IBM

Line Signaling and FEC Performance Comparison for 25Gb/s 100GbE

IEEE 802.3 100 Gb/s Backplane and Cable Task Force
Chicago, September 2011

Troy Beukema, Mounir Meghelli

Supporters and Contributors

- Mike Dudek, Qlogic Corporation
- Mark Bugg, Molex
- Peerouz Amleshi, Molex
- Myles Kimmitt, Emulex
- Ziad Hatab, Vitesse
- Frank Chang, Vitesse
- Iain Robertson, Texas Instruments
- Scott Kipp, Brocade
- Roy Cideciyan, IBM
- Barry Barnett, IBM
- Peter Pepeljugoski, IBM
- Jeffrey Lynch, IBM
- David Stauffer, IBM

Objectives

1) Determine / propose Optimal Line Code for 100G= 4x25Gb/s over backplane among (NRZ, PAM4) candidates

2) Determine / propose an efficient FEC to both increase maximum loss handling capability and increase immunity to crosstalk/reflections

Example Backplane Interconnect Topology

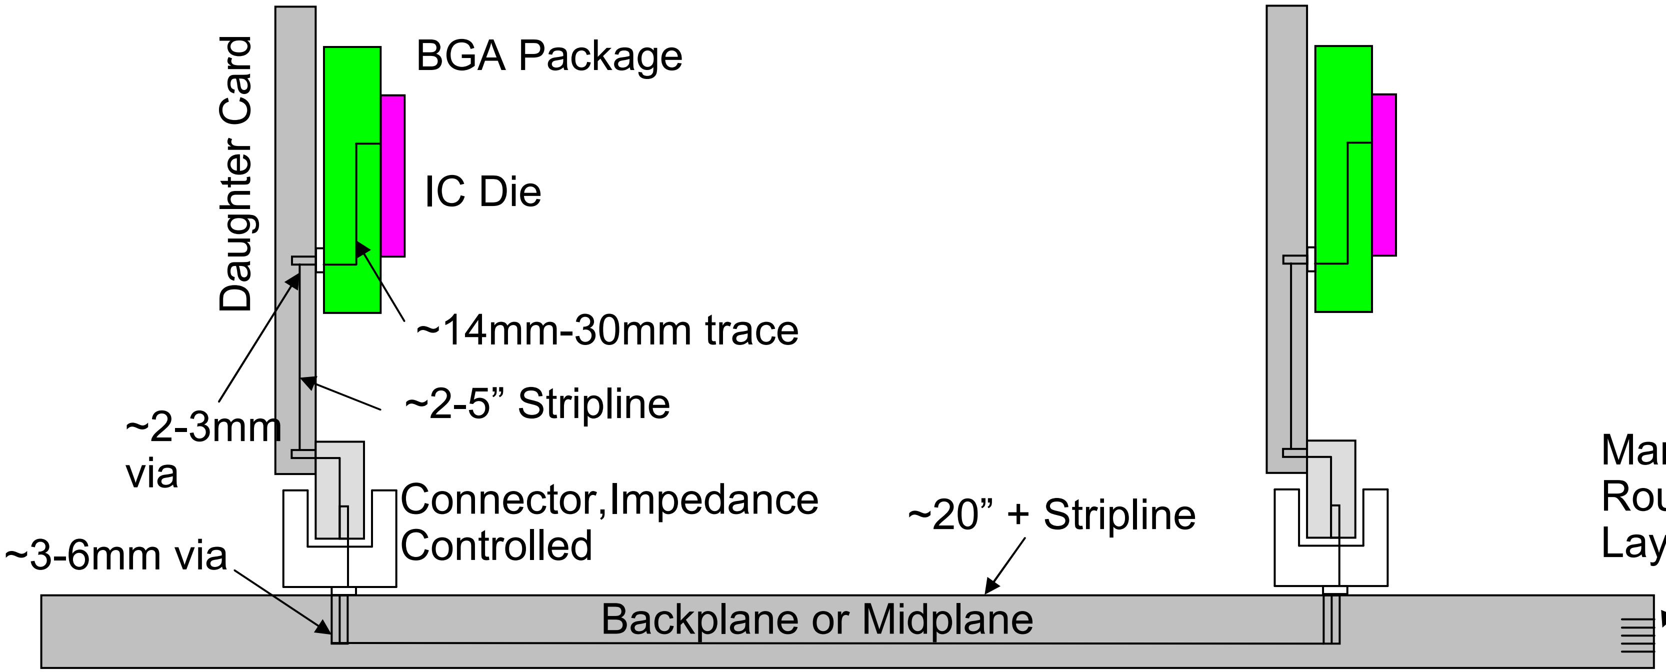

Line Signaling Simulations

(See Appendix for System and I/O Core Model)

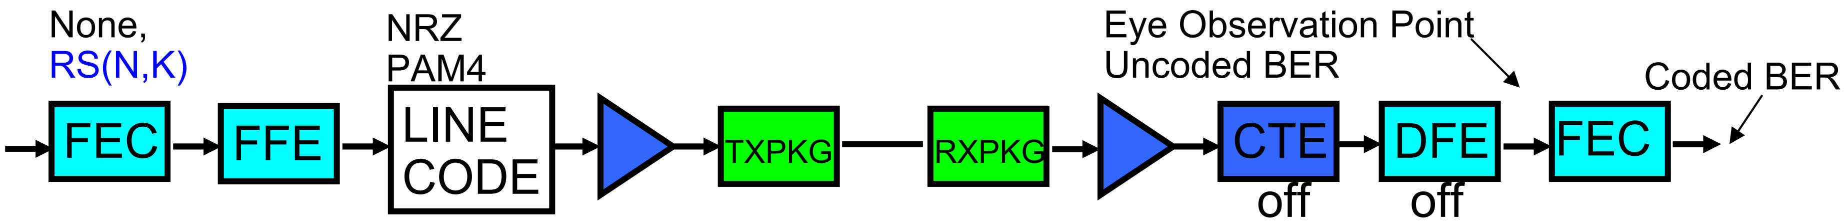

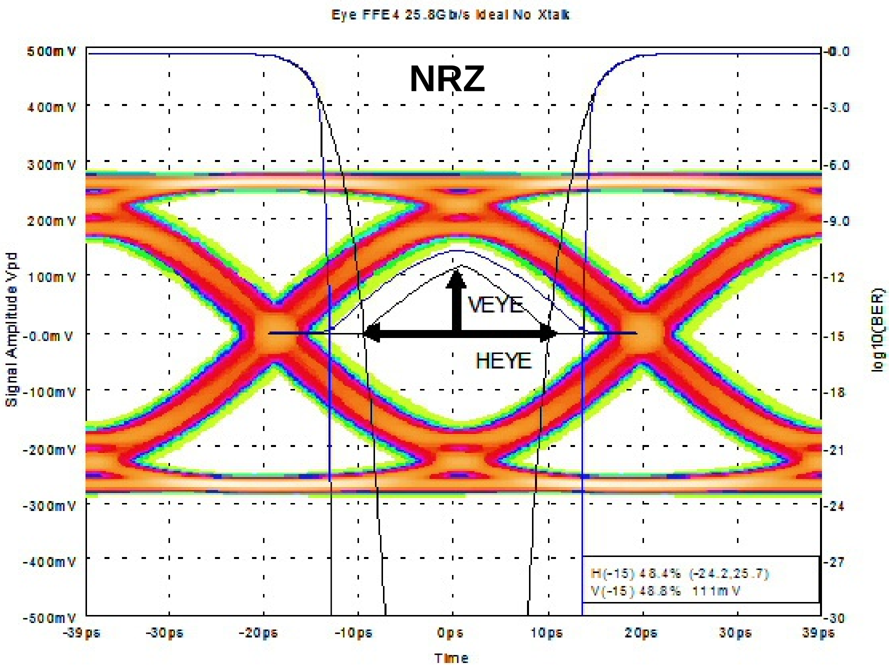

HEYEPP(1E-15)	49.9%
VEYE(1E-15)	111mV
HEYEPP(1E-15) CODED	67.5%
VEYE(1E-15) CODED	141mV

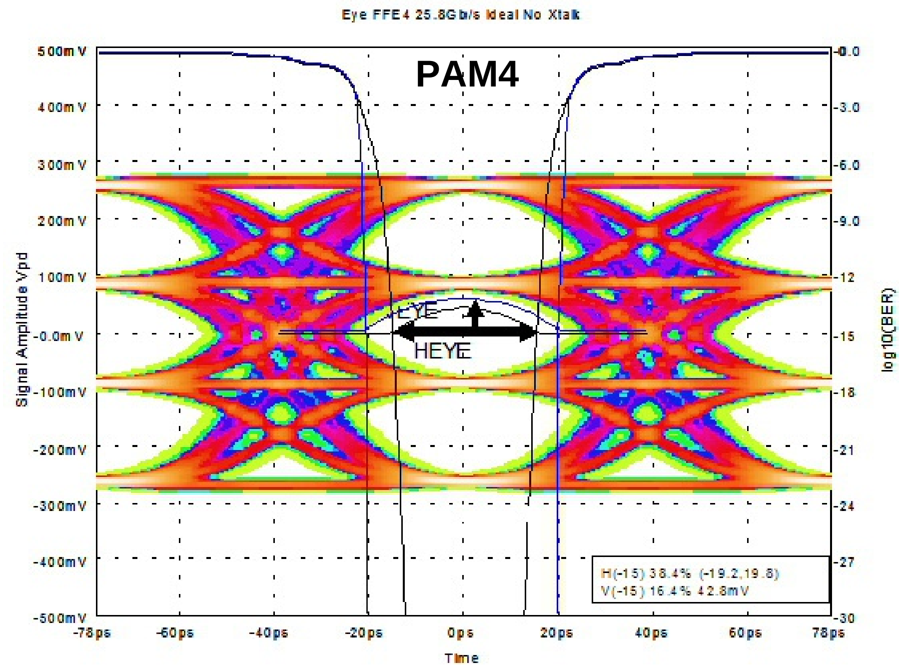

HEYEPP(1E-15)	39%
VEYE(1E-15)	42.8mV
HEYEPP(1E-15) CODED	51.4%
VEYE(1E-15) CODED	55.2mV

Simulated Block Codes

- Various FEC options have been simulated and compared
  - Highlighted FEC options are in particular of interest because of the low over-clocking penalty (0% or ≤3%)
- A DFE error propagation model has been used to determine the FEC coding gain (see Appendix)

