

University of California at Berkeley
College of Engineering
Department of Electrical Engineering and Computer Sciences

CS168
MIDTERM EXAMINATION
Monday, 20 October 2014

INSTRUCTIONS—READ THEM NOW! This examination is CLOSED BOOK/CLOSED NOTES. You will not require a calculator, iPhone®, laptop computer, or other calculation aid. Please put them away right now! You MAY use one 8.5"x11" double-sided crib sheet, packed with handwritten notes, formulas, and diagrams. All work should be done on the attached pages, and there are blank pages at the end for you to use as scratch paper. PLEASE DO NOT DETACH THE SCRATCH PAPER WHEN YOU TURN IN YOUR PAPER (doing so confuses Gradescope).

Note that the sub-questions within each question are arranged in increasing order of difficulty so if you find yourself stuck, consider moving on to the next question. In general, if something is unclear, write down your assumptions as part of your answer. If your assumptions are reasonable, we will endeavor to grade the question based on them. If necessary, of course, you may raise your hand, and a TA or the instructor will come to you. Please try not to disturb the students taking the examination around you.

Please write your SID on each page!

Name – Please Print!

(Signature)

--	--	--	--	--	--	--	--

SID


Discussion Section (Day/Time)

	Question	Possible Points	Points Obtained
1	General Multiple-Choice	20	
2	Design Decisions	10	
3	Discovering Path Properties	10	
4	BGP	10	
5	Reliable Communication	10	
6	Routing	10	
Total		70	

Fun Facts to Remember

The IP header, without options, is 20 bytes
 The TCP header, without options, is 20 bytes
 The UDP header is 8 bytes.

1 TB = 10^{12} bytes; 1 GB is 10^9 bytes; 1 MB is 10^6 bytes
 1Gbps is 10^9 bits/sec; 1Mbps is 10^6 bits/sec
 1 msec is 10^{-3} seconds



<p>Version</p> <p>Version of IP Protocol. 4 and 6 are valid. This diagram represents version 4 structure only.</p>	<p>Protocol</p> <p>IP Protocol ID. Including (but not limited to):</p> <table border="0"> <tr> <td>1 ICMP</td> <td>17 UDP</td> <td>57 SKIP</td> </tr> <tr> <td>2 IGMP</td> <td>47 GRE</td> <td>88 EIGRP</td> </tr> <tr> <td>6 TCP</td> <td>50 ESP</td> <td>89 OSPF</td> </tr> <tr> <td>9 IGRP</td> <td>51 AH</td> <td>115 L2TP</td> </tr> </table>	1 ICMP	17 UDP	57 SKIP	2 IGMP	47 GRE	88 EIGRP	6 TCP	50 ESP	89 OSPF	9 IGRP	51 AH	115 L2TP	<p>Fragment Offset</p> <p>Fragment offset from start of IP datagram. Measured in 8 byte (2 words, 64 bits) increments. If IP datagram is fragmented, fragment size (Total Length) must be a multiple of 8 bytes.</p>	<p>IP Flags</p> <table border="0"> <tr> <td>x</td> <td>D</td> <td>M</td> </tr> </table> <p>x 0x80 reserved (evil bit) D 0x40 Do Not Fragment M 0x20 More Fragments follow</p> <p>RFC 791</p> <p>Please refer to RFC 791 for the complete Internet Protocol (IP) Specification.</p>	x	D	M
1 ICMP	17 UDP	57 SKIP																
2 IGMP	47 GRE	88 EIGRP																
6 TCP	50 ESP	89 OSPF																
9 IGRP	51 AH	115 L2TP																
x	D	M																
<p>Header Length</p> <p>Number of 32-bit words in TCP header, minimum value of 5. Multiply by 4 to get byte count.</p>	<p>Total Length</p> <p>Total length of IP datagram, or IP fragment if fragmented. Measured in Bytes.</p>	<p>Header Checksum</p> <p>Checksum</p>																

1. General Multiple-Choice [20 points]

[Notes: Correct answers in red.]

