

Roadmap

- TCP congestion control - the problem
- TCP congestion control - solutions
- Performance of TCP congestion control
- The future of TCP

TCP congestion control: the problem

- Arguably the most crucial part of TCP
- Lots of research effort (over 25+ years)

Performance Objectives

- Link utilization
- Fairness
- Keep congestion down

Q: What is the right notion of fairness?

TCP congestion control: the problem

Nodes **do not know** about

- other nodes
- total or available capacity on any link
- the number of packets in a router buffer

So, how does a node **sense**

- **congestion?**

delays, losses,.....

- **absence of congestion?**

lower delays, fewer losses.....

Roadmap

- TCP congestion control - the problem
- **TCP congestion control - solutions**
- Performance of TCP congestion control
- The future of TCP

TCP congestion control – solutions

- “Packet loss indicates congestion”
- Sensing congestion: timeouts, duplicate ACKs
- Reacting to congestion: drastic decrease of sending rates
- How/when does a sender recover from congestion?
 - Continuous probing
 - Conservative increases


Q: Formula for increase and decrease ?

TCP congestion control – v1

Let's think in terms of rates:

- **Additive** (conservative) increase, **multiplicative** (drastic) decrease of sending rates (AIMD)
- Starting rate: **High? Low?**

Transmission rate



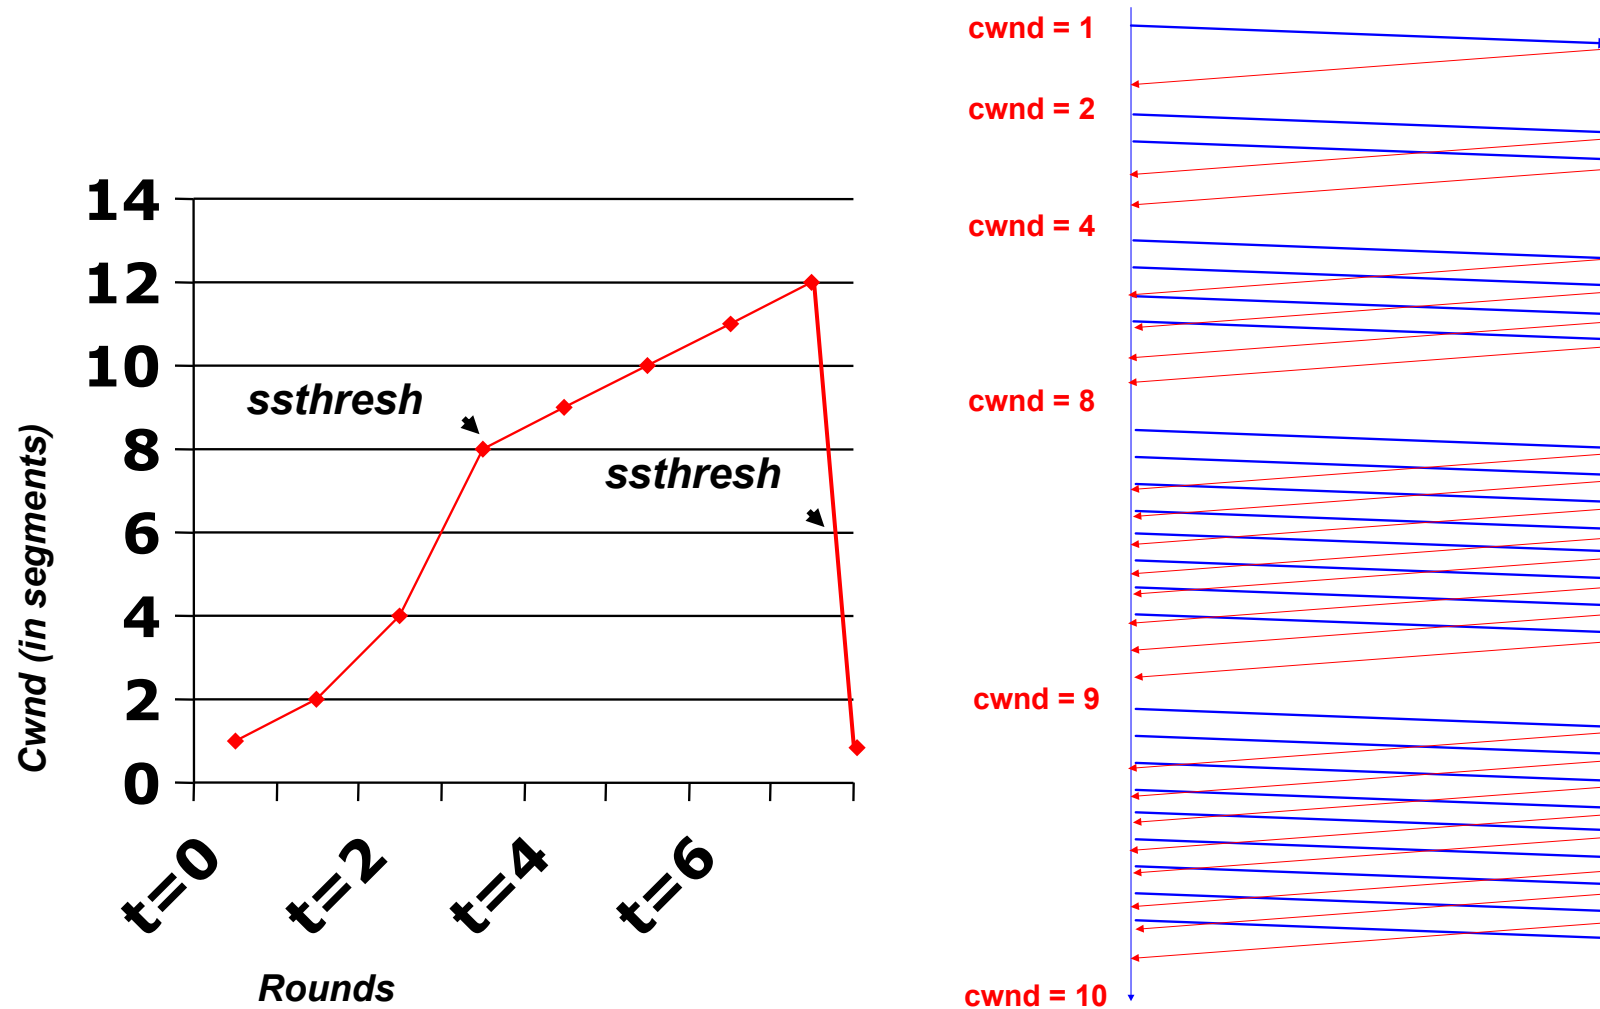
TCP congestion control – v2

TCP Tahoe (1988)

- Two modes: Aggressive probing (**Slow start**) and careful probing (**Congestion avoidance**).
- **Slow(!) start mode**: Multiplicative (aggressive) increase, multiplicative decrease of sending rates.
- **Congestion avoidance mode**: Additive (conservative) increase, multiplicative (drastic) decrease of sending rates.

