

# Chapter 3

## Transport Layer

Reti di Elaboratori

Corso di Laurea in Informatica

Università degli Studi di Roma "La Sapienza"

Canale A-L

Prof.ssa Chiara Petrioli

Parte di queste slide sono state prese dal materiale associato al libro  
*Computer Networking: A Top Down Approach*, 5th edition.

All material copyright 1996-2009

J.F Kurose and K.W. Ross, All Rights Reserved

Thanks also to Antonio Capone, Politecnico di Milano, Giuseppe Bianchi and  
Francesco LoPresti, Un. di Roma Tor Vergata

# TCP: controllo di congestione

- r Il TCP ha dei meccanismi di controllo della congestione
  - m il flusso dei dati in ingresso in rete è anche regolato dalla situazione di traffico in rete
  - m se il traffico in rete porta a situazioni di congestione il TCP riduce velocemente il traffico in ingresso
  - m in rete non vi è nessun meccanismo per notificare esplicitamente le situazioni di congestione
  - m il TCP cerca di scoprire i problemi di congestione sulla base degli eventi di perdita dei pacchetti



# TCP Congestion Control

- r end-end control (no network assistance)
- r sender limits transmission:  
 $\text{LastByteSent} - \text{LastByteAked} \leq \text{CongWin}$

r Roughly,

$$\text{rate} = \frac{\text{CongWin}}{\text{RTT}} \text{ Bytes/sec}$$

- r CongWin is dynamic, function of perceived network congestion

## How does sender perceive congestion?

- r loss event = timeout *or* 3 duplicate acks
- r TCP sender reduces rate (CongWin) after loss event

## three mechanisms:

- m AIMD
- m slow start
- m conservative after timeout events

# Starting a TCP transmission

- r A new offered flow may suddenly overload network nodes
  - m receiver window is used to avoid recv buffer overflow
  - m But it may be a large value (16-64 KB)
- r Idea: slow start
  - m Start with small value of cwnd
  - m And increase it as soon as packets get through
    - Arrival of ACKs = no packet losts = no congestion
- r Initial cwnd size:
  - m Just 1 MSS!
  - m Recent (1998) proposals for more aggressive starts (up to 4 MSS) have been found to be dangerous

# Detecting congestion and restarting

- r Segment gets lost
  - m Detected via RTO expiration
  - m Indirectly notifies that one of the network nodes along the path has lost segment
    - Because of full queue
- r Restart from  $cwnd=1$  (slow start)
- r But introduce a supplementary control: slow start threshold
  - $ssthresh = \max(\min(cwnd, window)/2, 2MSS)$
  - m The idea is that we now KNOW that there is congestion in the network, and we need to increase our rate in a more careful manner...
  - m Ssthresh defines the "congestion avoidance" region


# Congestion avoidance

- r If  $cwnd < ssthresh$ 
  - m Slow start region: Increase rate exponentially
- r If  $cwnd \geq ssthresh$ 
  - m **Congestion avoidance** region : Increase rate linearly
  - m At rate 1 MSS per RTT
    - Practical implementation:  $cwnd += MSS * MSS / cwnd$ 
      - Good approximation for 1 MSS per RTT
      - Alternative (exact) implementations: count!!
- r Which initial  $ssthresh$ ?
  - $ssthresh$  initially set to 65535: unreachable!

Corrisponde ad un segmento per finestra


*In essence, congestion avoidance is flow control imposed by sender while advertised window is flow control imposed by receiver*

