112Gbps LR COM Investigation (III)

Mike Li, Hsinho Wu, Masashi Shimanouchi, Adee Ran

For IEEE 802.3ck

Jan, 2019

Intel

Recap of the Basics Assumptions Used In Previous Investigations

- Investigation (I):http://www.ieee802.org/3/ck/public/18_09/li_3ck_02_0918.pdf
- Investigation (II): http://www.ieee802.org/3/ck/public/18_11/li_3ck_02a_1118.pdf
- Most of 100GEL/112Gbps LR Rx will be ADC-based designs, therefore
  - FFE-heavy/DFE(1-tap) (Config 0) should be the LR ref. receiver benchmark
    - TX:
      - Scale rise/fall time (T_r), jitter (A_DD, sigma_RJ)
      - Extended pre-tap 1 and post-tap 1 range to -0.3
    - RX:
      - eta_0 (8.2e-9 V^2/GHz),
      - CTLE: f_z, f_p1 (both scaled), f_p2 (56GHz), f_HP_PZ (0.7GHz)
      - Added FFE (3 pre-taps + 12 post-taps)
      - 1-tap DFE
  - Due to feedbacks and considerations in design complexity, IP protection, and complications in standardization, we explored 2 ref. Rx options.

Summary of Findings in Nov/2018 802.3 ck Presentation

(http://www.ieee802.org/3/ck/public/18_11/li_3ck_02a_1118.pdf)

- In the Nov/2019 802.3ck meeting, we proposed and examined 2 COM LR reference Rx
  - FFE-lite/DFE-heavy (Config 2)
    - Rx has FFE for pre-taps and DFE for post-taps
  - DFE-only (Config 1)
    - RX only has DFE
  - **Goal is to develop simple COM method yet provide reasonable performance**

Config #	Baseline Ref. RX	Performance (w.r.t. *Config 0*)	Complexity	Notes
0	FFE-heavy (3-pre/n-post)/DFE (1-tap)	**High**	**High**	• Good performance • Complex COM model and standardization
1	DFE-only	**Low**	**Low**	• Low performance • Uses existing COM methodology
2	FFE-lite (3-pre/0-post)/DFE-heavy (n-taps)	**High**	**Low/Medium**	• Good Performance • Use existing 802.3/OIF-CEI and COM methodology

Scope of This Work

- Recent 802.3ck discussions raised some concerns about the FFE-lite/DFE-heavy and DFE-only options
  - [1] Y. Hidaka and P. Sun, “COM Simulation for 100G KR/CR Channels” (http://www.ieee802.org/3/ck/public/adhoc/dec05_18/hidaka_3ck_adhoc_01_120518.pdf)
  - [2] P. Sun and Y. Hidaka, “Comparison of KR/CR Reference Receivers” (http://www.ieee802.org/3/ck/public/adhoc/dec05_18/sun_3ck_adhoc_01a_120518.pdf)
  - [3] Y. Lu, Z. Huang, and Y. Zhuang, “Investigation of Baseline Reference Receivers”, (http://www.ieee802.org/3/ck/public/adhoc/dec12_18/lu_3ck_adhoc_01a_121218.pdf)
  - [4] Y. Lu, Z. Huang, and Y. Zhuang, “*More Insights of IEEE 802.3ck Baseline Reference Receivers*”, (http://www.ieee802.org/3/ck/public/adhoc/dec12_18/lu_3ck_adhoc_01a_121218.pdf)
- Scope and Tasks
  - Simulate with more 802.3ck channels
  - Confirm/check observations from [1][2][3][4]
  - Study the risk and feasibility of enhanced DFE-only option (Config. 1)
  - Modify FFE-lite/DFE-heavy option (Config. 2) to address performance concerns

100GEL/112Gbps LR COM with DFE-only (Config 1) and FFE-lite/DFE-heavy (Config 2) options

