

Midterm 1 Solutions

(50 points)

Wednesday March 1, 2000

Professor Sengupta

Problem 1. Multiple choice questions. Circle the answer that is the most correct.

1.1 (1 point) If you were building a network for sources that generate point-to-point traffic at a steady rate of 64 kbps you would build a

- 1. a circuit switched network**
2. a packet switched network
3. a broadcast network
4. none of the above

1.2 (1 point) To maximize the pipelining gain in a packet switched network the ideal packet size

1. is as large as the data chunks generated by the applications
- 2. depends on the size of the protocol headers**
3. is as small as one bit ideally
4. is none of the above

1.3 (1 point) Statistical multiplexing gain is observed when

- 1. the users of a network transmit only a small fraction of the time they are connected**
2. the users of a network transmit all the time they are connected at their peak rate
3. when the users of the network send e-mail
4. none of the above

1.4 (1 point) When two packets from two hosts arrive simultaneously at an ethernet hub

1. both packets are recieved correctly
2. one packet is recieved correctly
3. the outcome is exactly like a switched ethernet
- 4. none of the above**

1.5 (1 point) If a system administrator replaces an ethernet hub with an ethernet switch

1. all of the hosts on the ethernet must replace their NIC's
2. some of the hosts on the ethernet must replace their NIC's
- 3. none of the hosts on the ethernet must replace their NIC's**
4. none of the above

1.6 (1 point) The best local area network for hosts that need to send traffic at regular intervals of time is

1. an ethernet
- 2. a token ring**
3. slotted aloha
4. none of the above

1.7 (1 point) An FDDI ring guarantees that each host will receive the token within

1. **2TTRT of releasing the token**
2. 0.5TTRT of releasing the token
3. no defined time limit
4. none of the above

1.8 (1 point) Ethernets are popular with system administrators because

1. **they are simple to install, easy to maintain, and provide satisfactory delay and throughput to hosts**
2. they provide deterministic access latencies to hosts
3. they provide special provisions for synchronous and asynchronous traffic
4. none of the above

1.9 (1 point) The information required by an IP protocol when handling an IP packet is contained in the

1. ethernet header
2. **IP header**
3. a separate ICMP message
4. none of the above

1.10 (1 point) When a DNS lookup is forwarded to the domain name server of the destination the reply is cached by the

1. root servers and the DNS of the source host
2. **DNS of the source host only**
3. all name servers involved in the lookup
4. none of the above

1.11 (1 point) The purpose of validation tags in http is to

1. **avoid fetching the content from the origin server unless it has changed since it was fetched by the proxy server**
2. avoid contacting the origin server all together
3. check if the data has expired
4. none of the above

1.12 (1 point) The HTTP header segment,

*Accept: test/plain; q=0.5, text/html
test/x-dvi; q=0.8, text/x-c*

is for

1. HTTP cache control
2. **HTTP content negotiation**
3. HTTP connection management
4. none of the above

1.13 (1 point) An ftp client on host A can control a data transfer between ftp servers on two host B and C remote from A because

1. the ftp control messages and data messages always flow through the same TCP connection
2. the ftp control connection runs on top of telnet
3. **the ftp control connection and data connections can be separated**
4. none of the above

1.14 (1 point) Exterior routing protocols such as BGP are designed primarily

1. to make link metrics consistent across all autonomous systems
2. to run the Dijkstra algorithm across the internet

3. to enable autonomous systems to cooperate while maintaining independent control of their routing policies

4. none of the above

1.15 (1 point) Each layer in the Internet protocol stack provides a service that is built on the

1. service provided by the layer immediately below it in the protocol stack

2. service provided by the layer immediately above it in the protocol stack

3. bit transport service provided by the physical layer

4. none of the above

1.16 (5 points) In figure 1 host N1.1 is sending a packet to host N2.3. Think of N1.1 as the network address of the host with N1 as the network part and 1 as the host part of the address. There should be four addresses in the packet header. Think of these addresses as a sequence s of four addresses in the order

s = source-ethernet-address, destination-ethernet-address, source-ip-address, destination-ip-address.

Let s(N1.1) denote the address sequence in the packet sent by the source host and s(Router) denote the address sequence in the corresponding packet forwarded by the router. The ethernet address of a NIC is denoted by m followed by an integer. For example m3 is the ethernet address of the NIC on port 1 of the router. Then the address sequences in the packet sent by N1.1 and that forwarded by the router are respectively:

1. s(N1.1) = m1,m3,N1.1,N2.3 and s(Router) = m4,m7,N1.1,N2.3

2. s(N1.1) = m1,m3,N1.1,N1.3 and s(Router) = m4,m7,N2.5,N2.3

3. s(N1.1) = m1,m7,N1.1,N2.3 and s(Router) = m1,m7,N1.1,N2.3

4. s(N1.1) = m1,m7,N1.1,N1.3 and s(Router) = m1,m7,N2.5,N2.3

Problem 2. (10 points)

An organization has a class C network 200.1.1 and wants to form subnets for four departments with hosts as follows:

A	72 hosts
B	35 hosts
C	20 hosts
D	18 hosts

There are 145 hosts in all.