# Simplified example (overall)



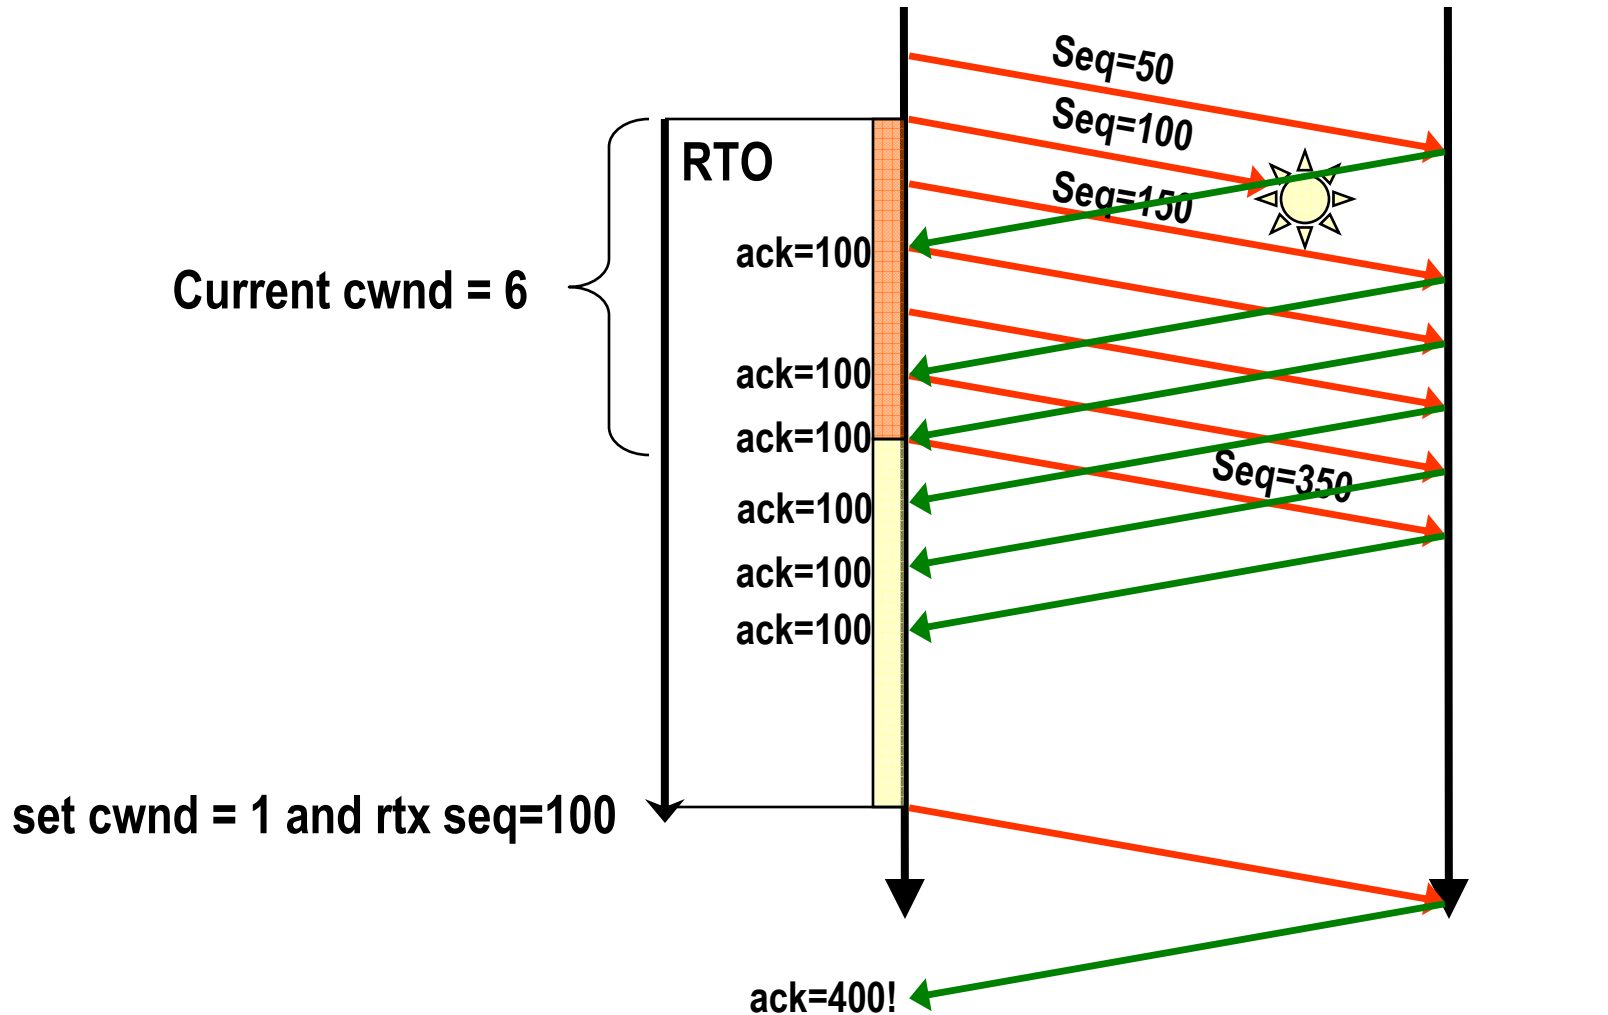
# The Fast Retransmit Algorithm

- Idea: use duplicate ACKs!
  - ⇒ Receiver responds with an ACK every time it receives an out-of-order segment
