TSN timing QoS mechanisms: what did we learn over the past 10 years?

Nicolas Navet
University of Luxembourg & Cognifyer

Outline

1. **Basics of TSN & TSN timing QoS**
2. **Main QoS policies: recap, use-cases, pros/cons, configuration issues**
    - Priorities
    - Preemption
    - Shaping policies : CBS and ATS
    - Time-triggered communications with TAS
3. **Conclusion & a look forward**

Caveat:

- This presentation is geared toward TSN in automotive, aerospace and industrial domains and may not be relevant for other application domains (data centers, telecom)
- Biased towards our own background & experience, other people in the TSN community will have different views

Time-Sensitive Networking in a nutshell

- TSN is the name of a “Task Group” that develops standards within IEEE 802.1: base standards – with 802.1Q being the main one, amendments eventually integrated in base standards and domain-specific profiles
- *“Time-Sensitive Networking (TSN) is a collection of features in IEEE 802.1 standards that provide the following”:*
  - Time synchronization among bridges and end stations [AS-2020]
  - Reduction in frame loss due to faults in network equipment or transmission [CB, Qci]
  - Bounded latencies and jitters, elimination of frame loss due to egress port congestion [802.1p, Qav, Qbv, Qcr, Qch, Qci]
- It is more than that! stream reservation, configuration mechanisms, network description models (Yang), …

TSN mechanisms offer a lot of flexibility, they

- **do not necessarily need to be implemented in HW**
- **can be used with non-TSN mechanisms**
- **can be applied only a subset of the network devices**
- **can be used with different parameters along a path**

TSN promise: mixed-criticality traffic on the same wire

- Toolbox of QoS mechanisms to address different timing requirements ☑
- But substantial network engineering needed for their selection and configuration!
- Additional complexity with fault-operational requirements: frame replication on disjoint paths, traffic policing, redundant master clocks, ..

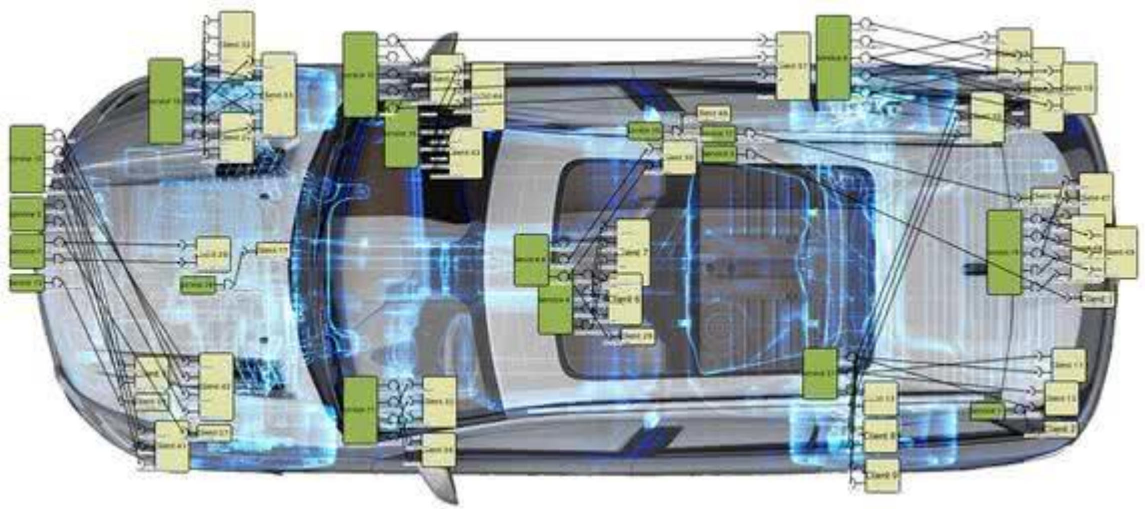

[RTaW-Pegase screenshot]

TSN is often used in complex systems such as vehicles that may have thousands of mixed-criticality functions/services communicating with each other, several displays, 10+ cameras, fault-operational requirements, …

Basics of TSN timing QoS

The TSN QoS machinery

Meeting Timing QoS = managing interfering traffic

Traffic Classes and Priorities

- Streams are grouped by users into max. 8 Traffic Classes (TC) e.g. based on functional domain, timing/safety/security constraints, …
- There is a priority associated to each traffic class (0 to 7, with 7 the highest)
- Each frame holds its priority, but network devices can change it locally

- At an egress port, the packets of a TC waiting for transmission are all stored in the same queue
- TSA: Transmission Selection Algorithm

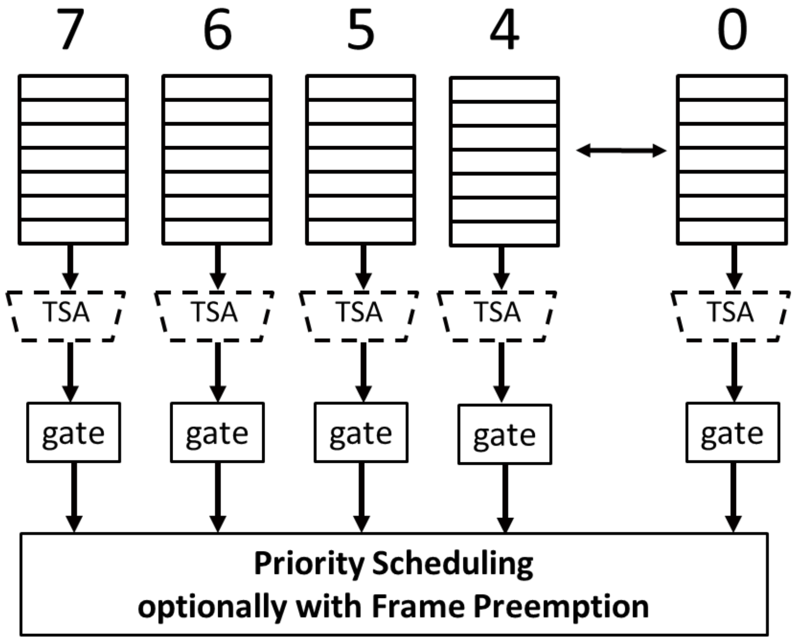

- HW implementation might differ

- In theory, there might be less queues than priority levels, but we have not seen that yet

Transmission Selection Algorithm (TSA)

- A “Transmission selection algorithm” (TSA) is associated to each queue of a port. It decides the frame from this queue - if any - that can be transmitted
- TSA can be either “Strict priority”, “Credit-based shaper”, “Enhanced Transmission Selection”, “ATS Transmission Selection”, or some vendor specific policy
- Each queue is controlled by a gate which is open or closed at any given time
- A priority scheduler selects for transmission the frame of the highest-priority traffic class that has a frame ready for transmission and whose gate is open

