

Network Meeting Report

This is a report on a series of three Network Working Group meetings at the Fall Joint Computer Conference, November 16, 17 and 18 in Houston, Texas. The meeting will be lumped together and ideas may or may not be identified as to their originator. The meetings were chaired by Steve Crocker.

The meetings began with a listing of topics of concern.

- 1) A site or group should be designated as protocol testers. As NCP's are implemented they should be subjected to comprehensive testing by an independent group.
- 2) The Host-Host protocol needs reworking in several areas: error control, overload conditions, allocation of resources, status information, and system crash problems.
- 3) The immediate need for specification of TELNET, the third level program which allows people to pass through their local hosts and use remote hosts. TELNET must provide facilities to log in at a distant site, run programs, transmit files, and call for help. This call for help is likely to mean getting a systems programmer at the remote site "taking control" of the user console.
- 4) The documentation of systems on the network must become available to all sites. This is to be done by the NIC with the cooperation of the other sites. Particularly useful will be on-line documentation. It is suggested that each site have an identical hard copy device (e.g. a line printer) suitable for reproducing documents.
- 5) The use of graphics consoles on the network will require a graphics protocol. People interested in this problem should write position papers on such a protocol. A meeting may be held between the authors of such papers if sufficient interest develops. The papers should be distributed as NWG/RFC's before 1 January 71.
- 6) Some sites must account for use of their computer resources, thus there must be some network accounting scheme. Sites can be categorized as Research Centers vs. Service Centers. The Service centers tend to have big machines, lots of users, and accounting problems; while the Research Centers tend to have specialized hardware, a small number of users, and no accounting at all.

- 7) Some people are interested in the network as an object of study. In particular UCLA-Computer Science, and BBN wish to perform measurements on the network. Is it appropriate to ask the NCP to keep statistics?

After this opening some discussion followed.

It was generally felt that changes to the protocol should be made in bunches and at about six-month intervals rather than a continuous stream of small changes. Also that a lead time of three months was not over optimistic. The proposed change to the IMP-Host protocol to get rid of marking was generally approved but it will not be implemented for some time since casual changes to the protocol are undesirable. Tom O'Sullivan suggested that perhaps new and old protocols could work together, that is the new protocol would support the old one as well as provide better mechanisms where possible. Steve Crocker suggested that a new protocol might be developed as a private experimental protocol between two or three sites.

It was stressed that it is necessary that the network be used to gain experience, and that we should get teletype-like console use of remote systems going before we get too involved in graphics. Perhaps the graphics protocol should be developed by a different set of people. The scheduling of a graphics protocol meeting was thus discouraged, but papers should still be written. Strong feelings were expressed in favour of first developing use of remote subsystems and file transmission instead of worrying about graphics at this stage. It was suggested that development of protocols at the higher levels be driven by applications.

Documentation will be a major concern for network users. Several people mentioned that users at their sites have already begun to inquire about the network. As Eric Harslem put it "What does the ARPA Network have to offer?" Some sites (Multics, SRI) keep system documentation on-line. It was suggested that the trillion bit store be used to keep on-line documentation of all systems.

At this point Doug Engelbart gave a presentation on the Network Information Center (NIC). The goals or services of NIC have not been well defined by anyone and have been left up to NIC to define. NIC has decided that one urgent task is to make information about the network and the host systems on the network available to users of the network. Doug has found that some people feel threatened by the revelation of their documentation inadequacy. Doug's project at SRI has built up a system that allows the user to create catalogs and indices into a collection of information. The system has a master catalog of all information files. Each user may have a number of private (or shared) catalogs. The system provides a means of examining on-line the catalogs

and amending them. The system also provides a means to examine any information file which happens to be on-line and for creating new information files on-line.

Several problems will delay the NIC from coming on the network. One of these is the switch from the XDS-940 to the PDP-10 (TENEX). The switch is being made because the 940 system is inadequate to handle the anticipated load. At first it was planned to offer service on the 940 and switch to the 10 when it came up, but too much effort would be required for a very small payoff.

Doug explained the working of the Network Dialoge System. At each site there is a communication agent and a technical liaison officer. The agents will be trained by NIC to use the facilities of NIC to get information about the Network and other sites. The agents will acquire from NIC documents of interest to users at the local site, be able to contact NIC at a toll free number, and should have an on-line console into the network (and therefore NIC). Thus the Network Dialoge System is a network of people (the agents).