  - ⇒ ACK value = last correctly received segment
- FAST RETRANSMIT algorithm:
  - ⇒ if 3 duplicate acks are received for the same segment, assume that the next segment has been lost. Retransmit it right away.
  - ⇒ Helps if single packet lost. Not very effective with multiple losses
- And then? A congestion control issue...





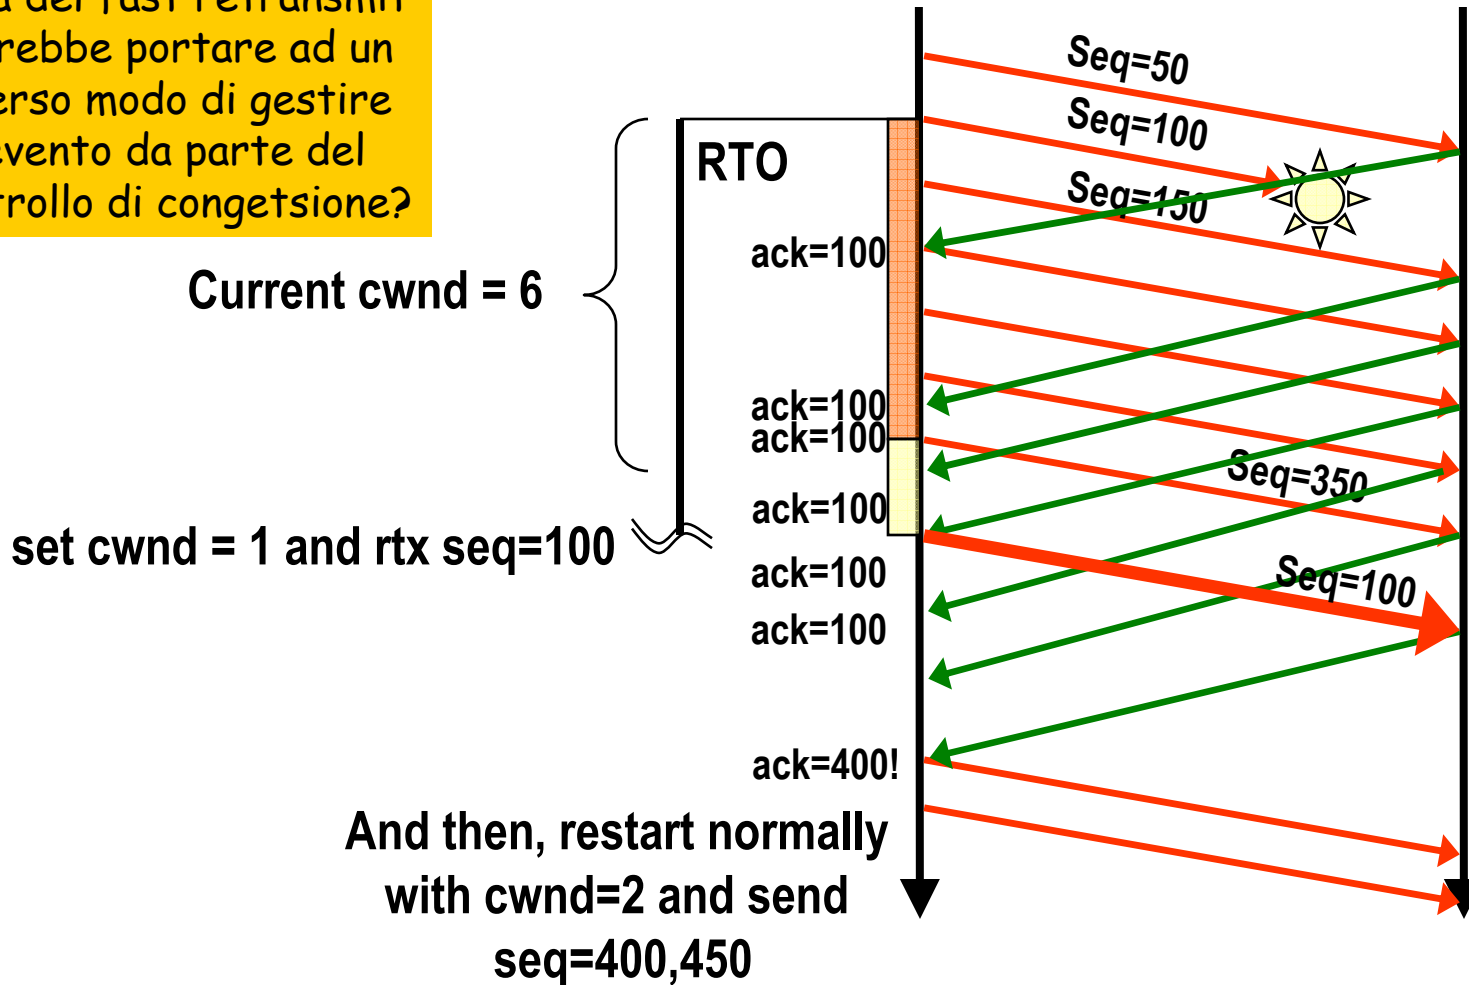
# What happens AFTER RTO? (without fast retransmit)



