Technische Universität Braunschweig

INSTITUT FÜR DATENTECHNIK UND KOMMUNIKATIONS-NETZE

Automotive Communication - From Buses to Switched Networks

Rolf Ernst
Institut für Datentechnik und Kommunikationsnetze
SIES, Porto 2013

Overview

- introduction - automotive embedded system architectures
- automotive network design
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

Overview

- **introduction - automotive embedded system architectures**
- automotive network design
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

Motivation – automotive networks

- **today's automotive networks use bus based complex networks**
  - **hundreds of functions**
  - **thousands of tasks**
  - **50+ ECUs**
  - **networked control**
  - **many suppliers**
  - **heterogeneous**

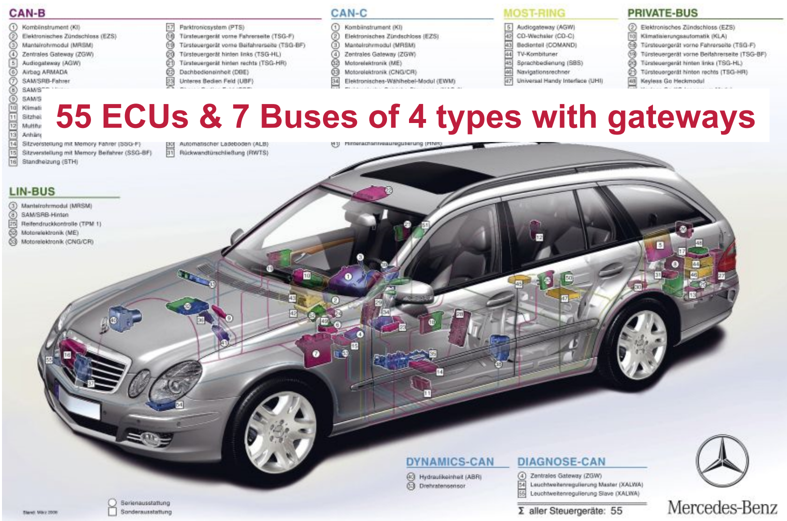

source: Daimler

- **networks are an efficient platform for systems integration**

Automotive function networks

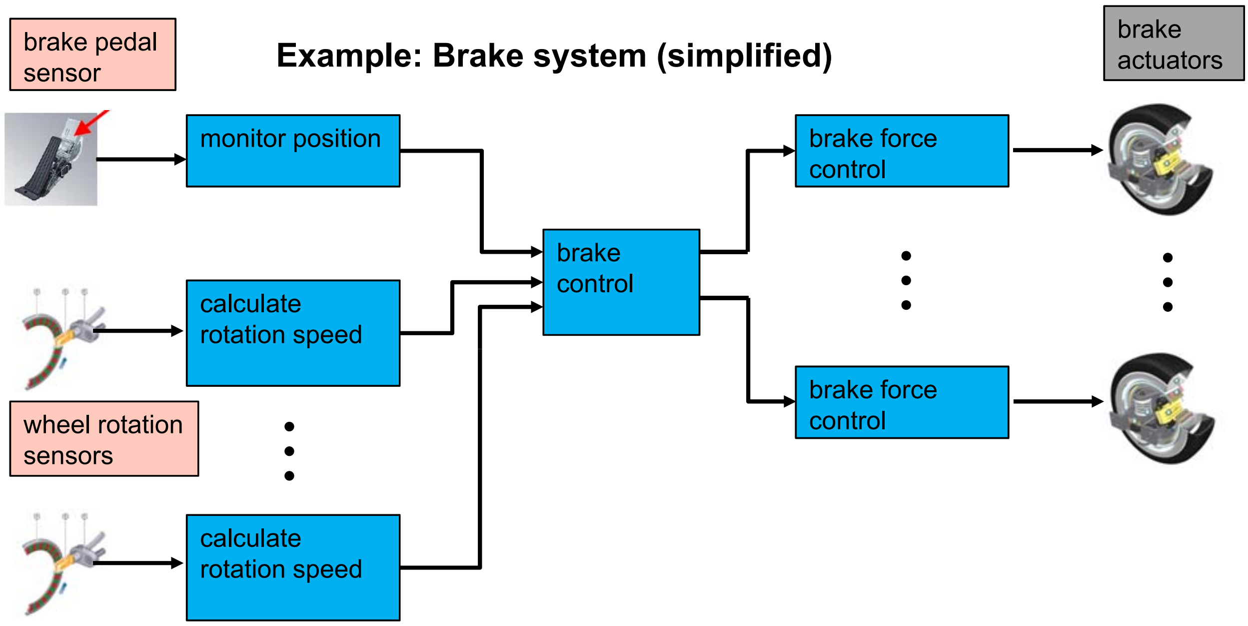

Functional architecture

- **periodic, sporadic or mixed task and communication activation**
- **main communication semantics**
  - **publish-subscribe**
  - **send-receive**
  - **client-server**
- **register and FIFO communication**
- **defined e.g. in automotive software standard AUTOSAR**

Register communication effects

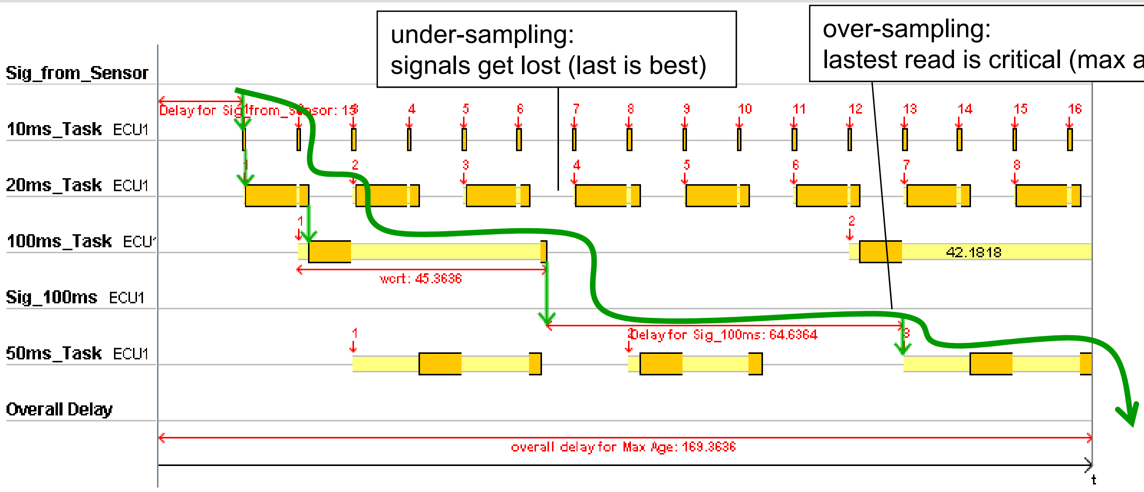

- only part of the messages utilized – „last is best“
- type of communication (register or FIFO) depends on application model – have to support both

