M.I.T. Laboratory for Computer Science
Computer Systems and Communications

November 25, 1980
Request for Comments No. 203

Comments on the Standard Ethernet

by J. H. Saltzer

On October 1, 1980, Digital Equipment Corporation, Intel Corporation, and Xerox Corporation dropped a long-awaited shoe by publishing version 1.0 of "The Ethernet: A local Area Network, Data Link Layer, and Physical Layer Specifications."[1] This three-company proposal offers a rare opportunity for cooperation and coordination in the computer inductry in an area where it can be especially important and helpful: intercomputer communication. It would be to everyone's advantage—manufacturer and user alike—if a standard such as this one were adopted not just by these three manufacturers, but by many more manufacturers of computers, data communication gear, and remote peripheral servers. The effect is a little like that of telephone manufacture: there is an obvious advantage to every telephone being able to talk with every other telephone, rather than there being several arbitrarily different kinds, each of which can talk only to others of its own type unless someone devises a translation gadget. The ability to create a hardware communication path between any two computers simply by plugging standard components together can eliminate one unnecessary obstacle in arranging for useful data communications.

For all these reasons, the standard Ethernet is of considerable interest, and it should be reviewed by as many people as possible to gain insight on two questions:

1) Is the proposed standard suitable as it stands, or does it raise objections that could inhibit widespread adoption?

2) What out of this standard can be equally well adopted by implementors of technologies other than the Ethernet, so as to minimize unnecessary differences?

The remainder of this memorandum raises several technical points of possible concern about the standard Ethernet not because of any objection that this system should become a standard, but rather to encourage making the standard as widely usable and used as possible. The comments are in two categories: major issues, and minor observations noticed along the way. Each major comment is followed by a specific suggestion, so as to open discussion on a constructive note.

Major issues

1. *Increases in scale and complexity.*

The standard Ethernet differs in size and scale with the widely used experimental Ethernet in the following four quantifiable ways:

	experimental Ethernet	standard Ethernet	factor of increase
bit rate	2.9 Mbit/sec.	10 Mbit/sec.	X 3.4
maximum length span	1 Km.	2.5 Km.	X 2.5
maximum number of stations	256	1000	X 4
maximum bits exposed to collision	15	450	X 30

In addition, there are at least three other unquantifiable but complexity-increasing differences: branching, use of repeaters, and constraints on cable lengths and tap positions. The experimental Ethernet has received extensive field test, and appears to be a robust design. However, a scale and complexity increase in any system almost always discloses effects that were not predictable by simple extrapolation. These quantifiable and unquantifiable changes in going from the experimental Ethernet to the standard Ethernet introduce a variety of potential problems:

a) A ten Mbits/sec. data rate with Manchester code leaves timing slots of only 50 ns in length compared with 170 ns in the experimental Ethernet. After subtracting fixed timing margins for transmit and receive threshold variations and for cable-length-dependent charging effects, the time left to detect signals and collisions amid echoes and noise is actually much less than implied by the factor of 3.4. The range of signal levels that must be accomodated by a receiver must be greater than in the experimental Ethernet because cable attenuation increases with frequency, and wave forms will be more distorted because velocity dispersion across the bandwidth from D.C. to 20 Mhz is more pronounced. The standard provides three counter-measures to these effects: standard tap spacings and coaxial cable lengths to control echo buildup, limiting length of individual cable segments to 500 meters (half the experimental Ethernet limit) to reduce charging and attentuation effects, and specification of a very high quality coaxial cable, to minimize dispersion. These three measures certainly should help, but there must still be a net, unknown increase in engineering difficulty.

b) Because repeaters are used to filter out echoes and noise, the maximum length span incease has as its primary effect increasing propagation delays, which increases the exposure of packets to collision, discussed in point d) below. It has a minor effect on maintenance and operations, in that a problem detected in the net at one point may actually be caused by a trouble requiring diagnosis on oscilloscope attached to a point 2.5 km away. Generally, by calling for a tightly-coupled network that can span a region 2.5 km long by 2.0 km wide, one would expect that trouble isolation will require more travel than in the experimental Ethernet, where tightly coupled regions were limited to a stripe 1 km long by 100 meters wide.

