DesignCon 2019

Skew metric and BER testing correlation for NRZ/PAM4 signaling

Mauro Lai, Intel Corporation
Jeremy Stephens, Intel Corporation
Joel Ficke, Intel Corporation
Prasanthi Yelamarthy, Intel Corporation
Kyle Roberts, Intel Corporation
Devon Jenson, Intel Corporation
Milana Karthik, Intel Corporation

Abstract

TDR and frequency domain methods are used to characterize skew between P/N legs of a differential pair. Those values are then correlated to silicon measured NRZ and PAM4 BER. A weighted, frequency domain, skew metric has been introduced to close identified gaps in skew characterization and to improve correlation to BER measured on 50Gbps PAM4 SerDes. Key items such as signal power spectral density as well as TX/RX bandwidths are introduced in the metric to estimate skew limit thresholds. Extension to 100Gb/s PAM4 signaling is discussed.

Authors Biography

Mauro Lai is a Senior Analog Engineer in the Intel/DCG group. His work has been focusing on package and board signal integrity on DDR, Fabric and PCIe interfaces. Mauro has a MS in Electromagnetic Compatibility from University of Missouri Rolla and a BS in Electrical Engineering from Universita' degli Studi di L'Aquila, Italy.

Jeremy Stephens is a Principal Engineer at Intel Corporation. His work has encompassed link layer, physical layer, signal integrity, and system design solutions for HPC applications. Jeremy has a ME in Electrical Engineering from Rensselaer Polytechnic Institute and a BS from Columbia University.

Joel Ficke is an Engineering Manager at Intel Corporation. He has spent over 30 years in the HPC industry with roles in device integration, SRAM and PLL design, and most recently HSIO test and integration. Joel has a MSEE in Electrical Engineering from University of Minnesota and a BS in Chemistry from University of Wisconsin – Eau Claire.

Prasanthi Yelamarthy is an Electrical Validation Engineer at Intel Corporation. Her work focuses on signal integrity and firmware validation on High Speed Fabric I/O. Prasanthi has a MS in Electrical Engineering from Arizona State University and BS from GITAM University, India.

Devon Jenson is an Electrical Engineer at Intel Corporation. His work focuses on signal integrity and electrical validation of high-speed I/O. Devon has MS and BS degrees in Electrical Engineering from the University of Minnesota – Twin Cities.

Kyle Roberts is an Electrical Engineer at Intel in the DCG group. His work involves a variety of signal integrity and EV related issues spanning the high speed electrical interfaces of Fabric, 802.3, and PCIe systems. Kyle has a MS and BS in Electrical Engineering from the University of Minnesota Twin Cities.

Milana Karthik is an Electrical Validation Engineer at Intel Corporation (DCG group). Her work includes electrical validation, channel simulation and signal integrity study of High Speed I/O interfaces. Milana has MS in Electrical Engineering from University of Southern California and Bachelor of Engineering in Electronics and Communications from PES University, India.

1. Introduction

As the industry moves to higher line rates and modulation levels to satisfy modern bandwidth demands, the effects of skew between positive and negative (P/N) legs of a differential pair play an ever increasing role in BER impairment. This paper demonstrates these effects across 25Gbps NRZ and 50/100Gbps PAM4 receivers over a variety of channel types. As part of that analysis, it was observed that the BER impact from skew inherent in cable assemblies is substantially more benign than skew injected with phase shifters. To explain that inconsistency, detailed explanations of skew mechanisms and skew measurement methods are explored. It is confirmed that channel operating margin does not explain observed PAM4 receiver skew sensitivity. To explain these deficiencies in current understanding a new skew figure of merit (FOM) is defined. This FOM is compared to existing methods and is proven to better enable the trade-off of skew with other interconnect impairments. More work will be required from the industry to test these conclusions across a large variety of interconnect and receiver designs.

2. Skew sensitivities of NRZ and PAM4 signaling

Skew can be injected into an interconnect channel using a pair of Optical Phase Shifter (OPS) structures[1], allowing additional time-domain delay to be added to either side of the differential pair. Figures 1 and 2 show how OPS induced skew impacts measured bit error ratio (BER) on short (10dB insertion loss) and long (30dB) reach channels, respectively. The lab setup for these measurements is described in Section 11.

