Name_______________ Student ID____________________ Department/Year__________

Midterm Examination

Introduction to Computer Networks
Class#: 901 E31110
Fall 2007

9:30-11:10 Tuesday
November 27, 2007

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 16 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 are 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.

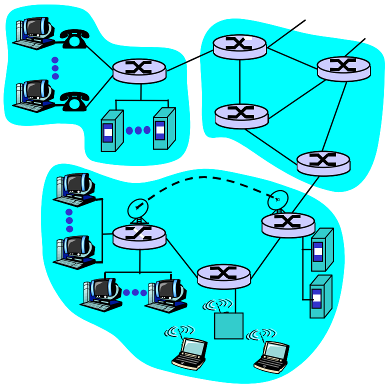

   (1) Based on how you understand what the Internet edge, core, and access network is, circle out the computers belonging to the edge of the network, the computers belonging to the core of the network, and the links belonging to the access network. (5%)

   (2) Consider two sets of terms, A={ workstations, laptops} and B={routers, switches}. Which set contains the terms to indicate different end systems use at the Internet edge? Which set to indicate different devices to route/forward data in the core of a network? (5%)

Sample Solution:

   (1) As depicted

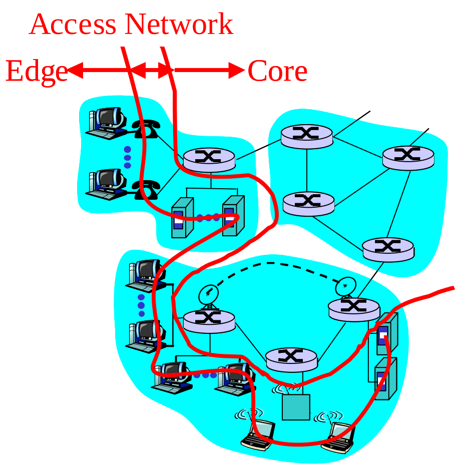

   (2) A: edge, B: core

2. (Overview) Packet switching and circuit switching are two ways of architecting data transfer in the core of a network. Suppose we have a simple 2-node network with a link in the middle connecting the 2 nodes. *L* users send data from one node to another through the link in the middle. The bandwidth of the link is *M* bps. Each user is active sending data P% of the time. Let *p* be P/100. When the user is active, it sends data in a constant rate *R* bps. Try if you can answer the following questions.

   (1) How do packet switching and circuit switching networks differ in general? (10%)
   (2) Give formula for the number of users *N* such a circuit-switching network can serve (in terms of *M* and *R*). (5%)
   (3) Give formula for the probability *Pr* of less than or equal to *N* users active simultaneously in such a packet-switching network (in terms of *N*, *L*, and *p*). (5%)

Sample Solution:

(1) Packet switching:

data sent in discrete chunks,
no call setup,
multiplexing,
congestion,
need additional mechanism to provide quality of service

Circuit switching:

one dedicated circuit per call,
call setup,
reserved resource (idle when there're no data within the call),
no congestion,
easier to provide quality of service

(2) $N = \dfrac{M}{R}$

(3) $Pr = \sum_{i=0}^{N} \binom{L}{i} p^{i} (1-p)^{L-i}$

3. (Application) Web clients and servers may send their HTTP requests and replies using the (1) non-persistent, (2) persistent without pipelining, or (3) persistent with pipelining mode. Suppose we have a web page which contains 10 objects, a .html main file, 8 .jpg image, and a .mp3 background music. The 10 objects are all so small that (1) the transmission time is negligible and (2) each object can be completely transferred in one TCP packet. Assume there are no packet losses, the round-trip time between the server and client is *RTT*, the time to set up a TCP connection are T_s, and the time to close down a TCP connection is T_c. In the non-persistent HTTP mode, the time required to complete the transfer of the web page is $10T_{s+}$ $10RTT + 10T_c$ and the communication between the client and server is depicted on the right. Based on the same assumptions, try if you can address the following questions.

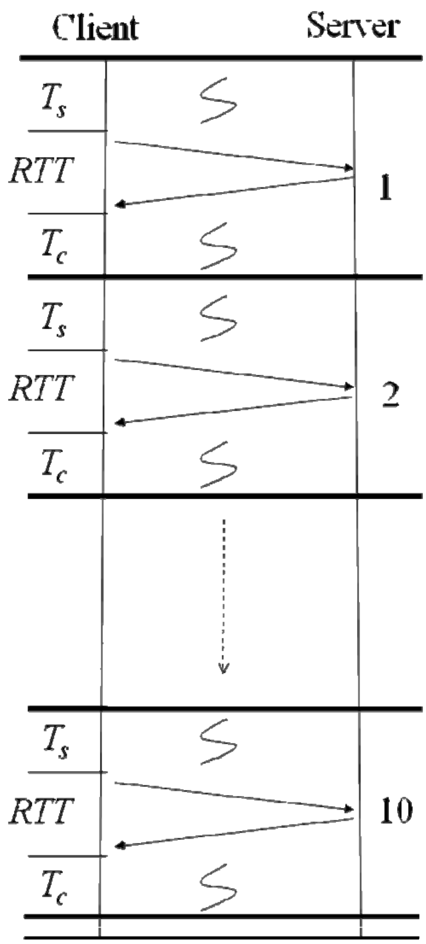

(1) What will be the time it takes to complete transferring the web page running in the persistent without pipelining mode? (5%)

