CS 421: Computer Networks **FALL 2005**

MIDTERM II
December 7, 2005
120 minutes

Name: ______________________________

Student No: ___________________________

Show all your work very clearly. Partial credits will only be given if you carefully state your answer with a reasonable justification.

1)

a) (10 pts) Execute the Dijkstra algorithm at node A for the network shown below by filling in the following table. In the table, you need to give both the distance *D(v)* and previous node *p(v)*.

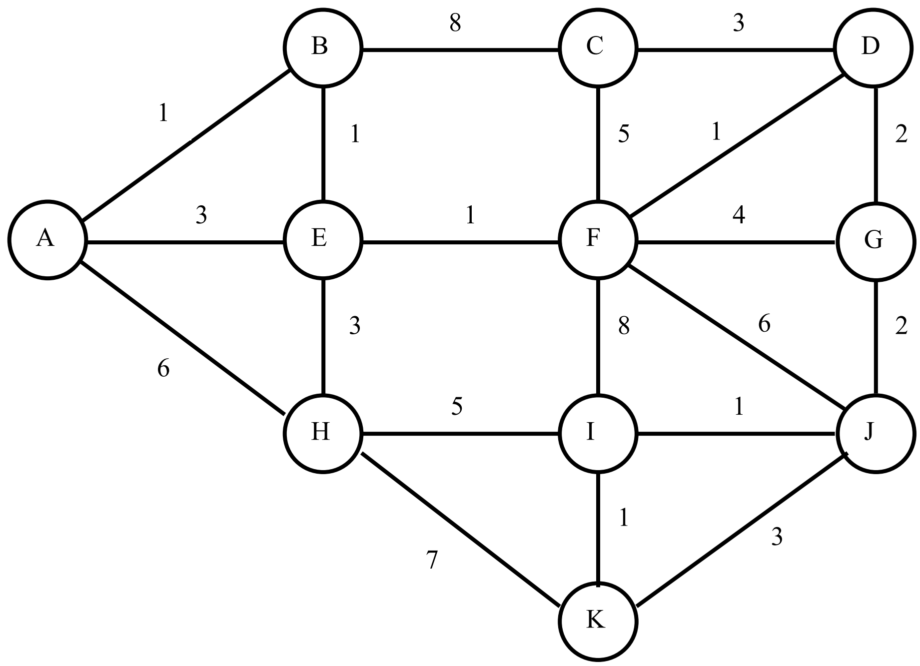

iter.	*N*	*D(B), p(B)*	*D(C), p(C)*	*D(D), p(D)*	*D(E), p(E)*	*D(F), p(F)*	*D(G), p(G)*	*D(H), p(H)*	*D(I), p(I)*	*D(J), p(J)*	*D(K), p(K)*

b) The network below uses the distance-vector routing algorithm. Assume the following:

- Links have the same cost in both directions.
- Nodes exchange their routing info once every second, in perfect synchrony and with negligible transmission delays. Specifically, at every t = i, i = 0, 1, 2, 3,..., each node sends and receives routing info instantaneously, and updates its routing table; the update is completed by time t=i+0.1.
- At time t = 0, the link costs are as shown below and the routing tables have been stabilized. At time t = 0.5, the cost of the link (1,3) becomes 11. There are no further changes in the link costs.
- Route advertisements are **only exchanged periodically**, i.e., there are no immediate route advertisements after a link cost change. Hence the first route advertisement after the link cost change at t = 0.5 sec occurs at t = 1.0 sec. *Note:* However, whenever a link cost change occurs, the

two nodes at the endpoints of this link immediately make corresponding changes in their distance tables.

i. (10 pts) Assume that the distance vector algorithm **does not use poisoned reverse**. Give the evolution of the distance tables with respect to destination 4. Specifically, give the distance table entries for destination 4 at nodes 1-3, for t = 0.1, 0.5, 1.1, 2.1, ..., **until** all distance vectors stabilize. Present your final answer in the table given below where $D^i(j)$ is the distance vector denoting the distance from i to j .

Time, t	$D^1(4)$ via		$D^2(4)$ via		$D^3(4)$ via		
	2	3	1	3	1	2	4
0.1							
0.5							
1.1							
2.1							
3.1							
4.1							
5.1							
6.1							
7.1							
8.1							
9.1							

ii. (10 pts) Redo part i. assuming that the distance vector algorithm **uses poisoned reverse**.

Time, t	$D^1(4)$ via		$D^2(4)$ via		$D^3(4)$ via		
	2	3	1	3	1	2	4
0.1							
0.5							
1.1							
2.1							
3.1							
4.1							
5.1							

iii. (6 pts) During the convergence of the distance vector algorithm, does it exist a routing loop regarding destination 4 for any t when poisoned reverse is not used, i.e., part i? Answer the same question for the case when poisoned reverse is used, i.e., part ii.

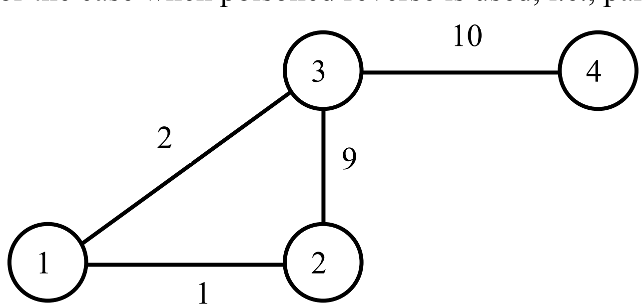

2)

a) (9 pts) Suppose host A transmits a 1796 byte IP packet (including the 20 Byte IP header) over a 2-hop path to host B. The MTU of the first link (A to router) is 900 bytes (IP header plus data), and the MTU of the second link (router to B) is 500 bytes (IP header plus data). Assuming that IP header does not contain any options, indicate the length (in bytes), more flag, and offset field values **(specify the offset values in units of 8 bytes)** of the fragment(s) transmitted over each link in the tables below.