And then, restart normally with **cwnd=2** and send **seq=400,450**

# TCP RENO (with fast retransmit)

Idea del fast retransmit  
Dovrebbe portare ad un  
Diverso modo di gestire  
L'evento da parte del  
Controllo di conegtsione?

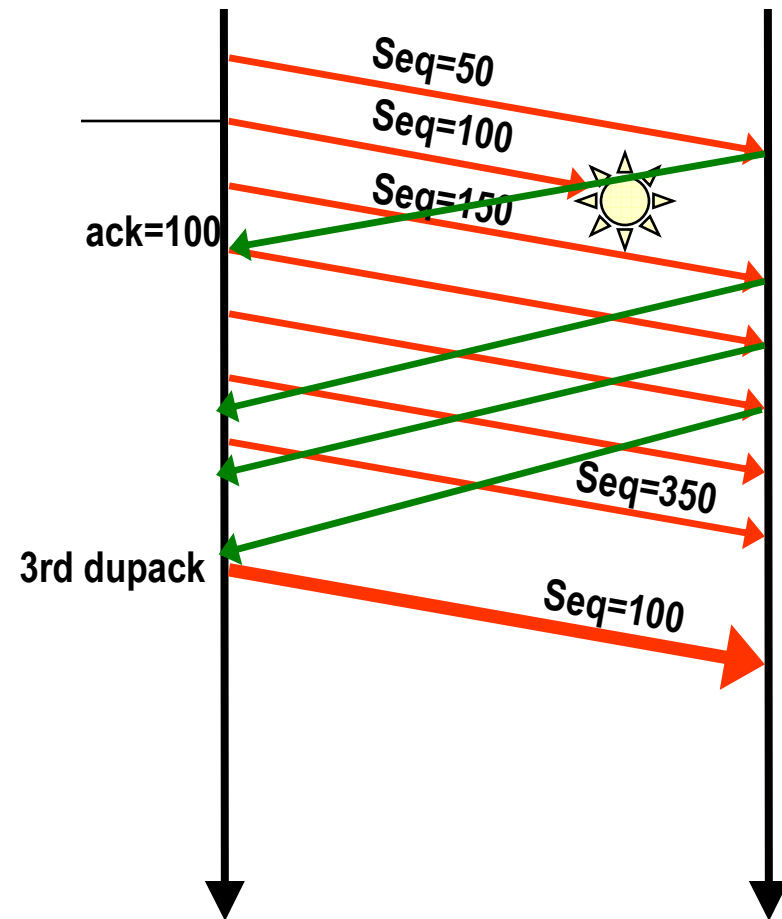


*Same as before, but shorter time to recover packet loss!*

# Motivations for fast recovery

## FAST RECOVERY:

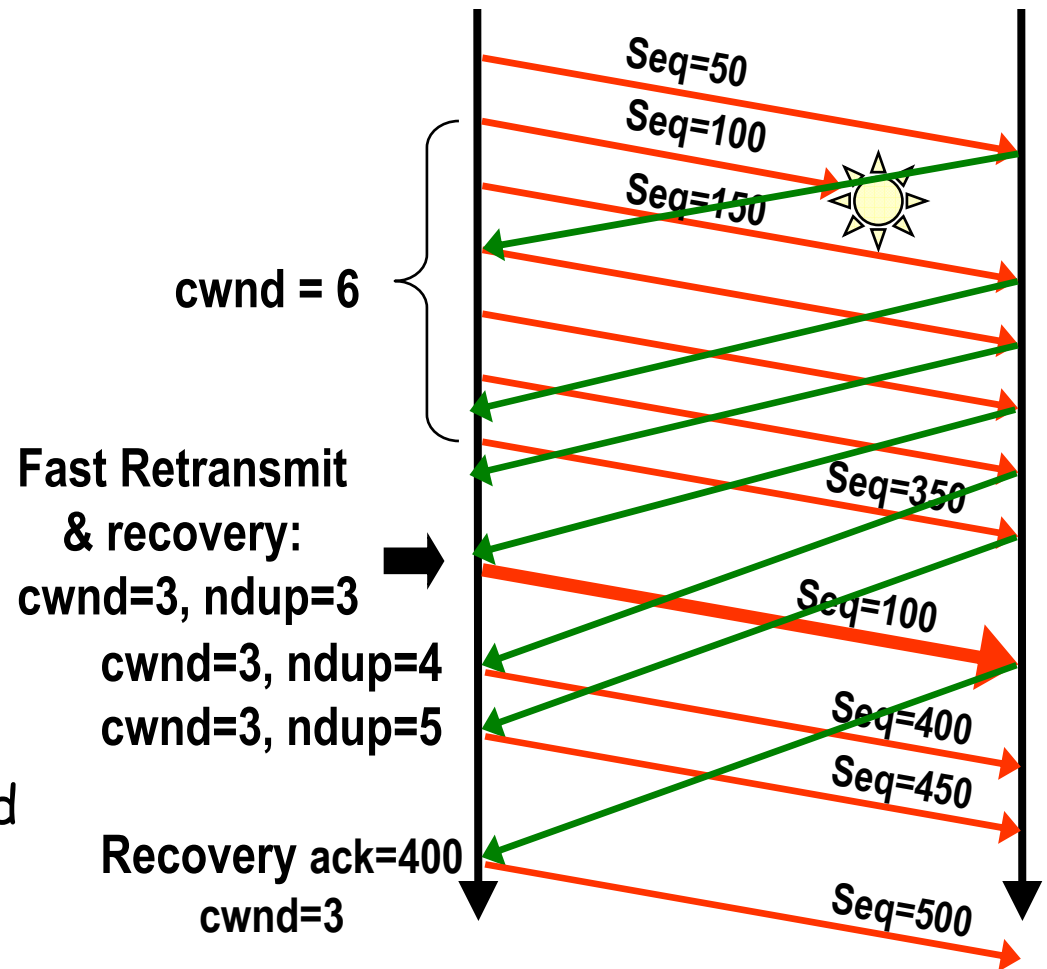
- ⇒ The phase following fast retransmit (3 duplicate acks received)
- ⇒ TAHOE approach: slow start, to protect network after congestion
- ⇒ However, since subsequent acks have been received, no hard congestion situation should be present in the network: slow start is a too conservative restart!