- Baseline
  - 802.3cd COM
    - http://www.ieee802.org/3/cd/public/adhoc/archive/mellitz_080217_3cd_01_adhoc.pdf
    - http://www.ieee802.org/3/cd/public/adhoc/archive/mellitz_3cd_0817_COM.zip
- Change Summary
  - TX
    - Scale: TX rise/fall time (T_r), jitter (A_DD, sigma_RJ)
    - TX EQ
      - Config 0 and Config 2: 2 pre-taps + 1 post-tap, 2.5% step size, pre-tap1 and post-tap1 min. = -0.3
      - Config 1: Same as above or 3 pre-taps + 1 post-tap
    - Same R_{LM}, SNR_{TX}
  - RX
    - RX input referred noise (eta_0): *8.2e-9 V^2/GHz*
    - Equalization
      - CTLE
        - » Scale *f_z, f_p1, f_p2*
        - » *f_p2 = 56 GHz (= baud rate)**
        - » *f_HP_PZ: 0.7 GHz (=f_b / 80)*
      - Config 1: DFE-only
        - DFE Configuration
          - *16 or 20* post-taps
          - DFE tap coef. Tap 1 ≤ *0.7 or 0.85*, others ≤ *0.2*
      - Config 2: FFE-lite/DFE-heavy
        - DFE Configuration
          - *16 or 20* post-taps
          - DFE tap coef. Tap 1 ≤ *0.7 or 0.6 or 0.5*, others ≤ *0.2*
        - FFE Configuration
          - 3 pre-taps
          - FFE tap coef: Main cursor: 1, Pre-tap 1: ≤*0.5*, other taps: ≤*0.125*
- Package / TX/RX Capacitance and Termination
  - Length: max *20 or 30mm*
  - T-Line: No change
  - Cd: *130fF*
  - Cp: *110fF*
  - Rd: *50 Ohms*

**: In 802.3cd COM, f_p2 is 2x baud rate (112GHz) which we believe it might be too high for today's CMOS technology. Further, f_r (RX noise filter) is set at 0.75x baud rate which will reduce the effectiveness of higher f_p2.*

Proposed 112Gbps LR COM Spreadsheet

Config 1: DFE-only

Table 93A-1 parameters			
Parameter	**Setting**	**Units**	**Information**
f_b	56	GBd	
f_min	0.05	GHz	
Delta_f	0.01	GHz	
C_d	[1.3e-4 1.3e-4]	nF	[TX RX]
z_p select	[1 2 3 4 5]		[test cases to run]
z_p (TX)	[12 20 30 20 30]	mm	[test cases]
z_p (NEXT)	[12 12 12 12 12]	mm	[test cases]
z_p (FEXT)	[12 20 20 30 30]	mm	[test cases]
z_p (RX)	[12 20 20 30 30]	mm	[test cases]
C_p	[1.1e-4 1.1e-4]	nF	[TX RX]
R_0	50	Ohm	
R_d	[50 50]	Ohm	[TX RX] or selected
f_r	0.75	*fb	
c(0)	0.6		min
c(-1)	[-0.3:0.025:0]		[min:step:max]
c(-2)	[0:0.025:0.1]		[min:step:max]
c(1)	[-0.3:0.025:0]		[min:step:max]
g_DC	[-20:1:0]	dB	[min:step:max]
g_DC_HP	[-6:1:0]		[min:step:max]
f_HP_PZ	0.7	GHz	
f_z	22.4	GHz	
f_p1	22.4	GHz	
f_p2	56	GHz	
A_v	0.41	V	tdr selected
A_fe	0.41	V	tdr selected
A_ne	0.6	V	tdr selected
L	4		
M	32		
N_b	16	UI	
b_max(1)	0.7		
b_max(2..N_b)	0.2		
sigma_RJ	0.01	UI	
A_DD	0.02	UI	
eta_0	8.20E-09	V^2/GHz	
SNR_TX	32.5	dB	tdr selected
R_LM	0.95		
DER_0	1.00E-04		
Operational control			
COM Pass threshold	3	dB	
Include PCB	0	Value	0, 1, 2

I/O control		
DIAGNOSTICS	0	logical
DISPLAY_WINDOW	0	logical
Display frequency domain	0	logical
CSV_REPORT	1	logical
RESULT_DIR	.\results\D1p2_{date}\	
SAVE_FIGURES	0	logical
Port Order	[1 3 2 4]	
RUNTAG	v165_d1p0a	
Receiver testing		
RX_CALIBRATION	0	logical
Sigma BBN step	5.00E-03	V
IDEAL_TX_TERM	0	logical
T_r	0.006160714	ns
FORCE_TR	1	logical

RX FFE		
ffe_enable	0	logical
ffe_pre_tap_len	3	UI
ffe_post_tap_len	0	UI
ffe_tap_step_size	0	UI
ffe_main_cursor_min	1	
ffe_pre_tap1_max	0.5	
ffe_post_tap1_max	0.5	
ffe_tapn_max	0.125	

Table 93A–3 parameters		
Parameter	**Setting**	**Units**
package_tl_gamma0_a1_a2	[0 1.734e-3 1.455e-4]	
package_tl_tau	6.141E-03	ns/mm
package_Z_c	90	Ohm (tdr sel)

Table 92–12 parameters		
Parameter	**Setting**	
board_tl_gamma0_a1_a2	[0 4.114e-4 2.547e-4]	
board_tl_tau	6.191E-03	ns/mm
board_Z_c	110	Ohm
z_bp (TX)	151	mm
z_bp (NEXT)	72	mm
z_bp (FEXT)	72	mm
z_bp (RX)	151	mm

