Computer Networks

Problems and solutions

Chapter one

Communication & Computer Network key-Lab of GD

School of Computer Science & Engineering

South China University of Technology

Questions:

3、4、11、12、13、14、18、19、20、22、24、26、28、33、36

Question3: The performance of a client-server system is influenced by two network factors: the bandwidth of the network (how many bits/sec it can transport) and the latency (how many seconds it takes for the first bit to get from the client to the server). Give an example of a network that exhibits high bandwidth and high latency. Then give an example of one with low bandwidth and low latency.

Answer: A transcontinental fiber link might have many gigabits/sec of bandwidth, but the latency will also be high due to the speed of light propagation over thousands of kilometers. In contrast, a 56-kbps modem calling a computer in the same building has low bandwidth and low latency.

Question4: Besides bandwidth and latency, what other parameter is needed to give a good characterization of the quality of service offered by a network used for digitized voice traffic?

Answer: A uniform delivery time is needed for voice, so the amount of jitter in the net-work is important. This could be expressed as the standard deviation of the delivery time. Having short delay but large variability is actually worse than a somewhat longer delay and low variability.

Question11: What are two reasons for using layered protocols?

Answer: Among other reasons for using layered protocols, using them leads to breaking up the design problem into smaller, more manageable pieces, and layering means that protocols can be changed without affecting higher or lower ones.

Question12: The president of the Specialty Paint Corp. gets the idea to work with a local beer brewer to produce an invisible beer can (as an anti-litter measure). The president tells her legal department to look into it, and they in turn ask engineering for help. As a result, the chief engineer calls his counterpart at the other company to discuss the technical aspects of the project. The engineers then report back to their respective legal departments, which then confer by telephone to arrange the legal aspects. Finally, the two corporate presidents discuss the financial side of the deal. Is this an example of a multilayer protocol in the sense of the OSI model?

Answer: No. In the ISO protocol model, physical communication takes place only in the lowest layer, not in every layer.

Question13: What is the principal difference between connectionless communication and connection-oriented communication?

Answer: Connection-oriented communication has three phases.In the establishment phase a request is made to set up a connection.Only after this phase has been successfully completed can the data transfer phase be started and data transported.Then comes the release phase.Connectionless communication does not have these phases.It just sends the data.

Question14: Two networks each provide reliable connection-oriented service. One of them offers a reliable byte stream and the other offers a reliable message stream. Are these identical? If so, why is the distinction made? If not, give an example of how they differ.

Answer: Message and byte streams are different.In a message stream,the network keeps track of message boundaries.In a byte stream,it does not.For example,suppose a process writes 1024 bytes to a connection and then a little later writes another 1024 bytes.The receiver then does a read for 2048 bytes. With a message stream,the receiver will get two messages,of 1024 bytes each.With a byte stream,the message boundaries do not count and the receiver will get the full 2048 bytes as a single unit.The fact that there were originally two distinct messages is lost.

Question18: Which of the OSI layers handles each of the following:

(a) Dividing the transmitted bit stream into frames.

(b) Determining which route through the subnet to use.

Answer: (a)Data link layer.(b)Network layer.

Question19: If the unit exchanged at the data link level is called a frame and the unit exchanged at the network level is called a packet, do frames encapsulate packets or do packets encapsulate frames? Explain your answer.

Answer: Frames encapsulate packets. When a packet arrives at the data link layer, the entire thing, header, data, and all, is used as the data field of a frame. The entire packet is put in an envelope(the frame),so to speak(assuming it fits).

Question20: A system has an n-layer protocol hierarchy. Applications generate messages of length M bytes. At each of the layers, an h-byte header is added. What fraction of the network bandwidth is filled with headers?

Answer: With n layers and h bytes added per layer,the total number of header bytes per message is hn,so the space wasted on headers is hn.The total message size is M+nh,so the fraction of bandwidth wasted on headers is hn/(M+hn).

Question22: What is the main difference between TCP and UDP?

Answer: TCP is connection oriented,whereas UDP is a connectionless service.

Question24: The Internet is roughly doubling in size every 18 months. Although no one really knows for sure, one estimate put the number of hosts on it at 100 million in 2001. Use these data to compute the expected number of Internet hosts in the year 2010. Do you believe this? Explain why or why not.

Answer: Doubling every 18 months means a factor of four gain in 3 years.In 9 years,the gain is then 43or 64,leading to 6.4 billion hosts.My intuition says that ismuch too conservative,since by then probably every television in the worldand possibly billions of other appliances will be on home LANs connected to the Internet. The average person in the developed world may have dozens of Internet hosts by then.

Question26: Why does ATM use small, fixed-length cells?

Answer: Small,fixed-length cells can be routed through switches quickly,and completely in hardware. Small, fixed-size cells also make it easier to build hardware that handles many cells in parallel.Also,they do not block transmission lines for very long,making it easier to provide quality-of-service guarantees.

Question:28 An image is 1024 x 768 pixels with 3 bytes/pixel. Assume the image is uncompressed. How long does it take to transmit it over a 56-kbps modem channel? Over a 1-Mbps cable modem? Over a 10-Mbps Ethernet? Over 100-Mbps Ethernet?

Answer: The image is 1024×768×3 bytes or 2,359,296 bytes.This is 18,874,368 bits.At 56,000 bits/sec,it takes about 337.042 sec.At 1,000,000 bits/sec,it takes about 18.874 sec.At 10,000,000 bits/sec,it takes about 1.887 sec.At 100,000,000 bits/sec,it takes about 0.189 sec.

Question33: Make a list of activities that you do every day in which computer networks are used. How would your life be altered if these networks were suddenly switched off?

Answer: Omitted

Question36: Go to IETF's Web site, www.ietf.org, to see what they are doing. Pick a project you like and write a half-page report on the problem and the proposed solution.

