1. a. Draw the layers of the ISO OSI network architecture. (5)

    Application
    Presentation
    Session
    Transport
    Network
    Data Link
    Physical

    b. Draw the layers of the ARPA Internet network architecture. Describe how each of the layers relates to the layers of the OSI network architecture. (5)

Application	OSI Application, Presentation, and Session layers
Transport	Transport
Network	Network
Physical	Data Link and Physical

2. For each of the following activities, indicate into which layer of the OSI architecture it fits: (8)

    a. Routing packets between networks.   Network

    b. Handling an HTTP request.   Application

    c. Sending bytes out a serial port.   Physical

    d. Retransmitting a packet if no acknowledgement is received.
                                    Data Link or Transport

3. a. Suppose a communications link has a bandwidth of 10 MHz and a measured signal-to-noise ratio of 255. What is the theoretical maximum data transmission rate (in bits/sec) over the link? Show your work. (4)

$C = B \log 2 (1 + S/N)$ $C = $ Capacity, $B = $ Bandwidth, $S/N = $ Signal to noise ratio

$= 10^7 * \log_2 (256) = 80$ Mbps

b. Now suppose the cable used for the link is shielded and transmission power is increased, so the signal-to-noise ratio is improved to 1023. What effect does this have on the maximum data transmission rate – is it four times as much? (2)

$C = 10^7 * \log_2 (1024) = 100$ Mbps

Even though the S/N increase by a factor of four, the capacity is the log of the S/N, so it only increased by 25 %

4. Given the following clock signal and the data bits 0011010, draw the waveforms if the sender uses NRZ, NRZI, and Manchester encoding. (6)
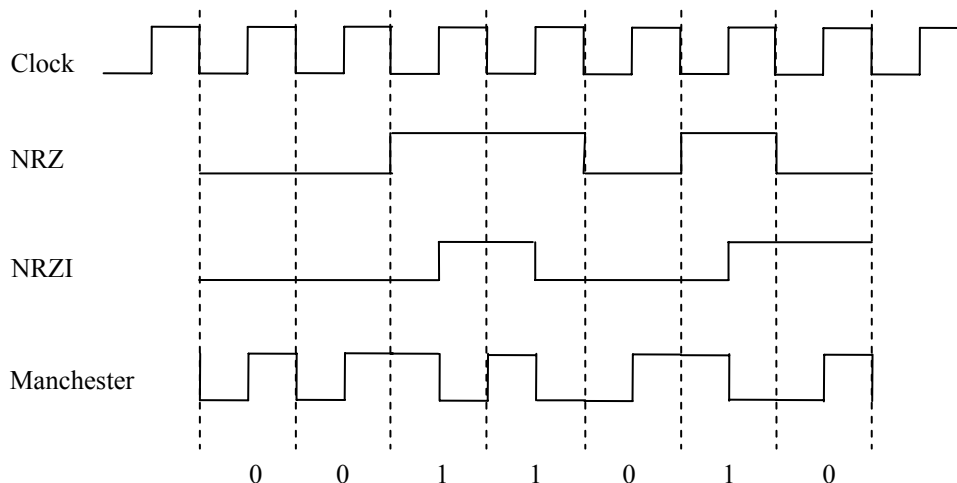
5. Why do the bit-oriented protocols in the HDLC family stuff a zero bit after runs of five one bits? (5)

Mainly to prevent data bits from being mistaken for the flag between frames (01111110). If NRZI encoding is used, also allows for clock recovery and prevents level shifts.

6. Explain how exponential backoff works in Ethernet. (10)

If node detects a collision when it transmits, it chooses a random 1-bit number (0 or 1) and multiplies it by 51.2 μsec. It waits this long before trying to transmit again. If the channel is clear and it can transmit, but it detects another collision, it chooses a random 2-bit number (0 to 3) and multiplies it by 51.2 μsec again to determine how long to wait. In general, after the nth collision, the sender choose a value between 0 and $2^n - 1$ as the multiplier. The value of n is capped at 10, but this process continues until either the sender knows the packet got through, or it has retransmitted 16 times. If that occurs, the packet is discarded as undeliverable.

7. a. Suppose you want to network a cluster of database servers with several client workstations using standard 10BaseT Ethernet, with 10BaseT hubs. If each of the workstations sends an average of one 1000 byte request to one of the servers each second, and the server responds to each request with a reply that averages 2750 bytes, estimate an upper limit on the number of workstations you could attach to the network before it starts to be overloaded by the traffic. Assume that the database servers can keep up with the load, and you have as many ports on the hubs as you want. (Hint: remember the difference between the theoretical throughput of a 10BaseT network and the practical throughput). (10)

Data load per workstation = 1000 bytes (request) + 2750 bytes (response) = 3750 bytes/wkstn
* 8 bits / byte = 30,000 bits/sec / wkstn

Network bandwidth for 10BaseT = 10 Mbps, but practical bandwidth is around 3 Mbps.

Number workstations = 3 Mbps / 30,000 bits/sec per workstation = 100 workstations.

b. Now suppose you get some extra money and can buy Ethernet switches instead of hubs. Would this change your estimated maximum number of workstations? Explain why or why not. If the number would change, would it go up or down? (5)

Yes, the number of workstations could increase. Switches eliminate collisions, and they may make it possible to send output on more than one port at a time. This would make the overall network capacity higher, so more workstations could share it.

