Network Systems (201600146/201600197), Test 3

March 23, 2018, 13:45–15:15

- 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!***
- 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.
- Write your answers to open questions on this paper, in the provided boxes, and hand this in.
- Questions marked with (MC) must be answered on the separate multiple-choice form, at the number indicated in the circle.
  Since the multiple-choice form will not be available at the exam review session, we recommend to *also* mark the MC answers on this paper.
- Total number of pages: 6.
- Total number of points: 29.

Your name:

(please underline your family name (i.e., the name on your student card), so that we know how to sort)

Your student number:

1. Addressing and routing

3 pt (a) For each of the following IPv6 address notations, choose one of the following answers:

A. This is not a valid IPv6 address.
B. This is a valid IPv6 address, but it does not belong to the `0100:0:8000::/33` subnet.
C. This is a valid IPv6 address, and it belongs to the `0100:0:8000::/33` subnet.

MC**01** `::1`

MC**02** `100::0:1::0`

MC**03** `100::8000:1:0`

MC**04** `0100:8000::1:0`

MC**05** `0100:80000:1:0`

MC**06** `100:0:face::1:0`

MC**07** `0100:0:1337::1:0`

MC**08** `100::8000:1:2:3:4:5`

MC**09** `01::8000:1:2:3:4:5`

1 pt (b) Can the HD ratio in an operational network be above 1 (i.e., above 100%)?

MC10

A. No, the upper limit for the HD ratio is about 87%.
B. Yes, the HD ratio can exceed 1 because it is a ratio of logarithms.
C. Yes, that means there is a very urgent need to expand the address space.
D. No, that would mean there are more addresses assigned than totally available.

1 pt (c) Does it make sense to apply the HD ratio to 48-bit MAC addresses?

MC11

A. No, because they are not globally unique.
B. No, because they are not used for routing.
C. Yes, and it would show that there is no shortage of them.
D. No, because they are not assigned hierarchically to networks.
E. No, because only routers have scalability problems, bridges don't.

1 pt (d) Since the first 4 bits of both the IPv4 and IPv6 header are the version number, someone suggests the following way of gradually transitioning from IPv4 to IPv6: when you send an IPv6 packet, just put '4' rather than '6' in that version field, so that old routers and hosts, that do no yet know about IPv6, can just treat the packet as an IPv4 packet. Will this work?

MC12

A. No, this cannot work at all.
B. This works fine for old *hosts*, but not for old *routers*.
C. This works fine for old *routers*, but not for old *hosts*.
D. This works fine for both old *routers* and old *hosts*.

1 pt (e) Is multicast easier to do with IPv6 than with IPv4?

MC13

A. Yes, since IPv6 has more addresses available.
B. No, since the maximum packet length was not increased in IPv6.
C. No, just like in IPv4, the difficulty of finding good routes remains.
D. Yes, since the IPv6 header has more space for the destination address.

---

2. 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 they use a distance vector routing algorithm, with all nodes doing their updates simultaneously; for now, no split horizon or poisoned reverse is used.

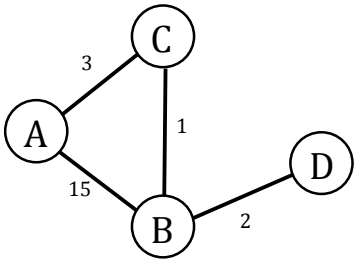

Assume that initially, the nodes know themselves and their immediate neighbours.

1 pt (a) What is the initial distance vector sent by node A to node C?

1 pt (b) What is the distance vector sent by node A to node C in the next iteration?

1 pt (c) What is the distance vector sent by node A to node C when the network has converged?

1 pt (d) Same question, but assuming split horizon is used.

Next, assume the link between A and C breaks.

2 pt (e) Assuming no split horizon / poison reverse is used, how many iterations will it take until the network has converged again?

1 pt (f) Now suppose that split horizon was used. Would this speed convergence up compared to your answer in the previous question?

MC14

A. No, since there is a loop in this network.
B. No, speeding up convergence requires poison reverse.
C. Yes, since D would know about the correct path to A via B.
D. Yes, since C would not initially think it can reach A via B at a low cost.
E. No, since B has not one but two links over which counting to infinity could occur.

1 pt (g) We know that distance vector may converge slowly if a link breaks. Can this problem also occur if a new link is added, e.g., a link with cost 1 between node C and D?

MC15

A. No, since adding the link can only lower the cost of best paths.
B. No, since the distance vector algorithm would simply not use this link.
C. Yes, just like in the case of link failures, outdated information may still bounce around in the network.
D. Yes, in fact the problem will become worse as there are more links on which counting-to-infinity can occur.