(2) If the web client learns the URL of the 2nd to the 10th objects in the page after fully downloading the main html object, what will the entire page download time be using the persistent with pipelining mode? (5%)

Sample Solution:

(1) T_{s+} $10RTT + T_c$

(2) T_{s+} $2RTT + T_c$

4. (Application) Let's assume a 15-node Napster network and a 15-node KaZaA network as follow. In the Napster network, 14 regular peers connect to a super peer that provides the directory service. In the KaZaA network, 3 super peers are interconnected. Each of the super peers connect to 4 regular peers.

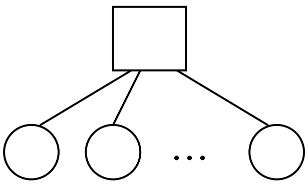

Napster

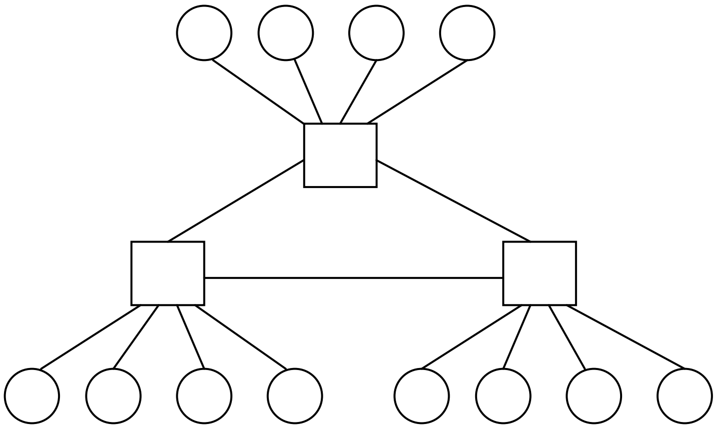

KaZaA

(1) Suppose each peer has N mp3 files to share. Information about each of the mp3 files requires K bytes to store. How much disk space is required to store information about all mp3 files in the Napster super peer? (5%)

(2) Following (1), how much disk space is required for a super peer in the KaZaA network to store information for itself and the regular peers connected to it? (5%)

(3) Assume the searching time is negligible. Suppose the time it takes for an mp3 file query to travel a link is the same. Say T seconds. How much time does it need for a peer in the Napster network to find an answer for a query? (5%)

(4) Suppose the desired mp3 files are evenly distributed. I.e., the chance of an mp3 file owned by each peer is the same. How much time does it need for a peer in the KaZaA network to find an answer for a query? (5%)

(5) If the super peer in the Napster network breaks down, how many % of the peers remain connected? (5%)

(6) If a super peer in the KaZaA network breaks down, how many % of the peers remain connected? (5%)

Sample Solution:

(1) 15NK bytes

(2) 5NK bytes

(3) 2T seconds

(4) (4/14)*2T+(10/14)*4T = 48T/14 = 24T/7seconds

(5) 0%

(6) 66.67% or 2/3

5. (Transport) TCP is a transport layer service that provides a number of major functions: (1) reliable data delivery, (2) connection management, (3) flow control, and (4) congestion control. Based on your understanding of TCP, try if you can address the following questions:
   (1) In classic reliable data delivery mechanisms, packet losses are detected by timer timeout. Setting the timeout interval is a challenging task. Can you describe what will happen when the timeout interval is set too short and what will happen when the timeout interval is set too long? (10%)
   (2) In TCP, packet losses may be detected by a timeout or by receiving three duplicate acknowledgements. Both events will trigger retransmissions. Why would the TCP use duplicate acknowledgements to trigger retransmissions instead of simply waiting for timeout? Is it OK to remove the timeout mechanism and use three duplicate acknowledgements as the only way to detect packet loss? (10%)
   (3) Again, packet losses may be detected by a timeout or by receiving three duplicate acknowledgements. TCP's congestion control mechanism reduces window size to 1 when the loss is detected by timeout and to its half when the loss is detected by duplicate acknowledgements. Why do you think TCP is designed with such difference? (10%)
   (4) The flow control and congestion control mechanisms in TCP are both to control the sending window size of the source. Which is the one that prevents the source from overwhelming the receiver's buffer? Which is the one that prevents the source from overwhelming the routers' buffer in the network? (10%)
   (5) Is the 3-way handshake mechanism for TCP's connection establishment or connection teardown? Why isn't it 2-way or 4-way? (10%)

