

AD-A099 598

HAWAII UNIV AT MANOA HONOLULU DEPT OF ELECTRICAL ENG--ETC F/6 17/2
DESIGN ISSUES FOR HIGH SPEED LOCAL NETWORK PROTOCOLS, (U)

MAY 81 F F KUO

N00014-76-C-0498

UNCLASSIFIED

CCIS-11

NL

[]
[]
[]

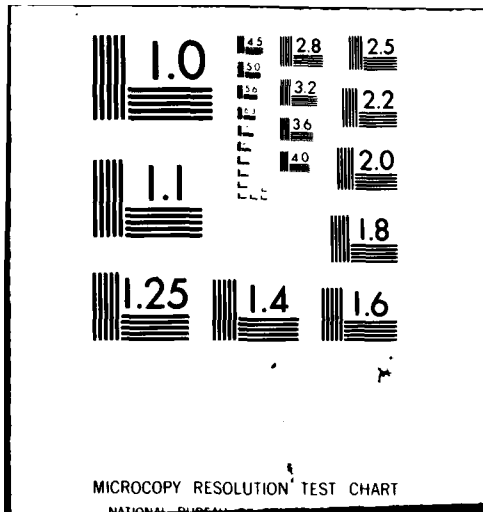


END

DATE
FILMED

6-81

DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

LEVEL II

12

AD A 099 598

DESIGN ISSUES FOR HIGH SPEED LOCAL
NETWORK PROTOCOLS*

by

DTIC
ELECTE
JUN 2 1981
S D
E

Franklin F. Kuo
Department of Electrical Engineering
University of Hawaii at Manoa
Honolulu, Hawaii 96822

EXHIBIT A
Approved for public release;
Distribution Unlimited

*This work was partially supported by the U.S. National Bureau of Standards
and by the Office of Naval Research under contract N00014-78-C-0498.✓

DTIC FILE COPY

81 6 01 185

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 14 CCIS-11	2. GOVT ACCESSION NO. AD-A099	3. RECIPIENT'S CATALOG NUMBER 598
4. TITLE (and Subtitle) 6 Design Issues for High Speed Local Network Protocols		5. TYPE OF REPORT & PERIOD COVERED
7. AUTHOR(s) 10 Franklin F. Kuo		6. PERFORMING ORG. REPORT NUMBER
8. PERFORMING ORGANIZATION NAME AND ADDRESS Dept. of Electrical Engineering University of Hawaii at Manoa Honolulu, Hawaii 96822		9. CONTRACT OR GRANT NUMBER(s) N00014-C-78-0498
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Code 437 Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 15 N00014-78-C-0498
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 12 11		12. REPORT DATE May 1981
		13. NUMBER OF PAGES 09
		15. SECURITY CLASS. (of this report) Unclassified
		18a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Local Area Networks Local Network Protocols High-Level Protocols		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) With recent advances in optical fiber technology, it is now feasible to consider data communication systems with speeds up to 1000 Mb/s. Because of their high bandwidth, low delay, and low error rate characteristics, optical fiber communication systems seem to provide an almost ideal transmission medium for high speed local networks. This paper deals with the implications of the use of fiber optics in the design of high speed local network protocols. 7		

444 637

[Handwritten signature]

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist and/or	
Special	
	A

DESIGN ISSUES FOR HIGH SPEED LOCAL NETWORK PROTOCOLS

Franklin F. Kuo

University of Hawaii at Manoa

With recent advances in optical fiber technology, it is now feasible to consider data communication systems with speeds up to 1000 Mb/s. Because of their high bandwidth, low delay, and low error rate characteristics, optical fiber communication systems seem to provide an almost ideal transmission medium for high speed local networks. This paper deals with the implications of the use of fiber optics in the design of high speed local network protocols.

1. TRANSMISSION MEDIA

With recent advances in optical fiber technology, data transmission at speeds up to 1000 Mb/s is now technically feasible. Optical fibers may be used in a variety of applications where twisted copper wire-pairs, coaxial cables, and waveguides are now used for the transmission of data; these applications range from short data links and equipment connections within a building and between buildings, to long trunk circuits between cities. This report deals with the implications of the use of fiber optics for local network applications upon local network communication protocols. Using fiber optics technology, high speed local networks with throughputs of 10 Mb/s to 100 Mb/s are possible. The discussions here will examine protocol implication of local nets with data rates in the range cited above.

Without going into the details of fiber optics technology the transmission medium can be described as having the following characteristics:

- low transmission delays
- wide transmission bandwidth
- low error rates

These characteristics give rise to the possibility of "ideal" transmission medium for local networks--low delay, high throughput, and almost error free (error rates of 1 bit in 10^{20} are commonly cited). What do these ideal characteristics imply from the standpoint of computer protocols? In the following are discussions of some specific issues.

