

Network Meeting Notes

The following summary was transcribed from notes I took at three network meetings held in Houston during the 1970 Fall Joint Computer Conference. Although I have tried to be objective, unavoidably these notes present a biased view of the meetings. This is due in part to my preoccupation with certain topics and possible misunderstanding of various discussions. While I have tried to accurately paraphrase the statements of the attendees, the import of some may have been twisted.

Attendees of Monday Meeting

Dick Benjamin	MITRE
Jack Bouknight	UI-CAC
Al Cocanower	MIRUT
Steve Crocker	UCLA
Dough Engelbart	SRI
Richard Greenblatt	MIT-MAC
Eric Harslem	RAND
Frank Heart	BBN
Allen Joseph	ORNL (Oak Ridge)
Peggy Karp	MITRE
William B. Kehl	UCLA
Bob Long	SDC
Jim Madden	UI-CAC
Bob Metcalfe	MIT-MAC
Edwin Meyer	MIT-MAC
Ari Olrikainen	UCLA
Tom O'Sullivan	Raytheon
Jon Postel	UCLA
Chris Reeve	MIT-MAC
Tjaart Schipper	UCAL-CCN
Michael S. Sher	UI-CAC
Bob Sundberg	Harvard
Hal van Zoeren	CMU
Albert Vezza	MIT-MAC
Alfred H. Vorhaus	MITRE
Clark Weissman	SDC

Network Meeting
8:05 PM Monday, 11/16/70

Crocker: Not everybody is here, so lets talk until more people get here. is everybody satisfied with the agenda in my announcement ?

Meyer: We should talk about logger protocol. Operational usage of the net- work, as opposed to experiments, depends on its implementation.

Introductions to all around.

Crocker: I have an agenda, but want suggestions for topics.

- 1) I will make introductory remarks.
- 2) I will list topics of concern.
- 3) Englebart will talk about the Network Information Center
- 4) I will review the status of sites.

Introductory remarks

- 1) ARPA will not pay for the coffee and pastry being served, so please chip in to help me pay for it.
- 2) I am going to devote full time to network coordination in an official capacity. My goals are: (a) to build up usability of the network. (b) to establish protocol levels, (c) ?

Areas of importance

- 1) Some site of coalition of sites should prepare a method by which a site's NCP could be checked out.
- 2) Reworking of NCP protocol. Some issues could be solved better: (a) error control, (b) flow control, (c) overloading - loosing network states, (d) simplification and relayering of protocol.
- 3) Telnet system console interaction, or logger protocol. How to get into the system and how to get help when in trouble.
- 4) Documentation of individual hosts. Network Info Center involved. Perhaps each site could be provided with a facsimile device.
- 5) More sophisticated consoles, particularly graphics consoles, to be attached through network. There should be a working group to formulate and workout a format for handling sophisticated consoles. There will be a graphics meeting in January in Colorado or Utah. The price of admission is to write a proposal. I expect up to 30 people. I will pick a small subset to develop specifications.

- 6) Accounting - In the 2nd half of 1971 more sites will come on where accounting is important. (They want to send bills.) Larry Roberts says that there will be a kind of banking system with bills passed around. Two types of sites: billing sites, and free but limited access research sites. I see no fundamental problems. What happens when a research site talks to a billing site? I think it is do-able.
- 7) Measurements - the network is a tool, but it is also a model that is better than a simulation package. Various people want to make measurements. This could be supported by keeping statistics in NCP's What about increasing the NCP's to include these?

Long: Putting accounting and measuring into NCP's costs space. Keep additions to a minimum.

Weissman: What about scheduled availability of various systems?

Crocker: This has to be coordinated with each individual system

? : What happens to connections when a system goes down?

Crocker: What about graphics proposals? I will write my own paper as a proposal. It uses the DEC 340 as a model. Modes assumes scope system a memory. Both output-input are included in standards making. I want a competent protocol to be developed out of the working group.

Crocker: What about documentation?

Meyer: Documentation on how to use other systems are a must. Only this can motivate operational use of the network.

--: What about putting documents on-line at each site, or at least abstracts.

