Duke University
CompSci 356 Final
Spring 2018

Name (Print):	______________________ (Family name), ______________________ (Given name)
Student ID Number:	______________________

Date of Exam:	*May 3, 2018*
Time Period:	9am-Noon
Number of Exam Pages:	14 (including this cover page)
Exam Type:	*closed book/notes*
Additional Materials Allowed:	*Two letter-size sheets of your own notes*

Marking Scheme:

Question	Score
1	/25
2	/10
3	/10
4	/10
5	/20
6	/10
7	/10
Extra credit	5/5
Total	/100

For your convenience, this page includes figures that describe packet formats for Ethernet, ARP, IP, and ICMP.

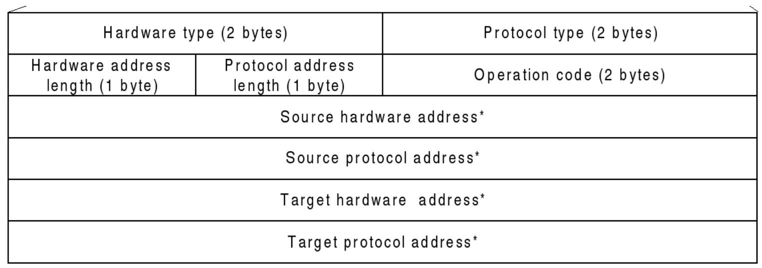

(a) ARP request or reply format.

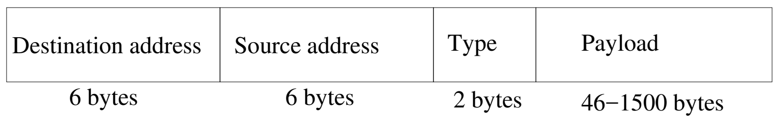

(b) Ethernet packet format.

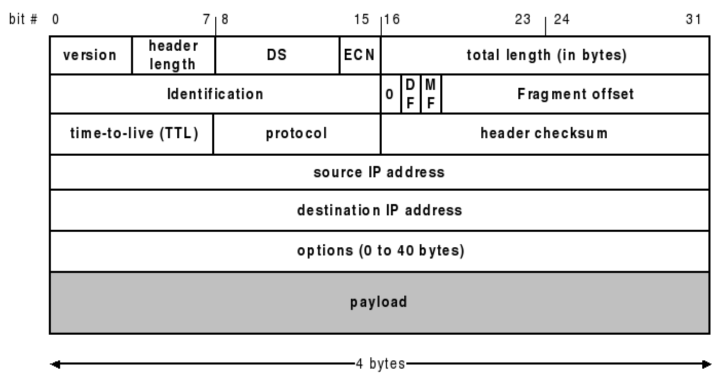

(c) IP packet format.

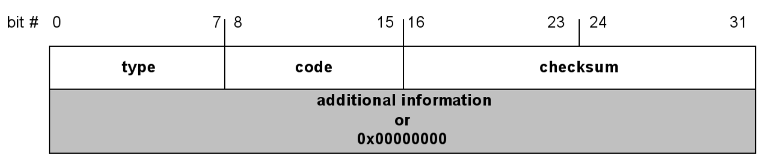

(d) ICMP header format.

1. (25 pts) Internet Basics.

   (a) (2 pts) Which of the following layers are **NOT** part of the Internet architecture? (Circle all that apply.)

      A. Link layer
      B. Network layer
      C. Transport layer
      D. Session layer
      E. Presentation layer
      F. Application layer

   (b) (4 pts) Name two or more transport layer protocols used on the Internet.

   (c) (4 pts) Name two or more popular application layer protocols used in the Internet. (Hint: application protocols are different from application programs.)

   (d) (4 pts) We learned both distance-vector and link-state based routing protocols in this class. If you are the network operator of a large ISP such as AT&T, which routing protocol will you use for your network? Explain why.