Proposed 112Gbps LR COM Spreadsheet

Config 2: FFE-lite/DFE-heavy

Table 93A-1 parameters			
Parameter	**Setting**	**Units**	**Information**
f_b	56	GBd	
f_min	0.05	GHz	
Delta_f	0.01	GHz	
C_d	[1.3e-4 1.3e-4]	nF	[TX RX]
z_p select	[1 2 3 4 5]		[test cases to run]
z_p (TX)	[12 20 30 20 30]	Mm	[test cases]
z_p (NEXT)	[12 12 12 12 12]	Mm	[test cases]
z_p (FEXT)	[12 20 20 30 30]	mm	[test cases]
z_p (RX)	[12 20 20 30 30]	mm	[test cases]
C_p	[1.1e-4 1.1e-4]	nF	[TX RX]
R_0	50	Ohm	
R_d	[50 50]	Ohm	[TX RX] or selected
f_r	0.75	*fb	
c(0)	0.6		min
c(-1)	[-0.3:0.025:0]		[min:step:max]
c(-2)	[0:0.025:0.1]		[min:step:max]
c(1)	[-0.3:0.025:0]		[min:step:max]
g_DC	[-20:1:0]	dB	[min:step:max]
g_DC_HP	[-6:1:0]		[min:step:max]
f_HP_PZ	0.7	GHz	
f_z	22.4	GHz	
f_p1	22.4	GHz	
f_p2	56	GHz	
A_v	0.41	V	tdr selected
A_fe	0.41	V	tdr selected
A_ne	0.6	V	tdr selected
L	4		
M	32		
N_b	16	UI	
b_max(1)	0.5		
b_max(2..N_b)	0.2		
sigma_RJ	0.01	UI	
A_DD	0.02	UI	
eta_0	8.20E-09	V^2/GHz	
SNR_TX	32.5	dB	tdr selected
R_LM	0.95		
DER_0	1.00E-04		
Operational control			
COM Pass threshold	3	dB	
Include PCB	0	Value	0, 1, 2

I/O control		
DIAGNOSTICS	0	logical
DISPLAY_WINDOW	0	logical
Display frequency domain	0	logical
CSV_REPORT	1	logical
RESULT_DIR	.\results\D1p2_{date}\	
SAVE_FIGURES	0	logical
Port Order	[1 3 2 4]	
RUNTAG	v165_d1p0a	
Receiver testing		
RX_CALIBRATION	0	logical
Sigma BBN step	5.00E-03	V
IDEAL_TX_TERM	0	logical
T_r	0.006160714	ns
FORCE_TR	1	logical

RX FFE		
ffe_enable	1	logical
ffe_pre_tap_len	3	UI
ffe_post_tap_len	0	UI
ffe_tap_step_size	0	UI
ffe_main_cursor_min	1	
ffe_pre_tap1_max	0.5	
ffe_post_tap1_max	0.5	
ffe_tapn_max	0.125	

Table 93A–3 parameters		
Parameter	**Setting**	**Units**
package_tl_gamma0_a1_a2	[0 1.734e-3 1.455e-4]	
package_tl_tau	6.141E-03	ns/mm
package_Z_c	90	Ohm (tdr sel)

Table 92–12 parameters		
Parameter	**Setting**	
board_tl_gamma0_a1_a2	[0 4.114e-4 2.547e-4]	
board_tl_tau	6.191E-03	ns/mm
board_Z_c	110	Ohm
z_bp (TX)	151	mm
z_bp (NEXT)	72	mm
z_bp (FEXT)	72	mm
z_bp (RX)	151	mm

Channel Data

CH #	Description	Reference Document	Notes
1-3	Initial Host 30dB Backplane Channel Models	heck_100GEL_01_0118.pdf	
4-6	Best Case 3”, 13”, 18” Tachyon Backplane	mellitz_100GEL_adhoc_01_010318.pdf	
7-8	Synthesized CR Channels (2.0m and 2.5m 28AWG Cable)	mellitz_100GEL_adhoc_01_021218.pdf	
9-10	Orthogonal or Cabled Backplane Channels	tracy_100GEL_03_0118.pdf	
11-26	16/20/24/28dB Cabled Backplane Channels	heck_3ck_01_1118.pdf	
27-32	24/28/32dB Cabled Backplane Channels including Via	mellitz_3ck_adhoc_02_081518.pdf	
33-51	Measured Traditional Backplane Channels	kareti_3ck_01a_1118.pdf	
52-70	Measured Cabled Backplane Channels		
71-85	Measured Orthogonal Backplane Channels		