[1] API Optical Communications Phase Shifter DC – 50 GHz PN: OPS-0002

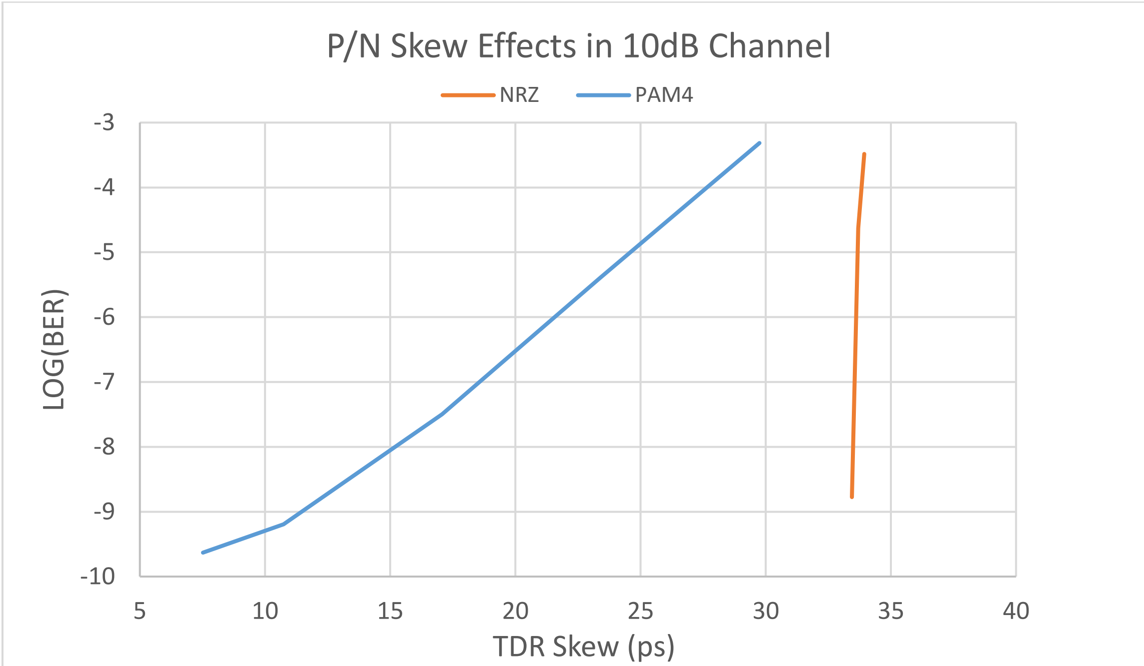

Figure 1 BER obtained on 10dB channel with OPS structure added in series

It is observed for PAM4 signaling that time-domain skew as small as 5ps can start impacting long reach channels and at 18ps BER can exceed 1E-4, revealing a fragile resilience to the impairment. On the short reach channel, skew of 10ps starts to noticeably impact BER and at 27ps BER hits 1E-4. NRZ signaling, on the other hand, is less sensitive to the same magnitude skew. NRZ BER is not impacted until 27ps on the long reach channel while skews in excess of 30ps can be tolerated on the short reach channel. As described in Section 11, the receiver was retuned for each separate skew magnitude in the experiments of Figures 1 and 2, but it is important to note that a lack of proper re-tuning will make the effects of skew more adverse.

