Ethernet & IP @ Automotive Technology Week
3 Nov 2021 - 4 Nov 2021

IEEE SA
STANDARDS ASSOCIATION

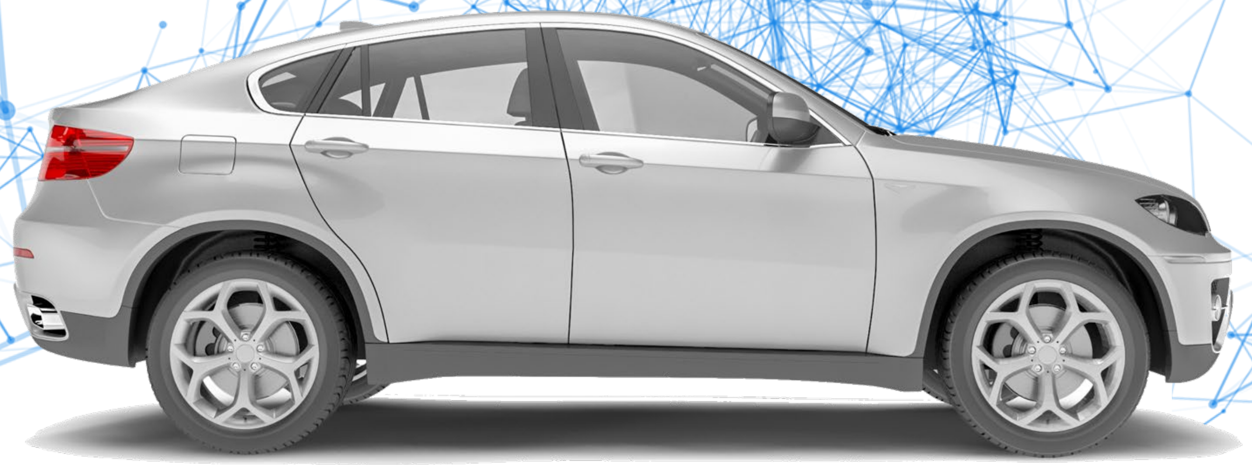

The "Free Rider Principle" for Low-Bandwidth Flows in High Line-Rate Networks

Max Turner(1) and Jean Walrand(2)

Berkeley UNIVERSITY OF CALIFORNIA (2)

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

Agenda

- Zonal Architecture Introduction
- Bandwidth Discussion
- Burstiness and Latency
- Designing the Network
- Proposing a Solution
- Conclusions

What drives Zonal Architecture?

- Weight of wiring-harness
  (communication and power)
- Complexity of wiring-harness (power)
- Electronic fuses for power distribution
  (simplify automotive "network management (NM)[1]")
- Hardware abstraction
- Number of ECU-boxes
  (housings, power conversion, ...)
- Compute, RAM and NVM scalability (pay for overhead only once)
- Simplified (over the air) update: fewer update targets, less double NVM overhead

