CS 5229
ADVANCED COMPUTER NETWORKS

Dr. Chan Mun Choon

School of Computing

National University of Singapore

Motivation

- Consider file transfer
- Sender sends a stream of packets representing fragments of a file
- Sender should try to match rate at which receiver and network can process data
  - can't send too slow or too fast
- Too slow
  - Resource under-utilize, unnecessary delay
- Too fast
  - Packet loss, retransmission, long delay

- A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer, *K. K. Ramakrishnan, R. Jain (Proc. SIGCOMM `88, Stanford, CA, August 1988, Vol. 18, No. 4)*

Terminologies

- End-to-end flow control looks at “selfish” control function
- Congestion control/avoidance address the “social” problem
- Congestion control vs. congestion avoidance

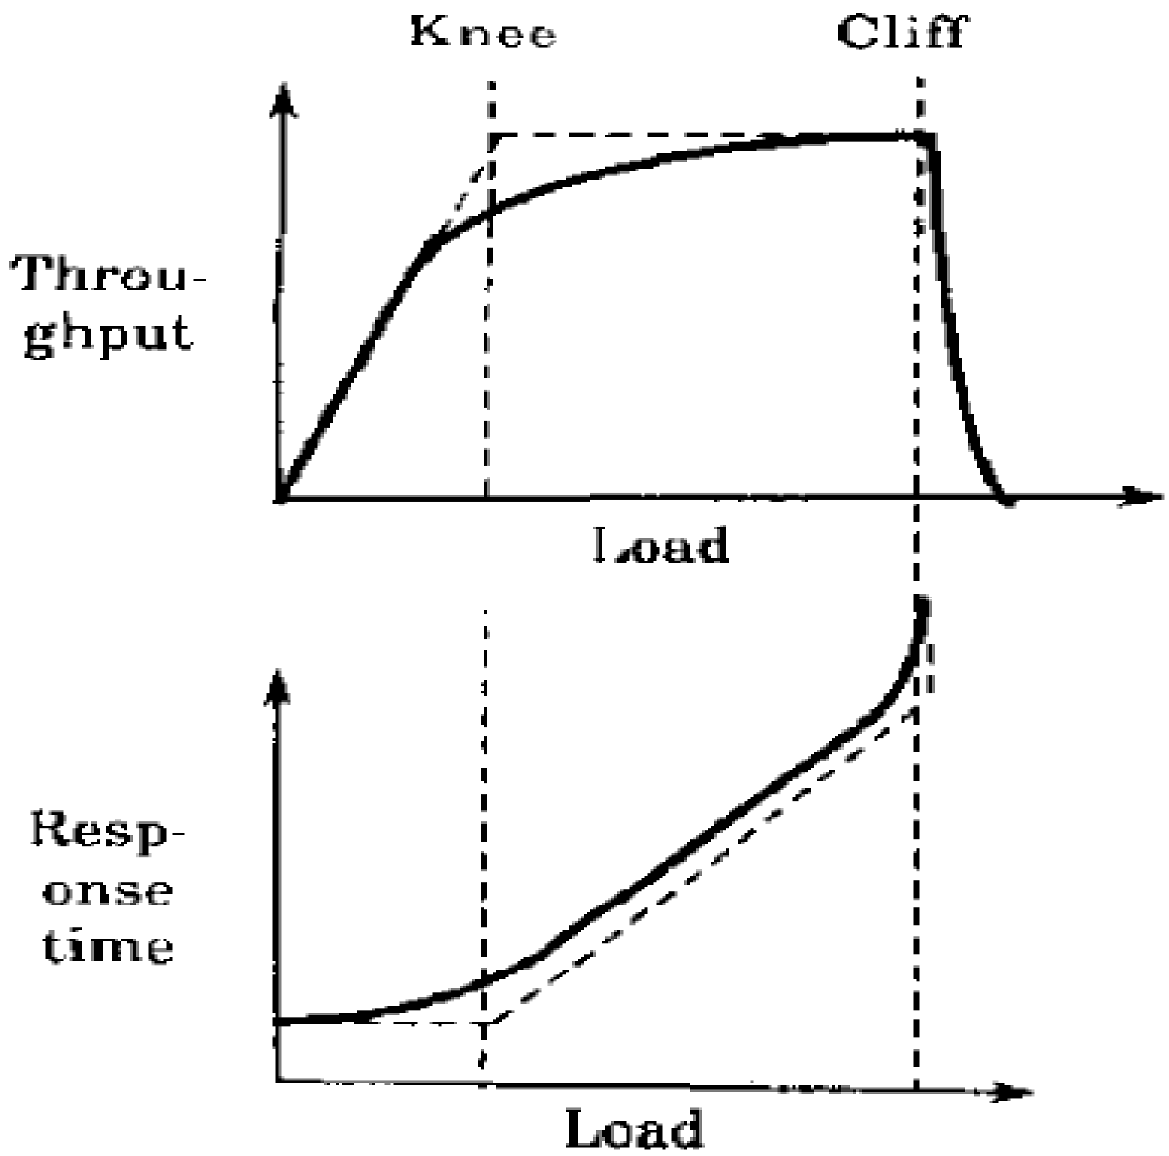
Knee
Cliff
Throu-
ghput
Load
Resp-
onse
time
Load

