Congestion Control

CS 168 Discussion

Week 7

RECAP: How does TCP set rate?

- How much data can be outstanding?
  - min{RWND, CWND}
- RWND: do not overload the receiver
  - Maintained by receiver
  - Advertised to sender
- CWND: do not overload the network
  - Maintained by sender with Slow Start and AIMD

Managing CWND: Ver.1

- Let's ignore timeout and Slow Start for now.
  - We are in Congestion Avoidance phase
- A new ACK: CWND += 1/CWND
  - Additive increase
- Three duplicate ACKs: CWND /= 2
  - Multiplicative decrease
  - Quiz: why three?

Worksheet 1(a)

t (s)	ACK (due to)	CWND	xmit
1.0	102 (101)	10+1/10	111
1.2	102 (103)	10+1/10	-

Worksheet 1(a)

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	-

Worksheet 1(a)

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

Fast Retransmit

Worksheet 1(a)

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 (110)	5	-

Cannot send any more packets, since # of unACKed > CWND

Worksheet 1(a)

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 (110)	5	-
2.4	112 (102)	5+1/5	112-116

A new ACK triggers a burst

Managing CWND: Ver.1

- No activity, followed by a sudden burst
  - Burst is bad!
    - Especially when CWND is large (e.g., 1000s of packets)
  - For others: congestion at a small time scale
  - For you: packets are more likely to be dropped

- Why does this happen?
  - We overestimate the # of in-flight packets!

Worksheet 1(a)

t (s)	ACK (due to)	CWND	xmit	# of unACKed	Actually in-flight
1.0	102 (101)	10+1/10	111	10	10
1.2	102 (103)	10+1/10	-	10	9
1.3	102 (104)	10+1/10	-	10	8
1.4	102 (105)	5	102	10	7
1.5	102 (106)	5	-	10	6
1.6	102 (107)	5	-	10	5
1.7	102 (108)	5	-	10	4
1.8	102 (109)	5	-	10	3
1.9	102 (110)	5	-	10	2
2.0	102 (110)	5	-	10	1
2.4	112 (102)	5+1/5	112-116	5	5

Managing CWND: Ver.2

- Idea: an ACK implies that the receiver got an packet
  - The packet is not in-flight any longer.
  - Give temporary credits so as to keep the packets flow!

- On triple dupACKs, enter “Fast Recovery”.
  - ssthresh = cwnd / 2
  - cwnd = ssthresh + 3
  - On dupACK, cwnd = cwnd + 1
  - On new ACK, exit Fast Recovery and set cwnd=ssthresh

Worksheet 1(b)

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

8 = 5 + 3

Worksheet 1(b)

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

Packets starts to flow

Worksheet 1(b)

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

Back to Congestion Avoidance

Worksheet 1(b)

t (s)	ACK (due to)	CWND	xmit	# of unACKed	Actually in-flight
1.0	102 (101)	10+1/10	111	10	10
1.2	102 (103)	10+1/10	-	10	9
1.3	102 (104)	10+1/10	-	10	8
1.4	102 (105)	8	102	10	7
1.5	102 (106)	9	-	10	6
1.6	102 (107)	10	-	10	5
1.7	102 (108)	11		1	5
1.8	102 (109)	12		2	5
1.9	102 (110)	13	114	13	5
2.0	102 (111)	14	115	14	5
2.4	112 (102)	5	116	5	5

