Transport Bandwidth Allocation

10/26/2009

Admin.

- Questions on programming assignment 2?
- Date of exam 1?

Recap: TCP/Reno Queueing Dynamics

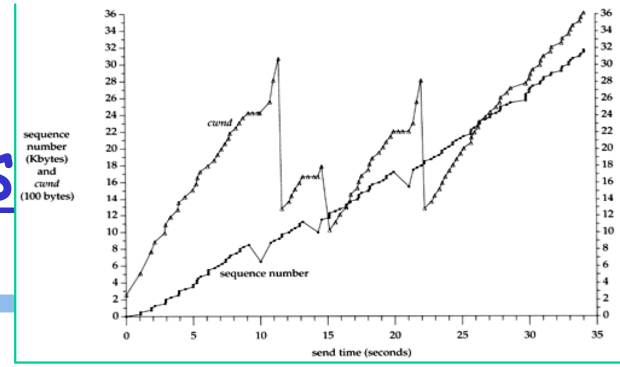

- Consider congestion avoidance only

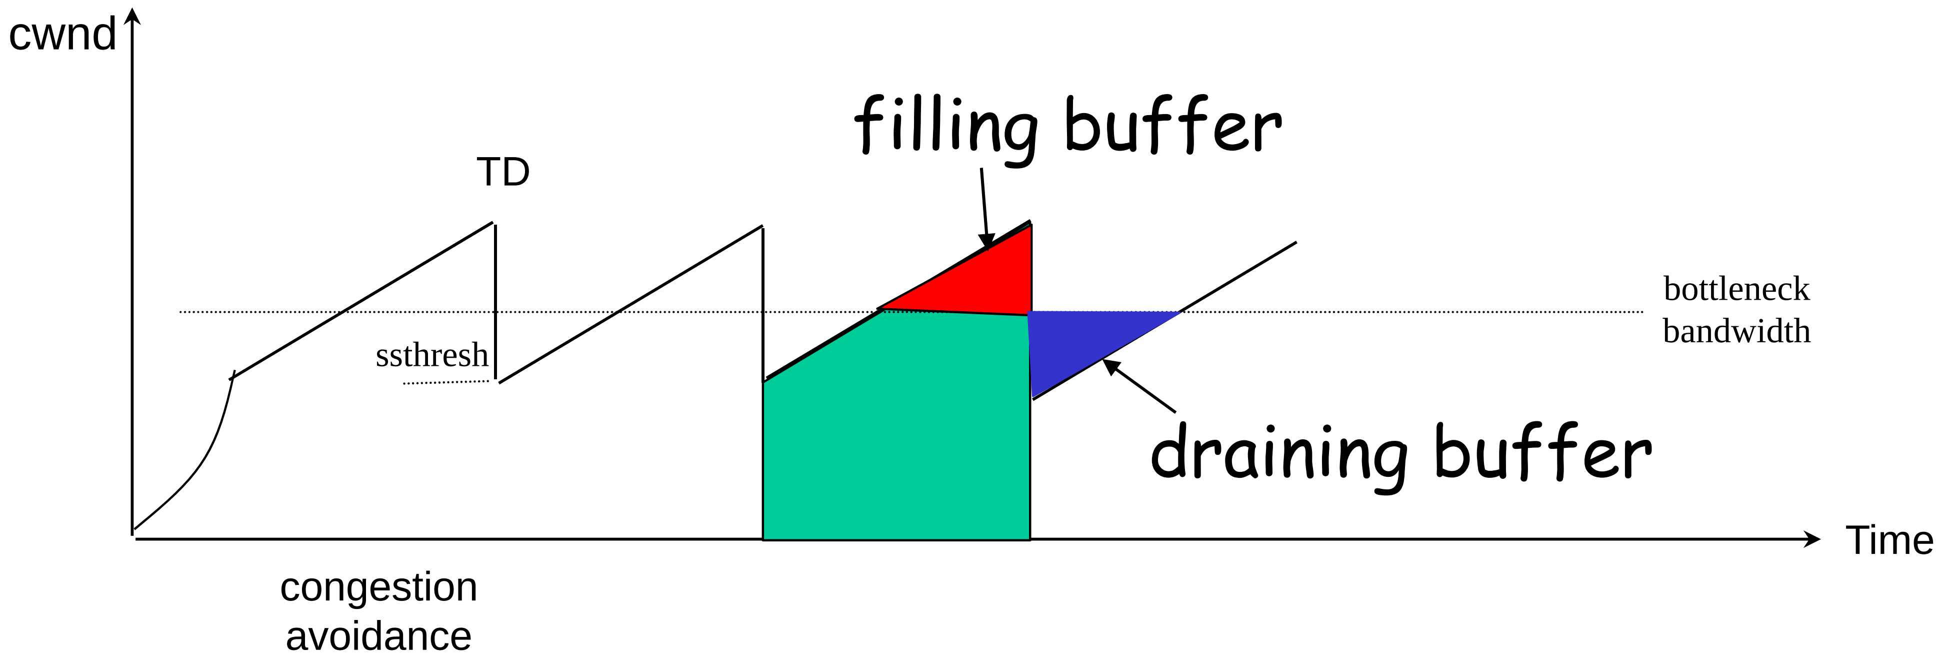

There is a filling and draining of buffer process for each TCP flow.

