

Cohen

ISI 121

J. Postel  
ISI  
25 October 1979

Internet Meeting Notes - 10, 11, 12 & 13 September 1979

I. INTRODUCTION TO UCL - Jones

Ron welcomed us to UCL and discussed the arrangements for facilities, refreshments, and lunch.

II. OVERVIEW AND OBJECTIVES - Cerf

Vint reviewed the need for a working internet system and cited the performance of the system as a key issue. The development of procedures to measure and test the performance of the internet are now important. Vint has pointed out that several conceptual design issues need to be resolved, for example, protocol support for voice conferencing in the internet.

III. STATUS REPORTS

A. DARPA - Cerf

Vint reported that the Director of the ARPA Information Processing Techniques Office, Col. Dave Russell retired on 31 August. The new Director is Dr. Robert Kahn.

The standardization of IP and TCP for use as the DOD networks protocol is progressing. Dr. Robert Lyons at DCEC is the Executive Agent. The only changes that are necessary are to incorporate mechanisms to pass the security and precedence information between the application program and TCP, and between TCP and IP. There already is an IP option to carry that information in the internet.

There is a need for a way to test TCP implementations and applications that use TCP without taking a service machine down. For TOPS20s this will be satisfied by using BBNF. BBNF is currently only on the BBN RCCNET but will be connected to the ARPANET for October through December 1979.

Vint acknowledged that there have been hardware delivery delays that have had a substantial impact on the internet program. The delivery schedules from manufactures are very long.

The program managers in the ARPA office are coordinating to have all new applications use internet protocols.

B. BBN - Haverty, Plummer, Strazisar, McNeill, Flood Page

Jack reported that the BBN-Unix TCP is in regular use. Currently work is in progress to provide a user interface to the IP layer. The IP support is available as a subroutine package. The interface will be similar to the TOPS20 interface documented in IEN 108. BBN is also working on a VAX UNIX with the goal of having the current 11/70 UNIX and VAX UNIX be compatible at the applications level.

Bill reported on the plan to attach BBNF to the ARPANET for ease in testing TCP and TCP using applications. In doing this testing it is preferred to have one person at a time to simplify the identification of problems. Note that one will have to use new 1822 (96-bit leaders) to access BBNF since it will be on IMP 67. The latest IP and TCP and Telnet are installed at ISID and ISIE TOPS20 Systems. In the last month or two several bugs have been identified and fixed. The next major step is to complete the TENEX version of TCP-4 and to eliminate any remaining instances of TCP-2.5. It is assumed that both the Packet Radio Project and UCL have completed their conversion to version 4, and that as soon as the SATNET project completes its conversion version 2.5 can be eliminated. The LDRSRV program was converted to use IP in a few hours. An FTP on TCP, based on ARPA FTP, is in debugging.

Ginny reported that several gateways are running, most are MOS based. It is planned to phase out the remaining ELF based gateways.

MOS Gateways

- at UCL between UCLNET and SATNET
- at NDRE between ARPANET and SATNET
- at BBN between ARPANET and SATNET
- at SRI between PRNET and ARPANET

ELF Gateways

- at Ft. Bragg between PRNET and ARPANET

It is planned to install a gateway at COMSAT between SATNET and the COMSAT local net. There is a new monitoring system. A configuration with a port expander and a gateway in the same machine is planned, it will be a tight fit.

Dale reported on the SATNET conversion to use IP. There are still

a number of programs to be converted, but it could be completed in about a month. Most of the remaining programs are used for maintenance or error recover. A weekend test was conducted using IP only. The major event impacting SATNET will be the removal of the NORWAY-LONDON ARPANET line at the end of September. This means all the LONDON ARPANET traffic will be routed via SATNET.

David noted that the scheduled GMCC demonstration will be a presentation, since there is nothing to demonstrate.

#### C. COMSAT - Mills

Dave reported that the ETAM and GOONHILLY PSPs are in the field and the TANUM PSP will be shipped soon. There have been some minor problems, for example, the software checksum and the hardware checksum do not agree. It is expected that the PSPs will be on the air next month.

Much of the COMSAT activity during the last few months has been directed to the NTC Demo in November and development of the Demo Terminal software. Currently, we have made arrangements for a suitable room and are now planning installation of lines, terminals and other gear. The demo will include speech, fax and the usual interactive and bulk transfers between ARPANET hosts and the Demo Terminal.

The Demo Terminal software has been checked out at UCL using an LSI-11, floppy disk, LPCM and Port Expander connected to UCLNET (net 13). TCP/Telnet was tested with ISIE via UCLNET, SATNET and ARPANET. FTP was checked out in loopback mode via the UCL Gateway. The LPCM was operated in expensive dictaphone mode only, since support is not yet complete for SATNET streams. This software will play with the Demo Terminal software at COMSAT and at the NTC Demo.

At COMSAT a Dacom 450 FAX machine has been delivered. This will be connected to the Demo Terminal and taken to the NTC site. Software to support this machine is now in frenzied development. When complete an update will be distributed to UCL. For local testing we are trying to arrange a direct connection to ARPANET (avoiding the SATNET for the time being) for debugging.

#### D. DCA - Cain

Ed reported that DCEC has a small (2-page) "C" program running in an 11/34 with two 1822 interfaces which implements a gateway.

## E. DOD - McFarland

Ray noted that mid October is the target date for the draft DOD Standards IP/TCP document. These documents will be rewritten to military document conventions.

## F. ISI STATUS - Postel

Jon noted that ISI is not implementing any IP, TCP or Gateways. ISI did attempt to write a program that used TCP on TOPS20. This program did not work (TOPS'20 crashed). Debugging is in progress with Bill Plummer and the ISI systems staff. ISI is working on an Internet Message Program. This is in an early testing phase and will be using TCP for inter module communication.

## G. LL - Forgie

Jim noted that the main Lincoln effort has been on the ST protocol to be discussed later and described in IEN 119.

Experiments with ARPANET-SATNET speech have been conducted with one source/destination in the ARPANET and another in the SATNET. There have been unexpectedly large delays in the SATNET streams, and in matching the ARPANET to the SATNET in the special purpose Gateway.

## H. MIT - Clark, Chiappa

Dave discussed the status of the LSCNET. There is a 5 node system up and running. The development of a version 2 interface is in progress. There has been difficulty getting an LSCNET/ARPANET gateway working due to low level hardware problems. A Trivial File Transfer has been implemented directly on IP. In Multics the IP layer will act as a gateway. MIT will be involved in the XEROX Grant program. Dave will be involved in integrating the Xerox equipment into the LCS environment. A PDP-11 based gateway between LSCNET and the ETHERNET is contemplated.

