Computer Networks (CS 453), Spring 2011
Homework 2
Instructor: V. Arun

Assigned: 2/14/11, Due: 2/23/11

Note: Please show as much of your work as you can. *Whenever possible, use variable names before plugging in numerical values*. Even if you get the answer wrong, you can get partial credit if you show your approach clearly. It will help us tell you where you made a mistake. If you plug in numbers right upfront and your answer is wrong, you will not get partial credit.

Problem 1: Quickies (36 points)

A. What are two fundamental advantages of peer-to-peer (P2P) systems over client-server systems.
B. What is one similarity and one difference between connection management in HTTP/1.0 and FTP?
C. POP3 does not maintain state across sessions but IMAP does. True or False? Explain your answer.
D. What protocol is being used between your browser and a mail server such as Hotmail or Gmail?
E. HTTP/1.1 reduces the response time compared to HTTP/1.0 most for web pages with [many/few] [large/small] objects. Pick the best of the four possible options.
F. List two benefits of Web caching that do not have analogues in caching in memory hierarchies on a single machine.
G. Skype uses P2P techniques for two key functions. What are they?
H. Does DNS use TCP or UDP? Why?
I. What information is used by a process running on a host to identify a process running on another host?
J. Kazaa is a pure P2P system while BitTorrent is a hybrid P2P system. True or false? Explain.

Problem 2 (Web caching, 24 points): Consider the figure used to illustrate Web caching (Slide 2-42). Suppose that the average object size is S=850,000 bits and that the average request rate from the institution's browsers to the origin servers is A = 16 requests per second. Also suppose that the amount of time it takes from when the router on the Internet side of the access link forwards an HTTP request until it receives the response is T = 3 seconds on average. Model the total average response time perceived by the institutional users as the sum of the average access delay (that is, the delay from the Internet router to institution router) and the average Internet delay. For the average access delay, use D/(1-DB), where D is the average time required to send an object over the access link and B is the arrival rate of objects to the access link. (Note that B = A if there is no caching. Ignore details of HTTP 1.0 vs. 1.1.)

A. Find the total average response time.
B. Now suppose a cache is installed in the institutional LAN. Suppose the miss rate is p=0.3. Find the total response time.

Problem 3 (DNS, 10 points): Suppose within your Web browser you click on a link to obtain a Web page. The IP address for the associated URL is not cached in your local host, so a DNS lookup is necessary to obtain the IP address. Suppose that *n* DNS servers are visited before your host receives the IP address from DNS; the successive visits incur an RTT of T_1, T_2, …, T_n, Further suppose that the Web page associated with the link contains exactly one object, consisting of a small amount of HTML text. Let T_0 denote the RTT between the local host and the server containing the object. Assuming zero transmission time of the object, how much time elapses from when the client clicks on the link until it receives the object?

Problem 4 (Digging around, 30 points):

A. Problem P15 from Chapter 2 in the textbook.
B. Problem P19 from Chapter 2 in the textbook.
C. Problem P35 from Chapter 2 in the textbook.