

Carnegie Mellon
Computer Science Department.
15-744 Spring 2015 Midterm Exam

Name: _____

Andrew ID: _____

INSTRUCTIONS:

- There are 7 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.
- It is better to partially answer a question than to not attempt it at all.
- Be clear and concise. Limit your answers to the space provided.

Section	Earned	Possible
A		20
B		19
C		20
D		11
E		11
F		19
Total		100

A Miscellaneous

1. IPv4 contained a header checksum that is checked at each router hop while IPv6 does not. Based on the end-to-end argument, (a) how can you justify the presence of a checksum in IPv4, and (b) give one reason why it could be eliminated in IPv6.

Solution: (a) The end-to-end principle suggests that checksums should only be done by the end points of a communication unless it can provide an some performance benefit. This was done in IPv4 since error rates were high at the time. (+4pts)

(b) These end points will be using a higher-level protocol than IP (e.g., TCP) to do the error control if they want error control to be done. Thus, the checksum at the IP level is not needed. (+3pts)

Some other reasons:

Furthermore, it will not provide an optimization since most lower-level protocols (e.g., Ethernet) already have checksums.

Network links today are much more reliable than they used to be, i.e. fewer bit errors.

Bits in the header are expensive: they incur an overhead in every packet sent for many years to come.

2. Which of the following is true about Internet topologies:
 - A. The degree distribution of the AS topology graph follows a power law distribution.
 - B. The HOT topology paper shows that the router topology graph does not have a power law degree distribution.
 - C. The failure of high fanout nodes in PL random graphs and preferential connectivity graphs are likely to cause significant failures.
 - D. The failure of high fanout nodes in HOT graphs are likely to cause significant failures.
 - E. HOT topologies have low likelihood but have high bandwidth.

Solution: a, c, e

-2 for circling something incorrect, +2 for a correct circle, max of 5, min of 0

3. What is one advantage of virtual output queuing over standard input queuing?

Solution:

eliminates head of line blocking (4pts)

4. What is one advantage of virtual output queuing over standard output queuing?

Solution:

reduces the need for speedup in the switch (4pts)

B Routing

5. Which of the following is true about BGP (circle all that apply):
- (a) An AS does not export announcements to its peering ASes.
 - (b) a BGP router periodically resends route announcements to its neighbors since it is a soft-state protocol
 - (c) BGP uses path vector instead of distance vector to remove routing loops.
 - (d) BGP prefers less specific prefixes to more specific ones.
 - (e) For two prefixes of the same length, BGP will always pick the prefix with the shortest AS Path length.

Solution: C

-1 for circling something incorrect, +5 for a correct circle, max of 5, min of 0

6. An Autonomous System (AS) claims that it “owns” an IP CIDR block such as 128.2.0.0/16 by advertising a path to 128.2.0.0/16 that has only the single number in its AS Path. For example, CMU’s routers claim to own the block 128.2.0.0/16 by advertising a route to this block with AS Path 9 (9 is CMU’s AS number). There is nothing, however, to prevent two different ASes from both claiming to own the same IP address, although one of the ASes must be mistaken.
- (a) Suppose that a router in AS 9 (Carnegie Mellon) mistakenly claims that it owns CIDR block 18.0.0.0/8 (M.I.T.) by advertising a route to 18.0.0.0/8 with AS Path 9. What are the consequences of this? Recall that one of the principles of BGP is that a router should only advertise the path that it considers the best path to a destination.

Solution: Some traffic to MIT might be forwarded to CMU, where it is dropped. (5pts)

- (b) Suppose AS 9 thinks that AS 3 drops too many packets. Using only BGP, is it possible for AS 9 to implement a policy stating that “traffic outbound from my AS should not cross AS 3” Why or why not?

Solution: Yes, assuming it has an alternative route. Prefer paths that don’t contain AS 3. Of course, if the only path to a destination contains AS 3, it can not reach that destination without going through AS 3. (5pts)

- (c) Now suppose AS 9 thinks that AS 3 generates a lot of illegal file sharing traffic. Using only BGP, is it possible for AS 9 to implement a policy stating that, “I don’t want to carry traffic from 3 to my customers” Why or why not? Assume that AS 9 does not want to deny transit to traffic from any other AS.

Solution: Not in general. Traffic from a neighbor might be from both good ASs and AS 3. BGP can only accept all the traffic by advertising a route or deny all of it by not advertising. (4pts)

C TCP, Congestion Control, and Queue Management

7. Recall that TCP's throughput depends on the RTT, the packet size (MSS), and the loss rate (p). Two connections experience the same loss rate and use the same packet size. Connection A has RTT 5ms, and connection B has RTT 10ms. Express the throughput of connection B ($tputB$) in terms of the throughput achieved by connection A ($tputA$), or indicate if there is no relationship between the two.

Solution: $tputB = \frac{tputA}{2}$ (4pts)

8. At time t , a TCP connection has a congestion window of 4000 bytes. The maximum segment size used by the connection is 1000 bytes. What is the congestion window after it sends out 4 packets and receives ACKs for all of them

- (a) If the connection is in slow-start?

Solution: 8000 bytes. In slow start, the sender increases its window for each byte successfully received. (3pts)

- (b) If the connection is in congestion avoidance (AIMD mode)?

Solution: 5000 bytes. The sender increases its window by one segment each window. (3pts)

9. When a packet arrives at a router with a full queue, the router's queueing discipline tells it which packet to drop. Two simple disciplines are drop-tail (the router drops the arriving packet) and drop-head (the router drops the packet at the head of the queue). Routers today use drop-tail queueing.

What is an advantage of drop-head queueing compared to drop-tail?

