Analysis of worst case latencies in an 10 Mbit Ethernet network with PLCA

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Which effect has PLCA on the latencies?

- PLCA (PHY Level Collision Avoidance) as layer 1 arbitration mechanism has huge impact on the network latencies because of its Round Robin nature
- Typical for automotive networks: network delay of messages of a high priority stream must not exceed a low number of milliseconds e.g. 1ms
- Unfortunately a Round Robin scheme contradicts strict priority schemes.
- **PLCA causes bad worst case latencies for high priority streams.**

Illustrative histogram of a potential latency distribution:

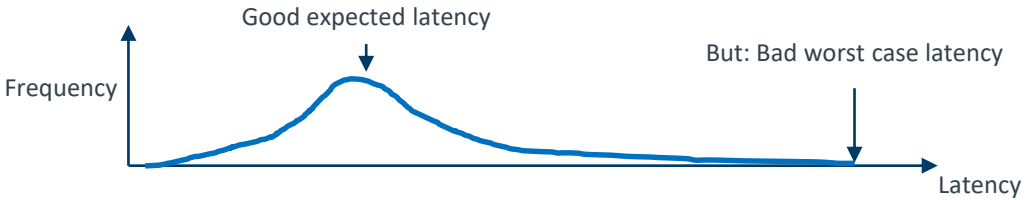

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

A worst case analysis is necessary

- To realize guarantees on latencies a worst case analysis is required.
  - Assumptions for a worst case analysis: Everything goes wrong!
    - All participants want to send all their messages at the same time
    - The execution order is the worst possible one
    - Everyone just misses his slot

- Notes:
  - This analysis is just about latencies which are caused by the arbitration mechanism (Transfer time of the message itself is not included)
  - Analysis is done with a prototype tool (not production ready) and worst case not formally proven yet

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Example use cases: microphone array

- Connecting up to 3 microphones (or other sensors) to a master device
- Every microphone produces a latency sensitive data stream about 1.9 Mbit/s (e.g. 2 channels @ 48kHz with 16 bit per sample)
  - IPv6+RTP based and strictly periodic with short cycle time (1,33 ms)
  - Worst case latency must not exceed the cycle time (1,33 ms)
- For every microphone there is a bidirectional control channel to the master device with about 120 kbit/s
  - IPv6+TCP based and bursty (up to 5 packets with 1500 Byte)

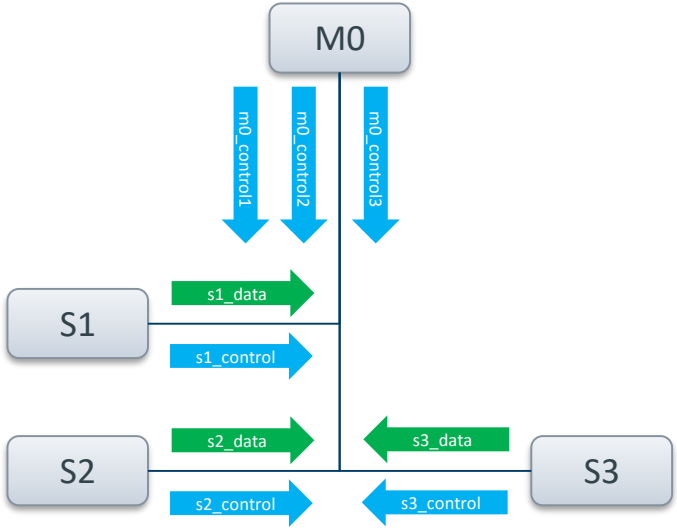

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Microphone array: worst case latency analysis result

Stream	Characteristics	Worst Case Latency	Test
s1_data, s2_data, s3_data	1 x 312 Byte every 1,33 ms	13,22 ms	FAIL
s1_control_m0, s2_control_m0, s3_control_m0	5 x 1500 Byte every 500 ms	89,25 ms	PASS
m0_control_s1, m0_control_s2, m0_control_s3	5 x 1500 Byte every 500 ms	28,87 ms	PASS

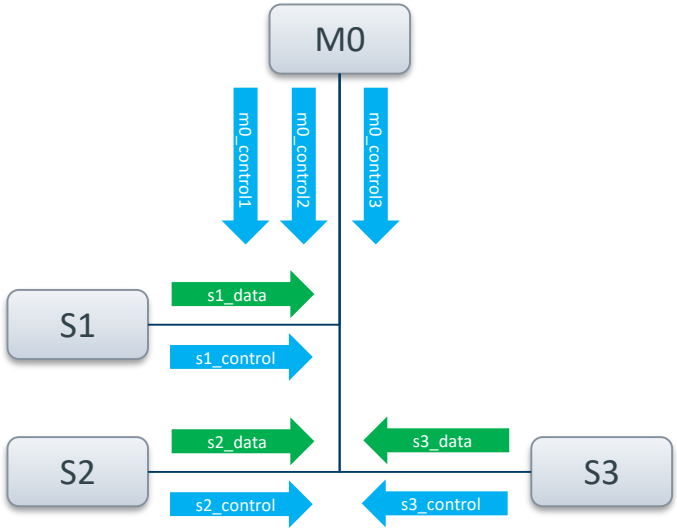

- Worst case latency caused by the PLCA arbitration is about 1 magnitude to high. (Unfeasible Solution)
- Reason: PLCA “allows” low priority TCP traffic to “interrupt” high priority data streams.

Assumptions:
- bitrate = 10 Mbit/s
- bus idle = 20 bit
- Overhead per Frame = 240 bit
- Overhead per Frame = 240 bit
- Priority of data stream > priority of control stream
- Queue weights = 1

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Proposal for improvement: Extend the PLCA mechanism to consider frame priority

Basic idea:

- A node listens to all frames of all other nodes and remembers the priority information (802.1Q) of the most recently sent frame for every node
- A node just uses its PLCA slot to send a frame if it would send a frame of at least the same priority like the recently sent frame of the other nodes. Otherwise it skips its slot on purpose.

Effect:

- Frames of high priority interrupt Streams of low priority which means they are transferred earlier (better worst case latencies)

Alexander Meier
alexander.meier@volkswagen.de

Priority based PLCA: How does it work?

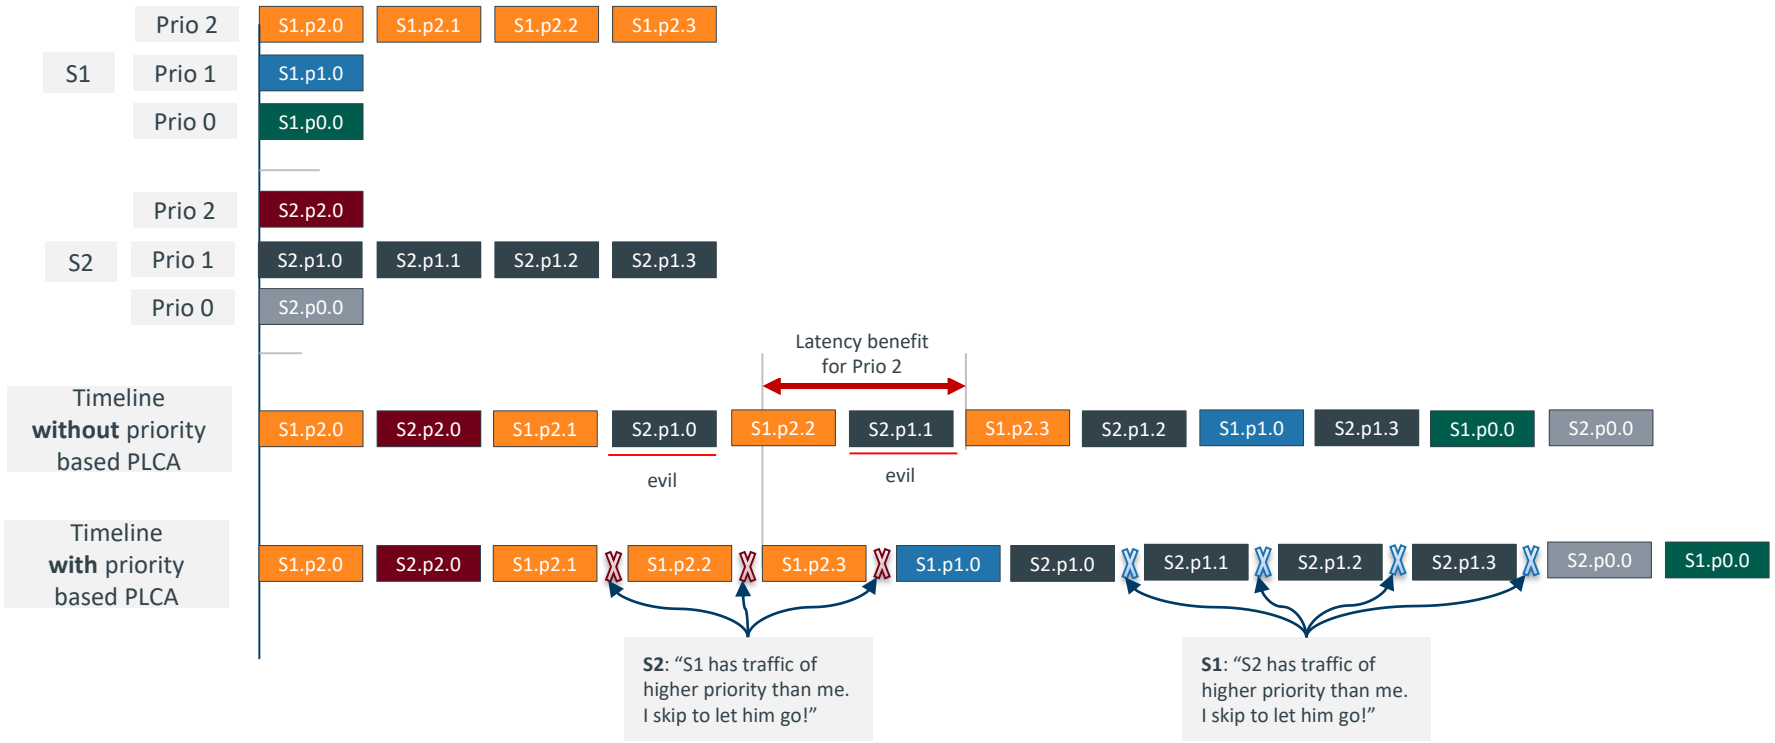

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Microphone array: worst case latencies with priority based PLCA

Stream	Characteristics	Worst Case Latency	Test
s1_data, s2_data, s3_data	1 x 312 Byte every 1,33 ms	3,01 ms*	FAIL
s1_control_m0, s2_control_m0, s3_control_m0	5 x 1500 Byte every 500 ms	90,32 ms	PASS
m0_control_s1, m0_control_s2, m0_control_s3	5 x 1500 Byte every 500 ms	46,59 ms	PASS

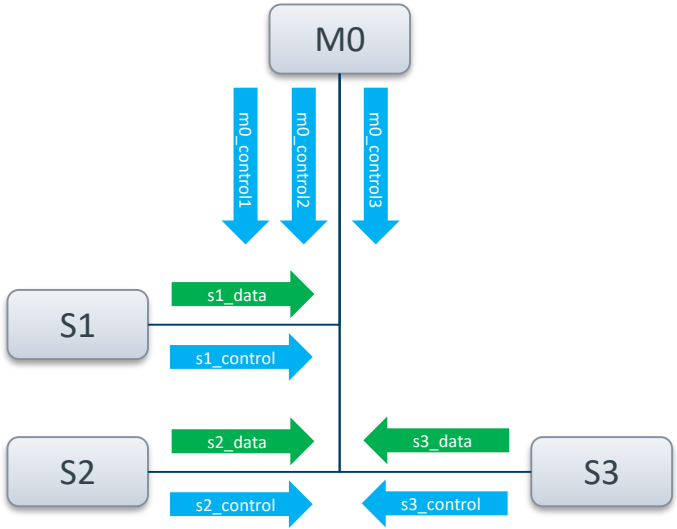

- Much better worst case latency for high priority data streams, but still FAIL

Assumptions:
- bitrate = 10 Mbit/s
- bus idle = 20 bit
- Overhead per Frame = 240 bit
- Overhead per Frame = 240 bit
- Priority of data stream > priority of control stream
- Queue weights = 1

*value depends on the schedule order of the PLCA (worst is shown)

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Proposal for further improvement: Introduce a very efficient low level segmentation mechanism

Problem:

- Huge frames in low speed networks lock the bus for a long time (about 1,2 ms for 1500 Byte frames)
- almost automatically given if you use TCP:
  - E.g. 8 nodes, each with at least 1 TCP stream → 8 * 1,2ms = 9,6 ms as “base worst case latency”

Basic idea:

- Cut frames into small pieces and reassemble them on receiver side
  - e.g. 64 Byte segments with small Overhead per segment

Effect:

- Low priority frames can be interrupted very early to prefer high priority streams

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Microphone array: worst case latencies with priority based PLCA and low level segmentation

Stream	Characteristics	Worst Case Latency	Test
s1_data, s2_data, s3_data	6 x 66 Byte** every 1,33 ms	1,04 ms*	PASS
s1_control_m0, s2_control_m0, s3_control_m0	120 x 66 Byte** every 500 ms	119,83 ms*	PASS
m0_control_s1, m0_control_s2, m0_control_s3	120 x 66 Byte** every 500 ms	163,89 ms	PASS

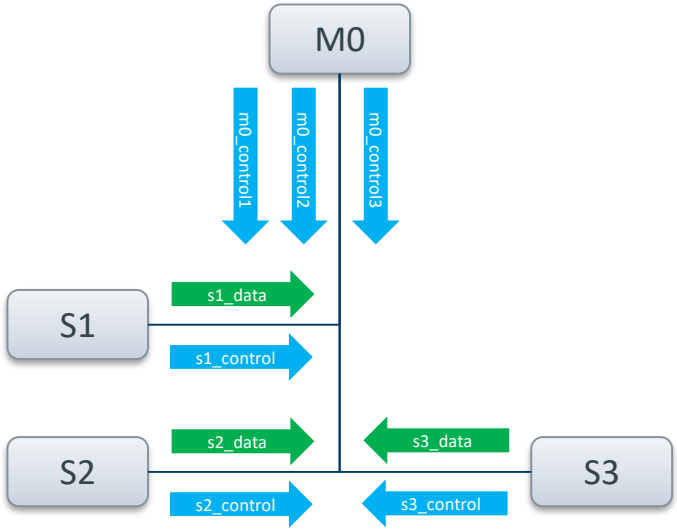

- Very good worst case latencies (1st setup that meets the requirements)

Assumptions:
- bitrate = 10 Mbit/s
- bus idle = 20 bit
- Overhead per Frame = 240 bit
- Overhead per Frame = 240 bit
- Priority of data stream > priority of control stream
- Queue weights = 1

*value depends on the schedule order of the PLCA (worst is shown)

**with padding to full segments

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)

Summary

- PLCA, due to its weighted round robin characteristics, has bad worst case latencies
- PLCA needs to be extended to consider frame priorities to meet automotive requirements and use cases (worst case latencies)
  - The need for this scales with the number of nodes

- Huge frames (1500 Byte) in low speed networks (10 Mbit Eth) ruin worst case latencies
- Introduction of an efficient low level segmentation mechanism is strongly recommended to meet automotive requirements and use cases

Alexander Meier
alexander.meier@volkswagen.de
17.01.2018 (V4)