[1] Do not confuse with the IT terminology around SNMP

Zonal Architecture reduces Wiring

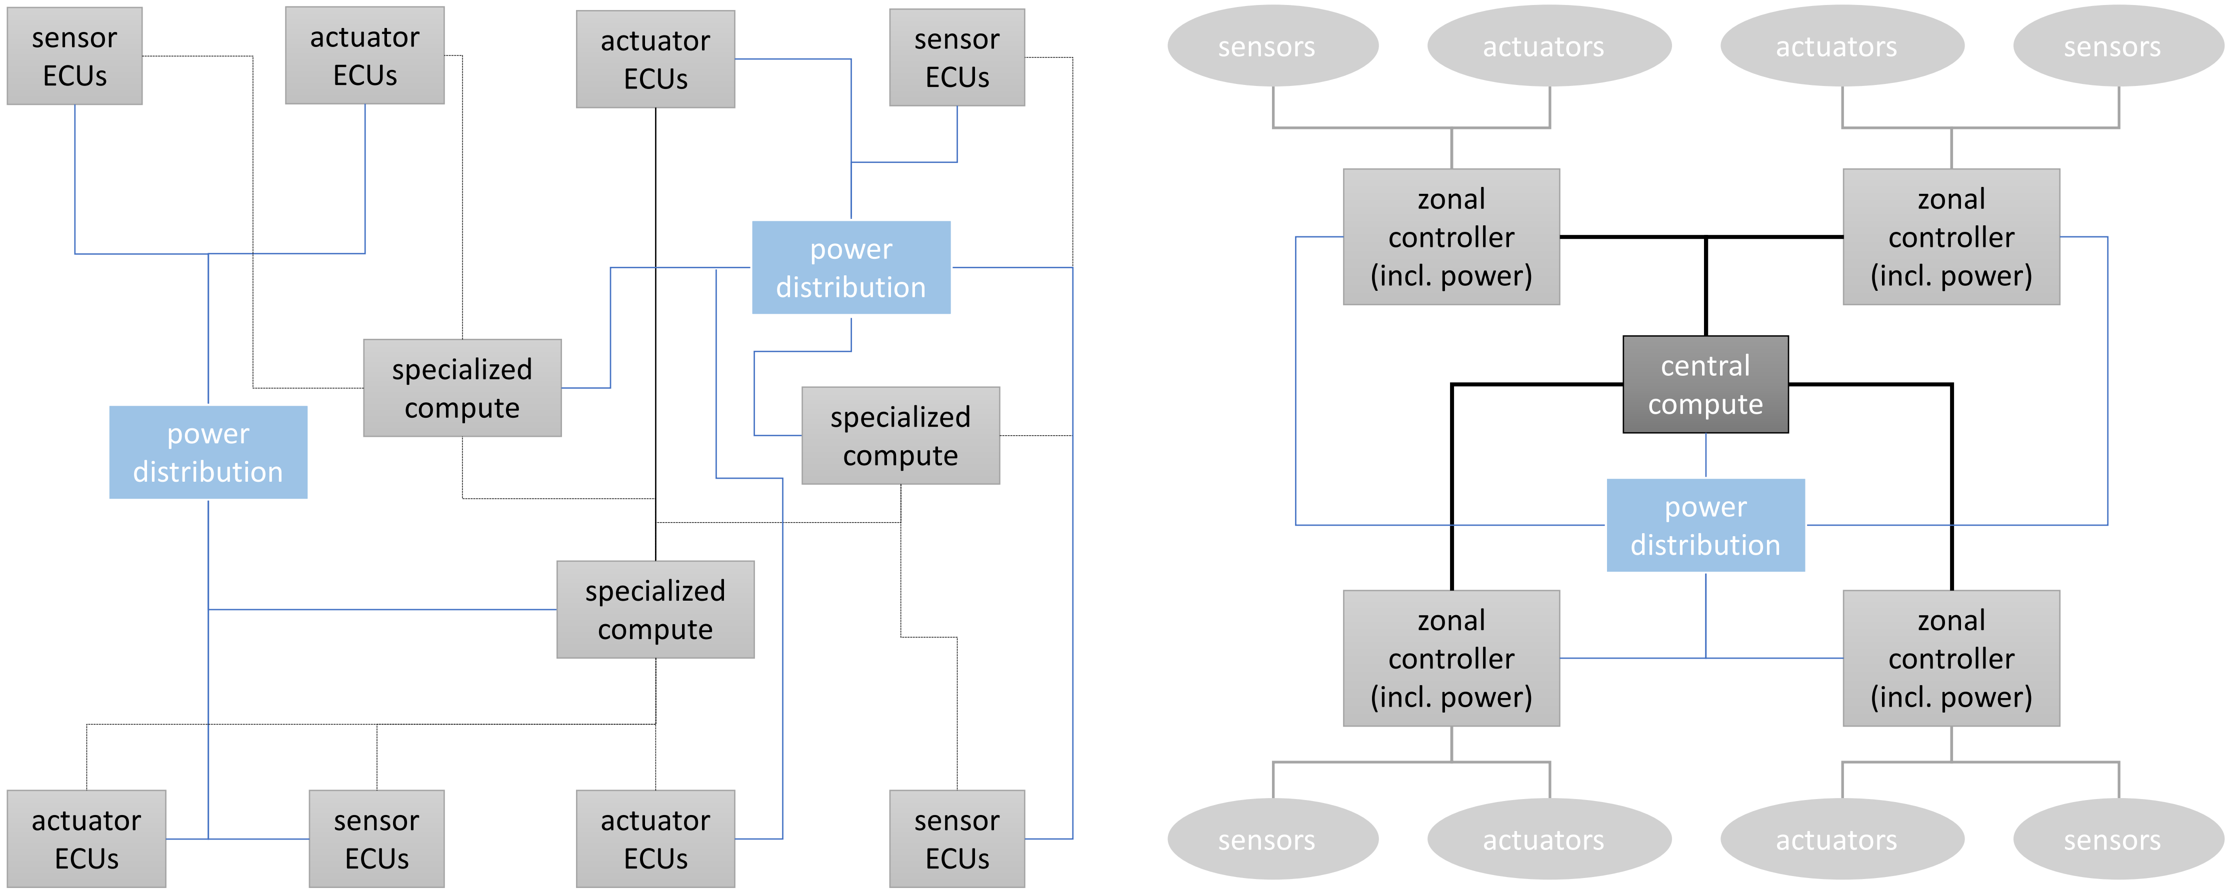

2021 Ethernet & IP @ Automotive Technology Week, M. Turner (Ethernovia) & J. Walrand (UC Berkeley)

Control-Loop Concerns

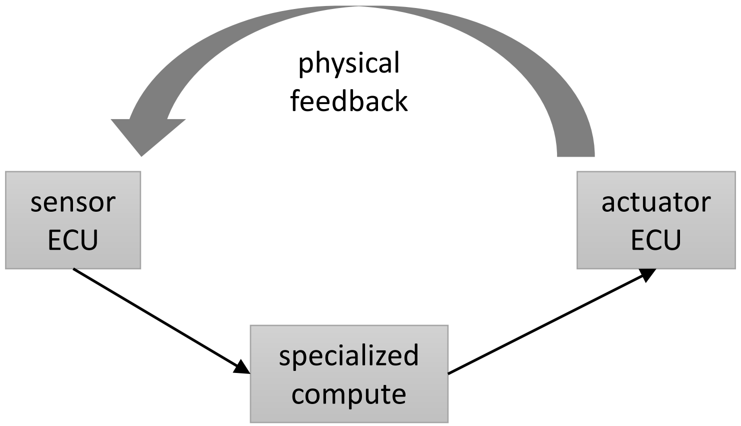

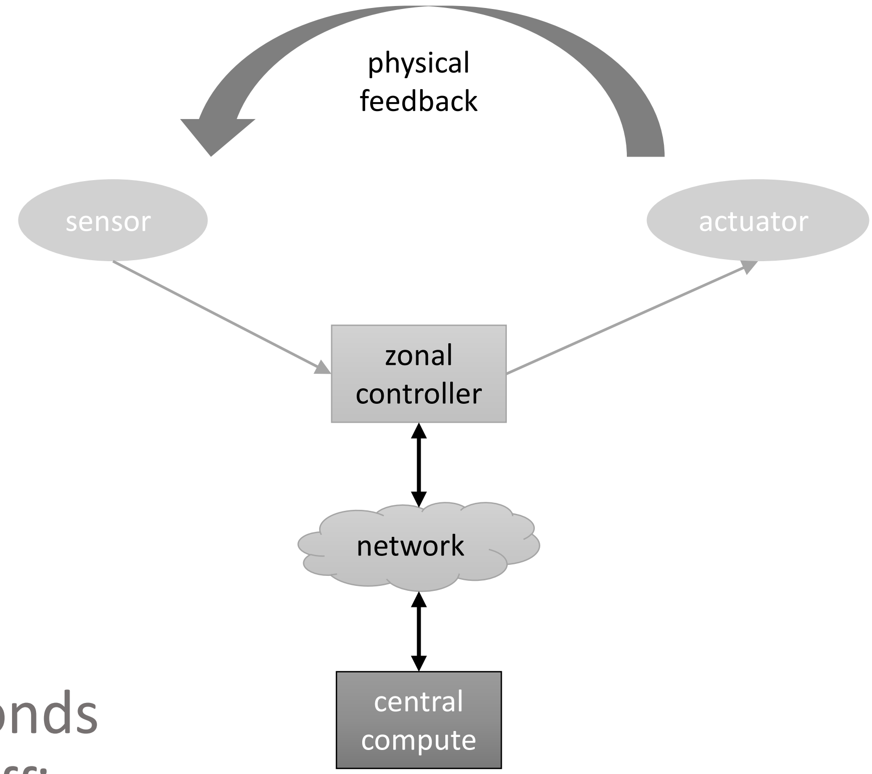

Control-loop latency in the order of milliseconds would drive higher line-rates, but control traffic often only has low bandwidth requirement - this is not economically viable

What does (not) drive Bandwidth?

