

### Problem 3

Note, wrap around if overflow.

$$\begin{array}{r} 0\ 1\ 0\ 1\ 0\ 0\ 1\ 1 \\ +\ 0\ 1\ 1\ 0\ 0\ 1\ 1\ 0 \\ \hline 1\ 0\ 1\ 1\ 1\ 0\ 0\ 1 \\ \\ 1\ 0\ 1\ 1\ 1\ 0\ 0\ 1 \\ +\ 0\ 1\ 1\ 1\ 0\ 1\ 0\ 0 \\ \hline 0\ 0\ 1\ 0\ 1\ 1\ 1\ 0 \end{array}$$

One's complement = 1 1 0 1 0 0 0 1.

To detect errors, the receiver adds the four words (the three original words and the checksum). If the sum contains a zero, the receiver knows there has been an error. All one-bit errors will be detected, but two-bit errors can be undetected (e.g., if the last digit of the first word is converted to a 0 and the last digit of the second word is converted to a 1).

### Problem 4

- a) Adding the two bytes gives 11000001. Taking the one's complement gives 00111110.
- b) Adding the two bytes gives 01000000; the one's complement gives 10111111.
- c) First byte = 01010100; second byte = 01101101.

### Problem 14

In a NAK only protocol, the loss of packet  $x$  is only detected by the receiver when packet  $x+1$  is received. That is, the receiver receives  $x-1$  and then  $x+1$ , only when  $x+1$  is received does the receiver realize that  $x$  was missed. If there is a long delay between the transmission of  $x$  and the transmission of  $x+1$ , then it will be a long time until  $x$  can be recovered, under a NAK only protocol.

On the other hand, if data is being sent often, then recovery under a NAK-only scheme could happen quickly. Moreover, if errors are infrequent, then NAKs are only occasionally sent (when needed), and ACK are never sent – a significant reduction in feedback in the NAK-only case over the ACK-only case.

### Problem 15

It takes 12 microseconds (or 0.012 milliseconds) to send a packet, as  $1500 \times 8 / 10^9 = 12$  microseconds. In order for the sender to be busy 98 percent of the time, we must have

$$util = 0.9 = (0.012n) / 30.012$$

or  $n$  approximately 2251 packets.

## Problem 22

- a) Here we have a window size of  $N = 4$ . Suppose the receiver has received packet  $k-1$ , and has ACKed that and all other preceding packets. If all of these ACK's have been received by sender, then sender's window is  $[k, k+N-1]$ . Suppose next that none of the ACKs have been received at the sender. In this second case, the sender's window contains  $k-1$  and the  $N$  packets up to and including  $k-1$ . The sender's window is thus  $[k-N, k-1]$ . By these arguments, the sender's window is of size 3 and begins somewhere in the range  $[k-N, k]$ .
- b) If the receiver is waiting for packet  $k$ , then it has received (and ACKed) packet  $k-1$  and the  $N-1$  packets before that. If none of those  $N$  ACKs have been yet received by the sender, then ACK messages with values of  $[k-N, k-1]$  may still be propagating back. Because the sender has sent packets  $[k-N, k-1]$ , it must be the case that the sender has already received an ACK for  $k-N$ . Once the receiver has sent an ACK for  $k-N$ , it will never send an ACK that is less than  $k-N$ . Thus the range of in-flight ACK values can range from  $k-N$  to  $k-1$ .


## Problem 23

In order to avoid the scenario of Figure 3.27, we want to avoid having the leading edge of the receiver's window (i.e., the one with the "highest" sequence number) wrap around in the sequence number space and overlap with the trailing edge (the one with the "lowest" sequence number in the sender's window). That is, the sequence number space must be large enough to fit the entire receiver window and the entire sender window without this overlap condition. So - we need to determine how large a range of sequence numbers can be covered at any given time by the receiver and sender windows.

Suppose that the lowest-sequence number that the receiver is waiting for is packet  $m$ . In this case, its window is  $[m, m+w-1]$  and it has received (and ACKed) packet  $m-1$  and the  $w-1$  packets before that, where  $w$  is the size of the window. If none of those  $w$  ACKs have been yet received by the sender, then ACK messages with values of  $[m-w, m-1]$  may still be propagating back. If no ACKs with these ACK numbers have been received by the sender, then the sender's window would be  $[m-w, m-1]$ .