Binary Feedback

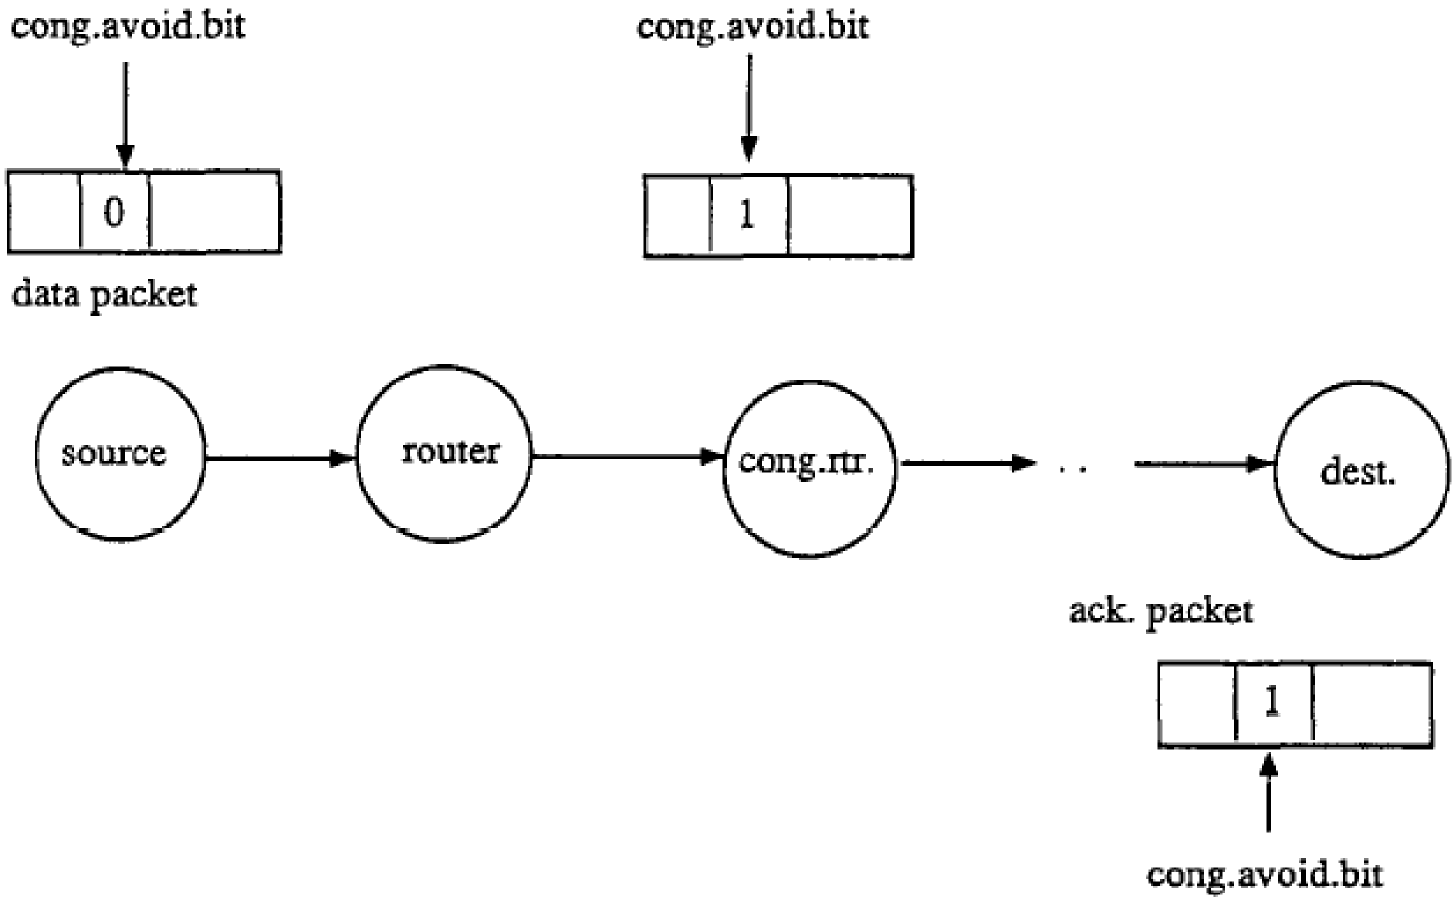

Model

- Network as a feedback control system
- Each user controls the amount of traffic, multiple users have to coordinate
- Instantaneous network state varies dynamically, feedback is noisy and subjected to delay
- Policies needed on the router and user
- Workload: each source is considered to have packets to send at all times
- Window based control

Objective

$$Power = \frac{Throughput^{\alpha}}{Response\ time}, \qquad \text{where} \quad 0 < \alpha < 1.$$

When $\alpha = 1$, power is maximized at the “knee”

Fairness

- Jain's Index

$$f = \frac{\left(\sum_{i=1}^{N} x_i\right)^2}{\left(n \sum_{i=1}^{N} x_i^2\right)} \qquad \text{where} \quad x_i = \frac{A_i}{D}.$$

1. Index is between 0 and 1
2. Independent of scale
3. Continuous value
4. If only k of n users are allocated resource, index is k/n

Examples

1. {0,60,60,60,60,60,60,60,60,60}
   - Index = 0.9
2. {10,60,60,60,60,60,60,60,60,60}
   -
3. {10,20,30,40,50,60,70,80,90,100}
   -
4. {10,11,13,15,17,20,25,33,100}
   -

What to monitor on router?

- Router set the congestion indication bit on arriving packet when the average number of packets at the router is greater or equal to 1
- What is measured?

How long to average?