Crocker: What sites have documents on-line? (MIT and Harvard) How do the sites feel about keeping documents on some foreign system?

Crocker: What about reworking the protocol?

Harslem: We have logged into the UCSB system and are debugging cooperatively.

Harslem: We are impressed with eliminating marking and padding (per RFC 67).

Crocker: We discussed this with the sites. Most seemed to accept it, but some reservations. What about changes to the basic protocol. I'm Meyer has something to say.

Meyer: The position at Project MAC is that at this point we are opposed to changes other than critical fixes. Time spent on changes is time that won't be spent on developing other necessary and interesting protocols and systems. And we at Multics have a long lead time for creation and installation of changes.

Weissman: I prefer to put in changes in one chunk, say at 6 month intervals. rather than in bits pieces.

O'Sullivan: Can't current and new systems work simultaneously?

Crocker: If the changes involve the IMP, no, because all IMPs want to operate the same system.

Meyer: The feeling at M.I.T. is that to be a success, the network needs desperately to be used operationally. If another year passes without significant operational use, it might go down the drain.

--: And documentation is critical towards motivating operational usage.

Engelbart: Perhaps we should put off graphics several months so as not to delay typewriters. Typewriters are important.

--: But would that be sufficiently impressive for DOD people?

Engelbart: But if it turns out to be a can of worms in two years...

--: But do the two (typewriters and graphics) development groups interact?

Veza and Engelbart: Yes.

Crocker: Let's hear more about this.

Harslem: We want to be able to access files.

Crocker: Then perhaps the graphics effort would dilute typewriter development. Is it the consensus of this group that we shouldn't have a graphics meeting?

Veza: Newcomers should work on graphics, not established people. Prohibit current people from going to this meeting.

Meyer: That would be very frustrating.

Benjamin: Why not solicit position papers (but have no meeting).

Weissman: Character transmission is easier than graphic transmission
More experiments needed for graphics. The lead time for
developing a graphic protocol is much longer than for typewriters.

Veza: I agree.

Crocker: There will be more meetings in the next few days to work on
problems of getting useful work over the network.

Intermission

9:15 PM


Crocker: Engelbart will speak on Network Information Center.

Engelbart: NIC grew as an ad hoc thing, with no specific directives
from ARPA. What kinds of things were envisioned? (1)
Sophisticated query systems, (2) Basic information about systems
at each site. Everyone feels very vulnerable about the state of
documentation at his own site. Everyone agrees: better documents
necessary. We see ourselves as providing the following services:
1) collecting hard-copy material; 2) on-line querying of catalogs
and indices of these; 3) giving access to this material. We
decided to go hard copy rather than on-line, perhaps on
microfiche.

Engelbart: As 940 was to be used for the documentation system,
expandable as usage increase. We are switching from a 940 to a
10X to better expand service capacity. Amount of capacity goes up
considerably. This has held up work on other facets. A conscious
gamble. We are worried about getting off the ground. We are short
on funds for more secondary storage and are interested in using
other hosts for tertiary storage. The cost of implementing the
protocol on the 940 was too high for potential gains, so it was
given up. Few sites would be up by January when our 940 was to be
shipped out.

Engelbart: We have created a Network Dialogue System. This is a
network of human agents. At each site there is: a) technical
communications agent (secretary) and b) a technical liaison
person. We are encouraging agents to talk to us and have created
"Enterprise" phone numbers so they can talk toll free.

Engelbart: We are at first sending out a tiny kit to each agent, a growing collection of network reference information. One person (agent) at each site is to be trained to handle the set of documents and retrieve information of contact another site's technical liaison. This involves a public dialogue, keeping a record of the documents passing back and forth. This is a sort of "human IMP" network, structured as follows:



- 1) Master collection has all material.
- 2) Each local collection has a subset considered most useful.

--: What about restricting access to documents?

Engelbart: All files are public files in this system.

Veza: You can send a private memo rather than use the NIC service.

Engelbart: The master collection contains books and other documents. Cataloged on-line. Hard copy stuff can be duplicated. For information that passes the value test, the service is to store, catalog, index, and provide access to documents. We will support a number of different terminal. We are prepared to go a long time with hard copy items, but can establish a hard copy to on-line transcription service for a price.