Networks used for systems integration

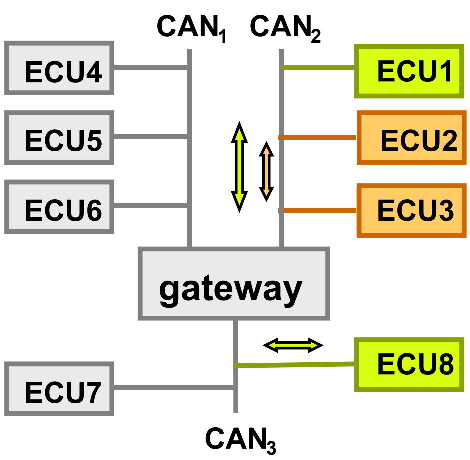

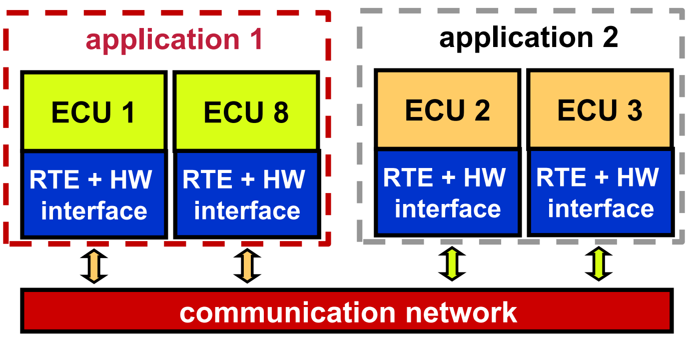

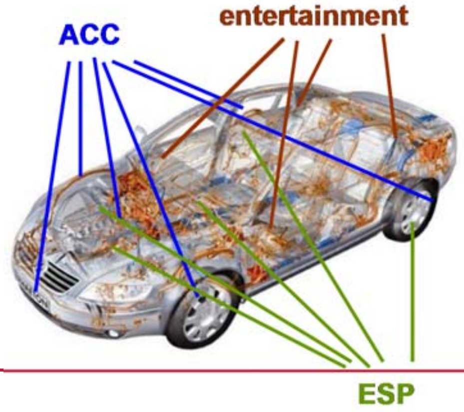

Automotive RTE - AUTOSAR

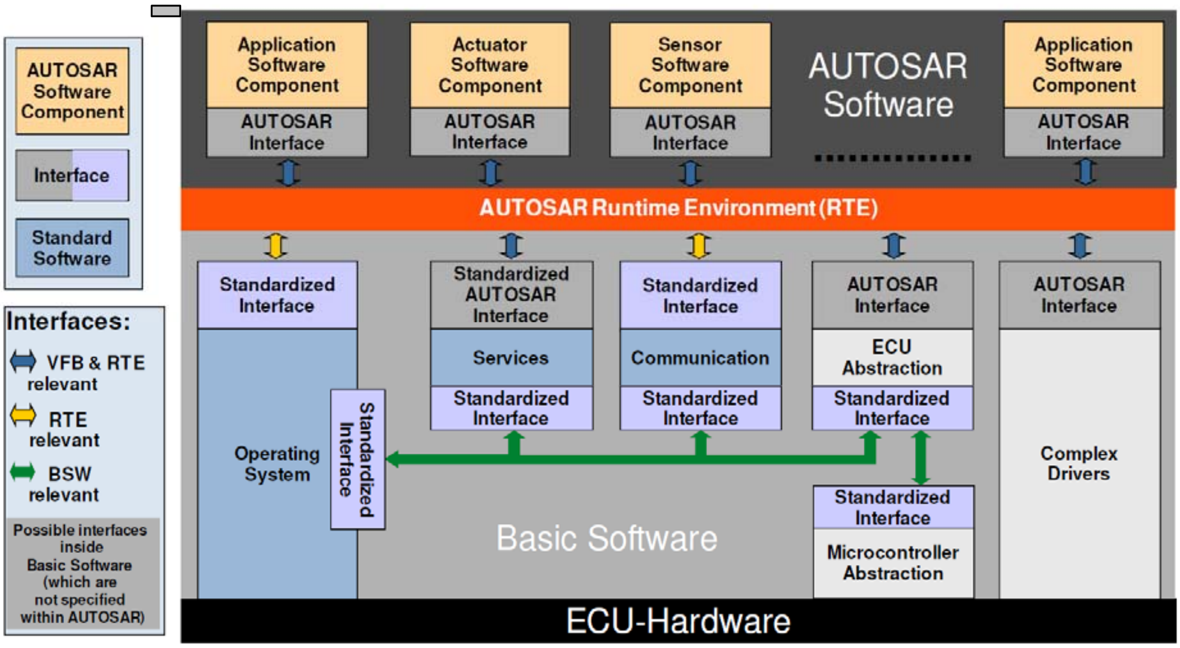

Source: Simon Fürst, BMW, et al., www.autosar.org

Application and Session Layers in AUTOSAR

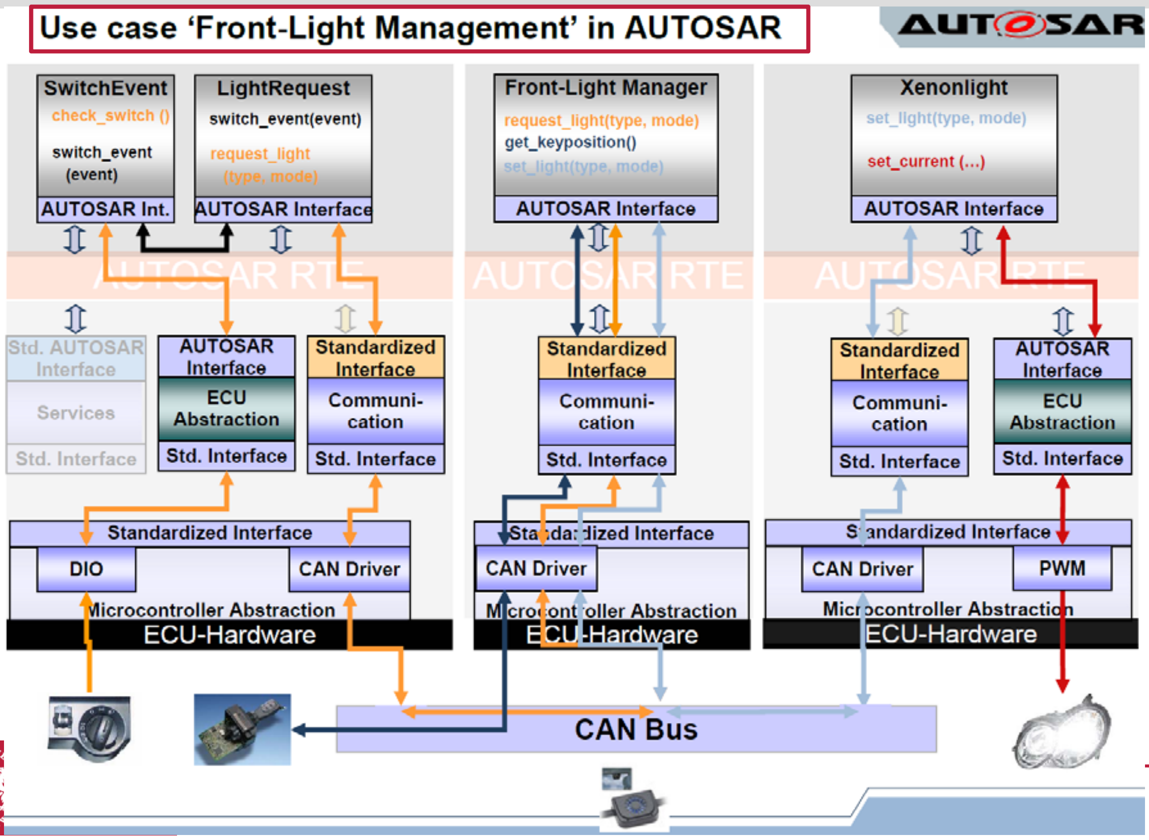

Source: www.autosar.org

Function integration - Mixed criticality

- **safety standards** require „**sufficient independence**“ (IEC 61508) or „**freedom from interference**“ (ISO26262) between safety critical and non critical functions
- end-to-end requirements include **function *and* timing**
- network and RTE have to guarantee independence
  - guarantee must be provided at **highest level of criticality** on network

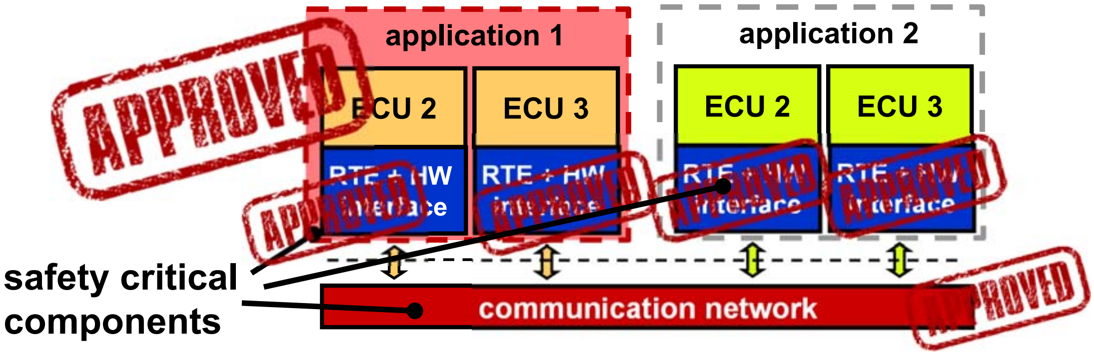

Overview

- introduction - automotive embedded system architectures
- **automotive network design**
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

SoA: Bus-based communication

- straightforward support of publisher-subscriber mechanism
- several application specific standards, CAN, FlexRay, LIN, …
  - < 100kbit … 10Mbit (FlexRay, CAN FD) data rate
- predictable scheduling: fixed priority or TDMA or slotted ring (MOST)
- routing by dedicated gateway
  - low speed allows SW implementation

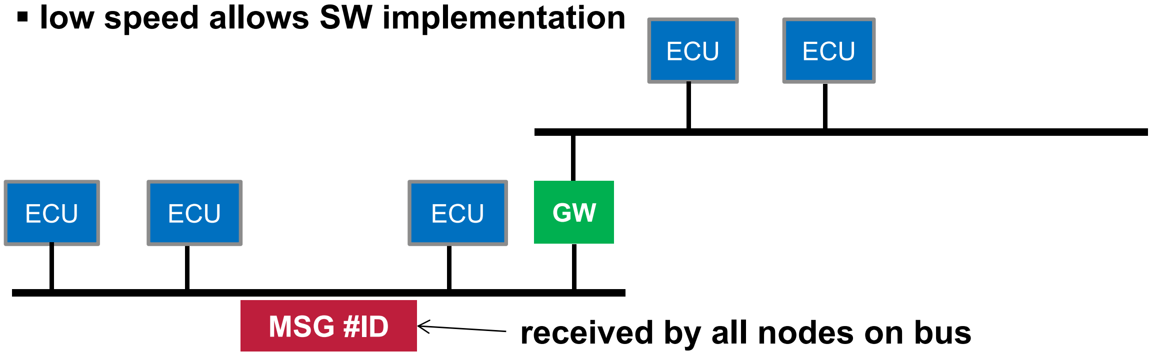

received by all nodes on bus
sender is not aware of receivers

Data Link Layer – signal packaging

- smaller „signals“ packed in larger frames to reduce com. overhead
  - frame activation pattern periodic or sporadic (i.e. immediate)
  - sent to all potential receivers (no bus overhead)
- tradeoff: activation frequency vs. overhead per signal

