200G/Lane AUI FEC Direction

IEEE P802.3df Task Force
IEEE 802.3 November 2022

Mark Gustlin – Cisco
Kent Lusted – Intel
Xinyuan Wang – Huawei
Adee Ran – Cisco

Supporters

Brian Welch – Cisco
Mike Dudek – Marvell
Rich Mellitz – Samtec
Shawn Nicholl – AMD
Gary Nicholl – Cisco
Dave Ofelt – Juniper
Jeff Slavick – Broadcom
Mike Li – Intel
Howard Heck – Intel
Bill Simms - Nvidia
Weiqiang Cheng - China Mobile
Liav Ben Artsi – Marvell
Ali Ghiasi - Ghiasi Quantum

Ted Sprague – Infinera
Xiang He – Huawei
Hao Ren – Huawei
Yu Xu – Huawei
Peter Stassar – Huawei
Phil Sun – Credo
Eugene Opsasnick – Broadcom
Kapil Shrikhande – Marvell
Eric Maniloff – Ciena
Matt Brown – Huawei
Haojie Wang - China Mobile
Adam Healey - Broadcom

Agenda

- A quick review many of the presentations that have preceded this and relate to the FEC choice for 200G/lane
- A proposal to choose RS(544,514,10) as the FEC for 200G/lane AUIs (C2M and C2C)

Review Previous Work - gustlin_3df_01a_220517

- These are from the 802.3df adopted logic architecture slides
- RS544 can operate as the FEC1/2/4