All channel data are from IEEE 802.3ck Task Force Tools & Channels page: http://www.ieee802.org/3/ck/public/tools/index.html

112Gbps Test Channel COM Results

w/20mm Packages w/o Crosstalk*: DFE-only (Config 1, 16 post-taps)

Config 1 (Original: TX pre-2,& 2.5% step size)

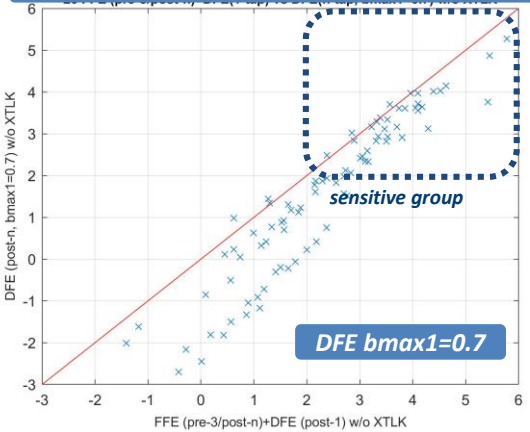

	Value	Note
Mean Diff	0.81	
Sensitive Group Mean Diff	0.53	
False Pass	1	1.18%
False Fail	10	11.76%

Original Config 1 is too pessimistic w.r.t. Config 0

Config 1 (w/ TX pre-3 & 2% step size)

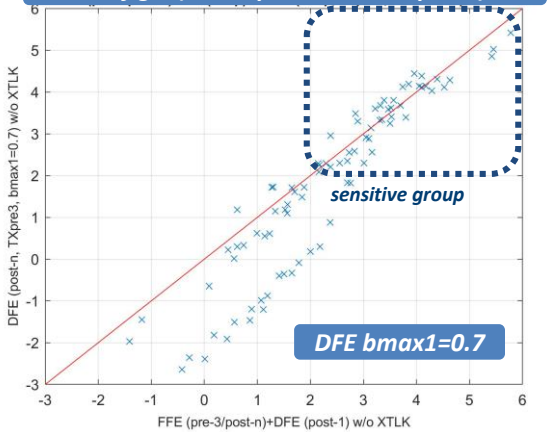

	Value	Note
Mean Diff	0.53	
Sensitive Group Mean Diff	0.13	
False Pass	2	2.35%
False Fail	4	4.71%

Config 1 (w/ TX pre-3 & 2% step size)

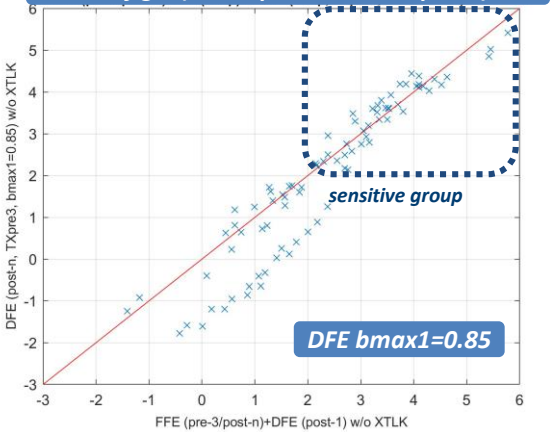

	Value	Note
Mean Diff	**0.30**	
Sensitive Group Mean Diff	**0.04**	
False Pass	**2**	2.35%
False Fail	**3**	3.53%

Config 1 w/ $b_{max}(1)$=0.85 and TX w/ 3 pre-taps matches better with Config 0

112Gbps Test Channel COM Results

w/ 20mm Packages w/o Crosstalk*: DFE-only (Config 1, 20 post-taps)

Config 1 (Original: TX pre-2,& 2.5% step size)

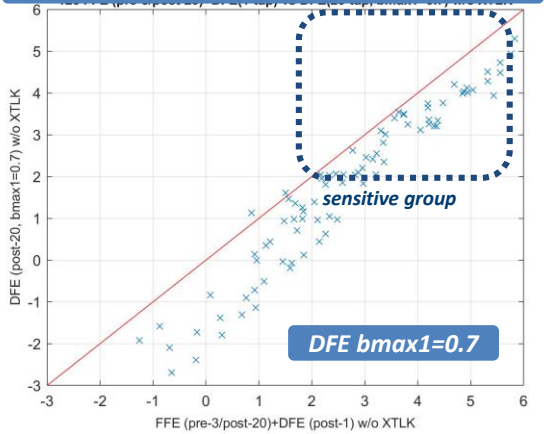

	Value	Note
Mean Diff	0.90	
Sensitive Group Mean Diff	0.76	
False Pass	0	0.00%
False Fail	6	7.06%

Original Config 1 is too pessimistic w.r.t. Config 0

Config 1 (w/ TX pre-3 & 2% step size)

