Ten years of assisting customers in TSN configuration

September 10, 2024

Jörn Migge, C.T.O. RTaW
jorn.migge@realtimeatwork.com

The very beginning ….

RTaW
RealTime-at-Work

It all started with CAN (Controller Area Network)…
We had
- a mathematical formula for upper bounds on CAN bus transmission delays
- the believe that it could
  - allow to make CAN networks safer and optimized
  - be made usable by network engineers (software tool)
- and the will to try to make this happen ….

We also had
- algorithms for optimal CAN id assignments
- algorithms for optimal frame offset assignments

But there where quite some difficulties to overcome …

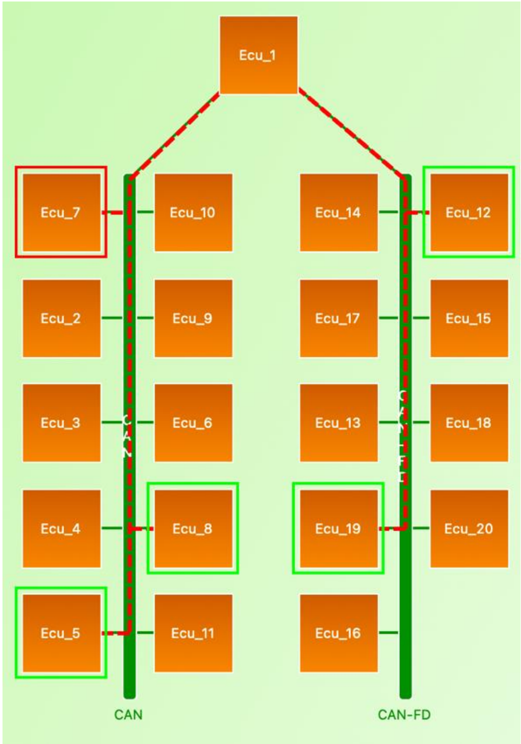

Difficulties to overcome

Name	Id(hex)	PayLoad	Period	MinDela	TxMode	Sender	Bound	Deadline	Slack
Frame2	0x2	4 byte		10 ms	E	ECU2	2,408 ms	2 ms	-0,408 ms
Frame3	0x3	32 byte	100 ms	20 ms	P+E	ECU3	0,661 ms	4 ms	3,339 ms
Frame5	0x5	3 byte	50 ms		P	ECU6	3,874 ms	10 ms	6,126 ms
Frame5	0x5	3 byte	50 ms		P	ECU6	0,840 ms	10 ms	9,160 ms
Frame1	0x1	5 byte	30 ms		P	ECU1	1,720 ms	6 ms	4,280 ms
Frame1	0x1	5 byte	30 ms		P	ECU1	1,720 ms	3 ms	1,280 ms
Frame4	0x4	8 byte	100 ms		P	ECU3	3,024 ms	20 ms	16,976 ms
Frame4	0x4	8 byte	100 ms		P	ECU3	0,840 ms	20 ms	19,160 ms
Frame4	0x4	8 byte	100 ms		P	ECU3	0,840 ms	20 ms	19,160 ms

Upper bounds?

- CAN was used since a while and the simple rule of a **maximal bus load of 30%** seemed to work fine …
- In general, because of over-approximation, no explanation can be given ..

