TELEMATICS and INFORMATICS
Vol. 5, No. 1, pp. 65–73, 1988

0736-5853/88 $3.00 + .00

SERVICE ACCESS PROCEDURE (SAP) FOR A TRANSPORT LAYER PROTOCOL

Asok Ray
Shashi Phoha

Abstract—The article presents the concept and design of a Service Access Procedure (SAP) for establishing transparent and reliable connections between the top three layer protocols of the ISO Open System Interconnection (OSI) and the corresponding bottom four layer protocols. The SAP allows a user to utilize network authentication services and accommodates for a single network login while establishing an active connection. Of the bottom four layers, the transport and network protocols are often selected to be the U. S. Department of Defense (DoD) standard Transmission Control Protocol (TCP) and Internet Protocol (IP), respectively. The upper layer protocols are usually a part of the host processor's software whereas the TCP/IP and the subsequent lower layer protocols might be resident in a physically distinct device such as an interface unit (IU). The rationale for residency of the TCP/IP in IUs, the concept and design of a TCP SAP, and the basis for selection of the lower layer protocols that constitute the host-front-end protocol set are presented in this article.

In a distributed data communication network, heterogeneous computers and intelligent terminals are usually interconnected via local area and wide area networks. The ISO Open System Interconnection (OSI) reference model (Day & Zimmerman, 1983) provides a layered protocol architecture for modularity of both software and hardware, and has been largely accepted as the standard in both commercial and military applications. A functional description of the seven layers of the ISO OSI reference model and their significance relative to commercially available network architectures like Systems Network Architecture (SNA) (IBM, 1982) and Digital Network Architecture (Digital Equipment Corporation, 1983) are provided by Schwartz (1987). Examples of network architectures which have been developed on the basis of the ISO OSI reference model are the Manufacturing Automation Protocol (Society of Manufacturing Engineers, 1986), the Technical and Office Protocols (TOP) (Boeing Company, 1985), and WWMCCS Information System (WIS) Protocols (Bernosky & Blankertz, 1983).

For long-haul communications, the transport and network layers (i.e., the fourth and third layers within the seven-layered ISO OSI reference model) are critical for transparent and reliable connections between peer level upper layer protocols (i.e., the session, presentation, and application layers). The MIL-STD-1778 Transmission Control Protocol (TCP) (U.S. Department of Defense [DoD], 1983a) and the MIL-STD-1777 Internet Protocol (IP) (DoD, 1983b) are the DoD standards for the transport and

Dr. Ray is an Associate Professor in the Mechanical Engineering Department at The Pennsylvania State University, University Park, PA 16802. Dr. Phoha is a Senior Scientist at ITT Defense Communication Division, 492 River Road, Nutley, NJ 07110.

network layers, respectively. The U.S. Department of Defense has decided to accept, as a standard, the ISO Transport Protocol (TP) (Information Processing Systems, 1984; Groenbaek, 1986) which is more versatile than the TCP, and is compatible with the IP. However, recent trends in the networking market (Petrosky, 1987) indicate that the TCP/IP is being adopted or evaluated for adoption in wide area networking by large corporations. The rationale for using the TCP/IP as an end-to-end service protocol is that it provides a common base for interconnecting heterogeneous computers to transfer data files (DoD, 1984a) and mail messages (DoD, 1984b) as well as interfacing between terminal devices (DoD, 1984c) without encountering serious equipment incompatibility problems.

A variety of protocols could reside within different physical components of a distributed network and serve to establish reliable communication paths between the participating pair of hosts. (All subscribers, including hosts, workstations, and intelligent terminals, are referred to in this article as hosts.) The TCP/IP and their supporting lower layer protocols are responsible for establishing error-free and reliable communication paths between hosts. An important issue is where the end points of a transport connection should be physically located, i.e., should the TCP/IP software be a part of the host processor's operating system, should it be used to provide end-to-end connectivity between the respective interface units (IUs), or should it be placed in gateways to support reliable end-to-end communication across a long-haul network that interconnects the different local area networks within the distributed communication system. So, the alternatives for residence of the TCP/IP residence are (1) Host Processors, (2) Interface Units (IUs), and (3) Gateways. These alternatives were analyzed with respect to the performance, design, and testing requirements of distributed data communication networks. On the basis of a trade-off analysis, the TCP/IP was placed in the IUs.

