

Internet Meeting Notes -- 7-8-9 October 1980

I. INTRODUCTION

The meeting was held at the Royal Signals and Radar Establishment in Malvern, England. John Laws was the host. Alan Fox gave a welcome to and explanation of RSRE. John Laws described the arrangements.

II. OBJECTIVES

Vint Cerf described the objectives of the meeting. The main points were to exchange information on the status of various projects, to discuss plans, and to review performance and design issues.

III. STATUS REPORTS

A. ARPA

Vint reported on some changes in the ARPA IPT office. People: (1) Dr. R. Kahn was married, (2) Col. J. Hammett is leaving the office for an assignment in Europe, (3) Dr. B. Leiner is joining the office to work on the Packet Radio Program, and (4) Lt. Col. Duane Adams is now Executive Assistant to the Director of IPTO. Equipment: Much of the equipment on the 7th floor will be moved to the 8th floor, the current 316 TIP will be replaced by a C30 IMP plus 316 Terminal Access Mini-Host, the XGP will be replaced by a Penguin, and a RAPICOM 450 facsimile machine will be installed.

An important goal for protocol implementations is to eliminate the use of NCP and completely switch over to TCP by January 1983.

ARPA is providing technical support to the National Science Foundation (NSF) in the planning for a Computer Science Network (CSNET) to be operated by a consortium of Universities. L. Landwebber at the University of Wisconsin is the coordinator of the consortium. The current plan is to use a combination of existing networks to provide the physical support for CSNET.

ARPA intends that by the end of 1982 all existing 516 and 316 IMPs will be replaced by C30 IMPs. The first few will go to (not in this order) SRI, MIT, ISI, ARPA, BBN, BRAGG, and Berkeley.

B. BBN

Dale McNeill reported on the status of SATNET. The PSP terminal provided by COMSAT has been installed and the SPADE terminal has been removed. The ARPANET line 41 between NORSAR and SDAC which is supported by a satellite is now using a military circuit and availability has been noticeably poorer.

CATENET gateways attempt to split the load when routing information indicates two or more paths of equal cost. This means that traffic from the UCL gateway over the SATNET to the ARPANET will be split between the BBN gateway and the NDRE gateway. When it turns out that the ARPANET line 41 is down, half of the traffic goes into a black hole. The NDRE gateway can talk to its IMP and so reports it is connected to the ARPANET, but the ARPANET is partitioned due to line 41 being down. Due to this problem, it was agreed that communication between the NDRE Gateway and the UCL Gateway be permanently broken to avoid packet loss due to gateway load splitting when line 41 is down.

There has also been an increase in the packets lost in SATNET due to an increase in the bit error rate which is due to satellite power being reduced by 1db since June 1st.

Dale also reported that the VAN Gateway (an ARPANET-TELENET gateway) work has progressed to testing the Frame Level protocol. There seem to be some problems with the interface due to differing interpretations of X.25. This gateway is implemented in LSI-11.

Ginny Strazisar reported on the recent development for the Gateways. The use of XNET has converted to version 4. A problem with the Redirect Message has been fixed. The implementation of fragmentation has not been completed.

A problem with reinitializing the UCL gateway was discussed. It seems as if some combination of using XNET, the robustness card and LDSRV should be able to handle this case. UCL, SRI, and BBN will investigate.

Jack Haverty reported that new TCP implementations are now just beginning for the VAX Unix (v.7), the HP 3000, and the new TIP (mini-host attached to a the C30 IMP). Design notes on these will be distributed as IENs. A description of the performance measurement tools developed for DCA should be available about 1 November. A local net based on the Ungerman-Bass NET-ONE equipment is being tested. A design for a TIP Login system is in progress. A gateway between the RCCNET and the ARPANET and a

local net based on C30s is being installed. IEN 158 was issued which describes the XNET-4 protocol.

David Flood Page reported that the CMCC log files are now primarily event driven and summaries are produced. If you want to receive the daily statistics, send a message to David. IEN 157 was issued which discusses some new performance measurement CMCC message formats.

C. COMSAT

Dave Mills reported on activities at COMSAT. COMSAT is working on a revised version of INTELPOST which will use IP and TCP. Dave described the configuration of equipment at COMSAT which consists of a number of small hosts, mainly LSI-11s. This network uses the logical host field of the Internet Address to identify the host. Some bugs were found in some ARPANET hosts treatment of this field. The local net protocol is IP.

Dave has been particularly active in testing TCPs and IPs, and will discuss some of these issues in another session (see section V). COMSAT has implemented programs to send and receive computer mail on the "playpen" host. These are written in C. COMSAT has also used NIFTP to transmit files between their hosts and ISIE. The NIFTP software was provided by UCL. COMSAT and ISI have exchanged facsimile images.

COMSAT plans to create a full CATENET gateway and to arrange a permanent connection to the ARPANET.

