Name_____________ Student ID___________________ Department/Year_________

Midterm Examination

Introduction to Computer Networks
Class#: 901 31110
Fall 2004

9:20-11:00 Tuesday
November 16, 2004

Prohibited

1. You are not allowed to write down the answers using pencils. Use only black- or blue-inked pens.
2. You are not allowed to read books or any references not on the question sheets.
3. You are not allowed to use calculators or electronic devices in any form.
4. You are not allowed to use extra sheets of papers.
5. You are not allowed to have any oral, visual, gesture exchange about the exam questions or answers during the exam.

Cautions

1. Check if you get 10 pages (including this title page), 6 questions.
2. Write your **name in Chinese**, student ID, and department/year down on top of the first page.
3. There are in total 150 points to earn. You have 100 minutes to answer the questions. Skim through all questions and start from the questions you feel more confident with.
4. Use only English to answer the questions. Misspelling and grammar errors will be tolerated, but you want to make sure with those errors your answers will still make sense.
5. If you have any extra-exam emergency or problem regarding the exam questions, raise your hand quietly. The exam administrator will approach you and deal with the problem.

1. (Overview) We have seen the figure on the right for many times in class. It illustrates a typical composition of a network on the Internet.

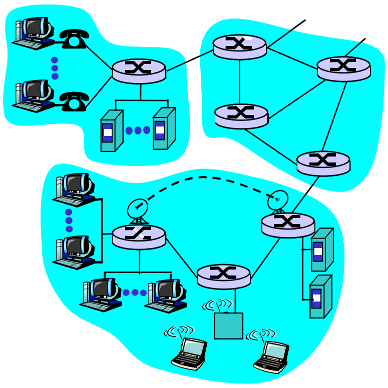

   (1) Based on how you understand what the Internet edge vs. core is, circle out the computers belonging to the edge of the network. (2%)
   (2) Continued from (1), circle out the computers belonging to the core of the network. (2%)
   (3) State the difference between the two communication models running in between computers on the edge of the network, client-server vs. peer-to-peer models. (5%)
   (4) State the difference between the two data transmission models running among the computers in the core of the network, circuit-switched vs. packet-switched model. (5%)

Sample Solution:

(1) (2)

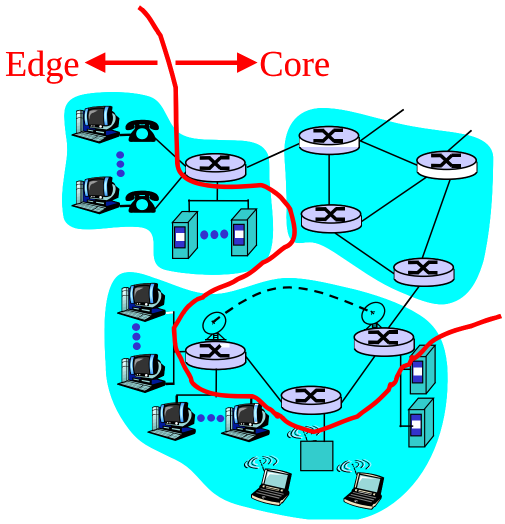

(3)

In the client-server model, clients request for services and servers provide services.

In the peer-to-peer model, all computers are both clients and servers. They may request for services and in the meantime provide services.

(4)

In the circuit-switched model, network bandwidth is divided into pieces. A certain pieces of bandwidth will be reserved and dedicated for a call. The reserved bandwidth will be sitting in idle when it is not used by the call.

In the packet-switched model, the data is divided into packets. All packets coming from all places will be sharing the bandwidth (statistically multiplexing).

2. (Overview) Consider a simple network as follows. The bandwidth of the link is R bits per second and the packets always come in the size of P bytes. The link speed is S meter per second and the length of the link is D meters.

