CMPT 371: Final Exam

August 15, 2006

1 Short Answer Questions

1. For each of the top four layers of the protocol stack, name one specific protocol that operates at that level.
2. True or False:
   (a) A cookie is the information stored on a web-server about its clients.
   (b) Binary files can be sent unencoded over SMTP.
   (c) A Top-Level Domain (TLD) DNS server can make a query to an Authoritative DNS server.
   (d) The receive function in java's interface for UDP never times out.
   (e) UDP's checksum sometimes makes mistakes.
   (f) Ethernet's checksum sometimes makes mistakes.
   (g) Temporary congestion in a network will have a strong long-term effect on the TCP timeout values of its hosts.
   (h) IP datagram fragmentation occurs because of link-layer limitations.
   (i) The Distance-Vector algorithm always passes fewer messages than the Link-State algorithm.
   (j) An ISP cannot prevent IP datagrams from passing through its routers on the way to their final destinations.
   (k) It is sometimes preferable to implement reliable data transfer at the link layer rather than the transport layer.
   (l) Pure Aloha is more efficient than slotted Aloha.
   (m) PPP is a multiple access protocol.
   (n) CMPT 371 is the best course ever!
3. Name two ways to recycle IP addresses so that every host doesn't need its own permanent address.
4. (a) Is it possible to build a large network with no routers, only switches?
   (b) What would be the advantages of a large, switch-only network?
   (c) What is the advantage of using routers?
5. Recall that in BGP a router may learn of multiple routes to the same network prefix. Name three rules BGP might use for eliminating multiple paths.
6. Briefly describe what IP tunneling is for and how it works.
7. If two TCP connections share a certain link, under what conditions will their throughputs along that link converge to the same value? How does this convergence happen?

2 Longer Questions

1. Consider the ethernet configuration below. Let the propagation rate of the links be 2×10^8 meters/sec. Let the transmission rate be 100 Mbps. Also assume each link has length 100 meters.

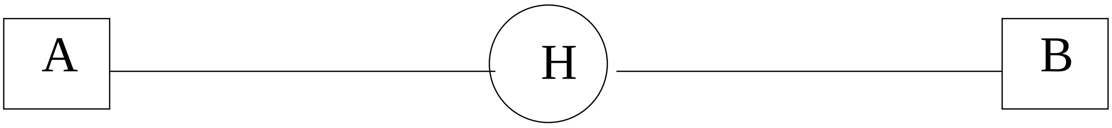

   (a) How long is a "bit-time" in seconds?

   (b) If node A transmits a bit, how long does it take to propagate to node B (in bit-times)?

   (c) If A and B each start transmitting a frame of length 512 bits at the same time, will they notice a collision before they are done transmitting?

   (d) If A and B transmit at the same time and detect collisions at the same time before they are finished transmitting and they pick different values of K in the first round of retransmission, will one of them succeed in retransmitting during that round?

   (e) How many rounds of retransmitting are necessary before one of the two has succeeded with probability at least 99%?

2. This question is about TCP congestion control. Assume there is a sender and receiver connected by a link with transmission rate R bits/second. The sender has a continuous supply of segments to transmit. The roundtrip time (from when a segment is finished being transmitted by the sender until an ACK is received by the sender) is RTT. Each segment has size L bits.

   (a) Let the sender's window be of size W segments. Recall that the sender's throughput is the average number of bits the sender sends per second. If $WL/R < RTT + L/R$, then what is the sender's throughput?

   (b) If $WL/R \geq RTT + L/R$, what is the sender's throughput?

   (c) Assume there is a value W_{loss} such that, whenever the sender's window reaches size W_{loss} a segment gets dropped and whenever the sender's window is smaller than W_{loss} no segments are dropped. Draw a graph of the sender's window size W (y-axis) versus time (x-axis) that includes both the slow-start and AIMD phases of congestion control. Label important points on each axis, including units.

   (d) Briefly explain why it should be the case that $W_{loss}L/R \leq RTT + L/R$.

   (e) What is the sender's throughput during the AIMD phase from part (c)?

   (f) What is the sender's throughput during the slow-start phase from part (c).

3. In this question, you will design a network application. Most unix systems have a "live-chat" application called `talk`. User 1 can invoke `talk` by typing `talk user2@remotehost`. If User 2 is logged in to `remotehost`, then a message appears on her console saying that User 1 is requesting a talk. If User 2 accepts, then the two users enter a talk session. During the talk session, whatever user 1 types should appear as quickly as possible on both user 1's screen and user 2's screen; likewise for user 2. There should never be any difference between what user 1 sees and what user 2 sees on their screens. Either user can end the talk session by typing 'ctrl-D'.

(a) What is the difference between the client and server in this application?

