C. Noronha
Winter 98-99

EE384A: Network Protocols and Standards
Midterm Solutions
Closed Book / Closed Notes
Thursday, February 11, 1999
8:00 AM – 9:15 AM

NAME: _____

Honor code observed: _____
(signature)

Assigned Grade:

1) _____ **/ 8**
2) _____ **/ 15**
3) _____ **/ 6**
4) _____ **/ 11**
5) _____ **/ 10**

Total _____ **/ 50**

For questions 1, 2 and 3, consider the topology depicted in Figure 1. Company X has three locations, each one with a LAN (A, B, and C). The LANs in these three locations need to be interconnected; the company then deploys three Gateway devices, interconnecting the LANs as shown.

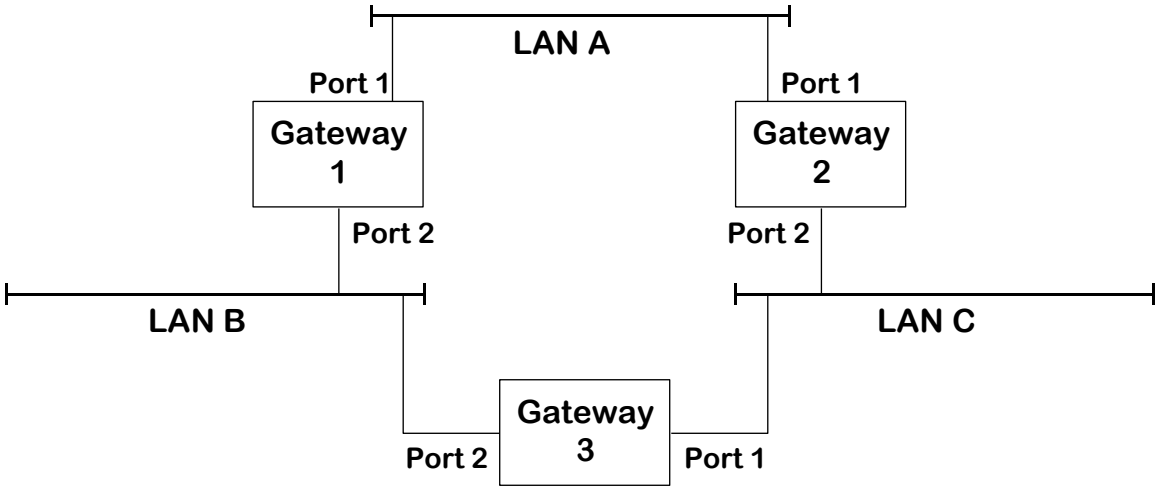

Figure 1: Company X Network Topology

Problem 1: Transparent Bridging (8 points)

For this problem, assume that Gateways 1, 2, and 3 are transparent bridges, and their ID numbers are the ones shown (i.e., 1 for Gateway 1, 2 for Gateway 2, etc).

a. Assume that the spanning three algorithm has reached steady state. In the table below, indicate the state of each port in the network. Use ***DP*** for Designated Port, ***RP*** for Root Port, and ***B*** for Blocking. (3 points)

Bridge	**Port**	**State**
Gateway 1	Port 1	***DP***
	Port 2	***DP***
Gateway 2	Port 1	***RP***
	Port 2	***DP***
Gateway 3	Port 1	***B***
	Port 2	***RP***

b. Assume that each LAN segment is a shared-medium LAN, with a perfect contention mechanism, i.e., the entire capacity of the LAN segment may be used. The capacity of each LAN segment is C bits/sec. The stations connected to each segment generate T bits/sec of traffic (i.e., the traffic that originates at ***each*** segment is T bits/sec). Assume that the traffic on each segment is uniformly addressed, i.e., out of the T bits/sec offered, $T/3$ bits/sec are addressed to stations in the same segment, and $T/3$ bits/sec are addressed to stations on each of the other two segments. Derive the maximum value of T as a function of C under these conditions. (5 points)

