Network Calculus
Regulators : From Integrated Services to Deterministic Networking

Jean-Yves Le Boudec[1]

2018 International Workshop on Network Calculus and Applications (NetCal 2018)

7 September 2018, Vienna, Austria

https://people.epfl.ch/105633/research

Abstract:

Flow constraints are the heart of network calculus, as they are allow to derive deterministic delay and backlog bounds. At the origin is Cruz's arrival curve constraint [Cruz PhD Dissertation 1987], which can easily be expressed with min-plus algebra and was used by the Internet integrated services. Other constraints that cannot be cast as arrival curves are the length-rate quotient rule and more generally Chang's g-regularity, expressed with max-plus algebra. More recently, IEEE TSN and IETF Detnet use a traffic specification based on packet rates which does not fit either of these formalisms. In this talk, we describe a new formalism for flow constraints, called Pi-regularity, which subsumes all of these. We define a new general concept of minimal regulators, which subsumes greedy shapers and minimal g-regulators. We show that appending a minimal regulator after any arbitrary FIFO system does not increase the per-flow worst case delay. Last, we review the concept of interleaved regulator, which acts on a serialized multiplex of flows without using per-flow queues. We show that appending a minimal regulator after any arbitrary FIFO system does not increase the per-class worst case delay. We explain how this feature can be used to analyze TSN and Detnet systems of any size and complexity.

[Le Boudec 2018] Le Boudec, Jean-Yves, "A Theory of Traffic Regulators for Deterministic Networks with Application to Interleaved Regulators", *arXiv preprint arXiv:1801.08477*.

Contents

1. FIFO Per-Class Networks

FIFO per class are commonly used in Time Sensitive Networking (IEEE 802.1 TSN, IETF Detnet).

Computing backlog and delay bounds is hard [Bennett et al 2002]:

- burstiness of every flow increases at every hop as a function of other flows burstiness
- Increased burstiness causes increased burstiness (cascade).

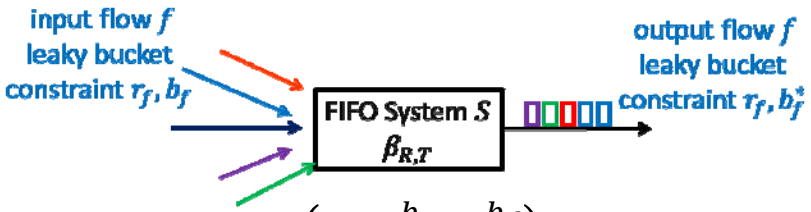

$$b_f^* = b_f + r_f\left(T + \frac{b_{tot} - b_f}{R}\right), \quad b_{tot} = \sum_i b_i$$

[Le Boudec-Thiran 2001, Section 6.4]

Several techniques find improved bounds [Amari et al 2016] [Boyer et al 2012] [Bouillard-Stea 2015][Bondorf et al 2017][Bouillard et al 2018][Rizzo-Le Boudec 2008].

Avoiding Burstiness Cascade

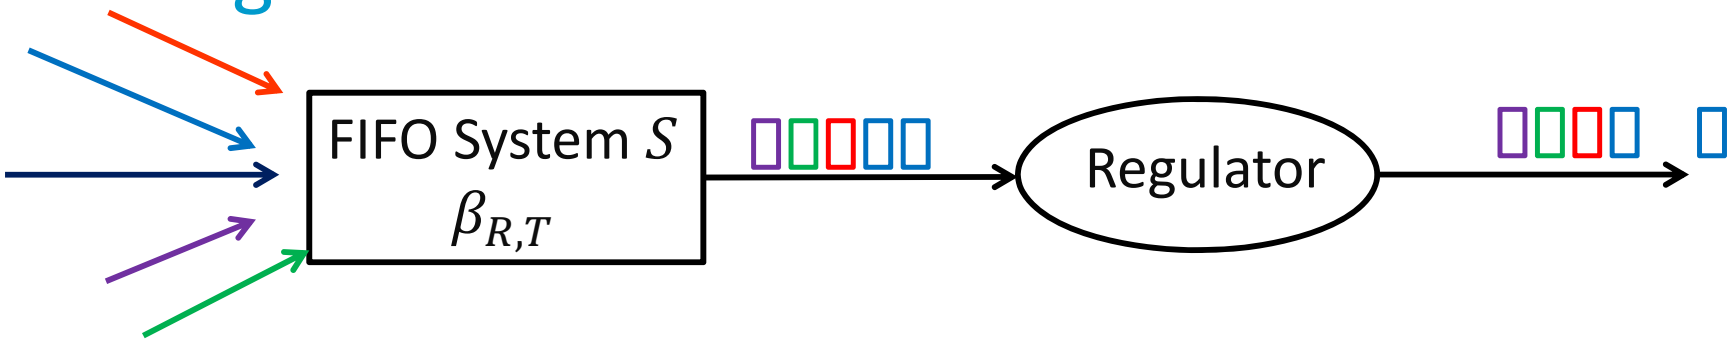

Regulate every flow at every hop (per-flow shaping)

Issue 1: how to perform this without per-flow queue

[Specht-Samii 2016] “Urgency Based Scheduler”, now called “Asynchronous Traffic Shaping” at IEEE TSN

Issue 2: latency due to regulator

2. Flow Regulation: Arrival Curve

One flow; packets of lengths $L_1, L_2, \ldots$ arrive at times $A_1 \leq A_2 \leq \cdots$

Arrival function:

$R(t) = \sum_{n=1,2,\ldots} L_n 1_{\{A_n<t\}}$

σ is arrival curve

$\Leftrightarrow R(t) - R(s) \leq \sigma(t-s)$ for $s \leq t$

$\Leftrightarrow R \leq R \otimes \sigma$ (min-plus convolution)

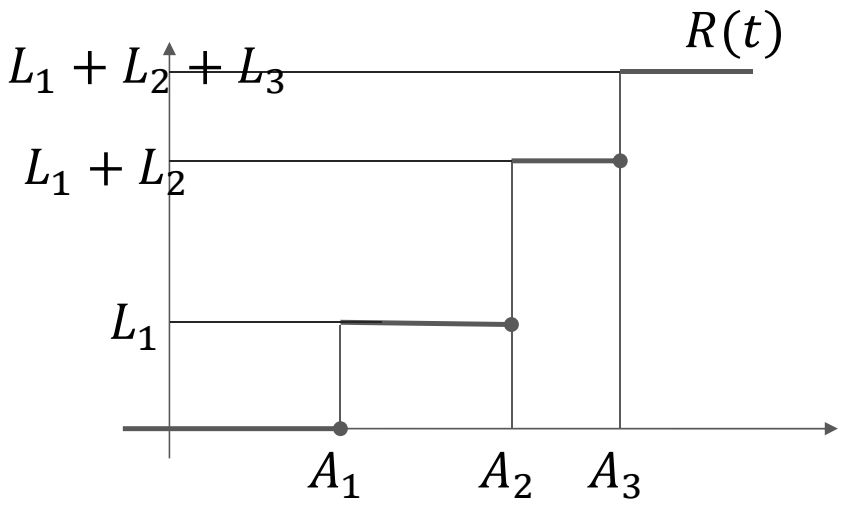

Superposition : if flow i has arrival curve σ_i then the superposition has arrival curve $\sigma = \sum_i \sigma_i$

[Cruz PhD Dissertation 1987]

Examples of Arrival Curves

Affine (leaky bucket): $\sigma(t) = rt + b$ (b is called burstiness)

IETF Intserv Traffic Specification:

$\sigma(t) = \min(pt + M, rt + b)$

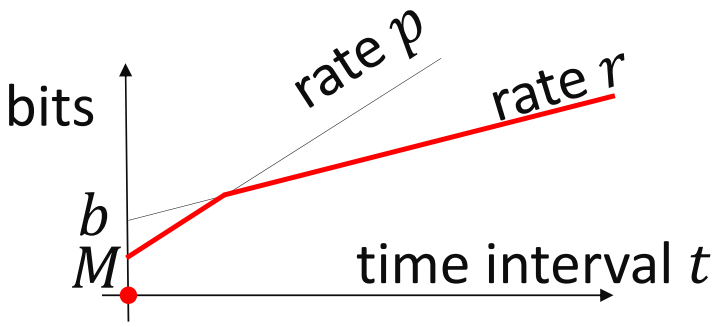

Staircase function

at most b bits in any τ second

$\sigma(t) = b\left\lceil \frac{t}{\tau} \right\rceil$

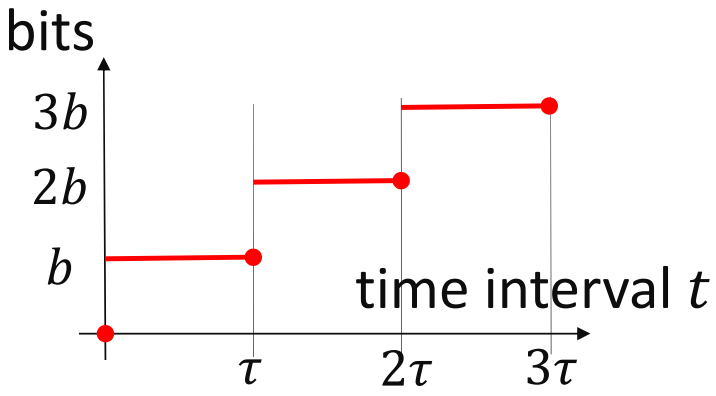

Greedy Shapers

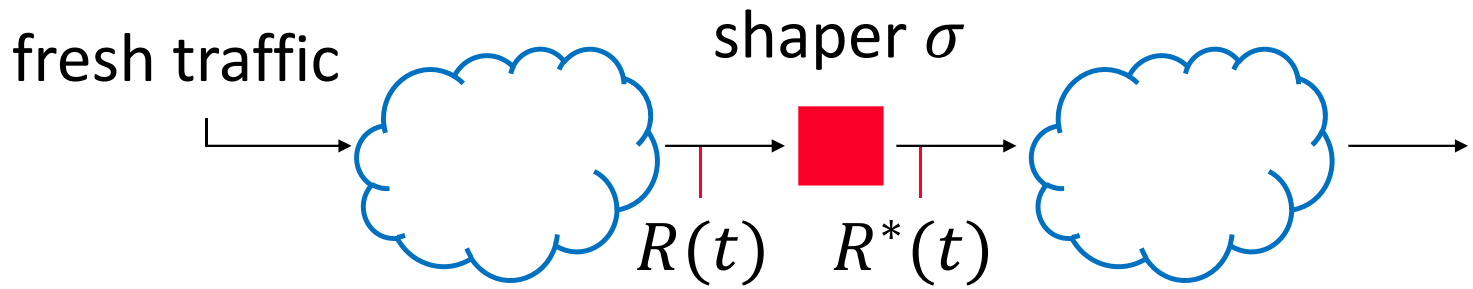

A shaper forces output to have σ as arrival curve

A shaper with output $R^*(t) \leq$ is maximal if $R^*(t) \geq R'^*(t)$ for any other shaper with output $R'^*(t)$.

There exists a maximal shaper (greedy shaper) given by

$$R^*(t) = (\sigma \otimes R)(t)$$

Properties of Greedy Shapers

Re-shaping does not increase end-to-end delay bound with per-flow scheduling and service curve elements

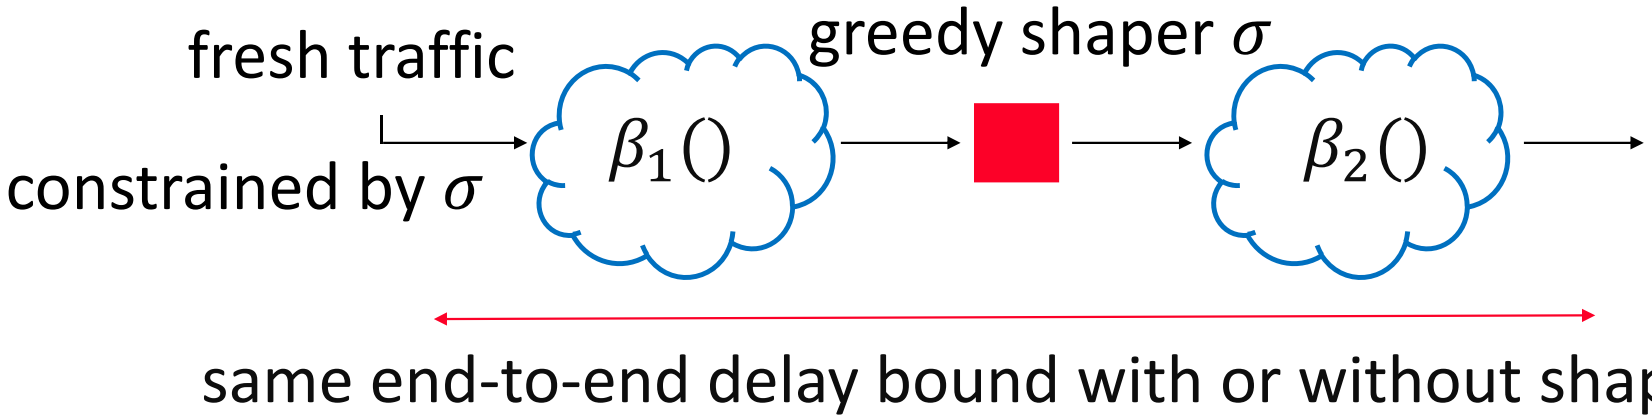

[Le Boudec Thiran 2001, Section 1.5]

Flow Regulation: LRQ

Length-Rate Quotient rule $LRQ(r)$ [Specht and Saami 2016] – used in the context of IEEE TSN

$$A_{n+1} - A_n \geq \frac{L_n}{r}$$

Chang's g-regularity

$$A_n - A_m \geq g(L_m + \cdots + L_{n-1}), \qquad m \leq n$$

$LRQ(r)$ is an instance of g-regularity with $g(\ell) = \frac{\ell}{r}$

g-regularity leads to max-plus convolution [Chang-Lin 1998]

is not equivalent to an arrival curve constraint

Flow Regulation: $TSN(\tau, K)$ Traffic Specification

At most K packets in any interval of duration τ

$\Leftrightarrow$ number of packets seen in interval of duration t is $\leq K \left\lceil \frac{t}{\tau} \right\rceil$

(staircase function)

Similar to an arrival curve

but counting in packets not in bytes

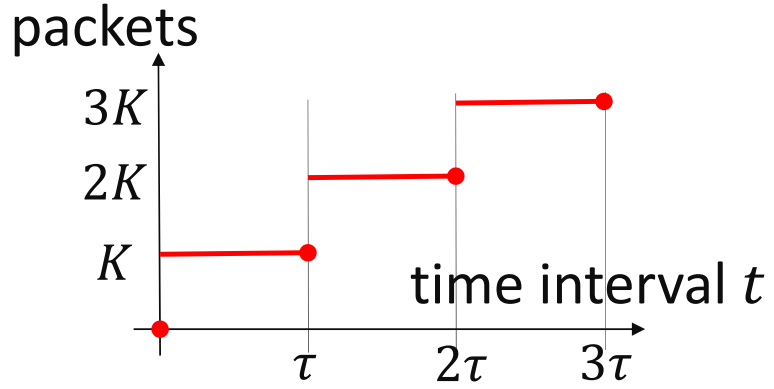

Flow Regulation: Packet Burstiness

Packet Burstiness $PB(\rho, K)$ [Le Boudec 2018] defined by:

number of packets seen in interval of duration t is $\leq \rho t + K$

Superposition property: if flow i satisfies $PB(\rho_i, K_i)$ then the superposition satisfies $PB(\sum_i \rho_i, \sum_i K_i)$.