The QoS mechanisms are applied in a hierarchical manner

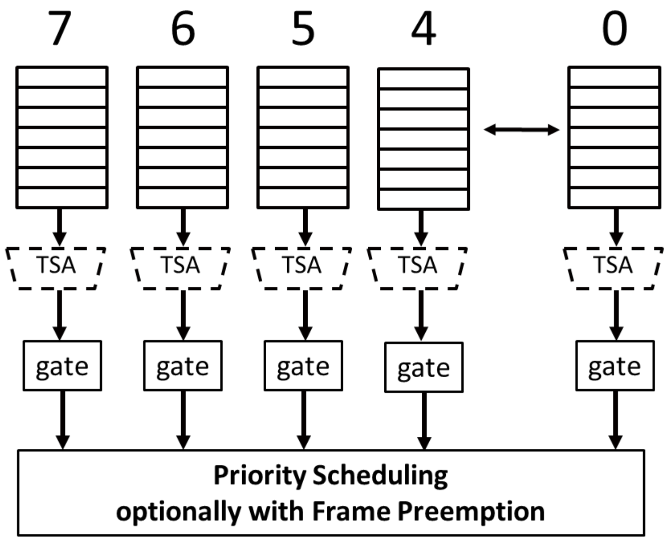

Gate Control List (GCL)

state	time interval
cco---cc	52 ns
cco---co	104 ns
cco---co	1230 ns
coo---co	800 ns
ooo---cc	52 ns
ooo---cc	250ns
ooo---cc	840 ns

Meeting Timing QoS = managing interfering traffic

Interferences experienced by a packet can be caused by:

– Lower-priority traffic

– Same-priority traffic

– Higher-priority traffic

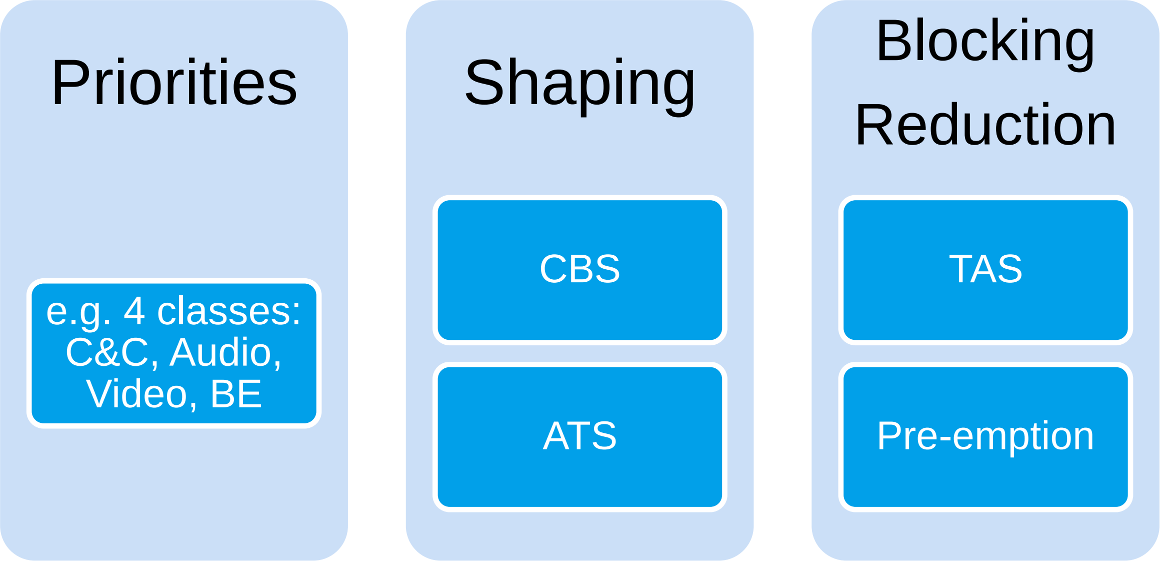

Priorities

e.g. 4 classes: C&C, Audio, Video, BE

Shaping

CBS

ATS

Blocking Reduction

TAS

Pre-emption

Each TSN classes of mechanisms combats one or several types of interferences:
- priorities reduces lower-priority interference
- shaping reduces higher-priority and same-priority interferences
- blocking reduction reduces lower-priority interference

Selection of the mechanisms based on max contributors to delays

Ex: using preemption on an automotive ring backbone (3-hops max typically) at 1+Gbit/s may not bring much benefits in terms of latencies for high-priority flows

The 3 classes of timing QoS mechanisms

Legend:

- CBS: Credit-Based Shaper [Qav]
- TAS: Time-Aware-Shaper [Qbv]
- Pre-emption [Qbu]
- ATS: Asynchronous Traffic Shaper [Qcr]

Not in the picture:

- CQF: Cyclic Queuing and Forwarding [Qch]
- ETS: Enhanced Transmission Selection [Qaz]