Noel commented that he saw a need for an effort to be made to convince groups to use IP rather than invent their own protocols.

## I. NDRE - Stensby

NDRE has been assisting in the preparations for the speech demo. The protocol status is nearly the same as last reported. This is due to VDH connection problems. Much work has been done on debugging the connection. Possible bugs have been very difficult to track down because we lack control of vital parts of the system. A couple of bugs have been found and corrected in the

NORD-10 VDH software without any substantial effect. The real problem seems to be that at certain times when the TIP has sent a packet out of sequence, it keeps retransmitting this packet instead of retransmitting the packet that is really missing.

J. RSRE - Davies

Brian reported that RSRE has programmed a gateway (based on IEN 30) and attempted to bring it up as a catenet gateway between UCLNET and RSRENET. RSRE will be bringing up a host with IP and TCP in about a month.

K. SRI - Kunzelman, Mathis

Ron reported that SRI is now operating two PRNETs in the San Francisco bay area, and one PRNET at Ft. Bragg, North Carolina. The net at Ft. Bragg is now eight terminals on two TIUs, and will grow to forty terminals.

Jim reported that the LDRSRV at SRI-KA which now uses TCP-2.5 will be changed to use version 4 soon, tested at BBNF, and installed at ISID and ISIE. Jim is also working on a Port Expander and Mini Gateway combined.

L. UCL - Kirstein, Jones, Treadwell, Cole, Bennett, Higginson

Peter gave a brief overview of the UCL effort, then let various members of the UCL group give reports on their activities.

Ron discussed the problems of including network software in a small Unix system. There was also some discussion of port expanders and X.25/IP interfaces.

Steve reported on the situation with the FAX experiment. The main problems seem to be with the timing constraints of the device, and matching these to the flow control constraints of TCP. This system currently uses TCP-2.5 but will convert to version 4 soon.

Bob reported on the development of a "C" to MOS linkage.

Chris discussed the prototype NIFTP service that allows FTPs from EPSS to Tenex and vice versa. A Unix version of NIFTP has been implemented.

Peter discussed the many flavors of X.25 one finds when building X.25/XXX interfaces.

M. IRIA & CELAR - Zimmerman

Hubert said that IRIA and CELAR would like to explore the possibility of cooperative research on the interconnection of networks. IRIA is working on local networks and interconnection with TRANSPAC and CYCLADES.

N. UNIVERSITY OF ROCHESTER - Lantz

Keith gave a brief overview of the environment at the University of Rochester and indicated the interest in becoming part of the internet system.

IV. ACTION ITEMS

A. DOCUMENTATION STATUS - Postel

Jon reviewed the IENs issued since the last meeting.

IEN	Author	Title
108	Plummer	Internet User Queues
		Describes the TOPS20 user interface to the IP.
109	Strazisar	How to Build a Gateway
		Describes the messages Gateways exchange with each other and with hosts, i.e., the Gateway-Gateway protocol. The main focus is on the computation of connectivity and routing information. Replaces IEN 30.
110	Cerf	Internet Addressing and Naming in a Tactical Environment
		Presents several addressing problems, particularly the "flying packet radio" problem.
111	Postel	Internet Protocol
112	Postel	Transmission Control Protocol
		The latest editions of the IP and TCP specifications. Not intended to be a change to the protocols but a better documentation of them. Replaces IENs 80 and 81.

- 114 Postel Protocol Options  
Summarizes the option available with IP and TCP. Replaces IEN 92.
- 115 Postel Address Mappings  
Shows how the addresses of various networks are carried in the Internet Address format. Replaces IEN 91.
- 116 Postel Internet Name Server  
Specifies the Internet Name Server, including some new features suggested by John Pickens. Replaces IEN 89.
- 117 Postel Assigned Numbers  
Lists the assigned number for various protocol fields, e.g. Networks, Sockets or Ports. Replaces IEN 93.
- 118 Postel Internet Protocol Handbook  
Table of Contents  
A one page list of the IENs that specify the current internet protocols. Replaces IEN 94.
- 119 Forgie ST - A Proposed Internet  
Stream Protocol

The description of a Internet level protocol for voice conferencing. Includes features for establishing multi-addressee connections.

The discussion that followed suggested several things that should be documented. For example, the higher level protocols (Telnet, FTP) should be in the internet protocol handbook, the procedure for determining parameters (e.g. retransmission timeout) should be documented, what to do in exception conditions should be described, testing procedures should be discussed (especially an acceptance test).

Concern was expressed that the new editions of the IP and TCP specification were not suitably reviewed by the implementors. It was suggested that a version with changes marked would be useful (see Appendix B).

## B. EXPECTED DOCUMENTS

1. IP Generic Services - Do the two documents 1) the IP Spec by Jon (IEN 111), and 2) the Gateway Spec by Ginny (IEN 109), supply the information needed?
2. SATNET IP Services - Dale McNeill is working with Dave Mills (currently the only user). When the current experiments are completed a document will be produced.
3. Transition Plan - Vint will discuss the issues in a small group meeting and appoint a task force to prepare the plan.
4. Host IP Services - Bill Plummer will prepare a document based on his talk on How to Build a Host IP.

C. XNET - It seems that XNET is not yet fixed to work with IP due to a technical disagreement between Jim Mathis and Ray Tomlinson. The solution to this is to be reported on at the next meeting.

D. Hack IP at SRI-KA and LDRSRV - This is working according to Jim Mathis.

## V. INTERNET TRAFFIC EXPERIENCE - Davies

RSRE has conducted some experiments sending messages looped back to themselves, with the loopback being located at various places, and using several input speeds and block sizes.

At RSRE there is a multi-terminal TIU with some modifications to perform measurements. The TIU measures the time between the TCP segment first being sent and the acknowledgment arriving, a kind of round-trip time.

The path is TIU via 2.4Kb line to PIXIE to PORT Expander to TIP into ARPANET. The input was either typing, a 1 octet segment traffic generator, or a 128 octet segment traffic generator. Histograms of performance under various inputs and loopback points were presented.

The results indicate lower than expected performance and the speculation was that the limiting factor was the program interrupt 1822 interface (either in PIXIE or the Port Expander).

It was noted that when two traffic generators were run sending segments to each other, ACKs were piggybacked 100% of the time both at the TCP level and at the X.25 level (recall the RSRE to PIXIE line operates the X.25 level 2 link protocol).



## VI. HOW TO BUILD A GATEWAY - Strazisar

Ginny discussed the way gateways determine how to do routing. The actual messages used are documented in IEN 109.