ECC	N	K	m	T	Transcode	Line Rate[4]	Rate/156.25	Over clocking
RS[3]	**272**	**260**	**10**	**6**	**64/65**	**26.5625**	**170**	**3%**
RS[1]	224	208	10	8	64/65	27.34375	175	6.1%
RS[3]	280	260	10	10	64/65	27.34375	175	6.1%
RS[2]	**240**	**228**	**9**	**6**	**512/513**	**26.36719**	**168 + 3/4**	**2.3%**
RS[2]	244	228	9	8	512/513	26.80664	171 + 9/16	4%
RS[2]	248	228	9	10	512/513	27.24609	174 + 3/8	5.6%
RS[2]	**352**	**342**	**12**	**5**	**512/513**	**25.78125**	**165**	**0%**

(1) bhoja_01_0911.pdf
(2) cideciyan_01_0911.pdf
(3) Proposed by John Ewen, IBM
(4) Line Rate = N / K / Transcode * 25.0

IBM Experimental Test Fixture Backplane Channels

25dB Loss Channel

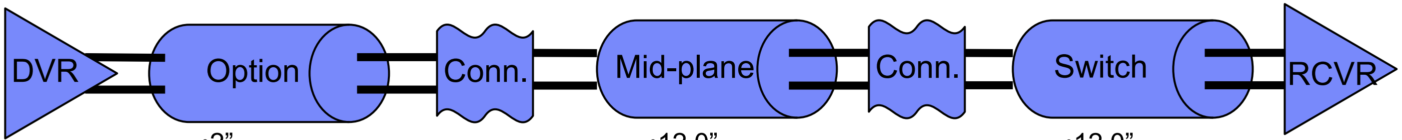

THRU.s4p
FEXT1.s4p FEXT5.s4p
FEXT2.s4p FEXT6.s4p
FEXT3.s4p FEXT7.s4p
FEXT4.s4p FEXT8.s4p

Option:
- 2"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110 ohm

Mid-plane:
- 12.0"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 90 ohm

Switch:
- 12.0"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110ohm

30dB Loss Channel

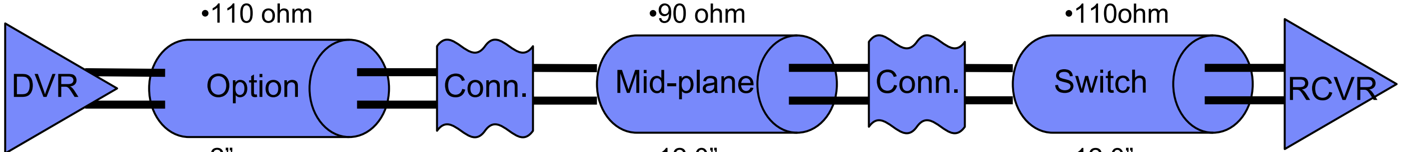

THRU.s4p
FEXT1.s4p FEXT5.s4p
FEXT2.s4p FEXT6.s4p
FEXT3.s4p FEXT7.s4p
FEXT4.s4p FEXT8.s4p

Option:
- 2"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110 ohm

Mid-plane:
- 12.0"
- IS415
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 90 ohm

Switch:
- 12.0"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110ohm

35dB Loss Channel

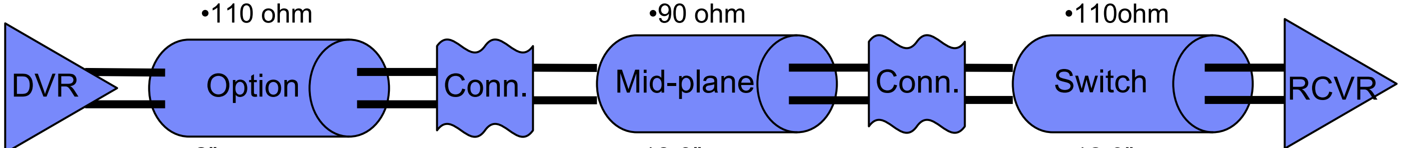

THRU.s4p
FEXT1.s4p FEXT5.s4p
FEXT2.s4p FEXT6.s4p
FEXT3.s4p FEXT7.s4p
FEXT4.s4p FEXT8.s4p

Option:
- 2"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers

Mid-plane:
- 18.0"
- IS415
- 220mil thick,
- 6/6/6 mils,
- 28 layers

Switch:
- 12.0"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers

40dB Loss Channel

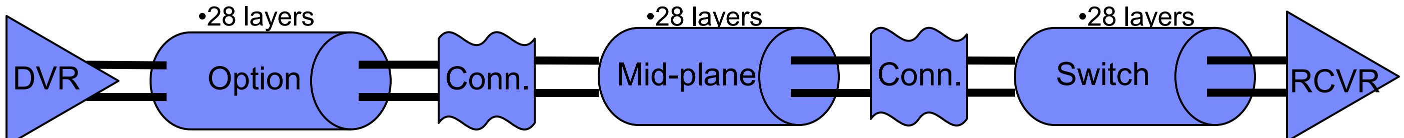

THRU.s4p
FEXT1.s4p FEXT5.s4p
FEXT2.s4p FEXT6.s4p
FEXT3.s4p FEXT7.s4p
FEXT4.s4p FEXT8.s4p

Option:
- 2"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110 ohm

Mid-plane:
- 18.0"
- IS415
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 90 ohm

Switch:
- 18.0"
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 110ohm

Typical Production Design Build Construction: 10% impedance tolerance, Standard Copper Foil, Backdrill 10mil +/- 10mil

High Channel Loss Eye Diagrams (30dB)

NRZ

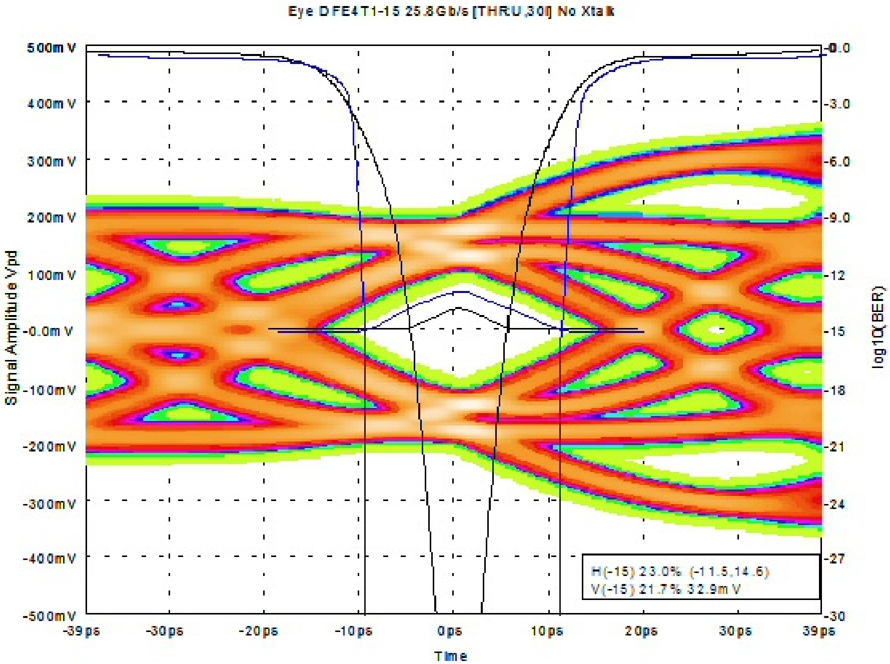

BAUD/2 LOSS (CHAN/LINK)	30/38dB
HEYEPP(1E-15)	26.1%
VEYE(1E-15)	32.9mV
BAUD/2 LOSS RS(255,239)	31/41dB
HEYEPP(1E-15) RS(255,239)	53.3%
VEYE(1E-15) RS(255,239)	71.6mV

PAM-4

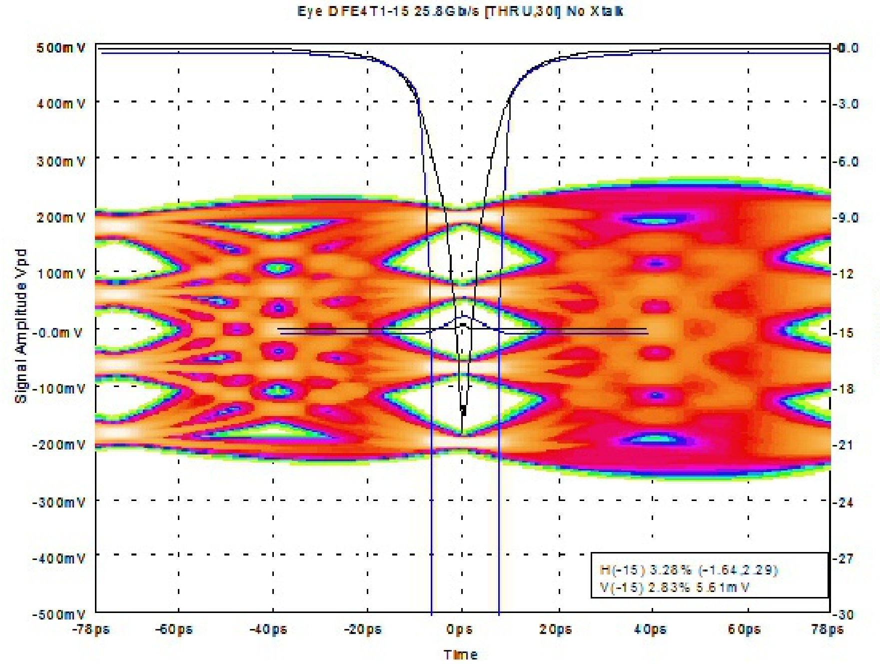

BAUD/2 LOSS (CHAN/LINK)	17/20dB
HEYEPP(1E-15)	4.0%
VEYE(1E-15)	4.9mV
BAUD/2 LOSS RS(255,239)	18/21dB
HEYEPP(1E-15) RS(255,239)	18.1%
VEYE(1E-15) RS(255,239)	25.2mV

High Channel Loss Eye Diagrams (~35dB)