D. DCA/DCEC

Ed Cain reported on the status of the EDN. Fragment reassembly has not yet been implemented. The Host and IMPs have been converted to 96 bit leaders. The gateway has implemented most of the CMCC measurements reports. One problem is that many gateways and hosts do not have the EDN in their tables.

E. DOD

Ray McFarland noted that Ken Shotting has done further work on a formal analysis of IP and a report will be forthcoming.

F. ISI

Jon Postel reported on the installation of the PENGUIN laser printer and the communication of data to it via IP, UDP, and TFTP. The program that manages the input and output of facsimile data

has been modified to use either the new or old format. Jon reported that the work on protocol verification is progressing and that Carl Sunshine may have a report at the next meeting.

Jon also reviewed the IENs and RFCs produced since the last meeting. Since the whole ARPANET is to convert to IP and TCP it is appropriate for protocol specifications and implementation related memos be RFCs while research and design related memos be IENs.

G. Linkabit

Estil Hoversten noted that Linkabit recently merged with another company and will be setting up a private corporate satellite network. Linkabit is doing the final testing of the ESI and expects send it to ISI in mid-October. The expected review of ST protocol will not be presented.

H. LL

Cliff Weinstein described the structure and status of the WBCNET. The initial four stations (LL, ISI, SRI, DCEC) all have their antennas installed. Work is underway to bring the net into operation in the next few months. LL is developing a local network called LEXNET, some digital voice terminals (VTs) and traffic emulator and traffic concentrator hosts. An 11/45 is being programmed to be an WBCNET-LEXNET gateway. This will be a ST gateway but not initially a full IP gateway.

I. MIT-LCS

Dave Clark described the complicated collection of networks and hosts at MIT. There are several families of protocols in use and several interconnections. Things are complicated enough that Corbitto has been designated czar of networks at MIT. Some things in progress are: Mail Service Programs, FTP programs, an NIFTP, X.25 interconnections for Telnet and Tymnet. The Nu machines are coming along; an IP for the Nu is being written. Multics implements an Echo and Echo.Reply to do PING with COMSAT. MIT still needs an additional IMP port. Interested in IP and TCP implementations for super small machines.

J. NAVELEX

Frank Deckelman noted the interest of the Navy Electronics Laboratories in the ARPA Internet.

Internet Meeting Notes

K. NDRE

Yngvar Lundh reported on the local net development at NDRE which is based on protocol processors implemented on Z80s with 65 kb of memory. Each processor has a kernel and a specialized module such as a speech packetizer, an LPCM, a TCP, or MVP. The implementation language is PLZ. A highly formalized graphical description language is being used to design the TCP.

L. RSRE

Brian Davies reported on the activities at RSRE. RSRE has been very active in measuring TCP performance and will report some results in another session. Some suggestions for changes to IP and TCP are: (1) If the option code could indicate by a type or class bit if the option were to be copied or not on fragmentation things would be easier for the gateway; (2) also on fragmentation it might be better to break the datagram into equal sized fragments rather than a maximum sized fragment and the left over; (3) TCP should focus more attention on the critical impact of timeouts on performance, especially noting that different destinations may vary in round trip times by an order of magnitude.

The line between RSRE and UCL was recently upgraded to 4800 baud. RSRE particularly noticed the lost traffic due to the problems with line 41.

M. SRI

Ron Kunzelman reviewed the status of the Packet Radio networks, and discussed a problem with terminal access from Ft. Bragg TIU users at ISID. The normal TOPS-20 character-at-a-time remote echo seems to saturate the bandwidth available via TCP, the gateway, or the networks. A line-at-a-time local echo mode has been developed to reduce the traffic load.

Ron also described some modifications to ACC controllers for the 1822-LSI11 interfaces:

- (1) 18 bit addresses
- (2) address incrementing
- (3) switch to tell (?)
- (4) cycle shortening
- (5) last bit on gather write
- (6) last input character
- (7) substitution of a part

A PDP 11/44 has arrived. This will run Unix (v.7) and will be used for a measurement host. The antenna for WBC is installed, but some concern has been expressed about safety to passers-by.

Jim Mathis reported that TIU's now do reassembly of datagram fragments.

Holly Nelson described some new software development tools for use on UNIX. Also noted that the Port Expander has all the features that were discussed at the last meeting.

N. UCL

Robert Cole described the activities at UCL. Fragmentation is implemented for sending datagrams and reassembly for incoming fragments. Chris Bennett is developing a Unix based mail server based on MMDF and MH using TCP and IP for one underlying communication medium, and X.25 for another; there is also interest in interworking with CSNET for computer mail. Steve Treadwell has developed some programs to decode and smooth facsimile data.

A Cambridge Ring local network is being installed, but the protocols have not been chosen as yet. Connections via X.25 to IPSS and EURONET are being tested.

UCL has also been affected by the poor availability of SATNET.

