Ethernet & IP @ Automotive Technology Week

3 Nov 2021 - 4 Nov 2021

IEEE SA STANDARDS ASSOCIATION

Arbitrating the fight between 802.1Q TSN shapers

Max Turner[(1)] , Anna Engelmann[(2)], and Jean Walrand[(3)]

(1) ETHERNOVIA TRANSFORMING HOW CARS OF THE FUTURE ARE BUILT

(2) CARIAD A VOLKSWAGEN GROUP COMPANY

(3) Berkeley UNIVERSITY OF CALIFORNIA

Agenda

- Problem Statement: Why do Shaping? Why TSN?
- Introduction to the Cruz Network Calculus Model[1]
- The Credit Based Shaper (CBS)[IEEE Std. 802.1Qav]
- The Asynchronous Traffic Shaper (ATS)[IEEE Std. 802.1Qcr]
- Comparing CBS with ATS[IEEE Std. 802.1Q-rev]
- The Role of Priority[IEEE Std. 802.1Q]
- Conclusions

[1] *R.L. Cruz, "A Calculus for Network Delay", Part I&II, IEEE Transactions on Information Theory, Vol. 37, No. 1, Jan. 1991*

Problem Statement

Why do Shaping? Why TSN?

IEEE 802.1 Time Sensitive Networking

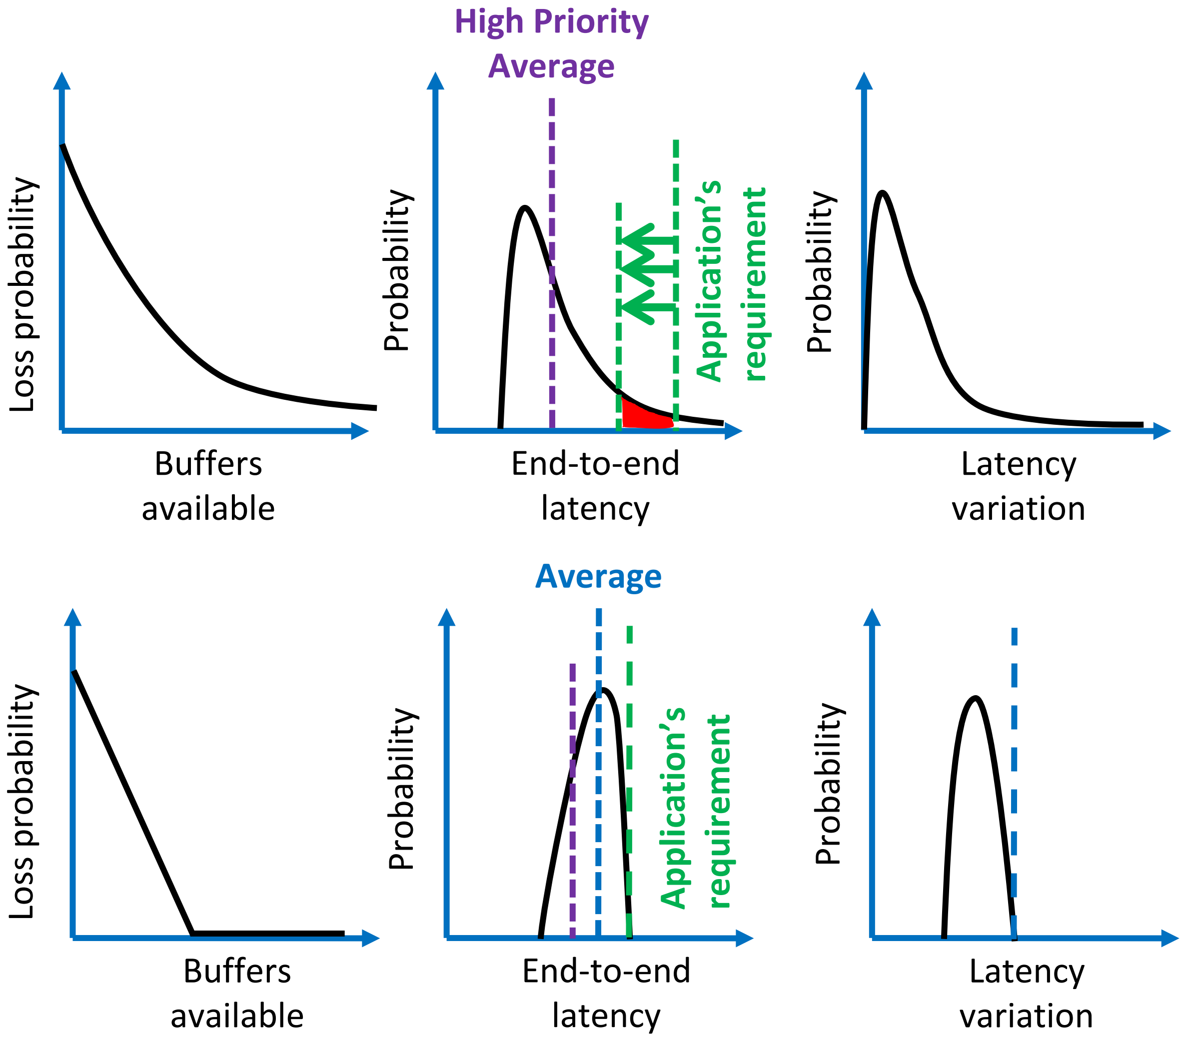

graphics from: https://www.ieee802.org/1/files/public/docs2018/detnet-tsn-farkas-tsn-basic-concepts-1118-v01.pdf
thanks to János Farkas of Ericsson Research

Without TSN:

- Many or even most frames arrive with a very low latency
- Some frames may suffer a very high latency or get dropped due to buffer overrun

Goals of TSN:

- All frames arrive within a pre-defined maximum, but higher latency
- Buffer occupancy is predictable and therefore no frames get dropped due to buffer overrun
- **How and for which flows?**

Goals of Shaping

- Shaping reduces the '**burstiness**' of flows by introducing gaps between frames, as this allows a more equal chance to access the bandwidth
- Shaping thereby **delays** some flows, but reduces the **latency** of others
- Re-transmissions due to frame loss make bandwidth predictions almost impossible, one therefore needs to prevent **buffer** overruns in bridges (congestion loss)
- Shaping mechanisms differ in:
  - Effectiveness at reducing burstiness of flows
  - Latency they introduce and buffer they require
  - Complexity of configuration
  - Sensitivity to characteristics of ingress flows
- **Goal of this analysis: Meet latency targets for ALL flows in the network using one or few simple mechanisms from the TSN toolbox and prevent congestion loss**

The Cruz Network Calculus Model

R.L. Cruz, "A Calculus for Network Delay", Parts I&II, IEEE Transactions on Information Theory, Vol. 37, No. 1, Jan. 1991

What is Burstiness?