NRZ

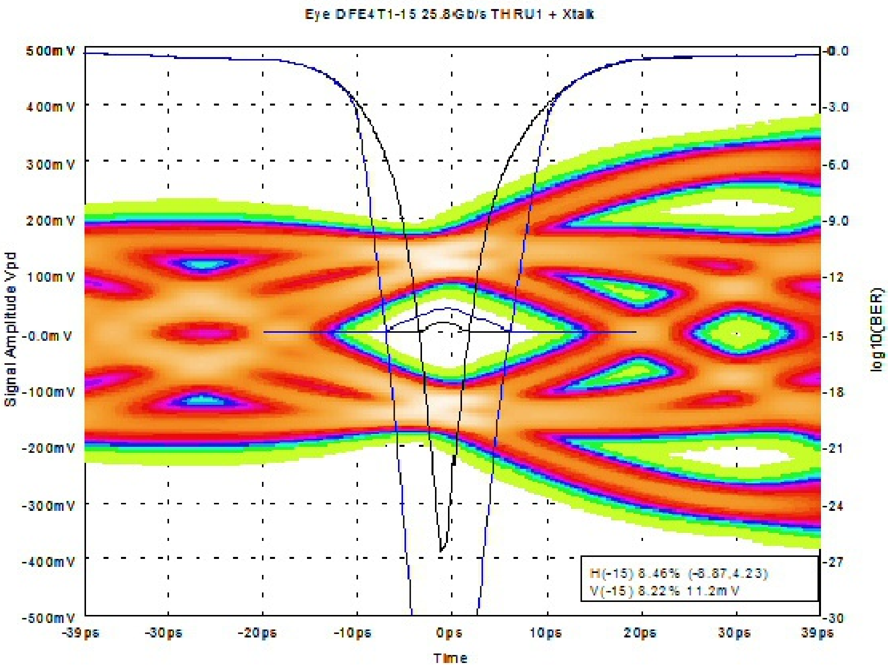

BAUD/2 LOSS (CHAN/LINK)	37/45dB
HEYEPP(1E-15)	13.1%
VEYE(1E-15)	11.2mV
BAUD/2 LOSS RS(352,342)	37/45dB
HEYEPP(1E-15) RS(352,342)	33.2%
VEYE(1E-15) RS(352,342)	38.8mV

PAM-4

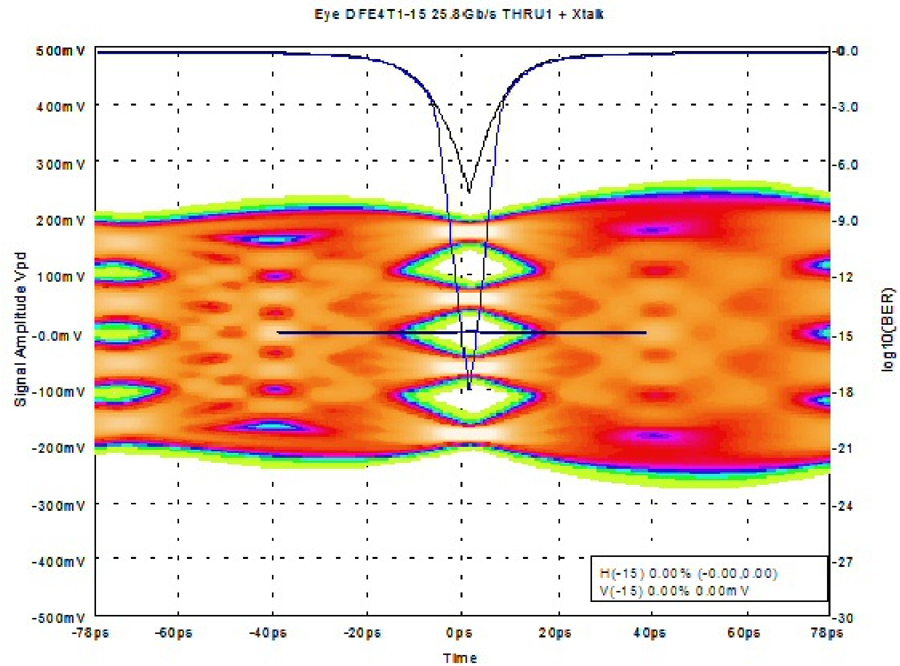

BAUD/2 LOSS (CHAN/LINK)	19/22dB
HEYEPP(1E-15)	0.0%
VEYE(1E-15)	0mV
BAUD/2 LOSS RS(352,342)	19/22dB
HEYEPP(1E-15) RS(352,342)	3.8%
VEYE(1E-15) RS(352,342)	1.8mV

