A Short Course on Network Calculus

Jean-Yves Le Boudec
Patrick Thiran
ICA, EPFL

CH-1015 Ecublens
leboudec@epfl.ch
http://icawww.epfl.ch

The Network Calculus Book

- available on-line (free download) from
  Le Boudec's or Thiran's home page

  or

  ica1www.epfl.ch/PS_files/NetCal.htm

- this presentation : based on chapters 1, 5 and 6 of the online version

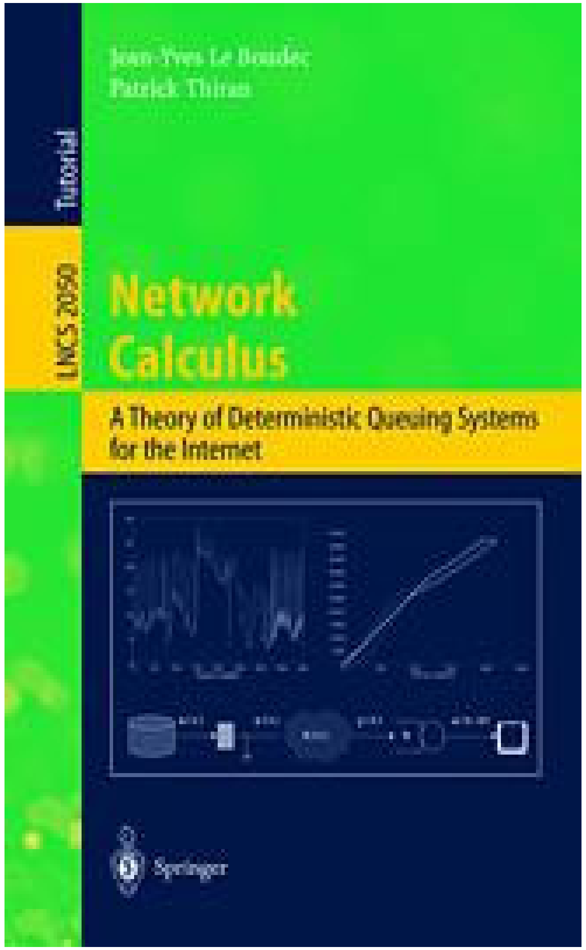

Contents

1. What is Network Calculus ?

- Deterministic analysis of queuing / flow systems arising in communication networks
- Abstraction of schedulers
- uses min, max as binary operators and integrals (min-plus. max-plus algebra)

A simple example

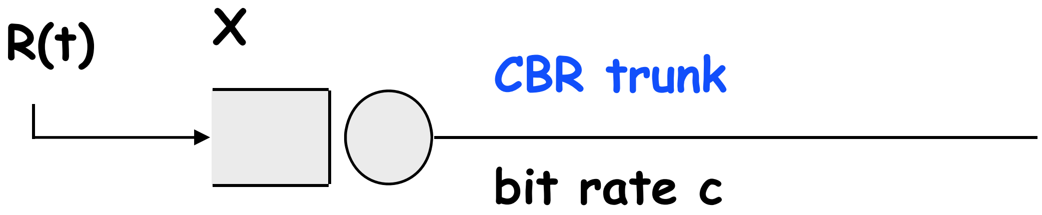

- assume R(t) = sum of arrived traffic in [0, t] is known
- required **buffer** for a bit rate c is

  $\sup_{s \leq t} \{R(t) - R(s) - c\,(t-s)\}$

2. Preliminary Concepts Arrival and Service Curves

- Internet integrated services use the concepts of *arrival curve* and *service curves*

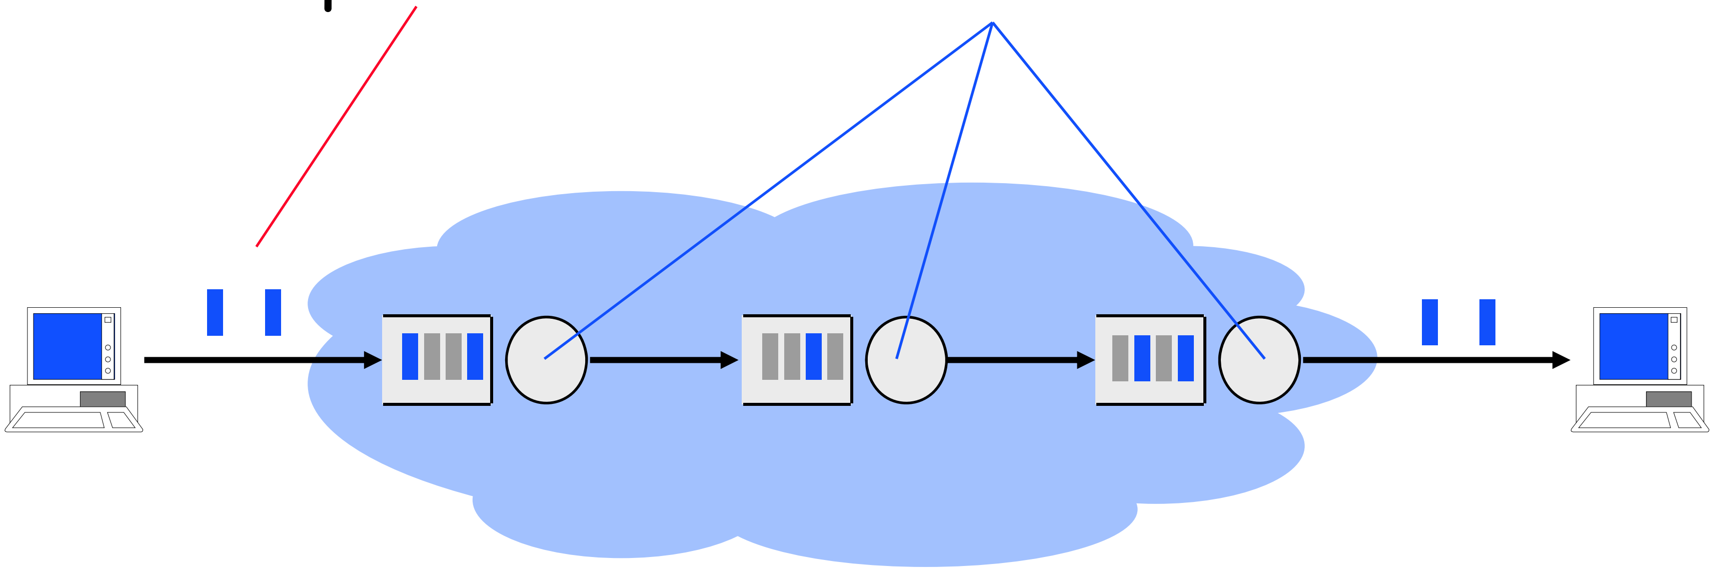

Arrival Curves

- Arrival curve α: $R(t) - R(s) \leq \alpha(t-s)$

Examples:

- leaky bucket $\alpha(u) = ru+b$
- standard arrival curve in the Internet
  $\alpha(u) = \min(pu+M, ru+b)$

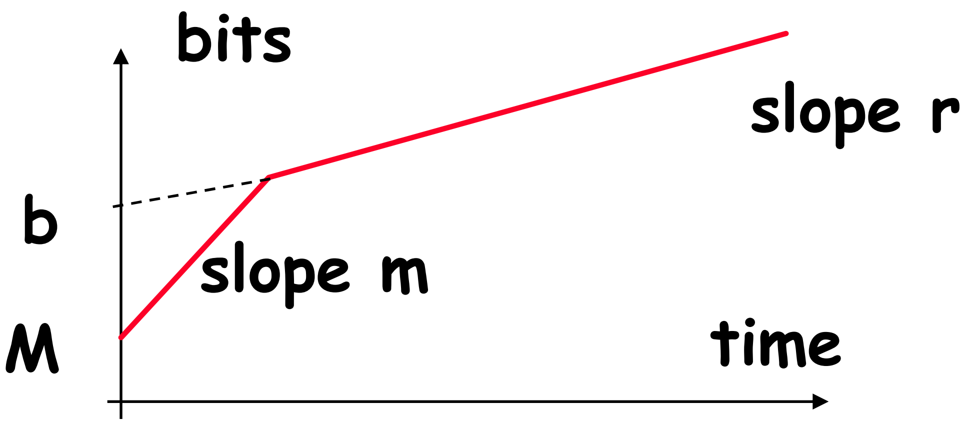

Arrival Curves can be assumed sub-additive

- **Theorem:** α can be replaced by a *sub-additive* function
- sub-additive: $\alpha(s+t) \leq \alpha(s) + \alpha(t)$
- concave => subadditive

Service Curve

- System S offers a service curve β to a flow iff for all *t* there exists some *s* such that

$$R^*(t) - R(s) \geq \beta(t-s)$$

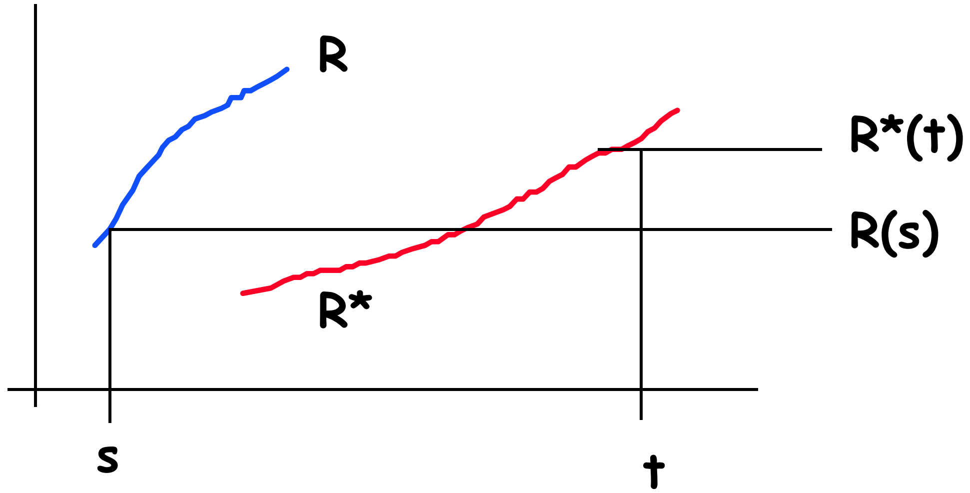

The constant rate server has service curve $\beta(t)=ct$

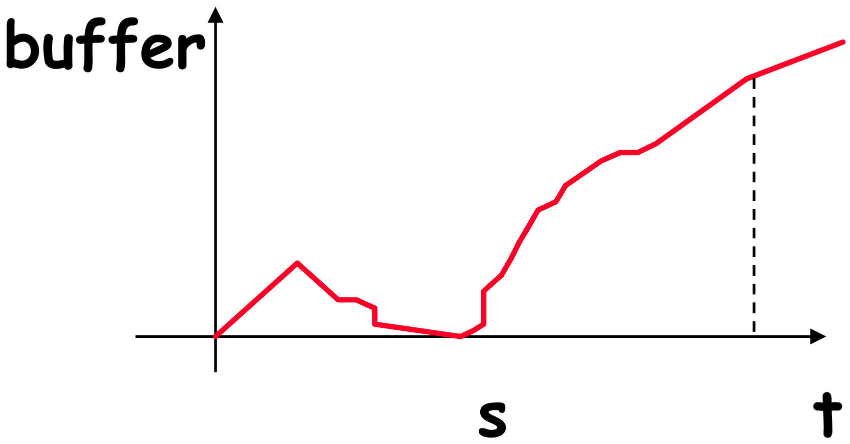

Proof: take s = beginning of busy period:

