IEEE 802.3
10BASE-T1S Multidrop Latency Enhancements

Call For Interest Consensus Presentation

Markus Jochim
Yong Kim

DRAFT version: April 16, 2026

CFI Objectives

- To gauge the interest in starting a Study Group for

  Differentiated Services Enhancements on Improving 10BASE-T1S Multidrop (10SPE) Access Latency

- This meeting will **NOT**
  - Fully explore the problem
  - Debate possible solutions
  - Choose a solution
  - Create a project documentation (PAR, CSD, or Objectives)
  - Create a standard or specification.

- Anyone in the room may vote or speak

10SPM Access Latency Improvement CFI Motivation

- 10SPM is widely adopted and ready to be deployed. However, the latency is too long to serve controls applications that requires bounded and shorter latencies.
- Target applications include
  - Controls applications in automotive, robotics, industrials, etc
  - Factory and Building automations, etc.
- Topics to address
  - Deterministic and bounded latency
  - Deterministic and bounded latency for a given traffic class

Draft Agenda

- Overview
- Presentation Content
  - What is 10BASE-T1S PLCA (“single pair multi-drop”)
  - The Latency Bounds and its Problem for System Builders
  - Technical Feasibility
  - Market Relevance and Potential
  - Why Now?
- Q&A
  - Panel: ?
- Straw Polls

Overview, Problem Statement, and the Needs

Why Are We Here?

- 802.3cg includes PLCA (Physical Layer Collision Avoidance) operation using 10BASE-T1S in shared medium. And it is deployed on production vehicles.
- PLCA provides fair and equal access to all nodes.
- Its Maximum Latency is bounded, by number of nodes, max packet sizes, number of packets in a burst. But….
- It’s too long (for control loop applications) and not deterministic because of it (variability). Even for a small number of nodes in a segment.

But first, what is PLCA?

noting only the relevancies to this presentation…

- It's entirely in Reconciliation Sub-layer (RS)
- It works with CSMA/CD MAC (Clause 4)
- Signals bus cycle start ('beacon') and
  then all nodes send its packets in round-robin.
  by every node by ID. Packet burst by configuration.
- When all nodes are operating in PLCA, max latency is bounded.
  - Theoretical: 255 nodes, 1536 octets, 256 packet burst = 80.2 sec.
  - Practical: 8 nodes, 1536 octets, 1 packet burst = 9.83 msec.
    Practical: 8 nodes, 256 octets, 1 packet burst = 1.76 msec
  - Single packet transit latency (1536 octets): 1.22 msec….

  And system designer then start to play with max burst for the server-side port, max MTU less than 1536 octets, number of nodes on the segment, etc, to achieve the acceptable latency.

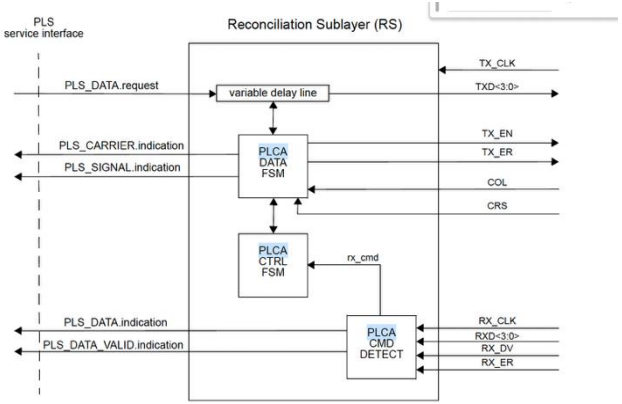
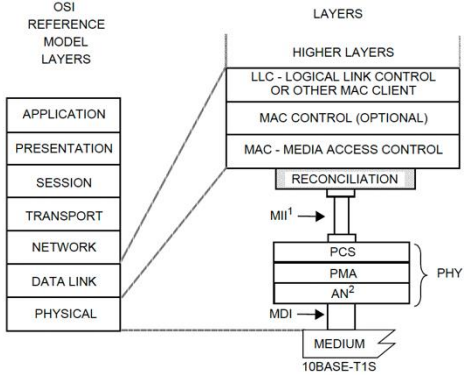


Figure 148–2—PLCA functions within the Reconciliation Sublayer (RS)



MDI = MEDIUM DEPENDENT INTERFACE
MII = MEDIA INDEPENDENT INTERFACE

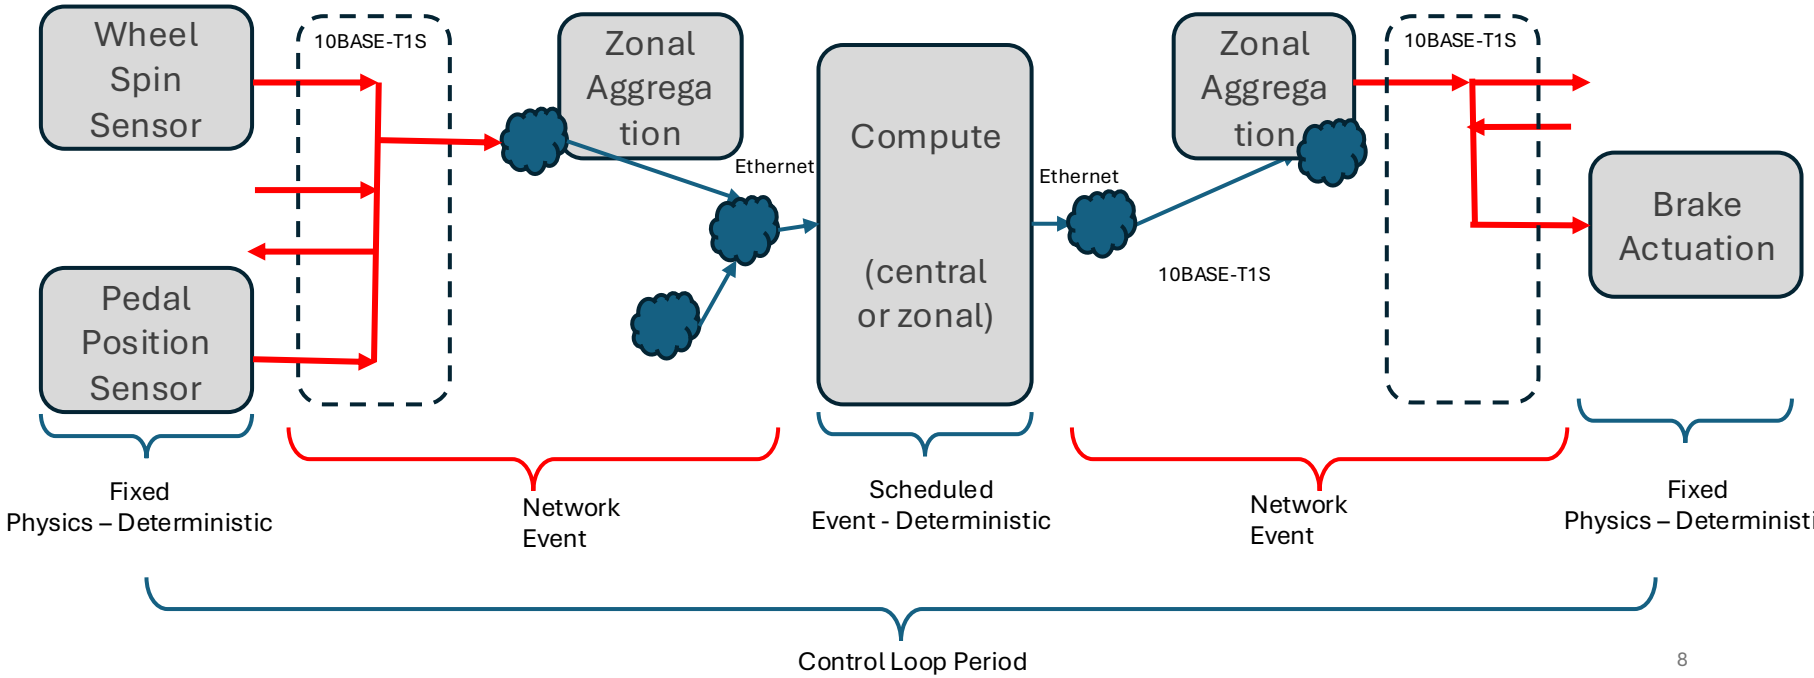
NOTE 1—MII is optional
NOTE 2—Auto-Negotiation is optional

PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PHY = PHYSICAL LAYER DEVICE
AN = AUTO-NEGOTIATION

Figure 148–1—Relationship of PLCA Reconciliation Sublayer to the ISO/IEC OSI reference model and the IEEE 802.3 Ethernet model

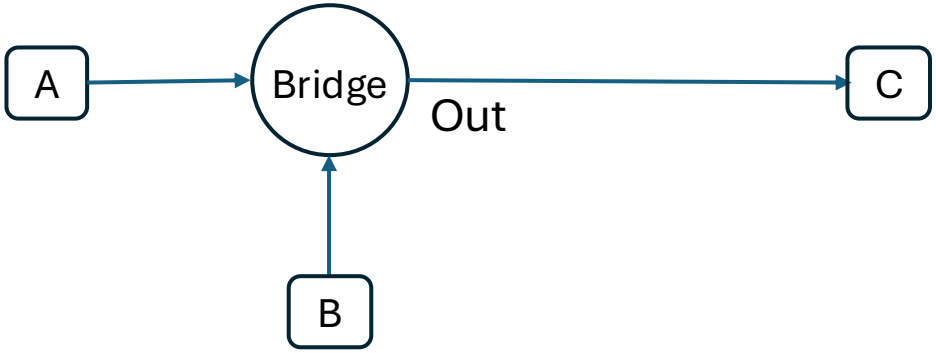
Control Loops (or Controls Applications)

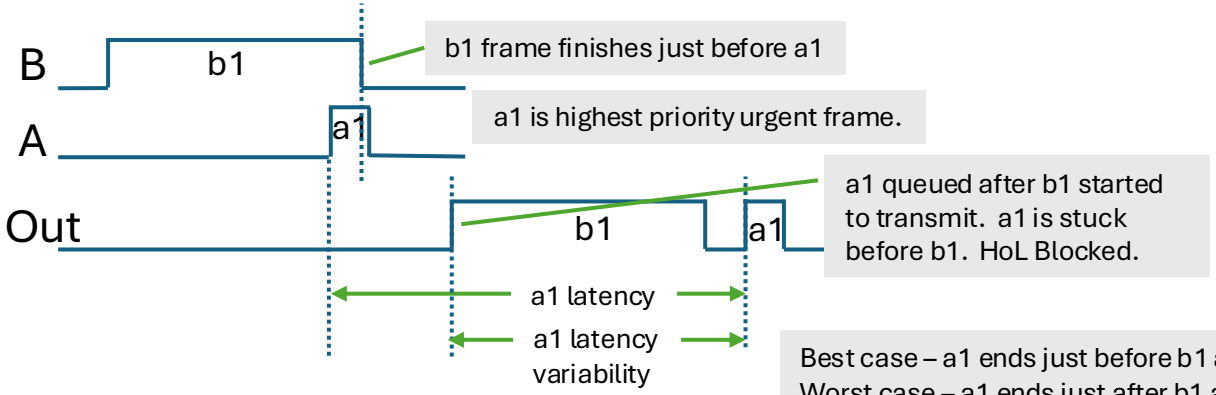
- [Sense] → [Compute] → [Actuate] control loop.
- The worst case end-to-end latency often determine the robustness and feature of the controls applications.



Let's illustrate through HoLB latency Analogy

Showing a single 802.1Q Bridge hop, typical of point-to-point PHYs





b1 frame finishes just before a1

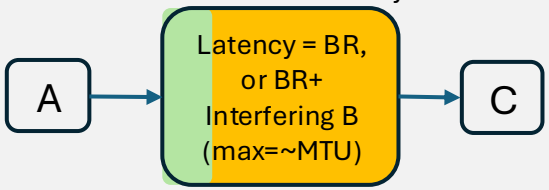
a1 is highest priority urgent frame.

a1 queued after b1 started to transmit. a1 is stuck before b1. HoL Blocked.

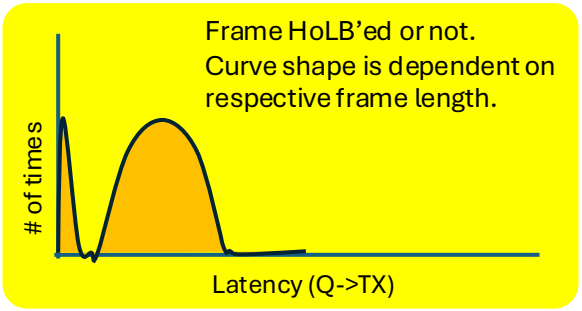
Best case – a1 ends just before b1 and queued.
Worst case – a1 ends just after b1 and HoL Bloked

