Congestion Control In The Internet Part 1: Theory

JY Le Boudec

2019

Contents

1. What is the problem; congestion collapse
2. Efficiency versus Fairness
3. Definitions of fairness
4. Additive Increase Multiplicative Decrease (AIMD)
5. Slow start

Textbook

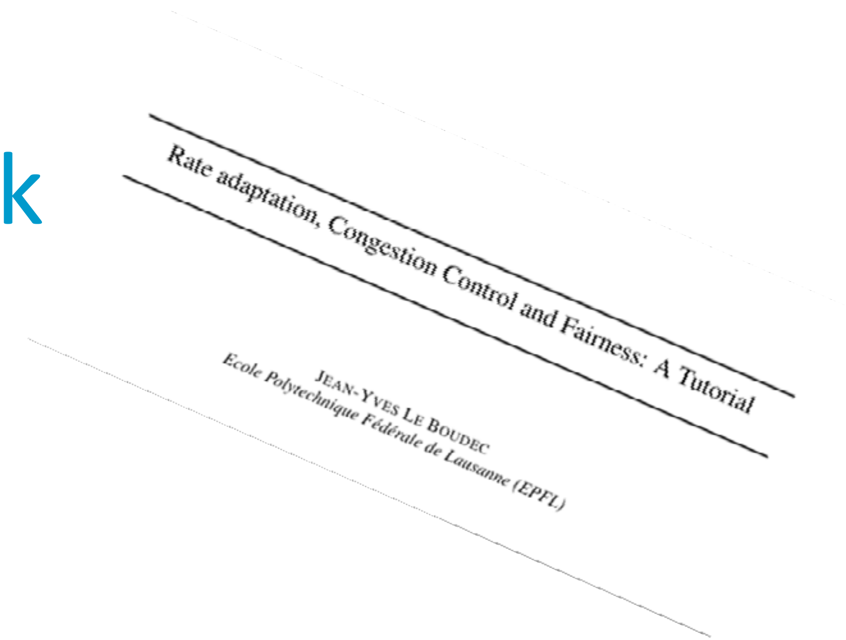

1. Congestion Collapse

In October of '86, the Internet had the first of what became a series of 'congestion collapses'. During this period, the data throughput from LBL to UC Berkeley (sites separated by 400 yards and three IMP hops) dropped from 32 Kbps to 40 bps. Mike Karels[1] and I were fascinated by this sudden factor-of-thousand drop

Jacobson, Van. "Congestion avoidance and control." *ACM SIGCOMM computer communication review*. Vol. 18. No. 4. ACM, 1988.

How much will S_1 send to D_1 ? S_2 to D_2 ?

Network may lose some packets;

Assume greedy sources (i.e. send as much as they want)

Assume loss is proportional to submitted traffic and links can be fully utilized

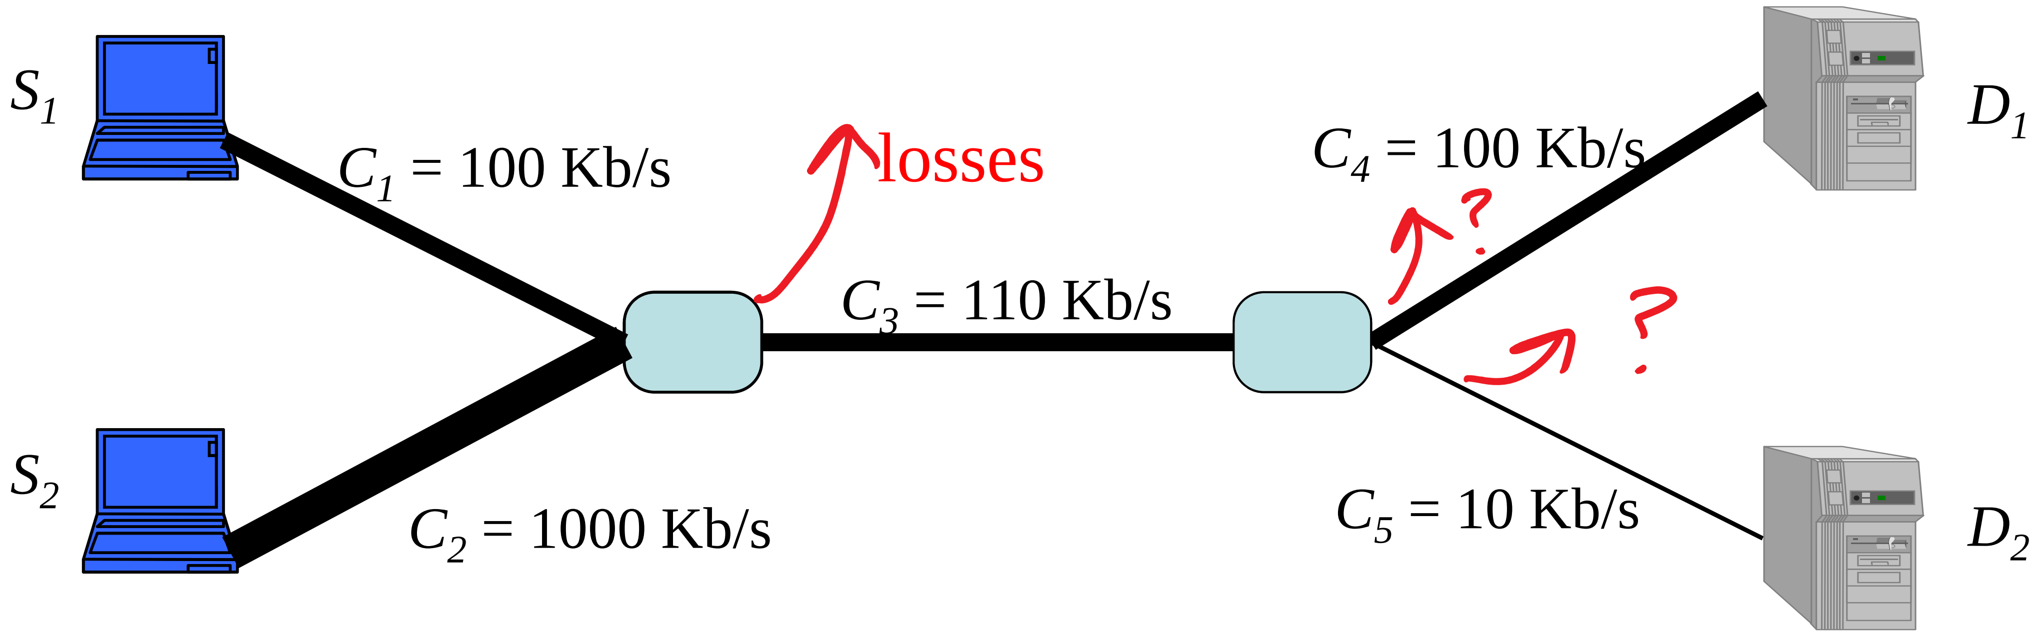

S_1 to D_1 rate is ...

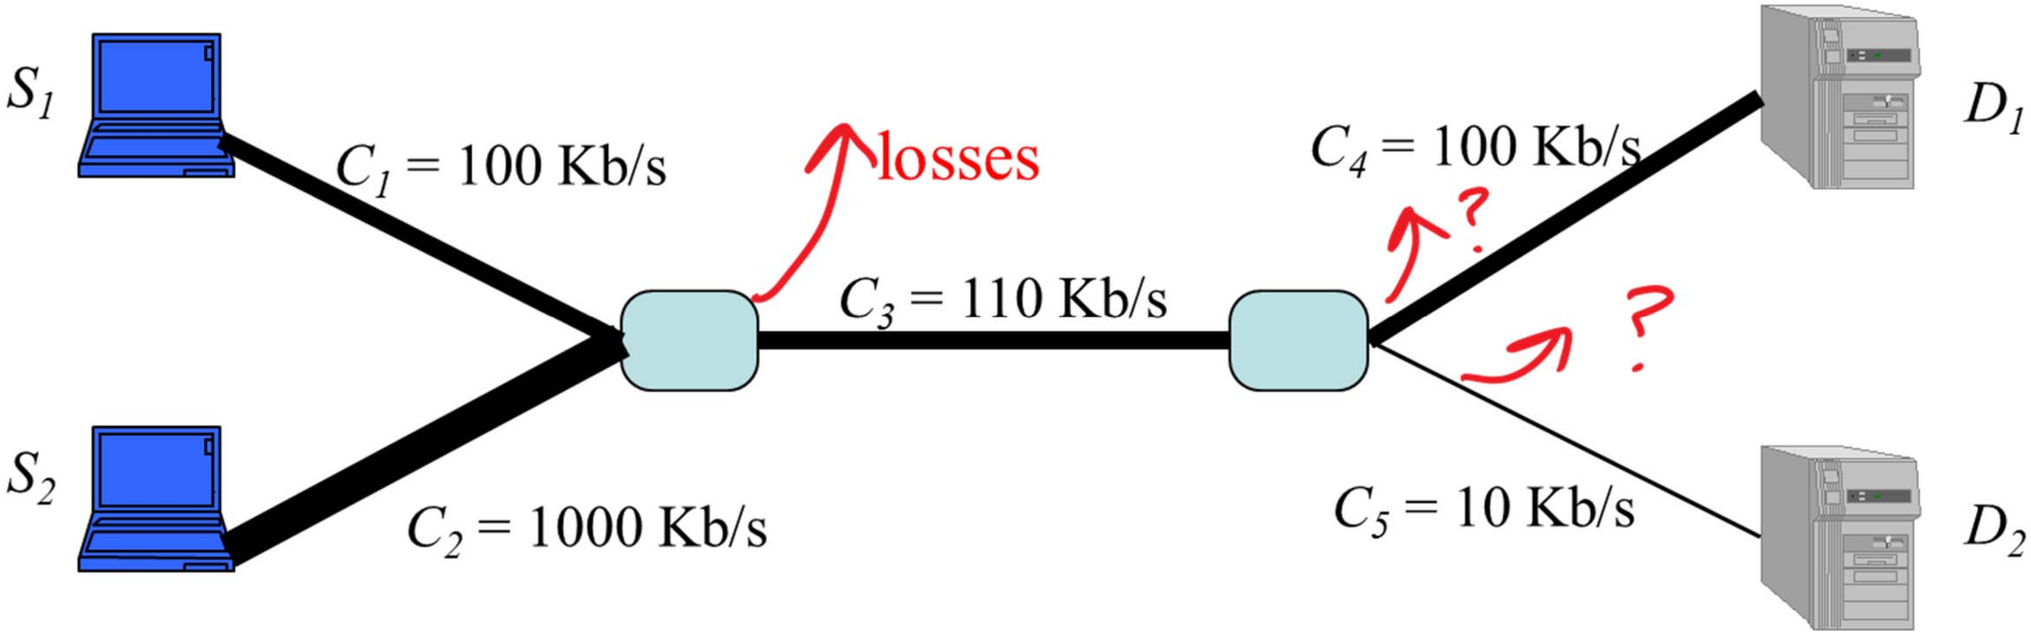

A. 10 kb/s

B. 50 kb/s

C. 100 kb/s

D. I don't know

Solution

Answer A: Ratio of accepted traffic at 1 is $\frac{110}{100+1000} = 10\%$

Ratio of accepted traffic at 2 is 100%

Ratio of accepted traffic at 3 is $\frac{10}{100} = 10\%$

Both sources send at 10 kb/s !

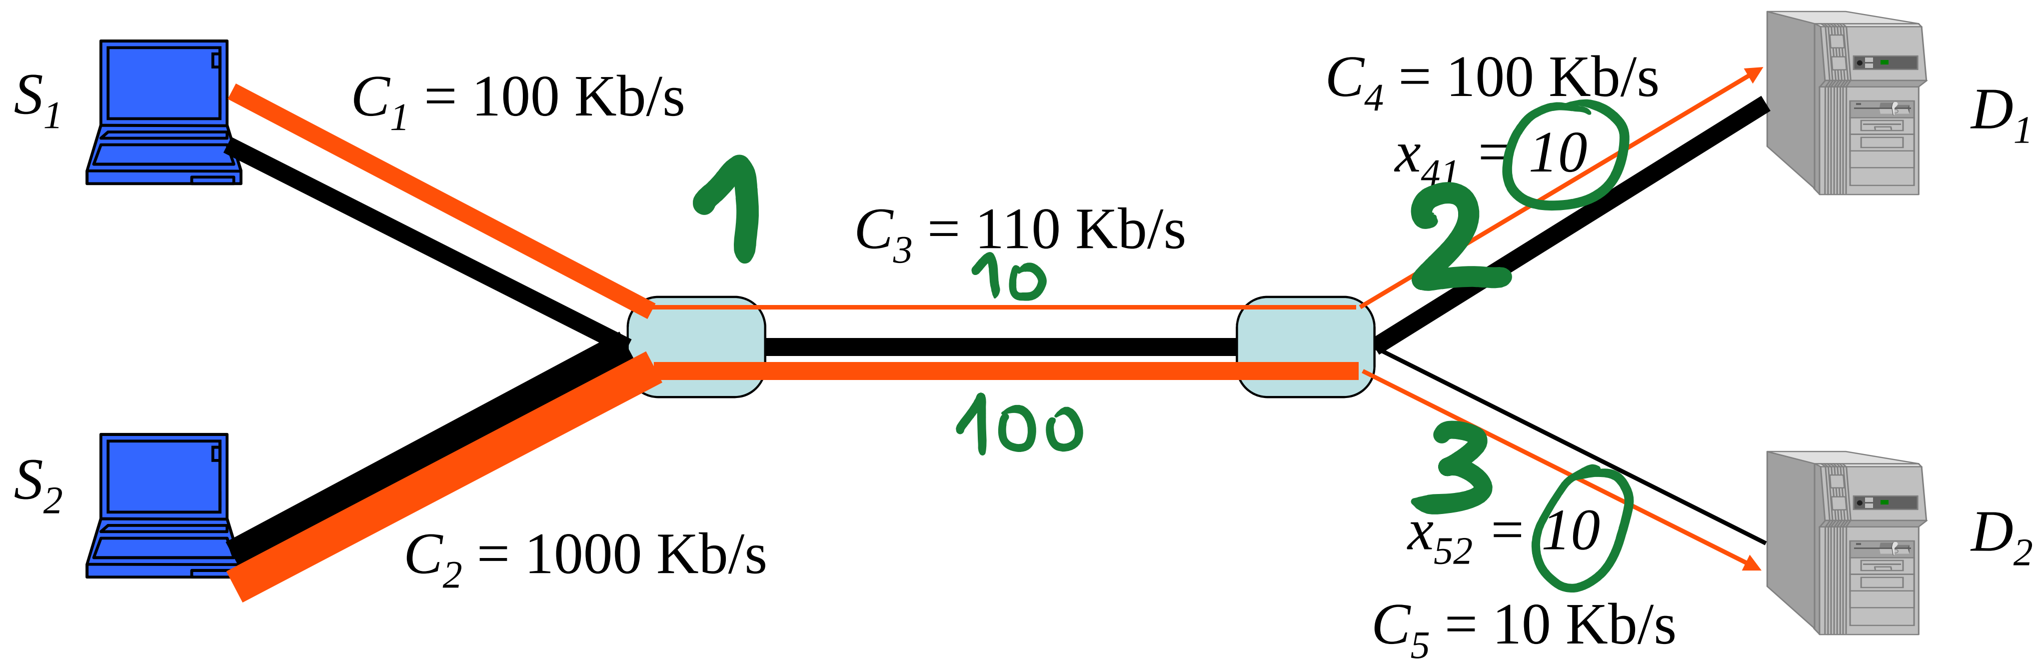

Take home message
Greedy Sources May Be Inefficient

A better allocation is:

S_1: 100 kb/s

S_2: 10 kb/s

The problem was that S_2 sent too much (but did not know)

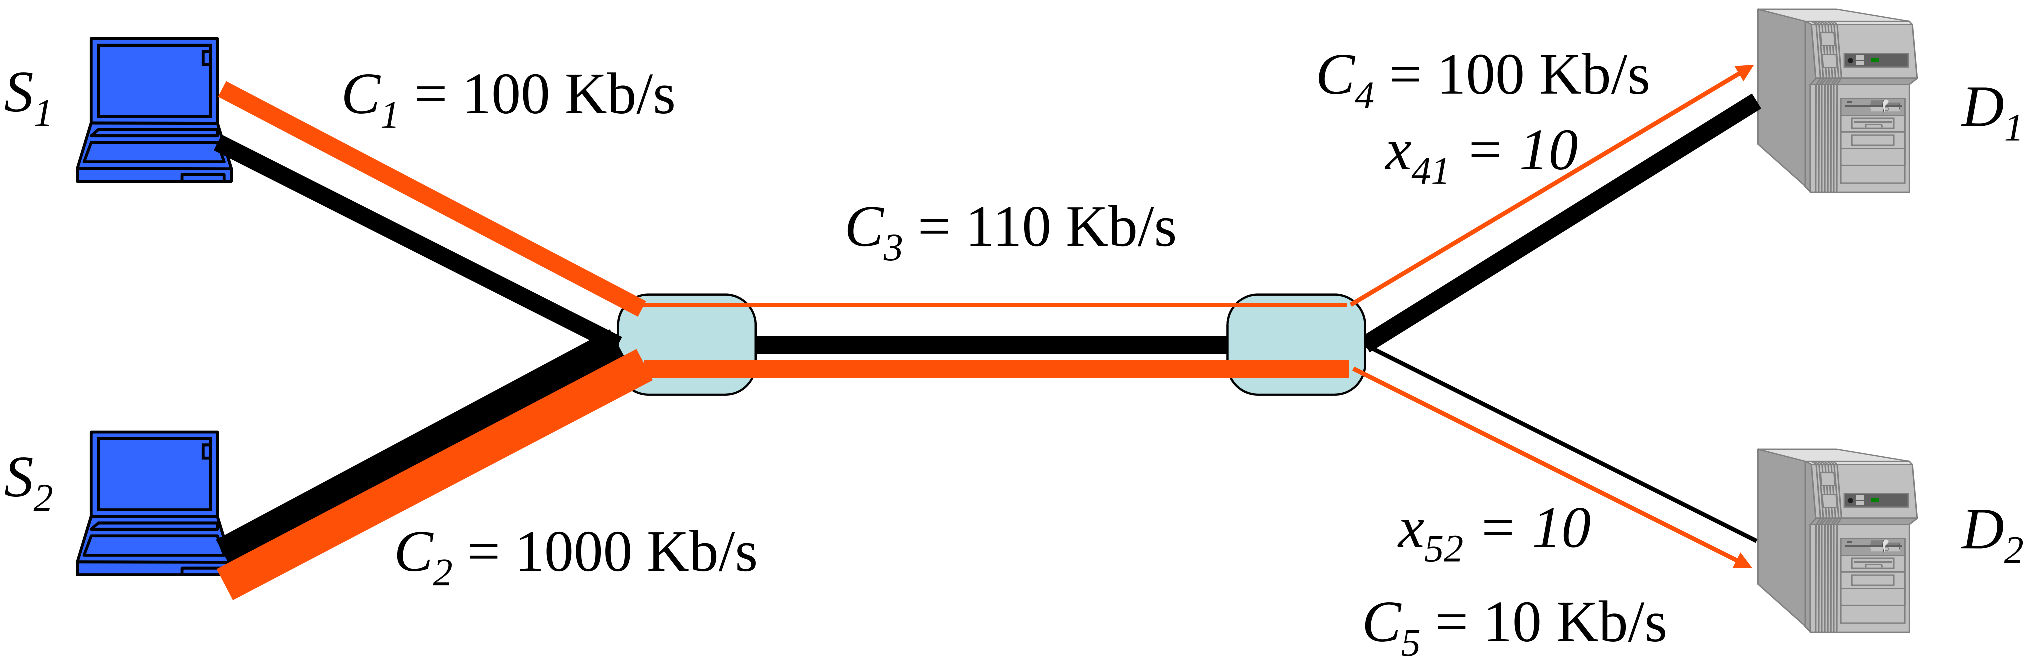

How much can node i send to its destination ?

Source i uses two links, all of same capacity c

At every node there is a source – all sources at same rate

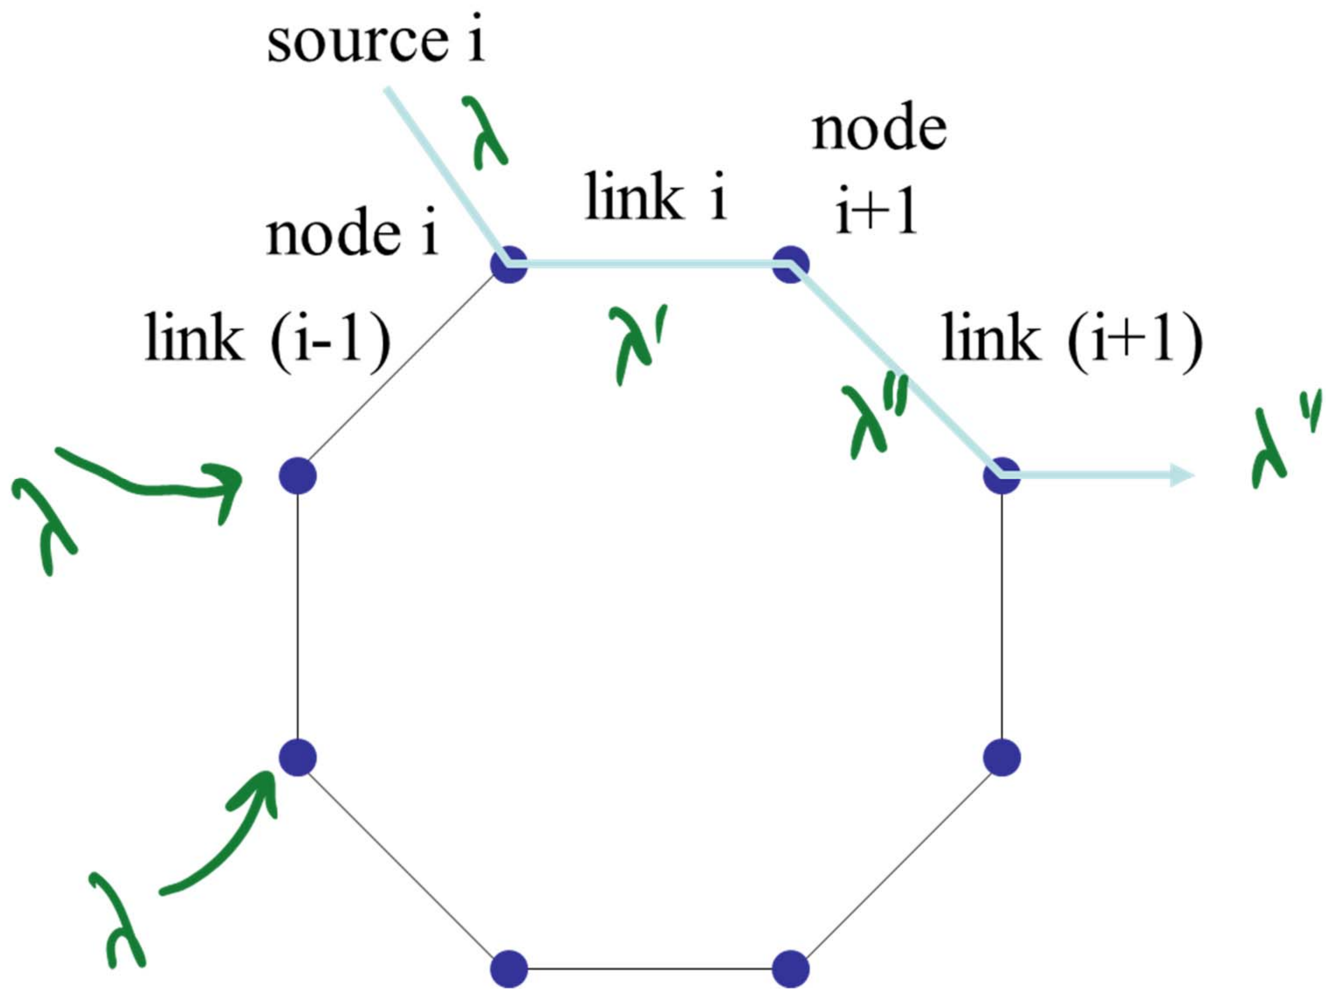

How much can node i send to its destination ?

