A Joint Proposal for PAM2 Training PHY Frame and InfoField

January 14, 2019

Mike Tu tum@broadcom.com
Tom Souvignier tom.souvignier@broadcom.com
Peter Wu xingwu@marvell.com
Brett McClellan bmc@marvell.com

Proposed Baseline

- Following 1000BASE-T1 approach
- One data mode super-frame = 4 RS-FEC frames for 10G/5G/2.5G
- Select PAM2 training PHY frame = 4 RS-FEC frames
  - Frame boundary aligned with data mode super-frame
  - 7200 PAM2 symbols per frame
  - Partitioned into 16 partial frames (PF)
  - 96-bit InfoField at the start of the 16th partial frame
  - First bit of each partial frame is inverted as synchronization markers
- 1000BASE-T1 style state transitions
  - Each side maintains a partial frame counter (PFC) which is embedded in InfoField
  - Tx switching from PAM2 (SEND_T) to PAM4 (SEND_N) is based on announced PFC values

PAM2 PHY Frame Alignment with Data Mode Super-Frame

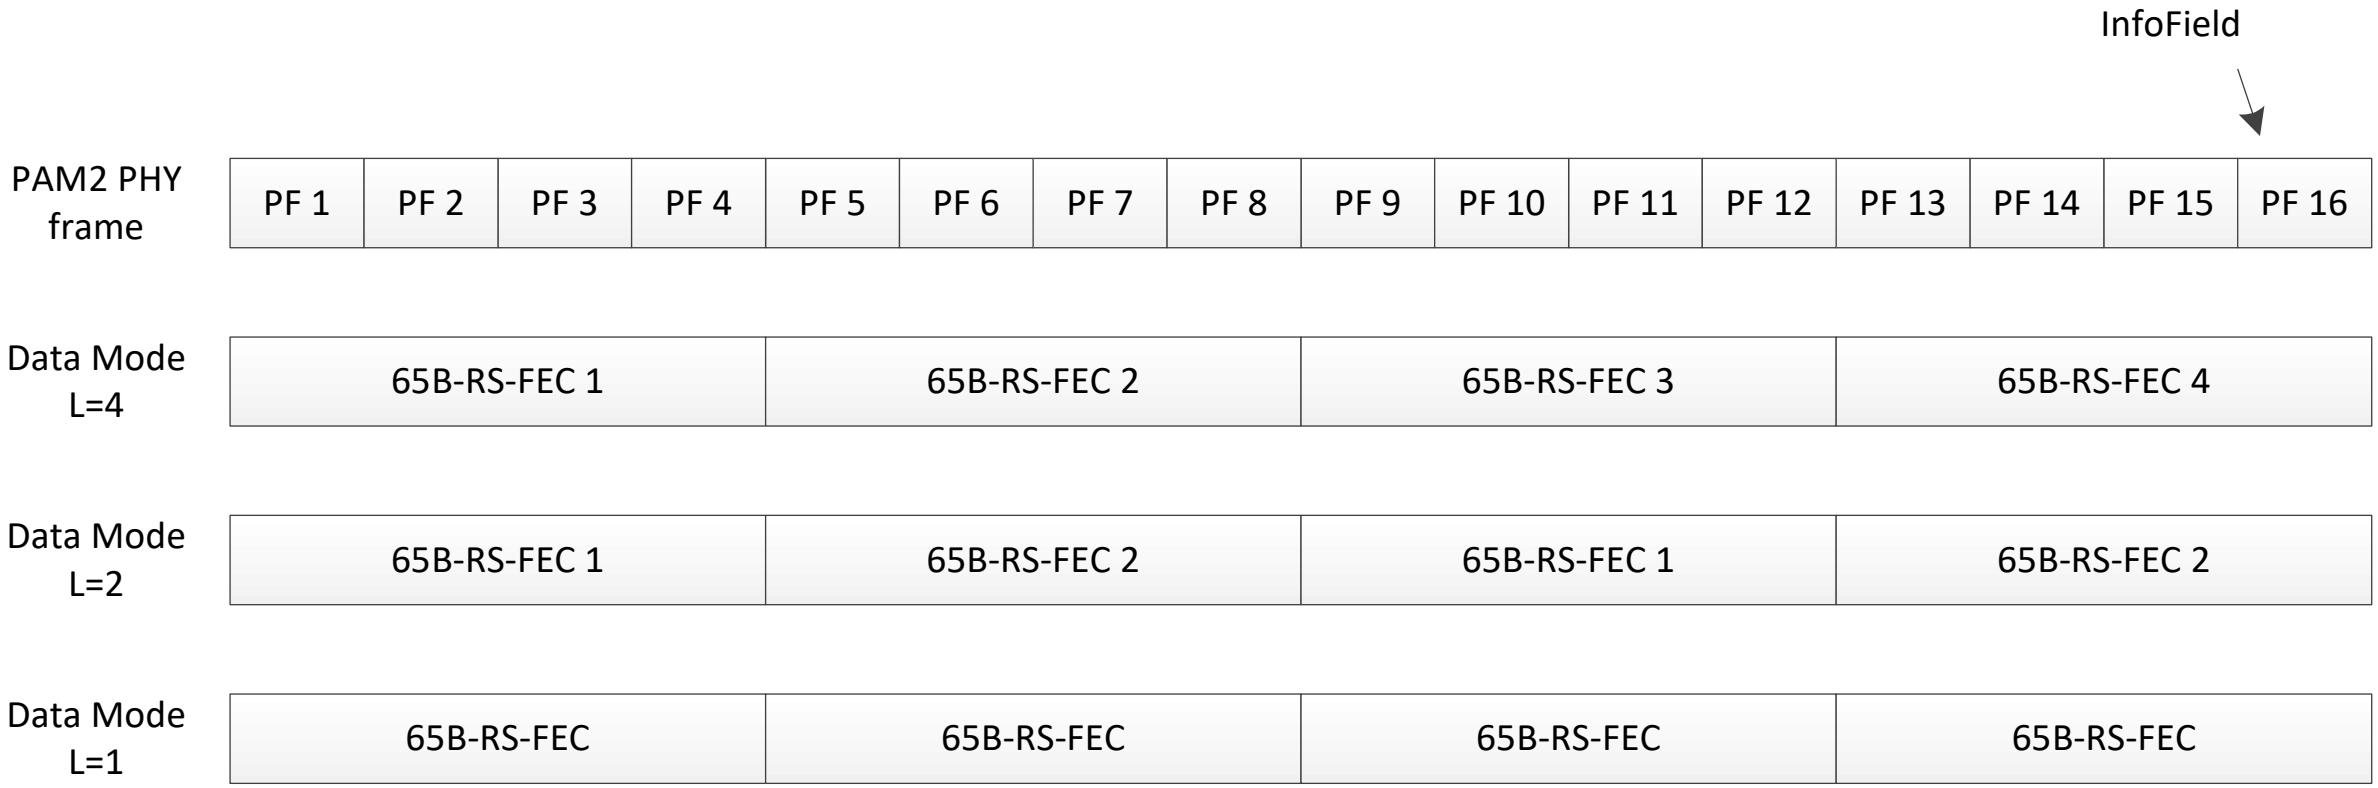

- Note: one data mode super-frame = 4 RS-FEC frames for 10G/5G/2.5G, regardless of interleaving depth.

