ICA

EPFL ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Congestion Control for Best Effort

Prof. Jean-Yves Le Boudec
Prof. Andrzej Duda
Prof. Patrick Thiran
ICA, EPFL

CH-1015 Ecublens
Andrzej.Duda@imag.fr
http://icawww.epfl.ch

Contents

- Congestion control for best effort
  - Objectives
    - effciency
    - fairness
  - Max-min fairness
  - Proportional fairness
  - Additive increase, multiplicative decrease
  - Different forms of congestion control

3

Congestion Control - example

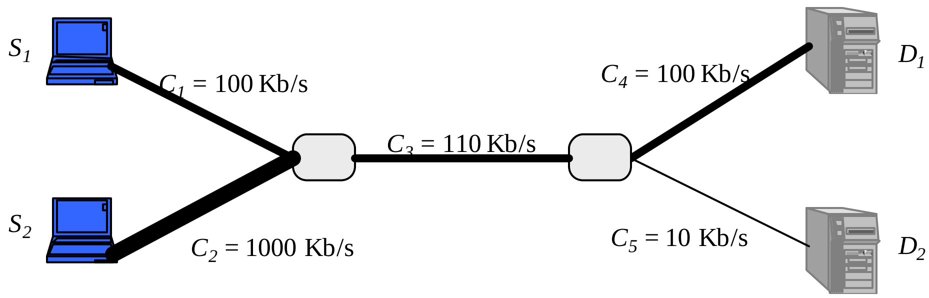

- Allocation of link bit rates
  - if the offered traffic exceeds the capacity of a link, all sources see their traffic reduced in proportion of their offered traffic
  - approximately true if FIFO in routers

4

Bit rate allocation

- Rate x_{ls}: source s on link l
- Traffic λ_s : generated by source s
- Allocation

$x_{11} = \min(\lambda_1, C_1)$

$x_{22} = \min(\lambda_2, C_2)$

$x_{3i} = \min(x_{ii}, C_3\, x_{ii}/(x_{11} + x_{22}))$

$x_{41} = \min(x_{31}, C_4)$

$x_{52} = \min(x_{32}, C_5)$

throughput $\theta = x_{41} + x_{52}$

Our example :

x_{11} = ?

x_{22} = ?

x_{31} = ?

x_{32} = ?

x_{41} = ?

x_{52} = ?

throughput θ = ? Kb/s

5

Bit rate allocation

- Rate x_{ls}: source s on link l
- Traffic λ_s : generated by source s
- Allocation

$x_{11} = \min(\lambda_1, C_1)$
$x_{22} = \min(\lambda_2, C_2)$
$x_{3i} = \min(x_{ii}, C_3\, x_{ii}/(x_{11} + x_{22}))$
$x_{41} = \min(x_{31}, C_4)$
$x_{52} = \min(x_{32}, C_5)$
throughput $\theta = x_{41} + x_{52}$

Our example :

$x_{11} = 100$
$x_{22} = 1000$
$x_{31} = 110\times100/1100=10$
$x_{32} = 110\times1000/1100=100$
$x_{41} = 10$
$x_{52} = 10$
throughput $\theta = 20$ Kb/s

6

Congestion Control - example

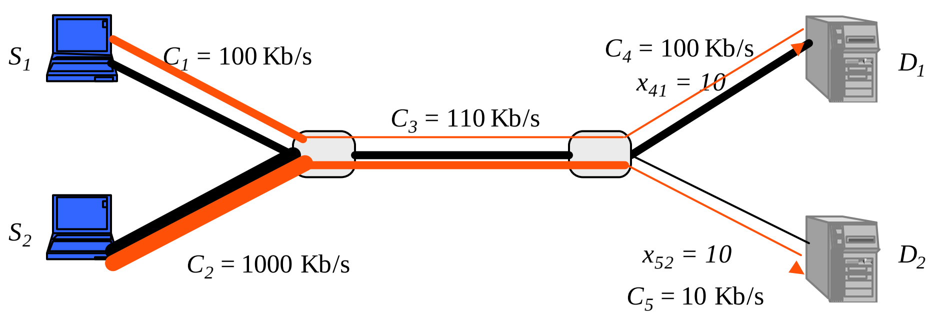

- S_1 sends 10 Kb/s because it is competing with S_2 on link 3
- S_2 is limited on link 5 anyway

Congestion Control - exemple

- How to increase throughput?
  - if S_2 is aware of the global situation and if it would cooperate
  - S_2 reduces x_{22} to 10 Kb/s, because anyway, it cannot send more then 10 Kb/s on link 5
  - x_{31} = 100 Kb/s and x_{41} = 100 Kb/s without any penalty for S_2
  - q = 110 Kb/s