$$R^*(t) - R^*(s) = c\ (t-s)$$

$$R^*(t) - R(s) = c\ (t-s)$$

The guaranteed-delay node has service curve δ_T

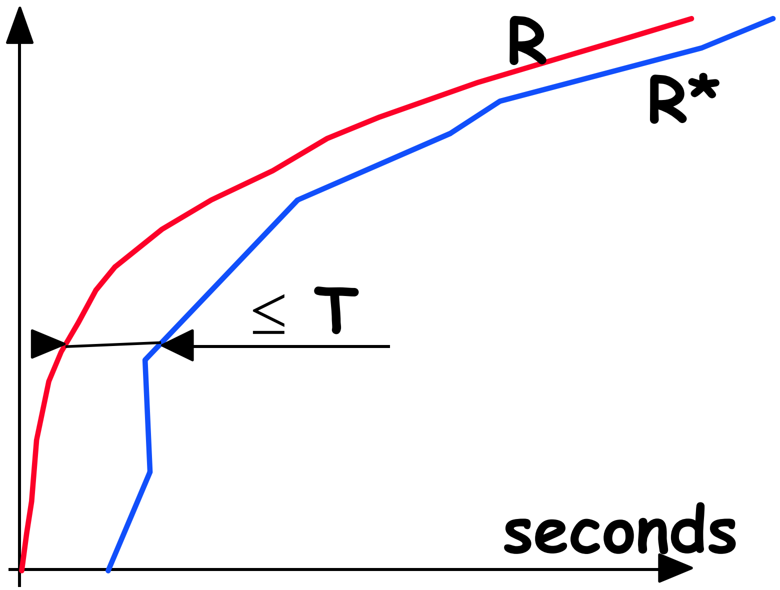

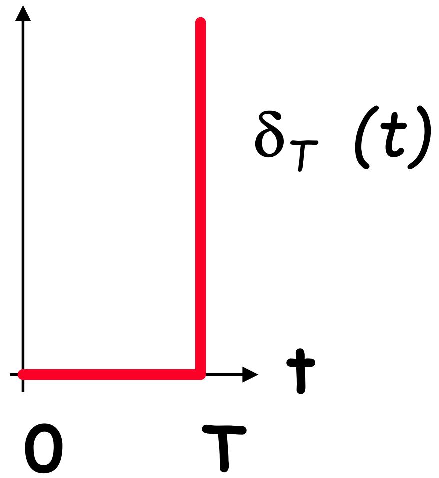

Function δ_T

The standard model for an Internet router

- rate-latency service curve

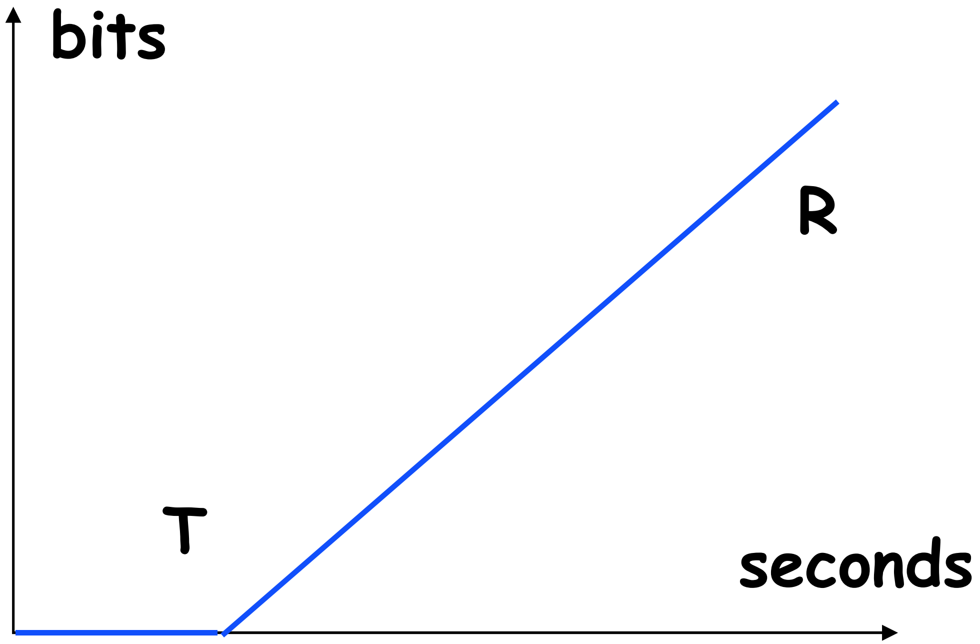

Tight Bounds on delay and backlog

If flow has arrival curve α and node offers service curve β then

- backlog $\leq$ sup $(\alpha(s) - \beta(s))$
- delay $\leq h(\alpha, \beta)$

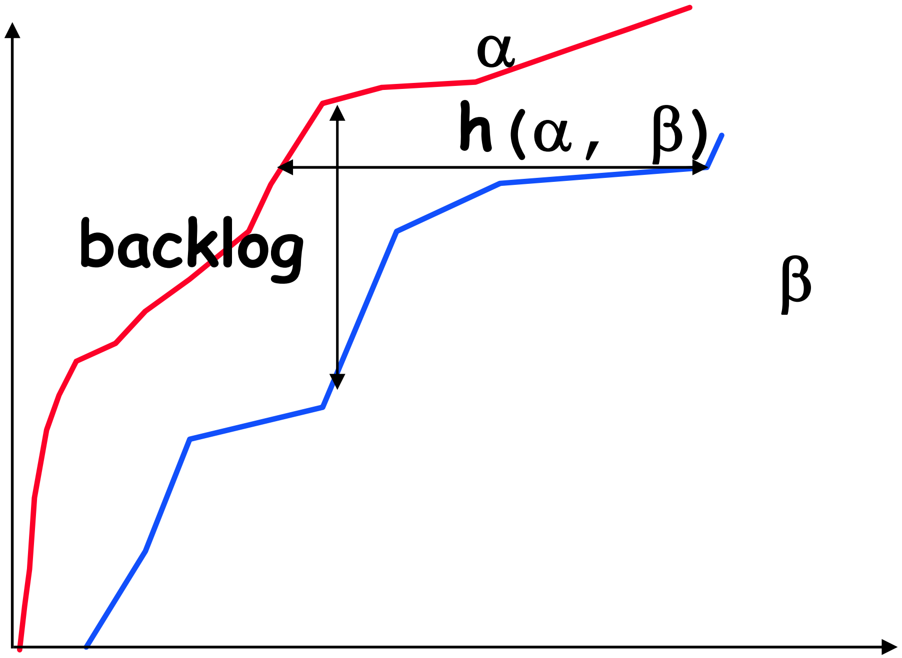

Example

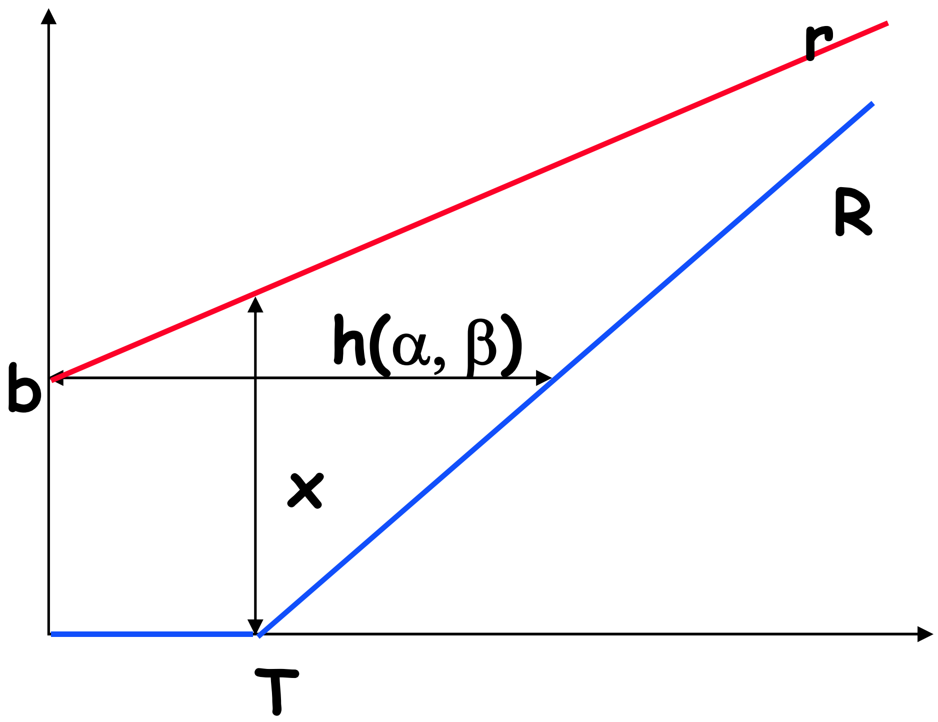

- delay bound: $b/R + T$
- backlog bound: $b + rT$

3. Another linear system theory: Min-Plus

- Standard algebra: $\mathbb{R}, +, \times$

  $$a \times (b + c) = (a \times b) + (a \times c)$$

- Min-Plus algebra: $\mathbb{R}, \min, +$

  $$a + (b \wedge c) = (a + b) \wedge (a + c)$$

Min-plus convolution

- Standard convolution:

$$(f*g)(t)=\int f(t-u)g(u)du$$

- Min-plus convolution

$$f \otimes g\,(t) = \inf_u \{\, f(t-u) + g(u) \,\}$$

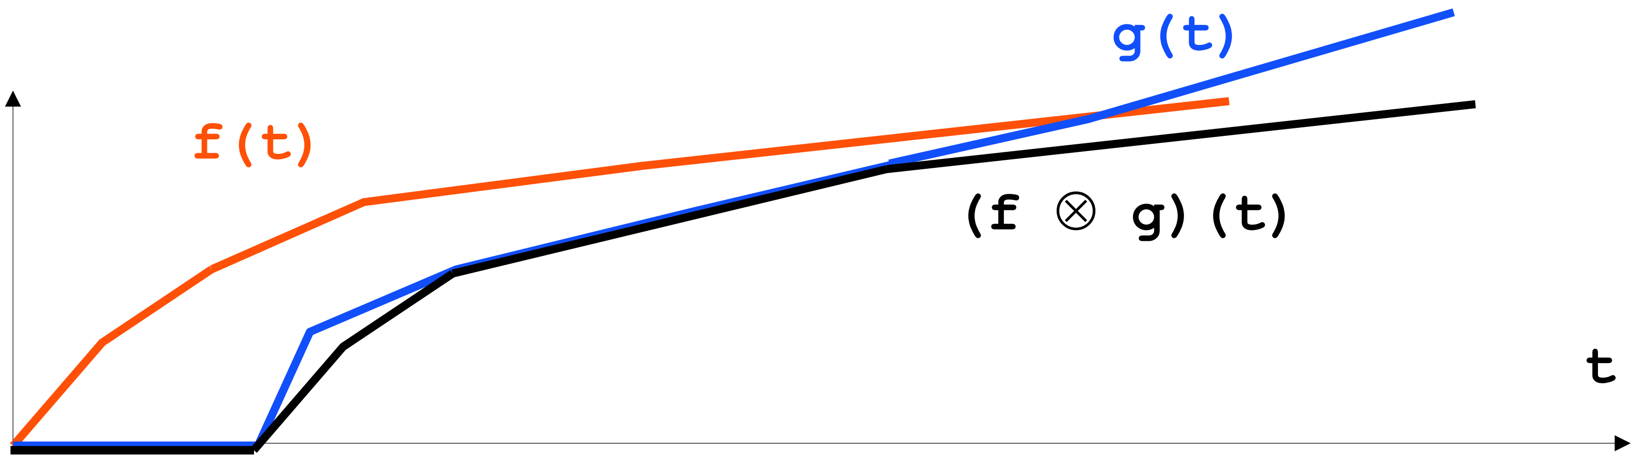

Examples of Min-Plus convolution

- $f \otimes \delta_T (t) = f(t-T)$
- convex piecewise linear curves, put segments end to end with increasing slope

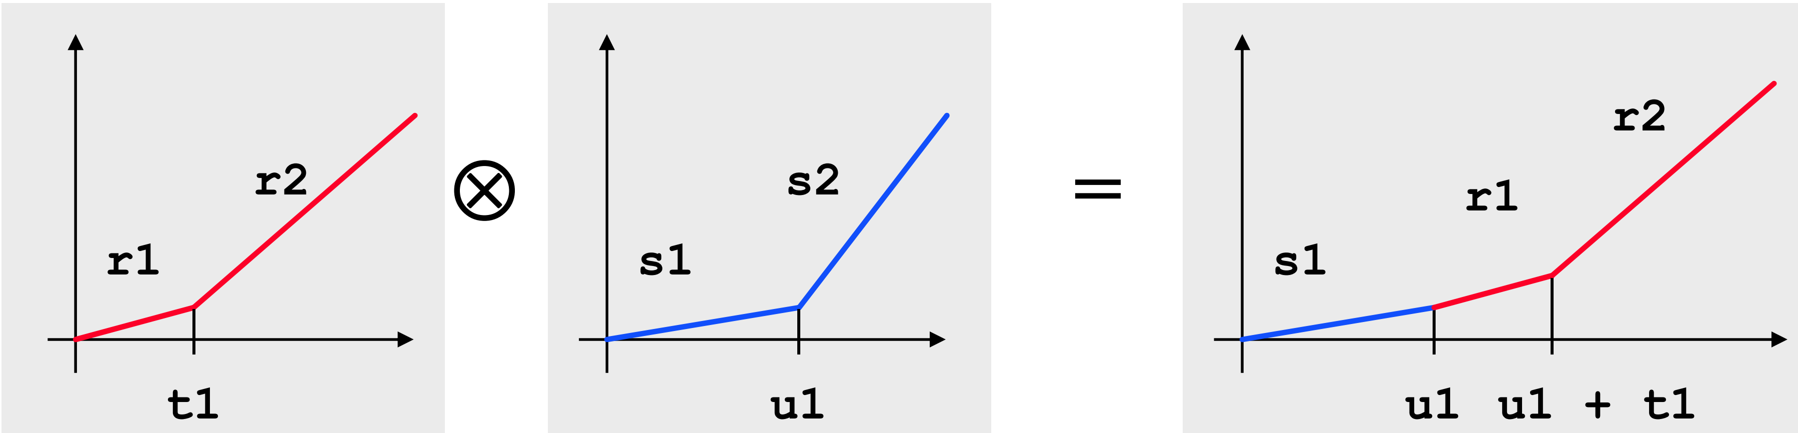

4. We can express arrival and service curves with min-plus

- Arrival Curve property means

$$R \leq R \otimes \alpha$$

- Service Curve guarantee means

$$R^* \geq R \otimes \beta$$

The composition theorem

- **Theorem**: the concatenation of two network elements each offering service curve β_i offers the service curve $\beta_1 \otimes \beta_2$

