

Network Systems (201600146/201600197), Test 2

April 9, 2020, 13:45–16:45

(Taken together with test 3, online at home because of the corona pandemic
— and eventually declared invalid because of large-scale fraud)

1. Switching and IPv4

1 pt

Q 1.1 Consider a virtual circuit network. What does the number of distinct VCIs required in such a network depend on?


- A. the number of nodes in the network
- B. the number of links in the network
- C. the total number of VCs in the network
- D. the number of VCs on the busiest link in the network
- E. the length of the longest path in the network

1 pt

Q 1.2 It is said that in a virtual circuit network, a routing algorithm is still needed. Why is that?

- A. to fill the VC tables of the switches
- B. to be able to send the VC connection request to the correct next-hop switch
- C. to determine the correct path from incoming interface to outgoing interface inside a switch
- D. for packets for which no outgoing VC is known
- E. to route between the different VCs on a link

Please consider the following network configuration. Host A is connected to the Internet, through some node X of which we do not know the functionality. We would like to deduce the functionality of node X by inspecting packets that traverse node X from A to the Internet, i.e., by looking at the same packet (with possibly somewhat modified headers) on point 1 and point 2.



1 pt

Q 1.3 Suppose the source MAC address of the packet we inspect is different at point 1 and point 2. What can we deduce?

- A. node X is a switch.
- B. node X is a router.
- C. node X is a NAT
- D. node X is a switch or a router
- E. node X is a switch or a NAT
- F. node X is a router or a NAT
- G. nothing (none of the above)

1 pt

Q 1.4 Suppose the source IP address of the packet we inspect is different at point 1 and point 2. What can we deduce?


- A. node X is a switch.
- B. node X is a router.
- C. node X is a NAT
- D. node X is a switch or a router
- E. node X is a switch or a NAT
- F. node X is a router or a NAT
- G. nothing (none of the above)

1 pt

Q 1.5 Suppose the destination TCP port number of the packet we inspect is the same at point 1 and point 2. What can we deduce?

- A. node X is a switch.
- B. node X is a router.
- C. node X is a NAT
- D. node X is a switch or a router
- E. node X is a switch or a NAT
- F. node X is a router or a NAT
- G. nothing (none of the above)

Please consider the network of switches below.



2 pt

Q 1.6 After running the spanning tree algorithm in this network, which links are not used anymore for forwarding packets on?

- A. S1 - S4
- B. S1 - S6
- C. S2 - S3
- D. S2 - S4
- E. S3 - S5
- F. S4 - S5
- G. S4 - S6
- H. S4 - S7
- I. S5 - S7

2 pt

Q 1.7 For some reason, the switch S3, depicted in the figure above has accidentally also been assigned identity S1 (so, there are now 2 switches S1). After running the spanning tree algorithm in this network, which links are not used anymore for forwarding packets on?

- A. S1 - S4
- B. S1 - S6
- C. S2 - S3 (now also numbered S1)
- D. S2 - S4
- E. S3 (now also numbered S1) - S5
- F. S4 - S5
- G. S4 - S6
- H. S4 - S7
- I. S5 - S7

Suppose a router is applying longest prefix matching and has the following table:

Prefix / Length	Next Hop
130.89.130.144/28	Interface 0
130.89.130.144/30	Interface 1
130.89.130.160/27	R2
130.89.130.0/24	R3
129.98.96.0/20	R4
0.0.0.0/0	R5

For each of the following addresses indicate to which Next Hop a packet with this address will be forwarded, or select 'drop' if the packet should be dropped instead of forwarded.