Tx/Rx Frame Alignment Considerations

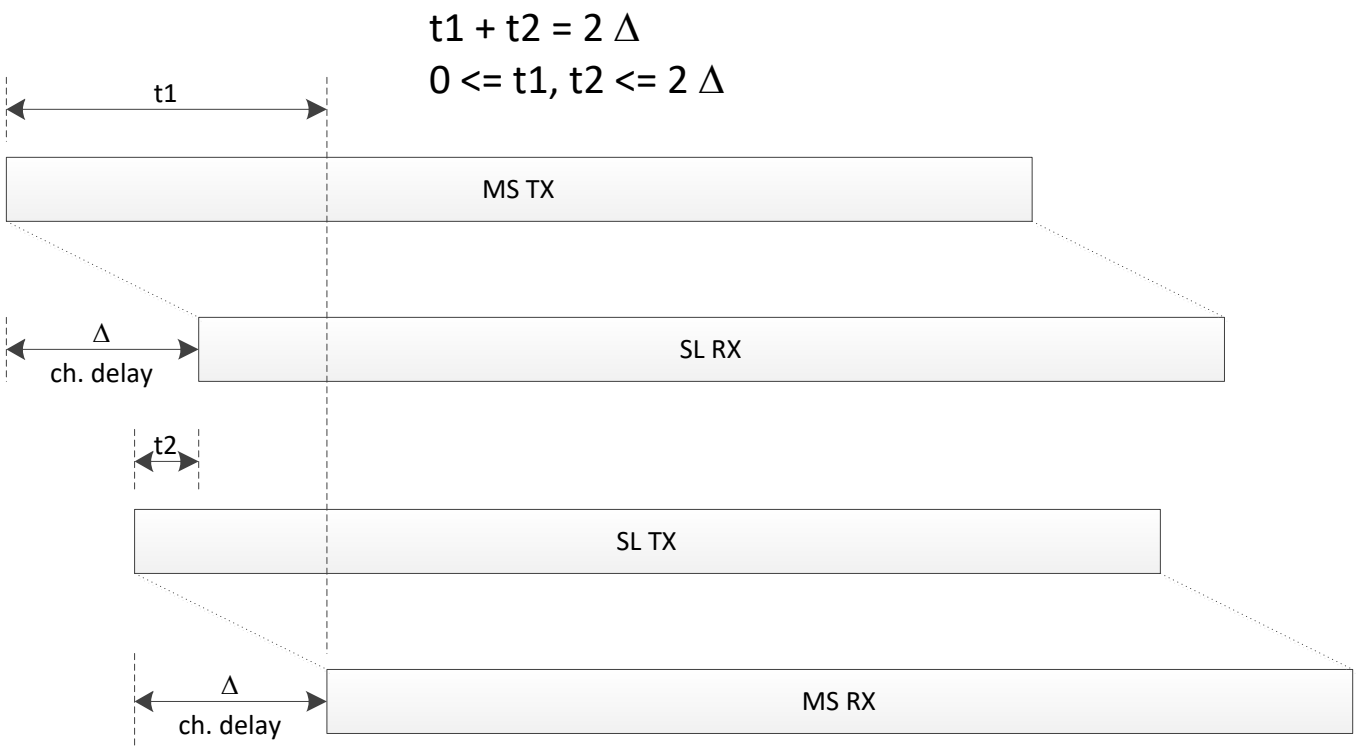

- 15m cable channel delay Δ ~= 75 nsec
- 1 partial frame
  - 320 nsec for 2.5G
  - 160 nsec for 5G
  - 80 nsec for 10G
- MS and SL Tx/Rx frame alignment requirement
  - 320 nsec
  - <=1 partial frame for 2.5G
  - <=2 partial frames for 5G
  - <=4 partial frames for 10G
  - Note: for 1000BASE-T1, this requirement is 240 nsec

Comparison of PHY Frame and Partial Frame Sizes

PAM2 PHY Frame	10GBASE-T1	5GBASE-T1	2.5GBASE-T1	1000BASE-T1
Frame duration	1.28 usec	2.56 usec	5.12 usec	3.6 usec
# RS FEC frames	4	4	4	1
# PAM2 symbols	7200	7200	7200	2700
# partial frames	16	16	16	15
# bits/PF	450	450	450	180
PF duration	80 nsec	160 nsec	320 nsec	240 nsec
Tx/Rx mis-alignment limit	4 PF 320 nsec	2 PF 320 nsec	1 PF 320 nsec	1 PF 240 nsec