- Should be adaptive

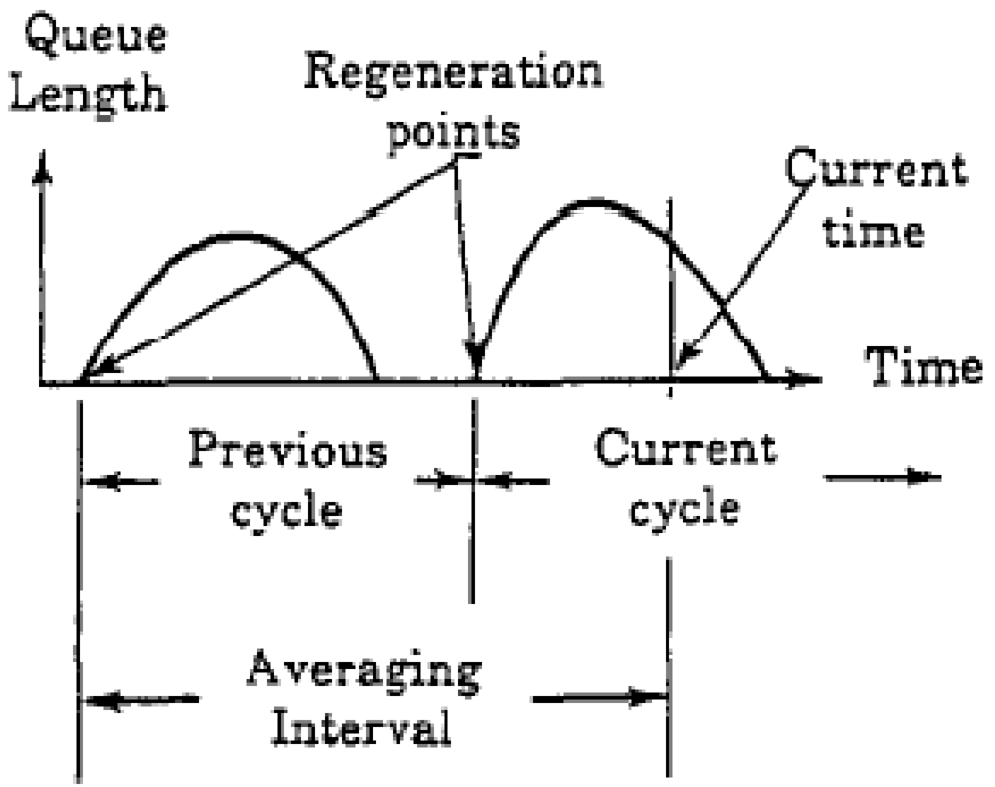

Decision Frequency

- If decision is made too often (say every packet), there may be considerable oscillation
- Wp = window size before update
- Wc = window size after update
- Frequency = Wp+Wc

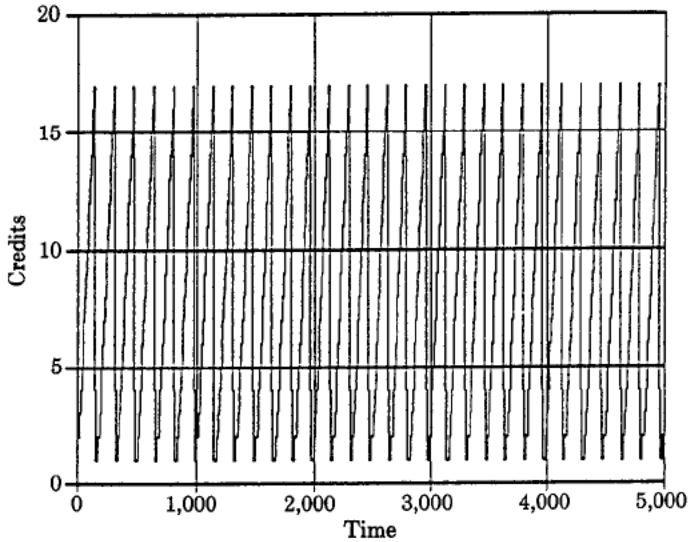

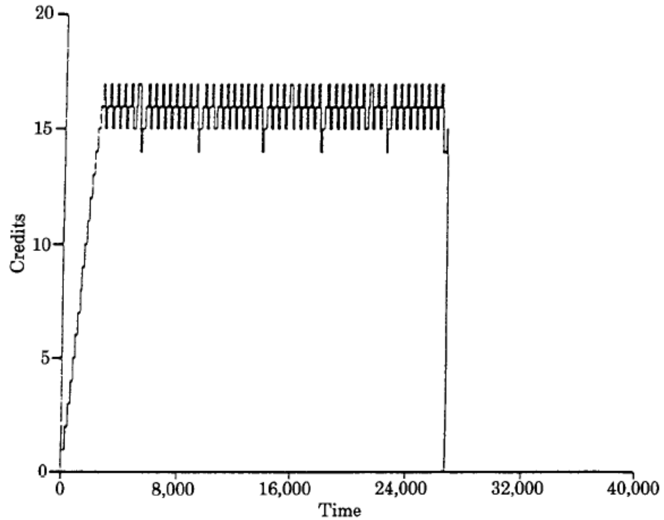

Information Use

- Receive Wp+Wc bits between updates
- Use only the last Wc bits
- Act when at least 50% of bits are set
- Try with M/M/1 model

Decision Function

- How to increase or decrease window size?

- DM Chiu, R Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," Computer Networks and ISDN systems, 1989.

Assumptions

- Feedback and control loop for all users is synchronous
- Single bottleneck
- Binary feedback
- Congestion state is determined by the number of packets in the system