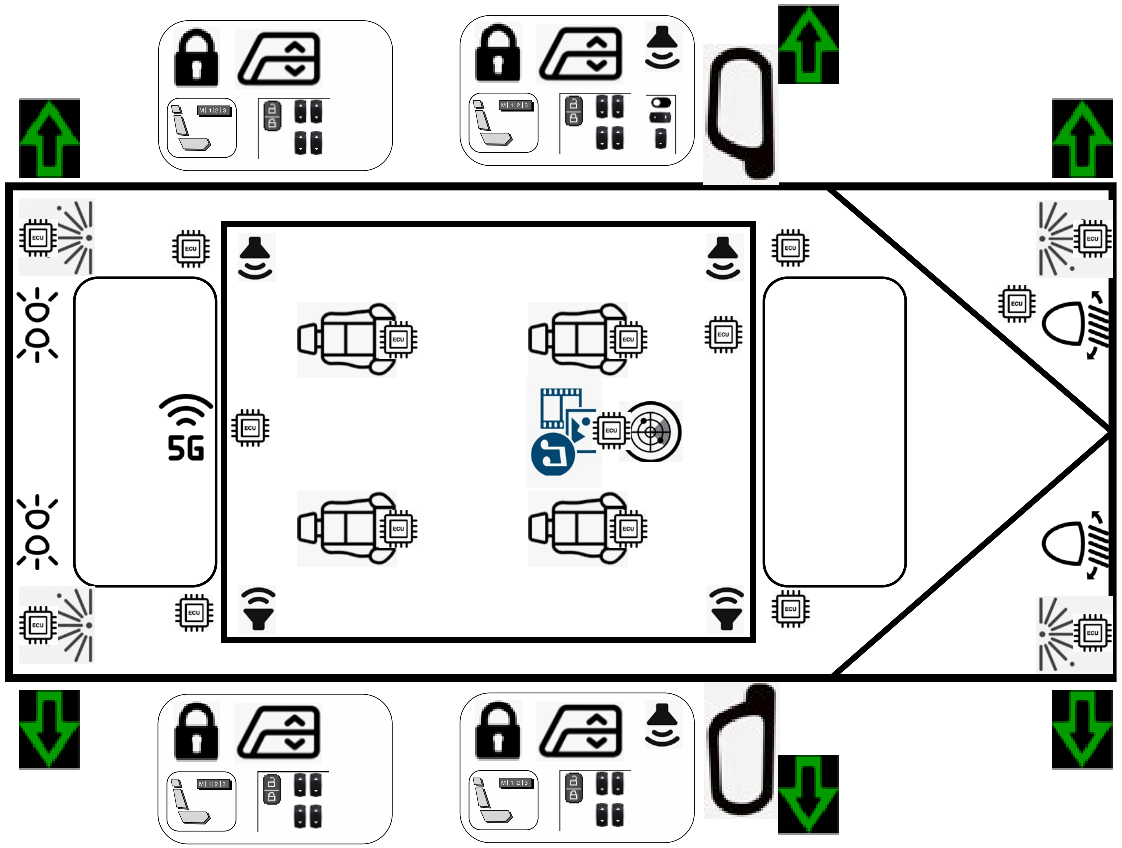

- Audio ↔ 200Mbit/s
- 5G cell ↔ 1Gbit/s
- Control ↔ 100Mbit/s
- Smart Sensors ↔ 2Gbit/s
- …

Σ < 5Gbit/s

Ethernet Video Bandwidth

Name	Resolution	Bandwidth	Link Usage
VGA (0.3Mpx)	640x480 24bit @60Hz	0.44Gbit/s	47% of 1Gbit/s
UXGA (1.9Mpx)	1600x1200 24bit @30Hz	1.4Gbit/s	59% of 2.5Gbit/s
QXGA (3.1Mpx)	2048x1536 24bit @30Hz	2.3Gbit/s	49% of 5Gbit/s
4k UHD (8.3Mpx)	3840x2160 16bit @30Hz	4.0Gbit/s	85% of 5Gbit/s

Cameras drive higher line-rates

- 100Byte overhead (MAC, IP or 1722, ...)
- max. frame size 1500Byte
- no blanking transmitted

