

Carnegie Mellon
Computer Science Department.
15-744 Spring 2017 Midterm Exam

Name: _____

Andrew ID: _____

INSTRUCTIONS:

- There are 11 pages (numbered at the bottom). Make sure you have all of them.
- Please write your name on this cover and put your initials at the top of each page in this booklet **except the last**.
- If you find a question ambiguous, be sure to write down any assumptions you make.
- Be clear and concise. Limit your answers to the space provided.
- This exam has a total of 80 points.

Question	Points	Score
1	12	
2	5	
3	5	
4	3	
5	8	
6	4	
7	7	
8	8	
9	12	
10	8	
11	7	
12	1	
Total:	80	

A Miscellaneous

1. (12 points) Please circle True or False for each statement below, and for any “false” items briefly explain why or give a counterexample. For “true” items make a relevant short statement.

[True, False] Distance vector routing protocols converge more quickly than link state routing protocols.

Solution: F: DV suffers from count-to-infinity type issues

[True, False] NFV techniques require SDN since they need strict control over routing to work correctly.

Solution: F: NFV can be deployed independently from SDN.

[True, False] Portland eliminates the use of broadcast in ARP for IP-MAC address translation.

Solution: T

[True, False] Virtual Output Queues (VOQs) solves/avoids the problems of both head-of-line blocking and high switch fabric speedup requirements.

Solution: T

[True, False] The iSlip crossbar algorithm uses an extra randomization stage to ensure fairness

Solution: F: it uses round-robin service to ensure fairness

[True, False] One advantage of the CoMb middlebox approach is that it reduces expenses by reducing aggregate processing requirements.

Solution: T

2. TCP Slow start can quickly probe for network bandwidth by doubling the sender's window each round of transmission. However, it still takes many RTTs to reach the optimal window size. Google proposes to change the initial window from 4 packets (about 6KB) to 16 packets (about 24KB).

(a) (2 points) How much time will this save on a typical transfer (in terms of RTTs)?

Solution: 2 RTTs

(b) (3 points) Why would this be useful (especially) for Google?

Solution: Google serves a lot of http requests which are short-lived TCP connections. For these connections, a larger initial TCP window allows for much lower end-to-end latency from the applications' perspective.

3. Consider the equation-based congestion control (e.g., TFRC) compared to window-based congestion control (e.g., vanilla TCP).

(a) (2 points) Give one advantage of using equation-based congestion control.

Solution: The major advantage and disadvantage of equation-based congestion control is that it is slow to react to changes in network bandwidth. Depending on the perspective of the application, this can be good or bad.

(b) (3 points) What applications should use TFRC?


Solution: For applications like VOIP or live video streaming that require a constant bit-rate, sudden decreases in path bandwidth don't cause a sudden decrease in application bit-rate. Inversely, for high throughput applications (e.g., bulk data downloads) equation-based congestion control is very slow to grab new bandwidth when it is available.

4. (3 points) Are current link-layer retransmission schemes a violation of the E2E argument, why or why not?

Solution: No. It helps to improve performance and they don't attempt perfect reliability

B Routing

5. (a) (4 points) Consider the following AS topology:



List all routes that F would advertise to B through BGP, and explain why.

Solution: F and G-F (Peers don't advertise routes from other peers or providers; they only advertise routes to their own prefixes and to their customers.)

- (b) (4 points) Why do we need to infer ISP relationships? Why is this information not publicly available?

Solution: ISPs don't release this information publicly because it leaks confidential information to other companies about business relationships. For example, you could learn which companies are paying each other more money or where (geographically) companies are currently expanding.

C Router Design

6. (4 points) RED drops packets before the queue gets full. List two advantages that RED provides over drop-tail queuing.

Solution:

- Prevents burst losses
- Prevents lock-out / shares b/w more fairly
- Prevents full queue problems / shorter queuing delays

7. Answer the following about Valiant Load Balancing (VLB) as used in systems such as the passive optical switch system discussed in class. Assume that the link speed is R and the switch has N ports. A VLB switch is composed of two switch stages, each with N input ports and N output ports.

