Very Low Latency Packet Delivery Requirements and Problem Statements

Yong Kim (ybkim at broadcom dot com)

Agenda

- The Requirements
- Use Cases that Drove Requirements
- The “Latency” Problem
- A Generic Solution for the Problem
- The IEEE 802.1/802.3 Functional Model
- Common Feedbacks to “Pre-emption” Proposal
- Summary and Conclusions

The Requirements

- Use Case: Low-latency real-time Control Loop use cases in Automotive and Industrial, converged onto rest of Ethernet network infrastructure. Payload size and bandwidth limited.

- Automotive -- 100 uS over five bridge hops @ 100 Mb/s and above –
  (from March & Sept 2011, 802.1:
  in public area: http://www.ieee802.org/1/files/public/docs2011)
  - new-avb-KimNakamura-automotive-network-requirements-0311.pdf
  - new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf (revised)
- Industrial -- <5 uS per hop, ~32 bridge hops @ 1000 Mb/s and above –
  (from January 2011, 802.1). 125 uS over 32 hops desired.
  - new-goetz-avb-ext-industrcom-0113-v01.pdf
  - ba-goetz-industrial-profile-0509.pdf ← /docs2009

- Problem Statement:
  (Ignoring the bridge and other delay for the moment)
  - Max Length Ethernet Frame @ 100 Mb/s =~120 uS – greater than automotive requirements.
  - Max Length Ethernet Frame @ 1000 Mb/s =~12 uS – greater than industrial requirements.
  - "Head of Line" blocked behind Max Length Frame exceeds the requirements above.

Use Case – Automotive 1

from new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf

Example next-generation automotive network architecture

Ethernet AVB applied to automotive control data transmission between domain gateways and in powertrain/chassis domains

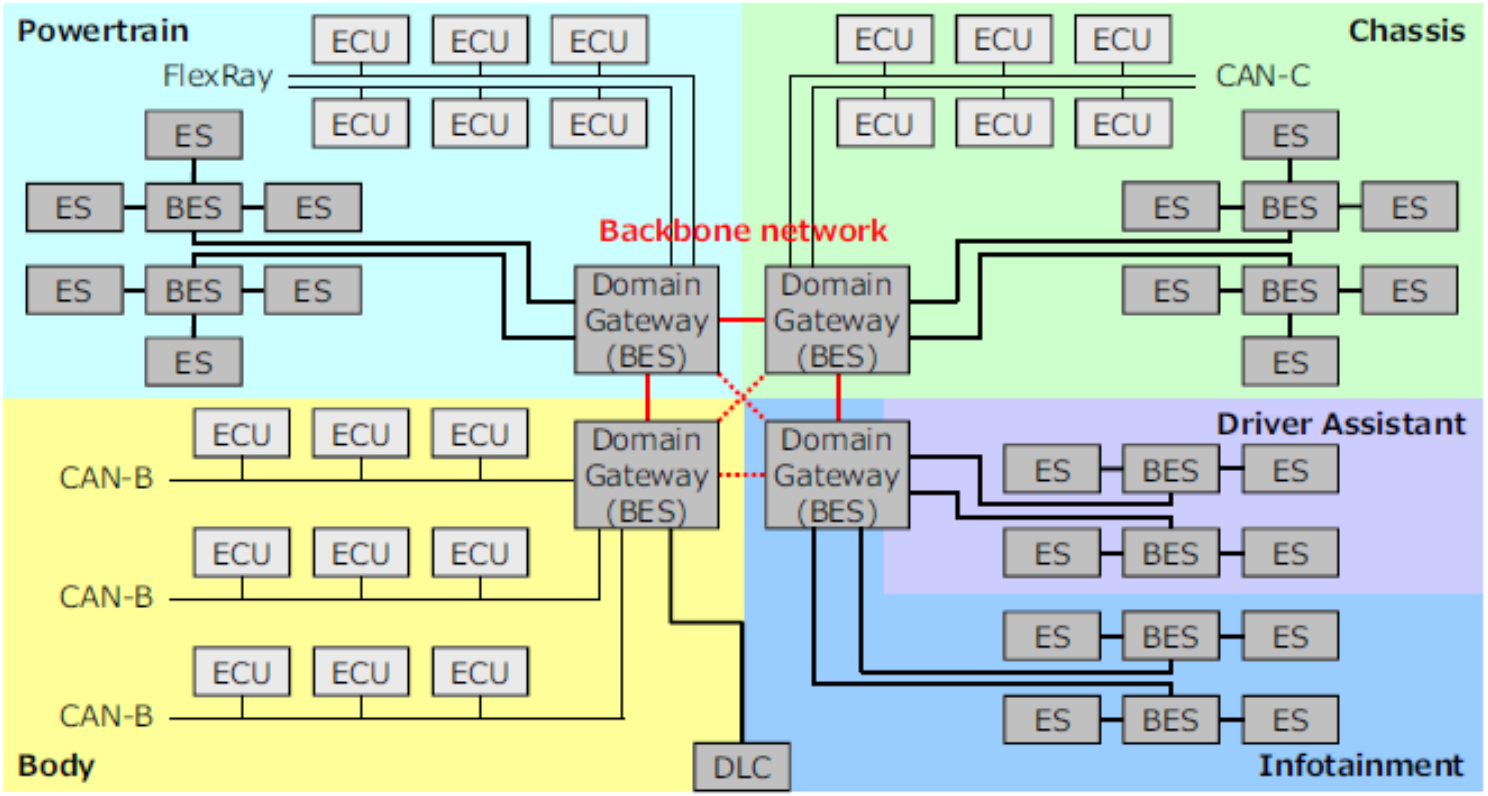

Spatial

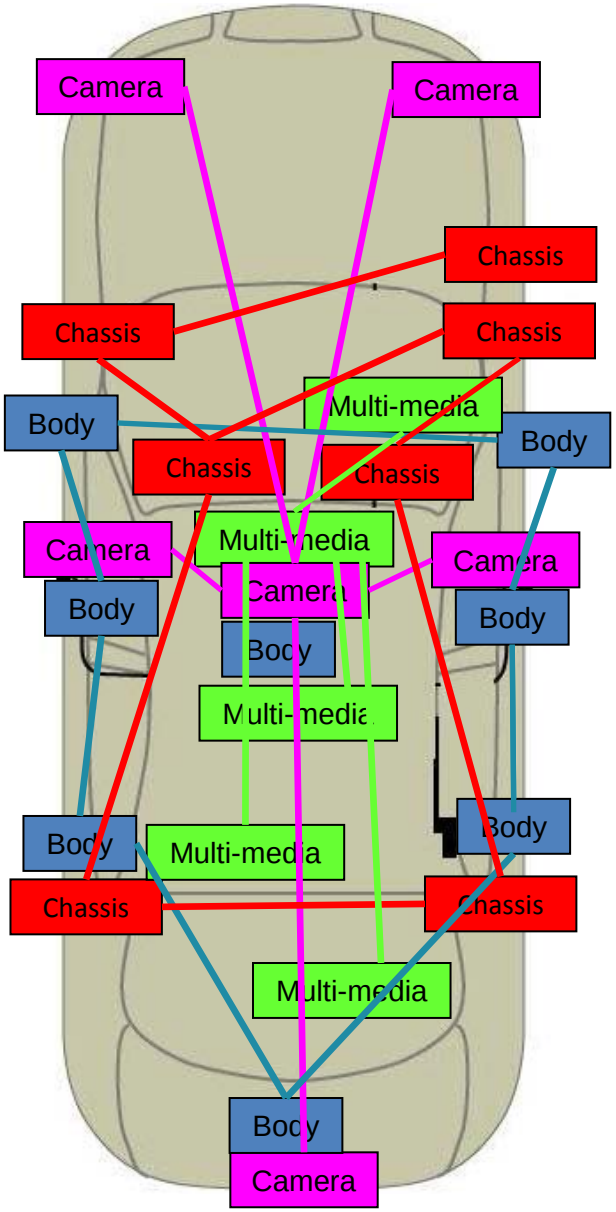

40~60 uC per mid-range car

Each dedicated domain Buses over an entire car

Ethernet Backbone "Network"