InfoField General Format

octet 1	octet 2	octet 3	octet 4/5/6	octet 7	octet 8/9/10	octet 11/12
0xBB	0xA7	0x00	PFC24	Message	MSG24	CRC16

- 96 bits InfoField embedded at the start of the 16th partial frame
- Octet 1,2,3 are fixed header = 0xBBA700.
- Octet 4,5,6 PFC24 is the local 24-bit partial frame counter, counting the number of partial frames transmitted.
  - PFC24 will be increased by 16 after each PAM2 PHY frame.
  - The valid value of PFC24 is 16n – 1, where n indicates the sequence number of current PAM2 PHY frames.
- Octet 7 is the Message field. It includes current PMA state and PHY status. For MS and SL Message fields, see the following slides.
- Octet 8,9,10 MSG24 contents depend on the current PMA state.
- Octet 11 and 12 contains CRC16 for the InfoField.

Octet 7 Message Field and PMA States

- Octet 7<7:6> indicates the PMA states
- There are two states: “Training” and “Count Down”
- Training state
  - Indicate local PHY status
  - Indicate local Rx interleaving depth: L=1 for 2.5G, L=1,2 for 5G, L=1,2,4 for 10G
  - Indicate local Rx precoder selection: bypass, 1+D, 1-D, $1-D^2$
  - Exchange OAM and EEE capability
    - Enabled only if both sides advertise the capability
- Count Down state
  - Indicate the exact time when local Tx will switch from PAM2 to PAM4 data mode
  - No need to exchange data mode scrambler seeds, as the same deg=33 training mode scramblers will continue to run into data mode

Valid Message Field Settings

InfoField message field valid MASTER settings						
PMA_state <7:6>	loc_rcvr_status <5>	en_slave_tx <4>	reserved <3>	reserved <2>	reserved <1>	reserved <0>
00 (training)	0	0	0	0	0	0
00 (training)	0	1	0	0	0	0
00 (training)	1	1	0	0	0	0
01 (count down)	1	1	0	0	0	0

InfoField message field valid SLAVE settings						
PMA_state <7:6>	loc_rcvr_status <5>	timing_lock_OK <4>	reserved <3>	reserved <2>	reserved <1>	reserved <0>
00 (training)	0	0	0	0	0	0
00 (training)	0	1	0	0	0	0
00 (training)	1	1	0	0	0	0
01 (count down)	1	1	0	0	0	0

Octet 8/9/10 in Training State (octet 7<7:6>=00)

octet 1	octet 2	octet 3	octet 4/5/6	octet 7	octet 8/9/10	octet 11/12
0xBB	0xA7	0x00	PFC24	Message	UsrCfgCap	CRC16

- User configurations and capabilities
- EEE enabled only if both PHYs set EEEen = 1.
- OAM enabled only if both PHYs set OAMen = 1.
- InterleaveDepth
  - Oct10<2:1> = 00 ⇔ L=1
  - Oct10<2:1> = 01 ⇔ L=2
  - Oct10<2:1> = 10 ⇔ L=4
  - Oct10<2:1> = 11 ⇔ reserved
  - Tx must support the requested interleaving depth
- PrecodeSel
  - Oct10<4:3> = 00 ⇔ precoder bypass
  - Oct10<4:3> = 01 ⇔ 1-D
  - Oct10<4:3> = 10 ⇔ 1+D
  - Oct10<4:3> = 11 ⇔ $1-D^2$
  - Tx must support the selected precoder

octet 8								octet 9								octet 10							
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
Reserved															EEEen	OAMen	InterleaveDepth			PrecodeSel		Reserved	Reserved	Reserved