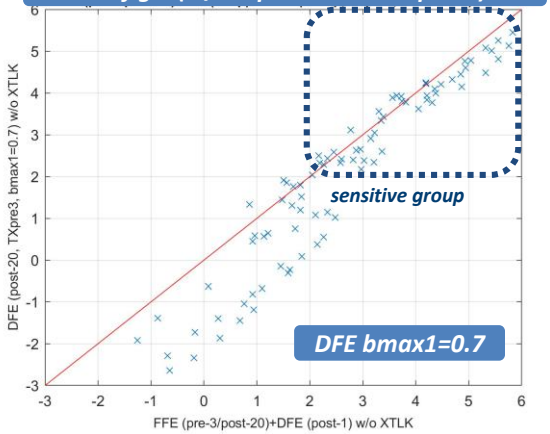

	Value	Note
Mean Diff	0.60	
Sensitive Group Mean Diff	0.34	
False Pass	1	1.18%
False Fail	4	4.71%

Config 1 (w/ TX pre-3 & 2% step size)

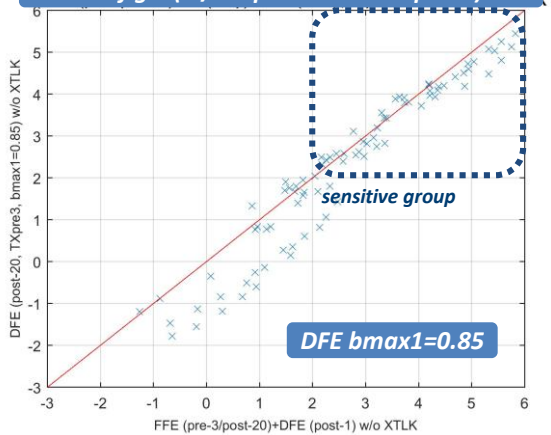

	Value	Note
Mean Diff	**0.37**	
Sensitive Group Mean Diff	**0.23**	
False Pass	**1**	1.18%
False Fail	**4**	4.71%

Config 1 w/ $b_{max}(1)$=0.85 and TX w/ 3 pre-taps matches better with Config 0

112Gbps Test Channel COM Results

w/ 20mm Packages w/o Crosstalk*: DFE-only (Config 1, 16/20 post-taps)

Config 1 (Original: TX pre-2,& 2.5% step size)

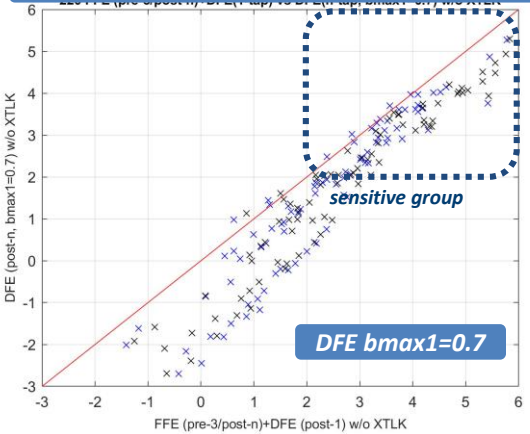

	Value	Note
Mean Diff	0.85	
Sensitive Group Mean Diff	0.65	
False Pass	1	0.59%
False Fail	16	9.41%

Original Config 1 is too pessimistic w.r.t. Config 0

Config 1 (w/ TX pre-3 & 2% step size)

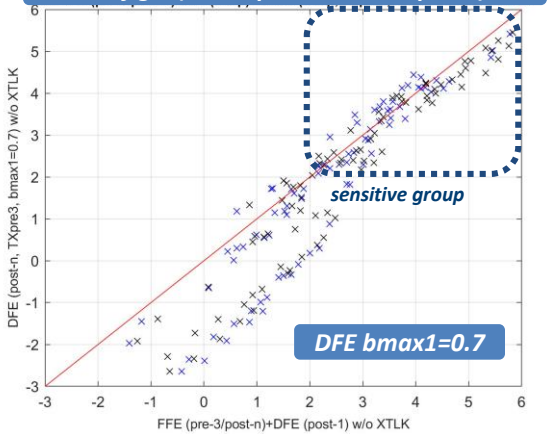

	Value	Note
Mean Diff	0.56	
Sensitive Group Mean Diff	0.24	
False Pass	3	1.76%
False Fail	8	4.71%

Config 1 (w/ TX pre-3 & 2% step size)