If $\lambda < \frac{c}{2}$ there is no loss (in this simple model) and $\lambda'' = \lambda$

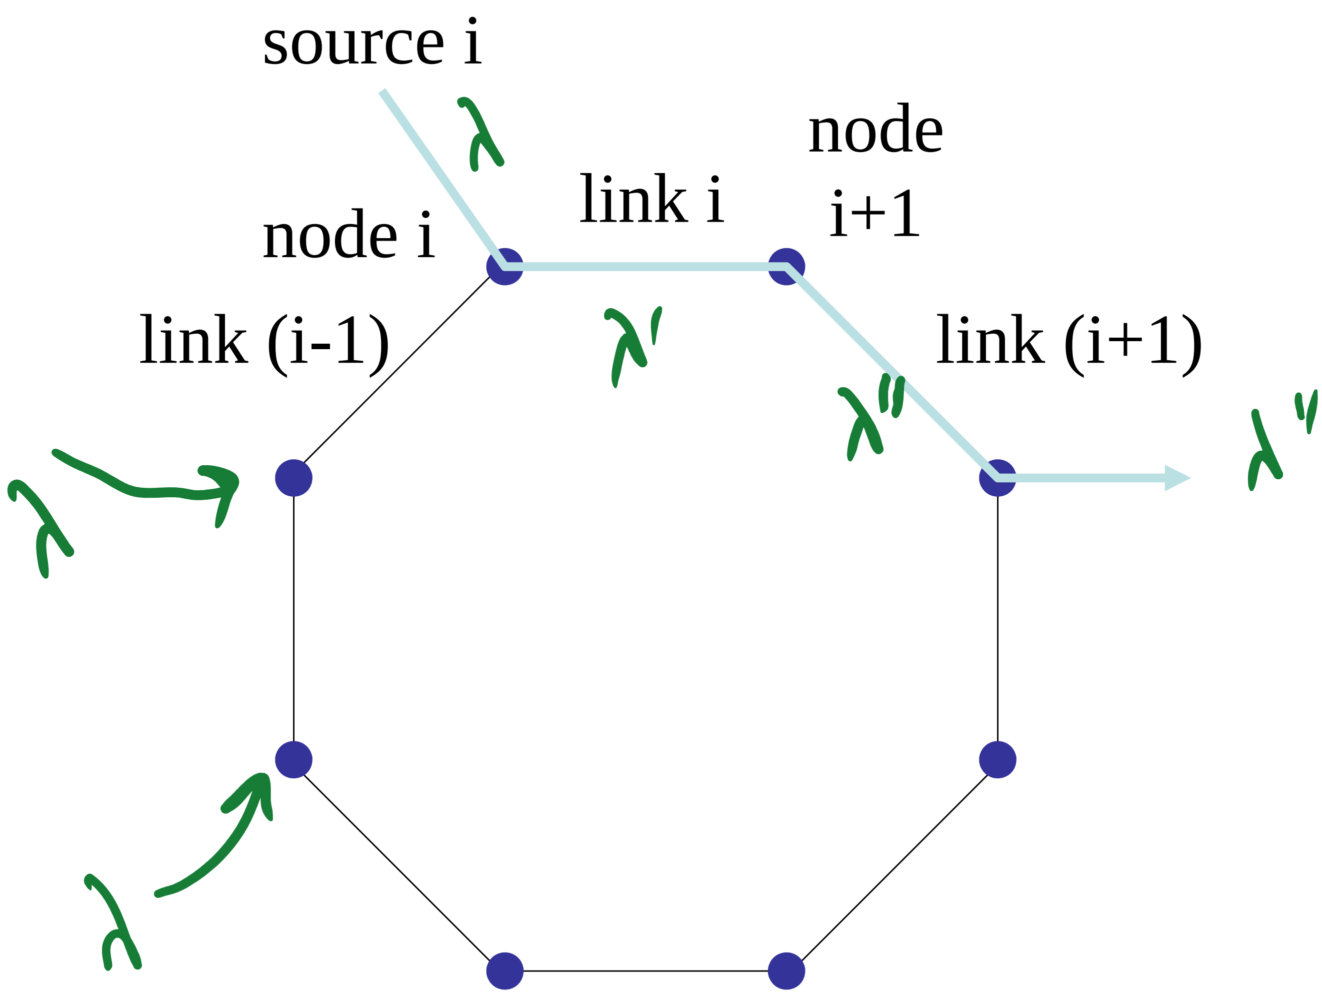

How much can node i send to its destination ?

If $\lambda > \frac{c}{2}$ there is some loss (in this simple model)

Ratio of accepted traffic at one node is $\frac{c}{\lambda+\lambda'}$ (capacity / offered traffic)

Therefore

$$\lambda' = \frac{c}{\lambda+\lambda'}\lambda \quad (1)$$

$$\lambda'' = \frac{c}{\lambda+\lambda'}\lambda' \quad (2)$$

Solve for λ',λ'' from (1),(2)

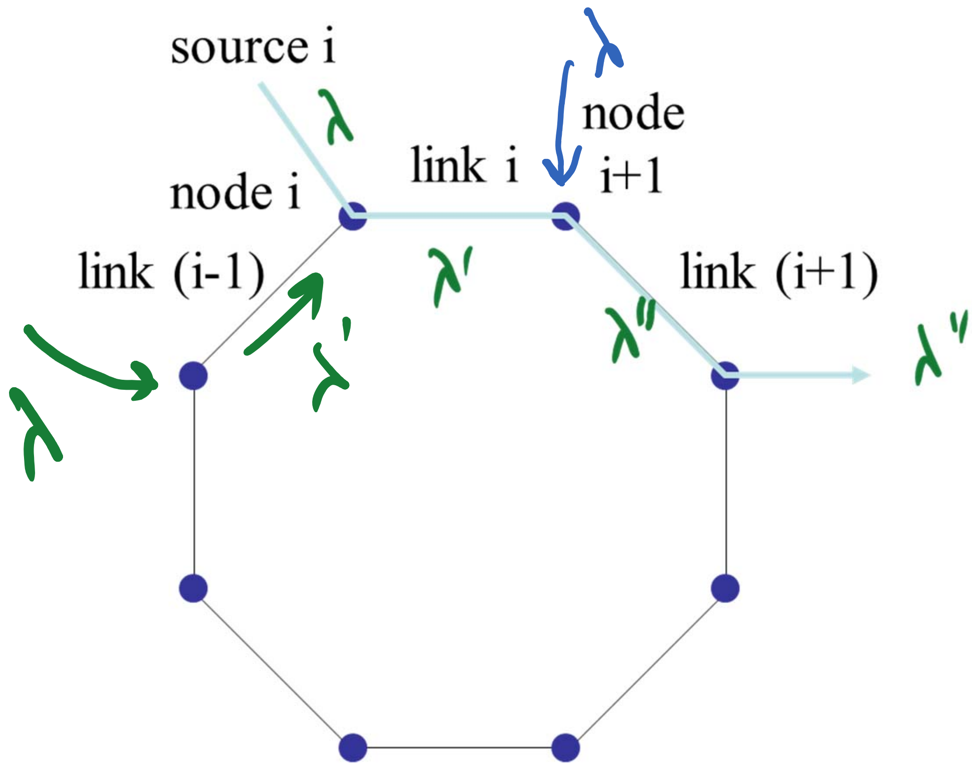

How much can node i send to its destination ?

We obtain, for $\lambda > \frac{c}{2}$:

$$\lambda'' = c - \frac{\lambda}{2}\left(-1 + \sqrt{1 + \frac{4c}{\lambda}}\right)$$

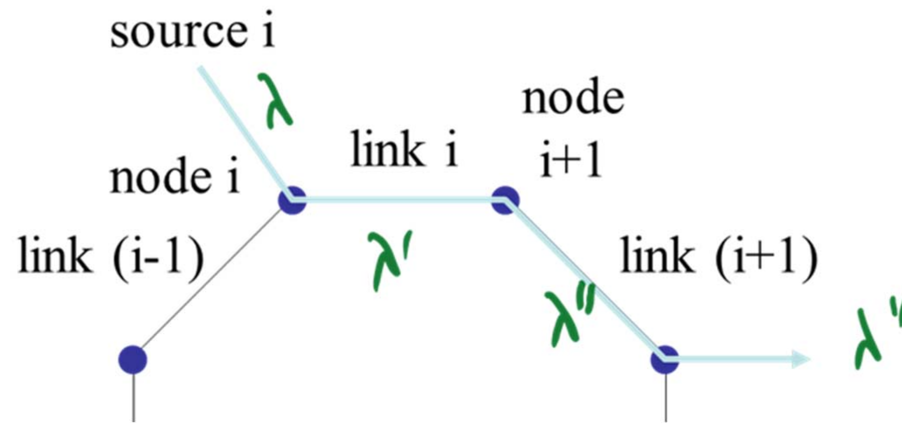

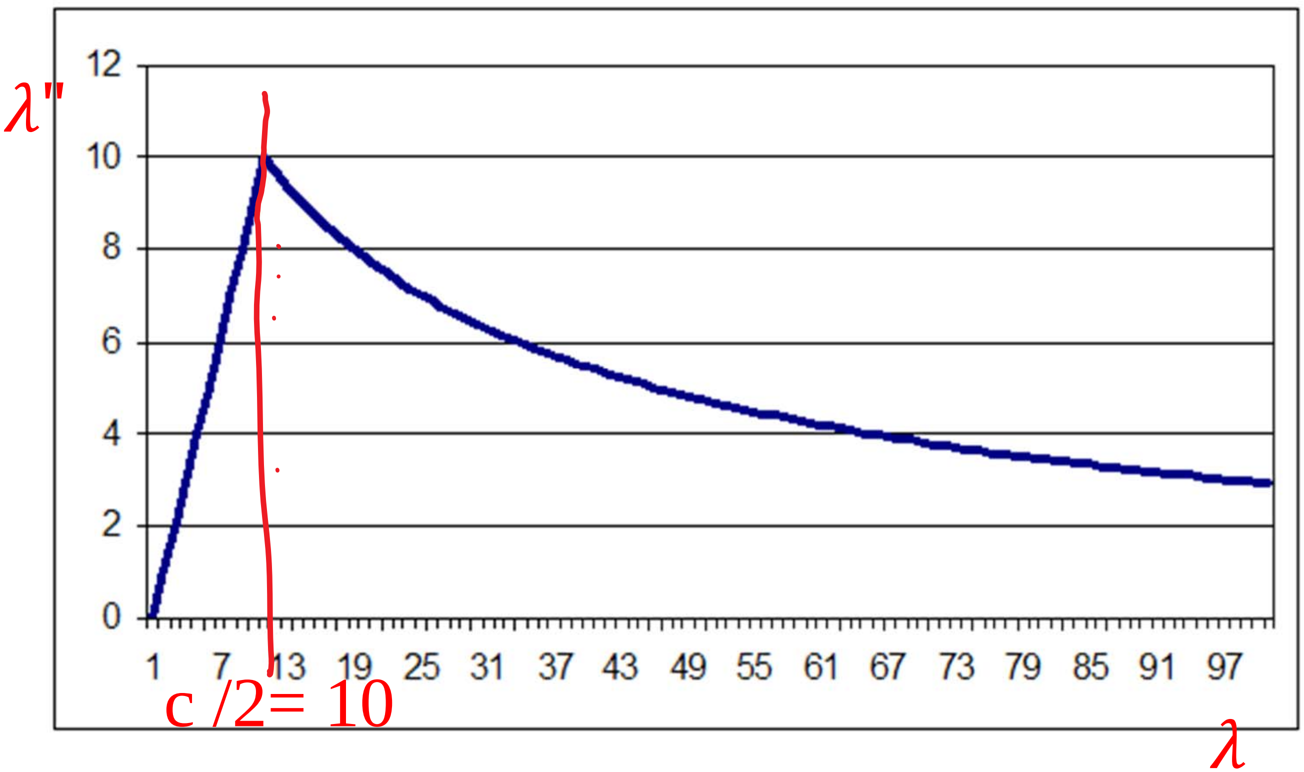

For large offered traffic λ, the limit of useful work is 0

We obtain, for $\lambda > \frac{c}{2}$:

$$\lambda'' = c - \frac{\lambda}{2}\left(-1 + \sqrt{1 + \frac{4c}{\lambda}}\right)$$

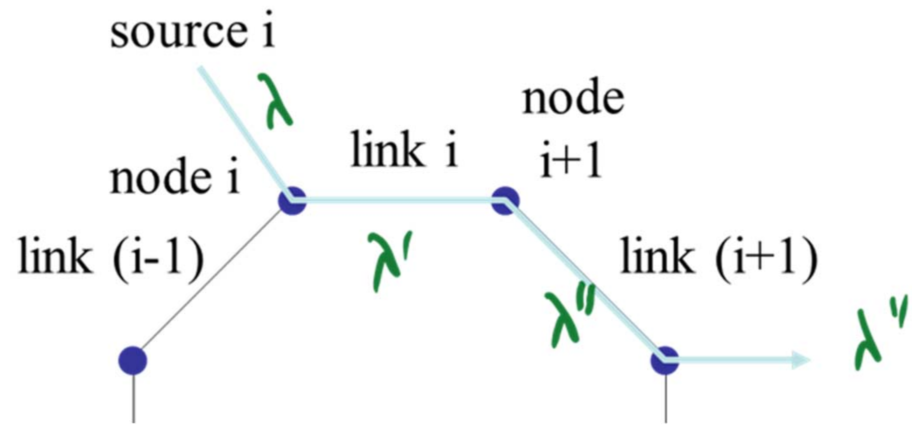

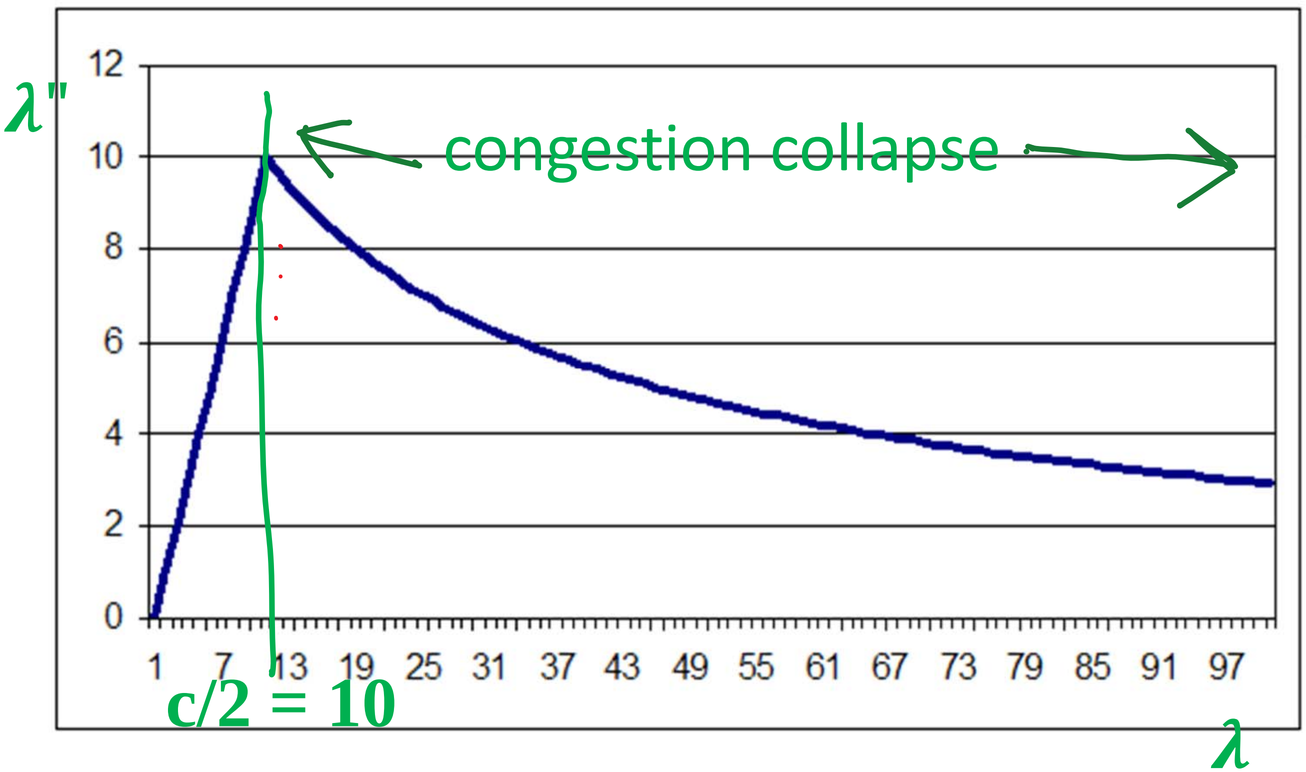

$$\sqrt{1+u} = 1 + \frac{1}{2}u - \frac{1}{8}u^2 + o(u^2)$$

$$\lambda'' = \frac{c^2}{\lambda} + o(\frac{1}{\lambda})$$

Take-Home Message 2

This *is* congestion collapse, i.e. : as the offered load increases, the total throughput decreases

Sources should limit their rates to adapt it to the network condition

Otherwise inefficiency or congestion collapse may occur

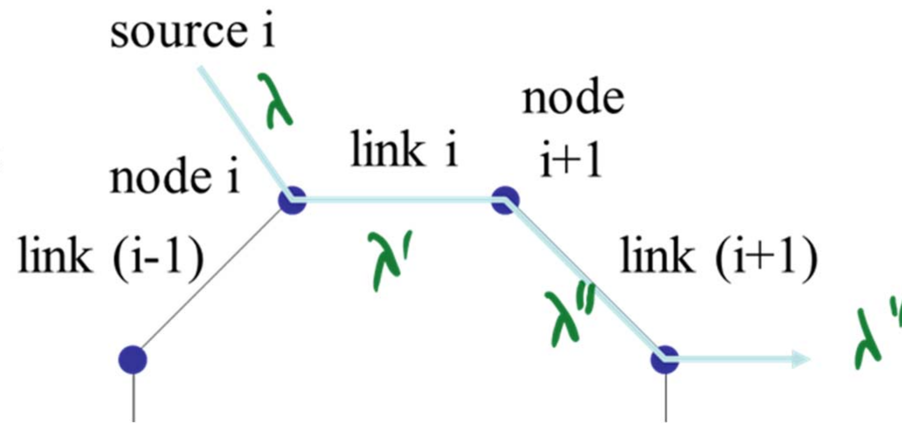

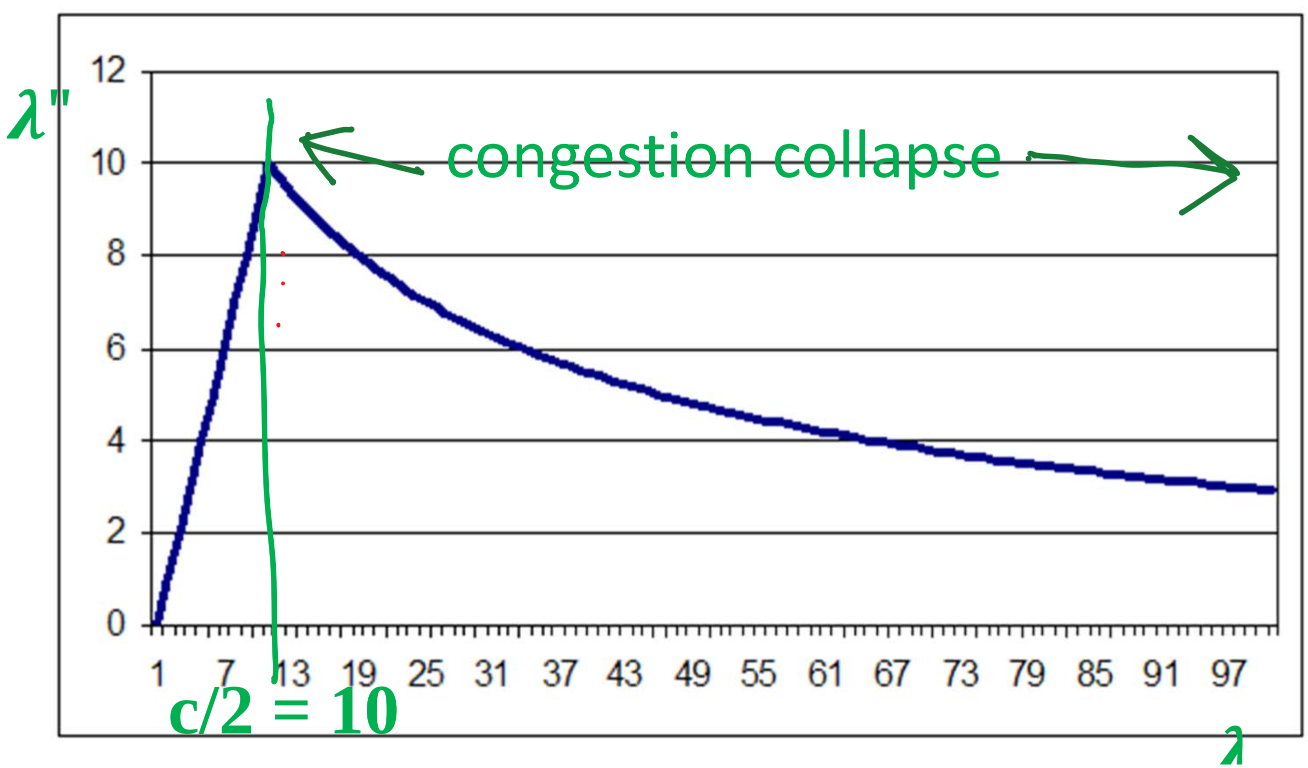

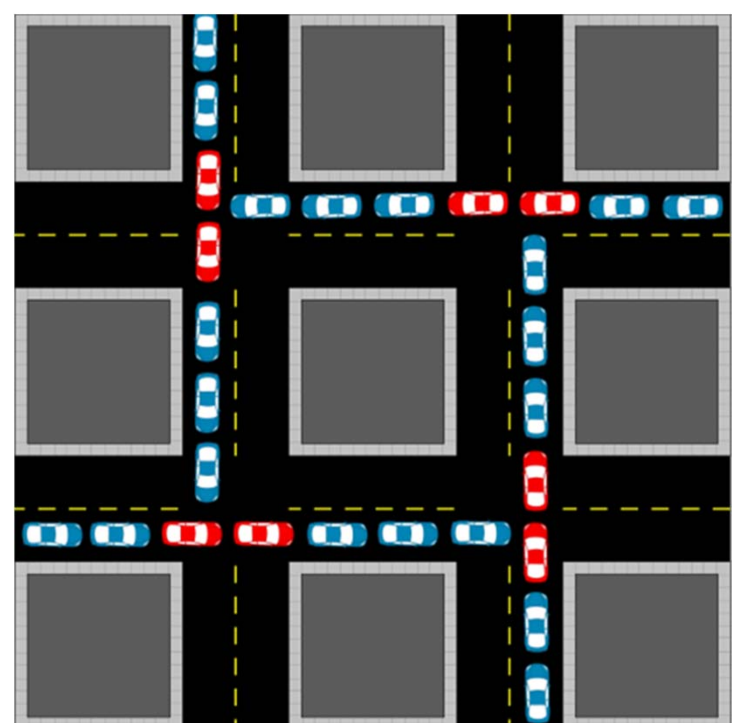

2. Efficiency vs Fairness

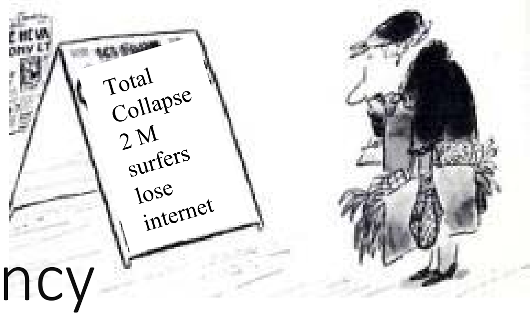

A network should be organized so as to avoid inefficiency

However, being maximally efficient may be a problem

Example : what is the maximum total throughput in this network ?

