Your Name ___________________________________

Sample Mid-term Exam #2
CptS/EE 455
November 4, 2016

Write your name on your paper first!

You may have one 8 1/2 x 11 sheet of notes (both sides, typewritten or handwritten). Otherwise, this is a **closed book, closed notes, closed neighbor** exam. Calculators are allowed. There are N numbered questions on M pages (K sheets front and back). Make sure that you have them all. Answer the questions in the space provided. This is a **50 minute** exam, 10:10 – 11:00. Exams will be collected promptly at 11:00.

For best chance at partial credit, show all calculations.

1. IP Network addressing

If a host's IPv4 address and netmask are given as 13.2.0.64 and 255.255.252.0:

a) (5 pts) How many bits make up the subnet part of the address? Explain.

22 bits: 8+8+6

b) (5 pts) How many bits make up the host part of the address? Explain.

10: whatever is not subnet is host (so 32-22); or 10: 8 bits corresponding to the low-order byte of 0s and 2 more corresponding to the zeroes in the low order bits of the next byte.

c) (5 pts) How many usable host addresses are there on this subnet? Explain.

Since there are 10 bits of host address with two reserved addresses (all zero and all 1 host part), 2^10 - 2

d) (5 pts) is 13.2.3.137 a valid address on this same network? Why or why not?
The question is does the bit pattern agree on the subnet part. Yes it does.

(additional network addressing topics: CIDR notation; NAT; IPv6 addresses)

2. DHCP is used to automatically configure networking on hosts.
a) (5 pts) In order to configure a host to operate on a network what 2 values must DHCP supply at a minimum?
The host address and the network mask

b) (5 pts) What are two additional values that DHCP could (and often does) supply to host that is joining a network?

The default gateway and the address of a DNS server. Other possible answers include: a hostname

3a) (10 pts) Complete the table below to show how Dijkstra’s link-state algorithm would build a routing table at node E for the following network. Note: the first row is **not** complete—be sure to fill in the missing entries.

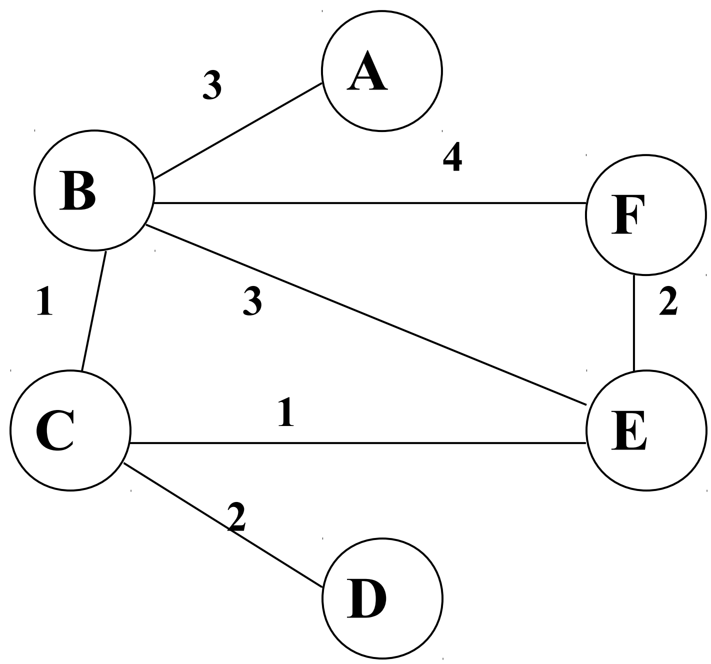

Table I: compute the distance and predecessor of each destination. At each step the N’ column should contain the list of nodes whose final distance is known at that step.

Step	N’	D(A), p(A)	D(B), p(B)	D(C), p(C)	D(D), p(D)	D(F), p(F)
0	E	***INF,-***	3,E	***1,E***	INF, -	***2,E***
1	***EC***	***INF,-***	***2,C***		***3,C***	***2,E***
2	***ECB***	***5,B***			***3,C***	***2,E***
3	***ECBF***	***5,B***			***3,C***	
4	***ECBFD***	***5,B***				
5	***ECBFDA***					

3b) (5 pts) Based on your Table I, what path would be used for a packet traveling from node E to node A? (No credit if your answer is not correctly derived from the table you computed in part a of this question.)
Use the table to find the predecessors: p(A) = B, p(B) = C, p(C)=E, so the path is ECBA.

4. (10 pts) Suppose that a router using a distance-vector algorithm currently holds the following routing table (not for the same network as the previous problem).

