

TCP congestion control

Recall:

$$\text{EffectiveWindow} = \text{MaxWindow} - (\text{LastByteSent} - \text{LastByteAcked})$$

where

$$\text{MaxWindow} = \min\{\text{AdvertisedWindow}, \text{CongestionWindow}\}$$

Key question: how to set `CongestionWindow` which, in turn, affects ARQ's sending rate?

→ linear increase/exponential decrease

→ AIMD

TCP congestion control components:

(i) Congestion avoidance

→ linear increase/exponential decrease

→ additive increase/exponential decrease (AIMD)

As in Method B, increase `CongestionWindow` linearly,
but decrease exponentially

Upon receiving ACK:


$$\text{CongestionWindow} \leftarrow \text{CongestionWindow} + 1$$

Upon timeout:

$$\text{CongestionWindow} \leftarrow \text{CongestionWindow} / 2$$


But is it correct...

“Linear increase” time diagram:



→ results in exponential increase

What we want:



→ increase by 1 every window

Thus, linear increase update:

$$\begin{aligned} \text{CongestionWindow} &\leftarrow \text{CongestionWindow} \\ &\quad + (1 / \text{CongestionWindow}) \end{aligned}$$

Upon timeout and exponential backoff,

$$\text{SlowStartThreshold} \leftarrow \text{CongestionWindow} / 2$$

(ii) Slow Start

Reset `CongestionWindow` to 1

Perform exponential increase

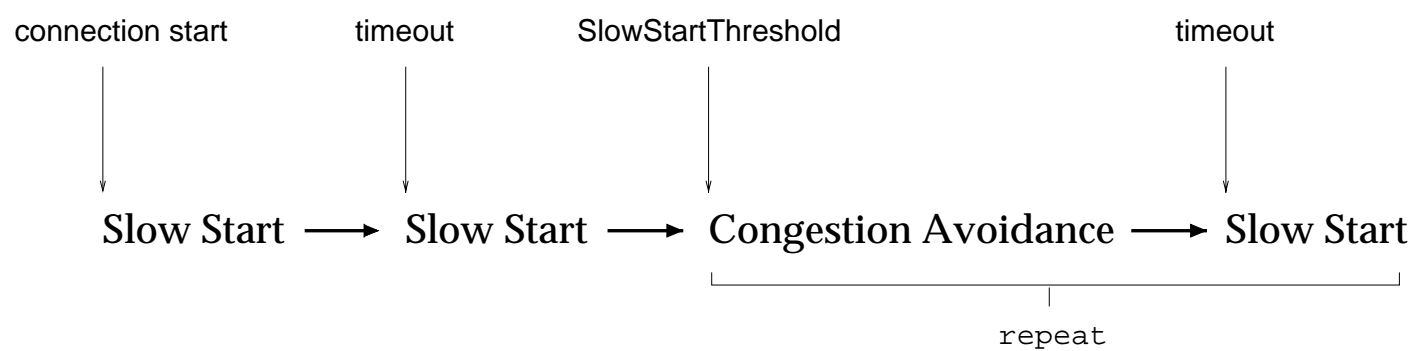
$\text{CongestionWindow} \leftarrow \text{CongestionWindow} + 1$

- Until timeout at start of connection
 - rapidly probe for available bandwidth
- Until `CongestionWindow` hits `SlowStartThreshold` following Congestion Avoidance
 - rapidly climb to safe level
 - “slow” is a misnomer
 - exponential increase is super-fast


Basic dynamics:

→ after connection set-up

→ before connection tear-down



CongestionWindow evolution:



(iii) Exponential timer backoff

$\text{TimeOut} \leftarrow 2 \cdot \text{TimeOut}$ if retransmit

(iv) Fast Retransmit

Upon receiving three duplicate ACKs:

- Transmit next expected segment
 - segment indicated by ACK value
- Perform exponential backoff and commence Slow Start
 - three duplicate ACKs: likely segment is lost
 - react before timeout occurs

TCP Tahoe: features (i)-(iv)

(v) Fast Recovery

Upon Fast Retransmit:

- Skip Slow Start and commence Congestion Avoidance
→ dup ACKs: likely spurious loss
- Insert “inflationary” phase just before Congestion Avoidance

TCP Reno: features (i)-(v)

→ pre-dominant form

Many more versions of TCP:

→ NewReno w/ SACK, w/o SACK, Vegas, etc.

→ wireless, ECN, multiple time scale

→ mixed verdict; pros/cons

Given sawtooth behavior of TCP's linear increase/exponential backoff:

Why use exponential backoff and not Method D?

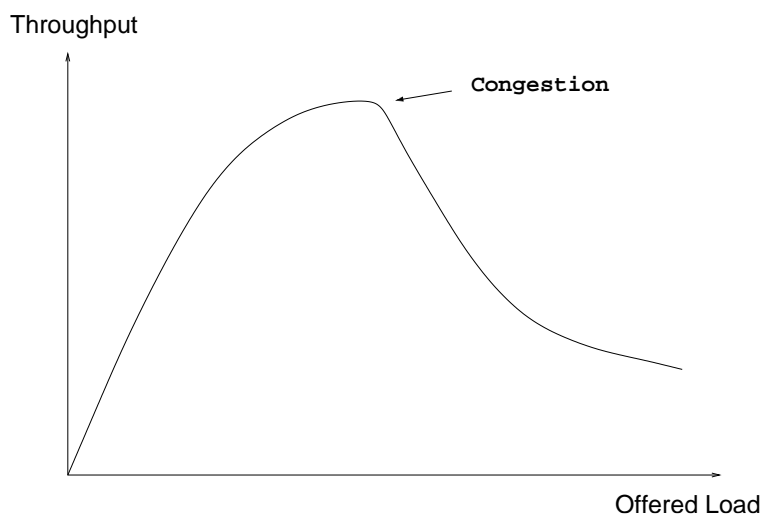
- For multimedia streaming (e.g., pseudo real-time), AIMD (Method B) is not appropriate
→ use Method D
- For unimodal case—throughput decreases when system load is excessive—story is more complicated
→ asymmetry in control law needed for stability