1 pt (h) Which feature of AODV makes it so suitable for ad-hoc networks with much mobility?

MC16

A. It uses path vectors.
B. It uses distance vectors.
C. It only looks for routes when needed.
D. The RREQ and RREP packets are short.

1 pt (i) If AODV is so good, why don't we use it in the fixed, wired Internet as well?

MC17

A. AODV only works if nodes are mobile.
B. AODV would not scale because of the RREQ flooding.
C. AODV would not scale because of storage requirements.
D. AODV can only work with MAC addresses, not with IP addresses.

You may want to use this diagram, copied from the book, for problem 3:

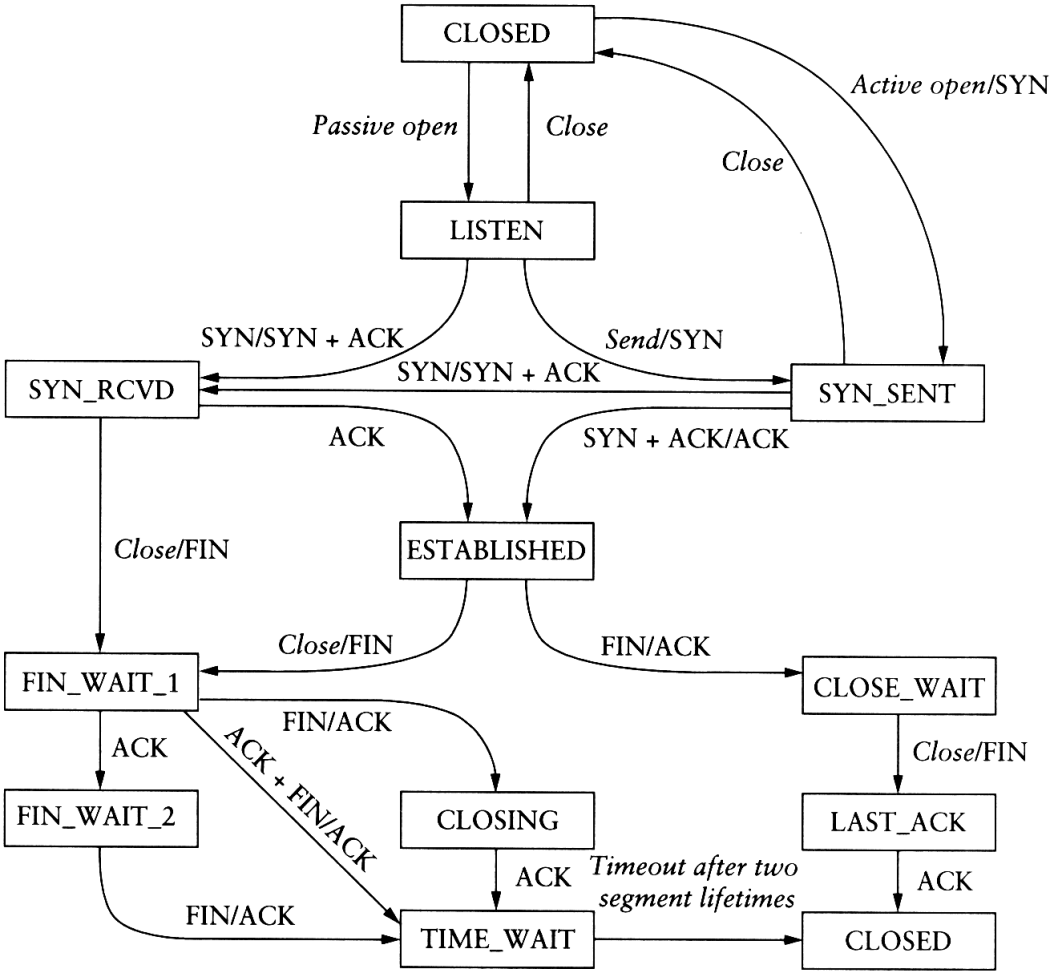

Continued on next page...

3. Transport layer protocols

1 pt (a) NTP is the Network Time Protocol, by which hosts can synchronize their clocks, basically by asking another host what it thinks te current time is. Is TCP or UDP more suitable as a transport-layer protocol for NTP?

MC18

A. TCP, because the messages are short.
B. UDP, because the messages are short.
C. TCP, to avoid timing inaccuracy due to retransmission.
D. UDP, to avoid timing inaccuracy due to retransmission.
E. TCP, to ensure that these important packets are not lost.
F. UDP, to ensure that these important packets are not lost.
G. Either, but only if it has the timestamp header option.

1 pt (b) What is a "well-known" port number?

