Early-stage Bottleneck Identification and Removal in TSN Networks

Nicolas NAVET, Uni. Luxembourg / Cognifyer.ai

Hoai Hoang BENGTSSON, Volvo Car Corporation

Jörn MIGGE, RealTime-at-Work (RTaW)

Automotive Ethernet Congress
February 12-13, 2020 | WESTIN Grand Munich

Early-stage design choices on virtual platforms

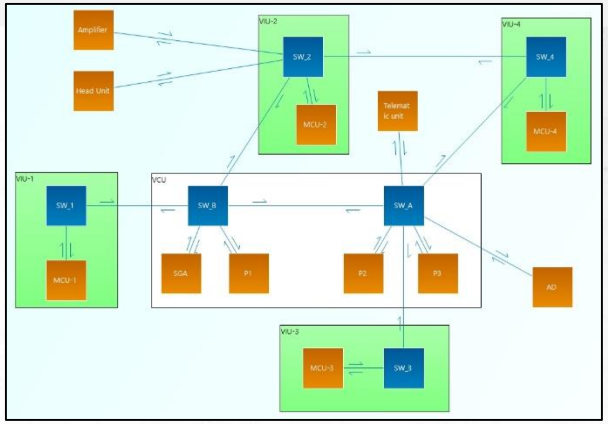

Candidate solution

1. Robustness
2. Total capacity
3. Bottlenecks
4. Cost-optimize
5. Reliability

1. Sensitivity to future additional low-priority traffic?
2. Quantify network capacity wrt TSN technological options
3. Identify bottlenecks in E/E architecture and remove them
4. Cost-optimize by reducing link speeds and # of ECUs
5. Assess and optimize communication reliability

Solution Refinement

Designing next-generation E/E architectures: Volvo Centralized and Zone-based Architecture

VIU : gateway from the Mechatronic Rim to the Core Network

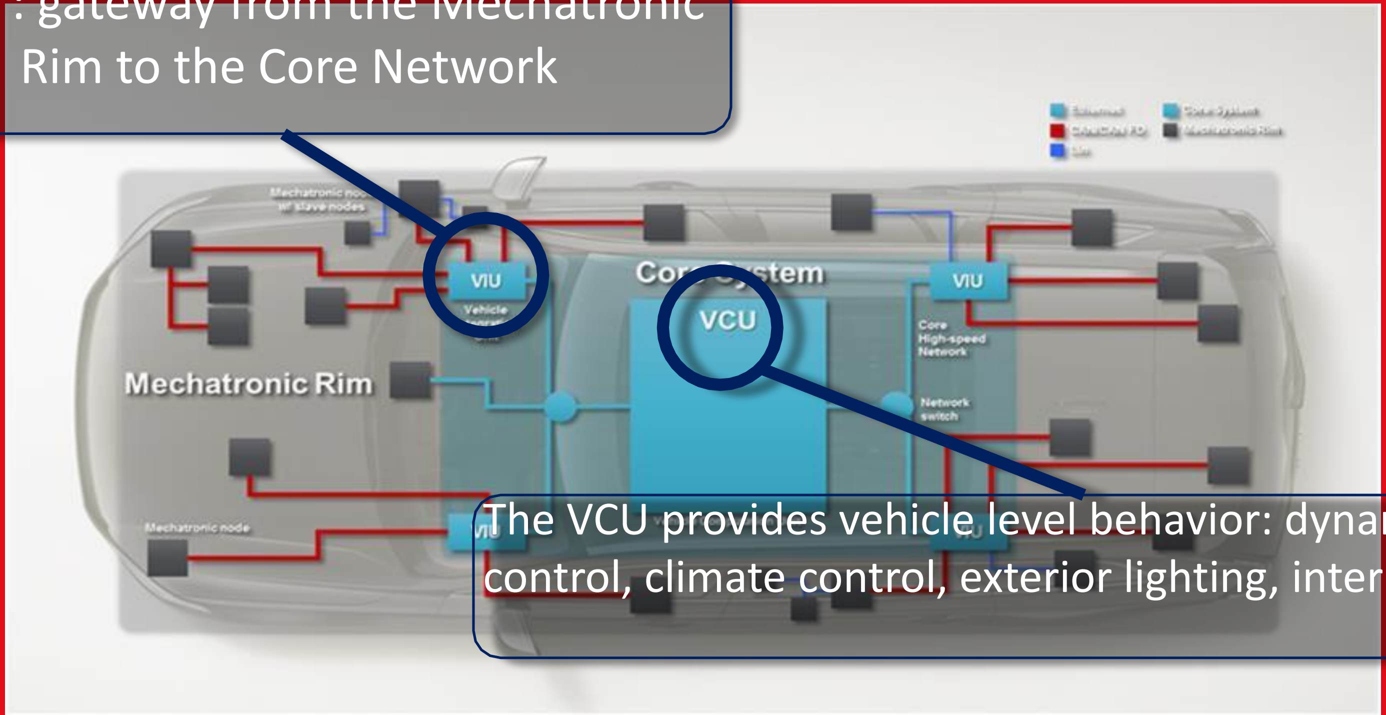

The VCU provides vehicle level behavior: dynamics, propulsion control, climate control, exterior lighting, interior lighting, ...

Volvo Core TSN Network

4 Vehicle Interface Units (VIU)

10 ECUs / CPUs incl. Head-Unit, Autonomous Driving (AD), Telematic Unit, …

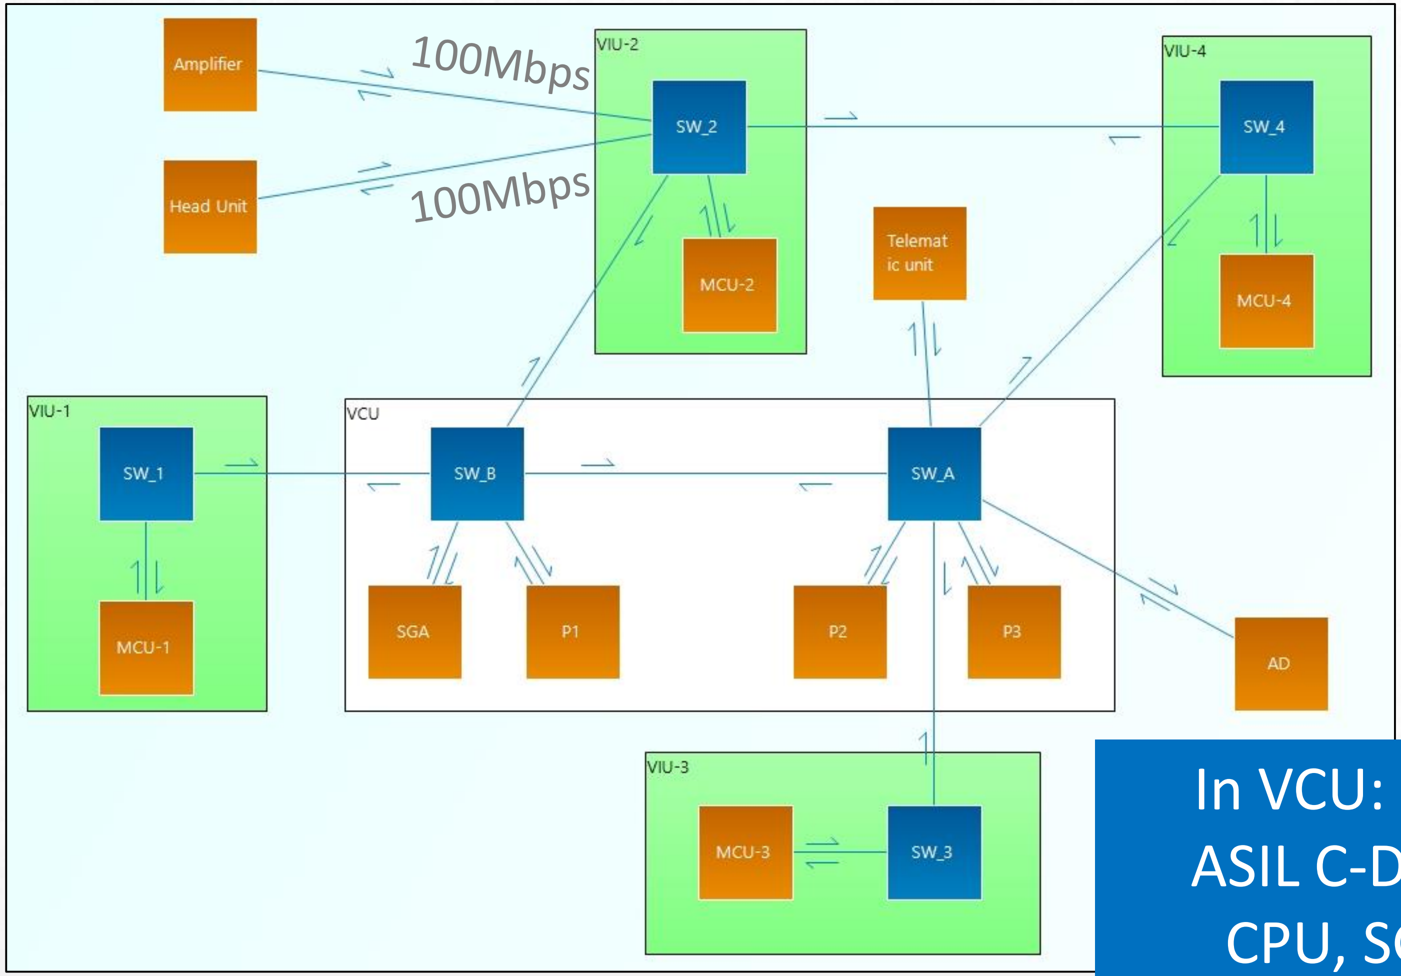

[RTaW-Pegase screenshot]

#Nodes	10
#Switches	6
#control streams (baseline traffic)	25 deadlines: 200us to 2ms
# TFTP streams	6 throughput: 5Mbps
Link data rates	1 Gbps (16 links) 100Mbps (2 links)

In VCU: P1 and P2 run redundant ASIL C-D functions, P3 is a high-perf. CPU, SGA is the Security Gateway

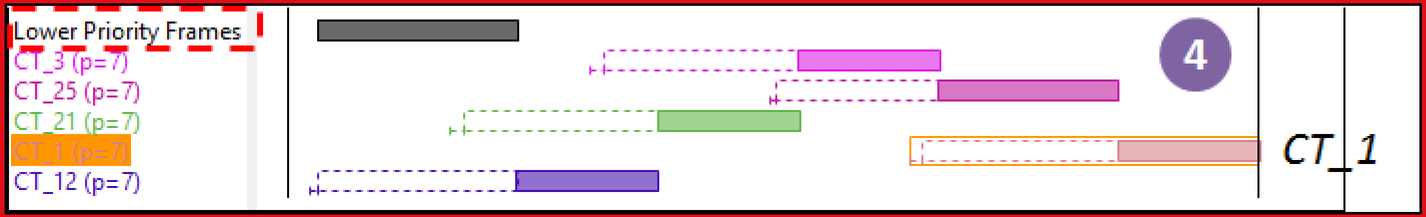

Design Question #1 : Sensitivity to future, unknown yet, Low-Priority Streams

UC#1: Impact of Additional Low-Priority Traffic

Design questions:

- ✓ Robustness: *can we ignore lowest priority streams at design time, without jeopardizing timing constraints for the rest of traffic ?*
- ✓ Evolutivity: *can we can safely consider adding functions inducing low-priority streams at a later evolution without caring about their exact characteristics ?*

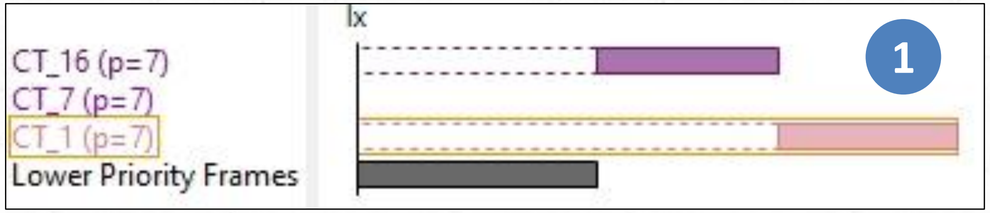

CT_1

- ✓ Blocking factor: a packet can be delayed at each hop by one low-priority frame being transmitted → increase delays and memory usage
- ✓ Higher- or same-priority packets can arrive during blocking factor (e.g., purple packet above) increasing thus further the delay…
- ✓ "Low-priority flooding" in RTaW-Pegase → scenario with worst-case interference from low-priority frames

