Fairness & stability of TCP with a single bottleneck link

- consider a bottleneck link of capacity c, all other links have plentiful capacity
- intuitively, a congestion-control mechanism is fair if on average each connection gets an equal share of the bottleneck link bandwidth
- graphical illustration of TCP evolution for 2 connections with same MSS & RTT

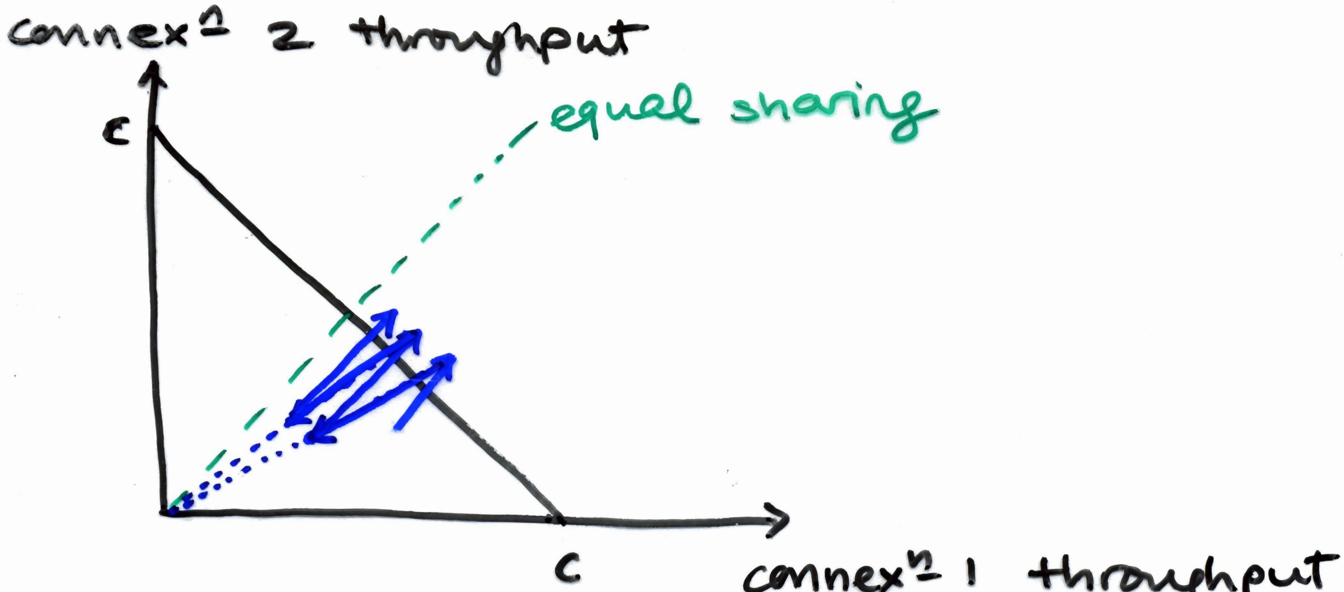

- the throughput of the 2 connections eventually fluctuates along the equal sharing line, regardless of starting point

- however, if the TCP connections have different RTTs, since the additive increase rate is 1 packet per RTT, connections with longer RTTs ↑ their window more slowly

Network-assisted rate-based control — ATM's Available Bit Rate (ABR) Congestion control

- ABR is 1 of 5 service classes in ATM
  - elastic data transfer service (like TCP)
- resource management (RM) cells used to convey congestion information
  - interspersed with data cells
  - sent from source to destination, turned around at destination & sent back to source
- based on info in returned RM cells, source explicitly calculates its maximum sending rate via complicated rules (Jain 96)
  - each RM cell contains a 2-byte explicit rate (ER) field — a congested switch can ↓ the value in a passing RM cell, setting the field to the minimum supportable rate of all switches on the source-destination path
  - each RM cell also contains a no increase (NI) bit & a congestion indication (CI) bit

- each data cell contains an <u>explicit forward congestion indication</u> (EFCI) bit
- a switch indicates mild congestion by setting the NI bit of an RM cell, & more serious congestion by setting the CI bit of an RM cell or the EFCI bit of a data cell
- destination sets the CI bit of an RM cell if the preceding data cell's EFCI bit was set

Bandwidth sharing & fairness in multi-hop networks

- when a single link is the only limiting resource, equal sharing is efficient & intuitively fair
- In a general multi-hop network, requiring equal rates for each session may not result in full bandwidth utilization, & other notions of fairness may be desirable
- Eg