Thus, the lower edge of the sender's window is  $m-w$ , and the leading edge of the receiver's window is  $m+w-1$ . In order for the leading edge of the receiver's window to not overlap with the trailing edge of the sender's window, the sequence number space must

thus be big enough to accommodate  $2w$  sequence numbers. That is, the sequence number space must be at least twice as large as the window size,  $k \geq 2w$ .



### Problem 32

a)

Denote  $EstimatedRTT^{(n)}$  for the estimate after the  $n$ th sample.

$$\begin{aligned} EstimatedRTT^{(4)} &= xSampleRTT_1 + \\ &\quad (1-x)[xSampleRTT_2 + \\ &\quad (1-x)[xSampleRTT_3 + (1-x)SampleRTT_4]] \\ &= xSampleRTT_1 + (1-x)xSampleRTT_2 \\ &\quad + (1-x)^2 xSampleRTT_3 + (1-x)^3 SampleRTT_4 \end{aligned}$$

b)

$$\begin{aligned} EstimatedRTT^{(n)} &= x \sum_{j=1}^{n-1} (1-x)^{j-1} SampleRTT_j \\ &\quad + (1-x)^{n-1} SampleRTT_n \end{aligned}$$

c)

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

The weight given to past samples decays exponentially.

### Problem 36

Suppose packets  $n$ ,  $n+1$ , and  $n+2$  are sent, and that packet  $n$  is received and ACKed. If packets  $n+1$  and  $n+2$  are reordered along the end-to-end-path (i.e., are received in the order  $n+2$ ,  $n+1$ ) then the receipt of packet  $n+2$  will generate a duplicate ack for  $n$  and would trigger a retransmission under a policy of waiting only for second duplicate ACK for retransmission. By waiting for a triple duplicate ACK, it must be the case that *two*

packets after packet  $n$  are correctly received, while  $n+1$  was not received. The designers of the triple duplicate ACK scheme probably felt that waiting for two packets (rather than 1) was the right tradeoff between triggering a quick retransmission when needed, but not retransmitting prematurely in the face of packet reordering.

### Problem 40

- a) TCP slowstart is operating in the intervals [1,6] and [23,26]
- b) TCP congestion avoidance is operating in the intervals [6,16] and [17,22]
- c) After the 16<sup>th</sup> transmission round, packet loss is recognized by a triple duplicate ACK. If there was a timeout, the congestion window size would have dropped to 1.
- d) After the 22<sup>nd</sup> transmission round, segment loss is detected due to timeout, and hence the congestion window size is set to 1.
- e) The threshold is initially 32, since it is at this window size that slow start stops and congestion avoidance begins.
- f) The threshold is set to half the value of the congestion window when packet loss is detected. When loss is detected during transmission round 16, the congestion window size is 42. Hence the threshold is 21 during the 18<sup>th</sup> transmission round.
- g) The threshold is set to half the value of the congestion window when packet loss is detected. When loss is detected during transmission round 22, the congestion window size is 29. Hence the threshold is 14 (taking lower floor of 14.5) during the 24<sup>th</sup> transmission round.
- h) During the 1<sup>st</sup> transmission round, packet 1 is sent; packet 2-3 are sent in the 2<sup>nd</sup> transmission round; packets 4-7 are sent in the 3<sup>rd</sup> transmission round; packets 8-15 are sent in the 4<sup>th</sup> transmission round; packets 16-31 are sent in the 5<sup>th</sup> transmission round; packets 32-63 are sent in the 6<sup>th</sup> transmission round; packets 64 – 96 are sent in the 7<sup>th</sup> transmission round. Thus packet 70 is sent in the 7<sup>th</sup> transmission round.
- i) The threshold will be set to half the current value of the congestion window (8) when the loss occurred and congestion window will be set to the new threshold value + 3 MSS. Thus the new values of the threshold and window will be 4 and 7 respectively.
- j) threshold is 21, and congestion window size is 4.
- k) round 17, 1 packet; round 18, 2 packets; round 19, 4 packets; round 20, 8 packets; round 21, 16 packets; round 22, 21 packets. So, the total number is 52.

### Problem 43

In this problem, there is no danger in overflowing the receiver since the receiver's receive buffer can hold the entire file. Also, because there is no loss and acknowledgements are returned before timers expire, TCP congestion control does not throttle the sender. However, the process in host A will not continuously pass data to the socket because the send buffer will quickly fill up. Once the send buffer becomes full, the process will pass data at an average rate of  $R \ll S$ .

<https://web.ugreen.cloud/web/#/share/e5f0800867eb4b69a8fd24460c378714> 提取码: T5CV