HoLB = Head of Line Blocking

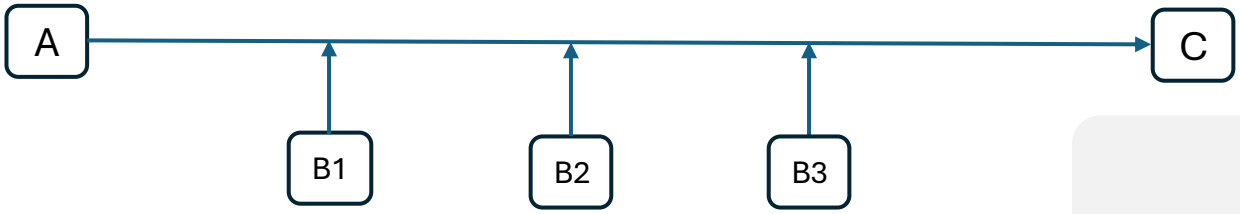
A→C network latency

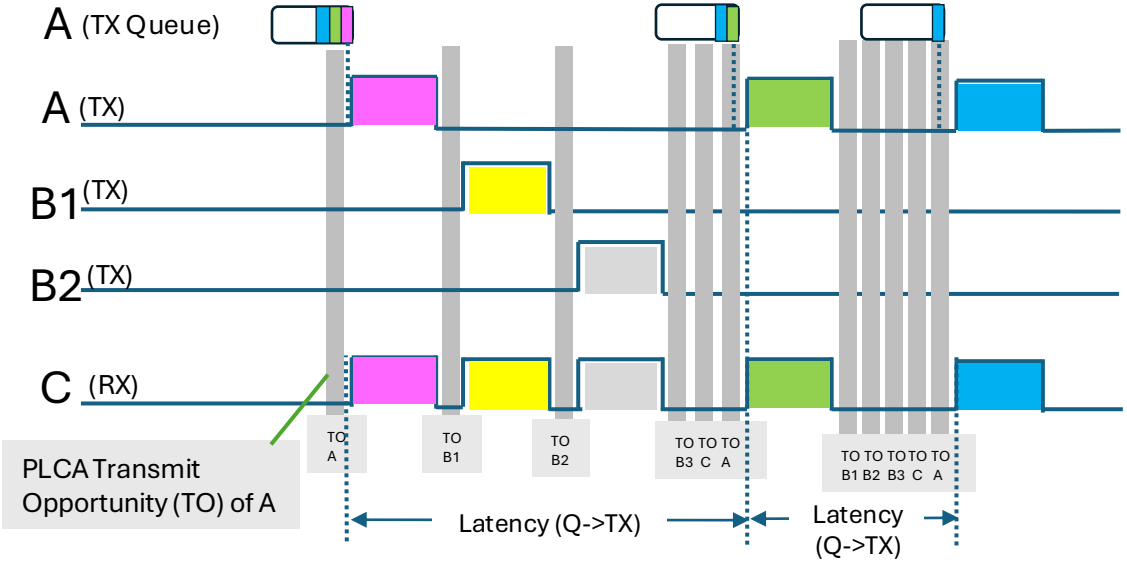


Either the minimum bridge transit delay, or additional max MTU interferer frame length delay.

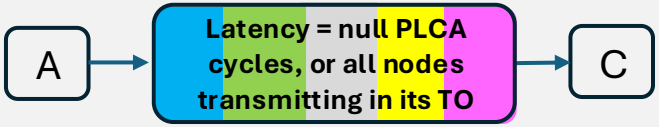


PLCA Segment latency, A→C

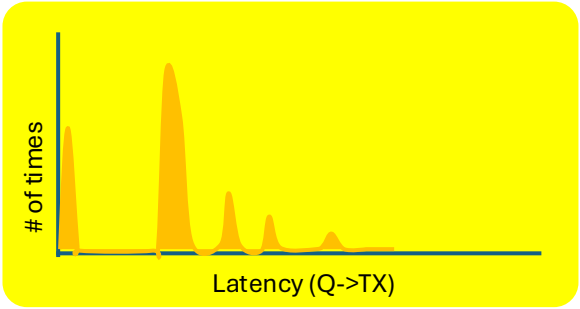




A→C network latency

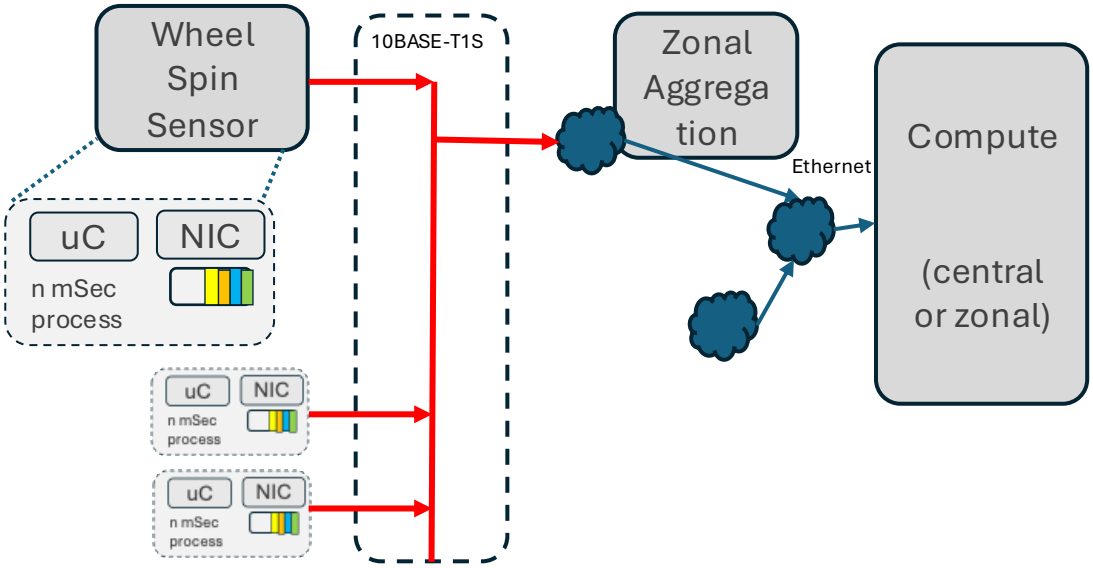


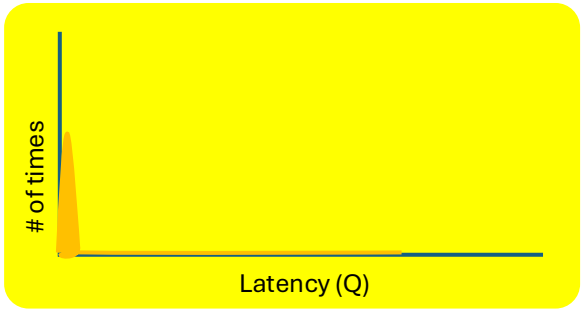
Either the minimum PLCA cycle delay, or additional frame transmissions from all the other nodes.

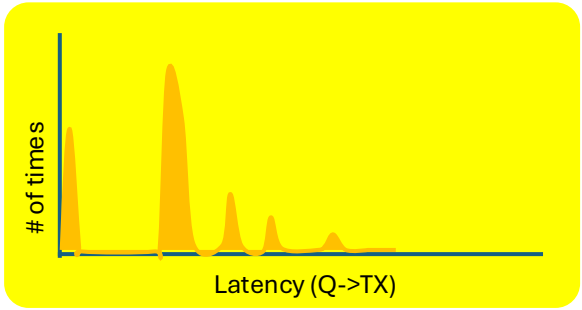


Engineered around 10BASE-T1S Latencies

- End-node operates on a controls application cycle.
- Interactions with other 10BASE-T1S nodes in PLCA cycles discrete blocking delays.