Congestion Control: Selfishness, Stability and Optimality

- to be, or not to be, selfish ...
- noncooperative game theory
- John von Neumann, John Nash, ...

Congestion and “tragedy of commons”:

- Garrett Hardin, '68



- if everyone acts selfishly, no one wins
 - in fact, everyone loses
- can this be prevented?

Two-party congestion control setting:

- Prisoner's Dilemma game
- both cooperate (stay silent): 1 year each
- both selfish (rat on the other): 5 years each
- one cooperative/one selfish: 9 vs. 0 years

When cast as congestion control game:

		<i>Bob</i>	
		C	N
<i>Alice</i>	C	5, 5	1, 9
	N	9, 1	3, 3

- (a, b) : throughput (Mbps) achieved by Alice/Bob
- what may happen?
- what do “rational” (w.r.t. selfishness) players do?

Outcome of game with cooperative players?

- configuration (C,C) with payoff (5,5)
- system optimal: $5 + 5 = 10$ (sum of payoffs)
- note: (1,9) and (9,1) are also system optimal
- also Pareto optimal

Def. (Pareto optimality): A system state or configuration is Pareto optimal if total system payoff cannot be improved without sacrificing one (or more) player's payoff.

- improvement requires “sacrificial lamb”
- welfare notion of overall goodness
- note: system optimal \Rightarrow Pareto optimal (trivial)
- (5,5), (1,9), (9,1): Pareto optimal
- (3,3): not Pareto optimal

Outcome of game with noncooperative (i.e., selfish) players?

- (N,N) with payoff (3,3)
- notion of stability: Nash equilibrium

Def. (Nash equilibrium): A configuration is a Nash equilibrium (NE) if no selfish player has an incentive to unilaterally change his/her action.

- (N,N) with payoff (3,3) is NE
- Alice, alone, changing N to C: $(N,N) \mapsto (C,N)$
- (C,N) has payoff (1,9): bad for Alice
- Nash equilibrium is a rest point
- i.e., stable fixed-point (equilibrium)
- idea: due to John Nash
- key contribution: dynamics under selfishness

Is congestion control game NE (3,3) system optimal?

- no: (1,9) and (9,1): total payoff 10 (vs. 6)
- in fact: system optimal (5,5) is better for both
- in general, NE need not be system optimal
- also NE need not be Pareto optimal

Puts a damper on Adam Smith's postulate:


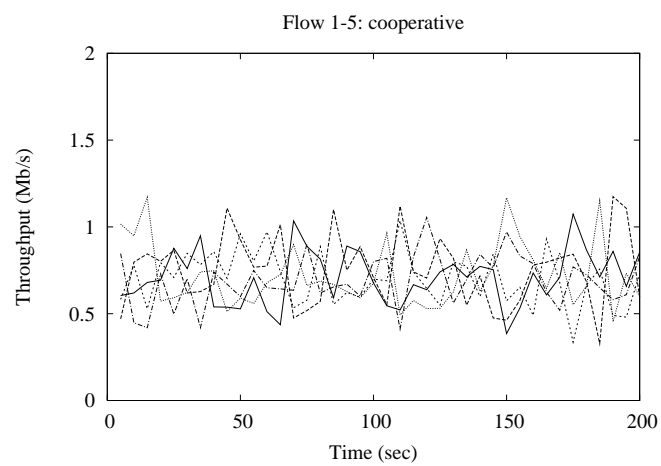
- wise, efficient "invisible hand" (i.e., "market")
- economy of selfish users self-organizes efficiently
- rarely true: Achilles' heel of "pure" capitalism

Karl Marx & communism?

- fantasy & wishful thinking
- evolution (hereto) has put premium on selfishness
- vulnerable to selfish elements
- Marx & Confucius: both more harm than good?



5 regular (cooperative) TCP flows:

→ share 11 Mbps WLAN bottleneck link



4 regular (cooperative) TCP flows and 1 noncooperative TCP flow:

→ same benchmark set-up



Remarks:

- NE, in general, are neither efficient nor fair
 - \exists special cases: strong rules/penalties
- in fact, in general, a Nash equilibrium need not exist
 - system subject to oscillation
 - circular “chain reaction”
- Nash’s main result (game theory): finite noncooperative games with mixed strategies—choose action probabilistically—always possess equilibrium
 - vs. pure strategy (more in tune with reality)
 - pure strategy games: hard problem
- congestion pricing
 - penalize those who congest: e.g., usage pricing
 - in the States: flat pricing (dominant)
 - not skimpy like the rest of the world!

- repeated/evolutionary games
 - e.g., iterated Prisoner's Dilemma
 - rob bank/get caught, again and again ...
 - what should the prisoners do then?
 - tit-for-tat, grim trigger: can be optimal
 - most relevant for greedy TCP