MC19

A. An agreed-upon port number on which a client listens.
B. An agreed-upon port number on which a server listens.
C. An agreed-upon port number from which a client connects.
D. An agreed-upon port number from which a server connects.
E. A port number memorized by the network administrator.
F. A port number chosen by the network administrator.
G. A port number that is easy to remember.

5 pt (c) Fill in the dots in this diagram: sequence/acknowledgement numbers and flags of the packets, and TCP state names along the left and right side of the diagram.

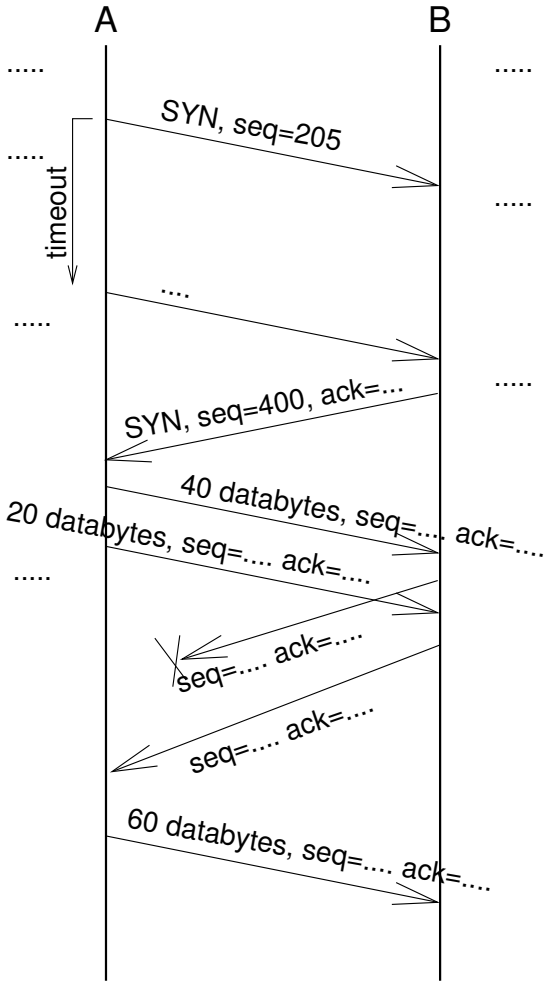

1 pt (d) Suppose TCP has most recently sent a segment with ACK=500, and it then receives a segment with SEQ=300 containing 50 bytes of data. What should it do?

MC20

A. Send a segment with ACK=301.
B. Send a segment with ACK=350.
C. Send a segment with ACK=500.
D. Send a segment with ACK=551.
E. Send a segment with PSH=1.
F. Send a segment with RST=1.
G. Send nothing.

1 pt (e) Consider the text "*Close*/FIN" at the arrow from `ESTABLISHED` to `FIN_WAIT_1` in the TCP state transition diagram (see page 4). What does this text mean?

MC21

A. When we receive a Close, we send a FIN.
B. When we receive a FIN, we send a Close.
C. When the application indicates it want to close, we send a FIN.
D. When we receive a FIN, we tell the application the connection is to be closed.
E. When we receive either a Close or a FIN, we go from `ESTABLISHED` to `FIN_WAIT_1`.
F. When we receive both a Close and a FIN, we go from `ESTABLISHED` to `FIN_WAIT_1`.
G. When we go from `ESTABLISHED` to `FIN_WAIT_1`, we can either close the connection or send a FIN.

1 pt (f) Suppose our side of a TCP connection is in the state `FIN_WAIT_2`, while the other side is in the state `ESTABLISHED`. Which statement is true?

MC22

A. Neither side can send data anymore.
B. The other side can still send data, our side can't.
C. Our side can still send data, the other side can't.
D. Both sides can still send data.
E. This is impossible; if we are in `FIN_WAIT_2`, the other side must be in `CLOSE_WAIT`.
F. This is impossible; if we are in `FIN_WAIT_2`, the other side must be in `LAST_ACK`.
G. This is impossible; if we are in `FIN_WAIT_2`, the other side must be in `CLOSING`.

1 pt (g) What problem does the Window Scaling extension solve?

MC23

A. Lack of bits in window field.
B. Lack of bits in sequence numbers.
C. Uniqueness of sequence numbers.
D. Mismatch of send and receive window.
E. TCP scalability to large numbers of hosts.

1 pt (h) In what situation would the use of the PAWS (Protection Against Wrapped Sequence numbers) extension be beneficial even if the Window Scaling extension is not used?

MC24

A. Never.
B. On a connection with small bandwidth and a small round-trip time.
C. On a connection with a large bandwidth but small round-trip time.
D. On a connection with a small bandwidth but large round-trip time.
E. On a connection with a large bandwidth and a large round-trip time.

End of this exam.