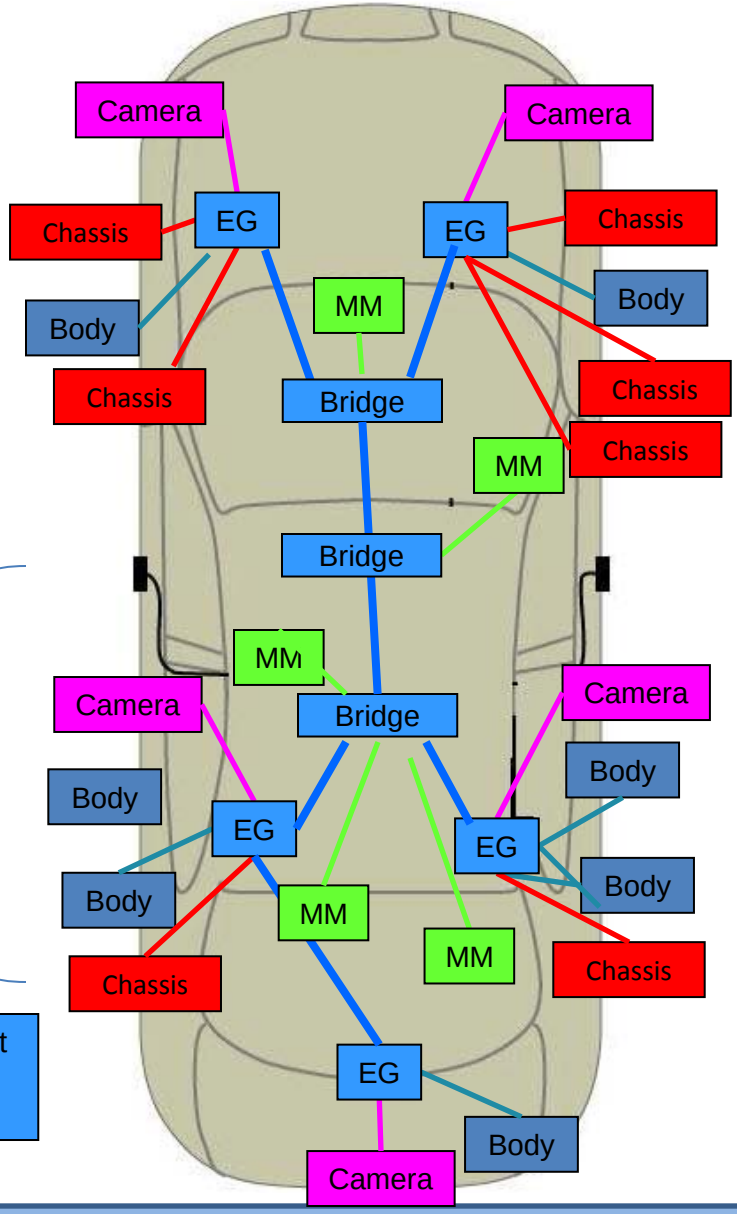

MM: Multi-media

EG: Ethernet [Bus] Gateway

Use Case – Automotive 2

from new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf

Consideration of AVB hop counts

Assumption for bridged end stations

- A bridged end station consists of a bridge and a end station (talker/listener).
- One AVB hop needs to be counted inside of bridged end stations.

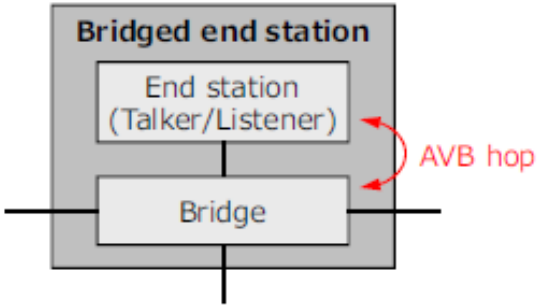

Example control data transmission in a backbone network

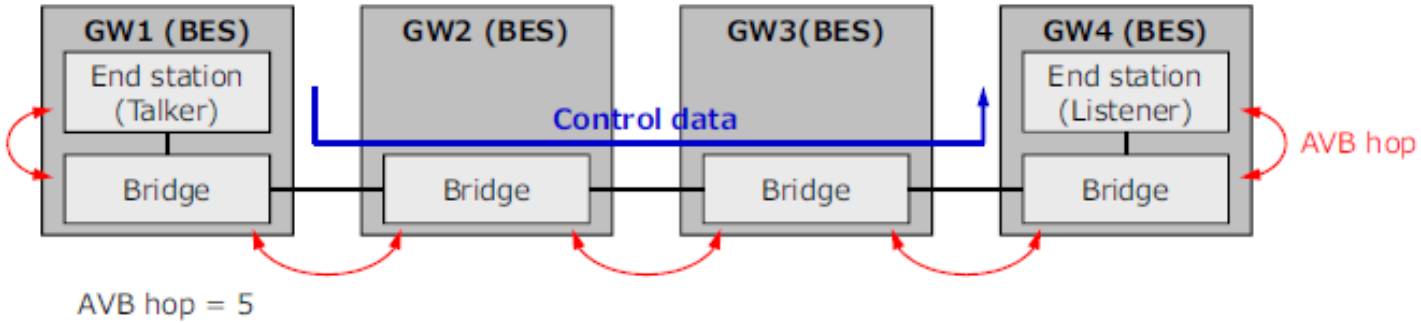

Use Case – Automotive Requirements

from new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf

QoS requirements for automotive control data class

Performance requirements for automotive control data class

- Maximum latency: 100 us / 5 AVB hops
  - Guaranteed latency
  - Topology independent
  - Automotive control data class to have higher priority than SR classes
    - Maximum 2 priority classes (e.g. Control data class and SR class A)

Preconditions for performance requirements

- Network type: Dedicated network in a vehicle
- Network attributes
  - Maximum AVB hop count: 7
  - Maximum number of nodes (bridged end station & end stations): 32
  - Maximum cable length: 24 m
  - Maximum end-to-end cable length: 30 m
- Automotive control data class attributes
  - Maximum data size (payload size): 128 bytes @FE ~ 256 bytes @GE
  - Maximum number of simultaneous transmission: 8 @FE ~ 32 @GE
  - Transmission period: 500 us
- Payload size for other/lower traffic classes: 256 bytes @FE ~ 1500 bytes @GE

These are our best estimates derived from multiple assumptions of the current and future automotive applications.