8. If you receive the message 1101010100, and the divisor polynomial you are using is $C(x) = x^3 + x + 1$, does the message pass the CRC check? Show your work. (5)

```
                        1 1 1 1 0 1 1
        1011  |  1 1 0 1 0 1 0 1 0 0
                 1 0 1 1
                   1 1 0 0
                   1 0 1 1
                     1 1 1 1
                     1 0 1 1
                       1 0 0 0
                       1 0 1 1
                         1 1 1 0
                         1 0 1 1
                           1 0 1 0
                           1 0 1 1
                           0 0 0 1
```

Since remainder is not zero, this message fails the CRC check.

9. a. Suppose a link is capable of carrying 1 Mbps, and the round trip time of the link is 6 msec. If a node at one end of the link is sending packets to the other end using a stop and wait protocol, and the packets are all 1000 bits long, what is the maximum possible utilization of the link? (5)

Delay – bandwidth product = 10^6 bits/sec * 6 * 10^{-3} sec/trip = 6,000 bits/trip

Utilization = 1000 bits/msg / 6000 bits/ trip = 1/6

b. If you decided to use a sliding window protocol, how many frames would you want to have in flight to efficiently utilize the link (i.e. how big should the window be)? How many bits would you need for sequence numbers in the frame headers for this size of window? (5)

Send 6 frames without waiting for an ACK to keep the link fully loaded – so window size is at least 6.

This requires ceiling($\log_2(2 * 6 + 1)$) = 4 bits for sequence numbers.

10. a. Consider the following network, consisting of several Ethernet LANs connected by learning bridges. After the network has been up and running for a while, if all computers attached to the networks have sent several packets, show the forwarding table in each bridge. (5)
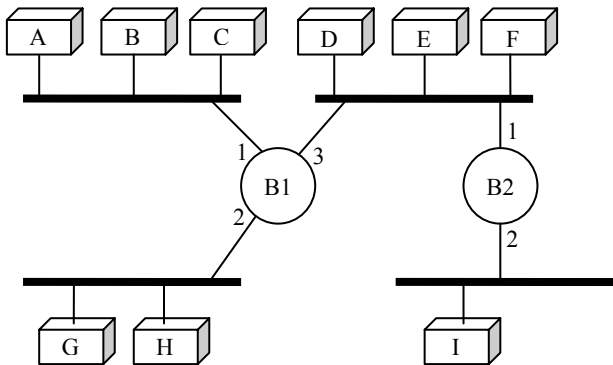
B1 1   A, B, C   B2   1   A, B, C, D, E, F, G, H
    2   G, H   2   I
    3   D, E, F, I

b. Explain what happens if node A sends a packet to node B, to node F, and to node I – on which ports does each bridge forward the packet? (5)

B1 receives packet to B on port 1 – this is B's port in forwarding table, so B1 does not forward packet.

B1 receives packet to F on port 1, forwards to port 3, since that's the port for F in the forwarding table.

B1 recieves packet to I on port 1, forwards to port 3. B2 receives packet to I on port 1, forwards to port 2.

c. If B1 is reset and now has an empty forwarding table, will that make it impossible for node A to send packets to node I? Explain what will happen when node A attempts to send a packet to I. (5)

A can still send to I – if B1 cannot find a node in its forwarding table, it just forwards the packet out all ports except the one on which it was received. So it would forward the packet on ports 2 and 3. B2 would do the same thing – it would forward a message to I on port 2. I would get the message, and in the process both B1 and B2 would learn a port for A.

11. Consider the following network
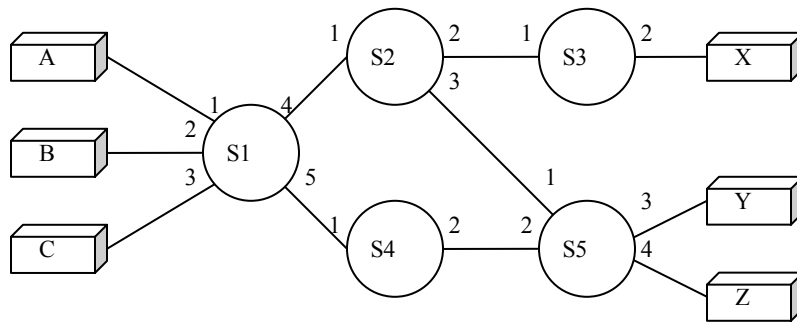
Suppose the switches have the following forwarding tables:

	S1	S2	S3	S4	S5
A	1	1	1	1	2
B	2	1	1	1	2
C	3	1	1	1	2
X	4	2	2	2	1
Y	5	3	1	2	3
Z	5	3	1	2	4

All the switches and nodes start out with empty virtual circuit tables, and then connections are created the following sequence:

1. A connects to X
2. A connects to Y
3. X connects to Y
4. Y connects to A

If all of these virtual circuits are still connected, show the virtual circuit table for each of the switches at this point – assume that each switch or node generates a virtual circuit ID by using the index of the next available entry in its virtual circuit table. (10)

	Input VCI	Input Port	Output VCI	Output Port
S1 (A → X)	1	1	1	4
(A → Y)	2	1	1	5
(Y → A)	3	5	3	1
S2 (A → X)	1	1	1	2
(X → Y)	2	2	2	3
S3 (A → X)	1	1	1	2
(X → Y)	2	2	2	1
S4 (A → Y)	1	1	1	2
(Y → A)	2	2	3	1
S5 (A → Y)	1	2	1	3
(X → Y)	2	1	2	3
(Y → A)	3	3	2	2