2. PACKET SIZE

It is well known that low delay requires short messages, short queues, and few control messages, while high throughput requires long messages, long queues, and low overhead. Are these performance trade-offs applicable to high speed local networks? According to McQuillan and Cerf (1), delay is defined as "time between transmission of the first bit and delivery of the first bit". Its components are:

- speed of light ~ distance
- transmission delay ~ message size/circuit rate
- processing delay
- queuing delay ~ system load.

On the other hand throughput is defined as "number of bits sent divided by time between transmission of the first bit and delivery of the last bit". The components of throughput are:

- effective bandwidth of processing equipment
- effective bandwidth of transmission media.

For fiber optics local networks the distances are short, the transmission delays are negligible, and the effective bandwidth of the transmission medium is very wide. However, the other components of delay and throughput cited above have nothing to do with the characteristics of the transmission medium, but are related to the inherent processing power of the devices attached to the local network. Since some of these devices can be terminals or microprocessors with little processing power, the throughput/delay trade-offs of a high-speed local network are limited by the processing power of these least capable members of the network. The implications as far as packet size is concerned are these:

- a. Packet sizes should be kept short in order to match the sizes of the buffers of the terminals and/or microprocessors on the net.

- b. Because bandwidth is inexpensive and readily available on a local network, there is little motivation to keep the size of the header or overhead bits down on a packet. Thus a highly inefficient packet in ARPANET consisting of a single character or word of data together with multiple words of administrative information would be regarded as acceptable in a high speed local net.

- c. To keep node processing down, a standard header format with fields in fixed locations is recommended. Multiple packet types, optional fields, etc. all tend to increase communications processing at the node. Metcalfe (2), in fact, recommends more spacious fields than one would think necessary since "try as you may, one field or another will always turn out to be too small."

In summary, small packet sizes are recommended for high-speed local networks in order to match the buffer capacities of the microprocessors and terminals attached to the network. However, since bandwidth is inexpensive, larger header overhead can be tolerated.

3. FLOW CONTROL

Because of the wide bandwidth available, congestion control on the transmission medium is not necessary. Thus, link level flow control is not a major requirement for high speed local networks. However, buffer management for the various hosts attached to the network is still an important issue since different hosts generate and absorb data at varying rates. Thus, a transport-level flow control mechanism is needed. Since flow control is closely tied to buffer availability, the amount of buffer space needed for efficient operation under different circumstances is an important factor in protocol performance.

A simple dynamic flow control mechanism that does not function efficiently in conventional long-haul packet switching networks might work well in a high speed local network. This mechanism is based upon a simple start-stop command. The receiver issues start and stop commands that place the sender in a transmission-allowed or transmission-blocked state. Since in high-speed local nets, these commands suffer little delay in transit, the flow control mechanism that would require somewhat more processing power is based upon granting "credits" for transmission. The receiver grants credits for a certain amount of data to the sender so that both know exactly how much data will be exchanged. The number of credits provided by the receiver is frequently

called the "window size" when credits are expressed relative to packet sequence numbers. In a high speed local network where control commands can be transmitted with very little delay, the start-stop flow control scheme seems to be preferable because of its simplicity to the "credit" flow-control scheme.

4. ERROR CONTROL

Since the transmission medium is almost error free, it means that link-level error control is unimportant. Since the link-level error control feature is found in almost every data link control protocol in use today (HDLC, SDLC, ADCCP, BISYNC, etc.), it raises the question whether to adopt one of the present data link control procedures unchanged and apply link-level error control, even when it is not needed, or to adapt or develop another version of the data link control procedures to better suit the characteristics of the transmission media.

It is clear, however, that end-to-end error control is needed in a high speed local network. This is because of the fact that the devices (computers, peripherals, and terminals) are inherently more error-prone than the fiber optics transmission medium. This means that error control at the transport protocol level is needed. As to the best way to implement the transport level error control, it appears that error detection with retransmission (ARQ) is probably preferable to forward error correction. This is because low transmission delays will produce fast acknowledgements, so the efficiency of an ARQ scheme is much greater in a high speed local network than for conventional transmission media. Moreover, for a local network many of the devices attached to it will not have much processing power, so the requirement of forward error correction might place an undue burden upon the device

5. DATAGRAM VS. VIRTUAL CIRCUIT PROTOCOL

In view of the fact that a local network might have considerably more simple hosts such as intelligent terminals and microprocessors than long haul networks, the low-level transmission protocols should be kept at a level that the simple hosts can handle. This argues strongly for a datagram protocol to be the basic transmission protocol. That datagram protocol would provide the basic responsibility of delivering a single addressed packet to one or more of its destinations. Above the datagram layer, however, should exist a virtual circuit layer for those hosts capable of supporting a transport station which can multiplex a number of virtual circuits.