Industrial Use Case 1 - Applications

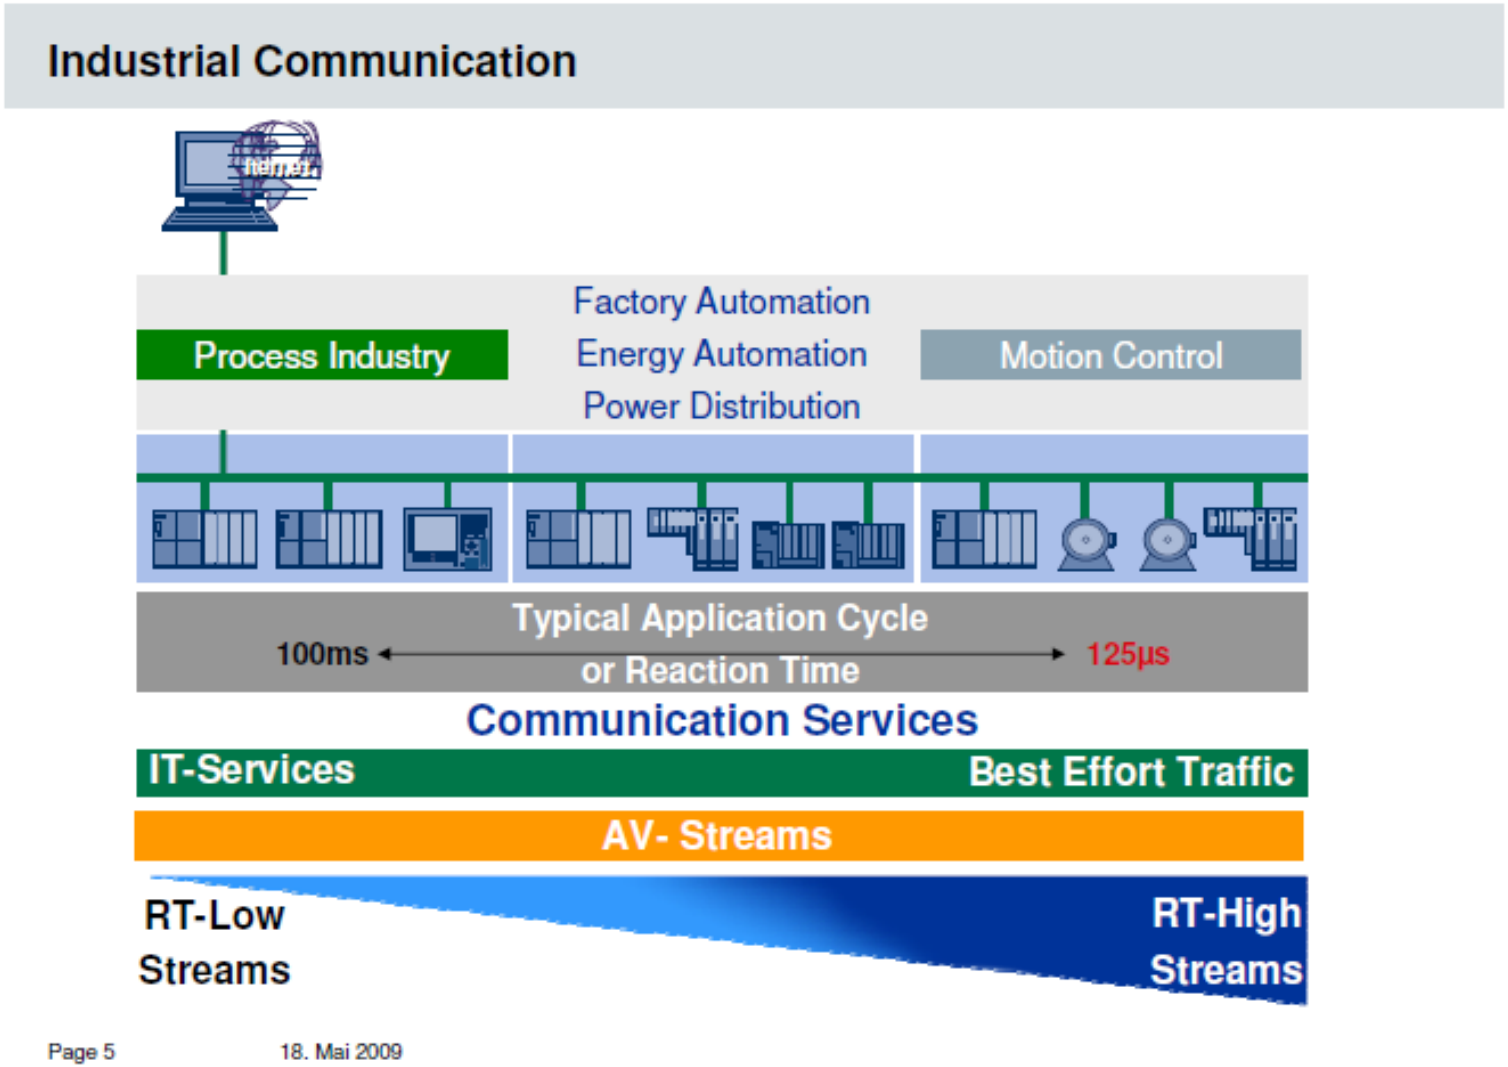

Industrial Use Case 2 - Topology

Typical Topology for Bridged LANs in Industry

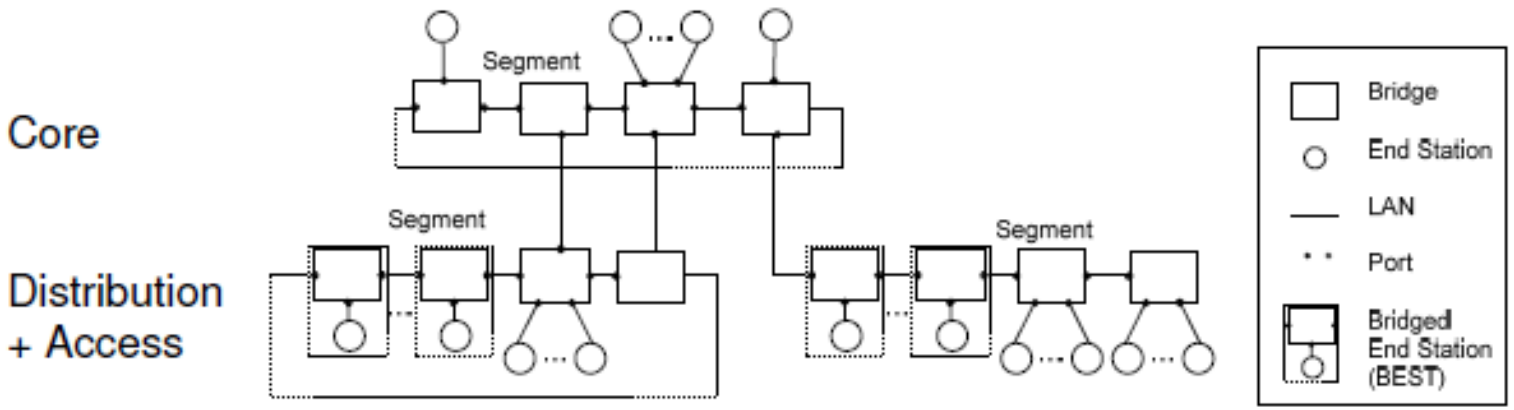

Linear Topology (Line, Ring)

