

SOLUTION MANUAL
OF
COMMUNICATION NETWORKS
BY
ALBERTO LEON GARCIA & INDRA WIDJAJA

Solutions to Chapter 1

1a. Describe the step-by-step procedure that is involved from the time you deposit a letter in a mailbox to the time the letter is delivered to its destination. What role do names, addresses and mail codes (such as ZIP codes or postal codes) play? How might the letter be routed to its destination? To what extent can the process be automated?

Solution:

The steps involved in mailing a letter are:

1. The letter is deposited in mailbox.
2. The letter is picked up by postal employee and placed in sack.
3. The letter is taken to a sorting station, where it is sorted according to destination, as determined by the mail code and grouped with other letters with the same destination mail code. (If there is no mail code, then it is determined by the largest geographical unit, for example, country (if specified), otherwise state (if specified), otherwise city (if specified).)
4. The letter is shipped to the post office that handles the mail for the specific mail code (or country or city).
5. The letter is then sorted by street address.
6. The letter is picked up at the post office by the postal worker responsible for delivering to the specified address.
7. The letter is delivered according to the number and street.

The name is not really used, unless the street address is missing or incorrect. The name is at the destination to determine who the letter belongs to. (Unless of course the letter is being sent to a small town, where most inhabitants are known to the postal worker.)

The mail delivery process can be automated by using optical recognition on the mail code. The letter can then be sorted and routed to the destination postal station, and even to the destination neighborhood, depending on the amount of geographical detail built into the mail code.

1b. Repeat part (a) for an e-mail message. At this point, you may have to conjecture different approaches about what goes on inside the computer network.

Solution:

The steps involved in e-mailing a message are:

1. The message is sent electronically by clicking 'Send'. (In Chapter 2 we see that the Simple Mail Transfer Protocol (SMTP) is used to do this.)
2. The mail provider of the sender sends a request to a name server for the network address of the mail provider of the recipient. The mail provider is determined by the information following the @ symbol.
3. If the mail provider finds the network address of the recipient's mail provider, then it sends the message to that address.

4. Upon receipt of the message the mail provider either informs the recipient automatically that there is mail or waits until mail is requested by the recipient, depending on how the mail service is set up. The recipient is determined by the information before the @ symbol.

- 1c. Are the procedures in parts (a) and (b) connection-oriented or connectionless?

Solution:

The transfer service of letters in a postal system is connectionless. Users do not set up connections prior to the transfer of each letter. (Note however that the postal system may have regularly scheduled shipments of bundles of mail from city to city, because the volume of mail between cities is predictable. These regular shipments can be viewed as pre-arranged "connections.")

E-mail transfer is also a connectionless service in that the *user* does not set up end-to-end connections to send a message. We will see in Chapter 2, however, that e-mail *protocols* use connection-oriented procedures to exchange messages.

NOTE: For parts (a) and (b) other step-by-step procedures are possible.

- 2a. Describe what step-by-step procedure might be involved inside the network in making a telephone connection.

Solution:

1. The telephone number specifies an "address" where the receiver is located. In the North American telephone numbering system the first three digits are the area code which specifies the main geographical region of the receiver; the next three digits specify a particular telephone office in the North American network. The final four digits identify the particular location of the receiver.
2. When a telephone number is dialed equipment at the other end of the telephone line uses the sequence of dialed numbers to determine a route across the telephone network from the call-originating phone to the destination phone. A circuit is established between the originating and destination phone along the identified route.
3. A ringing tone is then applied at the destination to indicate that there is an incoming call. If the destination party (a person, an answering machine, or some other device) is ready to answer, the call is completed through the lifting of the phone set or some equivalent action. The setting up of telephone calls is discussed in Chapter 4.

- 2b. Now consider a personal communication service that provides a user with a personal telephone number. When the number is dialed, the network establishes a connection to wherever the user is located at the given time. What functions must the network now perform in order to implement this service?

Solution:

1. The key difference here is that the personal telephone number is not tied to a specific location. Instead the personal number is associated with one or more pieces of equipment, for example, a cell phone, that can request service from various points in the network. This necessitates the translation of the personal telephone number to a number that corresponds to a specific location in the network at a given time.
2. When the telephone number is dialed, a message requesting a connection setup is sent to the "home" location of the personal number.

3. The home location must somehow be able to redirect the connection setup process to the current location of the user. For example, the user may register one or more forwarding "addresses" to which calls are to be redirected. The signaling system is used to locate the user. In cellular telephony, for example, requests for connections to a given mobile telephone are broadcast over specific signaling channels.
 4. Once the location of the destination is identified a connection is established.
3. Explain how the telephone network might modify the way calls are handled to provide the following services:

Solutions follow questions:

- a. Call Display: the number and/or name of the calling party is listed on a screen before the call is answered.

Along with the request for connection setup, the system sends identifying information regarding the originator's phone.

- b. Call Waiting: a special sound is heard when the called party is on the line and another user is trying to reach the called party.

When a second caller tries to access a busy line, the system applies an audible sound, for example, a "beep" or a tone, to the voice signal that is being sent to the called party. The tone alerts the called party that there is an incoming call and provides the choice of answering the second call without disconnecting the first call.

- c. Call Answer: if the called party is busy or after the phone rings a prescribed number of times, the network gives the caller the option of leaving a voice message.

If the call is not answered after so many rings, the system establishes a connection to another number (that is, forwards the call) that is associated with the answering system. The call is then automatically answered by an answering system.

- d. Three-Way-Calling: allows a user to talk with two other people at the same time.

After the establishment of a call between two parties, the caller can then indicate to the telephone system that he or she wants to call a third party. Without disconnecting the first pair, the system sets up a connection to the second called party, and once established, transmits both conversations to all parties simultaneously. Special equipment is required to combine and distribute the voice signals of the three parties.

- 4a. Suppose that the letter in problem 1 is sent by fax. Is this mode of communications connectionless or connection-oriented? Real-time or non-real time?

Solution:

In order to send the letter by fax, a telephone connection must first be established. Therefore the mode of communications is connection-oriented. The transfer of information across the network occurs in real-time.

- 4b. Repeat part (a) if a voice mail message is left at a given telephone.

Solution:

Again, the process of leaving a voice mail message is connection-oriented, since there must be an

end-to-end connection between the caller and the receiver with the answering service. The leaving of voice mail is also done in real-time. However, the mode of communication between the person that leaves the voice mail and the person that listens to the voice mail is connectionless and non-real-time. The people involved do not need to establish a simultaneous connection with each other in order to communicate; also, they record and listen to the messages at different times.

5. Suppose that network addresses are scarce, so addresses are assigned so that they are not globally unique; in particular suppose that the same block of addresses may be assigned to different organizations. How can the organizations make use of these addresses? Can users from two such organizations communicate with each other?

Solution:

To make the example concrete suppose that two organizations are assigned the same set of telephone numbers. Clearly, users within each organization can communicate with each other as long as they have a unique address within the organization. However, communications outside an organization poses a problem since any given address is no longer unique across multiple organizations.

A possible approach to enabling communications between users in different organizations is to use a two-step procedure as follows. Each organization has a special gateway to communicate outside the organization. Internal users contact the gateway to establish calls to other organizations. Gateways have procedures to establish connections with each other. This enables gateways to establish connections between their internal users and users in other networks.

6a. Describe the similarities and differences in the services provided by (1) a music program delivered over broadcast radio and (2) music delivered by a dedicated CD player.

Solution:

Both broadcast radio and a dedicated CD player provide users with similar types of information and with almost the same performance. However the broadcast radio offers its service in a real-time fashion with no interaction with users. A CD player, on the other hand, stores its information on a CD and delivers it on demand.

6b. Describe how these services might be provided and enhanced by providing them through a communications network.

Solution:

Broadcast radio is typically transmitted in real time using radio waves "over the air." Broadcast audio programs can also be transmitted over any communication network. Furthermore, when stored in servers that are attached to a network, "broadcast" audio programs can be retrieved for listening at a later time. Indeed a server can also store CD audio material and retrieve it on demand. If the network and server are sufficiently responsive, it may also be possible to provide the interactivity of a personal CD player through a network-based service.

7a. Use the World Wide Web to visit the sites of several major newspapers. How are these newspapers changing the manner in which they deliver news services?

Solution:

Traditionally, newspapers deliver news in printed "hard-copy" form in various editions during the day. The information is primarily in text form, supplemented by still pictures and graphics. The information in a newspaper is organized in several sections: main page, business, local, sports, classifieds, etc.

Some of the revenue of newspapers is from subscriptions, but most of the newspaper revenue is usually from advertising.

In the traditional model, getting the news from a particular newspaper required getting the physical printed form. The WWW and the Internet allow newspapers to deliver information in electronic form and at any time that a user requests it. Newspaper information on the Web still consists mostly of text, images, and graphics that can be retrieved and displayed readily by a Web browser. In most cases, subscription is not required and advertising is displayed along with news information. In order to be competitive, newspapers can no longer work in terms of editions, especially for late-breaking news. Instead, news items are updated continuously day and night. The use of a Web browser allows a user to quickly find the news items of interest. In addition, newspapers can provide more in-depth coverage in their web page on topics of special interest than is possible in print medium.

7b. Now visit the websites of several major television networks. How are they changing the manner in which they deliver news over the Internet? What differences, if any, are there to the approach taken by television networks and newspapers?

Solution:

Traditionally, television networks deliver news using audio, video, and graphics, with very little text. News is traditionally delivered in programs that are broadcast at various times of the day. Again, the information is organized in several sections: international, business, sports, entertainment, etc. The revenue in television news is almost entirely from advertising. The WWW and the Internet allow users to retrieve news information in electronic form and at any time. However, the transmission of audio and video over the Internet poses a challenge, so this information must be transmitted in compressed form and in relatively low quality. Consequently, more of the information is provided in text form. The dependence on audio and video also required the development of applications for the decompression and display of such information. Again, in order to be competitive, news networks update their Web pages continuously day and night as news develops. The higher cost involved in providing video service has led to the introduction of subscription fees.

8. Discuss the advantages and disadvantages of transmitting fax messages over the Internet instead of the telephone network.

Solution:

The comparison of fax transmission over the Internet and over the telephone network reflect the differences between e-mail and telephone calls. The transmission of a fax message as an attachment to e-mail means that the transfer of the message is not in real time and that delivery is not necessarily confirmed. The transmission of a fax message over the telephone network, on the other hand, is real time and is confirmed with fairly high certainty. On the other hand, the cost of transmission of a fax message over the Internet is very inexpensive and is not distance-sensitive. The transmission of fax messages over telephone networks may involve long distance charges.

9a. Suppose that an interactive video game is accessed over a communication network. What requirements are imposed on the network if the network is connection-oriented? connectionless?

Solution:

We suppose that the game involves the interaction between a player and a server across a network. To support an interactive video game over a communications network, the network, whether connection-oriented or connectionless, must provide real-time delivery of the player's commands to the server, and of the server's responses to the player. With a connection-oriented network, connections between the player and the servers transfer the sequence of commands and responses

throughout the game with very little delay. In a connectionless network, user commands may be delivered to the other end with variable delay, out-of-sequence, or not at all. The user's network software is responsible for ensuring the ordered and correct delivery of game commands. In-time delivery of commands cannot be assured.

- 9b. Repeat part (a) if the game involves several players located at different sites.

Solution:

The requirements on the network depend on how the game is implemented. In the centralized approach the players interact through a central server that processes the commands from all of the players, maintains a view of the state of the overall system, and issues appropriate responses to all the players. Alternatively, the game could be implemented in a decentralized fashion, where each player receives commands from some or all of the players, maintains a local view of the system state, and transmits responses to some or all of the players.

In the centralized approach, the network requirements are essentially the same as those in part (a). In the decentralized approach, the real-time response requirement may apply only to a subset of players when they happen to be interacting with each other. If players are located at different sites, the network could provide multicast capability so messages can be exchanged among the server and all the players.

- 9c. Repeat part (b) if one or more of the players is in motion, for example, kids in the back of the van during a summer trip.

Solution:

If one or more players is in motion, then the network must also be able to locate each mobile player and then deliver and receive information to and from such player. Additional delay and loss issues come into play when a mobile user is "handed-off" from one radio coverage area to another.

10. Discuss the similarities between the following national transportation networks and a communications network. Is the transportation system more similar to a telephone network or to a packet network?

Solution:

Transportation networks are designed to transfer people and goods; communications networks are designed to transfer information. Like communication networks, all transportation systems involve links, in the form of roads, rails, or air corridors, and switching points, in the form of stations, airports, and highway interchanges. Addressing and geographical names are used in transportation networks to identify destinations, and routing of various forms is required to direct goods to their destinations. Transportation resources, in the forms of cars or airplanes, are shared or "multiplexed" among various goods that traverse common parts of the network.

- a. Railroad network.

Railroad network: Goods or people arrive at train stations and are loaded for transport to specific places. Tickets are purchased for a given destination, usually guaranteeing that there will be a place in a car from the starting station all the way to the destination station. This is similar to establishing a connection across a telephone network. Alternatively, a passenger may choose to purchase a ticket at each station along the way. This corresponds more closely to a connectionless packet network model.

- b. Airline network.

Airline network: In this case passengers purchase tickets that guarantee a seat all the way to the destination, even if transfers are made at intermediate airports. This is similar to the establishment of connections across a telephone network.

c. Highway system.

Highway systems: Trucks or cars enter the highway without making reservations ahead of time and without informing any central authority of their destination or route. This mode of operation corresponds closely to the operation of a connectionless packet network.

d. Combination of (a), (b), and (c).

Combination of transportation systems: The combination of air, rail, and highway transportation systems can and are used jointly for the transfer of people and goods. Each transportation system is different in how it is organized and how transfers are accomplished. Nevertheless, it is possible to use these systems in combination to provide a higher degree of connectivity between sources and destinations of people and goods. We will see that the Internet plays a role similar to the combined transportation systems in that the Internet enables the transfer of information across multiple dissimilar networks that may differ in how they are organized and how they operate.

11. In the 1950s, standard containers were developed for the transportation of goods. These standard containers could fit on a train car, on a truck, or in specially designed container ships. The standard size of the containers makes it possible to load and unload them much more quickly than using non-standard containers of different sizes. Draw an analogy to packet switching communications networks. In your answer identify what might constitute a container and speculate on the advantages that may come from standard-size information containers.

Solution:

In a packet-switching communications network, the component that is analogous to the container is a constant-size packet (preferably of short length) that can be used for the transport of information. Transmission systems and networks of various types can be designed to transfer information of the given standardized size, much like trucks, trains, and ships can be designed to carry standard containers. Packing and unpacking of fixed-size units is simpler than packing and unpacking of variable-length units. Consequently, it is simpler to schedule the transfer of packets across switches that use constant-size packets than across switches that make use of variable-length packets.

12. The requirements of world commerce led to the building of the Suez and Panama canals. What analogous situations might arise in communication networks?

Solution:

The Suez and Panama canals were built to provide a shorter path between geographical areas with major trade relationships. The analogous situation arises in communication networks when the path between two areas in the network are too "long" either in terms of excessive delay or inadequate capacity to transfer the required volumes of traffic. The addition of network resources to provide shorter paths for lower delay or wider paths for higher capacity corresponds to the building of a canal.

13. Two musicians located in different cities wish to have a jam session over a communications network. Find the maximum possible distance between the musicians if they are to interact in real-time, in the sense of experiencing the same delay in hearing each other as if they were 10 meters apart. The speed of sound is approximately 330 meters/second. Assume that the network transmits the sound at the speed of light in cable, 2.3×10^8 meters/second.

Solution:

The first step is to find the delay for the sound when the musicians are 10 meters apart:

$$t_{10} = 10/330 = 30.30 \text{ milliseconds}$$

The maximum distance is the time required for a real-time 'experience' times the cable speed:

$$d = (2.3 \times 10^8) \times (30.30 \times 10^{-3}) = 6,969,000 \text{ meters} = 6969 \text{ kilometers}$$

14. The propagation delay is the time that is required for the energy of a signal to propagate from one point to another.

a. Find the propagation delay for a signal traversing the following networks at the speed of light in cable (2.3×10^8 meters/second):

- a circuit board 10 cm
- a room 10 m
- a building 100 m
- a metropolitan area 100 km
- a continent 5000 km
- up and down to a geostationary satellite 2×36000 km

Solution:

To find the propagation delay, divide distance by the speed of light in cable. Thus we have:

- a circuit board $t_{prop} = 4.347 \times 10^{-10}$ seconds
- a room $t_{prop} = 4.3478 \times 10^{-8}$ seconds
- a building $t_{prop} = 4.3478 \times 10^{-7}$ seconds
- a metropolitan area $t_{prop} = 4.3478 \times 10^{-4}$ seconds
- a continent $t_{prop} = 0.02174$ seconds
- up and down to a geostationary satellite $t_{prop} = 0.31304$ seconds

14b. How many bits are in transit during the propagation delay in the above cases, if bits are entering the above networks at the following transmission speeds: 10,000 bits/second; 1 megabit/second; 100 megabits/second; 10 gigabits/second.

Solution:

The number of bits in transit is obtained by multiplying the transmission rate R by the propagation delay:

Distance (m)	10 Kbps	1 Mbps	100 Mbps	10 Gbps
0.1	4.347×10^{-6}	4.347×10^{-4}	0.04347	4.3478
10	4.3478×10^{-4}	0.043478	4.3478	434.780
100	4.3478×10^{-3}	0.43478	43.478	4347.800
100 000	4.3478	434.78	43478	4.3478×10^6
5 000 000	217.4	21740	2174000	2.174×10^8
72 000 000	3130.4	313040	31304000	3.1304×10^9

15. In problem 14, how long does it take to send an L -byte file and to receive a 1-byte acknowledgment back? Let $L=1, 10^3, 10^6,$ and 10^9 bytes.

Solution:

The total time required to send a file and receive an acknowledgment of its receipt is given by:

$$t_{total} = L_{message}/R + L_{ack}/R + 2 * t_{prop} = L_{message}/R + L_{ack}/R + 2*d/c$$

where $L_{message}$ is the message length in bits, L_{ack} is the acknowledgment length in bits, R is the transmission bit rate, d is the distance traversed, and c is the speed of light. The above equation shows that there are two main factors that determine total delay:

1. Message and ACK transmission time, which depends on the message length and the transmission bit rate;
2. Propagation delay, which depends solely on distance.

When the propagation delay is small, message and ACK transmission times determine the total delay. On the other hand, when the bit rate becomes very large, the propagation delay provides a delay component that cannot be reduced no matter how fast the transmission rate becomes.

The tables below show the two main components of the total delay in microseconds. The message transmission time is shown in red and the propagation delay is shown in blue. The entries in the total delay tables are colored according to which delay component is dominant.

Distance (meters)	2 * prop. delay (microseconds)	Message + ACK delay @10 kbps	Message + ACK delay @1 Mbps	message +ACK delay @100 Mbps	message +ACK delay @10 Gbps
		1600	16	0.16	0.0016
		Total delay @10 kbps	total delay @1 Mbps	total delay @100 Mbps	total delay @10 Gbps
0.1	0.00087	1600.00087	16.0008696	0.16087	0.00247
10	0.086957	1600.086957	16.0869565	0.246957	0.088557
100	0.869565	1600.869565	16.8695652	1.029565	0.871165
100000	869.5652	2469.565217	885.565217	869.7252	869.5668
5000000	43478.26	45078.26087	43494.2609	43478.42	43478.26
72000000	626087	627686.9565	626102.957	626087.1	626087

Table 1: Message length = 1 byte

		Message + ACK delay @10 kbps	message + ACK delay @1 Mbps	Message +ACK delay @100 Mbps	message +ACK delay @10 Gbps
		800800	8008	80.08	0.8008
Distance (meters)	2 * prop. delay (microseconds)	Total delay @10 kbps	total delay @1 Mbps	Total delay @100 Mbps	total delay @10 Gbps
0.1	0.00087	800800.0009	8008.00087	80.08087	0.80167
10	0.086957	800800.087	8008.08696	80.16696	0.887757
100	0.869565	800800.8696	8008.86957	80.94957	1.670365
100000	869.5652	801669.5652	8877.56522	949.6452	870.366
5000000	43478.26	844278.2609	51486.2609	43558.34	43479.06
72000000	626087	1426886.957	634094.957	626167	626087.8

Table 2: Message length = 1000 bytes

		Message + ACK delay @10 kbps	message + ACK delay @1 Mbps	message +ACK delay @100 Mbps	message +ACK delay @10 Gbps
		8.00E+08	8.00E+06	80000.08	800.0008
distance (meters)	2 * prop. delay (microseconds)	Total delay @10 kbps	total delay @1 Mbps	total delay @100 Mbps	total delay @10 Gbps
0.1	0.00087	8.00E+08	8.00E+06	80000.08	800.0017
10	0.086957	8.00E+08	8.00E+06	80000.17	800.0878
100	0.869565	8.00E+08	8.00E+06	80000.95	800.8704
100000	869.5652	8.00E+08	8.00E+06	80869.65	1669.566
5000000	43478.26	8.00E+08	8.04E+06	123478.3	44278.26
72000000	626087	8.01E+08	8.63E+06	706087	626887

Table 3: Message length = 10⁶ bytes

		Message + ACK delay @10 kbps	message + ACK delay @1 Mbps	message +ACK delay @100 Mbps	message +ACK delay @10 Gbps
		8.00E+11	8.00E+09	8.00E+07	8.00E+05
Distance (meters)	2 * prop. delay (microseconds)	Total delay @10 kbps	total delay @1 Mbps	total delay @100 Mbps	total delay @10 Gbps
0.1	0.00087	8.00E+11	8.00E+09	8.00E+07	8.00E+05
10	0.086957	8.00E+11	8.00E+09	8.00E+07	8.00E+05
100	0.869565	8.00E+11	8.00E+09	8.00E+07	8.00E+05
100000	869.5652	8.00E+11	8.00E+09	8.00E+07	8.01E+05
5000000	43478.26	8.00E+11	8.00E+09	8.00E+07	8.43E+05
72000000	626087	8.00E+11	8.00E+09	8.06E+07	1.43E+06

Table 4: Message length = 10⁹ bytes

16. Use your web browser to access a search engine and retrieve the article "A Brief History of the Internet," by Leiner, Cerf, Clark, Kahn, Kleinrock, Lynch, Postel, Roberts, and Wolff. Answer the following questions:

Solutions follow questions:

a. Who was J. Licklider, and what was his "Galactic Network" concept?

J. Licklider, of MIT, was the first head of DARPA. In 1962 he envisioned a global network of computers where users could access data or programs from any site.

b. Who coined the term packet?

Donald Davies and Roger Scantlebury of NPL (National Physics Laboratory) in the United Kingdom first coined the term.

c. What (who?) is an IMP?

IMP stands for Interface Message Processor. An IMP is a packet switch and was a key component in the design of the ARPANET.

d. Did the ARPANET use NCP or TCP/IP?

The ARPANET used a host-to-host protocol called Network Control Protocol or NCP.

e. Was packet voice proposed as an early application for Internet?

In the early development of TCP, work on packet voice in particular prompted the separation of TCP into two components: TCP and IP. IP was to provide addressing and forwarding; TCP provided special services such as recovery from packet losses (important for packet voice).

- f. How many networks did the initial IP address provide for?

The initial IP addressing scheme provided for 256 networks.

17. Use your web browser to access a search engine and retrieve the following presentation from the CM 97 conference: "The Folly Laws of Predictions 1.0" by Gordon Bell. Answer the following questions:

Solutions follow questions:

- a. At what rate have processing, storage, and backbone technologies improved from 1950 to 2000? How does this rate compare to advances in telephony?

Processing, storage, and backbone technologies have improved at annual rates of between 20% and 40%. Telephony has advanced at a rate of 17%.

- b. What is Moore's Law?

Microelectronic technologies improve by a factor of 4 every 3 years.

- c. What's the point of making predictions?

To provide a vision and challenge for the future.

- d. What is the difficulty in anticipating trends that have exponential growth?

The trouble with exponential growth is that "you can't see them coming!" At an early stage, the trend is not visible; at a late stage the trend is unmistakable.

- e. Who was Vannevar Bush and why is he famous?

Vannevar Bush posited the "memex" in 1945 that in a sense provided the vision for the World Wide Web.

- f. What is the size in bytes of each frame in this presentation? What is the size in bytes of the audio clip for a typical frame? What is the size of the video clip for a typical scene?

The graphics-only presentation is approximately 1 MB; the graphics and audio presentation is 7 MB; the video presentation for 100 kbps is 14MB.

18. The introduction of new communications services usually impacts other services through substitution. Describe how substitution takes place in the following cases.

Solutions follow questions:

- a. E-mail, facsimile, and postal mail.

E-mail is used for most of the correspondence previously handled by postal mail. Documents sent by facsimile are also transferred using E-mail as attachments. Hardcopies can be scanned for electronic transmission.

- b. E-mail, local, and long-distance phone service.

E-mail is an inexpensive and convenient alternative for most of the communication in which real-time interaction is not essential. Instant-messaging is faster than email and more closely approaches the real-time experience of the telephone.

- c. Cell phone, local, and long-distance phone service.

Cell phone is used for local or long distance calls mostly because users can be reached even if they are not in a specific location such as home or office. The steep drop in the cost of long-distance relative to the cost of cellular airtime has enable “call anywhere” cellular service offerings.

- d. Peer-to-peer file exchange and commercial CD recording.

Peer-to-peer file exchange provides an easy means to access and share commercial recordings. Individual songs rather than entire albums can be acquired. Users can readily arrange their personalized song mixes.

19. Use your web browser to access news website and play a news video clip. Speculate about how the information is being transported over the Internet. How does the quality of the audio and video compare to that of broadcast or cable television?

Solution:

The video information is sent over the Internet, so a packet mode of information transfer is used. The video signal is placed in a stream of packets that is sent from the video server to the PC. A steady stream of video information needs to be supplied to the video player application in the PC. For this reason the player first buffers a certain amount of video information in order to avoid running out of material to play out during periods when the packets encounter congestion and experience excessive delay. The Internet is still subject to congestion and so packets are delayed and lost resulting in poor audio and video quality. Many streaming audio and video applications use TCP instead of UDP.

20. Use your web browser to access the IETF web page (currently at www.ietf.org) and learn the Internet Standards process documented in RFC 2026.

Solutions follow questions:

- a. What are the different types of Requests for Comments (RFC)s?

The status of Internet protocol and service specifications is summarized periodically in an RFC entitled "Internet Official Protocol Standards".

Some RFCs document Internet Standards. These RFCs form the 'STD' subseries of the RFC series.

Some RFCs standardize the results of community deliberations about statements of principle or conclusions about what is the best way to perform some operations or IETF process function. These RFC form the BCP (Best Current Practices) subseries.

Not all specifications of protocols or services for the Internet should or will become Internet Standards or BCPs. Such non-standards track specifications are not subject to the rules for Internet standardization. Non-standards track specifications may be published directly as "Experimental" or "Informational" RFCs at the discretion of the RFC Editor in consultation with the IESG.

b. What is an Internet Draft?

During the development of a specification, draft versions of the document are made available for informal review and comment by placing them in the IETF's "Internet-Drafts" directory, which is replicated on a number of Internet hosts. This makes an evolving working document readily available to a wide audience, facilitating the process of review and revision.

c. What are the difference among Proposed Standard, Draft Standard, and Standard?

Proposed Standard specification is the entry level for the IETF standards track. A Proposed Standard specification is generally stable, has resolved known design choices, is believed to be well-understood, has received significant community review, and appears to enjoy enough community interest to be considered valuable. However, further experience might result in a change or even retraction of the specification before it advances.

A specification from which at least two independent and interoperable implementations from different code bases have been developed, and for which sufficient successful operational experience has been obtained, may be elevated to the "Draft Standard" level. Elevation to Draft Standard is a major advance in status, indicating a strong belief that the specification is mature and will be useful. A Draft Standard may still require additional or more widespread field experience.

A specification for which significant implementation and successful operational experience has been obtained may be elevated to the Internet Standard level. An Internet Standard is characterized by a high degree of technical maturity and by a generally held belief that the specified protocol or service provides significant benefit to the Internet community.

d. Which group in the IETF approves a certain specification for standards-track?

A "standards action" -- entering a particular specification into, advancing it within, or removing it from, the standards track -- must be approved by the the Internet Engineering Steering Group (IESG).

e. How are disputes on working group documents resolved?

A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s), who may involve other members of the Working Group (or the Working Group as a whole) in the discussion.

If the disagreement cannot be resolved in this way, any of the parties involved may bring it to the attention of the Area Director(s) for the area in which the Working Group is chartered. The Area Director(s) shall attempt to resolve the dispute.

If the disagreement cannot be resolved by the Area Director(s) any of the parties involved may then appeal to the IESG as a whole. The IESG shall then review the situation and attempt to resolve it in a manner of its own choosing.

If the disagreement is not resolved to the satisfaction of the parties at the IESG level, any of the parties involved may appeal the decision to the Internet Architecture Board (IAB). The IAB shall then review the situation and attempt to resolve it in a manner of its own choosing.

The IAB decision is final with respect to the question of whether or not the Internet standards procedures have been followed and with respect to all questions of technical merit.

Solutions to Chapter 2 (Note: solution to Problem 62 to be added)

1. Explain how the notion of layering and internetworking make the rapid growth of applications such as the World Wide Web possible.

Solution:

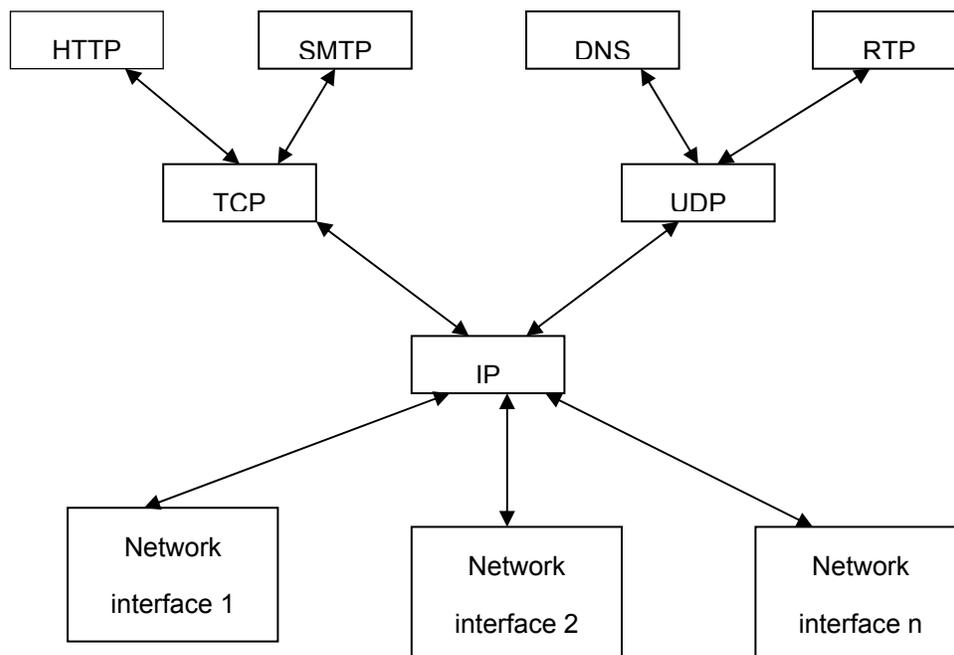
Internetworking allows many component networks each with different underlying technology and operation to work together and form one large network. As new network technologies are introduced they can be readily incorporated into the Internet. This provides the ubiquitous connectivity for applications like WWW.

The layering concept hides the specific underlying network technology from the upper layers and provides a common networking platform. Using the communication service provided by the layers below, new applications can be introduced independently and at a rapid rate.

2. (a) What universal set of communication services is provided by TCP/IP?

Solution:

The TCP/IP protocol stack provides two basic types of communications services through its two transport layer protocols: TCP provides reliable connection-oriented transfer of a byte stream; UDP provides for best-effort connectionless transfer of individual messages. TCP/IP provides with globally unique logical addressing that enables machines connected to the Internet to access these two services. The IP addressing scheme is very scalable because of its hierarchical structure.



2. (b) How is independence from underlying network technologies achieved?

Solution:

The two basic communications services provided by TCP and UDP are built on the connectionless packet transfer service provided by the Internet Protocol (IP). Many network interfaces are defined to

support IP. The salient part of the above figure is that all of the higher layer protocols access the network interfaces through IP. This is what provides the ability to operate over multiple networks.

2. (c) What economies of scale result from (a) and (b)?

Solution:

Once a network interface for IP is defined for a given network technology, then hosts connected using the given network technology can connect to the Internet. This allows the reach of the Internet to grow rapidly, leveraging multiple coexisting networks technologies. Thus investment in new network technologies extends the reach of the Internet.

3. What difference does it make to the network layer if the underlying data link layer provides a connection-oriented service versus a connectionless service?

Solution:

If the data link layer provides a connection-oriented service to the network layer, then the network layer must precede all transfer of information with a connection setup procedure. If the connection-oriented service includes assurances that frames of information are transferred correctly and in sequence by the data link layer, the network layer can then assume that the packets it sends to its neighbor traverse an error-free pipe.

On the other hand, if the data link layer is connectionless, then each frame is sent independently through the data link, probably in unconfirmed manner (without acknowledgments or retransmissions). In this case the network layer cannot make assumptions about the sequencing or correctness of the packets it exchanges with its neighbors.

The Ethernet local area network provides an example of connectionless transfer of data link frames. The transfer of frames using "Type 2" service in Logical Link Control (discussed in Chapter 6) provides a connection-oriented data link control example.

4. Suppose transmission channels become virtually error-free. Is the data link layer still needed?

Solution:

The data link layer is still needed for framing the data and for flow control over the transmission channel. In a multiple access medium such as a LAN, the data link layer is required to coordinate access to the shared medium among the multiple users.

5. Why is the transport layer not present inside the network?

Solution:

Some of the functions provided by the transport layer can be provided inside the networks, but other functions cannot. For example, the transport layer provides functions at the end-system to compensate for the limitations and impairments of the network layer, in order to meet requirements (e.g. QoS) of the upper layer. For example in TCP/IP, IP provides only best effort service. To provide the reliable service required by some applications - that is, to compensate for the shortcomings of best effort service - TCP establishes connections and implements error control on an end-to-end basis. One can imagine that a service provider could incorporate this error control function at the edge of its network. On the other hand, one of the main purposes of the transport layer is to allow multiple processes in the end systems to share a network service. This cannot be achieved inside the network.

6. Which OSI layer is responsible for the following?

(a) Determining the best path to route packets.

The network layer is concerned with the selection of paths across the network.

(b) Providing end-to-end communications with reliable service.

The transport layer is concerned with providing reliable service on an end-to-end basis across the network.

(c) Providing node-to-node communications with reliable service.

The data link layer provides for the reliable transfer of information between adjacent nodes in a network.

7. Should connection establishment be a confirmed service or an unconfirmed service? What about data transfer in a connection-oriented service? Connection release?

Solution:

In general, the establishment of a connection needs to be confirmed before information transfer can commence across a connection. Therefore connection establishment should be a confirmed service.

A connection-oriented service is usually reliable so confirmation of data delivery is not necessary. In certain situations, however, it is possible that the transfer across a connection is not reliable; in this case confirmation of correct data transfer may be required.

In general it is desirable that the release of a connection be confirmed by the parties involved. We will see in Chapter 8, section 5, that sometimes it is not easy to confirm that a connection has been closed. Consequently, many protocols attempt to confirm the closing of a connection several times, and then give up and simply stop transmitting.

8. Does it make sense for a network to provide a confirmed, connectionless packet transfer service?

Solution:

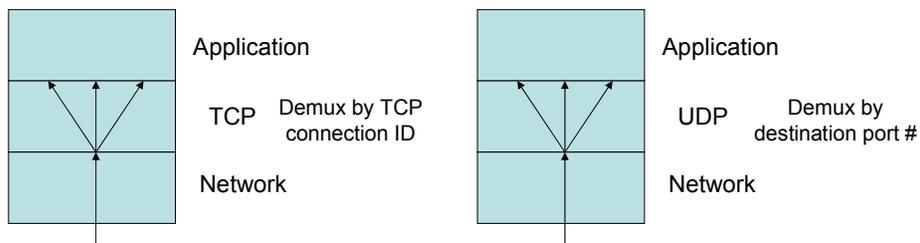
Yes. Connectionless packet transfer is often unreliable, that is, packets may be lost or discarded inside a network. Certain applications, for example, signaling in connection setup, require confirmation to acknowledge the receipt of packets.

9. Explain how the notion of multiplexing can be applied at the data link, network, and transport layers. Draw a figure that shows the flow of PDUs in each multiplexing scheme.

Solution:

Transport Layer: Multiple application layer processes can share the service provided by UDP. When a UDP PDU arrives from the network layer, the destination port number in the PDU is used to deliver the SDU to the appropriate application layer process. Multiple application layer processes also share the service provided by TCP. In this case, when a TCP segment arrives, the TCP connection ID, consisting of (source port #, source IP address, destination port #, destination IP address), is used to determine which application process to deliver the SDU to.

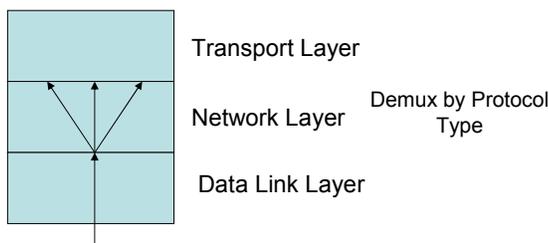
Transport Layer De-multiplexing



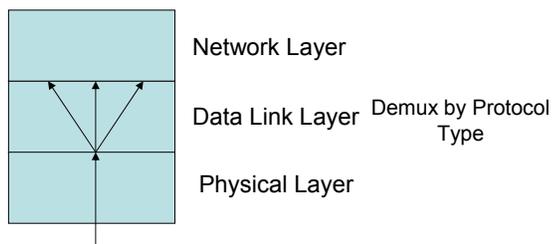
Network Layer: The packet transfer service provided by IP can be used by all transport layers operating in a machine. Each transfer layer passes SDUs to the IP layer which prepares IP packets with appropriate source and destination IP addresses for transfer across the Internet. Upon receiving an IP packet, a machine examines the protocol type field to determine which transport layer service to deliver the SDU to. We can also view all transport layer PDUs as sharing the IP packet transfer service between a source machine and a destination machine.

Data Link Layer: Network layer packets from different protocols (IP, IPX, Appletalk, etc) can share a data link (such as PPP or Ethernet). We can also view packet flows that traverse a data link between two routers as sharing the link.

Network Layer De-multiplexing



Data Link Layer De-multiplexing



10. Give two features that the data link layer and transport layer have in common. Give two features in which they differ. Hint: Compare what can go wrong to the PDUs that are handled by these layers.

Solution:

Features they have in common:

- Both layers can provide recovery from transmission errors.
- Both layers can provide flow control.
- Both layers can support multiplexing.

Features in which they differ:

- The transport layer is end to end and involves the interaction of peer processes across the network. The data link layer involves the interaction of peer-to-peer processes that are connected directly. In general, the time that elapses in traversing a data link is much smaller than the time traversing a network, where packets can become trapped in temporary routing loops. Consequently, transport layer protocols must be able to deal with out-of-sequence PDUs and a much larger backlog of PDUs than data link layers.
- The data link layer is concerned with framing and the transport layer is not.
- The data link layer may be concerned with medium access control, the transport layer does not have this concern.

11(a). Can a connection-oriented, reliable message transfer service be provided across a connectionless packet network? Explain.

Solution:

Yes. To provide a connection-oriented service, the transport layer can establish a logical connection across the connectionless packet network by setting up state information (for example, packet sequence number) at the end systems. During the connection setup, the message is broken into separate packets, and each packet is assigned a sequence number.

Using the sequence numbers, the end-system transport-layer entities can acknowledge received packets, determine and retransmit lost packets, delete duplicate packets, and rearrange out-of-order packets. The original message is reassembled as packets arrive at the receiving end.

For example, TCP provides a connection-oriented reliable transfer service over IP, a connectionless packet transfer service.

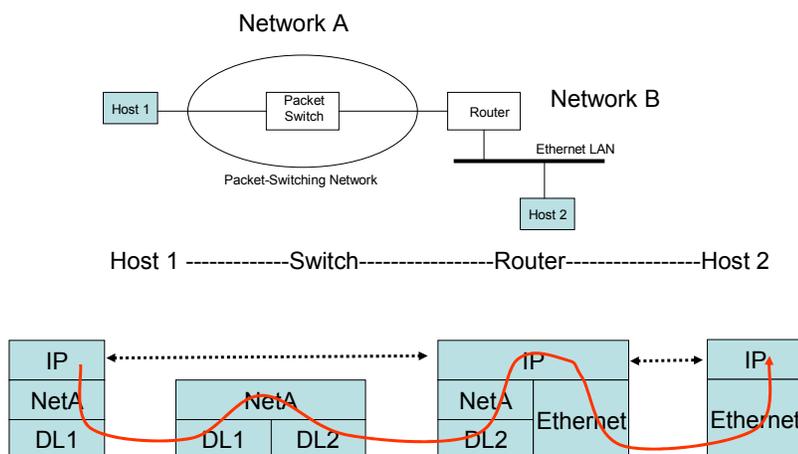
11b. Can a connectionless datagram transfer service be provided across a connection-oriented network?

Solution:

Yes. The connectionless datagram transfer service can be implemented by simply setting up a connection across the network each time a datagram needs to be transferred. Alternatively, all nodes can have permanent connections to a “connectionless server” that has the function of relaying datagrams in connectionless fashion.

12. An internet path between two hosts involves a hop across network A, a packet-switching network, to a router and then another hop across packet-switching network B. Suppose that packet switching network A carries the packet between the first host and the router over a two-hop path involving one intermediate packet switch. Suppose also that the second network is an Ethernet LAN. Sketch the sequence of IP and non-IP packets and frames that are generated as an IP packet goes from host 1 to host 2.

Solution:



The IP layer in Host 1 generates an IP packet addressed to the destination host on the destination network and sends it to the router. The network interface in the host encapsulates the IP packet into the packet PDU used by network A. This packet is encapsulated in a frame that traverses data link 1 to the packet switch. The packet is recovered and then forwarded inside a frame along data link 2. The data link at the router recovers the Network A packet, and the IP network interface at the router recovers the IP packet and determines that the next hop is on Network B. The router encapsulates the IP packet into an Ethernet frame, puts the host 2 Ethernet physical address in the frame and sends it to the LAN. The Ethernet card on the host captures the frame and extracts the IP packet and passes it to the host.

13. Does Ethernet provide connection-oriented or connectionless service?

Solution:

Ethernet provides connectionless transfer service of information frames.

14. Ethernet is a LAN so it is placed in the data link layer of the OSI reference model.

(a) How is the transfer of frames in Ethernet similar to the transfer of frames across a wire? How is it different?

The transfer of frames in Ethernet occurs directly over a transmission medium and in this sense is similar to direct transmission over a wire. The sequence of frames into Ethernet arrive in the same order they are transmitted. However multiple stations can transmit in Ethernet which differs from direct transmission over a wire.

(b) How is the transfer of frames in Ethernet similar to the transfer of frames in a packet-switching network? How is it different?

Ethernet supports the transfer of frames among multiple end systems and in this sense is similar to a packet switching network. Ethernet does not involve routing which is a feature of packet switching. Ethernet depends on broadcasting and/or bridging which differs from packet networks.

15. Suppose that a group of workstations is connected to an Ethernet LAN. If the workstations communicate only with each other, does it make sense to use IP in the workstations? Should the workstations run TCP directly over Ethernet? How is addressing handled?

Solution:

Ethernet supports the exchange of frames between stations and can support the direct exchange of information. Using Ethernet without IP would result in an inflexible and difficult to manage system. Ethernet addresses are fixed and tables need to be changed whenever a machine is moved, while IP addresses are logical and can be changed whenever a machine is moved. A TCP connection uses the IP addresses in its connection ID so Ethernet addresses could not be used.

16. Suppose two Ethernet LANs are interconnected by a box that operates as follows. The box has a table that tells it the physical addresses of the machines in each LAN. The box listens to frame transmissions on each LAN. If a frame is destined to a station at the other LAN, the box retransmits the frame onto the other LAN, otherwise the box does nothing.

Solutions follow questions:

- a. Is the resulting network still a LAN? Does it belong in the data link layer or the network layer?

The resulting network is a local area network that has been extended. The extended LAN transfers frames, and so it still belongs in the data link layer.

- b. Can the approach be extended to connect more than two LANs? If so, what problems arise as the number of LANs becomes large?

Yes, more than two LANs can be connected using the above approach to form an extended LAN. As the number of LANs becomes large, the number of physical addresses stored in the bridge grows and becomes unmanageable. Each time a machine is added the addresses in all the boxes need to be updated. Serious problems arise if boxes are connected so that loops can occur.

17. Suppose all laptops in a large city are to communicate using radio transmissions from a high antenna tower. Is the data link layer or network layer more appropriate for this situation?

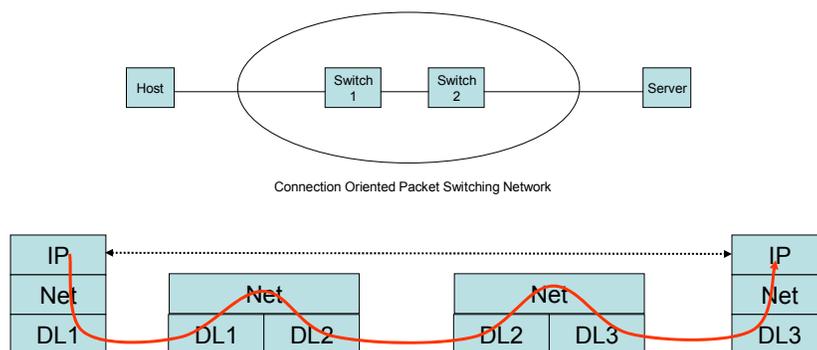
Solution:

The data link layer is concerned with the transfer of frames of information across a single hop. The network layer involves the transfer of information across a network using multiple hops per path in general. The connection from a radio antenna to the laptops is direct, and thus a data link layer protocol is more suitable for this situation.

Now suppose the city is covered by a large number of small antennas covering smaller areas. Which layer is more appropriate?

A number of areas each covered by small antennas can be interconnected using the "bridging" approach of problem 16, which remains in the data link layer. However, the network layer may be more appropriate because it provides for the transfer of data in the form of packets across the communication network. A key aspect of this transfer is the routing of the packets from the source machine to the destination machine, typically traversing a number of transmission link and network nodes where routing is carried out.

18. Suppose that a host is connected to a connection-oriented packet-switching network and that it transmits a packet to a server along a path that traverses two packet switches. Suppose that each hop in the path involves a point-to-point link, that is, a wire. Show the sequence of network layer and data link layer PDUs that are generated as the packet travels from the host to the server.

Solution:

Assume that a network connection has already been set up between the host machine and the network machine. When the host generates an IP packet for transfer to the server, the IP packet will be transferred using the network connection as follows.

- The IP packet is encapsulated into a network packet that has a connection ID in its header. The packet may then be encapsulated into a frame that traverses data link 1 and arrives at switch 1.
- The network packet is recovered from the data link 1 frame. The connection ID in the packet is used to determine the outgoing port from switch 1. The connection ID may need to be mapped into a corresponding connection ID over data link 2. The packet is encapsulated into a frame that traverses data link 2.
- The network packet is recovered from the data link 2 frame. The connection ID in the packet determines the outgoing port from switch 1 and the next connection ID. The packet is encapsulated into a frame that traverses data link 3.
- The network packet is recovered from the data link 3 frame. The connection ID in the arriving packet indicates that this is the destination node. The IP packet is recovered.

The connection-oriented network in this example could correspond to ATM or to frame relay.

19. Suppose an application layer entity wants to send an L -byte message to its peer process, using an existing TCP connection. The TCP segment consists of the message plus 20 bytes of header. The segment is encapsulated into an IP packet that has an additional 20 bytes of header. The IP packet in turn goes inside an Ethernet frame that has 18 bytes of header and trailer. What percentage of the transmitted bits in the physical layer correspond to message information, if $L = 100$ bytes, 500 bytes, 1000 bytes?

Solution:

TCP/IP over Ethernet allows data frames with a payload size up to 1460 bytes. Therefore, $L = 100$, 500 and 1000 bytes are within this limit.

The message overhead includes:

- TCP: 20 bytes of header
- IP: 20 bytes of header
- Ethernet: total 18 bytes of header and trailer.

Therefore

$L = 100$ bytes, $100/158 = 63\%$ efficiency.

$L = 500$ bytes, $500/558 = 90\%$ efficiency.

$L = 1000$ bytes, $1000/1058 = 95\%$ efficiency.

20. Suppose that the TCP entity receives a 1.5 megabyte file from the application layer and that the IP layer is willing to carry blocks of maximum size 1500 bytes. Calculate the amount of overhead incurred from segmenting the file into packet-sized units.

Solution:

$1500 - 20 - 20 = 1460$ bytes

$1.5 \text{ Mbyte} / 1460 \text{ byte} = 1027.4$, therefore 1028 blocks are needed to transfer the file.

Overhead = $((1028 \times 1500 - 1.5\text{M})/1.5\text{M}) \times 100 = 2.8\%$

21. Suppose a TCP entity receives a digital voice stream from the application layer. The voice stream arrives at a rate of 8000 bytes/second. Suppose that TCP arranges bytes into block sizes that result in a total TCP and IP header overhead of 50 percent. How much delay is incurred by the first byte in each block?

Solution:

Assume the stream is segmented as shown below, where the white cells represent data and the shaded cells represent the TCP header overhead.



Therefore, block size = 80 bytes and the payload size = 40 bytes.

Assume zero processing delay due to data arrangement and segmenting.

The delay incurred by the first byte of each block = $40/8000 = 0.5$ ms.

22. How does the network layer in a connection-oriented packet-switching network differ from the network layer in a connectionless packet-switching network?

Solution:

The network layer in connection-oriented networks maintains state information about every connection. It can allocate resources at the switches through admission control. The network layer in connectionless networks has no knowledge of "connections", and instead deals independently with each packet.

The network layer in connection-oriented networks performs routing on a per connection basis. Each packet is routed based on a connection identifier of some sort and packets of the same connection have the same identifier value. In a connectionless network, routing is performed on per packet basis; each packet is routed independently based on information carried in the packet header, for example, the destination address.

In connection-oriented networks, the network layer forwarding table is set up by a signaling procedure during the connection establishment. In connectionless networks, the routers may execute a distributed algorithm to share network state information and dynamically calculate the routing table continuously.

In case of failure, the connection must be re-established in connection-oriented networks, whereas in connectionless networks, the packets are re-routed. The network layer in connectionless networks is more robust against failures.

Summary of differences:

Connection-oriented	Connectionless
Maintain state information about every connection	No knowledge of the "connection"
Allocate resources to connections at switches	No resource allocation
Admission control	No admission control
Per connection routing	Per packet routing
Route packet based on identifier	Route packet based on destination address.
Forwarding table specifies the output port and outgoing identifier value as function of the incoming identifier value	Routing table specifies the output port depending on the destination address
Forwarding table set up by signaling during connection establishment.	Router executes distributed algorithm to share network state information and dynamically calculate the routing table
Connection must be re-established in cases of failure	Packets are rerouted around failures, robust against failures

23. Identify session layer and presentation layer functions in the HTTP protocol.

Solution:*Presentation layer functions:*

The request message and the response message headers include information about the content type of the documents (e.g. text/html, image/gif).

Session layer functions:

The HTTP protocol defines the client/server interaction in three steps:

1. Client sends the request for a file
2. Server replies with the file or error message if file is not found.
3. Server closes the TCP connection after some timeout period.

24. Suppose we need a communication service to transmit real-time voice over the Internet. What features of TCP and what features of UDP are appropriate?

Solution:

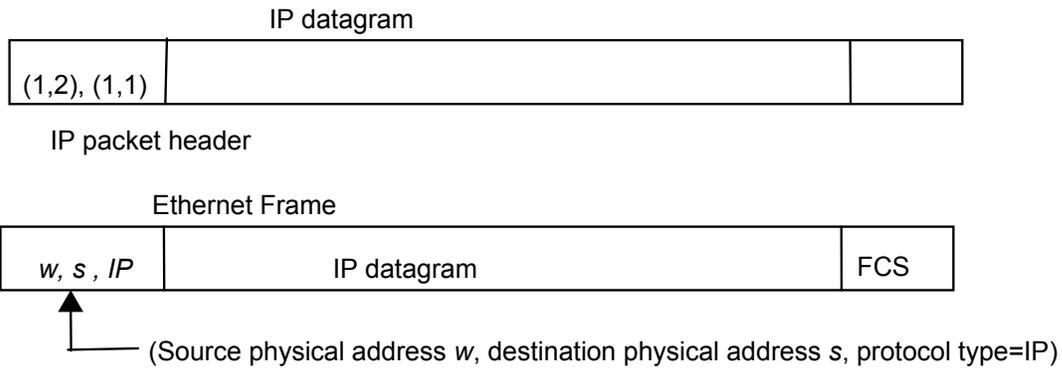
TCP is desirable in that it provides a connection for the transfer of a stream of information, which characterizes a digital voice stream. However, to provide reliable service TCP uses acknowledgments and retransmissions that result in packet delay and jitter that can not be tolerated by real-time traffic.

UDP provides connectionless service and delivers packets quickly. In case of packet loss, UDP does not provide retransmission, but some degree of packet loss can be tolerated by voice.

25. Consider the end-to-end IP packet transfer examples in Figure 2.15. Sketch the sequences of IP packets and Ethernet and PPP frames that are generated by the three examples of packet transfers: from the workstation to the server; from the server to the PC, and from the PC to the server. Include all relevant header information in the sketch.

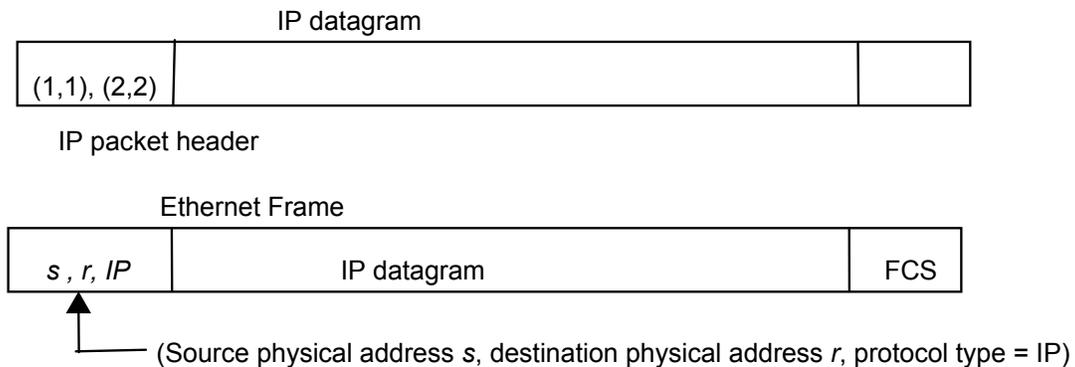
Solution:

Workstation to Server:

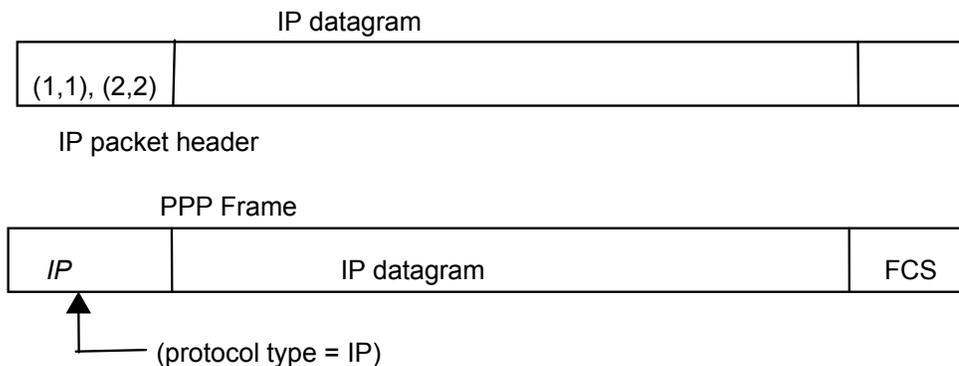


The Ethernet frame is broadcast over the LAN. The server's NIC card recognizes that the frame is intended for its host, so it captures the frame and examines it. It finds that the protocol type is set to IP, so it passes the IP datagram up to the IP entity.

Server to PC:



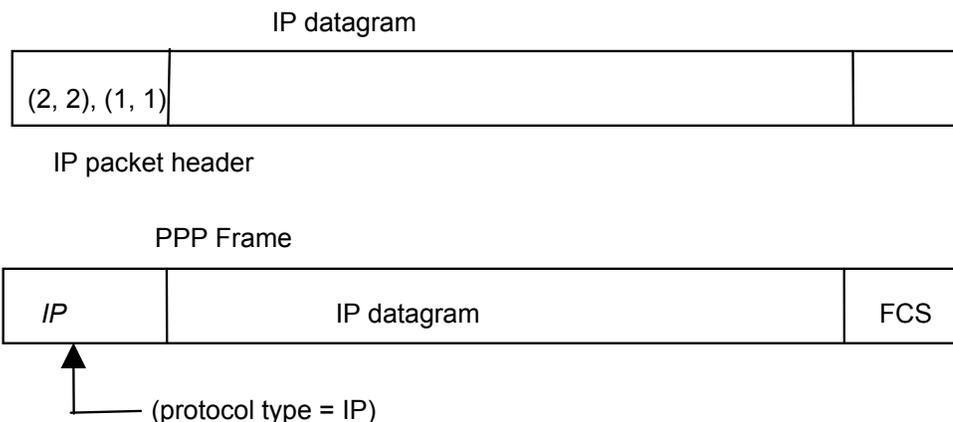
The Ethernet frame is broadcast over the LAN. The router examines frame and passes IP datagram to its IP entity which discover that the IP datagram is not for itself, but is to be routed on. The routing tables at the router show that the machine with address (2,2) is connected directly on the other side of the point-to-point link. The router encapsulates the IP datagram in a PPP frame.



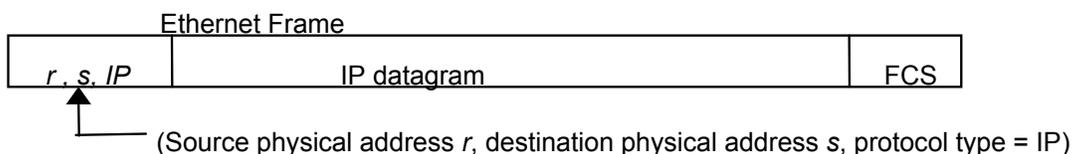
The PPP receiver at the PC receives the frame, checks the protocol type field and passes the IP datagram to its IP entity.

PC to Server:

The PC IP entity generates the IP packet shown below. The PPP transmitter at the PC encapsulates the IP packet into a PPP frame sends it to the point-to-point link. There's no need for a physical address specification



The router examines the PPP frame and passes the IP datagram to its IP entity which discovers that the IP datagram is not for itself, but is to be routed on. The routing table at the router shows that the machine with address (1,1) is connected in the other side of the Ethernet network. The router then encapsulates the IP datagram into an Ethernet frame that is broadcast in the LAN.



The server's NIC card recognizes that the frame is intended for its host, so it captures the frame and examines it. It finds that the protocol type is set to IP, so it passes the IP datagram up to the IP entity.

26. Suppose a user has two browser applications active at the same time, and suppose that the two applications are accessing the same server to retrieve HTTP documents at the same time. How does the server tell the difference between the two applications?

Solution:

A client application generates an ephemeral port number for every TCP connection it sets up. An HTTP request connection is uniquely specified by the five parameters: (TCP, client IP address, ephemeral port #, server IP address, 80). The two applications in the above situations will have different ephemeral port #s and will thus be distinguishable to the server.

27. Consider the operation of non-persistent HTTP and persistent HTTP.
- (a) In non-persistent HTTP (version 1.0): Each client-server interaction involves setting up a TCP connection, carrying out the HTTP exchange, and closing the TCP connection. Let T be the time that elapses from when a packet is sent from client to server to when the response is received. Find the rate at which HTTP exchanges can be made using non-persistent HTTP.

- (b) In persistent HTTP (version 1.1) the TCP connection is kept alive. Find the rate at which HTTP exchanges can be made if the client cannot send an additional request until it receives a response for each request.
- (c) Repeat part (b) if the client is allowed to pipeline requests, that is, it does not have to wait for a response before sending a new request.

Solution:

(a) Each HTTP exchange involves: 1. a three-way handshake to set up the TCP connection; 2. an HTTP request-response interaction; and 3. a TCP close. The client can send its request after the first two handshakes in part 1 (which takes up T seconds). The request and response then take an additional T second. A new request can be initiated with an associated new TCP connection even while the previous TCP connection is being closes. Thus a maximum of one HTTP exchange per $2T$ seconds is possible.

(b) Since each exchange is completed in T seconds, after the connection is setup, the exchange rate is $1/T$.

(c) The rate depends on how long it takes to send a request and how late it takes to compose a response. Considering the maximum of these to be t seconds, exchange rate can be up to $1/t$.

28. What is the difference between a physical address, a network address, and a domain name?

Solution:

The physical address is the unique hardware address that identifies an interface of a machine on a physical network such as a LAN. Physical addresses are used in the data link layer.

A network address is a machine's logical address on a network. The network address is used in the network layer. The network address used on the Internet is the IP address.

Domain names are used as an aid to identify hosts and networks in the Internet, since names are easier to remember than numbers. The DNS system is used to translate between domain names and IP addresses. The domain name for the network address 128.100.132.30 is toronto.edu.

29. Explain how a DNS query proceeds if the local name server does not have the IP address for a given host when the following approaches are used. Assume an example where four machines are involved in ultimately resolving a given query.

- (a) When a machine B cannot resolve an address in response to a query from A, machine B sends the query to another machine in the chain. When B receives the response, it forwards the result to B.
- (b) When a machine B cannot resolve an address in response to a query from A, machine B sends a DNS reply to A with the IP address of the next machine in the chain, and machine A contacts that machine.

Solution:

(a) Host A sends a query to a name server B. B cannot resolve an address, therefore sends the query to C. C cannot resolve an address either, and send the query to D. Similarly, D cannot resolve an address and sends the query to E, where finally an address is resolved and returned to D. D replies the address to C, C replies it to B, and finally B passes it to the host. In this scenario each server should remember the state of the query and its source.

(b) Host A sends a query to name server B. B cannot resolve an address, replies to A with the IP address of C. Host A send a query to C this time. C cannot resolve an address, and replies to A with the IP address of D. A sends a query to D. D cannot resolve an address, and replies with the IP address of E. A sends a query to E, E finally resolves an address and returns it to A. In this scenario

the servers do not need to remember the queries.

30. Suppose that the DNS system used a single centralized database to handle all queries. Compare this centralized approach to the distributed approach in terms of reliability, throughput (volume of queries/second that can be processed), query response delay, and maintainability.

Solution:

The centralized approach suffers from scalability, reliability, and security issues. Given the central role of DNS, a centralized system needs to be ultra-reliable whereas reliability is provided by redundancy in the distributed approach. Maintainability is easier in the distributed approach because individual servers can be brought down without interrupting service. The response delay will increase dramatically in the centralized approach beyond a certain level of traffic whereas the distributed approach can be scaled up by adding more servers. Security is a serious issue in both approaches but the centralized system is more vulnerable to complete service interruption.

31. What is wrong with the following methods of assigning host id addresses?

Solutions follow questions:

(a) Copy the address from the machine in the next office.

There is an address conflict. The host id must be unique to each machine.

(b) Modify the address from the machine in the next office.

The resulting address may be an existing address and result in address conflict, or the address may not be recognizable by the routers.

(c) Use an example from the vendor's brochure.

The address has different network and subnetwork ids, and is not recognized by the routers.

32. Suppose a machine is attached to several physical networks. Why does it need a different IP address for each attachment?

Solution:

The IP address dictates through which network the packets are sent to and from the machine. Therefore each network connection must have a different address. The use of the IP address for the two attachments could also confuse nearby routers and introduce routing loops.

33. Suppose a computer is moved from one department to another. Does the physical address need to change? Does the IP address need to change? Does it make a difference if the computer is a laptop?

Solution:

The physical address does not change. It is globally unique to the computer's NIC card.

The IP address may need to be changed to reflect a new subnetwork id and host id.

The situation is the same for laptops.

34. Suppose the population of the world is 6 billion, and that there is an average of 1000 communicating devices per person. How many bits are required to assign a unique host address to each communicating device? Suppose that each device attaches to a single network and that each network on average has 10000 devices. How many bits are required to provide unique network ids to each network?

Solution:

$$\log_2 (6 \times 10^9 \times 10^3) = 42.44$$

⇒ 43 bits are required to assign a unique host address to each communicating device.

$$\log_2 ((6 \times 10^9 \times 10^3) / 10,000) = 29.2$$

⇒ 30 bits are required to provide unique network ids to each network.

35. Can IP be used to run a homogeneous packet-switching network, that is, a network with identical packet switches interconnected with point-to-point links?

Solution:

Yes. For a homogeneous packet-switching network, the network interface function in each switch will be the same and will operate over the point-to-point links.

36. Is it possible to build a homogeneous packet-switching network with Ethernet LANs interconnecting the packet switches? If so, can connection-oriented service be provided over such a network?

Solution:

A homogeneous packet-switching network can be built where Ethernet LANs are used to interconnect packet switches. In the most common example the packet switches are routers running IP.

A connection-oriented service can be provided over such a packet-switching network in several ways. If the packet-switching network operates in connectionless manner, then additional functions can be added at the ingress and egress to the network to provide a connection-oriented transfer service. Alternatively, the packet-switching network itself could be designed to operate in connection-oriented fashion. In this case the packet switches might use a layer above Ethernet to ensure reliable and sequenced transfer of frames between packet switches. Such a packet-switching network can provide connection-oriented service.

37. In telephone networks one basic network is used to provide worldwide communications. In the Internet a multiplicity of networks are interconnected to provide global connectivity. Compare these two approaches, namely, a single network versus an internetwork, in terms of the range of services that can be provided and the cost of establishing a worldwide network.

Solution:

At a national level, the telephone network has a hierarchical structure. Each new telephone line should be connected to a local telephone center which is hierarchically connected to other telephone centers in different levels such that full connectivity can be achieved. At an international level national networks are interconnected through gateways. Scaling this structure is costly, primarily because the network must keep track of every single connection.

In the case of internetworking, new networks can be connected to existing networks through routers. The Internet can be scaled up by building larger routers that use higher speed links. The connectionless nature of the Internet keeps the operation of these routers simpler and hence more readily scalable.

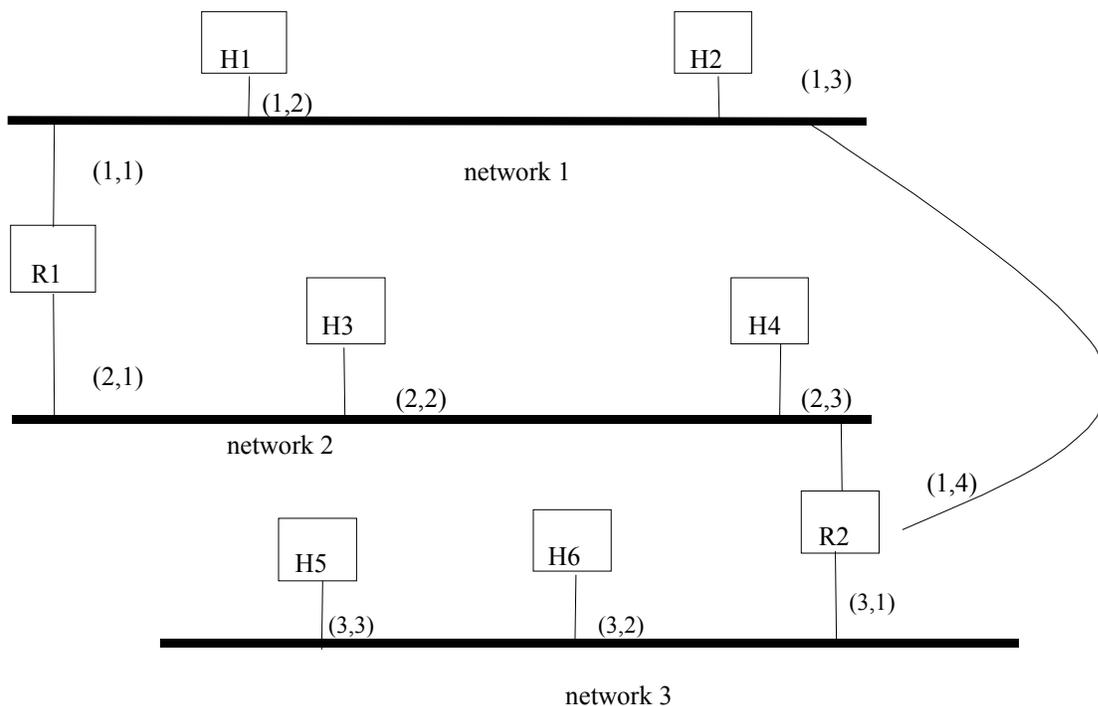
The Internet can support existing and new services through the well-defined TCP and UDP communications services. The telephone network on the other hand, provides a limited number of services and new services cannot easily be introduced.

38. Consider an internetwork architecture that is defined using gateways/routers to communicate across networks but that uses a connection-oriented approach to packet switching. What functionality is required in the routers? Are there any additional constraints imposed on the underlying networks?

Solution:

The routers must be able to setup and release connections across the internetwork. A connection must be established so that routers can forward packets along a path in the network. The underlying networks may or may not operate in connection-oriented fashion. Therefore it is still possible that packets may get out of sequence while traversing a given network. If we require that packets always traverse the end-to-end path in order, then either the underlying networks must be connection-oriented or protocols must operate above each network to ensure sequenced transfer of information.

39. The internet below consists of three LANs interconnected by two routers. Assume that the hosts and routers have the IP addresses as shown.



(a) Suppose that all traffic from network 3 that is destined to H1 is to be routed directly through router R2, and all other traffic from network 3 is to go to network 2. What routing table entries should be present in the network 3 hosts and in R2?

H5		H6		R2	
Destination	Next hop	Destination	Next hop	Destination	Next hop
Default	(3,1)	default	(3,1)	(1,2)	(1,4)
				(1,0)	(2,1)
				(2,0)	(2,4)
				(3,0)	(3,1)

- (b) Suppose that all traffic from network 1 to network 3 is to be routed directly through R2. What routing table entries should be present in the network 1 hosts and in R2?

R2		R1		H1		H2	
Destination	Next hop						
(1,0)	(2,1)	(1,0)	(1,1)	(1,0)	(1,2)	(1,0)	(1,3)
(2,0)	(2,4)	(2,0)	(2,1)	(2,0)	(1,1)	(2,0)	(1,1)
(3,0)	(3,1)	(3,0)	(2,4)	(3,0)	(1,4)	(3,0)	(1,4)

40. Explain why it is useful for application layer programs to have a "well-known" TCP port number?

Solution:

The TCP layer entity uses the port number to determine which application program the packets belong to. In the TCP connection setup process it is very convenient to have a unique well-known port number, otherwise some protocol or procedure would be required to find the desired number.

41. Use a web browser to connect to cnn.com. Explain what layers in the protocol stack are involved in the delivery of the video newscast.

Solution:

The delivery of a video newscast over the Internet involves the transfer of a long stream of information without assurance of delivery or protection from data loss. The current practice is to use TCP to send the video stream. This can be checked by doing an Ethereal packet capture while accessing a video clip. By observing the video display window it is apparent that some sort of protocol particular to video streaming is in operation. After the connection request, the video display application buffers a certain amount of information before initiating display. This buffering is done in an attempt to ensure a steady supply of information to feed the audio and video decoder. Running out of information would result in a freezing of the picture image and loss of the audio signal. The protocols used in video streaming are discussed in Chapter 10.

42. Use a web browser to connect to an audio program, say www.rollingstone.com/radio/ (Rolling Stone Radio) or www.cbc.ca (CBC Radio). Explain what layers in the protocol stack are involved here. How does this situation differ from the delivery of video in problem 41?

Solution:

The delivery of audio information is quite similar to that of video information. A significant difference is that the volume of information that has to be transferred for audio is much less than that required by video. In addition the video application must be concerned with the synchronization of the display of audio and video information, otherwise "lip synch" will not be achieved. For this reason, many newspaper web sites prefer to combine images of a speaker along with audio commentary instead of full-blown video.

43. Which of the TCP/IP transport protocol (UDP or TCP) would you select for the following applications: packet voice, file transfer, remote login, multicast communication (i.e., multiple destinations).

Solution:

Packet Voice - This example involves the transfer of a stream of information in real time across the network. At first, it may appear that TCP is suitable because of its connection orientation. However the acknowledgment and retransmission mechanisms in TCP introduce too much delay in the transfer of packets, and so UDP is the preferred approach to transferring a real-time voice stream across the network.

File Transfer - In general, file transfer requires reliable transfer and so TCP is preferred.

Remote Login - TCP is preferable because it provides for the reliable transfer of the stream of keystrokes that forms the basis for a remote login application.

Multicast Communication - In multicast services, a source sends information to a subset of destinations attached to the network. It is easy to imagine multicast applications that require reliable transfer of a stream of information to a set of destinations, and multicast applications that require only best effort transfer of individual messages. Therefore neither TCP nor UDP is preferred. A more pertinent point is that providing reliable multicast stream transfer service is quite difficult to implement, and TCP is not designed for this.

44. (a) Use the Telnet program to send an e-mail by directly interacting with your local mail server. The SMTP server has port 25. You can find the list of commands for the SMTP protocol in RFC 2821, which can be downloaded from www.ietf.org.

Solution:

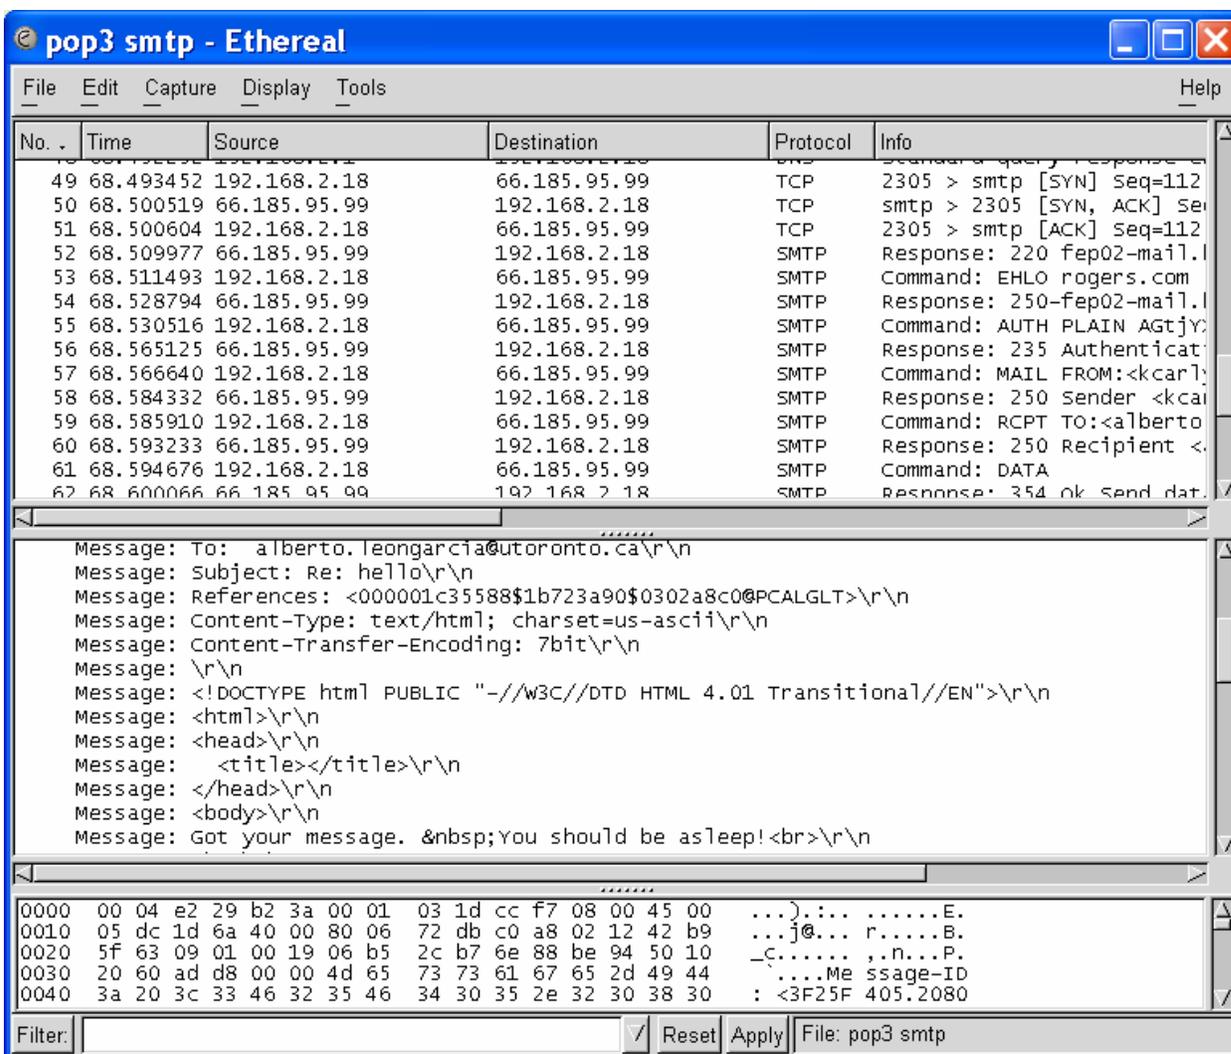
```
telnet <domain name> 25
```

(follow Table 2.3 to send the e-mail)

(b) Use Ethereal to capture and analyze the sequence of messages exchanged. Identify the various types of addresses for Ethernet, IP, and TCP PDUs. Examine the data in the Telnet messages to determine whether the login name and password are encrypted.

Solution:

The following screen capture shows some of the SMTP messages exchanged in the sending of email. The sequence begins with the TCP connection setup to the SMTP server. This is followed by authentication, and then by a command to send an email.



45. (a) Use the Telnet program to retrieve an e-mail by directly from your local mail server. The POP3 server has port 110. You can find the list of commands for the POP3 protocol in RFC 1939, which can be downloaded from www.ietf.org.

Solution:

Many servers do not allow access to email using telnet. The Ethereal screen capture below shows the first few lines in the interaction with a POP3 server to retrieve email. The sequence of frames is followed by an authentication phase with an exchange of user name and password.

pop3 smtp - Ethereal

File Edit Capture Display Tools Help

No.	Time	Source	Destination	Protocol	Info
5	24.619118	192.168.2.18	66.185.95.101	TCP	2303 > pop3 [SYN] Seq=112492730 Ack=0 win=8192 Le
6	24.634432	66.185.95.101	192.168.2.18	TCP	pop3 > 2303 [SYN, ACK] Seq=1939399055 Ack=1124927
7	24.634517	192.168.2.18	66.185.95.101	TCP	2303 > pop3 [ACK] Seq=112492731 Ack=1939399056 wi
8	24.683232	66.185.95.101	192.168.2.18	POP	Response: +OK InterMail POP3 server ready.
9	24.684800	192.168.2.18	66.185.95.101	POP	Request: AUTH

Frame 8 (88 bytes on wire, 88 bytes captured)

- Ethernet II, Src: 00:04:e2:29:b2:3a, Dst: 00:01:03:1d:cc:f7
- Internet Protocol, Src Addr: 66.185.95.101 (66.185.95.101), Dst Addr: 192.168.2.18 (192.168.2.18)
- Transmission Control Protocol, Src Port: pop3 (110), Dst Port: 2303 (2303), Seq: 1939399056, Ack: 112492731, Len: 88
- Post Office Protocol
 - Response: +OK
 - Response Arg: InterMail POP3 server ready.

```

0000  00 01 03 1d cc f7 00 04 e2 29 b2 3a 08 00 45 00  .....).:..E.
0010  00 4a 1c 63 40 00 36 06 c3 72 42 b9 5f 65 c0 a8  .J.c@.6. .rB_..
0020  02 12 00 6e 08 ff 73 98 e1 90 06 b4 80 bb 50 18  ...n..s. ....P.
0030  f0 00 97 c0 00 00 2b 4f 4b 20 49 6e 74 65 72 4d  ....+OK InterM
0040  61 69 6c 20 50 4f 50 33 20 73 65 72 76 65 72 20  ail POP3 server
    
```

Filter: / Reset Apply File: pop3 smtp

46. The nslookup program can be used to query the Internet domain name servers. Use this program to look up the IP address of www.utoronto.ca.

Solution:

```
nslookup www.utoronto.ca
```

```
Address: 128.100.132.30
```

The Ethereal capture below shows the sequence of DNS messages sent by the PC. The middle pane shows the contents of the final response, including the names of the authoritative servers.

The screenshot shows the 'nslookup ut - Ethereal' window with a packet capture table and a detailed view of a DNS response packet (Frame 32).

No.	Time	Source	Destination	Protocol	Info
28	6.884772	192.168.2.1	192.168.2.2	DNS	Standard query response, No such name
29	6.885377	192.168.2.2	192.168.2.1	DNS	Standard query A www.utoronto.ca.WorkGroup
30	7.102276	192.168.2.1	192.168.2.2	DNS	Standard query response, No such name
31	7.102530	192.168.2.2	192.168.2.1	DNS	Standard query A www.utoronto.ca
32	7.141562	192.168.2.1	192.168.2.2	DNS	Standard query response CNAME info.utcc.utoronto.ca

Frame 32 (293 bytes on wire, 293 bytes captured)

- Ethernet II, Src: 00:04:e2:29:b2:3a, Dst: 00:07:e9:c1:8e:0a
- Internet Protocol, Src Addr: 192.168.2.1 (192.168.2.1), Dst Addr: 192.168.2.2 (192.168.2.2)
- User Datagram Protocol, Src Port: domain (53), Dst Port: 1818 (1818)
- Domain Name System (response)
 - Transaction ID: 0x0003
 - Flags: 0x8180 (Standard query response, No error)
 - Questions: 1
 - Answer RRs: 2
 - Authority RRs: 4
 - Additional RRs: 6
 - Queries
 - www.utoronto.ca: type A, class inet
 - Name: www.utoronto.ca
 - Type: Host address
 - Class: inet
 - Answers
 - www.utoronto.ca: type CNAME, class inet, cname info.utcc.utoronto.ca
 - info.utcc.utoronto.ca: type A, class inet, addr 128.100.132.30
 - Authoritative nameservers
 - utcc.utoronto.ca: type NS, class inet, ns bay.cs.utoronto.ca
 - utcc.utoronto.ca: type NS, class inet, ns gpu.utcc.utoronto.ca
 - utcc.utoronto.ca: type NS, class inet, ns ns1.tpc.int
 - utcc.utoronto.ca: type NS, class inet, ns ns2.utoronto.ca
 - Additional records

Hex dump:

```

0000  00 07 e9 c1 8e 0a 00 04 e2 29 b2 3a 08 00 45 00  .....).:..E.
0010  01 17 f5 aa 00 00 40 11 fe d7 c0 a8 02 01 c0 a8  .....@.....
0020  02 02 00 35 07 1a 01 03 c3 d9 00 03 81 80 00 01  .....5.....
0030  00 02 00 04 00 06 03 77 77 77 08 75 74 6f 72 6f  .....w ww.utoro
0040  6e 74 6f 02 63 61 00 00 01 00 01 c0 0c 00 05 00  nto.ca.....

```

47. (a) Use PING to find the round-trip time to the home page of your university and to the home page of your department.

Solution:

```
ping www.utoronto.ca
```

(b) Use Ethereal to capture the ICMP packets exchanged. Correlate the information in the packet capture with the information displayed by the PING result.

Solution:

The sequence of ICMP packets exchanged is shown in the packet capture below.

No. .	Time	Source	Destination	Protocol	Info
1	0.000000	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
2	0.047257	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply
3	1.000054	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
4	1.073118	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply
5	2.001029	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
6	2.054577	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply
7	3.002215	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
8	3.053365	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply

.....

Frame 2 (74 bytes on wire, 74 bytes captured)

Ethernet II, Src: 00:04:e2:29:b2:3a, Dst: 00:07:e9:c1:8e:0a

Internet Protocol, Src Addr: 128.100.132.30 (128.100.132.30), Dst Addr: 192.168.2.2 (192.168.2.2)

Internet Control Message Protocol

Type: 0 (Echo (ping) reply)

Code: 0

Checksum: 0xe45b (correct)

Identifier: 0x0200

Sequence number: 6f:00

Data (32 bytes)

.....

```
0000 00 07 e9 c1 8e 0a 00 04 e2 29 b2 3a 08 00 45 00  .... .).:..E.
0010 00 3c d1 6f 40 00 f4 01 ee 23 80 64 84 1e c0 a8  .<.o@... .#.d...
0020 02 02 00 00 e4 5b 02 00 6f 00 61 62 63 64 65 66  ....[. o.abcdef
0030 67 68 69 6a 6b 6c 6d 6e 6f 70 71 72 73 74 75 76  ghijklmn opqrstuv
0040 77 61 62 63 64 65 66 67 68 69                    wabcdefg hi
```

Filter: / Reset Apply File: ping uoft

48. (a) Use netstat to find out the routing table for a host in your network.

Solution:

When you run the following command following a DOS prompt, such as in Windows 98 or Windows XP.

```
netstat -r
```

you will obtain the active routing table and the active TCP connections. The routing table has columns for IP address, network mask, gateway address, and network interface.

(b) Use netstat to find the IP statistics for your host.

Solution:

See Table 2.28 in textbook.

Other interesting netstat options are:

- a displays all active TCP connections & the TCP and UDP ports on which the computer is listening
- e displays Ethernet statistics
- p shows connections for the protocol specified by protocol, e.g. tcp, udp, icmp, ip, ...
- s displays stats by protocols

49. Suppose regularly spaced PING packets are sent to a remote host. What can you conclude from the following results?

Solutions follow questions:

- (a) No replies arrive back.

Possibilities are: the remote host is down; the remote host or the network is extremely congested; the remote host is set up not to reply.

- (b) Some replies are lost.

Some packets are discarded due to congestion at the remote-host listening-buffer or congestion at the network routers.

- (c) All replies arrive but with variable delays.

The packets traverse network routes that have different path length or traffic load.

- (d) What kind of statistics would be useful to calculate for the round-trip delays?

Time(reply packet arrival) - Time(the echo packet is sent).

50. Suppose you want to test the response time of a specific Web server. What attributes would such a measurement tool have? How would such a tool be designed?

Solution:

Retrieving document from a web server involves the establishment of a TCP connection, the sending of an HTTP request by the client, and the reply from the web server.

We define response time as the time elapsed from the time the client requests a document (GET command) to when the client receives the server's reply.

The measurement tool can make use of Telnet to access the web server. One would Telnet to port 80, after the TCP connection is set up, and then measure the time elapsed from sending the request to receiving a reply.

51. A denial-of-service attack involves loading a network resource to the point where it becomes non-functional.

- (a) Explain how PING can be used to carry out a denial-of-service attack.

Solution:

Using the PING program to send out a flood of packets to the network resource (for example, a server) increases the load on the server until it becomes nonfunctional.

- (b) On October 21, 2002 the 13 DNS root servers were subject to a distributed denial-of-service attack. Explain the impact of the attack on the operation of the Internet if some of the servers are brought down; if all of the servers are brought down.

Solution:

If all of the DNS root servers are brought down, then the root servers will not be available to handle the queries from local name servers resulting in severe disruption of the Internet. In practice, DNS servers make extensive use of caching and so can resolve a majority of queries without referring to the root servers.

52. (a) Use a web browser to retrieve a file from a local web server.

Solution:

We retrieved the main page from www.yahoo.com

- (b) HTTP relies on ASCII characters. To verify the sequence of messages shown in Table 2.1, use the Telnet program to retrieve the same file from the local web site.

Solution:

The Ethereal screen capture below shows the sequence of TCP segments exchanged by accessing www.yahoo.com using a telnet in MS Windows. Frames 4, 5, and 6 in the figure show the three-way handshake that transpires after the URL is typed into the telnet window. Telnet sends each character in a separate TCP segment. Frame 7 carries the first “g” in the GET HTTP command. The middle window shows that this TCP segment carries 1 byte of payload. The third window shows the payload in HEX 0x67 and in text “g”. The server TCP acknowledges each character individually. Frames 9, 11, 13, and 15 acknowledge g, e, t, (space), respectively. Frame 16 carries the \r\n that ends the HTTP request. Frame 17 begins the transfer of the yahoo web page to the client. The subsequent segments transfer information from the server with acknowledgements from the client.

No. .	Time	Source	Destination	Protocol	Info
4	13.001678	192.168.2.18	216.109.125.79	TCP	2498 > http [SYN] Seq=147142992 Ack=0 win=8192 L
5	13.039151	216.109.125.79	192.168.2.18	TCP	http > 2498 [SYN, ACK] Seq=2183346772 Ack=147142
6	13.039221	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142993 Ack=2183346773 W
7	18.472270	192.168.2.18	216.109.125.79	HTTP	Continuation
9	18.622879	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142994 W
10	18.734094	192.168.2.18	216.109.125.79	HTTP	Continuation
11	18.905241	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142995 W
12	18.924462	192.168.2.18	216.109.125.79	HTTP	Continuation
13	19.078556	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142996 W
14	19.244070	192.168.2.18	216.109.125.79	HTTP	Continuation
15	19.369356	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142997 W
16	19.692438	192.168.2.18	216.109.125.79	HTTP	Continuation
17	19.755278	216.109.125.79	192.168.2.18	HTTP	Continuation
18	19.756467	216.109.125.79	192.168.2.18	HTTP	Continuation
19	19.756515	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183349693 W
20	19.758513	216.109.125.79	192.168.2.18	HTTP	Continuation
21	19.843349	216.109.125.79	192.168.2.18	HTTP	Continuation
22	19.843467	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183352613 W
23	19.845612	216.109.125.79	192.168.2.18	HTTP	Continuation
24	19.846719	216.109.125.79	192.168.2.18	HTTP	Continuation
25	19.846764	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183355533 W
26	19.899221	216.109.125.79	192.168.2.18	HTTP	Continuation
27	19.900426	216.109.125.79	192.168.2.18	HTTP	Continuation
28	19.900508	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183358453 W
29	19.902280	216.109.125.79	192.168.2.18	HTTP	Continuation
30	19.911624	216.109.125.79	192.168.2.18	HTTP	Continuation
31	19.911745	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183361373 W
32	19.916041	216.109.125.79	192.168.2.18	HTTP	Continuation
33	19.917275	216.109.125.79	192.168.2.18	HTTP	Continuation

Frame 7 (55 bytes on wire, 55 bytes captured)

- Ethernet II, Src: 00:01:03:1d:cc:f7, Dst: 00:04:e2:29:b2:3a
- Internet Protocol, Src Addr: 192.168.2.18 (192.168.2.18), Dst Addr: 216.109.125.79 (216.109.125.79)
- Transmission Control Protocol, Src Port: 2498 (2498), Dst Port: http (80), Seq: 147142993, Ack: 2183346773
- Hypertext Transfer Protocol

Data (1 byte)

```

0000 00 04 e2 29 b2 3a 00 01 03 1d cc f7 08 00 45 00  ...).:.. ..E.
0010 00 29 80 77 40 00 80 06 61 e0 c0 a8 02 12 d8 6d  ..).w... a.....m
0020 7d 4f 09 c2 00 50 08 c5 39 51 82 23 3a 55 50 18  }o...P.. 9Q.#:UP.

```

The Tools menu in Ethereal include a Follow TCP Stream option to see the data carried in the TCP payloads. The text below shows the GET from the client and the first few characters in the file returned from the client.

```
get
```

```
<html><head>
```

```
<title>Yahoo!</title>
```

```
<meta http-equiv="PICS-Label" content="(PICS-1.1 "http://www.icra.org/ratingsv02.html"
```

53. Use Ethereal to capture the sequence of PDUs exchanged in problem 52 parts (a) and (b).

(a) Identify the Ethernet, IP, and TCP addresses of the machines involved in the exchange.

See middle pane for Ethernet addresses of machines in local LAN.

(b) Are there any DNS queries?

First few frames contain DNS.

(c) Identify the TCP connection setup.

Three-way handshake is shown.

(d) Examine the contents of the HTTP GET and response messages.

We show the first segment in the HTTP response.

(e) Examine how the TCP sequence numbers evolve over time.

The segment shown typically carry 1460 bytes from the server to the client.

The screenshot displays a network traffic capture in Ethereal. The main pane shows a list of packets with columns for No., Time, Source, Destination, Protocol, and Info. Packet 17 is highlighted, showing an HTTP continuation packet from 216.109.125.79 to 192.168.2.18. The packet details pane shows the structure of the packet, including Ethernet II, Internet Protocol, Transmission Control Protocol, and Hypertext Transfer Protocol. The packet bytes pane shows the raw data in hexadecimal and ASCII, including HTML tags like <h1>Yahoo</h1> and <meta charset="utf-8">.

No.	Time	Source	Destination	Protocol	Info
1	0.000000	Intel_c1:8e:0a	Broadcast	ARP	who has 192.168.2.3? Tell 192.168.2.2
2	12.992506	192.168.2.18	192.168.2.18	DNS	Standard query A www.yahoo.com
3	13.001008	192.168.2.1	192.168.2.18	DNS	Standard query response CNAME www.yahoo.ak
4	13.001678	192.168.2.18	216.109.125.79	TCP	2498 > http [SYN] Seq=147142992 Ack=0 win=
5	13.039151	216.109.125.79	192.168.2.18	TCP	http > 2498 [SYN, ACK] Seq=2183346772 Ack=
6	13.039221	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142993 Ack=218334
7	18.472270	192.168.2.18	216.109.125.79	HTTP	Continuation
8	18.600842	00000000.0001031dccf7	00000000.Broadcast	IPX SAP	General Query
9	18.622879	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=14714
10	18.734094	192.168.2.18	216.109.125.79	HTTP	Continuation
11	18.905241	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=14714
12	18.924462	192.168.2.18	216.109.125.79	HTTP	Continuation
13	19.078556	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=14714
14	19.244070	192.168.2.18	216.109.125.79	HTTP	Continuation
15	19.369356	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=14714
16	19.692438	192.168.2.18	216.109.125.79	HTTP	Continuation
17	19.755278	216.109.125.79	192.168.2.18	HTTP	Continuation
18	19.756467	216.109.125.79	192.168.2.18	HTTP	Continuation
19	19.756515	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=218334
20	19.758513	216.109.125.79	192.168.2.18	HTTP	Continuation
21	19.843349	216.109.125.79	192.168.2.18	HTTP	Continuation

Packet 17 (1514 bytes on wire, 1514 bytes captured)
 Ethernet II, Src: 00:04:e2:29:b2:3a, Dst: 00:01:03:1d:cc:f7
 Destination: 00:01:03:1d:cc:f7 (3com_1d:cc:f7)
 Source: 00:04:e2:29:b2:3a (Smc_Netw_29:b2:3a)
 Type: IP (0x0800)
 Internet Protocol, Src Addr: 216.109.125.79 (216.109.125.79), Dst Addr: 192.168.2.18 (192.168.2.18)
 Transmission Control Protocol, Src Port: http (80), Dst Port: 2498 (2498), Seq: 2183346773, Ack: 147142999,
 Hypertext Transfer Protocol
 Data (1460 bytes)

0030 ff ff 3e ef 00 00 3c 68 74 6d 6c 3e 3c 68 65 61 ..>...<h tm|><hea
 0040 64 3e 0a 0a 3c 74 69 74 6c 65 3e 59 61 68 6f 6f d>.<tit le>Yahoo
 0050 21 3c 2f 74 69 74 6c 65 3e 0a 3c 6d 65 74 61 20 !</title >.<meta
 0060 68 74 74 70 2d 65 71 75 69 76 3d 22 50 49 43 53 http-egu iv="PICS
 0070 2d 4c 61 62 65 6c 22 20 63 6f 6e 74 65 6e 74 3d -Label content=
 0080 27 28 50 49 43 53 2d 31 2e 31 20 22 68 74 74 70 (PICS-1 .1 "http
 0090 3a 2f 2f 77 77 77 2e 69 63 72 61 2e 6f 72 67 2f ://www.i cra.org/
 00a0 72 61 74 69 6e 67 73 76 30 32 2e 68 74 6d 6c 22 ratingsv 02.html"
 00b0 20 6c 20 72 20 28 63 7a 20 31 20 6c 7a 20 31 20 1 r (cz 1 lz 1
 00c0 6e 7a 20 31 20 6f 7a 20 31 20 76 7a 20 31 29 20 nz 1 oz 1 vz 1)
 00d0 67 65 6e 20 74 72 75 65 20 66 6f 72 20 22 68 74 gen true for "ht
 00e0 74 70 3a 2f 77 77 77 2e 79 61 68 6f 6f 2e 63 tp://www .yahoo.c
 00f0 6f 6d 22 20 72 20 28 63 7a 20 31 20 6c 7a 20 31 om" r (c z 1 lz 1
 0100 20 6e 7a 20 31 20 6f 7a 20 31 20 76 7a 20 31 29 nz 1 oz 1 vz 1)
 0110 20 22 68 74 74 70 3a 2f 2f 77 77 77 2e 72 73 61 "http://www.rsa
 0120 63 2e 6f 72 67 2f 72 61 74 69 6e 67 73 76 30 31 c.org/ra tingsv01
 0130 2a 68 74 6d 6c 22 20 6c 20 72 20 28 69 20 30 20 html" r (c z 1 lz 1)

54. Discuss the similarities and differences between the control connection in FTP and the remote control used to control a television. Can the FTP approach be used to provide VCR-type functionality to control the video from a video-on-demand service?

Solution:

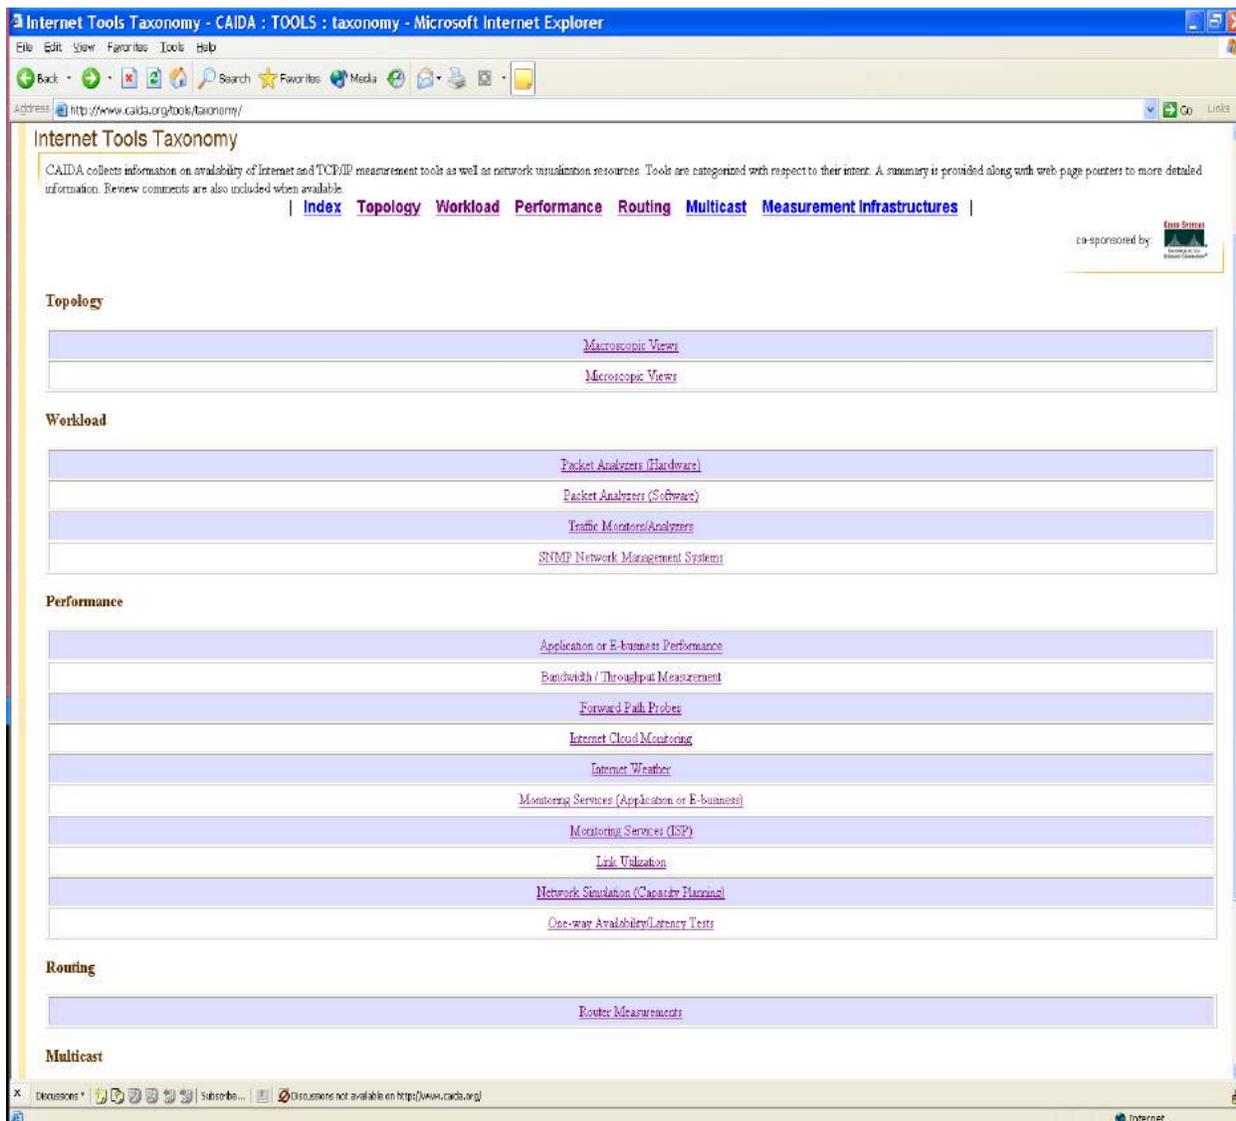
The FTP control connection and the TV's remote control are similar in that both are used to send commands to specify information about the data channel (TV display) and the data being requested (displayed). In both cases, the control channel is established upon the client/user's initiation. The client/user request for the closing of the control channel, and the server/TV is responsible for terminating the control and data channels.

Unlike the FTP control connection where a reply is generated for every command and is sent back on the control channel, for the TV remote control, either no reply is generated or the reply is sent on the data channel (TV display).

Yes, the FTP approach can be used to control video from video-on-demand service. The control channel can be used to provide VCR-type functionality (play, forward, reverse, stop) to control the video data sent on the data channel.

55. Use a web browser to access the Cooperative Association for Internet Data Analysis (CAIDA) web page (<http://www.caida.org/Tools/taxonomy.html>) to retrieve the CAIDA measurement tool taxonomy document. You will find links there to many free Internet measurement tools and utilities.

Solution:



56. Use traceroute to determine the path from your home PC to your university's main web page, while capturing the packets using Ethereal.

(a) Using the output from traceroute, try to identify how many different networks and service providers are traversed.

```
Tracing route to info.utcc.utoronto.ca [128.100.132.30]
over a maximum of 30 hops:
  0  1 ms    <1 ms   <1 ms   192.168.2.1
  1  7 ms    3 ms    10 ms   10.202.128.1
  2  21 ms   66 ms   30 ms   gw01-vlan961.ym.phub.net.cable.rogers.com [66.185.93.1]
  3  44 ms   25 ms   88 ms   gw02.ym.phub.net.cable.rogers.com [66.185.80.214]
  4  43 ms   23 ms   16 ms   gw02.bloor.phub.net.cable.rogers.com [66.185.80.138]
  5  14 ms   20 ms   30 ms   gw02.wlfdle.phub.net.cable.rogers.com [66.185.80.14]
  6  *       *       *       Request timed out.
  7  39 ms   54 ms   33 ms   p15-0.core01.yyz01.atlas.cogentco.com [154.54.1.173]
  8  43 ms   7 ms    11 ms   g0-1.na01.b011027-0.yyz01.atlas.cogentco.com [66.250.14.230]
  9  49 ms   *       32 ms   UniversityOfToronto.demarc.cogentco.com [38.112.2.34]
 10  32 ms   61 ms   63 ms   mcl-gateway.gw.utoronto.ca [128.100.96.101]
 11  27 ms   121 ms  26 ms   dbs.utcc.utoronto.ca [128.100.132.30]
Trace complete.
```

(b) Verify the operation of traceroute by examining the contents of the packets.

Solution:

The following Ethereal capture shows the last in the series of ICMP messages associated with the above traceroute. The IP address of www.utoronto.ca can be seen in the middle pane.

