COMP 361 Computer Communications Networks

Spring Semester 2004

Final Examination – Solution Key

Date: May 24, 2004. Time: 4:30pm – 7:30pm. Venue: Rm LG4204

Name:______________________Student ID:________________Email:_______________

Instructions:

1. This examination paper consists of 17 pages and 10 questions
2. Please write your name, student ID and Email on this page.
3. For each subsequent page, please write your student ID at the top of the page in the space provided.
4. Please answer all the questions within the space provided on the examination paper. You may use the back of the pages for your rough work.
5. Please read each question very carefully and answer the question clearly and to the point. Make sure that your answers are neatly written, readable and legible.
6. Show all the steps you use in deriving your answer, where ever appropriate.
7. For each of the questions assume that the concepts are known to the graders. Concentrate on answering to the point what is asked. Do not define or describe the concepts unless specifically asked to.

Question	Points	Score
1	8	
2	15	
3	6	
4	15	
5	15	
6	8	
7	8	
8	5	
9	5	
10	15	
TOTAL	100	

1) Answer the following true/false questions by circling either T or F. (8 points)

a) OSPF is the only Intra-AS routing protocol (an Intra-AS protocol is one that routes *within* Autonomous Systems) in use today T **F**

b) Bridge forwarding is *cut-through.* T **F**

c) Slotted Aloha is more efficient than pure Aloha **T** F

d) PPP uses a CSMA access-control protocol T **F**

e) ATM is an end-to-end (desktop-to-desktop) network protocol **T** F

f) In Selective Repeat protocols the loss of one packet can force the retransmission of many correctly received packets. T **F**

g) Domain Name Servers don't use caching to improve performance. They perform full name resolution each time they receive a name resolution request. T **F**

h) A two-dimensional parity check scheme can correct all two-bit errors in the original data. T **F**

2) (15pts)

a) What is meant by the **Hidden Terminal** problem? Why is this a problem?

b) Draw a diagram illustrating the Hidden Terminal Problem (you must explain how your diagram illustrates the problem.)

c) The 802.11 Wireless LAN protocol suite uses CSMA/CA (Collision Avoidance) to help solve the hidden terminal problem.

Explain how CSMA/CA works and why it solves the hidden terminal problem.

When explaining how CSMA/CA works you should assume a scenario in which there is a wireless network with (i) a sender, (ii) a receiver and (iii) other nodes. You should describe the actions taken by *all* of the members of the network when implementing the protocol.

(Note: in your explanation of how CSMA/CD works you do not have to write the formal technical names of all of the packets sent, but you do need to describe their functions.)

a) The hidden terminal problem arises in wireless LANS.

The problem is that (two) nodes A and B might not be able to hear each other's transmissions because they are blocked (hidden) from each other. They are both, though, in line of transmission with a third node C with which they are trying to exchange messages. The difficulty is that, since A and B can't hear each others' transmissions, they are unable to sense any collision between packets that they are trying to send to C. Therefore, Collision Detection algorithms can not work in this situation and a CSMA/CD protocol is unworkable here.

b) See p. 485 of the textbook or p 94 of the chapter 5 notes.

c)

1) sender senses channel for DISF time. If channel is idle, sender transmits a short RTS (request to send) packet (which indicates length of transmission being requested). If channel is not idle, sender performs a backoff procedure.

2) receiver waits a SIFS mount of time and then, if channel is idle, replies with short CTS (clear to send) packet. As well as informing sender that it can now transmit this has added effect of notifying (possibly hidden) other nodes that the transmission will take place. CTS also "reserves" channel for a requested amount of time by saying how long the expected message will be.

3) Sender then sends actual data packet

4) Receiver then sends an ACK

5) Other nodes (who have heard the RTS or CTS) know that a transmission is taking place and will refrain from broadcasting during the reserved amount of time.
Amount of time that these other nodes have calculated to wait is their network allocation vector (NAV).

This solves the hidden terminal problem since the nodes whose transmission could collide with the current transmission will ***avoid*** *broadcasting during the NAV period.*
That is why this protocol is described as ***collision avoidance.***

Student ID___________________

3) (6pts) Consider the picture below. Suppose Host A sends a datagram to host B. Assume that the datagram sent is small enough to fit into one link-layer frame.