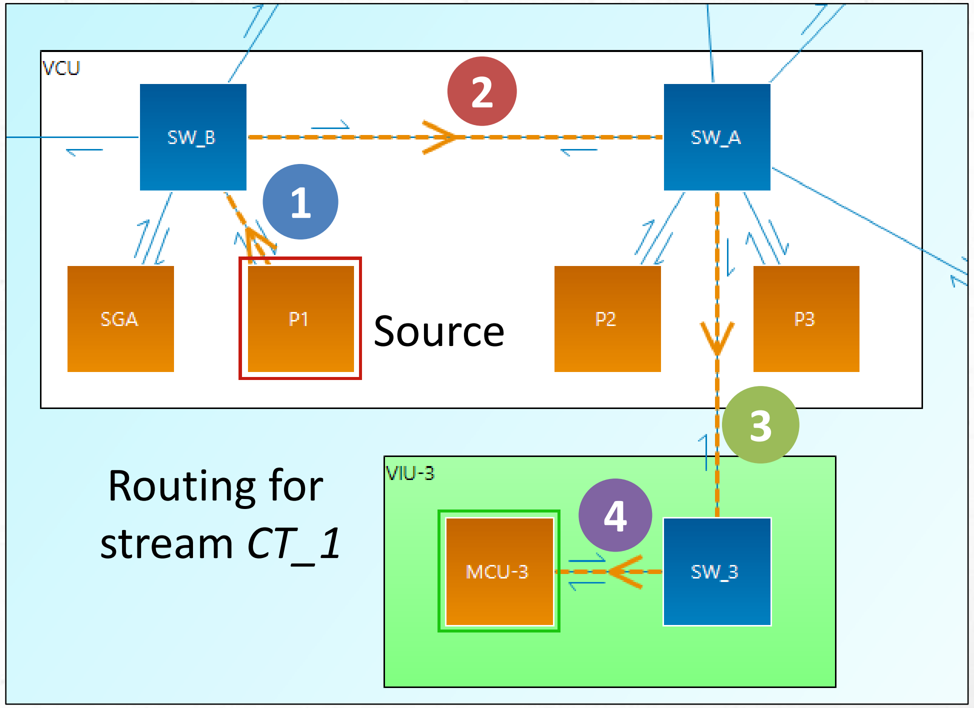

With / without low-priority frames interference: example

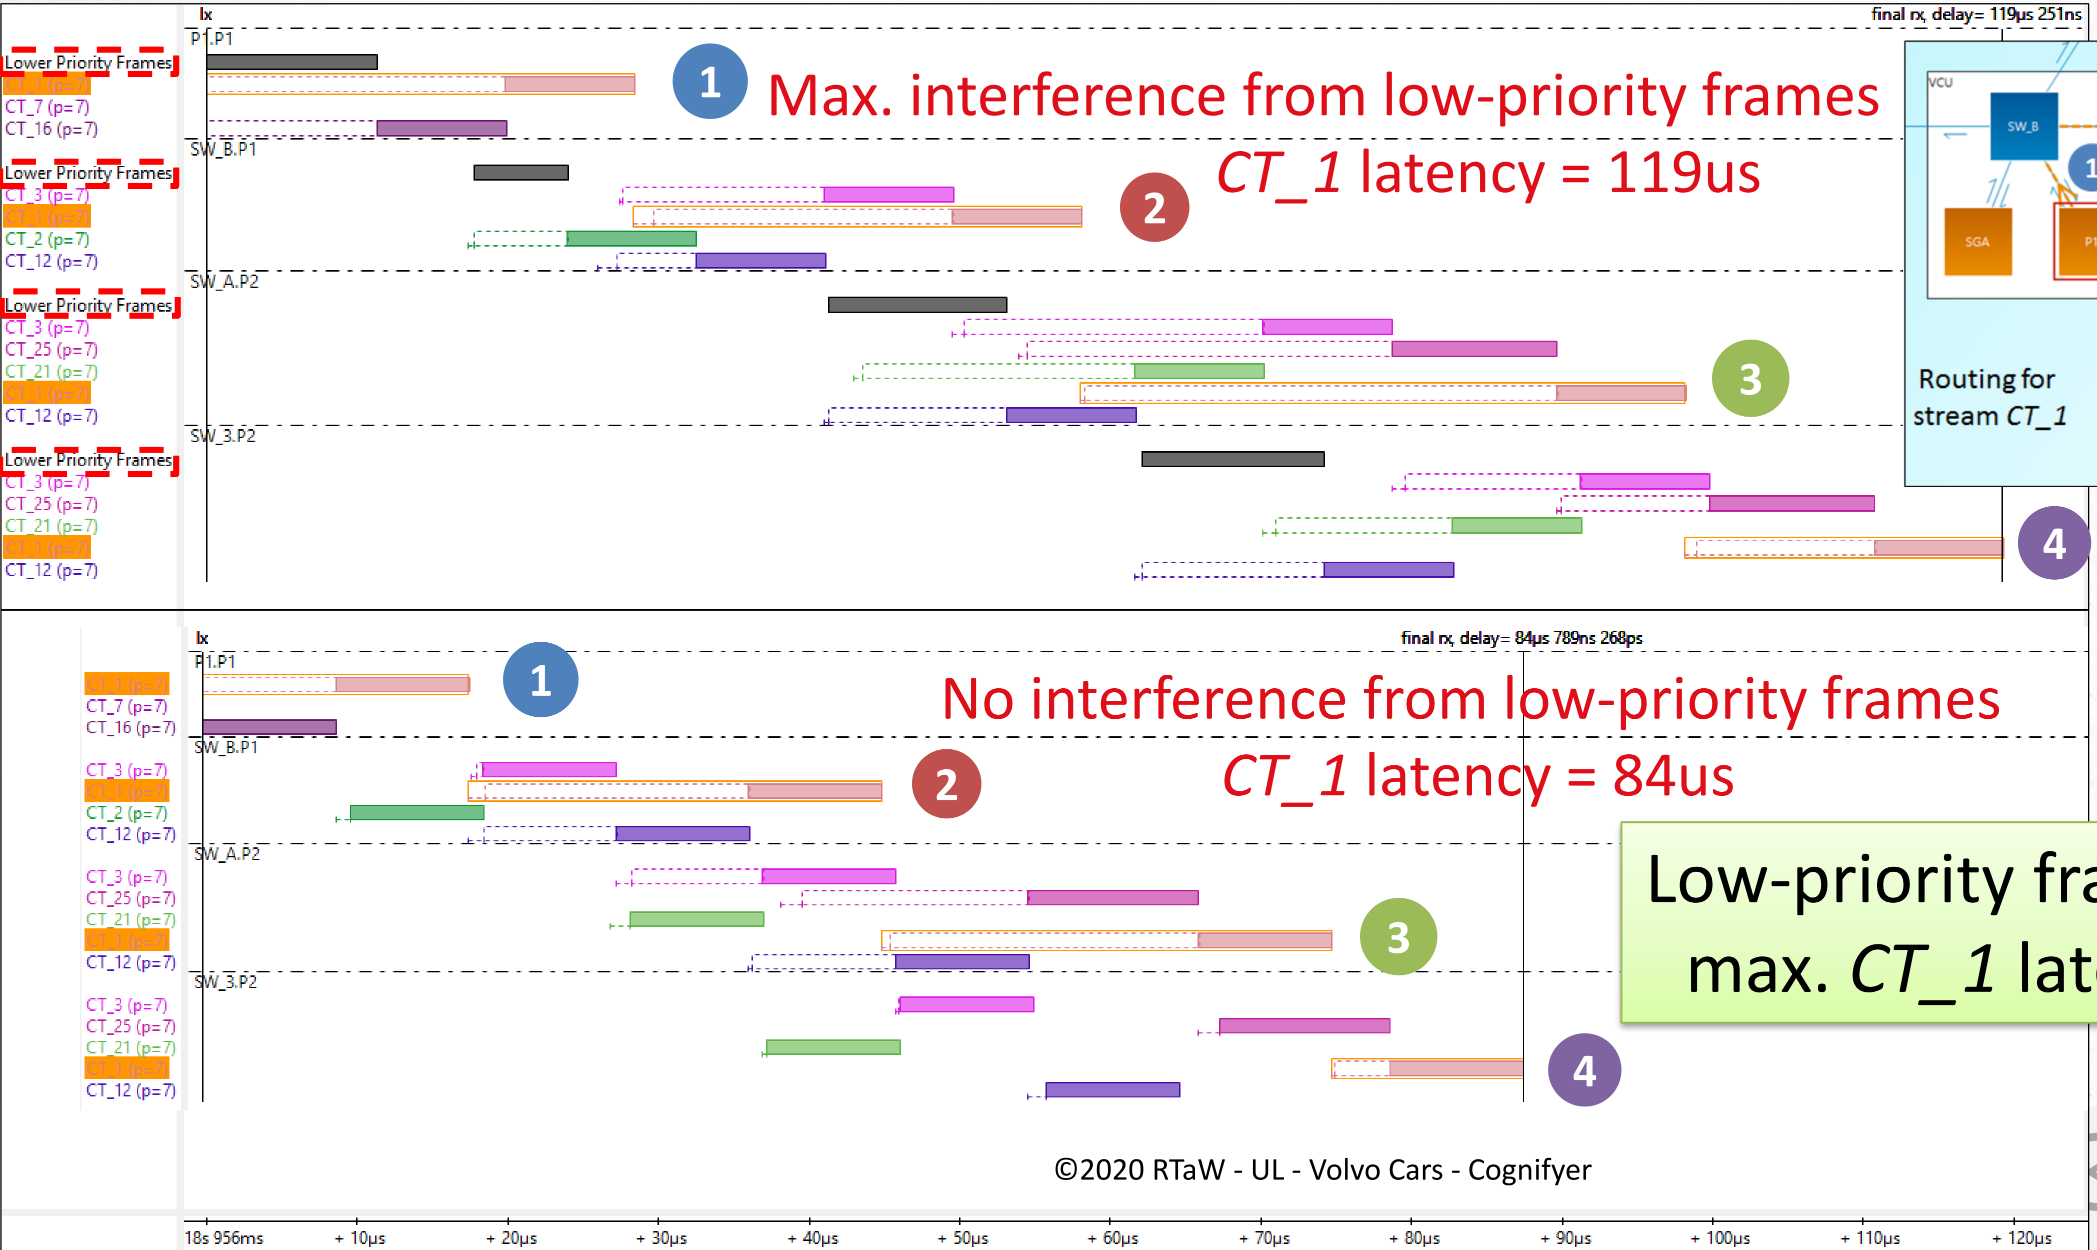

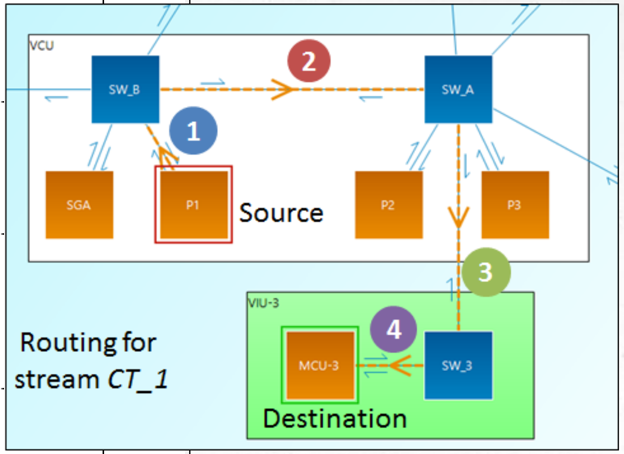

Low-priority frames increase max. *CT_1* latency by 40%

Latencies with TSN policies: W/O "low-priority flooding"

Worst-case latencies (ms)

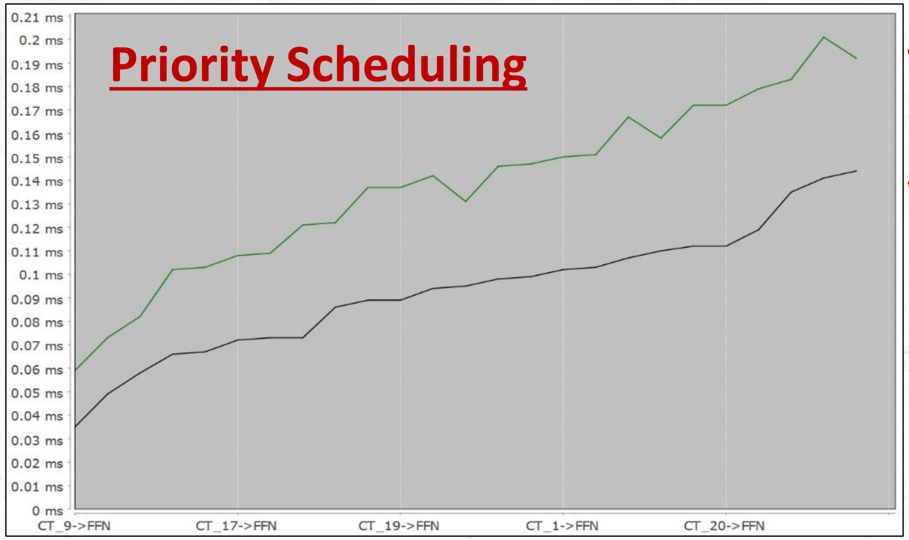

With

Without

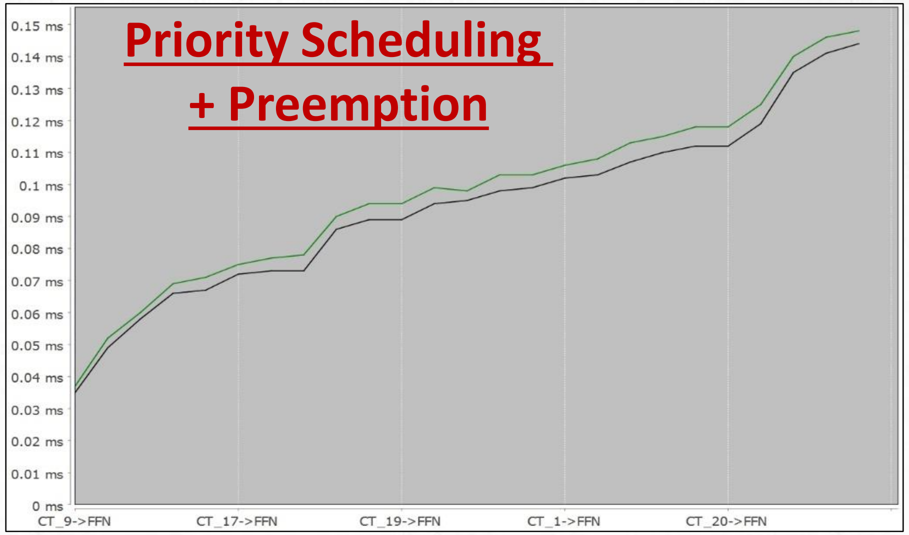

Command & Control streams

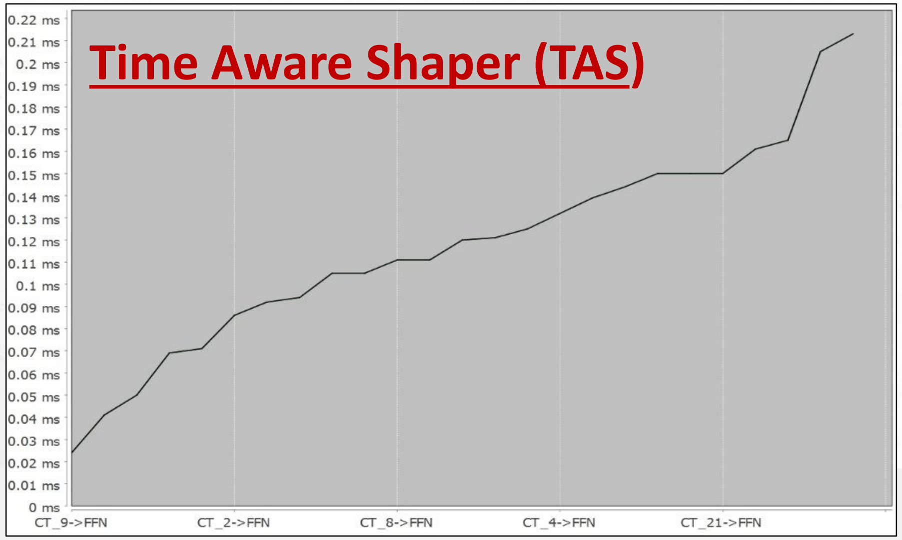

- ✓ Priority: communication latencies severely degraded by low-priority flooding: +49% on average
- ✓ Priority + Preemption: additional delays brought by low-priority frames reduced: +5% on average
- ✓ TAS: exact same performance W/O low-priority frames

→ TAS, and to a lesser extent preemption, are future-proof wrt adding low priority streams

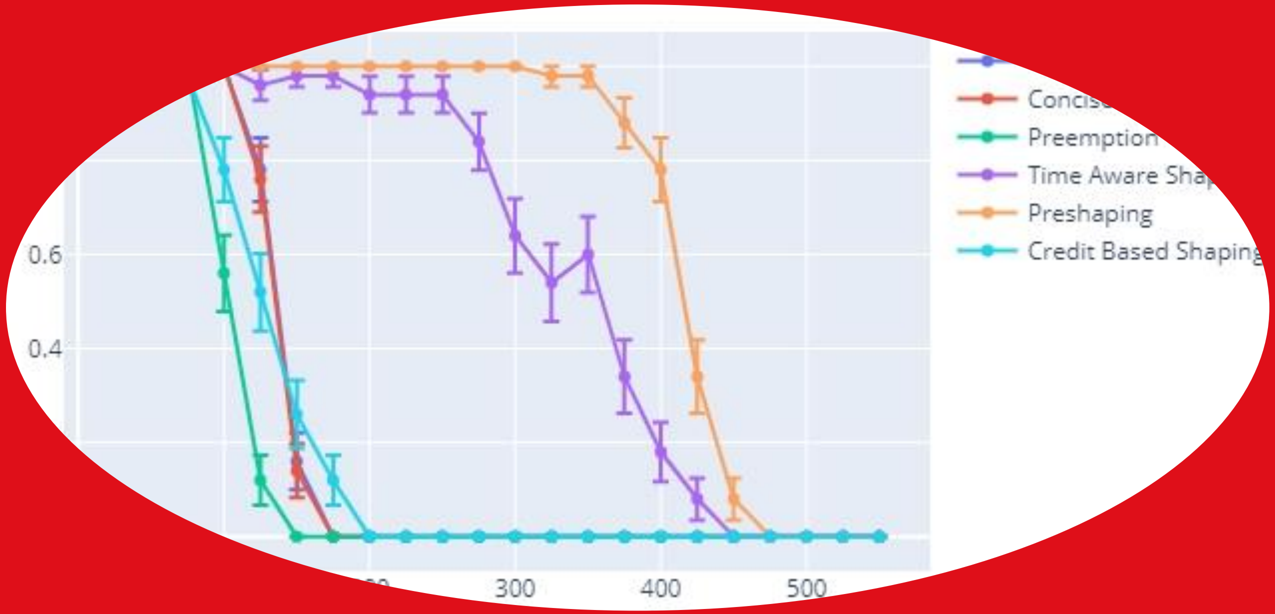

Design Question #2 : Evaluating the Total Capacity of the TSN Network with different TSN Technological Options

Total network "capacity": KPI of extensibility

How many streams will be successfully scheduled by the network with a given probability ("assurance level") ?