# Fast recovery rules


## FAST RECOVERY RULES:

- ⇒ Retransmit lost segment
- ⇒ **Set  $cwnd = cwnd/2$**
- ⇒ **Restart with congestion avoidance (linear)**
- ⇒ start fast recovery phase:
  - ⇒ Set counter for duplicate packets  $ndup=3$
  - ⇒ Use "inflated" window:  
 $w = cwnd + ndup$
  - ⇒ Upon new  $dup\_acks$ , increase  $ndup$ , not  $cwnd$  (and send new data)
  - ⇒ Upon recovery ack, "deflate" window setting  $ndup=0$



# Idle periods


- After a long idle period (exceeding one RTO), reset the congestion window to one.



# Further TCP issues

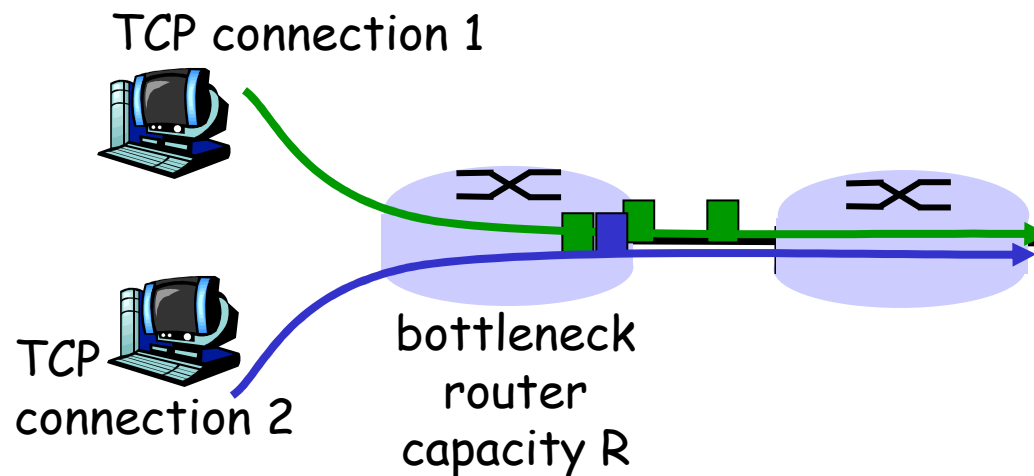
Timeout = packet loss occurrence in an internal network router  
TCP (both Tahoe & Reno) does not AVOID packet loss  
Simply REACTS to packet loss

*CONCLUSION: a TCP able to AVOID packet loss should be much better.....*



# TCP Fairness


**Fairness goal:** if  $K$  TCP sessions share same bottleneck link of bandwidth  $R$ , each should have average rate of  $R/K$



# Why is TCP fair?

Two competing sessions:

- r Additive increase gives slope of 1, as throughput increases
- r multiplicative decrease decreases throughput proportionally






# Fairness with UDP traffic

- r A serious problem for TCP
  - m in heavy network load, TCP reduces transmission rate. Non congestion-controlled traffic does not.
  - m Result: in link overload, TCP throughput vanishes!

*This is why we still live in a World Wide Wait time  
(Webcams are destroying TCP traffic)*

# Mixing TCP & UDP traffic



# Fairness (more)

## Fairness and UDP

- r Multimedia apps often do not use TCP
  - m do not want rate throttled by congestion control
- r Instead use UDP:
  - m pump audio/video at constant rate, tolerate packet loss
- r Research area: TCP friendly

## Fairness and parallel TCP connections

- r nothing prevents app from opening parallel connections between 2 hosts.
- r Web browsers do this
- r Example: link of rate  $R$  supporting 9 connections;
  - m new app asks for 1 TCP, gets rate  $R/10$
  - m new app asks for 11 TCPs, gets  $R/2$  !