

Spring 24, CS 3611: Computer Networks

Solution to Homework 3

Solution to problem 1

Yes, both segments will be directed to the same socket. For each received segment, at the socket interface, the operating system will provide the process with the source IP addresses to determine the origins of the individual segments.

Solution to problem 2

a) ~ d):

	source port numbers	destination port numbers
a) A → S	4467*	23
b) B → S	5513*	23
c) S → A	23	4467*
d) S → B	23	5513*

* Note here A and B can have any possible port numbers larger than 1024.

e) Yes.

f) No, it's impossible that A and B have the same port number.

Solution to problem 3

UDP has better control over what and when is sent in a segment. The reasons that some applications choose to run over UDP includes but not limited to:

- no connection establishment (which can add delay)
- no retransmission (which can add delay)
- simple: no connection state at sender, receiver
- small segment header: 8 bytes
- no flow control and congestion control: UDP can blast away as fast as desired

Solution to problem 4

The 32-bit sequence number can have up to 2^{32} sequence numbers. So in TCP case, the whole sequence number space can send up to 2^{32} bytes before wrapping around.

a) If the line is 112-kbps, the time to transmit the 2^{32} bytes is: $2^{32} * 8 / 192k = 49.7$ hours

b) If the line is 10-Mbps, the time to transmit the 2^{32} bytes is: $2^{32} * 8 / 20M = 28.6$ minutes

c) If the line is 4-Gbps, the time to transmit the 2^{32} bytes is: $2^{32} * 8 / 2G = 17.2$ seconds

d) Suppose a packet can stay in Internet for 1 minute at most, 32-bit sequence number is not enough for 2Gbps network, i.e., there might be two different TCP segment with same sequence number. To deal with this problem, TCP introduces a **TIMESTAMP** field in TCP option. Sequence number and Timestamp together uniquely identify a segment.

Solution to problem 5

a) Consider sending an application message over a transport protocol. With TCP, the application writes data to the connection send buffer and TCP will grab bytes without necessarily putting a single message in the TCP segment; TCP may put more or less than a single message in a segment. UDP, on the other hand, encapsulates in a segment whatever the application gives it; so that, if the application gives UDP an application message, this message will be the payload of the UDP segment. Thus, with UDP, an application has more control of what data is sent in a segment.

b) With TCP, due to flow control and congestion control, there may be significant delay from the time when an application writes data to its send buffer until when the data is given to the network layer. UDP does not have delays due to flow control and congestion control.

Solution to problem 6

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.

Solution to problem 7

a)

1. GoBackN:

A sends 10 segments in total. They are initially sent segments 1, 2, 3, 4, 5, 6 and later re-sent segments 3, 4, 5 and 6.

B sends 9 ACKs. They are ACKS with sequence numbers 1, 2, 2, 2, 2, and then ACKS with sequence numbers 3, 4, 5, 6.

2. Selective Repeat:

A sends 7 segments in total. They are initially sent segments 1, 2, 3, 4, 5, 6 and later re-sent segments 3.

B sends 6 ACKs. They are 5 ACKS with sequence number 1, 2, 4, 5, 6. And there is one ACK with sequence number 3.

3. TCP:

A sends 7 segments in total. They are initially sent segments 1, 2, 3, 4, 5, 6 and later re-sent segments 3.

B sends 6 ACKs. They are 5 ACKS with sequence number 2, 3, 3, 3, 3, and one ACK with sequence numbers 7. Note that TCP always send an ACK with expected sequence number.

b) TCP. This is because TCP uses fast retransmit without waiting until time out.

Solution to problem 8 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 or $R \ll S$.