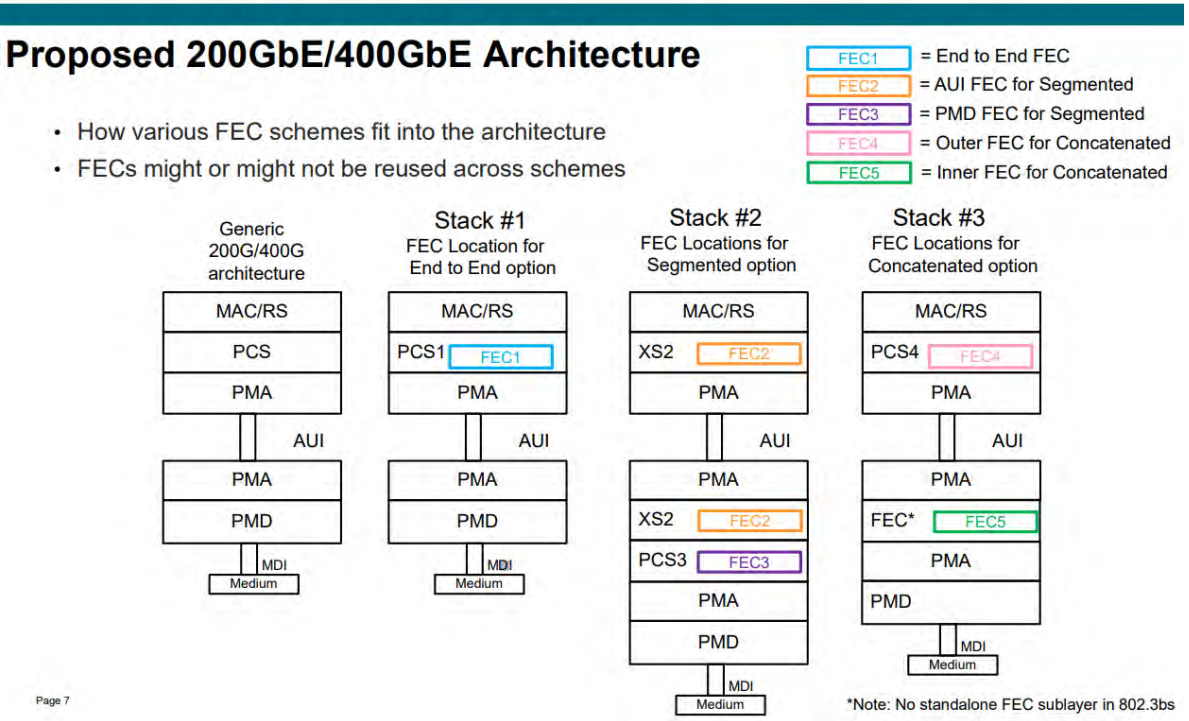

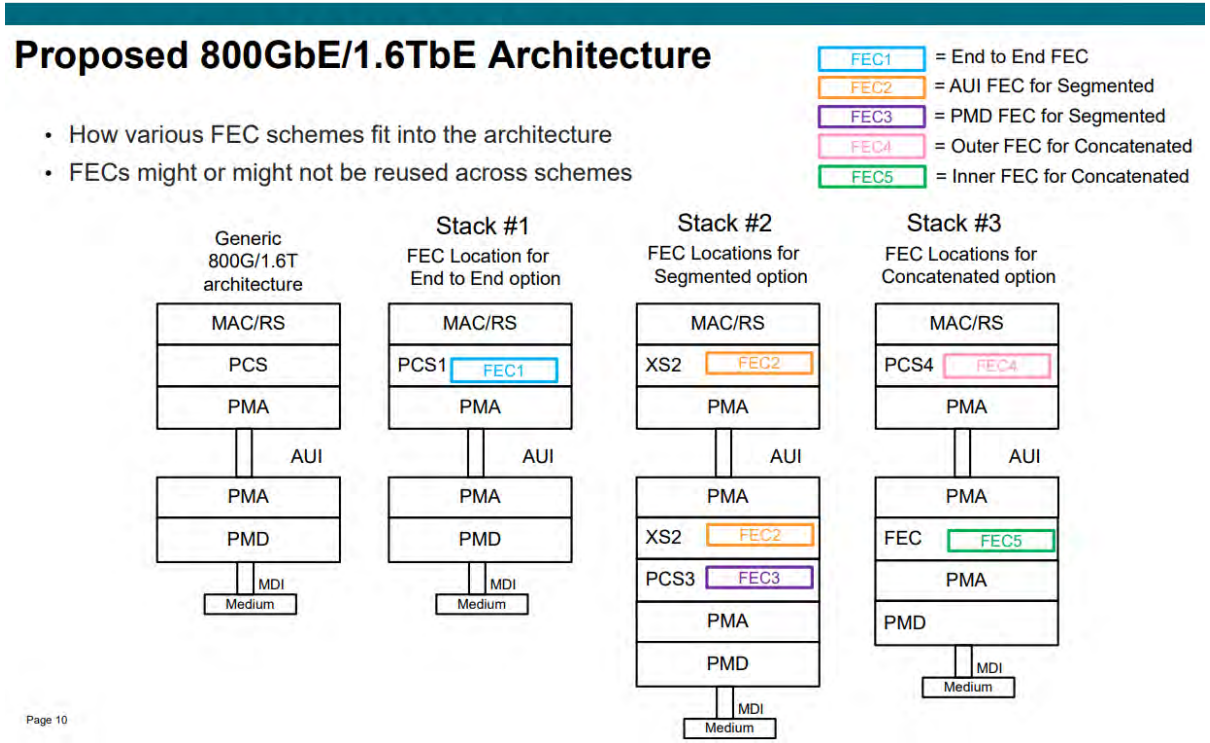

gustlin_3df_01a_220517.pdf

Review Previous Work - welch_3df_logic_220425

- Concept of two modes, end to end or segmented dependent on AUI loss
- The optical segment might be a concatenated FEC itself
- Each end determines its own operational mode (terminated or pass-through)

“End to Segmented” FEC

Brian Welch (Cisco)
Mark Gustlin (Cisco)

Link Configurations

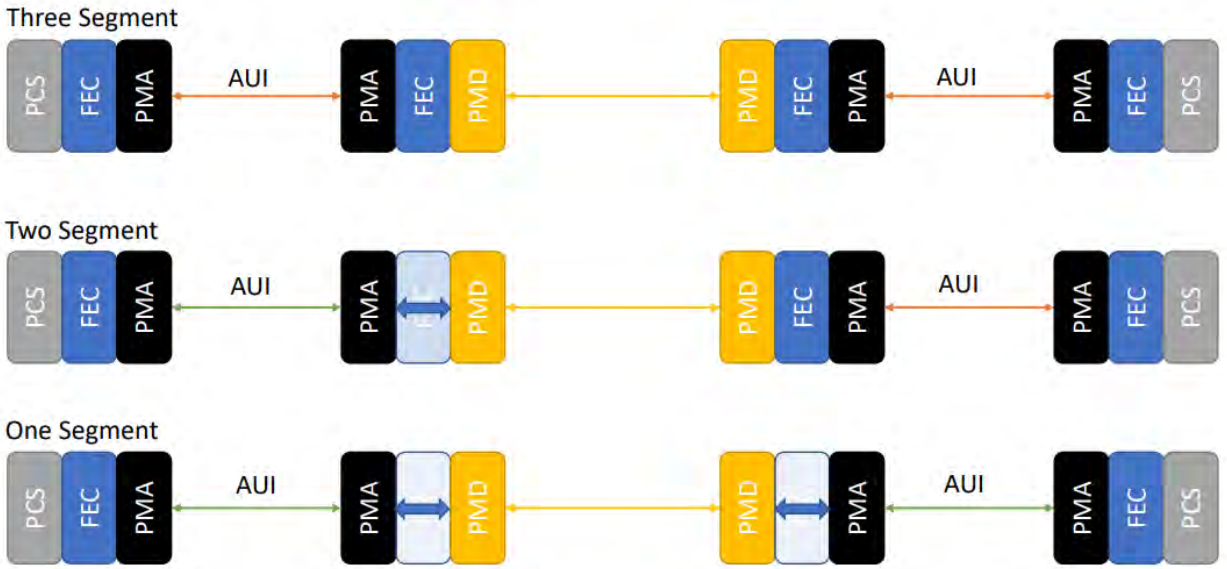

IEEE P802.3df Task Force 5

High/Low Loss C2M AUI (Conceptual)

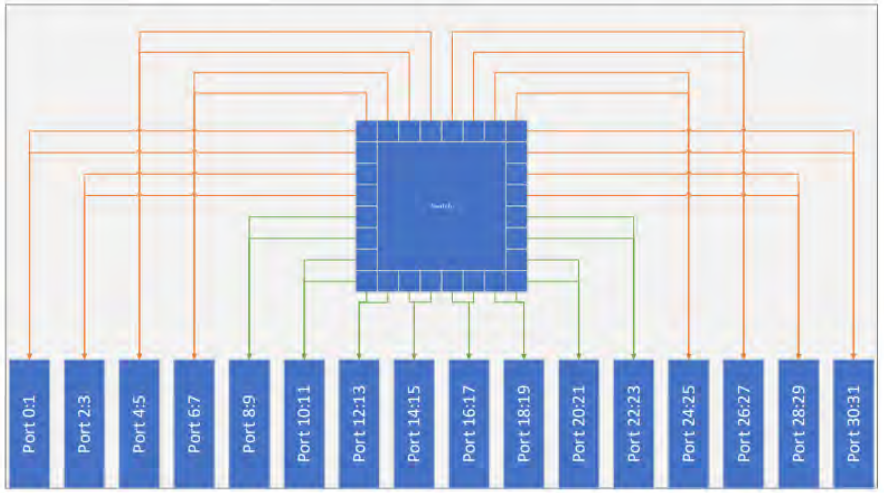

High Loss C2M: Segmented FEC Required
Low Loss C2M: Module FEC bypass allowed

IEEE P802.3df Task Force

welch_3df_logic_220425.pdf

Review Previous Work - gustlin_3df_logic_220411

- Presentation proposed to adopt RS544 FEC on the large ASIC as soon as reasonable…
  - For the AUIs
  - Now might be the right time?
- Still some stuff to figure out: bit muxing vs. symbol muxing, 1.6Tb architecture etc.