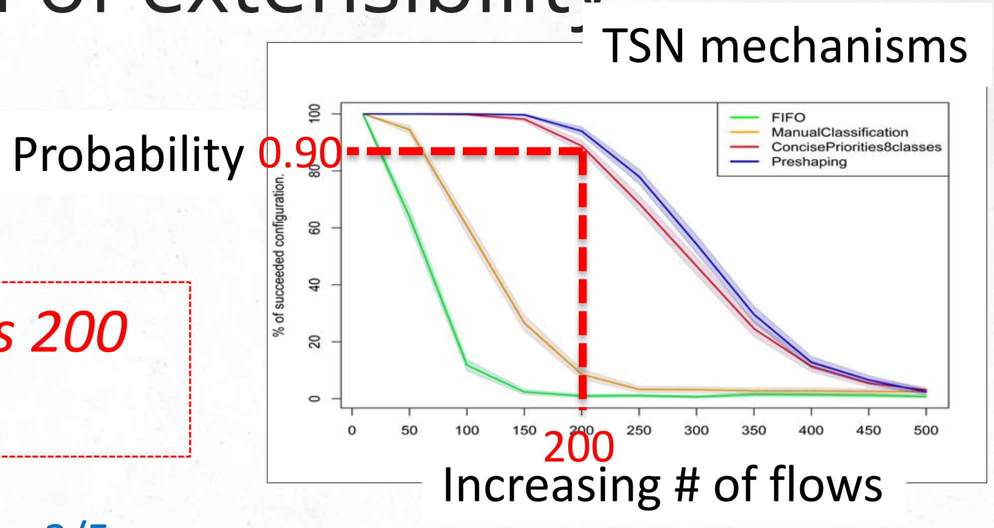

E.g. "With priority scheduling, the network capacity is 200 flows at the 90% assurance level"

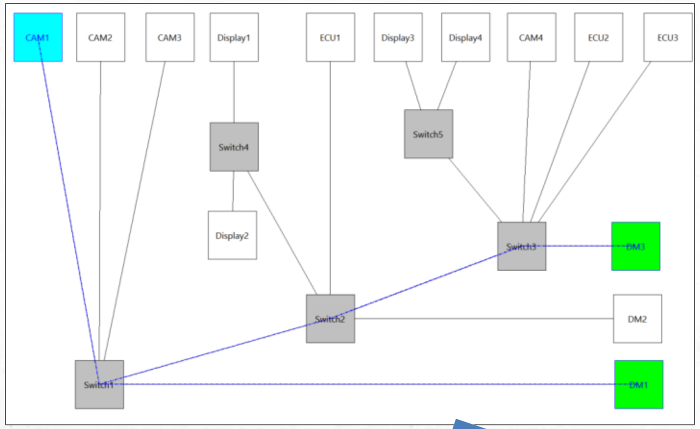

Command & Control (CC)	✓ 32 streams, 256 to 1024 byte frames ✓ 5ms to 80ms period and deadlines ✓ Hard deadline constraints	2/5
Audio Streams	✓ 8 streams: 128 and 256 byte frames ✓ 1.25ms period and deadline ✓ deadline constraints (soft)	1/10
Video streams	✓ 3 streams (vision): 30x1400 byte frames every 33ms – deadline = 33ms ✓ 4 streams (ADAS): 15x1000bytes frames every 33ms – deadline = 10ms ✓ hard and soft deadline constraints	1/5
Best-effort streams	✓ Bulk data = from 64K to 1MB transfers or ✓ 100ms periodic PDUs data, e.g. from CAN networks	3/10

Scheduling Mechanisms

- ☑ Without priorities
- ☑ User priorities
- ☑ Concise priorities
- ☑ PreShaping
- ☐ Credit Based Shaping (CBS)
- ☑ Time Aware Shaping (TAS)
- ☑ Preemption

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

Topology Stress Test (TST): Monte-Carlo Simulation on Synthetic Networks

Create

1 Random yet realistic network configurations over the load range

Configure

2 Configure generated networks for all selected TSN mechanisms

Evaluate

3 Check performance requirements by simulation and worst-case analysis

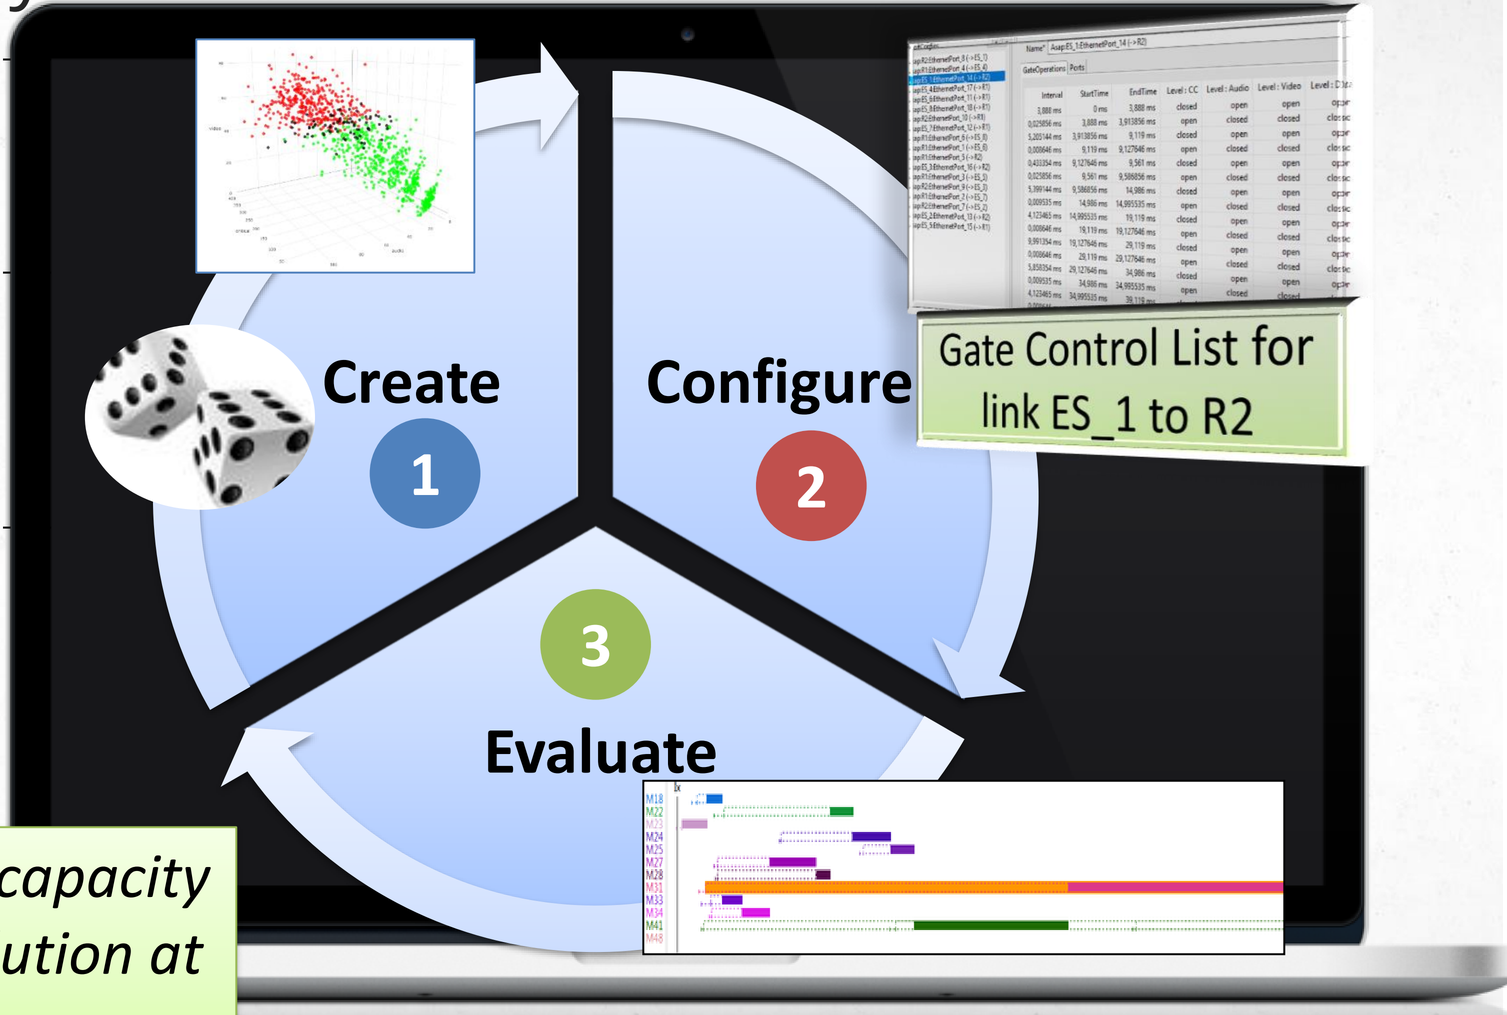

Collect statistics on the total network capacity with each selected TSN scheduling solution at the different load levels

TSN QoS Mechanisms Considered

- **#1** User priorities: **4 traffic classes** created by designer according to criticality and deadlines of flows
- **#2** Flows assigned to max. **8 traffic classes** 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** User priorities + **AVB/Credit-Based-Shaper** with SR-A and SR-B for two top-priority traffic classes
- **#6** User priorities + **Pre-shaping**: inserting "well-chosen" idle times between packets **for all segmented messages**

Traffic on the core TSN network

- ✓ Baseline traffic: 6 TFTP infotainment streams with 5Mbit/s throughput constraint between head-unit, telematic unit and security gateway at the lowest priority level
- ✓ Traffic added on top of baseline traffic:

Traffic Class	User Priorities	Type of traffic and constraints	Proportion
Infotainment	4	✓ Audio AVTP/IEC61883 – period: 2.5ms, (latency constraint: 2ms) ✓ Video AVTP/H.264 – 30 FPS, segmented (latency constraint: 33ms)	**8%**
Fusion & ADAS video	5	✓ ADAS video – 30 FPS streams, segmented (latency constraint: 10ms) ✓ Fusion data – Period: 50ms, segmented (latency constraint: 10ms)	**22%**
Command & Control	6	✓ 6 distinct types of streams ✓ Periods from 500us to 5ms, deadlines from 200us to 2ms	**70%**

How many such streams can be added?
How long will the architecture be able to evolve and support new functions?

Preliminary analysis : % Overloaded Networks VS # of flows

Overloaded network = the load of one link or more is higher than 100%
→ no TSN policy can meet timing constraints

