100G MMF Reach Objective

John Petrilla: Avago Technologies

March 2012

Outline & Acknowledgements

- Conclusions
- Block Diagrams
- Economic Feasibility
- Scaling 40GBASE-SR4 for a low-cost-low-power 100G SR4
- Reach and equalization

Acknowledgements

- The Kolesar Kalculator (Kolesar_Kalculator_2012_02_29.xls) was used to provide comparative total data center costs for various scenarios of relative module costs and supported reaches.
- The cable plant model and transceiver relative cost model was a collaboration with Doug Coleman (Corning).
- The link model, 32GFC 3T FFE v1a, was developed by David Cunningham (Avago Technologies).

Conclusions

•A low-cost-low-power transceiver based on unretimed interfaces and no optical link equalization does not appear technically feasible at this time.

•Transceivers based on retimed interfaces and some level of optical link equalization supporting OM4 reaches in the range of 70 m to 100 m with relative cost factors in the range of 0.7x to 1.0x SR10 appear to offer a total cost advantage relative to the incumbent case of SR10 and LR4 for server-to-switch segments with no intermediate reach SMF option and for all segments if an intermediate reach SMF option such as PSM4 is available and included.

•OM4 reaches above 50 m are expected to require some level of equalization while above 100 m may not be equalized by a 3-tap FFE or DFE(3,1).

•The MMF reach objective should not be more than 100 m of OM4.

•The effect of FEC implemented outside the transceiver in not included in the above conclusions. If available and implemented outside the transceiver, the link would appear to have a larger signal budget and noise penalties would be reduced. Beneficial effects would include pushing out the noise walls that are limiting the optical link equalizers.

100G 25G/Lane Parallel MM Transceiver: Description

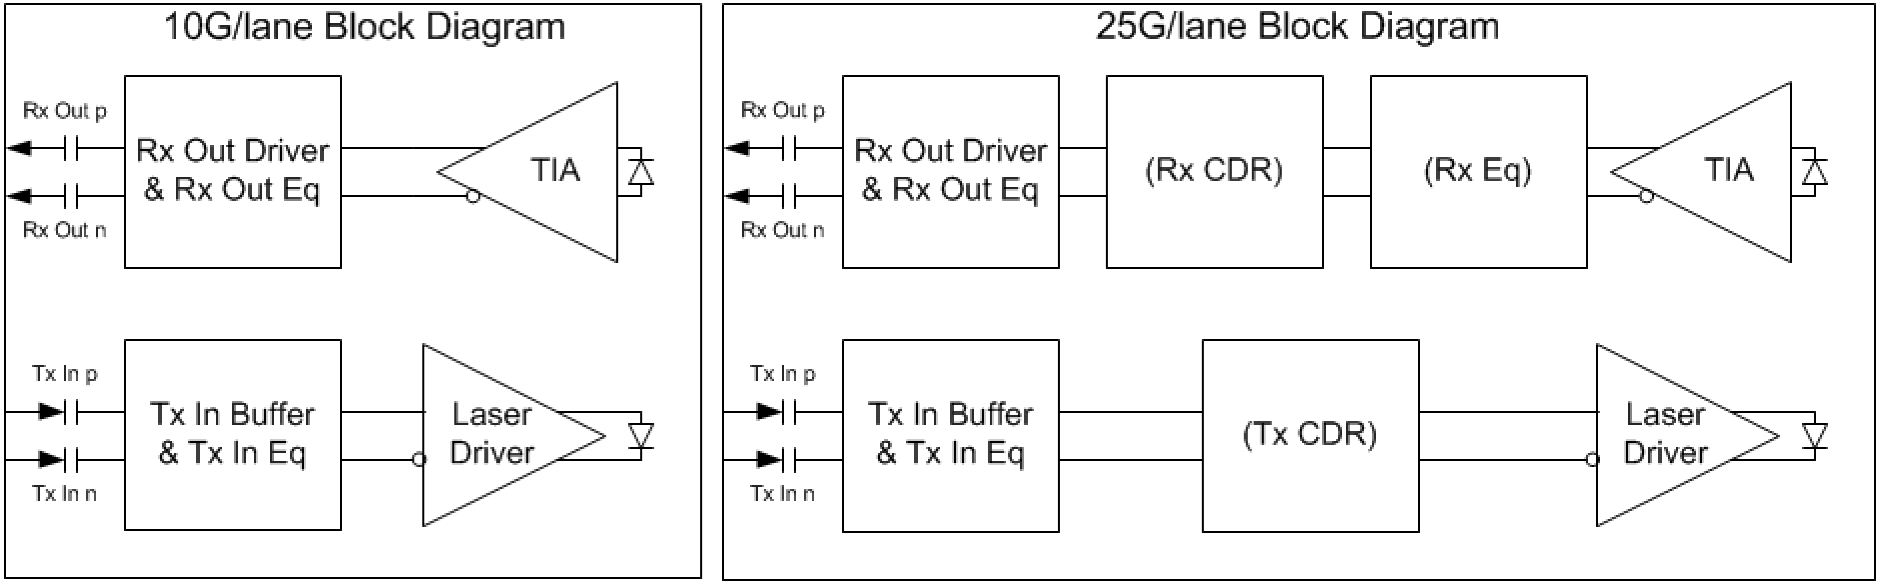

- NRZ modulation and 64b/66b encoding are assumed for both electrical and optical signals.
- MPO connectors are assumed for the optical interface.
- At 10G simple Tx input equalization and/or Rx output equalization (de-emphasis) may appear.
- At 25G electrical interfaces are expected to require equalization and, at least initially, retiming.
- At 25G some level of equalization is expected for the optical channel elements.