Congestion Control - exemple

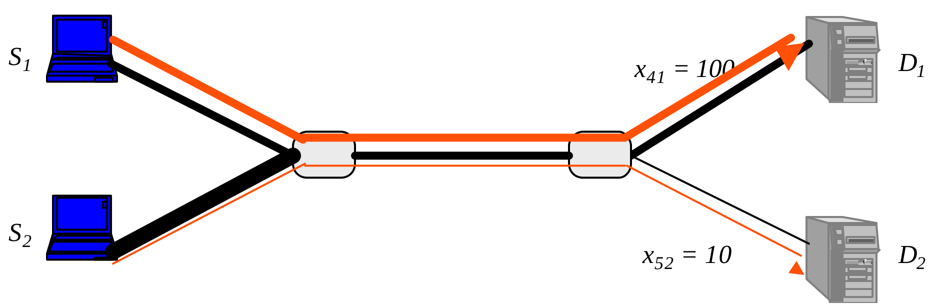

- Optimal use of resources

Throughput vs. offered load

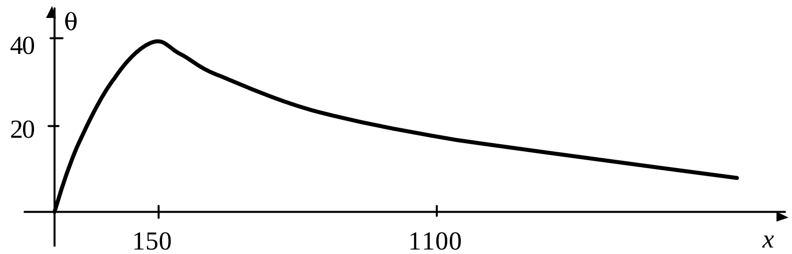

- same example ; sources increase their rate in parallel, but at different speeds
- $l_1 = l$, $l_2 = l^2/10$, l - a parameter
- $l_1(1) = 1$, $l_2(1) = 1/10$
- $l_1(10) = 10$, $l_2(10) = 10$
- $l_1(100) = 100$, $l_2(100) = 1000$
- offered load $x = l_1 + l_2$

Efficiency criterion

- In a packet network, sources should limit their sending rate by taking into consideration the state of the network. Ignoring this may put the network into congestion collapse.
- One objective of congestion control is to avoid such inefficiencies.

11

Efficiency versus Fairness

- Parking lot scenario
  - link capacity : C
  - n_i sources, rate x_i, $i = 1, ..., I$
  - traffic on link i : $n_0 x_0 + n_i x_i$

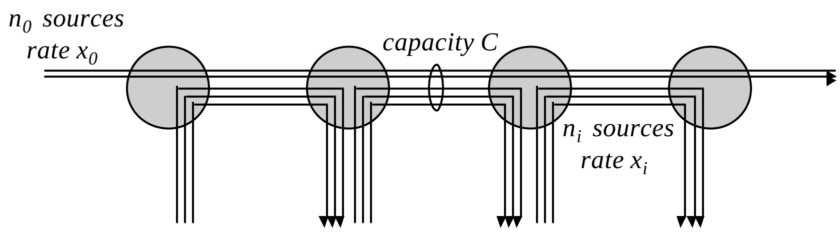

12

Maximal throughput

- For given n_0 and x_0, maximizing the throughput requires that
  - $n_i x_i = C - n_0 x_0$
- Total throughput, measured at the network output
  - $q = n_0 x_0 + \Sigma n_i x_i = n_0 x_0 + \Sigma(C - n_0 x_0) =$
    $= n_0 x_0 + I(C - n_0 x_0) = IC - (I - 1) n_0 x_0$

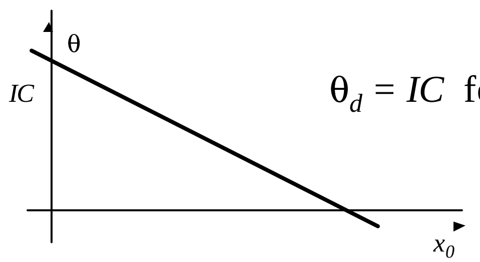

$\theta_d = IC$ for $x_0 = 0$!

