1 x 400G vs. 4 x 100G FEC performance

Pete Anslow, Ciena

IEEE P802.3bs Task Force, Waikoloa, July 2015

Supporters (for the architecture conclusion)

Thananya Baldwin	Ixia
Mark Gustlin	Xilinx
Martin Langhammer	Altera
Mike Peng Li	Altera
Jeffery Maki	Juniper
Gary Nicholl	Cisco
David Ofelt	Juniper
Jerry Pepper	Ixia
Steve Trowbridge	Alcatel-Lucent

Introduction

Both anslow_3bs_03_0515 and anslow_3bs_04_0715 discuss the performance of RS(544,514) FEC with the assumption:

> “assume a 1 x 400G FEC architecture with the advantage of diluting the errors from the worst lane with 15 other lanes rather than with 3 other lanes for the 4 x 100G case”

This presentation therefore contains analysis of the performance difference between these two candidate FEC architectures and concludes with a high level summary of the issues and a recommendation for the FEC architecture choice.

The measurement data used in this contribution was kindly supplied by Jonathan King, courtesy of Finisar.

400GBASE-SR16

The two cases considered for 400GBASE-SR16 are shown below:

4 x 100G

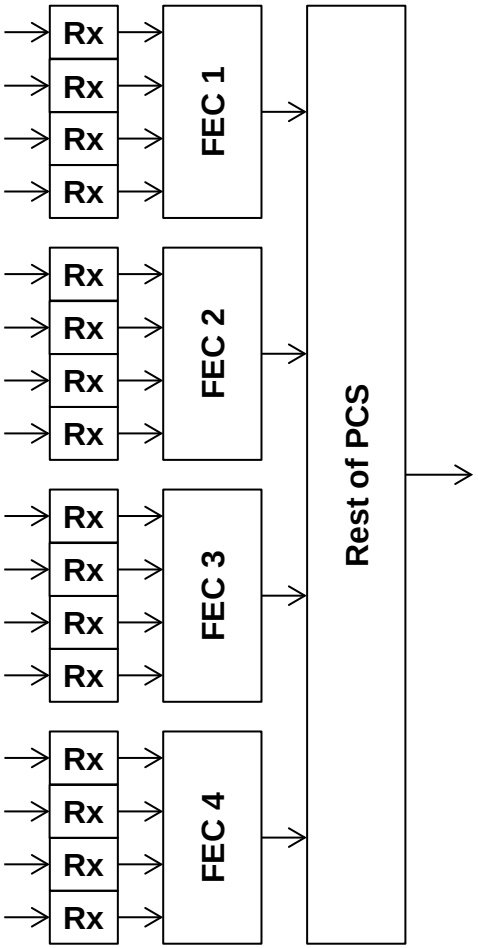

1 x 400G

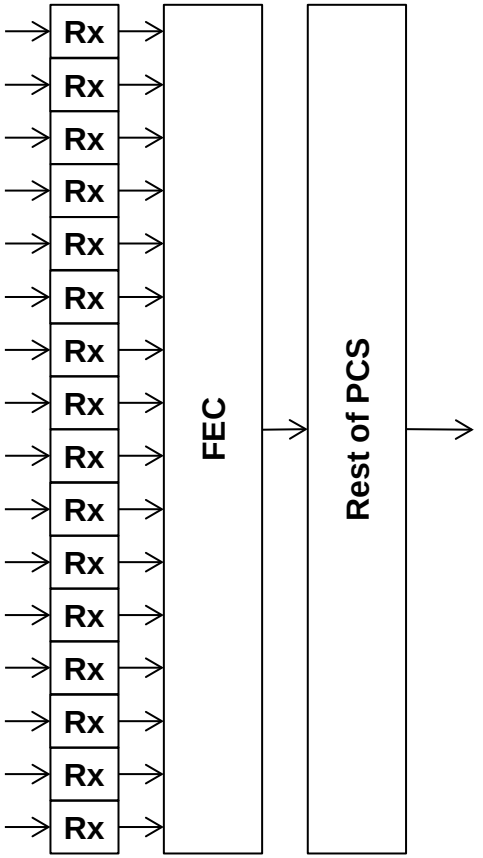

Monte Carlo analysis method

To analyse the difference between these two cases, existing test data for multi-lane receivers was used to create 500 sets of 16 relative receiver sensitivities.

For each set of relative receiver sensitivities a Gaussian model was used to produce a set of 16 input BERs. These BERs were then used as inputs to a 4 x 100G FEC model as illustrated on the left side of the next slide.

The 16 equal value attenuators in the model were then adjusted to the value required to produce a BER of 1E-13 for the 400G aggregate after FEC processing.

The same set of 16 BERs was then applied to the 1 x 400G FEC model as per the right side of the next slide and the resulting FEC output BER calculated.

The above process was then repeated for each of the 500 sets of relative sensitivities.

400GBASE-SR16 analysis

The analysis method used for 400GBASE-SR16 is illustrated below:

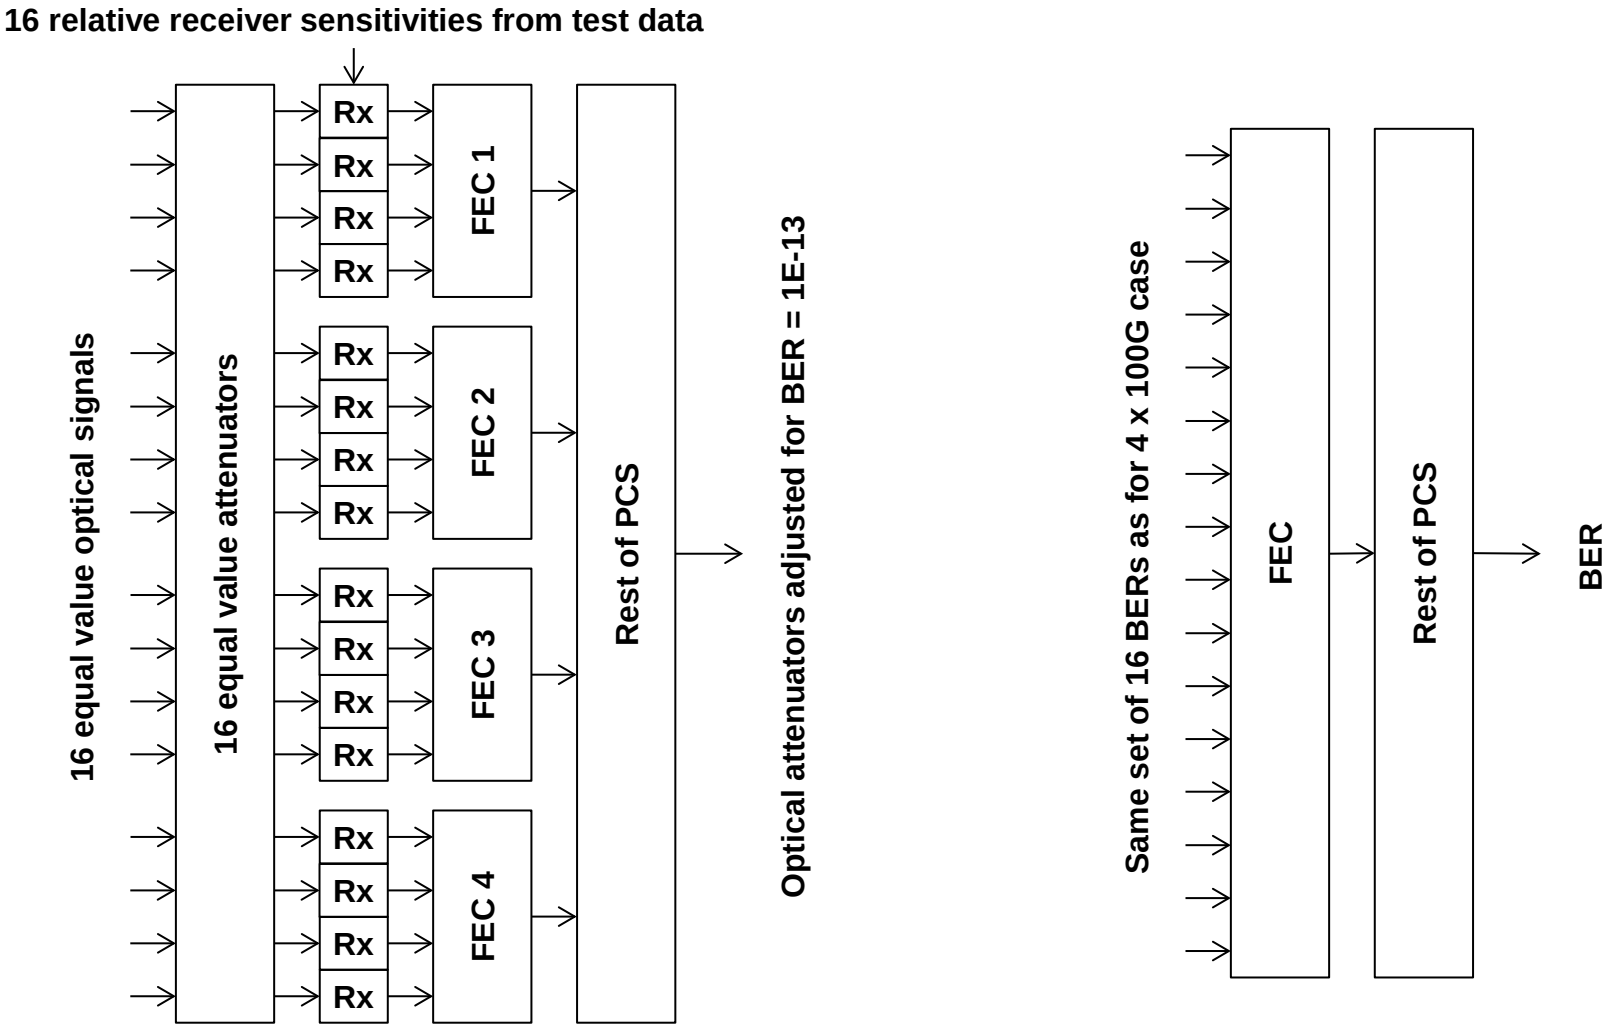

400GBASE-SR16 results

The distribution of BERs resulting from the 400GBASE-SR16 analysis is shown below. The output BER has a median value of 4.1E-17.

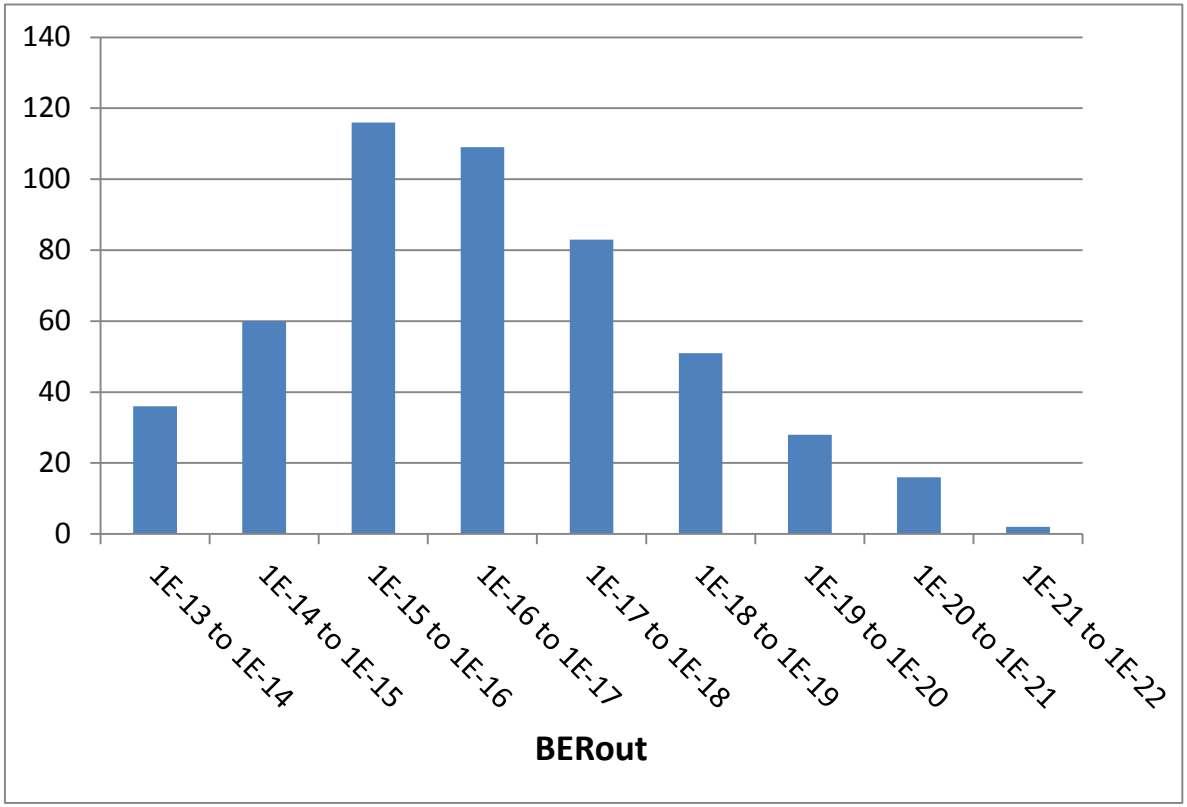

400GBASE-LR8 analysis

The analysis method used for 100GBASE-LR8 is illustrated below:

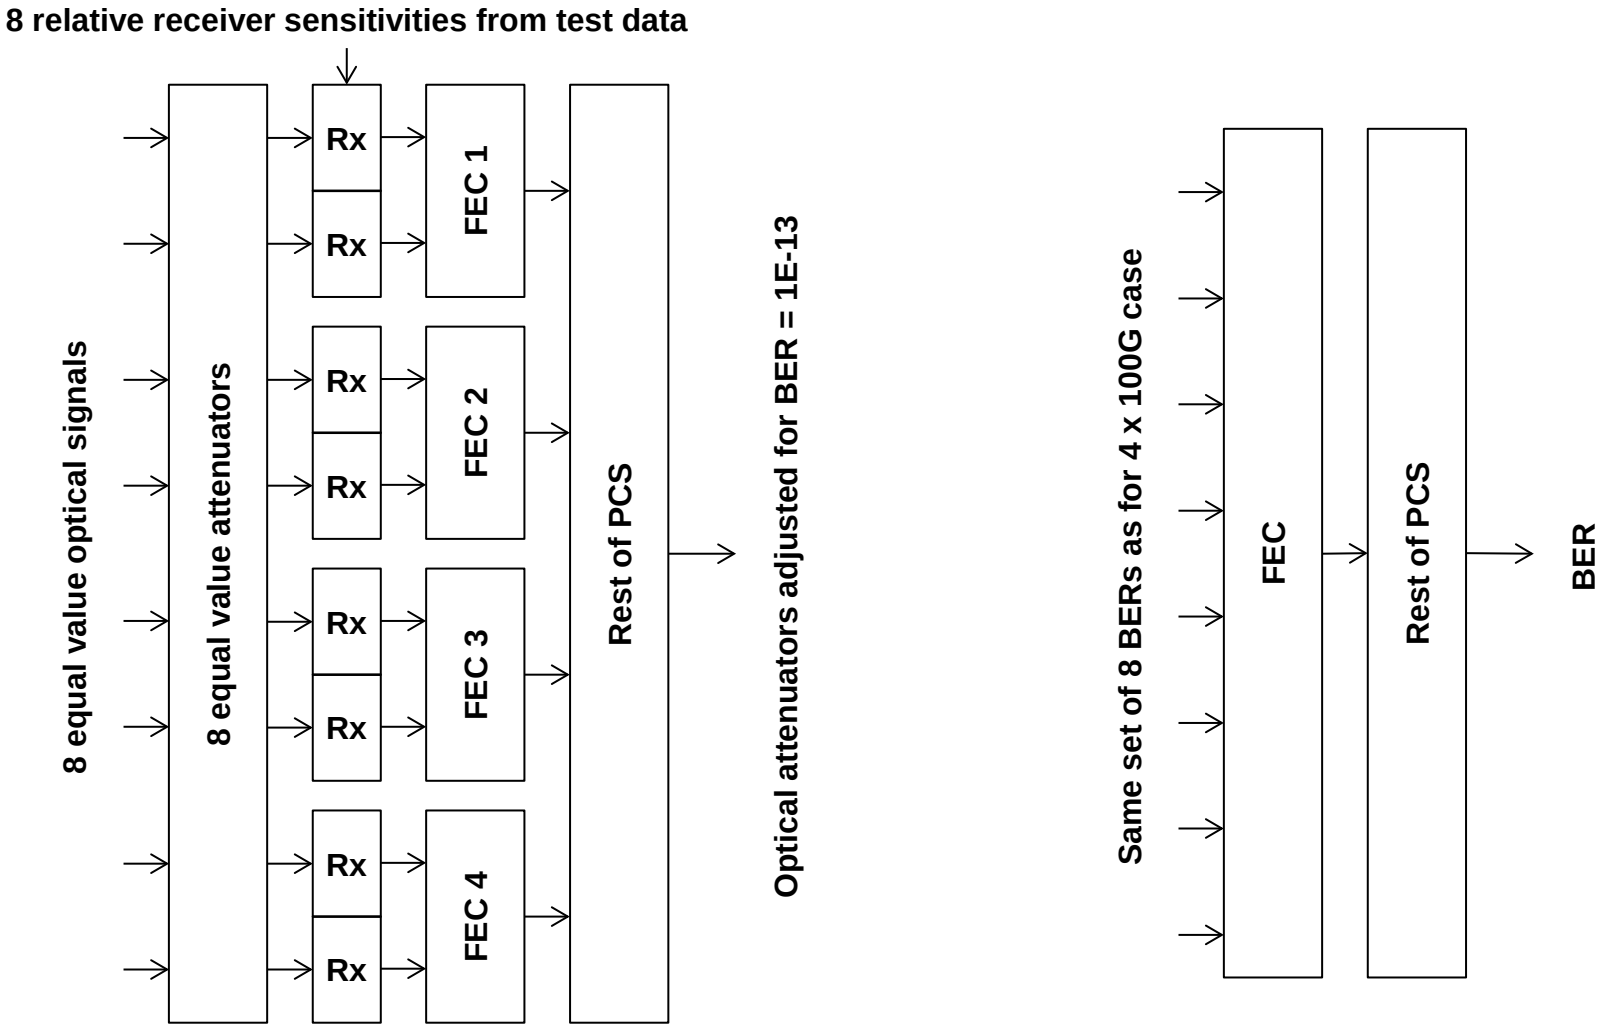

400GBASE-LR8 results

The distribution of BERs resulting from the 400GBASE-LR8 analysis is shown below. The output BER has a median value of 2.8E-18.