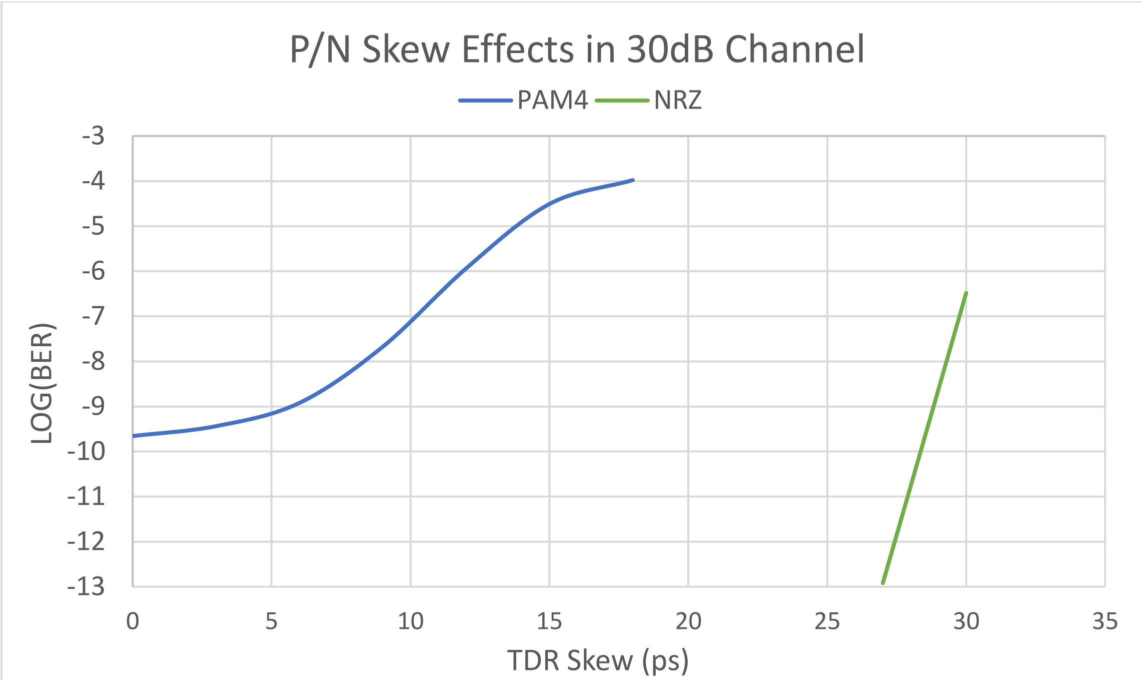

Figure 2 BER obtained on 30dB channel with OPS structure added in series

A QSFP28 cable assembly with approximately 30dB insertion loss was characterized with a VNA and each lane's BER was measured according to the procedure in Section 11. The time domain skew for each lane in the cable assembly is plotted against BER in Figure 3. Observe that skew in excess of 40ps has less than four decades of impact on measured PAM4 BER magnitude.

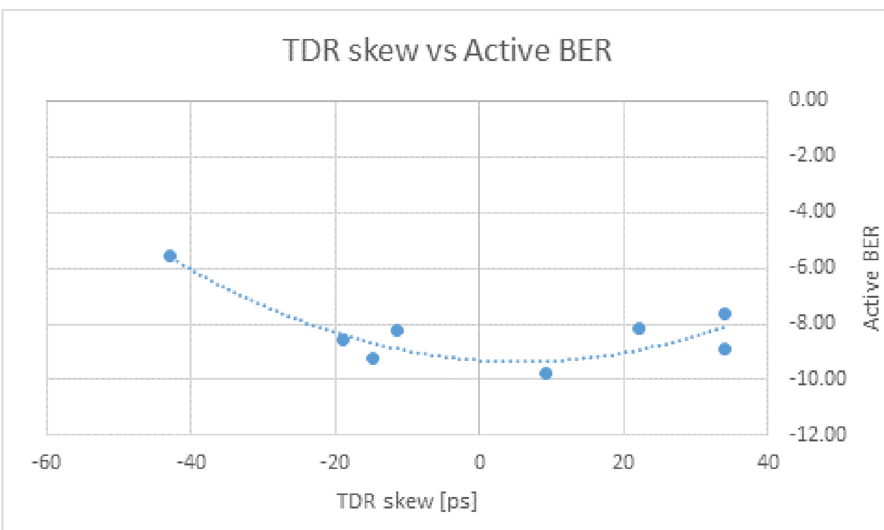

Figure 3 BER obtained on channel with a 3.8m cable assembly with inherent skew variation across differential pairs

A skew of 40ps injected with OPS structures would be fatal for the long reach PAM4 channel, yet it is relatively benign if the skew is due to the cable assembly of Figure 3. In order to explain that inconsistency between Figures 2 and 3, the next three sections provide more detailed definitions of skew measurement methods and skew mechanisms.

3. Time-domain skew measurement

A common method to measure skew in the time domain involves the injection of a voltage step response and observation of the delayed reflection arrival times (TDR). Defining skew with TDR poses challenges for non-ideal systems having mismatched impedance or coupling [1], [2]. Figures 4 and 5 demonstrate how the measured magnitude of skew between the P/N legs of a differential pair depends on the reference voltage level chosen to calculate the time delta for both cable assemblies and OPS structures. Figure 4.b shows that the measured time-domain skew of the cable assembly varies wildly with reference voltage.

