Network Systems (201600146/201600197), Test 2

March 10, 2017, 13:45–15:15

Answers

1. Physical media and framing

(a) Answer C.
Twisting copper wires is done to cancel the effect of external electric/magnetic fields, but glass fibers are not sensitive to those.

(b) Answer D.
Due to the increased distance, the power from the transmitter has spread out over a larger space, so the receive antenna will capture less of it, hence lower signal strength.

(c) Answer C.
Due to the larger index of refraction difference, the critical angle decreases, so even light whose direction deviates quite far from the straight line can still reflect totally. Such a light ray bounces back and forth many times over the same length of fiber and thus has a large delay compared to the straight path.

(d) Answers G and L.
This is a matter of carefully first applying bit-stuffing (inserting a 0 after every five 1s), then putting flags at the beginning and end, and the applying NRZI-S encoding (change of level for every 0 bit, not for a 1 bit). By starting with + and with -, you get answers G and L.

---

2. Medium access control

(a) Here's the diagram, where the red area indicates the collision:

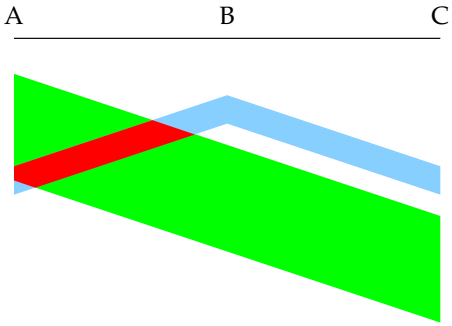

Variations are possible, but the following points are crucial:

- A must start transmitting *before* B's transmission reaches it, otherwise it wouldn't transmit (because of CSMA).
- B must stop transmitting *before* A's transmission reaches it, otherwise B would detect the collision.
- A's transmission must continue until *after* B's transmission reaches A, otherwise A does not detect the collision.

(b) The active node gets 1 out of N slots: $R/N = 10^7$ bits/s =10 Mbits/s

(c) The active node will get the entire capacity, because there is only a single transmitter, and there are no collisions: $R = 10^8$ bits/s =100 Mbits/s

(d) Sending a frame of Q bits takes the transmission time of that frame (Q/R) plus the time to poll all nodes $(N \cdot d_{poll})$; the resulting throughput is

$$\frac{Q}{Q.R + N \cdot d_{poll}} = \frac{10^5}{10^5/10^8 + 10 \cdot 0.0001} = \frac{10^5}{0.002} = 5 \cdot 10^7 \text{ bits/s} = 50 \text{ Mbits/s}$$

3. (Inter)Networking

(a) Answer A.
There is no mechanism in IP for asking for a retransmission of a lost fragment; so a packet of which one fragment is lost, is effectively lost entirely.

(b) Answer A.
C makes some sense, but the question asked for the *main* reason. Packet loss is relatively rare, and if it occurs, a simple retransmission of the packet will solve it. A failing server is much worse, hence answer A being preferred.

(c) Answer D (only).
The /24 notation means 8 of the 32 IPv4 address bits are available for addresses within the network, so it contains 2^8 addresses; in a /16 network, this is 2^{16}, which is more.

(d) Answer C and D.
The IP addresses identify end nodes (hosts) in the network, so there's no reason to change them as the packet travels from through the network.
The MAC addresses have local significance, on e.g. a single ethernet. As the packet travels over consecutive LANs, on each LAN it carries the MAC addresses of the sender and receiver on that network. See e.g. the example in the slides of lecture 9.

(e) Numbers 1 and 2 go together, as do numbers 3, 4 and 5.
Furthermore, number 7 is a complete packet by itself (offset=0, moreflag=0).

Many students tried to combine numbers 6 and 10, but that's not possible: they have different source addresses, so they cannot be fragments from the same packet.

4. Dijkstra's algorithm

(a)

Step	Confirmed	Tentative
1	(A,0,-)	(B,2,B), (D,6,D), (E,8,E)
2	(A,0,-), (B,2,B)	(C,4,B), (D,6,D), (E,3,B)
3	(A,0,-), (B,2,B), (E,3,B)	(C,4,B), (D,5,B)
4	(A,0,-), (B,2,B), (E,3,B), (C,4,B)	(D,5,B)
5	(A,0,-), (B,2,B), (E,3,B), (C,4,B), (D,5,B)	

(b) No.
Give e.g. link CE a cost of –2. This link played no role in question (a): it was never on the 'Tentative' list, so changing its cost will not change the outcome of the algorithm, so the algorithm will reach node E via the path A-B-E, with cost 3. But if the cost of CE is –2, then the path A-B-C-E has a total cost of 2, so should be preferred over A-B-E; however, the Dijkstra algorithm won't find it, since it has already put E on the 'confirmed' list by the time it considers C's links.
In fact, the situation is even worse. Once you have a link with a negative cost, each jump back and forth over that link lowers the cost. So by inserting many C-E-C steps into a path, you can make its cost arbitrarily low (negative). Thus, there is no meaningful "best" path anymore, if one or more link costs are negative.