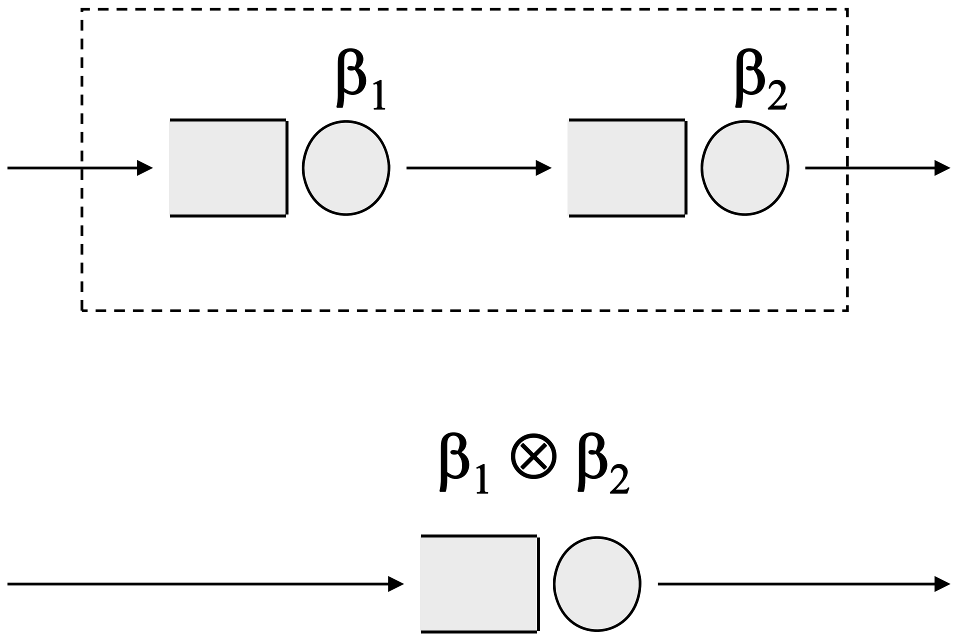

Example

- tandem of routers

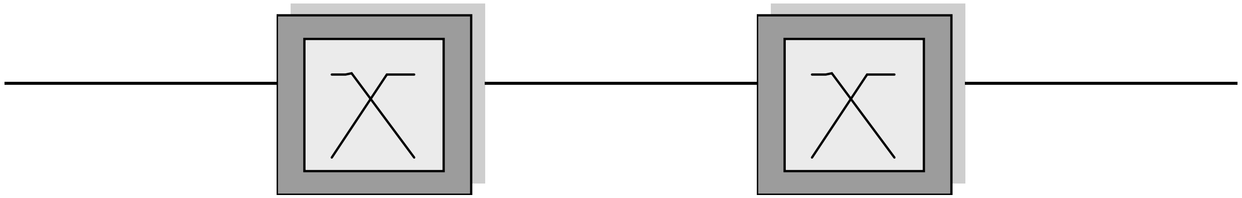

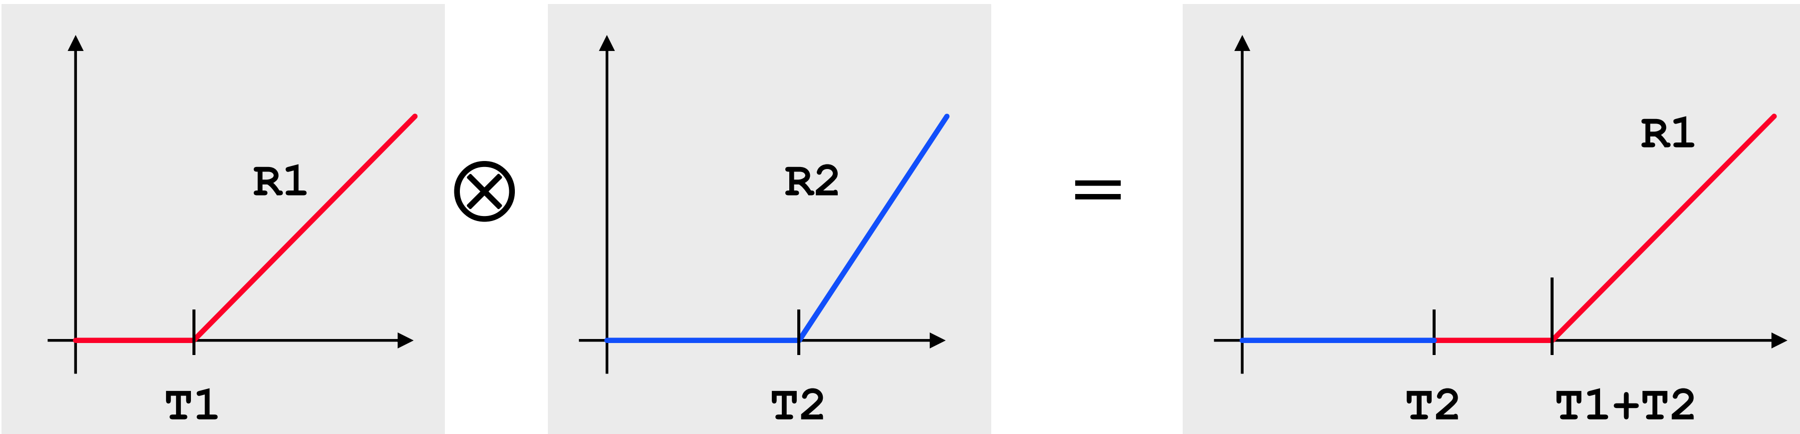

Pay Bursts Only Once

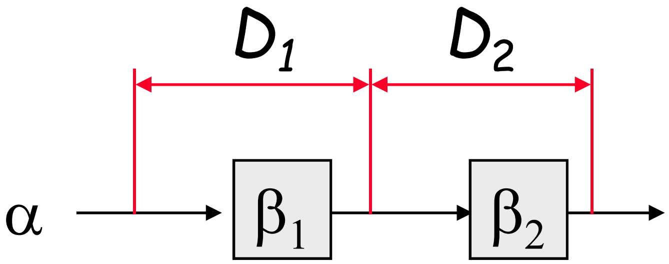

$$D_1 + D_2 \leq (2b + RT_1)/R + T_1 + T_2$$

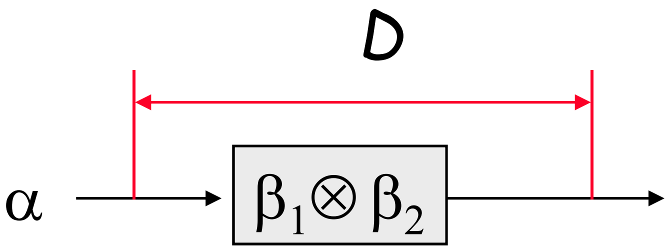

$$D \leq b/R + T_1 + T_2$$

end to end delay bound is less

5. Linear Characterization of Shapers

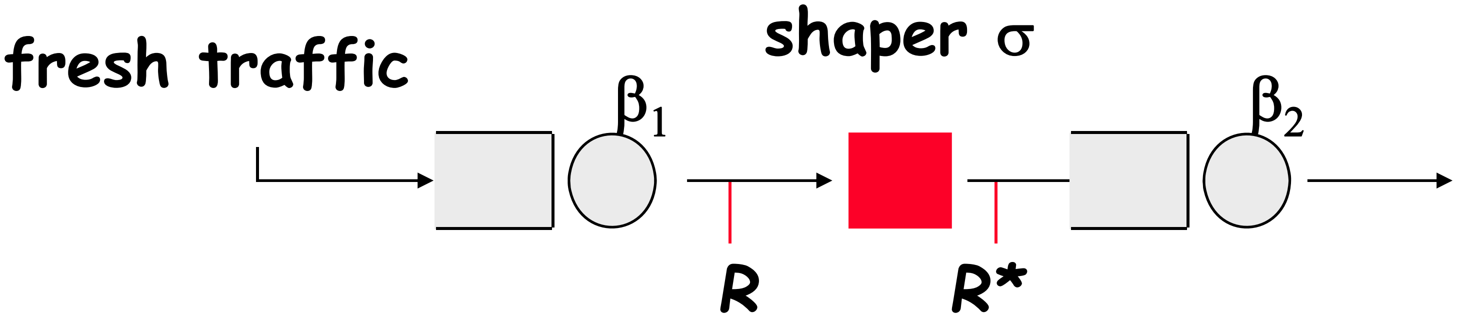

Definition: shaper

- forces output to be constrained by σ
- stores data in a buffer if needed

Theorem

- Output of shaper is $R^* = R \otimes \sigma$

Proof of Shaper Theorem is a typical Min-Plus result

- $R^*()$ is the maximum function $x()$ such that
  - (1) $x \le x \otimes \sigma$
  - (2) $x \le R$
- Solution: $x^0 = R$

  $x^i = x^{i-1} \otimes \sigma$

  R^* is given by $R^* = \inf \{x^0, x^1, \ldots, x^i, \ldots\}$
- Here: $\sigma \otimes \sigma = \sigma$

  $x^i = R \otimes \sigma$ for all $i \ge 1$

  $R^* = R \otimes \sigma$

Application: Re-shaping does not increase worst-case delay

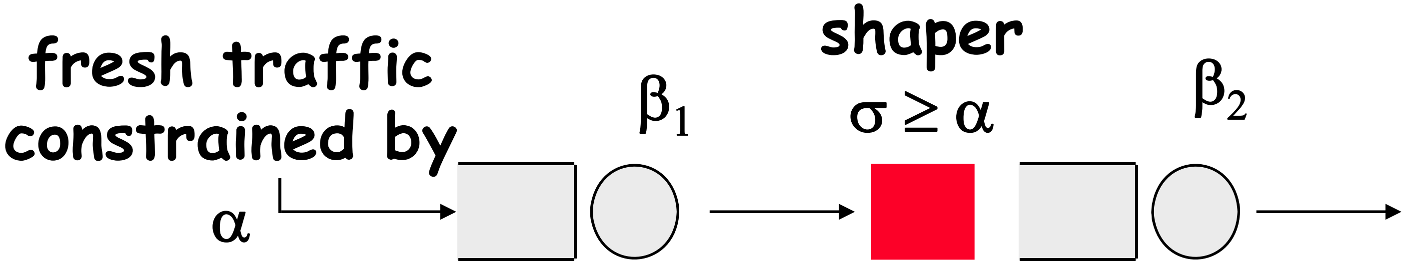

- Re-shaping is for free:
- **Proof** :

$$h(\alpha, \beta_1 \otimes \sigma \otimes \beta_2) = h(\alpha, \sigma \otimes \beta_1 \otimes \beta_2) = h(\alpha, \beta_1 \otimes \beta_2)$$

Shaper Keeps Arrival Constraints

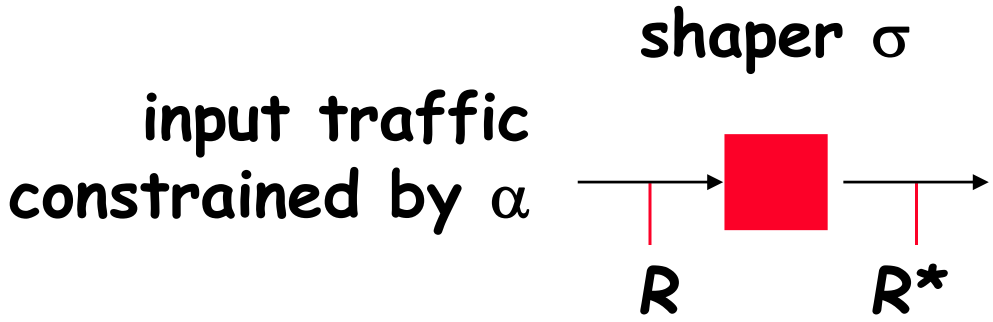

- The output of the shaper is still constrained by α
- **Proof**

  $R^* = R \otimes \sigma \le (R \otimes \alpha) \otimes \sigma = (R \otimes \sigma) \otimes \alpha = R^* \otimes \alpha$

6. Optimal Smoothing with pre-fecthing

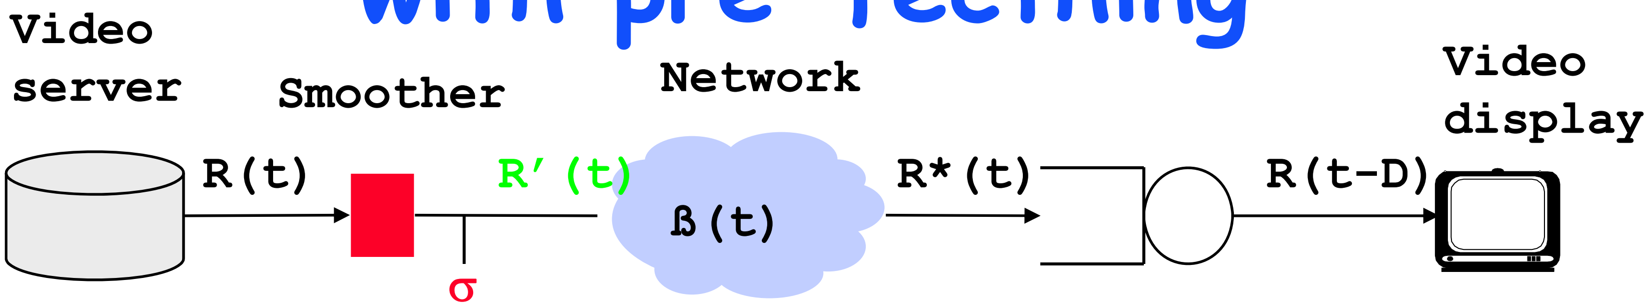

- smoother is a scheduler which produces a σ-smooth output
- may look-ahead
- Q: given signal R(t), σ and β
  what is the minimum playback delay D ?

System Equations

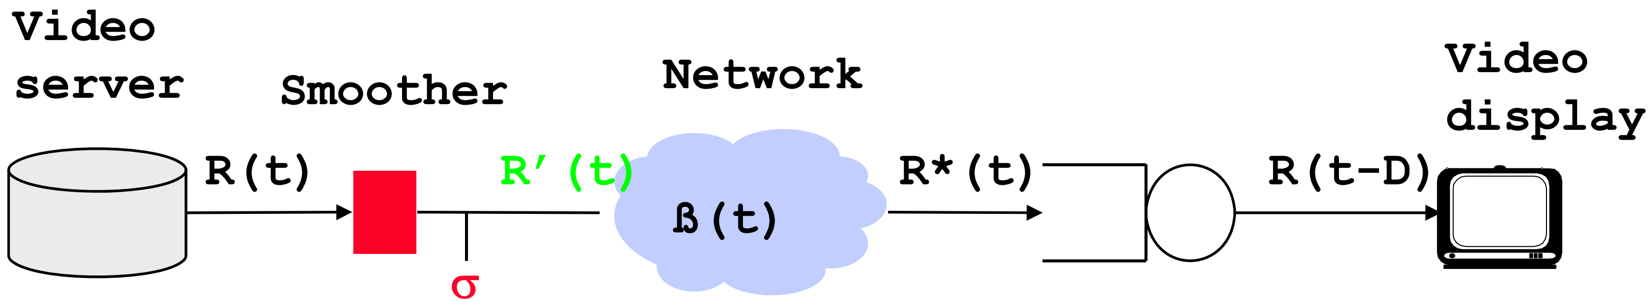

- (1) R' is σ-smooth
- (2) $(R' \otimes \beta)(t) \geq R(t-D)$
- $R'(t) = 0$ for $t \leq 0$

- Define min-plus deconvolution
  $(a \oslash b)(t) = \sup_{s \geq 0} [a(s+t) - b(s)]$
- $x \leq y \otimes \beta$ <-> $x \oslash \beta \leq y$