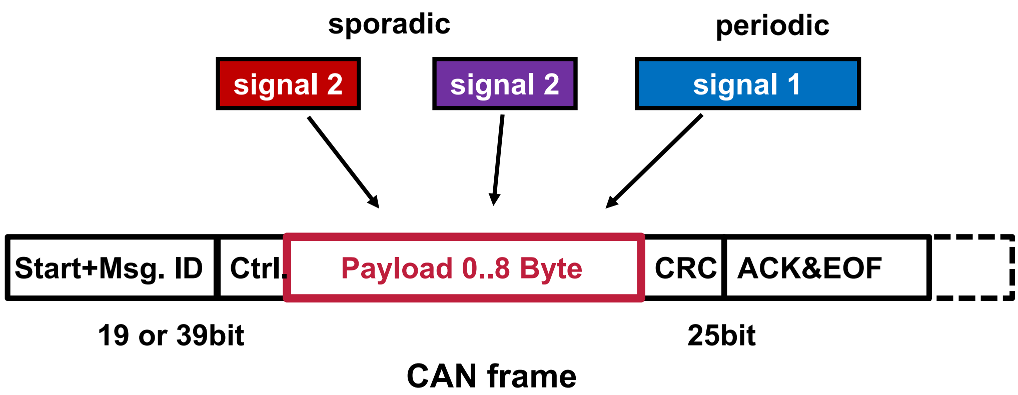

Current automotive network – protection

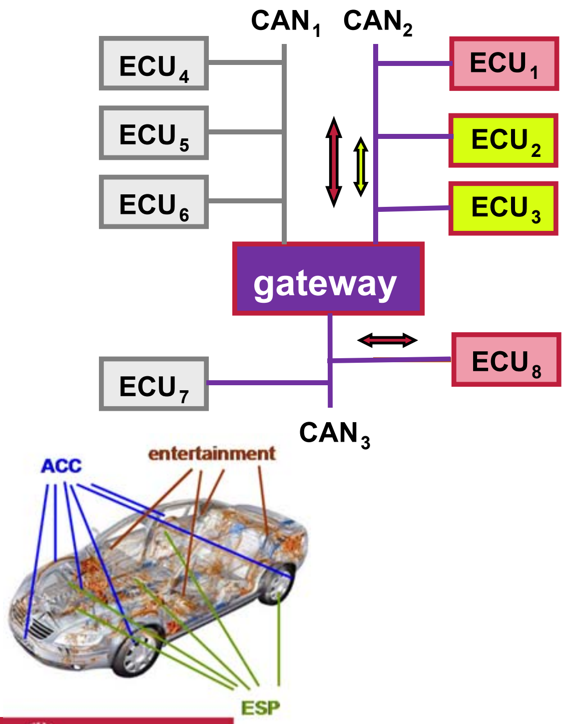

- application of safety standard affects large part of the system
  - *„freedom from interference!“*
  - single criticality on safety critical ECUs
  - mixed criticality on other parts
- *all ECUs on mixed critical buses must be trusted!*
  - no secure sender identification
  - no „babbling idiot“ (or denial-of-service-attack) on mixed critical bused allowed
  - hard requirement on classical buses

„Sufficient isolation“ on CAN

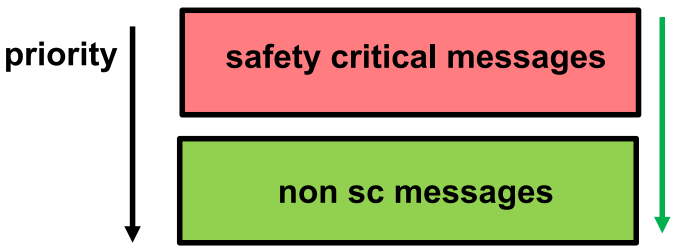

- WC design process
- WC frame timing independent of lower priority messages

- WC or other design process

- Msg. IDs used for arbitration - Static Priority Non Premptive (SPNP)
- asymetric separation using „criticality as a priority“
  - high priority critical messages are „sufficiently independent“
- separation leads to non-optimal priority assignment
- allow „occasional“ loss of non (time) critical frames due to overload
  - requires end-to-end protection on higher protocol level for event driven comm.
  - „less than worst case design“ using soft deadlines

„Less than worst case design“ on CAN – Example

FGA USE CASE

C-CAN Distribution Analysis:

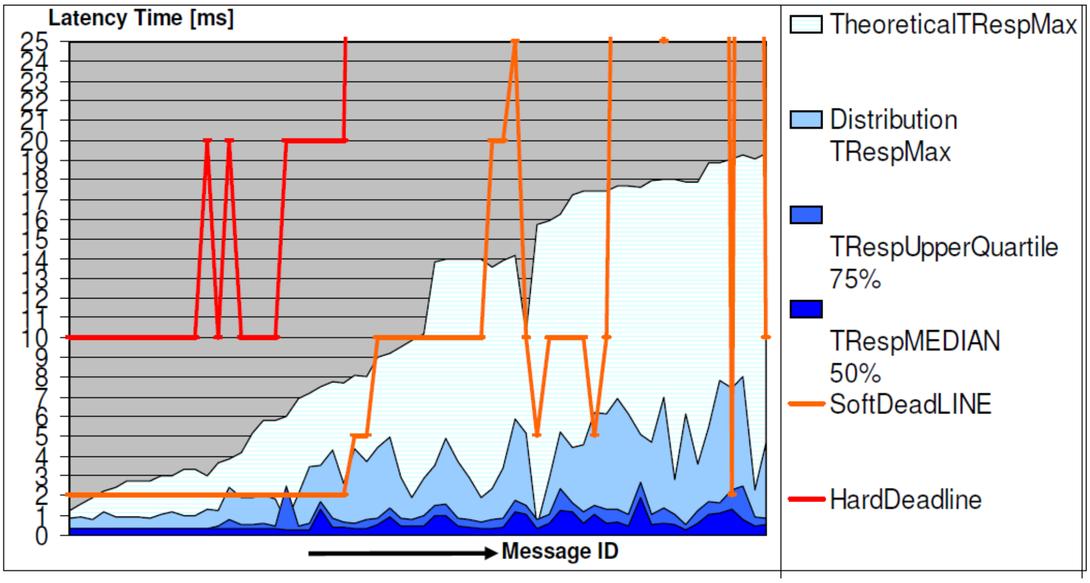

Source: IIF. AMATO, M.MELANI, "Timing & Distribution analysis with SymTA/S at FIAT Group Automobiles", at workshop on 'Timing Analysis in Verification & Validation for real time embedded systems' Turin, 9.6.2011

Overview

- introduction - automotive embedded system architectures
- automotive network design
- **trends and requirements to future automotive networks**
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

Future automotive network requirements

- current trends
  - Trend 1: growing network complexity due to highly interactive functions
    - *latency critical and safety critical traffic*
    - *some best effort traffic (e.g. diagnosis)*
  - Trend 2: steep increase of network bandwidth requirements with new traffic patterns from infotainment and driver assistance systems – streaming patterns
    - *primarily throughput critical traffic, moderate latency critical*
    - *some traffic safety critical*
  - Trend 3: networks with IP traffic via car-to-X communication
    - *primarily best effort*

Trend 1: Growing network complexity

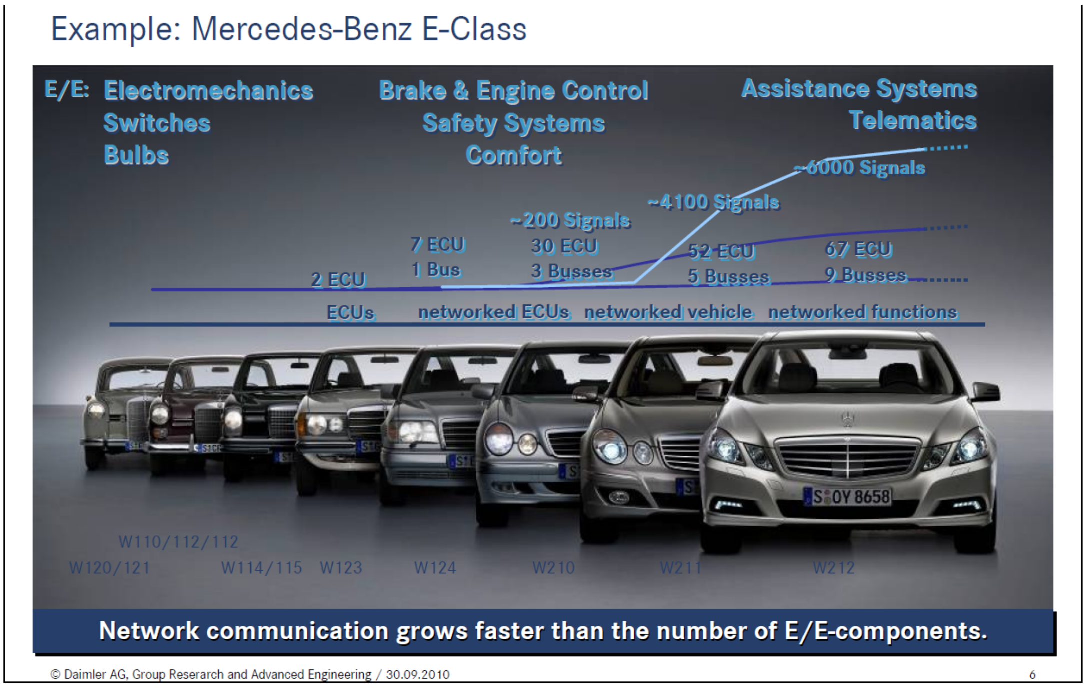

source: T. Bone, Daimler

Trend 2: ADAS using complex sensor systems

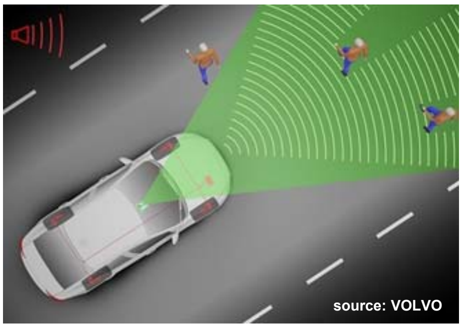

- advanced driver assistance systems
  - collision avoidance, lane keeping assist, blind spot assist, ...
  - fastest growing automotive segment
  - Combines radar, several cameras
  - Need for high data rate real-time communications…
- new requirement: safety + high performance + real-time !