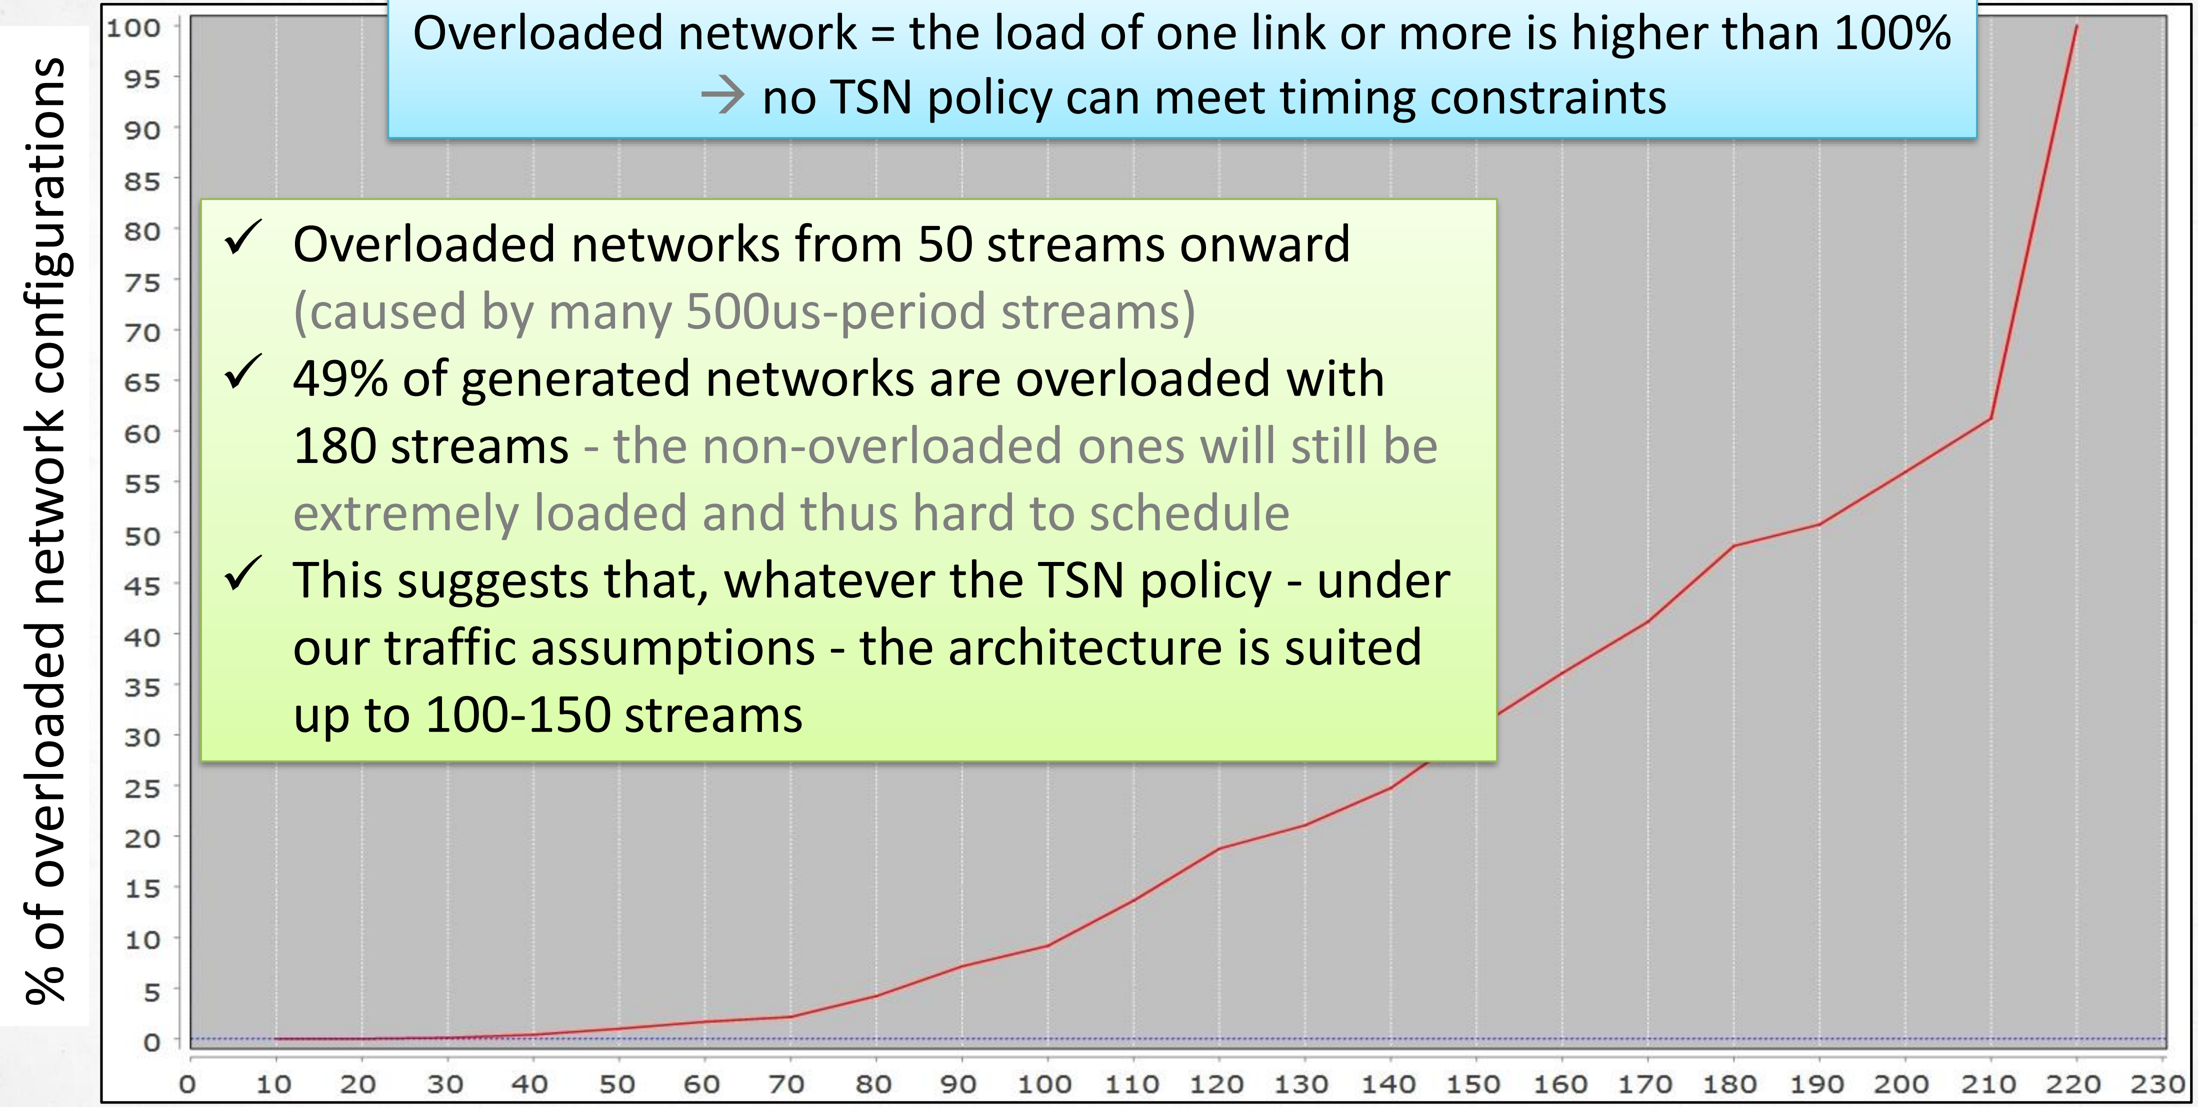

- ✓ Overloaded networks from 50 streams onward (caused by many 500us-period streams)
- ✓ 49% of generated networks are overloaded with 180 streams - the non-overloaded ones will still be extremely loaded and thus hard to schedule
- ✓ This suggests that, whatever the TSN policy - under our traffic assumptions - the architecture is suited up to 100-150 streams

\# of flows from 10 to 220

Network Capacity for ≠ TSN QoS options

Capacity of the network at the 90% assurance level is 100 flows with Concise Priorities (CP) (8 priority levels – no shaping)

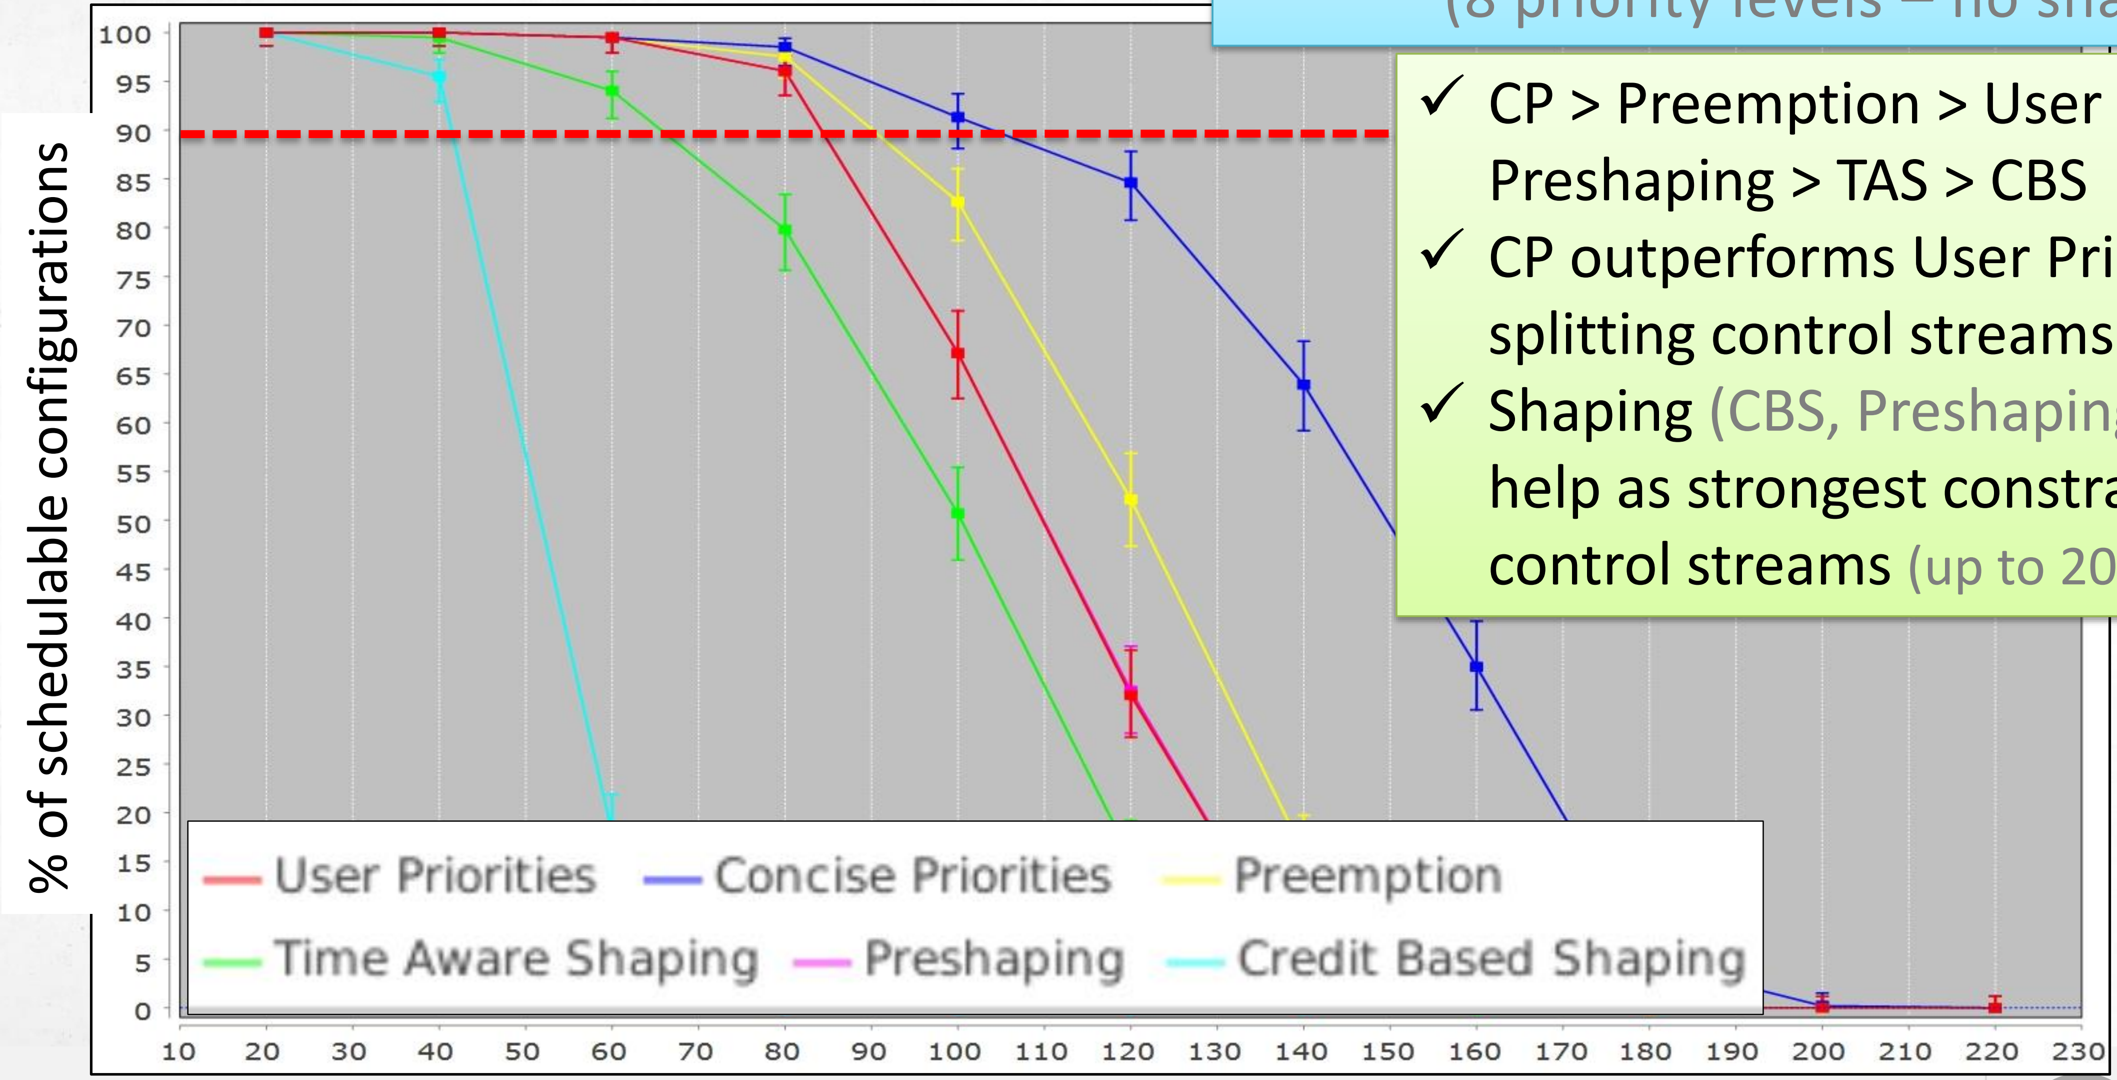

# of flows from 20 to 220

- ✓ CP > Preemption > User Priorities = Preshaping > TAS > CBS
- ✓ CP outperforms User Priorities by splitting control streams into 2 classes
- ✓ Shaping (CBS, Preshaping) does not help as strongest constraints are on control streams (up to 200us deadline!)

Statistical Guarantees on Network Capacity

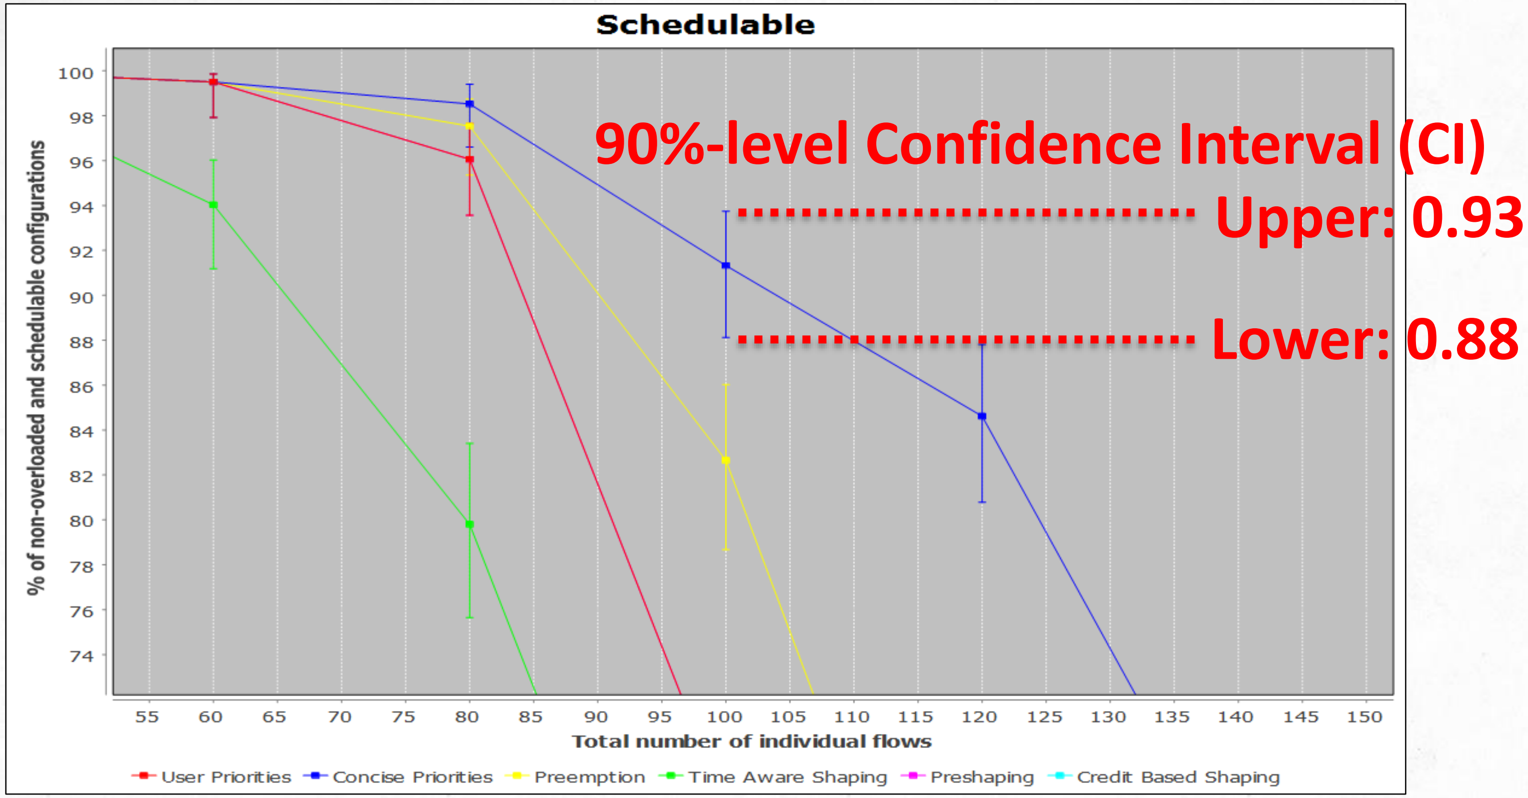

"The probability that the true network capacity is within the confidence Interval [0.88,0.93] is 90% (=confidence level)"

The # of synthetic networks analysed per load level allows to control the size of the CI, e.g.:
250 config. → CI size < 0.1 1000 config. → CI size < 0.05 4000 config. → CI size < 0.025

Design Question #3 : Bottlenecks in the architecture ? 2-step approach

1) Bottleneck traffic class analysis: constraint violations in which traffic class? Which additional TSN mechanisms would help ?

2) Bottleneck resources analysis: which resources in the network contribute the most to latencies ? How to improve on that ?