- **Intuitively:** uneven spacing of frames

same average rate *a*
for both flows

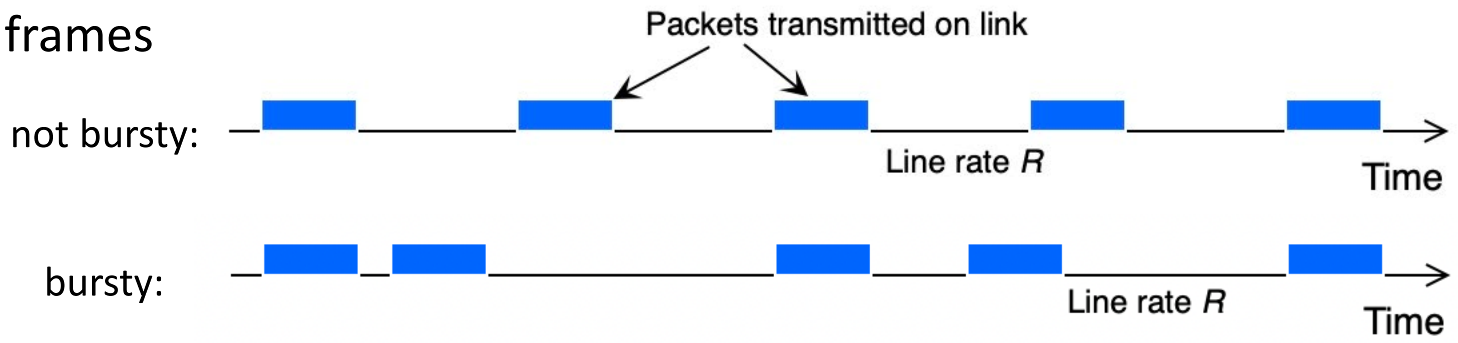

- **Mathematically**: burstiness *A* is the max. backlog in a buffer that serves the flow at its average rate a: If at most *A + at* bits arrive in any t seconds[1], we say that the flow type is (A, a)

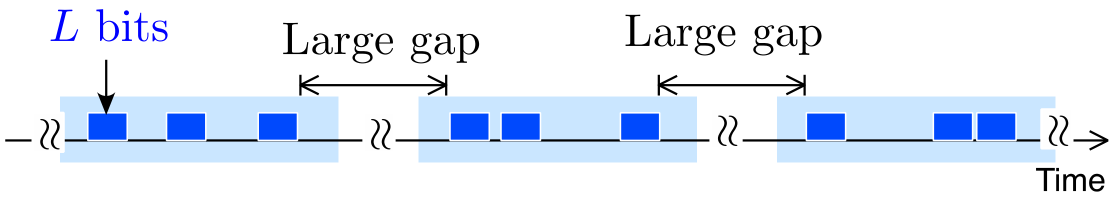

burstiness A = *3L*, max. back-to-back bits = *L*

Note: burstiness A is not the largest number of back-to-back bits!

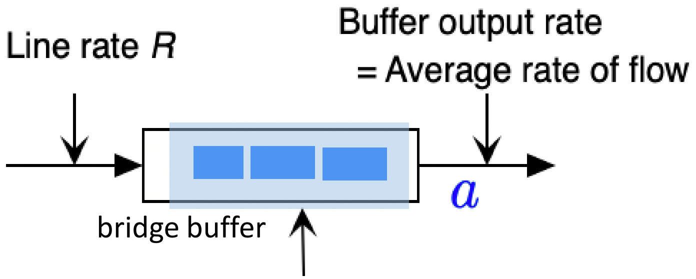

burstiness A = max. backlog

[1] This is a linear upper bound on the 'arrival curve' of the flow

Egress Burstiness on a Shared FIFO Port

Intuition:

During B/R seconds, the egress port sends B red bits, while the blue flow accumulates aB/R bits, then the egress port sends $A+aB/R$ blue bits.

