CNT-5008 Midterm 2 Grading Key

Date: Tuesday November 5, 2019

Name: ……

Instructions:

- This exam is open book and open notes. Textbooks and notes on tablet devices are acceptable but they must be put into airplane mode. No device with a keyboard is acceptable.
- It is recommended that you use a pencil, such that you can make corrections. Do not use highlighters, and don't use red colored pens.
- Allotted time is 180 minutes.
- Note that the points add up to 100 + 20 bonus points.

Problem 1 – Reliable data transfer protocol (20 pts)

In protocol rdt3.0 we developed in class, the ACK packets don't have sequence numbers. However, they do have an ACK field that contains the sequence number of the packets they are acknowledging. Why it is that the ACK packets do not require sequence numbers?

To best answer this question, consider why we needed sequence numbers in the first place. We saw that the sender needs sequence numbers so that the receiver can tell if a data packet is a duplicate of an already received data packet. In the case of ACKs, the sender does not need this info (i.e., a sequence number on an ACK) to tell detect a duplicate ACK. A duplicate ACK is obvious to the rdt3.0 receiver, since when it has received the original ACK it transitioned to the next state. The duplicate ACK is not the ACK that the sender needs and hence is ignored by the rdt3.0 sender.

Problem 2 ACK vs NAK (20pts)

Consider a reliable data transfer protocol that uses only negative acknowledgement (NAK). Suppose the sender sends data only infrequently. Would a NAK only protocol be preferable to a protocol that uses ACKs? Why?

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 receivers 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.

Now suppose that the sender has a lot of data to send and the end-to-end connection experiences losses only very rarely. In this second case, would a NAK-only protocol be preferable to one that uses ACKs? Why?

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 3 TCP acks (20 pts)

Host A and B are communicating over a TCP connection and Host B has already received from A all bytes up through byte 126. Suppose Host A then sends two segments to Host B back-to-back. The first contains 80 bytes of data, and second segment contains 40 bytes of data, respectively. In the first segment, the sequence number is 127, the source port number is 302, and the destination port number is 80. Host B sends an acknowledgement whenever it receives a segment from Host A.

a) In the second segment send from A, what are the sequence number, source port number and destination port number?

In the second segment from Host A to B, the sequence number is 207, source port number is 302 and destination port number is 80.

b) If the first segment arrives before the second segment, in the acknowledgment of the first arriving segment, what is the acknowledgment number, the source port number and the destination port number?

If the first segment arrives before the second, in the acknowledgement of the first arriving segment, the acknowledgement number is 207, the source port number is 80 and the destination port number is 302.

c) If the second segment arrives before the first segment, in the acknowledgment of the first arriving segment, what is the acknowledgement number?

If the second segment arrives before the first segment, in the acknowledgement of the first arriving segment, the acknowledgement number is 127, indicating that it is still waiting for bytes 127 and onwards.

d) Support the two segments sent by A arrive in order at B. The first acknowledgment is lost and the second acknowledgment arrives after the first timeout interval. Draw a timing diagram, showing these segments and all the other segments and acknowledgements sent. Assume that there is no additional packet loss. For each segment in your figure, provide the segment number and the number of bytes of data; for each acknowledgment you add, provide the acknowledgment number.