First link

Fragment	Length	Offset	Flag
1			
2			
3			
4			
5			
6			

Second link

Fragment	Length	Offset	Flag
1			
2			
3			
4			
5			
6			

b) (10 pts) The following is a forwarding table at router R, which uses Classless Interdomain Routing (CIDR).

Destination Network	Next Hop
139.179.39.0 / 25	A
139.179.39.128 / 25	B
139.179.72.0 / 26	C
196.101.153.64 / 26	D
139.179.0.0 / 16	E
196.0.0.0 / 8	F

Suppose packets with the following destination IP addresses arrive at router R. Determine to what next hop each of these packets will be delivered (Give **only one** answer for each destination.)

(i) 139.179.39.179
(ii) 139.179.72.68
(iii) 196.101.153.136
(iv) 139.179.39.108
(v) 196.101.153.115

c) (10 pts) In the network shown below, ISP A serves eight customer subnets. Suppose ISP A wants to implement a policy that allows Internet traffic to subnets 198.35.80.0/23 through 198.35.86.0/23, i.e., the first 4 subnets, to be routed only through ISP C but not through ISP B and the traffic to rest of the subnets to be routed from either ISP B or ISP C. What are the route advertisements that ISP A should make to ISP B and ISP C in order to implement this policy? **In answering this question, use the smallest possible number of advertisements that are sent to ISP B and ISP C after route aggregation.**

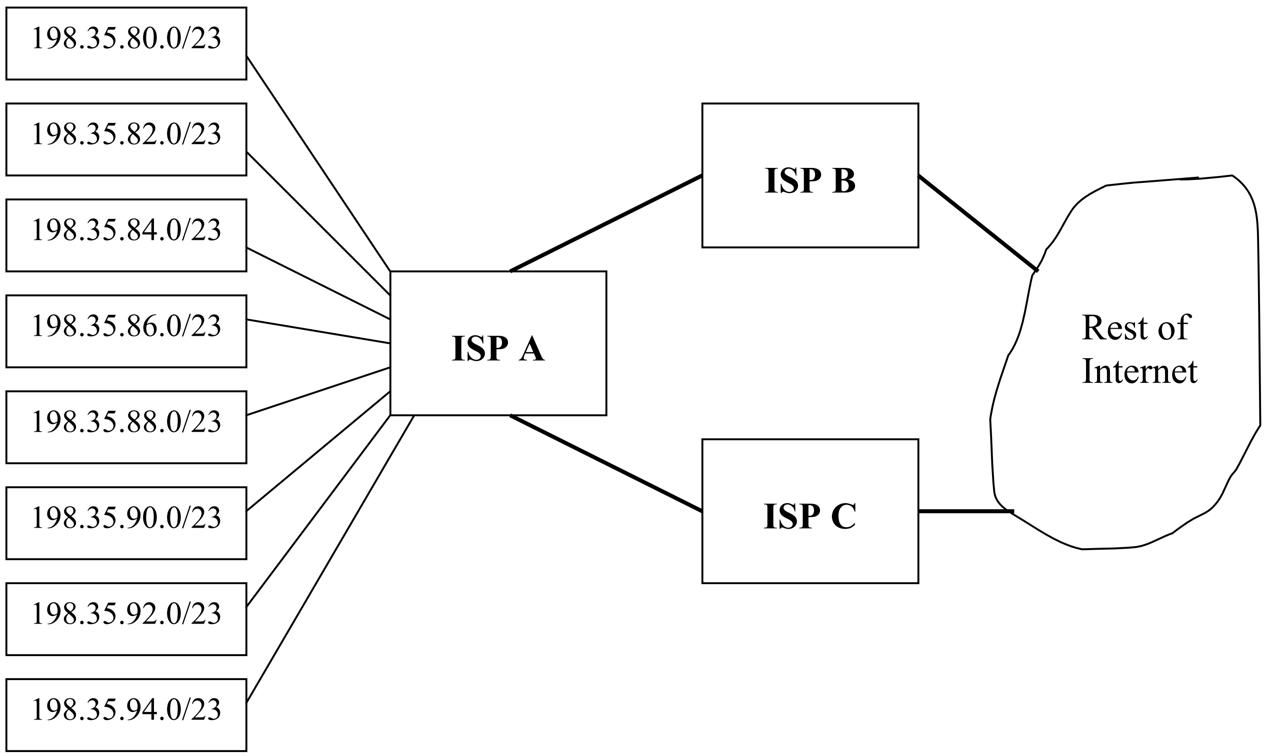

3)

a) (5 pts) What are the two most important reasons that the Internet is organized as a hierarchy of networks for routing purposes? **(Use at most two sentences.)**

b) (5 pts) Name two methods/algorithms/protocols that help in making more efficient use of IP addresses.

c) (5 pts) IPv4 has been designed to reassemble fragments only at the destination, but that is just one possibility. It would be possible to alternatively require IP to reassemble fragments at every hop. There could be reasons for choosing either approach. (a) Give a reason for reassembling fragments of an IP datagram at each router (hop). (b) Give a reason for reassembling fragments only at the destination. **(Use at most two sentences.)**

d) (5 pts) Give two reasons why TCP/IP performs error checking using checksums both at the transport and network layers?

e) (5 pts) Give three reasons why an IP router may discard an IP packet.

f) (5 pts) Which is easier to deploy in the Internet: a new Network Layer protocol or a new Transport layer protocol? Why?

g) (5 pts) State two modifications of IPv6, from the IPv4, that allow a router to process a packet faster.