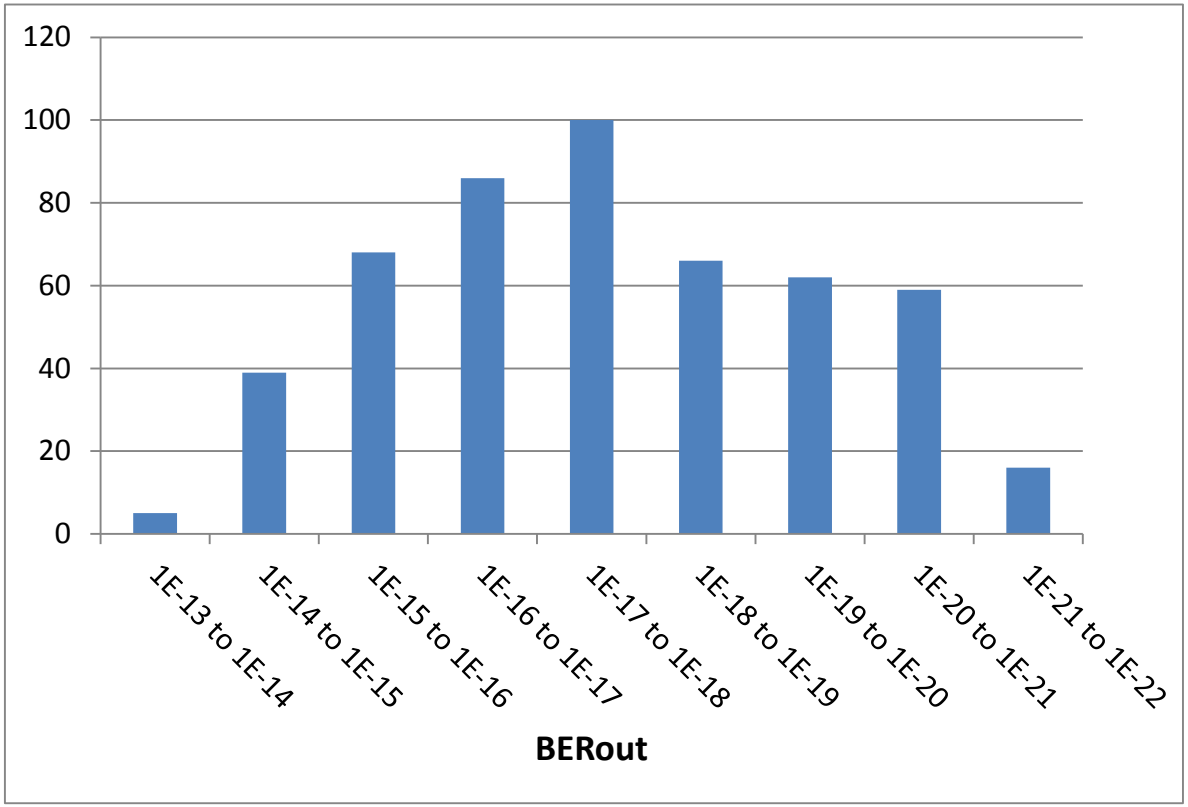

400GBASE-?R4 analysis

The analysis method used for 100GBASE-?R4 is illustrated below:

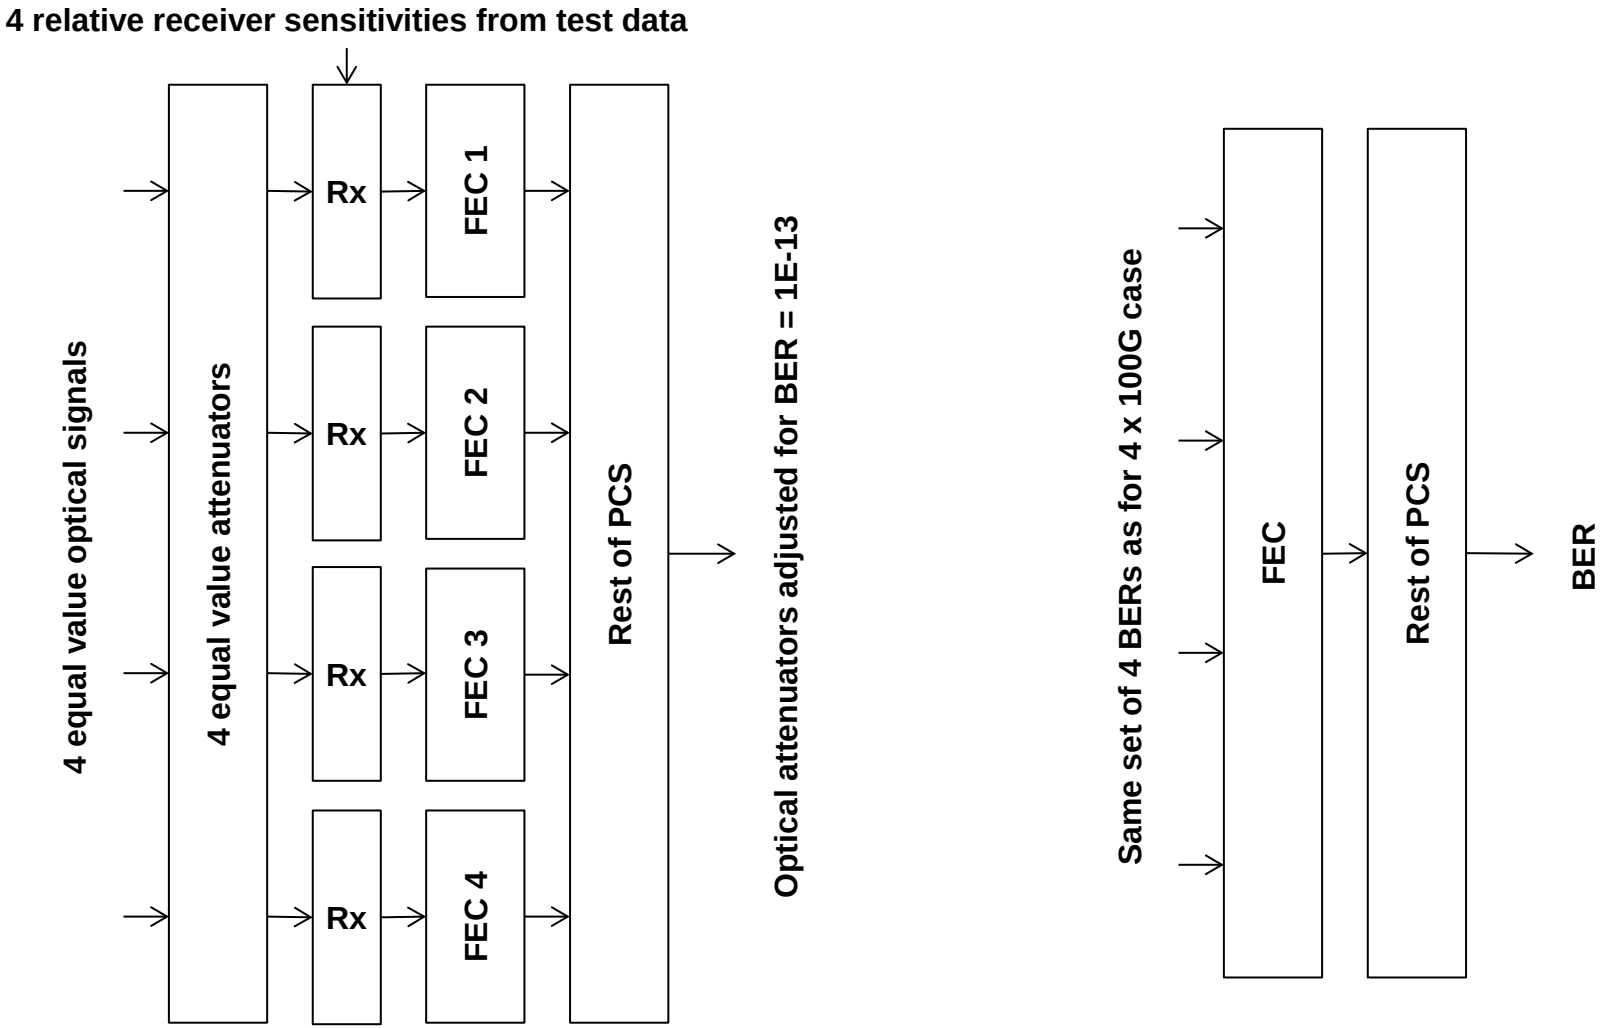

400GBASE-?R4 results

The distribution of BERs resulting from the 400GBASE-?R4 analysis is shown below. The output BER has a median value of 1.3E-19.

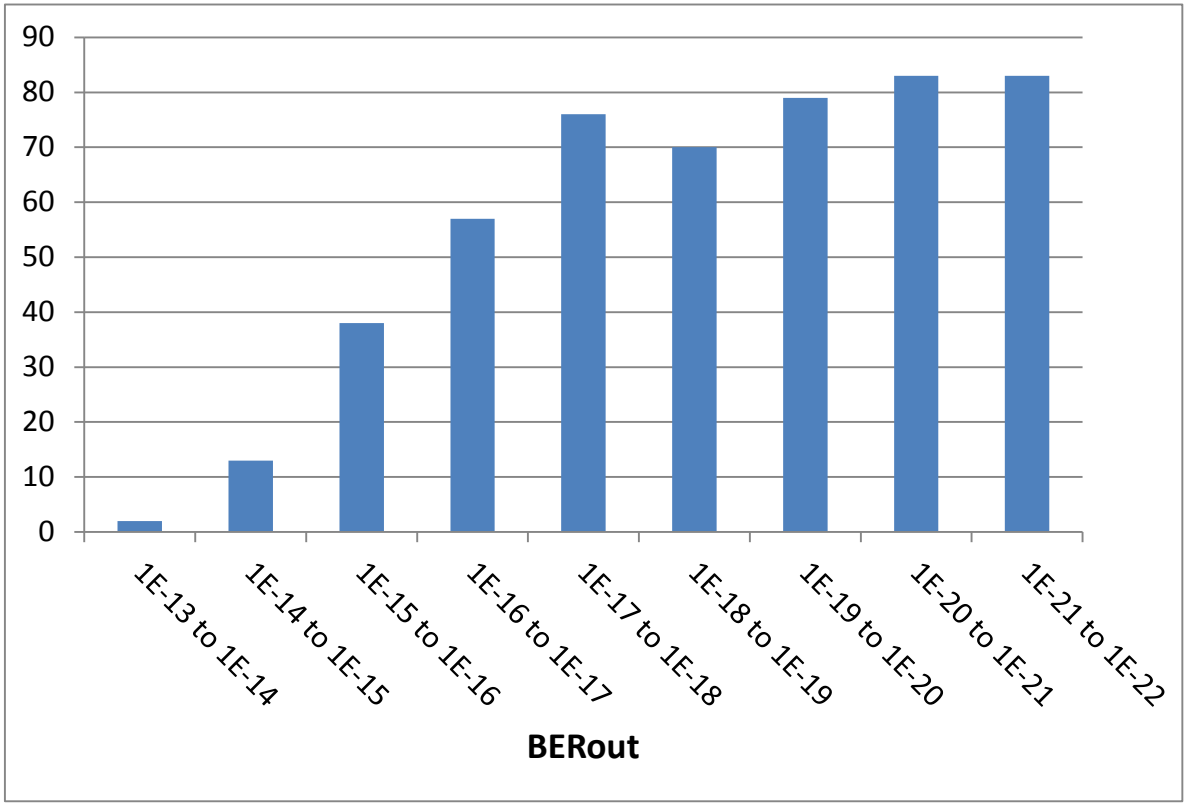

What does this mean for a PMD layer spec?

For the 1 x 400G FEC case, the PMD layer specification can be written (as 100GBASE-SR10 was) to place a limit on the average BER across all of the PCS lanes so that the FEC output BER should be 1E-13 or better (or the FLR equivalent).

For the 4 x 100G case, things are more complicated as the PMD layer (the module) does not know which lanes will be sent to which of the 4 FEC instances. (FOM does not solve this problem).

Option 1. The simplest way to write the PMD specification for the 4 x 100G case is therefore to limit the BER of each lane to that required for the FEC output BER to be 1E-13 or better (or the FLR equivalent).

Option 2. The specification could limit the average BER over the 4 worst PCS lanes to that required for the FEC output BER to be 1E-13 or better (or the FLR equivalent).

The effect of implementing these two alternatives is analysed in the following slides.