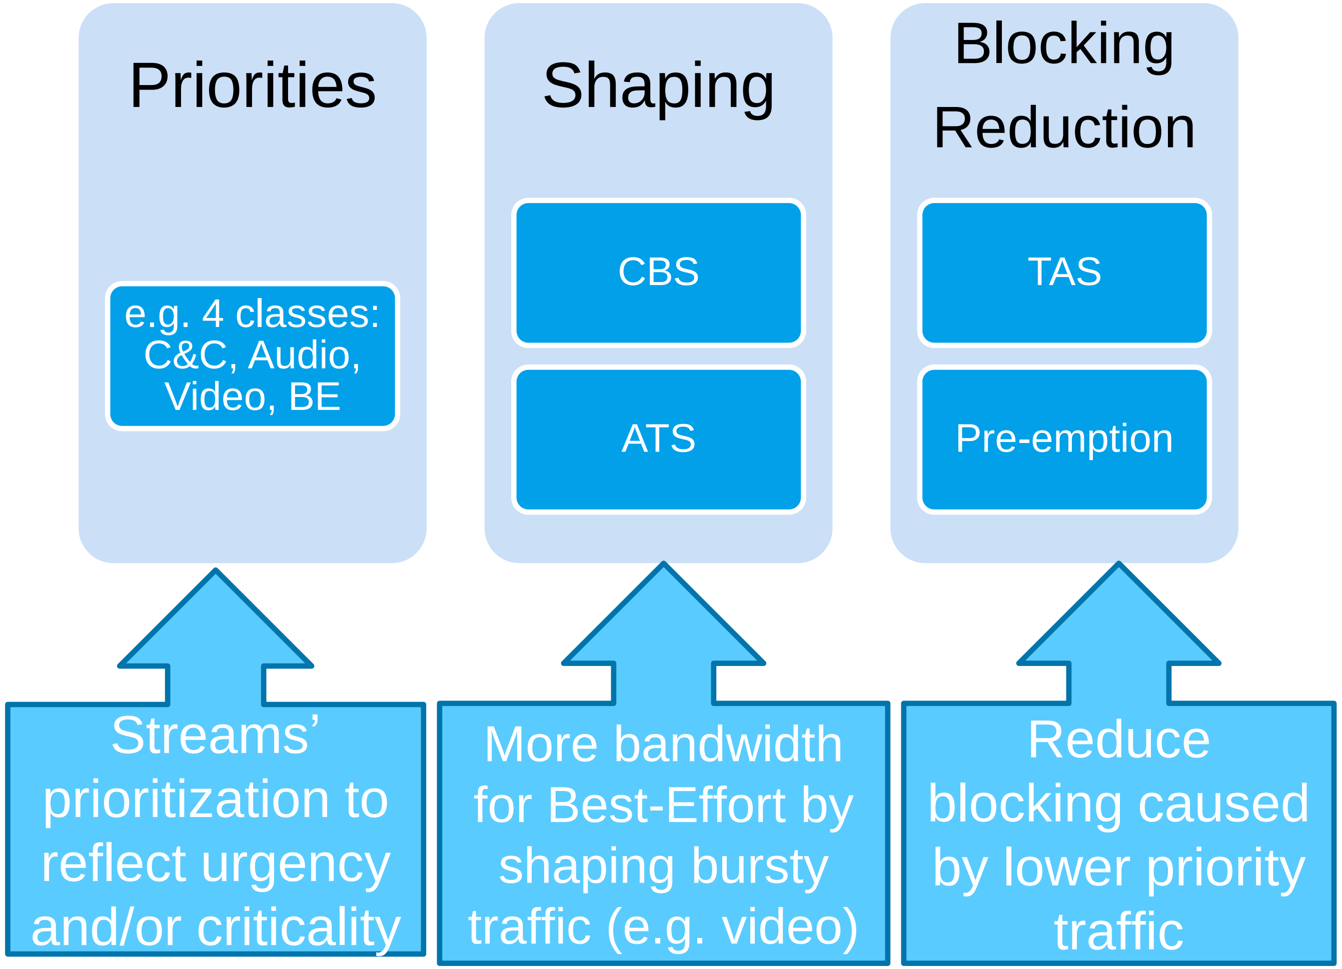

Priorities

e.g. 4 classes: C&C, Audio, Video, BE

Shaping

CBS

ATS

Blocking Reduction

TAS

Pre-emption

Streams' prioritization to reflect urgency and/or criticality

More bandwidth for Best-Effort by shaping bursty traffic (e.g. video)

Reduce blocking caused by lower priority traffic

Constraints typically addressed:

deadlines, safety

throughput, memory

strong deadlines, synchronization

Relevant TSN standards

Standard’s name	Also known as	nb
Strict Priority	IEEE 802.1p – 1998	–
Forward and Queueing for Time-Sensitive Streams (FQTSS)	IEEE 802.1Qav – 2009	Credit Based Shaper (CBS)
Enhancement for Scheduled Traffic	IEEE 802.1Qbv – 2015	Time Aware Shaper (TAS)
Frame Preemption	IEEE 802.1Qbu – 2016 & IEEE 802.3br – 2016	–
Cyclic Queuing & Forwarding (CQF)	IEEE 802.1Qch – 2017 / peristaltic shaper	–
Asynchronous Traffic Shaping (ATS)	IEEE 802.1Qcr – 2020	–

All mechanisms in IEEE 802.1Q-2022

- **Can be used in a combined manner, a certain QoS mechanism is applied to each TC**
- May have an important impact on memory usage

Criteria in the choice of a TSN QoS solution

- Meeting timing (deadlines, jitters, synchronization) & reliability constraints for each of the flow
- Max memory available in network devices
- Meeting network interfaces & bridge ports capabilities: delays, TSN support, memory (# of GCL), ….
- CPU overhead in end-systems in case of SW implemented mechanisms
- Network robustness wrt to departure from assumptions
- Evolutivity: e.g., adding streams without global re-configuration
- …

Multi-dimensional problem that can be solved with DSE algorithms (e.g., RTaW's ZeroConfig-TSN). Further research could lead to optimal solutions.

The use of priorities

The cornerstone of timing QoS

Assigning static priorities to flows

Pros

– Simple to use

– Efficient at the highest priority level for deadlines

– Can be combined with shaping & blocking reduction techniques

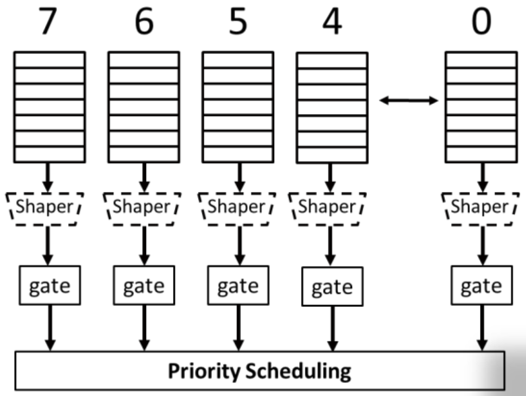

Limits

– Only 8 priority levels, not fine-grained enough for certain applications

– “Starvation” can happen at lowest priority levels with high priority bursty traffic

– Does not protect against some of the lower-priority interferences (“blocking factor”)

Configuration

– Optimal (schedulability-wise) priority assignment algorithm in O(n^2) ✓

– Extensions for:

  – safety/security constraints,

  – reserved priority levels,

  – streams from same functional domain at same priority level

  – …

The use of dynamic priorities, technically feasible, is still mostly unexplored

UNIVERSITÉ DU LUXEMBOURG

Frame Preemption

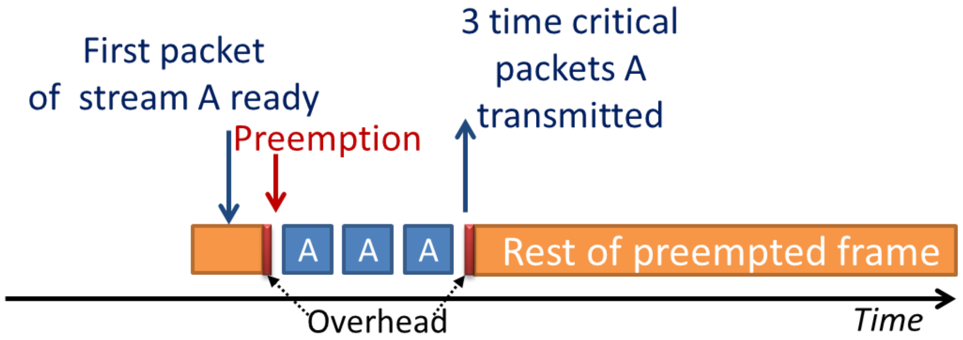

- Is a blocking reduction strategy
- Reduce by a factor 10 the “blocking factor” at each hop (*i.e.*, max. delay caused by a low priority frame under transmission)
- Provide a performance boost for all TSN scheduling solutions
- Less effective in high-speed networks with a small # of hops

Shaping Fundamentals

Why do Shaping?

What can be expected from Shaping?