A Zonal Architecture Example

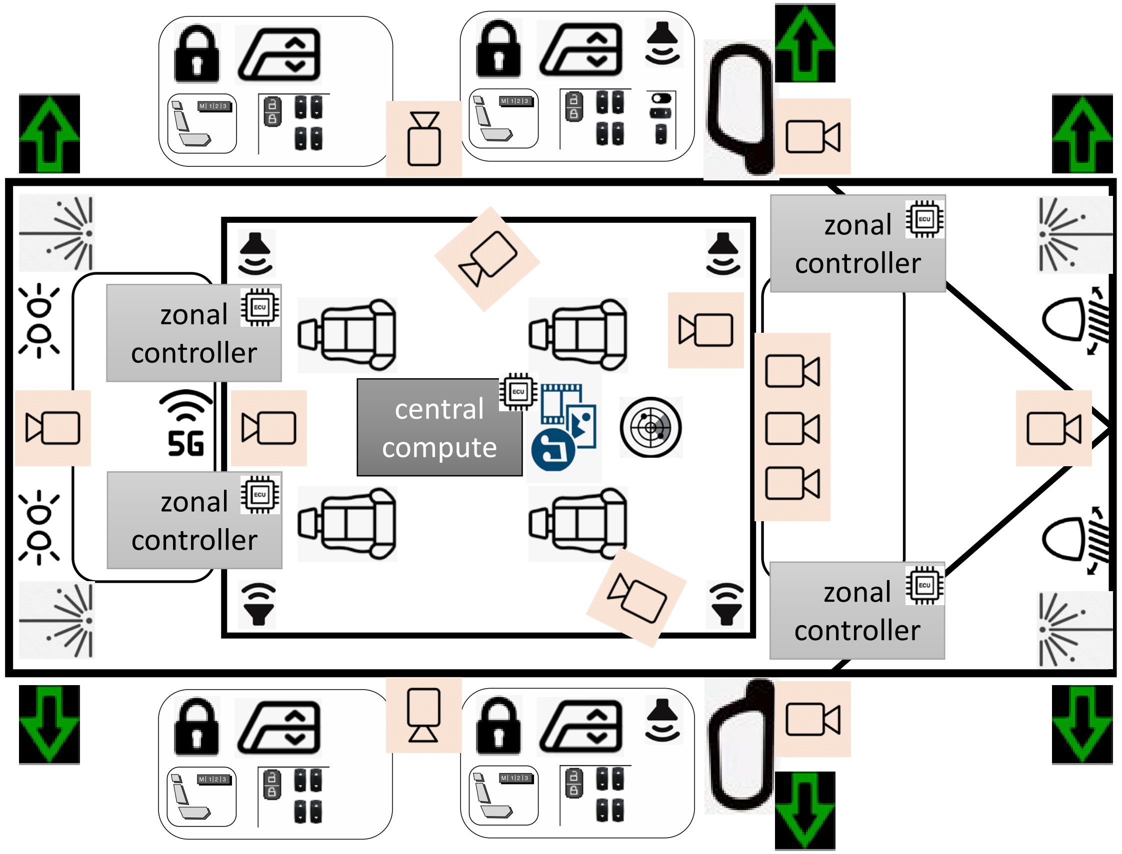

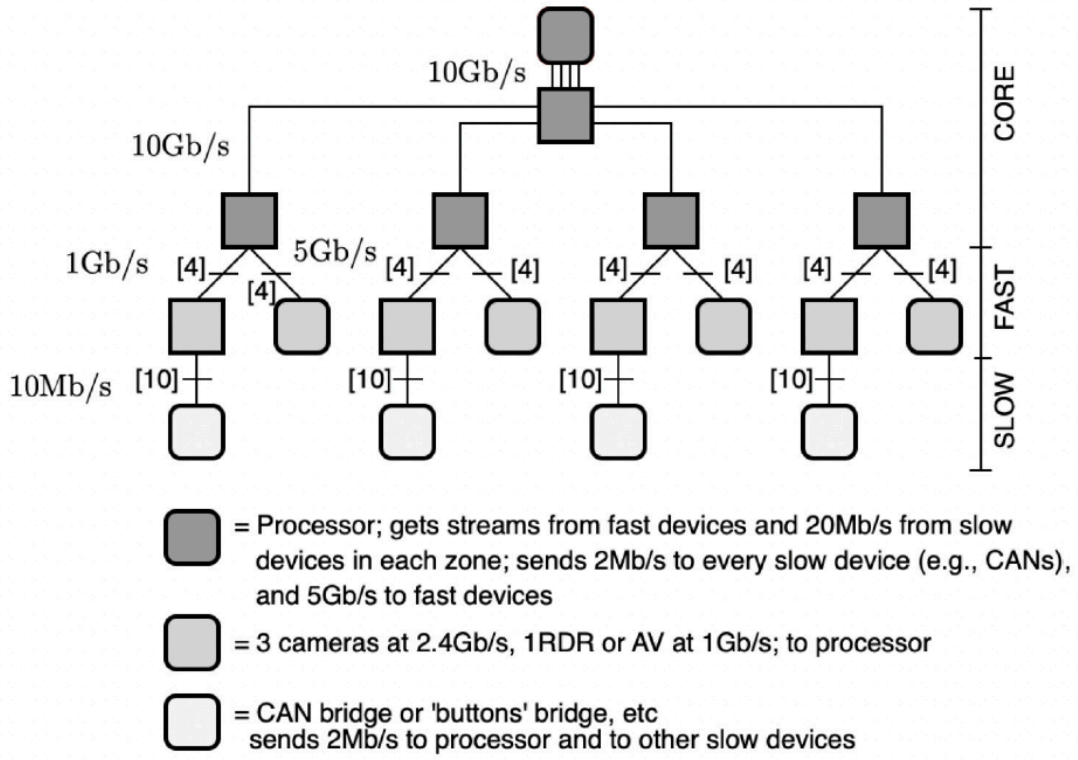

2021 Ethernet & IP @ Automotive Technology Week, M. Turner (Ethernovia) & J. Walrand (UC Berkeley)

Traffic Categories

- Sensors usually generate periodic data due to sampling (temperature, RPM, ...) or scan-rate (camera, radar, lidar, ...)
- Safety relevant (control) communication is usually periodic for loss (of application/input) detection (incl. counters, CRC)
- Event driven traffic is low in bandwidth, meaning small frames transmitted only infrequently
- TCP traffic, e.g. from/to the internet (5G), may use large frames, but has low bandwidth compared to the video/sensor traffic

➢ For shapers to deliver predictable (per hop) latency, the ingress must be well defined(1)

➢ Re-transmissions due to losses introduce unpredictable bandwidth needs

➢ In order to avoid losses for all streams:

- ✓ Shape and police(2) **all** flows
- ✓ Shape and police at the talkers(3)

(1) See next slides on Burstiness
(2) Policing: Check ingress rates and drop frames upon overload
(3) E.g. in a middleware, the communication stack or NIC HW

Burst Accumulation on a "Daisy-Chain"

