EE 122 2nd Midterm Exam

1. **Question 1 (20 pt)**

(a) (10 pt) What is Red? Explain briefly how it works. Use no more than five phrases for the explanation.
(b) (10 pt) Name two goals of Red. Explain. Use no more than three phrases to describe each goal.

Note: For (a) you need to explain only the main ideas. You don't need to give the exact formula.

Answer:

(a) RED is a buffer management scheme that works in conjunction with a FIFO scheduler. RED decides which packet to drop/enqueue based on three main variables: the average queue length (avgLen), a minimum threshold (minThresh), and a maximum threshold (maxThresh):

(1) if (avgLen < minThresh) enqueue packet
(2) if (avgLen > maxThresh) drop packet
(3) otherwise, compute a probability p that increases linearly from 0 to Pmax (where Pmax < 1), and drop arrival packet with probability p

(b) RED goals:

(1) Avoid synchronization of multiple TCP flows. Because packets are dropped randomly before the buffer overflows, different flows will experience losses at different times. This breaks the synchronization amongst the competing flows.

(2) Avoid bursts of packet losses. With the drop-tail buffer management, multiple packets from the same TCP window can be dropped (when the buffer overflows). This may lead to retransmission timeouts. RED avoids this by dropping packets long before the buffer overflows.

2. **Question 2 (20 pt)**

(a) (5 pt) What is the goal of Fair Queueing?
(b) (15 pt) Assume a link of capacity 10 Mbps that is traversed by four flows with arrival rates of 6, 4, 2, and 1 Mbps, respectively. Compare the fair rate along the congested link. What bandwidth will each flow get? (Show all work.)

Answer:

(a) Fair queueing is a scheduling discipline that provides protection amongst competing flows. In particular, each flow receives no more than the fair rate, no matter how much traffic it sends. In the simplest case in which there are N flows competing on a link of capacity C and each flow has enough demand, each flow gets C/N bandwidth.

> (1) N = 4; C/N = 10 Mbps/ 4 = 2.5 Mbps; However, two flows need less than 2.5 Mbps, i.e., one needs 2 Mbps, and another one needs only 1 Mbps. As a result these flows will get 2 Mbps, and 1 Mbps, respectively.
> (2) The remaining bandwidth is distributed amongst the remaining flows: N = 2; (C – 3 Mbps)/N = 7 Mbps/2 = 3.5 Mbps. Since both flows have the arrival rate larger than 3.5 Mbps, the fair rate f of the link is 3.5 Mbps.

(b) As a result each flow i receives min (ri, f) bandwidth, where ri represents the arrival rate of the flow, i.e., min(6,3.5) = 3.5 Mbps; min (4,3.5) = 3.5 Mbps; min (2, 3.5) = 2 Mbps, and min (1, 3.5) = 1 Mbps.

3. **Question 3 (20 pt)**

Consider a TCP flow that sends 7 packets. Assume the TCP source experiences exactly one packet loss before finishing the transmission. Let d be the **one**-way propagation delay between the source and the receiver, and let RTO be the retransmission timeout. Let T denote the total time it takes the source to transmit all packets, i.e., the difference between the time when the source receives the acknowledge for the last (7^{th}) data packet and the time when the first data packet was sent.

Determine which packet has to be dropped to result in maximum T, and which packet has to be dropped to result in minimum T (if there are more than one possibilities in each case, present only one of them). Compute T in both cases as a function of d and RTO. For both cases draw the time diagram of the data packets and acknowledgements exchanged by the sender and the receiver. Upon each data packet arrival compute the congestion window size.

Note: The TCP implements both fast recovery and fast retransmission; fast recovery is triggered by 3 duplicate acknowledgements.

Answer:

(a) Worst case delay = 8d + RTO. This happens when the first packet is lost. Note that after the first packet is lost, TCP switches to additive increase. (This is not the only solution; for example, another solution is when packet 7 is lost).

(b) Best case delay = 8d. This happens when the 4^{th} packet is lost. This packet is retransmitted by using fast retransmission upon receiving the 3^{rd} duplicate acknowledge.

(a) (b)