A Service Access Procedure (SAP) has been developed to provide transparent and reliable connections between the IU-resident TCP/IP and the upper layer protocols (ULPs) that are contained in the host processor. The major contribution of the article is the conceptual development and evaluation of the SAP.

TCP/IP RESIDENCE

The three alternatives for residence of the TCP/IP (mentioned earlier) were analyzed with respect to a number of performance, design, and test considerations. The results of analysis are summarized below.

End-to-end (processor-to-processor) user performance. The end-to-end user performance is dependent upon the residence of specific protocols in the source and destination host processors. In general, the data rate between the host and the IU is at least an order of magnitude smaller than that between the IU and the medium; and the protocol overhead and the processing burden of the TCP/IP are significantly larger than that of point-to-point connections required in the host-front-end protocols. Therefore, if the TCP/IP is resident in the IU as opposed to the host processor, the end-to-end performance such as data latency and throughput will be improved. On the other hand, if the TCP/IP is resident in the gateways, end-to-end intranet connectivity between local hosts may not exist and gateways would have to be designed for a larger capacity to maintain the network performance. This option does, however, have the advantage of potentially higher throughput and smaller delays for local connections due to the reduced protocol overhead.

Availability of network services. If the TCP/IP is resident in the host processor, their software will be linked with the host processor's operating system. In that case, Mean-Time-To-Repair (MTTR) for the TCP/IP and its accessories is expected to be larger than if they are resident in IUs and gateways. On the other hand, if the TCP/IP are resident in gateways, no end-to-end error recovery services between local hosts would be available and, therefore, the network reliability may suffer.

Impact of changes in protocols upon host processor performance. Modifications in the TCP/IP protocols should have as little bearing as possible on the performance of host processors. In this respect, if the TCP/IP are resident in IUs, then the requirements on the host's operating system to support network functions are reduced. On the other hand, for the TCP/IP's residence in gateways, it is the host processor which will have to determine, during each call set-up, whether the connection is local or remote.

Standardization of processor interfaces. There are many versions of TCP/IP implementations in use. Unless a standard implementation is accepted, there will be potential difficulties with respect to access control, security requirements, and testing and certification. Selection of a standard TCP/IP and the supporting lower layer protocols in each IU eliminates this problem whereas it is difficult to mandate this requirement on differing implementations on individual hosts. When the TCP/IP are resident in gateways, the host processors directly interface with the gateway software and, consequently, testing difficulties are increased.

Based on the above trade-off analysis, it was decided to place TCP/IP in the IU (Ray, Angevine, & Haughney, 1985). The next step was to develop the concept for procedures to provide access to the IU from the respective host.

FEATURES OF THE SERVICE ACCESS PROCEDURE (SAP)

The Service Access Procedure (SAP) should provide a standard, evolutionary method for hosts to access the full range of communication services offered by the network system, thereby preserving their processing power for applications. (A discussion of other SAP developments is given in Lilenkamp, Mandell, & Padlipsky [1984] and Padlipsky [1984]). The SAP is designed to provide a transparent interface between the upper layer protocols (ULPs) within a host and the transport and (possibly) lower-layer protocols resident in the attached IUs, i.e., the ULPs will be unaware that these transport and lower-layer services are not coresident. The SAP is structured to utilize standard lower protocols that provide reliable, ordered transfer of data between a host and its IU.

The SAP also isolates the host processors from software modifications that may be required in the future for planned, evolutionary improvements in the network communication services. Such services may include provisions for a possible migration to other connection-oriented transport layer protocols such as TP (DoD, 1983b) and the capability of a host to access multiple networks through a single interface.