Shapers basics (1/2)

- Shapers insert time gaps between packets of bursty streams, thereby increasing their latencies
- Lower priority packets can take advantage of the gaps to get transmitted, thereby decreasing their average latencies

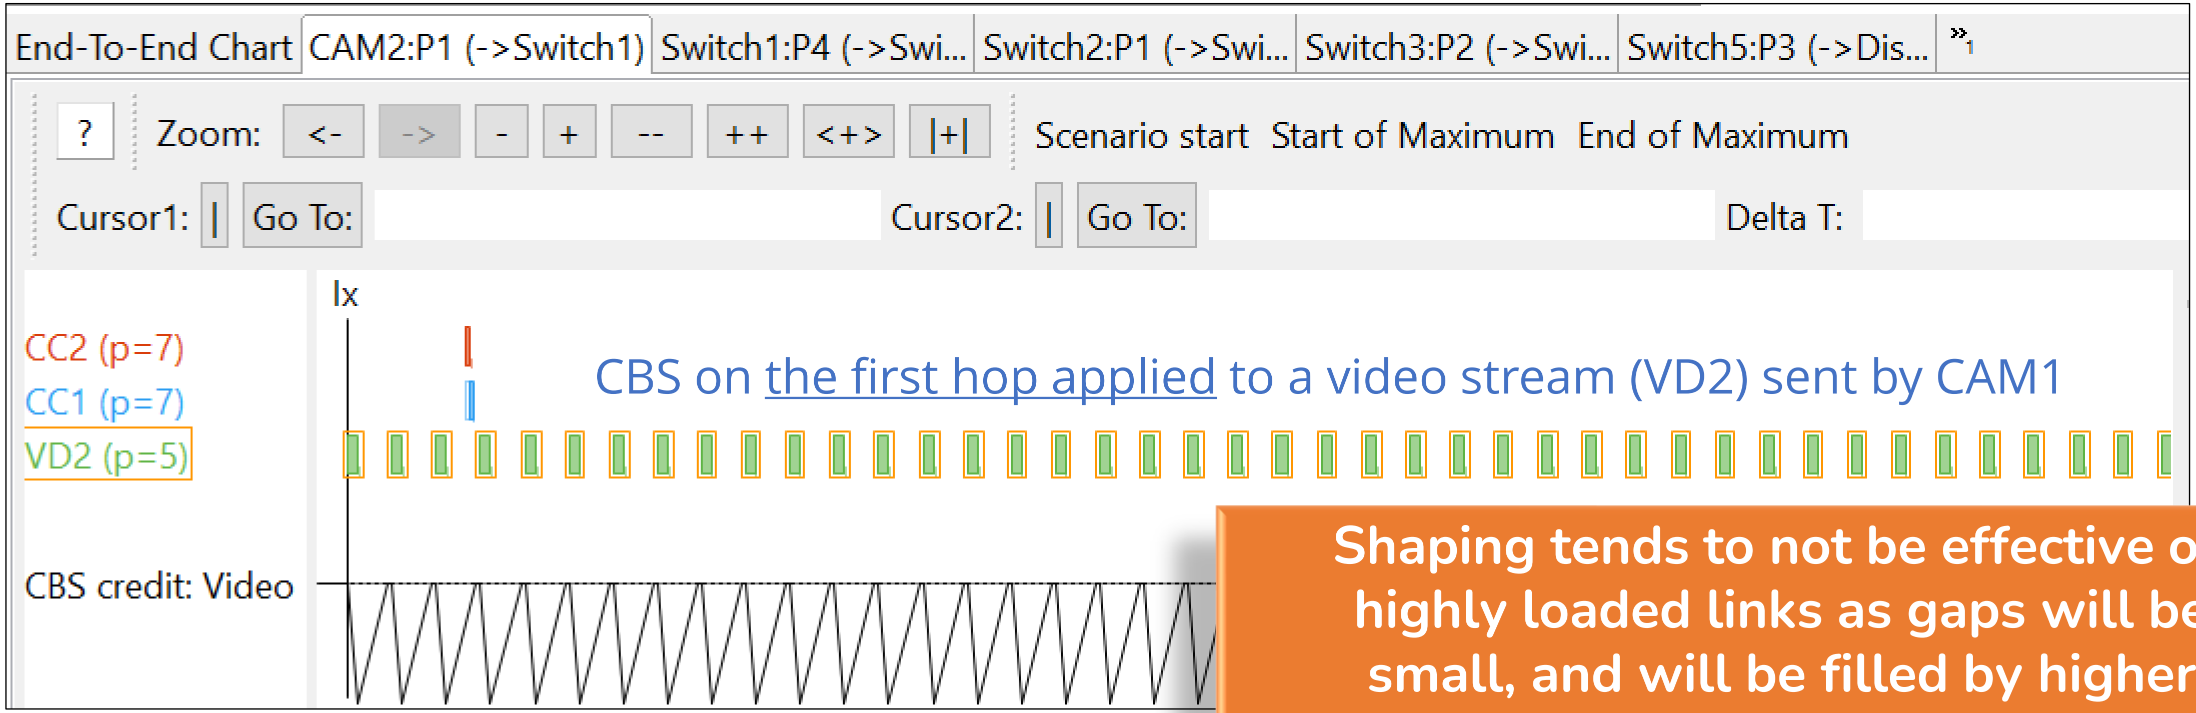

CBS on the first hop applied to a video stream (VD2) sent by CAM1

[RTaW-Pegase screenshot]

Shaping tends to not be effective on highly loaded links as gaps will be small, and will be filled by higher priority packets

Impact of shaping on max. memory usage

Illustration on a approx. 1000 streams helicopter avionics network

Max memory usage (bytes)