Recap: TCP/Reno Throughput Modeling: Relating W with Loss Rate p

- Consider congestion avoidance only

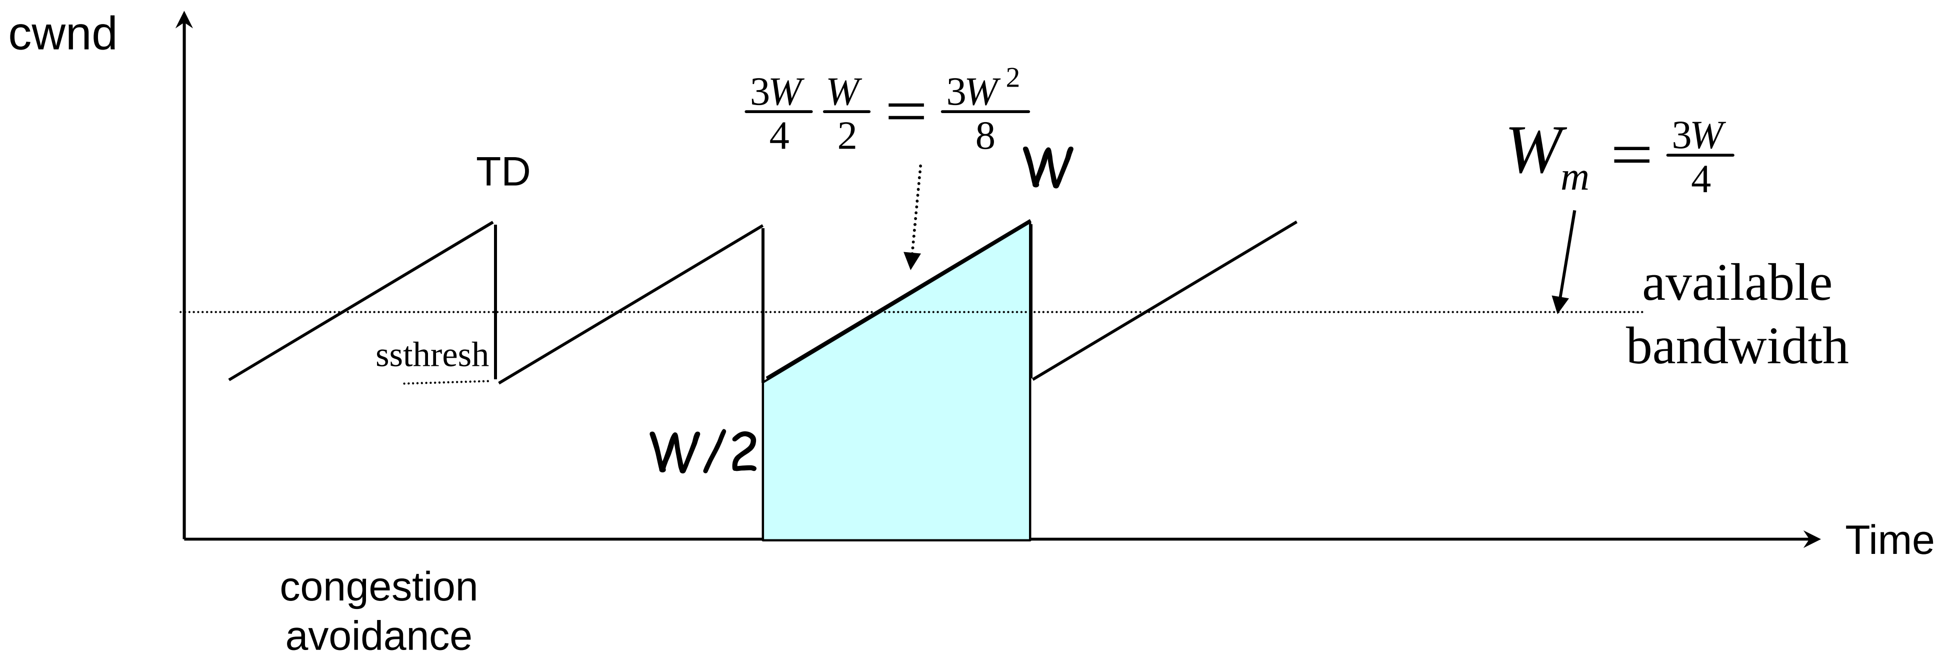

Assume one packet loss (loss event) per cycle
Total packets send per cycle = (W/2 + W)/2 * W/2 = $3W^2/8$
Thus p = $1/(3W^2/8) = 8/(3W^2)$

$$W = \frac{\sqrt{8/3}}{\sqrt{p}} = \frac{1.6}{\sqrt{p}} \quad \Rightarrow throughput = \frac{S}{RTT}\frac{3}{4}\frac{1.6}{\sqrt{p}} = \frac{1.2S}{RTT\sqrt{p}}$$

Recap: TCP/Reno Throughput Modeling

$$\Delta W = \begin{cases} \frac{1}{W} & \text{if the packet is not lost} \\ -\frac{W}{2} & \text{if packet is lost} \end{cases}$$

$$mean\ of\ \Delta W = (1-p)\frac{1}{W} + p(-\frac{W}{2})$$

=> $$mean\ of\ W = \sqrt{\frac{2(1-p)}{p}} \approx \frac{1.4}{\sqrt{p}}, \text{ when } p \text{ is small}$$

=> $$throughput \approx \frac{1.4S}{RTT\sqrt{p}}, \text{ when } p \text{ is small}$$

Recap: TCP/Vegas CA algorithm

maintain a *constant* number of packets in the bottleneck buffer

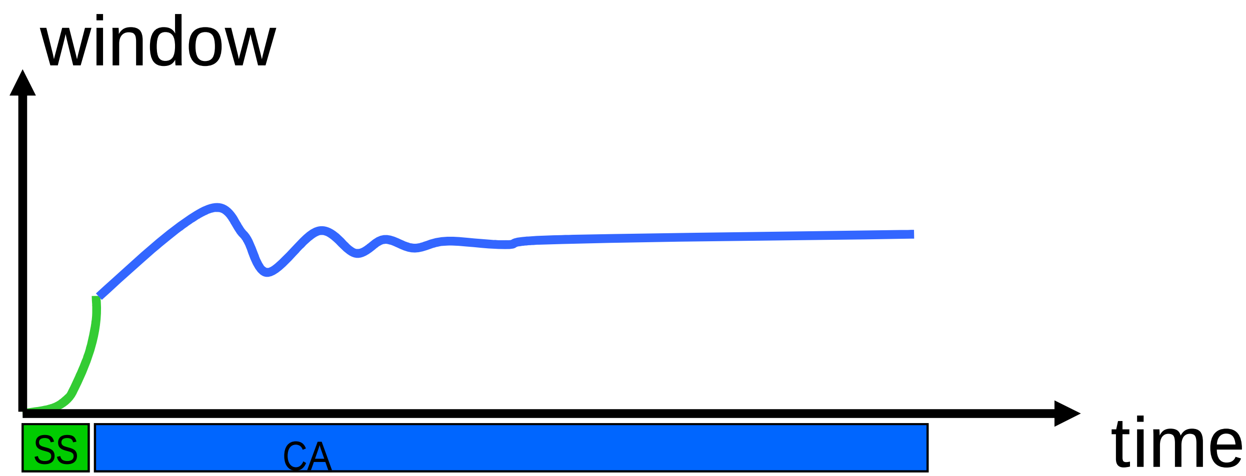


```
for every RTT
{    if  W – W/RTT RTT_min < α  then W ++
     if  W – W/RTT RTT_min > α  then W --
}
for every loss
        W := W/2
```

queue size

TCP/Vegas Dynamics

$$\Delta w_{\mathrm{RTT}} \approx -(w - xRTT_{\min} - \alpha)$$

$$\Delta w_{\text{unit-time}} = -\left(\frac{w}{RTT} - \frac{x}{RTT} RTT_{\min} - \frac{\alpha}{RTT}\right) = \frac{x}{RTT} RTT_{\min} + \frac{\alpha}{RTT} - x$$

$$\Delta x = \frac{\Delta w_{\text{unit-time}}}{RTT} = \frac{x}{RTT^2}\left(RTT_{\min} + \frac{\alpha}{x} - RTT\right)$$

TCP/Reno vs. TCP/Vegas

	TCP/Reno	TCP/Vegas
Congestion signal	loss rate p	queueying delay $T_{queueing}$
Dynamics	$\Delta x = \frac{1}{RTT^2} - p\frac{1}{2}x^2$	$\Delta x = \frac{x}{RTT^2}(RTT_{\min} + \frac{\alpha}{x} - RTT)$
Equilibrium	$x_{reno} = \frac{\alpha_{reno}}{RTT\sqrt{p}}$	$x_{vegas} = \frac{\alpha_{vegas}}{T_{queueing}}$

Discussion: Why and why not TCP/Vegas?

Interpreting Congestion Measure

- A congestion measure (loss/delay) is a signal from the network to the flows reflecting congestion

- Another way to think of congestion measure is to think of it as "price"
  - price goes up as the rate to a link is getting close to capacity
  - the higher the "price", the lower the rate

