CS/ECE 438: Computer Networks	**Spring 2026**

Homework 3

Handed Out: March 24, 2026 — *Due: 11:59pm, April* 10, 2026

TA: Xiaojuan Ma

- Homework assignments must be submitted online through GradeScope. No handwritten solutions will be accepted. LaTeX formatting is recommended, but not enforced.
- Each question is worth the same number of points (e.g., 10 points). Show your work for full credit. You will not receive points for a correct answer with no work shown. Partial credit may be awarded for correct reasoning even if the final answer is incorrect.
- Please come to office hours if you have questions about the homework.
- While we encourage discussion within and outside of the class, cheating and copying is strictly prohibited. Copied solutions will result in the entire assignment being discarded from grading at the very least and a report filed in the FAIR system.

1 IP Fragmentation – Part 1

Consider two hosts, A and B. A is connected to an Ethernet LAN with MTU= 1300 bytes and B is connected to a Wi-Fi WLAN with MTU= 1250 bytes. In addition to these LAN's, the route connecting host A to host B through the Internet contains an additional hop over a point-to-point link between a router on A's Ethernet and a second router on B's Ethernet. The point-to-point link has MTU= 470 bytes.

Recall that MTU is the maximum amount of data that can be sent in a frame at the physical layer and thus includes all TCP and IP headers (assume each of which occupies 20 bytes). Also recall that IP fragmentation breaks its data along 8 byte boundaries.

a. An application on host A passes 2800 bytes of data to TCP. How many bytes are delivered to the network layer protocol at host B? *(Hint: For each link in the route, draw each packet traversing it as a sequence of labeled fields showing the header(s) and payload size in bytes. The sketch will not be graded but will help you arrive at the answer.)*

b. If the probability that any IP datagram crossing any link arrives intact (without error) is given by p, calculate the probability that the entire 2800 bytes sent in part (a) arrives without the need for retransmission.

c. Calculate the probability of retransmission for each fragment given p = 9/10, where p is as defined in part (b).

d. Most IP datagram reassembly algorithms have a timer to avoid having a lost fragment tie up reassembly buffers forever. Suppose a datagram is fragmented into four fragments. The first three fragments arrive, but the last one is delayed. Eventually the timer goes off and the three fragments in the receiver's memory are discarded. A little later, the last fragment stumbles in. What happens to this last fragment at the receiver?

2 IP Fragmentation – Part 2

Suppose an IP packet is fragmented into 7 fragments, each with a 2% (independent) probability of loss.

a. What is the probability of losing the whole packet due to at least one fragment being lost?

Based on the previous answer, what is the probability of net loss of the whole packet if the packet is transmitted three times…

b. assuming that all received fragments must belong to the same transmission to recover the original packet (where one "transmission" = one full attempt to send all 7 fragments)?

c. assuming any fragment from any transmission can be used to reconstruct the original packet?

d. Explain how use of the *Ident* field might be applicable here.

3 Forwarding and Classless Inter-domain Routing (CIDR)

a. Consider a router that interconnects three subnets: Subnet A, Subnet B, and Subnet C. Suppose all of the interfaces in each of these three subnets are required to have the prefix `172.16.2.0/23`. Also suppose that Subnet A is required to support at least **240** interfaces, Subnet B must support at least 90 interfaces, and Subnet C must support at least 50 interfaces. Provide three network addresses (of the form a.b.c.d/x) that satisfy these constraints (Note: Allocate the subnets in the order of A, B, then C, starting from the lowest available address).

Subnet	Requirement	Block Size	Usable Addresses	Network Address
A	240			
B	90			
C	50			

b. Suppose a router has built up the routing table shown below. For CIDR addresses, "/20" indicates a mask of 20 1's followed by 12 0's.

Network Prefix / Mask	**Next Hop**
128.174.192.0 / 18	Router 3 (R3)
128.174.224.0 / 20	Router 2 (R2)
128.174.236.0 / 23	Interface 2 (I2)
128.174.238.128 / 26	Interface 1 (I1)
128.174.238.16 / 28	Router 1 (R1)
128.174.238.7	Interface 3 (I3)
default	Router 4 (R4)

The router can deliver packets directly over interfaces 1, 2, or 3, or it can forward to routers 1, 2, 3, or 4. Specify the next hop for each of the following destinations. Recall that if a destination matches more than one IP prefix in the table, the *longest match* is used.