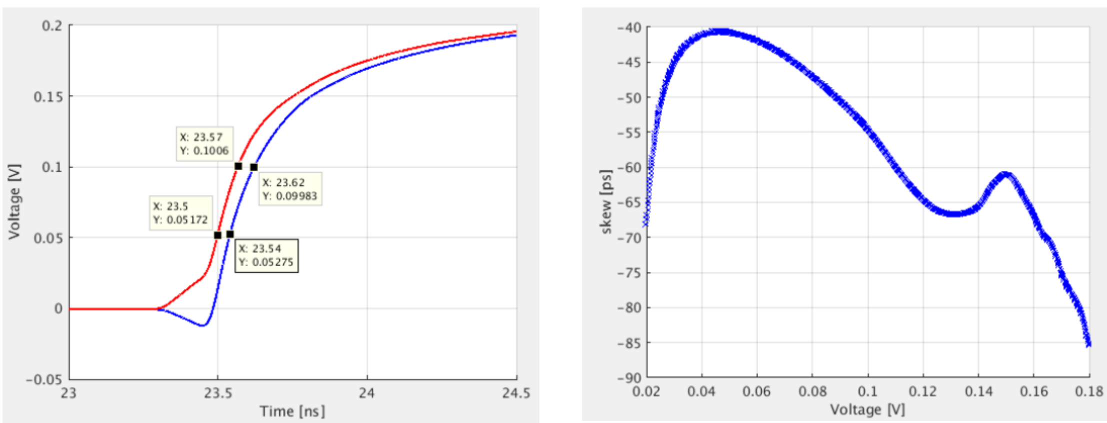

(a) (b)

Figure 4 TDR of differential pair in a 3.8m cable and skew vs voltage trend.

Compare how the measured skew magnitude varies as the reference voltage is selected within the range of 0.02-0.1V in Figure 4 to the OPS skewed channel in Figure 5. The OPS system allows for a reliable skew estimate from a TDR measurement because the impedance of the structure is more uniform and there is no coupling between the true and complement of the differential pair (more details in Section 5).

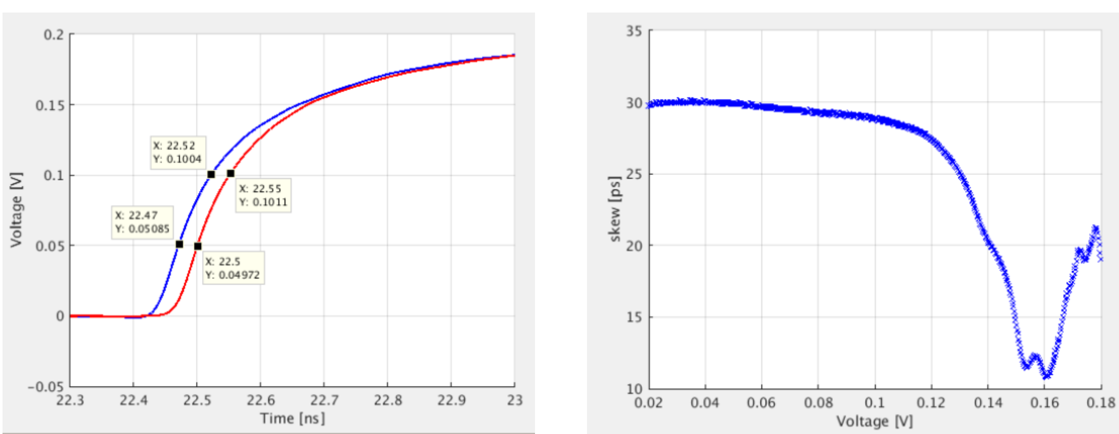

(a) (b)

Figure 5 TDR of OPS skewed channel and skew vs voltage trend

It may be concluded that time-domain skew from a TDR measurement is not a reliable metric for skew measurements across a large set of interconnects, cabled channels being the most problematic. A frequency domain approach to quantify skew is defined in the next section to address weaknesses of time-domain skew measurements.

4. Frequency-domain skew measurement

Frequency domain measurements of skew are more reliable [3], [4] because conversion to partial mixed-mode s-parameters allows for the effects of reflections as well as far end coupling to be considered. Equation (1) defines a frequency-domain skew metric skew(f), with units of seconds. SCD21(f) and SDD21(f) are both available from 4-port s-parameter measurements or simulation and are typically used to calculate insertion loss IL(f).

$$skew(f) = \frac{\angle\{SDD21(f) + SCD21(f)\}}{2\pi f\sqrt{2}} - \frac{\angle\{SDD21(f) - SCD21(f)\}}{2\pi f\sqrt{2}} \quad (1)$$

Skew(f) spectra for OPS structures and for a QSFP28 cable assembly are shown in Figure 6. The two types of spectra are qualitatively different in that the OPS skew is frequency-independent out to 25GHz while the cable assembly skew reveals its frequency-dependent nature well before 5GHz. The cable assembly skew is large at low frequency, but begins to decay as frequency increases and eventually oscillates with small amplitude, as described in [2].