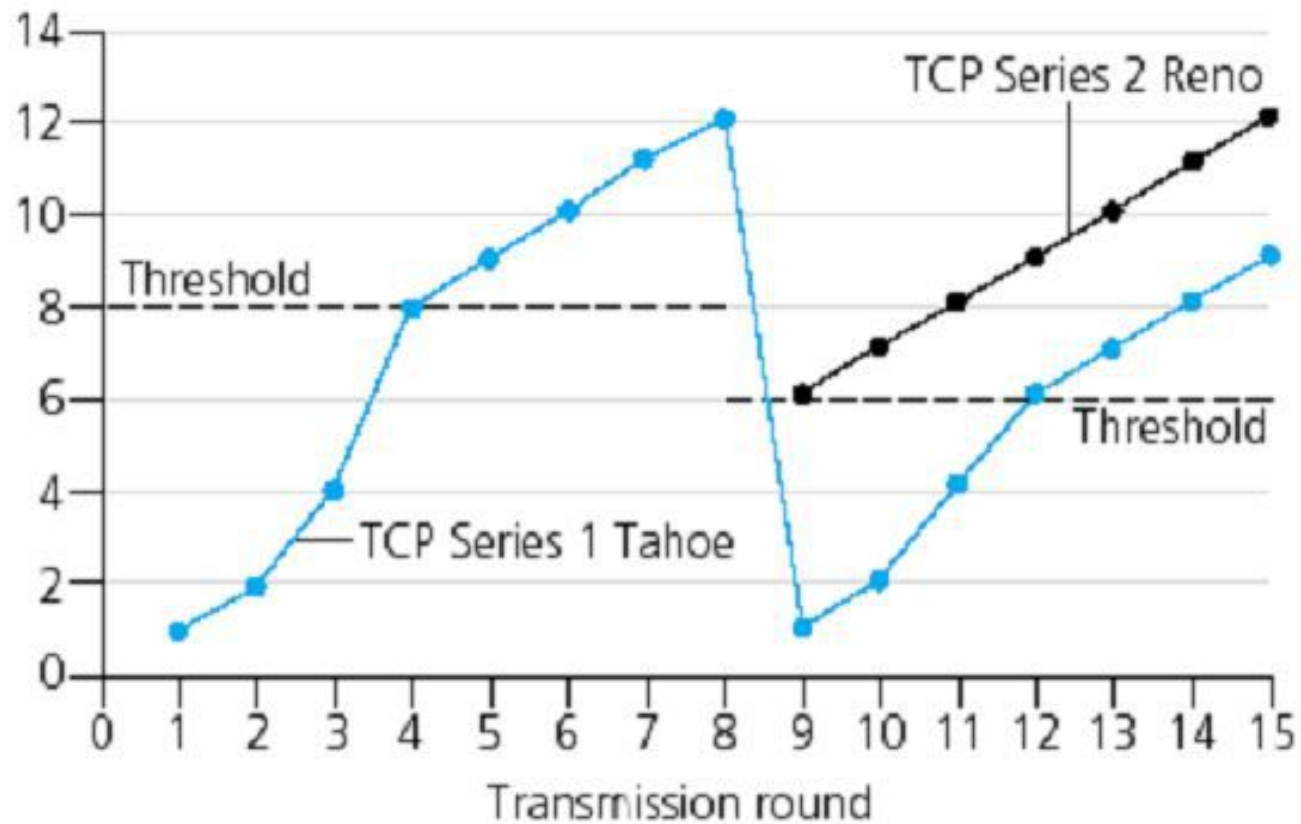
TCP Tahoe illustration

Assume that *ssthresh* = 8



A further improvement

- TCP Reno(1990) and Fast recovery
- Intuition: Distinguish between loss and 3-dup ACK



Windows vs rates

- TCP uses windows (number of packets sent without waiting for ACKs)
