CS31006: Computer Networks (Theory)
Class Test – II SOLUTIONS

Time: 1 Hour **Marks: Highest**

Question 1

(1) A large number of consecutive IP addresses 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. [3]

Solution
Marks division: *1 mark is awarded for each correct answer, i.e. if someone gets 3 correct out of 4, he/she gets full marks.*
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 – 198.16.23.255 written as 198.16.16.0/21
C: 198.16.32.0 – 198.16.47.255 written as 198.16.32.0/20
D: 198.16.64.0 – 198.16.95.255 written as 198.16.64.0/19

(2) A router has the following (CIDR) entries in its routing table: [3]

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? (Answer in one sentence each)
(a) 135.46.63.10
(b) 135.46.52.2
(c) 192.53.40.7

Solution
The packets are routed as follows:
(a) Interface 1
(b) Router 2
(c) Router 1

(3) Based on the network shown below, state a sequence of events that will cause a count-to-infinity problem in Distance Vector Routing. (Assume that no corrective measure, e.g. poison reverse, is used). [2]

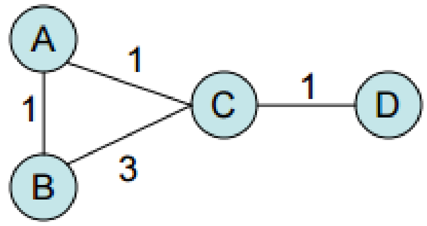

Solution

C-D fails. C routes to A for D, A continues to route to C for D. Routing loop is created. Distance vector cost increases in each round, until it reaches infinity.

(4) Describe two major differences between the ECN method and the RED method of congestion avoidance. [2]

Solution

(1) In ECN, router sends an explicit signal to the receiver that it is experiencing congestion. In RED, there is no explicit congestion signal – packet drops act as implicit/hidden signals.

(2) In ECN, packets are never dropped – rather they are marked if congestion is experienced. In RED, packets are dropped randomly.

Question 2

(1) Consider an instance of TCP's Additive Increase Multiplicative Decrease (AIMD) algorithm where the window size at the start of the slow start phase is 2 MSS (maximum segment size) and the threshold at the start of the first transmission is 8 MSS. Assume that a time-out occurs during the fifth transmission. Find the congestion window size at the end of the tenth transmission. [3]

Solution

<u>Answer</u>: 7 MSS

<u>Explanation</u>:

Since slow start is used, window size is increased by the number of segments successfully sent. This happens until either threshold value is reached or time-out occurs. When threshold is reached, window size increases linearly. When time-out occurs, threshold is halved.

Given, initial threshold = 8

- Time=1, congestion window size = 2 (1st transmission, slow start phase)
- Time=2, congestion window size = 4 (Double the no. of acks.)

- Time=3, congestion window size = 8 (Threshold is met)
- Time=4, congestion window size = 9, (Increase linearly after threshold)
- Time=5, congestion window size = 10, (Transmits, but time-out occurs)
    - New threshold = Current window size/2 = 10/2 = 5 MSS
- Time=6, congestion window size = 2 (Slow start phase)
- Time=7, congestion window size = 4 (Double the no. of acks.)
- Time=8, congestion window size = 5 (Since threshold is 5)
- Time=9, congestion window size = 6 (Increase linearly after threshold)
- Time=10, congestion window size = 7 (Increase linearly after threshold)

Hence at the end of the 10th transmission, size of congestion window is 7 MSS.

(2) Answer the following in one sentence each: [3x1=3]

(a) What is jitter?

(b) How is jitter controlled in the real-time transport protocol (RTP)?

(c) What is marshalling in a remote procedure call (RPC)?

Solution

(a) The variation of delay between packets reaching the receiver is called jitter.

(b) Jitter can be controlled by buffering packets at the receiver before they are played out (also called "playout with buffering").

(c) In a RPC, when a client invokes a *client stub* procedure with parameters, the parameters are *marshalled*; it includes converting the representation of the parameters into a standard format, and copying those into a message.

(3) For a host machine that uses the token bucket algorithm for congestion control, the token bucket has a capacity of 1 MB and the maximum output rate is 20 MBps. Tokens arrive at a rate to sustain output at a rate of 10 MBps. The token bucket is currently full and the machine needs to send 12 MB of data. What is the minimum time (in seconds) required to transmit the data? [4]

Solution

<u>Answer</u>: 1.2 seconds; (also accept if someone answers 1.1 seconds)

<u>Explanation</u>:

Capacity C = 1 MB, Max output rate M = 20 MBps, Arrival rate ρ = 10 MBps

Min. time required to transmit date, $S = C/(M - \rho) = 1/(20 - 10) = 1/10 = 0.1$ sec

Total time for 12 MB = 12S = 1.2 secs

<u>Alternative explanation</u>: Someone may assume that initial output rate is 20MBps, and therefore first 2 MB data takes (1/20-10) = 0.1 sec. This is also correct, so 1.1 seconds is also a valid answer.