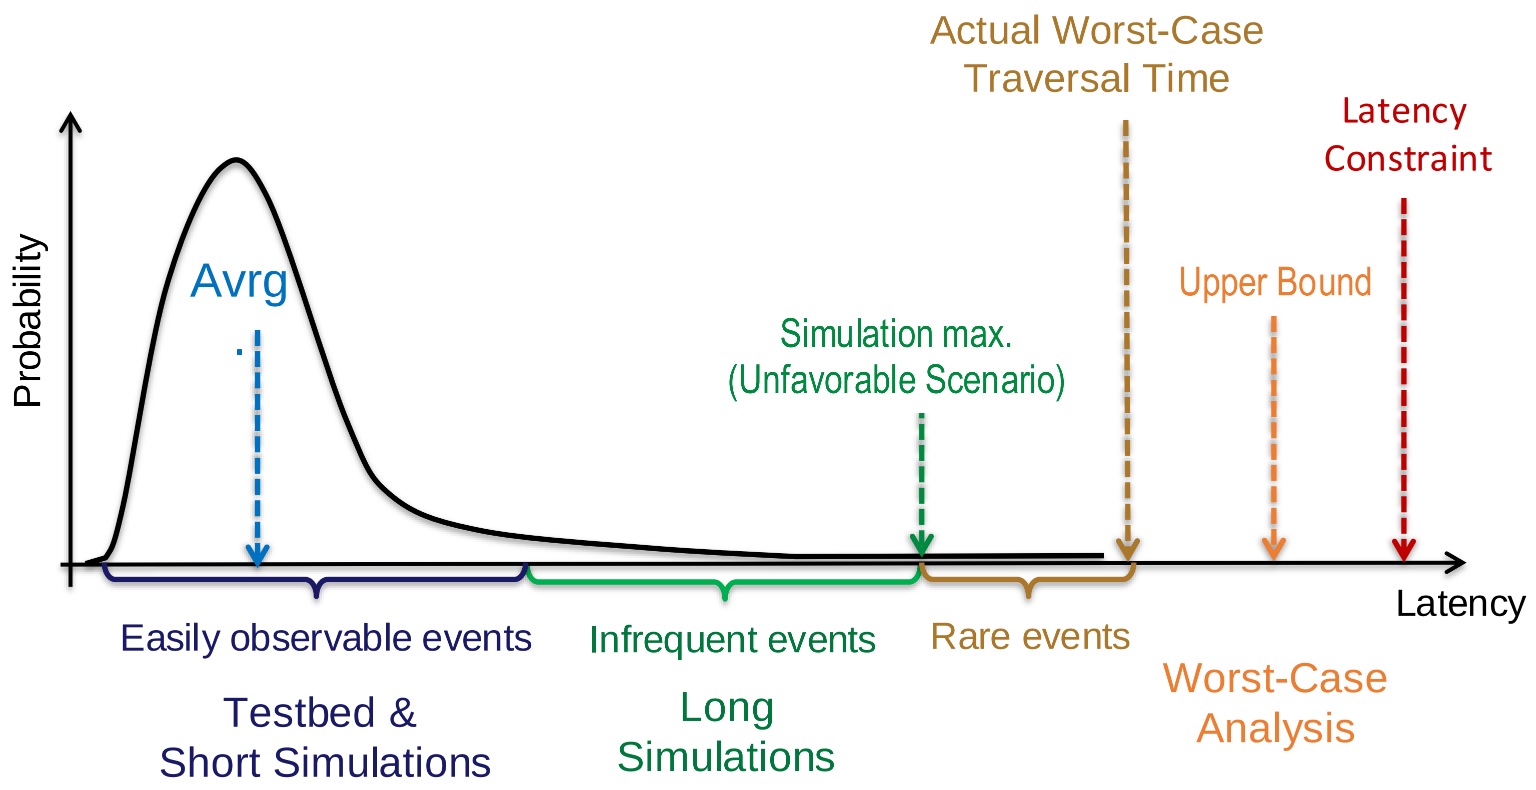

Difficulties to overcome

We added Simulation

- as complementary evaluation technique
- and with the possibility to re-simulate and **illustrate as Gantt** chart the scenarios where the **simulation maximum** has been seen: **observed unfavorable scenario**