The SAP presented in this article has been primarily designed to provide an interface between host-resident ULPs and the IU-resident TCP, hereinafter referred to as TCP SAP. The TCP SAP allows the host to access all essential TCP services including the security and precedence options. This provides compatibility with standard DoD ULPs including File Transfer Protocol (FTP) (DoD, 1984a), Simple Mail Transfer Protocol

(SMTP) (DoD, 1984b), and TELNET (DoD, 1984c). The TCP SAP should also have the potential for supporting other connection-oriented ULPs that may be developed in the future.

The TCP SAP also permits a user to utilize network authentication and name services while establishing an active connection. A Network Authentication Service (NAS) can be a function of the centralized access control of the network system. NAS allows a user to perform a single network login with his/her identifier and password to establish a network session and to be preauthorized for all ULP connections while the network has other active connections. The Network Name Service (NNS) provides the capability for a user to denote the destination port and internet address as a symbolic name prior to establishing an active connection. The TCP SAP has a symmetrical architecture similar to the ISO OSI reference model as shown in the host-front-end protocol layout in Figure 1. The TCP SAP is implemented either orthogonal to the transport layer or interposed between the transport and session layers in peer modules. The host-resident part is called "Client SAP" and the IU-resident part "Server SAP." The TCP Client and Server SAPs exchange information over a channel level interface performing the functions defined in the lower three layers of the ISO OSI reference model.

The TCP Client SAP, implemented in each host, allows connection-oriented ULPs to access the network services by performing the following functions:

1. Package and format the ULP's requests for TCP services, including the NAS and NNS parameters into TCP SAP messages;
2. Exchange TCP SAP messages with the IU via the channel, level interface; and
3. Provide interfaces to allow the ULP to interact with the TCP and utilize NAS and NNS.

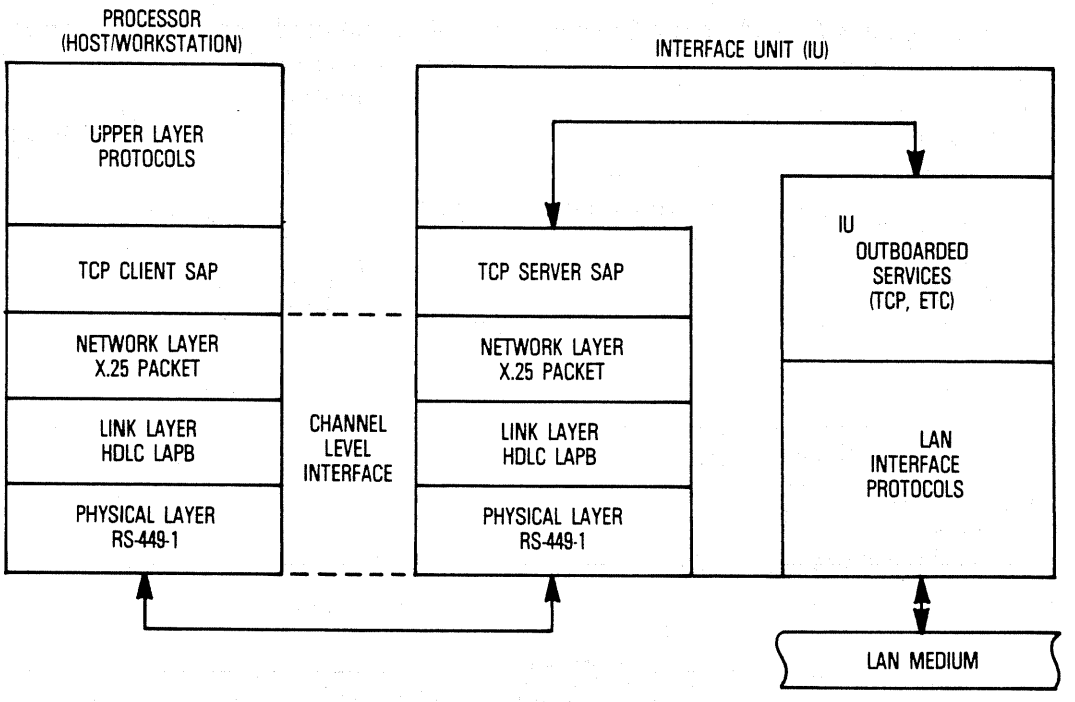

Figure 1. TCP SAP Architecture.