   (e) (6 pts) Four hosts (H1, H2, H3, and H4) are connected by an Ethernet switch and there is no router in the local area network. H1's IP address is 152.3.110.7/16; H2's IP address is 152.3.7.33/24; H3's IP address is 152.3.7.43/16; H4's IP address is 152.3.168.2/28. Circle all host pairs that can ping each other. (Hint: a successful ping involves two packets. Since there is no router present, each host does not have a default route, but it knows its directly connected network via its IP address configuration.)

      A. H1-H2
      B. H1-H3
      C. H1-H4
      D. H2-H3
      E. H2-H4
      F. H3-H4

(f) (5 pts) Describe at least two solutions that can alleviate the IPv4 address depletion problem, and compare their pros and cons.

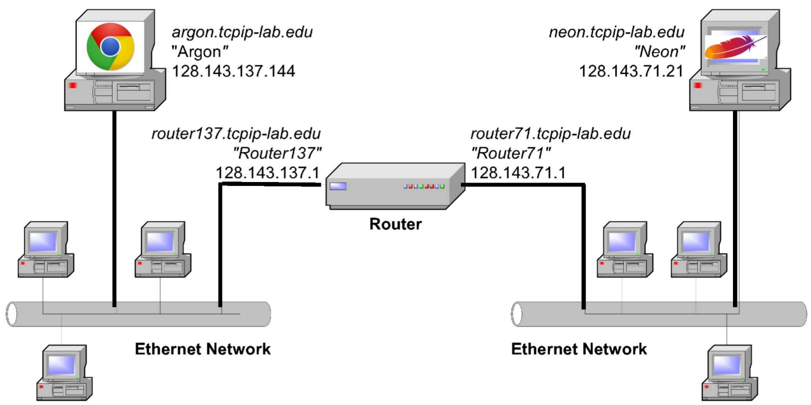

Figure 1: **A simple network topology.**

2. (10 pts) Figure 1 shows a simple network topology. Suppose a user using a web browser running on the computer Argon to access the page `http://neon.tcp-lab.edu/index.html`. As we have learned in the class, a sequence of networking events must happen in order for the web browser running on Argon to display the web page. Consider the first TCP packet sent from Argon to Neon. Please select all the events that will happen to send this first packet, and list them in the correct order. Hint: first select the relevant events and then order them.

   (a) Argon sends an ARP request for the IP address 128.143.71.21.
   (b) On Neon, the ether_input() function is called to receive an Ethernet frame.
   (c) On Argon, the ether_output() function is called to send an Ethernet frame.
   (d) On Argon, tcp_output() is called to pass a TCP packet to the IP layer.
   (e) On Argon, the web browser calls the connect() function.
   (f) Router looks up the IP address 128.143.71.1 in its forwarding table.
   (g) Argon sends an ARP request for the IP address 128.143.71.1
   (h) Argon sends an ARP request for the IP address 128.143.137.1
   (i) On Argon, the ip_output() function is called to pass an IP packet to the Ethernet layer.
   (j) On Neon, the tcp_input() function is called for the TCP layer to receive a packet from the IP layer.
   (k) Argon sends an Ethernet frame with the Ethernet address of 128.143.137.144 as the source address, and the Ethernet address of 128.143.137.1 as the destination address.
   (l) Argon sends an Ethernet frame with the Ethernet address of 128.143.137.144 as the source address, and the Ethernet address of 128.143.71.1 as the destination address.
   (m) Router sends an ARP request for the IP address 128.143.71.21.
   (n) Router looks up the IP address 128.143.137.21 in its forwarding table.
   (o) On Neon, the ip_input() function is called for the IP layer to receive a packet from the Ethernet layer.
   (p) Router looks up the IP address 128.143.137.1 in its forwarding table.
   (q) Argon sends an Ethernet frame with the Ethernet address of 128.143.137.144 as the source address, and the Ethernet address of 128.143.71.21 as the destination address.