100G 10G & 25G/Lane MMF Transceivers Comparison

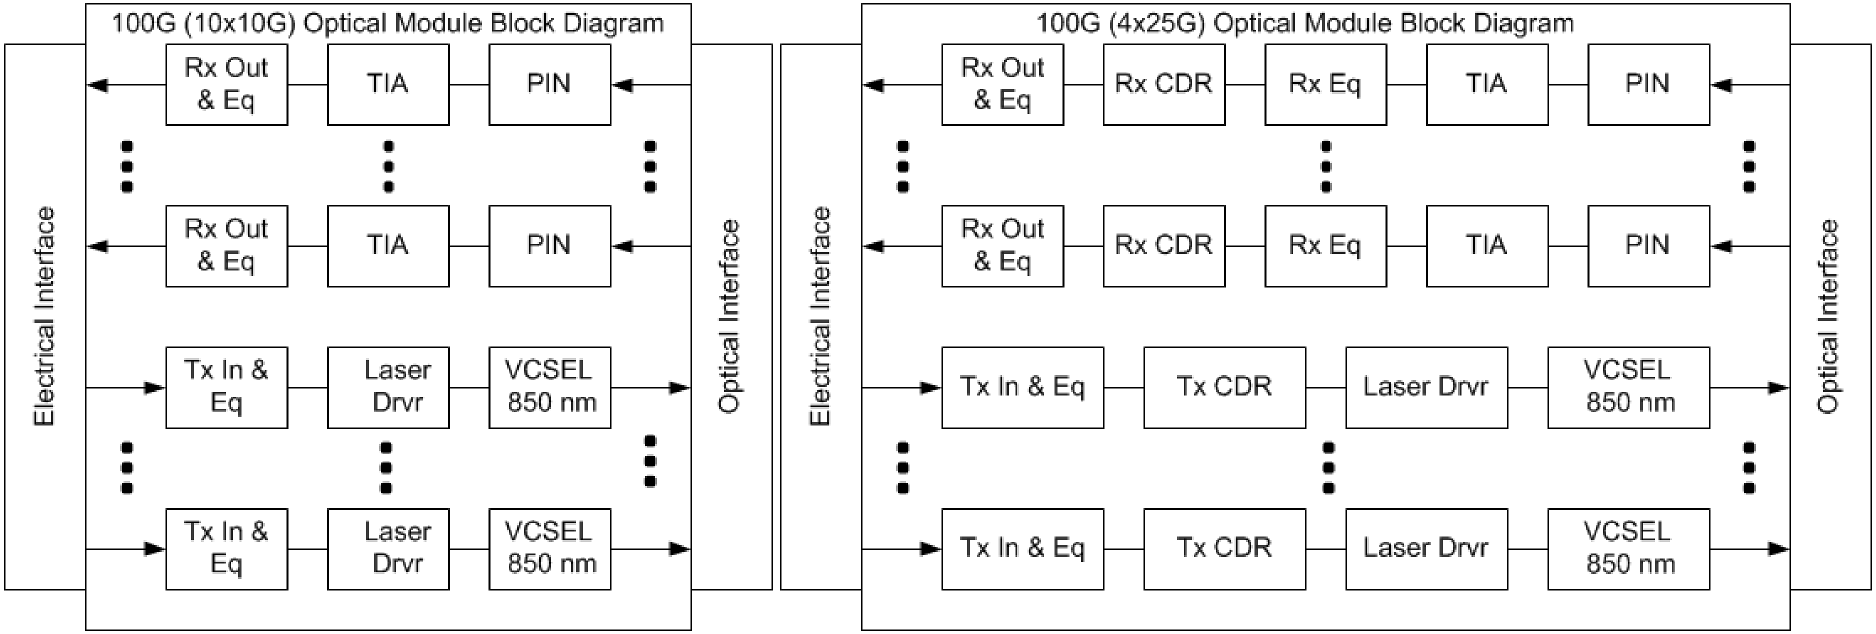

100GBASE-SR10 (10x10G) is shown to assist comparisons.

CXP is the assumed form factor for 100GBASE-SR10.

QSFP+ or CFP4, depending on power consumption and thermal management, are the assumed form factors for 100GBASE-SR4.

Comparative factors relative to 100GBASE-SR10 include lane count decrease, signal rate increase and added features.

The lower lane/channel count is expected to simplify host board layout and enable higher aggregate channel count.

Kolesar Kalculator Analysis (1 of 5)

Case		Relative 100G Ch. Cost	Reach		
#	Items		OM3	OM4	OS2
1	SR10 LR4	1 10	100 m	150 m	 10 km
2	SR4 LR4	1.3 10	70 m	100 m	 10 km
3	SR4 LR4	1.0 10	70 m	100 m	 10 km
4	SR4 LR4	0.70 10	70 m	100 m	 10 km

Relative cost factors in the above table are not estimates but provided as 'what-if' factors.

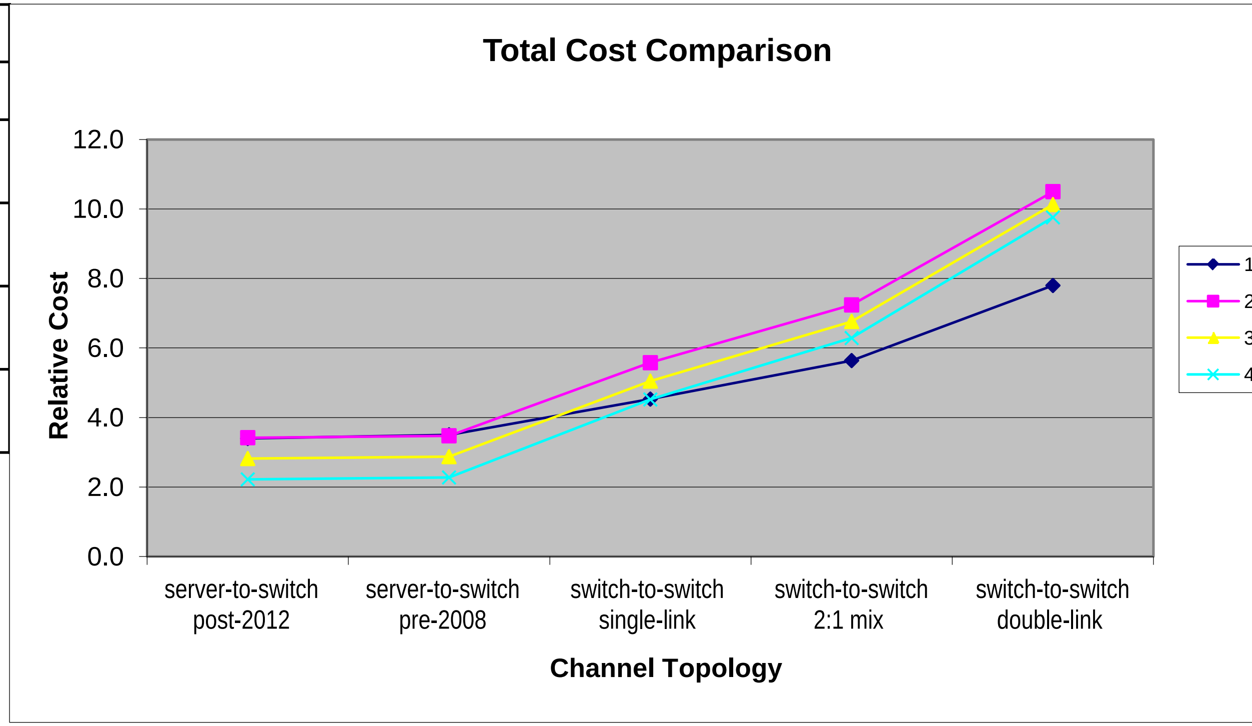

- The above chart shows data center costs comparisons generated with the Kolesar Kalculator for the reach distributions built into the calculator. Four 100G/Ch. cases were analyzed. See above table for definitions.
- Case 1 is the incumbent case where only 100GBASE-SR10 and 100GBASE-LR4 exist.
- Cases 2, 3 & 4 replace SR10 in Case 1 with SR4. SR4 OM3 and OM4 reaches are in discussion. The values used here are for example. Here relative costs of SR4 are considered.
- For the server-to-switch segments where the relative cost of LR4 is not a factor, SR4 offers little value if its relative cost approaches 1.3x SR10 for the example OM3 & OM4 reaches.