200G/Lane AUI FEC Proposal

Location of RS544 FEC

Use RS(544,514,10) for the AUIs, for both 200G/lane and 100G/lane

FEC Schemes : End-to-End, Concatenated, Segmented

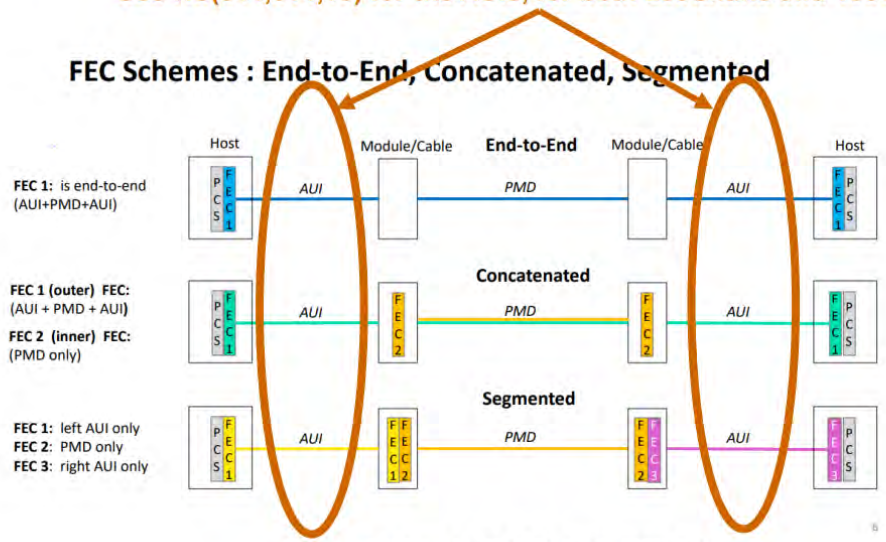

From: shrikhande_3df_01a_220203.pdf

EEE P802.3df Task Force
IEEE 802.3 April 2022
Logic ad hoc

Adopt RS544 FEC in the Large ASIC for 200G/lane

Mark Gustlin – Cisco

- Reasons **to** do this soon:
  - Gives guidance to the industry for developing large ASICs
  - RS544 FEC is very likely to be able to support at least one AUI and is therefore viable for at least one segment
  - Fully evaluating the raw BER targets for optical and electrical links is a long pole, but giving a good target to the task force to coalesce around is valuable
  - We don't want a higher overhead FEC on the AUI anyhow, impact on NCG due to speed increases reaches diminishing returns
  - Having a common RS(544,514,10) strategy across 100G/Lane for all interfaces (AUIs and PMDs) and for the AUIs for 200G/Lane is a big benefit
- Reasons **not** to do this soon:
  - What if it does not work? Then we adjust…
  - It does not address the FEC requirement for CR links (but we don't know those are at this point)

gustlin_3df_logic_220411.pdf

Review Previous Work - wang_3df_01_220215

- This Presentation shows RS544 as the outer FEC code in a concatenated overall FEC scheme

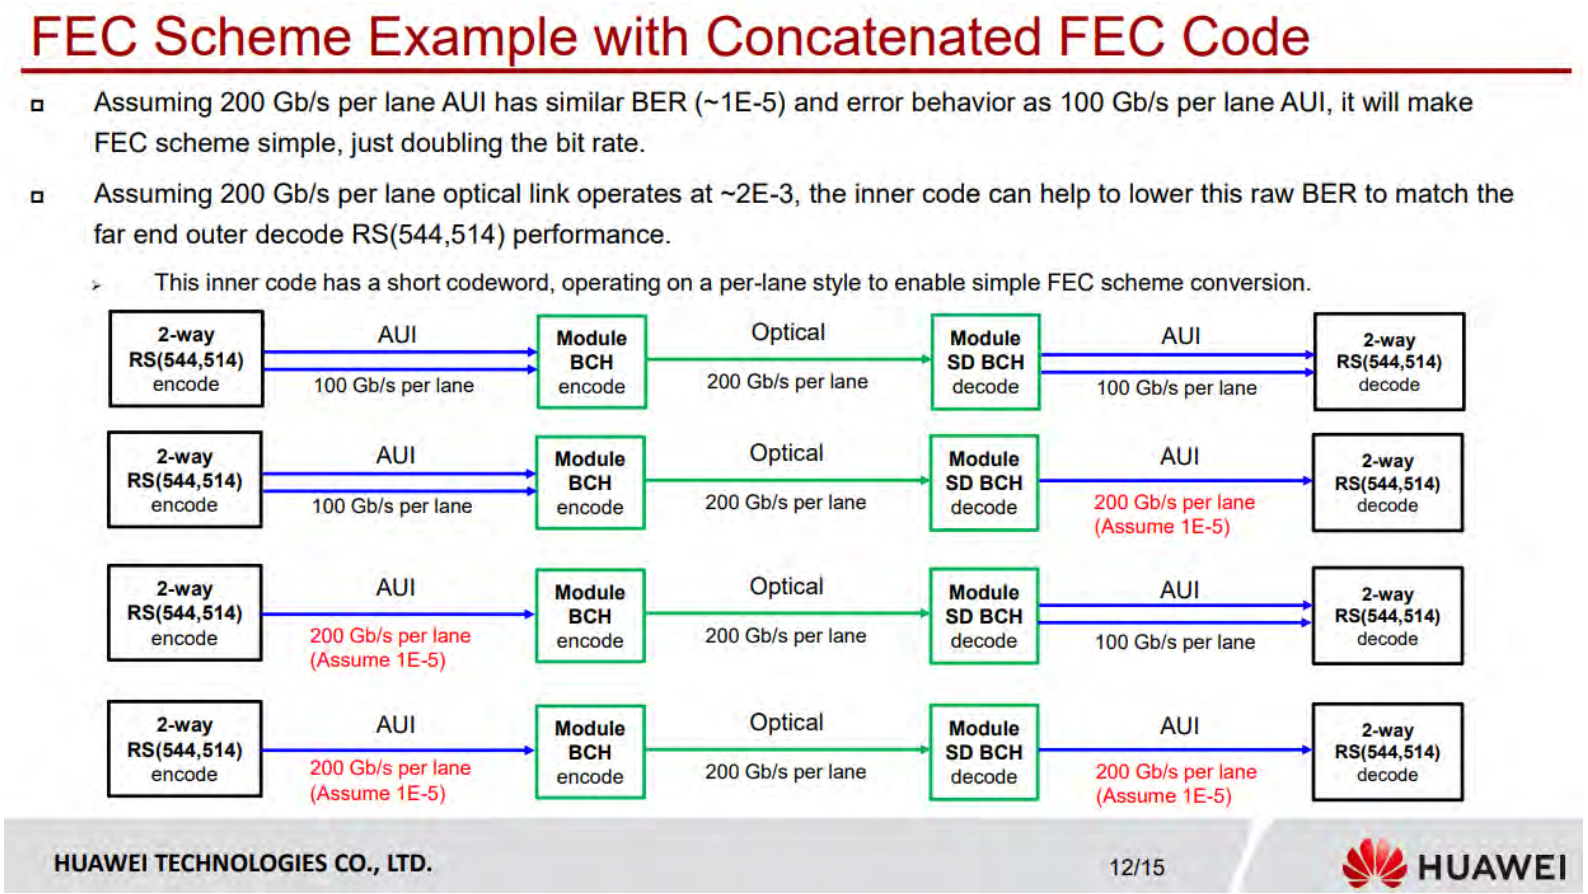

wang_3df_01_220215.pdf

Review Previous Work - lusted_3df_01_220927

- This Presentation proposed two AUI specifications based on loss targets
- Dominant choice of the straw poll favors having both medium and higher loss options