Fairness

- Maximizing network throughput as a primary objective may lead to large unfairness
  - some sources may get a zero throughput
- Fairness criterion
  - allocate the same share to all sources, eg. for $n_i = 1$
    - $x_i = C/2$
    - $\theta_{fair} = (I+1)C/2$

Optimizing fairness

- Consider the parking lot scenario for general values of n_i
  - fair allocation on link i
    - $x_i = C/(n_0 + n_i)$, $i = 1, ..., I$
  - decrease x_0 to increase q
    - x_0 = min $C/(n_0 + n_i)$,
  - example
    - I = 2, $n_0 = n_1 = 1$, $n_2 = 9$
    - link 2 : $x_2 = C/(1 + 9) = 0.1\ C$
    - link 1 : $x_1 = C/(1 + 1) = 0.5\ C$
    - x_0 = min $(0.5\ C, 0.1\ C) = 0.1\ C$

15

Example

- Problem
  - link 1 : 0.6 C
    - underutilized
  - link 2 : 1 C

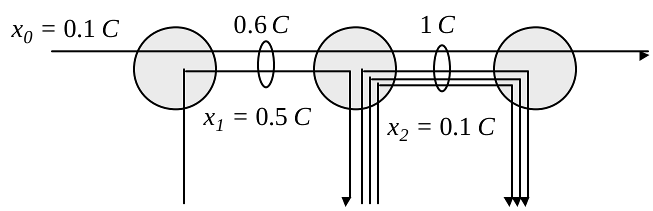

16

Max-Min Fairness

- We can increase x_1 without penalty for other flows
  - $x_0 = 0.1\,C, x_1 = 0.9\,C, x_2 = 0.1\,C$

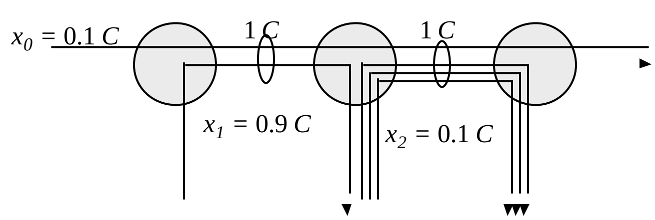

Max-Min Fairness

- Allocating resources in an equal proportion is not a good solution since some sources can get more that others without decreasing others' shares.
  - Max-Min allocation
    - Min : for fairness on bottleneck links
    - Max : for increase throughput whenever possible

Progressive filling

- Bottleneck link *l* for source *s*
  - link *l* is saturated : $\Sigma\ x_i = C$
  - source *s* on link *l* has the maximum rate among all sources using that link
- Progressive filling allocation
  - $x_i = 0$
  - increase x_i equally until $\Sigma\ x_i = C$
  - rates for the sources that use this link are not increased any more (all the sources that are stopped have a bottleneck link)
  - continue increasing the rates for other sources

Example

- Parking lot scenario
  - $x_i = 0$
  - $x_i = d$ until $n_0 x_0 + n_i x_i \leq C$
  - bottleneck link for $d_1 = \min(C/(n_0 + n_i))$, source 0 or i
    - $x_0 = \min(C/(n_0 + n_i))$
  - increase other sources
    - $x_i = (C - n_0 x_0)/ n_i$
- In our example
  - $x_0 = 0.1\, C$, $x_2 = 0.1\, C$
  - $x_1 = 0.9\, C$

Proportional Fairness

- An allocation of rates x_s is ``proportionally fair" if and only if, for any other feasible allocation y_s we have

$$\sum_{s=1}^{S} \frac{y_s - x_s}{x_s} \leq 0$$

- Any change in the allocation must have a negative average change
- Parking lot example with $n_s = 1$
  - max-min fair allocation $x_s = C/2$ for all s
  - let decrease x_0 by δ: $y_0 = C/2 - \delta$, $y_s = C/2 + \delta$, $s = 1, \ldots, I$
  - average rate of change is

$$\left(\sum_{s=1}^{I} \frac{2d}{C} \right) - \frac{2d}{C} = \frac{2(I-1)d}{C}$$

  - not proportionally fair for $I \geq 2$!

Proportional Fairness

- There exists one unique proportionally fair allocation. It is obtained by maximizing

$$J(\vec{x}) = \sum_s \ln(x_s)$$

over the set of feasible allocations

Parking lot example

- For any choice of x_0 we should set x_i such that

  $n_0 x_0 + n_i x_i = C, i = 1, \ldots, I$
- Maximize

$$f(x_0) = n_0 \ln(x_0) + \sum_{i=1}^{I} n_i (\ln(C - n_0 x_0) - \ln(n_i)$$

over the set $0 \le x_0 \le C/n_0$.
- The maximum is for

$$x_0 = \frac{C}{\sum_{i=0}^{I} n_i} \qquad x_i = \frac{C - n_0 x_0}{n_i}$$

- If $n_i = 1$, $x_0 = C/(I+1)$, $x_i = CI/(I+1)$
- Max-min allocation is $C/2$ for all rates - sources of type 0 get a smaller rate, since they use more network resources

