CS 168 Section 7: Advanced Congestion Control

1. TCP Fast Recovery (30 mins)

Consider a TCP connection, which is currently in Congestion Avoidance (AIMD).

- The last ACK sequence number was 101 (the receiver expects the packet #101 for next).
- The CWND size is 10 (in packets).
- The packets #101-110 were sent at t = 0, 0.1, …, 0.9 (sec), respectively.
- The packet #102 is lost for its first transmission.
- RTT is 1 second.

Fill in the tables below, until the sender transmits the packet #116.

(a) Without Fast Recovery

- *On new ACK,* ***CWND = CWND +*** $\lfloor 1/CWND \rfloor$
- *On triple duplicate ACK,* ***SSTHRESH = CWND/2****, then* ***CWND = SSTHRESH***

t (s)	ACK (due to)	CWND	xmit
1.0	102 (101)	10+1/10	111
1.2	102 (103)	10+1/10	-
1.3	102 (104)	10+1/10	-
1.4	102 (105)	5	102
1.5	102 (106)	5	-
1.6	102 (107)	5	-
1.7	102 (108)	5	-
1.8	102 (109)	5	-
1.9	102 (110)	5	-
2.0	102 (111)	5	-
2.4	112 (102)	5+1/5	112-116

(b) Fast Recovery

- *On triple duplicate ACK, **SSTHRESH = CWND/2,** then **CWND = SSTHRESH + 3**, and enter fast recovery*
- *In fast recovery, **CWND += 1** on every duplicate ACK*
- *Exit fast recovery on new ACK, setting **CWND = SSTHRESH***

t (s)	ACK (due to)	CWND	xmit
1.0	102 (101)	10+1/10	111
1.2	102 (103)	10+1/10	-
1.3	102 (104)	10+1/10	-
1.4	102 (105)	8	102
1.5	102 (106)	9	-
1.6	102 (107)	10	-
1.7	102 (108)	11	112
1.8	102 (109)	12	113
1.9	102 (110)	13	114
2.0	102 (111)	14	115
2.4	112 (102)	5	116

(Feel free to skip this and let students do it at home, if you don't have enough time)

(c) Consider a scenario where two packets are lost: #102 and #107. What would happen, if we have Fast Recovery or not?

Without Fast Recovery, the sender gets stuck and #107 is retransmitted due to TCP timeout.

With Fast Recovery, #107 is retransmitted due to triple dupACK, without having to wait for retransmission timeout.

2. TCP Throughput Equation **(20mins)**

The following equation provides a simple way to estimate the throughput of a TCP connection, as a function of the loss probability (p) and the round-trip time (RTT).

$$\text{Throughput} = \sqrt{\frac{3}{2}}\frac{1}{RTT\sqrt{p}}$$

(a) Derive the above equation yourself! (give students enough time!)

Slides.

(b) Alice wants to send a large amount of data to Bob, over a network path with RTT=100ms, p=0.01, and MSS=10,000bits. What is the expected throughput in Mbps? ($\sqrt{\frac{3}{2}} \approx 1.22$)

1.22 Mbps

(c) Alice has two options to improve the throughput: halving either the RTT or the loss probability (p). If the both options cost the same, which one is more cost effective?

Halving the RTT ($2 > \sqrt{2}$)

(d) Food for thoughts: Considering how the equation is derived, what assumptions does it need for accurate prediction? When is it possible that they may not hold in reality?

- e.g., RTT is predictable. ← RTT may fluctuate at a short time scale, for example, due to queueing delay.
- The loss probability p is predictable. ← It can be dynamic depending on the level of congestion.
- All loss indications must be exclusively triple dupACKs. ← TCP retransmission timeout leads to slow start, deviating from the "sawtooth" pattern.
- Only the network is the bottleneck. ← The receiver may also throttle the rate with Flow Control.
- ...