Weissman: What about distributing OCR Selectric balls to sites?

--: Will NIC take what is sent or actively search it out?

Engelbart: More or less what comes our way. A system will exist in Spring 1971, to allow an agent to insert items into a catalog. The dialogue that goes on will determine which way the data base grows. We are pretty sure that eventually SRI will have to charge because of many potential users not at primary sites seeking limit resources.

--: What about an NCP for your 10X.

Engelbart: If BBN's NCP is ready by February 1971, we'll use it.

Crocker: How do people get access?

Engelbart: Each site is registered. Any person who gets in on a site's account has its access. We won't worry about accounting until saturation occurs. We would like to encourage use of the agent system to create and use a survey of resources at each site. Some subgroup should talk about this.

Crocker: When can people meet to discuss this? (Tomorrow morning)

Engelbart: We have nice facilities for developing mailing lists, private bibliographies, personnel profiles, but it depends on the interest of the network people.

Engelbart: Agents have been set up with MIT, UCLA, RAND, UI, Utah, etc. A good percentage of sites

Veza: Many sites are sending out stuff 3rd and 4th class. Takes too much time.

Crocker: Site status report. ILLIAC IV not to be operational until mid-71, on network later (72?). Other possible sites: RADC, AWS, NCAR. Currently up: UCSB, RAND. Imminent (January 71): MIT BBN, Harvard, UCLA, Utah, LL, SDC. Some percentage by end of year rest in January.

Heart: A brand new IMP system (major change) goes in tomorrow. Some more sites are thinking about coming on. The network will grow considerably beyond what's already on board. We too are interested in site resource information. No long term interest, but we will put information on paper to help ARPA.

Crocker: A lot of people are yawning. What about meeting schedules? During FJCC's? 1 day vs. 2 day meetings? What about dual East and West coast meetings?

End of Meeting

Network Meeting
9:15 AM Tuesday, 11/17/70

Crocker: Engelbart will talk in more detail. Later may discuss logger protocol and file transfer.

Engelbart: Basic thing is a collection of documents with a catalog to describe it. Entry has lots of data items, including where to find it. Techniques for adding and updating entries. We do it now, but would want to give capability to other sides, partly because we can't determine what's of value. (Displayed 3 types of printout.) 1) Catalog listing, by ordinal index in collection and NIC index. for inventory control, finding out what's there. 2) Compacted format on one line. 3) Sorted by author-one line per entry. We Will have procedures where an untrained user can manage a collection.

Meyer: How are these systems implemented?

Engelbart: We have a compiler-compiler on the 940. Our subsystems are written in specialized high level language. We are moving this over to the 10X.

Heart: How many people can the 10X support in rough figures?

Engelbart: Perhaps 100-1000 collections.

--: Perhaps people could supply own DEC tapes for additional storage.

Engelbart: Could, but requires on-site operator. Slow access. We don't have money for more storage, but are considering shipping files down to UCSB. We provide on-line querying of on-line data. Willing to worry about data management, whether or not we store it.

Crocker: Please describe the various subsystems. (Description by Engelbart follows.)

Heart: Have people tried to use it over the network?

Engelbart: No. Don't have an NCP on the 940. Decided against putting it in a system that is going away. The biggest hang up is when 10X gets an NCP. Bobrow developing it, but it is slipping.

Heart: Who is going to get at it (SRI) early?

UI: Illinois can access only SRI to begin with.

Postel, UCLA: We plan to use it.

Heart: Would be a significant task if someone would take it as a goal to get into Engelbart's system.

MITRE: We're going to use other systems form BBN's 10X.

Engelbart: We are trying to isolate essential subsystems for people to use easily. Files are organized hierarchically and will fill out as years go by. Documents are referenced by pathnames. (Discussion of systems follows.)

Crocker: How does one get into the system? (Engelbart describes entry sequence to T0das.)

Crocker: How does one get registered on the system?

Engelbart: Ultimately by personal entry, but currently there is one user id per site.