A gateway gathers information on network connectivity by polling its network interfaces and known neighbor gateways. The polling is done by sending a message, currently every 30 seconds. There is an algorithm to decide if a net or gateway is up or down based on the results of the polling. The routing algorithm is based on a distance measure, a directly connected network is zero distance, an unreachable network is infinite distance. The gateways exchange with neighbors their distance tables (substituting infinity in entries where the distance to the destination is greater than the distance from the neighbor). Routing information is sent only when something changes. A new gateway can join the system as long as old gateways can add a row to their tables dynamically.

Note that "smart gateways" do all this routing and connectivity work, "dumb gateways" do not. Smart gateways try to route things to other smart gateways, but if that can't be done they may fall back to using compiled in fixed routes to dumb gateways.

Danny Cohen suggested that there may be problems with this algorithm as the number of networks gets large.

## VII. HOW TO BUILD A HOST IP - Plummer

Bill presented the general flow of processing in a Host IP. In general datagrams arrive from the connected network interfaces. Any fragmented datagrams are reassembled. Complete datagrams are entered into a queue. Datagrams addressed to other hosts may be forwarded if this host is acting as a gateway, otherwise they are discarded. Datagrams addressed to this host are delivered to waiting protocol modules or user processes. User processes or protocol modules may create datagrams to be sent, these are queued and based on the destination address sent via one of the connected network interfaces.

During the processing of arriving datagrams several checks must be made:

1. check the IP version number
2. verify the checksum
3. see that the length field and the amount received are consistent (amount received may be slightly greater due to padding)
4. check time to live
5. assemble fragments

## Internet Meeting Notes

When sending a datagram the IP fills in several fields and makes a routing decision:

1. Insert source address
2. Insert any IP options called for
3. Insert time to live
4. Insert checksum
5. Make a routing decision, select an output interface, generate local net header and interpret type of service.
6. Queue for transmission

The routing decision can be based on a table of network - gateway correspondences. For each network a gateway on a directly connected network is listed. If the host is notified to use a different gateway for that destination network the table can be updated. If a gateway breaks and the host can tell (e.g., receives a "destination dead" response) then switch to an alternate gateway. It is suggested that a heuristic for switching destination gateway is to do it when the higher level protocol (e.g., TCP) retransmits a segment. (However it is not clear the IP can tell a new transmission from a retransmission).

In the area of congestion control Bill suggests that it is the rate that must be controlled which may be considered to determine an internet datagram interval. If a gateway tells a host to slow down, increase the internet datagram interval to, say, 130% of its former value. To avoid being stuck at a too large value for the internet datagram interval, reduce it to 95% of its former value every 30 seconds or so.

There is an issue of fairness in any congestion control. How can the IP properly allocate the available transmission capacity among the higher level protocol modules and users?

The discussion suggested the need for a "Trace Datagram" that causes each gateway and the destination host to send back a note, a kind of "multi-echo."

IP modules should provide a handle on the user side to explicitly cause a new routing to be calculated.

The congestion control parameters may be very critical and sensitive. Dave Clark has done some simulation experiments and will prepare an IEN on this topic. There may be some interaction between source-destination host pairs in this area.

## VIII. STREAMS AND CONFERENCING - Forgie

Jim presented the proposed ST protocol which is documented in IEN 119. The goals of ST are: (1) to provide a good base for point-to-point and conference packet speech, (2) to support research on effective traffic control techniques, (3) to support a demonstration of cost-effective speech, and (4) to operate as an extension of internet protocol.

Jim indicated four requirements and corresponding approaches for meeting them: (1) guaranteed data rate - know requirement in advance and assign loads to links statistically, (2) controlled delay (predictable dispersion) - prevent congestion by controlling access on a call basis, (3) small quantity of speech per packet - use abbreviated header and aggregate small packets for efficiency, and (4) efficiency equal to or better than circuit switching without TASI (packet efficiency \* link utilization > 45%) - abbreviated headers for packet efficiency and high link utilization with effective traffic control.

One of the key decisions is to consider ST connections to be full duplex. The claim is made that this allows faster connection setup for conferences. There was some discussion on this point and the issue will be examined further.

Another issue is the sharing of transmission capacity between speech (ST) and data (IP) in the internet.

Also discussed was the identification of the route by a list of gateways rather than by a list of networks. The resources to be reserved are controlled by networks not gateways.

One very interesting feature of ST is its mechanism for flexible multi-addressing and routing of messages such that each addressee receives at most one copy of a message.

Jim presented the message formats and worked through an example of connection setup. There was some concern expressed about the effects of delayed duplicate control messages, and how ST operates under failure conditions. Also concern that several conference setups in competition may result in a deadlock similar to reassembly lock up.

## IX. SOURCE ROUTING &amp; PROTOCOL MULTIPLEXING - Cohen

Danny reviewed the proposed source routing procedure described in IEN 95. This mechanism provides a list of addresses for the datagram to pass through. Normal routing is used to move the datagram between these points. The route information for sending a datagram from A through B and C to D is expressed by A as TO: B, SR: C,D. At B this is changed to TO: C, SR: D, and at C it is changed to TO: D. The key is that the next address must be interpretable at the point it is used. This means that local (rather than internet) addresses may be used.

The discussion focused around the issue of non internet addresses in the source route and how (if at all) to mark them. It seems desirable to let a source route contain any of the following in any order:

1. Internet Address.
2. Local Address (prefixed with an escape code and a length).
3. Network Address (Network part only).
4. Here - an address value that means this host, this net (or any host, or every host).

Alternately one could have each address in the source route have a prefixed "op-code" that indicates what kind of address it is.

Danny did not discuss the multiplexing of protocols (due to time pressures) but urged everyone to review IEN 90. It appeared that the multiplexing protocol was accepted (at least in principle).

## X. GATEWAY MONITORING &amp; CONTROL CENTER - Flood Page

David described the GMCC organization and the types of information that will be reported.

The operation is much like the ARPANET NCC except that gateways are the objects of interest rather than IMPs.

Gateways will communicate with a set of monitoring and control processes (running on a TOPS20 somewhere). These monitoring and control processes will update a central data base. The data in this data base may be accessed by Terminal Processes (i.e., User Interface Programs).

A terminal process may take control of a gateway, but a gateway may only be controlled by one terminal process at a time.

Three types of controls may be exercised: start or stop reports, inquiries, enable or disable traps. The reports give throughput

statistics and interface up/down status. The inquiries give the gateway description (name, connectivity), echo response, or routing data. The traps give changes in the interface up/down status, or neighbor gateway up/down status.

The information stored for each gateway in the central data base includes: the latest regular report, the interface address, the neighbor gateways, the reporting and trap status, the gateways own up/down status, and the process currently controlling the gateway (if any).