Indicate your answer by circling

- 1) In a reliable transport protocol, the sender must send packets in sequential order.
 - a) True
 - b) False**

- 2) Which of the following mechanisms does TCP use to avoid overloading the receiver?
(circle one)
 - a) Flow control**
 - b) Congestion control
 - c) None of the above

- 3) The [10.0.0.0/8](#) network is twice as big as [10.0.0.0/16](#) in terms of the potential number of hosts in the network.
 - a) True
 - b) False**

- 4) Routers that use virtual output queuing will not suffer from head-of-line blocking.
 - a) True**
 - b) False

- 5) Which of the following prefixes is contained in the CIDR prefix 201.10.0.0/21?
 - a) 201.10.4.0/24**
 - b) 201.10.8.0/23

- 6) Which of the following statements about the link-state algorithm is/are true? (circle all that apply)
 - a) Each router sends its routing table to its neighbors.
 - b) Each router needs to learn the complete network topology.**
 - c) Route update messages are flooded.**
 - d) The link-state algorithm suffers the count to infinity problem.

- 7) The IPv4 checksum field is computed over the IP header and payload.
 - a) True
 - b) False**

- 8) Suppose the one-way delay for packets going from Berkeley to Paris is roughly 100milliseconds and the bandwidth available to a connection is roughly 10Mbps. Assuming you have infinite amounts of data to send, what is the minimum window size of the sender to fully utilize the bandwidth?
 - a) 50000 bits (5×10^4)

- b) 100000 bits (10^5)
 - c) 1000000 bits (10^6)
 - d) 2000000 bits (2×10^6)**
 - e) 20000000 bits (2×10^7)
- 9) Which of the following is *not* a benefit of a layered communication model?
- a) Reduced system design complexity
 - b) Easier to achieve high performance**
 - c) Interoperability among devices, protocols, and applications
- 10) The end-to-end argument provides a rationale for the adoption of a best-effort service at the network layer.
- a) True**
 - b) False
- 11) Where are forwarding lookups done in a typical router? (circle one)
- a) Route/Control Processor
 - b) Input linecards**
 - c) Interconnect fabric
 - d) Output linecards
- 12) Sliding window is better than simple Stop-and-wait, most especially when the path is: (circle one)
- a) low bandwidth, low latency
 - b) low bandwidth, high latency
 - c) high bandwidth, low latency
 - d) high bandwidth, high latency**
- 13) Which of the following statements about CIDR (classless inter-domain routing) are true? (circle all that apply)
- a) CIDR simplifies the forwarding lookup logic at routers.
 - b) CIDR reduces the forwarding table size at routers.
 - c) CIDR allows more efficient use of the IP address space.**
 - d) CIDR requires an explicit prefix length for an IP address range.**
- [Notes: b) could be true and false. Since it is ambiguous, we accept both choices for this option.]
- 14) The purpose of a TCP handshake is: (circle all the apply)
- a) Connection establishment**
 - b) Exchange of initial sequence number**
 - c) Transfer of client's request data
 - d) To indicate a packet loss
- 15) TCP fast retransmit can be triggered by: (circle all that apply)
- a) Too small TCP timeout value

- b) **Packet loss**
 - c) **Packet reordering**
 - d) TCP SYNs
- 16) For an IPv4 router, which of the following header fields is typically examined or updated during packet forwarding? (circle all that apply)
- a) Source IP address
 - b) **Destination IP address**
 - c) Protocol number (e.g., 6=TCP, 17=UDP)
 - d) **Checksum**
- 17) Which of the following is provided by TCP, but not by UDP? (circle all that apply)
- a) **Reliable transfer**
 - b) **Congestion control**
 - c) Mux and demux from/to application processes
 - d) **Byte-stream abstraction**
- 18) The MF (More Fragments) field in the IPv4 header would be unnecessary if: (circle one)
- a) There was no packet reordering in the network.
 - b) There was no packet loss in the network.
 - c) **There was a field that specifies the length of the original IP packet.**
 - d) None of the above
- 19) The potential benefits of multi-homing include: (circle all that apply)
- a) **High availability**
 - b) **High throughput**
 - c) Compact routing tables
- 20) According to the end-to-end principle, link-level retransmission does not provide any benefits if reliable data transfer is already implemented by end hosts.
- a) True
 - b) **False**

2. Design Decisions [10 points]

[Notes: We assigned points for correctly selecting or omitting each option as appropriate. Additional explanation inserted as needed.]