Future automotive network requirements

- trends will trigger transition from communication over buses &gateways to switched networks
  - automotive networks – strongest candidate: switched Ethernet
    - 2-wire physical medium (OPEN alliance)

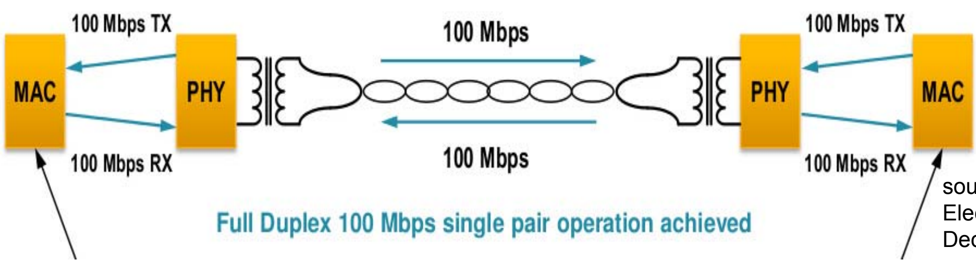

source:
Electronic Design
Dec 13, 2011

From buses w. gateways to switched networks 1/2

- expectations
  - keep heterogeneous networks
    - cost efficiency and experience in CAN, LIN networks
    - transmit existing network traffic over switched network
    - replace MOST – infotainment streams
  - include new traffic – camera streams
    - compressed or raw, different frame sizes and rates
  - keep design process with many variants, annual „facelifts“, modular product lines
  - cost sensitive solution - avoid high component and qualification cost

From buses w gateways to switched networks 2/2

- constraints
  - time critical functions
    - control end-to-end timing
  - safety critical functions
    - provide sufficient independence („freedom from interference“) from non critical parts with lower quality
    - provide sufficient reliability of the affected components
  - security
    - avoid illegal access and denial-of service attacks
      - ensure managing access rights and secure sender identifcation (buses?)

Overview

- introduction - automotive embedded system architectures
- automotive network design
- trends and requirements to future automotive networks
- **Switched Ethernet as an automotive backbone**
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

Ethernet as a Backbone Network

- current generation introduces Ethernet as infotainment network
- future use as **backbone**
  - inherits requirements and constraints of existing buses

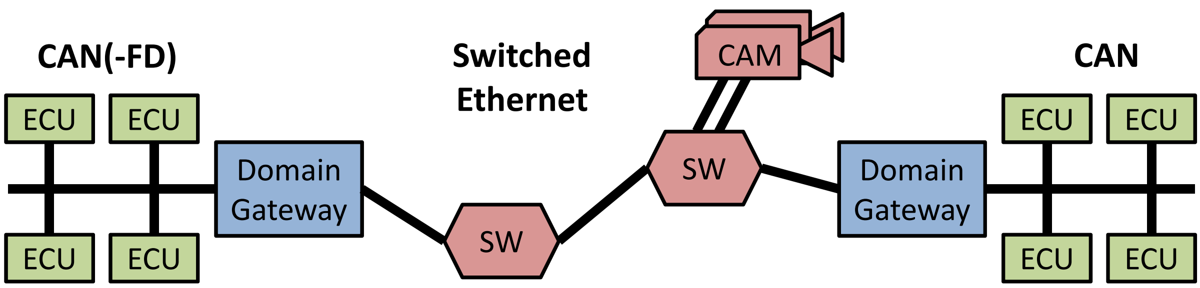

Exploiting Ethernet

- data link level: what frame sizes, packeting, error handling
- network level: what network protocol (higher levels?) switch scheduling and arbitration, ...
- handling of very different traffic requirements (bandwidth, latency, QoS, message sizes, ...)
- network topologies
- lossy or lossless?
  - publisher subscriber/register semantics – typically single message may be lost
  - send-receive, client-server: no loss, must be protected

Ethernet as a backbone network – Switch

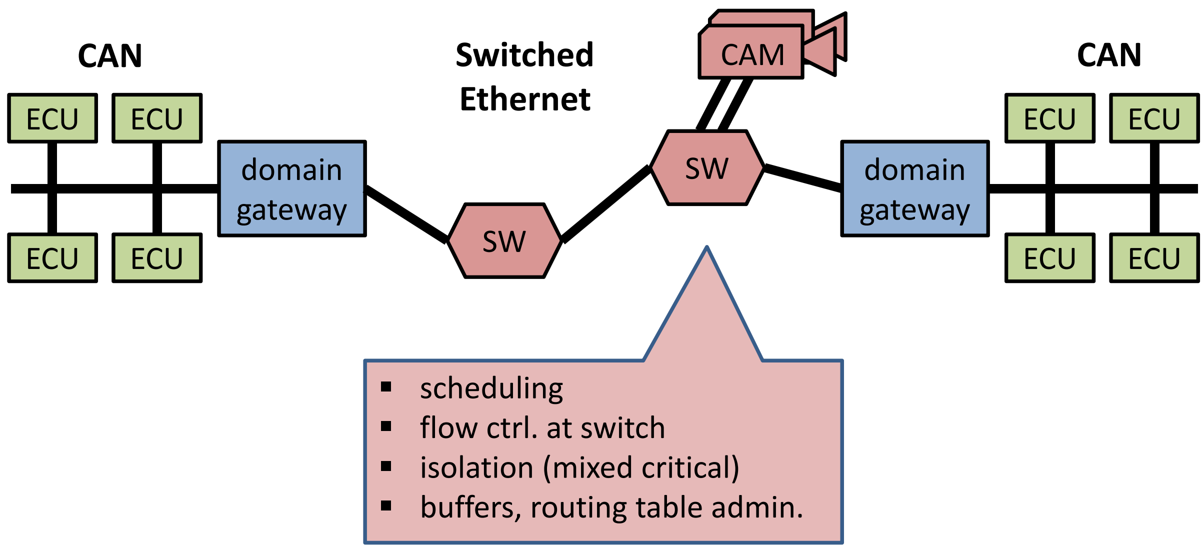

Ethernet frame

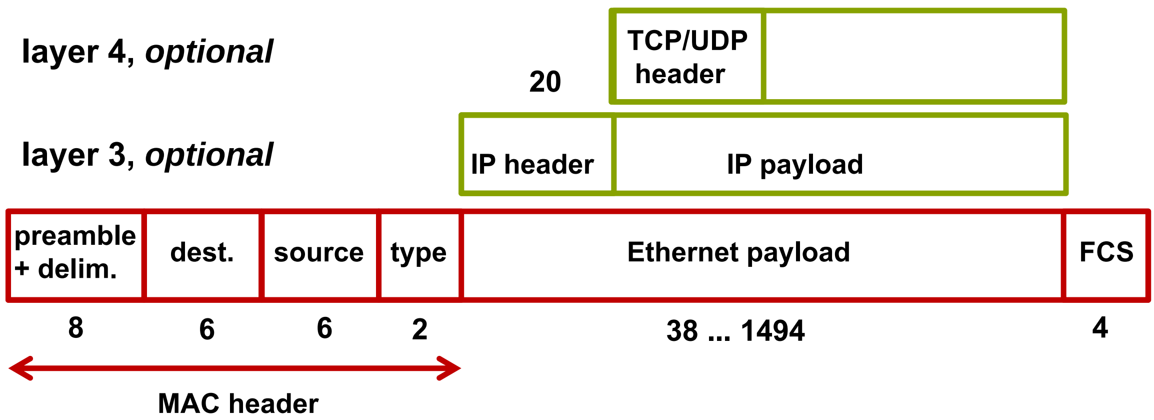

- 22 bit MAC header generally sufficient for routing in local network
- IP protocol e.g. for external communication or use of TCP/UDP
- layer 4 protocol for end-to-end control incl. sequence and port addr.
  - TCP builds up connection - handles packet losses (additional overhead)
  - UDP is connectionless – no error handling (best for loss free conections)

Switched Ethernet - consequences

- many more options and parameters than in current networks
- heterogenous protocol combinations possible and useful, e.g.
  - „plain“ Ethernet or UDP/IP for fast transmission with worst case design along packet route (no losses)
  - TCP/IP for secured communication over (potentially) lossy routes
- no synchronization needed
- point-to-point communication allows secure sender identification in switch supporting safety firewalls
  - improvement for mixed critical designs

Ethernet – Economic issues

- low cost – try to use established high volume standard
  - avoid cost of ownership in automotive industry
  - exploit competitive market
- many switch functions implemented in HW (performance)
  - less flexibility when using standard switch ICs
  - bus side of domain gateway ctrl. can be SW

Ethernet – overheads

- Overheads
  - Ethernet: 42 B incl. IF gap
  - IP: 20 Bytes
  - UDP: 8 Bytes
  - TCP: 20 Bytes
- Min. transmission times per hop:
  - 100Mbit/s
    - Min: 6.72 µs
    - Max: 122 µs
  - 1Gbit/s
    - Min: 0.672 µs
    - Max: 12.2 µs