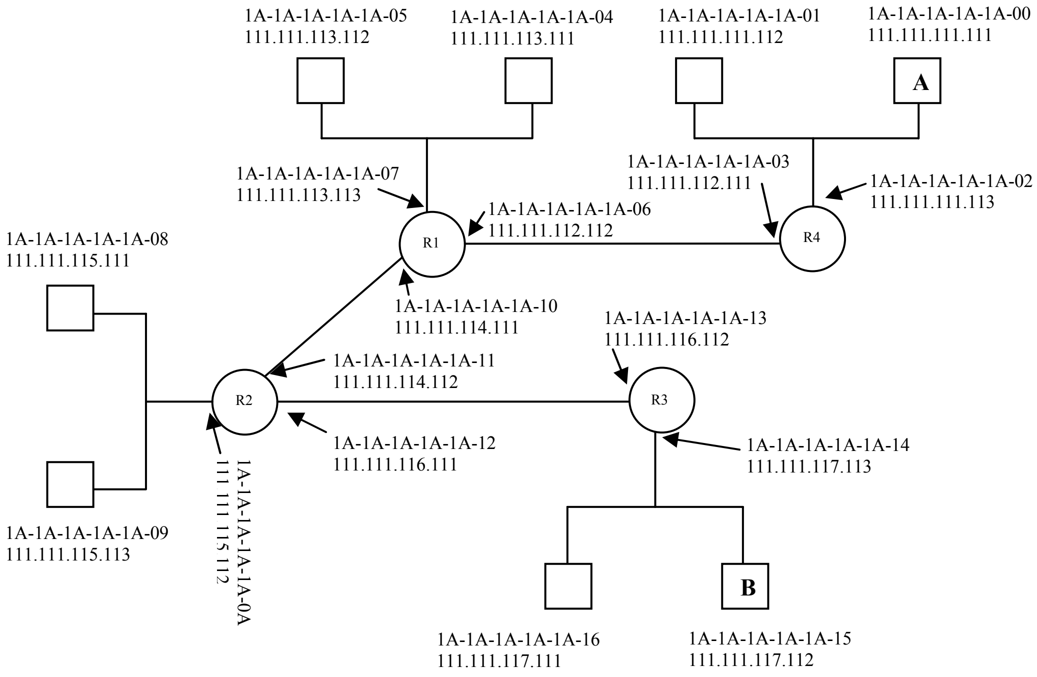

During the transmission of this datagram from A to B a frame is sent from R4 to R1, a frame is sent from R1 to R2 and a frame from R2 to R3. (Other frames might be sent as well but in this question we are not concerned with them).Each of these frames contains a

(1) frame source (MAC) address, (2) a frame destination (MAC) address

and an encapsulated datagram containing a

(3) source IP address and a (4) destination IP address.

In the table below the diagram fill in the values of these four items for each of the 3 frames indicated

Frame sent on link from	Frame Source (MAC) Address	Frame Destination (MAC) Address	IP Datagram SourceAddress	IP Datagram Destination Address
R4 to R1	**1A-1A-1A-1A-1A-03**	**1A-1A-1A-1A-1A-06**	**111.111.111.111**	**111.111.117.112**
R1 to R2	**1A-1A-1A-1A-1A-10**	**1A-1A-1A-1A-1A-11**	**111.111.111.111**	**111.111.117.112**
R2 to R3	**1A-1A-1A-1A-1A-12**	**1A-1A-1A-1A-1A-13**	**111.111.111.111**	**111.111.117.112**

4) (15pts)

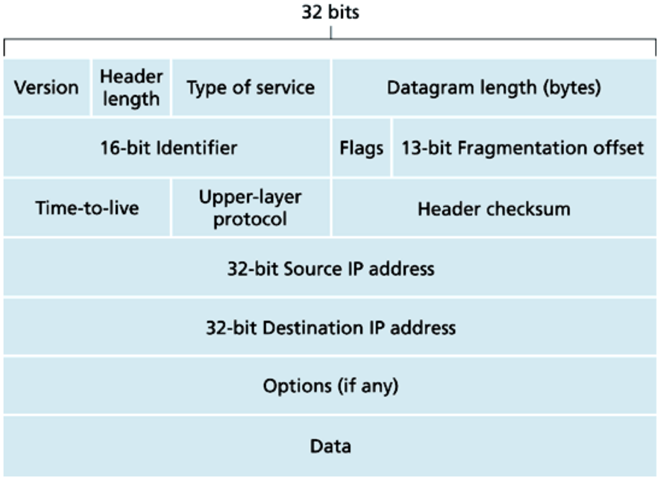

The picture above describes the format of an IPv4 datagram. The diagram below illustrates Router R1 sending a datagram to host H through Router R2.

