

Communication Networks

Prof. Laurent Vanbever


Exercise 10 – TCP and Congestion Control

10.1 TCP Warm-up (Exam Question 2019)

For the following questions answer either with true or false and give an explanation for your decision. Note that in the actual exam you only had to pick true, false or do not give an answer at all (as wrong answers result in point deductions).


- a) In contrast to the GBN protocol used in the project, TCP's sequence number often increases by more than one between two consecutive data packets.
- b) A client having an ongoing TCP connection to a server (IP 1.2.3.4, port 80) is not able to start a second TCP connection towards 1.2.3.4:80.
- c) A TCP packet with a value of 9 in its header length field (HdrLen, sometimes also called data offset) indicates 16 bytes of TCP options.
- d) Consider an ongoing TCP flow. Whenever the congestion window increases, the sender can transmit additional data segments.

For the following four questions, we consider the congestion window (CWND) evolution observed for a flow f and depicted in the figure below.



- e) Flow f was in the slow start phase exactly twice.
- f) Flow f experiences at least one packet loss between time A and B.
- g) In the future, f will never be able to experience a higher CWND size than B.
- h) Consider another flow f_2 starting at exactly the same time as f and traversing the exact same path, then f_2 would experience the exact same CWND evolution.

10.2 Fairness



Are you getting a fair share?

Consider the situation in which two hosts, A and B, are concurrently using a 1 Mbps link with a Maximum Segment Size (MSS) of 100 kb.


Assuming that B starts with 500 kbps and A with 200 kbps (see left picture). Describe the evolution of the throughput of the two hosts when:

- A and B rely on Additive Increase Multiplicative Decrease (AIMD).
- A and B rely on Multiplicative Increase Additive Decrease (MIAD).

Assume now that only A is malicious, and wants to cheat congestion control to get more throughput. Describe two distinct ways A could do so and what would be the net effect on B's throughput.

10.3 Congestion Window

Consider the following plot which depicts the evolution of the size of the TCP congestion window of the sender.



What kind of network conditions is this flow seeing?

Describe briefly:

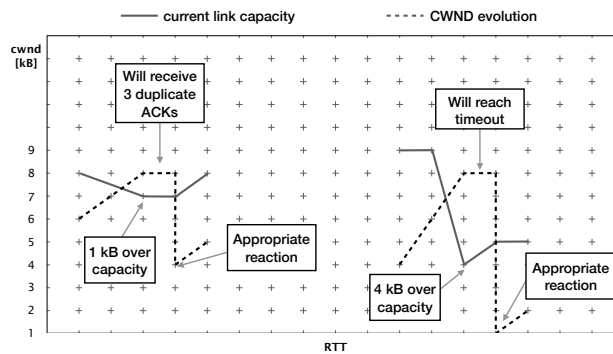
- What happens at point B?
- Does the event happening at point B require the network to discard packets? Why or why not?
- What happens at point E?
- Does the event happening at point E require the network to discard packets? Why or why not?

Consider that the Maximum Segment Size (MSS) of the connection is 1 kB and the Round-Trip Time (RTT) between the two end points is 100 milliseconds. The sender opens the connection at time $t = 0$. Transmission delay in this network is negligible, so you should only consider the propagation delay in the following.

- e) How much time has elapsed at point A?
- f) How much time has elapsed *between* point C and D?
- g) How much time has elapsed *between* point F and point G?

Briefly explain how come point D is higher than point B. Would you expect this to happen often?

10.4 Drawing practice (Exam Question 2018)



Reaction of the CWND (dashed line) if the current link capacity (continuous line) is exceeded by at most 2 kB (left) or by more than 2 kB (right).

In this task, you will draw the Congestion Window (CWND) evolution in reaction to the available capacity of a link in a network. The CWND follows the well-known TCP congestion control algorithm using slow-start. Whenever the CWND value exceeds the current link capacity, the CWND algorithm will react in the following way:

1. The current CWND value is kept for the entire next RTT (no increase or decrease);
- 2a. If the current link capacity was exceeded by at most 2 kBs, the CWND algorithm will observe three duplicate ACKs during the next RTT and will react appropriately (see the figure above on the left);
- 2b. If the current link capacity was exceeded by more than 2 kBs, the CWND algorithm will reach its timeout during the next RTT and will react appropriately (see the figure above on the right right).

Draw the CWND evolution directly into the figure below in reaction to the link capacity indicated with the continuous line. Start at the bottom left corner (RTT 1, CWND 1 kB) and assume that the CWND corresponds to a flow that just started, e.g. you are in the slow-start phase. You can stop once you reach RTT 22. To help you, a correct portion of the CWND is plotted between RTT 11 and 14 (dashed line).