NRZ HEYE and VEYE vs. Channel Loss, 64b/65b Transcode

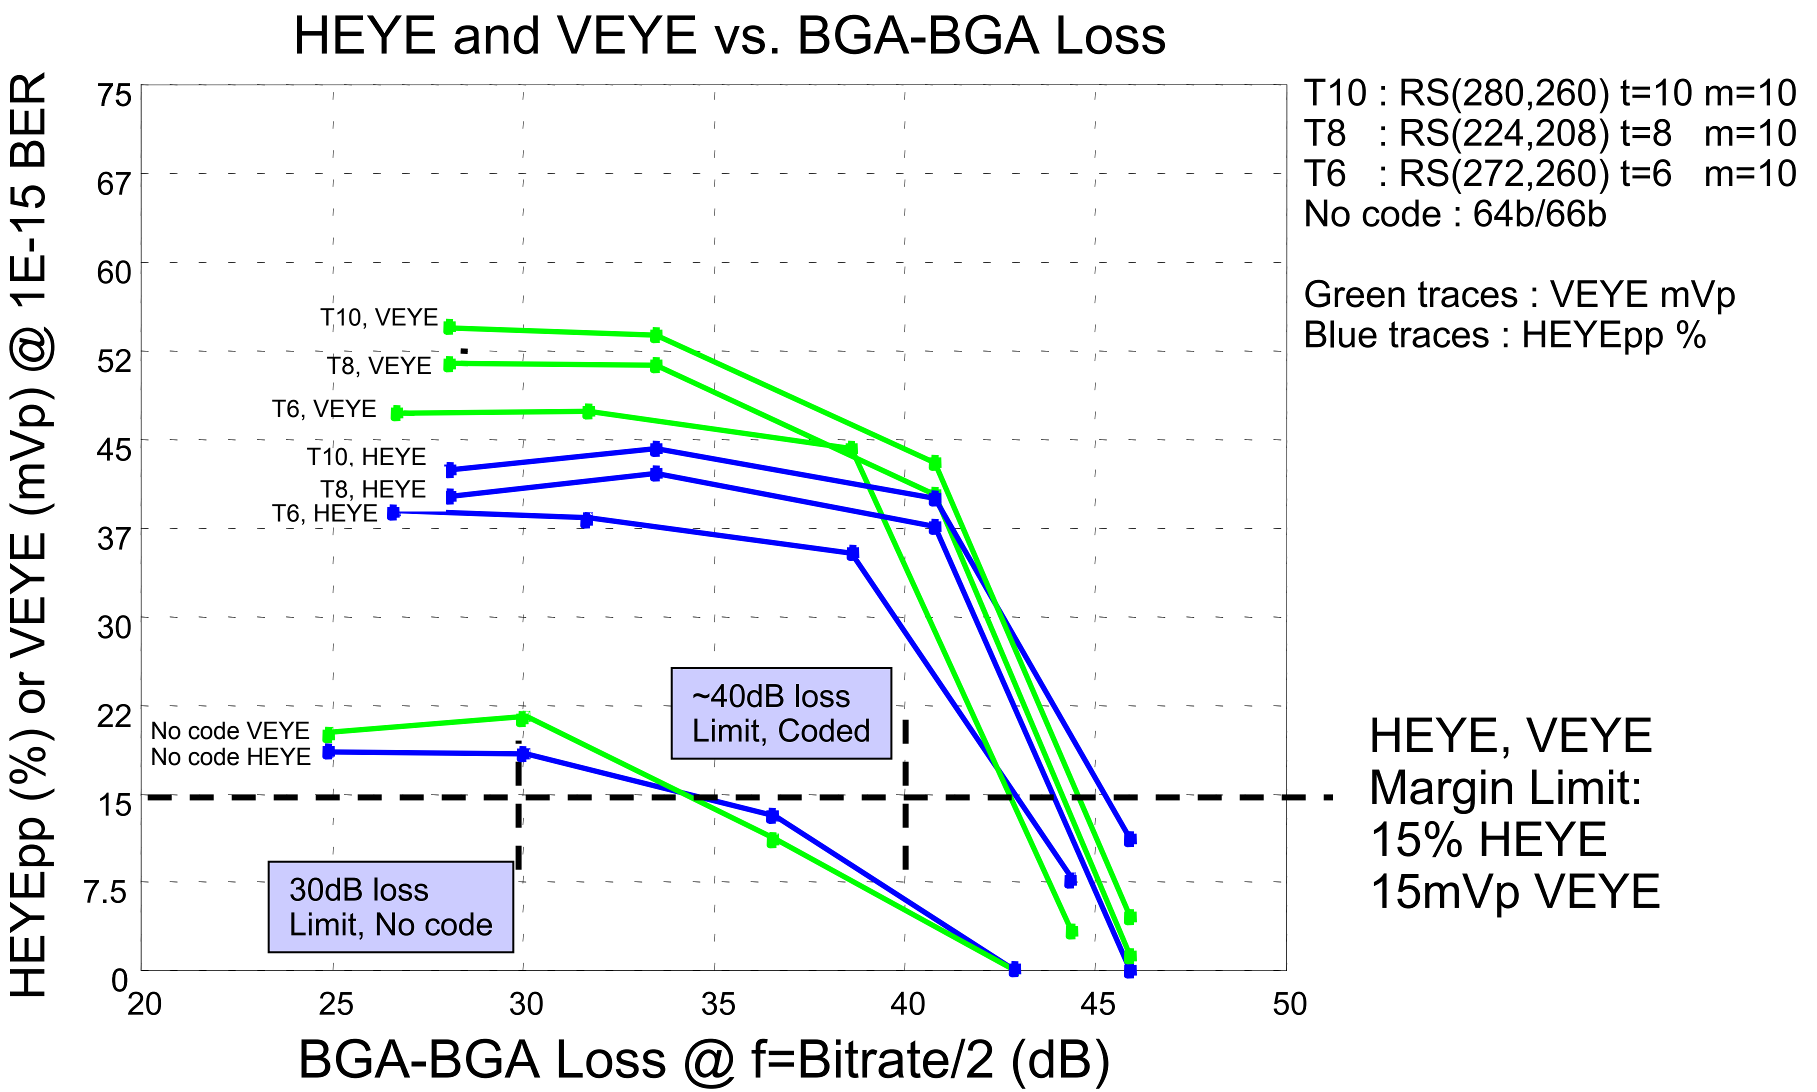

NRZ HEYE and VEYE vs. Channel Loss, 512b/513b Transcode

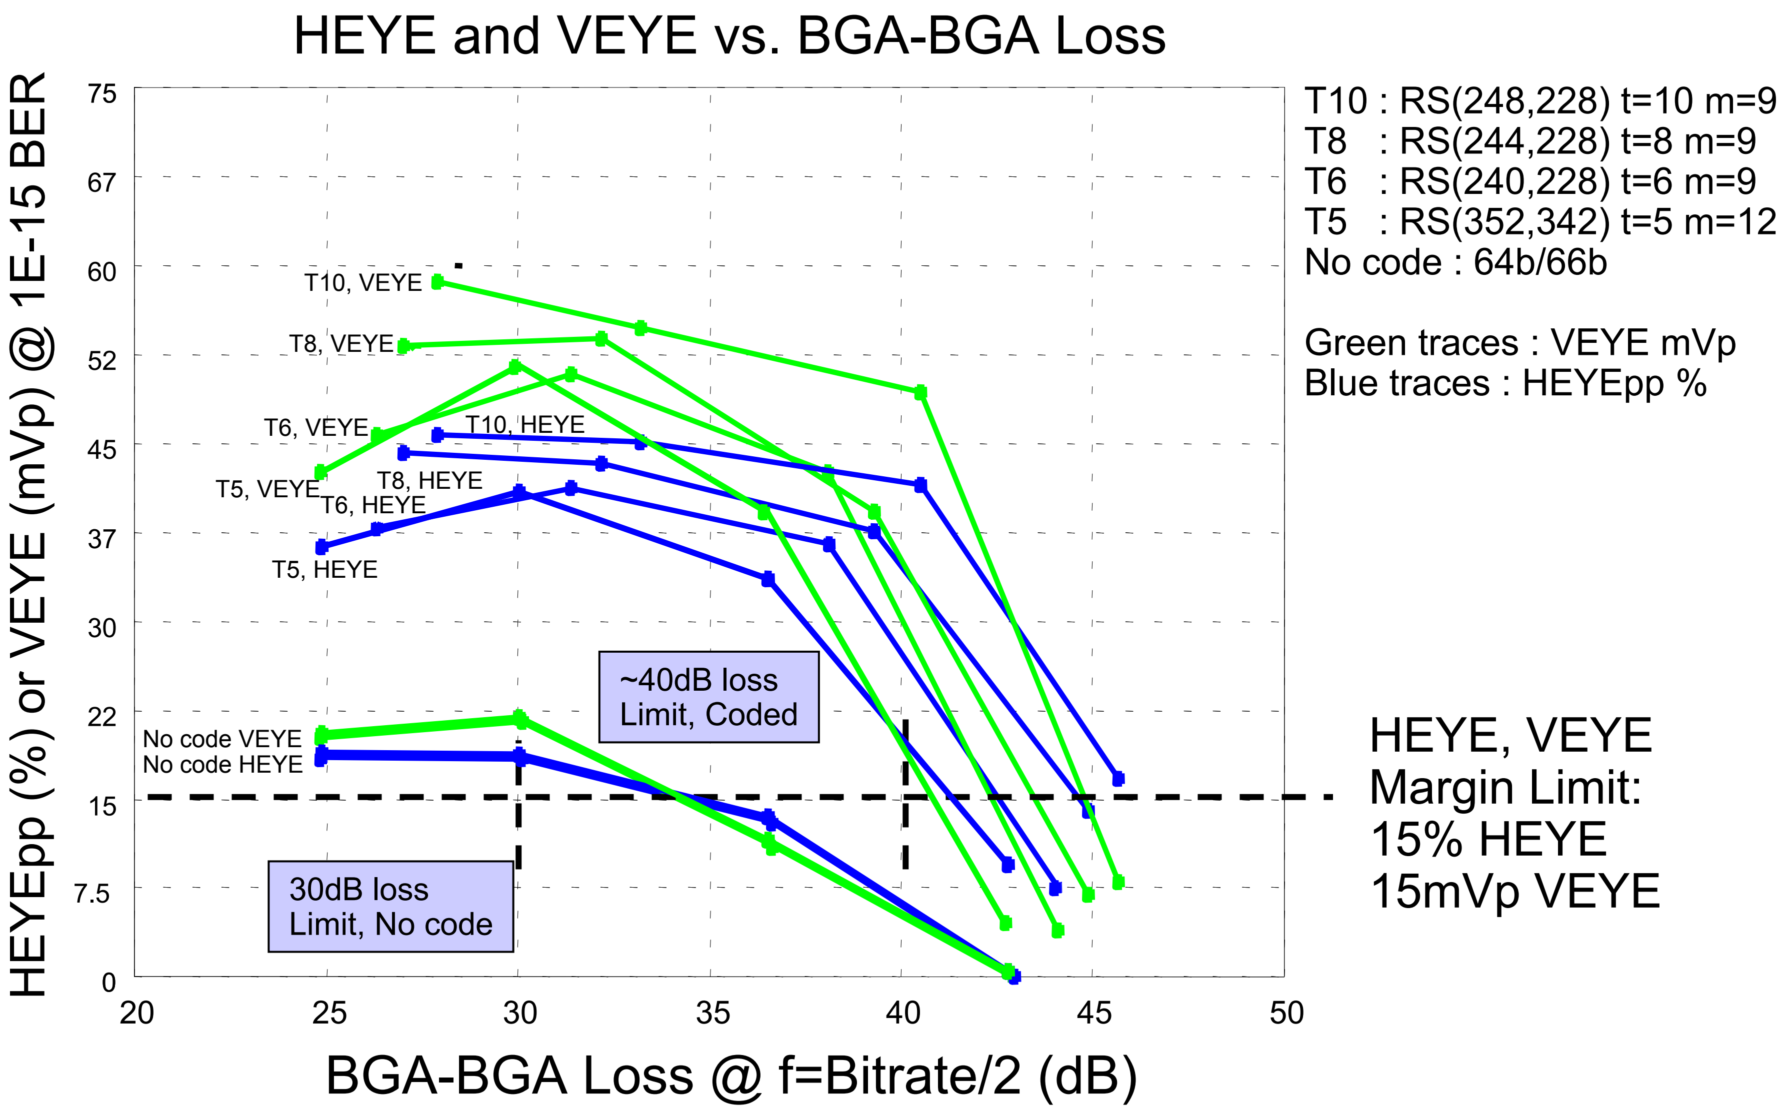

PAM4 HEYE and VEYE vs. Channel Loss, 64b/65b Transcode

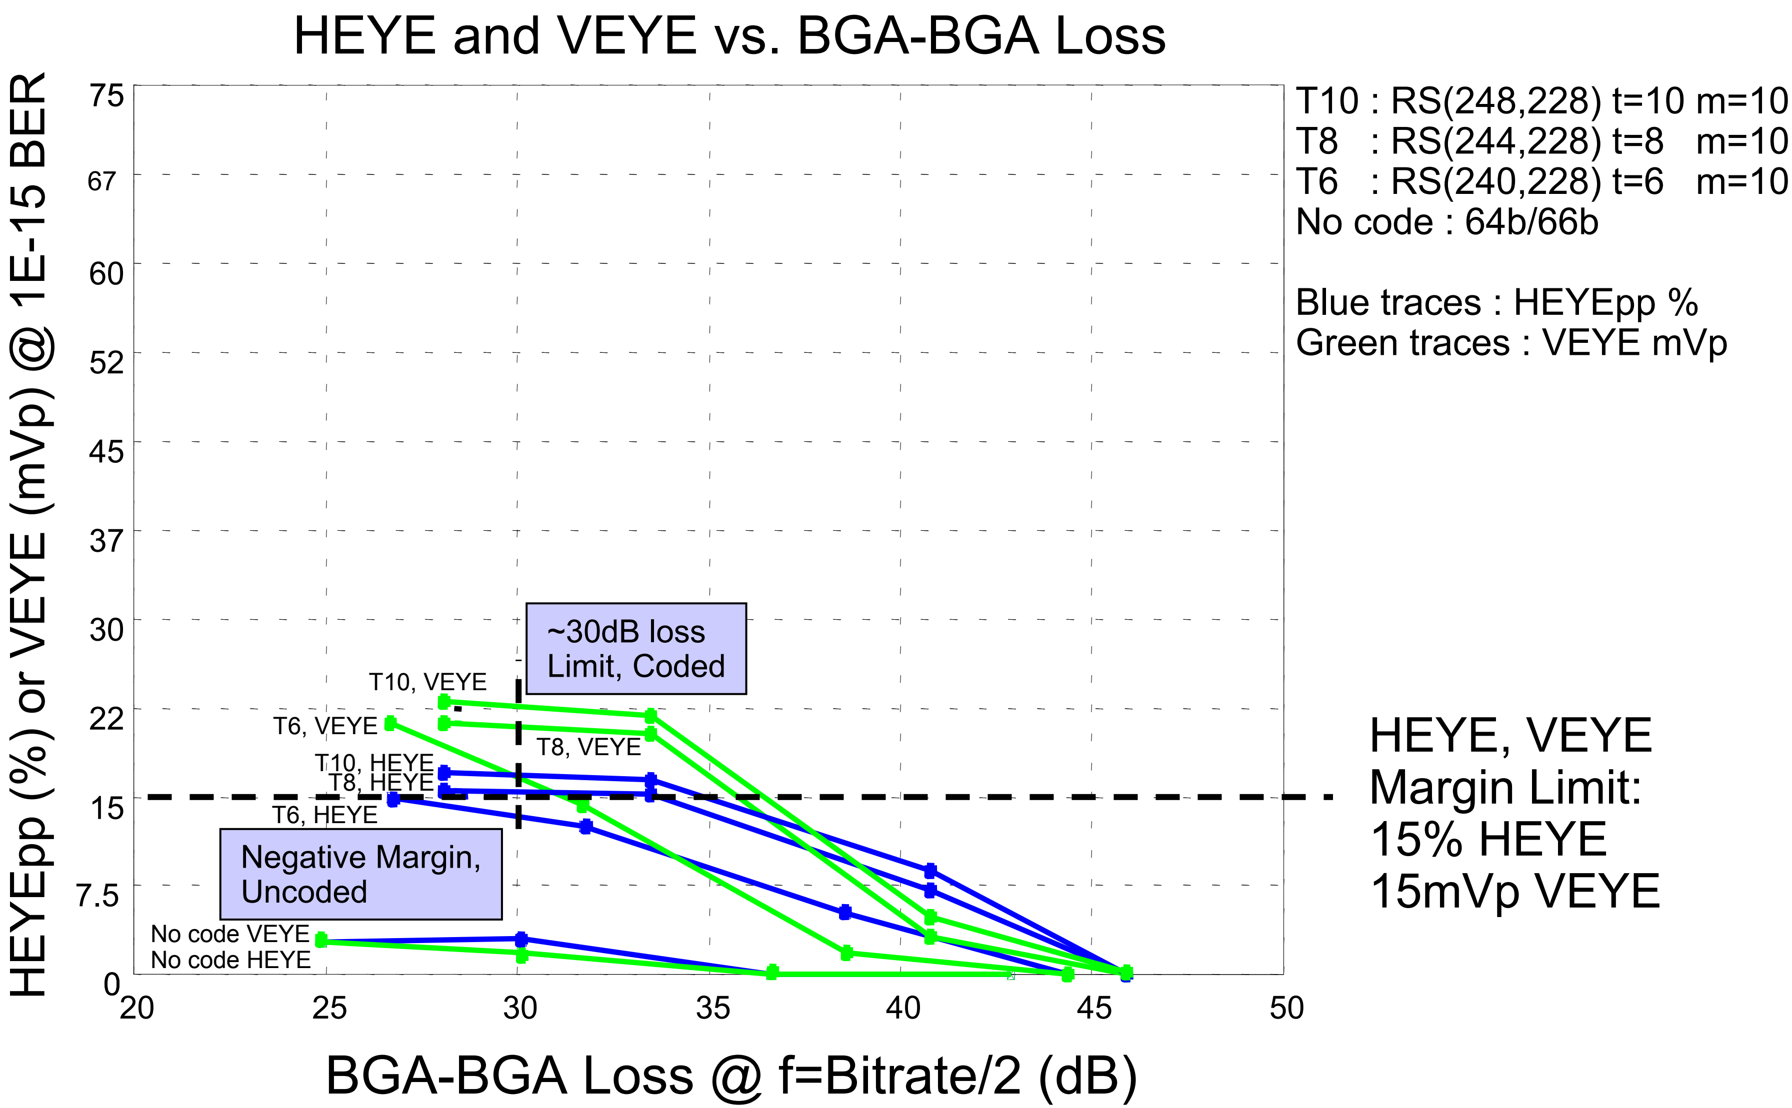

Summary/Conclusions

- Signal Integrity simulation results showing that NRZ line signaling is far superior to PAM4 line signaling up to BGA-BGA channel losses of ~40dB
  - ✓ **NRZ line signaling is proposed for the 100GbE BP/Cable PHY**
- Uncoded NRZ operates to about 30dB BGA-BGA channel loss limit. To increase loss handling capability, a T=5 RS Code is sufficient to enable operation on high loss (>30dB, <40dB Loss at bitrate/2) backplane channels
  - ✓ **Medium-strength RS codes with low (<3%) or no over-clocking (such as T=5 or T=6 RS Code with 512b/513b transcoding) is proposed for the 100GbE BP/Cable FEC**
- Reduced latency FEC encoding can be made optional or bypassed to eliminate power draw when operating on easy channels