The TCP Server SAP, implemented in each IU, provides an interface between the host and the network services by performing the following functions:

1. Package and format the TCP interaction primitives and NAS parameters into the TCP SAP messages;
2. Exchange TCP SAP messages with the host via the channel level interface; and
3. Provide interfaces to interact with the TCP, NAS, and NNS that are resident in the IU and Network Control Center (NCC).

The initial implementation of the TCP SAP interfaces the hosts and their attached IUs with the CCITT X.25 channel layer (Meijer & Peeters, 1982; Folts, 1980), HDLC LAPB link layer (Meijer & Peeters, 1982; Rybczynski, 1980), and the EIA-RS-449 physical layer (Electronic Industries Association, 1977), as the lower three ISO OSI reference model layers (channel level interface). The TCP SAP messages (header and text) become the information field of the X.25 packets that are exchanged across the channel level interface. The X.25 packet header and HDLC header and trailer, including flags, form the data frame that is exchanged between a host and its IU. The structure of the data frame is shown in Figure 2. The channel level interface functions and their selection criteria are summarized below.

1. Channel Layer: This layer provides multiplexing of connections over a single data link, connection management including flow control, error detection and recovery, and sequenced data delivery. The CCITT X.25 protocol was selected for the channel layer because it provides point-to-point, i.e., host-to-IU, control, has the capability of multiplexing several TCP connections over a single link normally provided by the TCP at the transport layer, and is commercially available.
2. Data Link Layer: This layer serves to directly connect two points, i.e., host to IU, and manage the physical circuit to support information exchange. It also provides for data transparency, packet scheduling, and error detection and recovery. The HDLC based Link Access Procedure Balanced (LAPB) was selected because of its data reliability, compatibility with the X.25 network layer, and commercial availability.
3. Physical Layer: This layer provides the full-duplex, bit-serial, ordered delivery of information between the end points (ports) of the physical connection. EIA-RS449 (Electronic Industries Association, 1977) circuits with the balanced electrical characteristics of MIL-STD-188-114 (DoD, 1985) was selected as the physical layer standard because of its clarity and installation flexibility (cable lengths in excess of 1000 meters are supported).

FUNCTIONAL DESCRIPTION OF THE TCP SAP

The TCP SAP provides full access to the TCP by supporting all essential TCP connection management and data transport services and the Quality of Service features of the IP. There is a one-to-one relationship between the TCP SAP messages and their mandatory fields, and the TCP service request and response primitives and their mandatory fields. The TCP SAP provides an extended service response to pass the user identification (UID) to the host processor upon completion of a passive connection request.

All TCP SAP messages are designed to have a standard format as illustrated in Figure 2, and provide all information required by the corresponding TCP service request or response primitive. All parameter fields have also standard formats with defined values