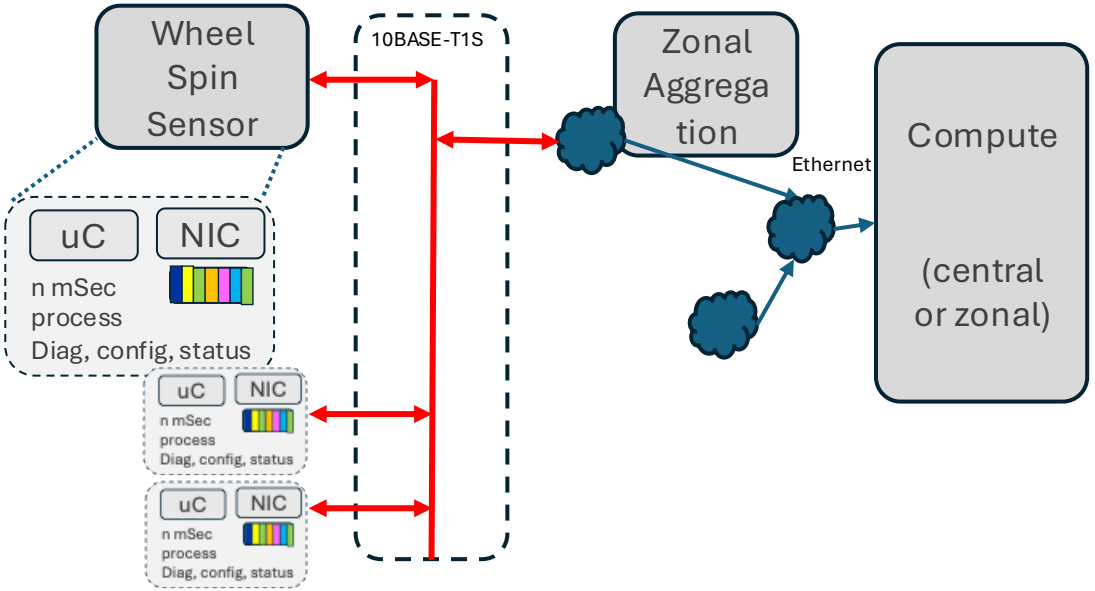


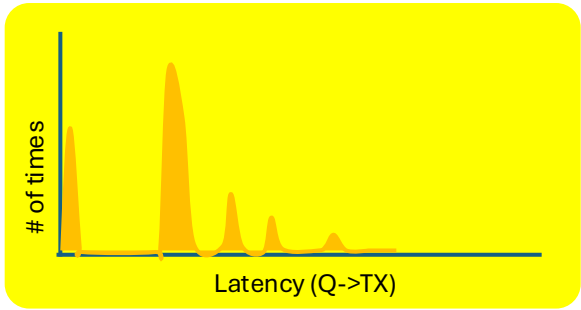




What about other workload traffic?

- End-node operates on a controls application cycle.
- Interactions with other 10BASE-T1S nodes in PLCA cycles discrete blocking delays.
- What happens with Diag, Config, Status messages in the same network? Interactions with additional controls applications traffic within the same PLCA cycles with the same end-nodes.

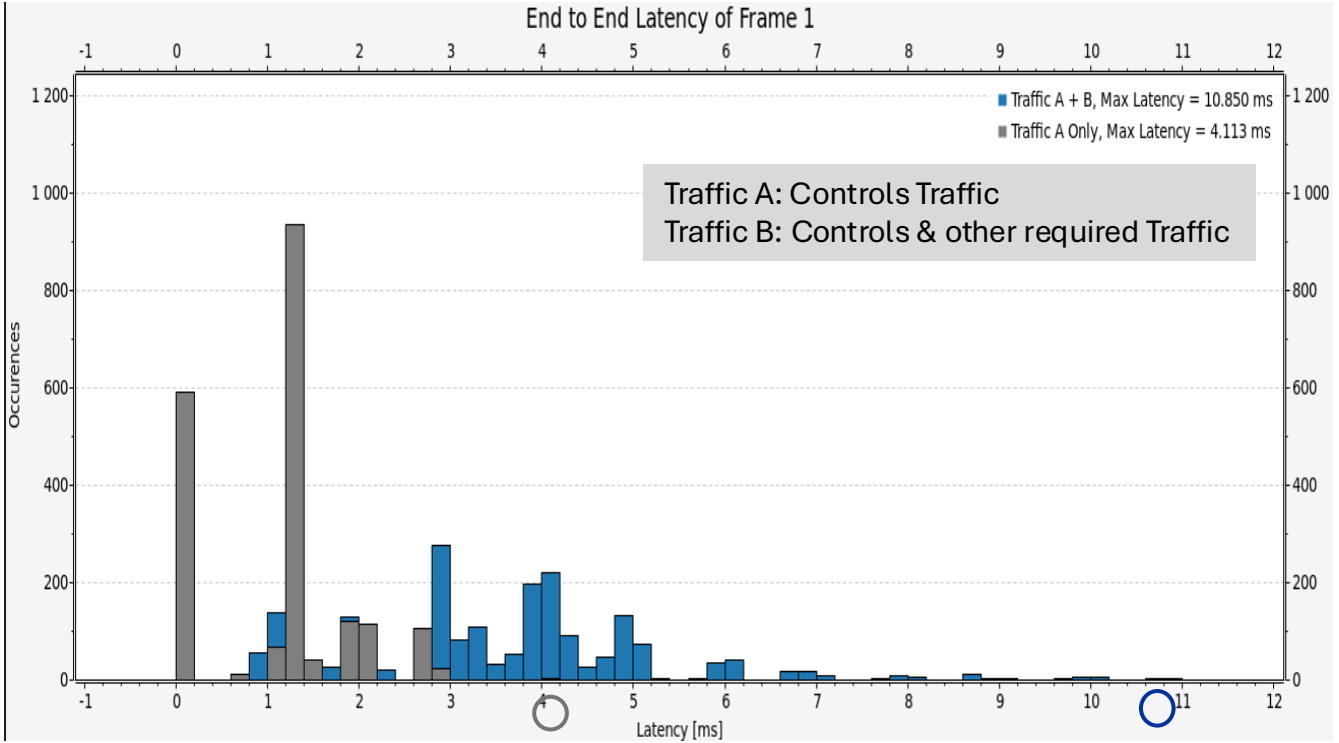




?

Simulation from 10BASE-T1S network

- Discrete latency distribution caused by variable length PLCA cycles.
- Challenging to engineer network for to meet latency requirements for controls traffic A.
- After adding traffic B: Near impossible to assure traffic A frames still meet latency requirements!
- Any change to the engineered network requires re-validation.



Parameters for the Simulation Slides

OMNEST

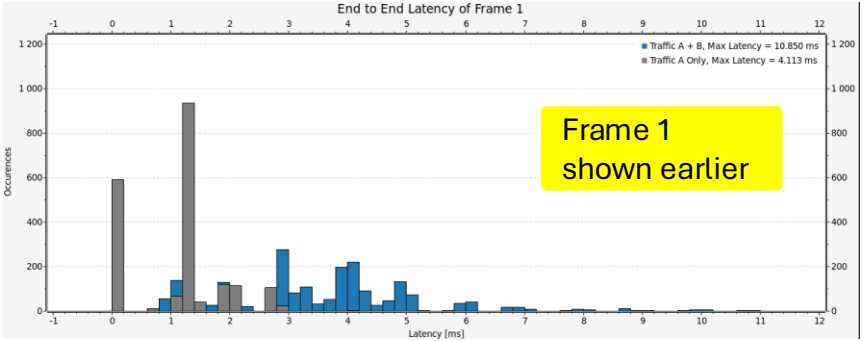
Link		4 nodes, N0 - N3				
		# periodic frames	min period (ms)	max period (ms)	min length (bytes)	max length (bytes)
Traffic A:						
	N0	22	1	1000	64	617
	N1	14	5	1000	93	1305
	N2	11	10	1000	70	98
	N3	11	10	1000	70	105
Traffic B:						
	N0	2	10	25	133	311
	N1	9	2	100	93	128
	N2	4	25	50	94	106
	N3	2	25	50	99	106