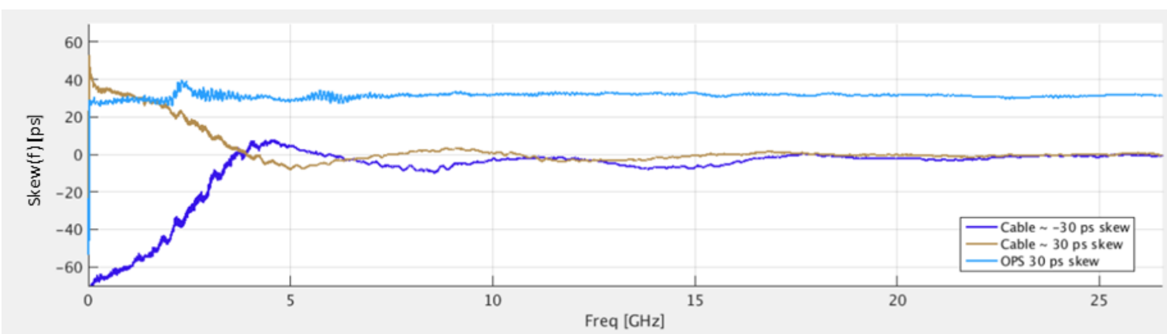

Figure 6 Skew(f) for Cable with +30ps skew (--), -30ps skew (--) and OPS structure (--)

By comparison of the skew magnitudes of Figure 6 to those of Figures 4 and 5, it is apparent that the TDR skew measurement method of the coupled differential line emphasizes low frequency skew. However, it is not clear why low frequency skew would be detrimental to a communication channel's BER. In fact the cable assembly's skew(f) spectrum is shown to be benign in a PAM4 channel, implying that low frequency skew is tolerable in an operating PAM4 system. This is an extension of the results shown in [2] where the high level of coupling between two signal paths results in a smaller effective skew.

IEEE 802.3 [5] defines a cable assembly skew limit in the form of a frequency-domain mask for the spectra defined by the following Equation (2).

$$Conversion\_loss(f) - IL(f) \geq \begin{cases} 10 & 0.01 \leq f < 12.89 \\ 27 - (29/22)f & 12.89 \leq f < 15.7 \\ 6.3 & 15.7 \leq f \leq 19 \end{cases} \tag{2}$$

where

f	is the frequency in GHz
$Conversion\_loss(f)$	is the cable assembly differential to common-mode conversion loss
$IL(f)$	is the cable assembly insertion loss

Figure 7 shows the spectra of the difference between the insertion loss and conversion loss. This metric will be referred to as SCD21-SDD21 and is plotted for both the OPS structures and example QSFP28 cable assemblies along with the mask from [5]. The OPS assembly clearly violates the mask at a much lower frequency than the cable assembly, supporting the observations from Section 2.

The spectra of SCD21-SDD21 overtly demonstrates why OPS skew impacts SerDes performance so drastically, as mode conversion is substantial across the entire frequency range. Extensive research has been performed on ways to mitigate skew in PCBs [6] and cables [7], along with ways to provide models for “unbalanced lines” [8]. However, this paper’s main focus is to provide a quantitative performance impact caused by such imbalances.

