Network Systems (201600146/201600197), Test 4

April 5, 2019, 13:45–15:15

- This is an open-book test: 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 test, 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 test review session, we recommend to *also* mark the MC answers on this paper.
- Total number of pages: 6.
- Total number of points: 28.

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:

Continued on next page...

1. Congestion control

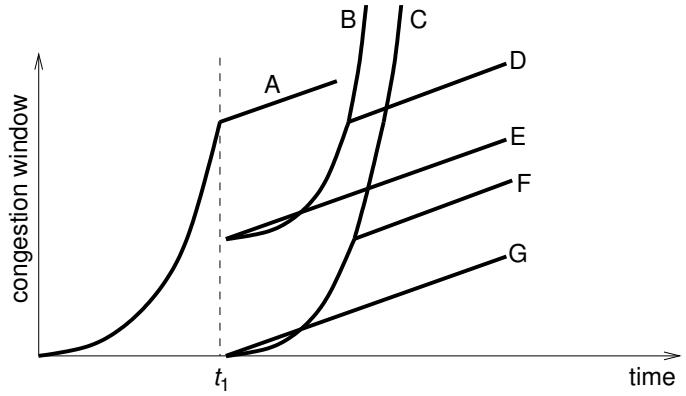

The above figure shows schematically several (im)possibilities of how the TCP congestion window (in standard TCP, i.e., New-Reno) evolves over time, after some event happens at the moment t_1, marked by a dashed line. (Note that some lines cross (e.g., C and D), and that some lines partially coincide (e.g., B and D).)
Which line represents the congestion window in each of the following cases?

1 pt (a) MC**01** At t_1 a packet loss was detected by timeout.

1 pt (b) MC**02** At t_1 a packet loss was detected by triple duplicate ack.

1 pt (c) MC**03** At t_1 the congestion window became equal to the slow-start threshold.

1 pt (d) How does the standard (a.k.a. New-Reno) TCP congestion control algorithm work on connections with a large bandwidth-delay product?

MC**04**

A. It works fine, since congestion is impossible on such links.
B. It works fine, it will automatically and quickly find the right congestion window to fully utilise the link.
C. It won't work well since such a link is by definition always congested.
D. It will have trouble quickly making the congestion window large enough to fully utilise the link.
E. It will make its congestion window too large and thus become unfair to other TCP connections using the same link.

1 pt (e) Suppose someone wants to claim a larger part of the bandwidth of a shared link by increasing the "additive increase" parameter of his TCP's congestion control implementation (while the competitors do not change their parameters). Will this work?

MC**05**

A. Yes, she will get a higher throughput, while the competitors' throughput does not change.
B. Yes, she will get a higher throughput, and his competitors get a lower throughput.
C. Yes, she will get almost all of the bandwidth of the link, while the competitors starve.
D. No, she will just get more packet loss.
E. No, because TCP uses AIMD which is always fair.
F. No, because routers will detect that she is cheating.
G. No, she would also need to modify the receiving side.

1 pt (f) Which statement about ECN (Explicit Congestion Notification) is true?

MC06

A. ECN uses UDP packets for congestion notifications.
B. ECN uses ICMP packets for congestion notifications.
C. ECN does not need any modifications in the endhosts.
D. ECN notifies endhosts of congestion by dropping a packet.
E. ECN notifies endhosts of congestion without dropping a packet.
F. ECN makes the network work as if the routers have infinite buffers.
G. ECN can only work if *all* endhosts support it, even the ones not causing congestion.

1 pt (g) What is the use of the "TCP-friendly" formula?

MC07

A. It tells TCP applications at what rate they should send at least.
B. It tells TCP applications precisely at what rate they should send.
C. It tells TCP applications at what rate they should send at most.
D. It tells UDP applications at what rate they should send at least.
E. It tells UDP applications precisely at what rate they should send.
F. It tells UDP applications at what rate they should send at most.
G. It tells how many TCP flows can share a link fairly.

1 pt (h) Consider an end host running TCP, transmitting data through a router over a link; assume it is the only (active) user of the link. It uses a delay-measurement based congestion control algorithm, for which it measures the round trip time at several sending window sizes:

window size	measured round trip time
1 packet	50 ms
2 packets	50 ms
10 packets	60 ms
20 packets	80 ms

Each packet is 1000 bits long.
Compute the speed (in bits/second) of the link; explain your answer.

1 pt (i) Continuing the previous question: compute how many of the packets are in the queue in the case of a window of 20 packets; explain your answer.

2. QoS

1 pt (a) Suppose we have movie encoded as a file of 3 gigabytes, and we want to transfer this file as quickly as possible. Is this an elastic or a non-elastic application?

MC08

A. Elastic, because the size of the file is given, not variable.
B. Elastic, because it can use as much bandwidth as is available.
C. Elastic, because the file contains audio and video information.
D. Non-elastic, because the size of the file is given, not variable.
E. Non-elastic, because it can use as much bandwidth as is available.
F. Non-elastic, because the file contains audio and video information.
G. That depends on whether the file is sent using TCP or UDP.

2 pt (b) Consider a source whose traffic is described by a token bucket with r=1000 tokens/s and B = 200 tokens, and counting 1 token per bit, and suppose this traffic is sent through a router which guarantees a transmission speed of at least 2000 bits/s.
Calculate the maximum delay the traffic can incur at this router; show your calculation.

1 pt (c) Consider the situation of the previous question, but assume there are now three such sources all feeding that single router. What is the maximum delay then?

MC09

A. 1/3 of the previous question.
B. Same as in the previous question.
C. 3 times that of the previous question.
D. 10 ms more than in the previous question.
E. 100 ms more than in the previous question.
F. 1000 ms more than in the previous question.
G. No maximum delay can be guaranteed in that situation.