400GBASE-SR16 PMD spec option 1

The analysis method used for option 1 is illustrated below:

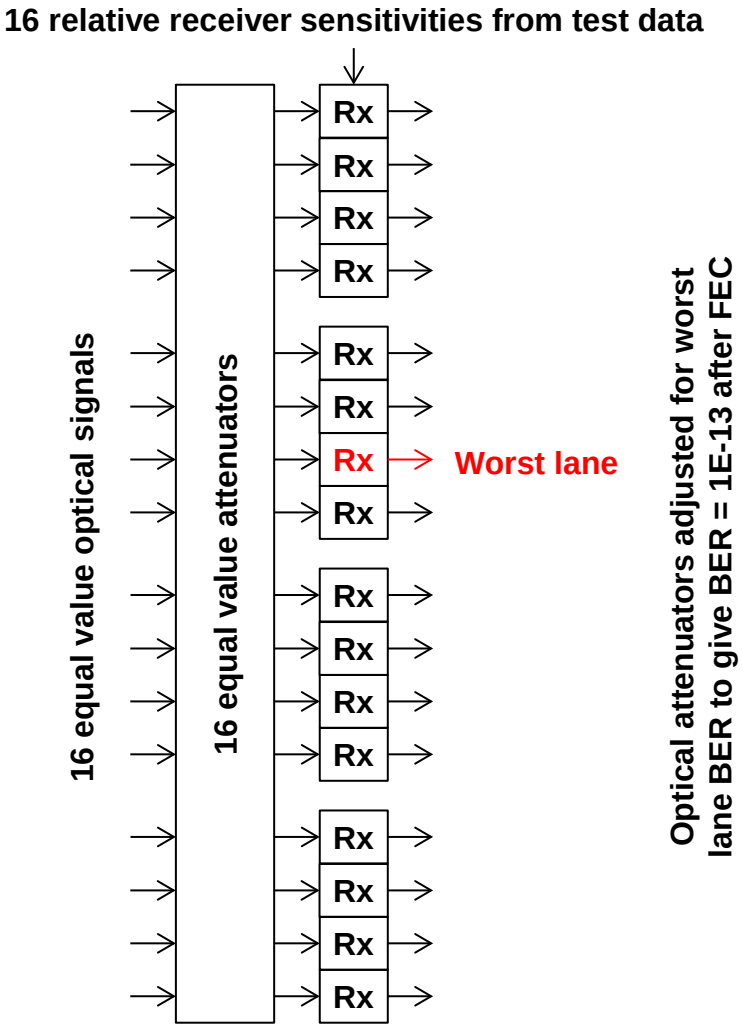

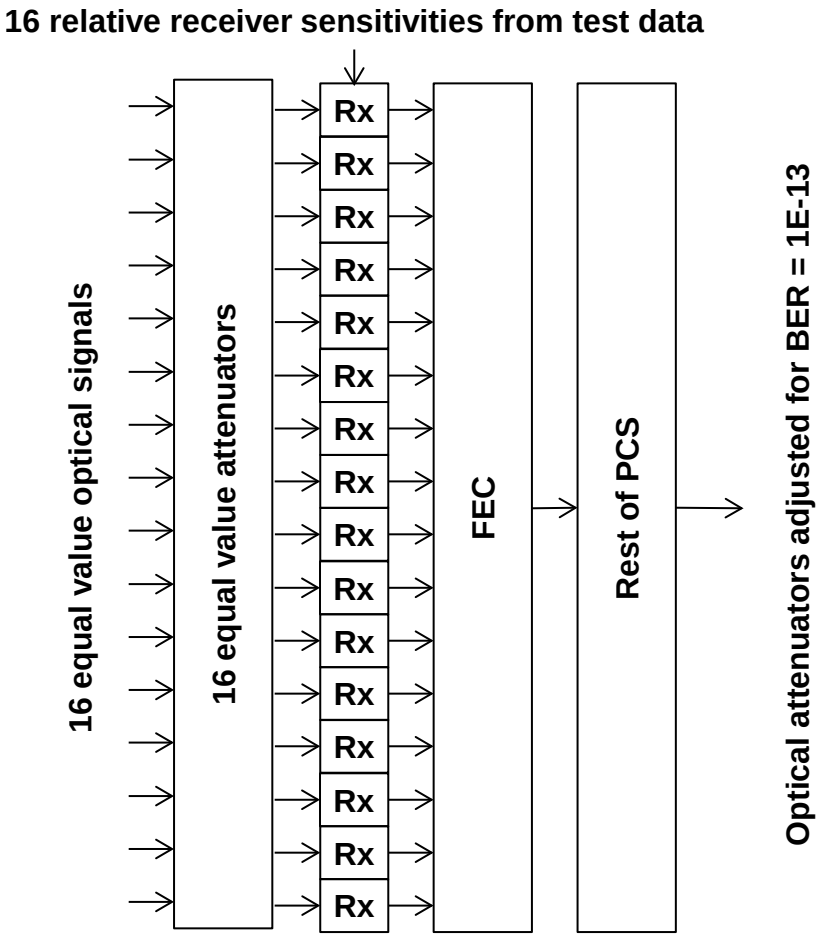

400GBASE-SR16 PMD spec option 1 results

The distribution of attenuator setting differences resulting from the option 1 analysis is shown below. The median value is 0.64 dB.