- Applications mostly feature distributed linear dimension, i.e. production lines
- The network is tailored to the application, i.e. line topologies
- To reduce total cost, 2-port-bridges are integrated into the end stations
- For enhanced availability, lines are closed to rings
- Homogeneous segments (100MBit or GBit)

Industrial Use Case 3 - Requirements

Typical Topology for Bridged LANs in Industry

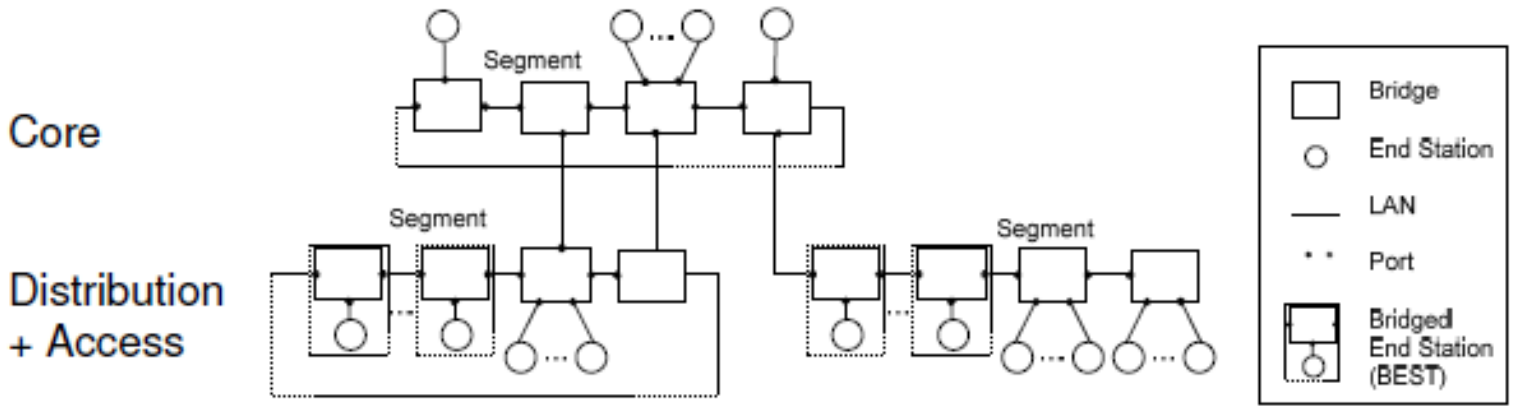

Linear Topology (Line, Ring)

- Applications mostly feature distributed linear dimension, i.e. production lines
- The network is tailored to the application, i.e. line topologies
- To reduce total cost, 2-port-bridges are integrated into the end stations
- For enhanced availability, lines are closed to rings
- Homogeneous segments (100MBit or GBit)

The "Latency" Problem

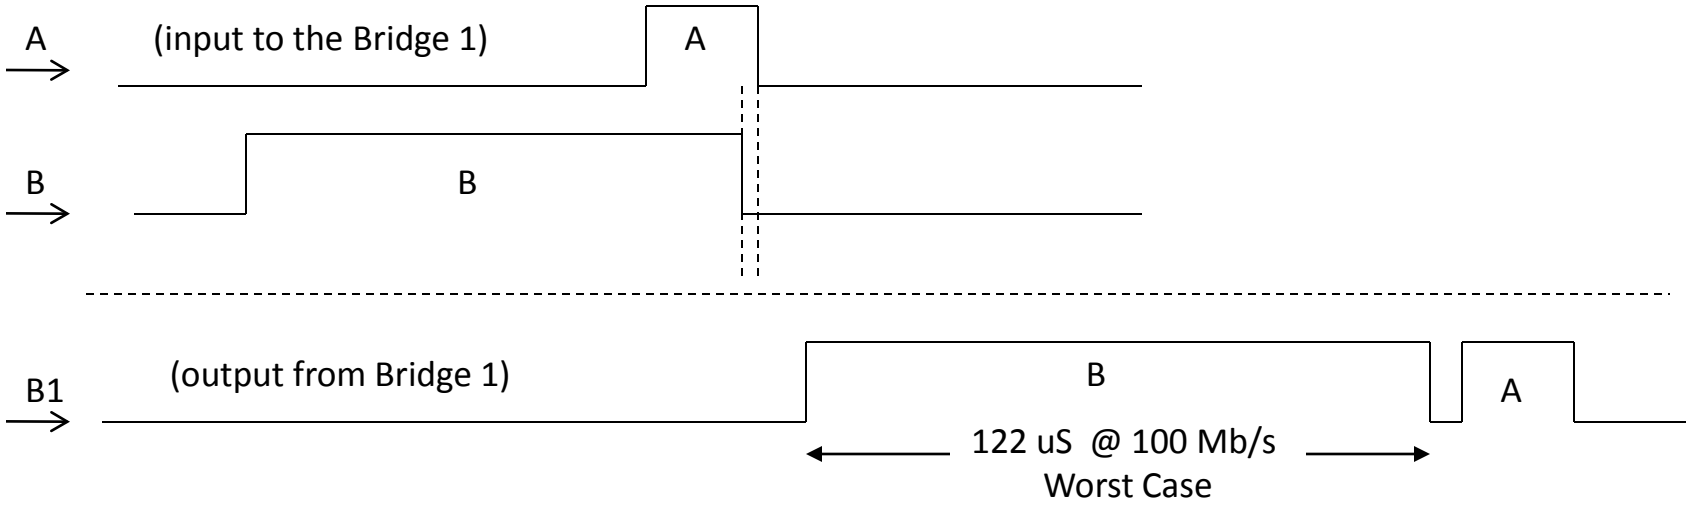

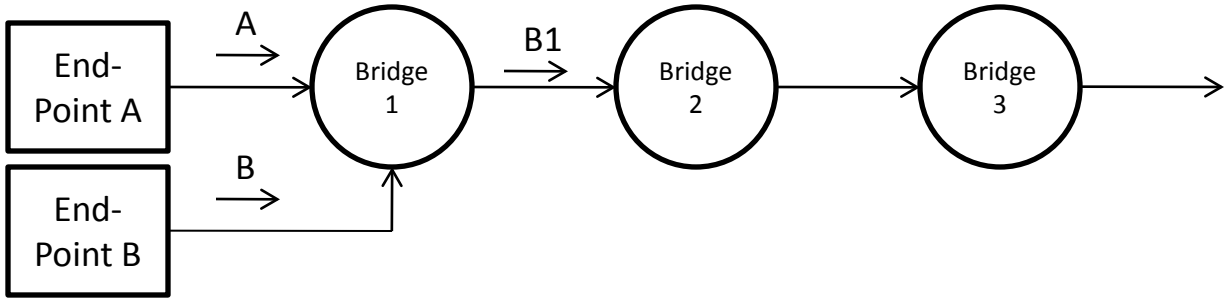

- The Classic "Head of Line" problem. The worst case is urgent frame A is scheduled behind best-effort maximum length frame B. But we desire urgent frame to get through (e.g. 5 hops @ 100 Mb/s in 100 uS).
- One maximum sized frame @ 100 Mb/s ahead of low-latency control frame exceeds desired latency over 5 bridge hops.

