Lecture 12, Review Question 2:

- Assume that a TCP sender *A* is connected to router *X* by a 100 Mb/s link, that the corresponding receiver *B* is connected to router *Y* by a 100 Mb/s link and that the link connecting *X* and *Y* is 10 Mb/s. Also, assume that the roundtrip propagation time between *A* and *B* is 50 ms, that the MSS is 1250 bytes and that *ssthresh*=64 KB. Suppose that *A* starts in the slow-start state. At what rate is *A* sending after 400 ms?

Nx1250kB >= 64 kBytes
N = 52
We are looking for
cwnd = 52

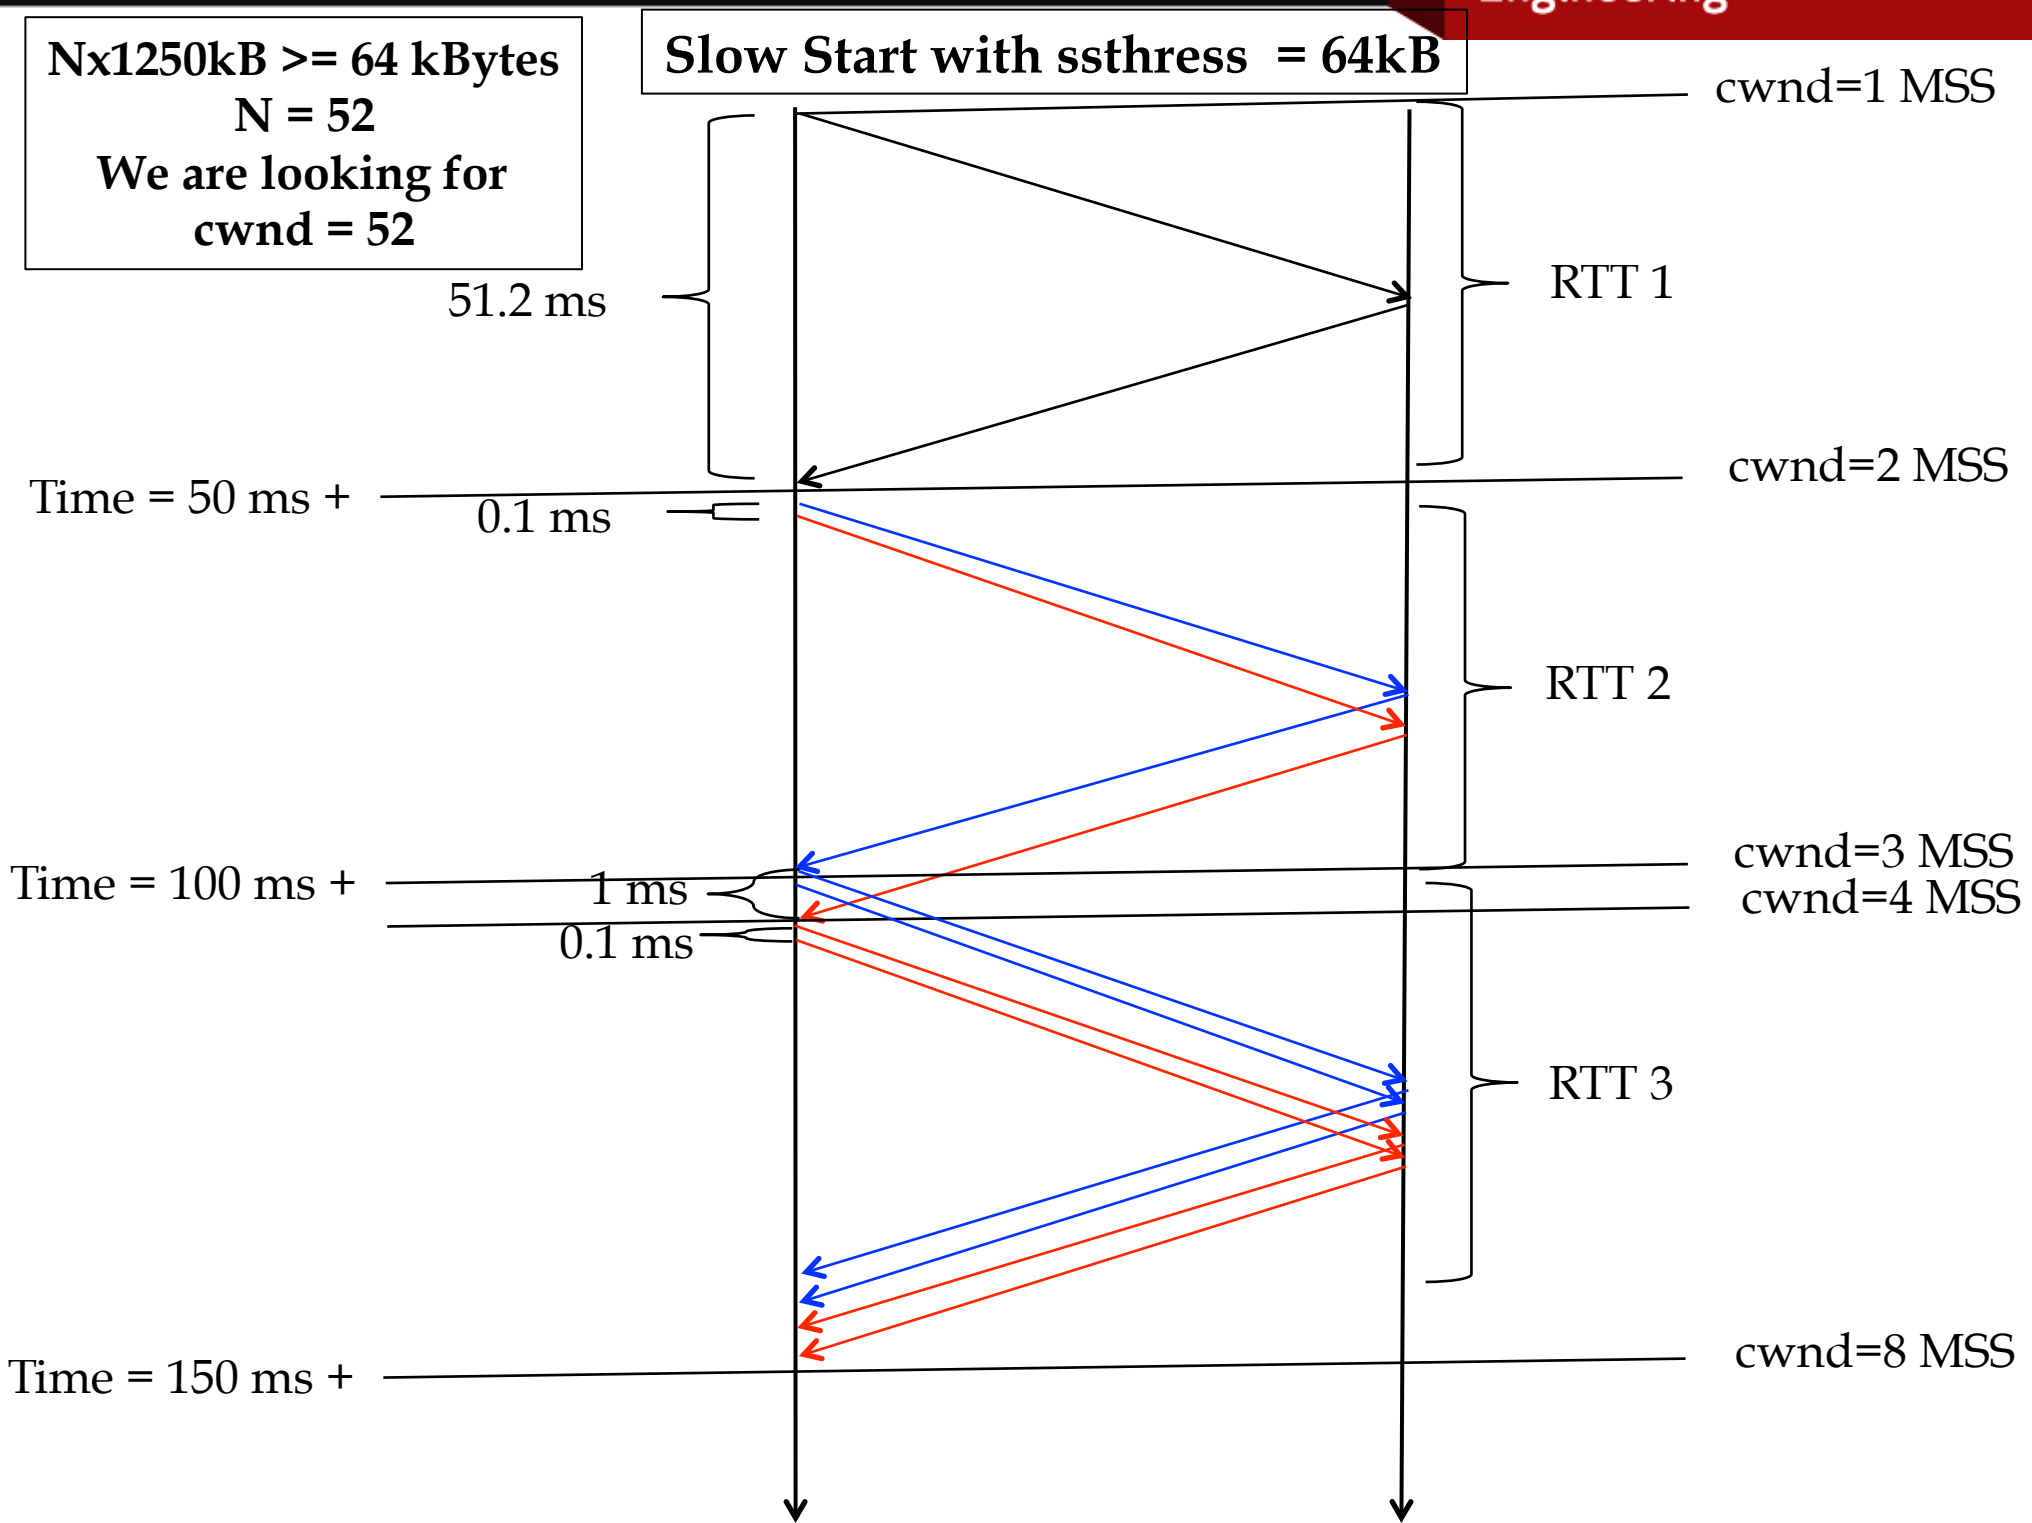

Nx1250kB = 64 kBytes
N = 52
We are looking for
cwnd = 52
Then what happens after that?

