Name________________ Student ID______________________ Department/Year___________

Midterm Examination

Introduction to Computer Networks
Class#: 901 E31110
Fall 2008

9:30-11:10 Tuesday
November 11, 2008

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 figures similar to the one on the right. It illustrates a typical composition of a network on the Internet.

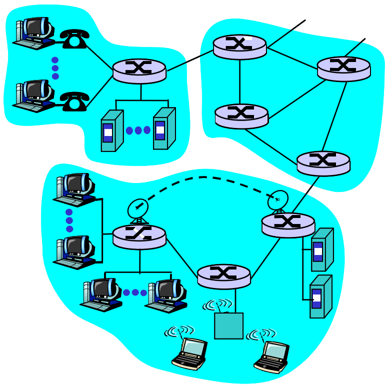

   (1) Based on how you understand what the Internet edge, core, and access network is, circle out the computers belonging to the edge of the network, the computers belonging to the core of the network, and the links belonging to the access network. (5%)
   (2) Name two kinds of Access Network. (5%)
   (3) Describe the Access Network you use at home or student dorm. Be specific about the uplink and downlink bandwidth and the approximate (monthly/yearly) cost. (5%)

Sample Solution:

   (1) As depicted

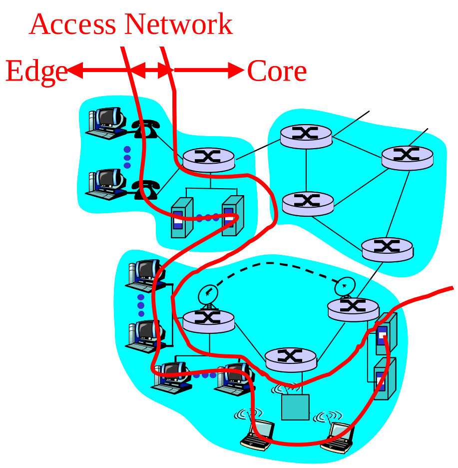

   (2) ADSL, Cable, Ethernet, WiFi, 3G
   (3) Depending on what you have, describe what you know.

2. (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.

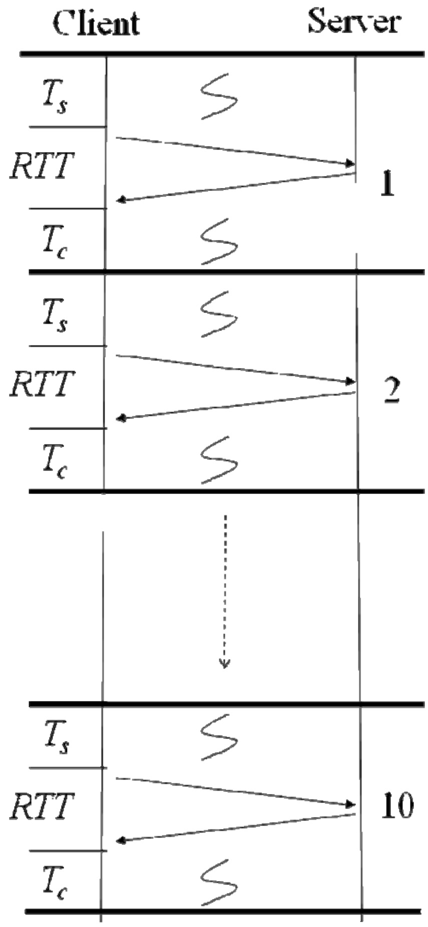

(1) What will be the time it takes to complete transferring the web page running in the persistent without pipelining mode? (5%)

(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%)

(3) Continue from (2). Now suppose the html main page, 8 jpg image and the mp3 music objects are located on 10 different Web servers. The Web client can only have one download connection at a time. Assume the delay from the client to all these different Web servers are the same – *RTT*. What will be the entire page download time? (5%)

(4) Continue from (3). If the main page and 7 of the jpg image objects are from the same server, the other two objects are from another server, what will be the entire page download time? (5%)

(5) Continue from (4). If the main page and 7 of the jpg image objects are from the same server, the other two are from two other different servers, what will be the entire page download time? (5%)

(6) Continue from (5). To generalize the page download time for N objects spreading across K servers, let's denote the number of objects on server i as N_i. I.e., $N=N_1+N_2+N_3\ldots+N_k$. Let $N_i > 1$ for all i. Can you derive the formula for the entire page download time using T_s, T_c, RTT, K and N_i. (10%)

(7) Continue from (6). If you are to create a page that includes a number of foreign objects, would you like to download these foreign objects first to your host Web server? Why? (5%)

Sample Solution:

(1) $T_s + 10RTT + T_c$

(2) $T_s + 2RTT + T_c$

(3) $10T_s + 10RTT + 10T_c$

(4) $2T_s + 3RTT + 2T_c$

(5) $3T_s + 4RTT + 3T_c$

(6) For the main server, $T_s + 2RTT + T_c$. For each of the other servers, $T_s + RTT + T_c$. Total: $KT_s + (K+1)RTT + KT_c$.