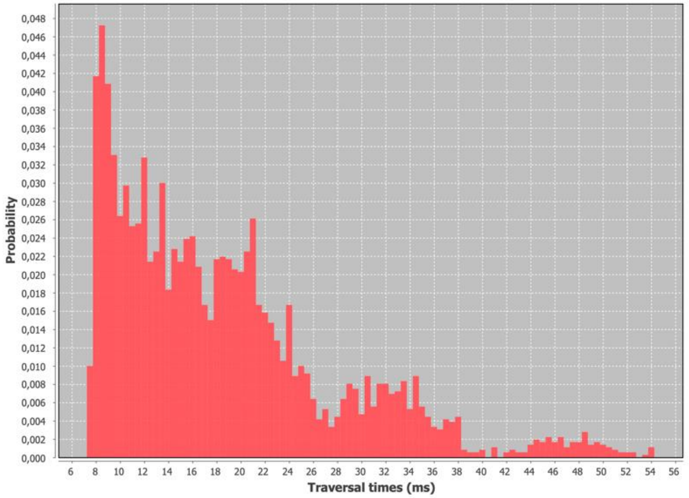

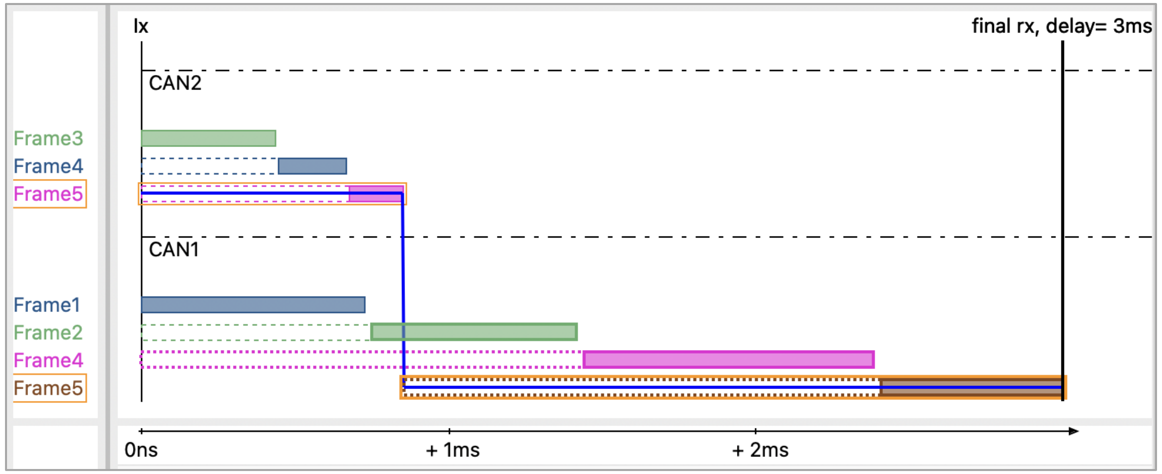

How to make simulation max become closer to the worst-case is a research topic:

P. Keller, N. Navet, "Approximation of Worst-Case Traversal Times in Real-Time Ethernet Networks: Exploring the Potential of Many-Objective Optimization for Simulation Aggregation", Proc. 20th IEEE International Workshop on Factory Communication System (WFCS 2024), Toulouse, France, April 17-19, 2024.

Difficulties to overcome

- **CAN ids**: there were more used as a composed message identifier for filtering than ass used as priorities
  - We developed an algorithm that changes only a minimal number of CAN IDs
- **Transmission offsets**: some ECU cannot support them or only with a higher granularity for the offsets
  - We adapted and extended our algorithm accordingly

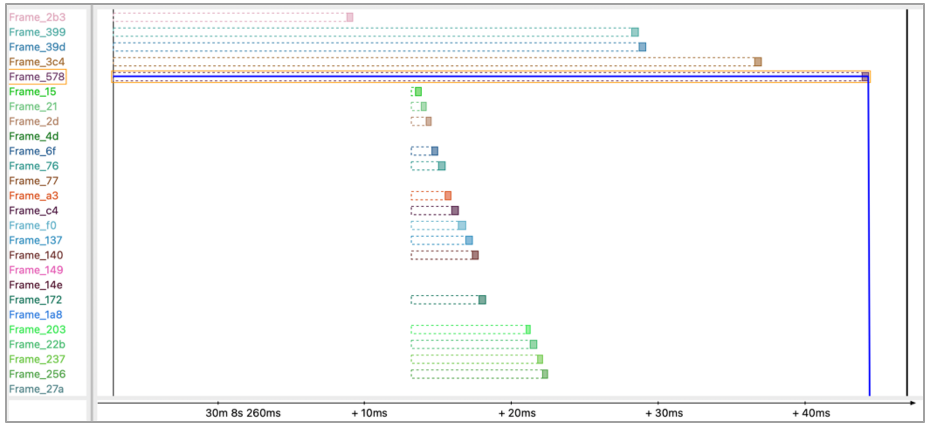

no transmission offsets

Unfavorable Scenarios

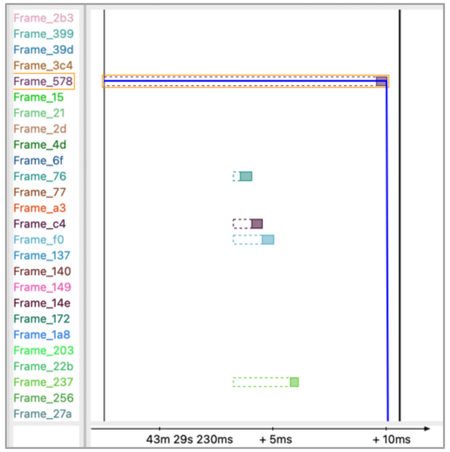

good transmission offsets

Difficulties to overcome

- The mathematical model makes precise assumptions about the physical system, that need to be checked
  - We added trace analysis
    - Priority inversion
    - Actual periods and occurrences of frames
    - Frame sizes and Bit stuffing estimations

