EE122: Communication Networks		**SP'10**
	Final — May 11, 2010	
Lecturer: Jean C. Walrand		*GSI: Victor Wen*

SOLUTIONS

Problem 1. (10%)

Consider the Ethernet network shown in Figure 1. Each rectangle is an Ethernet switch.

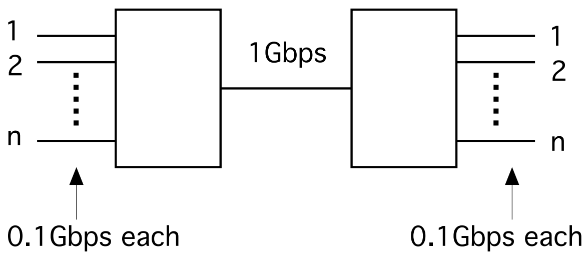

Figure 1. Ethernet Network for Problem 1.

The links are full-duplex with the rates shown. Assume that one host is attached to each 0.1Gbps link and that each host sends data at an average rate of Rbps, with equal fractions to all the other $2n-1$ hosts. (Thus, host i sends data to host j at an average rate $R/(2n-1)$ for all $j \neq i$.)

(a) Find the maximum value of R.

(b) For what values of n does the 1Gbps link become the bottleneck?

(a) First, $R \leq 0.1$ since the rate leaving a host cannot exceed the line rate. Second, consider the rate that arrives to host 1. This rate is $(2n-1) \times (R/(2n-1)) = R$. Thus, $R \leq 0.1$. Third, the rate that goes from the left switch to the right switch is $nR \times (n/(2n-1))$. Thus we need $nR \times (n/(2n-1)) \leq 1$. Hence, the maximum possible value of R is

$$\max\{0.1, \frac{2n-1}{n^2}\}.$$

(b) The 1Gbps link becomes the bottleneck when

$$\frac{2n-1}{n^2} < 0.1, \text{ i.e., } n \geq 20.$$

Problem 2. (12%)

This problem studies the throughput of a Reservation ALOHA network that works as follows (see Figure 2). There are N stations (not shown in the figure). Time is divided into alternating reservation and transmission phases. A reservation phase has K slots. The N stations transmit independently with probability p during each reservation slot, to make a reservation for later transmission. A reservation is successful when exactly one station sends a reservation during a reservation slot. After the reservation phase, the stations (if any) that were successful in making a reservation transmit their packets in turn, one after the other. After these transmissions, a new reservation phase starts, and so on. Assume that the average transmission time of a packet lasts the same as M reservation slots.

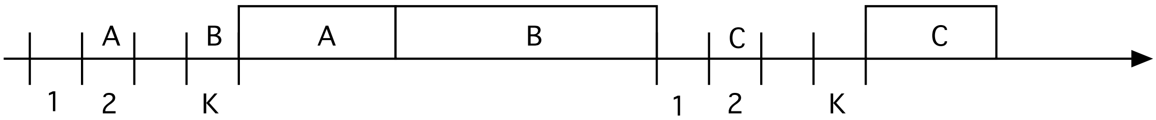

Figure 2. Reservation ALOHA Network for Problem 2.

(a) What is the fraction of time η that the stations transmit packets (not reservation messages).

(b) Assume that N is fixed. What is the values of p that maximizes the above fraction of time? What is the corresponding value of η when N is large?

(c) Assume that N and p are fixed. What is the average number A of slots between transmissions of packets by one given station?

(a) The probability that a success occurs during one reservation slot is $Np(1-p)^{N-1}$. Consequently, the average number of successful reservations during the reservation phase is $KNp(1-p)^{N-1}$. Thus, a reservation phase of K slots is followed by $MKNp(1-p)^{N-1}$ slots of packet transmission, on average. Accordingly, the average fraction of time that the stations transmit packets is

$$\eta = \frac{MKNp(1-p)^{N-1}}{MKNp(1-p)^{N-1} + K} = \frac{MNp(1-p)^{N-1}}{MNp(1-p)^{N-1} + 1}.$$

(b) The best value of p maximizes the probability of success, i.e., $p = 1/N$. Then $\eta \approx M/(M+e)$.