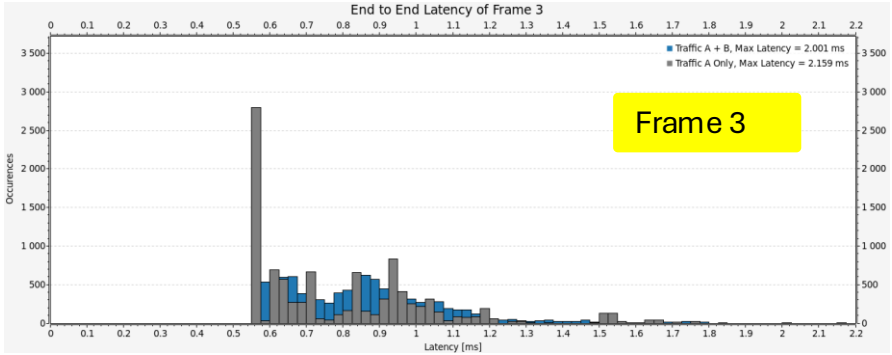
Total # of messages = 75

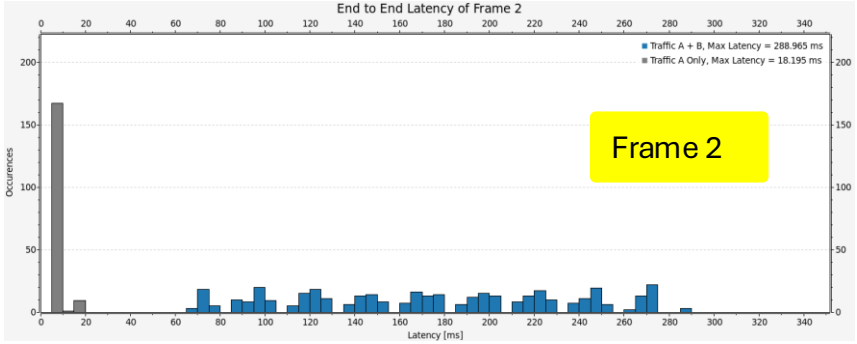
Adding traffic affects frame latencies in unpredictable ways

Grey: Traffic Profile A, 58 periodic frames

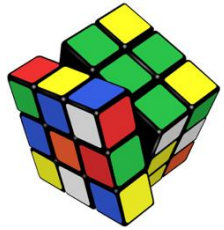
Blue: Traffic Profile A+B, 58 + 17 periodic frames







	Latency impact of added traffic
Frame 1	2.6x
Frame 2	15.9x
Frame 3	1.0x



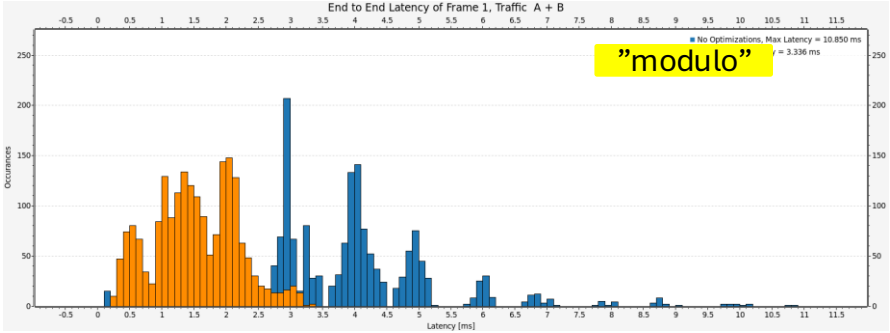
Controls and PLCA Latency Observations

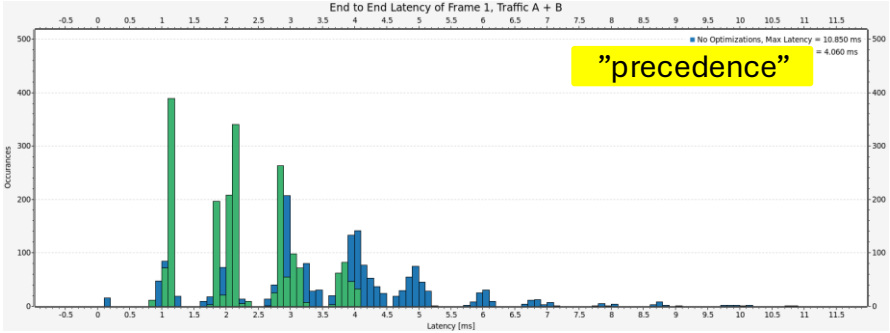
- Controls application (sensor → compute, and compute → actuator) have their own cyclic periods.
- 10BASE-T1S PLCA cycles insert variable latencies ahead of the message on hand (top of the queue).
- Observed latency over 10BASE-T1S PLCA network is highly dependent on the other traffic from the other nodes.
- It is possible to engineer a message latency over 10BASE-T1S PLCA network to a desired range. Every one of the latency sensitive message must be engineered and verified.
- It is exceedingly difficult to hold all the message's latencies within the desired range when add/delete/move/changes… wait…

… we have a need, so **some of us invented**… (next slide)

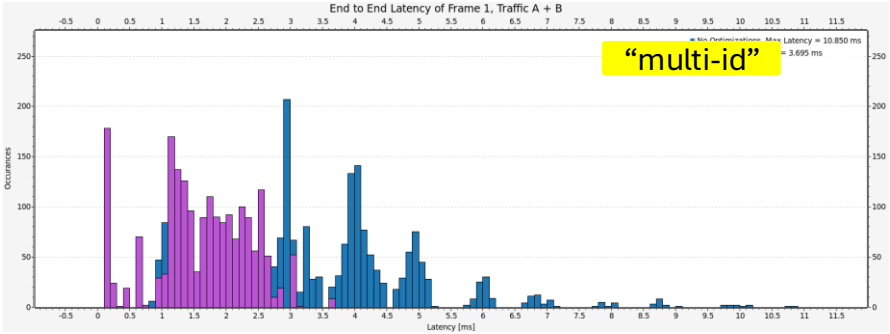
Technical Feasibility – existence there of

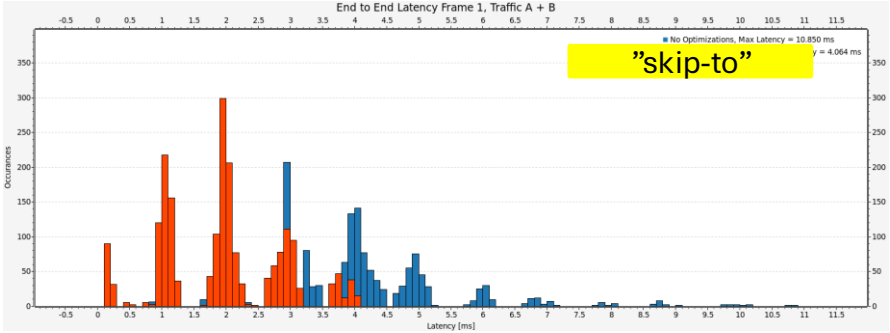
One or more of these method may be used to achieve desired determinism and latencies.





Caveat. These diagrams are for the latencies of one message. What about the rest of 74 messages?





Problem Statement (again)

- Controls Applications have strict sense, compute, accurate cycles.
- Variable cycle lengths cause worst-case latencies that are way longer than what low latency control applications can tolerate.
  - Current bounded latency is way too long.
  - We need a better way to constrain the max. latency of controls traffic.
- Adding traffic to a link impacts the latency of existing traffic in unpredictable ways. We need deterministic latencies for critical control traffic.
- Proprietary solutions are emerging to help solve some aspects of the latencies in system design.

What do we need to study?