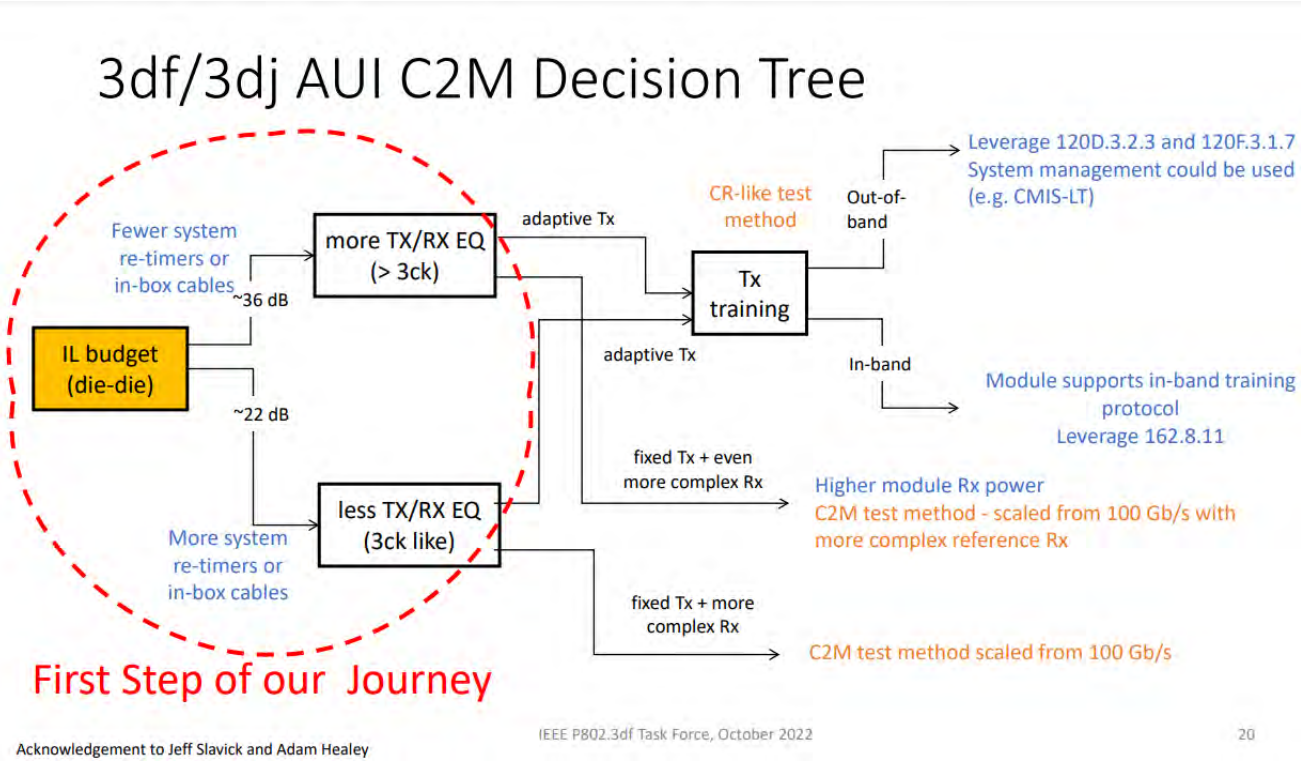

lusted_3df_01_220927.pdf

Straw Poll #1

For the front panel pluggable use case, I am interested in 200 Gbps/lane AUI C2M specifications for:

A. medium loss only (e.g. up to ~22 dB IL die-die per lusted_3df_01_220927)
B. higher loss only (e.g. up to ~36 dB IL die-die per lusted_3df_01_220927)
C. both medium and higher loss
D. need more information

pick one

Results: A: 17, B: 11, C: 49, D: 12

motions_3df_221004.pdf

Review Previous Work - ran_3df_elec_01b_220921

- This Presentation shows with certain COM choices most channels can be supported with a DER ~ 1e-4
- Other shorter channels may operate at 1e-5

COM results at operating point

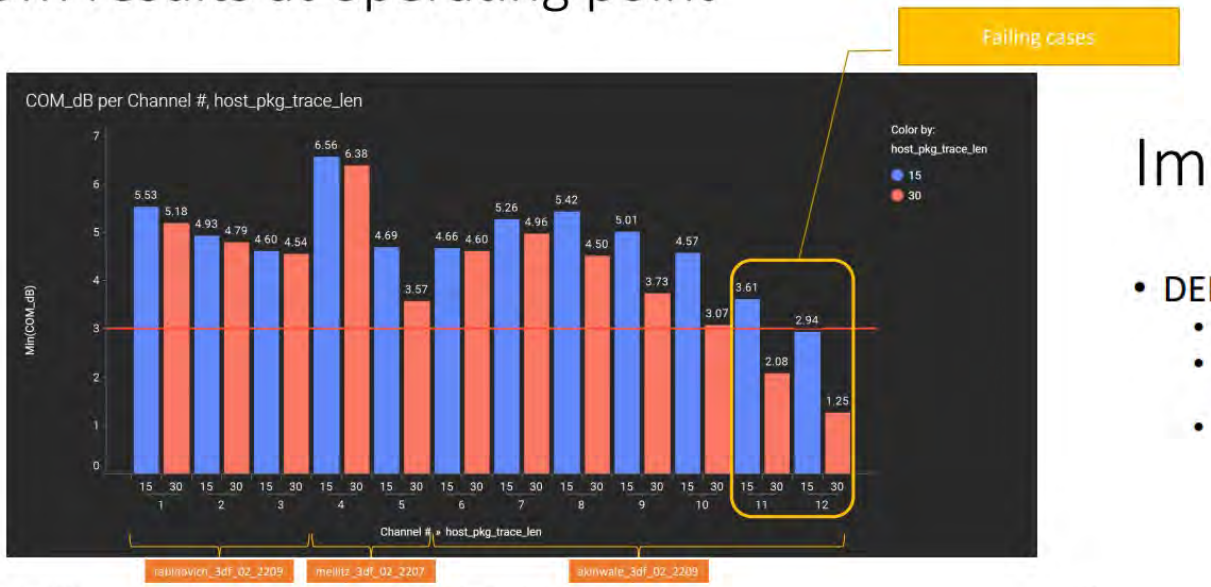

21 September 2022 IEEE P802.3df Electrical ad hoc meeting 10

Implications

- DER0=1e-4 is close to the full RS544 correction capability
  - We need to enable error correction for the electrical segment alone
  - Operating with 1e-5 may be possible with lower loss channels – in these cases it may be possible to bypass error correction
  - Possible solution: **Flexible segmented / concatenated architecture**

ran_3df_elec_01b_220921.pdf

Review Previous Work - li_3df_01a_2

- This Presentation shows with certain COM choices a DER ~ 1e-5

Summary & Conclusions

COM and Link Simulation Result Summ
212.5Gbps PAM4 (DER =10^{-5})