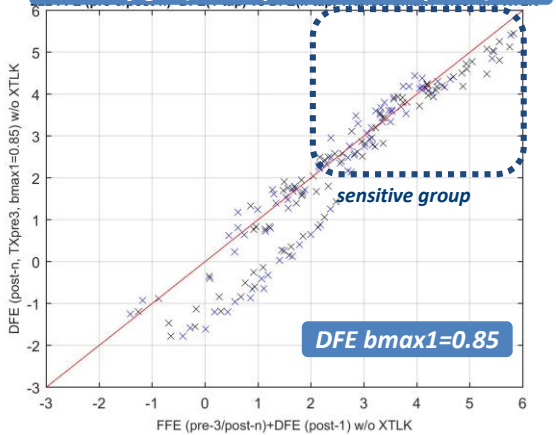

	Value	Note
Mean Diff	**0.33**	
Sensitive Group Mean Diff	**0.14**	
False Pass	**3**	1.76%
False Fail	**7**	4.12%

Config 1 w/ $b_{max}(1)$=0.85 and TX w/ 3 pre-taps matches better with Config 0

Concerns with DFE-only (Config 1) with Large DFE-Tap Limit

- Observed same trend in [1][2] that DFE tap 1 with $b_{max}(1)=0.85$ (Config 1) yields closer results with ADC-based design (Config 0)
- Concerns
  - Most of 100/112Gbps transceivers are ADC-based with FFE-heavy/DFE-lite (Config 0) configuration
    - e.g. RX FFE with 3 pre-taps and long post taps and 1-tap DFE
  - DFE with increased tap coefficient range, e.g. $b_{max}(1)$=0.85, is known to
    - Exacerbate burst errors and reduce FEC efficiency
    - Shift post-cursor equalization away from TX FIR and CTLE
      - » Due to "enhanced" DFE tap weight, CTLE output will likely be poorly equalized and, hence, make clock recovery more difficult
      - » In contrary to today's RX design and operating condition
  - TX with additional pre-tap
    - Adding additional burden to TX
      - » Mismatch with ADC-based design and increase COM computation time

112Gbps Test Channel COM Results

w/ 20mm Packages w/ Crosstalk: FFE-lite/DFE-heavy (Config 2, 16 post-taps)

Config 2 (Original: DFE bmax1=0.7)

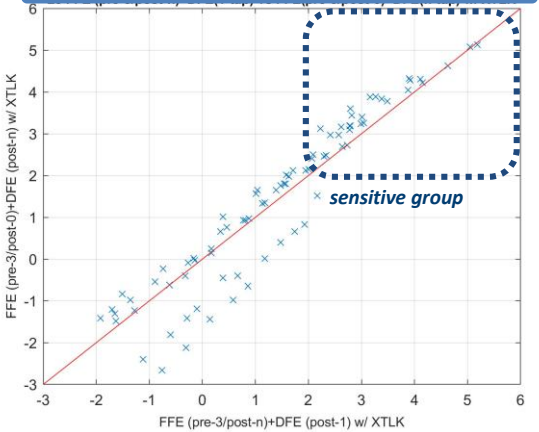

	Value	Note
Mean Diff	0.04	
Sensitive Group Mean Diff	-0.30	
False Pass	8	9.41%
False Fail	0	0.00%

Original Config 2 is optimistic w.r.t. Config 0

Config 2 (w/ DFE bmax1=0.6)

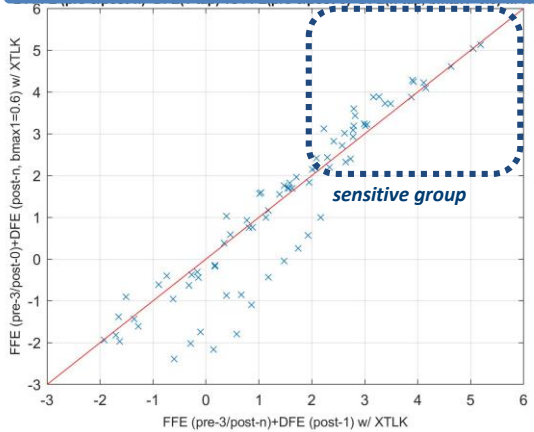

	Value	Note
Mean Diff	0.30	
Sensitive Group Mean Diff	-0.19	
False Pass	7	8.24%
False Fail	0	0.00%

Config 2 (w/ DFE bmax1=0.5)

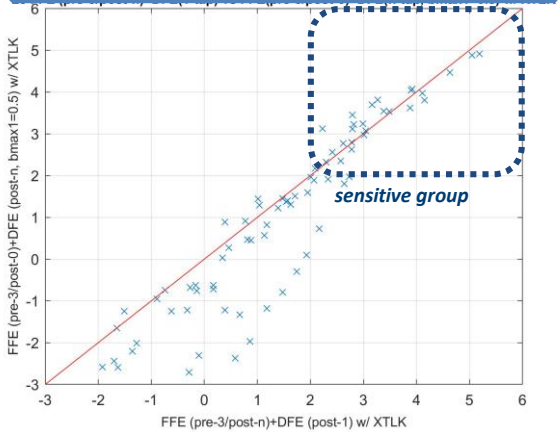

	Value	Note
Mean Diff	**0.97**	
Sensitive Group Mean Diff	**0.03**	
False Pass	**5**	5.88%
False Fail	**1**	1.18%