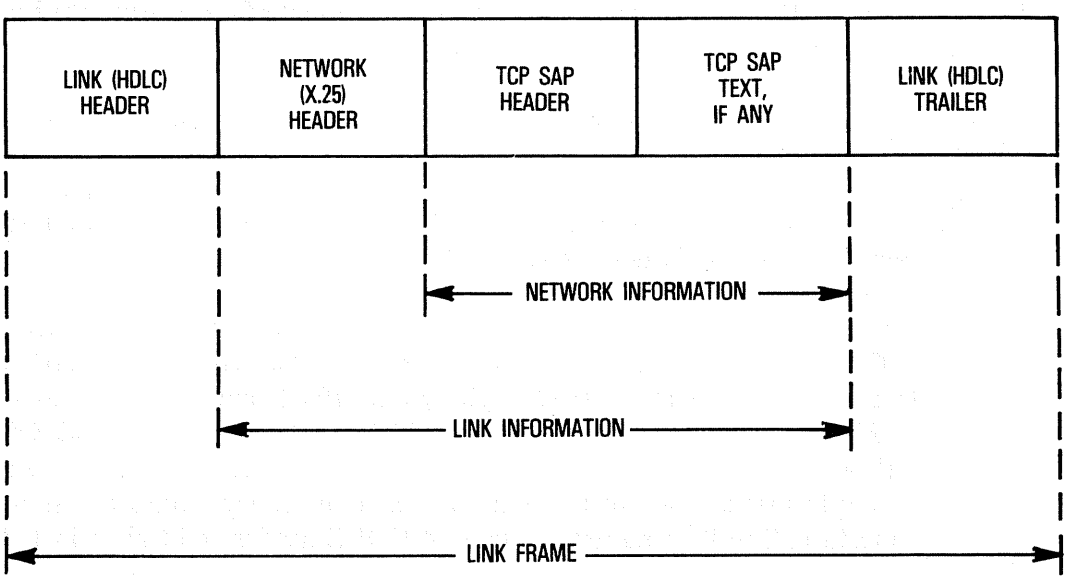

Figure 2. Channel Level Interface Data Structure.

for unused fields. The TCP SAP text field (see Figure 2) in the send and receive data messages is the only variable length field, and it has a maximum length to avoid data fragmentation in the network layers at all network interfaces.

A brief description of the TCP SAP messages, categorized into connection establishment, maintenance, termination, and data transport functions is given below. With the exception of the "extended" messages, there is a one-to-one mapping of TCP SAP messages and their parameters to TCP service request and response primitives.

1. TCP SAP Connection Establishing Messages. These allow a host processor to establish TCP connections for ULPs and utilize the Network Authentication Service (NAS) and Network Name Service (NNS) which have been defined earlier. The important SAP primitives are described below.
   a. Unspecified Passive Open Request: Defining security range and precedence level for a potential connection, and establishing connections with an unnamed ULP.
   b. Fully Specified Passive Open Request: Defining security range and precedence level for a potential connection, and establishing connections with a named ULP.
   c. Active Open Request. Defining security range and precedence level for a potential connection, and initiating and establishing connections with a named ULP.
   d. Extended Active Open Request. Functions of the active open request, specified above as well as passing a preauthenticated user identification (UID) to the destination host, and selectively defining the destination port and internet address of a named ULP as a symbolic name.
   e. Open ID Response. Informing a host of the local connection name assigned by the outboarded TCP service to a previously requested connection.
   f. Open Success Response. Informing a host of the successful completion of any connection request (including the complete address of the peer host as parameters).

g. Extended Open Success Response. Informing a host of the successful completion of a passive connection request and an authenticated UID is provided by the NAS to the host processor's operating environment.
h. Open Failure Response. Informing a host of the failure of any connection request and the reason(s) for failure using an error code.

2. TCP SAP Connection Maintenance Messages. These allow a host to monitor the status of an established connection.
   a. Error Response. Informing (via an unsolicited response) a host of predefined error conditions occurring on an established connection.
   b. Status Request. Querying the outboarded TCP of the current status of an established connection.
   c. Status Response. Informing a host, in response to a previous status request, of the current TCP status of an established connection.
3. TCP SAP Connection Termination messages. These allow a host to terminate, gracefully or by preemption, a previously requested or established connection.
   a. Close Request. Initiating an ordered termination of a named connection such that no data is lost.
   b. Closing Response. Informing a host ULP that an ordered termination of an established connection has been requested by the other host ULP.
   c. Abort Request. Initiating the forced, unilateral termination of a named connection.
   d. Terminate Response. Informing a host that a named connection has been terminated.
4. TCP SAP Data Transport Messages. This allows a host ULP to exchange information (including data stream push and urgent data signalling) with a peer-level ULP over an established connection.
   a. Send Request. Transmitting data to another host over an established connection, i.e., to support a ULP write function.
   b. Deliver Response. This allows a host to receive data from another host over an established connection, i.e., to support a ULP read function.