Interpreting Congestion Measure

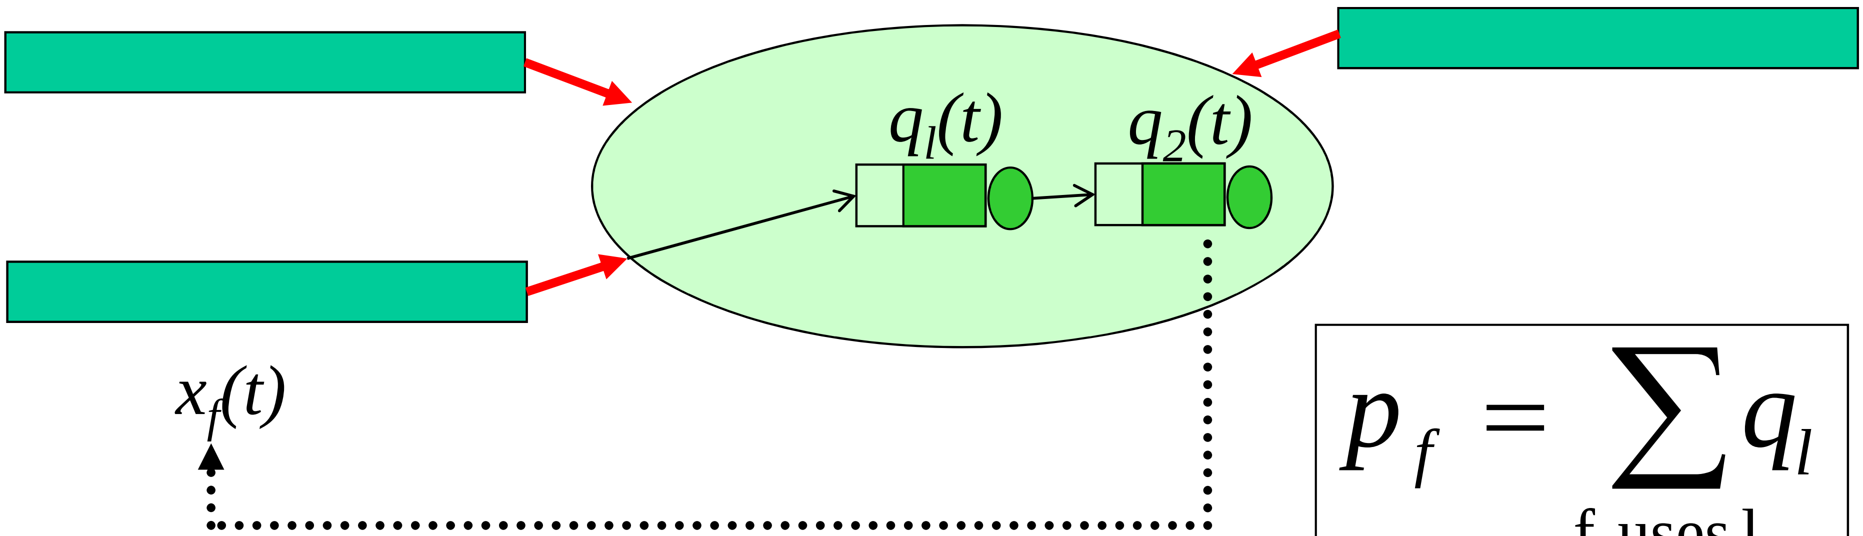

$$p_f = \sum_{\text{f uses l}} q_l$$

TCP/Reno: $\Delta x = \frac{1}{RTT^2} - p\frac{1}{2}x^2 \quad = \frac{1}{2}x^2\left(\frac{2}{RTT^2x^2} - p\right)$

TCP/Vegas: $\Delta x = \frac{x}{RTT^2}(RTT_{\min} + \frac{\alpha}{x} - RTT) \quad = \frac{x}{RTT^2}(\frac{\alpha}{x} - T_{queueing})$

Outline

- Recap
- Network bandwidth allocation framework
  - motivation

Motivation

- So far our discussion is implicitly on a network with a single bottleneck link; this simplifies design and analysis:
  - efficiency/optimality (high utilization)
    - fully utilize the bandwidth of the link
  - fairness (resource sharing)
    - each flow receives an ***equal*** share of the link's bandwidth

Network Resource Allocation

- It is important to understand and design protocols for a general network topology
  - how will TCP allocate resource in a general topology?
  - how should resource be allocated in a general topology?