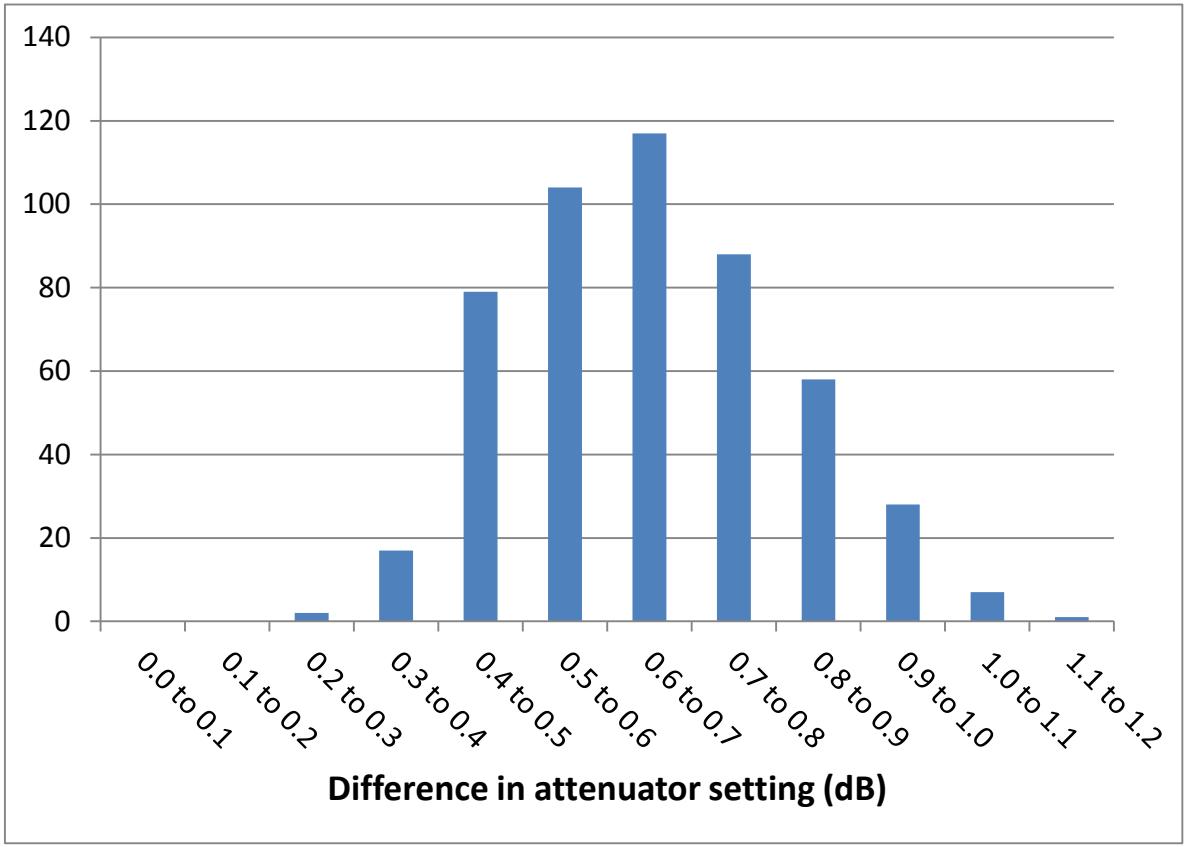

400GBASE-SR16 PMD spec option 2

The analysis method used for option 2 is illustrated below:

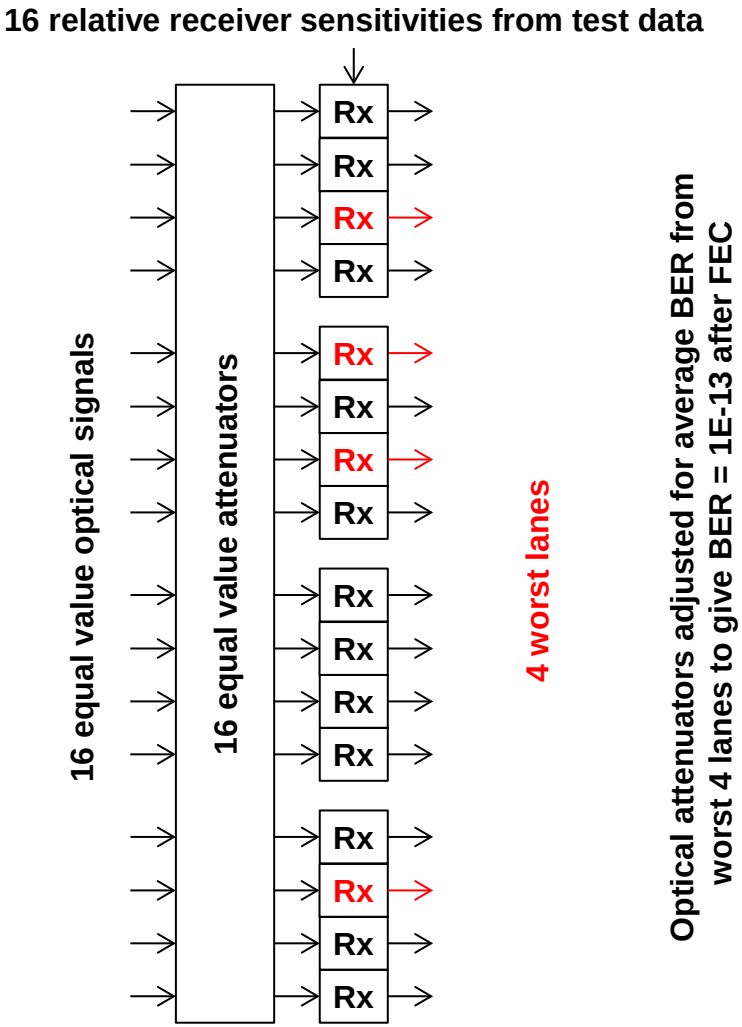

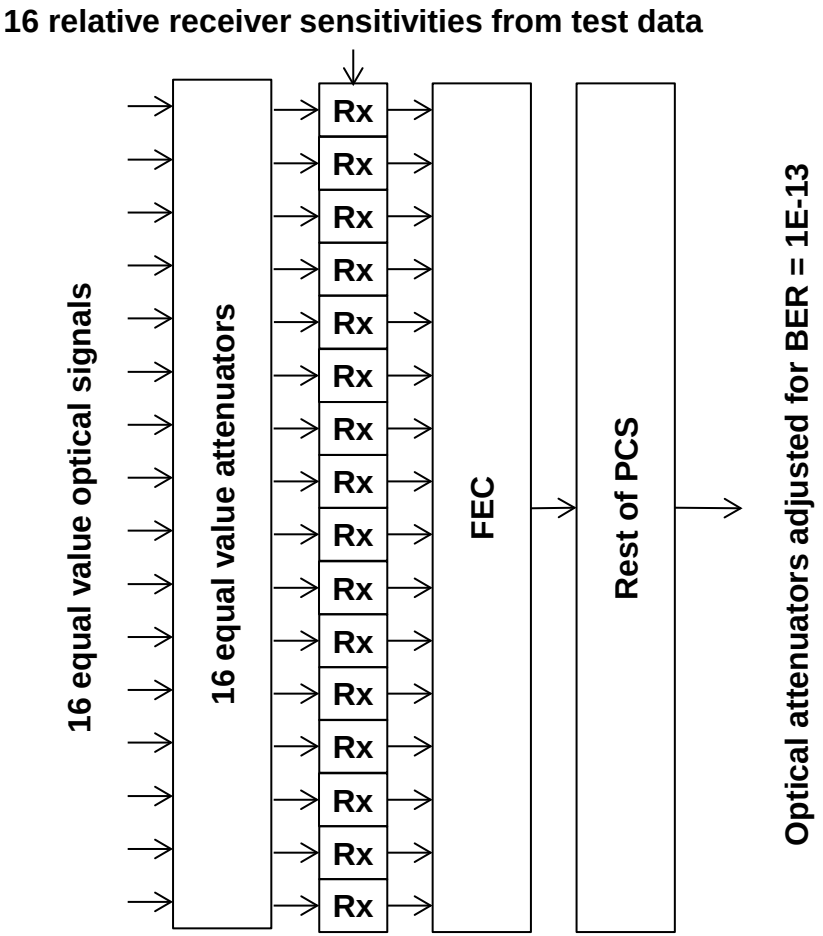

400GBASE-SR16 PMD spec option 2 results

The distribution of attenuator setting differences resulting from the option 2 analysis is shown below. The median value is 0.43 dB.