c) The large increase in maximum number of stations, surprisingly, appears to have the smallest engineering impact, because of the use of repeaters between groups of no more than 200 stations. The primary effect seems to be the minor one of increased effort of trouble isolation when components are four times as numerous.

d) The maximum number of bits at the front of a message that were exposed to collision (because of the time it might take for the first bit to propagate to the most distant station) was about 15 in the experimental Ethernet. In the proposed standard, use of repeaters, increased data rate, and increased maximum length span multiply their effects, so this number increases to about 450 bits. Since one might anticipate that collision frequencies increase combinatorially or exponentially (rather than linearly) with such parameters as bit exposure, number of stations, and traffic load, there seems to be some chance that the negligible collision rate of the experimental Ethernet could transform to a performance-limiting problem in the standard, especially if early applications include many short packets, which amplify the effect. As a specific example, suppose we have a configuration that allows an average internode distance of 2000 meters to realized and the Ethernet is utilized only for minimum length packets, as in, for example, typed character-at-a-time input. This Ethernet will act as a pure Aloha system, with a maximum carrying capacity of only 18% of the nominal channel data rate, or about 1.8 Mbit/sec. This Aloha effect would occur when the presented load neared 3500 500-bit packets per second, which would occur if the 1000 stations originate an average of only 3.5 packets per second each. While the particular configuration and application assumed in that example may be a little exotic or even pathological, it illustrates that the standard Ethernet has the potential to operate in modes substantially different from the experimental Ethernet, with perhaps uninvestigated side effects.

e) The experimental Ethernet allows no branching, but the standard allows branching in an unrooted tree configuration. Branching must make trouble isolation a more tedious job.

f) The experimental Ethernet was a single piece of passive cable. The standard allows repeaters to extend the area coverage. Repeaters, being active, add a reliability and trouble isolation factor to a previously very simple system. They also contribute a lot to the propagation delay that causes point d), above, to be so significant. Finally, each repeater swallows eight of the bits of the preamble, so the presence of repeaters leads to a need for a long preamble.

g) The experimental Ethernet had no constraints on cable lengths between splices or on tap position, while the standard recommends specific tap spacing and cable splice intervals. These recommendations add to the installation and maintenance burden. One may not be able simply to push up the false ceiling outside an office to install a tap at the most convenient place the cable passes by: the nearest unused tap point may be 5 meters away, behind an HVAC duct.

The real concern about these changes of scale and complexity is that they all lead in the same direction: increased difficulty of engineering and operation. There seems little doubt that any one of the changes is technically feasible by itself, but the cumulative impact of all simultaneously seems to take the designer to a completely new and untested arena containing questions about bit error rate, collision rate, throughput, reliability, availability, maintainability, and ease of trouble isolation and administration. (Answering these questions would be an engaging research project, but the standard does not suggest that research is required to build an Ethernet.)

It would be much more reassuring if there were actually running somewhere an Ethernet that stretches all the scale factors to the limit simultaneously. Without such a demonstration system, there must be some doubt as to whether or not it is technically possible to build a network that can meet the standard, and whether or not the standard Ethernet preserves the attractive properties of the experimental Ethernet.

Suggestion: Reduce ambitions by lowering data rate to 5 Mbits/sec., which would not interfere with many applications. A distance limit of 1.5 km would cover most buildings, and might then be achievable without repeaters. Finally, eliminate branching. The other complexity-increasing points diminish on their own if these three suggestions are followed.