1 pt (d) Which statement about Integrated Services (IntServ) and Differentiated Services (DiffServ) is true?

MC10

A. DiffServ is more secure than IntServ.
B. IntServ is for sending large files, while DiffServ for sending short messages.
C. DiffServ runs on datagram networks, while IntServ runs on virtual-circuit networks.
D. IntServ requires routers to know about each real-time flow individually, while DiffServ doesn't.
E. DiffServ uses DNS to look up the service class of a packet, while for IntServ this is stored in the IP header.
F. IntServ is based on the integral of the packet arrival rate, while DiffServ is based on the derivative of the packet arrival rate.
G. DiffServ gives different service to real-time flows than to non-real-time flows, while IntServ gives all flows the same service.

1 pt (e) Consider a queue with first-in, first-out scheduling. Which of the following is true?

MC11

A. A packet has to wait for packets which arrive after itself.
B. A packet has to wait for packets which arrive before itself.
C. Short packets incur less queueing delay than long packets.
D. Short packets incur more queueing delay than long packets.
E. The average delay is always less than in a queue with priority scheduling.
F. The average delay is always more than in a queue with priority scheduling.

1 pt (f) Why is bit-by-bit round-robin scheduling not used in practice?

MC12

A. It would give flows with shorter packets a smaller share of the bandwidth.
B. It would give flows with longer packets a smaller share of the bandwidth.
C. The receiver would not know which bits belong to which flow.
D. The receiver would be expensive because it needs two receive buffers.
E. It would require a special clock to calculate the packet transmission times.
F. Its calculations need making a sketch so can only be done by humans.

1 pt (g) Consider Fair Queueing scheduling for sharing a 100 Mbit/s link among two flows. One flow (let's call it red) sends at 90 Mbit/s, the other (blue) at 20 Mbit/s. What will happen?

MC13

A. Red gets 20 Mbit/s, blue gets 20 Mbit/s.
B. Red gets 50 Mbit/s, blue gets 50 Mbit/s.
C. Red gets 80 Mbit/s, blue gets 20 Mbit/s.
D. Red gets 90 Mbit/s, blue gets 10 Mbit/s.
E. Red gets 90 Mbit/s, blue gets 20 Mbit/s.
F. Depends on the packet lengths.
G. Depends on the packet arrival times.

1 pt (h) Same, but red sends at 60 Mbit/s, and blue at 30 Mbit/s. What will happen?

MC14

A. Red gets 30 Mbit/s, blue gets 30 Mbit/s.
B. Red gets 50 Mbit/s, blue gets 30 Mbit/s.
C. Red gets 50 Mbit/s, blue gets 50 Mbit/s.
D. Red gets 60 Mbit/s, blue gets 30 Mbit/s.
E. Red gets $66\frac{2}{3}$ Mbit/s, blue gets $33\frac{1}{3}$ Mbit/s.
F. Depends on the packet lengths.
G. Depends on the packet arrival times.

---

3. Security

1 pt (a) Suppose you want an IPsec-based VPN which allows you to connect your laptop from outside your home network, to your entire home network, so that all the devices in your home network are available to your laptop as if you were at home. You don't want other people listening in on the contents of your connections. Which of the following do you need?

MC15

A. IPsec AH in Transport mode.
B. IPsec AH in Tunnel mode.
C. IPsec ESP in Transport mode.
D. IPsec ESP in Tunnel mode.

1 pt (b) Suppose you're in a coffeeshop, which offers an open Wifi network. Unfortunately, your IPsec-based VPN is blocked on this network. However, you need to send a strictly confidential e-mail to a colleague, so only that colleague is able to open and read it. How do you do this?

MC16

A. You use WEP on the coffeeshop wireless network.
B. You log into your webmail on HTTPS.
C. You calculate a MAC and append it to your message.
D. You download a certificate from a CA and encrypt your message using that certificate.
E. You use the public PGP key of your colleague to encrypt the message.
F. You use the symmetric PGP key of your colleague to encrypt the message.

1 pt (c) Which of the following statements is true:

MC17

A. SSH runs over TLS.
B. TLS runs over SSH.
C. SSH was developed by Netscape.
D. TLS offers more integrity than SSH.
E. TLS always authenticates both the user and the server side of the connection.
F. SSH always authenticates both the user and the server side of the connection.
G. None of the above is true.

1 pt (d) Name the two modes of WPA2.

2 pt (e) Describe the differences between the two modes of WPA2.

2 pt (f) Suppose you are at work and you urgently need to transfer money. Your banking website is available over both HTTP and HTTPS. Your company firewall only allows outgoing connections over TCP ports 22 and 80. Can you still access the bank website securely, and if so, describe the steps and equipment necessary to accomplish this.

1 pt (g) What is a DDoS attack?

MC18

A. An attack in which DNS responses are intercepted and modified.
B. An attack on old computers running Disk Operating System software.
C. An attack in which packets with incorrect IP source addresses are sent.
D. An attack in which many machines overwhelm a single computer with traffic.
E. An attack in which a single computer causes other computers to become unavailable.

1 pt (h) Suppose you want to configure a firewall such that it blocks incoming packets to TCP port 1234, unless an outgoing connection with 1234 as the source port has been set up.

MC19

A. This can be done with either a stateless or a stateful firewall.
B. This requires a stateful firewall.
C. This requires a stateless firewall.
D. This cannot be done with either type of firewall.
E. This is meaningless as a firewall can only block outgoing packets.
F. This is meaningless as port 1234 is never used for incoming packets.
G. This is meaningless as outgoing connections don't have a source port number.

End of this test.