

Minutes of TCP Meeting, March 12, 1977, Washington, D.C.
Prepared by Carl Sunshine

- I. Vint Cerf introduced participants, and outlined the agenda for the meeting (distributed in advance via ARPANET mail). Copies of the Version 2 TCP Specification were distributed.
- II. Bob Kahn suggested that defining basic terms like "network" and "gateway" is surprisingly difficult, and that this problem deserved further attention (see point VI).
- III. TCP Version 2 - Vint Cerf
 - A. Three way handshake connection establishment for reliability as before.
 - B. Sequencing, fragment reassembly, duplicate detection as before.
 - C. Flow control (window) as before.
 - D. Simplified resynchronization does not need additional states, just normal ACK.
 - E. Simplified TCP-TCP error messages.
 1. "Destination unreachable" eliminated because it is not reliable information. (Sunshine and Kirstein objected to this.) Suggested that this could be provided by gateway-gateway level exchange outside of TCP.
 2. "Nonexistent connection," "inappropriate SYN," and RESET combined into single RESET message -- receiver can tell which is meant. Question: Does this work with encryption?
 - F. "Graceful" connection termination.
 1. Eliminated flushing with FIN.
 2. FIN must be acknowledged by receiver so sender is certain connection closed.
 3. Question: how to close gracefully when reconnection by an intermediary is involved?
 - G. Interrupt has had flush semantics removed, now just "pokes" receiving process.

- H. The FLUSH control function has been eliminated.
- I. A new Abort command has been defined to immediately terminate a connection (erase TCB and optionally send a single RESET).
- J. A new ARQ control function has been added for reliably opening zero windows. Question: Does it work if both windows are zero and receivers have "taken back" some window space?
- K. Packet format changes (see New Spec. page 72)
 - 1. An "Internet" (e.g. datagram) header has been defined to precede (the rest of) the TCP header.
 - 2. Header length and version number fields allow variable "option" fields (e.g. for Resynch, Timestamps) to follow the header. Question: Should internet options be after TCP header? How to checksum options, if desired?
 - 3. TCP (host) address field = 3 bytes, port field = 4 bytes.
 - 4. Other changes in line with A - J (e.g. control bits).
- L. Letter handling and EOL bit.
 - 1. Letter marking optional by sender (with each send command).
 - 2. At most one letter per buffer on receiving.
 - 3. Discussion: Should TCP have letters at all? If so, what should they mean?

Uses of EOL:

- a) Producer says send all data up to this point (at the end of a "processable" unit).
- b) To tell consumer when a "processable" amount of data has arrived.
- c) To mark logical records.

Suggestions (by Clark, Reed): Letter markings belong in data stream, not TCP control. More than one kind of "flag" in data stream may be desirable (as in their interrupt facility). Receiver may not agree with or care about "processing points" indicated by sender.

Suggestion: If EOL used at all, receiver need only remember most recent EOL (and only its presence, not its place) in the data stream, since several pending "process now" marks are no better than one.

Suggestion (by Kahn): Need synchronization of time among multiple connections (see below). Should interrupt serve this purpose?

Suggestion (by Kirstein): Receiver may want to issue special signals as well as sender (e.g. to stop output).

IV. Comments on scope, approach (spontaneous)

- A. Crocker: Document needs "purpose" section, and a metric for evaluating goodness of alternatives.
- B. Cohen: Should design be narrowly constrained as in current specification, or remain flexible to allow various modes of operation for different applications? (See point V)

V. Packet Speech - Danny Cohen

- A. Presented overview of packet speech architecture.
- B. Real time speech needs guaranteed bandwidth, low delay, can sacrifice some reliability.
- C. Levels of protocol for voice in internet system:
User (speaker); vocoder (e.g. LCM); NVP; end-end virtual circuit or datagram (e.g. security, end-end ACK, sequence); TC or CP (globally used internet header, address, fragmentation); gateway-gateway; local net end-end; local net internal.
- D. Main Point: CP level should allow requesting various modes of operation to support different applications.
Aspects of CP (communication profile):
 - 1. Mainly reliability, delay, bandwidth. These involve flow control, acknowledgement, error control, sorting, priority, routing, security, size, cost (limits on feasible combinations), sequencing, timing.
 - 2. These are not all orthogonal to each other (e.g. flow control needs sequencing).

3. A particular aspect is relevant at all levels (point C), but levels may not be orthogonal (e.g. end-end priority requires lower level priority.)
4. Desirable to avoid binding decisions in protocol design to retain flexibility for varying needs.

E. Ordering vs. Timing

1. Sequence number sufficient for ordering, detecting missing items, but not for timing.
2. Time stamp adequate for ordering and timing, but cannot detect missing items. "Duration" info also needed.

F. Example multi-connection application with differing needs: parallel voice, graphics, terminal channels. Should timing be done by transmission protocol, or by application? Should synchronization (in time) between parallel connections be provided?

