Carnegie Mellon
Computer Science Department

15-441 Spring 2009
Midterm

Name: ______________________________

Andrew ID: ______________________________

INSTRUCTIONS:

There are **??** pages (numbered at the bottom). Make sure you have all of them.

Please write your name on this cover and at the top of each page in this booklet **except the last.**

If you find a question ambiguous, be sure to write down any assumptions you make.

It is better to partially answer a question than to not attempt it at all.

Be clear and concise. Limit your answers to the space provided.

A	B	C	D	E	F	G	H	I	Total
/ 8	/ 10	/ 6	/ 9	/ 6	/ 10	/ 20	/ 8	/ 3	/ 80

A Multiple Choice

1. This question tests your understanding of the layers in the Internet Protocol Suite. For each item in the list on the right, enter a letter (A–F) from the list on the left to indicate the corresponding layer.

A. application
B. transport
C. Internet
D. link
E. physical
F. none of the above

____ HTTP
____ TCP
____ UDP
____ IRC router
____ IP
____ ethernet
____ carrier pigeon
____ Manchester encoding

Solution:

HTTP: (A) Application
TCP: (B) Transport
UDP: (B) Transport
IRC router: (A) Application
IP: (C) Internet
Ethernet: (E) Physical
Carrier Pigeon: (E) Physical or (F) None of the above
Manchester Encoding: (E) Physical

B True/False

____ UDP uses a 3-way handshake

____ SONET has a method for resolving clock skew

____ The cyclic redundancy check for the IP packet header is used primarily to prevent routing loops

____ OSPF never has routing loops

____ Sometimes it is neccessary to generate multiple tokens in a token ring

____ CSMA uses exponential backoff

____ The ethernet spanning tree protocol in guaranteed to maintain shortest paths between nodes in a switched network

____ ARP can be used to find the IP adresses of computers in an ethernet network

____ Fragmented packets are reassembled whenever possible by routers

____ There are thousands of root nameservers

Solution: False: UDP has no connection setup phase.
True: it does have a method.
False: CRC is mainly used to check packet integrity.
False: Loops can always occur while the network is converging.
False: It is never necessary to do so.
True: CSMA uses exponential backoff for collision reduction and fairness.
False: It only maintains shortest path to the root.
False: ARP is used to go from IP to MAC.
False: Fragments are only reassmbled at the destination host.
False: The number of root name servers is small due to heiarchy.

C Short Answer

2. Give an example of a network that is configured to use DHCP, but does not perform any sort of NAT.

Solution: CMU's NetReg and Wireless Andrew systems, for example, use DHCP to provide globally routable addresses.

Give an example of a network that is configured to use both DHCP and NAT.

Solution: A standard consumer "home router" uses such an arrangement by default.

What network equipment and configuration would one use to provide Internet access to several computers, including one that acts as a public web server, from a home connection with a single external IP address?

Solution: Use a consumer NAT router, and either reserve an address via DHCP or manually assign one for the web server. Then, configure the router to forward the HTTP port through NAT.

D DNS

3. For answering the following question, consider the following name servers and the entries their translation tables contain:

local name server

Name	Value	Type	Class
com	ns1.nstld.com	NS	IN
ns1.nstld.com	192.5.6.32	A	IN

name server ns1.nstld.com

Name	Value	Type	Class
funsite.com	dns.funsite.com	NS	IN
dns.funsite.com	128.112.129.15	A	IN

name server dns.funsite.com

Name	Value	Type	Class
lotsafun.funsite.com	128.112.136.72	A	IN
manatee.funsite.com	128. 112.155.166	A	IN

(a) Suppose client1 queries the local name server for lotsafun.funsite.com. The local name server will recursively resolve the request and cache the results with a TTL of 10 minutes. Make a diagram of what will happen involving the client and the three nameservers, and writing down all DNS requests and responses.

(b) DNS generally uses UDP, and each request includes a transaction ID. After issuing a request, a name server listens on the UDP socket for a response that matches the transaction ID. Danny has noticed that it takes some time for the local name server to resolve his requests, and can't resist the temptation to cause trouble. He wants to send a single request for lotsafun.funsite.com to his local name server, and then trick it into storing (and returning to him) an incorrect translation. If Danny is willing to spoof a source address and blast traffic at his local name server, how might he pull this off?

Solution: The DNS packet the isp asks for can come from Danny. He asks to resolve page, before the websites DNS server can responds, he responds as though he is that website's DNS server and gives an incorrect result, the actual result from the isp server is thrown out by the isp server. must send a lot of requests and responses to the isp server to get a transaction id match.

