

Massachusetts Institute of Technology
6.829 Quiz (Fall 2016)

There are **10 questions** and **11 pages** in this quiz booklet. Answer each question according to the instructions given. You have **85 minutes** to answer the questions. The maximum score is 100.

If you find a question ambiguous, be sure to write down any assumptions you make. **Be neat and legible.** If we can't understand your answer, we can't give you credit!

Use the empty sides of this booklet if you need scratch space. You may also use them for answers, although you shouldn't need to. **If you use the blank sides for answers, please say so clearly.**

Write your name in the space below and your initials at the bottom of each page of this booklet. Please do this now.

You may use your “cheat sheets” but no other materials. Calculators are allowed.

Do not write in the table below.

1	2	3	4	5	6	7	8	9	10
/12	/4	/6	/20	/10	/10	/10	/8	/6	/14

Total score:

Name:

Congestion control

1. [12 points] Consider a network with two long-running TCP Reno flows sharing a single bottleneck router that runs the PIE queue management algorithm. Each flow increases its congestion window by $1 / \text{cwnd}$ on each ACK and reduces by one-half on a packet drop. The bottleneck link capacity is $C = 1500$ packet/s, and the minimum RTT (in the absence of queueing) for the two flows are 10 ms and 30 ms. The target queueing delay for PIE is set to 10 ms, which it achieves on average. The bottleneck link is fully utilized. You may use the fact that TCP's throughput depends on the packet loss rate and RTT, with a constant of proportionality equal to $\sqrt{3/2}$.
 - a. [3 points] What is the average number of queued packets at the bottleneck link?
 - b. [3 points] What is the ratio of the average throughput of the two flows?
 - c. [6 points] What is the average packet drop probability in this network? (You may express your answer as a fraction.)

Congestion control (cont.)

2. [4 points] PIE periodically adjusts the drop probability as


$$p \leftarrow p + \alpha \times (d - d_t),$$

where d is the current queueing delay and d_t is the target queueing delay. (In this equation, we have ignored PIE's derivative term for simplicity.) Suppose N long-lived TCP flows share a bottleneck link that runs PIE. Circle True or False for each statement below:

- a. True / False The average queueing delay increases with N .
 - b. True / False The average drop probability increases with N .
3. [6 points] The window increase rule in TCP Cubic increases by a certain amount depending on the elapsed real time, rather than increasing a little bit on each ACK. Why did the designers of Cubic adopt such a real-time-based increase rule?
- a. True / False This rule is to reduce the throughput bias against long-RTT connections compared to TCP Reno.
 - b. True / False Together with multiplicative decrease (halving the window on congestion), this rule ensures that Cubic's throughput has the same dependence on packet loss rate as TCP Reno.
 - c. True / False This rule ensures that Cubic performs better than TCP Reno on short flows.

Fair queueing

4. [20 points] The network shown below has two links and three backlogged TCP flows. Weighted fair queueing is used at each link. Flow 1 traverses both links; the other two flows traverse one link each. On the first link, flows 1 and 2 each have weight 1. On the second link, flow 1 has weight X and flow 3 has weight 1. The steady-state rates (throughputs) of the three flows are r_1 , r_2 , and r_3 , respectively.



- a. [8 points] Derive an expression for r_1 as a function of X . Show your work below.
In addition, sketch it on the graph below, showing the shape of the curve clearly.



- b. [4 points] As a function of r_1 and X , which of these is correct? (Circle all that apply.)

- | | | |
|------|-----------------|-----------------------|
| i. | $r_2 = r_1$ | $r_3 = r_1 / (X + 1)$ |
| ii. | $r_2 = 1 - r_1$ | $r_3 = 1 - r_1$ |
| iii. | $r_2 = 0.5$ | $r_3 = 1 - r_1$ |
| iv. | $r_2 = 1 - r_1$ | $r_3 = r_1 / (X + 1)$ |

- c. [8 points] For what value(s) of X does the throughput of the flows in this network achieve **proportional fairness**? (Explain your answer in the space below.)