Kolesar Kalculator Analysis (2 of 5)

Case		Relative 100G Ch. Cost	Reach		
#	Items		OM3	OM4	OS2
1	SR10 LR4	1 10	100 m	150 m	10 km
2	SR4 LR4	1.3 10	70 m	100 m	10 km
3	SR4 LR4	1.3 10	100 m	140 m	10 km
4	SR4 LR4	1.3 10	140 m	200 m	10 km

Relative cost factors in the above table are not estimates but provided as 'what-if' factors.

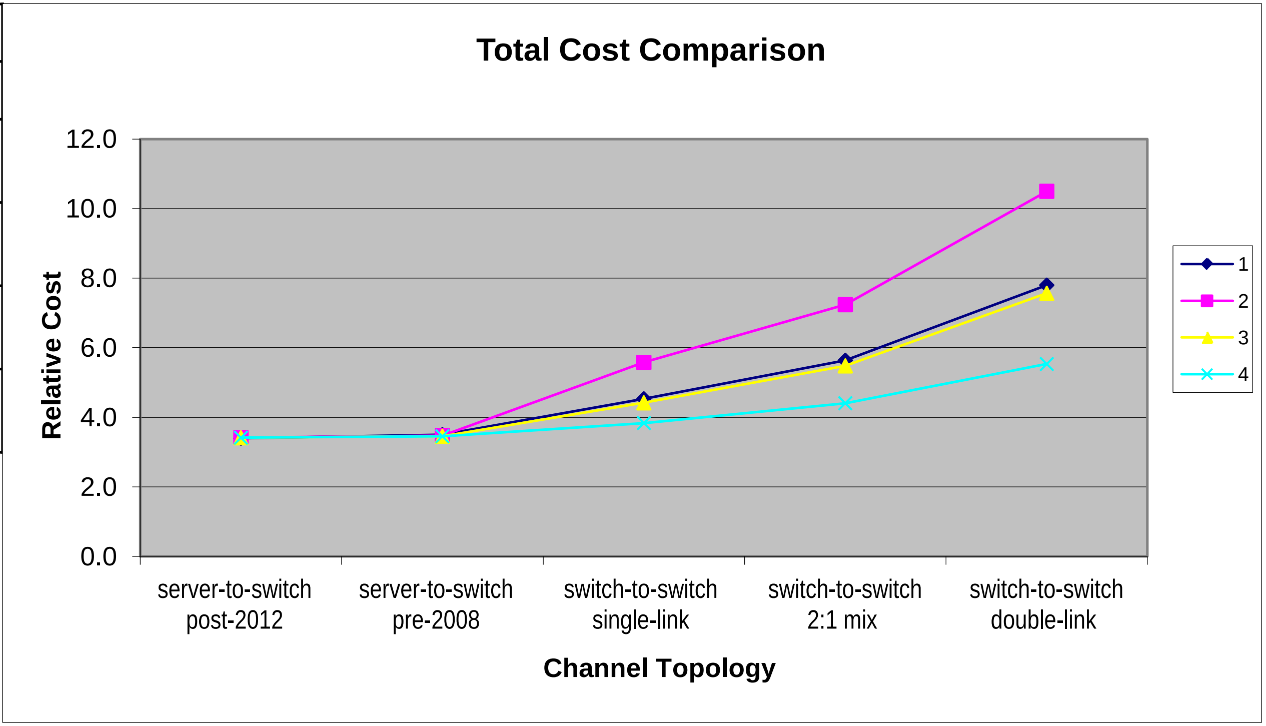

The above chart shows data center costs comparisons generated with the Kolesar Kalculator for the reach distributions built into the calculator. Four 100G/Ch. cases were analyzed. See above table for definitions.

Case 1 is the incumbent case where only 100GBASE-SR10 and 100GBASE-LR4 exist.

Cases 2, 3 & 4 replace SR10 in Case 1 with SR4. Here SR4 reaches are considered.

Parity with the incumbent case is reached for an SR4-LR4 system for an SR4 relative cost of 1.3x supporting 100 m of OM3 and 140 m of OM4.

Kolesar Kalculator Analysis (3 of 5)

Case		Relative 100G Ch. Cost	Reach		
#	Items		OM3	OM4	OS2
1	SR10 LR4	1 10	100 m	150 m	 10 km
2	SR4 LR4	0.7 10	50 m	70 m	 10 km
3	SR4 LR4	1.0 10	70 m	100 m	 10 km
4	SR4 LR4	1.3 10	100 m	140 m	 10 km

Relative cost factors in the above table are not estimates but provided as 'what-if' factors.

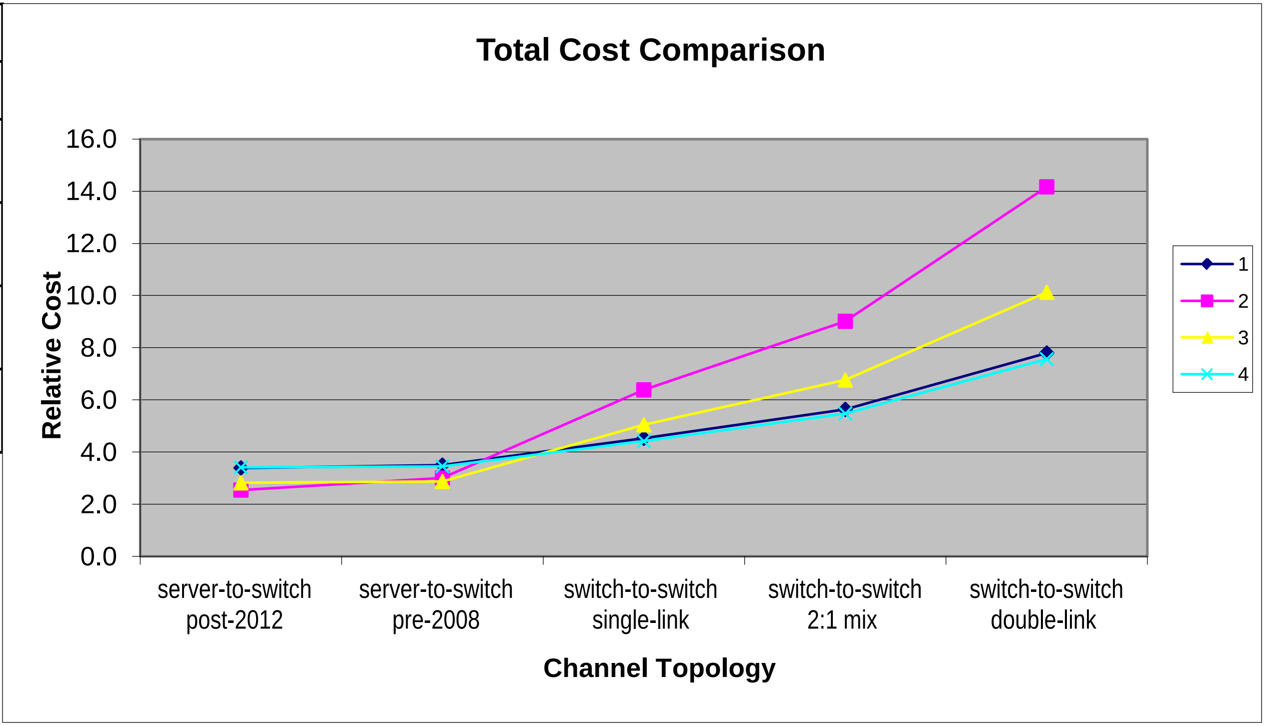

The above chart shows data center costs comparisons generated with the Kolesar Kalculator for the reach distributions built into the calculator. Four 100G/Ch. cases were analyzed. See above table for definitions.