3. (10 pts) Figure 2 shows a portion of the Internet2 topology. Each square represents a router in a given city. Suppose all routers run the OSPF routing protocol, and the numbers are link costs. Answer the following questions.

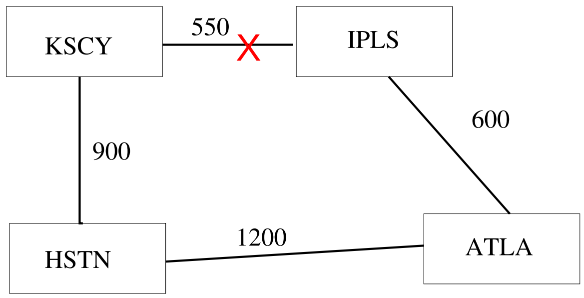

Figure 2: A portion of the Internet2's topology.

(a) (2 pts) Suppose the link between KSCY and IPLS fails at time t_1. Before the link fails, what is the next hop from KSCY for packets destined to IPLS? What is the next hop from HSTN for packets destined to IPLS?

(b) (3 pts) After the link between KSCY and IPLS fails, suppose the router KSCY updates its routing table at time t_2, where $t_2 > t_1$, and the router HSTN updates its routing table at time t_3, where $t_3 > t_2$. The router HSTN updates its routing table later than the KSCY router because it takes time to send link state messages and to recompute routing tables. What would happen to packets sent from HSTN and destined to IPLS between the time interval $[t_2, t_3)$?

(c) (2 pts) The packet delivery problem you noticed above is caused by the asynchronous update order of routing tables. In this example, what would be the correct order of routing table updates that can avoid the above problem?

(d) (3 pts) Can you generalize your answer to an arbitrary topology? That is, when a link fails in a network, what would be the correct order of routing table updates that can avoid the packet delivery problem you observed in (b)?

4. (10 pts) Figure 3 shows an AS-level topology. Assume all ASes use the common BGP policies to select forwarding paths. Answer the following questions.

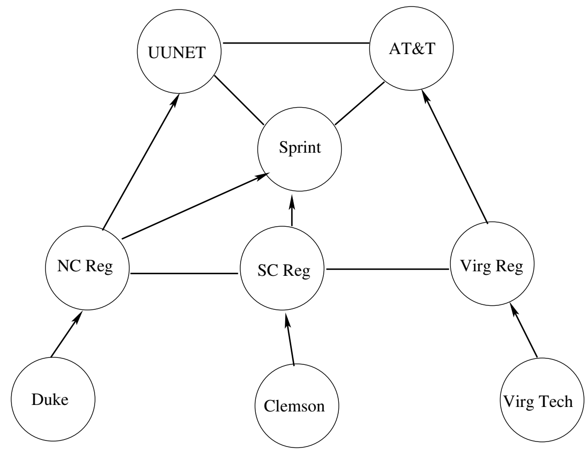

Figure 3: **An AS-level network topology. Each circle represents an AS. A solid line represents a peering relationship between two ASes, and an arrowed line represents a customer-provider relationship with the arrow pointing to the provider.**

(a) (5 pts) What are the policy-allowed AS paths for packets sent from Duke to Virginia Tech? List all such paths. For each path, specify each AS on the path.

(b) (5 pts) According to the common BGP policies, what is the best path between Duke and Clemson?

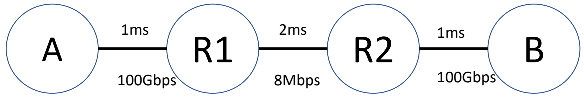

Figure 4: A four-node network topology.

5. (20 pts) Understanding the relation between delay-bandwidth product and router buffer size. Figure 4 shows a simple network topology. Node A is sending a large file to Node B using TCP. Suppose the maximum MSS is 1K bytes and the ACK size is 40 bytes. Suppose the router R1 keeps a buffer size of Buf. Let $Buf = 4MSS$. The initial slow start threshold *ssthresh* is set to a very large value (i.e., infinite).

