Department of Computer Science and Engg, IIT Bombay
CS 348: Computer Networks Midsem: 10th Sept 2012, 11:00 – 13:00
Handwritten notes permitted. No photocopies. Weightage: 30%, Max Marks – 50

Q1: [10 Marks]

a) Hosts A and B are each connected to a router R via 10Mbps links (A ----- R ------B). The propagation delay on each link is 20 microseconds. R is a store and forward device i.e it doesn't transmit the packet on the R-B link until the whole packet is received on the A-R link. Suppose R forwards a packet 35 microseconds (processing delay) after it has finished receiving it. Calculate the total time required to transmit 10,000 bits from A to B, in the following two cases: (i) as a single packet, and (ii) as two 5000 bit packets sent one right after the other.

Solution:

(i) Per-link transmit delay is 10^4 bits / 10^7 bits/sec = 1000 μs.

Total transmission time = 2 × 1000 + 2 × 20 + 35 = 2075 μs.

(ii) When sending as two packets, here is a table of times for various events:

T=0 start
T=500 A finishes sending packet 1, starts packet 2
T=520 packet 1 finishes arriving at S
T=555 packet 1 departs for B
T=1000 A finishes sending packet 2
T=1055 packet 2 departs for B
T=1075 bit 1 of packet 2 arrives at B
T=1575 last bit of packet 2 arrives at B

Expressed algebraically, we now have a total of one switch delay and two link delays; transmit delay is now 500 μs: 3 × 500 + 2 × 20 + 1 × 35 = 1575 μs.

b) You have a link of bandwidth 10 Mbps and latency of 10 ms. Your friend has a link of bandwidth 1 Mbps and latency 5 ms. Which link do you consider is better? Why?
Solution:
For large messages (file downloads), bandwidth dominates latency. For small messages (VoIP), latency dominates bandwidth. So, it depends on the type of applications commonly used by you or your friend. If you use applications having both types of messages, 10 Mbps with 10ms is better.

(If any student has given a different answer but it makes logical sense, give suitable marks)

c) In Ethernet (802.3), the back-off window (BW) is exponentially increased on each unsuccessful collision (i.e., window of size 0, 2, 4, 8, 16, 32 .. 1024). After a successful transmission, back-off window is reset back to 0 immediately. What undesirable behavior may result due to this aspect of the algorithm of setting BW=0 after a successful transmission?
Solution:
Fairness is affected, since stations that were lucky enough to not collide will immediately transmit, so other stations will find the medium busy. Also, delay for colliding stations will increase, since they will have to count down from a larger window size before they can successfully transmit and get their BW reset to 0.

d) Give an example of a 4-bit error that would not be detected by a two-dimensional parity scheme. Explain your answer.
Solution: If we flip the bits corresponding to the corners of a rectangle in the 2-D layout of the data, then all parity bits will still be correct. Furthermore, if four bits change and no error is detected, then the bad bits must form a rectangle: in order for the error to go undetected, each row and column must have no errors or exactly two errors.

Q2: **[10 Marks]**

We have looked at ARQ protocols (stop & wait, sliding window) assuming full-duplex links, i.e data can be transmitted in both directions of the link simultaneously (You could be sending ACKs at the receiver while simultaneous receiving data sent by the sender). Now consider half-duplex links, transmission can happen in only one direction at a time (Eg. Walkie-talkie).

Consider a half-duplex link with a one-way propagation delay of P seconds and a bandwidth of B bits/sec. The sender and the receiver decide to share the link using Time Division Multiplexing i.e. the sender sends for T_S time and then the receiver sends for T_R time and so on. Refer to the figure below. Assume that the sender sends data in packets of size F bits and receiver sends ACK in packets of size A bits.

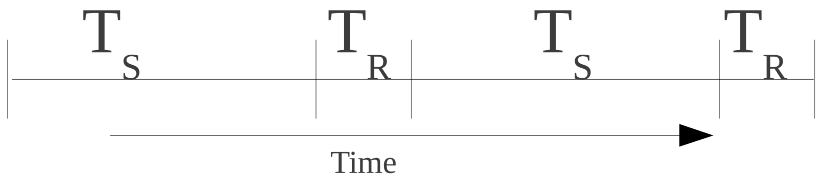

a) Suppose you are implementing the stop and wait protocol. Let $T_S = P + F/B$ and let $T_R = P + A/B$ What is a natural value for sender timeout for retransmission? Explain your answer.

b) Does the sender need to number the packets? Explain why or why not.

c) Does the receiver need to number the ACKs? Explain why or why not.

d) What is the efficiency of this system in terms of link usage i.e. effective throughput? In other words, if this reliability is being implemented at the transport layer, how many bits of application layer data are you able to get across from sender to receiver per second. (Assume no packet losses)

e) To improve efficiency, we change the implementation to a sliding window. If the sender window size is limited to X frames, how would you set T_S and T_R to get maximum link efficiency? What is the resulting link-efficiency? (Assume no packet losses).

Q2 Solution

a) Ts+Tr
b) Yes. Consider the case when the ack is lost. The transmitter will retransmit the original packet while the receiver is expecting the next packet.
c) No. An ack always acknowledges only the packet sent in the same slot.
d) F/B over (Ts + Tr)
*e) Ts = XF/B + P, Tr = P + A/B. X*F/B over (Ts + Tr)*

Q3: [10 Marks]

Suppose you want to transfer an N-byte file along a path having five switches between the source and the destination. Each link along the path has a propagation delay of 2ms, bandwidth 10Mbps and that the switches support both packet and circuit switching. Thus, you can either break the file up into packets or set up a circuit through the switches and send the file as contiguous bit stream.
Suppose packets have 20 bytes of packet header information and 1000 bytes of payload. Suppose that the circuit setup requires a 1000 byte message to make one **round-trip** on the path.

