Homework Assignment #3

Due 11:59pm, November 20th

EE122 Fall 2011

Please submit your solutions using BSpace (https://bspace.berkeley.edu/). Log in to BSpace with your CalNet ID, click on the EL ENG 122 Fa11 tab, and click on Assignments under Course Tools. Assignments should be submitted in one of the following formats: .txt, .pdf, or .ps.

1. TCP Basics: [10 points]

i) A TCP connection has been established between hosts A and B. A receives a packet from B with the following field values shown below:

Sequence: 1001
Acknowledgement: 5001
Window size: 3000
[TCP Payload Size: 52]

Note: the size of the TCP payload is not a field, but is the number of bytes in the TCP payload, which the receiver can determine from the header lengths and total length, so we treat it here as something that can be determined from the headers.

Which of the following are possible valid responses from A? (i.e., which represent TCP packets A might generate immediately after receiving this packet?)

Answer: ___ **a,c** _____ (there can be more than one valid response)

a) Sequence: 5001
Acknowledgement: 1053
Window size: 2000
[TCP Payload Size: 1000]

b) Sequence: 1053 **Wrong sequence number (in wrong direction)**
Acknowledgement: 5001
Window size: 3000
[TCP Payload Size: 1000]

c) Sequence: 6001
Acknowledgement: 1053
Window size: 2000
[TCP Payload Size: 1000]

d) Sequence: 8001 **Beyond window size**
Acknowledgement: 1053
Window size: 2000
[TCP Payload Size: 1000]

e) Sequence: 4001 **Already ACK'd**
Acknowledgement: 1053
Window size: 2000
[TCP Payload Size: 1000]

ii) Client A initiates a Telnet session with server S, using an ephemeral source port of 45644. Client B initiates a HTTP session with server S, using an ephemeral source port of 52125. Fill in the source and destination ports for the following TCP segments:

Sent from A to S: source port_ **45644** _, destination port__**23**____
Sent from S to A: source port___**23**___, destination port_**45644**_
Sent from B to S: source port_**52125**_, destination port___**80**___
Sent from S to B: source port___**80**___, destination port_**52125**_

iii) Host A and B are communicating over a TCP connection, and Host B has already received from A all bytes up through byte 6500. A now sends two segments to B back-to-back. The first segment has sequence number 6501 and contains 200 bytes of data. The second contains 400 bytes of data.

What is sequence number of the second segment? _**6701**_

If the segments arrive in order, and B sends ACKs to both packets, what are acknowledgement numbers in the first_**6701**_ and second_**7101**_ACKs?

If the second segment arrives before the first segment, what are acknowledgement numbers in the first_**6501**_ and second_**7101**_ ACKs?

If the first segment is lost, what is the acknowledgement number is the ACK packet sent in response to the arrival of the second segment? **6501**__

iv) If the average of the RTT is estimated to be 100msec and the average of the deviation of the RTT is estimated to be 20msec, what is the resulting RTO?__ **180**_

2. TCP Congestion Control [15 points]

Use the approximate equation (derived in class) for throughput as a function of drop rate:

$$throughput = \frac{\sqrt{1.5}MSS}{RTT\sqrt{p}}$$

Assume an RTT of 40msec and an MSS of 1000bytes. In the questions to follow, give answers to three significant figures and ignore IP and TCP headers in your calculations.

i). What drop rate *p* would lead to a throughput of 1Gbps?_ **6×10^{-8}**_

ii) What drop rate *p* would lead to a throughput of 10 Gbps?_ **6×10^{-10}**_

iii) If the connection is sending data at a rate of 10Gbps, how long on average is the time interval between drops?__ **1333sec**___

iv) What window size W (measured in terms of MSSes) would be required to maintain a sending rate of 10Gbps?_ _(rounded down to the nearest integer)

There are two acceptable answers. The first, is that the minimum fixed window needed to sustain a certain throughput T must satisfy the equation T = W*MSS/RTT. That gives the answer of 50,000.