The TCP allocate service request primitive was not implemented because the memory management functions that it supports are not essential for a host processor. If a host processor has insufficient memory resources, the flow control features of X.25 will prohibit additional data from being received and thus prevent data loss.

The TCP SAP is designed to work with any channel level interface that provides the functions of ISO OSI reference model's lower three layers. The initial implementation has been defined to utilize the three X.25 layers and to map the TCP connections on a one-to-one basis to X.25 Permanent Virtual Circuits (PVCs) or Virtual Calls (VCs). This allows the TCP SAP to take advantage of the X.25 flow control and error detection/recovery features on a host to IU information exchange. These features of X.25 provide ordered, error-free data delivery and prevent a single TCP connection (an FTP process for example) from monopolizing the host-to-IU link.

PERFORMANCE EVALUATION OF THE TCP SAP

The performance of the TCP SAP and other protocols that are resident in the hosts and IUs were evaluated by both analysis and discrete-event simulation (Pegden, 1985) in terms of throughput, end-to-end packet delay, data reliability, path availability, and system security. The significant results are summarized below.

Since the TCP SAP is primarily concerned with packet transfer between a host and the associated IU, i.e., there are no routing issues, the impact on throughput is due to the number of error-free frames the network and link layers can accommodate via their flow control mechanism. Given that the IU is designed to have ample processing capacity, the throughput is constrained by the data rate of the host to IU.

The delay introduced by the TCP SAP and its supporting protocols in the packet transmission path consists of additional processing delays in the host and IU. This occurs in addition to the queueing and transmission delays at the host/IU interface, which are directly related to the average packet size, throughput, and data rate, and are unavoidable as long as the host and IU are physically distinct devices. The processing delays at the host and IU are dependent on the protocol complexity and implementation as well as on the respective processors' operating systems and hardware. The simulation results have shown that the additional processing delay (about 2 to 3 milliseconds in an IU) due to the TCP SAP and its supporting functions is not significant in comparison to the end-to-end one-way delay (in the order of 100 milliseconds) between a pair of hosts when the communication path does not include a bridge or a gateway. The existence of a bridge or a gateway is expected to cause an additional delay in the range of 10 to 50 milliseconds.

Data reliability is related to the undetected frame error rate which, in turn, depends on the raw-bit error rate at the physical link and the error detecting algorithm at the link layer. (The use of HDLC LAPB with the TCP SAP provides a 16-bit CRC.) The probability of undetected frame errors in the host/IU interface was analytically determined to be at least one order less in magnitude than the corresponding value (typically one in a million) for a host-to-host path.

The data path from the source host to the destination host includes, among other protocols, two pairs of client and server TCP SAPs, one each at the source and destination end. The TCP SAP constitutes only a minor part of hardware and software in a host and an IU. The impact of TCP SAP on the host-to-host path availability was analyzed and was found to have no major significance.

The TCP SAP has no direct bearing on the centralized security system because it simply transfers security-related parameters between ULPs and NAS and TCP.

Other indirect effects of the TCP SAP include system development costs and logistic delays in the project schedule. Although the TCP SAP has to be accommodated in each host and IU, it does not have a significant bearing on the overall cost of the network because of the following reasons:

1. The hardware and software requirements for the TCP SAP are small in comparison to those related to other protocols, and
2. The TCP SAP hardware and software have a modular structure that, once developed, could be ported in a variety of hosts and IUs.

The logistic delays can be avoided by appropriate planning and scheduling, which is a subject of project management and is beyond the scope of this research.

SUMMARY AND CONCLUSIONS

The TCP SAP presented in this article provides a flexible, evolutionary means for access to the various network communication services. The TCP SAP allows for a transparent interface between upper layer protocols (ULPs) in a host and the network transport layer, namely TCP, that is resident in a physically separate interface unit (IU). The TCP SAP architecture permits incorporation of additional transparent interfaces

to other protocols such as IP without any significant modifications in the software structure. The host-front-end protocol consists of the TCP SAP overlayed on a defined implementation of the CCITT X.25 layer 3 through 1 protocol set.