Flow α accumulates bursts due to bursts from interfering streams[1] β and γ

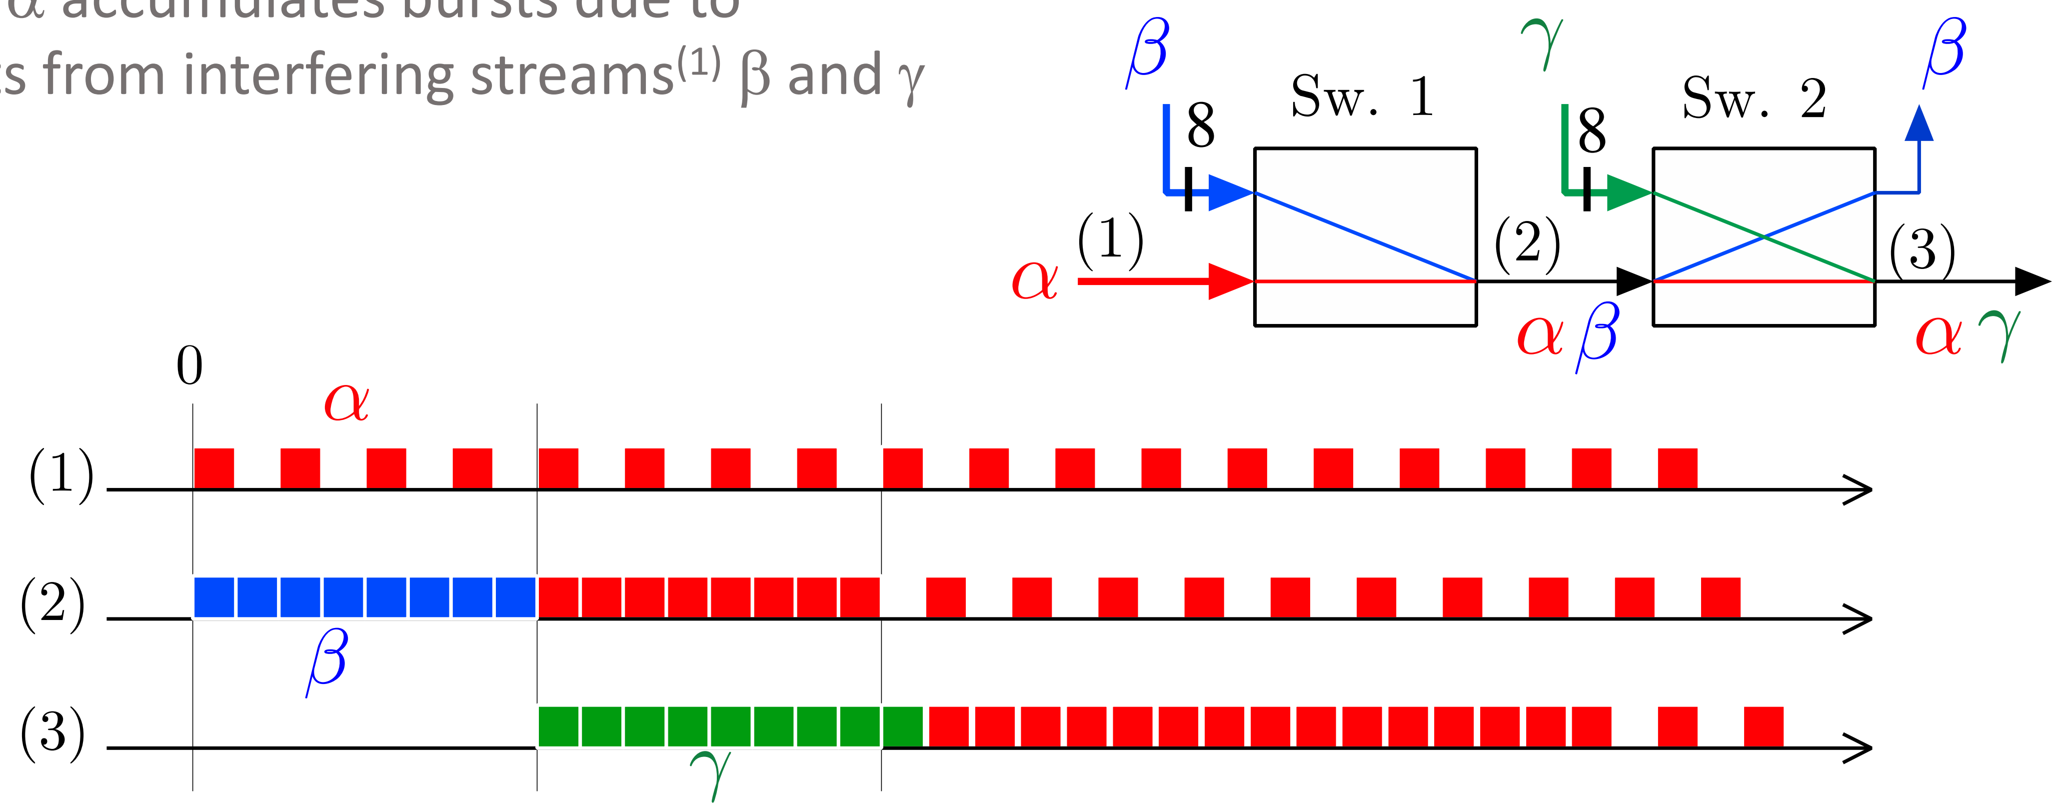

[1] Streams β and γ consist of 8 flows each

What is Burstiness?

- **Intuitively:** uneven spacing of frames

same average rate *a*
for both flows

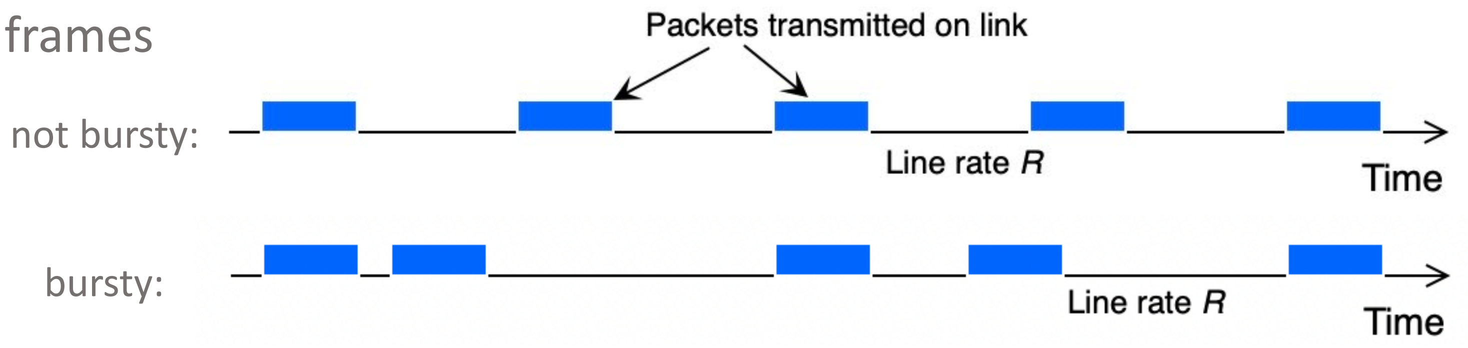

- **Mathematically**[1]: burstiness *A* is the max. backlog in a buffer that serves the flow at its average rate a: If at most *A + at* bits arrive in any t seconds(2), we say that the flow type is (A, a)