Channel	COM EH	COM VEC	Simulation Eye Opening Height	Simulation VEC
CH6	9.19 mV	11.405 dB	3.80 mV	14.32 dB
CH7	10.16mV	10.648 dB	4.49 mV	13.15 dB

- Preliminary COM and time-domain simulations with Intel/Amphenol latest C2M channels suggest good 212.5Gbps C2M TP1a performance/solution space with PAM4 modulation scheme.

li_3df_01a_2207.pdf

Review Previous Work - rabinovich_3df_elec_01b_220921

- This Presentation shows with certain C2M channel choices with a DER ~ 1e-5/5e-5

200G PAM4 C2M Via Length Effect Study

Structures and COM Configurations

- Three Via Lengths
  - ✓ 19 mil – 67 mil – 93 mil
- Two Breakouts
  - ✓ Parallel
  - ✓ Orthogonal
-
- Two Filters
  - ✓ Butterworth
  - ✓ Raised Cosine (starts @ 42.5 GHz, ends @ 80 GHz)
- Two Equalization Strengths *:

Equalization	DER_0	SNR_TX	eta_0	Float. Taps
Less Aggressive	1.00E-05	32.5	2.05E-08	NO
More Aggressive	5.00E-05	34	2.05E-09	YES

* Slides 17 and 18 - Mellitz_3df_elec_01_220621.pdf

KEYSIGHT TECHNOLOGIES

Page 12

200G PAM4 C2M Via Length Effect Study

COM Results

2 F[illegible]s - 1 NEXT - Small Package

Orthogonal Breakout

Case #	Via Length	C2M Configuration	Filter	DER_0	SNR_TX	eta_0	Float. Taps	EH (mV)	VEC (dB)	COM (dB)	ICN	ERL (dB)	DER
1	19 mil	Less Aggressive	Butterwort	1.00E-05	32.5	2.05E-08	NO	18.7	9.06	3.77	1.47	18.1	1.05E-12
2	19 mil	More Aggressive	Butterwor	5.00E-05	34	2.05E-09	YES	25.7	5.59	6.47	1.47	18.6	6.70E-26
3	19 mil	Less Aggressive	Raised Cos	1.00E-05	32.5	2.05E-08	NO	19.3	8.19	4.29	1.47	18.1	6.75E-15
4	19 mil	More Aggressive	Raised Co	5.00E-05	34	2.05E-09	YES	30.3	5.16	6.98	1.47	18.6	7.10E-34
5	67 mil	Less Aggressive	Butterwo	1.00E-05	32.5	2.05E-08	NO	9.2	12.79	2.26	2.04	17.5	8.27E-09
6	67 mil	More Aggressive	Butterw	5.00E-05	34	2.05E-09	YES	19.9	7.30	4.91	2.04	18.1	3.01E-14
7	67 mil	Less Aggressive	Raised C	1.00E-05	32.5	2.05E-08	NO	15.8	9.78	3.40	2.04	17.5	2.11E-11
8	67 mil	More Aggressive	Raised C	5.00E-05	34	2.05E-09	YES	24.1	5.99	6.05	2.04	18.1	5.35E-21
9	93 mil	Less Aggressive	Butterw	1.00E-05	32.5	2.05E-08	NO	5.6	18.36	1.12	2.27	15.4	5.09E-07
10	93 mil	More Aggressive	Butterw	5.00E-05	34	2.05E-09	YES	17.0	8.23	4.26	2.27	15.9	6.62E-12
11	93 mil	Less Aggressive	Raised C	1.00E-05	32.5	2.05E-08	NO	11.5	12.19	2.45	2.27	15.4	3.57E-09
12	93 mil	More Aggressive	Raised C	5.00E-05	34	2.05E-09	YES	19.9	6.65	5.43	2.27	15.9	1.15E-16

Parallel Breakout

Case #	Via Length	C2M Configuration	Filter	DER_0	SNR_TX	eta_0	Float. Taps	EH (mV)	VEC (dB)	COM (dB)	ICN	ERL (dB)	DER
1	19 mil	Less Aggressive	Butterw	1.00E-05	32.5	2.05E-08	NO	15.2	10.18	3.22	1.79	18.3	5.66E-11
2	19 mil	More Aggressive	Butterw	5.00E-05	34	2.05E-09	YES	25.8	6.23	5.82	1.79	18.8	3.06E-19
3	19 mil	Less Aggressive	Raised C	1.00E-05	32.5	2.05E-08	NO	19.8	8.66	4.00	1.79	18.3	1.99E-13
4	19 mil	More Aggressive	Raised C	5.00E-05	34	2.05E-09	YES	26.0	5.53	6.54	1.79	18.8	4.97E-26
5	67 mil	Less Aggressive	Butterw	1.00E-05	32.5	2.05E-08	NO	10.1	13.27	2.12	2.36	17.9	1.39E-08
6	67 mil	More Aggressive	Butterw	5.00E-05	34	2.05E-09	YES	20.6	8.19	4.29	2.36	18.5	2.75E-12
7	67 mil	Less Aggressive	Raised Co	1.00E-05	32.5	2.05E-08	NO	16.2	10.36	3.14	2.36	17.9	1.12E-10
8	67 mil	More Aggressive	Raised Cos	5.00E-05	34	2.05E-09	YES	23.4	6.49	5.58	2.36	18.5	7.34E-18
9	93 mil	Less Aggressive	Butterwor	1.00E-05	32.5	2.05E-08	NO	7.9	16.38	1.43	2.62	15.0	1.81E-07
10	93 mil	More Aggressive	Butterwort	5.00E-05	34	2.05E-09	YES	17.8	9.19	3.70	2.62	15.5	1.73E-10
11	93 mil	Less Aggressive	Raised Cosin	1.00E-05	32.5	2.05E-08	NO	13.6	12.20	2.45	2.62	15.0	3.48E-09
12	93 mil	More Aggressive	Raised Cosine	5.00E-05	34	2.05E-09	YES	24.1	7.09	5.07	2.62	15.5	3.29E-15