A Generic Preemption Solution

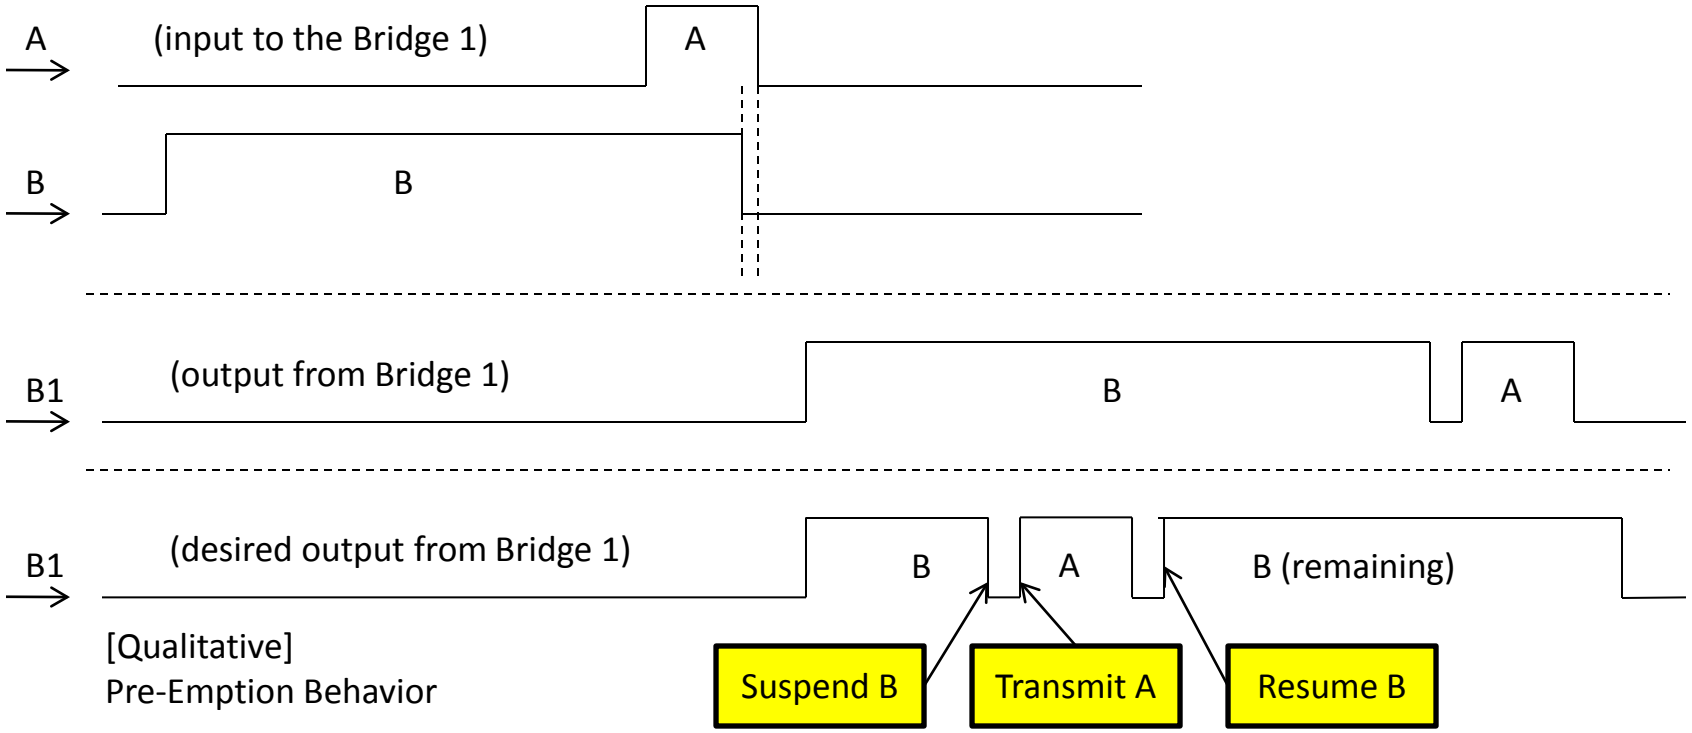

- There are few other solution to “head of line blocking” delay behind a max frame length packet.
- A generic solution is to suspend the max frame length (“B” in this example) packet, transmit urgent frame, and resume previous. Note: Other completion options besides resume-previous is retransmit B entirely, or always fragment B (regardless of existence of A) are not desirable).

Summary of "Preemption Function"

- Easy to explain – easy to understand.
- Fairly straight forward to implement, pending selection of many [functionally] equivalent proposals.
- Observation – "Everyone knows what 'preemption' means, but everyone has different ideas on how it may be implemented". Let's hold off on this.

- But the MAC Service interface (boundary between 802.1 and 802.3) as defined (particularly .3 side) is not friendly to "Pre-emption" considerations.
  - [Next Slide] Entire 802.1 Frame is sent atomically to 802.3 MAC, as an example. No way to offer preempting frame to 802.3 from 802.1