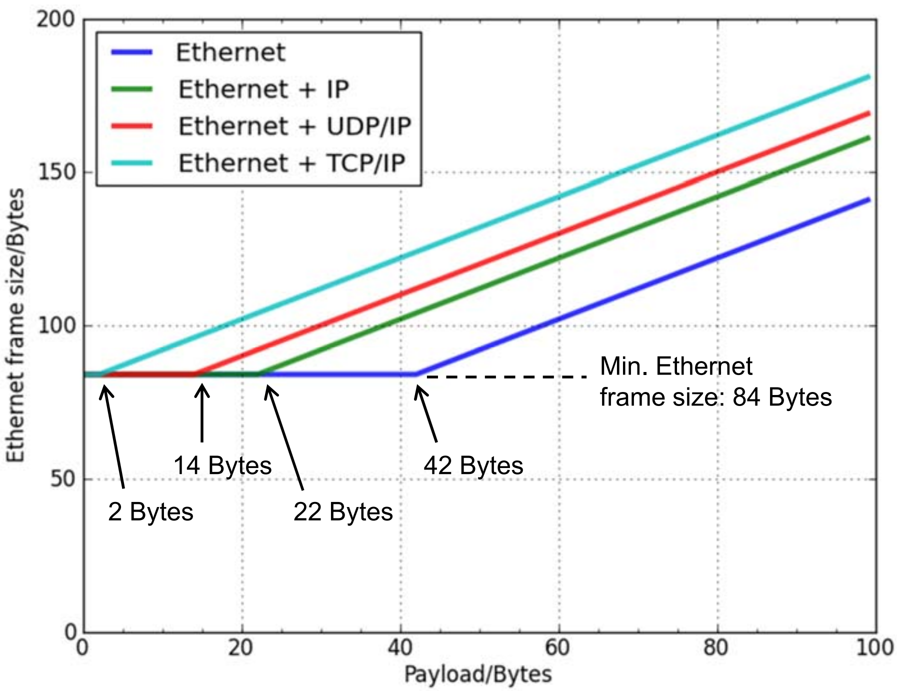

Ethernet as a Backbone Network for CAN Traffic

- **Worst-case CAN frame size (incl. stuffing bits): 77 + 10 * payload_in_bytes [bit]**
- **evaluation**
  - **max speed CAN: 1 Mbit/s**
  - **Send traffic from 100% utilized CAN bus (peak load) over Ethernet**

	100 Mbit/s Ethernet		1Gbit/s Ethernet	
	"Raw", IP, UDP/IP	TCP/IP	"Raw, IP, UDP/IP"	TCP/IP
8 Bytes CAN payload (max.)	4.1 Mbit/s 6.72 µs/fr	4.4 Mbit/s 7.2 µs/fr	4.1 Mbit/s 0.672 µs/fr	4.4 Mbit/s 0.72 µs/fr
1 Byte CAN payload	7.4 Mbit/s 6.72 µs/fr		7.4 Mbit/s 0.672 µs/fr	

Camera data transmission via Ethernet

- **example setup: stereo camera system**
  - **2 cameras for front view applications and depth extraction**
  - **Image data rate e.g. 1000x1000: 28,6 MBytes/s * 2 cameras (long Ethernet packets: 1500Bytes/fr)**
- **significant data rate challenges**
  - **system requirements:**
    - **100 BaseT: 12,5 MBytes/s, utilization: 462 %, 120µs/fr**
    - **1000 BaseT: 125 MBytes/s, utilization: 46,2% 12µs/fr**
  - **camera-based systems require Gigabit Ethernet or compression**

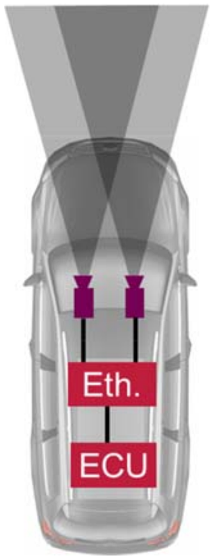

Frame sizes and packaging - consequences

- several complete CAN traffics can be transmitted over 100Mbit/s and 1Gbit/s Ethernet, in different packeting schemes
  - delay times are < 10 μs/hop for CAN traffic
  - switched Ethernet suitable, optimization only required to decrease cost
- camera traffic
  - mainly useful for 1Gb/s, otherwise compression needed
  - for 1Gbit/s: delay < 12μs/hop

Overview

- introduction - automotive embedded system architectures
- automotive network design
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- **related developments**
- AVB as top candidate? – challenges and possible directions
- conclusions

AFDX - network

- 100Mbit/s, UDP/IP
- switches: fixed forwarding table, FIFO buffering
- „virtual links“ (VL) provide multicast communication to implement „virtual“ legacy bus communication - uses IP protocol
- sender traffic shaping to control network load and VL timing interference
  - Lmax: maximum length of Ethernet Frames on this VL
  - BAG: Bandwidth Allocation Gap: Min. dist. between VL frames
  - no shaping in switches
    - timing dependencies through FIFO scheduling
    - requires performance analysis based on BAG data
    - *timing analysis e.g. group of Christian Fraboul, Toulouse*
- all network certified on highest level (DAL A)

AFDX VL and BAG - Example

- VL_2 and VL_3 implement multicasting

BAG

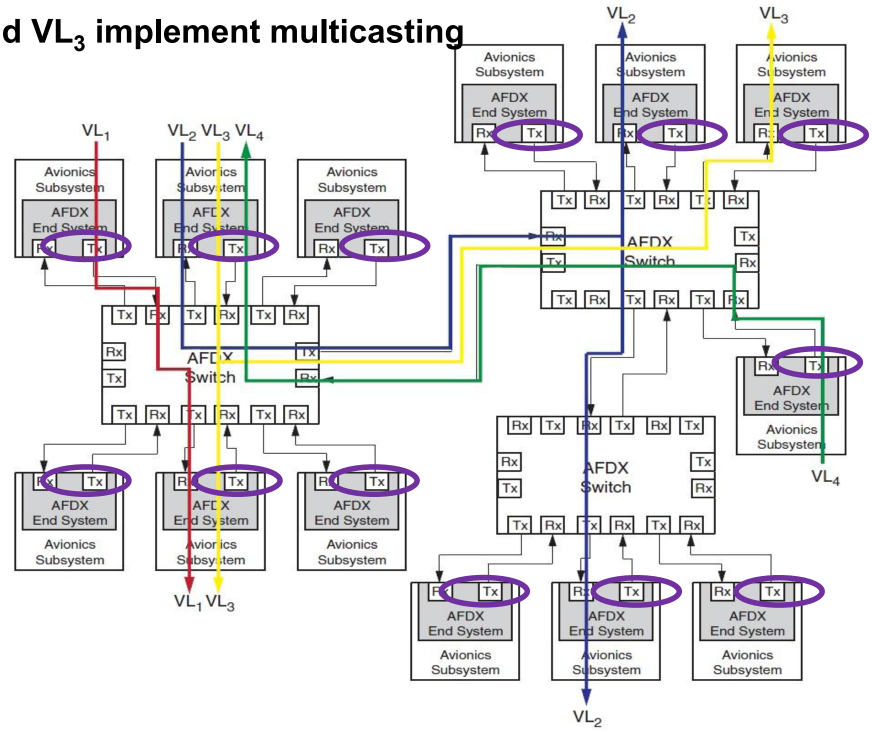

TTEthernet (SAE AS6802)

- multiple IEEE 802.3 speeds (100Mbit, 1Gbit)
  - introduces precise time service for network synchronization
- introduces 3 traffic classes in priority scheduled Ethernet
  - time triggered (TT) frames sent periodically at highest priority
    - non TT frames are held back if transmission can interfere with TT frame transmission - reaches high precision
    - feature can be turned off to optimize throughput
  - rate constraint (RC) frames have rate limitations like AFDX BAG
  - best effort (BE) frames carry other traffic at lowest priority
- achieves sufficient independence
  - *but for TT frames only*

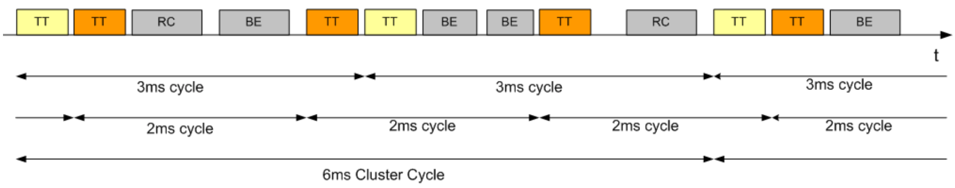

source: TTTech white paper

PROFINET

- time triggered (TDMA) communication with segments as in FlexRay
  - uses precise time base IEEE 1588 (PTP)
- 3 segments for 3 communication classes
  - Non Real-Time class NRT – TCP/IP
  - Real-time class RT
    - synchronous and asynchronous transmission
    - uses priorities
  - Isochronous Real-time IRT
    - unbuffered globally time synchronous frames
    - no higher level protocols (level 2: MAC addresses only)

PROFINET frames

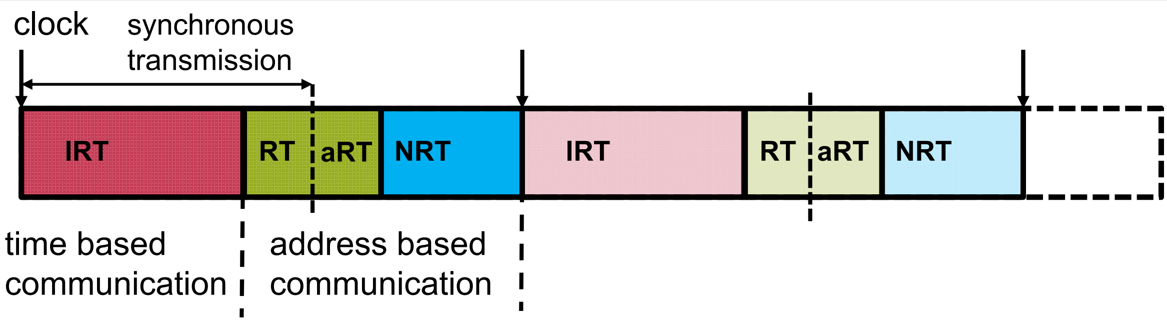

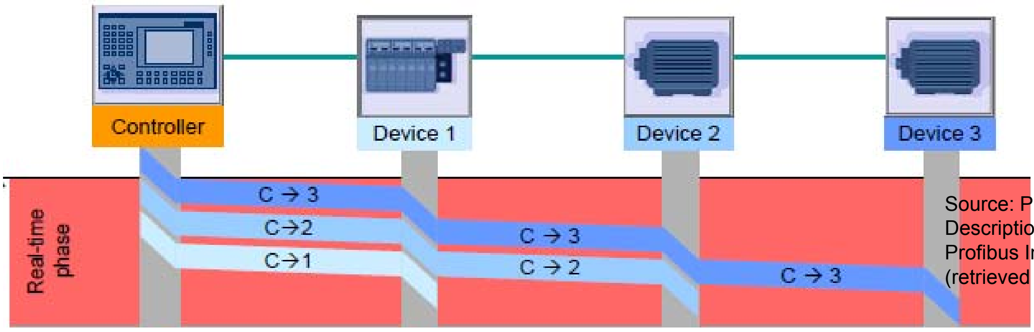

Source: Profinet System Description, Profibus International (retrieved 2011)