In the future the GMCC will provide summary reports on traffic and up/down status.

Other future topics are Performance Measures, Higher Level Fault Diagnosis, Accounting Information, and inclusion of information from NCCs.

To use the GMCC facilities one would login to the GMCC host and run a terminal process. There was some discussion of the access control on people taking control of gateways. Also some concern about the artifact introduced into the measurement by accessing the GMCC via the gateway being measured, and also the effect on other gateways (or their statistics) of relaying the measurement messages to the GMCC host.

David will prepare reports on: (1) GMCC Message Formats, (2) How to use a Terminal Process.

#### XI. DISCUSSION OF SPEECH DEMO - Forgie

Jim described the configuration for the speech demo. There are four sites involved: Lincoln Laboratory and ISI on the ARPANET and NDRE and UCL on the SATNET. There is a special purpose gateway at BBN. This gateway is a PDP 11/34 ELF system. The basic data rate for one site is 5 packets/sec, at this rate the gateway must handle 15 packets/sec.

In the ARPANET section of the conference the mode of operation is centralized control and distributed data. In the SATNET section the mode is distributed control and distributed data. The SATNET section uses a SATNET shared stream.

#### XII. DISCUSSION OF FAX SYSTEM - Treadwell

Steve described the configuration and operation of the FAX-Network system. The FAX machine is a DACOM 6000. It is connected to an LSI-11 (MOS operating system). The LSI-11 contains an interface to

the DACOM, a TCP, and a controlling program. The controlling program interacts with a user at a display terminal.

There is also a server FAX program on BBNC that accepts input and stores it in the file system, or on request sends stored FAX data from the file system.

The operation is for the user to type on the controlling programs terminal a request to send FAX data to a file. A TCP connection is established between the LSI-11 and FAX server at BBNC.

Pages are input through the DACOM and data is sent to BBNC and stored in the file. To move FAX data from the file to paper, the user enters a retrieve request on the controlling programs terminal.

This system currently uses TCP-2.5. The next step will be to convert to using version 4. Also to change the address so the transmission will be through an internet gateway.

There have been several problems in getting this working. Primarily the timing requirements of the DACOM and the flow control of TCP. A typical page results in 300,000 bits (with wide variance). The DACOM sends 585 bit blocks and if not accepted within 6 seconds aborts the page. The total transmission path must support at least 100 bits/sec on the long term average.

ISI is getting a RAPICOM 450 FAX machine and will use it in a similar way. (A RAPICOM 450 and a DACOM 6000 are identical machines). At ISI the FAX machine will be treated as a device (a peculiar terminal on TOPS20) rather than as a host.

### XIII. ADDRESSING ISSUES - Cerf

Vint described several situations where a host may have several possible addresses and may wish to change between them dynamically without losing ongoing communications.

A host may be dual (or multiply) homed, i.e., have two (or more) interfaces to the same net. A host may be connected to more than one net. Or a host may change nets.

This last case is illustrated by a flying packet radio, passing over a series of ground PRNETs each connected to the internet via a gateway. The flying packet radio is for a time in the range of PRNET-1, then passes out of that range into the range of PRNET-2, etc.

Another problem is partitioning. That is a network breaks into two (or more) parts which continue to function independently. Can a host

in one part communicate with a host in another part via gateways and other networks?

Does source routing solve the partitioning problem? Or is a scheme needed to dynamically assign unique addresses to each partition so that the ordinary gateway dynamic routing will solve the problem?

Another proposal is to drop the requirement that connection be identified by the total address. By using only the pair of ports to identify a connection one could accept as a continuation of a conversation messages from a different host if it had the same pair of ports (and the expected sequence number).

#### XIV. PARTITIONING - Plummer

Bill discussed the problems of multiply connected hosts. If a host is connected to two networks, which are also interconnected via gateways, the host has alternate routes to use in sending messages. The internet gateways can exchange connectivity information. A multiply connected host cannot take full advantage of this without being a gateway itself and calling the host a network. Another problem is for a destination to answer a message from a multi-connected host. The multi-connected host would like to allow any of its interfaces to be used to receive the reply. To do this it could use a pseudo network number and act as if those interfaces were gateway connections.

This use of network numbers for pseudo nets may use up the network numbers too quickly. Bill proposes what he calls host-group routing which uses a 32 bit address as if it were a network number. A host chooses one of its internet addresses as its "name". Gateways treat that name as a network name for connectivity and routing purposes. A mask can be used to allow groups of hosts to be routed based on one entry. Bill gave an example of how this might work and showed a routing table for a pseudo net at BBNC.

This topic is to be discussed again at the next meeting.

#### XV. RIG - Lantz

Keith reported on the network environment at the University of Rochester. The key is a multi-machine multi-net system started in 1974. All communication between processes in this system is via messages. The system is based on Feldman's experience and other papers on message communication.

The equipment involved includes: 4 Altos, 2 Eclipses, an Ethernet, 9 terminals, 1 PDP-10, 1 VAX, and an ARPANET. The applications include: editing, compiling, file management. An important

development is the virtual terminal model, which allows several processes to interact with a terminal using multiple possibly overlapping windows.

The process level communication involves a name to address conversion by lookup in a name server, and an address to route conversion supported by a local kernel and network communication modules. Three communication styles are supported: atomic transactions, connections or streams, and asynchronous messages.

One problem area is protocols. There are too many: U of R Ethernet, PARC Ethernet, ARPANET, DECNET, internet. Can all these live in harmony?

#### XVI. SUMMARY OF SMALL GROUP MEETINGS

##### 1. PERFORMANCE MEASUREMENT SMALL GROUP MEETING - Kunzelman

This group discussed the need for performance measurements of the internet protocols at all levels. The group reviewed the current work, discussed performance metrics, measurement tools, and documentation of measurement methods and results.

The group recommends that specifications of performance measures for IP and TCP be written, that a set of tests for IP, TCP, and higher level protocols be defined, and that a bibliography on performance measurement be prepared.

Please see Appendix C for further detail on this group meeting.

##### 2. UNIX SMALL GROUP MEETING - Cain

Three topics were addressed in the UNIX small group meeting:

1. Use of network UNIX on small PDP-11's (specifically, the 11/34 and the 11/40).
2. MOS work on UNIX systems.
3. Use of UNIX on VAX machines.

For PDP-11's which do not support the separation of instruction and data space, incorporating the widely distributed NCP programs into the UNIX kernel usually produces an object which is barely small enough to boot, even after greatly reducing various system parameters. Processes like TCP and Telnet have to share the scarce remaining space, and are bound to perform poorly because of swapping. Noel Chiappa is about to obtain, from NOSC, a version of UNIX which has no disk buffers in the kernel data space, and will attempt to install the NCP kernel software. Performance of this



version of UNIX is uncertain, but a mailing list (see Appendix A) has been established for those interested in the outcome.