Case 1 is the incumbent case where only 100GBASE-SR10 and 100GBASE-LR4 exist.

Cases 2, 3 & 4 replace SR10 in Case 1 with SR4. Here SR4 combinations of relative cost and reaches are considered.

Parity with the incumbent case is reached for an SR4-LR4 system for an SR4 relative cost of 1.3x supporting 100 m of OM3 and 140 m of OM4. While cases 2 & 3 show an advantage for the server-to-switch segments, the advantage is lost for the switch-to-switch segments.

Kolesar Kalculator Analysis (4 of 5)

Case		Relative 100G Ch. Cost	Reach		
#	Items		OM3	OM4	OS2
1	SR10 LR4	1 10	100 m	150 m	 10 km
2	SR4 LR4	1.3 10	100 m	140 m	 10 km
3	SR4 PSM4 LR4	1.2 3.5 10	70 m	100 m	 1 km 10 km
4	SR4 PSM4 LR4	1 3.5 10	70 m	100 m	 1 km 10 km

Relative cost factors in the above table are not estimates but provided as 'what-if' factors.

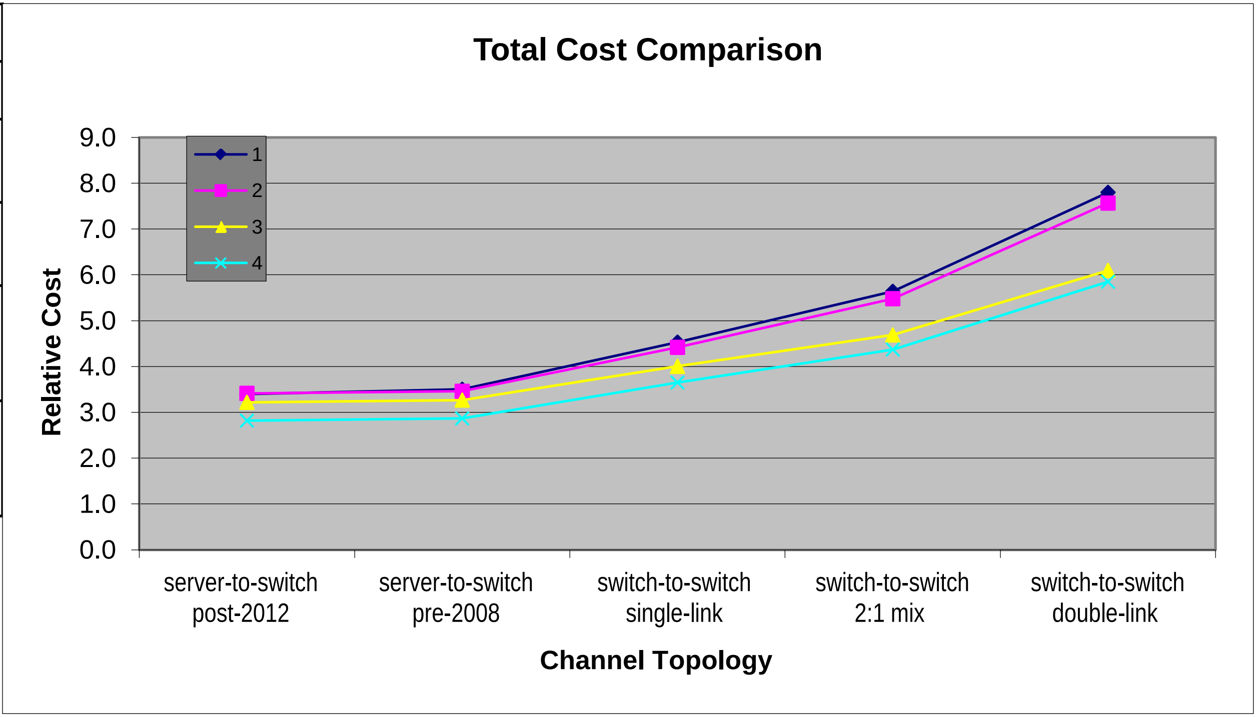

The above chart shows data center costs comparisons generated with the Kolesar Kalculator for the reach distributions built into the calculator. Four 100G/Ch. cases were analyzed. See above table for definitions.

Case 1 is the incumbent case where only 100GBASE-SR10 and 100GBASE-LR4 exist.

Case 2 replaces SR10 in Case 1 with SR4 at factors providing parity with the incumbent case.

Cases 3 & 4 explores the inclusion of PSM4 with SR4 at a shorter reach than the parity case but with lower cost factors.

PSM4 in combination with a lower-cost-short-reach SR4 enables lower total costs for all segments.

Kolesar Kalculator Analysis (5 of 5)

Case		Relative 100G Ch. Cost	Reach		
#	Items		OM3	OM4	OS2
1	SR10 LR4	1 10	100 m	150 m	 10 km
2	SR4 LR4	1.3 10	100 m	140 m	 10 km
3	SR4 PSM4 LR4	1.0 3.5 10	70 m	100 m	 1 km 10 km
4	SR4 PSM4 LR4	0.7 3.5 10	50 m	70 m	 1 km 10 km

Relative cost factors in the above table are not estimates but provided as 'what-if' factors.