The screenshot shows the 'traceroute - Ethereal' window. The main pane displays a table of network hops:

No.	Time	Source	Destination	Protocol	Info
44	10.591376	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
45	15.091375	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
49	19.591408	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
50	19.630449	154.54.1.173	192.168.2.2	ICMP	Time-to-live exceeded
51	19.630588	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
52	19.685270	154.54.1.173	192.168.2.2	ICMP	Time-to-live exceeded
53	19.685410	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
54	19.718935	154.54.1.173	192.168.2.2	ICMP	Time-to-live exceeded
55	20.687053	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
56	20.730768	66.250.14.230	192.168.2.2	ICMP	Time-to-live exceeded
57	20.730925	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
58	20.738835	66.250.14.230	192.168.2.2	ICMP	Time-to-live exceeded
59	20.738972	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
60	20.750235	66.250.14.230	192.168.2.2	ICMP	Time-to-live exceeded
61	21.739794	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
62	21.789228	38.112.2.34	192.168.2.2	ICMP	Time-to-live exceeded
63	21.789363	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
64	26.091363	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
65	26.124106	38.112.2.34	192.168.2.2	ICMP	Time-to-live exceeded
66	27.093297	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
67	27.125539	128.100.96.101	192.168.2.2	ICMP	Time-to-live exceeded
68	27.125646	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
69	27.186954	128.100.96.101	192.168.2.2	ICMP	Time-to-live exceeded
70	27.187059	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
71	27.250232	128.100.96.101	192.168.2.2	ICMP	Time-to-live exceeded
72	28.188023	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
73	28.215480	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply
74	28.215585	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
75	28.337215	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply
76	28.337358	192.168.2.2	128.100.132.30	ICMP	Echo (ping) request
77	28.363731	128.100.132.30	192.168.2.2	ICMP	Echo (ping) reply

The bottom pane shows details for Frame 77 (106 bytes on wire, 106 bytes captured):

- Ethernet II, Src: 00:04:e2:29:b2:3a, Dst: 00:07:e9:c1:8e:0a
- Internet Protocol, Src Addr: 128.100.132.30 (128.100.132.30), Dst Addr: 192.168.2.2 (192.168.2.2)
- Internet Control Message Protocol
 - Type: 0 (Echo (ping) reply)
 - Code: 0
 - Checksum: 0x8fff (correct)
 - Identifier: 0x0200
 - Sequence number: 6e:00
 - Data (64 bytes)

The hex dump at the bottom shows the raw data of the frame:

```

0000 00 07 e9 c1 8e 0a 00 04 e2 29 b2 3a 08 00 45 00  ....E.
0010 00 5c 25 75 40 00 f4 01 99 fe 80 64 84 1e c0 a8  .\%u@...d...
0020 02 02 00 00 8f ff 02 00 6e 00 00 00 00 00 00 00  .....n.....
0030 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
0040 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
  
```

57. Run the UDP client and server programs from the Berkeley API section on different machines, record the round-trip delays with respect to the size of the data, and plot the results.

Solution:

To be added.

58. In the TCP example from the Berkeley API section, the message size communicated is fixed regardless of how many characters of actual information a user types. Even if the user wants to send only one character, the programs still sends 256 bytes of messages - clearly an inefficient method. One possible way to allow variable-length messages to be communicated is to indicate the end of a message by a unique character, called the sentinel. The receiver calls `read` for every character (or byte), compares each character with the sentinel value, and terminates after this special value is encountered. Modify the TCP client and server programs to handle variable-length messages using a sentinel value.

Solution:

All strings in C end with the null character '\0'. Use the null character as the sentinel value.

Note: modifications to the program are in **bold**.

Client program:

```
-----
#include <string.h>

#define MAX_BUFLEN 256 /* buffer length */
#define SENTINEL '\0'

int main (int argc, char **argv)
{
    .
    .
    .
    char *host, *bp, rbuf[MAX_BUFLEN], sbuf[MAX_BUFLEN];
    int BUFLen; /* actual message length <= MAX_BUFLEN */
    .
    .
    .
    printf("Transmit:\n");
    gets(sbuf); /* reads char into sbuf from stdin */
    BUFLen = strlen(sbuf); /* get length of the message */
    if (BUFLen > MAX_BUFLEN) {
        fprintf(stderr, "Data is too big\n");
        exit(1);
    }
    write(sd, sbuf, BUFLen+1); /* send to server - including the sentinel
character '\0'*/

    printf("Receive:\n");
    bp = rbuf;
    bytes_to_read = 1; /* read one character at a time */
    while ( ((n = read(sd, bp, bytes_to_read)) > 0)
        && (*bp != SENTINEL) ) {
        bp += n;
    }
    .
    .
    .
} /* main */
```

Server program:

```
-----
int main (int argc, char **argv) {
    .
    .
    .
    while (1) {
        client_len = sizeof(client);
        if ((new_sd = accept(sd, (struct sockaddr*)&client, &client_len ))
            == -1) {
            fprintf(stderr, "Can't accept client\n");
            exit(1);
        }
    }
}
```

```

    bp = buf; /* bp is the traversing pointer of buf */

    bytes_to_read = 1;
    while ( ((n= read(new_sd, bp, bytes_to_read)) > 0) &&
            (*bp != SENTINEL)) {
        bp += n; /* move the pointer bp within buf */
    }

    BUFLen = strlen(buf);
    write(new_sd, buf, BUFLen+1);
    close(new_sd);
} /* while */

.
.
.
} /* main */

```

59. Another possible way to allow variable-length messages to be communicated is to precede the data to be transmitted by a header indicating the length of the data. After the header is decoded, the receiver knows how many more bytes it should read. Assuming the length of the header is two bytes, modify the TCP client and server programs to handle variable-length messages.

Solution:

Note: modifications to the program are in **bold**.

Client Program:

```

-----
#include <string.h>

#define MAX_BUFLen 256 /* buffer length */
#define HEADER_TYPE short
#define HEADER_SIZE sizeof(HEADER_TYPE)
.
.
.
int main (int argc, char **argv) {

    char *host, *bp, rbuf[MAX_BUFLen], sbuf[MAX_BUFLen];
    short BUFLen; /* actual message length <= MAX_BUFLen */
    HEADER_TYPE Header;
    .
    .
    .
    printf("Transmit:\n");
    gets(sbuf); /* reads char into sbuf from stdin */
    BUFLen = strlen(sbuf) + 1; /* get length of the message including the NULL
char */

    if (BUFLen-1 > MAX_BUFLen) {
        fprintf(stderr, "Data is too big\n");
        exit(1);
    }

    Header = (HEADER_TYPE)BUFLen;

```

```

    write(sd, (char *)&Header, HEADER_SIZE);
    write(sd, sbuf, BUFLen); /* send to server - including the sentinel
character '\0'*/

    printf("Receive:\n");
    /* Read Header */
    bytes_to_read = HEADER_SIZE;
    bp = (char *)&Header;
    while ( (n=read(sd, bp, bytes_to_read)) > 0) {
        bp += n;
        bytes_to_read -= n;
    }

    /* Read Message */
    bp = rbuf;
    bytes_to_read = (int)Header; /* read one character at a time */
    while ( (n = read(sd, bp, bytes_to_read)) > 0) {
        bp += n;
        bytes_to_read -= n;
    }
    .
    .
    .
} /* main */

```

Server program:

```

-----
int main (int argc, char **argv) {
    .
    .
    .
    while (1) {
        .
        .
        .
        /* Read Header */
        bytes_to_read = HEADER_SIZE;
        bp = (char *)&Header;
        while ( (n=read(new_sd, bp, bytes_to_read)) > 0) {
            bp += n;
            bytes_to_read -= n;
        }

        /* Read message */
        bytes_to_read = (int)Header;
        bp = buf; /* bp is the traversing pointer of buf */
        while ( (n= read(new_sd, bp, bytes_to_read)) > 0){
            bp += n; /* move the pointer bp within buf */
            bytes_to_read -= n;
        }

        BUFLen = strlen(buf)+1;
        Header = (HEADER_TYPE)BUFLen;

        write(new_sd, ((char *)&Header), HEADER_SIZE);
    }
}

```

```

    write(new_sd, buf, BUFLLEN);
    close(new_sd);
} /* while */
.
.
.
} /* main */

```

60. The UDP client program in the example from the Berkeley API section may wait forever if the datagram from the server never arrives. Modify the client program so that if the response from the server does not arrive after a certain timeout (say, 5 seconds), the `read` call is interrupted. The client then retransmits a datagram to the server and waits for a new response. If the client does not receive a response after a fixed number of trials (say, 10 trials), the client should print an error message and abandon the program. Hint: use the `sigaction` and `alarm` functions.

Solution:

Note: modifications to the program are in **bold**.

Client program:

```

-----
#include <signal.h>
#include <unistd.h>

#define MAX_NUM_TRIAL 10
#define TIMEOUT 5

void retransmit(); /* alarm handler function */
int NUM_TRIAL = MAX_NUM_TRIAL;
int ServerReply = 0; /* flag indicating whether server has
replied, used to determined whether to retry */

int main (int argc, char **argv) {
    .
    .
    .
    struct sigaction action;
    /* ----- */
    /* set up signal action */
    /* ----- */
    action.sa_handler = retransmit;
    sigemptyset(&action.sa_mask);
    action.sa_flags = 0;
    sigaction(SIGALRM, &action, NULL);
    .
    .
    .
    gettimeofday(&start, NULL); /* start delay measure */

    ServerReply = 0;
    while ( (ServerReply == 0) && (NUM_TRIAL > 0) ) {

        /* Transmit data */
        server_len = sizeof(server);
        if (sendto(sd, sbuf, data_size, 0, (struct sockaddr *)&server,
            server_len) == -1) {

```

```

    fprintf(stderr, "Sendto error\n");
    exit(1);
}

/* Receive data */
ServerReply = 1;
alarm(TIMEOUT);
if (recvfrom(sd, rbuf, MAXLEN, 0, (struct sockaddr *)&server,
            &server_len) < 0) {
    /* Note: Sigaction handler function returns here, to
       continue retry, do not exit on recvfrom error */
}
alarm(0);
} /* while */

if (ServerReply) {
    gettimeofday(&end, NULL); /* end delay measure */
    printf("\nRound-trip delay = %ld ms.\n", delay(start, end));
    if (strncmp(sbuf, rbuf, data_size) != 0)
        printf("Data is corrupted\n");
}
else
{
    printf("client: no reply from server at %s: Abort.\n", host,
          MAX_NUM_TRIAL);
}

close(sd);
return(0);

} /* main */

/* sigaction handler function */
void retransmit() {
    NUM_TRIAL--;
    printf("client: receive time out - retry #%d\n", MAX_NUM_TRIAL-NUM_TRIAL);
    ServerReply = 0;
}

```

61. Modify the UDP client to access a date-and-time server in a host local to your network. A date-and-time server provides client programs with the current day and time on demand. The system internal clock keeps the current day and time as a 32-bit integer. The time is incremented by the system (every second). When an application program (the server in this case) asks for the date or time, the system consults the internal clock and formats the date and time of day in human-readable format. Sending any datagram to a date-and-time server is equivalent to making a request for the current date and time; the server responds by returning a UDP message containing the current date and time. The date-and-time server can be accessed in UDP port 13.

Solution:

Make the following modification to the original UDP client program (changes noted in **bold**).

```

-----
int main (int argc, char **argv) {
    .
    .
    .
    /* Receive data */
    if (recvfrom(sd, rbuf, MAXLEN, 0, (struct sockaddr *)&server,

```

```
        &server_len) < 0) {
    fprintf(stderr, "recvfrom error\n");
    exit(1);
}
printf("%s system time is %s", host, rbuf);
    .
    .
    .
} /* main */
```

Run the program and specify the server port to be 13.

62. Write a file transfer application that runs over UDP. Assume the transfer occurs over a local area network so reliable transfer is not a concern. Assume also that UDP will accept at most 500 bytes/datagram. Implement a server that opens a socket and listens for incoming data at a particular port number. Implement a client that reads a file from the file system and transfers the file to the server. When the server receives the client's data, it writes this data to a file. Include a means for the client to indicate the end of transmission.

Solution:

TO BE ADDED.

Solutions to Chapter 3

1. Suppose the size of an uncompressed text file is 1 megabyte.

Solutions follow questions:

- a. How long does it take to download the file over a 32 kilobit/second modem?

$$T_{32k} = 8 (1024) (1024) / 32000 = 262.144 \text{ seconds}$$

- b. How long does it take to take to download the file over a 1 megabit/second modem?

$$T_{1M} = 8 (1024) (1024) \text{ bits} / 1 \times 10^6 \text{ bits/sec} = 8.38 \text{ seconds}$$

- c. Suppose data compression is applied to the text file. How much do the transmission times in parts (a) and (b) change?

If we assume a maximum compression ratio of 1:6, then we have the following times for the 32 kilobit and 1 megabit lines respectively:

$$T_{32k} = 8 (1024) (1024) / (32000 \times 6) = 43.69 \text{ sec}$$

$$T_{1M} = 8 (1024) (1024) / (1 \times 10^6 \times 6) = 1.4 \text{ sec}$$

2. A scanner has a resolution of 600 x 600 pixels/square inch. How many bits are produced by an 8-inch x 10-inch image if scanning uses 8 bits/pixel? 24 bits/pixel?

Solution:

The number of pixels is $600 \times 600 \times 8 \times 10 = 28.8 \times 10^6$ pixels per picture.

With 8 bits/pixel representation, we have: $28.8 \times 10^6 \times 8 = 230.4$ Mbits per picture.

With 24 bits/pixel representation, we have: $28.8 \times 10^6 \times 24 = 691.2$ Mbits per picture.

3. Suppose a computer monitor has a screen resolution of 1200 x 800 pixels. How many bits are required if each pixel uses 256 colors? 65,536 colors?

Solution:

The number of pixels is $1200 \times 800 = 960000$ pixels per screen.

With 256 colors, we have: $9.6 \times 10^5 \times \log_2 256 = 9.6 \times 10^5 \times 8 = 7.68 \times 10^6$ bits per screen.

With 65536 colors, we have: $9.6 \times 10^5 \times \log_2 65536 = 9.6 \times 10^5 \times 16 = 15.36 \times 10^6$ bits per screen.

4. Explain the difference between facsimile, GIF, and JPEG coding. Give an example of an image that is appropriate to each of these three methods.

Solution:

Facsimile scans a black-and-white image into rows of black and white dots, which are then compressed. The scanning process introduces distortion in the sense that the scanned image is not the same as the original image.

GIF takes image data, in binary form, and applies a lossless data compression scheme.

JPEG involves scanning an image and applying a lossy compression scheme that is designed so that the distortion is not visible.

5. A digital transmission system has a bit rate of 45 megabits/second. How many PCM voice calls can be carried by the system?

Solution:

$$\text{PCM channels} = (45 \times 10^6 \text{ bits/sec}) / (64 \times 10^3 \text{ bits/sec channel}) = 703 \text{ channels.}$$

6. Suppose a storage device has a capacity of 1 gigabyte. How many 1-minute songs can the device hold using conventional CD format? using MP3 coding?

Solution:

A stereo CD signal has a bit rate of 1.4 megabits per second, or 84 megabits per minute, which is approximately 10 megabytes per minute. Therefore a 1 gigabyte storage will hold 1 gigabyte/10 megabyte = 100 songs.

An MP3 signal has a lower bit rate than a CD signal by about a factor of 14, so 1 gigabyte storage will hold about 1400 songs.

7. How many high-quality audio channels could be transmitted using an HDTV channel?

Solution:

An audio channel is about 1.4 Mbps and an HDTV channel is about 38 Mbps, so an HDTV channel can carry about 27 such audio channels. The number of audio channels increases by a factor of 10 if more compression (e.g. MP3) is applied to the audio signals.

8. How many HDTV channels can be transmitted simultaneously over the optical fiber transmission systems in Table 3.3?

Solution:

Suppose that an optical fiber carries 1600×10^9 bps, and an HDTV channel is about 38 Mbps, then the fiber can carry about $1600000/38 = 40,000$ HDTV channels.

9. Comment on the properties of the sequence of frame images and the associated bit rates in the following examples.

Solutions follow questions:

a. A children's cartoon program.

The frames consist of relatively simple objects and static backgrounds with little fast motion, so the differences between consecutive frames are small. Therefore these sequences are highly compressible.

b. A music video.

The frames consist of relatively complex images that may consist of the overlay/combination of multiple images, and there can be fast motion and sudden changes in scenes, so the differences between consecutive frames can be high. Therefore these sequences are not highly compressible.

c. A tennis game; a basketball game.

Both tennis and basketball have fast motion. In tennis there are fewer moving objects than in basketball, and in basketball the background can be more complex. Both sequences are not easy to compress.

- d. A documentary on famous paintings.

The sequences are relatively slowly changing, with much panning and zooming. The degree of detail/complexity of the images can be high. The bit rate depends on the detail in the images.

10. Suppose that at a given time of the day, in a city with a population of 1 million, 1% of the people are on the phone.

- a. What is the total bit rate generated by all these people if each voice call is encoded using PCM?

The number of people on the phone at a given time is $10^6 \times 10^{-2} = 10^4$, so the total bit rate is $10^4 \times 64 \times 10^3 = 640$ megabits per second.

- b. What is the total bit rate if all of the telephones are replaced by H.261 videoconferencing terminals?

If the videoconferencing bit rate is 64 kbps, then the total bit rate generated by the city is unchanged. However if high-quality videoconferencing is used and the bit rate is increased to 1.5 Mbps, then the total bit rate increases to 15 Gbps.

11. Consider an analog repeater system in which the signal has power σ_x^2 and each stage adds noise with power σ_n^2 . For simplicity assume that each repeater recovers the original signal without distortion but that the noise accumulates. Find the SNR after n repeater links. Write the expression in decibels: SNR dB = $10 \log_{10}$ SNR.

Solution:

After n stages, the signal power is σ_x^2 and the noise power is $n\sigma_n^2$, so the SNR is:

$$\text{SNR dB} = 10 \log_{10} \sigma_x^2 / n\sigma_n^2 = 10 \log_{10} \sigma_x^2 / \sigma_n^2 + 10 \log_{10} 1/n = 10 \log_{10} \sigma_x^2 / \sigma_n^2 - 10 \log_{10} n$$

12. Suppose that a link between two telephone offices has 50 repeaters. Suppose that the probability that a repeater fails during a year is 0.01, and that repeaters fail independently of each other.

- a. What is the probability that the link does not fail at all during one year?

Let p be the probability that a repeater fails during a year, then $1 - p$ is the probability that it does not fail, and the probability that all 50 repeaters do not fail is $(1 - .01)^{50} \approx e^{-50(.01)} = 0.605$ where we have used the approximation $(1 - p)^n \approx e^{-np}$ which is valid for large n and small p .

- b. Repeat (a) with 10 repeaters; with 1 repeater.

The probability that all 10 repeaters do not fail is $(1 - .01)^{10} \approx e^{-10(.01)} = 0.905$, and the probability that a single repeater does not fail is 0.99.

The moral of the calculations is that a system that requires the functioning of a large number of relatively reliable components may be fairly unreliable. In terms of repeaters, this implies that minimizing the number of repeaters needed in a link is important from the point of view of reliability. Of course this also reduces the cost expended to install and maintain the repeaters.

13. Suppose that a signal has twice the power as a noise signal that is added to it. Find the SNR in decibels. Repeat if the signal has 10 times the noise power? 2^n times the noise power? 10^k times the noise power?

Solution:

$$\text{SNR dB} = 10 \log_{10} \sigma_x^2 / \sigma_n^2 = 10 \log_{10} 2 = 3.01 \text{ dB}$$

$$\text{SNR dB} = 10 \log_{10} 10 = 10 \text{ dB}$$

$$\text{SNR dB} = 10 \log_{10} 2^n = 10n \log_{10} 2 = 3.01n \text{ dB}$$

$$\text{SNR dB} = 10 \log_{10} 10^k = 10k \text{ dB}$$

14. A way of visualizing the Nyquist theorem is in terms of periodic sampling of the second hand of a clock which makes one revolution around the clock every 60 seconds. The Nyquist sampling rate here should correspond to 2 samples per cycle, that is, sampling should be done at least every 30 seconds.

- a. Suppose we begin sampling when the second hand is at 12 o'clock and that we sample the clock every 15 seconds. Draw the sequence of observations that result. Does the second hand appear to move forward?

0 15 30 45 0 15 30 45 0 ... Yes, the second hand appears to move forward.

- b. Now suppose we sample every 30 seconds. Does the second hand appear to move forward or backward? What if we sample every 29 seconds?

0 30 0 30 0 30 0 30 0 30 ... No. It's hard to tell if it moves forward or backward because either direction will give the same sequence.

0 29 58 27 56 25 54 23 ... Yes, the second hand appears to move forward.

- c. Explain why a sinusoid should be sampled at a little more than twice its frequency.

When a sinusoid is being sampled at exactly twice its frequency, it is possible that all samples have value of zero. Therefore, the sampling rate should be a little more than twice its frequency.

- d. Now suppose that we sample every 45 seconds. What is the sequence of observations of the second clock hand?

One sampling every 45 seconds -> 0 45 30 15 0 45 30 15 0 The second hand appears to be moving backward.

- e. Motion pictures are made by taking a photograph 24 times a second. Use part (c) to explain why car wheels in movies often appear to spin backward while the cars are moving forward!

The spinning wheel is being sampled at 24 times per seconds. Depending on the rotation speed of the wheel, the wheel may appear either as spinning backward while actually spinning forward.

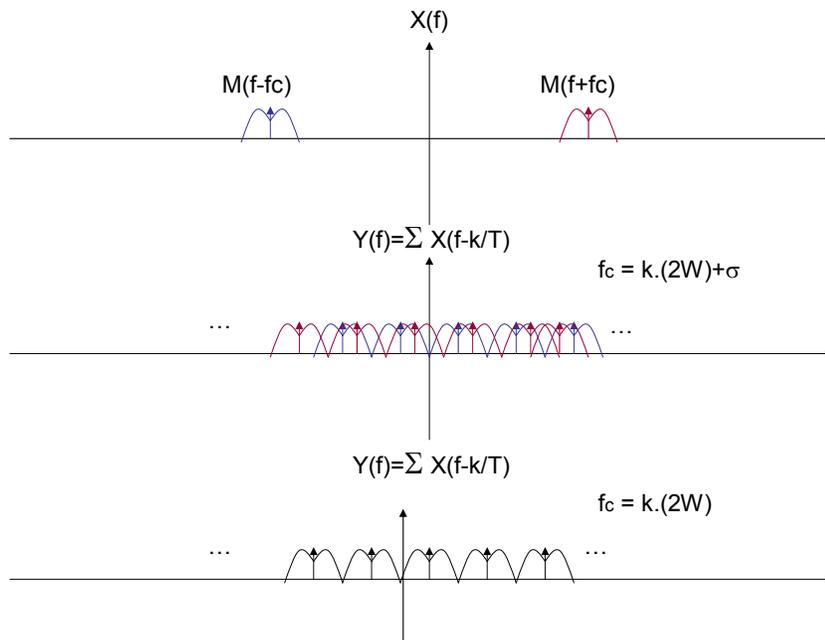
15. "Software radios" are devices that can demodulate and decode any radio signal regardless of format or standard. The basic idea in software radio is to immediately convert the transmitted radio signal into digital form so that digital signal processing software can be used to do the particular required processing. Suppose that a software radio is to demodulate FM radio and television. What sampling rate is required in the A/D conversion? The transmission bandwidth of FM radio is 200 kHz and the transmission bandwidth of television is 6 MHz.

Solution: The sampling rate must be at least $2 \times 6 = 12$ MHz.

16. An AM radio signal has the form $x(t) = m(t)\cos(2\pi f_c t)$, where $m(t)$ is a low-pass signal with bandwidth W Hz. Suppose that $x(t)$ is sampled at a rate of $2W$ samples per second. Sketch the spectrum of the sampled sequence. Under which conditions can $m(t)$ be recovered from the sampled sequence? (Hint: See Appendix 3C.)

Solution:

As shown in the figure, the spectrum $X(f)$ consists of two components at f_c and $-f_c$. After sampling $x(t)$ at a rate of $2W$, the spectrum $Y(f)$ consists of repeated versions of $X(f)$, which includes repeated versions of both of the two components. The repeated versions of each one of the components of $X(f)$ will be apart from each other because the sampling rate is $2W$ however the repeated versions of the two components will interfere, resulting in a distorted spectrum if f_c is not a multiple of $2W$. Therefore the condition for the recovery of $m(t)$ is $f_c = k(2W)$, where k is an integer number.



17. A black-and-white image consists of a variation in intensity over the plane.

Solutions follow questions:

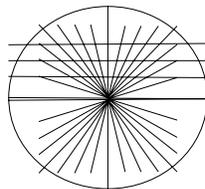
- a. By using an analogy to time signals, explain spatial frequency in the horizontal direction; in the vertical direction. Hint: Consider bands of alternating black and white bands. Do you think there is a Nyquist sampling theorem for images?

Consider a picture that consists of alternating black and white vertical bands with equal width. We take a horizontal strip of the picture and draw a signal that represents the color of the strip as we traverse the vertical bands. This will generate a square wave whose frequency is the spatial frequency of the picture in horizontal direction. If the width of the bands is smaller the spatial frequency of the picture will be higher. Similarly spatial frequency in the vertical direction can be described considering similar bands in the horizontal direction and a vertical strip.

Sampling this picture in either horizontal or vertical direction is equivalent to the sampling of the above mentioned signals. Nyquist sampling theorem applies to these signals even though the signals are in space domain rather than time domain. Considering the above example again, sampling each strip of the picture will preserve all the information if at least one sample is taken from each band. Since each period of the signal involves two samples, the sampling rate will be twice the frequency of the signal.

- b. Now consider a circle and select a large even number N of equally spaced points around the perimeter of the circle. Draw a line from each point to the center and color alternating regions black and white. What are the spatial frequencies in the vicinity of the center of the circle?

The spatial frequency will be higher as we get closer to the center of the circle, and aliasing may become manifest. In practice an image will not have absolute precision, that is, will be bandlimited and aliasing is avoided by sampling at a sufficiently high frequency. An image can be made bandlimited by blurring it a little bit. Thus photographers often intentionally blur the image of racing cars to avoid the effects of aliasing.



18. A high-quality speech signal has a bandwidth of 8 kHz.

- a. Suppose that the speech signal is to be quantized and then transmitted over a 28.8 kbps modem. What is the SNR of the received speech signal?

$$W = 8 \text{ kHz and } R = 28.8 \times 10^3 \text{ bits/sec}$$

$$2W = 16 \text{ K samples/sec}$$

$$m = R/2W = 28.8/16 \approx 2 \text{ bits/sample}$$

$$\text{SNR} \approx 6m - 7.2\text{dB} \approx 5 \text{ dB (very approximate)}$$

- b. Suppose that instead a 64 kbps modem is used? What is the SNR of the received speech signal?

$$R/2W = 64/16 = 4 \text{ bits/sample}$$

$$\text{SNR} \approx 24 - 7.2\text{db} \approx 17 \text{ dB (very approximate)}$$

- c. What modem speed is needed if we require an SNR of 40 dB?

$$\text{SNR} = 40 \text{ dB} = 6m - 7.2 \text{ dB}$$

$$6m = 47.2$$

$$m = 8$$

$$R = 8 \text{ bits/sample} \times (16 \times 10^3 \text{ samples/sec}) = 128 \text{ kbps}$$

19. An analog television signal is a lowpass signal with a bandwidth of 4 MHz. What is the bit rate required if we quantize the signal and require an SNR of 60 dB?

Solution:

$$W = 4 \times 10^6 \text{ Hz}$$

$$2W = 8 \times 10^6 \text{ samples/sec}$$

$$\text{SNR} = 60 \text{ dB} = 6m - 7.2 \text{ dB}$$

$$6m = 67$$

$$m \approx 11 \text{ bits}$$

$$R = 2W \times 11 = 88 \text{ Mbps}$$

20. An audio digitizing utility in a PC samples an input signal at a rate of 44 kHz and 16 bits/sample. How big a file is required to record 20 seconds?

Solution:

$$R = 44 \times 10^3 \times 16 = 704 \times 10^3 \text{ bps}$$

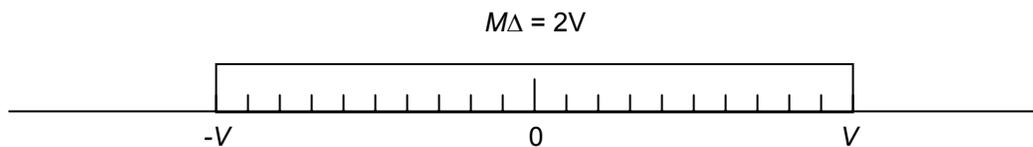
$$\text{Number of bits generated in 20 seconds} = 20 \times 704 \times 10^3 \text{ bits} = 14080 \times 10^3$$

$$\text{The file size} = (14080 \times 10^3) / (8 \times 1024) \approx 1719 \text{ K bytes} \approx 1.7 \text{ M bytes}$$

21. Suppose that a signal has amplitudes uniformly distributed between $-V$ and V .

Solutions follow questions:

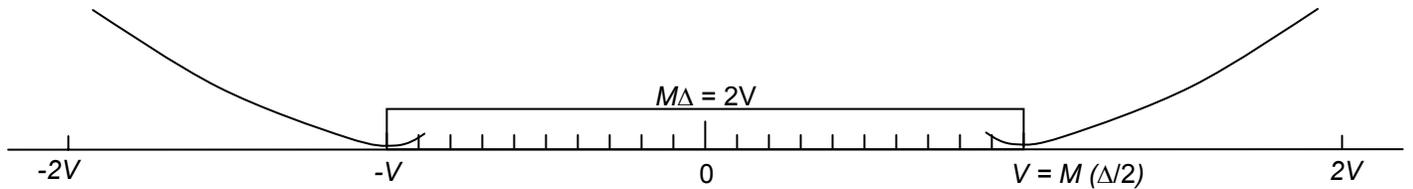
a. What is the SNR for a uniform quantizer that is designed specifically for this source?



$$\sigma_x^2 = \frac{1}{2V} \int_{-V}^V x^2 dx = \frac{1}{V} \int_0^V x^2 dx = \frac{V^3}{3V} = \frac{V^2}{3}$$

$$\text{SNR} = \frac{V^2/3}{\Delta^2/12} = \frac{4V^2}{(2V/M)^2} = M^2 = 2^{2m}$$

b. Suppose that the quantizer design underestimates the dynamic range by a factor of 2; that is, the actual dynamic range is $-2V$ to $2V$. Plot the quantization error vs. signal amplitude for this case. What is the SNR of the quantizer?



$$\sigma_x^2 = \frac{1}{4V} \int_{-2V}^{2V} x^2 dx = \frac{1}{2V} \int_0^{2V} x^2 dx = \frac{4V^2}{3}$$

$$\text{Additional noise power} = \frac{1}{2V} \int_V^{2V} \underbrace{\left(x - \left(V - \frac{\Delta}{2} \right) \right)}_y^2 dx$$

$$= \frac{1}{2V} \int_{\frac{\Delta}{2}}^{V + \frac{\Delta}{2}} y^2 dy = \frac{1}{6V} \left[\left(V + \frac{\Delta}{2} \right)^3 - \left(\frac{\Delta}{2} \right)^3 \right] = \left(V + \frac{\Delta}{2} \right)^2 + \frac{3\Delta^2}{16}$$

$$\text{SNR}' = \frac{\frac{4V^2}{3}}{\frac{\Delta^2}{12} + \left(V + \frac{\Delta}{2} \right)^2 + \frac{3\Delta^2}{16}}$$

22. A telephone office line card is designed to handle modem signals of the form $x(t) = A \cos(2\pi f_c t + \phi(t))$. These signals are to be digitized to yield an SNR of 40 dB using a uniform quantizer. Due to variations in the length of lines and other factors, the value of A varies by up to a factor of 100.

Solutions follow questions:

a. How many levels must the quantizer have to produce the desired SNR?

$$\text{SNR} = 40 = 10 \log_{10} \left(\frac{(A/100)^2 / 2}{\Delta^2 / 12} \right) = 10 \log_{10} \left(\frac{(A/100)^2 / 2}{(2A/2^m)^2 / 12} \right) = 10 \log_{10} \left(\frac{1.5x2^{2m}}{10000} \right) = 1.76 - 40 + 6m$$

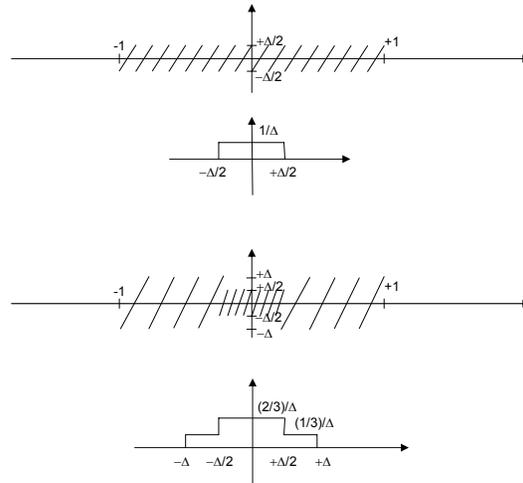
$$m = 13$$

b. Explain how an adaptive quantizer might be used to address this problem.

An adaptive quantizer may use larger intervals when the signal level is high and smaller intervals when the signal level is low. This way the number of bits will be less.

23. The basic idea in companding is to obtain robustness with respect to variations in signal level by using small quantizer intervals for small signal values and larger intervals for larger signals values. Consider an eight-level quantizer in which the inner four intervals are Δ wide and the outer four intervals are 2Δ wide. Suppose the quantizer covers the range -1 to 1 . Find the SNR if the input signal is uniformly distributed between $-V$ and V for $\frac{1}{2} < V < 1$. Compare to the SNR of a uniform quantizer.

Solution:



In the case of nonuniform quantizer: $\Delta = 1/12$

$$\sigma_x^2 = \frac{1}{2V} \int_{-V}^V x^2 dx = \frac{1}{V} \int_0^V x^2 dx = \frac{V^3}{3V} = \frac{V^2}{3}$$

$$\sigma_e^2 = \frac{1}{3} \left(\frac{1}{\Delta} \int_{-\Delta/2}^{\Delta/2} x^2 dx \right) + \frac{2}{3} \left(\frac{1}{2\Delta} \int_{-\Delta}^{\Delta} x^2 dx \right) = \frac{1}{18} \Delta^2 + \frac{7}{36} \Delta^2 = \frac{1}{4} \Delta^2$$

$$SNR = \frac{\sigma_x^2}{\sigma_e^2} = \frac{V^2/3}{\Delta^2/4} = \frac{4}{3} \frac{V^2}{\Delta^2}$$

$$SNR = \frac{4}{3} \frac{V^2}{(1/12)^2} = 192V^2$$

In the case of uniform quantizer considering the same 8 level and three bits: $\Delta=1/8$

$$SNR = \frac{\sigma_x^2}{\sigma_e^2} = \frac{V^2/3}{\Delta^2/12} = 4 \frac{V^2}{\Delta^2}$$

$$SNR = 4 \frac{V^2}{(1/8)^2} = 256V^2$$

24. Suppose that a speech signal is A/D and D/A converted four times in traversing a telephone network that contains analog and digital switches. What is the SNR of the speech signal after the fourth D/A conversion?

Solution:

At each A/D and D/A conversion the noise introduced in the quantization process is added to the noise level in previous conversions. Assuming that the noise level in all conversions is the same, in the final stage the average noise power will be four times larger:

$$\text{SNR} = 10 \log_{10}(\sigma_x^2 / 4\sigma_e^2) = 10 \log_{10}(\sigma_x^2 / \sigma_e^2) - 10 \log_{10}4$$

Considering SNR = 38 dB for each A/D and D/A conversion in the telephone system, the SNR after the fourth conversion will be 38 – 6 = 32 dB.

25. A square periodic signal is represented as the following sum of sinusoids:

$$g(t) = \frac{2}{\pi} \sum_{k=0}^{\infty} \frac{(-1)^k}{2k+1} \cos(2k+1)\pi t$$

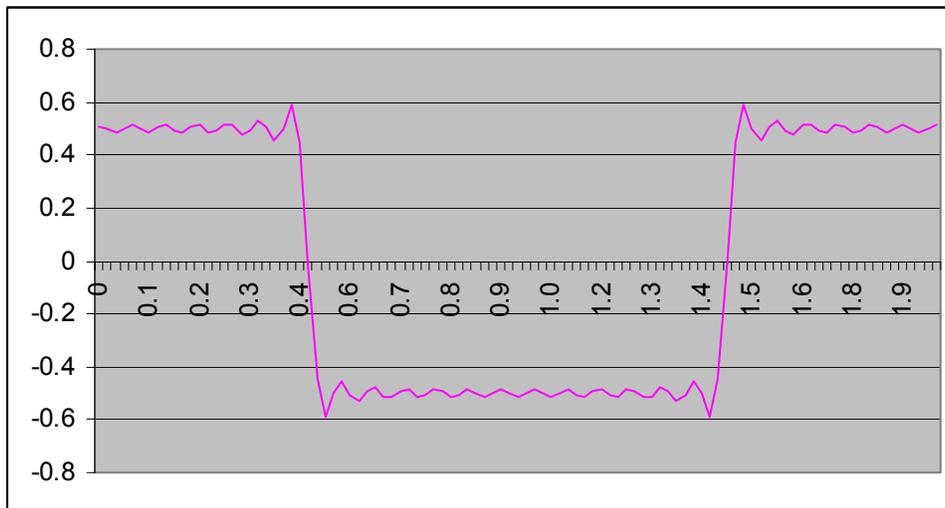
a. Suppose that the signal is applied to an ideal low-pass filter with bandwidth 15 Hz. Plot the output from the low-pass filter and compare to the original signal. Repeat for 5 Hz; for 3 Hz. What happens as W increases?

If we expand the above series to obtain the first few terms of $g(t)$, we get:

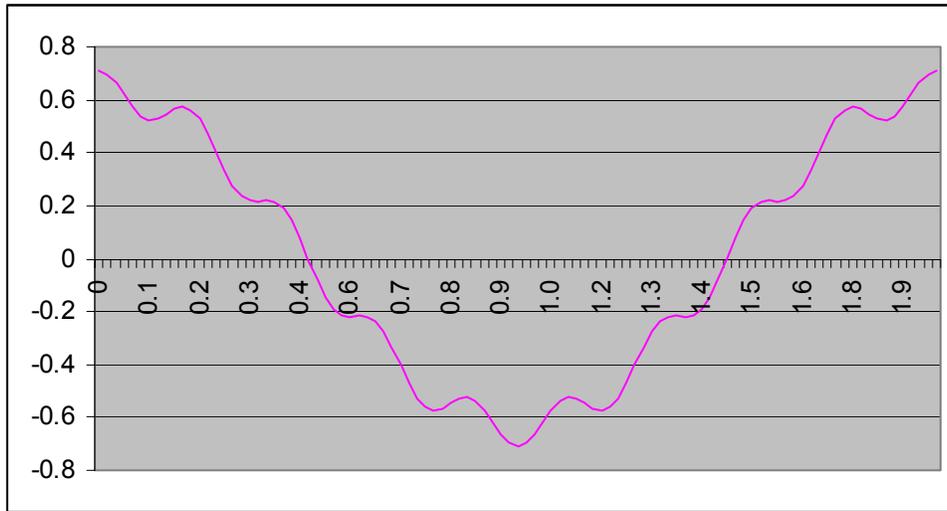
$$g(t) = 2/\pi \{ \cos\pi t - (1/3) \cos3\pi t + (1/5) \cos5\pi t - (1/7) \cos7\pi t + (1/9) \cos9\pi t - \dots \}$$

So frequencies of the components in Hz are 1/2, 3/2, 5/2, 7/2, 9/2, and so on.

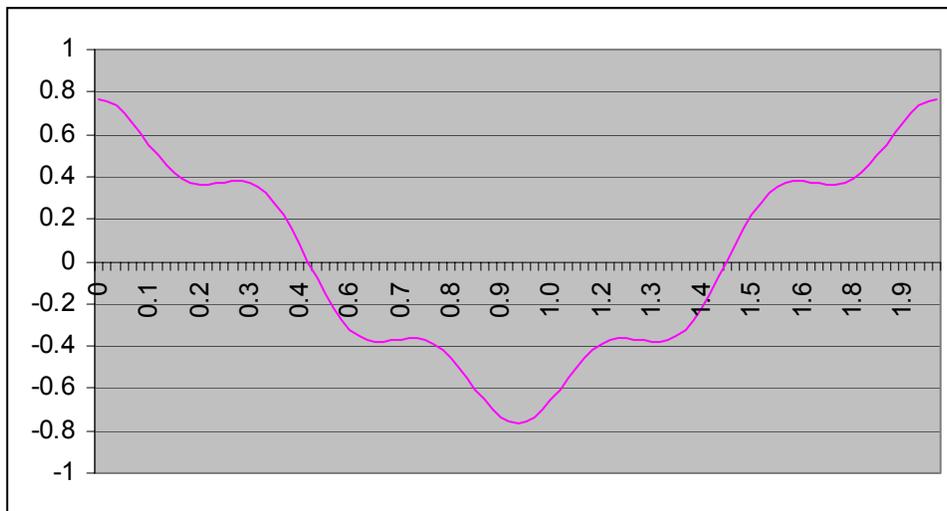
If the signal is applied to a low pass filter with bandwidth 15 Hz, then the output will consist of the sum of all components up to frequency 14.5 Hz, that is, up to component $(2k + 1)/2 = 14.5$, which gives $k = 14$. The resulting signal is shown below.



If the lowpass filter has bandwidth 5 Hz, then the first five components are passed ($k=4$). The resulting signal is shown below:

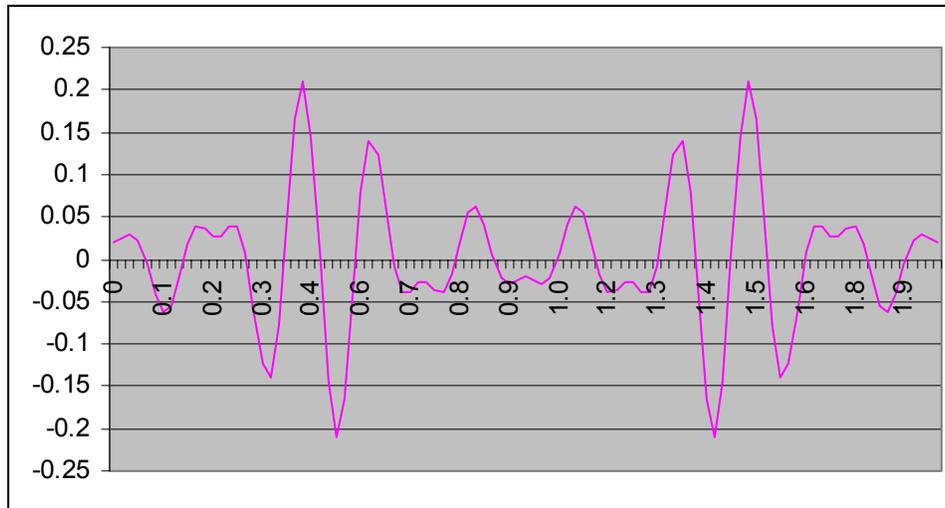


If the lowpass filter has bandwidth 3 Hz, then the first three components are passed ($k=2$). The resulting signal is shown below:



b. Suppose that the signal is applied to a bandpass filter which passes frequencies from 5 to 9 Hz. Plot the output from the filter and compare to the original signal.

In this case the filter passes the components corresponding to $k = 5, 6, 7,$ and 8 . The resulting signal is shown below:



While the three signals in part (a) are seen to be converging to a given signal, the signal obtained in part (b) shows no clear relationship to the signals in part (a)

26. Suppose that the 8 kbps periodic signal in Figure 3.25 is transmitted over a system that has an attenuation function that is equal to 1 for all frequencies, and a phase function that is equal to -90° for all frequencies. Plot the signal that comes out of this system. Does it differ in shape from the input signal?

Solution:

Note that the Fourier series in the book has a typographical error, the coefficients should be divided by n :

$$-0.5 + (4/\pi) \{ \sin(\pi/4) \cos(2\pi 1000t) + \sin(2\pi/4) \cos(2\pi 2000t)/2 + \sin(3\pi/4) \cos(2\pi 3000t)/3 + \dots \}$$

The signal that results from transmission over the above system changes the phase of each function by -90° or equivalently $-\pi/2$ radians. The series for this signal is:

$$-0.5 + (4/\pi) \{ \sin(\pi/4) \cos(2\pi 1000t - \pi/2) + \sin(2\pi/4) \cos(2\pi 2000t - \pi/2)/2 + \sin(3\pi/4) \cos(2\pi 3000t - \pi/2)/3 + \dots \}$$

Or:

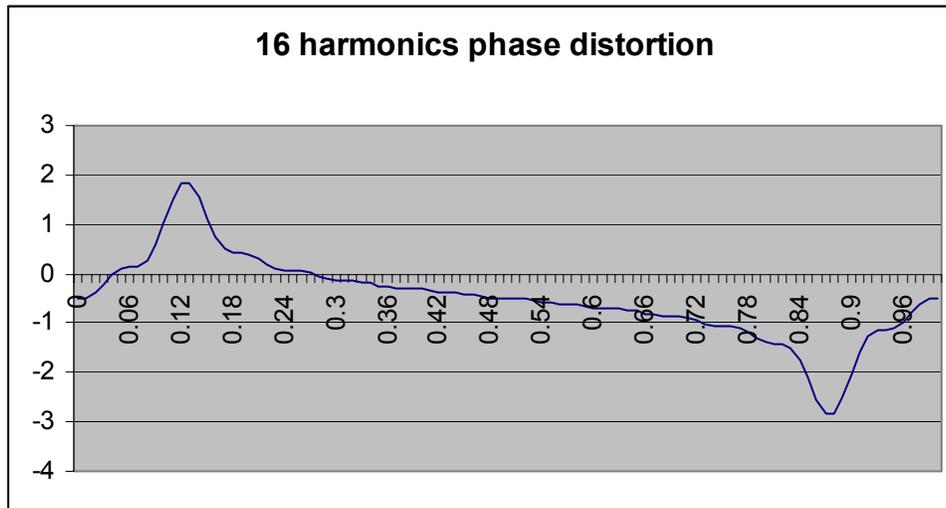
$$-0.5 + (4/\pi) \{ \sin(\pi/4) \cos 2\pi(1000t - 1/4) + \sin(2\pi/4) \cos 2\pi(2000t - 1/4)/2 + \sin(3\pi/4) \cos 2\pi(3000t - 1/4)/3 + \dots \}$$

Or:

$$-0.5 + (4/\pi) \{ \sin(\pi/4) \cos 2\pi 1000(t - 1/4000) + \sin(2\pi/4) \cos 2\pi 2000(t - 1/8000)/2 + \sin(3\pi/4) \cos 2\pi 3000(t - 1/12000)/3 + \dots \}$$

As it is seen in the equation, different components get different time shifts in time domain for the same phase shift in the frequency domain.

The plot of the signal is shown below. It can be seen that the signal differs considerably from that in Figure 3.15.



27. A 10 kHz baseband channel is used by a digital transmission system. Ideal pulses are sent at the Nyquist rate and the pulses can take 16 levels. What is the bit rate of the system?

Solution:

Nyquist pulses can be sent over this channel at a rate of 20000 pulses per second. Each pulse carries $\log_2 16 = 4$ bits of information, so the bit rate is 80000 bits per second.

28. Suppose a baseband transmission system is constrained to a maximum signal level of ± 1 volt and that the additive noise that appears in the receiver is uniformly distributed between $[-1/16, 1/16]$. How many levels of pulses can this transmission system use before the noise starts introducing errors?

Solution:

If two adjacent signal levels are separated by more than $2/16$ then the noise cannot translate one adjacent signal into the next. The maximum range that the signal can span is $+1 - (-1) = 2$, so the maximum number of levels is $2/(1/8) = 16$.

29. What is the maximum reliable bit rate possible over a telephone channel with the following parameters?

Solutions follow questions:

a. $W = 2.4$ kHz SNR = 20 dB

An SNR of 20 dB corresponds to a value of 100. The channel capacity formula then gives

$$C = 2400 \log_2(1 + 100) = 15979 \text{ bps.}$$

b. $W = 2.4$ kHz SNR = 40 dB

$$C = 2400 \log_2(1 + 10000) = 31890 \text{ bps.}$$

c. $W = 3.0$ kHz SNR = 20 dB

$$C = 3000 \log_2(1 + 100) = 19974 \text{ bps.}$$

d. $W = 3.0$ kHz SNR = 40 dB

$$C = 3000 \log_2(1 + 10000) = 39863 \text{ bps.}$$

30. Suppose we wish to transmit at a rate of 64 kbps over a 3 kHz telephone channel. What is the minimum SNR required to accomplish this?

Solution:

We know that $R = 64$ kbps and $W = 3$ kHz. What we need to find is SNR_{\min} . The channel capacity is:

$$C = W \log_2(1 + \text{SNR}), C \geq C_{\min} = 64 \text{ kbps}$$

$$C_{\min} = W \log_2(1 + \text{SNR}_{\min}) \Rightarrow \log_2(1 + \text{SNR}_{\min}) = 64/3 \Rightarrow 1 + \text{SNR}_{\min} = 2^{64/3}$$

$$\Rightarrow \text{SNR}_{\min} = 2.64 \times 10^6$$

$$\text{in dB: } \text{SNR}_{\min} = 10 \log_{10}(2.64 \times 10^6) = 64.2 \text{ dB} \therefore \text{ a very clean channel}$$

31. Suppose that a low-pass communications system has a 1 MHz bandwidth. What bit rate is attainable using 8-level pulses? What is the Shannon capacity of this channel if the SNR is 20 dB? 40 dB?

Solution:

Nyquist pulses can be sent over this system at a rate of 2 million pulses per second. Eight-level signaling carries 3 bits per pulse, so the bit rate is 6 Mbps.

The Shannon capacities are:

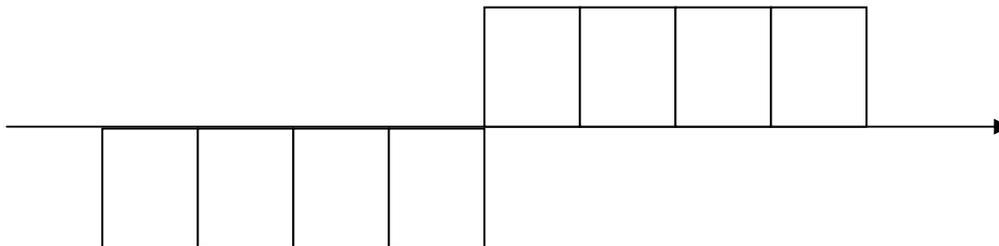
$$C = 1000000 \log_2(1 + 100) = 6.6 \text{ Mbps.}$$

$$C = 1000000 \log_2(1 + 10000) = 13.3 \text{ Mbps.}$$

32. Most digital transmission systems are “self-clocking” in that they derive the bit synchronization from the signal itself. To do this the systems use the transitions between positive and negative voltage levels. These transitions help define the boundaries of the bit intervals.

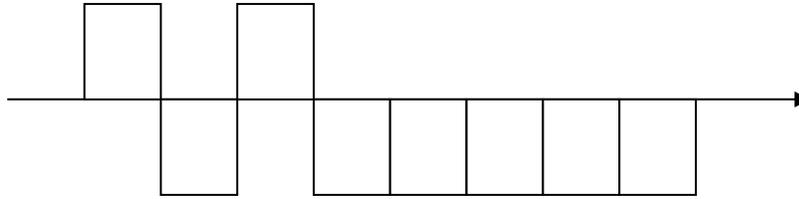
Solutions follow questions:

- a. The nonreturn-to-zero (NRZ) signaling method transmits a 0 with a +1 voltage of duration T , and a 1 with a -1 voltage of duration T . Plot the signal for the sequence n consecutive 1s followed by n consecutive 0s. Explain why this code has a synchronization problem.



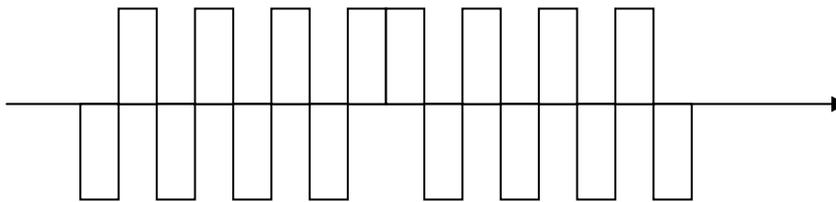
The above figure shows a sequence of 4 1s followed by 4 0s. A long sequence of 1s or a long sequence of 0s produces a long period during which there is no change in the signal level. Consequently, there are no transitions (“zero crossings”) that help a synchronization circuit determine where the boundary of each signaling interval is located.

- b. In differential coding the sequence of 0s and 1s induces changes in the polarity of the signal; a binary 0 results in no change in polarity, and a binary 1 results in a change in polarity. Repeat part (a). Does this scheme have a synchronization problem?



The occurrence of a “1” induces a transition and helps synchronization. However sequences of “0s” still result in periods with no transitions.

- c. The Manchester signaling method transmits a 0 as a +1 voltage for $T/2$ seconds followed by a -1 for $T/2$ seconds; a 1 is transmitted as a -1 voltage for $T/2$ seconds followed by a +1 for $T/2$ seconds. Repeat part (a) and explain how the synchronization problem has been addressed. What is the cost in bandwidth in going from NRZ to Manchester coding?



Every T -second interval now has a transition in the middle, so synchronization is much simpler. However, the bandwidth of the signal is doubled, as pulses now are essentially half as wide, that is, $T/2$ seconds.

33. Consider a baseband transmission channel with a bandwidth of 10 MHz. What bit rates can be supported by the bipolar line code and by the Manchester line code?

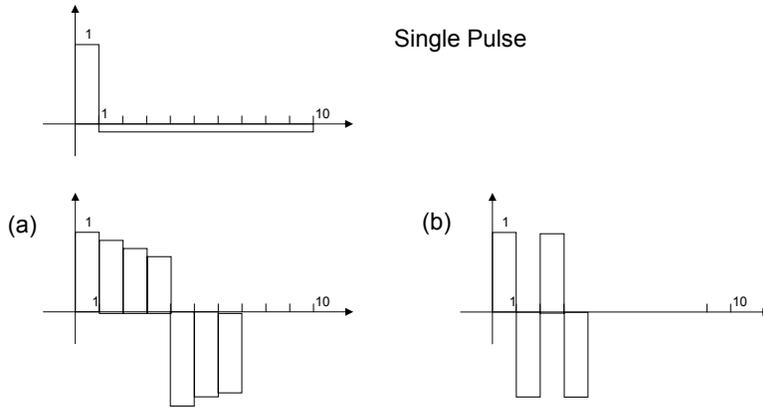
Solution:

From Figure 3.26 we see that a bipolar code with pulses T -seconds wide occupies a bandwidth of $W = 1/T$ Hz. Therefore a 10 MHz bandwidth allows a signaling rate of 10 megabits/second.

From the figure it can also be seen that a Manchester code occupies twice the bandwidth. Hence a 10 MHz bandwidth allows a signaling rate of 5 megabits/second.

34. The impulse response in a T-1 copper-wire transmission system to the input pulse has the idealized form where the initial pulse is of amplitude 1 and duration 1, and the afterpulse is of amplitude -0.1 and duration 10.

- a. Let $\delta(t)$ be the narrow input pulse in Figure 3.27. Suppose we use the following signaling method: Every second, the transmitter accepts an information bit; if the information bit is 0, then $-\delta(t)$ is transmitted, and if the information bit is 1, then $\delta(t)$ is transmitted. Plot the output of the channel for the sequence 1111000. Explain why the system is said to have “dc” or baseline wander.



“DC” stands for direct current and indicates a constant signal level. In a sequence of pulses of the same polarity we would expect to obtain a constant signal level. Instead we get a level in which the signal level drifts in the direction of opposite polarity. For this reason, the term “dc wander” is used.

- b. The T-1 transmission system uses bipolar signaling in the following fashion: If the information bit is a 0, then the input to the system is $0 \cdot \delta(t)$; if the information bit is a 1, then the input is $\delta(t)$, for an even occurrence of a 1, and $-\delta(t)$ for an odd occurrence of a 1. Plot the output of the channel for the sequence 1111000. Explain how this signaling has solved the “dc” or baseline wander problem.

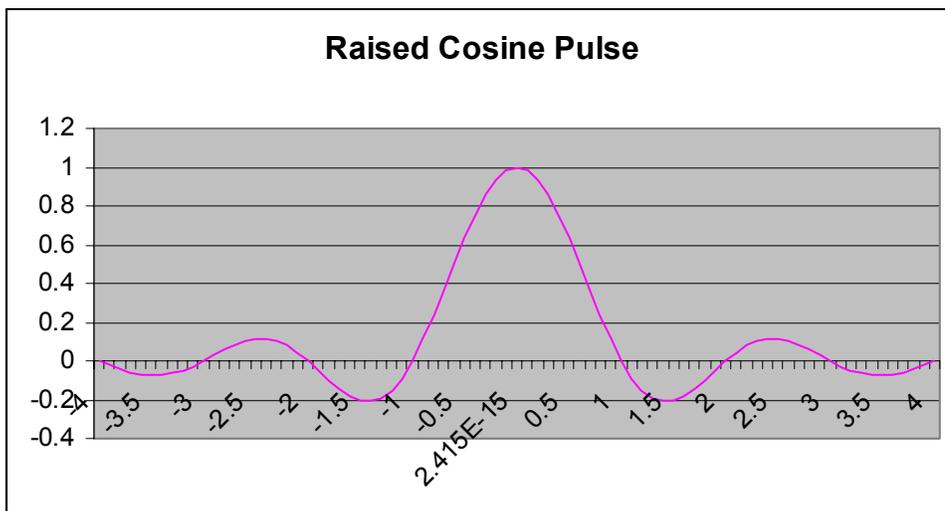
Each pair of “1” produces “tails” of opposite polarity that cancel out.

- 35. The raised cosine transfer function, shown in Figure 3.31, has a corresponding impulse response given by:

$$p(t) = \frac{\sin(\pi/T)}{\pi/T} \frac{\cos(\pi\alpha t/T)}{1-(2\alpha t/T)^2}$$

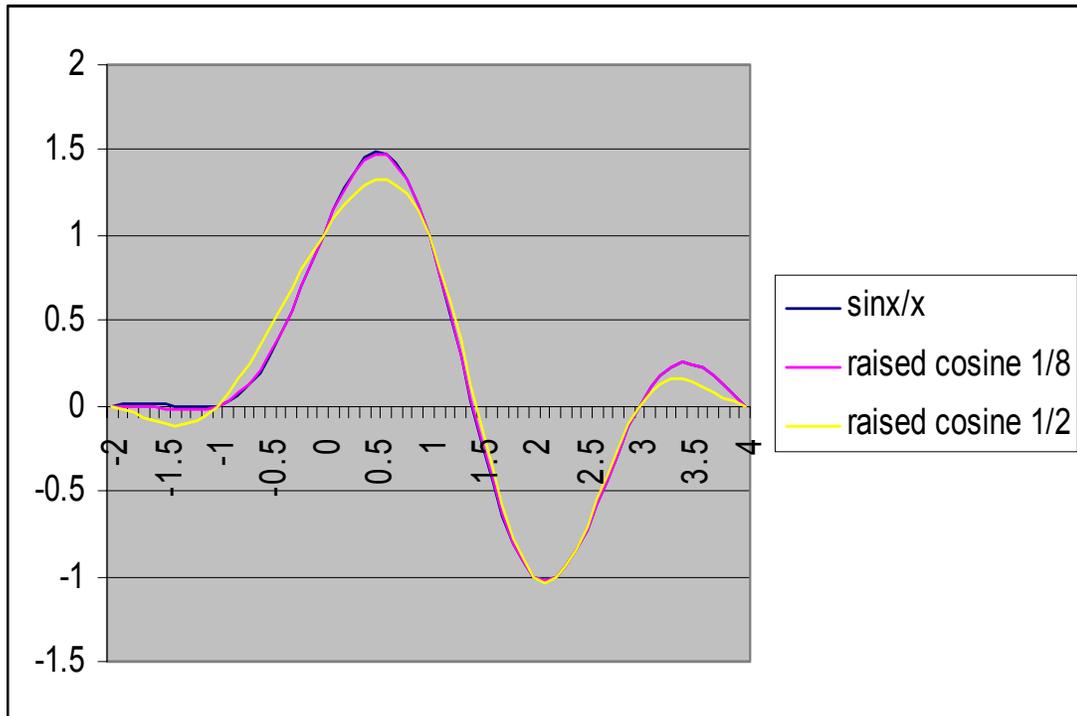
Solutions follow questions:

- a. Plot the response of the information sequence 1010 for $\alpha = 1/2$; $\alpha = 1/8$.



- b. Compare this plot to the response, using the pulse in Figure 3.27.

The plot below for the sequence 110 compares the $\sin x / x$ pulse with the raised cosine pulse with $\alpha = 1/8, 1/2$. It can be seen that the signal levels are 1 or -1 at the appropriate time instants. It can also be seen that the $\alpha = 1/2$ pulse lower amplitudes in its oscillation which is indicative of its lower intersymbol interference at non-integer values of T. This improved performance is obtained at the cost of 50% more bandwidth.



36. Suppose a CATV system uses coaxial cable to carry 100 channels, each of 6 MHz bandwidth. Suppose that QAM modulation is used.

QAM modulation with 2^m point $\Rightarrow m$ bits/pulse

$T =$ pulse duration $= 1/W \Rightarrow$ pulse rate $=$ baud rate $= W = 6 \times 10^6$

- a. What is the bit rate/channel if a four-point constellation is used? eight-point?

$R = mW$, if $m = 4 \Rightarrow R = 24$ Mbps. If $m = 8$, then $R = 48$ Mbps

- b. Suppose a digital TV signal requires 4 Mbps. How many digital TV signals can each channel handle for the two cases in part (a)?

$R_{DTV} = 4$ Mbps $\Rightarrow N_{DTV} = R/R_{DTV}$, where N is the number of digital TV signals per channel

If $m = 4 \Rightarrow N_{DTV} = 6$; if $m = 8 \Rightarrow N_{DTV} = 12$.

37. Explain how ASK was used in radio telegraphy. Compare the use of ASK to transmit Morse code with the use of ASK to transmit text using binary information.

Solution:

In ASK, the amplitude of the transmitted sinusoidal signal corresponds to the signal value. In binary amplitude shift keying to transmit text, the presence or absence of the sinusoidal

signal corresponds to the transmission of a “1” or a “0”. In this approach the duration of a “1” pulse and a “0” pulse are the same.

Morse code also uses the presence of absence of a tone, but in addition also uses the length of the tone to convey a “dot” (short tone) and a “dash” (long tone). The absence of tone amounts to having a third symbol, “space”.

38. Suppose that a modem can transmit 8 distinct tones at distinct frequencies. Every T seconds the modem transmits an arbitrary combination of tones (that is, some are present, and some are not present).

Solutions follow questions:

a. What bit rate can be transmitted using this modem?

Each tone is either present or absent, hence there are 2^8 possible combinations of tones that can be transmitted every T seconds. The corresponding transmitted bit rate is $8/T$ bps.

b. Is there any relationship between T and the frequency of the signals?

Yes, there is a relationship. T must be long enough that enough of each sinusoid can be observed to determine its frequency. This implies that the periods of all the sinusoids must be less than T .

39. A phase modulation system transmits the modulated signal $A\cos(2\pi f_c t + \phi)$ where the phase ϕ is determined by the 2 information bits that are accepted every T -second interval:

for 00, $\phi = 0$; for 01, $\phi = \pi/2$; for 10, $\phi = \pi$; for 11, $\phi = 3\pi/2$.

Solutions follow questions:

a. Plot the signal constellation for this modulation scheme.

The transmitted signals corresponding to the phase values are as follows:

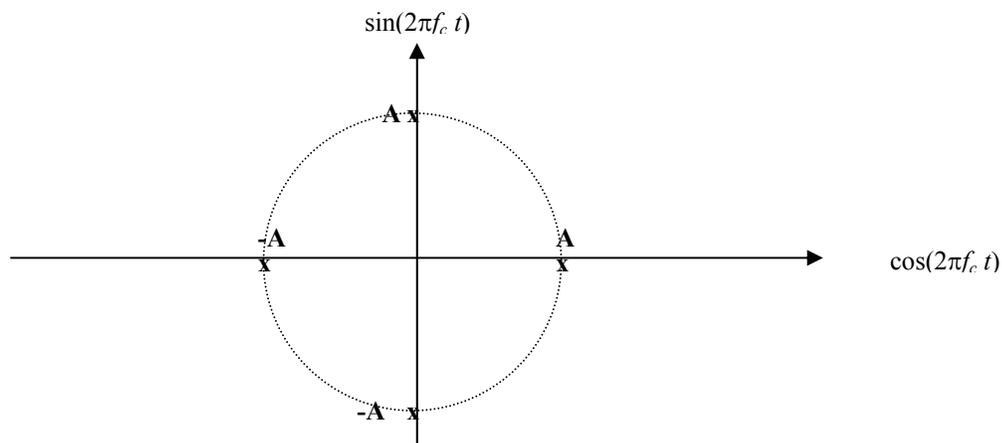
for 00 $\phi = 0$, so $x(t) = A\cos(2\pi f_c t)$

for 01 $\phi = \pi/2$, so $x(t) = A\cos(2\pi f_c t + \pi/2) = -A\sin(2\pi f_c t)$

for 10 $\phi = \pi$, so $x(t) = A\cos(2\pi f_c t + \pi) = -A\cos(2\pi f_c t)$

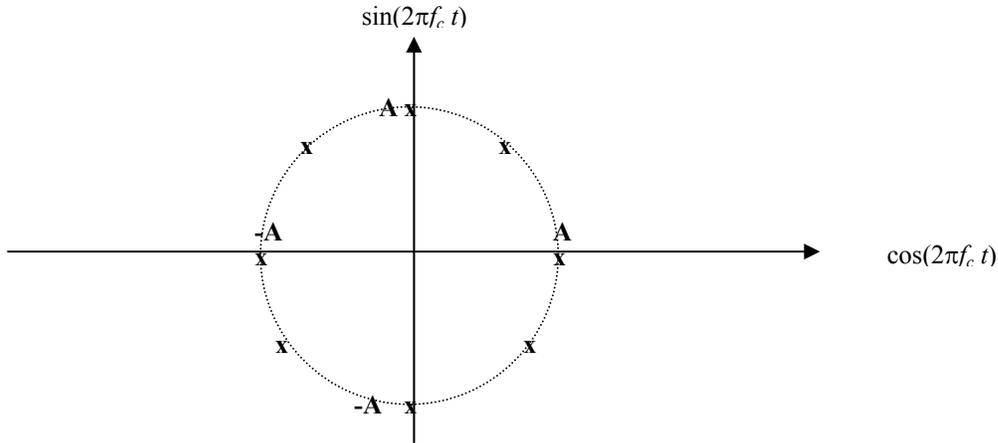
for 11 $\phi = 3\pi/2$, so $x(t) = A\cos(2\pi f_c t - \pi/2) = A\sin(2\pi f_c t)$

The signal constellation is shown below:



b. Explain how an eight-point phase modulation scheme would operate.

The generalization to an eight-point constellation is straightforward. In the above figure we can see that the four constellation points are placed at equidistant points in a circle about the origin. The figure below shows how eight points can be placed in a circle with angle $\pi/4$ between them.



40. Suppose that the receiver in a QAM system is not perfectly synchronized to the carrier of the received signal; that is, it multiplies the received signal by $2\cos(2\pi f_c t + \phi)$ and by $2\sin(2\pi f_c t + \phi)$ where ϕ is a small phase error. What is the output of the demodulator?

Solution:

The transmitted signal in QAM is $y(t) = A_k \cos(2\pi f_c t) + B_k \sin(2\pi f_c t)$

The upper multiplier in Figure 3.42 computes the following:

$$y(t) 2\cos(2\pi f_c t + \phi) = 2\cos(2\pi f_c t + \phi) A_k \cos(2\pi f_c t) + 2\cos(2\pi f_c t + \phi) B_k \sin(2\pi f_c t)$$

$$= A_k \{ \cos(\phi) + \cos(4\pi f_c t + \phi) \} + B_k \{ -\sin(\phi) + \sin(4\pi f_c t + \phi) \}$$

The lowpass filter removes the double-frequency component, so the output of the upper demodulator circuit is: $A_k \cos(\phi) - B_k \sin(\phi)$. When the phase error ϕ is small, then $\cos \phi$ is approximately 1, and $\sin \phi$ is approximately 0, so the phase error causes a small error in the demodulator output. As the phase error increases however, the desired signal A_k becomes harder to discern because the cosine term decreases and the sine term increases.

It can be similarly shown that the output of the lower demodulator is: $B_k \cos(\phi) + A_k \sin(\phi)$.

41. In differential phase modulation the binary information determines the *change* in the phase of the carrier signal $\cos(2\pi f_c t)$. For example, if the information bits are 00, the phase change is 0; if 01, it is $\pi/2$; for 10, it is π ; and for 11, it is $3\pi/2$.

a. Plot the modulated waveform that results from the binary sequence 01100011. Compare it to the waveform that would be produced by ordinary phase modulation as described in problem 39.

Assume that the initial phase is 0, so the initial signal is $\cos(2\pi f_c t)$.

Bit pair	phase change	net phase	modulated signal
01	$\pi/2$	$\pi/2$	$\cos(2\pi f_c t + \pi/2) = -\sin(2\pi f_c t)$
10	π	$3\pi/2 = -\pi/2$	$\cos(2\pi f_c t - \pi/2) = \sin(2\pi f_c t)$
00	0	$-\pi/2$	$\cos(2\pi f_c t - \pi/2) = \sin(2\pi f_c t)$
11	$3\pi/2 = -\pi/2$	$-\pi$	$\cos(2\pi f_c t - \pi) = -\cos(2\pi f_c t)$

The modulated signal looks the same as the signal produced by the scheme in problem 39. The difference is in the manner in which the information bits determine the transmitted signal.

b. Explain how can differential phase modulation be demodulated.

The received signal is demodulated by finding the phase difference between the signals in two consecutive intervals. For example, we can use a delayed version of the signal in an interval (say $\cos(2\pi f_c t + \pi/2)$ in the first interval) and use it to generate two reference signals that multiply the signal in the second interval ($\cos(2\pi f_c t - \pi/2)$) we obtain:

$$\cos(2\pi f_c t + \pi/2) \cos(2\pi f_c t - \pi/2) = 1/2 \{ \cos\pi + \cos(4\pi f_c t) \}$$

$$\sin(2\pi f_c t + \pi/2) \cos(2\pi f_c t - \pi/2) = 1/2 \{ \sin\pi + \sin(4\pi f_c t) \}$$

The pair of values $\cos\pi = -1$ and $\sin\pi = 0$ uniquely identify the phase difference as being π . The demodulator can then decode the binary pair 10. To check that you understand the solution try demodulating the fourth interval.

42. A new broadcast service is to transmit digital music using the FM radio band. Stereo audio signals are to be transmitted using a digital modem over the FM band. The specifications for the system are the following: Each audio signal is sampled at a rate of 40 kilosamples/second and quantized using 16 bits; the FM band provides a transmission bandwidth of 200 kiloHertz.

Solutions follow questions:

a. What is the total bit rate produced by each stereo audio signal?

The bit rate for each signal is 40 ksamples/sec x 16 bits/sample = 640 kbps. The bit rate for the pair of signals is then 1.28 Mbps.

b. How many points are required in the signal constellation of the digital modem to accommodate the stereo audio signal?

A transmission bandwidth of 200 kHz allows 200 kilopulses/second. To obtain a bit rate of 1.28 Mbps, we need to send $1280/200 = 6.4$ bits/pulse. If we use a $2^7 = 128$ point constellation, we can then meet the desired bit rate.

43. A twisted-wire pair has an attenuation of 0.7 dB/kilometer at 1 kHz.

Solutions follow questions:

a. How long can a link be if an attenuation of 20 dB can be tolerated?

$$20 \text{ db} / (0.7 \text{ db/km}) = 28 \text{ km}$$

b. A twisted pair with loading coils has an attenuation of 0.2 dB/kilometer at 1 kHz. How long can the link be if an attenuation of 20 dB can be tolerated?

$$20 \text{ db} / (0.2 \text{ db/km}) = 100 \text{ km.}$$

44. Use Figure 3.47 and Figure 3.50 to explain why the bandwidth of twisted-wire pairs and coaxial cable decreases with distance.

Solution:

The bandwidth is the range of frequencies where the channel passes a significant proportion of the power in the input signal. Both figures show that the attenuation when measured in dB/km increases with higher frequency. For example, the attenuation for 19-gauge wire at 100 kHz is about 5 dB/km and at 1 MHz it is about 15 dB/km. This implies that the relative attenuation between a lower frequency and a higher frequency increases with distance. Using the same example, we have that at 1 km and at 10 km, the attenuation at 100 kHz is 5 dB and 50 dB respectively, but at 1 MHz it is 15 dB and 150 dB respectively. Thus at longer distances higher frequencies are attenuated much more severely and consequently the bandwidth decreases with increasing distance.

45. Calculate the bandwidth of the range of light covering the range from 1200 nm to 1400 nm. How many Hz per person are available if the population of the world is six billion people? Repeat for 1400 nm to 1600 nm. Keep in mind that the speed of light in fiber is approximately 2×10^8 m/sec.

Solution:

Frequency and wavelength are related by the expression: $f = v/\lambda$, where $v = 2 \times 10^8$. The frequencies corresponding to 1200 nm and 1400 nm are $f_1 = v/\lambda_1 = 2(10^8)/1.2(10^{-6})$ and $f_2 = v/\lambda_2 = 2(10^8)/1.4(10^{-6})$ so the bandwidth is $BW = f_1 - f_2 = 23.8$ THz.

The Hz/person is $23.8 \times 10^{12} / 6 \times 10^9 = 5.95 \times 10^3$ Hz or approximately 6 KHz per person.

For the band from 1400 nm to 1600 nm we get $BW = f_1 - f_2 = 2(10^8)(1/1.4 - 1/1.6)10^6 = 17.8$ THz. The Hz per person is $17.8 \times 10^{12} / 6 \times 10^9 = 2.96 \times 10^3$ Hz or approximately 3 KHz per person.

46. Suppose that we wish to delay an optical signal by 1 nanosecond. How long a length of optical fiber is needed to do this? How much is the signal attenuated? Repeat for 1 millisecond.

Solution:

The length of fiber for 1 nanosecond delay is: $(2 \times 10^8) \times (1 \times 10^{-9}) = 0.2$ meter. Considering 0.2 dB per kilometer attenuation, the signal is attenuated by $0.2 \times (2 \times 10^{-4})$ or 4×10^{-5} dB.

For 1 millisecond delay the length of the fiber is: $(2 \times 10^8) \times (1 \times 10^{-3}) = 200$ Km. The signal is attenuated by $0.2 \times 200 = 40$ dB.

47. Compare the attenuation in a 100 km link for optical fibers operating at 850 nm, 1300 nm, and 1550 nm.

Solution:

At 850 nm the attenuation is about 2 dB per km, so the attenuation in 100 km is 200 dB. At 1300 nm, the attenuation is $0.5(100) = 50$ dB. At 1500 nm, the attenuation is $0.2(100) = 20$ dB.

48. The power of an optical signal in dBm is defined as $10 \log_{10}P$ where P is in milliwatts.

a. What is the power in milliwatts of a -30 dBm signal? 6dBm signal?

Solution:

For a -30 dBm signal: $10 \log_{10}P = -30$ dBm or $\log_{10}P = -3$ then $P = 10^{-3}$ milliwatts

For a 6 dBm signal: $10 \log_{10} P = 6 \text{ dBm}$ or $\log_{10} P = 0.6$ then $P = 10^{0.6} = 3.98$ milliwatts

b. What is the power in dBm of 1 microwatt signal?

Solution:

$P = 1 \text{ microwatt} = 10^{-3} \text{ milliwatt}$ So the power in dBm is $10 \log_{10} 10^{-3} = -30 \text{ dBm}$

c. What is the power of an optical signal if initially it is 2 mW and then undergoes attenuation by 10 dB?

Solution:

Initial power = $10 \log_{10} 2 = 3 \text{ dBm}$

Power after attenuation = $3 - 10 = -7 \text{ dBm}$

49. A 10 dBm optical signal propagates across N identical devices. What is the output signal power if the loss per device is 1 dB? (See problem 3.48.)

Solution:

The output signal power is: $10 - N$

50. Suppose that WDM wavelengths in the 1550 nm band are separated by 0.8 nm. What is the frequency separation in Hz? What is an appropriate bit rate for signals carried on these wavelengths? Repeat for 0.4 nm and 0.2 nm.

Solution:

For $\Delta\lambda = 0.8 \text{ nm}$: $B = (v \cdot \Delta\lambda) / \lambda_c^2 = (2 \times 10^8 \times 0.8 \times 10^{-9}) / (1550 \times 10^{-9})^2 = 66.6 \times 10^9 \text{ Hz} = 66.6 \text{ GHz}$

Bit rates up to 66 Gbps can be carried on these wavelengths. An appropriate bit rate could be 40 Gbps or 10 Gbps.

For $\Delta\lambda = 0.4 \text{ nm}$: $B = 66.6 / 2 = 33.3 \text{ GHz}$, good for bit rates up to 33 Gbps. An appropriate bit rate could be 10 Gbps or 2.4 Gbps.

For $\Delta\lambda = 0.2 \text{ nm}$: $B = 33.3 / 2 = 16.6 \text{ GHz}$, good for bit rates up to 16 Gbps. An appropriate bit rate could be 2.4 Gbps.

51. Can WDM be used for simultaneous transmission of optical signals in opposite directions?

Solution:

WDM can be used for simultaneous transmission in opposite directions, however the regenerators and amplifiers along the path should support both directions.

52. Explain how prisms and prismlike devices can be used in WDM systems.

Solution:

Prisms and prism-like devices can be used for combining and splitting different wavelengths at the source, destination, or intermediate amplification or regeneration stages.

53. Compare the transition from analog repeaters to digital regenerators for copper-based transmission systems to the current transition from single-wavelength digital regenerator optical systems to multiwavelength optically amplified systems? What is the same and what is different? What is the next transition for optical transmission systems?

Solution:

The use of digital regenerators in copper-based transmission systems eliminates the accumulation of noise that takes place in analog systems, therefore makes it possible to transmit to longer distances. This results in lower overall system cost. Similarly, the use of optical amplifiers in optical systems makes it possible to transmit to longer distances which also results in lower overall system cost. Multiwavelength optically amplified systems provide for transmission of DWDM signals over long distances, which allows transmission and amplification of several wavelengths. This results in more savings because DWDM regenerators require splitting the wavelengths, regenerating individual wavelengths, and combining them again, which is more expensive.

EDFA optical amplifiers amplify the optical signal, while the digital regenerators actually regenerate and retransmit the digital signal. Also, unlike the digital repeaters, the number of EDFA optical amplifiers that can be cascaded is limited. Eventually the optical signal needs to be converted to electrical form and regenerated.

The next transition for optical transmission systems is the introduction of all-optical regeneration.

54. Suppose a network provides wavelength services to users by establishing end-to-end wavelengths across a network. A wavelength converter is a device that converts an optical signal from one wavelength to another wavelength. Explain the role of wavelength converters in such a network.

Solution:

Assume multiple end-to-end connections have to share the same path in a link across the network. If two or more of the connections use the same wavelength, only one of them can be established through the path. This limitation can be eliminated by using wavelength converters at the nodes of the network.

55. A satellite is stationed approximately 36,000 km above the equator. What is the attenuation due to distance for the microwave radio signal?

Solution:

The attenuation for free space is proportional to $2(10)\log_{10}d = 2(10)\log_{10}36000000 = 151$ dB.

56. Suppose a transmission channel operates at 3 Mbps and that it has a bit error rate of 10^{-3} . Bit errors occur at random and independent of each other. Suppose that the following code is used. To transmit a 1, the codeword 111 is sent; To transmit a 0, the codeword 000 is sent. The receiver takes the three received bits and decides which bit was sent by taking the majority vote of the three bits. Find the probability that the receiver makes a decoding error.

Solution:

The receiver makes a decoding error if two or more out of the three bits are in error. Therefore,

$$P_{error} = 3p^2(1 - p) + p^3 = 3(10^{-3})^2(1 - 10^{-3}) + (10^{-3})^3 \approx 3(10^{-6})$$

57. An early code used in radio transmission involved using codewords that consist of binary bits and contain the same number of 1s. Thus, the 2-out-of-5 code only transmits blocks of 5 bits in which 2 bits are 1 and the others 0.

Solutions follow questions:

a. List the valid codewords.

11000	01010
10100	01001
10010	00110
10001	00101
01100	00011

b. Suppose that the code is used to transmit blocks of binary bits. How many bits can be transmitted per codeword?

There are 10 possible codewords. Three bits per codeword can be transmitted if eight codewords are used.

c. What pattern does the receiver check to detect errors?

Each received codeword should have exactly two bits that are ones and three bits that are zeros to be a valid codeword.

d. What is the minimum number of bit errors that cause a detection failure?

A valid codeword can be changed into another valid codeword by changing a 1 to a 0 and a 0 to a 1. Therefore, two bit errors can cause a detection failure.

58. Find the probability of error-detection failure for the code in the problem 3.57 for the following channels:

Solutions follow questions:

a. The random error vector channel.

Each error vector can convert a code word to a different word. 10 of the 32 resulted codes are the 10 code words caused by 10 of the error vectors, excluding the all zero vector that converts the code word to itself. For example, for the first code word, 11000, the nine error vectors include the vectors that convert one 1 to 0 and one 0 to 1, 6 vectors in total including 10100, 10010, 10001, 01100, 01010, and 01001, and the vectors that convert two 1s to 0 and two 0s to 1s, 3 in total including 11110, 11101, and 11011, adding up to 9. Therefore 9/32 is the fraction of error vectors that cause detection failure.

b. The random bit error channel.

The set of error patterns that cause detection failure are the same as in part a. The probability of occurrence of each one of the first 6 patterns is $p^2(1-p)^3$ and the probability of occurrence of each one of the next 3 patterns is $p^4(1-p)^1$. So the probability of detection failure is $6p^2(1-p)^3 + 3p^4(1-p)$.

59. Suppose that two check bits are added to a group of $2n$ information bits. The first check bit is the parity check of the first n bits, and the second check bit is the parity check of the second n bits.

Solutions follow questions:

a. Characterize the error patterns that can be detected by this code.

If we rearrange the $2n$ information bits and the 2 parity bits into two codewords, each consisting of n information bits and a parity bit, we see that in effect we have divided the overall codeword into two subcodewords of length $n + 1$. An error pattern is detectable if the error pattern in the first codeword and in the second codeword are both detectable. An error pattern in each subcodeword is detectable if the number of errors in the subcodeword is odd. Therefore an error pattern is detectable if the number of errors in both subcodewords is odd.

- b. Find the error detection failure probability in terms of the error-detection probability of the single parity check code.

An error detection failure occurs is either the first subcodeword or the second subcodeword or both fail. Let $P_{\text{detect}}(n + 1)$ be the probability of successful error detection in the single parity code of length $n + 1$, then

$$\begin{aligned} P[\text{detection failure in code of length } 2n + 2] &= \\ &= 1 - P[\text{successful detection in both codes of length } n + 1] \\ &= 1 - P_{\text{detect}}(n + 1) P_{\text{detect}}(n + 1) \end{aligned}$$

- c. Does it help to add a third parity check bit that is the sum of the all the information bits?

Note that the above scheme fails to detect the case where there is an even number of errors in the first subcodeword and an even number of errors in the second subcodeword. An overall parity bit would not detect this error as the number of errors in the overall code word is still even and not detectable by a single parity bit.

60. Let $g(x)=x^3+x+1$. Consider the information sequence 1001.

Solutions follow questions:

- a. Find the codeword corresponding to the preceding information sequence.

Using polynomial arithmetic we obtain:

$$\begin{array}{r} 1011 \quad \left| \begin{array}{r} 1010 \\ 1001000 \\ \underline{1011} \\ 01000 \\ \underline{1011} \\ 00110 \end{array} \right. \end{array}$$

Codeword = 1001110

- b. Suppose that the codeword has a transmission error in the first bit. What does the receiver obtain when it does its error checking?

$$\begin{array}{r} 1011 \quad \left| \begin{array}{r} 0001 \\ 0001110 \\ \underline{1011} \\ 101 \end{array} \right. \end{array}$$

CRC calculated by Rx = 101 → error

61. ATM uses an eight-bit CRC on the information contained in the header. The header has six fields:
 First 4 bits: GFC field
 Next 8 bits: VPI field
 Next 16 bits: VCI field
 Next 3 bits: Type field

63. Let $g_1(x) = x + 1$ and let $g_2(x) = x^3 + x^2 + 1$. Consider the information bits (1,1,0,1,1,0).

- a. Find the codeword corresponding to these information bits if $g_1(x)$ is used as the generating polynomial.

$$\begin{array}{r}
 11 \quad \left| \begin{array}{r}
 100100 \\
 \hline
 1101100 \\
 11 \\
 \hline
 0011 \\
 11 \\
 \hline
 0000
 \end{array} \right.
 \end{array}$$

Codeword = 1101100

- b. Find the codeword corresponding to these information bits if $g_2(x)$ is used as the generating polynomial.

$$\begin{array}{r}
 1101 \quad \left| \begin{array}{r}
 100011 \\
 \hline
 110110000 \\
 1101 \\
 \hline
 01000 \\
 1101 \\
 \hline
 1010 \\
 1101 \\
 \hline
 111
 \end{array} \right.
 \end{array}$$

Codeword = 110110111

- c. Can $g_2(x)$ detect single errors? double errors? triple errors? If not, give an example of an error pattern that cannot be detected.

Single errors can be detected since $g_2(x)$ has more than one term. Double errors *cannot* be detected even though $g_2(x)$ is primitive because the codeword length exceeds $2^{n-k}-1=7$. An example of such undetectable error is 1000000010. Triple errors cannot be detected since $g_2(x)$ has only three terms.

- d. Find the codeword corresponding to these information bits if $g(x) = g_1(x) g_2(x)$ is used as the generating polynomial. Comment on the error-detecting capabilities of $g(x)$.

$$\begin{array}{r}
 10111 \quad \left| \begin{array}{r}
 111101 \\
 \hline
 1101100000 \\
 10111 \\
 \hline
 11000 \\
 10111 \\
 \hline
 11110 \\
 10111 \\
 \hline
 10010 \\
 10111 \\
 \hline
 010100 \\
 10111 \\
 \hline
 0011
 \end{array} \right.
 \end{array}$$

Codeword = 1101100011

The new code can detect all single and all odd errors. It cannot detect double errors. It can also detect all bursts of length $n - k = 4$ or less. All bursts of length 5 are detected except for the burst that equals $g(x)$. The fraction $1/2^{n-k} = 1/16$ of all bursts of length greater than 5 are detectable.

64. Take any binary polynomial of degree 7 that has an even number of nonzero coefficients. Show by long hand division that the polynomial is divisible by $x+1$.

Solution:

Let $p(x) = x^7 + x^5 + x^2 + 1$, then $p(x) = (x + 1)(x^6 + x^5 + x + 1)$.

65. A repetition code is an $(n,1)$ code in which the $n - 1$ parity bits are repetitions of the information bit. Is the repetition code a linear code? What is the minimum distance of the code?

Solution:

Yes this is a linear code in which: $c_2 = c_1, c_3 = c_1, \dots, c_n = c_1$, where c_1 is the information bit. This code has two codewords: $(0,0,\dots,0)$ and $(1,1,\dots,1)$ so the minimum distance is $d_{min} = n$.

66. A transmitter takes groups of K groups of k information bits and appends a single parity bit to each group. It then appends a block parity check word in which the j th bit in the check word is the modulo 2 sum of the j th components in the K codewords.

Solutions follow questions:

- a. Explain why this is a $((K + 1)(k + 1), Kk)$ linear code.

Every binary linear code is represented by two parameters n and k . n is the total number of bits transmitted (codeword), and k is the number of information bits. The transmitter takes K groups of k information bits, having a total of Kk information bits (k). The code appends one check bit to each row of k information bits, and appends one check bit to each $k + 1$ column, transmitting a total of $(k + 1)(K + 1)$ codeword bits (n). The result is a $[(k + 1)(K + 1), Kk]$ linear code, in which the set of check bits are derived from the Kk information bits by $K+k+1$ equations.

- b. Write the codeword as a $(k + 1)$ row by $(K + 1)$ column array in which the first K columns are the codewords and the last column is the block parity check. Use this code to show how the code can detect all single, double, and triple errors. Give an example of a quadruple error that cannot be detected.

The $(k + 1)(K + 1)$ encoded matrix satisfies the pattern that all the rows and columns have even parity. If one, two, or three bit errors occur anywhere in the matrix of bits, then at least one row or column parity check will detect the errors. A set of four errors placed so that they form a rectangle in the array will result in column check bits and row check bits that are zero and hence fail to be detected.

- c. Find the minimum distance of the code. Can it correct all single errors? If so explain how the decoding can be done.

The minimum distance is $d_{min} = 4$. Yes, the code can correct all single errors. If a single error occurs in row i and column j , then the check bits in row i and column j will fail, indicating that the bit in that location should be complemented.

- d. Find the probability of error-detection failure for the random bit error channel.

A pattern with exactly i bit errors occurs with probability $p^i(1 - p)^{n-i} = (1 - p)^n (p/(1 - p))^i$. For $p < 1/2$, $p / (1 - p) < 1$, so the above probability decreases with increasing value of i . Therefore the error terms with $i = d_{min}$ errors will have the largest probability. We can therefore estimate the probability of error-detection failure by the probability that all rectangular error patterns occur. We first count the number of possible rectangular error patterns:

The number of ways of picking the two rows is $\binom{k+1}{2} = (k+1)k/2$.

The number of ways of picking the two columns is $\binom{K+1}{2} = (K+1)K/2$.

Therefore the total number of rectangular error patterns is $(k+1)k(K+1)K/4$, and the probability of undetectable error pattern is approximately:

$$P_{\text{detection failure}} \approx Kk(K+1)(k+1)p^4(1-p)^{(K+1)(k+1)-4}$$

67. Consider the $m = 4$ Hamming code.

a. What is n , and what is k for this code?

$$n = 2^m - 1 = 15; \quad k = n - m = 11$$

(15,11) Hamming code

b. Find parity check matrix for this code.

$$\mathbf{H} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 \end{bmatrix}$$

c. Give the set of linear equations for computing the check bits in terms of the information bits.

$$\begin{aligned} b_{12} &= b_5 + b_6 + b_7 + b_8 + b_9 + b_{10} + b_{11} \\ b_{13} &= b_1 + b_2 + b_3 + b_8 + b_9 + b_{10} + b_{11} \\ b_{14} &= b_2 + b_3 + b_4 + b_6 + b_7 + b_{10} + b_{11} \\ b_{15} &= b_1 + b_3 + b_4 + b_5 + b_7 + b_9 + b_{11} \end{aligned}$$

d. Write a program to find the set of all codewords. Do you notice anything peculiar about the weights of the codewords?

The minimum weight is 3.

For every w between 1 and 15 there is at least one code word with the weight w except for 1, 2, 13, and 14 which shows a symmetric pattern especially if we note that the number of code words with the weight w is the same as the number of code words with the weight $15-w$:

$$N_0 = N_{15} = 1$$

$$N_3 = N_{12} = 35$$

$$N_4 = N_{11} = 105$$

$$N_5 = N_{10} = 168$$

$$N_6 = N_9 = 280$$

$$N_7 = N_8 = 435$$

The following is a short program written in C.

```

#include <stdio.h>

main()
{
    int a1=0,a2=0,a3=0,a4=0,a5=0,a6=0,a7=0,a8=0,a9=0,a10=0,a11=0,x,y,u[16],d;
    char yn,enter,yyy='y';

    for (d=0; d<2048; d++) {
        u[0]=0; /*reset matrix values*/
        u[1]=a1;
        u[2]=a2;
        u[3]=a3;
        u[4]=a4;
        u[5]=a5;
        u[6]=a6;
        u[7]=a7;
        u[8]=a8;
        u[9]=a9;
        u[10]=a10;
        u[11]=a11;
        u[12]=0;
        u[13]=0;
        u[14]=0;
        u[15]=0;
        /*operations for Hamming encoding of the word (u)*/
        /*with the Generator matrix (v=uG)*/
        u[12]=u[5]+u[6]+u[7]+u[8]+u[9]+u[10]+u[11];
        u[13]=u[1]+u[2]+u[3]+u[8]+u[9]+u[10]+u[11];
        u[14]=u[2]+u[3]+u[4]+u[6]+u[7]+u[10]+u[11];
        u[15]=u[1]+u[3]+u[4]+u[5]+u[7]+u[9]+u[11];

        for (x=12; x<16; ++x)
            { /*if value =3 then in binary it should be =1*/
                if (u[x]>1) /*if value =2 then in binary it should be =0*/
                {
                    if (u[x]==3)
                        u[x]=1;
                    else
                        u[x]=0;
                }
            }
        printf("\nCODEWORD (b1-b11 b12-b15) : %i%i%i%i%i%i%i%i%i%i%i %i%i%i%i",
            u[1],u[2],u[3],u[4],u[5],u[6],u[7],u[8],u[9],u[10],u[11],u[12],u[13],u[14],
            u[15]);
        printf(" w = %i",u[1]+u[2]+u[3]+u[4]+u[5]+u[6]+u[7]+u[8]+u[9]+u[10]+u[11]+
            u[12]+u[13]+u[14]+u[15]);
        a1=a1+1;

        if (a1>1){
            a1=0;
        }
        if (a1==0){
            a2=a2+1;
        }
        if (a2>1){
            a2=0;
        }
        if (a2==0){
            a3=a3+1;
        }
        if (a3>1){
            a3=0;
        }
        if (a3==0){
            a4=a4+1;
        }
        if (a4>1){

```


Since the columns in the Hamming code are always distinct, the sum of two columns will never be zero so at least three columns are required. By adding zeros to the above columns we obtain example of three columns that add to zero for any Hamming code.

69. Suppose we take the (7,4) Hamming code and obtain an (8,4) code by adding an overall parity check bit.

a. Find the **H** matrix for this code.

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

b. What is the minimum distance?

$d_{\min} = 4$ since the least-weight codeword not has 4 bits.

c. Does the extra check bit increase the error correction capability? the error-detection capability?

The (8,4) extended Hamming code maintains its single error-correcting capability, and increases its error-detection capability to triple-error detecting.

70. A (7,3) linear code has check bits given by

$$\begin{aligned} b_4 &= b_1 + b_2 \\ b_5 &= b_1 + b_3 \\ b_6 &= b_2 + b_3 \\ b_7 &= b_1 + b_2 + b_3 \end{aligned}$$

a. Find the **H** matrix.

$$\mathbf{H} = \begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 \end{bmatrix}$$

b. Find the minimum distance.

$$d_{\min} = 4$$

c. Find the set of all codewords. Do you notice anything peculiar about the set of codewords?

000 0000000	100 1001101
001 0010111	101 1011010
010 0101011	110 1100110
011 0111100	111 1110001

Every nonzero codeword has a weight of 4.

71. An error-detecting code takes k information bits and generates a codeword with $2k + 1$ encoded bits as follows:

The first k bits consist of the information bits
 The next k bits repeat the information bits
 The next bit is the XOR of the first k bits.

a. Find the check matrix for this code.

$$c_{k+1}=c_1, c_{k+2}=c_2, \dots, c_{2k}=c_k, \text{ and } c_{2k+1}=c_1 + \dots + c_k$$

For $k = 4$, we have

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

b. What is the minimum distance of this code?

It can be seen that we need to add three columns to obtain the all zero column, so $d_{\min} = 3$.

c. Suppose the code is used on a channel that introduces independent random bit errors with probability 10^{-3} . Estimate the probability that the code fails to detect an erroneous transmission.

The error pattern with three errors has the largest probability, so we count how these undetectable triple errors occur: pick 1 column from the first k , then there is one pair from the second k that cause the sum to be zero. There are k ways of picking the column and exactly one corresponding pair, therefore:

$$P_{\text{detectionfailure}} = kp^3(1-p)^{2k+1-3} \approx kp^3.$$

72. A (6,3) linear code has check bits given by

$$\begin{aligned} b_4 &= b_1 + b_2 \\ b_5 &= b_1 + b_3 \\ b_6 &= b_2 + b_3 \end{aligned}$$

a. Find the check matrix for this code.

$$\mathbf{H} = \begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}$$

b. What is the minimum distance of this code?

$$d_{\min} = 3$$

c. Find the set of all codewords.

000000 100110 010101 001011 110011 101101 011110 111000

73. (Appendix 3A). Consider an asynchronous transmission system that transfers a sequence of N bits between a start bit and a stop bit. What is the maximum value of N if the receiver clock frequency is within 1 percent of the transmitter clock frequency?

Solution:

Let X = transmitter pulse duration and T = receiver pulse duration. Suppose $X = 0.99T$ (transmitter is slower), then we need:

$$1.5T + NT < (N + 2)X = (N + 2).99T$$

which implies $N < 48$.

Solutions to Chapter 4

1. A television transmission channel occupies a bandwidth of 6 MHz.

Solutions follow questions:

- a. How many two-way 30 kHz analog voice channels can be frequency-division multiplexed in a single television channel?

$$6 \times 10^6 / 30 \times 10^3 = 200 \text{ channels}$$

- b. How many two-way 200 kHz GSM channels can be frequency-division multiplexed in a single television channel?

$$6 \times 10^6 / 200 \times 10^3 = 30 \text{ channels}$$

- c. Discuss the tradeoffs involved in converting existing television channels to cellular telephony channels?

The biggest advantage of using existing television channels to provide cellular telephony channels is the very large bandwidth that they occupy. In theory, one could divide each television channel into many cellular telephone channels, as shown in parts (a) and (b). The frequency reuse aspect of cellular networks would multiply these channels many times. This would be at the cost of television channels that cover a broad region and presumably provide service to a large audience.

2. A cable sheath has an inner diameter of 2.5 cm.

Solutions follow questions:

- a. Estimate the number of wires that can be contained in the cable if the wire has a diameter of 5 mm.

Ignoring empty space between the wires:

$$\frac{\pi \left(\frac{25}{2}\right)^2}{\pi \left(\frac{5}{2}\right)^2} = 25 \text{ wires}$$

- b. Estimate the diameter of a cable that holds 2700 wire pairs.

Ignoring empty space between the wires:

$$\frac{\pi \left(\frac{d}{2}\right)^2}{\pi \left(\frac{5}{2}\right)^2} = 2 \times 2700 \text{ wires} \rightarrow d \cong 368 \text{ mm}$$

3. Suppose that a frequency band W Hz wide is divided into M channels of equal bandwidth.

- a. What bit rate is achievable in each channel? Assume all channels have the same SNR.

Each user uses W/M bandwidth. Using Shannon's Channel Capacity formula:

$$\text{Bit rate} = \left(\frac{W}{M}\right) \log_2(1 + SNR) \text{ bps}$$

- b. What bit rate is available to each of M users if the entire frequency band is used as a single channel and TDM is applied?

In this case, the total bit rate afforded by the W Hz is divided equally among all users:

$$\text{Bit rate} = \frac{W \log_2(1 + SNR)}{M} \text{ bps}$$

- c. How does the comparison of (a) and (b) change if we suppose that FDM requires a guard band between adjacent channels? Assume the guard band is 10% of the channel bandwidth.

Because of the guard band we expect that the scheme in (b) will be better since the bit rate in (a) will be reduced. In (a), the bandwidth usable by each channel is $0.9W/M$. Thus, we have:

$$\text{Bit rate} = \left(0.9 \frac{W}{M}\right) \log_2(1 + SNR) \text{ bps}$$

4. In a cable television system (see Section 3.8.2), the frequency band from 5 MHz to 42 MHz is allocated to upstream signals from the user to the network, and the band from 550 MHz to 750 MHz is allocated for downstream signals from the network to the users.

Solutions follow questions:

- a. How many 2 MHz upstream channels can the system provide? What bit rate can each channel support if a 16-point QAM constellation modem is used?

The system can provide:

$$\frac{(42 - 5) \text{ MHz}}{2 \frac{\text{MHz}}{\text{Channel}}} = 18 \text{ upstream channels.}$$

For bandpass channels, we have:

$$R = (2 \times 10^6) \frac{\text{pulses}}{\text{second}} * 4 \frac{\text{bits}}{\text{pulse}} = 8 \text{ Mbps}$$

- b. How many 6 MHz downstream channels can the system provide? What bit rates can each channel support if there is an option of 64-point or 256-point QAM modems?

Similarly, the system can provide:

$$\frac{(750 - 550)MHz}{6 \frac{MHz}{Channel}} = 33 \text{ upstream channels.}$$

$$R_{64\text{-point}} = (6 \times 10^6) \frac{\text{pulses}}{\text{second}} * 6 \frac{\text{bits}}{\text{pulse}} = 36 \text{ Mbps}$$

$$R_{256\text{-point}} = (6 \times 10^6) \frac{\text{pulses}}{\text{second}} * 8 \frac{\text{bits}}{\text{pulse}} = 48 \text{ Mbps}$$

5. Suppose a radio transmission system has a large band of available bandwidth, say 1 GHz, which is to be used by a central office to transmit and receive from a large number of users. Compare the following two approaches to organizing the system:

Solutions follow questions:

a. A single TDM system.

The channel will provide a maximum bit rate, say b bps. Each user will use b/n bps where n is the number of users. The number of users will be limited by the SNR on the channel, the time guard band required so collisions do not occur, and the rate at which users can switch their transmitters on and off.

b. A hybrid TDM/FDM system in which the frequency band is divided into multiple channels and time division multiplexing is used within each channel.

Although in theory, the total bit rate provided by the system is the same, the hybrid solution may be practically better because, if we assume again n users, each user timeslot is at a lower bit rate and for a longer duration. Therefore, TDM synchronization and switching speed problems are less significant than in the pure TDM version. By leveraging both the time and frequency domain, the technological requirements posed by either one are lessened.

6. Suppose an organization leases a T-1 line between two sites. Suppose that 32 kbps speech coding is used instead of PCM. Explain how the T-1 line can be used to carry twice the number of calls.

Solution:

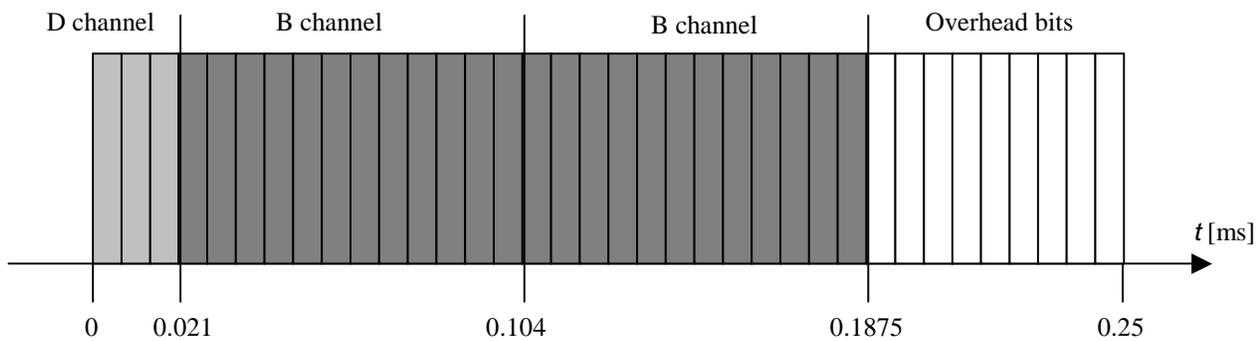
If a 32 kbps speech coding (such as Adaptive Differential Pulse Code Modulation (ADPCM)) is used, the bit rate required for each user is reduced to half of that required for PCM (64 kbps). The format in which a T-1 system carries information needs to be carried. For example a frame could now carry 48 4-bit samples.

Aside: Chapter 12 explains how advanced speech coding schemes reduce the bit rate while maintaining speech quality.

7. A basic rate ISDN transmission system uses TDM. Frames are transmitted at a rate of 4000 frames/ second. Sketch a possible frame structure. Recall that basic rate ISDN provides two 64 kbps channels and one 16 kbps channel. Assume that one-fourth of the frame consists of overhead bits.

Solution:

Assuming that the 16 kbps “D” channel is followed by the two 64 kbps “B” channels and then the overhead (which is not necessarily the case), in the time domain, the frame could look as follows:



8. The T-1 carrier system uses a framing bit to identify the beginning of each frame. This is done by alternating the value of the framing bit at each frame, assuming that no other bits can sustain an alternating pattern indefinitely. Framing is done by examining each of 193 possible bit positions successively until an alternating pattern of sufficient duration is detected. Assume that each information bit takes a value of 0 or 1 independently and with equal probability.

Solutions follow questions:

- a. Consider an information bit position in the frame. Calculate the average number of times this bit position needs to be observed before the alternating pattern is found to be violated.

Define the random variable X as the number of times the bit position needs to be observed before the alternating pattern is violated. Assume the information bits are equally likely to be 0 or 1. Without loss of generality assume first observed bit is zero.

$P[X \leq 1] = 0$	since at least two observations are required
$P[X = 2] = 1/2$	observe 0
$P[X = 3] = 1/4$	observe 11
$P[X = 4] = (1/2)^3$	observe 10

$$P[X = n] = (1/2)^{(n-1)}$$

$$\begin{aligned}
 E[X] &= \sum_{i=2}^{\infty} iP[X = i] = \sum_{i=2}^{\infty} ip^{i-1} = \sum_{i=2}^{\infty} \frac{d}{dp} p^i = \frac{d}{dp} \sum_{i=2}^{\infty} p^i = \frac{d}{dp} [\sum_{i=0}^{\infty} p^i - 1 - p] \\
 &= \frac{d}{dp} [\frac{1}{1-p} - 1 - p] = \frac{1}{(1-p)^2} - 1 = \frac{1}{(1-1/2)^2} - 1 \\
 &= 3
 \end{aligned}$$

Note that two frame times elapse in the time required to make three observations.

- b. Now suppose that the frame synchronizer begins at a random bit position in the frame. Suppose the synchronizer observes the given bit position until it observes a violation of the alternating pattern. Calculate the average number of bits that elapse until the frame synchronizer locks onto the framing bit.

On average, if the frame synchronizer starts at a random location, it will have to examine $193/2 = 96.5$ incorrect bit locations before locking onto the framing bit. For each such incorrect bit

location, on average, a non-alternating pattern will be discovered on the third observation, at which time the synchronizer will move one bit forward to the next bit.

Thus, the average number of bits elapsed until the synchronizer observes the framing bit is:

$$\begin{aligned} \text{Average number of bits elapsed until frame bit found} \\ = 96.5 \times [(2 \times 193) + 1] = 37345 \text{ bits} \end{aligned}$$

At this point, the synchronizer will begin observing the framing bit. It can't be sure that it has located the framing bit until an alternating pattern has been observed for a sufficient number of bits, say n . The synchronizer can never be sure that it has locked onto the frame bit, so n must be chosen sufficiently large so that the synchronizer is reasonably confident that it has indeed found the frame bit.

