Problem Set 0: Review of transport and routing

V. Arun

CS 653, Fall 2010
09/14/10

1. (**Protocol multiplexing**):

   (a) If a web server has 100 open connections, how many sockets and ports does it have active?

   (b) What tuple uniquely identifies a TCP socket? A UDP socket?

   (c) Suppose you wrote a web server and client application program using just UDP implementing reliability yourself. If the server were serving data to 100 users simultaneously, how many sockets and ports would it have active?

   (d) Is the transport protocol being used part of the transport header or network header? Why?

2. **(Reliable transport)**:

   (a) Can you implement a reliable transport protocol without sequence numbers? Why or why not? What if protocol only had to tolerate corruption but not loss?

   (b) When are the pros and cons of cumulative acknowledgments, negative acknowledgments (NACKs), selective acknowledgments (SACKs) respectively? What does TCP use?

   (c) Timeouts are needed to tolerate loss, but not to tolerate only corruption. True/false?

   (d) The link, network, and transport headers each have a checksum-like error detection capability. Explain why a packet is not guaranteed to be correct even if it satisfies all three checks. How do you reconcile this design decision with the end-to-end principle?

3. **(Pipelined transport)**:

   (a) What is the window size required to saturate a 10Gbps link with a 50ms RTT assuming 1500KB packets?

   (b) What is the main benefit of Go-back-N (GBN) protocols compared to Selective Repeat (SR) protocols?

   (c) Consider the GBN protocol with a sender window size of 3 and a sequence number range of 1024. Suppose that at time T, the next in-order packet that the receiver is expecting has a sequence number of K. Assume that the medium does not reorder messages. Answer the following questions:

   (i) What are the possible sets of sequence numbers inside the senders window at time t? Justify your answer.

   (ii) What are all possible values of the ACK field in all possible messages currently propagating back to the sender at time T? Justify your answer.

(d) What is the relationship between the size of the sequence number space and window size in order for the SR protocol to operate correctly?

4. **(TCP)**:

   (a) How does TCP estimate the timeout after which it should retransmit a packet?

   (b) What do triple duplicate ACKs and fast retransmit mean in TCP?

   (c) If a TCP sends a packet of size 500B with a sequence number of 122 and an acknowledgment sequence number of 4344, what is the largest sequence number up to which it has received the byte stream reliably? Could it have received bytes beyond that number? If the packet is correctly received by the TCP at the other end and it responds immediately, what are the sequence numbers in the response packet?

   (d) Why is it important for a TCP server to use a special initial sequence number in the SYNACK instead of just starting from 0?

   (e) For estimating the RTT, why do you think TCP exclude the sampled RTTs of retransmitted packets?

5. **(Congestion control)**:

   (a) Consider a network with N sender-receiver pairs. Suppose every link in the network has greater capacity than the sum of the rates at which the applications at each of the N senders is generating data. If the data has to be delivered reliably, is congestion control necessary in this scenario?

   (b) Can congestion collapse happen over a single link network? Over a chain network?

   (c) Consider the following plot of TCP window size as a function of time. Assuming TCP Reno is the protocol, answer the following questions.

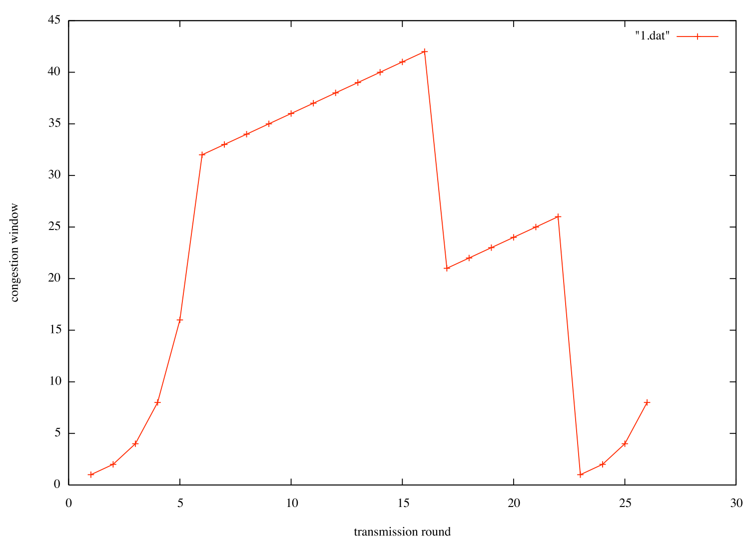

Figure 1: TCP window dynamics

      i. Identify the intervals of time when slow start is operating.

      ii. Identify the intervals of time when TCP congestion avoidance is operating.

iii. After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

iv. What is the initial value of `ssthresh` at the first transmission round?

v. What is the value of `ssthresh` at the 18th and 24th transmission round?

vi. During what transmission round is the 70th segment sent?

vii. Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congestion window size and of `ssthresh`?

(d) Consider the following simple model of TCP that varies the sending rate from $W/(2 \cdot RTT)$ to W/RTT. Assume that exactly one packet is lost at the very end of the period.

i. Show that the loss rate (fraction of packets lost) L is equal to $L = \frac{1}{\frac{3}{8}W^2+\frac{3}{4}W}$

ii. Use the result above to show that if a connection has loss rate L, then its average rate is approximately given by $\approx \frac{1.22 \cdot MSS}{RTT\sqrt{L}}$.

6. **(Quickies)**:

(a) Does a virtual circuit network maintain connection state information? If a VC is required to use the same VC number on all links along the path, how could the number be assigned?

(b) Give an example of how load-dependent link-state routing can lead to oscillations.

(c) Derive an upper bound on the maximum number of iterations required for distance vector routing to converge in terms of the size of the underlying graph.

7. Consider a datagram network using 8-bit host addresses. Suppose a router uses longest prefix matching and has the forwarding table below. For each of the four interfaces, give the associated range of destination host addresses and the number of addresses in the range.

Prefix match	Interface
1	0
11	1
111	2
otherwise	3

Table 1: Forwarding table using longest prefix matching.

8. Consider the network shown below. Suppose the cost of the link BC increases to 42. Assume that nodes exchange distance vectors once every 30 seconds.

(a) Show how A's cost to C increases in the next 3 minutes.

(b) Do the same assuming poisoned reverse is being used.

9. You are hired by OIT to setup two routers, R1 and R2, as shown below. You have the network 128.119.248.0/21 at your disposal. Give an example of how this address space can be distributed in the domain by presenting the network identifiers for each of the four networks plus the forwarding tables for the two routers.

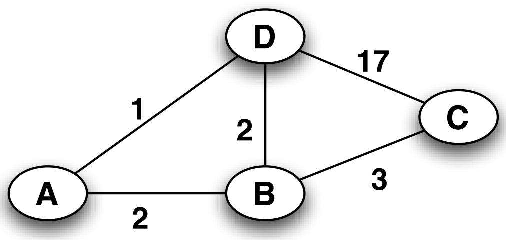

Figure 2: Distance vector routing dynamics

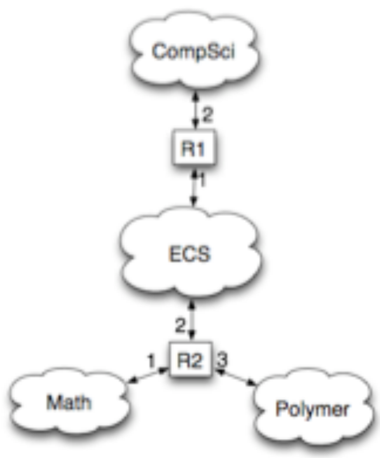

Figure 3: IP address assignment to subnets