(c) The N stations transmit η packets per slot, on average. Thus, one given station transmits η/N packets per slot, on average. Since each packet transmission takes M slots, on average, it must be that the average time A between two packet transmissions by one given station is such that

$$\frac{M}{M+A} = \frac{\eta}{N}.$$

Hence,

$$A = M(\frac{N}{\eta} - 1) \approx M(N-1) + Ne \text{ if } p = \frac{1}{N}.$$

Problem 3. (12%) *In this problem, we investigate the effect of packet losses due to wireless transmission errors on the throughput of a TCP connection. The model is as follows. A source of a TCP connection increases its transmission rate by ρbps every RTT. The source multiplies its rate by $\alpha < 1$ every time that it notices that a packet is lost.*

(a) We assume that a packet is lost exactly every T seconds, independently of the transmission rate (for simplicity of analysis). Say that the source starts with rate $R(0) = 0$. Plot the rate of the source as a function of time. Thus, the rate increases with rate ρ/RTT for T seconds, then gets multiplied by α, then increases with rate ρ/RTT for T seconds, and so on.

(b) Does the average rate of the source converge to some limit? If so, what is the limit when $\rho = 10^4, RTT = 0.1, T = 1, \alpha = 0.5$?

(a) Here is a plot:

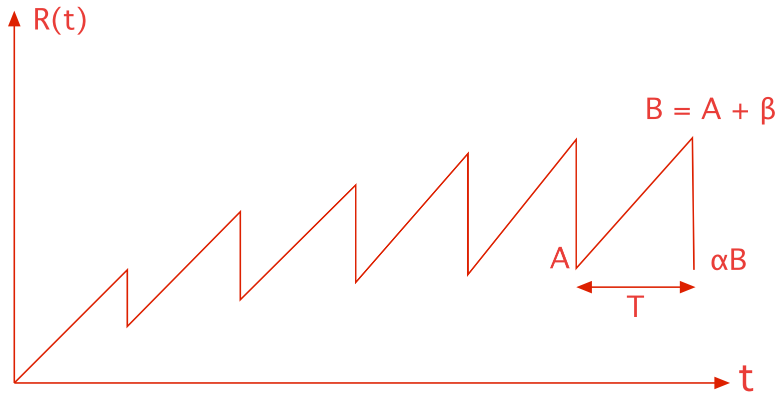

(b) The plot shows that $R(t)$ eventually repeats a cycle where it grows from A to $B := A + \beta$ in T seconds, with $\beta = \rho T/RTT$, and then drops to αB. Thus, we need $A = \alpha B = \alpha(A + \beta)$, which implies that $A = \alpha\beta/(1 - \alpha)$. The average rate is then

$$R = A + \frac{\beta}{2} = \frac{1 + \alpha}{2(1 - \alpha)}\beta.$$

For the numerical example, we find that the average rate is $R = 150$kbps.

Problem 4. (10%)

Consider a TCP connection that is in congestion-avoidance phase for a long time. Its rate increases from $R/2$ to R, then drops to $R/2$, and this cycle repeats. [R is in packets/second.] The rate increases at rate $1/T^2$ where $T = RTT$ (in packets/second^2). Assume that one packet is lost in each cycle.

(a) What is the duration S of a cycle , in terms of R and T?

(b) What is the average throughput λ of the connection, in terms of R?

(c) What is the number N of packets sent in a cycle, in terms of R and T?

(d) What is the loss rate ρ of the connection, in terms of R and T?

(e) Express λ in terms of T and ρ.

(f) Assume two TCP connections with different RTTs share a single bottleneck link. Do you think they have the same loss rate? Explain.

(a) $S = (R/2)/(1/T^2) = R.T^2/2$ since S is the time needed for the rate to increase from $R/2$ to R.

(b) $\lambda = 3R/4$.

(c) $N = \lambda S = 3R^2.T^2/8$.

(d) $\rho = 1/N = 8/(3R^2.T^2)$.