$TSN(\tau, K) \Rightarrow PB(\rho, K)$ with $\rho = \frac{K}{\tau}$

Packet Burstiness is a tractable proxy to TSN traffic specification

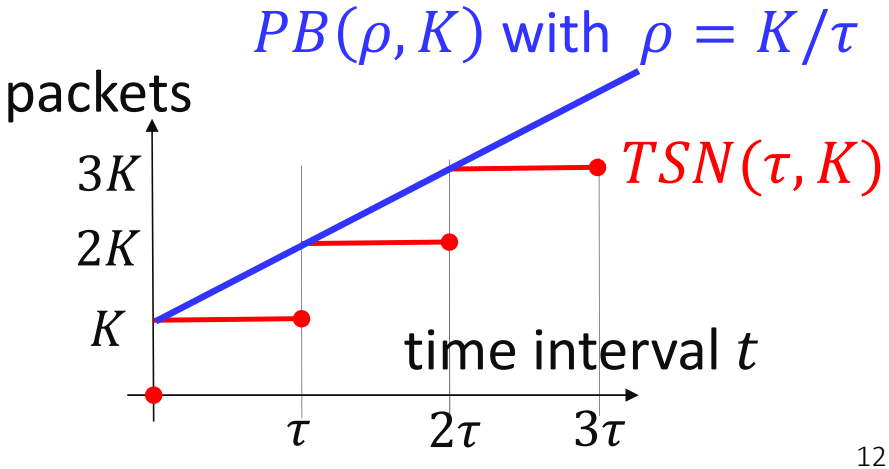

Flow Regulation: (λ, ν) Constraint

[Jiang 2018]

$A_n - A_m \geq \frac{n-m-\nu}{\lambda}$ for $m \leq n$

Similar to affine g-regularity

but counting in packets not in bytes

Min-plus versus Max-plus formulation

Based on byte/ packet counts

- affine arrival curve
- staircase arrival curve
- $TSN(\tau, K)$ (at most K packets in τ seconds)
- $PB(\rho, K)$ (at most $\rho t + K$ packets in any t seconds)

Min-plus formulation

Based on arrival times

- g-regulation
- $LRQ(\tau)$
- (λ, ν) constraint (packet based constraint)

Max-plus formulation

Recap

Many different flow constraints

- some count bytes, some count packets
- some min-plus, some max-plus

Can we have a unified theory ?

Can we understand their associated regulators ?

Do they enjoy properties similar to greedy shapers ?

3. Equivalence of Viewpoints

Theorem: [Thm 1, Le Boudec 2018]

The conditions are equivalent

1. Flow has arrival curve constraint σ
2. $A_n - A_m \geq \sigma^{\downarrow}(L_m + \cdots + L_n)$ for all $m \leq n$

where $\sigma^{\downarrow}$ is the lower pseudo-inverse of σ

Lower Pseudo-Inverse [Liebeherr 2017]

$$f^{\downarrow}(x) = \inf\{t, f(t) \geq x\}$$

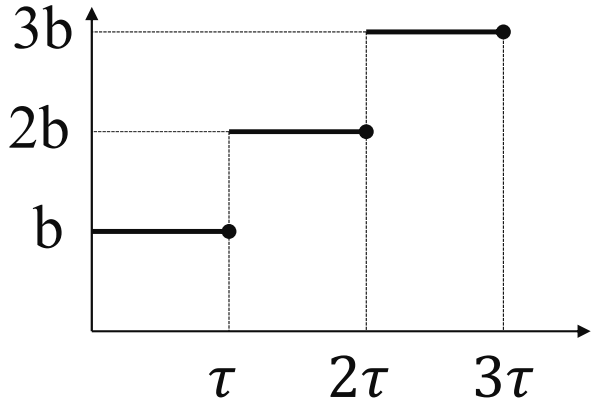

$$\sigma(t) = b\left\lceil \frac{t}{\tau} \right\rceil$$

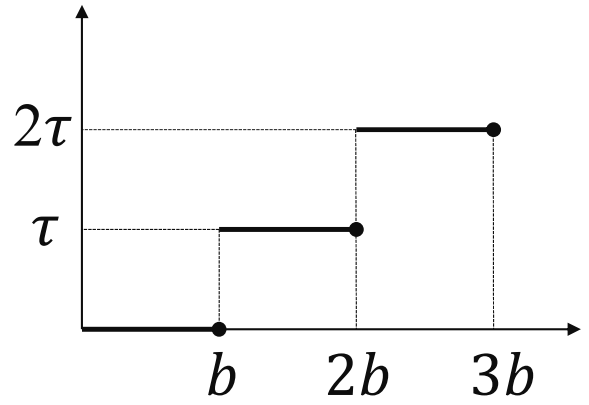

$$\sigma^{\downarrow}(x) = \tau\left\lceil \frac{x}{b} - 1 \right\rceil$$

for $x > 0$

Equivalent Formulations

	Original Definition	Equivalent Definition by Application of Theorem
affine arrival curve (leaky bucket)	$R(t) - R(s) \leq r(t-s) + b$	$A_n - A_m \geq \frac{L_m + \cdots + L_n - b}{r}$
staircase arrival curve (at most b bits in τ seconds)	$R(t) - R(s) \leq b \left\lceil \frac{t}{\tau} \right\rceil$	$A_n - A_m \geq \tau \left\lceil \frac{L_m + \cdots + L_n - b}{b} \right\rceil$

$R(t)$ = number of bits seen in $[0, t]$; A_m = arrival time for packet n

Equivalence, Packet based Constraint

Apply theorem with $L_k = 1$:

The conditions are equivalent

1. Number of packets in any interval of duration t is $\leq f(t)$
2. $A_n - A_m \geq f^{\downarrow}(n - m + 1)$ for all $m \leq n$

Equivalent Formulations

$TSN(\tau, K)$ (at most K packets in τ seconds)	$P(t) - P(s) \leq K \left\lceil \frac{t-s}{\tau} \right\rceil$	$A_n - A_m \geq \tau \left\lceil \frac{n-m+1-K}{K} \right\rceil$
$PB(\rho, K)$ (at most $\rho t + K$ packets in any t seconds)	$P(t) - P(s) \leq \rho(t-s) + K$	$A_n - A_m \geq \frac{n-m+1-K}{\rho}$
(λ, ν) constraint = $PB(\lambda, \nu+1)$	$P(t) - P(s) \leq \lambda(t-s) + \nu + 1$	$A_n - A_m \geq \frac{n-m-\nu}{\lambda}$

$P(t)$ = number of packets seen in $[0, t]$; A_m = arrival time for packet n

4. Pi-Regulation

A single packet flow (A, L)

$A = (A_1, A_2, \dots)$ packet arrival times, $A \in \mathcal{F}_{inc}$

$L = (L_1, L_2, \dots)$ packet lengths

Π a mapping $(A, L) \to \Pi(A, L) = (E_1, E_2, \dots) \in \mathcal{F}$ (eligibility times)

Definition [Le Boudec 2018]

This flow is Π-regular $\Leftrightarrow A \geq \Pi(A, L)$ i.e. $A_n \geq E_n$

We require that Π is causal ($E_n = \Pi(A, L)_n$ depends only on $A_1, \dots, A_{n-1}$ and $L_1, \dots, L_n$), homogeneous (invariant by change of time origin), and isotone (if $A \geq A'$ then $\Pi(A, L) \geq \Pi(A', L)$).

Examples of Pi-regulation

All flow constraints shown before are instances of Pi-regulation

- affine arrival curve $\Leftrightarrow A_n \geq \underbrace{\max_{m \leq n-1} \left(A_m + \frac{L_m + \cdots + L_n - b}{r} \right)}_{\Pi^{LB(r,b)}(A,L)_n}$
- $LRQ(r) \Leftrightarrow A_n \geq \underbrace{A_{n-1} + \frac{L_{n-1}}{r}}_{\Pi^{LRQ(r)}(A,L)_n}$
- $TSN(\tau, K) \quad \Leftrightarrow A_n \geq \underbrace{\max_{m \leq n-1} \left(A_m - \tau \left\lceil \frac{n-m+1-K}{K} \right\rceil \right)}_{\Pi^{TSN(\tau,K)}(A,L)_n}$

Minimal Pi-Regulator

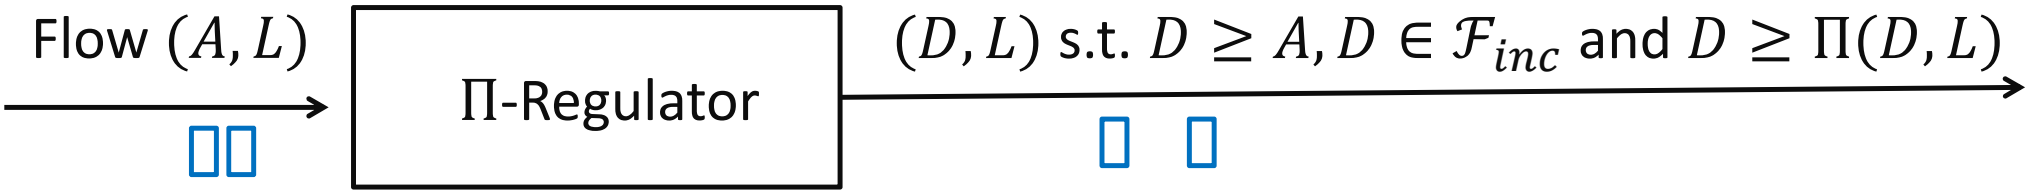

Definition: Π-Regulator for a flow is any FIFO system that transforms this flow into a Π-regular flow.