1. (7 points) Give a possible arrangements of the subnet masks to make this possible.

2. (3 points) Suggest what the organization might do if the department D grows to 34 hosts.

Solution 2.

Department	Mask	IP Address
A	255.255.255.128	200.1.1.128
B	255.255.255.192	200.1.1.64
C	255.255.255.224	200.1.1.32
D	255.255.255.224	200.1.1.0

If department D grows to 34 hosts we could split A into 3 subnets with 32 hosts each as follows and give the extra subnet to department D.

Department	Mask	IP Address
A1	255.255.255.224	200.1.1.224
A2	255.255.255.224	200.1.1.192
A3	255.255.255.224	200.1.1.160
B	255.255.255.192	200.1.1.64
C	255.255.255.224	200.1.1.32
D1	255.255.255.224	200.1.1.0
D2	255.255.255.224	200.1.1.128

Problem 3. (20 points)

Consider the network in figure 2. the link speeds are given in Megabits bet second.

1. (5 points) Assume that the cost of each link is 1/link speed. Not that this means the link cost does not change even if the traffic in the network changes. Assume that the distance vector algorithm is used to find the shortest paths in this network. Assume a synchronous execution with all the routers re-computing distance vectors every 1 second. Fill out the initial distance vectors (at time T = 0sec) at all routers using table 1. For example the entry in row A and column C should be the initial distance estimate of A to C. The first row is filled out for your convenience. Fill out the distance vectors at T = 1sec in table 1. How many seconds would it take for the distance vector algorithm to converge? Would the final distance vectors have been different if the distance vector algorithm was run asynchronously?

It takes the distance vector algorithms 2 seconds to converge.

The final distance vectors would not be different in an asynchronous execution.

T = 0	A	B	C	D
A	0	2/3	2	inf
B	2/3	0	2/3	2/3
C	2	2/3	0	2
D	inf	2/3	2	0
T = 1	A	B	C	D
A	0	2/3	4/3	4/3
B	2/3	0	2/3	2/3
C	4/3	2/3	0	4/3
D	4/3	2/3	4/3	0

Table 1: Tables for T = 0 and T = 1.

2. (3 points) Fill out the next hop for each destination in table 2 for the final routing table of A.

Destination Name	Next Node
B	B
C	B
D	B

Table 2: Final routing table at node A.

3. (3 points) Note that the router A has two buffers (one for each outgoing link) as shown in figure 2. If the routing table of A is as in the previous part then describe the number of bits in the output buffers A1 and A2

as functions of time. Think of the buffers for the time being as very large. Assume that all the network traffic flows from router A to D at a steady rate of 2 Mbps. Explain using your equations why one of the output buffers must eventually overflow.

A1: $2t - 1.5t = 0.5t$ Mbits, because all the traffic flows along the ABD and the routing is static.

A2: 0 Mbits

A1 will eventually overflow because the bits in A1 grows linearly with time. Thus for any finite buffer size there will come a time when the buffer is full.

4. (5 points) In this part we assume that the link between B and C is absent. We also assume that the link cost is the number of bits in the output buffer at the head of the link, i.e., the cost of the link (A, B) at time t is the number of bits in the buffer A1 at time t . Note that unlike the first three parts of this question the link cost is now dynamic. It depends on the queue in the buffer. All traffic is flowing from router A to D at 2 Mbps. We also assume that at $T = 0$ all the buffers at all the routers are empty and that A forwards all traffic for D along the faster path. Graph the bits in output buffers A1 and A2 as functions of time in figure 3 assuming that every one second A re-routes its traffic along the shorter of the two paths. Assume that if the two paths have equal cost then A stays with the path in the previous epoch. For your convenience the buffer occupancies is plotted for the first epoch.

5. (2 points) Can you recommend a buffer size in bits for the buffers A1 and A2?

A1: No less than 1.5 Mbits from figure 3.

A2: No less than 2.0 Mbits from figure 3

6. (2 points) Would you raise or lower your size recommendation if A re-routed every 0.5 second?

I would lower my recommended buffer size because if the routing switches faster the buildup of bits in each buffer will be less.