The second acceptable answer is that in AIMD, the peak window size that yields a certain throughput satisfies the equation T=(3/4) W*MSS/RTT (i.e. AIMD would sawtooth between W/2 and W). This yields 66,666.

v) If a connection suffered a drop upon reaching 10Gbps, how long would it take for it to return to 10Gbps (after undergoing a fast retransmit)?_ ___(in seconds, rounded down to the nearest second)

Depending on which answer you gave above, the acceptable answers are: 1,000sec or 1333sec.

vi) Consider two TCP connections whose throughput obeys the TCP throughput equation listed above. The first TCP connection has the following parameters:

MSS = 1000 bytes, RTT = .2msec, drop rate = .5%

The second TCP connection has the following parameters:

MSS = 2000 bytes, RTT = .1msec, drop rate = 8%.

What is the ratio of throughputs (the throughput of the first TCP connection divided by the throughput of the second TCP connection)?___ **1** ____

vii) Consider the plot (on the following page) of CWND versus time for a TCP connection. At each of marked marked points along the timeline in the figure on the next page, indicate what event has happened, or what phase of congestion control TCP is in (as appropriate), from the following set: (a) Slow-Start, (b) Congestion

Avoidance, (c) Fast Retransmit, (d) Timeout.

A.__ **a**___
B.__ **c**___
C.__ **b**___
D.__ **c**___
E.__ **b**___
F.__ **d**___
G.__ **a**___

Assume CWND is 10,000 right before F, what is the value of SSTHRESH at G? _**5000** _

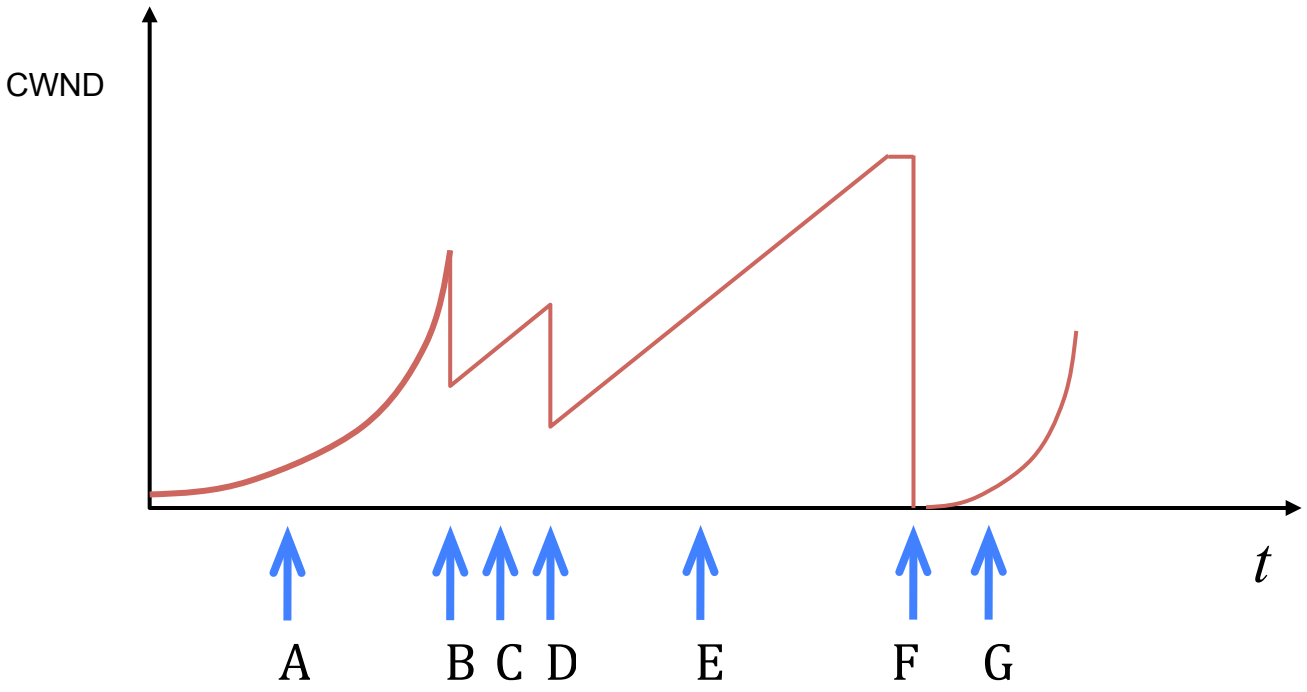
CWND
t
A
B C D
E
F G

3. DNS Basics [5 points]

Consider Host A in domain berkeley.edu, with a local DNS server (call it B. The rest of the DNS infrastructure includes:

a) Root
b) Tree
c) TLD server for .com
d) TLD server for .edu
e) TLD server for .org
f) Authoritative server for foo.com
g) Authoritative server for berkeley.edu
h) mine.foo.com
i) mine.berkeley.edu