Bottleneck Traffic Class: illustration with Preemption

Nbr frames	Not Schedulables	Infotainement	Control Traffic	BE	ADAS Video / Fusion
20	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
40	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
60	0.5 %	0.0 %	0.0 %	0.0 %	0.0 %
80	2.5 %	0.0 %	100.0 %	0.0 %	0.0 %
100	17.4 %	0.0 %	100.0 %	0.0 %	0.0 %
120	47.9 %	0.0 %	98.7 %	1.3 %	0.0 %
140	84.1 %	0.0 %	98.7 %	12.8 %	0.0 %

Considering networks with 120 flows scheduled under User Priorities + Preemption:

- ✓ on average 47.9% will not meet their timing constraints
- ✓ 98.7% of the non-feasible networks have at least one control stream that does not meet its performance constraint (deadline)
- ✓ 1.3% of the non-feasible networks have at least one Best-Effort stream (TFTP here) that does not meet its performance constraint (throughput)

→ Bottleneck traffic class: Control traffic

→ More priority levels would be beneficial to support mixed deadlines in the class

→ This explains why Concise Priorities outperforms User Priorities + Preemption

Let's focus on the best TSN QoS solution: Concise Priorities

Bottleneck traffic class analysis: constraint violations across all traffic classes → adding TSN mechanisms would not help here - gains may come from removing bottlenecks resources

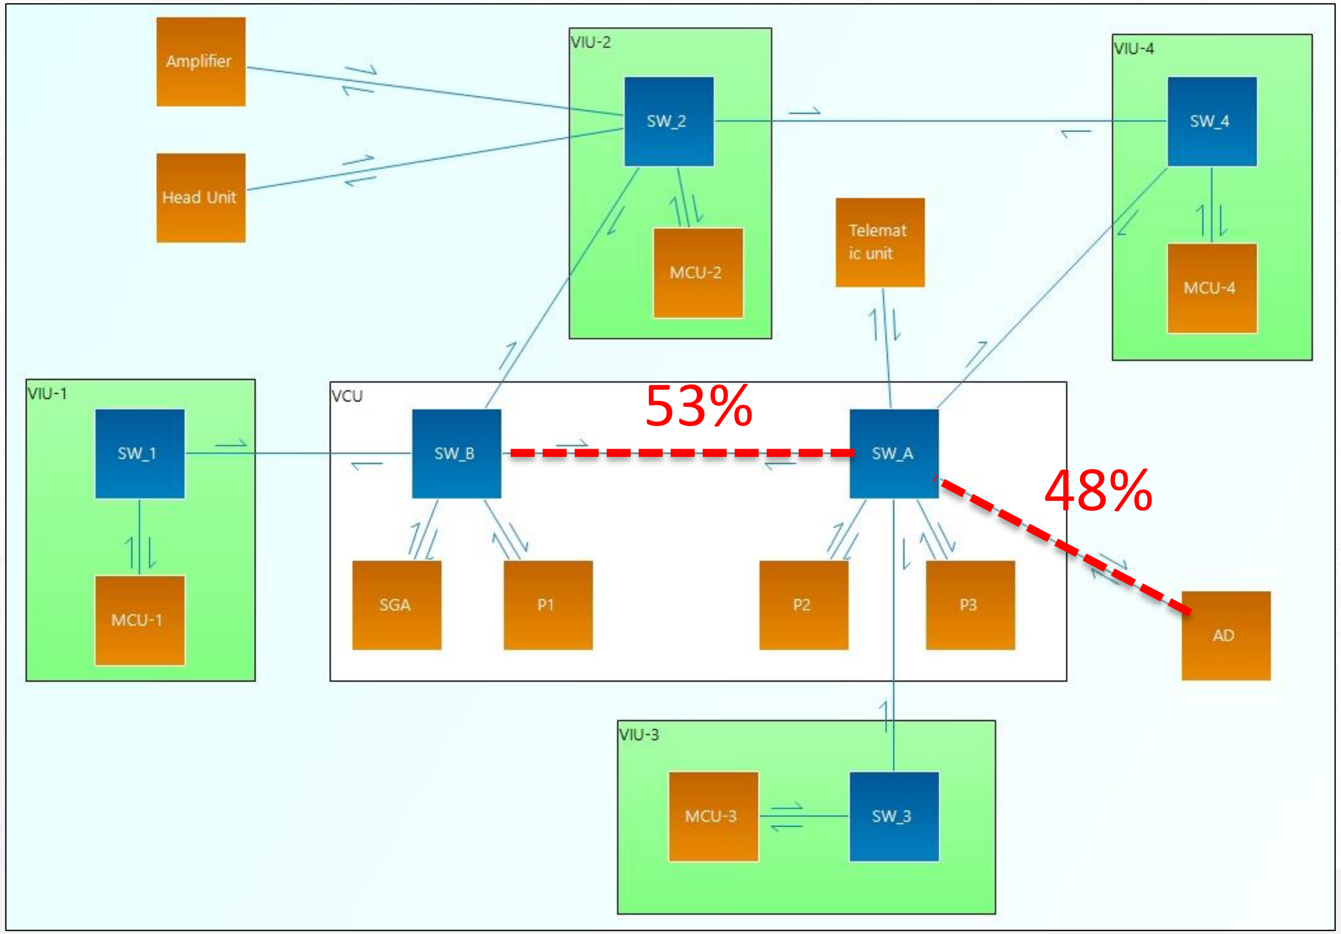

Metric for bottleneck resource analysis: contribution of a "hop" to the overall latency of the streams that are not meeting their performance requirements

2 largest contributors to latencies:
- *link SW_A → SW_B : 53%*
- *link AD → SW_A : 48%"*

Improving the Architecture

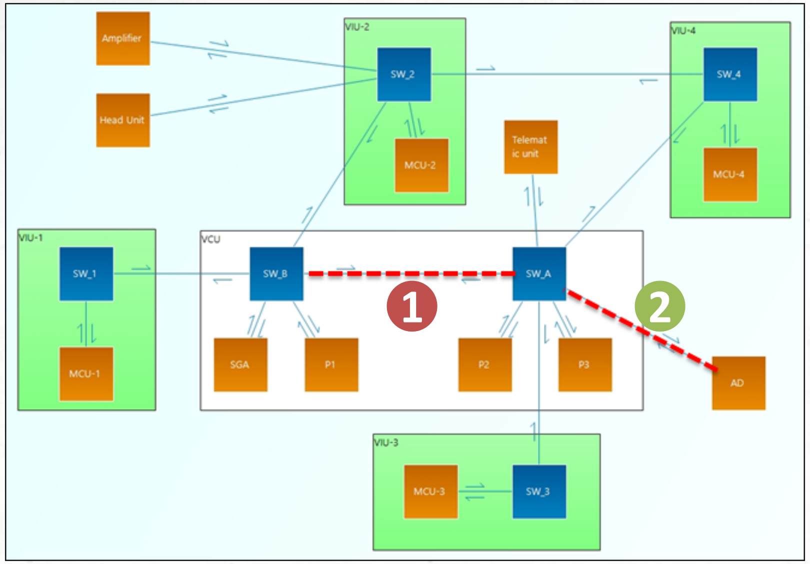

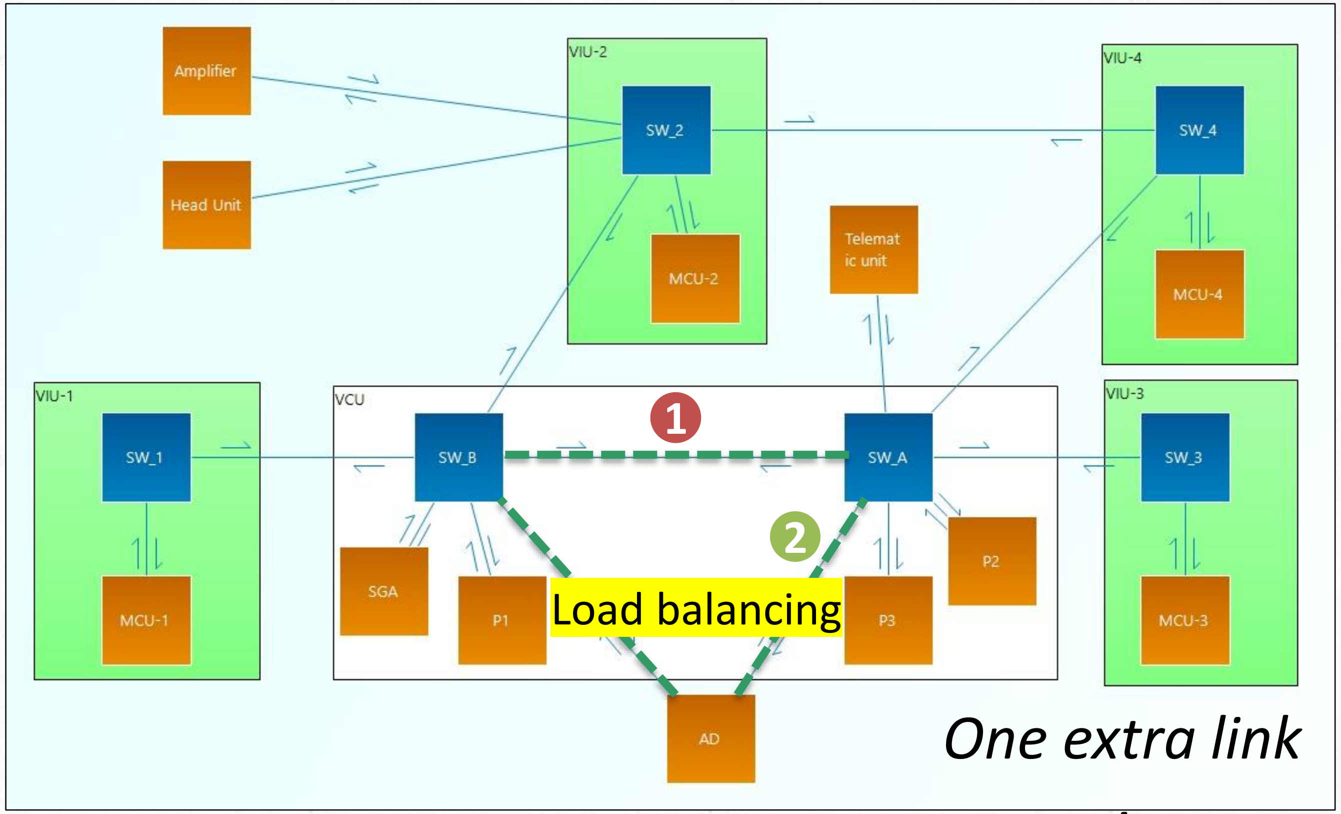

+

One network interface

- ✓ Reduces traffic on SW_A ↔ SW_B 1
- ✓ Reduces traffic on AD → SW_B 2
- ✓ May be used to replicate AD streams

Improved E/E Architecture: Total Capacity with Concise-Priorities

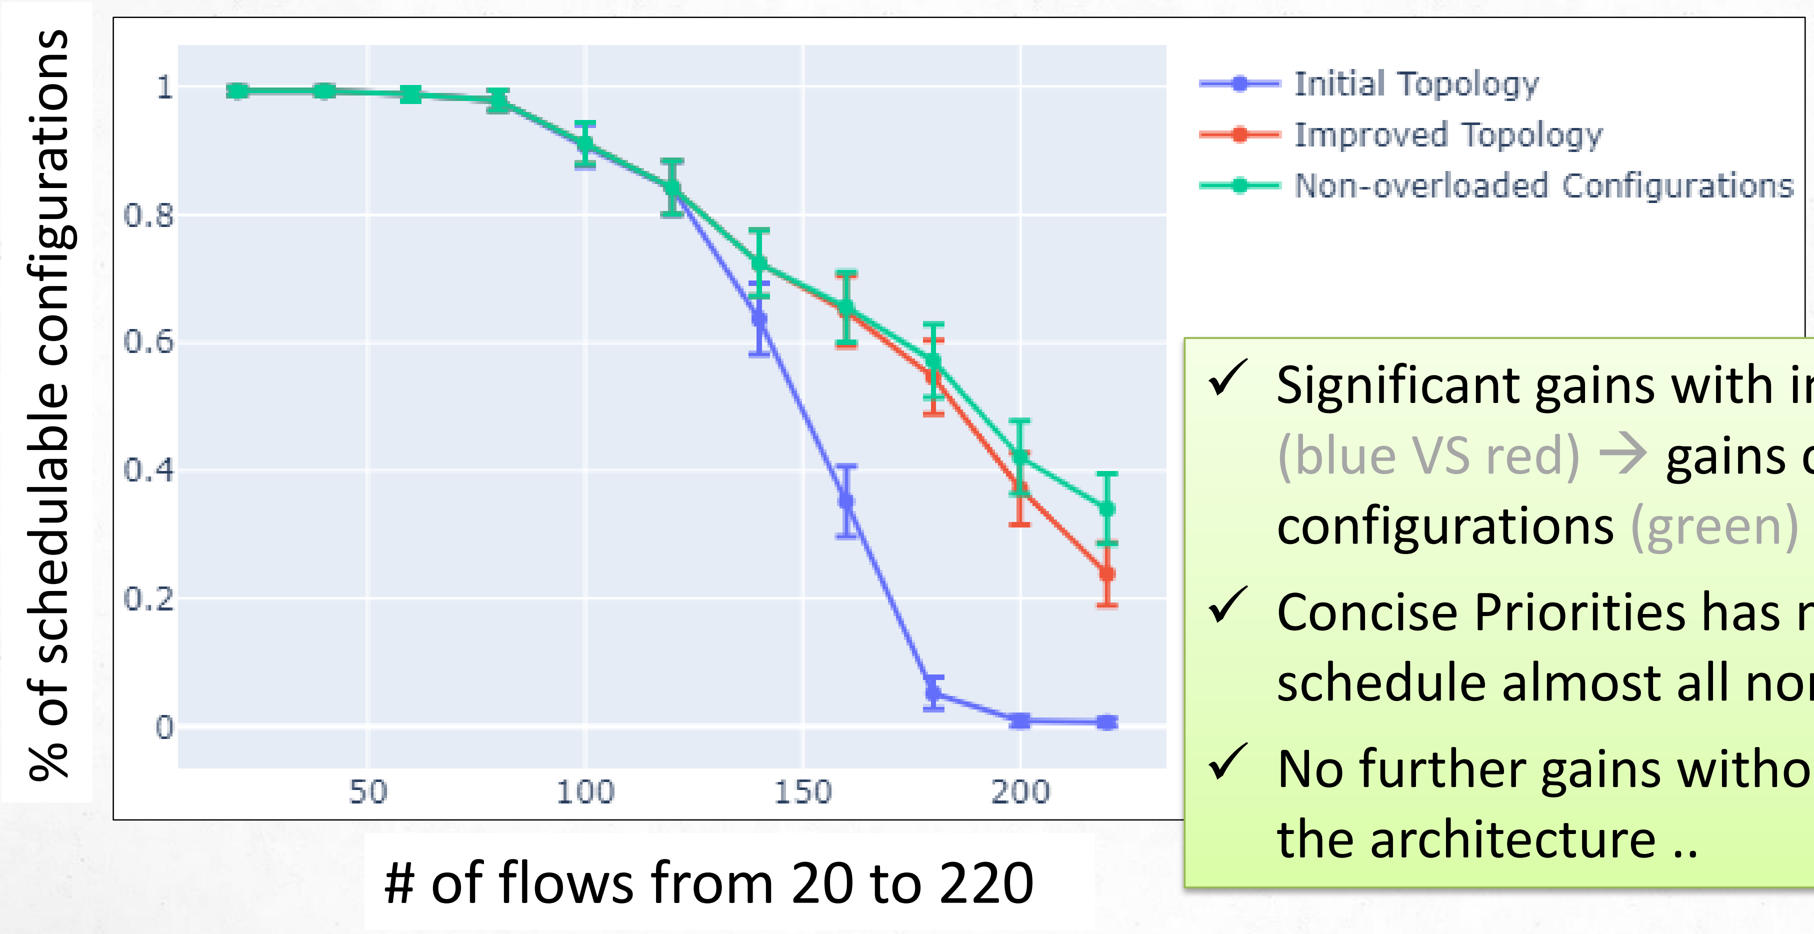

- ✓ Significant gains with improved E/E architecture (blue VS red) → gains capped by overloaded configurations (green)
- ✓ Concise Priorities has maximal efficiency as it can schedule almost all non-overloaded configurations
- ✓ No further gains without more drastic changes to the architecture ..