Definition: A Π-Regulator is minimal if it delivers packets no later than any other Π-Regulator.

Theorem [Le Boudec 2018]: There is one Minimal Π-Regulator; it is defined by $D_1 = A_1$ and

$$D_n = \max\{A_n, D_{n-1}, \Pi(D)_n\}$$

This the central result in this theory, as we see next !

Examples of Minimal Pi-Regulators

Flow regulation	*Minimal Regulator*
Arrival curve	Packetized greedy shaper
Chang's g-regulation	Chang's g-regulator
$TSN(\tau, K)$ (at most K packets in τ seconds)	Packet based spacer - similar to ATM spacer-controller [Guillemin et al 1992]
$PB(\rho, K)$ (at most $\rho t + K$ packets in any t seconds)	Packet based leaky bucket controller (counts only packets, not bits)

Universal Property of Minimal Pi-Regulators

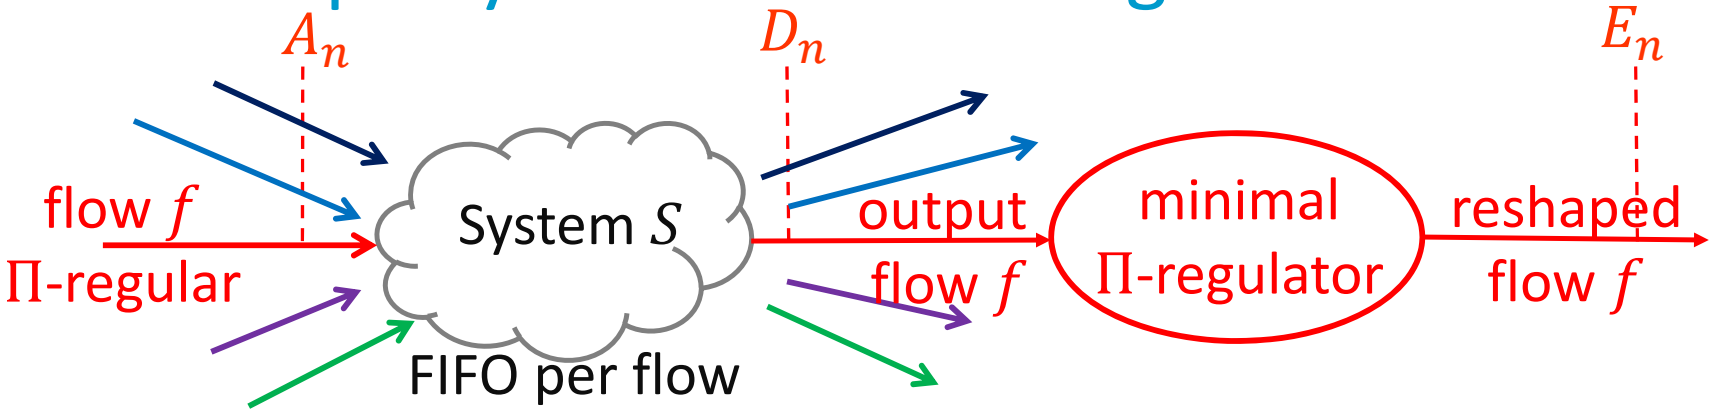

One flow f goes through a system S; system S is FIFO for flow f

- flow f is $\Pi-$ regular at input to S
- output flow f is reshaped through a minimal Π-regulator

Theorem [Le Boudec 2018]: The worst case delay of flow f is not increased: $\sup_n(E_n - A_n) = \sup_n(D_n - A_n)$

Re-shaping is for free !

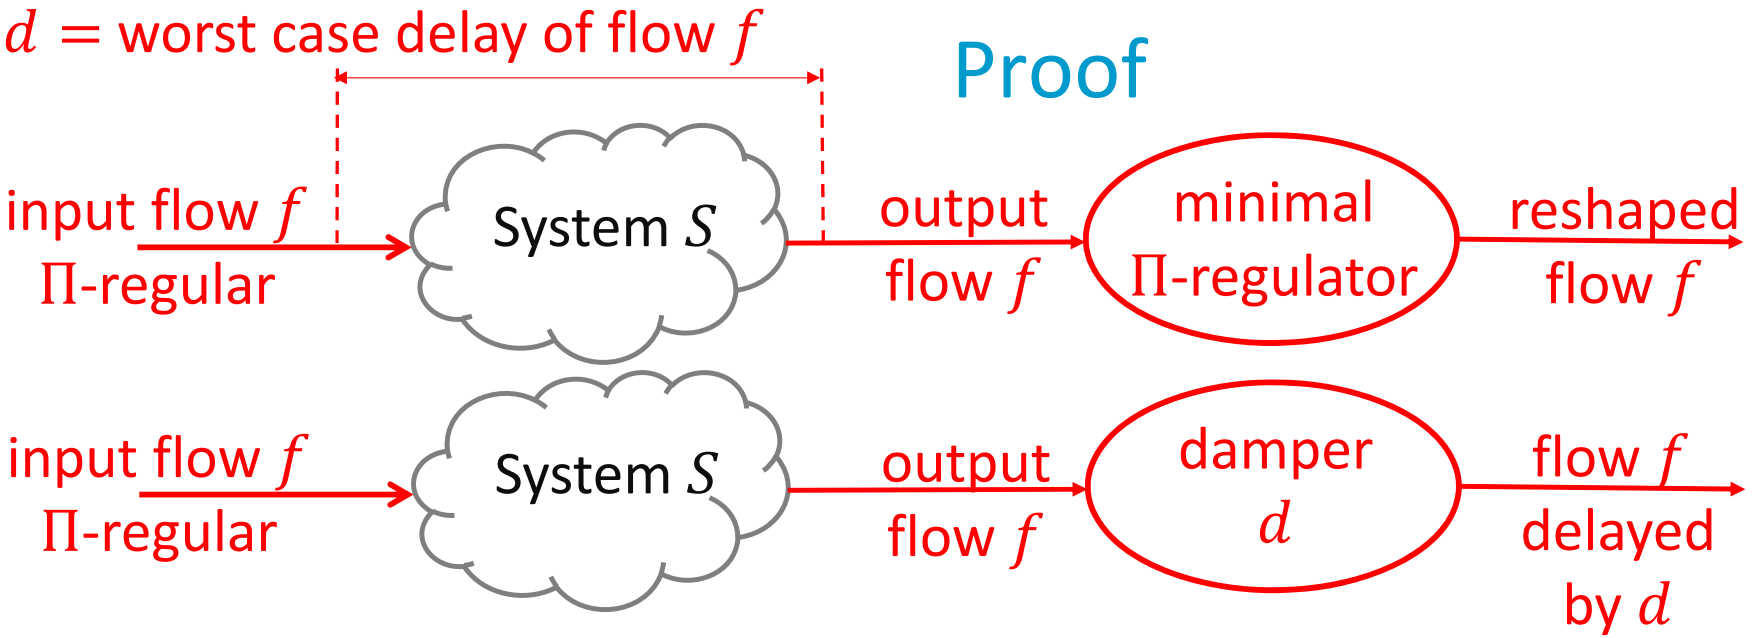

- Replace minimal Π −regulator by **damper** [Verma et al 1991]: Damper forces total delay of flow f to be exactly d; Damper is causal if d is $\geq$ worst-case delay through S
- Output of damper is input flow f time-shifted by d $\Rightarrow$ is Π −regular $\Rightarrow$ Damper is a Π −regulator $\Rightarrow$ (Minimal property) flow f delayed by d is no earlier than reshaped flow f

Packetized Greedy Shaper

Greedy shaper = minimal regulator when constraint is arrival curve and packet can be split into infinitesimal bits

Greedy shapers don't increase delay bounds given by service curve elements

Packetized greedy shaper (PGS) = minimal Pi-regulator when constraint is arrival curve σ. When σ is concave, PGS is concatenation of greedy shaper and packetizer ⇒ PGS does not increase one-hop delay bound for a FIFO service curve element.