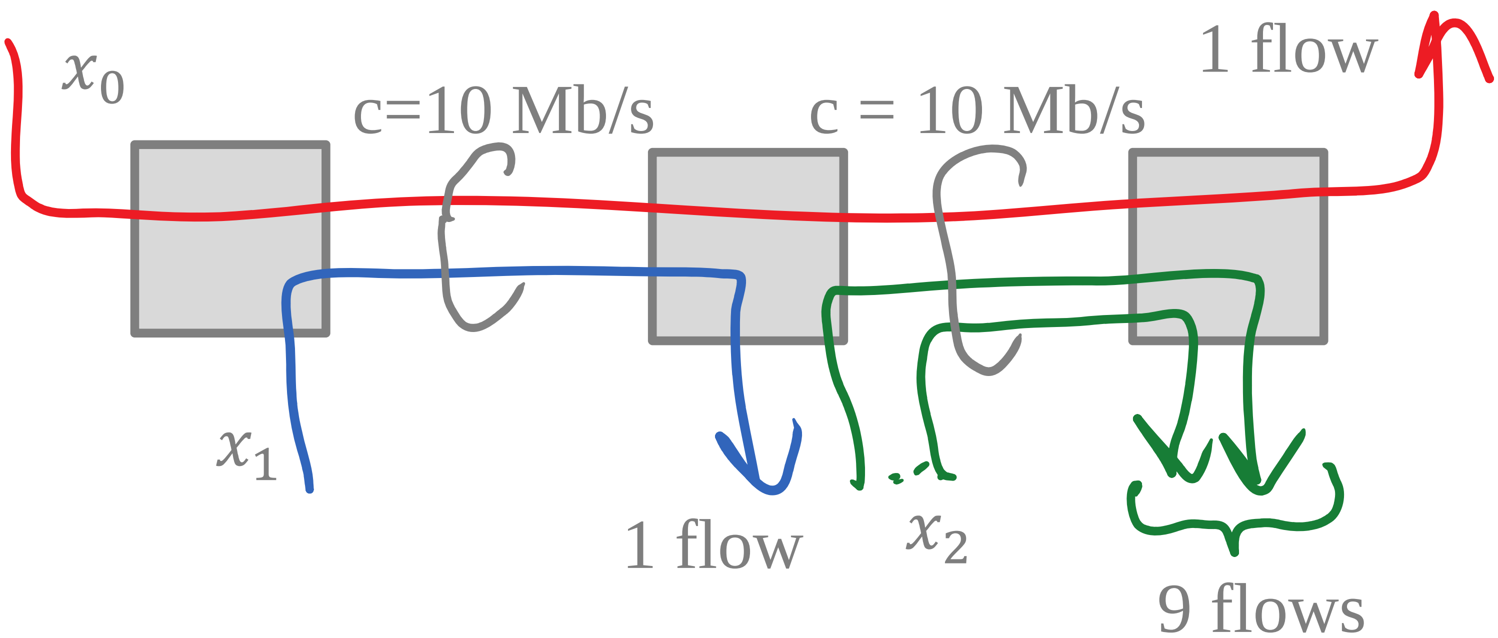

The maximum throughput is …

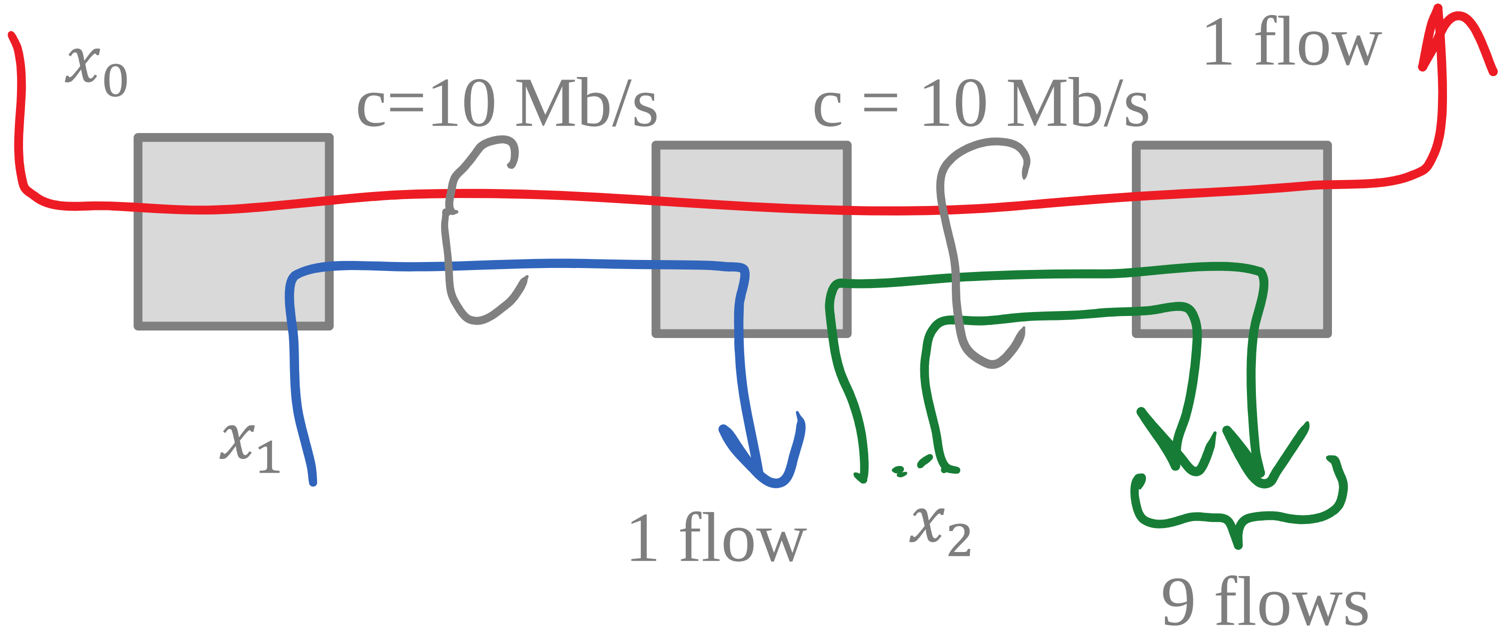

A. 5 Mb/s
B. 10 Mb/s
C. 20 Mb/s
D. None of the above
E. I don’t know

Solution

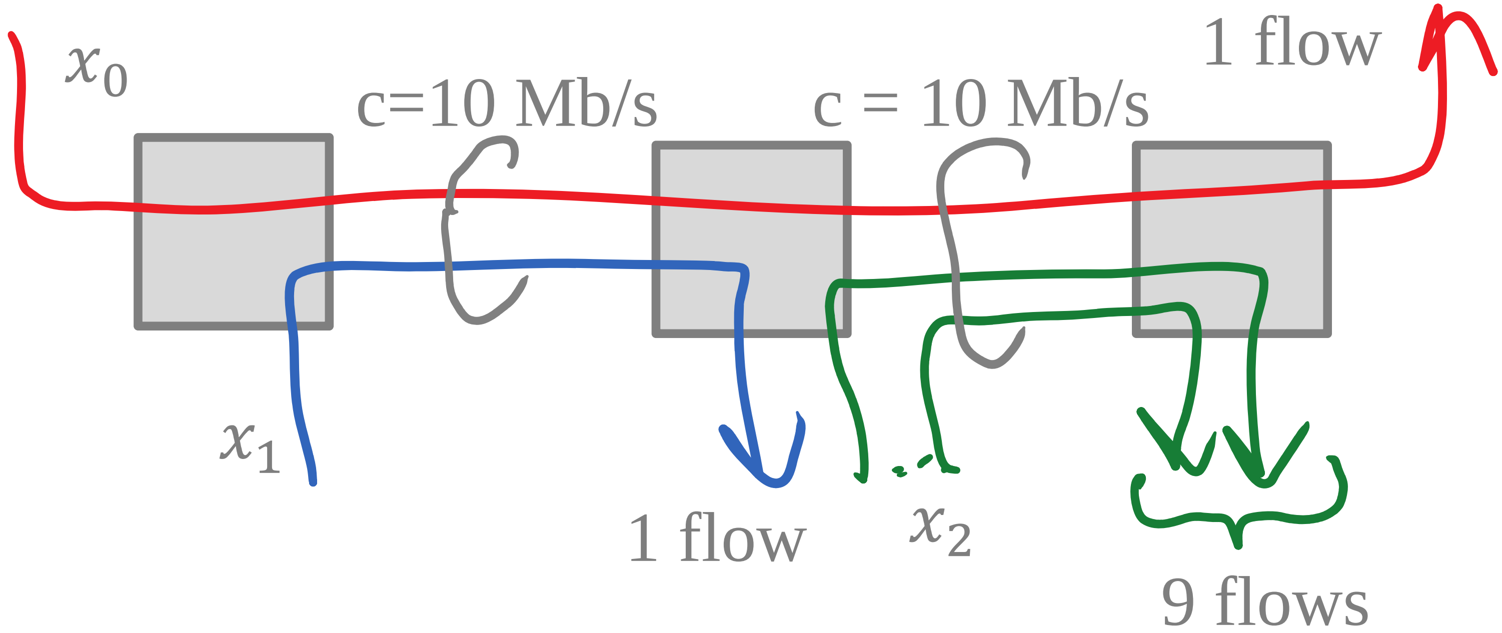

Answer C

Total throughput $\theta = x_0 + x_1 + 9x_2$

Maximize $\theta = x_0 + x_1 + 9x_2$

subject to $x_0 + x_1 \leq 10, x_0 + 9x_2 \leq 10$

over $x_0 \geq 0, x_1 \geq 0\ x_2 \geq 0$

The max can be obtained by linear programming, or directly here by inspection:

$\theta \leq\ 20$ because $x_0 + x_1 \leq 10$ and $9\ x_2 \leq 10$

$\theta = 20$ is achieved with $x_1 = 10$ and $x_2 = \frac{10}{9}$

therefore the max is 20 Mb/s

The value of x_0 when the maximum throughput is attained is …

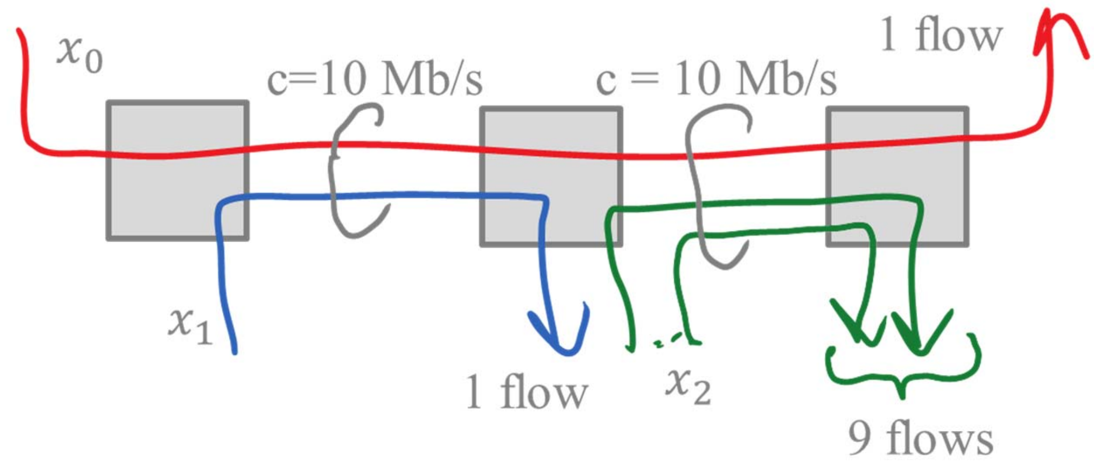

A. 1 Mb/s

B. $\frac{10}{9}$ Mb/s

C. 2 Mb/s

D. None of the above

E. I don't know

Solution

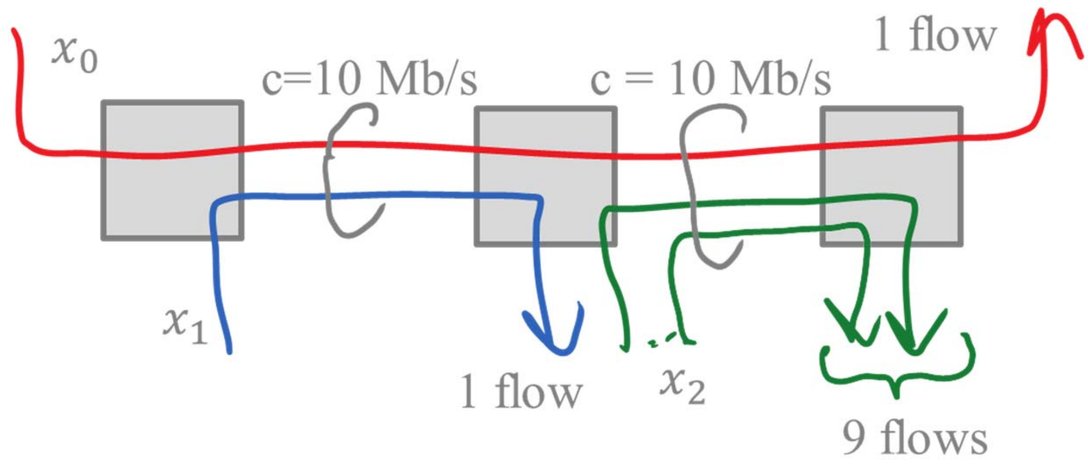

Answer D

Find all $x_0 \geq 0$ $x_1 \geq 0$ $x_2 \geq 0$ subject to

$$x_0 + x_1 \leq 10 \quad (1)$$
$$x_0 + 9x_2 \leq 10 \; (2)$$
$$x_0 + x_1 + 9x_2 = 20 \; (3)$$

By (1) and (3): $9x_2 \geq 10$

Compare to (2): $9x_2 = 10$

Thus $x_0 = 0$ (and $x_1 = 10, x_2 = 10/9$)

Pareto Efficiency

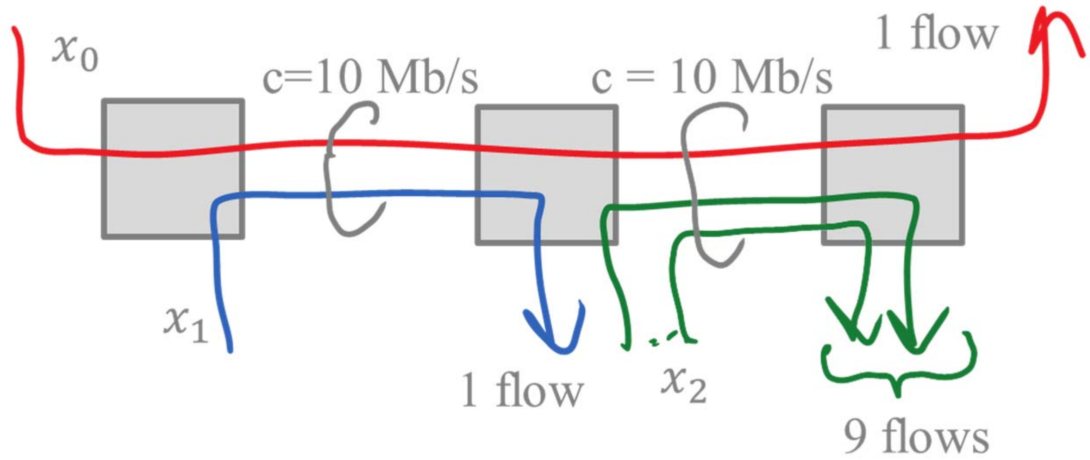

A feasible allocation of rates $\vec{x}$ is called Pareto-efficient (synonym: Pareto-optimal) iff increasing one source must be at the expense of decreasing some other source

i.e. $\vec{x}$ is Pareto-efficient iff :

$$\text{for any other feasible } \vec{x}', \exists i: x_i' > x_i \Rightarrow \exists j: x_j' < x_j$$

$\Leftrightarrow$ every source has a bottleneck link (i.e. for every source i there exists a link, used by i, which is saturated)

Which allocation is Pareto-Efficient ?

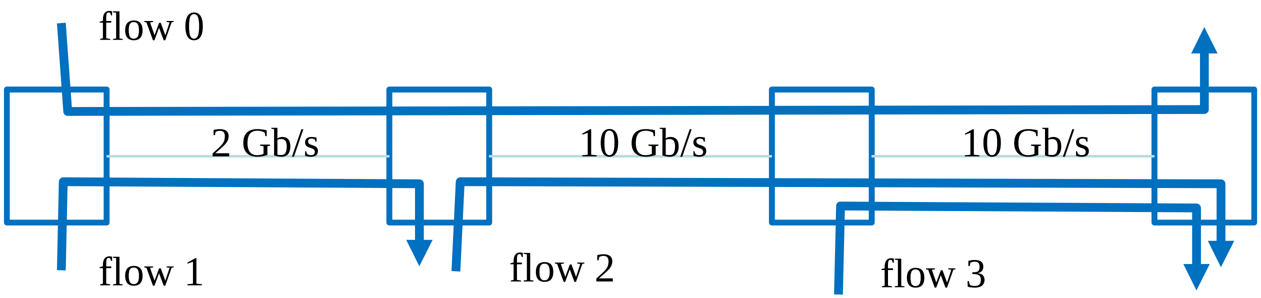

A. $x_0 = 1, x_1 = 1, x_2 = 2, x_3 = 7$
B. $x_0 = 1, x_1 = 1, x_2 = 4.5, x_3 = 4.5$
C. Both
D. None
E. I don't know

Solution

Answer C

A: first link is saturated ⇒ flows 0 and 1 have a bottleneck

third link is saturated ⇒ flows 0, 2 and 3 have a bottleneck

All flows have a bottleneck, A is Pareto-efficient

Same for B.

Which allocation is Pareto-Efficient ?

A. $x_0 = 0, x_1 = 10, x_2 = 10/9$
B. $x_0 = 0.55, x_1 = 9.45, x_2 = 1.05$
C. $x_0 = 1, x_1 = 9, x_2 = 1$
D. A and B
E. A and C
F. B and C
G. All
H. None
I. I don't know

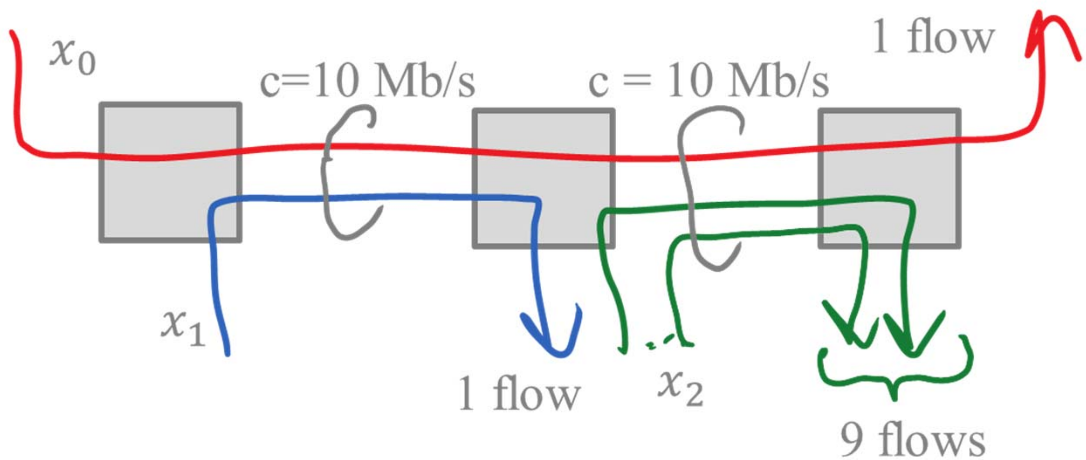

Solution

Answer G

The feasible $\vec{x}$ are such that $x_0 + x_1 \leq 10 \;\; (1), x_0 + 9x_2 \leq 10 \; (2)$

The Pareto-efficient $\vec{x}$ are such that $x_0 + x_1 = 10$ and $x_0 + 9x_2 = 10$

Indeed, if there is strict inequality in (1), we can increase x_1 without modifying the other values, so $\vec{x}$ is not Pareto-efficient. Similarly, if there is strict inequality in (2), we can increase x_2 without modifying the other values. So, to be Pareto-efficient, it is necessary to have both equalities. Conversely, if both equalities hold, any increase in one variable implies a decrease in some other variable and $\vec{x}$ is Pareto-efficient.

Here, allocations A,B and C satisfy the condition and are all Pareto efficient.

The Pareto efficient allocations are the ones that use the resources maximally.

Solution

Alternatively, we can see that every source has a bottleneck.

A: first link is a bottleneck for source 0 and source 1. Second link is a bottleneck for sources 2.

Same for B and C.

Which allocation is Pareto-Efficient ?

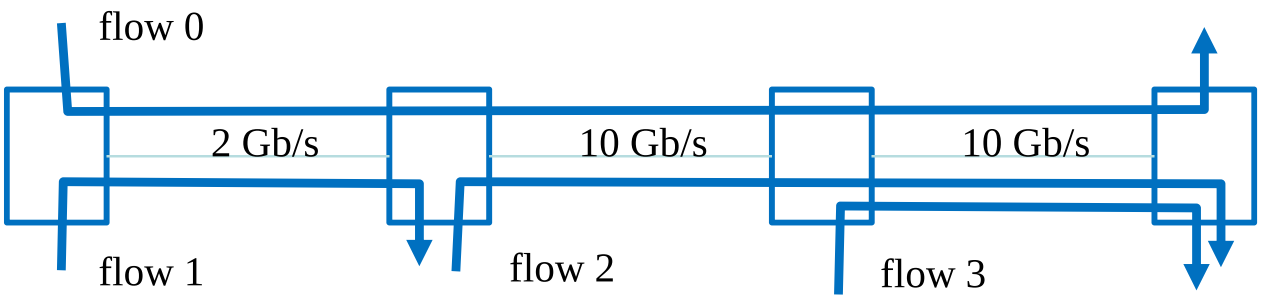

A. $x_0 = 1, x_1 = 1, x_2 = 2, x_3 = 7$
B. $x_0 = 1, x_1 = 1, x_2 = 4.5, x_3 = 4.5$
C. Both
D. None
E. I don't know

Solution

Answer C
A: first link is saturated ⇒ flows 0 and 1 have a bottleneck
third link is saturated ⇒ flows 0, 2 and 3 have a bottleneck
All flows have a bottleneck, A is Pareto-efficient

Same for B.

Take Home Message

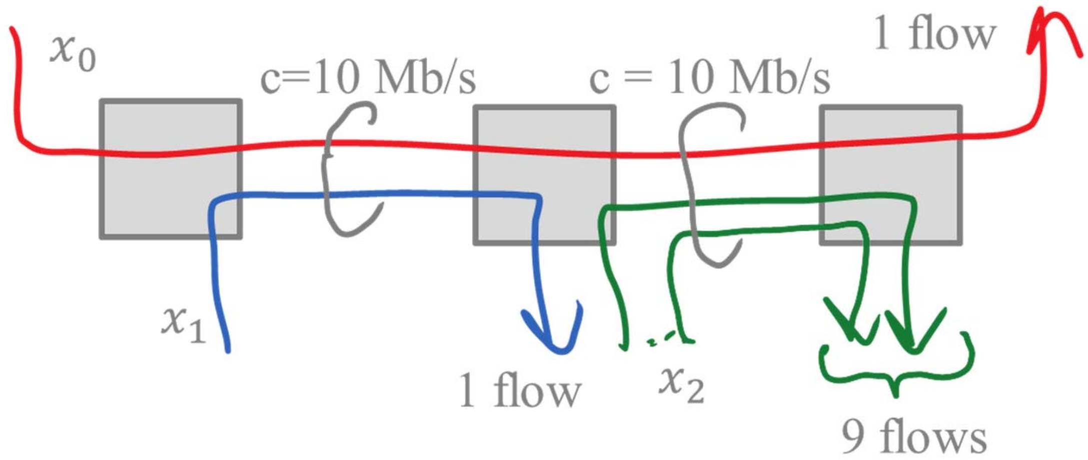

Maximal efficiency means Pareto optimality.