Ron Jones presented measurements on Port Expander and 1822 interface performance. The maximum achievable packet rate through the PE from a source to a sink was 65 pkts/sec (with the source and sink directly connected the rate was 255 pkts/sec). The 1822 interface performance was found by looping packets of various sizes through the PE. The interface bit rate for a PI interface connected to a DMA interface was 71 kbits/sec and for two DMA interfaces was 325 kbits/sec. The latter experiment also yielded 5.7 ms as the time for the PE to switch a packet.

IV. ACTION ITEMS

A. Internet Name Server

SRI is developing an internet name server consistent with the specification in IEN 116. It was expected to be demonstrable at this meeting. The development has not reached such a state and this item is deferred until the next meeting.

B. Interim Internet Mail System

ISI is developing a prototype mail program for supporting text mail in the internet and between NCP and TCP hosts. It was expected to be able to demonstrate these programs at this meeting. The design of these programs has only recently been developed (see section XII). This item is deferred until the next meeting.

C. FTP and MAIL Memo

ISI has produced a memo on text mail for the Internet. It is RFC 772. RFCs 771 and 773 offer motivation and discussion of this approach (see section XII).

D. IP and GGP Error Report Memo

ISI was to produce a memo on error handling in IP and GGP. This was not done and this item is continued to the next meeting (see section XV).

E. XNET Specification

BBN has produced a memo on the XNET protocol. It is IEN 158.

V. TAKE OFF SUMMARY

Dave Mills reported on the sessions for testing the fragmentation and reassembly implementation in the IP protocol modules. The results indicate that such testing is useful in that a few problems were found and corrected, but disappointing in that a number of hosts and gateways don't implement fragmentation as yet.

VI. TINY PIPE NETS

Dave Mills discussed some performance problems with networks composed of low speed lines (e.g., 1200-1900 baud) and small hosts with minimal buffers. In such "tiny pipe" networks, attention must be paid to the performance effects of TCP segmentation decisions, acknowledgment strategy, window strategy, and retransmission timeout

selection. These performance parameters must be chosen with respect to the partner at the other end of the TCP connection. For example, a connection between a tiny host in a tiny net and a large host in a large net (e.g., the playpen in COMSAT net and ISIE in ARPANET) has very different characteristics than connections between roughly equal hosts.

VII. INTERNET PERFORMANCE MEASUREMENT

Zaw-Sing Su presented a detailed description of some measurement ideas.

The aspects of performance measurement can be identified as follows:

- Artificial Traffic Generation
- Data Registration
- Data Collection and Transport
- Experiment Preparation and Control
- Data Reduction and Presentation
- User Interface

Measurements could be made at several levels:

- TCP Performance
- TCP Analysis
- IP Performance
- IP Analysis
- Internet Component Performance

Where performance is as perceived by the user and analysis is by mechanism.

The types of things studied in performance measurements are:

- Traffic Arrival and Departure Statistics
- Throughput and Delay Performance

The types of things studied in analysis measurements are the mechanism used by the protocol. For TCP:

- Relative data efficiency
- Retransmission sent and received
- Out of order arrivals
- Duplicates
- Ack Only segments
- Ack Delay
- Zero Window periods

Internet Meeting Notes

For IP:

Fragmentation and Reassembly statistics
Use of options statistics

For GGP:

Routing Decisions
Congestion Statistics

Zaw-Sing suggested an internet delay measurement capability to be implemented as an optional header field for both TCP and IP.

The TCP option would have source and a destination timestamps of 32 bits each plus the option code and length octets for a total length of 10 octets. The time would be in milliseconds.

The IP option would have the type and length octets, a 16 bit id, and N stamps, where each stamp is a 32 bit IN address and a 32 bit time (in milliseconds). The maximum number of stamps would be 4 due to the limitations on the header size.

Questions were raised about combining this option with the Source-Route or Return-Route options.

The consensus seemed that rather than an option in the IP header the timing data ought to be accumulated in the data part of the datagram in a protocol like GGP.

VIII. CATENET PERFORMANCE CONSIDERATIONS

Bruce Hitson discussed the need for one-way delay measurements in the Catenet. It was stressed that these measurements would be particularly critical for performance measurement in an internet environment. The very large physical distances involved and the non-participatory (at least in performance measurement) nature of the member networks are examples of why one-delays would be very useful. Round trip measurements are of questionable value since the network may change quite a bit during the course of a single round trip (up to 10 seconds in some cases). To be able to make one-way delay measurements, synchronized clocks are needed.

Bruce described several clever ways of synchronizing clocks with third party time sources. The best choice seems to be some devices that listen to a radio station (i.e., WWVB in the United States) and have a computer interface (e.g., RS 232) to report the time to the computer on request.

Internet Meeting Notes

ISI (as noted in the notes of the last meeting) has already ordered 4 such devices for use in WBC and CATENET measurements. Design of experiments that will make use of these devices is being considered. For further details, see Appendix B.