UCL, MIT, and BBN are all interested in doing MOS work in a UNIX environment. The problem is that there is a variety of file types (.a.out, .obj, .rel, .lda) which must somehow be combined into an object which is loadable into an LSI-11, and is understandable by DDT or a similar debugger. UCL has "C" interfaces to MOS system calls and some utilities which Noel Chiappa will attempt to put in a "nice" UNIX environment, and make them available for others. A mailing list (see Appendix A) was established for those who wish to follow his progress.

Keith Lantz discussed a meeting with ARPA, Rochester, CMU, and others associated with ARPA-funded research involving VAX machines. At the meeting, the decision was made to use UNIX on the VAX machines, specifically the version of UNIX modified at Berkeley to take advantage of the VAX paging hardware. The source of support for the VAX UNIX is not yet named. Rochester and CMU are both interested in implementing TCP and IP on their VAX machines. A mailing list (see Appendix A) was established for those interested.

### 3. TCP VERSION 4 AND TIU SMALL GROUP MEETING - Haverty

The small group meeting on the status of the new TCP implementation and TIU also covered some general issues of TCP support for MOS-based systems in general.

Jim Mathis reported that recent work on the MOS TCP version 4 implementation involved separating the code into individual modules, to isolate the network- and operating-system-dependent sections of the code. The major part of the TCP code is now in the TV4PRO module. Operating system support code exists in TV4MOS or TV4ELF, depending on the operating system involved. The "released" versions of this implementation are available on the SRI-KL machine in the MOS-DEVEL directory, but they will be moved soon to PKT-LSI-11.

Ongoing work with the TCP involves addition of an internet layer which will be accessible to user processes using entry points to be defined. This will occur in a later release of the software, which is not expected to involve any changes to the TCP interface to user processes. The internet layer will also implement some of the gateway-gateway protocol functions.

In addition, the software can be extended to support multiple hardware interfaces, to permit multi-homing of hosts on multiple networks. This work however is not yet in progress.

The discussions which followed identified a need for two new mechanisms within the user community. The first is simply an announcement mechanism, to inform users of the SRI software of new releases, the changes which have been made, functions added, and location of the code itself. A second mechanism involves integration of changes developed at user sites into the SRI-maintained sources. For example, software to support other network types must be done at the site running the network, but should then be integrated into the standard sources.

Because of the need for sites other than SRI to implement network-interface software, the internal interface between the TCP and internet layer and the local-net module must be clearly defined and stabilized. MIT is likely to be the first site to implement a new local-net module.

Further discussion explored details of the functionality of the existing TCP implementation, and requirements which various user projects expect to have in the future. Some of the facilities not yet implemented include:

- o ability to have multiple connections in one MOS process
- o reassembly of internet fragments
- o full processing of EOL
- o processing of "rubber EOL"
- o user access to URGENT functions
- o full functionality at the Telnet user interface

The last area explored identified the need for additional work on the Telnet implementation, to, for example, provide the ability to send remote RESETs.

A short discussion of "coming attractions" followed, including the following expected changes:

- o internet layer in Bliss, probably by the end of 1979
- o use of the internet name server by Telnet
- o XNET server within MOS, to debug MOS processes remotely

The TCP software is currently written in MACRO-11. It is a candidate for conversion to Bliss, but this will not be done until the Port-expander software is converted.

One possible performance limitation was identified during the discussion. Currently, data is ACKed when it is accepted by a user process. Since this can be significantly later than the actual receipt of the data from the network, throughput might be improved by modifying the TCP to ACK data on receipt from the network.

A short discussion of TIU issues identified the need for determining how many terminals can be supported by a TIU at once, in two cases. In a heavy-use situation, the limitation on number of active, high-throughput terminals should be determined. This number is probably limited by the raw speed of the LSI-11, and may be increased by conversion to faster hardware such as the PDP-11/23 when it is available. A second case involves the limitation on support of low-usage terminals. This limit involves primarily how many serial interfaces can be connected to the LSI-11 and configured in the packaging.

The meeting also demonstrated a need for better interchange of information among the participants and general community of users of MOS, TIUs, and related hardware and software. A mailing list is being set up, called MOS-TCP, for this kind of activity (see Appendix A).

#### 4. LONDON-NORSAR LINE AND SATNET SMALL GROUP MEETING - McNeill

Dale said that loading of the Goonhilly SIMP and the London TIP over the satellite channel is operational. Monitoring and control paths exist for all SIMPs; sufficient backup paths exist through use of the satellite channel and gateways connected to other SIMPs.

Vint announced that beginning January 1, 1980, the NORSAR-SDAC circuit will become a military circuit for operational data only. Seismic data traffic from the NORSAR TIP will continue to be sent on this line.

Funding for the London-NORSAR circuit ceases October 1, 1979. Peter anticipates that the ARPANET direct connection circuit via SATNET (line 77 between London and SDAC) will be needed to provide ARPANET connectivity to the London TIP until some time after October 1980. This circuit will continue to provide ARPANET service to the Rutherford Lab, EPSS, and the London TIP. The long-term goal is to have internet service to London only using the UCL gateway attached to SATNET. Prerequisite is a change of the London TIP to a host (TCP TIP). The mechanism of internet loading and support of London machines must be developed.

In conflict with the maintenance of operational service to London is the introduction and checkout of PSP Terminals, the checkout and release of new SIMP software, and the checkout and performance of the NTC demonstration. Hardware and software development will be restricted to after 1400 EDT, while operational service will be maintained from 0200 to 1400 EDT, daily.

## Internet Meeting Notes

## 5. NTC DEMO PLANNING SMALL GROUP MEETING - Mills

The principal results from that meeting were a revised plan for network connectivity and a change in the way speech would be handled. Specifically, it was decided that Jim Forgie's software would be used in the BBN and UCL gateways with the LPCM operated from Washington via appropriate data lines to the BBN Gateway. Jim has already made provision for this in the LPCM design, but it has not been operated in this mode before.

Datagram traffic for the Demo is planned to be sent via Clarksburg, with a backup to Etam via a backdoor connection via the Clarksburg SIMP internal gateway and the RCC or other suitable hack. UCL will participate in the Demo from London with live speech and fax. Participation by NDRE will depend on a PSP Terminal being installed there.

Jim mentioned a couple of minor problems with the LPCM hardware. Correcting these will probably involve sending ROMs to the various sites. There is a question of Gateway reliability which leaves all of us nervous

## 6. IP/TCP FLOW CONTROL SMALL GROUP MEETING - Clark

This meeting was a fairly informal discussion among the implementors of TCP and gateways. Several topics were discussed.