2. *Ethernets may be hazardous*

The specification calls for the shield of the coaxial cable to be *ungrounded*, apparently to insure that there are not *multiple* grounds, which can make it difficult to protect sensitive electronics from lightning-induced transients. Unfortunately, ungrounded coaxial cable is a well-known safety hazard, because of the possibility that a power line conductor might fall across the cable, wear through or damage the sheath, and contact the shield. Such an accidental contact raises the potential of the coaxial shield, and in the case of the Ethernet, the connectors and part of the electronics at every station tap. If a person in contact with an earth ground now touches a coaxial connector anywhere in the Ethernet (perhaps 500 meters away from the actual fault) he can obtain a lethal shock. (The particular operation of connecting two pieces of coax with a barrel connector is especially hazardous, since one normally grasps one connector in each hand firmly, and if a potential is present the hands will involuntarily clamp around the connectors and the flexible coaxial cable will make it difficult to let go.)

The Ethernet appears to slip between the provisions of the 1981 National Electrical Code[2], so one cannot find a requirement demanding that its coaxial shield be grounded. (The closest requirements are for CATV and radio antenna coax, which must be grounded if they enter or leave a building or could be exposed to accidental crossing with conductors carrying 300 volts or more.) However, it seems to be in the spirit of the electrical code to require such systems as the Ethernet to be grounded or to include a sensor that grounds the cable if its potential rises.

Suggestion: One possible approach is to require that the shields of adjacent pieces of cable be bonded to each other by a wire that is permanently attached at installation time and that does not disconnect when the cable connector is detached. In addition, this shield interconnect system should be grounded (or protected by a device that provides a ground when it detects a potential) at exactly one point. Whether or not that specific suggestion is followed, this area of the Ethernet specification should be reviewed by an electric safety specialist, and a strategy that meets the grounding requirements of both the Ethernet signalling system and personnel safety should be devised.

3. *An internet interconnection plan is needed*

The standard supplies two facts about a network addressing plan: addresses are to be 48 bits in length, and each station is to be able to respond to some unspecified number of multicast addresses in addition to its own assigned address and messages that are labelled as "broadcast". These two ideas on the one hand imply, and on the other hand may constrain, a plan for interconnecting several local networks into a larger internetwork arrangement. There are several cases where an internetwork plan is essential:

- a single site spans a distance of more than 2.5 Km. or has more than 1000 nodes.
- a single organization has two or more sites separated by a public carrier.
- some other network technology must be used among some set of computers, perhaps because they have a special requirement that the Ethernet cannot handle (e.g., a 100 Mbit/sec. data rate) or their manufacturer has chosen an alternative technology for interconnect.

The standard approach so far used both within Xerox and outside is to install gateways, which are forwarding nodes that have attachments to two or more networks. The Ethernet addressing plan raises several interesting questions about how these gateways should be expected to operate:

- when a gateway receives a packet with a 48-bit unstructured destination, how is it expected to determine where to forward it? It seems likely that a catalog of all in-use 48-bit addresses would be a burden to maintain.
- is it intended that the broadcast feature be propagated by gateways? If so, some scheme is needed to guarantee that cycles of local networks do not rebroadcast such packets forever.
- is it intended that the multicast feature be propagated by gateways? If so, what tables are implied for the gateways?

The point here is that there is a potential interaction between the internetwork interconnection plan and the addressing features of the local network itself; one should have some specific, workable internetwork plan in mind when specifying the particular features of the local network. It is not clear, for the case of the standard Ethernet what internetwork plan the proposers have in mind. (Since Xerox has a lot of experience in internetwork interconnection with the experimental Ethernet, it seems unlikely that no thought was given to this topic.)

Suggestion: An internetwork plan compatible with the standard Ethernet should also be proposed, for discussion as to its features and practicality. Only then can one draw firm conclusions about some of the features of the standard Ethernet.