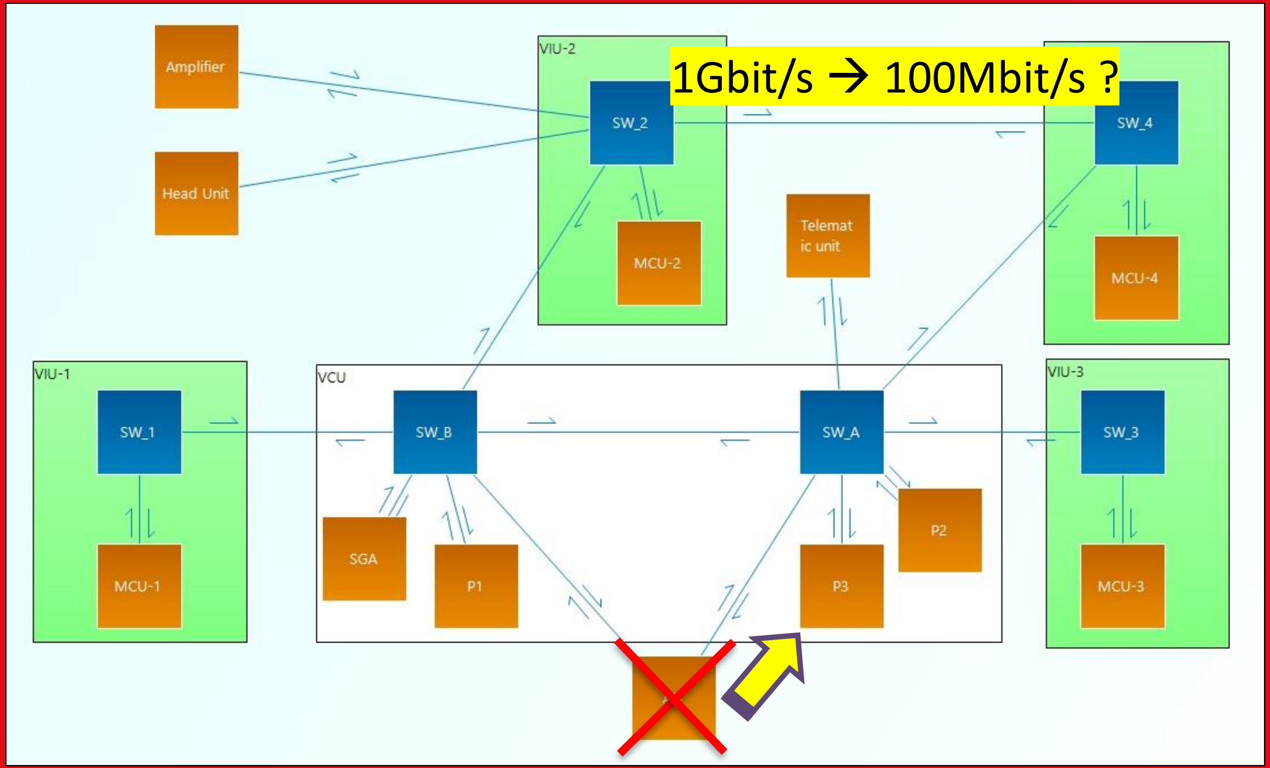

Design Question #4 : Ways to reduce costs with “acceptable” loss in network capacity?

1) Reduce link speeds: from 1GBbit/s to 100Mbit/s
2) Remove ECUs by relocating their functions into other ECUs

Cost Reduction: from 1Gbit/s to 100MBit/s

- ✓ Objective: reduce the speed of as many links as possible with a loss in network capacity of max. 15%
- ✓ Baseline capacity: 160 streams (65% assurance level) with Concise priorities → cost-optimized architecture must support more than 136 streams (at 65% assurance level)

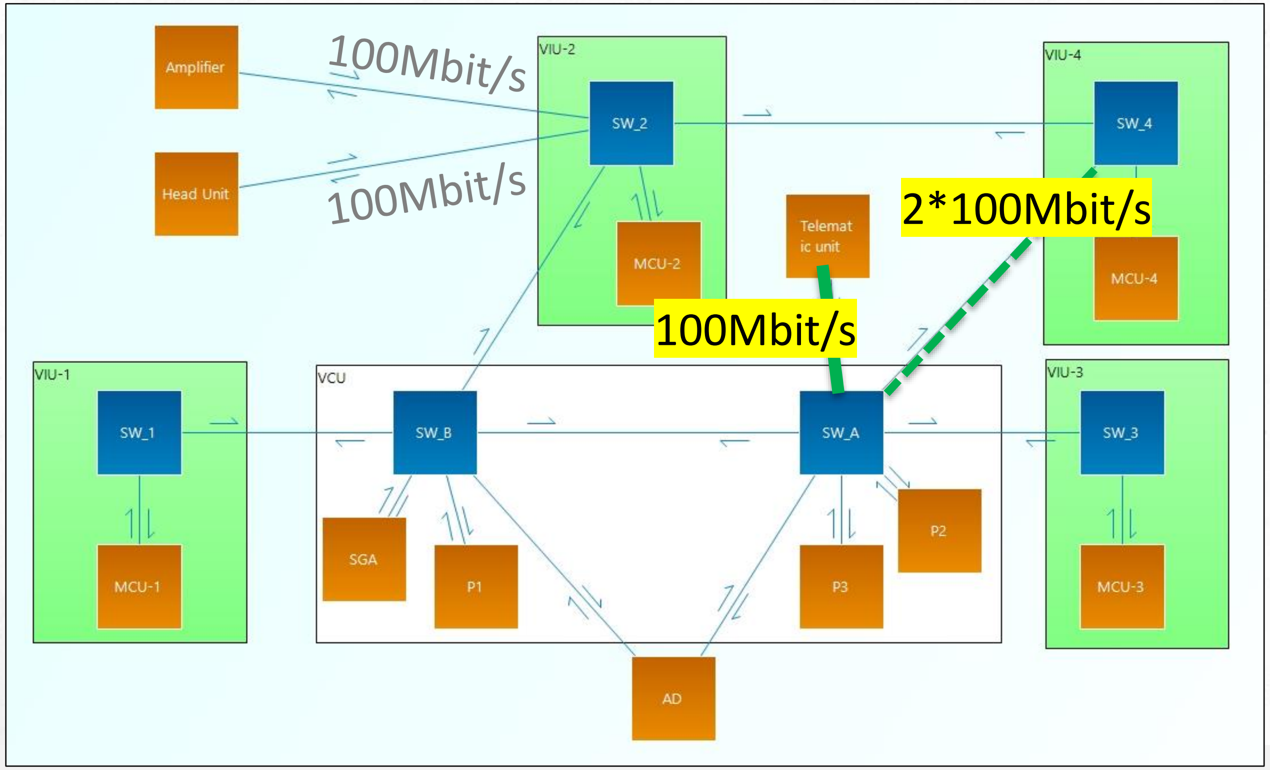

- ✓ 1Gbit/s link SW_A ↔ Telematic Unit can go down to 100Mbit/s
- ✓ 1Gbit/s Link SW_A ↔ SW can be replaced by two 100Mbit/s links
- ✓ Speed of both links can be reduced at the same time
- ✓ **1Gbit/s technology is required for all but 4 links in this architecture**

Cost Reduction: from 1Gbit/s to 100MBit/s technology

- ✓ <u>Constraint:</u> link speed reduction must result in a loss in network capacity no larger than 15% at 65% assurance level → 136 streams

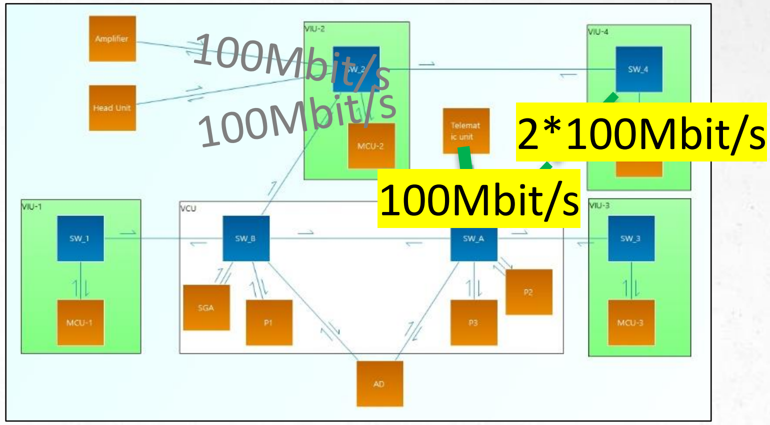

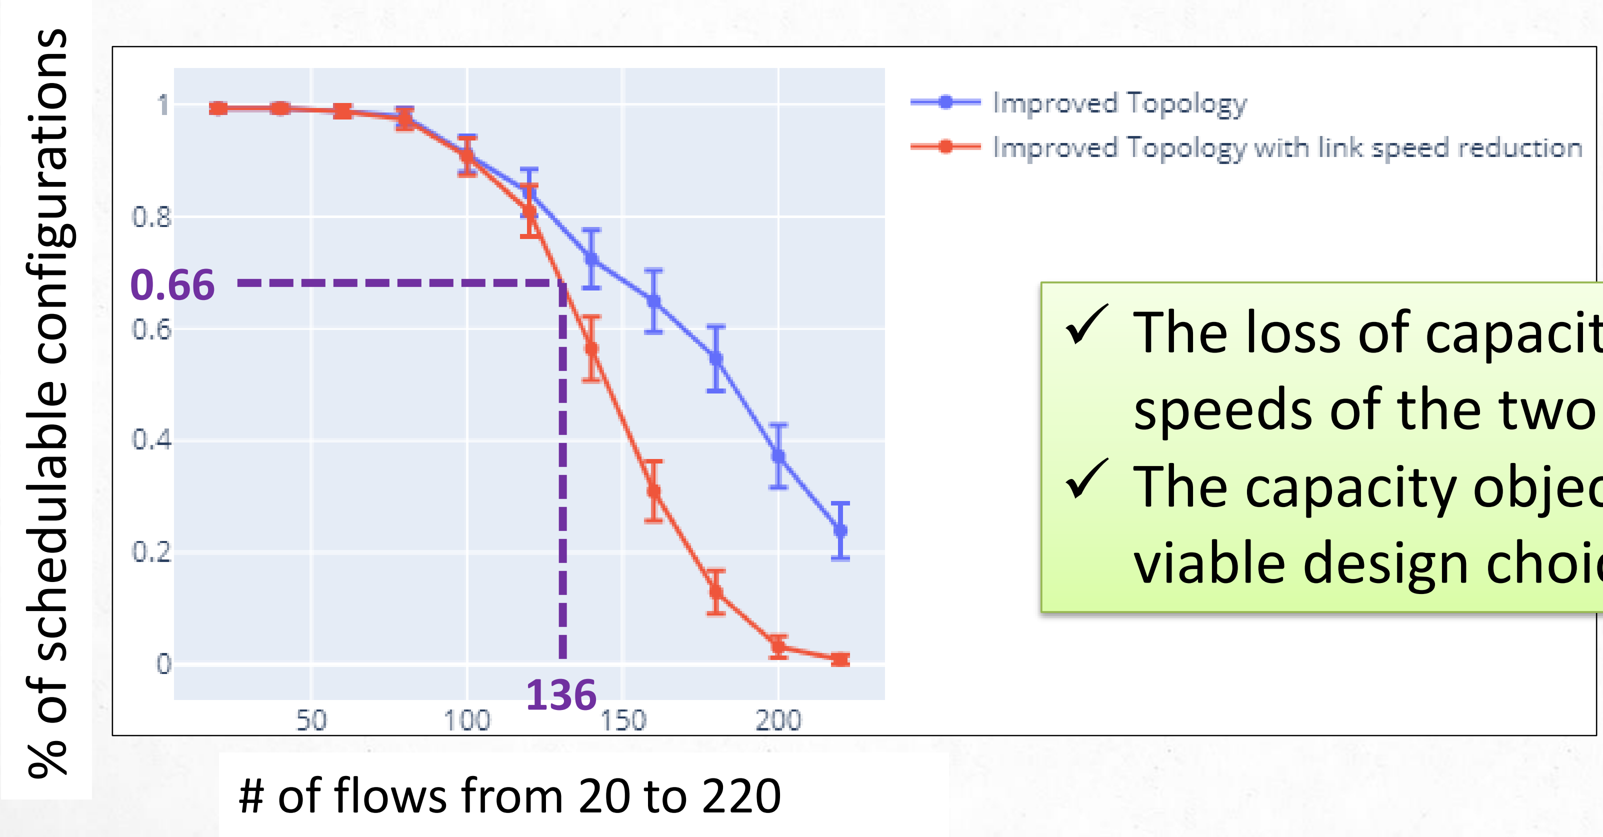

- ✓ The loss of capacity is substantial with the speeds of the two links reduced
- ✓ The capacity objective is however met → viable design choice

Cost Reduction: ECU removal

- ✓ Objective: reduce the # of ECUs by relocating their functions into other ECUs
- ✓ Constraint: timing requirements must be met and *total network capacity at the 65% assurance level must remain above 136 streams*
- ✓ Caveat: here only communication requirements are considered

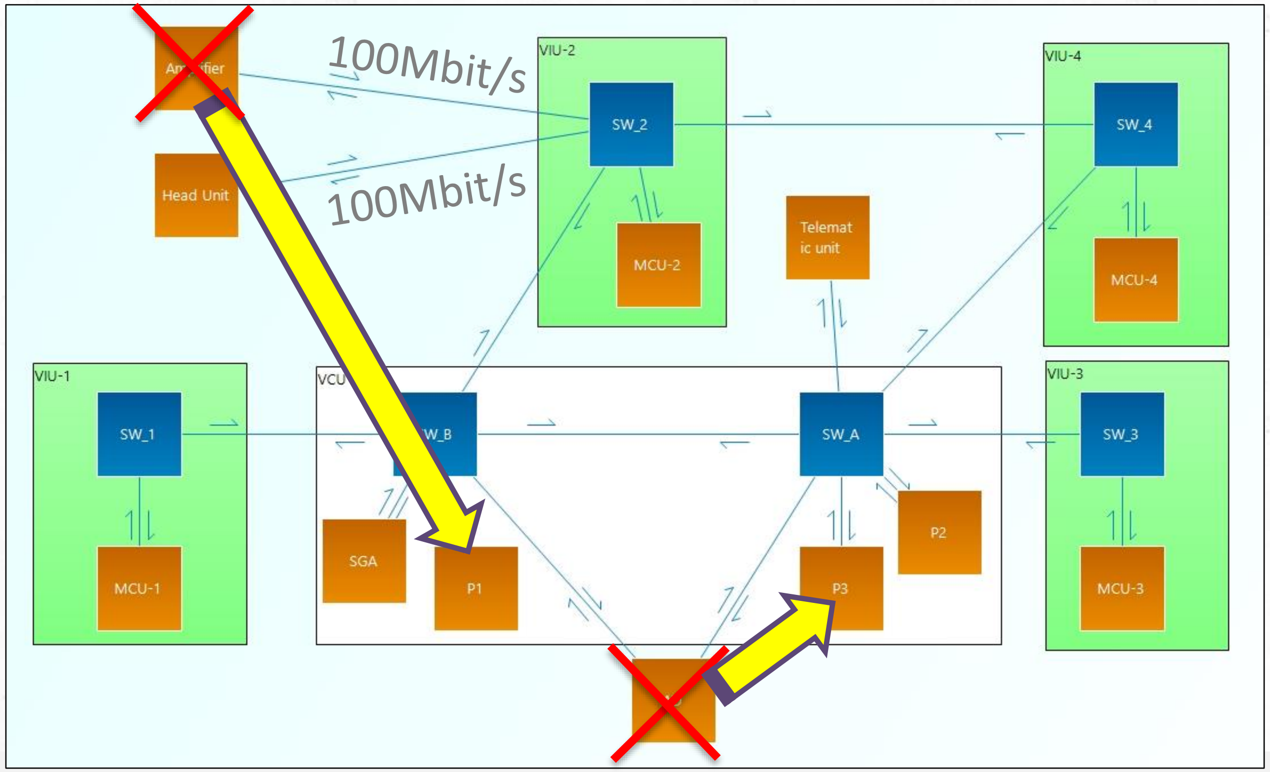

ECUs candidate for removal:
- ✓ *AD*: moved into *P3*
- ✓ *Amplifier*: moved into *P1*