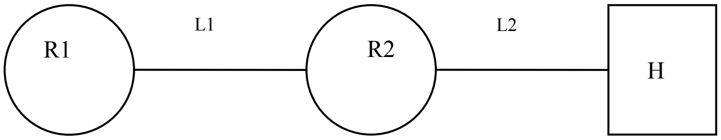

Link L1 only permits a MTU of 1000 bytes. Link L2 permits a MTU of 1500 bytes.
(MTU= Maximum Transfer Unit)

A is an IP datagram which

i) Has size 4000 bytes (the size of a datagram includes its header)

ii) Is not using any of the option fields in its header.

Because A is larger than the MTU of Link L1, A is *fragmented* when it is sent over L1.

a) Into how many IP datagrams is **A** fragmented when it is sent from R1 to R2 over L1? What is the size (in bytes) of each of these smaller fragments?

b) For some fields of the IP header, the fragments created all contain the same value. For some fields, the fragments contain different values. In particular the *fragmentation offset* and *offset-flag bit* are different in different fragments. For each of the fragments described in (**a**)   give the the value of the *fragmentation offset* and *offset-flag bit.*

c) In order for Host H to receive the data in **A**, R2 must also send some datagrams to H. Describe the datagrams that R2 sends to H. How many are there and what are their sizes?

d) Explain how the *fragmentation offset* field and *offset-flag bit* are used to reconstruct datagrams.

e) IPv6 does not support fragmentation. Give one reason why the designers of IPv6 decided not to support fragmentation.

a) and b) 5 fragments.
Note that A has 20 bytes of header and therefore only has 3980 bytes of data.

Fragment	*Size (header + data)*	*flag*	*offset*
1	*1000 (20+980)*	*1*	*0*
2	*1000 (20+980)*	*1*	*980*
3	*1000 (20+980)*	*1*	*1960*
4	*1000 (20+980)*	*1*	*2940*
5	*80 (20+60)*	*0*	*3920*

c) Since fragments are only reassembled at the destination (H) R2 only has the responsibility of forwarding the 5 fragments it receives from R1 to H. Since all 5 fragments have size less than 1500 (the MTU of L2) R2 simply forwards the 5 fragments it receives, as is, without further fragmenting them.

d) The receiving host can tell (i) which fragment is the first (offset =0), which fragment is the last (flag = 0) and, given two fragments B and C, whether C comes immediately after B (if offset B + length B = offset C).

The receiving host then takes all fragments it receives with the same ID and checks to see if they can be put together from first to last without any gaps. If they can, that's the reconstructed datagram. If not, it behaves as if the full datagram was lost.

e) Fragmentation (and reassembly) are time consuming operations that can slow down IP-forwarding in the network. Removing them from the router can speed up IP-forwarding in the network.

Student ID___________________

5) (15pts) Consider the following linear network.

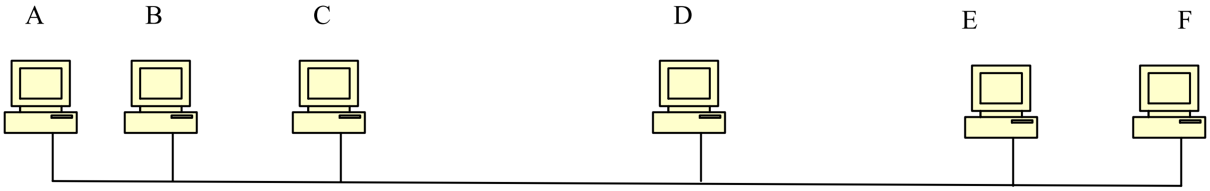

Assume that propagation speed of electromagnetic waves sent over the medium is $2*10^8$ m/s and the transmission rate of the network is 10Mbps.

Also, assume that we are using a CSMA/CD protocol with minimum frame size of 400 bits.

What is the *maximum* length of the linear network , i.e., distance from A to F, that will ensure that CSMA/CD will work properly for this network? Explain your reasoning.

The solution to this problem is a slightly modified form of the presentation in the supplementary notes of week 13.