9. The CEPT-1 carrier system uses a framing *byte* at the beginning of a frame.

- a. Suppose that all frames begin with the same byte pattern. What is the probability that this pattern occurs elsewhere in the frame? Assume that each information bit takes a value of 0 or 1 independently and with equal probability.

The probability that random information matches the framing pattern is quite small:

$$p = (1/2)^8 = 1/256$$

- b. Consider an arbitrary information bit position in the frame. Calculate the average number of times that the byte beginning in this bit position needs to be observed before it is found to not be the framing byte.

The length N of a CEPT frame is 32 bytes, so $N=256$. The average number of frames M observed before a location mismatches the framing pattern is:

$$\begin{aligned} M &= 0(1-p) + 1p(1-p) + 2p^2(1-p) + \dots = p(1-p)\{1 + 2p + 3p^2 + \dots\} \\ &= p(1-p) \frac{1}{(1-p)^2} = \frac{p}{1-p} = \frac{(1/2)^8}{1-(1/2)^8} \end{aligned}$$

The first observation usually indicates that this is not the right position.

- c. Now suppose that the frame synchronizer begins at a random bit position in the frame. Suppose the synchronizer observes the byte beginning in the given bit position until it observes a violation of the alternating pattern. Calculate the average number of bits that elapse until the frame synchronizer locks onto the framing byte.

The average number of bits observed until the beginning of frame is found is then:

$$\frac{N}{2}(MN + 1) = \frac{2^8}{2} \left(\frac{1/2^8}{1-1/2^8} N + 1 \right) = \frac{2^8}{2} \left(\frac{2-1/2^8}{1-1/2^8} \right) \approx 256 \text{ bits}$$

10. Suppose a multiplexer has two input streams, each at a nominal rate of 1 Mbps. To accommodate deviations from the nominal rate, the multiplexer transmits at a rate of 2.2 Mbps as follows. Each group of 22 bits in the output of the multiplexer contains 18 positions that always carry information bits, nine from each input. The remaining four positions consist of two flag bits and two data bits. Each flag bit indicates whether the corresponding data bit carries user information or a stuff bit because user information was not available at the input.

- a. Suppose that the two input lines operate at exactly 1 Mbps. How frequently are the stuff bits used?

In this case, the stuff bits are always used because the information bits alone only provide an aggregate bit rate of 1.8 Mbps.

- b. How much does this multiplexer allow the input lines to deviate from their nominal rate?

This multiplexer provides either 9 or 10 bits for each stream per 22-bit frame. Thus, it allows either of the two input streams to transmit as low as 0.9 Mbps and as high as 1.0 Mbps.

11. Calculate the number of voice channels that can be carried by an STS-1, STS-3, STS-12, STS-48, and STS-192. Calculate the number of MPEG2 video channels that can be carried by these systems.

Solution:

A conventional voice channel is 64 kbps. An MPEG2 video signal can vary in bit rate. Assuming a bit rate of 5 Mbps, the results are summarized in the table below:

	Bit rate of signal (Mbps)	Voice Channels	MPEG2 Video
STS-1	51.84	810	10
STS-3	155.52	2,430	31
STS-12	622.08	9,720	124
STS-48	2488.32	38,880	498
STS-192	9953.28	155,520	1,991

12. SONET allows positive or negative byte stuffing to take place at most once every four frames. Calculate the minimum and maximum rates of the payload that can be carried within an STS-1 SPE.

Solution:

STS-1 rate = 51.84 Mbps

Payload rate = 50.112 Mbps

Frame = 125 μ sec

Drifts = (8 bits) / (4 x 125 μ sec) = 0.016 Mbps

Maximum payload rate = 50.112 + 0.016 = 50.128 Mbps

Minimum payload rate = 50.112 – 0.016 = 50.096 Mbps

13. Consider a SONET ring with four stations. Suppose that tributaries are established between each pair of stations to produce a logical topology. Find the capacity required in each hop of the SONET ring in the following three cases, assuming first that the ring is unidirectional and then that the ring is bidirectional.

Note: The solution to this problem is currently being reviewed and may be revised.

Solutions follow questions:

- a. The traffic between each pair of stations is one STS-1.

Unidirectional: STS-6
 Bidirectional: STS-3

- b. Each station produces three STS-1s worth of traffic to the next station in the ring and no traffic to other stations.

Unidirectional: STS-3
 Bidirectional: STS-3

- c. Each station produces three STS-1's worth of traffic to the furthest station along the ring and no traffic to other stations.

Unidirectional: STS-6
 Bidirectional: STS-6

14. Consider a set of 16 sites organized into a two-tier hierarchy of rings. At the lower tier a bidirectional SONET ring connects four sites. At the higher tier a bidirectional SONET ring connects the four lower-level SONET rings. Assume that each site generates traffic that requires an STS-3.

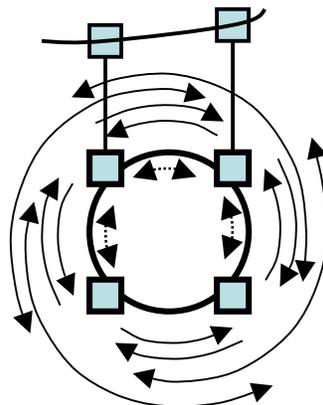
The high-tier ring must carry all inter-ring traffic from the lower tiers. The capacity required on the high-tier ring is thus dependent on the volume of inter-ring traffic generated from the lower rings and the distribution of such traffic. For this question, let us assume that the traffic generated is distributed evenly throughout the high-tier network – that is, each low-tier ring sends an equal amount of traffic to each of the other low-tier rings.

The lower-tier ring must accommodate 100% of the traffic sent from each of its sites. Since each site will receive as much traffic as it will generate, each site will receive STS-3 of bandwidth.

Solutions follow questions:

- a. Discuss the bandwidth requirements that are possible if 80% of the traffic generated by each site is destined to other sites in the same tier ring.

If 80% of the traffic is to stay within the same ring, then the required bandwidth on the higher tier ring will be reduced as most traffic is confined locally to the lower tier rings.



In the lower tier rings the number of links to be traversed by the average message is between one and two. Assuming that the traffic is distributed evenly on all of the links in the lower tier ring, the bandwidth required on each link will be four times the average traffic between each two links,

which is $4 \times \frac{1}{3} \times 80\%$ STS-3 or $\frac{16}{15}$ STS-3. Adding the other 20% of the traffic which is destined to the higher tier link and has to traverse some of the lower tier links in both directions, the average link capacity in the lower tier links should be $\frac{16}{15}$ STS-3 + $2 \times 20\%$ STS-3 or $\frac{22}{15}$ STS-3, which is almost 1.5 STS-3 on average.

On the higher tier ring, the traffic from each subring will be routed equally to each other subring. Thus the bandwidth required on each link will be four times the average traffic between two subrings or $4 \times \frac{1}{3}$ of the total traffic from each subring. Because each ring produces $4 \times 0.2 \times$ STS-3 amount of traffic, the bandwidth required on each link of the higher tier ring is on the order of $\frac{4}{3} \times 0.8 \times$ STS-3 \approx STS-3.

- b. Discuss the bandwidth requirements that are possible if 80% of the traffic generated by each site is destined to sites in other rings.

If 20% of the traffic is to stay within the same ring, then the required bandwidth on the higher tier ring will be increased. The required bandwidth on the lower tier rings will not change significantly, as each source still produces and receives STS-3 of traffic.

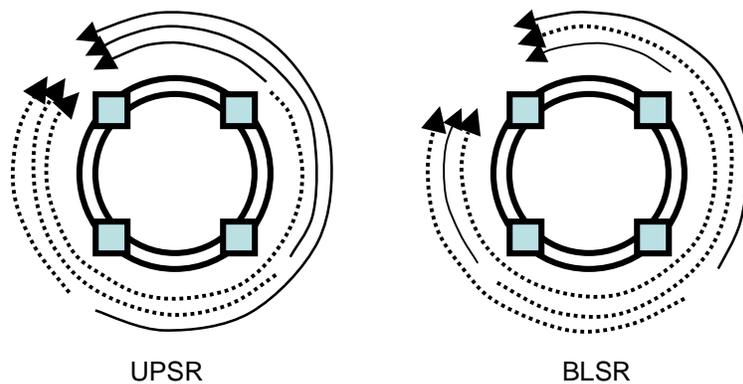
The higher tier ring will now support 4 times more traffic. Thus, the bandwidth requirement is around STS-12.

15. Compare the efficiency of BLSR and UPSR rings in the following two cases:

- a. All traffic originating at the nodes in the ring are destined for a given central node.

Considering that the total capacity of working and protection lines in both BLSR and UPSR is the same, the total bandwidth utilized in both BLSR and UPSR in this scenario is the same. The difference is the distribution of the traffic and unused capacity across the ring in the two cases.

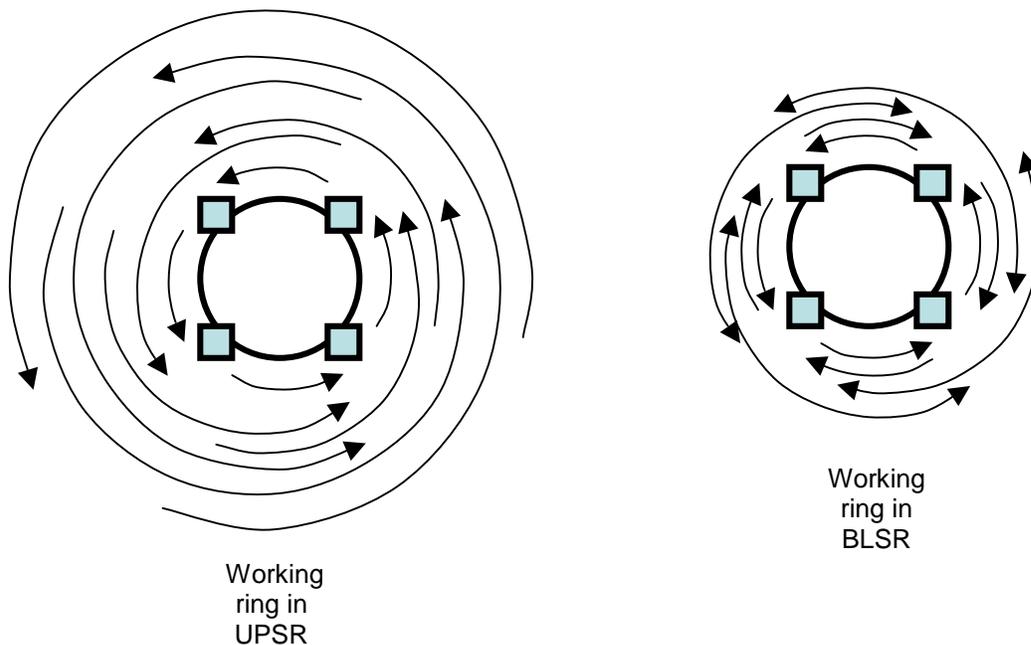
As an example we consider the ring shown in the figure below, consisting of four nodes. We first consider the UPSR scenario. The last hop before the central node is the bottleneck in this case limiting the traffic in the other hops on the working ring shown with solid lines. Assuming equal distribution of bandwidth among the nodes, if the last hop is fully utilized the hop prior to the last hop will be utilized up to $\frac{2}{3}$ of its capacity and the one before that will be utilized up to $\frac{1}{3}$ of its capacity. The remaining hop will not be utilized at all. The lines on the protection path will be utilized similarly but in the opposite order, shown with dotted lines. As a result the sum of working and protection lines together will be utilized only to half of the total capacity all across the ring.



Now we consider the BLSR scenario. In this case the working traffic can be sent in both directions but the protection traffic will be sent in the opposite direction and as a result as seen in the figure the total of working and protection traffic all across the ring will be the same as the other case as well as the unused capacity of the links.

- b. Each node originates an equal amount of traffic to all other nodes.

In both BLSR and UPSR lines can be shared by traffic from other nodes. The more nodes share the same line the less the line can be utilized by the node attached to the line. In the given scenario each line is shared by the traffic from twice the number of nodes in UPSR compared to the BLSR. As a result BLSR is more efficient in this case.



16. Consider the operation of the dual gateways for interconnecting two bidirectional SONET rings shown in Figure 4.29. The primary gateway transmits the desired signal to the other ring and simultaneously transmits the signal to the secondary gateway which also routes the signal across the ring and then to the primary gateway. A service selector switch at the primary gateway selects between the primary and secondary signals. Explain how this setup recovers from failures in the link between the primary gateways.

Solution:

Two logical connections are maintained between the primary gateways. The first is a direct physical connection via the primary gateway link. The second is connected indirect via the secondary gateway as described in the question.

This redundancy provides a backup in case of a link failure. If the primary gateway link fails, the service selector switch can select signals for the alternate indirect path until the primary gateway link is restored.

17. Consider the synchronous multiplexing in Figure 4.15. Explain how the pointers in the outgoing STS-1 signals are determined.

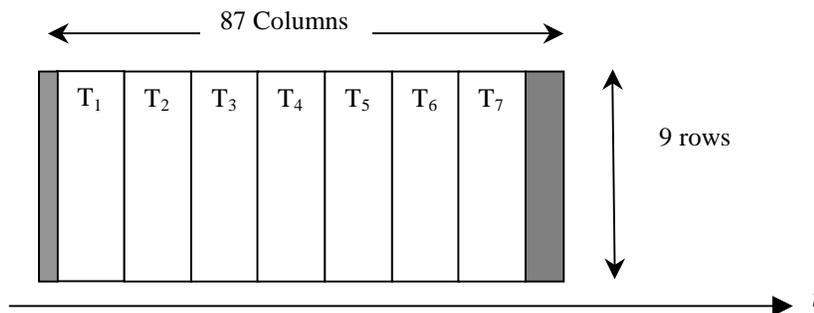
Solution:

The section and line overhead of each incoming STS-1 signal is terminated and its payload is mapped into a new STS-1 frame that is synchronous to a local clock. The outgoing pointers are all adjusted according to a common clock such that all the outgoing STS-1 frames are synchronized with respect to each other.

18. Draw a sketch to explain the relationship between a virtual tributary and the synchronous payload envelope. Show how 28 T-1 signals can be carried in an STS-1.

Solution:

Seven tributaries are multiplexed in stream. Each tributary carries four T-1 channels and occupies twelve columns of the SPE, as shown below:



The first column is path overhead. The next 84 columns are the seven tributaries and the last two columns are unused.

19. Do a web search for DWDM equipment available from telecommunications equipment vendors. What specifications do you see in terms of number of channels, spacing between channels, bit rates and formats of the signals that can be carried?

Solution:

Some specs for WDM and DWDM multiplexers:

Wavelength: 632- 820 nm, 980-1064 nm, 980-1480 nm, 980-1550 nm, 1064-1310 nm, 1310-1550 nm, 1310-1625 nm, 1480-1550 nm, 1550-1625 nm, etc.

Channel Spacing: 50 GHz (0.1 nm spacing between channels), 100 GHz (0.2 nm spacing between channels), 200 GHz (0.4 nm spacing between channels), etc.

Per channel bit rates: 100 Mbps, 1.2 Gbps, 2.4 Gbps, 10 Gbps, etc.

Number of Channels: 4, 8, 16, 32, 80, 160

20. Consider WDM systems with 100, 200, and 400 wavelengths operating at the 1550 nm region and each carrying an STS-48 signal.

Solutions follow questions:

- a. How close do these systems come to using the available bandwidth in the 1550 nm range?

One STS-48 signal carries 2488.32 Mbps of traffic. Thus, the aggregate bit rate of each of these systems is:

System	Bit rate
100 wavelengths	248.832 Gbps
200 wavelengths	497.664 Gbps
400 wavelengths	995.328 Gbps

The available bandwidth of this window is 12.49 THz. Thus, none of these systems even comes close to using all of the available bandwidth.

- b. How many telephone calls can be carried by each of these systems? How many MPEG2 television signals?

The number of telephone calls (64 Kbps) and MPEG2 video signals (5 Mbps) that can be carried each system is shown below:

System	Telephone Calls	MPEG2 Video Signals
100 wavelengths	3.888×10^6	49 766
200 wavelengths	7.776×10^6	99 532
400 wavelengths	15.552×10^6	199 065

Although, none of the systems comes close to using all of the available bandwidth, they can still carry huge amounts of traffic. This shows how WDM provides the potential for huge amounts of bandwidth and radically expands the bandwidth of existing optical links.

21. Consider the SONET fault protection schemes described earlier in the chapter. Explain whether these schemes can be used with WDM rings.

Solution:

The fault protection schemes described in problem 15 are applicable to any ring network structure. To use these schemes in WDM, the exact same procedures can be carried out except that in the case of WDM, instead of switching streams, wavelengths are switched. Furthermore, the switching can be done in optical domain by all optical protection switches that detect the signals and switch to protection signal instantaneously. Optical protection switching can be faster and more efficient.

22. Calculate the spacing between the WDM component signals in Figure 4.8. What is the spacing in hertz and how does it compare to the bandwidth of each component signal?

Solution:

In Figure 4.8 there are 16 wavelengths distributed equally over the spectrum from 1545 nm to 1557 nm. Thus the spacing between each wavelength is:

$$\frac{1557 - 1545}{16} = 0.75 \text{ nm}$$

The spacing in Hertz is:

$$\text{Frequency Bandwidth} = \frac{(3 * 10^8)(0.75 * 10^{-9})}{(1550 * 10^{-9})^2} = 93.65 \text{GHz}$$

Thus, the maximum achievable bit rate that each wavelength can achieve using on-off modulation is $2 \times 93.65 = 187$ Gbps. Each signal carries 2.5 Gbps of traffic, so we see that because of the relatively large guard bands needed, WDM wavelengths have a very large spacing compared to their width.

23. How does WDM technology affect the hierarchical SONET ring topology in Figure 4.29? In other words, what are the consequences of a single fiber providing large number of high-bandwidth channels?

Solution:

WDM allows for much more aggregate bandwidth on each fiber. Thus, in the context of the hierarchical SONET ring topology, the high-tier ring can use WDM in order to increase its bandwidth and the number of channels it can support and, in turn, the number of lower-tier rings that it can service. WDM allows it to achieve this increased service ability without the need to lay down more physical fibers over its links. On the other hand, the use SONET ring topology mean that each wavelength is associated with a separate SONET ADM, and this creates a problem in terms of interconnecting traffic across rings.

24. WDM and SONET can be used to create various logical topologies over a given physical topology. Discuss how WDM and SONET differ and explain what impact these differences have in the way logical topologies can be defined.

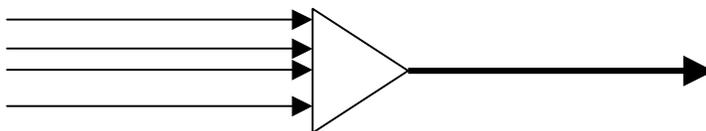
Solution:

Both WDM and SONET define logical topologies by setting up semi-permanent paths between nodes in the network. WDM creates these paths by assigning lightpaths between nodes. WDM without wavelength conversion has an inherent restriction imposed on the logical topologies that it can create, because no two lightpaths that share the same link can use the same wavelength. In SONET, since tributaries are electrically defined, the only restriction imposed on the topologies is the total bandwidth utilized on each link.

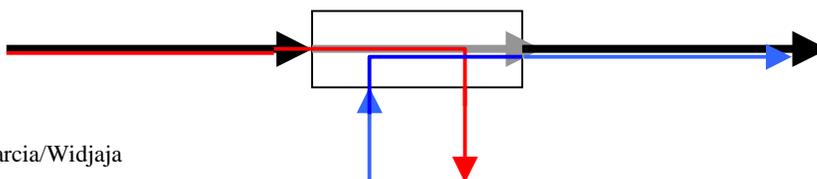
25. Compare the operation of a multiplexer, an add-drop multiplexer, a switch, and a digital cross-connect.

Solution:

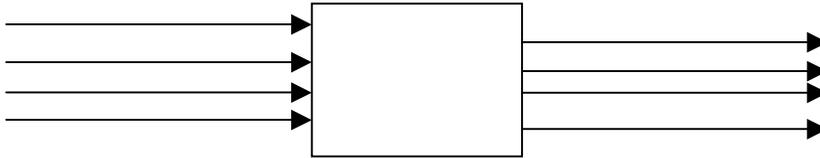
A multiplexer is a 1:N device. It takes N separate signals on N different inputs and combines them into one higher rate signal on one output port.



An add-drop multiplexer takes in N signals on one input port and replaces one of them with a new signal from a separate input port. The new aggregate signal is routed to an output port and the signal that was replaced is "dropped" to a separate output port.



A switch takes in N inputs and routes them to N different outputs. Frequently each input to a switch contains a number of multiplexed connections. These connections can be demultiplexed and routed to different output ports where they are aggregated prior to exiting the switch. Generally, switches are configured using signaling that establishes paths across the network.



A digital cross-connect is similar to a switch except that it is semi-permanent, usually configured by network operators rather than signaling processes. Their configuration is also done on a larger time scale (days or weeks). Digital cross-connects provide a basic network topology on which routing can be applied.

26. Consider a crossbar switch with n inputs and k outputs.

Solutions follow questions:

- a. Explain why the switch is called a concentrator when $n > k$. Under what traffic conditions is this switch appropriate?

This switch is a concentrator because traffic comes in on n lines and is concentrated onto k lines. Traffic on the output lines is higher than that of the input lines by a factor of n/k on average. The switch has inherent multiplexing functionality.

- b. Explain why the switch is called an expander when $n < k$. Under what traffic conditions is this switch appropriate?

This switch is an expander because the number of lines is expanded from n to k . This switch is appropriate when the outgoing lines carry less traffic than the incoming lines. One such example is at the egress of a backbone network to an access network. The switch has inherent demultiplexing functionality.

- c. Suppose an $N \times N$ switch consists of three stages: an $N \times k$ concentration stage; a $k \times k$ crossbar stage; and a $k \times N$ expansion stage. Under what conditions is this arrangement appropriate?

The number of crosspoints in an $N \times N$ crossbar switch is N^2 . If the traffic load on the inputs is relatively low, this configuration may save on hardware costs. The number of crosspoints needed in the three-stage switch described is:

$$k^2 + 2kN$$

Depending on the values of N and k , $k^2 + 2kN$ may be less than N^2 , especially if N is much greater than k .

- d. When does the three-stage switch in part (c) fail to provide a connection between an idle input and an idle output line?

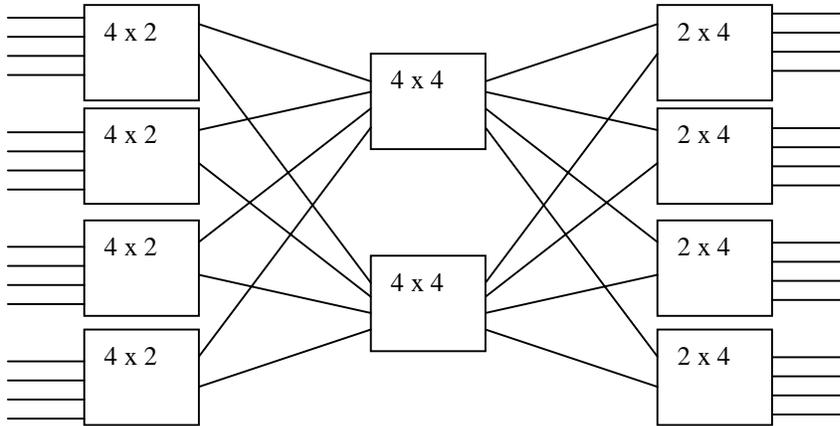
Assuming unicast traffic, if more than k inputs require a path to one of the third stage's n outputs at the same time, blocking will occur in the first stage, since it only has k outputs.

27. Consider the multistage switch in Figure 4.35 with $N = 16$, $n = 4$, $k = 2$.

Solutions follow questions:

- a. What is the maximum number of connections that can be supported at any given time? Repeat for $k = 4$ and $k = 10$.

For $N = 16$, $n = 4$ and $k = 2$, we have the following switch architecture:



Thus, the second stage is the bottleneck, and blocking can occur in the first stage. Thus, eight connections can be supported at a time.

If $k = 4$, then blocking will occur if we are not allowed to rearrange connections. It can be shown that in this case blocking can be avoided if we are allowed to rearrange the connection pattern every time a new connection request is made.

If $k = 10$, then there are ten 4 x 4 switches in the second stage. Since there are only 16 inputs and 16 outputs, the switch can accommodate any set of connections without blocking.

- b. For a given set of input-output pairs, is there more than one way to arrange the connections over the multistage switch?

As shown in the picture in part (a), it is clear that each input-output pair can be connected through any one of the k second-stage switches. Thus, there are k ways to arrange the connections over a multi-stage switch.

28. In the multistage switch in Figure 4.35, an input line is busy 10% of the time.

Solutions follow questions:

- a. Estimate the percent of time p that a line between the first and second stage is busy.

Although, the busyness of each output line of the first stage depends on the switch control, on average, we can estimate that each of the output lines will have n/k the amount of traffic as the input lines. Thus

$$p = (n/k) 10\%$$

- b. How is p affected by n and k ?

p is directly proportional to the ratio of n to k .

- c. How does this p affect the blocking performance of the intermediate crossbar switch?

Blocking occurs at the second switches when two of its inputs are to be routed to the same output. Although the control and routing algorithm of the switch can minimize the chance of this occurring, the greater the value of p , the greater the chance that such blocking will occur.

- d. Supposing that the blocking probability of the intermediate crossbar is small, what is the proportion of time p' that a line between the second and third stage is busy?

Assuming that the blocking probability of the second stage switches is small, the percentage of time that the output lines on a switch in the second stage are busy is equal to the percentage of time that its input lines are busy. This is the case because the second stage switches are $(N/n) \times (N/n)$ components. Thus,

$$p' = p = (n/k) 10\%$$

- e. For a given input and output line, what is the probability that none of the N/n paths between the input and output lines are available?

In this question N/n should be k , since there are k possible paths from each input to each output (one different path possible through each second-stage switch). Assuming independence of all lines, which is not the case generally, for any one of the k paths between a given input and the output to be unavailable, one of the 2 jumps must be unavailable.

$$P[\text{path not available}] = 1 - P[\text{path is available}] = 1 - (1-p')(1-p') \approx 2p$$

For none of the paths to be available, all k must be busy. Since we assumed that all are independent of each other,

$$P[\text{none of the paths available}] = (2p)^k$$

29. Consider the multistage switch in Figure 4.35 with $N = 32$. Compare the number of crosspoints required by a nonblocking switch with $n = 16$, $n = 8$, $n = 4$, and $n = 2$.

Solution:

For any switch to be non-blocking, we require $k_{nb} = 2n - 1$. The total number of crosspoints is $2Nk + k(N/n)^2$. The resulting number of crosspoints necessary for different values of n is shown below.

N	n	k	Number of Crosspoints needed
32	16	31	2108
32	8	15	1200
32	4	7	896
32	2	3	960

For a one-stage $N \times N$ switch with $n = 32$, we would require 1032 crosspoints. Thus we see that, just as was noted in question 26c, for k much less than n , multistage switches can provide good hardware economy while remaining non-blocking.

30. A multicast connection involves transmitting information from one source user to several destination users.

Solutions follow questions:

- a. Explain how a multicast connection may be implemented in a crossbar switch.

In a crossbar switch, multicasting can be implemented by copying the multicast input to multiple outputs.

- b. How does the presence of multicast connections affect blocking performance of a crossbar switch? Are unicast calls adversely affected?

Unlike unicast switches, flow in multicast switches is not conserved. Generally, if multicasting is implemented in a switch, the total traffic flow out of the switch is more than the flow into the switch. Thus the probability of blocking is increased because the traffic inside the switch is higher.

Unicast traffic is also affected by this increased flow because fewer resources in the switch are available.

- c. Explain how multicast connections should be implemented in a multistage switch. Should the multicast branching be done as soon as possible or as late as possible?

In a multistage switch with a given multicast traffic profile, each branching procedure represents an increase in traffic flow beyond the branch point. Thus, to keep traffic load in the switch to a minimum, branching should be done as late as possible.

- 31.** Do a web search for crosspoint switching chips. What specifications do you find for number of ports, bit rate/port? Are dualcasting and multicasting supported? Do the crosspoints need to be changed together or can they be changed independently? What considerations determine how fast the crosspoints can be changed?

Solution:

Some specs for Crosspoint fabrics chips:

Switch size: 2x2, 4x4, 17x17, 34x34, 36x37, 68x69, 64x64, 72x72, 144x144

Port speed: 1.6 Gbps, 2.5 Gbps, 3.2 Gbps, 3.6 Gbps

Dualcasting and multicasting is usually supported by crosspoint chips, as well as masking particular ports, and also programming a single port or a subset of ports independently.

The speed of changes depend on the clock recovery and signal synchronization at the port interfaces (receivers and transmitters at input and output ports of the crosspoint fabric), and also the speed of programming the I/O Table registers.

- 32.** What is the delay incurred in traversing a TSI switch?

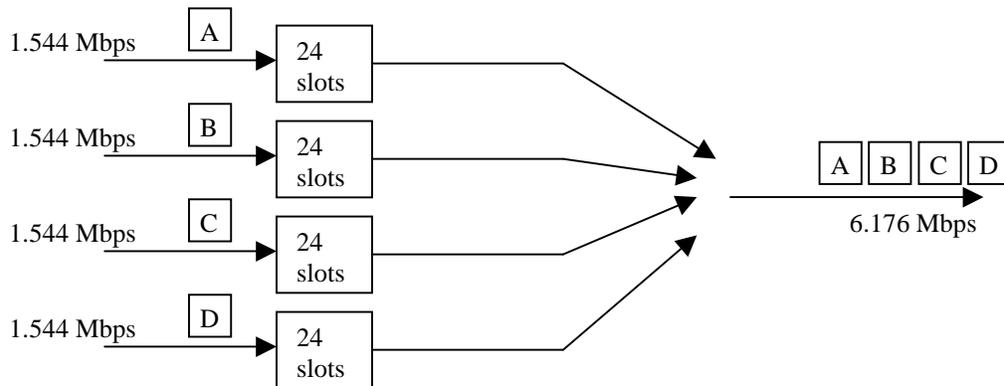
Solution:

The arriving frame is written onto the switch register and after the frame is completely written in, it needs to be read out in permuted order.

- 33.** Explain how the TSI method can be used to build a time-division multiplexer that takes four T-1 lines and combines them into a single time-division multiplexed signal. Be specific in terms of the number of registers required and the speeds of the lines involved.

Solution:

Each input line carries T-1 traffic. The input frame size is $n = 24$ slots. Each frame that is output is of length $4n = 96$ slots. After a full set of 96 slots is written into the 4 registers, the slots are read out in the outgoing line.



- 34.** Suppose that the TDM frame structure is changed so that each frame carries two PCM samples. Does this affect the maximum number of channels that can be supported using TSI switching?

Note: The solution to this problem is currently being reviewed and may be revised.

Solution:

Yes. Now each slot requires two memory reads and two memory writes. Thus, the maximum number of channels that can be supported is:

$$\text{Max} = \frac{125\mu\text{s}}{4 \times \text{mem}}$$

In other words, one half as many slots per frame can be supported as compared to when each slot requires one PCM sample.

- 35.** Examine the time-space-time circuit-switch architecture and explain the elements that lead to greater compactness, that is, smaller physical size in the resulting switching system.

Solution:

In a space switch, n inputs are switched simultaneously in a switching time of T seconds. The space-time circuit switch switches each of the n inputs in succession with a switching time of T/n each, for a total switching time of T .

By switching the inputs one at a time instead of simultaneously, the switching matrix needed at any given time is reduced in size accordingly and greater hardware compactness is achieved. The implicit tradeoff, however, is that higher line and switch speeds are needed.

- 36.** Consider the three-stage switch in problem 4.26c. Explain why a space-time-space implementation of this switch makes sense. Identify the factors that limit the size of the switches that can be built using this approach.

Solution:

A space-time-space implementation would have the middle stage replaced by a TSI switch stage, and thus reduce hardware costs. The maximum number of slots in a TSI switch, say M , is limited

by the length of the memory cycle. This limitation means that $N/n \leq M$, which in turn places a limitation on N , the number of inputs to the switch (since n has a maximum that is determined by technological constraints). The limitation on the number of slots M suggests that the approach is suitable when concentration can be used in the input stage because of light loading in the input lines.

37. Consider n digital telephones interconnected by a unidirectional ring. Suppose that transmissions in the ring are organized into frames with slots that can hold one PCM sample.

Solutions follow questions:

- a. Suppose each telephone has designated slot numbers into which it inserts its PCM sample in the outgoing direction and from which it extracts its received PCM sample from the incoming direction. Explain how a TSI system can be used to provide the required connections in this system.

At some point in the ring, TSI is carried out to interchange the information in the time slots between all pairs of connected users

- b. Explain how the TSI system can be eliminated if the pairs of users are allowed to share a time slot?

If all user pairs share a time slot, for any connected pair of users we can assign the same time-slot to both. To communicate, each of the users inserts outgoing PCM samples on the slot and extracts incoming PCM samples from the same slot.

- c. How would connections be established in parts (a) and (b)?

In the system described in part (a) the TSI switch monitors all of the slots continuously for connection request made by the users. If user A requests a connection to user B, the central switch controller makes the connection if the slot corresponding to user B is available.

In the system described in part (b) a mechanism is needed to coordinate the sharing of slots between all the pairs of users. The most efficient mechanism would have half as many timeslots as users. The central switching mechanism would assign each pair of users a timeslot to communicate over. Call setup would require all users to communicate with the central switch on a setup slot that would have to be accessed via a MAC protocol.

An alternative simple distributed mechanism would involve each user monitoring their own time-slot. Assuming that the users have time slot tunability, a call could be initiated by sending a sample on the intended receivers time slot (provided it is free). Communication would occur on the receiver's time slot in this scenario. The drawback of this mechanism is that no more than half of the available bandwidth would ever be used.

38. Consider the application of a crossbar structure for switching optical signals.

- a. Which functions are the crosspoints required to implement?

The crosspoints must be able to open, close, hold open, and hold closed connections.

- b. Consider a 2×2 crossbar switch and suppose that the switch connection pattern is $(1 \rightarrow 1, 2 \rightarrow 2)$ for T seconds and $(1 \rightarrow 2, 2 \rightarrow 1)$ for T seconds. Suppose it takes τ seconds to change between connection patterns, so the incoming optical signals must have guardbands to allow for this. Calculate the relationship between the bit rate R of the information in the optical signals, the number of bits in each "frame", and the values T and τ ? For R in the range from 1

gigabit/second to 1 terabit/second and τ in the range of 1 microsecond to 1 millisecond, find values of T that yield 50% efficiency in the use of the transmission capacity.

The transition time is τ seconds, and the hold time is T seconds. The optical transmission rate is R bps and the actual information bit rate is $R_{\text{effective}}$, and letting n be the number of bits per frame (including overhead), we have:

$$R_{\text{effective}} = (\text{total bit rate of connection})(\text{proportion of useful bits})$$

$$= R * \left(\frac{T}{T + \tau} \right) = \frac{n}{T + \tau}$$

The value of T that yields 50% efficiency is $T = \tau$.

For $T = \tau = 1$ microsecond, we have $n = RT$ which gives frame size of $n = 1000$ bits at 1 Gbps and $n = 1000000$ bits at 1 Tbps. For $T = \tau = 1$ millisecond, the frame sizes are 1000 times longer.

39. Suppose an optical signal contains n wavelength-division multiplexed signals at wavelengths $\lambda_1, \lambda_2, \dots, \lambda_n$. Consider the multistage switch structure in Figure 4.35 and suppose that the first stage switches consist of an element that splits the incoming optical signal into n separate optical signals each at one of the incoming wavelengths. The j th such signal is routed to j th crossbar switch in the middle stage. Explain how the resulting switch structure can be used to provide a rearrangeable optical switch.

Solution:

In the resulting architecture, the second-stage inputs to any given switching element have the same wavelength. Thus, each single wavelength signal input to the second stage can be directed to any one of the third stage elements. Each third stage element corresponds to a single output port. It is assumed that the third stage consists of elements that recombine the n incoming signals, each with a different wavelength, into a WDM signal. The operation of this switch is very similar to that of the switch in figure 4.35. However, it differs in that there is only one possible route for each input-output-wavelength combination.

40. Consider the equipment involved in providing a call between two telephone sets.

Solutions follow questions:

- a. Sketch a diagram showing the various equipment and facilities between the originating telephone through a single telephone switch and on to the destination telephone. Suppose first that the local loop carries analog voice; then suppose it carries digital voice.

In an analog connection, the telephone set is connected by a pair of wires to the telephone office. The wires are terminated in a line card that digitizes the voice signal and transfers the digital signal to a digital switching system. The digital signal travels along a fixed path from the telephone switch in the originating office to the telephone switch in the destination office. The two or more offices are interconnected by digital transmission lines.

If the telephone set were to generate digital voice then a digital signal would be transmitted along the local loop that connects the telephone to the central office. The operation of the line card would differ, as the arriving signal would already be in digital form.

- b. Repeat part (a) in the case where the two telephone calls are in different LATAs.

If the two telephones are in different LATAs then the digital voice signal must traverse tandem switches that connect the two LATAs.

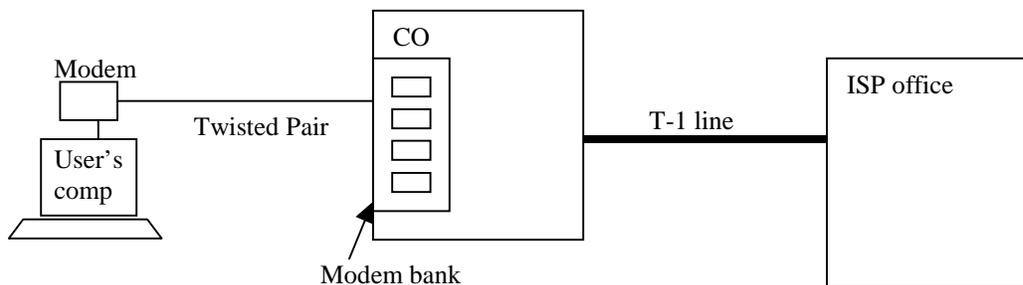
- c. In parts (a) and (b) identify the points at which a call request during setup can be blocked because resources are unavailable.

Call requests can be blocked when there are insufficient resources to complete the call. These resources can include time slots and crosspoints in the switches and time slots in the digital lines. Blocking can also result from unavailability of signaling processing equipment.

41. Suppose that an Internet service provider has a pool of modems located in a telephone office and that a T-1 digital leased line is used to connect to the ISP's office. Explain how the 56K modem (that was discussed in Chapter 3) can be used to provide 56 kbps transfer rate from the ISP to the user. Sketch a diagram showing the various equipment and facilities involved.

Solution:

The ISP office is connected to the user's local central office, which is in turn connected to the user. The ISP can use the T-1 line to send digital information at a rate of 56 kbps into and across the telephone network to the line card that connects to the user. The line card can then deliver information at 56 kbps to the user. Recall that the user cannot send at 56 kbps because of limitations on the available SNR.



42. Why does a conventional telephone still work when the electrical power is out?

Solution:

The telephone company supplies each of its telephone lines with power at the central office. This power is stored in the form of wet batteries that can alternately be charged by a backup battery in the event of a power failure at the central office. These huge batteries occupy entire floors in telephone offices.

43. In Figure 4.44b, how does the network know which inter-exchange carrier is to be used to route a long distance call?

Solution:

The inter-exchange carrier that handles the long-distance calls for a specific user is specified a priori by the user's service contract. The user "signs-up" with one long distance plan over another, and it is the responsibility of the local service provider to route the call over the appropriate inter-exchange carriers.

44. ADSL was designed to provide high-speed digital access using existing telephone facilities.

- a. Explain how ADSL is deployed in the local loop.

An ADSL modem is required at each end of the user's twisted pair connection. Since ADSL coexists with regular phone service, it uses higher frequencies for data transmission. Bandwidth-limiting elements, such as loading coils, must be removed. ADSL can only be offered within certain distances of the central office.

- b. What happens after the twisted pairs enter the telephone office?

At the telephone office, the voice and data are separated/filtered and the demultiplexed signals are routed accordingly (for example, voice may be routed over the phone network and data over the Internet).

- c. Can ADSL and ISDN services be provided together? Explain why or why not.

No, ADSL and ISDN use overlapping parts of the signal spectrum so they cannot operate simultaneously.

45. In this problem we compare the local loop topology of the telephone network with the coaxial cable topology of cable television networks (discussed in Chapter 3).

- a. Explain how telephone service may be provided using the cable television network.

Since the television network is inherently a broadcast, shared medium, a method such as the cable modem method described in chapter 3 would be needed to accommodate the traffic of predominantly point-to-point bidirectional telephone service. The mechanism used would include a MAC protocol of some form to provide a fair method of partitioning the total available bandwidth among callers.

- b. Explain how cable television service may be provided using the local loop.

It would be very difficult to provide cable television on the local loop network due to the huge bandwidth currently provided by the cable broadcast network. The only effective means of providing full cable service over the local loop would likely involve sending a single channel to each user. Changing channels would require an upstream request to the local switch, which would then replace the originally transmitted channel with the requested one. By moving the pedestal that connects to the home, it may be possible to provide sufficiently high speeds to offer several simultaneous channels to a home.

- c. Compare both topologies in terms of providing Internet access service.

Both these networks currently provide Internet access. Cable modems have the advantage of providing huge peak bandwidth via that large bandwidth (up to 30 Mbps) of the cable network. However, as has been discovered by many subscribers to this service, since the cable network is a shared broadcast medium, the bandwidth provided does not scale well to many users. During peak hours, the total bandwidth of the network is divided among all of the users connected to the final branching point of the network. This leads to noticeable and significant decrease in service.

Access to the Internet via the phone network using DSL modems has the advantage of providing a dedicated access line to its users. As a result, each user's service access rate is independent of the number of users. The tradeoff is the lower bandwidths available. Currently ADSL typically provides bandwidth of up to 1 Mbps.

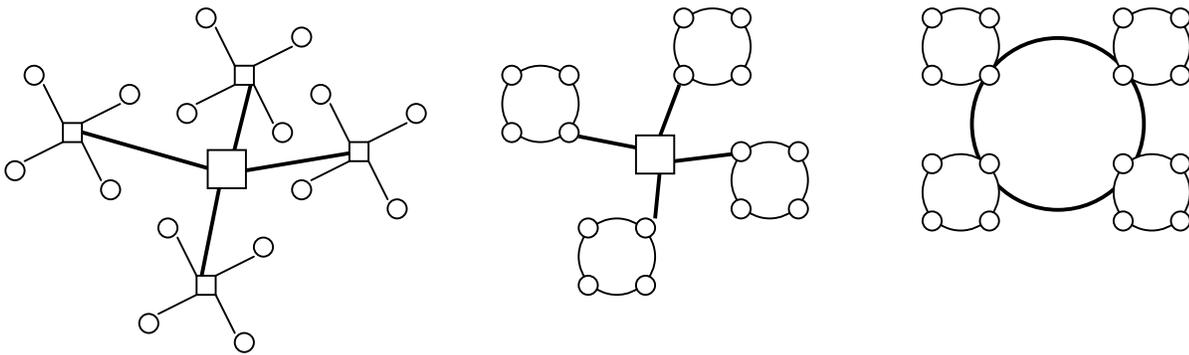
46. The local loop was described as having a star topology in the feeder plant and a star topology in the distribution plant.

- a. Compare the star-star topology with a star-ring topology and a ring-ring topology. Explain how information flows in these topologies and consider issues such as efficiency in use of bandwidth and robustness with respect to faults.

Star topologies generally provide efficient use of bandwidth because their hierarchical routing structure provides near shortest path routing. However, in situations with low load, ring networks are more efficient. When a user is not transmitting any information in a star topology, the bandwidth of the line to the user is wasted. In a ring topology, since the bandwidth is shared, active users can make use of the extra available bandwidth provided by the inactive users.

In general, pure star topologies provide no mechanism for fault line recovery because there is not backup path. Thus, this architecture should be used if the connections are very reliable.

The three topologies under examination are shown below.



The star-star topology is efficient when the load on each of the lines is high. If one of the lines on a high tier star fails, service for many users is lost. Thus, such architecture is only good if the lines are very reliable.

The star-ring topology is good when the traffic produced by each user is relatively sporadic, but the traffic in the “star” links is high. This situation could arise if the bandwidth provided in “star” links is limited in capacity. It still has the same vulnerability as the star-star network in term of large scale loss of service if one of the high-tier links fails.

The ring-ring topology is very robust to faults because every path can be backed up by forming the backup path around the ring in the other direction. In every link in the network, bandwidth is shared and therefore used efficiently.

- b. What role would SONET transmission systems play in the above topologies?

SONET transmission systems could be used in the architectures described. It would be particularly useful in the rings to provide OAM functions and fault protection.

47. Suppose that the local loop is upgraded so that optical fiber connects the central office to the pedestal and twisted pair of length at most 1000 feet connects the user to the pedestal.

- a. What bandwidth can be provided to the user?

Using ADSL a bandwidth of 8 Mbps can be provided. Since the twisted pair length is only 1000 feet, VDSL may be used. In this case, a bandwidth of 55 Mbps can be provided.

- b. What bit rate does the optical fiber have to carry if each pedestal handles 500 users?

If ADSL is used, and each fiber handles 500 users, the total bandwidth on the fiber must be at least $500 \times 8 \text{ Mbps} = 4 \text{ Gbps}$. If VDSL is used, the bandwidth increases to $55 \text{ Mbps} \times 500 = 27.5 \text{ Gbps}$.

- c. How can SONET equipment be used in this setting?

SONET can be used at the physical layer in the optical fiber so that the multiplexing of the 500 lines is simplified using tributaries. This functionality could be extended to the pedestal to provide OAM control of the large aggregate bandwidth connecting it to the central office.

48. Let's consider an approach for providing fiber-to-the-home connectivity from the central office to the user. The telephone conversations of users are time-division multiplexed at the telephone office and *broadcast* over a "passive optical network" that operates as follows. The TDM signal is broadcast on a fiber up to a "passive optical splitter" which transfers the optical signal to N optical fibers that are connected to N users. Each user receives the entire TDM signal and retrieves its own signal.

Solutions follow questions:

- a. Trace the flow of PCM samples to and from the user. What bit rates are required if $N=10$, 100? Compare these to the bit rate that is available.

Because a passive splitter is used, the fiber must have enough bandwidth to support N connections. Thus,

$$N = 10 \rightarrow 640 \text{ kbps minimum capacity}$$

$$N = 100 \rightarrow 6.4 \text{ Mbps}$$

Currently bit rates of 2.5 Gbps on a single wavelength fiber is common and 10 Gbps is available.

- b. Discuss how Internet access service might be provided using this approach.

Internet access could be very simply provided by replacing the PCM samples with Internet traffic.

- c. Discuss how cable television service might be provided using this approach.

Since MPEG2 television signals are 2-5 Mbps, on the order of 1000 channels could be provided by this network to each user. Each TDM slot could correspond to one channel. The users could change channels by resynchronizing their receivers to the appropriate time slot.

- d. What role could WDM transmission play in the connection between the central office and the optical splitter? in the connection all the way to the user?

If WDM is used, the aggregate bandwidth of the optical fiber would be greatly increased. Because the architecture is a passive broadcast system, if WDM were used between the central office and the passive splitter, it would also be on the fiber to the user.

The additional bandwidth might still be leveraged in a number of ways. The additional wavelengths provide an additional degree of multiplexing flexibility. For example, the users could tune into one wavelength at a time (with tuning done on a relatively large time scale) and each wavelength could carry a "cable package" that includes a subset of the provided channels.

Although WDM provides increased bandwidth, a WDM receiver typically represents significantly higher hardware costs. If the huge amounts of available bandwidth were used to carry heterogeneous traffic, the hardware costs required would be offset by the benefits. The resulting network could carry television, telephone, and Internet traffic and, in effect, replace all of the currently existing access networks.

49. Explain where the following fit in the OSI reference model:

Solutions follow questions:

- a. A 4 kHz analog connection across the telephone network.

Physical Layer: the actual 4 kHz analog signal exists only in the physical layer of the OSI reference model.

- b. A 33.6 kbps modem connection across the telephone network.

Data-Link Layer: a 33.6 kbps modem uses framing, flow-control, and error correction to connect a user to the switch.

- c. A 64 kbps digital connection across the telephone network.

Physical Layer: the digital link across the network is controlled by many higher layer functions, but the 64 kbps signal that carries user information is analogous to the 4 kHz signal used over the twisted pair that runs to the user's premises.

50. For the following examples of connections, explain the signaling events that take place inside the network as a connection is set up and released. Identify the equipment and facilities involved in each phase of the connection.

Solutions follow questions:

- a. A voice call in the telephone network.

SSP directs the signaling information to signaling network. SCP sends required commands to individual switches to setup the path. The same procedure takes place to disconnect the call.

- b. A frame relay connection.

Frame relay is a connection oriented network. Once the request for a connection is made all the nodes are programmed to direct the frames to appropriate route.

- c. A SONET connection in metropolitan network as shown in Figure 4.29.

Connection setup requires finding available tributary or STS-n in local rings as well as the metro ring, and regional ring if required. The path is setup by programming the nodes across the path to use the selected tributary or STS-n streams for the connection. Once this is done the transmission can start. To disconnect all of the assigned streams are released.

- d. An optical connection across the continent.

To setup an optical connection across the continent an available wavelength is to be found for every piece of the route from source to destination. Once the wavelength or set of wavelengths is found at each node the switch is programmed or manually set up to direct the signal in the right output according to the path information. To disconnect the wavelengths are set to be free.

51. Sketch the sequence of events that take place in the setting up of a credit-card call over the intelligent network. Identify the equipment involved in each phase.

Solution:

Before a connection is made, a recorded message requests the credit card number for the call. The user inputs the digits and the tones are converted to numbers and the message sent to a database in the appropriate SCP. The credit card number is verified and may be used to check for the authorization of funds. The results of the query are sent to the SCP. The appropriate recorded message is sent back to the user. Based on the query, the call is accepted and a connection made, or the call is denied.

52. Explain how the intelligent network can provide the following services:

Before a connection is made every call has its signaling information in the control plane routed through an SCP that stores the destination phone's profile (that is, identifies the services that have been subscribed).

- a. Caller identification – A display on your phone gives the telephone number or name of the incoming caller.

When the ring signal is sent to the destination phone, if it subscribes to caller ID, then the number used for routing can also be forwarded, via the signaling network to the receiver phone. The receiver phone must have the capability to interpret the signal.

- b. Call forwarding – Allows you to have calls transferred from one phone to another where you can be reached.

When the call request reaches the SCP and call forwarding is activated, the database is referenced to find the forwarding number and the call is rerouted accordingly. Although the call can be incrementally rerouted from this point, bandwidth may be conserved by rerouting the connection again from the start to the forwarding number.

- c. Call answer – Takes voice mail message if you are on the phone or do not answer the phone.

While the phone is ringing, the number of rings must be counted somewhere in the signaling network. After a certain number of rings (specified by the user a priori and stored in an SCP database), if the call is not answered, it is forwarded by a mechanism similar to call forwarding to a voice mailbox center, which handles the caller accordingly.

In another scenario, when a user is on the phone and a call arrives (assuming no call waiting), instead of generating a busy signal, the call is forwarded automatically to the voice mailbox.

- d. Call waiting – If you are on the phone, a distinctive tone informs that you have an incoming local or long-distance call. You can place your current call on hold, speak briefly to the incoming party, and return to the original call.

Generally, if a line is free, a ring is sent to the dialer and receiver. If a line is busy, a busy signal is sent back to the dialer. If call waiting is activated, a tone from the local switch is sent to the receiving party (which briefly interrupts the current conversation).

The control plane treats this tone as a ring (that is, it sends a ring back to the dialing party and voicemail and call forwarding may still be activated). If the user accepts the new call, the current connection is “held” one hop upstream from the user at the local switch, and a second connection is made from the dialing party to the user.

At this point, both calls share the line that makes up the last hop to the user. Thus, the user can only talk to one party at a time. When one of the two connections is being used, the other is maintained in its incomplete state while the caller is on hold.

- e. Called-party identification – Each member of a family has a different phone number. Incoming calls have a different ring for each phone number.

In the signaling network, multiple numbers can all be routed to the same phone. For each number, however, a different ring signal can be sent to the phone. The phone may have only one ringer, so to differentiate the numbers each ring can have a different duration with different length pauses, similar to Morse code. Thus, one number may be one long ring, while the other number may be two short rings.

53. Consider a 64 kbps connection in ISDN.

Solutions follow questions:

- a. Sketch the layers of the protocol stack in the user plane that are involved in the connection. Assume that two switches are involved in the connection.

Only the physical layer in the user plane of the end systems and the two switches are involved in the 64 kbps connection.

- b. Sketch the layers of the protocol stack in the control plane that are involved in the setting up of the connection. Assume that the two switches in part (a) are involved in the call setup.

The physical, data link, network and application layers of the control plane are involved in the setting up of the connection.

54. Explain how the control plane from ISDN can be used to set up a virtual connection in a packet-switching network. Is a separate signaling network required in this case?

Solution:

In order to setup a virtual connection in a packet switched network, the ISDN control plane must first reserve resources during a call setup phase. The circuit does not exist as a physical entity, but rather a sequence of entries in forwarding tables in the packet switched network. As such, each packet in the call follows the same route through the network.

55. Identify the components in the delay that transpires from when a user makes a request for a telephone connection to when the connection is set up. Which of these components increase as the volume of connection requests increases?

Solution:

After the user requests a connection by dialing the phone number, the call setup procedure involves finding a path from the source to the destination. As shown in figure 4.44a, if the two phones are attached directly to the same switch, assuming the switch is non-blocking, the connection can be made and the delay is equal to the processing time in the switch and the short propagation delay of the signal.

If the two phones are not connected to the same switch, a path must be selected for the call using the signaling network. The delay from this stage includes, processing time in each of the signaling nodes involved, propagation delay of the signals through the network, and any routing delay overhead. Before a connection can be made, the receiving user has to pick up the phone. This also contributes to the overall connection setup time.

The amount of time required for the user to answer the call does not depend on the volume of connection requests in the network. However, if there are many simultaneous requests for connections, the signaling network will have to juggle more requests simultaneously and the routing procedure will take longer. Although there is a fixed minimum propagation delay, as call

requests increase, the propagation delay may also increase because the shortest path route may not be available.

56. Discuss the fault tolerance properties of the STP interconnection structure in the signaling system in Figure 4.52.

Solution:

The mesh structure means that any STP is connected to any other STP by more than one path. Thus, if a fault occurs on one link, the signaling traffic can be automatically rerouted over another path. If an STP fails, paths are available to a backup STP.

57. A set of trunks has an offered load of 10 Erlangs. How many trunks are required to obtain a blocking probability of 2%? Use the following recursive expression for the Erlang B blocking probability:

$$B(c, a) = \frac{aB(c-1, a)}{c + aB(c-1, a)} \quad B(0, a) = 1$$

Solution:

Given $a = 10$ Erlangs and $P_B = 2\% = 0.02$, 16 trunks are required to obtain $P_B = 2\%$.

C	a	B(c,a)
0	10	1.000
1	10	0.909
2	10	0.820
3	10	0.732
4	10	0.647
5	10	0.564
6	10	0.485
7	10	0.409
8	10	0.338

c	a	B(c,a)
9	10	0.273
10	10	0.215
11	10	0.163
12	10	0.120
13	10	0.084
14	10	0.057
15	10	0.036
16	10	0.022
17	10	0.013

58. Compare the blocking probabilities of a set of trunks with offered load $a = 9$ and $c = 10$ trunks to a system that is obtained by scaling up by a factor of 10, that is, $a = 90$ and $c = 100$. Hint: Use the recursion in problem 57 in a spreadsheet or program.

Solution:

To find the blocking probability of trunks with offered load $a = 9$ and $c = 10$, we recursively calculate $B(10,9)$ starting from $B(0,9)$. We use a similar method for the case when $a = 90$ and $c = 100$. The results of this process are summarized in the table below.

c	a	B(c,a)
0	9	1.000
1	9	0.900
2	9	0.802
3	9	0.706
4	9	0.614
5	9	0.525
6	9	0.441

c	a	B(c,a)
0	90	1.000
1	90	0.989
2	90	0.978
3	90	0.967
4	90	0.956
5	90	0.945
6	90	0.934

7	9	0.362
8	9	0.289
9	9	0.224
10	9	0.168

7	90	0.923
.	.	.
.	.	.
100	90	0.027

As shown, a lower probability of blocking is achieved by scaling up the system.

59. Calls arrive to a pool of 50 modems according to a Poisson process. Calls have an average duration of 25 minutes.

Solutions follow questions:

- a. What is the probability an arriving calls finds all modems busy if the arrival rate is two calls per minute?

For this system, the utilization $a = 2/(1/25) = 50$ and $c = 50$. The probability of blocking is:

$$P_b = \frac{a^c}{c!} / \sum_{j=0}^c \frac{a^j}{j!} = 0.105$$

- b. What is the maximum arrival rate that can be handled if the maximum acceptable blocking probability is 1%? 10%?

If the maximum acceptable blocking probability is 10% and 1%, the maximum handled arrival rates, given a service rate of 1/25 calls per minute are:

P_b	Load
1%	1.516 calls/minute
10%	1.982 calls/minute

60. Consider dynamic nonhierarchical routing (DNHR).

Solutions follow questions:

- a. Explain how DNHR can be used to exploit the time differences between different time zones in a continent.

In the morning in the east, the business day begins and long-distance calls are made between cities on the East Coast. The west-coast business day has not yet begun, so the links from east to west are relatively under-utilized.

Eastern traffic can be re-routed through the western free lines to provide additional capacity for the traffic between eastern cities. The opposite can occur in the evening after the business day in the east ends and the under-utilized eastern trunks are used for western traffic.

- b. Explain how DNHR can be used to exploit different business and residential activity patterns during the day.

For local calls during business hours, traffic may be re-routed through relatively non-busy lines in residential area to free up resources in the business areas. The opposite can be done in the evening when most people are in the residential areas and the telephone traffic in the business area is low volume.

61. Suppose that setting up a call requires reserving N switch and link segments.

- a. Suppose that each segment is available with probability p . What is the probability that a call request can be completed?

For a call to be completed, every segment in the circuit must be free. Thus, assuming that each switch and link has a probability p of being available, and assuming that this probability is independent of all others,

$$P[\text{every circuit free}] = P[1^{\text{st}} \text{ free}]P[2^{\text{nd}} \text{ free}] \dots P[N^{\text{th}} \text{ free}] = p^N$$

- b. In allocating switch and transmission resources, explain why it makes sense to give priority to call requests that are almost completed rather than to locally originating call requests.

For calls that are almost completed, resources are already locked up, so if the call is denied, there is a large overhead penalty. Thus they should be given priority. To illustrate, suppose locally originating calls are given a higher priority. One can then envision a situation where very few completed connections exist, but all resources are used by partially completed calls.

62. Consider a cellular telephone system with the following parameters: B is the total bandwidth available for the system for communications in both directions; b is the bandwidth required by each channel, including guard bands; R is the re-use factor; and a is the fraction of channels used for set up.

Solutions follow questions:

- a. Find an expression for the number of channels available in each cell.

$$\text{Number of channels available in each cell} = \frac{\frac{B}{b}(1-a)}{R} \text{ channels/cell}$$

- b. Evaluate the number of channels in each cell for the AMPS system.

In AMPS, there are 416 total channels, 21 of which are used for call setup. The reuse factor is 7. Thus the number of channels per cell is:

$$(416 - 21)/7 = 56 \text{ channels per cell}$$

63. Consider the AMPS system in problem 62.

Solutions follow questions:

- a. How many Erlangs of traffic can be supported by the channels in a cell with a 1% blocking probability? 5%?

The probability of blocking is given by the equation:

$$P_b = \frac{a^c}{c!} \bigg/ \sum_{j=0}^c \frac{a^j}{j!}$$

where in the AMPS system $c = 56$ and a is the number of Erlangs of traffic in a cell. If the blocking probability is 10% and 1%, the maximum number of Erlangs that can be handled is:

P_b	Load [Erlangs]
1%	43.342
10%	56.059

- b. Explain why requests for channels from handoffs should receive priority over requests for channels from new calls. How does this change the Erlang load calculations?

In telephone conversations, service interruption is less acceptable than denial of service due to a busy network.

The calculations in part (a) assumed equal priority for all calls. If priority is given to current calls that are being handed off, the Erlang load calculations are more complicated. Newly attempted calls in this scenario have a lower priority and, thus, a higher P_b .

64. Suppose that an analog cellular telephone system is converted to digital format by taking each channel and converting it into n digital telephone channels.

Solutions follow questions:

- a. Find an expression for the number of channels that can be provided in the digital system using the parameters introduced in problem 62.

Assuming that there will be n times more control channels in this new scheme, the number of channels in each cell is:

$$\frac{\frac{B}{b}(1-a)}{R} n$$

- b. Consider the AMPS system and assume that it is converted to digital format using $n = 3$. How many Erlangs of traffic and the new system support at 1% blocking probability? 5%?

This calculation is exactly the same as that in problem 4.62, except that now the number of trunks is tripled, so $c = 160$.

To support 1% blocking or 5% blocking, the traffic can be:

P_b	Load [Erlangs]
1%	141.167
5%	157.046

65. Suppose that a CDMA system has the same number of channels as the digital system in problem 64, but with a reuse factor of 1.

Solutions follow questions:

- a. How many Erlangs of traffic can be supported in each cell by this system at 1% blocking probability? 5%?

This CDMA system has 416 channels. Assuming that 21 channels are still used for call setup and the reuse factor is 1, each cell has 395 channels, so we have:

$$c = 395$$

$$a = \text{Erlangs of traffic}$$

The probability of blocking in this system is given by:

$$P_b = \frac{a^c}{c!} / \sum_{j=0}^c \frac{a^j}{j!}$$

Thus we have:

P_b	Load [Erlangs]
1%	370.410
5%	401.856

- b. Suppose the per capita traffic generated in a city is 0.10 Erlangs during the busiest hour of the day. The city has a population of 1 million residents and the traffic is generated uniformly throughout the city. Estimate the number of cells required to meet the cities traffic demand using the system in part (a).

For a probability of blocking of 1%,

Total number of cells in city

$$\begin{aligned} &= (0.10 \text{ Erlangs/person})(1\,000\,000 \text{ people}) / (370.410 \text{ Erlangs/cell}) \\ &= 270 \text{ Cells} \end{aligned}$$

For a probability of blocking of 5%,

Total number of cells in city

$$\begin{aligned} &= (0.10 \text{ Erlangs/person})(1\,000\,000 \text{ people}) / (401.856 \text{ Erlangs/cell}) \\ &= 249 \text{ Cells} \end{aligned}$$

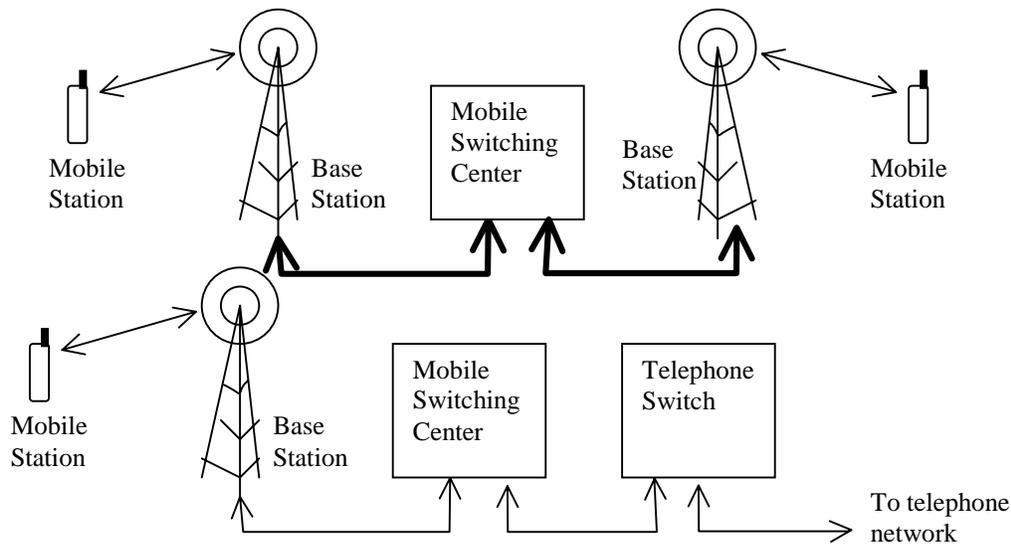
66. Consider the equipment involved in providing a call between mobile and wireline telephones.

Solutions follow questions:

- a. Sketch a diagram showing the various equipment and facilities between an originating mobile telephone to a wireline destination telephone. Suppose first that the mobile carries analog voice; then suppose it carries digital voice.

The mobile phone or mobile subscriber unit (MSU) transmits and receives signals to and from the cell base station (BS) in the area where the phone is located. The cell base station exchanges the signals with the mobile switching center (MSC) located at the mobile telephone switching office (MTSO). All cell base stations are connected to the mobile telephone switching office through wire or microwave channels. The mobile switching center is also connected to the local public switched telephone network (PSTN). The mobile switching center switches the calls among the base stations and also between the base stations and the wireline telephone network.

In analog cellular networks, the signal from the mobile phone arrives in analog form and is then digitized for switching and transmission. In digital cellular systems the signal from the mobile user arrives in digital form. However cellular phones use high performance speech coding formats that require conversion PCM prior to switching in the digital telephone network.



- b. Repeat part (a) in the case where the two telephone calls are mobile.

In this case once the mobile telephone switching office received the call it determines the cell where the requested mobile phone is located and connects the call to the corresponding base station. If the two mobile phones are in the area covered by the mobile telephone switching office, then the signals do not traverse the public telephone network. If the two mobiles are covered by different telephone switches then the signals must be transmitted over digital lines, and possibly across the public telephone network.

- c. In parts (a) and (b) identify the points where a call request can be blocked during call setup because resources are unavailable.

The call can be blocked between the mobile user and the base station if no frequency channels are available. The call cannot likely be blocked between the Base station and the mobile switching center since it is assumed that this link has a capacity at least equal to that of the aggregate traffic from the base station. For this same reason, it is unlikely that the call can be blocked between the mobile switching center and the telephone switch (in part a). The cellular service provider leases a link to the ground-based telephone network, so it may be possible that a call can be blocked at this point if the entire network experiences a lot of traffic.

67. Explain the signaling events that take place when call is set up and released in a cellular telephone system. Identify the equipment and facilities involved in each phase of the call. Consider the following cases.

Solutions follow questions:

- a. The source and destination mobile phones are in the same cell.

Whether or not the destination phone is in the same cell as the source phone, the procedure for call setup may still be the same, because the base station does not know a priori whether the destination phone is in its cell.

- Mobile station sends a request in reverse setup channel along with user information such as phone number, serial number, and authentication
- Base station verifies information with MSC
- MSC sends call request to all base stations

- Base Stations broadcast request to all forward setup channels
- Receiving phone replies on reverse setup channel
- All base stations that receive this reply forward it to the MSC along with other information such as the power level of the received signal. This information helps the MSC to assign the connection channels
- MSC identifies which base station is best suited to service the call for the receiving phone
- The base station assigns a forward and reverse channel and the receiving phone is rung

b. The source and destination mobile phones are in different cells, but in the same MSC.

See answer for part (a).

c. The source and destination mobile phones are in different cellular networks.

If the number is not stored in the MSC database, it must go through the wire-line phone network which in turn will interface with the receiving phone's cellular network as described above via the MSC and the closest base station. At the cellular-wire-line network interface, the data must be reformatted to accommodate the signaling systems of the downstream network.

68. Explain the signaling events that take place when call is handed off from one cell to another cell. Suppose first that the two cells are under the same MSC, and then that they are under different MSCs.

Solution:

The signal level must be periodically monitored when a call is in progress. When it drops below a certain threshold, the base station instructs the mobile station to transmit on a specific setup channel. The base station also signals the MSC, which signals all neighboring base stations to monitor the setup channel. The signal strength and, possibly, other information about the signal is recorded at each base station and sent to the MSC, which uses it to decide which cell is best suited to receive the handoff. The selected base station is informed and the old connection from the MSC to mobile station via the original base station is broken and a new connection via the new base station is established. The service to the mobile phone is briefly interrupted during this process.

The process is centralized if the mobile station moves to a cell that is connected to the same MSC as the original cell. Otherwise it is possible that a more distributed control mechanism is used. If, after receiving the information about the signal quality at each base station, the MSC decides that none of the cells are suitable to receive the handoff, the MSC communicates to neighboring MSCs to locate the mobile station. If the mobile station is found in a cell of one of these neighboring MSCs, the base station is informed and the new connection is made through the new MSC.

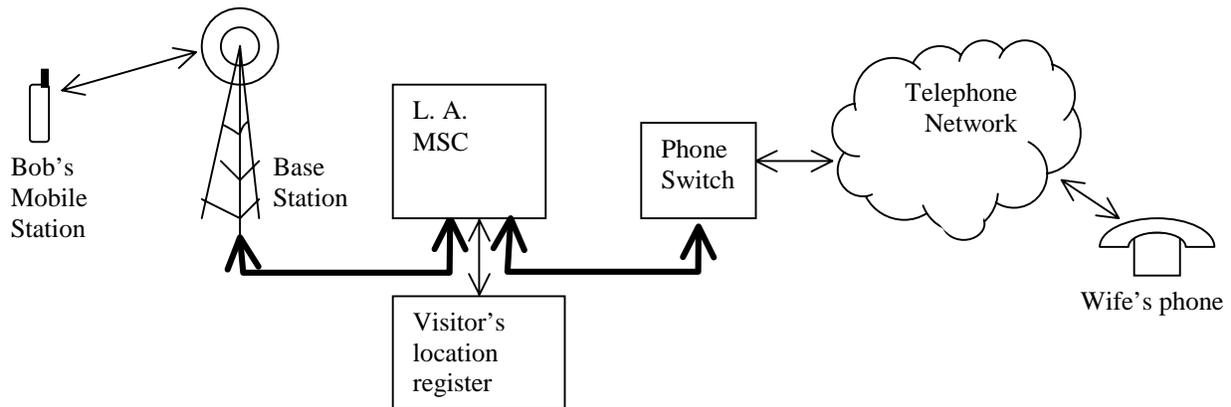
69. Bob Smiley, star salesman, has just arrived in Los Angeles from his home base in Chicago. He turns on his cell phone to contact his Los Angeles client.

a. Sketch the sequence of events that takes place to set up his call. Assume he subscribes to roaming service.

Bob Smiley turns on his phone, which sends user information to the MSC via the nearest base station. The MSC contacts Bob's home location register (in Chicago), and copies Bob's database information into its own visitor location register (in LA). Further, Bob's home MSC is not aware that his calls should be forwarded to LA. Now Bob can use his phone in the new network area.

b. Next he calls home in Chicago to inform his wife that he forgot to take out his son's hockey gear from the trunk of his car and to give her the parking spot where he left the car in the

airport (“somewhere on the third level of parking lot A”). Sketch the sequence of events that take place to set up this call. (Don’t concern yourself with the specifics of the conversation.)



- c. In the meantime, Bob’s college roommate, Kelly, who now works in Hollywood calls Bob’s cell phone. Note that Bob’s cell phone begins with the Chicago area code. Sketch the sequence of events that take place to set up this call. Should this call be billed as local or long distance?

Kelly’s long-distance service provider sees Bob’s Chicago area code and proceeds with connecting her to Bob’s Chicago cellular network. When the call arrives at Bob’s cellular network (in the form of signaling traffic), Bob’s home MSC checks its home location register and sees that Bob is in LA (recall that it found out about his whereabouts in part a)). The call is then forwarded back to LA where the connection can be made via Bob’s temporary MSC and its visitors’ location register.

Whether Bob or Kelly should be billed long distance depends on the setup of Bob’s “roaming service”. If the three networks involved (two cellular and one wireline) are coordinated enough, Kelly’s final connection could be routed directly from Hollywood to Bob’s mobile station directly. In this case it could be possible to bill them both locally.

However, it is more likely that the connection will actually be made from Kelly to Bob’s home network and then back to L. A. By dialing long distance, Kelly implicitly agrees that she will pay the long distance charges for a call to Chicago. By using his phone in L. A. Bob expects to pay long distance fees for all calls he receives (even local ones) because if someone from Chicago calls him, the call must be carried by a long distance carrier. In this case, both Bob and Kelly will pay long distance charges. She will pay the same as if she were dialing a regular Chicago number (which she is), and Bob will pay as if he is receiving a call from a Chicago caller, say, his wife.

70. Compare cellular wireless networks to the local loop and to the coaxial cable television system in terms of their suitability as an integrated access network. In particular comment on the ability to support telephone service, high-speed Internet access, and digital television service. Consider the following two cases:

In question 45, the local loop and coaxial cable network were compared in terms of their ability to act as integrated access networks. Here, we examine the ability of cellular wireless networks to provide these services.

Each type of traffic on an integrated access network has very different characteristics. Digital television requires high bandwidth and low loss. Although the signal can be compressed using

MPEG encoding, the compressed signal still requires a minimum of 1-2 Mbps. Service can most efficiently be provided in a broadcast network because it lacks the interactivity of phone and Internet, and each user can be sent the same information.

Telephone service is less constrained by bandwidth. Because the vast majority of calls are between two users, point-to-point networks are well suited for carrying phone traffic.

Current high speed Internet access has medium bandwidth requirements. IP traffic can also be adaptive, in the sense that if need be, the connection can usually be slowed down. This affects the quality of service, but unlike the other two (phone and TV), it does not interrupt service completely.

- a. The band of frequencies available spans 300 MHz.

Current cellular networks have only moderate bandwidth. Cellular networks share the same scalability problem as the coaxial television network. If the band of frequencies in the network span 300 MHz (about 10x the current band), the bandwidth available to each user is inversely proportional to the number of users. Of course cells can be made small, and CDMA can be used so that the reuse factor is 1, but a bottleneck still exists. Clearly, the cellular network can support telephone traffic, since it currently already does. Successful efforts have been made to provide Internet service on mobile cellular devices, but due to the bandwidth restriction mentioned, the service quality is moderate. New protocols are currently being defined that reduce bandwidth for Internet services that run on devices of these networks.

It would be difficult to provide digital television using the above-described cellular network. Because of the high bandwidth required, by partitioning the network into a number of parallel channels, a broadcast network could provide digital television service.

- b. The band of frequencies available spans 2 GHz.

If the band of frequencies available spanned 2 GHz, all of the above obstacles would be significantly reduced. Since the aggregate bandwidth of this system is 100 times that of the current cellular networks, most of the problems mentioned above would melt away. With a rough calculation, assuming each cell has on the order of 100 users, and the frequency reuse factor is 1, we see that a bandwidth on the order of 10 Mbps can be provided to each user (or to each channel in a 100-station digital television network). This is comparable to current ADSL in the local loop network, and would likely be sufficient to serve as an integrated access network. A difficulty with this scenario is that the region of spectrum where 2 GHz is available would be in the 28 GHz band or beyond. These frequencies tend to be suitable to point-to-point or point-to-multipoint but not to omnidirectional communications.

Solutions to Chapter 5

Version 1.01 (list of changes on last page)

1. Explain the difference between connectionless unacknowledged service and connectionless acknowledged service. How do the protocols that provide these services differ?

Solution:

In an acknowledged connectionless service, reliable delivery can be achieved through the use of ACK and NAK control messages. Such protocols are suited for communication over networks in which higher layers are sensitive to loss and the underlying network is inherently unreliable with a significant probability of loss or error. For example, HDLC provides for unnumbered acknowledgment service for connection setup and release.

Unacknowledged networks provide simpler and faster communication for networks that are inherently reliable or provide service to higher layers that can tolerate information loss or have built-in error recovery mechanisms.

2. Explain the difference between connection-oriented acknowledged service and connectionless acknowledged service. How do the protocols that provide these services differ?

Solution:

The use of acknowledgments can provide reliable transfer over links or networks that are prone to error, loss, and or resequencing. In a connection-oriented service, a setup phase between the sending user and receiving user establishes a context for the transfer of information. In a connection-oriented acknowledged service acknowledgments are provided to the sending user for all SDUs.

In a connectionless service, there is no prior context provided for the transfer of information between the sending user and the receiving user. The sender passes its SDU to the underlying layer without prior notice. In an acknowledged connectionless service, the sending user requires an acknowledgment of delivery of its SDU.

The protocols that provide these services are very different. Connection-oriented acknowledged service requires the use of stateful protocols that keep track of sequence numbers, acknowledgments, and timers. Connectionless services use much simpler protocols that are stateless in nature. Connectionless acknowledged service does require that the transmitting protocol track the acknowledgment of a PDU. In the simplest instance, the receiver would be required to send an ACK for correctly received PDU and the transmitter would keep a timer. If an ACK was not received in time, the transmitter would inform the user of a failure to deliver.

3. Suppose that the two end systems α and β in Figure 5.6 communicate over a connection-oriented packet network. Suppose that station α sends a 10-kilobyte message to station β and that all packets are restricted to be 1000 bytes (neglect headers); assume that each packet can be accommodated in a data link frame. For each of the links, let p be the probability that a frame incurs errors during transmission.

Solutions follow questions:

- a. Suppose that the data link control just transfers frames and does not implement error control. Find the probability that the message arrives without errors at station β .

Let p be the probability that a frame incurs errors during transmission. We know the following:

Message length = 10,000 bytes

Maximum packet size = 1000 bytes
 Number of packets for transmission = 10

The probability of a packet arriving error free at end system $P_{packet} = (1 - p)^3$. The probability that all packets arrive error free at end system β is $P_{error} = [(1 - p)^3]^{10} = (1 - p)^{30} \approx e^{-30p}$.

- b. Suppose that error recovery is carried out end to end and that if there are any errors, the entire message is retransmitted. How many times does the message have to be retransmitted on average?

The average number of required transmissions = $1/P_{error} = e^{30p}$.

- c. Suppose that the error recovery is carried out end to end on a packet-by-packet basis. What is the total number of packet transmissions required to transfer the entire message?

The average number of transmissions per packet = $1/P_{packet} = e^{3p}$. The total number of packet transmissions is then $10/P_{packet} = 10e^{3p}$.

As an example suppose $p = .01$, then the message retransmission approach requires 1.35 message transmissions. The packet transmission approach requires 1.03 message transmissions. Clearly packet-by-packet retransmission is better.

4. Suppose that two peer-to-peer processes provide a service that involves the transfer of discrete messages. Suppose that the peer processes are allowed to exchange PDUs that have a maximum size of M bytes including H bytes of header. Suppose that a PDU is not allowed to carry information from more than one message.

Solutions follow questions:

- a. Develop an approach that allows the peer processes to exchange messages of arbitrary size.

To exchange messages of arbitrary size, large messages must be segmented into parts of $M-H$ bytes each in length to be transmitted in multiple PDUs. Small messages must be placed in a single PDU.

- b. What essential control information needs to be exchanged between the peer processes?

The peer processes need to communicate information that allows for the reassembly of messages at the receiver. For example, the first PDU may contain the message length. The last PDU may contain an end-of-message marker. Sequence numbers may also be useful to detect loss in connection oriented networks and to help in reconstruction of the messages in connectionless networks. Lastly, since variable size PDUs are permitted, the size of the PDU payload must be transmitted in the PDU header.

- c. Now suppose that the message transfer service provided by the peer processes is shared by several message source-destination pairs. Is additional control information required, and if so, where should it be placed?

In this case, in addition to all of the header information mentioned in b, each PDU must be labeled with a stream ID, so that the receiver can treat each stream independently when reassembling messages. This stream ID may be avoided if the source and destination operate so that they handle the transfer of a single message at a time. For example, this approach is used by AAL5 in ATM.

5. Suppose that two peer-to-peer processes provide a service that involves the transfer of a stream of bytes. Suppose that the peer processes are allowed to exchange PDUs that have a maximum size of M bytes, including H bytes of header.

Solutions follow questions:

- a. Develop an approach that allows the peer processes to transfer the stream of bytes in a manner that uses the transmission line efficiently. What control information is required in each PDU?

The streams should be segmented into $M - H$ size blocks and transmitted in the PDUs. The best possible efficiency in line usage occurs when every PDU is of maximum size. If the byte stream arrives at a steady rate, full PDUs can be constructed by waiting until a sufficient number of bytes to fill a PDU payload have arrived.

The control information should include sequence numbering, payload size, and error checking. If reliability is required, the control information should also allow for ACK control messages.

- b. Suppose that the bytes in the stream arrive sporadically. What is a reasonable way to balance efficiency and delay at the transmitter? What control information is required in each PDU?

If the bytes stream arrives sporadically, arriving bytes can be buffered while the transmitter waits to fill a PDU. A timer can be used to place a bound on the delay incurred waiting.

- c. Suppose that the bytes arrive at a constant rate and that no byte is to be delayed by more than T seconds. Does this have an impact on the efficiency?

Yes. Because the header length is constant, larger PDUs provide more efficient data transfer. If T is very small, partially filled packets must always be sent. If T is large, larger PDUs can be transmitted so more efficient transmission can be achieved. Because the bytes arrive at a constant rate, the PDUs will all be the same length. Thus, there is no need for a PDU length field in the header.

- d. Suppose that the bytes arrive at a variable rate and that no byte is to be delayed by more than T seconds. Is there a way to meet this requirement?

A timer is required that counts down from T seconds. When the first byte arrives the timer starts. When a PDU is full or when the timer expires (whichever occurs first), the PDU is transmitted and the timer is restarted upon the arrival of the next byte.

6. Suppose that two peer-to-peer processes provide a service that involves the transfer of a stream of bytes. Develop an approach that allows the stream transfer service to be shared by several pairs of users in the following cases:

Solutions follow questions:

The protocol that provides the service must be able to multiplex information from different user pairs for transfer. A basic design decision is whether the PDUs should carry information from multiple users or be constrained to carry information from a single user. The latter approach is simpler in that only a single user needs to be identified per PDU. The latter approach is complicated by the need to identify several users (actually multiplexing IDs) per PDU.

- a. The bytes from each user pair arrive at the same constant rate.

If the bytes arrive at a constant rate, it may be possible to arrange a PDU structure that transfers the merged byte stream without extensive multiplexing ID information. For example, the position of a byte in the payload could identify which user it belongs to. Alternatively, bytes for each user can be buffered until they can fill a minimum sized payload. Note that the first approach will involve less delay than the second approach.

- b. The bytes from the user pairs arrive sporadically and at different rates.

Mixing bytes from the different streams poses a greater challenge in this case. The bytes from different streams need to be buffered separately until they can fill a minimum sized PDU in the

approach where each PDU carries information from a single user. Alternatively, the bytes from a stream could fill a minimum-size sub-payload that could then be carried with sub-payloads from other users on shared PDUs. The choice of minimum size payload will tradeoff efficiency against delay.

7. Consider the transfer of a single real-time telephone voice signal across a packet network. Suppose that each voice sample should not be delayed by more than 20 ms.

Solutions follow questions:

- a. Discuss which of the following adaptation functions are relevant to meeting the requirements of this transfer: handling of arbitrary message size; reliability and sequencing; pacing and flow control; timing; addressing; and privacy, integrity and authentication.

Message size is important because in real-time voice, samples are delayed while waiting to fill a packet. A portion of the 20 ms delay must be apportioned to the packetization delay, which in turn determines the packet payload size. The handling of arbitrary message size is not as important as long as the desired packet size for voice can be handled.

Sequencing is important because the voice samples need to be played back in the same sequence that they were generated.

Reliability is moderately important since voice transmission can tolerate a certain level of loss and error.

Pacing and flow control are not as important because the synchronous nature of the voice signal implies that the end systems will be matched in speed.

Timing, for real-time voice transfer is important because this adaptation function helps to control the jitter in the delivered signal.

Addressing is only during the connection setup phase if we assume some form of virtual circuit packet switching method. Addressing the form of multiplexing ID may be required if multiplexing is involved.

Privacy, integrity, and authentication have traditionally not been as important as the other issues discussed above. However there will surely be applications where it is essential that the voice information be kept confidential (privacy), that the information be tamper free (integrity), and that imposters be detected and deterred (authentication.)

- b. Compare a hop-by-hop approach to an end-to-end approach to meeting the requirements of the voice signal.

Sequencing and timing are the most important requirements for real-time voice. These requirements are better met in a hop-by-hop approach than in an end-to-end approach because delay performance is critical. Resequencing on an end-to-end basis may lead to excessive delay. Hop-by-hop controls on transfer delay may be critical to achieving real-time transfer.

8. Suppose that a packet network is used to transfer *all* the voice signals that arrive at the base station of a cellular telephone network to a telephone office. Suppose that each voice sample should not be delayed by more than 20 ms.

Solutions follow questions:

- a. Discuss which of the following adaptation functions are relevant to meeting the requirements of this transfer: handling of arbitrary message size; reliability and sequencing; pacing and flow control; timing; addressing; and privacy, integrity and authentication.

The following table summarizes parts (a) and (b):

Adaptation function	Relevant in Upstream Direction (part a)	Relevant in Downstream Direction (part b)
Handling of arbitrary message size	No	No
Reliability and sequencing	Yes	Yes
Pacing and flow control	No	No
Timing	Yes	Yes
Addressing	Yes	Yes
Privacy, integrity and authentication	Yes	No

Because traffic in a cellular network consists of constant rate streams of information, the message size need not be arbitrary, and the timing and sequence of packets is essential for service.

Addressing is necessary to identify the two end callers, and authentication may only be necessary in the upstream direction, although if it is done at the base station, it need not be repeated. Pacing and flow control are not necessary because cellular networks service constant bit rate traffic and only admit calls that can be accommodated at that constant bit rate.

The key difference from the case of a single voice call (considered in Problem 7) is the multiplexing aspect of handling of multiple voice streams. The considerations in multiplexing multiple byte streams of Problem 6 apply. The separate voice streams could be carried in separate PDUs or they could be combined in the payload of the PDUs. The first approach is simpler but involves greater delay. The second approach can lead not only to lower delays but also to higher efficiency because larger payloads are possible using multiple voice streams.

Mobility is another aspect that needs to be considered. As users move from cell to cell, the network must track their movement and perform handoffs as cell boundaries are traversed. Changes in addresses and routing are involved in these handoffs.

b. Are the requirements the same in the opposite direction from the telephone office to the base station?

The set of requirements are essentially the same in both directions. However, the direction from the users upstream into the network is more difficult in general because of the need to deal with multiple incoming streams.

c. Do the answers to parts (a) and (b) change if the signals arriving at the base station include e-mail and other short messages?

If the signals include email and other short messages in addition to voice, further adaptation functions are required. The handling of arbitrary message sizes is needed, since email messages are variable in length. Alternatively, fixed small message sizes may be specified, as in Short Message Service (SMS). Pacing and flow control may now be needed depending on the nature of the messages and applications. Text email is low in bandwidth and likely would not require flow control, but very large messages, particularly those transmitted in the downstream direction to the cellular user would require flow control.

9. Suppose that streaming video information is transferred from a server to a user over a packet network.

Solutions follow questions:

a. Discuss which of the following adaptation functions are relevant to meeting the requirements of this transfer: handling of arbitrary message size; reliability and sequencing; pacing and flow control; timing; addressing; and privacy, integrity and authentication.

Streaming video involves the delivery of video information across a network. The term “streaming” usually implies a non-real-time situation where significant delay can be tolerated. The following table summarizes the results of part (a).

Adaptation function	Required	Discussion
Handling of arbitrary message size	No	A stream has no inherent required message size. However video and audio formats do have specific structures, which could be relevant to how transfer is carried out.
Reliability and sequencing	Yes	Reliability is important for video & audio quality. Sequencing is important but can be relaxed as buffering delay is increased.
Pacing and flow control	Maybe	See note below
Timing	Yes	Video has strict timing jitter requirements. Synchronization of different media is also important, but not an issue when the media are combined before packetization.
Addressing	Maybe	Required for point-to-point or point-to-multipoint transfer, but not in broadcast system. Required for different media components.
Privacy, integrity and authentication	Maybe	Required for point-to-point or point-to-multipoint transfer, but <i>may</i> not be required in broadcast system

Whether explicit pacing and flow control is necessary depends on the nature of the receiver. If the receiver is dedicated for receiving video, it should be able to handle the incoming signal. If other applications may be received at the same time, flow control might be needed.

If the signal is constant-rate, uncompressed video, flow control is not possible. For a compressed signal, flow control may be implemented by using MPEG, which transmits video in layers that can be selectively discarded to reduce the bandwidth to the receiver.

- b. Suppose that the user has basic VCR features through control messages that are transferred from the user to the server. What are the adaptation requirements for the control messages?

The following table summarizes the results of part (b).

Adaptation function	Required	Discussion
Handling of arbitrary message size	No	Control messages are short so this is not an issue.
Reliability and sequencing	Yes	Necessary to allow user to perform many operations in a short time period
Pacing and flow control	No	Control messages would not likely produce high bandwidth.
Timing	Yes	Interactive control has timing requirements.
Addressing	Yes	The control traffic must be addressed to the video server, but a generalized control protocol for multiple devices is possible.
Privacy, integrity and authentication	Yes	Authentication is important that only the users have control over the content they receive. Privacy for control information may be provided, but would likely not be a strict requirement.

10. Discuss the merits of the end-to-end vs. hop-by-hop approaches to providing a constant transfer delay for information transferred from a sending end system to a receiving end system.

Solution:

Jitter in networks has two primary sources – variation in queuing delay and variation in propagation delay. The former is caused because traffic intensity at nodes varies with time, so the buffering that is required at each node is variable. The latter exists primarily in connectionless networks and results because each packet in a flow can traverse a different path through the network to its destination.

Hop-by-hop approaches can be used to deal with queuing delay variation, and end-to-end approaches are required to deal with path-length variation. By setting up a path prior to transmission, a constant propagation delay can be ensured. To deal with queuing delay variation, the scheduler at a node can give delay sensitive traffic priority, ensuring that its delay is kept below some maximum

amount. This can be based on header information in the packet. If the path for a flow has been previously specified, the number of nodes in the path is known, so an overall queuing delay maximum is, thus, insured.

End-to-end approaches can also provide constant transfer delay, for example, by attaching timestamps to packets obtained from a common clock, e.g. Global Positioning System, and buffering information at the receiving end until a target delivery time.

11. Consider the Stop-and-Wait protocol as described in the chapter. Suppose that the protocol is modified so that each time a frame is found in error at either the sender or receiver, the last transmitted frame is immediately resent.

Solutions follow questions:

a. Show that the protocol still operates correctly.

The protocol will operate correctly because the only difference is that frames are retransmitted sooner than otherwise. The detection of errors in an arriving frame at the receiver will cause an ACK to be sent sooner, possibly causing the transmitter to retransmit sooner. The detection of errors in an arriving frame at the transmitter will cause an immediate retransmission of the current information frame.

b. Does the state transition diagram need to be modified to describe the new operation?

The state transition diagram remains the same.

c. What is the main effect of introducing the immediate-retransmission feature?

The main effect is a speedup in the error recovery process.

12. In Stop-and-Wait ARQ why should the receiver always send an acknowledgment message each time it receives a frame with the wrong sequence number?

Solution:

The sender cannot send the next frame until it has received the ACK for the last frame so, if the receiver gets a frame with the wrong sequence it has to be a retransmission of the previous frame received. This means that the ACK was lost so the receiver has to ACK again to indicate the sender that it has received the frame.

13. Discuss the factors that should be considered in deciding whether an ARQ protocol should act on a frame in which errors are detected.

Solution:

If a frame is in error, then all of the information contained in it is unreliable. Hence any action taken as a result of receiving an erroneous frame should not use the information inside the frame. A viable option when an erroneous frame is received is to do nothing, and instead to rely on a timeout mechanism to initiate retransmission. However error recovery will be faster if we use a NACK message to prompt the sender to retransmit. The inherent tradeoff is between the bandwidth consumed by the NACK message and the faster recovery.

14. Suppose that a network layer entity requests its data link layer to set up a connection to another network layer entity. To setup a connection in a data link, the initiating data link entity sends a SETUP frame, such as SABM in Figure 5.47. Upon receiving such a frame, the receiving data link entity sends an acknowledgment frame confirming receipt of the SETUP frame. Upon receiving this acknowledgment the initiating entity can inform its network layer that the connection has been setup and is ready to transfer information. This situation provides an example of how *unnumbered acknowledgments* can arise for confirmed services.

Solutions follow questions:

- a. Reexamine Figure 5.10 and Figure 5.11 with respect to error events that can take place and explain how these events are handled so that connection setup can take place reliably.

Fundamentally, the problem involves getting the receiver state to change from an idle state to a connected state. This is the same as getting the state to go from state (0,0) to state (1,1) in Figure 5.10. Therefore transmissions of the SETUP message must be accompanied by the start of an associated timer to trigger retransmissions. Also, the receiver must acknowledge every SETUP frame that it receives. The sender should ignore redundant SETUP acknowledgments and the receiver should ignore redundant SETUP frames. In order to allow multiple connections, each flow's SETUP and acknowledgment frames should be indexed. In the absence of connection indexing, the link should only handle one SETUP at a time until the SETUP is confirmed through an acknowledgment. Otherwise ambiguities in terms of what frame an ACK corresponds to will arise.

- b. To terminate the connection, either data link layer can send a DISC frame that is then acknowledged by an unnumbered acknowledgment. Discuss the effect of the above error events and how they can be dealt with.

This problem is very similar to problem a. except that in the SETUP problem the actual transfer of frames cannot begin until a SETUP acknowledgment is received. In the disconnect case, the sender can stop transmitting after a certain number of retransmissions of the DISC message.

- c. Suppose that an initiating station sends a SETUP frame twice but that the corresponding ACK frames are delayed a long time. Just as the ACK frames from the original transmissions are about to arrive, the initiating station gives up and sends a DISC frame followed by another SETUP frame. What goes wrong if the SETUP frame is lost?

If the first delayed ACK frame arrives after the initiating station sends the DISC frame and the second delayed ACK frame arrives after it sends the SETUP frame, the initiating station will assume the connection is established. Meanwhile the receiver receives the DISC frame and disconnects the connection and because the last SETUP frame is lost it remains in this state.

To resolve this problem the initiating station should not send out SETUP messages with DISC messages outstanding.

15. A 1 Mbyte file is to be transmitted over a 1 Mbps communication line that has a bit error rate of $p = 10^{-6}$.

Solutions follow questions:

The file length $n = 8 \times 10^6$ bits, the transmission rate $R = 1$ Mbps, and $p = 10^{-6}$.

- a. What is the probability that the entire file is transmitted without errors? Note for n large and p very small, $(1 - p)^n \approx e^{-np}$.

$$\begin{aligned} P[\text{no error in the entire file}] &= (1 - p)^n \approx e^{-np}, \text{ for } n \gg 1, p \ll 1 \\ &= e^{-8} = 3.35 \times 10^{-4} \end{aligned}$$

We conclude that it is extremely unlikely that the file will arrive error free.

- b. The file is broken up into N equal-sized blocks that are transmitted separately. What is the probability that all the blocks arrive correctly without error? Does dividing the file into blocks help?

A subblock of length n/N is received without error with probability:

$$P[\text{no error in subblock}] = (1 - p)^{n/N}$$

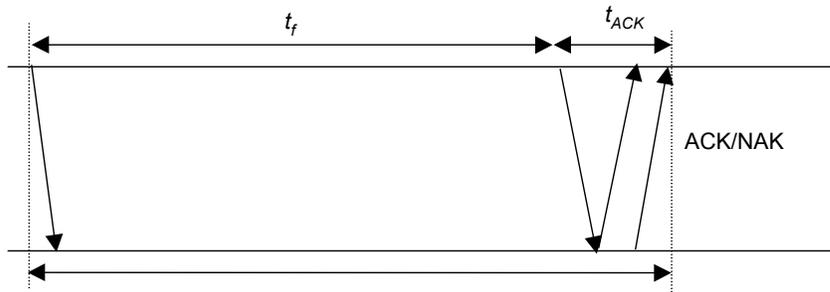
A block has no errors if all subblocks have no errors, so

$$P[\text{no error in block}] = P[\text{no errors in subblock}]^N = ((1 - p)^{n/N})^N = (1 - p)^n$$

So simply dividing the blocks does not help.

- c. Suppose the propagation delay is negligible, explain how Stop-and-Wait ARQ can help deliver the file in error-free form. On the average how long does it take to deliver the file if the ARQ transmits the entire file each time?

Refer to the following figure for the discussion.



We assume the following:

- t_0 = basic time to send a frame and receive the ACK/NAK $\approx t_{\text{timeout}}$
- t_{total} = total transmission time until success
- n_f = number of bits/frame
- n_a = number of bits per ACK
- n_t = number of transmissions
- P_f = probability of frame transmission error

$$t_0 = t_f + t_{ACK} = n_f/R + n_a/R \quad (t_{\text{prop}} \approx 0).$$

$$P[n_t = i] = P[\text{one success after } i - 1 \text{ failure}] = (1 - P_f) P_f^{i-1}$$

$$t_{\text{total}} | i \text{ transmissions} = i \cdot t_0$$

$$E[t_{\text{total}}] = \sum_{i=1}^{\infty} i t_0 P[n_t = i] = t_0 (1 - P_f) \sum_{i=1}^{\infty} i P_f^{i-1} = t_0 (1 - P_f) / (1 - P_f)^2 = t_0 / (1 - P_f)$$

Here, $n_f = n \gg n_a$ thus $t_0 \approx t_f = n/R$; and $P_f = 1 - P[\text{no error}] = 1 - e^{-np}$

$$E[\text{total}] = n/R (1 - P_f) = n/[R e^{-np}] = 8 / (3.35 \times 10^{-4}) = 23,847 \text{ seconds} = 6.62 \text{ hours!}$$

The file gets through, but only after many retransmissions.

- d. Now consider breaking up the file into N blocks. (Neglect the overhead for the header and CRC bits.) On the average how long does it take to deliver the file if the ARQ transmits the blocks one at a time? Evaluate your answer for $N = 80, 800, \text{ and } 8000$.

For 1 block $P_f = 1 - P_b = 1 - (1 - p)^{n/N}$ and $n_f = n/N$

if $t_{\text{prop}} \approx 0$ and $n_a \ll n/N$: $t_0^b = n_f/R = n/NR$

$$T_b = E[t_{\text{total}}^b] = t_0^b / (1 - P_f) = n(1 - p)^{-n/N} / NR \quad \text{average time to transmit one block}$$

$$T = E[t_{\text{total}}] = N T_b = n(1 - p)^{-n/N} / R = 8 (1 - p)^{-n/N} = 8 e^{np/N} \quad \text{if } n/N \gg 1, p \ll 1$$

- $N = 80 \Rightarrow T \approx 8 e^{0.1} = 8.84 \text{ sec}$
- $N = 800 \Rightarrow T \approx 8 e^{0.01} = 8.08 \text{ sec}$
- $N = 8000 \Rightarrow T \approx 8 e^{0.001} = 8.008 \text{ sec}$

Each subblock has a higher probability of arriving without errors, and so requires fewer retransmissions to deliver error free. The overall delay is reduced dramatically.

e. Explain qualitatively what happens to the answer in part (d) when the overhead is taken into account.

As N increases, the effect of overhead becomes more significant because the headers constitute a bigger fraction of each subblock.

16. Consider the state transition diagram for Stop-and-Wait ARQ in Figure 5.12. Let P_f be the probability of frame error in going from station A to station B and let P_a be probability of ACK error in going from B to A. Suppose that information frames are two units long, ACK frames are one unit long, and propagation and processing delays are negligible. What is the average time that it takes to go from state (0,0) to state (0,1)? What is the average time that it then takes to go from state (0,1) to state (1,1)? What is the throughput of the system in information frames/second?

Solution:

We know that P_f is the probability of frame error and P_a is the probability of ACK error. We assume that:

- X is the random variable that represents the number of trials before a successful transmission of a frame. Each unsuccessful trial requires a timeout for retransmission. We assume that the timeout time is set to be equal to frame transmission time plus ACK transmission time.
- Y is the random variable that represents the number of trials before a successful transmission of an ACK. An ACK error will require a new successful retransmission of the frame for next ACK transmission. An ACK is not sent until a new retransmitted frame arrives at the receiver.
- X and Y follow a geometric random-variable distribution.

$$T1 = \text{Average time to go from } (0,0) \text{ to } (0,1) = (T_f + T_a) E(X) + T_f$$

$$T1 = (2 + 1) \frac{P_f}{1 - P_f} + 2 = 2 \left[\frac{P_f}{1 - P_f} + 1 \right] + \frac{P_f}{1 - P_f} = \frac{2}{1 - P_f} + \frac{P_f}{1 - P_f} = \frac{2 + P_f}{1 - P_f}$$

$$T2 = \text{Average time to go from } (0,1) \text{ to } (1,1) = T_1 E(Y) + T_a$$

$$T2 = T_1 \left[\frac{P_a}{1 - P_a} \right] + T_a = \left[\frac{2 + P_f}{1 - P_f} \right] \left[\frac{P_a}{1 - P_a} \right] + 1$$

$$\text{Throughput} = \text{Frame Time} / \text{Expected Total Transmission Time} = 2 / (T1 + T2)$$

P_f	P_a	$T1$	$T2$	Throughput
0.2	0.1	2.75	1.3055	0.4931
0.02	0.01	2.06	1.0208	0.6492
0.002	0.001	2.006	1.0020	0.6649
0	0	2.0000	1.0000	0.6667

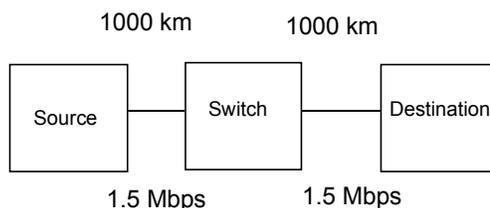
17. Write a program for the transmitter and the receiver implementing Stop-and-Wait ARQ over a data link that can introduce errors in transmission. Assume station A has an unlimited supply of frames to send to station B. Only ACK frames are sent from station B to station A. Hint: Identify each event that can take place at the transmitter and receiver and specify the required action.

Solution:

Stop and Wait Transmitter				
Current State	Event	Sequence Number	Action	Next State
Ready	Request from upper layer		Prepare frame and send; start timer	Wait
Wait	Arrival of error-free ACK frame	Correct sequence number $R_{next} = S_{last} + 1$	Increment Send Sequence Number $S_{last} = R_{next}$	Ready
Wait	Request from upper layer		New requests are blocked	Wait
Wait	Arrival of error-free ACK frame	Incorrect sequence number	Discard frame	Wait
Wait	Arrival of erroneous frame		Discard frame	Wait
Wait	Timeout Expires		Resend frame; start timer	Wait

Stop and Wait Receiver				
Current State	Event	Sequence Number	Action	Next State
Ready	Arrival of error-free frame	Expected Sequence Number $S_{last} = R_{next}$	Accept frame, increment R_{next} , $R_{next} = R_{next} + 1$, and send ACK, deliver packet to higher layer	Ready
Ready	Arrival of error-free ACK frame	Incorrect sequence number	Discard frame; send ACK with R_{next}	Ready
Ready	Arrival of erroneous frame		Discard frame	Ready

18. A 64-kilobyte message is to be transmitted from the source to the destination. The network limits packets to a maximum size of two kilobytes, and each packet has a 32-byte header. The transmission lines in the network have a bit error rate of 10^{-6} , and Stop-and-Wait ARQ is used in each transmission line. How long does it take on the average to get the message from the source to the destination? Assume that the signal propagates at a speed of 2×10^5 km/second.



Solution:

Message Size 65536 bytes
 Max Packet Size 2048 bytes
 Packet Header 32 bytes
 Available for info 2016 bytes
 # of packets needed 32.51 packets
 Total 33 packets

bit error rate 1E-06
 bits/packet 16384
 Probability of error in packet 0.016251 $1 - (1 - \text{bit_error_rate}) ^ (\text{bits/packet})$
 Propagation speed 2E+05 Km/s
 Distance 1000 Km
 Bandwidth 1.5 Mb/s

We assume that the ACK error, the ACK time, and processing time are negligible.

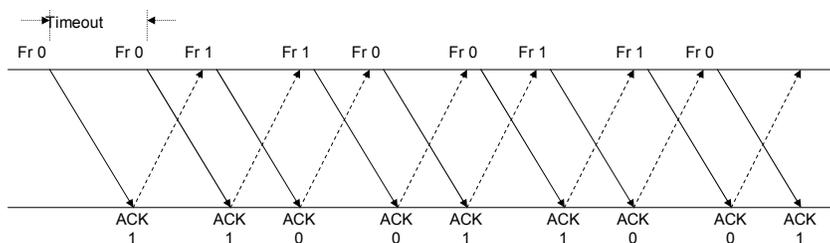
$T_{prop} = \text{distance} / \text{propagation speed} = 0.0050 \text{ s}$
 $T_f = \text{packet size} / \text{bandwidth} = 0.0109 \text{ s}$
 $T_0 = T_{prop} + T_f = 0.0159 \text{ s}$
 $P_f = \text{probability of error in packet} = 0.016251$

$E[T_{total}] = T_0 / (1 - P_f) = 0.0162$

There is pipelining effect that occurs as follows: After the first packet arrives at switch 1, two transmissions take place in parallel. The first packet undergoes stop-and-wait on the second link while the second packet undergoes stop-and-wait in the first link. The packet arriving at the switch cannot begin transmission on the next link until the previous packet has been delivered, so there is an interaction between the transmission times of the two packets. We will neglect this effect. The time to send every packet over two links is then the initial packet transmission time + 33 additional packet times, and so the average time is $E[T_{total}] * 34 = 0.522$ seconds.

19. Suppose that a Stop-and-Wait ARQ system has a time-out value that is less than the time required to receive an acknowledgment. Sketch the sequence of frame exchanges that transpire between two stations when station A sends five frames to station B and no errors occur during transmission.

Solution:



20. The Trivial File Transfer Protocol (RFC 1350) is an application layer protocol that uses the Stop-and-Wait protocol. To transfer a file from a server to a client, the server breaks the file into blocks of 512 bytes and sends these blocks to the client using Stop-and-Wait ARQ. Find the efficiency in transmitting a 1 MB file over a 10 Mbps Ethernet LAN that has a diameter of 300 meters. Assume the transmissions are error free and that each packet has 60 bytes of header attached.