IRT segment – Slot assignment

Ethernet AVB

- **Based on strict priority Ethernet (IEEE 802.1Q)**
- **“Real-time” streams register bandwidth via reservation protocol**
- **AVB Classes A + B prioritized: A > B > non-Real-Time**
- **FIFO queuing in each class**
- **Traffic Shaping at Class A+B queues**

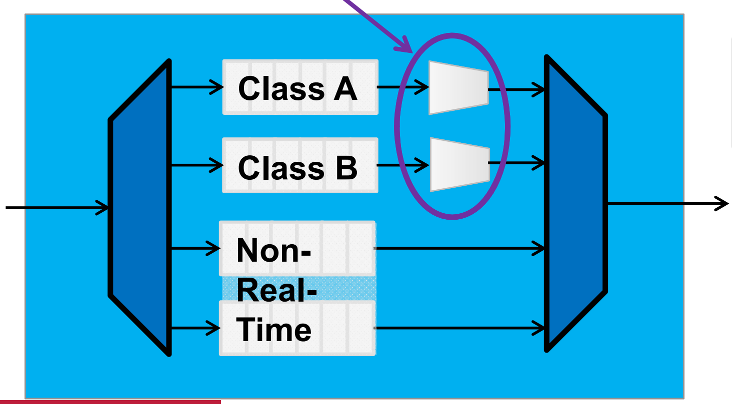

Traffic Shaper:
→ Throttles high-priority streams
→ Allows bursts after blocking

SoA - Consequences

- AFDX target is closest to needs, but design process requirements and costs are high, high determinism does not match automotive mixed criticality and flexibility requirements
- AVB will be widely used as real-time Ethernet standard
  - highest expected volume for shared cost across industries
  - flow control mechanisms and no. priorities very limited
  - access control and packaging must be handled at boundary gateways
  - „time sensitive network“ (TSN) extensions with time synchronization under standardization when needed
- AVB most likely candidate
  - *flow control and scheduling to be revisited*?

Overview

- introduction - automotive embedded system architectures
- automotive network design
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- conclusions

Ethernet as a backbone network – AVB

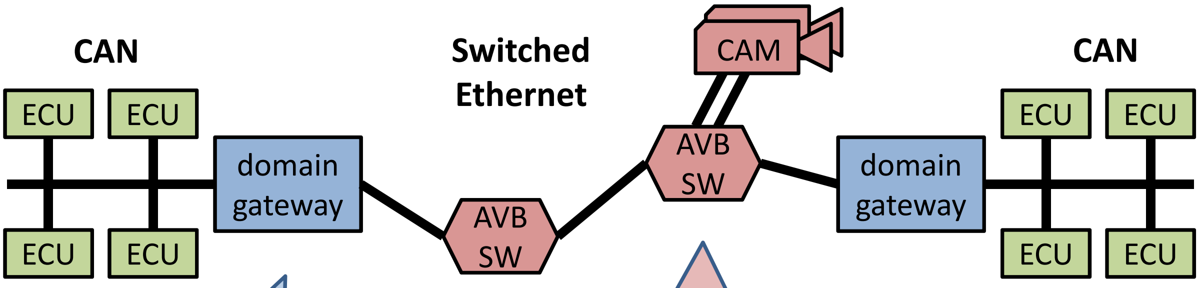

- traffic mapping and packeting
- access ctrl.

- scheduling, flow ctrl. at switch
- isolation (mixed critical)
- buffers, routing tables

CAN-Ethernet Gateway – mapping to few priorities

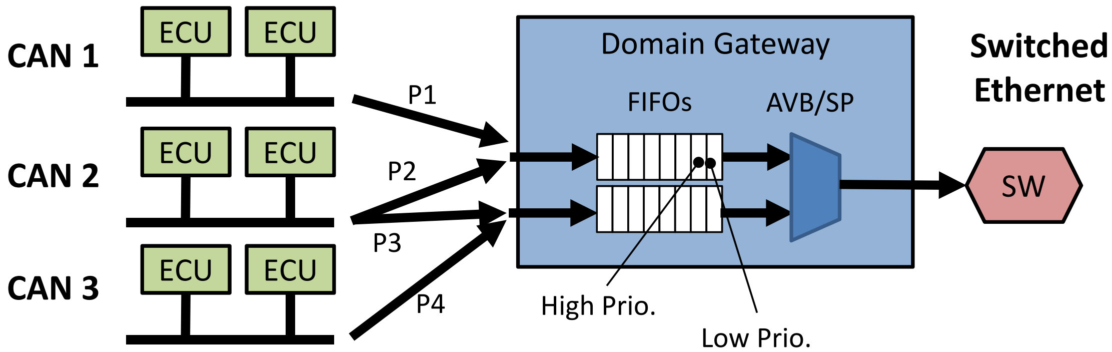

- **many CAN IDs must be mapped to few Ethernet priorities**
  - **CAN supports up to 2^{11} or 2^{29} IDs and as many priorities**
  - **Ethernet supports max. 8 priorities**
- **Head of line blocking at FIFOs in gateway output port**
  - ***FIFOs can be reordered by fast domain gateway processor - not in switch***

Mixed critical traffic mapping in AVB

- asymmetric isolation on CAN can be preserved
  - but limited by few priority levels
  - all messages mapped to one queue must have same criticality level
    - avoids uncontrolled line blocking in switch

Ethernet AVB - traffic shaping

- AVB Classes A + B prioritized: A > B > non-Real-Time
- non real-time can use further classes
- FIFO queuing in each class
- Traffic Shaping at Class A+B queues

Traffic Shaper:
→ Throttles high-priority streams
→ Allows bursts after blocking

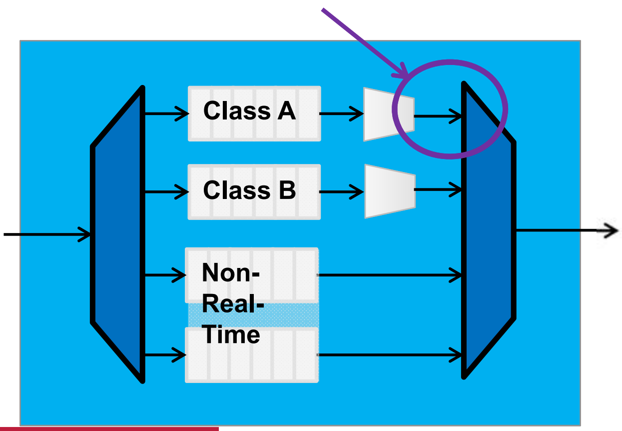

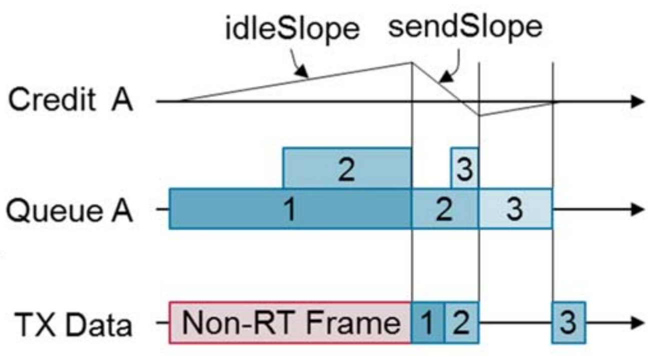

Star topology results for varying over-reservation

- experiment: constant number of 12 I/O nodes
- same priority blocking in class FIFOs dominates timing
  - direct shaping result
- latency increase for both A *and* B classes
  - shaped class B latency cannot profit from class A shaping
- similar observations for other topologies and traffic
- *details SIES 2012 and IECON 2012*

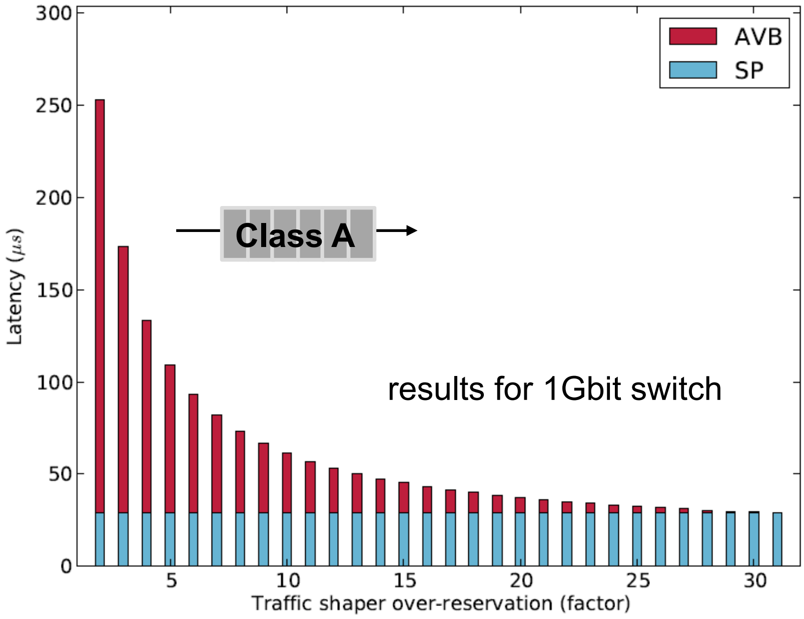

Large network delays possible

- larger industrial example as AVB benchmark (“Deggendorf”) (IECON 12)