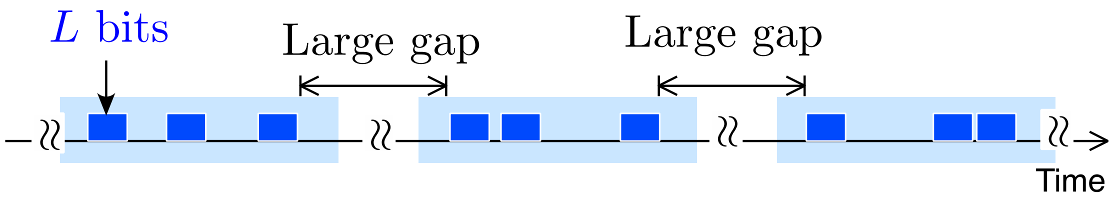

burstiness A = *3L*, max. back-to-back bits = *L*

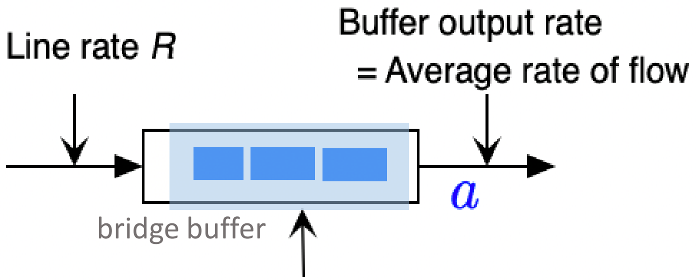

burstiness A = maximum backlog

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

(2) This is a linear upper bound on the 'arrival curve' of the flow

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

Classifying Flows

Accumulation of α *(A, a)* during transmission of β *(B, b):*
= (rate of α) * (transmission time of burst of β)
= (rate of α) * (size of burst of β)/(line-rate R)
= B * a/R

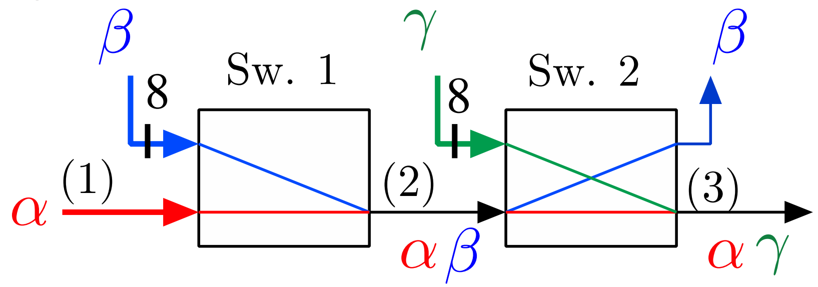

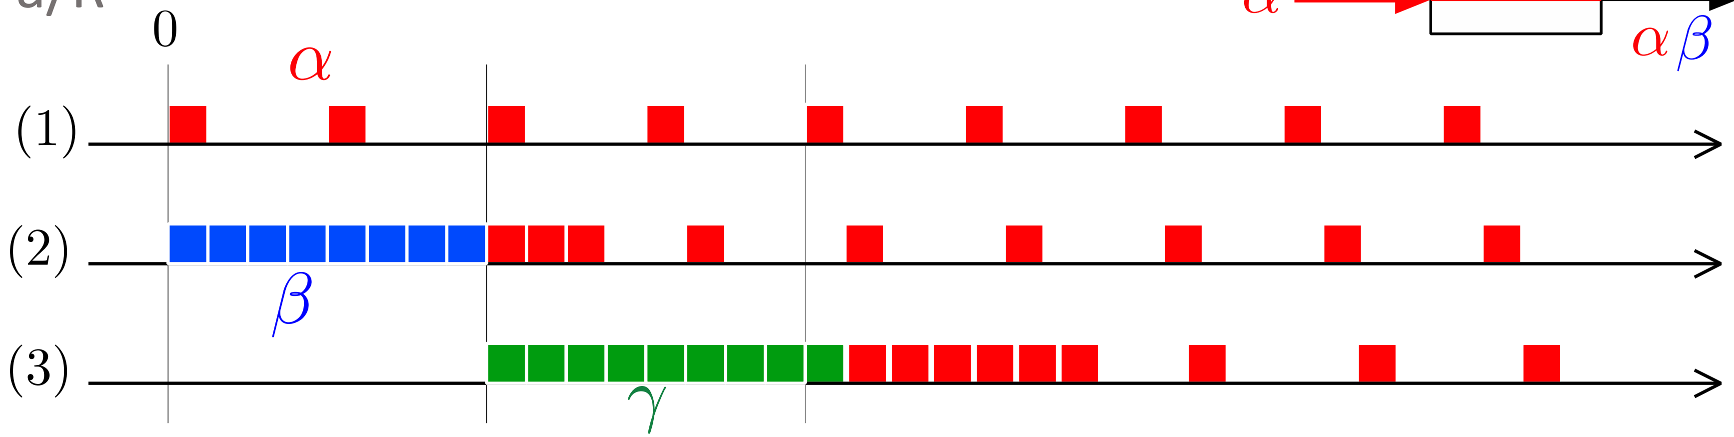

We call a flow ***(A, a)* a ‘fast flow’, if a/R ≈ 1** <
We call a flow ***(A, a)* a ‘small flow’, if a/R <<1**

Burst Accumulation for Fast Flows

A small number of interfering streams (β & γ) cause little to no accumulation in a fast flow (α) when sharing few hops