Answer: Omitted

Computer Networks
Problems and solutions

Chapter two

Communication & Computer Network key-Lab of GD
School of Computer Science & Engineering
South China University of Technology

Questions:

1、2、3、4、5、7、9、18、22、23、24、26、28、29、30、31、33、34、37、40、42、43、50、51、53、57

Question1: Compute the Fourier coefficients for the function f(t) = t (0<=t<= 1).

Answer: $a_n = \frac{-1}{\pi n}, b_n = 0, c = 1.$

Question2: A noiseless 4-kHz channel is sampled every 1 msec. What is the maximum data rate?

Answer: A noiseless channel can carry an arbitrarily large amount of information,no matter how often it is sampled.Just send a lot of data per sample.For the 4kHz channel,make 8000 samples/sec.If each sample is 16 bits,the channel can send 128 kbps.If each sample is 1024 bits,the channel can send 8.2 Mbps.The key word here is‘‘noiseless.’’With a normal 4 kHz channel,the Shannon limit would not allow this.

Question3: Television channels are 6 MHz wide. How many bits/sec can be sent if four-level digital signals are used? Assume a noiseless channel.

Answer: Using the Nyquist theorem, we can sample 12 million times/sec. Four-level signals provide 2 bits per sample, for a total data rate of 24 Mbps.

Question4: If a binary signal is sent over a 3-kHz channel whose signal-to-noise ratio is 20 dB, what is the maximum achievable data rate?

Answer: A signal-to-noise ratio of 20 dB means S/N=100.Since log2101 is about 6.658,the Shannon limit is about 19.975 kbps.The Nyquist limit is 6 kbps. The bottleneck is therefore the Nyquist limit,giving a maximum channel capacity of 6 kbps.

Question5: What signal-to-noise ratio is needed to put a T1 carrier on a 50-kHz line?

Answer: To send a T1 signal we need Hlog2 (1+S/N)=1.544×106 with H=50,000.This yields S/N=230?1,which is about 93 dB.

Question7: How much bandwidth is there in 0.1 micron of spectrum at a wavelength of 1 micron?

Answer:

$$f = \frac{c}{\lambda} \qquad \frac{df}{d\lambda} = \frac{c}{\lambda^2}$$

$$df = \frac{c}{\lambda^2} d\lambda \qquad \Delta f = \frac{c}{\lambda^2} \Delta\lambda$$

$$c = 3 \times 10^8 \qquad \lambda = 10^{-6} m$$

$$\Delta\lambda = 0.1 \times 10^{-6} = 10^{-7} m$$

$$\Delta f = \frac{3 \times 10^8}{(10^{-6})^2} \times 10^{-7} = 30 \times 10^{12} Hz = 30THz$$

Question9: Is the Nyquist theorem true for optical fiber or only for copper wire?

Answer: The Nyquist theorem is a property of mathematics and has nothing to do with technology.It says that if you have a function whose Fourier spectrum does not contain any sines or cosines above f,then by sampling the function at a frequency of 2f you capture all the information there is.Thus,the Nyquist

theorem is true for all media.

Question18: A simple telephone system consists of two end offices and a single toll office to which each end office is connected by a 1-MHz full-duplex trunk. The average telephone is used to make four calls per 8-hour workday. The mean call duration is 6 min. Ten percent of the calls are long-distance (i.e., pass through the toll office). What is the maximum number of telephones an end office can support? (Assume 4 kHz per circuit.)

Answer: Each telephone makes 0.5 calls/hour at 6 minutes each.Thus,a telephone occupies a circuit for 3 minutes/hour.Twenty telephones can share a circuit, although having the load be close to 100%(ρ=1 in queueing terms)implies very long wait times).Since 10%of the calls are long distance,it takes 200 telephones to occupy a long-distance circuit full time.The interoffice trunk

Question22: A modem constellation diagram similar to Fig. 2-25 has data points at the following coordinates: (1, 1), (1, -1), (-1, 1), and (-1, -1). How many bps can a modem with these parameters achieve at 1200 baud?

Answer: There are four legal values per baud,so the bit rate is twice the baud rate.At 1200 baud,the data rate is 2400 bps.

Question23: A modem constellation diagram similar to Fig. 2-25 has data points at (0, 1) and (0, 2). Does the modem use phase modulation or amplitude modulation?

Answer: The phase shift is always 0,but two amplitudes are used,so this is straight amplitude modulation.

Question24: In a constellation diagram, all the points lie on a circle centered on the origin. What kind of modulation is being used?

Answer: If all the points are equidistant from the origin,they all have the same amplitude,so amplitude modulation is not being used.Frequency modulation is never used in constellation diagrams,so the encoding is pure phase shift keying.

Question26: An ADSL system using DMT allocates 3/4 of the available data channels to the downstream link. It uses QAM-64 modulation on each channel. What is the capacity of the downstream link?

Answer: There are 256 channels in all,minus 6 for POTS and 2 for control,leaving 248 for data.If 3/4 of these are for downstream,that gives 186 channels for downstream.ADSL modulation is at 4000 baud,so with QAM-64(6 bits/baud)we have 24,000 bps in each of the 186 channels.The total bandwidth is then 4.464 Mbps downstream.

Question28: Ten signals, each requiring 4000 Hz, are multiplexed on to a single channel using FDM. How much minimum bandwidth is required for the multiplexed channel? Assume that the guard bands are 400 Hz wide.

Answer: There are ten 4000 Hz signals.We need nine guard bands to avoid any interference.The minimum bandwidth required is 4000×10+400×9=43,600 Hz.

Question29: Why has the PCM sampling time been set at 125 μsec?