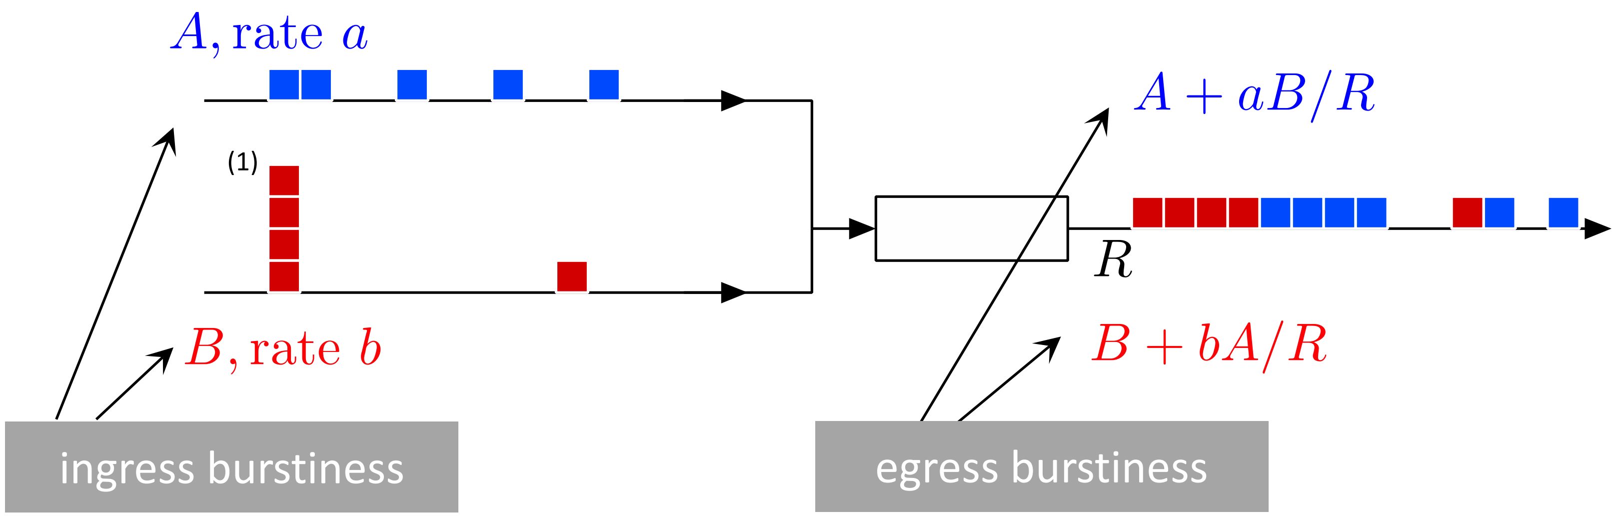

To get **predictable egress burstiness**, one needs **predictable ingress burstiness** - this requires ingress **policing** and **shaping at the talkers**

Sharing a port (FIFO) w/o shaping or priority increases the burstiness of each stream Using TSN shaping instead can prevent this.

(1) Each red frame accumulating to B, arrives simultaneously on a different ingress port

The Credit Based Shaper (CBS)

[IEEE Std. 802.1Qav]

Definition of the CBS

CBS with configurable egress rate *a* is designed to give **minimum burstiness** while guaranteeing an **upper bound on latency**

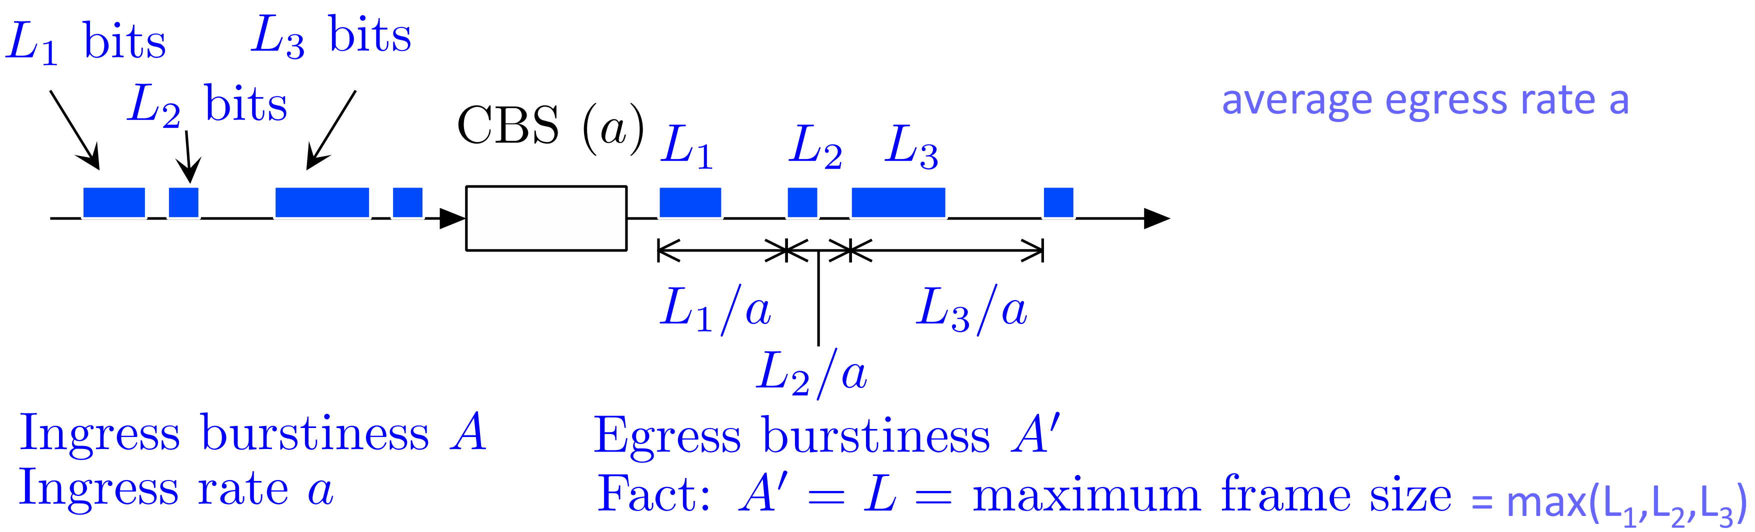

CBS separates the frames as much as possible given the configured average rate *a*, thereby reducing the flow's burstiness (L<A)

CBS with Multiple Ingress Flows

Flow *(A, a)*: Burstiness *A*, Average rate *a*

Flow *(B, b)*: Burstiness *B*, Average rate *b*

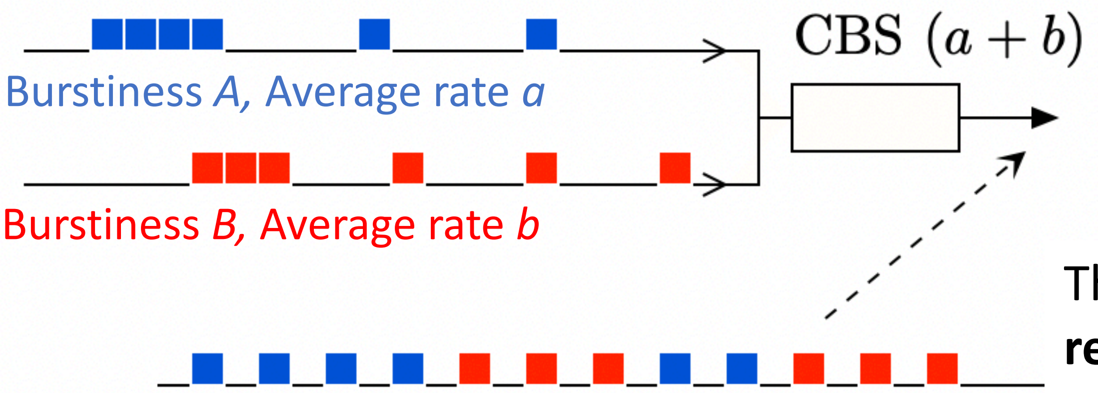

Aggregate flow is *(L, a + b)*

The **burstiness of the aggregate flow is reduced**, but that of the **sub-flows** (red and blue) is **not**

