

March 11, 1980

Environment Considerations for Campus-Wide Networks

by Jerome H. Saltzer

"The Campus Environment" is a name proposed here to identify a particular set of physical properties, geographical extents, data communication requirements, administrative relationships, and needs for flexibility that characterize our university campus. With only minor exceptions they equally apply to a corporate site, a government complex, or another university. This note discusses seven characteristic properties of this campus environment. These seven properties provide a basis for design decisions for a data communication network to span a campus. As will be seen, the properties of this environment are quite different from those of a single building, or of a nation-wide, common-carrier-based network.

Seven Properties of the Campus Environment

- 1) It has a geographical extent beyond a single building, but within a single political and administrative boundary that permits transmission media to be installed without resort to a common carrier.

This first property is essential, so as to allow exploitation of low-cost, high-bandwidth communication technology. With current technology and prices the difference in costs between communicating over privately installed equipment and using common carrier facilities can be a factor between 10 and 100.

This note is an informal working paper of the M.I.T. Laboratory for Computer Science. It should not be reproduced without the author's permission, and it should not be cited in other publications.

- 2) Within this geographical area, a large number of nodes--that is, computers, data sources, and data sinks--require interconnection. Today the number of such nodes may be in the range of ten to one hundred. Looking ahead to the advent of desktop computers, one may be faced with from a few hundred to several thousand nodes by the end of the next decade.

The combination of the previous two properties seems to make it inevitable that local interconnect technologies such as the ETHERNET, CHAOSNET, L.C.S. Ring net, HYPERCHANNEL, or MITRENET cannot by themselves completely accomplish the required interconnection, since all such technologies that have so far been demonstrated have limitations on distance on the order of a thousand meters and limitations on node count on the order of a hundred nodes. Thus one would expect to use those technologies to attach clusters of nodes into subnetworks, for example all the nodes in a single building, and then install interconnections (gateways) among these subnetworks. For our own campus, one might envision by 1990 as many as 100 subnetworks each comprising an average of, say, 100 nodes. Subnetworks and gateways introduce the problem of how to route a message from a source node through a series of subnetworks and gateways, so that it ends up at a desired target node.

- 3) Administratively, there exist forces both for commonality and for diversity of network attachment strategies. The primary force for commonality is a desire to be able easily to set up communications between any pair of nodes on the campus. The primary force for diversity is that the choice of a computer, data source, or data sink typically pre-determines the technology of the network to which it must be attached, because off-the-shelf network hardware for that node may be

available in only one technology. Further, some applications may have special requirements for some connections (e.g., high bandwidth) that can be met only with a particular network supplier's equipment, yet still need occasional "ordinary" connections to nodes elsewhere. Thus the emerging diversity of local networks will continue, and probably increase, rather than decrease, with time.

- 4) The worldwide academic, commercial, and regulatory community has not yet reached anything resembling a consensus on how networks should be organized, how protocols should be layered or how functions should be divided. Arguments range over issues ranging from obscure matters of taste, through fundamental technical disagreements about which requirements should have priority in design, to alternative opinions of the directions that communication technology is moving. Many different and competing standards have been proposed, and one can find in the literature a good technical case against any one of them. One must anticipate that these arguments will be reflected internally in the campus environment, in the form of a diversity of protocols and standards, and particularly in the requirement that any mutually consenting set of nodes be able to carry on communication with one another using a protocol that no one else has ever heard of, much less agreed to.*

This fourth requirement suggests strongly that any network interconnection strategy that must be implemented today should have a campus-wide lowest layer of protocol that accomplishes datagram passing between any two nodes while

* Imagery borrowed from a Chaosnet working paper by David Moon.

making an absolute minimum number of assumptions about the nature of the higher-level communications that are taking place or the policy of network administration. Some typical assumptions that should be avoided unless an unusual opportunity is obvious are: what level of reliability/delay tradeoff is appropriate; how routing should be optimized; fragmentation/reassembly strategy; flow control requirements; addressing plan; and particular network topology.

- 5) Because a data communication network is a campus-wide service, there will be no single user or user group with a wide-enough interest to administer the entire network. This means that network administration will either be done by a haphazard confederation of special interest groups or else by a chronically underfunded central service organization modeled on the one whose role is to minimize telephone costs.

In either case, this property places a requirement on the network interconnection technology that it be robust and self-surviving to every extent imaginable. Trouble isolation must be easy to accomplish and easy for individual users to participate in if they are so inclined, because trouble isolation and repair may involve multiple administrations. Simplicity of operation of gateways is important, so that operation can be completely unattended for long stretches of time. A network design approach that requires close monitoring is undesirable.

- 6) The topology of subnetwork interconnection will be administered partly with central planning and partly without. This property arises from two needs: First, a "dependable" set of gateways that one can expect to exhibit predictable and stable properties is an essential backbone to a

useful service. A centrally planned and administered set of gateways would provide this dependability. Second, whenever a node finds that for some reason it is attached to two subnetworks, it may find that it is useful in some of its applications to serve also as a gateway between the subnetworks; yet it may not want to take on the official responsibility of being a publicly available gateway. Another example of a gateway that is not centrally administered may arise if some particular application needs, and has purchased the gateway equipment to support, a path through the network with special properties of delay, reliability, bandwidth, or privacy. The person or organization that has purchased the special gateway equipment may not be prepared or willing to allow public use of it. Alternatively, a user may wish to avoid use of a sometimes troublesome gateway that is claimed by its owner to be perfectly operating.

- 7) External networks such as TELENET, the ARPANET, TYMNET, XTEN, SBS, or A.C.S., may be attached to some nodes, and some of those nodes will serve as gateways between the campus network and the external networks. In some cases, the external network will be used simply as a "long link" in the campus net. In other cases, facilities within the campus net will set up communication paths to services having no other connection with or knowledge of the campus net. Both kinds of cases require careful consideration of the interactions between internal and external network properties.

Note that the campus environment has all these properties only if we assume the technological opportunity mentioned in point one: that low-cost hardware and media can provide communication paths in the range from 1 to 10 Mbits/sec.

between any two points within the campus. Availability of interconnect media and subnetworks with this bandwidth has been demonstrated in several forms. Gateways that operate with such bandwidths may be harder to construct, and that concern is one of the considerations involved in developing a campus-wide net. Individual nodes that can sustain these data rates for very long are likely to be rare; software often limits the rate at which a node can act as either a data source or data sink. Instead, the high bandwidth technology is to be exploited in two ways:

- 1) to provide enough capacity to handle the aggregate demand of many lower-bandwidth sources and sinks of data.
- 2) non-optimal strategies that are relatively simple to implement or administer can be considered; it is not a requirement that every bit of the available bandwidth be optimally utilized.

The availability of high bandwidth, together with lack of a requirement to use that bandwidth efficiently, is probably the most fundamental technical difference between the "campus-wide network" and the commercial long-haul data communication network, a difference that can lead to significantly different design decisions. Future notes in this series will explore some of these specific technical design consequences.