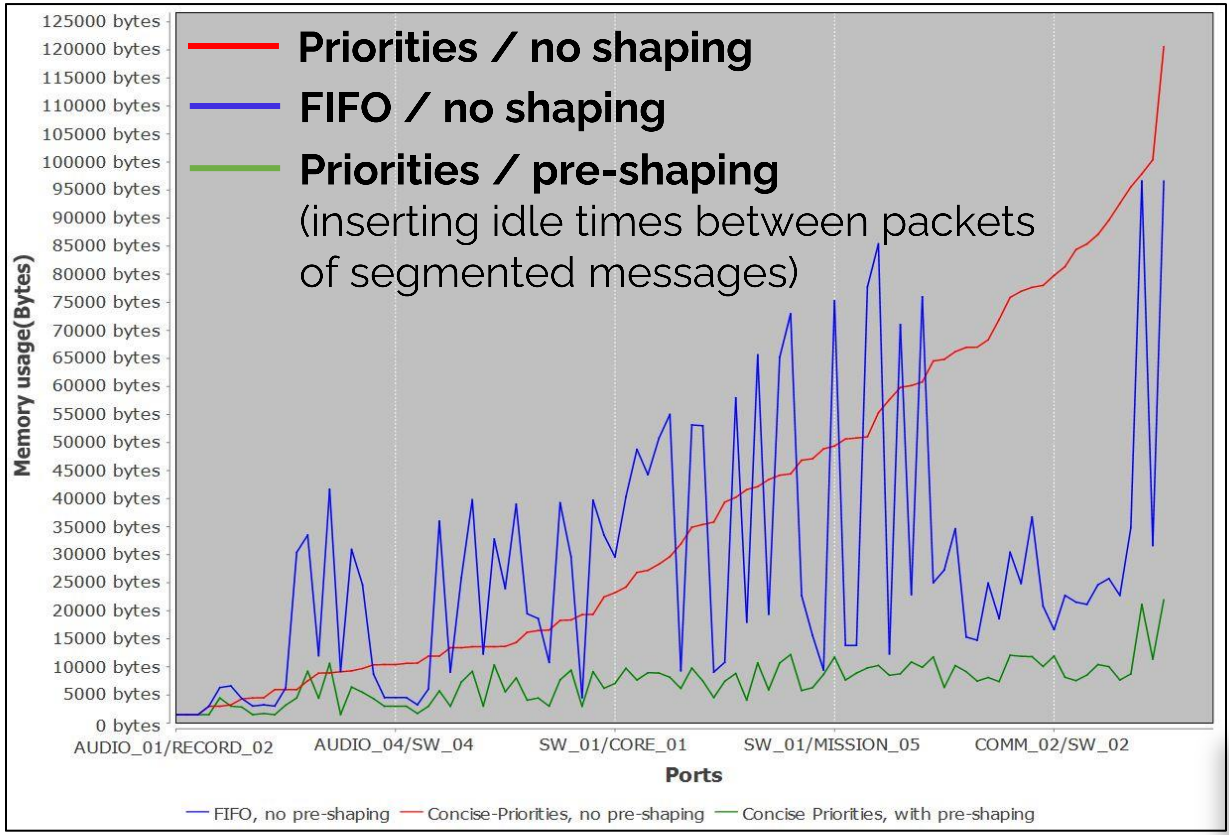

Priorities / no shaping

FIFO / no shaping

Priorities / pre-shaping (inserting idle times between packets of segmented messages)

Egress ports (interfaces, switches)

DASC 2021 - © Airbus Helicopters, RTaW, UL & Cognifyer

[RTaW-Pegase screenshot]

- **Shaping**, SW-implemented per-stream here, **reduces average memory usage by 80% here**
- Total memory usage per switch: up to 568KB without shaping vs 168KB with shaping
- Priorities do not reduce overall memory usage over FIFO
- CBS performs well wrt memory too

Shaping can be effective even if not done end-to-end and system-wide! For instance, it can be done only in end-nodes (as here) or only in bridges

Credit Based Shaper (CBS)

aka Forwarding and Queuing Enhancements for Time-Sensitive Streams (FQTSS)

Introduced for Audio/Video Bridging (AVB)

CBS scheduling rules

Each TC has a transmission credit

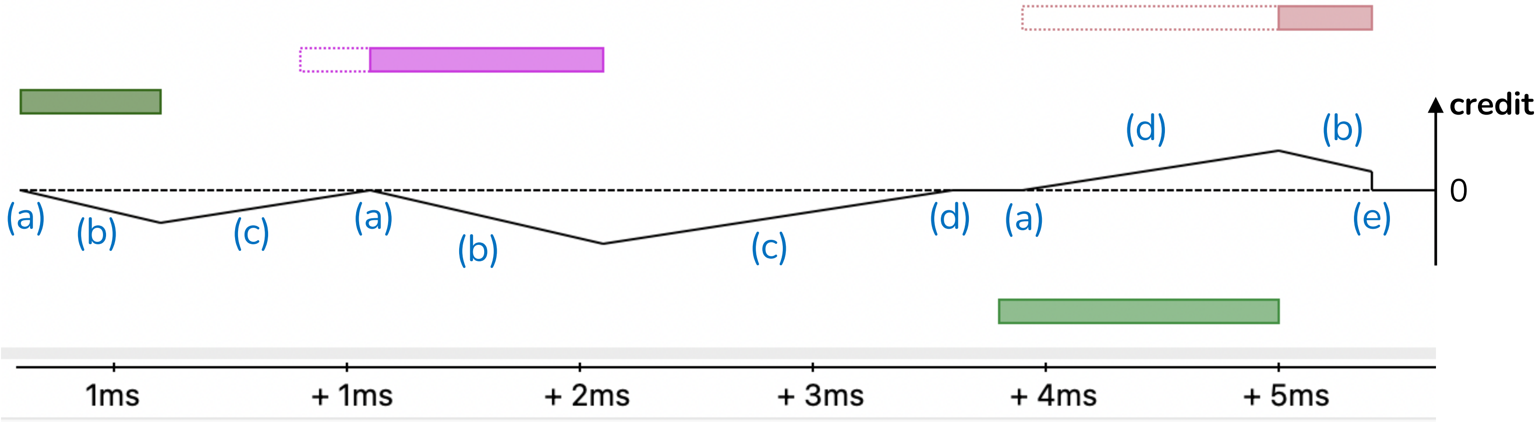

The 3 frames on top belong to the same CBS-shaped class

a) Head of the queue frame becomes eligible as soon as credit ≥ 0.

b) During transmission credit decreases with sendSlope = idleSlope – lineRate.

c) While credit is negative, it increases with idleSlope.

d) While frames are waiting and no frame of the class is being transmitted, the credit increases at rate idleSlope.

e) If credit > 0, no frames are waiting, and no frame is being transmitted, the credit is reduced to 0.

Slide from "What are the relevant differences between Asynchronous (ATS) and Credit Based (CBS) Shaper?", M. Turner, J. Migge, TSN/A Conference, 2023. Re-used with permission.