A asks B to resolve hostname mine.foo.com. Assume there are no cached entries relevant to this request. Write down the steps taken to resolve foo.com and respond to A. Use the letters above for your answers.

Step 1: B contacts _**a**__ and gets the IP address of what host?_ **c**_

Step 2: B contacts _**c**_ and gets the IP address of what host?__ **f**_

Step 3: B contacts _**f**__ and gets the IP address of what host?__ **h**_

Step 4: B contacts A and gives it the IP address of what host?__ **h**__

4. HTTP [10 Points]

Consider a client accessing a base HTML page with three URL references to images to be displayed in the page; the images reside on the same server as the base page, and the path to the server has latency L (that is its one-way propagation delay) and bandwidth B. None of the images can be requested until the entire base page is retrieved. Ignore congestion from other traffic, and server processing time (i.e., assume the web server can process requests as fast as they come in). Assume all HTTP request packets, as well as TCP SYNs and ACKs, are tiny (compared to the returning responses), the base page has size P and each image has size M. Further, assume that the TCP window is large enough so that the connections can fill the pipe (i.e., achieve bandwidth B). Note that the next TCP connection can start before the previous TCP connection has terminated (i.e., HTTP does not need to wait until for the exchange of FINs to complete before issuing its next request).

For each of the scenarios below, describe the total time (in terms of L, B, P, and M) it takes to retrieve the base page and its embedded images. Pick from one of the following options for each:

a) 10L + P/B + 3M/B
b) 5L + (P+M)/B
c) 6L + 4M/B
d) 4L + P/B + M/B
e) 16L + P/B + 3M/B
f) 8L + P/B + 6M/B
g) 8L + 2P/B + 3M/B
h) 8L + 3P/B + 2M/B
i) 8L + P/B + 3M/B
j) 4L + P/B + 3M/B
k) 4L + 3(P+B)/M
l) 6L + P/B + 3M/B

i) Sequential (one-at-a-time) requests with nonpersistent TCP connections:_**e**__

ii) Concurrent requests with nonpersistent TCP connections:__ **i** __

iii) Sequential requests within a single persistent TCP connection:__ **a** __

iv) Pipelined requests within a single persistent TCP connection:_ **l** ___

5. BGP [10 points]

i) Consider the network in Figure 1 (on the next page), with ASes A, B, C, D. Each AS has some number of routers (labeled as A1, A2, etc.) and the domains are connected internally and with each other by the links depicted in the figure. Assume that eBGP and iBGP are used for interdomain routing, and that ASes A and D are using RIP for intradomain routing while ASes B and C are using OSPF for intradomain routing. Prefix x hangs off an interface on router C3. For the following answer, use one of these options:

a) OSPF
b) RIP
c) eBGP
d) iBGP

Router D3 learns about prefix x from which routing protocol?__ **c** __

Router D1 learns about prefix x from which routing protocol? __ **d**__

Router A3 learns about prefix x from which routing protocol? __ **c**__

Router A1 learns how to reach router A3 from which routing protocol? __ **b**__

Will router A1 use interface 1 or interface 2 to reach prefix x? _ **1**__

ii) Consider an interdomain network with domains A through F. For simplicity, assume that destinations in this problem are domains, not prefixes. Recall that routes are expressed in terms of the series of domains: e.g., [A-B-C] denotes a route that started with domain A and went to domain B and then to domain C (which is the destination). Domains always advertise the route to themselves (i.e., domain X advertises the path [X] to all peers, customers, and providers). The following connectivity/business relationships exist:

- B is a customer of A
- C is a customer of A
- D is a customer of B
- E is a customer of B
- F is a customer of C
- G is a customer of C
- B and C are peers
- E and F are peers