IX. Loader Server MEASUREMENTS

Jim Mathis reported the results of some round trip measurements made using the LDSRV program. A round trip in the ARPANET between DARCOM-KA and the Ft. Bragg Gateway took about 600 ms. Similar measurements between DARCOM-KA and a host in a local PRNET took 170 ms., and between DARCOM-KA and a host in the Ft. Bragg PRNET took 800 ms. For further details see Appendix C.

X. CATENET MEASUREMENTS

Brian Davies discussed some suggestions for performance improvements based on the experience at RSRE.

The use of an adaptive retransmission timeout seems to be very helpful. RSRE has experimented with one based on the following:

1. For each segment record the sequence number and time sent.
2. For each acknowledgment determine the round trip time (RTT) of the sequence number thereby acknowledged.
3. Compute an Integrated Ack Time (IAT) as follows:

$$\text{IAT} = (\text{ALPHA} * \text{IAT}) + \text{RTT}$$

4. Compute a Retransmission Time Estimate (RTE) as follows:

$$\text{RTE} = \text{Min} [\text{BOUND}, (\text{BETA} * \text{IAT})]$$

Where BOUND is an upper bound on the retransmission time and BETA is an adjustment to the IAT to account for variation in the delay.

RSRE currently uses $\text{ALPHA} = 31/32$ and $\text{BETA} = 1.33$.

[Dave Clark noted that MIT-MULTICS uses the same algorithm but with $\text{ALPHA} = 4/5$ and $\text{BETA} = 1.5$.]

Brian presented some graphs of the actual RTTs and resulting RTE for TCP connections between RSRE and ISIE.

Andy Bates presented data from recent measurements of double round

trip times between RSRE and UCL, BBN Gateway, and ISIE. There are two aspects of the measurement to be explained: The absolute delay, and the variation in delay. The absolute delay (for the least delayed packets) seems to be in agreements with the expected delay given the speed of the lines, the protocols used, etc. For example, one second time for a separate reservation is required (which is the current strategy). The variation in delay is not yet satisfactorily explained. Some arises in SATNET, and a good deal arises in the ARPANET. One suggestion (by Danny Cohen) is that if the packets are far enough apart in time then new Source IMP to Destination IMP reservations are needed for each packet, this could add a large variable delay.

A goal for the next meeting is to figure out the cause of the variation in the delay.

XI. NBS PROTOCOL STANDARDS DEVELOPMENT

Mike Wingfield discussed the standards work on protocols sponsored by NBS. Mike first noted that many standards groups are actively working on standards in the communications area. The work NBS is sponsoring at BBN is to develop service specifications, feature analysis, formal specification, and a test implementation for a transport protocol.

BBN has developed a formal description language which is an extension of a finite state machine description. The language is entirely textual and so can be easily manipulated with computer tools.

The basic element is a transition description which is composed of a next state, a current state, an event, and a procedure. The procedure is written in "C". The following tools for working with the language have been developed: a syntax checker, a state checker, a special editor, and a test generator.

Additional work involves a feature analysis of X.75, IP, and PUP as internet protocols. NBS is also interested in the ECMA proposal for a Transport Protocol and BBN is analyzing it.

XII. MAIL TRANSITION PLAN

Vint Cerf described the problem arising in the exchange of computer mail between old NCP only hosts and new TCP only hosts. The plan is to provide relay or forwarding hosts that implement both transport protocols and relay the mail.

Jon Postel explained some of the details of this plan. A key decision is to provide a new MAIL SERVER and a new mail protocol.

These will be very similar to the existing protocols and servers but will have an important new feature: the address passed in the MAIL command will include the destination host as well as the user.

Jon also discussed the need for hosts to add to their host tables some indication of the status of each other host: old table NCP, new table NCP, or TCP (all TCP hosts are to have new tables). These three kinds of host mean there are 9 combinations of mail exchange. These were discussed on a case by case basis. The difficult case was the old NCP host to the TCP host. In this case a relay host could receive the mail using the old mechanisms but would have no idea of which host to forward it to. One proposal is to have a list of registered users available to this relay host.

Vint also discussed the idea that the hosts have unique identifiers independent of network, and even that organization names could be used instead of host names.

The discussion indicated some doubts about the practicality of registered users, especially in resolving name conflicts now resolved by host names. Also, concern was expressed about the globally unique host names.

RFCs 771, 772, and 773 present the plan, the mail protocol, and a discussion of the issues. Comments are solicited.

XIII. CONTROLLED ROUTING

Danny Cohen discussed routing to avoid undesirable networks or to limit a route to known networks. The existing source routing option may be called loose source routing (LSR) since it allows the gateway to send the datagram on any route they choose between the specified addresses.

Danny proposes another option called strict source routing (SSR), which is like LSR except that a gateway may only route a datagram to the next specified address through the network specified in the NET part of that address.

This is very restrictive and would highly control the route. It would also provide a way to route things out of one network, through a second network, and back into the first network, to get around a partition or use special transmission capabilities.

Vint Cerf proposed an alternative of simply a list of acceptable networks. This would be less restrictive, but might lead to problems in routing.