A challenge in networking is that networks are built by, and for, parties with very different concerns and incentives. In this question, we'll consider system design decisions from the viewpoint of different parties. In each case, circle the design decisions you'd agree with.

1) Viewpoint: Router vendor

Recall that router hardware includes linecards and a general-purpose CPU (the "control processor"). A router vendor is interested in building an ultra low-cost router and their analysis reveals that linecards constitute 90% of the cost of building a router while the control processor represents the remaining 10%. The router vendor thus decides to focus on reducing the cost of linecards. Which of the following strategies would help the vendor achieve this goal? (*Circle all that apply.*) Keep in mind that high-speed memory is expensive!

- a) Develop a faster implementation of Dijkstra's algorithm
[No. A faster implementation of Dijkstra's algorithm may reduce the cost of CPUs, but not linecards.]

- b) Successfully lobby the IETF (the Internet Engineering Task Force, the standards organization that develops Internet standards) to eliminate support for fragmentation in IPv4**
[Yes. As linecards take care of packet fragmentation, we can reduce the cost of linecards by removing support of fragmentation.]

- c) Successfully lobby the IETF to eliminate the use of the MED attribute in BGP
[No. Control plane messages are handled by the control processor.]

- d) Successfully lobby the IETF to abolish multi-homing**
[Yes. By abolishing multi-homing, we can reduce the size of routing table, and thus save the cost of high speed memory.]

- e) Successfully convince Microsoft to modify their operating systems to implement traffic "pacing" such that an end host will not generate bursts of back-to-back packets**
[Yes. Linecards take care of queueing. If end hosts do not generate bursts, we can reduce the queue size on linecards, and thus save the cost of high speed memory.]

2) Viewpoint: Content provider

You are running a TCP-based movie streaming service called TheyTube. Several of your users are complaining about poor performance! You call in your team of five engineers to examine the problem. They all observe that the access link to the TheyTube server is a 10Gbps link and it is only 25% utilized. In addition, they all observe that the TheyTube server's CPU is only 15% utilized. Each engineer then independently spends a few days diagnosing the problem and reports back with the following conclusions and suggestions for improvements. Which of the following recommendations do you believe are correct? Circle all that apply.

- a) Engineer Bob: Given the low utilization numbers observed, Bob concludes that the poor performance users are experiencing must stem from problems at the user's access link to the network and cannot be improved by any actions the TheyTube team take.
[No. It is not sufficient to reach this conclusion.]
- b) Engineer Alice: Alice noticed that a number of packets are being retransmitted multiple times. Alice thus recommends that they modify the TCP implementation in the TheyTube server to reduce the timeout period and to reduce the fast retransmit threshold from 3 to 2 dupAcks (duplicate Acknowledgments). While this may lead to even more retransmissions, this is OK since the server's access link is only 25% utilized. With these changes, users will receive lost packets faster and hence will definitely see improved performance.
[No. It may lead to more severe congestion in the network.]
- c) Engineer John: John runs *traceroute* from the TheyTube server to impacted users and finds that a router R1 along the path repeatedly drops his traceroute messages. John concludes that the high packet drop rates at R1 are the cause of their performance problems.
[No. It is not reliable to use *traceroute* for this purpose, because *traceroute* sends ICMP packets, and some routers are configured to drop these packets while allowing other packets.]
- d) **Engineer Jane: Jane observes that TheyTube is currently using AT&T as its network provider and that Verizon is currently offering similar network connectivity at half the price. She recommends they switch to using Verizon as their network provider. This may or may not improve user performance but will at least lower TheyTube's network costs.**
[Yes.]
- e) **Engineer Joe: Joe examines the IP addresses of the users reporting problems and discovers that all these IP addresses belong to the CIDR block allocated to France Telecom. Joe believes the problem may originate due to high loss rates involving France Telecom's network. Joe recommends they contact France Telecom and ask them to diagnose and fix the problem.**
[Yes.]

3. Discovering Path Properties [10 points]

Consider a next-generation network in which routers do not implement fragmentation. Instead, when a router receives a packet that exceeds the link MTU (maximum transmission unit), it drops the packet and sends an error message back to the source of the packet; the error message contains the value of the link MTU.

1) Alice and Bob are end-hosts. Briefly sketch (in words) a simple test Alice might devise to discover the MTU of the path between Alice and Bob.