Answer: A sampling time of 125μsec corresponds to 8000 samples per second. According to the Nyquist theorem,this is the sampling frequency needed to capture all the information in a 4 kHz channel,such as a telephone channel. (Actually the nominal bandwidth is somewhat less,but the cutoff is not sharp.)

Question30: What is the percent overhead on a T1 carrier; that is, what percent of the 1.544 Mbps are not delivered to the end user?

Answer: The end users get 7×24=168 of the 193 bits in a frame.The overhead is therefore 25/193=13%.

Question31: Compare the maximum data rate of a noiseless 4-kHz channel using

(a) Analog encoding (e.g., QPSK) with 2 bits per sample.

(b) The T1 PCM system.

Answer: In both cases 8000 samples/sec are possible.With dibit encoding,two bits are sent per sample.With T1,7 bits are sent per period.The respective data rates are 16 kbps and 56 kbps.

Question33: What is the difference, if any, between the demodulator part of a modem and the coder part of a codec? (After all, both convert analog signals to digital ones.)

Answer: A coder accepts an arbitrary analog signal and generates a digital signal from it.A demodulator accepts a modulated sine wave only and generates a digital signal.

Question34: A signal is transmitted digitally over a 4-kHz noiseless channel with one sample every 125 μsec. How many bits per second are actually sent for each of these encoding methods?

(a) CCITT 2.048 Mbps standard.

(b) DPCM with a 4-bit relative signal value.

(c) Delta modulation.

Answer: (a)64 kbps.(b)32 kbps.(c)8 kbps.

Question37: In Fig. 2-37, the user data rate for OC-3 is stated to be 148.608 Mbps. Show how this number can be derived from the SONET OC-3 parameters.

Answer: Of the 90 columns,86 are available for user data in OC-1.Thus,the user capacity is 86×9=774 bytes/frame.With 8 bits/byte,8000 frames/sec,and 3 OC-1 carriers multiplexed together,the total user capacity is 3×774×8×8000,or 148.608 Mbps.

Question40: What is the available user bandwidth in an OC-12c connection?

Answer: The OC-12c frames are 12×90=1080 columns of 9 rows.Of these, 12×3=36 columns are taken up by line and section overhead.This leaves an SPE of 1044 columns.One SPE column is taken up by path overhead, leaving 1043 columns for user data.Since each column holds 9 bytes of 8 bits,an OC-12c frame holds 75,096 user data bits.With 8000 frames/sec,the user data rate is 600.768 Mbps.

Question42: Compare the delay in sending an x-bit message over a k-hop path in a circuit-switched network and in a (lightly loaded) packet-switched network. The circuit setup time is s sec, the propagation delay is d sec per hop, the packet size is p bits, and the data rate is b bps. Under what conditions does the packet network have a lower delay?

Answer: With circuit switching,at t=s the circuit is set up;at t=s+x/b the last bit is sent;at t=s+x/b+kd the message arrives.With packet switching,the last bit is sent at t=x/b.To get to the final destination,the last packet must be retransmitted k?1 times by intermediate routers,each retransmission taking p/b sec,so the total delay is x/b+(k?1)p/b+kd.Packet switching is faster if s>(k?1)p/b.

Question43: Suppose that x bits of user data are to be transmitted over a k-hop path in a packet-switched network as a series of packets, each containing p data bits and h header bits, with x >>p + h. The bit rate of the lines is b bps and the propagation delay is negligible. What value of p minimizes the total delay?

Answer: The total number of packets needed is x/p,so the total data+header traffic is (p+h)x/p bits.The source requires(p+h)x/pb sec to transmit these bits. The retransmissions of the last packet by the intermediate routers take up a total of(k?1)(p+h)/b sec.Adding up the time for the source to send all the bits,plus the time for the routers to carry the last packet to the destination, thus clearing the pipeline,we get a total time of(p+h)x/pb+(p+h)(k?1)/b sec.Minimizing this quantity with respect to p,we find p= √ hx/(k?1).

Question50: Suppose that A, B, and C are simultaneously transmitting 0 bits, using a CDMA system with the chip sequences of Fig. 2-45(b). What is the resulting chip sequence?

Answer: The result is obtained by negating each of A,B,and C and then adding the three chip sequences.Alternatively the three can be added and then negated. The result is(+3+1+1?1?3?1?1+1).

Question51: In the discussion about orthogonality of CDMA chip sequences, it was stated

that if S•T = 0 then is also 0. Prove this.

Answer: By definition

$$\mathbf{S} \bullet \mathbf{T} \equiv \frac{1}{m} \sum_{i=1}^{m} S_i T_i$$

If T sends a 0 bit instead of 1 bit,its chip sequence is negated,with the i-th element becoming?Ti.Thus,

$$\mathbf{S} \bullet \mathbf{T} \equiv \frac{1}{m} \sum_{i=1}^{m} S_i(-T_i) = -\frac{1}{m} \sum_{i=1}^{m} S_i T_i = 0$$

Question53: A CDMA receiver gets the following chips: (-1 +1 -3 +1 -1 -3 +1 +1). Assuming the chip sequences defined in Fig. 2-45(b), which stations transmitted, and which bits did each one send?

Answer: Just compute the four normalized inner products:

˜(1 +1˜3 +1˜1˜3 +1 +1) *˜(1˜1˜1 +1 +1˜1 +1 +1)/8 = 1

˜(1 +1˜3 +1˜1˜3 +1 +1) * ˜(1˜1 +1˜1 +1 +1 +1˜1)/8 = 1

˜(1 +1˜3 +1˜1˜3 +1 +1) * ˜(1 +1˜1 +1 +1 +1˜1˜1)/8 = 0

˜(1 +1˜ 3 +1˜ 1˜ 3 +1 +1) * ˜(1 +1˜ 1˜ 1˜ 1 +1˜ 1)/8 = 1

Question57: Using the spectral allocation shown in Fig. 2-48 and the information given in the text, how many Mbps does a cable system allocate to upstream and how many to downstream?

