200G/Lane Host to Module Short Channels

Femi Akinwale, Kusuma Matta, Oleg Kashurkin, Cesar Mendez-Ruiz, Howard Heck, Francisco Olguin Tellez

IEEE P802.3df 200Gb/s, 400Gb/s, 800Gb/s and 1.6Tb/s Ethernet Task Force

Contributors

- Michael Rowlands (Amphenol)
- Sam Kocsis (Amphenol)
- Ali Hammoodi (Amphenol)

Outline

- Objective
- Channel Description
- Channel Model Overview
- Channel Characteristics
- COM Setup and Results
- Summary

Objective

- Provide a set of channels representative of “server” and “nic” designs
  - Focus is on short (<5in) and low-cost solution
  - Optimization is ongoing and updates will be made

Architectural changes to ToRs due to reduced physical VSR reach

- Hypothetical Example:
  - 25.6T, 256 x 100G
  - 1RU box, Single ASIC (ToR design profile, also used as virtual chassis, aka “Fixed Box”)
  - Can be used with all optical IO in a spine application (common practice today in hyperscale datacenters)
  - 32 x 800G module cages, all front panel IO
- Using Rosemont budget proposal from Jane Lim:
  - http://www.ieee802.org/3/100GEL/public/18_03/lim_100GEL_01b_0318.pdf
  - [~ 5” Host trace supported for VSR channels]
  - Approximately 12 / 32 module cages cannot accommodate the proposed host budgets (VSR or CR), requiring either intermediate retimers, or intra-box cabling