* Pass: EHmin = 10 mV; VECmax = 12.5; ERLmi[illegible] = 10

KEYSIGHT TECHNOLOGIES

Page 13

rabinovich_3df_elec_01b_220921.pdf

Review Previous Work -mellitz_3df_01_2207

- This Presentation shows with certain C2C channel choices with a DER ~ 1e-5

COM 3.8beta example (w/ crosstalk

FOR TEST BOARD WIRED TO TEST CHIP ATTACH

Table 93A-1 parameters			2
Parameter	Setting	Units	Information
f_b	106.25	GBd	
f_min	0.05	GHz	
Delta_f	0.01	GHz	
C_d	[0.4e-4 0.9e-4 1.1e-4 ; 0.4e-4 0.9e-4 1.1e-4]	nF	[TX RX]
L_s	[.13 .15 .14; .13 .15 .14]	nH	[TX RX]
C_b	[.3e-4 .3e-4]	nF	[TX RX]
z_p select	[1 2]		[test cases to run]
z_p (TX)	[12 31; 1.8 1.8]	mm	[test cases]
z_p (NEXT)	[12 29; 1.8 1.8]	mm	[test cases]
z_p (FEXT)	[12 31; 1.8 1.8]	mm	[test cases]
z_p (RX)	[12 29; 1.8 1.8]	mm	[test cases]
C_p	[0.5e-4 0.5e-4]	nF	[TX RX]
R_0	50	Ohm	
R_d	[45 45]	Ohm	[TX RX]
A_v	0.408	V	
A_fe	0.408	V	
A_ne	0.608	V	
L	4		
M	32		
filter and Eq			
f_r	0.75	*fb	
c(0)	0.6		min
c(-1)	[-0.34:0.02:0]		[min:step:max]
c(-2)	[0:0.02:0.2]		[min:step:max]
c(-3)	[-0.1:0.02: 0]		[min:step:max]
c(1)	[-0.1:0.02:0]		[min:step:max]
N_b	12	UI	
b_max(1)	0.85		
b_max(2..N_b)	0.3		
b_min(1)	-0.85		
b_min(2..N_b)	-0.3		
g_DC	[-20:1:0]	dB	[min:step:max]
f_z	42.5	GHz	
f_p1	42.5	GHz	
f_p2	106.25	GHz	
g_DC_HP	[-8:1:0]		[min:step:max]
f_HP_PZ	1.0625	GHz	
param.RC_Start	42500000000	HZ	
param.RC_end	79687500000	HZ	
Butterworth	0		logical
Raised_Cosine	1		logical

I/O control		
DIAGNOSTICS	1	logical
DISPLAY_WINDOW	1	logical
CSV_REPORT	1	logical
RESULT_DIR	.\results\200G_C2C_{date}\	
SAVE_FIGURES	0	logical
Port Order	[1 2 3 4]	logical
RUNTAG	R200_eval	
COM_CONTRIBUTION	0	logical
Operational		
C[illegible]ass threshold	3	dB
[illegible] Pass threshold	10.5	dB
DER_0	1.00E-05	
[illegible]T_r	3.29E-03	ns
FO[illegible]		logical
Local Search	2	
SAVE_CONFIG2MAT	0	
TDR and ERL options		
TDR	1	logical
ERL	1	logical
ERL_ONLY	0	logical
TR_TDR	0.01	ns
N	3000	
beta_x	0	
rho_x	0.618	
fixture delay time	[0 0]	[port1 port2]
TDR_W_TXPKG	0	
N_bx	0	UI
Z_t	[45 50]	ohm
Receiver testing		
RX_CALIBRATION	0	logical
Sigma BBN step	5.00E-03	V
Noise, jitter		
sigma_RJ	0.01	UI
A_DD	0.02	UI
eta_0	4.10E-09	V^2/GHz
SNR_TX	34	dB
R_LM	0.95	

Table 93A-3 parameters	
Parameter	Setting
package_tl_gamma0_a1_a2	[0 8.4e-4 1.1e-4]
package_tl_tau	6.14E-03
package_Z_c	[87.5 87.5 ; 92.5 92.5]
Table 92-12 parameters	
Parameter	Setting
board_tl_gamma0_a1_a2	[0 6.44084e-4 3.6036e-05]
board_tl_tau	0.00579
board_Z_c	100
z_bp (TX)	40
z_bp (NEXT)	40
z_bp (FEXT)	40
z_bp (RX)	40
C_0	[0.2e-4]
C_1	[0.1e-4]
Include PCB	0
Floating Tap Control	
N_bg	4
N_bf	3
N_f	60
bmaxg	0.2
B_float_RSS_MAX	0.2
N_tail_start	25
ICN parameters	
f_v	0.676
f_f	0.676
f_n	0.676
f_2	79.688
A_ft	0.600
A_nt	0.600

COM results

TP0 TO TP5 EXAMPLE

- Channel ERL
  - ERL for Z_t 45 Ω : 13.1 dB
  - ERL for Z_t 50 Ω : 12.2 dB
- Package 1 (12 mm)
  - COM : 4.34 dB
- Package 2 (31 mm)
  - COM : 3.57 dB

Review Previous Work - welch_3df_01a_221011

- This Presentation shows FEC options for the DR/FR PMDs at 200G/lane
- RS544 might be viable for these optical links
  - An added BCH code might not give you too much NCG, depending on various component BWs etc.