Answer: The upstream bandwidth is 37 MHz.Using QPSK with 2 bits/Hz,we get 74 Mbps upstream.Downstream we have 200 MHz.Using QAM-64,this is 1200 Mbps.Using QAM-256,this is 1600 Mbps.

Computer Networks

Problems and solutions

Chapter three

Communication & Computer Network key-Lab of GD

School of Computer Science & Engineering

South China University of Technology

Questions:

2、3、4、5、6、8、9、10、11、12、14、15、16、17、18、21、22、23、24、29、31、36

Question2: The following character encoding is used in a data link protocol: A: 01000111; B: 11100011; FLAG: 01111110; ESC: 11100000 Show the bit sequence transmitted (in binary) for the four-character frame: A B ESC FLAG when each of the following framing methods are used:

(a) Character count.

(b) Flag bytes with byte stuffing.

(c) Starting and ending flag bytes, with bit stuffing.

Answer: The solution is

(a)00000100 01000111 11100011 11100000 01111110

(b)01111110 01000111 11100011 11100000 11100000 11100000 01111110 01111110

(c)01111110 01000111 110100011 111000000 011111010 01111110

Question3: The following data fragment occurs in the middle of a data stream for which the byte-stuffing algorithm described in the text is used: A B ESC C ESC FLAG FLAG D. What is the output after stuffing?

Answer: After stuffing,we get A B ESC ESC C ESC ESC ESC FLAG ESC FLAG D.

Question4: One of your classmates, Scrooge, has pointed out that it is wasteful to end each frame with a flag byte and then begin the next one with a second flag byte. One flag byte could do the job as well, and a byte saved is a byte earned. Do you agree?

Answer: If you could always count on an endless stream of frames,one flag byte might be enough.But what if a frame ends(with a flag byte)and there are no new frames for 15 minutes.How will the receiver know that the next byte is actually the start of a new frame and not just noise on the line?The protocol is much simpler with starting and ending flag bytes.

Question5: A bit string, 0111101111101111110, needs to be transmitted at the data link layer. What is the string actually transmitted after bit stuffing?

Answer: The output is 011110111110011111010.

Question6: When bit stuffing is used, is it possible for the loss, insertion, or modification of a single bit to cause an error not detected by the checksum? If not, why not? If so, how? Does the checksum length play a role here?

Answer:

Question8: To provide more reliability than a single parity bit can give, an error-detecting coding scheme uses one parity bit for checking all the odd-numbered bits and a second parity bit for all the even-numbered bits. What is the Hamming distance of this code?

Answer: Making one change to any valid character cannot generate another valid character due to the nature of parity bits.Making two changes to even bits or twochanges to odd bits will give another valid character,so the distance is 2.

Question9: Sixteen-bit messages are transmitted using a Hamming code. How many check bits are needed to ensure that the receiver can detect and correct single bit errors? Show the bit pattern transmitted for the message 1101001100110101. Assume that even parity is used in the Hamming code.

Answer: Parity bits are needed at positions 1,2,4,8,and 16,so messages that do not extend beyond bit 31(including the parity bits)fit.Thus,five parity bits are sufficient.The bit pattern transmitted is 011010110011001110101

Question10: An 8-bit byte with binary value 10101111 is to be encoded using an even-parity Hamming code. What is the binary value after encoding?

Answer: The encoded value is 101001001111.

Question11: A 12-bit Hamming code whose hexadecimal value is 0xE4F arrives at a receiver. What was the original value in hexadecimal? Assume that not more than 1 bit is in error.

Answer: If we number the bits from left to right starting at bit 1,in this example,bit 2 (a parity bit)is incorrect.The 12-bit value transmitted(after Hamming encoding)was 0xA4F.The original 8-bit data value was 0xAF.

Question12: One way of detecting errors is to transmit data as a block of n rows of k bits per row and adding parity bits to each row and each column. The lower-right corner is a parity bit that checks its row and its column. Will this scheme detect all single errors? Double errors? Triple errors?

Answer: A single error will cause both the horizontal and vertical parity checks to be wrong.Two errors will also be easily detected.If they are in different rows, the row parity will catch them.If they are in the same row,the column parity will catch them.Three errors might slip by undetected,for example,if some bit is inverted along with its row and column parity bits.Even the corner bit will not catch this.

Question14: What is the remainder obtained by dividing x7 + x5 + 1 by the generator polynomial x3 + 1?

Answer: The remainder is x2 +x+1.

Question15: A bit stream 10011101 is transmitted using the standard CRC method described in the text. The generator polynomial is x3 + 1. Show the actual bit string transmitted. Suppose the third bit from the left is inverted during transmission. Show that this error is detected at the receiver's end.

Answer: The frame is 10011101.The generator is 1001.The message after appending three zeros is 10011101000.The remainder on dividing 10011101000 by 1001 is 100.So,the actual bit string transmitted is 10011101100.The received bit stream with an error in the third bit from the left is 10111101100. Dividing this by 1001 produces a remainder 100,which is different from zero. Thus,the receiver detects the error and can ask for a retransmission.

Question16: Data link protocols almost always put the CRC in a trailer rather than in a header. Why?

Answer: The CRC is computed during transmission and appended to the output streamas soon as the last bit goes out onto the wire.If the CRC were in the header, it would be necessary to make a pass over the frame to compute the CRC before transmitting.This would require each byte to be handled twice—once for checksumming and once for transmitting.Using the trailer cuts the work in half.

Question17: A channel has a bit rate of 4 kbps and a propagation delay of 20 msec. For what range of frame sizes does stop-and-wait give an efficiency of at least 50 percent?