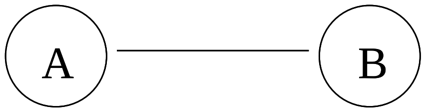

(1) Give the transmission delay of a packet (in terms of R, P, S, or D). (5%)
(2) Give the propagation delay of the link (in terms of R, P, S, or D). (5%)
(3) Suppose node A has infinite amount of buffer space and the data come into node A in a constant rate, 2R. Give the amount of data waiting to be transmitted in node A at time T seconds (in terms of T, R, P, S, or D). (5%)
(4) Suppose the processing delay is negligible. Give the nodal delay of the packet arriving at node A at time T (in terms of T, R, P, S, or D). (10%)

Sample Solution:

(1) 8P/R
(2) D/S
(3) RT
(4) 8P/R + D/S +RT

3. (Application) Address the following questions concerning cookies.

Credit: NABiSCOworld.com

   (1) Have you ever had Oreo cookies (or any kind of cookies) before? (2%)
   (2) Have you ever been given cookies in exchange of your home phone number when you were a child? (2%)
   (3) How do the HTTP cookies work in general? (5%)
   (4) Can the HTTP cookies be used to exchange for personal information? (5%)
   (5) Have you ever been told not to take cookies from the strangers when you are a child? (2%)
   (6) Would you advise others not to take HTTP cookies from the unknown websites? And why? (5%)

Sample Solution:

(1) Yes or no, whatever your experience is.

(2) Yes or no, whatever your experience is.

(3)
   1. A user visits a Web site
   2. The Web server sets up cookies with the information known about the user at the Web client
   3. When the user uses the same Web client to revisit the Web site, the cookies with the known information about the user will be sent to the Web server

(4) Yes. A user may give his/her name, email address, and etc in order to log in, to shop or to obtain files from the Web site. As a result, the Web site will be able to know your identity even if you do not log in specifically the next times.

(5) Yes or no, whatever your experience is.

(6) Yes or no, whatever you think (for example, yes to protect your privacy).

4. (Application) Let us try to understand the bandwidth cost of query flooding in Gnutella peer-to-peer system. Suppose the algorithm of query flooding is defined as follows:

```
If (self is source)
        forward query Q to all neighbors
Upon receiving a query Q from neighbor N {
        if (Q is new)
                forward Q to all neighbors except N
        else
                discard Q
}
```

(1) Consider a simple string network with 4 nodes. Node 1 is the query source. The source sends a query packet to all neighbors. When Node 2 receives the query packet, it reacts based on the 'algorithm of flooding' defined above. Node 2 will forward the query packet only to neighbor Node 3. After one more iteration, the query packet reaches Node 4. Node 4 will simply discard the query packet and the flooding process of a query is completed. During the flooding process, 3 copies of the query packet are transmitted on the network entirely. State the number of copies of query packets transmitted in the query flood process for the n-node string network. (5%)

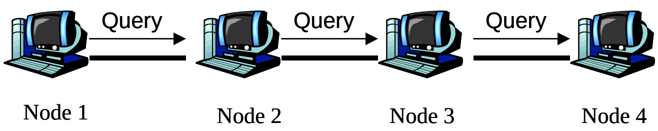

(2) Now consider a simple grid network of 2 x 2 nodes. Node 1 is again the query source. The source sends a query packet to all neighbors. When Node 2 and 3 receive the query packet, they react based on the 'algorithm of flooding'. Node 2 and 3 will both forward the query packet to Node 4. Then, Node 4 will forward a copy to Node 3 if it receives the query packet received from Node 2 first. Otherwise, it will forward a copy to Node 2 for the query packet received from Node 3. Node 3 or Node 2 has already seen the query packet before they will simply discard the query packet. There are 5 copies of the query packet being transmitted in the query flood process. Follow the algorithm. State the number of copies of query packets transmitted in the process for the 3 x 3 node case. (10%)

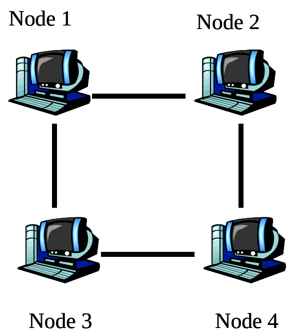

(3) Generalize the case and derive the communication cost of query flood for the n x n grids. (15%)

Sample Solution:

(1) n-1

(2) 16

(3) $3n^2-4n+1$

L: number of links

N: number of nodes

If all nodes pass to all their neighbors, the cost will be 2L

But all nodes, except the source, pass the packet to all-1 neighbors.

