Chapter 2 Review Questions

1. The browser program running in the user’s host (desktop, laptop, tablet, smartphone, and so on) and the Web server program running in the Web server host.

2.

3. Web, FTP, Telnet, e-mail etc.

4. No. In a P2P file-sharing application, the peer that is receiving a file is typically the client and the peer that is sending the file is typically the server.

5.

6. You would use UDP. With UDP, the transaction can be completed in one roundtrip time (RTT) - the client sends the transaction request into a UDP socket, and the server sends the reply back to the client's UDP socket. With TCP, a minimum of two RTTs are needed - one to set-up the TCP connection, and another for the client to send the request, and for the server to send back the reply.

7.

8. a) Reliable data transfer

TCP provides a reliable byte-stream between client and server but UDP does not.

b) A guarantee that a certain value for throughput will be maintained

Neither

c) A guarantee that data will be delivered within a specified amount of time

Neither

d) Confidentiality (via encryption)

Neither

9.

10. HTTP defines how Web clients request Web pages from Web servers and how servers transfer Web pages to clients. Because an HTTP server maintains no information about the clients, HTTP is said to be a stateless protocol. For example, if a particular client asks for the same object twice in a period of a few seconds, the server does not respond by saying that it just served the object to the client; instead, the server resends the object, as it has completely forgotten what it did earlier.

11. The applications associated with those protocols require that all application data be received in the correct order and without gaps. TCP provides this service whereas UDP does not.

12. When the user first visits the site, the server creates a unique identification number, creates an entry in its back-end database, and returns this identification number as a cookie number. This cookie number is stored on the user's host and is managed by the browser. During each subsequent visit (and purchase), the browser sends the cookie number back to the site. Thus the site knows when this user (more precisely, this browser) is visiting the site.

13. HTTP has a mechanism that allows a cache to verify that its objects are up to date. This mechanism is called the conditional GET. A message is a so-called conditional GET message if (1) the request message uses the GET method and (2) the request message includes an If-Modified-Since header line.

14.

15. The server must keep track of the user's current directory as the user wanders about the remote directory tree.

16. The message is first sent from Alice's host to her mail server over HTTP. Alice's mail server then sends the message to Bob's mail server over SMTP. Bob then transfers the message from his mail server to his host over POP3.

17.

18.

19. Yes an organization's mail server and Web server can have the same alias for a host name. The MX record is used to map the mail server's host name to its IP address.

20.

Chapter 2 Problems

Problem 1

a) T
b) F
c) F
d) F
e) T

Problem 3

(i) The types of messages exchanged, for example, request messages and response Messages. (ii) The syntax of the various message types, such as the fields in the message and how the fields are delineated. (iii) The semantics of the fields, that is, the meaning of the information in the fields. (iv) Rules for determining when and how a process sends messages and responds to messages.

Problem 7

The total time to get the IP address is $kD_1 + (n\text{-}k)D_2$. After this, 2RTT_0 is required to set up each of the *m* TCP connections and to request and receive each object. Thus, the total response time is $kD_1 + (n\text{-}k)D_2 + 2m\ RTT_0$.

Problem 8

a) $2D_1 + D_2 + 12\ RTT_0$

The total time to get the IP address is $2D_1 + 1D_2$. After this, 2RTT_0 is required to set up each of the *5* TCP connections and to request and receive each HTTP object. Thus, the total response time is $2D_1 + 1D_2$ (1%) + 1 RTT_0 (TCP handshaking) + 1 RTT_0 (HTTP request/response for web page) (1%) +5* [1 RTT_0 (parallel TCP handshaking) +1 RTT_0 (HTTP request/response for 5 objects)] (1%) = $2D_1 + D_2 +$

12 RTT_0 (2%)

b) $2D_1 + D_2 + 4\ RTT_0$

$2D_1 + 1D_2$ + 1 RTT_0 (TCP handshaking) + 1 RTT_0 (HTTP request/response for web page) + 1 RTT_0 (parallel TCP handshaking) +1 RTT_0 (HTTP request/response for 5 objects) = $2D_1 + D_2 + 4\ RTT_0$

c) $2D_1 + D_2 + 3\ RTT_0$

$2D_1 + 1D_2$ + 1 RTT_0 (TCP handshaking) + 1 RTT_0 (HTTP request/response for web page) + 1 RTT_0 (HTTP request/response for 5 objects) = $2D_1 + D_2 + 3\ RTT_0$

Problem 9

a) The time to transmit an object of size *L* over a link or rate *R* is *L/R*. The average time is the average size of the object divided by *R*:

Δ= (960,000 bits)/(16,000,000 bits/sec) = 0.06 sec

The traffic intensity on the link is given by $\beta\Delta$=(15 requests/sec)(.06 sec/request) = 0.9. Thus, the average access delay is (0.06 sec)/(1 - 0.9) = 0.6 seconds. The total average response time is therefore .6 sec + 2 sec = 2.6 sec.

b) The traffic intensity on the access link is reduced by 60% since the 60% of the requests are satisfied within the institutional network. Thus the average access delay is (0.06 sec)/[1 – (0.4)(0.9)] = 0.094 seconds. The response time is approximately zero if the request is satisfied by the cache (which happens with probability .6); the average response time is 0.094 sec + 2 sec = 2.094 sec for cache misses (which happens 40% of the time). So the average response time is (0.6)(0 sec) + (0.4)(2.094 sec) = 0.838 seconds. Thus the average response time is reduced from 2.6 sec to 0.838 sec.

Problem 13

The control connection is used for sending control information between the two hosts - information such as user identification, password, commands to change remote directory, and commands to “put” and “get” files. The data connection is used to actually send a file.

Problem 16

With POP3 access, one cannot arrange emails in folders in the server. One has to download in local machine and then one can delete messages, move messages across folders, and search for messages (by sender name or subject). But this paradigm—namely, folders and messages in the local machine—poses a problem for the nomadic user, who would prefer to maintain a folder hierarchy on a remote server that can be accessed from any computer. This problem is solved in IMAP protocol.

Problem 20

(1) The user machine runs the client side of the DNS application.
(2) The browser extracts the hostname, www.someschool.edu, from the URL and passes the hostname to the client side of the DNS application.
(3) The DNS client sends a query containing the hostname to a DNS server.
(4) The DNS client eventually receives a reply, which includes the IP address for the hostname.
(5) Once the browser receives the IP address from DNS, it can initiate a TCP connection to the HTTP server process located at port 80 at that IP address.