Maximizing total throughput is Pareto optimal, but it means shutting down flow 0; this is at the expense of fairness.

Are there Pareto efficient allocations that are fair ?

What is fairness ?

3. Let us try a first definition (Egalitarianism) : allocate as much as possible but same to all.
In this example, what is the allocation according to this definition ?

A. $x_0 = x_1 = x_2 = 0.5\ Mb/s$
B. $x_0 = x_1 = x_2 = 1\ Mb/s$
C. $x_0 = x_1 = x_2 = \frac{10}{9}\ Mb/s$
D. None of the above
E. I don't know

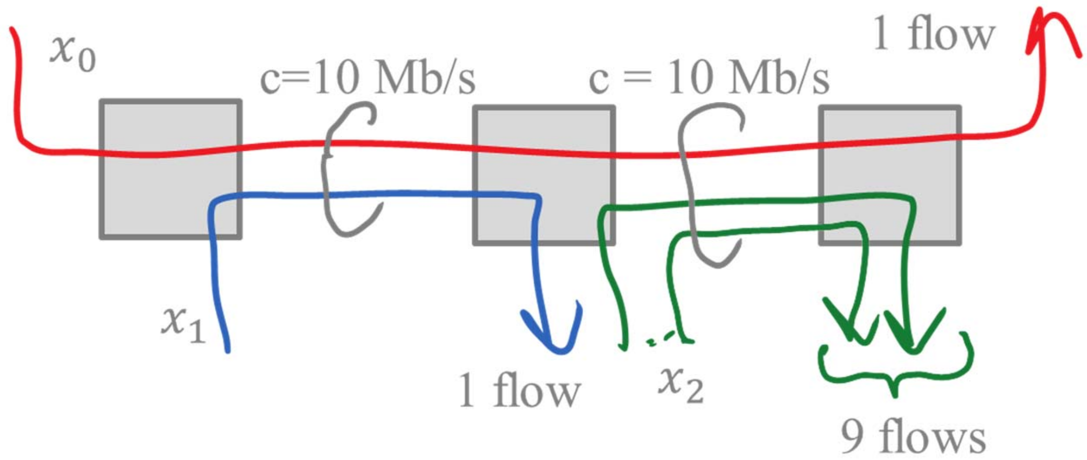

Solution

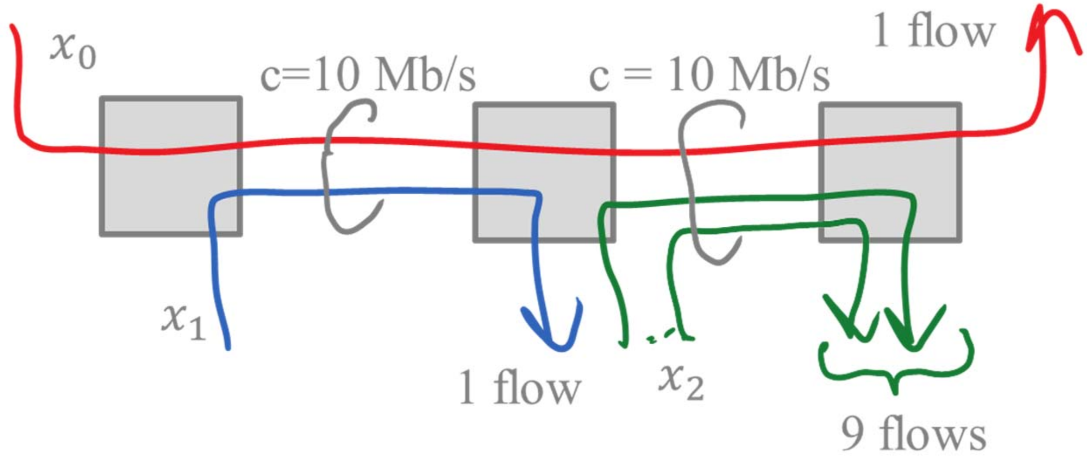

Maximize $x = x_0 = x_1 = x_2$ subject to

$$2\,x \leq 10$$
$$10\,x \leq 10$$

with $x \geq 0$

The solution is $x = 1$ Mb/s

Answer B

Egalitarianism is not Pareto-efficient

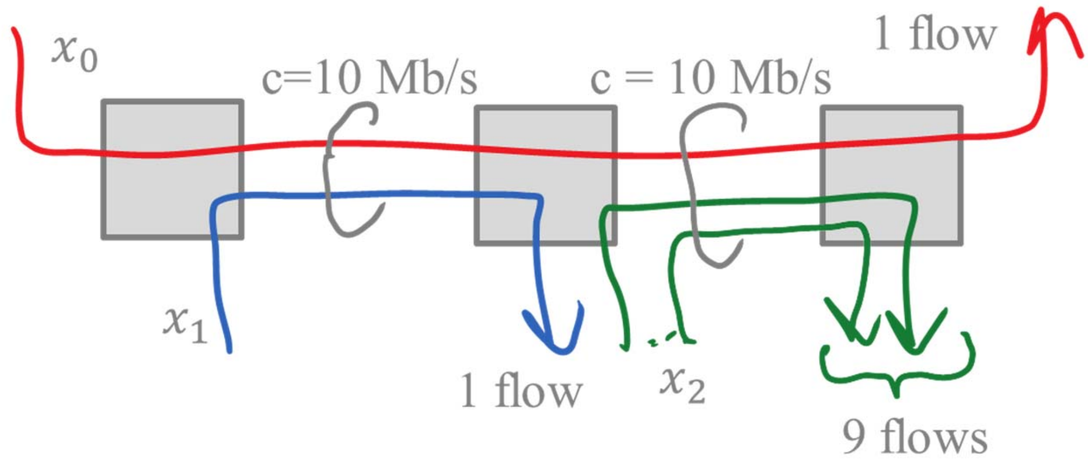

Egalitarianism gives $x = 1$ Mb/s to all
this is stupid, we could give more to x_1 without hurting anyone

A better allocation is:

$$x_0 = 1, \qquad x_1 = 9, \qquad x_2 = 1$$

It is Pareto-efficient (all constraints are satisfied with equality) and is fair, since it gives to every one at least as much as egalitarianism

This is the max-min fair allocation for this example.

Max-Min fairness

We say that a feasible allocation $\vec{x}$ is *max-min fair* iff for any other feasible $\vec{x}'$, $(\exists i\colon\ x_i' > x_i \Rightarrow \exists j\colon\ x_j' < x_j$ and $x_j \leq x_i)$

i.e. an allocation is max-min fair if the following is true

For every source i, increasing its rate must force the rate of some other (not richer) source j to decrease

In other words: an allocation is max-min fair if ***any rate increase contradicts fairness***

Which allocations are max-min fair ?

A

$x_0 = 0$ Mb/s

$x_1 = 10$ Mb/s

$x_2 = \frac{10}{9}$ Mb/s

B

$x_0 = 1$ Mb/s

$x_1 = 9$ Mb/s

$x_2 = 1$ Mb/s

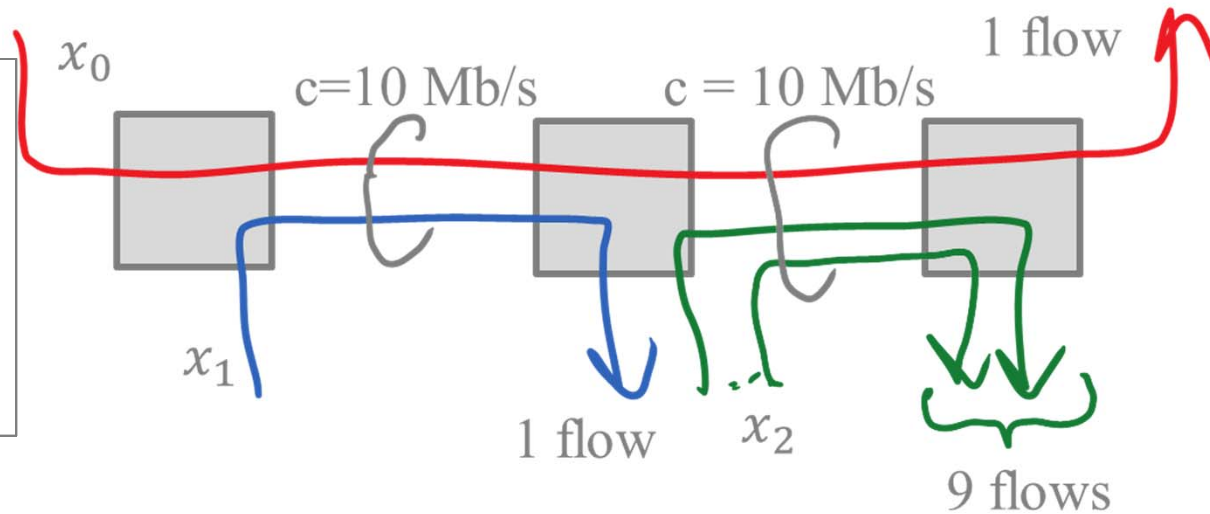

A. A

B. B

C. A and B

D. None

E. I don’t know

Solution

A
$x_0 = 0$ Mb/s
$x_1 = 10$ Mb/s
$x_2 = \frac{10}{9}$ Mb/s

B
$x_0 = 1$ Mb/s
$x_1 = 9$ Mb/s
$x_2 = 1$ Mb/s

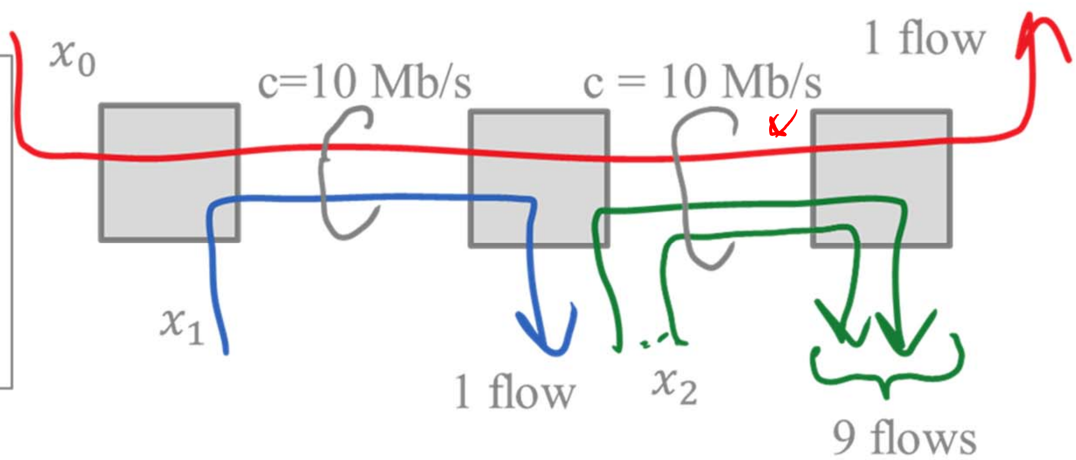

Allocation A

I can increase x_0 (e.g. $x_0 \leftarrow 1$) and decrease x_1 ($x_1 \leftarrow 9$) and x_2 ($x_2 \leftarrow 1$); this does not contradict fairness because x_1 and x_2 are larger than x_0

There exists one increase that does not contradict fairness

A is not max-min fair

Allocation B

If I increase x_0 I must decrease x_2 $\Rightarrow$ contradicts fairness

If I increase x_1 I must decrease x_0 $\Rightarrow$ contradicts fairness

If I increase x_2 I must decrease x_0 $\Rightarrow$ contradicts fairness

Any increase contradicts fairness

B is max-min fair

The Maths of Max-Min Fairness

Given a set of constraints for the rates

- If it exists, the max-min fair allocation is *unique*
- There *exists* one max-min fair allocation if the set of feasible rates is convex (this is the case for networks, as we have linear constraints)
- The max-min fair allocation is Pareto-efficient (converse is not true)

For a set of feasible rates as in our case (the sum of the rates on every link is upper bounded), the (unique) max min fair allocation is obtained by *water-filling*

1 mark all sources as non frozen
2 **Do**
3 increase the rate of all non frozen sources to the largest possible common value
4 mark sources that use a saturated link as frozen
5 **Until** all sources are frozen

Water-Filling Example

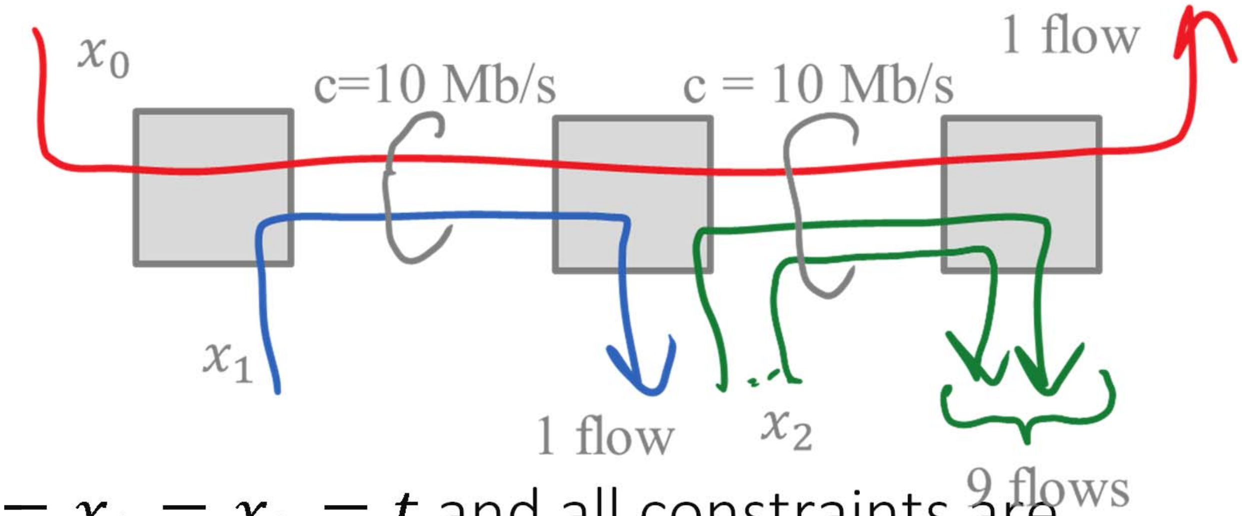

Step 1:

- maximize t such that $x_0 = x_1 = x_2 = t$ and all constraints are satisfied; we find $t = 1$, hence $x_0 = x_1 = x_2 = 1$;
- link 2 is saturated, is used by sources 0 and 2 ⇒ mark sources 0 and 2 as frozen

Step 2 :

- maximize t such that $x_1 = t$, with $x_0 = 1, x_2 = 1$ and all constraints are satisfied; we find $t = 9$, hence $x_0 = x_2 = 1$ and $x_1 = 9$
- link 1 is saturated, is used by sources 0 and 1 ⇒ mark source 1 as frozen; all sources are frozen, STOP.

The max-min fair allocation is $x_0 = x_2 = 1$ and $x_1 = 9$

What is the max-min fair allocation ?

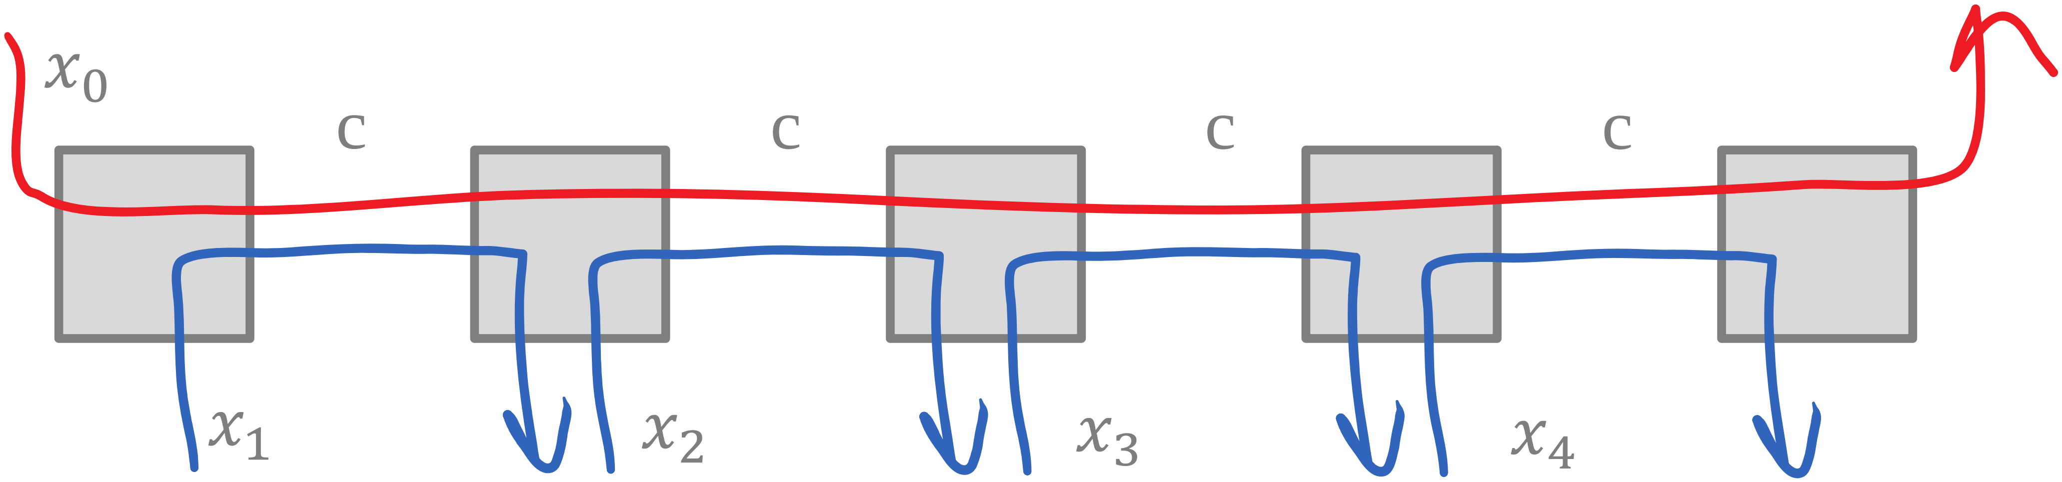

A. $x_i = \frac{c}{2} \ \forall i$

B. $x_0 = \frac{c}{3}, \quad x_i = \frac{2c}{3} \ \forall i \neq 0$

C. $x_0 = \frac{c}{4}, \quad x_i = \frac{3c}{4} \ \forall i \neq 0$

D. $x_0 = \frac{c}{5}, \quad x_i = \frac{4c}{5} \ \forall i \neq 0$

E. None of the above

F. I don’t know

Solution

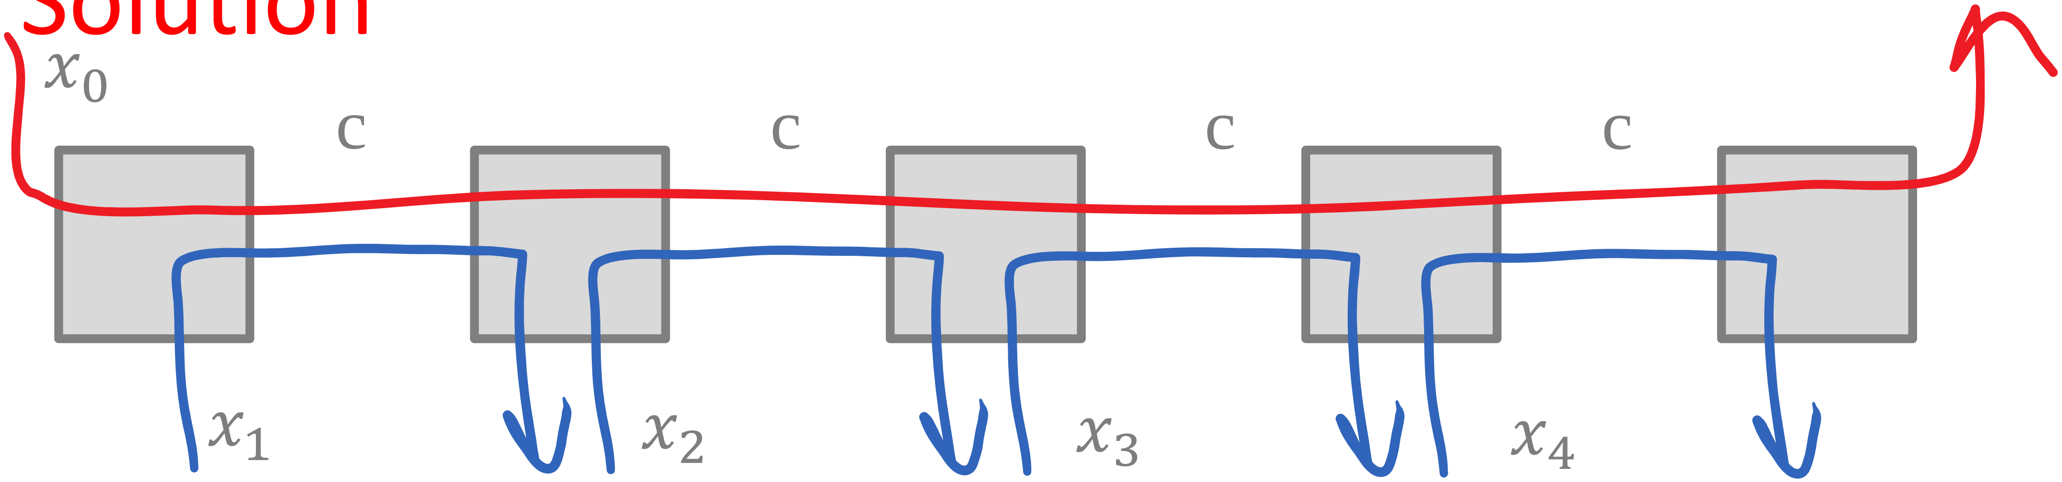

Answer A

The max-min fair allocation gives the same rate $\frac{c}{2}$ to all sources

Some people think that x_0 should be penalized as it uses more of the network

This is what led to the definition of proportional fairness

Definition of Proportional Fairness

We say that a feasible allocation $\vec{x}$ is *proportionally fair* iff for any other feasible $\vec{x}'$, $\sum_i \frac{x_i' - x_i}{x_i} \leq 0$

In other words: an allocation is proportionnally fair if ***any change contradicts proportional fairness***, in the sense that:

the total (or average) *rate of change* $\sum_i \frac{\Delta x_i}{x_i}$ is negative

Two effects

Relative shares matter, not absolute

Global effect

Which allocations are proportionally fair ?

