Computer Networks
End-Semester Examination

Time: 2 hours **Full Marks: TBA**

(1) Answer the following questions briefly (within 20 words): [1 x 10 = 10]

(a) Define DC component in a signal.

If a signal includes a component of zero frequency, that component is a direct current (dc) component.

(b) A network on the Internet has a subnet mask of 255.255.240.0. What is the maximum number of hosts that it can accommodate?

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 can be accommodated.

(c) What is a peer-to-peer network?

A peer-to-peer network consists of a large connection of computers, without central control where typically each node has some information of interest.

(d) What is attenuation?

Attenuation is the reduction of signal strength at higher frequencies.

(e) In an (n, k) block ECC, what do n and k represent?

A (n, k) block ECC encodes *k* data bits into *n*-bit codewords.

(f) What is *p* in p-persistent CSMA?

In p-persistent CSMA, *p* is the probability that a station transmits when the medium is idle.

(g) Convert the IP address whose hexadecimal representation is C22F1582 to dotted decimal notation.

C2.2F.15.82. Convert each segment from hexadecimal to decimal. Answer is **194.47.21.130**

(h) What is the function of an Internet Daemon (inetd)?

An internet daemon attaches itself to multiple ports and waits for the first connection request, then forks to that service.

(i) Give one scenario where the urgent flag in the TCP header is used.

When a user presses Ctrl+C to break-off a remote computation, the urgent flag is set.

(j) What is meant by SLA (Service Level Agreement)?

When a traffic flow is set up, the user and the network (i.e., the customer and the provider) agree on a certain traffic pattern (i.e., shape) for that flow, which is called the SLA.

(2) Answer the following questions briefly (within 30 words). Draw neat diagrams where applicable. [2 x 5 = 10]

(a) How is anycast routing different from multicast routing?

Some applications require sending the same message to multiple recipients who form a well-defined group. The routing algorithm used in such a case is called multicast routing.

In anycast routing, a packet is delivered only to the nearest member of a group.

(b) What does a IP pseudo-header contain and why is it included in the checksum?

The IP pseudo-header contains the 32-bit IPv4 addresses of the source and destination machines, the protocol number for the transport layer protocol (UDP/TCP), and the byte count for the segment (including the header).

Including the pseudo-header in the checksum computation helps detect mis-delivered packets with greater reliability.

(c) What is the advantage of sliding-window flow control compared to stop-and-wait flow control?

Sliding-window flow control is potentially much more efficient than stop-and-wait flow control. With sliding-window flow control, the transmission link is treated as a pipeline that may be filled with frames in transit. In contrast, with stop-and-wait flow control, only one frame may be in transit at a time.

(d) What is the difference between *strict source routing* and *loose source routing*?

The *Strict source routing* option gives the complete path from source to destination as a sequence of IP addresses. The datagram is required to follow that exact route.

The *Loose source routing* option requires the packet to traverse the list of routers specified, in the order specified, but it is allowed to pass through other routers on the way.

(e) What are the assigned ports for the following protocols: (i) SSH, (ii) HTTP? What assigned ports do the following protocols map to: (i) 443, (ii) 25?

(i) SSH – 22, (ii) HTTP – 80; (i) 443 – HTTPS, (ii) 25 – SMTP.

(3) Answer the following questions:

(a) Given the narrow (usable) audio bandwidth of a telephone transmission facility, a nominal SNR of 56dB (400,000), and a certain level of distortion: [2+2]

(i) What is the theoretical maximum channel capacity (kbps) of traditional telephone lines?

The theoretical maximum capacity is obtained using Shannon's formula:
$C = B \log_2 (1 + SNR) = 3000 \log_2 (1+400000) =$ **56 kbps**

(ii) What can we say about the actual maximum channel capacity?

The actual maximum channel capacity will be much lower due to:
(1) The formula assumes white noise (thermal noise).
(2) Impulse noise is not accounted for.
(3) Attenuation distortion or delay distortion (mentioned in the question) is also not accounted for.

(4) Even in an ideal environment, encoding issues, such as coding length and complexity exist.

(b) Given a receiver with an effective noise temperature of 294K and a 10-MHz bandwidth, what is the thermal noise level at the receiver output? [2]

Noise level N in decibel-watts = 10 log (k) + 10 log (T) + 10 log (B)
= - 228.6 + 10 log (294) + 10 log (10000000)
= - 228.6 + 24.7 + 70
= **- 133.9 dbW**

(c) Consider an audio signal with spectral components in the range 300 to 3000 Hz. Assume that a sampling rate of 7000 samples per second will be used to generate a PCM signal. [2+2]

(i) For SNR = 30 dB, what is the number of uniform quantization levels needed?