Sample Solution:

(1) If too short: unnecessary retransmission
    If too long: take a long time to recover the loss
(2) To speed up the retransmission
    No. When the congestion window size is small, there might not be enough packets following the lost one to trigger the sending of 3 duplicate acknowledgements
(3) In the duplicate acknowledgements case, the receiver can still receive a good amount of packets (at least 3) following the lost one. Therefore, the source slows down but not to 1.
    In the timeout case, the receiver has trouble getting enough packets after the loss. It is a stronger indication of the network being busier. Therefore, the source shall slow down faster to 1.

(4) Flow control
Congestion control

(5) Connection establishment
Why not 2-way: TCP is full-duplex. Both sides need to communicate with the others the initial window size. This information is carried in the initial packet (SYN packet) and should be acknowledged (ACK packet). There should at least be 4 exchanges of the packets
Why not 4-way: the 1st ACK and the 2nd SYN can be piggybacked to save time.

6. (Transport) When the cwnd is smaller than or equal to the slow start threshold (ssthresh), the TCP source is in the slow start state. When the cwnd is larger than the ssthresh, the TCP enters the congestion avoidance state. TCP sends the packets exponentially fast when it is in the slow start state vs. linearly fast in the congestion avoidance state. Upon receiving a new acknowledgement, the pseudo code for the increment of the cwnd is provided as follows:

If (cwnd ≤ ssthresh)
then cwnd = cwnd + 1
else cwnd = cwnd + 1/cwnd

sliding window size = round(cwnd)

Suppose there is a connection with initial cwnd = 1 and ssthress = 2, as shown in the figure on the right. There are only 8 packets to send. The cwnd will become 2 when the source node receives the acknowledgement for packet 1. As a result, the source will send packet 2 and 3 at once. When the source receives the acknowledgement for packet 2, the cwnd will be 3. The source, in turn, sends packet 4 and 5. When it receives the acknowledgement for packet 3, the cwnd will be 3+1/3. The source then sends packet 6.

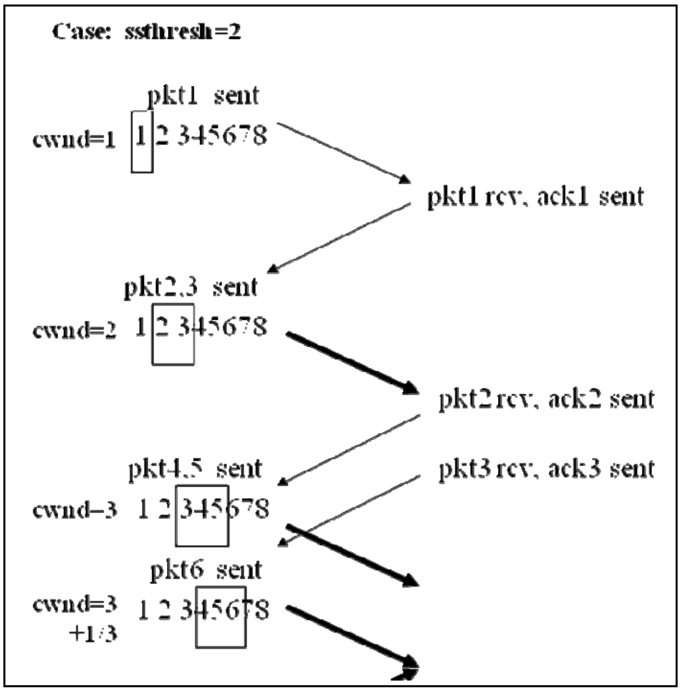

(1) Suppose the initial cwnd = 1 and ssthresh = 4. Can you tell us what the cwnd will be when the acknowledgement for packet 3 is received at the source and which packet(s) will be sent next? (10%)
(2) Following question (1), if packet 6 is lost, Can you tell us what the cwnd will be after the loss is detected and which packet(s) will be sent next? Assume that TCP receiver does not buffer packets out of order and retransmits with 3 duplicate acks. (10%)

(3) Following question (1), if packet 4 is lost, Can you tell us what the cwnd will be after the loss is detected and which packet(s) will be sent next? Assume that TCP receiver does not buffer packets out of order and retransmits with 3 duplicate acks. (10%)

Sample Solution:

(1) cwnd=4, packet 6, 7
(2) cwnd=1, packet 6
(3) cwnd=2, packet 4, 5