1 pt **Q 1.8** 130.89.130.144

- A. Interface 0
- B. Interface 1
- C. R2
- D. R3
- E. R4
- F. R5
- G. drop

1 pt **Q 1.9** 130.89.130.191

- A. Interface 0
- B. Interface 1
- C. R2
- D. R3
- E. R4
- F. R5
- G. drop


1 pt **Q 1.10** 129.98.115.33

- A. Interface 0
- B. Interface 1
- C. R2
- D. R3
- E. R4
- F. R5
- G. drop

2. Routing

Link state routing

Please consider the following network, in which all routers (vertices) are labelled by a capital letter, and all (bi-directional) links (edges) between routers are labelled with the cost of that link. Assume that the network is running a link-state routing protocol, using Dijkstra's shortest path algorithm for route calculation.



Calculate the optimal routes from the point of view of node A using the Dijkstra algorithm. We recommend you do the algorithm on paper first.

Below, we ask you to fill in the "confirmed" and the "tentative" list after every step. Use the same format as in the book, so a sequence of (X,n,Y), where X is the destination, n the cost of the currently cheapest known path to X, and Y the next hop on that path.

We start with the confirmed list in the first step, which is clearly (A,0,-).

Q2.1—10
total 6 pt

- Q 2.1** In the next step, the first non-empty tentative list is calculated. Please give this list.
- Q 2.2** Subsequently, the confirmed list is updated. Please give this list (according to the numbering of steps in the book, this is step 3).
- Q 2.3** Give the next tentative list, after recalculating it (step 4 according to the book's numbering).
- Q 2.4** Give the updated confirmed list (step 5).
- Q 2.5** Give the recalculated tentative list (step 6).
- Q 2.6** Give the next confirmed list (step 7).
- Q 2.7** Give the next tentative list (step 8).
- Q 2.8** Give the next confirmed list (step 9).
- Q 2.9** Give the next tentative list (step 10).
- Q 2.10** Give the final confirmed list, from which the entries in the forwarding table can be derived.

Distance Vector routing

We consider a distance vector routing algorithm without split horizon or poison reverse. We consider a node A, which has links to two other nodes: a link to node B with link cost 2, and a link to node C with cost 1. Let us now assume that node A receives a packet from B containing the following list of costs: (B, 0), indicating that B currently only knows a path to itself at cost 0. Likewise, A will receive from C: (C, 0).

1 pt

Q 2.11 Which is the list of costs that A will send back to B and C? Please use as notation for this list of costs a list of (X,c), where X is the name of the node and c is the cost to X just like it is done in the rest of this question.

Subsequently, A receives a new list of costs from B: (B,0)(A,2)(C,100)(D,5). Furthermore, it receives from C: (C,0)(A,1)(B,100)(D,200).

1 pt

Q 2.12 Which is the list of costs that A will now send to B and C?

1 pt

Q 2.13 Which list of costs will B send out to A (and its other neighbours), after it has received A's list from the previous question?

Now, let us suppose that A detects that its link to C has disappeared. A sends a message to B with a list of costs reflecting this change.

1 pt

Q 2.14 How many packets with lists of cost (each resulting in an update from A to B) will B send to

A **before** the real distance from B to C is given in B's list of costs?

1 pt **Q 2.15** Please explain how you got to that answer:

1 pt **Q 2.16** Why does the BGP protocol include path information in its routing messages? (mark all that apply)

- A. to avoid count-to-infinity
- B. to allow administrators to control which authoritative systems their packets are going through
- C. to use for source routing
- D. to use in the OSPF (Open Shortest Path First) protocol
- E. for routing acks to find their way back to the source

3. IPv6 and addressing

Check the following IPv6 addresses (or lookalikes) for correctness and if correct, for being in subnets.