Temporary credits…

reflect the right CWND

The Big Picture

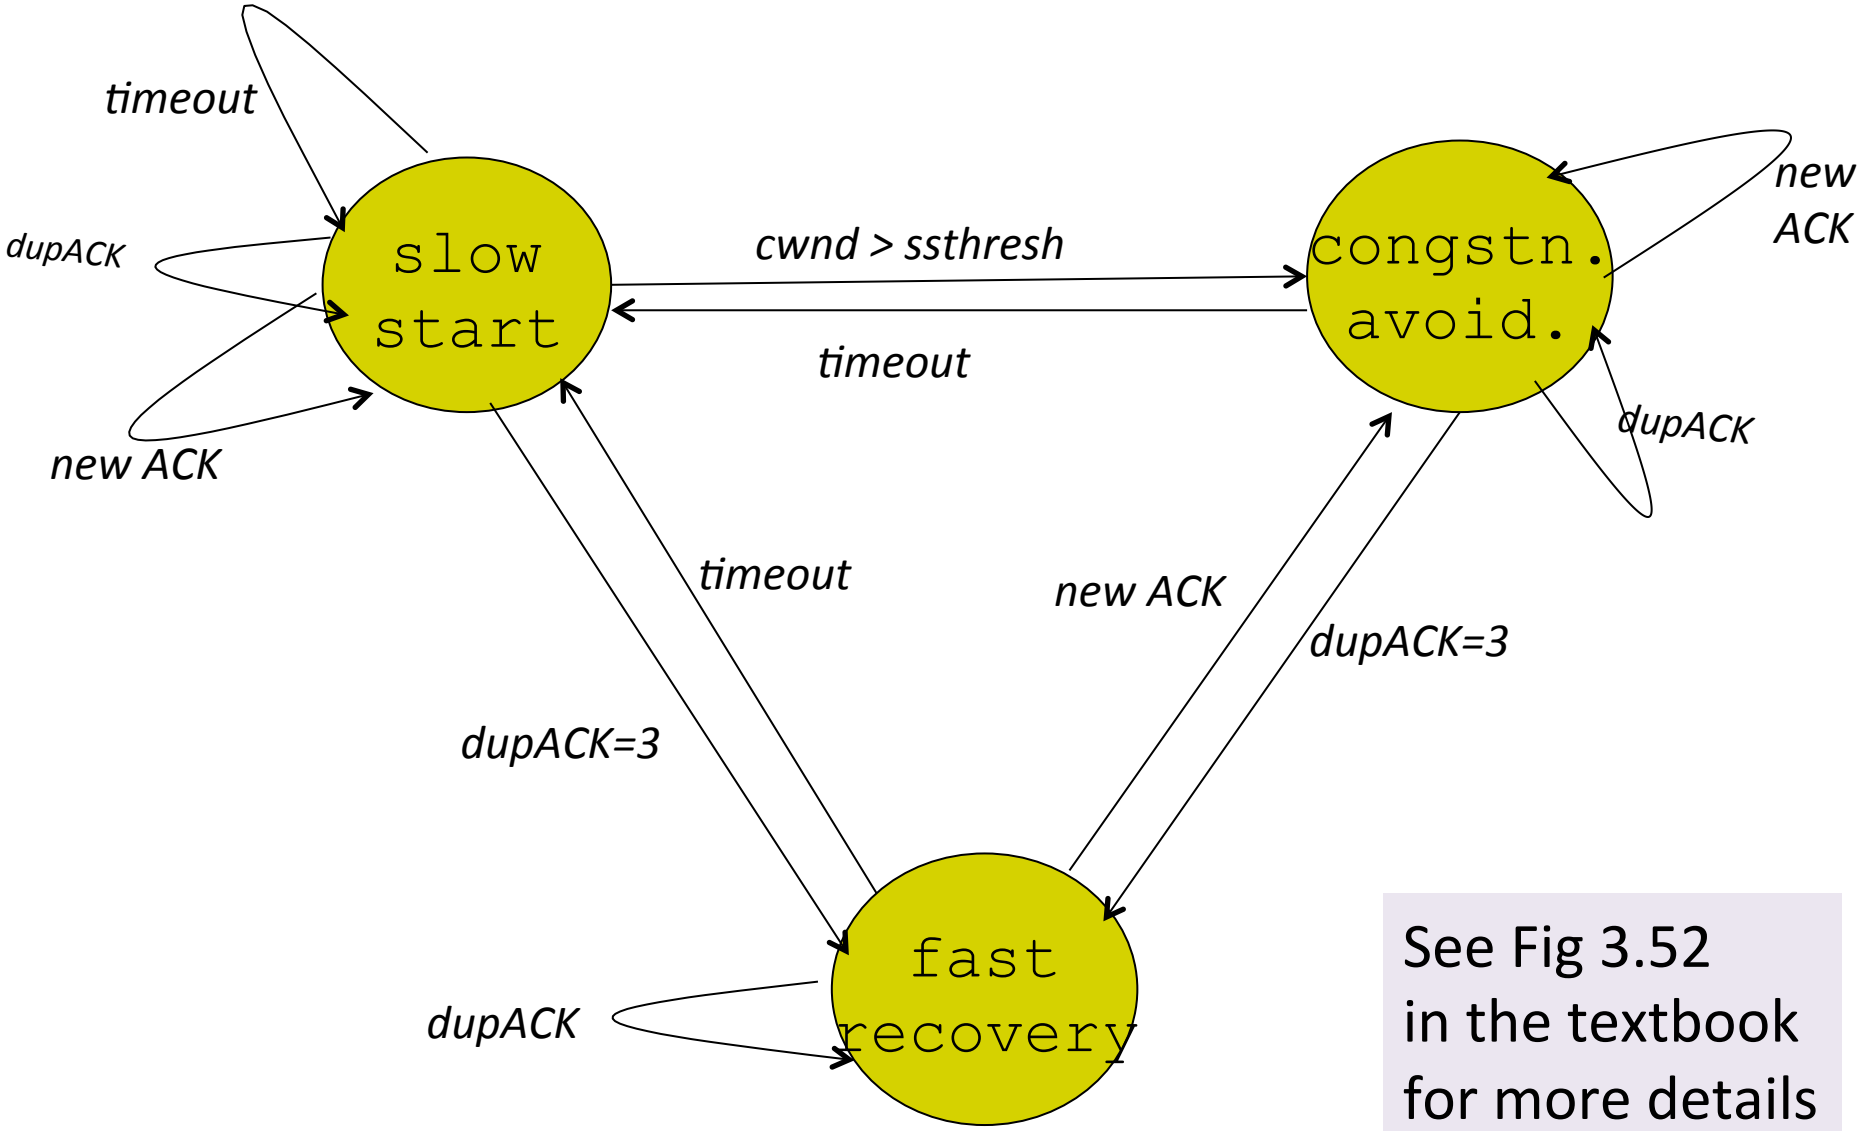

See Fig 3.52
in the textbook
for more details

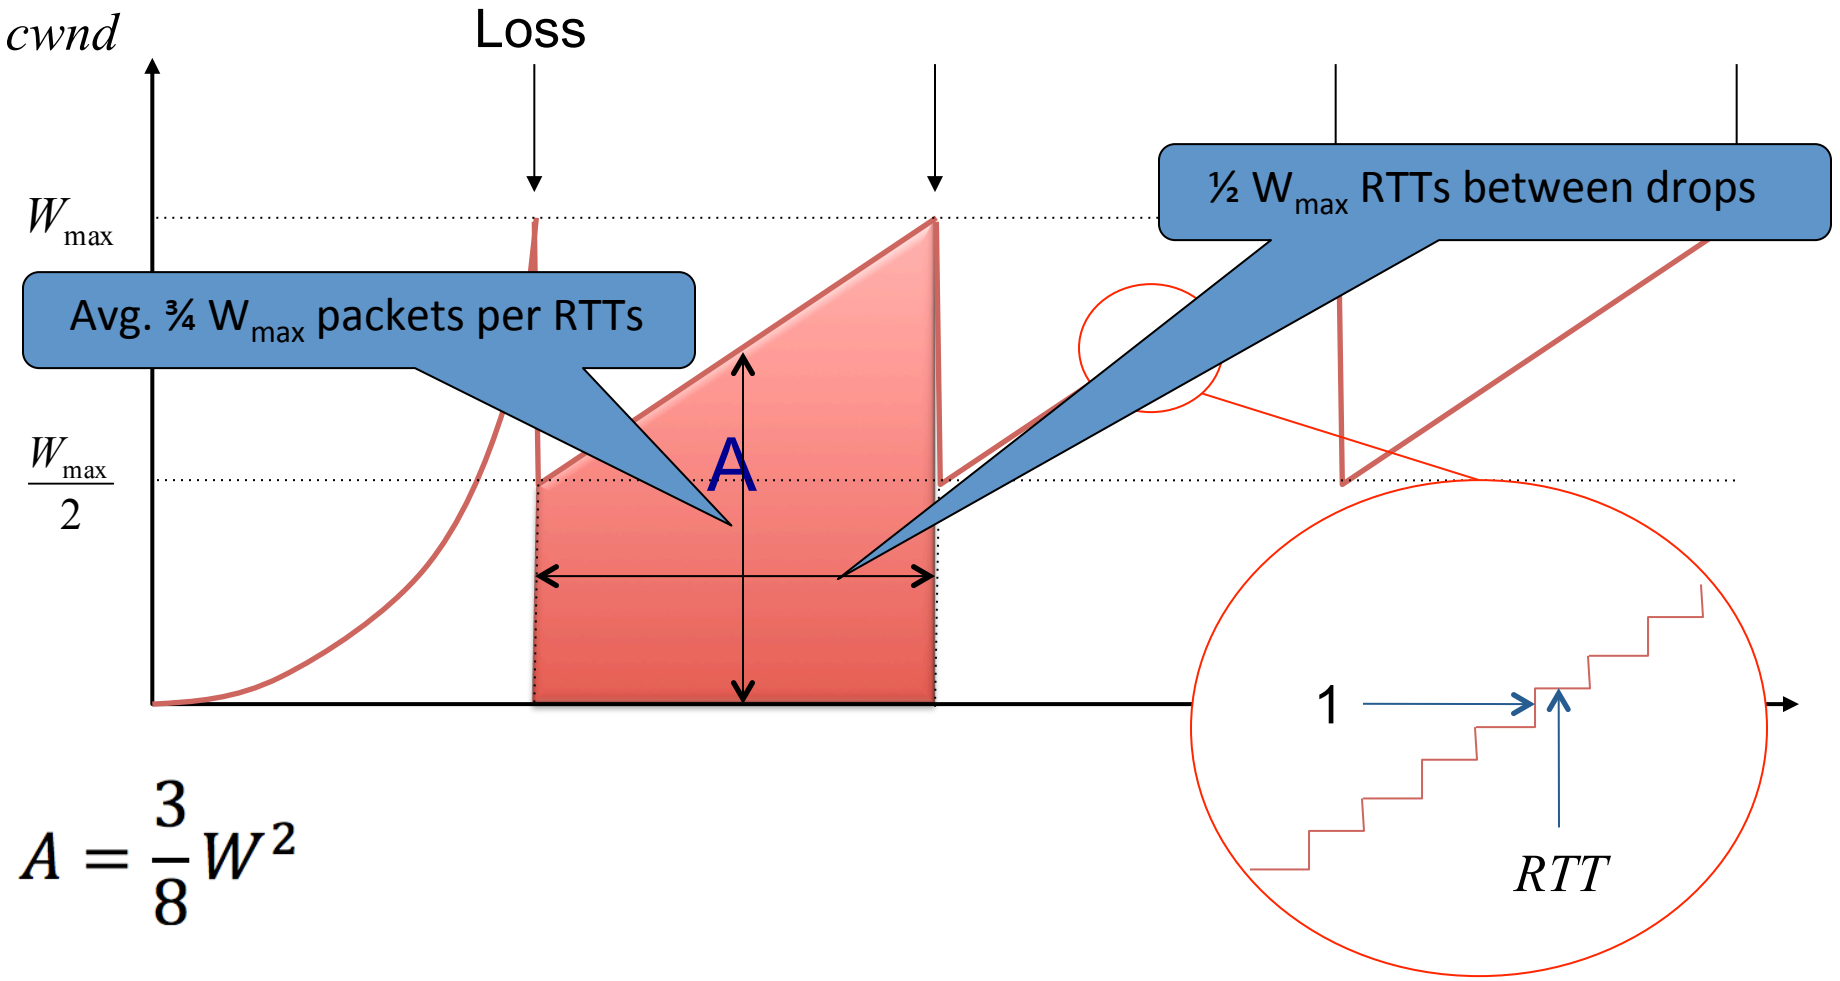

$$A = \frac{3}{8}W^2$$

$$p = \frac{1}{A}$$

$$Throughput = \frac{\#\ of\ packets}{time} = \frac{A}{\left(\frac{W}{2}\right)RTT} = \sqrt{\frac{3}{2}}\frac{1}{RTT\sqrt{p}}$$

Worksheet 2(b) and 2(c)

- (b) 1.22Mbps

  1.22 * 10^4 / (0.1 * 0.1)

- (c) RTT

  2 > sqrt(2)

Worksheet 2 (d)

- The loss probability p is predictable.
  - It can be dynamic depending on the level of congestion.
- Loss is indicated by triple dupACK.
  - TCP retransmission timeout deviates from “sawtooth”
- Only the network is the bottleneck.
  - What if the receiver is the bottleneck at a given time?
- …