CS 5480/6480: Computer Networks – Spring 2012
Homework 1
Due by 9:00 AM MT on January 31st 2012

Important:
- **No cheating will be tolerated.**
- **No extension.**

CS 5480 total points = 32
CS 6480 total points = 44

Question 1 (Packet Switching versus Circuit Switching) ***5 points*****:** Suppose 40 users share a 1Mbps link. Also suppose that each user alternates (independently of the other users) between periods of activity, when the user generates data at a constant rate of 100 Kbps and periods of inactivity when the user generates no data. Suppose further that the user is active (independent of the other users) only 10 percent of the time.

1. (*3.5 points*) When circuit switching is used to allocate resources on the shared link, 10 users can be supported. When packet switching is used, what is the probability that 11 or more users are active simultaneously? (You must show the steps for full credit.)
2. (*1.5 points*) Do you expect the result of part 1 to change when all the 40 users synchronize in their use of the shared link (i.e. all the 40 users generate data or are idle exactly at the same time)? Explain with arguments. No calculations are required for this part of the question.

Question 2 (Internet Delay Components) ***3 points*****:** The figure below shows a router and two links which are part of a bigger network.

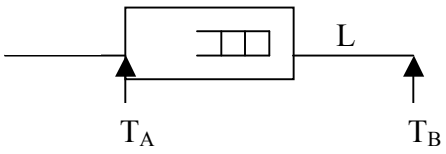

Every time a packet is received at the router it is time-stamped (T_A). It is also times-stamped when it is about to leave link L (T_B). When 5 packets of the size 500 bytes are transmitted through the link L, ($T_B - T_A$) is measured to be 10 ms, 2.8 ms, 2.4 ms, 4 ms, and 5.5 ms. When 5 packets of the size 1000 bytes are transmitted through the link L, ($T_B - T_A$) is measured to be 11.0 ms, 10 ms, 2.8 ms, 3.0 ms, and 5.5 ms. Assume that processing delay at the router is negligible. What is the average queuing delay experienced by the 1000 byte packets? What are the reasonable estimates of transmission and propagation delays experienced by a packet of size 600 bytes sent through the link L?

Question 3 (Traceroute) *6 points*:

(a) (*3 points*) Execute the *traceroute* command to two destinations of your choice, at least 12 hops away, from a source. Compute the average delay (averaged over the three delay values) for each hop and plot it in a graph with x-axis showing the hop number and y-axis the average delay corresponding to that hop. Run the same experiment a few hours later and show the new results on the same graph. In all, your graph should have four curves. How many hops are common along the paths to the two destinations? Attach the *traceroute* outputs. (If the traceroute command does not work on your machine, try using the service at traceroute.org.)

(b) (*2 points*) Suppose one of the three *traceroute* delay values between the source and a given router hop turns out to be unusually high. What are *two* possible causes for this unusually high delay?

(c) (*1 point*) How would you change the *traceroute* program to find the IP address of every third hop instead of every hop (i.e., it finds the address of the 3rd hop, the 6th hop, …, you can assume that the destination is a multiple of 3 hops away from the source)?

Question 4 (Web):

(a) (*2.5 points, 0.5 point for each question*) Consider the following string of ASCII characters that were captured by Wireshark when the browser sent an HTTP GET message (i.e., this is the actual content of an HTTP GET message). The characters *<cr><lf>* are carriage return and line-feed characters (that is, the italized character string *<cr>* in the text below represents the single carriage-return character that was contained at that point in the HTTP header). Answer the following questions, indicating where in the HTTP GET message below you find the answer.

GET /cs453/index.html HTTP/1.1*<cr><lf>*Host: gaia.cs.umass.edu*<cr><lf>*User-Agent: Mozilla/5.0 (Windows;U; Windows NT 5.1; en-US; rv:1.7.2) Gecko/20040804 Netscape/7.2 (ax) *<cr><lf>*Accept:ext/xml, application/xml, application/xhtml+xml, text/html;q=0.9, text/plain;q=0.8,image/png,*/*;q=0.5*<cr><lf>*Accept-Language: en-us,en;q=0.5*<cr><lf>*Accept-Encoding: zip,deflate*<cr><lf>*Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7*<cr><lf>*Keep-Alive: 300*<cr><lf>*Connection:keep-alive*<cr><lf><cr><lf>*

a. What is the URL of the document requested by the browser?
b. What version of HTTP is the browser running?
c. Does the browser request a non-persistent or a persistent connection?
d. What is the IP address of the host on which the browser is running?
e. What type of browser initiates this message? Why is the browser type needed in an HTTP request message?

(b) (*2 point, 0.5 point for each question*) The text below shows the reply sent from the server in response to the HTTP GET message in the question above. Answer the following questions, indicating where in the message below you find the answer.

