CEN 4007C Computer Networks Fundamentals
Instructor: Prof. A. Helmy

Homework 4: Network Layer

Assigned: Nov. 28th, 2011. Due Date: Dec 8th, 2011 (to the TA)

1. (3 points) What are the 2 most important network-layer functions in a datagram network? What are the 3 most important network-layer functions in a VC network?

Datagram-based network layer: forwarding; routing. Additional function of VCbased network layer: call setup.

2. (4 points) What is the difference between routing and forwarding?

Forwarding is about moving a packet from a router's input link to the appropriate output link. Routing is about determining the end-to-routes between sources and destinations. Usually routing runs continuously to establish the routing tables before the packets are forwarded and in anticipation of packets.

3. (7 points) Describe some hypothetical services that the network layer can provide to a single packet. Do the same for a flow of packets. Are any of your hypothetical services provided by the Internet's network layer? Are any provided by ATM's CBR service model? Are any provided by ATM's ABR service model?

Single packet: guaranteed delivery; guaranteed delivery with guaranteed maximum delay. Flow of packets: in-order packet delivery; guaranteed minimal bandwidth; guaranteed maximum jitter. None of these services is provided by the Internet's network layer. ATM's CBR service provides both guaranteed delivery and timing. ABR does not provide any of these services.

4. (3 points) List some applications that would benefit from ATM's CBR service model.

Interactive live multimedia applications, such as IP telephony and video conference, could benefit from ATM CBR's service, which maintains timing.

5. (3 points) List and briefly describe three types of switching fabrics.

Switching via memory; switching via a bus; switching via an interconnection Network

6. (4 points) What is HOL blocking? Does it occur in input ports or output ports?

HOL blocking – a queued packet in an input queue must wait for transfer through the fabric because it is blocked by another packet at the head of the line. It occurs at the input port.

7. (3 points) An application generates chunks of 40 bytes of data every 20msec, and each chunk gets encapsulated in a TCP segment and then an IP datagram.
What percentage of each datagram will be overhead, and what percentage will be application data?

TCP header=20 bytes, IP header=20 bytes, so there's 50% overhead

8. (4 points) Suppose you purchase a wireless router and connect it to your cable modem. Also suppose that your ISP dynamically assigns your connected device (that is, your wireless router) one IP address. Also suppose that you have five PCs at home that use 802.11 to wirelessly connect to your wireless router. How are IP addresses assigned to the five PCs? Does the wireless router use NAT? Why or why not?

Typically the wireless router includes a DHCP server. DHCP is used to assign IP addresses to the 5 PCs and to the router interface. Yes, the wireless router also uses NAT as it obtains only one IP address from the ISP.

9. (6 points) Compare and contrast link-state and distance-vector routing algorithms.

Link state algorithms (used in OSPF): Computes the least-cost path between source and destination using complete, global knowledge about the network. Distance-vector routing (used in RIP): The calculation of the least-cost path is carried out in an iterative, distributed manner. A node only knows the neighbor to which it should forward a packet in order to reach a given destination along the least-cost path, and the cost of that path from itself to the destination.

10. (5 points) Discuss how a hierarchical organization of the Internet has made it possible to scale to millions of users.

Routers are aggregated into autonomous systems (ASs). Within an AS, all routers run the same intra-AS routing protocol. Special gateway routers in the various ASs run the inter-autonomous system routing protocol that determines the routing paths among the ASs. The problem of scale is solved since an intra-AS router need only know about routers within its AS and the gateway router(s) in its AS.
It also allows prefix aggregation (per AS) to be able to represent a large number of routes with one entry in the routing tables.

11. (2 points) Is it necessary that every autonomous system use the same intra-AS routing algorithm? Why or why not?

No. Each AS has administrative autonomy for routing within an AS.

12. (6 points) Compare and contrast the advertisements used by RIP and OSPF.

With OSPF, a router periodically broadcasts routing information to all other routers in the AS, not just to its neighboring routers. This routing information sent by a router has one entry for each of the router's neighbors; the entry gives the distance from the router to the neighbor. A RIP advertisement sent by a router contains information about all the networks in the AS, although this information is only sent to its neighboring routers.

13. (3 points) Fill in the blank: "RIP advertisements announce the number of hops to various destinations. BGP updates, on the other hand, announce the _____________ to the various destinations".

"sequence of ASs on the routes"

14. (2 points) When a host joins a multicast group, must it change its IP address to that of the multicast group it is joining?