Assuming that each domain's routing policies follow normal business practice, and that BGP has converged,

- What routes does A advertise to B? _ **[A] [A-C] [A-C-F] [A-C-G]**

- What routes does C advertise to B? _**[C] [C-F] [C-G]**
- What routes does E advertise to B? _**[E]**
- What routes does F advertise to E? ___**[F]**___
- What path do packets from E take to F? __**E-F**__
- What path do packets from D take to F? ___**D-B-C-F**___
- What path do packets from D take to G? ___**D-B-C-G**

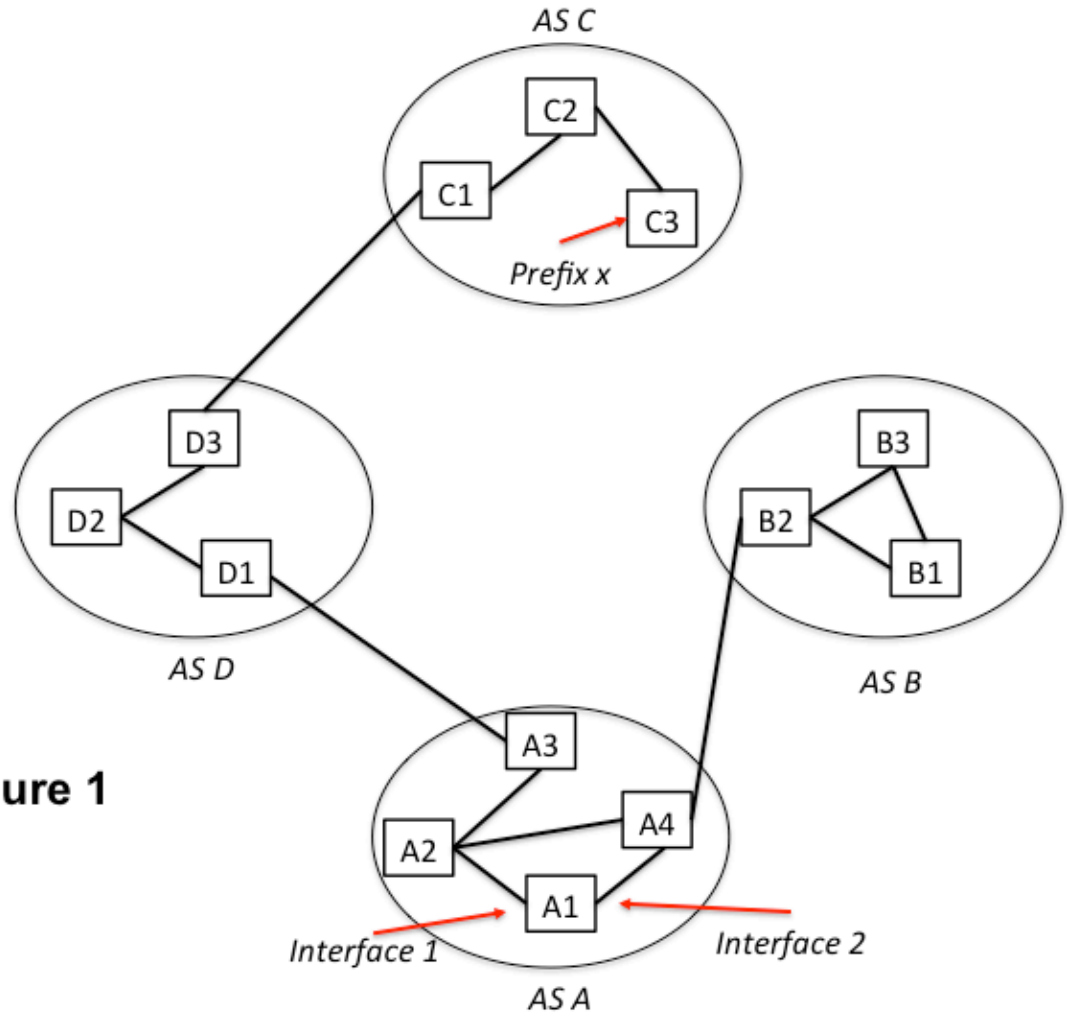

Figure 1