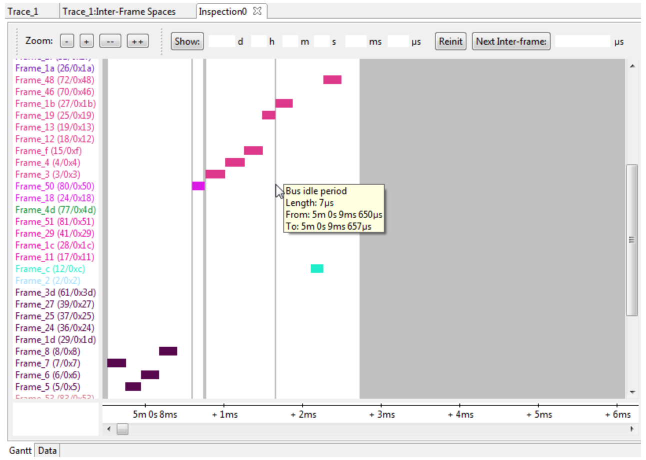

- Engineers started to appreciate our tool when the maximal bus load rule reached 50%
  - With the optimization of priorities, the introduction of offsets, the computation of upper bounds it is safer to go beyond 50 % ….

AFDX

RTaW
RealTime-at-Work

- Our Ethernet journey actually started with **AFDX**, at a time when **BMW** was already investigating seriously if **Ethernet** could become a solution for the automotive domain
- **Network Calculus** provided formulas for computing upper bounds on network traversal times, but they had the reputation to provide **pessimistic bounds**
- A research and technology transfer project called "PEGASE" was initiated by our academic partner ONERA in order to **improve the precision of Network Calculus delay bounds**. Our role was to implement the new formulas in a software tool, to make them available to industrial users.
- This was the starting point for a long-term collaboration with ONERA that has allowed to extend Network Calculus and the support of **TSN** in our tool .

TSN

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	–
Asynchronous Traffic Shaping (ATS)	IEEE 802.1Qcr – 2020	–

- They have an important impact not only on delays but also on memory requirements
- Can be used in a combined manner

Configuration of TSN mechanisms

RTaW
RealTime-at-Work

CBS

PortOwner	Port	Destination	TxRate	IdleSlope = 'Audio'	Load : 'Audio'	LoCredit : 'Audio'	HiCredit : 'Audio'	IdleSlope = 'Video'	Load : 'Video'
R1	EthernetPort_1	ES_6	100 Mbit/s	8,686 Mbit/s	7,405 Mbit/s	-4252 bits	772 bits	38,544 Mbit/s	22,707 Mbit/s
ES_6	EthernetPort_11	R1	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	0 Mbit/s	0 Mbit/s
R1	EthernetPort_2	ES_7	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	0 Mbit/s	0 Mbit/s
ES_7	EthernetPort_12	R1	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	17,494 Mbit/s	4,172 Mbit/s
R2	EthernetPort_7	ES_2	100 Mbit/s	6,458 Mbit/s	5,868 Mbit/s	-2440 bits	1123 bits	14,502 Mbit/s	11,37 Mbit/s
ES_2	EthernetPort_13	R2	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	18,409 Mbit/s	4,138 Mbit/s
R2	EthernetPort_8	ES_1	100 Mbit/s	[illegible]	0 Mbit/s	0 bits	0 bits	0 Mbit/s	0 Mbit/s
ES_1	EthernetPort_14	R2	100 Mbit/s	1[illegible]	[illegible]2 Mbit/s	-3948 bits	2270 bits	53,046 Mbit/s	34,076 Mbit/s
R1	EthernetPort_3	ES_5	100 Mbit/s	7,[illegible]	[illegible]4 Mbit/s	-4299 bits	1384 bits	71,808 Mbit/s	16,618 Mbit/s
ES_5	EthernetPort_15	R1	100 Mbit/s	2,152 Mbit/s	1,957 Mbit/s	-2552 bits	265 bits	0 Mbit/s	0 Mbit/s
R2	EthernetPort_9	ES_3	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	14,502 Mbit/s	11,37 Mbit/s
ES_3	EthernetPort_16	R2	100 Mbit/s	2,152 Mbit/s	1,957 Mbit/s	-2552 bits	303 bits	18,409 Mbit/s	4,138 Mbit/s
R1	EthernetPort_4	ES_4	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	0 Mbit/s	0 Mbit/s
ES_4	EthernetPort_17	R1	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	17,494 Mbit/s	4,172 Mbit/s
R1	EthernetPort_5	R2	1000 Mbit/s	2,152 Mbit/s	1,957 Mbit/s	-2603 bits	34 bits	0 Mbit/s	0 Mbit/s
R2	EthernetPort_10	R1	1000 Mbit/s	13,068 Mbit/s	10,896 Mbit/s	-4596 bits	199 bits	75,363 Mbit/s	30,982 Mbit/s
R1	EthernetPort_6	ES_8	100 Mbit/s	6,534 Mbit/s	5,448 Mbit/s	-4352 bits	763 bits	14,96 Mbit/s	11,303 Mbit/s
ES_8	EthernetPort_18	R1	100 Mbit/s	0 Mbit/s	0 Mbit/s	0 bits	0 bits	0 Mbit/s	0 Mbit/s