No.	**Destination IP**	**Next Hop**
i.	128.174.240.7	
ii.	128.174.238.7	
iii.	128.174.237.7	
iv.	128.174.225.7	
v.	128.174.190.7	
vi.	128.174.238.20	

c. For each of the entries in the routing table provided in Part (b), calculate how many individual IP addresses (usable hosts) match that prefix. You must show your work by identifying the number of host bits available and applying the appropriate formula.

(Note: For this question, ignore any overlapping address space between prefixes; calculate the capacity for each entry as an independent subnet. Add space if needed.)

Next Hop	**Reasoning / Calculation**	**Usable Hosts**
R3 (/18)		
R2 (/20)		
I2 (/23)		
I1 (/26)		
R1 (/28)		
I3 (Host)		
R4 (Default)		

4 Networking Utilities

Show the commands that you use to solve the problem and the output you get. You must show your work to receive full credit.

a. The **ifconfig** utility is used to assign and examine network interface parameters. Read the man page on ifconfig and study the different options. Use the command to determine the Internet broadcast address for a machine that you have logged on.

b. The **arp** utility can be used to display the Internet-to-Ethernet address translation table in the host it is running on. Consult the man page for arp, and use it to determine the Ethernet address of 3 different machines.
(Hint: Since entries in the address table time-out, the table may be empty sometimes. To fill out the table, send a ping to the broadcast address you obtained in part (a). The host will inspect the return ping packets to build up its address translation table.)

5 BGP Policy

Fig. 1 represents the relationship between ASes. The vertices are individual ASes and edges are links between them and IP is the prefix from AS E. Also suppose that arrows represent customer-provider relationships where the customer points to its provider. An edge without arrows represents a link between peers.

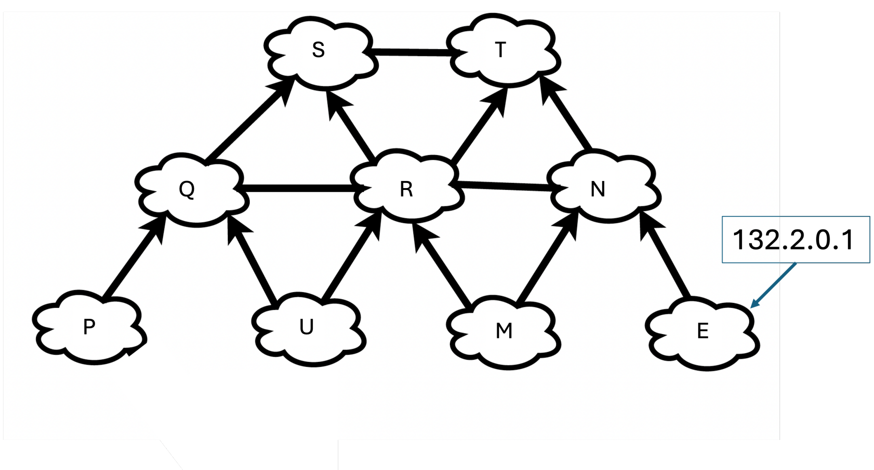

Figure 1: BGP: Relationship between ASes.

a. Suppose all ASes follow local preference rules that enforce *valley-free* paths: any path must follow a sequence of zero or more provider links, followed by at most one peer link, followed by a sequence of customer links. An AS will route through the *valley-free* path with the least number of hops. List the routes that each AS will follow to reach E in a *valley-free* manner.

b. Suppose AS S does not like AS R. Using only BGP, is it possible for AS S to implement a policy stating that "traffic outbound from my AS should not cross R"? If it is possible,

AS	Path to E
N	
R	
M	
T	
S	
U	
Q	
P	

Table 1: BGP: Table for part (a)

show that S can still reach all ASes using *valley-free* paths that do not cross R. If it is not possible, show that there exists an AS such that any *valley-free* path from S must go through R.

c. Suppose AS R does not like AS P, and therefore decides to not forward any traffic from P. Can AS P deal with this change? If it can, show that P can find *valley-free* paths to all ASes that do not cross R. If it cannot, show that there exists an AS such that any *valley-free* path from P must go through R.