0.5 pt **Q 3.1** ::5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet ::0/32 and also not in ::0/128
- C. valid IPv6 address, in the subnet ::0/32 but not in ::0/128
- D. valid IPv6 address, in the subnet ::0/128 but not in ::0/32
- E. valid IPv6 address, both in the subnet ::0/32 and in ::0/128

0.5 pt **Q 3.2** 0001:0001::5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet ::0/32 and also not in ::0/128
- C. valid IPv6 address, in the subnet ::0/32 but not in ::0/128
- D. valid IPv6 address, in the subnet ::0/128 but not in ::0/32
- E. valid IPv6 address, both in the subnet ::0/32 and in ::0/128

0.5 pt **Q 3.3** 0100::0200:5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet 0100::0/32 and also not in 0100::/128
- C. valid IPv6 address, in the subnet 0100::0/32 but not in 0100::/128
- D. valid IPv6 address, in the subnet 0100::0/128 but not in 0100::/32
- E. valid IPv6 address, both in the subnet 0100::0/32 and in 0100::/128

0.5 pt **Q 3.4** 0100::0200::5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet 0100::0/32 and also not in 0100::/128
- C. valid IPv6 address, in the subnet 0100::0/32 but not in 0100::/128
- D. valid IPv6 address, in the subnet 0100::0/128 but not in 0100::/32
- E. valid IPv6 address, both in the subnet 0100::0/32 and in 0100::/128

0.5 pt **Q 3.5** 987f:0200::5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet 987f::0/18 and also not in 987f::/22
- C. valid IPv6 address, in the subnet 987f::0/18 but not in 987f::/22
- D. valid IPv6 address, in the subnet 987f::0/22 but not in 987f::/18
- E. valid IPv6 address, both in the subnet 987f::0/18 and in 987f::/22

0.5 pt **Q 3.6** 987f:0200:1:2:3:4:5005

- A. not a valid IPv6 address
- B. valid IPv6 address, not in the subnet 987f::0/18 and also not in 987f::/22
- C. valid IPv6 address, in the subnet 987f::0/18 but not in 987f::/22

- D. valid IPv6 address, in the subnet 987f::0/22 but not in 987f::/18
- E. valid IPv6 address, both in the subnet 987f::0/18 and in 987f::/22

Q 3.7 Consider a network using 24 bit addresses, in which 1 million hosts are allocated. Calculate the HD ratio. Outcome:

Q 3.8 Give your calculation:

1 pt **Q 3.9** What does it mean when the HD ratio in a network is 0.50 ?

- A. half of all addresses are in use.
- B. half of the address bits are in the subnet mask.
- C. half of the subnets have at least one computer on them.
- D. half of the addresses are allocated to hosts, the other half to routers.
- E. half as many hosts could be accommodated if structured address assignment was used.
- F. half as many address bits would suffice, if the addresses wouldn't need structured assignment.

1 pt **Q 3.10** How is the compatibility between IPv4 and IPv6 ?

- A. IPv4 clients can connect to IPv6 servers, and IPv6 clients can connect to IPv4 servers.
- B. IPv4 clients can connect to IPv6 servers, but IPv6 clients cannot connect to IPv4 servers.
- C. IPv4 clients cannot connect to IPv6 servers, but IPv6 clients can connect to IPv4 servers.
- D. IPv4 clients cannot connect to IPv6 servers, and IPv6 clients cannot connect to IPv4 servers.

Suppose an IPv6 packet containing a TCP segment gets fragmented into two fragments. What would the headers in the fragments be?

1 pt **Q 3.11** First fragment:

- A. First the IPv6 header, then the fragmentation header, then the TCP header.
- B. First the IPv6 header, then the TCP header, then the fragmentation header.
- C. First the IPv6 header, then the TCP header (and no fragmentation header).
- D. First the IPv6 header, then the fragmentation header (and no TCP header).
- E. Only the fragmentation header (and no IPv6 nor TCP header).
- F. Only the IPv6 header (and no fragmentation nor TCP header).

1 pt **Q 3.12** Second fragment:

- A. First the IPv6 header, then the fragmentation header, then the TCP header.
- B. First the IPv6 header, then the TCP header, then the fragmentation header.
- C. First the IPv6 header, then the TCP header (and no fragmentation header).
- D. First the IPv6 header, then the fragmentation header (and no TCP header).
- E. Only the fragmentation header (and no IPv6 nor TCP header).
- F. Only the IPv6 header (and no fragmentation nor TCP header).