Case 3: Comparisons to FEC Options

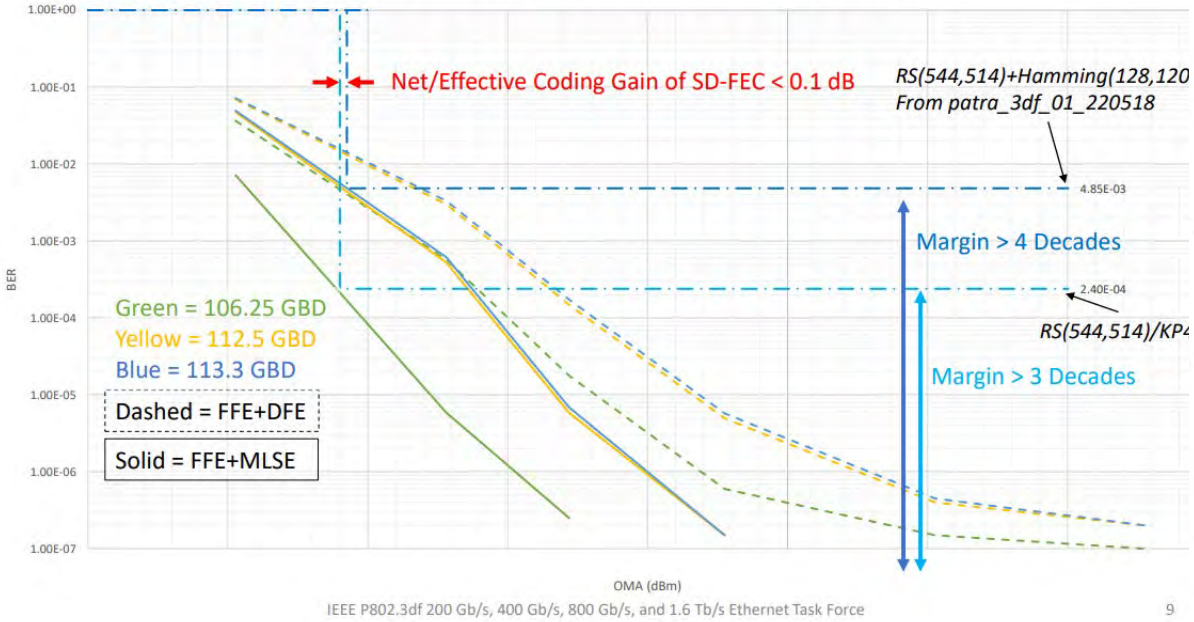

Summary and Observations

- **More Analysis Needed**
  - Investigate different filter types/responses for creating max TDECQ
- Increasing baud rate impacts pre-FEC BER curve
  - Also expected to impact module power, more study needed.
- RS(544,514) may be viable for 200G/L
  - Degradation vs. 100G/L may be ≤ 2dB
  - Margin to noise floor ≥ 3 decades
- Net/effective coding gain of SD-FEC reduced by increase of baud rate
  - Assuming the same transmitter (same EQ & SER applied for TECQ calibration) and receiver (same bandwidth).

Review Previous Work - simms_3df_01_221005

- This Presentation shows an RS544 FEC options working for end-to-end FEC at 200G/lane
  - Assumes AUI at 1e-5 or better

SIMULATION RESULTS

- A plot of BER vs. OMA (at the Rx photo-diode) is shown for 212.5 Gbps:
  - Baseline EQ : FFE + 1-tap DFE (Blue)
  - Advanced EQ : FFE + MLSE (Black)
  - Advanced EQ : FFE + MLSE + Reduced RIN (Magenta)
  - 800G MSA shown for reference (Red dotted)
- BER target of 2 x 10^{-4} (Red) can be achieved under given conditions
- We ca[illegible]-to-end KP4 RS(544,514) FEC in the ho[illegible] under the usua[illegible]assumptions that the AUI links can m[illegible]et BER of 10^{-5} or [illegible]etter
- With this example the link budget is closed with -6.4dBm

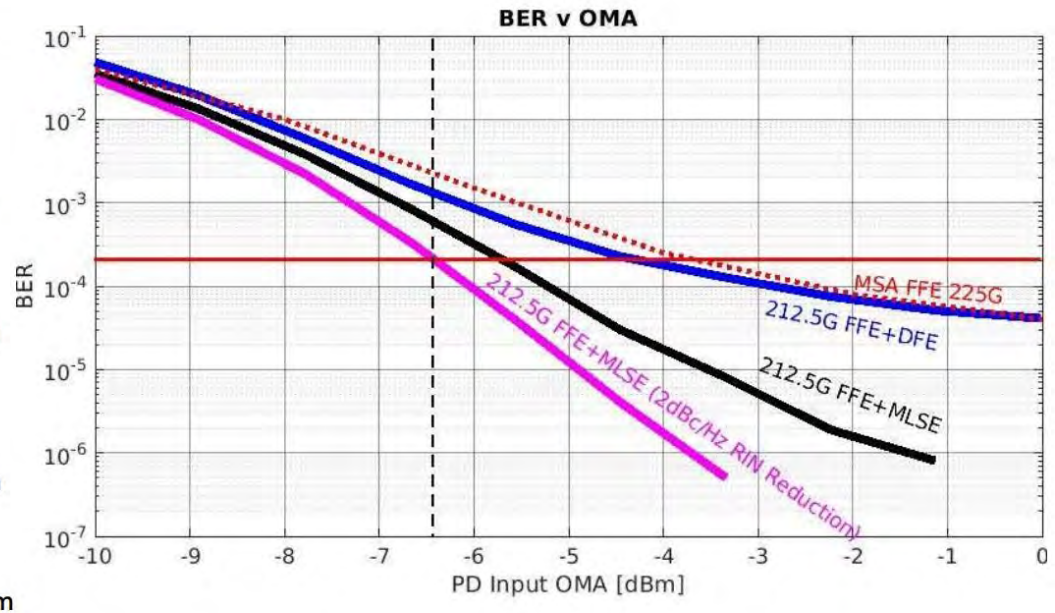

*5: MSA FFE 225G curve extracted from 800G Pluggable Multi-Source Agreement White Paper 200g-per-lane-for-future-800g-and-16t-modules

8

simms_3df_01_221005.pdf

Proposal

- Adopt RS544 now as the FEC for the 200G/lane AUI
- Reasons to do this now:
  - Having a common RS(544,514,10) FEC strategy across 100G/Lane for all interfaces (AUIs and PMDs) and for the AUIs for 200G/Lane is a big benefit
  - Gives guidance to the industry for developing large ASICs
    - The FEC structures are quite important for power and area in these devices
  - RS544 FEC looks to be able to support the 200G/lane AUIs on one side of a link
    - When dedicating the gain just for the AUI
  - Fully evaluating the raw BER targets for optical and electrical links is a long pole, but giving a good target to the task force to coalesce around is valuable
  - We don’t want a higher overhead FEC on the AUI anyhow, impact on NCG due to speed increases reaches diminishing returns

What More Needs to be Defined?

- **Need to decide on the format of the RS544 on the AUI at 200G/lane**
  - Bit muxing?
    - Up to N:1, where N depends on the rate and number of lanes
    - Might require precoding to reduce the burst error impacts
    - Simplest option
  - Symbol muxing?
    - Maintains burst error tolerance better
    - Slightly more complicated
  - Some other option?
  - Still being studied, leave for a future decision

What More Needs to be Defined? Cont..