- If the flows **remain together downstream**, only the **aggregate burstiness** matters, not that of the sub-flows
- **Shape together flows that remain together** (shape by destination)

M flows with destination *1*

N flows with destination *2*

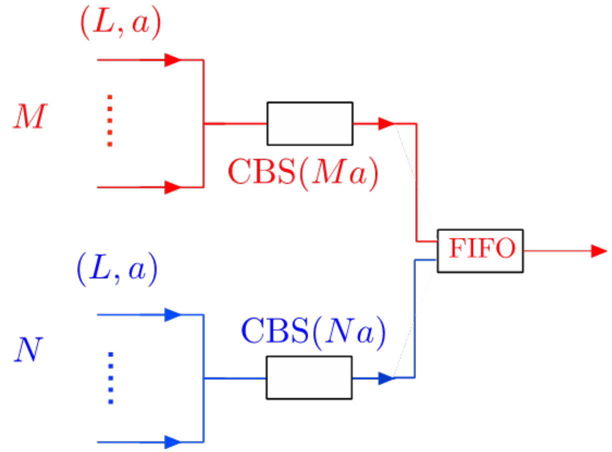

The Asynchronous Traffic Shaper (ATS)

[IEEE Std. 802.1Qcr]

Definition of the ATS

ATS implements **per-flow leaky-bucket shapers** with max. credit *B* and average rate *a*.

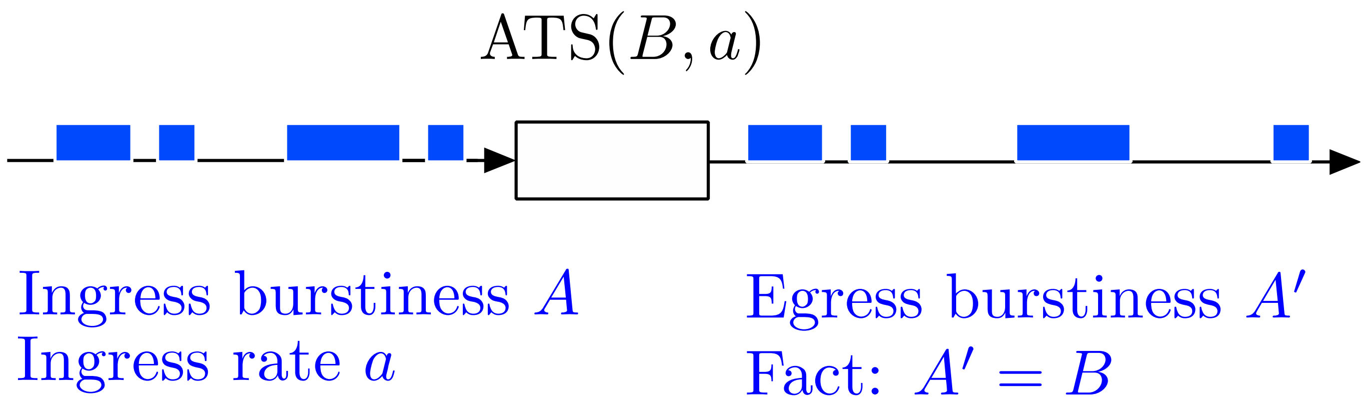

The **burstiness** ***A'=B*** **of the ATS shaper is a configurable parameter**, it can be adjusted to be more bursty to achieve a lower shaping delay for this flow, but will thereby increase the latency of other flows

ATS with Multiple Ingress Flows

The leaky bucket shapers reduces the egress burstiness of the sub-flows

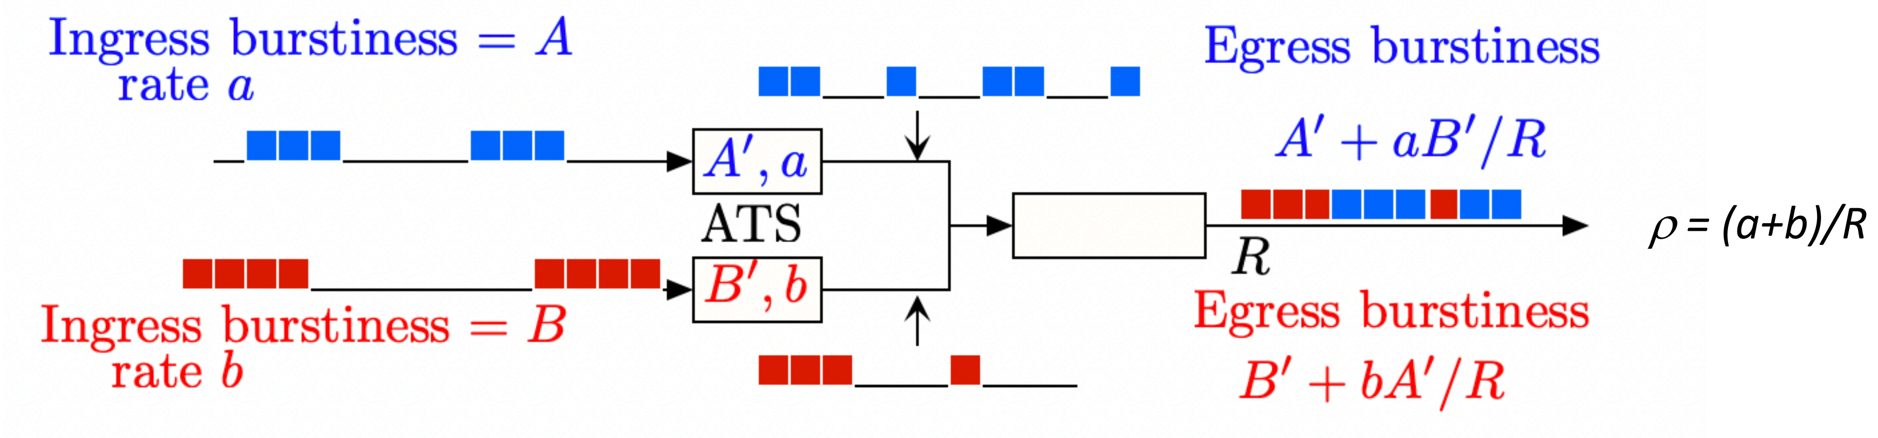

For a network with many (*M >> 1)* bursty (A) flows (*frame length L < A) and a high link utilisation ($\rho = Ma/R$), one gets the least latency and burstiness by* ***configuring A'=L***