(7) Yes. That will keep all objects on the host Web server and at least save the extra T_s and T_c time required for establishing and closing connections to the foreign servers.

3. (Application) Let's assume a network of 7 users and 1 source. The source has the file that the users desire. To distribute this file to all users, the simplest way is to have all 7 users download the file one by one from the source. This client-server way of file sharing is depicted in the left figure below. Assume (1) each user allows one upload at a time and (2) the file download time between any two nodes is the same, denoted as *T*. The time to distribute the file to all users will be 7*T*.

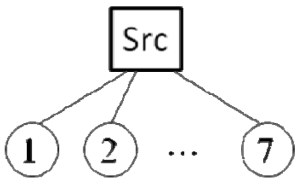

Client-Server

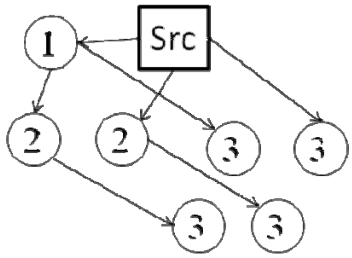

Peer-to-Peer

The other way is to form a peer-to-peer network as the figure on the right to speed up the distribution of the file to all users. In that, after the first user completes the download from the source, it can help distribute the file to other users. After time *T*, two users can download simultaneously from the source and the first users, and so on. After time *3T*, 7 of the users will be able to complete the download the file.

(1) Suppose there are now 15 users. How much time will be necessary to complete the distribution too all users taking the client-server approach and how much time will be taking the peer-to-peer approach? (5%)
(2) Generalize the above cases to a network of 2^K-1 users. How much time is required to distribute the file taking the client-server approach and how much time taking the peer-to-peer approach? (5%)
(3) Can you comment on which approach will be better in terms of the distribution time when the size of the network grows? (5%)

Sample Solution:

(1) 15*T* vs. 4*T*
(2) (2^K -1)*T* vs. *KT*
(3) If the network size is N, the required distribute time in client-server model is at the scale of O(N) and in the peer-to-peer model O(lgN). The latter is more scalable.

4. (Application) Continue from Problem 3. Instead of downloading a whole file each time, divide the file into 4 equal segments and download one segment at a time. The time to download a segment will be $T/4$. Taking the peer-to-peer approach, now pipeline the sending of segments from the source to the 15 users.
   (1) How much time does it take to distribute the file over all 15 users? (5%)
   (2) Divide the file into S segments. Pipeline the segments. How much time does it take to distribute the file over all 15 users? (5%)
   (3) Consider pipelining S segments over 2^K-1 users. How much time does it require? (5%)
   (4) Continue from (3). Comment on the performance in terms of download time when S is large. (5%)
   (5) In practice, the time to download a segment is not quite T/S. There is usually extra connection establish delay which makes the segment download time $T/S+C$. If this is the case, how much time does it take to distribute the file in S segments over 2^K-1 users? (5%)
   (6) Continue from (5). Comment on the performance in terms of download time when S is large. (5%)

Sample Solution:

(1) For the first segment to reach all: 4 * $T/4$
For the following 3 segments to reach the last nodes: 3 * $T/4$
Total: $7T/4$

(2) For the first segment to reach all: 4 * T/S
For the following S-1 segments to reach the last 8 nodes: (S-1) * T/S
Total: $(4+S-1)T/\mathrm{S}$

(3) For the first segment to reach all: K * T/S
For the following S-1 segments to reach the last 8 nodes: (S-1) * T/S
Total: $(K+S-1)T/\mathrm{S}$

(4) $(K+S-1)T/S = (K-1)T/S+T$
If S increases, the distribution time will decrease. Dividing a file into more segments benefits the distribution time.

(5) For the first segment to reach all: K * ($T/S+C$)
For the following S-1 segments to reach the last 8 nodes: (S-1) * ($T/S+C$)
Total: ($K+S$-1) * ($T/S+C$)

(6) When S increases, the distribution time might not decrease. Take S towards infinity. The distribution time tends to infinity. In practice, dividing a file into unnecessarily many segments does not help.

