

## Assignment #3 – Transport Layer

Laurent Charignon, Avner May, Mathias Lécuyer (TAs) A. Chaintreau (instructor),

**How to read this assignment :** Exercise levels are indicated as follows

- ( $\rightarrow$ ) “elementary”: the answer is not strictly speaking obvious, but it fits in a single sentence, and it is an immediate application of results covered in the lectures.

*Use them as a checkpoint: it is strongly advised to go back to your notes if the answer to one of these questions does not come to you in a few minutes.*

- ( $\curvearrowright$ ) “intermediary”: The answer to this question is not an immediate translation of results covered in class, it can be deduced from them with a reasonable effort.

*Use them as a practice: how far are you from the answer? Do you still feel uncomfortable with some of the notions? which part could you complete quickly?*

- ( $\nrightarrow$ ) “tortuous”: this question either requires an advanced notion, a proof that is long or inventive, or it is still open.

*Use them as an inspiration: can you answer any of them? does it bring you to another problem that you can answer or study further? It is recommended to work on this question only AFTER you are done with the rest!*

**Exercise 1: True or false? (5 pt)**

Justify your answer and correct any false statement to the most precise answer.

1. ( $\rightarrow$ ) ”The 16-bit port numbers mean that you can only have 65,535 clients connect to the same web server at the same time.”
2. ( $\rightarrow$ ) “Suppose host A sends over a TCP connection to host B one segment with sequence number 38 and 4 bytes of data. In this same segment, the acknowledgement number is necessarily 42.”
3. ( $\rightarrow$ ) “The size of the TCP RcvWindow never changes throughout the duration of the session.”
4. ( $\rightarrow$ ) “There is no relationship between the variable LastByteRcvd in Section 3.5.5 and the variable y in Section 3.5.4?”
5. ( $\rightarrow$ ) “There is no known technique for a server to protect against a large number of malicious TCP connection requests.”

## Exercise 2: Are sequence numbers needed? (12 pt, including 1pt for question 4)

**Motivation:** Imagine you give your credit card number on the phone but your cell phone is noisy. It's annoying to give a sequence number to every digit. That's why many of us like to say "I'll repeat" in case something seemed wrong, to avoid rewriting the same number again. In this exercise, you will see if the same idea works between computers.

In order to detect duplicates while avoiding using sequence numbers, one may think of adding a different field in the header of any packet. Two alternatives are:

**retransmission flagging:** A unique bit is added to the header of the packet. This bit is a "0" if this is the first time this particular data is sent, or a "1" if this is not the first time it is sent.

**retransmission counting:** An integer value is added to the header of the packet, that indicates the number of previous transmissions of this data. This number will be "0" when this is the first time. Later, in future retransmissions it will be changed to "1", "2" etc.

You may conceptually think of it as a "sequence number" associated with all packets transmitting the same data. Note that we assume that the receiver, having no particular information on the packet it acknowledges, simply answer either with a "ACK" or a "NACK" at the reception of a packet.

1. ( $\rightarrow$ ) Prove using a simple counter example that none of these schemes can correctly operate if packets can get reordered.

For the following three questions, we assume that packets are not reordered. However, they may be delayed or loss. The sender implements a time-out to avoid deadlock.

2. ( $\curvearrowright$ ) Prove that **retransmission counting** does not correctly operate with packet losses. To do that, you will provide a counter-example that uses the minimum number of packets and packet loss.

For the following two questions, we assume that no packet loss occurs. However, packets can be corrupted. We assume that any corruption is detected using a checksum

3. ( $\curvearrowright$ ) Prove that **transmission flagging** does not correctly operate even in this case?
4. ( $\leftrightarrow$ ) Do you think that **retransmission counting** can correctly operate if no packets are lost?

## Exercise 3: Are ACK needed? (0 pt, not graded, only proposed as midterm practice)

In normal conversation, we generally only ask for clarification if something went wrong. Similarly, can we imagine operating a data reliable transmission protocol only using NACK?

1. ( $\rightarrow$ ) Using only 2 packets and no packet loss, prove that no scheme can operate using only NACK and no sequence number.
2. ( $\curvearrowright$ ) We now consider designing a scheme using sequence-number and NACK. When will the loss of the  $n$ -th packet be detected at the receiver? Can timeout protect against some losses?
3. ( $\curvearrowright$ ) In which context does a NACK-only protocol make the most sense? To help you, consider two applications: one is always on and send data frequently, such as a critical sensing operation. The other application is sporadic and asks for a burst of packets, such as multiple web sessions.

**Exercise 4: Additive Increase Multiplicative Decrease (8 pt)**

Refer to Figure 3.56 that in section 3.7.1 illustrating the convergence of TCP’s additive increase, multiplicative decrease algorithm. Suppose that instead of a multiplicative decrease, TCP decreased the window size “linearly”, that is by a number that does not depend on the current window size. We consider only two connections and make the same assumption as in this section.

1. ( $\curvearrowright$ ) Let us first assume that all connection decrease their window by the same amount and have the same RTTs. Would the resulting additive increase additive decrease converge to an equal share algorithm from any initial condition? Justify your answer using a diagram similar to Figure 3.56.
2. ( $\curvearrowright$ ) We now assume that, due to different RTTs etc., the effective throughput decreases linearly by a different amounts for the two connections. What will happen then? Does it depend on the initial condition


**Exercise 5: The TCP Throughput formula (10 pt)**

Consider the macroscopic description of TCP throughput. Note that, in the period of time in which the connections rate varies from  $\frac{W}{2RTT}$  to  $\frac{W}{RTT}$  only a single packet is lost (at the very end of the period).

1. ( $\curvearrowright$ ) Show that the loss rate (fraction of packets lost) is equal to  $L = \frac{1}{\frac{3}{8}W^2 + \frac{3}{4}W}$ .  
Hint: Consider the number of packets sent during a period.
2. ( $\curvearrowright$ ) Use the result above to show that if a connection has a small loss rate  $L$ , then its average rate is approximately given by  $\frac{1.22 \times MSS}{RTT \times \sqrt{L}}$ .

**Exercise 6: Max-min Fairness (5 pt)**

Consider the network above containing 6 flows ( $f_1$  through  $f_6$ ) passing through links with bandwidth capacities 10, 6, 6, and 2.



1. ( $\curvearrowright$ ) For each allocation below, indicate whether the allocation is max-min fair for the above network. If an allocation is not max-min fair, explain why it is not by showing that one flow’s rate can be increased by a small amount  $\delta$  without “unfairly” penalizing another flow.

Each allocation is written as  $(a_1, a_2, a_3, a_4, a_5, a_6)$ , where  $a_i$  is the bandwidth allocated to flow  $f_i$ .

- (a) (1, 1, 1, 1, 1, 1)
- (b) (4, 3, 3, 0, 1, 1)
- (c) (6, 2, 2, 2, 2, 0)
- (d) (6, 2, 2, 2, 1, 1)
- (e) (8, 1, 1, 4, 1, 1)