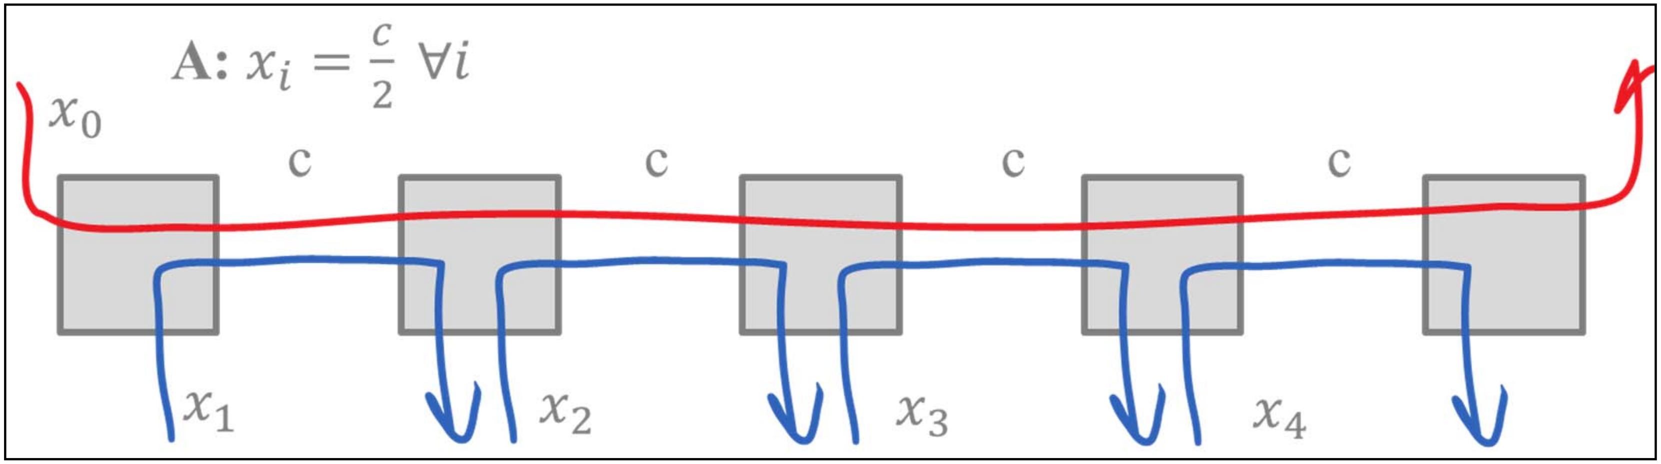

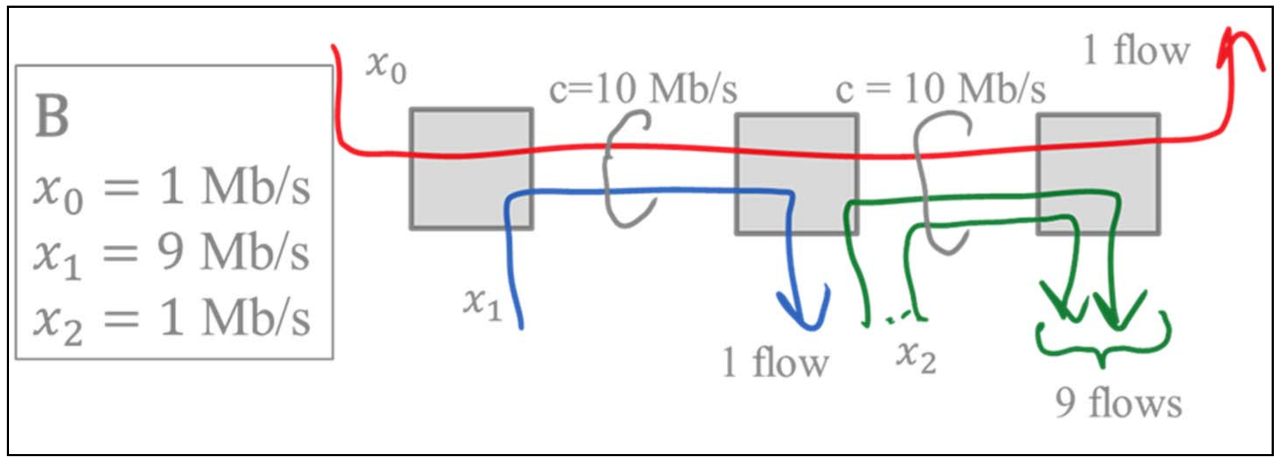

A. A
B. B
C. A and B
D. None
E. I don't know

Solution

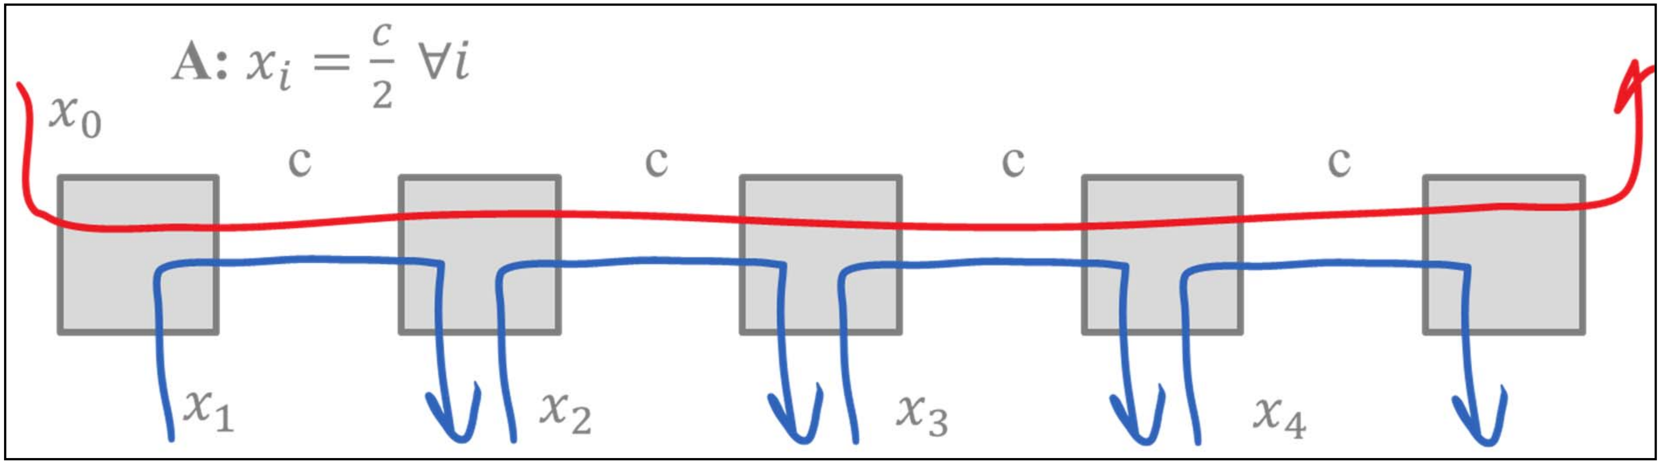

Answer D

I can do $x_i \leftarrow x_i + \delta$ for $i = 1 \ldots 4$ and $x_0 \leftarrow x_0 - \delta$; if δ is small enough (i.e. $0 < \delta \leq \frac{c}{2}$) the new allocation is feasible

The total rate of change is $-\frac{\delta}{\frac{c}{2}} + 4\,\frac{\delta}{\frac{c}{2}} > 0$

We could change the allocation without contradicting proportional fairness

A is not proportionally fair

Solution

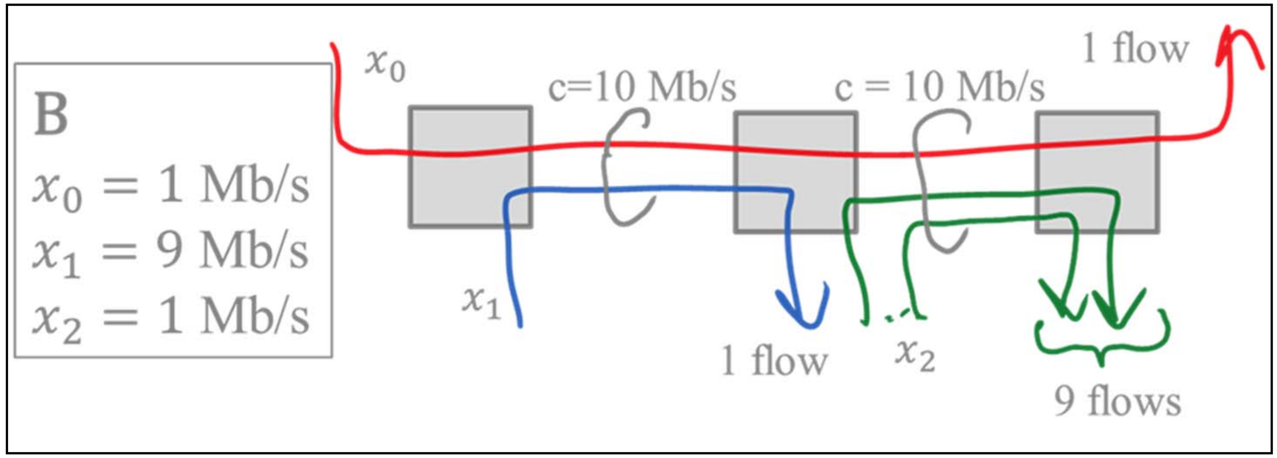

I can do $x_2 \leftarrow x_2 + \delta$ and $x_0 \leftarrow x_0 - 9\delta$; $x_1 \leftarrow x_1 + 9\delta$; if δ is small enough (i.e. $0 < \delta \leq \frac{1}{9}$) the new allocation is feasible

The total rate of change is $-\frac{9\delta}{1} + \frac{9\delta}{9} + 9\ \frac{\delta}{1} = \delta > 0$

We could change the allocation without contradicting proportional fairness

B is not proportionally fair

The Maths of Proportional Fairness

Given a set of constraints for the rates that is convex:

The proportionally fair allocation *exists* and is *unique*

It is obtained by maximizing

$$J(\vec{x}) := \sum_i \log x_i$$

over all feasible allocations

Intuitive explanation:

$$dJ\,(\vec{x}) = \sum_s \frac{dx_i}{x_i}$$

A proportionally fair allocation is Pareto-efficient.

Let us compute the proportionally fair allocation

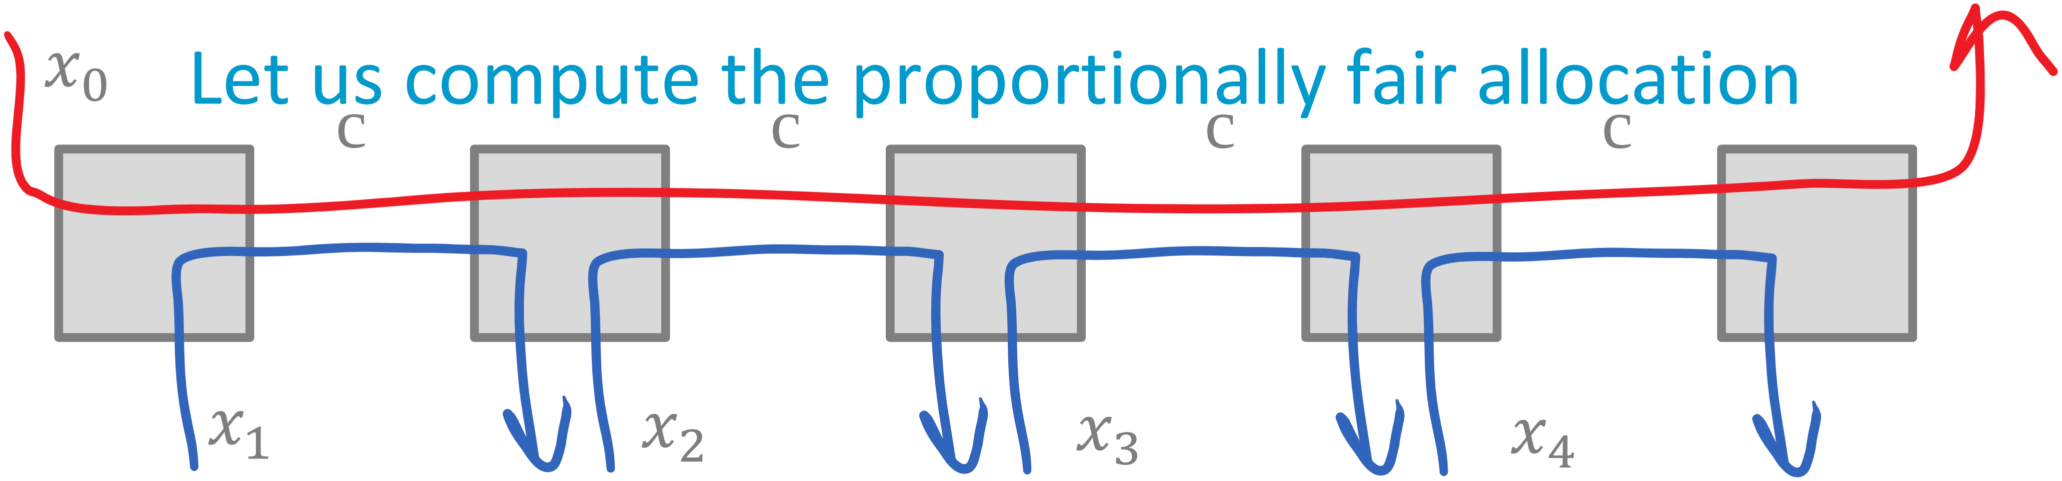

We have to solve the optimization problem:

$$\max U = \log x_0 + \log x_1 + \log x_2 + \log x_3 + \log x_4$$

subject to

$$x_0 + x_1 \le c$$
$$x_0 + x_2 \le c$$
$$x_0 + x_3 \le c$$
$$x_0 + x_4 \le c$$

We can use convex programming, but here we can also do a direct solution

Let us compute the proportionally fair allocation

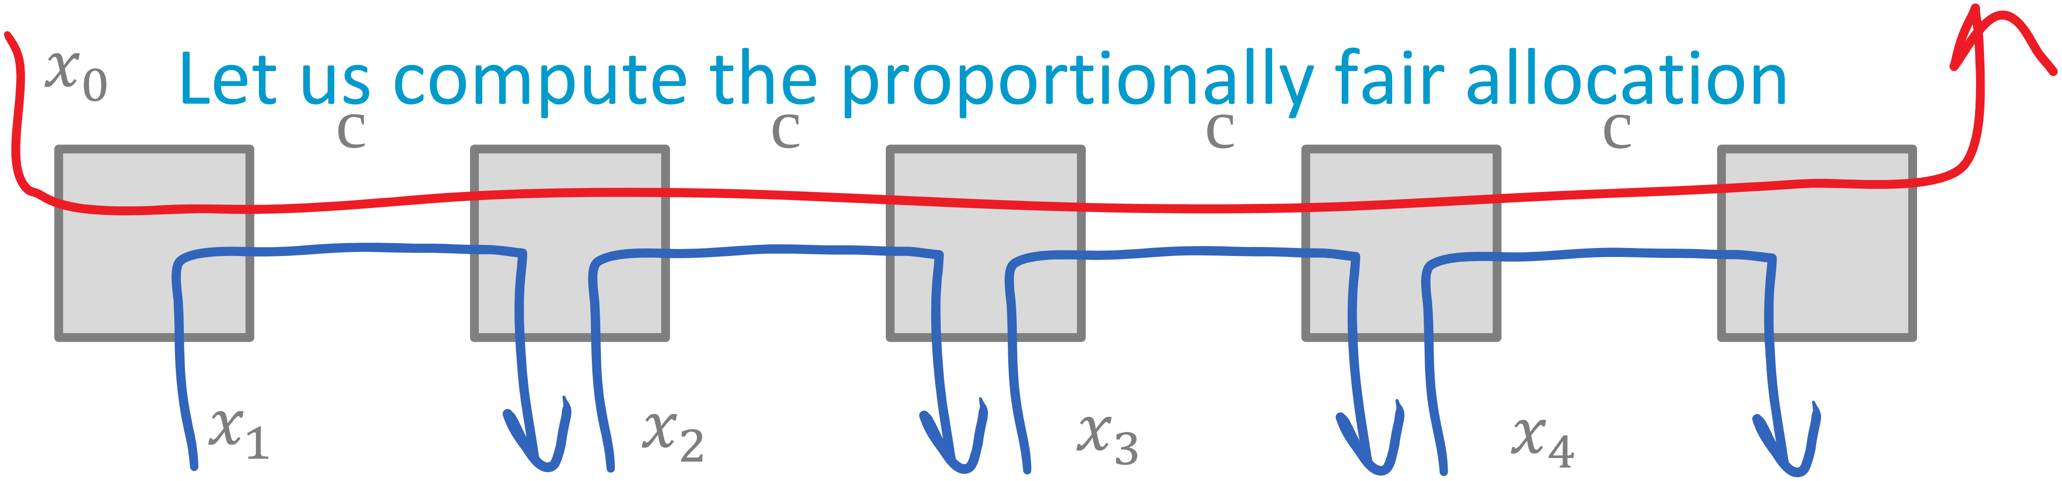

We have to solve the optimization problem:

$$\max U = \log x_0 + \log x_1 + \log x_2 + \log x_3 + \log x_4$$

subject to

$$x_0 + x_1 \le c$$
$$x_0 + x_2 \le c$$
$$x_0 + x_3 \le c$$
$$x_0 + x_4 \le c$$

We must have equality everywhere otherwise we can increase x_i $(i \neq 0)$ and increase U

Therefore we must have $x_1 = x_2 = x_3 = x_4 = c - x^*$ and $x_0 = x^*$ for some x^*

Let us compute the proportionally fair allocation

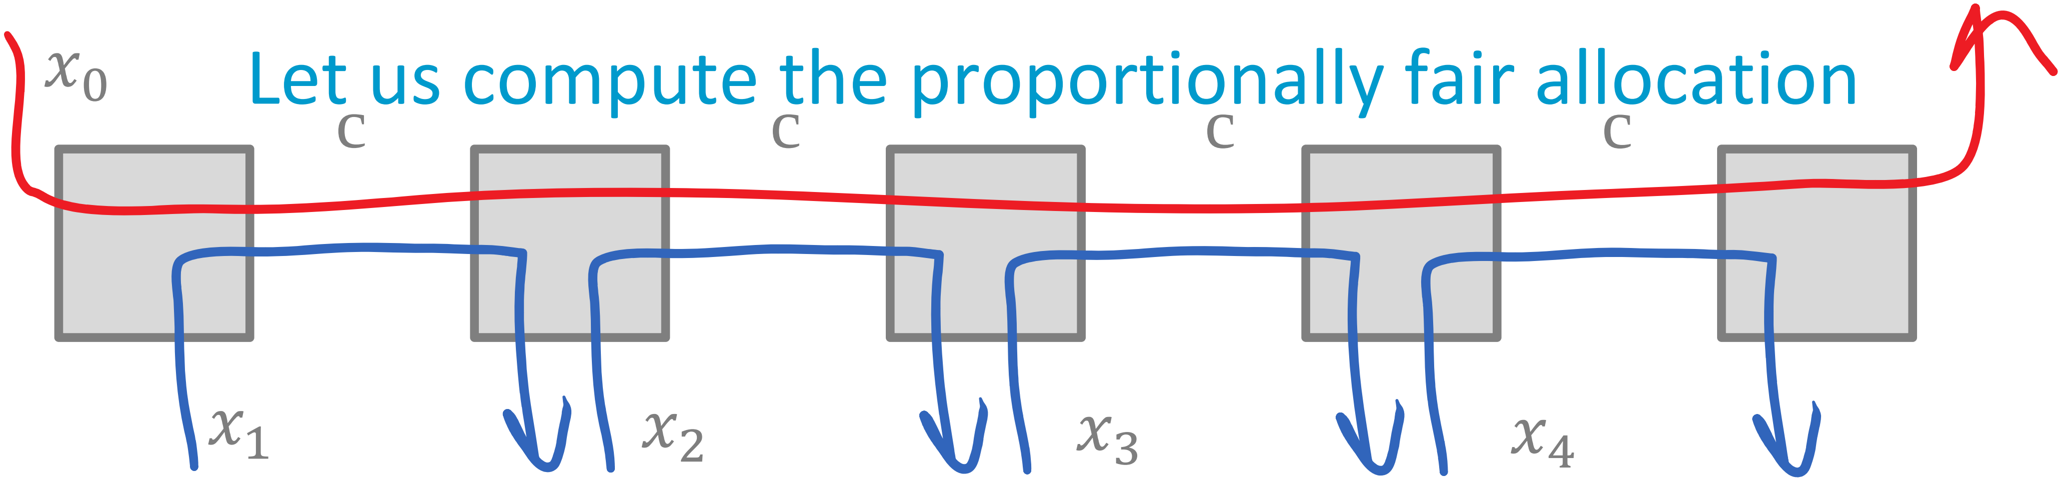

We have to solve the optimization problem:

$$\max U = \log x^* + 4\log(c - x^*)$$
$$\text{subject to } 0 < x^* < c$$

This is a 1d problem, can be solved by computing the derivative

We find $\frac{dU}{dx^*} = \frac{1}{x^*} - \frac{4}{c-x^*}$

There is a maximum for $x^* = \frac{c-x^*}{4}$ i.e. $x^* = \frac{c}{5}$

The proportionally fair allocation is

$$x_0 = \frac{c}{5}, \qquad x_1 = x_2 = x_3 = x_4 = \frac{4c}{5}$$

Which one is the proportionally fair allocation ? (in Mb/s)

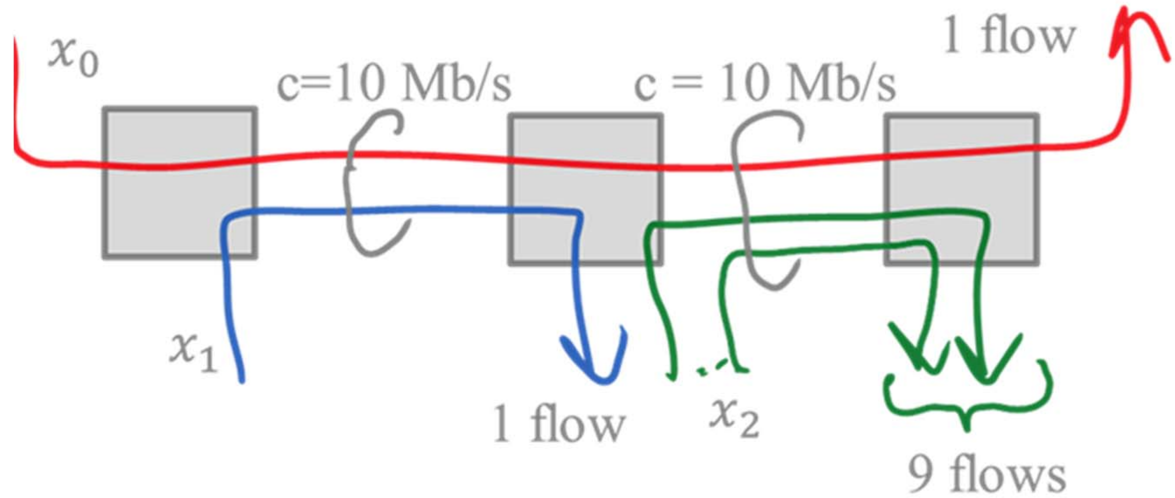

A. $x_0 = 1, x_1 = 9, x_2 = 1$

B. $x_0 = 0.909, x_1 = 9, x_2 = 1.010$

C. $x_0 = 1.009, x_1 = 8.991, x_2 = 0.999$

D. $x_0 = 0.909, x_1 = 9.091, x_2 = 1.010$

E. I don't know

Solution

A. $x_0 = 1, x_1 = 9, x_2 = 1$

B. $x_0 = 0.909, x_1 = 9, x_2 = 1.010$

C. $x_0 = 1.009, x_1 = 8.991, x_2 = 0.999$

D. $x_0 = 0.909, x_1 = 9.091, x_2 = 1.010$

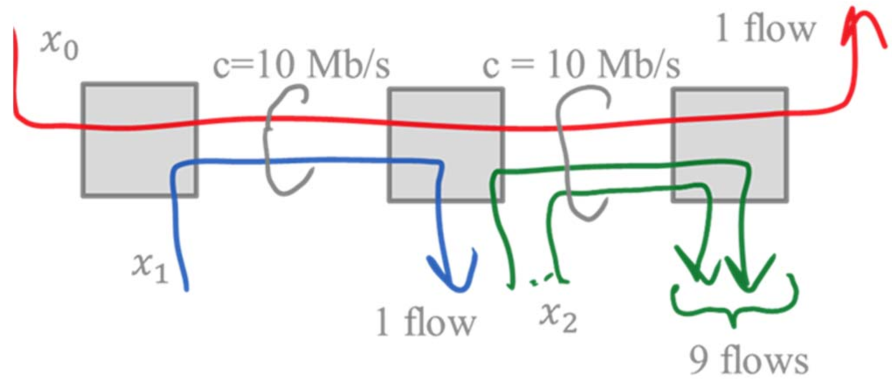

Answer D

We know A is not proportionally fair

B is not Pareto-efficient (you can increase x_1 only) therefore is also not proportionally fair

C goes in the wrong direction (gives more to 0 than to 2) and is probably not proportionally fair

D is probably the correct answer

Solution

A. $x_0 = 1, x_1 = 9, x_2 = 1$

B. $x_0 = 0.909, x_1 = 9, x_2 = 1.010$

C. $x_0 = 1.009, x_1 = 8.991, x_2 = 0.999$

D. $x_0 = 0.909, x_1 = 9.091, x_2 = 1.010$

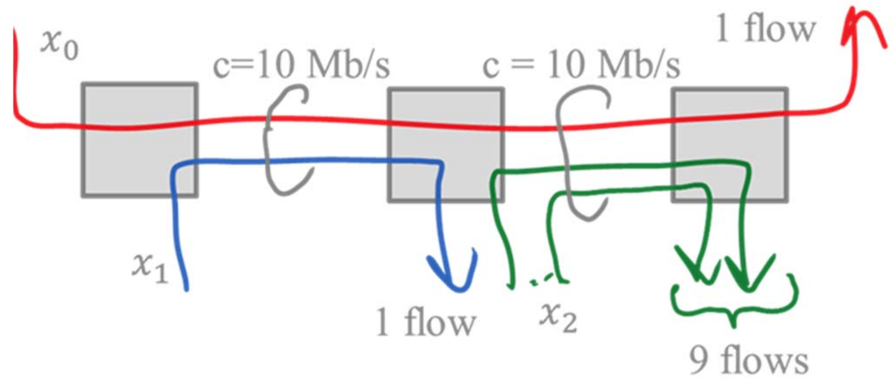

We can compute the proportionally fair allocation with the same method as before; and obtain

$$x_0 = x^*$$
$$x_1 = 10 - x^*$$
$$x_2 = \frac{10 - x^*}{9}$$

with x^* that maximizes $\log x + \log(10 - x) + 9 \log \frac{10-x}{9}$

$x^* = \frac{10}{11}$ Mb/s

This is allocation D

Utility Fairness

One can *interpret* proportional fairness as the allocation that maximizes a global utility $\sum_i U_i(x_i)$ with $U_i(x_i) = \log x_i$.