Octet 8/9/10 in Count Down State (octet 7<7:6>=01)

octet 1	octet 2	octet 3	octet 4/5/6	octet 7	octet 8/9/10	octet 11/12
0xBB	0xA7	0x00	PFC24	Message	DataSwPFC24	CRC16

- The PHY shall wait for the link partner to indicate loc_rcvr_status=1, before entering the Count Down state
- DataSwPFC24
  - Select a time in the future for the local Tx to switch from PAM2 to PAM4, including:
    - tx_mode = SEND_N
    - Interleaver enabled
    - Tx precoder enabled
  - The transition time is represented as a 24-bit partial frame count value
  - Align at PAM2 PHY frame boundary ➔ DataSwPFC24 shall be an integer multiple of 16

Generating PAM2 Training PHY Frame

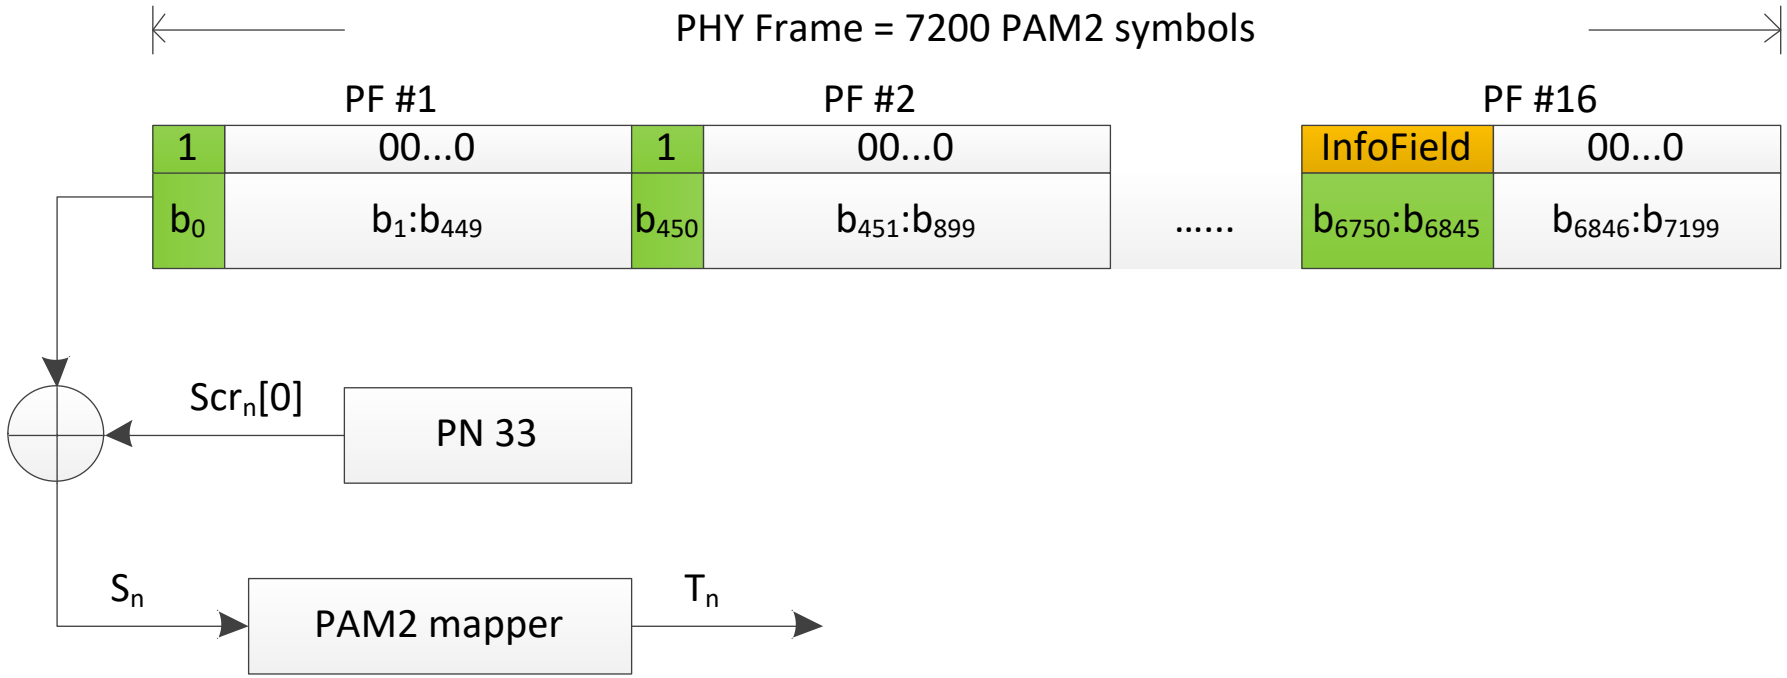

- First bit of each partial frame is inverted, for PF1 to PF15.
- First 96 bits of PF16 are the InfoField.
- The rest of the PHY frame bits are 0.
- PHY frame is XOR'ed with PN33 scrambler output $Scr_n[0]$, and then mapped into PAM2 symbols.