- What is the FEC structure on the PMD
  - Possibly also RS544, and possibly with a BCH inner FEC
    - At least for the shorter PMDs (500m/2km)
    - Typically, with symbol muxing and or a level of permutation to spread out errors
  - This can be decided in the future, does not constrain the AUI FEC choice
    - If it is RS544 (with or without a BCH), then you can support the multiple modes as shown in a couple of slides

What this Presentation Proposes

- We are only proposing that we decide to use RS(544,514,10) as the FEC for 200G/lane AUIs
- Distribution is still TBD
- # of codewords is not part of this proposal, but:
  - Assume we would reuse 2CWs for 200/400GE
  - And 4CWs for 800GE
  - TBD for 1.6TbE

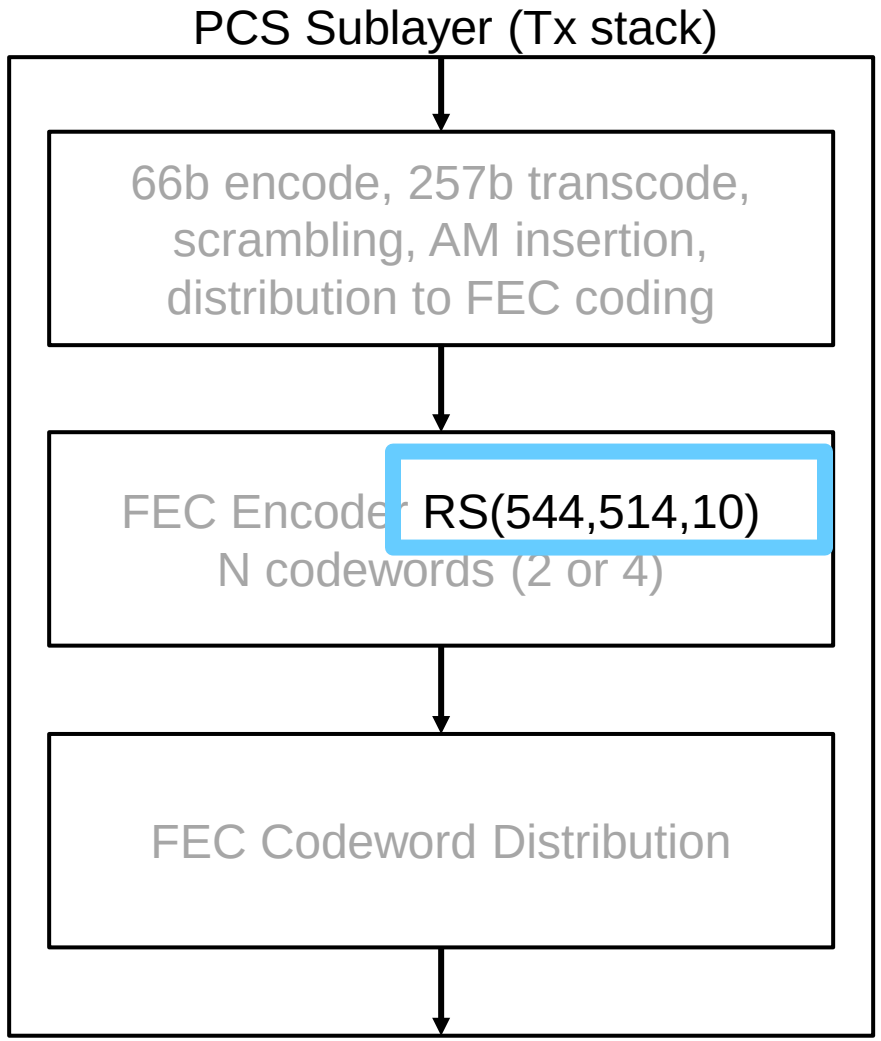

Enabling Multiple Modes

- We will likely have multiple AUI reach/losses
- We should support multiple FEC modes as shown below (if there is synergy with the PMD FEC)
- The RS544 FEC is either terminated or not depending on the correction needed for a given AUI

Segmented FEC (3 segments), concatenated code on the PMD

Partially segmented (2 segments), concatenated code on the PMD)

End to end FEC with a concatenation in the PMD

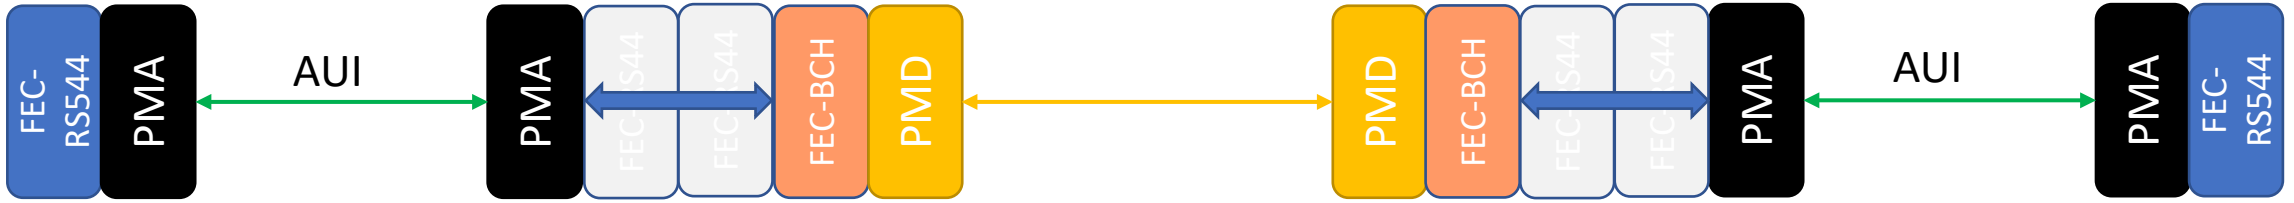

* Diagrams are conceptual and not compliant with adopted architecture

** A BCH FEC might or might not be required for the PMD span

Summary

- It is clear from the bulk of the presentations that we should be retaining RS544 as the FEC for 200G/lane AUIs
- Other details can be figured out soon
- Making this decision now will focus the task force
- This choice does not stop us from adopting a FEC scheme that might be different for copper cables (if needed)

Possible Straw Poll

- I would support adopting RS(544,514,10) as the FEC code for the 200G/lane AUIs (C2M and C2C)
  - Y
  - N
  - Need more information

Thanks!