Our new result extends this to 1) the worst-case delay of any FIFO per-flow system 2) any regulation constraint

5. Back to: Avoiding Burstiness Cascade

Solution 1: re-shape every flow at every hop (per-flow minimal regulator). Solves the problem but defeats the purpose of per-class network.

Solution 2: Interleaved Regulator

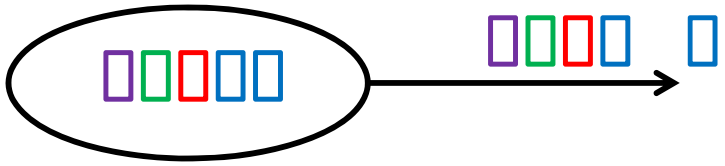

- FIFO queue of all packets of all flows in class
- packet at head of queue is examined versus traffic regulation of its flow; this packet is delayed if it came too early
- packets not at head of queue wait for their turn to come

Invented by [Specht-Samii 2016] as “Urgency Based Scheduler”, now called “Asynchronous Traffic Shaping” at IEEE TSN

Interleaved Regulator

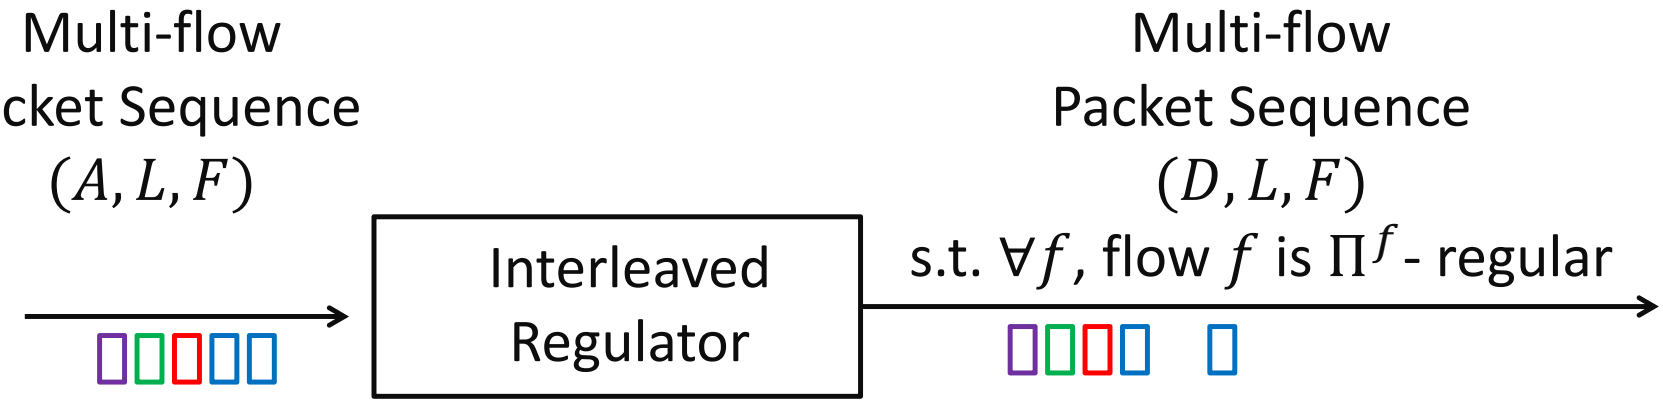

A_n: arrival time of packet n; L_n: length; F_n: flow id of packet n

An Interleaved regulator is a FIFO system such that every output flow f is Π^f - regular

$$D^f \geq \Pi^f\left(D^f, L^f\right)$$

where D^f is the subsequence of D obtained by keeping only dates that correspond to packets of flow f

Minimal Interleaved Regulator

Theorem [Le Boudec 2018]: There is one minimal interleaved regulator (i.e. such that $D_n \leq D'_n$ for any other interleaved regulator).

It is given by $D_1 = A_1$ and

$$D_n = \max\left\{A_n, D_{n-1}, \Pi^{F_n}\left(D^{F_n}, L^{F_n}\right)_{I(n)}\right\}$$

where $I(n)$ is the index of packet n in its flow.

Implementation of Minimal Interleaved Regulator

$$D_n = \max\left\{A_n, D_{n-1}, \Pi^{F_n}\left(D^{F_n}, L^{F_n}\right)_{I(n)}\right\}$$

Eligibility Time of packet at head of queue

- One FIFO queue for all packets of all flows.
- Packet at head of queue is examined and delayed until it can be released while satisfying the regulation of its flow.
- Other packets wait until their turn comes.

[Specht-Samii 2016]

Minimal Interleaved Regulator Does Not Increase Worst Case Delay

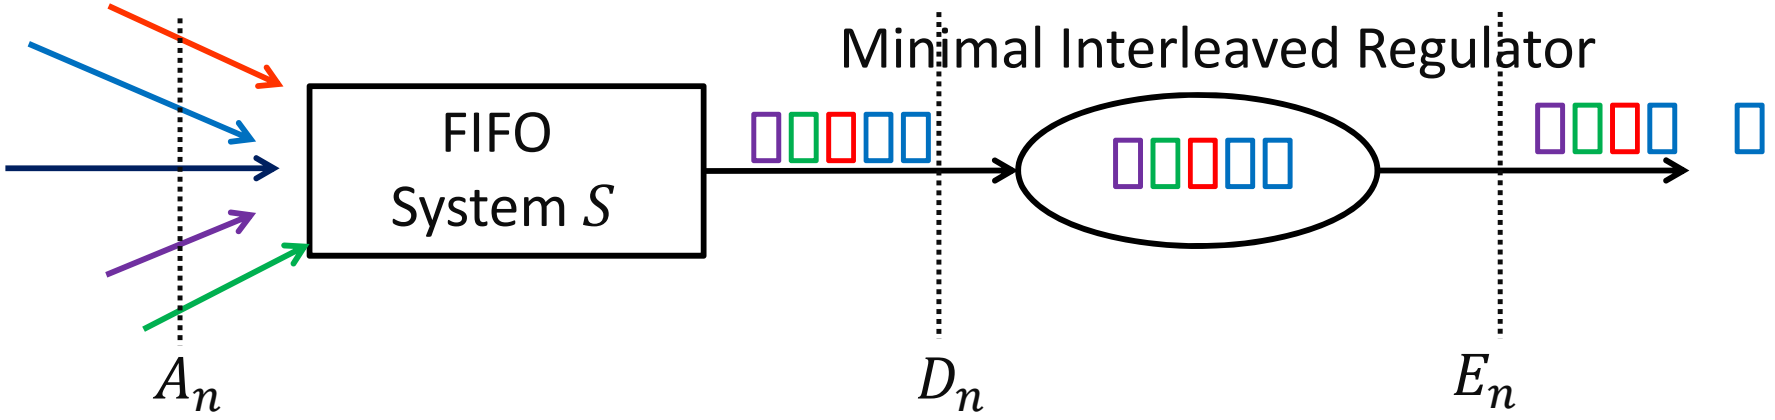

Every flow f is Π^f regular before input to S

Output of S is fed to interleaved regulator with regulator Π^f for flow f

Theorem: $\sup_n(D_n - A_n) = \sup_n(E_n - A_n)$

Interleaved Regulator is for free !