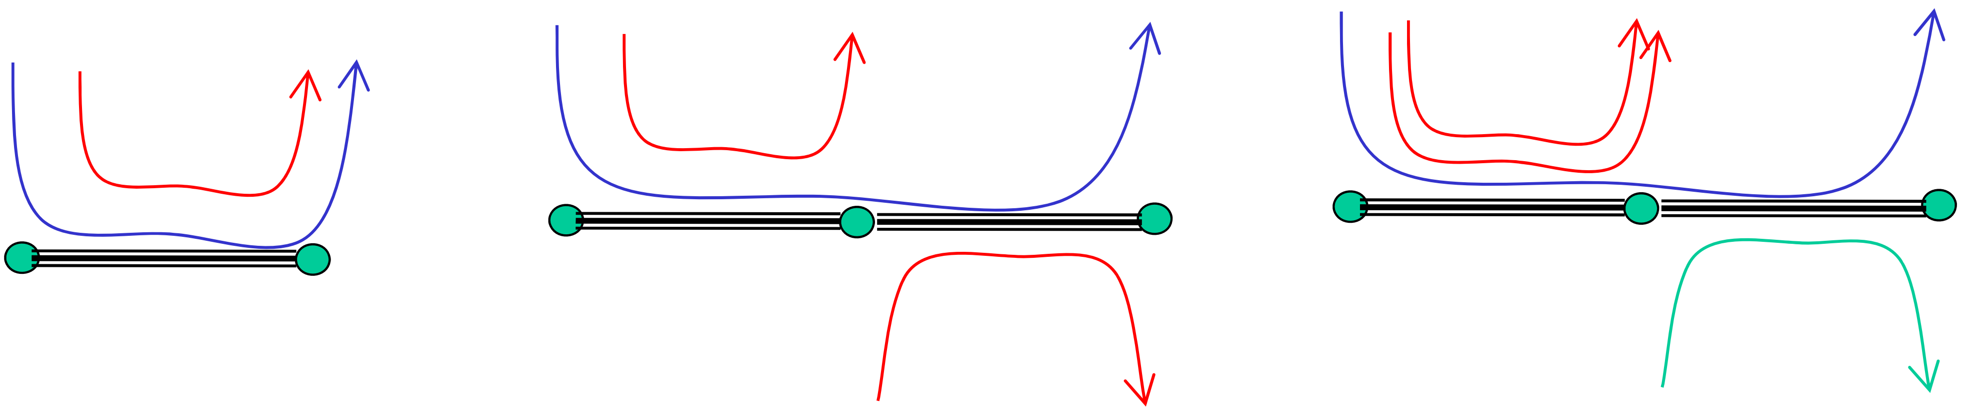

Example: TCP/Reno Rates

- Rate: $x_1 = 0.26$
  $x_2 = x_3 = 0.74$

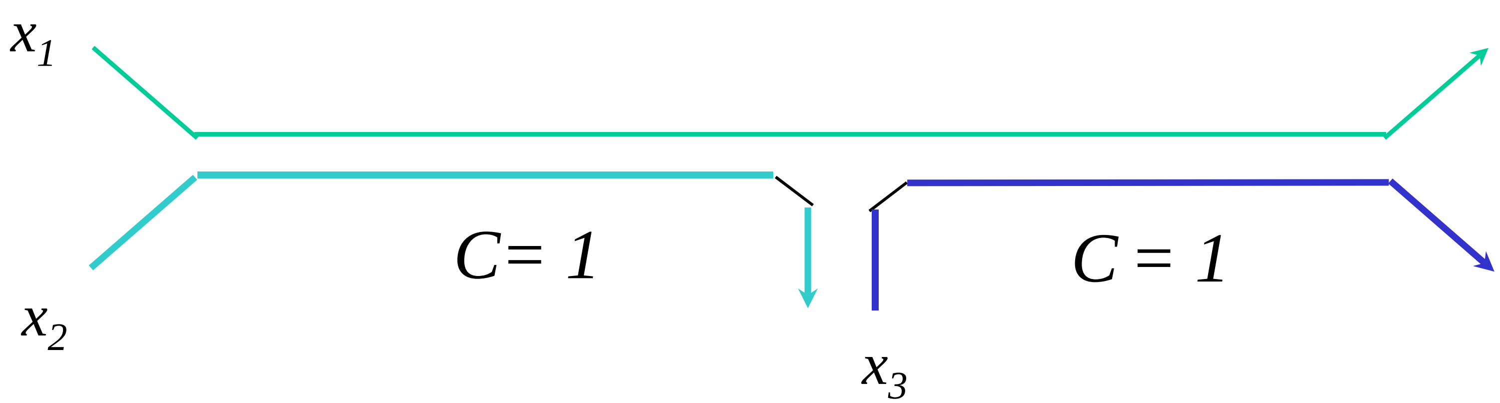

Example: TCP/Vegas Rates

- Rates : $x_1 = 1/3$
  $x_2 = x_3 = 2/3$

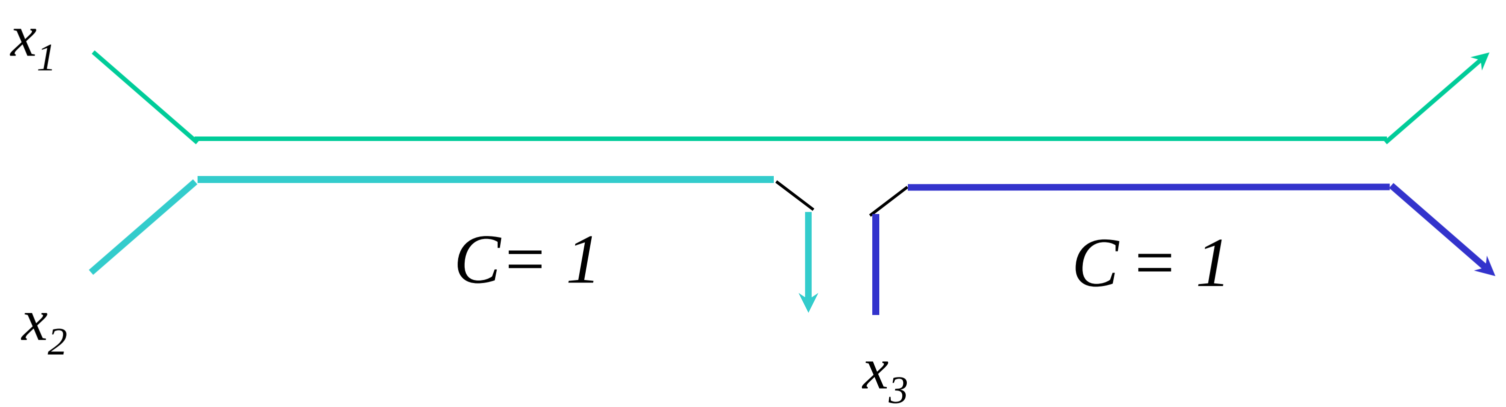

Example: Maximize Throughput

$$\max_{x_f \geq 0} \sum_f x_f$$

subject to

$$x_1 + x_2 \leq 1$$

$$x_1 + x_3 \leq 1$$

- Optimal: $x_1 = 0$
  $x_2 = x_3 = 1$

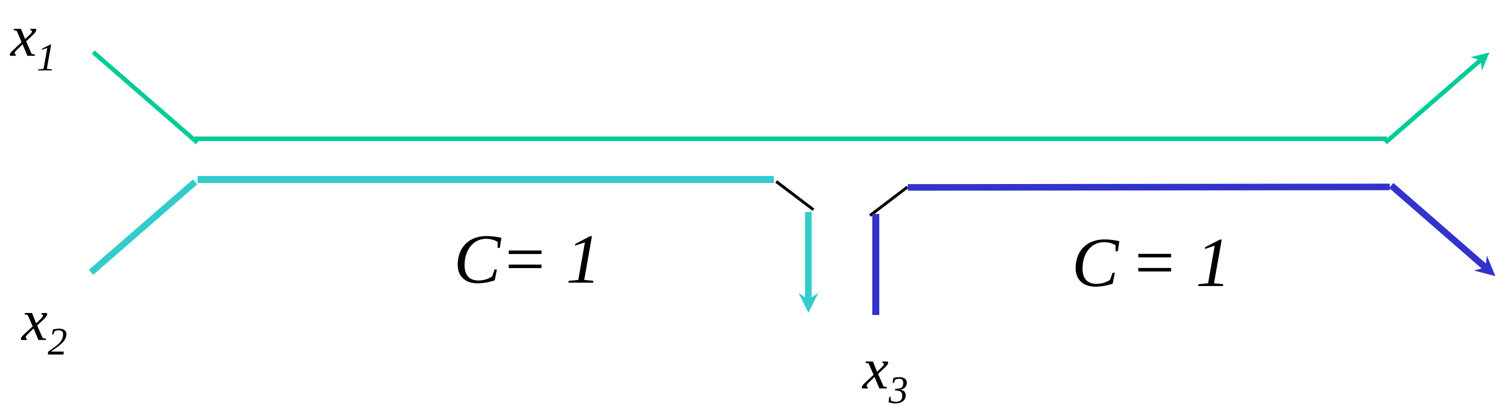

Example: Max-min Fairness

- Max-min fairness: maximizes the throughput of the flow receiving the minimum (of resources)
  - Justification: John Rawls, *A Theory of Justice* (1971)
    - http://en.wikipedia.org/wiki/John_Rawls
  - This is a resource allocation scheme used in ATM and some other network resource allocation proposals

Example: Max-Min

$$\max_{x_f \geq 0} \quad \min \{x_f\}$$

subject to

$$x_1 + x_2 \leq 1$$

$$x_1 + x_3 \leq 1$$

- Rates: $x_1 = x_2 = x_3 = 1/2$

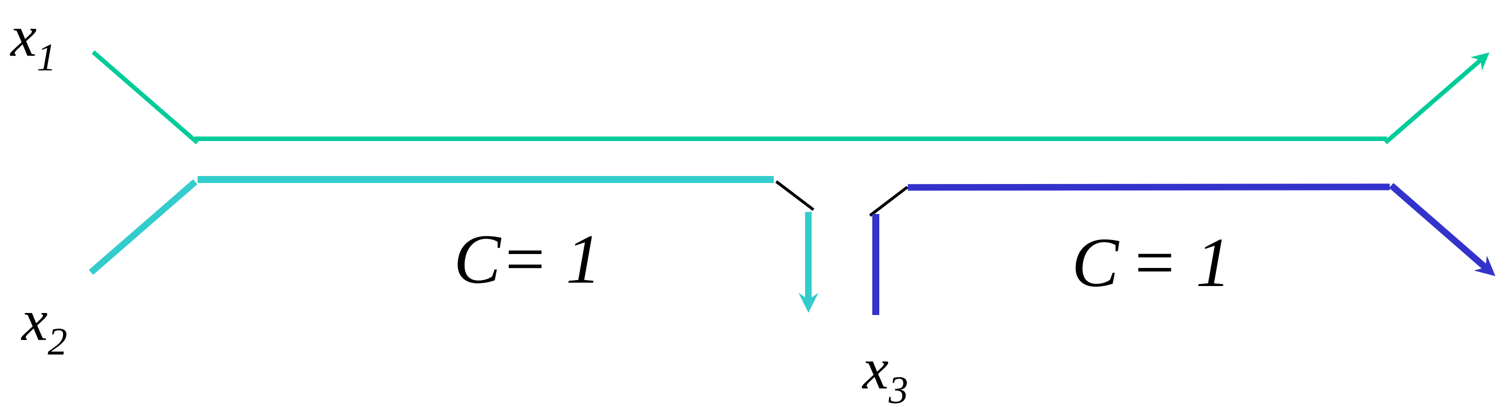

Network Resource Allocation Using Utility Functions

- A set of flows F
- If x_f is the rate of flow f, then the utility to flow f is $U_f(x_f)$, where $U_f(x_f)$ is a concave utility function.
- Maximize aggregate utility, subject to capacity constraints

$$
\begin{array}{ll}
\max & \sum_{f \in F} U_f\left(x_f\right) \\
\text{subject to} & \sum_{f:f \text{ uses link } l} x_f \le c_l \text{ for any link } l \\
\text{over} & x \ge 0
\end{array}
$$

