

**SCHDF0018 - HIGHER DIPLOMA IN COMPUTER SCIENCE EXAMINATION
ARBDF0015 – THIRD YEAR ARTS EXAMINATION**

COMPUTER SCIENCE

**COMPP303: Networks and Internet Systems
COMP3616: Networks and Internet Systems**

**OUTLINE SOLUTIONS WINTER 2004
Dr. J. Murphy**

Question 1

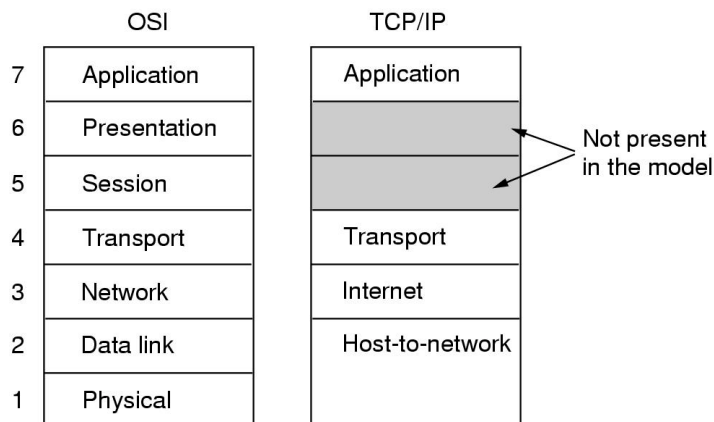
(1-a) IP address: 154.33.7.220
Mask: 255.255.255.192 (255 means all-1's)

Sub-network address is Boolean-AND of these: 154.33.7.192 The destination IP address is **154.33.7.x** with **x=11011100** Since this is a Class B address, the Net-id is 154.33, the 10-bit Subnet-id is 0000011111, and the 6-bit Host-id is 011100.

(1-b) packet transmission time $\text{pkt_trans} = (800)/(1000000) = 0.8$ millisecond. Therefore total travel time for a packet = $(\text{pkt_trans}) + 3 \times (4 \times \text{pkt_trans}) = (13 \times 0.8) = 10.4$ milliseconds.

- (1-c) Address Resolution Protocol (ARP):
- each host maintains an *ARP cache* of mappings between IP addresses and physical addresses
 - ARP cache entries *time out* after (typically) 15 minutes
 - if destination's IP address not in sending host's ARP cache, sending host broadcasts *ARP query* on local network
 - ARP query asks for physical address for destination's IP address
 - ARP query includes sending host's IP address and physical address (so other hosts can enter this in their ARP caches)
 - if destination is a host on local network, it sends its physical address to sending host in an *ARP reply* if destination not on the local network, ARP server replies with its own physical address: *proxy ARP*

(1-d) in the Lecture Notes (or any computer networks textbook)



Question 2

- (2-a) TRANSF = 400 microseconds;
TIMEOUT = TRANSA + 2×PROP, since TIMEOUT chosen optimally and PROC=0
= 40 + 2×20 = 80 microseconds;

therefore $\text{throughput}_{\text{SW}} = (1 - 0.2)/((400+80) \times 10^{-6}) = \mathbf{1,667 \text{ packets/second}}$ (*not frames/second*)

and $\text{throughput}_{\text{GBN}} = (1 - 0.2)/((400+(0.2 \times 80)) \times 10^{-6}) = \mathbf{1,923 \text{ packets/second}}$

Since we wanted a throughput of at least 1,600 packets/sec, Go-back-n and Stop-and-wait can be used.

- (2-b) (i) Sequence number starvation is probably happening. We are running out of numbers to put in the frames.
(ii) Make bigger frames if we have only control over the frame size.

Question 3

(3-a) ***Datagram packet-switching:***

- Each packet is treated individually within the network, so successive packets may follow different routes through the network. Each packet contains the receiver's address and a sequence number (so that receiver can put them into correct order). Network nodes are routers, which have routing tables telling them which output link to use for each possible destination. No connection set-up needed.
- Flexible routing possible (e.g. if a router crashes).
- Network resources are *not shared at the same time*: each packet monopolises a link during its transmission, after which the link is available for other packet transmissions.
- Ideal for *short-lived* bursty traffic; less suitable for *long-lived* &/or *interactive* bursty traffic.

Virtual circuit packet-switching:

- A route is set up in the network between sender and receiver (by making appropriate entries in the routing tables).
 - Resources may or may not be reserved for this route. If resources need to be reserved and are not available, the connection request is blocked.
 - Each packet contains its virtual circuit identifier.
 - Routers have routing tables telling them which output link to use for each established virtual circuit.
- Connection set-up required, which may involve significant delay.
- Network resources are *not shared at the same time*: each packet monopolises a link during its transmission, after which the link is available for other packet transmissions.

- (3-b) (i) ***distance-vector:*** each router exchanges information about the entire network with neighbouring routers at regular intervals. Neighbouring routers = connected by a direct link (e.g. a LAN); regular intervals: e.g. every 30 seconds. Information exchanged = routing tables (details in lecture notes).

link-state: each router exchanges information about its neighbourhood with all routers in the network when there is a change. Neighbourhood of a router = set of neighbour routers for this router; each router's neighbourhood information is flooded through the network; change: e.g. if a neighbouring router does not reply to a status message. Information exchanged = link-state packets (details in lecture notes).

(ii) A:	Advertiser	Network	Cost	Neighbour
	A	12	1	B
	A	10	3	C

B:	Advertiser	Network	Cost	Neighbour
	B	12	3	A
	B	90	1	C
	B	3	1	D

C:	Advertiser	Network	Cost	Neighbour
	C	10	1	A
	C	90	2	B
	C	6	1	--

D:	Advertiser	Network	Cost	Neighbour
	D	3	2	B

Question 4

(4-a) token = special frame that is passed from one node to the next. A node may only transmit Frame(s) when it has the token. When a node receives the token, it removes it from the ring (“removing the token” means changing it into the start of a data Frame) and transmits one Frame, then puts the token back on the ring (i.e. transmits it to next node). Since only one node can hold the token at any given moment, no collisions can occur.

(4-b) (i) The limit is when **TRANSF** goes to zero:

$$\text{throughput} = 1 / (5.4 \times \text{PROP})$$

(ii) However a limitation on this limit is that there is a minimum frame size in Ethernet and therefore can not reach this limit and also no data is transmitted.

Question 5

(5-a) TCP: retransmissions can lead to high delay and delay jitter; doesn't support multicast; slow start congestion control mechanism not suitable for continuous media.

UDP: traditionally, no retransmissions or congestion control; supports multicast.

(5-b)	Transmission number	Sender's Congestion Window (kB)	Threshold (kB)
	0	3	50
	1	6	50
	2	12	50
	3	24	50
	4	48	50
	5	50	50
	6	53	50
	7	56	50
	8	59	50
	9	62	50
	10	3	31
	11	6	31
	12	12	31
	13	24	31
	14	31	31
	15	34	31
	16	3	17