Meyer: I think we're ignoring unsolved problems in the typewriter interfacing. For example, the entry sequence to T0das, where the user type one or two characters and the system types out the remaining chars of a key word, will be frustrating to use form half-duplex system like Multics. Our system will not recognize an input line until a new line is typed.

Various: Discussion of 1/2 duplex communication. Brings out distinction between a) Full duplex systems where system echoes input vs. 1/2 duplex where input typed locally, and b) systems where each character is recognized as it is typed in vs. systems where entire line is recognized only after EOL char.

Crocker: Isn't Multics the only half-duplex line-oriented system on the net- work?

Meyer: I can't believe this. Don't the IBM systems operate like that?

Engelbart: We could have a 1/2 duplex interface on our system (SRI). Is it the Multics hardware that enforces this restriction?

Meyer: Yes, the input-output controller.*

* The Multics IO controller's typewriter adaptor is 1/2 duplex, but can accept break characters other than the "new line" character.

Engelbart: Each system should have a preprocessor to talk with other systems. We're going to put a graphics interface onto the network.

Meyer: How do you view these interfaces? Do these adhere to some network standard, or ill each system construct an interface to you?

Engelbart: Standard network protocol.

Crocker: Let's move on to other things.

O'Sullivan: What about 2741's on your (CMU) 10X system. Do you have serious interfacing problems? (CMU's 2741's go through a software package that transforms them into TTY 37's. No serious difficulty.)

Various: Brief discussion on how Multics handles input.

Sundberg, HARVARD: Out 10X can take char-oriented input, but our higher level subsystems prefer line-oriented input.

--: What about the efficiency of transmitting messages through the network one character at a time?

Crocker: There is more output, which goes packed, than input, so the input inefficiency is negligible.

Engelbart: We plan to have several different ports into our system. If each system had an NIC module; it could communicate with us without the necessity of a login. We prefer a batch-type system, where a site sends spooled batch of edit requests, gets stuff back, and frees ports. The typewriter transmission by line problem could be handled similar to spooled-up requests. We encourage spooling, but will support interactive users. We can support more batch than inter- active people.

* The Multics IO controller's typewriter adaptor is 1/2 duplex, but can accept break characters other than the "new line" character.

Veza: Do people feel that the full and 1/2 duplex issue is a problem? Let all the people go back and find out about this. M.I.T. with a full and 1/2 duplex system 20 feet apart can help here.

O'Sullivan: There seem to be 2 issues: (1) echoing (full duplex) vs. 1/2 duplex. (2) single character vs. full line transmission.

Crocker: Two definitions: Serving host - provides computation; using host - parasitic, manages user's terminal. This view network usage as a link between local user and foreign server.

Veza: What about 1/2 duplex - full duplex interconnection if some full duplex systems echo other than what was input.

Crocker: Two independent possibilities. Let's diagram:

	"2741" hard wire local echo computer does not echo	"33, 35, 37" 2 separate lines all printed
Process each character	hard	X
Process only after EOL	X	easy

Crocker: I claim there are really only two possibilities (marked by X's).

Postel: Why about a system where echoing is done at such a low level that it can't be deleted.

Crocker: If so, it's like non-echoing.

Van Zoeren: Our system thinks we have full duplex TTY's, but our 2741's are attached via a software transformation box.

Meyer: What happens when non-echoing systems are attached to echoing systems via the network? I type my input line, then the echoing system responds with my input, then some output. My system can't filter this because there is no way of differentiating echo from output.

Crocker: This isn't necessarily a bad thing. I type command abbreviation to SRI; then next output line is an expanded form of the input command.

Meyer: Our goal should be one common protocol rather than a bunch of kludgy schemes to implement communication between specific pairs of hosts.

Long, SDC: We prefer to receive a full line fed through the network.

Crocker: Let's differentiate between research centers and service centers. Only the service centers are concerned with a half-duplex interface. (lower left hand X on chart). These include SRI, BBN, Multics.

O'Sullivan: What about research center?

Crocker: They can call up service centers, but may themselves be hard to use.

Illinois: Then the ILLIAC IV will have to be half-duplex.

Postel: I think half-duplex, line-oriented is weaker (than full-duplex, character-oriented protocol).