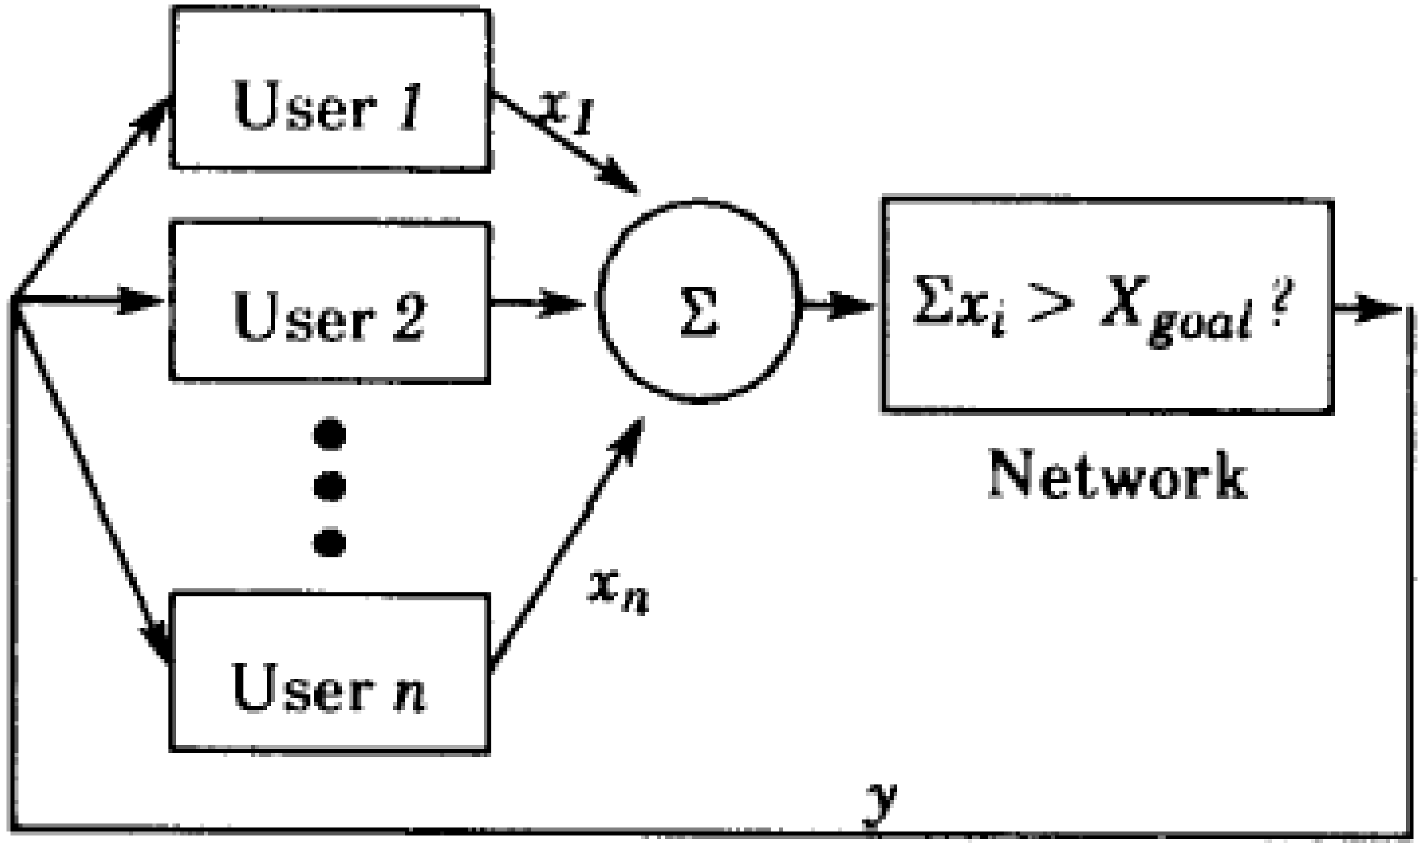

Fig. 2. A control system model of n users sharing a network.

Linear control

$$x_i(t+1) = \begin{cases} a_\mathrm{I} + b_\mathrm{I} x_i(t) & \text{if } y(t) = 0 \Rightarrow \text{Increase}, \\ a_\mathrm{D} + b_\mathrm{D} x_i(t) & \text{if } y(t) = 1 \Rightarrow \text{Decrease}. \end{cases}$$

- a -> additional
- b -> multiplicative
- a_I, b_I
- a_D, b_D

4 options

1. Multiplicative Increase/Multiplicative decreases
   - $b_I > 1$, $0 < b_D < 1$
2. Additive Increase/Additive decreases
   - $a_I > 0$, $a_D < 0$, $b = 1$
3. Additive Increase/Multiplicative decreases
   - $a_I > 0$, $0 < b_D < 1$, $b_I = 1$
4. Multiplicative Increase/Additive decreases
   - $b_I > 1$, $a_D < 0$, $b_D = 1$

Criteria for Selection

- How to choose the “best” option?

1. Efficiency
   - Sum of rates is close to target rate
2. Fairness
   - Jain’s index
3. Distributedness
   - Minimum feedback, no global information
4. Convergence

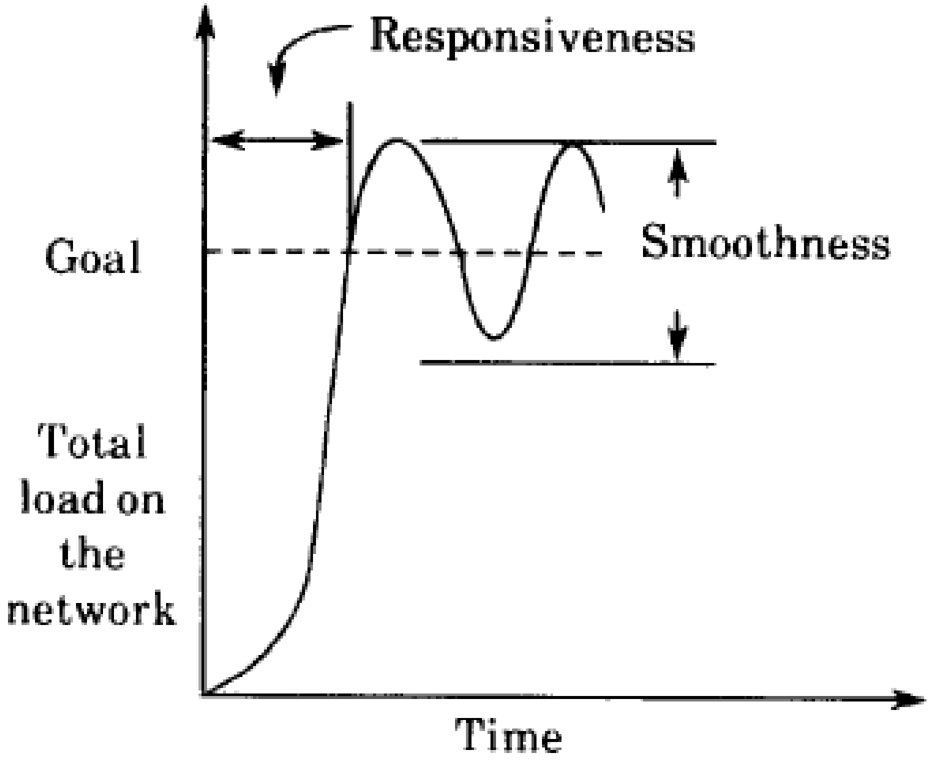

Fig. 3. Responsiveness and smoothness.