Example: Proportional Fairness

$$\max_{x_f \geq 0} \sum_f \log x_f \qquad U_f(x_f) = \log(x_f)$$

subject to

$$x_1 + x_2 \leq 1$$

$$x_1 + x_3 \leq 1$$

- Optimal: $x_1 = 1/3$
  $x_2 = x_3 = 2/3$

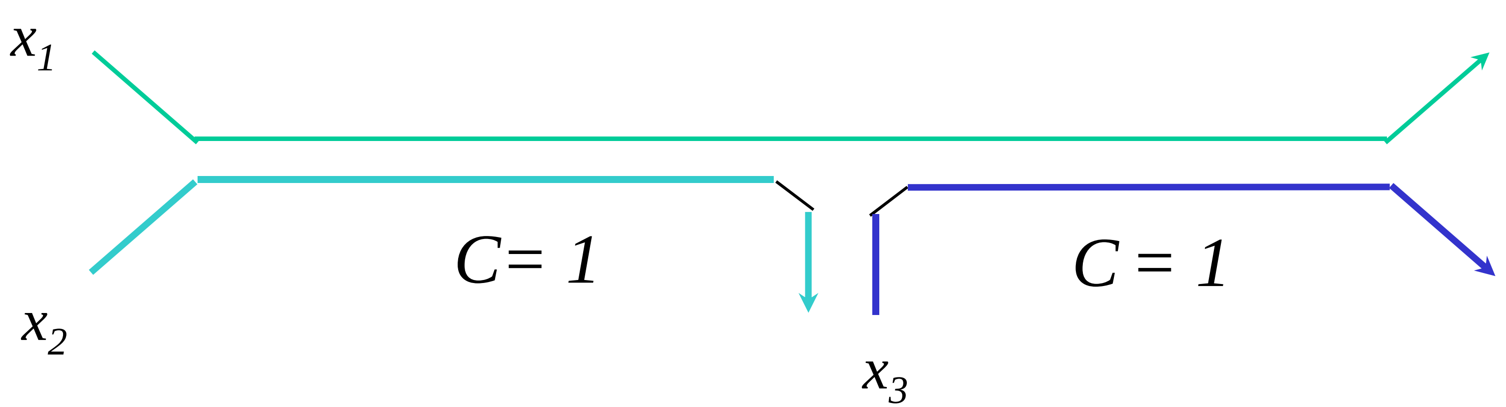

Example 3: a Utility Function

$$\max_{x_f \geq 0} \quad -\frac{1}{4x_1} - \frac{1}{x_2} - \frac{1}{x_3}$$

subject to

$$x_1 + x_2 \leq 1$$

$$x_1 + x_3 \leq 1$$

$$U_f(x_f) = -\frac{1}{RTT^2 x_f}$$

- Optimal: $x_1 = 0.26$
  $x_2 = x_3 = 0.74$

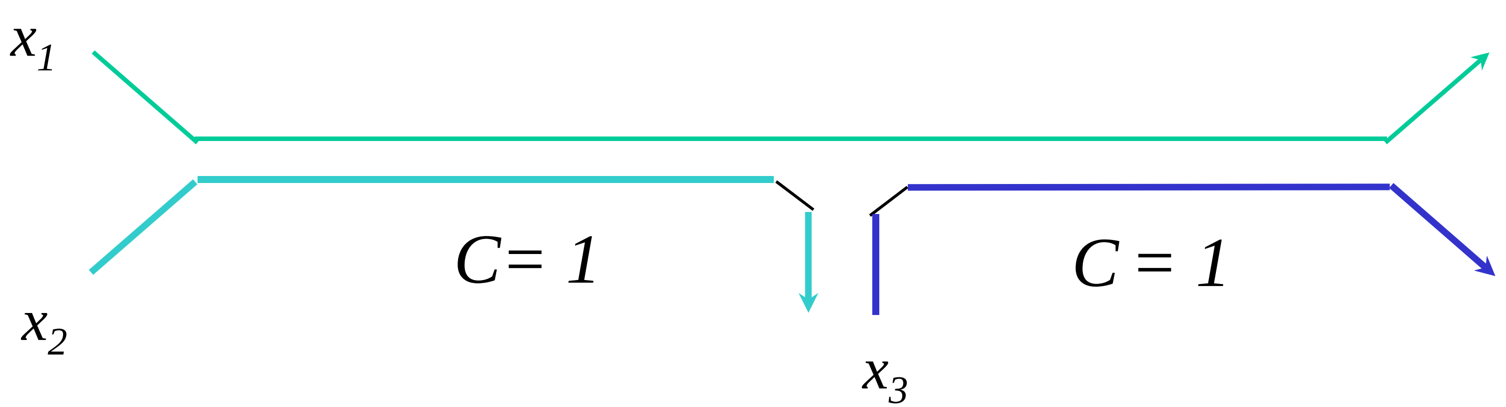

Summary: Allocation

Objective	Allocation (x1, x2, x3)		
TCP/Reno	0.26	0.74	0.74
TCP/Vegas	1/3	2/3	2/3
Max Throughput	0	1	1
Max-min	½	½	½
Max sum log(x)	1/3	2/3	2/3
Max sum of -1/(RTT² x)	0.26	0.74	0.74

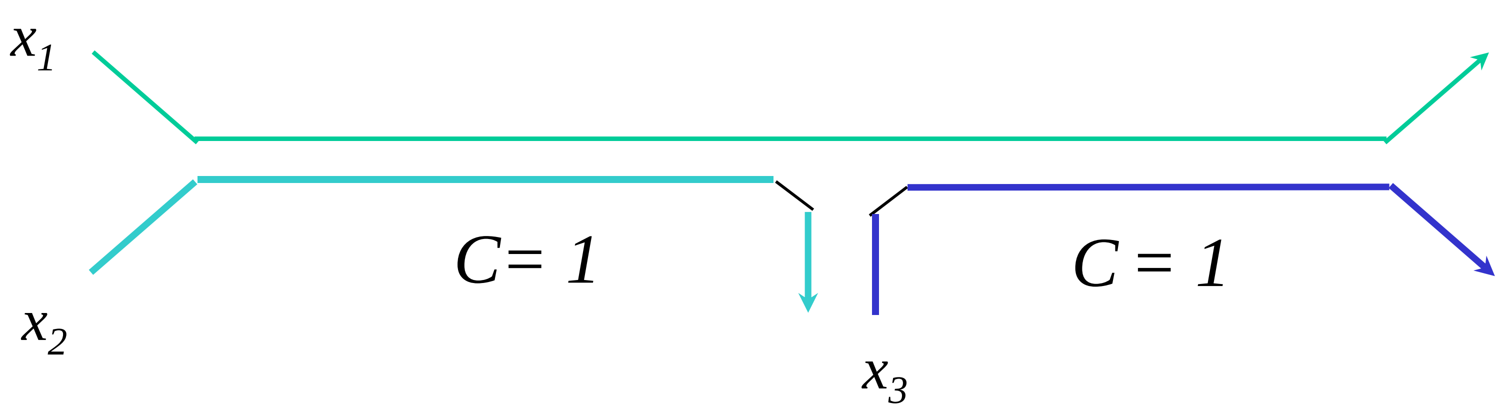

Questions

$$\begin{array}{ll} \text{max} & \sum_{f \in F} U_f(x_f) \\ \text{subject to} & \sum_{f: f \text{ uses link } l} x_f \le c_l \text{ for any link } l \\ \text{over} & x \ge 0 \end{array}$$

- Forward engineering: which allocation objective when allocation by optimization?

- Reverse engineering: what are the objectives of TCP/Reno, TCP/Vegas?

Objective	Allocation (x1, x2, x3)		
TCP/Reno	0.26	0.74	0.74
TCP/Vegas	1/3	2/3	2/3
Max throughput	0	1	1
Max-min	½	½	½
Max sum log(x)	1/3	2/3	2/3
Max sum of $-1/(RTT^2 x)$	0.26	0.74	0.74

Outline

- Recap
- Bandwidth allocation framework
  - Motivation
  - Nash Bargaining Solution (NBS)

Network Bandwidth Allocation Using Nash Bargain Solution (NBS)

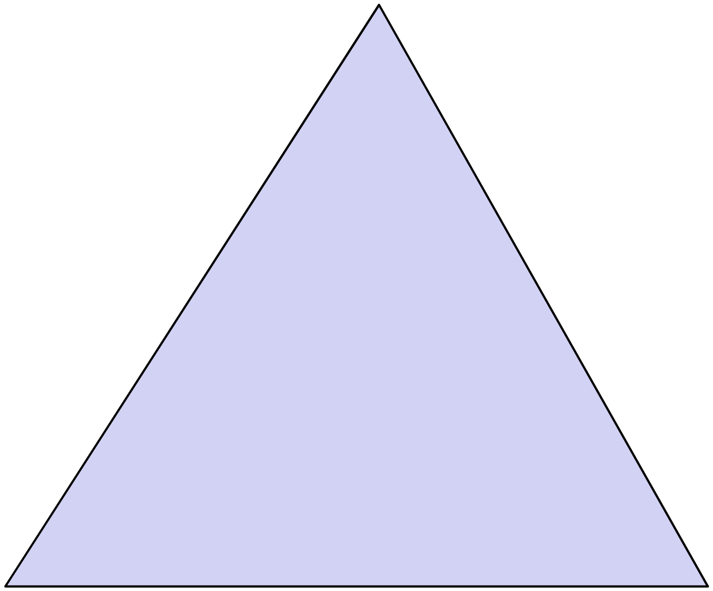

- High level picture
  - given the feasible set of bandwidth allocation, we want to pick an allocation point that is efficient and fair
- The determination of the allocation point should be based on "first principles" (axioms)