A max-plus model

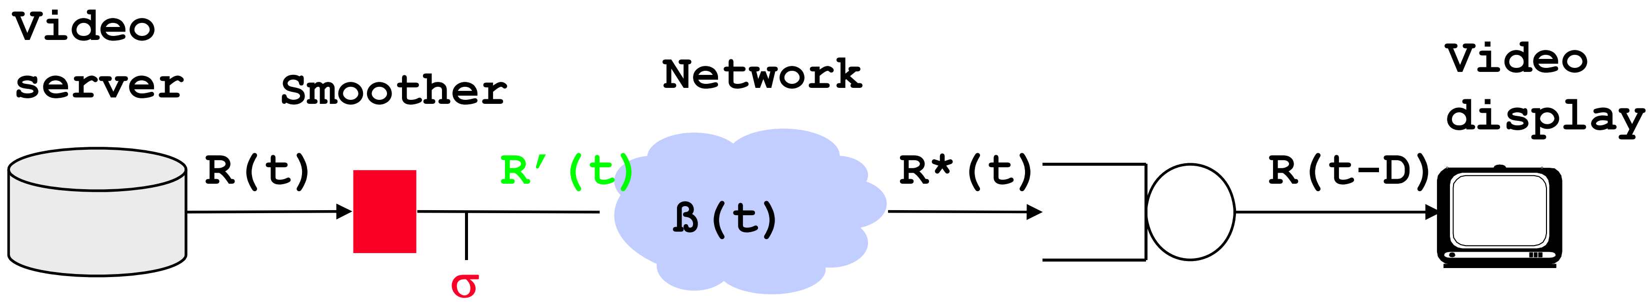

- (1) $R' \geq R' \oslash \sigma$
- (2) $R' \geq (R \oslash \beta)(t-D)$
- this is a max-plus problem, with *minimum* solution

  $x^* = \inf \{x^0, x^1, \ldots, x^i, \ldots\}$

  $x^0(t) = (R \oslash \beta)(t-D)$

  $x^i = x^{i-1} \oslash \sigma$
- thus $R' = (R \oslash \beta) \oslash \sigma\,(t-D) = R \oslash (\beta \otimes \sigma)\,(t-D)$

Example

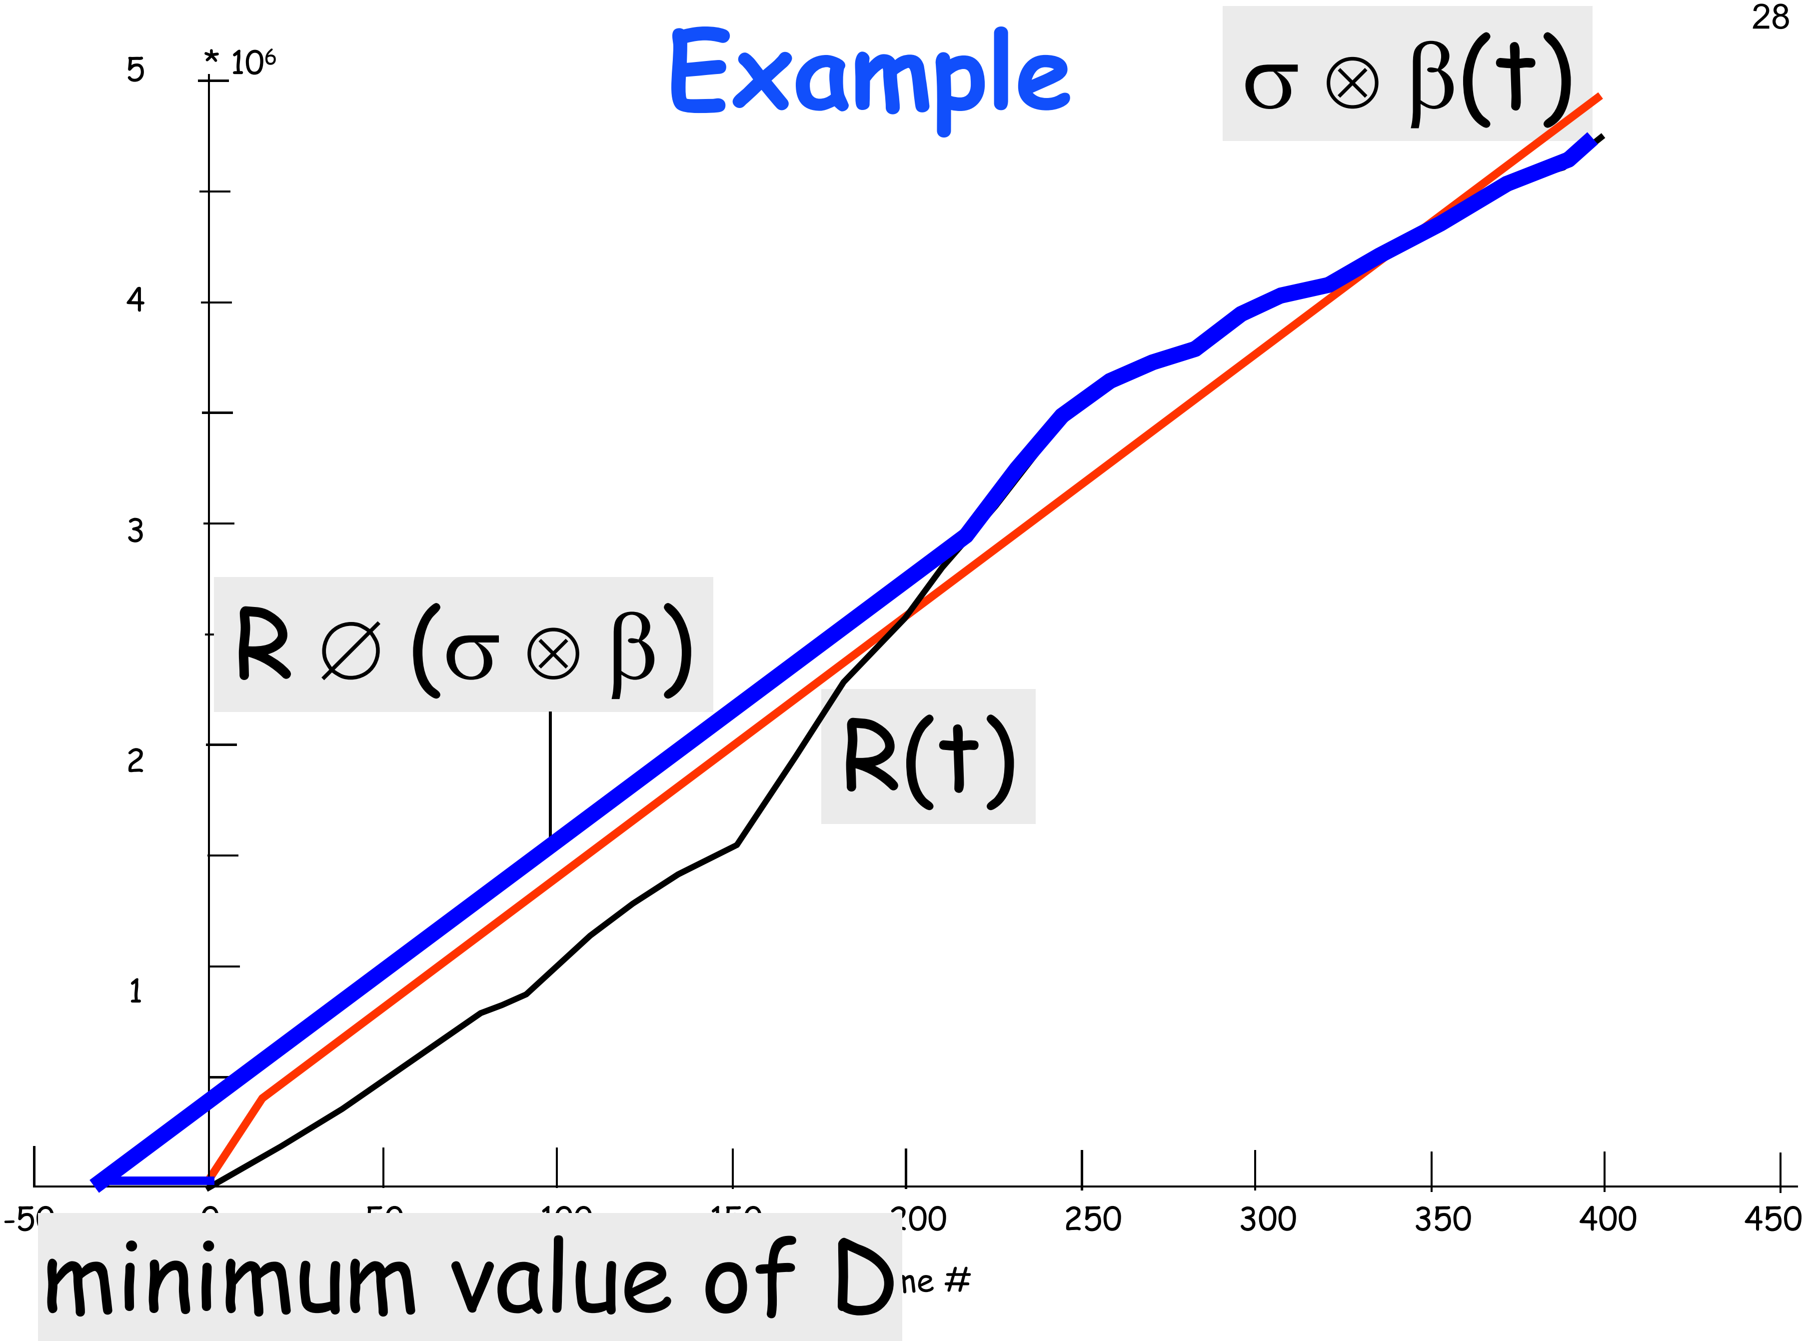

$\sigma \otimes \beta(t)$
* 10^6
5
4
3
2
1
$R \oslash (\sigma \otimes \beta)$
$R(t)$
-50
0
50
100
150
200
250
300
350
400
450
minimum value of D
ne #

Minimum Playback Delay

- D must satisfy :

