Meeting the AVB Gen2 Latency Requirements

Christian Boiger
christian.boiger@hdu-deggendorf.de

IEEE 802 Plenary
November 2011
Atlanta, GA

Automotive and Industrial Latency Requirements

Copied from the assumptions document [1]:

- Automotive:
  - 100 µs w/5 FE hops
  - 128 bytes payload
  - max 7 hops

- Industrial:
  - < 3 µs / hop (Interfering Frames (includes other same PCP frames) + Bridge Latency (not including Store Forward Latency))
  - max 64 hops (daisy chain)

[1] http://www.ieee802.org/1/files/public/docs2011/avb-dolsen-gen2-assumptions-0911-v6.pdf

Implications of the Industrial Latency Requirements (1)

Latency = tDevice + tInterferingNonUl + tInterferingUl + tUl

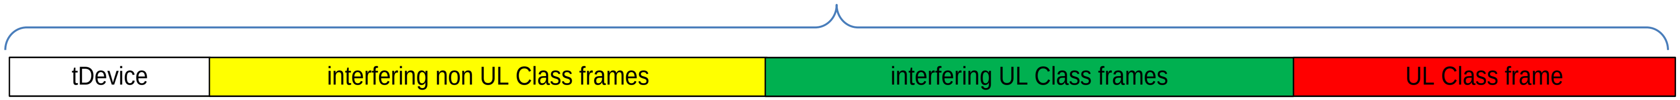

Industrial Goal:

tDevice + tInterferingNonUl + tInterferingUl < 3µs

tDevice = 0.512 µs – 1.024 µs

⇒ tInferferingNonUl + tInterferingUl < 1.976µs – 2.48µs

⇒ InterferingNonUl + InterferingUl < 247 bytes – 310 bytes

Implications of the Industrial Latency Requirements (2)

InterferingNonUl + InterferingUl < 247 bytes – 310 bytes

Assuming preemption with a minimum framelet length of 128 bytes

⇒ InterferingNonUl = 128 byte (payload incl. pre-emption overhead)

⇒ 150 bytes frame

⇒ 158 bytes packet

⇒ 170 bytes packet and IPG

⇒ InterferingUl = 77 bytes – 140 bytes

(minimum size packet + IPG = 84 bytes)

⇒ **It has to be guaranteed that UL class frames do not interfere with other UL class frames**

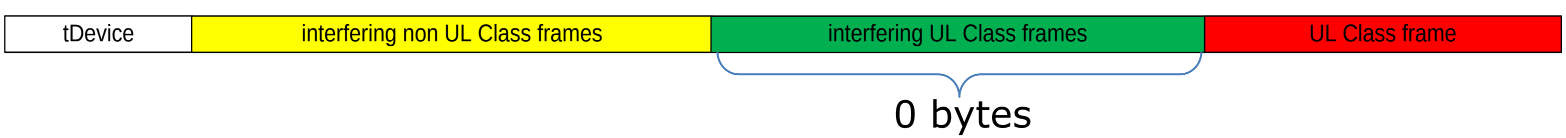

Implications of the Automotive Latency Requirements (1)

Latency = tDevice + tInterferingNonUl + tInterferingUl + tUl

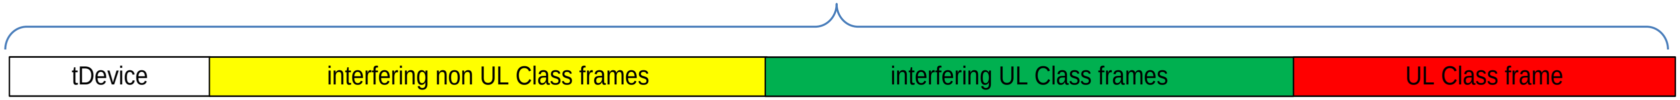

Automotive Goal:

tDevice + tInterferingNonUl + tInterferingUl + tUl < 20 µs (5 FE Hops)

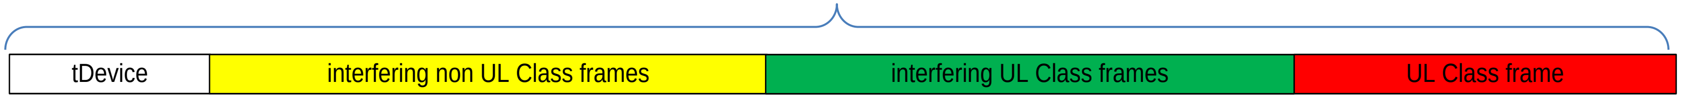

tDevice = 5.12 µs

tUl = 12.64µs (158 bytes)

⇒ tInferferingNonUl + tInterferingUlFrames = 2.24 µs

⇒ InterferingNonUl + InterferingUl = 28 bytes

⇒ **28 bytes are not enough to send a minimum size frame**

⇒ **Pre-emption alone is not enough to achieve this goal**

Implications of the Automotive Latency Requirements (2)

tDevice + tInterferingNonUl + tInterferingUl + tUl < 33.3 μs (3 FE Hops)

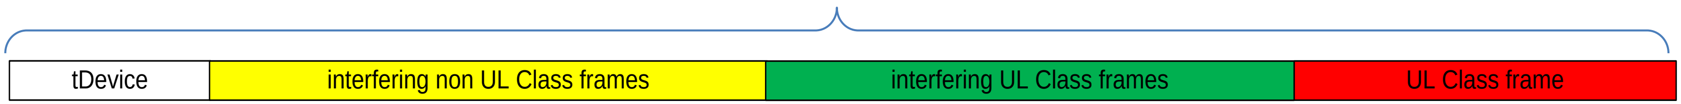

tDevice = 5.12 μs

tUl = 12.64μs (158 bytes)

⇒ tInferferingNonUl + tInterferingUlFrames = 15.54 μs

⇒ InterferingNonUl + InterferingUl = 194 byte

⇒ InterferingUl = 24 byte (assuming a minimum framelet size of 170 bytes)

⇒ **Also in this case interfering UL traffic has to be avoided**

Possibilities to Guarantee No Interference with UL Traffic

- Coordinating the talker
  - TDMA scheduler in the talker
  - The talker is sending UL class packets at a specified timeslot
  - The scheduler in the talker has to be configured by network management which makes sure that there is no interference of two UL streams within the network (assuming all low latency streams are well known and engineered)

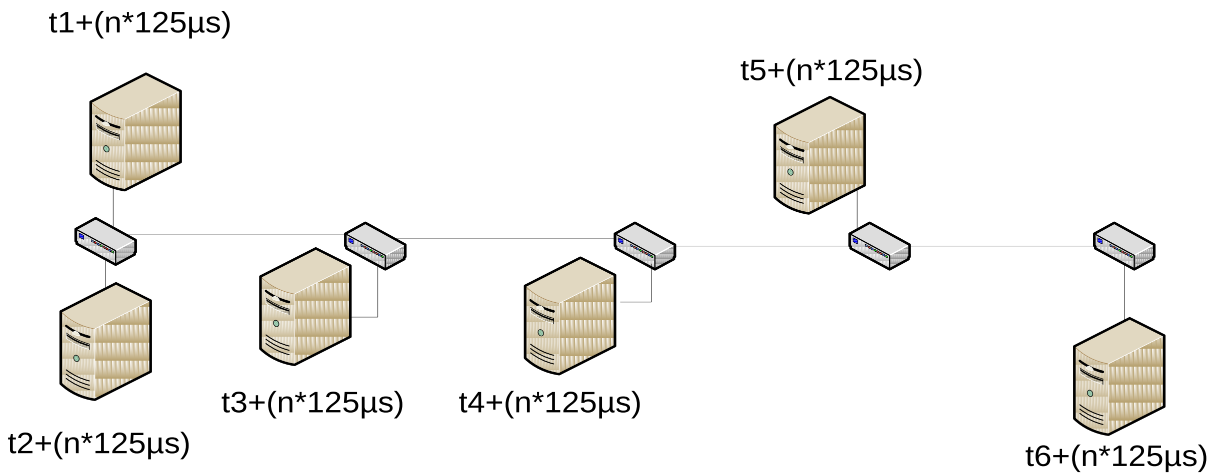

Possibilities to Guarantee No Interference with UL Traffic

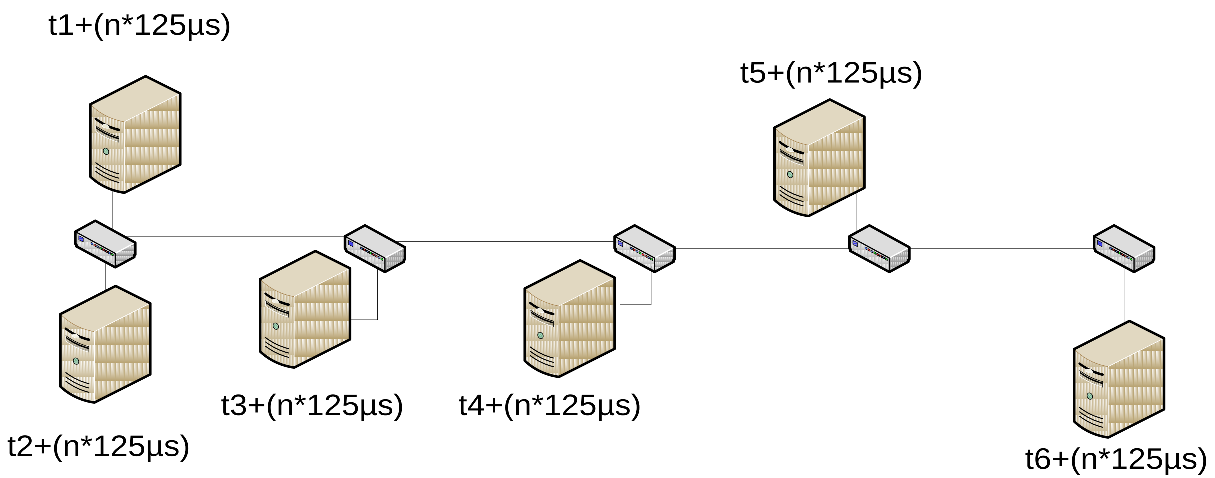

t1 – t6 have to be chosen in such a way that the UL packets will never interfere with each other

Is this possible?

Possibilities to Guarantee No Interference with UL Traffic

"Schedule" of an bridge egress port:

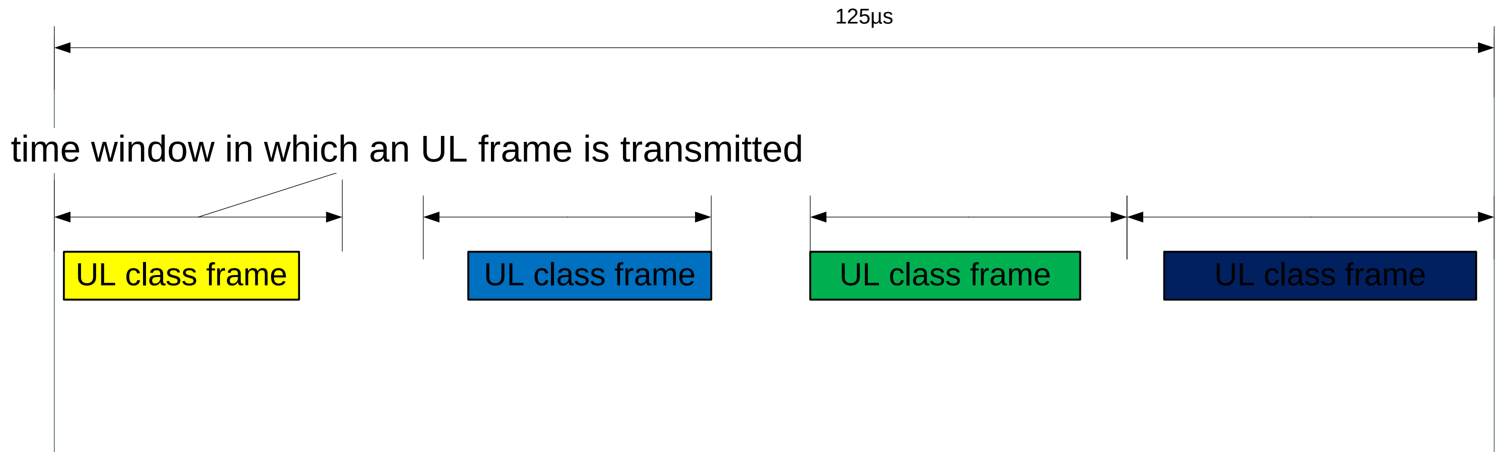

How big has the time window to be?

UL Packet Jitter (GigE)

- The time slot calculation which has to be done by management has to consider the UL packet jitter

- Jitter is added at every hop

- MaxJitter = MaxLatency - MinLatency
- MaxJitter = tDevice + tInterference + tUl – (tDevice + tUl)
- MaxJitter = tInterference
- assuming a maximum interference of 170 (min size framelet + IPG):

  MaxJitter = 170 bytes * 8 bit/byte / 1000 Mbit/s = 1.36 µs

  There might be additional forms of jitter which have to be added to MaxJitter e.g. the jitter of the clock (in the talker)

Consequences of UL Packet Jitter in Big GigE Networks

- MaxJitter/hop = 1.36 µs
- Jitter in a daisy chain with 64 hops

  1.36 µs / hop * 64 hop = 87.04 µs

  i.e. in the case of an stream from the first node of a daisy chain to the last one, at the last hop a window of 87.04 µs has to be reserved, as the packets can arrive at any point within this window

$\Rightarrow$ TDMA only in the talker works only in small GigE networks

$\Rightarrow$ in bigger networks this approach is unrealistic (especially for small transmission periods (down to 31.25 µs [1]))

[1] http://www.ieee802.org/1/files/public/docs2011/avb-dolsen-gen2-assumptions-0911-v6.pdf

UL Packet Jitter (FE)

- MaxJitter = tInterference
- Assuming a maximum interference of 170 bytes (minimum size framelet + IPG):

  MaxJitter = 170 bytes * 8 bit/byte / 100 Mbit/s = 13.6 µs
- Jitter at the 5th hop: **68 µs**
- Jitter at the 7th hop: **95.2 µs**
- Transmission period: 500 µs [1]
- Max number of simultaneous transmissions: 8 [1]
- => 500 µs / 8 = 62.5µs
- => Maximum allowed jitter = 62.5 µs – 13.6 µs = **48.9 µs** (assuming equally spaced frames within the 500 µs transmission period)
- => **Interference between UL frames**

[1] http://www.ieee802.org/1/files/public/docs2011/avb-dolsen-gen2-assumptions-0911-v6.pdf

Would Cut-Through Reduce the Jitter?

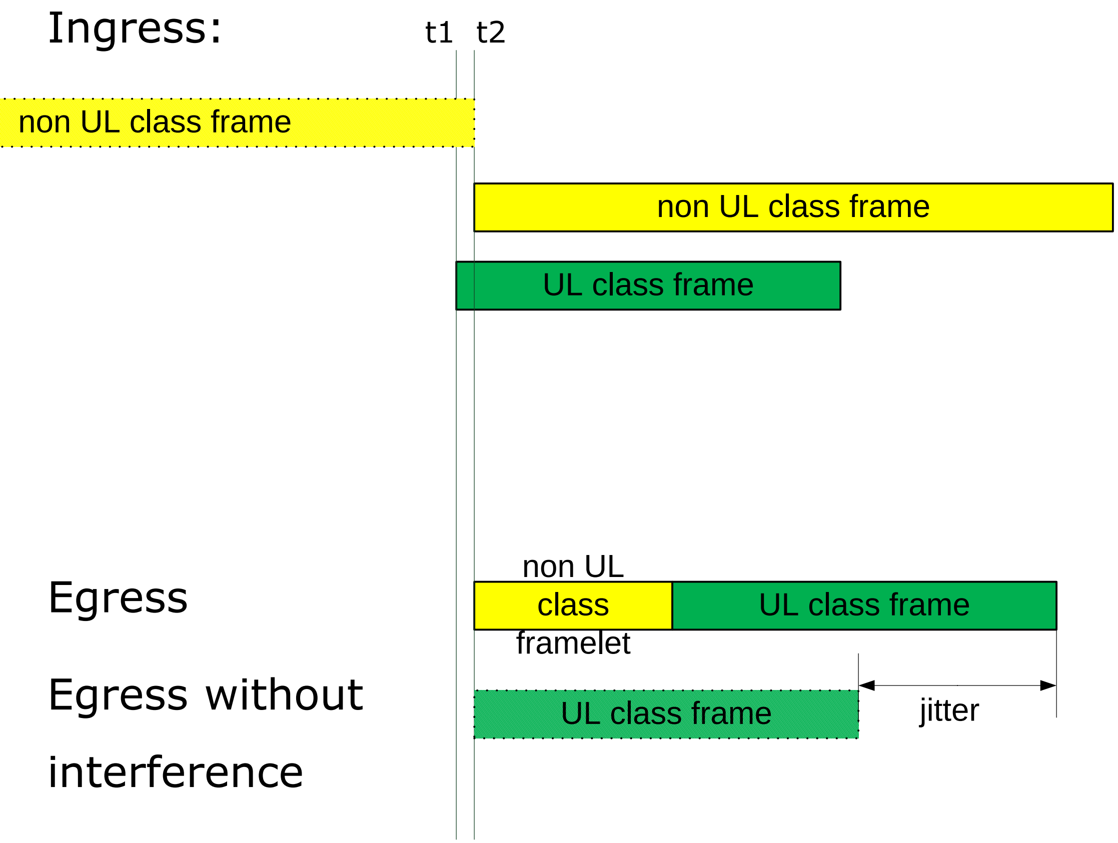

t1: first bit of UL class frame is received

t2 – 1 clk cycle:

non UL class frame is ready for transmission

t2: UL class stream is ready for transmission

Striped non UL class frame: store and forward for non UL class traffic

Striped UL class frame: frame in the case with no interference

Would Cut-Through Reduce Jitter?

- No, the jitter is the same.
- The jitter depends on the max interfering frame. As the size of the maximum interfering frame is not changing with cut-through, cut-through would only decrease latency.

=> Additional mechanisms are necessary to improve latency for GigE networks. Pre-emption only is not enough, especially in bigger networks.

Possibilities to Reduce Jitter

- In order to reduce the jitter, the time of arrival of a packet has to be less variable
- It has to be guaranteed, that a UL class frame is not interfering with any other frame (including framelets).

- Possible Solutions:
  – Banning legacy traffic
  – TDMA in bridges
  – Any other?

- The Time Aware Shaper approach shown in the March meeting [1] would be a possible TDMA solution.

[1] http://www.ieee802.org/1/files/public/docs2011/new-pannell-latency-options-0311-v1.pdf

Combining Pre-emption with TABS and TADS

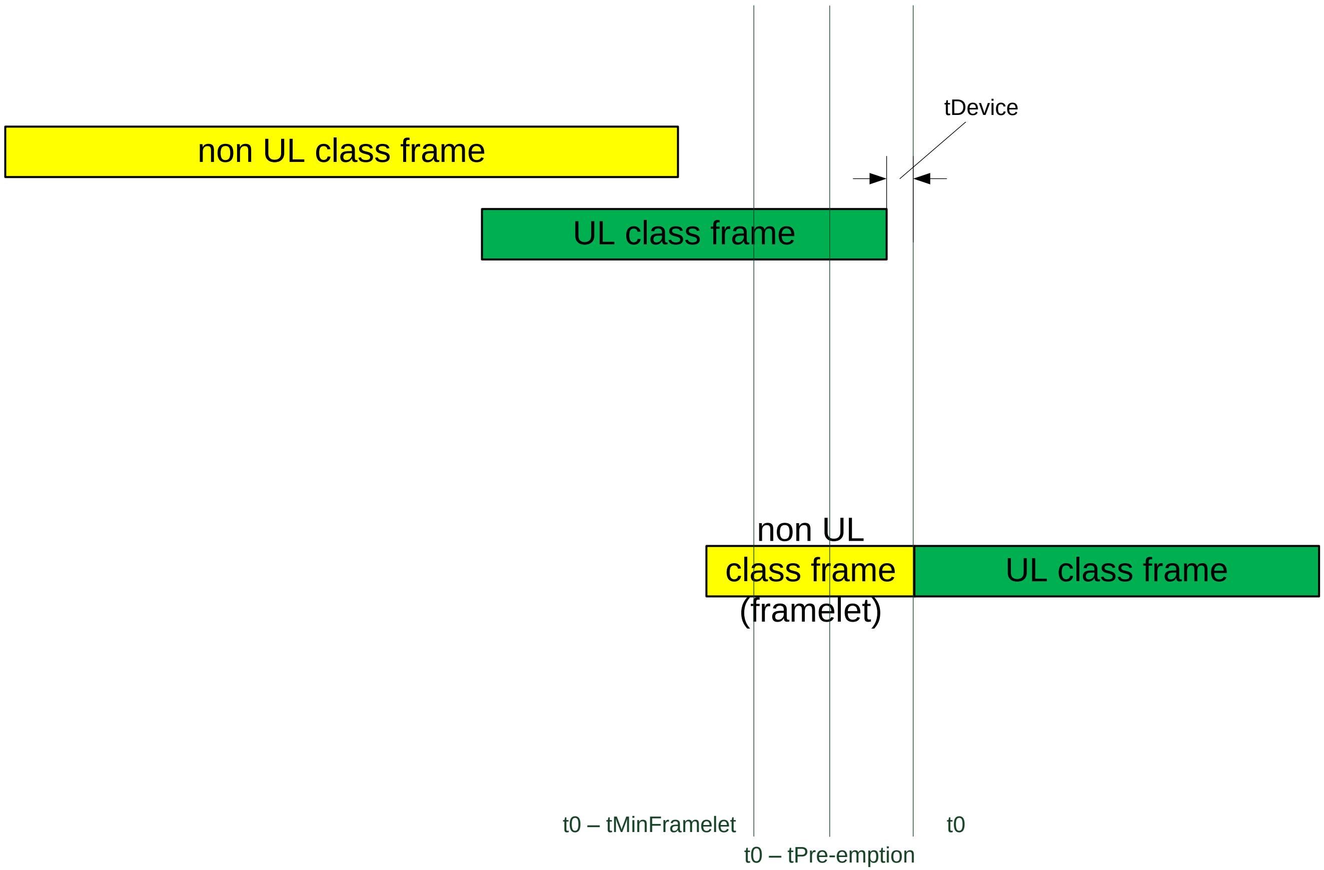

t0: UL class frame is ready to be transmitted and UL class queue is de-blocked

t0-tMinFramelet: Egress port is blocked for non UL class

t0-tPre-emption: Pre-emption mechanism is activated
(t0-tMinFramelet = t0-tPre-emption ?)

Advantages of Pre-emption with TABS and TADS

- For Fast Ethernet Networks

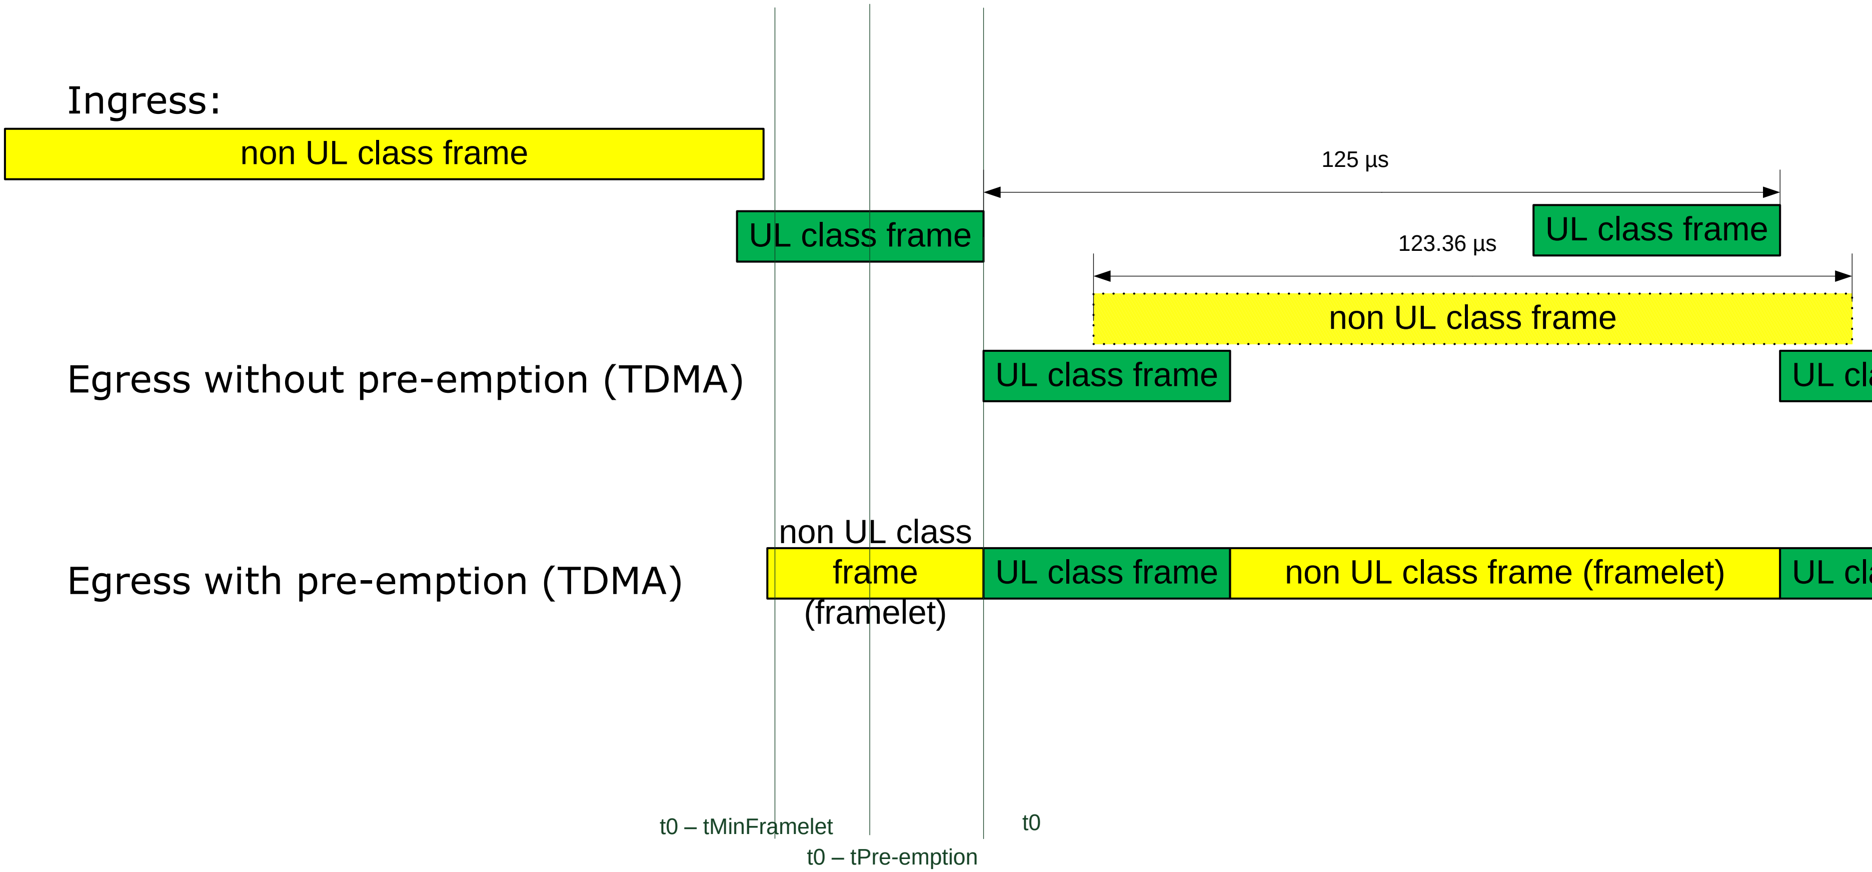

t0: UL class frame is ready to be transmitted and UL class queue is de-blocked

t0-tMinFramelet: Egress port is blocked for non UL class

t0-tPre-emption: Pre-emption mechanism is activated

Advantages of Pre-emption with TABS and TADS

- For Fast Ethernet Networks (Max legacy MTU = 256 bytes)

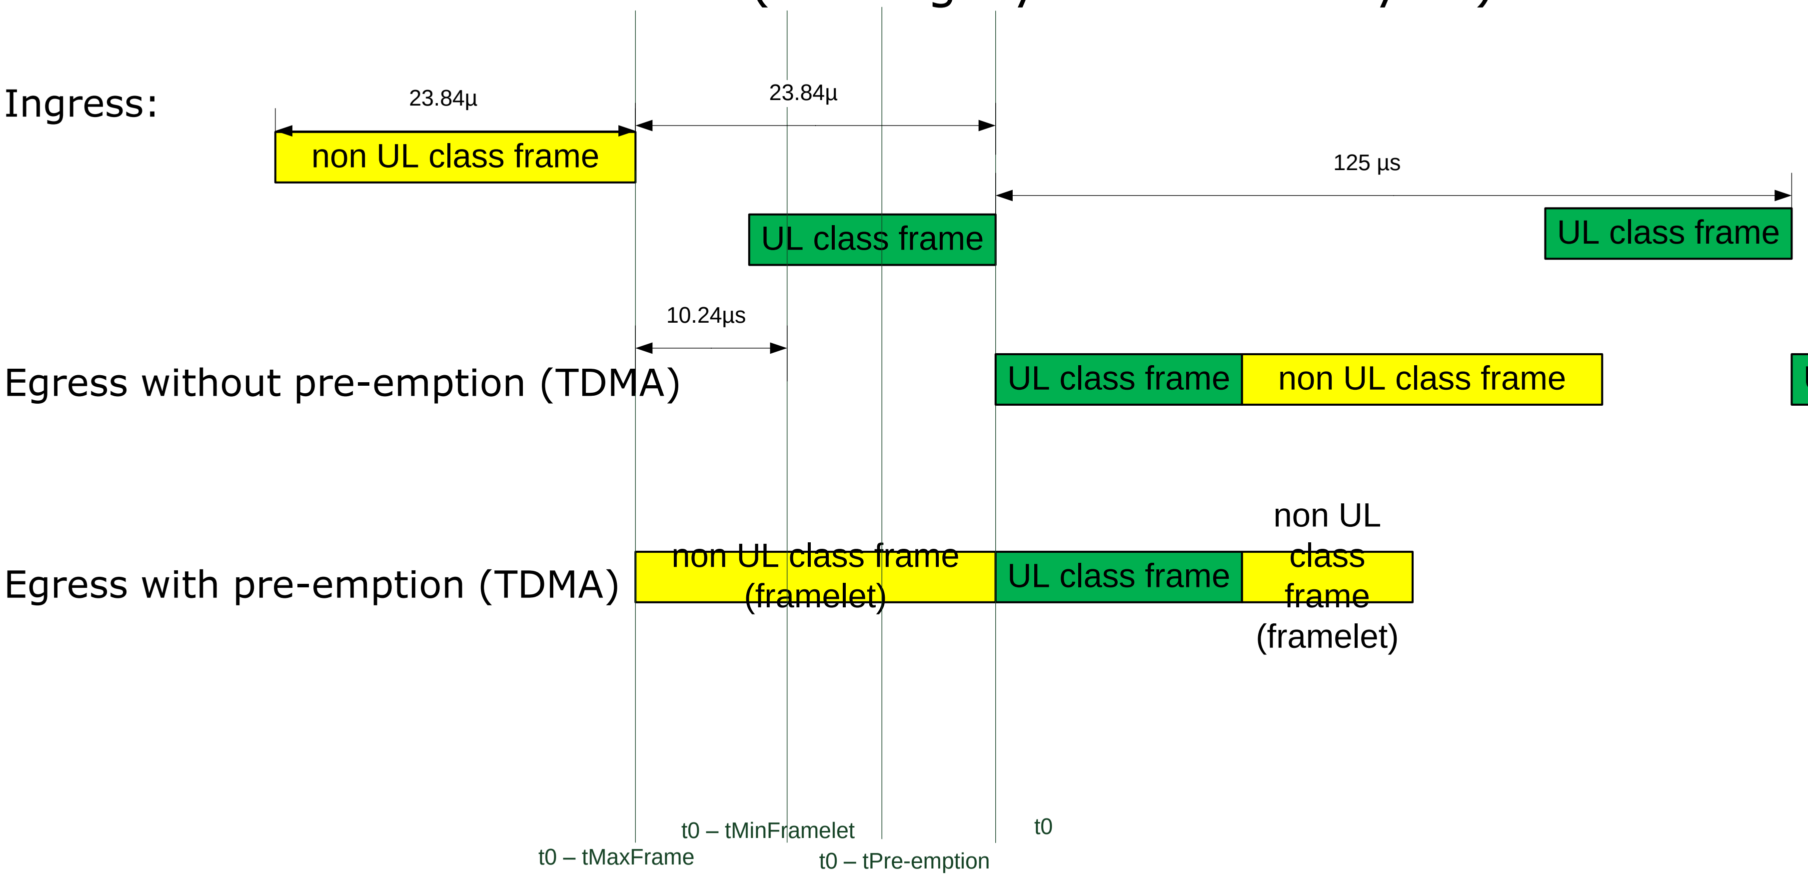

t0: UL class frame is ready to be transmitted and UL class queue is de-blocked

t0-tMinFramelet: Egress port is blocked for non UL class

t0-tPre-emption: Pre-emption mechanism is activated

Advantages of Pre-emption with TABS and TADS

- For Gigabit Ethernet Networks

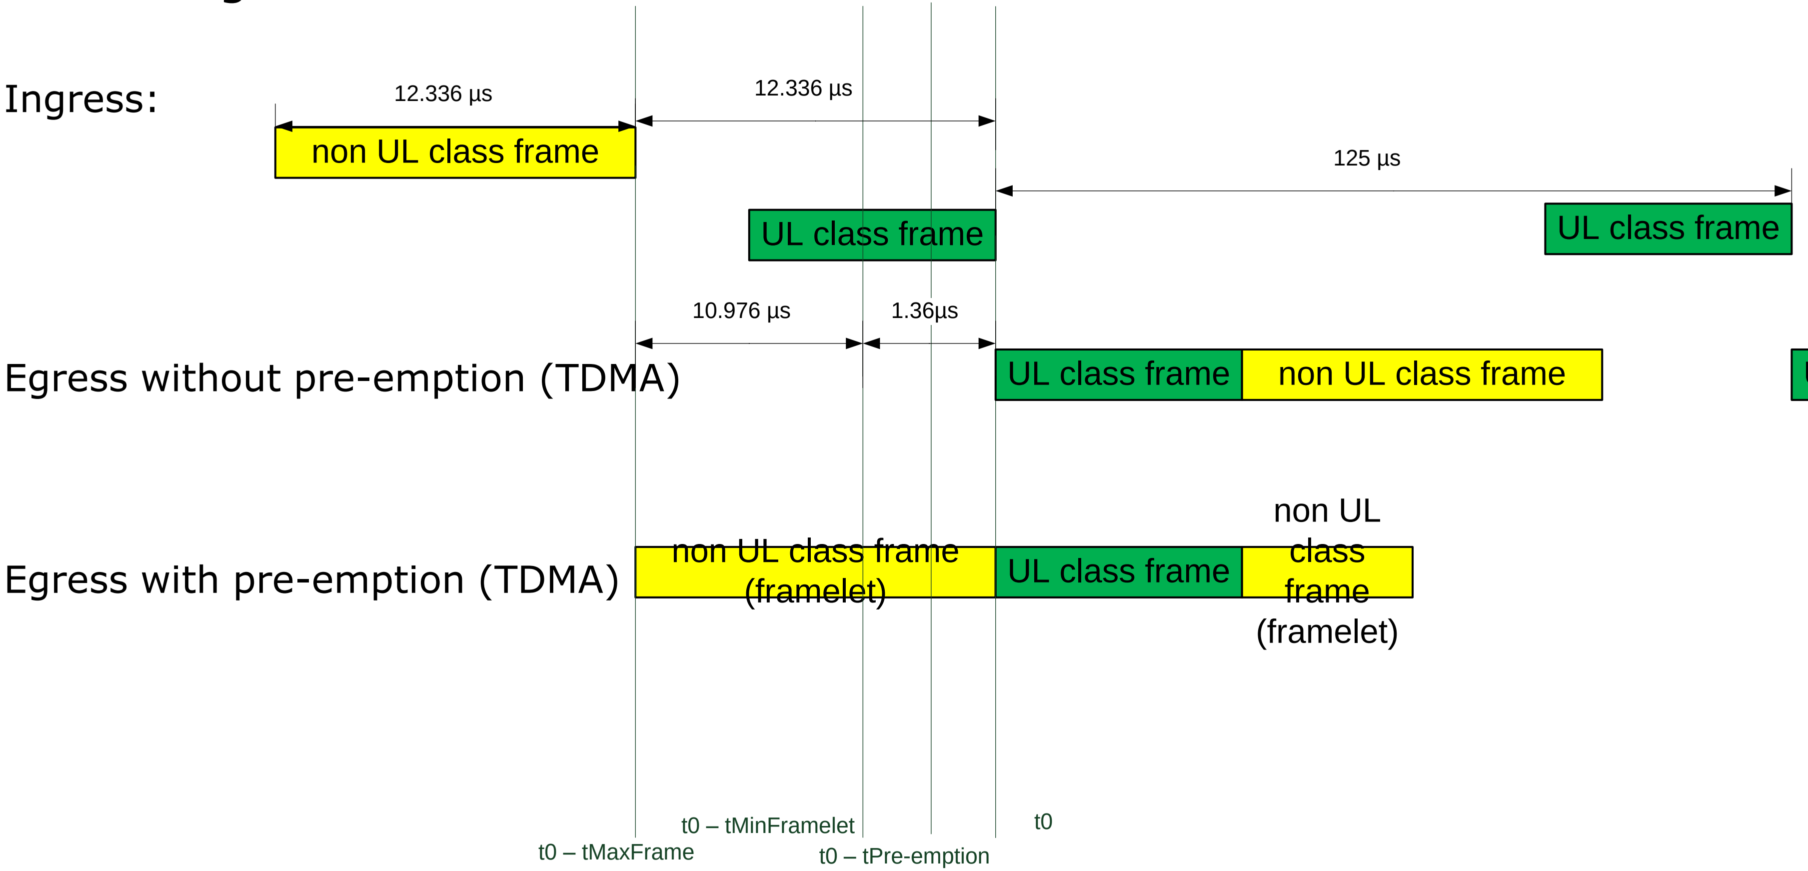

t0: UL class frame is ready to be transmitted and UL class queue is de-blocked

t0-tMinFramelet: Egress port is blocked for non UL class

t0-tPre-emption: Pre-emption mechanism is activated

Resulting Latency

- Latency with pre-empted non UL class frame ahead of the UL Class frame with no interference with other UL class frames:

  Latency = tDevice + tInterferingNonUL + tUl

- Resulting latency with no interfering frame or framelet (TDMA with pre-emption

  Latency = tDevice + tUl

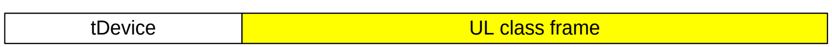

Resulting Latency

- Automotive:
  - Max latency = 100 µ / 5 hops
  - Latency(TDMA + Pre-emption) = tDevice + tUl = 5.12 µs + 12.64 µs
  - Latency(TDMA + Pre-emption) = 17.76 µs
  - => 88.8 µs / 5 hops
- Industrial:
  - < 3 µs / hop (Interfering Frames (includes other same PCP frames) + Bridge Latency (not including Store Forward Latency))
  - Latency(TDMA + Pre-emption) = tDevice = 1.024 µs

TDMA with Pre-Emption

- TDMA in bridges would guarantee the automotive and industrial latency requirements
- TDMA with pre-emption would allow Fast Ethernet networks with small transmission periods
- Pre-emption would increase the bandwidth utilization for Fast Ethernet and Gigabit Ethernet networks with TDMA

Thank You