CIS 553: Networked Systems April 21, 2020

Problem Set 2

Instructor: Prof. Vincent Liu *Department of CIS, University of Pennsylvania*

Turn in your solutions in on **April 27, 2020** by 10:00 p.m. via Canvas. Grading will be lenient, as long as you give an honest effort to answer every question. **Please work alone** and do not discuss answers with other students.

Note that while all of these questions are taken from prior exams, the content on the exam may be different. The fact that the exam will be a take-home exam means that the question style will be different as well. Instead, treat this as an opportunity to review some of the material and debug any basic misunderstandings about the basic ideas covered so far.

1. True or False

   (a) The 5-tuple that identifies distinct flows is usually defined as (src IP, dst IP, src port, dst port, protocol #).

   (b) One of the primary reasons why TCP is used for video streaming is that many firewalls block UDP packets.

   (c) The Pakistan hijacking case occurred because Pakistan's ISP spoofed low-hop advertisements to remote ISPs.

   (d) Streaming video is video that can begin playout before downloading the entire file.

   (e) AIMDs Multiplicative Decrease is a form of BEB.

   (f) Explicit proxies are proxies that act on behalf of a few servers, but handle all clients.

   (g) In wireless networks, multipath effects cause self-interference.

   (h) Users in censoring countries can defend against TCP RST injection by modifying their laptop network configurations.

   (i) One of the enablers of DDoS attacks is the fact that the Internet implements best-effort routing.

2. Imagine that you are watching streaming video from `unoptimizedvideo.com` over HTTP/1.0 and are experiencing video playback stuttering. For each of the following optimizations:

   i. Do you expect the technique to improve the smoothness of video playback? If so, how? If not, why not?

   ii. Do you expect the technique to be more effective on video than a typical webpage? Why or why not?

   (a) Pipelined/persistent connections

   (b) CDNs

   (c) HTTP/2

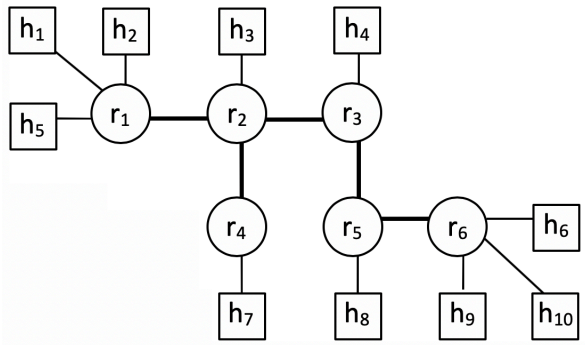

3. Consider the network above. Each host-to-router link has capacity 100 Mb/s and each router-to-router link has capacity 200 Mb/s.

   Given the following concurrent TCP connections (src→dst; demand), what are the steady-state max-min fair share allocations of each of them on the given network? What is the bottleneck link of the connection (e.g., $r_5 \rightarrow r_6$)?

   (a) $h_10 \rightarrow h_5$; 80 Mb/s

   (b) $h_8 \rightarrow h_3$; 50 Mb/s

   (c) $h_9 \rightarrow h_1$; 200 Mb/s

   (d) $h_7 \rightarrow h_2$; 40 Mb/s

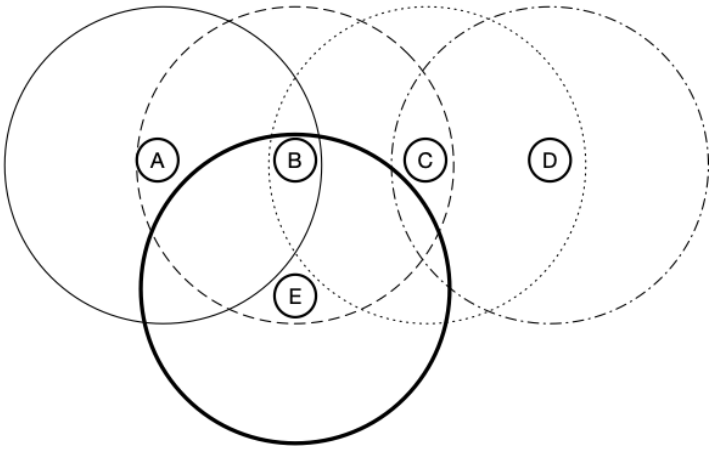

4. For the following questions, consider the wireless topology above, comprised of 5 nodes: A, B, C, D, and E. All of them have an effective transmission range as shown. For example, D can only speak to C (C is the only node in its circle), and can also only hear from C (D is in Cs circle and no others).

   (a) Why are acknowledgements used in 802.11, but not in wired Ethernet?

   (b) Assume A is about to transmit to B. Which terminals transmissions can cause hidden terminal problems? List in terms of src→dst. What about exposed terminal problems?

   (c) Repeat (a) for B→C. Which terminals transmissions can cause hidden or exposed terminal problems?

5. Fair Queuing

   (a) Why does even bit-level fair queuing fail to provide max-min fairness in some scenarios?

   (b) Using python-like pseudocode, write an implementation of packet-level fair queuing. You can base your implementation around two functions: `recv(packet, size, flowID)`, which is called on every new packet arrival, and `send()`, which is called whenever the output port is idle. `send()` should return the next packet to send or NULL if there are no packets queued. You can define any global state you need.