CBS effects on streams that come in shaped

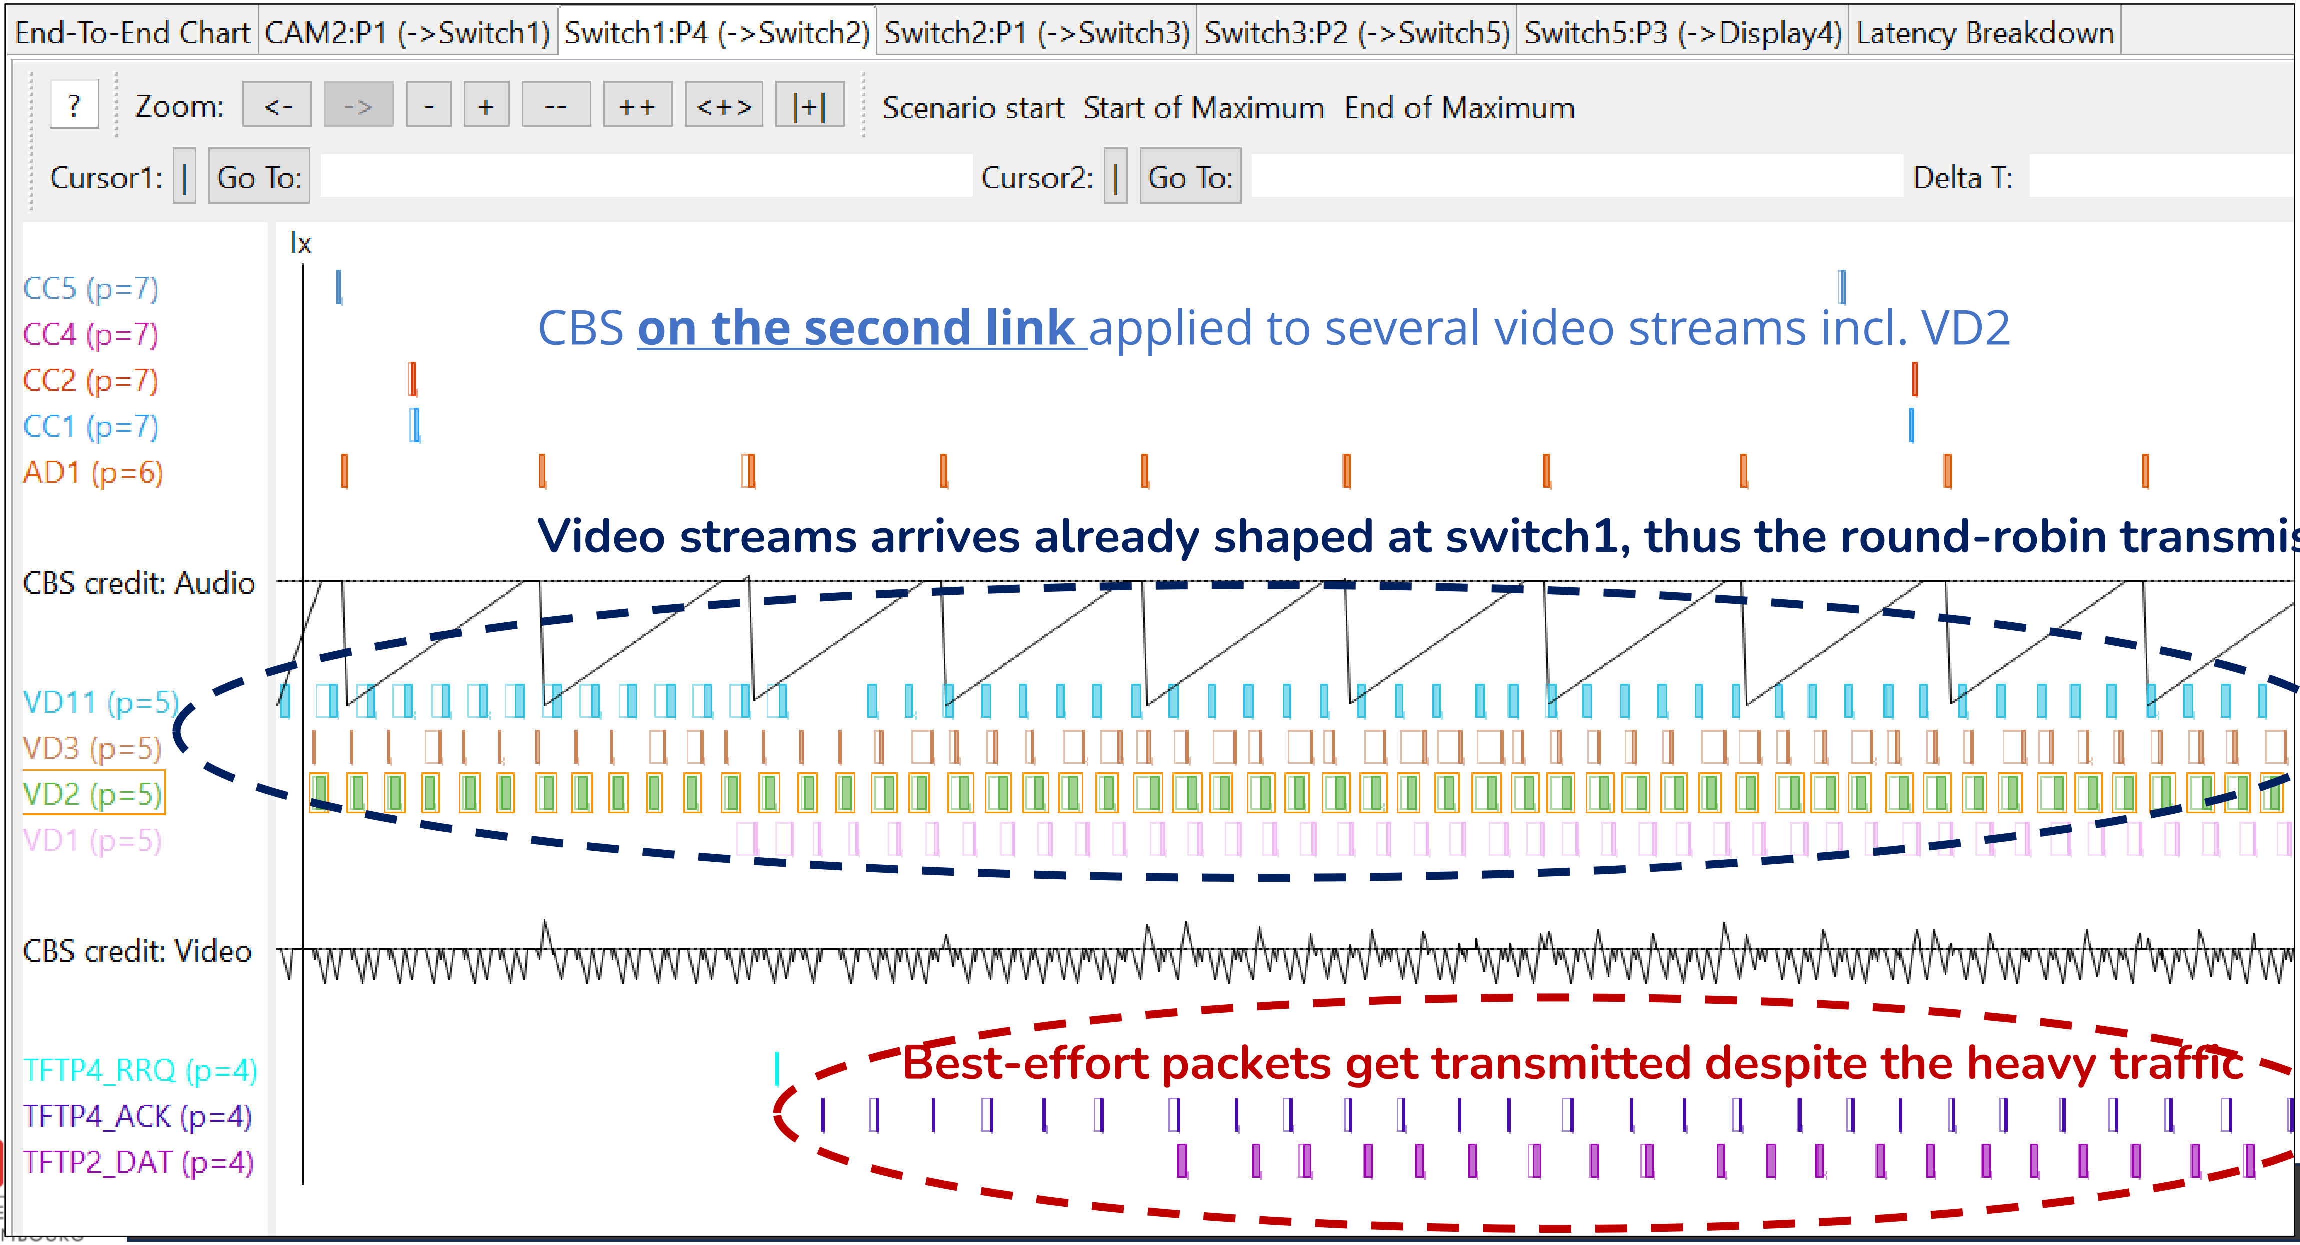

CBS **on the second link** applied to several video streams incl. VD2

Video streams arrives already shaped at switch1, thus the round-robin transmission