No. A host can join multiple multicast groups at once, and still have its own IP address used for unicast.

15. (5 points) What is the difference between a group-shared tree and a source-based tree in the context of multicast routing?

In a group-shared tree, all senders send their multicast traffic using the same routing tree. With source-based tree, the multicast datagrams from a given source are routed over s specific routing tree constructed for that source; thus each source may have a different source-based tree and a router may have to keep track of several source-based trees for a given multicast group.

16. (3 points) A bare-bones forwarding table in a VC network has 4 columns. What is the meaning of the values in each of these columns? A bare-bones forwarding table in a datagram network has 2 columns. What is the meaning of the values in each of these columns?

For a VC forwarding table, the columns are : Incoming Interface, Incoming VC Number, Outgoing Interface, Outgoing VC Number. For a datagram forwarding table, the columns are: Destination Address, Outgoing Interface.

17. (5 points) Consider the network setup in Fig 4.22 (the figure below). Suppose that the ISP instead assigns the router the address 126.13.89.67 and that the network address of the home network is 192.168/16.

a. (2 points) Assign addresses to all interfaces in the home network.
b. (3 points) Suppose each host has two ongoing TCP connections, all to port 80 at host 128.119.40.86. Provide the six corresponding entries in the NAT translation table.

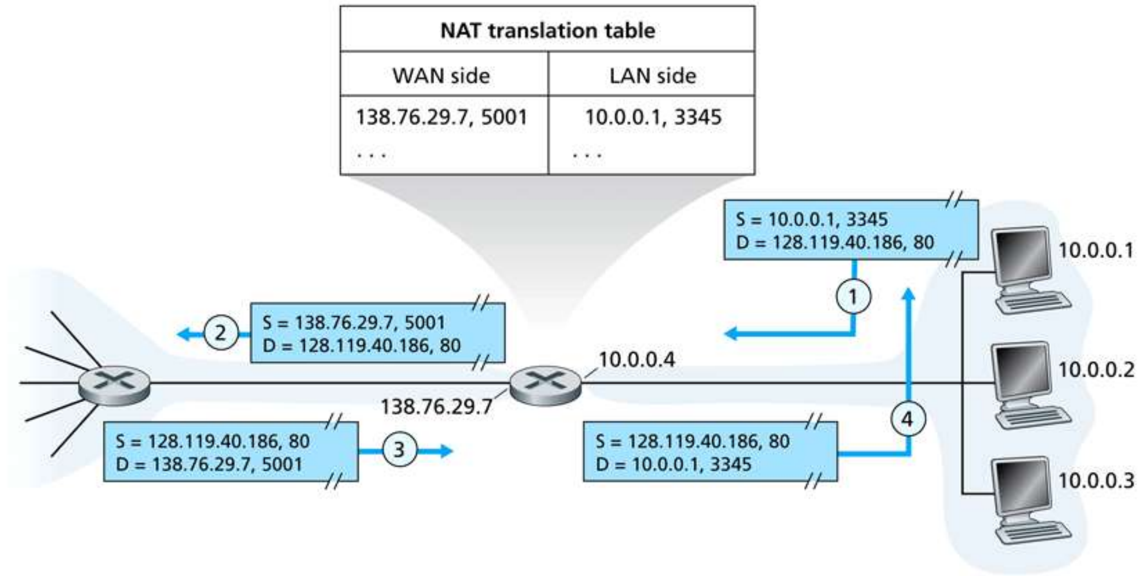

Figure 4.22 ♦ Network address translation

a) Home addresses: 192.168.0.1, 192.168.0.2, 192.168.0.3 with the router interface being 192.168.0.4

b)

NAT Translation Table

WAN Side	LAN Side
126.13.89.67, 4000	192.168.0.1, 3345
126.13.89.67, 4001	192.168.0.1, 3346
126.13.89.67, 4002	192.168.0.2, 3445
126.13.89.67, 4003	192.168.0.2, 3446
126.13.89.67, 4004	192.168.0.3, 3545
126.13.89.67, 4005	192.168.0.3, 3546

18. (3 points: Extra) In this problem we'll explore the impact of NATs on P2P applications:
Suppose a peer A discovers through query that peer B has a file it wants to download. Also, suppose that A and B are both behind a NAT. Is there a way to allow A to establish a TCP connection with B without application-specific NAT configuration? Discuss why.

