. Mathis (PITT/SUP), Y. Rekhter (T.J. WATSON), and J. Yu (MERIT/NSFNET), *Application of the Border Gateway Protocol in the Internet*, June 1990.

- RFC 1165: Crowcroft, J. (UCL), and J. Onions (Nottingham University), *Network Time Protocol (NTP) Over the OSI Remote Operations Service*, June 1990.
- RFC 1166: Kirkpatrick, M. Stahl, M. Recker, *Internet Numbers*, July 1990.
- RFC 1167: Cerf, V., *Thoughts on the National Research and Education Network*, CNRI, July 1990.
- RFC 1168: Westine, A., A. DeSchon, J. Postel, and C. Ward, *Intermail and Commercial Mail Relay Services*, USC/ISI, July 1990.
- RFC 1169: Cerf, V. (IAB), and K. Mills (NIST), *Explaining The Role of GOSIP*, August 1990.
- RFC 1171: Perkins, D., *The Point-to-Point Protocol for the Transmission of Multi-Protocol Datagrams Over Point-to-Point links*, UC Davis, July 1990.
- RFC 1172: Perkins, D., R. Hobby, *The Point-to-Point Protocol (PPP) Initial Configuration Options*, UC Davis, July 1990.
- RFC 1173: VanBokkelen, J., *Responsibilities of Host and Network Managers A Summary of the 'Oral Tradition' of the Internet*, FTP Software, Inc., August 1990.
- RFC 1174: Cerf, V., *IAB Recommended Policy on Distributing Internet Identifier Assignment and IAB Recommended Policy Change to Internet 'Connected' Status*, CNRI, August 1990.
- RFC 1175: Bowers, K. (CNRI), T. LaQuey (UTEXAS), J. K. Reynolds (ISI), K. Roubicek (BBNST), M. Stahl (SRI), and A. Yuan (MITRE), *FYI on Where to Start - A Bibliography of Internetworking Information*, August 1990.
- RFC 1176: Crispin, M., *Interactive Mail Access Protocol - Version 2*, University of Washington, August, 1990.
- RFC 1177: Malkin, G. (FTP Software, Inc.), A. Marine (SRI), and J. Reynolds (ISI), *FYI on Questions and Answers - Answers to Commonly Asked 'New Internet User's Questions*, August 1990.
- RFC 1178: Libes, D., *Choosing a Name for Your Computer*, Integrated Systems Group/ NIST, August 1990.
- RFC 1179: McLaughlin, L., Ed., *Line Printer Daemon Protocol*, The Wollongong Group, August 1990.
- RFC 1181: Blokzijl, R., *RIPE Terms of Reference*, RIPE/NIKHEF, September 1990.
- RFC 1183: Everhart, C. (TRANSARC), L. Mamakos (UNIV. MARYLAND), R. Ullman (PRIME COMPUTER), and P. Mockapetris (ISI), *New DNS RR Definitions*, October 1990.