Answer: Efficiency will be 50%when the time to transmit the frame equals the round-trip propagation delay.At a transmission rate of 4 bits/ms,160 bits takes 40 ms.For frame sizes above 160 bits,stop-and-wait is reasonably efficient.

Question18: A 3000-km-long T1 trunk is used to transmit 64-byte frames using protocol 5. If the propagation speed is 6 μsec/km, how many bits should the sequence numbers be?

Answer: To operate efficiently,the sequence space(actually,the send window size) must be large enough to allow the transmitter to keep transmitting until the first acknowledgement has been received.The propagation time is 18 ms.At T1 speed,which is 1.536 Mbps(excluding the 1 header bit),a 64-byte frame takes 0.300 msec.Therefore,the first frame fully arrives 18.3 msec after its transmission was started.The acknowledgement takes another 18 msec to get back,plus a small(negligible)time for the acknowledgement to arrive fully. In all,this time is 36.3 msec.The transmitter must have enough window space to keep going for 36.3 msec.A frame takes 0.3 ms,so it takes 121 frames to fill the pipe.Seven-bit sequence numbers are needed.

Question21: If the procedure between in protocol 5 checked for the condition a <=b<= c instead of the condition a <=b < c, would that have any effect on the protocol's correctness or efficiency? Explain your answer.

Answer: The protocol would be incorrect.Suppose that 3-bit sequence numbers are in use.Consider the following scenario:

A just sent frame 7.

B gets the frame and sends a piggybacked ACK.

A gets the ACK and sends frames 0–6,all of which get lost.

B times out and retransmits its current frame,with the ACK 7.

Look at the situation at A when the frame with r.ack=7 arrives.The key variables are AckExpected=0,r.ack=7,and NextFrameToSend=7.The modified between would return true,causing A to think the lost frames were being acknowledged.

Question22: In protocol 6, when a data frame arrives, a check is made to see if the sequence number differs from the one expected and no_nak is true. If both conditions hold, a NAK is sent. Otherwise, the auxiliary timer is started. Suppose that the else clause were omitted. Would this change affect the protocol's correctness?

Answer: Yes.It might lead to deadlock.Suppose that a batch of frames arrived correctly and were accepted.Then the receiver would advance its window. Now suppose that all the acknowledgements were lost.The sender would eventually time out and send the first frame again.The receiver would send a NAK.Suppose that this were lost.From that point on,the sender would keep timing out and sending a frame that had already been accepted,but the receiver would just ignore it.Setting the auxiliary timer results in a correct acknowledgement being sent back eventually instead,which resynchronizes.

Question23: Suppose that the three-statement while loop near the end of protocol 6 were removed from the code. Would this affect the correctness of the protocol or just the performance? Explain your answer.

Answer: It would lead to deadlock because this is the only place that incoming acknowledgements are processed.Without this code,the sender would keep timing out and never make any progress.

Question24: Suppose that the case for checksum errors were removed from the switch statement of protocol 6. How would this change affect the operation of the protocol?

Answer: It would defeat the purpose of having NAKs,so we would have to fall back to timeouts.Although the performance would degrade,the correctness would not be affected.The NAKs are not essential.

Question29: Frames of 1000 bits are sent over a 1-Mbps channel using a geostationary satellite whose propagation time from the earth is 270 msec. Acknowledgements are always piggybacked onto data frames. The headers are very short. Three-bit sequence numbers are used. What is the maximum achievable channel utilization for

(a) Stop-and-wait.

(b) Protocol 5.

(c) Protocol 6.

Answer: Let t=0 denote the start of transmission.At t=1 msec,the first frame has been fully transmitted.At t=271 msec,the first frame has fully arrived.At t=272 msec,the frame acknowledging the first one has been fully sent.At t=542 msec,the acknowledgement-bearing frame has fully arrived.Thus, the cycle is 542 msec.A total of k frames are sent in 542 msec,for an efficiency of k/542.Hence

(a)k=1,efficiency=1/542=0.18%

(b)k=7,efficiency=7/542=1.29%

(c)k=4,efficiency=4/542=0.74%

Question31: Consider an error-free 64-kbps satellite channel used to send 512-byte data frames in one direction, with very short acknowledgements coming back the other way. What is the maximum throughput for window sizes of 1, 7, 15, and 127? The earth-satellite propagation time is 270 msec.

Answer: The transmission starts at t=0.At t=4096/64000 sec=64 msec,the last bit is sent.At t=334 msec,the last bit arrives at the satellite and the very short ACK is sent.At t=604 msec,the ACK arrives at the earth.The data rate here is 4096 bits in 604 msec or about 6781 bps.With a window size of 7 frames,transmission time is 448 msec for the full window,at which time the sender has to stop.At 604 msec,the first ACK arrives and the cycle can start again.Here we have 7×4096=28,672 bits in 604 msec.The data rate is 47,470.2 bps.Continuous transmission can only occur if the transmitter is still sending when the first ACK gets back at t=604 msec.In other words,if the window size is greater than 604 msec worth of transmission,it can run at full speed.For a window size of 10 or greater,this condition is met,so for any window size of 10 or greater(e.g.,15 or 127),the data rate is 64 kbps.

Question36: PPP is based closely on HDLC, which uses bit stuffing to prevent accidental flag bytes within the payload from causing confusion. Give at least one reason why PPP uses byte stuffing instead.

Answer: PPP was clearly designed to be implemented in software,not in hardware as HDLC nearly always is.With a software implementation,working entirely with bytes is much simpler than working with individual bits.In addition, PPP was designed to be used with modems,and modems accept and transmit data in units of 1 byte,not 1 bit.

