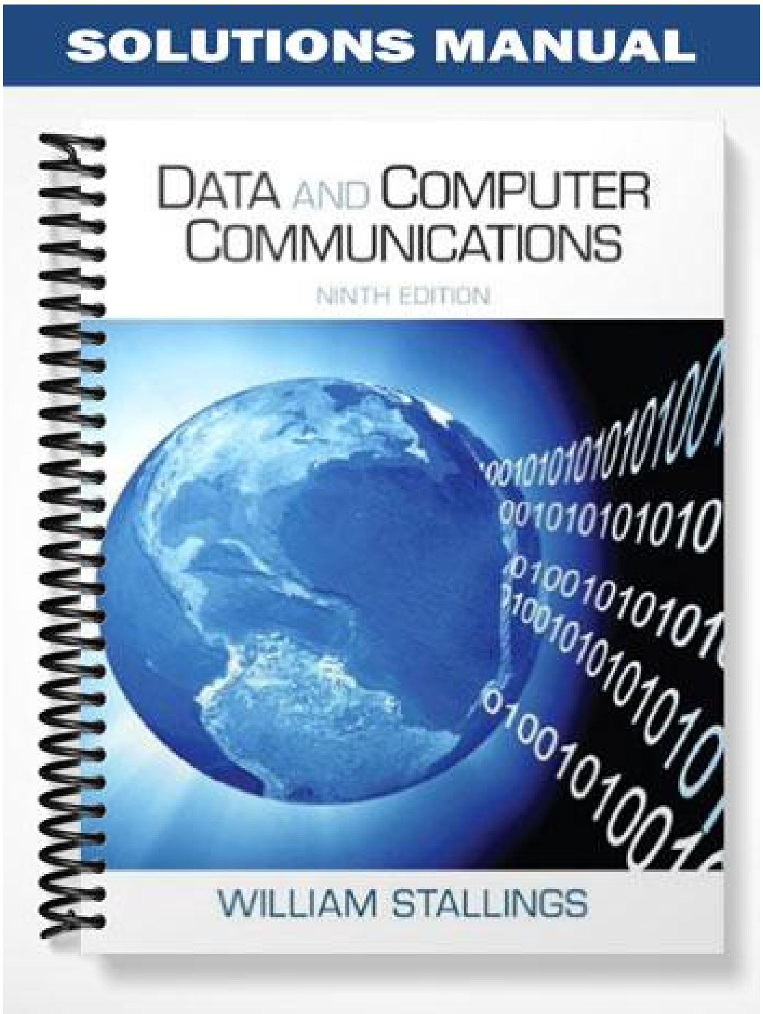

SOLUTIONS MANUAL

DATA AND COMPUTER COMMUNICATIONS

NINTH EDITION

WILLIAM STALLINGS

CHAPTER 2 PROTOCOL ARCHITECTURE

ANSWERS TO QUESTIONS

2.1 The network access layer is concerned with the exchange of data between a computer and the network to which it is attached.

2.2 The transport layer is concerned with data reliability and correct sequencing.

2.3 A protocol is the set of rules or conventions governing the way in which two entities cooperate to exchange data.

2.4 A PDU is the combination of data from the next higher communications layer and control information.

2.5 The software structure that implements the communications function. Typically, the protocol architecture consists of a layered set of protocols, with one or more protocols at each layer.

2.6 Transmission Control Protocol/Internet Protocol (TCP/IP) are two protocols originally designed to provide low level support for internetworking. The term is also used generically to refer to a more comprehensive collection of protocols developed by the U.S. Department of Defense and the Internet community.

2.7 Layering decomposes the overall communications problem into a number of more manageable subproblems.

2.8 A router is a device that operates at the Network layer of the OSI model to connect dissimilar networks.

2.9 IPv4.

2.10 No, other transport layer protocols, such as UDP, are also used. Some traffic uses no transport protocol, such as ICMP.