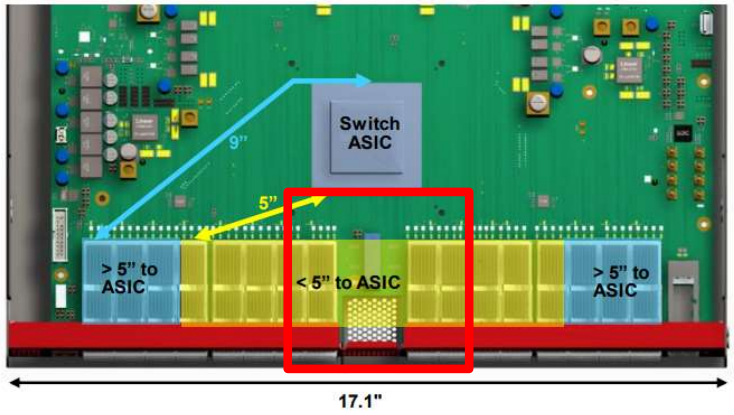

Source: stone_3ck_01a_0518

Channel Description

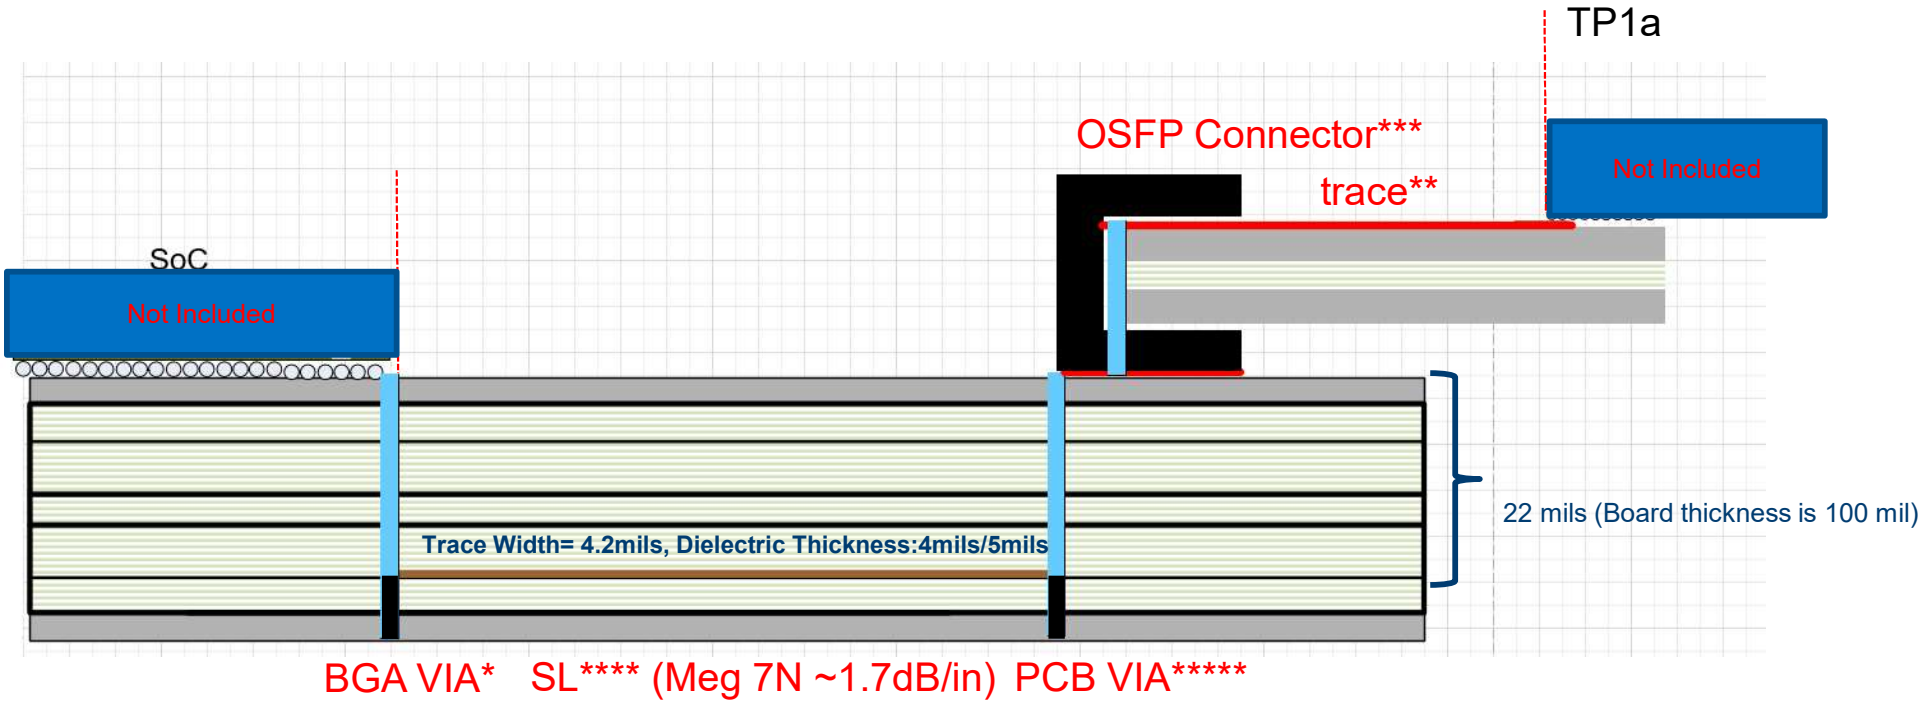

*BGA footprint included in the channel
**Module Loss is 3.5dB @ 53.125 GHz

***Connector loss is 2.2dB @53.125GHz
****PCB Loss is Max 7dB @53.125GHz (93 ohms)
*****Vias are staggered microVia
*****Assumed 1 FEXT and 2 NEXT aggressors

Channel Model Overview

- Assumed PCB C2M Topology
  - TP0 to TP1a Loss Targets: 10dB, 11dB, 12dB, and 13dB
- Synthesized channel with OSFP200G Connector (Concept from Amphenol)
- Assumed nominal design.
- Assumed use of staggered uVia technology.
- Transmission line loss is 1.7dB/in @ 53.125GHz
  - Skip Layer is not assumed
  - Trace Width is 4.6mils with Intra-pair spacing of 7.2mils