F. Writeup is available from Cohen on this material.

VI. Bob Kahn (distributed copies of his slides)

- A. Outline of difficulty in defining basic terms like network gateway, process, port.
- B. What needs an internet address: Hosts, processes, files, records, words, users?
- C. What is structure of gateway and where are interfaces? "Pure" or "core" gateway functions vs. local net interface functions. Transit gateway vs. gateway associated with TCP.
- D. User - TCP interface option examples
 1. Inhibit error control (sequencing, retransmission).
 2. Order, but don't wait for missing.
 3. Timestamp and reproduce timing.
- E. Interface between "pure" gateway and TCP portion of gateway needs better definition, study.
- F. Broadcasting
 1. Unreliable allows class routing.

2. Reliable (need ACK from all): must know ID's or at least number of recipients.
3. Flow control ill-defined. When rates vary among participants, slow ones may have data dropped, or be temporarily "de-addressed." But key parties (i.e. the commander) must not be dropped.
4. Point-point replies are often adequate after broadcast sending.
5. Approach to conferencing: User's local net "conference controller" (CC) broadcasts to CC's associated with other nets. Each CC lines up local resources, replies to initiating CC.

VII. Error and Fault Isolation - Vint Cerf

- A. How to tell if and where an error has occurred?
Diagnostic facilities: loopback, timestamps, etc.
- B. What should we test for? Some things we assume normally work, and we leave for external test procedures.
Likely errors we want instrumented.
- C. Granularity of error and status info should match the service provided (e.g. TCP users get connection info, not internal local net problems).
- D. Providers of service need more info than users of service, and may be unwilling to disclose it to outsiders.

VIII. Extended addressing and gateways - Steve Crocker

- A. Strongly vs. weakly connected nets. The former are the permanent solution, but the latter keep cropping up as additions are made.
- B. Weak connection requires
 1. Source (or his agent) must know full path to destination.
 2. Different sources have different paths to same destination (i.e. different "names").
 3. A way of learning of paths (outside of normal channels?)

- C. These features always bother people until they realize it's the only way to do weak connection, that is until the new net can be integrated into an existing strongly connected net.
 - D. Loose internet addressing should be part of internet protocol below TCP.
 - E. Details of technique in INWG Note #133 and SIGCOMM Quarterly Review, January 1977, by Sunshine. Also copy of slides handed out.
 - F. Gateways implementing weak connection do not need to exchange routing data with any other gateways, or have other than local routing information. They just do address transformation and embedding.
 - G. How to describe gateway structure: Is a three-way gateway a single entity, or three gateways connected to an internal net? (See Sunshine paper in Computer Networks Journal).
- IX. Flow Control, Sequencing -- Bill Plummer
- A. Desire to store data in user buffers as packets arrive. Can't now because:
 - 1. Missing packets may have an unknown number of control bytes which take up sequence numbers.
 - 2. End-of-Letter bit effectively fills rest of buffer.
 - B. Hence
 - 1. Separate sequencing control from data, while maintaining their relation. "Floating point" approach, where control is fractional between integer data sequence numbers. (Bob Kahn proposes another way to do this which he will write up).
 - 2. Establish fixed buffer size, and EOL consumes any remaining sequence space at sender. This essentially converts TCP to letter based flow control. Is this desirable?
- X. MIT Local Net, DSP
- A. Decided TCP was more complex than necessary and did not satisfy addressing needs. Wanted to avoid building in any higher level functions not needed.

B. DSP Features

1. Broad interpretation of address (no host field).
 2. Avoid reuse of ports. Hence port name serves as incarnation ID to avoid need for complex initial sequence number selection, Resynch, and three-way-handshake.
 3. Interrupt ("Urgent" flag) always associated with at least one byte of data in the data stream to allow specifying many kinds of signals.
 4. Generic names used for addressing which spawn a new instance (port) each time a connection is established.
 5. "Flush" means preceding data won't be coming.
 6. EOS for checksumming.
 7. No EOL (left to higher level).
 8. Full details in MIT Local Net Notes #3, 5.
 9. Comparison with TCP being prepared.
- C. DCS and ether rings may both be implemented (as "subnets") with simple gateway (e.g. for speed matching) in between.
- D. Comment by Kahn: concentrate on designing higher level protocols so they can undo things, so that lower level protocols can be simpler, less reliable than even DSP.

XI. Action Items

- A. Comments on version 2 specification welcome.
- B. Revise TCP-DSP comparison - Clark.
- C. Write up sequencing technique - Kahn.
- D. Prepare note on gateways emphasizing interfaces to local nets - Cerf.
- E. Revise communication profile note - Cohen.
- F. Prepare note on security implications for TCP, particularly encryption - MIT or Sunshine.
- F. Comments on desirable version 3 features welcome.
- G. Plan next meeting.