Sundberg: Harvard can go either way, but prefer line-oriented system.

Engelbart: Graphics terminals harder to put on network because of non-standard input.

Harslem: You're thinking of the keys as function keys rather than input keys.

Engelbart: I'm worried about people who want to use graphics.

O'Sullivan: We haven't spoken to the problem of what kind of protocol should be established.

Crocker: That's not a difficult technical thing. We'll get to that later and make a decision.

Meyer: I'm not authorized to make any decision. I'm to report back to the MAC group.

Crocker: Okay. Then, a proposal, to be accepted through the normal mechanism. Intermission

Crocker: I will propose how to handle X'ed boxes, ignoring hard and ease boxes:

Line-Oriented Input - 8 bit ascii including End of Line character:

n, C1, ..., Cn;

Cn=EOL

120>n>>_1 n is the character count in an 8-bit field. The character count precedes the line so as to give the software system the same efficiency as the hardware system, the computer doesn't have to scan for the EOL.

Veza: Don't you get the length information with the IMP message?

Crocker: My philosophy is that IMP message boundaries should be completely invisible.

Long: I object to splitting typewriter messages into two separate chunks.

Crocker: What is your objection, 1) Lines beginning on message boundaries or 2) message not beginning on a line boundary?

Long: Both.

Engelbart: Each host should write an interface to handle the most common terminal types.

Crocker: The official protocol does not allow IMP message boundaries to have any significance.

Engelbart: I don't want to worry about IMP message boundaries. The network should be invisible (at this level).

Veza, Long: We'll concede, we'll go along.

Meyer: I'd like to change the restriction. The last character in the line packet need not be an EOL (as when an output does not advance to a ne line), but an EOL cannot appear in the midst of a packet.

Van Zoeren: I don't like this restriction.

Meyer: Count tells us that any EOL is at the end, we need not scan.

Crocker: The EOL is the character that tells the system to take action.

Harslem: Our system has 46 function keys, not just one EOL.

Crocker: How about if C, E {breakset}; i=n. This s more complex, because have to transmit a breakset. I'll propose this in a moment. How about this: message oriented (1/2 duplex) connection

between User and Server hosts for console interaction. Local echo, no server echo. This is for line oriented service systems. These are slight generalizations of Multics conventions.

Meyer: I'm sure other systems other than Multics use it. It's not as bad as you seem to think.

Engelbart: Upper management should know that it is bad.

Meyer: That's not clear. There are efficiency questions.

Van Zoeren: I don't want to have to transmit files this way.

Crocker: This is for consoles, not file transmission.

Engelbart: We need a unified scheme for data transmission.

O'Sullivan: (For consoles) we're to devise a way to tell a system where its interrupt should be simulated.

Crocker: There is a general problem of data transmission for tapes and files.

O'Sullivan: But we have the specific problem of implementing typewriter communications.


Engelbart: But what we need is a general way of sending stuff through the network (so it is invisible) and have the host interpret it as it wants to.

Meyer: There should be one console interface for the network, not several at each site.

Crocker: This problem is perhaps overblown.

Engelbart, Meyer, O'Sullivan: Discussion about supporting specific terminal types.

Engelbart: I'll draw a graph of systems vs. terminal type. The intersection of a system and a terminal that is accepted by that system is marked by a dot. Network communication problem is one of finding a terminal at local host that is also supported at the target host.



Crocker: There is a general problem of a subsystem reacting to input. Let's propose that input should be sent as a full message or in multiples of 8-bits.

Veza: Are we constraining too much?

Meyer: Why is it necessary to have 8-bit multiples?

Crocker, Engelbart: Okay throw that away.

End of Second Meeting

Network Meeting
8:20 PM Wednesday, 11/18/70

(The following notes are greatly condensed and attempt only to present the major themes discussed at this meeting.)

Crocker: Let's meet at the SJCC with more prior organization. Let's have several day meetings at 2-3 month intervals. We've got a lot of good discussion on the next level protocol. Let a subgroup work it out.

(Harslem volunteers to redraft the logger protocol proposed in RFC 66. Meyer will revise proposal in RFC 46.)

