

Early-stage topological and technological choices for TSN-based communication architectures


Nicolas NAVET, University of Luxembourg

Josep VILLANUEVA, Groupe Renault

Jörn MIGGE, RealTime-at-Work (RTaW)



**2019 IEEE Standards Association (IEEE-SA)
Ethernet & IP @ Automotive Technology Day
September 24-25, 2019 | Detroit, Mi, USA**




Designing next-generation E/E architectures: Renault FACE service-oriented architecture



FACE : Future Architecture for Computing Environment




- Zonal EE Architecture
 - PCU : Physical Computing Unit
 - Intelligence
 - PIU : Physical Interface Unit
 - Analog I/O
 - Interface to legacy networks
- Service Communication (SOA)
- Ethernet Backbone
 - Mixed domains (Body, chassis, ADAS, ...)
 - Mixed safety constraints (QM, ASIL-B, ...)
 - Mixed Security levels
 - Mixed QoS requirements (C&C, Video, Audio, Reprog, ...)




MAIN GOAL : Scalability, capability to add new features

Which TSN protocols to meet requirements ? Which are the limits of the architecture ?





CAN Snapshot packets aggregate all data from CAN buses




Ps : Snapshot Period (1ms, 5ms, 10ms, ...)
 Tb : TSN Backbone transfer time
 Tec : Ethernet – CAN bridging time

$$Tgtw = Ps + Tb + Tec$$

CAN Snapshot: simplified and predictable gatewaying strategy

Simulation Process




New Use-Case?
Start over again, need to create a new custom suit!

Network configuration process might be the limiting factor for scalability !



Finding the limits of an E/E Architecture




New Use-Case?

Use pre-defined configuration



MAIN GOAL : Network configuration process ready for scalability






Early-stage design choices for TSN networks




Topological and technological choices




Design choices based on evidence at a time when all communication requirements are not known ?

1. Network dimensioning to add functions & services during car's lifetime?
KPI for network extensibility
2. Identifying and removing bottlenecks? KPI for resource congestion




Synthetic data captures what is known and foreseen about communication needs



Artificial data: all possible communication requirements

- ✓ Based on past vehicle projects and what can be foreseen for the current project
- ✓ Assumptions made on the streams and their proportion

- ✓ *Stream characteristics overall well known*
- ✓ *Stream proportion more uncertain → several scenarios may be considered*





- We have a candidate E/E architecture and a baseline traffic, the objective is to
1. Estimate the **max. # of SOME/IP services** that can be supported with each TSN protocol
 2. Identify and **remove architecture's bottlenecks**



Total network “capacity”: KPI of extensibility

Probability that the network will successfully meet the performance constraints of a given number of streams

KPI: “the network capacity is 200 flows at the 90% threshold when no other mechanism than priorities is used”



Command & Control (CC)	<ul style="list-style-type: none"> ✓ 32 streams, 256 to 1024 byte frames ✓ 5ms to 80ms period and deadlines ✓ Hard deadline constraints 	2/5
Audio Streams	<ul style="list-style-type: none"> ✓ 8 streams: 128 and 256 byte frames ✓ 1.25ms period and deadline ✓ deadline constraints (soft) 	1/10
Video streams	<ul style="list-style-type: none"> ✓ 3 streams (vision): 30x1400 byte frames every 33ms – deadline = 33ms ✓ 4 streams (ADAS): 15x1000bytes frames every 33ms – deadline = 10ms ✓ hard and soft deadline constraints 	1/5
Best-effort streams	<ul style="list-style-type: none"> ✓ Bulk data = from 64K to 1MB transfers, or ✓ 100ms periodic PDUs data, e.g. from CAN networks 	3/10

Scheduling Mechanisms	
<input checked="" type="checkbox"/> Without priorities	<input checked="" type="checkbox"/> User priorities
<input checked="" type="checkbox"/> Concise priorities	<input type="checkbox"/> Credit Based Shaping (CBS)
<input checked="" type="checkbox"/> PreShaping	<input checked="" type="checkbox"/> Preemption
<input checked="" type="checkbox"/> Time Aware Shaping (TAS)	

➤ Topology-stress-test (topology, traffic assumption, TSN policies)

Monte-Carlo simulation on synthetic networks

Create

1

Random yet realistic communication needs with increasing load

Configure

2


Configure generated networks

Evaluate

3

Check performance requirements by simulation and worst-case analysis

Gather statistics on the capacity of the architecture with each selected TSN protocol at the different load levels



Specifying characteristics of streams : example of a video stream class


The screenshot shows a configuration window for 'Traffic #1'. It features a 'PacketCharacteristics' section with a list of four classes: C&C (Priority=7, Weight=20), Audio (Priority=6, Weight=20), Video (Priority=5, Weight=20), and BE (Priority=4, Weight=40). The 'Video' class is selected. Below this, the 'Video' class configuration is shown with fields for 'ClassName*' (Video), 'Priority*' (5), and 'Weight*' (20). A sub-section for 'PacketCharacteristics' shows 'Period = [33.333 ms]' and 'Burst = [15,...,30]'. Further down, 'Weight*' is set to 20, 'NamePrefix' is empty, 'Afdx*' is unchecked, 'AllowedSenders' and 'AllowedReceivers' are dropdown menus, 'Period' is expanded to show 'Min*' (1000 byte), 'Max*' (1500 byte), and 'StepSize' (100 byte). 'Latencies' and 'ReceiversPerFrameFlow' are also visible.