(e) Hence,

$$\lambda = \frac{\sqrt{3/2}}{T\sqrt{\rho}}.$$

(f) No. We saw earlier that the rate R of a source is proportional to $1/T^2$ where $T = RTT$ of that source. Thus, the loss rate ρ is proportional to RTT^2 (in view of (d) above). Intuitively, a source with a small RTT sends packets fast and drops one every S seconds, where S does not depend on RTT. One might be tempted to think that the loss rate is the same for all the sources, which would imply that the throughput is inversely proportional to RTT. Unfortunately, that is not true. Assuming that the loss rate is the same for all connections appears to be a common misconception in the literature. This misconception underlies the strange notion of "TCP-friendliness" for a transport protocol that exhibits the relationship between throughput and loss rate given in (e).

Problem 5. (12%) *Consider a wireless LAN (WLAN) operating at 11 Mbps that follows the 802.11 MAC protocol with the parameters given below. There is only one user present in this WLAN who wants to transmit a file with 10^6 bytes of data using UDP.*

The drawings in Figure 3 (not drawn to scale) show the sequence of channel activities for a successful transmission and for a transmission with channel errors.

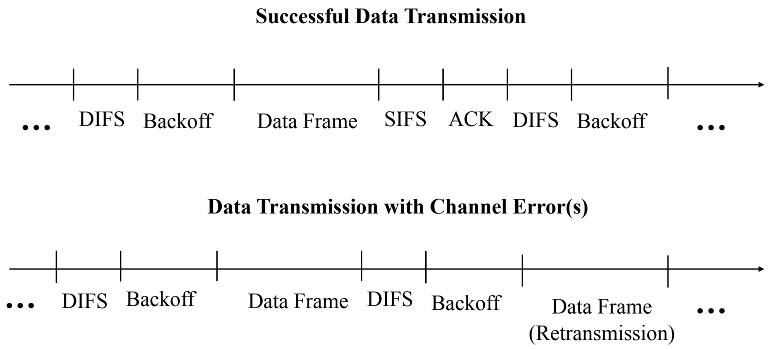

Figure 3. Plot for Problem 5.

Recall that transmission of each data or acknowledgement frame incurs the Physical Layer overhead as indicated below.

The parameters for the WLAN are as follows:

$Slot = 20\mu s$
$DIFS = 2.5$ *Slots*
$SIFS = 0.5$ *Slots*
$ACK = 10.5$ *Slots (including Physical Layer overhead)*
Physical Layer overhead for data frames = 10 Slots
Collision Windows grow as 31, 63, 127, ... Slots.

Assume the following:

- *All data frames are 1100 byte long which includes 100 byes of overhead for MAC and upper layers;*

- *Propagation delay is negligible;*

- *An acknowledgement is always received without any errors.*

With the above parameters and assumptions, answer the following questions:

(1) How long will it take transmit the file if data frames are always transmitted without errors?

Each data frame carries 1000 bytes of payload. Hence, 1000 data frames need to be transmitted. Each data frame on average takes

$$(2.5 + 15.5 + 10 + 0.5 + 10.5) \times 20 + (1100 \times 8)/(11 \times 10^6)\mu s = 1580\mu s.$$

Hence, the file will take $1000 \times 1580\mu s$ *or 1.58 seconds.*

(2) How long will it take to transmit the file on average if each new data frame is received correctly either after the original transmission or after one retransmission with probabilities 0.9 and 0.1, respectively?

A data frame needing retransmission will take

$$(2.5 + 15.5 + 10) \times 20 + (1100 \times 8)/(11 \times 10^6) + (2.5 + 31.5 + 10 + 0.5 + 10.5) \times 20 \\ + (1100 \times 8)/(11 \times 10^6)\mu s = 3260\mu s.$$

Hence, the file will take $1000 \times (0.9 \times 1580 + 0.1 \times 3260)\mu s = 1000 \times 1748\mu s$ *or 1.748 seconds.*

(3) If each data frame transmission (an original transmission or a retransmission) has the probability 0.9 of being received correctly, how will the time to transmit the file compare to the estimate in (2) above? Explain your answer. (Only a qualitative answer is needed for this part.)