If we take some other utility function we have what we call a utility fairness

It can be shown that max-min fairness is the limit of utility fairness when the utility function converges to a step function

but max-min fairness cannot be expressed exactly as a utility fairness

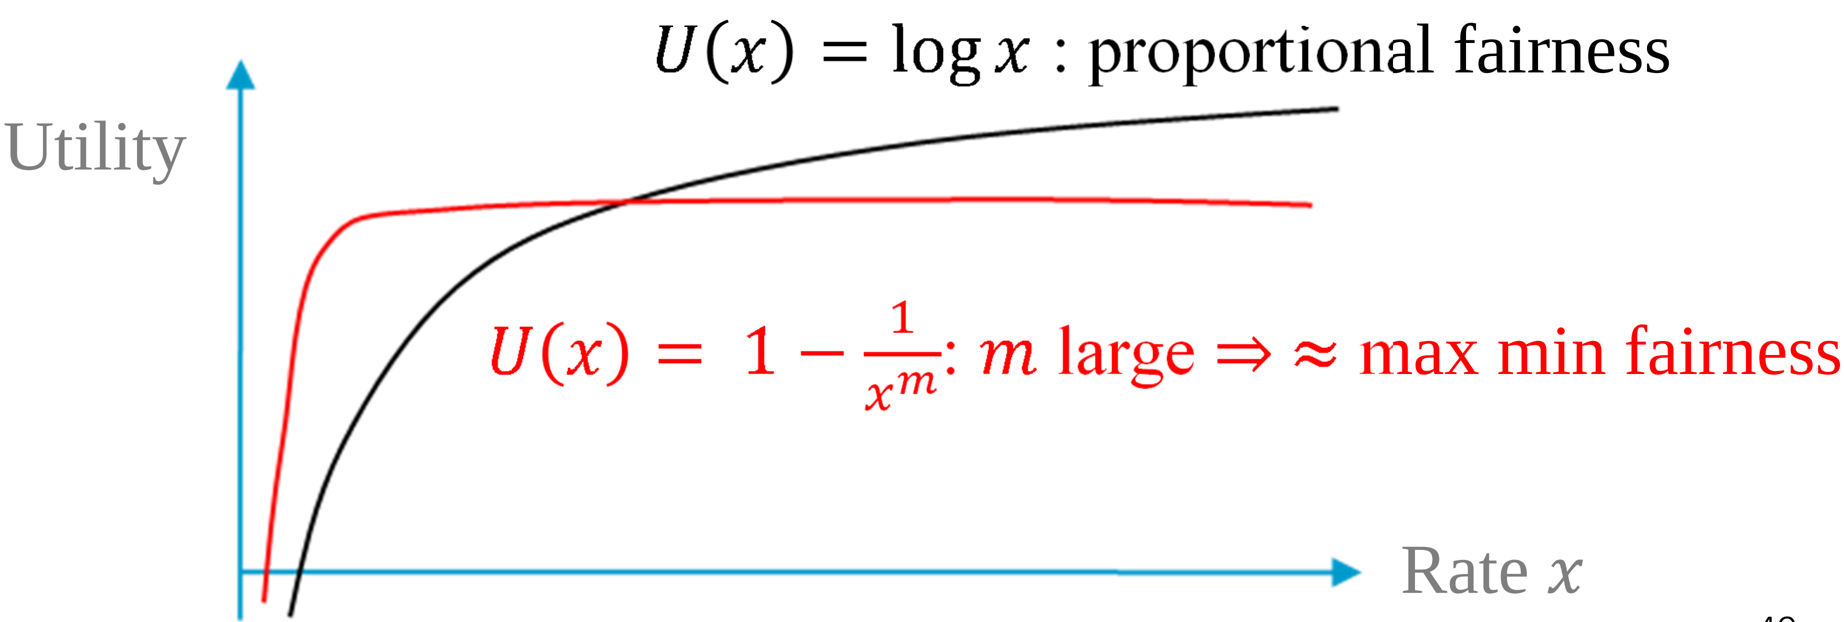

Take Home Message

- Sources should adapt their rate to the state of the network in order to avoid inefficiencies and congestion collapse
  - This is called "congestion control"
- A rate adaptation mechanism should target some form of fairness
  - E.g. max-min fairness or proportional fairness

4. Additive Increase Multiplicative Decrease

How can congestion control be implemented ?

Explicit (rate based): tell every host how fast it can send
ATM ABR
MPLS networks (smart grid)
Cellular networks

Hop by hop = backpressure: STOP/GO signals sent upstream
Gigabit LAN switches

Fair Queuing per Flow: One queue per flow / per user, served round robin -> Cellular networks

End-to-end: hosts taste the water and increase or decrease their sending rate using a host congestion control algorithm
The solution in the Internet

Additive Increase Multiplicative Decrease (AIMD)

The first congestion control algorithm deployed in the Internet and before that, in Decnet (the "Decbit").

Still widely deployed today

We have designed a scheme that allows a network to operate at its knee. As shown in Figure 3, the scheme uses one bit called the **congestion avoidance bit** in the network layer header of the packet for feedback from the subnet to the users. A source clears the congestion avoidance bit as the packet enters the subnet. All routers in the subnet monitor their load and if they detect that they are operating above the knee, they set the congestion avoidance bit in the packets belonging to users causing overload. Routers operating below the knee pass the bit as received. When the packet is received at the destination the network layer passes the bit to the destination transport, which takes action based on the bits.

Raj Jain

Raj Jain, K.K. Ramakrishnan, and Dah-Ming Chiu. Congestion avoidance in computer networks with a connectionless network layer. Technical Report DEC-TR-506, Digital Equipment Corporation, August 1987.

A Simple Network Model

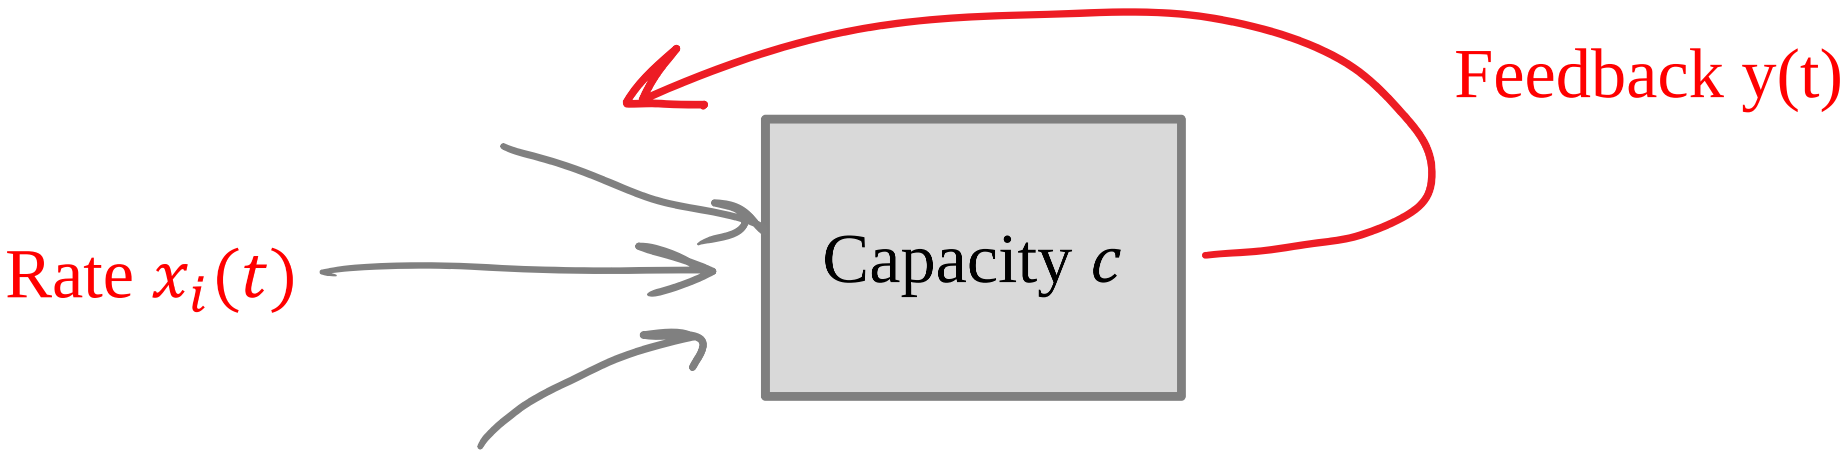

Network sends a one bit feedback :

$y(t) = 0$ if $\sum_i x_i(t) \leq c$, $y(t) = 1$ if $\sum_i x_i(t) > c$

Sources reduce rate $x_i(t+1)$ if $y(t) = 1$, increase otherwise

Question: what form of increase/decrease laws should one pick?

Linear Laws

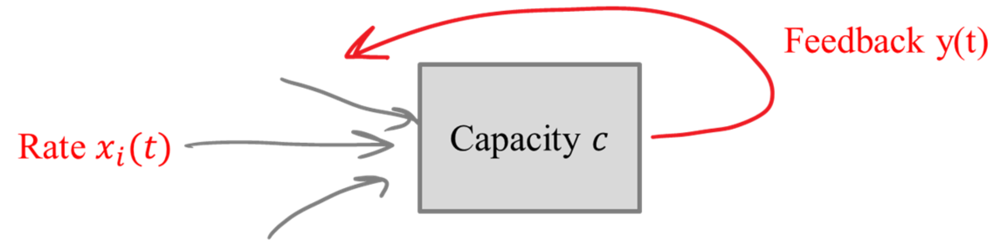

We consider linear laws

if $y(t) = 1$ then $x_i(t+1) = u_1\, x_i(t) + v_1$
if $y(t) = 0$ then $x_i(t+1) = u_0\, xi(t) + v_0$

We want to decrease when $y(t) = 1$, so

$u_1 \leq 1$ and $v_1 \leq 0$ and at least one inequality must be strict

Multiplicative decrease factor

Additive decrease term

We want to increase when $y(t) = 0$, so

$u_0 \geq 1$ and $v_0 \geq 0$ and at least one inequality must be strict

Multiplicative increase factor

Additive increase term

Example

$u_1 = 0.5, v_1 = 0$ (multiplicative decrease)
$u_0 = 1, v_0 = 1$ (Mb/s) (additive increase)

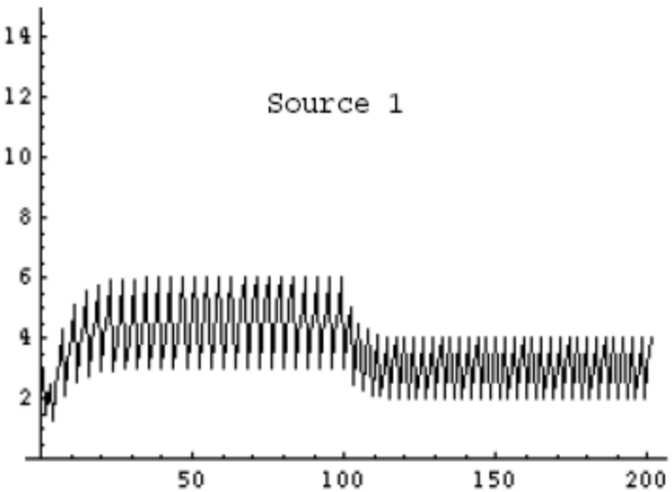

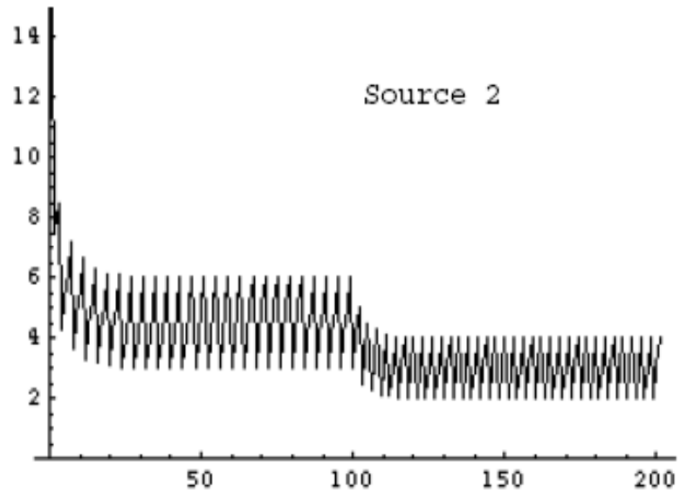

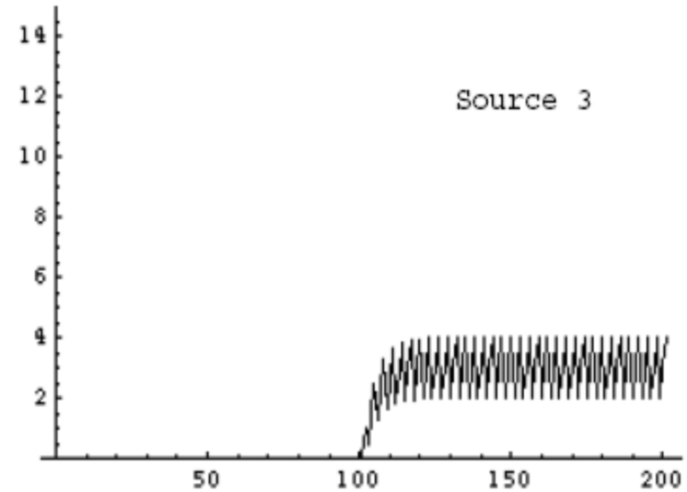

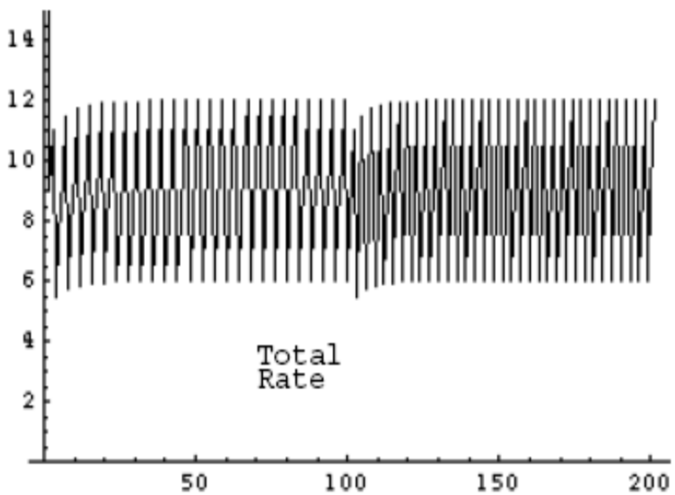

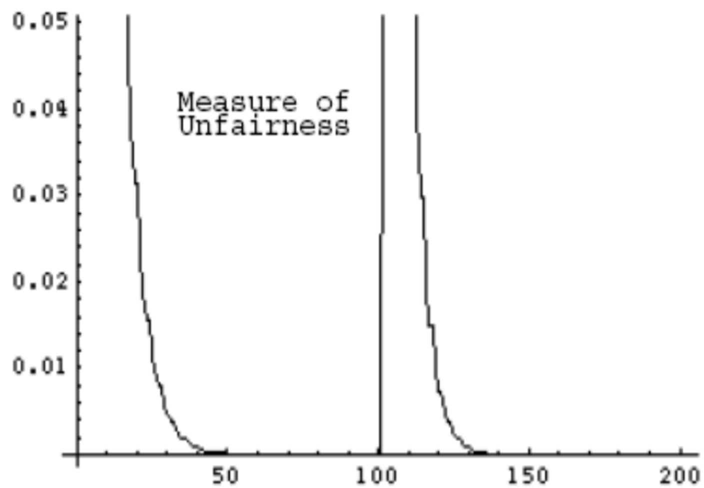

Analysis of Linear Control Schemes

We want to achieve efficiency and fairness

We could target either max-min fair or proportionally fair allocations

Here they are the same

We will now analyze the impact of each of the four coefficients u_0, u_1, v_0 and v_1.

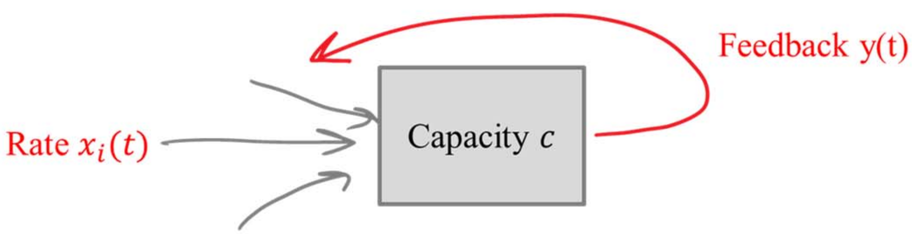

- We consider linear laws

  if $y(t) = 1$ then $x_i(t+1) = u_1\, x_i(t) + v_1$

  if $y(t) = 0$ then $x_i(t+1) = u_0\, xi(t) + v_0$

- We want to decrease when $y(t) = 1$, so

  $u_1 \leq 1$ and $v_1 \leq 0$ and at least one inequality must be strict

  Multiplicative decrease factor (u_1); Additive decrease term (v_1)

- We want to increase when $y(t) = 0$, so

  $u_0 \geq 1$ and $v_0 \geq 0$ and at least one inequality must be strict

  Multiplicative increase factor (u_0); Additive increase term (v_0)