For PCM, $SNR_{dB} = 20 \log_2 n + 1.76$ dB = 6.02 n + 1.76 dB
=> 30 = 6.02 n + 1.76 => n = 4.69
Nearest higher integral value of n = 5.
Therefore, number of levels = 2^5 = **32**

(ii) What data rate is required?

Required data rate = 7000 samples/sec x 5 bits/sample = **35 kbps**

(4) Answer the following questions:

(a) Suppose that for an ISDN (integrated services digital network) with a 64-kbps channel, 1 frame with undetected error is expected per day. Assuming that the frame length is 1000 bits: [2+2]

(i) What is the probability that a frame is received with an undetected error?

The number of frames that can be transmitted in a day comes out to 5.529×10^6, which yields a frame error rate of $1 / (5.529 \times 10^6)$ = $\mathbf{0.18 \times 10^{-6}}$

(ii) If the actual bit error rate is 10^{-6}, is it possible to achieve a probability close to that achieved in (i)? Compute.

If BER = 10^{-6}, then P1 = $(0.999999)^{1000}$ = 0.999, and P2 = (1 – P1) = 0.001 = $\mathbf{10^{-3}}$. This value is 3 orders of magnitude larger than that obtained in (i).

(b) Two communicating devices are using a single-bit even parity check for error detection. The transmitter sends the byte 10101010 and, because of channel noise, the receiver gets the byte 10011010. Will the receiver detect the error? Why or why not? [1]

No. The number of 1's in both messages is the same and will produce the same parity bit.

(c) 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. [1+3]

As check bits are needed at position 1, 2, 4, 8, 16 to check a 16-bit message, 5 check bits are needed to ensure that the receiver can detect and correct single bit errors.

The bit pattern:

C1 | C2 | 1 | **C4 |** 101 | **C8 |** 0011001 | **C16 |** 10101

(Assume that '•' = XOR)

C1 = 1•1•1•0•1•0•1•1•1•1 = 0

C2 = 1•0•1•0•1•0•1•0•1 = 1

C4 = 1•0•1•1•0•0•1•0•1 = 1

C8 = 0•0•1•1•0•0•1 = 1

C16 =1•0•1•0•1 = 1

So, the message is transmitted as follows: 011110110011001110101

(d) Calculate the Hamming distances among the following pairs of code-words: [1]
(i) (00000, 10101) (iii) (01010, 10101)

Answers: (i) 3, (ii) 5

(5) Answer the following questions:

(a) The distance from earth to a distant planet is approximately 9×10^{10} m. [2+2]

(i) What is the channel utilization if a stop-and-wait protocol is used for frame transmission on a 64 Mbps point-to-point link? Assume that the frame size is 32 KB and the speed of light is 3×10^8 m/s.

Each period is: Transmission time + prop. time + ACK prop. time = Transmission time + RTT
Now, transmission time = 32KB/64Mbps = 0.004 seconds
RTT = 2 * ($9*10^{10}$m) / ($3*10^8$m/s) = 600 seconds
Utilization = 0.004 / (600 + 0.004) = $6.667 x 10^{-6}$ = $6.667 x 10^{-4}$ %

(ii) Suppose a sliding window protocol is used instead. For what send window size, will the link utilization be 100%? You may ignore the protocol processing times at the sender and the receiver.

To get to 100% utilization, one needs to keep transmitting packets while waiting for the ACKs, which means for the whole period.
Period = 600.004 seconds
1 Transmission time = 0.004 seconds
For non-stop transmission, one needs to transmit 600.004/0.004 packets each period, therefore 150001 should be the window size.

(b) A large population of ALOHA users manages to generate 50 requests/sec, including both originals and retransmissions. Assume that time is slotted in units of 40 msec. [1+1+1]

(i) What is the chance of success on the first attempt?

Frame transmission time (X) is given to be 40msec.
Total load in requests/s is given to be 50 requests/s implying that:
Total arrival rate G in frames/X seconds = 50 requests/s × X = 50 × (40×10^{-3}) = 2 requests/X seconds.
Prob. of successful transmission on the 1st attempt = P[0 frames arriving in the datalink layer in X seconds] = $e^{-G}G^k$ / k! (with k=0) = e^{-G} = 0.1353

(ii) What is the probability of exactly k collisions and then a success?

Prob. of successful tx. in (k + 1) attempts = P[collisions in k attempts] × P[success in k+1 attempts]
Based on (a), P[collisions in k attempts] = (1 − P[success in the 1st attempt])^k = $(1 - e^{-G})^k$ Hence, Prob. of successful tx. in (k + 1) attempts = $(1 - e^{-G})^k \times e^{-G} = 0.8647^k$ (0.1353)

(iii) What is the expected number of transmission attempts needed?