Existing 802.1 and 802.3 Service Interface

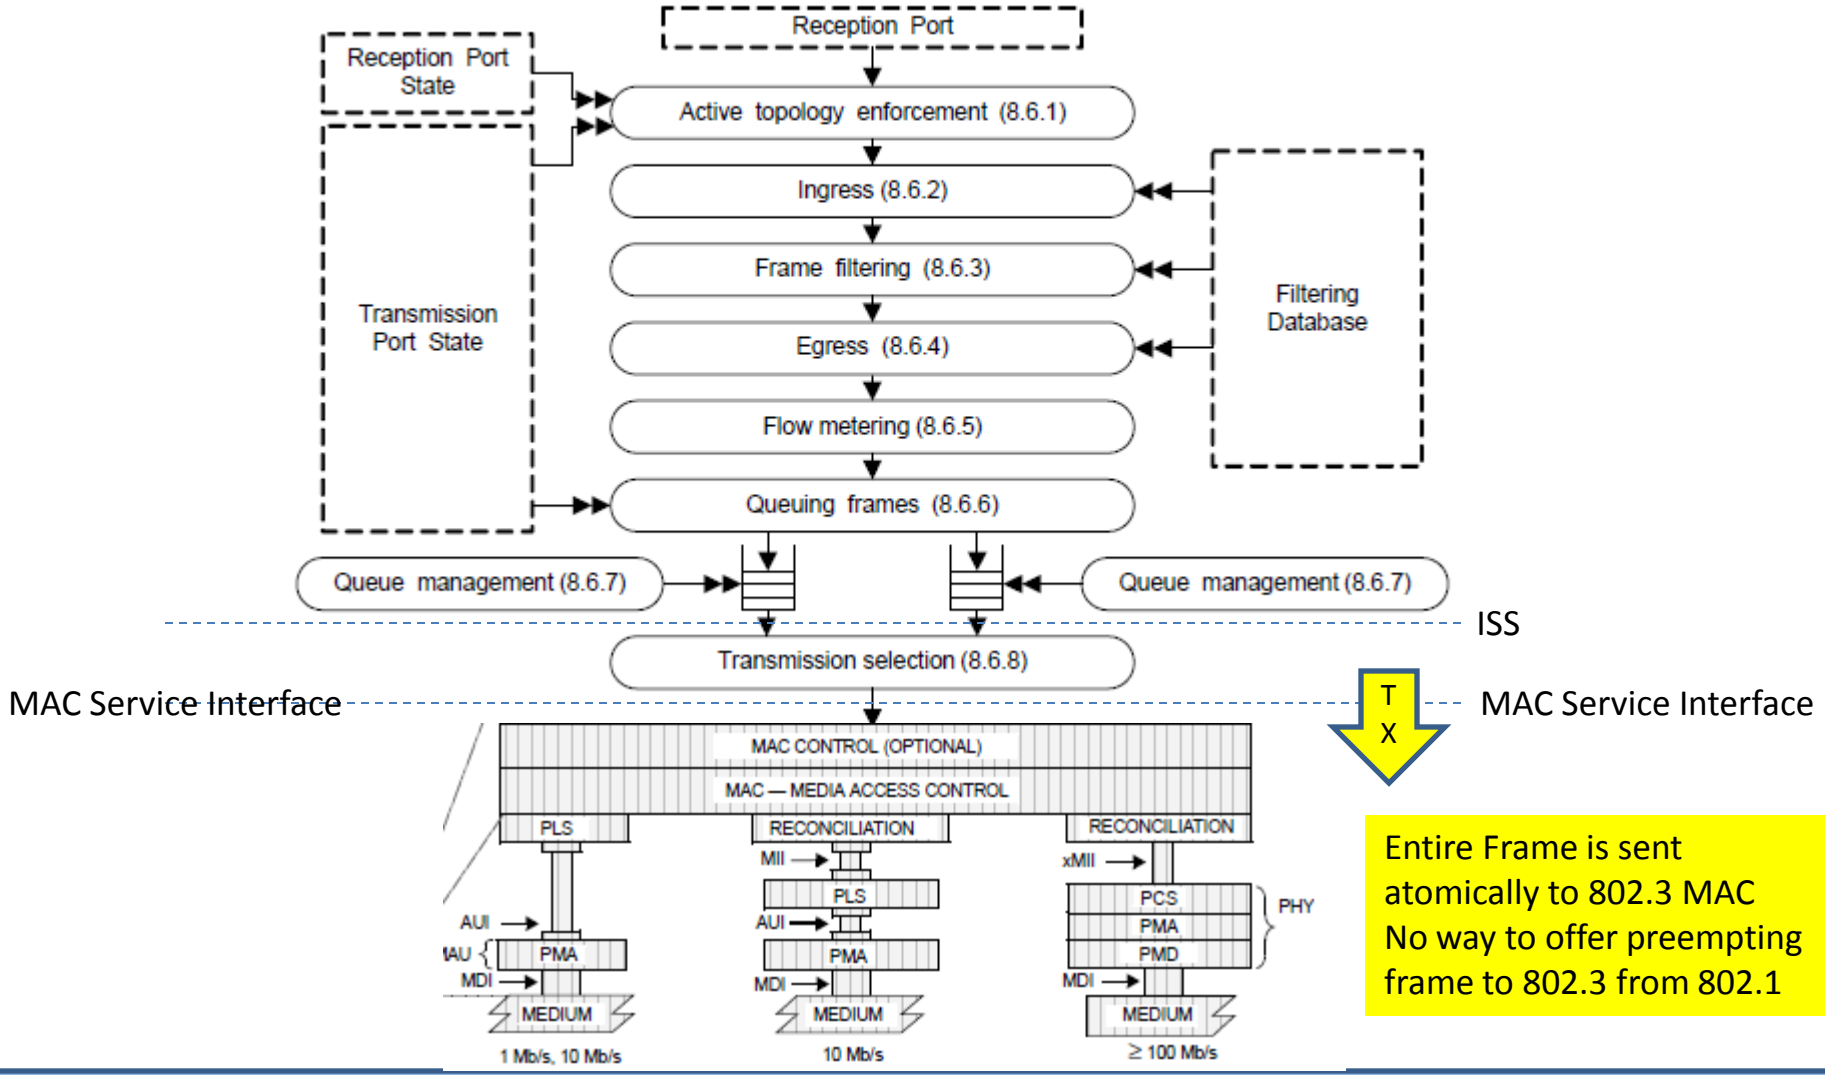

Entire Frame is sent atomically to 802.3 MAC
No way to offer preempting frame to 802.3 from 802.1

IEEE 802.3 Clause 4 or Annex 4A

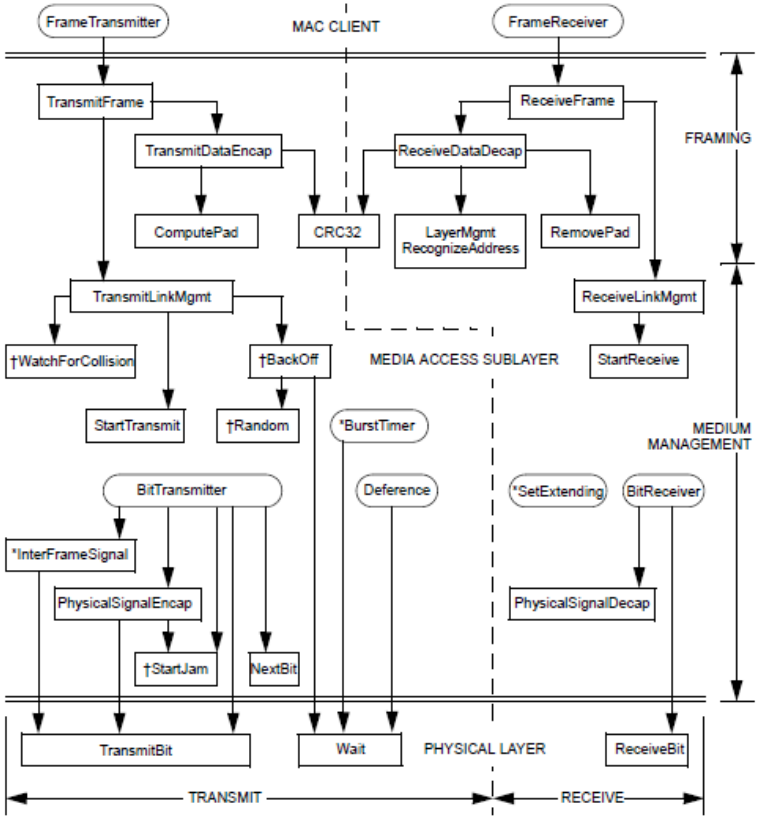

† Not applicable to full duplex operation.
* Applicable only to half duplex operation at 1000 Mb/s.

