1. Some organizations provided (and continue to provide) alternative DNS root services. These services allow the resolution of non-standard TLDs alongside the classic TLDs (i.e., .com, .net, etc.). These services were useful before ICAAN started allowing people to reserve gTLDs (e.g., .google, .microsoft, etc.), and continue to be useful to people who for one reason or another do not want to use the "standard" DNS root servers.

   (a) Sketch out what an implementation of an alternative DNS root service would look like.

   (b) What instructions would you need to provide to users who are interested in using your service?

   (c) What challenges around getting large-scale adoption does an alternative DNS root suffer from?

2. A piece of advice given to web designers to speed up page loading is to combine multiple CSS files on a page into a single one.

   (a) Why would this speed up the loading process?

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

3. Many web 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.

   (c) Compare the benefits of using DNS-based geo-blocking over IP address-based blocking and vice-versa. Discuss things like scalability, robustness of the filtering, ease of implementation, etc.

4. 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. Under what circumstances is that claim true? What are the disadvantages of doing this?

5. 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?

6. DHTs are overlay networks, which means the topology of the DHT may not match the underlying network. For example, neighbors in the DHT might actually be very distant from each other. Explain how a mismatch in these topologies could affect the performance of the DHT. Propose a solution to address this issue.