Expected number of transmissions = Σ k [Prob of success in k transmissions], which reduces to $e^G = 7.3891$

(c) Two CSMA/CD stations are each trying to transmit long (multi-frame) files. After each frame is sent, they contend for the channel, using the binary exponential back-off algorithm. What is the probability that the contention ends on round k, and what is the mean number of rounds per contention period? [1.5+1.5]

Attempt i is distributed among 2^{i-1} slots. Thus, the probability of a collision on attempt i is $1/2^{i-1}$. The prob. that first k attempts will fail is then:
$P_k = [1 - (1/2^{k-1})]\ \pi\ (1/2^{i-1}) = [1 - (1/2^{k-1})]\ .\ [1 / 2^{(k-1)(k-2)/2}]$

Mean number of rounds = expectation = Σ k. P_k

(6) Answer the following questions:

(a) Consider the following arrangement for implementation of the *Chord* protocol (larger circles represent the actual machines):

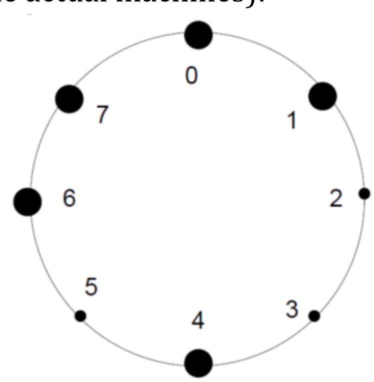

(i) Construct the finger table for the nodes 0, 1, 4, 6, 7 (assume that the node identifier consists of 3 bits). [2.5]

The finger tables are as follows:

Finger table at node 0

Start	Address of successor
1	1
2	4
4	4

Finger table at node 1

Start	Address of successor
2	4
3	4
5	6

Finger table at node 4

Start	Address of successor
5	6
6	6
0	0

Finger table at node 6

Start	Address of successor
7	7
0	0
2	4

Finger table at node 7

Start	Address of successor
0	0
1	1
3	4

(ii) Show the visited nodes and intermediate steps when one searches for key value 6. Assume that the search has started from node 0. [1.5]

Looking up key 6 at node 0:
(1) Nearest pred. to 6 is 4 in 0's table. Send to 4.
(2) Nearest pred. to 6 is 6 in 4's table. Send to 6.
(3) 6 knows that key 6 is with itself. Sends back its IP address to 0.

(b) State three problems of Network Address Translation (NAT). [3]

The following are the problems with NAT. Any 3 of the following are required:
(1) Architectural model of IP where each IP is identified by a single machine
(2) Internet from connection-less to connection-oriented
(3) NAT violates the most fundamental rule of protocol layering
(4) Processes on the Internet are not required to use TCP or UDP
(5) Some application insert IP address in their body – (FTP) – NAT doesn't know about it.

(c) Consider the network shown below. 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 cost of the links from C 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 cost. [3]

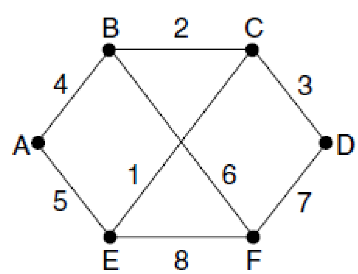

Going via B gives (11, 6, 14, 18, 12, 8).
Going via D gives (19, 15, 9, 3, 9, 10).
Going via E gives (12, 11, 8, 14, 5, 9).
Taking the minimum of each destination except C gives (11, 6, 0, 3, 5, 8) for outgoing links (B, B, --, D, E, B).

(7) Answer the following questions:

(a) What is a care-of-address in the case of routing for mobile hosts? Explain how packets are routed from a sender to a mobile host who has moved out from his home network. [1+2]

When a mobile host shows up at a foreign site, it obtains a new IP address at the foreign site, which is known as the care-of-address. The mobile tells the home agent where it is now by giving it its care-of address.

The steps for packets getting routed are as follows:
(1) Packet is sent to the mobile host's home address.
(2) Packet is tunnelled to the foreign agent.
(3) Sender is given foreign agent's address.
(4) Subsequent packets are tunnelled to the foreign agent.

(b) 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? [3]

S = C/(M-R) = 8/(6 – 1) = 1.6 seconds.