Strictly speaking, there is a small probability that the file will not be transmitted fully. This is because there is a small probability that after a certain number of attempts, transmission of that data frame may be given up on. Ignoring this effect, the transmission time of the file will be higher since for a given data frame this scenario has the probability of needing higher number of retransmissions as compared to the scenario in (2).

Problem 6. (12%)

Suppose a router has built up the routing table shown in the following table. The router can deliver packets directly over interfaces 0 and 1, or it can forward packets to routers R2, R3, or R4. Assume the router does the longest prefix match. Describe which subnet mask the router uses with a packet addressed to each of the following destinations and which interface or router does it forward the packet to:

Subnet Number	Subnet Mask	Next Hop
128.96.170.0	255.255.254.0	Interface 0
128.96.168.0	255.255.254.0	Interface 1
128.96.166.0	255.255.254.0	R2
128.96.164.0	255.255.252.0	R3
Default		R4

a *128.96.171.92*

Ans. *Applying the subnet mask 255.255.254.0, we get 128.96.170.0. Use the interface 0 as the next hop.*

b *128.96.167.151*

Ans. *Applying the subnet mask 255.255.254.0, we get 128.96.166.0 (next hop is Router 2). Applying the subnet mask 255.255.252.0, we get 128.96.164.0 (next hop is Router 3). However, 255.255.254.0 is a longer prefix. Use Router 2 as next hop.*

c *128.96.163.151*

Ans. *None of the subnet number entries match, hence use default Router R4.*

d *128.96.169.192*

Ans. *Applying subnet mask 255.255.254.0, we get 128.96.168.0. Use interface 1 as the next hop.*

e *128.96.165.121*

Ans. *Applying subnet mask 255.255.252.0, we get 128.96.164.0. Use Router 3 as the next hop.*

Problem 7. (12%)

Suppose we have the forwarding tables shown in the following table for nodes A and F, in a network where all links have cost 1. Give a diagram of the smallest network consistent with these tables.

A

Node	*Cost*	*NextHop*
B	*1*	*B*
C	*1*	*C*
D	*2*	*B*
E	*3*	*C*
F	*2*	*C*

F

Node	*Cost*	*NextHop*
A	*2*	*C*
B	*3*	*C*
C	*1*	*C*
D	*2*	*C*
E	*1*	*E*

The following is a sample topology.

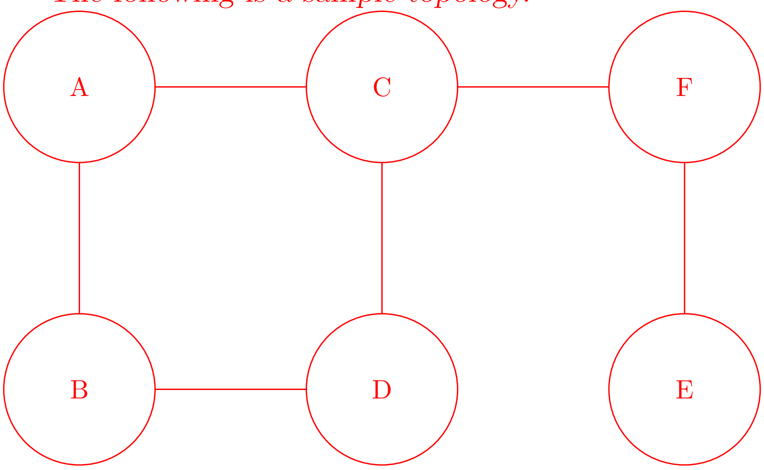

Problem 8. (20%)

1 *Let A be the number of autonomous systems on the Internet, and let D (for diameter) be the maximum AS path length.*

a *Give a connectivity model for which D is of order $logA$.*

A regular tree.

b *Give one for which D is of order $\sqrt{A}$.*

A arrangement of routers in a grid that is $\sqrt{A}$ by $\sqrt{A}$.