- (a) (2 points) In each stage, how fast are the links that connect an input port to an output port?

Solution: R/N

- (b) (3 points) What does the first stage do with the input traffic

Solution: It randomly mixes the input traffic,


- (c) (2 points) In combination, what valuable property does the VLB switch accomplish with the two stages?

Solution: It is able handle forwarding non-uniform traffic even on a uniform switch interconnect

D Congestion Control

8. The picture above shows the famous TCP saw tooth behavior. We are assuming that fast retransmit and fast recovery always work, i.e. there are no timeouts and there is exactly one packet lost at the end of each “tooth”. We are assuming that the flow control window is large and that the sender always has data to send, i.e. throughput will be determined by TCP congestion control.

In the picture, W represents the congestion window size at which a congestion packet loss occurs (expressed in maximum transfer units). You can assume that W is large, so feel free to approximate $(W-1)$ or $(W+1)$ by W . STT represents the “saw tooth time” expressed in seconds.



In this question, you will calculate the average throughput T for this connection as a function of the roundtrip time (RTT), the maximum transfer unit size (MTU), and packet loss rate p for this connection. Please use the notation suggested by the figure, i.e. W and STT , as intermediate values if you need them.

- (a) (2 points) Express STT in terms of RTT and W

Solution: $STT = RTT * W/2$

- (b) (2 points) Based on the shape of saw tooth, how many units can be sent within each STT ?

Solution: $\frac{1}{2}(W/2 + W) * W/2 = \frac{3}{8}W^2$

- (c) (2 points) Since exactly one packet (unit) loss event happens at the end of each STT, express p in terms of W ?

Solution: $p = \frac{8}{3*W^2}$

- (d) (2 points) TCP throughput (T) is in generally defined as number of bytes divided by the duration of time to send them. Given the results in the previous questions, please calculate the average TCP throughput as a function of RTT, MTU and p .

Solution: $T = \frac{\frac{3}{8}W^2*MTU}{RTT*W/2} = \frac{3}{4} \frac{W*MTU}{RTT} = \sqrt{\frac{3}{2}} \frac{MTU}{\sqrt{p*RTT}}$

E Data Centers

9. Answer the following questions about data center traffic and DCTCP.

(a) (2 points) What type of workloads result in Incast?

Solution: synchronized response with a barrier sync stage

(b) (2 points) Why does this workload result in poor TCP performance?

Solution: Timeout of one flow... cause the rest to all wait..

(c) (2 points) Briefly describe how a DCTCP endpoints react differently to congestion than standard TCP.

Solution: DCTCP reacts to amount of congestion by counting ECN marks

(d) (2 points) Give one reason that Clos Networks and Fat Trees topologies are commonly used in data centers

Solution: - Avoids single point of failure
- Use relatively cheap hardware to achieve high aggregate bandwidth
- Helps to achieve equal bandwidth at every bisection

(e) (4 points) Give two reasons that data center networks are large Ethernet networks and don't use IP routing

Solution: - Easy to migrate VMs among physical machines
- Easy to configure and setup

F SDN

10. The 4D paper proposes a clean slate design of network control and management. Please list the four planes in 4D architecture, and use one sentence to briefly describe what each plane does.

(a) (2 points) -----

Solution: Decision

(b) (2 points) -----

Solution: Discovery

(c) (2 points) -----

Solution: Dissemination

(d) (2 points) -----

Solution: Data

11. Answer the following about programmable data-planes.

(a) (4 points) OpenFlow processing is often called match-action processing. List 3 possible actions that OpenFlow can take on a packet

Solution: Forward packet to port(s)
Encapsulate and forward to controller
Drop packet
Send to normal processing pipeline

(b) (3 points) What are two advantages of using P4 compared to using original OpenFlow interface?

Solution:

- Can flexibly define headers to be used.
- Allows parallel processing of different match+action stages.

Anonymous Feedback

12. (1 point) Tear this sheet off to ****receive the free points****. We'd love it if you handed it in either at the end of the exam or, if time is lacking, to the course TA later today.

(a) Name one topic that you would like added in the second half.

(b) Name one topic you wish were cut from either the first half or second half.