- 1) How to manage RFNMs on internet link.
- 2) A simple simulation of a host responding to quench messages.
- 3) How to perform quenching experiments, given the current level of internet operation.
- 4) Should the TCP-IP or application-TCP interface be specified/standardized with respect to flow control?

While the discussion was fruitful, no substantial conclusions were drawn.

## 7. X.25 GATEWAY SMALL GROUP MEETING - Binder

The purpose of this meeting was to review various factors impacting the development of an ARPANET/X.25 gateway using an LSI-11 to provide a 50 Kbps full duplex service. The meeting basically covered only the issue of hardware interfacing to the X.25 network. Three different interfaces were discussed: a byte-interrupt device, a non-DMA packet-interrupt device, and a full DMA device.

The byte-interrupt interface and related level 2 and 3 X.25 software has been developed by UCL and is now available for use. Although precise throughput measurements have not yet been made,

experience with the LSI-11/MOS system by meeting attendees indicated that a byte-interrupt configuration would not achieve 50 Kbps full duplex operation. In particular, measurements by SRI have indicated a maximum of about 50 Kbps through a looped-back 1822 byte-interrupt interface on a MOS system running only a message generator (no gateway or second network interface).

The packet-interrupt interface is being developed by RSRE, and along with level 2 software is expected to be checked out in 6-8 weeks. This interface contains 4K bytes of buffering, shared among input and output. Operation consists of copying packets between the interface memory and main memory, with a single interrupt associated with each packet.

UCL plans to develop a full DMA interface, but this is not expected to be available until anywhere from 8 to 12 months from now.

A framing issue presently exists in that both IPSS and Telenet X.25 networks now use bi-sync framing. However, both are expected to make HDLC bit-oriented framing available this fall.

#### 8. PORT EXPANDER SMALL GROUP MEETING NOTES - Davies

##### 1. Hardware Delivery Outlook.

SRI has fifteen systems on its order book and enough components to build only two complete ones at this time. Some of the hardware bottlenecks are:

- a) Boxes from DEC - deliveries up to January '80.
- b) Low power Schottky chips for 1822 DMA boards.
- c) Cable connectors for 1822 units - Amphenol.
- d) Control Panels.

Delivery of port expander units would continue at a slow rate at least until the end of the year with Vint Cerf nominating the recipient of each system as it rolls off the production line.

A postscript on the hardware discussion concerned the robustness card which may have to be reprogrammed for automatic loading from nets which are not on the ARPANET. The robustness card uses UV erasable PROMS. Mathis said that there would be documentation supplied with the robustness card indicating how this should be done. The cards themselves were just coming back from the printed circuit facility and although one or two types of chip for this board were in short supply, it was hoped that board delivery would begin at the end of September.

## 2. Software Developments.

Most of the discussion centered around the new facilities which would be offered with the SRI supported BLISS-11 version of the PORT Expander code. Some of these features are:

- a) Provision of NOP handling.
- b) Provision of IMP padding.
- c) Flow control facility.
- d) Facility to simulate RFNMs internally.

There were still some open ended problems like what should a port expander do when the NCP host crashes. One possible approach is to offer two different solutions to this problem and select one at configuration time.

Possible release dates for the BLISS-11 version of the port expander were mid October if Mathis can get an IMP port for debugging or up to a month after depending on the priority of the work.

Finally Strazisar joined the meeting for a discussion on integrating the mini-gateway code and port expander code into the same machine. Clark suggested that this integration problem would become almost trivial if someone wrote a null 1822 driver under MOS which allowed buffers to be passed across the driver without copying. Strazisar volunteered to take a debugged copy of Mathis' port expander code and integrate it into the mini-gateway with an optimistic delivery date of the End of November.

## 9. TCP IMPLEMENTATION SMALL GROUP MEETING - Cerf

The discussion focused on the changes to IP required by DCA to satisfy the DOD security and precedence needs. The existing IP option for S/P/T does most of what's needed, but in addition the TOS field will be redefined slightly to be:

```

0 1 2 3 4 5 6 7
+-----+
|      |S|  |S|S|
| PRC  |/|Rel|/|P|
|      |D|  |R|D|
+-----+
```

That is:

- Bits 0-2: Precedence
- Bit 3: Stream or Datagram
- Bits 4-5: Reliability
- Bit 6: Speed or Reliability

## Bit 7: Speed

This information has to be passed up and down the layers of protocol. Enforcement of the security and precedence is up to each host. There needs to be a specification of how to interpret and implement these functions. Preemption rules must be specified too, for the higher level protocols.

There was some discussion of the unauthorized use of security and precedence features. The basic rule is to act on it speedily now and to log the header of the message so one can ask later if the originator was authorized.

## 10. TCP APPLICATIONS SMALL GROUP MEETING - Postel

This group discussed problems in programming application programs that use TCP. Postel had a problem figuring out how to use the user/TCP interface (JSYSS) to create a successful TCP using program. The problem may be with the documentation. Haverty has the problem that programs don't work but don't crash either. Clark said his TCP can't tell the user anything more specific than "it didn't work."

What is needed is better documentation, fault isolation tools, and user level debugging aids. Also a memo on what is available, e.g., a traffic generator at BBN-UNIX.

## 11. ST PROTOCOL SMALL GROUP MEETING - Hoversten

This group first made up a detailed agenda of issues to discuss. In general the topics included:

- Connectivity
- Routing
- Stability
- Protocol Layering
- Resource Allocation
- Information Exchange

Further discussion of the role of ST in resource allocation is needed. It is agreed that ST provides a structure to support multi-addressing and a resource allocation policy.

The relationship of ST to IP is that ST is experimental, will be implemented at a small number of sites, and will have IP imbedded in it.

12. 1822 BOARD PROBLEM SMALL GROUP MEETING - Mathis

The problem with 1822 boards is being worked on at SRI. Software for the TIU which circumvents the problem will be made available. A hardware fix is being studied.

There was some discussion of the unauthorized use of security and precedence features. The basic rule is to act on it as quickly as possible and to log the header of the message so one can ask later if the originator was authorized.

There was some discussion of the unauthorized use of security and precedence features. The basic rule is to act on it as quickly as possible and to log the header of the message so one can ask later if the originator was authorized.

10. TCP APPLICATIONS SMALL GROUP MEETING - Postel

This group discussed problems in programming application programs that use TCP. Postel had a problem figuring out how to use the user/TCP interface (LIBS) to create a successful TCP using program. The problem was with the documentation. Haverly has the problem that programs don't work but don't crash either. Clark said his TCP can't tell the user anything more specific than "it didn't work".