Channel Characteristics

10dB with staggered microVia

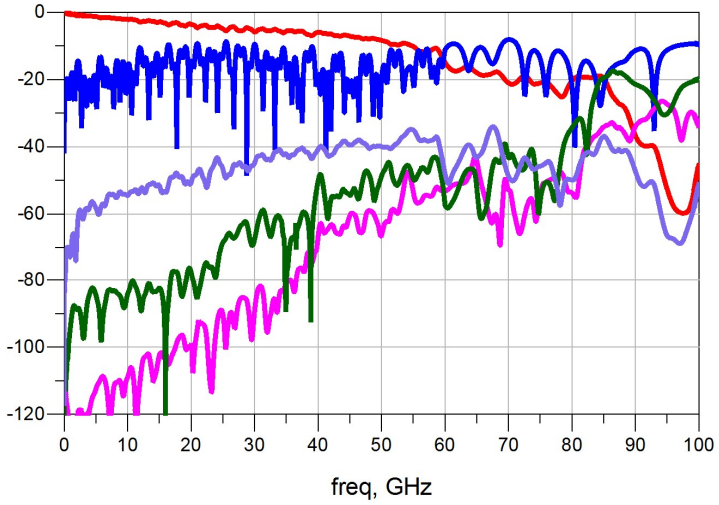

Insertion Loss
Return Loss
NEXT1
NEXT2
FEXT

11dB with staggered microVia

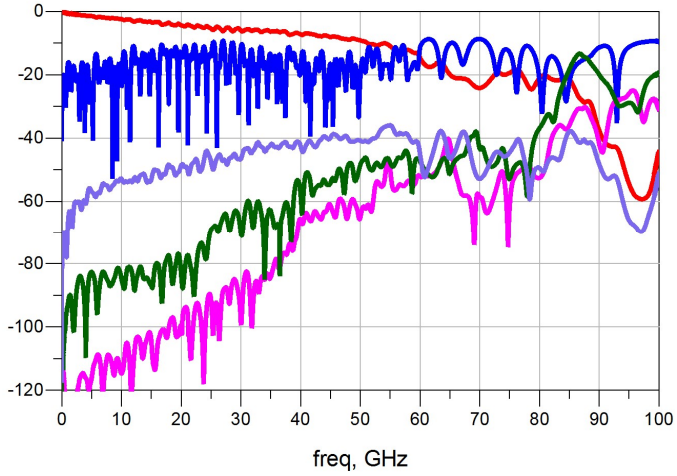

12dB with staggered microVia

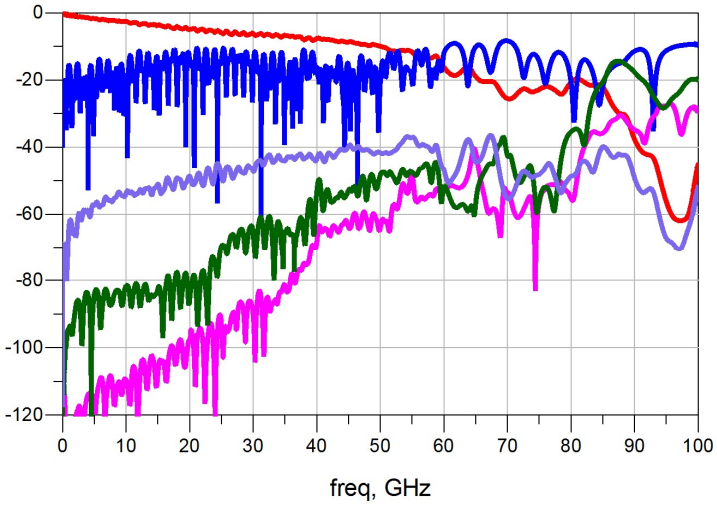

13dB with staggered microVia

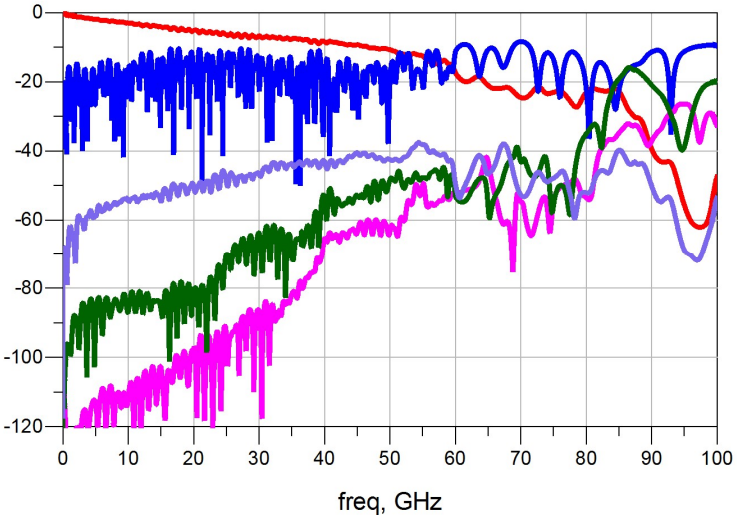

COM Setup

Table 93A-1 parameters			
Parameter	**Setting**	**Units**	**Information**
f_b	106.25	GBd	
f_min	0.05	GHz	
Delta_f	0.01	GHz	
C_d	[0.7e-4 0]	nF	[TX RX]
L_s	[0.12 0]	nH	[TX RX]
C_b	[0.3e-4 0]	nF	[TX RX]
z_p select	[1 2]		[test cases to run]
z_p (TX)	[15 31; 1.8 1.8]	mm	[test cases]
z_p (NEXT)	[0 0 ; 0 0]	mm	[test cases]
z_p (FEXT)	[15 31; 1.8 1.8]	mm	[test cases]
z_p (RX)	[0 0 ; 0 0]	mm	[test cases]
C_p	[0 0]	nF	[TX RX]
R_0	50	Ohm	
R_d	[50 50]	Ohm	[TX RX]
A_v	0.413	V	vp/vf=.694
A_fe	0.413	V	vp/vf=.694
A_ne	0.45	V	
L	4		
M	32	Samp/UI	
samples_for_C2M	32	Samp/UI	
T_O	0	mUI	
AC_CM_RMS	0	V	[test cases]
filter and Eq			
f_r	0.75	*fb	
c(0)	0.54		min
c(-1)	[-0.34:0.02:0.2]		[min:step:max]
c(-2)	[-0.1:0.02:0.1]		[min:step:max]
c(-3)	[-0.1:0.02:0.1]		[min:step:max]
c(-4)	[-0.04:0.02:0.04]		[min:step:max]
c(1)	[-0.1:0.02:0.2]		[min:step:max]
N_b	8	UI	
b_max(1)	**0.85**		**As/dffe1**
b_max(2..N_b)	**[0.3 0.3 0.2*ones(1,5)]**		**As/dfe2..N_b**
b_min(1)	**0**		**As/dffe1**
b_min(2..N_b)	**[-0.05 -0.05 -0.03*ones(1,5)]**		**As/dfe2..N_b**
g_DC	[-13:1:-2]	dB	[min:step:max]