1. If cost is a direct function of the total number of bytes sent across the network, for what file size (N bytes) does it cost less to use circuit switching over packet switching? You may assume that file size is a multiple of 1000 bytes. Ignore overheads other than those mentioned above. [Explain all sub-components that make up the cost.]

2. If cost is a direct function of the total latency incurred before the entire file arrives at the destination, for what file size does it cost less to use circuit switching over packet switching.

Q3 Solution:

1. Circuits pay an up-front penalty of 1000 bytes being sent on one round trip for a total data count of 2* 1000 + n. Packets pay an ongoing per packet cost of 20 bytes for a total count of $1020 \times n/1000$. Thus for files 100,000 bytes or longer, using packets results in more cost i.e total data sent on the wire.
2. Packet switching: Made up of 3 components
   - transmission of file = (n/1000) packets * (1020*8/10Mbps)tx time = 0.000816 * n (ms)
   - sum of tx times across the switches = (020*8/10 Mbps) * 5 = 4.080ms
   - propagation delay = 6 * 2ms = 12 ms
     [The important thing to note is the pipelining, as the first packet traverses the links, the other packets are being transmitted by the sender]

   Circuit switching also has 3 components
   - total transmission time: 8 * n/10Mbps = 8n * 10^-4 (ms)
   - Total propogation delay = 6 * 2ms = 12 ms
   - setup delay = (0.8+2) * 12 ms = 33.6ms

Solve for n, where packet switching cost > circuit switching cost

Q4: [5 Marks]

Consider a string topology of 4 nodes, A,B,C and D. Assume bi-directional wireless links between A and B, B and C, and C and D. Each node uses a RTS-CTS based access mechanism during transmission. Assume node A wants to send a frame to node B and a little later (when A's transmission is ON) node C wants to transmit a frame to node D. Draw the time-sequence diagram from the start to the end of communication.

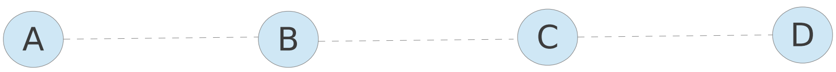

Q4 Solution:

RTS-by A; CTS-by B; C overhears CTS and sets NAV for duration of transmissoin (keeps quiet till A's frame-exchange-sequence RTS-CTS-DATA-ACK completes; RTS-by C; CTS-by D; C-Data; D-Ack.

Some of them may have assumed that “A's transmission is ON” implies that “when A is sending the RTS”. If so, they will get an RTS collision and then both A and C will back off and retry as above.

Q5: [5 Marks]

Consider the extended LAN shown in Fig 1 below. Show the main steps of the spanning tree algorithm being executed by the Bridges. Show which ports of which bridges are not selected in the final configuration after completion of the spanning tree algorithm.

Q5 Solution: B2:A; B5:B; B5:F;B6:I

<u>Q6:</u> [10 Marks]

Consider a network where all the links have a cost of 1. The routing tables for nodes A and F are shown in Fig 2 below.

1. Give a diagram of the smallest network consistent with these tables.
2. Show the distance-vector routing table for all the nodes *in the initial state,* when each node knows the distance only to its immediate neighbours.
3. Show the distance-vectors after *one round of RIP updates* with its immediate neighbours.

<u>Q6 Solution:</u>

1. The cost=1 links show A connects to B and D; F connects to C and E. F reaches B through C at cost 2, so B and C must connect. F reaches D through E at cost 2, so D and E must connect. A reaches E at cost 2 through B, so B and E must connect.

```
 A – B – C
 |   |   |
 D – E - F
```

2. Values in Black below.
3. Changed Values are shown in Red below.

	A	B	C	D	E	F
A	0	1	-, 2	1	-, 2	-
B	1	0	1	-, 2	1	-, 2
C	-, 2	1	0	-	-, 2	1
D	1	-, 2	-	0	1	-, 2
E	-, 2	1	-, 2	1	0	1
F	-	-, 2	1	-, 2	1	0

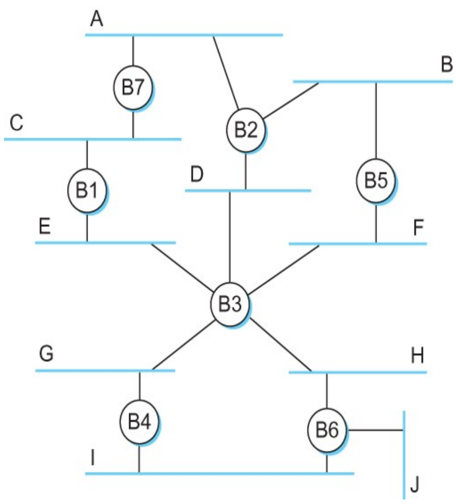

Figure 1 - for Q5: Spanning Tree Algorithm

Table for A				Table for F		
Node	**Cost**	**Nexthop**		**Node**	**Cost**	**Nexthop**
B	1	B		A	3	E
C	2	B		B	2	C
D	1	D		C	1	C
E	2	B		D	2	E
F	3	D		E	1	E

Figure 2 - for Q6 – Routing Tables