  $R \oslash (\beta \otimes \sigma)(-D) \geq 0$

- this is equivalent to

  $D \geq h(R, \beta \otimes \sigma)$

Back to our example

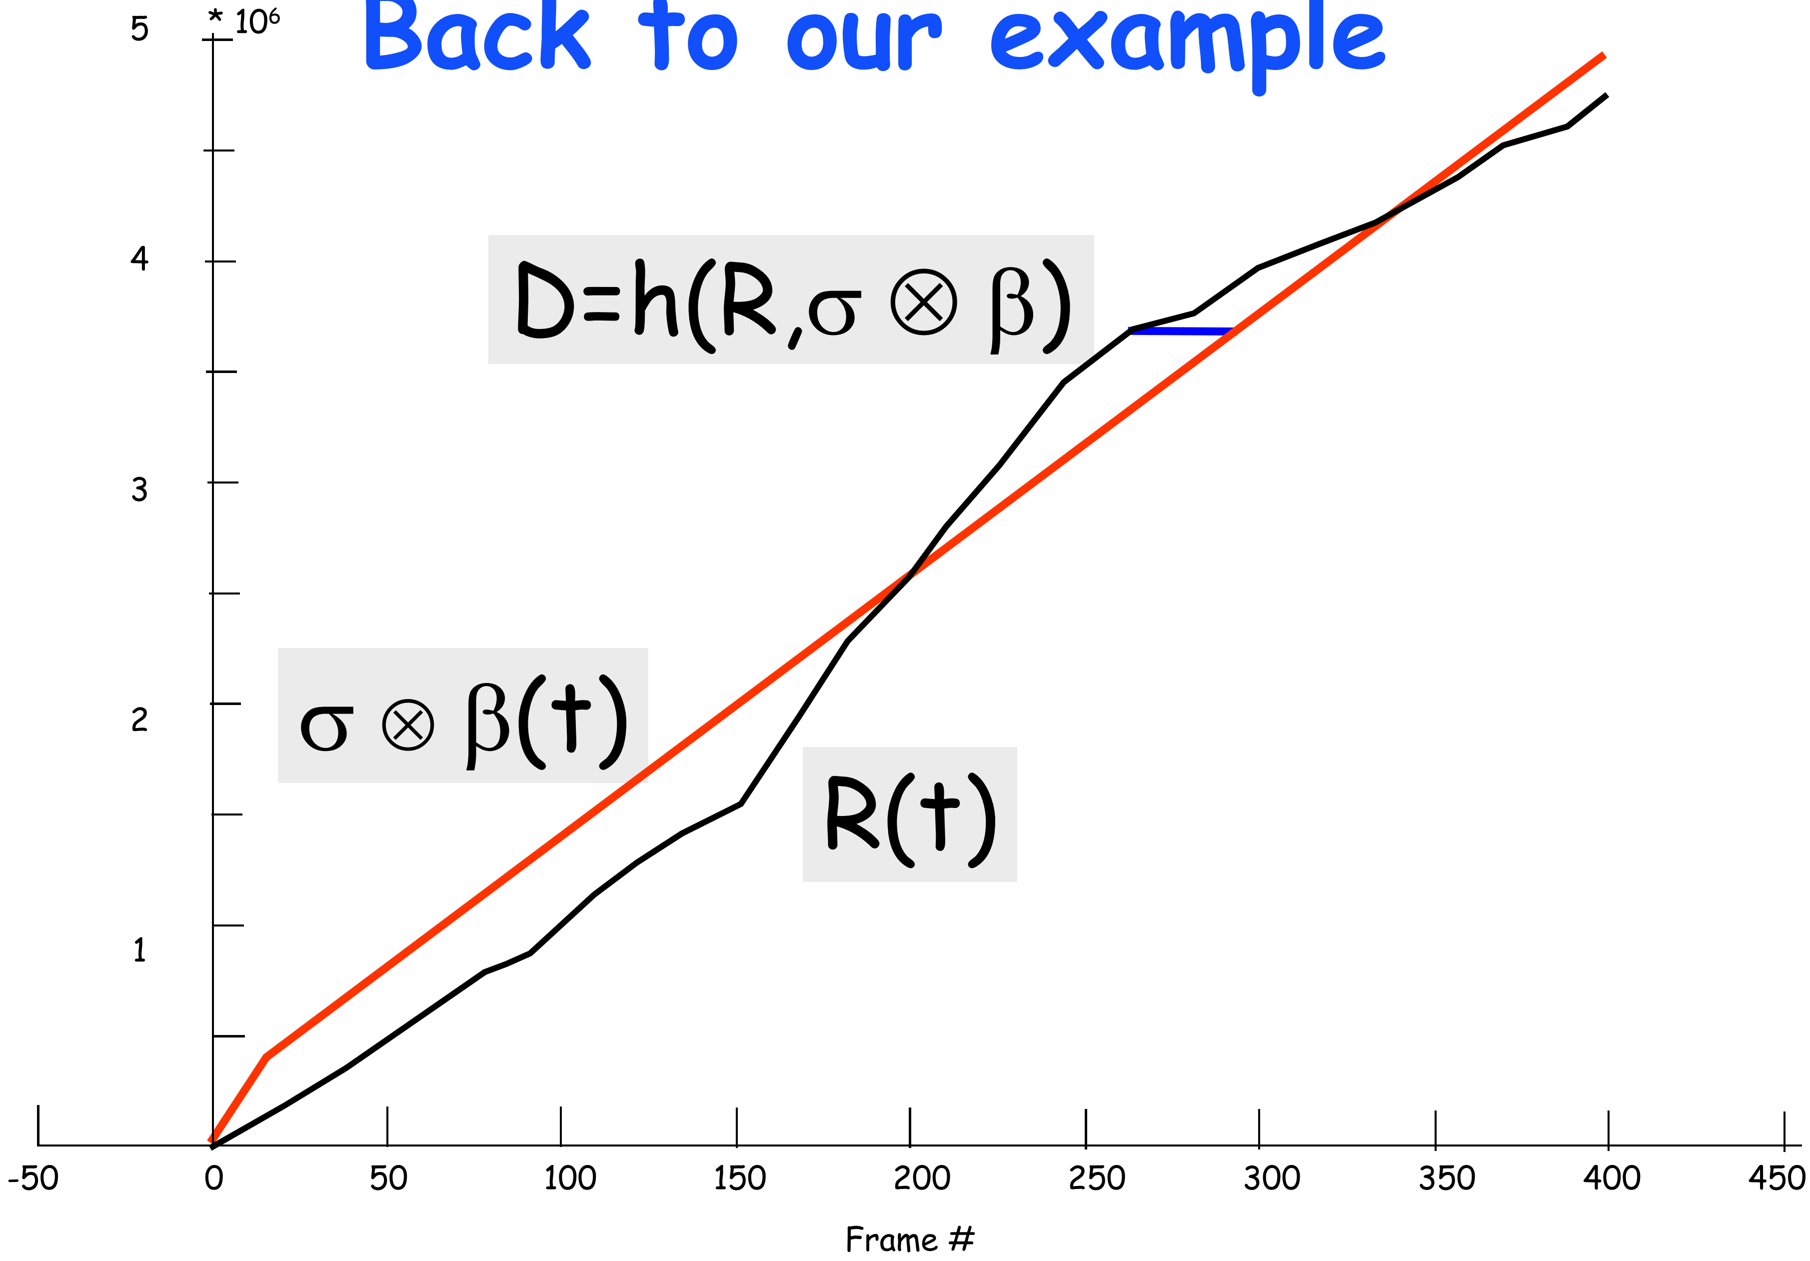

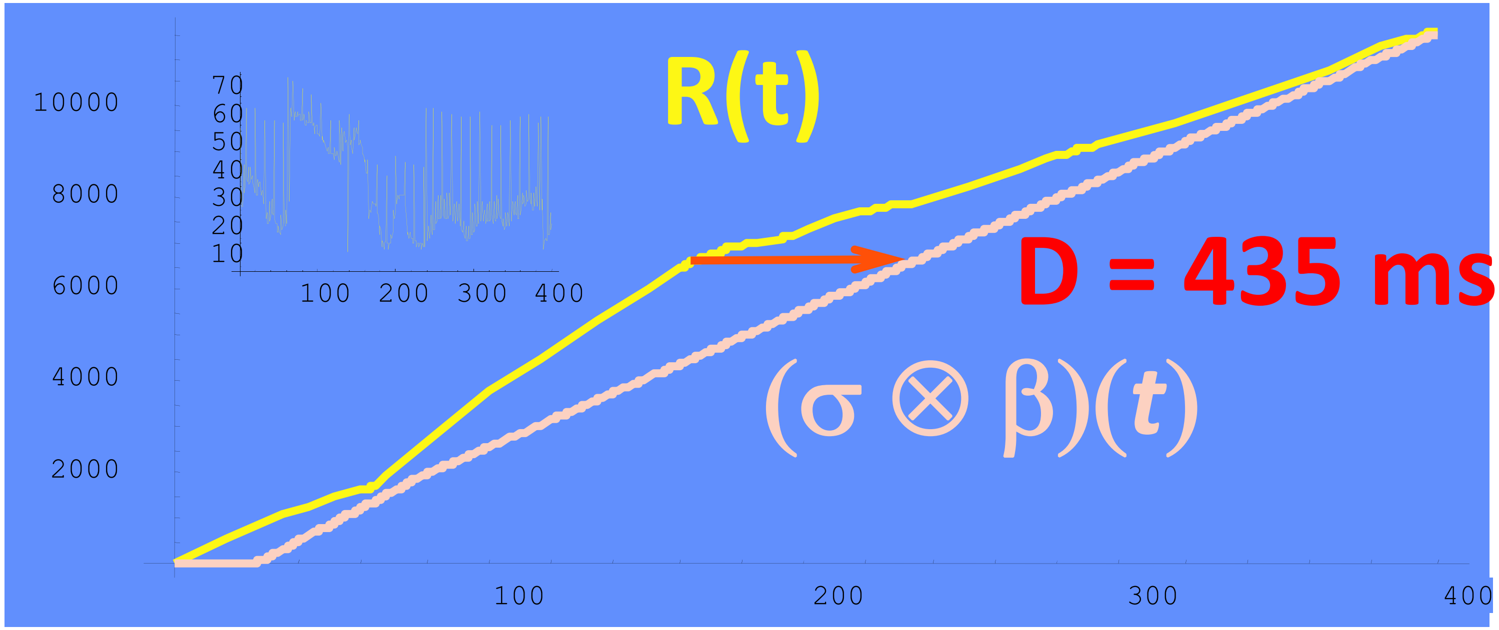

R(t)
D = 435 ms
(σ ⊗ β)(t)
10000
8000
6000
4000
2000
100
200
300
400
70
60
50
40
30
20
10
100 200 300 400