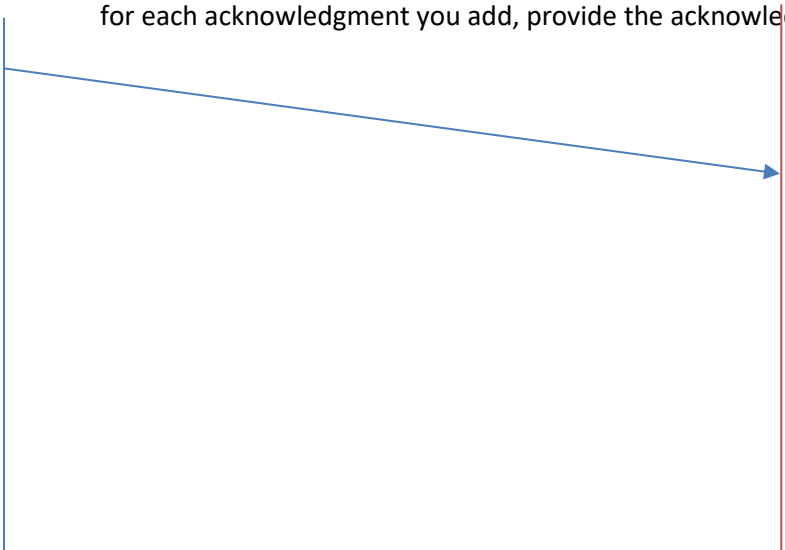

Solution for d)

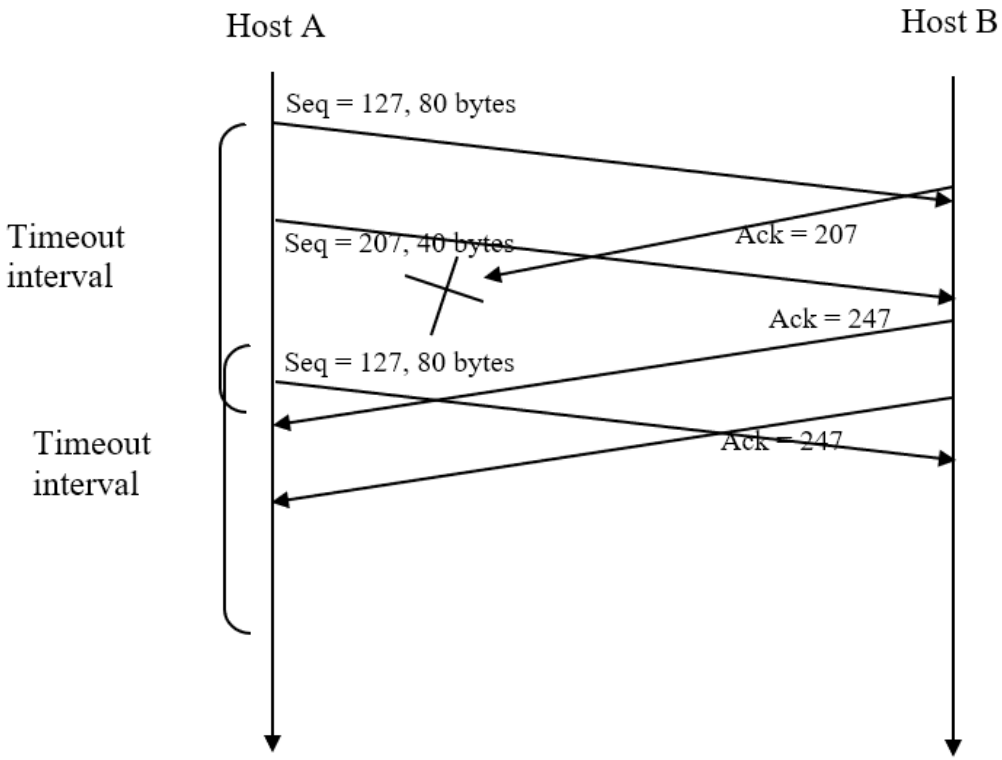

Problem 4: Flow control (20pts)

Hosts A and B are directly connected with a 100 Mbps link. There is one TCP connection between the two hosts, and Host A is sending to Host B a very large file over this connection. Host A can send application data into its TCP socket at a rate as high as 120 Mbps but Host B can read out of its TCP receive buffer at a maximum rate of 20Mbps. Describe the effect of TCP flow control.

