Department of Computer Science and Engg, IIT Bombay

CS 348: Computer Networks Endsem: 19th Nov 2012, 14:00 – 17:00

Handwritten notes permitted. No photocopies. Weightage: 50%, Max Marks – 80

Q1: [15 Marks]

a) When using CRC, how many bits are in the checksum when using the generator polynomial "x^4+x+1"?
4 bits (if the generator is of degree k, then there are k bits appended to the message)

b) A channel has a bit-rate of 4Kbps and a propogation delay of 20msec. For what range of frame sizes does stop-and-wait ARQ give an efficiency of 50%?
Efficiency = Transmit/(RTT+transmit). So for Efficiency of 0.5, Transmit should be >= RTT.
Frame size = Transmit * bit-rate = RTT * bit-rate = 2 * delay * bit-rate = 160 bits.
(if some have assumed 1k= 1024 treat it as correct).

c) What is the network number, subnet number and host number for the address 135.104.192.100, mask 255.255.128.0?
Since it is a class B address, take first 16 bits as network number, so network number is 135.104. Since 128 in binary is 10000000, and 192 is 11000000, only 1 bit is used for subnet, so subnet id is 128, (some may have written 135.104.128.0, which is also correct). Host id is last 8 bits, so 100.

d) Let N be the number of nodes, E be the number of edges. For Link state routing, how many message exchanges (roughly, expressed as a function of the above parameters) need to be exchanged to build the topology at each and every node in the network?
For Link-State: Messages are of the order N * E (each node floods on every edge)

e) In public-key cryptography, suppose A is sending a message to B. How should the message be encrypted so that it is only understood by B, and B is sure that the message was indeed sent by A?
A encrypts with own private key for authentication. Resulting cipher text is encrypted again using B's public key, for confidentiality. B decrypts with own private key, then with A's public key.

Q2: [20 Marks] Explain the difference between:

a) p-Persistent and non-Persistent CSMA.
P-Persistent: Sense the channel. IF the channel is idle, THEN with probability p transmit and with probability (1-p) delay for one time slot and start over. IF the channel is busy, THEN delay one time-slot and start over.
non-Persistent: Sense the channel. IF the channel is idle, THEN transmit. If the channel is busy, THEN wait a random amount of time and start over.

b) Persist Timer and Retransmit Timer in TCP.
Sender starts the Persist Timer upon receipt of an ACK having Receiver's advertised window as 0. When the Timer expires, sender sends a 1-byte packet to probe if Advertised window has opened.
Sender starts the Retransmit Timer upon sending a packet.
When the Timer expires and ACK for the packet has not yet been received, sender retransmits.

c) Recursive and Iterated Queries in DNS.
See slides 14-16 in http://www.cse.iitb.ac.in/~sri/cs348/cs348-lec31-DNS-2012.pdf

d) Idempotent and non-Idempotent operations in RPC.
Idempotent: Has no additional effect if called more than once with the same input parameters. Used with at-least-once semantics, i.e., a request to perform an operation is guaranteed to happen at least once, but might also happen more than once. If the operation is idempotent, then there is no harm in performing the operation two or more times. Example: Remote calculator.
Non-Idempotent: Multiple executions will change the result each time. Example: Append to a list.

Q3: [5 Marks]

On an 10Mbps Ethernet, there is a rogue station that detects transmissions and disrupts them by beginning a competing transmission as soon as it hears the beginning of the transmitted frame, thereby causing a collision. Suppose your machine, which is on this Ethernet, detects collision during the transmission of its 12th bit on the wire (including any preamble). Given that the speed of the signal in the wire is 2 x 10 ^8 meter/second, how far away (in meters) is the rogue station from you machine?

If we see the collision after 12 bit times, then that means that the rogue host sent its competing transmission after 6 bit times. This means that the distance is: 6 bits/10 Mbps * (2 * 10^8 m/s) = 120 meters

Q4: [10 Marks]

Consider an 802.11 WLAN operating at 54 Mbps, in DCF mode and having several stations (STA). Each STA runs a voice over IP (VoIP) application which generates a 240 byte VoIP packet, every 20 milliseconds. The VoIP packet is sent using UDP over IP to the destination. Compute the maximum number of VoIP calls that can be supported.

Assume (VoIP header + UDP header + IP header + Checksum) = 44 bytes; (MAC header + FCS) = 33 bytes; ACK size = 11 bytes; RTS Threshold = 800 bytes. PHY header = 20μs; SlotTime = 20μs; SIFSTime = 10μs; (Hint: First compute the minimum overhead for one VoIP packet in DCF mode, including various headers, ACK and channel access times. Then compute how many such packets can be accomodated in one second. Finally compute the maximum number of calls, keeping in mind that each VoIP call is duplex.)

Time taken to transmit VoIP packet and its ACK = 160 microsec; [DIFS + PHY + (MAC+Other headers+data) * 8/54] + [SIFS + PHY + (MAC+Other headers+ACK) * 8/54].

6250 such packets per second => 6250/50 connections (since one connection generates 500 packets per second) => 6250/100 calls (since it needs a duplex).