With binary feedback, system does not generally converge to a single steady state but oscillates around the “goal” at equilibrium

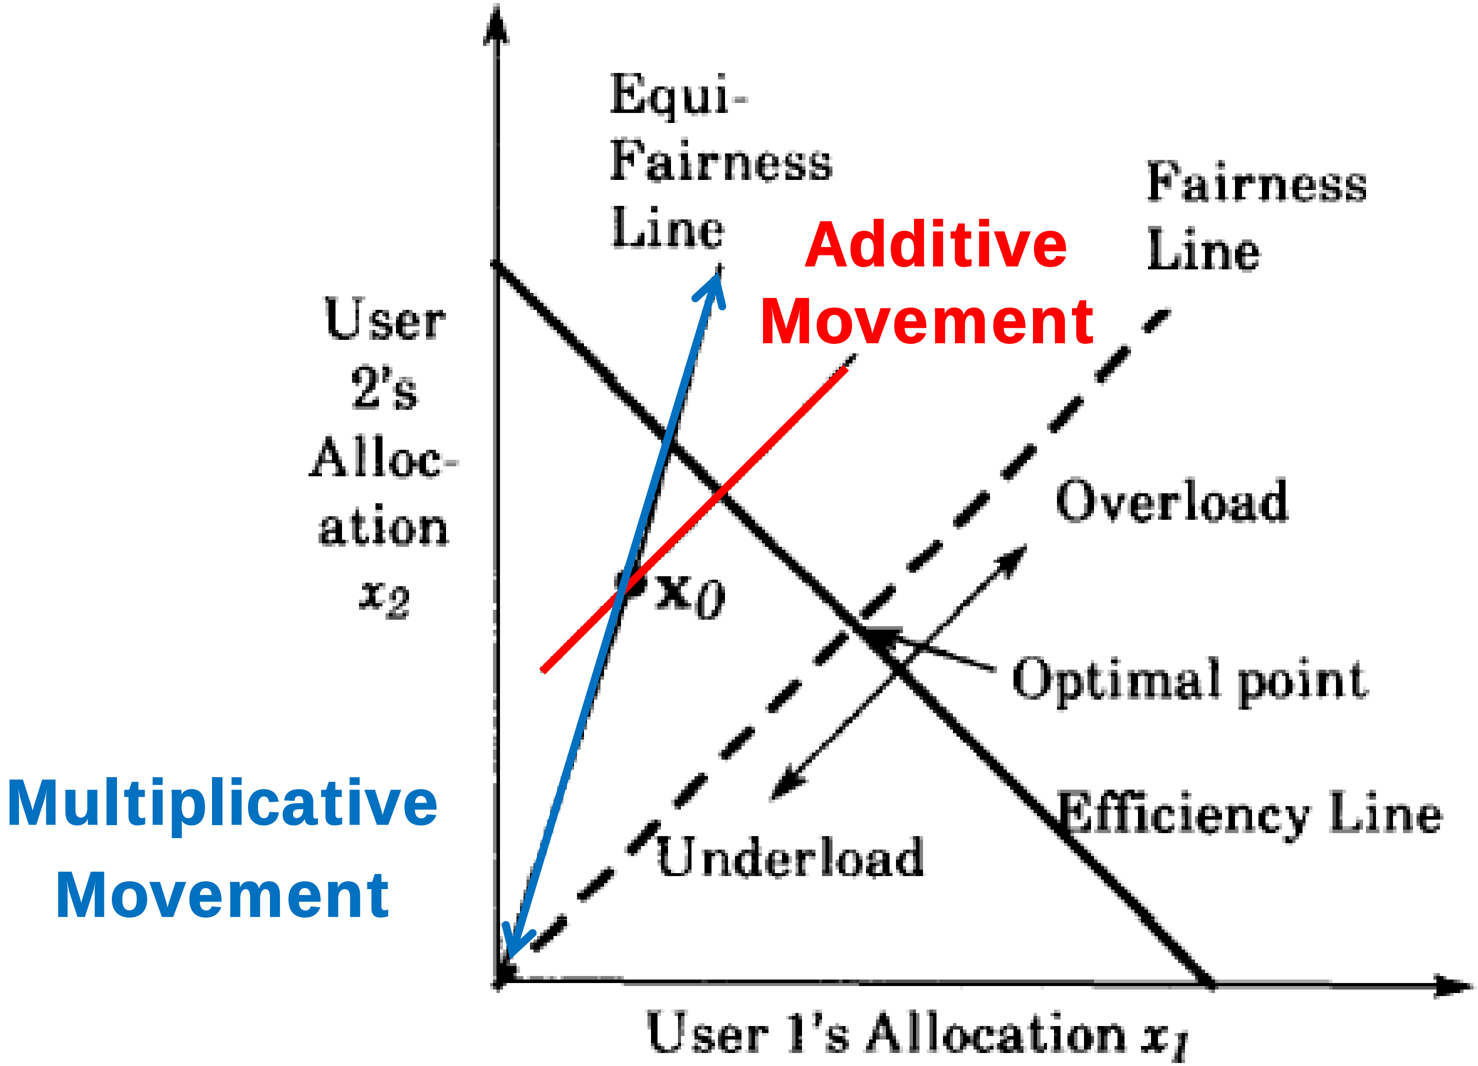

Fig. 4. Vector representation of a two-user case.