Danny's proposal is described in IEN 156.

XIV. PROTOCOL ISSUES

Dave Clark discussed some problems in IP and TCP implementation in the MIT environment.

One situation is that several hosts at MIT are on two or more networks, and have as many internet addresses. The problem is that while one interface is down, the host and the other interface may be up, but datagrams sent to the down interface address will be discarded rather than routed to the other interface address.

Dave presented a scheme that would allow a two-connected host to avoid this problem with the help of nearby gateways. The gateways in the networks the hosts is connected to would have tables to map a special form of address into one of the regular address and to choose between regular addresses based on knowledge of which interfaces were up. The hosts sending to the two-connected host would always send to the special address.

This scheme was not well received. The consensus seemed to be that it was too specialized and involved too much special case mechanisms. It is agreed to be an important problem, but a better solution is needed.

Another situation involves very small hosts and the desire to implement TCP in them. For example, a TCP in 2K bytes of memory. What corners can be cut? For example, no data with SYN or FIN, simple ACK and window policies.

One suggestion is for an option to specify a maximum receive segment size to be sent only with a SYN. This would have different effects than a consistently small window because it may allow many segments to be in transit at a time even though each is small. The consensus seemed to be that this was a good idea.

XV. IP AND GGP ERROR REPORTS

Jon Postel reviewed the discussion from the last meeting of error reporting in IP and GGP. They are different in that in IP an option is used and in GGP the data area is used to carry the error information. They are redundant in that most of the errors mentioned in IP are also mention in GGP.

Jon proposed to put all these error reports into a new protocol and use a GGP style mechanism. The consensus seemed to be that we should continue to use GGP, and eliminate the IP error option.

XVI. NEXT MEETING

The next meeting will be hosted by ISI and will be held 28-29-30 January 1981. Attendance will be limited so if you plan to attend please clear that with Vint Cerf (CERF@ISIA). and notify Victor Brown (BROWN@ISIF). Jon Postel will be the host.

A tentative schedule for subsequent meetings is: May 81 at COMSAT, September 81 at UCL, January 82 at SRI, May 82 at BBN, and September 82 at IABG/DFVLR.

AGENDA ITEMS:

1. Provision for source routing in the interface between TCP and IP and in the TCP tables. How will the reverse route information be handled?
2. Further discussion on internet mail transition. Focus on alternatives to the unique name requirement which led to potential misrouting.
3. More on performance observations within the Catenet, in the Hosts.
4. MITRE report on their new 350 kb/s TCP that runs in a Z-8000.
5. Front-ending of TCP/IP.
6. Progress reports on FTP where ever it is being implemented.
7. Progress reports on the (interim internet) message system.
8. Progress reports on the protocol verification. ISI, BBN, SDC, and DoD.
9. Discussion of "worm" programs, XEROX.
10. Discussion of the IIN problem, the inter-inter-net. How and at what levels can we interoperate with other nets?
11. Report on On-Tym and Telemail relative to ARPANET mail and Internet Mail.

ACTION ITEMS:

1. New text for the TCP and IP specification changes.

APPENDIX A: INTERNET DESIGN PHILOSOPHIES & THEIR IMPLICATIONS

The third day of the Internet Meeting was a seminar on the TCP/IP and Yellow Book approaches to internetting. The following notes are provided by Brian Davies of RSRE.

INTRODUCTION - John Laws (RSRE).

John Laws started the informal Seminar by giving some background as to its genesis: namely that with so many TCP/IP experts having come to the UK and it being the case that the UK public network users are promoting an apparently competing solution, then it seemed an ideal opportunity for parties to exchange views. It was not expected or intended that there should be any dramatic conversions of the parties vis-a-vis datagram and virtual call issues. However, it was hoped that areas of common approach might be identified even though expressed in a different language. In addition, the reasons for divergence might be mutually identified and appreciated.

An "impartial" Chairman, Derek Barber (Microprocessor Applications Project, Dept. of Industry) was introduced. Many presentations of views were scheduled and firm discipline, of both presenters and the audience, was administered by the Chairman. A summary of presentations and ensuing discussion follows.

OPERATIONAL REQUIREMENTS - Sqn Ldr David Stupples (RSRE).

The user requirement for communications in a typical air defence scenario was presented. The data communications were provided by a common bearer network and mobile tactical networks of the JTIDS type. The data-bases and their functions were described to illustrate why they were distributed. The mobility of the users not only in a single net but across multiple tactical nets was emphasized with particular reference to the problems of maintaining connections to their associated databases.

PHILOSOPHY AND ARCHITECTURE OF YELLOW BOOK TRANSPORT SERVICE (YBTS) - Peter Linington (Data Communications Protocol Unit, Dept. of Industry).