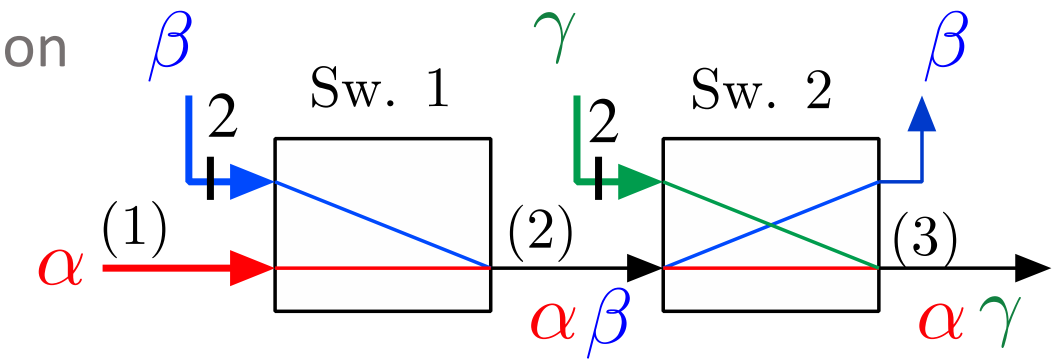

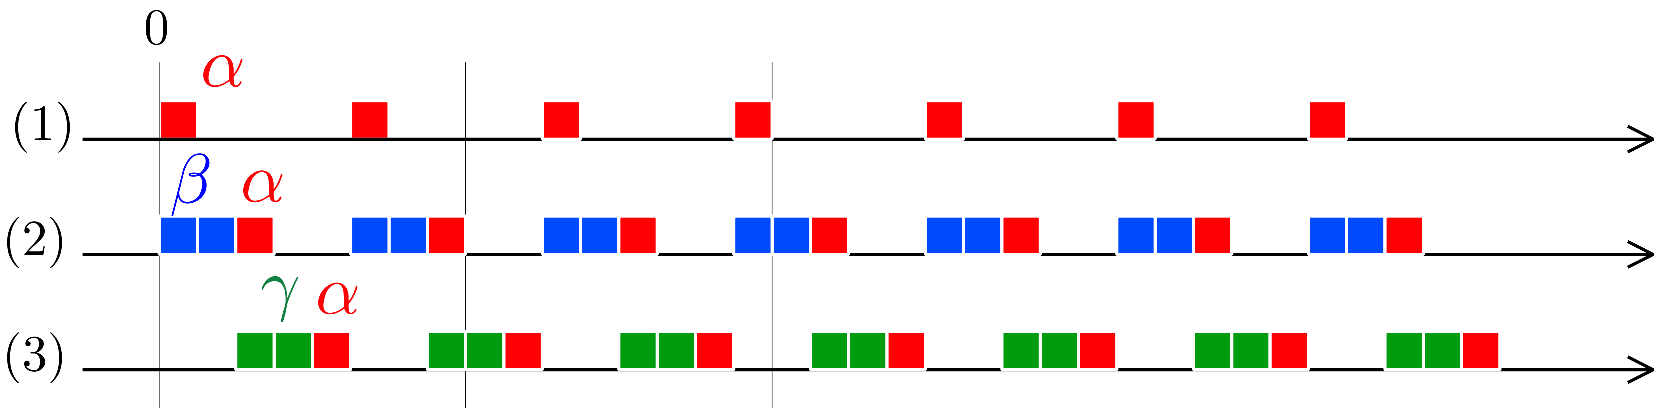

Burst Accumulation for Small Flows

There is only little burstiness accumulation for small flow α, if flow rate a << R line-rate