Steve Crocker then brought us up to date on the status of the network. He drew a picture of what is connected and what is proposed. He discussed the level of implementation at various sites. Eric Harslem mentioned that RAND and UCSB had conducted tests of their NCP implementations last week (10 Nov 70) and that things worked well.

Frank Heart announced that BBN was planning the development of a "Terminal" IMP. The Terminal IMP would support some large number of a wide range of consoles as well as provide the normal IMP functions.

At this point we broke and scheduled to reconvene Tuesday morning.

The Tuesday meeting started with Doug giving another pass at explaining the SRI system at a more detailed level. The basic thing to deal with is the collection. The user can query over the collection or over sub collections. The user can obtain bibliographic references of three kinds: a) full references, b) first line, c) author indexed. Information files of the collection may be on-line, in tape libraries, or only in hard copy. It is suggested that much data could be kept at other network sites, for example the trillion bit store and before that perhaps on disk at UCSB. If files are kept at other sites then the system must be able to retrieve them automatically when they are requested. The subsystem to be used is called TODAS. TODAS is an evolving program and the documentation of TODAS is inadequate. In TODAS, files are organized hierarchically, each paragraph is numbered, and it is possible to do context analysis on the text.

Doug then mentioned some things about the console interaction. This raised a question about half vs. full duplex and line oriented vs. character oriented systems. The remainder of the meeting revolved around this issue.

I shall try to define the terms as I understand them for purpose of clarity in the following. Half duplex is the situation where the console, a peripheral processor or some very low level software, echos the character entered. The console can not be used to input data while output is in progress. Full duplex is the situation where the character typed is echoed by software, and input can be done at the same time as output. In line oriented systems the user enters a complete line terminated by an extra sensitive and of line character (e.g. carriage return). Often the keyboard is then locked until after the next output. In character oriented systems each character the user enters is interpreted by software before it is echoed and the echo may be different from the character entered. In particular after a few character the software may guess what the user wants and complete the line for him. The following chart will be used for clarity.

	Half Duplex	Full Duplex
Character Oriented	type1	type2
Line Oriented	type3	type4

It was discovered that many people don't really know where their own systems fit in this chart and are very vague about what it means to interact with a system in a different than their own. Doug stated that NIC has a system of type 2 but would try to provide service to all types of systems. The following table shows systems with their interaction type and categorization as to Research vs. Service Center.

System	Interaction Type	Categorization
UCLA - Sigma-7	2 - char, full	Research
UCLA - 360/91	3 - line, half	Service
MIT - Multics	3 - line, half	Service
SDC	3 - line, half	?
RAND	3 or 4 - line, ?	?
SRI	2 - char, full	?

Al Vezza promised to study this problem and to circulate his results as an NWG/RFC. It was pointed out that line oriented systems usually allow line editing of the form "delete last character" (back space) and "delete line", however this feature does not alter their classification as to interaction type. Concern arose over what do line oriented systems expect to receive from the network for a connection acting as console input to a subsystem. Steve Crocker made the suggestion that when using a line oriented system transmission be in lines. More precisely that transmission be in strings of the following form.

n c1 c2 ...cn

where $1 \leq n \leq 120$ (n is eight bits)

and if c_i is an "end of line" character then $i = n$

This suggestion was not immediately accepted and some discussion took place regarding the significance of Host-Imp-Host message boundaries. Doug brought up file transmission and the problem of finding the end of the file, which provoked more discussion. At this point the meeting broke up with a third session scheduled for 8:00 p.m. Wednesday evening.

The Wednesday meeting began with the suggestion that at future xJCC's there be an official ARPA Network hotel with a block of rooms on one floor and a nearby meeting room for networkers. This suggestion was favored by all.

Steve Crocker asked how people felt about these meetings. The general feeling was that the meetings were very useful and should occur about 3 months apart. Al Vezza pointed out that meetings this size (15 - 30 people) are good for bringing up problems but not for putting them down. Steve proposed that 3 or 4 people be designated to solve particular problems. Al responded that 3 people can't legislate. That any such

solution must be considered in the same way as a proposal by an individual.

Steve persuaded Peggy Karp to act as NWG/RFC editor. This is a job independent of cataloging RFC's or assigning numbers (functions now performed by NIC). The RFC editor will only categorize RFC as "hot issues", current, out of date, or superseded.