Let d = distance between hosts X and Y.
*Let R = 2*10^8 m/s be the propagation speed of the network*
Suppose host X starts transmitting a frame of size F.
Then, before hearing the first bit of the frame, host Y starts transmitting a frame.
Since Y transmits BEFORE hearing X this can be at most d/R time after X starts transmitting.
X hears the beginning of Y's transmission d/R time after that.

Therefore, the MAXIMUM amount of time that can pass from the time that X starts transmitting until X hears the first bit from Y is 2d/R.

Therefore we must have F/10Mbps >= 2d/R or FR/20Mbps >=d.

Plugging in all of the values we find that ***4km = 4 X 10^3 m >=d.***

Student ID___________________

6) (8 points)
Consider the network below with the given link costs.

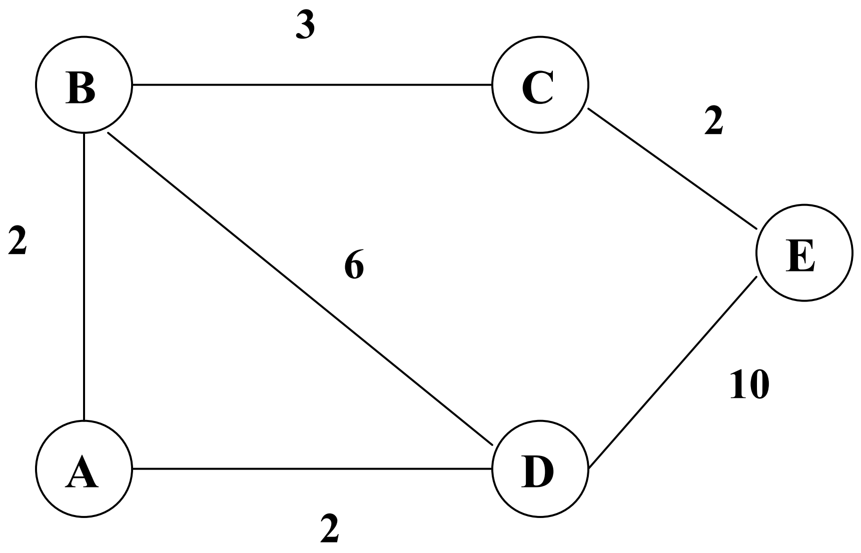

Fill in the final values in C's distance table that will result after running the *distance vector routing algorithm* with *Poisoned Reverse*.

		cost to dest via	
	D^C ()	B	E
dest	A	5	inf
	B	3	inf
	D	7	inf
	E	inf	2

7) (8pts) In this question you must calculate CRC bits.

The input data is the string of 4 bits **D = 1 0 0 1.**

Set r=3. The r+1 generator bit pattern used by the algorithm will be **G = 1 0 1 0.**

Given **D** and **G** as above, find the r CRC bits generated by the CRC algorithm.
For full credit you must show your calculations.

$D(x) = x^6 + x^3$
$G(x) = x^3 + x$

Dividing (modulo base 2) we get
$D(x) = G(x)\,(x^3 + x + 1) + (x^2 + x)$

so
$R(x) = x^2 + x$

And (since R is supposed to have 3 bits)

$R=110$

Student ID___________________

8) (5pts) Consider the following Inter-AS routing Scenario

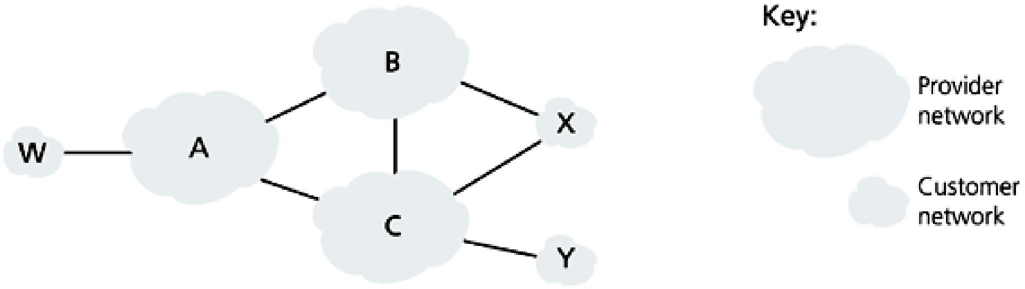

Customer X is dual-homed, i.e., it is connected to two different service providers B and C. Customer X *does not* want to route datagrams from Network C to Network B.