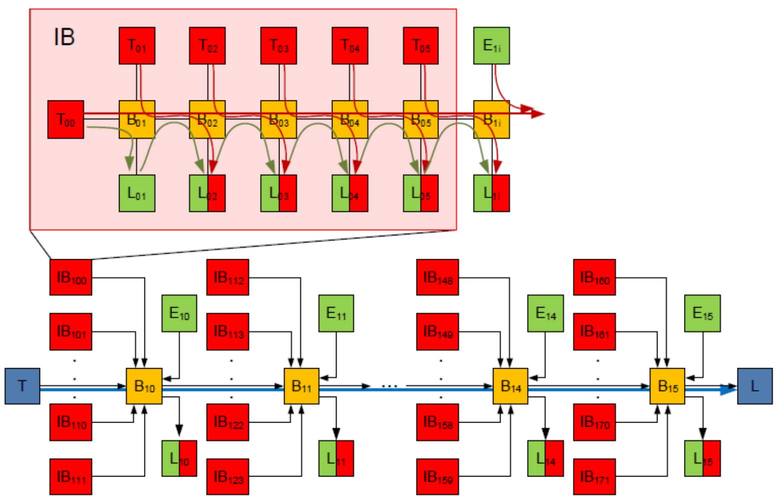

Scenario	Frames in burst	Bridge Delay	Latency
Sim. (no drops)	7	893.76 us	5.493 ms
Sim. (w. drops)	11 (12 effective)	1.434 ms	8.733 ms
Analysis	11 (12 effective)	1.422 – 1.566 ms	8.975 ms

Replace Strict Priority (SP) by WRR scheduling ?

- very complex traffic scenario – SP less suitable?
- Weighted Round Robin (WRR) strategy
  - approximation of Weighted Fair Queuing
  - serves a predefined maximum number (weight) of packets from each queue in every scheduling round
  - empty queues are skipped
  - weights do no accumulate (unused slots expire)
- WRR rationale
  - available in many commercial switches
  - minimum guarantees on all levels
    - can be used for isolation
  - scheduling can be adapted to predictable traffic volume

Weighted Round Robin – principle

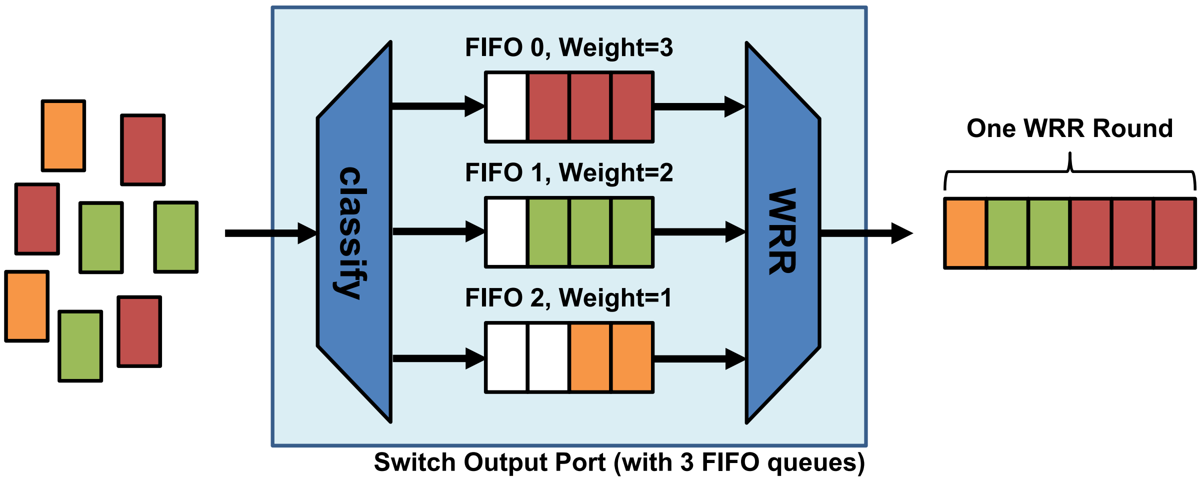

- improved worst-case response time analysis algorithm allows direct comparison with static priority scheduling
- *will be published with experiments at CODES 2013*

Weighted Round Robin – Experiment

- star topology (CODES paper also line)
  - one controller node, multiple I/O nodes

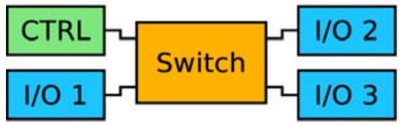

- traffic description
  - two classes, HP has more share than LP (prioritization)
  - bidirectional communication using the same traffic characterization
  - higher jitter reflects queuing in switch

Class	HP	LP
Period	250 us	250 ms
Jitter (low and high)	250 us and 2250 us	250 ms and 2250 ms
Payload	10 Bytes	10 Bytes
WRR Weight	300	120

Weighted Round Robin – result for star topology

- number of I/O nodes are increased from 1 to 23
- Y-axis shows worst-case end-to-end latencies for each scenario
- schedulers
  - SP
  - WRR
  - TDMA

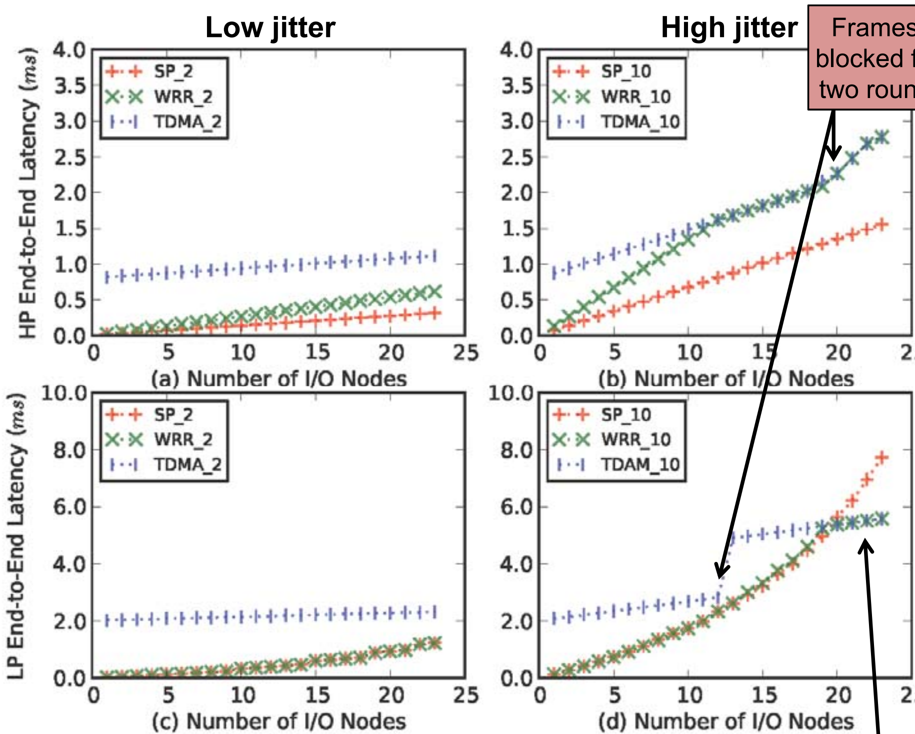

Frames blocked for two rounds

In high-load scenarios WRR (and TDMA) give tighter latency guarantees for LP traffic than SP

Weighted Round Robin – Results and Challenges

- prioritization of traffic streams by weight ratios
  - higher priority -> larger weight
- WRR does not help to reduce latency
  - improves low priority only in high load situations at the cost of high priority traffic
- complex dependencies between latency and weight ratios
  - If weights are too small, a frame might be blocked for multiple rounds
  - Best-effort frames (if available) contribute to every scheduling round
- WRR would be challenging in design process
  - weights must likely be adapted in case of changes
- WRR appears less useful than SP to serve latency critical traffic

AVB switch scheduling - consequences

- switch queuing + shaping is main reason for delays and traffic dependencies
- AVB is optimized for streaming, not for latency guarantees
  - SP provides better results for latency critical traffic
- possible solution
  - selectively turn shapers off for high priority latency critical traffic – class A
  - use access ctrl. for this traffic
  - keep shapers for high and guaranteed throughput traffic – class B
  - use non real-time queues for other classes
    - must avoid starvation as consequence of non worst case design
      - *lossy design with end-to-end protection (TCP)?*
    - include sufficient independence where needed (WRR?)