2 *The reason(s) that routing is decomposed into two levels is/are: [Mark the correct answer(s) with an X]*

a *Scalability X*

b *The design of IP*

c *Historic*

d *Economic/Business X*

e *All of the above*

3 *State the key difference in path advertisements between transit and peering agreements between ISPs.*

In a transit agreement, the ISP advertises all the paths it learns. In a peering agreement, the ISP only advertises its local paths.

4 *Consider a token bucket of bucket size B bits and rate of tokens α bits/s that is full at time $t = 0$. A source starts transmitting data at r bits/s at time t=0 and is regulated by this token bucket. Let t be the time until the output rate of data from the token bucket regulator is r bits/s. Then [Mark the correct answer(s) with an X]*

a $t = B/(\alpha - r)$ if $\alpha > r$

b $t = \infty$ if $\alpha > r \leftarrow X$

c $t = B/(r - \alpha)$ if $r > \alpha \leftarrow X$

d $t = B/r$ irrespective of α

5 *Suppose a switch performs processor sharing on 3 queues. Let packets of sizes 3, 6 and 5 units respectively arrive at the three queues at time t=0. What are the finish times of the respective packets if the outgoing link rate is 1 unit/s? [Mark the correct answer with an X]*

a *3s, 6s, 5s*

b *7s, 14s, 13s*

c *3s, 9s, 14s*

d *8s, 9s, 5s*

e *9s, 14s, 13s X*

6 *Assume a TCP connection starts in slow start (an later stays in that phase) with an initial window size equal to 10KBytes and with an RTT equal to 0.2s. Assume also that no packet gets lost. Approximately how long does it take for a 1MByte file to reach the destination? [Circle the best answer.]*

$$19s, 1.9s, 2.6s, (1.3s), 0.6s.$$

7 *To minimize interference with neighboring cells in networks based on OFDMA, one can deploy [Mark the correct answer(s) with an X]*

a *Fractional frequency reuse. X*

b *Interference averaging using pseudo-random mapping of subcarriers to subchannels. X*

c *Coordinated use of subchannels in neighboring cells. X*

d *Careful spectrum planning to avoid neighboring cells using the identical frequency channels. X*

8 *LTE allows for the signaling plane prioritization (e.g., for preemption) based on [Mark the correct answer(s) with an X]*

a *Maximum Bit Rate.*

b *Allocation and Retention Priority. X*

c *Guaranteed Bit Rate.*

d *Aggregate Maximum Bit Rate.*

9 *LTE has chosen to use SC-FDMA in the UL direction mainly to reduce [Mark the correct answer(s) with an X]*

a *Delay Jitter.*

b *Interference.*

c *Peak-to-Average Power Ratio. X*

d *Inter-Symbol Interference.*

10 *Base Station in WiMAX can provide UL opportunity for transmitting [Mark the correct answer(s) with an X]*

- a *Bandwidth Request. X*
- b *Data. X*
- c *HARQ Feedback. X*
- d *End-to-End Performance Indication.*

11 *The key feature that makes the subcarriers in OFDM orthogonal is [Mark the correct answer(s) with an X]*

- a *Number of subcarriers is equal to the number of symbol times in a frame.*
- b *Symbol time is reciprocal of the subcarrier spacing. X*
- c *Number of subcarriers in a subchannel is fixed.*
- d *Multiplexing is allowed both in frequency and time dimensions.*

12 *In WiMAX, Bandwidth Request can be sent [Mark the correct answer(s) with an X]*

- a *On contention basis. X*
- b *Piggybacked with data. X*
- c *Upon receiving appropriate allocation from the Base Station. X*
- d *Out of band.*

13 *In LTE, handovers based on the principle of make-before-break is [Mark the correct answer(s) with an X]*

- a *Done away with. X*
- b *A key requirement.*
- c *Allowed per Base Stations discretion.*
- d *Negotiated at the time of setting up the bearer.*

14 *The primary duplex modes for LTE and WiMAX are [Mark the correct answer(s) with an X]*

- a *FDD for both.*
- b *TDD for both.*
- c *FDD and TDD, respectively. X*
- d *TDD and FDD, respectively.*