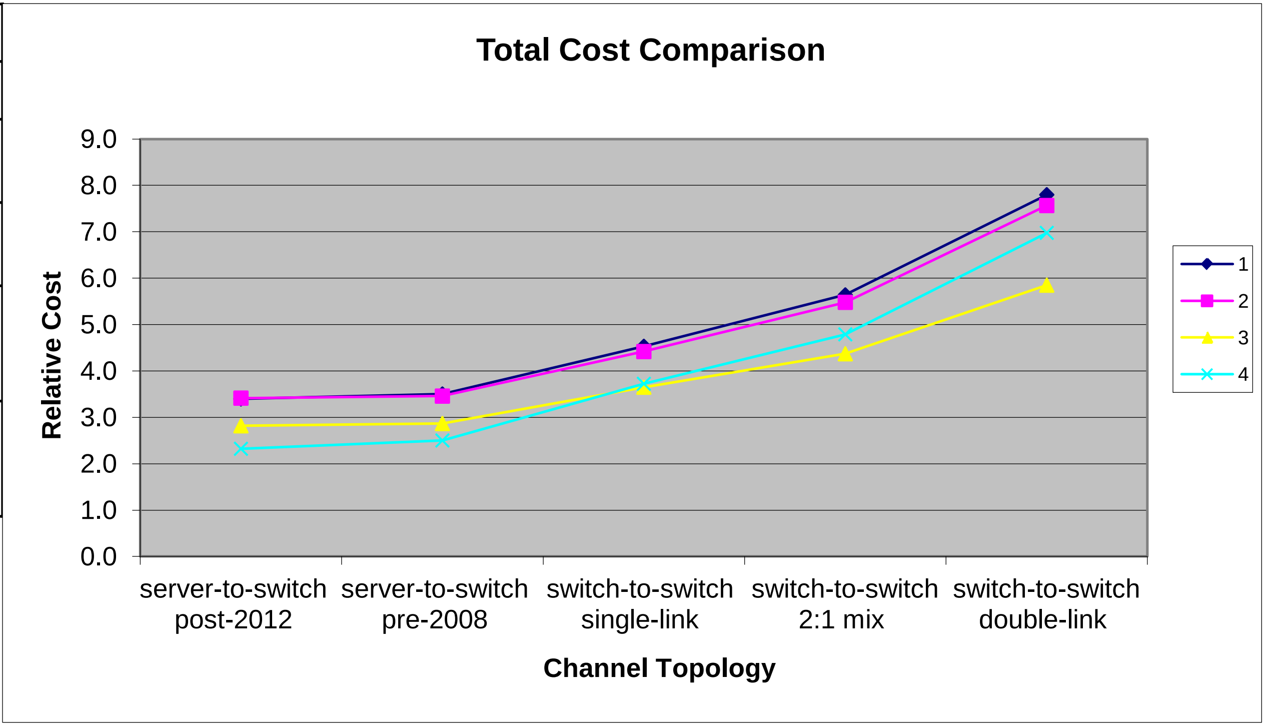

The above chart shows data center costs comparisons generated with the Kolesar Kalculator for the reach distributions built into the calculator. Four 100G/Ch. cases were analyzed. See above table for definitions.

Case 1 is the incumbent case where only 100GBASE-SR10 and 100GBASE-LR4 exist.

Case 2 replaces SR10 in Case 1 with SR4 at factors providing parity with the incumbent case.

Cases 3 & 4 explores the inclusion of PSM4 with SR4 for combinations of reach and cost factors.

While PSM4 in combination with a lower-cost-short-reach SR4 enables lower total costs for all segments, there is a trade-off between optimizing for costs and reach.

100G Crossover Analysis – Relative Cable Plant & Transceiver Cost (1)

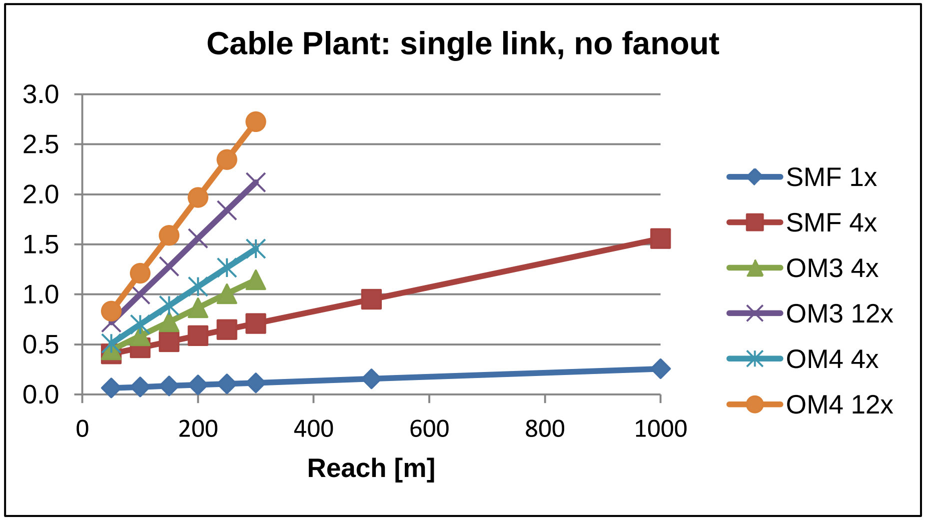

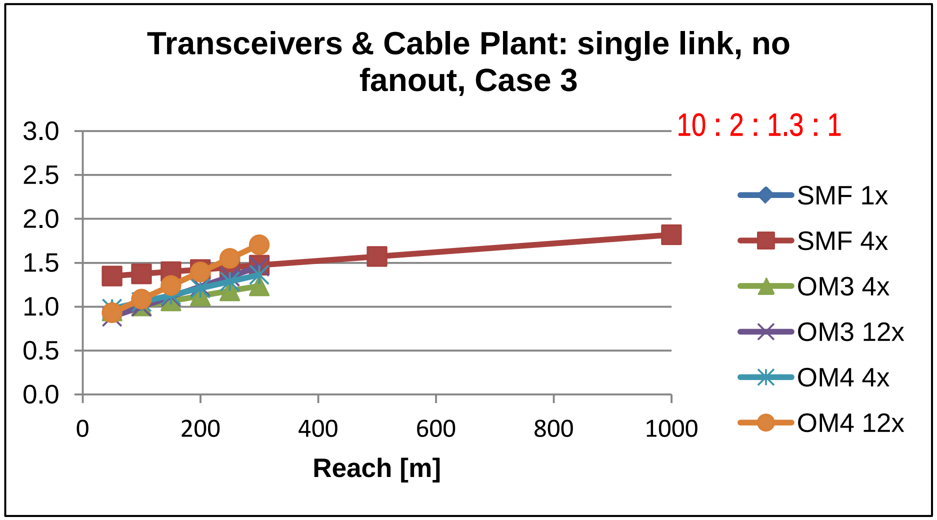

The charts shows relative cable and transceiver costs for a 100G channel based on estimates from Corning and petrilla_02a_0112_NG100GOPTX (page 12) and is the same cable plant used in petrilla_01_0312_NG100GOPTX.

This scenario reflects the simplest cable plant structure . Transceiver cost factors are shown in the charts, e.g. 10 : 2 : 1.3 : 1 corresponds to LR4 : PSM4 : SR4 : SR10 and SMF1x = LR4, SMF4x = PSM4, OMn4x = SR4, and OMn12x = SR10.

While parity between SR4 and SR10 costs can be seen at 50 m for the 1.3 SR4/SR10 factor, SR4 shows an advantage at longer reaches and lower cost ratios.