What is needed is better documentation, fault isolation tools, and user level debugging aids. Also a memo on what is available, e.g., a traffic generator at BBN-UNIX.

11. ST PROTOCOL SMALL GROUP MEETING - Haverly

This group first made up a detailed agenda of issues to discuss. In general the topics included:

- Connectivity
- Routing
- Stability
- Protocol Layering
- Resource Allocation
- Information Exchange

Further discussion of the role of ST in resource allocation is needed. It is agreed that ST provides a structure to support multi-addressing and a resource allocation policy.

The relationship of ST to IP is that ST is experimental, will be implemented at a small number of sites, and will have IP imbedded in it.



## XVII. AGENDA FOR NEXT TIME

The next meeting is February 4-6 at SRI International in Menlo Park. Ron Kunzelman is the contact.

## Agenda Items:

1. Kirstein, Davies, Frankel - Internet Performance Experience
2. Flood Page - Longterm Gateway Statistics
3. Cerf - Schedule & Milestones Review
4. Forgie - ST Protocol
5. Cohen - ST Performance
6. McQuillan - Internet Routing
7. Cain, McFarland, Cerf - IP/TCP Acceptance Testing
8. Plummer, Tomlinson - Higher Level Protocols FTP & Integration with TCP
9. Postel - Internet Message Handling Interim FTP Based Multi Media Supporting
10. Deutsch - BBN - FAX & Text Message System
11. Cerf - Multi Homing & Partitioning
12. Plummer - Host Group Routing
13. Demos:
  - A. GMCC Monitoring - Flood Page
  - B. Fault Isolation in TIU - Mathis
  - C. Internet Message Transport - Postel

## Action Items:

1. Milestones - each contractor is to provide a list of milestones to Cerf. The intention is to quantize tasks and identify increments in capability. The goal is to make progress easily detectable. The combined milestones will be distributed to the internet group, and will be reviewed at the next meeting.

## Internet Meeting Notes

2. Monthly Reports - each contractor is to provide a monthly report to Postel. The report should note accomplishments, progress on milestones, unusual events, problems. A summary of all reports will be distributed to the internet group.
3. XNET - to be converted to version 4 by Jim Mathis and Ray Tomlinson.
4. Host IP Document - Bill Plummer is to write an IEN on how to build a host IP.
5. GMCC - David Flood Page will write an IEN on GMCC message formats and an IEN on how to use a terminal process .
6. Congestion Control - Dave Clark will write an IEN on simulation experiments in IP congestion control.

## Future Meetings Schedule

1980

FEB - SRI  
 MAY - MIT  
 SEP - RSRE

1981

JAN - ISI  
 MAY - ARPA  
 SEP -UCL

## XVIII. DOCUMENTS DISTRIBUTED

IEN	Author	Title
---	-----	-----
109	Strazisar	How to Build a Gateway
110	Cerf	Internet Addressing and Naming in a Tactical Environment
119	Forgie	ST - A Proposed Internet Stream Protocol

## XIX. ATTENDEES

Vint Cerf	ARPA	Cerf@ISIA
Robert E. Kahn	ARPA	KAHN@ISIA
Dick Binder	BBN	BINDER@BBNE
Jack Haverty	BBN	JHAVERTY@BBND
Dale McNeill	BBN	DMCNEILL@BBNE
David Flood Page	BBN	DFLOODPAGE@BBNE
William Plummer	BBN	Plummer@BBNA
Virginia Strazisar	BBN	STRAZISAR@BBNA
Hoi Y. Chong	COMSAT	Mills@ISIE
David Mills	COMSAT	Mills@ISIE
Chip Bumgardner	CTEC	CHIPS@BBNC
Jack Hammett	DARPA-IPT	Hammett@ISIA
Ed Cain	DCA	Cain@EDN-UNIX
Ray McFarland	DoD	McFARLAND@ISIA
Michel L. Audren	French Ministry	Observer
Dominique A. Truchetet	of Defense	Observer
Hubert Zimmerman	IRIA	HZim@BBNC
Danny Cohen	ISI	Cohen@ISIB
Jon Postel	ISI	Postel@ISIB
Estil Hoversten	Linkabit	Hoversten@ISIA
Noel Chiappa	MIT	JNC@MIT-AI
Dave Clark	MIT	Clark@MIT-MULTICS
Jim Forgie	Lincoln Lab	FORGIE@BBN
Frank Deckelman	NAVELEX	DECKELMAN@ISIA
Aage Stensby	NDRE	AAGE@SRI-KA
Andrew Bates	RSRE	RSRE-T4@ISIA
Brian Davies	RSRE	RSRE-T4@ISIA
Allan J. Fox	RSRE	RSRE-T4@ISIA
John Laws	RSRE	RSRE-T4@ISIA
Ron Kunzelman	SRI	KUNZELMAN@SRI-KL
Jim Mathis	SRI	Mathis@SRI-KL
Robert Cole	UCL	UKSAT@ISIE
Sunil K. Das	UCL	PKT-UCL@SRI-KL
Peter L. Higginson	UCL	UKSAT@ISIE
Andrew Hinchley	UCL	UKSAT@ISIE
Ron Jones	UCL	UKSAT@ISIE
Peter Kirstein	UCL	Kirstein@ISIA
Alan J. Mayne	UCL	UKSAT@ISIE
Steve Treadwell	UCL	UKSAT@ISIE
David Lloyd	UCL/UKMOD	LLOYD@ISIA
Keith A. Lantz	U of R	LANTZ@SRI-KL

APPENDIX A

The following special interest group mailing lists have been set up by Noel Chiappa at MIT-ML. To use them, send to <group>-PEOPLE at MIT-ML [or just <group> (DON'T include the "<"'s)].

If you want to be on one or more of these mailing lists please send a note to Noel, JNC@MIT-AI (not to the whole group).

If you get a message from Person and to a mailing list, DO NOT send your reply to the mailing list AND Person; the MIT-ML mailer isn't smart enough to notice and will send duplicate copies to Person.

- MOS-UNIX           MOS support on UNIX
- SMALL-UNIX        UNIX on small machines
- VAX-UNIX           UNIX on VAX
- MOS-TCP           MOS, TIU and the MACRO-11 TCP
- MOS-PE            Port expander / Minigateway
- NET-UNIX          Network support on any UNIX

## APPENDIX B

## IP and TCP Specification Differences

The following is a brief description of the difference between IEN-80 and IEN-111 (IP), and between IEN-81 and IEN-112 (TCP).

