15-440/640 Distributed Systems
Homework 1

Due: September 29, 2015, 12:00 PM (in class)

Name:
Andrew ID:

September 16, 2015

1. Blogger Bob decides he wants to set up streaming video. You can treat this video as though it were already stored (like an image). What is the max video bit rate (in bits/s) he can get given the following conditions? Assume the best case transmission time, where there is no processing being done at the packet's destination. Also assume that RTT is symmetric, so if RTT = 200 ms, the time taken in each direction is 100 ms.

   1 Mbps = 10^6 bits/s

   1 Kbytes = 2^{10} bytes

   (a) RTT (Round Trip Time) = 200 ms. Packet size = 10 Kbytes, Bandwidth = 2 Mbps. Packets are sent continuously, as in UDP.

   (b) RTT = 200 ms. Packet size = 10 KBytes, Bandwidth = 2 Mbps. After each packet is sent we must wait one RTT (waited for an ACK/Acknowledgement).

   (c) RTT = 200 ms. Packet size = 10 KBytes, Bob is provided with infinite bandwidth, but only up to 20 packets can be sent per RTT. Packets are sent continuously.

   (d) RTT = 200 ms. Packet size = 10 KBytes, Bandwidth = 2 Mbps. Video filesize = 100 MB. Bob is experimenting with a new way of sending data. He sends 1 packet and waits for an ACK, then 2 packets and waits for an ACK, then 3 packets, and so on until all the data is sent. He stops increasing number of packets after max bandwidth is reached.

   (e) What is the minimum number of packets that Bob could send before waiting for an ACK and still achieve the full bandwidth? RTT = 200 ms, Packet size = 10 KBytes, Bandwidth = 2 Mbps. Packets are sent continuously.

2. Jimmy, a student in distributed systems at Carnegie Mellon, wants to build a distributed file system. Aware that this is a major undertaking, he decides to spend some time designing the system before implementing. For his first version, Jimmy assumes that the system initially doesn't contain any files, all files are stored in the same directory, and clients should be unaware of files created by other clients. Being a good friend, you decide to help Jimmy.

   a) In the following client-server interaction, please write down where unexpected behavior could occur. Also describe how Jimmy should build the system so that clients are unaware of each other's files.

   1. Client A creates and opens file1 for writing
   2. Client A writes to file1
   3. Client A loses connection to the server
   4. Client A writes to file1

5. Client A regains connection to the server
6. Client A closes file1
7. Client A opens file2 for reading
8. Client B creates and opens file2 for writing
9. Client B's system crashes, and is restarted

b) Suppose Jimmy decides that he wants clients to be able to read files created by other clients, but not to write to them. How might your design change?

3. Every distributed system faces trade offs between consistency, availability, and partition tolerance or resistance.

a) Define the three terms.

b) It is said you can satisfy at most two out of the three aforementioned goals. Why is this the case?

c) Which of these traits might an online banking system optimize? Why?

d) Which of these traits might facebook optimize? Why?

e) Which of these traits might an airport departure system optimize? Why?

4. In the diagram below, the links are labeled with (latency (in ms), bandwidth (in Mbps)). Also assume that data transmission occurs without loss of data, and sources transmit at rates such that no queueing results.

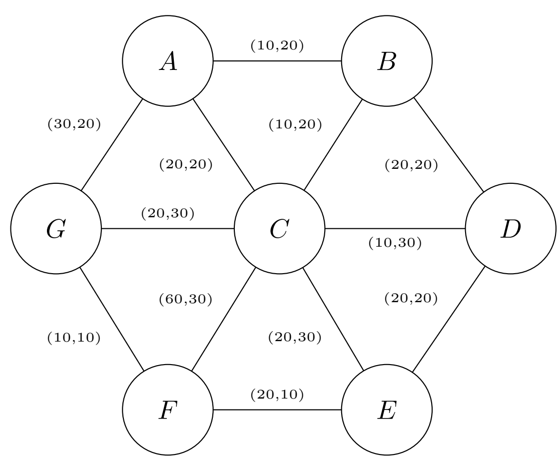

a) Given this network topology, what path gives the lowest possible end-to-end latency from node B to node F?

b) What is the bandwidth of this route?

c) What would be the bandwidth of this route if data is also being sent from node A to node F (along the route with lowest latency)? Assume that, as in TCP, bandwidth gets allocated equally to both data streams.

d) What is the latency of this route if data is also being sent from node A to node F?

e) What path gives the highest bandwidth for transmitting data from node B to F?

f) What is the latency of this path?

g) What path gives the lowest possible latency between nodes B and F if C fails?

h) What is the bandwidth of this path?

i) What is easier to improve, bandwidth or latency? Give at least two reasons for your choice.

5. In order to store data from a client to a server, the message needs to cross a a sea of routers. Typically a router doesn't store any messages which means they do not ensure that subsequent routers receive the message. This requires the client to write logic to figure out whether or not the message it sent actually reached its destination (TCP or UDP). Imagine you are an angel investor and a start-up proposes the concept of reliable routers have the responsibility of making sure the data they receive is transmitted to the next reliable router (and so on until the message reaches its destination). The selling point of this start-up is that their product allows for a simplification of client code. Is this a good investment? Why or why not?

6. a) The RPC paradigm is useful in many situations. What are some of the motivations for using RPCs?

   b) RPC's are supposed to be transparent. Describe two ways in which this is not true.

   c) Describe two ways in which an application may need to handle this lack of transparency.