Suggestions from automotive industry 1/2

Shaping latency vs. arbitration latency

Jonas Diemer, Daniel Thiele and Rolf Ernst, **"Formal Worst-Case Timing Analysis of Ethernet Topologies with Strict-Priority and AVB Switching"** in *7th IEEE International Symposium on Industrial Embedded Systems (SIES12)*, Juni 2012

source: Stein, Racu: Ethernet Technology Day, Regensburg, 2012

Assumptions:
- Single Switch (1 Hop)
- Gbit Ethernet
- No Jumbo Frames
- Communication to 1 sink
- Sending of 2 minimum sized frames
- Shaper Overreservation: 2x

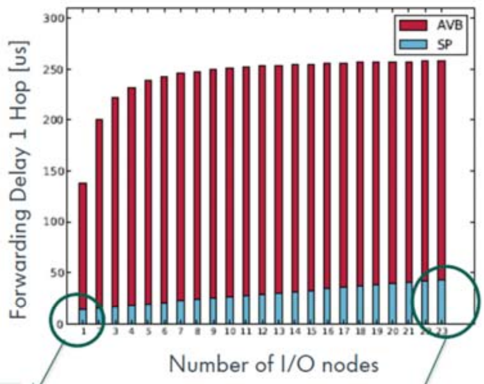

~130us for 100MBit

~432us for 100MBit

EEIN – EE-Architektur und Netzwerke
ELEKTRIK/ELEKTRONIK ENTWICKLUNG

Stein, Racu, Ethernet QoS @ Volkswagen, Ethernet Technology Day, Regensburg, 20.09.2012

Suggestions from automotive industry 1/2

Putting it all together – QoS Concept

source: Stein, Racu: Ethernet Technology Day, Regensburg, 2012

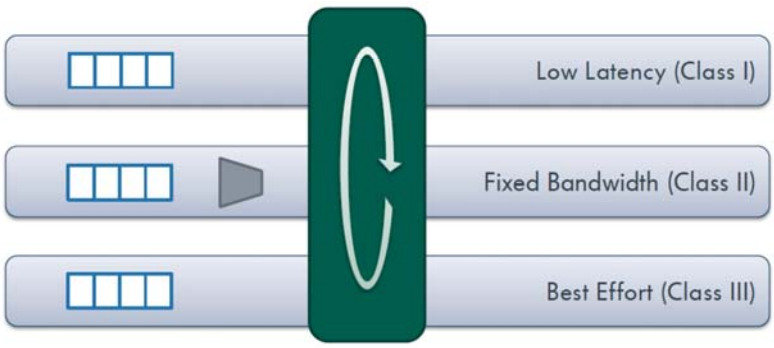

EEIN – EE-Architektur und Netzwerke
ELEKTRIK/ELEKTRONIK ENTWICKLUNG

Stein, Racu, Ethernet QoS @ Volkswagen, Ethernet Technology Day, Regensburg, 20.09.2012

Switch buffering and tables

- implementation **less restricted**
- **buffers**
  - overflow can break isolation if back pressure applied to whole switch
  - dynamic buffer assignment can break isolation and defy predictable response times
- **routing tables**
  - dynamic table management can break isolation and response times
  - table management must be static or at least predictable
- **switch back plane**
  - switch back plane must have predictable upper bound timing

Ethernet as a backbone network – terminal gateway

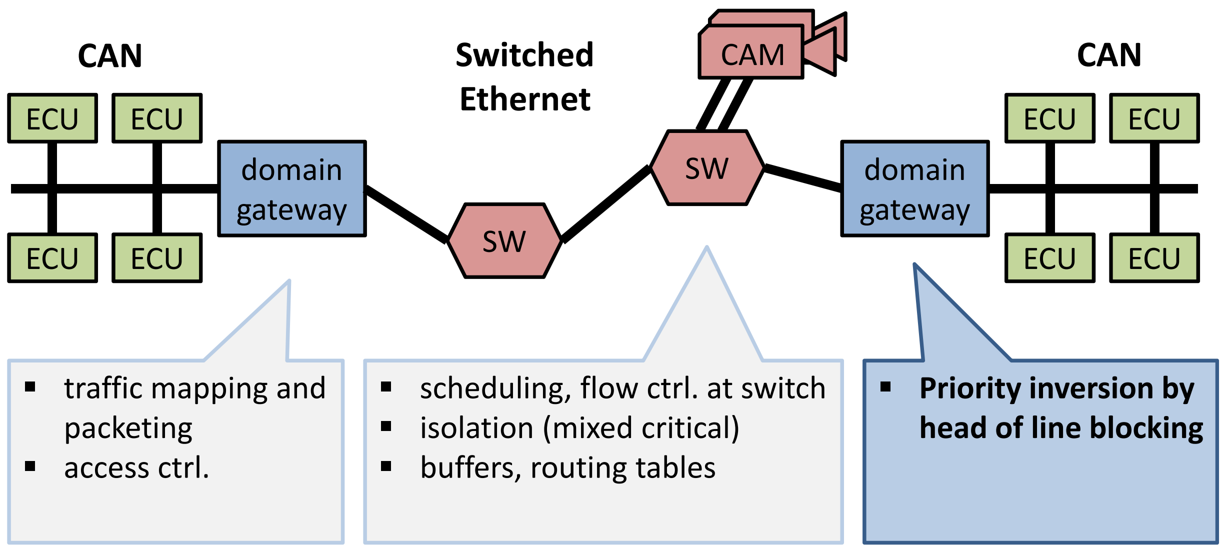

Head of Line Blocking at Ethernet-CAN Gateway

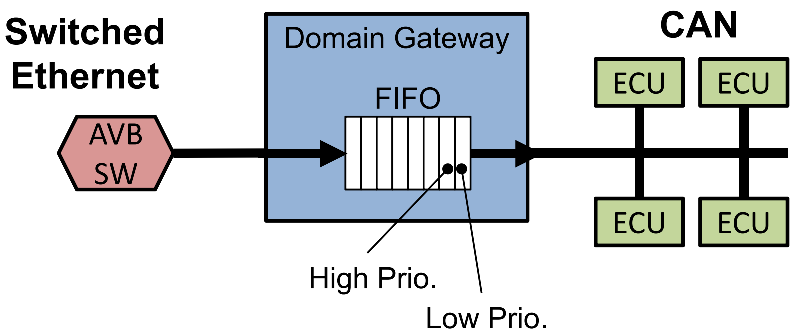

- low priority frame waits for CAN arbitration
- **high priority frame is stuck in FIFO** and inherits priority of low priority frame
- **order can be changed** by fast Domain Gateway processor
  - *not possible if inversion queues back into backbone network*
  - *queuing into network must be avoided*

Many more interesting issues – e.g. frame packing

- **publisher-subscriber** communication corresponds to **broadcast or multicast** on Switched Ethernet
- many frame packing alternatives
  - packaging for one or for multiple target domains
  - packaging for equal deadlines
  - …
- *optimization problem*

Future: End-to-end switched networks

- the future automotive network is part of an overall end-to-end system of switched networks
  - system level switched networks
  - ECU level networks such as PCIe where used
  - NoCs on multicore and manycore ICs

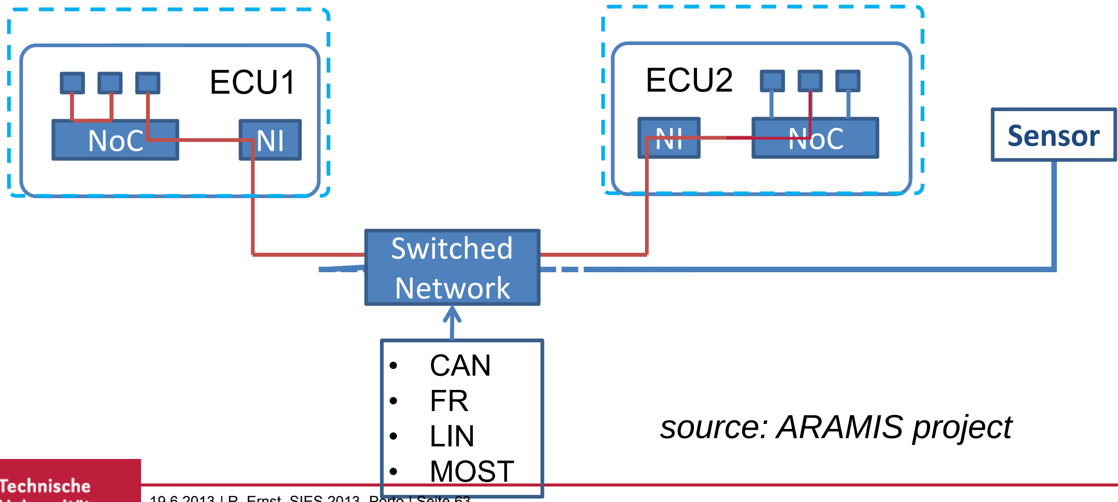

source: ARAMIS project

Overview

- introduction - automotive embedded system architectures
- automotive network design
- trends and requirements to future automotive networks
- Switched Ethernet as an automotive backbone
- related developments
- AVB as top candidate? – challenges and possible directions
- **conclusions**

Conclusion

- future generation automotive networks will use Switched Ethernet
  - must be tailored to automotive function and traffic trends
  - high bandwidth promises seemingly simple solution - *fallacy*
- different real-time requirements and efficient „freedom from interference“ under mixed criticality are major challenges
- Switched Ethernet requires global network view
  - network security, safety and timing are tightly interrelated
  - solutions possible, even incremental development of network functionality
  - *formal analysis* will be a cornerstone for safe and efficient design

Thank you!

Acknowledgement and references

- Acknowledgements
  - Thanks go to Daniel Thiele, Henning Sahlbach, Philip Axer, and Jonas Diemer for contributions to the slides

- References (selected)
  - for AVB timing analysis
    - Jonas Diemer, Daniel Thiele und Rolf Ernst, "Formal Worst-Case Timing Analysis of Ethernet Topologies with Strict-Priority and AVB Switching" *in Proc. 7th IEEE International Symposium on Industrial Embedded Systems, SIES 2012,* Karlsruhe, June 2012. Invited Paper.
    - Jonas Diemer, Jonas Rox, Rolf Ernst, Feng Chen, Karl-Theo Kremer, Kai Richter "Exploring the Worst-Case Timing of Ethernet AVB for Industrial Applications". *Proc. IEEE IECON 2012,* Montreal, Oct. 25 to 28, 2012.
  - for WRR analysis
    - Daniel Thiele, Jonas Diemer, Philip Axer, Rolf Ernst, Jan Seyler. "Improved Formal Worst-Case Timing Analysis of Weighted Round Robin Scheduling for Ethernet". *Appears: Proc. International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), Montreal, Oct. 2013.*

Related major projects

- CERTAINTY (FP7, start Nov. 1, 2011)
  " CErtification of Real Time Applications desIgNed for mixed criticaliTY "
  - formal composition and integration of networked mixed-critical systems
- ARAMIS (BmBF, start Dec. 2011)
  „ Automotive, Railway and Avionic Multicore Systems"
  - mixed criticality in the context of other non-functional requirements (security, …)