Comments

- Using only additive control
  - Can converge to efficiency line, but not to fairness line (though fairness value changes)
- Using only multiplicative control
  - Can converge to efficiency line, but has no effect on fairness

Convergence to Efficiency

$$y(t) = 0 \;\Rightarrow\; \Sigma x_i(t+1) > \Sigma x_i(t),$$
$$y(t) = 1 \;\Rightarrow\; \Sigma x_i(t+1) < \Sigma x_i(t).$$

$$b_{\mathrm{I}} > 1 - \frac{n a_{\mathrm{I}}}{\Sigma x_i(t)}$$
$$b_{\mathrm{D}} < 1 - \frac{n a_{\mathrm{D}}}{\Sigma x_i(t)} \quad \forall n \text{ and } \forall \Sigma x_i(t). \qquad (3)$$

Convergence to Fairness

$$F(x(t+1)) = \frac{\left(\Sigma x_i(t+1)\right)^2}{n\left(\Sigma x_i^2(t+1)\right)} \qquad (4)$$

$$\text{where } c = a/b \qquad (6)$$

$$= F(x(t)) + (1 - F(x(t)))$$

$$\times \left(1 - \frac{\Sigma x_i^2(t)}{\Sigma\left(c + x_i(t)\right)^2}\right). \qquad (7)$$

As long as c >0, fairness improves over time

Cont'd

$$\frac{a_I}{b_I} \geqslant 0 \quad \text{and} \quad \frac{a_D}{b_D} > 0 \tag{8}$$

or

$$\frac{a_I}{b_I} > 0 \quad \text{and} \quad \frac{a_D}{b_D} \geqslant 0. \tag{9}$$

- Fairness can improve during increase or decrease
- All parameters must be positive

Distributedness

- Equation (3) won't work because it needs sum of rates and number of users
- Translate condition for sum of flows (global) into conditions for individual flow (local)

$$\begin{aligned} y(t) = 0 \quad &\Rightarrow \quad x_i(t+1) > x_i(t) \ \forall i, \\ y(t) = 1 \quad &\Rightarrow \quad x_i(t+1) < x_i(t) \ \forall i. \end{aligned} \tag{11}$$

Result

$$a_{\mathrm{I}} + (b_{\mathrm{I}} - 1)x_i(t) > 0 \quad \forall x_i(t) \geqslant 0,$$

$$a_{\mathrm{D}} + (b_{\mathrm{D}} - 1)x_i(t) < 0 \quad \forall x_i(t) \geqslant 0.$$

This implies further constraining equation (10) to be

$$\begin{aligned} &a_{\mathrm{I}} > 0, \qquad b_{\mathrm{I}} \geqslant 1, \\ &a_{\mathrm{D}} = 0, \qquad 0 \leqslant b_{\mathrm{D}} < 1. \end{aligned} \tag{12}$$

Stated as proposition 1

Convergences to Fairness

- C > 0, determines the rate of convergence
  - Large c means faster convergence
- C = a/b, to increase C
  - Choose large a or small b
- Smallest b is 1
  - Additive increases is good for fairness convergence

Main Result

Proposition 3. *For both feasibility and optimal convergence to fairness, the increase policy should be additive and the decrease policy should be multiplicative.*