In both cases the new documents are intended to simply be better documents, and no significant changes to the protocols are intended. There are many minor changes of wording, punctuation, or spacing, so that a source compare would catch many many paragraphs in which there is no significant change. It is intended that the text added (or in some cases deleted) lead to an easier to understand description of the protocols.

## IP Differences

1. A much more specific description of a fragmentation and reassembly procedure is included.
2. Some Options are changed or added:
  - a. Three options are added for use by the BCR protocol.
  - b. A Stream-ID option is added.
  - c. The Source Route option is changed to conform with IEN-95.
  - d. The Return Route option is added to conform with IEN-95.
  - e. The Error Report option is expanded to have several specific error codes., and a standardized IP header for error reporting.

## TCP Differences

1. Two additional States are introduced in the connection closing sequence.
2. The EOL sequence number fixup procedure is changed to be based on remembering the sequence number of the beginning of the most recent buffer rather than the initial sequence number of the connection.
3. In many cases the RST is not required to acknowledge that segment that caused the reset to be generated, and in many cases it is not necessary to check the ACK value to verify a RST.

## APPENDIX C

## Minutes of Internet Subgroup Meeting for Performance Measurement

Present: Ron Kunzelman (SRI) (Chairman), Noel Chiappa (MIT), Brian Davies (RSRE), Stephen Edge (UCL), David Flood Page (BBN), Andrew Hinchley (UCL), Ron Jones (UCL), Bob Kahn (ARPA), Peter Kirstein (UCL), Alan Mayne (UCL) (Minute-taker), Ray McFarland (DOD)

Kunzelman pointed out that there is need for performance measurement at different levels of protocol, including measurements on Telnet and FTP at the user level, TCP at the transport level, and datagrams at the internet level. The meeting then considered various specific aspects of measurement.

## Measurement Work Now Being Done or Being Planned

BBN (Flood Page) will do some performance measurements at the datagram level on traffic passing through gateways; no tests are envisaged. The BBN gateway will monitor IP traffic going through the gateway. BBN is looking at CPU utilization, and can also extract interface information by address pairs. BBN (Wingfield) is doing some performance tests on TCP. Throughput and delays have been investigated.

SRI is measuring Telnet throughput (bits/sec) from TOPS20 and TENEX hosts.

UCL is trying to time FAX transmissions, and NIFTP transmissions at the user level. UCL would like to measure individual performances end-to-end, to see if there is any destructive interference. For ARPANET-SATNET-ARPANET transmissions, there is a need to know the best ARPANET performance, the best gateway performance, and the best Satnet performance, and then compare the combination of these with the best actual ARPANET-SATNET-ARPANET performance, to see whether they are comparable or differ by an order of magnitude.

RSRE (Davies) has previously measured with old PIXIE, and would now like to repeat some of these measurements with new PIXIE. Measures have been made of round-trip TCP ack times. The measurement techniques consist mainly of turning the traffic generator on for about five minutes at a time, then obtaining a delay histogram.

### Performance Metrics

The following metrics have been used to measure ARPANET performance:

1. ARPANET has problems of low bandwidth, hence uses throughput metrics such as packets/sec and user-bits/sec.
2. SATNET has long delays, hence measures mean transmission time.
3. PRNET is liable to high loss, hence measures the percentage of missing packets.
4. Packet efficiency is the ratio of information packets received to total packets received (the difference being due to control and retransmission overhead).

Kirstein stressed the need for better individual metrics, and then for developing algorithms for combining them. For example, both bits/sec and packets/sec are important.

He also raised the question of how the performance varies when going through different combinations of networks.

### Measurement Tools

Tools currently being used in Arpanet include:

1. At the user level, FTP + "stopwatch" for bandwidth and Telnet + "stopwatch" for measurements of remote echo time and bandwidth.
2. Timestamps in pickup packets.
3. Echo server in gateways.
4. Traffic generators and sinks.

There is also a possibility of using synchronized clocks to measure one-way times and asymmetric delays as well as round trip times. Hinchley is looking at interval times between successive packet streams, as a method of measuring delays.

### Measurement Needs

Kunzelman pointed out that there is currently no measurement at the transport level in the ARPANET. Statistics are needed for the distribution of user traffic, which can be measured on entry to

the internet system. There is also a need to have statistics for the distribution of packet sizes.

More measurement is needed of round-trip TCP ack times: so far, this has been done only by RSRE. Also in TCP, measurements are needed of retransmissions and checksum errors. TCP-specific measures should be separated from general network and gateway measures, although such separation might be difficult to achieve in practice.

Mayne would be grateful to anyone who would send him empirical data that would be useful to help validate his proposed simulation and theoretical studies, especially on problems of joint ARPANET-SATNET operation. It was pointed out that McQuillan has a large amount of empirical data.

#### Some Problem Areas

Kahn pointed out the need to isolate problem areas, such as optimization and tuning, for which measurements are needed. Some of these problems relate to research and development about networks, including obtaining insights about their behavior, others are concerned with the operational control of networks, and there is also the problem of persuading some network users that they actually have problems!

Kirstein said that a big problem is to put uniform methods of measurement into the internet environment; this is very difficult and requires much thought.

Kahn asked if we could do absolute performance measurements, or whether they are all relative.

Kahn said that measurements could contribute to the understanding of various problem areas, for example: performance tuning and other parameter optimization, retransmission strategies, alternative routing strategies, minimum delays for interactive-type traffic, cost considerations, robustness and recoverability.

Mayne mentioned two important incentives for carrying out comprehensive measurements:

1. Validation of simulation runs and theoretical models.
2. Obtaining insights into what is happening, especially in network operating problems of crucial practical importance.



### Documentation

Mayne said that it would be valuable to prepare a comprehensive list of all documents relevant to network measurement, including papers on methods and tools, measurement software and write-ups of that software, and collections and summaries of actual performance data. He suggested that each organization participating in the internet, should contribute relevant entries, and he himself will prepare an annotated bibliography of relevant UCL literature. He is also willing to act as a mailbox for collecting information from the other organizations.

It was pointed out that important documents in this area have been written by Kleinrock, McQuillan, and Wingfield, and also at UCL.

### Recommendations for Further Action

1. Study network performance in relation to types of network subsystem, including network as a whole, gateways, hosts, operating systems.
2. Write specifications of performance measures for TCP (it was suggested that this might be done by Wingfield), user level protocols, and IP.
3. Define a set of tests for IP, Datagram, TCP, higher level protocols, etc.
4. More work needs to be done on IP, for example, using GGP and echo packets and measuring round trip times, also investigating loss rates including numbers of duplicate packets and packets out of order.
5. Prepare an adequate bibliography of documents on performance measurements, and keep it up to date.