Network Bandwidth Allocation: Feasible Region

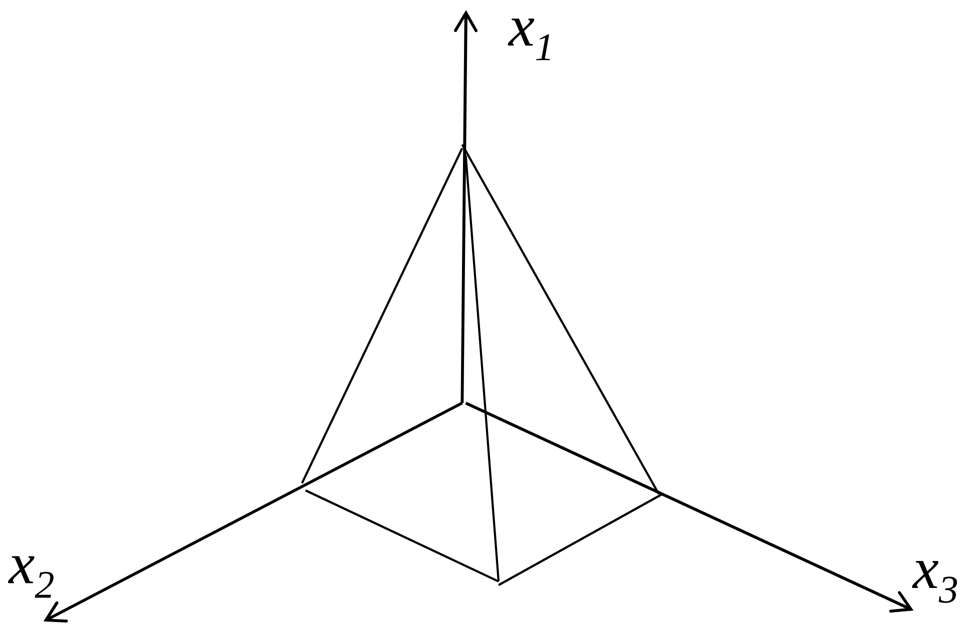

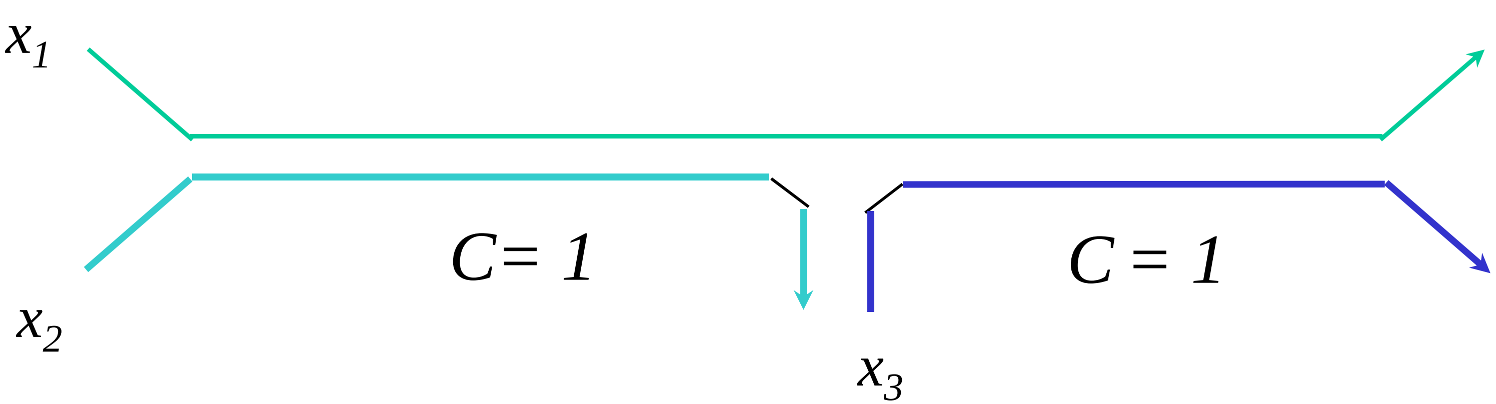

Nash Bargain Solution (NBS)

- Assume a finite, convex feasible set
- Axioms
  - Pareto optimality
    - impossibility of increasing the rate of one user without decreasing the rate of another
  - symmetry
    - a symmetric feasible set yields a symmetric outcome
  - invariance of linear transformation
    - the allocation must be invariant to linear transformations of users' rates
  - independence of irrelevant alternatives
    - assume s is an allocation when feasible set is R, $s \in T \subset R$, then s is also an allocation when the feasible set is T

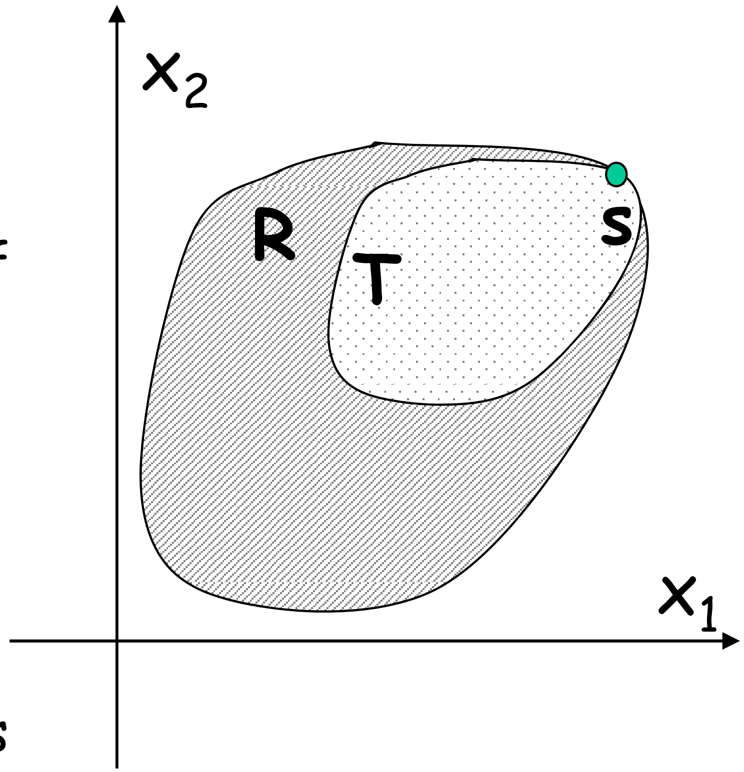

Nash Bargain Solution (NBS)

- Surprising result by John Nash (1951)
  - the rate allocation point is the feasible point which maximizes

$$x_1 x_2 \cdots x_F$$

- This is equivalent to maximize

$$\sum_f \log(x_f)$$

- In other words, assume each flow f has utility function $\log(x_f)$
- I will give a proof for F = 2
  - think about F > 2

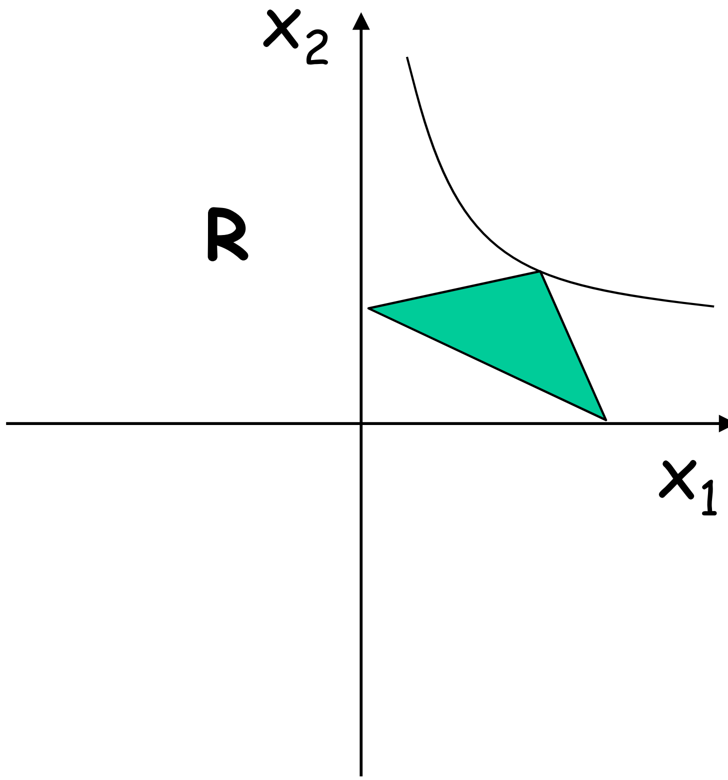

Nash Bargain Solution

- Assume s is the feasible point which maximizes x1 * x2
- Scale the feasible set so that s is at (1, 1)
- Question: after the transformation, is s still the point maximizing x1*x2?

