CSCI-GA.2620-001 final (80 points)

May 9, 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. If you can't simplify some mathematical expressions, it's OK to leave them unsimplified; however, please provide sufficient explanation in text.
5. We have a makeup exam next week. So please do not discuss the exam's contents with anyone.

1. (3 points) What is one shortcoming of loss-based congestion control (e.g., TCP Cubic) in networks with switches containing large buffers?

2. (3 points) How does BBR fix this shortcoming?

3. (4 points) Why does BBR use a time-windowed maximum for its bottleneck bandwidth estimate and a time-windowed minimum for its RTT estimate? What would happen if it used a different technique (e.g., time-windowed average or time-windowed minimum for bottleneck bandwidth) for each of these two quantities?

4. (5 points) Recall that the XCP algorithm requires several floating point operations to be carried out by the switch on every packet to compute the congestion feedback to XCP senders. What if the router's data plane is incapable of such complex operations? How can XCP still be implemented in such a network?

5. (5 points) What does instability mean in the consequence of Internet inter-domain routing? Give an example of instability. Explain the Gao-Rexford condition for stable Internet routing.

6. (3 points) Explain how DCTCP extracts multi-bit information on the *extent* of congestion from single-bit information on the *presence* of congestion. How is the single-bit information on the presence of congestion obtained?

7. (3 points) What is the difference in the congestion window decrease rule between DCTCP and standard TCP?

8. (4 points) Why does VL2 carry out flow-level load balancing instead of packet-level load balancing? What is the main shortcoming of flow-level load balancing relative to packet-level load balancing? Why is this shortcoming not as relevant within the datacenter, as opposed to the wide-area Internet?

9. (3 points) What are the two types of addresses that VL2 uses? Why does VL2 use two sets of addresses? Who translates between the two addresses?

10. (4 points) What shortcomings of OpenFlow did the P4 language attempt to address?

11. (5 points) We'll assume a simplified model of the RMT pipeline for this question. As part of the programmable action unit in RMT, let's assume the action unit within a single stage of the pipeline can add, subtract, multiply, or divide two packet header fields and write it into a third header field. Let's also assume you can create as many header fields as you like to store temporary results of computations between pipeline stages. In this model of the switch pipeline, how would you implement the following operation: $pkt.e = (pkt.a * pkt.b)/(pkt.c - pkt.d)$? Here pkt refers to a packet and a, b, c, d, and e are fields within the packet. Use a diagram to clearly specify how many pipeline stages you would need and what portion of this operation would run in each pipeline stage.

12. (4 points) Recall that the RMT architecture provides the ability to match against user-defined packet headers and carry out actions on these packet headers based on the result of the matching process. By using such match-action rules, explain how a programmer can carry out regular expression and string matching (i.e., in the style of grep and CTRL-F) on packets at line rate.

13. (3 points) Why does RMT need both a programmable ingress and a programmable egress pipeline? Why can't it perform all packet processing solely in the ingress pipeline?

14. (3 points) Give two examples of private data that can be extracted using the Heartbleed vulnerability, and explain how the Heartbleed vulnerability is used to extract this data.

15. (3 points) In the Heartbleed paper, the authors realized they had a bug in their data collection program after collecting data. Why did they have to estimate the impact of this bug using a back-of-the-envelope calculation instead of just recollecting their data?

16. (3 points) Tor is a system that enables anonymous communication. However, anonymity can still be compromised on Tor if Tor is used without a suitable encryption scheme (like HTTPS) for encrypting data flowing between the client and the server. Explain how this might happen with an example.

17. (2 points) How was the earliest version of Tor blocked by censors when it was used for censorship circumvention?

18. (5 points) Recall that the Sprout protocol was designed explicitly for cellular networks that provide isolation between users using per-user queues (e.g., using fair queueing). What would happen if you ran Sprout on a network without such per-user queues (e.g., on a network with FIFO queues)? Answer this question separately for two cases. One, if multiple Sprout sender-receiver pairs are running on the same network. Two, if a Sprout sender-receiver pair is sharing the same FIFO network as a TCP Cubic sender-receiver pair.

19. (4 points) LTEye relies on the control channel being unencrypted in order to measure various characteristics of the cellular network. Why is the control channel unencrypted? Given that it's unencrypted, how is secure communication achieved in a cellular context?

20. (3 points) LTEye runs a fingerprint matching algorithm using the Hungarian algorithm. What is the purpose of this matching algorithm?

21. (5 points) Imagine a multi-user MIMO setup with two transmit antennas on a single AP and two receivers with a single antenna each. Hence, in total, we have two transmit antennas and two receive antennas; we'll refer to them using the indices 1 and 2. The channels between antennas 1 and 2 on the transmitter and receiver are: $h_{11} = 1$, $h_{12} = 2$, $h_{21} = 3$, and $h_{22} = 4$. Let's say the AP wants to transmit the sample $x_1 = 10$ to receiver 1 and the sample $x_2 = 20$ to receiver 2. What should the AP *actually* transmit so that the receivers are able to successfully receive the samples intended for them?

22. (3 points) In the previous question, what if the channel values h_{11}, h_{12}, h_{21}, and h_{22} were all set to the same quantity y? Can the AP transmit something that would allow successful reception at the receivers? If yes, what should it transmit? If no, explain why. Hint: the answer does not depend on y.