To ensure that the parties to the debate were speaking a common language a set of "basic" and "obvious" axioms for internetworking were presented. The properties of a global transport service were described. Computability for internetworking would be achieved by enhancing each underlying network to the common transport service standard. This enhancement would be within hosts and gateways; the enhancement being achieved by encapsulation of the specific

network features within appropriate network specific protocols. Thus only the service attributes have to be agreed globally; hosts only have to implement their local network enhancing protocols; gateways at worst have only to implement neighbour network enhancing protocols.

ADDRESS AND ROUTING IN THE YELLOW BOOK TRANSPORT SERVICE -Mike Guy
(Computing Centre, Cambridge University).

The point was made that diverse network address naming authorities will always exist and will always assert their technical need to be different. A naming and addressing mechanism appropriate for this environment was described. This featured use of an "address" string whose component parts are addresses, titles, generic names. Traversal of connected networks was driven by successive "activation", with possible address/title expansion/transformation, of the component parts appropriate to the current address/naming domain: in effect akin to source routing. In particular, the use of titles/generic names allowed the user to be unconcerned with changes in remote network addresses or interconnection architectures. Depending on the name/address conventions of the neighbour networks, gateways would be required to have varying degrees of intelligence about name/address mappings.

TRANSPORT SERVICE ARCHITECTURE FOR THE LAYMAN - Vince Hathway
(National Physical Laboratory).

Vince detailed his varied and long experience of implementing and working with protocols in an internetworking environment. This environment had gateways to both datagram and virtual call networks. Both approaches had used hop-by-hop techniques for some of the levels of protocols. In addition one gateway also passed high levels of protocol transparently. Neither approach could be declared as "better". The community of users associated with each of the external networks had different requirements and this had significantly determined the internetworking solutions. In analogy, domestic cars were ideal for many transportation requirements, but there are times when the virtues of tanks are appreciated!

TCP AS A BASIS FOR YELLOW BOOK TRANSPORT SERVICE REALISATION -
Chris Bennett (UCL)

In order to incorporate an existing network into the Yellow Book framework, a 'Yellow Book protocol' must be defined locally to that network which builds on the existing network interface to bring it in line with the Yellow Book service. In the Catenet,

the appropriate starting point is TCP which already offers the basic virtual call interface. The data stream is structured into arbitrary length data and control messages. Each control message is allied with TCP actions, where possible, to gain the appropriate effect. The TCP-based Yellow Book protocol is not complex and is mostly required to allow the transmission of connection information beyond the Catenet. Details may be found in IEN 154.

A USER REQUIREMENT FOR STANDARDS - Ken Heard (Joint Networks Team, Science Research Council, Rutherford Labs.).

Background was given as to some of the functions/responsibilities of the Joint Network Team; this being advising/aiding the selection/implementation of an architecture and protocols for internetworking between diverse networks sited at UK Universities and Research Establishments. The point was made that the productive functioning of this internetworking in the immediate future was urgent and, that time and resources could not be wasted on re-inventing the wheel. Hence the adoption and rapid implementation of protocols emerging as "interim" standards for use on the UK PTT public network. A review of protocol/architecture standardisation activities within national/international bodies (AFNOR, BSI, CCITT, ECMA, ISO, NBS,...) was given. As a final comment attention was drawn to the influence on standards of the pioneering research networks such as ARPANET and the European Informatics Network.

THE INTERNET ARCHITECTURE - Vint Cerf (DARPA).

The wide range of diverse network implementations and environments incorporated in the ARPA Catenet was described. The ARPANET itself, packet radio networks, broadcast packet satellite networks, and a variety of local networks all interwork using IP as the universal protocol. Some of the experiments and demonstrations conducted were discussed.

INTERNET PROTOCOL DESIGN - Jon Postel (ISI).

The architecture of the ARPA Catenet protocol family was briefly described and the services provided by the IP and by the TCP were described. The point that both protocols are available to the users and can be used to build very different kinds of applications was emphasised.

END-TO-END VS HOP-BY-HOP - Danny Cohen (ISI).

The difficulty of building effective communication systems by

relying on plug-to-plug compatibility on a hop-by-hop basis was discussed.

SUMMARY OF ISSUES - Brian Davies (RSRE)

The advantages of harmonization of transport protocols in the civil and military fields were presented. However, the differences in emphasis in the user requirements would probably preclude this standardization in the foreseeable future. These differences included availability of resources in the face of threat, mobility of users, and the military requirement to accommodate a dynamically changing internet topology. In addition, the relative importance and costs of such properties as security, precedence and survivability had to be considered. Some of the particular architectural issues involved in the Yellow Book and TCP/IP approaches were then presented. In particular the implications of the addressing strategies, connection control and economies of use were highlighted.