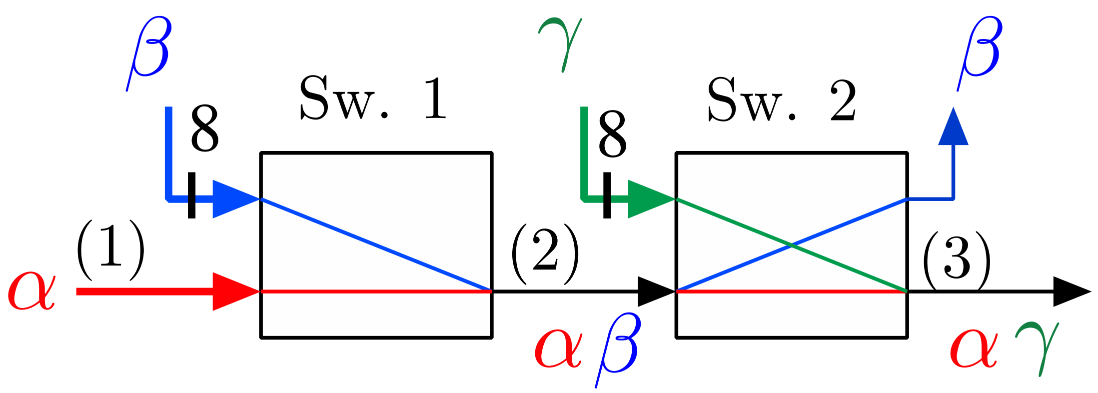

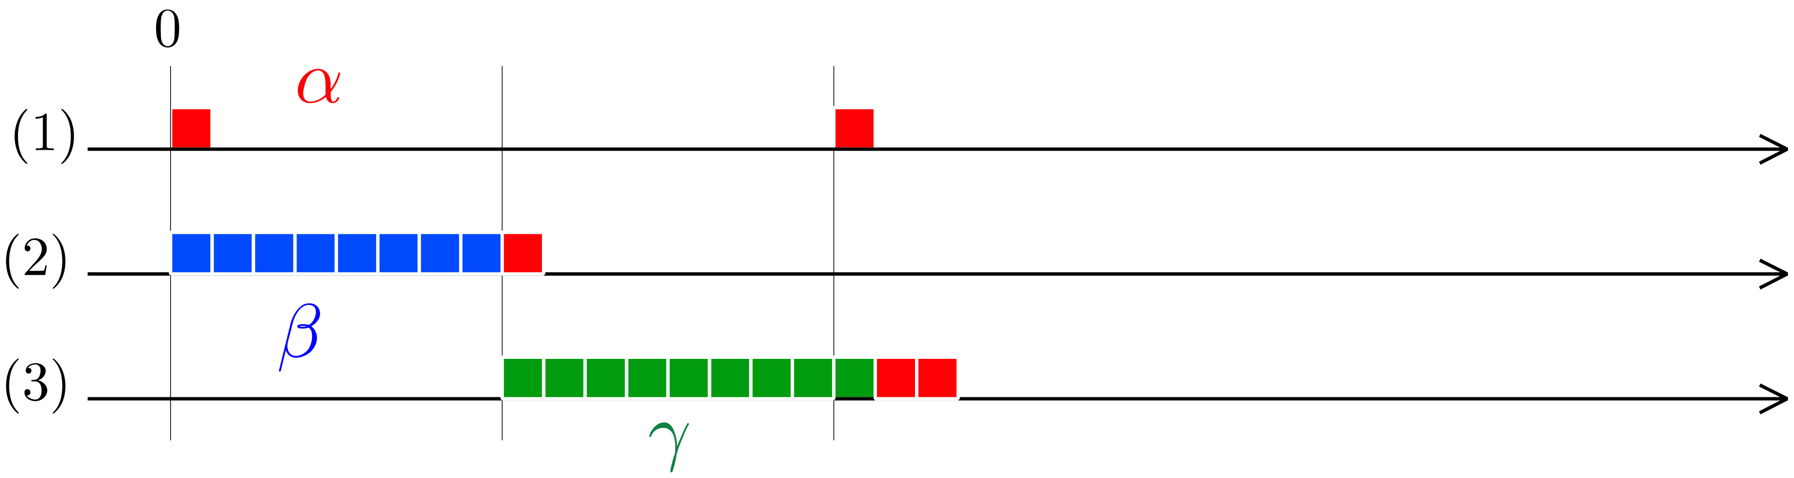

The Free Rider Principle for Small Flows

Consider two flows α=A+a$*$t (max. frame size[1] L_A) and β=B+b$*$t that share an output port of line-rate R:

We say flow α is small, if a/R << 1, then

$$A' \leq A + (B+L_A)*a/R$$

$$\text{i.e. } A' \approx A$$

meaning the burstiness of α remains constant

We call A'=A the small flow approximation, which enables the

Free Rider Principle:

Small flows do not suffer additional latency if combined with fast flows on high line-rate links

[1] including L_A turns Delay into Latency

Configuration Complexity

- CBS - [IEEE Std. 802.1Qav]
  - per class shaping is easy to configure, especially if only done only at the talkers
  - per class shaping can lead to buffer “residue” and constant delays (IEEE Std. 802.1Q: L.3.1.3 Permanent delay) in daisy chain networks
  - per flow policing is desirable to prevent intra-class interference and guarantee burst sizes

- ATS - [IEEE Std. 802.1Qcr]
  - per flow is easy to configure if only done only at the talkers, but means significant effort in the network
  - per destination shaping is a viable option
  - per flow policing is desirable to prevent interference and guarantee burst sizes

- TAS - [IEEE Std. 802.1Qav]
  - bus mode is very inefficient in larger networks
  - phased mode is very complex to configure (NP-hard)

Solution Proposal

- If the network is sufficiently small (low number of hops) and link speeds increase from the leaves to the root
  - One can use credit-based shaping only at the talkers to guarantee milisecond network transport latency
  - No re-shaping within the network is required
  - Policing at the Talkers is required to ensure deterministic behaviour

- Notes
  - Shaping at sources is simple: depends only on source
  - Combine with priority where needed
  - Prevent loss of data for all flows, especially best effort traffic, which may be retransmitted

Details: J. Walrand, M. Turner, and R. Myers, "An Architecture for In-Vehicle Networks," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 7, pp. 6335–6342, Jul. 2021.

Conclusions

- **Camera traffic aggregated in zonal controllers** provides a viable requirement to implement
  - few **high line-rate links** and
  - **few hops**
- This enables the use of the **Free Rider Principle** to transport
  - **low bandwidth control traffic at millisecond latency** through the network with
  - **shaping only at the sources** and thereby allowing for
  - a **low configuration complexity**
- Removing the need for frame packing efficiency (nPDU feature) further reduces latency and configuration complexity
- Policing of all flows[(1)] at the sources ensures QoS for all other flows

[(1)] Policing- check ingress rates and drop frames upon overload - can be done e.g. in a middleware, the communication stack or NIC HW

Ethernet & IP @ Automotive Technology Week
3 Nov 2021 - 4 Nov 2021

IEEE SA STANDARDS ASSOCIATION

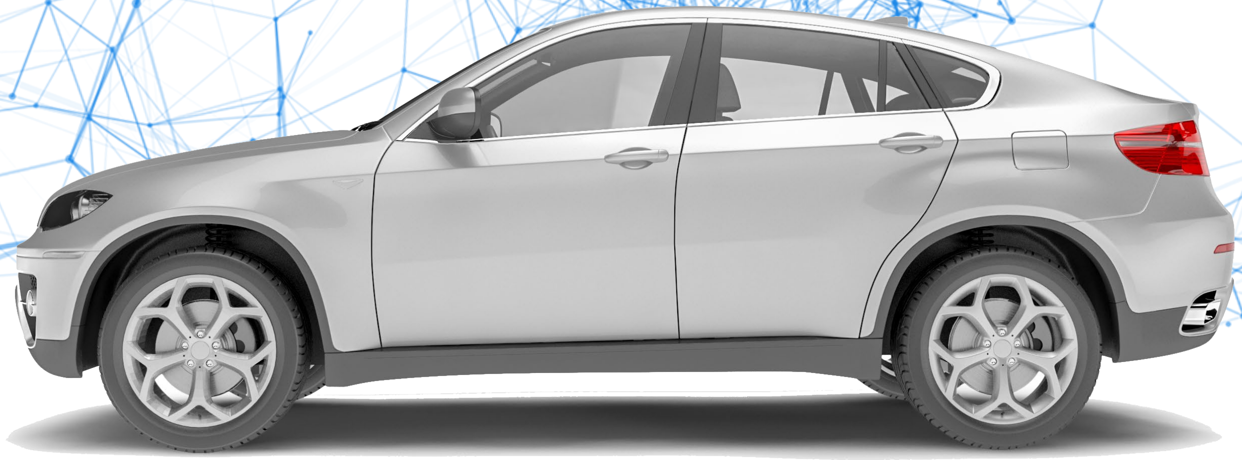

ThAnk you

walrand@berkeley.edu | max.turner@ethernovia.com

Berkeley UNIVERSITY OF CALIFORNIA

eTHERNOVIA
TRANSFORMING HOW CARS OF THE FUTURE ARE BUILT