Appendix

- I/O System Model
- Reference Model I/O Core parameters
- Reference Package Model
- Measured 30dB Channel Response
- DFE Error Propagation Impact
  - Pc vs. Pb transfer functions incorporating error propagation

I/O System Model

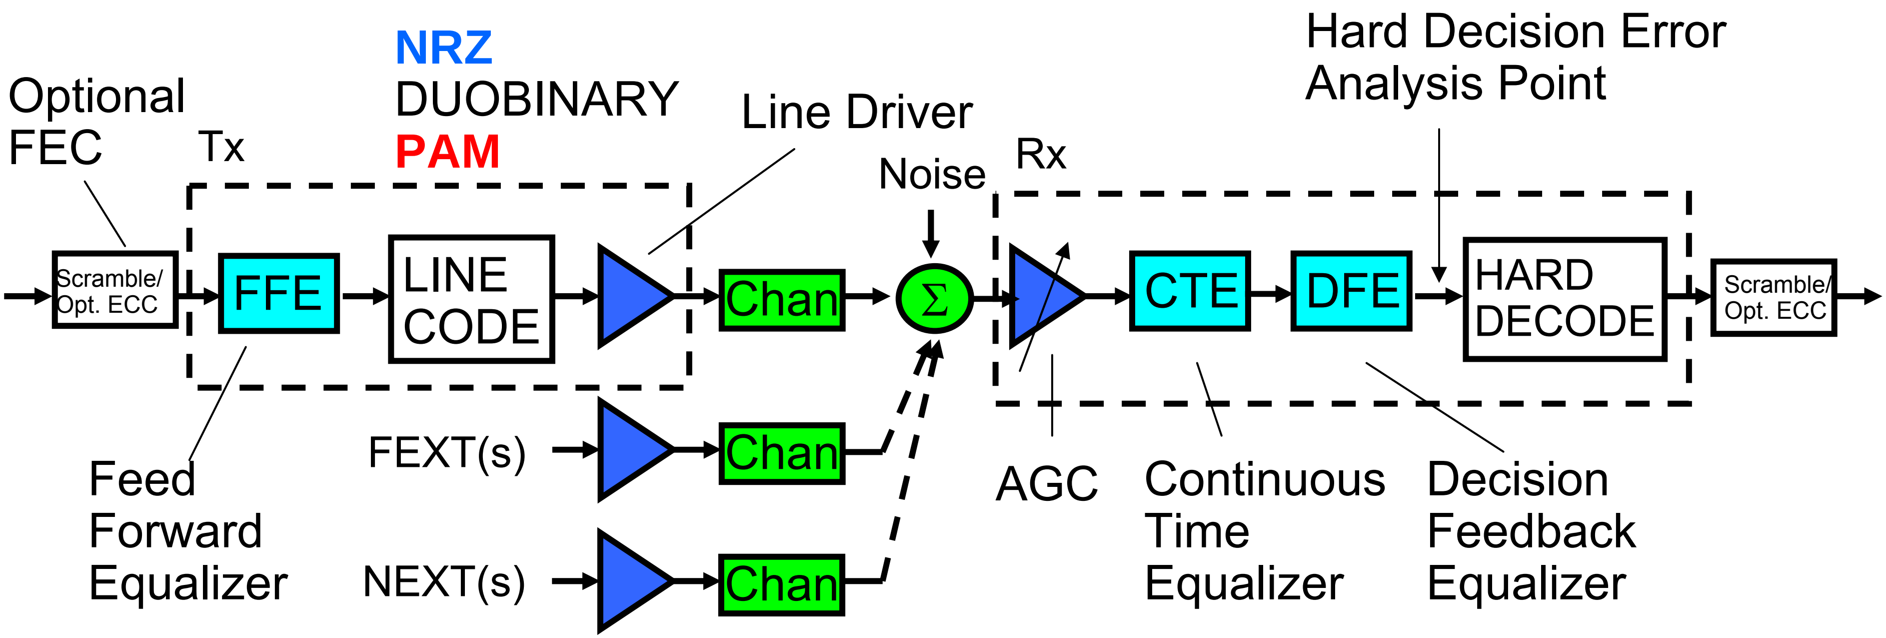

TX + RX POWER EFFICIENCY TARGETS : 20-25mW / Gb/s

EXAMPLE I/O POWER FOR HIGH DENSITY APPLICATION :

128 FULL DUPLEX I/O

20mW/Gb/s * 25Gb/s per I/O Lane = 500mW/Lane

500mW/Lane * 128 Lanes = <u>64W for I/O ALONE</u>

Reference Model I/O Core Parameters

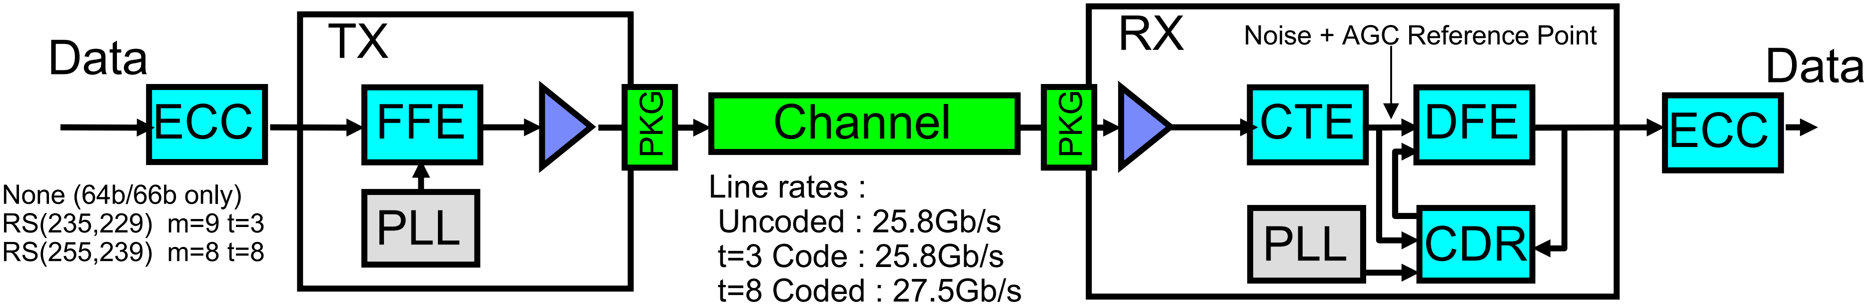

PARAMETER	VALUE NRZ	VALUE PAM4
PEAK SWING	1000mVppd	1000mVppd
RJ	350fs RMS	350fs RMS
DCD	1.6% (49.2:50.8)	0
SJ	5% UI	5%
BW	-1.5dB@13GHz 2 pole Bessel	-1.5dB @13GHz 2 pole Bessel
PKG	-2.6dB@13GHz	-2.6dB @13GHz
FFE	4 tap 2 precursor	4 tap 2 precursor

- Simplified T & R model
- Parameters selected to approximate "real" hardware realization performance
- Set up to "favor" PAM4 : 2x complex DFE
- Target BER = 1E-15
- E/L CDR active for NRZ & PAM4