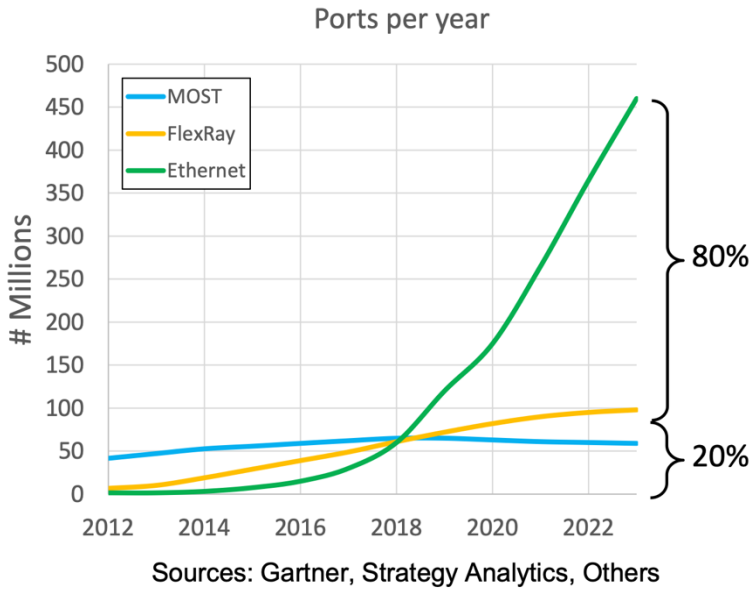
- Can the maximum latency be constrained such that both bandwidth and latency are deterministic on this shared bus?
- In this context:
  Latency variability to be bounded, and absolute latency value for critical traffic to be much smaller value.
- Unequal access (differentiated services) over fair access must avoid starvation of lower class of service.
- Evaluate and study proprietary schemes and help avoid fragmentation and encourage interoperable standard based solution.

Market Needs

Automotive Market Size

- From 802.3cg CFI (did not include multidrop), 2016 https://www.ieee802.org/3/cfi/0716_1/CFI_01_0716.pdf, slide 38.

Automotive Market Size



Sources: Gartner, Strategy Analytics, Others

- ~4 Billion IVN ports in 2019
- Not targeting
  - 2 Billion CAN
  - 1.3 Billion LIN
- 10Mb/s market leveraging from
  - FlexRay
  - New Sensor Applications
  - New ECU developments
  - Where 100Base-T1 exceeds requirements
- Gartner expects an 80:20 ratio between 100Mb/s and 1Gb/s PHYs
- 10Mb/s can expect to be greater or equal to 1Gb/s automotive market

Multidrop Targets.

Inclusion in 2018.

38

IVN Market Evolved Around Single Pair Multidrop

- 2016 • Key market application for the Single Pair Ethernet operating at 10 Mb/s short-reach was for automotive networks when 802.3cg task force was formed.
- 2018 • 10BASE-T1S single pair multidrop (PLCA) was added to 802.3cg to serve the automotive Ethernet needs.
- 2020 • The automotive market for 10BASE-T1S has evolved around multidrop networks. Automotive application is primarily in mechatronics controls applications, body, chassis, etc controls, often replacing or augmenting CAN and LIN buses. Appropriate for the low-speed aspect of 10 Mbps Ethernet. OPEN Alliance helped to accelerate adoption.
- 2022 • Very few 10BASE-T1S full-duplex point-to-point installations exist due to component availability.
- 2022 • Vehicles with 10BASE-T1S networks are in development and in production planning in very near future.
- 2022 • Many early adopters of production systems ran into non-deterministic latencies
- 2024 • Many proprietary scheme are already available that expands the ‘toolkit’ for the engineered network design work. Some of these schemes address determinism, not just relative priority.
- 2026 • **Time to start a study group to address this need and provide interoperable standardized solution to this market.**

Q&A

Expert Panel: <names, affiliation>

Contributors & Supporters

<To be compiled and inserted.>

- Systems Integrators
- Tier 1 Suppliers
- Tier 2 Suppliers
- Users

Straw Polls

Call-for-Interest Consensus

- Should a study group be formed for “10SPE Multidrop Latency Enhancements”?

- Yes ____; Abstain ____;
- Room count: _____

Participation

- I would participate in a “10SPE Multidrop Latency Enhancements” study group in IEEE 802.3.
  - Tally: ___
- My company would support participation in a “10SPE Multidrop Latency Enhancements” study group.
  - Tally: ___

Thank You for Joining

Backup

Variable cycle length problem…



Nodes & Traffic:

- N0: Three periodic low latency frames LLa, LLb, LLc (64 bytes, every 2.5 ms)
- N1 to N7: Two periodic best effort frames Ba, Bb per node (each 430 bytes, every 20 ms)
- Assume application task on N0 creates data for LL frames at t=0, t=2.5, t=5, . . .

Scenario 1 (best case)

Applications on N1,…,N7 make all data for Ba and Bb frames available after LL3 frame was sent):

LL1, LL2, and LL3 data available

Cycle 1								Cycle 2								Cycle 3								Cycle 4								Cycle 5								Cycle 6								
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0
LL1	-	-	-	-	-	-	-	LL2	-	-	-	-	-	-	-	LL3	Ba1	-	Ba3	Ba4	Ba5	-	Ba7	-	Bb1	Ba2	-	Bb4	Bb5	Ba6	-	-	-	-	Bb3	-	-	Bb6	Bb7	-	-	Bb2	-	-	-	-	-	-

t=0.15 ms

Scenario 2 (worst case)

Applications on N1,…,N7 make all data for Ba and Bb frames available before LL3 was sent:

LL1, LL2, and LL3 data available

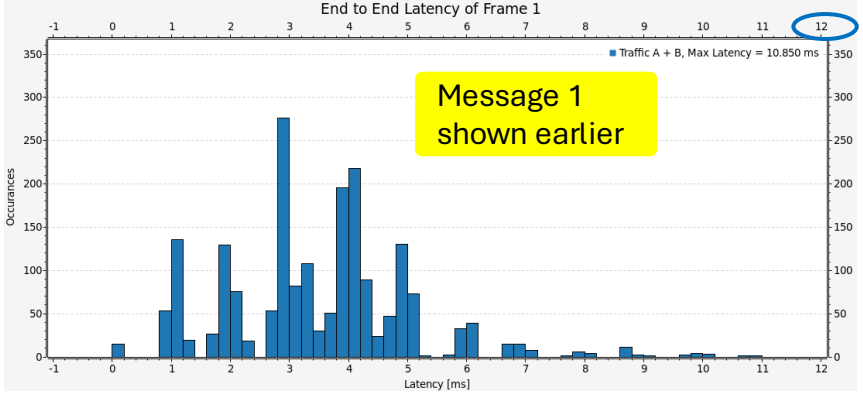
Cycle 1								Cycle 2								Cycle 3				
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4
LL1	Ba1	Ba2	Ba3	Ba4	Ba5	Ba6	Ba7	LL2	Bb1	Bb2	Bb3	Bb4	Bb5	Bb6	Bb7	LL3	-	-	-	-