(c) Assume for the next part that each name server can connect to the client or any other nameserver with RTT = 1.6s and bandwidth 100Mbps. Also, let all DNS packets have size 1 Kbit. For this problem assume the local server has the result for dns.funsite.com from NS1 cached. There are 65536 possible transaction IDs that can be associated with a request, and suppose that the local nameserver always choses from among 100 UDP ports to listen for a response. Suppose that both the transaction ID and the UDP port are chosen at random. Given that Danny sends packets as fast as possible to the local name server, what is the probablity he will trick it?

Solution: Danny can send at 10^5 kb/s for 1.6 seconds so he can send 160,000 Kb. So he has 160,000 chances to compromise the server in the 1.6 seconds, $160000/(100 * 65536) = .02$.

E Routing

4. Switches and routers are both responsible for forwarding data between networks and network segments, but they operate on different layers and are as a result very different. Please discuss the following differences:

   (a) How do the addresses differ?

   Solution: Hierarchical IP addresses versus flat (with respect to address look up) 48 bit IEEE addresses.

   (b) How does the structure of the forwarding tables differ?

   Solution: Switches use a simple look up of the flat addresses, e.g. using a content addressable memory. Routers need to deal with CIDR address look up, which usually uses tree.

   (c) What is different about how the forwarding tables are filled in?

   Solution: Routers use full blown routing protocols such OSPF or BGP. Switches uses a spanning tree protocol.

5. For this problem, consider the following network topology:

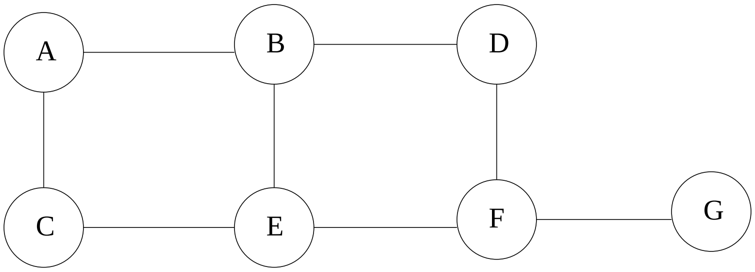

Suppose that the routers are running RIP, and are all turned on simultaneously. Recall that in RIP, the weight assigned to each link is one. You can also assume that if RIP identifies multiple paths of equal cost, it picks one randomly.

The initial routing table at node A is:

Destination	Cost	Next Hop
B	1	B
C	1	C
D	∞	—
E	∞	—
F	∞	—
G	∞	—

Fill in the following tables to show the initial routing table at node F:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

Solution: 3pts, -1 per mistake

Destination	Cost	Next Hop
A	∞	—
B	∞	—
C	∞	—
D	1	D
E	1	E
G	1	G

6. Now show the contents of the routing table after each iteration of the algorithm:

(a) After iteration 1:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

Solution: 4pts, -1 per mistake

Destination	Cost	Next Hop
A	∞	
B	2	D or E
C	2	E
D	1	D
E	1	E
G	1	G

(b) After iteration 2:

Destination	Cost	Next Hop
A		
B		
C		
D		
E		
G		

Solution: 4pts, -1 per mistake

Destination	Cost	Next Hop
A	3	D or E
B	2	D or E
C	2	E
D	1	D
E	1	E
G	1	G

7. In some failure situations, the administrator notices that it takes an exceptionally long time for the routing protocol to stabilize in this network.

   (a) What problem with RIP is the cause?

   Solution: 3pts
   Count to infinity

   (b) The administrator is told that BGP does not suffer from this problem. What prevents BGP from having this problem?

   Solution: 2pts
   Path vector - passing path info with distance eliminates loops

F Physical and MAC Layers

8. Show the NRZ, Manchester, and NRZI encodings for the bit pattern shown in the figure below. Assume that the NRZI signal starts out low.