PARAMETER	VALUE NRZ	VALUE PAM4
NOISE@ SLICER	2.75mV RMS	2mV RMS
Sensitivity@ SLICER	20mVpd	20mVpd
AGC LEVEL	280mVpd	280mVpd
AGC GAIN MAX	3	3
RJ	350fs RMS	350fs RMS
SJ	5% UI	0%
BW	-1.5dB @13GHz 4 pole Bessel	-1.5dB @13GHz 4 pole Bessel
PKG	-2.7dB @ 13GHz	-2.7dB @ 13GHz
CTE	12dB peak @ 13GHz 3 pole 2 zero	12dB peak @ 13GHz pole 2 zero
DFE	15 tap	15 tap (2X NRZ)

Reference Package Model

PKG Identifier	Trace Length mm	Zo ohm	Loss 25G dB
PKG25	21	100 T 100 R	2.6

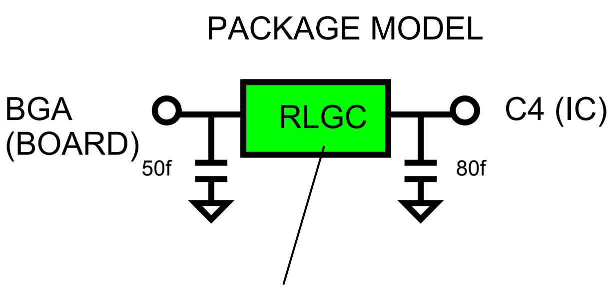

SINGLEWIRE RLGC MODEL :

Ro = 1 ohm/m
Lo = 350 nH/m
Go = 0
Co = 140pF/m
Rf = 5e-3
Gf = 25e-12

R(f) = Ro + Rf * sqrt(f)
L(f) = Lo
G(f) = Go + Gf * f
C(f) = Co

Note : correct causality manually

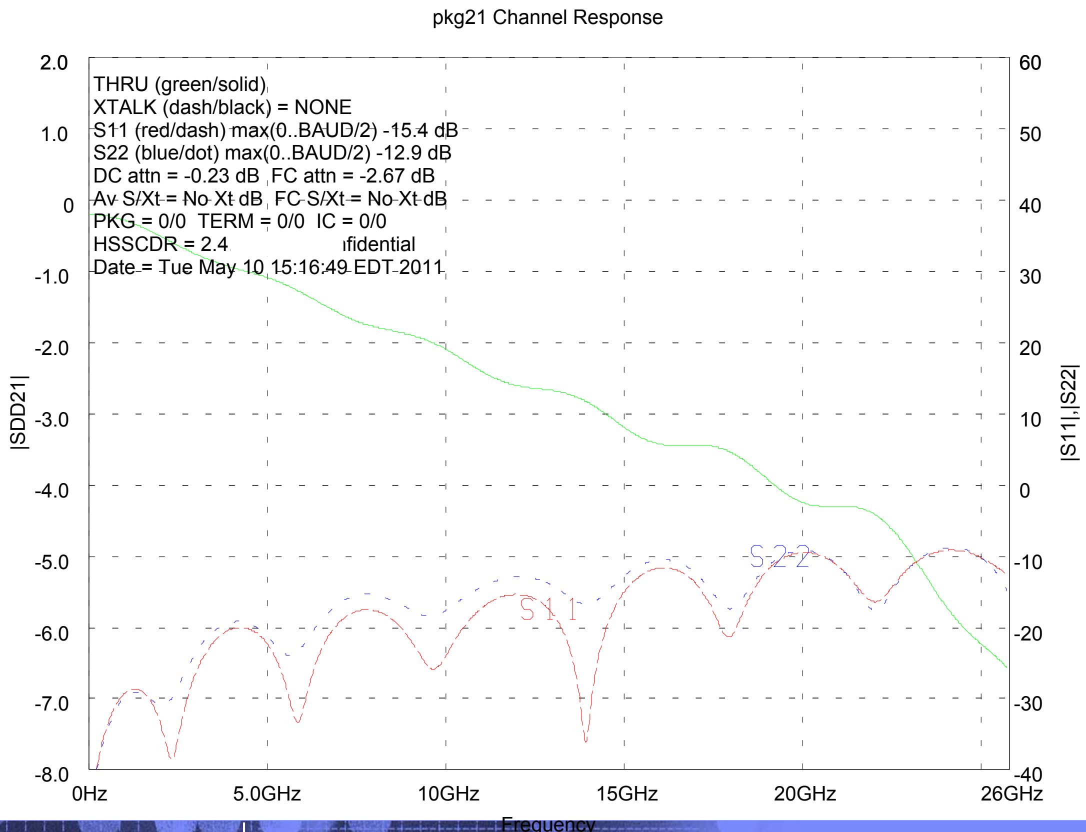

Measured 30dB Channel Response

30dB Loss Channel

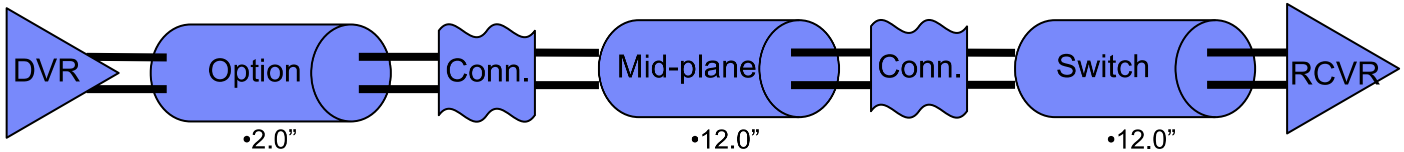

Adds Crosstalk + Reflections

Option:
- 2.0”
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 12 layers
- 110 ohm

Mid-plane:
- 12.0”
- IS415
- 220mil thick,
- 6/6/6 mils,
- 28 layers
- 90 ohm

Switch:
- 12.0”
- Meg 6
- 220mil thick,
- 6/6/6 mils,
- 14 layers
- 110ohm

XFER

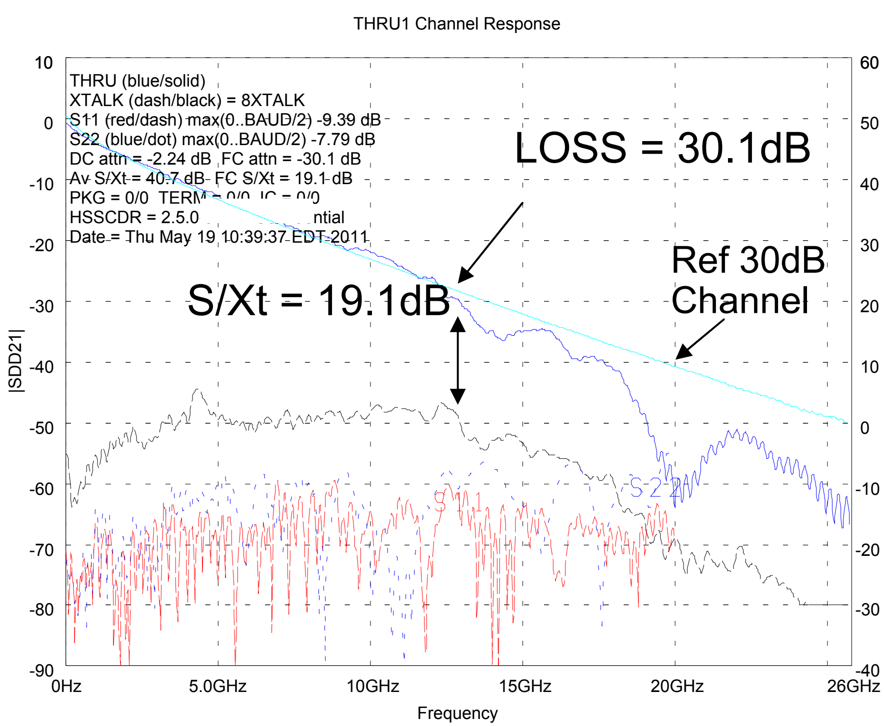

Impulse Response

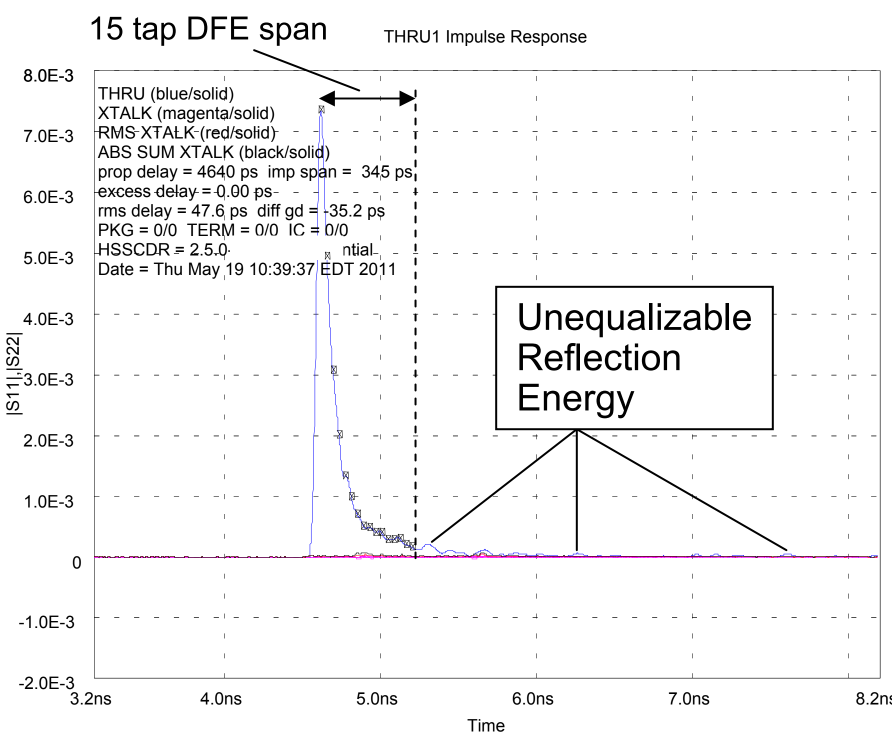

DFE Error Propagation Impact

Error propagation model :

DFE h1 tap at 0.65, single DFE-tap approximation
$p(L) =\sim (0.47)^{L-1}$ (1)
where p(L) is probability of a burst-error length of L bits

DFE error propagation impact :

RS(280,260) m=10 T=10 approximated by RS(152,140) m=10 T=6

RS(224,208) m=10 T=8 approximated by RS(130,120) m=10 T=5

RS(272,260) m=10 T=6 approximated by RS(178,170) m=10 T=4

RS(352,342) m=12 T=5 approximated by RS(198,192) m=12 T=3

(1) see “DFE Burst Errors” Slide in “FEC Proposal for NRZ Modulation”, S. Bhoja, et. al., 100Gb/s Backplane and Cable Task Force, IEEE 802.3, Chicago, Sept. 2011