Solution:

The propagation delay in an Ethernet LAN is negligible compared to the total transmission time of a packet from start to finish. Ignoring processing time and using the terminology in the chapter, we have:

$$t_o = t_f + t_{ack} = \frac{8(512 + 60)}{10 \times 10^6} + \frac{64}{10 \times 10^6} = 4.64 \times 10^{-4}$$

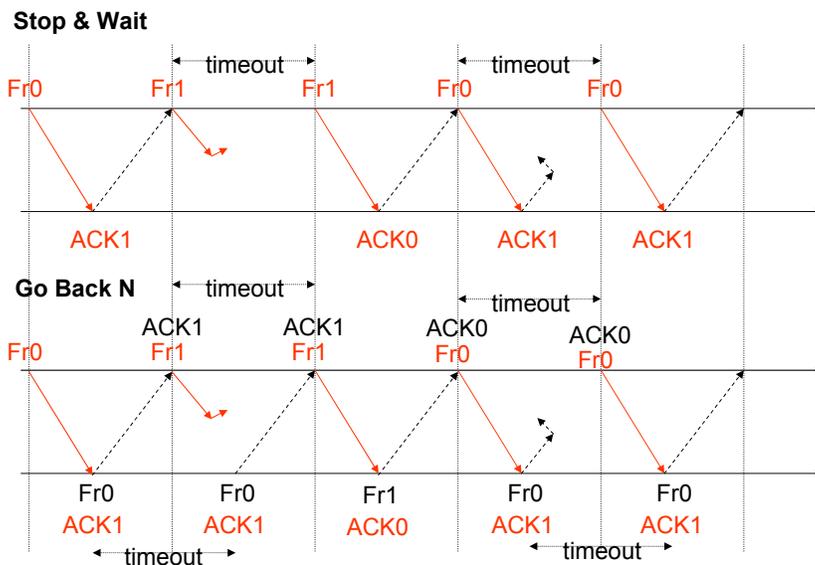
$$\eta_o = \frac{R_{eff}^0}{R} = \frac{\frac{n_f - n_o}{t_o}}{R} = \frac{8 \times 512}{4.64 \times 10^{-4} \times 10 \times 10^6} = 0.8828 = 88.3\%$$

One more source of overhead occurs because the last packet is not full. However, this additional overhead accounts for a very small fraction of the total overhead and does not affect the above result.

21. Compare the operation of Stop-and-Wait ARQ with bidirectional Go-Back-N ARQ with a window size of 1. Sketch out a sequence of frame exchanges using each of these protocols and observe how the protocols react to the loss of an information frame and to the loss of an acknowledgment frame.

Solution:

The figure below shows that the bidirectional Go-Back-N ARQ recovers from errors in the same time that Stop-and-Wait:



22. Consider the various combinations of communication channels with bit rates of 1 Mbps, 10 Mbps, 100 Mbps, and 1 Gbps over links that have roundtrip times of 10 msec, 1 msec, and 100 msec.

Solutions follow questions:

- a. Find the delay-bandwidth product for each of the 12 combinations of speed and distance.

Delay-bandwidth (Megabits)

Bit Rate Mbps	Round Trip Time (msec)		
	100	10	1
1	0.1	0.01	0.001
10	1.0	0.10	0.010
100	10.0	1.00	0.100
1000	100.0	10.00	1.000

- b. Suppose 32-bit sequence numbers are used to transmit blocks of 1000 bytes over the above channels. How long does it take for the sequence numbers to wrap around, that is, to go from 0 up to 2^m ?

Block 1000 bytes
 Sequence 32 bits 4294967296

Time for the sequence number to wrap around (sec)
 $\text{Frame Time} * 2^{32} = 4294967296 * 8 * 1000 / R$ seconds

Bit Rate Mbps	Round Trip Time (msec)		
	100	10	1
1	34359738.37	34359738.37	34359738.37
10	3435973.84	3435973.84	3435973.84
100	343597.38	343597.38	343597.38
1000	34359.74	34359.74	34359.74

- c. Now suppose the 32-bit sequence numbers are used to count individual transmitted bytes. How long does it take for the sequence numbers to wrap around?

Time for the sequence number to wrap around (sec)
 $\text{Byte Time} * 2^{32} = 4294967296 * 8 / R$ seconds

Bit Rate Mbps	Round Trip Time (msec)		
	100	10	1
1	34359.74	34359.74	34359.74
10	3435.97	3435.97	3435.97
100	343.60	343.60	343.60
1000	34.36	34.36	34.36

The sequence number wraps around in much shorter time. At 1 Gbps the sequence number wraps around in only 34.36 seconds.

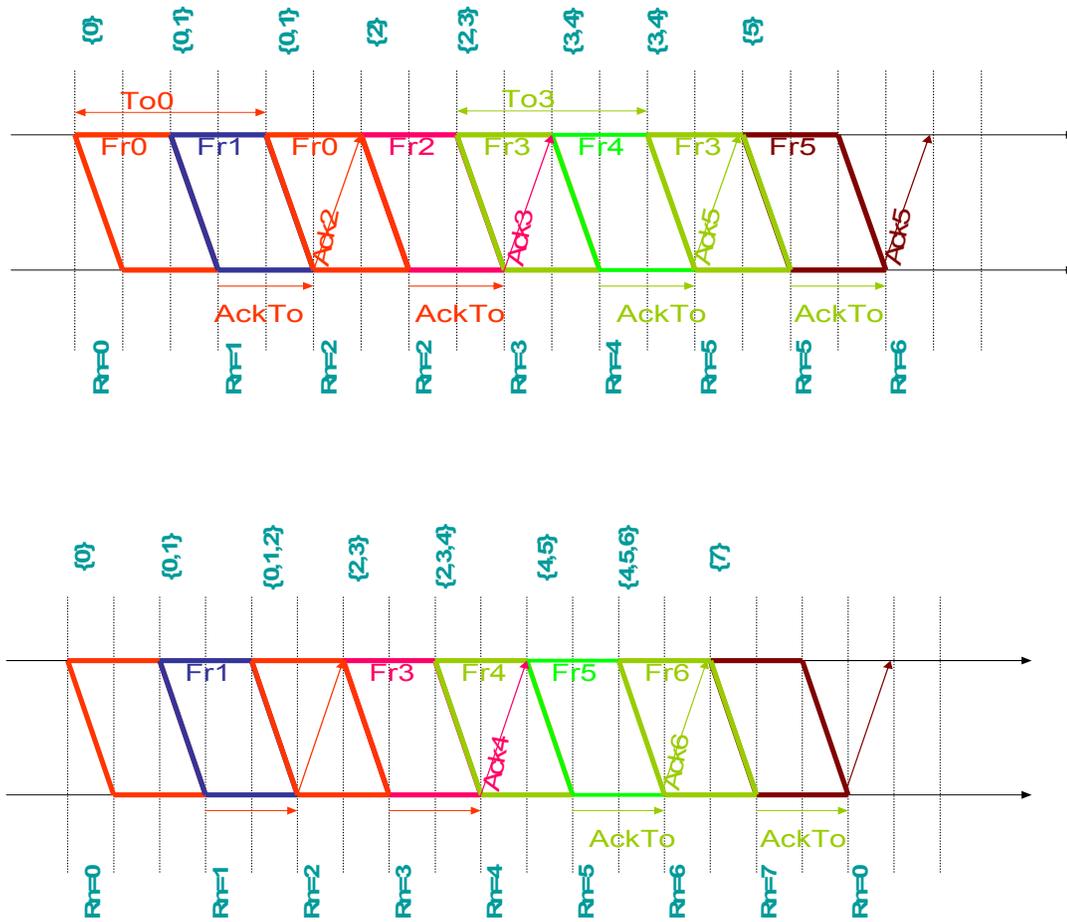
23. Consider a bidirectional link that uses Go-Back-N with $N = 7$. Suppose that all frames are one unit long and that they use a time-out value of 2. Assume the propagation is 0.5 unit and the processing time is negligible. Assume the ACK timer is one unit long. Assuming stations A and B begin with their sequence numbers set to zero, show the pattern of transmissions and associated state transitions for the following sequences of events:

Solutions follow questions:

Go-Back-N with $N = 7$, $t_{timeout} = 2$, $t_{prop} = 0.5$, $t_{ACK} = 1$, $t_f = 1$. Each tick represents one half of a unit of time. The ACK timer delays the sending of an ACK to provide an opportunity for piggybacking.

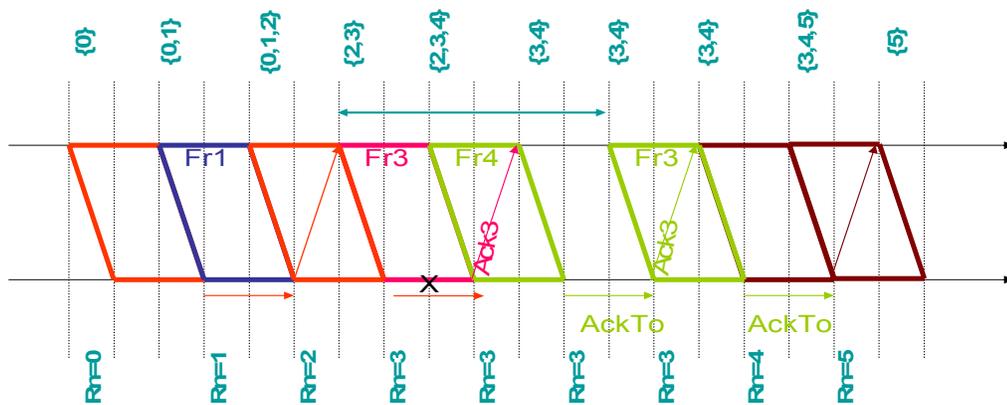
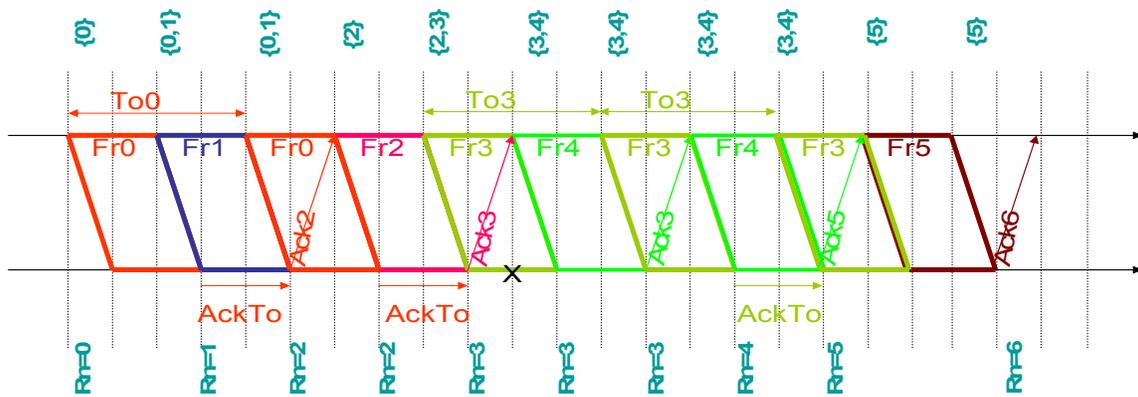
- a. Station A sends six frames in a row, starting at $t = 0$. All frames are received correctly.

The first figure below shows that the timeout value is too short and cause needless retransmissions. The second figure below shows that by increasing the timeout value to 3 units, transmissions occur smoothly and efficiently.



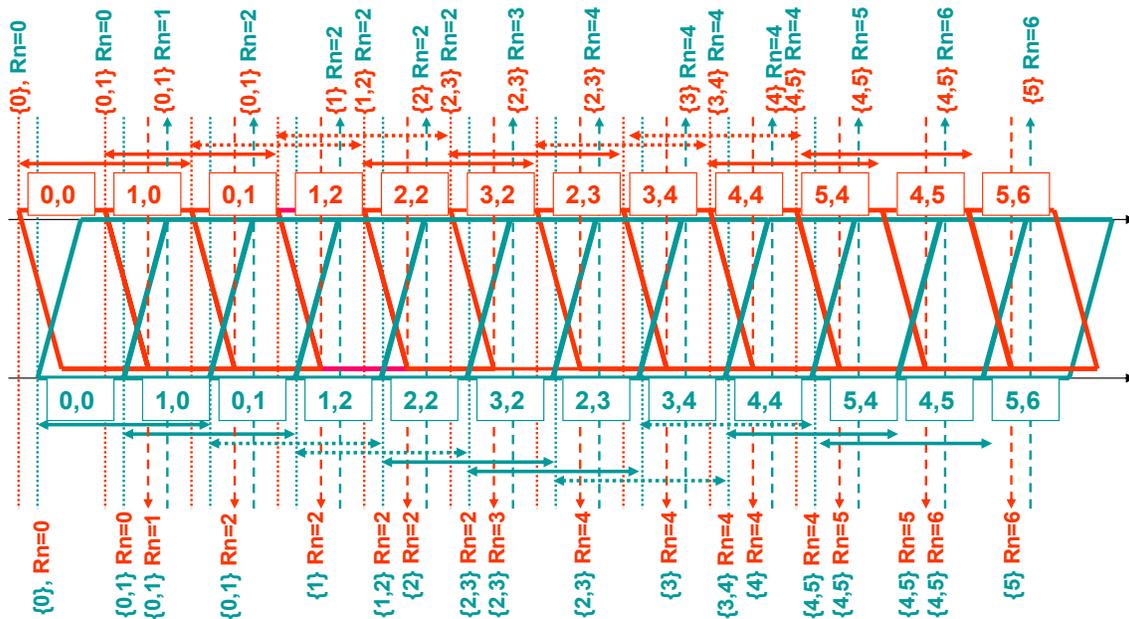
- b. Station A sends six frame in a row, starting at $t = 0$. All frames are received correctly, except frame 3 is lost.

The following two figures show the sequence of frame exchanges with timeout values of 2 and 3 units. The system with the shorter timeout value initiates the recovery from error sooner, but cannot advance the window fast enough due to the short timeout value.



- c. Station A sends six frames in a row, starting at $t = 0$. Station B sends six frames in a row starting at $t = 0.25$. All frames are received correctly.

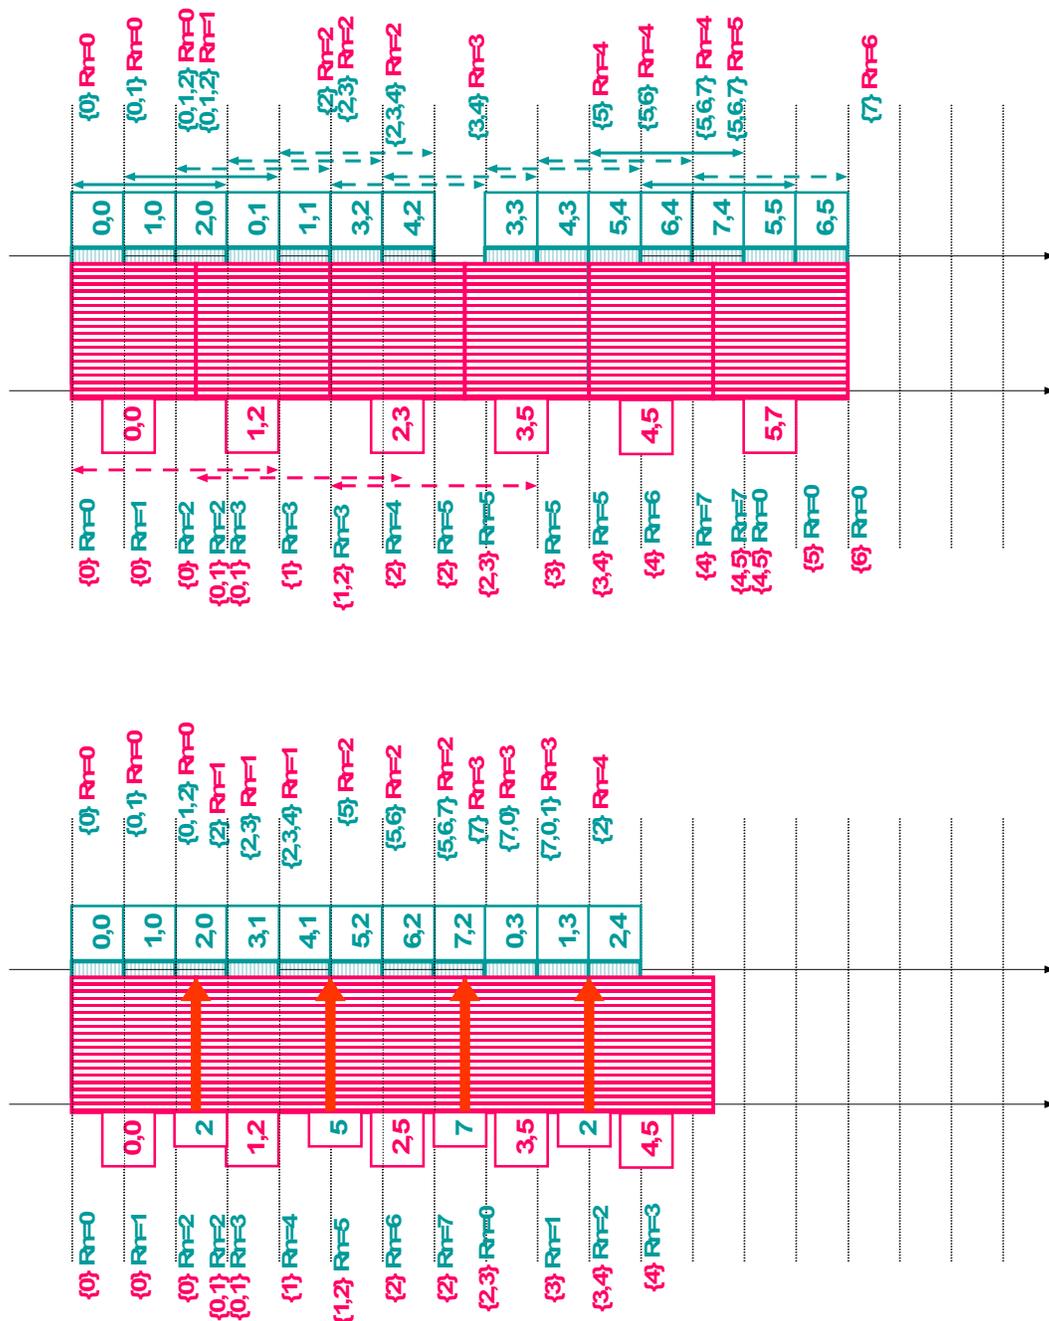
This problem shows that when piggybacking is used the timeout has to be somewhat longer. In the following figure we see that timeouts occur repeatedly causing needless retransmission of frames that have already arrived at the receiver. A timeout value of 4 is required to assure a smooth sliding forward of the transmitter's send window.



24. Consider a bidirectional link that uses Go-Back-N with $N = 3$. Suppose that frames from station A to station B are one unit long and that they use a time-out value of 2, and that frames in the opposite directions are 2.5 units long and that they use a time-out value of 4. Assume propagation and processing times are negligible, that the stations have an unlimited number of frames ready for transmission, and that all ACKs are piggybacked onto information frames. Assuming stations A and B begin with their sequence numbers set to zero, show the pattern of transmissions and associated state transitions that result if there are no transmission errors.

Solution:

The following shows the sequence of frames exchanged by stations A and B, as well as the set of unacknowledged frames, and the next expected frame at each station. Note that acknowledgments from station A to station B are frequent because of the short frame length and so station B sends frames in a continuous efficient manner. This is not the case for station A which receives acknowledgments after large delay and so frequently spends its time carrying out needless retransmissions. The solution to this problem is to insert ACK messages from station B to station A in between the long message transmissions. The second figure below shows that transmissions from A to B then flow smoothly.



25. Consider the Go-Back-N ARQ protocol.

Solutions follow questions:

a. What can go wrong if the ACK timer is not used?

When no traffic arrives at a receiver during bidirectional Go-Back-N ARQ, and the receiver has to send an ACK, it usually sends the ACK after the ACK timer expires. If the ACK timer is not used, there are only two options remaining:

- 1) The ACK must be sent immediately (that is, use piggybacking only if frame already available)

Although this will function correctly, it is an inefficient use of bandwidth in the general case.

- 2) The ACK must only be sent if it can be piggybacked

This is problematic if traffic arrives sporadically. The sender will wait a long time until a piggyback opportunity arises.

- b. Show how the frame timers can be maintained as an ordered list where the time-out instant of each frame is stated relative to the time-out value of the previous frame.

Assume that the timer counts down from $t_{timeout}$. In order to have a separate timer for each frame, we need not implement N timers. Only the oldest frame can timeout. The system can save, for each frame, an arrival offset time that is related to the frame that preceded it and place these offsets in an ordered list based on the frame sequence numbers. If an ACK for the oldest frame arrives, the system simply increments the timer by the offset of the following frame in the list. If an ACK for any other frame arrives, the timer is incremented by the sum of all of the offsets in the list that are up to this newly acknowledged frame.

- c. What changes if each frame is acknowledged individually instead of by using a cumulative acknowledgment (R_{next} acknowledges all frames up to $R_{next} - 1$)?

If each frame needs to be acknowledged individually, then the number of ACK messages will increase and the rate at which the transmission window can be increased will be reduced.

26. Suppose that instead of Go-Back-N ARQ, N simultaneous Stop-and-Wait ARQ processes are run in parallel over the same transmission channel. Each SDU is assigned to one of the N processes that is currently idle. The processes that have frames to send take turns transmitting in round-robin fashion. The frames carry the binary send sequence number as well as an ID identifying which ARQ process it belongs to. Acknowledgments for *all* ARQ processes are piggybacked onto *every* frame.

Solutions follow questions:

- a. Qualitatively, compare the relative performance of this protocol with Go-Back-N ARQ and with Stop-and-Wait ARQ.

For simplicity assume that the time between consecutive frame transmissions in Stop-and-Wait corresponds to N consecutive transmissions without stopping. The parallel Stop-and-Wait procedure described above is an effective way to fill the transmission pipe without the additional complexity of Go-Back-N ARQ.

Vs. Go-Back N. Go-Back-N delivers frames in order. The parallel Stop-and-Wait protocol does not deliver frames in order, so additional processing is required if frames must be delivered in sequence. Because all the processes are independent, this protocol retransmits erroneous frames individually. In contrast, the Go-Back-N protocol retransmits a group of N frames. In this sense, the parallel protocol seems to perform similarly to a Selective Repeat process.

Vs. Stop-And-Wait. If Stop-and-Wait is used, the effective bit rate, without errors, will be N times less than the protocol described here. In fact, the larger N , the more efficient the protocol described is. At its worst case, where $N = 1$, it reduces to Stop-And-Wait.

- b. How does the service offered by this protocol differ from that of Go-Back-N ARQ?

The parallel Stop-and-Wait protocol does not deliver frames in order, unless augmented by a frame resequencing scheme.

27. Write a program for the transmitter and the receiver implementing Go-Back-N ARQ over a data link that can introduce errors in transmission.

Solutions follow questions:

- a. Identify what variables need to be maintained.

R_{next} : frame expected by receiver

S_{last} : oldest outstanding frame (back of window) at the transmitter

S_{recent} : most recently transmitted frame

- b. The program loops continuously waiting for an event to occur that requires some action to take place. Identify the main events that can occur in the transmitter. Identify the main events that can occur in the receiver.

Go Back N Transmitter			
Event	Condition	Action	Next State
Request from upper layer	Not the last sequence number in nonempty send window	Prepare frame with S_{recent} and send, increment send sequence number; start timer	Ready
Request from upper layer	Last sequence number in nonempty send window	Prepare frame and send, increment send sequence number; start timer	Blocking
Arrival of error-free ACK frame	Correct sequence number, R_{next} is between S_{last} and S_{recent}	Slide window forward, $S_{last} = R_{next}$, max Sequence number = $S_{last} + W_S - 1$	Ready
Arrival of error-free ACK frame	Sequence number R_{next} is NOT between S_{last} and S_{recent}	Discard frame	Ready
Timeout Expires		Resend all frames from S_{last} onwards, reset timer	Ready
Arrival of erroneous frame		Discard frame	Ready
Request from upper layer		Requests are blocked	Blocking
Arrival of error-free ACK frame	Correct sequence number, R_{next} is between S_{last} and $S_{last} + W_S - 1$	Slide window forward, $S_{last} = R_{next}$, max Sequence number = $S_{last} + W_S - 1$	Ready
Arrival of error-free ACK frame	Sequence number R_{next} is NOT between S_{last} and S_{recent}	Discard frame	Blocking
Timeout Expires		Resend all frames from S_{last} onwards, reset timer	Blocking
Arrival of erroneous frame		Discard frame	Blocking

Go Back N Receiver				
Current State	Event	Sequence Number	Action	Next State
Ready	Arrival of error-free frame	Expected sequence number $S_{last} = R_{next}$	Accept frame, increment R_{next} , $R_{next} = R_{next} + 1$, and send ACK, deliver packet to higher layer	Ready
Ready	Arrival of error-free frame	Incorrect sequence number	Discard frame; send ACK with R_{next}	Ready
Ready	Arrival of erroneous frame		Discard frame	Ready

28. Modify the program in problem 5.27 to implement Selective Repeat ARQ.

Solution:

Selective Repeat Transmitter				
Current State	Event	Sequence Number	Action	Next State
Ready	Request from upper layer	Not the last sequence number in nonempty send window	Prepare frame with S_{recent} and send, increment send sequence number; start timer	Ready
Ready	Request from upper layer	Last sequence number in nonempty send window	Prepare frame and send, increment send sequence number; start timer	Blocking
Ready	Arrival of error-free ACK frame	Correct sequence number, R_{next} is between S_{last} and S_{recent}	Slide window forward, $S_{\text{last}} = R_{\text{next}}$, max sequence number = $S_{\text{last}} + W_S - 1$	Ready
Ready	Arrival of error-free NAK frame	Correct sequence number, R_{next} is between S_{last} and S_{recent}	Slide window forward, $S_{\text{last}} = R_{\text{next}}$, max Sequence number = $S_{\text{last}} + W_S - 1$, retransmit frame with R_{next}	Ready
Ready	Arrival of error-free ACK/NAK frame	Sequence number R_{next} is NOT between S_{last} and S_{recent}	Discard frame	Ready
Ready	Timeout Expires		Resend frame corresponding to timer, reset associated timer	Ready
Ready	Arrival of erroneous frame		Discard frame	Ready
Blocking	Request from upper layer		Requests are blocked	Blocking
Blocking	Arrival of error-free ACK frame	Correct sequence number, R_{next} is between S_{last} and $S_{\text{last}} + W_S - 1$	Slide window forward, $S_{\text{last}} = R_{\text{next}}$, max sequence number = $S_{\text{last}} + W_S - 1$	Ready
Blocking	Arrival of error-free NAK frame	Correct sequence number, R_{next} is between S_{last} and S_{recent}	Slide window forward, $S_{\text{last}} = R_{\text{next}}$, max sequence number = $S_{\text{last}} + W_S - 1$, retransmit frame with R_{next}	Ready
Blocking	Arrival of error-free NAK frame	Correct sequence number, $R_{\text{next}} = S_{\text{last}}$	retransmit frame with R_{next}	Blocking
Blocking	Arrival of error-free ACK/NAK frame	Sequence number R_{next} is NOT between S_{last} and S_{recent}	Discard frame	Blocking
Blocking	Timeout expires		Resend frame corresponding to timer, reset associated timer	Blocking
Blocking	Arrival of erroneous frame		Discard frame	Blocking

Selective Repeat Receiver				
Current State	Event	Sequence Number	Action	Next State
Ready	Arrival of error-free frame	Sequence number is R_{next}	Accept frame; slide window forward $R_{\text{next}} = R_{\text{next}} + k$, and send ACK, deliver k packets to higher layer	Ready
Ready	Arrival of error-free frame	Sequence number is in the receive window, that is, between $R_{\text{next}} + 1$ and $R_{\text{next}} + W_R - 1$	Accept and buffer frame; and send ACK	Ready
Ready	Arrival of error-free frame	Sequence number is outside the receive window,	Discard frame, send ACK with R_{next}	Ready
Ready	Arrival of erroneous frame		Discard frame	Ready

29. Three possible strategies for sending ACK frames in a Go-Back-N setting are as follows: send an ACK frame immediately after each frame is received, send an ACK frame after every other frame is received, and send an ACK frame when the next piggyback opportunity arises. Which of these strategies are appropriate for the following situations?

Solutions follow questions:

- a. An interactive application produces a packet to send each keystroke from the client; the server echoes each keystroke that it receives from the client.

Since each keystroke is echoed, there will always be a piggyback opportunity. Thus, the piggyback method should be used. Indeed, the echo packet constitutes an acknowledgment.

- b. A bulk data transfer application where a server sends a large file that is segmented in a number of full-size packets that are to be transferred to the client.

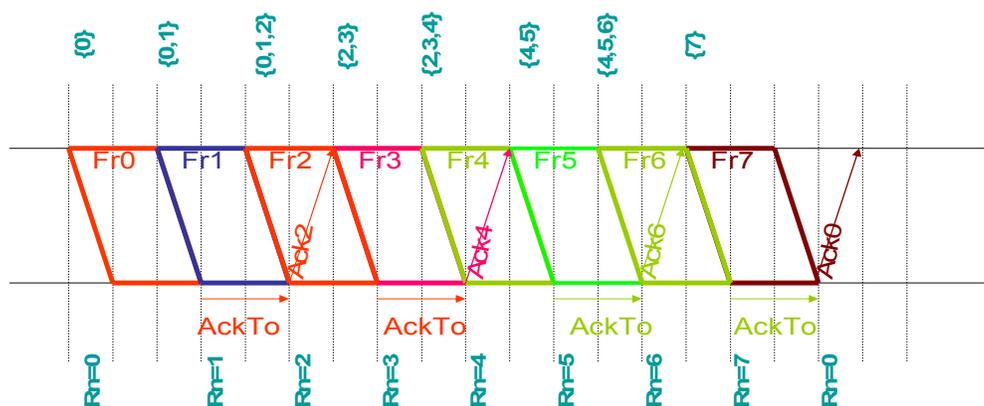
The upstream traffic to a server is generally much less than the downstream traffic. Thus, the piggybacking method is non-ideal in this case. If the channel has a low probability of error, the alternating ACK method is better, as it saves bandwidth. However, if the connection causes frequent errors, every frame should be acknowledged. Additional overhead traffic is caused by the ACK frames, but will be compensated by the bandwidth savings that will arise when the errors are discovered more quickly.

30. Consider a bidirectional link that uses Selective Repeat ARQ with a window size of $N = 4$. Suppose that all frames are one unit long and use a time-out value of 2. Assume that the one-way propagation delay is 0.5 time unit, the processing times are negligible, and the ACK timer is one unit long. Assuming station A and B begin with their sequence numbers set to zero, show the pattern of transmissions and associated state transitions for the following sequences of events:

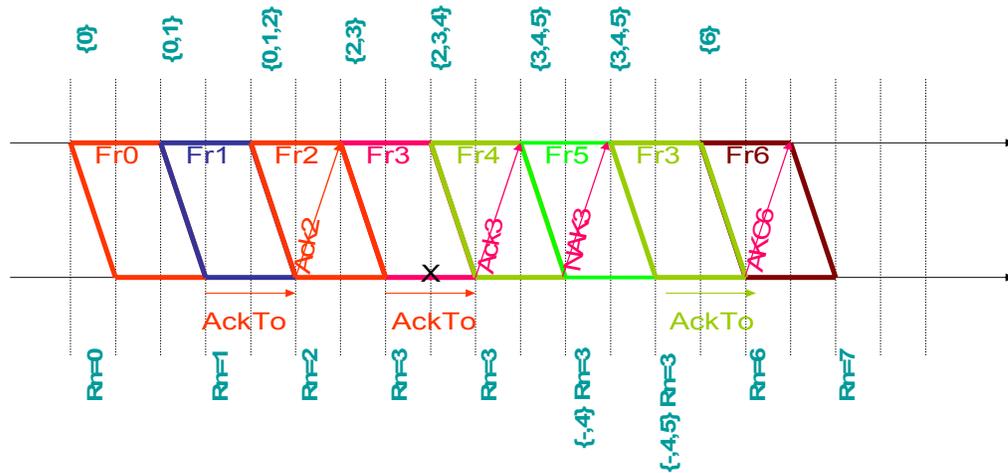
Solutions follow questions:

- a. Station A sends six frames in a row, starting at $t = 0$. All frames are received correctly.

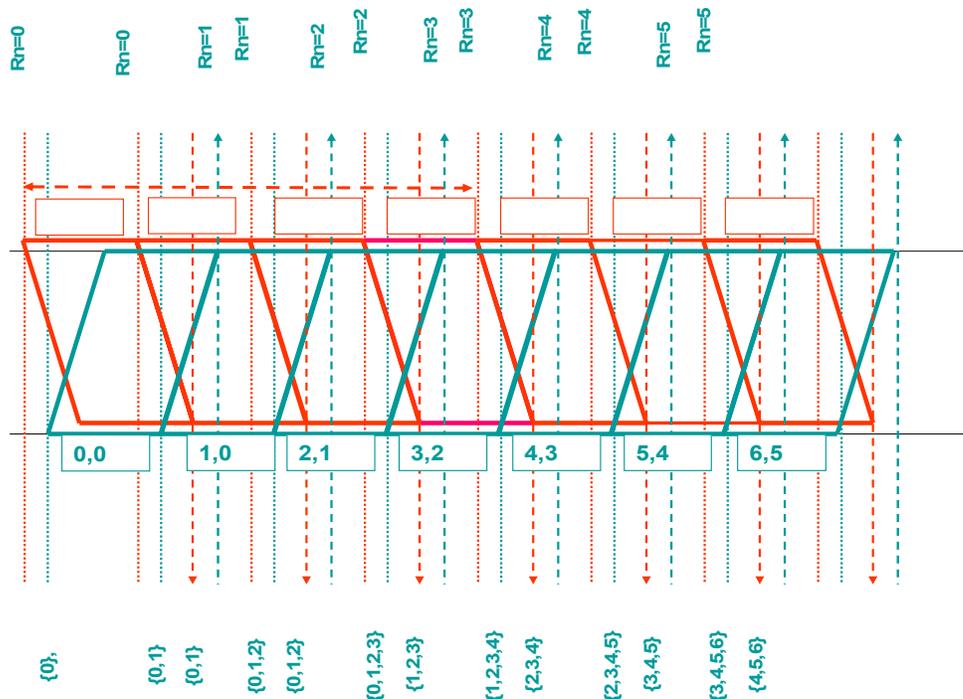
A timeout value of 2 causes the transmitter to resend frames before the window size of 4 is exhausted. We will assume a timeout value of 4 in the following solutions.



- b. Station A sends six frame in a row, starting at $t = 0$. All frames are received correctly, except frame 3 is lost.



- c. Station A sends six frames in a row, starting at $t = 0$. Station B sends six frames in a row starting at $t = 0.25$. All frames are received correctly.



31. In the chapter we showed that if the transmit and receive maximum window sizes are both equal to half the available sequence number space, then Selective Repeat ARQ will work correctly. Rework the arguments presented in the chapter to show that if the sum of the transmit and receive maximum window sizes equals the available sequence number space, then Selective Repeat ARQ will work correctly.

Solution:

Assume $W_s + W_R = 2^m$ and assume that the current send window is 0 to $W_s - 1$. Suppose also that the receive window is 0 to $W_R - 1$. Now suppose that frame 0 is received correctly but that the

acknowledgment for frame 0 is lost. The transmitter can transmit new frames only up to frame $W_S - 1$. Depending on which transmissions arrive without error, R_{next} will be in the range between 1 and W_S while $R_{next} + W_R - 1$ will be in the range of 1 to $W_R + W_S - 1$. The maximum value of R_{next} occurs when frames 0 through $W_S - 1$ are received correctly, in which case the value of R_{next} is W_S and the value of $R_{next} + W_R - 1$ increases to $W_R + W_S - 1$. Crucially, the receiver will not receive frame $W_R + W_S$ until the acknowledgment for frame 0 has been received at the transmitter. Any receipt of frame 0 prior to frame $W_R + W_S$ indicates a duplicate transmission of frame 0. Therefore, the sum of the maximum window sizes is 2^m .

32. Suppose that Selective Repeat ARQ is modified so that ACK messages contain a list of the next m frames that it expects to receive.

Solutions follow questions:

a. How does the protocol need to be modified to accommodate this change?

First, the frame header needs to be modified to accommodate the list of frames to receive. It can be a fixed or a variable number of slots. NAK won't be necessary because the receiver explicitly indicates which frames need to be transmitted. Second, the transmitter operation must change to retransmit frames according to the received list. If the list contains the m oldest frames that are yet to be received, then the list can be used to skip retransmissions of frames that have already been received.

b. What is the effect of the change on protocol performance?

The performance will increase in cases with high error rate or in cases where the delay is high. A single frame can ask for the retransmission of several frames. The drawback is the overhead in the header and the increased protocol complexity relative to pure Selective-Repeat ARQ.

33. A telephone modem is used to connect a personal computer to a host computer. The speed of the modem is 56 kbps and the one-way propagation delay is 100 ms.

Solutions follow questions:

a. Find the efficiency for Stop-and-Wait ARQ if the frame size is 256 bytes; 512 bytes. Assume a bit error rate of 10^{-4} .

First we have the following:

$$P_f = 1 - (1 - 10^{-4})^{n_f}$$

$$n_f = 256 \times 8 = 2048 \quad \text{or} \quad n_f = 512 \times 8 = 4096$$

$$t_{prop} = 100 \text{ ms}$$

$$n_o = 0$$

$$n_a = 64 \text{ bits}$$

$$t_{proc} = 0$$

Using the results in Equation 5.4,

$$\eta = (1 - P_f) \frac{1 - \frac{n_o}{n_f}}{1 + \frac{n_a}{n_f} + \frac{2(t_{prop} + t_{proc})}{n_f} R}$$

$$= 0.125 \quad (n_f = 2048)$$

$$= 0.177 \quad (n_f = 4096)$$

- b. Find the efficiency of Go-Back-N if three-bit sequence numbering is used with frame sizes of 256 bytes; 512 bytes. Assume a bit error rate of 10^{-4} .

Given that $W_S = 2^3 - 1 = 7$, we can calculate that the window size is:

$$\frac{n_f \times W_S}{R} = 256ms$$

Since this is greater than the round trip propagation delay, we can calculate the efficiency by using the results in Equation 5.8.

$$\begin{aligned} \eta &= (1 - P_f) \frac{1 - \frac{n_o}{n_f}}{1 + (W_S - 1)P_f} \\ &= 0.385 \quad (n_f = 2048) \\ &= 0.220 \quad (n_f = 4096) \end{aligned}$$

34. A communications link provides 1 Mbps for communications between the earth and the moon. The link is used to send color images from the moon. Each image consists of $10,000 \times 10,000$ pixels, and 16 bits are used for each of the three-color components of each pixel.

Solutions follow questions:

- a. How many images per second can be transmitted over the link?

The number of images that can be transmitted per second is:

$$1 \times 10^6 \frac{\text{bits}}{\text{sec}} \Big/ (10000^2 \times 16 \times 3 \frac{\text{bits}}{\text{image}}) = 2.1 \times 10^{-4} \text{ images / second}$$

- b. If each image is transmitted as a single block, how long does it take to get an acknowledgment back from earth? The distance between earth and the moon is approximately 375,000 km.

The total time to get an acknowledgment from earth, assuming that $t_{ACK} \ll t_f$ is:

$$\begin{aligned} t_O = t_f + 2t_{prop} &= \frac{(10000^2 \times 16 \times 3 \frac{\text{bits}}{\text{image}})}{1 \times 10^6 \frac{\text{bits}}{\text{sec}}} + 2 \frac{375000 \times 10^3 \text{ m}}{3 \times 10^8 \frac{\text{m}}{\text{sec}}} \\ &= 4800 + 2.50 = 4802.5 \text{ sec / image} \end{aligned}$$

Note that if each image is transmitted in a single block, t_{prop} becomes insignificant compared to t_f .

- c. Suppose that the bit error rate is 10^{-5} , compare Go-Back-N and Selective Repeat ARQ in terms of their ability to provide reliable transfer of these images from the moon to earth. Optimize your frame size for each ARQ protocol.

The difficulty here is that if we transmit an entire image as a single block, then the average number of errors in a block is $48 \times 10^8 \times 10^{-5} = 48000$. Thus every transmission fails. Clearly we need to transmit using a smaller frame size, say n . Qualitatively, the following happens as we vary n . If n is very small, then the probability of frame error is small, but the overhead due to the header n_0 will become significant. If n becomes too large, then the efficiency drops because of frequent retransmissions. Thus there must be an intermediate value of n that optimizes efficiency. We will find this value by trial and error. We will assume a header overhead of 64 bits in a frame of size n greater than 64 bits.

Each frame carries n bits of information, so the required window size for a propagation delay of 2.5 seconds is then $W_s = 2.5 \times 10^6 / n + 1$. The probability of frame error is then

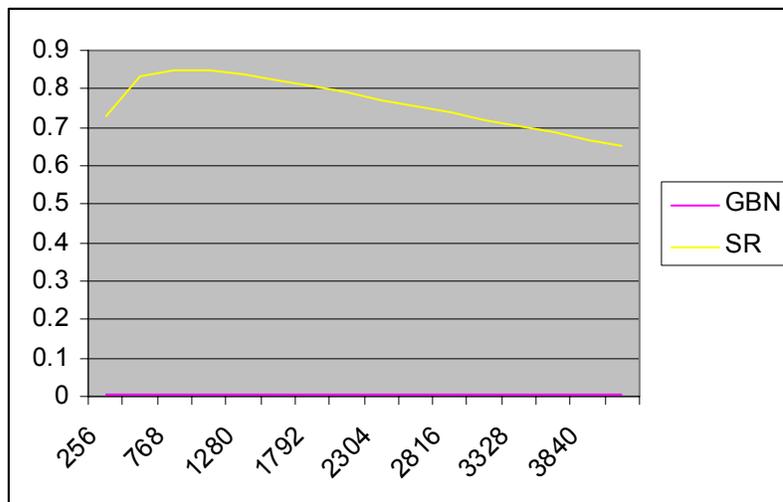
$$P_f = 1 - (1 - 10^{-5})^n \approx 1 - e^{-n10^{-5}}$$

The efficiency for Go-Back-N is given by:

$$\eta_{GBN} = \frac{(1 - P_f)(1 - \frac{64}{n})}{1 + (W_s - 1)P_f} = \frac{(1 - \frac{64}{n})e^{-n10^{-5}}}{1 + (\frac{2.5 \times 10^6}{n})(1 - e^{-n10^{-5}})}$$

Using similar assumptions, the efficiency of Selective Repeat ARQ is given by

$$\eta_{SR} = (1 - P_f)(1 - \frac{64}{n}) = (1 - \frac{64}{n})e^{-n10^{-5}}$$



The figure above shows the efficiency of Go-Back-N and Selective Repeat ARQ as a function of frame size. Go-Back-N has very low efficiency (always below 1%) for all values of n . Selective Repeat achieves a maximum of about 85% efficiency at around $n=768$ bits.

Note that the optimum value of n is dependent on the value of n_0 .

35. Two computers are connected by an intercontinental link with a one-way propagation delay of 100 ms. The computers exchange 1-Megabyte files that they need delivered in 250 ms or less. The transmission lines have a speed of R Mbps and the bit error rate is 10^{-8} . Design a transmission system by selecting the bit rate R , the ARQ protocol, and the frame size.

Solution:

Using Selective Repeat ARQ for efficiency's sake and assuming that the overhead $n_o = 64$ bits, we need to select n_f and R to ensure delivery in 250 ms or less.

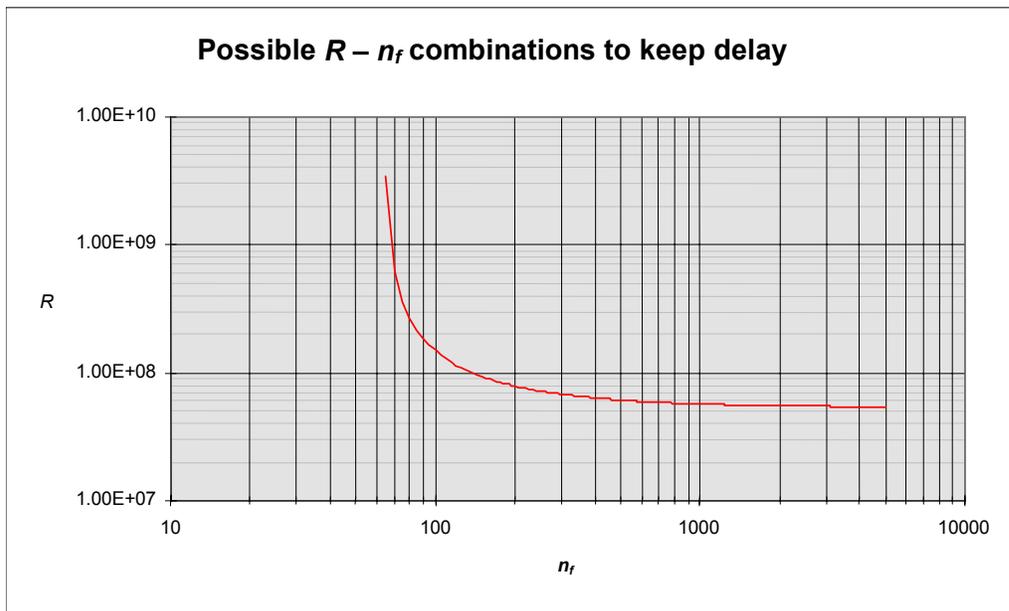
$$\text{Total delay} = 2t_{prop} + \frac{8 \times 10^6}{R_{eff}} = 200 \times 10^{-3} + \frac{8 \times 10^6}{\eta R} < 250 \times 10^{-3}$$

$$\eta R > 6.25 \times 10^{-9}$$

From Equation 5.11 we see that

$$(1 - 10^{-8})^{n_f} \left(1 - \frac{64}{n_f}\right) R > 53.33 \times 10^6$$

The allowable region is shown below as the area under the graph.



36. Find the optimum frame length n_f that maximizes transmission efficiency for a channel with random bit errors by taking the derivative and setting it to zero for the following protocols:

Solutions follow questions:

- a. Stop-and-Wait ARQ.

$$\eta = (1 - P_b)^{n_f} \frac{n_f - n_o}{n_f + n_a + 2(t_{proc} + t_{prop})R}$$

Let $a = (1 - P_b)$

Let $b = 2(t_{proc} + t_{prop})R + n_a$

$$\frac{d}{dn_f} \left[a^{n_f} \frac{n_f - n_o}{n_f + b} \right] = \frac{a^{n_f} [\ln a(n_f - n_o) + 1](n_f + b) - a^{n_f} (n_f - n_o)}{(n_f + b)^2} = 0$$

$$a^{n_f} \{ \ln a [n_f^2 + (b - n_o)n_f - n_o b] + n_f + b - n_f + n_o \} = 0$$

$$[n_f^2 + (b - n_o)n_f - n_o b] + \frac{b + n_o}{\ln a} = 0$$

$$n_f = \frac{(n_o - b) \pm \sqrt{(b - n_o)^2 - 4 \left[\frac{b + n_o}{\ln a} - bn_o \right]}}{2}$$

b. Go-Back-N ARQ.

$$\eta = (1 - P_b)^{n_f} \frac{n_f - n_o}{n_f + n_f \left[1 - (1 - P_b)^{n_f} \right] \left[\frac{2R(t_{proc} + t_{prop})}{n_f} \right]}$$

Let $a = (1 - P_b)$

Let $b = 2R(t_{proc} + t_{prop})$ and use the approximation $1 - (1 - p)^n \approx np$, (which is valid when $np \ll 1$), then

$$\eta = \frac{a^{n_f} (n_f - n_o)}{n_f + n_f P_b b} = \frac{a^{n_f} (n_f - n_o)}{n_f [1 + P_b b]}$$

Taking the derivative and equating the numerator to zero, we find that:

$$0 = \frac{d\eta}{dn_f} = \frac{d}{dn_f} \left[\frac{a^{n_f} (n_f - n_o)}{n_f} \right] = \frac{(a^{n_f} \ln a (n_f - n_o) + a^{n_f}) n_f - a^{n_f} (n_f - n_o)}{n_f^2}$$

which leads to the quadratic equation

$$n_f^2 - n_o n_f + n_o / \ln a = 0$$

which gives the solution:

$$n_f = \frac{n_o \pm \sqrt{n_o^2 - 4 \frac{n_o}{\ln a}}}{2}$$

c. Selective Repeat ARQ.

Borrowing the result from problem 5.34c:

$$\eta = (1 - P_b)^{n_f} \left(1 - \frac{n_o}{n_f} \right)$$

Except for a scale factor, this equation is the same as the approximation for η found in part (b). The solution for the optimum frame size has the same form.

- d. Find the optimum length for a 1 Mbps channel with 10 ms reaction time, 25-byte overhead, 25-byte ACK frame, and $p = 10^{-4}, 10^{-5},$ and 10^{-6} .

$$R = 1 \times 10^6,$$

$$n_o = 8 \times 25 = 200,$$

$$n_a = 8 \times 25 = 200,$$

$$P_b = 10^{-4}, 10^{-5}, 10^{-6}$$

$$a = 1 - P_b, \ln a = -1 \times 10^{-3}, -1 \times 10^{-4}, -1 \times 10^{-5}$$

$$t_{prop} + t_{proc} = 10 \text{ ms}, \quad b = 2 (t_{proc} + t_{prop}) R + n_a = 2 \times 10^{-2} \times 1 \times 10^6 + 200 = 20200$$

P	10^{-4}	10^{-5}	10^{-6}
Stop and Wait ARQ	1156	7552	36303
Go back N ARQ	558	1517	4573
Selective Repeat	558	1517	4573

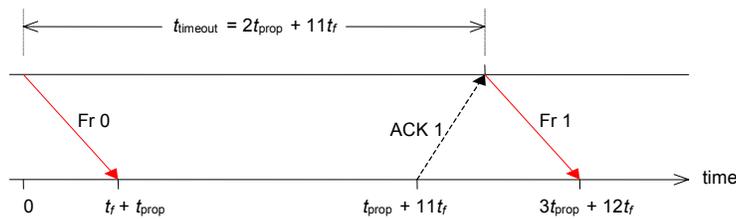
37. Suppose station A sends information to station B on a data link that operates at a speed of 10 Mbps, and that that station B has a 1-Megabit buffer to receive information from A. Suppose that the application at station B reads information from the receive buffer at a rate of 1 Mbps. Assuming that station A has an unlimited amount of information to send, sketch the sequence of transfers on the data link if Stop-and-Wait ARQ is used to prevent buffer overflow at station B. Consider the following cases:

- One-way propagation delay is 1 microsecond.
- One-way propagation delay is 1 ms.
- One-way propagation delay is 100 ms.

Note: The solution to this problem is currently being reviewed and may be revised.

Solution:

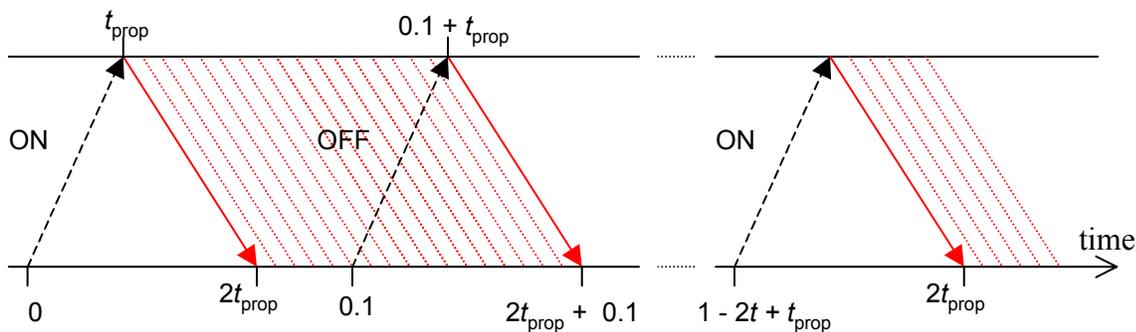
The issue in this problem is that the receiver clears its buffer at 1/10 the rate that the channel can put information into the buffer. Consequently, acknowledgments must be withheld from the transmitter to allow the receiver time to clear the buffer. For the sake of simplicity assume a frame is 1 Mbit long, then it takes 1 Mbit/10 Mbps = 100 ms to transmit the frame, which arrives in its entirety at the receiver at time 100 ms + t_{prop} as shown in the figure below.



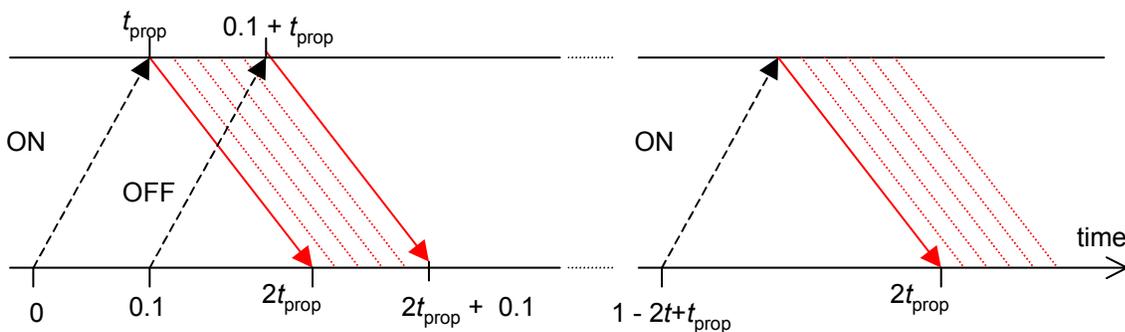
38. Redo problem 5.37 using Xon/Xoff flow control.

Solution:

If t_{prop} is less than 50ms (parts (a) and (b)), the following picture holds.



If t_{prop} is greater than 50ms, the following picture is more indicative of the sequence of events:

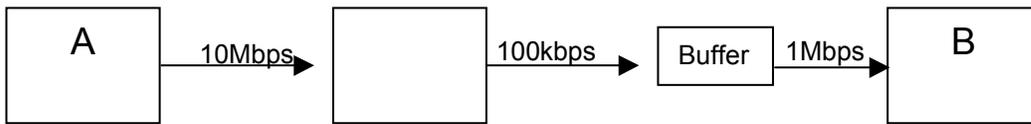


39. Suppose station A sends information to station B over a two-hop path. The data link in the first hop operates at a speed of 10 Mbps, and the data link in the second hop operates at a speed of 100 kbps. Station B has a 1-Megabit buffer to receive information from A, and the application at station B reads information from the receive buffer at a rate of 1 Mbps. Assuming that station A has an unlimited amount of information to send, sketch the sequence of transfers on the data link if Stop-and-Wait ARQ is used on an end-to-end basis to prevent buffer overflow at station B.

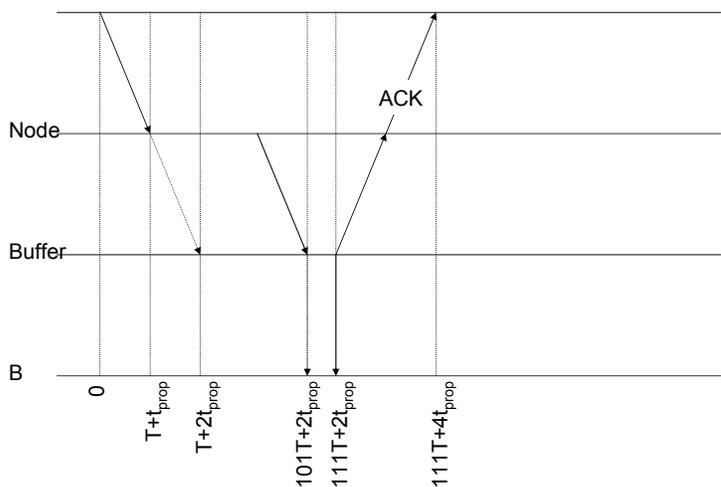
Note: The solution to this problem is currently being reviewed and may be revised.

- a. One-way propagation delay in data link 1 and in data link 2 are 1 ms.
- b. One-way propagation delay in data link 1 and in data link 2 are 100 ms.

Solution:



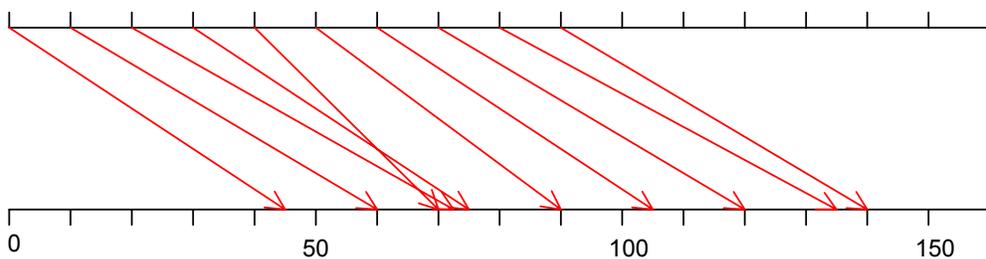
Assume that a frame is n_f bits long, then the frame transmission time in link 1 is $t_1 = n_f/10^7$ seconds, and $t_2 = n_f/10^5$ seconds on the second link as shown in the figure below. Therefore the time to transmit a frame over the second link is 100 times greater than the time to transmit over the first link. For example if $n_f = 10000$ bits, then $T=t_1 = 1$ ms and $t_2 = 100$ ms = $100T$, and similarly if $n_f = 1000000$ bits, then $t_1 = 100$ ms and $t_2 = 10$ seconds. Consequently, the node between link 1 and link 2 must be capable of buffering the bits that arrive in link 1 but which must wait for transmission over link 2. If we assume that the frame is buffered until it has completely arrived, then it will take $t_3 = n_f/10^6 = 10T$ seconds to read out. In the above examples, we have $t_3 = 10$ ms and 1 second respectively. The ACK message for Stop-and-Wait can then be sent. Let $t_f = t_1 + t_2 + t_3 = 111T$. The sequence of events is then as follows: a frame that begins transmission at time 0 will be received in its entirety at node B at time $t_f + 2t_{prop}$; an ACK message is sent thereafter, and the message arrives at the transmitter at approximately time $t_f + 4t_{prop}$.



40. A sequence of fixed-length packets carrying digital audio signal is transmitted over a packet network. A packet is produced every 10 ms. The transfer delays incurred by the first 10 packets are 45 ms, 50 ms, 53 ms, 46 ms, 30 ms, 40 ms, 46 ms, 49 ms, 55 ms, 51 ms.

Solutions follow questions:

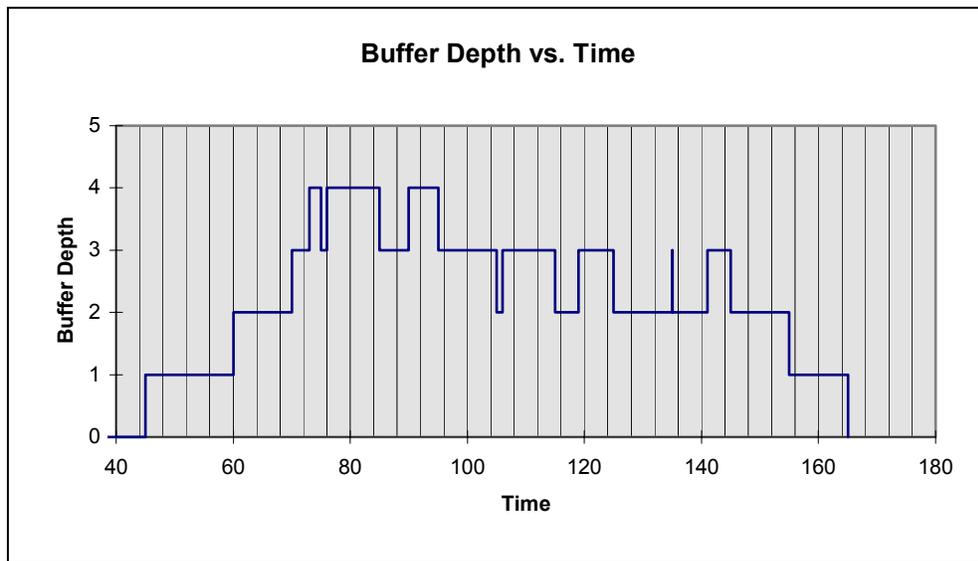
a. Sketch the sequence of packet transmission times and packet arrival times.



b. Find the delay that is inserted at the receiver to produce a fixed end-to-end delay of 75 ms.

Packet #	0	1	2	3	4	5	6	7	8	9
Transmission Time	0	10	20	30	40	50	60	70	80	90
Delay	45	50	53	46	30	40	46	49	55	51
Arrival Time	45	60	73	76	70	90	106	119	135	141
Desired playout time	75	85	95	105	115	125	135	145	155	165
Delay Inserted	30	25	22	29	45	35	29	26	20	24

c. Sketch the contents of the buffer at the receiver as a function of time.



41. Consider an application in which information that is generated at a constant rate is transferred over a packet network so timing recovery is required at the receiver.

Solutions follow questions:

a. What is the relationship between the maximum acceptable delay and the playout buffer?

There are several components in the delay experienced by the information: (1) packetization delay, the time the first byte in a packet payload waits until the packet is full; (2) network transfer delay, the time to traverse the network; and (3) playout delay, the additional delay inserted at the receiver. The sum of these three components should not exceed the maximum acceptable delay. In particular, the length of the playout buffer cannot be longer than the product of the maximum delay and the bit rate of the information.

b. What is the impact of the bit rate of the application information stream on the buffer requirements?

The maximum delay incurred by the buffering is explicitly equal to

$$\text{Max buffer delay} = \text{buffer size (in bits)} / \text{playout rate}$$

Since the playout rate is equal to the bit rate of the application information stream, we see that for a given acceptable maximum playout delay, the buffer size is directly determined by the playout rate.

c. What is the effect of jitter on the buffer size design?

In general, to reduce playback delay, the buffer length should be kept as small as possible. However, the buffer must be able to accommodate the jitter of the incoming traffic stream.

Jitter has the effect of making the incoming data rate variable. If the stream has high jitter (i.e. high delay variations), then the buffer must be made large to accommodate times when the incoming data rate may be higher than the playback rate. In addition the playout delay should ensure that enough packets are buffered so at no time will the buffer become empty and cause “starvation” of the playout system.

42. A speech signal is sampled at a rate of 8000 samples/second. Packets of speech are formed by packing 10 ms worth of samples into each payload. Timestamps are attached to the packets prior to transmission and used to perform error recovery at the receiver. Suppose that the time stamp is obtained by sampling a clock that advances every Δ seconds. Is there a minimum value that is required for Δ ? If so what is it?

Solution:

The clock is sampled every 10 ms to produce a timestamp. Clearly, we require that the clock advance at least once between samples, that is, $\Delta < 10$ ms, in order to ensure a unique timestamp for each packet. In addition if timestamps consist of m bits, we do not want the clock to advance so quickly that the timestamp counter wraps around in one sample time, that is, $2^m \times \Delta \gg 10$ ms.

The timestamps can be used to resequence packets at the receiver. However if the timestamps are to be used for error detection, a Δ of just under 10 ms is not sufficient. To illustrate, suppose that the clock advances every $\Delta = 9$ ms. Then the times at which the clock advances will be as follows:

Clock tick:	0	9	18	27	36	45	54	63	72	81	90
Count:	0	1	2	3	4	5	6	7	8	9	10

The time stamps that result from sampling every 10 ms will be:

Sample:	0	10	20	30	40	50	60	70	80	90	100
Stamp:	0	1	2	3	4	5	6	7	8	10	11

We see that the sample at time 90 will have a time stamp of 10, which is 2 units larger than the time stamp at time 80. If a packet is lost, a gap in the timestamps will be detected by the receiver. However for the example, the receiver may not be able to decide whether a packet was lost or whether the phenomenon above occurred.

While it is true that this problem can be overcome because of the periodic nature of the above phenomenon, a more reliable solution is to place the constraint that $\Delta < 5$ ms. With this constraint in place each timestamp will be offset by 2Δ or possibly 3Δ (in the case of an extra skipped clock pulse). If a packet is lost, a timestamp offset of 4Δ or possibly 5Δ will be detected at the receiver, which is completely unambiguous.

43. Suppose that UDP is used to carry the speech signal in the previous problem. What is the total bit rate consumed by the sequence of UDP PDUs? What is the percent overhead?

Solution:

Number of data bits in each PDU = $64 \times 10^3 \times 10 \times 10^{-3} = 640$ bits

UDP header size = 8×8 bytes = 64 bits

$$\text{Total bit rate} = (640 + 64) / (10 \times 10^{-3}) = 704 \times 10^2 \text{ bps} = 70.4 \text{ Kbps}$$

$$\text{Overhead} = [64 / (640 + 64)] \times 100 = 9.1\%$$

44. Suppose that PDUs contain both timestamps and sequence numbers. Can the timestamps and sequence numbers be combined to provide a larger sequence number space? If so, should the timestamps or the sequence numbers occupy the most significant bit locations?

Solution:

The combination of timestamps and sequence numbers help protect a receiver against sequence number wraparound in very high speed networks. If the transmission speed is sufficiently high, it is possible for a transmitter to re-use a sequence number in a relatively short period of time. The timestamp can be used to distinguish between frames (segments) with the same sequence number. Here is how: the sequence numbers can be viewed as being obtained by looking at the time in a wall clock; the timestamp can correspond to the date. The combination of date and clock time allow the receiver to distinguish between a new segment (with a given sequence number) and an old retransmitted segment (with the same sequence number). In order for this to work, the date should advance at least once in the time it takes for the sequence to wrap around, i.e. the date should advance every 24 hours. This combination of timestamps and sequence numbers is called the PAWS algorithm for "protection against wrapped sequence numbers."

45. Consider the timestamp method for timing recovery discussed in the chapter.

Solutions follow questions:

a. Find an expression that related the difference frequency Δf to the number of cycles M and N .

$$\Delta f = f_n - f_s = f_n \left(1 - \frac{M}{N}\right)$$

b. Explain why only M needs to be sent.

f_n is globally known and N is agreed on beforehand (as a standard or else during connection setup). Thus, only M needs to be sent.

c. Explain how the receiver uses this value of M to control the playout procedure.

The procedure plays out frames at a rate:

$$f_r = f_n - \Delta f = f_n - f_n \left(1 - \frac{M}{N}\right) = \frac{M}{N} f_n$$

46. In Figure 5.32, determine the number of bytes exchanged between client and server in each direction.

Solution:

12 bytes from server to client in frame 4, and 3 bytes from client to server in frame 5. All other exchanges involve only control information in the header.

47. Suppose that the delays experienced by segments traversing the network is equally likely to be any value in the interval [50 ms, 75 ms].

Solutions follow questions:

- a. Find the mean and standard deviation of the delay.

The delay lies between interval [50ms, 75ms] and is a uniform random variable. The mean is:

$$E[X] = (50 + 75) / 2 = 62.5 \text{ ms.}$$

The standard deviation of delay is:

$$STD[X] = VAR[X]^{1/2} = [(75 - 50)^2 / 12]^{1/2} = 7.217$$

- b. Most computer languages have a function for generating uniformly distributed random variables. Use this function in a short program to generate random times in the above interval. Also, calculate t_{RTT} and d_{RTT} and compare to part (a).

See below for sample program written in C.

```

/* Communication Networks - Chapter 5      */
/* Question 47 (b)                        */
/* Description - Generate a random value  */
/* between 50 to 75 ms. Calculate t_RTT   */
/* and d_RTT                              */
/* The min, max and avg value of t_RTT   */
/* and d_RTT are also recorded            */

#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main (void)
{
    int i;
    float temp, t_n, t_rtt_new, t_rtt_old;
    float d_rtt_new, d_rtt_old;
    float t_rtt_min, t_rtt_sum, t_rtt_max;
    float d_rtt_min, d_rtt_sum, d_rtt_max;

    const float alpha = 0.875;
    const float beta = 0.25;

    srand (time(NULL));
    t_rtt_old = 0;
    d_rtt_old = 0;

    t_rtt_sum = t_rtt_max = 0;
    d_rtt_sum = d_rtt_max = 0;

    t_rtt_min = d_rtt_min = 500;

    for (i = 0; i < 500; i++)
    {

        /* Generate a random value between 0 to 1 */
        temp = (float) rand() / RAND_MAX;

        /* Scale the random value to fit between 50 to 75 */
        t_n = temp * 25 + 50;

```

```

/* Calculate t_RTT and d_RTT */
t_rtt_new = (alpha * t_rtt_old) + ((1 - alpha) * t_n);
d_rtt_new = (beta * d_rtt_old) + ((1 - beta) * fabs(t_n - t_rtt_old));

if (t_rtt_new < t_rtt_min)
    t_rtt_min = t_rtt_new;
if (t_rtt_new > t_rtt_max)
    t_rtt_max = t_rtt_new;

if (d_rtt_new < d_rtt_min)
    d_rtt_min = d_rtt_new;
if (d_rtt_new > d_rtt_max)
    d_rtt_max = d_rtt_new;

t_rtt_sum += t_rtt_new;
d_rtt_sum += d_rtt_new;

printf ("t_RTT: %f d_RTT: %f\n", t_rtt_new, d_rtt_new);
t_rtt_old = t_rtt_new;
d_rtt_old = d_rtt_new;
}
printf ("t_RTT min: %f t_RTT max: %f t_RTT avg: %f\n",
        t_rtt_min, t_rtt_max, (t_rtt_sum / 500.0));
printf ("d_RTT min: %f d_RTT max: %f d_RTT avg: %f\n",
        d_rtt_min, d_rtt_max, (d_rtt_sum / 500.0));
}

```

According to the preceding program the average value of $t_{RTT} = 61.6924$ and $d_{RTT} = 7.1139$. These values are averaged from a sample of over 500 values.

48. Suppose that the advertised window is 1 Mbyte long. If a sequence number is selected at random from the entire sequence number space, what is the probability that the sequence number falls inside the advertised window?

Solution:

If the sequence number field is 32 bits in length and the advertised window is 1Mbyte long, the probability that the sequence number falls inside the advertised window is:

$$P = (1 \times 10^6) / 2^{32} = 2.33 \times 10^{-4}$$

49. Explain the relationship between advertised window size, RTT, delay-bandwidth product, and the maximum achievable throughput in TCP.

Solutions follow questions:

- Plot the maximum achievable throughput versus delay-bandwidth product for an advertised window size of 65,535 bytes.

First consider delay-bandwidth product, $DBP = R \cdot 2t_p$. Here delay $2t_p$ is the propagation time that elapses from when a bit is sent by a source to the destination to when the bit is returned back to the source. This is the minimum time that elapses from when a packet leaves a source to when the acknowledgment is received. The delay-bandwidth product DBP is then the number of bits (or bytes) that are in the network from when the source transmits continuously at the maximum rate to when the bits return immediately back to the source.

The round-trip time RTT is the time that actually elapses from when a packet is sent to when its acknowledgment is received. RTT includes not only the propagation delay, but also queueing and processing delays. The advertised window W places a limit on the amount of information that a source can have outstanding in the network.

Consider the time from when a byte leaves the source to when its acknowledgment is received (that is, consider a RTT). In that time, the source will have transmitted at most a window-full of bytes into the network. Therefore the window size divided by the RTT places a limit on the throughput r , that is, the rate at which information can be transmitted into the network: $r < W/RTT$.

The throughput cannot exceed the maximum bit rate $R = DBP/2t_p$ that is available for the source to transmit into the network. Therefore, the throughput increases as the window size is increased, but cannot exceed the bit rate R :

$$\text{Throughput} = r = \min\{R, W/RTT\} = \min\{DBP/2t_p, W/RTT\}$$

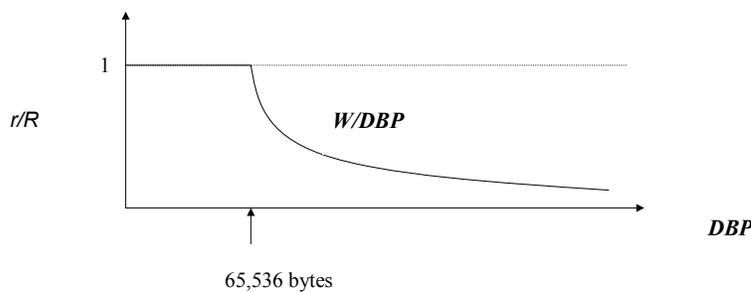
Suppose that the window size is less than the delay-bandwidth product. We then expect that the source cannot transmit at the maximum bit rate R . Indeed, we have that:

$$r < W/RTT < W/2t_p.$$

Therefore we have that:

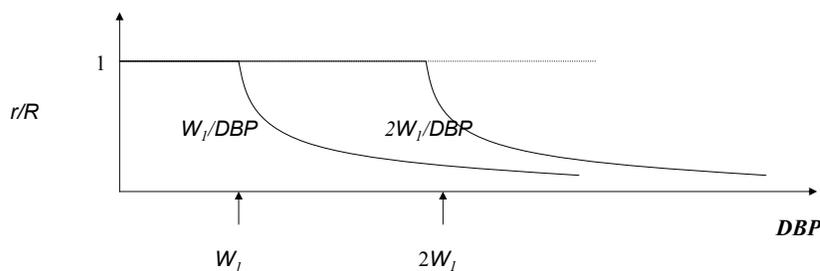
$$r/R < W/(R*2t_p) = W/DBP.$$

We conclude that the ratio of the maximum achievable throughput to R is less than the ratio of the window size to the DBP , as shown in the figure below.



- b. In the above plot include the maximum achievable throughput when the above window size is scaled up by a factor of 2^K , where $K = 4, 8, 12$.

The following figure shows the case where the window size is doubled.



- c. Place the following scenarios in the plot obtained in part (b): Ethernet with 1 Gbps and distance 100 meters; 2.4 Gbps and distance of 6000 km; satellite link with speed of 45 Mbps and RTT of 500 ms; 40 Gbps link with distance of 6000 km.

The window sizes (in bytes) of interest are: a. $65,536=2^{16}$; b. $2^{20}=10^6$; c. $2^{24}=16.7 \times 10^6$; d. $2^{28}=2.68 \times 10^8$.

Case	DBP implied
Ethernet $R = 1 \text{ Gbps}$ $D = 100 \text{ m}$	$T_p = 100 / 2.5 \times 10^8 = 4 \times 10^{-7}$ $DBP = 2 * T_p * R = 8 \times 10^{-7} * 1 \times 10^6$ $DBP = 8 \times 10^{-1}$ DBP < 1 bit Scenario a
OC-48 Link $R = 2.4 \text{ Gbps}$ $D = 6000 \text{ km}$	$T_p = 6 \times 10^3 / 2.5 \times 10^5 = 2.4 \times 10^{-1}$ $DBP = 2.4 \times 10^{-1} * 2.4 \times 10^9$ $DBP = 2.4 \times 10^8$ DBP = 2400 Mbits (300 Mbytes) Scenario slightly beyond d
Satellite link $R = 45 \text{ Mbps}$ $RTT = 500 \text{ ms } (5 \times 10^{-1} \text{ sec})$	$DBP = RTT * R$ $DBP = 5 \times 10^{-1} * 45 \times 10^6$ $DBP = 225 \times 10^5 \text{ bits}$ DBP = 21 Mbits (< 2.6 Mbytes) Scenario c
OC-768 link $R = 40 \text{ Gbps}$ $D = 6000 \text{ km } (6 \times 10^6 \text{ m})$	$T_p = 6 \times 10^3 / 2.5 \times 10^5$ $T_p = 2.4 \times 10^{-1}$ $DBP = 2 * T_p * R$ $DBP = 2 * 2.4 \times 10^{-1} * 40 \times 10^9$ $DBP = 192 \times 10^8$ DBP = 18000 Mbits (2288 Mbytes) Scenario beyond d

50. Use a network analyzer to capture the sequence of packets in a TCP connection. Analyze the contents of the segments that open and close the TCP connection. Estimate the rate at which information is transferred by examining the frame times and the TCP sequence numbers. Do the advertised windows change during the course of the connection?

Solution:

The packet capture below is from a session with the HTTP server for yahoo.com. The initial sequence number for segments from the server can be seen from frame 5 to be 2183346772, so the first data byte has sequence number 2183346773. The FIN sent from the server in frame 56 has sequence number 2183382510. Thus the server sent 35,737 bytes. These bytes were received in the period between frame 7 (time=18.47 seconds) and frame 54 (time =21.27 seconds), so the elapsed time was about 2.8 seconds. The bit rate is then about 101 kbps.

The client window size remains fixed at 8760 bytes and the server window size remains fixed at 65535 bytes.

No.	Time	Source	Destination	Protocol	Info
2	12.992506	192.168.2.1	192.168.2.1	DNS	Standard query A www.yahoo.com
3	13.001008	192.168.2.1	192.168.2.1	DNS	Standard query response CNAME www.yahoo.akadns.net A 216.109.125
4	13.001678	192.168.2.1	216.109.125.79	TCP	2498 > http [SYN] Seq=147142992 Ack=0 Win=8192 Len=0
5	13.039151	216.109.125.79	192.168.2.1	TCP	http > 2498 [SYN, ACK] Seq=2183346772 Ack=147142993 win=65535 Le
6	13.039221	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142993 Ack=2183346773 win=8760 Len=0
7	18.472270	192.168.2.1	216.109.125.79	HTTP	Continuation
8	18.600842	00000000.0001031d	00000000.Broadca	IPX SA	General query
9	18.622879	216.109.125.79	192.168.2.1	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142994 win=65535 Len=0
10	18.734094	192.168.2.1	216.109.125.79	HTTP	Continuation
11	18.905241	216.109.125.79	192.168.2.1	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142995 win=65535 Len=0
12	18.924462	192.168.2.1	216.109.125.79	HTTP	Continuation
13	19.078556	216.109.125.79	192.168.2.1	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142996 win=65535 Len=0
14	19.244070	192.168.2.1	216.109.125.79	HTTP	Continuation
15	19.369356	216.109.125.79	192.168.2.1	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142997 win=65535 Len=0
16	19.692438	192.168.2.1	216.109.125.79	HTTP	Continuation
17	19.755278	216.109.125.79	192.168.2.1	HTTP	Continuation
18	19.756467	216.109.125.79	192.168.2.1	HTTP	Continuation
19	19.756515	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183349693 win=8760 Len=0
20	19.758513	216.109.125.79	192.168.2.1	HTTP	Continuation
21	19.843349	216.109.125.79	192.168.2.1	HTTP	Continuation
22	19.843467	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183352613 win=8760 Len=0
23	19.845612	216.109.125.79	192.168.2.1	HTTP	Continuation
24	19.846719	216.109.125.79	192.168.2.1	HTTP	Continuation
25	19.846764	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183355533 win=8760 Len=0
26	19.899221	216.109.125.79	192.168.2.1	HTTP	Continuation
27	19.900426	216.109.125.79	192.168.2.1	HTTP	Continuation
28	19.900508	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183358453 win=8760 Len=0
29	19.902280	216.109.125.79	192.168.2.1	HTTP	Continuation
30	19.911624	216.109.125.79	192.168.2.1	HTTP	Continuation
31	19.911745	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183361373 win=8760 Len=0
32	19.916041	216.109.125.79	192.168.2.1	HTTP	Continuation
33	19.917275	216.109.125.79	192.168.2.1	HTTP	Continuation
34	19.917335	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183364293 win=8760 Len=0
35	19.942667	216.109.125.79	192.168.2.1	HTTP	Continuation
36	19.943834	216.109.125.79	192.168.2.1	HTTP	Continuation
37	19.943914	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183367213 win=8760 Len=0
38	19.959414	216.109.125.79	192.168.2.1	HTTP	Continuation
39	19.962971	216.109.125.79	192.168.2.1	HTTP	Continuation
40	19.963078	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183370133 win=8760 Len=0
41	19.964934	216.109.125.79	192.168.2.1	HTTP	Continuation
42	19.965690	216.109.125.79	192.168.2.1	HTTP	Continuation
43	19.965733	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183373053 win=8760 Len=0
44	19.987210	216.109.125.79	192.168.2.1	HTTP	Continuation
45	19.994098	216.109.125.79	192.168.2.1	HTTP	Continuation
46	19.994170	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183375973 win=8760 Len=0
47	20.013318	216.109.125.79	192.168.2.1	HTTP	Continuation
48	20.016425	216.109.125.79	192.168.2.1	HTTP	Continuation
49	20.016538	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183377433 win=8760 Len=0
50	20.018301	216.109.125.79	192.168.2.1	HTTP	Continuation
51	20.018342	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183377433 win=8760 Len=0
52	20.034934	216.109.125.79	192.168.2.1	HTTP	Continuation
53	20.035032	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183377433 win=8760 Len=0
54	21.275248	216.109.125.79	192.168.2.1	HTTP	Continuation
55	21.275415	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183382510 win=8760 Len=0
56	21.346468	216.109.125.79	192.168.2.1	TCP	http > 2498 [FIN, ACK] Seq=2183382510 Ack=147142999 win=65535 Le
57	21.346585	192.168.2.1	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183382511 win=8760 Len=0
58	23.239449	192.168.2.1	216.109.125.79	TCP	2498 > http [FIN, ACK] Seq=147142999 Ack=2183382511 win=8760 Len=0

51. Explain why framing information is required even in the case where frames are of constant length.

Solution:

To detect frame boundaries based on the frame length, the first frame must be detected correctly, which requires synchronization information. Also the receiver must have exact timing of the transmitter which is not practical and as a result the bit synchronization may be lost. The synchronization may be lost because of bit errors or loss of data, or loss of signal in the link.

52. Perform the bit stuffing procedure for the following binary sequence: 110111111011111110101.

Solution:

The inserted stuff bits are underlined.

110111111011111110101 → 110111110110111110110101

53. Perform bit de-stuffing for the following sequence: 11101111101111100111110.

Solution:

The removed stuff bits are indicated by a '-'.
 11101111101111100111110 → 111011111-11111-011111-

54. Consider the PPP byte stuffing method. What are the contents of the following received sequence of bytes after byte destuffing:

0x7D 0x5E 0xFE 0x24 0x7D 0x5D 0x7D 0x5D 0x62 0x7D 0x5E

Solution:

0x7D 0x5E 0xFE 0x24 0x7D 0x5D 0x7D 0x5D 0x62 0x7D 0x5E
 → 0x7E 0xFE 0x24 0x7D 0x7D 0x62 0x7E

55. Suppose that GFP operates over an error-free octet-synchronous physical layer. Find the average time the GFP receiver spends in the hunt state.

Solution:

At the beginning of its operation, the GFP receiver takes 4 bytes and checks to see if the second two bytes correspond to the check sum of the first pair of bytes. There is exactly one 16-bit pattern that corresponds to this checksum and so the probability that this pattern occurs at random is $2^{-16} = 1.5 \times 10^{-5}$, so most of the time the receiver immediately determines a given phase is incorrect. If the frame size is F, then on average F/2 bytes will be examined until the beginning of the frame is identified.

56. For GFP framing find the probability that the PLI and cHEC are not consistent. Assume that bit errors occur at random with probability p . Find the probability that the PLI and cHEC are not consistent in two consecutive frames.

Solution:

Considering that single bit errors are corrected:

$$P_{error} \approx P[errors \geq 2] \approx \binom{32}{2} p^2 (1-p)^{30} \approx 496 p^2$$

$$\text{Probability of error in two consecutive frames} = (P_{error})^2 \approx 246016 p^2$$

For example, if $p = 10^{-6}$, then probability of error in two consecutive frames is approximately 2×10^{-7} .

57. What is the maximum efficiency of GFP with respect to header overhead?

Solution:

The header is at least 4 bytes and the payload is at most 2^{16} bytes.

$$\text{Maximum efficiency} = 2^{16} / (2^{16} + 4) = 0.99994$$

58. Can GFP be used to carry video signals? If so, how would you handle the case where each video image produces a variable length of compressed information? How would you handle the case where each video frame produces a fixed length of compressed information?

Solution:

GFP can be used to carry video signals. In the case of variable-length information, frame-mapped mode is used in which variable-length GFP frames are sent. In the case of fixed-length information, transparent mode is used in which fixed-length GFP frames are sent.

59. Suppose a server has a 1 Gbps Ethernet link to a GFP transmitter. The GFP has an STS-1 SONET connection to another site. The GFP transmitter has a limited amount of buffering to accommodate frames prior to transmission. Explain how the GFP transmitter may use flow control on the Ethernet link to prevent buffer overflow. What type of flow control is preferred if the server and GFP transmitter are co-located? If the server and transmitter are 5 km apart?

Solution:

The GFP transmitter sets two threshold levels in the buffer for low and high buffer fill levels. Once the buffer fill level exceeds the high buffer fill threshold, the flow control mechanism stops the server. Once the buffer fill level crosses the low buffer fill threshold, the flow control mechanism signals the server to start transmission.

If the server and GFP transmitter are co-located X-ON X-OFF is preferred. If the server and GFP transmitter are not co-located sliding-window flow control is preferred.

60. An inverse multiplexer combines n digital transmission lines of bit rate R bps and makes them appear as a single transmission line of $n \times R$ bps. Consider an inverse multiplexer that combine multiple STS-1 SONET lines. Find an appropriate value of n to carry a 1 Gbps Ethernet stream. What is the improvement in efficiency relative to carrying the Ethernet stream on an OC-48 link?

Solution:

$$1 \times 10^9 / 50 \times 10^6 = 20, \text{ Efficiency} = 100\%$$

Efficiency if OC-48 is used: $1 \times 10^9 / 2.4 \times 10^9 = 42\%$

Improvement = $100 - 42 = 58\%$

61. Suppose that a 1-Megabyte message is sent over a serial link using TCP over IP over PPP. If the speed of the line is 56 kbps and the maximum PPP payload is 500 bytes, how long does it take to send the message?

Solution:

Assume the overhead in one packet is equal to 8 bytes for the PPP header plus 20 bytes for the IPv4 header and 20 bytes for the TCP header. Thus, the total overhead in bits is $8 \times (8 + 20 + 20) = 384$ bits. Thus,

Time to send 8×10^5 bits = (time for 1 packet)(# of packets needed)

$$\begin{aligned} &= \left(\frac{n_f}{R}\right) \left(\frac{8 \times 10^6}{n_f - n_o}\right) \\ &= \left(\frac{8 \times 500 \text{ bits}}{56 \times 10^3 \text{ bps}}\right) \left[\frac{10^6 \text{ bytes}}{(500 - 70) \text{ bytes}}\right] \\ &= 0.07143 \frac{\text{sec}}{\text{packet}} \times 2326 \text{ packets} \\ &= 166.14 \text{ seconds} \end{aligned}$$

62. Compare PPP and GFP. Identify situations where one is preferable to the other.

Solution:

PPP uses character-based framing which requires byte stuffing. As a result the frame size varies depending on the character pattern. GFP resolves this problem by using frame length based framing. GFP can support synchronous services, while PPP cannot. However PPP provides useful capabilities through link and network control protocols.

63. Use the Ethereal packet capture tool to analyze the exchange of packets during PPP setup using ad dial up connection to a local ISP. Start the Ethereal packet capture and then initiate the connection to the ISP. The number of captured packets will stop increasing after some point.

Solutions follow questions:

- Analyze the options fields of the packets during the LCP negotiations. You may need to consult RFC 1662 to interpret some of the packet contents.
- Examine the packet contents during the PAP exchange. Repeat the packet capture using CHAP.

The screen capture below shows a sequence of LCP and NCP negotiations for PPP. The two stations exchange LCP messages where they propose and reject various configuration options. The highlighted frame (9) shows the final ACK accepting a configuration. It can be seen from the middle pane that PAP authentication will be used. Frame 10 sends the password request, and frame 11 sends the password.

PPP LCP and NCP Negotiation - Ethereal

No.	Time	Source	Destination	Protocol	Info
1	0.000000	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Request
2	2.999526	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Request
3	3.130440	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Configuration Reject
4	3.130495	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Request
5	3.243457	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Configuration Ack
6	5.096025	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Configuration Request
7	5.096072	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Reject
8	5.220084	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Configuration Request
9	5.220127	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Ack
10	5.220155	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP PAP	PPP PAP Authenticate-Request
11	5.423283	20:52:45:43:56:00	20:52:45:43:56:00	PPP PAP	PPP PAP Authenticate-Ack
12	5.423367	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
13	5.423390	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP CCP	PPP CCP Configuration Request
14	5.428998	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Request
15	5.429038	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Ack
16	5.558729	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Reject
17	5.558785	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
18	5.564373	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Protocol Reject
19	5.699896	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Nak
20	5.699938	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
21	5.846675	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Ack

Frame 9 (38 bytes on wire, 38 bytes captured)

- Ethernet II, Src: 20:53:45:4e:44:00, Dst: 20:53:45:4e:44:00
- PPP Link Control Protocol
 - Code: Configuration Ack (0x02)
 - Identifier: 0x83
 - Length: 24
 - Options: (20 bytes)
 - Async control character map: 0x000a0000 (DC1 (XON), DC3 (XOFF))
 - Authentication protocol: 4 bytes
 - Authentication protocol: Password Authentication Protocol (0xc023)
 - Magic number: 0x218ad821
 - Protocol field compression
 - Address/control field compression

```

0000  20 53 45 4e 44 00 20 53 45 4e 44 00 c0 21 02 83   SEND. S END...!..
0010  00 18 02 06 00 0a 00 00 03 04 c0 23 05 06 21 8a   .....#...!..
0020  d8 21 07 02 08 02                                  !.....
    
```

- c. Analyze the IPCP packet exchange. What IP addresses are configured during the exchange? What other options are negotiated? You may wish to consult RFC 1332 for this part.

The screen capture below shows a sequence the final IPCP ACK message confirming the use of protocol field compression and various IP addresses.

The screenshot displays the Wireshark interface for a capture titled "PPP LCP and NCP Negotiation - Ethereal". The packet list pane shows the following entries:

No.	Time	Source	Destination	Protocol	Info
9	5.220127	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP LCP	PPP LCP Configuration Ack
10	5.220155	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP PAP	PPP PAP Authenticate-Request
11	5.423283	20:52:45:43:56:00	20:52:45:43:56:00	PPP PAP	PPP PAP Authenticate-Ack
12	5.423367	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
13	5.423390	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP CCP	PPP CCP Configuration Request
14	5.428998	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Request
15	5.429038	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Ack
16	5.558729	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Reject
17	5.558785	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
18	5.564373	20:52:45:43:56:00	20:52:45:43:56:00	PPP LCP	PPP LCP Protocol Reject
19	5.699896	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Nak
20	5.699938	20:53:45:4e:44:00	20:53:45:4e:44:00	PPP IPCP	PPP IPCP Configuration Request
21	5.846675	20:52:45:43:56:00	20:52:45:43:56:00	PPP IPCP	PPP IPCP Configuration Ack

The packet details pane for frame 21 shows the following structure:

- Frame 21 (42 bytes on wire, 42 bytes captured)
- Ethernet II, Src: 20:52:45:43:56:00, Dst: 20:52:45:43:56:00
- PPP IP Control Protocol
 - Code: Configuration Ack (0x02)
 - Identifier: 0x03
 - Length: 28
 - Options: (24 bytes)
 - IP compression protocol: 6 bytes
 - IP compression protocol: VJ compressed TCP (0x2d)
 - Data (2 bytes)
 - IP address: 64.228.125.69
 - Primary DNS server IP address: 206.47.244.42
 - Secondary DNS server IP address: 198.235.216.126

The packet bytes pane shows the raw data for the selected frame:

```

0000  20 52 45 43 56 00 20 52 45 43 56 00 80 21 02 03  REC.V. R ECV..!..
0010  00 1c 02 06 00 2d 0f 01 03 06 40 e4 7d 45 81 06  .....}.E..
0020  ce 2f f4 2a 83 06 c6 eb d8 7e  ./.*.~
  
```

64. Explain the differences between PPP and HDLC.

Solution:

HDLC can support various data transfer modes, supports multipoint links and point to point links, and can implement error control and flow control mechanisms. PPP uses HDLC-like frames but does not use error control and flow control protocols. Instead PPP supports powerful link and network control protocols.

PPP is character based and can be implemented on any physical layer, HDLC is bit based and can be implemented only on bit synchronous physical layer.

65. A 1.5 Mbps communications link is to use HDLC to transmit information to the moon. What is the smallest possible frame size that allows continuous transmission? The distance between earth and the moon is approximately 375,000 km, and the speed of light is 3×10^8 meters/second.

Solution:

The round trip propagation delay is:

$$2t_{prop} = 2 \frac{(375 \times 10^6 \text{ m})}{3 \times 10^8 \text{ m/s}} = 2.50 \text{ sec}$$

To allow for continuous transmission, we must use Go-Back-N or Selective Repeat.

Go-Back-N:

$$\text{If } N = 7 \rightarrow \frac{7n_f}{1.5\text{Mbps}} = 2.5s \rightarrow n_f = 535\,715 \text{ bits}$$

$$\text{If } N = 127 \rightarrow \frac{127n_f}{1.5\text{Mbps}} = 2.5s \rightarrow n_f = 29\,528 \text{ bits}$$

Selective Repeat:

$$\text{If } N = 4 \rightarrow \frac{4n_f}{1.5\text{Mbps}} = 2.5s \rightarrow n_f = 973\,500 \text{ bits}$$

$$\text{If } N = 64 \rightarrow \frac{64n_f}{1.5\text{Mbps}} = 2.5s \rightarrow n_f = 58\,594 \text{ bits}$$

66. Suppose HDLC is used over a 1.5 Mbps geostationary satellite link. Suppose that 250-byte frames are used in the data link control. What is the maximum rate at which information can be transmitted over the link?

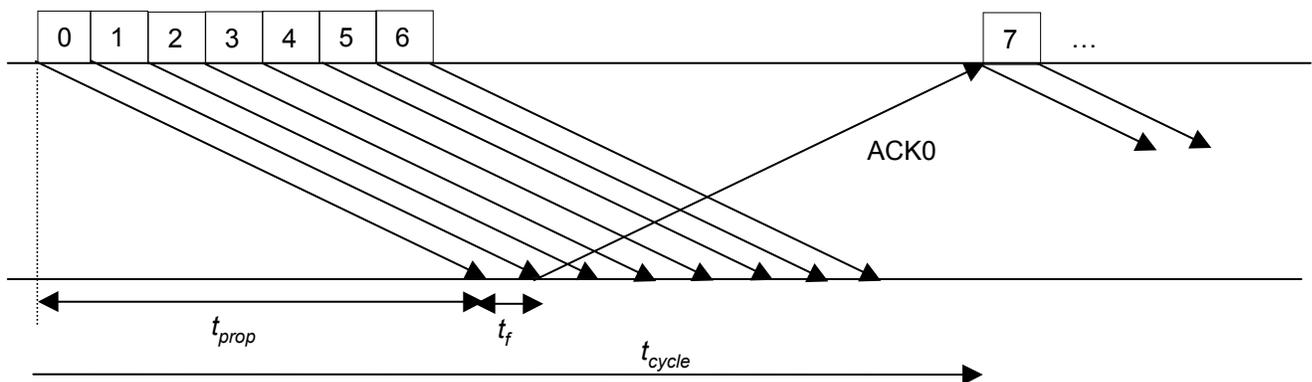
Solution:

$R = 1.5 \text{ Mbps}$, and $n_f = 250 \text{ bytes}$ or 2000 bits (250×8). The distance that the information must travel is the earth-to-satellite distance, or $d \approx 36,000 \text{ km}$. The speed of light c is 3×10^8 . We can calculate the propagation delay and processing rate as follows:

$$t_{prop} = d/c = 36 \times 10^6 / 3 \times 10^8 = 120 \text{ ms}$$

$$t_f = n_f/R = 2000/1.5 \times 10^6 = 1.33 \text{ ms}$$

We can use either Go-Back-N or Selective Repeat ARQ. The default window size is $N = 7$ (with a 3-bit sequence number).



The maximum information rate is achieved with no error, and hence, no retransmission.

$$t_{cycle} = \text{minimum time to transmit a group of } N \text{ packets}$$

$$= t_f + 2 t_{prop} = 1.33 + 2 \times 120 = 241.33 \text{ ms}$$

$$n = \text{no. of bits transmitted in a cycle} = N \cdot n_f = 7 \times 2000 = 14,000 \text{ bits}$$

$$R_{max} = \text{no. of bits sent in a cycle} / \text{minimum cycle time} = n / t_{cycle} = 58 \text{ kbps}$$

If the extended sequence numbering option (7-bit) is used, the maximum send window size would be $N = 2^7 - 1 = 127$, and hence, the maximum information rate is:

$$R_{max} = N \cdot n_f / t_{cycle} = 127 \times 2000 / (241.33 \times 10^{-3}) = 1.052 \text{ Mbps}$$

67. In HDLC how does a station know if a received frame with the fifth bit set to 1 is a P or an F bit?

Solution:

If the station is a secondary station, the bit is a 'P' (and the station is being *Polled* for more frames). If it is a primary station, the bit is a 'F' bit (indicating the *Final* frame of the current transmission).

68. Which of the following statements are incorrect?

Solutions follow questions:

For this question, one must just keep in mind that all frames always contain the address of the secondary station.

a. A transmitting station puts its own address in command frames.

Incorrect. Command frames are destined for secondary stations, so this transmitter would put the address of the secondary station in the frame.

b. A receiving station sees its own address in a response frame.

Incorrect. Responses come from secondary stations to primary stations. Thus the address in the frame is that of the sender in this case.

c. A response frame contains the address of the sending station.

TRUE. Response frame originate in secondary stations, so they will have contain the address of the sender.

69. In HDLC suppose that a frame with P = 1 has been sent. Explain why the sender should not send another frame with P = 1. What should be done to deal with the case where the frame is lost during transmission?

Solution:

The poll bit is like a token that is passed to a secondary so it can transmit. At the end of transmission the secondary must return it in the form of final bit.

If the poll frame is lost then the token no longer exists and no station can transmit. The server must create a new token (poll frame) only after it is sure that the previous poll frame is lost and not just delayed. Consequently it must wait for certain period of time.

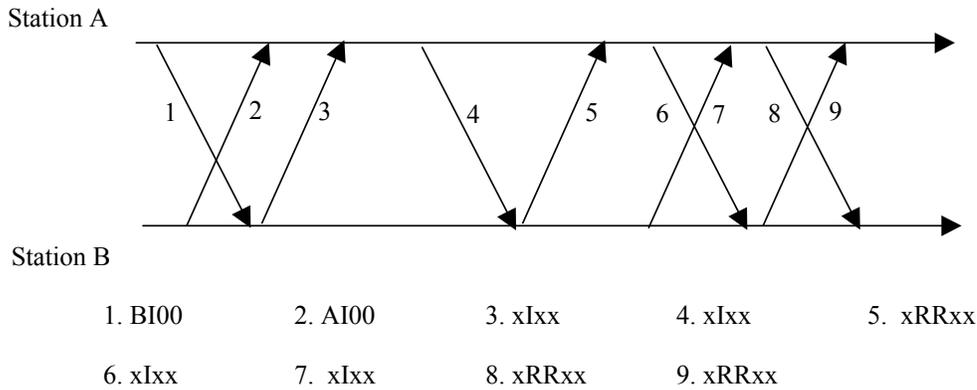
70. HDLC specifies that the N(R) in a SREJ frame requests the retransmission of frame N(R) and also acknowledges all frames up to N(R) - 1. Explain why only one SREJ frame can be outstanding at a given time.

Solution:

Suppose two outstanding SREJ frames exist. Let frame A have N(R) = m and frame B have N(R) = n . Without loss of generality, suppose $n > m$.

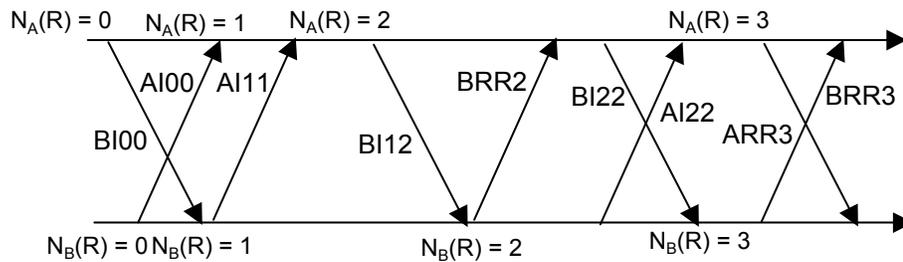
Since each SREJ frame with value $N(R)$ implicitly acknowledges all previous frames up to $N(R) - 1$, frame A indicates that frame m has not yet been received and frame B indicates that frame m has been received. Thus, if two SREJ are allowed to be outstanding at the same time, contradictory information will be sent to the receiver.

71. The following corresponds to an HDLC ABM frame exchange with no errors.

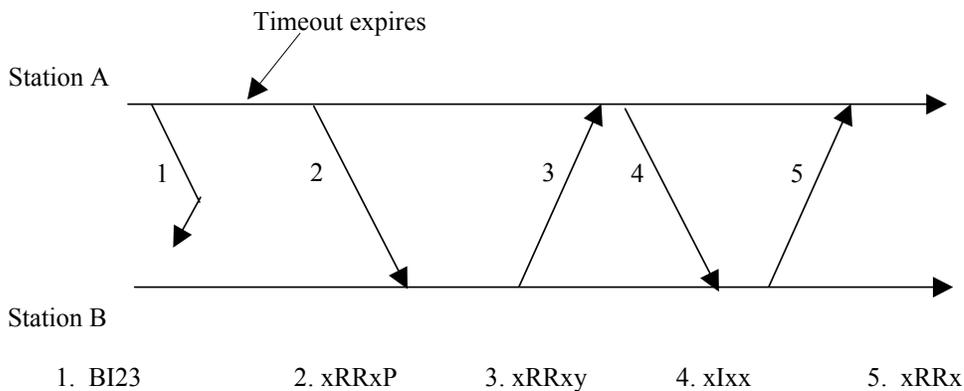


- Complete the diagram by completing the labeling of the frame exchanges.
- Write the sequence of state variables at the two stations as each event takes place.

Solution:

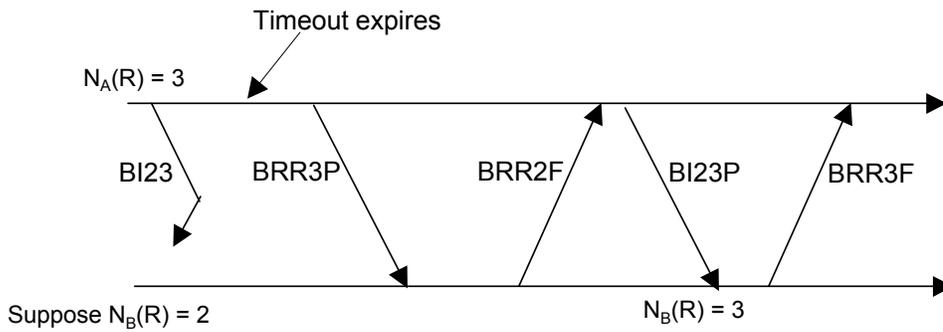


72. Assume station B is awaiting frame 2 from station A.

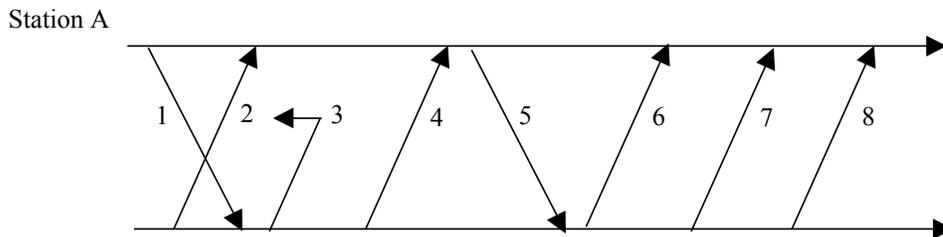


- Complete the diagram in HDLC ABM by completing the labeling of the frame exchanges.
- Write the sequence of state variables at the two stations as each event takes place.

Solution:



73. The following corresponds to an HDLC ABM frame exchange.

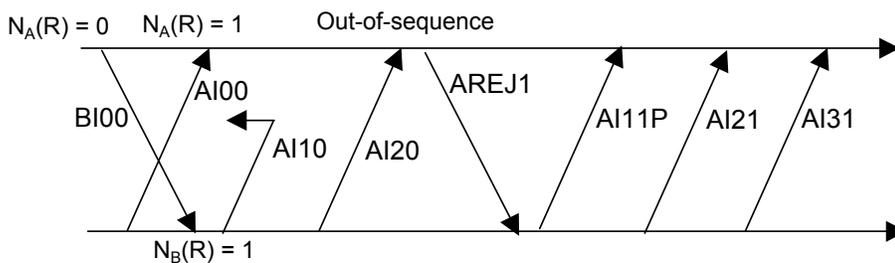


Station B

- | | | | |
|----------|---------|---------|---------|
| 1. BI00 | 2. AI00 | 3. xIxx | 4. xIxx |
| 5. xREJx | 6. xIyx | 7. xyxx | 8. xyxx |

- Complete the diagram by completing the labeling of the frame exchanges.
- Write the sequence of state variables at the two stations as each event takes place.

Solution:



74. Suppose that packets arrive from various sources to a statistical multiplexer that transmits the packets over a 64 kbps PPP link. Suppose that the PPP frames have lengths that follow an exponential distribution with mean 1000 bytes and that the multiplexer can hold up to 100 packets at a time. Plot the average packet delay as a function of the packet arrival rate.