- 1. Pick a sufficiently large MTU.**
- 2. Send a packet to Bob using that MTU.**
- 3. If an error message is received, use the returned link MTU in the message, and repeat Step 2.**
- 4. Otherwise, the MTU of the path is found.**

[Notes: Starting at a small packet size and increasing upwards is inefficient. In this case, the number of steps is bound by the value of MTU, or by the logarithm of the value of MTU if you use binary search, rather than by the number of hops. We give partial credit for these answers.]

2) Now consider generalizing the above problem. Is there a test that Alice can always use to discover the MTU of each link along the path? If yes, briefly sketch such a test. If not, explain why using a counter-example.

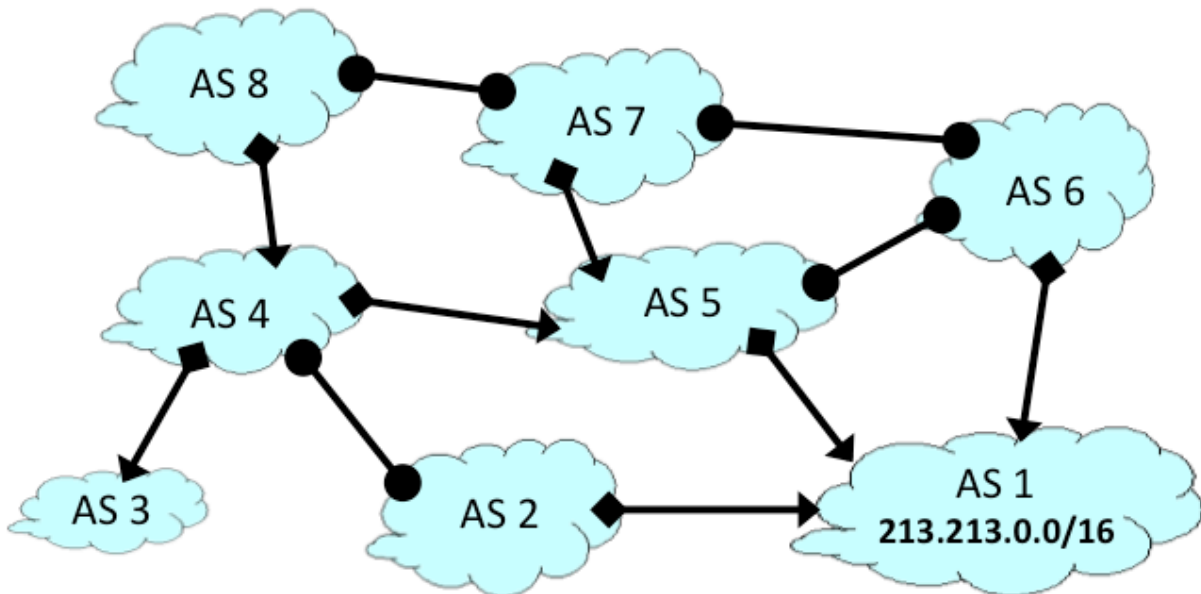

No.

Consider two links L1 and L2 between Alice and Bob, and the MTU of L1 is smaller than the one of L2. Alice cannot discover L2's MTU, because any packets reaching L2 from Alice are smaller than L2's MTU.

3) In practice, do end hosts need to discover the MTU of each link along a path for useful communication? Explain.

No. For useful communication, either routers need to support packet fragmentation, or end hosts need to discover the minimum MTU along the path.

4. BGP [10 points]



Consider the AS-level topology in the picture above. Say AS 1 starts by announcing its prefix, 213.213.0.0/16, into the network at time 0. At time 1, each of AS 1's neighbors – AS 2, AS 5 and AS 6 – propagate the route to their neighbors. And so on.

All routing messages exchanged in this network pertain to the prefix 213.213.0.0/16. You may assume that all ASes follow the Gao-Rexford rules.

(a) What route does **AS 5** send out at **time 1**? Give your answer in terms of the path-vector to the destination prefix (i.e., the sequence of ASes that AS 5 would take to reach the prefix, enumerated starting from the AS 5). Which of its neighbors does AS 5 send this path vector to?