Best-effort packets get transmitted despite the heavy traffic

[RTaW-Pegase screenshot]

CBS : strong for shaping, less so for real-time constraints

Pros

- Effective in reducing the latencies for lower priority traffic
- Memory usage reduction
- HW implementation in all TSN network devices
- No need for a global clock

Limits

- Not a QoS policy for traffic with strong deadlines
- Coarse-grained shaper: per class and not per stream filtering
- Additional per-stream pre-shaping (at sender only) such as T-Spec can be needed to avoid bursts from the same streams within the CBS classes

Configuration

- Near-optimal (shaping-wise) algorithms for most practical use-cases ☑
- Configuration with standard AVB parameters (SR A, SR B) suited for plug & play use-cases
- Precision of the response time analyses ?

CBS is seemingly simple to use!

Asynchronous Traffic Shaper (ATS)

A shaper that is more versatile than CBS

A primer on ATS

Configuration options:

- 1 queue for a single flow
- 1 queue for a single group (= 1 input port + 1 output port + 1 priority level)
- 1 queue for several flows of the same group

Token buckets determine the eligibility time of the packets

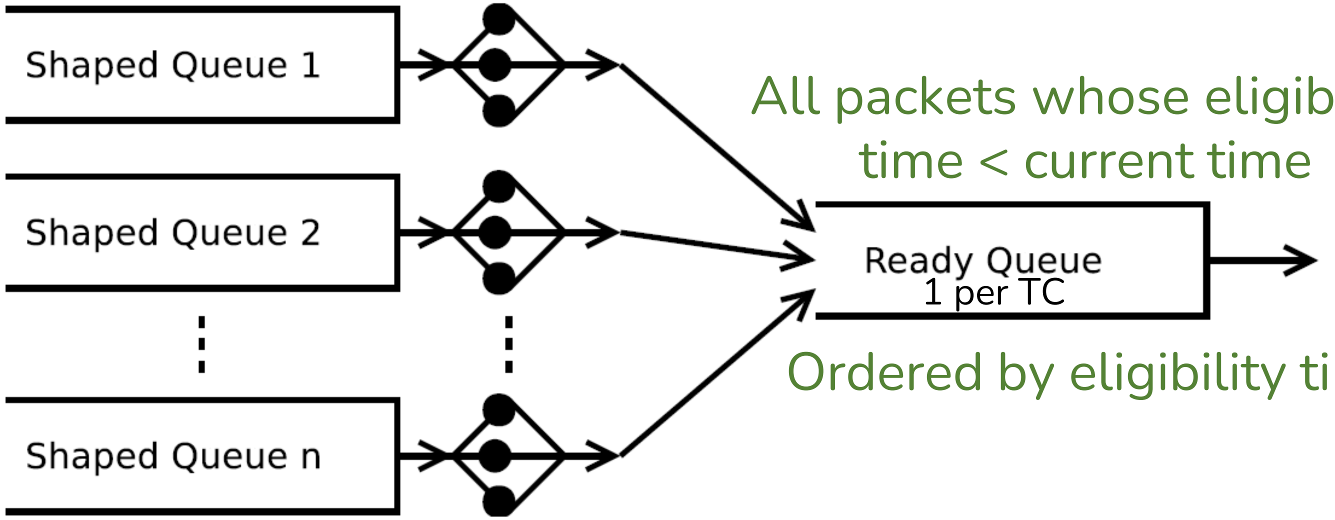

All packets whose eligibility time < current time

Ordered by eligibility times

The “theoretical model” [1] - HW implementation might differ

There is one “ATS scheduler” per shaped queue whose role it to compute the eligibility time

[1] From: Marc Boyer. Equivalence between the theoretical model and the standard algorithm of Asynchronous Traffic Shaping. 2022. hal-03788302

ATS Scheduling Rules – zoom in on single shaped queue

“Virtual Credit” of an ATS schedule = # of bits in the token bucket at time t

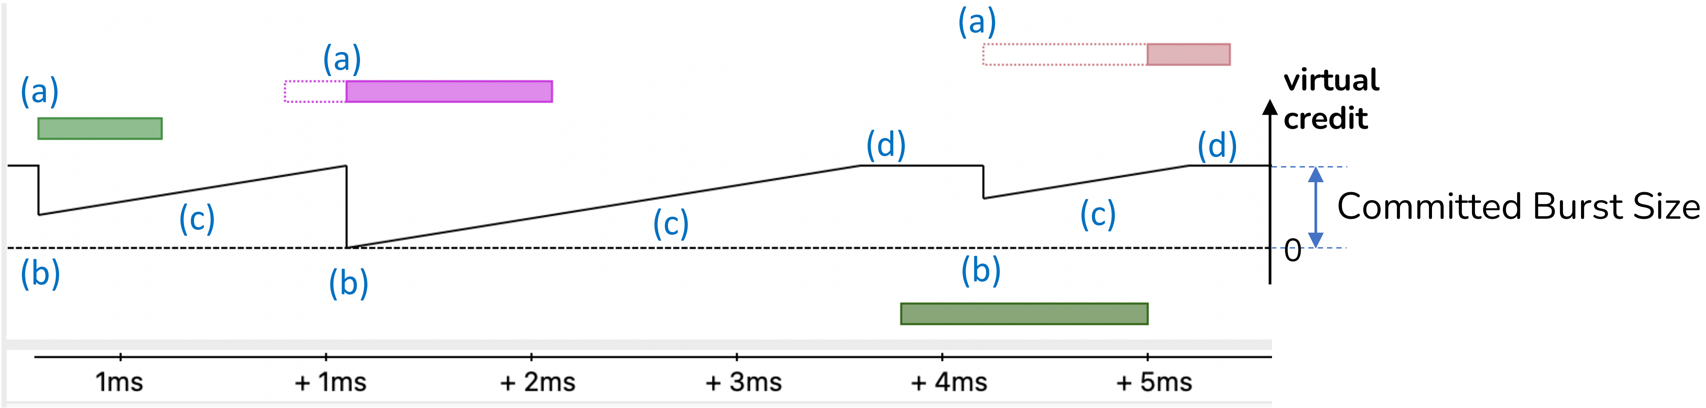

The 3 frames on top belong to the same ATS scheduler

a) The head of the line frame **is set eligible**, as soon as the **virtual credit** ≥ frame’s size

b) Virtual credit is **decreased at eligibility time (not transmission time)** by frame’s size

c) Virtual credit increases continuously with **C**ommitted **I**nformation **R**ate

d) if the credit reaches the **C**ommitted **B**urst **S**ize, it is **capped at Committed Burst Size**

Slide from “What are the relevant differences between Asynchronous (ATS) and Credit Based (CBS) Shaper?”, M. Turner, J. Migge, TSN/A Conference, 2023. Re-used with permission.