TAS

GateControlList + −

- Asap:ES_1:EthernetPort_14 (->R2)
- Asap:ES_2:EthernetPort_13 (->R2)
- Asap:ES_3:EthernetPort_16 (->R2)
- Asap:ES_4:EthernetPort_17 (->R1)
- Asap:ES_5:EthernetPort_15 (->R1)
- Asap:ES_6:EthernetPort_11 (->R1)
- Asap:ES_7:EthernetPort_12 (->R1)
- Asap:ES_8:EthernetPort_18 (->R1)
- Asap:R1:EthernetPort_1 (->ES_6)
- Asap:R1:EthernetPort_2 (->ES_7)
- Asap:R1:EthernetPort_3 (->ES_5)
- Asap:R1:EthernetPort_4 (->ES_4)
- Asap:R1:EthernetPort_5 (->R2)
- Asap:R1:EthernetPort_6 (->ES_8)
- Asap:R2:EthernetPort_10 (->R1)
- Asap:R2:EthernetPort_7 (->ES_2)
- Asap:R2:EthernetPort_8 (->ES_1)
- Asap:R2:EthernetPort_9 (->ES_3)

Name* Asap:ES_1:EthernetPort_14 (->R2) IsDefaultConfiguration*

GateOperations | Ports

1 selected GateOperations

Interval	StartTime	EndTime	TC : CC	TC : Audio	TC : Video	TC : Data
0,026000 ms	3,890000 ms	3,916000 ms	open	closed	closed	closed
5,204000 ms	3,916000 ms	9,120000 ms	closed	open	open	open
0,010000 ms	9,120000 ms	9,130000 ms	open	closed	closed	closed
0,432000 ms	9,130000 ms	9,562000 ms	closed	open	open	open
0,026000 ms	9,562000 [illegible]	[illegible]	open	closed	closed	closed
5,400000 ms	9,588000 [illegible]	[illegible]	closed	open	open	open
0,010000 ms	14,988000 [illegible]	[illegible]	open	closed	closed	closed
4,122000 ms	14,998000 ms	19,120000 ms	closed	open	open	open
0,010000 ms	19,120000 ms	19,130000 ms	open	closed	closed	closed
9,990000 ms	19,130000 ms	29,120000 ms	closed	open	open	open
0,010000 ms	29,120000 ms	29,130000 ms	open	closed	closed	closed
5,858000 ms	29,130000 ms	34,988000 ms	closed	open	open	open
0,010000 ms	34,988000 ms	34,998000 ms	open	closed	closed	closed
4,122000 ms	34,998000 ms	39,120000 ms	closed	open	open	open
0,010000 ms	39,120000 ms	39,130000 ms	open	closed	closed	closed

ATS

GroupId	Ingress	Egress	MaxResidence	SchedulerId	CBS	Max Size	CIR	Load
n/a	n/a	ES_1/EthernetPort_14	n/a	0	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
n/a	n/a	ES_1/EthernetPort_14	n/a	1	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
n/a	n/a	ES_1/EthernetPort_14	n/a	2	582 byte	562 byte	3,5 Mbit/s	3492,000 kbit/s
n/a	[illegible]	ES_1/EthernetPort_14	n/a	4	582 byte	562 byte	3,5 Mbit/s	3492,000 kbit/s
n/a	[illegible]	ES_1/EthernetPort_14	n/a	6	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
n/a	n/a	ES_3/EthernetPort_16	n/a	0	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
n/a	n/a	ES_5/EthernetPort_15	n/a	0	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
0	R2/EthernetPort_8	R2/EthernetPort_10	n/a	0	582 byte	562 byte	11 Mbit/s	10896,000 kbit/s
0	R1/EthernetPort_5	R1/EthernetPort_1	n/a	0	582 byte	562 byte	5,5 Mbit/s	5448,000 kbit/s
1	R1/EthernetPort_5	R1/EthernetPort_6	n/a	1	582 byte	562 byte	5,5 Mbit/s	5448,000 kbit/s
3	R1/EthernetPort_5	R1/EthernetPort_3	n/a	3	582 byte	562 byte	7 Mbit/s	6984,000 kbit/s
6	R2/EthernetPort_8	R2/EthernetPort_7	n/a	6	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
7	R2/EthernetPort_9	R2/EthernetPort_7	n/a	7	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
6	R1/EthernetPort_3	R1/EthernetPort_5	n/a	6	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
8	R2/EthernetPort_10	R2/EthernetPort_7	n/a	8	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s
7	R1/EthernetPort_3	R1/EthernetPort_1	n/a	7	326 byte	306 byte	2 Mbit/s	1956,000 kbit/s