Solution:

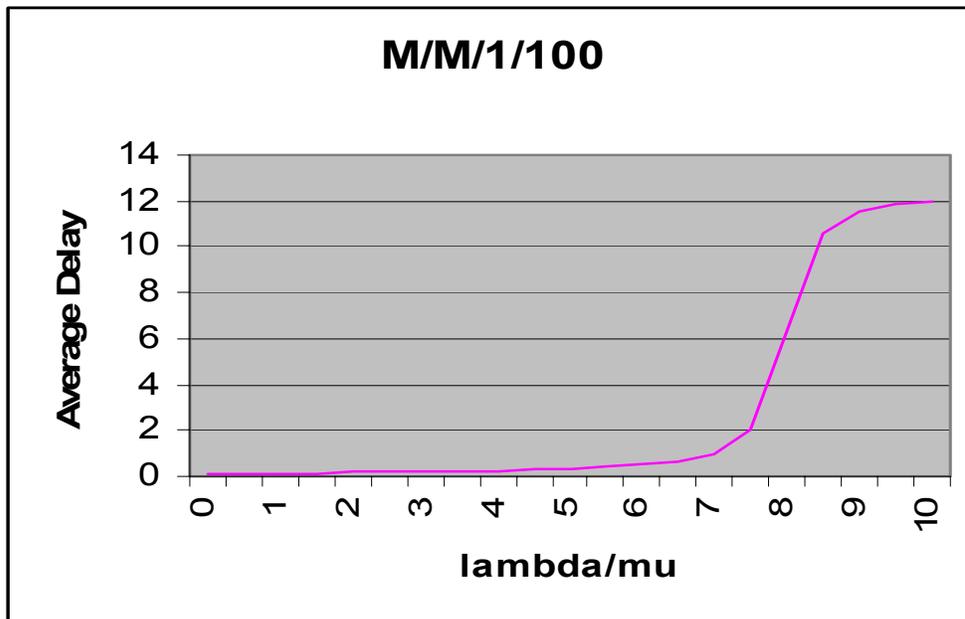
This is an M/M/1/K system in which:

Service rate $\mu = 64000 \text{ bps} / 8000 \text{ bits/packet} = 8 \text{ packets/sec}$
 Packet length $E[L] = 8000 \text{ bits}$
 Buffer size $K = 100 \text{ packets}$

Thus:

$$E[T] = \frac{E[N]}{\lambda(1 - P_L)}$$

$$\text{where } P_L = \frac{(1 - \frac{\lambda}{\mu})(\frac{\lambda}{\mu})^K}{1 - (\frac{\lambda}{\mu})^{K+1}} \text{ and } E[N] = \frac{\lambda}{\mu - \lambda} - \frac{(K + 1)\lambda^{K+1}}{\mu^{K+1} - \lambda^{K+1}}$$



75. Suppose that the traffic that is directed to a statistical multiplexer is controlled so that ρ is always less than 80%. Suppose that packet arrivals are modeled by a Poisson process and that packet lengths are modeled by an exponential distribution. Find the minimum number of packet buffers required to attain a packet loss probability of 10^{-3} or less.

Solution:

$$P_{loss} = \frac{(1 - \rho) \rho^k}{1 - \rho^{k+1}}$$

with $k = 24$ $P_{loss} = 0.000948$

76. Suppose that packets arrive from various sources to a statistical multiplexer that transmits the packets over a 1 Mbps PPP link. Suppose that the PPP frames have constant length of L bytes and that the multiplexer can hold a very large number of packets at a time. Assume that each PPP frame contains a PPP, IP, and TCP header in addition

to the user data. Plot the average packet delay as a function of the rate at which user information is transmitted for $L = 250$ bytes, 500 bytes, and 1000 bytes.

Solution:

Assume the overhead in one packet is equal to 8 bytes for the PPP header plus 20 bytes for the IPv4 header and 20 bytes for the TCP header.

$$R = 10^6 \text{ bps}$$

$$L = \text{constant (250, 500, and 1000 bytes)}$$

$$K \rightarrow \infty$$

$$n_o = 48 \text{ bytes}$$

$$\mu = R/(8*L) \text{ packets/sec} = 500 \text{ (L=250 B); } 250 \text{ (L=500 B); } 125 \text{ (L=1000)}$$

Because of the constant length packets, this is an M/D/1 system.

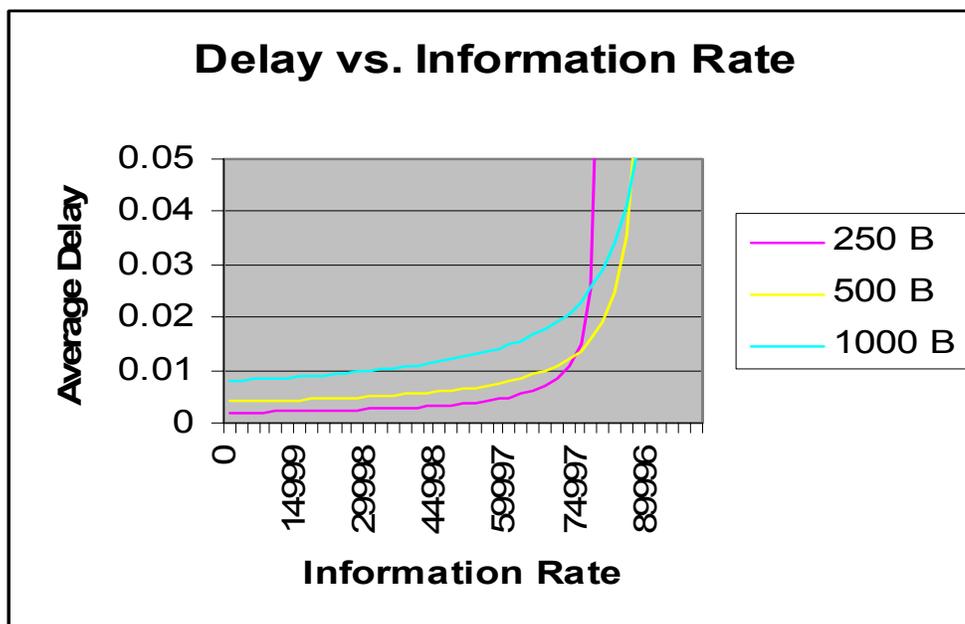
$$\text{Average delay, } E[T_D] = \left(1 + \frac{\lambda}{2(\mu - \lambda)}\right) \frac{1}{\mu}$$

Let γ be the rate in bits per second at which user information transferred.

$$\gamma = \lambda(8L) \frac{(L - n_o)}{L} = \lambda(8L) \left(1 - \frac{48}{L}\right) \text{ or } \lambda = \frac{\gamma}{8L - 384}$$

And so, the delay as a function of the user information bit rate is

$$E[T_D] = \left(1 + \frac{\frac{\gamma}{8L - 384}}{2\left(\frac{R}{8L} - \frac{\gamma}{8L - 384}\right)}\right) \frac{8L}{R}$$



77. Suppose that a multiplexer receives constant-length packet from $N = 60$ data sources. Each data source has a probability $p = 0.1$ of having a packet in a given T -second period. Suppose that the multiplexer has one line in which it can transmit eight packets every T seconds. It also has a second line where it directs any packets that cannot be transmitted in the first line in a T -second period. Find the average number of packets that are transmitted on the first line and the average number of packets that are transmitted in the second line.

Solution:

The probability that there are k packet arrivals in a T -second period is given by the binomial distribution with parameters $N = 60$ and $p = 0.1$. The average number of arrivals is $Np = 6$. The average number of arrivals that get transferred to the first line is given by:

$$\sum_{k=0}^8 k \binom{60}{k} (0.1)^k (0.9)^{60-k} = 4.59$$

The remainder of the packet arrivals is sent to the second line, so the average number sent to line 2 is $6 - 4.59 = 1.41$ packets per T -second period.

78. Discuss the importance of queueing delays in multiplexers that operate at bit rates of 1 Gbps or higher.

Solution:

By aggregating traffic, better delay performance is achieved, as illustrated in section 5.5. In multiplexers that operate at bit rates of 1 Gbps or higher, the queueing delay experienced by packets is much shorter than it would be if many lower bit rate systems were used.

Changes to this document

<i>Version No.</i>	<i>Problem No.</i>	<i>Change</i>
1.00 to 1.01	5.14a	Changed reference Figure 5.12 to 5.10 in solution text
	5.35	Changed reference Table 5.12 to Equation 5.11 in solution text.
	Throughout	Cosmetic changes resulting in changes in pagination

Solutions to Chapter 6

1. Why do LANs tend to use broadcast networks? Why not use networks consisting of multiplexers and switches?

Solution:

The computers in a LAN are separated by a short distance (typically < 100m) so high speed and reliable communication is possible using a shared broadcast medium. The cost of the medium is negligible and the overall cost is dominated by the cost of the network interface cards in each computer. In addition, the LAN users usually belong to the same group where all users are generally trusted, so broadcast does not pose much security danger.

The original reason for avoiding a multiplexer and switch approach to LANs is that a centralized, expensive "box" is required. The availability of Application Specific Integrated Circuits (ASICs) has reduced the cost of switching boxes and made switch-based LANs feasible, and in some environments the dominant approach.

2. Explain the typical characteristics of a LAN in terms of network type, bit rate, geographic extent, delay-bandwidth product, addressing, and cost. For each characteristic, can you find a LAN that deviates from the typical? Which of the above characteristics is most basic to a LAN?

Solution:

- Type: broadcast network in a bus, ring or star topology.
- Bit rate: from 1Mbps to 100 Mbps.
- Delay-bandwidth product: small.
- Addressing: flat.
- Geographical extent: up to 1000 m (small).
- Cost: low.
- 10 Gbps Ethernet deviates from the above characteristics in that it is non-broadcast, of much higher bit rate, large delay-bandwidth product, larger geographic extent and high cost.

The most basic characteristic of a LAN is small geographical extent.

3. Compare the two-channel approach (Figure 6.4) with the single-channel approach (Figure 6.5) in terms of the types of MAC protocols they can support.

Solution:

Figure 6.4 consists of two unidirectional channels, one outbound from a central node to secondary nodes, and another inbound from the secondaries to the central node. The bandwidth that is available in each direction is fixed. This arrangement can support polling protocols as well as contention protocols in the inbound direction. Figure 6.5 provides a single channel that is shared by all stations. This arrangement also supports polling and contention-type MAC protocols. The bandwidth available in each direction can be controlled dynamically.

4. Suppose that the ALOHA protocol is used to share a 56 kbps satellite channel. Suppose that frames are 1000 bits long. Find the maximum throughput of the system in frames/second.

Solution:

Maximum throughput for ALOHA = 0.184

Maximum throughput in frames/sec = (56000 bits/sec) x (1 frame/1000 bits) x 0.184 = 10.304

The maximum throughput is approximately 10 frames/sec.

5. Let G be the total rate at which frames are transmitted in a slotted ALOHA system. What proportion of slots goes empty in this system? What proportion of slots go empty when the system is operating at its maximum throughput? Can observations about channel activity be used to determine when stations should transmit?

Solution:

$$\text{Proportion of empty slots} = P[0 \text{ transmission}] = [G^0/0!]e^{-G} = e^{-G}$$

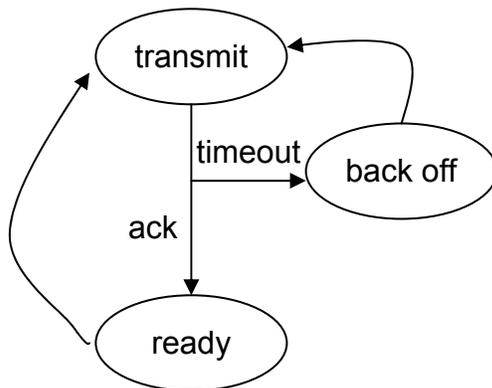
$$\text{Maximum throughput} = 0.368; G_{mt} = 1$$

$$\text{Proportion of empty slots at maximum throughput} = e^{-1} = 0.368$$

Any attempt to decrease the proportion of empty slots below e^{-1} is counterproductive as this action will push the throughput below its maximum value.

6. Modify the state transition diagram of Stop-and-Wait ARQ to handle the behavior of a station that implements the ALOHA protocol.

Solution:



7. Suppose that each station in an ALOHA system transmits its frames using spread spectrum transmission. Assume that the spreading sequences for the different stations have been selected so that they have low cross-correlations. What happens when transmissions occur at the same time? What limits the capacity of this system?

Solution:

The transmitted signals from different stations occupy the whole frequency band simultaneously. To each receiver, the aggregate of the other signals appears like noise after the receiver applies its spreading sequence in the demodulation. Consequently, the receiver can reliably extract its desired signal as long as the signal to noise ratio remains sufficiently high.

The system capacity is limited by the maximum amount of signal interference allowed at the receiver. Increasing the number of signals in the system increases the interference power level. As the number of signals is increased, the SNR decreases and the bit error rate in each receiver will increase.

8. Consider four stations that are all attached to two different bus cables. The stations exchange fixed-size frames of length 1 second. Time is divided into slots of 1 second. When a station has a frame to transmit, the station chooses either bus with equal probability and transmits at the beginning of the next slot with probability p . Find the value of p that maximizes the rate at which frames are successfully transmitted.

Solution:

To maximize the successful transmission rate is to maximize the probability of successful transmission.

$P(\text{success}) = (\text{number of stations}) \times P(\text{one station transmits on one bus}) \times P(\text{no other station transmit on the same bus})$

$$= 4\left(\frac{1}{2}p\right)\left(1 - \frac{1}{2}p\right)^3 = 2p\left(1 - \frac{1}{2}p\right)^3$$

Take the derivative with respect to p ,

$$\frac{\partial P(\text{success})}{\partial p} = 2\left(1 - \frac{1}{2}p\right)^3 - (3p)\left(1 - \frac{1}{2}p\right)^2$$

set it to 0 and find the value of p that maximizes $P(\text{success})$.

$$p = \frac{1}{2}$$

9. In a LAN, which MAC protocol has a higher efficiency: ALOHA or CSMA-CD? What about in a WAN?

Solution:

The maximum efficiency achieved by the Slotted ALOHA is 0.368. The efficiency of CSMA-CD is given by $1/(1 + 6.4a)$, and is sensitive to $a = t_{prop}R/L$, the ratio between delay-bandwidth product and frame length.

In a LAN environment, the end-to-end distance is around 100m and the transmission rates are typically 10Mbps, 100Mbps and 1Gbps (See Table 6.1). An Ethernet frame has a maximum length of 1500 bytes = 12,000 bits.

The table shows the efficiency of CSMA-CD at various transmission rate. Assume $L = 12,000$ bits and propagation speed of 3×10^8 .

	a	Efficiency
10 Mbps	3×10^{-4}	0.998
100 Mbps	3×10^{-3}	0.981
1 Gbps	3×10^{-2}	0.839

Note however that if shorter frame sizes predominate, e.g. 64 byte frames, then a increases by a factor of about 20. According to the above formula the efficiency of CSMA-CD at 1 Gbps then drops to about 0.7. The situation however is worse in that the minimum frame size at 1 Gbps needs to be extended to 512 bytes, as discussed in page 436 of the text.

In a WAN environment d is larger. Assuming 100 Km, a is larger by a factor of 10^3 resulting in an efficiency of 0.36, 0.05, and 0.005 respectively for 10 Mbps, 100 Mbps, and 1 Gbps transmission rates. In the case of 10 Mbps transmission rate the efficiency of CSMA-CD is close to the efficiency of ALOHA but in the other two cases it is much less than ALOHA.

10. A channel using random access protocols has three stations on a bus with end-to-end propagation delay τ . Station A is located at one end of the bus, and stations B and C are together located at the other end of the bus. Frames arrive at the three stations and are ready to be transmitted at stations A, B, and C at the respective times $t_A = 0$, $t_B = \tau/2$, and $t_C = 3\tau/2$. Frames require transmission times of 4τ . In appropriate figures, with time as the horizontal axis, show the transmission activity of each of the three stations for

Frame arrival times:

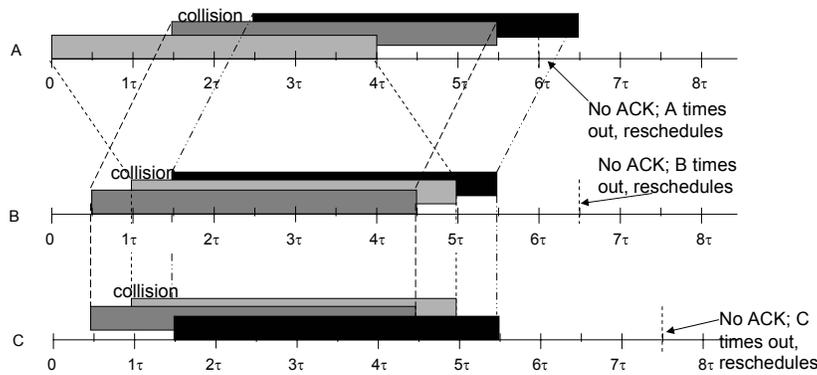
A: $t_A = 0$

B: $t_B = \tau/2$

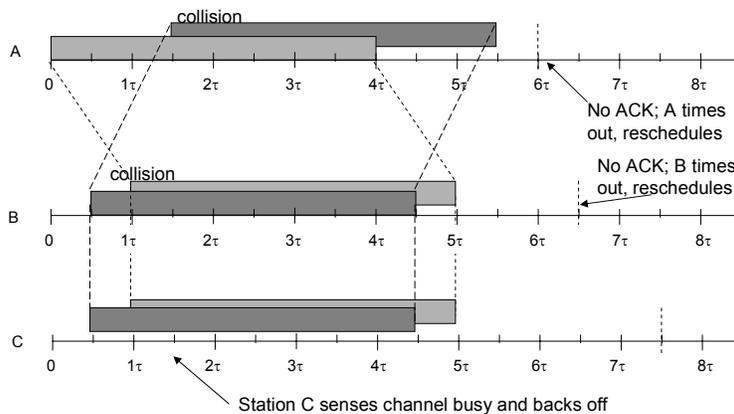
C: $t_C = 3\tau/2 = 1\ 1/2\ \tau$

$t_p = \tau$ and $X = 4\tau$

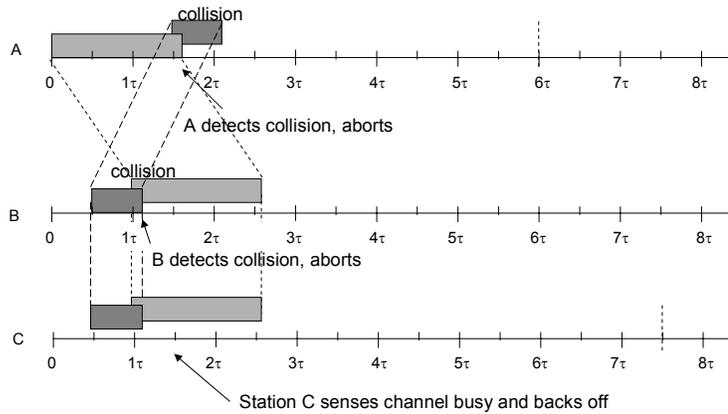
a. ALOHA



b. Non-persistent CSMA



c. Non-persistent CSMA-CD.



11. Estimate the maximum throughput of the CDPD system assuming packets are 1096 bytes in length.

Hint: What is a for this system?

Solution:

During the contention interval, it takes two microblock delays ($2t_{\text{ublock}}$) for a station to know if it has captured the reverse channel. Divide the time into minislots of $2t_{\text{ublock}}$.

As in the CSMA-CD system, the maximum probability of successful transmission is $P_{\text{success}}^{\text{max}} = \frac{1}{e}$, and

the average number of minislots in a contention period is $E[J] = \frac{1}{P_{\text{success}}^{\text{max}}} = e$.

The maximum throughput occurs when the channel time is spent in packet transmission followed by a contention period. After the packet transmission $E[X]$, each station spends one t_{ublock} to determine if the channel is idle or busy.

The maximum throughput is given by

$$\rho_{\text{max}} = \frac{E[X]}{E[X] + t_{\text{ublock}} + e(2t_{\text{ublock}})} = \frac{1}{1 + (2e + 1)a} \text{ where } a = \frac{t_{\text{ublock}}}{E[X]}.$$

Calculation:

$$t_{\text{ublock}} = \frac{60 \text{ bits}}{19,200 \text{ bps}} = 3.1 \text{ ms}$$

Assume the packet is already in HDLC format.

$$1096 \text{ bytes} = 8768 \text{ bits}$$

Packet length after segmentation and header insertion:

$$L = \frac{8768}{274} \times \frac{378}{54} \times (54 + 6) = 13440 \text{bits}$$

$$E[X] = \frac{13440 \text{bits}}{19200 \text{bps}} = 0.7 \text{s}$$

$$a = \frac{3.1 \text{ms}}{0.7 \text{s}} = 4.43 \times 10^{-3}$$

Therefore, the maximum throughput is

$$\rho_{\max} = \frac{1}{1 + (2e + 1)(4.43 \times 10^{-3})} = 0.97 .$$

12. Can the Digital Sense Multiple Access protocol, which is used by CDPD, also be used on the digital carrier of GSM? If yes, explain how.

Solution:

Digital Sense Multiple Access would require redesign of the GSM frame structure. In the downlink transmission broadcast from the base station, a version of the flag word would need to be inserted to indicate whether the medium is busy. Furthermore since the frame structure in GSM effectively divides the frame into 8 channels, a separate Digital Sense Multiple Access protocol would need to be in operation for each channel. This in turn would require the presence of a flag word per channel in the downlink frames.

13. M terminals are attached by a dedicated pair of lines to a hub in a star topology. The distance from each terminal to the hub is d meters, the speed of the transmission lines is R bits/second, all frames are of length 12500 bytes, and the signal propagates on the line at a speed of $2.5 (10^8)$ meters/second. For the four combinations of the following parameters $\{d = 25 \text{ meters or } d = 2500 \text{ meters}; R = 10 \text{ Mbps or } R = 10 \text{ Gbps}\}$, compare the maximum network throughput achievable when the hub is implementing: Slotted ALOHA; CSMA/CD.

Solution:

$$L = 12500 \times 8 \text{ bits}, t_{\text{prop}} = d / (2.5 \times 10^8 \text{ meters/sec}), a = t_{\text{prop}} R/L$$

Values for a :

R/d	2x25	2x2500
1.00E+07	2E-05	2E-03
1.00E+10	2E-02	2E+00

Maximum Throughput for Slotted ALOHA:

R/d	2x25	2x2500
1.00E+07	0.367879	0.367879
1.00E+10	0.367879	0.367879

Maximum throughput for CSMA-CD:

R/d	2x25	2x2500
1.00E+07	0.999872	0.98736
1.00E+10	0.886525	0.07246

14. Consider the star-topology network in problem 6.13 when the token-ring protocol is used for medium access control. Assume single-frame operation, eight-bit latency at each station, $M = 125$ stations. Assume a free token is three bytes long.

Solutions follow questions:

$$\begin{aligned} M &= 125 \\ b &= 8 \text{ bits} \\ L_{\text{token}} &= 3 \text{ bytes} \\ L_{\text{frame}} &= 12,500 \text{ bytes} \\ v &= 2.5 \times 10^8 \text{ m/sec} \end{aligned}$$

- a. Find the effective frame transmission time for the four combinations of d and R .

The distance from each terminal to the hub is d meters, so the total distance around the ring is then $M2d$. Assuming *single frame transmission token reinsertion*, then

$$\begin{aligned} X_{\text{eff}} &= \text{token transmission time} + \text{frame transmission time} + \text{ring latency} \\ &= \frac{L_{\text{token}} + L_{\text{frame}}}{R} + M \left(\frac{2d}{v} + \frac{b}{R} \right) \end{aligned}$$

- b. Assume that each station can transmit up to a maximum of k frames/token. Find the maximum network throughput for the four cases of d and R .

The maximum throughput occurs when all stations transmit k frames per token. After completing the transmission of k frames, each station waits one ring latency time and then transmits a free token into the ring.

$$\rho_{\text{max}} = \frac{M \times k \times X_{\text{frame}}}{M(k \times X_{\text{frame}} + X_{\text{token}} + \tau) + \tau} = \frac{1}{1 + \frac{X_{\text{token}} + \tau}{kX_{\text{frame}}} + \frac{\tau}{MkX_{\text{frame}}}}$$

where

$$\tau = M \left(\frac{2d}{v} + \frac{b}{R} \right) \text{ is the amount time spent in passing the token around the ring.}$$

Calculation:

- i. $d = 25$ m, $R = 10$ Mbps

$$\tau = 0.125 \text{ ms}, X_{\text{frame}} = 10 \text{ ms}, X_{\text{token}} = 2.4 \mu\text{sec},$$

$$X_{\text{eff}} = \frac{(3 + 12500) \times 8}{10M} + 125 \left(\frac{2 \times 25}{2.5 \times 10^8} + \frac{8}{10M} \right) = 10.1 \text{ ms}$$

$$\rho_{\max} \approx \frac{1}{1 + \frac{.125}{10k}} \approx 99\%$$

X_{token} and τ are negligible.

ii. $d = 25$ m, $R = 10$ Gbps

$$\tau = 25.1 \text{ } \mu\text{sec}, X_{\text{frame}} = 10 \text{ } \mu\text{sec}, X_{\text{token}} = 0.0024 \text{ } \mu\text{sec},$$

$$X_{\text{eff}} = 35.1 \text{ } \mu\text{sec}$$

$$\rho_{\max} \approx \frac{1}{1 + \frac{25.1}{10k}} = \frac{1}{1 + \frac{2.5}{k}}$$

$\rho_{\max} = 28\%$ for $k=1$ and improves as k increases.

iii. $d = 2500$ m, $R = 10$ Mbps

$$\tau = 2.6 \text{ ms}, X_{\text{frame}} = 10 \text{ ms}, X_{\text{token}} = 2.4 \text{ } \mu\text{sec},$$

$$X_{\text{eff}} = 12.6 \text{ ms}$$

$$\rho_{\max} \approx \frac{1}{1 + \frac{2.6}{10k}} = \frac{1}{1 + \frac{0.26}{k}}$$

$\rho_{\max} = 79\%$ for $k = 1$ and improves as k increases.

iv. $d = 2500$ m, $R = 10$ Gbps

$$\tau = 2.5 \text{ ms}, X_{\text{frame}} = 10 \text{ } \mu\text{sec}, X_{\text{token}} = 0.0024 \text{ } \mu\text{sec},$$

$$X_{\text{eff}} = 2.51 \text{ ms}$$

$$\rho_{\max} \approx \frac{1}{1 + \frac{2.5}{.010k}} = \frac{1}{1 + \frac{250}{k}}$$

$\rho_{\max} = 0.40\%$, nearly zero.

15. A wireless LAN uses polling to provide communications between M workstations and a central base station. The system uses a channel operating at 25 Mbps. Assume that all stations are 100 meters from the base station and that polling messages are 64 bytes long. Assume that frames are of constant length of 1250 bytes. Assume that stations indicate that they have no frames to transmit with a 64-byte message.

Solutions follow questions:

$d = 100$ m between the base station and the stations

$v = 3 \times 10^8$ m/sec

$t_{prop} = 100 / (3 \times 10^8) = 0.33$ μ sec

$R = 25$ Mbps

$X_{frame} = 10000/25\text{Mbps} = 400$ μ sec

$X_{poll} = 512/25\text{Mbps} = 20$ μ sec

$R = 2.5$ Gbps

$X_{frame} = 10000/2.5\text{Gbps} = 4$ μ sec

$X_{poll} = 512/2.5\text{Gbps} = 0.2$ μ sec

$X_{end} = X_{poll}$

- a. What is the maximum possible arrival rate that can be supported if stations are allowed to transmit an unlimited number of frames/poll?

$$\rho_{max} = 1 \text{ and } \rho = \lambda X$$

$$R = 25 \text{ Mbps}$$

$$\lambda_{max} = 1/400 \text{ } \mu\text{sec} = 2,500 \text{ frames/sec}$$

- b. What is the maximum possible arrival rate that can be supported if stations are allowed to transmit N frames/poll?

$$\rho_{max} = \frac{MN X_{frame}}{M \{NX_{frame} + X_{end} + X_{poll} + 2t_{prop}\}} = \frac{1}{1 + \frac{20 + 20 + 0.66}{400N}} \approx \frac{1}{1 + \frac{0.1}{N}}, \quad \lambda_{max} = \frac{\rho_{max}}{X} \text{ as } N$$

increases, ρ_{max} approaches 1.

$$R = 25 \text{ Mbps}$$

$$N = 10, \rho_{max} = 99\%$$

$$\lambda_{max} = 1/400\mu = 2,500 \text{ frames/sec}$$

- c. Repeat parts (a) and (b) if the transmission speed is 2.5 Gbps.

(a) $R = 2.5$ Gbps

$$\rho_{max} = 1$$

$$\lambda_{max} = 1/4 \text{ } \mu\text{sec} = 250,000 \text{ frames/sec}$$

(b) $R = 2.5$ Gbps

$$\rho_{max} = \frac{1}{1 + \frac{0.2 + 0.2 + 0.66}{4N}} \approx \frac{1}{1 + \frac{0.26}{N}}, \quad \lambda_{max} = \frac{\rho_{max}}{X}$$

$$N = 10,$$

$$\rho_{max} = 97.4\%$$

$$\lambda_{max} = 0.788/4\mu = 243,700 \text{ frames/sec}$$

16. A token-ring LAN network interconnects M stations using a star topology in the following way. All the input and output lines of the token-ring station interfaces are connected to a cabinet where the actual ring is placed. Suppose that the distance from each station to the cabinet is 100 meters and that the ring latency per station is eight bits. Assume that frames are 1250 bytes and that the ring speed is 25 Mbps.

Solutions follow questions:

$d = 100$ m from each station to the cabinet

$v = 2 \times 10^8$ m/sec

$b = 8$ bits

$L = 1250$ bytes = 10000 bits

$R = 25$ Mbps

$$X = \frac{10000}{25M} = 4 \times 10^{-4} \text{ sec}$$

$$\tau' = \frac{M2d}{v} + \frac{M8}{R} = \frac{M200}{2 \times 10^8} + \frac{M8}{25 \times 10^6} = 1.32 \times 10^{-6} M$$

$$a' = \frac{\tau'}{X} = \frac{1.32 \times 10^{-6} M}{4 \times 10^{-4}} = 3.3 \times 10^{-3} M$$

- a. What is the maximum possible arrival rate that can be supported if stations are allowed to transmit an unlimited number of frames/token?

When all stations are allowed to transmit an unlimited number of frames/token, $\rho_{\max} = 1$ and $\rho = \lambda X$.

$$\lambda_{\max} = \frac{1}{4 \times 10^{-4}} = 2500 \text{ frames/sec}$$

- b. What is the maximum possible arrival rate that can be supported if stations are allowed to transmit 1 frame/token using single-frame operation? Using multiple token operation?

Single-frame operation:

$$\rho_{\max} = \frac{1}{1 + a' \left(1 + \frac{1}{M}\right)}$$

$$\lambda_{\max} = \frac{1}{1 + (3.3 \times 10^{-3} M) \left(1 + \frac{1}{M}\right)} \cdot \frac{1}{4 \times 10^{-4}} = \frac{2500}{1.0033 + 0.0033M}$$

Multitoken operation:

$$\rho_{\max} = \frac{1}{1 + a'/M} = \frac{1}{1 + 3.3 \times 10^{-3} M/M} = 0.997$$

$$\lambda_{\max} = 0.997 \frac{1}{4 \times 10^{-4}} = 2492 \text{ frames/sec}$$

- c. Repeat parts (a) and (b) if the transmission speed is 2.5 Gbps.

$R = 2.5$ Gbps

$$X = \frac{10000}{2.5G} = 4 \times 10^{-6} \text{ sec}$$

$$\tau' = \frac{M2d}{v} + \frac{M8}{R} = \frac{M200}{2 \times 10^8} + \frac{M8}{2.5G} = 1 \times 10^{-6} M$$

$$a' = \frac{\tau'}{X} = \frac{1 \times 10^{-6} M}{4 \times 10^{-6}} = 0.25M$$

Unlimited number of frames per token:

$$\lambda_{\max} = \frac{1}{4 \times 10^{-6}} = 250000 \text{ frames/sec}$$

Single-frame operation:

$$\lambda_{\max} = \frac{1}{1 + (0.25M)(1 + \frac{1}{M})} \cdot \frac{1}{4 \times 10^{-6}} = \frac{250000}{1.25 + 0.25M} \text{ frames/sec}$$

Multitoken operation:

$$\lambda_{\max} = \frac{1}{1 + 0.25M/M} \cdot \frac{1}{4 \times 10^{-6}} = 0.2 \times 10^6 \text{ frames/sec}$$

17. Suppose that a LAN is to carry voice and packet data traffic. Discuss what provisions if any are required to handle the voice traffic in the reservation, polling, token ring, ALOHA and CSMA-CD environments. What changes if any are required for the packet data traffic?

Solution:

Voice traffic is delay sensitive, so the MAC protocols must ensure the delays experienced by voice data packets are sufficiently low.

The reservation, polling and token ring systems use scheduling approach to medium access control and so are better able to provide predictable delay performance. Random access approaches on the other hand have greater variability in delay.

Reservation: Reservation schemes can provide performance close to time-division multiplexing and hence can support voice traffic.

Polling: Polling schemes can be configured to provide a bound on the value on the cycle time. This provides a guaranteed delay bound for voice traffic.

Token Ring: Limit the transmission time per token to place a maximum value on the token rotation time. Avoid single-token or single-frame operation if ring latency is large. Priority token operation can also be used if there is non-voice traffic in the network.

The ALOHA and CSMA-CD systems use random access approach to medium access control. The system throughput and delay performance are affected by traffic load. At a high load, the collision rate is high, and the packets experience long delays. To meet the delay requirement of voice traffic, the maximum traffic load should be kept very low for Aloha and relatively low for CSMA-CD. The CSMA-CD systems can be modified to provide different levels of access priority. Lower delay performance can be provided by giving voice traffic higher priority.

Packet data traffic is usually not delay sensitive, but does require low loss. In carrying both voice and data, the MAC protocol must provide voice traffic with low delay and data with low loss. In general, this requires giving voice priority access over data. Otherwise surges in data traffic can introduce unacceptable delay for voice traffic.

18. A wireless LAN has mobile stations communicating with a base station. Suppose that the channel available has W Hz of bandwidth and suppose that the inbound traffic from the mobiles to the base is K times smaller than the outbound traffic from the base to the workstations. Two methods are considered for dealing with the inbound/outbound communications. In frequency-division duplexing the channel is divided into two frequency bands, one for inbound and one for outbound communications. In time-division duplexing all transmissions use the full channel but the transmissions are time-division multiplexed for inbound and outbound traffic.

Solutions follow questions:

- a. Compare the advantages and disadvantages of the two methods in terms of flexibility, efficiency, complexity, and performance.

Time Division Duplexing (TDD) is more flexible than Frequency Division Duplexing (FDD) because it can be allocated dynamically without changing hardware. FDD is less flexible because the allocation of bandwidth to channels cannot be changed. However this is not a problem if the parameter K is fixed.

TDD is a more efficient than FDD because the need for frequency guard bands is avoided.

The relative performance of TDD and FDD depends on the nature of the traffic. If the network is working with steady traffic then both systems have roughly the same performance. But if the network is working with bursty traffic then in FDD the whole bandwidth is not used, whereas in TDD the empty slots can be used by others.

FDD systems tend to be simpler because they operate at a per-station bit rate, whereas TDD systems need to operate at aggregate bit rate.

- b. How is the ratio K taken into account in the two methods?

For FDD, the amount of frequency bandwidth is allocated according the ratio K . In TDD the ratio K becomes apparent in the relative usage of time slots in either direction.

19. Consider the following variation of FDMA. Each station is allotted two frequency bands: a band on which to transmit and a band in which to receive reservations from other stations directing it to listen to a transmission from a certain station (frequency band) at a certain time. To receive a frame, a station tunes in the appropriate channel at the appropriate time. To make a reservation, a station transmits at the receiving stations reservation channel. Explain how transmitting and receiving stations can make use of the reservation channels to schedule frame transmissions.

Solution:

Suppose station A is the transmitting station, and station B is the receiving station. Each station has two unique frequency bands for transmission and receiving reservations:

- A: F1 for transmission
 F2 for receiving reservation
 B: F3 for transmission
 F4 for receiving reservation

Assume both stations have tunable transmitter and receiver, and know each other's reservation channel's frequency initially.

First, station A tunes its transmitter to F4, B's reservation channel, to notify station B the pending frame transmission and its transmission channel F1. After the reservation, station A transmits the frames on F1, and station B tunes its receiver to F1 to receive the frames. The procedure is straightforward except for the fact that contention can arise in the reservation channels when two or more stations wish to transmit to the same station. A medium access control protocol such as slotted Aloha is thus required in the reservation channel.

20. Compare FDMA, TDMA, and CDMA in terms of their ability to handle groups of stations that produce information flows that are produced at constant but different bit rates.

Solution:

Suppose the total bit rate supported by the transmission medium is R. In the typical FDMA, each station can transmit at a rate of R/M on its assigned frequency, where M is the total number of stations. The bit rate R/M is fixed for each station, and must satisfy the highest bit rate generated by the group. But for stations with lower bit rate, unused bandwidth is wasted. FDMA is inflexible and inefficient in handling flows with different bit rates. FDMA would need to be modified to allocate bands of different bandwidth to different users to accommodate differences in bit rate requirements.

In TDMA, each time slot gives an average bit rate of R/M, where M is the number of time slots. TDMA is more flexible than FDMA, because it can accommodate flows of different bit rates in two ways: assign multiple slots to each flow according to their bit rate, or allow the slot to be variable in duration. In the first method, the bit rate of the flow must be some multiple of a basic bit rate. In the second method, a means of identifying the endpoints of the variable length frame must be provided, and this adds overhead. TDMA is more flexible than FDMA in handling flows of various bit rates, but does not necessarily do this efficiently.

Unlike TDMA and FDMA where transmission bandwidth is statically allocated to different stations in terms of time or frequency, in CDMA, the signal of every flow occupies the entire frequency band at the same time. In general CDMA assumes that the bit rate of the information source is fixed, and the symbol rate of the spreading sequence is an integer multiple of the information bit rate. If the bit rate is reduced by some integer factor, for example by 2, then the integration period for each bit time (in Figure 6.33) is twice as long. This means that the signal-to-noise ratio at the receiver will be higher, and hence the transmitter can reduce its transmitted power. The capacity of the CDMA system is limited by the interference between different flows. Thus flows of lower bit rate require lower transmission power, and cause lower interference to other flows. The extra capacity can be used to handle more flows. CDMA can efficiently handle flows of different bit rates as long as the inter-flow interference is below an acceptable level.

21. Calculate the autocorrelation function of the pseudorandom sequence in Figure 6.30 as follows. Replace each 0 by -1 and each 1 by +1. Take the output sequence of the generator and shift it with respect to itself; take the product of seven (one period) symbol pairs and add. Repeat this calculation for shift values of 0, 1, ..., 7. In what sense does the result approximate the autocorrelation of a random sequence?

Solution:

$$\text{Autocorrelation } R(n) = \sum_i^m A0_i A_n_i$$

where $n = 0, \dots, 7$ is the shift value, $A0$ is the original sequence, A_n is the sequence after n shifts, and m is the number of bits in the pseudorandom sequence.

Shift to the right every time:

	1	2	3	4	5	6	7	
A0:	+1 (1)	+1 (1)	+1 (1)	-1 (0)	+1 (1)	-1 (0)	-1 (0)	R(0) = 7
A1:	-1	+1	+1	+1	-1	+1	-1	R(1) = -1
A2:	-1	-1	+1	+1	+1	-1	+1	R(2) = -1
A3:	+1	-1	-1	+1	+1	+1	-1	R(3) = -1
A4:	-1	+1	-1	-1	+1	+1	+1	R(4) = -1
A5:	+1	-1	+1	-1	-1	+1	+1	R(5) = -1
A6:	+1	+1	-1	+1	-1	-1	+1	R(6) = -1
A7:	+1	+1	+1	-1	+1	-1	-1	R(7) = 7

The normalized autocorrelation $R(n)/N$, where $N = 7$ is:

$$\frac{R(n)}{7} = \begin{cases} 1 & n = 0, 7, 14, \dots \\ -\frac{1}{7} & \text{otherwise} \end{cases}$$

This is similar to the autocorrelation of a random sequence

$$R(n) = \begin{cases} 1 & n = 0 \\ 0 & \text{otherwise} \end{cases}$$

However, the pseudorandom sequence is not truly, that is, $R(n)$ repeats itself after a period of N shifts. However as we go to longer pseudorandom sequences, i.e. N increases, $R(n)$ approaches the autocorrelation function of a truly random sequence.

22. Construct the Walsh orthogonal sequences of length 16.

Solution:

$$W_{16} = \begin{bmatrix} W_8 & W_8 \\ W_8 & W_8^c \end{bmatrix}$$

$$W_8 = \left[\begin{array}{cccc|cccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{array} \right]$$

$$W_{16} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$$

23. Decode the sum signal in Figure 6.33 using the Walsh sequence for channel 4. What do you get? Explain why.

Solution:

$$\mathbf{S}_4 = (-1, 1, 1, -1)$$

Sum signals:

$$\mathbf{r}_1 = (+1, -1, -1, -3)$$

$$\mathbf{r}_2 = (-1, +1, -3, -1)$$

$$\mathbf{r}_3 = (+1, -1, +3, +1)$$

$$\mathbf{S}_4 * \mathbf{r}_1 = -1 -1 -1 + 3 = 0$$

$$\mathbf{S}_4 * \mathbf{r}_2 = 1 + 1 -3 + 1 = 0$$

$$\mathbf{S}_4 * \mathbf{r}_3 = -1 -1 +3 -1 = 0$$

We obtain 0 for every information bit. This is because no channel 4 signal is present in the sum signal, and the Walsh sequence of channel 4 is orthogonal to other channel signals.

24. Compare IS-54 and GSM in terms of their handling of speech and the effect on spectrum efficiency.

Solution:

GSM provides higher a bit rate (22.8 kbps) to a full-rate speech channel than IS-54 which only provides 16.2 kbps. The high traffic channel bit rate decreases the spectrum efficiency of GSM.

However, much of the 22.8 kilobits are used in the error protection of the speech signal. This decreases the interference between adjacent cells and improves GSM's frequency reuse scheme. Compared to the frequency reuse factor of 7 of IS-54, GSM has a reuse factor of 3.

Both traffic channel bit rate and the frequency reuse factor affects the spectrum efficiency of the standards. In the end, GSM achieves a spectrum efficiency close to IS-54, but provides a higher speech quality.

25. Suppose that the A provider in the 800 MHz cellular band uses GSM and the B provider uses IS-95. Explain how a call from a mobile user in system B to a user in system A is accomplished.

Solution:

The GSM and IS-95 systems use different encodings for voice. At the very least, the formats between the two systems need to be converted in each direction. In practice the two systems are interconnected by the conventional telephone network, and so the signals from both systems are first converted to PCM from conventional telephony and then to the other format.

26. Suppose that a 1 MHz channel can support a 1 Mbps transmission rate. The channel is to be shared by 10 stations. Each station receives frames with exponential interarrivals and rate $\lambda = 50$ frames /second and frames are constant length $L = 1000$ bits. Compare the total frame delay of a system that uses FDMA to a system that uses TDMA.

Solution:

Given:

$$R = 1 \text{ Mbps}$$

$$L = 1000 \text{ bits}$$

$$M = 10 \text{ stations}$$

$$X = L/R = (1000 \text{ bits})/(1000000 \text{ bits/sec}) = 1 \times 10^{-3} \text{ sec}$$

$$\rho = (\lambda/M)(X \times M) = 50(10 \times 10^{-3}) = 0.5$$

For FDMA:

$$\begin{aligned} \text{Total Packet Delay} &= [\rho M / (2(1-\rho)) + M/2 + M] E[X] \\ &= [(0.5)(10) / (2(1-0.5)) + 10/2 + 10] [10^{-3}] \\ &= [20] [10^{-3}] = 0.02 \end{aligned}$$

For TDMA:

$$\begin{aligned} \text{Total Packet Delay} &= [\rho M / (2(1-\rho)) + M/2 + 1] E[X] \\ &= [(0.5)(10) / (2(1-0.5)) + 10/2 + 1] [10^{-3}] \\ &= [11] [10^{-3}] = 0.011 \end{aligned}$$

Both TDMA and FDMA are sensitive to the number of stations. We can observe that TDMA outperforms FDMA because of the faster frame transmission time.

27. Discuss how the delay and throughput performance of GPRS vary with the allocation in the number of access request channels, access grant channels, and data channels.

Solution:

Allocating more access request channels reduces the contention in each channel and results in lower delay and higher throughput. Allocating more access grant channels does not improve the performance assuming that the existing channel can service all the grants. Allocating more data channels provides more resources allowing more requests to be granted. This results in better delay and throughput performance.

28. Consider an “open concept” office where 64 carrels are organized in an 8 x 8 square array of 3m x 3m space per carrel with a 2m alley between office rows. Suppose that a conduit runs in the floor below each alley and provides the wiring for a LAN to each carrel.

Solutions follow questions:

- a. Estimate the distance from each carrel to a wiring closet at the side of the square office.

Distance from carrel to wiring room in meters:

	1	2	3	4	5	6	7	8
1	15	18	21	24	27	30	33	36
2	15	18	21	24	27	30	33	36
3	8	11	14	17	20	23	26	29
4	8	11	14	17	20	23	26	29
5	8	11	14	17	20	23	26	29
6	8	11	14	17	20	23	26	29
7	15	18	21	24	27	30	33	36
8	15	18	21	24	27	30	33	36

- b. Does it matter whether the LAN is token ring or Ethernet? Explain.

All distances to wiring room are less than 100m, well within the distance that can be handled by Ethernet.

A token ring LAN requires a wire to and from the wiring closet for each station. The length of the ring is then the sum of twice the sum of all distances from the stations to the closet: 2816 meters. It is clear that the token ring must deal with much greater distance than Ethernet.

- c. Discuss the merits of using a wireless LAN in this setting.

A wireless LAN can obviously save a lot of wiring and therefore reduce the overall cost. In addition, the distances from the stations to the base will be shorter and more uniform. However the bandwidth in a wireless LAN will be shared among all stations. In a switched Ethernet, it is possible to reduce the amount of bandwidth sharing by introducing separate collision domains.

29. Suppose that a LAN is to provide each worker in problem 6.28 with the following capabilities: digital telephone service; H.261 video conferencing capability; 250 ms retrieval time for a 1 Mbyte file from servers in the wiring closet; 10 e-mails/hour sent and received by each worker (90 percent of e-mails are short, and 10 percent contain a 100-kilobyte attachment).

Solutions follow questions:

- a. Estimate the bit rate requirements of the LAN.

Capability	Bit Rate
digital telephone service	64 kbps
H.261 videoconferencing, 2 x 64 kbps	128 kbps
large file retrieval, 8 Mb/250 ms =	32 Mbps
e-mails, $\frac{10 \times 0.1 \times 800k}{60 \text{ min/hr} \times 60 \text{ sec/min}} =$	222 bps
Total =	32+ Mbps

In this case, the fast file download requirement dominates the overall bit rate requirement. If one were to accommodate this “peak rate” requirement for all users simultaneously, then an aggregate peak bit rate of $64 \times 32 \text{ Mbps} = 2 \text{ Gbps}$ would be required. This peak bit rate can be reduced by taking into account the probability that various numbers of users simultaneously download files.

Another consideration is the impact of the peak bit rate events on the quality of the digital voice service. Unless the digital voice is given some sort of priority access, periods of high peak rate traffic will result in long delays for voice packets and associated loss in delivered voice quality.

- b. Is it worthwhile to assign the users to several different LANs and to interconnect these LANs with a bridge?

The bandwidth requirements are dominated by the fast retrieval time required for the extra large files. These retrievals are likely to be spurious in time, so consequently the bandwidth requirement for each user is extremely bursty. The digital telephone and videoconferencing applications are much lower in the bit rate requirement, but these are likely to be sustained over periods of minutes or even hours.

Suppose that 25% of the 64 users in the LAN are involved in either a digital telephone or a videoconference call. The required bit rate is then $.25 \times 64 \times 128 \text{ kbps} \approx 2 \text{ Mbps}$. Thus the users can be kept in one LAN, as long as there is enough spare bandwidth to provide the rapid response required for the retrieval of the long files.

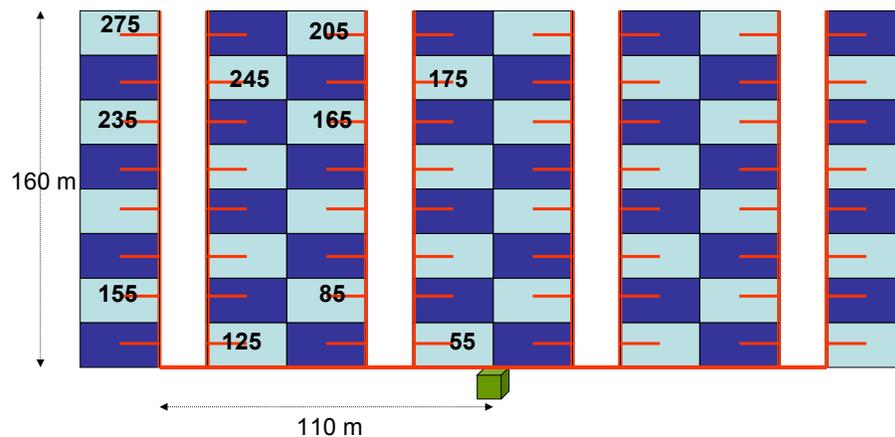
30. Consider a LAN that connects 64 homes arranged in rows of 8 homes on 20m x 30m lots on either side of a 10-meter-wide street. Suppose that an underground conduit on either side of the street connects the homes to a pedestal at the side of this rectangular array.

Solutions follow questions:

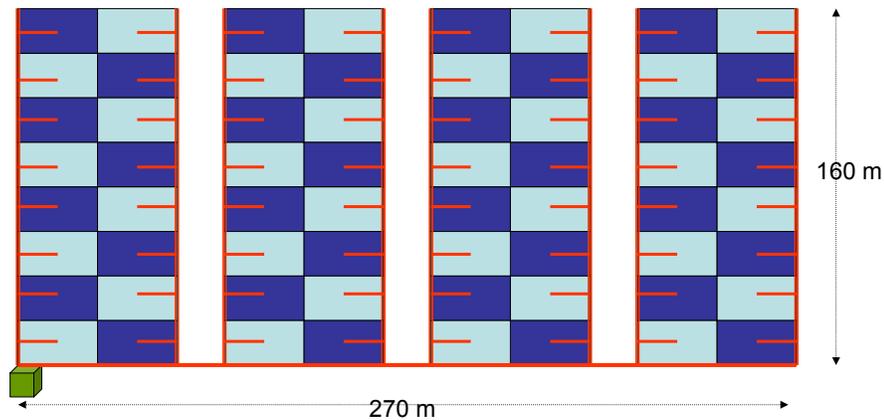
- a. Estimate the distance from each home to the pedestal.

Distance from center of homes to pedestal in meters:

Placement 1:



Placement 2:



- b. Estimate the bit rate requirements of the LAN. Assume two telephone lines, three MPEG2 televisions, and intense peak-hour Web browsing, say two Web page retrievals per minute, at an average of 20 kilobytes/page.

Web Pages: $20 \text{ Kb/page} \times (2 \text{ pages/min}) \times (1 \text{ min}/60 \text{ sec}) = 5.333 \text{ kbps}$

MPEG2: $(6 \text{ Mbps}) \times (3) = 18 \text{ Mbps}$

Telephone: $(64 \text{ Kbps}) \times (2) = 128 \text{ kbps}$

The total bit rate should be higher than $(18 \text{ Mbps}) \times 64 \text{ homes} = 1,152 \text{ Mbps}$

The bit rate is dominated by the MPEG2 streams. Let's examine the web page download requirement more closely. The average total bit rate for the 64 homes is 341 kbps which is quite low.

The problem does not consider the impact of peer-to-peer file sharing. Suppose one quarter of the homes are downloading 1 Mbyte files at the same time and that we would like each file to be downloaded in 1 minute. The required bit rate is: $64 \times \frac{1}{4} \times 8 \times 10^6 / 60 = 2 \text{ Mbps}$.

- c. Can a single LAN meet the service requirements of the group of 64 homes? Explain.

The average bit rate required for each home is about 20 Mbps. This bit rate can be met with a 100 Mbps Ethernet link that connects each home to an Ethernet switch at the pedestal, although some sort of priority mechanism is required for the voice traffic. The switch however needs to be capable of handling more than 1 Gbps aggregate throughput. This can be met by providing several 1 Gbps links into the network backbone.

31. Use HDLC and Ethernet to identify three similarities and three differences between medium access control and data link control protocols. Is HDLC operating as a LAN when it is used in normal response mode and multipoint configuration?

Solution:

Similarities: Both medium access control and data link protocols deal with the transfer of blocks of information across a "single hop" which can consist of a point-to-point line, a multipoint line, or a shared medium.

The frames in HDLC and Ethernet contain framing information that delineates the beginning and end of each frame. HDLC inserts flag bits, and Ethernet inserts a preamble to delineate their frames.

Both data link control and MAC protocols check the CRC in the received frames for errors. Frames with errors are discarded.

Differences: The main functionality of the MAC protocols is to coordinate the access of multiple nodes to a shared transmission medium. To do so, the MAC protocols implement contention resolution procedures.

The data link layer protocols implement error control and flow control functions to provide reliable transmission. The MAC protocols do not implement these functions.

The data link layer protocols such as HDLC are concerned with providing connection-oriented, connectionless, and acknowledged connectionless packet transfer services to the network layer. The MAC protocols only provide connectionless packet transfer.

The data link control protocol often handles packet transfer only over a point-to-point link between two nodes, where medium access control is not required. However, the multipoint-to-point configuration is an exception. When HDLC operates in normal response mode, it in fact is implementing polling to control access to the medium.

32. An application requires the transfer of network layer packets between clients and servers in the same LAN. Explain how reliable connection-oriented service can be provided over an Ethernet LAN. Sketch a diagram that shows the relationship between the PDUs at the various layers that are involved in the transfer.

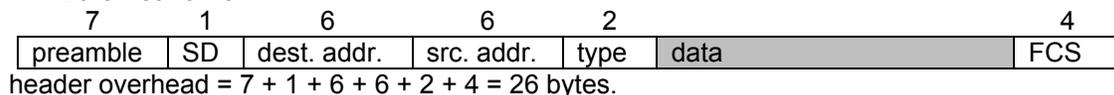
Solution:

A connection-oriented service has to be built on top of the connectionless service provided by an Ethernet LAN. The IEEE Logical Link Control layer can implement an HDLC-like error control protocol to provide this service. The relationship between layers is shown in Figure 6.48 in the text. The relationship between SDUs and PDUs can be visualized using Figure 6.50: the network layer SDU is passed to the LLC which prepares a PDU suitable for error control. The LLC PDU is then passed to the Ethernet layer.

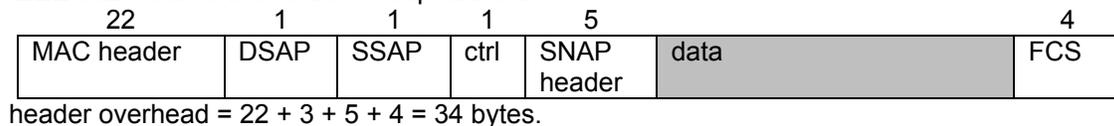
33. Calculate the difference in header overhead between a DIX Ethernet frame and an IEEE 802.3 frame with SNAP encapsulation.

Solution:

DIX Ethernet frame:



IEEE 802.3 frame with SNAP encapsulation:



34. Suppose that a group of 10 stations is serviced by an Ethernet LAN. How much bandwidth is available to each station if (a) the 10 stations are connected to a 10 Mbps Ethernet hub; (b) the 10 stations are connected to a 100 Mbps Ethernet hub; (c) the 10 stations are connected to a 10 Mbps Ethernet switch.

Solution:

- Assuming essentially 100% efficiency, the 10 Mbps are shared equally by the 10 stations, so each station can receive a maximum of 1 Mbps on average.
- Assuming essentially 100% efficiency, the 100 Mbps are shared equally by the 10 stations, so each station receives a maximum of 10 Mbps on average.
- The bit rate available to each station depends on the number of collision domains that are configured in the switch. In the best case, each station has nearly 10 Mbps to the Ethernet switch. Each station will have full access to the 10 Mbps if the switch capacity can handle the aggregate rate from all the stations.

35. Suppose that an Ethernet LAN is used to meet the requirements of the office in problem 3.28.

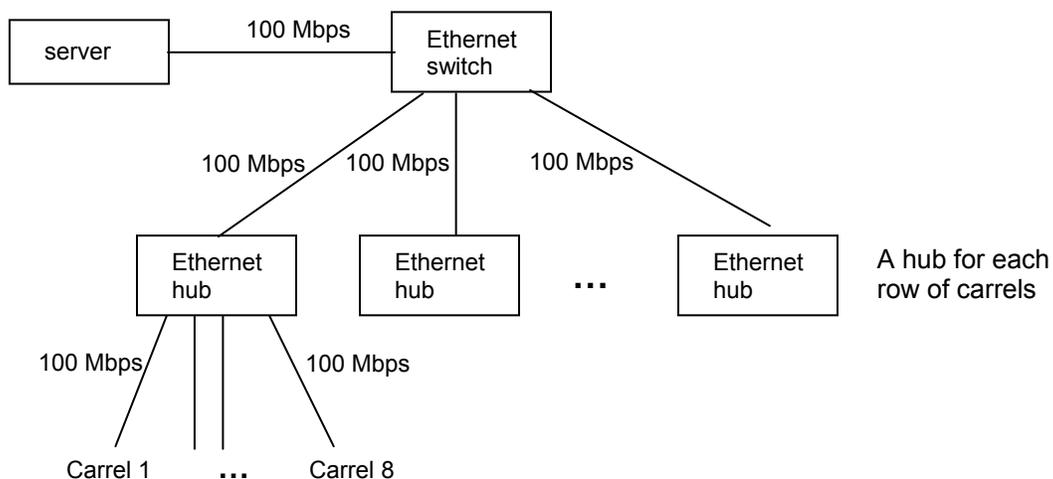
Solutions follow questions:

- Can the requirements of one row of carrels be met by a 10 Mbps Ethernet hub? By a 10 Mbps Ethernet switch?

No. The bit requirement for each worker is 33.6 Mbps (for example, a 1 Mbyte file retrieved in 250 ms). The maximum rates of both Ethernet hub and switch are both 10 Mbps. At 10 Mbps, the file transfer would take 800 ms.

- Can the requirements of the office be met by a hierarchical arrangement of Ethernet switches as shown in Figure 6.57?

For the arrangement shown below, each group of 8 users in a row shares 100 Mbps. This would allow up to three users to be downloading a file and still meet the 250 ms file download requirement.



36. Suppose that 80 percent of the traffic generated in the LAN is for stations in the LAN, and 20 percent is for stations outside the LAN. Is an Ethernet Hub preferable to an Ethernet switch? Does the answer change if the percentages are reversed?

Solution:

The difference between a hub and a switch is that in the hub frame are broadcast to all lines, while in a switch, frames are forwarded to another collision domain only if the destination is in that domain.

When 80% of the traffic is local, the switch will only forward 20% of frames to other collision domains, thus increasing the bandwidth available on those domains. If 80% of the traffic is to other collision domains, then the switch forward more traffic to other domains. If there is only one other domain, then the switch forwards almost as much traffic as a hub would and is thus ineffective in containing broadcast traffic. On the other hand, if the switch has multiple broadcast domains, then the amount of traffic forwarded from the switch will be less than that forwarded by a hub.

37. Calculate the parameter a and the maximum throughput for a Gigabit Ethernet hub with stations at a 100-meter distance and average frame size of 512 bytes; 1500 bytes; and 64,000 bytes.

$$d = 100$$

$$T_{\text{prop}} = 0.0000005$$

$$R = 1.00\text{E}+09$$

	512	1500	64000
a	0.12207	0.041667	0.000977
Throughput	0.56	0.788525	0.993754

38. Provide a brief explanation for each of the following questions:

- a. Under a light load, which LAN has a smaller delay: Ethernet or token ring?

Ethernet has smaller delay under a light load. In Ethernet under a light load, there is little or no contention for the channel, the delay incurred is close to the frame transmission time. In token ring, however, there is always the additional delay incurred from circulating the token around the ring.

- b. Under a high load, which LAN has a smaller delay: Ethernet or token ring?

Token ring has smaller delay under a high load. In Ethernet there is more contention for the channel, much of the time is spent in collision and backoff, so on average the frames experience longer delay and higher delay variability. In comparison, token ring provides each station with an orderly and round-robin access to the channel by passing the token around. When the number of frames transmitted per token is limited, and frames are kept at a fixed length, token ring can guarantee a maximum delay for each station.

39. Suppose that a token-ring LAN is used to meet the requirements of the office in problem 6.29

- a. Calculate the ring latency if all carrels are to be connected in a single ring as shown in Figure 6.58. Repeat for a ring for a single row of carrels.

$$M = 64$$

$$b = 2.5$$

$$d = 2816 \text{ m}$$

$$\tau' = (d/v) + M(b/R) = (2816/2 \times 10^8) + 64(2.5/R) = 14 \times 10^{-6} + (160/R)$$

$$\text{If } R = 16 \text{ Mbps}$$

$$\tau' = 14 \times 10^{-6} + (160/32 \times 10^6) = 14 \times 10^{-6} + 5 \times 10^{-6} = 24 \times 10^{-6} = 24 \mu\text{sec}$$

$$\text{If } R = 32 \text{ Mbps}$$

$$\tau' = 14 \times 10^{-6} + (160/32 \times 10^6) = 14 \times 10^{-6} + 5 \times 10^{-6} = 19 \times 10^{-6} = 19 \mu\text{sec}$$

$$M = 8$$

$$b = 2.5$$

$$d = 408 \text{ m}$$

$$\tau' = (d/v) + M(b/R) = (408/2 \times 10^8) + 8(2.5/R) = 2 \times 10^{-6} + (20/R)$$

If $R = 16$ Mbps

$$\tau' = 2 \times 10^{-6} + (20/32 \times 10^6) = 2 \times 10^{-6} + 0.6 \times 10^{-6} = 3.2 \times 10^{-6} = 3.2 \mu\text{sec}$$

$R = 32$ Mbps

$$\tau' = 2 \times 10^{-6} + (20/32 \times 10^6) = 2 \times 10^{-6} + 0.6 \times 10^{-6} = 2.6 \times 10^{-6} = 2.6 \mu\text{sec}$$

b. Can the requirements of one row of carrels be met by a 16 Mbps token ring?

The file download transfer time of 250 ms cannot be met by a 16 Mbps token ring. At this speed the download time is 500 ms. On the other hand, the requirements of the other applications can be met by a 16 Mbps ring.

c. Can the requirements of the office be met by an FDDI ring?

For FDDI, $R = 100$ Mbps and so

$$\tau' = 14 \times 10^{-6} + (160/100 \times 10^6) = 14 \times 10^{-6} + 1.6 \times 10^{-6} = 15.6 \times 10^{-6} = 15.6 \mu\text{sec}.$$

At $R = 100$ Mbps, up to three stations out of 64 can be busy downloading files while still meeting the 250 ms requirement. The other requirements of the office are met easily.

40. Suppose that a group of 32 stations is serviced by a token-ring LAN. For the following cases calculate the time it takes to transfer a frame using the three token reinsertion strategies: after completion of transmission; after return of token; after return of frame.

- 1000 bit frame; 10 Mbps speed; 2.5-bit latency per adapter; 50 meters between stations.
- Same as part (a) except 100 Mbps speed and 8-bit latency/adapter.
- Same as part (a) except 1 km distance between stations.

Solution:

$$M = 32$$

$$v = 2 \times 10^8 \text{ m/sec}$$

$$L = 1000 \text{ bits}$$

I. Reinsertion after completion of transmission

Assume the token transmission time is small and can be ignored. The one-hop walk time is for passing the token to the next station.

$$X_{\text{eff}} = \text{frame transmission time} + \text{one hop walk time}$$

$$= L/R + (d/v + b/R)$$

$$\text{a. } R = 10 \text{ Mbps, } b = 2.5 \text{ bits, } d = 50 \text{ m}$$

$$X_{\text{eff}} = 100.5 \mu\text{sec}$$

$$\text{b. } R = 100 \text{ Mbps, } b = 8 \text{ bits, } d = 50 \text{ m}$$

$$X_{\text{eff}} = 10.3 \mu\text{sec}$$

$$\text{c. } R = 10 \text{ Mbps, } b = 2.5 \text{ bits, } d = 1 \text{ km}$$

$$X_{\text{eff}} = 105 \mu\text{sec}$$

II. Reinsertion after return of token

$$X_{\text{eff}} = \max\{\text{frame transmission time, ring latency}\} + \text{one hop walk time}$$

$$= \max\{L/R, \tau'\} + (d/v + b/R)$$

$$\tau' = M \left\{ \frac{d}{v} + \frac{b}{R} \right\}$$

- a. $X_{\text{eff}} = 101 \mu\text{sec}$
- b. $X_{\text{eff}} = 11 \mu\text{sec}$
- c. $X_{\text{eff}} = 173 \mu\text{sec}$

III. Reinsertion after return of frame

$$X_{\text{eff}} = \text{frame transmission time} + \text{ring latency} + \text{one hop walk time}$$

$$= L/R + \tau' + (d/v + b/R)$$

- a. $X_{\text{eff}} = 117 \mu\text{sec}$
- b. $X_{\text{eff}} = 21 \mu\text{sec}$
- c. $X_{\text{eff}} = 273 \mu\text{sec}$

41. Suppose that an FDDI LAN is used to meet the packet voice requirements of a set of users. Assume voice information uses 64 kbps coding and that each voice packet contains 20 ms worth of speech.

- a. Assume that each station handles a single voice call and that stations are 100 meters apart. Suppose that the FDDI ring is required to transfer each voice packet within 10 ms. How many stations can the FDDI accommodate while meeting the transfer requirement?

packet length = 20 ms x 64 kbps = 1280 bits

header length = 28 x 8 = 224 bits

$X_{\text{frame}} = (1280 + 224) \text{ bits} / 100 \text{ Mbps} = 15.04 \mu\text{sec}$

token length = 11 bytes x 8 = 88 bits

$X_{\text{token}} = 88 \text{ bits} / 100 \text{ M} = 0.88 \mu\text{sec}$

$d = 100 \text{ m}$ between the stations

$b = 10 \text{ bits}$

$v = 2 \times 10^8 \text{ m/sec}$

ring length = 100m \leq 200 km (upper limit)

Every station handles a single call, and each voice packet must be transmitted within 10ms. This means the target token rotation time (TTRT) must be within $\frac{1}{2} \times 10\text{ms} = 5 \text{ ms}$.

TTRT = ring latency + M stations transmitting 1 busy token + 1 frame + 1 free token

$$= \left(\frac{100M}{2 \times 10^8} + \frac{10M}{100 \times 10^6} \right) + M(2 \times 0.88\mu + 15.04\mu)$$

$$= M(0.6 + 1.76 + 15.04) \times 10^6 = 1.74 \times 10^{-5} M \leq 0.005$$

$$M \leq 287$$

Therefore, the FDDI can accommodate maximum 287 stations.

- b. How many simultaneous calls can be handled if each station is allowed to handle up to 8 calls?

In part (a), note that the first component of the TTRT (0.6 μ sec) is due to propagation delay, the second component (1.76 μ sec) is due to station latency, and the last component is the frame time (15.04 μ sec). When a station handles multiple voice calls, we can imagine each station as sending 8 voice packets per token opportunity, so in effect the frame transfer time becomes $8 \times 15.04 = 120.32$ μ sec. To keep the TTRT to less than 5 ms, we then require that

$$M(0.6 + 1.76 + 120.32) < 5000$$

The maximum number of stations is 40, and so the maximum number of calls is $8 \times 40 = 320$.

42. Use IEEE 802.3 and IEEE 802.11 to discuss three differences between wired and wireless LANs.

Solution:

Error rate: Unlike wired LANs, wireless LANs have high error rate due to interference and noise. Wireless LANs need to implement ARQ and/or error correction to increase the reliability of the communication channel.

Station mobility: Unlike wired LANs where stations connected to the LANs are static, in wireless LANs, the stations can be mobile and portable. Wireless LAN protocols may have to implement dynamic traffic routing and service handoff when the station moves from one service area to another.

Collision detection: Collision detection is not effective in wireless LANs due to the hidden station problem. Consequently, the sender must wait for explicit acknowledgment (e.g. RTS/CTS) from the receiver to know whether or not a frame has been received. The wireless LAN protocol implements a collision avoidance algorithm rather than the collision detection in wired LAN, and the delay in the contention period is longer than three round-trip delay of $2t_{prop}$ of wired LAN because of waiting for the receiver's acknowledgment.

Other differences:

Security: In a wired LAN, the transmission medium is usually physically secure. In a wireless LAN, any device within the geographic transmission area can intercept the transmissions. To provide data security, wireless LANs need to implement encryption at the expense of higher cost and reduced performance.

Power consumption: Portable and mobile devices are usually battery powered, and thus have limited power capacity. The wireless LAN protocol must be designed to be power efficient.

All these issues are addressed in the IEEE 802.11 wireless LAN protocol.

43. For data packet radio networks, discuss the advantages and disadvantages of providing reliability by (a) implementing error correction at the physical layer, (b) implementing error control as part of the MAC layer, and (c) implementing error control at the LLC layer.

Solution:

Error correction at the physical layer uses up bandwidth in the form of check bits and adds complexity in terms of hardware for detecting and correcting errors. If the bit error rate is low, ARQ at a higher layer may be simpler and more efficient. However, if the bit error rate is high, error correction may become essential to be able to communicate at all over the radio medium.

Error control at the MAC layer is the preferred approach if the error rate is not too high. The number of check bits required for error detection is less than for error correction. Bandwidth is "wasted" only

on retransmissions, not in a large number of check bits in every transmission. The implementation of ARQ is much simpler than complex error correction.

Error control at the LLC layer provides flexibility when operating over a MAC sublayer. The MAC sublayer can provide connectionless service and the LLC can add reliability for those network layers that require it.

44. Consider the distributed coordination function in IEEE 802.11. Suppose that all packet transmissions are preceded by a RTS-CTS handshake. Find the capacity of this protocol following the analysis used for CSMA-CD.

Solution:

In IEEE 802.11 DCF with RTS-CTS handshake, the sender stations contend for the channel by sending a RTS frame to the receiver, and it successfully captures the channel only when it receives a CTS frame from the receiver.

The sender does not know if it has succeeded until time = DIFS + X_{RTS} + SIFS + X_{CTS}. By this time, if no CTS frame arrives, the sender knows it has failed and will execute a backoff. This duration is similar to 2t_{prop} in CSMA-CD. Therefore, in CSMA-CA, the time can be divided in contention slots of size (DIFS + X_{RTS} + SIFS + X_{CTS}).



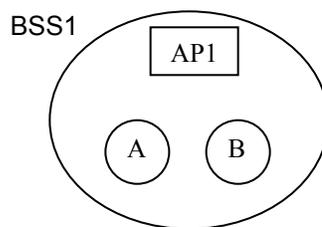
Similar to CSMA-CD, suppose all stations are contending for the channel and have a probability of p to transmit RTS during a contention slot. Then the maximum probability of success is $1/e$ as the number of station increases, and the average # of contention slots in a contention period is e .

The maximum throughput occurs when all of the channel time is spent in transmission period followed by contention intervals.

$$\rho_{\max} = \frac{X_{data}}{X_{data} + 2SIFS + X_{ack} + e(DIFS + X_{rts} + SIFS + X_{cts})}$$

45. Suppose one station sends a frame to another station in an IEEE 802.11 ad hoc network. Sketch the data frame and the return ACK frame that are exchanged, showing the contents in the relevant fields in the headers.

Solution:



Suppose station A sends to station B without replaying through AP1.

Data Frame:

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	$X_{data} + SIFS + X_{ack}$	B	A	BSS1	—

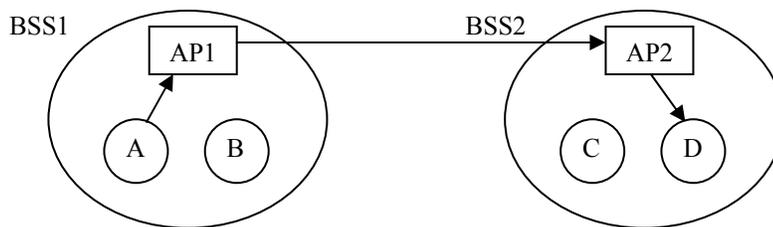
prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	data	data	0	0	0	—	—	—	—	—

ACK Frame:

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	0	A	B	BSS1	—

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	ctrl.	ack	0	0	0	—	—	—	—	—

46. Suppose one station sends a frame to another station in a different BSS in an IEEE 802.11 infrastructure network. Sketch the various data frames and ACK frames that are exchanged, showing the contents in the relevant fields in the headers.



Solution:

Suppose station A sends frame to station D. To exchange frames between BSS, the frames are relayed through the APs.

Data Frames:

1. A → AP1

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	$X_{data} + SIFS + X_{ack}$	BSS1	A	D	—

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	data	data	1	0	0	—	—	—	—	—

2. AP1 → AP2

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control	Addr. 4
	$X_{data} + SIFS + X_{ack}$	AP2	AP1	D	—	A

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	data	data	1	1	0	—	—	—	—	—

3. AP2 → D

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	$X_{data} + SIFS + X_{ack}$	D	BSS2	A	—

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	data	data	0	1	0	—	—	—	—	—

ACK Frames:

1. D → AP2

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	0	BSS2	D	A	—

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	ctrl	ack	1	0	0	—	—	—	—	—

2. AP2 → AP1

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control	Addr. 4
	0	AP1	AP2	A	—	D

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	ctrl	ack	1	1	0	—	—	—	—	—

3. AP1 → A

Frame control	Duration/ID	Addr. 1	Addr. 2	Addr. 3	Sequence control
	0	A	BSS1	D	—

prot. ver.	type	sub-type	To DS	From DS	More flag	Retry	Pwr mgt.	More data	WEP	Rsvd
00	ctrl	ack	0	1	0	—	—	—	—	—

47. Why is error control (ARQ and retransmission) included in the MAC layer in IEEE 802.11 and not in IEEE 802.3?

Solution:

Error control is required in the MAC in IEEE 802.11 because of the noise and interference in the wireless medium. In contrast, the wired medium for IEEE 802.3 protocol has very low error rates.

48. Consider the exchange of CSMA-CA frames shown in Figure 6.72. Assume the IEEE 802.11 LAN operates at 2 Mbps using frequency-hopping physical layer. Sketch a time diagram showing the frames transmitted including the final ACK frame. Show the appropriate interframe spacings and NAV values. Use Table 6.6 to obtain the appropriate time parameters. Assume that the data frame is 2000 bytes long.

Solution:

Assume the data frame includes the MAC header and the CRC field.

$$X_{\text{frame}} = (2000 \times 8) / 2M = 8 \times 10^{-3} \text{ sec}$$

Data frame transmission time

$$X_{\text{data}} = X_{\text{frame}} + X_{\text{preamble}} + X_{\text{PLCPheader}} = 8 \times 10^{-3} + (96 + 32) \times 10^{-6} = 8.128 \times 10^{-3} \text{ sec}$$

$$X_{\text{rts}} = (20 \times 8) / 2M = 80 \mu\text{sec}$$

$$X_{\text{cts}} = (14 \times 8) / 2M = 56 \mu\text{sec}$$

The ACK frame consists of the MAC header (30 bytes) and the CRC field (4 bytes). Its total length is 34 bytes.

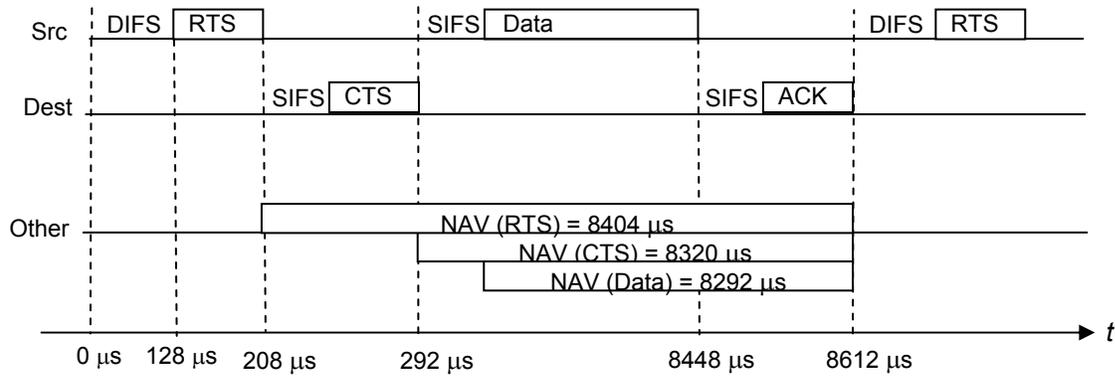
$$X_{\text{ack}} = (34 \times 8) / 2M = 136 \mu\text{s}$$

$$\text{slot time} = 50 \mu\text{s}$$

$$\text{SIFS} = 28 \mu\text{s}$$

$$\text{PIFS} = \text{SIFS} + 1 \text{ slot time} = 78 \mu\text{s}$$

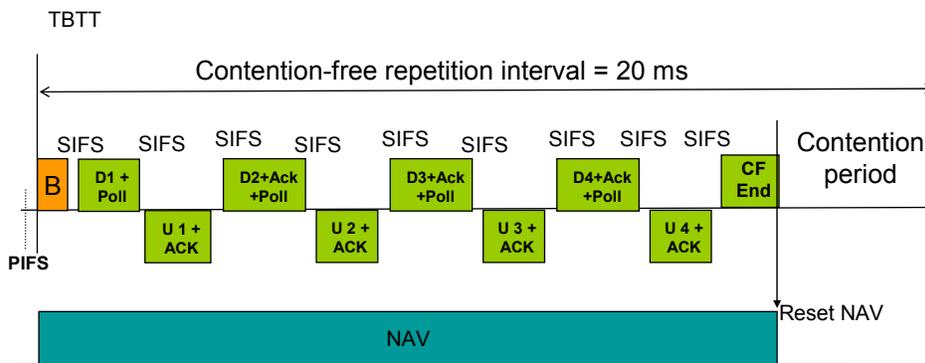
$$\text{DIFS} = \text{PIFS} + 1 \text{ slot time} = 128 \mu\text{s}$$



49. Suppose that four stations in an IEEE 802.11 infrastructure network are in a polling list. The stations transmit 20 ms voice frames produced by 64 kbps speech encoders. Suppose that the contention-free period is set to 20 ms. Sketch a point-coordination frame transfer with the appropriate values for interframe spacings, NAV, and data and ACK frames.

Solution:

The point-coordination frame transfer is shown in below.



- D1, D2, D3, D4 = frames sent by point coordinator
- U1, U2, U3, U4 = frames sent by polled stations
- TBTT = target beacon transmission time
- B = beacon frame

Every 20 ms, each voice source produces $64 \times 10^3 \times 20 \text{ ms} = 1280$ bits, which forms the payload for each voice frame. This payload is carried in a MAC frame that includes 30 bytes of header and 4 bytes of CRC, for a total of 272 bits of overhead. Assuming the parameters used in Problem 6.48, the bit rate is 2 Mbps, and the frame transmission time is:

$$X_{\text{frame}} = (1280 + 272) / 2 \text{ Mbps} = 776 \text{ } \mu\text{sec}.$$

Assume that each frame sent by the point coordinator also contains a voice frame. In this case, frames from the point coordinator have the same duration as frames from the users. Each MAC frame is transmitted within a PLCP frame of total duration:

$$X_{\text{data}} = 776 + 96 + 32 \text{ } \mu\text{sec} = 904 \text{ } \mu\text{sec}.$$

Each SIF interval is 28 μsec , and the CF control frame and the Beacon frame are 34 bytes long and of duration

$$X_{cf} = X_{beacon} = 34 \times 8/2\text{Mbps} = 136 \mu\text{sec.}$$

Therefore the duration of the contention free interval is:

$$8 \times X_{data} + 9X_{sifs} + X_{cf} + X_{beacon} = 8 \times 904 + 9 \times 28 + 2 \times 136 = 7756 \mu\text{sec.}$$

50. Can a LAN bridge be used to provide the distribution service in an IEEE 802.11 extended service set? If so, explain how the service is provided and give an example of how the frames are transferred between BSSs.

Solution:

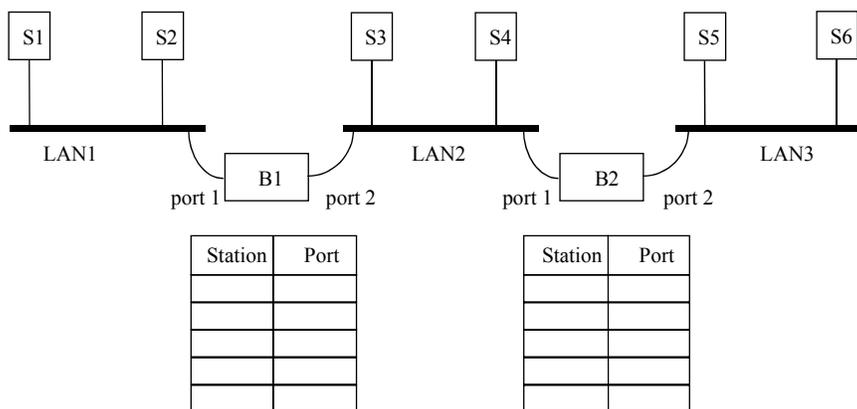
Yes, a LAN bridge can be used to provide the distribution service in an IEEE 802.11 LAN. A bridge is associated with each Access Point. The bridge learns the MAC addresses that should be forward on each port into the distribution system. The 802.11 frames emerging out of a BSS and onto a port must be encapsulated in a layer 2 tunnel and transferred to the input port of a bridge associated with the destination AP. The IEEE 802.11 frame is recovered at the destination AP and forwarded to the appropriate station.

51. Can a router be used to provide the distribution service in an IEEE 802.11 extended service set? If so, explain how addressing is handled and give an example of how the frames are transferred between BSSs.

Solution:

A router can be used to provide the distribution service in an IEEE 802.11 LAN. A bridge is associated with each Access Point. The bridge learns the MAC addresses that should be forward on each port into the distribution system. The 802.11 frames emerging out of a BSS and onto a port must be encapsulated in a layer 3 tunnel and transferred across an IP network to the input port of a bridge associated with the destination AP. The IEEE 802.11 frame is recovered at the destination AP and forwarded to the appropriate station.

52. Six stations (S1-S6) are connected to an extended LAN through transparent bridges (B1 and B2), as shown in the figure below. Initially, the forwarding tables are empty. Suppose the following stations transmit frames: S2 transmits to S1, S5 transmits to S4, S3 transmits to S5, S1 transmits to S2, and S6 transmits to S5. Fill in the forwarding tables with appropriate entries after the frames have been completely transmitted.

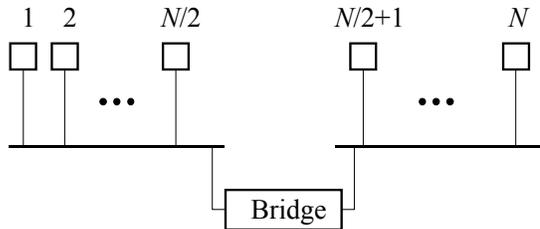


Solution:

Station	Port
S2	1
S5	2
S3	2
S1	1

Station	Port
S2	1
S5	2
S3	1
S1	1
S6	2

53. Suppose N stations are connected to an extended Ethernet LAN, as shown below, operating at the rate of 10 Mbps. Assume that the efficiency of each Ethernet is 80 percent. Also assume that each station transmits frames at the average rate of R bps, and each frame is equally likely to be destined to any station (including to itself). What is the maximum number of stations, N , that can be supported if R is equal to 100 kbps? If the bridge is replaced with a repeater, what is the maximum number of stations that can be supported? (Assume that the efficiency of the entire Ethernet is still 80 percent.)



Solution:

10 Mbps LAN
 $R = 100$ kbps
 efficiency = 0.80

Bridge:

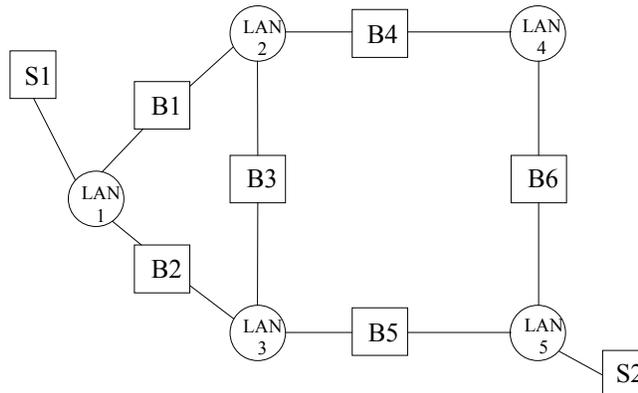
Each station is equally likely to transmit to any other stations in the extended LAN. Then $R/2$ traffic is local.

$$\frac{100\text{kbps} \left(\frac{1}{2}\right) N}{10\text{Mbps}} \leq 0.8 \quad N \leq 160$$

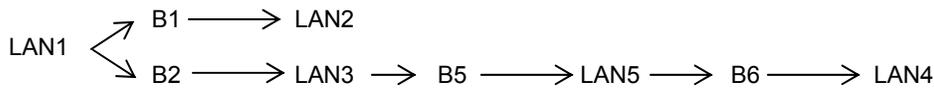
Repeater:

$$\frac{100\text{kbps} N}{10\text{Mbps}} \leq 0.8 \quad N \leq 80$$

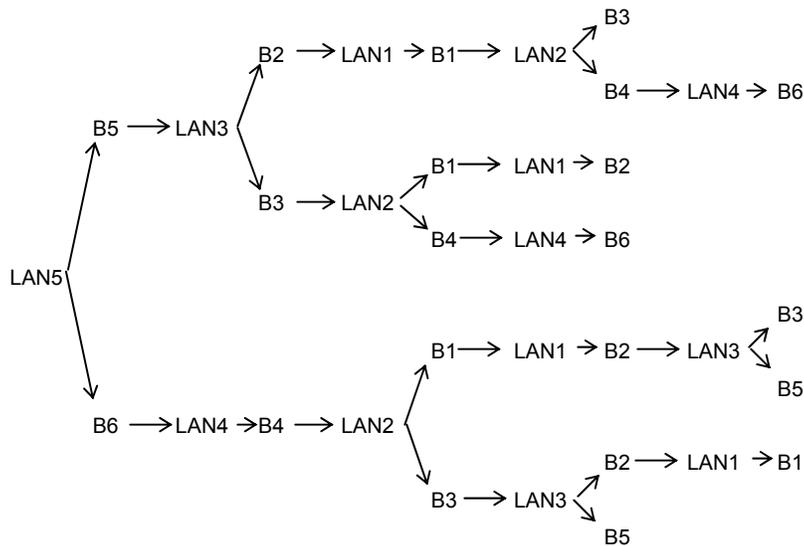
54. The LANs in the figure below are interconnected using source routing bridges. Assume that bridges 3 and 4 are not part of the initial spanning tree.



a. Show the paths of the single route broadcast frames when S1 wants to learn the route to S2.



b. Show the paths of all routes broadcast frames returned by S2.



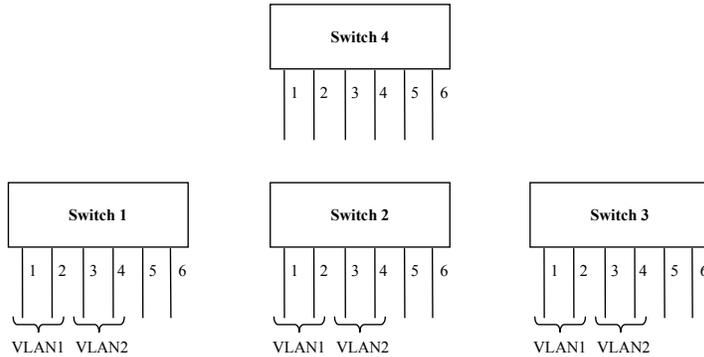
c. List all possible routes from S1 to S2 from part (b).

- LAN1 → B2 → LAN3 → B5 → LAN5
- LAN1 → B1 → LAN2 → B3 → LAN3 → B5 → LAN5
- LAN1 → B1 → LAN2 → B4 → LAN4 → B6 → LAN5
- LAN1 → B2 → LAN3 → B3 → LAN2 → B4 → LAN4 → B6 → LAN5

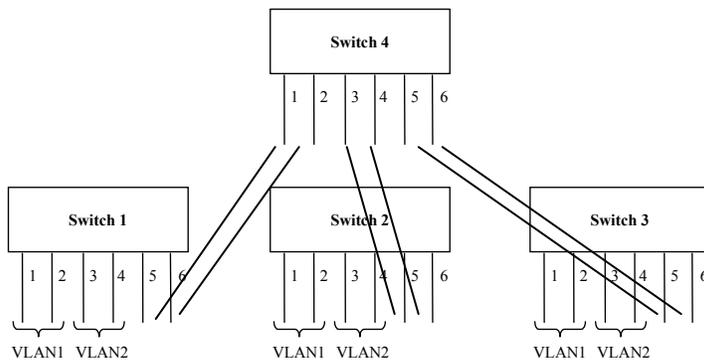
d. How many LAN frames are required to learn the possible routes?

Four frames.

55. Ports 1 through 4 in switches 1, 2, and 3, as shown in the figure below, provide port-based VLAN connectivity to stations attached to these ports. A network administrator needs to interconnect the other ports and configure the switches so that stations of the same VLAN connected to different switches can communicate. Use the association notation: switch a, VLAN n, port x, port y, ..., port z, to indicate that ports x, y, ..., z are associated to VLAN n in switch a, and the connectivity notation: (switch a, port x) to (switch b, port y) to indicate that port x in switch a is connected to port y in switch b. Show the association and connectivity for all the switches in the figure.



Solution:



- Switch 1, VLAN 1, Port 1, Port 2, Port 5
- Switch 1, VLAN 2, Port 3, Port 4, Port 6
- Switch 2, VLAN 1, Port 1, Port 2, Port 5
- Switch 2, VLAN 2, Port 3, Port 4, Port 6
- Switch 3, VLAN 1, Port 1, Port 2, Port 5
- Switch 3, VLAN 2, Port 3, Port 4, Port 6
- Switch 4, VLAN 1, Port 1, Port 3, Port 5
- Switch 4, VLAN 2, Port 2, Port 4, Port 6

Solutions to Chapter 7

7.1. Explain how a network that operates internally with virtual circuits can provide connectionless service. Comment on the delay performance of the service. Can you identify inefficiencies in this approach?

Solution:

The connection-oriented (co) network can present a network sublayer interface to the upper layer to make itself appear as a connectionless (cl) network. The upper layer sends packets in connectionless fashion, but the network resolves the packet destination's cl address into the corresponding co address (e.g. ATM address), and then establishes a virtual circuit between the source and destination. Any subsequent packets with the same source and destination are transmitted through this virtual circuit. The interface hides the internal operation of the network from the upper layer.

The cl-to-co address resolution is performed for the transmission of every packet, and hence incurs extra processing delay. Address caching can speed up the process, but cannot eliminate the delay. The first cl packet incurs extra delay because of the time required to set up the virtual circuit.

This cl-service-over-co-network approach is inefficient in that the QoS support of the connection-oriented network cannot be made available to the network layer.

ATM LAN emulation provides an example that involves providing connectionless service over a connection-oriented service.

7.2. Is it possible for a network to offer best-effort connection-oriented service? What features would such a service have, and how does it compare to best-effort connectionless service?

Solution:

Best-effort connection-oriented service would involve the transfer of packets along a pre-established path in a manner that does not provide mechanisms for dealing with the loss, corruption or misdelivery of packets. Best-effort connection-oriented service would require some means for establishing a path prior to the transfer of packets. Best-effort connectionless service would involve the transfer of packets in a datagram fashion, where routing decisions are made independently for each packet.

The path setup requirement makes connection-oriented service more complex than connectionless service. On the other hand, once a path is established, less processing is required to decide how a packet is to be forwarded. Connectionless service is more robust than connection-oriented service since connectionless service readily reroutes packets around a failure while VC service requires that new paths be established.

7.3. Suppose a service provider uses connectionless operation to run its network internally. Explain how the provider can offer customers reliable connection-oriented network service.

Solution:

To provide connection-oriented network service, an upper sublayer in the network layer at the edge of the network can establish logical connections across the connectionless network by setting up state information (for example, packet sequence number). The logical connection must be set up before packets can be transported, and each packet is assigned a sequence number. Using the sequence number, the upper sublayer entities can acknowledge received packets, determine and retransmit lost packets, delete duplicate packets, and rearrange out-of-order packets, hence providing reliable connection-oriented network service.

7.4. Where is complexity concentrated in a connection-oriented network? Where is it concentrated in a connectionless network?

Solution:

The complexity in connection-oriented networks revolves around the need to establish and maintain connections. Each node must implement the signaling required by the connection establishment process; each node must also maintain the state of the node in terms of connections already established and transmission resources available to accommodate new connections. End systems must be capable of exchanging signaling information with the network nodes to initiate and tear down connections. A connection oriented network must also include routing to select the paths for new connections.

Connectionless networks only require that nodes forward packets according to its routing tables. End systems only need to place network address information in the packet headers. The complexity of connectionless networks revolves around routing tables. Routing tables may be static and set up by a network administrator, or they may be dynamic and involve processing and exchange of link state information among nodes.

7.5. Comment on the following argument: Because they are so numerous, end systems should be simple and dirt cheap. Complexity should reside inside the network.

Solution:

This argument holds only if the computing resources are scarce and expensive as in the early days of computers. For example, telephone networks and terminal-oriented computer networks were designed based on this principle. Keeping the complexity inside the network is fine as long as the network operations are simple enough to keep the network manageable. But as the number of applications increases and the number of end systems grow, the complexity of the network increases drastically. In this case, keeping the complexity inside the network may not be scalable.

When the computing resources become plentiful and cheap, more applications and sophisticated communication services are demanded by the numerous end systems. This cannot be provided efficiently by one big monolithic network. Instead, by moving some of the complexity to the network edge and by utilizing the copious end-system computing resources, more and better network services can be deployed at a faster rate.

7.6. In this problem you compare your telephone demand behavior and your web demand behavior.

Solutions follow questions:

- (a) Arrival rate: Estimate the number of calls you make in the busiest hour of the day; express this quantity in calls/minute. Service time: Estimate the average duration of a call in minutes. Find the load that is given by the product of arrival rate and service time. Multiply the load by 64 kbps to estimate your demand in bits/hour.

One possibility:

During the busiest hour of the day

5 calls/hour = 5 calls/hour

Average call duration = 7 minutes/call

$5 \times 7 = 35$ minute/hour = 2100 sec/hour

Demand = 2100 sec/hour x 64 kbps = 134.4 Mbits/hour

- (b) Arrival rate: Estimate the number of Web pages you request in the busiest hour of the day. Service time: Estimate the average length of a Web page. Estimate your demand in bits/hour.

One possibility:

During the busiest hour of the day:

60 Web pages / hour

Average Web page size = 20 kbytes

Demand = 20k x 8 x 60 = 9.6 Mbits/hour

Note that if web pages are heavy in terms of graphics content, say 200 kbytes/page, then the demand can approach that of telephone calls.

- (c) Compare the number of call requests/hour to the number of web requests/hour. Comment on the connection setup capacity required if each web page request requires a connection setup. Comment on the amount of state information required to keep track of these connections.

The number of telephone call requests is much less than the number of Web requests per hour. The processing required to set up a call is much greater than that required for Web requests, especially because of the need to maintain billing records. On the other hand, Web requests are so much more frequent that its overall processing capability is much greater. Fortunately, the Web page request connection setup, for example, the TCP connection for each HTTP request, is implemented by the end systems in the network, a situation that is easier to scale than connection setups inside the network.

7.7. Apply the end-to-end argument to the question of how to control the delay jitter that is incurred in traversing a multi-hop network.

Solution:

Delay jitter is the variability in the packet delay experienced by a sequence of packets traversing the network. Many audio and video applications have a maximum tolerable bound on the delay jitter.

To control the delay jitter inside the network, every router must forward the packets within a certain time window by means of delaying the “early” packets or immediately forwarding the “late” packets. This requires significant amount of state information and processing at each router, and hence greatly complicates the network.

In the end-to-end approach, the applications at the end-systems have the knowledge of the tolerable range of jitter. A smoothing buffer can be used at the end system to collect the packets, and pass the packets to the application layer at a constant rate according to the application’s specification. Consequently, the jitter is controlled while the network service remains simple.

7.8. Compare the operation of the layer 3 entities in the end systems and in the routers inside the network.

Solution:

All layer 3 (network layer) entities both in the end systems and inside the network perform packet routing. However, the layer 3 entities at the end systems are only concerned with transporting packets from one source to one destination (end-system to network or network to end-system). Routers perform both routing and forwarding of packets.

The layer 3 entities in the end systems in some networks regulate the input flow for congestion control. In networks with QoS support, they also police and shape input traffic according to the pre-negotiated contract between the end-system and the network. Layer 3 entities may also segment packets for the data link layer network, and assemble packets to pass to the upper layer at the end system.

The routers inside the network switch packets from multiple sources to multiple destinations. They perform traffic management functions (priority scheduling, buffer management, and shaping) to provide QoS to different traffic flows.

7.9. Consider a “fast” circuit-switching network in which the first packet in a sequence enters the network and triggers the setting up of a connection “on-the-fly” as it traces a path to its destination. Discuss the issues in implementing this approach in packet switching.

Solution:

There are two major issues. The first is the delay that the setup procedure introduces to the first packet. As a result, if the transmission of the rest of the packets starts after the first packet is sent, all these packets will be stuck behind the first packet as it sets the path and moves forward hop-by-hop. The proposed approach in fast circuit switching is to delay the packets for some time after the first packet is sent. The second issue is the possibility of failure in the path setup in a route which requires the choice an alternative route. If this happens all the packets behind the first packet will be lost because rerouting is done for the path setup packet and not the data packets following it.

7.10. Circuit switching requires that the resources allocated to a connection be released when the connection has ended. Compare the following two approaches to releasing resources: (a) use an explicit connection release procedure where the network resources are released upon termination of the connection and (b) use a time-out mechanism where the resources allocated to a connection are released if a “connection refresh” message is not received within a certain time.

Solution:

In the first approach once a connections has ended, a disconnect procedure is initiated during which control messages are sent to release the resources across the connection path. In this approach the resources are released as soon as the connection is ended. The cost associated with this approach is in the signaling that is required to terminate the connection and release the resources allocated to the connection.

In the second approach the resources are released if a connection refresh message is not received within a certain time. In this approach the refresh message needs to be sent periodically which consumes some of bandwidth. On the other hand, the connection management procedure is simplified by not having to explicitly release resources. Also, it takes some time after the connection has ended before the resources are released. And finally, there is the possibility that the refresh messages are lost during the connection, which results in the release of the resources in some sections of the route. This latter problem is addressed by setting the refresh timer to be a fraction of the connection reservation time to allow for the possibility of lost refresh messages.

7.11. The oversubscription ratio can be computed with a simple formula by assuming that each subscriber independently transmits at the peak rate of c when there is data to send and zero otherwise. If there are N subscribers and the average transmission rate of each subscribers is r , then the probability of k subscribers transmitting data at the peak rate simultaneously is given by

$$P_k = \binom{N}{k} (r/c)^k (1-r/c)^{N-k}$$

The access multiplexer is said to be in the “overflow mode” if there are more than n subscribers transmitting data simultaneously, and the corresponding overflow probability is given by

$$Q_n = \sum_{k=n+1}^N \binom{N}{k} (r/c)^k (1-r/c)^{N-k}$$

If a certain overflow probability is acceptable, then n is the smallest value that satisfy the given overflow probability. Find the oversubscription ratio if $N=50$, $r/c=0.01$, and $Q_n < 0.01$.

Solution:

$$Q_n = 1 - \sum_{k=0}^n \binom{N}{k} (r/c)^k (1-r/c)^{N-k} \leq 0.01$$

$$Q_n = 1 - \sum_{k=0}^n \binom{50}{k} (.01)^k (.99)^{50-k} \leq 0.01$$

$$Q_n = 1 - (.99^{50}) \sum_{k=0}^n \binom{50}{k} (.01/.99)^k \leq 0.01$$

$$n = 3 \rightarrow Q_n = 1 - (0.605006067) 1.65023771 = 1 - 0.998403827 = 0.001596173$$

The table below shows Q_n as a function of n . Thus if we have 50 users that are busy only 1% of the time, then the probability that more than 3 subscribers are busy is 0.0016. The probability that 2 or more subscribers are busy is 0.0148.

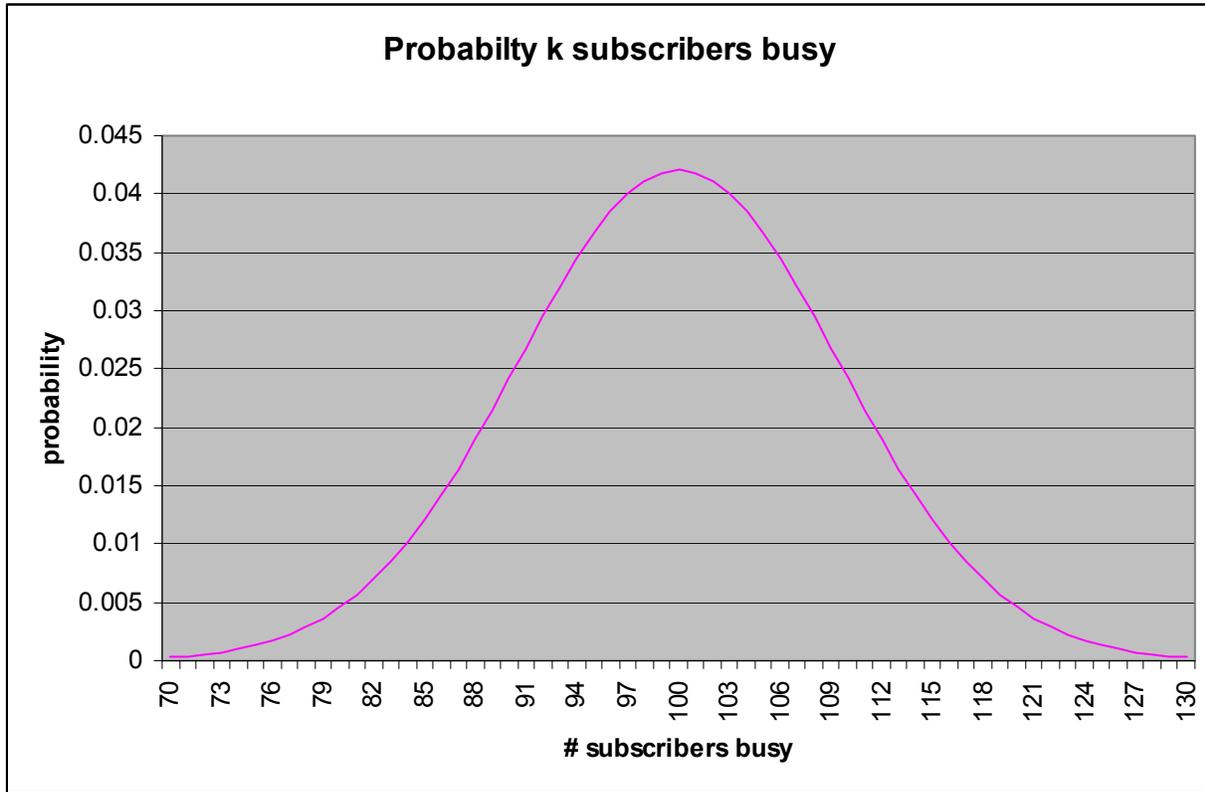
k	N		Q_n
	50		
	p		
	0.01		
		0	1
0	0.605006	0.605006	0.394994
1	0.305559	0.910565	0.089435
2	0.075618	0.986183	0.013817
3	0.012221	0.998404	0.001596
4	0.00145	0.999854	0.000146
5	0.000135	0.999989	1.09E-05
6	1.02E-05	0.999999	6.85E-07
7	6.48E-07	1	3.69E-08
8	3.52E-08	1	1.73E-09
9	1.66E-09	1	7.13E-11
10	6.87E-11	1	2.61E-12
11	2.52E-12	1	8.63E-14
12	8.29E-14	1	3.44E-15
13	2.45E-15	1	0

7.12. An access multiplexer serves $N = 1000$ subscribers, each of which has an activity factor $r/c = 0.1$. What is the oversubscription ratio if an overflow probability of 1 percent is acceptable? If $Np(1-p) \gg 1$, you may use an approximation technique (called the DeMoivre-Laplace Theorem), which is given by

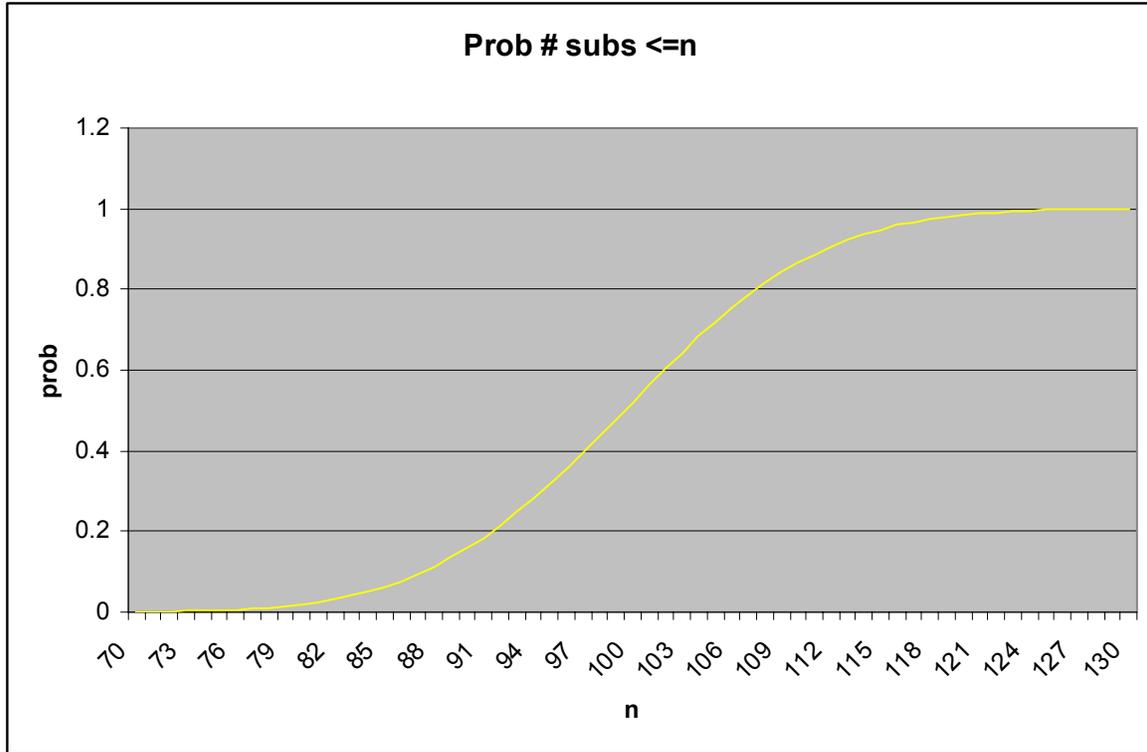
$$P_k = \binom{N}{k} p^k (1-p)^{N-k} \approx \frac{1}{\sqrt{2\pi Np(1-p)}} e^{-\frac{(k-Np)^2}{2Np(1-p)}}$$

Solution:

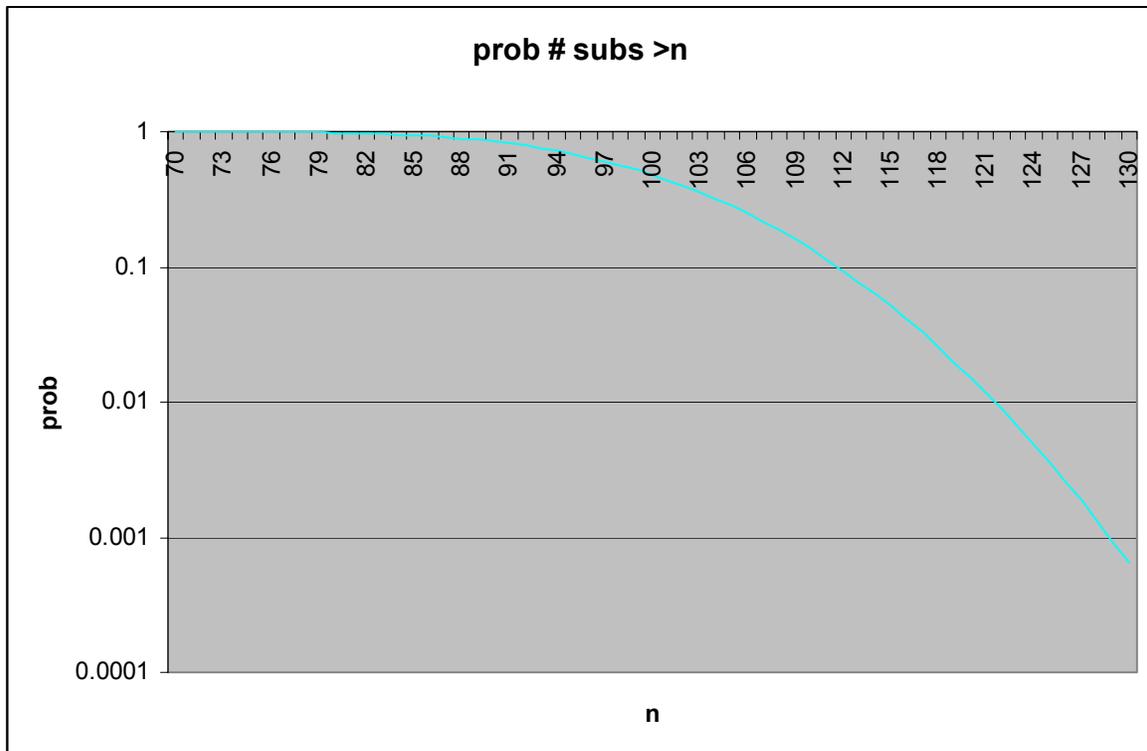
The average number of busy subscribers is $Np = 1000 \cdot 0.1 = 100$. The graph below shows the probabilities in the vicinity of $k = 100$ using the above approximation:



By adding the probabilities up to a certain number of subscribers, say n , we can find the probability that the number of subscribers is less than or equal to n . The graph below shows the result.