(c) The 14th byte (bit positions 105-112) of a TCP header consists of 8 flags of size 1-bit each. State any 4 of them and explain in one sentence each. [4]

The 8 flags are:
(1) CWR is set to signal Congestion Window Reduced from the TCP sender to the TCP receiver so that it knows the sender has slowed down and can stop sending the ECN-Echo.
(2) ECE is set to signal an ECN-Echo to a TCP sender to tell it to slow down when the TCP receiver gets a congestion indication from the network.
(3) URG is set to 1 if the Urgent pointer is in use. The Urgent pointer is used to indicate a byte offset from the current sequence number at which urgent data are to be found. This facility is in lieu of interrupt messages.
(4) The ACK bit is set to 1 to indicate that the Acknowledgement number is valid. If ACK is 0, the segment does not contain an acknowledgement, so the Acknowledgement number field is ignored.
(5) The PSH bit indicates PUSHed data. The receiver is requested to deliver the data to the application upon arrival and not buffer it until a full buffer has been received.
(6) The RST bit is used to abruptly reset a connection that has become confused due to a host crash or some other reason. It is also used to reject an invalid segment or refuse an attempt to open a connection.
(7) The SYN bit is used to establish connections. In essence, the SYN bit is used to denote both CONNECTION REQUEST and CONNECTION ACCEPTED, with the ACK bit used to distinguish between those two possibilities.
(8) The FIN bit is used to release a connection. It specifies that the sender has no more data to transmit.

(8) Answer the following questions:

(a) If the TCP round-trip time, RTT, is currently 30 msec and the following acknowledgements come in after 26, 32, and 24 msec, respectively, what is the new RTT estimate using Jacobson's algorithm for *Smoothed round-trip time* (SRTT) computation? Use $\alpha = 0.9$. [2]

Jacobson's algorithm: $SRTT = \alpha\ SRTT + (1 - \alpha)\ R$
The successive estimates are 29.6, 29.84, 29.256.

(b) 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. [2]

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

(c) Consider a telnet connection to an interactive editor that reacts on every keystroke. What would be the worst-case overhead of sending one character assuming TCP connection (consider only TCP and IP headers)? Briefly describe the algorithm which solves this problem. [2+2]

When a character arrives at the sending TCP entity, TCP creates a 21-byte TCP segment, which it gives to IP to send as a 41-byte IP datagram. At the receiving side, TCP immediately sends a 40-byte acknowledgement.. Later, when the editor has read the byte, TCP sends a window update, moving the window 1 byte to the right. This packet is also 40 bytes. Finally, when the editor has processed the character, it echoes the character as a 41-byte packet. In all, 162 bytes of bandwidth are used and four segments are sent for each character typed.

Solution: Nagle's algorithm.
When data come into the sender one byte at a time, just send the first byte and buffer all the rest until the outstanding byte is acknowledged. Then send all the buffered characters in one TCP segment and start buffering again until they are all acknowledged. The algorithm additionally allows a new packet to be sent if enough data have trickled in to fill half the window or a maximum segment.

(d) What are the 4 possible connection release scenarios in the transport layer? Support your answer with neat diagrams. [2]

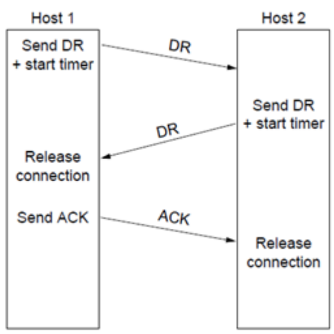

(1) Normal case of three-way handshake

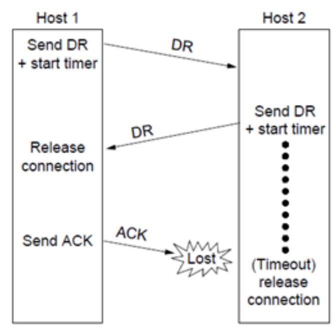

(2) Final ACK lost.

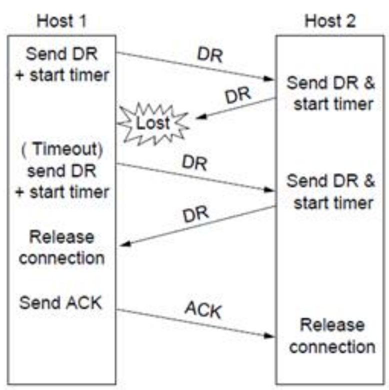

(3) Response lost

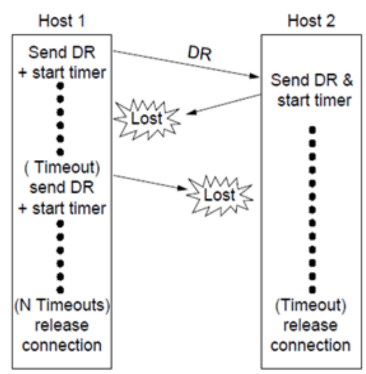

(4) Response lost and subsequent DRs lost.