Config 2 w/ $b_{max}(1)$=0.5 matches better with Config 0

112Gbps Test Channel COM Results

w/ 20mm Packages w/ Crosstalk: FFE-lite/DFE-heavy (Config 2, 20 post-taps)

Config 2 (Original: DFE bmax1=0.7)

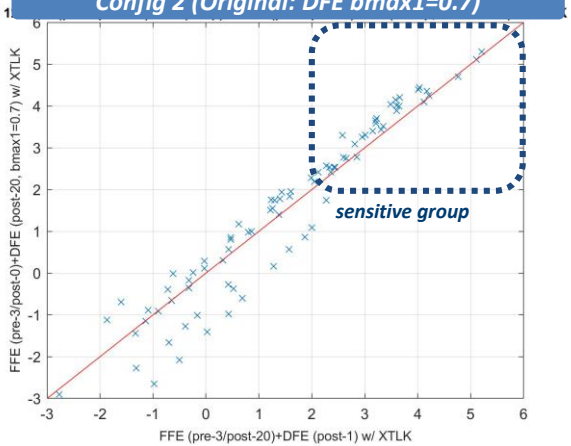

	Value	Note
Mean Diff	0.03	
Sensitive Group Mean Diff	-0.23	
False Pass	3	3.53%
False Fail	0	0.00%

Original Config 2 is optimistic w.r.t. Config 0

Config 2 (w/ DFE bmax1=0.6)

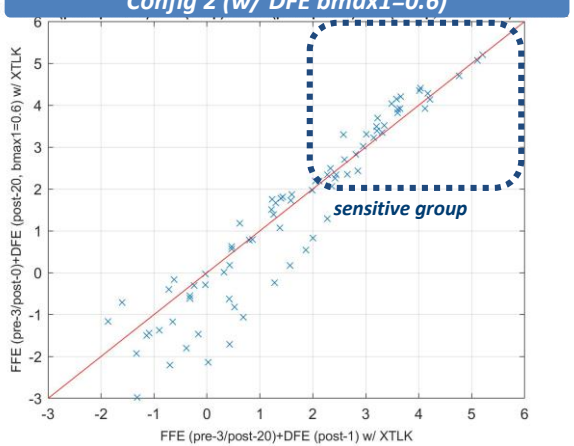

	Value	Note
Mean Diff	0.28	
Sensitive Group Mean Diff	-0.10	
False Pass	2	2.35%
False Fail	0	0.00%

Config 2 (w/ DFE bmax1=0.5)

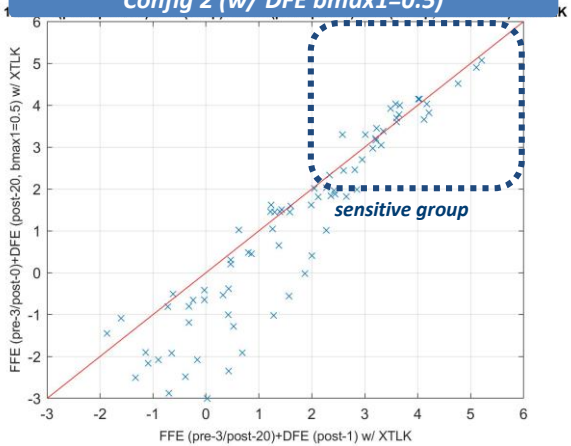

	Value	Note
Mean Diff	**0.66**	
Sensitive Group Mean Diff	**0.14**	
False Pass	**1**	1.18%
False Fail	**1**	1.18%

Config 2 w/ $b_{max}(1)$=0.5 matches better with Config 0

112Gbps Test Channel COM Results

w/ 20mm Packages w/ Crosstalk: FFE-lite/DFE-heavy (Config 2, 16/20 post-taps)

Config 2 (Original: DFE bmax1=0.7)