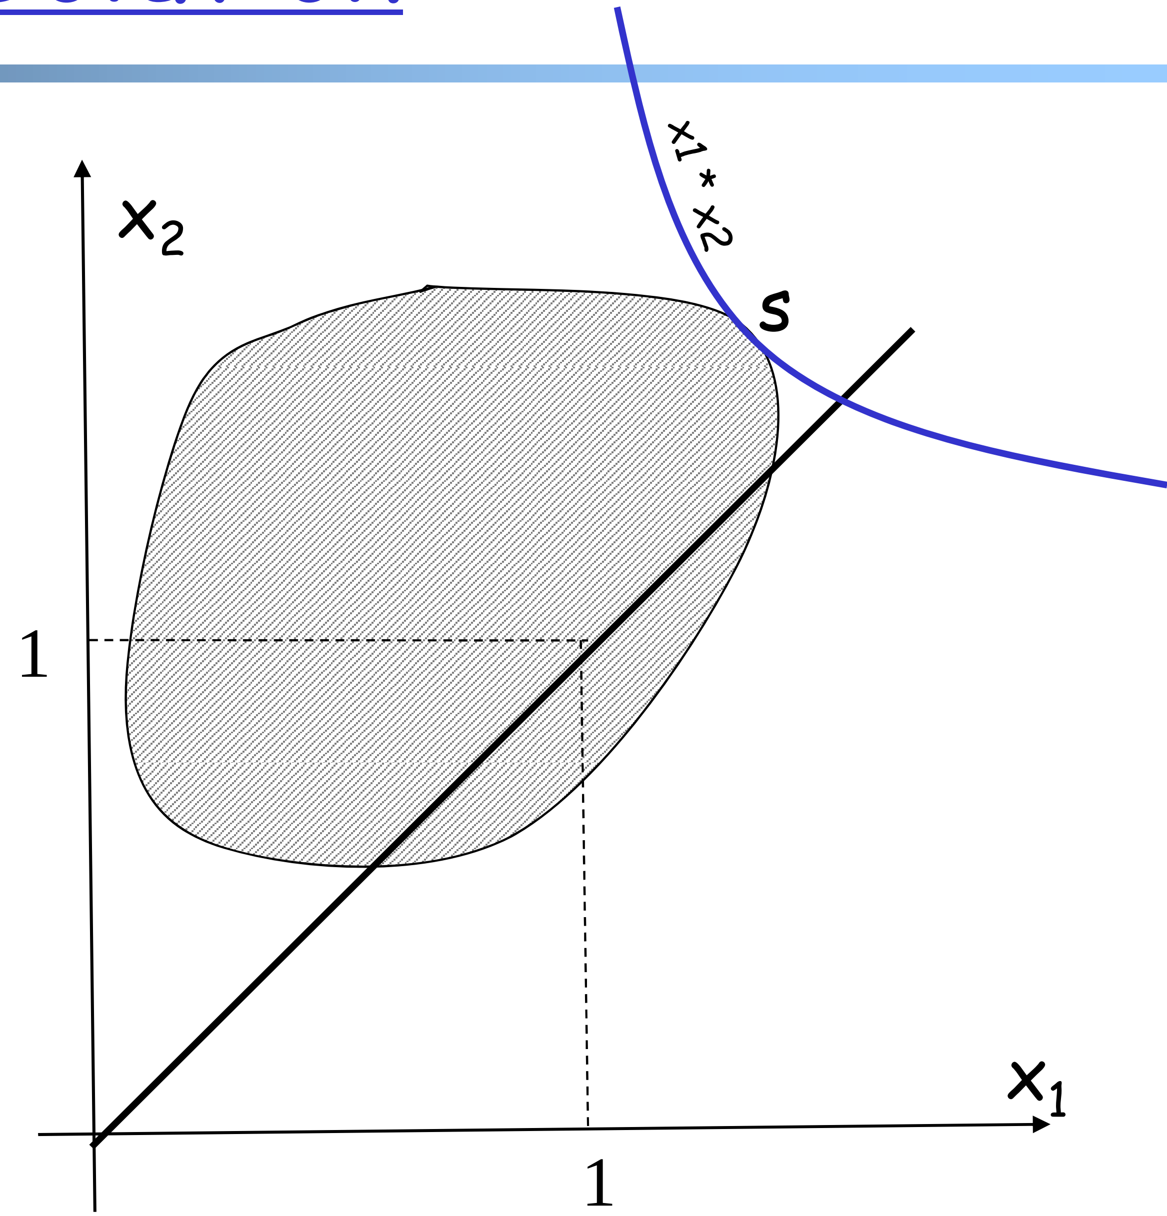

Nash Bargain Solution

Question: after the transformation, is there any feasible point with x1 + x2 > 2?

