CS/ECE 438: Communication Networks for Computers Spring 2010
Problem Set 5 Due: start of class, Wednesday, April 7th.

Internet Protocol (IP)

Assigned reading: Peterson and Davie: Chapter 4. All problems carry equal weight. Please show all your work.

1. Link-State Routing and BGP

(a) Show how the link-state algorithm builds the routing table for router A in AS 38 in the following network. Use the same format as in P&D (Table 4.9, page 282).

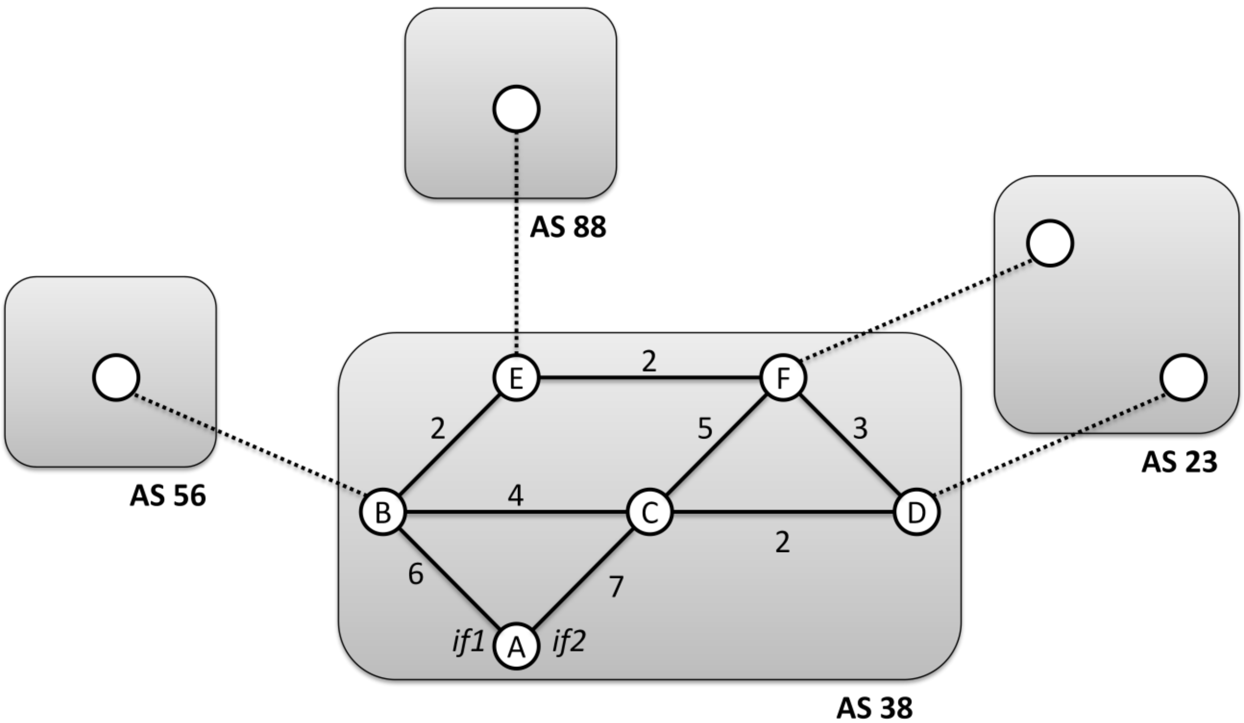

Step	Confirmed	Tentative
1	(A, 0, -)	
2	(A, 0, -)	(B, 6, B) (C, 7, C)
3	(A, 0, -) (B, 6, B)	(C, 7, C)
4	(A, 0, -) (B, 6, B)	(C, 7, C) (E, 8, B)
5	(A, 0, -) (B, 6, B) (C, 7, C)	(E, 8, B)
6	(A, 0, -) (B, 6, B) (C, 7, C)	(E, 8, B) (F, 12, C) (D, 9, C)
7	(A, 0, -) (B, 6, B) (C, 7, C) (E, 8, B)	(F, 12, C) (D, 9, C)
8	(A, 0, -) (B, 6, B) (C, 7, C) (E, 8, B)	(F, 10, B) (D, 9, C)
9	(A, 0, -) (B, 6, B) (C, 7, C) (E, 8, B) (D, 9, C)	(F, 10, B)
10	(A, 0, -) (B, 6, B) (C, 7, C) (E, 8, B) (D, 9, C)	(F, 10, B)
11	(A, 0, -) (B, 6, B) (C, 7, C) (E, 8, B) (D, 9, C) (F, 10, B)	