The following plots the probability that more than n subscribers are busy in a logarithmic scale.



We see that the probability of exceeding 122 is about 1%.

7.13. In Figure 7.5 trace the transmission of IP packets from when a web page request is made to when the web page is received. Identify the components of the end-to-end delay.

Solutions follow questions:

- (a) Assume that the browser is on a computer that is in the same departmental LAN as the server.

(Assume TCP connection to server has been established)
 Creation of IP packet at workstation to carry Web page request
 Frame transfer delay in the LAN
 Receipt of IP packet at server
 Processing of web page request
 Creation of IP packet(s) to return Web page
 Transfer of frame(s) across the LAN
 Receipt of IP packet(s) at workstation

- (b) Assume that the web server is in the central organization servers.

(Assume TCP connection to server has been established)
 Creation of IP packet at workstation to carry Web page request
 Frame transfer delay in the LAN
 Receipt of IP packet at departmental router
 Routing and forwarding of IP packet across campus backbone
 Receipt of IP packet at server farm router
 Routing and forwarding of IP packet to Web server
 Processing of web page request
 Creation of IP packet(s) to return Web page
 Transfer of frame(s) across the server LAN to server farm router
 Receipt of IP packet(s) at server router
 Routing and forwarding of IP packet(s) across campus backbone
 Receipt of IP packet(s) at departmental router
 Routing and forwarding of IP packet(s) to workstation
 Transfer of frame(s) across departmental LAN
 Receipt of IP packet(s) at workstation

- (c) Assume that the server is located in a remote network.

(Assume TCP connection to server has been established)
 Creation of IP packet at workstation to carry Web page request
 Frame transfer delay in the LAN
 Receipt of IP packet at departmental router
 Routing and forwarding of IP packet across campus backbone
 Receipt of IP packet at gateway router
 Routing and forwarding of IP packet across Internet to Web server
 Processing of web page request
 Creation of IP packet(s) to return Web page
 Transfer of frame(s) across the Internet to gateway router
 Receipt of IP packet(s) at gateway router
 Routing and forwarding of IP packet across campus backbone
 Receipt of IP packet(s) at departmental router
 Routing and forwarding of IP packet(s) to workstation
 Transfer of frame(s) across departmental LAN
 Receipt of IP packet(s) at workstation

7.14. In Figure 7.5 trace the transmission of IP packets between two personal computers running an IP telephony application. Identify the components of the end-to-end delay.

Solutions follow questions:

- (a) Assume that the two PCs are in the same departmental LAN.

IP packets flow simultaneously in both directions.
Consider the flow from PC1 to PC2:

Processing of sound to create digital signal
Creation of IP packet at PC1 to carry a voice packet
Frame transfer delay in the LAN
Receipt of IP packet at PC2
Processing of digitized signal to produce sound

The sources of delays are in the processing of the voice signal and in the transfer of frames across the LAN.

- (b) Assume that the PCs are in different domains.

IP packets flow simultaneously in both directions.
Consider the flow from PC1 to PC2:

Processing of sound to create digital signal
Creation of IP packet at PC1 to carry a voice packet
Frame transfer delay in LAN1
Receipt of IP packet at departmental router
Routing and forwarding of IP packet across campus backbone
Receipt of IP packet at gateway router
Routing and forwarding of IP packet across Internet to PC2
Receipt of IP packet at PC2
Processing of digitized signal to produce sound

The sources of delays are in the processing of the voice signal, in the transfer of frames across LAN, and most significantly in the transfer across multiple routers in the Internet. The latter component can introduce large and variable delays.

7.15. In Figure 7.5 suppose that a workstation becomes faulty and begins sending LAN frames with the broadcast address. What stations are affected by this broadcast storm? Explain why the use of broadcast packets is discouraged in IP.

Solution:

All workstations in the extended (departmental) LAN are affected by the broadcast storm. Each LAN interface recognizes the frame as a broadcast frame and hence is compelled to process it. If a workstation sends out IP packets with broadcast address, the routers will forward the packets beyond the departmental LAN, to the campus backbone network, wide area network, and eventually the broadcast storm would affect the entire network. Therefore, the IP routers usually have the broadcast option disabled.

7.16. Explain why the distance in hops from your ISP to a NAP is very important. What happens if a NAP becomes congested?

Solution:

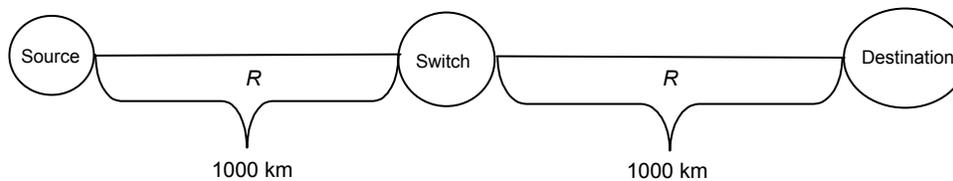
NAPs are the points where multiple networks interconnect to exchange IP traffic. Many IP transfers are therefore likely to involve traveling across a NAP. Thus the distance and hence the delay to other destinations is influenced strongly by the distance to a NAP.

If a NAP becomes congested, then the ISPs connected to it are limited in the traffic they can exchange. The ISPs have to route traffic through other NAPs and may cause further congestion elsewhere.

7.17. A 64-kilobyte message is to be transmitted over two hops in a network. The network limits packets to a maximum size of 2 kilobytes, and each packet has a 32-byte header. The transmission lines in the network are error free and have a speed of 50 Mbps. Each hop is 1000 km long. How long does it take to get the message from the source to the destination?

Solution:

The situation is shown in the figure below.



First let's determine the number of packets required for the message. Given:

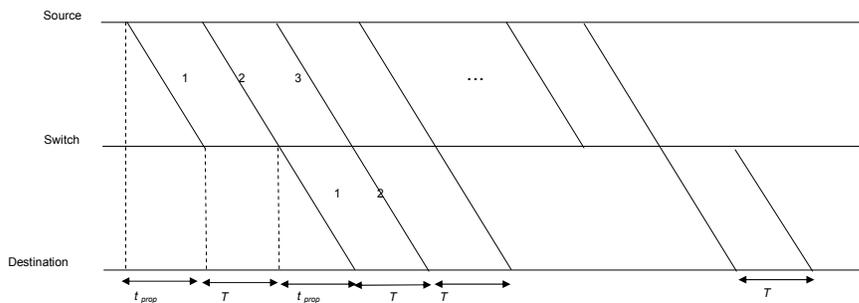
message size = 64 Kb

maximum packet size = 2 Kb, with a 32-byte header.

$R = 50$ Mbps

number of packets = $64 \times 1024 / (2 \times 1024 - 32) = 32.5 \cong 33$ packets.

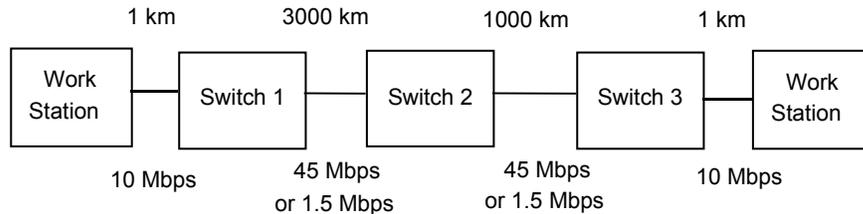
Assume that the signal travels at the speed of light, or $t_{prop} = 1000 \text{ km} / (3 \times 10^8 \text{ ms}) = 3 \frac{1}{3} \text{ ms}$. The transmission time of one packet is $T = 2 \text{ Kb} / 50 \text{ Mbps} = 2 \times 1024 \times 8 / 50 \times 10^6 = 327.68 \mu\text{sec}$. The figure below shows the basic situation.



Assume that queueing and processing delays at the intermediate switch are negligible.

$$\begin{aligned} \text{Total delay} &= 2 t_{\text{prop}} + T (\text{number of packets} + 1) \\ &= 2 (3 \frac{1}{3} \text{ ms}) + (327.68 \mu\text{sec}) (33 + 1) = 17.8 \text{ ms} \end{aligned}$$

7.18. An audiovisual real-time application uses packet switching to transmit 32 kilobit/second speech and 64 kilobit/second video over the following network connection.



Two choices of packet length are being considered: In option 1 a packet contains 10 milliseconds of speech and audio information; in option 2 a packet contains 100 milliseconds of speech and audio information. Each packet has a 40-byte header.

Solutions follow questions:

- (a) For each option find out what percentage of each packet is header overhead.

Option 1:

10 ms of speech gives 32 x 10 bits of information
 10 ms of audio gives 64 x 10 bits of information
 40-byte header produces 40 x 8 bits of header

Therefore, the packet size is 1280 bits, with an overhead of 25%.

Option 2:

100 ms of speech gives 32 x 10² bits of information
 100 ms of audio gives 64 x 10² bits of information
 40-byte header produces 40 x 8 bits of header

Therefore, the packet size is 9920 bits, with an overhead of 3.225%.

- (b) Draw a time diagram and identify all the components of the end-to-end delay. Keep in mind that a packet cannot be sent until it has been filled and that a packet cannot be relayed until it is completely received (that is, store and forward). Assume that bit errors are negligible.

The end-to-end delay contains the following components: 1. a packetization delay P_s consisting of the time it take a source to generate enough information to fill a packet; the first byte in the packet is delayed by this amount: 2. queueing delay, frame transmission delay, and propagation delay at each hop.

We assume that a message is the same as a packet and that

$$\text{delay} = P_s + P_1 + P_2 + P_3 + P_4 + T_1 + T_2 + T_3 + T_4$$

⏟
⏟
 Propagation delays Transmission through four hops

We sketch these delays in the following part.

- (c) Evaluate all the delay components for which you have been given sufficient information. Consider both choices of packet length. Assume that the signal propagates at a speed of 1 km /5 microseconds. Consider two cases of backbone network speed: 45 Mbps and 1.5 Mbps. Summarize your result for the four possible cases in a table with four entries.

Transmission Delays:

Bit Rate	1280 bits	9920 bits
1.5×10^6	853 μ sec	6613 μ sec
45×10^6	28,4 μ sec	220.4 μ sec

Propagation delays:

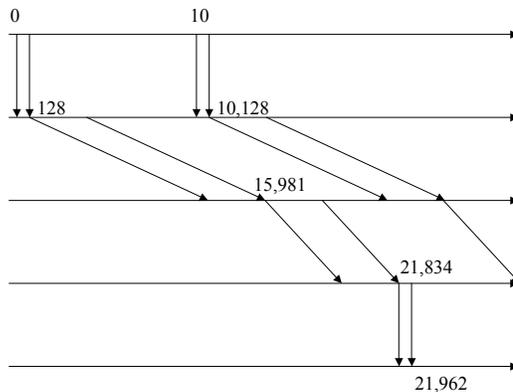
$P_1 = P_4 = 5 \times 10^{-6} \text{ sec} = .005 \text{ ms}$

$P_2 = 15 \text{ ms}$

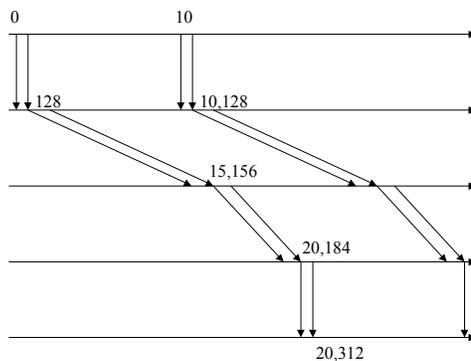
$P_3 = 5 \text{ ms}$

Total = 20 ms

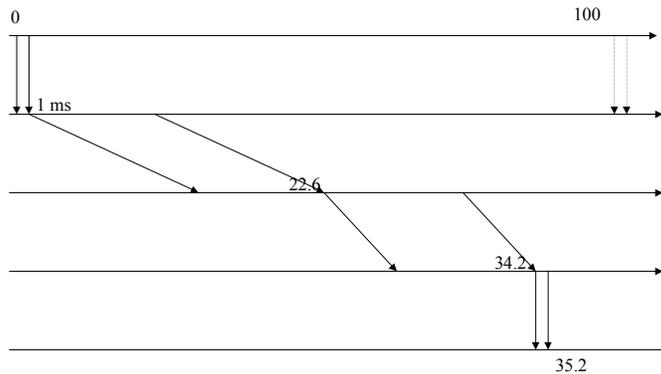
10 ms packetization, 1.5 Mbps: The packet transmission time is $1280/1.5\text{Mbps}=853 \mu\text{sec}$. In this case, the long propagation delay dominates the overall delay.



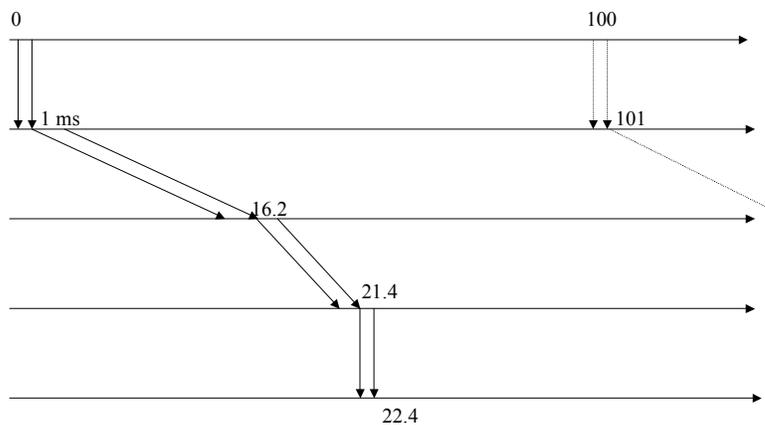
10 ms packetization, 45 Mbps: the packet transmission time is $1280/45\text{Mbps}=28.4 \mu\text{sec}$. The high-speed in the backbone portion of the network does little to reduce the overall delay.



100 ms packetization, 1.5 Mbps: The packet transmission time is now $9920/1.5\text{Mbps} = 6613 \mu\text{sec}$. The long propagation delay and the long packet size provide comparable delays.



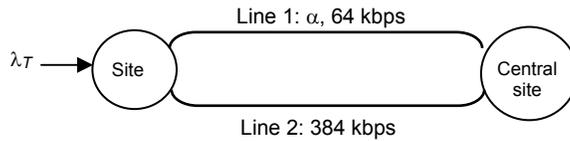
100 ms packetization, 45 Mbps: The high backbone speed makes the long packet transmission time negligible relative to the propagation delay.



(d) Which of the preceding delay components would involve queuing delays?

The transfer in switches 1, 2, and 3 involve random queuing delays.

7.19. Suppose that a site has two communication lines connecting it to a central site. One line has a speed of 64 kbps, and the other line has a speed of 384 kbps. Suppose each line is modeled by an M/M/1 queuing system with average packet delay given by $E[D] = E[X]/(1 - \rho)$ where $E[X]$ is the average time required to transmit a packet, λ is the arrival rate in packets/second, and $\rho = \lambda E[X]$ is the load. Assume packets have an average length of 8000 bits. Suppose that a fraction α of the packets are routed to the first line and the remaining $1 - \alpha$ are routed to the second line.

Solutions follow questions:

- (a) Find the value of
- α
- that minimizes the total average delay.

$$E[L] = 8000 \text{ bits}$$

$$\lambda_1 = \alpha\lambda_T, \lambda_2 = (1 - \alpha)\lambda_T$$

$$E[X_1] = 8000 \text{ bits} / 64 \text{ kbps} = 0.125 \text{ sec}$$

$$E[X_2] = 8000 \text{ bits} / 384 \text{ kbps} = 0.02083 \text{ sec}$$

$$E[D_T] = E[D_1] \alpha + E[D_2] (1 - \alpha), \alpha < 1$$

$$= \frac{1}{\frac{1}{E[X_1]} - \lambda_1} \alpha + \frac{1}{\frac{1}{E[X_2]} - \lambda_2} (1 - \alpha)$$

$$= \frac{1}{8 - \lambda_1} \alpha + \frac{1}{48 - \lambda_2} (1 - \alpha)$$

$$= \frac{1}{8 - \lambda_T \alpha} \alpha + \frac{1}{48 - \lambda_T (1 - \alpha)} (1 - \alpha)$$

Before doing the math to find the minimum in the above expression, it is worthwhile to consider a couple of extreme cases. The maximum capacity available to transfer packets across the combined links is 56 packets per second, 8 p/s across the first link and 48 p/s across the second link. Clearly when the arrival rate approaches capacity, we will send traffic across the two links. However consider the case when the arrival rate is very small. The delay across the first link is $1/8$ and the delay across the second link is $1/48$. Clearly at low traffic we want to send all the traffic on the fast link. As we increase the arrival rate, we eventually reach a point where it is worthwhile to route some packets along the slow link. Thus we expect α to be zero for some range of arrival rates, and then become nonzero after a specific value of arrival rate is reached.

Now take derivatives to find the minimum:

Let A be the first term in the delay expression and let B be the second term.

$$\frac{dE[D_T]}{d\alpha} = \frac{dA}{d\alpha} + \frac{dB}{d\alpha}$$

$$\frac{dA}{d\alpha} = \frac{d}{d\alpha} \frac{1}{8 - \lambda_T \alpha} \alpha = (-1)(8 - \lambda_T \alpha)^{-2} (-\lambda_T) \alpha + \frac{1}{8 - \lambda_T \alpha}$$

$$= \frac{\lambda_T \alpha}{(8 - \lambda_T \alpha)^2} + \frac{1}{8 - \lambda_T \alpha} = \frac{\lambda_T \alpha + 8 - \lambda_T \alpha}{(8 - \lambda_T \alpha)^2} = \frac{8}{(8 - \lambda_T \alpha)^2}$$

Similarly, we have

$$\begin{aligned} \frac{dB}{d\alpha} &= \frac{d}{d\alpha} \left[\frac{1}{48 - \lambda_T(1 - \alpha)} (1 - \alpha) \right] = (-1) [48 - \lambda_T + \lambda_T \alpha]^{-2} (\lambda_T)(1 - \alpha) + \frac{-1}{48 - \lambda_T(1 - \alpha)} \\ &= \frac{-\lambda_T(1 - \alpha)}{(48 - \lambda_T + \lambda_T \alpha)^2} + \frac{-1}{48 - \lambda_T + \lambda_T \alpha} = \frac{\lambda_T(1 - \alpha) - (48 - \lambda_T + \lambda_T \alpha)}{(48 - \lambda_T + \lambda_T \alpha)^2} = \frac{-48}{(48 - \lambda_T + \lambda_T \alpha)^2} \end{aligned}$$

To find the α that minimizes the delay, we set $\frac{dA}{d\alpha} + \frac{dB}{d\alpha} = 0$. We then obtain:

$$\frac{8}{(8 - \lambda_T \alpha)^2} + \frac{-48}{(48 - \lambda_T + \lambda_T \alpha)^2} = 0$$

$$\frac{1}{(8 - \lambda_T \alpha)^2} = \frac{6}{(48 - \lambda_T + \lambda_T \alpha)^2}$$

$$6(8 - \lambda_T \alpha)^2 = (48 - \lambda_T + \lambda_T \alpha)^2$$

$$\sqrt{6}(8 - \lambda_T \alpha) = 48 - \lambda_T + \lambda_T \alpha$$

$$8\sqrt{6} - \sqrt{6}\lambda_T \alpha = 48 - \lambda_T + \lambda_T \alpha$$

$$8\sqrt{6} - 48 + \lambda_T = \lambda_T \alpha + \sqrt{6}\lambda_T \alpha$$

$$\alpha = \frac{8\sqrt{6} - 48 + \lambda_T}{(1 + \sqrt{6})\lambda_T} = \frac{\lambda_T - 28.4}{3.45\lambda_T}$$

Thus we see that α must be zero for arrival rates less than 28.4.

- (b) Compare the average delay in part (a) to the average delay in a single multiplexer that combines the two transmission lines into a single transmission line.

With a single transmission line the link speed will be $64 + 384 = 448$ kbps. Therefore,

$$E[X] = 8000 \text{ bits} / 448 \text{ kbps} = 0.0179 \text{ sec}$$

$$E[D] = \frac{1}{\frac{1}{E[X]} - \lambda_T} = \frac{1}{\frac{1}{0.0179} - \lambda_T} = \frac{1}{56 - \lambda_T}$$

The average delay in the single multiplexer is smaller than the average delay in the system served by two servers. This reduction is due to the fact that some of the capacity in the two-server system can be idle when there is only one packet in the system, whereas the single multiplexer system is always working at full rate.

20. A message of size m bits is to be transmitted over an L -hop path in a store-and-forward packet network as a series of N consecutive packets, each containing k data bits and h header bits. Assume that $m \gg k + h$. The transmission rate of each link is R bits/second. Propagation and queuing delays are negligible.

Solutions follow questions:

We know from Figure 7.10 that if the message were sent as a whole, then an entire message propagation delay must be incurred at each hop. If instead we break the message into packets, then from Figure 7.12 a form of pipelining takes place where only a packet time is incurred in each hop. The point of this problem is to determine the optimum packet size that minimizes the overall delay.

We know that:

- message size = m bits = N packets
- one packet = k data bits + h header bits
- $m \gg k + h$
- there are L hops of store-and-forward
- transmission rate = R bits/sec
- t_{prop} and t_{qing} are negligible
- transmission time of one packet, T , equals the amount of bits in a packet / transmission rate, or $T = (k + h) / R$.

(a) What is the total number of bits that must be transmitted?

$N(k + h)$ bits are transmitted.

(b) What is the total delay experienced by the message (that is, the time between the first transmitted bit at the source and the last received bit at the destination)?

The total delay, D , can be expressed by

$$D = Lt_{prop} + LT + (N - 1)T.$$

Since the propagation delay is negligible, the first term in the above equation is zero. Substituting for T we get,

$$\begin{aligned} D &= L \frac{k+h}{R} + (N-1) \frac{k+h}{R} \\ &= (N-1+L) \frac{k+h}{R} \end{aligned}$$

(c) What value of k minimizes the total delay?

We determined D in the previous equation. We also know that $N = \left\lceil \frac{m}{k} \right\rceil$, where m is the ceiling.

Therefore,

$$D \approx \left(\frac{m}{k} - 1 + L \right) \frac{k+h}{R} = \left(\frac{m}{k} \right) \frac{k+h}{R} + (L-1) \frac{k+h}{R}$$

Note that there are two components to the delay: 1. the first depends on the number of packets that need to be transmitted, m/k , as k increases this term decreases; 2. the second term increases as the size of the packets, $k + h$, is increased. The optimum value of k achieves a balance between these two tendencies.

$$\begin{aligned} \frac{dD}{dk} &= \left[\frac{d}{dk} \left(\frac{m}{k} - 1 + L \right) \right] \frac{k+h}{R} + \left(\frac{m}{k} - 1 + L \right) \frac{d}{dk} \left(\frac{k+h}{R} \right) \\ &= \frac{-m}{k^2} \times \frac{k+h}{R} + \left(\frac{m}{k} - 1 + L \right) \frac{1}{R} \end{aligned}$$

If we set dD/dk to zero, we obtain:

$$\frac{-m}{k^2} \times \frac{k+h}{R} + \left(\frac{m}{k} - 1 + L \right) \frac{1}{R} = 0$$

$$\frac{m}{k^2} \left(\frac{k}{R} + \frac{h}{R} \right) = \frac{m}{kR} - \frac{1}{R} + \frac{L}{R}$$

$$\frac{m}{kR} + \frac{mh}{k^2R} = \frac{m}{kR} + \frac{L-1}{R}$$

$$\frac{mh}{k^2} = L-1$$

$$k^2 = \frac{mh}{L-1}$$

$$k = \sqrt{\frac{mh}{L-1}}$$

It is interesting to note that the optimum payload size is the geometric mean of the message size and the header, normalized to the number of hops (-1).

7.21. Suppose that a datagram packet-switching network has a routing algorithm that generates routing tables so that there are two disjoint paths between every source and destination that is attached to the network. Identify the benefits of this arrangement. What problems are introduced with this approach?

Solution:

There are several benefits to having two disjoint paths between every source and destination. The first benefit is that when a path goes down, there is another one ready to take over without the need of finding a new one. This reduces the time required to recover from faults in the network. A second benefit is that traffic from a source to a destination can be load-balanced to yield better delay performance and more efficient use of network resources. On the other hand, the computing and maintaining two paths for every destination increases the complexity of the routing algorithms. It is also necessary that the network topology be designed so that two disjoint paths are available between any two nodes.

7.22. Suppose that a datagram packet-switching network uses headers of length H bytes and a virtual-circuit packet-switching network uses headers of length h bytes. Use Figure 7.15 to determine the length M of a message for which the virtual-circuit switching delivers the packet in less time than datagram switching does. Assume packets in both networks are the same length.

Solution:

Assume the message M produces k packets of M/k bits each. With headers the datagram packet is then $H + M/k$ bits long, and the VC packet is $h + M/k$ bits long.

From Figure 7.16, the total delay for datagram packet switching (assuming a single path for all packets) is

$$T = L t_{prop} + L P + (k-1) P \text{ where } P = (H + M/k)/R.$$

For VC packet switching, we first incur a connection setup time of $2L t_{prop} + L t_{proc}$ for the processing and transmission of the signaling messages. The total delay to deliver the packet is then:

$$T' = 2L t_{prop} + L t_{proc} + L t_{prop} + L P' + (k-1) P' \text{ where } P' = (h + M/k)/R.$$

Datagram switching has higher total delay if $T > T'$, that is,

$$L t_{prop} + L P + (k-1) P > 2L t_{prop} + L t_{proc} + L t_{prop} + L P' + (k-1) P'$$

which is equivalent to:

$$2L t_{prop} + L t_{proc} < (L+k-1)(H-h)/R.$$

This condition is satisfied when the VC setup time is small and the difference between the lengths of the datagram and VC headers is large.

It is also interesting to consider VC packet switching with cut-through switching. In this case the overall delay is:

$$T'' = 2L t_{prop} + L t_{proc} + L t_{prop} + L h + k P' \text{ where } P' = (h + M/k)/R.$$

Now datagram switching will require more time if:

$$L t_{prop} + L P + (k-1) P > 2L t_{prop} + L t_{proc} + L t_{prop} + L h + (k-1) P'$$

which is equivalent to:

$$2L t_{prop} + L t_{proc} < (k-1)(H-h)/R - L(P-h).$$

Note that if the VC setup time can be avoided (for example if the VC has been set up by previous messages and if the network caches the VC for some time), then the delay for the VC approach improves dramatically.

7.23. Consider the operation of a packet switch in a connectionless network. What is the source of the load on the processor? What can be done if the processor becomes the system bottleneck?

Solution:

There are two sources of load on the processor. The first involves the processing involved in determining the route and next-hop for each packet. The second source of load to the processor is the implementation of routing protocols. This involves the gathering of state information, the

exchange of routing messages with other nodes, and the computation of routes and associated routing tables.

The first source of processing load can be addressed by implementing per-packet processing in hardware, and by distributing this processing to line cards in a switch. Distributing the computation of the routing tables in a switch between central processors and processors in the line cards can be done but is quite tricky.

7.24. Consider the operation of a packet switch in a connection-oriented network. What is the source of the load on the processor? What can be done if the processor becomes overloaded?

Solution:

The main source of load in the switch is in the processing required for the setting up and tearing down of connections. The processor must keep track of the connections that are established in the switch and of the remaining resources for accepting new connections. The processor must participate in the exchange of signals to setup connections and it must implement connection admission procedures to determine when a connection request should be accepted or rejected. The processor load increases with the calling rate. The processor must be upgraded when the calling rate approaches the call handling capacity of the processor. Alternatively, the processing can be distributed between central processors and processing in the line cards.

7.25. Consider the following traffic patterns in a banyan switch with eight inputs and eight outputs in Figure 7.21. Which traffic patterns below are successfully routed without contention? Can you give a general condition under which a banyan switch is said to be non-blocking (that is, performs successful routing without contention)?

Solutions:

- (a) Pattern 1: Packets from inputs 0, 1, 2, 3, and 4 are to be routed to outputs 2, 3, 4, 5, and 7, respectively.
- (b) Pattern 2: Packets from inputs 0, 1, 2, 3, and 4 are to be routed to outputs 1, 2, 4, 3, and 6, respectively.
- (c) Pattern 3: Packets from inputs 0, 2, 3, 4, and 6 are to be routed to outputs 2, 3, 4, 5, and 7, respectively.

Pattern a is routed without contention. Patterns b and c have contention. A banyan switch is nonblocking if the input packets are sorted according to their destination port numbers and are given to the switch entries starting from top with no gaps or duplicates.

7.26. Suppose a routing algorithm identifies paths that are “best” in the sense that: (1) minimum number of hops, (2) minimum delay, or (3) maximum available bandwidth. Identify conditions under which the paths produced by the different criteria are the same? are different?

Solution:

The first criterion ignores the state of each link, but works well in situations where the states of all links are the same. Counting number of hops is also simple and efficient in terms of the number of bits required to represent the link. Minimum hop routing is also efficient in the use of transmission resources, since each packet consumes bandwidth using the minimum number of links.

The minimum delay criterion will lead to paths along the route that has minimum delay. If the delay is independent of the traffic levels, e.g. propagation delay, then the criterion is useful. However, if the delay is strongly dependent on the traffic levels, then rerouting based on the current delays in the

links will change the traffic on each link and hence the delays! In this case, not only does the current link delay need to be considered, but also the derivative of the delay with respect to traffic level.

The maximum available bandwidth criterion tries to route traffic along pipes with the highest “cross-section” to the destination. This approach tends to spread traffic across the various links in the network. This approach is inefficient relative to minimum hop routing in that it may use longer paths.

At very low traffic loads, the delay across the network is the sum of the transmission times and the propagation delays. If all links are about the same length and bit rate, then minimum hop routing and minimum delay routing will give the same performance. If links vary widely in length, then minimum hop routing may not give the same performance as minimum delay routing.

Minimum hop routing will yield the same paths as maximum available bandwidth routing if links are loaded to about the same levels so that the available bandwidth in links is about the same. When link utilization varies widely, maximum available bandwidth routing will start using longer paths.

7.27. Suppose that the virtual circuit identifiers (VCIs) are unique to a switch, not to an input port. What is traded off in this scenario?

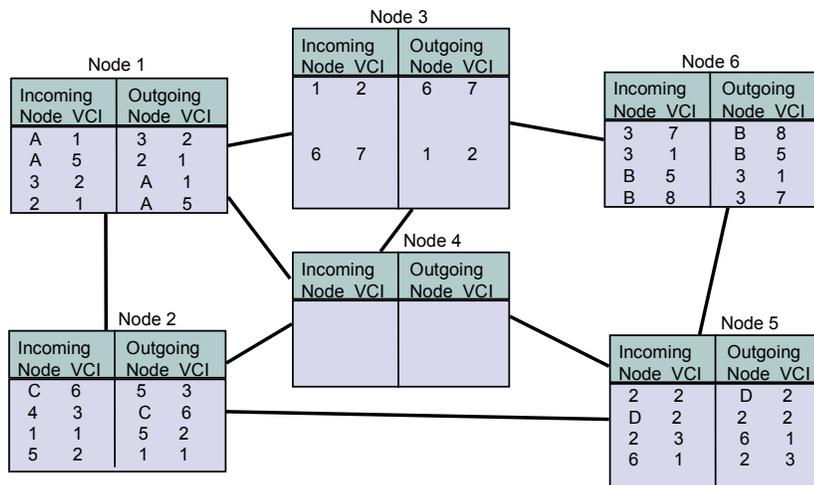
Solution:

When the VCIs are unique only within an input port, then the VCI numbers can be reused at all ports in the switch. Thus the number of VCIs that can be supported increases with the number of ports. In addition, VCI allocation is decentralized and can be handled independently by each input port. When the VCIs are unique only to the switch, then the number of VCIs that can be handled is reduced drastically. In addition VCI allocation must now be done centrally. Thus the per-switch VCI allocation approach may only be useful in small switch, small network scenarios.

7.28. Consider the virtual-circuit packet network in Figure 7.23. Suppose that node 4 in the network fails. Reroute the affected virtual circuits and show the new set of routing tables.

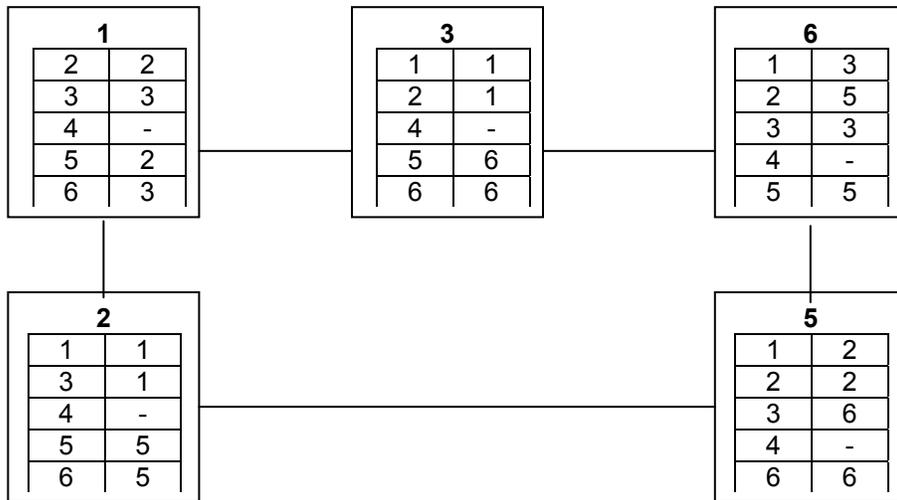
Solution:

When node 4 fails, all the affected connections need to be re-routed. In the example, the VC from A to D needs to be re-routed. One possible new route is A-1-2-5-D. The routing tables at nodes 1, 2, and 5 need to be modified. The VC from C to B also needs to be rerouted. A possible new route is C-2-5-6-D. The routing tables in 2, 5, and 6 need to be modified. The routing tables in node 3 have to be modified to remove the connections that no longer traverse it.

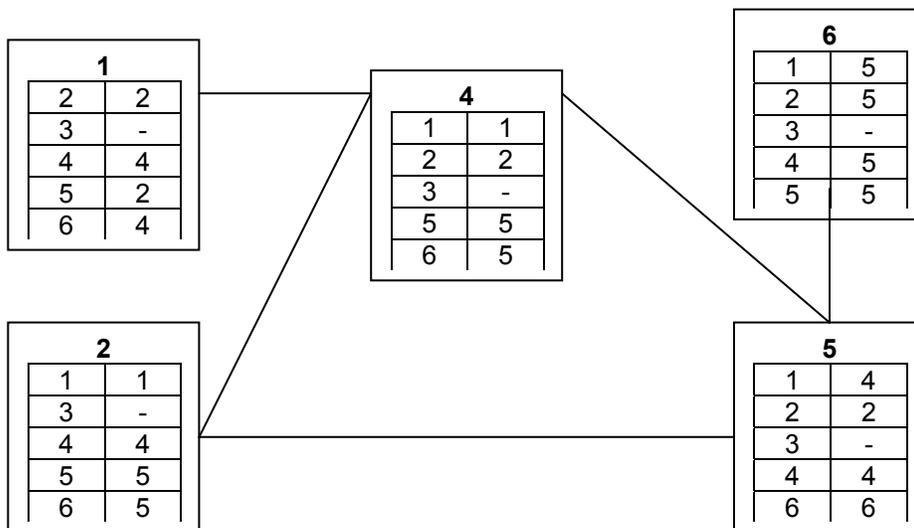


7.29. Consider the datagram packet network in Figure 7.25. Reconstruct the routing tables (using minimum hop routing) that result after node 4 fails. Repeat if node 3 fails instead.

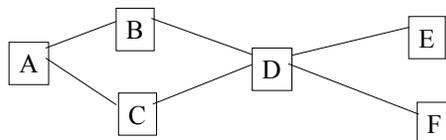
Solution:



If node 3 fails:



7.30. Consider the following six-node network. Assume all links have the same bit rate R .



Solutions follow questions:

- (a) Suppose the network uses datagram routing. Find the routing table for each node, using minimum hop routing.

destination	next node
B	B
C	C
D	B
E	B
F	B

destination	next node
A	A
C	D
D	D
E	D
F	D

destination	next node
A	A
B	A
D	D
E	D
F	D

destination	next node
A	B
B	B
C	C
E	E
F	F

destination	next node
A	D
B	D
C	D
D	D
F	D

destination	next node
A	D
B	D
C	D
D	D
E	D

(b) Explain why the routing tables in part (a) lead to inefficient use of network bandwidth.

Minimum-hop based routing causes all traffic flow through the minimum hop path, while the bandwidth of other paths of equal and longer distance are unused. It is inefficient in its usage of network bandwidth and likely to cause congestion in the minimum-hop paths.

For example, in the routing table of node A, the traffic flows from A to E and A to F both pass through node B, and leave the bandwidth of link AC and CD unused.

(c) Can VC routing give better efficiency in the use of network bandwidth? Explain why or why not.

Yes. During the VC setup, the links along the path can be examined for the amount of available bandwidth, if it below a threshold, the connection will be routed over alternate path. VC routing allows more even distribution of traffic flows and improves the network bandwidth efficiency.

(d) Suggest an approach in which the routing tables in datagram network are modified to give better efficiency. Give the modified routing tables.

Assign a cost to each link that is proportional to the traffic loading (number of connections) of the link. A higher cost is assigned to more congested links. Thus the resulting routing table avoids congested links and distributes traffic more evenly.

The changes in the routing table are in bold and italicized.

destination	next node
B	B
C	C
D	B
E	C
F	C

destination	next node
A	A
C	A
D	D
E	D
F	D

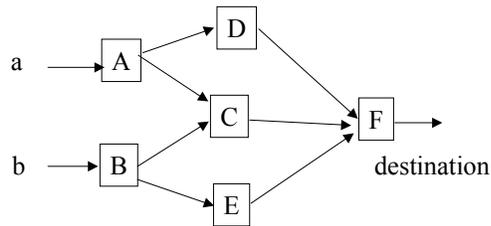
destination	next node
A	A
B	A
D	D
E	D
F	D

Destination	next node
A	B
B	B
C	C
E	E
F	F

destination	next node
A	D
B	D
C	D
D	D
F	D

destination	next node
A	D
B	D
C	D
D	D
E	D

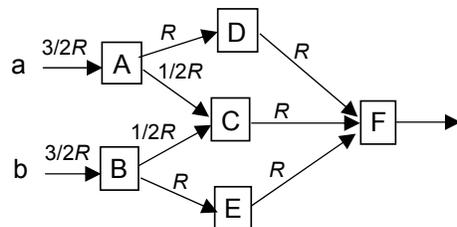
7.31. Consider the following six-node unidirectional network where flows a and b are to be transferred to the same destination. Assume all links have the same bit rate $R = 1$.



Solutions follow questions:

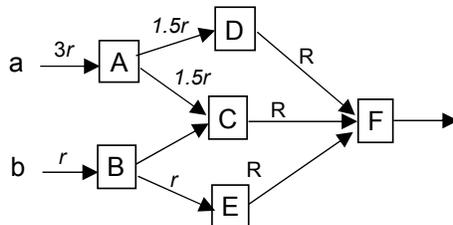
(a) If the flows a and b are equal, find the maximum flow that can be handled by the network.

The links to node F are the bottleneck links. A maximum flow of $3R$ can be handled.



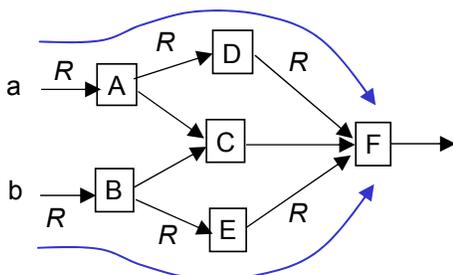
(b) If flow a is three times as large as flow b, find the maximum flow that can be handled by the network.

Let $b = r$ and $a = 3r$. Flow a sees a maximum cross-section of $2R$ flow across the network, and flow b sees a cross-section of $2R$ as well. Since flow a has the larger flow, its flow of $3r$ is limited to a maximum value of $2R$, that is, $r = 2R/3$. It then follows that b must be given by $r = 2R/3$.

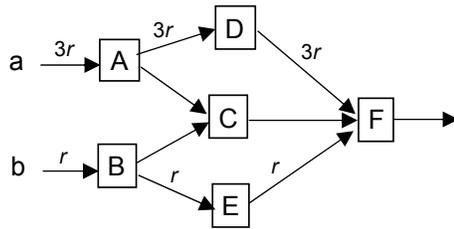


(c) Repeat (a) and (b) if the flows are constrained to use only one path.

Part (a): To have maximum flow, flow a must take path $A \rightarrow D \rightarrow F$. Similarly, flow b must take the path $B \rightarrow E \rightarrow F$. In this case, $a = R$ and $b = R$. So the total maximum flow is $2R$.



Part (b): $R = 3r$ and $r = 1/3R$. Therefore $a = R$ and $b = 1/3R$. The total maximum flow is $4/3R$.



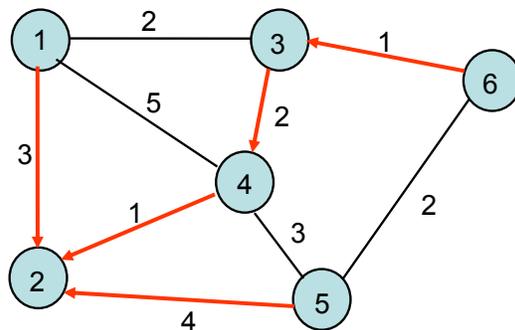
7.32. Consider the network in Figure 7.30.

Solutions follow questions:

(a) Use the Bellman-Ford algorithm to find the set of shortest paths from all nodes to destination node 2.

Iteration	Node 1	Node 3	Node 4	Node 5	Node 6
Initial	$(-1, \infty)$	$(-1, \infty)$	$(-1, \infty)$	$(-1, \infty)$	$(-1, \infty)$
1	(2, 3) (3, ∞) (4, ∞)	(1, ∞) (4, ∞) (6, ∞)	(1, ∞) (2, 1) (3, ∞) (5, ∞)	(2, 4) (4, ∞) (6, ∞)	(3, ∞) (5, ∞)
2	(2, 3) (3, ∞) (4, 6)	(1, 5) (4, 3) (6, ∞)	(1, 8) (2, 1) (3, ∞) (5, 7)	(2, 4) (4, 4) (6, ∞)	(3, ∞) (5, 6)
3	(2, 3) (3, 5) (4, 6)	(1, 5) (4, 3) (6, 7)	(1, 8) (2, 1) (3, 5) (5, 7)	(2, 4) (4, 4) (6, 8)	(3, 4) (5, 6)
4	(2, 3) (3, 5) (4, 6)	(1, 5) (4, 3) (6, 5)	(1, 8) (2, 1) (3, 5) (5, 7)	(2, 4) (4, 4) (6, 6)	(3, 4) (5, 6)

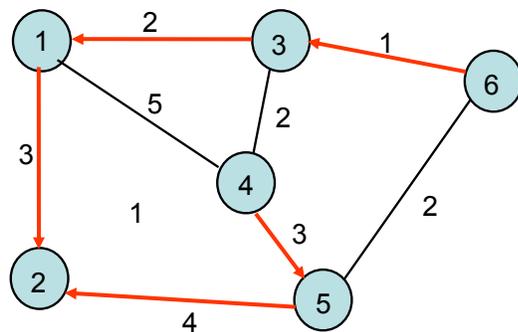
The set of paths to destination 2 are shown below:



Now continue the algorithm after the link between node 2 and 4 goes down.

Iteration	Node 1	Node 3	Node 4	Node 5	Node 6
Before break	(2, 3)	(4, 3)	(2, 1)	(2, 4)	(3, 4)
1	(2, 3) (3, 5) (4, 6)	(1, 5) (4, 3) (6, 5)	(1, 8) (3, 5) (5, 7)	(2, 4) (4, 4) (6, 6)	(3, 4) (5, 6)
2	(2, 3) (3, 5) (4, 10)	(1, 5) (4, 7) (6, 5)	(1, 8) (3, 5) (5, 7)	(2, 4) (4, 8) (6, 6)	(3, 4) (5, 6)
3	(2, 3) (3, 7) (4, 10)	(1, 5) (4, 7) (6, 5)	(1, 8) (3, 7) (5, 7)	(2, 4) (4, 8) (6, 6)	(3, 6) (5, 6)
4	(2, 3) (3, 7) (4, 12)	(1, 5) (4, 9) (6, 7)	(1, 8) (3, 7) (5, 7)	(2, 4) (4, 10) (6, 8)	(3, 6) (5, 6)

The new set of paths are shown below:

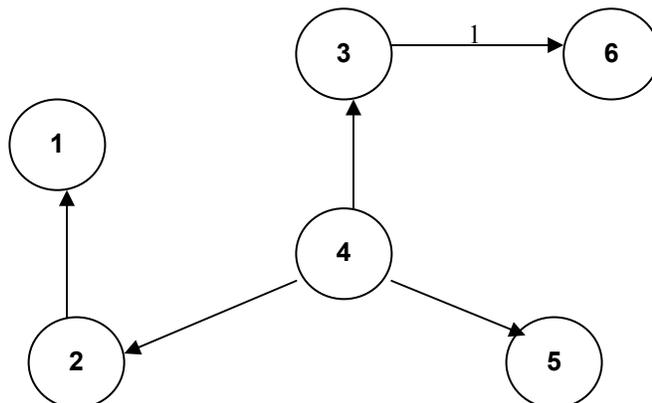


7.33. Consider the network in Figure 7.30.

Solutions follow questions:

(a) Use the Dijkstra algorithm to find the set of shortest paths from node 4 to other nodes.

Iteration	N	D ₁	D ₂	D ₃	D ₅	D ₆
Initial	{4}	5	1	2	3	∞
1	{2, 4}	4	1	2	3	∞
2	{2, 3, 4}	4	1	2	3	3
3	{2, 3, 4, 5}	4	1	2	3	3
4	{2, 3, 4, 5, 6}	4	1	2	3	3
5	{1, 2, 3, 4, 5, 6}	4	1	2	3	3



- (b) Find the set of associated routing table entries.

Destination	Next Node	Cost
1	2	4
2	2	1
3	3	2
5	5	3
6	3	3

7.34. Compare source routing with hop-by-hop routing with respect to (1) packet header overhead, (2) routing table size, (3) flexibility in route selection, and (4) QoS support, for both connectionless and connection-oriented packet networks.

Solution:

Packet header overhead is higher in source routing in both connectionless and connection-oriented networks because the ID of the nodes across the path is included in the header in source routing while in hop-by-hop routing only the destination in connectionless networks or the connection number in connection-oriented networks is specified in the header.

Source routing requires no routing table lookup in intermediate nodes while hop-by-hop routing requires a routing table with one entry per destination in connectionless networks or one entry per connection in connection-oriented networks. Source nodes however must have sufficient routing information to determine the path across the network, essentially the same information required in link state routing.

Source routing is more flexible because it can consider all possible routes from the source to destination while in hop-by-hop routing each node can choose only its neighbors. However, hop-by-hop routing can respond more quickly to changes in the network.

Source routing can incorporate QoS considerations more flexibly as it can consider all possibilities in routing, while hop-by-hop routing can incorporate the QoS considerations at each hop but not universally.

7.35. Suppose that a block of user information that is L bytes long is segmented into multiple cells. Assume that each data unit can hold up P bytes of user information, that each cell has a header that is H bytes long, and the cells are fixed in length, and padded if necessary. Define the efficiency as the ratio of the L user bytes to the total number of bytes produced by the segmentation process.

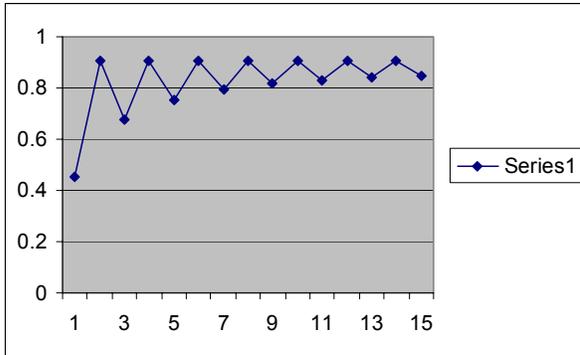
Solutions follow questions:

- (a) Find an expression for the efficiency as a function of L , H , and P . Use the ceiling function $c(x)$ which is defined as the smallest integer larger or equal to x .

$$eff = \frac{L}{L + H \times c\left(\frac{L}{P}\right)} = \frac{1}{1 + \left(\frac{H}{L}\right) \times c\left(\frac{L}{P}\right)}$$

- (b) Plot the efficiency for the following ATM parameters: $H = 5$, $P = 48$, and $L = 24k$ for $k = 0, 1, 2, 3, 4, 5$, and 6 .

0	0	0	8	0.905660377	0.90566
1	0.452830189	0.45283	9	0.81509434	0.815094
2	0.905660377	0.90566	10	0.905660377	0.90566
3	0.679245283	0.679245	11	0.830188679	
4	0.905660377	0.90566	12	0.905660377	
5	0.754716981	0.754717	13	0.84097035	
6	0.905660377	0.90566	14	0.905660377	
7	0.79245283	0.792453	15	0.849056604	



7.36. Consider a videoconferencing application in which the encoder produces a digital stream at a bit rate of 144 kbps. The packetization delay is defined as the delay incurred by the first byte in the packet from the instant it is produced to the instant when the packet is filled. Let P and H be defined as they are in problem 7.35.

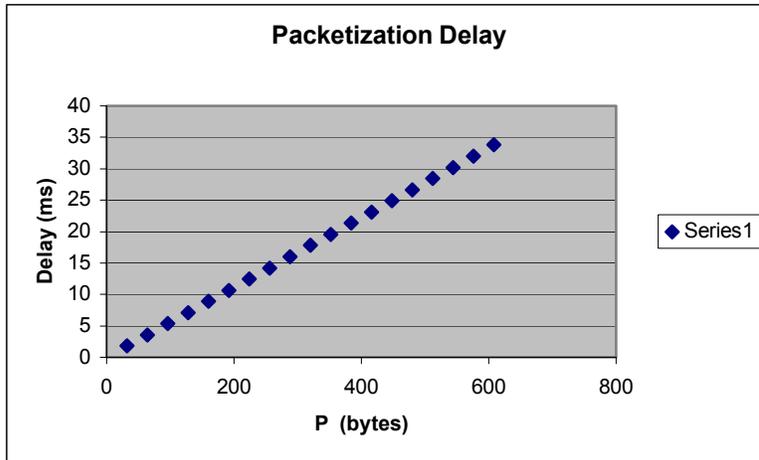
Solutions follow questions:

- (a) Find an expression for the packetization delay for this video application as a function of P .

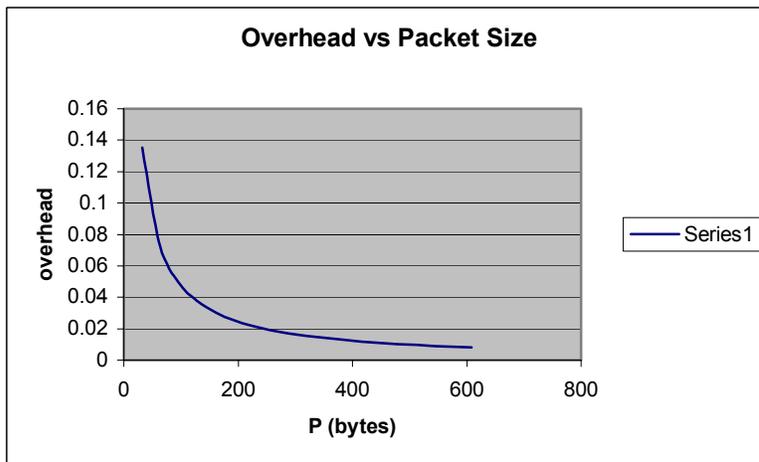
$$\text{packetization delay} = P \cdot 8 / (144 \text{ kbps})$$

- (b) Find an expression for the efficiency as a function of P and H . Let $H = 5$ and plot the packetization delay and the efficiency versus P .

$$\text{efficiency} = P / (P + H) = 1 / (1 + H/P)$$



overhead = 1 – efficiency



The two graphs show that there is a tradeoff between efficiency and packetization delay. The efficiency improves as P is increased; however the packetization delay also increases with P . An overly stringent delay requirement can lead to bandwidth inefficiency.

7.37. Suppose an ATM switch has 16 ports each operating at SONET OC-3 transmission rate, 155 Mbps. What is the maximum possible throughput of the switch?

Solution:

The maximum possible throughput of the switch is achieved when all of the input ports operate at maximum speed:

$$155 \text{ Mbps} \times 16 = 2.48 \text{ Gbps.}$$

This calculation assumes that the switch fabric is capable of switching the maximum amount of traffic that can arrive at the line cards. Achieving a capacity of 2.48 Gbps is trivial today, but at much higher rates (say, a terabit-per-second) many switch designs do not have the ability to switch all the traffic that can arrive at the input ports.

The student is warned that the networking industry literature tends to quote the capacity of a switch as the sum of the input and output bit rates of all the line cards. This practice gives a capacity that is at best twice the actual maximum throughput. As indicated above, it is not unusual for a switch fabric to not have enough capacity to transfer all the traffic that can arrive at its input ports.

7.38. Refer to the virtual-circuit packet-switching network in Figure 7.23. How many VCIs does each connection in the example consume? What is the effect of the length of routes on VCI consumption?

Solution:

Each connection uses the same number of VCIs as the number of hops in its path. However, since each VCI is defined local to each link, the use of a VCI on one hop does not deplete the available VCIs in another link.

7.39. Generalize the hierarchical network in Figure 7.26 so that 2^K nodes are interconnected in a full mesh at the top of the hierarchy, and so that each node connects to 2^L nodes in the next lower level in the hierarchy. Suppose there are four levels in the hierarchy.

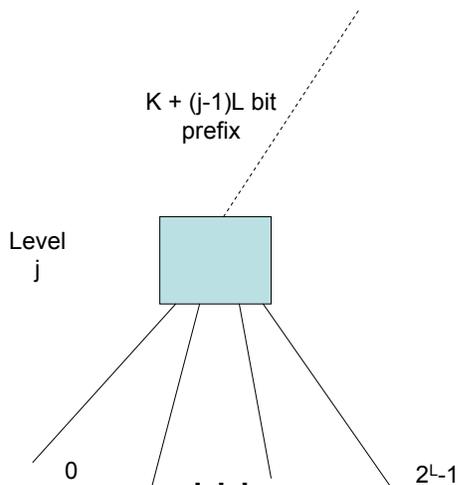
Solutions follow questions:

(a) How many nodes are there in the hierarchy?

$$2^K \times 2^L \times 2^L \times 2^L$$

(b) What does a routing table look like at level j in the hierarchy, $j = 1, 2, 3,$ and 4 ?

Each router is placed at a node in the tree network. The routes available either go up a single link in the hierarchy or down 2^L possible links down in the hierarchy. Packets that go down in the hierarchy have prefixes that correspond to the index of the node. The first L -bit symbol after said index indicates which link is to be taken down the hierarchy. All other packets go up the hierarchy by default.



(c) What is the maximum number of hops between nodes in the network?

6 hops: maximum 3 hops to the top level and maximum 3 hops to the lowest level

7.40. Assuming that the earth is a perfect sphere with radius 6400 km, how many bits of addressing is required to have a distinct address for every 1 cm x 1 cm square on the surface of the earth?

Solution:

Area of sphere = $4\pi r^2$, r is the radius.

$$\log_2 \left[\frac{4 \times \pi \times (6400 \times 10^5)^2}{1 \times 1} \right] = 63 \text{ bits}$$

This is fewer bits than the 128 bits of addressing used in IP version 6 discussed in Chapter 8.

7.41. Suppose that 64 kbps PCM coded speech is packetized into a constant bit rate ATM cell stream. Assume that each cell holds 48 bytes of speech and has a 5 byte header.

Solutions follow questions:

- (a) What is the interval between production of full cells?

$$\frac{48 \times 8 \text{ bits}}{64 \text{ Kb/sec}} = 6 \text{ milliseconds}$$

- (b) How long does it take to transmit the cell at 155 Mbps?

$$\frac{(48 + 5) \times 8 \text{ bits}}{155 \text{ Mbps}} = 2.73 \text{ microseconds}$$

- (c) How many cells could be transmitted in this system between consecutive voice cells?

$$\frac{.006}{2.73 \times 10^{-6}} = 2197 \text{ cells}$$

7.42. Suppose that 64 kbps PCM coded speech is packetized into a constant bit rate ATM cell stream. Assume that each cell holds 48 bytes of speech and has a 5 byte header. Assume that packets with silence are discarded. Assume that the duration of a period of speech activity has an exponential distribution with mean 300 ms and that the silence periods have a duration that also has an exponential distribution but with mean 600 ms. Recall that if T has an exponential distribution with mean $1/\mu$, then $P[T > t] = e^{-\mu t}$.

Solutions follow questions:

- (a) What is the peak cell rate of this system?

$$\text{peak cell rate} = \frac{64 \text{ kbps}}{48 \times 8 \text{ bits}} \approx 167 \text{ cells/sec}$$

- (b) What is the distribution of the burst of packets produced during an active period?

A speech activity period produces k packets if the duration of the burst is between $(k - 1)t_{\text{cell}}$ and $k t_{\text{cell}}$, where $t_{\text{cell}} = 6$ ms. Therefore the probability of a burst of size k is given by:

$$P[k] = P[(k - 1)0.006 \leq T \leq k0.006] = e^{-6(k-1)/300} - e^{-6k/300} = (1 - e^{-6/300})(e^{-6/300})^{k-1}$$

This is a geometric distribution with parameter $p = e^{-6/300}$.

- (c) What is the average rate at which cells are produced?

$$\frac{167 \text{ cells/sec} \times 300 \text{ ms}}{900 \text{ ms}} = 56 \text{ cells/sec}$$

7.43. Suppose that a data source produces information according to an on/off process. When the source is on, it produces information at a constant rate of 1 Mbps; when it is off, it produces no information. Suppose that the information is packetized into an ATM cell stream. Assume that each cell holds 48 bytes of data and has a 5 byte header. Assume that the duration of an on period has a Pareto distribution with parameter $\alpha = 1$. Assume that the off period is also Pareto but with general parameter α . If T has a Pareto distribution with parameter α , then $P[T > t] = t^{-\alpha}$ for $t > 1$. If $\alpha > 1$, then $E[T] = \alpha/(\alpha - 1)$, and if $0 < \alpha < 1$, then $E[T]$ is infinite.

Solutions follow questions:

- (a) What is the peak cell rate of this system?

$$\text{peak cell rate} = \frac{1 \text{ Mbps}}{48 \times 8 \text{ bits}} \approx 2604 \text{ cells/sec}$$

- (b) What is the distribution of the burst of packets produced during an on period?

The time to produce a cell is given by $t_{\text{cell}} = 1/2604 = 0.384 \text{ ms}$

$$\begin{aligned} P[k] &= P[(k-1)0.000384 \leq T \leq k0.000384] = 1 - (k \cdot 0.000384)^{-\alpha} - (1 - ((k-1) \cdot 0.000384)^{-\alpha}) \\ &= (2604/(k-1))^\alpha - (2604/k)^\alpha \end{aligned}$$

If $\alpha = 1$, then

$$P[k] = (2604/(k-1)) - (2604/k)$$

- (c) What is the average rate at which cells are produced?

If $\alpha > 1$, then the average length of a burst is $\alpha/(\alpha-1)$. Suppose that the off time has a Pareto distribution with parameter $\beta > 1$, and average length $\beta/(\beta-1)$. Thus the average rate at which cells are produced is the peak cell rate multiplied by the average time that the data source produces information:

$$\frac{\alpha/(\alpha-1)}{\alpha/(\alpha-1) + \beta/(\beta-1)} \times 2604$$

The Pareto distribution is characterized by high variability. For example the variance is infinite for $\alpha < 2$. Thus the on and off periods will exhibit high variability and are therefore very difficult to deal with in traffic management.

When α is less than 1, then the mean time in the on state is infinite. In this case, the source can stay in the on state forever and cells will arrive continuously at the peak rate.

7.44. An IP packet consists of 20 bytes of header and 1500 bytes of payload. Now suppose that the packet is mapped into ATM cells that have 5 bytes of header and 48 bytes of payload. How much of the resulting cell stream is header overhead?

Solution:

Total payload for ATM: 1520 bytes
 This implies 32 ATM frames
 Total ATM header bytes: 160
 Total Header bytes: 180
 Total bytes transmitted: 1696

Header overhead = $180 / 1696 = 10.61\%$

7.45. Suppose that virtual paths are set up between every pair of nodes in an ATM network. Explain why connection set up can be greatly simplified in this case.

Solution:

When two nodes need to communicate, each switch in the path does not have to be involved in the connection set up. Instead the switches at the ends of the VP assign an end-to-end VCI to each connection. When a cell traverses the network, only the VPI is used to forward it. The VCI is used only to direct the cell to the appropriate VC as the cell exits the network. The problem with this approach is that the number of virtual paths required grows as the square of the number of nodes in the ATM network.

7.46. Suppose that the ATM network concept is generalized so that packets can be variable in length. What features of ATM networking are retained? What features are lost?

Solution:

The capability of forwarding frames end-to-end across a virtual circuit using label switching is retained even if packets are variable in length. The ability to police, schedule and shape packet flows is also retained, but the fine degree of control afforded by the small and constant ATM cells is lost. The fast switching afforded by fixed-length blocks need not be lost since a switch can internally switch fixed-length blocks while still appearing as a variable length switch to the rest of the network. The header becomes more complex because of the variable packet length. On the other hand, the header processing rate is reduced because the time between header arrivals increases. For example a line card that handles ATM cells at 10 Gbps must process a cell every $53 \times 8 / 10^{10} = 42.4$ nanoseconds. The same line card handling 1500 byte MPLS packets must process a cell every $1500 \times 8 / 10^{10} = 1200$ nanoseconds, which is less demanding by a factor of nearly 30 times.

7.47. Explain where priority queueing and fair queueing may be carried out in the generic switch/router in Figure 7.19.

Solution:

Queueing and scheduling are used whenever there is contention for resources. In the ideal case, called "output-buffered queueing", the interconnection fabric has ample capacity to transfer all arriving packets virtually instantly from the input port to the output port. In this case, each output line card appears like a multiplexer that must exercise priority and fair queueing to decide the order in which packets are transmitted. Most switch designs depart from the ideal case and are limited in the rate at which packets can be transferred from input line cards to output line cards. Consequently, packets can accumulate in the input line card. In this case, scheduling is required to decide the order in which packets must be transferred across the fabric.

7.48. Consider the head-of-line priority system in Figure 7.42. Explain the impact on the delay and loss performance of the low-priority traffic under the following conditions:

Solutions follow questions:

The presence of high-priority packets prevents low-priority packets from accessing the transmission line. The periods of times when the high-priority buffer is non-empty are periods during which the low-priority buffer builds up without relief. The delay and loss performance of the low-priority traffic is therefore affected strongly by the pattern of bandwidth usage by the high-priority stream.

If high priority packets can preempt the transmission of low-priority packets, then low priority traffic has no effect whatsoever on high priority traffic. If high priority packets cannot preempt low priority traffic, then arriving high priority packets must wait for the completion of service of the low priority packet found upon arrival. If low priority packets can be long, then this impact can be significant.

- (a) The high-priority traffic consists of uniformly spaced, fixed-length packets.

Packets of the high-priority queue have constant arrival rate and service rate. Thus the low priority traffic is deprived of access to transmission bandwidth during periodically recurring time intervals. To the extent that the transmission line is sufficiently available to handle the low priority stream, the performance of low priority traffic is relatively steady and predictable.

- (b) The high-priority traffic consists of uniformly spaced, variable-length packets.

The high priority traffic has constant arrival rate but variable service rate due to the variable-length packets. The periods during which low priority traffic has access to transmission is now less predictable. Consequently, there is increased variability in delay and increased average loss and delay.

- (c) The high-priority traffic consists of highly bursty, variable-length packets.

The availability of bandwidth to the low priority traffic is now highly unpredictable. The delay and loss performance of the low priority traffic is the worst among all three cases.

7.49. Consider the head-of-line priority system in Figure 7.42. Suppose that each priority class is divided into several subclasses with different “drop” priorities. For each priority subclass there is a threshold that if exceeded by the queue length results in discarding of arriving packets from the corresponding subclass. Explain the range of delay and loss behaviors that are experienced by the different subclasses.

Solution:

Each priority class will experience delay performance better than that of lower priority classes and worse than that of higher priority classes.

Within each class, packets with higher drop preference will experience higher packet loss probability than packets from the same class but of lower drop preference. Packets with higher drop preference that are admitted into the buffer will also experience lower average delay than packets of lower drop precedence because on average they will find a smaller number of packets in the buffer.

7.50. Incorporate some form of weighted fair queueing in the head-of-line priority system in Figure 7.42 so that the low-priority traffic is guaranteed to receive r bps out of the total bit rate R of the transmission link. Explain why this may be a desirable feature. How does this impact the performance of the high-priority traffic?

Solution:

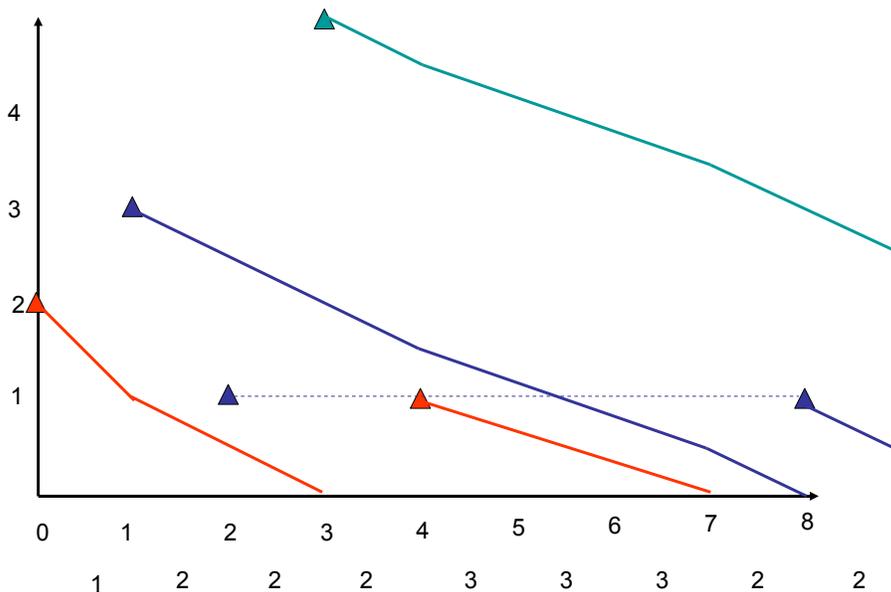
The trouble with the strict priority system in Figure 7.43 is that high priority traffic can starve the low priority traffic if no limits are placed on the volume of high priority traffic that is allowed into the system. An effective method for preserving the priority structure but preventing starvation of low priority traffic is to shape the high priority stream so that the average rate does not exceed some value, say $R - r$. In this case, r bps is guaranteed to become available to the low priority traffic. The performance of the high priority traffic is affected only insofar as the shaping introduces delay when packets need to be delayed to meet the shaping constraint. The guarantee of r bps to the low priority traffic can significantly improve the delay performance of the low priority traffic.

7.51. Consider a packet-by-packet fair queueing system with three logical buffers and with service rate of one unit/second. Show the sequence of transmissions for this system for the following packet arrival pattern. Buffer 1: arrival at time $t = 0$, length 2; arrival at $t = 4$, length 1. Buffer 2: arrival at time $t = 1$, length 3; arrival at $t = 2$, length 1. Buffer 3: arrival at time $t = 3$, length 5.

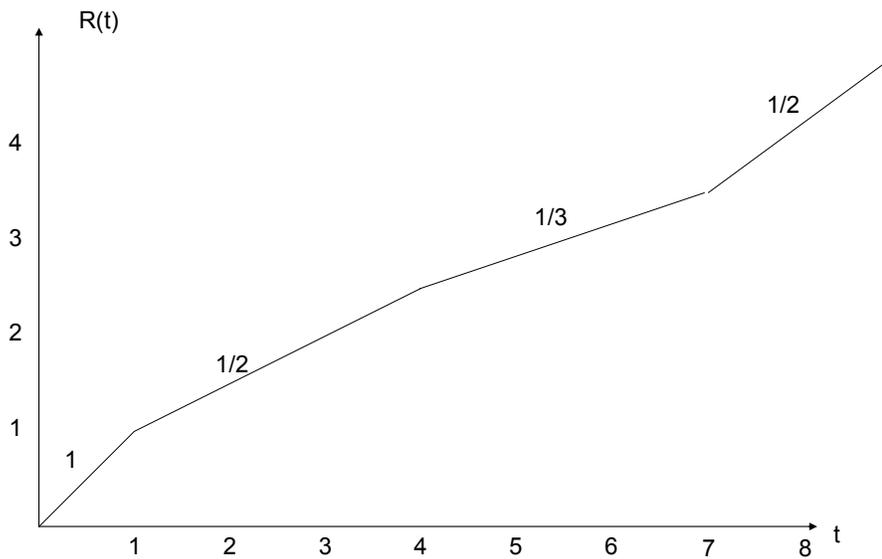
Solution:

The packet-by-packet fair queueing system selects the order of packet transmission according to the finish times in an emulated fluid flow system. The first figure below shows the times at which the various packets arrive and the length of the packets. The red triangles indicate packet arrivals at buffer 1; blue at buffer 2; and green at buffer 3. The rate at which packets are serviced is the reciprocal of the number of active buffers at the time. The second figure below shows how the round number $R(t)$ grows as a function of time. The slope of $R(t)$ is the reciprocal of the number of active stations. A given packet finishes service in the fluid flow system when the round number reaches the finish time computed for the given packet.

First Figure: packet arrivals and fluid flow service:



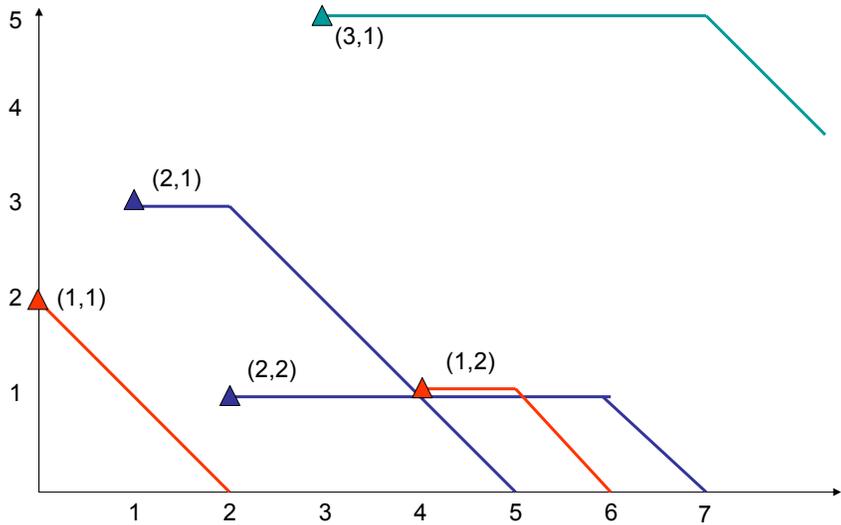
Second Figure: Round number vs. time



The table below gives the finish times for the various packet arrivals.

Buffer		$F(\text{buffer no.}, \text{packet no.})$	Order of transmission
B1	(1) $t = 0, L = 2$ (2) $t = 4, L = 1$	$F(1,1) = R(0) + 2 = 2$ $F(1,2) = R(4) + 1 = 3.5$	1 3
B2	(1) $t = 1, L = 3$ (2) $t = 2, L = 1$	$F(2, 1) = R(1) + 3 = 4$ $F(2,2) = \max\{F(2,1), 2\} + 1 = 5$	2 4
B3	(1) $t = 3, L = 5$	$F(3,1) = R(3) + 5 = 7$	5

The figure below shows the actual packet completion times for each of the packets.

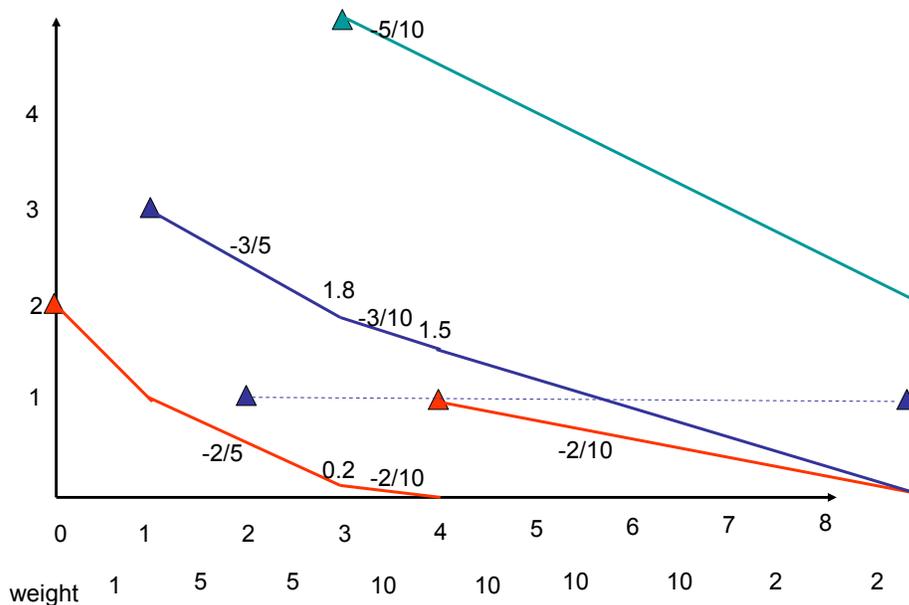


Note the difference in the order of packet completion times in the ideal fluid flow system and the packetized system. The first packet in buffer 2 is able to “sneak in” ahead of the second packet from buffer 1 in the time interval between 2 and 3 seconds.

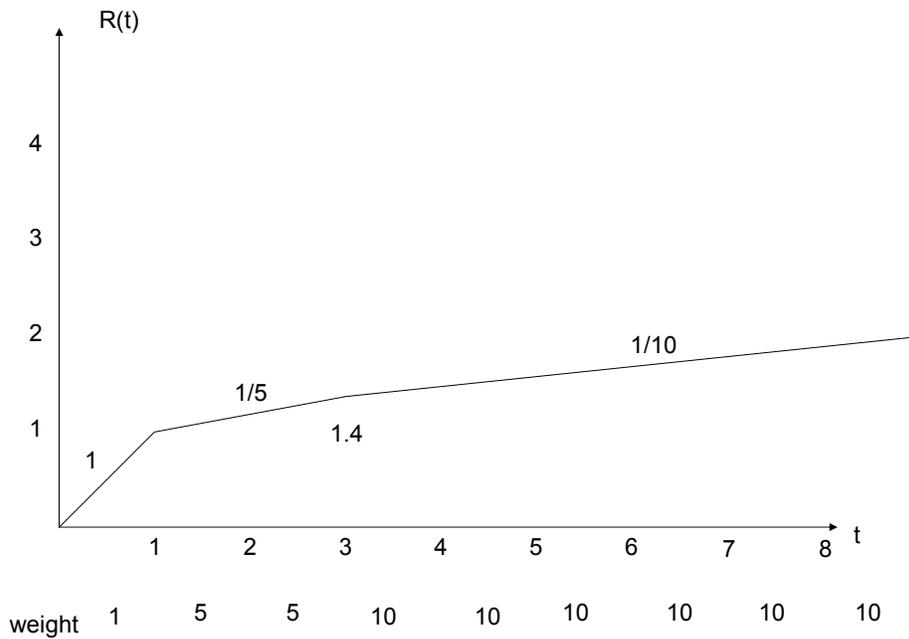
7.52 Repeat problem 7.51 if buffers 1, 2, and 3 have weights 2, 3, and 5, respectively.

Solution:

The following figure shows how the various packet arrivals are serviced in the fluid flow system. It can be seen that the packets from the first and second buffers take much longer to complete service because the lower weighting relative to the packet from buffer 3.



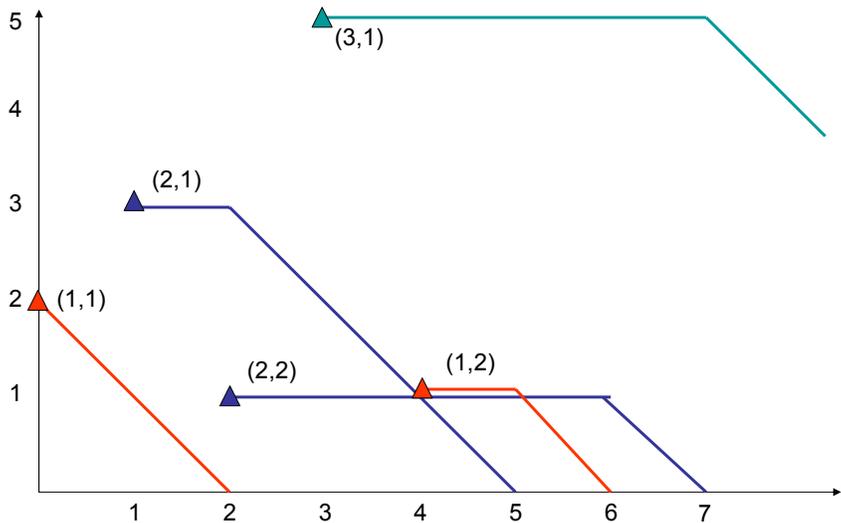
The generalization of the round number in the weighted fair queueing case has a slope that is the reciprocal of the sum of the weights of the active buffers at a given time. The round number vs. time is shown in the next figure.



The packet finish times are given in the table below:

Buffer		$F(\text{buffer no.}, \text{packet no.})$	Order of transmission
B1 $w = 2$	(1) $t = 0, L = 2$ (2) $t = 4, L = 1$	$F(1,1) = R(0) + (2/2) = 1$ $F(1,2) = R(4) + (1/2) = 2.$	1 3
B2 $w = 3$	(1) $t = 1, L = 3$ (2) $t = 2, L = 1$	$F(2, 1) = R(1) + (3/3) = 2$ $F(2,2) = \max\{2, 2\} + (1/3) = 2 \frac{1}{3}$	2 4
B3 $w = 5$	(1) $t = 3, L = 5$	$F(3,1) = R(3) + (5/5) = 2.4$	5

The following figure shows the actual packet completion times.



7.53. Suppose that in a packet-by-packet weighted fair-queueing system, a packet with finish tag F enters service at time t . Is it possible for a packet to arrive at the system after time t and have a finish tag less than F ? If yes, give an example. If no, explain why.

Solution:

Yes, the finish tag is a function of the arriving time of the packet to the system, but also the packet length. The early arriving packet may have a very long packet length while the late-arriving packet has a very short packet length. The late-arriving packet would then have a smaller F than the earlier packet.

In Problem 7.51, the first packet that arrives at buffer 2 ends up with a later finish time than the second packet that arrives at buffer 1. By the time the latter packet arrives, the former packet has already entered service in the real system.

7.54. Deficit round-robin is a scheduling scheme that operates as follows. The scheduler visits the buffers in round-robin fashion. A deficit counter is maintained for each buffer. When the scheduler visits a buffer, the scheduler adds a quantum of service to the deficit counter, and compares the resulting value to the length of the packet at the head of the line. If the counter is larger, the packet is served and the counter is reduced by the packet length. If not, the deficit is saved for the next visit. Suppose that a system has four buffers and that these contain packets of length 16, 10, 12 and 8 and that the quantum is 4 units. Show the deficit counter at each buffer as a function of time and indicate when the packets are transmitted.

Solution:

	R1	R2	R3	R4
BQ1: $L = 16$ $C = 0$	$C = 4$	$C = 8$	$C = 12$	$L = 0, Xmit$ $C = 0$
B2: $L = 10$ $C = 0$	$C = 4$	$C = 8$	$L = 0, Xmit$ $C = 2$	$L = 0$ $C = 0$
B3: $L = 12$ $C = 0$	$C = 4$	$C = 8$	$L = 0, Xmit$ $C = 0$	$L = 0$ $C = 0$
B4: $L = 8$ $C = 0$	$C = 4$	$L = 0, Xmit$ $C = 0$	$L = 0$ $C = 0$	$L = 0$ $C = 0$

The DRR scheduler visits each buffer in turn, adding a quantum of credits to the deficit if the buffer is nonempty. When a buffer has enough credits to transmit a packet, the scheduler sends the packet for transmission. The scheduler resumes visiting the buffers from this point when the packet has completed transmission.

7.55. Should packet-level traffic management be performed in the core of the network where the packet streams have been aggregated to multi-Gigabit/second bit rates? Discuss the alternative of using TDM circuits to carry packet streams in the core network.

Solution:

The purpose of packet-level traffic management is to ensure that packets are provided with the appropriate level of QoS, which typically involves packet delay and packet loss performance. As packet traffic is aggregated into higher bit-rate streams, it becomes more difficult and costly to do packet-level traffic management in the core. However, the much higher bit rates mean that the packet transmission times are much shorter, and hence queueing delays are shorter even if the number of packets in queue increases. Similarly, larger buffer sizes are affordable for higher bit rates so loss can be kept at relatively low levels. A reasonable approach is to multiplex packet streams into larger streams before entering the core of the network and to then transfer the streams efficiently switched using TDM or optical path circuits, thus eliminating the need for packet-level traffic management.

7.56. Queue management with random early detection (RED):

Solutions follow questions:

- (a) Explain why RED helps prevent TCP senders from detecting congestion and slowing down their transmission rates at the same time.

By randomly dropping the packets when the threshold is passed, only a subset of the sources is induced to reduce their traffic rates. Thus as congestion sets in, the community of TCP senders is gradually and progressively directed to slow down their transmission rates.

- (b) Discuss the effect of RED on network throughput.

RED maintains a relatively stable network throughput because it drops packets before congestion occurs and because it avoids the oscillations in throughput that can result when a large number of senders change their sending rate in unison. How close this throughput is to the network capacity is dependent to some extent on the min_{th} and max_{th} values set for RED.

- (c) Discuss the implementation complexity of the RED algorithm.

In the RED implementation, every queue must maintain two variables: min_{th} and max_{th} . A processing module that continuously takes the running average (to allow for occasional traffic bursts) of the queue length for every queue is also needed. A random number generator of some sort is required to drop packets with probabilities that depend on the running average of the queue length.

- (d) Discuss what would happen if instantaneous queue length were used instead of average queue length.

Using average queue length allows RED to tolerate small fluctuations around the threshold level. The use of instantaneous queue length would cause normal buffer fluctuations to cause RED to be initiated more frequently resulting in inefficiencies in the network.

- (e) Explore ways to find reasonable values for the RED parameters (i.e., min_{th} , max_{th} , and the packet drop probability when the average queue length reaches max_{th}).

The values of the RED parameters depend on a number of factors. The min_{th} and max_{th} selection depends on available buffer sizes. The difference between the two thresholds should be larger than the typical increase in average length in one round-trip time (RTT). The difference is also related to the transient burst sizes that should be tolerated when congestion occurs. A recommended setup for max_{th} is 2 to 3 times min_{th} . In general min_{th} should be lower if RTT is higher. The dependence on RTT relates the selection of the thresholds to link speed and propagation delay. The selection of min_{th} has an impact on average delay as well. For smaller delays, min_{th} should be smaller however the link utilization may be compromised if the threshold is too small.

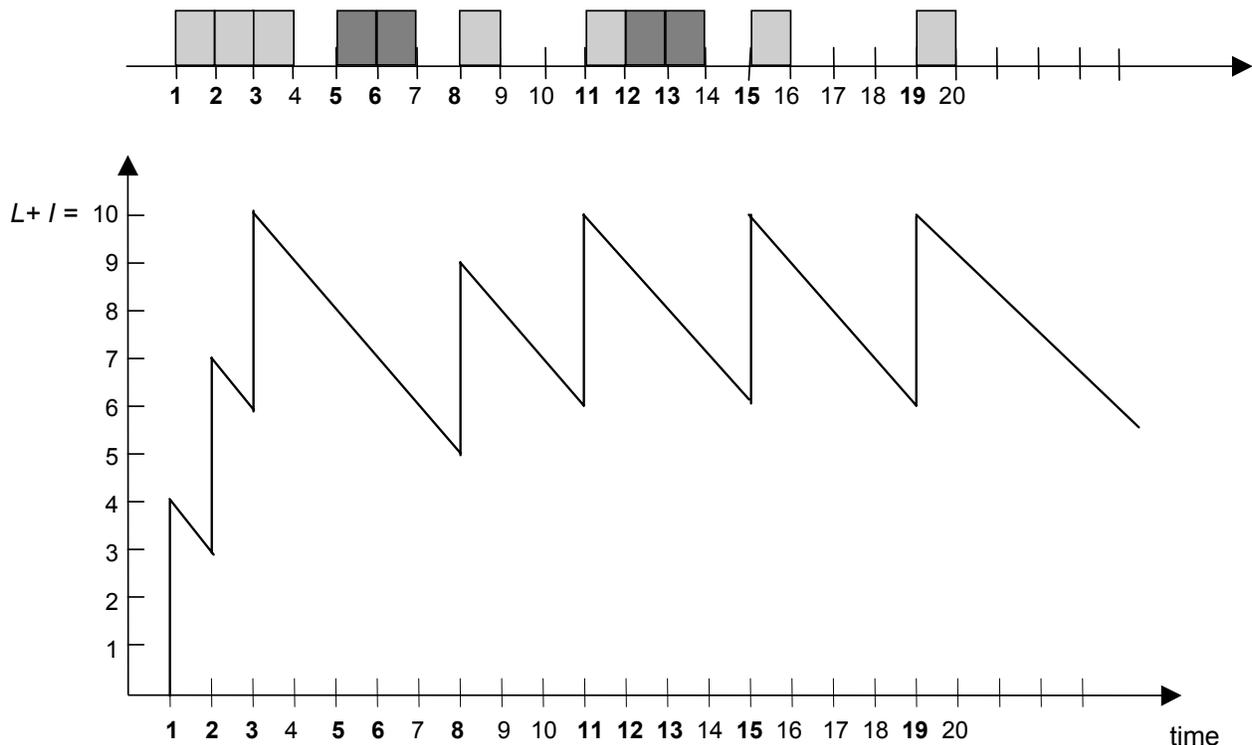
The selection of the packet drop probability should be related to the steady state packet drop in the network.

In weighted RED, the weight determines the time constant for the averaging for the average queue size. If the weight is too small the average queue size will respond too slowly. A large weight on the other hand will cause the average to follow the instantaneous queue size more closely and will result in fluctuations.

7.57. Suppose that ATM cells arrive at a leaky bucket policer at times $t = 1, 2, 3, 5, 6, 8, 11, 12, 13, 15,$ and 19 . Assuming the same parameters as the example in Figure 7.54, plot the bucket content and identify any nonconforming cells. Repeat if L is reduced to 4.

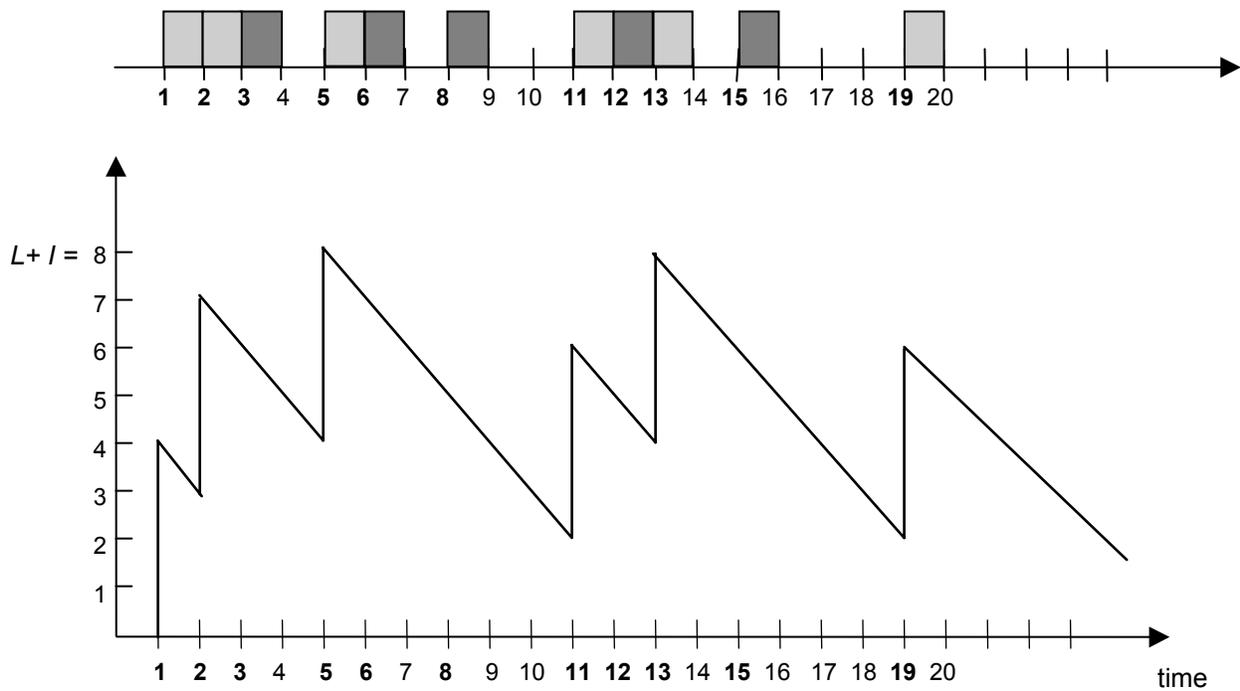
Solution:

1) $L = 4$ and $L = 6$. Drain rate = 1 unit/plot-time.



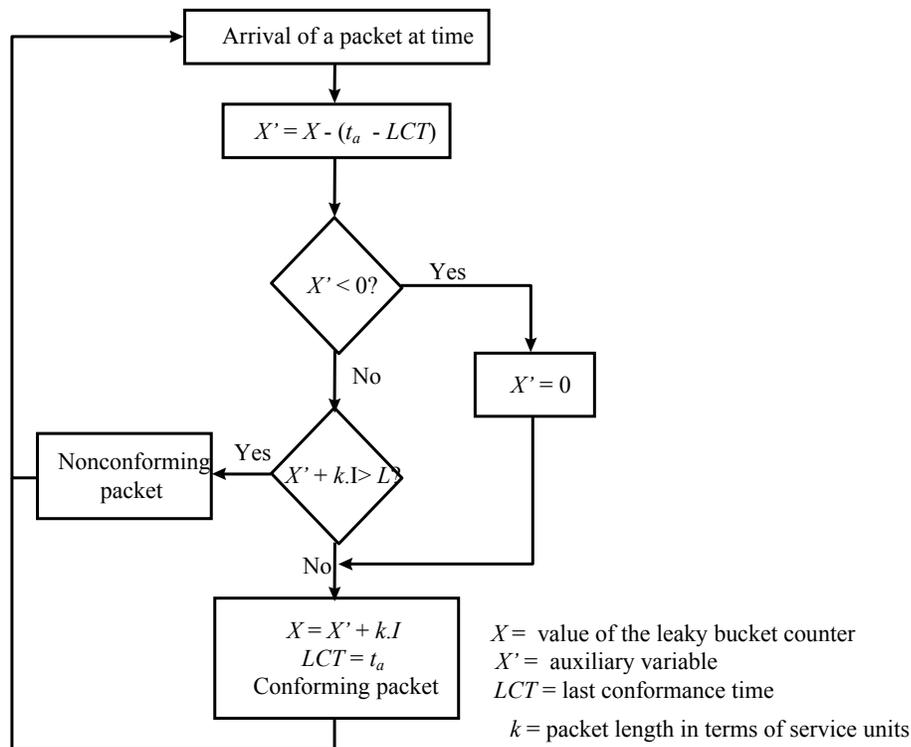
Cells 5, 6, 12, and 13 are nonconforming.

2) $L = 4$ and $L = 4$. Drain rate = 1 unit/plot-time.



7.58. Modify the leaky bucket algorithm in Figure 7.53 if packet length is variable.

Solution:



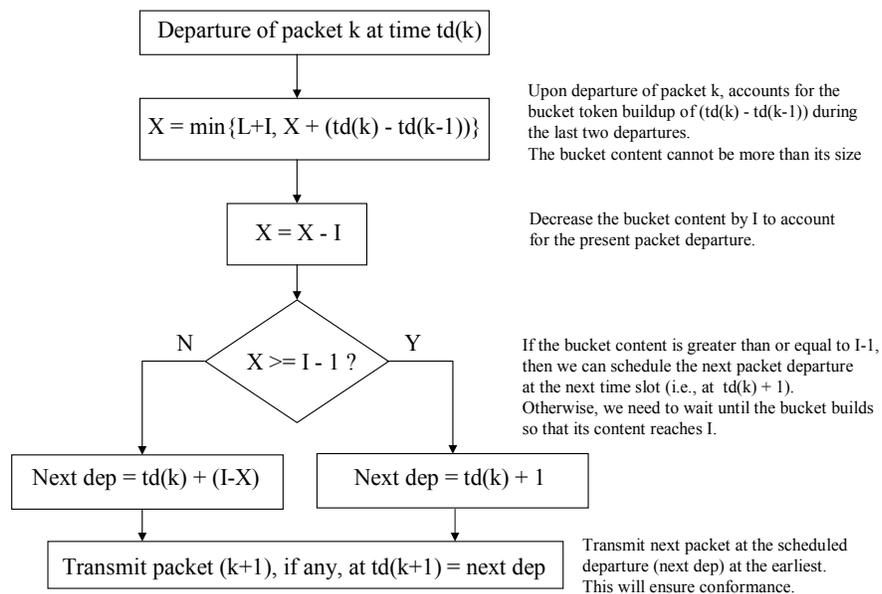
7.59. Explain the difference between the leaky bucket traffic shaper and the token bucket traffic shaper.

Solution:

In the leaky bucket the packet output rate is always constant. In the token bucket, the output rate can be variable and is controlled at one packet per token. If there are tokens in the bucket, the packets are sent out as soon as they arrive. The depth of the token bucket limits the traffic burstiness. If there are no tokens in the bucket, the packet has to wait for the arrival of a token, and the token arrival rate becomes the packet output rate. In the case of the empty token bucket where the tokens arrive at a constant rate, the token bucket essentially becomes a leaky bucket.

7.60. Show the algorithm for token bucket traffic shaper using a flow chart similar to the one shown in Figure 7.53 for policing. The flow chart begins with a departure of a packet k at time $td(k)$ and calculates the departure time for packet $k + 1$. Define the necessary parameters and assume that packet length is fixed.

Solution:



7.61. Explain where the policing device and the traffic shaping device should be located in the generic packet switch in Figure 7.19.

Solution:

The traffic policing device should be placed at the input port of the ingress line card. As packets arrive their usage history is updated and the packet is tagged and/or dropped if necessary. The traffic shaping device should be placed in the output port of egress line card, prior to transfer to another switch or network.

7.62. Which of the parameters in the upper bound for the end-to-end delay (Equation 7.24) are controllable by the application? What happens as the bit rate of the transmission links becomes very large?

Solution:

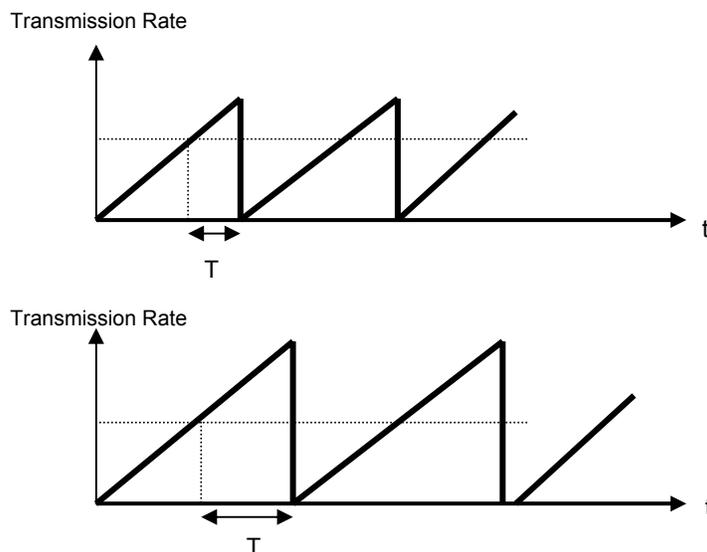
$$D \leq \frac{b}{R} + \frac{(H-1)m}{R} + \sum_{j=1}^H \frac{M}{R_j}$$

The maximum packet size of the given flow, m , can be controlled by the application. The reserved bandwidth allocation R for the flow and the maximum burst size b can be negotiated. If source routing is used the number of hops can be controlled to a limited extent.

As the transmission link speeds become very large, the last term becomes negligible and the performance is determined by b , H , R , and m .

7.63. Suppose a source with an unlimited amount of information to transmit uses a closed-loop control to regulate the transmission rate according to feedback information. If the feedback information indicates that there is no congestion, the source continuously increases its transmission rate in a linear fashion. If the feedback information indicates congestion along the path, the source sets its transmission rate to zero and then repeats the cycle by continuously increasing its transmission rate until congestion is detected once again. Assume that it takes T seconds for the feedback information to reach the source after congestion occurs. Sketch the transmission rate at the source versus time trajectory for a low and a high value of T . Explain how propagation delay T plays a role in closed-loop control.

Solution:

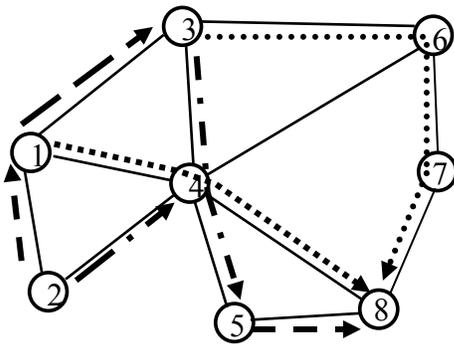


The rate at the source exhibits a sawtooth behavior. The rate starts at 0 and increases linearly until it reaches a peak and drops to 0 again. This is repeated periodically with a frequency which is proportional to the T , as the mechanism has no means of adapting to the available bandwidth. With a larger T the drop in the rate is delayed further from the time that congestion is detected, and so the source cannot respond quickly to the onset to congestion, and indeed may contribute to congestion as well.

7.64. Consider the network in Figure 7.63. Suppose that paths need to be set up in the following order: nodes 5 to 8, 1 to 8, 2 to 4, 3 to 8, 3 to 5, 2 to 1, 1 to 3, 3 to 6, 6 to 7, and 7 to 8. Assume that each link has a capacity of one unit and each path requires bandwidth of one unit.

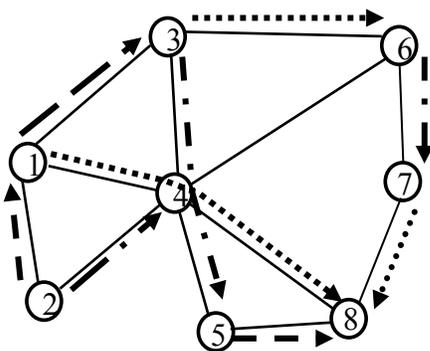
Solutions follow questions:

- (a) Use shortest-path routing to set up as many paths as possible. How many paths are blocked?



One connection (3 to 8) is blocked.

- (b) Use constraint shortest-path routing to set up as many paths as possible. How many paths are blocked?



Three connections (3 to 6, 6 to 7, and 7 to 8) are blocked.

- (c) Can you suggest an improvement to constraint shortest-path routing?

The number of connections that are blocked by each connection is tracked and if it is realized that the connection is blocking too many other connections the connection is rerouted.

Solutions to Chapter 8

8.1. The IP header checksum only verifies the integrity of IP header. Discuss the pros and cons of doing the checksum on the header part versus on the entire packet.

Solution:

Error checking in the header is more important because the packet is routed according to the header information. In addition, the delivery of the data at the destination to the higher layers also requires the header information. Thus error checking of the header protects against misdelivery of the information. Restricting the error checking to the header also simplifies the implementation in the nodes, requires less checksum bits, and prevents unnecessary packet discard. Some higher layers can tolerate some data errors, and higher layers also have the option of performing retransmission.

8.2. Identify the address class of the following IP addresses: 200.58.20.165; 128.167.23.20; 16.196.128.50; 50.156.10.10; 250.10.24.96.

Solution:

An IP address has a fixed length of 32 bits, where the most significant bits identify the particular class. Therefore, to identify the address class we need to convert the dotted-decimal notation back into its binary counterpart, and compare the binary notation to the class prefixes shown in Figure 8.5 in the text. (Recall that the dotted-decimal notation was devised to communicate addresses more readily to other people. In this notation, the 32 bits are divided into four groups of 8 bits – separated by periods – and then converted to their decimal counterpart.) The first few bits (shown in red) of the address can be used to determine the class.

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
128	64	32	16	8	4	2	1

200.58.20.165

11001000.00111010.00010100.10100101

Class C

128.167.23.20

10000000.10100111.00010111.00010100

Class B

16.196.128.50

00010000.11000100.10000000.00110010

Class A

150.156.10.10

10010110.10011100.00001010.00001010

Class B

250.10.24.96

11111010.00001010.00011000.01100000

Class E

8.3. Convert the IP addresses in Problem 8.2 to their binary representation.

Solution:

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
128	64	32	16	8	4	2	1

200.58.20.165
11001000.00111010.00010100.10100101

128.167.23.20
10000000.10100111.00010111.00010100

16.196.128.50
00010000.11000100.10000000.00110010

150.156.10.10
10010110.10011100.00001010.00001010

250.10.24.96
11111010.00001010.00011000.01100000

8.4. Identify the range of IPv4 addresses spanned by Class A, Class B, and Class C.

Solution:

The range of IPv4 addresses spanned by each class is:

Class A: 1.0.0.0 to 127.255.255.255
Class B: 128.0.0.0 to 191.255.255.255
Class C: 192.0.0.0 to 223.255.255.255

8.5. What are all the possible subnet masks for the Class C address space? List all the subnet masks in dotted-decimal notation, and determine the number of hosts per subnet supported for each subnet mask.