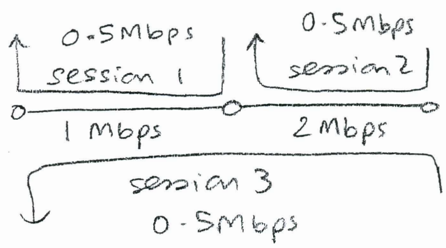

  - each session gets the same rate
  - however, the 2Mbps link has 1Mbps of idle capacity; allocating this to session 2 ↑s its rate to 1.5Mbps, while not making any other session worse off — such a change is a <u>Pareto improvement</u> over the original allocation
- <u>Defn</u>: An allocation of resources is <u>Pareto efficient</u> if there does not exist another allocation in which some individual is better off & no individual is worse off
- Eg

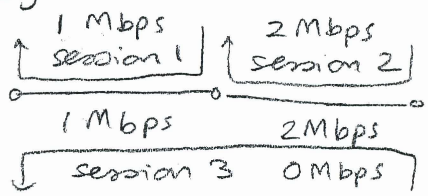

  - this allocation maximizes the sum of the 3 sessions' rates
  - but it is not fair to session 3
- Consider for simplicity a static network, where the set of sessions & the route taken by each session is fixed
  - each session has an infinite amount of data to send (<u>infinite backlog</u> assumption)
- Define
  - $\mathcal{L}$: set of directed links
  - $\underline{c} = (c_1, c_2, \ldots, c_{|\mathcal{L}|})$: vector of link capacities available for allocation (if there are dynamics eg. from dynamic congestion control algorithms or bursty sources, this should be some function of the actual capacity that is strictly less than the mean)
  - S: set of sessions, each with a fixed route
  - $\mathcal{L}_\sigma$: set of links used by session $\sigma \in S$
  - S_ℓ: set of sessions using link $\ell \in \mathcal{L}$

- $\underline{r} = (r_1, \dots, r_{|S|})$: vector of rates allocated to each session
- $f_\ell(\underline{r}) = \sum_{\sigma \in S_\ell} r_\sigma$: total flow through link $\ell \in \mathcal{L}$ under rate allocation $\underline{r}$
- $(\mathcal{L}, C, S)$ defines an instance of the bandwidth-sharing problem

• A rate vector $\underline{r} \geq 0$ is feasible for $(\mathcal{L}, C, S)$ if

$$f_\ell(\underline{r}) \leq C_\ell \quad \forall \ \ell \in \mathcal{L} \text{ — (1)}$$

& is Pareto efficient if for every session σ there is a link $\ell \in \mathcal{L}_S$ s.t.

$$f_\ell(\underline{r}) = C_\ell$$

• (1) gives a system of $|\mathcal{L}|$ linear inequalities, which defines a convex polyhedral set in the nonnegative orthant of $\mathbb{R}^{|S|}$

• Max-Min fair bandwidth sharing

- Def^{n}: A feasible rate vector is max-min fair (MMF) for $(\mathcal{L}, C, S)$ if it is not possible to increase the rate of any session σ while maintaining feasibility & not reducing the rate of another session ρ s.t. $r_\rho \leq r_s$
- Def^{n}: Given a feasible rate vector $\underline{r}$, a link ℓ is a bottleneck link for session j if
  - i) link ℓ is saturated ie $f_\ell(\underline{r}) = C_\ell$, &
  - ii) every session using ℓ has rate no more than r_j i.e. $r_\sigma \leq r_j \ \forall \ \sigma \in S_\ell$
- Proposition: A feasible rate vector $\underline{r}$ is MMF iff each session has a bottleneck link wrt $\underline{r}$

Proof: Suppose $\underline{r}$ is MMF but $\exists$ a session j with no bottleneck link

→ for each link $\ell \in \mathcal{L}_j$ st. $f_\ell(\underline{r}) = C_\ell$, $\exists$ some session σ st. $r_\sigma > r_j$

→ for each link $\ell \in \mathcal{L}_j$,

$$\delta_\ell = \begin{cases} C_\ell - f_\ell(\underline{r}) & \text{if } f_\ell(\underline{r}) < C_\ell \\ r_\sigma - r_j & \text{otherwise} \end{cases}$$

$$> 0$$

→ r_ℓ can be ↑'d by $\min_{\ell \in \mathcal{L}_j} \delta_\ell$ without ↓^{g} the rate of any session whose existing rate is at most r_ℓ

→ contradiction □