(b) Suppose that the gateway routers of AS 38 receive the following BGP advertisements from their BGP peers:

Network	**AS Path**

AS 56	
1.2.3.0/24	56 83 99
1.3.8.0/23	56 75
1.4.8.5/24	56 97

AS 88	
2.3.0.0/16	88 107 56 23
1.4.8.4/24	88 62 103

AS 23	
1.2.3.0/24	23 99
7.12.0.0/16	23 99 117

Show the routing table in router *A* that will be formed as a result of these advertisements. Assume that no routes are rejected due to local policy rules. Represent the routing table in this form:

Network	**Interface**
5.5.5.0/24	if1

Network	**Interface**	**Comments**
1.2.3.0/24	if2	Shorter AS path (23 has route shorter than 56), and Hot potato routing (F vs D)
1.3.8.0/23	if1	
1.4.8.4/23	if1	Aggregated prefix
2.3.0.0/16	if1	
7.12.0.0/16	if2	Hot potato routing (F vs D)

2. Distance-Vector Routing

Consider the following network configuration where the routers calculate shortest routes using the Distance Vector Routing Protocol.

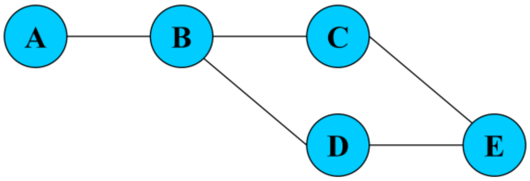

Initially, the router tables for routes to node A look like the following:

B		C		D		E	
Cost	Next Hop	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop
1	A	2	B	2	B	3	D

Now, assume node A goes down.

Given the following sequence of routing update messages, fill in the table for the routing entries for reaching A at each event, where the notation B → C indicates that node B sent a routing update to node C.

	B		C		D		E	
Event	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop	Cost	Next Hop
	1	A	2	B	2	B	3	D
Node A goes down	∞	-	2	B	2	B	3	D
B → C	∞	-	∞	-	2	B	3	D
E → C	∞	-	4	E	2	B	3	D
C → B	5	C	4	E	2	B	3	D
B → D	5	C	4	E	6	B	3	D
D → E	5	C	4	E	6	B	7	D
E → C	5	C	8	E	6	B	7	D
C → B	9	C	8	E	6	B	7	D
B → D	9	C	8	E	10	B	7	D

Analyze this carefully and try to see what's happening compared to what should happen ideally. So what, according to you, is the problem here? Why is it a problem?

3. IP Fragmentation

Consider two hosts, A and B, each on a separate shared Ethernet with MTU=1200 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=600 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 (each of which occupies 20 bytes). Also recall that IP fragmentation breaks data along 8 byte boundaries.

a. An application on host A passes 2000 bytes of data to TCP. Following the approach used to draw Figure 4.4 (page 241 of P&D) but including the TCP header, sketch the packets that cross each link in the route. How many bits are delivered to the network layer protocol at host B?
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 2000 bytes sent in part (a) arrives without the need for retransmission.
c. Calculate the average amount of data, including TCP and IP headers, and including all transmissions and retransmissions, that must be sent by host A in order to successfully deliver the 2000 bytes to an application on host B given $p = 9/10$, where p is as defined in part (b). Assume that host B will buffer any data that arrives at the TCP level.
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?

Sol:

a.