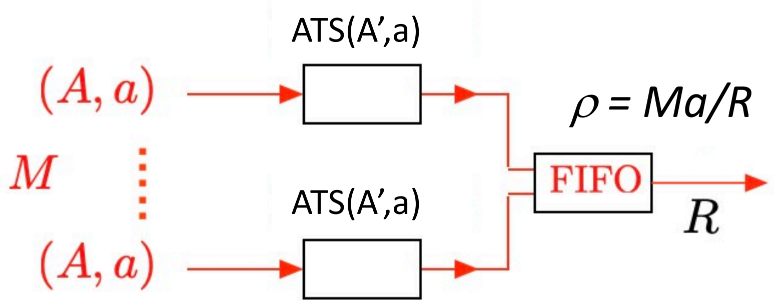

Then **ATS shows the best effects in networks of many hops with high link utilisation[1] and many flows**

[1] i.e. $\rho > 70\%\ R$

Comparing CBS with ATS

[IEEE P802.1Q-rev]

Effect on Sub-Flows and Effort

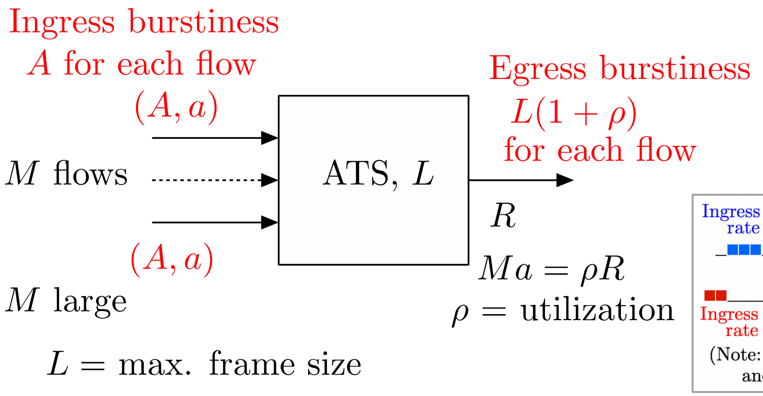

smooth sub-flows

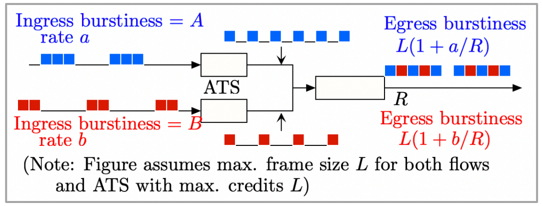

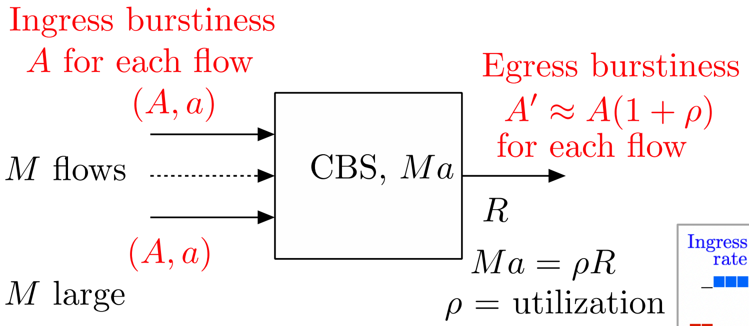

bursty sub-flows

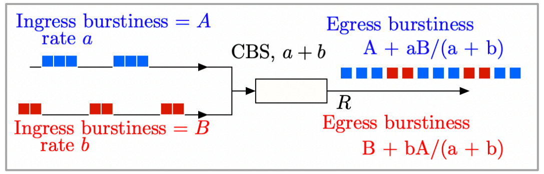

- For a single flow and an ATS burst size equal to the frame size, ATS is very similar to CBS
- For many ingress flows, **ATS reduces the burstiness of the sub-flows, CBS does not**
- ATS requires additional HW support for queue management, while CBS is limited by the number of traffic class egress queues
- ATS requires per stream configuration (IEEE Std. 802.1Qci), CBS is configured per class using priorities
- **ATS can improve highly utilised networks but increases configuration complexity and HW effort**

CBS or ATS for Split Flows

The reduction of sub-flows burstiness by ATS does not matter if the sub-flows stay together as they traverse the network