f_z	42.5	GHz	
f_p1	42.5	GHz	
f_p2	106.25	GHz	
g_DC_HP	[-3:0.5:-0]		[min:step:max]
f_HP_PZ	2.65625	GHz	
G_Qual	[-2 -13 ;-3 -12; -4 -11;-5 -10]	dB	ranges
G2_Qual	[0 -1 -2 -3]	dB	ranges
GDC_Min	0	dB	0 disables check. maybe different for each interface.

I/O control		
DIAGNOSTICS	1	logical
DISPLAY_WINDOW	1	logical
CSV_REPORT	1	logical
RESULT_DIR	.\results\100GEL_C2M_host_dfe8	
SAVE_FIGURES	0	logical
Port Order	[1 2 3 4]	
RUNTAG	C2M_eval_	
COM_CONTRIBUTION	0	logical
Local Search	2	
Operational		
VEC Pass threshold	12	db
EH_min	10	mV
ERL Pass threshold	7.3	dB
Min_VEO_Test	0	mV
DER_0	1.00E-05	
T_r	0.00375	ns
FORCE_TR	1	5
PMD_type	C2M	
BREAD_CRUMBS	0	logical
SAVE_CONFIG2MAT	1	logical
PLOT_CM	0	logical
TDR and ERL options		
TDR	1	logical
ERL	1	logical
ERL_ONLY	0	logical
TR_TDR	0.01	ns
N	800	
beta_x	0	
rho_x	0.618	
fixture delay time	[0 0.2e-9]	[port1 port2] different for each test fixture
TDR_W_TXPKG	1	
N_bx	0	UI
Tukey_Window	1	
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	2.05E-08	V^2/GHz
SNR_TX	32.5	dB
R_LM	0.95	

Table 93A–3 parameters		
Parameter	**Setting**	**Units**
package_tl_gamma0_a1_a2	[0 0.000644085 0.00018018]	
package_tl_tau	5.700E-03	ns/mm
package_Z_c	[87.5 87.5 ; 92.5 92.5]	Ohm
ICN & FOM_ILD parameters		
f_v	0.594	*Fb
f_f	0.594	GHz f_r specified in first column
f_n	0.594	GHz
f_2	80	GHz
A_ft	0.600	V
A_nt	0.600	V
Histogram_Window_Weigh	Gaussian	gaussian. triangle, rectangle
sigma_r	0.02	sigma in UI fo or gaus.. Wind