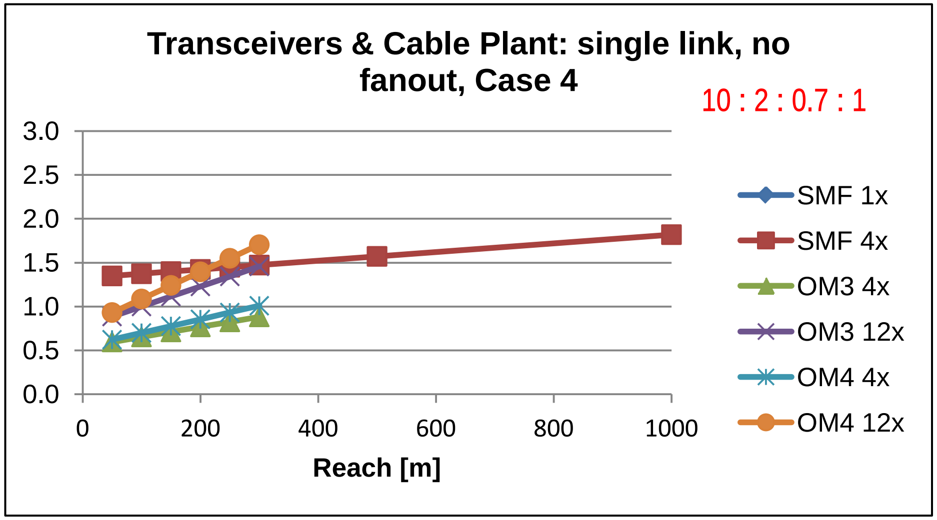

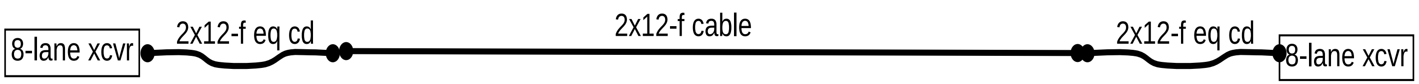

Diagram from kolesar_02_0911_NG100G)PTX

100G Crossover Analysis – Relative Cable Plant & Transceiver Cost (2)

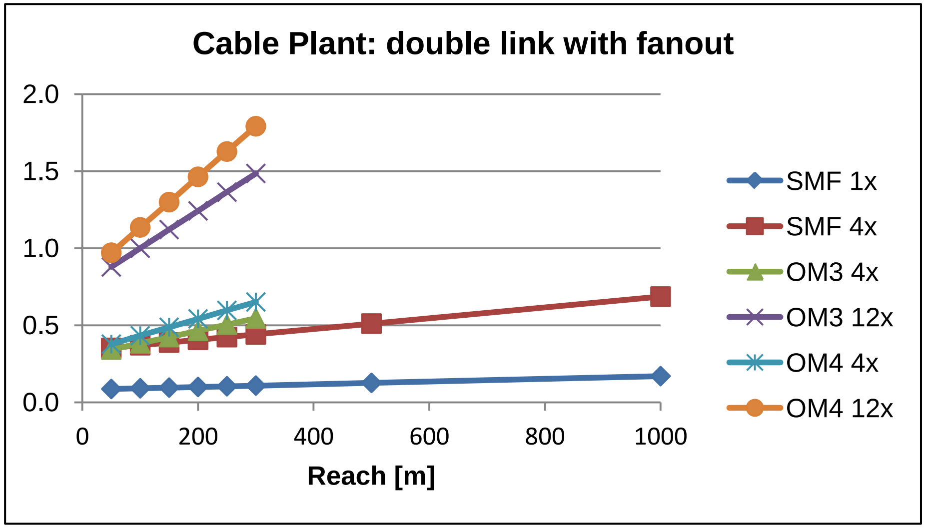

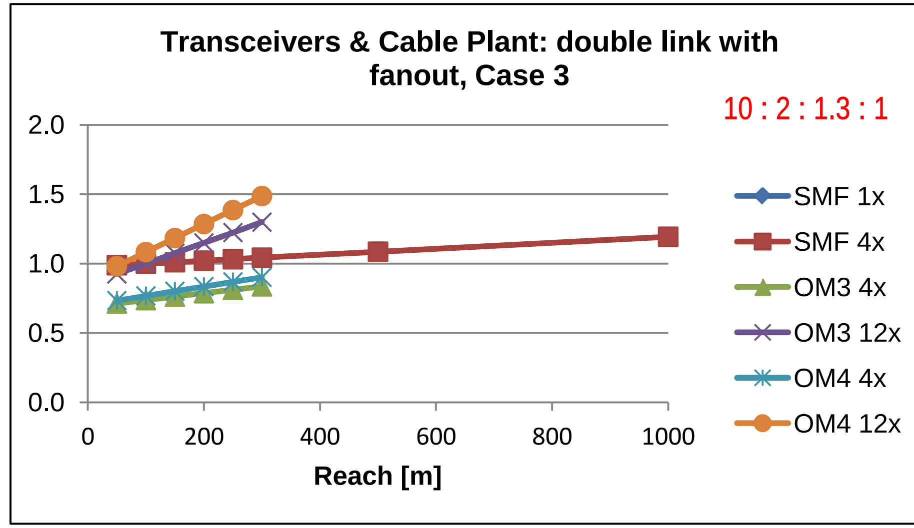

The charts shows relative cable and transceiver costs for a 100G channel based on estimates from Corning and petrilla_02a_0112_NG100GOPTX (page 12)

This scenario reflects the simplest cable plant structure . Transceiver cost factors are shown in the charts, e.g. 10 : 2 : 1.3 : 1 corresponds to LR4 : PSM4 : SR4 : SR10 and SMF1x = LR4, SMF4x = PSM4, OMn4x = SR4, and OMn12x = SR10.

Here the SR4 solution, perhaps since all trunk fibers are used, shows an advantage at all reaches for SR4/SR10 cost factors in the range of 0.7 to 1.3.

