CSC358 Tutorial 5

Julian Sequeira and KyoKeun Park

February 14, 2020

University of Toronto Mississauga

Question 1: Concept Review

(a) What are the pros and cons between Go-Back-N and Selective Repeat?

(b) What's the purpose of "delayed ACK" and "triple duplicate ACK"?

(c) Why is a 2-way handshake not enough for establishing a TCP connection but a 3-way handshake is?

(d) What's the difference between flow control and congestion control?

For the answers, review the lectures, books, go to office hours, and use the discussion board!

Question 2: Estimating RTT and the Exponential Moving Average

Recall the TCP's formula for estimating RTT:

$$\texttt{EstimatedRTT} = (1-\alpha)\cdot\texttt{EstimatedRTT} + \alpha\cdot\texttt{SampleRTT}$$

In this question, suppose that $\alpha = 0.1$. Let S_1 be the most recent sample RTT, let S_2 be the next most recent sample RTT, and so on.

(a) Suppose that there have been four packets acknowledged with the RTTs being S_1, S_2, S_3, and S_4. Express $\texttt{EstimatedRTT}$ in terms of these four sample RTTs.

(b) Generalize your formula for n sample RTTs.

(c) Let n approach infinity. Comment on why this averaging procedure is called an exponential moving average.

Question 2: Estimating RTT and the Exponential Moving Average (a)

Let us denote $EstimatedRTT^{(n)}$ for the estimate after n_{th} sample.

$$\begin{aligned} EstimatedRTT^{(4)} = \alpha \cdot SampleRTT_1 + (1-\alpha) \cdot [\alpha \cdot SampleRTT_2 + \\ (1-\alpha) \cdot [\alpha \cdot SampleRTT_3 + (1-\alpha) \cdot SampleRTT_4]] \end{aligned}$$

Let us denote $EstimatedRTT^{(n)}$ for the estimate after n_{th} sample.

$$
\begin{aligned}
EstimatedRTT^{(4)} &= \alpha \cdot SampleRTT_1 + (1-\alpha) \cdot [\alpha \cdot SampleRTT_2 + \\
&\quad (1-\alpha) \cdot [\alpha \cdot SampleRTT_3 + (1-\alpha) \cdot SampleRTT_4]] \\
&= \alpha \cdot SampleRTT_1 + (1-\alpha) \cdot \alpha \cdot SampleRTT_2 + \\
&\quad (1-\alpha)^2 \cdot \alpha \cdot SampleRTT_3 + (1-\alpha)^3 \cdot SampleRTT_4
\end{aligned}
$$

Question 2: Estimating RTT and the Exponential Moving Average (b)

In order to get generalized formula, we can take our answer from part (a) and examine the pattern.

Question 2: Estimating RTT and the Exponential Moving Average (b)

In order to get generalized formula, we can take our answer from part (a) and examine the pattern.

$$EstimatedRTT^{(n)} = \alpha \cdot \sum_{j=1}^{n-1} (1-\alpha)^{j-1} \cdot SampleRTT_j + (1-\alpha)^{n-1} \cdot SampleRTT_n$$

Definition of exponential moving average: giving higher weight to more recent events

As $n \longrightarrow \infty$, our generalized formula from part (b) becomes:

Definition of exponential moving average: giving higher weight to more recent events

As $n \longrightarrow \infty$, our generalized formula from part (b) becomes:

$$\begin{aligned} \mathit{EstimatedRTT}^{(\infty)} &= \frac{\alpha}{1-\alpha} \cdot \sum_{j=1}^{\infty} (1-\alpha)^j \cdot \mathit{SampleRTT}_j \\ &= \frac{1}{9} \cdot \sum_{j=1}^{\infty} 0.9^j \cdot \mathit{SampleRTT}_j \end{aligned}$$

Recall that $\mathit{SampleRTT}_1$ is the *most recent RTT*.

Definition of exponential moving average: giving higher weight to more recent events

As $n \longrightarrow \infty$, our generalized formula from part (b) becomes:

$$\begin{aligned} EstimatedRTT^{(\infty)} &= \frac{\alpha}{1-\alpha} \cdot \sum_{j=1}^{\infty} (1-\alpha)^j \cdot SampleRTT_j \\ &= \frac{1}{9} \cdot \sum_{j=1}^{\infty} 0.9^j \cdot SampleRTT_j \end{aligned}$$

Recall that $SampleRTT_1$ is the *most recent RTT*.

We observe that the weight given to past samples decays exponentially!