6 SDNs

Consider the SDN OpenFlow network shown below in Fig. 2. Suppose we want switch s2 to function as a firewall. Specify the flow table in s2 that implements the following firewall behaviors. Specify a different flow table for each of the four firewalling behaviors below. The flow table should only consider delivery of datagrams destined to h5 and h6. You do not need to specify the forwarding behavior in s2 that forwards traffic to other routers. The flow table should show the matching rule and the action taken. Add rows or leave rows empty as needed.

a. Only traffic arriving from hosts h1 and h2 should be delivered to hosts h5 or h6.

Match Rule	Action

b. Only TCP traffic is allowed to be delivered to hosts h5 or h6. (i.e., that UDP traffic is blocked.)

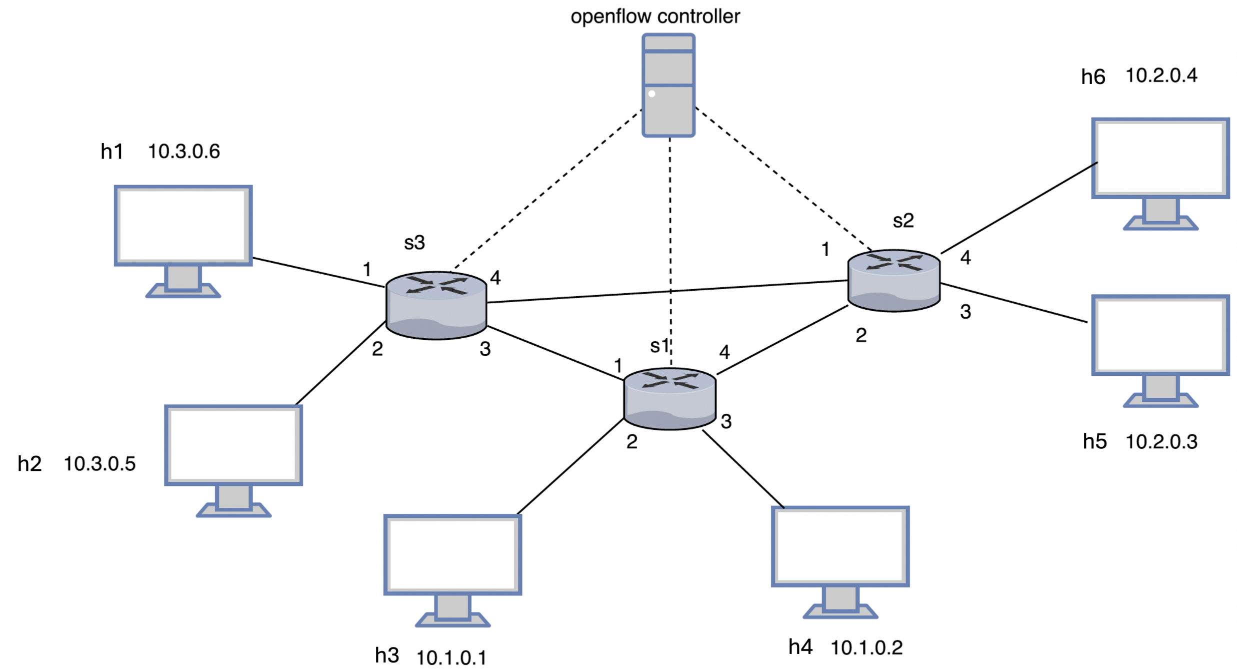

Figure 2: SDNs

Match Rule	**Action**

c. Only traffic destined to h5 is to be delivered (i.e., all traffic to h6 is blocked.)

Match Rule	**Action**

d. Only UDP traffic from h4 and destined to h6 is to be delivered. All other traffic is blocked.

Match Rule	**Action**