213.213.0.0/16: AS 5 \rightarrow AS 1

This path vector is sent to AS 4, AS 6, and AS 7. [We also accept AS 1]

(b) Show the path vector that **AS 7** sends out at **time 2**. Which of its neighbors does AS 7 send the path vector to?

213.213.0.0/16: AS 7 → AS 5 → AS 1

This path vector is sent to AS 8 and AS 6. [We also accept AS 5.]

[We also accept 213.213.0.0/16: AS 7 → AS 6 → AS 1 to AS 5.]

(c) A *valid* AS level path is one that is allowed by Gao-Rexford policies. List all possible, valid, AS level paths that AS 8 learns to the prefix 213.213.0.0/16.

AS 8 → AS 7 → AS 5 → AS 1

AS 8 → AS 4 → AS 5 → AS 1

(d) Is **AS 8 → AS 7 → AS 6 → AS 1** a valid AS level path from AS 8 to the prefix 213.213.0.0/16? Why or why not?

No. AS 7 will not advertise the path AS 7 → AS 6 → AS 1 to AS 8 according to Gao-Rexford policies.

(e) Which path will AS 8 use to route to the prefix 213.213.0.0/16? Why?

AS 8 → AS 4 → AS 5 → AS 1

AS 8 will prefer routes through its customers.

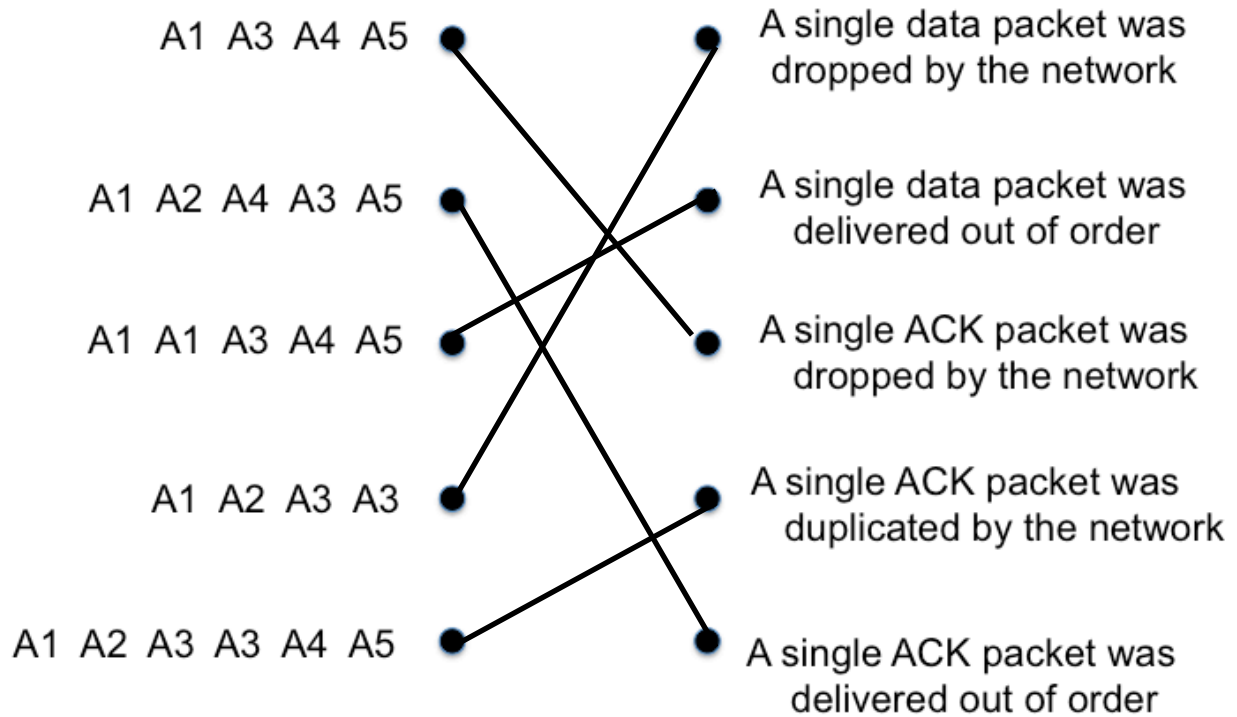
5. Sliding-Window: Anomalies [10 Points]

Consider a host using a sliding window protocol with a window size of 5 using cumulative ACKs. The host sends out the first five data packets and waits for the ACKs. Below we consider several cases where something has gone wrong (such as a packet drop, or packet reordering) and the resultant order in which the ACK packets arrive at the sending host. Match each observed ACK stream to exactly one of the potential causes. **The matching must be perfect – i.e., each possible cause is matched to exactly one ACK stream and vice versa.** The scenarios describe only the unusual events; assume all else functioned normally (i.e., packets are not dropped or reordered).