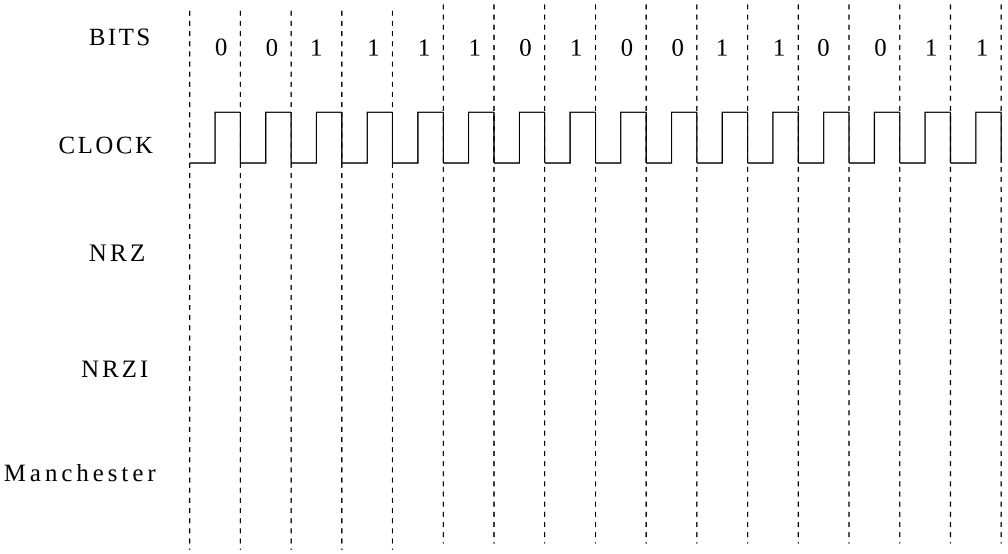

Charles, an ex-441 student, is given the task of building a new network link technology. Unfortunately, many of his beta-testers complain that their packets get corrupted when using his technology! He tracks the problem down to time synchronization problems between the sender and receiver on the link. Perhaps you can help Charles solve his problems by telling him a little about different encoding methods. Identify the problem and give a 1-2 sentence explanation about why this occurs.

For each of these sub-parts, identify whether the encoding can have problems with:

A. Long strings of 0s

B. Long strings of 1s

C. Both long strings of 1s or long strings of 0s

D. None of the above

(a) Manchester encoding

(b) NRZ

(c) NRZI

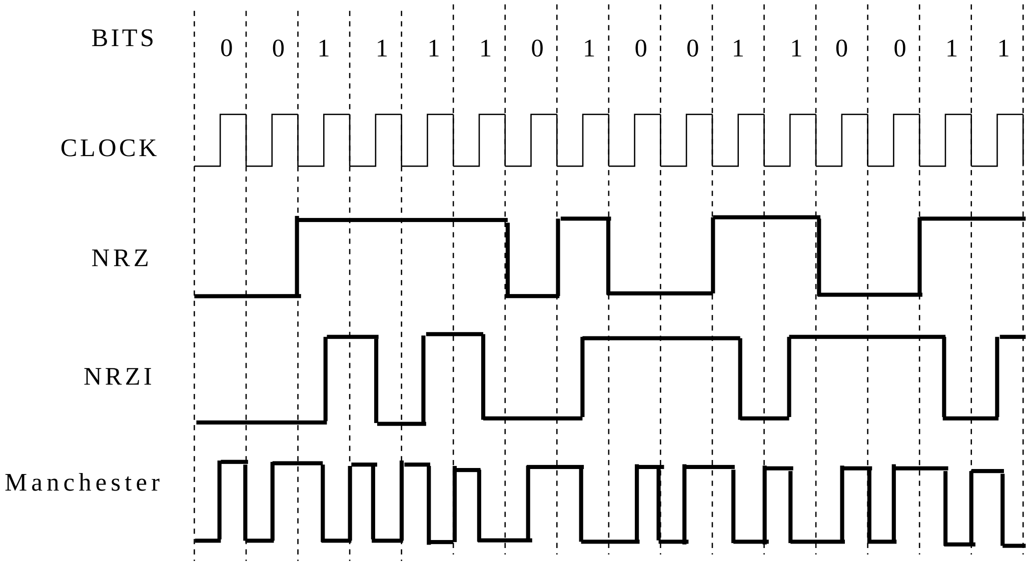

Solution:

(a) Manchester encoding: (D) Non-of the above. Manchester encoding will provide a signal change on every bit of data transferred, but has the down-side of making the data-rate only have the baud rate (i.e., rate at which the signal can change on the wire).

(b) NRZ: (C) For non-return to zero, the signal only changes when the data on the wire changes. Thus, it can lose clock synchronization with a sequence of either 1's or 0's.

(c) NRZI: (A) Non-return to zero inverted signals a 1 by making a transition, but signals a 0 by staying at the same signal. Thus, it has problems with long sequences of zeros.