4. Transport layer protocols

Suppose a web browser has fetched an HTML file from <http://www.utwente.nl/> using HTTP/1.0 (i.e., no persistent connections); the TCP connection used for this we'll call the 'first' request. Next, it fetches a picture file from that same HTTP server.

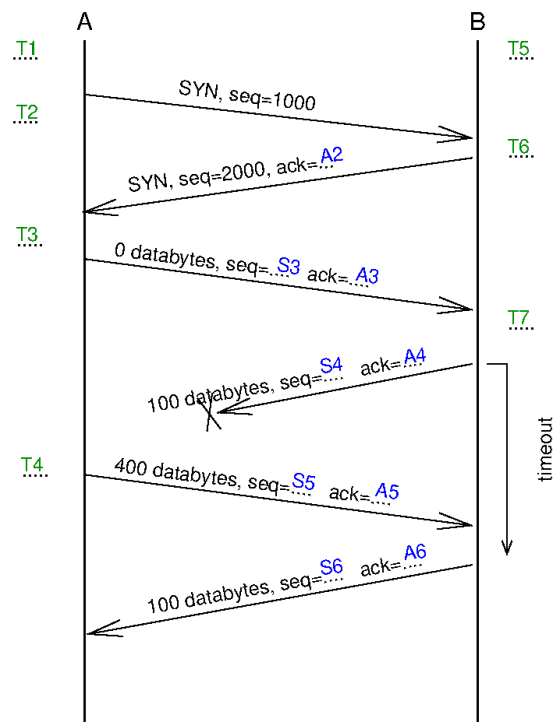
1 pt **Q 4.1** What will the port numbers in the *second* connection request be?

- A. source and destination both 80
- B. source = 80, destination is same as first request
- C. destination = 80, source is same as first request
- D. source = 80, destination is random and different from first request
- E. destination = 80, source is random and different from first request

1 pt

Q 4.2 In the response to that second connection request, what will the port numbers be?

- A. same as in the request
- B. interchanged compared to the request
- C. source = 80, destination = random
- D. destination = 80, source = random



5 pt

Q 4.3—18 The above shows a time-sequence diagram. Please fill in what should be written on the dots. T1...T7 are state names from the TCP state transition diagram. S3...S6 and A2...A6 are sequence and acknowledgement numbers. (Note: to avoid a lot of scrolling, it's probably easiest if you write them down on a piece of paper before filling them in in the input fields below.)

1 pt

Q 4.19 Why did Google choose to use UDP rather than TCP for the QUIC protocol?

- A. reliable delivery is not important for web browsing
- B. this decision was a mistake the Google engineers now regret
- C. retransmissions take too much time for real-time audio and video
- D. this is not correct, QUIC can also run on TCP rather than UDP, if desired
- E. so that fragments of different files that are being downloaded simultaneously don't need to wait for each other

1 pt

Q 4.20 There are several extensions to TCP for dealing with "long fat pipes". Why aren't there any such extensions for UDP?

- A. UDP is rarely used over long fat pipes
- B. UDP does not have sequence numbers
- C. long fat pipes just cause UDP to be even more unreliable (that's what the 'U' in UDP stands for anyway)
- D. the extensions designed for TCP can be used for UDP as well

1 pt

Q 4.21 When we studied the link layer, we learnt that the sequence number space need not be larger than sum of the send and receive window (SWS+RWS). However, the sequence number fields in the TCP header are much larger (32 bit) than the receive window field (16 bit). Why is this?

- A. The sequence numbers were made 32 bit in anticipation of the "window scaling" extension, which allows larger windows.
- B. TCP needs to deal with reordered packets, while at the link layer reordering does not happen.
- C. For compatibility with IPv4's 32-bit address length.
- D. TCP sequence numbers count bytes rather than packets.
- E. To compensate for the relatively weak checksum that TCP uses, compared to the CRC at the link layer.