ECU removal - Step1: AD into P3

✓ <u>Topology changes:</u> moving AD functions into P3, suppression of AD and the two links from AD, new link from P3 to SW_B

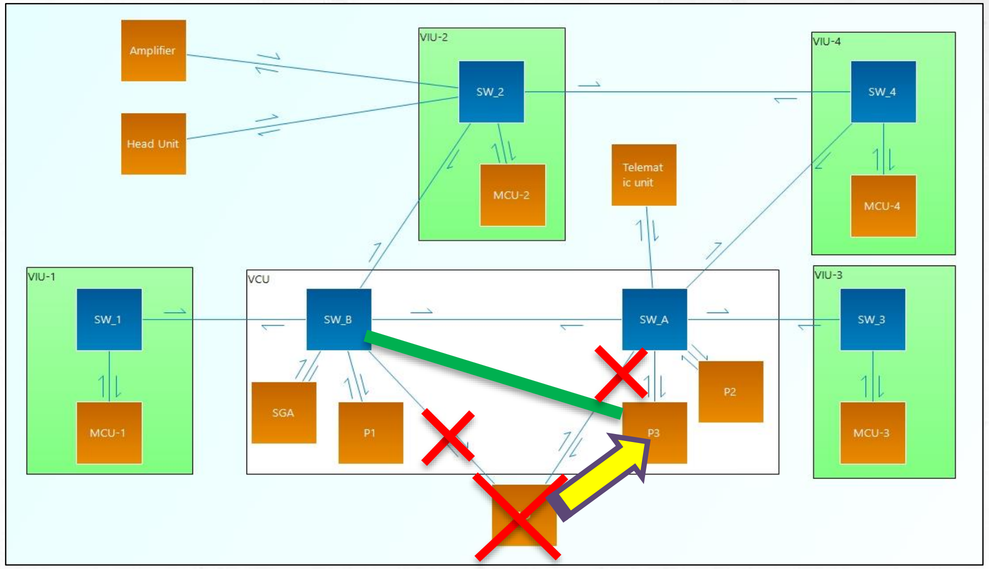

% of schedulable configurations

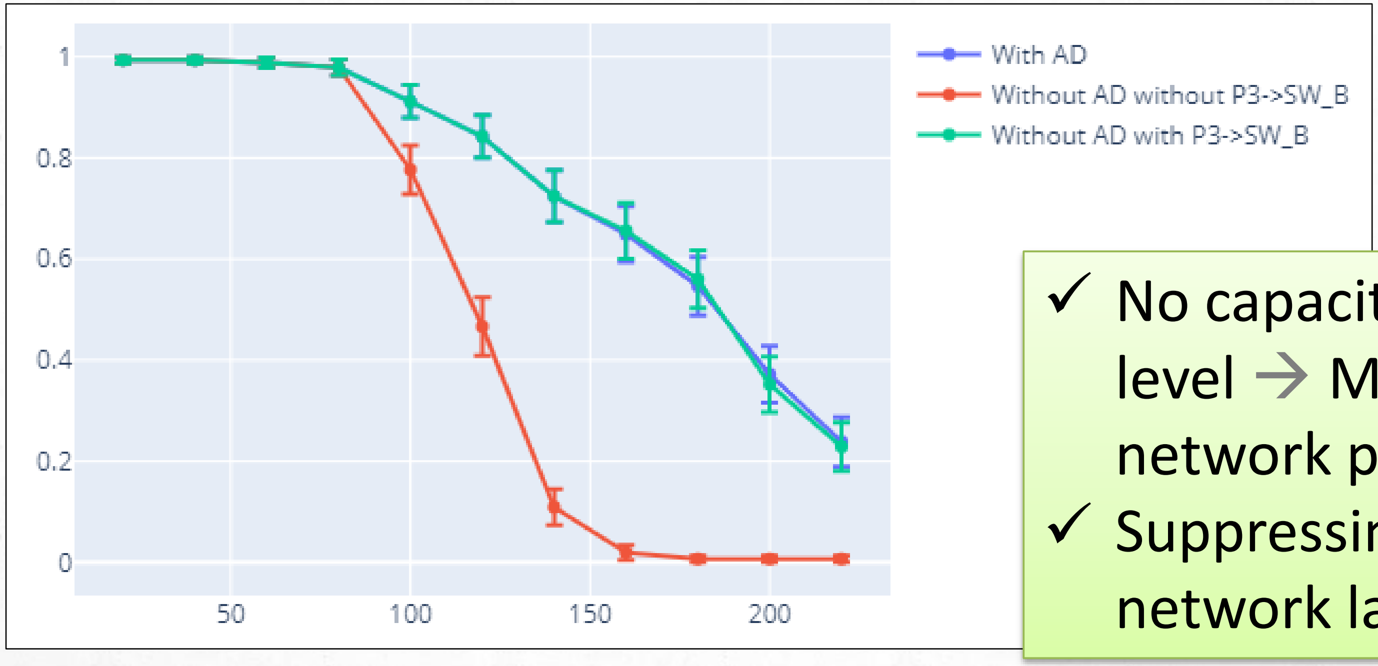

# of flows from 20 to 220

- ✓ No capacity penalty whatever the assurance level → Moving AD into P3 is feasible from network point of view
- ✓ Suppressing ECUs may require adjustments in network layout, here a new link

ECU removal – Step2: removing AMP in addition to AD

- ✓ <u>Topology changes:</u> topology at step 1, moving *Amplifier* functions into P1, suppression of *Amplifier* and link from *Amplifier*

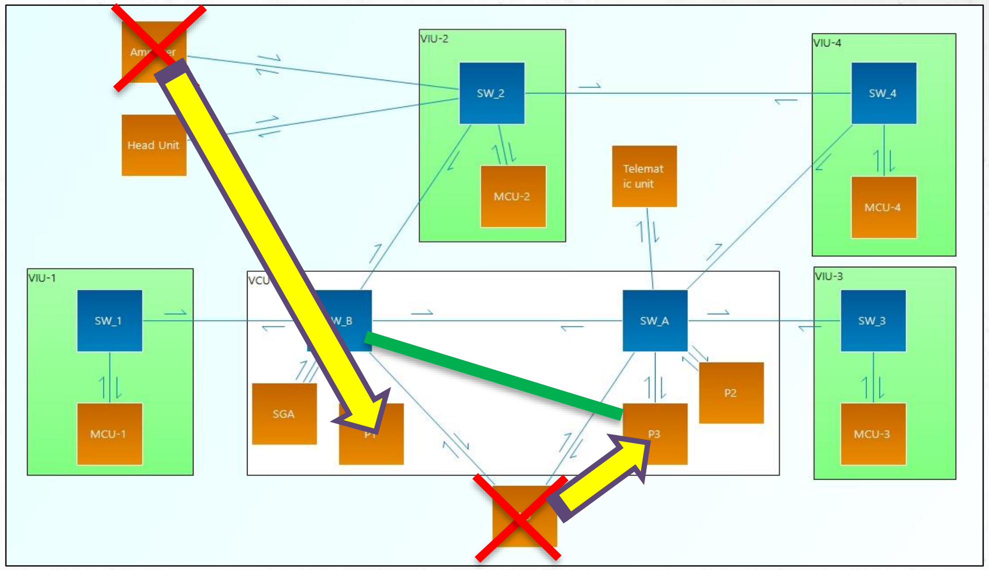

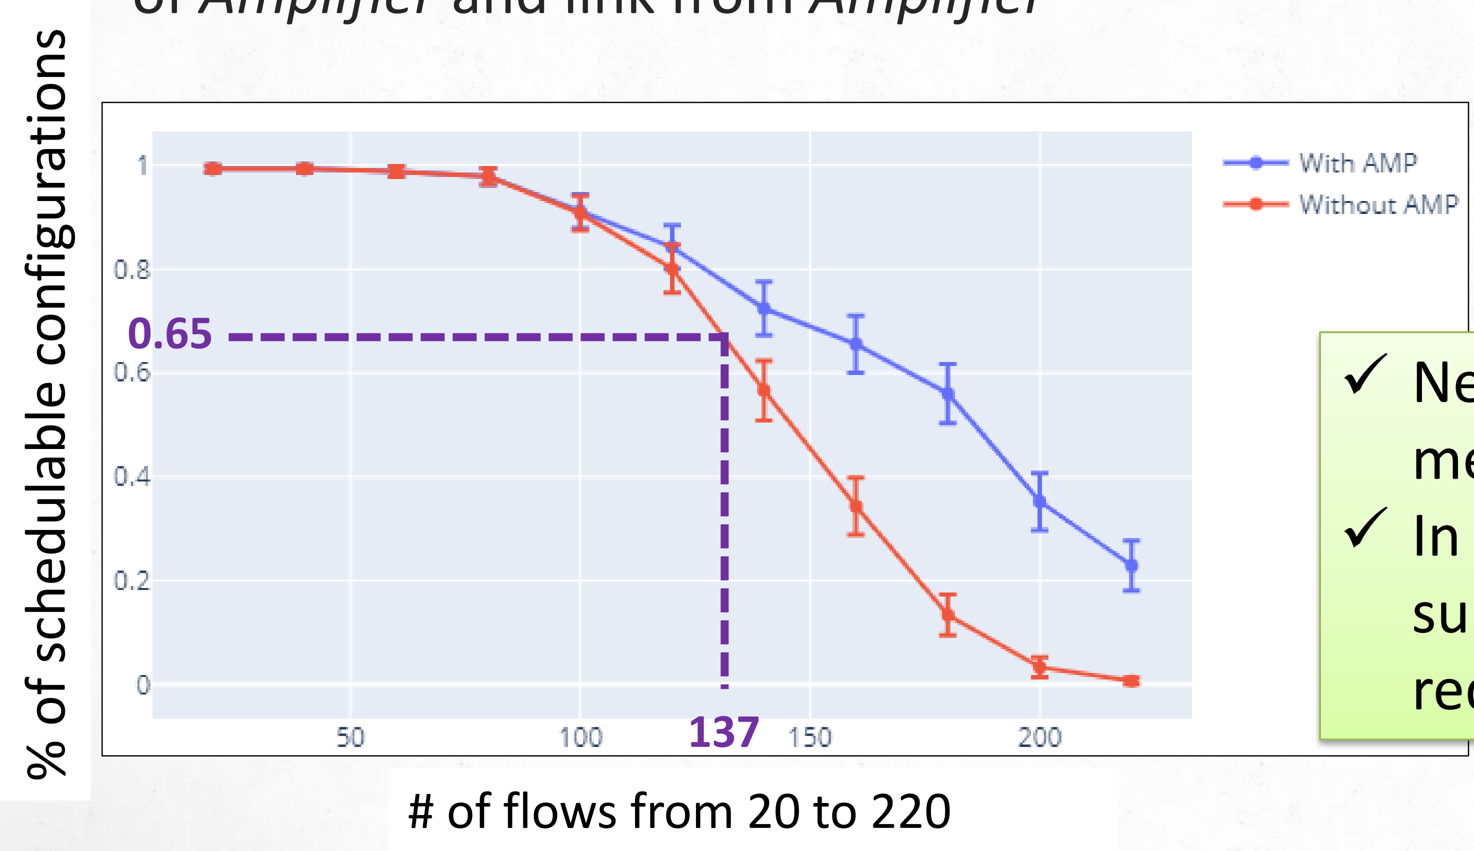

- ✓ Network capacity without *Amplifier* meets target objective of 136 streams
- ✓ In addition to *AD*, *Amplifier* ECU can be suppressed from the communication requirements point of view

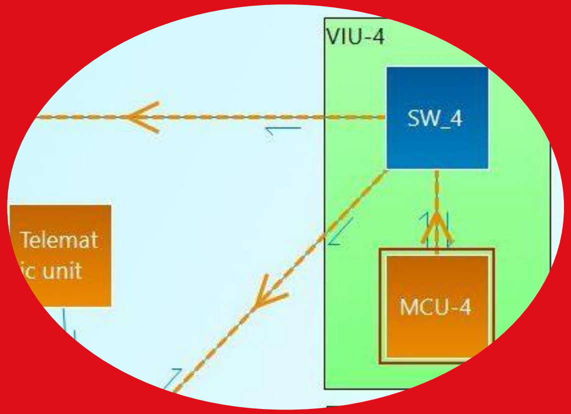

Design Question #5: Which gain in communication reliability with IEEE802.1CB ?

Cost-optimized topology: robustness to packet losses ?

- ✓ 136 streams including <u>8 critical streams from MCU-4 to SGA</u>,
- ✓ Bit Error Rate on all links: $1 \cdot 10^{-10}$ (*i.e.*, requirement for 1000BASE-T1 PHY in 802.3bp)
- ✓ How effective is 802.1CB here? Impact on timing constraints?

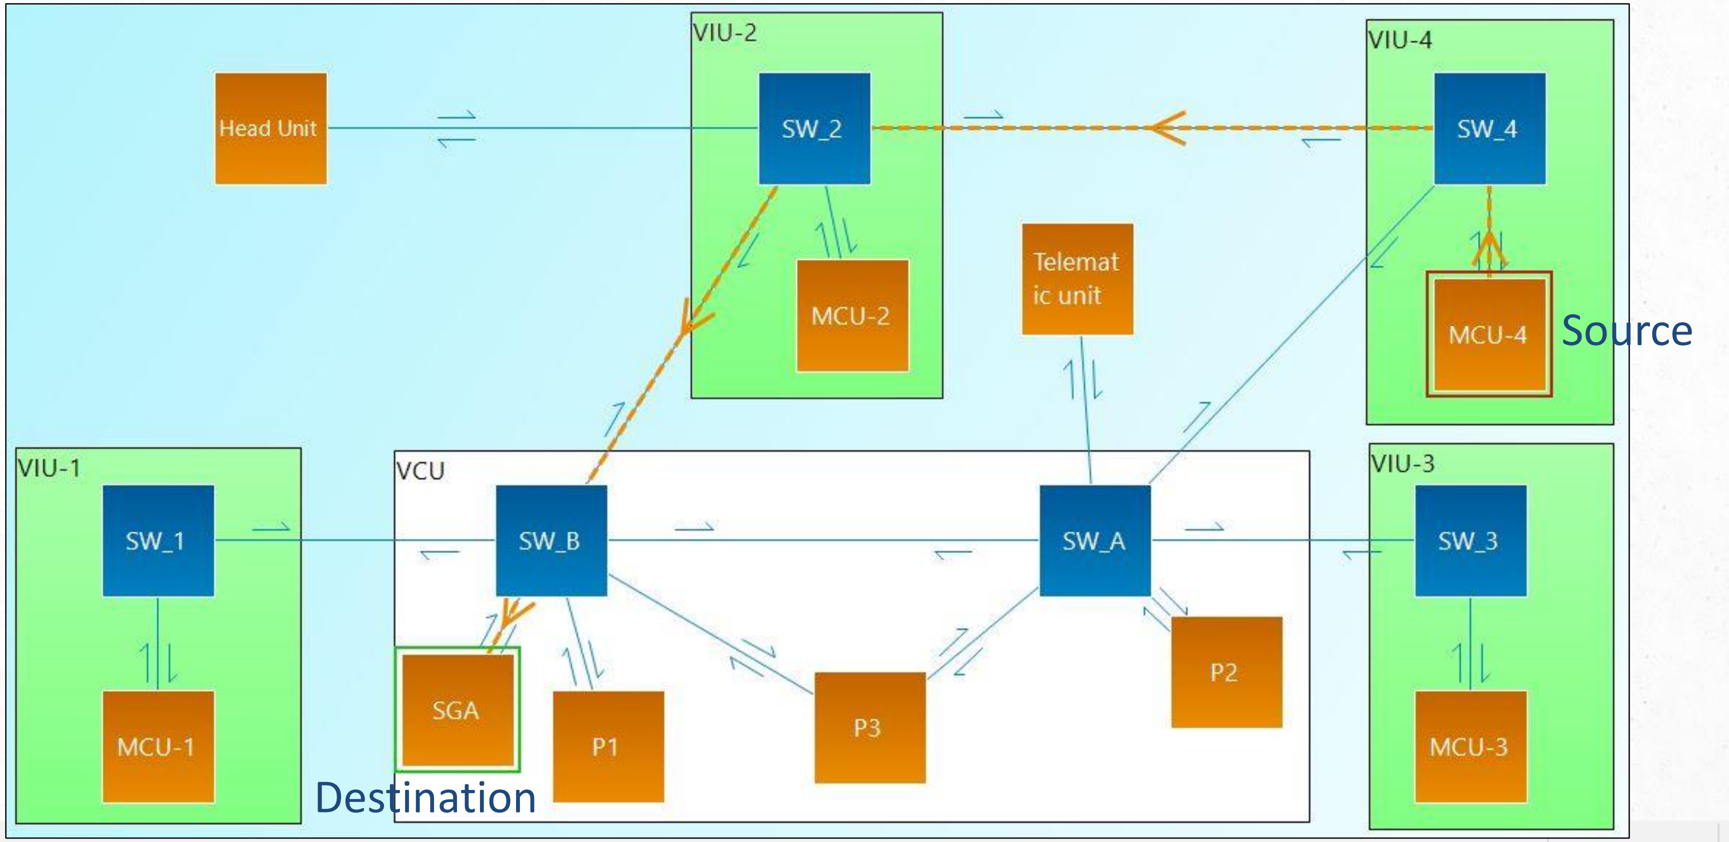

Implementing IEEE802.1CB

✓ Duplicated paths from MCU-4 to SGA

Critical Streams	Message loss rate Without CB	Message loss rate With CB
C_CT_21	1.82 * 10^-6	8.62 * 10^-7
C_CT_29	1.67 * 10^-6	8.77 * 10^-7
C_CT_41	1.69 * 10^-6	8.10 * 10^-7
C_CT_48	1.46 * 10^-6	7.47 * 10^-7
C_CT_51	1.41 * 10^-6	7.52 * 10^-7
C_CT_53	2.14 * 10^-6	1.05 * 10^-6
C_CT_56	1.40 * 10^-6	6.77 * 10^-7
C_CT_62	1.50 * 10^-6	7.00 * 10^-7