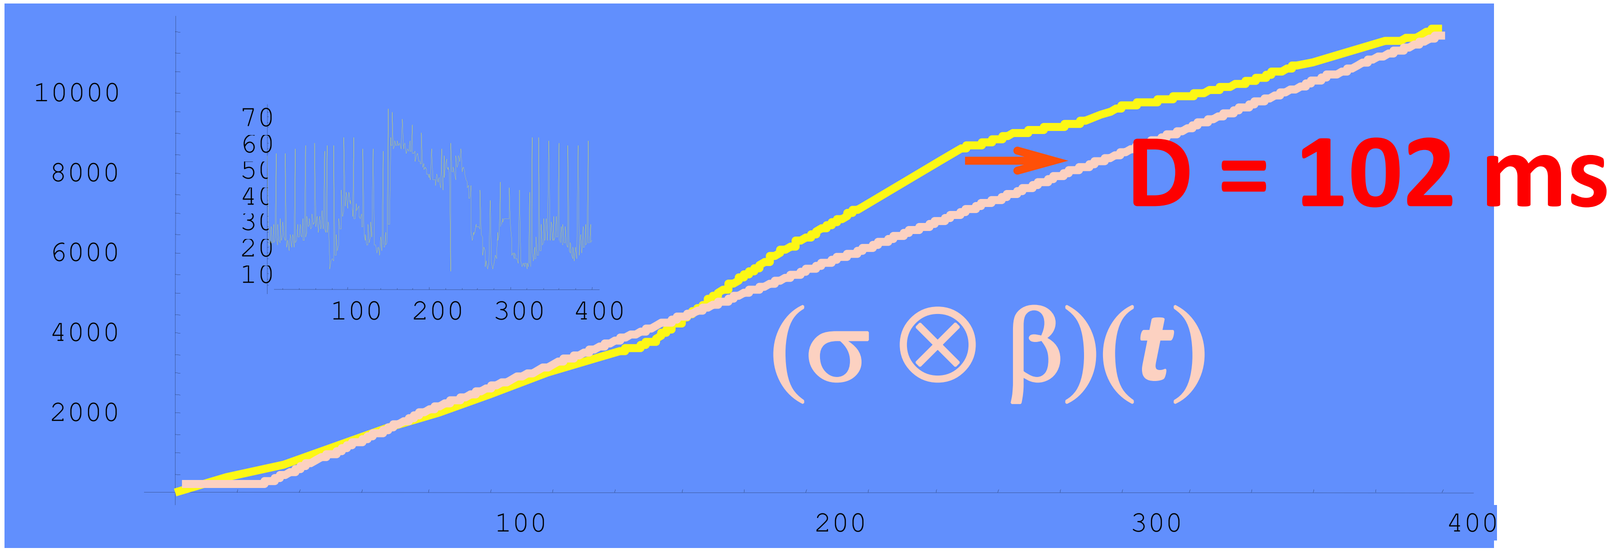

D = 102 ms
(σ ⊗ β)(t)
10000
8000
6000
4000
2000
100
200
300
400
70
60
10
100 200 300 400

Other examples

- Bounds for Diff-serv (chapter 2)
- Packet Scale Rate Guarantee (chapter 7)
  - used in the definition of EF (infocom 2001)

7. Diff-Serv Bounds Problem Definition

- Consider a network implementing aggregate scheduling
  - For example: Expedited Forwarding (EF)

- Problem: find a bound for the end-to-end delay variation that is
  - valid for any network topology
  - has closed form
  - does not require per flow information

We assume aggregate scheduling and leaky bucket constraint microflows

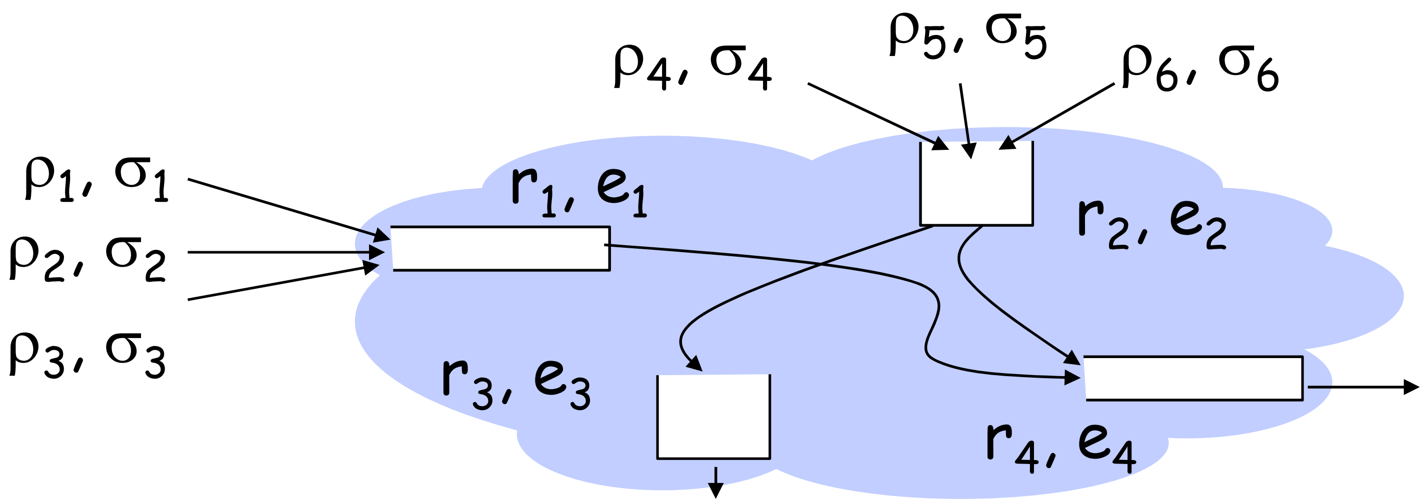

- microflows constrained by leaky buckets (ρ_i, σ_i)
- the aggregate receives a guaranteed service, with service curve $\beta_l(t) = r_l(t - e_l)^+$

State of the art

- Utilization factor $\nu = \max_l (\Sigma_f \rho_f / r_l)$
- Even if network is subcritical (ν <1), we don't know:
  - good bounds for a general topology
  - if network is stable
- M. Andrews presents an example of a network which is *unstable* for $\nu = 1$ and claims it is also true for some $\nu < 1$

Our main result is a delay bound valid for small utilization factors

$$\boxed{D = \frac{h}{1-(h-1)\nu}(e+\tau)}$$

$$\nu \le \nu_{crit} = \frac{1}{(h-1)}$$

- A better bound exists when peak rates C_m are known

Examples (h=10 hops)

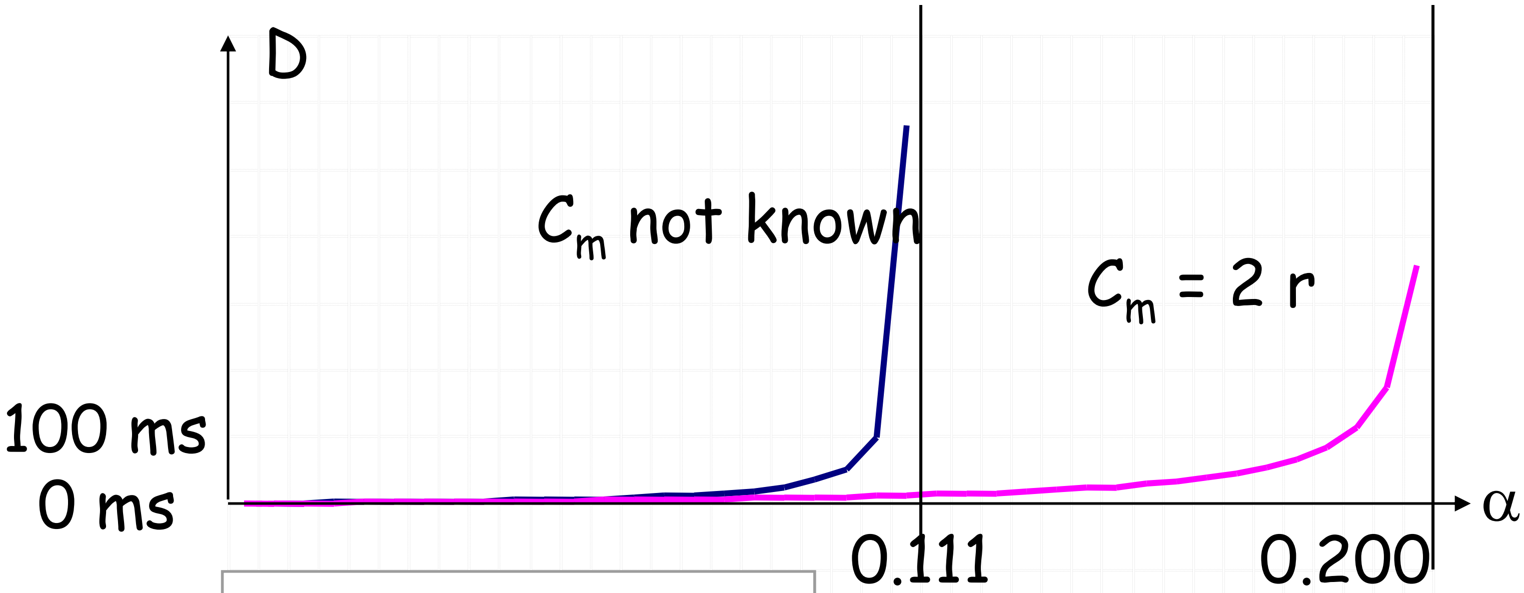

C_m not known

$\nu_{cri} = 1 / (h-1)$

$C_m = 2C$

$\nu_{cri} = 2 / h$

Derivation of the bound

- Assume nodes are GR (orFIFO-per aggregate rate latency service curve elements)

1) Assume delay bound hD on low delay traffic (EF) exists, where h = max number of hops, D = max delay bound per node

2) An arrival curve of aggregate traffic at node i

$$\alpha_i(t) = \Sigma_{m \ni i} (\rho_m t + (h-1)\rho_m D + \sigma_m) = \nu_i R_i t + (h-1) \nu_i R_i D + \nu_i R_i \tau_i$$

where $\nu_i = (\Sigma_{m \ni i} \rho_m)/R_i$ and $\tau_i = (\Sigma_{m \ni i} \sigma_m)/(\Sigma_{m \ni i} r_m)$

3) Compute horizontal distance between $\alpha_i(t)$ and $\beta_i(t)$:

$$D_i = e_i + (h-1) \nu_i D + \nu_i \tau_i$$

4) Deduce $D <= (e + \nu\tau)/(1 - (h-1)\nu)$ where $e = \max_i e_i$, $\nu = \max_i \nu_i$ and $\tau = \max_i \tau_i$

5) Show that finite bound exists at any time t, and let t -> ∞

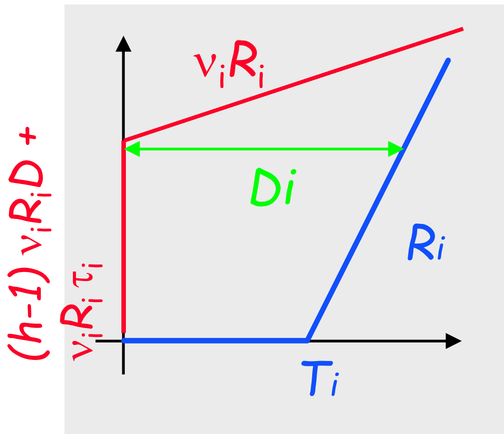

Is the bound tight ?