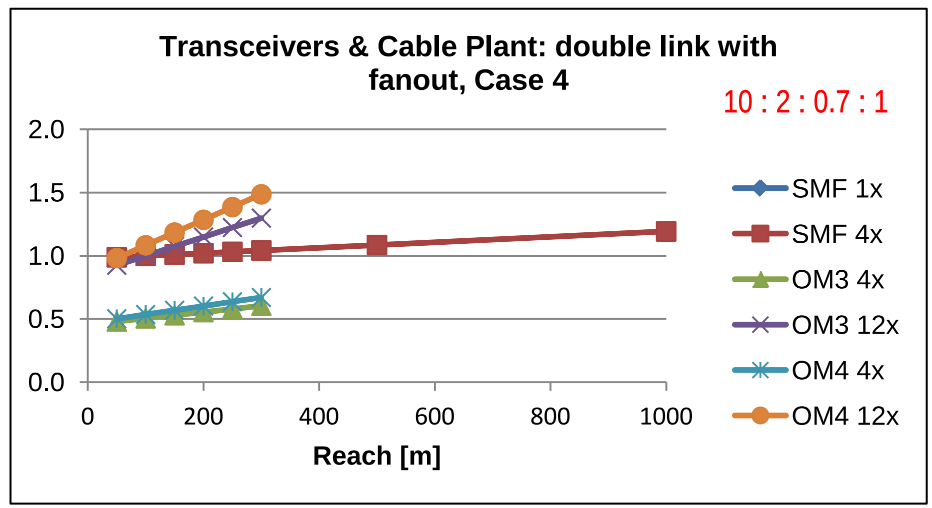

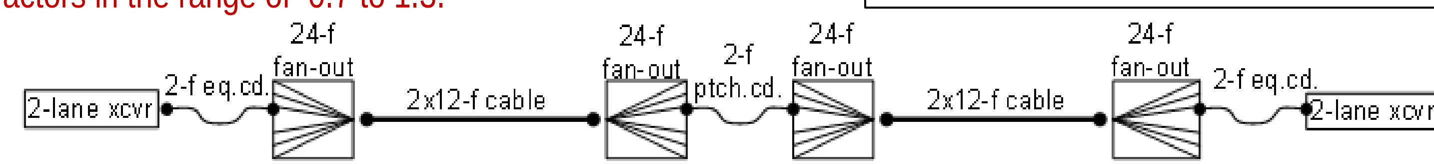

Diagram from kolesar_02_0911_NG100G)PTX

Recipe for low cost, low power SR XCVR (Iclp SR): Scaling 40G SR4

	40G SR4	Iclp SR	Risk	Comments
Signal Rate/Lane	10.3125 GBd	25.78125 GBd		2.5 x scaling
Tx output transition time	44 ps	17.6 ps	High	1/2.5 x scaling
Fiber effective modal BW	20 GHz	50 GHz	Low	2.5 x scaling, change 100m of OM3 to 88m of OM4
Chromatic Dispersion BW	26.5 GHz	68.5 GHz	High	2.5 x scaling, change fiber Disp. So from 0.10275 to 0.04526 ps/nm^{2}km,
Chromatic Dispersion BW	26.5 GHz	68.5 GHz	High	or change Tx Uw from 0.65 to 0.286 nm, or a combination, e.g. Uw = 0.50 nm & Disp So = 0.05883 ps/nm^{2}km,
Chromatic Dispersion BW	26.5 GHz	68.5 GHz	Low	or further reduce reach e.g. 67m of OM4 with Uw=0.50 nm and no change to Disp. So.
Rx BW w RxS = -11.3 dBm	7500 GHz	18750 GHz	High	2.5 x scaling, maintains 8.3 dB signal budget , no change in OMA
Rx BW w RxS = -9.3 dBm	7500 GHz	18750 GHz	Medium	Requires Min Tx OMA from -3.0 to -1.0 dBm for 8.3 dB signal budget
Rx BW w RxS = -7.3 dBm	7500 GHz	18750 GHz	Medium	Requires Min Tx OMA from -3.0 to +1.0 dBm for 8.3 dB signal budget
TP1 J2, J9, Qsq, Eye Mask Coordinates	0.17UI, 0.29UI, 45V/V, 0.11UI, 0.31UI, 95mV, 350mV	0.17UI, 0.29UI, 45V/V, 0.11UI, 0.31UI, 95mV, 350mV	High	Input conditions scaled from XLPPI, no CDR and no equalization except as found in 40GBASE-SR4.
TP4 J2, J9, Eye Mask Coordinates	0.42UI, 0.659UI, 0.29UI, 0.5UI, 150mV, 425mV	0.42UI, 0.659UI, 0.29UI, 0.5UI, 150mV, 425mV	High	Output requirements scaled from XLPPI, no CDR and no equalization except as found in 40GBASE-SR4

Known solutions – low risk
Early results – medium risk
Need verification – high risk

A low cost, low power 100G SR4 would presumably have unretimed interfaces and require no optical link equalization, essentially a scaled 40GBASE-SR4. The above table considers key attributes for such scaling and assigns a risk factor. The only low risk item appears to be scaling the optical media to reach the required 2.5x BW.

A 100G SR4 that does not require, at least, retimers does not seem feasible within the near future.

100G 25G/Lane MMF Transceiver: Reach & Equalization (1/3)