- Transmission rate = $W * \text{Packet size} / \text{RTT}$
RTT = round trip time
- Self-clocking

Roadmap

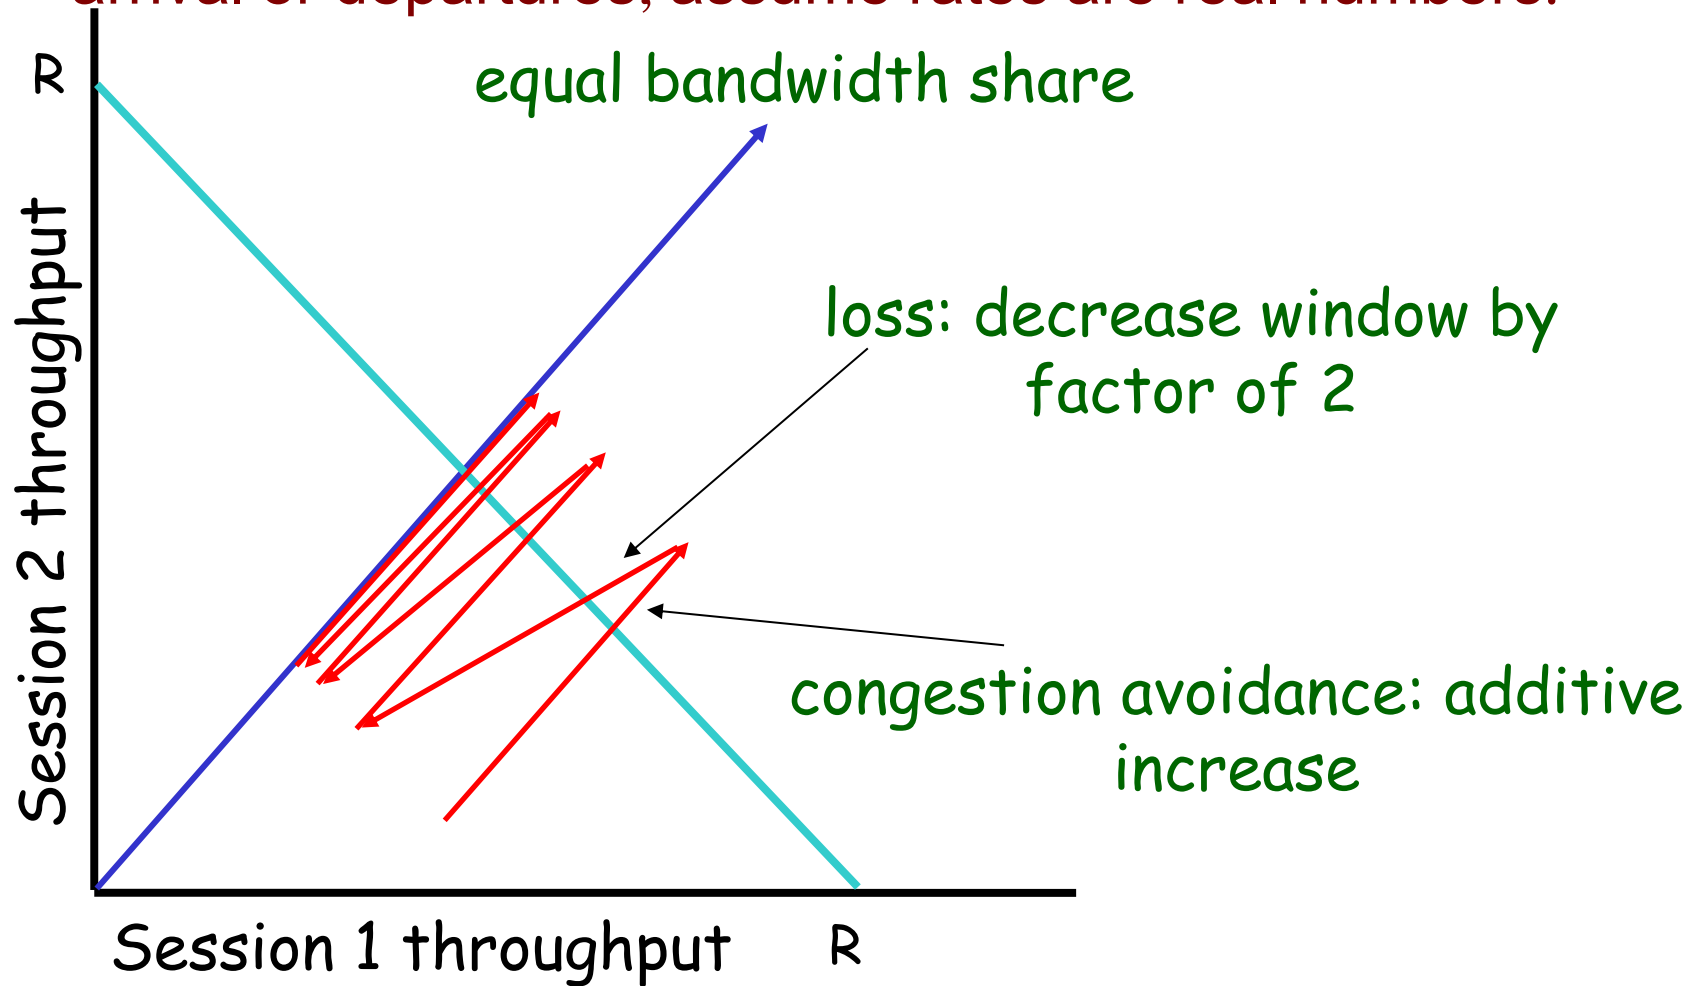
- TCP congestion control - the problem
- TCP congestion control - solutions
- **Performance of TCP congestion control**
- The future of TCP