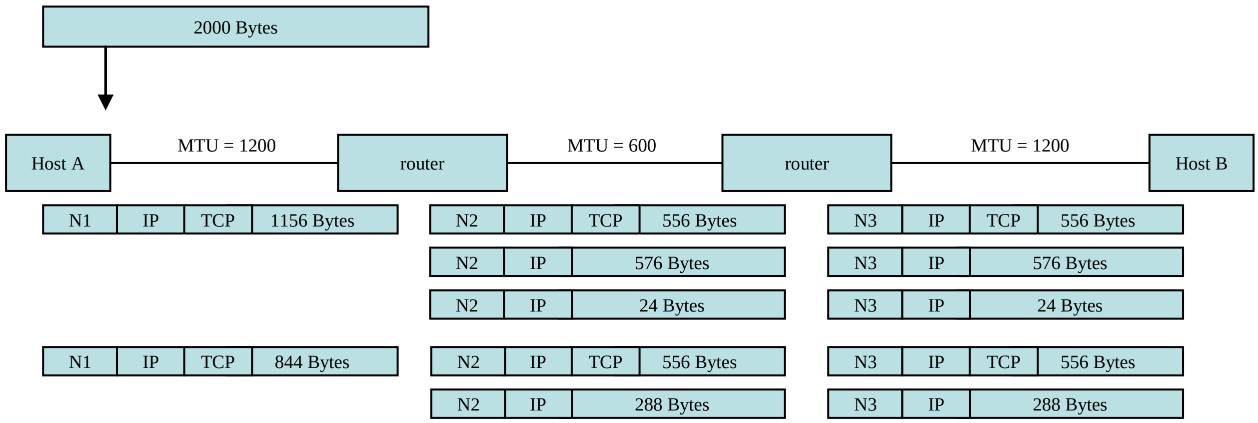

The network layer receives 5 IP headers, 2 TCP headers and 2000 bytes of data = 2140 bytes.

b. The probability of success with no retransmissions is equal to the probability that none of the fragments gets lost = p^{12}.

c. When IP reassembly fails due to fragment loss, retransmission occurs at the TCP level. The first TCP packet has a probability of p^7 to be reassembled per transmission, so $1/p^7$ = 2.09 transmissions are necessary on average. The second TCP packet succeeds with probability p^5, so only 1.69 transmissions are necessary on average. Host A must send an average of (1196 bytes x 2.09) + (884 bytes x 1.69) = 2499.64 + 1493.96 bytes = 3993.6 bytes.

d. As far as the receiver is concerned, this is a part of new datagram, since no other parts of it are known. It will therefore be queued until the rest show up. If they do not, this one will time out too.

4. Forwarding and Classless Interdomain Routing (CIDR)

Suppose a router has built up the routing table shown below. The first four lines are for CIDR addresses, with "/22" indicating a mask of 22 1's followed by 10 0's.

Net/Masklength	NextHop
128.174.248.0/22	R1
128.174.248.128/25	R2
128.174.252.0/22	R3
128.174.254.7	Interface1
128.174.254.32/27	Interface2
128.174.224.0/19	Interface3
Default	R4

a. How many individual IP addresses match each of the four Net/Masklength pairs?
b. The router can deliver packets directly over interfaces 0, 1 or 2, or it can forward to routers R1,R2, R3 or R4. Specify the next hop for each of the following destinations. Remember that if a destination matches more than one line of the table, the longest match is used.
   i. 128.174.255.7
   ii. 128.174.254.7
   iii. 128.174.253.7
   iv. 128.174.251.7
   v. 128.174.240.7
   vi. 128.174.220.7

a.

Net/Masklength	Number of Hosts

128.174.248.0/22	1023 – 2 (. . .255 & . . .0) – 126 (. . .128/25) = 895
128.174.248.128/25	127 – 1 (. . .255) = 126
128.174.252.0/22	1023 – 2 (. . .255 & . . .0) – 31 (. .254.32/27) = 990
128.174.254.7	1
128.174.254.32/27	31
128.174.224.0/19	8192 – 2 (. . .255 & . . .0) – 2043 (all the above) = 6147
Default	-

b.

i. 128.174.255.7 ⇨ R3
ii. 128.174.254.7 ⇨ I1
iii. 128.174.253.7 ⇨ R3
iv. 128.174.251.7 ⇨ R1
v. 128.174.240.7 ⇨ I3
vi. 128.174.220.7 ⇨ R4