1. Consider the following network topology where some residential hosts connect to the Internet via NAT devices:

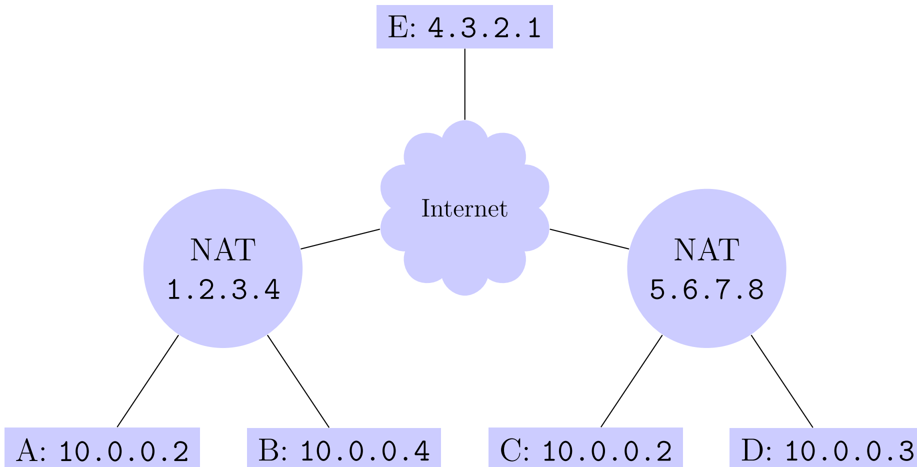

   (a) Host A has an active SSH connection (Port 22) to Host E. Give a plausible value for the (source address, destination address, source port, destination port) 4-tuple for packets in this connection as they exit Host A's computer. Indicate what elements of the tuple would be assigned at connection time (pick whatever value you want for these *ephermeral* ports).

   (b) What is a plausible value for the (source address, destination address, source port, destination port) 4-tuple for packets for the connection in (a) when they arrive at Host E? Again, indicate what elements would be assigned at connection time.

   (c) Suppose Host D wishes to serve web pages (on port 80) to the Internet. What must happen in order for devices on the Internet to access the server on Host D?

   (d) Host B attempts to connect Host D (i.e. by sending a packet to D). List out how the (source address, destination address, source port, destination port) 4-tuple evolves as it traverses the network from B to D. Indicate what elements would be assigned at connection time.

   (e) Host A and Host C have the same IP address. Is this an issue? Explain why it is or isn't.

2. Course Handout: Transport layer 15, 16, 17

3. Using a timing diagram to track the sequence of sent messages and acknowledgements, show that if messages can be reordered (i.e., a packet might be

delayed and received after a packet that is subsequently sent), the rdt 3.0 state machine using alternating bits will not work as designed.

4. For TCP, the lecture notes (slide 89) make the following statement: ”receivers CAN buffer out of sequence packets”.
   (a) Why might a receiver choose not to do this?
   (b) Does a sender need to know whether a receiver buffers out of sequence packets? Why? Why not?
5. Derive the TCP throughput equation on slide 199. Why is it important for TCP to measure the round trip time?
6. It is interesting to consider the consequences if TCP's behaviour was a bit different. Describe the consequences if the following were true:
   (a) TCP slow-start increased CWND by 1 for each ack
   (b) TCP does not reset the CWND to 1 after a timeout
   (c) TCP resets CWND to 1 after a duplicate ACK
   (d) TCP uses MIMD (i.e., CWND doubles every RTT) instead of AIMD
7. TCP has a few problems areas (slide 206)
   (a) Why are non-congestion losses bad for TCP? Give examples of how non-congestion losses might arise.
   (b) TCP has a tendency to fill up buffers. What problem do backed up queues present?
   (c) Give two reasons why short flows are inefficient for TCP. What kind of Internet activity can result in many short flows?