   (a) (1 pt) Which link is the bottleneck link in this network? Circle the correct answer.

      A. Link between A and R1

      B. Link between R1 and R2

      C. Link between R2 and B

   (b) (2 pts) Suppose A is in the slow start mode and its initial window size is 1 MSS. In the first RTT, A sends one MSS packet to B. How long does it take for A to receive an ACK from B? You may ignore all processing delay, and a packet and ACK's transmission delay on links A-R1 and R2-B.

   (c) (2 pts) After A receives the ACK, it will increase its congestion window size by 1 MSS. Suppose in the second RTT, A sends out a window size of packets (which is 2 MSS packets) back-to-back at 100Gbps. When B receives a data packet from A, it sends an ACK immediately. What is the time interval between the ACKs A receives from B? Suppose R2 immediately forwards a packet to B when it receives the packet from R1. Hint: packets cannot arrive at B faster than what R1 can send.

(d) (1 pt) If A sends at a speed faster than the link bandwidth between R1 and R2, what will happen?

(e) (2 pt) What's the effective delay bandwidth product of this connection? That is, how many MSS packets should A transmit in one window to keep the bottleneck link fully utilized?

(f) (1 pt) What's A's throughput when its congestion window size equals to the effective delay bandwidth product of this connection?

(g) (1 pt) After A's congestion window size exceeds the effective delay bandwidth product, packets will be queued at R1. Given router R1's buffer size is 4 MSS, what is the maximum TCP congestion window size before the TCP sender loses a packet?

(h) (1 pt) What's A's congestion window size after it detects a packet loss after three duplicate acknowledgment?

(i) (1 pt) What's A's throughput in the RTT right after it reduces its congestion window size after a packet loss? Can this window size keep the bottleneck link fully utilized?

(j) (1 pt) Now let's change R1's buffer size to 20 MSS. What's A's maximum window size before it suffers a packet loss?

(k) (2 pts) When $Buf = 20$ MSS, what's A's throughput in the RTT right after it reacts to a packet loss? Hint: can A fully utilize the bottleneck's bandwidth and can it achieve a higher throughput than the bottleneck's bandwidth?

(l) (1 pt) When $Buf = 20$MSS, and when A reaches its maximum window size, how long will a packet wait in R1's buffer before it is transmitted to the next hop? We refer to this wait time as queuing delay.

(m) (2 pts) From the above analysis, can you describe the pros and cons of a too large and a too small router buffer respectively?

(n) (2 pts) What will be the "optimal" router buffer size at R1 such that the node A can keep the bottleneck link fully utilized but does not suffer unnecessary queuing delay?

6. (10 pts) When a TCP sender suffers a packet loss, it reduces its congestion window size by half, and when there is no packet loss and the TCP is in congestion avoidance, it increases its congestion window size by one maximum segment size (MSS) per round trip time (RTT). Ben BitDiddle argued that this congestion avoidance mechanism makes the TCP window size oscillate two much, and hurts TCP's performance. He argues that instead, when a TCP sender suffers a packet loss, it should reduce its congest window size by 1/8 of its current window size. With this modification, can a TCP sender fully utilize a bottleneck link's bandwidth? If there are two TCP senders sharing the same bottleneck link, will this modification ensure that the two TCP senders can share the bottleneck fairly? Explain why. Use a phase plot where there are two axes ($sender_1$ and $sender_2$) and each axis represents one sender's sending rate to explain your answer.