Destination	Next-hop	Cost
A	E	4
B	B	4
C	E	6
D	E	7
E	E	2

This router now receives the following from its neighbor, B, to which it is attached by a cost 4 link.

Destination	Cost
A	1
C	1
D	2
E	2

Fill in the table below to reflect this router's routing table after processing the update above.

Destination	Next-hop	Cost
A	***E***	***4***
B	***B***	***4***
C	***B***	***5***
D	***B***	***6***
E	***E***	***2***

Consider each of the entries in the distance vector received from B. To get to A through B now costs 5; the existing route through E is better so it is unchanged; To get to C through B costs 5; this is better than the existing route. D through B costs 6, again better. E through B costs 6 which is worse

5. (5 pts) By definition *head-of-line blocking* occurs in a switch or router at
   a) output queues
   b) input queues
   c) switching fabrics
   d) memory

b) – input queues

6. (5 pts) Number the types of router architectures below from 1, least performance, to 3, greatest performance.

_______1___ Memory-based

_______3___ Crossbar-based

______2____ Bus-based

7. (10 pts) Consider a datagram network using 8-bit addresses. Suppose a router uses longest-prefix matching and has the following forwarding table. (Ignore the last column for a moment).

Prefix	Interface	# of hosts
00	***0***	***64***
010	***1***	***32***
01	***2***	***32***
1	***3***	***64***
10	***4***	***64***

a) (7 pts) in the 3rd column write the number of hosts that are reached via that interface (note that the prefixes cover overlapping address ranges and you have to consider that in your answers). Do not consider any of the addresses to be reserved and hence unusable.
There are 256 possible hosts in all. Work from longer to shorter prefixes: for interface 1 there are 5 bits of available host addresses, thus 32 hosts. For interface 2 there are 6 bits which would be 64 hosts, but 32 of these are captured by the longer prefix for interface 1, hence 32 for interface 2; For interfaces 0 and 4 there is no longer prefix, so 64; And for interface 3 there would be 7 bits, 128 hosts, except the prefix for interface 4 is longer and has already captured 64 of those, so 64 for interface 3.

b) (3 pts) Using the above forwarding table, what is the output interface for datagrams with the following source and destination addresses?

Source	Destination	Interface
0011 0001	0100 0010	1 (matches 010)
1000 0001	1111 0010	3 (matches 1 and does not match 10)
1100 0100	0111 1000	2 (matches 01 and doesn't mach 010)

Only the destination address matters!

8. (10 pts) Describe the idea of *tunneling* and describe how it can be used to carry IPv6 traffic across portions of the Internet that are not IPv6 enabled. Things to consider: does an IPv6 router at the start of a tunnel have to know the IPv4 path to the IPv6 router at the end of the tunnel? Does it have to know the IPv4 address of the router at the end of the tunnel? Do the routers at both ends of the tunnel have to be aware that tunneling is being used?
At the start of the tunnel, the router encapsulates the IPv6 packet into an IPv4 packet with its own address as the source and the tunnel end router as the destination. This packet is carried over the IPv4 to the end router, which unencapsulates the IPv6 packet and sends it on the IPv6 network. Both of the tunnel-end routers must know that

tunneling is being used, and the router at the start must know the IPv4 address of the tunnel at the end.

9. (5 pts) Why do VC routers typically change the VC identifier in packets each time the packet passes through a switch?

If hop-by-hop VC numbers were not used, VC numbers would have to be globally alllocated which would be an unnecessarily hard thing to do.

10. (5 pts) Why would a datagram network generally suffer less disruption from a link outage than would a VC network?

When a link outage occurs in a VC network data stops flowing until a new VC is established that does not involve the down link. In a datagram network, the routers near the down link update their routing tables to send data over a different path.

11. (5 pts) Explain why broadcast or multicast using the controlled flooding techniques (either sequence-number controlled flooding or reverse path forwarding) use more network resources (i.e. more total bits are sent on the network) than do techniques based on spanning trees.

When using spanning trees every packet sent over the network is useful. When using sequence number controlled flooding or reverse path forwarding packets are stopped from further forwarding only after they have made one useless trip over a link – i.e. a packet is stopped from being forwarded only after it arrives at a place it has already been received.

The above questions mainly pertain to the network layer; in addition to those topics there will also be questions about the link layer. Topics include: what is the advantage of CSMA/CD over CSMA? What's in an ARP table? ARP protocol: given a network configuration what steps occur when a host wants to send a network-layer message to another host on the same subnet? On a different subnet? What is the difference between switches and routers? How does a switch's forwarding table get populated and what does it contain? Why are switch forwarding table entries given a time-to-live?