2.11 IPv4 - 32 bits; IPv6 - 128 bits

ANSWERS TO PROBLEMS

2.1 The guest effectively places the order with the cook. The host communicates this order to the clerk, who places the order with the cook. The phone system provides the physical means for the order to be transported from host to clerk. The cook gives the pizza to the clerk with the order form (acting as a "header" to the pizza). The clerk boxes the pizza with the delivery address, and the delivery van encloses all of the orders to be delivered. The road provides the physical path for delivery.

2.2 **a.**

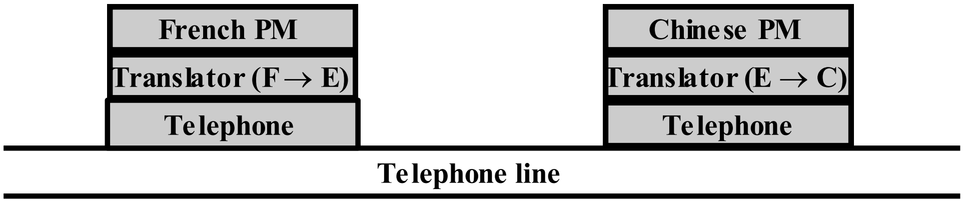

The PMs speak as if they are speaking directly to each other. For example, when the French PM speaks, he addresses his remarks directly to the Chinese PM. However, the message is actually passed through two translators via the phone system. The French PM's translator translates his remarks into English and telephones these to the Chinese PM's translator, who translates these remarks into Chinese.

b.

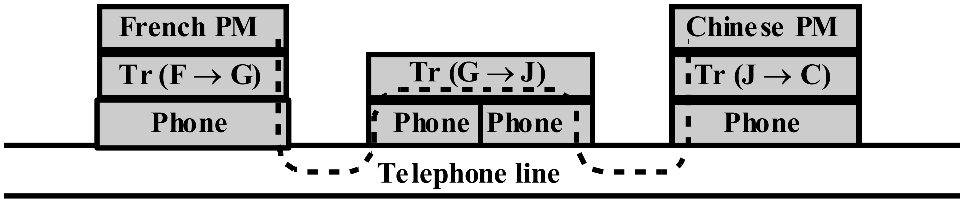

An intermediate node serves to translate the message before passing it on. Note that the intermediate node handles the message only up to the second level; a minister's level is not needed.

2.3 Perhaps the major disadvantage is the processing and data overhead. There is processing overhead because as many as seven modules (OSI model) are invoked to move data from the application through the communications software. There is data overhead because of the appending of multiple headers to the data. Another possible disadvantage is that there must be at least one protocol standard per layer. With so many layers, it takes a long time to develop and promulgate the standards.

2.4 No. There is no way to be assured that the last message gets through, except by acknowledging it. Thus, either the acknowledgment process continues forever, or one army has to send the last message and then act with uncertainty.

2.5 A case could be made either way. **First**, look at the functions performed at the network layer to deal with the communications network (hiding the details from the upper layers). The network layer is responsible for routing data through the network, but with a broadcast network, routing is not needed. Other functions, such as sequencing, flow control, error control between end systems, can be accomplished at layer 2, because the link layer will be a protocol directly between the two end systems, with no intervening switches. So it would seem that a network layer is not needed. **Second**, consider the network layer from the point of view of the upper layer using it. The upper layer sees itself attached to an access point into a network supporting communication with multiple devices. The layer for assuring that data sent across a network is delivered to one of a number of other end systems is the network layer. This argues for inclusion of a network layer.

In fact, the OSI layer 2 is split into two sublayers. The lower sublayer is concerned with medium access control (MAC), assuring that only one end system at a time transmits; the MAC sublayer is also responsible for addressing other end systems across the LAN. The upper sublayer is called Logical Link Control (LLC). LLC performs traditional link control functions. With the MAC/LLC combination, no network layer is needed (but an internet layer may be needed).