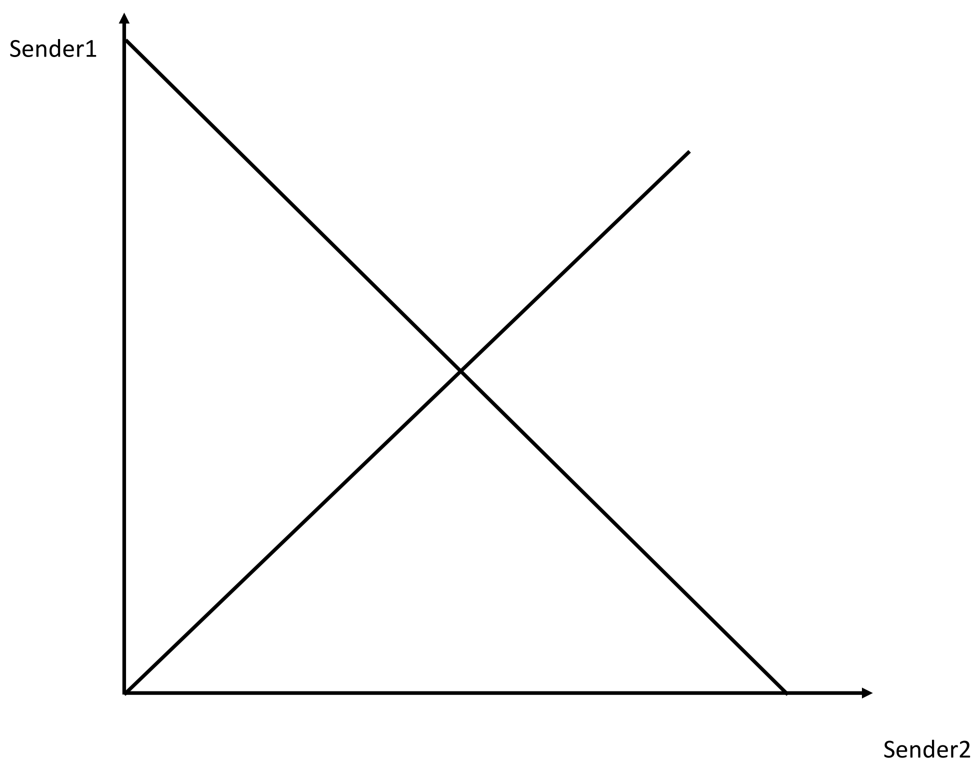

Figure 5: **Use this figure to explain whether the modified TCP can achieve efficiency and fairness.**

7. (10 pts) A non-recursive DNS query is sent to a root DNS server `@a.root-servers.net.` for the domain name `www.duke.edu`, and the following answer is received:

```
  ; <<>> DiG 9.10.3-P4-Ubuntu <<>> +norecurse @a.root-servers.net. www.duke.edu.
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 28793
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 6, ADDITIONAL: 8

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1472
;; QUESTION SECTION:
;www.duke.edu. IN A

;; AUTHORITY SECTION:
edu. 172800 IN NS a.edu-servers.net.
edu. 172800 IN NS c.edu-servers.net.
edu. 172800 IN NS d.edu-servers.net.
edu. 172800 IN NS f.edu-servers.net.
edu. 172800 IN NS g.edu-servers.net.
edu. 172800 IN NS l.edu-servers.net.

;; ADDITIONAL SECTION:
a.edu-servers.net. 172800 IN A 192.5.6.30
c.edu-servers.net. 172800 IN A 192.26.92.30
d.edu-servers.net. 172800 IN A 192.31.80.30
f.edu-servers.net. 172800 IN A 192.35.51.30
g.edu-servers.net. 172800 IN A 192.42.93.30
l.edu-servers.net. 172800 IN A 192.41.162.30
g.edu-servers.net. 172800 IN AAAA 2001:503:cc2c::2:36

;; Query time: 13 msec
;; SERVER: 198.41.0.4#53(198.41.0.4)
;; WHEN: Wed May 02 13:06:11 EDT 2018
;; MSG SIZE  rcvd: 276
```

(a) (5 pts) Which servers are the authoritative servers for the zone `edu.`?

(b) (5 pts) What additional non-recursive DNS queries need to be sent to resolve the IP address of `www.duke.edu`? Specify the servers to which the queries will be sent.

Scratch page