DISCUSSION

The initial discussion was used by a number of the presenters to clarify facts or points they had made in their talks. The following paragraphs are a selection of the varied topics discussed:

- a) The number of "grades" of service that might be required was discussed. In the YBTS approach it had been identified that the one grade of reliable connection-oriented service was a requirement overwhelming all others. In contrast, the TCP/IP approach had identified a need for a number of distinct "grades" of service to be realized by specific protocols (e.g. speech, user datagram etc.).
- b) Discussion showed that different emphasis of requirement had significant influences on the two architectures. The more commercial environment within which the YBTS has been formulated places great emphasis on the economic use of resources. However, the problems involved in maintaining connections across vulnerable or unreliable networks was an aspect not deeply considered within YBTS. The TCP/IP approach had given much greater attention to this military aspect. The two approaches could be seen to be non-competitive as they are solving different problems.
- c) The YBTS approach to addressing is based on the premise that all networking authorities will not agree to a global addressing strategy. The global addressing approach of the TCP/IP would seem to be in-line with the ISO recommendation on this issue.

APPENDIX B: SYNCHRONIZED CLOCKS

ISI is evaluating an absolute time clock for network measurement use. The clock is actually a radio receiver which is tuned to National Bureau of Standards station WWVB, which broadcasts on a frequency of 60kHz, and is located at Fort Collins, Colorado. WWVB provides both time and frequency information and its coverage area is effectively the continental United States.

The clock is the Model 8170, made by Spectracom Corporation, 320 N. Washington St., Rochester NY 14625. The 8170 is a rack-mountable unit, 17" wide by 5.25" high by 13.5" deep. Line power is 115/230 volt 50 or 60 Hz AC at 60 watts. An external antenna is required, which is a ferrite loop antenna in a tubular plastic housing 14.8" long and 2.7" in diameter with a threaded fitting in the center which accepts a standard 1" threaded plumbing pipe (for mounting purposes only). The antenna is connected to the clock receiver via 50 ohm (RG58) coaxial cable with BNC connectors on the ends.

For best accuracy, the antenna should be placed outside just above rooftop level (like a TV antenna). However, an inside location near a window facing Fort Collins may turn out to be acceptable.

Overall accuracy is specified as plus or minus (.1 millisecond + noise uncertainty + propagation delay setting error), where noise uncertainty is expected to be less than plus or minus .5 millisecond worst case. This accuracy should be more than sufficient for network measurements.

Time is output via a front-panel display (hours, minutes, seconds) and an RS-232 (D-15) jack on the rear panel. Also output on the back panel is a 1Hz, 10% duty cycle TTL signal called the "on time" pulse. The leading edge of the on time pulse occurs at the beginning of each second. In addition, a 1MHz standard frequency output is available on the back panel, which is a 3.4V rectangular positive pulse designed to drive a 93 ohm load, but which will drive a 50-ohm load with reduced amplitude. This 1MHz output is phase-locked to the WWVB carrier, and can be used for calibrating counters with high accuracy.

The user requests the time by sending the clock an ASCII capital "T". When it receives the "T", the clock waits until the start of the next second, and sends back an ASCII string containing the time. The rising edge of the start bit of the first character in the string occurs within a few microseconds of the rising edge of the on time pulse.

The time string is structured as:

(CR)(LF)I DDD HH:MM:SS TZ=XX(CR)(LF)

where:

I =space if time synch lamp is on,
 ? if time synch lamp is off
 DDD =Julian day of the year (000 to 366)
 HH =hours
 MM =minutes
 SS =seconds (at the start of this timeout)
 XX =time zone setting with respect to UTC
 (formerly called GMT). EST is 05, PST is 08.

The RS-232 port on the clock runs at a baud rate of 300 baud.

The synch lamp on the front panel will be lit if the time can be expected to be reliable within plus or minus 1 millisecond. The clock contains an internal 1MHz oscillator which is kept synchronized to WWVB. If the WWVB carrier is lost, the internal oscillator will keep running, and accuracy to plus or minus 1 millisecond will be maintained for about 10 minutes, at which time the synch lamp will be turned off, although the oscillator keeps running.

The back panel also contains thumb wheel switches for time zone setting, propagation delay setting, and a leap year switch. Yes, Virginia, a leap year switch. The purpose of the leap year switch is to maintain the day count properly if the WWVB carrier is lost near midnight of Dec. 30 (on a leap year) or Dec. 31 (on other years).

Price is \$2600 for the clock itself, \$190 for the loop antenna, and \$35 for a rack mount kit.

APPENDIX C: LOADER SERVER MEASUREMENTS

The numbers were average round-trip times calculated by the LDSRV. Each number is the average RT time seen during the load operation, which takes about 350 or so packets. RT times are not calculated for packets that are retransmitted.


The average RT time to ARPANET hosts on the same IMP is around 70 msec. There were some times somewhat less and a few a bit more. To MIT, the usual average RT time is about 450 msec. To the Ft. Bragg PE or gateway, the smallest average RT time is 600 msec. The largest average RT time was 1.2 sec.

The usual average RT time to SRI PRNET TIUs is 170 msec. The smallest time was not less than 100 msec and the largest is 400 msec.