Here 4 traffic classes,
including one video

30FPS camera
with an image
sent as a burst of
15 to 30 packets

Percentage of video
streams among all
streams

Packet size from
1000 to 1500bytes
by step 100




Case-study: Renault FACE architecture



Renault Ethernet prototype SOA

8 Physical Computing Units (PCU)

10 ECUs



#Nodes	33
#Switches	8
#streams	52
Services excluded	
Virtual Switch	1 with 4 VMs
Link data rates	100Mbit/s

15 Physical Interface Units (PIU) gateways to CAN and LIN buses

[RTaW-Pegase screenshot]

TSN QoS mechanisms considered

- #1 User priorities: priorities manually allocated to classes by designer according to criticality and deadlines
- #2 8 priority levels assigned to flows by *Concise Priorities* algorithm
- #3 User priorities + frame preemption for the top-priority traffic class
- #4 User priorities + Time-Aware-Shaping with exclusive gating for top-priority traffic class
- #5 AVB/Credit-Based-Shaper with SR-A and SR-B at the two top priority levels
- #6 User priorities + Pre-shaping: inserting “well-chosen” idle times between packets of segmented messages



Heterogeneous backbone traffic

Baseline traffic:
no services

Traffic Class	User Priorities	Type of traffic and constraints
TFTP	1	✓ TFTP (throughput constraints: 5Mbit/s and 9 Mbit/s).
Non-urgent Services & Short Files	3	✓ Services with medium latency constraints (30ms - 100ms)
Multimedia & ADAS	4	✓ Less urgent ADAS UC6.A.x (latency constraint: 33ms) ✓ Multimedia video UC6.B.x (latency constraint: 33ms)
Fusion & ADAS	5	✓ UC6.A.x ADAS Video (latency constraint: 15ms) ✓ UC7 Fusion (latency constraint: 10ms)
CAN snapshots & Urgent Services	6	✓ Services with short latency constraints (<30ms) ✓ UC8 CAN snapshot frames (2ms or 5ms)


Capacity of the network in terms of # of services ?

- ✓ SOME/IP traffic: both urgent and non urgent services
- ✓ Urgent (60%): periods from 5 to 30ms, deadlines = periods, size = 64bytes
- ✓ Non-Urgent (40%): periods from 30 to 100ms, deadlines = periods, size from 128 to 1500bytes



Scenario #1: CAN snapshots with a 2ms deadline

Capacity of the network at the 90% threshold is 370 services with Preshaping




- ✓ Priority alone is not enough, shaping is needed
- ✓ “CBS” quickly fails because the shorter CAN snapshot deadlines cannot be met with video at highest priority
- ✓ Preemption detrimental because its overhead decreases TFTP throughput



Scenario #2: CAN snapshots with a 5ms deadline

Capacity of the network at the 90% threshold is 435 services with CBS



- ✓ 5ms CAN snapshots deadlines less limiting for CBS
- ✓ "CBS" > Pre-shaping because it "repacketizes" the data into smaller packets
- ✓ Preemption still detrimental
- ✓ More than 5 priority levels not helpful



Where are the bottlenecks? Which traffic classes ?
Which constraints? Where in the network?



Bottleneck traffic class under CBS

Metric: % of the non-feasible configurations for which at least one stream of a traffic class does not meet its performance constraints




TFTP



Non-urgent Services
& Short Files



Multimedia
& ADAS



Fusion & ADAS



CAN snapshots
& Urgent Services

CBS
CAN deadline=2ms




Limiting constraint is
TFTP throughput



CBS
CAN deadline=5ms



Bottleneck traffic class under TAS




- ✓ Limiting traffic class → limiting constraints → missing TSN mechanisms - here shaping video streams would improve TFTP throughput
- ✓ The bottleneck traffic class may vary depending on the # of flows → use the # of flows corresponding to a fixed probability threshold – here 90%




Identifying bottleneck resources

Metric: contribution of a “hop” to the overall latency of the streams that are not meeting their performance requirements



Improvement: duplicating link SWa → SW1 and balancing load




Capacity of the network at the 90% threshold is now 670 services with CBS

+54% capacity improvement with a single 100Mbps link duplication!

- User Priorities
- Concise Priorities
- Preemption
- Time Aware Shaping
- Preshaping
- Credit Based Shaping






Conclusion and a look forward



Contributions




- 1** Design Space Exploration with artificial yet realistic data to support architectural and technological choices → Topology-Stress-Test in RTaW-Pegase
- 2** KPIs to 1) evaluate the evolutivity of a network and 2) measure resources congestion → link load is insufficient with performance constraints
- 3** Tool-supported approach to identify which performance constraints is the limiting factor and where the bottlenecks are in the network
- 4** On the FACE E/E architecture duplicating a single 100Mbit/s bottleneck link allows supporting 54% more services!
- 5** No “one fits all” TSN scheduling solution for TSN backbones, need the combined use of several TSN mechanisms → tool support helps keep up with complexity



A look forward: towards E/E architecture synthesis

- 1 Extensions:** better results explanation and support for combined TSN mechanisms (e.g., priorities+TAS+CBS+preemption), task allocation on ECU/cores/cloud
- 2 Propose incremental changes that allow a “minimum gain”** (# of flows, costs, safety, ...)




Sol A: duplicate 100Mbps link and balance load

Sol B: switch to 1Gbps

Sol C: increase switch memory by 150Kb

....



Thank you for your attention!



Questions ?