Question 3: Fairness of Additive Increase Additive Decrease

Refer to the lecturer slide that illustrates the convergence to fairness of the AIMD algorithm (Page 66 of Week 5). Now suppose that instead of multiplicative decrease, TCP does additive decrease, i.e., it decreases the window size by a constant amount each time. Would the resulting AIAD algorithm still converge to fairness? Justify your answer using a diagram similar to the one in the lecture slide. More specifically, consider the following two cases:

(a) The two connections decrease by the same constant each time, i.e., they linearly decrease with the same slope.

(b) They two connections decrease by different constants each time, e.g., one connection's constant is twice of the other connection's constant.

Question 3: Fairness of Additive Increase Additive Decrease

Let us first consider the case where two connections both linearly increase and decrease with the same slope.

Question 3: Fairness of Additive Increase Additive Decrease

Let us first consider the case where two connections both linearly increase and decrease with the same slope.

Let's say connection 1 had more time to linearly increase its bandwidth, compared to connection 2 initially. When bandwidth becomes full, both connection 1 and 2 will linearly decrease its bandwidth until total bandwidth is free enough for them to linearly increase again.

What would this look like? Would it be possible for connection 1 and 2 to reach an equal bandwidth?

Question 3: Fairness of Additive Increase Additive Decrease

Let us first consider the case where two connections both linearly increase and decrease with the same slope.

Let's say connection 1 had more time to linearly increase its bandwidth, compared to connection 2 initially. When bandwidth becomes full, both connection 1 and 2 will linearly decrease its bandwidth until total bandwidth is free enough for them to linearly increase again.

What would this look like? Would it be possible for connection 1 and 2 to reach an equal bandwidth?

The bandwidth of connections will end up just oscillating back and fourth between each other. Hence, it will never converge to fairness! (cannot guarantee equal bandwidth share)

Question 3: Fairness of Additive Increase Additive Decrease

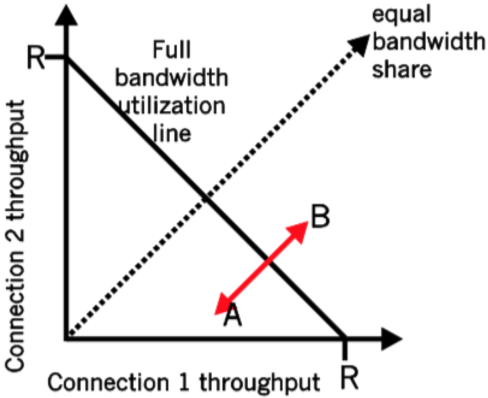

(a) linear increase, with equal linear decrease

Question 3: Fairness of Additive Increase Additive Decrease

Now, let us consider the case where one connection doubles its decrease rate relative to the other connection.

Think of a similar scenario as previous case, but instead, connection 1 decreases faster than connection 2.

Question 3: Fairness of Additive Increase Additive Decrease

Now, let us consider the case where one connection doubles its decrease rate relative to the other connection.

Think of a similar scenario as previous case, but instead, connection 1 decreases faster than connection 2.

This time, what would happen every time total bandwidth is full?

Question 3: Fairness of Additive Increase Additive Decrease

Now, let us consider the case where one connection doubles its decrease rate relative to the other connection.

Think of a similar scenario as previous case, but instead, connection 1 decreases faster than connection 2.

This time, what would happen every time total bandwidth is full?

Since connection 1 drops twice as much bandwidth as connection 2 when full, connection 2 will start “hogging” more bandwidth overtime.

Meaning, this method will converge to allocating entire bandwidth to connection 2!

Question 3: Fairness of Additive Increase Additive Decrease

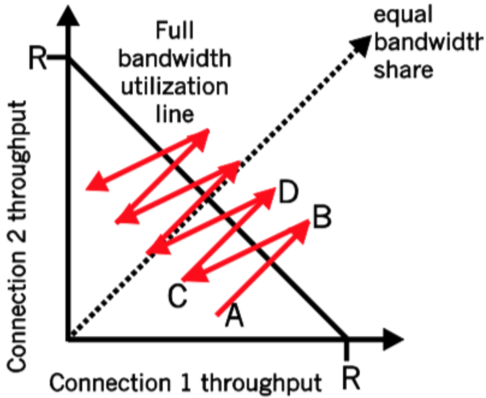

(b) linear increase, connection 1 decrease is twice that of connection 2

Question 4: Average Throughput of TCP

Recall the macroscopic description of TCP throughput. In the period of time from when the connection's rate varies from $W \cdot \frac{MSS}{2 \cdot RTT}$ to $W \cdot \frac{MSS}{RTT}$, only one packet is lost (right before the decrease).