The final spanning tree for this network is shown below, with the traffic flows:

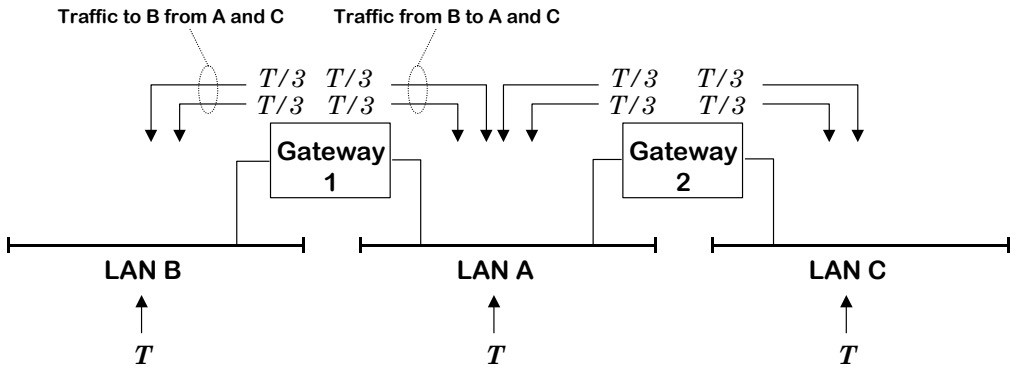

The traffic flow equation for segments B and C is:

$$T + 2T/3 \leq C, \text{ or } T \leq 3C/5$$

The traffic flow equation for segment A is:

$$T + 2T/3 + 2T/3 \leq C, \text{ or } T \leq 3C/7$$

Since both equations must be satisfied simultaneously, the overall traffic is limited by segment A (which intuitively must be the case, since it is being used to relay the traffic from B to C). Therefore, the final solution is $T \leq 3C/7$.

Problem 2: IP Routing (15 points)

For this problem, assume that Gateways 1, 2, and 3 are IP Routers. The network is connected to the Internet through a third port on Gateway 3, as shown in Figure 2 below.

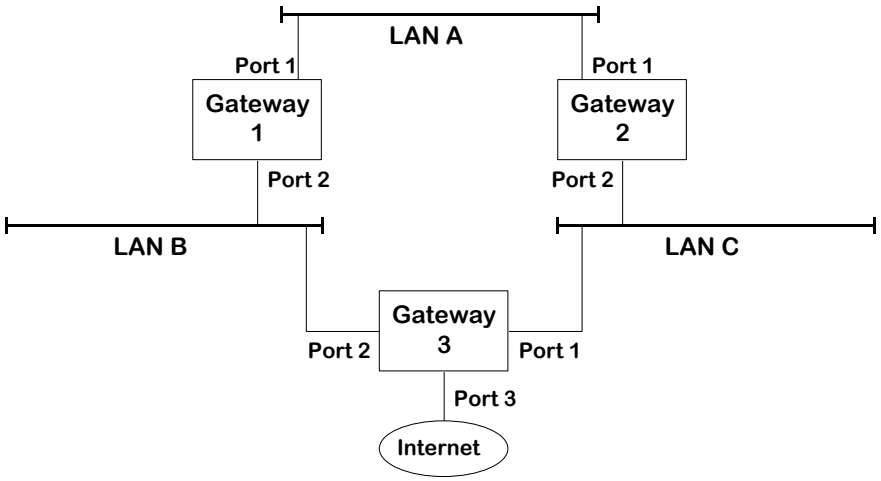

Figure 2: Connecting Company X to the Internet

a. Assume that the class-C network number 204.240.18.0 has been assigned to the company. Assume that there are 100 hosts on LAN A, 50 hosts on LAN B, and 50 hosts on LAN C. The organization uses subnetting, and each LAN should be a distinct subnet. It is required that all machines and all router interfaces be given IP addresses. You are asked to assign IP addresses to the stations and router interfaces, and to indicate the entries in each router's table needed to route IP datagrams destined to stations within the company network. The IP address of Port 3 of Gateway 3 is assigned by the organization's ISP on its backbone range; the assigned value is 38.100.20.78 (and it has no bearing to this problem). Write your answers in the tables provided.