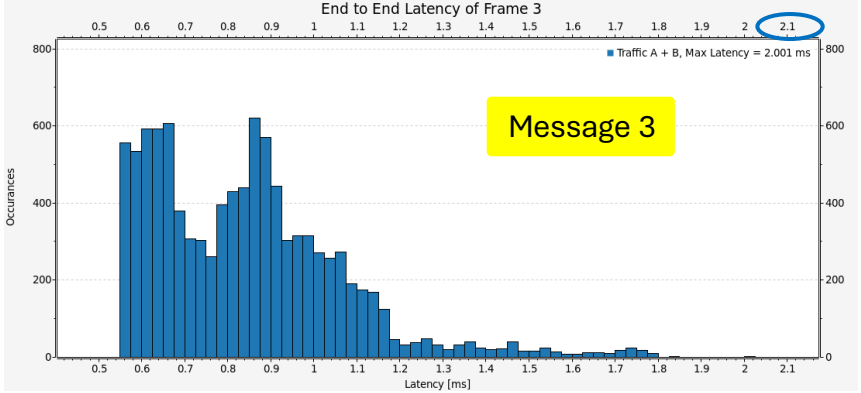
t= 4.3ms

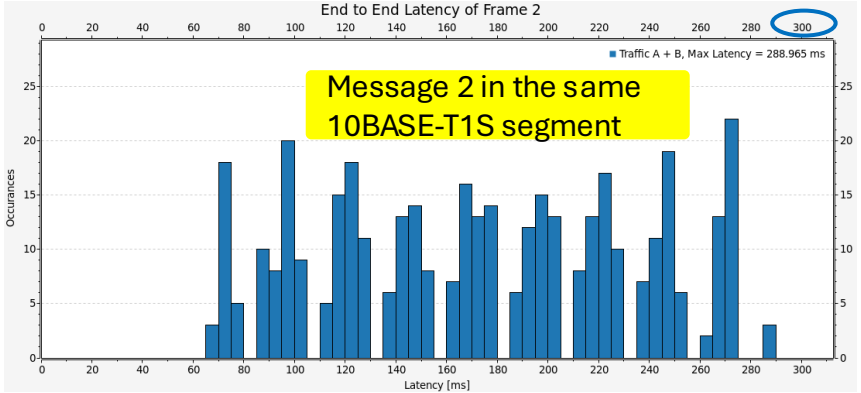
- **Controlling/coordinating the times at which data is created by application tasks on different nodes adds complexity.**
- **This is even more so the case, if some data sources are ‘remote nodes’ located somewhere in the switched part of the vehicle network.**
- **Introducing Differentiated Services to get the variable cycle length problem under control is the better option**

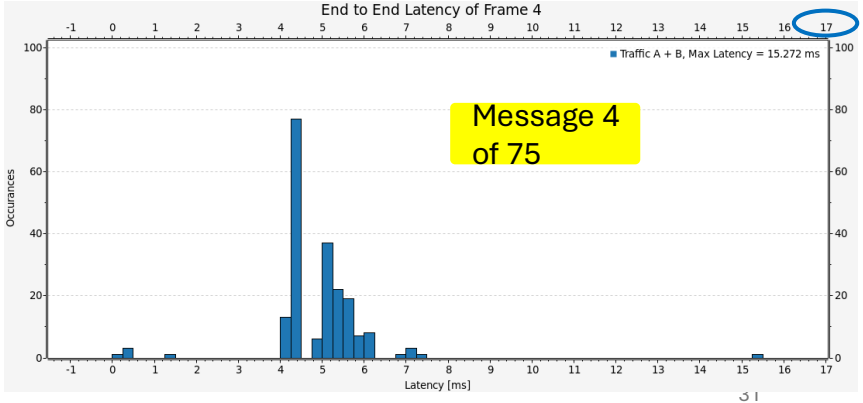
Illustration of latencies associated with three other messages in the same 10BASE-T1S network

Engineering one message's latency affects latencies of others. What to do… what to do….



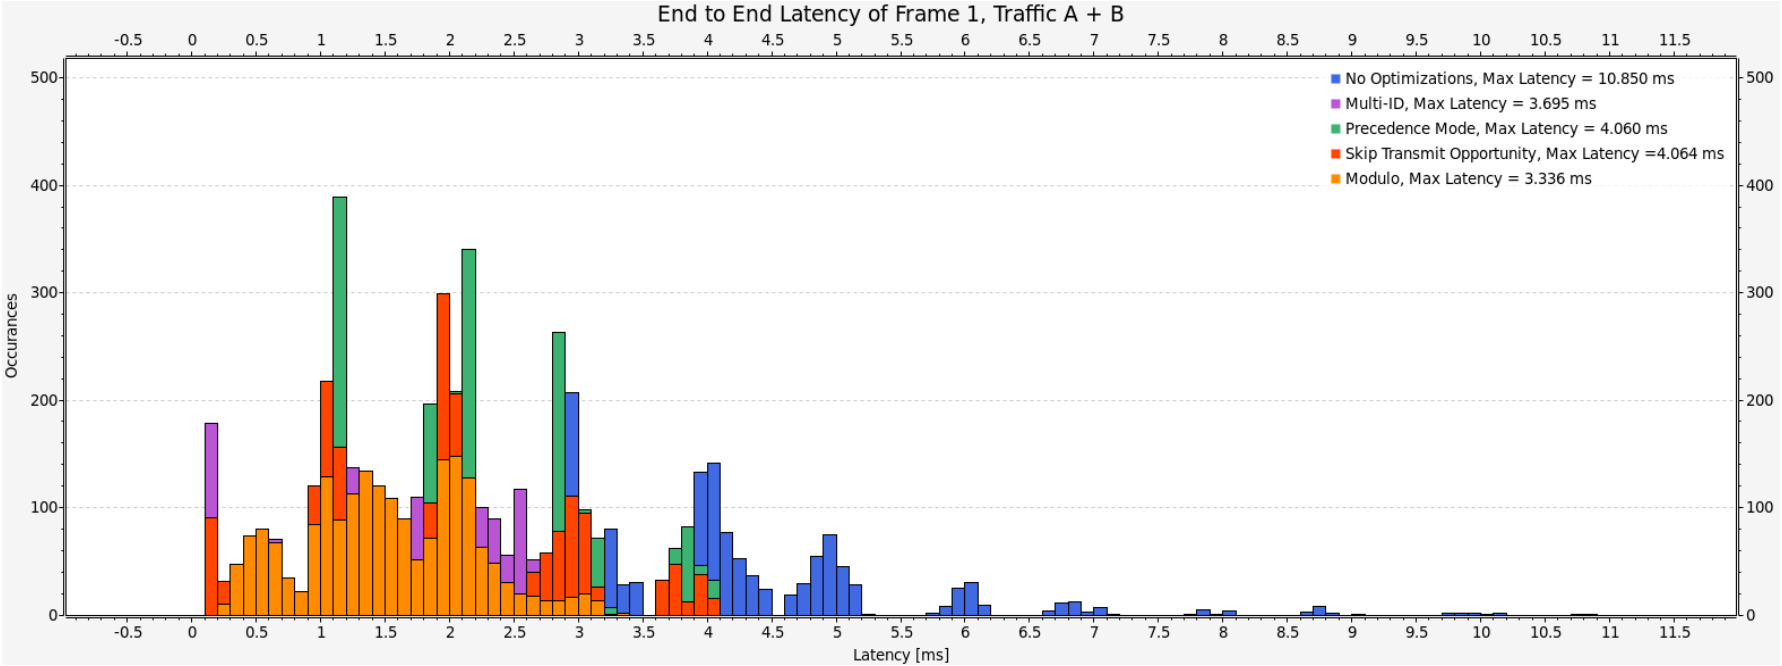






Technical Feasibility – Four Methods Illustrated.

Backup Slide. Same data on a same figure.
Still comparing different methods on that same message's ("Frame 1") latencies.



SP Multi-drop never had a CFI

History in 802.3cg

- Single Pair Multidrop Straw Poll, 2017-07
- https://www.ieee802.org/3/cg/public/July2017/multidrop_straw_polls_3cg_01_0717.pdf
- Single Pair Multidrop Discussion (P2MP, CSMA/CD), 2017-07
- https://www.ieee802.org/3/cg/public/July2017/brandt_cg_01a_0717.pdf
- PLCA proposal for 10BASE-T1S (multidrop), 2017-11
- https://www.ieee802.org/3/cg/public/Nov2017/beruto_3cg_01a_117.pdf
- https://www.ieee802.org/3/cg/public/Nov2017/8023cg_plca_overview_revA.pdf
- Adopted by motion #10, 2017-11, Y=30, N=2, Abstain=35
- https://www.ieee802.org/3/cg/public/Nov2017/motions_3cg_01_1117.pdf
- Objectives with Multidrop approved by 802.3, 2018-03
  PAR and 5C also modified but without addition of multidrop.
- https://www.ieee802.org/3/cg/objectives_3cg_0318.pdf

FAQ: (Relative) Priorities vs Determinism

- Relative Priorities (High, low, etc.) are great. But it does not help with bounded latency and determinism.
- If determinism is not given, everyone would request high priority. If everyone is important, no one is important.
- Determinism does not mean lowest latency. It means bounded latency and assured delivery.



FAQ: Multiple Node IDs (”Weighted Round Robin”)

- A single physical node may contain more than one PLCA node IDs such that each node ID gets a transmit opportunity.
- This allows for weighted round robin transmission onto the shared medium.
- How transmit queues (MAC client) use these multiple node IDs are highly implementation and systems integrator dependent.
- It’s a variation on priority scheme. Not determinism.

Call for Interest Draft

Differentiated Services Enhancements on 10BASE-T1S PLCA (toward Deterministic Network)

The Physical Layer Collision Avoidance (PLCA) layer developed as a part of 802.3cg provides a fair and equal access over shared medium while providing bounded maximum access latency. This method has been widely adopted in automotive Ethernet networks onto production vehicles. The maximum access latency is not sufficiently bounded in providing deterministic latency for many of the controls applications. In this context, control loop applications have Sense --> Compute --> Actuate --> <repeat> such as power train, suspension, power assisted lift gate, motion control, etc.

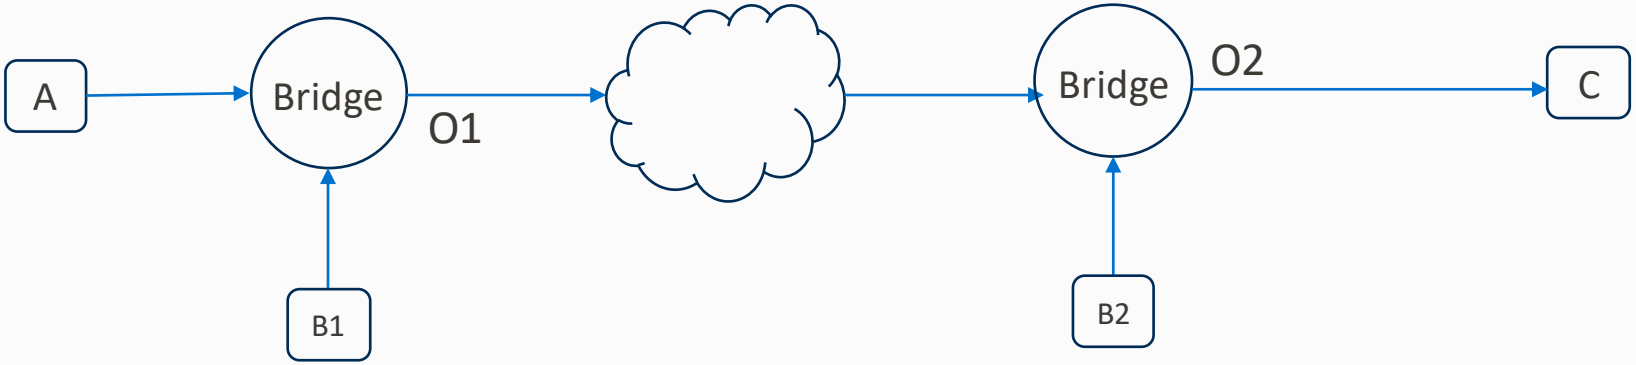
The PLCA access method provides fair and equal access. Sufficiently bounded access latency may be achieved through differentiated services enhancement to PLCA such that the maximum access latency is made to be deterministic. This need has been identified in prior projects in the past. We believe there are more than one schemes that may address this need.

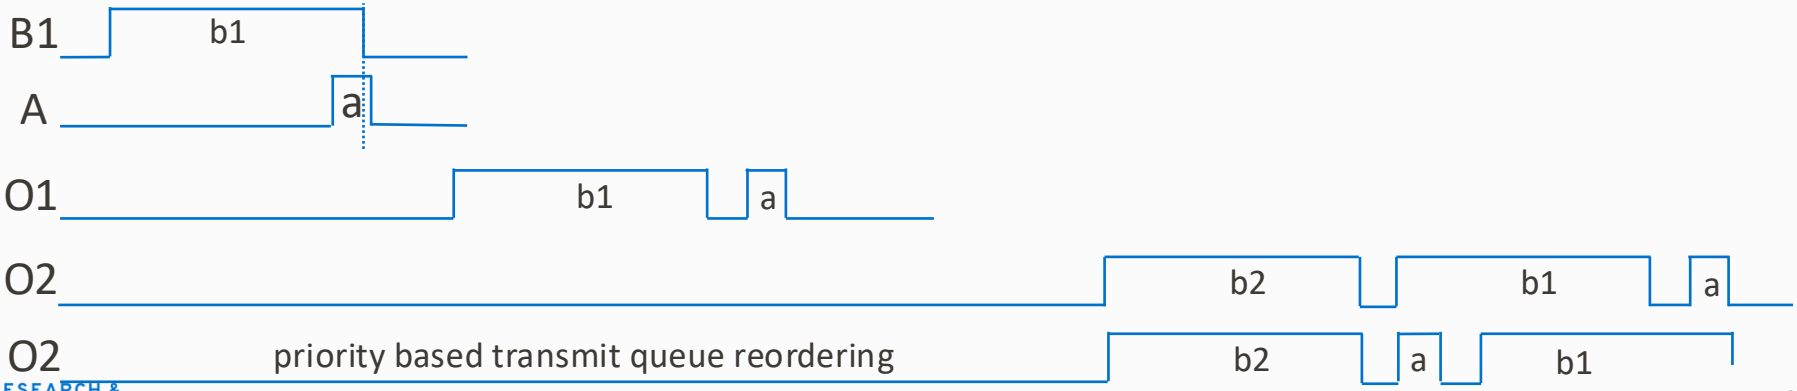
This is a call for interest to initiate a Study Group to explore the market requirements for an interoperable solution and to develop a PAR and CSD for Differentiated Services Enhancements on 10BASE-T1S PLCA.

Supporters

- Bhagyashri Katti, General Motors
- Steve DiBella, General Motors
- Hiroshi Ono, Analog Devices, Inc.
- Sean Mullings, Analog Devices, Inc.
- David O’Neill, Analog Devices, Inc.

A Simple Link with Head-of-Line Blocking





Control Loops (or Controls Applications)

- [Sense] → [Compute] → [Actuate] control loop.
- The worst case end-to-end latency often determine the robustness and feature of the controls applications.