Computer Networks

Problems and Solutions

Chapter four

Communication & Computer Network key-Lab of GD

School of Computer Science & Engineering

South China University of Technology

Questions:

2、3、4、9、16、17、18、21、23、24、42、43

Question2: A group of N stations share a 56-kbps pure ALOHA channel. Each station outputs a 1000-bit frame on an average of once every 100 sec, even if the previous one has not yet been sent (e.g., the stations can buffer outgoing frames). What is the maximum value of N?

Answer: With pure ALOHA the usable bandwidth is 0.184×56 kbps=10.3 kbps. Each station requires 10 bps,so N=10300/10=1030 stations.

Question3: Consider the delay of pure ALOHA versus slotted ALOHA at low load. Which one is less? Explain your answer.

Answer: With pure ALOHA,transmission can start instantly.At low load,no collisions are expected so the transmission is likely to be successful.With slotted ALOHA,it has to wait for the next slot.This introduces half a slot time of delay.

Question4: Ten thousand airline reservation stations are competing for the use of a single slotted ALOHA channel. The average station makes 18 requests/hour. A slot is 125 µsec. What is the approximate total channel load?

Answer: Each terminal makes one request every 200 sec,for a total load of 50 requests/sec. Hence G=50/8000=1/160.

Question9: A LAN uses Mok and Ward's version of binary countdown. At a certain instant, the ten stations have the virtual station numbers 8, 2, 4, 5, 1, 7, 3, 6, 9, and 0. The next three stations to send are 4, 3, and 9, in that order. What are the new virtual station numbers after all three have finished their transmissions?

Answer: When station 4 sends,it becomes 0,and 1,2,and 3 are increased by 1.When station 3 sends,it becomes 0,and 0,1,and 2 are increased by 1.Finally, when station 9 sends,it becomes 0 and all the other stations are incremented by 1.The result is 9,1,2,6,4,8,5,7,0,and 3.

Question16: What is the baud rate of the standard 10-Mbps Ethernet?

Answer: The Ethernet uses Manchester encoding,which means it has two signal periods per bit sent.The data rate of the standard Ethernet is 10 Mbps,so the baud rate is twice that, or 20 megabaud.

Question17: Sketch the Manchester encoding for the bit stream: 0001110101.

Answer: The signal is a square wave with two values,high(H)and low(L).The pattern is LHLHLHHLHLHLLHHLLHHL.

Question18: Sketch the differential Manchester encoding for the bit stream of the previous problem. Assume the line is initially in the low state.

Answer: The pattern this time is HLHLHLLHHLLHLHHLHLLH.

Question21: Consider building a CSMA/CD network running at 1 Gbps over a 1-km cable with no repeaters. The signal speed in the cable is 200,000 km/sec. What is the minimum frame size?

Answer: For a 1-km cable,the one-way propagation time is 5μsec,so 2τ=10μsec. To make CSMA/CD work,it must be impossible to transmit an entire frame in this interval.At 1 Gbps,all frames shorter than 10,000 bits can be completely transmitted in under 10μsec,so the minimum frame is 10,000 bits or 1250 bytes.

Question23: Ethernet frames must be at least 64 bytes long to ensure that the transmitter is still going in the event of a collision at the far end of the cable. Fast Ethernet has the same 64-byte minimum frame size but can get the bits out ten times faster. How is it possible to maintain the same minimum frame size?

Answer: The maximum wire length in fast Ethernet is 1/10 as long as in Ethernet.

Question24: Some books quote the maximum size of an Ethernet frame as 1518 bytes instead of 1500 bytes. Are they wrong? Explain your answer.

Answer: The payload is 1500 bytes,but when the destination address,source address, type/length,and checksum fields are counted too,the total is indeed 1518.

Question42: Briefly describe the difference between store-and-forward and cut-through switches.

Answer: A store-and-forward switch stores each incoming frame in its entirety,then examines it and forwards it.A cut-through switch starts to forward incoming frames before they have arrived completely.As soon as the destination address is in,the forwarding can begin.

Question43: Store-and-forward switches have an advantage over cut-through switches with respect to damaged frames. Explain what it is.

Answer: Store-and-forward switches store entire frames before forwarding them.After a frame comes in,the checksum can be verified.If the frame is damaged,it is discarded immediately. With cut=through,damaged frames cannotbe discarded by the switch because by the time the error is detected,the frame is already gone.Trying to deal with the problem is like locking the barn door after the horse has escaped.

Computer Networks

Problems and Solutions

Chapter five

Communication & Computer Network key-Lab of GD

School of Computer Science & Engineering

South China University of Technology

Questions:

2、3、7、9、11、20、25、26、27、28、37、38、39、40、41、42、43、44、46、48、50、51、52、53、57

Question2: Are there any circumstances when connection-oriented service will (or at least should) deliver packets out of order? Explain.

Answer: Yes.Interrupt signals should skip ahead of data and be delivered out of equence.A typical example occurs when a terminal user hits the quit(kill) key.The packet generated from the quit signal should be sent immediately and should skip ahead of any data currently queued up for the program,i.e.,data already typed in but not yet read.

Question3: Datagram subnets route each packet as a separate unit, independent of all others. Virtual-circuit subnets do not have to do this, since each data packet follows a predetermined route. Does this observation mean that virtual-circuit subnets do not need the capability to route isolated packets from an arbitrary source to an arbitrary destination? Explain your answer.

Answer: Virtual circuit networks most certainly need this capability in order to route connection setup packets from an arbitrary source to an arbitrary destination.

Question7: Consider the network of Fig. 5-7, but ignore the weights on the lines. Suppose that it uses flooding as the routing algorithm. If a packet sent by A to D has a maximum hop count of 3, list all the routes it will take. Also tell how many hops worth of bandwidth it consumes.

Answer: It will follow all of the following routes:ABCD,ABCF,ABEF,ABEG, AGHD,AGHF,AGEB,and AGEF.The number of hops used is 24.