HTTP/1.1 200 OK*<cr><lf>*Date: Tue, 07 Mar 2008 12:39:45GMT*<cr><lf>*Server: Apache/2.0.52 (Fedora)<cr><lf>Last-Modified: Sat, 10 Dec2005 18:27:46 GMT*<cr><lf>*ETag: "526c3-f22-a88a4c80"*<cr><lf>*Accept-Ranges: bytes*<cr><lf>*Content-Length: 3874*<cr><lf>* Keep-Alive: timeout=max=100*<cr><lf>*Connection:Keep-Alive*<cr><lf>*Content-Type: text/html; charset=ISO-8859-1*<cr><lf><cr><lf>*<!doctype html public "-//w3c//dtd html 4.0 transitional//en"><*lf*><html><*lf*><head><*lf*> <meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1"><*lf*> <meta name="GENERATOR"

content="Mozilla/4.79 [en] (Windows NT 5.0; U) Netscape]"><*lf*> <title>CMPSCI 453 / 591 /NTU-ST550A Spring 2005 homepage</title><*lf*></head><*lf*><*much more document text following here (not shown)*>

a. Was the server able to successfully find the document or not? What time was the document reply provided?
b. When was the document last modified?
c. How many bytes are there in the document being returned?
d. What are the first 5 bytes of the document being returned? Did the server agree to a persistent connection?

Question 5 (DNS) (*3.5 points, 0.5 point for each part*):
a. What is a *whois* database?
b. Use various whois databases on the Internet to obtain the names of two DNS servers. Indicate which whois databases you used.
c. Use nslookup on your local host to send DNS queries to three DNS servers: your local DNS server and the two DNS servers you found in part (b). Try querying for Type A, NS, and MX reports. Summarize your findings.
d. Use nslookup to find a Web server that has multiple IP addresses. Does the Web server of your institution (school or company) have multiple IP addresses?
e. Use the ARIN whois database to determine the IP address range used by your university.
f. Describe how an attacker can use whois databases and the nslookup tool to perform reconnaissance on an institution before launching an attack.
g. Discuss why whois databases should be publicly available.

Question 6 (P2P) :
(a) (*4 points)* Consider distributing a file of F bits to N peers using a client-server architecture. Assume a fluid model where the server can simultaneously transmit to multiple peers, transmitting to each peer at different rates, as long as the combined rate does not exceed u_s.
a. Suppose that $u_s/N \le d_{\min}$. Specify a distribution scheme that has a distribution time of NF/u_s.
b. Suppose that $u_s/N \ge d_{\min}$. Specify a distribution scheme that has a distribution time of $F/ d_{\min}$.
c. Conclude that the minimum distribution time is in general given by $\max\{NF/u_s, F/ d_{\min}\}$.

(b) (*4 points*) Consider distributing a file of F bits to N peers using a P2P architecture. Assume a fluid model. For simplicity, assume that $d_{\min}$ is very large, so that peer download bandwidth is never a bottleneck.
a. Suppose that $u_s \le (u_s + u_1 + ... + u_N)/N$. Specify a distribution scheme that has a distribution time of F/u_s.
b. Suppose that $u_s \ge (u_s + u_1 + ... + u_N)/N$. Specify a distribution scheme that has a distribution time of $NF/(u_s + u_1 + ... + u_N)$.
c. Conclude that the minimum distribution time is in general given by $\max\{F/u_s, NF/(u_s + u_1 + ... + u_N)\}$

(c) (*2 points*) Consider an overlay network with N active peers, with each pair of peers having an active TCP connection. Additionally, suppose that the TCP connections pass through a total of M routers. How many nodes and edges are there in the corresponding overlay network?

Question 7 (required for CS6480, extra credit for CS5480) *12 points*: Read the following paper – "Internet Indirection Infrastructure" by I. Stoica, D. Adkins, S. Zhuang, S. Shenker and S. Surana, ACM Sigcomm Conference, August 2002.
(http://eng.utah.edu/~cs5480/readings/i3-sigcomm.pdf)
Answer the following questions that are based on this paper:

1. (*3 points*) Consider a model in which data sent from two senders is added together before being sent to a receiver. Show and describe how this could be implemented using the *i3* architecture.
2. (*3 points*) Consider a model in which multimedia MPEG data sent from a sender must be converted to the JPEG format before it is received at a receiver. The conversion operation is specified by the sender. The receiver wants to ensure that all the data it receives comes through a firewall. Show and describe how this could be implemented using the *i3* architecture.
3. (*3 points*) Explain how mobility can be supported in the *i3* architecture.
4. (*3 points*) Can you use the *i3* architecture for preventing Denial-of-service attacks on servers? Discuss.