CSCI-GA.2620-001 final (55 points)

May 6, 2019

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.
4. The questions are roughly in the order of the lectures.

1. (4 points) Derive the TCP throughput equation that relates the packet loss probability and the minimum round-trip time on a link to the throughput achieved by AIMD TCP over this link.

Solution: Look at xcp_motivation.pdf in Lecture 3.

2. (3 points) Use the TCP throughput equation you derived above to estimate the packet loss probability at which TCP throughput falls to 1% of the link's capacity. Assume a link with a minimum round-trip time of 100 ms and a link capacity of 1 Gbit/s.

Solution: Plug in numbers into the equation from the previous answer.

3. (2 points) Give two scenarios in which BBR outperforms TCP Cubic. For each of these scenarios, what is the metric on which BBR outperforms TCP Cubic?

Solution:

1. Wide-area data transfers (Fig 7 of BBR paper).
2. Improvement in latency that directly affects YouTube's quality-of-experience (Fig 9 of BBR paper).

4. (2 points) What performance property does WFQ provide that XCP doesn't? In what way is a router implementation of WFQ more challenging than a router implementation of XCP?

Solution: WFQ provides isolation (or equivalently packet-level fairness) where a flow is completely insulated from another misbehaved flow because the misbehaved flow can't grab a large share of the link's capacity. WFQ is challenging because it requires per-flow state to track the virtual times for each flow.

5. (2 points) In BGP, route advertisements are implicitly trusted, i.e., all routers assume that all other routers are honest. Give an example of what could happen if routers lied in their advertisements. There isn't one right answer to this question. I'll accept anything that's well justified.

Solution: One example is the route hijacking incident in 2008 due to which YouTube was down for most of the Internet due to an incorrect routing advertisement from Pakistan telecom.

6. (3 points) Assume you're an AS/ISP with two paths to the Internet—maybe one through Comcast and another through Verizon. How would you use BGP to advertise to the rest of the Internet that you would prefer to receive traffic through Comcast and would only want to use Verizon as a backup?

Solution: By using path prepending to artificially make the Verizon route's path length look bigger. The BGP lecture notes in lecture 4 and the SDX paper from lecture 5 both explain this.

7. (2 points) Give one example of a routing policy that is easy to achieve using SDX but much more cumbersome with BGP.

Solution: Application-specific peering/routing where a different route should be used depending on the application.

8. (2 points) In the VL2 paper, why is randomized load balancing performed at the granularity of flows instead of packets? What is the drawback of performing load balancing at the level of flows instead of packets?

Solution: It's performed at the level of flows to avoid reordering packets within a flow. Reordered packets cause retransmissions in TCP and also cause TCP to reduce its congestion window. By performing load balancing at the level of flows, the load balancing is likely to be less even because of difference in flow sizes, e.g., a large flow is assigned to one path and a small flow is assigned to another path.

9. (1 point) How is the ECN marking required for DCTCP implemented using an off-the-shelf non-programmable switch?

Solution: By repurposing the RED algorithm and setting its min and max queue thresholds to the same value (Section 3.1).

10. (2 points) In what ways is a wired Ethane switch simpler than a commercial wired Ethernet switch? If the switch itself is simpler, where does all the complexity reside in an Ethane network?

Solution: An Ethane switch doesn't need support for virtual LANs, address learning, access-control lists, and network-address translation. Most of the complexity resides in the controller of the Ethane network.

11. (2 points) Give three examples of the kinds of flexibility that RMT provides relative to a fixed-function Ethernet/OpenFlow switch. Give an example of the kind of flexibility that RMT does not provide.

Solution: Flexibility to add new header fields, flexibility to match on arbitrary sets of header fields, and flexibility to carry out new actions composed out of a primitive set of actions. RMT does not allow deep packet inspection or more generally processing of the packet's payload.

12. (3 points) What groups of people benefit from a technology like RMT? Why?