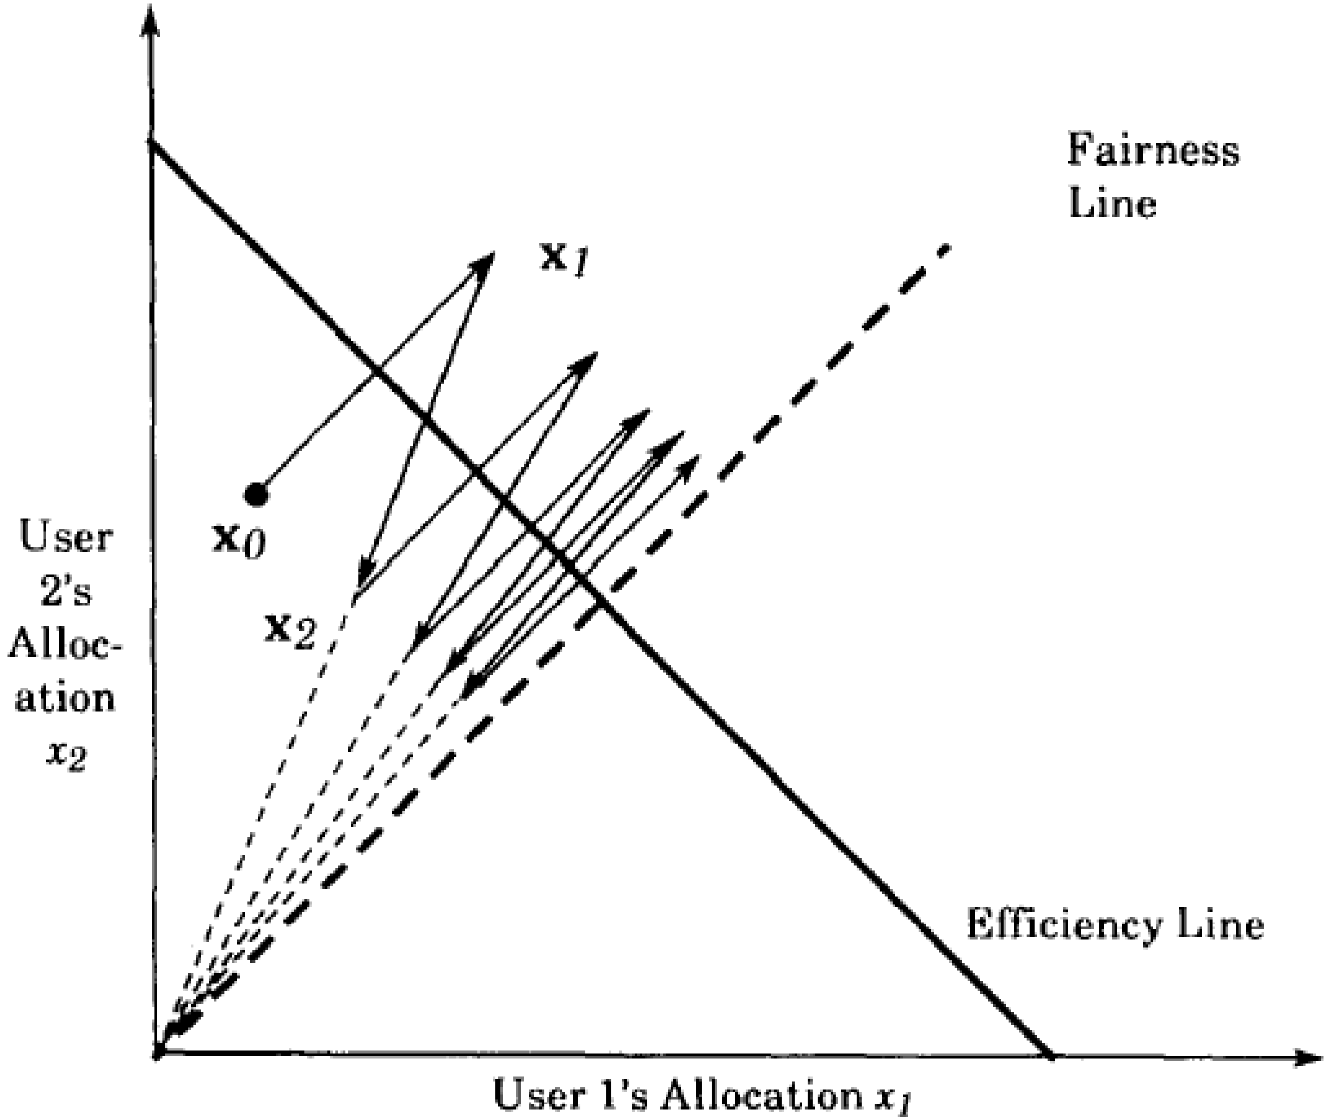

Fig. 5. Additive Increase/Multiplicative Decrease converges to the optimal point.

Discussions

- What about variable delays?
- What is user actions are asynchronous?
- What if router does not provide feedback?
- Is it worthwhile to estimate number of users?

- Congestion Avoidance and Control, *V. Jacobson (Proc. SIGCOMM `88, Stanford, CA, August 1988, Vol. 18, No. 4)*

Fairness

- Scheduling discipline *allocates* a *resource*
- An allocation is fair if it satisfies *some notion of fairness*
- Intuitively
  - each connection gets what it “deserves”

Fairness (contd.)

- Fairness is *intuitively* a good idea
- But it also provides *protection*
  - traffic hogs cannot overrun others
  - automatically builds *firewalls* around heavy users
- Fairness is a *global* objective, but congestion avoidance (as presented before) is local
- Each endpoint must restrict its flow to the smallest fair allocation

Notion of Fairness

- What is "fair" in resource sharing?
  - Everybody gets what they need?
  - How about excess resources?
  - What if there is insufficient resource?
- Example:
  - A "flat" tax system whereby everybody pays the same tax rate.
  - A "progressive" tax system whereby people who has larger income pay at a higher tax rate.
- Factors to consider
  - How does fairness relate to ability to use resource?
  - How does fairness affects overall resource utilization?

Fairness

- Equal Share
  - Resources are shared among all users independent of user requirements and resource utilization
  - Is it a good model for resource sharing?
- Jain's index use equal share as the objective

Max-Min Fairness

- Maximizes the minimum share of a resource whose demand is not fully satisfied
- Intuitively:
  - each connection gets no more than what it wants
  - the excess, if any, is equally shared