So the cost is 2L – (N-1)

L for n x n grid is 2n(n-1)

N for n x n grid is n^2

5. (Transport) Compare and contrast Go-Back-N and Selective Repeat based on the following criteria.
   (1) How is the semantic of the Acknowledges being different? (5%)
   (2) How is the semantic of timeout being different? (5%)
   (3) How is the amount of timers required at the sending node being different? (5%)
   (4) How is the memory requirement at the receiving node being different? (5%)
   (5) How is the actual amount of packets in-flight being different? (5%)
   (6) How is the amount of packet retransmission being different? (5%)
   (7) Suppose we know that most of the packet losses are due to congestion on the Internet and that when there is congestion, consecutive packets tend to be dropped altogether due to buffer overflow at the intermediate routers. Which of G-Back-N and Selective Repeat is better for the Internet? And why? (10%)

Sample Solution:

(1)

GBN: accumulative ack, acking for all packets with lower sequence number

SR: individual ack, acking for a specific packet each

(2)

GBN: a timeout means a whole window of packets are lost. Retransmit from the beginning of the window

SR: a timeout means a specific packet is lost. Retransmit only the packet

(3)

GBN: 1 per window

SR: 1 per packet in window

(4)

GBN: no need

SR: the window size to be safe. Not-in-sequence packets need to buffered until the gaps are filled to ensure ordering to the application layer

(5)

GBN: the whole window

SR: could be only a part of the window. The window base does not advance unless the lowest in-sequence unacked packet is acked

(6)

GBN: each timeout triggers retransmission from the beginning of the window. Some of the packets might actually have arrived but will be retransmitted anyway

SR: retransmission will only happen to packets that are truly lost.

(7)

Your choice. Whatever reason justifies your choice.
(For example: I choose GBN. GBN is better in terms of criteria compared in (3)(4)(5) and worse in (6). However if the network tends to drop packets consecutively, then GBN will not retransmit redundant packets too much, which means the problem in (6) is less critical and yet has the advantages in (3)(4)(5). Therefore, GBN is more suitable for the Internet.)

6. (Transport) Suppose the change of the congestion window size of a TCP connection is as depicted below.

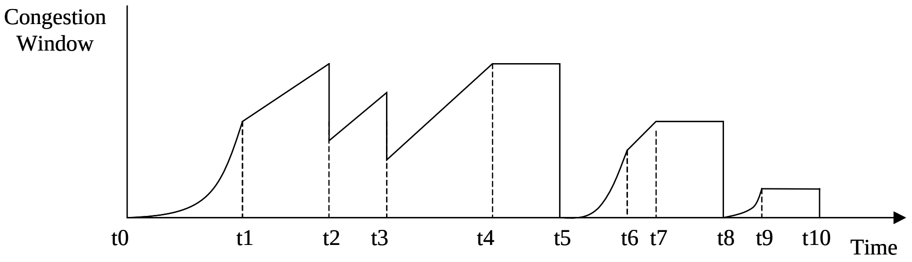

(1) State the time periods that the TCP connection is in the slow start phase. (5%)
(2) State the time periods that the TCP connection is in the congestion avoidance phase. (5%)
(3) State the time points that the TCP connection is experiencing a timeout. (5%)
(4) State the time point that the TCP connection is doing a Fast Retransmission. (5%)

Sample solution:

(1) t0-t1, t5-t6, t8-t10
(2) t1-t5, t6-t8
(3) t5, t8, t10
(4) t2, t3