6. INTERNETTING CONSIDERATIONS

In order to allow for growth and evolution, local networks should be designed to allow interconnection with other local networks and long-haul networks. For other local networks, the concept of a bridge is useful. As described by Clark et al.

(3) a bridge contains

"two network interfaces, one appropriate to each of the subnetworks it interconnects, a limited amount of packet buffer memory, and a control element, which implements an appropriate filter function to decide which messages to "pull off" and buffer until it has an opportunity to retransmit it to the other subnetwork."

If there is substantial speed disparity between the two local networks, the bridge must have either extensive buffer capacity or the ability to regulate the flow of information from the higher speed network. The long distance bridge concept which Clark et al. describes might be adapted for this purpose (3).

For interconnection with a long-haul network, a gateway is required. The simplest gateway would be for the case where the long-haul network offered a datagram interface to the local network. However, for most commercial long-haul networks the interface offered is a virtual circuit (e.g. X.25). Therefore, a virtual circuit protocol should be implemented on the local network which is as close as possible to the long-haul virtual circuit model. Since the local network probably will not require the full range of functions available on a long-haul virtual circuit model, a subset of the long-haul virtual circuit protocol could be implemented for the local network. However, according to Clark et al. (3), the compatibility of the virtual circuit protocols between long-haul and local area network does not answer the question of how the features as flow control, buffering and speed matching should be implemented. Standards should be proposed for gateways between local networks and X.25-based long-haul networks.

7. ADDRESSING AND ROUTING

For high speed local nets, routing is not an important issue. Since delay is so low, optimal routes are not significantly different from suboptimal routes. Routing is only important at a gateway or bridge, and internet packets must be sent to the correct gateway or bridge for forwarding.

In local area datagram networks, message exchanging between two cooperating entities such as ports requires that each entity

knows the network address of the other. Sometimes entities are known by their names rather than their network addresses. In such cases, it is advantageous "to maintain, as a network service, a facility which will take the name of a desired entity and give back its network address" (3). In a high-speed local network with many simple hosts, the network directory service can perform very effectively since the delay between queries and responses can be quite low due to the high speed transmission media. In general, central network services become more efficient when they can be accessed through wideband, high speed transmission media. The only efficiency constraint then becomes their own processing power in dealing with queries and responses.

8. HIGH-LEVEL PROTOCOLS

High level protocols are those protocols primarily concerned with performing remote operations across a network. The low-level virtual circuit and datagram protocols discussed in previous sections are "communications protocols" whereas the high-level protocols discussed here are "resource sharing" protocols. In terms of the ISO Open System Interconnection Architecture Model (ISO/TC97/SC16), the high-level protocols discussed here are at the session or presentation level (4). We will concentrate on three high level protocols: terminal, file transfer, and remote job entry protocols, which provide basic services for the users of a local network. Terminal protocols establish mechanisms that allow efficient and flexible terminal access to networks. Terminal protocols not only allow a user to access a time-sharing service through the local network, but can also be used as a character-oriented network interprocess communication facility. File transfer protocols allow users to manipulate remote file systems and to transfer files from one host system to another. Remote job entry protocols provide users with a mechanism for submitting jobs to various batch services on a network. Many of the problems encountered in these protocols recur in more complex forms in more sophisticated protocols (e.g. network mail protocols, distributed data-base protocols) which may be built on top of them.

In this discussion we are primarily concerned with heterogeneous network protocols, i.e. those protocols which deal with networks of heterogeneous computers, terminals, peripherals, etc. The common problem that all three high level protocols share is that they require substantial network software effort to implement. Moreover, in long-haul networks each host or terminal offering the remote service is required to have a copy of the high level protocol software within its own physical memory space. Since

a heterogeneous computer network can have many different varieties of computers and terminals, the task of programming each of the devices for high level protocols is expensive and time-consuming. For long-haul networks, it is not feasible to have a central facility which acts as a mediator between two entities wishing to engage in a cooperative task, such as remote job entry. This is because of the excessive time-delays encountered in having such a central facility. However, in high speed local networks, where time delays are very short, such central facilities are feasible and even advantageous. Not only do they provide a centrally supported software facility, but they also reduce the requirements that each user terminal or host retain a full copy of the high level protocol within its own memory space. The use of the central high level protocol facility for each specific protocol will be discussed in subsequent sections. However, in each case, it can be described as shared facility, somewhat like a re-entrant compiler on a time-shared computer. Note that a re-entrant compiler can be shared among a number of users without the necessity of each user having a copy of the compiler in his own working space. Similarly, a central high level protocol facility can be shared by a number of processes without each possessing a complete copy of the protocol in its own file space.