(a) Show that the loss rate (fraction of packets lost) is equal to

$$L = \frac{1}{(3/8) \cdot W^2 + (3/4) \cdot W}$$

(b) Use the result above to show that if a connection has loss rate L, then its average throughput is approximately given by

$$\text{Average throughput} \approx \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}}$$

You may assume that W is very large so that $W^2 >> W$.

(c) Assuming, realistically, an MSS of 1500 bytes and a RTT of 100 milliseconds, in order to achieve a throughput of 10 Mbps, what's the requirement on the loss rate L? How about achieving a 10 Gbps throughput? Discuss the potential issues of the current version of TCP.

Recall: Loss rate (L) is the ratio of the number of packets lost over the number of packets sent.

Since we are said to be losing one packet every cycle, by definition of loss rate, the numerator makes sense. Now let's consider the number of packets sent in a cycle (denominator):

$$\frac{W}{2} + (\frac{W}{2} + 1) + ... + W = \sum_{n=0}^{W/2}(\frac{W}{2} + n)$$

Question 4: Average Throughput of TCP (a)

Recall: Loss rate (L) is the ratio of the number of packets lost over the number of packets sent.

Since we are said to be losing one packet every cycle, by definition of loss rate, the numerator makes sense. Now let's consider the number of packets sent in a cycle (denominator):

$$\frac{W}{2} + (\frac{W}{2} + 1) + ... + W = \sum_{n=0}^{W/2} (\frac{W}{2} + n)$$

$$= (\frac{W}{2} + 1) \cdot \frac{W}{2} + \sum_{n=0}^{W/2} n$$

Question 4: Average Throughput of TCP (a)

Recall: Loss rate (L) is the ratio of the number of packets lost over the number of packets sent.

Since we are said to be losing one packet every cycle, by definition of loss rate, the numerator makes sense. Now let's consider the number of packets sent in a cycle (denominator):

$$\begin{aligned}
\frac{W}{2} + (\frac{W}{2} + 1) + ... + W &= \sum_{n=0}^{W/2} (\frac{W}{2} + n) \\
&= (\frac{W}{2} + 1) \cdot \frac{W}{2} + \sum_{n=0}^{W/2} n \\
&= (\frac{W}{2} + 1) \cdot \frac{W}{2} + \frac{W/2 \cdot (W/2 + 1)}{2}
\end{aligned}$$

Question 4: Average Throughput of TCP (a)

Recall: Loss rate (L) is the ratio of the number of packets lost over the number of packets sent.

Since we are said to be losing one packet every cycle, by definition of loss rate, the numerator makes sense. Now let's consider the number of packets sent in a cycle (denominator):

$$
\begin{aligned}
\frac{W}{2} + (\frac{W}{2} + 1) + ... + W = \sum_{n=0}^{W/2}(\frac{W}{2} + n) \\
&= (\frac{W}{2} + 1) \cdot \frac{W}{2} + \sum_{n=0}^{W/2} n \\
&= (\frac{W}{2} + 1) \cdot \frac{W}{2} + \frac{W/2 \cdot (W/2 + 1)}{2} \\
&= \frac{W^2}{4} + \frac{W}{2} + \frac{W^2}{8} + \frac{W}{4} \\
&= \frac{3}{8} \cdot W^2 + \frac{3}{4} \cdot W
\end{aligned}
$$

Question 4: Average Throughput of TCP (b)

Since W is large, we can say that $\frac{3}{8} \cdot W^2 >> \frac{3}{4} \cdot W$. In other words, $\frac{3}{4} \cdot W$ starts to mean a lot less as W increases.

Since W is large, we can say that $\frac{3}{8} \cdot W^2 >> \frac{3}{4} \cdot W$. In other words, $\frac{3}{4} \cdot W$ starts to mean a lot less as W increases.

This means that $L \approx \frac{8}{3} \cdot W^2 \implies W \approx \sqrt{\frac{8}{3 \cdot L}}$. So, just with some substitution, we have:

Question 4: Average Throughput of TCP (b)

Since W is large, we can say that $\frac{3}{8} \cdot W^2 >> \frac{3}{4} \cdot W$. In other words, $\frac{3}{4} \cdot W$ starts to mean a lot less as W increases.