Zoom on 2 Sources

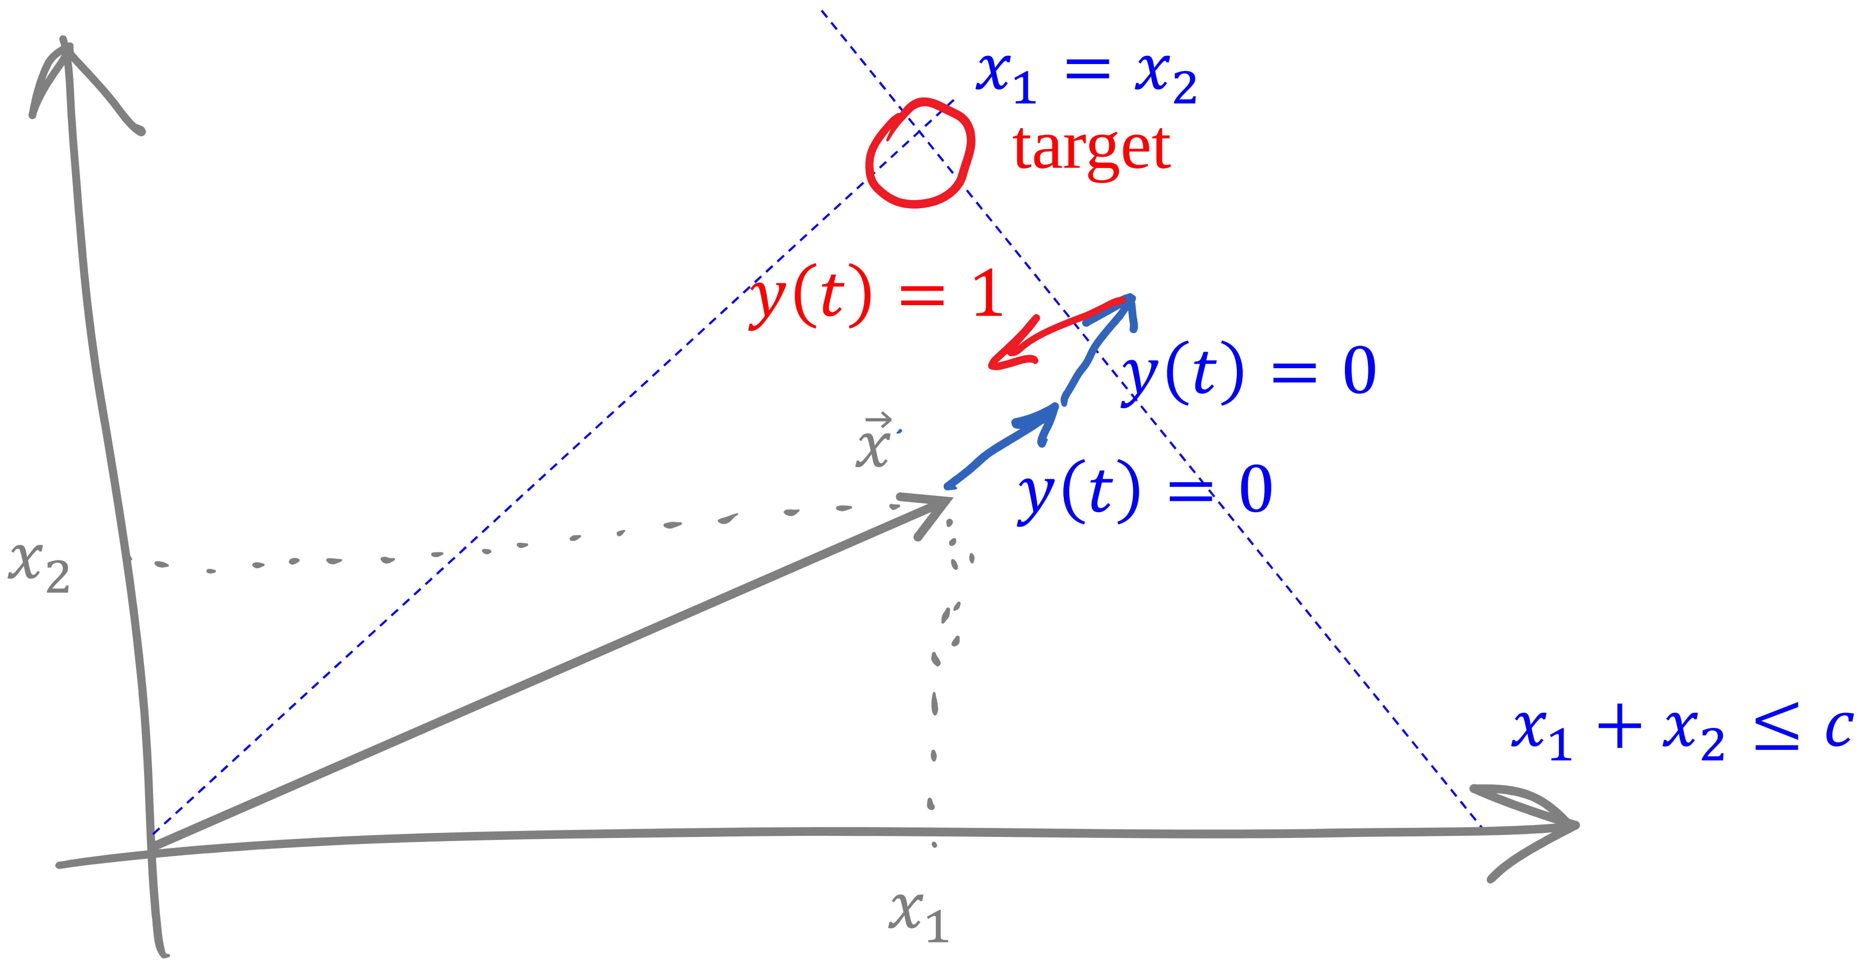

Zoom on 2 sources ; say what is true

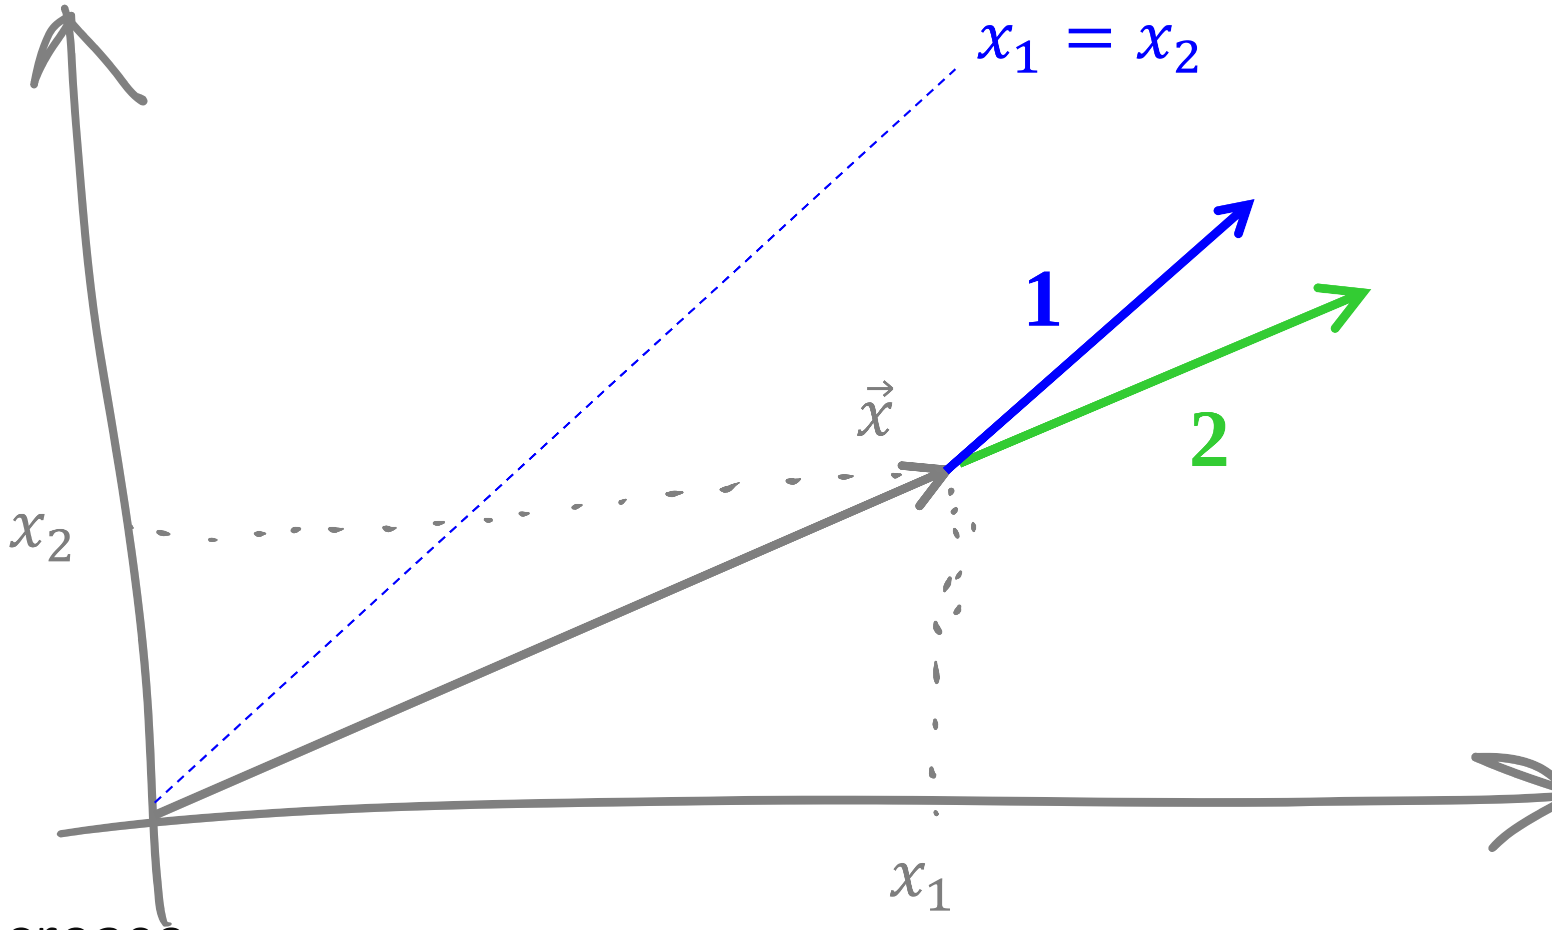

A. 1 = additive increase,
   2 = multiplicative increase

B. 1 = multiplicative increase
   2 = additive increase,

C. None of the above

D. I don’t know

Solution

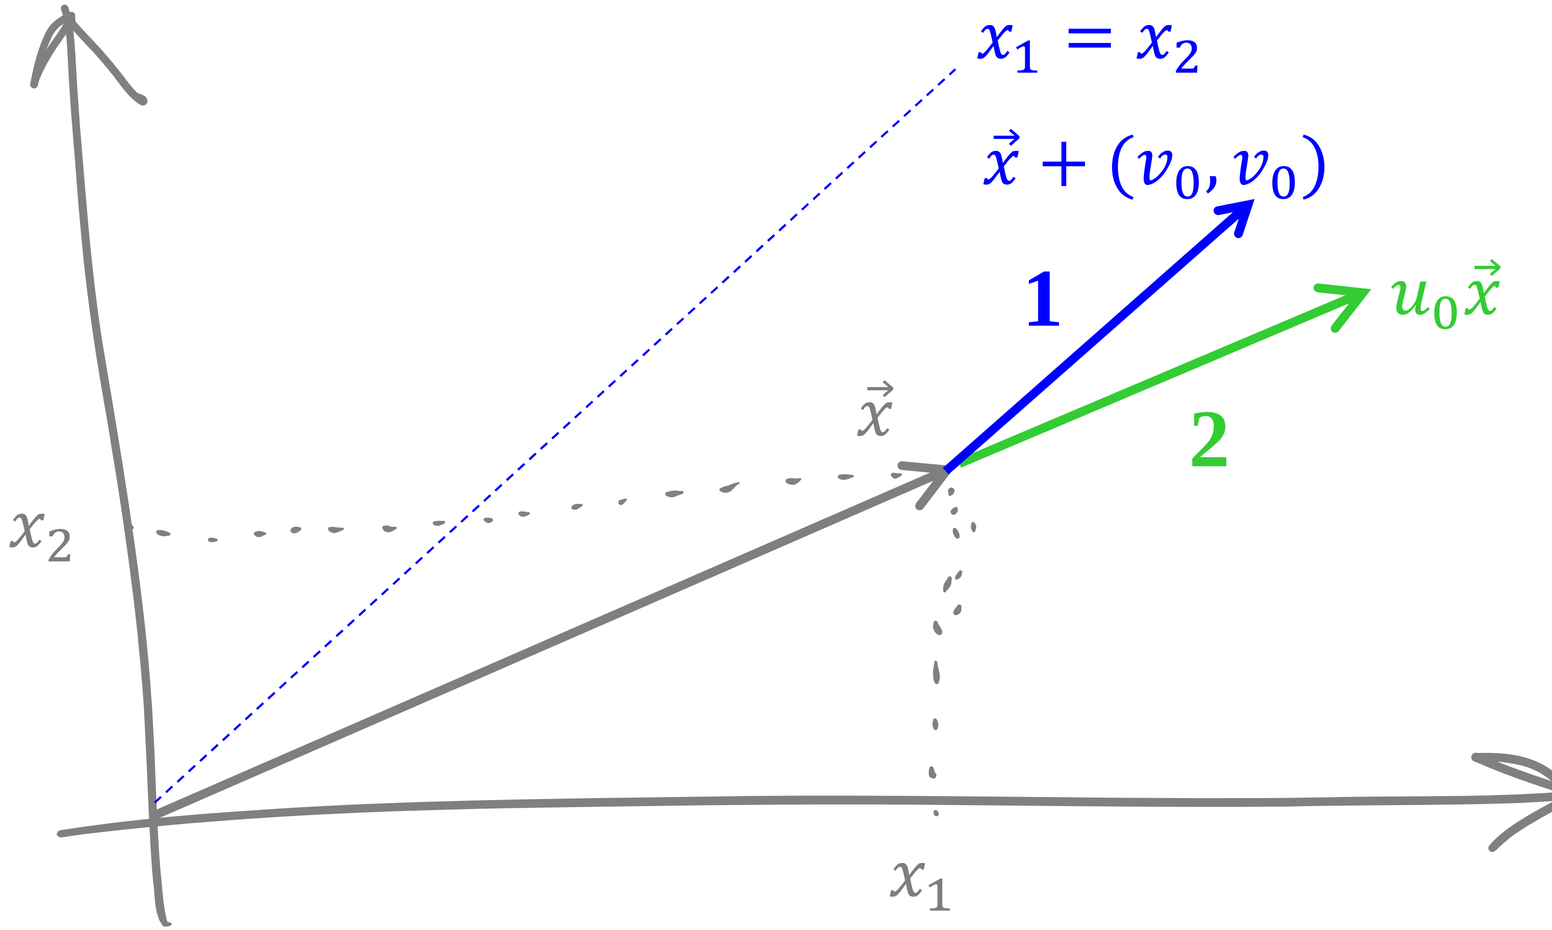

1 = additive increase,
2 = multiplicative increase

Answer A

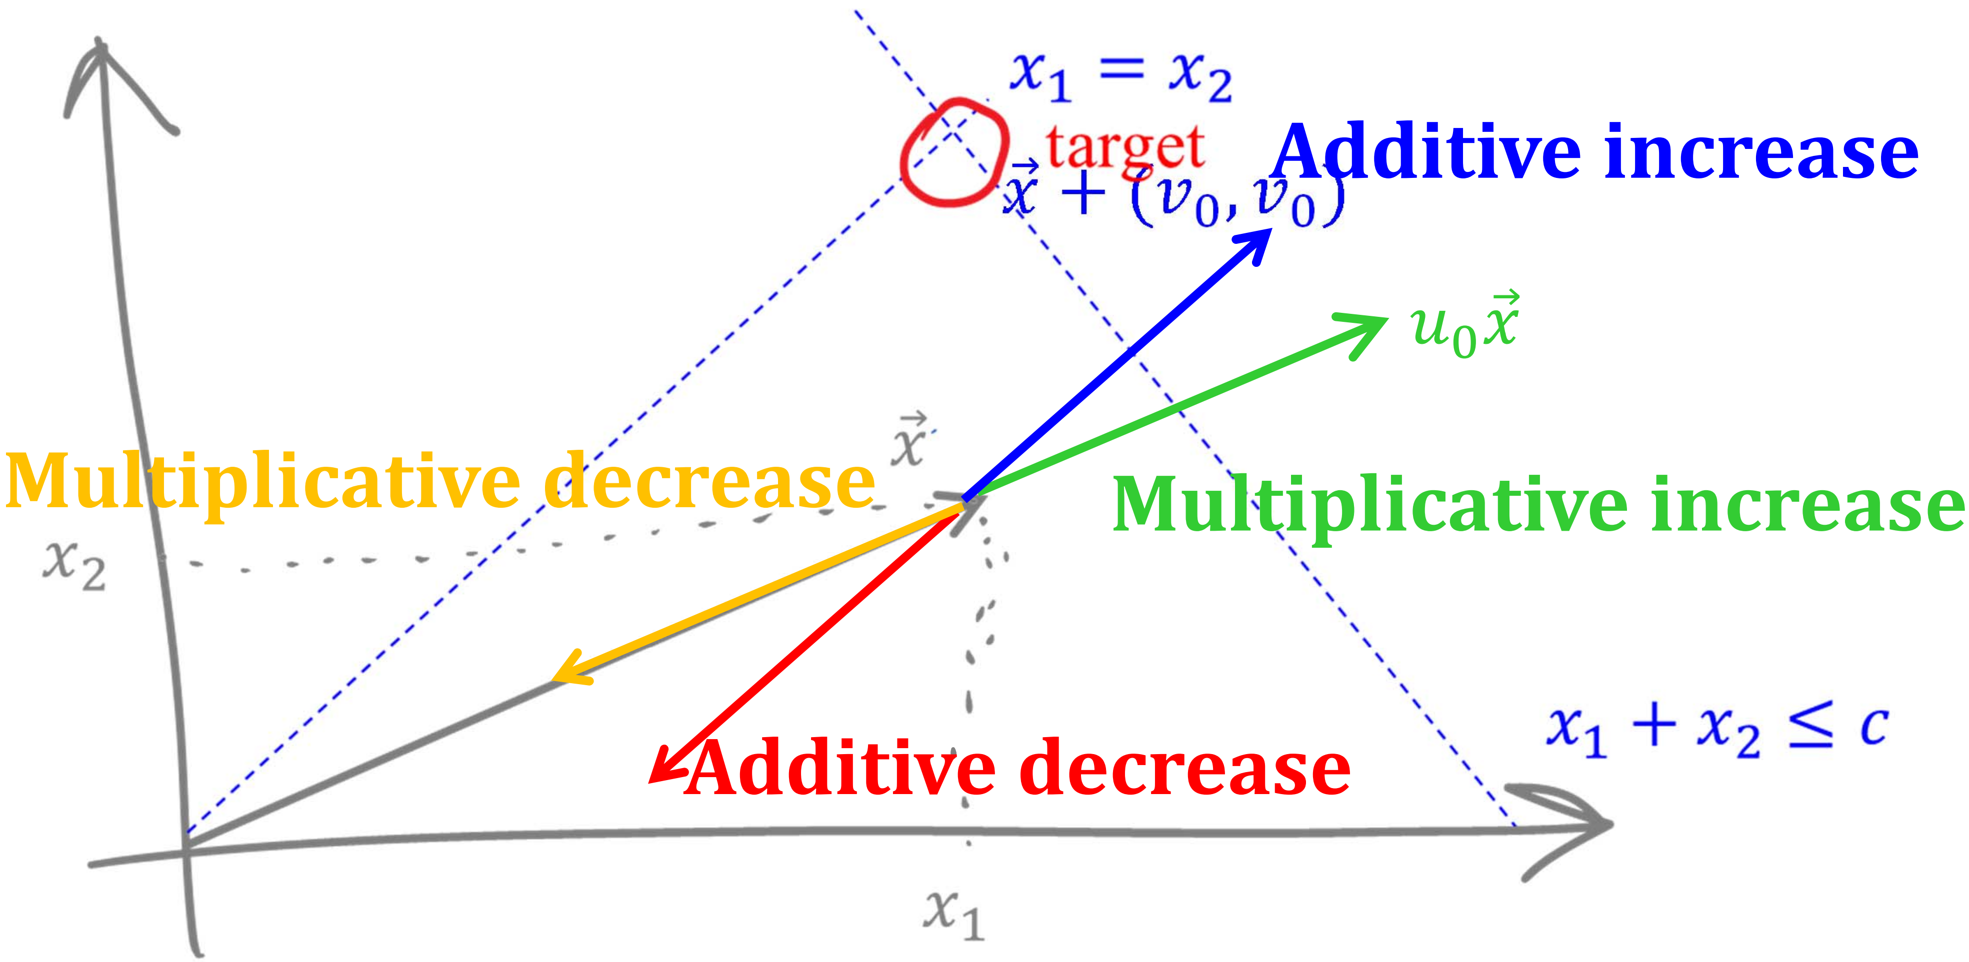

1. *Additive decrease* worsens fairness (goes away from $x_1 = x_2$) and should be avoided ⇒ decrease should be mutiplicative
2. *Additive increase* is the only move that increases fairness and should be therefore be included ⇒ increase should be additive

Why AIMD

Among the linear controls, only additive increase – multiplicative decrease (AIMD) tends to bring the allocation towards fairness and efficiency.

This is what was implemented in the Internet after the first congestion collapses.

In a more complex network setting, does AIMD distribute rates according to max-min fairness or proportional fairness ?

A. Max-min
B. Proportional
C. None of the above
D. I don't know

Fairness of AIMD

Answer C

AIMD with: additive increase $+r$, multiplicative decrease $\times (1 - \eta)$ and one update per time unit implements utility fairness, with utility of flow i given by

$$U(x_i) = \log \frac{x_i}{r + \eta x_i}$$

with x_i = rate. The fairness of AIMD is between maxmin fairness and proportional fairness, closer to proportional fairness.

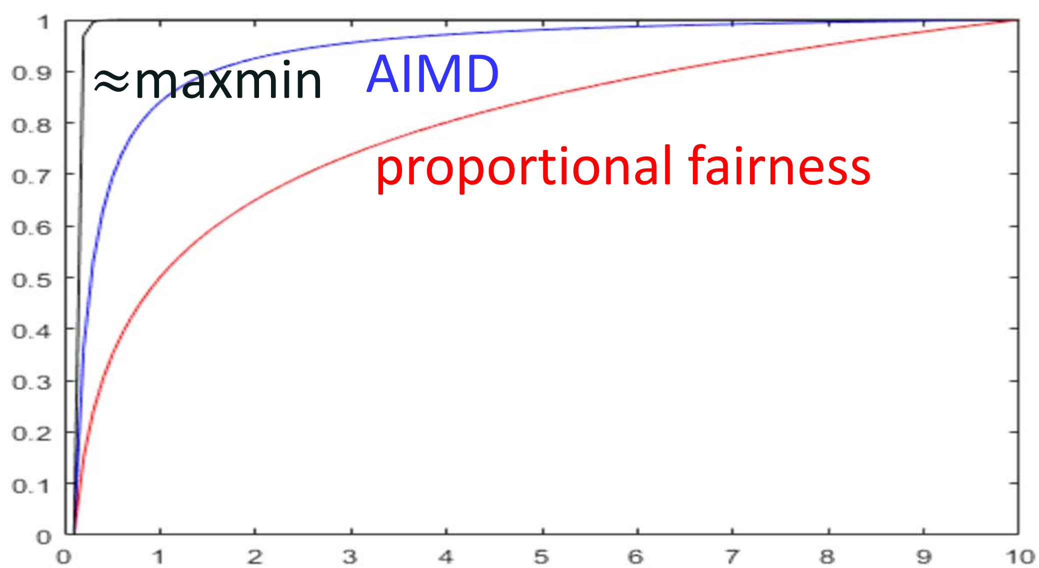

rescaled utility functions;
AIMD is for $\eta = 0.5$ $r = 1$MSS per 100 ms
maxmin approx. is $U(x) = 1 - x^{-5}$

5. Slow Start

AIMD convergence can be accelerated when initial conditions are very different

Slow start is an additional method, added to AIMD

Used at beginning of connection and at losses detected by timeout

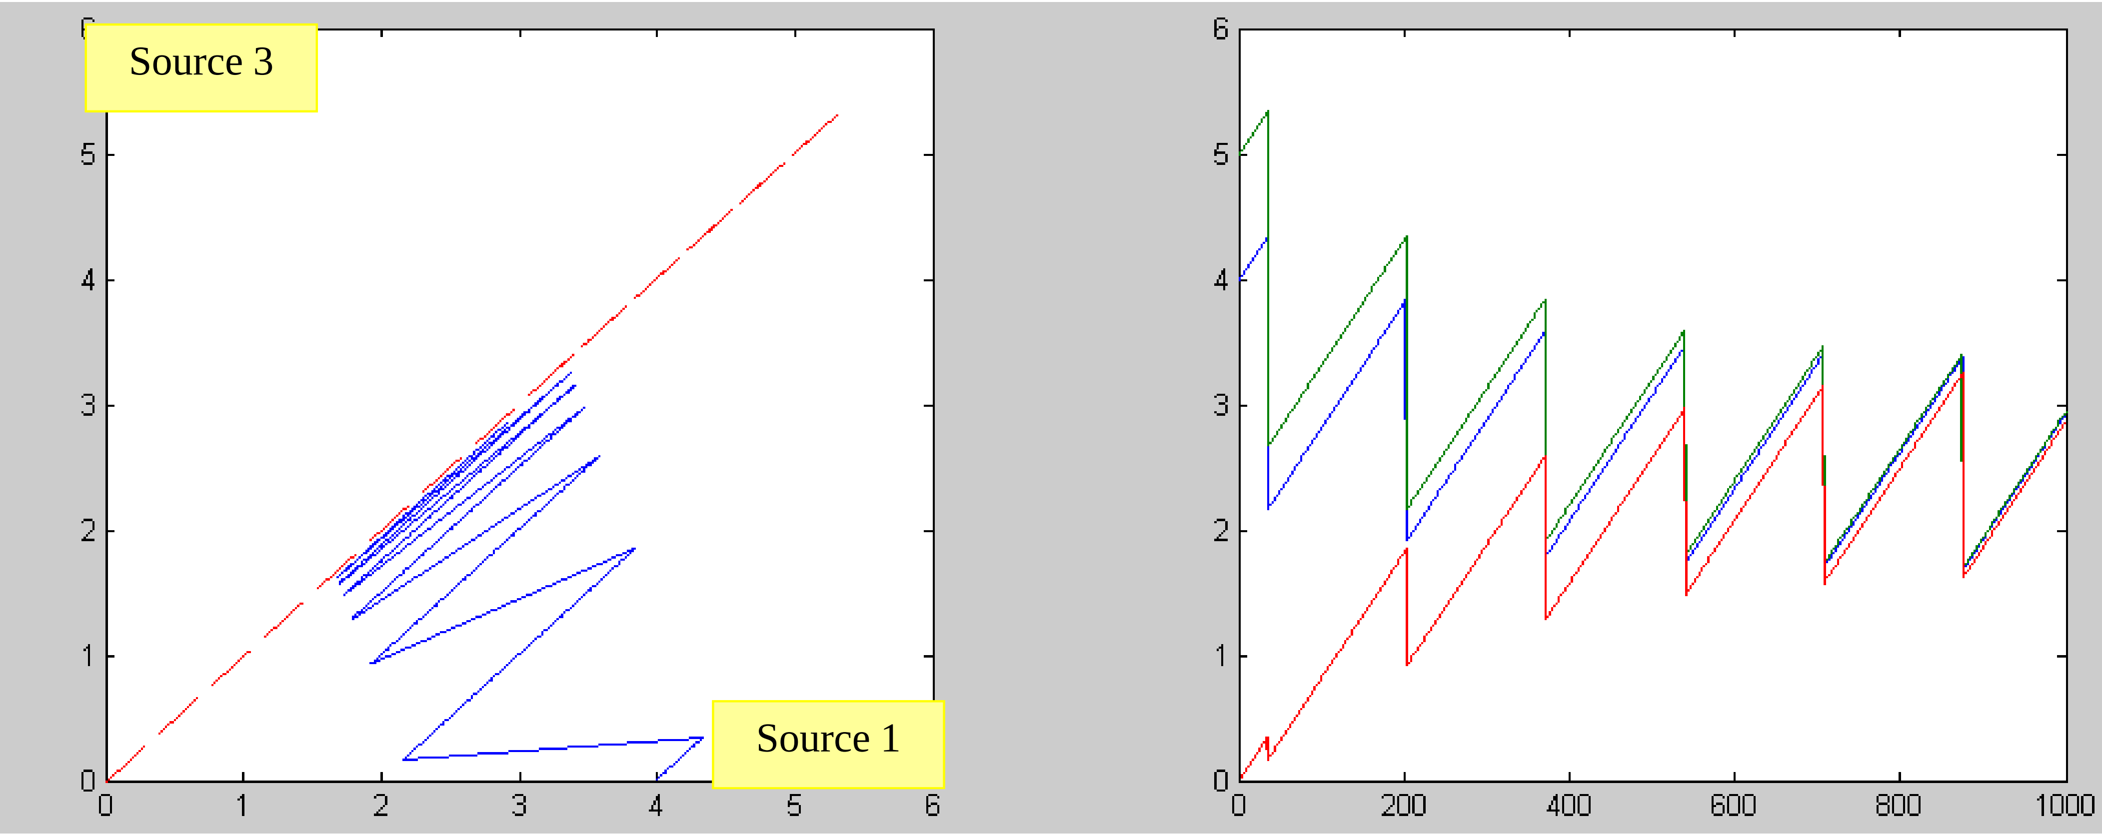

Slow Start

Increase the rate multiplicatively

(by w_0, e.g. $w_0 = 2$) until a target rate is reached or negative feedback is received