Subnetworks: (2 points) Note: IP Address Range refers to the assigned range of IP addresses for each of the three subnets.

Since there are only three networks, we just need three subnets. Note that the first and last addresses of each subnet are reserved; this applies to all subnets. The total number of addresses shown below is the size of the block assigned to the subnet; the actual number of available addresses is the number shown minus two.

	IP Address Range		**Subnet Mask**	**Total # of Addresses**
Subnet	**From IP**	**To IP**		
LAN A	*204.240.18.128*	*204.240.18.255*	*FF.FF.FF.80*	*128*
LAN B	*204.240.18.0*	*204.240.18.63*	*FF.FF.FF.C0*	*64*
LAN C	*204.240.18.64*	*204.240.18.127*	*FF.FF.FF.C0*	*64*

Assigned Router Interface IP Addresses: (1 point)
(Remember that the first and last addresses are reserved; we use the next two addresses for the gateway.)

Router	**Port**	**IP Address**
Gateway 1	Port 1	*204.240.18.129*
	Port 2	*204.240.18.1*
Gateway 2	Port 1	*204.240.18.130*
	Port 2	*204.240.18.65*
Gateway 3	Port 1	*204.240.18.66*
	Port 2	*204.240.18.2*
	Port 3	38.100.20.78

Assigned Host IP Addresses: (1 point)

Subnet	**From IP**	**To IP**
LAN A	*204.240.18.131*	*204.240.18.230*
LAN B	*204.240.18.3*	*204.240.18.52*
LAN C	*204.240.18.67*	*204.240.18.116*

Routing Tables (the supplied tables may have more rows than required; you do not necessarily need to use all the rows). For Output Port, indicate Port 1 or Port 2 (or Port 3, in case of Gateway 3). For next Hop IP, use the word "Direct" when the packet is destined to a directly attached network. (6 pts)

*(Since the problem indicated that only the entries for routing packets inside the company network are needed, you could skip the default entry. However, we present it here for completeness. Note that, for routes marked with an *, there are two equivalent shortest paths, and one is chosen at random. Finally, note that the entries are sorted by the length of the mask – from the more specific to the less specific.)*

Routing Table for Gateway 1			
Destination IP	Mask	Output Port	Next Hop IP
204.240.18.0	*FF.FF.FF.C0*	*Port 2*	*Direct*
204.240.18.64	*FF.FF.FFC0*	*Port 1**	*204.240.18.130*
204.240.18.128	*FF.FF.FF.80*	*Port 1*	*Direct*
0.0.0.0	*0.0.0.0*	*Port 2*	*204.240.18.2*

Routing Table for Gateway 2			
Destination IP	Mask	Output Port	Next Hop IP
204.240.18.64	*FF.FF.FF.C0*	*Port 2*	*Direct*
204.240.18.0	*FF.FF.FF.C0*	*Port 1**	*204.240.18.129*
204.240.18.128	*FF.FF.FF.80*	*Port 1*	*Direct*
0.0.0.0	*0.0.0.0*	*Port 2*	*204.240.18.66*

For this table, use the word "ISP" in the Next Hop IP column when the output port is Port 3.