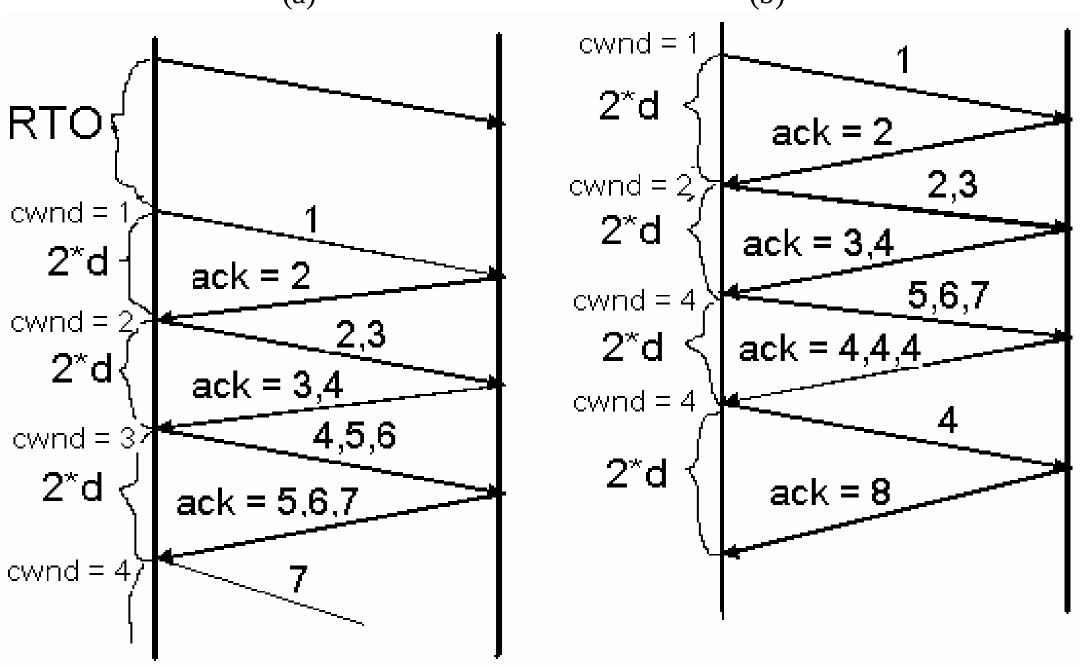

4. **Question 4 (20 pt)**

Consider the network below. Assume that the network implements the Distance Vector routing protocol.

(a) (5 pt) Write down the routing table of each node.
(b) (15 pt) Assume the cost of the link between nodes A and C decreases from 7 to 1. Write down the routing table of each node at every step until routing tables converge. Assume that both A and C see that the change of the link (A, C)'s cost at the same time, and that exchange of routing information and routing table updates are synchronous (i.e., they happen at the same time at all nodes)

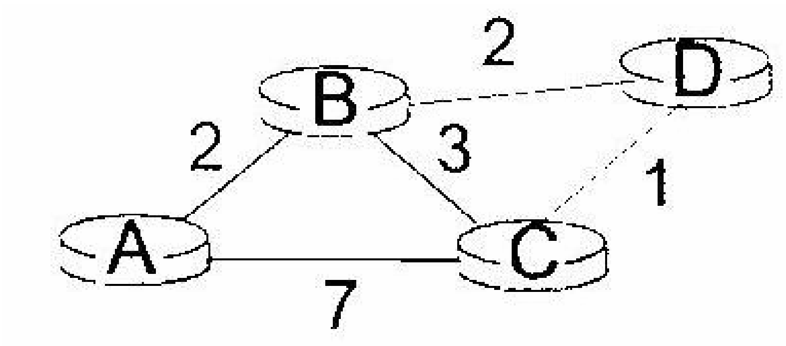

Answer: (a) See first column; (b) see the last three columns. Link costs that change during each iteration are in red

D^A	B	C
B	2	10
C	5	7
D	4	8

D^A	B	C
B	2	4
C	5	1
D	4	2

D^A	B	C
B	2	4
C	5	1
D	4	2

D^A	B	C
B	2	4
C	5	1
D	4	2

D^B	A	C	D
A	2	8	6
C	7	3	3
D	6	4	2

D^B	A	C	D
A	2	8	6
C	7	3	3
D	6	4	2

D^B	A	C	D
A	2	4	6
C	3	3	3
D	6	4	2

D^B	A	C	D
A	2	4	4
C	3	3	4
D	4	4	2

D^C	A	B	D
A	7	5	5
B	9	3	3
D	11	5	1

D^C	A	B	D
A	1	5	5
B	3	3	3
D	5	5	1

D^C	A	B	D
A	1	5	5
B	3	3	3
D	5	5	1

D^C	A	B	D
A	1	5	3
B	3	3	3
D	3	5	1

D^D	B	C
A	4	6
B	2	4
C	5	1

D^D	B	C
A	4	6
B	2	4
C	5	1

D^D	B	C
A	4	2
B	2	4
C	5	1

D^D	B	C
A	4	2
B	2	4
C	5	1

.

5. **Question 5 (20 pt)**

Consider a multicast group consisting of nodes Y, X and T in the figure below. Assume Y is the source, and X and T are the receivers.

(a) (5 pt) Draw the multicast tree in the case of DVRMP. What is the maximum path length from source to any of the receivers?

