1. If you didn't complete 16(c) from last week's assignment, please make sure to do that this week.

2. Show that to achieve a steady-state throughput of 10 Gbps, a TCP session with a Round-Trip-Time (RTT) of 100 ms and a Maximum-Segment-Size (MSS) of 1500 bytes can tolerate a packet loss probability of less than 2×10^{-10}. What would the tolerable loss be for a 100 Gbps connection?

3. Consider the following plot of CWND size over the duration of a TCP connection.

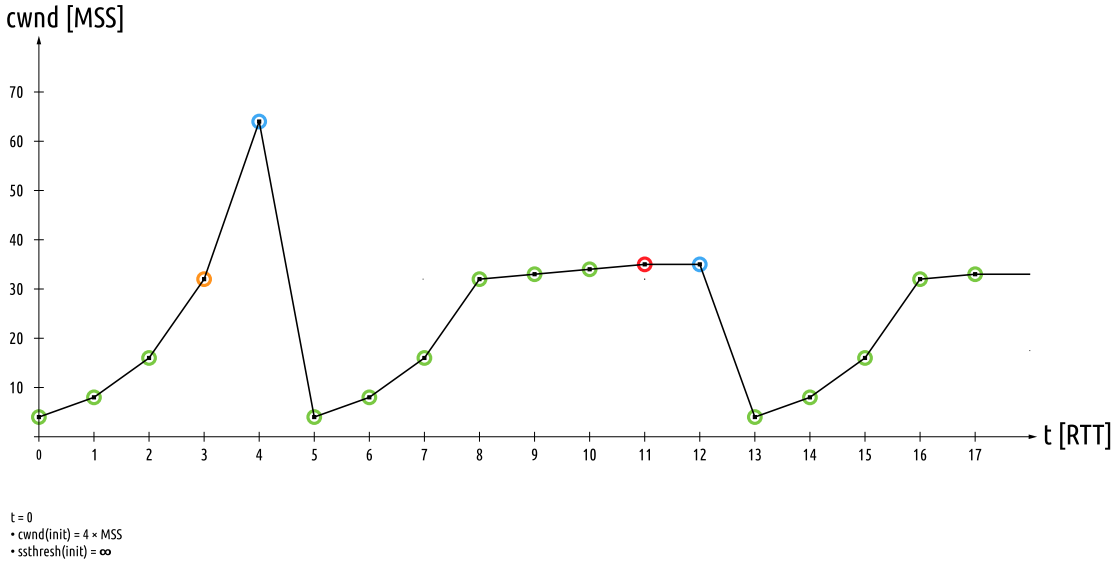

   (a) What has occurred during the events marked by blue circles?

   (b) What has occurred during the event marked by the red circle?

   (c) What range(s) of times correspond to slow-start?

   (d) What is `ssthresh` at t=5?

   (e) What is `ssthresh` at t=13?

4. It is interesting to consider the consequences if TCP's behaviour was a bit different. Describe the consequences if the following were true:

   (a) TCP slow-start increased CWND by 1 for each ack

   (b) TCP does not reset the CWND to 1 after a timeout

   (c) TCP resets CWND to 1 after a duplicate ACK

5. TCP has a few problems areas (slide 259)

   (a) Why are non-congestion losses bad for TCP? Give examples of how non-congestion losses might arise.

   (b) Give two reasons why short flows are inefficient for TCP. What kind of Internet activity can result in many short flows?

6. Explain how you would set up alternative DNS servers that allows *both* standard DNS top-level domains (since you don't want users of your DNS server to not be able to reach well-known Internet sites) as well as custom TLDs that you've created to be resolved?

7. Cascading Style Sheets (CSS) are a way to describe the appearance of items on a webpage. CSS data is often stored in separate files that are referenced from a webpage. A piece of advice given to web designers to speed up page loading is to combine the contents of multiple CSS files page into a single one.

   (a) Why does this advice speed up the loading process?

   (b) If the CSS file is very small, another recommendation is to include the CSS data directly in the HTML file itself. However, this is not recommended when the CSS file is larger. What is the rationale for this recommendation?

8. Many streaming services implement "geo-blocking", which restricts content to users coming from a specific region or country.

   (a) Given what you know about the structure of the Internet, how does checking a client's IP give a coarse estimate about where a client is located?

   (b) Interestingly, for some websites, simply changing the DNS server one uses, rather than changing one's IP address, is enough to bypass the blocks. Hypothesize about how blocking is being implemented in this case.

9. A friend claims that the fastest way to retrieve all of the contents of a web page is to open lots of simultaneous TCP connections to download all of the resources in the page in parallel. Under what circumstances is that claim true? Under what circumstances is this false? What are the disadvantages of doing this?

10. Your computer is connected via a 50 Mbps link to server C, which is a HTTP cache server. Server C is connected via a 1 Gbps link to the web server with the actual content. Assuming the version of TCP you are using allows an MSS of unlimited size and the TCP SYN/ACK and HTTP request packets have negligible size:

    (a) How long does it take for you to retrieve a 125KB file from C if the file is already in the cache if the RTT between your computer and C is 40ms and RTT between C and the web server is 400ms.

    (b) How long does it take for you to retrieve the same file if the file is not present in the cache?

11. Key-value lookup is a frequently encountered activity in computing. Using a file's identifying details (key) to locate peers that have a copy of a file (value) is one example. Finding the mapping from an IP address (key) to a MAC address (value) is another. Of course, computers do not use DHTs to do MAC lookups and ARP-like algorithms are not used to locate file sharing peers. Describe two ways where finding peers having a file is a *different* problem from doing a MAC lookup on a LAN. What characteristics of the peer-finding problem motivates the use of a DHT over an ARP-like protocol?