Network Systems (201300179), Test 2

March 7, 2014, 15:45–17:15

Brief answers

1. Physical media, encoding, and medium access

2 pt (a)

Lower, otherwise the phenomen of total internal reflection cannot occur which is required to keep the light inside the core.

2 pt (b)

Total internal reflection cannot occur anymore, so at every ”bounce” some light will leak into the cladding, leading to a rather high loss of the signal in the core. It may still be usable over short distances, but not over long.

2 pt (c)

Higher dispersion means pulses are lengthened (smeared out) more, so fewer pulses can be sent per second without overlap, so lower bandwidth. S/N is same because of same equipment and same attenuation. Formula then indicates that C is lower.

2 pt (d)

It prevents the flag pattern (which contains 6 consecutive 1s) from occurring within the frame, by inserting a 0 after every 5 consecutive 1s in the data.

2 pt (e)

Bitstuffing ensures never more than 6 consecutive 1s occur, but long rows of 0s can still occur; sending this with NRZI-S, which has a transition for each 0 bit, then ensures that the signal changes frequently (never more than 6 times in a row the same signal level).

4 pt (f)

One obvious proposal would be to use something similar to DOCSIS from the cable TV networks; this is quite efficient (not many collisions), but requesting access to the medium takes time.
Another good choice would be to use a different optical wavelength (“colour”) for each home station. Advantage is the high capacity per home, while a disadvantage is that the central node has to have a separate optical receiver for each home.

Definitely **not** correct is e.g. CSMA, since that requires that the home nodes can hear each other’s signals, and it was given that they don’t.

---

2. Link-state routing

1 pt (a)

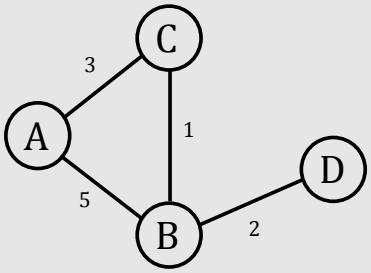

4 pt (b)

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

2 pt (c)

LSPs will immediately be dropped by the receiver, so that routers will not learn about links available in the network. Routing will fail completely, and network will not be able to forward any traffic.

2 pt (d)

Initially, LSPs will be correctly flooded through the network, and correct routes will be calculated. If changes in the topology occur, and new LSPs are created, these are not flooded though the network, because they do not have a new sequence number. This means that the routing will not adapt to the changing topology. So, initially, the network will function correctly, but as the topology changes, packets may not take the shortest route, or be dropped completely. After the TTL of the initial LSPs expire, these will be dropped, and new LSPs could flood successfully, resulting in a period of successful routing and packet delivery again.

2 pt (e)

Flooding of LSPs will completely overload the network, since, if there is a cycle in the network, each LSPs will loop in that cycle, continuously incrementing its sequence number until the time-to-live expires. Using the received LSPs, correct routes can be calculated, but the same routes have to be calculated over and over again, as updates of the LSPs arrive. The network could deliver some datagrams, but the network will be overloaded due to the circulating LSPs.

2 pt (f)

Initial flooding will be successful, and routes can be correctly calculated, although routers have a wrong notion of the current sequence number for LSPs. If topology changes occur that lead to shorter paths in the network, the corresponding LSPs will be successfully flooded, and changes will be correctly reflected in the route calculations. If topology changes occur that lead to longer paths (e.g., because links fail), the corresponding LSPs will be dropped during flooding, and the old (broken) routes will remain, resulting in non-delivered packets. Over time, new LSPs with increased sequence numbers will be flooded successfully, and correct routes can be calculated again. So, if topology changes do not occur too often, the network could function at some level.

3. Naming and addressing

2 pt (a)

The root domain name server resolves top-level domain names such as .nl and .com, and returns the names and addresses of the top-level domain name server that is able to resolve names within those TLDs. The TLDs will store more resource records, because there are far more names within a TLD than TLDs.

2 pt (b)

Local domain name servers resolve queries for many hosts. Because of the recursive querying, they get to see the results, which they can cache, and use to (partly) resolve queries from other hosts

3 pt (c)

Each time you hear an ARP query from one of the two IP addresses (let's say A), asking for the MAC address belonging to the other IP address (let's say B), you either

- after a small delay, reply with an ARP response to this query with *your* MAC address, so that the receiver believes you have the requested IP address (B), or
- broadcast another ARP query, having your MAC and IP address B as source, asking for the MAC address belonging to IP address A.

In both cases, node A will add an entry to its ARP table, associating your MAC address to the IP address of B, and will as a result send packets destined to B to you. You should do the same for queries from B for the MAC address belonging to IP address A.

End of this exam.