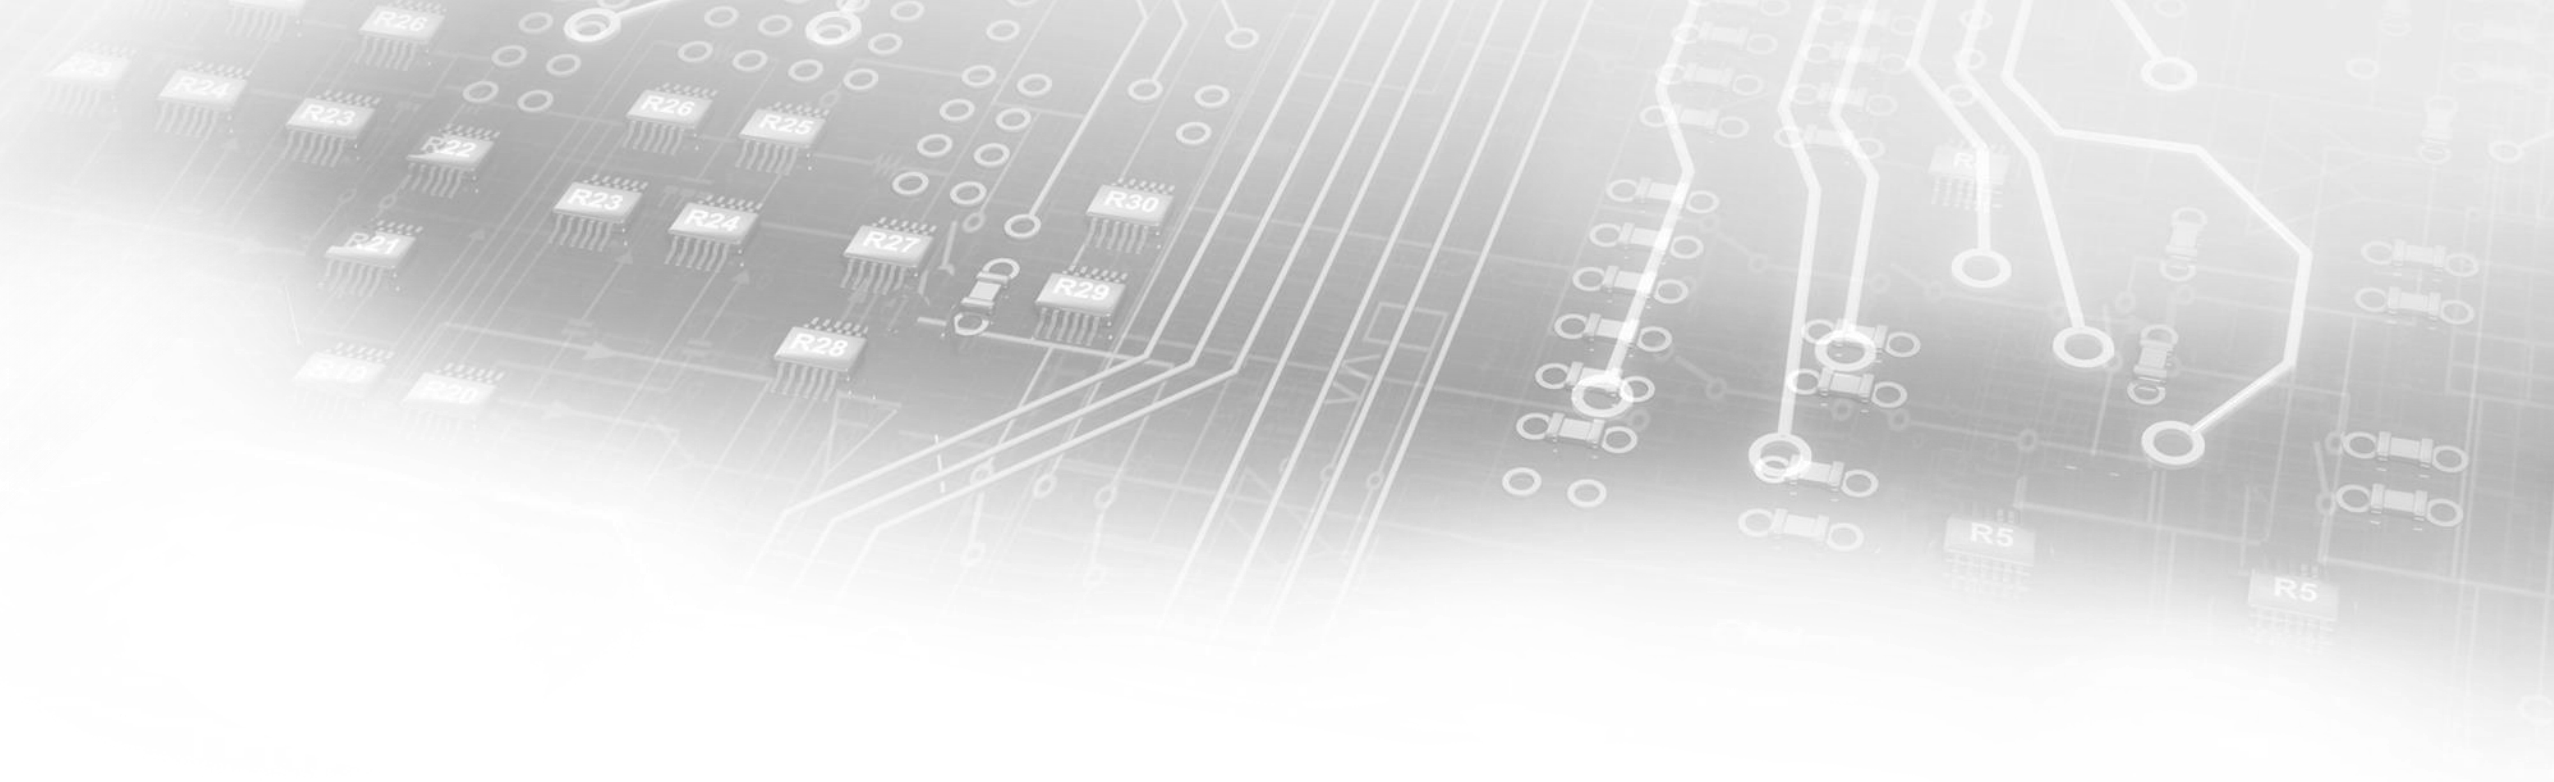

THANK YOU