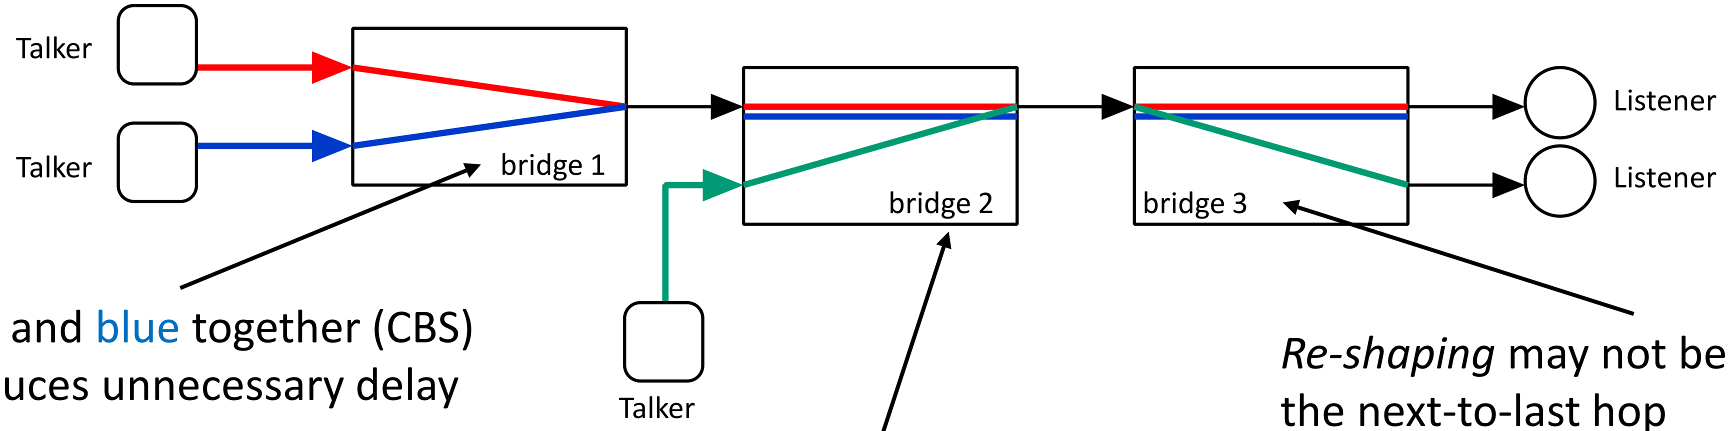

Shape red and blue together (CBS)
ATS introduces unnecessary delay

Re-shaping may not be needed in the next-to-last hop

*R*e-shape red & blue together and green separately - use ATS per flow or CBS per class depending on network size

Shape together flows that remain together!

The Role of Priority

[IEEE Std. 802.1Q]

Priority’s Effect on Burstiness & Latency

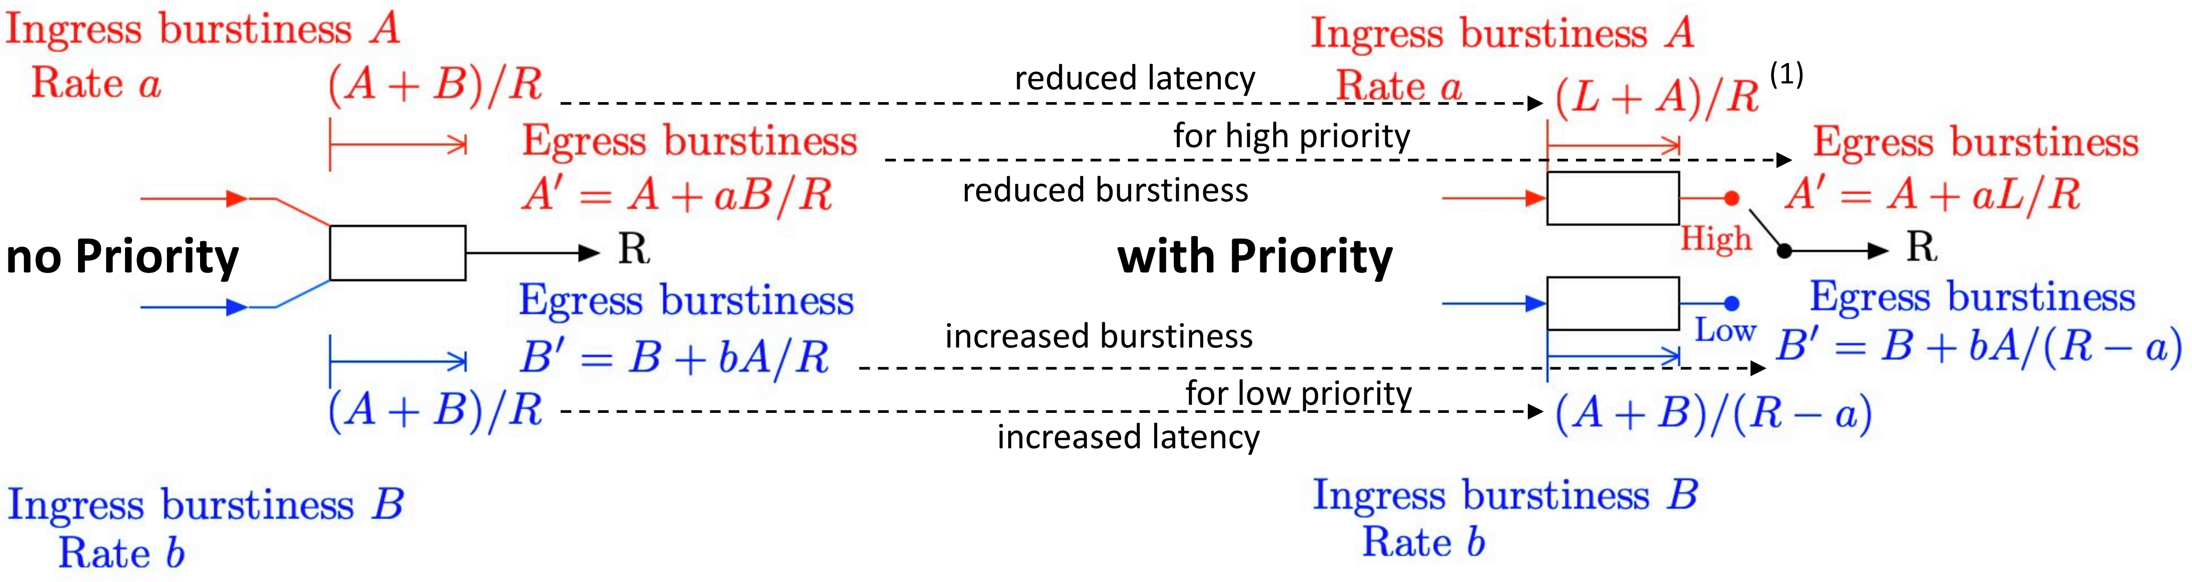

- Priority scheduling **limits the latency and burstiness of the high-priority flow**
- Priority scheduling **increases the latency and burstiness of the low-priority flow**(s)
- In IT & telecom networks low priority traffic may be dropped and re-transmitted
- For the automotive restricted resource network, TSN must aim to **prevent congestion loss for all flows**, limiting the applicability of priority