Apply multiplicative decrease (by u_1, e.g. $u_1 = 0.5$) to target rate if negative feedback is received

Exit slow start when target rate is reached

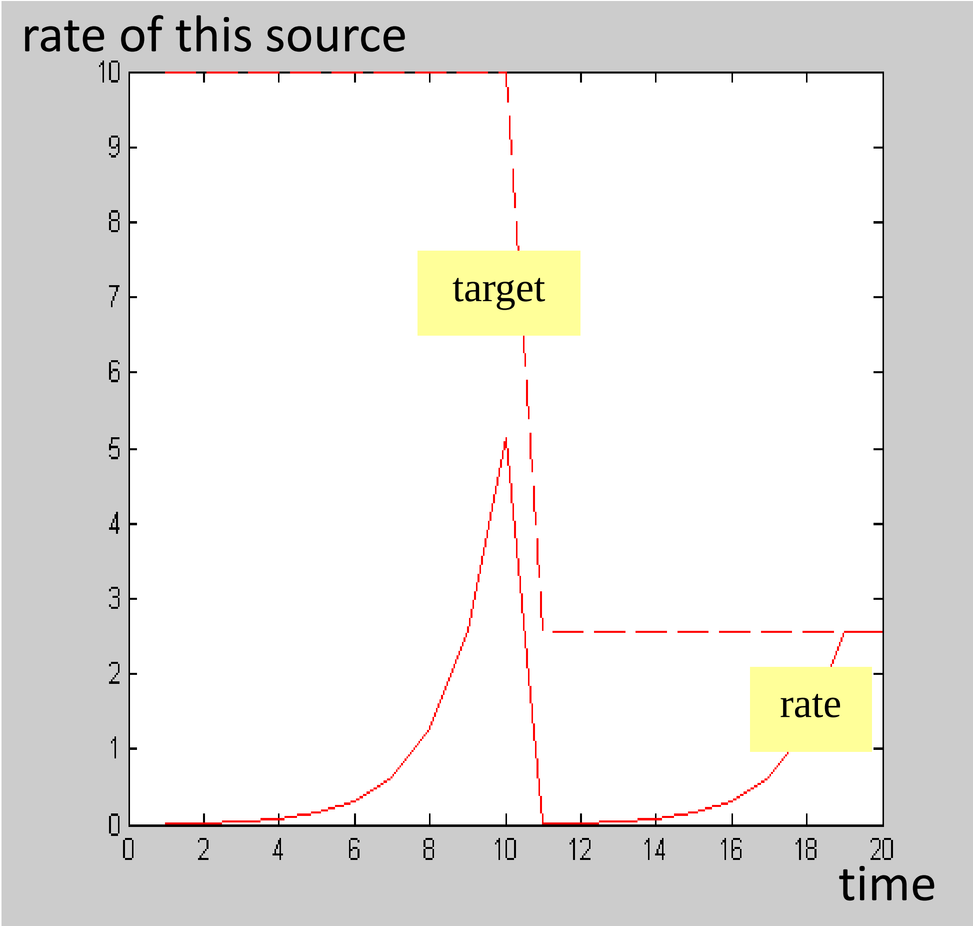

Algorithm 2 Slow Start with the following parameters: AIMD constants multiplicative increase factor $w_0 > 1$; maximum rate $r_{\max} > 0$.

rate $\leftarrow v_0$
targetRate $\leftarrow r_{\max}$
do forever
receive feedback
if feedback is positive **then**
    rate $\leftarrow w_0 \cdot$ rate
    if rate $\geq$ targetRate **then**
        rate $\leftarrow$ targetRate
        exit do loop
    end if
else
    targetRate $\leftarrow \max(u_1 \cdot \text{rate}, v_0)$
    rate $\leftarrow v_0$
end if
end do

$w_0 = 2$

$u_1 = \frac{1}{2}$

With Slow Start

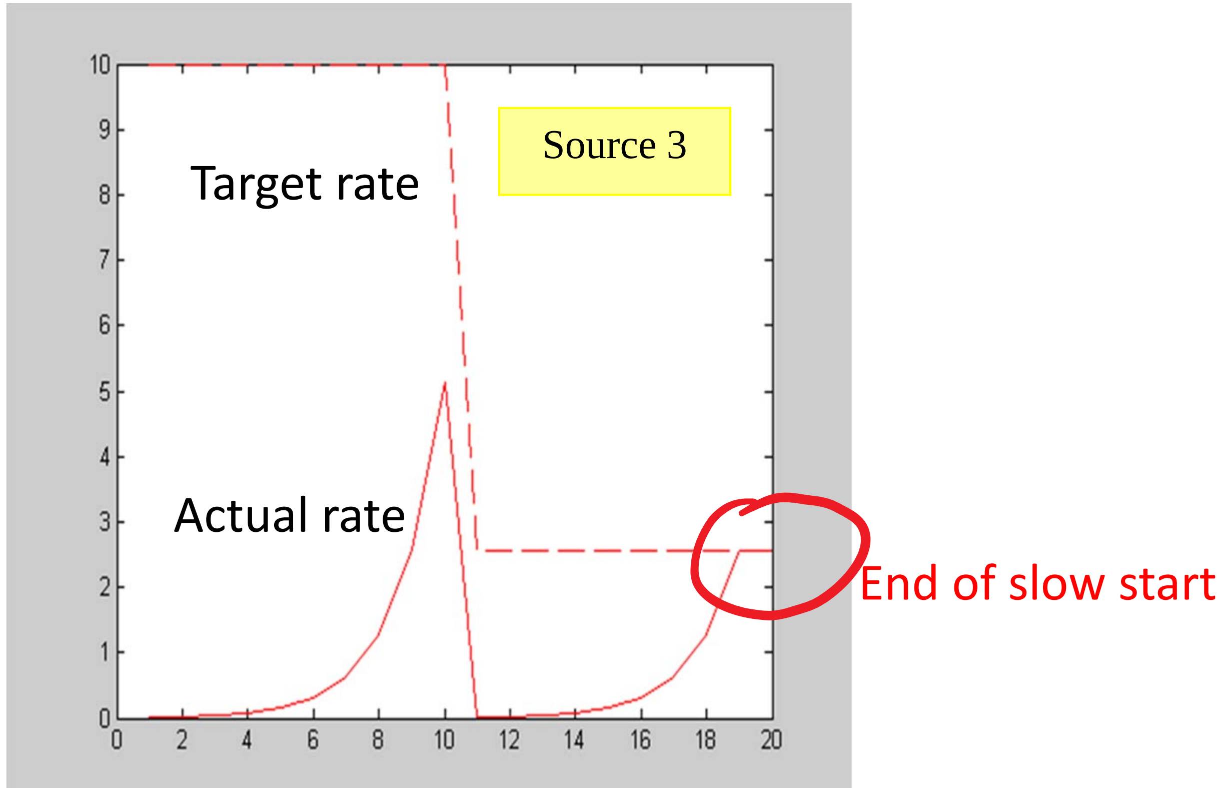

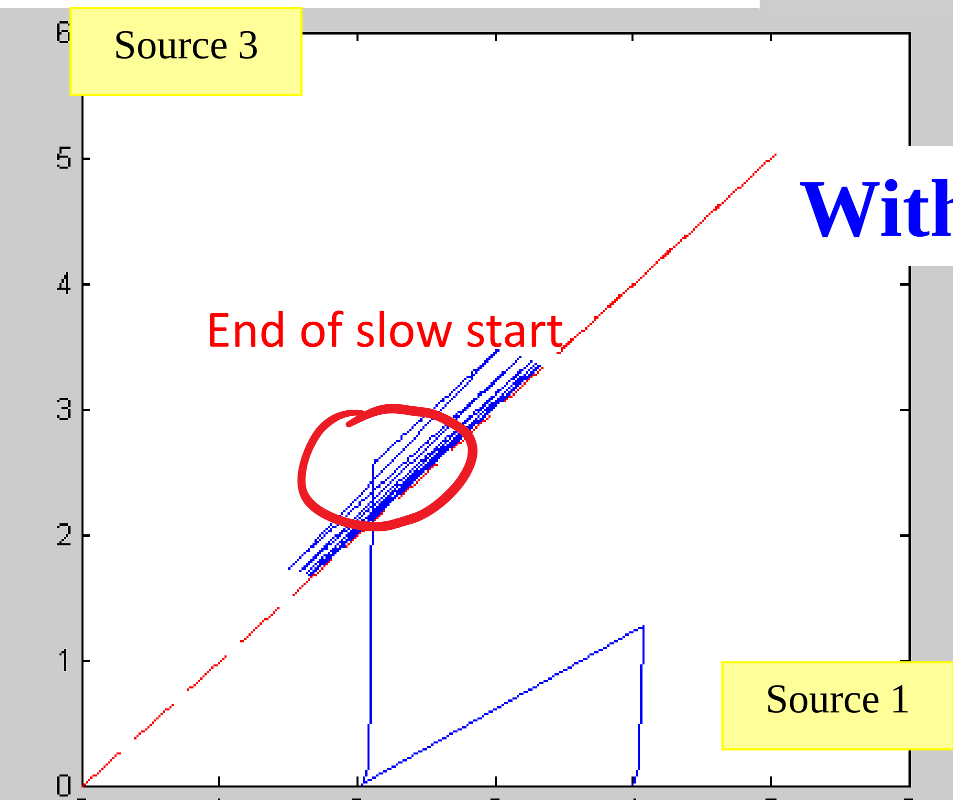

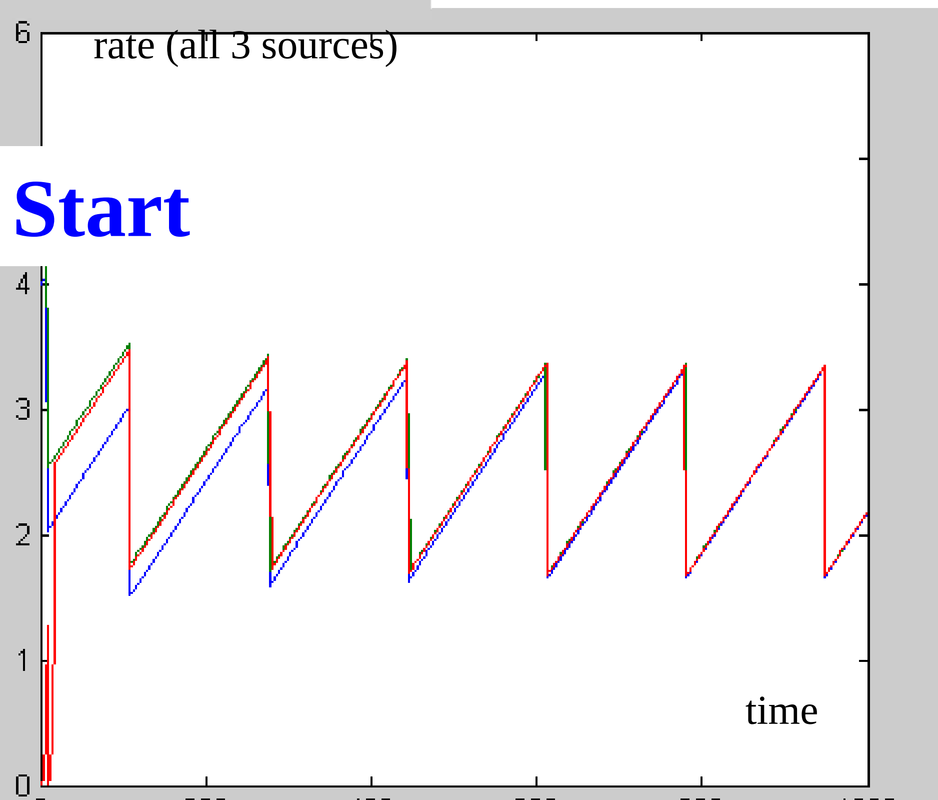

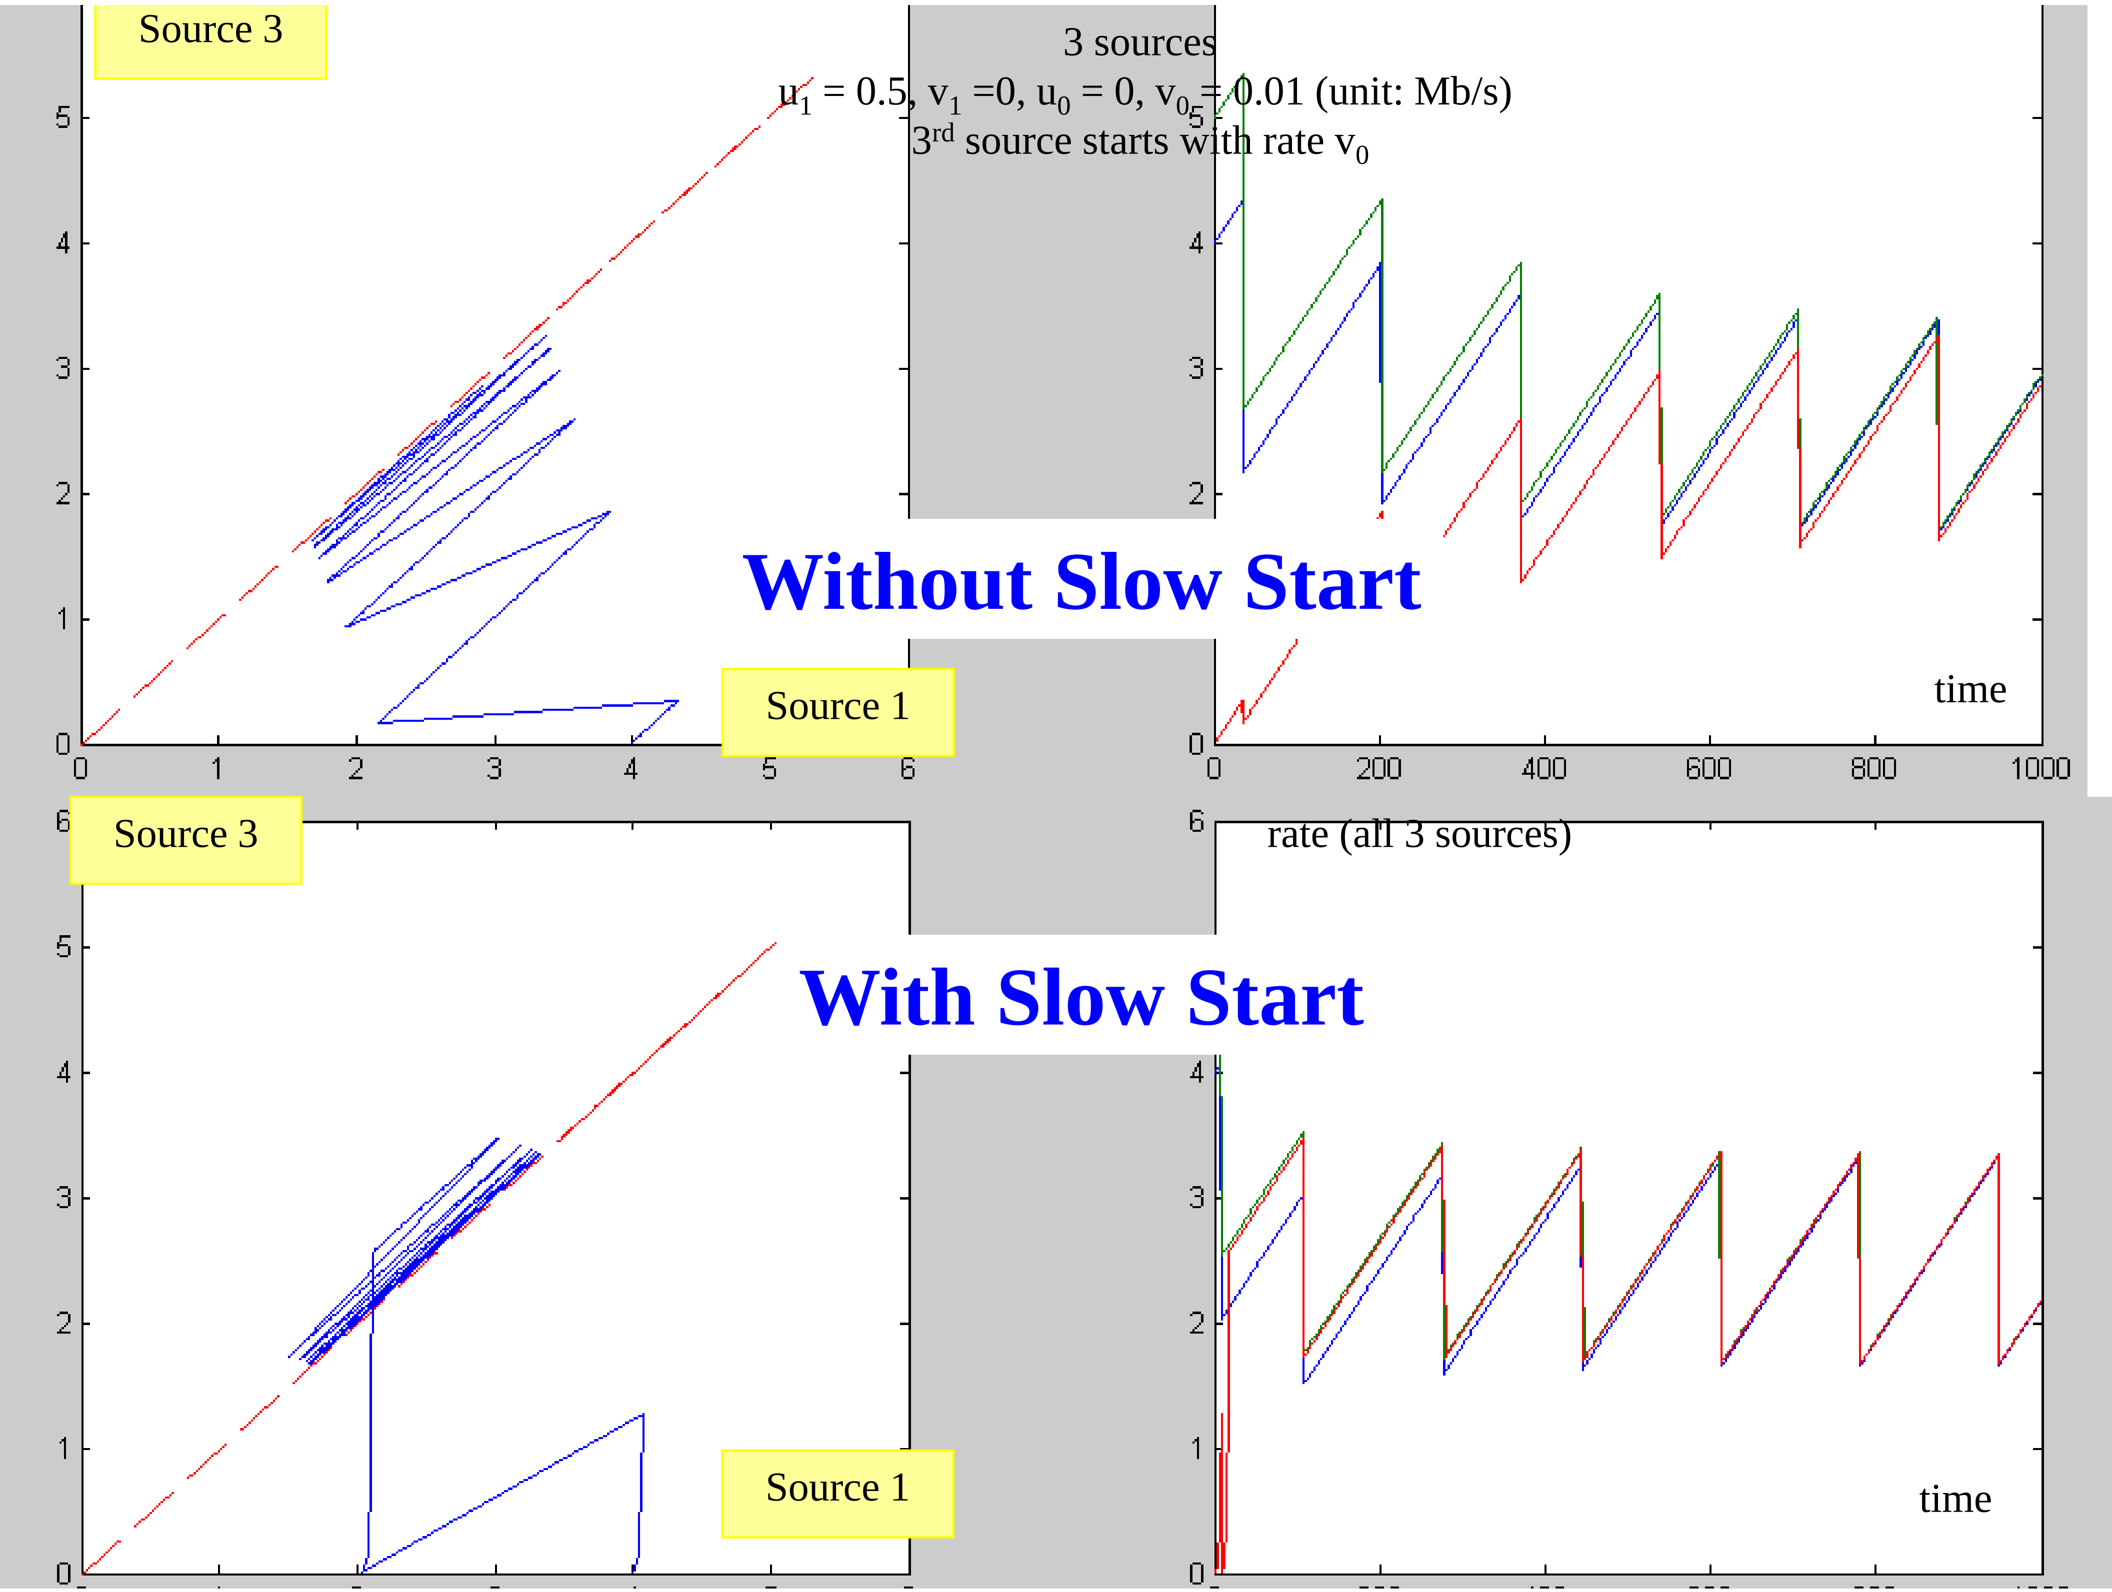

3 sources
u1 = 0.5, v1 =0, u0 = 0, v0 = 0.01 (unit: Mb/s)
3rd source starts with rate v0
Without Slow Start
With Slow Start
Source 3
Source 1
Source 3
Source 1
time
time
rate (all 3 sources)

Summary

Congestion control is necessary to avoid inefficiencies and collapses

A congestion control scheme aims at allocating rates according to some form of fairness

In the internet, we use end-to-end congestion control with

- AIMD
- Slow Start
- and other refinements – see part 2.