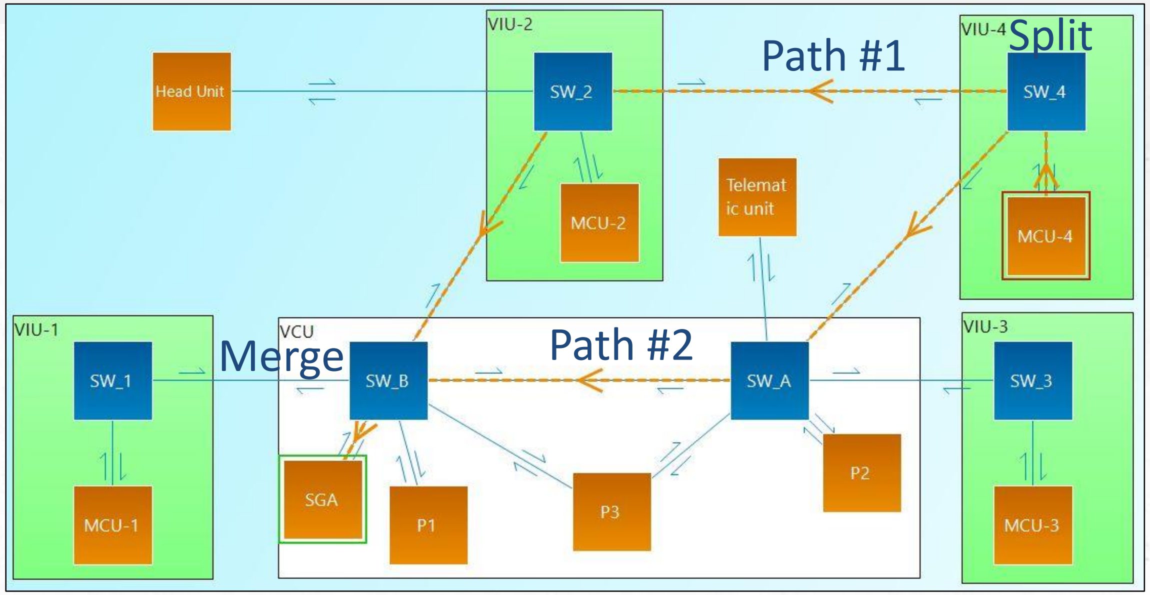

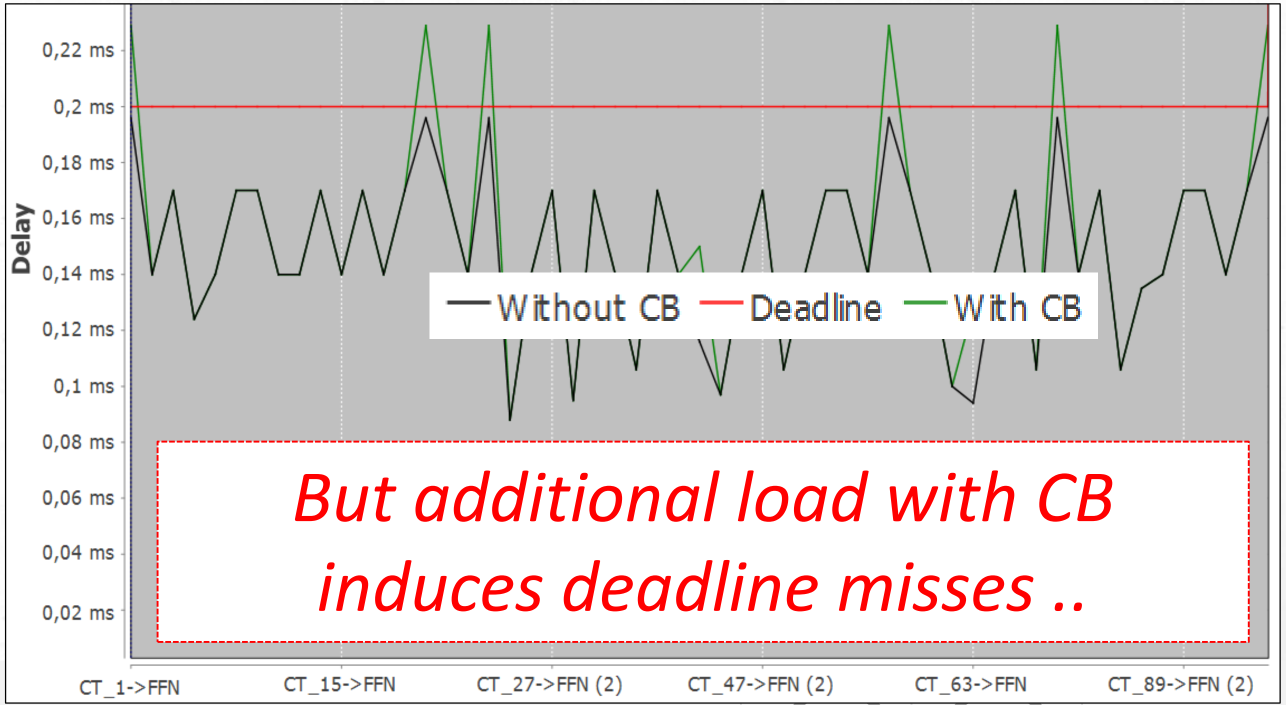

- ✓ Without CB, avg. loss rate ≈ 1 every 600 000
- ✓ With CB, avg. loss rate ≈ 1 every 1 200 000
- ✓ Gains with CB capped here by the 2 non-replicated links: before *split* and after *merge* points

Solution: reworking traffic classes & priority allocation

- ✓ With 8 streams replicated with 802.1CB, no feasible solution with 4 traffic classes
- ✓ <u>Constraint:</u> critical streams must remain all at highest priority level
- ✓ *Concise Priorities* algorithm returns a feasible solution with 5 priority levels

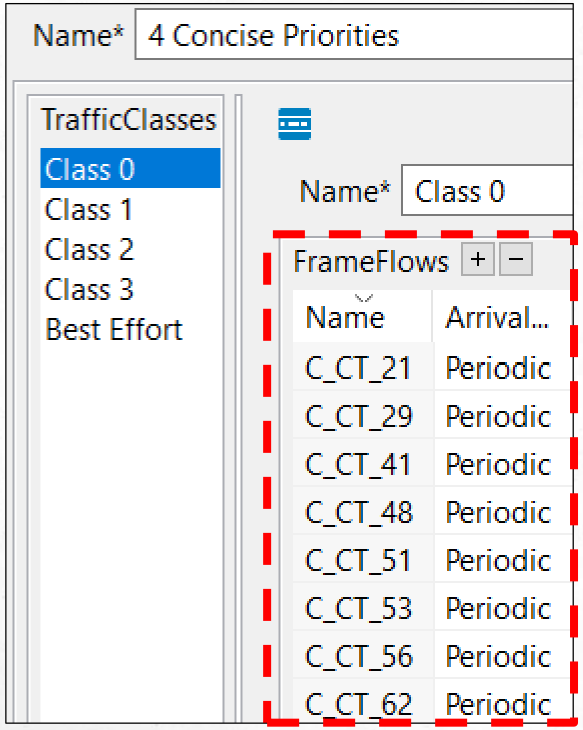

ClassConfigurations

TrafficClass	Priority
Class 0	7
Class 1	6
Class 2	5
Class 3	4
Best Effort	3

802.1CB increases reliability but may require TSN scheduling solutions to be revisited!

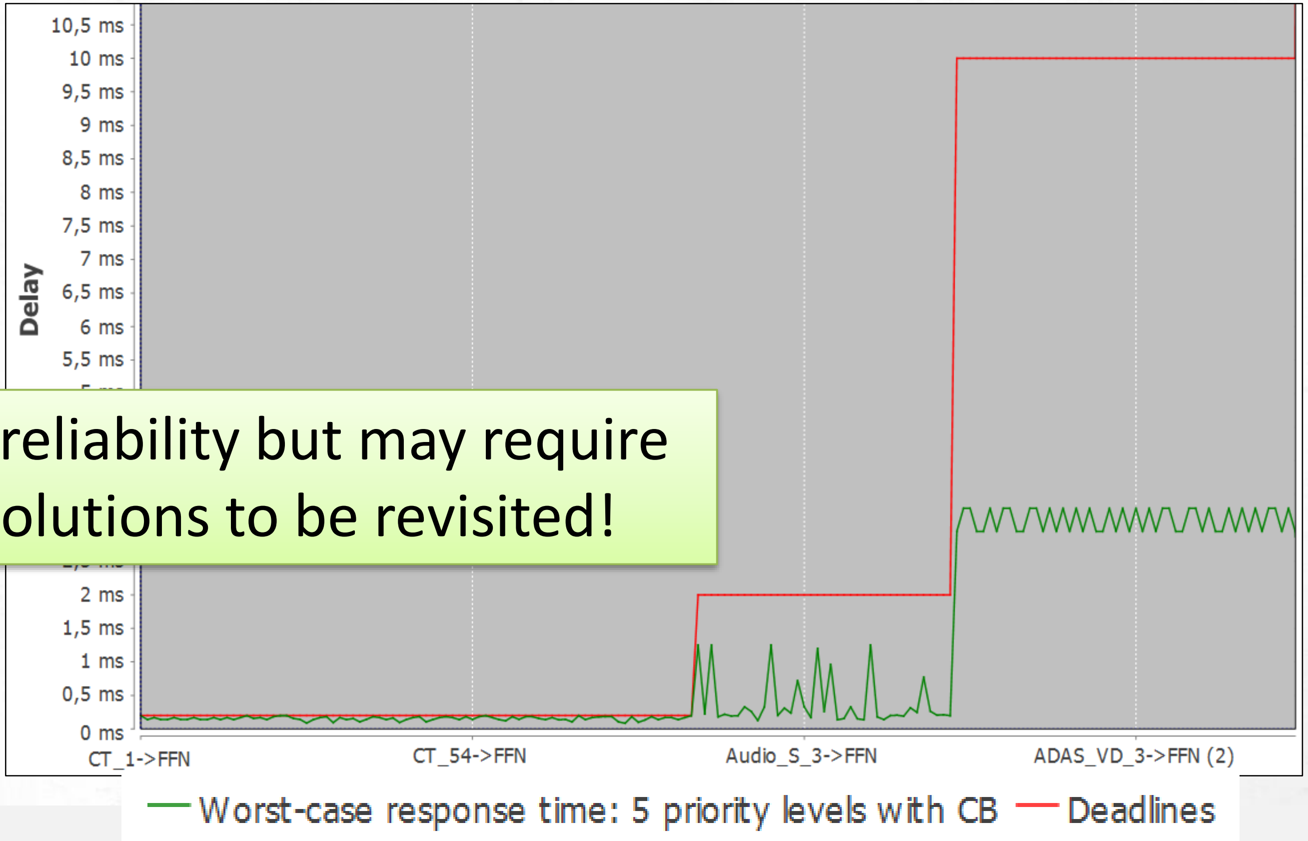

Conclusion and a look forward

Designing future-proof TSN-based architectures

Application to Volvo's centralized architecture

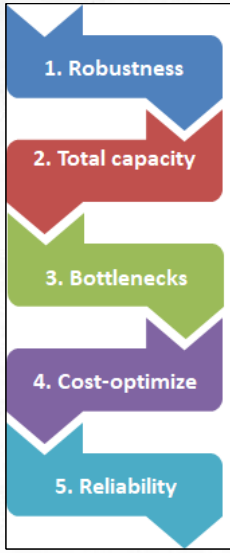

1. Impact of additional low-priority traffic ?
2. Quantify network capacity wrt TSN technological options
3. Identify bottlenecks in E/E architecture and remove them
4. Cost-optimize by reducing link speeds and # of ECUs
5. Assess and optimize communication reliability

Progress factor by taking advantage of the capabilities of today's design space exploration techniques for the problems that they can solve better and faster – implemented as Topology-Stress-Test and Topology-Optimizer modules in RTaW-Pegase

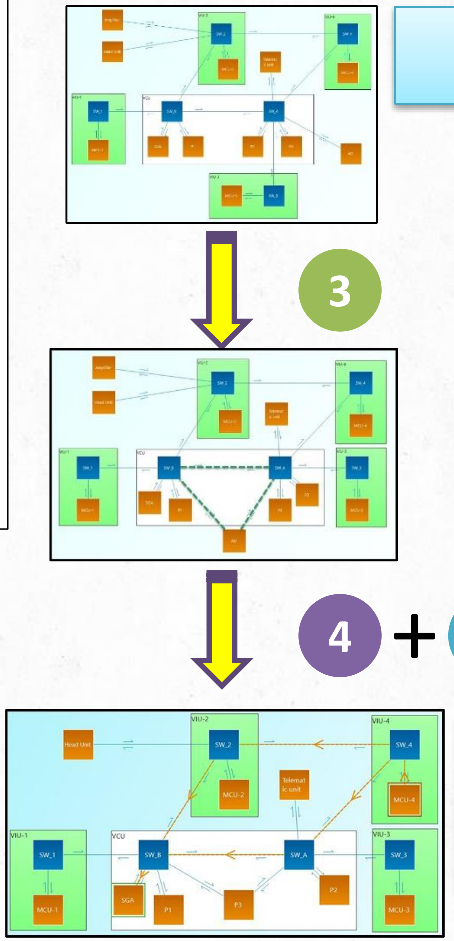

Baseline architecture

1 + 2

3

Enhanced-capacity architecture

4 + 5

Cost- and reliability-optimized architecture

Thank you for your attention!

uni.lu
UNIVERSITÉ DU LUXEMBOURG

VOLVO

RTaW
RealTime-at-Work