Network Systems (201300179), Test 3

March 21, 2014, 15:45–17:15

- This is an open-book exam: you are allowed to use the book by Peterson & Davie and the reader that belongs to this module. Furthermore, use of a dictionary is allowed. Use of a simple (non-graphical) calculator is allowed.
- Other written materials, and laptops, tablets, graphical calculators, mobile phones, etc., are not allowed. ***Please remove any such material and equipment from your desk, now!***
- Although the questions are stated in English, you may answer in English or Dutch, whichever you are more comfortable with.
- You should always explain or motivate your answers, with so much detail that the grader can judge whether you understand the material; so just saying "yes" or giving a formula without explanation is not enough.
- Visiting the toilet without explicit permission of the supervisor is not allowed. During the last 30 minutes of the exam, no toilet visits are allowed.

---

1. Distance-vector routing

Consider the following network, where the nodes represent routers, and the labelled links represent links between routers with their associated link costs. We assume the use a distance vector routing algorithm without split horizon or poisoned reverse.

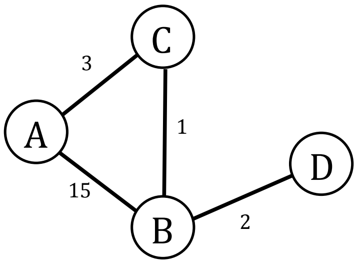

In the following, you will be asked to give the distances stored at each node in the form of a table. This table has the following form (this is also the format used in Table 3.10 in Peterson & Davie):

Information stored at node	Distance to reach node			
	A	B	C	D
A				
B				
C				
D				

3 pt (a) Give, in table form, the distances stored when each router only knows the distance to its immediate neighbours.

3 pt (b) Give, in table form, the distances stored when each router has reported the information from the previous subquestion to its immediate neighbours, i.e., after one iteration of the distance vector algorithm

Let us now assume that the distance vector algorithm has converged, and all nodes know the shortest path and distance to all other nodes.

1 pt (c) Give, in table form, the distances stored at this moment.

Now, the link between B and C disappears (breaks), which both B and C get to know.

3 pt (d) Give, in table form, the distances stored after B and C have recalculated[1] their routing table, but before they have sent any updates to their neighbours.

[1] you may assume they have kept a copy of the last distance vectors they received from their neighbours

2 pt (e) How many updates will node B send to node D before node D learns the real cost to node C? Explain your answer.

Continued on next page...

2. Addressing issues

IPv6 addresses are 128 bit long, sufficient for 2^{128} computers, but just about every computer also has a unique MAC address, and there are only 2^{48} of those, so we cannot have more than 2^{48} computers.

2 pt (a) Why does it still make sense to have much more than 2^{48} IPv6 addresses if we never expect more than 2^{48} computers to be connected?

3 pt (b) Which of the following are correct resp. incorrect notations for IPv6 addresses? Why? Which of these addresses is/are in the `2001:0600::/23` range?

```
2001:0600::1
2001::0600:1
2001::0600::1
2001:0700::abcd:1
::1
```

3 pt (c) Consider an organization which has a /20 block of IPv4 addresses. Using the Host-Density ratio, argue how many computers the network of this organization could comfortably contain.

3 pt (d) Suppose an internet service provider has run out of IPv4 addresses and decides to use NAT (Network Address Translation) to solve this problem. Can an unlimited number of computers share a single IPv4 address using NAT, or are there practical limitations? If you do find one or more limitations, discuss it/them both qualitatively (i.e., what causes the limitation) and quantitatively (how many computers).

Continued on next page...

3. TCP

Figure 5.7 of Peterson & Davie displays the state-transition diagram of TCP. For your convenience, this figure is reprinted below.

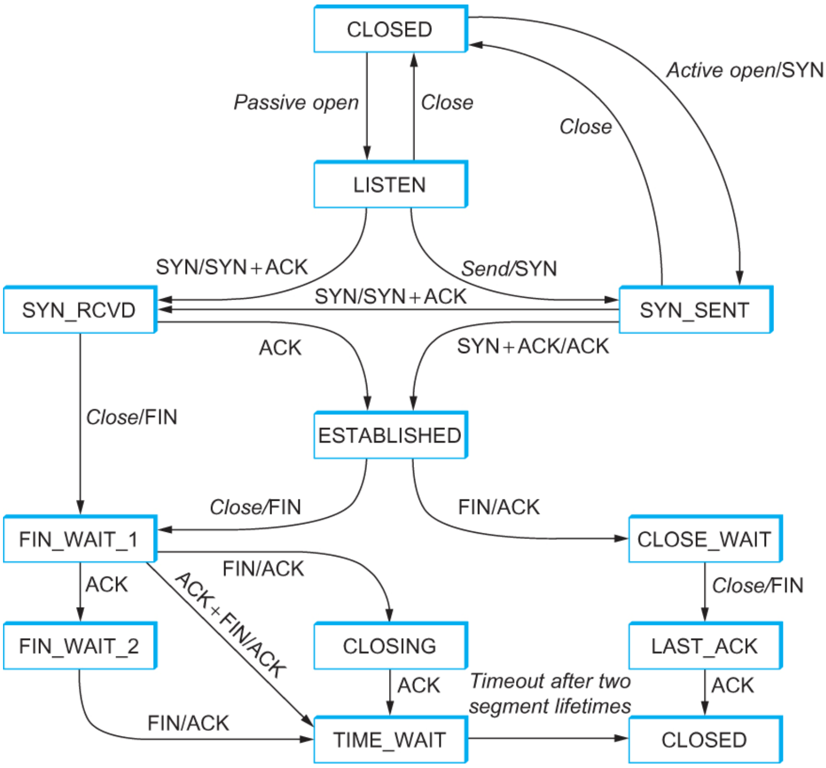

3 pt (a) Give a time-sequence diagram (similar to Figure 5.6 of Peterson Davie) that shows how a connection can be closed in such a way that the local side goes from the state `ESTABLISHED` to the state `CLOSED` using the diagonal transition labelled `ACK+FIN/ACK`.

Note: draw a time-sequence diagram where the local side is depicted on the left, where the messages transferred are labelled with the type of message (or the flag set). Denote on both on the outer side of both time-lines all states the local (left) or remote (right) side is going through.

3 pt (b) In which of the states in the figure above can a node receive new user data and pass that data to the application. Explain your answer.

3 pt (c) In its acknowledgements to the sender, the TCP receiver sets the `AdvertisedWindow` to the value

$$\texttt{AdvertisedWindow} = \texttt{MaxRcvBuffer} - ((\texttt{NextByteExpected} - 1) - \texttt{LastByteRead})$$

This is basically the difference between the capacity of the buffer (`MaxRcvBuffer`) and the bufferspace already filled with data. However, some out-of-sequence data may also be present in the buffer (in that case `NexByteExpected` - 1 does not equal `LastByteRcvd`). Why does that not play a role in the setting of `AdvertizedWindow` (i.e., why can `LastbyteRcvd` not be found in the formula above)?

End of this exam.