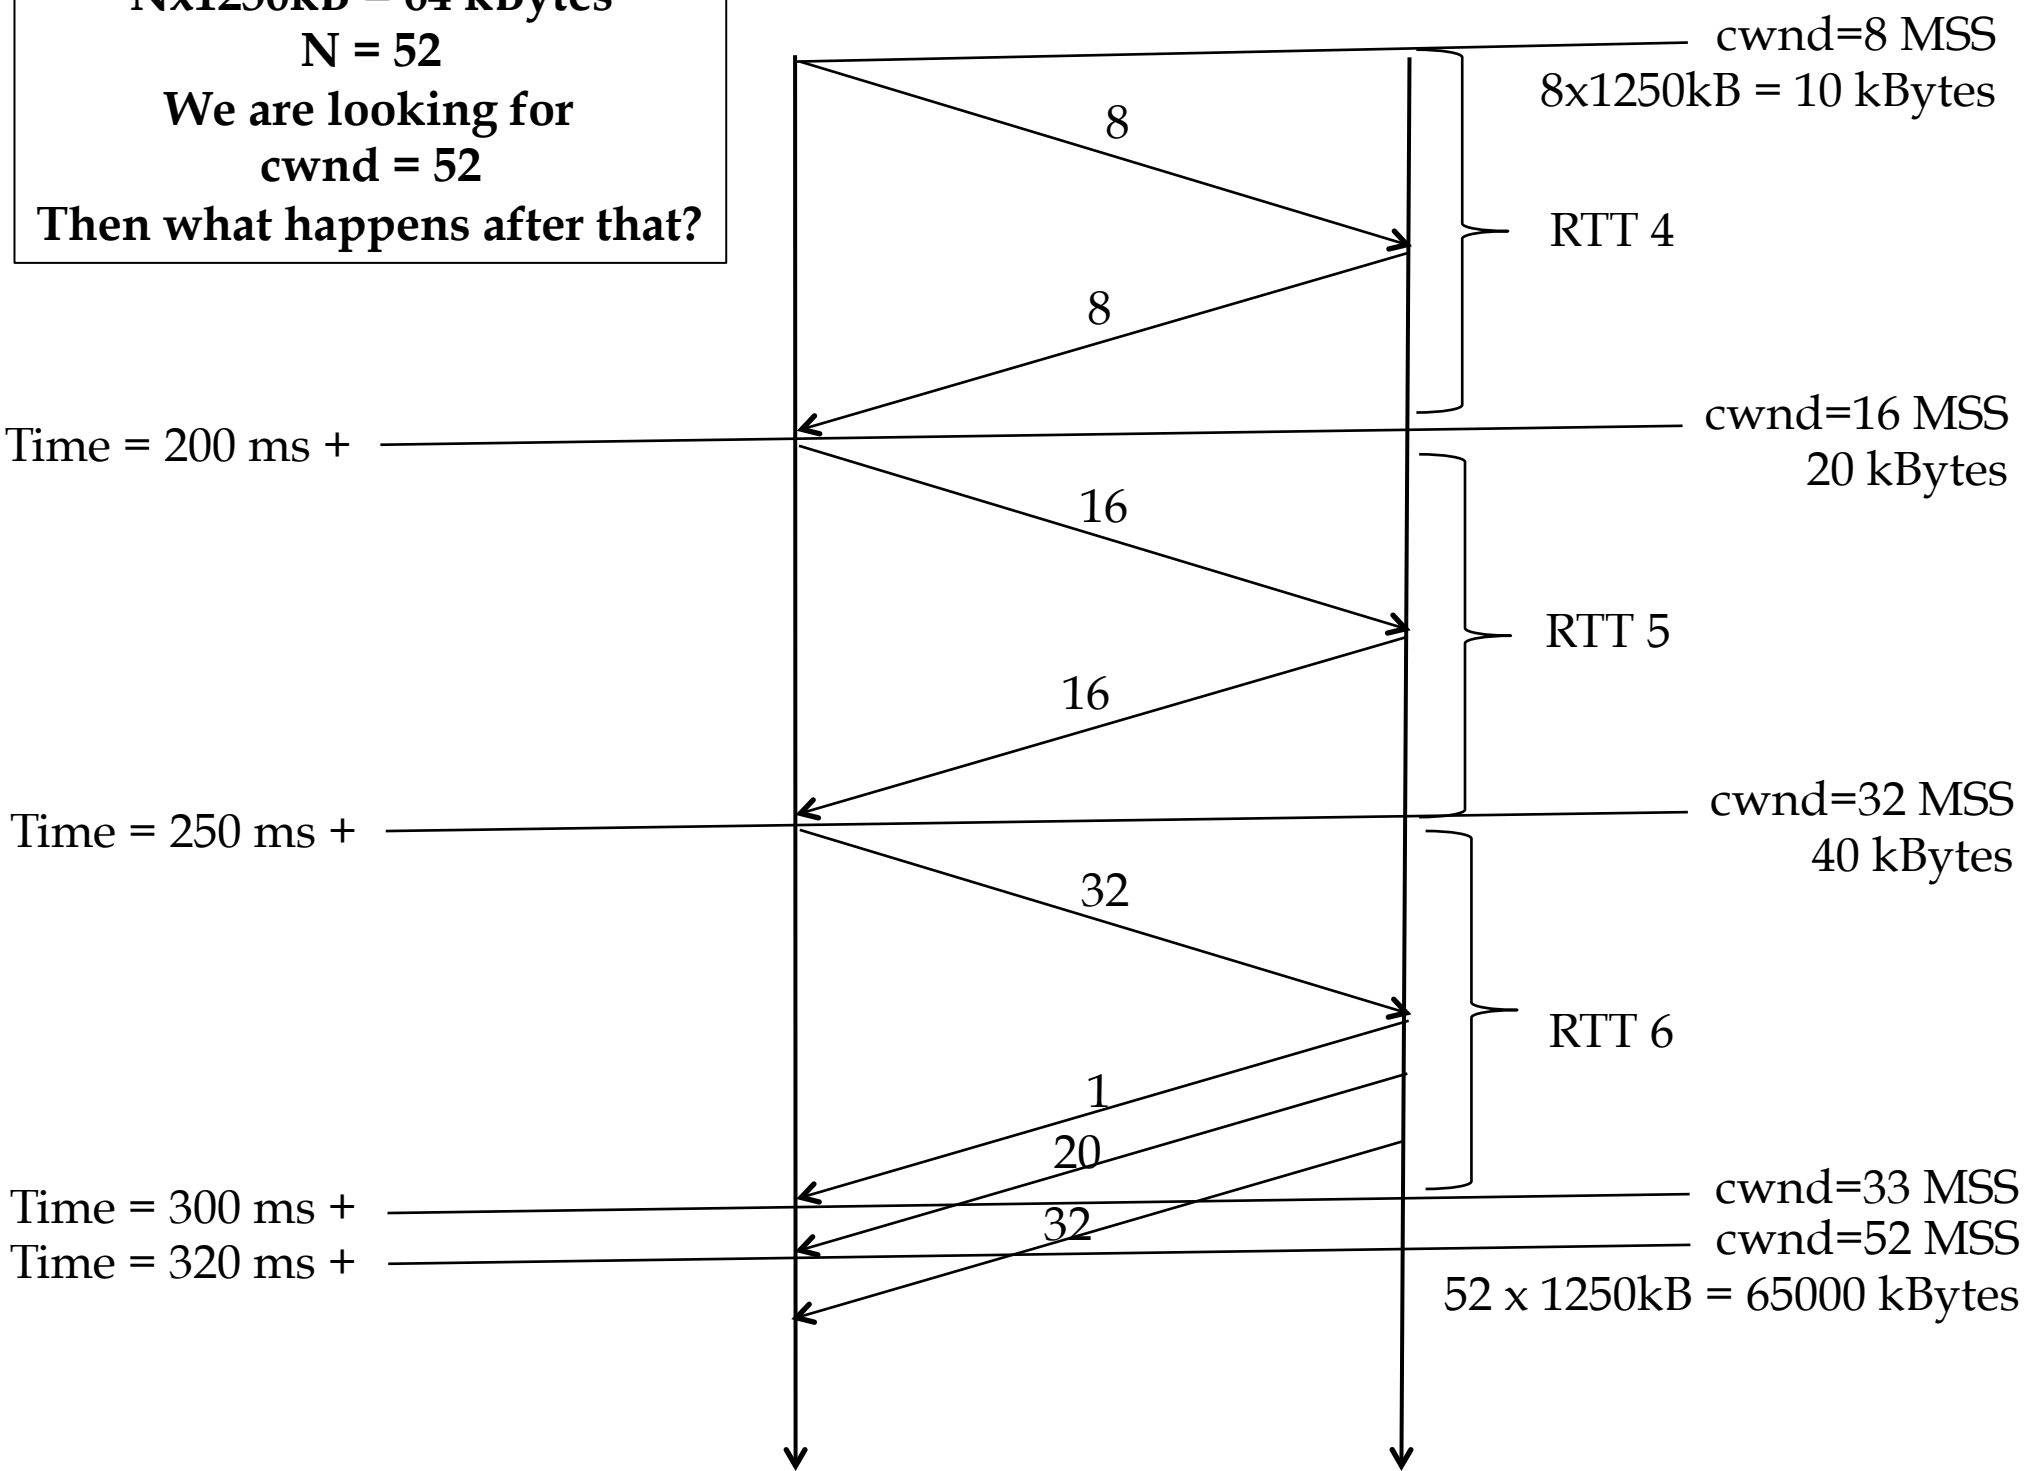

After that…

- What happens once we have reached 64KB?
  - » Congestion control is geared toward controlling the sending rate when it is the network limiting the sender.
  - » What if it is the application limiting it?
    - This is when we don't have enough data from the application on the sending host so that the send buffer empties before we run out of *cwnd* sending space. In this case we are application limited and *cwnd* will not be incremented for each ACK.
  - » What if it is the receive window limiting it?
    - This is when *cwnd* >= *rwnd* and it is no longer *cwnd* that controls the sending rate of the sender. Again, *cwnd* should not continue to be incremented if it is not *cwnd* that limits our sending rate.
      - I'm still looking for a definitive reference for this…
- So, if we assume that the receive window has the default maximum of 64KBytes that *cwnd* will not increment higher than that.
  - » The second part of question turns off that assumption…