CSCI-UA.0480-009 2018 midterm (47 points)

October 18, 2021

Instructions

1. Write your name and N number on top.
2. Provide concise and clear explanations for all your answers.
3. Use the other side of each sheet if you need more space. If you need even more space, ask us for additional sheets.
4. Some useful conversions: 1 Gbit = 10^9 bits, 1 Mbit = 10^6 bits, 1 kbit = 10^3 bits
5. Questions are roughly in order of the lectures.
6. You should have 9 pages with questions.

1. (3 points) What are the units of throughput, queueing delay, window size, capacity, RTT, and Bandwidth-Delay Product?
   Answer: Throughput - bits/second Window - packets/bits Capacity - bits/second Bandwidth Delay Product - bits RTT - seconds Queuing Delay - seconds

2. (3 points) What are the 5 layers of the Internet stack? What is each responsible for?
Answer: Application Layer - tailoring to specific application needs (HTTP) Transport Layer - ensuring reliability and encapsulation from end host (TCP) Routing/Network Layer - global routing and forwarding (IP) Link Layer - local transport between network nodes/hosts (Ethernet) Physical Layer - encoding bits into electrical signals

3. (2 points) How does the idea of layering achieve the Internet's goal of generality?
Answer: Layering maintains a consistent and clean interface between the different layers of the stack and allows a diverse set of protocols to run smoothly on top of a single protocol in a different layer.

4. (3 points) Recall assignment 1 where you had to implement a relay that replaced bad words in a string that was being sent from the sender to the receiver. The service provided by the relay, which is malware detection and removal, is commonly provided today by devices called middleboxes. Middleboxes are deployed inside the networks of companies to keep the employees' desktops and laptops free from malware. Do you think the presence of these middleboxes violates the end-to-end principle? Why or why not?

   Answer: Two answers are acceptable. The use of middleboxes doesn't violate the end-to-end principle because middleboxes are typically implemented on server machines, which are really just end hosts. Because the end-to-end principle allows complicated functionality to run on end hosts, the end-to-end principle is not violated.

   The use of middleboxes does violate the end-to-end principle because middleboxes are devices that are inspecting traffic that does not belong to them. In that sense they are not end-hosts but are really part of the network, even if they happened to be implemented on server machines. The end-to-end principle argues that the network should be as simple as possible, so putting malware detection into the network is against the end-to-end principle.

5. (2 points) Unlike circuit switching, why does packet switching need end hosts to insert the destination address into the packet header?

   Answer: Circuit switching sets up a dedicated channel beforehand which provides a path to the endhost. Packet switching however is broken up into "packets" and each of these can take a different route to the end host, so each packet must contain a header to allow it to be routed along independent paths as the network conditions dictate.

6. (2 points) How do private IP addresses solve the problem of exhaustion of IP addresses?
   Answer: Private IP addresses allow many internet connected devices to reuse the same IP address across independent routing domains while using only one public IP address.

7. (4 points) If you were given the task to choose either between TCP or UDP for four applications (video conferencing, file downloads, instant messaging, and live video streaming), what would you choose and why?
   Answer: Video Conferencing - UDP (want interactivity) File Download - TCP (want the whole file) Instant Messaging - TCP (want the whole file) Live Stream - UDP (want realtime display)

8. (2 points) Give two examples of failure conditions that a reliable transport protocol like the sliding-window or stop-and-wait protocol is designed to handle.
Answer: Dropped packet or failure of one of the routers in the calculated path.

9. (2 points) Give two examples of failure conditions that a reliable transport protocol like the sliding-window or stop-and-wait protocol is NOT designed to handle.
Answer: End host disconnected, DDoS attack.

10. (2 points) What are the benefits/drawbacks of using a larger/smaller α in the computation of the mean RTT in TCP's retransmission timer?
Answer: Larger alpha = more quick to respond but overreacts to minor changes. Smaller alpha = slower to respond but has a more robust estimator.

11. (2 points) Why does TCP use an adaptive algorithm to compute a retransmission timeout (i.e., the rtt_mean + 4.rtt_var algorithm)? What would be the drawback of using a fixed retransmission timeout?
Answer: A fixed timeout value might either be too small or too high depending on the actual rtt_min of the network. Too high would imply waiting too long to retransmit lost packet, while too low would risk flooding the network with too many retransmissions, leading to a congestion collapse in the process.

12. (3 points) Let's say you were running TCP on a network where there were no losses whatsoever because the queues could grow infinitely large. Let's assume you were using the TCP congestion control algorithm (Slow Start + Congestion Avoidance) for dynamically adjusting your window on this network. Would your window ever decrease in size? Why or why not?
Answer: Yes. Eventually the delay would become so great as a result of the buildup of the queues that the packets would be counted as lost even though they were not dropped from the queue memory. When this happens, the window would decrease in size.

13. (4 points) What is the queueing delay (i.e., the round-trip time minus the minimum round-trip time) when a sliding window sender's window size exceeds the network's bandwidth-delay product? Express your answer in terms of the window size W, the connection's bottleneck bandwidth C, and minimum round-trip time along the path RTT_{min}, and justify your answer.
Answer: Queuing Delay = (W/C) - RTT_{min}. Each packet's transmission delay is (1/C) and it will have to wait behind W - C*RTT_{min} packets because W packets are unacknowledged, of which $C * RTT_{min}$ are "in flight" in the network, and the rest are sitting in queues. The total queuing delay is the number of packets in the queue (W - C*RTT_{min}) multiplied by the transmission delay of each packet.

14. (1 point) What is the queueing delay in the previous question if the window size is below the bandwidth-delay product?
Answer: Zero delay.

15. (3 points) The Internet is often known for "triangle inequality" violations, where the path from host A to host B on the Internet has longer latency than the sum of the latency from host A to host C and the latency from host C to host B. In other words, $Latency_{AC} + Latency_{CB} < Latency_{AB}$. Here latency between two points P1 and P2 (i.e., $Latency_{P1P2}$) is the sum of the transmission delay and the propagation delay from P1 to P2. Notably, latency here does not include queueing delay.

    Give one plausible reason for these triangle inequality violations by providing a network topology/graph that explains this violation. The nodes in this topology should be autonomous systems or domains, and you should clearly indicate which domain hosts A, B, and C belong to. Hint: You should think in terms of how inter-domain routing's goals might be different from a user's goals.

    Answer: C's domain might be a customer of both A's domain and B'domain, allowing C to use a short path to both A and B. However, A and B's domains may not have a peering relationship (for business reasons) and might have to use a circuitous path to speak to each other. An example of domains that do not have a peering relationship (and hence use a circuitous path) is the Kenyan ISP and the South African ISP from the lecture's example on boomerang routing.

16. (3 points) Explain the two kinds of relationships between domains that are relevant to Internet inter-domain routing.
Answer: Peering relationship: multiple domains agree to carry each other's traffic. and transit relationship: a provider sends internet traffic to a customer.

17. (3 points) Recall that BGP has two variants: external BGP and internal BGP. What purpose do each of these two variants serve in the context of inter-domain routing?
Answer: The External BGP protocol is responsible for discovering paths that cross multiple domains. Internal BGP protocol handles the propagation of this interdomain information to routers within each domain.

18. (3 points) What are the two planes (or classes of functionality) implemented by routers within a network? In response to what kinds of events does each plane get executed?
Answer: The Data Plane deals with the actual forwarding of packets using the routing tables. The Control Plane calculates the routes the packets will take and updates the forwarding tables.