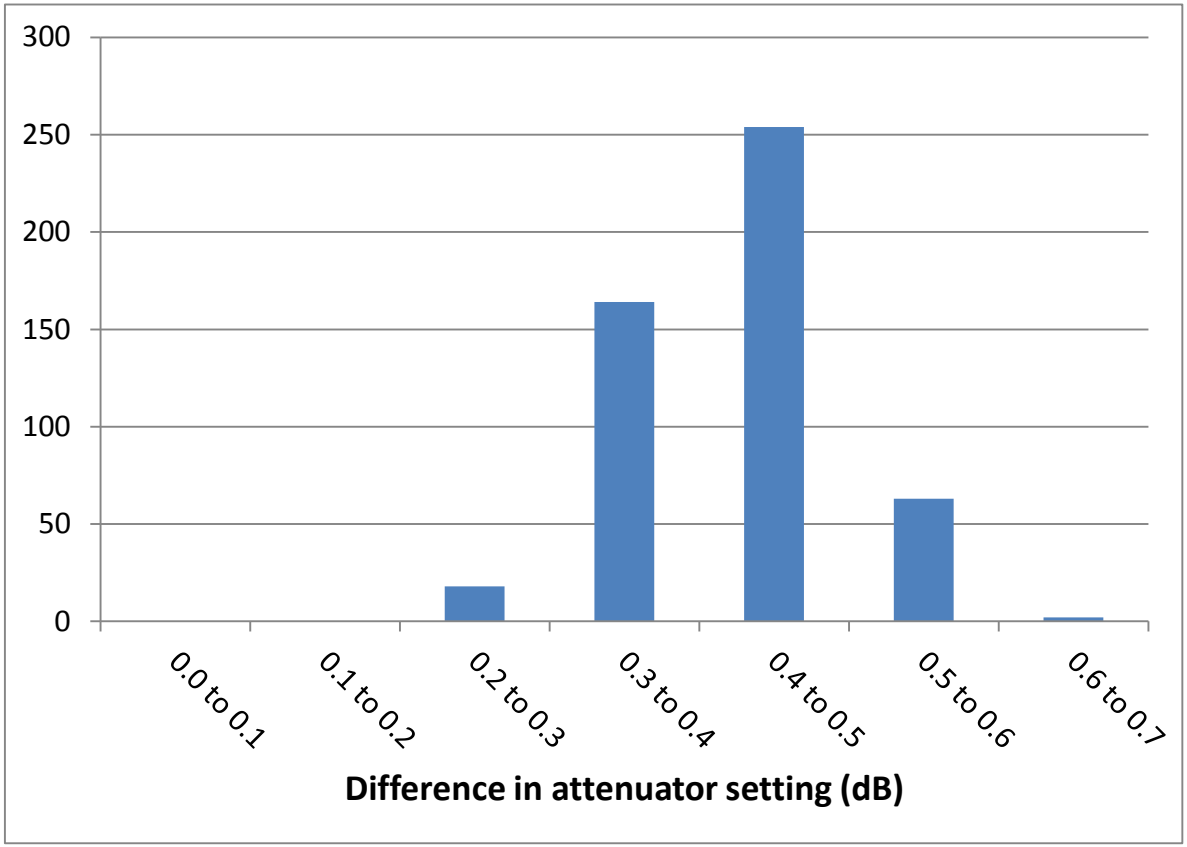

Performance conclusions 1

For a 16 lane link with a realistic Rx sensitivity distribution and equal Tx power per lane, if the powers are such as to give a 1E-13 BER with a 4 x 100G FEC scheme, the median BER for the same link with a 1 x 400G FEC scheme would be 4.1E-17.

For a link as above, but with 8 physical lanes instead of 16 the 1 x 400G FEC scheme would have a BER of 2.8E-18. Invoking FOM on the 4 x 100G FEC scheme would reduce the 1 x 400G FEC value to 4.1E-17.

For a link as above, but with 4 physical lanes instead of 16 the 1 x 400G FEC scheme would have a BER of 1.3E-19. Invoking FOM on the 4 x 100G FEC scheme would reduce the 1 x 400G FEC value to 4.1E-17.

Note: the difference in performance shown above is very conservative as it only takes receiver sensitivity variation into account. The additional variation in Tx output power and link loss across lanes will make the difference significantly larger.

Performance conclusions 2

If the FEC input BER needed for an output BER of 1E-13 is BERlim:

Option 1. For a 16 lane PMD with a realistic Rx sensitivity distribution, the difference between a PMD spec where each lane is limited to BERlim (4 x 100G) and one where the average lane BER is limited to BERlim (1 x 400G) is 0.64 dB

Option 2. For a 16 lane PMD with a realistic Rx sensitivity distribution, the difference between a PMD spec where the average of the worst 4 lanes is limited to BERlim (4 x 100G) and one where the average lane BER is limited to BERlim (1 x 400G) is 0.43 dB

Architecture conclusions

There have been many presentations on the topic of the relative implementation difficulties of 1 x 400G FEC vs. 4 x 100G FEC. Nothing has emerged from this discussion as a compelling reason to choose one scheme over the other with current technology, and it is expected that in a few years time either scheme will be easily implementable.

The simplifying assumption that all lanes have the same BER is clearly wrong for any realistic link. This presentation has shown a substantial performance advantage for a 1 x 400G scheme due to this effect and this advantage will remain in the long term.

In the presence of burst errors, 4 x 100G FEC still has significantly worse performance than 1 x 400G FEC unless a restriction is placed on the lane multiplexing so that lanes from a single FEC instance are never multiplexed together (FOM). As shown in anslow_3bs_04_0715 there are alternative methods (precoding or a modest restriction on DFE taps) to improve the performance of 1 x 400G FEC (should that be needed) that retain the freedom to route any PCS lane over any physical lane.

Consequently, it is proposed that a 1 x 400G FEC scheme is selected for 400 Gb/s Ethernet.

Thanks!