Meyer: Let's go back, discuss these issues, write proposals. Later we have an open meeting to decide on a formal proposal.

Crocker: Small group is better, perhaps I'll pick a subset.

Veza: It's true that things aren't settled here. Major proposals should be on paper preparatory to a meeting. We can't legislate what a small group does. It has no more authority than an individual.

(Karp of MITRE volunteers to produce a bibliography of network documents, perhaps by January.)

(Who has implemented logger protocol? UCSB and UCLA mod 91 have or are planning. SDC may have it by 21/1, found it awkward, willing to change.)

(Discussion of file transmission. Crocker proposes that a future protocol change might attach a byte size such as 8, 32, 36 bits to a connection.)

(Regarding control links, everything is transmitted in 8-bit bytes except ECO, ERP, ERR commands, No objection was voiced to changing the protocol so that they also must be multiples of 8-bit bytes.)

(Discussion of how to specify the end of a file. Prior transmission of bit count, or send EOR character at end? Suggestion that we want global solution to the general problem of sending an arbitrary length message, rather than just file transmission.)

(Discussion of "transaction units" or record sizes. What is an optimum transaction unit size? IMP message boundaries are invisible (by protocol fiat) and are not connected with this discussion. Multics block size was brought up. Nearest thing is page size, 1024 words.)

(How to specify end of file. Engelbart says send data packets, then EOF packet. Crocker suggests that CLSing connection can act as EOF. Veza suggests that IMP message boundaries be used to determine end. If less than full IMP message, this is last part of file. Meyer suggests use of two connections, data channel and control channel, over which all control messages, such as file name, bit length, etc. are passed.)


(Discussion concerning different situations in which whole file, part of the file, or the whole file in arbitrary chunks was wanted.)

Meyer: Why not defer this, and talk about typewriter communications, which is most critical.

Veza: Engelbart wants a clean general solution.

Crocker: If we get an ad hoc solution now, it may interfere with implementing a general solution later.

(Crocker proposes format for transmitting a file of arbitrary length records of fixed sized bytes of 8-26 bits. A record is less than 10^5 bytes. Each record is headed by a count byte.)



O'Sullivan: Does this model fit a terminal which has character and graphic modes?

(Discussion about differences between keyboard and file transmission. Uncertainty as to whether a global solution would fit both.)

(Who wants to ship files through the network? Multics and 6-10, RAND to UCLA, MITRE using BBN.)

Crocker: Let's go away thinking about this and propose solutions later.

(Harslem proposes format for transmitting data with operation codes. Each record consists of: <opcode> <length> <data>. Gives the opportunity to send many type of status info.)

(Discussion regarding sending data and control information intermixed or on separate connections. Issues of pollution of data vs. synchronization and race problems. Claimed that synchrony problems are easily overcome.)

(Suggestion that we really don't know much about this area. We should go off and write.)

Intermission

Crocker: What has to be done before we can log onto other systems?

Meyer: 3 issues: 1) ho to establish the connection, 2) what is the character set, 3) what is the mode of transmission (relating to full and 1/2 duplex problem).

(Discussion of orienting standard protocol towards service systems which generally are line-oriented and 1/2 duplex. Any systems offering services to have a 1/2 duplex interface.)

(Discussion of whether possible or desirable for the logger protocol to allow transmission of partial lines in a IMP message. Less efficient to take partial lines, reasonable to send full line. Pointed out that NCP protocol disallows any meaning to IMP message boundaries, so systems must be prepared to accept lines straddling IMP message boundaries. However, best to send complete line.)

(Discussion of whether line-oriented protocol should bend so as to accept single character transmission from full duplex systems. Seems that we are coming up with a protocol to allow any system to use a line-oriented system. To use a char-oriented system form other systems is more difficult and requires a separate protocol.)

Heart: I am in favor of an immediate solution.

Postel: Once something goes in, it will be hard to change it.

Crocker: I think these meetings will turn out to be more important than we ever wanted them to be. I am more concerned with the long term effects than the starting date.

Van Zoeren: If we don't decide it, somebody else will decide it the bad way.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Gottfried Janik 2/98]