Explain how BGP allows customer X to avoid routing datagrams from Network C to Network B.

X will not send any BGP advertisements to Network C advertising a route to Network B.

Therefore, C will never know that a route to B via X even exists.

9) (5 pts)

a) What is the difference between classful addressing (the original scheme used in the Internet) and classless addressing (CIDR)?

b) What is the reason that CIDR was introduced?

a) IP addresses are 32 bits long. They are split into a (prefix) network address followed by a (suffix) host address. All machines on the same network will share the same network address.

Classful addressing essentially required that the network address component be one of 8 (class A), 16 (class B), or 24 (class C) bits. Those were the only options allowed.

*In classless addressing the network part of an IP address can be **any** number of bits long. This is often denoted by the **dotted decimal** form a.b.c.d/x where x is the number of (leading) bits that compose the network component of the address.*

b) A reason the CIDR was introduced is that classful addressing led to inefficient use of the IP address space and the internet was running out of addresses.

More specifically, Class C networks only allow a very small number (254) of network addresses so they were too small for many organizations. Most small and medium size organizations therefore requested Class B networks.. But these have an address space of 65,634 addresses, which is more than most of these organizations needed. Therefore, much of the IP address space was going to waste.

CIDR permits organizations to request (and pay for) an address space that more closely matches their needs, permitting more efficient use of the IP address space.

Student ID__________________

10) (15 pts) Bridges are *self-learning* (plug-and-play) devices that implement filtering/forwarding of frames on the same LAN. A bridge maintains a *bridge table*, which is a mapping between MAC addresses of nodes in the LAN and its (the bridge's) interfaces.

a) When the bridge is first turned on, its bridge table is empty.
Explain how a bridge learns the information that it uses to fill in its bridge table.

b) When a bridge detects a frame on one of its interfaces it filters the frame and decides whether the frame should be forwarded to other interfaces and, if so, to which interface(s). Describe the algorithm that the bridge uses to do this filtering and forwarding.

a) Whenever a new frame is received, bridge "learns" location (=the interface from which the frame enters) of sender and records sender/location pair in bridge table

b)

```
When bridge receives a frame:
index bridge table using MAC dest address
if entry found for destination
then{
    if dest on segment (interface) from which frame arrived
      then drop the frame
          else forward the frame on interface indicated
      }
else (i.e., no entry found)  flood by forwarding  on all  interfaces except the one on  which the frame arrived
```

c) Consider the following diagram of a LAN composed of three LAN segments connected by a bridge. Assume that the bridge table starts off empty and the given sequence of frames is sent, in the given order (frame 1, frame 2, etc.). As the bridge filters the frames it is using the standard bridge learning algorithm to fill in its table. For each frame write down to which interfaces (if any) the Bridge forwards the frame (the answer to the first one is filled in for you) Note that it is possible that in some cases the Bridge will drop the frame and not forward it. If this happens you should say so.

You should assume that the entries in the bridge table TTL fields are large enough so that no bridge table entry times-out while these frames are being sent.

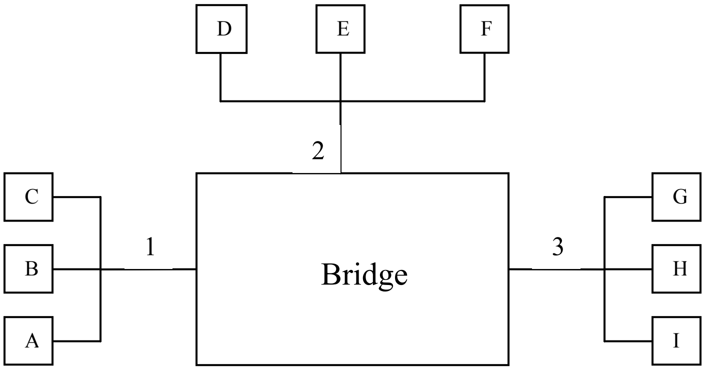

Frame	Source node	Destination Node	Bridge forwards frame to interfaces #
1	A	D	2,3
2	H	D	***1,2***
3	C	H	***3***
4	G	H	***Nothing (drop frame)***
5	E	D	***1,3***
6	B	E	***2***