To Ft. Bragg PRNET TIUs, the smallest average RT time is a bit more than 600 msec. The usual RT time is around 800 msec and average RT times as large as 2.1 sec have been recorded. Most of the loads had average RT times less than 1.1 sec.

In the graphs, the y-axis is the number of loads in which the average round trip time was within the time bucket.


ARPANET Hosts



SRI PRNET Hosts



Ft. Bragg PRNET Hosts



Internet Meeting Notes

APPENDIX D: RECENT DOCUMENTS

Author	Subject	Number
	IENs	
Postel	Internet Meeting Notes - 14 & 15 May 1980	145
Perlman	Flying Packet Radios and Network Partitions	146
Perlman	Utilizing Internet Routes as Expressways Through Slow Nets	147
Postel	Telnet Protocol Specification	148
Postel	File Transfer Protocol Specification	149
Plummer	TCP JSYS Calling Sequences	150
Cerf	Final Report of the Stanford University TCP Project	151
Cerf	DoD Protocol Standardization	152
Bennett	Realization of the Yellow Book Transport Service Above TCP	153
Bennett	Realization of the Yellow Book Transport Service Above TCP (supersedes IEN 153)	154
Bennett	The Yellow Book Transport Service: Principles and Status	155
Cohen	Controlled Routing in the Catenet Environment	156
Flood Page	CMCC Performance Measurement Message Formats	157
Haverty	XNET Formats for Internet Protocol Version 4	158
	RFCs	
Postel	Internet Protocol Handbook TOC	766
Postel	Structured format for Multimedia Mail	767
Postel	User Datagram Protocol	768
Postel	RAPICOM 450 Facsimile Format	769
Postel	Assigned Numbers	770
Cerf	Mail Transition Plan	771
Sluizer	Mail Transfer Protocol	772
Cerf	Comments on the NCP/TCP Mail Service Strategy	773
Postel	Internet Protocol Handbook TOC	774

APPENDIX E: ATTENDEES

Name	Affiliation	Nationality	Mailbox
Vint Cerf	ARPA	USA	Cerf@ISIA
Don Allen	BBN	USA	Allen@BBND
David Flood Page	BBN	British	DFloodPage@BBNE
Jack Haverty	BBN	USA	JHaverty@BBND
Bruce Hitson	BBN	USA	BHitson@BBND
Dale McNeill	BBN	USA	DMcNeill@BBNE
Virginia Strazisar	BBN	USA	Strazisar@BBNA
Mike Wingfield	BBN	USA	Wingfield@BBND
Gustav Fickenscher	BWB.MoD	Germany	---
Hoi Chong	COMSAT	Rep. China	Chong@ISIE
Dave Mills	COMSAT	USA	Mills@ISIE
Ed Cain	DCEC	USA	Cain@BBNB
Hans Dodel	DFVLR	Germany	---
Helmut Lang	DFVLR	Germany	---
J. Majus	DFVLR	Germany	---
Gerry Amey	DND/CANADA	Canada	---
Ray McFarland	DoD	USA	McFarland@ISIA
Romy Fratilla	ELEX	USA	Deckelman@ISIA
Horst Clausen	IABG	Austria	Clausen.IABG@MIT-Multics
Danny Cohen	ISI	USA	Cohen@ISIB
Jon Postel	ISI	USA	Postel@ISIF
Cliff Weinstein	Lincoln Lab	USA	cjw@LL-11
Estil Hoversten	Linkabit	USA	Hoversten@ISIA
Dave Clark	MIT	USA	Clark@MIT-Multics
Frank Deckelman	NAVELEX	USA	Deckelman@ISIA
Oyvind Hvinden	NDRE	Norway	Oyvind@DARCOM-KA
Yngvar Lundh	NDRE	Norway	Yngvar@DARCOM-KA
Paal Spilling	NDRE	Norway	Spilling@SRI-KL
Vince Hathway	NPL	British	---
Andrew Bates	RSRE	British	T45@ISIE
Trevor Benjamin	RSRE	British	T45@ISIE
Brian Davies	RSRE	British	T45@ISIE
Roger Edwards	RSRE	British	T45@ISIE
Alan Fox	RSRE	British	T45@ISIE
Silvia Giles	RSRE	British	T45@ISIE
John Laws	RSRE	British	T45@ISIE
Paul Masterman	RSRE	British	T45@ISIE
Jo Penley	RSRE	British	T45@ISIE
Gill Roberts	RSRE	British	T45@ISIE
Ron Kunzelman	SRI	USA	Kunzelman@SRI-KL
Jim Mathis	SRI	USA	Mathis@SRI-KL
Holly Nelson	SRI	USA	HNelson@SRI-KL
Zaw Sing Su	SRI	Canada	ZSu@SRI-KL

Internet Meeting Notes

Chris Bennett	UCL	British	UKSAT@ISIE
Robert Cole	UCL	British	UKSAT@ISIE
Ron Jones	UCL	British	UKSAT@ISIE
Steve Treadwell	UCL	British	UKSAT@ISIE