The performance of TCP SAP and its supporting protocols was evaluated by analysis and simulation of the network protocols. The additional functions that are required to be performed (due to the TCP SAP and its supporting protocols) do not apparently have a major impact on the network system performance with respect to throughput, end-to-end delay, data reliability, path availability, system security, and cost.

Although the TCP SAP is specifically designed for a given network configuration, it can be adapted to other distributed data communication systems, and its concept is applicable to local area and long-haul networks with diverse architectures.

REFERENCES

Bernosky, L., & Blankertz, W. (1983). The LAN as a transition tool in the WWMCCS modernization. *Government Data Systems, 12*(6), 340–356.

Boeing Company. (1985, November). *Technical and office protocols specification version 1.0.* Seattle, WA: Author.

Day, J. D., & Zimmermann, H. (1983). The OSI reference model. *IEEE Proceedings, 71*, 1334–1345.

Digital Equipment Corporation. (1983, May). *DECnet, digital network architecture, routing layer specification*, Version 2.0.0. Maynard, MA: Author.

Electronic Industries Association. (1977, December [Bulletin]; 1980, January [Addendum]). EIA Standard RS-449, *Interface for data terminal equipment and data circuit-terminating equipment.*

Folts, H.C. (1980). X.25 transaction-oriented features—datagram and fast-select. *IEEE Trans. on Comm., COM-28*(4), 496–499.

Groenbaek, I. (1986). Conversion between the TCP and ISO transport protocols as a method of achieving interoperability between data communication systems. *IEEE Journal on Selected Areas in Communications, SAC-4*(2), 288–296.

IBM. (1982). *Systems network architecture: Technical overview* (GC30-3073-0). Research Triangle Park, NC: Author.

Information processing systems—open systems interconnection—transport protocol specification (ISO DIS 8073). Revised, September 1984.

Lilenkamp, J., Mandell, R., & Padlipsky, M. (1984, December). Proposed host-front-end protocol. Network Working Group, RFC 929.

Meijer, A., & Peeters, P. (1982). *Computer network architectures.* Rockville, MD: Computer Science Press.

Padlipsky, M. (1984, December). Introduction to proposed DoD standard HFP. Network Working Group, RFC 928.

Pegden, C. D. (1985, July). *Introduction to SIMAN.* State College, PA: Systems Modeling Corporation.

Petrosky, M. (1987, June 1). TCP/IP Glue for multivendor nets. *Network World, 4*(22), p. 1.

Ray, A., Angevine, R. M., & Haughney, J. F. (1985, October). A service access procedure in support of the WIS architecture. *IEEE Military Communications Conference.* Boston, MA.

Rybczynski, A. (1980). X.25 interface to end-to-end virtual circuit service characteristics. *IEEE Trans. on Comm., COM-28*(4), 500–510.

Schwartz, M. (1987). *Telecommunication networks: Protocols, modeling and analysis.* Reading, MA: Addison-Wesley.

Society of Manufacturing Engineers. (1986, August). *Manufacturing automation protocol reference specification 2.1A and 2.2.* Dearborn, MI: Author.

U. S. Department of Defense. (1983a, August 12). *MIL-STD-1778 Transmission control protocol (TCP).* Washington, DC: Author.

U. S. Department of Defense. (1983b, August 12). *MIL-STD-1777 Internet protocol (IP).* Washington, DC: Author.

U. S. Department of Defense. (1984a, May 10). *MIL-STD-1780 File transfer protocol (FTP).* Washington, DC: Author.

U. S. Department of Defense. (1984b, May 10). *MIL-STD-1781 Simple mail transfer protocol (SMTP).* Washington, DC: Author.

U. S. Department of Defense. (1984c, May 10). *MIL-STD-1782 TELNET Protocol.* Washington, DC: Author.

U. S. Department of Defense. (1985, September 30). *MIL-STD-188-114 Electrical characteristics of digital interface circuits.* Washington, DC: Author.

Acknowledgement—The authors gratefully acknowledge the technical contributions of the first author's former colleagues R. M. Angevine and J. F. Haughney at GTE Strategic Systems Division.