The subject of Logger protocol -- that is, how to get the first connection -- needs to be officially defined. NWG/RFC #66 suggests one way. Eric Harslem will revise this and send it out as proposed official protocol. Ed Myer will also send out a proposal.

Steve then opened up discussion of the topics of the previous meeting by suggesting we talk about the following: Message boundaries, half duplex vs. full duplex, line oriented vs. character oriented, file transmission, byte counts in messages, byte sizes and transactional units. It was proposed that transactions on the command link (i.e. between NCP's) be always in multiples of eight bits. This means that the length field in the ECO, ERP, and ERR commands will always have three low order zeroes. This was approved. Steve then proposed that connections could be established with a declared byte size and a maximum record length in bytes. Transactional units on this type of connection would be of the form

$$n \text{ } c_1 \text{ } c_2 \text{ } c_3 \text{ } \dots \text{ } c_n$$

where $0 \leq n \leq \text{max record length}$

if $n = 0$ then the transactional unit acts like a semaphore. Steve suggested that we should look into the theory of information exchange, particularly along the lines of Richard Kaline (NWG/RFC #60). Perhaps for each information unit sent there should be some status response.

The next question was on file transmission. In particular, how do you find the end? Frank Heart suggested that with each portion there be a flag indicating "this is not the end" until in the last portion the flag is switched to indicate "this is the end". Eric Harslem suggested that each portion should have an "opcode" field, a length field, and the text which is length bits (bytes?) long. This appears to be like the data types proposed at the Lincoln Lab meeting last spring. Ed Myer proposed that two connections be used, one for the file transmission and the other to control it. The file control connection would specify the data connection and indicate that transmission is about to start. After the sender had completed the file transmission he would send on the file control link the total number of bits sent. The receiver would then know how many bits to receive and exactly where the end of the file should be. Bob Metcalfe was concerned that some of the proposals mixed

control information with data and felt that perhaps this mixing should be avoided.

Steve asked if anybody could suggest an advisor we might talk about these problems. Bob Metcalfe suggested Anatol Holt. Bob Sundberg suggested George Mealy. Eric Harslem and Peggy Karp suggested that people who worked on the COIN System might be helpful. Frank Heart suggested that no one has solved these problems.

Steve proposed that Service Centers offer line oriented interaction with no echoing of the input. Any simple editing (e.g. back space) would be done at the using site. Ed Meyer suggested that there be official protocols for both line oriented and character oriented interaction. Steve promised to write a NWG/RFC clarifying the issues and laying out the arguments on full transactions, byte counts, and accumulating data on the receive side.

It was felt that these were hard problems that needed more thought. Thus the meeting was adjourned with the request that people circulate any ideas or proposals as NWG/RFC's. Ed Myer took notes and agreed to also prepare a NWG/RFC summarizing these meetings.

Network Meeting Attendance List 16 - 18 Nov. 70 Houston

Name	Site	Sessions
1. Dick Benjamin	MITRE	1
2. Jack Bouknight	Illinois - CAC	1,2
3. Al Cocanower	MERIT	1,3
4. Steve Crocker	UCLA - SPADE	1,2,3
5. Doug Engelbart	SRI - ARC	1,2,3
6. Wayne Fischer	MERIT	3
7. Richard Greenblatt	MIT - AI	1
8. Eric Harslem	RAND	1,2,3
9. Frank Heart	BBN	1,2,3
10. Allen Joseph	ORNL	1
11. Peggy Karp	MITRE	1,2,3
12. William Kehl	UCLA - CCN	1
13. Bob Long	SDC	1,2,3
14. Jim Madden	Illinois - CAC	1,2
15. Bob Metcalfe	MIT - DM	1,3
16. Edwin Myer	MIT Multics	1,2,3
17. Ari Olrikainen	UCLA - SPADE	1,2,3
18. Tom O'Sullivan	Raytheon	1,2,3
19. Jon Postel	UCLA - SPADE	1,2,3
20. Chris Reeve	MIT - DM	1,3

Network Meeting Attendance List 16 - 18 Nov. 70 Houston

Name	Site	Sessions
21. Tijaart Schipper	UCLA - CCN	1
22. Michael Sher	Illinois - CAC	1
23. Bob Sundberg	Harvard	1,2,3
24. Hal Van Zoeren	CMU	1,2,3
25. Albert Vezza	MIT - DM	1,2,3
26. Alfred Vorhaus	MITRE	1
27. Clark Weissman	SDC	1

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Gottfried Janik 02/98]