2.6 **a.** No. This would violate the principle of separation of layers. To layer (N – 1), the N-level PDU is simply data. The (N – 1) entity does not know about the internal format of the N-level PDU. It breaks that PDU into fragments and reassembles them in the proper order.

b. Each N-level PDU must retain its own header, for the same reason given in (a).

2.7 Data plus transport header plus internet header equals 1820 bits. This data is delivered in a sequence of packets, each of which contains 24 bits of network header and up to 776 bits of higher-layer headers and/or data. Three network packets are needed. Total bits delivered = 1820 + 3 × 24 = 1892 bits.

2.8 UDP provides the source and destination port addresses and a checksum that covers the data field. These functions would not normally be performed by protocols above the transport layer. Thus UDP provides a useful, though limited, service.

2.9 In the case of IP and UDP, these are unreliable protocols that do not guarantee delivery, so they do not notify the source. TCP does guarantee delivery. However, the technique that is used is a timeout. If the source does not receive an acknowledgment to data within a given period of time, the source retransmits.

2.10 UDP has a fixed-sized header. The header in TCP is of variable length.

2.11 Suppose that A sends a data packet k to B and the ACK from B is delayed but not lost. A resends packet k, which B acknowledges. Eventually A receives 2 ACKs to packet k, each of which triggers transmission of packet (k + 1). B will ACK both copies of packet (k + 1), causing A to send two copies of packet (k + 2). From now on, 2 copies of every data packet and ACK will be sent.

2.12 TFTP can transfer a maximum of 512 bytes per round trip (data sent, ACK received). The maximum throughput is therefore 512 bytes divided by the round-trip time. Source: [STEV94].

2.13 The "netascii" transfer mode implies the file data are transmitted as lines of ASCII text terminated by the character sequence {CR, LF}, and that both systems must convert between this format and the one they use to store the text files locally. This means that when the "netascii" transfer mode is employed, the file sizes of the local and the remote file may differ, without any implication of errors in the data transfer. For example, UNIX systems terminate lines by means of a single LF character, while other systems, such as Microsoft Windows, terminate lines by means of the character sequence {CR, LF}. This means that a given text file will usually occupy more space in a Windows host than in a UNIX system.

2.14 If the same TIDs are used in twice in immediate succession, there's a chance that packets of the first instance of the connection that were delayed in the network arrive during the life of the second instance of the connection, and, as they would have the correct TIDs, they could be (mistakenly) considered as valid.

2.15 TFTP needs to keep a copy of only the last packet it has sent, since the acknowledgement mechanism it implements guarantees that all the previous packets have been received, and thus will not need to be retransmitted.

2.16 This could trigger an "error storm". Suppose host A receives an error packet from host B, and responds it by sending an error packet back to host B. This packet could trigger another error packet from host B, which would (again) trigger an error packet at host A. Thus, error messages would bounce from one host to the other, indefinitely, congesting the network and consuming the resources of the participating systems.

2.17 The disadvantage is that using a fixed value for the retransmission timer means the timer will not reflect the characteristics of the network on which the data transfer is taking place. For example, if both hosts are on the same local area network, a 5-second timeout is more than enough. On the other hand, if the transfer is taking place over a (long delay) satellite link, then a 5-second timeout might be too short, and could trigger unnecessary retransmissions. On the other hand, using a fixed value for the retransmission timer keeps the TFTP implementation simple, which is the objective the designers of TFTP had in mind.

2.18 TFTP does not implement any error detection mechanism for the transmitted data. Thus, reliability depends on the service provided by the underlying transport protocol (UDP). While the UDP includes a checksum for detecting errors, its use is optional. Therefore, if UDP checksums are not enabled, data could be corrupted without being detected by the destination host.

2.19 **a.** The internet protocol can be defined as a separate layer. The functions performed by IP are clearly distinct from those performed at a network layer and those performed at a transport layer, so this would make good sense.

b. The session and transport layer both are involved in providing an end-to-end service to the OSI user, and could easily be combined. This has been done in TCP/IP, which provides a direct application interface to TCP.

CHAPTER 3 PROTOCOL ARCHITECTURE