

The University of Queensland
School of Information Technology and Electrical Engineering
Semester Two, 2011

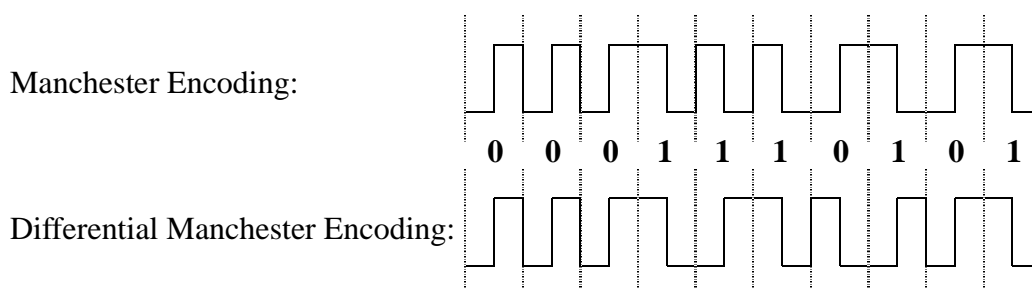
COMS3200 – Tutorial 7 - Solutions

MAC:

Question 1

Sketch the Manchester and differential Manchester encodings for the bit stream 0001110101.

(Note - the differential Manchester encoding actually depends on what comes before- we make an assumption here.)



Question 2

Ethernet frames must be at least 64 bytes long to ensure the transmitter is still going in the event of a collision at the far end of the cable. Fast Ethernet has the same 64 byte minimum frame size but can get bits out 10 times faster. How is it possible to maintain the same minimum frame size?

The 64 byte (512 bit) limit is based on a 2.5km long 10base5 cable. The 512 bits (equals 51.2 microseconds) is based on twice the end-to-end propagation time (about 5microseconds per km) plus time to propagate through repeaters. Fast Ethernet has shorter cable runs (100m of cable implies about 0.5microseconds, or about 50 bits at 100Mbps). (Note that there are only two transmitters on many Fast Ethernet segments and many of the standards support full-duplex transmission so you don't actually get collisions on a segment.)

Question 3

Consider building a CSMA/CD network running at 1 Gbps over a 1km cable with no repeaters. The signal speed (signal propagation time) in the cable is 200,000 km/sec. What is the minimum frame size.

For a 1 km cable, the one way propagation time is 5 μ sec, so $2a = 10 \mu$ sec. To make CSMA/CD work the station must keep transmitting for at least 10 μ sec. At 1 Gbps the station can transmit 10,000 bits, therefore the minimum frame would be 10,000 bits (1250 bytes).

Internetworking:

Question 1

If connection-oriented protocols are used for packet/frame forwarding they are based on the same principle for routers and switches. Using the example network given in Figure 1, give the forwarding tables for all the switches after each of the following connections is established. Assume that the sequence of connections is cumulative, that is, the first connection is still up when the second connection is established, and so on. Also assume that the connection identifier (CI) assignment always picks the lowest unused CI on each link, starting with 0.

a) Host A connects to host G, b) Host C connects to host G, c) Host E connects to host I, d) Host D connects to host B, e) Host F connects to host J, f) Host H connects to host A.

Table 1: Question 1 answer

Exercise part	Switch	Input port	Input CI	Output port	Output CI
a	1	2	0	1	0
	2	3	0	0	0
	3	0	0	3	0
b	1	3	0	1	1
	2	3	1	1	0
	4	3	0	1	0
c	2	2	0	0	1
	3	0	1	2	0
d	1	0	0	1	2
	2	3	2	0	2
	3	0	2	3	1
e	2	1	1	0	3
	3	0	3	1	0
	4	2	0	3	1
f	1	1	3	2	1
	2	1	2	3	3
	4	0	0	3	2

Question 2

Consider the connection-oriented (virtual circuit) switches in Figure 2. Table 1 lists, for each switch, what <port, CI> (or <CI, interface>) pairs are connected to what other. Connections are bidirectional. List all endpoint-to-endpoint connections.

In the following $S_i[j]$ represents j th entry in the table for switch S_i .

A connects to D via $S_1[1] - S_2[1] - S_3[1]$

A connects to B via $S_1[2]$

B connects to E via $S_1[3] - S_2[2] - S_3[2]$

Question 3

Consider the arrangement of learning bridges shown in Figure 3. Assuming all are initially empty, give the forwarding tables for each of the bridges B1-B4 after the following transmissions: A sends to C, C sends to A, D sends to C. Identify ports with the unique neighbour reached directly from that port; that is, the ports for B1 are to be labelled "A" and "B2".

When A sends to C, all bridges see the packet and learn where A is. However, when C sends to A, the packet is routed directly to A and B4 does not learn where C is. Similarly, when D sends to C, the packet is routed by B2 towards B1 only, and B1 does not learn where D is.

B1: A-interface: A B2-interface: C (not D)

B2: B1-interface: A B3-interface: C B4-interface: D

B3: B2-interface: A,D C-interface: C

B4: B2-interface: A (not C) D-interface: D

Question 4

Suppose a 10Mbps Ethernet hub (repeater) is replaced by a 10Mbps switch, in an environment where all traffic is between a single server and N clients. Because all traffic must still traverse the server-switch link, nominally there is no improvement in bandwidth.

a) Would you expect any improvement in bandwidth? If so, why?

b) What other advantages and drawbacks might a switch offer versus a hub?

a) After the upgrade the server-switch link is the only congested link. For a busy Ethernet the contention interval is roughly proportional to the number of stations contending, and this has now been reduced to two. So performance should increase, but only slightly.

c) A switch makes it impossible for a station to eavesdrop on traffic not addressed to it. On the other hand, switches tend to cost more than hubs, per port.