Mechanisms may be combined

Name* Classes Configuration   TAS Config TAS Gate Control Lists

ATS Config ATS parameters   CBS Config CBS Idle Slopes

4 ClassConfigurations + −

TrafficClass	Priority	Express	SchedulingPolicy
C&C	7	true	FIFO
Audio	6	false	CBS
Video	5	false	ATS
Data	4	false	FIFO

Priorities

⇒ a lot of parameters that must be configured “well”: **correct** and **optimized**

Configuration & Optimization Goals

RTaW
RealTime-at-Work

- 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 configuration robustness wrt to departure from assumptions
- Evolutivity: e.g., adding streams without global re-configuration
- …

⇒ Multi-dimensional problem

Configuration of TSN mechanisms

Our configuration algorithms

- use **worst-case analysis to validate** the potential solution (Network Calculus, …)
- are based on **optimality results** if possible
- or **heuristics**, based on our **deep insight on how scheduling mechanism work** and what they allow to achieve

Remarks:

- we have done experiments on CBS configuration with **genetic algorithms**, but the solutions where not better and took several hours to converge, compared to less than a minute for the heuristics
- we are also looking into AI approaches, but for know we have not yet found any efficient approach for the scheduling problems themselves

What can be achieved with TSN mechanisms?

RTaW
RealTime-at-Work

Interferences experienced by a packet can be caused by:

- Lower-priority traffic
- Same-priority traffic
- Higher-priority traffic

TSN scheduling mechanisms may reduce one or several types of interferences.

We are using TAS in a particular way:

- time triggered scheduling for critical a class
- bandwidth partitioning with respect to all other traffic classes