Proof

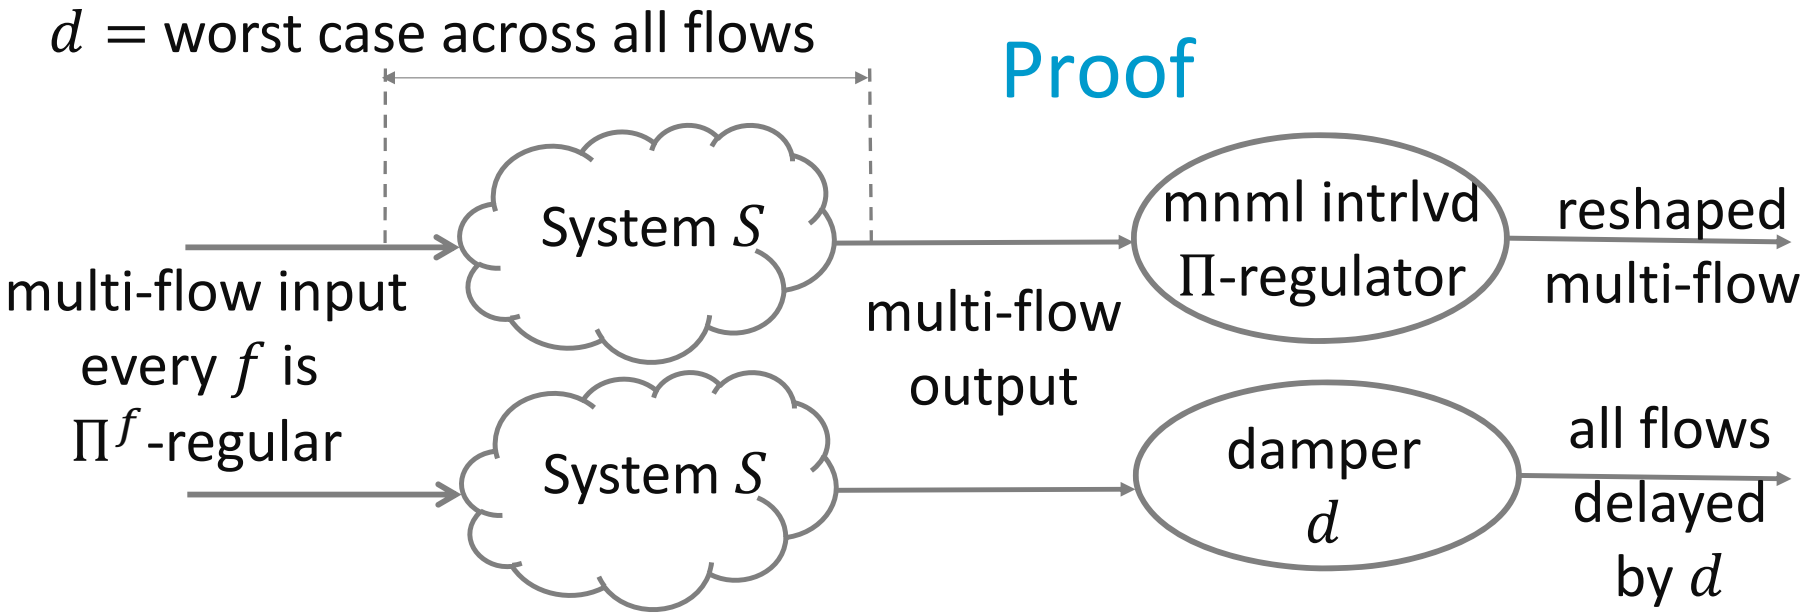

- Replace minimal Π –regulator by **damper** [Verma et al 1991]: Damper forces total delay of input to be exactly d; Damper is causal if d is $\geq$ worst-case delay through S
- Damper is an interleaved Π –regulator $\Rightarrow$ multi-flow output delayed by d is no earlier than reshaped multi-flow

Per-Flow versus Interleaved Minimal Regulation

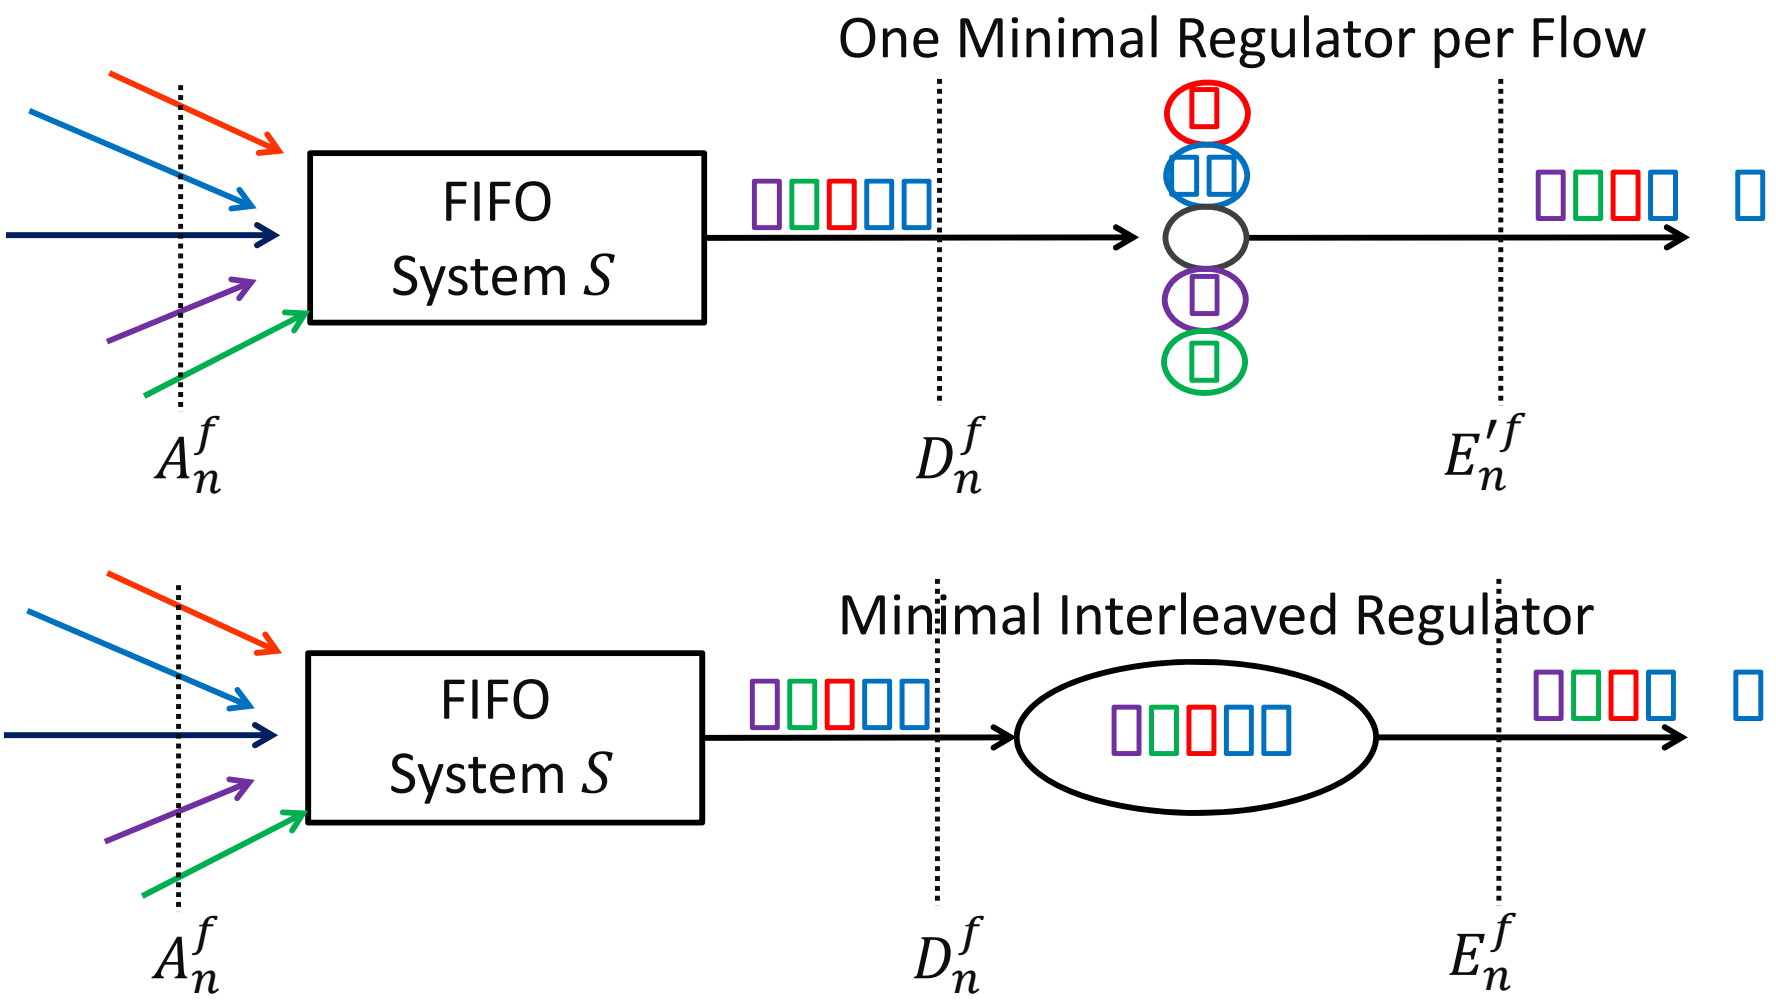

- Interleaved does not require per-flow queue
- Do they give the same delay?
  I.e., do we have $E'^{f}_{n} = E^{f}_{n}$?