Routing Table for Gateway 3			
Destination IP	Mask	Output Port	Next Hop IP
204.240.18.0	*FF.FF.FF.C0*	*Port 1*	*Direct*
204.240.18.64	*FF.FF.FF.C0*	*Port 2*	*Direct*
204.240.18.128	*FF.FF.FF.80*	*Port 1**	*204.240.18.65*
0.0.0.0	*0.0.0.0*	*Port 3*	*ISP*

b. Similarly to question 1.b, assume that each LAN segment is a shared-medium LAN, with a perfect contention mechanism. The capacity of each LAN segment is C bits/sec. The stations connected to each segment generate T bits/sec of traffic (i.e., the traffic that originates at ***each*** segment is T bits/sec). Assume that the traffic on each segment is uniformly addressed, i.e., out of the T bits/sec offered, $T/3$ bits/sec are addressed to stations in the same segment, and $T/3$ bits/sec are addressed to stations on each of the other two segments. For the purposes of this question, there is no traffic to or from the Internet. Derive the maximum value of T as a function of C under these conditions. (5 points)

In this case, the traffic from a segment to each one of its two neighbors travels one hop, through the directly attached router. The traffic flow is symmetric (i.e., the same in all segments) and is illustrated below:

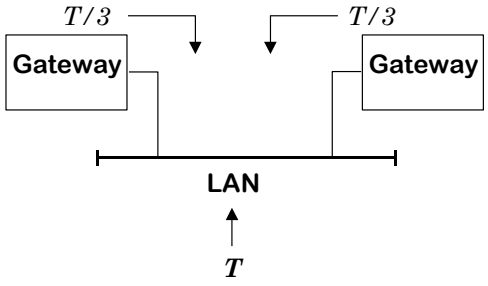

From the figure, it is clear that the traffic flow equation is:

$$T + 2T/3 \leq C, \text{ or } T \leq 3C/5$$

Which is the final solution.

Problem 3: Comparison between Bridges and Routers (6 points)

Based on your answers to Problems 1 and 2 and your general knowledge of how transparent bridges and routers work, compare the two approaches to interconnecting the three locations of Company X. You should be able to identify at least one advantage and one disadvantage for each approach.

Bridged Solution	*Routed Solution*
Advantages: – *Protocol-Independent (works for any protocol)* – *Hosts can be freely moved between segments without the need for reconfiguration* *Disadvantages:* – *As shown in 1.b and 2.b, the overall network throughput is lower by a ratio of 5/7 because one of the links is not used.* – *Broadcast traffic gets sent to every segment, further decreasing throughput.*	*Advantages:* – *Better throughput (as shown in 1.b and 2.b) because all the links are used simultaneously and the traffic goes through the shortest path.* – *Broadcast traffic is limited to each LAN.* *Disadvantages:* – *Works only for IP; other protocols are not transported.* – *Hosts need to be reconfigured (get a new IP address) if they are moved from LAN to LAN.*

Problem 4: GARP/GMRP/GVRP (11 points)

a. GARP is said to be a scalable protocol (i.e., it can scale well to large numbers of devices). What makes it scalable? What prevents the number of protocol messages from growing with the number of devices? (2 pts)

GARP is scalable because the number of protocol messages does not grow with the number of devices. This is achieved by the following:

- *Messages are multicast and are received by all the participants.*
- *Participants will track the state of the registrations, and will not emit any messages if they wish to declare an attribute they know it is already registered.*
- *When a device wants to send a response, it will wait for a random interval. If another device responds within that time period, the device suppresses its response.*
- *When a device wishes to withdraw a declaration, it will do so silently if it knows that other devices wish to keep it.*

b. The table below shows all the possible states of the GARP Applicant (this is table 12-2 from IEEE P802.1D/D15). Note that there is no Leaving Passive Member State. Why is that? What happens when a Quiet Passive Member (QP) decides to withdraw a declaration? Does it send any messages? If yes, which one(s)? (3 pts)

	Very Anxious	Anxious	Quiet	Leaving
Active Member	VA	AA	QA	LA
Passive Member	VP	AP	QP	--
Observer	VO	AO	QO	LO

A Passive Member is a device that wishes to declare an attribute, but has not sent any messages to that effect (because some other member did it first). Therefore, it knows that there is at least one more device in the same segment that has declared that attribute. If a Quiet Passive Member wishes to withdraw a declaration, all it needs to do is to transition to the Quiet Observer (QO) state, without sending any messages. This is to save on protocol exchanges and improve scalability; if the device were to send a Leave message, it would force the other devices interested in this attribute (which are known to exist) to re-declare it. In other words, it suppresses the Leave message because it knows that there are other devices that wish to keep the attribute declared.