Solution: Researchers: to try out new algorithms Switch equipment vendors: to fix bugs in firmware instead of software Network operators: to customize networks to their needs.

13. (3 points) What security properties does TLS provide? Describe these properties as precisely as possible.

Solution: Confidentiality, authentication, and integrity. https://cs.nyu.edu/courses/fall18/CSCI-UA.0480-009/lectures/lec22.pdf has a detailed explanation of each.

14. (2 points) Describe the Heartbleed vulnerability.

Solution: The ability to read remote memory using the SSL heartbeat extension. Also look at Section 2 of the Heartbleed paper.

15. (2 points) How can the Heartbleed vulnerability be fixed?

Solution: Adding a bounds check. Section 2.3 of the Heartbleed paper.

16. (2 points) Give examples of unauthorized data that the Heartbleed vulnerability allows an attacker to access.

Solution: Private keys, private user data, passwords, etc.

17. (2 points) In the paper on censorship circumvention by Ensafi et al., what is the key weakness of the obfs2 protocol? How does obfs3 fix this?

Solution: It is easy to detect obfs2 because it works by sending a key and then decrypting the subsequent text with that key. 3.2 of Ensafi et al.'s paper describes this.

18. (2 points) Why did the "Low cost traffic analysis" authors configure TCP_NODELAY when performing their experiments? Why not just use UDP instead?

Solution: TCP_NODELAY is used to prevent delays in sending out probe data. As a matter of implementation, Tor nodes only accept TCP connections from other nodes, which is why the authors could not use UDP.

19. (2 points) In the Sprout paper, why does Saturator try to maintain a relatively high target delay (i.e., between 750 and 3000 ms)?

Solution: To ensure that the queues are always non-empty and hence the link rarely starves for data. This is important to ensure we get an accurate estimate of the link's capacity.

20. (2 points) Why does Saturator use two phones when attempting to saturate the cellular network? What would happen if Saturator instead tried to saturate both the uplink and the downlink of a cellular network using a single phone?

Solution: The second phone provides a low-delay feedback channel. If they had used a single phone, the feedback would have arrived on the same high-delay channel (recall Saturator tries to keep the delay high to ensure queues are non-empty), which would have made achieving a target queueing delay higher. In general, the higher the feedback delay, the harder it is to control a system to hit a particular target.

21. (3 points) In the LTEye paper, what is the C-RNTI? How is it used? When is the C-RNTI reassigned? How does the LTEye system handle these reassignments?

Solution: The C-RNTI is a temporary identifier to track mobile users attached to a base station. It is reassigned if the user is idle or the user moves to a different base station. LTEye uses a matching algorithm (Section 5) to match C-RNTIs across different traces to identify the same user/phone across different traces.

22. (3 points) Describe the difference between regular MIMO and multi-user MIMO using a diagram that illustrates where the transmit and receive antennas are located in each case. Where does the channel matrix inversion happen in each case?

Solution: In regular MIMO, the transmitter and receiver both have multiple antennas. In multi-user MIMO, the transmitter has multiple antennas, while the receivers typically have a single antenna each. Figure 1a of the MegaMIMO paper describes multi-user MIMO.

23. (2 points) What is the difference between the multi-user MIMO problem and the distributed MIMO problem, i.e., the problem that MegaMIMO solves? As a consequence of this difference, what is the key technical contribution of the MegaMIMO paper?

Solution: In MegaMIMO the antennas that are being coordinated are on different access points. In multi-user MIMO, they are all on the same access point. The key technical contribution is the ability to synchronize time, phase, and frequency between different access points to make the distributed MIMO setting as close as possible to the multi-user MIMO setting.

24. (2 points) What is the conceptual difference between analog and digital self-interference cancellation in the full duplex paper? (Besides the fact that one is analog and the other is digital!)

Solution: The digital portion is model based, while the analog portion (because it attempts to cancel out noise which is random) is not model based. Instead it tunes itself to match the transmitted signal as well as possible (using the attenuation coefficients and delay lines).