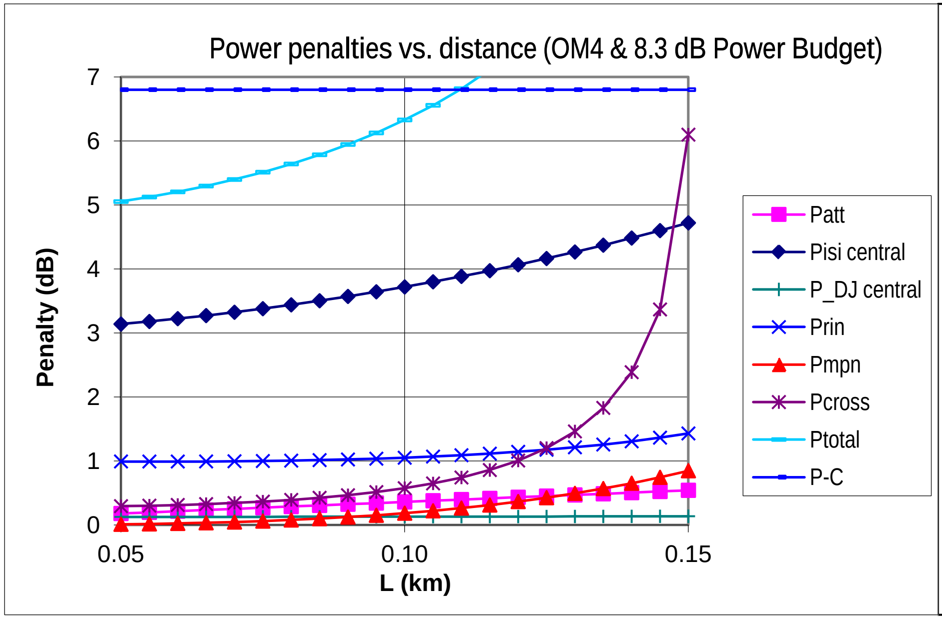

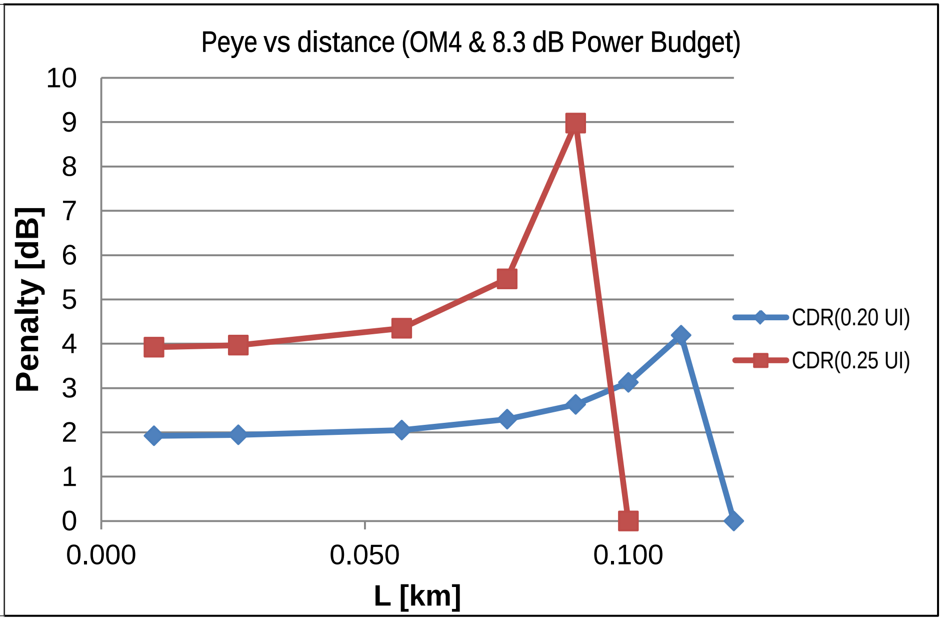

- The above charts are generated using a 10GEPBud3_1_16a link model with no optical link equalization where jitter at TP1 (RJ = 0.1 UI and DJ = 0.1 UI) and at TP4 (TJ = 0.8 UI) represent the jitter generation and tolerance of CDR included in the optical transceiver.
- The chart on the left is provided by the link model which unfortunately does not include the power penalty required to support a desired eye opening. The chart on the right provides that information for two CDR cases: CDR(0.20 UI) for a CDR with a 0.2 UI jitter generation and a 0.8 UI jitter tolerance, and CDR(0.25 UI) for a CDR with a 0.25 UI jitter generation and a 0.75 UI jitter tolerance. The power penalty, Peye, from the right should be added to Ptotal on the left.
- With CDRs of the quality of CDR(0.20 UI) and power budgets of 8.3 dB, OM4 links up to 50 m may not require equalization.

100G 25G/Lane MMF Transceiver: Reach & Equalization (2/3)

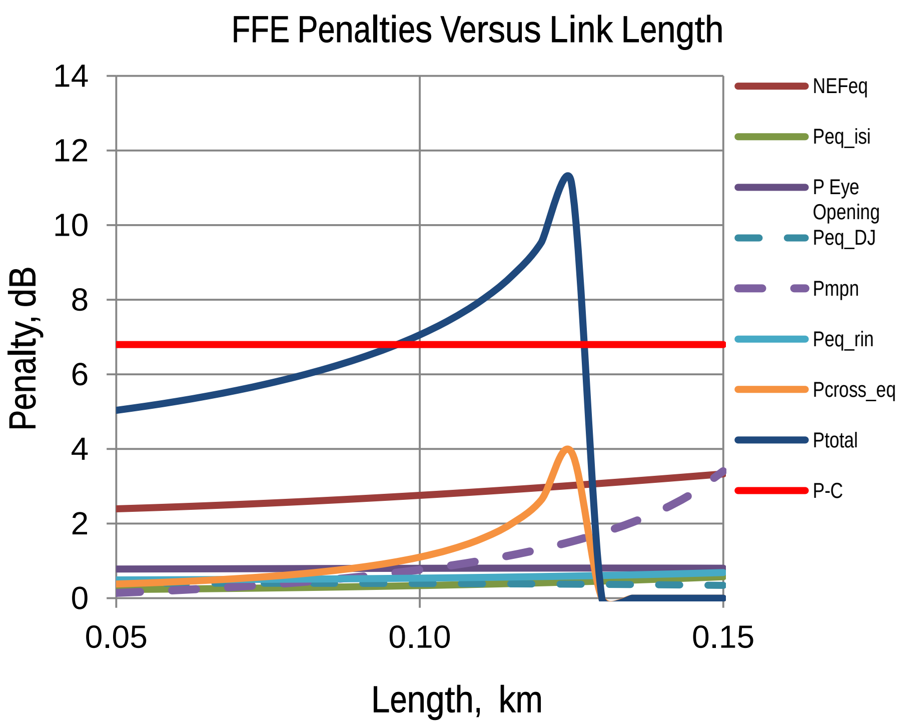

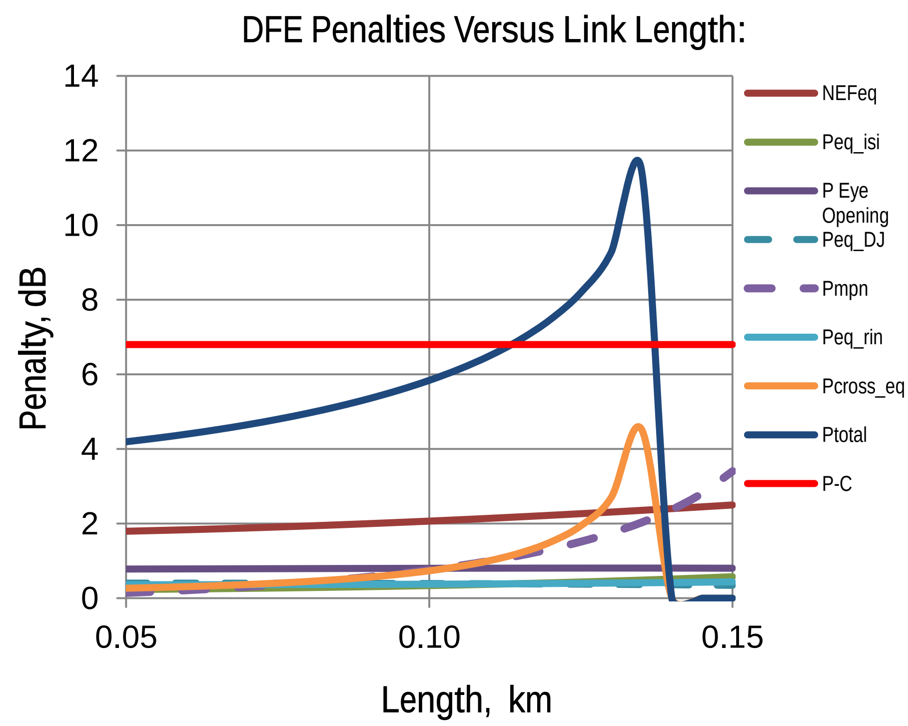

- The above charts are generated using a link model, 32GFC 3T FFE v1a, that includes the affects of equalization for a link with the same attributes as that on the previous page, i.e. power budget of 8.3 dB with input and output CDRs of the quality of CDR(0.20 UI). As on the previous page the penalty for opening the output eye, Peye, is not included.
- The link model, 32GFC 3T FFE v1a, was developed by David Cunningham in the FC-PI-6 project and is available in T11/12-044v1. Additional references are available on the following page.
- Here a 3-tap FFE appears to run into a noise induced wall near 125m of OM4 and a DFE(3,1) runs into that wall near 135 m of OM4. Further, if Peye is considered, 100m of OM4 will be a challenge for a DFE(3,1) as will 80 m of OM4 be for a 3-tap FFE.

100G 25G/Lane MMF Transceiver: Reach & Equalization (3/3)

- The DFE represented by the spreadsheet is roughly a 3-Tap FFE (T spaced) followed by 1 DFE tap which is termed DFE(3,1). In presentations to FC-PI-6 , David Cunningham has shown that for the optical link a DFE(3,1) is close to optimum. Longer DFE's do not help the optics.
- The document number for the 32FGC Spreadsheet is: T11/12-044v1
- The document numbers and titles for the T11 presentations on the spreadsheet are:
- T11/12-042v0 "Modifications To The Mode Partition Noise Penalty Calculation For Equalised 32GFC Links"
- T11/11-502v0 "Optical Link Model Modifications Required To Include A Short Equaliser"
- These can be accessed by logging in to the T11 (http://www.t11.org/index.html) site as a guest, clicking on Docs, clicking on access by document number and entering the number.