(1) L is the maximum frame size for the blue low priority flow(s), therefore: $L \le B$

Combining CBS and Priority

Latency of a red flow

$A/a + L/R$ with priority
same without priority

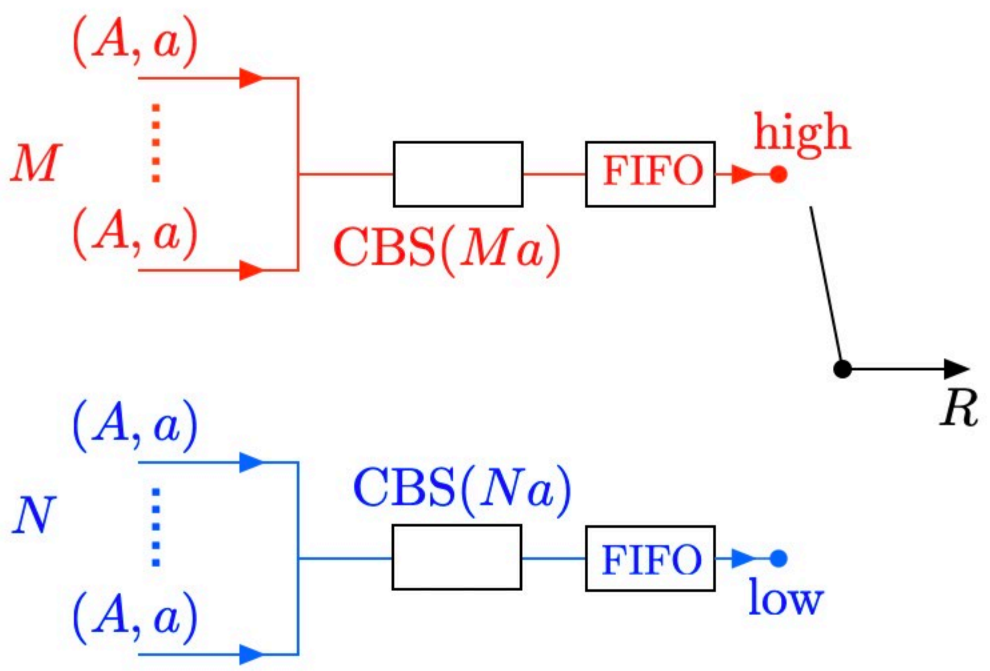

If all queues are CBS shaped, priority is not useful, as the gaps between frames guarantee that no frame has to wait for a burst from the other queue

When another unshaped queue is added, higher priority for the CBS shaped traffic is required

The gaps between CBS shaped frames still guarantee that the unshaped queue will be able to access the shared bandwidth, reducing the average latency for all flows

Combining ATS and Priority

Latency of a red flow

$(A - L)/a + L/R$ w. priority

$(A - L)/a + NL/R$ without

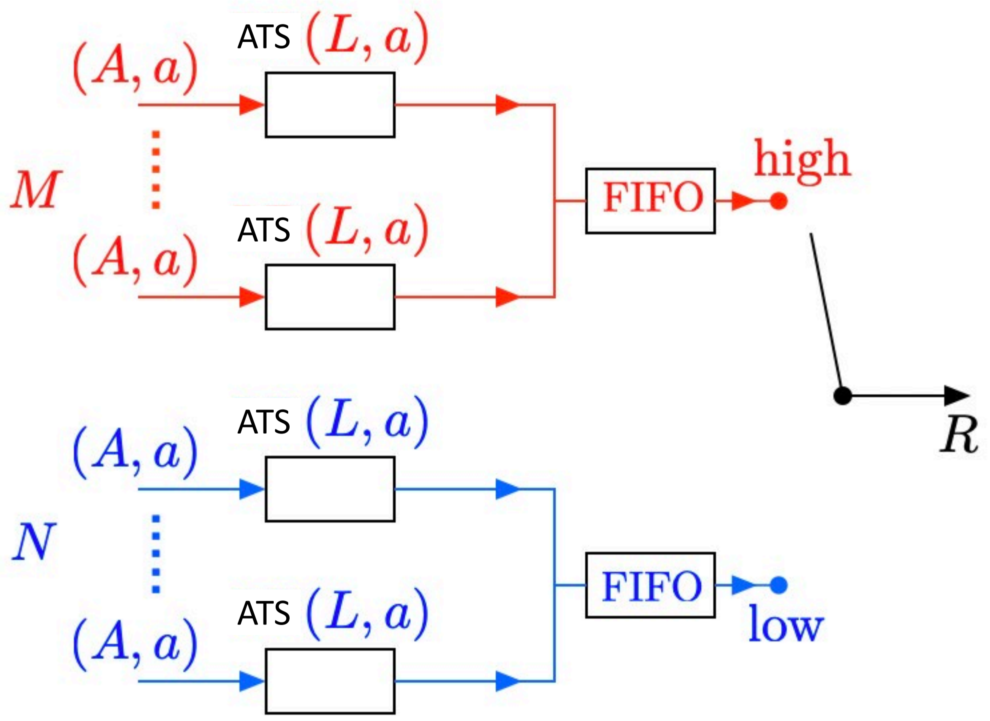

If all queues are ATS shaped, priority is still useful if there are many flows, to limit waiting for a burst of blue bits

When another unshaped queue is added, higher priority for the ATS shaped traffic is required

The gaps between ATS shaped frames still guarantee that the unshaped queue will be able to access the shared bandwidth, reducing the average latency for the ATS shaped flows at the cost of increasing latency for the low-priority flows