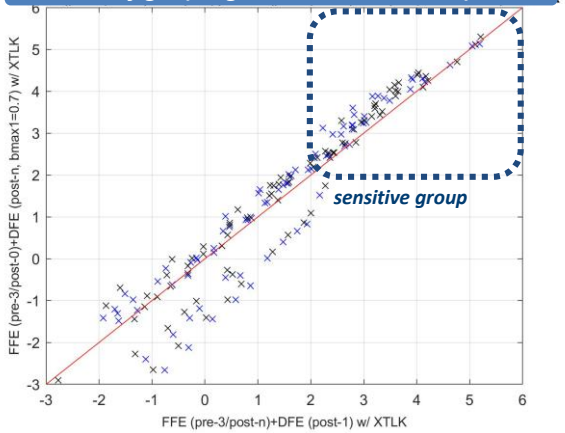

	Value	Note
Mean Diff	0.03	
Sensitive Group Mean Diff	-0.27	
False Pass	11	6.47%
False Fail	0	0.00%

Original Config 2 is optimistic w.r.t. Config 0

Config 2 (w/ DFE bmax1=0.6)

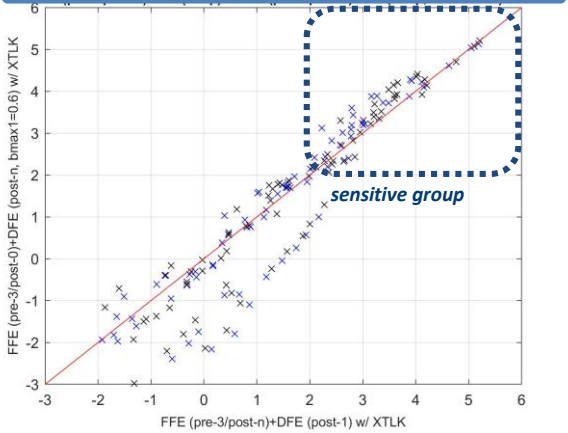

	Value	Note
Mean Diff	0.29	
Sensitive Group Mean Diff	-0.14	
False Pass	9	5.29%
False Fail	0	0.00%

Config 2 (w/ DFE bmax1=0.5)

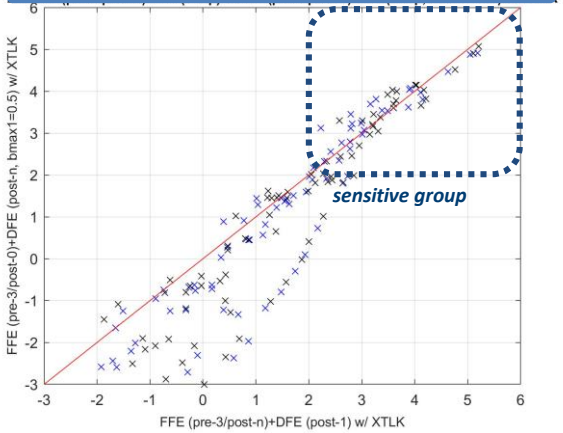

	Value	Note
Mean Diff	**0.81**	
Sensitive Group Mean Diff	**0.09**	
False Pass	**6**	3.32%
False Fail	**2**	1.18%

Config 2 w/ $b_{max}(1)$=0.5 matches better with Config 0

FFE-lite/DFE-heavy (Config 2) Adjustment

- Observed same trend in [1][2][3][4] that FFE-lite/DFE-heavy (Config 2) is seen to be more optimistic than FFE-heavy/DFE-lite (Config 0)
- In ADC-based design, DFE tap coefficient 1 is unlikely to be high because
  - 1st post-cursor ISI is jointly equalized by TX FIR, CTLE, RX FFE/DFE
- Propose adjustment to FFE-lite/DFE-heavy (Config 2)
  - Reduce DFE tap 1 coefficient limit, e.g. $b_{max}(1)$, from 0.7 to *0.55*
    - Align with FFE-heavy/DFE-lite (Config 0) performance
    - Reduce burst errors and improve FEC efficiency
    - CTLE output will be better optimized for CDR operation

Summary and Conclusions

- Goal clarifications of COM LR reference RX model
  - Provide reasonable performance of mainstream RX design, e.g. ADC-based
  - Simple to be specified in COM and 100GEL specifications
  - Likely to get consensus/agreement from IC vendors
- FFE-lite/DFE-heavy (Config 2, with $b_{max}(1)=0.55$) is shown to be the best candidate

Config #	Baseline Ref. RX	Performance (w.r.t. *Config 0*)	Complexity	Notes
0	FFE-heavy (3-pre/n-post)/DFE (1-tap)	**High**	**High** 	• Good performance and align with mainstream RX design • Complex COM model and standardization
1	DFE-only	**Low** 	**Low**	• Not align with mainstream RX design • Uses existing COM methodology
2	FFE-lite (3-pre/0-post)/DFE-heavy (n-taps)	**High** 	**Low/Medium**	• Good Performance and better match with mainstream RX design • Use existing 802.3/OIF-CEI and COM methodology

Thank You