RS(280,260) m=10 T=10 Coding Transfer with DFE Error Propagation

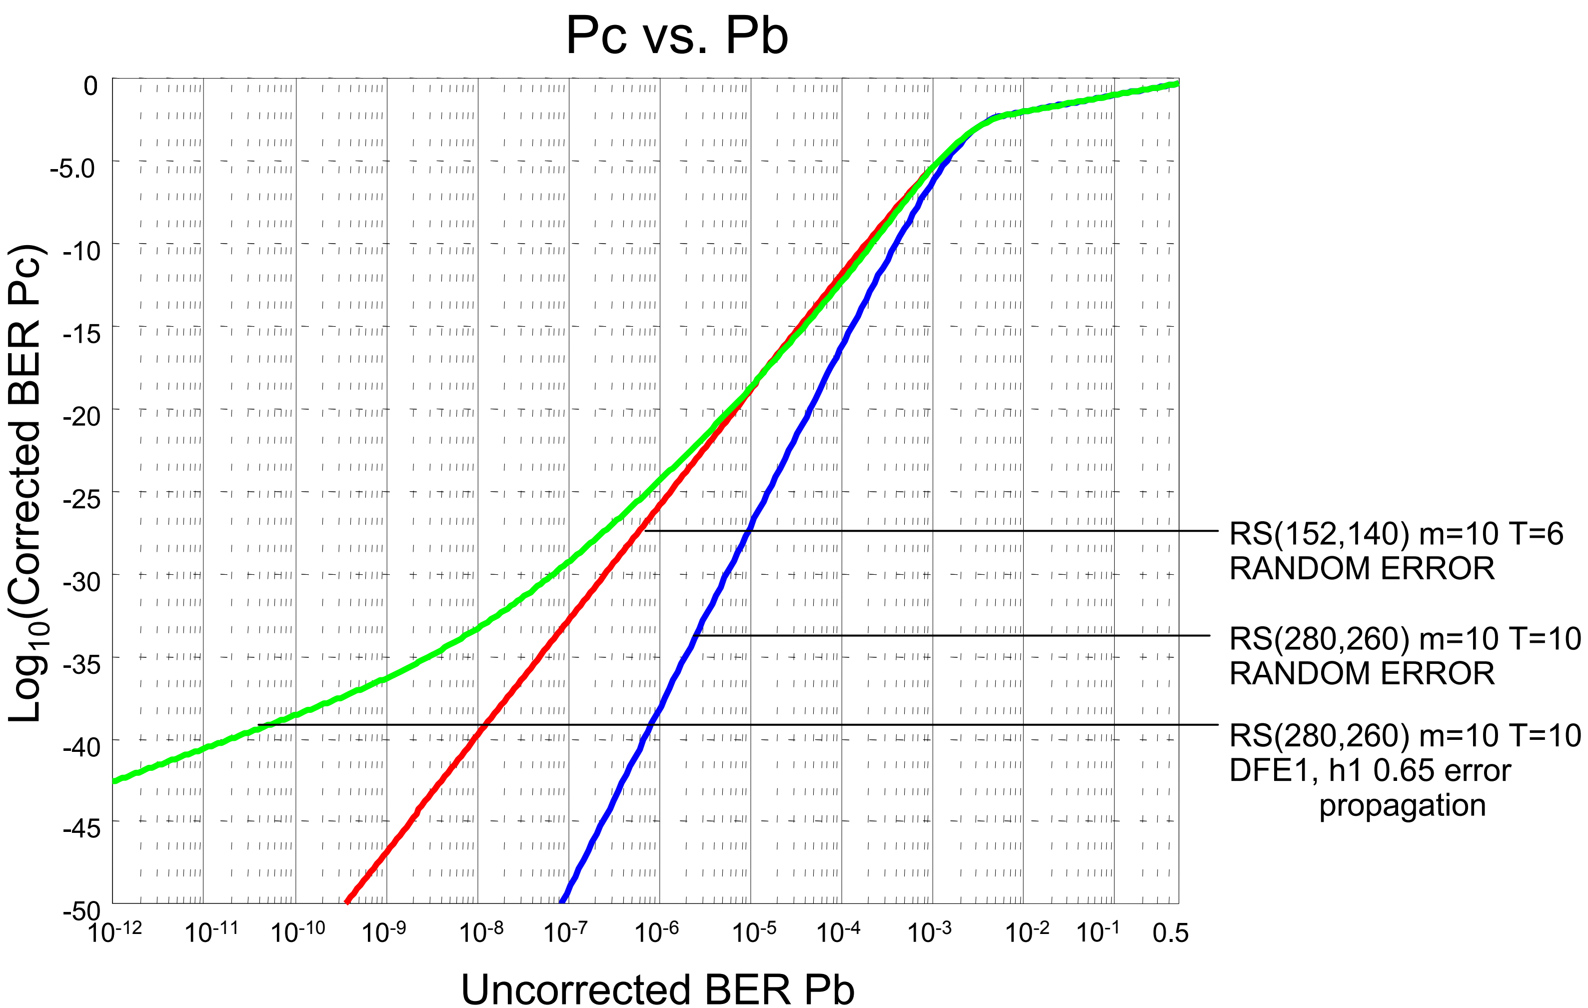

Error propagation curve derived by John Ewen, IBM

RS(224,208) m=10 T=8 Coding Transfer with DFE Error Propagation

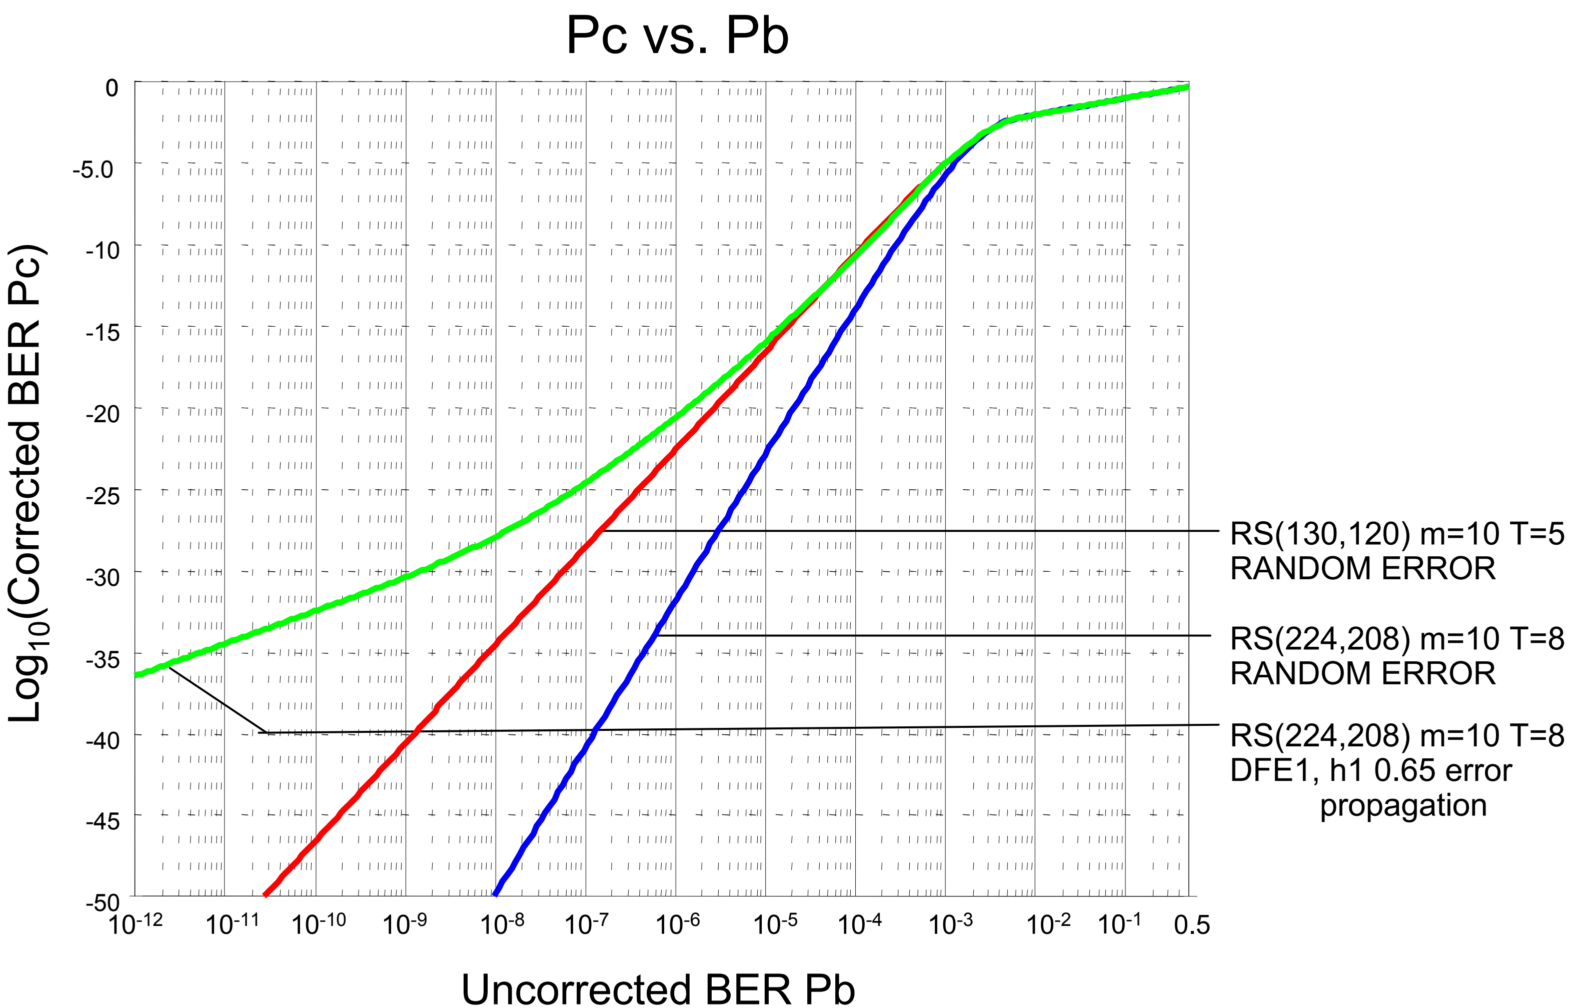

Error propagation curve derived by John Ewen, IBM