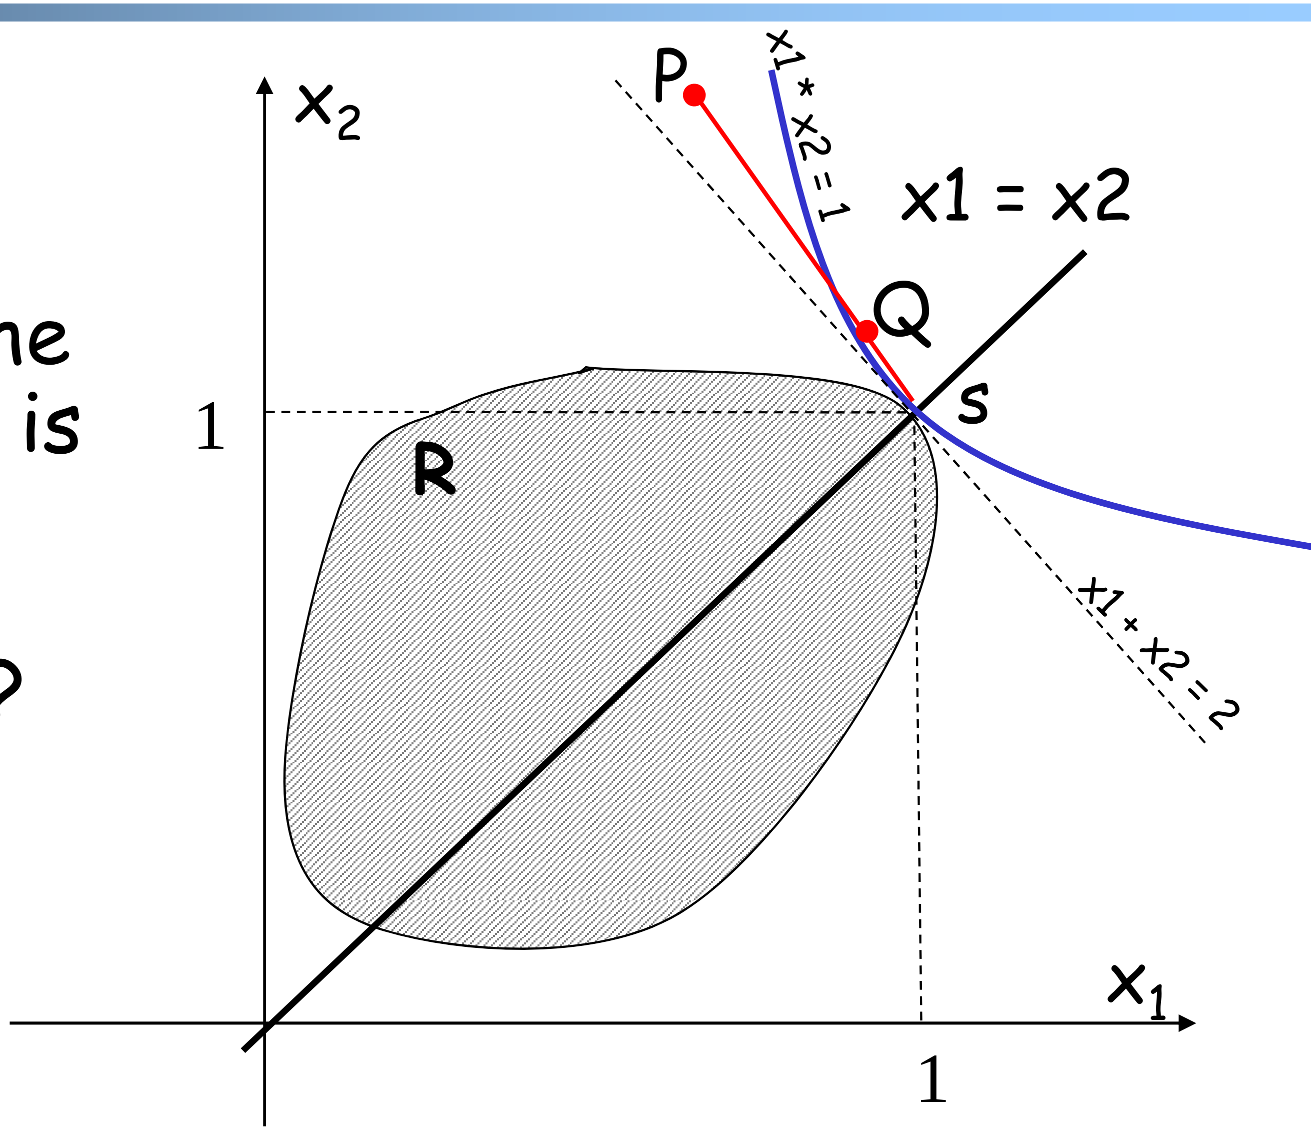

Nash Bargain Solution

- Consider the symmetric rectangle U containing the original feasible set

-> According to symmetry and Pareto, s is the allocation when feasible set is U

- According to independence of irrelevant alternatives, the allocation of R is s as well.

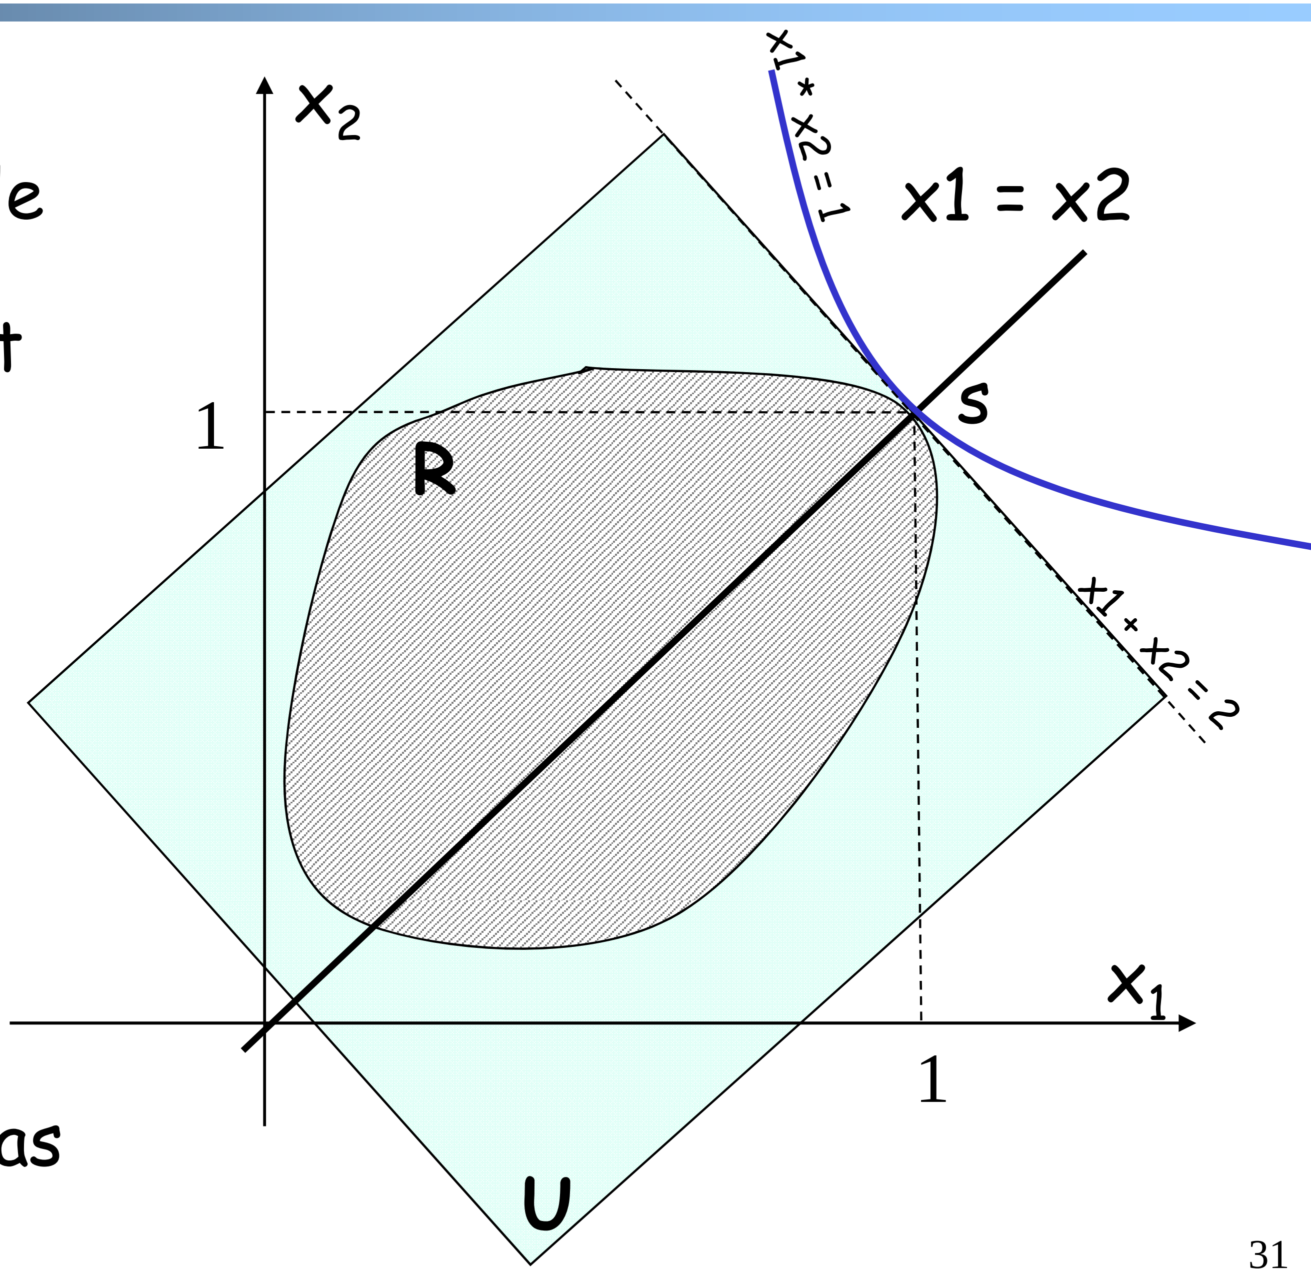

NBS $\Leftrightarrow$ Proportional Fairness

- Allocation is proportionally fair if for any other allocation, aggregate of proportional changes is non-positive, e.g. if x_f is a proportional-fair allocation, and y_f is any other feasible allocation, then require

$$\sum_f \frac{y_f - x_f}{x_f} \leq 0$$

Summary: Allocation Schemes

- Max throughput
- Max-min
- Proportional fair
  - NBS