Notation (**read carefully**): The notation A_x is used to mean that the ACK packet is acknowledging the receipt of all packets *up to and including data packet x* . That is, A_5 is acknowledging the receipt of packet 5 and all those before it; *to be clear, the notation does not mean that the receiver is expecting packet 5 as the next data packet.*

Observed ACK stream

Possible causes



[A1 A3 A4 A5: A2 was dropped.]

[A1 A2 A4 A3 A5: A3 was out of order.]

[A1 A1 A3 A4 A5: The first three delivered packets are 1, 3, 2.]

[A1 A2 A3 A3: Packet 4 was dropped.]


[A1 A2 A3 A3 A4 A5: A3 was duplicated.]

6. Routing [10 points]

A certain router B in a network receives the following sequence of link state updates from other nodes in the same network. You may assume that link state updates from a node arrive in order at B.

Update from	Neighbor/cost
C	B/8, A/4
A	B/2, C/4, G/10
C	B/10, A/4
F	B/1, G/6
G	F/6, A/10

(a) Show the network topology, with nodes, edges and link costs, as constructed by router B after it receives the above link state packets.



(b) Suppose that the A—B link goes down. Say, A and B quickly re-compute shortest paths to other network nodes, but their link state updates are yet to reach their neighbors. Is there a possibility of packets looping in the network? If not, why? If yes, given an example of a potential loop.

Yes. Let's say C wants to send a packet to B. From C, the shortest path to B is $C \rightarrow A \rightarrow B$. From A, it is $A \rightarrow C \rightarrow B$. Thus, the loop will be $C \rightarrow A \rightarrow C \rightarrow A \rightarrow \dots$

(c) Consider again the topology from part (a) (i.e., before the link A-B fails) but now assume we

are running distance-vector routing over the topology. What is the path from C to F?

C → A → B → F

(d) Continuing with the topology from (c), assume we are not using poison reverse in our implementation of distance-vector and the link B-F fails. After how many rounds of exchanging routing updates with its neighbors will B have a loop-free path to F again? Recall that in one round of exchanges, B sends a route update to all its neighbors and receives route updates from each of its neighbors. You may assume that, in each round, route updates from B are received and processed before those from A and that updates from A are received and processed before C's. You may further assume that no route updates are lost or reordered.

- I. Zero. Node B has a loop-free alternate path immediately available.
- II. 1 round of route updates
- III. 2 rounds of route updates
- IV. More than 2 rounds of exchanging routing updates**

Initially, the cost from A to F is 3, from G to F is 6. After the link B-F fails, A will not choose G as the next hop to F, until the cost of A-B-F exceeds 16. After each round of exchanging updates with B, the cost of A-B-F increases by 4. Therefore, it takes $(16-3)/4 \approx 4$ rounds.

(e) Now assume we are using poison reverse in our implementation of distance-vector and the link B-F fails. After how many rounds of exchanging routing updates with its neighbors will B have a loop-free path to F again? (As in part (d), you may assume that in each round, route updates from B are received and processed before those from A and that updates from A are received and processed before C's. No route updates are lost or reordered.)

- I. Zero. Node B has a loop-free alternate path immediately available.
- II. 1 round of route updates**
- III. 2 rounds of route updates
- IV. More than 2 rounds of exchanging routing updates

Initially, the cost of path from C via A to F is 7.

After the link B-F fails, B updates its path cost to 17, and sends an update to A, C, F.

A learns that it should use a cost-16 path via G.

C, F do not update its routing table.

A sends its update to B, C, G.

B, G do not update its routing table.

C learns that it should use a cost-20 path via A.

C sends its update to A and B.

A does not update its routing table.

B learns that it should use a cost-18 path via A.