It is not possible to devise such a technique. In order to establish a direct TCP connection between Arnold and Bernard, either Arnold or Bob must initiate a connection to the other. But the NATs covering Arnold and Bob drop SYN packets arriving from the WAN side. Thus neither Arnold nor Bob can initiate a TCP connection to the other if they are both behind NATs.

19. a, b, c (9 points: Extra)
Consider the network shown in the figure below. Using Dijkstra's algorithm, and showing your work using a table for each of a, b and c, do the following:

a. Computer the shortest path from *s* to all network nodes.

b. Computer the shortest path from *t* to all network nodes.
c. Computer the shortest path from *u* to all network nodes.

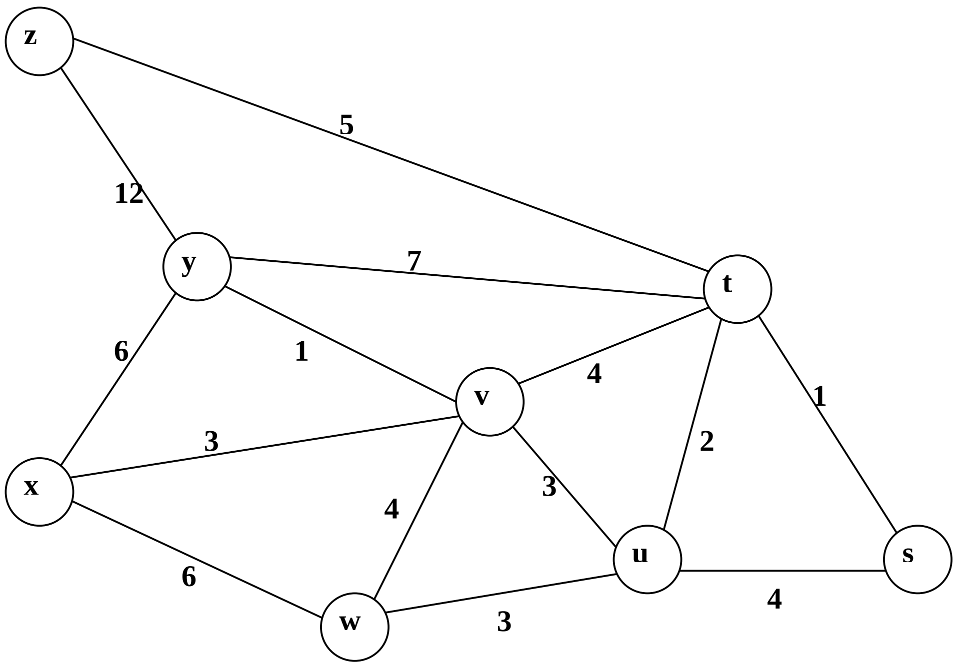

a.

Step	*N'*	*D(x), p(x)*	*D(t),p(t)*	*D(u),p(u)*	*D(v),p(v)*	*D(w),p(w)*	*D(y),p(y)*	*D(z),p(z)*
0	s	∞	1,s	4,s	∞	∞	∞	∞
1	st	∞	1,s	3,t	5,t	∞	8,t	6,t
2	stu	∞	1,s	3,t	5,t	6,u	8,t	6,t
3	stuv	8,v	1,s	3,t	5,t	6,u	6,v	6,t
4	stuvy	8,v	1,s	3,t	5,t	6,u	6,v	6,t
5	stuvyz	8,v	1,s	3,t	5,t	6,u	6,v	6,t
6	stuvyzw	8,v	1,s	3,t	5,t	6,u	6,v	6,t
7	stuvyzwx	8,v	1,s	3,t	5,t	6,u	6,v	6,t

b.

Step	*N'*	*D(x), p(x)*	*D(s),p(s)*	*D(u),p(u)*	*D(v),p(v)*	*D(w),p(w)*	*D(y),p(y)*	*D(z),p(z)*
0	t	∞	1,t	2,t	4,t	∞	7,t	5,t
1	ts	∞	1,t	2,t	4,t	∞	7,t	5,t
2	tsu	∞	1,t	2,t	4,t	5,u	7,t	5,t
3	tsuv	7,v	1,t	2,t	4,t	5,u	5,v	5,t
4	tsuvw	7,v	1,t	2,t	4,t	5,u	5,v	5,t
5	tsuvwy	7,v	1,t	2,t	4,t	5,u	5,v	5,t
6	tsuvwyz	7,v	1,t	2,t	4,t	5,u	5,v	5,t
7	tsuvwyzx	7,v	1,t	2,t	4,t	5,u	5,v	5,t