This means that $L \approx \frac{8}{3} \cdot W^2 \implies W \approx \sqrt{\frac{8}{3 \cdot L}}$. So, just with some substitution, we have:

$$\begin{aligned} \text{average throughput} &= \frac{3}{4} \cdot \sqrt{\frac{8}{3 \cdot L}} \cdot \frac{MSS}{RTT} \\ &= \frac{3}{4} \cdot \sqrt{\frac{8}{3}} \cdot \frac{MSS}{RTT \cdot \sqrt{L}} \end{aligned}$$

Where did $\frac{3}{4}$ come from?

Question 4: Average Throughput of TCP (b)

Since W is large, we can say that $\frac{3}{8} \cdot W^2 >> \frac{3}{4} \cdot W$. In other words, $\frac{3}{4} \cdot W$ starts to mean a lot less as W increases.

This means that $L \approx \frac{8}{3} \cdot W^2 \implies W \approx \sqrt{\frac{8}{3 \cdot L}}$. So, just with some substitution, we have:

$$\begin{aligned}
\text{average throughput} &= \frac{3}{4} \cdot \sqrt{\frac{8}{3 \cdot L}} \cdot \frac{MSS}{RTT} \\
&= \frac{3}{4} \cdot \sqrt{\frac{8}{3}} \cdot \frac{MSS}{RTT \cdot \sqrt{L}} \\
&= \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}}
\end{aligned}$$

We are now given $MSS = 1500\text{bytes}$ and $RTT = 100\text{ms}$. Let us first consider the throughput of 10 Mbps. Since we are given everything except L, we can simply substitute the variables and solve for L:

$$\text{throughput} = \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}}$$

We are now given $MSS = 1500\text{bytes}$ and $RTT = 100\text{ms}$. Let us first consider the throughput of 10 Mbps. Since we are given everything except L, we can simply substitute the variables and solve for L:

$$\text{throughput} = \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}}$$

$$\sqrt{L} = \frac{1.22 \cdot MSS}{RTT \cdot \text{throughput}}$$

We are now given $MSS = 1500\text{bytes}$ and $RTT = 100\text{ms}$. Let us first consider the throughput of 10 Mbps. Since we are given everything except L, we can simply substitute the variables and solve for L:

$$\begin{aligned}
\text{throughput} &= \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}} \\
\sqrt{L} &= \frac{1.22 \cdot MSS}{RTT \cdot \text{throughput}} \\
\sqrt{L} &= \frac{1.22 \cdot 12000 bits}{0.1\text{s} \cdot 10000000 bps} \quad \text{(Be careful with the units!)}
\end{aligned}$$

We are now given $MSS = 1500\text{bytes}$ and $RTT = 100\text{ms}$. Let us first consider the throughput of 10 Mbps. Since we are given everything except L, we can simply substitute the variables and solve for L:

$$\begin{aligned}
\text{throughput} &= \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}} \\
\sqrt{L} &= \frac{1.22 \cdot MSS}{RTT \cdot \text{throughput}} \\
\sqrt{L} &= \frac{1.22 \cdot 12000 bits}{0.1s \cdot 10000000 bps} \quad \text{(Be careful with the units!)} \\
L &\approx 2 \cdot 10^{-4}
\end{aligned}$$

Hence, to achieve 10 Mbps, the loss rate must be at most $2 \cdot 10^{-4}$

Now let's consider the case where throughput is 10 Gbps instead. We can follow the same step except with 10 Gbps instead of 10 Mbps as the throughput.

Now let's consider the case where throughput is 10 Gbps instead. We can follow the same step except with 10 Gbps instead of 10 Mbps as the throughput.

Following the same logic, we get $L = 2 \cdot 10^{-10}$ as our upper bound for loss rate to achieve 10 Gbps!

Is this loss rate realistic? What does this say about TCP?

"Hi, I'd like to hear a TCP joke."
"Hello, would you like to hear a TCP joke?"
"Yes, I'd like to hear a TCP joke."
"OK, I'll tell you a TCP joke."
"Ok, I will hear a TCP joke."
"Are you ready to hear a TCP joke?"
"Yes, I am ready to hear a TCP joke."
"Ok, I am about to send the TCP joke. It will last 10 seconds, it has two characters, it does not have a setting, it ends with a punchline."
"Ok, I am ready to get your TCP joke that will last 10 seconds, has two characters, does not have an explicit setting, and ends with a punchline."
"I'm sorry, your connection has timed out. ... Hello, would you like to hear a TCP joke?"