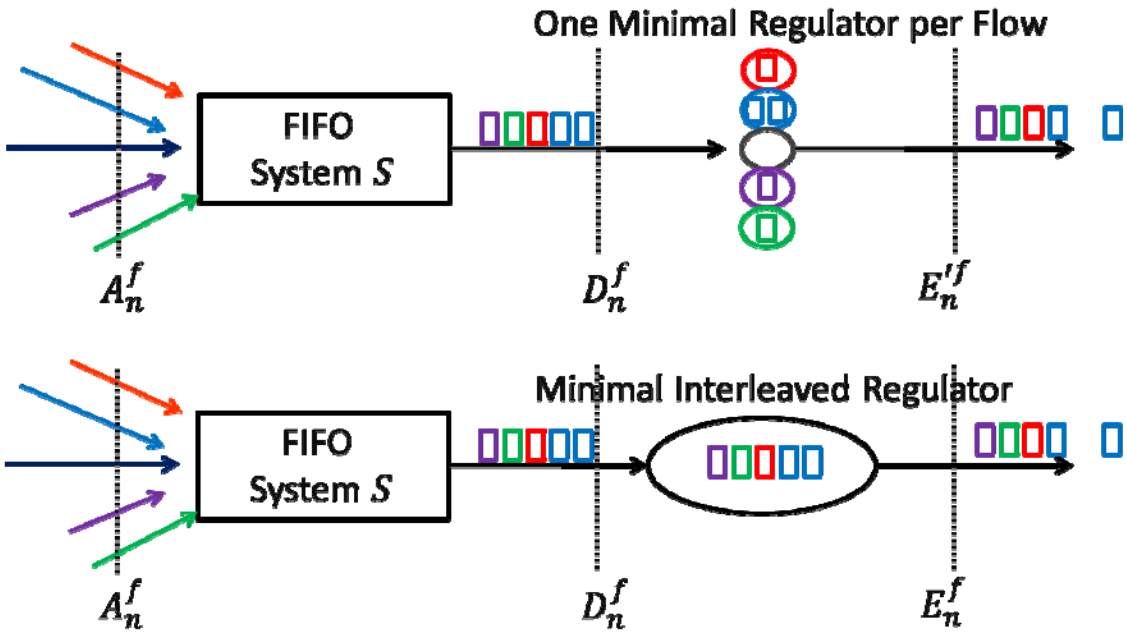

Minimal Interleaved Regulator is a Regulator for output flow f
$\Rightarrow E^{f}_{n} \geq E'^{f}_{n}$

In general, it is possible that $E^{f}_{n} > E'^{f}_{n}$ for some packet n and some flow f (i.e., interleaved regulator may delay some flows more than per-flow regulator)

Worst-case delay at S for flow f

$$d^f = \sup_n \left(D_n^f - A_n^f \right)$$

Worst-case delay at S

$$d = \sup_f d^f = \sup_{n,f} \left(D_n^f - A_n^f \right)$$

In general $d^f < d$ for some flows

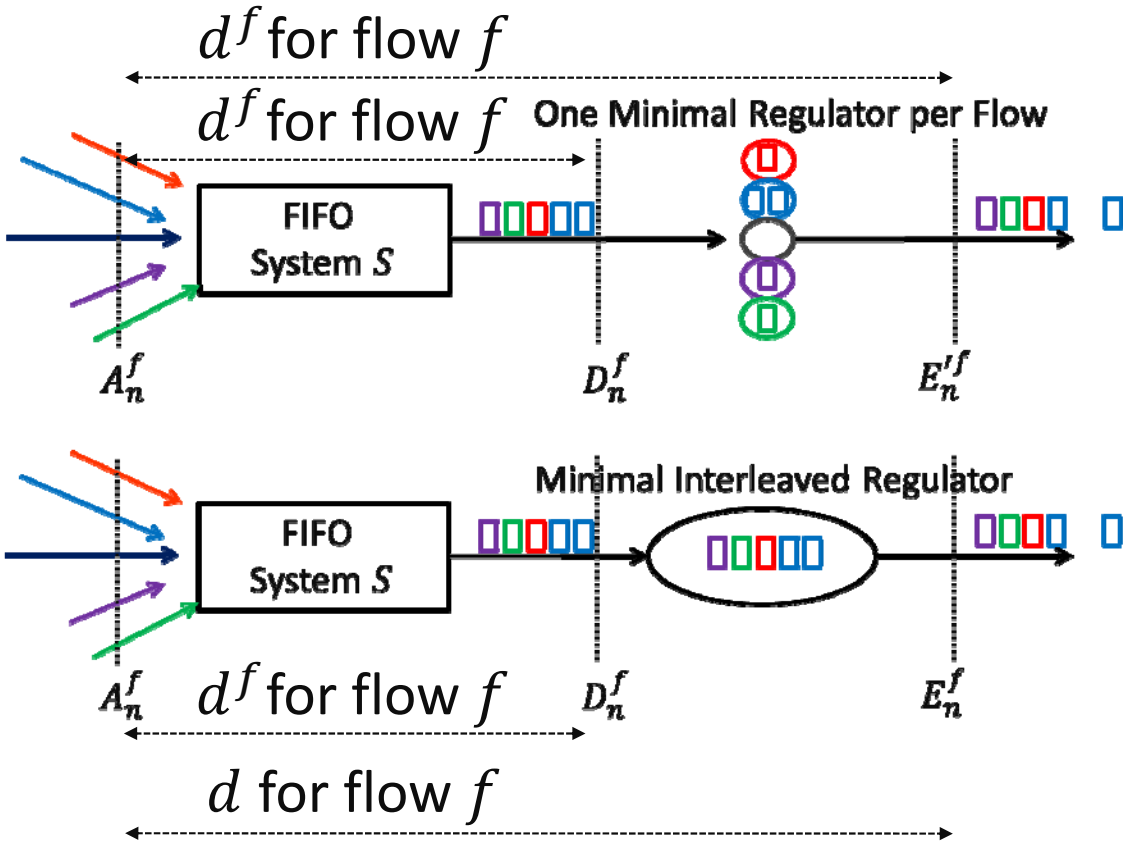

Minimal Interleaved Regulator might force delay d to a flow f that has $d^f < d$.