Figure 4–1—Relationship among CSMA/CD procedures

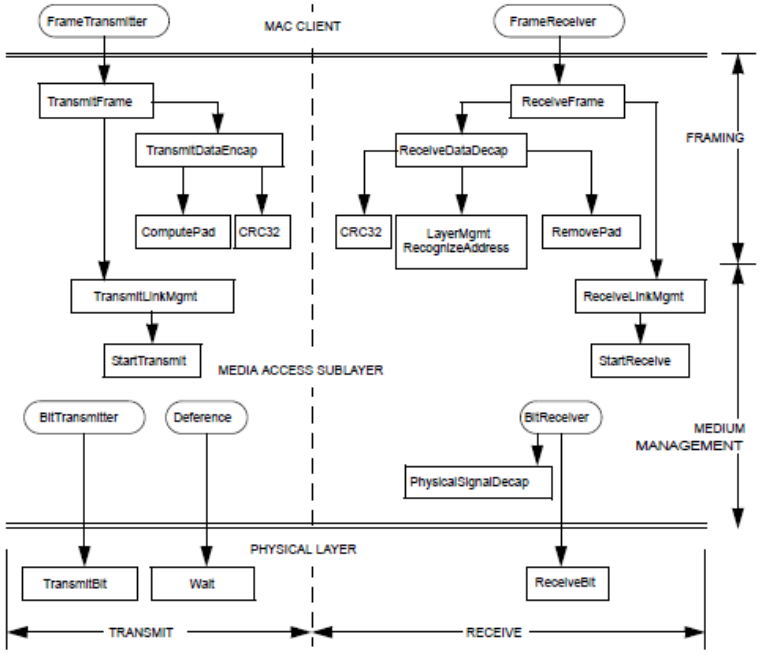

Figure 4A–1—Relationship among MAC procedures

Annex 4A (Full-Duplex only) may be the best clause to consider for preemption function.

Leave Claus 4 alone.

Proposed 802.1 & 802.3 Service Interface

- Problem – A packet is sent from 802.1 to 802.3 in zero time. 802.3 transmits the frame.
- Generic Solution –
  1. 802.1/802.3 MAC Service Interface needs to be augmented to convey a second preempting frame.
     - Additional information may be in form of “preempt-indication”, or second “transmit-urgent-request” or any other TBD during standardization.
  2. 802.1 Services determine frames that are preemptable, preempting, and not preemptable.
     - Effect of AVB shaper, transmit selection, congestion management (if relevant) TBD during standardization.
     - 802.1 could complete its preemption so that it works with any MACs that supports preemption capability.
  3. 802.3 MAC to consider preemption related enhancements and handle preemption point and preemption framing.
     - Depends on the latency objectives and fragmentation header formats. TBD during standardization.

Common Feedbacks on Pre-emption Proposal

Common Feedback 1a

- Why don't Automotive just use higher speed?
- Answer: EMC is the bottleneck (from Pg 10, BMW presentation @)
  - http://www.freescale.com/files/ftf_2010/Americas/WBNR_FTF10_AUT_F0558.pdf

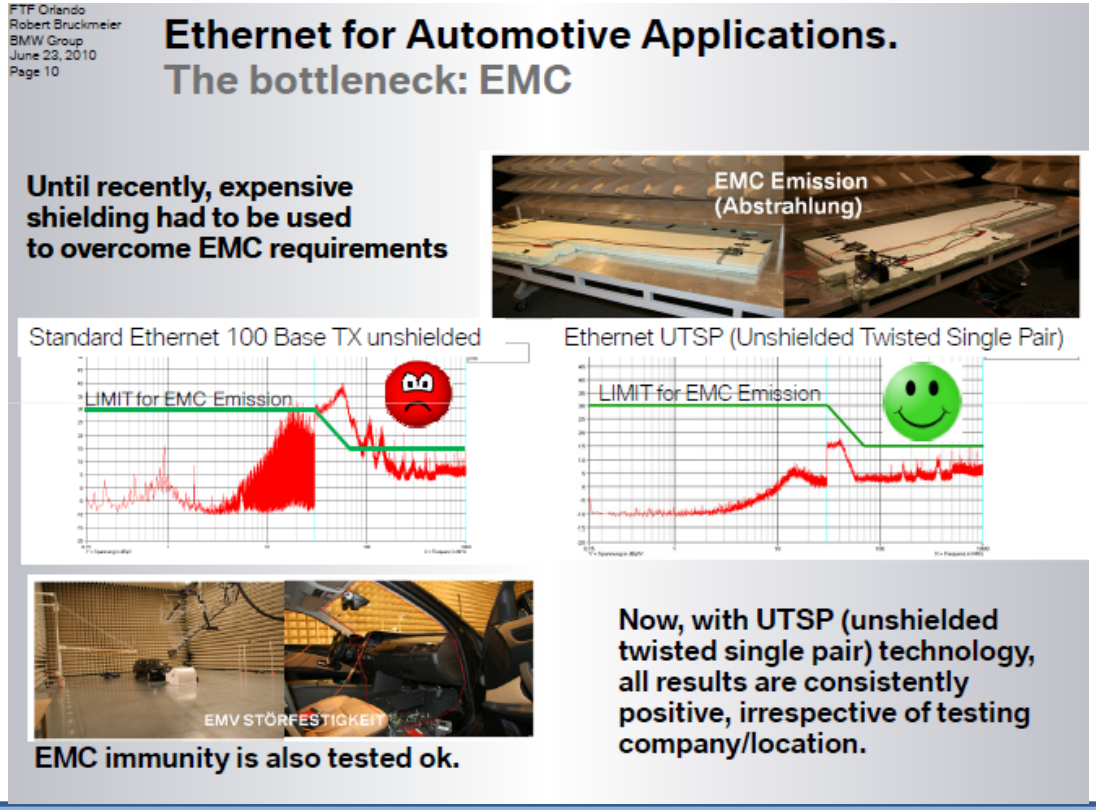

Common Feedback 1b

- Why don't Industrial just use higher speed?
- Answer: It does (adopts Gigabit Speeds), and the dominating latency comes from the electronics (bridges, PHYs) when cascading. There is limit on how much latency you could improve.

Common Feedback 2

Why can't you just use shorter max size frames?

- Answer: It could and it is an interim solution.

Why Interim solution?

- Answer: Limits converged infrastructure of data, streaming, control unless every device in the network needs changes to adopt new max size
- Answer: Payload efficiency goes down significantly.
  100 Mbps link with 128 byte max frame versus 1500 byte would be 41% versus 95% of link SW.
  Any additional header overhead make this efficiency worse.