Additive increase, Multiplicative decrease

- End-to-end congestion control
  - binary feedback
  - adaptation mechanism of additive increase, multiplicative decrease
- Modeling
  - I sources, rate $x_i(t)$, $i = 1, \ldots, I$
  - link capacity : c
  - discrete time, feedback cycle = one time unit
  - during one time cycle, the source rates are constant, and the network generates a binary feedback signal $y(t) \in \{0, 1\}$
  - sources: increase the rate if $y(t) = 0$ and decrease if $y(t) = 1$
  - feedback

$$y(t) = [if\ (\sum_{i=1}^{I} x_i(t) \leq c)\ then\ 0\ else\ 1]$$

Linear adaptation algorithm

- Find constants u_0, u_1, v_0, v_1, such that

$$x_i(t+1) = u_{y(t)}\, x_i(t) + v_{y(t)}$$

  - converge towards a fair allocation
  - one single bottleneck, so all fairness criteria are equivalent
  - we should have $x_i = c/I$
  - the total rate

$$f(t) = \sum_{i=1}^{I} x_i(t)$$

  should oscillate around c: it should remain below c until it exceeds it once, then return below c

25

Necessary conditions

$$f(t+1) = u_{y(t)} f(t) + v_{y(t)}$$

- we must have

$$u_0 f + v_0 > f, \quad \text{increase}$$

$$u_1 f + v_1 < f, \quad \text{decrease}$$

- this gives the following conditions

$$u_1 < 1 \text{ and } v_1 \leq 0$$

or

$$u_1 = 1 \text{ and } v_1 < 0$$

and

$$u_0 > 1 \text{ and } v_0 \geq 0$$

or

$$u_0 = 1 \text{ and } v_0 > 0$$

26

Ensure fairness

- we need to measure how much a given rate allocation deviates from fairness
- measure of unfairness:
  - the distance between the rate allocation $\vec{x}$ and its nearest fair allocation $\Pi(\vec{x})$, where $\Pi(\vec{x})$ is the orthogonal projection on the set of fair allocations, normalized by the length of the fair allocation

27

Effect on fairness of an increase or a decrease

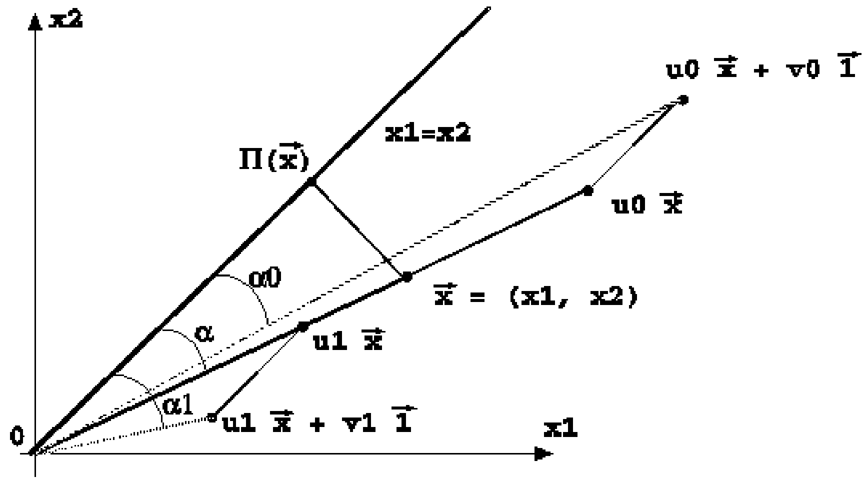

28

Ensure fairness

- when we apply a multiplicative increase or decrease, the unfairness is unchanged
- an additive increase decreases the unfairness, whereas an additive decrease increases the unfairness
- to obtain that unfairness decreases or remains the same, and such that in the long term it decreases

$v_1 = 0$	decrease must be multiplicative
$v_0 > 0$	increase must be additive

Additive increase, Multiplicative decrease

- Fact
  - In order to satisfy efficiency and convergence to fairness, we must have a multiplicative decrease (namely, $u_1 < 1$ and $v_1 = 0$ and a non-zero additive component in the increase (namely, $u_0 \geq 1$ and $v_0 > 0$).
  - If we want to favour a rapid convergence towards fairness, then the increase should be additive only (namely, $u_0 = 1$ and $v_0 > 0$).

Facts to Remember

- In a packet network, sources should limit their sending rate by taking into consideration the state of the network
- Maximizing network throughput as a primary objective may lead to gross unfairness
- Objective of congestion control is to provide both efficiency and some form of fairness
- Fairness can be defined in various ways: max-min, proportional
- Congestion control may use different solutions: rate based, hop by hop, end to end
- End-to-end congestion control in packet networks is based on binary feedback and the adaptation mechanism of additive increase, multiplicative decrease.