Question9: Consider the subnet of Fig. 5-13(a). Distance vector routing is used, and the following vectors have just come in to router C: from B: (5, 0, 8, 12, 6, 2); from D: (16, 12, 6, 0, 9, 10); and from E: (7, 6, 3, 9, 0, 4). The measured delays to B, D, and E, are 6, 3, and 5, respectively. What is C's new routing table? Give both the outgoing line to use and the expected delay.

Answer: Going via B gives(11,6,14,18,12,8).Going via D gives(19,15,9,3,9,10).Going via E gives(12,11,8,14,5,9). Taking the minimum for each destination except C gives(11,6,0,3,5,8). The outgoing lines are(B,B,–,D,E,B).

Question11: In Fig. 5-14 the Boolean OR of the two sets of ACF bits are 111 in every row. Is this just an accident here, or does it hold for all subnets under all circumstances?

Answer: It always holds.If a packet has arrived on a line,it must be acknowledged.If no packet has arrived on a line,it must be sent there.The cases 00(has not arrived and will not be sent)and 11(has arrived and will be sent back)are logically incorrect and thus do not exist.

Question20: Consider the Chord circle of Fig. 5-24. Suppose that node 10 suddenly goes on line. Does this affect node 1's finger table, and if so, how?

Answer: The node in entry 3 switches from 12 to 10.

Question25: The byte-counting variant of the leaky bucket algorithm is used in a particular system. The rule is that one 1024-byte packet, or two 512-byte packets, etc., may be sent on each tick. Give a serious restriction of this system that was not mentioned in the text.

Answer: It is not possible to send any packets greater than 1024 bytes,ever.

Question26: An ATM network uses a token bucket scheme for traffic shaping. A new token is put into the bucket every 5 μsec. Each token is good for one cell, which contains 48 bytes of data. What is the maximum sustainable data rate?

Answer: With a token every 5μsec,200,000 cells/sec can be sent.Each cell holds 48 data bytes or 384 bits.The net data rate is then 76.8 Mbps.

Question27: A computer on a 6-Mbps network is regulated by a token bucket. The token bucket is filled at a rate of 1 Mbps. It is initially filled to capacity with 8 megabits. How long can the computer transmit at the full 6 Mbps?

Answer: The naive answer says that at 6 Mbps it takes 4/3 sec to drain an 8 megabit bucket.However,this answer is wrong,because during that interval,more tokens arrive.The correct answer can be obtained by using the formula S=C/(M?ρ).Substituting,we get S=8/(6?1)or 1.6 sec.

Question28: Imagine a flow specification that has a maximum packet size of 1000 bytes, a token bucket rate of 10 million bytes/sec, a token bucket size of 1 million bytes, and a maximum transmission rate of 50 million bytes/sec. How long can a burst at maximum speed last?

Answer: Call the length of the maximum burst interval?t.In the extreme case,the bucket is full at the start of the interval(1 Mbyte)and another 10?t Mbytes come in during the interval.The output during the transmission burst contains 50?t Mbytes.Equating these two quantities,we get 1+10?t=50?t. Solving this equation,we get?t is 25 msec.

Question37: Suppose that instead of using 16 bits for the network part of a class B address originally, 20 bits had been used. How many class B networks would there have been?

Answer: With a 2-bit prefix,there would have been 18 bits left over to indicate the network.Consequently,the number of networks would have been 218 or 262,144.However,all 0s and all 1s are special,so only 262,142 are available.

Question38: Convert the IP address whose hexadecimal representation is C22F1582 to dotted decimal notation.

Answer: The address is 194.47.21.130.

Question39: A network on the Internet has a subnet mask of 255.255.240.0. What is the maximum number of hosts it can handle?

Answer: The mask is 20 bits long,so the network part is 20 bits.The remaining 12 bits are for the host,so 4096 host addresses exist.

Question40: A large number of consecutive IP address are available starting at 198.16.0.0. Suppose that four organizations, A, B, C, and D, request 4000, 2000, 4000, and 8000 addresses, respectively, and in that order. For each of these, give the first IP address assigned, the last IP address assigned, and the mask in the w.x.y.z/s notation.

Answer: To start with,all the requests are rounded up to a power of two.The starting address,ending address,and mask are as follows:

A:198.16.0.0–198.16.15.255 written as 198.16.0.0/20

B:198.16.16.0–198.23.15.255 written as 198.16.16.0/21

C:198.16.32.0–198.47.15.255 written as 198.16.32.0/20

D:198.16.64.0–198.95.15.255 written as 198.16.64.0/19

Question41: A router has just received the following new IP addresses: 57.6.96.0/21, 57.6.104.0/21, 57.6.112.0/21, and 57.6.120.0/21. If all of them use the same outgoing line, can they be aggregated? If so, to what? If not, why not?

Answer: They can be aggregated to 57.6.96/19.

Question42: The set of IP addresses from 29.18.0.0 to 19.18.128.255 has been aggregated to 29.18.0.0/17. However, there is a gap of 1024 unassigned addresses from 29.18.60.0 to 29.18.63.255 that are now suddenly assigned to a host using a different outgoing line. Is it now necessary to split up the aggregate address into its constituent blocks, add the new block to the table, and then see if any reaggregation is possible? If not, what can be done instead?

Answer: It is sufficient to add one new table entry:29.18.0.0/22 for the new block.If an incoming packet matches both 29.18.0.0/17 and 29.18.0.0./22,the longest one wins.This rule makes it possible to assign a large block to one outgoing line but make an exception for one or more small blocks within its range.

Question43: A router has the following (CIDR) entries in its routing table:

Address/mask	Next hop
135.46.56.0/22	Interface 0
135.46.60.0/22	Interface 1
192.53.40.0/23	Router 1
default	Router 2

For each of the following IP addresses, what does the router do if a packet with that address arrives?