Packet Length	64	128	256	1536
IPG	12	12	12	12
Preample	8	8	8	8
MAC Header	16	16	16	16
L2 Payload	28	92	220	1500
L2 Payload/Pkt Size	**44%**	**72%**	**86%**	**98%**
IPv4 HDR	20	20	20	20
L3 Payload	8	72	200	1480
L3 Payload/Pkt Size	**13%**	**56%**	**78%**	**96%**
TCP HDR	20	20	20	20
L4 Payload	-12	52	180	1460
L4 Payload/Pkt Size	**NM**	**41%**	**70%**	**95%**

Common Feedback 3

Why can't you just use star-wired network?

- Answer: Industrial – it does when it could. Legacy cabling is combination of star, ring, overlapping rings, and daisy-chains. One of the attractiveness of Ethernet is that it could supports existing wiring. Rings offer attractive low-cost and complexity redundant paths.
- Answer: Automotive – One of the attractiveness of Ethernet is reduction of wiring (400~600 lb of cabling in a car). Combination of star and daisy chain topology allow for ease of wire routing and managing optional package connections.

Summary and Conclusions

- Both Industrial and Automotive control system require low-latency support beyond AVB.
- Several projects in 802.1 to enable Bridge use in both system platform.
- Pre-emption is deemed to be most optimal method for switched (versus bused) network.
- Start a parallel project in 802.1 and 802.3 to serve these markets.
- Draft PAR for 802.1 project is available.

Thank you!

Two backup slides

Recent 802.1 bridge models

Relevant 802.1 documents.

Provider Backbone Baggie Pants Model

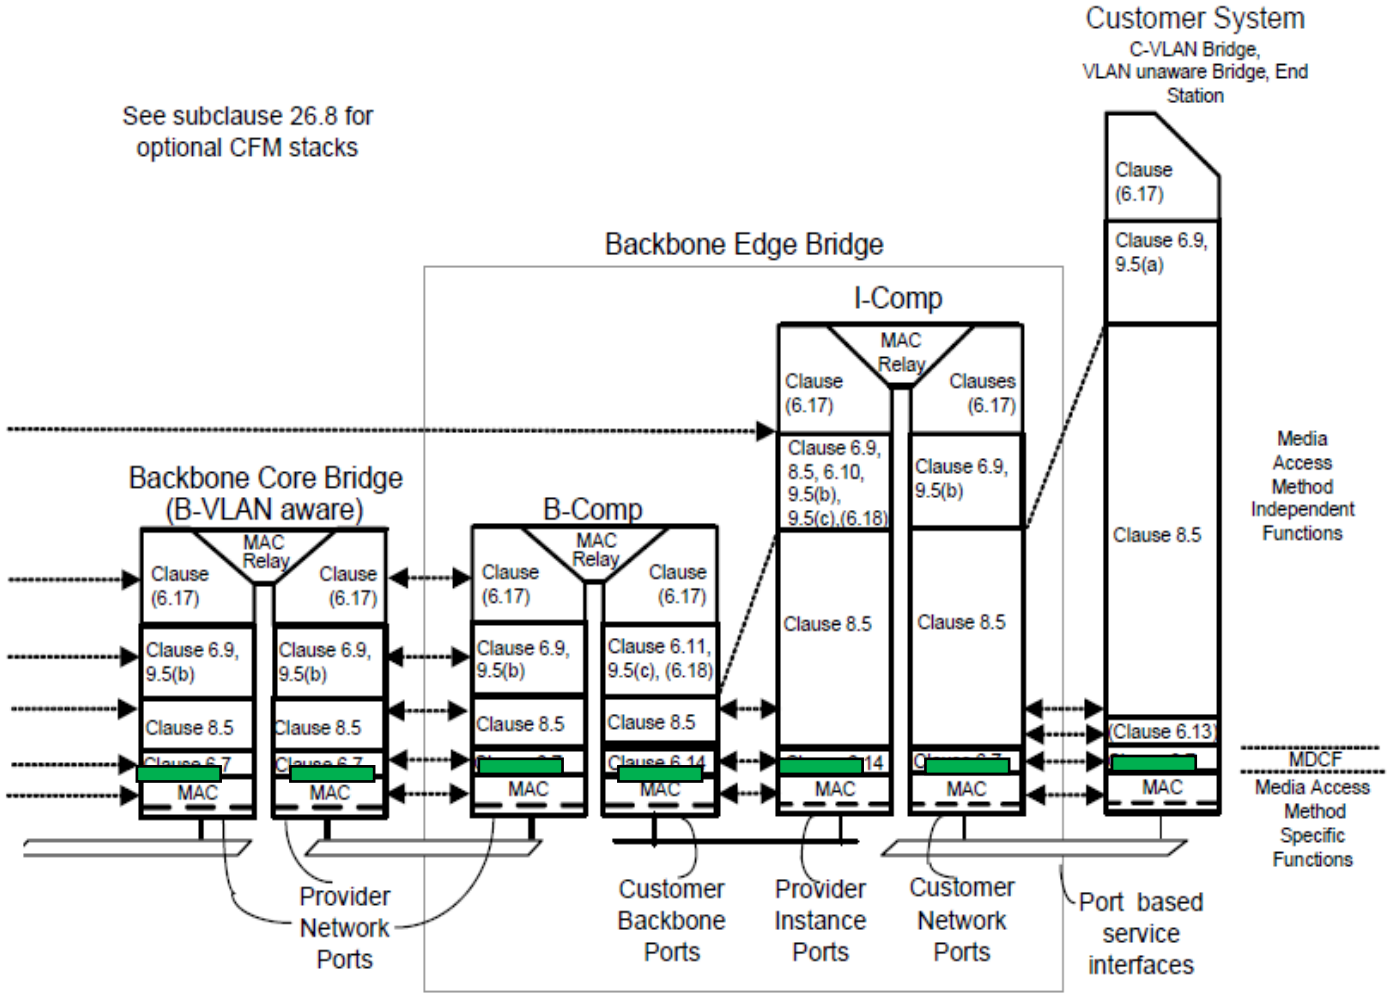

Figure 25-4—Port-based service interface

Suggested preemption Q-Rev work

Relevant 802.1 documents

All located in: http://www.ieee802.org/1/files/public/docs2011, except noted.

- Automotive
  - new-avb-KimNakamura-automotive-network-requirements-0311.pdf
  - new-avb-nakamura-automotive-backbone-requirements-0907-v02.pdf
- Industrial
  - ba-goetz-industrial-profile-0509.pdf ← /docs2009
  - new-goetz-avb-ext-industrcom-0113-v01.pdf
  - at-klein-kleineberg-avb-redundancy-continuation-0711.pdf
- Preemption Related
  - new-avb-kim-8021-8023-Preemption-Problem-Statements-0711-v03.pdf
  - new-avb-kim-8021-Preemption-DRAFT-PAR-5C-0711-v04.pdf
  - new-imtiaz-goetz-fragmentation-0511.pdf
  - new-kim+goetz-Ultra-Low-Latency-Switching-v5.pdf
  - avb2-cgunther-sr-class-preemption-0711-v05.pdf
  - new-avb-boiger-ultra-low-latency-shaper-0711.pdf

Automotive Wiring (Partial)