Solution:

255.255.255.128 supports 126 hosts (not including the broadcast address)

255.255.255.192 supports 62 hosts

255.255.255.224 supports 30 hosts

255.255.255.240 supports 14 hosts

255.255.255.248 supports 7 hosts

255.255.255.252 supports 3 hosts

255.255.255.254 and 255.255.255.255 are not practically usable.

8.6. A host in an organization has an IP address 150.32.64.34 and a subnet mask 255.255.240.0. What is the address of this subnet? What is the range of IP addresses that a host can have on this subnet?

Solution:

Address: 10010110 00100000 01000000 00100010

Mask: 11111111 11111111 11110000 00000000

Subnet: **10010110 00100000 01000000 00000000**

Host:

From: 10010110 00100000 01000000 **00000001**

To: 10010110 00100000 0100**1111 11111110**

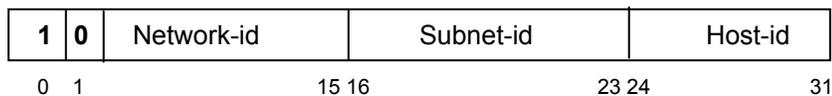
8.7. A university has 150 LANs with 100 hosts in each LAN.

Solutions follow questions:

(a) Suppose the university has one Class B address. Design an appropriate subnet addressing scheme.

A Class B address has 14 bits for the network ID and 16 bits for the host ID. To design an appropriate subnet addressing scheme we need to decide how many bits to allocate to the host ID versus the subnet ID. We can choose either 7 bits or 8 bits to identify the hosts.

If we allocate 8 bits for to identify the host, as shown below, then there are sufficient subnet-id bits to cover up to $2^8=256$ LANs and enough host-id bits to cover up to 256 hosts for each LAN. The subnet mask in this case is 255.255.255.0



Subnet mask: 255.255.255.0

If we allocate 7 bits for to identify the host, as shown below, then there are sufficient subnet-id bits to cover up to $2^9=512$ LANs and enough host-id bits to cover up to 128 hosts for each LAN. The subnet mask in this case is 255.255.255.128.

The choice between 7 or 8 bits to represent the hosts depends on which is likely to grow more, the number of subnets or the number of hosts in a LAN. Alternatively a variable-length prefix scheme using 7-bit host addresses, and grouping these form larger subnets provides greater flexibility in accommodating future changes.

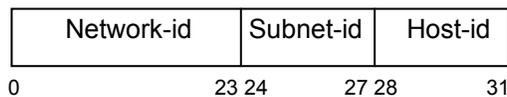
(b) Design an appropriate CIDR addressing scheme.

CIDR addressing scheme involves devising a prefix length that indicates the length of the network mask. In this case, 8 bits are required to identify each LAN (since $127 < 150 < 255$) and 7 bits are required to identify each host in each LAN (since $63 < 100 < 127$). Therefore a CIDR address would use a 17-bit prefix, and thus have an address of the form address/17.

8.8. A small organization has a Class C address for seven networks each with 24 hosts. What is an appropriate subnet mask?

Solution:

A Class C address requires 21 bits for its network ID, leaving 8 bits for the host ID and subnet ID to share. One possible scheme would assign 4 bits to the host and 4 to the subnet ID, as shown below. The number of bits assigned to the host can be increased to 5 as well.



Subnet mask: 255.255.255.224

8.9. A packet with IP address 150.100.12.55 arrives at router R1 in Figure 8.8. Explain how the packet is delivered to the appropriate host.

Solution:

The packet with IP address 150.100.12.55 arrives from the outside network. R1 has to know the next-hop router or host to send the packet to. The address corresponds to the binary string 10010110.01100100.00001100.00110111. R1 knows that a 9 bit subnet field is in use so it applies the following mask to extract the subnetwork address from the IP address.
11111111.11111111.11111111.10000000

The resulting IP address is 10010110.01100100.00001100.00000000 and corresponds 150.100.12.0. This indicates that the host is in subnet 150.100.12.0, so the router transmits the IP packet on this (attached) LAN.

8.10. In Figure 8.8 assign a physical layer address 1, 2, ... to each physical interface starting from the top row, moving right to left, and then moving down. Suppose H4 sends an IP packet to H1. Show the sequence of IP packets and Ethernet frames exchanged to accomplish this transfer.

Solution:

- Send IP packet from H4 to R1:
Source address 150.100.12.55 to destination IP address 150.100.12.176
Source Ethernet 4 to Receive Ethernet 6
- Forward IP packet from R1 to H1
Source address 150.100.12.55 to destination IP address 150.100.12.176
Source Ethernet 3 to Receive Ethernet 2

8.11. ARP is used to find the MAC address that corresponds to an IP address; RARP is used to find the IP address that corresponds to a MAC address. True or false?

Solution:

True, ARP is used to find the MAC address for a given IP address.

True, Reverse ARP is used by a device to find its IP address given its MAC address.

8.12. Perform CIDR aggregation on the following /24 IP addresses: 128.56.24.0/24; 128.56.25.0/24; 128.56.26.0/24; 128.56.27.0/24.

Solution:

```
128.56.24.0/22 = 10000000.00111000.00011000.00000000
128.56.25.0/22 = 10000000.00111000.00011001.00000000
128.56.26.0/22 = 10000000.00111000.00011010.00000000
128.56.27.0/22 = 10000000.00111000.00011011.00000000
mask           = 11111111.11111111.11111100.00000000
The resulting prefix is 128.56.24.0/22.
```

8.13. Perform CIDR aggregation on the following /24 IP addresses: 200.96.86.0/24; 200.96.87.0/24; 200.96.88.0/24; 200.96.89.0/24.

Solution:

```
200.96.86.0/20 = 11001000.01100000.01010110.00000000
200.96.87.0/20 = 11001000.01100000.01010111.00000000
200.96.88.0/20 = 11001000.01100000.01011000.00000000
200.96.89.0/20 = 11001000.01100000.01011001.00000000
mask           = 11111111.11111111.11110000.00000000
The resulting prefix is 200.96.80.0/20.
```

8.14. The following are estimates of the population of major regions of the world: Africa 900 million; South America 500 million; North America 400 million; East Asia 1500 million; South and Central Asia 2200 million; Russia 200 million; Europe 500 million.

Solutions follow questions:

- (a) Suppose each region is to be assigned 100 IP addresses per person. Is this possible? If not, how many addresses can be assigned per person? Repeat for IPv6.

The total number of IPv4 addresses is: 2^{32} which is approximately 4.29 billion. The above world population estimate totals 6.2 billion, so it is not possible to assign an individual address to each person. IPv6 is required to provide 100 addresses per person.

- (b) Design an appropriate CIDR scheme to provide the addressing in part (a).

Regions	(millions)	(millions)	IPV6 128 bits	
	Population	IP/Person	Net-id	Host-id
Africa	900	90000	91	37
South America	500	50000	92	36
North America	400	40000	92	36
East Asia	1500	150000	90	38
South C. Asia	2200	220000	90	38
Russia	200	20000	93	35
Europe	500	50000	92	36

8.15. Suppose four major ISPs were to emerge with points of presence in every major region of the world. How should a CIDR scheme treat these ISPs in relation to addressing for each major region?

Solution:

The networks of these ISPs will span across national and geographical boundaries and will be connected in a non-hierarchical manner. These large global ISPs constitute major transit routing domains, so it makes sense to assign them blocks of unique IP addresses and to require that domains attached to them should begin with the transit domain's prefix. The blocks of addresses should be managed so that fragmentation of the CIDR block does not take place. For example, when a customer of the ISP changes to another ISP, the customer is required to return the block of addresses. This policy enables CIDR to be effective in controlling the size of routing tables.

8.16. Discuss the difficulties with using actual time in the TTL field.

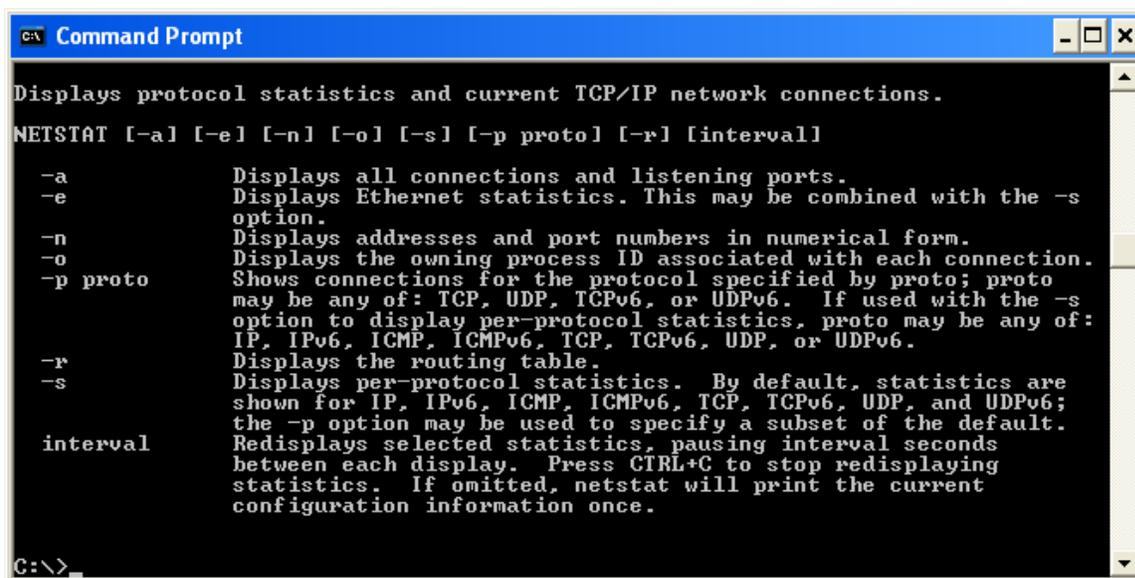
Solution:

Unlike the number of hops, which is predictable if a packet is routed correctly, the actual time that it takes to go through the route is not predictable. Therefore the amount of time that a packet stays in the network is not necessarily an indication of misrouting. To allow an upper limit for delay across the network TTL field would become a very large number. It is also more complex to track and update the TTL according to actual time as the packet traverses the network.

8.17. Lookup the `netstat` command in the manual for your system. Use the command to display the routing table in your host. Try the different command options.

Solution:

The exact command depends on your computing environment. In a DOS (Windows) environment, the command may be `netstat -r`. In a SUN/UNIX environment, the command may be `netstat -a`. See the answer to problem 45 in Chapter 1 for an example routing table. The figure below shows a screen capture of the parameters for the command in Windows XP.



```

C:\> netstat /?

Displays protocol statistics and current TCP/IP network connections.

NETSTAT [-a] [-e] [-n] [-o] [-s] [-p proto] [-r] [interval]

-a           Displays all connections and listening ports.
-e           Displays Ethernet statistics. This may be combined with the -s
            option.
-n           Displays addresses and port numbers in numerical form.
-o           Displays the owning process ID associated with each connection.
-p proto     Shows connections for the protocol specified by proto; proto
            may be any of: TCP, UDP, TCPv6, or UDPv6. If used with the -s
            option to display per-protocol statistics, proto may be any of:
            IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, or UDPv6.
-r           Displays the routing table.
-s           Displays per-protocol statistics. By default, statistics are
            shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6;
            the -p option may be used to specify a subset of the default.
interval    Redisplays selected statistics, pausing interval seconds
            between each display. Press CTRL+C to stop redisplaying
            statistics. If omitted, netstat will print the current
            configuration information once.

C:\>

```

8.18. Suppose a router receives an IP packet containing 600 data bytes and has to forward the packet to a network with maximum transmission unit of 200 bytes. Assume that the IP header is 20 bytes long. Show the fragments that the router creates and specify the relevant values in each fragment header (i.e., total length, fragment offset, and more bit).

Solution:

Given:

IP packet = 600 data bytes

MTU = 200 bytes

IP header = 20 header bytes

Maximum possible data length per fragment = MTU – IP header = 200 – 20 = 180 bytes.

The data length of each fragment must be a multiple of eight bytes; therefore the maximum number of data bytes that can be carried per fragment is $22 \times 8 = 176$.

The data packet must be divided into 4 frames, as shown by the following calculations:

$$176 + 176 + 176 + 72 = 600$$

$$\begin{array}{cccc} 20 & 20 & 20 & 20 \\ \hline 196 & 196 & 196 & 92 \end{array}$$

The sequence of frames and packet headers is shown below:

	Total length	Id	Mf	Fragment Offset
Original Packet	620	x	0	0
Fragment 1	196	x	1	0
Fragment 2	196	x	1	22
Fragment 3	196	x	1	44
Fragment 4	92	x	0	66

8.19. Design an algorithm for reassembling fragments of an IP packet at the destination IP.

Solution:

- I. Set data = Null
- II. Verify that all fragments for id = x have arrived
- III. Sort fragments in ascending order based on fragment offset
- IV. For each fragment starting with fragment offset = 0, move data = data + data-in-fragment,
- V. Data contains the reassembled information

8.20. Does it make sense to do reassembly at intermediate routers? Explain.

Solution:

No, because the packet may be de-fragmented again, and all the time required to wait for all fragments and to reassemble the packet will be wasted. Also it is not guaranteed that all fragments go through the same path and arrive at the same node in a datagram network such as IP.

8.21. Describe the implementation issues of IPv4 packet processing.

Solution:

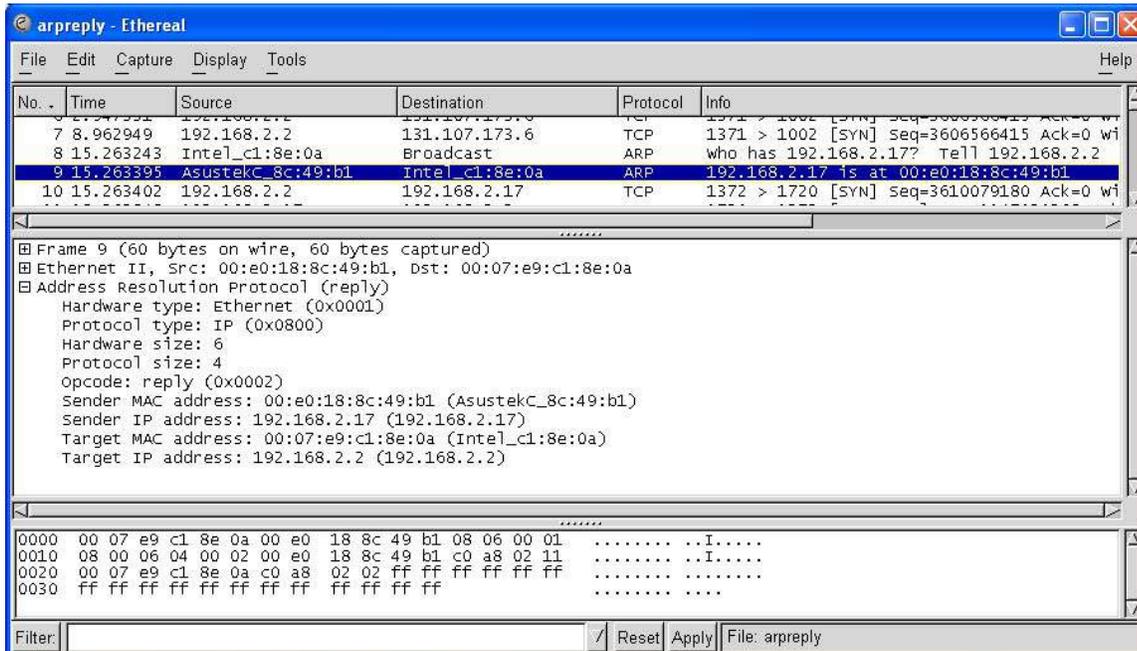
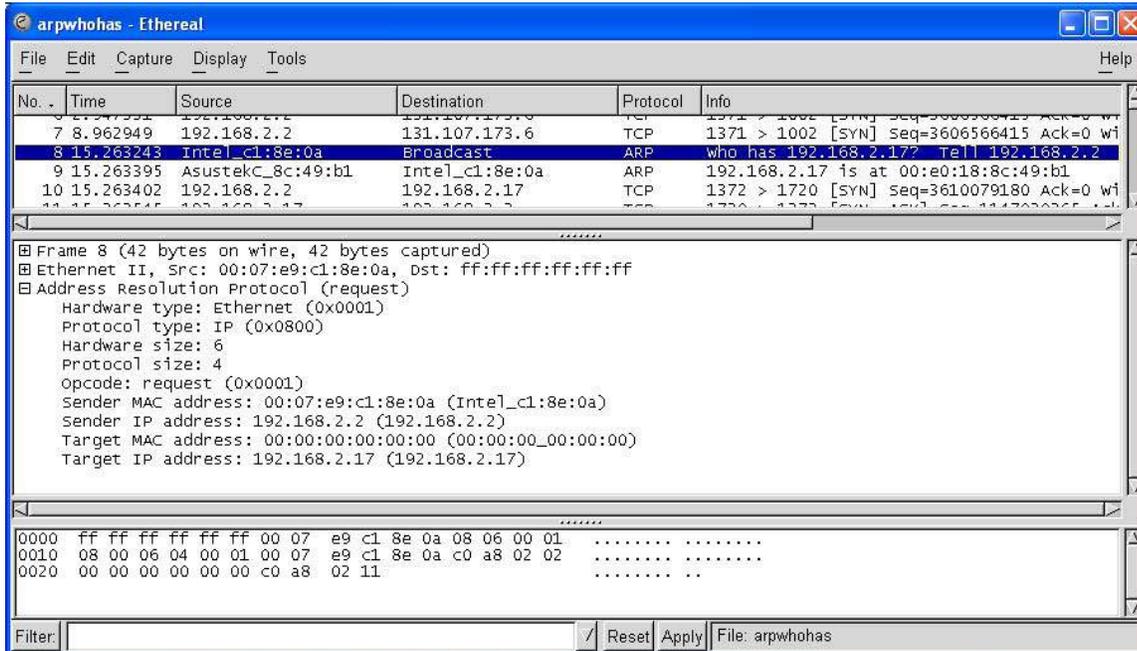
The router performs error checking: the header checksum is computed; the version and total length fields are checked for valid values. The router finds the next-hop by consulting its routing table. The

router then updates various fields, including the TTL and checksum fields. The packet is then forwarded.

8.22. Use Ethereal to capture ARP packets to find the MAC addresses in a LAN.

Solution:

The following screen captures show an ARP request and the corresponding ARP reply.



8.23. Abbreviate the following IPv6 addresses:

Solutions follow questions:

(a) 0000:0000:0F53:6382:AB00:67DB:BB27:7332

::F53:6382:AB00:67DB:BB27:7332

(b) 0000:0000:0000:0000:0000:0000:004D:ABCD

::4D:ABCD

(c) 0000:0000:0000:AF36:7328:0000:87AA:0398

::AF36:7328:0:87AA:398

(d) 2819:00AF:0000:0000:0000:0035:0CB2:B271

2819:AF::35:CB2:B271

8.24. What is the efficiency of IPv6 packets that carry 10 ms of 64 kbps voice? Repeat if an IPv6 packet carries 1 frame of 4 Mbps MPEG2 video, assuming a frame rate of 30 frames/second.

Solution:

10 ms of 64 kbps voice = $10 \times 10^{-3} \times 64 \times 10^3 = 640$ bits = 80 bytes
Header = 40 bytes; Efficiency = $80 / (80 + 40) = 2/3 = 0.6666 = 66.7\%$.

1 frame of video is: $4 \times 10^6 / 30 = 133,333$ bits = 16666 bytes
Efficiency = $16666 / (16666 + 40) = 99.76\%$.

8.25. Why does IPv6 allow fragmentation at the source only?

Solution:

The task of fragmenting a packet uses processing resources in a router. By requiring that all fragmentation be done at the source, routers are relieved of the fragmentation processing load, and hence they can operate faster on the basic routing task.

8.26. Assuming the population estimates in problem 8.14, how many IP addresses does IPv6 provide per capita?

Solution:

Based on the estimates in problem 10, there are $6,200 \times 10^6$ humans which in turn means 5.4×10^{28} IPv6 addresses per capita.

8.27. Suppose that IPv6 is used over a noisy wireless link. What is the effect of not having header error checking?

Solution:

The transmission over the noisy wireless link will introduce errors in the transmitted frames. If the frames do not contain error-checking, then assuming the frame is recognizable, erroneous packets may be passed to the router and unpredictable behavior may ensue. However, error checking (and retransmission) is included in most noisy wireless links; thus the effect of transmission errors is to trigger retransmission of frames in the link layer. Only packets that arrive in frames that pass error-checking are transferred to the IP layer.

8.28. Explain how the use of hierarchy enhances scalability in the following aspects of Internet:

Solutions follow questions:

(a) Domain name system

The use of hierarchy helps to speed up the translation of a domain name into an internet address. The search starts from the highest level (for example, .com, .org, .net) and eventually down to the specific hostname information. The hierarchy also helps to organize the database architecture of the DNS. Moreover, with a hierarchical system, it is not necessary for each DNS server to contain every single domain name in the network. There are different levels of DNS servers each containing the essential information for its own domain.

(b) IP addressing

Classful IP addressing uses hierarchy to arrange the address space in several discrete classes of addresses that correspond to networks of different sizes. CIDR IP addressing uses a variable-length prefix and a subnet mask to represent networks at a finer granularity of network size. In doing so, CIDR addressing increases the utilization of the address space. When combined with address allocation policies that aggregate routes, CIDR makes it possible to reduce the size of the routing tables required in router.

(c) OSPF routing

OSPF uses a two-level hierarchy that allows an AS to be partitioned into several groups called areas each interconnected by a central backbone area. This localization reduces the amount of routing information that needs to be maintained by individual routers. It also reduces the number of routing messages that need to be exchanged within the network.

(d) Interdomain routing

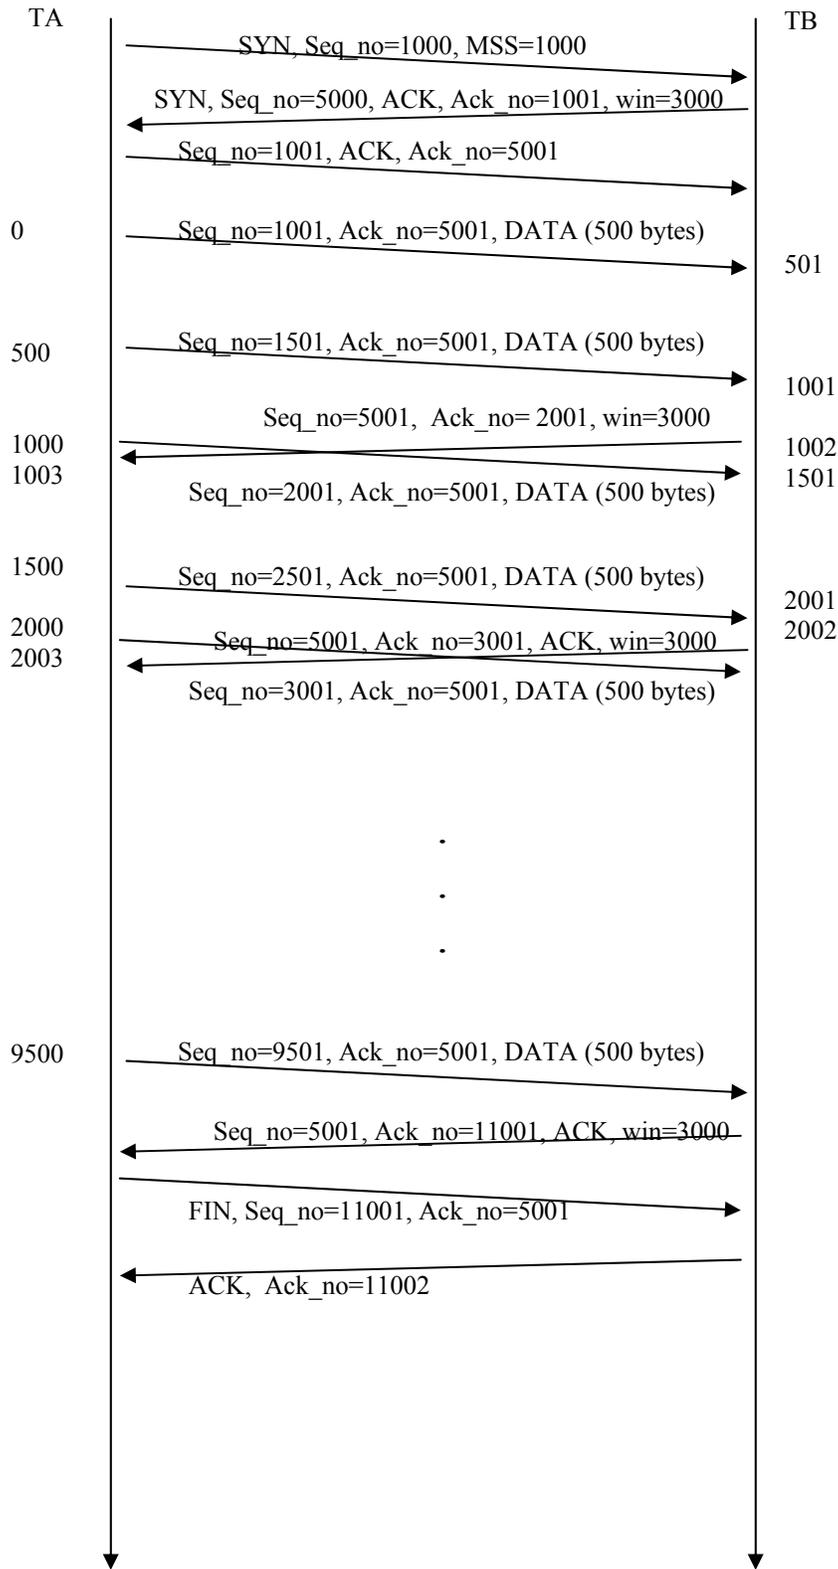
Interdomain routing uses the Border Gateway Protocol (BGP) to exchange routing information between AS's that in turn allows IP packets to flow across the AS border. Thus interdomain routing enables the scalability of the overall Internet by enabling various AS's to become interconnected.

8.29. The TCP in station A sends a SYN segment with ISN = 1000 and MSS = 1000 to station B. Station B replies with a SYN segment with ISN = 5000 and MSS = 500. Suppose station A has 10,000 bytes to transfer to B. Assume the link between stations A and B is 8 Mbps and the distance between them is 200 m. Neglect the header overheads to keep the arithmetic simple. Station B has 3000 bytes of buffer available to receive data from A. Sketch the sequence of segment exchanges, including the parameter values in the segment headers, and the state as a function of time at the two stations under the following situations:

Solutions follow questions:

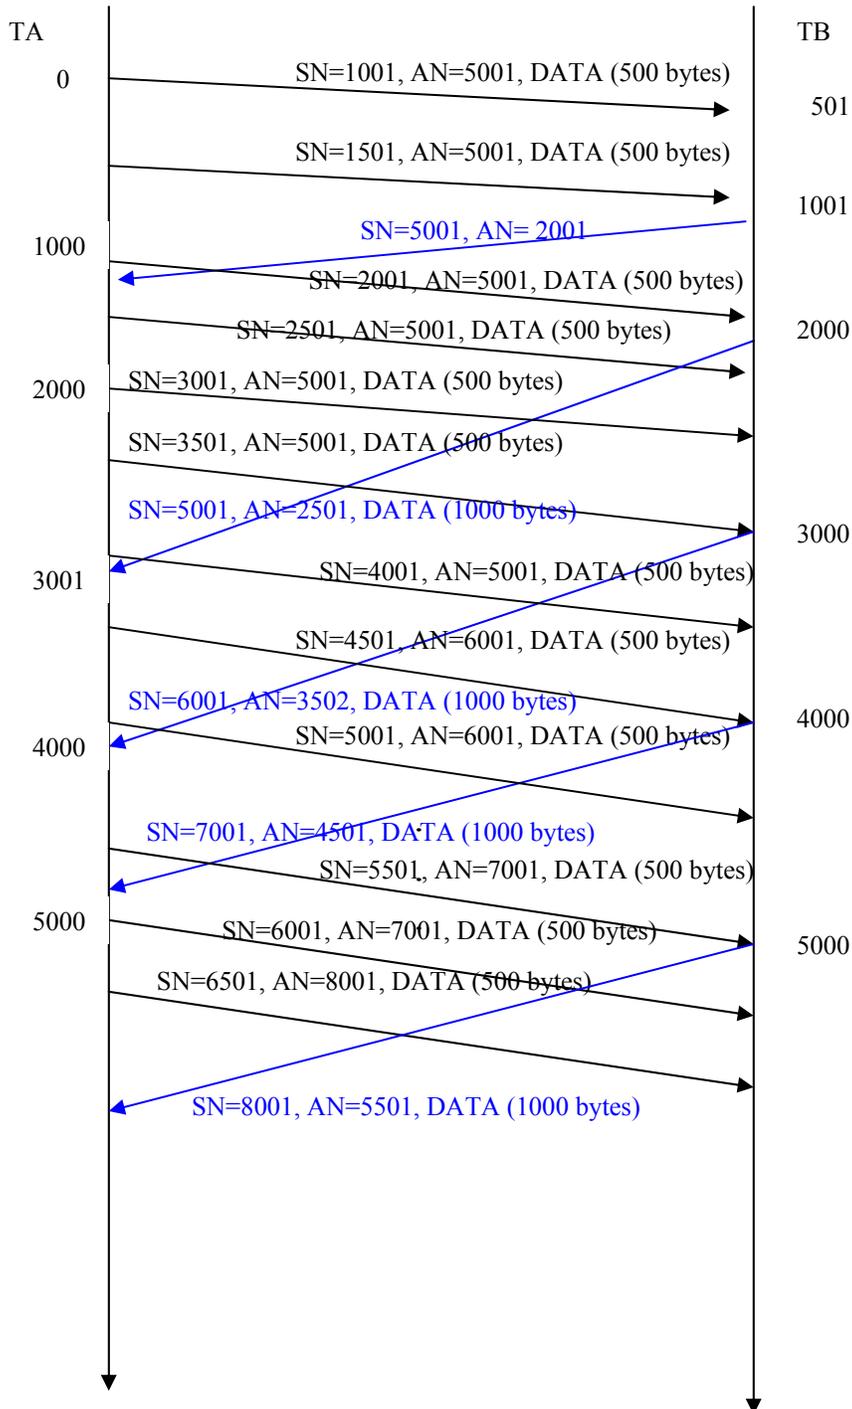
(a) Station A sends its first data segment at $t = 0$. Station B has no data to send and sends an ACK segment every other frame.

At a transmission rate of 8 megabits per second, a single byte has a transmission time of $8 \text{ bits} / 8 \times 10^6 \text{ bits/second} = 1 \text{ microsecond}$. A distance of 200 meters in optical fiber has a propagation time of $200 \text{ meters} / 2 \times 10^8 \text{ meters/second} = 1 \text{ microsecond}$. Therefore a segment of 500 bytes requires 501 microseconds to arrive completely at the receiver. In the following we also assume that the send window is replenished by the receiver as soon as it receives a segment. The time scale below is in microseconds.



- (b) Station A sends its first data segment at $t=0$. Station B has 6000 bytes to send, and it sends its first data segment at $t = 2$ ms.

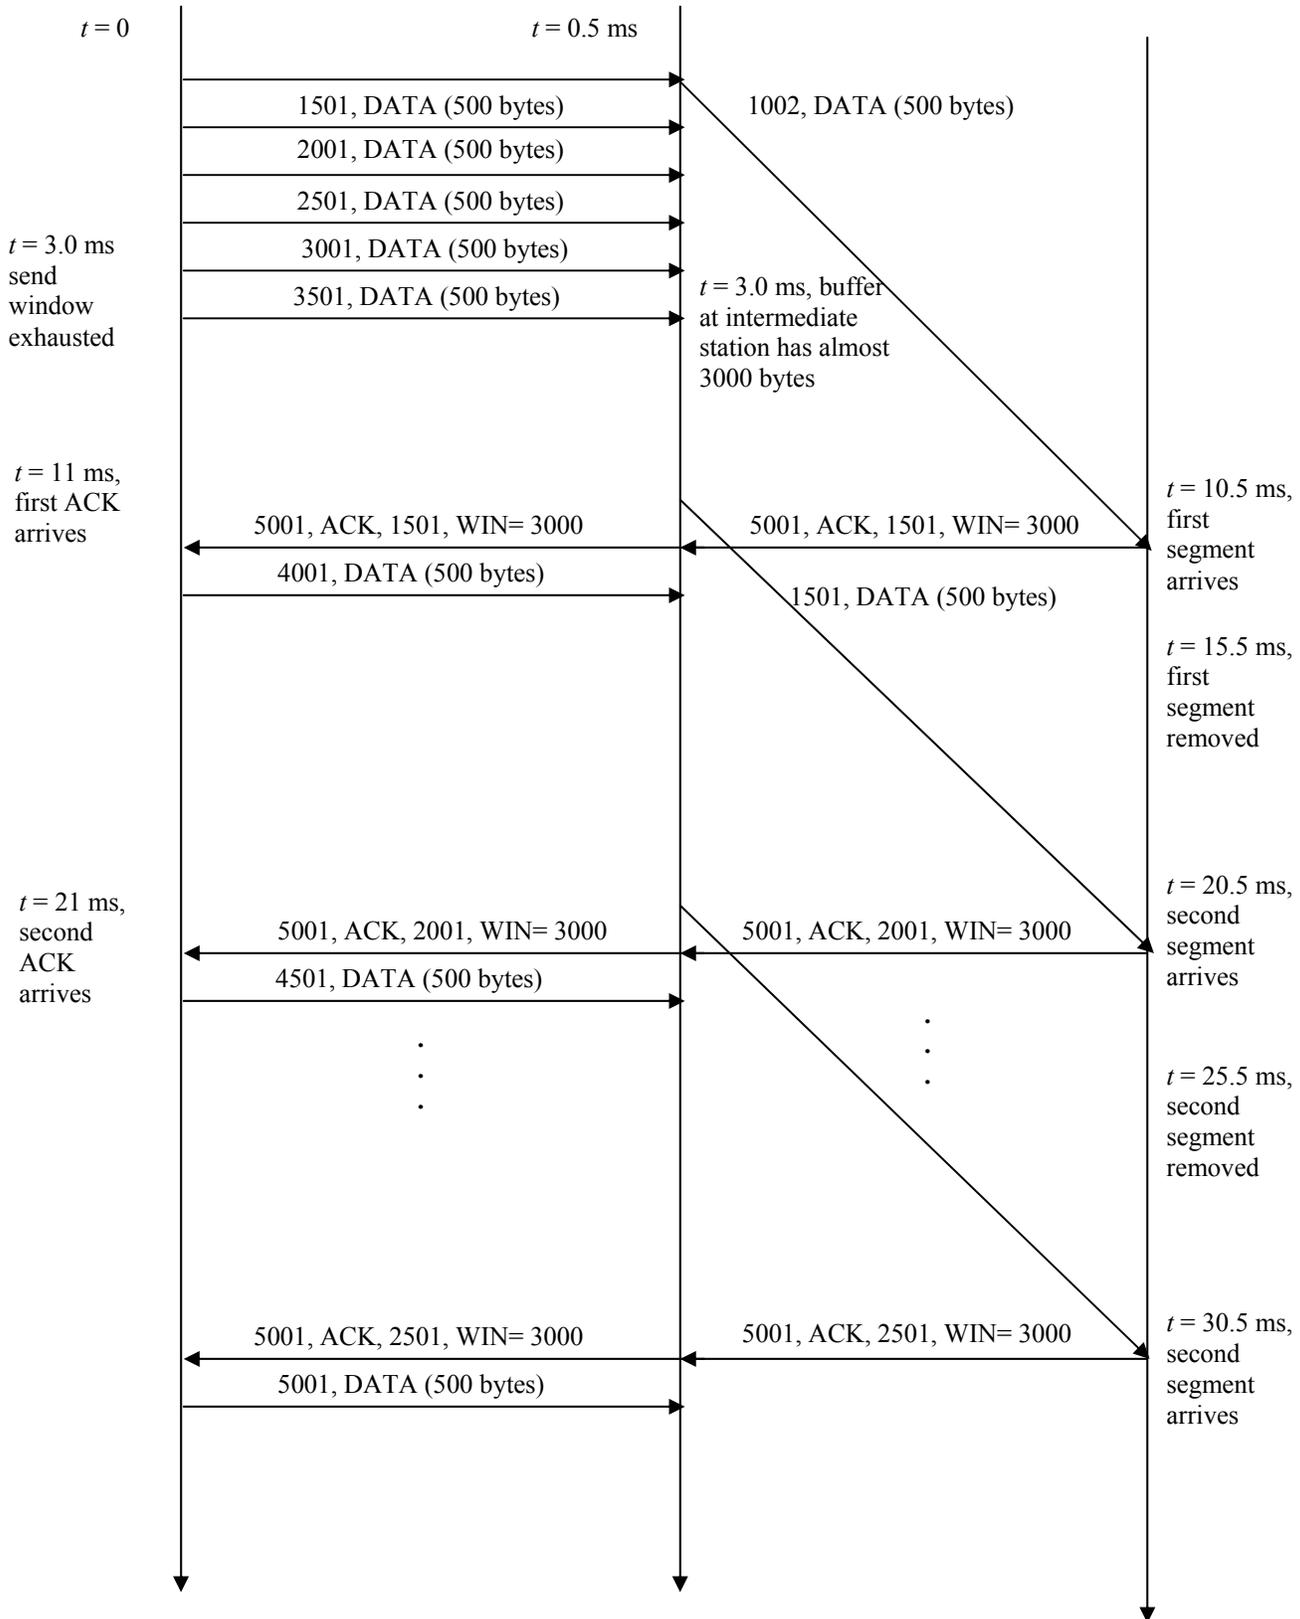
The main feature of this problem is that the acknowledgments are delayed longer because of the long segments that are transmitted from B to A.



8.30. Suppose that the TCP in station A sends information to the TCP in station B over a two-hop path. The data link in the first hop operates at a speed of 8 Mbps, and the data link in the second hop operates at a speed of 400 kbps. Station B has a 3 kilobyte buffer to receive information from A, and the application at station B reads information from the receive buffer at a rate of 800 kbps. The TCP in station A sends a SYN segment with ISN = 1000 and MSS = 1000 to station B. Station B replies with a SYN segment with ISN = 5000 and MSS = 500. Suppose station A has 10,000 bytes to transfer to B. Neglect the header overheads to keep the arithmetic simple. Sketch the sequence of segment exchanges, including the parameter values in the segment headers, and the state as a function of time at the two stations. Show the contents of the buffers in the intermediate switch as well as at the source and destination stations.

Solution:

It takes 500 microseconds to transmit 500 bytes from station A to the intermediate station, but it takes 10 milliseconds to send the same 500 bytes from the intermediate station to station B. Hence, segments will accumulate at the intermediate station until station A exhausts its send window of 3000 bytes. Eventually station A will receive acknowledgments that allow it to resume transmission. Note that the rate at which acknowledgments are returned to station A is controlled by the rate at which segments arrive at station B from the bottleneck at the intermediate node.



8.31. Suppose that the delays experienced by TCP segments traversing the network is equally likely to be any value in the interval [50 ms, 75 ms]. (See Equations 5.17 to 5.20.)

Solutions follow questions:

(a) Find the mean and standard deviation of the delay.

The delay lies between interval [50ms, 75ms] and is a uniform random variable. The mean is:

$$E[X] = (50 + 75) / 2 = 62.5 \text{ ms.}$$

The standard deviation of delay is:

$$STD[X] = VAR[X]^{1/2} = [(75 - 50)^2 / 12]^{1/2} = 7.217$$

(b) Most computer languages have a function for generating uniformly distributed random variables. Use this function in a short program to generate random times in the above interval. Also, calculate t_{RTT} and d_{RTT} and compare to part (a).

See below for sample program written in C.

```

/* Communication Networks - Chapter 8      */
/* Question 25 (b)                        */
/* Description - Generate a random value  */
/* between 50 to 75 ms. Calculate t_RTT  */
/* and d_RTT                              */
/* The min, max and avg value of t_RTT   */
/* and d_RTT are also recorded            */

#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main (void)
{
    int i;
    float temp, t_n, t_rtt_new, t_rtt_old;
    float d_rtt_new, d_rtt_old;
    float t_rtt_min, t_rtt_sum, t_rtt_max;
    float d_rtt_min, d_rtt_sum, d_rtt_max;

    const float alpha = 0.875;
    const float beta = 0.25;

    srand (time(NULL));
    t_rtt_old = 0;
    d_rtt_old = 0;

    t_rtt_sum = t_rtt_max = 0;
    d_rtt_sum = d_rtt_max = 0;

    t_rtt_min = d_rtt_min = 500;

    for (i = 0; i < 500; i++)
    {

        /* Generate a random value between 0 to 1 */
        temp = (float) rand() / RAND_MAX;

        /* Scale the random value to fit between 50 to 75 */
        t_n = temp * 25 + 50;

```

```

/* Calculate t_RTT and d_RTT */
t_rtt_new = (alpha * t_rtt_old) + ((1 - alpha) * t_n);
d_rtt_new = (beta * d_rtt_old) + ((1 - beta) * fabs(t_n - t_rtt_old));

if (t_rtt_new < t_rtt_min)
    t_rtt_min = t_rtt_new;
if (t_rtt_new > t_rtt_max)
    t_rtt_max = t_rtt_new;

if (d_rtt_new < d_rtt_min)
    d_rtt_min = d_rtt_new;
if (d_rtt_new > d_rtt_max)
    d_rtt_max = d_rtt_new;

t_rtt_sum += t_rtt_new;
d_rtt_sum += d_rtt_new;

printf ("t_RTT: %f d_RTT: %f\n", t_rtt_new, d_rtt_new);
t_rtt_old = t_rtt_new;
d_rtt_old = d_rtt_new;
}
printf ("t_RTT min: %f t_RTT max: %f t_RTT avg: %f\n",
        t_rtt_min, t_rtt_max, (t_rtt_sum / 500.0));
printf ("d_RTT min: %f d_RTT max: %f d_RTT avg: %f\n",
        d_rtt_min, d_rtt_max, (d_rtt_sum / 500.0));
}

```

We ran the preceding program and obtained the average values of $t_{RTT} = 61.6924$ and $d_{RTT} = 7.1139$. These values are averaged from a sample of 500 values.

8.32. Suppose that the advertised window is 1 Mbyte long. If a sequence number is selected at random from the entire sequence number space, what is the probability that the sequence number falls inside the advertised window?

Solution:

If the sequence number field is 32 bits in length and the advertised window is 1Mbyte long, the probability that the sequence number falls inside the advertised window is:

$$P = (1 \times 10^6) / 2^{32} = 2.33 \times 10^{-4}$$

8.33. Explain the relationship between advertised window size, RTT, delay-bandwidth product, and the maximum achievable throughput in TCP.

Solutions follow questions:

- (a) Plot the maximum achievable throughput versus delay-bandwidth product for an advertised window size of 65,535 bytes.

First consider delay-bandwidth product, $DBP = R \cdot 2t_p$. Here delay $2t_p$ is the propagation time that elapses from when a bit is sent by a source to the destination to when the bit can be returned back to the source. This is the minimum time that elapses from when a packet leaves a source to when the acknowledgment is received. The delay-bandwidth product DBP is then the number of bits (or bytes) that are in the network when the source transmits continuously at the maximum rate and when the bits return immediately back to the source.

The round-trip time RTT is the time that actually elapses from when a packet is sent to when its acknowledgment is received. RTT includes not only the propagation delay, but also queueing and

processing delays. The advertised window, W , places a limit on the amount of information that a source can have outstanding in the network.

Consider the time from when a byte leaves the source to when its acknowledgment is received (that is, consider a RTT). In that time, the source will have transmitted at most a window-full of bytes into the network. Therefore the window size divided by the RTT places a limit on the throughput r , that is, the rate at which information can be transmitted into the network: $r < W/RTT$.

The throughput cannot exceed the maximum bit rate $R = DBP/2t_p$ that is available for the source to transmit into the network. Therefore, the throughput increases as the window size is increased, but cannot exceed the bit rate R :

$$\text{Throughput} = r = \min\{R, W/RTT\} = \min\{DBP/2t_p, W/RTT\}$$

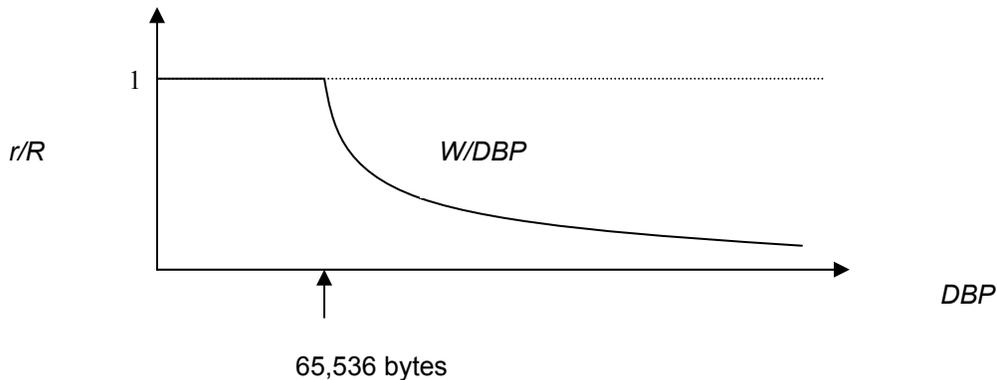
Suppose that the window size is less than the delay-bandwidth product. We then expect that the source cannot transmit at the maximum bit rate R . Indeed, we have that:

$$r < W/RTT < W/2t_p.$$

Therefore we have that:

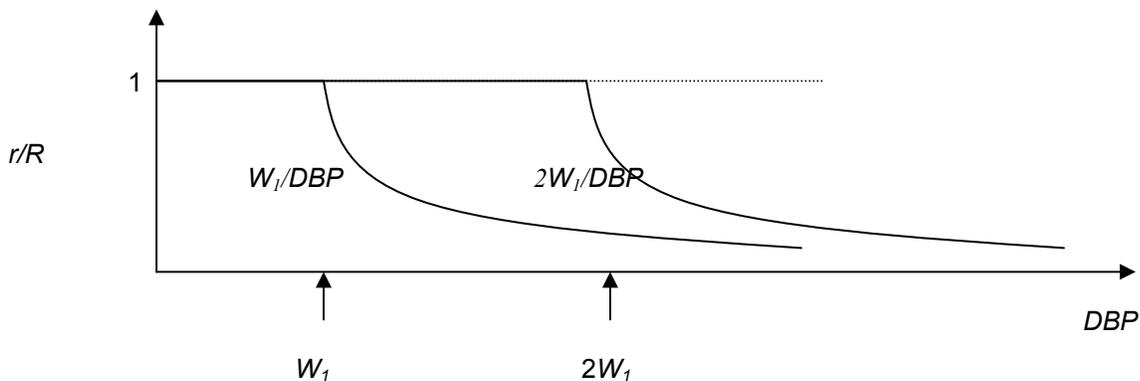
$$r/R < W/(R \cdot 2t_p) = W/DBP.$$

We conclude that the ratio of the maximum achievable throughput to R is less than the ratio of the window size to the DBP , as shown in the figure below.



(b) In the above plot include the maximum achievable throughput when the above window size is scaled up by a factor of $2K$, where $K = 4, 8, 12$.

The following figure shows the case where the window size is doubled.



- (c) Place the following scenarios in the plot obtained in part (b): Ethernet with 1 Gbps and distance 100 meters; 2.4 Gbps and distance of 6000 km; satellite link with speed of 45 Mbps and RTT of 500 ms; 40 Gbps link with distance of 6000 km.

Case	DBP implied
Ethernet $R = 1$ Gbps $D = 100$ m	$T_p = 100 / 2.5 \times 10^8$ $T_p = 4 \times 10^{-7}$ $DBP = 2 * T_p * R$ $DBP = 8 \times 10^{-7} * 1 \times 10^9$ $DBP = 8 \times 10^2$ DBP = 800 bits
Link $R = 2.4$ Gbps $D = 6000$ km	$T_p = 6 \times 10^3 / 2.5 \times 10^8$ $T_p = 2.4 \times 10^{-2}$ $DBP = 2 * T_p * R$ $DBP = 2 * 2.4 \times 10^{-2} * 2.4 \times 10^9$ $DBP = 11.52 \times 10^7$ DBP = 115.2 Mbits (14.4 Mbytes)
Satellite link $R = 45$ Mbps $RTT = 500$ ms (5×10^{-1} sec)	$DBP = RTT * R$ $DBP = 5 \times 10^{-1} * 45 \times 10^6$ $DBP = 225 \times 10^5$ bits DBP = 22.5 Mbits (2.85 Mbytes)
link $R = 40$ Gbps $D = 6000$ km (6×10^6 m)	$T_p = 6 \times 10^3 / 2.5 \times 10^8$ $T_p = 2.4 \times 10^{-2}$ $DBP = 2 * T_p * R$ $DBP = 2 * 2.4 \times 10^{-2} * 40 \times 10^9$ $DBP = 192 \times 10^7$ DBP = 1.92 Gbits (240 Mbytes)

8.34. Consider the three-way handshake in TCP connection setup.

Solutions follow questions:

- (a) Suppose that an old SYN segment from station A arrives at station B, requesting a TCP connection. Explain how the three-way handshake procedure ensures that the connection is rejected.

In a three-way handshake procedure, one must ensure the selection of the initial sequence number is always unique. If station B receives an old SYN segment from A, B will acknowledge the request based on the old sequence number. When A receives the acknowledgment segment from B, A will find out that B received a wrong sequence number. A will discard the acknowledgment packet and reset the connection.

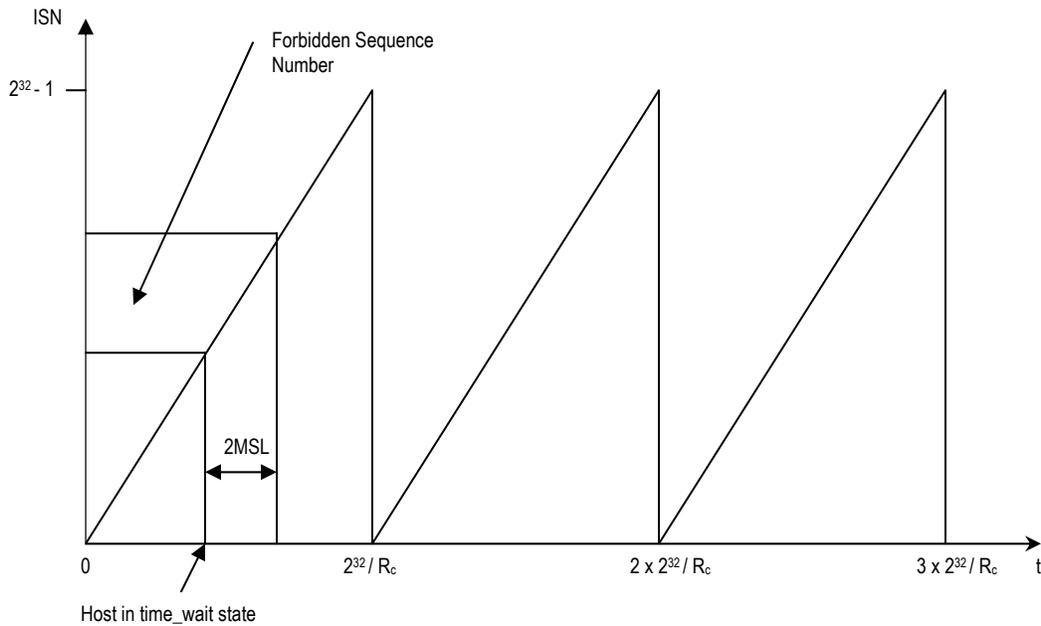
- (b) Now suppose that an old SYN segment from station A arrives at station B, followed a bit later by an old ACK segment from A to a SYN segment from B. Is this connection request also rejected?

If an old SYN segment from A arrives at B, followed by an old ACK segment from A to a SYN segment from B, the connection will also be rejected. Initially, when B receives an old SYN segment, B will send a SYN segment with its own distinct sequence number set by itself. If B receives the old ACK from A, B will notify A that the connection is invalid since the old ACK sequence number does not match the sequence number previously defined by B. Therefore, the connection is rejected.

8.35. Suppose that the Initial Sequence Number (ISN) for a TCP connection is selected by taking the 32 low-order bits from a local clock.

Solutions follow questions:

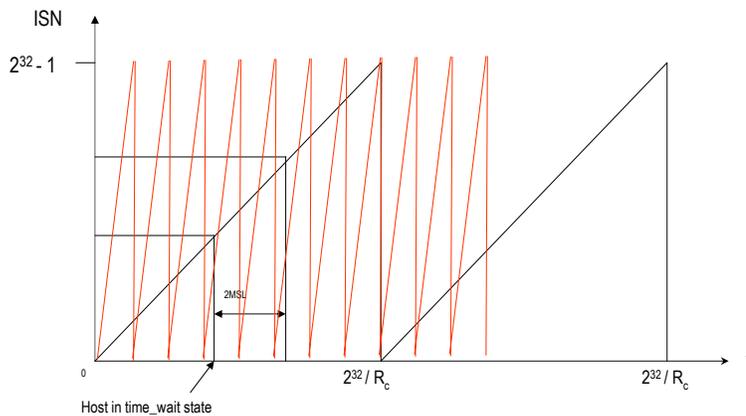
- (a) Plot the ISN versus time assuming that the clock ticks forward once every $1/R_c$ seconds. Extend the plot so that the sequence numbers wrap around.



- (b) To prevent old segments from disrupting a new connection, we forbid sequence numbers that fall in the range corresponding to $2MSL$ seconds prior to their use as an ISN. Show the range of forbidden sequence numbers versus time in the plot from part (a).

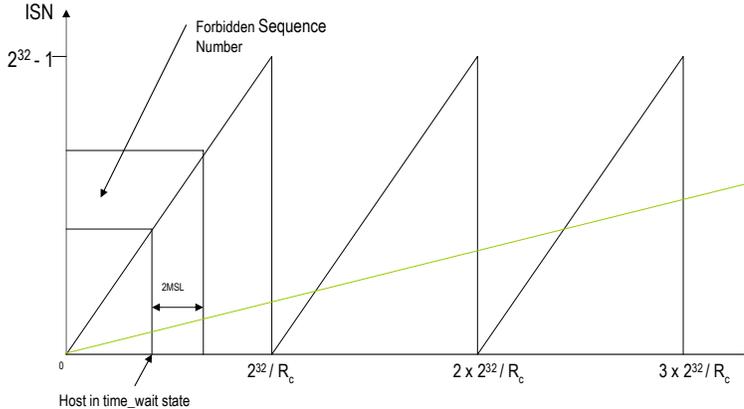
See above graph.

- (c) Suppose that the transmitter sends bytes at an average rate $R > R_c$. Use the plot from part (b) to show what goes wrong.



If the transmitter sends data at an average rate $R > R_c$, the ISN will lag behind the transmitter's sequence number. In particular, if the source uses all the sequence numbers in less than 2 MSL, then all the sequence numbers would be forbidden as ISNs for the next connection.

- (d) Now suppose that the connection is long-lived and that bytes are transmitted at a rate R that is much lower than R_c . Use the plot from part (b) to show what goes wrong. What can the transmitter do when it sees that this problem is about to happen?

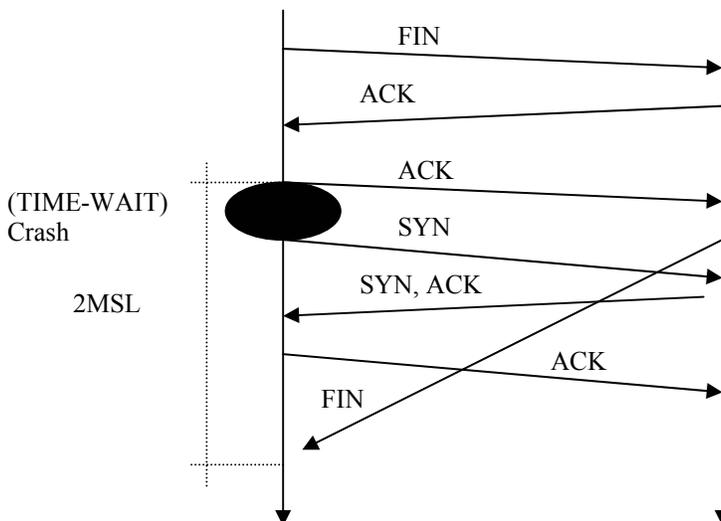


When the connection is long-lived and $R < R_c$, the ISN sequence number will wrap around at a much faster rate than the transmitter sequence number. When a new connection is established, it is possible for the ISN to be selected within the range of sequence numbers used by the slow connection.

8.36. Suppose that during the TCP connection closing procedure, a machine that is in the TIME_WAIT state crashes, reboots within MSL seconds, and immediately attempts to reestablish the connection using the same port numbers. Give an example that shows that delayed segments from the previous connections can cause problems. For this reason RFC 793 requires that for MSL seconds after rebooting TCP is not allowed to establish new connections.

Solution:

A delayed FIN from the earlier connection causes the new connection to be closed prematurely.



8.37. Are there any problems if the server in a TCP connection initiates an active close?

Solution:

As TCP is defined, no problems should arise if the server initiates an active close. Recall from Figure 8.36 that the side that does the active close (by issuing the first FIN segment) will enter the TIME_WAIT state. The side that does the passive close does not. Therefore, when the server does the active close, it will go into the TIME_WAIT state and hence will not be able to be restarted with its same (well-known) port number, because this is part of the parameters that are set aside during the 2MSL wait.

8.38. Use a network analyzer to capture the sequence of packets in a TCP connection. Analyze the contents of the segments that open and close the TCP connection. Estimate the rate at which information is transferred by examining the frame times and the TCP sequence numbers. Do the advertised windows change during the course of the connection?

Solution:

The following sequence of packet captures was obtained by connecting to www.yahoo.com using telnet. The TCP open and close can be observed at the beginning and end of the packet sequence. Advertised window sizes, acknowledgments, and sequence numbers are also shown. (The data was obtained from Ethereal using the print-to-file option.)

No.	Time	Source	Destination	Protocol	Info
1	0.000000	Intel_c1:8e:0a	Broadcast	ARP	Who has 192.168.2.3? Tell 192.168.2.2
2	12.992506	192.168.2.18	192.168.2.1	DNS	Standard query A www.yahoo.com
3	13.001008	192.168.2.1	192.168.2.18	DNS	Standard query response CNAME www.yahoo.akadns.net A 216.109.125.79 A
216.109.125.72	A	216.109.125.69	A	216.109.117.205	A 216.109.125.78 A 216.109.125.71 A 216.109.125.64 A 216.109.118.64
4	13.001678	192.168.2.18	216.109.125.79	TCP	2498 > http [SYN] Seq=147142992 Ack=0 Win=8192 Len=0
5	13.039151	216.109.125.79	192.168.2.18	TCP	http > 2498 [SYN, ACK] Seq=2183346772 Ack=147142993 Win=65535 Len=0
6	13.039221	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142993 Ack=2183346773 Win=8760 Len=0
7	18.472270	192.168.2.18	216.109.125.79	TCP	2498 > http [PSH, ACK] Seq=147142993 Ack=2183346773 Win=8760 Len=1
8	18.600842	00000000.0001031d	cccf7 00000000.	Broadcast	IPX SAP General Query
9	18.622879	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142994 Win=65535 Len=0
10	18.734094	192.168.2.18	216.109.125.79	TCP	2498 > http [PSH, ACK] Seq=147142994 Ack=2183346773 Win=8760 Len=1
11	18.905241	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142995 Win=65535 Len=0
12	18.924462	192.168.2.18	216.109.125.79	TCP	2498 > http [PSH, ACK] Seq=147142995 Ack=2183346773 Win=8760 Len=1
13	19.078556	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142996 Win=65535 Len=0
14	19.244070	192.168.2.18	216.109.125.79	TCP	2498 > http [PSH, ACK] Seq=147142996 Ack=2183346773 Win=8760 Len=1
15	19.369356	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142997 Win=65535 Len=0
16	19.692438	192.168.2.18	216.109.125.79	TCP	2498 > http [PSH, ACK] Seq=147142997 Ack=2183346773 Win=8760 Len=2
17	19.755278	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183346773 Ack=147142999 Win=65535 Len=1460
18	19.756467	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183348233 Ack=147142999 Win=65535 Len=1460
19	19.756515	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183349693 Win=8760 Len=0
20	19.758513	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183349693 Ack=147142999 Win=65535 Len=1460
21	19.843349	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183351153 Ack=147142999 Win=65535 Len=1460
22	19.843467	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183352613 Win=8760 Len=0
23	19.845612	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183352613 Ack=147142999 Win=65535 Len=1460
24	19.846719	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183354073 Ack=147142999 Win=65535 Len=1460
25	19.846764	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183355533 Win=8760 Len=0
26	19.899221	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183355533 Ack=147142999 Win=65535 Len=1460
27	19.900426	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183356993 Ack=147142999 Win=65535 Len=1460
28	19.900508	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183358453 Win=8760 Len=0
29	19.902280	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183358453 Ack=147142999 Win=65535 Len=1460
30	19.911624	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183359913 Ack=147142999 Win=65535 Len=1460
31	19.911745	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183361373 Win=8760 Len=0
32	19.916041	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183361373 Ack=147142999 Win=65535 Len=1460
33	19.917275	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183362833 Ack=147142999 Win=65535 Len=1460
34	19.917335	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183364293 Win=8760 Len=0
35	19.942667	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183364293 Ack=147142999 Win=65535 Len=1460
36	19.943834	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183365753 Ack=147142999 Win=65535 Len=1460
37	19.943914	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183367213 Win=8760 Len=0
38	19.959414	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183367213 Ack=147142999 Win=65535 Len=1460
39	19.962971	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183368673 Ack=147142999 Win=65535 Len=1460
40	19.963078	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183370133 Win=8760 Len=0
41	19.964934	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183370133 Ack=147142999 Win=65535 Len=1460
42	19.965690	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183371593 Ack=147142999 Win=65535 Len=1460
43	19.965733	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183373053 Win=8760 Len=0
44	19.987210	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183373053 Ack=147142999 Win=65535 Len=1460
45	19.994098	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183374513 Ack=147142999 Win=65535 Len=1460
46	19.994170	192.168.2.18	216.109.125.79	TCP	2498 > http [ACK] Seq=147142999 Ack=2183375973 Win=8760 Len=0
47	20.013318	216.109.125.79	192.168.2.18	TCP	http > 2498 [ACK] Seq=2183375973 Ack=147142999 Win=65535 Len=1460

```

48 20.016425 216.109.125.79 192.168.2.18 TCP http > 2498 [ACK] Seq=2183378893 Ack=147142999 Win=65535 Len=1460
49 20.016538 192.168.2.18 216.109.125.79 TCP 2498 > http [ACK] Seq=147142999 Ack=2183377433 Win=8760 Len=0
50 20.018301 216.109.125.79 192.168.2.18 TCP http > 2498 [ACK] Seq=2183380353 Ack=147142999 Win=65535 Len=1460
51 20.018342 192.168.2.18 216.109.125.79 TCP 2498 > http [ACK] Seq=147142999 Ack=2183377433 Win=8760 Len=0
52 20.034934 216.109.125.79 192.168.2.18 TCP http > 2498 [FIN, PSH, ACK] Seq=2183381813 Ack=147142999 Win=65535 Len=697
53 20.035032 192.168.2.18 216.109.125.79 TCP 2498 > http [ACK] Seq=147142999 Ack=2183377433 Win=8760 Len=0
54 21.275248 216.109.125.79 192.168.2.18 TCP http > 2498 [ACK] Seq=2183377433 Ack=147142999 Win=65535 Len=1460
55 21.275415 192.168.2.18 216.109.125.79 TCP 2498 > http [ACK] Seq=147142999 Ack=2183382510 Win=8760 Len=0
56 21.346468 216.109.125.79 192.168.2.18 TCP http > 2498 [FIN, ACK] Seq=2183382510 Ack=147142999 Win=65535 Len=0
57 21.346585 192.168.2.18 216.109.125.79 TCP 2498 > http [ACK] Seq=147142999 Ack=2183382511 Win=8760 Len=0
58 23.239449 192.168.2.18 216.109.125.79 TCP 2498 > http [FIN, ACK] Seq=147142999 Ack=2183382511 Win=8760 Len=0
59 23.273437 216.109.125.79 192.168.2.18 TCP http > 2498 [ACK] Seq=2183382511 Ack=147143000 Win=65535 Len=0

```

8.39. Devise an experiment to use a network analyzer to observe the congestion control behavior of TCP. How would you obtain Figure 8.37 empirically? Run the experiment and plot the results.

Solution:

Congestion involves the buildup of packets in a buffer and can be triggered by the sustained arrivals of packets from a high speed network, e.g. a LAN, to a router feeding a slow-speed network, e.g. a dialup modem to an ISP. Thus one way of observing congestion is to send a stream of packets from a home LAN onto the Internet via a dialup modem. Congestion can also occur when multiple users send packets on multiple inputs to the same output port on a router. Thus a second way of generating congestion is to have several machines simultaneously send a stream of packets to the Internet via a home router that connects to a DSL or cable modem. A packet capture tool such as Ethereal can be used to track the evolution of sequence numbers and segment retransmissions over time. However, the congestion window is controlled by TCP which operates in the kernel of the OS. To obtain traces of `cwnd` versus time a tool such as `tcpdump` with debug option, or TCP instrumentation tools (such as provided by the Web100 project for Linux) need to be used.

8.40. A fast typist can do 100 words a minute, and each word has an average of 6 characters. Demonstrate Nagle's algorithm by showing the sequence of TCP segment exchanges between a client, with input from our fast typist, and a server. Indicate how many characters are contained in each segment sent from the client. Consider the following two cases:

Solutions follow questions:

The typist types 100 words per minutes, averaging 6 characters per word. This is equivalent to 600 characters per minute or 10 characters per second. Therefore, the typist can type a character every 100 ms.

(a) The client and server are in the same LAN and the RTT is 20 ms.

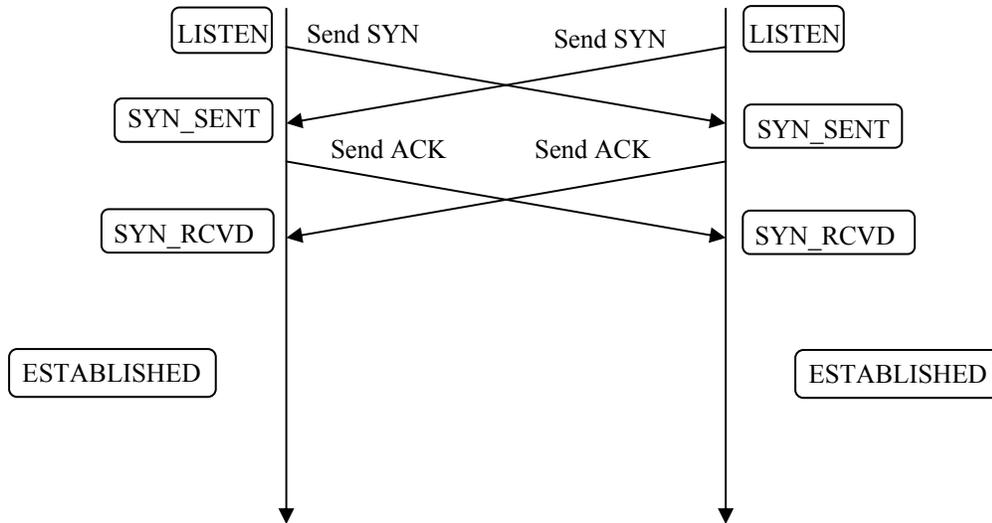
In this case Nagle's Algorithm is not activated since acknowledgments arrive before the next character is typed. Each client segment is 41 bytes long assuming IP and TCP headers are 20 bytes each.

(b) The client and server are connected across a WAN and the RTT is 100 ms.

In this case one or two characters are typed before an acknowledgment is received. Therefore, segments are either 41 or 42 bytes long.

8.41. Simultaneous Open. The TCP state transition diagram allows for the case where the two stations issue a SYN segment at nearly the same time. Draw the sequence of segment exchanges and use Figure 8.36 to show the sequence of states that are followed by the two stations in this case.

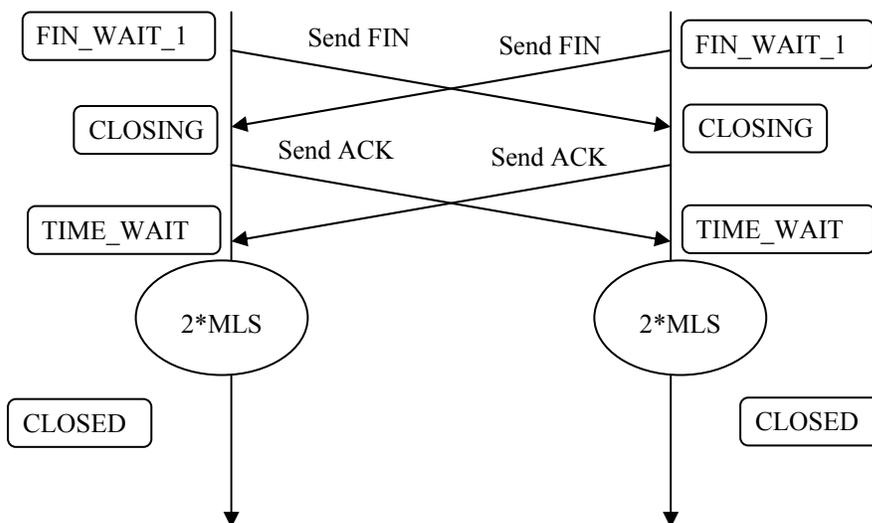
Solution:



8.42. Simultaneous Close. The TCP state transition diagram allows for the case where the two stations issue a FIN segment at nearly the same time. Draw the sequence of segment exchanges and use Figure 8.36 to show the sequence of states that are followed by the two stations in this case.

Solution:

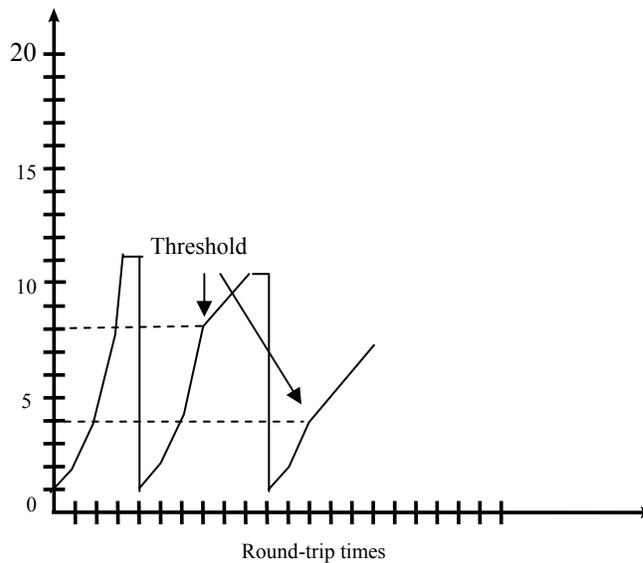
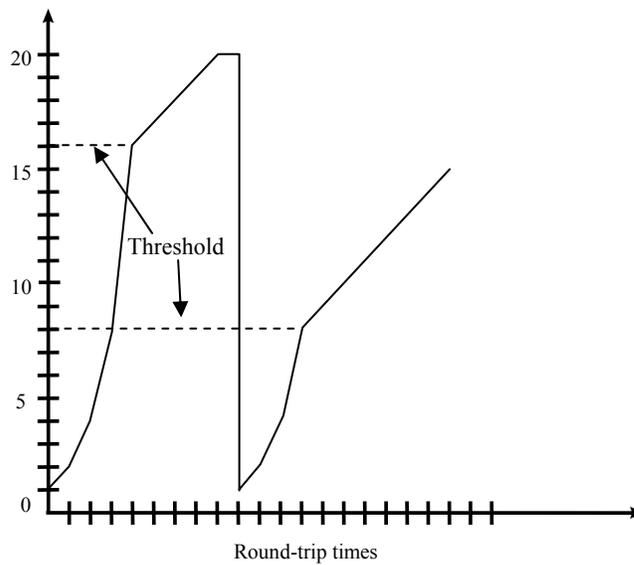
The sequence of the state transition will be the same for two hosts. After FIN sent, it moves from **ESTABLISHED** to **FIN_WAIT_1**. Once the host receives the FIN from other host, it sends an **ACK** and move from **FIN_WAIT_1** to **CLOSING**. Finally, when the host received **ACK** from each other, it moves from **CLOSING** to **TIME_WAIT**. After **2*MSL**, both hosts transit back to **CLOSED** state.



8.43. Suppose that a TCP source (with unlimited amount of information to transmit) begins transmitting onto a link that has 1 Mbps in available bandwidth. Sketch congestion window versus time trajectory. Now suppose that another TCP source (also with unlimited amount of information to transmit) begins transmitting over the same link. Sketch the congestion window versus the time for the initial source.

Solution:

Initially the TCP source has about 1 Mbps of available bandwidth, so its congestion avoidance behavior will begin as the rate approaches 1 Mbps. When the available bandwidth drops to 500 kbps, the congestion window behavior will begin sooner. The following two figures show the corresponding congestion window trajectories.



8.44. What is the maximum width of a RIP network?

Solution:

The maximum width of an RIP network is 15 nodes. Since the maximum number of hops in RIP network is 15, node 16 represents infinity.

8.45. Let's consider the bandwidth consumption of the RIP protocol.

Solutions follow questions:

(a) Estimate the number of messages exchanged per unit time by RIP.

A router implementing RIP sends an update message every 30 seconds. Assume that the typical node has D neighbors and that the number of nodes in the network is N . The number of message exchanges is then $ND/30$ messages per second.

(b) Estimate the size of the messages exchanged as a function of the size of the RIP network.

A RIP message consists of a four-byte header plus 20-byte per entry and up to 25 entries per message. Therefore, each node will send out a message of at most $20(N - 1) + 4$ bytes every 30 seconds.

(c) Estimate the bandwidth consumption of a RIP network.

The bandwidth consumption is at most $[ND [20(N - 1) + 4]] / 30$ bytes per second.

8.46. RIP runs over UDP, OSPF runs over IP, and BGP runs over TCP. Compare the merits of operating a routing protocol over TCP, UDP, IP.

Solution:

RIP is a protocol in which the routers operate in a highly distributed fashion following a distance vector algorithm. Message exchanges occur only between neighbors and at periodic intervals or triggered by specific events. UDP is suitable for the exchange of individual messages but without delivery guarantees. However the operation of RIP makes allowances for the lack of such guarantees.

OSPF relies on the use of a reliable flooding procedure to distribute link-state information to all the routers. This reliable flooding procedure requires close coordination with the operation of the routers, and hence direct operation over IP instead of over a transport layer protocol is preferred.

BGP peers exchange the entire BGP routing table initially and incremental updates are sent instead of periodic updates to reduce the bandwidth consumption. A small periodic KEEPALIVE message is used to determine that the BGP peers are alive. Reliable delivery of the routing information is required to minimize the bandwidth consumption. For this reason, TCP is chosen to provide the reliable delivery required by BGP.

8.47. Compare RIP and OSPF with respect to convergence time and the number of messages exchanged under several trigger conditions, that is, link failure, node failure, link coming up.

Solution:

Link failure – OSPF has a faster convergence time than RIP. When a link fails, the corresponding OSPF routers send a link-state update to all peer routers. The peer routers then update their databases. This process allows the link failure state information to be propagated to other routers

quickly. In contrast, when a link fails in RIP, the corresponding RIP router updates its own distance vector and sends the link update message to its neighbor. The neighboring nodes update their own routing tables, calculate the new distance vector values and send the updated distance vector to their neighbors. The processing and distribution overhead in RIP slows down the convergence time and routing loops may be created while the algorithm is converging. The faster convergence time of OSPF is at the cost of flooding the network with update messages. RIP is based on the exchange of messages between adjacent nodes only.

Node failure – OSPF converges faster than RIP. OSPF sends a HELLO packet every 10 seconds, compared to RIP's update message every 30 seconds. In case of node failure, OSPF can detect failure of a node within the range of 10 seconds. On the other hand, RIP requires a period of 180 seconds (worst case) to detect a node failure. The grace period of 180 seconds in RIP is due to the fact that RIP is running over UDP which cannot be relied upon to deliver messages consistently. Again, OSPF requires more message exchanges (HELLO packet every 10 seconds) than RIP (update message every 30 seconds).

Link coming up – OSPF converges faster than RIP. When a link is coming up, the OSPF routers attached to this link start sending HELLO packets. Next, these router pairs exchange link-state database description packets and send link-state request packets for those LSA headers that are not in their respective link-state databases. After the databases are updated and synchronized, these OSPF routers send the updated database description packet to their own neighbors and execute the routing algorithms to find out the shortest path. In contrast, a RIP router must first perform a distance vector routing calculation. Next, the router sends the update distance vector to its own neighbor. The time it takes for RIP to distribute the new link information is slower than OSPF due to the processing overhead imposed on each RIP router.

8.48. Consider the OSPF protocol.

Solutions follow questions:

(a) Explain how OSPF operates in an autonomous system that has not defined areas.

When OSPF operates in an AS that has no defined area, the broadcast packets (link-state update and HELLO) must flow all over the AS. If the network (AS) is too large, this approach consumes too many network resources and it does not scale well. Also, each OSPF router link-state database and routing table size increase dramatically when the number of nodes within the AS increases.

(b) Explain how the notion of area reduces the amount of routing traffic exchanged.

When area is used in OSPF, the number of routers within an area decreases as compared to the previous case. The broadcast packets only need to flow within an area. Therefore, the network traffic is reduced and utilization of the network resources is increased. Also, each router has a smaller link-state database and routing table.

(c) Is the notion of area related to subnetting? Explain. What happens if all addresses in an area have the same prefix?

The idea behind subnetting is to add another hierarchical level within a particular class of IP address. The idea behind an area is similar in that it involves the use of hierarchy to simplify routing, but there is no direct relationship between area and subnetting. In an area, multiple classes of IP addresses can exist. Within each class of IP addresses, different ways of subnetting can be done. Therefore, subnetting is just another level of hierarchical level to allow a network administrator to better manage a particular class of IP address.

If all addresses in an area have the same prefix, the area border router (ABR) will advertise or exchange a simple summary to indicate all the addresses with this prefix belong to this particular

area. The ABR contains the whole network topology and performs appropriate routing within its area upon the transmission between different areas.

8.49. Assume that there are N routers in the network and that every router has m neighbors.

Solutions follow questions:

(a) Estimate the amount of memory required to store the information used by the distance-vector routing.

Each node needs the distance to each neighbor and the distance from each neighbor to all destinations which is $m(N - 1)$ entries. Assuming E bytes for each entry the amount of memory is $m(N - 1)E$.

(b) Estimate the amount of memory required to store the information by the link-state algorithm.

Each node needs the information for all links across the network. Each node is connected to m links and there are N nodes in the network. Therefore the total number of links in the network is $(1/2)Nm$ entries. Assuming E bytes per entry the amount of memory is $(1/2)mNE$.

8.50. Suppose a network uses distance-vector routing. What happens if the router sends a distance vector with all 0s?

Solution:

A distance vector with all zeros means that the node has distance 0 to all other nodes. This will prompt all neighbors to route all their packets through the given router. Eventually all packets in the network will be routed to this router, resulting in what can be characterized as a “black hole.”

8.51. Suppose a network uses link-state routing. Explain what happens if:

Solutions follow questions:

(a) The router fails to claim a link that is attached to it.

The link will be eventually omitted from the routing tables and as a result will not be utilized. This will result in some routing decisions that are not the shortest (optimal) paths. Loss of connectivity is also possible. Moreover, a routing loop may occur due to confusion among routers as to which links they are attached to.

(b) The router claims to have a link that does not exist.

If a router claims to have a link that does not exist, it reports false topology information to other routers within the network. If this non-existing link belongs to a shortest path, all the packets that are sent via this path will be lost. This will affect the network performance severely.

8.52. Consider a broadcast network that has n OSPF routers.

Solutions follow questions:

(a) Estimate the number of database exchanges required to synchronize routing databases.

In the worst case, for an n -router OSPF network, n^2 database exchanges are required to synchronize all of the routers databases.

(b) What is the number of database exchanges after a designated router is introduced into the network?

If a designated router is introduced, the number of database exchanges will reduce to n . This is because each router must only communicate with the designated router to obtain link-state information.

(c) Why is the backup designated router introduced? What is the resulting number of database exchanges?

If the primary designated router fails, the information exchange will function properly when the backup designated router is introduced. This mechanism is designed to protect against failure and to provide fast recovery. The number of database exchanges is $2n$ since every router must update both the primary and backup designated router with link-state information.

8.53. Suppose n OSPF routers are connected to a non-broadcast multi-access network, for example, ATM.

Solutions follow questions:

(a) How many virtual circuits are required to provide the required full connectivity?

For a full connectivity, ATM requires an $n(n - 1)$ virtual circuit connections.

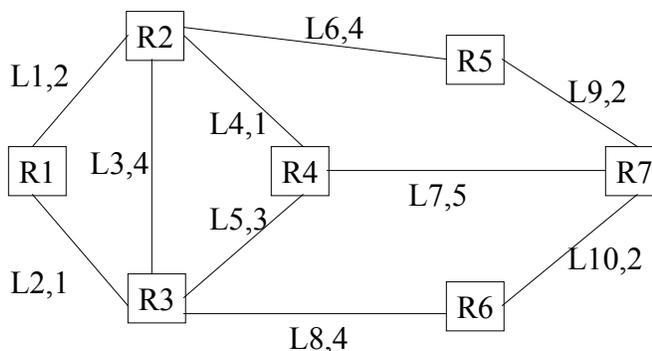
(b) Does OSPF function correctly if a virtual circuit fails?

Yes, OSPF will function correctly if one of the VC fails. This is because the failed VC becomes a failed link in the IP layer and the packet can re-route via a different link (VC) to reach the destination.

(c) Is the number of required virtual circuits reduced if point-to-multipoint virtual circuits are available?

If point-to-multipoint VCs are used, the number of VCs required will be n .

8.54. The figure below shows seven routers connected with links that have the indicated costs. Use the Hello protocol to show how the routers develop the same topology database for the network.



Solution:

In the first stage, every router sends a HELLO packet to all the links that it is attached to.

- R1 sends HELLO packets to R2 and R3 via L1 and L2 respectively.
- R2 sends HELLO packets to R1, R3, R4 and R5 via L1, L3, L4 and L6 respectively.
- R3 sends HELLO packets to R1, R2, R4 and R6 via L2, L3, L5 and L8 respectively.
- R4 sends HELLO packets to R2, R3 and R7 via L4, L5 and L7 respectively.

- R5 sends HELLO packets to R2 and R7 via L6 and L9 respectively.
- R6 sends HELLO packets to R3 and R7 via L8 and L10 respectively.
- R7 sends HELLO packets to R4, R5 and R6 via L7, L9 and L10 respectively.

In the next stage, when a router receives a HELLO packet, it replies with a HELLO packet containing the router ID of each neighbor it has currently seen. For example, when R1 receives HELLO from R2, it will send a HELLO packet to R2. The HELLO packet will inform R2 that the neighbors R1 currently sees are R2 and R3, given that R3's HELLO packet has been received by R1 before R2's HELLO packet.

Once all the routers know their neighbor routers, routers exchange Database Description packets to check if their link-state databases are in agreement. If the databases are not in agreement, the routers exchange link-state request and update packets to synchronize their databases.

8.55. Consider the exchange of Hello messages in OSPF.

Solutions follow questions:

- (a) Estimate the number of Hello messages exchanged per unit time.

A HELLO packet is sent periodically every 10 seconds. Assume the number of links within a network is L . The total number of HELLO packet exchanged per 10-second unit time is $2L$.

- (b) Estimate the size of the Hello messages.

The size of the HELLO packet is 20 bytes for the header plus an entry for each neighbor the router sees. Therefore the size depends on the number of neighbors a router is connected to. Assume the average degree for a router is m . Each router ID is 4 bytes. The size of a HELLO packet = 20 + $4m$ bytes.

- (c) Estimate the bandwidth consumed by Hello messages.

The bandwidth consumed by the HELLO packets is $2L [(20 + 4m) \times 8] / 10$ bps.

8.56. Consider the notion of adjacency in OSPF.

Solution:

- (a) Explain why it is essential that all adjacent routers be synchronized.

Adjacent routers must be synchronized to ensure that routers use the same topology when running the routing algorithm, and ultimately to avoid loops and unnecessary packet dropping in the network.

- (b) Explain why it is sufficient that all adjacent routers be synchronized, that is, it is not necessary that all pairs of routers be synchronized.

For every adjacent router pair, the routers that are connected to this particular adjacent router pair must also be synchronized (according to the rule of adjacency). As a result, all the routers within the network will eventually synchronize.

8.57. Consider the robustness of OSPF.

Solutions follow questions:

(a) Explain how the LSA checksum provides robustness in the OSPF protocol.

The LSA checksum provides error detection for the entire content of the LSA except the link-state age. It gives a second level of error detection to ensure each individual LSA entry carries correct information. If an error is detected in an LSA entry, the router can discard the specific entry while continuing to use other entries in the LSA update packet.

(b) An OSPF router increments the LS Age each time it inserts the LSA into a link-state update packet. Explain how this protects against an LSA that is caught in a loop.

The router increments the LS age each time it inserts the LSA into a link-state update packet. Therefore, whenever a router receives an LSA that is caught in a loop, the router can verify the validity of the LSA with its current LS age. If the LSA is too old, the router will ignore the LSA packet.

(c) OSPF defines a minimum LS update interval of 5 seconds. Explain why.

OSPF uses flooding to distribute the LSA packets. If the LS update interval is too short, the router may not be able to distinguish the order of the LSA packets it receives. A too-short update interval may consume a lot of processing power just to handle the link-state database and routing algorithm. Therefore, choosing the interval of 5 seconds will allow router to exchange LSA twice with every HELLO packet interval. This would allow the network to respond faster in case of any link-state change while keeping the processing overhead low.

8.58. Assume that for OSPF updates occur every 30 minutes, an update packet can carry three LSAs, and each LSA is 36 bytes long. Estimate the bandwidth used in advertising one LSA.

Solution:

An OSPF update occurs every 30 minutes = 1800 seconds. An update can carry three LSAs. Each LSA is 36 bytes long. The OSPF common header is 24 bytes. The size of the LSA field in link-state update consumes 4 bytes. Therefore, if OSPF updates occur every 1800 seconds containing 3 LSAs, the bandwidth for advertising one LSA is equal to $\frac{[(24 + 4 + 3 \times 36) \times 8] / 1800}{3} = 0.201$ bits/second.

8.59. Identify elements where OSPF and BGP are similar and elements where they differ. Explain the reasons for similarity and difference.

Solution:

Similarities - OSPF and BGP are used to exchange routing information including active routes, inactive routes and error conditions within the network in general. OSPF handles the information exchange within an AS while the BGP handles the information exchange between different ASs. Both OSPF and BGP allow a router to construct the network topology based on the link-state and path vector information respectively. OSPF and BGP use HELLO and KEEPALIVE messages respectively to determine the presence of peers.

Differences - OSPF runs over IP while BGP runs over a TCP connection. OSPF is a link-state protocol and BGP uses a path-vector protocol. BGP can enforce policy by affecting the selection of different paths to a destination and by controlling the redistribution of routing information. BGP only requires an incremental update of the database. However, OSPF requires a periodic update (refresh every 30 minutes).

8.60. Discuss the OSPF alternate routing capability for the following cases:

Solutions follow questions:

(a) Traffic engineering, that is, the control of traffic flows in the network.

In OSPF, the link-state metric exchange between routers can represent the volume of flow of traffic for each individual link. Therefore, each router can use these link-state metrics (traffic flow) to determine the best route for a particular IP packet given its type-of-service (TOS) field. However, the processing required in each router increases dramatically. One possible solution is to make use of explicit routing by pre-establishing a path that is suitable for each class of service.

(b) QoS routing, that is, the identification of paths that meet certain QoS requirements.

OSPF can use a QoS parameter or metric for each link to determine the optimal path for QoS routing. All routers within the network can use a set of routing algorithms to send packets along a particular route that satisfies a certain QoS requirement, e.g. delay, bandwidth, or bit-error rate requirement.

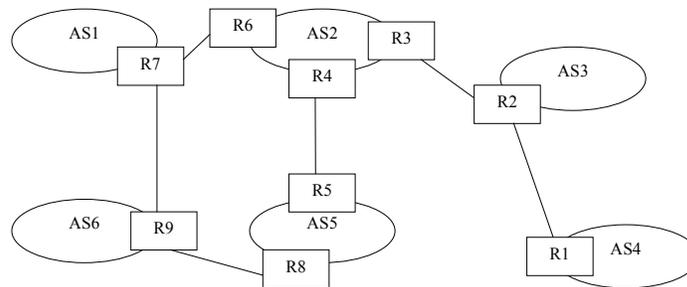
(c) Cost-sensitive routing, that is, the identification of paths that meet certain price constraints.

In cost-sensitive routing, the link-state metric can represent the cost of each link. The router determines the best route that meets a certain price constraint for a particular packet. An ingress router needs to have the information to determine the minimum cost paths to all destinations.

(d) Differential security routing, the identification of paths that provide different levels of security.

In differential security routing, the link-state metric involves identification of the security level of each link. All routers need to come up with a standard definition of a security metric for links. The router identifies the best route to satisfy each individual packet's security requirement.

8.61. Consider the autonomous systems and BGP routers in the following figure.



Solutions follow questions:

(a) Suppose that a certain network prefix belongs to AS4. Over which router pairs will the route to the given network be advertised?

If a route in AS4 requires advertising, R1 will run an eBGP protocol and establish a TCP connection to R2 to exchange the routing information. R2 will relay the message to R3 through a TCP connection based on eBGP. R3 forwards this information to R4 and R6 via iBGP. R6 advertises the information to R7 and R4 advertises to R5 via eBGP. R7 advertises the information to R9 via eBGP. Finally, R5 sends the information to R8 via iBGP and R9 sends the information also to R8 via eBGP.

- (b) Now suppose the link between R1 and R2 fails. Explain how a loop among AS1, AS2, AS6, and AS5 is avoided.

With the use of path-vector routing, a loop occurs whenever a BGP router receives an update message with an AS path attribute that contains its own AS number. The BGP router ignores the route and discards the information immediately.

- (c) Suppose that R9 is configured to prefer AS1 as transit and R6 is configured to prefer AS1 as transit. Explain how BGP handles this situation.

The use of path-vector routing will again allow routers R9 and R6 to avoid the potential loop that can result from these preferences.

8.62. Why does BGP not exchange routing information periodically like RIP?

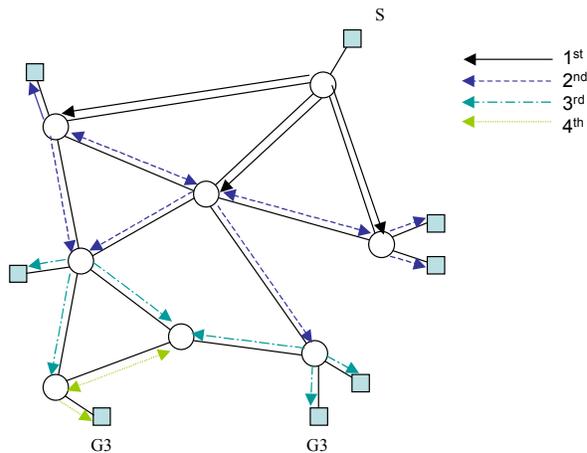
Solution:

BGP runs over TCP, which provides reliable service and simplifies BGP significantly by allowing the protocol to assume the availability of reliable information. In contrast, RIP runs over UDP, which is an unreliable protocol that may experience packet loss. RIP therefore requires periodic exchange of routing information to ensure routing information is correct and up-to-date.

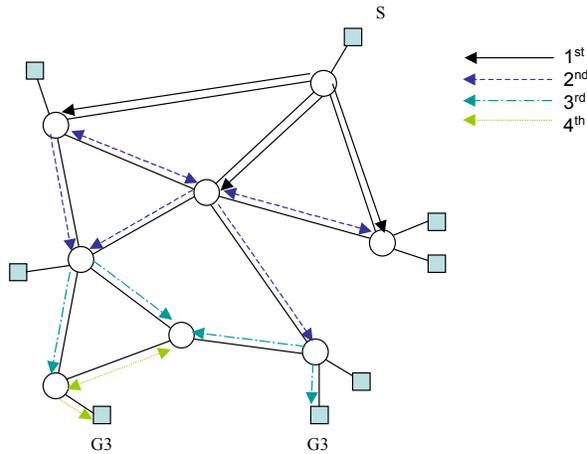
8.63. Consider the network shown in Figure 8.60. Suppose that a source connected to router 7 wishes to send information to multicast group G3.

Solutions follow questions:

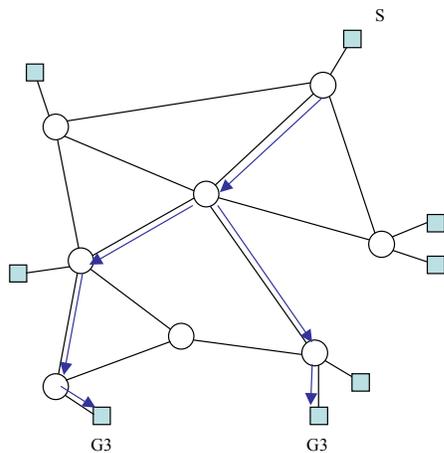
- (a) Find the set of paths that are obtained from reverse-path broadcasting.



(b) Repeat for truncated reverse-path broadcasting.



(c) Repeat for reverse-path multicasting.



8.64. Discuss the operation of the reverse-path multicasting in the following two cases:

Solutions follow questions:

(a) The membership in the multicast group in the network is dense.

If the membership in the multicast group in the network is dense, then there will not be extensive use of prune messages and operation will be similar to reverse-path broadcasting. The utilization of bandwidth for multicasting packets will be relatively efficient.

(b) The membership in the multicast group in the network is sparse.

If the membership in the multicast group in the network is sparse, there will be extensive pruning of paths. The multicast tree will be “thin” relative to the overall network, and the utilization of bandwidth for multicasting packets will be inefficient.

8.65. Suppose an ISP has 1000 customers and that at any time during the busiest hour of the day, the probability that a particular user requires service is .20. The ISP uses DHCP. Is a class C address enough so that the probability is less than 1 percent that there is no IP address available when a customer places a request?

Solution:

A class C address space provides 254 IP addresses. The distribution of the number of active customers is given by a binomial distribution with parameters $n = 1000$ customers and $p = .20$. Let X be the number of active customers. The probability that $X = k$ customers are active at a given time is given by:

$$P[X = k] = \binom{1000}{k} (0.2)^k (0.8)^{1000-k}$$

The random variable X has mean equal to $m = np = 1000 \cdot 0.2 = 200$ and variance given by $\sigma^2 = np(1-p) = 1000(.2)(.8) = 160$.

For large values of n , the binomial distribution can be approximated by a Gaussian distribution with the same mean and variance. In particular we have that:

$$P[X > k] \approx P[X_{Gauss} > k] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_k^{\infty} e^{-(x-m)^2 / 2\sigma^2} dx = \frac{1}{\sqrt{2\pi}} \int_{\frac{k-m}{\sigma}}^{\infty} e^{-x^2 / 2} dx$$

The value of the above integral at the point $(k - m)/\sigma = (254-200)/40 = 1.35$ is 0.088=8.8%. This is the probability that the ISP does not have enough IP addresses to serve customer connection requests.

8.66. Compare mobile IP with the procedures used by cellular telephone networks (Chapter 4) to handle roaming users.

Solutions follow questions:

(a) Which cellular network components provide the functions of the home and foreign agent?

The home location register in cellular network provides the equivalent functions of the home agent in mobile IP network. The visitor location register in cellular network provides the functions of the foreign agent in mobile IP network.

(b) Is the handling of mobility affected by whether the transfer service is connectionless or connection-oriented?

The basic handling of the mobile users is the same in terms of the operation of the home location and visitor location registers. The manner in which information is transferred is of course different given that one is connectionless and the other connection-oriented.

8.67. Consider a user that can be in several places (home networks) at different times. Suppose that the home networks of a user contain registration servers where users send updates of their location at a given time.

Solutions follows questions:

- (a) Explain how a client process in a given end system can find out the location of a given user in order to establish a connection, for example, Internet telephone, at a given point in time.

A client process can locate a given user by sending a request to a designated registration server. This registration server contains the most up-to-date location of a given user. Once the client process finds the location, it can establish a connection based on the information. The way to locate the registration server can be predefined by the network administrator or based on the request IP address prefix. In the latter approach, the client can use the address prefix of the given user to locate the corresponding registration server.

- (b) Suppose that proxy servers are available, whose function is to redirect location requests to another server that has more precise location information about the callee. For example, a university might have such a server, which redirects requests for prof@university.edu to departmental servers. Explain how a location request for engineer@home.com might be redirected to a.prof@ece.university.edu.

When a location request for engineer@home.com is sent to the proxy server, the server will forward the request to the home.com server. The home.com server can send a redirect message informing the proxy server that the engineer@home.com can be redirected to a.prof@ece.university.edu. Therefore, whenever the client sends messages to the engineer@home.com via the proxy server, the server will automatically redirect the message to a.prof@ece.university.edu.

Solutions to Chapter 9

9.1. Suppose that instead of ATM, BISDN had adopted a transfer mode that would provide constant-bit-rate connections with bit rates given by integer multiples of 64 kbps connections. Comment on the multiplexing and switching procedures that would be required to provide this transfer mode. Can you give some reasons why BISDN did not adopt this transfer mode?

Solution:

In terms of multiplexing, flows could be aggregated into integer multiples of 64 kbps. This type of multiplexing is easy to implement and is widely deployed in telephone networks. The difficulty is with switching, where new designs are required to provide connections that are multiple integers of 64 kbps. In these switches various inputs and outputs are connected by pipes that are integer multiples of 64 kbps. Time-slot interchange techniques do not lend themselves readily to this type of transfer.

In addition, this proposed transfer mode lacks flexibility to adapt to variable bit rate traffic and traffic rates that don't easily fit into an integer multiple of 64 kbps. ATM was designed to be able to efficiently handle all types of traffic, whether fixed- or variable-rate. In addition, a multirate transfer mode would require over-provisioning to handle bursty and VBR traffic.

9.2a. Compare the bandwidth management capabilities provided by ATM virtual paths to the capabilities provided by SONET networks.

Solution:

SONET uses the SONET multiplexing hierarchy, virtual tributaries and add-drop multiplexers to route information flows that have bit rates that are integer multiples (and submultiples) of OC-1.

ATM provides bandwidth management using virtual connections that are grouped together to form virtual paths. The multiplexing mechanism is much more flexible than that of SONET. Each VP can contain anywhere from 1 to 2^{16} VCs, and up to 2^8 VPs can occupy the same link.

Further, since VP numbers and VC numbers have only local significance, each link can take full advantage of the VC and VP ID space.

9.2b. Can ATM virtual paths be adapted to provide the fault tolerance capabilities of SONET rings? Explain your answer.

Solution:

Yes, ATM VPs can be used to provide fault tolerance. For example, one method of providing SONET-like fault tolerance is to create a logical ring topology using VPs. Backup VPs can be allocated to produce self-healing rings similar to those of SONET.

The 50% bandwidth efficiency of the SONET self-healing ring topology is unattractive. Another method of providing good fault tolerance is to take advantage of the ATM control plane and provide fault protection in a mesh topology network.

(Ring vs mesh approaches to providing fault tolerance have received much attention in the context of emerging optical networks.)

9.3. In Chapter 6 we saw that the performance of MACs for LANs can depend strongly on the delay-bandwidth product of the LAN. Consider the use of ATM in a LAN. Does the performance depend on the delay-bandwidth product? If yes, explain how; if no explain why not? Does the same conclusion apply to any LAN that involves the switching of packets, rather than the use of a broadcast medium?

Solution:

In an ATM network we do not have a shared medium. Instead, we have a network controlled by a switch, or a group of switches (ATM switches). When communication between two nodes is needed, then a connection is established between them and the data flows in a bi-directional virtual channel. ATM transmits one cell after each other without waiting for an ACK before continuing. Since there is no contention for the channel, the channel can be “always” in use. Delay-bandwidth product does not directly affect performance here.

However, there are other factors that can affect the performance of an ATM LAN, such as cell loss or congestion control. These problems are minimized to the extent that ATM connections reserve sufficient buffer and bandwidth to accommodate the actual traffic flow. However when ATM attempts to handle unscheduled traffic flows, as in ABR service, then congestion control becomes a necessity. In this case, delay-bandwidth effects need to be considered in the design of feedback congestion control mechanisms.

The general observations for ATM apply to virtual-circuit oriented packet-switched LANs. However, the variable-length nature of packets reduces the fine-ness of the control that can be exercised on buffer and bandwidth allocation.

9.4. Does the performance of ATM depend strongly on the delay-bandwidth product of a wide-area network? Specifically, consider the performance of ATM congestion control in a network with large delay-bandwidth product. Does the size of the buffers in the ATM switches influence the severity of the problem? Give an order of magnitude for the amount of buffering required in the switches.

Solution:

ATM uses an admission control mechanism to accept connections. This prevents congestion in the network due to certain types of services (for example, CBR and UBR [Section 7.8.1]). For ABR a closed-loop control mechanism is used for congestion control [Section 7.8.2]. In this case explicit feedback is used through special control cells called Resource Management (RM) cells.

If the delay-bandwidth product is large, the information about the congestion state of the network will arrive after a long delay to the destination, that is, many cells will already been sent to the congested switch before the RM cells arrive to the source. Similarly there will be a waste of bandwidth when the congestion is reduced.

Appropriate sizing of the buffer in the ATM switch can mitigate the problem produced by the delay of the RM cells. If cell loss is to be avoided, sufficient buffering must be available to buffer the cells that are transmitted from the onset of congestion to time when the control signal take effect. This buffering is proportional to a few times the delay-bandwidth product, measure in cells.

9.5. Compare a conventional TDM leased line with an ATM PVC from the user’s point of view and from the network operator’s point of view. What features of PVCs make them attractive from both points of view?

Solution:

TDM connections waste bandwidth when the information flow is not continuous. They are always on, regardless of whether or not the user is generating traffic. Unlike TDM connections PVC connections only use bandwidth when they are carrying traffic. Thus, PVCs are attractive to service providers because, when the PVC is idle, the resulting free bandwidth can be used to carry other traffic. It also allows the service provider to charge the user less during such times, which is attractive to the user. Thus, PVCs are attractive to both users and service providers because they are more flexible and, therefore, more efficient. This efficiency translates into savings that benefit both parties.

9.6. An inverse multiplexer is a device that takes as input a high-speed digital stream and divides it into several lower-speed streams that are then transmitted over parallel transmission lines to the same destination. Suppose that a very high speed ATM stream is to be sent over an inverse multiplexer. Explain what requirements must be met so that the inverse demultiplexer produces the original ATM stream.

Solution:

To reproduce the original stream at an inverse demultiplexer, all streams must arrive at nearly the same time for re-assembly. Although a certain amount of delay variation and jitter can be tolerated by buffering the incoming streams before reassembly, the average data rates of all the individual streams must be identical.

In terms of QoS, the paths taken by the individual streams must all conform to the service class(es) of the original stream. For example, if the original stream carried voice traffic, the above-mentioned jitter must be kept low in all individual streams.

9.7. Suppose an ATM switch has 32 input ports and 32 output ports.

Solutions follow questions:

- (a) Theoretically, how many connections can the switch support?

Virtual Path (VPI):

UNI	NNI
8 bits	12 bits
$2^8 = 256$	$2^{12} = 4096$

Virtual Channel (VCI):

16 bits
$2^{16} = 65,536$ per VP

	Size of VPI (bits)	Size of VCI	Max VP	Max VC per VP	Max VCC Per LINK	32 LINKS
UNI	8	16	256	65,536	16,777,216	5.37×10^8
NNI	12	16	4,096	65,536	2.68×10^8	8.59×10^9

- (b) What are the table lookup requirements for supporting a large number of connections? Do they limit the practical size on the number of connections that can actually be supported?

The switch does not have to maintain an entry in a table for each virtual channel connection. It has to maintain information for every virtual path and information of the virtual channels contained in the virtual paths that end at the switch. In the case of an NNI the maximum size of the virtual path lookup table is 4,096, and there must be another table containing the information for the connections contained in the VPs that terminate. This table can be very large: 65,536 in the worst case for every VP.

- (c) (Optional) Do a web search on content addressable memories (CAMs) and explain how these can help in addressing the table lookup problem.

Do a web search for Sibercore, a company that produces CAMs for this type of application.

In table lookup, in order to be able to use the incoming VC number as an address to access its corresponding information in the table, the table needs to have one entry for every possible VC

number covering entire VC space. In a CAM this is not the case. Only VC numbers that are in use are stored in the table, reducing the size of the table to maximum possible simultaneous VCs. In a CAM a VC number that is in use is stored in the CAM as one of the portions of the stored data for that VC. Once the VC is looked up the CAM locates the entry in the table whose VC number portion matches the number that is being looked up. The matching is done in hardware and is very fast.

9.8. Explain how the header error checksum can be used to synchronize to the boundary of a sequence of contiguous ATM cells. What is the probability that an arbitrary five-octet block will satisfy the header error checksum? Assume bits are equally likely to be 0 or 1. What is the probability that two random five-octet blocks that correspond to two consecutive headers pass the error check?

Solution:

The stream of data is passed through a sliding window of 40 bits. The window slides one bit at a time. Every time the content of the window is examined for error checksum. If it passes it is assumed that the content of the window at that instant is the header of a new cell, thus synchronizing the boundary of the corresponding ATM cell. Once the synchronization is achieved the checksum is applied only to headers of the cells and the sliding window detection is used only if the synchronization is lost for any reason. This mechanism is based on the assumption that the probability of any arbitrary 40 bit other than the header to pass the checksum is low, as well as the probability that the header itself fail the checksum in normal circumstances. In the next part of the solution we will find the probability that a false synchronization happen and the probability that it persist for two consecutive cells.