9. TERMINAL PROTOCOLS

In order to accommodate the wide variety of different kinds of terminals a virtual terminal protocol (VTP) is commonly used in which a network virtual terminal (NVT) is defined as the network standard. The terminal side of a connection maps the output of its terminal into the NVT format for transmission to the host. The host then maps the NVT format into its local form. Each host of the network then only needs to support one terminal type (the NVT). In order to allow not only terminal to process communication, but also process-process and terminal-terminal interactions, the NVT software should reside in both sides of a connection, thus leading to a symmetrical view. Such VTP schemes are expensive in that network resources must be dedicated at every host and terminal to support the NVTs. For a high-speed local network environment, a centralized NVT is feasible so that each side can access a single virtual data structure that performs the functions of data translation, option negotiation, echoing, and interrupt signaling. Such centralized NVTs were not previously possible in long-haul networks because of response-time limitations. However, in high speed local networks, their use can be highly advantageous, not only in conserving memory space at each node, but also because of software support for a single facility is more efficient than for multiple computer and terminal-types.

10. FILE TRANSFER PROTOCOLS

A File Transfer Protocol (FTP) defines the set of rules for the transfer of files from the file system on one host to the file system on another. In heterogeneous networks the purpose of an FTP is to establish a network virtual file system (NVFS) which allows a process in a local host to access data stored on any remote host as if the data were stored locally. To accomplish this, canonical or virtual file formats are defined in an FTP so that the source will map the file into the proper virtual format and the destination will map the received virtual form into the proper local format.

Other very important functions of the FTP are maintaining access control and directory information across a network. Access controls are required for reasons of security and privacy. Directory information is needed for efficient file management.

Recent FTP designs partition the FTP into two separate protocols: a data transfer portion and a file management portion. This partitioning allows the FTP to be partially off-loaded onto a network front-end, so that the front-end performs the data-transfer tasks and the host performs the file management tasks. In a high speed local network, a central FTP facility which performs the front-end data transfer tasks for all hosts can be quite advantageous. In addition to performing data translations from a local host format to the NVFS format, the central facility can also act as a security filter for access control and maintain current directories of important network files. The central facility will also relieve a host from some (but not all) of its data and file conversion tasks of going from its local form to the network canonical form, but more importantly will relieve the host of all network access control burdens. It should be emphasized that the efficiency of the central FTP facility is strongly dependent upon how it is implemented, and not simply dependent upon the high-throughput character of the local network. However, because the wideband transmission available, the central FTP concept should be carefully considered in future local network architectural designs.

11. REMOTE JOB ENTRY PROTOCOLS

Remote job entry (RJE) protocols are, after remote terminal protocols, the most important network requirement in the current data-processing environment. Remote Job Entry Protocols (RJEP) provide the network user with the flexibility to use a single implementation of RJE software with a variety of batch systems. The RJE protocol sits halfway between the file transfer protocol (FTP) where high throughput is a requirement, and the virtual

terminal protocol (VTP) where low delay is a primary goal. As such it could draw upon facilities of each protocol without having to develop a completely unique set of functions. For economy reasons, the RJE protocol could be implemented easily if it is co-located with the central FTP and VTP facilities. We will not dwell on this point further.

12. CONCLUSIONS AND ACKNOWLEDGEMENTS

In this paper we have presented a number of issues concerning the design of protocols for high speed local networks. With increasing attention being paid to the application of fiber optics links to local networks, the paper addresses the key question of whether conventional communication protocols are appropriate for transmission media with speeds of 10 to 100 Mb/s.

This work was partially supported by the U.S. National Bureau of Standards and by the Office of Naval Research under contract N00014-78-C-0498.

13. REFERENCES

- (1) McQuillan, J.M. and Cerf, V.G., *A Practical View of Computer Communications Protocols*, IEEE Computer Society, 1978, pp. 4-5.
- (2) Metcalfe, R.M. and Boggs, D.R., *Ethernet: Distributed Packet Switching for Local Computer Networks*, Comm. ACM, July 1976, pp. 395-404.
- (3) Clark, D.D., Pograd, K.T. and Reed, D.P., *An Introduction to Local Area Networks*, Proc. IEEE, November 1978, pp. 1497-1517.
- (4) International Standards Organization (ISO), *Reference Model of Open Systems Interconnection Architecture*, Version 3, November 1978.