Solution: Sending hosts will be notified of congestion sooner (because they will notice the missing packet sooner). (3pts)

10. A network uses routers with fair queueing. This is *not* weighted—all weights are the same.

- (a) Two connections share the same congested gateway and have no other congested gateways. Connection A has RTT 5ms, connection B has RTT 10ms. Express the throughput of connection B ($tputB$) in terms of the throughput achieved by connection A ($tputA$), or indicate if there is no relationship between the two.

Solution: $tputA = tputB$ (4pts)

- (b) Two connections traverse the same congested gateway, but also traverse some other unshared congested gateways. Express the throughput of connection B ($tputB$) in terms of the throughput achieved by connection A ($tputA$), or indicate if there is no relationship between the two.

Solution: Without knowing the other flows in the network, you cannot find a relationship between the two. (3pts)

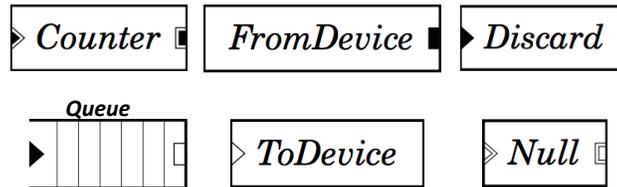
D Software Routers

11. According to the Click paper, what is the key advantage of a software router?
- Lower cost
 - Higher performance
 - Extensibility
 - Ease of configuration

Solution: C

-3 for circling something incorrect, +5 for a correct circle, max of 5, min of 0

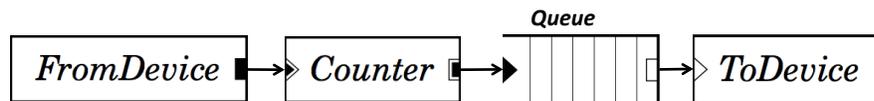
12. Use the minimum required subset of the following Click elements to show the configuration diagram of a “toy” Click router that receives a packet, updates the count of observed packets, and sends the packet out.



Recall that in each element:

- triangular ports are inputs and rectangular ports are outputs;
- solid black ports are push, solid white ports are pull, and agnostic ports are shown as push or pull ports with a double outline.

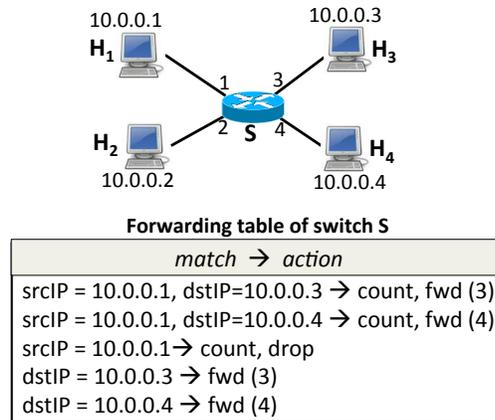
Solution: We need to accommodate the constraint that a push port and a pull port cannot be connected to each other.



+6pts, -2pts per minor error

E SDN/Verification

13. In the following software-defined network, the controller (not shown) has already populated the forwarding table of switch S as shown underneath the figure:



CORRECTION: Note that the addresses in the figure should all be 10.*

The forwarding rules in the table are in the prioritized order with the higher rules having higher priorities.

Circle True/False:

- T F Host H_1 may send traffic to host H_2
- T F Host H_1 may send traffic to host H_3
- T F All hosts can send traffic to host H_4
- T F All traffic from host H_1 is monitored (i.e., counted)
- T F All traffic to host H_3 is monitored (i.e., counted)

Solution: False, True, True, True, False

1pt each, +1 bonus if all 5 are correct (6 total)

14. Suppose we use Header Space Analysis (HSA) in a network and find a routing loop. Which of the following pieces of information can HSA provide?
- (a) What packet headers get in the loop
 - (b) Which switch forwarding rules are involved in the loop
 - (c) Which routing protocols have caused the loop
 - (d) Whether the loop is persistent

Solution: A, B

-3 for circling something incorrect, +3 for a correct circle, max of 5, min of 0

F Data Center Networks

Using a single Ethernet LAN for a large number of hosts results in two scaling issues: 1) large forwarding tables and 2) too much broadcast traffic.

15. Give one reason for connecting all hosts in a data-center to a single Ethernet instead of using IP subnetting.

Solution:

Simplifies configuration and makes it easy to move hosts/VM

4pts

16. How does the PortLand design eliminate some key broadcast traffic?

Solution:

It eliminates ARP broadcasts by intercepting ARP requests and directing them to a fabric manager.

4pts

17. Which of the following techniques does PortLand use to reduce forwarding table size
- (a) The design uses manually assigned MAC addresses that are hierarchical (i.e. like subnets)
 - (b) Edge switches rewrite the MAC address in the packet with a more easily routable address.
 - (c) End hosts must be modified to use special Ethernet VLAN labels that route packets to the right destination
 - (d) Switches forward based on MAC addresses that are based on the location of node in the data center network
 - (e) The design uses source routing techniques that encode a path through a Fat-tree topology in each packet.

Solution:

b,d

-2 for circling something incorrect, +3 for a correct circle, max of 5, min of 0

18. The coflow paper considers a number of different scheduling heuristics for deciding what flow should go next: Smallest-First, Narrowest-First, Smallest-Effective-Bottleneck-First and Shortest-First. Which one does it say is the right choice for coflow scheduling and give either one example or one reason that it makes this choice.

Solution:

Smallest-Effective-Bottleneck-First (3pts)

it considers length, width and size or a schedule that shows benefit over any one alternative (3pts)