CS 536 Fall 2008 - Homework 1

Problem 1 – 15pts

Statistical multiplexing. Suppose users share 1 Mbps link. Also suppose each user requires 100 Kbps when transmitting, but each user transmits only 10% of the time.

- When circuit switching is used, how many users can be supported ?
- For the remainder of this problem, suppose packet switching is used. Find the probability that a given user is transmitting.
- Suppose now there are 40 users. Find the probability that at any given instant, exactly n users are transmitting. (Hint: Use the binomial distribution)
- Find the probability that there are 11 ore more users transmitting simultaneously.

Problem 2 – 10pts

Queuing delay. Consider the queuing delay in a router buffer (preceding an outbound link). Suppose all packets are L bits, the transmission rate is R bps, and that N packets simultaneously arrive at the buffer every LN/R seconds. Find the average queuing delay of a packet. (Hint: The queuing delay for the first packet is zero; for the second packet L/R; for the third packet it is $2L/R$. The Nth packet has already been transmitted when the second batch of packets arrive.)

Problem 3 – 15pts

Web caching. Consider an institutional network that is connected to the Internet with an access link of capacity 15 Mbps. Suppose that the average object size is 900,000 bits and that the average request rate from the institution's browsers to the origin servers is 15 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 two seconds on average. Model the total average response time as the sum of the average access delay (i.e., the delay from the Internet router to the institution router) and the average Internet delay. For the average access delay, use $\delta/(1-\delta\beta)$, where δ is the average time required to send an object over the access link and β is the arrival rate of objects to the access link.

- Find the total average response time.
- Now suppose a cache is installed in the institutional LAN and the hit-rate of the cache is 0.4, find the total response time.

Problem 4 – 10pts

HTTP. Consider a short, 10 meter link, over which a sender can transmit at a rate of 150 bits/sec in both directions. Suppose that packets containing data are 100,000 bits long, and packets containing only control (eg. ACK or handshaking) are 200 bits long. Assume that N parallel connections each get 1/N of the link bandwidth. Now consider the HTTP protocol, and suppose that the object is 100Kbits long, and the initial downloaded object contains 10 referenced objects from the same sender. Would parallel downloads via parallel instances of non-persistent HTTP make sense in this case ? Now consider persistent HTTP. Do you expect significant gains over the non-persistent case ? Justify and explain your answer.

Problem 5 – 20pts

Web access and DNS. 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 5 DNS servers need to be visited before your host receives the IP address from DNS; the successive visits incur an RTT of $10, 20, ..., 50$ milliseconds. Further, suppose that the Web page associated with the link contains exactly one object, consisting of a small amount of HTML text.

- Let the RTT between the local host and the server containing the object be say 5 seconds. Assuming zero transmission time of the object, how much time elapses from when the client clicks on the link until the client receives the object.
- Suppose there are 100 embedded objects within the web page that need to be fetched, and these objects need to be fetched from exactly 10 different Web servers. Suppose none of domains names are cached in the local DNS server and all of them need to go through the same path of DNS servers as before to get resolved. How much time elapses from when the client clicks on the link until the entire Web page is rendered. Assume that the browser does not use either persistent connections or pipelining in this

case. Also, each object requires exactly 10 seconds to download once a connection is established with the server. Assume the network is not the bottleneck for downloading these objects.

- If the browser uses persistent connections, how much time will elapse before the page is completely rendered.
- If the browser uses persistent connections with a maximum of up to 5 parallel downloads, how much time would elapse.

Problem 6 – 15pts

Peer to Peer File Sharing. Consider distributing a file of $F = 10$Gbits to N peers. The server has an upload rate of $u_s = 20Mbps$ and each peer has a download of $d_i = 1Mbps$ and an upload rate of u. For $N = 10$,100, and $1,000$ and $u = 200$Kbps, 600 Kbps and 1 Mbps, give the minimum distribution times for each combination of N and u for both client-server distribution and P2P distribution.

Problem 7 – 15pts

P2P Query Flooding. Consider query flooding as we have discussed in the class. Suppose that each peer is connected to at most 10 neighbors in an overlay network that contain about 10,000 nodes. Also suppose controlled flooding is employed with each query stopped after 5 nodes.

- What is the upper bound on the number of messages that are sent into the overlay network for each query that a node makes ?
- What is the lower bound ?
- Suppose we use super nodes which have the knowledge of all objects in the network. And, let's say about 1 in 100 nodes are designated super nodes. A query can be satisfied by an ordinary node with a probability of 1 in 10,000. What is the expected number of messages that need to be injected into the p2p network before a query can be answered.