Can AVB Gen 1 be used for Automotive Control ?

Johannes Specht, University of Duisburg-Essen
Markus Jochim, General Motors Research

Background

- Future: Automotive Control based on AVB Gen 2 (= TSN)
  - Our expectation: AVB Gen 2 (= TSN) will be designed to fullfil automotive requirements for (potentially safety critical) control applications.
  - Desired characteristics like:
    Low Latency, High determinism, Support for adequat traffic classes, Redundancy support, Fault tolerant clock sync, Ingress policing, . . .
  - Some of these topics are currently "work in progress" in 802.1, others may still need to be added.

- However… it will take some time before . . .
  - AVB2 (= TSN) will be specified.
  - implementations will be available.

Background

- The key question is:
  Can Gen 1 serve as an intermediate solution / a migration path to TSN for automotive control?

- Gen 1 certainly offers some properties & guarantees that are of interest to automotive control.
  Example: Bandwidth guarantees, Latency guarantees, Fairness guarantees.

- However, we have several questions related to some of the claimed properties of Gen 1.

The intention of this presentation is to raise these questions to the 802.1 TSN group.

Main Questions (Overview Slide)

1. **Maximum number of Streams**
   Is Gen 1 able to support a sufficient number of streams?

2. **Traffic Specification**
   What is the real meaning of the class measurement interval ?
   What types of bandwidth guarantees are actually in place?

3. **Sending frames at rates lower then once per 250µs**
   SRP does not allow this. Do we see any risk associated with manually configuring the IdleSlope to support lower transmission frequencies?

4. **Latency Guarantees**
   Do we really have a hard and proven E2E latency guarantee of 2ms over 7 hops? Is it worth the effort to define a framework for calculating guaranteed latency values for a given network / topology.
   (Very obviously the guarantees depend on several parameters like the number of hops, the number of ports per bridge etc.)

Maximum number of Streams

How many Class A streams can be reserved?

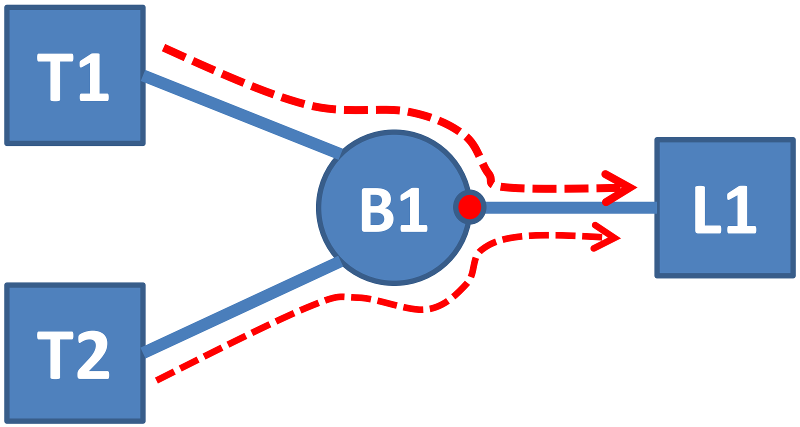

Assume:
Multiple streams forwarded over a common egress port.

AVB bandwidth guarantees express themselves in terms of a T-Spec. T-Spec defines:

– How many frames per class measurement interval (125 µs) ?
– Maximum size of such a frame?

Maximum number of Streams

Calculation of the maximum number of streams:

- Assume:
  - Class A: 125µs class measurement interval
  - Only 1 frame per interval (SRP MaxIntervalFrames = 1)
  - Minimum sized frames (MaxFrameSize = 42 bytes)

This is not even a worst case assumption. It is almost the "best case" !

- Time required to transmit a frame with 42 bytes payload:
  - 64 Bytes + 7 Byte Preamble + 1 Byte Start of Frame Delimiter + 12 Byte Interframe Gap = 84 Byte.
  - @100 Mbit/s: $84 * 8 / 10^{8} = 6.72*10^{-6}$ s. => 6.72 µs

Maximum number of Streams

– Within a window of 125 µs, we can certainly not guarantee the transmission of more than floor(125 / 6.72) = 18 frames
– Considering the recommendation

  “Maximum of 75% reserved traffic in order to allow for a minimum of 25% best effort traffic”

  we can accommodate approximately 13 to 14 streams.

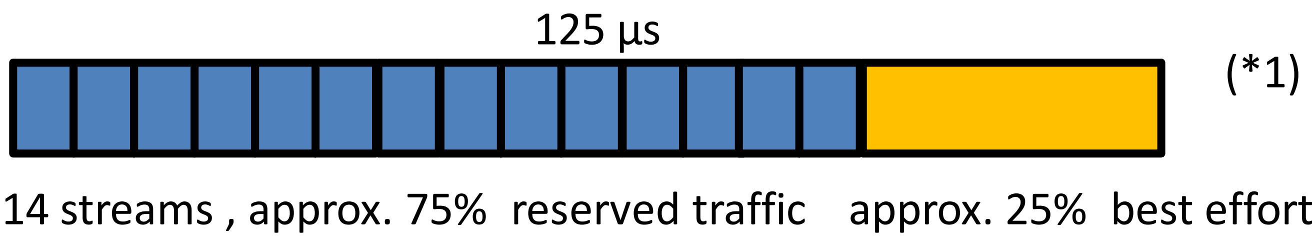

13 to 14 is really not very much!

(*1): The diagram does not intend to imply that the best effort traffic is send at the end of the window

T-Spec & Class Measurement Intervals

We are also wondering over which time period bandwidth guarantees can actually be established!

Simple observation:
Without preemption, a single best effort frame with a payload of 1500 Byte is „standing in the way“ of high priority traffic for 123µs.

Not much can be achieved in the remaining 2 µs ;-)

=> The purpose of the TSpec and the class measurement interval is to limit the talker rather than to guarantee that the talker will actually be guaranteed to be able to transmit *MaxIntervalFrames* frames with a maximum length of *MaxFrameSize* Bytes during a period as short as 125µs!

But over which time period would reserved bandwidth then actually be guaranteed to the talker?

Sending frames at lower rates

- Conclusions from the previous slides:
  - Only a very small number of class A streams can be reserved.
  - Class A and class B transmission periods are unreasonably short for many applications
    (Of course we do not need to send in every transmission interval and of course the unused bandwidth will be available for best effort traffic, BUT it will not be available for additional reservations.)
  - There seems to be no transmission guarantee for a period as short as a class measurement interval anyway!

- These observations may lead to the conclusion that, for control applications, we can ignore class A and class B and instead exclusively configure the system via *adminIdleSlope*.

Any concerns related to this idea? What are the side effects / implications that we may have overlooked ?

Thank you for your Attention!

Johannes Specht
Dipl.-Inform. (FH)

Dependability of Computing Systems
Institue for Computer Science and
Business Information Systems (ICB)
Faculty of Economics and
Business Administration
University of Duisburg-Essen

Schützenbahn 70
Room SH 502
45127 Essen
GERMANY
T +49 (0)201 183-3914
F +49 (0)201 183-4573

specht@dc.uni-due.de
http://dc.uni-due.de

Markus Jochim

ECS Architectures and Vehicle Connectivity
Electrical & Controls Systems Research Lab
General Motors Research & Development
Warren, Michigan USA

markus.jochim@gm.com
www.gm.com