ATS : more powerful than CBS

Pros

- Fine-grained shaper: possibly per stream shaping at each hop
- Offer new trade-offs between real-time and shaping capabilities
- No need for a global clock

Limits

- Complex!
- First HW implementations available and no industry ROI yet
- Less effective than CBS for reducing lower-priority traffic latencies because of (small) bursts of ATS-shaped packets
- Cannot fully emulate CBS

Configuration

- Practical use-cases known so far covered by heuristics
- Precision of the response time analyses ?

ATS is an emerging technology and the TSN community is still learning how to use it

UNIVERSITÉ DU LUXEMBOURG

Time-Aware Shaper (TAS)

Time-triggered communication in TSN on a per-class (not per-flow) basis

Blocking reduction through reserved transmission windows

A primer on TAS

- The state of the transmission gate associated with each queue (“Open” or “Closed”) determines whether a frame from a given queue can be selected for transmission

- The Gate Control List (GCL) associated with each port contains an ordered list of gate operations. It defines a schedule that is executed in cycle at run-time

“Exclusive gating”: some TCs with critical flows are given exclusive link access during their transmission windows

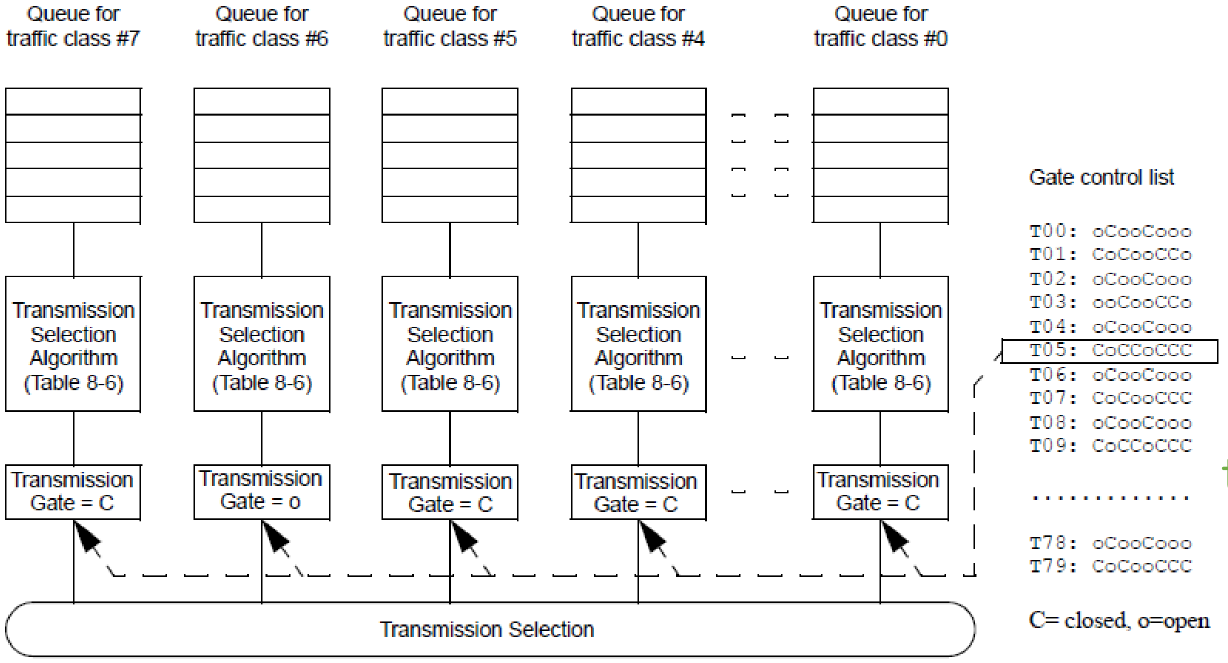
Queue for traffic class #7
Queue for traffic class #6
Queue for traffic class #5
Queue for traffic class #4
Queue for traffic class #0
Transmission Selection Algorithm (Table 8-6)
Transmission Selection Algorithm (Table 8-6)
Transmission Selection Algorithm (Table 8-6)
Transmission Selection Algorithm (Table 8-6)
Transmission Selection Algorithm (Table 8-6)
Transmission Gate = C
Transmission Gate = o
Transmission Gate = C
Transmission Gate = C
Transmission Gate = C
Transmission Selection

Gate control list

```
T00: oCooCooo
T01: CoCooCCo
T02: oCooCooo
T03: ooCooCCo
T04: oCooCooo
T05: CoCCoCCC
T06: oCooCooo
T07: CoCooCCC
T08: oCooCooo
T09: CoCCoCCC

.............

T78: oCooCooo
T79: CoCooCCC
```

C= closed, o=open

Figure 8-16—Transmission selection with gates

From [802.1Q §8.6.8.4]

A head of queue frame is eligible for transmission if the transmission gate is in the open state and there is sufficient time left to transmit the entirety of that frame

TAS : only as good the configuration algorithm

Pros

- Tight timing constraints can be met (deadline, synchronization)
- Timing verification is easier, timing behavior is close to deterministic
- Support bandwidth-sharing use-cases as well (e.g., 1ms every 10ms is reserved to an Android partition)

Limits

- Per-class shaping and not per stream shaping → jitters
- Manual construction of comm. schedules unfeasible except for a few flows
- Task schedule must be tailored to communication schedule for best performances
- Clock synchronization needed
- Bandwidth utilization not optimal

Configuration

- NP-hard problem
- Heuristics reasonably good for exclusive gating

TAS is well suited for certain use-cases (e.g., control loops) but at the expense of a strong coupling between SW and communication. Other use-cases may emerge.

TSN QoS for embedded networks: a consolidating landscape

Ongoing work: Cut-Through Forwarding, a generic speed-up feature like preemption

Use-cases well served by the different QoS policies are progressively identified .. still configuration remains a challenge for the practitioners and requires complex tooling

We see 2 approaches to QoS configuration:

— Automated black-box approaches: from goals & constraints to device config. files

— “Explainable” configuration based on best practices and scheduling theory results

Topics for further research:

DSE, ATS, local use of TSN mechanisms, per-hop shaping mechanisms, HW heterogeneity, dynamically evolving QoS parameters, verification techniques for complex traffic patterns (e.g., TCP)…