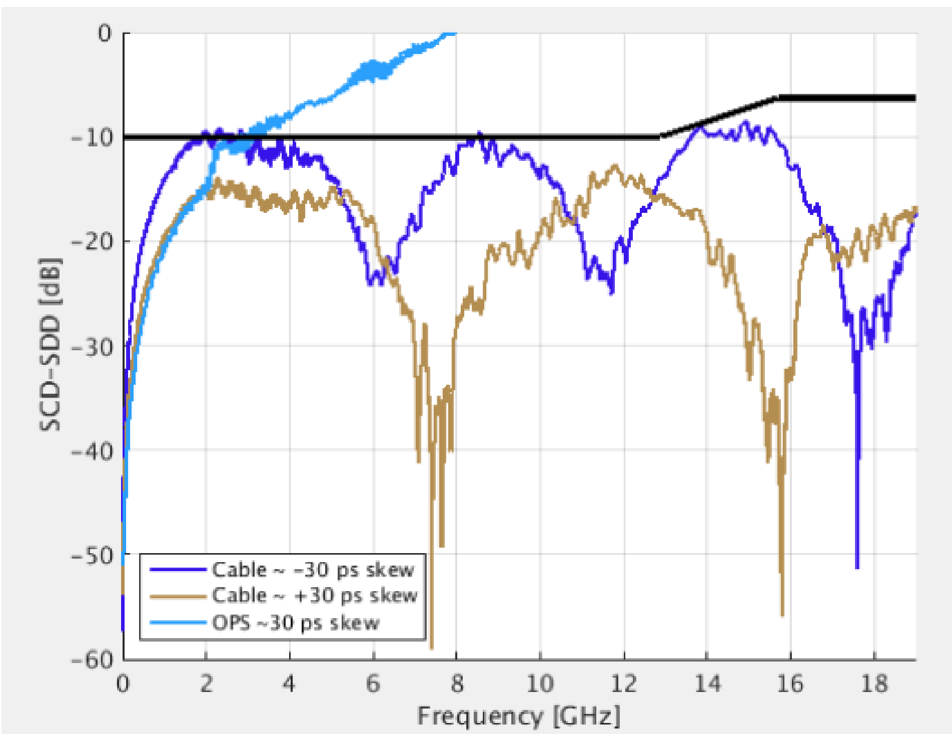

Figure 7 SCD21-SDD21 spectra for cable with +30ps skew (--), -30ps skew (--) and OPS structure (--)

Quantifying the BER performance impact of a mask violation is not trivial. A violation in a small frequency band may be harmless in some cases if it is compensated for with substantial margin in other wide frequency bands. In other cases, however, such a violation maybe fatal.

In other words, frequency-domain masks do not allow for the trade-off of one frequency band's violation against another frequency band's margin. Nor does the mask allow for trading off SCD21-SDD21 mask violations for margin in other interconnect impairments. A new skew figure of merit is therefore introduced in Section 7.

5. Skew Mechanisms

As mentioned in Section 3, skew in OPS structures can be quantified with the TDR method (Figure 5), but cable assemblies represent a more complicated scenario (Figure 4) since there is a wide variation in the measured skew depending on the voltage reference level [4].

To compare the skew mechanisms, two models have been generated. Skew model type 1 injects skew with conductor length mismatch between P/N legs of the differential pair, and is shown in Figure 8.a. An ideal time delay element similar to the OPS structure is inserted in series with the differential pair. The resulting 'uncoupled' skew(f) in Figure 8.b is frequency-independent and is referred to as such for the remainder of this document.

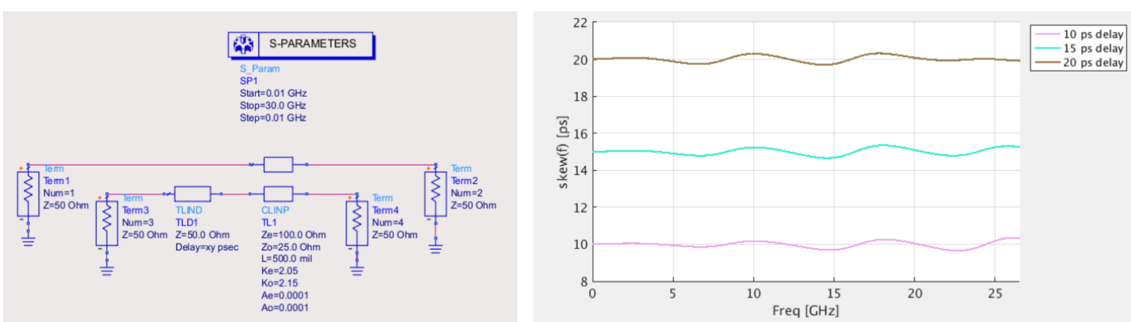

(a) (b)

Figure 8 Skew model type 1. ADS circuit schematic for the conductor length mismatch (a) and frequency-independent skew(f) profiles for 10, 15 and 20ps skews.