Max-Min Flow Control

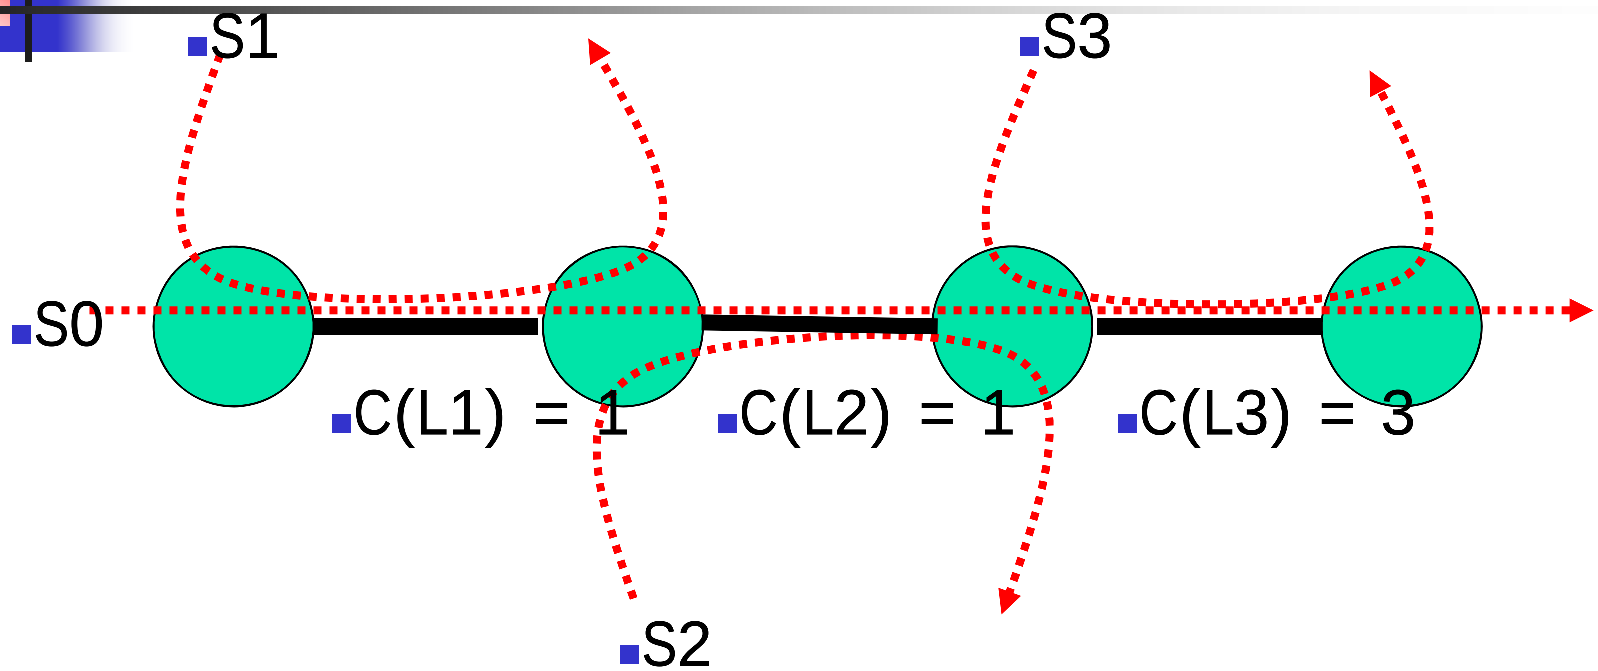

- How much rate should be allocated to S0, S1, S2 and S3?
- Two possibilities:
- {0.5,0.5,0.5,0.5} but L3 is under-utilized
- {0.5,0.5,0.5,2.5} S3 gets more bw with no impact on others

Max-Min Flow Control

- A rate allocation is max-min fair if no rate can be increased without decreasing another rate with a smaller or equal value

All link capacities are 1

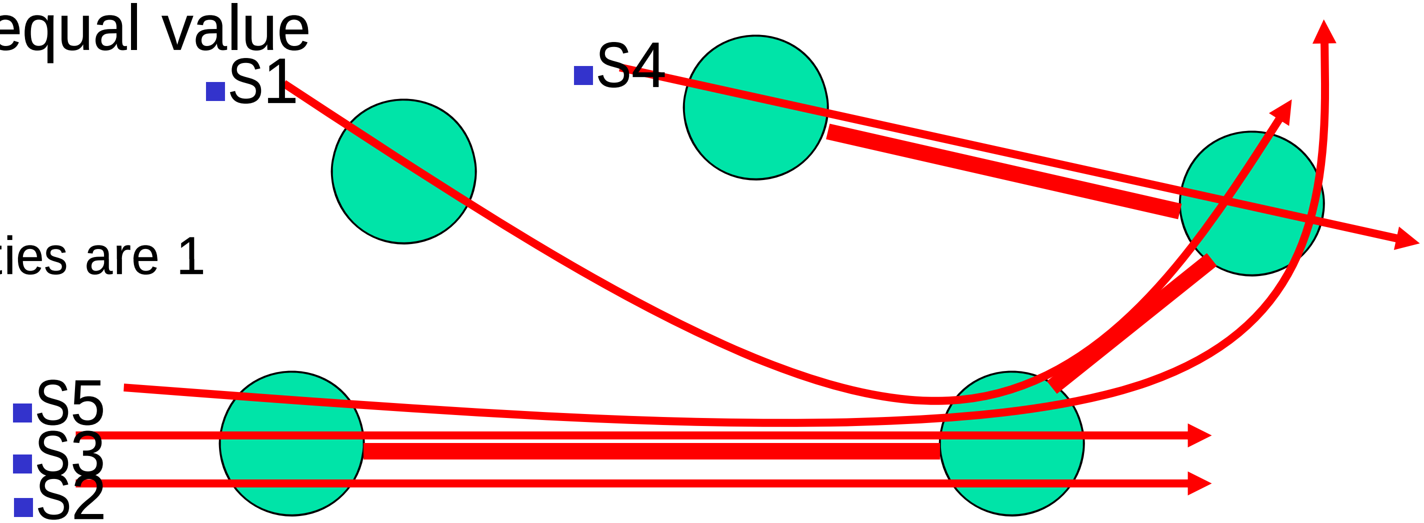

1. {1/3, 1/3, 1/3, 1/3, 1/3}
2. {2/3, 1/3, 1/3, 2/3, 1/3}
3. {2/3, 1/3, 1/3, 1, 1/3}

Max-Min Allocation

- Apply max-min allocation to a single resource
- Interesting case is when demand is greater than capacity
- Given users with demands {2,2.6,4,5} and capacity 10. Total demand = 13.5.

1. {2.5, 2.5, 2.5, 2.5} {0.5,-0.1,-1.5,-2.5} excess=0.5
2. {2,2.66,2.66,2.66} {0,0.06,1.34,2.34} excess=0.06
3. {2,2.6,2.7,2.7} {0,0,1.3,2.3}

Proportional Fair (PF)

- Maximize sum of utility (a function of the allocated rate), a reasonable utility function is log()
- A PF allocation x_i satisfies $\Sigma(y_i - x_i)/x_i <= 0$ for any feasible allocation y
- The allocation below would be
  - Max-Total: {0,1,1,1}. Total = 3. Utility = ?
  - Max-Min Fair: {0.5,0.5,0.5,0.5}. Total = 2. Utility = ?
  - Proportional Fair: {0.25,0.75,0.75,0.75}. Total = 2.5. Utility = ?