5. (Transport) Let us compare and contrast Go-Back-N (GBN) and Selective Repeat (SR) quantitatively. Let the window size be W.
   (1) How many timers are required at the sender side of GBN and how many timers at the sender side of SR? (5%)
   (2) How much buffer is required at the receiver side of GBN and how much at the receiver side of SR? (5%)
   (3) Suppose W is an even number. If packet W/2 is lost, what packet(s) will be retransmitted in GBN and what packet(s) will be retransmitted in SR? (5%)
   (4) Suppose W is an even number. If packet W/2 is lost, how much time does it take to have the first W packets arriving at the receiver in GBN and how much time in SR? Let packet transmission time be T, timeout interval be TO, and round-trip time be RTT. Assume the path is symmetric where RTT equals twice the one-way delay. Assume also TO is greater than the time to send the whole window of packets WT plus the RTT. The transmission time of acknowledge packets is negligible. Denote the time to have all first W packets arriving at the receiver by the packet transmission time (T), timeout (TO), and round-trip time (RTT). (5%)
   (5) If memory is expensive and network bandwidth is abundant, which mechanism, GBN or SR, will you choose to transfer data? Why? (5%)
   (6) If your users care the most about download delay, which mechanism, GBN or SR, will you choose to transfer data? Why? (5%)

Sample Solution:

(1) GBN: 1 timer
    SR : W timers
(2) GBN: 1 packet
    SR : W packets
(3) GBN: W/2, W/2+1, … W
    SR : W/2
(4) GBN: (W/2-1) T + RTT + TO +1 T +W/2 T + RTT/2
    The timer for packet W/2 starts when the ack for W/2-1 arrives: (W/2 -1)T+RTT
    After timeout, the sender retransmits packet W/2 and the rest of the packets.
    These retransmitted packets will arrive after RTT/2 time.
    SR : W/2 T + TO +1 T + RTT/2
    The timer for packet W/2 starts after the packet W/2 is sent: W/2T. After timeout, retransmit packet W/2 which will arrive after half the round-trip time
(5) GBN: memory requirement on the end hosts is very small
(6) SR: the delay is shorter when there are losses

6. (Transport) For pipelined reliable transfer protocols, one needs to be careful in setting the window size. Let the end-to-end delay be D seconds, the bandwidth B bps, and the packet size P bytes.

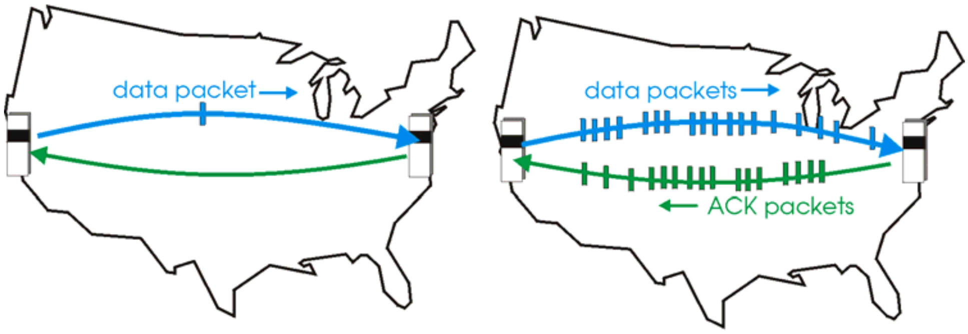

(a) a stop-and-wait protocol in operation (b) a pipelined protocol in operation

(1) The stop-and-wait protocol is essentially setting the window size to 1. Is it good to set the window size to 1 packet? Why? (5%)
(2) To maximize the utilization, one can set the window size to be very large. Is it good to set the window size to be infinitely large? Why? (5%)
(3) What value will be appropriate to set for the window size if there is no other traffic on the way? Why? (5%)
(4) What value will be appropriate to set for the window size if there will be some other traffic traveling the same path? Why? (5%)

Sample Solution:

(1) Not good in terms of utilization. The utilization of the bandwidth will be low.
(2) Not good either. Although the utilization will be high but the network cannot accommodate that much data which will result in packet drops.
(3) When there is no other traffic, the window size can be set of the product of the end-to-end delay and the bandwidth (BM/8P). This way the utilization will be 1 and there will be no drops incurred.
(4) When there is other traffic, the window size should be set of the bandwidth-delay product minus the traffic to make sure the utilization is high and there will be no drops. A better mechanism should dynamically adapt the window size depending on the amount of cross traffic.