Referring to Section 3.8.6, if the received polynomial is $R(x) = b(x) + e(x)$, where $b(x)$ is the transmitted polynomial and $e(x)$ is the error polynomial, if $e(x)$ is divisible by the generator polynomial $g(x) = x^8 + x^2 + x + 1$, then the error pattern is undetectable.

Thus, for an undetected error we require:

$$e(x) / g(x) = 0 \rightarrow g(x)m(x) = e(x)$$

where $m(x)$ is any polynomial. If we find the number of solutions for $m(x)$, we will have found the number of received bit patterns that are not detected as errors. Because $g(x)$ is 9 digits long and $e(x)$ is 40 digits long, the above equation is satisfied for:

$$m(x) \leq 2^{40-9}$$

which implies that there are 2^{31} such $e(x)$ s. Assuming that for each bit 1 and 0 are equally probable, each of the 2^{40} possible bit patterns is equally probable. Thus:

$$P[\text{undetectable error}] = 2^{31} / 2^{40} = 2^{-9} = 1.95 \times 10^{-3}$$

Thus the probability that two such random five-octet blocks pass the error check:

$$\begin{aligned} P[2 \text{ blocks pass error check}] &= P[\text{undetectable error}]^2 \\ &= 3.81 \times 10^{-6} \end{aligned}$$

9.9. We saw that ATM provides a GFC field in the UNI ATM header. Suppose that several terminals share a medium to access an ATM network.

Solutions follow questions:

- (a) Explain why flow control may be required to regulate the access of traffic from the terminals into the network. Explain how the GFC field can be used to do this task?

In order to avoid collisions and to ensure fair access for every terminal, each terminal's traffic flow must be regulated. Because the GFC field is not currently used, its bits could be used to provide indexing for some form of polling or controlled access to the medium.

- (b) Explain how the GFC field can be used as a subaddress to provide point-to-multipoint access to an ATM network. Does this conflict with the flow control requirement?

The GFC field could hold the subaddress of the particular subgroup in the ATM network to which the traffic is destined.

Since the GFC field is only 4 bits long, it is unlikely it can be effectively partitioned for both flow control and multicast subaddressing. However, assuming that all point-to-multipoint traffic flows downstream from the server to the terminals of the access network, the GFC field could be used for flow control in the upstream direction from the terminals to the ATM server and for multicast subaddressing (in the downstream direction) from the server to the terminals.

9.10. The purpose of the header error control (HEC) field is to provide protection against errors in the header that may result in misaddressing and misdelivery of cells. The CRC used in the ATM header can correct all single errors and can detect (but not correct) all double errors that occur in the header. Some, but not all, multiple errors in excess of two can also be detected.

Solutions follow questions:

- (a) Suppose that bit errors occur at random and that the bit error rate is p . Find the probability that the header contains no errors, a single error, a double error; more that two errors. Evaluate these probabilities for $p=10^{-3}$, 10^{-6} , and 10^{-9} .

We assume:

- Bit error rate = p
- L = Header Length = 5 bytes = 40 bits
- X = Number of errors in Header

$$P[X = 0] = (1 - p)^{40}$$

$$P[X = 1] = \binom{40}{1} (1 - p)^{39} p = 40(1 - p)^{39} p$$

$$P[X = 2] = \binom{40}{2} (1 - p)^{38} p^2 = 780(1 - p)^{38} p^2$$

$$P[X > 2] = 1 - (P[X = 0] + P[X = 1] + P[X = 2])$$

p	1.00E-03	1.00E-06	1.00E-09
$P[X=0]$	0.960770211	0.999960001	0.999999996
$P[X=1]$	3.85E-02	4.00E-05	4.00E-08
$P[X=2]$	7.51E-04	7.80E-10	7.80E-16
$P[X>2]$	9.61E-06	1.11E-14	0

(b) Relate the calculations in part (a) to the CMR experienced in such an ATM network.

CMR - Cell Misinsertion Rate (Section 9.3.3)

The probability for a cell misinsertion is bounded by $P[X > 2]$ because the CRC in the Header can detect double and single errors and some multiple errors. If the error is detected then the cell is discarded, but if the error is not detected then the cell can be delivered to an erroneous destination.

(c) In practice the bit errors may not always be random, and so to protect against bursts of errors the following adaptive procedure may be used.

Normally the receiver is in the “correction” mode. If a header is error free, the receiver stays in this mode; if the receiver detects a single error, the receiver corrects the error and changes to the “detection” mode; if the receiver detects a multiple error, the receiver changes to the HEC error detection mode. If the receiver is in “detection” mode and one or more errors are detected in the header, then the cell is discarded and the receiver stays in the detection mode; otherwise, the receiver changes to correction mode.

Explain why this procedure protects against bursts of errors. If the bit errors are independent, what is the probability that two consecutive headers contain single errors and hence that a cell is discarded unnecessarily?

When an isolated error occurs, the receiver detects the single error, and the next cell is very likely to be error-free. Thus the receiver will go into detection mode and then back to correction mode.

When a burst of errors occurs, several bits are changed hence one or more codewords can be affected. If a single error is detected, then the cell goes into detection mode, and then there is a fairly high probability that errors will be detected in the next cell. Similarly, if multiple errors are detected, then the receiver must go into detection mode as the single-error case has already been eliminated.

If an error was detected in cell n and another error is detected in cell $n + 1$ then it is possible that a burst of errors has occurred. This is why the cell is discarded even when there is only a single error.

$$P[1 \text{ error in two consecutive Headers}] = P[X = 1] * P[X = 1] = [40 (1 - \rho)^{39} \rho]^2$$

	1.00E-03	1.00E-06	1.00E-09
$P[\text{error in 2 consecutive cells}]$	1.48E-03	1.6E-09	1.6E-15

9.11. What is the difference between CER and CLR? Why is one negotiated during connection setup and the other is not?

Solution:

CER stands for Cell Error Rate. It is dependent on the intrinsic physical characteristics of the network, such as channel noise. As a result, CER is not under control of the network operator.

CLR stands for Cell Loss Rate. It is dependent on topology, routing and traffic distribution within the network and occurs when packets are dropped during periods of congestion due to buffer overflow or the predicted inability to accommodate their QoS requirements. In the latter case, the buffers may not be full, but the switch chooses to discard the packet because, for example, it has incurred too much delay to be useful to the receiver. CLR can be controlled with the granularity as small as one VC in ATM by ensuring that the VC does not travel over a congested path and by assigning higher priority to its packets at each node in its path.

9.12. Why does the calculation of CER and CMR exclude blocks that are counted in the SECBR?

Solution:

The block of cells that are counted in the SECBR are in error due to relatively infrequent, unpredictable and atypical situations, such as re-routing during a link failure. These errors tend to occur in relatively large blocks and are not indicative of the CER during normal operation. For this reason, they are not included in the calculation of CER and CMR statistics.

9.13. Explain how weighted fair queuing scheduling can be used to affect the CLR and CTD experienced by cells in an ATM connection.

Solution:

If a VC has very strict CLR and CTD service requirements, its queue may be given priority in the scheduler. In weighted fair queuing this corresponds to giving the queue a large weight. Since queues are serviced at a rate proportional to their weight, during congestion, this VC's queue will be serviced more quickly than others. As a result, the queue length will never be very long, ensuring low CTD. Its queue is also unlikely to overflow, ensuring low CLR.

9.14. Explain the effect of a single cell loss from a long packet in a situation that uses an end-to-end ARQ retransmission protocol. Can you think of strategy to deal with such losses? Can you think of a way to ameliorate the impact of packet retransmission?

Solution:

If a single cell is lost, then the entire packet must be retransmitted. Thus each cell that is transmitted after the given cell loss will be wasting transmission bandwidth. For this reason, many ATM switches implement procedure that discards all cells that belong to the same packet (as determined from AAL5) after an erroneous cell has been detected.

9.15. Consider a sequence of ATM cells carrying PCM voice from a single speaker.

Solutions follow questions:

- (a) What are the appropriate traffic descriptors for this sequence of cells and what is an appropriate leaky bucket for policing this stream?

Since the PCM voice traffic arrives at a constant rate, only the PCR needs to be specified. In the case of voice traffic, the PCR = 64 kbps. To police this stream, one could use a leaky bucket algorithm with a depth of 1 cell and a drain rate equal to the PCR.

- (b) Suppose that an ATM connection is to carry the cell streams for M speakers. What are appropriate traffic descriptors for the resulting aggregate stream, and how can it be policed?

The average rate of the combined cell streams is $M \cdot 64$ kbps. However, the cell stream may or may not be constant rate depending on the relative times at which the streams produce their cells. If we have no control over how the cells are produced, then the PCR = M (64kbps). A simple leaky bucket algorithm could be used to police the stream with a bucket depth equal to M bytes (in case all of the packets arrive simultaneously) and a drain rate of M (64Kbps).

9.16. Consider a sequence of ATM cells carrying PCM voice from a single speaker, but suppose that silence suppression is used.

Solutions follow questions:

- (a) What are the appropriate traffic descriptors for this sequence of cells and what is an appropriate leaky bucket(s) arrangement for policing this stream? What situation leads to nonconforming cells?

PCM requires 64 kbps. Let's assume that in this case silence detection reduces the needed bandwidth by half, that is, up to 32 kbps.

$$\begin{aligned} 64 \text{ kbps} / [(48 \text{ bytes/cell}) (8 \text{ bits/byte})] &= 166.6 \text{ cells/s} \\ 32 \text{ kbps} / [(48 \text{ bytes/cell}) (8 \text{ bits/byte})] &= 83.3 \text{ cells/s} \end{aligned}$$

$$\begin{aligned} \text{Peak Cell Rate (PCR)} &= 166.6 \\ \text{Sustainable Cell Rate (SCR)} &= 83.3 \\ I &= 1 / 83.3 \end{aligned}$$

(See Section 7.8.1)

Let's assume that a person can speak continuously without interruption for at most M seconds. Then:

$$\text{MBS} = 166.6 M$$

$$L = (\text{MBS} - 1) \left(\frac{1}{83.3} - \frac{1}{166.6} \right) = \frac{(166.6M - 1)}{166.6} = M - \frac{1}{166.6}$$

Dual leaky buckets (Figure 7.56 Section 7.8.1) are appropriate for policing this stream because PCR and SCR need to be checked.

If the speaker talks continuously for a long period of time without pauses then there will not be silence periods and the connection contract can be violated, in other words MBS and SCR will be exceeded.

- (b) Suppose that an ATM connection is to carry the cell streams for M speakers. What are appropriate traffic descriptors for the resulting aggregate stream and how can the stream be policed?

In general the speakers are talking to parties at different locations, so all the cell streams of the M speakers are independent of each other. The on-off manner in which the cells are produced will tend to randomize and spread out the times at which cells are produced. The number of cells produced in a time slot is a binomial random variable with parameter M and probability p , where p is about 0.4. The sustained rate is therefore approximately $0.4M \cdot 64$ kbps. The standard deviation is $(0.4(0.5)M)^{1/2}$, so the MBS should be some multiple of this value. The sustained rate will be $M \cdot 32$ kbps.

9.17. Suppose that constant-length packets (of size equal to M cells) arrive at a source to be carried by an ATM connection and that such packets are separated by exponential random times T . What are the appropriate traffic descriptors for this sequence of cells and what is an appropriate leaky bucket(s) arrangement for policing this stream? What situation leads to non-conforming cells?

Solution:

Assume that the arrival process is Poisson. (i.e. the distribution of the inter-arrival times is $P[T \leq t] = 1 - e^{-\lambda t}$). The traffic can be described by:

$$\begin{aligned} \text{SCR} &= \lambda M \text{ cells/sec} \\ \text{MBS} &= M \text{ cells} \\ \text{PCR} &= R \end{aligned}$$

The leaky bucket should have the following parameters:

$$\begin{aligned} \text{Drain Rate} &= \lambda M \text{ cells/sec} \\ \text{Depth of Bucket} &= L \text{ cells} \end{aligned}$$

Each arriving packet fills the bucket with M cells. Because the arrival process is random, there is always a nonzero probability that the bucket will overflow when back-to-back packets arrive with zero or very small interarrival time. L must be chosen large enough so that the probability of this occurring is below some predetermined threshold. By choosing the drain rate equal to λM cells/sec, we ensure that the long-term average rate of the connection cannot be above λM cells/second. The choice of L basically defines how large a variance the arrival process can possess. The maximum burst size that is allowed is determined by L . If more than the expected number of back-to-back packets arrive, the bucket will overflow leading to nonconforming cells.

9.18. Explain why each specific set of traffic descriptors and QoS parameters were selected for each of the ATM service categories.

Solution:

Service Category	Traffic Descriptor(s)	QoS Parameter(s)	Explanation
CBR	PCR	CTD, CDV, CLR	<i>QoS parameters:</i> Provides service to applications that require a virtual leased line. Delay, jitter and loss must be controlled. <i>Traffic Descriptors:</i> The bit rate is constant, so only PCR needs be specified.
rt-VBR	PCR, SCR, MBS	CTD, CLR, CDV	<i>QoS parameters:</i> Provides service to applications that require real-time delivery and low loss. Thus, delay, jitter and loss must be controlled. <i>Traffic Descriptors:</i> Bit rate can be variable. Thus, need PCR, SCR, and MBS specified.
nrt-VBR	PCR, SCR, MBS	CLR	<i>QoS parameters:</i> Provides service to applications that require low loss. Thus, only loss must be controlled (no timing requirements). <i>Traffic Parameters:</i> as described above.
ABR	PCR, MCR	Maybe CLR	<i>QoS parameters:</i> No timing requirements. Low loss is preferable but is dependent on how well the source responds to feedback flow-control. <i>Traffic Parameters:</i> Can specify maxima based on the application, but otherwise rate is completely adaptive so no specified MBS, SCR.
UBR	Maybe PCR	None	<i>QoS parameters:</i> Lowest priority traffic. No guarantees and not even an attempt to accommodate traffic until traffic from all other classes is accommodated. <i>Traffic Parameters:</i> PCR may be limited, but generally source may transmit at whatever rate can be supported in the network.

9.19. Suppose that IP packets use AAL5 prior to transmission over an ATM connection. Explain the transfer-delay properties of the IP packets if the ATM connection is of the following type: CBR, rt-VBR, nrt-VBR, ABR, or UBR.

Solution:

If we use AAL5, the system adds a trailer to the IP packet with some information about it. Then it is segmented into 48-byte blocks and these blocks are inserted as the cell's payload [Section 9.4.4].

IP does not use QoS parameters (currently), but it could make use of the advantages of ATM to satisfy certain requirements; for example, a video transmission can use a CBR service to provide a service with low delay and low-delay variation.

If we use rt-VBR, the transmission will guarantee that the packet will comply with certain parameters of quality especially low delay and low-delay variation. IP Telephony can use this service.

In the case of nrt-VBR the transmission will guarantee an average transmission speed. Transaction processing or file transfer can use this type of service.

Similarly, ABR will have some guarantees in the delivery service, depending on the availability of bandwidth in the network. In the case of UBR the IP packet will be served as if it were on a traditional network (that is, a non-ATM network) depending on the congestion.

9.20. Proponents of ATM argue that VBR connections provide a means of attaining multiplexing gains while providing QoS. Proponents of IP argue that connectionless IP routing can provide much higher multiplexing gains. Can you think of arguments to support each of these claims? Are these claims conflicting or can they both be correct?

Solution:

ATM: VBR can use statistical multiplexing to use bandwidth efficiently. If an application is not CBR, the VBR service may represent no "noticeable" loss of service quality when compared to a dedicated line. Further, since it is a connection-oriented service, it can provide mechanisms to make strong guarantees on delay properties.

IP: In connection-oriented networks, the path of a flow is fixed, so the ability to make multiplexing gains is limited. In IP, the network implicitly multiplexes over many paths (though they are not explicit) to take full advantage of the network bandwidth.

Both of the arguments are correct, but each points to different strengths and weaknesses. IP is more efficient in principle, as described, but it is difficult to provide QoS. As DiffServ (Chapter 10) shows, some QoS mechanisms can be offered by marking packets at the ingress with their QoS requirements and treating them accordingly at each node in the network. However, connection oriented service, such as ATM can make much tighter service guarantees. Further, the explicit paths setup by the ATM network can increase bandwidth efficiency by reducing packet loss and ensuring near-shortest-path routing.

9.21a. Consider a link that carries connections of individual voice calls using PCM. What information is required to perform call admission control on the link?

Solution:

Peak rates of the connections and their delay and loss requirements are required. Each connection uses 64 kbps of bandwidth with strict delay requirements. To perform call admission control, the network needs to determine if available bandwidth (64 kbps) exists in the link and if the desired delay

requirement can be met. The delay depends on the type of scheduling that is used in the switches, as well as on the number of calls already in place.

9.21b. Now suppose that the link carries connections of individual voice calls using PCM but with silence suppression. What information is required to do call admission control?

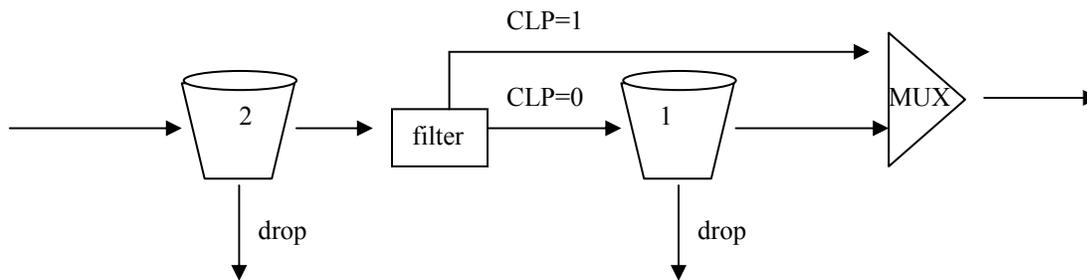
Solution:

With silence suppression, the problem becomes a statistical one. Although, on average, each one-way flow is silent for just over 50% of the call, the duration and spacing of these silent periods is random and variable bit rate traffic (of the on-off variety) results. In this case in addition to the peak rate, sustained rate and maximum burst size should also be known such that the delay and loss probabilities can be guaranteed. A call should be admitted if the chance of a packet in the new connection being dropped or delayed too long due to insufficient resources is below some predetermined threshold. In the case of voice traffic, this threshold must be very low in order to guarantee acceptable service quality.

9.22. Suppose that an ATM traffic stream contains cells of two priorities, that is, high-priority cells with CLP = 0 in the headers and low-priority cells with CLP = 1.

Solutions follow questions:

- (a) Suppose we wish to police the peak cell rate of the CLP = 0 traffic to p_0 as well as the peak rate of the combined CLP = 0 and CLP = 1 traffic to p_{0+1} . Give an arrangement of two leaky buckets to do this policing. Nonconforming cells are dropped.

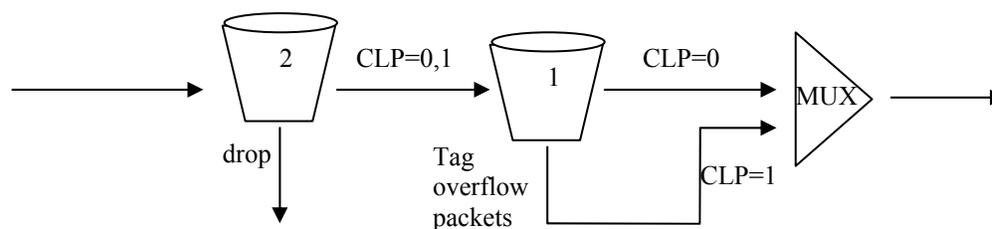


Leaky Bucket 1: drain rate = p_0 , depth = 1
 Leaky Bucket 2: drain rate = p_{0+1} , depth = 1.

- (b) Compare the following policing schemes: (1) police CLP = 0 traffic to peak rate p_0 and police CLP = 1 traffic to peak rate p_1 , (2) police the combined CLP = 0 and CLP = 1 traffic to peak rate $p_0 + p_1$. Which approach is more flexible?

The second approach is more flexible because CLP = 1 traffic can fluctuate to $p_0 + p_1$ as long as the total traffic remains below $p_0 + p_1$. In this way, each traffic type can make use of unused bandwidth when traffic volume of the other type the other is low.

- (c) Repeat part (a) if CLP = 0 cells that do not conform to the p_0 are tagged by changing the CLP bit to 1. Cells that do not conform to p_{0+1} are dropped.

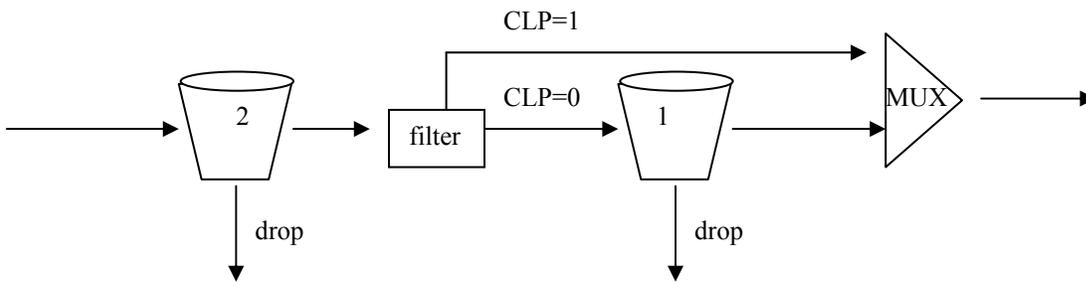


Leaky Bucket 1: drain rate = ρ_0 , depth = 1
 Leaky Bucket 2: drain rate = ρ_{0+1} , depth = 1

9.23. Suppose that an ATM traffic stream contains cells of two priorities, that is, high-priority cells with CLP = 0 in the headers and low-priority cells with CLP = 1.

Solutions follow questions:

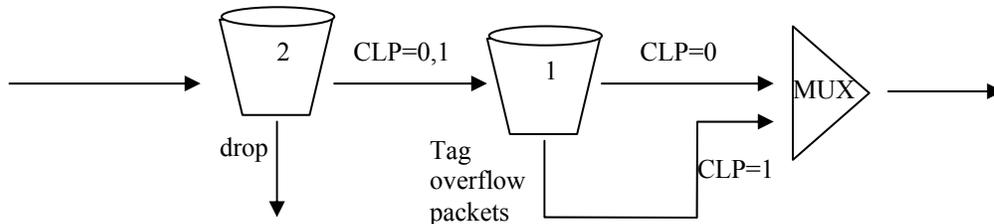
- (a) Suppose we wish to police the sustainable cell rate of the CLP = 0 traffic to SCR_0 and MBS as well as the peak rate of the combined CLP = 0 and CLP = 1 traffic to p_{0+1} . Give an arrangement of two leaky buckets to do this policing. Nonconforming cells are dropped.



Leaky Bucket 1: drain rate = SCR_0 , depth = $L \rightarrow MBS = 1 + \frac{L}{\frac{1}{SCR_0} - T}$

Leaky Bucket 2: drain rate = ρ_{0+1} , depth = 1

- (b) Repeat part (a) if cells that do not conform to SCR_0 and MBS are tagged by changing the CLP bit to 1. Cells that do not conform to p_{0+1} are dropped.

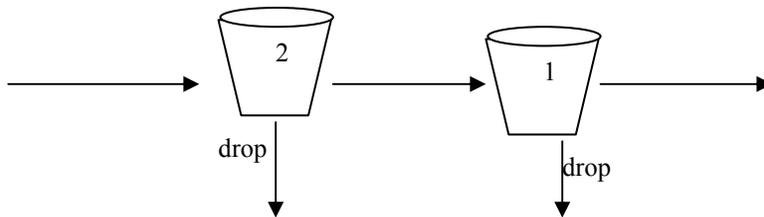


Leaky Bucket 1: drain rate = SCR_0 , depth = $L \rightarrow MBS = 1 + \frac{L}{\frac{1}{SCR_0} - T}$

Leaky Bucket 2: drain rate = ρ_{0+1} , depth = 1

9.24. Suppose that an ATM traffic stream contains cells of two priorities, that is, high-priority cells with CLP = 0 in the headers and low priority cells with CLP = 1. Suppose we wish to police the peak cell rate and the sustainable cell rate of the combined CLP = 0 and CLP = 1 traffic. Give an arrangement of two leaky buckets to do this policing. Nonconforming cells are dropped.

Solution:



Leaky Bucket 1: drain rate = SRC_{o+1} , depth = L (L very large)
 Leaky Bucket 2: drain rate = ρ_{o+1} , depth = 1

Note that this policing can be done using only leaky bucket number 1 if the maximum burst size of the source is known. In this case L could be chosen small enough to accommodate no more than a burst of the maximum size at the peak rate.

9.25. Explain how weighted fair queuing might be used to combine the five ATM service categories onto a single ATM transmission link. How are the different service categories affected as congestion on the link increases?

Solution:

Weighted fair queuing provides a single parameter, namely the weight, to control the service provided to a given flow. If we use weighted fair queuing then each connection must have its own queue and its own weight. The weight must be assigned according to the traffic descriptors and QoS parameters; for example, in a CBR connection the weight must guarantee that the cells in a flow receive the desired bandwidth within the appropriate delay requirements.

Adjustments in weight may be required as new connections are set up. CBR and rt-VBR maintain a constant weight while nrt-VBR and ABR weights vary according to the traffic conditions. The weight of nrt-VBR will depend on traffic congestion and parameters such as PCR, SCR and QoS parameters. The weight for ABR will depend on the bandwidth left free by the CBR and VBR services. UBR should weigh zero, or a very small value.

9.26a. Discuss what is involved in calculating the end-to-end CLR, CTD, and CDV.

Solution:

All of these end-to-end parameters are calculated as the sum of the individual parameters at each node in the VC's path.

The CLR, CTD, and CDV in each of the nodes are based on how the scheduler treats the packets of the VC in question. With the possible exception of highest priority CBR traffic, as congestion at a node increases, the CLR and CTD will increase depending on the priority assigned or the bandwidth allocated to the flow by the scheduler. CDV can be caused by the variation in traffic volume at the node. Because traffic characteristics are random, any calculations of CDV, CTD, and CLR will be statistical in nature.

9.26b. Compare the following two approaches to allocating the end-to-end QoS to per-link QoS: equal allocation to each link; unequal allocation to various links. Which is more flexible? Which is more complex?

Solution:

Unequal allocation of QoS is clearly more flexible, as it represents a superset of the equal allocation method. More specifically, it is more flexible because one can put more strict QoS constraints at nodes with higher traffic flow, since low traffic nodes will be more easily able to provide sufficient service without extensive QoS provisioning.

Unequal allocation is also more complex. As mentioned in part (a), the calculations of end-to-end QoS parameters are generally additive random calculations. These calculations are more complex if the network is non-homogeneous (i.e. if the QoS policies at the nodes are not identical).

9.27. Suppose that an application uses the reliable stream service of TCP, which in turns uses IP over ATM over AAL5.

Solutions follow questions:

- (a) Compare the performance seen by the application if the AAL5 uses CBR connection; nrt-VBR connection; ABR connection; UBR connection.

For a CBR connection the network will do everything possible to deliver tight bounds on delay, jitter and loss by allocating buffer space and bandwidth in the network. The TCP connection will appear as a fixed-rate pipe across the network.

For nrt-VBR, the network will allocate resources to provide the application with a low loss rate. The cells in the flow may experience variations in the delay across the network and if sufficiently long, these delays may trigger the retransmit mechanisms in TCP.

For ABR, if the source responds well to flow control feedback messages that it receives from the network, the application can expect service with a low loss rate. However, the delay across the network can vary over a wide range, and these delay variations will interact with the retransmission and rate-adjustment mechanisms of TCP.

For UBR, no service guarantees are made. In terms of bandwidth the application will receive unpredictable service. The service is completely dependent on the network congestion. Thus, it can be extremely good in times of low congestion or extremely poor (or non-existent) in times of high network congestion. This is the scenario which TCP was developed to operate in.

- (b) Discuss the effect on the performance seen by the application if the ATM connection encounters congestion somewhere in the network.

If the ATM connection encounters congestion somewhere in the network, the change in performance is dependent of the service class of the VC. CBR traffic would get top priority. It implicitly guarantees good isolation, so ideally the performance should remain unchanged during times of congestion. Nrt-VBR is sensitive to loss, but not delay. Thus, it is likely that the network will add delay to the VC via extensive buffering in order to ensure low loss (and to help it meet the delay requirement on other CBR VCs in the node).

ABR connections would get a feedback message to alert the source of the congestion and request that it lower its bit rate accordingly. If it complies, it can expect to experience low loss. Thus, the service of ABR VCs experience lower throughput if congestion is encountered at a node.

UBR traffic is the first to be sacrificed at a node that experiences congestion. Thus the performance of UBR flows is radically degraded in this situation.

9.28. Suppose that an ATM connection carries voice over AAL1. Suppose that the packetization delay is to be kept below 10 ms.

Solutions follow questions:

- (a) Calculate the percentage of overhead if the voice is encoded using PCM.

In AAL1 47 bytes are needed to fill the payload. PCM encodes at 64 Kbps

How long does it take to fill the payload?

$$\frac{(47 \text{ bytes})(8 \text{ bits / byte})}{64 \text{ Kb / s}} = 5.88 \times 10^{-3}$$

The cell is packetized in less than 10 ms then

$$\text{Overhead} = \frac{5 + 1}{53} = 0.1132$$

- (b) Calculate the percentage of overhead if the voice is encoded using a 12 kbps speech encoding scheme from cellular telephony.

How long does it takes to fill the payload?

$$\frac{(47 \text{ bytes})(8 \text{ bits / byte})}{12 \text{ Kb / s}} = 31.33 \times 10^{-3}$$

This is greater than 10 ms. How many bytes are encoded in 10 ms?

$$\frac{12 \text{ Kb / s}}{8 \text{ bits / byte}} \cdot 10 \text{ ms} = 15 \text{ bytes}$$

$$\text{overhead} = \frac{5 + 1 + (47 - 15)}{53} = 0.7169$$

9.29. Explain how the three-bit sequence number in the AAL1 header can be used to deal with lost cells and with misinserted cells.

Solution:

In connection-oriented networks packets always arrive in order. Thus, if a sequence number is out of order, the receiver immediately knows that the previous packet was lost or the current cell was misinserted.

9.30. How much delay is introduced by the two interleaving techniques used in AAL1?

Solution:

In one technique, 15 cells are interleaved and one cell of check bits is added. Thus the delay incurred is 15 cells because, after a cell arrives, the receiver must wait 15 more cells to verify the checksum before the first cell can be forwarded.

In the other technique, 124 cells are transmitted in columns of an array. Four columns of check bits are added. Thus, a 124-cell delay is incurred.

9.31a. How many low-bit-rate calls can be supported by AAL2 in a single ATM connection?

Solution:

The channel header in the CPCS header is 1 byte = 8 bits. Thus, $2^8 = 256$ calls can be supported in a single connection.

9.31b. Estimate the bit rate of the connection if an AAL2 carries the maximum number of calls carrying voice at 12 kbps.

Solution:

Assume the packet size is n bits. For each packet a header of 3 bytes is added.

Maximum number of calls is 2^8 and the packet size for each call is $n + 3 \times 8$ bits. The total number of bits is $2^8(n + 3 \times 8)$.

Considering that there is an overhead of 6 bytes for each cell, the packets are transmitted in

$$2^8 \frac{n+3(8)}{47(8)} \text{ cells.}$$

Thus, for these packets, we transmit:

$$2^8 \left(\frac{n+3(8)}{47(8)} \right) \text{ ATM Cells}$$

Thus, the total bit rate is:

$$\begin{aligned} R &= \left[\frac{12000 \text{ bits/sec}}{1 \text{ call}} \frac{1 \text{ packet}}{n \text{ information_bits}} \right] \left[2^8 \frac{n+24}{47 \times 8} \text{ cells} \right] [53 \times 8 \text{ bits/cell}] \\ &= (3.46417 \times 10^6) \left[\frac{n+24}{n} \right] \text{ bps} \end{aligned}$$

For example, a packetization delay of 10 ms gives $n = 120$ bits. In this case, the bit rate is 4.157 Mbps.

9.31c. What is the percentage of overhead in part (b)?

Solution:

The percentage of overhead =

$$\left(2^8 \times 3 \times 8 + 2^8 \times \frac{n+24}{47 \times 8} \times 6 \times 8 \right) / \left(2^8 \times \frac{n+24}{47 \times 8} \times 53 \times 8 \right) = \left(24 + \frac{n+24}{47/6} \right) / \left(\frac{n+24}{47/53} \right). \text{ Thus the overhead is 26\%.}$$

9.32. Compare the overhead of AAL3/4 with that of AAL5 for a 64K byte packet.

Solution:

For AAL3/4, from Figure 9.15, we see that:

$$\text{Overhead} = 8 + 9 \left(\left\lfloor \frac{65536 + 8}{44} \right\rfloor + 1 \right) + \left(\left(\left\lfloor \frac{65536 + 8}{44} \right\rfloor + 1 \right) \times 44 - (65536 + 8) \right) = 13418 + 16 = 13434 \text{ bytes}$$

For AAL5, from Figure 9.18, we see that:

$$\text{Overhead} = 8 + 5 \left(\left\lfloor \frac{65535 + 8}{48} \right\rfloor + 1 \right) + \left(\left(\left\lfloor \frac{65535 + 8}{48} \right\rfloor + 1 \right) \times 48 - (65535 + 8) \right) = 6838 + 25 = 6863 \text{ bytes}$$

9.33. Discuss the purpose of the error checking that is carried out at the end systems and in the network for an ATM connection that carries cells produced by AAL3/4. Repeat for AAL5.

Solution:

In AAL3/4, a 10-bit CRC is used to detect bit errors. The SN (sequence number) field is 4 bits and is used to detect cell loss and misinsertion, as described in the solution to problem 29.

In AAL5, a 32-bit CRC is used. There is no sequence number field. Instead, AAL5 uses its powerful CRC checksum to detect lost, misinserted or out-of-sequence cells.

9.34. Suppose that in Figure 9.17 packets from A and B arrive simultaneously and each produce 10 cells. Show the sequence of SPDUs produced including segment type, sequence number, and multiplexing ID.

Solution:

ST	SN	MID	Payload	LI	CRC
10	0000	a	-	101100	-
10	0000	b	-	101100	-
00	0001	a	-	101100	-
00	0001	b	-	101100	-
00	0010	a	-	101100	-
00	0010	b	-	101100	-
⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮
00	1000	a	-	101100	-
00	1000	b	-	101100	-
01	1001	a	-	001010	-
01	1001	b	-	001010	-

9.35. Consider SSCOP, the AAL protocol for signaling. Discuss the operation of the Selective Repeat ARQ procedure to recover from cell losses. In particular, discuss how the protocol differs from the Selective Repeat ARQ protocol introduced in Chapter 5.

Solution:

The destination knows that if an out of order SDU arrives, then a packet loss has occurred. The sender must learn this information in order to know which SDUs to retransmit. Rather than learning this from constant ACKs from the receiver, in SSCOP, the sender periodically polls the destination to learn of the receive window status. Based on the response, it will retransmit the necessary packets. This protocol uses less bandwidth than Selective Repeat ARQ discussed in Chapter 5 at the cost of a slower response time to PDU loss.

9.36. Can the SSCOP AAL protocol be modified to provide the same reliable stream service that is provided by TCP? If yes, explain how; if no, explain why not.

Solution:

Both SSCOP and TCP provide error-free delivery, in the proper order, without duplication. However, the service provided by each is very different. Because TCP operates over connectionless networks it needs an explicit method to protect against packet duplication (which occurs when a particular packet is delayed a long time) and packets arriving out of order. SSCOP was designed to run on ATM, so out-of-order packet delivery and duplication are not possible and need not be addressed directly.

To adapt SSCOP to operate over a connectionless network, it could be modified to emulate TCP by giving it a mechanism to exchange window information. Since the SSCOP packet has no field to provide this information in the header, it could be done with the use of control packets.

9.37. The “cells in frames” proposal attempts to implement ATM on workstations attached to a switched Ethernet LAN. The workstation implements the ATM protocol stack to produce cells, but the cells are transmitted using Ethernet frames as follows. The payload of the Ethernet frame consists of a four-byte CIF header, followed by a single ATM header, and up to 31 ATM cell payloads.

Solutions follow questions:

- (a) Find the percentage of overhead of this approach and compare it to standard ATM.

31 ATM cells, each of length 48 bytes, means that the payload of this “cells in frame” proposal would be 1488 bytes. Thus, the percentage of overhead is:

$$\text{cell in frame overhead} = 9 \text{ bytes header} / 9 + 1488 \text{ bytes total} = 0.6\%$$

The overhead of conventional ATM cells is

$$\text{ATM overhead} = 5 \text{ bytes header} / 53 \text{ bytes total} = 9.4\%$$

- (b) The workstation implements ATM signaling and the NIC driver is modified to handle several queues in order to provide QoS. Discuss the changes required in the Ethernet switch so that it can connect directly to an ATM switch.

To connect directly to an ATM switch, the Ethernet switch needs to receive and send ATM traffic. Therefore, it needs an AAL that converts incoming cells-in-frame Ethernet frames into 48-byte blocks. It also needs an ATM layer to add the 5 bytes of appropriate header information to the 48-byte blocks for the AAL in order to make ATM cells. The PC and the modified NIC can provide signaling for the transmission of these cells.

9.38. Compare the size of the address spaces provided by E-164 addressing, AESA addressing, IPv4 addressing, IPv6 addressing, and IEEE 802 MAC addressing.

Solution:

Ignoring the partition of address space fields, which reduces the address space, the address space of an address field is 2^n addresses, where n is the number of bits or 10^m , where m is the number of base 10 digits used to represent the address. Thus, we have:

$$\text{E-164(10digits) raw address space} = 10^{15}$$

$$\text{AESA(18bytes) raw address space} = 2^{(18 \times 8)} = 2.23 \times 10^{43}$$

$$\text{DCC: 2 bytes + 10 bytes + 6 bytes} = (2^{(2 \times 8)}) (2^{(10 \times 8)}) (2^{(6 \times 8)}) = 2.23 \times 10^{43}$$

$$\text{ICD: 4 digits + 10 bytes + 6 bytes} = (10^4) (2^{(10 \times 8)}) (2^{(6 \times 8)}) = 3.40 \times 10^{42}$$

$$\text{E.164: 15 digits + 4 bytes + 6 bytes} = (10^{15}) (2^{(4 \times 8)}) (2^{(6 \times 8)}) = 1.21 \times 10^{39}$$

$$\text{IPv4(32 bits) raw address space} = 2^{(32)} = 4.29 \times 10^9$$

$$\text{Class A: 7 bits + 24 bits} = (2^7) (2^{24}) = 2.15 \times 10^9$$

$$\text{Class B: 14 bits + 16 bits} = (2^{14}) (2^{16}) = 1.07 \times 10^9$$

$$\text{Class C: 21 bits + 8 bits} = (2^{21}) (2^8) = 5.37 \times 10^8$$

$$\text{Class D: 28 bits} = (2^{28}) = 2.68 \times 10^8$$

$$\text{Class E: 28 bits} = (2^{28}) = 2.68 \times 10^8$$

$$\text{IPv6(128 bits) raw address space} = 2^{(128)} = 3.4 \times 10^{38}$$

$$\text{IEEE 802 MAC} = 24 \text{ bits} + 24 \text{ bits} = (2^{24})(2^{24}) = 2.8 \times 10^{14}$$

This consists of a combination of 2^{24} global \times 2^{24} vendor-specific addresses including multicast addresses.

9.39. Could IP addresses be used in ATM? Explain why or why not?

Solution:

IP addresses cannot be used directly in ATM because they each use a totally different address space. Methods of IP over ATM must perform a mapping of one space to the other. For example, Classical IP over ATM (CLIP) described in Section 10.2 performs address resolution using ATM Address Resolution Protocol (ATM ARP).

9.40. Identify the components that contribute to the end-to-end delay experienced in setting up an ATM connection using PNNI.

Solution:

During setup of a PNNI connection, the setup message must propagate from the source to the destination node. At each node generic connection admission control (GCAC) must be used to select a path that is likely to meet the QoS requirements of the connection. Each time there is a shift in the path from one level in the hierarchy to another, a source based routing calculation must be performed using the link state information. After, the destination is reached a connect message must propagate back to the source.

Thus, assuming the path is n -hops long, the total delay incurred is:

- $2n$ one-hop propagation times
- + m route processing times (once per hierarchical shift)
- + n GCAC calculation times
- + delay incurred from any crankback that is performed

9.41. Describe the sequence of DTLs that are used in setting up a connection from A.1.3 to A.2.2 in Figure 9.26. Repeat for a connection from B.4 to A.1.2.

Solution:

A.1.3 to A.2.2	B.4 to A.1.1
At A.1.3 → DTL: [A.1.3, A.1.2] pointer-2 → DTL: [A.1, A.2] pointer-1	At B.4 → DTL: [B.4, B.1] pointer-2 → DTL: [B, A] pointer-1
At A.1.2 → DTL: [A.1, A.2] pointer-2	At B.1 → DTL: [B, A] pointer-2
At A.2.1 → DTL: [A.2.1, A.2.2] pointer-2 → DTL: [A.1, A.2] pointer-2	At A2.4 → DTL: [A2.4, A2.3, A.2.1] pointer-2 → DTL: [A.2, A.1] pointer-1 → DTL: [B, A] pointer-2
	At A2.3 → DTL: [A.2.4, A2.3, A.2.1] pointer-3 → DTL: [A.2, A.1] pointer-1 → DTL: [B, A] pointer-2
	At A2.1 → DTL: [A.2, A.1] pointer-2 → DTL: [B, A] pointer-2

9.42a. Discuss the differences and similarities between PNNI and OSPF.

Solution:

Although both use hierarchical routing, a major difference is that OSPF is restricted to two levels, while PNNI is multi-level. Although both distribute link-state information to calculate source-based trees, PNNI can use complex, multi-metric QoS based link-state information, while OSPF bases link-state on only one metric (namely, bandwidth).

Another difference is that PNNI reserves resources in a path to the destination upon a call request. As a result, PNNI performs routing during the call request, but once a path that guarantees the QoS is reserved, routing computations are no longer done, as packets are forwarded automatically at each node over the path. OSPF performs routing calculations using shortest path algorithms and then sets tables that are used to independently route each arriving packet based on the address alone.

9.42b. Can PNNI be modified to provide QoS routing in the Internet? Explain.

Solution:

Yes, but not easily. PNNI provides QoS by reserving resources and setting up a connection. This is not possible in connectionless IP networks unless modifications are made. Further, in the Internet there is no context for the QoS parameters that are currently needed to perform PNNI routing calculations. To use PNNI in the Internet to provide QoS, Internet based QoS parameters must be defined and the ability to provide connection-oriented service must be added to the current Internet architecture.

This notion that a version of the Internet that provides QoS requires connections motivated the development of connection-based approaches such as RSVP-IntServ and MPLS.

9.43. Compare the hierarchical features of BGP4 and OSPF combined with the features of PNNI.

Solution:

OSPF has a two level hierarchy in which the routers within an autonomous system are grouped in areas that are interconnected by a central backbone area. BGP4 allows multiple autonomous systems to be interconnected creating a third level of hierarchy on top of the autonomous systems. To route a packet from a node in one area in one of the autonomous systems to a destination in an area in another autonomous system the source needs to route the packet to a backbone router within the area in the first level of hierarchy, to the autonomous system boundary router in the second level of hierarchy, and to the destination autonomous system in the third level of hierarchy. This is similar to the hierarchical structure of PNNI with the difference that the hierarchy is created logically in the PNNI and that it is not limited to three levels and can be extended to multiple levels.

9.44. Which aspects of the ATM network architecture depend on the fixed-length nature of ATM cells? What happens if ATM cells are allowed to be variable in length?

Solution:

The fixed-length packets of ATM were selected because of the simplicity in building high-speed switches. Most aspects the ATM architecture extend to variable length packets. Indeed high-speed packet switches have evolved to use cell-like structures to do switching internally while adopting variable length packet protocols on a network basis.

9.45. Which aspects of ATM traffic management change if ATM connections must belong to one of a limited number of classes and if QoS is guaranteed not to individual connections but the class as a whole? Can VPs play a role in providing QoS to these classes?

Solution:

If only a finite number of classes and QoS guarantees exist, the state information in each switch is dramatically reduced. Instead of implementing a queue for each circuit, one queue could be implemented for each of the classes. This allows for better scalability in the switch architecture as the bandwidth and number of connections per incoming link grows.

The tradeoff that comes with making such QoS definitions arises because it is impossible to provide isolation between the flows in one class because the switch can't differentiate between them. As a result, a high QoS class could receive worse performance than a low QoS class solely because a large proportion of the switch traffic is of that class. A solution to this problem is to store state information for each queue in the form of the number of flows aggregated in the class. Resources could be provided accordingly.

VCs in a VP are all treated identically. Thus, VPs can be used to implement the above aggregating architecture by grouping all traffic destined to a particular output interface with a particular service class into one VP. Resources would then be allocated on a granularity of VPs.

9.46. Explain how the ATM architecture facilitates the creation of multiple virtual networks that coexist over the same physical ATM network infrastructure but that can be operated as if they were separate independent networks. Explain how such virtual networks can be created and terminated on demand.

Solution:

The VP capability of ATM allows multiple points in the network to exchange information as if they were connected by dedicated transmission lines. In addition multiple VPs can traverse an ATM switch and interact minimally (if appropriate resources have been reserved.) Thus the switches and links of an ATM network can be partitioned into groups that support different groups of end-systems, that is, into multiple virtual networks. Each such virtual network can be set up and released through a process that generalizes connection setup.

Solutions to Chapter 10

10.1. Can all traffic in an IntServ router be of guaranteed service type? Explain.

Solution:

All the traffic can be of guaranteed service only if all the flows are shaped and if they strictly comply with the traffic specifications. However in practice the total amount of traffic that can be admitted will have some limit. The reason is that to bound the delay, the maximum queue length in the router must be below a certain threshold. If all of the flows are well behaved and without bursts, the queue size will always be reasonably small. However if the flows are bursty, the queue may grow to levels that can cause large unacceptable delays. To guarantee the delay bounds, extra bandwidth should be assigned based to accommodate the burstiness of the flows. This overbooking can sometimes lead to the queue becoming empty, and hence result in underutilization of the output link. Thus, it is reasonable to fill the output bandwidth during such periods with best effort traffic. So in summary, to efficiently utilize the bandwidth and achieve higher throughputs it is necessary to admit best effort traffic alongside guaranteed service traffic.

10.2. Develop a packet-scheduling approach to provide guaranteed service, controlled load service, and best effort service in a router. Can any mix of these three traffic types be supported? Explain.

Solution:

Separate queues must be provided for the packets from each type of service in order for their packets to be serviced appropriately. The set of queues belonging to guaranteed service should always receive service with a rate R , which should be larger than the total guaranteed service traffic. The aggregate rate R is guaranteed and controlled by giving priority to the guaranteed service queues and controlling the service they receive by a leaky bucket shaper with a sustained rate of R . Individual queues within the guaranteed service group are serviced in a round-robin fashion while for each queue its service is controlled by a dedicated leaky bucket shaper with the parameters set based on the requirements of the flow assigned to that queue.

The set of controlled load service queues are serviced during the periods that guaranteed service queues are not serviced. The capacity in excess of R is available to controlled load traffic. Flows are assigned to different queues and weighted fair queueing can be used within the group to service the queues depending on their requirements.

Best effort traffic is given the lowest priority. The queue is serviced only if the other two sets of queues are empty.

10.3. RSVP signaling is very different from ATM signaling. Discuss the differences and list the advantages and disadvantages of each protocol.

Solution:

RSVP uses Receiver-Initiated Reservation rather than ATM's per-hop path request to reserve network resources. RSVP is intended for use with multicasting and is receiver oriented. ATM is designed for unicast traffic and is sender oriented.

RSVP receiver-initiated reservation

Advantage: interoperates with non-RSVP routers; soft state provides robustness with respect to faults, requests can be merged reducing the reservation load on the routers.

Disadvantage: soft state – requires periodic updates, receiver needs to know the path for connection.

ATM per-hop path request

Advantage: hard state – reserve only once.

Disadvantage: Individual VCC is required for each host in multicast application. Source requests the connections so it needs to know the requirements of the receiver.

10.4. Consider the merging of reservations in RSVP. Suppose that there is a single sender and 2^n receivers that are reached through a path that consists of a depth n binary tree that has a router located at each node. What is the total number of reservation request messages generated by the receivers?

Solution:

The total number of reservation request messages generated by the receivers and routers are equivalent to finding the branches of a binary tree with a depth of n . Let N_n be the number of branches in a binary tree with depth n . At level n , there 2^n new branches, and so we have the recursion $N_n = N_{n-1} + 2^n$. (Draw a picture of a tree to verify this.) It is easy to then verify that N_n , total number of request messages, is $2^{n+1} - 1$.

10.5. Suppose RSVP is used to set up a virtual private network to interconnect n users across an Internet. Discuss the appropriateness of the three RSVP reservation styles for this application.

RSVP has three reservation styles, Wildcard Filter, Fixed Filter and Shared Explicit.

Wild card filter (WF) creates a single reservation shared by all senders. It is designed for applications that are unlikely to have multiple senders transmitting simultaneously. It is not appropriate for VPN applications because WF cannot provide the receiver requests if more than one sender is transmitting at the same time.

Fixed filter (FF) creates a distinct reservation for each sender. It is an appropriate solution for VPN requirements. Each sender (VPN) has its own resource reservation and this resource is shared by number of receivers (hosts located in different location belong to the same VPN).

Shared explicit (SE) creates a single reservation shared by a set of explicit senders. It may also be appropriate in VPN applications under certain conditions. The merging of the reservation between senders can result in a more efficient use of network resources.

10.6. Select and justify a choice an RSVP reservation style that is appropriate for the following situations:

Solutions follow questions:

- (a) A multiuser application where a video stream from each participant must be delivered to every other participant.

This situation is a multipoint to multipoint application so the Fixed Filter (FF) is the appropriate choice. FF creates a distinct reservation for each participant (sender) and it can merge downstream requests and forward the sum of flow spec upstream. It is an appropriate reservation style because each participant (sender) will only receive a single reservation request that has been merged by a multiple of other participants. This simplifies the processing overhead for each sender and allows the processing power to focus on accepting multiple video streams for other participants.

- (b) An audioconferencing application that is self-limiting in that more than two participants are unlikely to speak simultaneously.

This audio conferencing application is best suited with the Wild card Filter (WF) style. The WF-style is suitable for applications that are unlikely to have multiple senders transmitting simultaneously. In this application the self-limiting feature prevents two participants from speaking simultaneously.

- (c) A multiuser application in which each participant must receive explicit permission before transmitting information.

The Shared Explicit (SE) reservation style is best suited for this multiuser application. The SE style creates a single reservation shared by a set of explicit senders. In this application, only the participant that receives an explicit permission is allowed to join a single reservation. Any participant that does not receive permission is not allowed to use this single reserved resource.

- 10.7.** What information needs to be carried by Resv messages in RSVP to set up guaranteed IP service with a specific bandwidth and a given delay-bound requirement? How does a router process this information?

Solution:

The information that must be carried by Resv message in RSVP to set up guaranteed IP service with a specific bandwidth and a given delay bound is: the common header, SESSION, RSVP_HOP, TIME_VALUES, STYLE, and Fixed Filter flow descriptor along with flow spec and filter spec. When a router receives this Resv message, it reserves the necessary resources based on the flow specification, which in turn governs the bandwidth and delay-bound requirements. Next, the router exams the filter spec to determine the data packets that correspond to the defined flow spec.

- 10.8.** Discuss the interplay between a QoS routing algorithm, RSVP, and integrated services IP.

Solution:

Integrated services IP provides QoS packet delivery between the source and destination pair. The QoS parameters are dependent on the source requirements, for example, bandwidth requirement and delay-bound requirement. To provide integrated services, it is necessary for the routing protocol (network layer) to provide QoS routing. QoS routing provides for the selection of paths with specific attributes that can support specified QoS requirements. RSVP provides the means for reserving resources along paths identified by QoS routing.

- 10.9.** Explain how the soft-state feature of RSVP allows it to adapt to failures in the network.

Solution:

The RSVP soft-state approach requires each node to refresh its reservation periodically by using Path and Resv messages. When a state is not refreshed within a certain time-out period, the state is deleted. Upon a link failure within the network, a PathErr message is sent to the ingress node (source node). The ingress node simply defines a new route by refreshing the Path message to a different next hop. The old failure path will eventually tear down after the timer is expired. This approach provides error recovery upon failure, however, it is necessary for the ingress node to keep updating the status of the link-state topology and determine the next possible route in case of the primary route fails.

- 10.10.** Discuss whether RSVP provides a scalable means to set up connections to a session that consists of the broadcasting of a given TV program.

Solution:

Broadcasting a TV program is a point to multipoint multicast application. RSVP does provide a scalable means to set up connections to a session in multicast application. RSVP performs the resource reservation for multicast applications, adapting dynamically to changing group membership and changing route. Also, RSVP requires the receiver to initiate and maintain the resource reservation rather than obtaining the bandwidth requirement from each receiver. This process avoids an implosion problem and minimizes the processing overhead to the sender. Finally, the RSVP routers provide different reservation styles and the requests can merge in several ways according to

the reservation style. This prevents duplicate resource reservations if multiple downstream receivers belong to the same multicast group.

10.11. Explain what happens if RSVP is used to set up small-bandwidth flows on a large-bandwidth link? How well does RSVP scale in this situation?

Solution:

If RSVP is used to set up small-bandwidth flows on a large-bandwidth link, the router will have to update the available resource (bandwidth) every time there are small-bandwidth flow request for resources. Moreover, the router needs to keep track of the available resources, and reject resource reservation requests when the available bandwidth cannot support a particular requirement. The router has to keep track of every single flow resource usage during the setup and tear down procedure. All of these activities consume a lot of processing power when the link has large bandwidth being shared by large numbers of small-bandwidth flows.

10.12. Obtain a packet capture file for RSVP either in the lab or downloaded from a web site. Examine the details of Path, Resv, and refresh messages.

Solution:

The following text shows the details of the first PATH message for MPLS path setup shown in Figure 2.27. The file is obtained by printing the capture data from Ethereal. We highlight in color the main fields and interesting information. In particular note the explicit route identified in the message and the LABEL REQUEST. The TIME VALUES field is set to 30000 ms, that is, 30 second refresh period.

Frame 1 (306 bytes on wire, 306 bytes captured)

Arrival Time: Feb 10, 2000 08:49:03.806994000
 Time delta from previous packet: 0.000000000 seconds
 Time relative to first packet: 0.000000000 seconds
 Frame Number: 1
 Packet Length: 306 bytes
 Capture Length: 306 bytes

Ethernet II, Src: 00:90:92:9d:94:01, Dst: 00:d0:63:c3:b8:47

Internet Protocol, Src Addr: 17.3.3.3 (17.3.3.3), Dst Addr: 16.2.2.2 (16.2.2.2)

Resource ReserVation Protocol (RSVP): PATH Message. SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303. SENDER TEMPLATE: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.

RSVP Header. **PATH Message**.

RSVP Version: 1
 Flags: 00
 Message Type: PATH Message. (1)
 Message Checksum: 0xdb58 (correct)
 Sending TTL: 254
 Message length: 264

SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303.

Length: 16
 Class number: 1 - SESSION object
 C-type: 7 - IPv4 LSP
 Destination address: 16.2.2.2 (16.2.2.2)
 Tunnel ID: 1
 Extended Tunnel ID: 285410051 (17.3.3.3)

HOP: IPv4, 210.0.0.1

Length: 12
 Class number: 3 - HOP object
 C-type: 1 - IPv4
 Neighbor address: 210.0.0.1

Logical interface: 0

TIME VALUES: 30000 ms

Length: 8

Class number: 5 - TIME VALUES object

C-type: 1

Refresh interval: 30000 ms (30 seconds)

EXPLICIT ROUTE: IPv4 210.0.0.2, IPv4 204.0.0.1, IPv4 207.0.0.1, ...

Length: 60

Class number: 20 - **EXPLICIT ROUTE object**

C-type: 1

IPv4 Subobject - **210.0.0.2, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 210.0.0.2

Prefix length: 32

IPv4 Subobject - **204.0.0.1, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 204.0.0.1

Prefix length: 32

IPv4 Subobject - **207.0.0.1, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 207.0.0.1

Prefix length: 32

IPv4 Subobject - **202.0.0.1, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 202.0.0.1

Prefix length: 32

IPv4 Subobject - **201.0.0.1, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 201.0.0.1

Prefix length: 32

IPv4 Subobject - **200.0.0.1, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 200.0.0.1

Prefix length: 32

IPv4 Subobject - **16.2.2.2, Strict**

Strict Hop

Type: 1 (IPv4)

Length: 8

IPv4 hop: 16.2.2.2

Prefix length: 32

LABEL REQUEST: Basic: L3PID: IP (0x0800)

Length: 8

Class number: 19 - **LABEL REQUEST object**

C-type: 1

L3PID: IP (0x0800)

SESSION ATTRIBUTE: SetupPrio 0, HoldPrio 0, SE Style, [sys17-3_t1]

Length: 20

Class number: 207 - SESSION ATTRIBUTE object

C-type: 7 - IPv4 LSP

Setup priority: 0

Hold priority: 0

Flags: 0x04

.... ..0 = Local protection not desired

.... ..0. = Label recording not desired

.... .1.. = SE style desired

.... 0... = Bandwidth protection not desired

...0 = Node protection not desired

Name length: 10

Name: sys17-3_t1

SENDER TEMPLATE: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.

Length: 12

Class number: 11 - SENDER TEMPLATE object

C-type: 7 - IPv4 LSP

Sender IPv4 address: 17.3.3.3 (17.3.3.3)

Sender LSP ID: 1

SENDER TSPEC: IntServ: Token Bucket, 625000 bytes/sec.

Length: 36

Class number: 12 - SENDER TSPEC object

C-type: 1 - Integrated Services

Message format version: 0

Data length: 7 words, not including header

Service header: 1 - Traffic specification

Length of service 1 data: 6 words, not including header

Token Bucket TSpec: Rate=625000 Burst=1000 Peak=625000 m=0 M=0

Parameter 127 - Token bucket

Parameter 127 flags: 0x00

Parameter 127 data length: 5 words, not including header

Token bucket rate: 625000

Token bucket size: 1000

Peak data rate: 625000

Minimum policed unit [m]: 0

Maximum packet size [M]: 0

ADSPEC

Length: 84

Class number: 13 - ADSPEC object

C-type: 2

Message format version: 0

Data length: 19 words, not including header

Default General Parameters

Service header 1 - Default General Parameters

Break bit not set

Data length: 8 words, not including header

IS Hop Count - 1 (type 4, length 1)

Path b/w estimate - 1250000 (type 6, length 1)

Minimum path latency - 0 (type 8, length 1)

Composed MTU - 1500 (type 10, length 1)

Guaranteed Rate

Service header 2 - Guaranteed Rate

Break bit not set

Data length: 8 words, not including header

End-to-end composed value for C - 169500 (type 133, length 1)
 End-to-end composed value for D - 1200 (type 134, length 1)
 Since-last-reshaping point composed C - 169500 (type 135, length 1)
 Since-last-reshaping point composed D - 1200 (type 136, length 1)
 Controlled Load
 Service header 5 - Controlled Load
 Break bit not set
 Data length: 0 words, not including header

The following text shows the details of the first RESV message in the opposite direction for the MPLS path setup shown in Figure 10.27. In particular note the last entry which assigns the label 16 to the path.

Frame 2 (146 bytes on wire, 146 bytes captured)

Arrival Time: Feb 10, 2000 08:49:03.909463000
 Time delta from previous packet: 0.102469000 seconds
 Time relative to first packet: 0.102469000 seconds
 Frame Number: 2
 Packet Length: 146 bytes
 Capture Length: 146 bytes

Ethernet II, Src: 00:d0:63:c3:b8:47, Dst: 00:90:92:9d:94:01

Internet Protocol, Src Addr: 210.0.0.2 (210.0.0.2), Dst Addr: 210.0.0.1 (210.0.0.1)

Resource ReserVation Protocol (RSVP): RESV Message. SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303. FILTERSPEC: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.

RSVP Header: RESV Message.

RSVP Version: 1
 Flags: 00
 Message Type: RESV Message. (2)
 Message Checksum: 0x130b (correct)
 Sending TTL: 255
 Message length: 108

SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303.

Length: 16
 Class number: 1 - SESSION object
 C-type: 7 - IPv4 LSP
 Destination address: 16.2.2.2 (16.2.2.2)
 Tunnel ID: 1
 Extended Tunnel ID: 285410051 (17.3.3.3)

HOP: IPv4, 210.0.0.2

Length: 12
 Class number: 3 - HOP object
 C-type: 1 - IPv4
 Neighbor address: 210.0.0.2
 Logical interface: 0

TIME VALUES: 30000 ms

Length: 8
 Class number: 5 - TIME VALUES object
 C-type: 1
 Refresh interval: 30000 ms (30 seconds)

STYLE: Shared-Explicit (18)

Length: 8
 Class number: 8 - STYLE object
 C-type: 1
 Flags: 0x00
 Style: 0x000012 - Shared-Explicit

FLOWSPEC: Controlled Load: Token Bucket, 625000 bytes/sec.

Length: 36

Class number: 9 - FLOWSPEC object
 C-type: 2
 Message format version: 0
 Data length: 7 words, not including header
 Service header: 5 - Controlled Load
 Length of service 5 data: 6 words, not including header
 Token Bucket: Rate=625000 Burst=1000 Peak=1.#INF m=0 M=0
 Parameter 127 - Token bucket
 Parameter 127 flags: 0x00
 Parameter 127 data length: 5 words, not including header
 Token bucket rate: 625000
 Token bucket size: 1000
 Peak data rate: 1.#INF
 Minimum policed unit [m]: 0
 Maximum packet size [M]: 0

FILTERSPEC: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.

Length: 12
 Class number: 10 - FILTER SPEC object
 C-type: 7 - IPv4 LSP
 Sender IPv4 address: 17.3.3.3 (17.3.3.3)
 Sender LSP ID: 1

LABEL: 16

Length: 8
 Class number: 16 - LABEL object
 C-type: 1 (Packet Label)
 Label: 16

The following text shows the details of the first refresh message sent about 23 seconds after the first frame. The contents of the RSVP message are identical to the previous RESV message.

Frame 3 (146 bytes on wire, 146 bytes captured)

Ethernet II, Src: 00:d0:63:c3:b8:47, Dst: 00:90:92:9d:94:01

Internet Protocol, Src Addr: 210.0.0.2 (210.0.0.2), Dst Addr: 210.0.0.1 (210.0.0.1)

Resource ReserVation Protocol (RSVP): RESV Message. SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303. FILTERSPEC: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.

RSVP Header: RESV Message.

RSVP Version: 1
 Flags: 00
 Message Type: RESV Message. (2)
 Message Checksum: 0x130b (correct)
 Sending TTL: 255
 Message length: 108

SESSION: IPv4-LSP, Destination 16.2.2.2, Tunnel ID 1, Ext ID 11030303.

Length: 16
 Class number: 1 - SESSION object
 C-type: 7 - IPv4 LSP
 Destination address: 16.2.2.2 (16.2.2.2)
 Tunnel ID: 1
 Extended Tunnel ID: 285410051 (17.3.3.3)

HOP: IPv4, 210.0.0.2

Length: 12
 Class number: 3 - HOP object
 C-type: 1 - IPv4
 Neighbor address: 210.0.0.2
 Logical interface: 0

TIME VALUES: 30000 ms

Length: 8
 Class number: 5 - TIME VALUES object
 C-type: 1
 Refresh interval: 30000 ms (30 seconds)
STYLE: Shared-Explicit (18)
 Length: 8
 Class number: 8 - STYLE object
 C-type: 1
 Flags: 0x00
 Style: 0x000012 - Shared-Explicit
FLowsPEC: Controlled Load: Token Bucket, 625000 bytes/sec.
 Length: 36
 Class number: 9 - FLOWSPEC object
 C-type: 2
 Message format version: 0
 Data length: 7 words, not including header
 Service header: 5 - Controlled Load
 Length of service 5 data: 6 words, not including header
 Token Bucket: Rate=625000 Burst=1000 Peak=1.#INF m=0 M=0
 Parameter 127 - Token bucket
 Parameter 127 flags: 0x00
 Parameter 127 data length: 5 words, not including header
 Token bucket rate: 625000
 Token bucket size: 1000
 Peak data rate: 1.#INF
 Minimum policed unit [m]: 0
 Maximum packet size [M]: 0
FILTERSPEC: IPv4-LSP, Tunnel Source: 17.3.3.3, LSP ID: 1.
 Length: 12
 Class number: 10 - FILTER SPEC object
 C-type: 7 - IPv4 LSP
 Sender IPv4 address: 17.3.3.3 (17.3.3.3)
 Sender LSP ID: 1
LABEL: 16
 Length: 8
 Class number: 16 - LABEL object
 C-type: 1 (Packet Label)
 Label: 16

10.13. What RSVP reservation style is appropriate for distributed network gaming?

Solution:

Distributed network gaming is similar to video conferencing in nature. The signals from each of the participants are received by all of the other participants. Therefore the fixed-filter style is appropriate in this case. The shared-explicit style may also be appropriate if the number of simultaneous signals is not too large.

10.14. What RSVP reservation style is appropriate for instant messaging? For stock price notification?

Solution:

For instant messaging the fixed-filter style is appropriate because the messages are usually exchanged between the two participants.

For stock price notification the wildcard-filter style is appropriate because the notification can be received by any number of receivers similar to a multicast session. A customized notification however could be different and may require fixed-filter style for individual notifications.

10.15. A sender transmits layered video information. Provide appropriate reservation messages for the receivers in Figure 10.5 if Rx1 is on a 802.11 wireless LAN, Rx2 is on a 3G cell phone, and Rx3 is a workstation on a fast Ethernet LAN.

Solution:

Rx1 sends a reservation request message (Resv) for a shared-explicit style connection in which the flowspec indicates a flow rate of 2 Mbps. The request can be represented by SE(Rx1,{2 Mbps}). Similarly Rx2 sends the request message SE(Rx2,{144 Kbps}) and Rx3 sends the request message SE(Rx3,{100 Mbps}).

10.16. Derive the state lifetime L in RSVP.

Solution:

A local node should be able to tolerate losses of the refresh signal as long as at least one out of K consecutive messages gets through. In the worst case this message could be the last one in a window of K messages. Therefore L must be larger than K times the refresh period. The refresh period is randomized number with a uniform distribution in the range of $[0.5R, 1.5R]$. In the worst case the refresh period could be $1.5R$. As a result to tolerate the worst-case scenario, L must be larger than $K*1.5*R$.

10.17. Discuss possible scheduling mechanisms (e.g., head-of-line priority and weighted fair queueing) that you could use to implement EF PHB, AF PHB group, and DE PHB. Discuss the trade-offs between performance and complexity.

Solution:

If head-of-line (HOL) priority is used to implement EF PHB, AFPHB and DE PHB, at least three separate queues will be defined with different priority classes. The EF PHB assigns the highest priority followed by AF PHB. DE PHB has the lowest priority. HOL priority is easy to implement to provide different levels of service to different classes (PHB). However, HOL priority does not allow for providing some degree of guaranteed access to the transmission bandwidth for the lower priority classes. Another problem is that it does not discriminate users of the same priority. A fairness problem arises when certain users hog the bandwidth by sending an excessive number of packets. In general, HOL priority may be able to provide EF PHB service requirements. On the other hand, it is necessary to have a large buffer queue for AF PHB to achieve low packet drop rate due to the fairness problems. As a result, HOL is not the most suitable traffic management for different PHB applications.

In weighted fair queueing each PHB still has its own queue, but each PHB also has a weight to determine the relative service share it receives. Therefore, EF PHB gets a somewhat larger share of the bandwidth to enable it to meet its delay requirements; and DE PHB only some minimum share of the bandwidth. Weighted fair queueing is a more complicated traffic management approach than the HOL. However, it can prevent the fairness problem. Weighted fair queueing ensures both EF and AF PHB to have a fair share amount of bandwidth to provide high assurance while maintaining EF to have low latency and assured bandwidth. Complex policy control is necessary to determine the right amount of bandwidth for each PHB.

A practical approach to scheduling involves a combination of the above mechanisms. EF PHB traffic is subjected to shaping and then placed in a queue that receives priority over all other traffic. The shaping places a limit on the sustained rate, and hence assures other traffic of a certain amount of bandwidth. AF and DE PHB traffic are subjected to weighted fair queueing in order to assure fairness to the remaining bandwidth.

10.18. Perform a web search of router equipment that implements differentiated services IP. Find out what PHBs are supported by the routers, and if possible, find out what scheduling mechanisms are used to implement the PHBs.

Solution (Based on CISCO 7400 series (7401):

Routers that implement differentiated services support at least real-time (EF PHB) and non real-time forwarding, which are either only best effort or both best effort and guaranteed (AF PHB). Multiple classes within the AF PHB are also supported in these routers.

The implementation of differentiated services and support for different PHBs is achieved in these routers through a combination of different coordinated mechanisms:

- Classification engines, which classify the applications according to differentiated services specifications at the source
- Class-based traffic shaping and bandwidth policing mechanisms
- Class-based buffer management and congestion control mechanisms such as class-based weighted random early detection (CBWRED)
- Class-based scheduling such as class-based weighted fair queueing (CBWFQ)

On top of all of these mechanisms the router OS software is the main element that enables the router to use the mentioned mechanisms to support the differentiated services.

10.19. Perform a web search for microelectronic chips that implement packet scheduling. What QoS mechanisms are supported by these chips? At what bit rates? Do the chips require external memory and if so how much? Explain the memory requirements.

Solution (Based on Vitesse VSC871, VSC881, Terachip TCF16x10, Agere P1-40SAX, BCM-832, TSX-1200, TSX-1400):

Packet scheduling chips are designed either to be used in line cards or jointly with the switching fabric. The chips usually come as part of a chip set. The scheduling chips support QoS by providing multiple buffering queues that can be grouped in different classes. A mechanism is usually provided to schedule the service among the classes. Also within each class an internal mechanism is provided to schedule different queues. This basic structure can be easily used to achieve different qualities of service. This is done through the control software by assigning the traffics to appropriate queues and by setting the parameters for the scheduling mechanisms appropriately.

Some scheduling chips support only two queues for real-time and non-real-time traffics. Other chips support 4 queues for real-time and multiple classes of non-real-time traffic. There are also chips that support a larger number of queues, such as 16 per class, and the two-level scheduling mechanism. The scheduling mechanisms include strict priority to distinguish real-time traffic from best-effort in the simplest case, strict priority in first level and round-robin among the non-real-time to achieve fairness, weighted round-robin to provide differentiation in addition to fairness, and finally weighted fair queueing.

The bit rates vary from several mega bits per second to hundreds of mega bits per second depending on the target type of the routers that can use the schedulers. For example chips intended for use in the core routers support multiple 3.125 Gbps interfaces that can be combined in groups of four to support 10 Gbps ports. The schedulers usually need external buffering for buffering and queue management.

10.20. Provide arguments for and against the statement: “It is essential that Ethernet provide QoS.”

Solution:

Ethernet is a link layer protocol that is primarily used in local area networks. The local traffic in a LAN can be exchanged with low delay and low loss probability. Using LAN switches, partitioning the LAN into smaller clusters interconnected through switches or bridges, and also using higher speeds (10/100/1000) keeps the delay low as the LAN scales and as the traffic increases. Requiring the support for QoS in LAN environment while it is not a critical issue would make the equipment more complex and more expensive. However, the application of Ethernet is not limited to LAN environment. In larger networks such as access networks the distances are larger, the aggregate traffic is greater, and the potential for congestion high. In order to use Ethernet in larger networks while guaranteeing the quality of service, one can argue that it is essential that Ethernet provide QoS.

10.21. Can MPLS make use of DS information to determine the forwarding experienced by different traffic flows? If yes, explain how. If no, explain why not.

Solution:

Yes. At the ingress MPLS router, the router examines the DS field of the IP packet and maps the required service into the label. The MPLS header has the EXP bits available for determining service class. The label thus states the service level in the forwarding of the packet within the MPLS domain. Within the MPLS domain, each router examines the label and performs the forwarding requirement. At the egress node, the router pops the label and forwards the packet to the appropriate next hop. Traffic shaping may be required at the egress node to ensure that traffic conforms to the policing action in the subsequent network. The use of MPLS reduces the overhead required by a conventional router to examine the packet header and determine the required service at each hop.

10.22. What is an appropriate traffic conditioner for EF PHB packets? AF PHB packets?

Solution:

EF PHB packets require low delay and low jitter treatment. These requirements are more readily and efficiently met for traffic flows that are relatively smooth. Therefore EF flows should be subjected to shaping that ensures that the packet flow is smooth flow. Packets that violate the smoothness requirement should be discarded as they may end up not receiving the required low delay low jitter treatment. AF flows are likely to be more bursty, and hence sustained rate and maximum allowable bursting are more important. Packet buffering, marking, and dropping are appropriate traffic conditioning mechanisms for AF flows.

10.23. Suppose that two RSVP-capable users are each connected to an intranet which in turn connects them through an Internet backbone that provides differentiated services IP. Suppose that the intranets provide integrated services IP. Explain the functions that must be implemented in the intranet routers in order to provide controlled-load service and guaranteed service.

Solution:

The router in the intranet must be able to map the controlled-load service and guaranteed service into the differentiated services domain. Since the controlled-load service is very similar to best effort delivery with guaranteed resources, AF PHB is appropriate in the DS backbone. Therefore, if the flow

conforms to the specification (service level agreement), the flow will experience low delay and low loss. Guaranteed service requires bandwidth guarantee and low delay and jitter and hence should be mapped to EF PHB.

10.24. Compare integrated services IP and differentiated services IP in terms of their ability to protect themselves against theft of service.

Solution:

Integrated services IP provides better protection against theft of service than differentiated services IP. Integrated services IP requires admission control before transporting any data packets. As a result, each flow that is admitted should have enough network resources as long as each flow follows its flow specification. If any flow does not conform and exceed its traffic specification, this particular flow will be treated as best effort service to ensure fairness. To prevent theft of service, the admission control process also needs to be secure. Differentiated services IP accords all packets with the same DSCP the same PHB. Thus protection against theft of service involves controlling the setting of DSCP values in the IP packet headers. This suggests strict control by edge routers at the ingress to a DS domain.

10.25. Explain why shaping may be required at the egress node of a differentiated services domain.

Solution:

DS domains that belong to different administrations are likely to use traffic shape parameters in negotiating the exchange of traffic between the domains. A shaper is required at the egress node to ensure the traffic will conform to the policing action in the subsequent network. Therefore, a shaper is important to ensure end-to-end QoS across multiple networks.

10.26. With reference to Figure 10.12, explain how the various types of routers are configured to support packet transport services that require end-to-end low-delay and low-jitter, i.e. voice over IP.

Solution:

A packet transfer service that requires low delay and jitter should be assigned EF PHB. To ensure the low delay and jitter at each DS router, the aggregate arrival rate of packets with EF PHB should be less than the aggregate minimum allowed departure rate. Therefore every DS router must be configured with a minimum departure rate for EF PHB independent of the other traffic. In addition, the aggregate arrival rate must be shaped and policed so that it is always less than the minimum configured departure rate. To ensure this requirement, the access DS router in the local DS domain needs to support traffic classification and traffic conditioning. The interior DS router uses buffer management and scheduling mechanisms to deliver the specific PHB according to the DS field of each packet. The border DS router needs to shape the flow of packets such that it conforms to the traffic conditioning agreement included in the SLA with the border DS router in the transit network. The border DS router in the transit network uses traffic conditioning, including metering, marking and shaping or dropping, to ensure that the arriving traffic behaves according to its agreement and to treat the packets accordingly.

10.27. A common IP-over-ATM overlay model uses an ATM permanent virtual connection (PVC) to connect each pair of IP routers thereby creating a fully meshed topology for the IP network. Thus $N(N-1)/2$ PVCs are needed if there are N routers. Suppose that the routers run a link-state protocol (e.g., OSPF) among themselves. How many advertisements the routers would receive if a given PVC is down. Repeat if a given router is down.

Solution:

If a given PVC is down, the two routers at the two ends detect the link failure and advertise it to all other routers. This generates $2N(N-1)$ advertisements. If a given router is down, each one of the

other routers will detect the failure and advertise the failure of their links to the failed router to all other routers. This will create $(N - 1)(N - 1)$ advertisements.

10.28. Discuss the advantages and disadvantages of different label assignment approaches (i.e., traffic-driven, topology-driven, and request-driven) in terms of number of connections and latency.

Solution:

Traffic-driven label assignment approach: In traffic driven label assignment the assignments follow the traffic pattern in the network. The number of connections will depend on the traffic pattern and its diversity. Monitoring the traffic pattern and adapting the connections will introduce latency.

Topology-driven label assignment approach: The assignment is slower as it may take time to collect updated topology information. Traffic engineering can be incorporated in the assignment of connections, which helps to optimize the resource utilization and increase the number of connections.

Request-driven label assignment approach: Request-driven label assignment can result in an efficient use of the network resources because in this case one single node determines the routes that are created. The node can optimize the route in terms of resource utilization. As a result the number of connections that can be established can be maximized. The procedure requires however updated state information from all nodes and processing, to optimize the routes that introduces higher latency.

10.29. MPOA uses flow detection to setup virtual circuits, whereas MPLS also relies on routing topology to establish virtual circuits. Discuss the advantages and disadvantages of the two approaches.

Solution:

MPOA requires flow detection to set up a VCC

- Advantage: quick adaptivity to flow patterns
- Disadvantages: time delay on VCC establishment, processing load due to churn in VCC setup and teardown

MPLS relies on routing topology to set up a VCC

- Advantage: less processing load than MPOA, traffic engineering capability
- Disadvantage: more complex and requires link state distribution or similar protocol to update the network topology.

10.30. Consider explicit routing in MPLS.

Solutions follow questions:

- (a) Can explicit routing in MPLS be used to select a route that meets the QoS requirements of a particular flow (e.g., guaranteed bandwidth and delay)? Explain.

Yes, explicit routing in MPLS is capable of selecting a route to meet certain QoS requirements. Explicit routing can establish a path between the source and destination that meets the QoS requirements. However, it is necessary that the link-state distribution protocol be able to distribute the required information.

- (b) Give an example at a fine level of granularity that can make use of the QoS explicit routing in part (a). Give another example at a coarser level of granularity.

An example at a fine level of granularity would be to define the QoS requirement on a per application flow basis between a source and destination pair. An example at a coarse level of granularity would be to define the QoS requirement between a set of network prefixes or between two network domains.

- (c) Compare the case where the first LSR computes an explicit path to the case where all LSRs participate in the selection of the path. What information does the first LSR require to select an explicit path?

If the first LSR computes an explicit path, the first LSR will need to have updated link-state information for the whole MPLS domain. On the other hand, when all LSRs participate, each LSR must have the signaling protocol for requesting, establishing, rejecting and tearing down the explicit path. Moreover, in the latter case, each route does not need updated link-state topology of the MPLS domain because each LSR can participate in the selection of path based on its currently known network topology.

10.31. In an MPLS domain, rank the following three flows in terms of their level of aggregation: (a) all packets destined to the same host; (b) all packets with the same egress router; (c) all packets with the same CIDR address.

Solution:

- (a) All packets to the same host – This level of aggregation focuses only on the packets from any source to the same destination by aggregating them into the same flow. This is a fine level of traffic aggregation.
- (b) All packets with the same egress router – All packets are aggregated into the same flow if they are destined to the same egress router. This is a less specific level of traffic aggregation compared to the previous case. More packets can aggregate into the flow as long as they are destined to the same egress router. The egress router distributes the arriving packet to the appropriate destination by examining each packet's IP header.
- (c) All packets with the same CIDR address – This is coarser level of flow aggregation. Basically all packets destined to the same CIDR address (a number of network domains) will aggregate to the same flow. A lot of packets can aggregate and the aggregation path is usually within the backbone network.

10.32. Consider MPLS for unicast traffic. Explain what information is used to specify labels when labels are determined according to (a) host pairs, (b) network pairs, (c) destination network, (d) egress router, (e) next-hop AS, (f) destination AS.

Solution:

- (a) Host pairs – the host pairs' IP address
 (b) Network pairs – the network pairs' network prefix (IP prefix)
 (c) Destination network – the network's IP address and the AS number (if available)
 (d) Egress router – the network or IP address that the egress router can reach.
 (e) Next hop AS – the next hop AS number and the destination the next hop AS can reach
 (f) Destination AS – the destination AS, or a network or IP prefix the AS contains

10.33. Explain why it may be useful for the MPLS header to carry a time-to-live value.

Solution:

A TTL field in an MPLS header can be useful for the same reasons it is used in IP: to terminate packets caught in a loop; to assist utilities such as `traceroute`; to limit the reach of certain packets.

10.34. Compare the overall scalability of the following three cases:

Solutions follow questions:

- (a) Layer 3 forwarding only: each router performs a longest-prefix match to determine the next forwarding hop;

Longest-prefix matching can take much processing time if the routing table is very large. This can limit the rate at which packets can be routed.

(b) Layer 3 forwarding and some Layer 2 MPLS forwarding;

The layer 2 forwarding helps to reduce the load of the layer 3 forwarding. This method is very scalable if the majority of packets use the layer 2 forwarding.

(c) Layer 2 MPLS forwarding only.

Layer 2 MPLS forwarding provides fast packet forwarding capabilities independent of the routing table size. The packets will experience minimal delay on each hop. However the node cannot handle IP packets that do not have MPLS labels.

10.35. Are MPLS label switching and packet filtering at firewalls compatible? Explain.

Solution:

Packet filtering at firewalls operates at or above the IP layer and MPLS label switching runs below IP layer. Filtering requires extracting specific fields from a packet header and applying rules to the values of the fields to determine the handling of a given packet. Network processors that can extract and process multiple fields from multiple headers in a packet are now becoming available. In particular MPLS is capable of establishing an explicit path with an egress node that has firewall capability. The firewall at the egress node decides whether to forward or deny the packet that is coming from an explicit path in the MPLS domain.

You may be interested in doing a web search for network processors. You will find Intel, IBM, and many others involved in the development of network processors.

10.36. Discuss what factors determine the level of computational load for label assignment and distribution for the following three cases: (a) topology-driven label assignment, (b) request-driven label assignment, (c) traffic-driven label assignment.

Solution:

(a) Topology-driven label assignment – The label assignment for this approach basically follows the shortest-path algorithms used in the IP network. As a result, label assignment is based on the shortest path found by the IP layer and the computational load is minimal.

(b) Request-driven label assignment – The ingress router requests a traffic flow within an MPLS domain. This requires each LSR to have the current status of the link-state network topology. The computational load for label assignment and distribution is higher than for the topology-driven approach.

(c) Traffic-driven label assignment – The router must monitor the traffic flow/pattern on each of its I/O ports before assigning any label to the link. Computational load is the highest among the three approaches because traffic monitoring requires extra processing overhead in this approach.

10.37. Compare MPLS label stacks and IP-over-IP encapsulation tunnels in terms of their traffic engineering capabilities.

Solution:

Traffic engineering involves path pre-establishment and re-routing in case of failure. MPLS label stacks are more flexible than IP-over-IP encapsulation in terms of traffic engineering. In MPLS label stack, path re-routing simply changes the label value along the routing path without altering the payload. However, in IP-over-IP encapsulation, the re-routing requires changing the IP address as

well as manipulating the IP header (including header checksum and possibly fragmentation offset). This results in complicated processing overhead for the router and poses delay on the packet transmission.

10.38. Compare the following two approaches to traffic engineering:

Solutions follow questions:

- (a) Offline computation: A central site knows the demand (total bandwidth requirement) between each pair of LSRs and computes the paths in the network.

The centralized database contains an updated and accurate network state topology. As a result, it provides the optimal path between any source and destination pair within the domain. However, scalability is a problem. As the network becomes large, the volume of requests becomes difficult to handle.

- (b) Online computation: An ingress LSR independently computes a path to a specified egress LSR as new demands arrive at the ingress LSR.

The processing load is distributed to individual ingress nodes. However, the ingress node must have the updated link-state network topology and this topology may be less accurate compared with the centralized database as multiple LSRs may assign paths at the same time.

10.39. Suppose that an MPLS network consisting of N routers supports C diffserv traffic classes. A trunk is defined as traffic from a single traffic class that is aggregated into an LSP.

Solutions follow questions:

- (a) What is the maximum number of trunks if a trunk is defined for every pair of ingress router and egress router?

If all of the N routers are considered as ingress/egress pair routers, the the maximum number of trunks for every pairs of ingress and egress router will be $N(N - 1) \times C$.

- (b) Suppose that the trunks in part (a) are merged if they have the same traffic class and the same egress routers. What is the maximum number of multipoint-to-point trees rooted in an egress router?

The trunks are merged if they have the same traffic class and the same egress router, the maximum number of trees rooted in an egress router is $N \times C$.

- (c) Can you think of other ways of merging trunks?

The trunks could be merged if they have the same egress router and even if they have a different traffic class as long as enough network resources are allocated. The maximum number of trees rooted in an egress router will become N .

10.40. Can an MPLS label correspond to all packets that match a particular integrated services filter specification? If yes, explain how such a label may be used. If no, explain why not.

Solution:

The MPLS label is primarily concerned with the forwarding of a packet. The forwarding information is accessed through the label in the MPLS header. The integrated services filter is concerned with determining which packets are allowed access to a given set of reserved resources. However the MPLS label can also be used to determine the type of scheduling that the associated packet flow is to

undergo. Therefore the MPLS label can be used to determine that a certain packet flow is to undergo a certain scheduling treatment that corresponds to the integrated service filter.

10.41. Suppose that network 2 in Figure 10.15 is a SONET network and that networks 1 and 3 are IP networks. Explain the overlay and peer-to-peer models to network interconnection. What changes if network 1 and network 3 are MPLS networks? What changes if network 2 uses GMPLS?

Solution:

In the overlay model, the server network operates an independent control plane from the client networks. The client networks receive the service from the server network in the form of logical links connecting their border nodes. In this case the border nodes B and F contain packet-over-SONET interfaces that connect to nodes C and E in the SONET network respectively. A SONET connection is established between nodes C and E by the control plane in network 2 to provide the transport service to the connection between networks 1 and 3.

In the peer-to-peer model, a single instance of the control plane is used for server and client networks. In this case the nodes in the network 2 have an IP forwarding layer that interact with the IP control plane in nodes in networks 1 and 3, providing a unified end-to-end IP control. The nodes in network 2 translate the routing result of the IP layer into SONET connections to transport IP flows between networks 1 and 3.

If networks 1 and 3 are MPLS networks, in both overlay and peer-to-peer cases the difference is in how traffic flows are handled within the client networks. Instead of a larger IP network, we have a larger combined MPLS network.

If network 2 supports GMPLS, then it is possible for GMPLS signaling to create a tunnel for the IP connections between networks 1 and 3 with SONET connectivity across network 2.

10.42. Suppose that network 2 in Figure 10.15 provides optical wavelength service and that networks 1 and 3 are IP networks. Explain the overlay and peer-to-peer models to network interconnection. What changes if network 1 and network 3 are MPLS networks? What changes if network 2 uses GMPLS?

Solution:

In the overlay model, nodes B and F contain IP over λ interfaces that connect to nodes C and E in network 2 respectively. Network 2 provides wavelength connectivity between nodes B and F. The connectivity is established by the control plane in network 2 with the interaction with the control planes in networks 1 and 3.

In the peer-to-peer model an IP forwarding layer on top of the wavelength service uses the IP forwarding results to switch wavelengths across the network 2 to create connectivity between networks 1 and 2.

If networks 1 and 3 are MPLS networks, instead of signaling for IP flows, MPLS connections are mapped into wavelengths, a process which simplifies the signaling considerably.

If network 2 is GMPLS, it can simply create a tunnel for the IP connections between networks 1 and 3 using wavelength connectivity across the network 2, which simplifies the interconnection problem.

10.43. Suppose that network 2 in Figure 10.15 provides optical wavelength services and that networks 1 and 3 are SONET networks. Explain the overlay and peer-to-peer models to network interconnection. What changes if all the networks use GMPLS?

Solution:

In overlay model, SONET connections are mapped into wavelengths in nodes B and F and transported across network 2 using wavelength connectivity created by the control plane in network 2.

In peer-to-peer model the nodes in network 2 interact with SONET control plane to establish the end-to-end SONET connections, which uses the wavelength connectivity in network 2 as part of its transport path.

If all the networks use GMPLS, a unified control plane establishes the end-to-end connection, while the connection is mapped on to the SONET transport in networks 1 and 3 and wavelength transport in network 3.

10.44. Does the notion of label merging apply to GMPLS? Explain.

Solution:

The notion of label merging cannot be applied to GMPLS at all levels in a similar way it is applied to MPLS. In the wavelength level for instance, multiple wavelengths cannot be merged into a new wavelength.

10.45. Suppose that MPLS traffic runs over a SONET network. Is protection and restoration required at both the SONET and MPLS levels? Explain.

Solution:

MPLS can provide both protection and restoration in an efficient way, as MPLS is the lowest level that is aware of the link states while it has a comprehensive view of network topology and at the same time has the means for traffic management. Protection at SONET level is done in the physical layer and thus it is much faster, but requires more equipment deployment that is more costly. Restoration relies on the network level routing protocols to calculate alternate paths. In general, MPLS does not require SONET protection underneath, unless the service restoration times cannot be met by MPLS mechanisms.

10.46. Consider a 4-node fully connected optical network. Suppose that each pair of nodes is connected by optical fiber that carries 100 wavelengths in each direction. Compare the scale required to do OSPF routing with and without link bundling.

Solution:

Without link bundling OSPF will distribute the link state for every wavelength as a separate routing adjacency. With link bundling all parallel adjacencies can be aggregated into one bundled link reducing the number of links and consequently the number of link state messages by a factor of 100.

10.47. Should RTP be a transport layer protocol? Explain why or why not.

Solution:

The transport layer is intended for the end-to-end transfer of data with between processes in end systems. This basic multiplexing capability is addressed by TCP and UDP. If RTP were a transport protocol it would unnecessarily duplicate this capability. By running RTP run on top of a transport layer protocol such as UDP, the focus can be kept on the additional capabilities that are required to provide real-time service.

10.48. Can RTP be used to provide reliable real-time communication? If yes, explain how. If no, explain why not.

Solution:

RTP cannot be used to provide reliable real-time communication. It can only deal with reordering of packets but cannot deal with packet losses. It can help to reduce the impact of losses and also delay and jitter incurred in the network. RTCP can also be used to help the source to adjust to the conditions of the network to reduce the impacts on loss, delay, and jitter. But nevertheless it cannot guarantee a completely reliable real-time communication.

10.49. Can RTP operate over AAL5? Explain.

Solution:

RTP can operate over AAL5 by mapping RTP PDUs directly onto AAL5 PDU's in point-to-point applications. RTP is designed for multicast scenarios, which is not readily handled by ATM.

10.50. Suppose that the marker bit in the RTP header is used to indicate the beginning of a talk spurt. Explain the role of the timestamps and sequence numbers at the receiver.

Solution:

The timestamps are used for synchronization and jitter cancellation. The sequence numbers are used to detect packet losses, and to recover packet sequences. In the case of the talk spurt, the timestamps and sequence numbers are used to regenerate the same talk spurt at the receiver with as much accuracy as possible. The marker indicates the sequence of packets that contain the talkspurt.

10.51. Discuss the relative usefulness of RTCP for unicast and multicast connections.

Solution:

RTCP can be more useful in unicast connections because in this case the source should adjust to only one route and one receiver. In the case of multicast connections RTCP will provide feedback on multiple routes and if the source adjusts itself to any one of the routes other routes will share the impact unnecessarily.

10.52. Run Microsoft NetMeeting or some other conferencing application that carries voice and/or video. Capture the packets that are exchanged during the setting up of the session and during the data transfer phase. Determine what protocols are used for session setup and for data transport. Examine details of the fields in each packet type. What type of media encoding is used?

Solution:

Setup Capture:

No.	Time	Source	Destination	Protocol Info	
6	18.904189	192.168.0.149	192.168.0.143	TCP 1748 > 1720 [SYN] Seq=1739645016 Ack=0 Win=16384 Len=0	TCP control packets
7	18.905196	192.168.0.143	192.168.0.149	TCP 1720 > 1748 [SYN, ACK] Seq=4252100644 Ack=1739645017 Win=17520 Len=0	
8	18.905366	192.168.0.149	192.168.0.143	TCP 1748 > 1720 [ACK] Seq=1739645017 Ack=4252100645 Win=17520 Len=0	
11	19.497846	192.168.0.149	192.168.0.143	H.225.0 CS: Setup-UUIE	H.225.0 setup
12	19.769449	192.168.0.143	192.168.0.149	TCP [Desegmented TCP]	
13	20.099818	192.168.0.149	192.168.0.143	TCP 1748 > 1720 [ACK] Seq=1739645217 Ack=4252100649 Win=17516 Len=0	
14	20.101044	192.168.0.143	192.168.0.149	H.225.0 CS: Alerting-UUIE	Alerting
16	20.501086	192.168.0.149	192.168.0.143	TCP 1748 > 1720 [ACK] Seq=1739645217 Ack=4252100688 Win=17477 Len=0	
20	29.091030	192.168.0.143	192.168.0.149	TCP [Desegmented TCP]	
21	29.329256	192.168.0.149	192.168.0.143	TCP 1748 > 1720 [ACK] Seq=1739645217 Ack=4252100692 Win=17473 Len=0	
22	29.330385	192.168.0.143	192.168.0.149	H.225.0 CS: Connect-UUIE	Connecting
23	29.400799	192.168.0.149	192.168.0.143	TCP 1749 > 1862 [SYN] Seq=1740980379 Ack=0 Win=16384 Len=0	
24	29.401781	192.168.0.143	192.168.0.149	TCP 1862 > 1749 [SYN, ACK] Seq=4253464033 Ack=1740980380 Win=17520 Len=0	
25	29.401944	192.168.0.149	192.168.0.143	TCP 1749 > 1862 [ACK] Seq=1740980380 Ack=4253464034 Win=17520 Len=0	
26	29.405685	192.168.0.149	192.168.0.143	TCP [Desegmented TCP]	
27	29.453530	192.168.0.143	192.168.0.149	TCP [Desegmented TCP]	
28	29.453752	192.168.0.149	192.168.0.143	H.245 TerminalCapabilitySet MasterSlaveDetermination	Negotiating channel usage
29	29.455958	192.168.0.143	192.168.0.149	H.245 TerminalCapabilitySet MasterSlaveDetermination	
30	29.465312	192.168.0.149	192.168.0.143	TCP [Desegmented TCP]	
31	29.471165	192.168.0.143	192.168.0.149	TCP [Desegmented TCP]	
32	29.471402	192.168.0.149	192.168.0.143	H.245 TerminalCapabilitySetAck MasterSlaveDeterminationAck	
33	29.472271	192.168.0.143	192.168.0.149	H.245 TerminalCapabilitySetAck MasterSlaveDeterminationAck	
34	29.678682	192.168.0.149	192.168.0.143	TCP [Desegmented TCP]	
35	29.679868	192.168.0.143	192.168.0.149	H.245 OpenLogicalChannel OpenLogicalChannel OpenLogicalChannel	Requesting bandwidth
43	30.532924	192.168.0.149	192.168.0.143	TCP 1749 > 1862 [ACK] Seq=1740980892 Ack=4253464654 Win=16900 Len=0	
45	35.545028	192.168.0.143	192.168.0.149	RTCP Receiver Report	
48	35.546773	192.168.0.149	192.168.0.143	RTCP Receiver Report	
50	1239.290373	192.168.0.149	192.168.0.143	RSVP PATH Message. SESSION: IPv4, Destination 192.168.0.143, Protocol 17, Port 49608.	
SENDER TEMPLATE: IPv4, Sender 192.168.0.149, Port 49608.					
51	1239.322672	192.168.0.143	192.168.0.149	RSVP RESV Message. SESSION: IPv4, Destination 192.168.0.143, Protocol 17, Port 49608.	

Data Transfer Capture:

No.	Time	Source	Destination	Protocol Info	
62	1240.429251	192.168.0.143	192.168.0.149	TCP 1862 > 1749 [ACK] Seq=4253464654 Ack=1740980896 Win=17004 Len=0	
63	1240.429470	192.168.0.149	192.168.0.143	H.245 OpenLogicalChannelReject OpenLogicalChannelAck OpenLogicalChannelAck	RTP video H.263 packet
MiscellaneousCommand MiscellaneousCommand OpenLogicalChannelConfirm					
65	1240.622962	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12977, Time=3017250	
66	1240.623218	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12978, Time=3017250	
71	1240.854456	192.168.0.143	192.168.0.149	TCP 1862 > 1749 [ACK] Seq=4253464654 Ack=1740980980 Win=16920 Len=0	
72	1240.854730	192.168.0.143	192.168.0.149	TCP 1503 > 1752 [ACK] Seq=4255085782 Ack=1742545885 Win=17315 Len=0	
74	1240.915746	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12980, Time=3030750	
75	1240.916004	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12981, Time=3030750	
76	1240.916239	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12982, Time=3030750, Mark	
87	1240.977683	192.168.0.149	192.168.0.143	TCP 1753 > 1503 [PSH, ACK] Seq=1742695460 Ack=4255248829 Win=17520 Len=25	
88	1240.979358	192.168.0.143	192.168.0.149	TCP 1503 > 1753 [PSH, ACK] Seq=4255248829 Ack=1742695485 Win=17495 Len=21	
89	1241.212546	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12983, Time=3044070	
90	1241.212801	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12984, Time=3044070, Mark	More control
91	1241.231672	192.168.0.149	192.168.0.143	TCP 1753 > 1503 [ACK] Seq=1742695485 Ack=4255248850 Win=17499 Len=0	
92	1241.231775	192.168.0.149	192.168.0.143	TCP 1749 > 1862 [ACK] Seq=1740980980 Ack=4253464658 Win=16896 Len=0	
93	1241.232617	192.168.0.143	192.168.0.149	H.245 MiscellaneousCommand	
97	1241.364228	192.168.0.143	192.168.0.149	RTP Payload type=ITU-T H.263, SSRC=4062428632, Seq=32093, Time=1885770, Mark	
98	1241.366210	192.168.0.149	192.168.0.143	TCP [Desegmented TCP]	
99	1241.380190	192.168.0.149	192.168.0.143	TCP 1753 > 1503 [PSH, ACK] Seq=1742695485 Ack=4255248850 Win=17499 Len=268	
100	1241.384497	192.168.0.143	192.168.0.149	TCP 1503 > 1753 [PSH, ACK] Seq=4255248850 Ack=1742695753 Win=17227 Len=106	
Time=51984, Mark					
163	1242.823618	192.168.0.143	192.168.0.149	RTP Payload type=ITU-T G.723, SSRC=756814963, Seq=32091, Time=52704	RTP audio G.723
164	1242.836785	192.168.0.149	192.168.0.143	TCP 1753 > 1503 [ACK] Seq=1742695794 Ack=4255249018 Win=17331 Len=0	
165	1242.836872	192.168.0.149	192.168.0.143	TCP 1754 > 1503 [ACK] Seq=1742808026 Ack=4255332966 Win=17168 Len=0	
166	1242.897739	192.168.0.143	192.168.0.149	RTP Payload type=ITU-T G.723, SSRC=756814963, Seq=32092, Time=53424	
167	1242.908675	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12996, Time=3120210	
168	1242.908963	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12997, Time=3120210	
173	1243.099425	192.168.0.143	192.168.0.149	RTP Payload type=ITU-T H.263, SSRC=4062428632, Seq=32109, Time=1963890, Mark	
174	1243.107713	192.168.0.143	192.168.0.149	RTP Payload type=ITU-T G.723, SSRC=756814963, Seq=32093, Time=54144	
175	1243.202816	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=12999, Time=3133620	
176	1243.203076	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=13000, Time=3133620	
177	1243.205246	192.168.0.149	192.168.0.143	RTP Payload type=ITU-T H.263, SSRC=1488263488, Seq=13001, Time=3133620, Mark	

10.53. Repeat the previous problem for a distributed network game.

Solution:

Pending

10.54. Suppose a router acts as a firewall and filters packets. Can the router identify RTP packets? Can the router identify RTP packets from different connections?

Solution:

RTP uses randomly selected port number. Firewalls normally restrict all port numbers except for a select set of well-known numbers. Consequently RTP packets can have difficulty traversing firewalls. Various fixes for different types of RTP-based applications and firewalls are available.

10.55. Suppose you want to implement an IP telephony system.

Solutions follow questions:

(a) Look up the RTP payload format for G.711. Discuss the transmission efficiency for the resulting system.

G.711 is used for encoding the 64 Kbps voice. Assuming 10 msec sampling frame size, the payload is 80 bytes RTP header is 40 bytes and the efficiency is 66%. Efficiency will increase if larger frame size is used. For 20 msec frame size, payload is 160 bytes and efficiency is 80%. For 30 msec frame size, payload is 240 bytes and efficiency is 85%.

(b) Repeat for G.729.

In this case, the signal rate is 8 kbps, and for 10 msec frame size the payload is only 10 bytes. Considering 40 bytes RTP header, the efficiency is 20%. Assuming 20 msec frame size, payload is 20 bytes and the efficiency is 33%. Assuming 30 msec frame size, payload is 30 bytes and efficiency is 43%.

10.56. Explain the relevance to scalability of the SIP protocol not having to maintain state. Contrast this with the case of traditional telephone networks.

Solution:

In the telephone network once a call is setup the network is responsible to maintain the state of the connection all across the network in order to keep the connection. In SIP protocol once the call is setup the end systems can continue the call without requiring maintaining the connection. The two end systems can communicate directly once they know each other's address. As a result SIP does not have scalability problem, any number of new users can join in and make any number of new calls. The only limitation is the network traffic. In contrast, in the telephone network scaling the network requires installing complex switches and other equipment proportional to the number of new users and new calls they are required to handle.

10.57. Consider whether and how SIP can be used to provide the following services:

Solutions follow questions:

- (a) Call Display: The number and/or name of the calling party is listed on a screen before the call is answered.

The invite message by the SIP client carries the caller information. The information can be extracted by the user agent server and displayed on the screen when the message is received.

- (b) Call Waiting: A special sound is heard when the called party is on the line and another user is trying to reach the called party.

The user agent server in the called party's system can indicate the new call request to the user by generating a special sound whenever it receives the new invite message while another call is in progress.

- (c) Call Answer: If the called party is busy or after the phone rings a prescribed number of times, the network gives the caller the option of leaving a voice message.

The called party's proxy server or the user agent server of the called party can return a message to the calling party giving the option of leaving a message in the case the called party is busy or does not answer the call. An agent can be used to record the message. The agent can directly communicate with the user agent client of the caller to record the message.

- (d) Three-Way-Calling: Allows a user to talk with two other people at the same time,

SIP supports multi party conference calls. Three-way calling can be handled as a three-party conference call. The user agent client of the caller can create two simultaneous invite messages to the two other people to initiate the call. Once the call is setup, the messages of each party are sent to both of the other two parties.

10.58. A professor has three locations at his university: his regular professorial office, his lab, and the administrative office for a research center he operates.

Solutions follow questions:

- (a) Explain how the *fork* capability of SIP may be used to direct an Internet phone call to the professor.

The addresses of the three locations are given to the proxy server managing the SIP calls for the professor. The proxy server forks the invite messages that it receives to all three locations.

- (b) Suppose that the professor spends Tuesdays working at a small startup, where he also has an Internet phone. Explain how a REGISTER message can be used to direct calls from the professor's office in the university to the startup.

The register method is used to send location information to a SIP server. In this method the user sends a REGISTER message to its SIP server giving it a temporary forwarding address. In this example the professor initiates the method using the IP telephone. The IP telephone then sends the register message to the SIP server automatically.

- (c) Suppose that on Wednesday, the professor forgets to change the forwarding instructions from the day before. Can a call placed to him at the startup reach him at the office? At the lab?

The calls placed at the startup cannot reach him at the office or lab until the proxy server locates him in one of the new locations unless the call forwarding method described in (a) is in use.

10.59. Suppose that a university campus is covered by 802.11 wireless LANs that are connected to the campus IP backbone. Explain how SIP, RTP and associated protocols can be used to provide mobile phone service across the campus. Explain how the service can be extended to provide service to other universities.

Solution:

SIP and RTP operate in the same manner as they would in a wired network environment, and the extension to other sites would follow the same approach as in wired networks. The tricky part is in authentication, the handling of IP addressing, the interconnecting of the various wireless LANs, and the handling of mobility and the handoff latency impacts on voice service. One extreme approach would have each wireless access point connected to a router and all communications thereafter handled in layer 3. Mobile IP could be used in this case to handle mobility. Current implementations may introduce large handoff delays. An opposite extreme would have all wireless LANs interconnected in a large layer-2 bridged network. Bridge learning algorithms could be used to determine the location of a given user. The learning time of current implementations are too large for effective handoff of voice calls, and so modifications of bridge learning algorithms may be required.

10.60. A recent proposal involves replacing the H.323 gateway with a trunking gateway that is responsible solely for media format conversion. Signaling to establish a call is handled by a call agent that talks to the client on the packet network side, to a signaling #7 gateway to the telephone network, and to the trunking gateway.

Solutions follow questions:

(a) Explain how a telephone call is set up between an Internet terminal and a telephone set.

Let's assume that the client on the packet network initiates a call. The client is connected to the call agent and sends messages indicating the request for the call and the desired destination. The call agent talks to an SS7 gateway to perform the signaling to establish the call in the telephone network. At the same time the call agent talks to the trunking gateway to setup a trunk for the call. Once the signaling is done, the call is setup between the internet terminal and the telephone set through the trunking gateway.

(b) Does this system have greater scalability than an H.323 gateway has?

H.323 involves complex signaling and requires heavy processing that can only be done by powerful end systems attached to packet network. The new system described here is more scalable because it eliminates the need for H.323 servers and replaces them with much simpler agents in the network.

(c) Can the call agents be used to move telephone intelligent network functionality to the Internet?

Yes the call agents can be used to provide some of the intelligent network functionality.