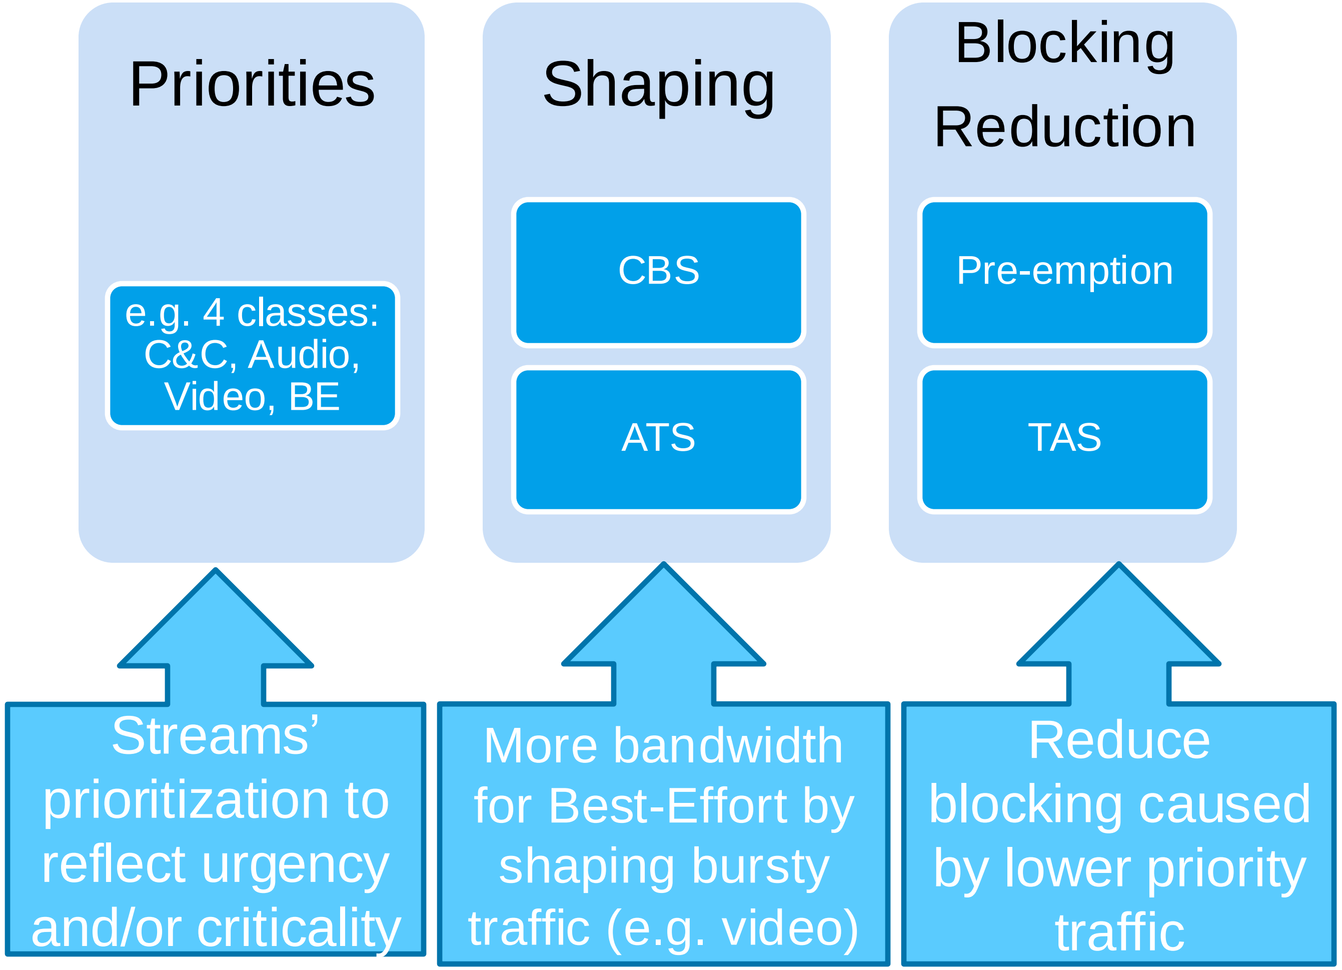

Constraints typically addressed:	deadlines, safety	throughput, memory	strong deadlines, synchronization

ZeroConfig-TSN (ZCT)

RTaW
RealTime-at-Work

- Rework Topology
- Revise Constraints

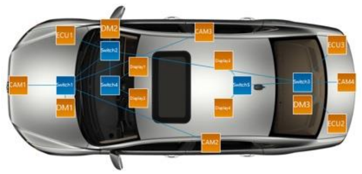

Communication Requirements
- Latency and Throughput

Network Topology
- Switches and Link Speeds

Allowed QoS mechanisms
- traffic classes, priorities
- shapers: CBS, TAS
- preemption, pre-shaping

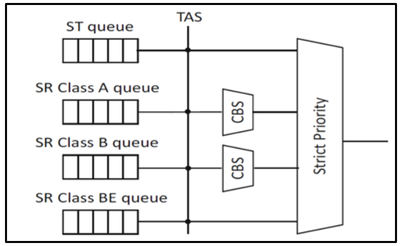

Zero Config TSN

Add more mechanisms

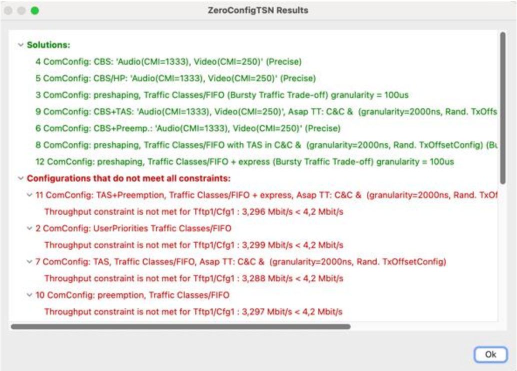

Configurations meeting all requirements = Verified Solutions

Configurations with some requirement violations

Training

We offer training that allows users to learn through our design tool
- how TSN mechanism work
- and when to apply them

Unfavorable Scenarios provide valuable insights!

Memory requirements in switch port with CSB shaping:

Without CBS shaping, 4 streams of continuously incoming frames accumulate:

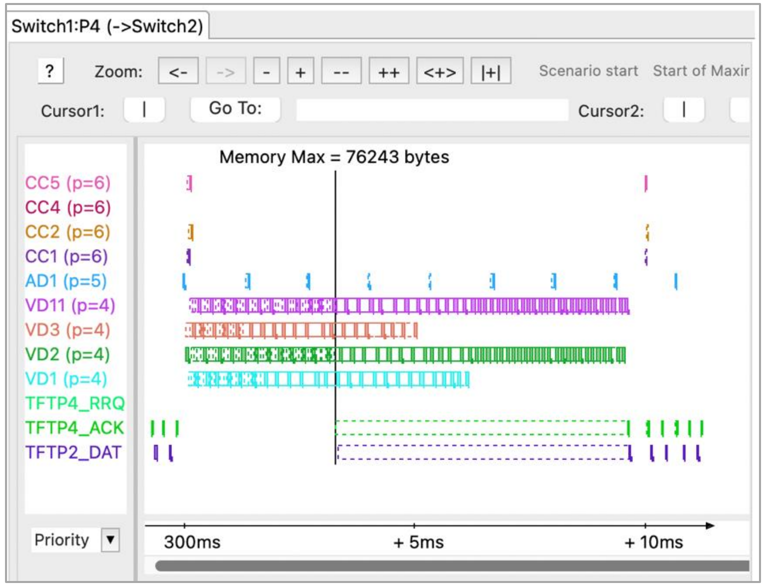

With CBS shaping, only few frames of the 4 streams arrive at the same time:

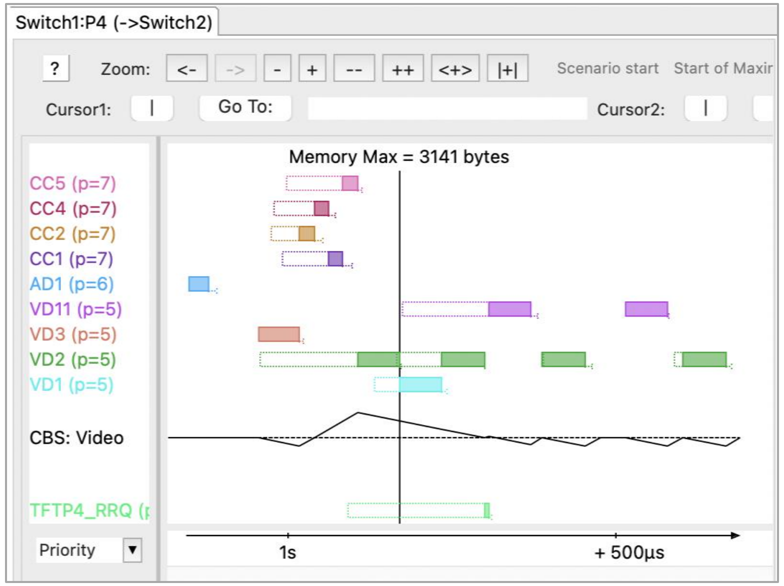

Illustrated definitions of shapers

RTaW
RealTime-at-Work

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 scheduling rules

Each TC has a transmission credit

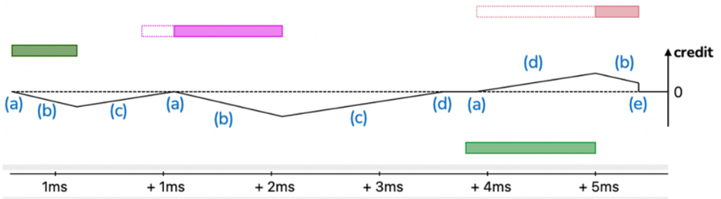

The 3 frames on top belong to the same CBS-shaped class

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.

ATS Scheduling Rules – zoom in on single shaped queue

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

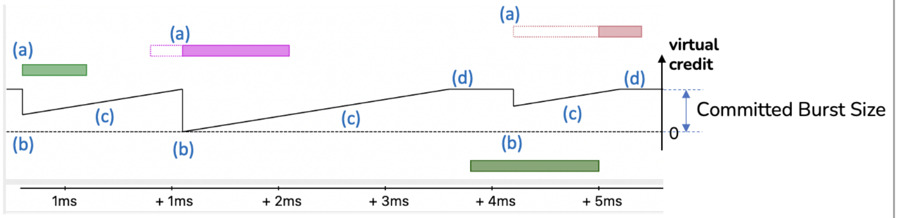

The 3 frames on top belong to the same ATS scheduler

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**

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

Conclusion and outlook

Conclusion

- Applying state of the art scheduling theory allows to improve designs
- Theories need to be adapted to take into account real word particularities
- Academic collaborations where and are crucial for meeting our initial goals

Outlook

- We listen to customer needs and try to support new scheduling mechanisms: HTB, PON, …
- We seek to increase of the number of criteria taken into account by the configuration algorithms
- We continuously seek to improve our configuration algorithms with respect to real-world requirements

"We help you build provably safe and optimized critical systems"

RTaW
RealTime-at-Work

Contact: jorn.migge@realtimeatwork.com
C.T.O. RTaW