(b) (15 pt) Explain how the multicast tree is constructed in the case of CBT. Assume the core/root can be any of the nodes Z, U, or V. Draw the best **and** the worst multicast trees amongst the three resulting multicast trees. (The metric used to compare two multicast trees is the length of the longest path from source to any of the destinations; the multicast tree with the smallest length is the best one, while the multicast tree with the longest length is the worst one.)

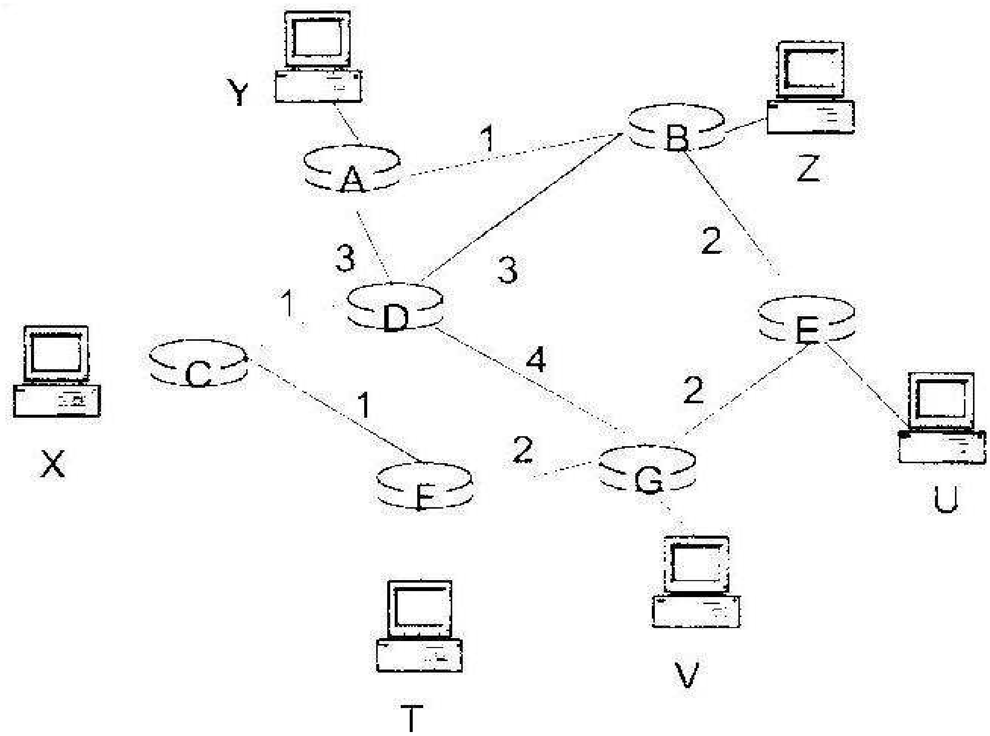

Answer: (a) DVRMP

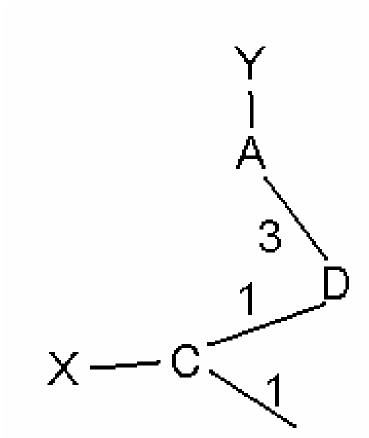

(b) With CBT each member of the multicast tree sends a join message towards a designated core/root. The multicast tree represents the union of all join messages. The CBT multicast tree is sensitive to the location of the core/root. Note that core/root does **not** take part in forwarding multicast packets or processing the join messages. In some cases even if the core fails (e.g., when the core is an end-host), multicast tree operations will not be affected. See the figures below for the best, respectively the worst multicast tree.

CBT-best

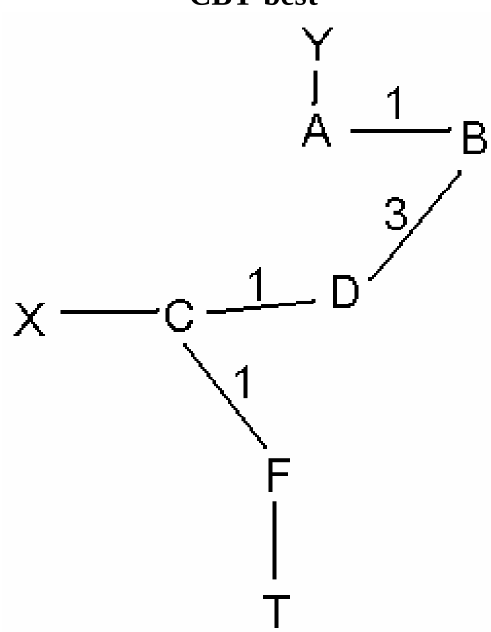

CBT-worst