Table 92–12 parameters		
Parameter	**Setting**	
board_tl_gamma0_a1_a2	[0 3.8206e-04 9.5909e-05]	
board_tl_tau	0.00579	ns/mm
board_Z_c	100	Ohm
z_bp (TX)	407	mm
z_bp (NEXT)	407	mm
z_bp (FEXT)	407	mm
z_bp (RX)	407	mm
C_0	0	nF
C_1	0	nF
Include PCB	0	logical

upodated for D3.1

Floating Tap Control		
N_bg	0	0 1 2 or 3 groups
N_bf	3	taps per group
N_f	40	span for floating taps
bmaxg	0.2	FE value for floating taps

Referenced: tli_3df_01b_220316
COM: version 3.7
Highlighted **Red**: Changed

Channel Results

Channel	Tx Package Length[mm]	FOM_ILD	MDNEXT_ICNmV	MDFEXT_ICN_mV	ICN_mV	ERL(1e-4)	fitted_IL_dB_at_Fnq	DER= 1e-4 COM[dB]	DER= 1e-5 COM[dB]
C2M_PCB_10dB	15	0.54	1.08	3.72	3.88	12.23	10.32	2.53	1.38
C2M_PCB_10dB	31	0.54	1.08	3.72	3.88	12.23	10.32	3.08	1.93
C2M_PCB_11dB	15	0.60	1.08	3.41	3.58	12.73	11.05	2.84	1.68
C2M_PCB_11dB	31	0.60	1.08	3.41	3.58	12.73	11.05	3.38	2.21
C2M_PCB_12dB	15	0.55	1.04	3.16	3.33	13.20	11.94	3.06	1.90
C2M_PCB_12dB	31	0.55	1.04	3.16	3.33	13.20	11.94	3.37	2.20
C2M_PCB_13dB	15	0.54	1.02	2.90	3.08	13.64	12.81	3.13	1.98
C2M_PCB_13dB	31	0.54	1.02	2.90	3.08	13.64	12.81	3.41	2.25

Channel	Tx Package Length[mm]	FOM_ILD	MDNEXT_ICN_mV	MDFEXT_ICN_mV	ICN_mV	ERL(1e-4)	fitted_IL_dB_at_Fnq	DER= 1e-4 COM[dB]	DER= 1e-5 COM[dB]
KEY_C2M_200G_120G_4p0HCB_022422_Thru	15	0.56	6.19	6.35	8.87	16.90	12.40	3.64	2.48
KEY_C2M_200G_120G_4p0HCB_022422_Thru	31	0.56	6.19	6.35	8.87	16.90	12.40	3.33	2.17
KEY_C2M_200G_120G_2p5HCB_022422_Thru	15	0.57	4.50	5.47	7.08	18.49	13.89	3.96	2.78
KEY_C2M_200G_120G_2p5HCB_022422_Thru	31	0.57	4.50	5.47	7.08	18.49	13.89	3.63	2.45

Channels from [rabinovich_3df_01a_220224]

Red indicates failing 3dB COM limit

COM: version 3.7

Summary

- Providing 4 channels representative of a realistic implementation
- Improved Equalization and FEC architecture are needed
  - No solution at DER = 1e-5.
  - Reduced eta_0 to half the value at 3ck
- Short channels paired with a short reference package are of particular concern
  - What is the minimum loss requirement?
- Crosstalk is dominated by OSFP connector and BGA verticals
- Next Steps
  - Explore additional minimum loss channels
  - Explore solutions options

Appendix

Channel Contribution

Thru	FEXT	NEXT1	NEXT2
C2M_PCB_10dB.s4p	C2M_PCB_10dB_FEXT1.s4p	C2M_PCB_10dB_NEXT1.s4p	C2M_PCB_10dB_NEXT2.s4p
C2M_PCB_11dB.s4p	C2M_PCB_11dB_FEXT1.s4p	C2M_PCB_11dB_NEXT1.s4p	C2M_PCB_11dB_NEXT2.s4p
C2M_PCB_12dB.s4p	C2M_PCB_12dB_FEXT1.s4p	C2M_PCB_12dB_NEXT1.s4p	C2M_PCB_12dB_NEXT2.s4p
C2M_PCB_13dB.s4p	C2M_PCB_13dB_FEXT1.s4p	C2M_PCB_13dB_NEXT1.s4p	C2M_PCB_13dB_NEXT2.s4p