(b) Briefly discuss why this application might want to use UDP.

(c) Briefly discuss why this application might want to use TCP.

(d) Choose either TCP or UDP. Describe the types of messages the application will need to send, including their format and content (*hint:* think about how `telnet` works). If you have more than one message type, describe each type separately.

4. This question is about different methods for implementing broadcast. In particular, we will consider two methods covered in lecture:
Method 1: The sender sends a copy of the packet individually to each recipient (via unicast).
Method 2: A spanning tree is constructed so that each node in the network knows which of its adjacent links are contained in the spanning tree. The sender sends out the packet on every adjacent link in the spanning tree. When a router receives a copy of the packet, it forwards copies on all adjacent spanning tree links except the one it came in on.

(a) Consider the following network. As usual, squares represent hosts and circles represent routers. Draw a spanning tree for this network.

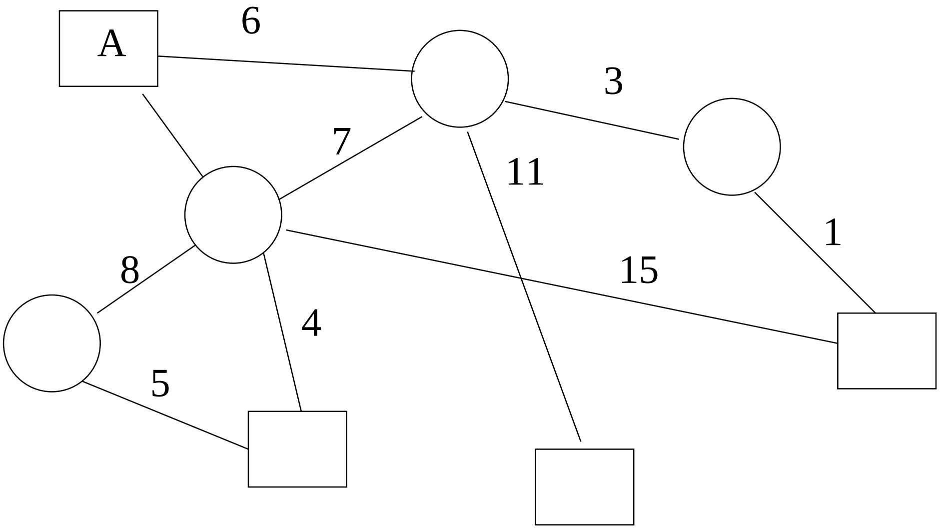

(b) Every time a copy of a packet gets sent across a link, we add the cost of that link to the total broadcast cost. What are the total broadcast costs of **Method 1** and **Method 2** (using your spanning tree) in the above network if node A broadcasts?

(c) For the remainder of the question, every link will have cost 1. Repeat parts (a) and (b) for the following network:

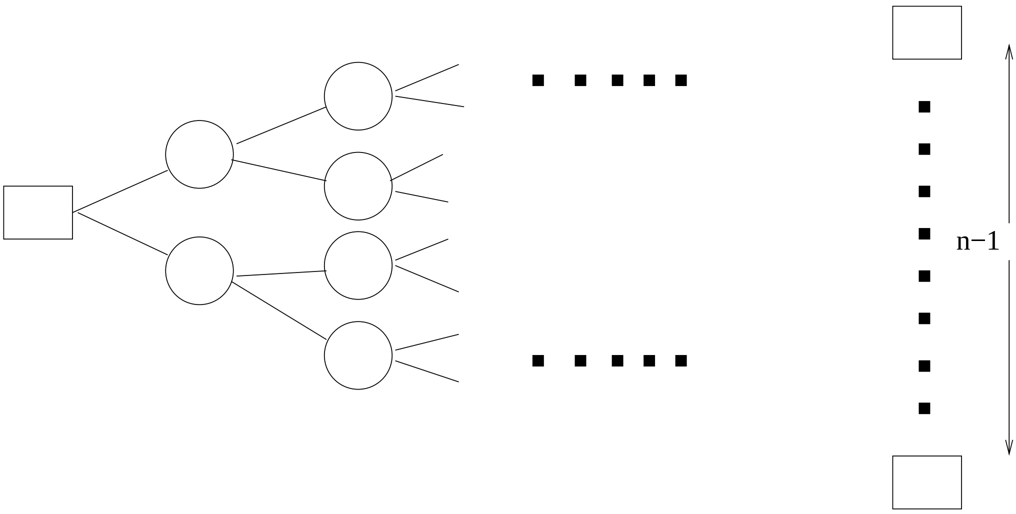

(d) Describe a network with n hosts where **Method 1** and **Method 2** have the same cost no matter which node is the sender. There is no limit on the number of links adjacent to any host.