NOTE TO TAs - PHY header is measured in time, not in bytes. Some may have taken it as bytes, assuming it is a typo. If so, grade accordingly. Give marks if they have got the basic idea even if final the ans is different. Since we are calculating max, the duration of a call does not matter. They can assume that a call lasts at least 1 second and get the upper bound.

Q5: [10 Marks]

Consider a network with links having different *capacities*, i.e., one link could have 100 Kbps capacity, and another could have 1 Mbps. We want to modify the Distance Vector (DV) protocol to compute paths which have the *maximum capacity* (instead of minimum cost). The capacity of a path is the minimum of the capacities of its constituent links (since the minimum capacity link is the bottleneck). From among the many paths between a Source and Destination, we want to choose the path that has maximum bottleneck capacity.

How would you modify the DV protocol to implement this feature? Specify clearly (a) What will a node advertise? (b) What does a node do on recepit of an advertisement from a neighbor?

1. In the DV table, we would now have the three columns as: <destination, capacityofmaxcapacitypath, nexthop>. And a node will advertise the first two entries.
2. When a node B receives a DV update from its neighbour A, specific to a given destination X, it does the following:
   - Capacity of path to X via A = min{capacity of B to A, capacity of A to X (specified in the advertisement)}
   - B updates its own DV entry for X if this newly calculated capacity is higher than the already stored capacity for destination X.
   - If such an update were to happen, the next hop for X would be marked as A

<u>Q6:</u> [5 Marks]

Three users X1, X2, X3 using three different computers within the IIT Bombay network serviced by a single mail server wish to send emails. X1 wishes to send email to Y1, X2 to Y2 and X3 to Y3, where Y1, Y2 and Y3 are three distinct users within IIT Kanpur network. What is the worst case and best case number of TCP connections that should be opened to enable this email exchange? Assume all the users use a non-HTTP email reader such as thunderbird. And the email should finally arrive in the inbox of the respective email readers. Specify all assumptions made clearly.

Best Case: Happens when all three users send email roughly the same time.
7 (1 each from X1, X2, X3 to their mail server, one from IITB mail server to IITK mail server, 1 each from Y1,Y2,Y3 to the IITK mail server)

Worst Case: Happens when they sent emails at different times.
9 (1 each from X1, X2, X3 to their mail server, one each from IITB mail server to IITK mail server for the 3 emails, 1 each from Y1,Y2,Y3 to the IITK mail server)

<u>Q7:</u> [5 Marks]

Consider a cell phone network. Assume that the geographical area covered by the network has been divided into a 2-dimensional array of square cells. When there is an incoming call for a given mobile, the network has to find the location of the mobile, i.e., it has to determine the cell in which the mobile is currently located. Consider the following location update strategy. A mobile updates its location after every k cell changes, for a given integer k. The network pages for the mobile in a group of cells near the mobile's last known location. For example: if k=1, then the mobile updates the location each time it changes a cell and no paging is required. If k is very large, then the network has to page all the cells. Let C(k) denote the number of cells the system must page for a given k. (i) What is the number of cells the system must page if k=3? (ii) What is the general expression for C(k)?

Draw the figure, count the cells and get the expression.
Cells in grid [k-1] will have to be paged, since kth hop will update. No of cells to be paged = [2(k-1)+1]^2
If mobile is in kth hop but update has not yet happened, answer of [2(k)+1]^2 may also be permitted.

Q8: **[10 Marks]**

Consider the topology shown in figure 2, where A, B, C and D are wireless nodes. The solid lines represent the transmission range of nodes A and D (centered at the nodes) and the dotted lines represent the transmission range of nodes B and C (centered at the nodes). Such an asymmetry in range can result when some nodes (A and D) transmit at the higher power than others (B and C). Answer the questions below, with appropriate justification for the answers.

(a) Suppose node C starts transmission of a broadcast packet. Nodes A, B and D also have their own packets to transmit. If all nodes employ carrier sensing, at this point, which nodes will defer their transmissions and which nodes actually transmit?

B and D will defer, A transmits

(b) Referring to above, given that some other nodes may transmit packets apart from node C, which nodes (among A, B and D) successfully receive the broadcast packet from C?

Only D will receive, B's packet would be corrupted by A.

(c) Suppose A wants to send large amount of data to B. At about the same time, C also wants to send large amount of data to D. Now, suppose the nodes are using ALOHA (do not implement carrier-sensing and transmit data as fast as they can, only limited by the physical capacity of the channel). What is the throughput achieved by the transfer from A toB and the transfer from C to D as a fraction of the wireless channel capacity? Assume that there are no wireless transmission errors.

A->B is 0 and C->D is 1. Since collision will happen at B but no collision at D.

(d) Referring to above, suppose all nodes employ CSMA (carrier sensing). Now calculate the throughput of A to B and C to D transfers as a fraction of the wireless channel capacity.

A->B is 1 and C->D is 0. Since C will sense A's transmission and keep quiet till A is done.

(e) Suppose all nodes employ RTS/CTS as well as a random backoff before transmission of any packet. Again calculate the throughput as a fraction of the wireless channel capacity.

A->B is 0.5 and C->D is 0.5. Whatever is sent will be received, however channel will be idle during backoff times. Assuming equal opportunity due to RTS/CTS and random backoff, each will transmit for half the time.