(a) 135.46.63.10

(b) 135.46.57.14

(c) 135.46.52.2

(d) 192.53.40.7

(e) 192.53.56.7

Answer: The packets are routed as follows:

(a)Interface 1

(b)Interface 0

(c)Router 2

(d)Router 1

(e)Router 2

Question44: Many companies have a policy of having two (or more) routers connecting the company to the Internet to provide some redundancy in case one of them goes down. Is this policy still possible with NAT? Explain your answer.

Answer: After NAT is installed,it is crucial that all the packets pertaining to a single connection pass in and out of the company via the same router,since that is where the mapping is kept.If each router has its own IP address and all traffic belonging to a given connection can be sent to the same router,the mapping can be done correctly and multihoming with NAT can be made to bwork.

Question46: ARP and RARP both map addresses from one space to another. In this respect, they are similar. However, their implementations are fundamentally different. In what major way do they differ?

Answer: RARP has a RARP server that answers requests.ARP does not have this. The hosts themselves answer ARP queries.

Question48: Most IP datagram reassembly algorithms have a timer to avoid having a lost fragment tie up reassembly buffers forever. Suppose that a datagram is fragmented into four fragments. The first three fragments arrive, but the last one is delayed. Eventually, the timer goes off and the three fragments in the receiver's memory are discarded. A little later, the last fragment stumbles in. What should be done with it?

Answer: As far as the receiver is concerned,this is a part of new datagram,since no other parts of it are known.It will therefore be queued until the rest show up. If they do not,this one will time out too.

‘

Question50: A person who lives in Boston travels to Minneapolis, taking her portable computer with her. To her surprise, the LAN at her destination in Minneapolis is a wireless IP LAN, so she does not have to plug in. Is it still necessary to go through the entire business with home agents and foreign agents to make e-mail and other traffic arrive correctly?

Answer: Yes.The fact that the Minneapolis LAN is wireless does not cause the packets that arrive for her in Boston to suddenly jump to Minneapolis.The home agent in Boston must tunnel them to the foreign agent on the wireless LAN in Minneapolis.The best way to think of this situation is that the user has plugged into the Minneapolis LAN,the same way all the other Minneapolis users have.That the connection uses radio instead of a wire is irrelevant.

Question51: IPv6 uses 16-byte addresses. If a block of 1 million addresses is allocated every picosecond, how long will the addresses last?

Answer: With 16 bytes there are 2^{128}or 3.4×10^{38} addresses.If we allocate them at arate of 10^{18}per second,they will last for 10^{13}years.This number is 1000times the age of the universe.Of course,the address space is not flat,so they are not allocated linearly,but this calculation shows that even with an allocation scheme that has an efficiency of 1/1000(0.1 percent),one will never run out.

Question52: The Protocol field used in the IPv4 header is not present in the fixed IPv6 header. Why not?

Answer: The Protocol field tells the destination host which protocol handler to give the IP packet to.Intermediate routers do not need this information,so it is not needed in the main header.Actually,it is there,but disguised.The Next header field of the last(extension)header is used for this purpose.

Question53: When the IPv6 protocol is introduced, does the ARP protocol have to be changed? If so, are the changes conceptual or technical?

Answer: Conceptually,there are no changes.Technically,the IP addresses requested are now bigger,so bigger fields are needed.

Question57: Use the traceroute (UNIX) or tracert (Windows) programs to trace the route from your computer to various universities on other continents. Make a list of transoceanic links you have discovered. Some sites to try are

www.berkeley.edu (California)
www.mit.edu (Massachusetts)
www.vu.nl (Amsterdam)
www.ucl.ac.uk (London)
www.usyd.edu.au (Sydney)
www.u-tokyo.ac.jp (Tokyo)
www.uct.ac.za (Cape Town)

Answer: Omitted

Computer Network

Problems and Solutionss

Chapter six

Communication & Computer Network key-Lab of GD

School of Computer Science & Engineering

South China University of Technology

Questions:

6、14、25、28、29

Question6: Imagine that a two-way handshake rather than a three-way handshake were used to set up connections. In other words, the third message was not required. Are deadlocks now possible? Give an example or show that none exist.

Answer: Deadlocks are possible. For example, a packet arrives at A out of the blue, and A acknowledges it. The acknowledgement gets lost, but A is now open while B knows nothing at all about what has happened. Now the same thing happens to Band both are open, but expecting different sequence numbers. Timeouts have to be introduced to avoid the deadlocks.

Question14: Why does UDP exist? Would it not have been enough to just let user processes send raw IP packets?

Answer: No. IP packets contain IP addresses, which specify a destination machine. Once such a packet arrived, how would the network handler know which process to give it to? UDP packets contain a destination port. This information is essential so they can be delivered to the correct process.

Question25: The maximum payload of a TCP segment is 65,495 bytes. Why was such a strange number chosen?

Answer: The entire TCP segment must fit in the 65,515-byte payload field of an IP packet.Since the TCP header is a minimum of 20 bytes,only 65,495 bytes are left for TCP data.

Question28: Consider the effect of using slow start on a line with a 10-msec round-trip time and no congestion. The receive window is 24 KB and the maximum segment size is 2 KB. How long does it take before the first full window can be sent?

Answer: The first bursts contain 2K,4K,8K,and 16K bytes, respectively. The next one is 24 KB and occurs after 40 msec.

Question29: Suppose that the TCP congestion window is set to 18 KB and a timeout occurs. How big will the window be if the next four transmission bursts are all successful? Assume that the maximum segment size is 1 KB.

Answer: The next transmission will be 1 maximum segment size. Then 2,4,and 8.So after four successes, it will be 8 KB.