c. What is the purpose of GMRP? Why can't GMRP-unaware bridges filter the multicast traffic in the same (transparent, learning) fashion as with unicast traffic? (2 pts)

The main purpose of GMRP is to convey multicast group membership information from the hosts to the bridges, so that the bridges can suppress the transmission of multicast packets on segments where no host is interested in them. They cannot be filtered in the same fashion as the unicasts because the multicast packet does not have any indication of who will receive it; the only information the bridge learns from the packet is the location of the source, which may not even be interested in the group. In other words, there is no way to determine which hosts are listening to the group by just monitoring the traffic.

d. What is the purpose of a VLAN? (Why is a VLAN useful?) (2 pts)

VLANs are used to segregate users in groups whose members "appear" to be in the same LAN. With VLANs, group membership is independent of the actual physical location of the host; it can be connected anywhere in the network, and it will still be in the same group; this is useful when group members move from location to location. Segregation of users may be useful from a traffic and security point of view (only the members of the group may access a certain resource); VLANs make this easier to realize, and independent of the network topology.

e. VLANs can be created manually through static management (i.e., the network administrator configures every bridge). What is the use of the GVRP protocol, which enables dynamic VLAN Registration? Give an example. (2 pts)

With GVRP, the registration is now dynamic. When a host is connected to the network, it can declare its VLAN, and the bridges will automatically start forwarding traffic to its location. One example: a company may have its engineering offices in one wing of a building, and the labs in another; an engineer can freely move with his laptop between the lab and the offices and still be in the engineering LAN.

Problem 5: ICMP and RIP (10 points)

a. ICMP messages are normally generated to indicate an error condition on a received packet. However, ICMP messages are not generated about errors on ICMP packets or on fragments other than the first fragment of a PDU. Why? (2 pts)

This is done in order to cut down on ICMP traffic. If you start generating error messages (ICMP messages) about error messages, you can quickly go into a loop where congestion causes error messages to be generated, and the excess of error messages keeps adding to the congestion. As for the fragmented packet, if there is a problem with any part of the packet, the IP layer will discard it all; therefore, a single ICMP message should be generated.

b. Explain split horizon routing. Indicate why something like that is needed, and what would happen if the routers did not do it. (3 pts)

In RIP, routers build their routing tables from advertisements from other routers, and advertise their own routes. Split horizon means that, if router A has a path to some destination through router B, router A will not advertise this path to router B. This is needed because, if B looses its path to the destination, it will start believing that it can now be reached through A, creating a routing loop. This routing loop will eventually be resolved as advertisements get exchanged, but it will take several message exchanges. Split horizon prevents this situation from happening, improving the error recovery.

c. What are triggered updates in RIP? What is the benefit of using them? (2 pts)

In RIP, routers periodically advertise the paths in their routing table. A triggered update happens when there is a change in the paths advertised. Instead of waiting for the end of the current period, the router sends the update immediately (i.e., the change "triggers the update"). This is done to speed up route convergence; even with split horizon, routing loops (involving three or more routers) are still possible.

d. Routing loops are always a problem in most environments. However, they are a much more severe problem in bridged networks than in IP routed networks. Explain why. (3 pts)

Routing loops are always a problem because, as the packet goes in its merry way around the loop, it is wasting network resources and not going anywhere. However, in bridged networks, the packet is forwarded "as is"; there is no change in the packet contents. Therefore, the n^{th} transmission of the packet looks just like the first; there is no intrinsic mechanism to break the loop. In IP routed networks, the packet cannot loop forever because there is a field in the IP header, called the Time To Live (TTL), which contains how many more hops the packet can go through. At every hop, this field is decremented, and when it gets to zero, the packet is discarded. This way, the packet does not loop forever.