RS(272,260) m=10 T=6 Coding Transfer with DFE Error Propagation

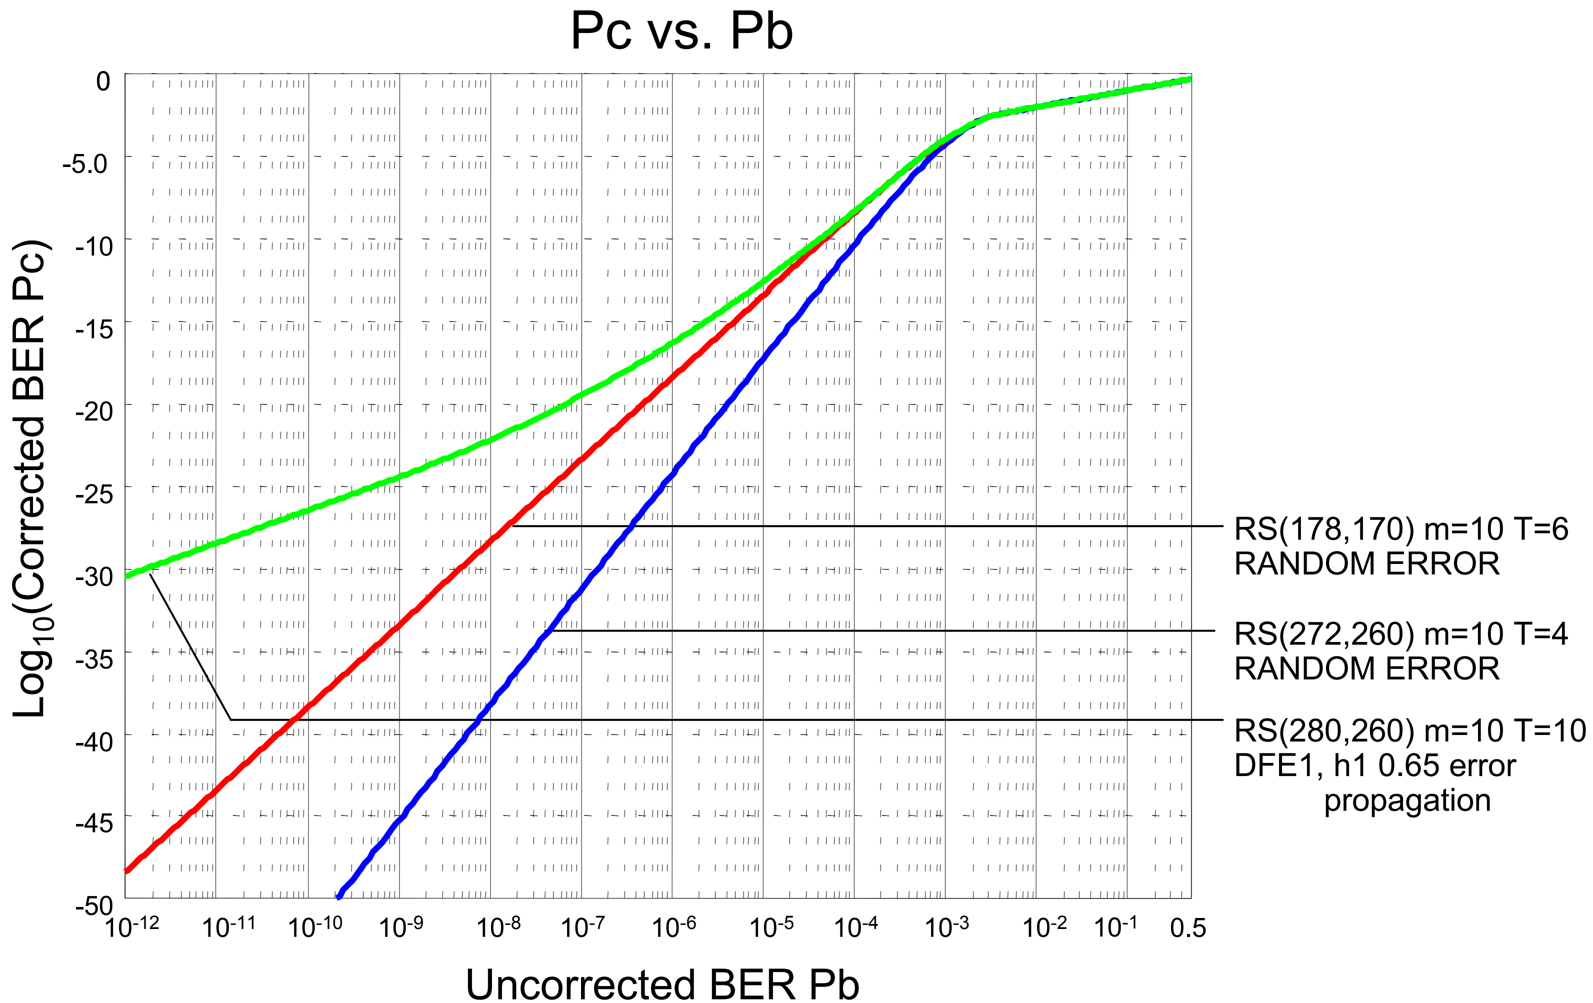

Error propagation curve derived by John Ewen, IBM

RS(352,342) m=12 T=5 Coding Transfer with DFE Error Propagation

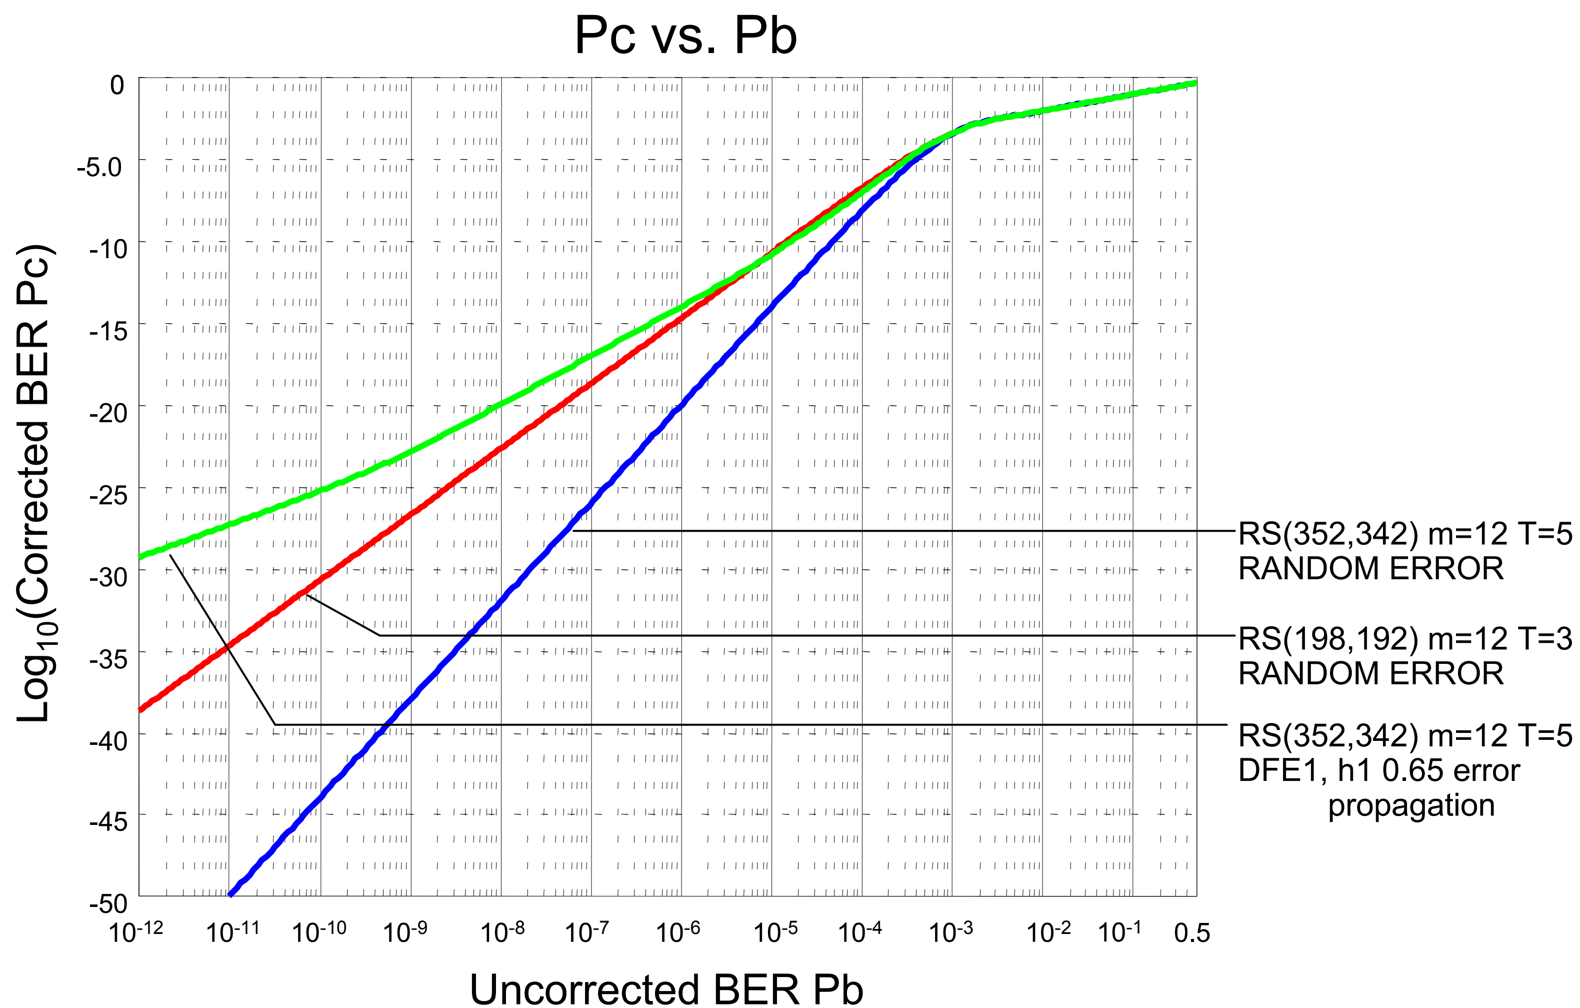

Error propagation curve derived by John Ewen, IBM