G TDMA/FDM

9. George wants to setup a wireless network in a reserved frequency band for only the 15-441 staff to use. **The frequency band reserved for the staff is 500MHz** wide, for which Peter, Bruce, David, George, and Jacob will share. George, who did not pay attention during lecture, is trying to figure out the best way to multiplex the channel such that each member can achieve the best throughput in the network. The only method he can currently remember is time division multiple access (TDMA). To test the network the staff will **perform a 1000KB file transfer** simultaneously. Ignore propagation delay and handshaking, and assume **each user has one transceiver that is tunable to any frequency and can transmit data at a maximum rate of 2Mbps, using 100MHz** of the channel. There is a **maximum packet size of 1000 bytes**, no header is used. Assume a noiseless channel (no loss will occur), such that no ACKing or collision detection scheme is needed. Assume kilo means 1000 in all cases (e.g., $1000KB = 1 * 10^6 bytes$).

(a) In the TDMA protocol that George develops, each staff member is allowed to transmit a single packet, to avoid interference 6ms of silence is required, followed by the next staff member's packet transmission. Therefore, a single "round" of transmissions looks like this, note the order: George, 6ms, Bruce, 6ms, David, 6ms, Peter, 6ms, Jacob, 6ms. Using this protocol, how long will it take George's 1000KB file transmission to finish?

Hint: Break the times down in this problem, such as the time to transmit a packet and a total round time (e.g., the time it takes from the start of George's first packet to the start of George's second packet).

Solution:

The total number of packets (also represents the total number of TDMA rounds) needed to complete the 1000KB file transfer is: $packets = data\_size/packet\_size = 1 * 10^6/1000 = 1000$.

A single packet transmission takes: $T_{packet} = 8 * 1000bits/2 * 10^6 \frac{bits}{sec} = 0.004seconds$.

There is an idle time of 6ms after each transmission (commonly referred to as a guard time), so the total time for a node to transmit is: $T_{node} = T_{packet} + T_{guard} = 0.004s + 0.006s = 0.01seconds$.

A total round time therefore takes: $T_{round} = nodes * T_{node} = 5 * 0.01 = 0.05seconds$.

For George to transmit the whole file, it takes 999 rounds and then only T_{packet} additional time, since George transmits at the start of every round. This covers all 1000 packet transmissions. Therefore: $T_{George} = 999 * T_{round} + T_{packet} = 999 * 0.05s + 0.004s = 49.954seconds$!

(b) TDMA is not the optimal way to multiplex the given channel. There is a more efficient way which allows all staff member's transmissions to finish faster. Name the technique and briefly describe how it works. How fast can George's file transmission finish now?

Solution:

The staff could achieve much better network performance by using Frequency Division Multiplexing (FDM) since the total channel bandwidth is 500MHz and each transmission only requires 100MHz. Therefore, each staff member could transmit at the same time in the own channel without interfering with each other. George's transmission could therefore complete in the total time it takes to transmit the data only: $T_{George(FDM)} = packets * T_{packet} = 1000 * 0.004 seconds = 4 seconds$.

H Link Layer

I Collision Detection

10. The first random media access control (MAC) protocol developed in 1970 was the ALOHA protocol developed for wireless networks, which was a major building block for the Ethernet protocol. ALOHA's basic "collision" detection technique was that, after every transmission, the transmitter waited for an acknowledgement packet from the receiver to know that the transmission was successful. If an ACK was not received within a time period, the sender decided to retransmit the packet. Answer the following short questions.

    (a) After collision detection was developed for Ethernet, its technique was not adopted by wireless networks. Why is it not possible to achieve true collision detection in wireless networks?

    Solution: It is not possible due to the hardware. The transmission power from the sender is so great that it cannot detect another lower powered transmission at its receiver.

    (b) Finally, in wireless and wired protocols alike, a backoff occurs once a collision is detected. What is the purpose of the backoff, and why is the backoff value chosen randomly?

    Solution: The purpose of the backoff and random value is that so the two senders do not collide again by immediately transmitting again at the same time. Also, by choosing random backoff values we introduce some amount of fairness in the network.

The End – Phew!

J 4 Free Points for Tearing Off Page: Anonymous Feeback

List one thing you liked about the *class* and would like to see more of or see continued (any topic - lectures, homework, projects, bboards, topics covered or not covered, etc., etc.):

List one thing you would like to have changed or have improved about the class: