MAC-xMII-PHY requirements Why latency and latency variation matters!

V03

Günter Steindl
Siemens AG

2024-12-16

Markets and Topologies

- Process automation
  - Often tree topology

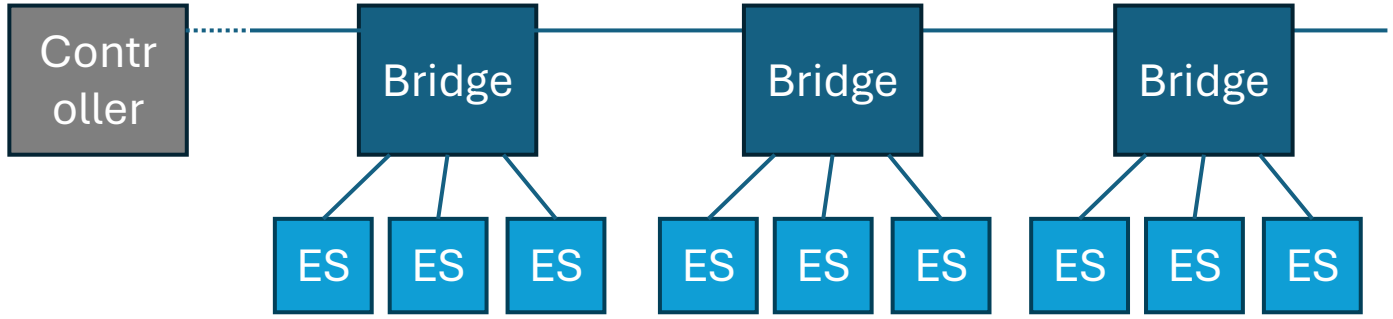

- Factory automation
  - Often linear or ring topology

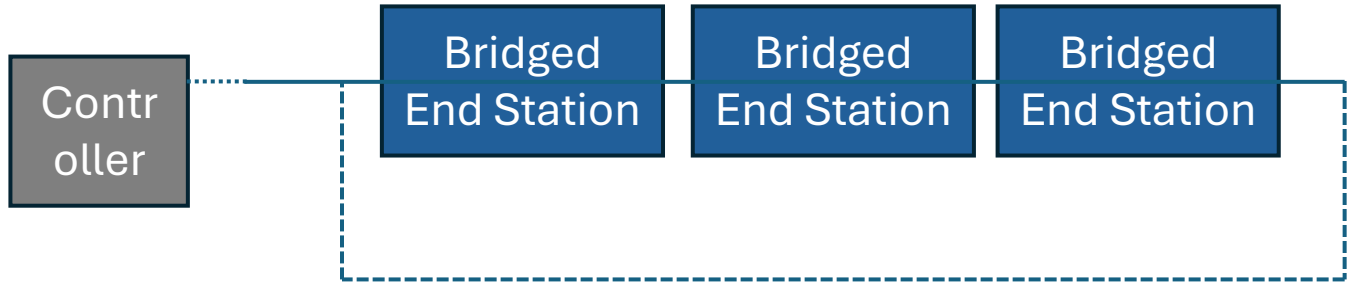

- Machine automation
  - Often linear or ring topology

ES := End station

Performance impact of PHY delays

- Tree topology
  - The PHY delay of an end station **only applies once** in the communication between end station and controller

- Linear or ring topology
  - The PHY delay of a bridged end station **applies for each hop** in the communication between end station and controller

  - Example: Linear topology with 64 bridged end stations connect to a controller

  ⇒The frame transmitted to the last end station is charged with a delay of

  ⇒63 hops, each with two PHY delays (receive and transmit)

Expected PHY delays
An example

- IEC61158 / IEC61784 family of fieldbus standards
  - Example “IEC61158 Type 10 – PROFINET”
    - The expected PHY delays (Table 88) are specified in the IEC61784-2-3:2023

Link Speed	RX delay [a]	TX delay [a]	Jitter
10 Mbit/s – 1 Tbit/s	Mandatory: < 1 µs Recommended:< 500 ns	Mandatory: < 1 µs Recommended:< 500 ns	< 4 ns
10 Mbit/s **Special case: 10BaseT1L**	Mandatory: < 5 µs [b] Recommended:< 500 ns	Mandatory: < 5 µs Recommended:< 500 ns	< 4 ns

[a] If IEEE 802.3 defines lower values, then these definitions apply. Lower values mean lesser latencies.

[b] 5 µs are equivalent to 6,25 octets at 10 Mbit/s.

NOTE – Latency variation influences the synchronization quality!

Customer expectations – for stations

- Existing machine and automation cell designs are not PHY dependent
- Introducing a new PHY for copper should not require changes in the machine and automation cell
- Roundtrip delays for motion or other high-speed applications shall still be achievable without new machine designs

Influencing factors - Model

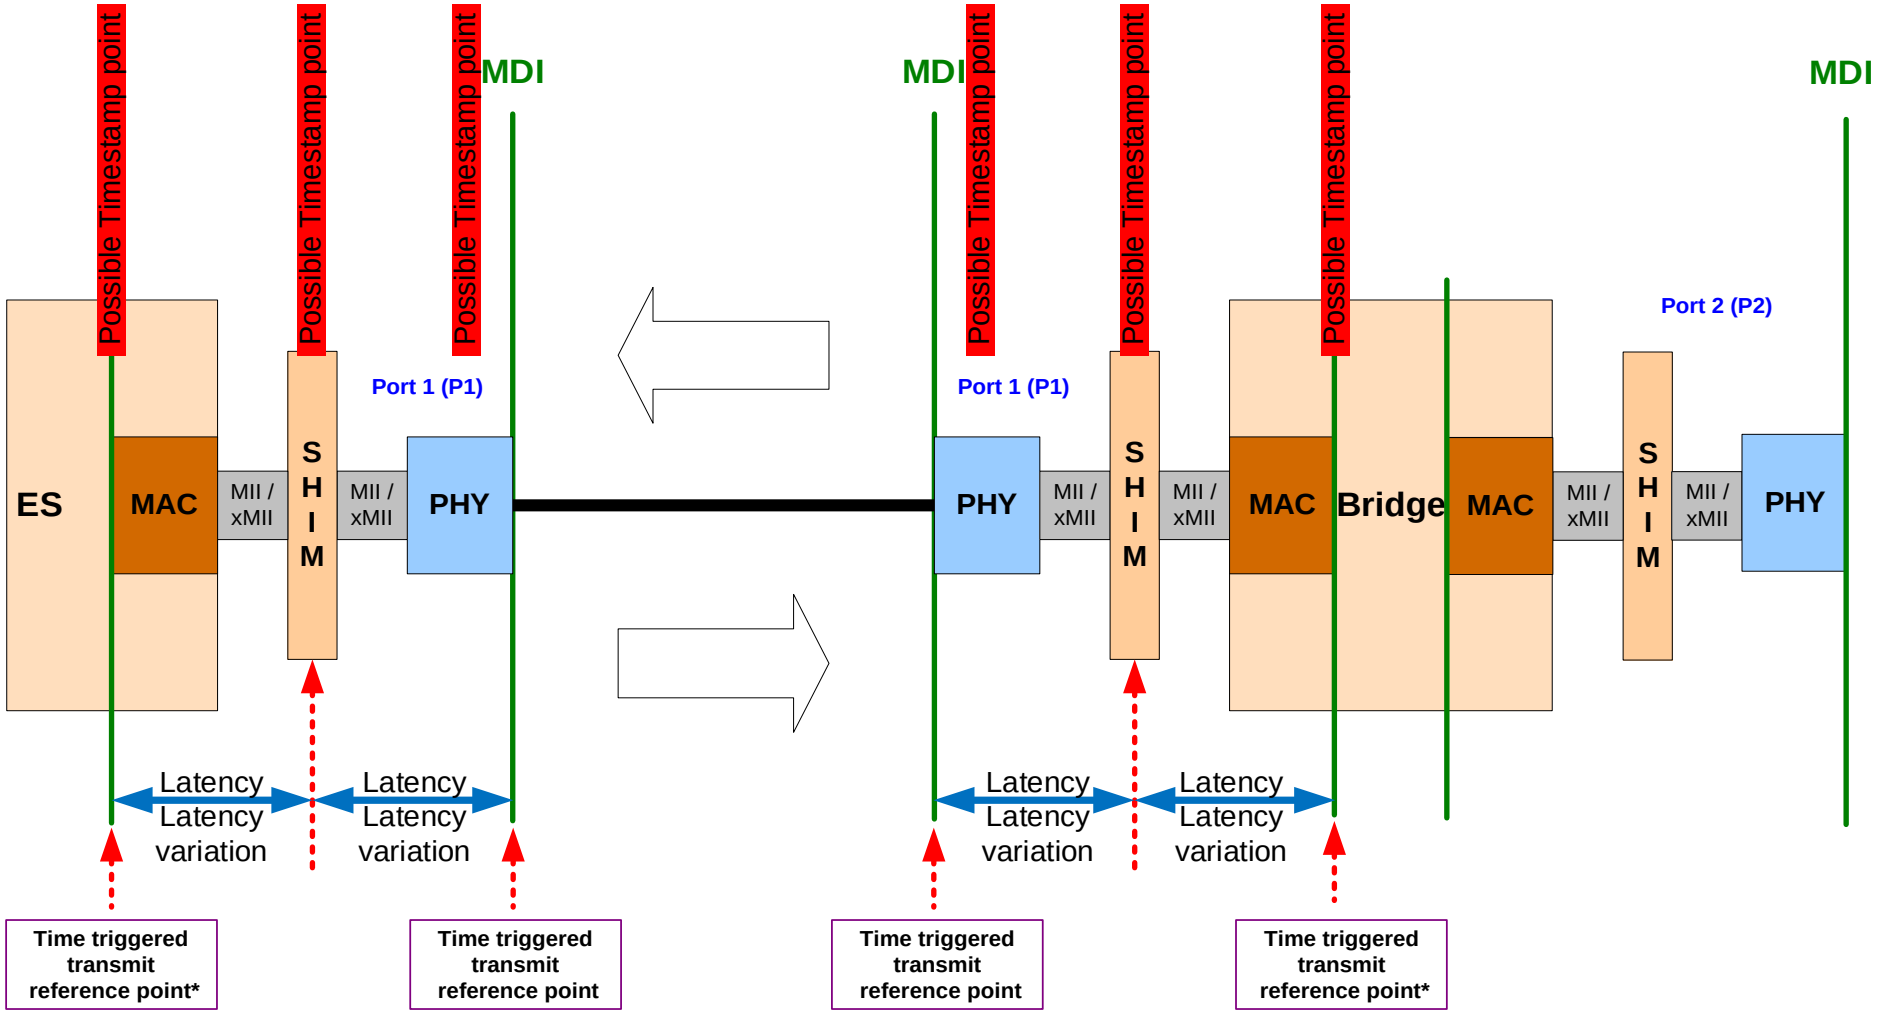

Time triggered transmit

- The configured time for time triggered transmit is based on time reference point at the MDI
- Time triggered transmit is execute at the MAC level (EM_UNITDATA.request / M_UNITDATA.request)
- Latency of MAC, xMII and PHY can be compensated (call xM_UNITDATA.request earlier)
- Latency variation of MAC, xMII and PHY can't be compensated at the station
- If combined with gate control, only engineering can compensate latency variation by adding safety margins to the size of the gates
- Thus, an increase in latency variation leads to a decrease in performance

=> Limit the allowed latency variation to reasonable values

Latency between before the MAC and MDI

- Latency variation of MAC, xMII and PHY can’t be compensated at the station
- The engineering may cover the latency variation with safety margins, e.g. for the size of the gates
- The latency variation of MAC, xMII and PHY may be more than cumulative due to different clock zones
  - Latency variation is a platform value and may differ for each link speed
  - Adding fiber transceivers increases the Latency variation, too
- Changes in latency and latency variation with every linkup shall be avoided

Values

Linkspeed	MII	RX (MAC, PHY) (Limits)	TX (MAC, PHY) (Limits)	Copper to Fiber (Limits)
10Mbit/s	MII, …	Latency: < 5µs Variation: < 1µs	Latency: < 5µs Variation: < 1µs	tbd
100Mbit/s*	MII, RMII, RGMII, …	Latency: < 1µs/500ns Variation: < 4ns	Latency: < 1µs/500ns Variation: < 4ns	tbd
1Gbit/s	GMII, RGMII, …	Latency: < 1µs/500ns Variation: < 4ns	Latency: < 1µs/500ns Variation: < 4ns	tbd
2,5Gbit/s	RGMII, XGMII, …	Latency: < 1µs/500ns Variation: < 4ns	Latency: < 1µs/500ns Variation: < 4ns	tbd
5Gbit/s	RGMII, XGMII, …	Latency: < 1µs/500ns Variation: < 4ns	Latency: < 1µs/500ns Variation: < 4ns	tbd
10Gbit/s	XGMII, …	Latency: < 1µs/500ns Variation: < 4ns	Latency: < 1µs/500ns Variation: < 4ns	tbd

*As an example – todays selected hardware: MII; RX: Latency ~200ns, Variation <4ns; TX: Latency ~100ns, Variation <4ns; Fiber: Latency ~8ns, Variation <4ns;

Influence of latency and latency variation

- Application / General
  - The sum of the PHY latencies, both transmit and receive, adds to the roundtrip delays
  - At 1Gbit/s with Cut Through (bridge delay ~ 1µs without PHYs) support, it starts to be the leading factor
  - The time from memory to MDI and vice versa reduces the available compute time

- End station
  - Fitting frames from a traffic class into a window of the gating cycle requires safety margins
  - Same for PSFP on receive
  - Latency variations that are not nullified by the timestamp negatively affect synchronization

- Bridge
  - Fitting frames from a traffic class into a window of the gating cycle requires safety margins
  - Same for PSFP on receive
  - Latency variations that are not nullified by the timestamp negatively affect synchronization

=> Limit the allowed latency and latency variation to reasonable values

Possible ways forward

- We should provide some information for the silicon vendors that high latency and high latency variation values are bad
- We should provide upper limits to avoid accidental development of almost unusable silicon
- We shall provide the real (latency and latency variation) values in the digital data sheet
  - 60802 YANG module providing per port and MAU type (or link speed) the minimum and maximum values for latency (offline and if different at runtime)

Thanks

Questions