4. *Need for a technology-independent compatibility interface.*

The standard Ethernet specifies two low-level compatibility interfaces, one on the coaxial wire and a second between the transceiver and the rest of the system. Both of these interfaces are, of course, very dependent on the nature of the Ethernet itself; one can, for example, see collisions across both interfaces. It is disappointing not to find a higher level compatibility interface also proposed, one that utilizes, for example, the same packet formats and addressing structure, but that hides the particular network technology and speed. Such a standard interface could vastly increase the potential inoperability of the standard, because any manufacturer who for some reason decided to adopt a different hardware technology but did so with this interface would still be able to accomplish a hardware interconnect with an Ethernet, too. Software compatibility is another matter, perhaps too hard to standardize at this instant, but the potential effect of hardware interconnectability alone should not be underestimated: many organizations can undertake software projects, while few are prepared to implement a special piece of hardware.

Suggestion: Investigate the feasibility of defining an X.21bis or RS-449-like interface with a packet buffer and complete network control between that interface and the Ethernet. This kind of technology-independent interface would not only allow a wider range of equipment to interconnect to an Ethernet, but it might also encourage more widespread use of those standard interfaces.

Minor observations

1. Although some attention seems to have been given to identifying an Underwriter's Laboratories listed cable insulation material so that the Ethernet can be strung through false ceilings used as air return plenums, there is no mention of a corresponding need for approval of the active electronics of the transceiver, which must be located within 5 cm. of the coaxial cable, and therefore will probably be also placed above the same false ceiling. Since the transceiver is a potential source of fire and smoke, a careful electrical inspector would probably object to such an installation unless the transceiver is UL-listed for the purpose.

2. Although collisions are detected and the backoff-retry algorithm is carefully specified, there is no specified requirement for monitoring the number of collisions. An essential part of any automatic recovery strategy should always be a logging-reporting mechanism, to insure that troubles get fixed, rather than ignored. If automatically recovered troubles are ignored, their causes will accumulate until they overwhelm the recovery mechanism, at which point the system fails, perhaps in a way more mysterious, troublesome, and difficult to repair than if the recovery mechanism hadn't been provided in the first place. For the case of collisions, of course some residual level of collision is expected; what is important is to detect that the collision frequency has climbed above the normal residual level, indicating perhaps that some station's transceiver is getting flaky or whatever.
A similar comment applies to the mechanisms that calculate cyclic redundancy checks and that discard runt packets.

3. The standard mandates a protocol type field, but offers no rationale or explanation of its intended use. It seems that unless there is a standard intention for its use, there is little point in providing a standard field.

4. The standard calls for a station to respond to some unspecified number of multicast addresses. This vagueness makes it difficult for a manufacturer to decide how many to implement, which in turn will lead to different decisions, which in turn may make the multicast feature unusable. It would seem some guidance on exactly what to implement would be a good idea.

5. One must analyze the standard like a Philadelphia lawyer to discover that the configuration flexibility implied by the specifications is not all available simultaneously, and that one must exercise some care in planning. For example, the 2.5 meter minimum tap spacing interacts with the maximum 50-meter transceiver-to-station distance. One might imagine that it is reasonable to run an Ethernet cable down just the center of a building wing, if it happens that all offices in the wing are within 50 meters of the center. However, it then might be impossible to place a cluster of ten stations in a single room that is at the building edge, because not enough tap points are contained within the section of the Ethernet cable that lies within 50 meters of the cluster room. Of course one might anticipate this problem by leaving coils of coax in some locations, but that approach shortens the overall geographical span of the Ethernet. (Perhaps a set of configuration planning guide lines could be developed that point out possible traps such as this one.)

Footnotes

1. Available from Digital Equipment Corporation at 1925 Andover Street, Tewksbury, MA., 01876, Intel Corporation, 3065 Bowers Avenue, Santa Clara, CA., 95051, and Xerox Corporation, 3333 Coyote Hill Road, Palo Alto, CA., 94304.

2. Ross, J.A., and Summers, W.I., *The National Electrical Code Handbook*, 2nd Edition, National Fire Protection Association, Boston, MA., 1980.