Combinations of ATS, CBS and Priority

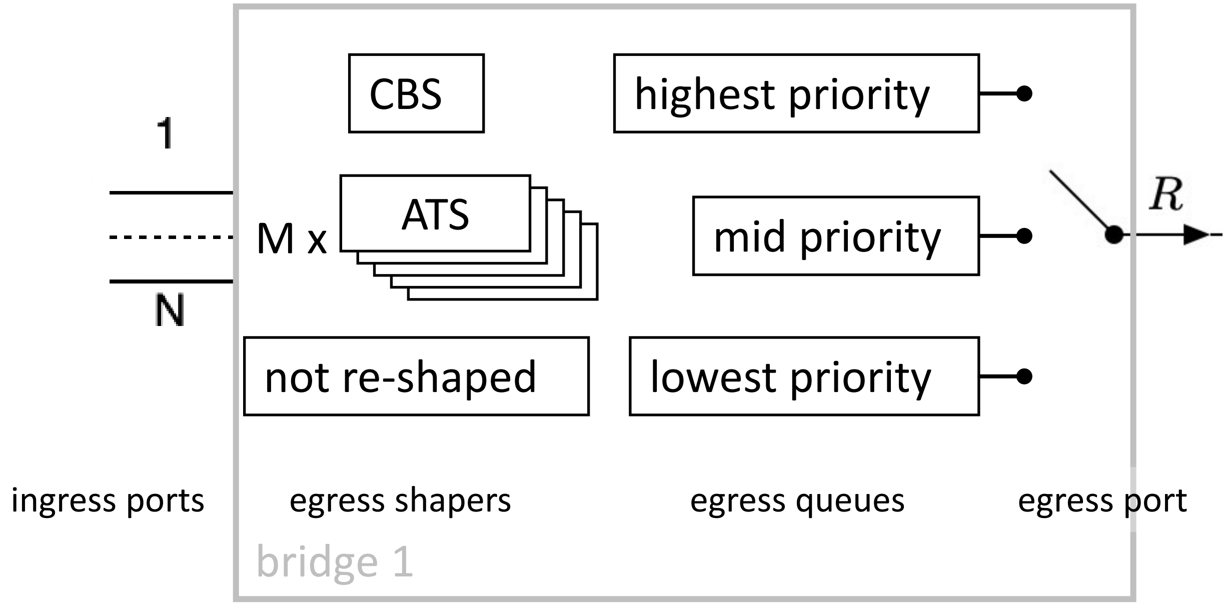

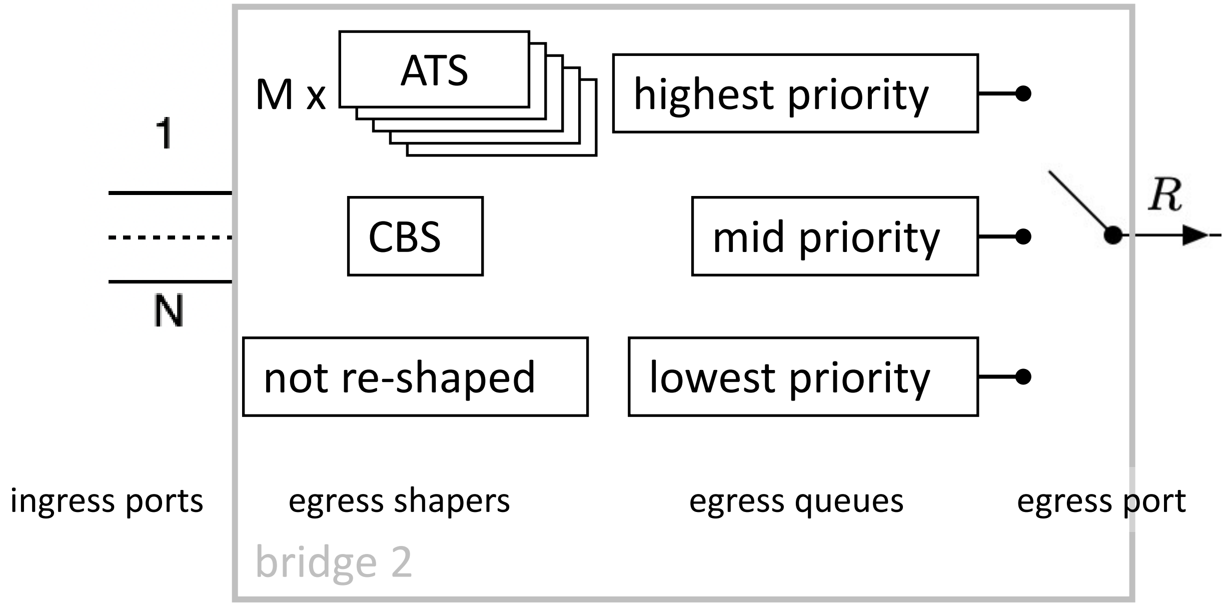

If there are very few high-priority flows the highest priority can be a CBS queue
Mid level priorities can use ATS or CBS, depending on how many flows are expected
The lowest priority need not be shaped in the bridges, but only in the talkers to prevent congestion loss (prevent retransmission!)

For many high-priority flows, it makes sense to use ATS on the highest priority queue
Mid level priorities should use at least CBS in order to create gaps for the lowest priority to transmit
The lowest priority need not be shaped in the bridges, but only in the talkers to prevent congestion loss (prevent retransmission!)

Conclusions

1. Shape flows at the talkers to limit interference, even for intermittent talkers.
2. Shape together flows that stay together: CBS shaping per destination, ATS for high flow count. Per-flow shaping is useless if flows stay together.
3. Use priority if methods 1 and 2 do not suffice to meet latency requirements: increase the latency of non-urgent flows for the benefit of critical flows.
4. Low priority enables best-effort flows to use the bandwidth not used by shaped traffic (e.g., TCP file transfers between distributed or redundant processors, with or without ECN[(1)]).
5. Analysis provides guaranteed upper bounds on storage and latency and provides a tool for optimizing the network design and configuration.

[(1)] Explicit Congestion Notification [IETF RFC 3168]

Thank you!

ETHERNOVIA
TRANSFORMING HOW CARS OF THE FUTURE ARE BUILT

CARIAD
A VOLKSWAGEN GROUP COMPANY

Berkeley
UNIVERSITY OF CALIFORNIA

max.turner@ethernovia.com anna.engelmann@cariad.technology walrand@berkeley.edu