- No
- Consider only ν_{cri}
- Example:
  - Unidirectional ring with ***n*** nodes
  - ***n*** flows, each of them going through ***n*** hops
  - $\nu_{cri} = 1$ [Tassiulas96]

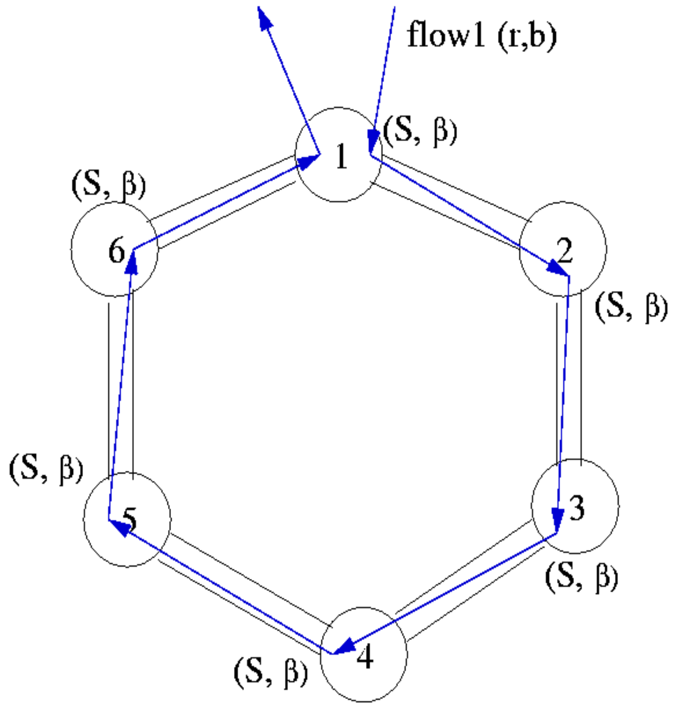

Is the bound tight ?

- 1. No
- 2. Yes
- However, the explosion may always occur at the point predicted by the bound
- More precisely, if $\nu > 1/(h-1)$, we can always construct a network where the worst case delay can be arbitrarily large

The elements in our toy network

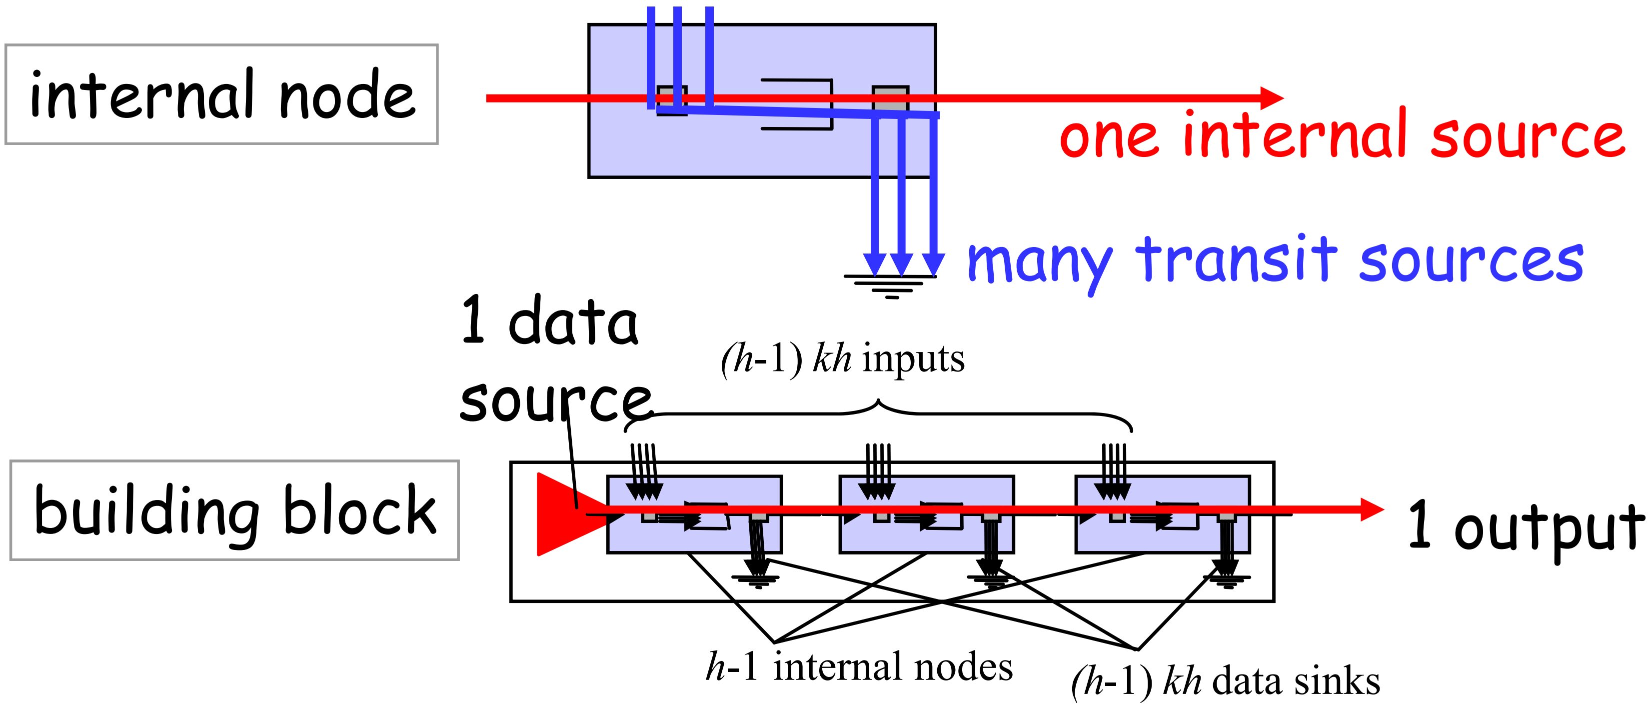

- If transit sources are able to send *b* packets back to back, then we can arrange the internal source to output a burst of *b*+1 packets

The output of level J-1 building blocks is a transit source for level J building blocks

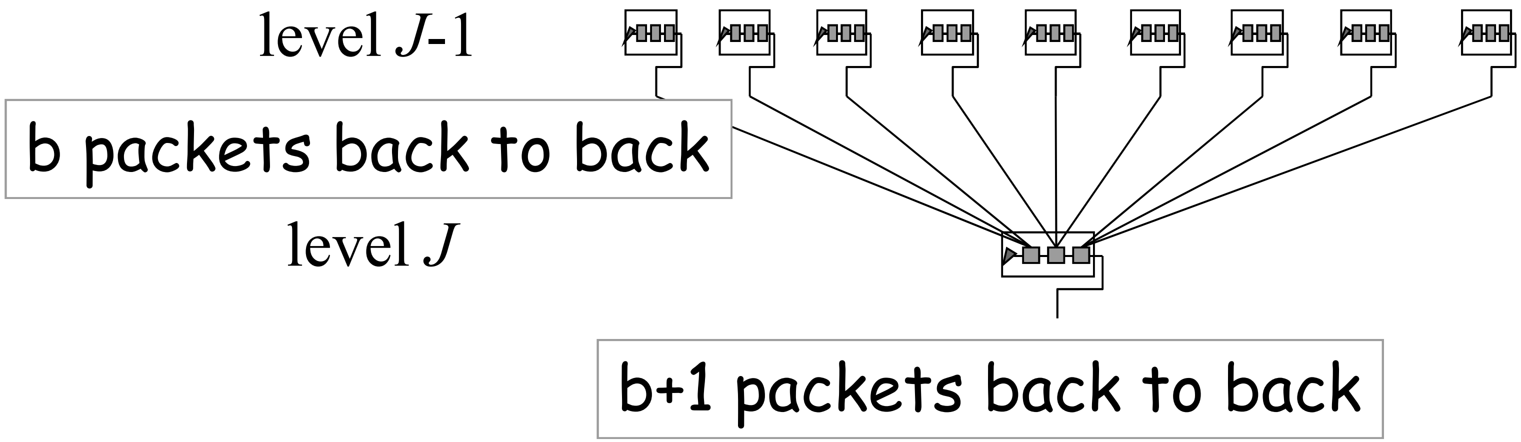

A network where $\alpha > \alpha_{cri}$ with arbitrarily large delay

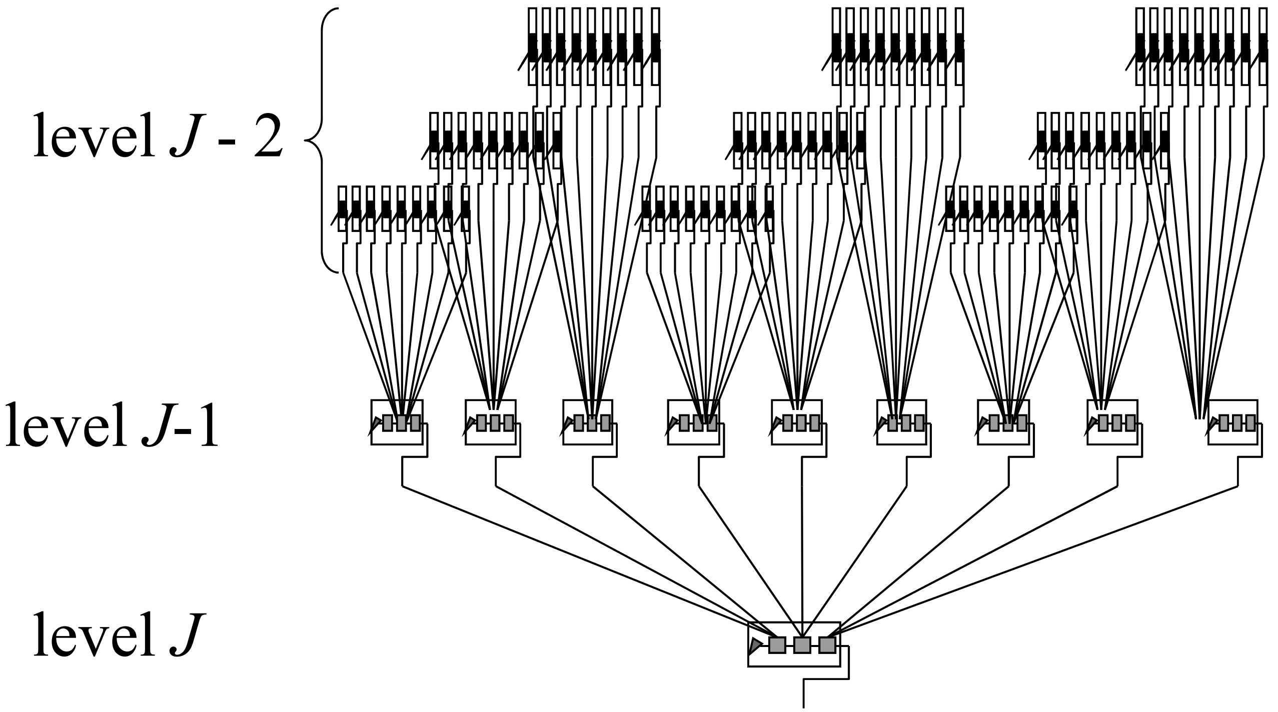

- The worst-case delay at level J is at least $J\tau$, with τ depending only on α and h

Conclusion

- Network Calculus is a set of tools and theories for the deterministic analysis of communication networks
- Application of min-plus algebra
- Explains and proves delay and backlog properties in networks
- Book and slides available online at Le Boudec's or Thiran's home page