FIFO Network With Interleaved Regulators

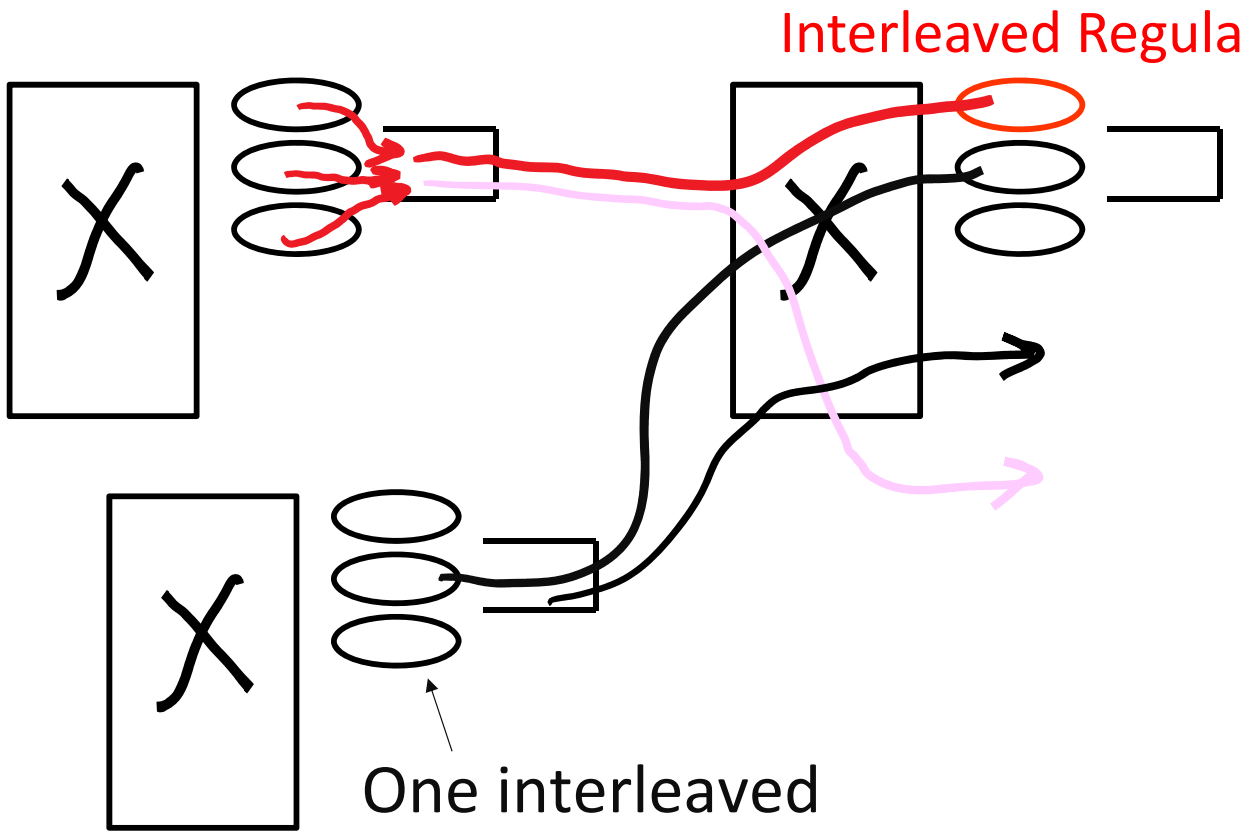

[Specht-Samii 2016] places one interleaved regulator per input port before output queue.

Output of interleaved regulator has known burstiness
⇒ no burstiness cascade

Delay Computations in IEEE TSN

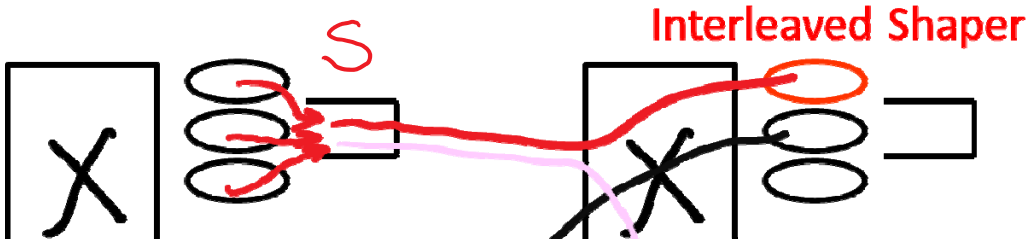

- Apply theorem on worst case delay where S = output scheduler at previous hop. Worst case end-to-end queuing delay can ignore interleaved regulators. Delay bound at one interleaved regulator is absorbed by delay at previous hop.
- Queuing delay at every scheduler S (without shaper) can be computed easily since traffic is regulated. [Next Session]
- Worst case delay at one node cannot ignore interleaved shaper.
  ⇒ Worst case end-to-end delay is generally less than sum of per-hop delays.

Conclusions

- Pi-regulation generalizes arrival curves, g-regulation, packet rate limitations.
- TSN's traffic spec uses constraint on packet rate with a staircase function. Can be replaced (for tractability) by an affine function (Packet Burstiness).
- Minimal regulator does not increase per-flow worst-case delay.
- Minimal interleaved regulator does not increase overall worst-case delay.
- Minimal Interleaved Regulators can be used to simplify and control FIFO networks.

References

[Amari et al 2016] Amari, Ahmed and Mifdaoui, Ahlem and Frances, Fabrice and Lacan, Jérôme. Worst-Case Timing Analysis of AeroRing- A Full Duplex Ethernet Ring for Safety-critical Avionics. In: 12th IEEE World Conference on Factory Communication Systems (WFCS), 3 May 2016 - 6 May 2016 (Aveiro, Portugal).

[Bennett et al 2002] Bennett, J.C., Benson, K., Charny, A., Courtney, W.F. and Le Boudec, J.Y., 2002. Delay jitter bounds and packet scale rate guarantee for expedited forwarding. *IEEE/ACM Transactions on Networking (TON)*, *10*(4), pp.529-540.

[Bondorf et al 2017], Bondorf, Steffen, Paul Nikolaus, and Jens B. Schmitt. "Quality and Cost of Deterministic Network Calculus: Design and Evaluation of an Accurate and Fast Analysis." *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1.1 (2017) and arXiv:1603.02094v3

[Bouillard et al 2018] Bouillard, A. Boyer, M. and Le Coronc, E. *Deterministic Network Calculus: from Theory to Practical Implementation,* ISTE Editions, to appear in 2018

[Bouillard-Stea 2015] Bouillard, A. and Stea, G., 2015. Exact worst-case delay in FIFO-multiplexing feed-forward networks. *IEEE/ACM Transactions on Networking (TON)*, *23*(5), pp.1387-1400.

[Boyer et al 2012] Boyer, Marc, Nicolas Navet, and Marc Fumey. "Experimental assessment of timing verification techniques for AFDX." *6th European Congress on Embedded Real Time Software and Systems*. 2012.

[Chang-Lin] Chang, C.S. and Lin, Y.H., 1998, May. A general framework for deterministic service guarantees in telecommunication networks with variable length packets. In *Quality of Service, 1998.(IWQoS 98) 1998 Sixth International Workshop on* (pp. 49-58). IEEE.

[Chang 2002] Chang, C.S., 2012. *Performance guarantees in communication networks*. Springer Science & Business Media.

[Guillemin et al 1992] Fabrice P. Guillemin, Pierre E. Boyer, Michel J. Servel. The spacer-controller: an efficient UPC/NPC for ATM networks. ISS '92, Session A9.3, volume 2, Oct 1992.

[Le Boudec 2018] Le Boudec, Jean-Yves, "A Theory of Traffic Regulators for Deterministic Networks with Application to Interleaved Regulators", *arXiv preprint arXiv:1801.08477*.

[Le Boudec-Thiran 2001] Le Boudec, Jean-Yves, and Patrick Thiran. *Network calculus: a theory of deterministic queuing systems for the internet*. Vol. 2050. Springer Science & Business Media, 2001, online at http://ica1www.epfl.ch/PS_files/NetCal.htm

[Liebeherr 2017] Liebeherr, J., 2017. Duality of the Max-Plus and Min-Plus Network Calculus. *Foundations and Trends® in Networking*, *11*(3-4), pp.139-282.

[Rizzo-Le Boudec 2008] Rizzo, G. and Le Boudec, J.Y., Stability and delay bounds in heterogeneous networks of aggregate schedulers. In INFOCOM 2008. The 27th Conference on Computer Communications. IEEE (pp. 1490-1498). IEEE.

[Specht-Samii 2016] Specht, J. and Samii, S., 2016, July. Urgency-based scheduler for time-sensitive switched ethernet networks. In *Real-Time Systems (ECRTS), 2016 28th Euromicro Conference on* (pp. 75-85). IEEE.

[Verma et al 1991] D. Verma, H. Zhang, and D. Ferrari. Guaranteeing delay jitter bounds in packet switching networks. Tricomm '91, Chapel Hill, pages 35–46, April 1991.