Inter-domain routing

5. [10 points] Four autonomous systems, W, X, Y, and Z run BGP with each other to exchange reachability information about p, a prefix in network X.
- * X advertises p to Y and Z but not to W.
 - * Y advertises p to W and Z.
 - * Z advertises p to W but not to Y.
- a. [2 points] In the space below to draw the topology and advertisements.


- b. [8 points] For each of the following statements, circle True or False.

- i. True / False W **MAY** be X's transit service provider.
- ii. True / False Y and Z **MUST** be in a peering relationship
- iii. True / False X and Z **MAY** be in a peering relationship.
- iv. True / False Y **MUST** be X's transit service provider.

Wireless

6. [10 points] The figure below shows a Roofnet wireless network. The probability of a successful bi-directional packet delivery (i.e., successful delivery of a packet **and** its link-layer ACK) for each link is shown in the figure. All links run at the same bit rate and the network uses the expected transmission count (ETX) metric. You are told that the ETX path from S to T goes through Y.

If the probability of successful bi-directional delivery on the S-X link is px , what is the **smallest** value for the probability of successful bi-directional delivery on link S-Y? The answer is a function of px . Explain your answer in the space below.



Mobility with Mobile IP

7. [10 points] A mobile host with IP address 6.8.2.9 is visiting a foreign network in which it obtains the address 192.168.2.9 from a foreign agent with care-of address 18.31.0.82. The home agent's IP address is 6.8.20.16. A correspondent host with IP address 2.0.1.6 has established a TCP connection with the mobile host.

Specify below the destination IP address and source IP address for packets as they traverse the path between the correspondent host and the mobile host via the home network (there is no route optimization). If a packet is being encapsulated inside so there are multiple IP headers, specify the destination and source addresses in each header.

a. Correspondent host → Home network


b. Home network → Foreign network

Datacenter networks

8. [8 points] Circle True or False for each of these statements about the VL2 architecture.
- a. True / False VL2 employs per-packet randomized load balancing.
 - b. True / False In a VL2 network, packets are only sent on paths with the minimum number of hops to the destination.
 - c. True / False A service transmitting uncontrolled UDP traffic can impede performance isolation in a VL2 network.
 - d. True / False In a VL2 network, migrating a service from one server to another does not require changes to the routing tables in the switches.
9. [6 points] What is the **maximum** number of servers for which an **4-level** Clos topology built from **24-port switches** can provide full bisection bandwidth? Each level is a layer of switches (e.g., TORs, aggregation switches, etc.).
- a. 331,776.
 - b. 20,736.
 - c. 41,472.
 - d. 292,008.
 - e. None of the above.

Chord

10. [14 points] A Chord network with 8-bit IDs for keys and nodes is shown below. Each node shown has correct successor and finger table entries.



- [2 points] Assuming that keys are uniformly distributed, which node ID is most likely to store the largest number of keys?
- [4 points] Node 7 calls `lookup(195)`. What is the sequence of node IDs traversed to answer the query, assuming that there is no caching of node or key information at any node?
- [2 points] Suppose Node 98 is attempting to join this Chord network. It knows about Node 201. What query (lookup) does Node 98 issue via Node 201 to join the network?

- d. [6 points] Which node(s) must invoke the stabilization procedure for Node 98 to be completely and correctly integrated into the Chord ring? The stabilization code is given below.

```
// called periodically. verifies n's immediate  
// successor, and tells the successor about n.  
n.stabilize()  
     $x = \text{successor.predecessor};$   
    if ( $x \in (n, \text{successor})$ )  
         $\text{successor} = x;$   
     $\text{successor.notify}(n);$   
  
// n' thinks it might be our predecessor.  
n.notify(n')  
    if ( $\text{predecessor is nil or } n' \in (\text{predecessor}, n)$ )  
         $\text{predecessor} = n';$ 
```

END OF QUIZ