Since the link capacity is only 100 Mbps, so host A's sending rate can be at most 100Mbps. Still, host A sends data into the receive buffer faster than Host B can remove data from the buffer. The receive buffer fills up at a rate of roughly 80Mbps. When the buffer is full, Host B signals to Host A to stop sending data by setting RcvWindow = 0. Host A then stops sending until it receives a TCP segment with RcvWindow > 0. Host A will thus repeatedly stop and start sending as a function of the RcvWindow values it receives from Host B. On average, the long-term rate at which Host A sends data to Host B as part of this connection is no more than 20Mbps.

Problem 5: Data plane / control plane (20 pts)

Explain the difference between the data plane and the control plane of the networking layer.

The main function of the data plane is packet forwarding, which is to forward datagrams from their input links to their output links. For example, the data plane's input ports perform physical layer function of terminating an incoming physical link at a router, perform link-layer function to interoperate with the link layer at the other side of the incoming link, and perform lookup function at the input ports.

The main function of the control plane is routing, which is to determine the paths a packet takes from its source to its destination. A control plane is responsible for executing routing protocols, responding to attached links that go up or down, communicating with remote controllers, and performing management functions.

Explain the difference between routing and forwarding in the networking layer.

The key differences between routing and forwarding is that forwarding is a router's local action of transferring packets from its input interfaces to its output interfaces, and forwarding takes place at very short timescales (typically a few nanoseconds), and thus is typically implemented in hardware. Routing refers to the network-wide process that determines the end-to-end paths that packets take from sources to destinations. Routing takes place on much longer timescales (typically seconds), and is often implemented in software.

What field in the IP header can be used to ensure that a packet is forwarded through no more than N routers?

Time-to-live.

Do routers have IP addresses? Discuss.

Yes. They have one address for each interface.

Problem 6: Forwarding (20 pts)

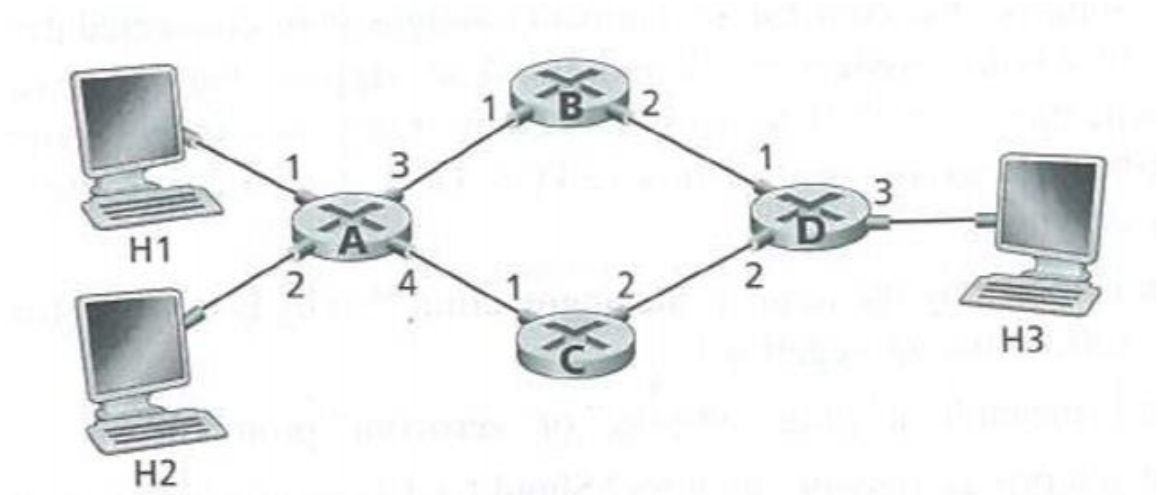

Consider the network above:

a) Show the forwarding table in router A, such that all traffic destined to Host H3 is forwarded through interface 3.

Destination Address	Link Interface
H3	3

b) Can you write down a forwarding table in router A, such that all traffic from H1 destined to H3 is forwarded through interface 3, while all traffic from H2 destined to H3 is forwarded through interface 4?

No, because forwarding rule is only based on destination address.