Performance of TCP congestion control

- Fairness
- Link utilization
- Algorithmic perspective?
- Note: Many other TCP versions have been proposed. No one version outperforms all others in all situations

TCP Performance – Fairness

[Chiu and Jain, 89] Two competing LONG sessions, no arrival or departures, assume rates are real numbers.




Lessons

- Caveat: In reality, both sessions may not detect congestion at the same time
- Once the trajectory hits $x=y$ line, it stays there
- For real-valued rates, reaches $x=y$ only as $t \rightarrow \infty$
- Will *leave* optimal point after reaching it (continuous probing)
- Generalizes to n sessions
- **MIMD, MIAD, AIAD** all unfair
- Can be shown to hold in the presence of arrivals and departures if all sessions are long [Edmonds, Datta, Dymond, 03]

TCP Performance – link utilization

- At least 25% capacity lost

Transmission rate



TCP Performance – algorithmic view

- How good an algorithm is it?
i.e. **Optimal? Near-optimal?**
- Very difficult question
 - Does not correspond to clean theoretical models.
- [Edmonds, Datta, Dymond, 03] TCP “performs well” (is competitive against a limited adversary) when each session is of a minimum specified size...

TCP Performance – problems

When/where does TCP not work well?

- **Multimedia** networking
 - Non-smooth (sawtooth) bit rate
 - Best-effort: no QoS guarantees
- **Wireless/hybrid** networks
 - “Packet loss indicates congestion” ?
 - Congestion and loss are not always correlated
- **Very high-speed** networks
- **Fairness** issues: parallel sessions, multiple bottlenecks, small sessions...

TCP/IP Design issues

- Co-operative algorithm
- Policing is hard to do
- Tracing of malicious hosts/users difficult

Roadmap

- A 2-minute introduction to the TCP/IP architecture
- The role of TCP
- TCP congestion control - the problem
- TCP congestion control - solutions
- Performance of TCP congestion control
- **The future of TCP**

The Future of TCP

- Active queue management schemes
 - Drop packets proactively to reduce congestion
 - RED, REM, BLUE,...
- TCP for ultra-high speed networks
 - FastTCP, HSTCP,....
 - Usually use delay for being more responsive to congestion
 - Less draconian
- TCP for wireless, hybrid networks
 - Treats wireless parts separately