Skew model type 2 injects skew with dielectric asymmetry as shown in Figure 9.a. The dielectric asymmetry may be interpreted as a difference in impedance in the P/N wires (similar to a twin differential pair). The resulting skew(f) shown in Figure 9.b is frequency-dependent, as described in [2]. Note that the period of oscillation varies based on model parameters selected as shown in Figures 9.b and 9.c and that this effect is also observed on characterized lab samples. Frequency-dependent skew is also referred to as 'coupled skew' in the literature, but will be referred to as frequency-dependent for the remainder of this document.

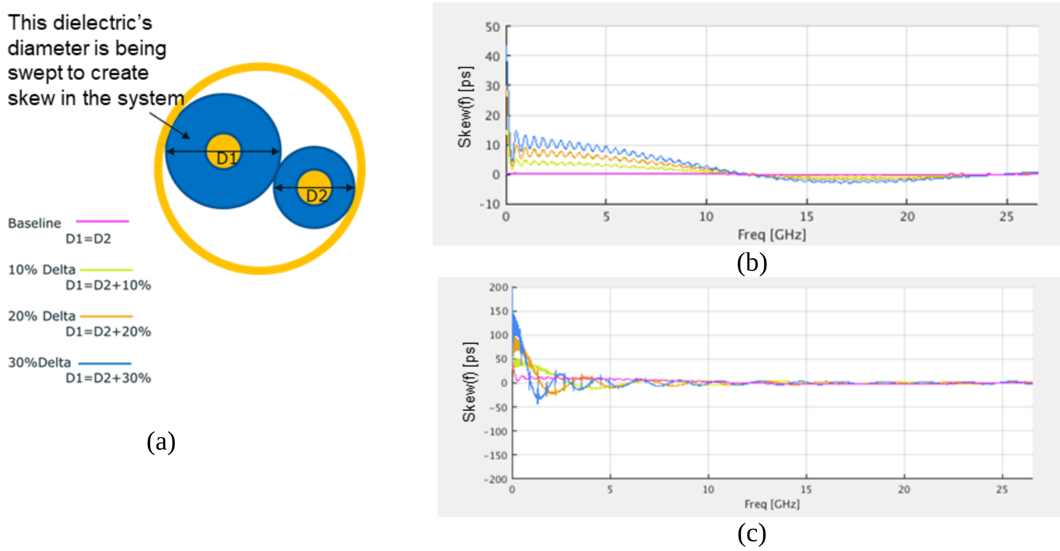

Figure 9 Skew model type 2. Dielectric asymmetry structure (a) and frequency-dependent skew(f) for different diameter mismatch (color coded for baseline 10%, 20%, 30%) and different cable length 0.25m (b) and 3m (c).

The observed skew(f) profiles for OPS structures and QSFP28 cable assemblies shown in Figure 6 are qualitatively different, but can be described well with this section's skew model types 1 and 2, respectively.

Lastly, note that skew in a real system channel may consist of some combination of these two effects. For example, length mismatch in a printed circuit board traces would introduce frequency-independent skew behavior. If those traces are connected to a cable assembly, frequency-dependent skew behavior could also be added to the same channel. The nature of how these types of skew combine has been subject to limited investigation in both lab and simulation, but is not yet fully understood from a mathematical perspective. In some cases a seemingly linear addition of the skews from separate structures is observed in the overall channel, however, other cases appear to suggest otherwise. More comprehensive studies on this topic will be required to make further conclusions that allow for application to interconnect design.

6. Skew impact on COM

As mentioned earlier, frequency mask type limits as shown in Section 4 and [5] do not enable any interconnect impairment trade-off analysis. For example, it is not clear how much insertion loss (IL) should be reduced to accommodate SDC21-SDD21 mask violations in Figure 7.

Recent high-speed interconnect standards rely on channel operating margin (COM) to specify compliant channels. COM is calculated according to [9] for end-to-end interconnect channels with OPS-injected skew. In order to compare COM results to measured silicon receiver BER, the BER target for COM is varied until the COM output score is equal to 1dB. In the lab, measured silicon BER is averaged over ten receiver tunes. Figure 10 compares the BER at which COM equals 1dB to the average measured silicon BER for several OPS-injected skew settings in a channel with a 1.5m QSFP28 cable assembly. The reader should note that the frequency-dependent skew of this cable assembly has been characterized as an insignificant magnitude to warrant consideration. BER at a COM of 1dB does degrade as skew is added to a channel, but it is much less sensitive than the measured silicon BER for the tested PAM4 receiver. In this case, COM clearly underestimates the impact of frequency-independent skew on the receiver BER. Although not shown here, a similar result for frequency-dependent skew cases exists.

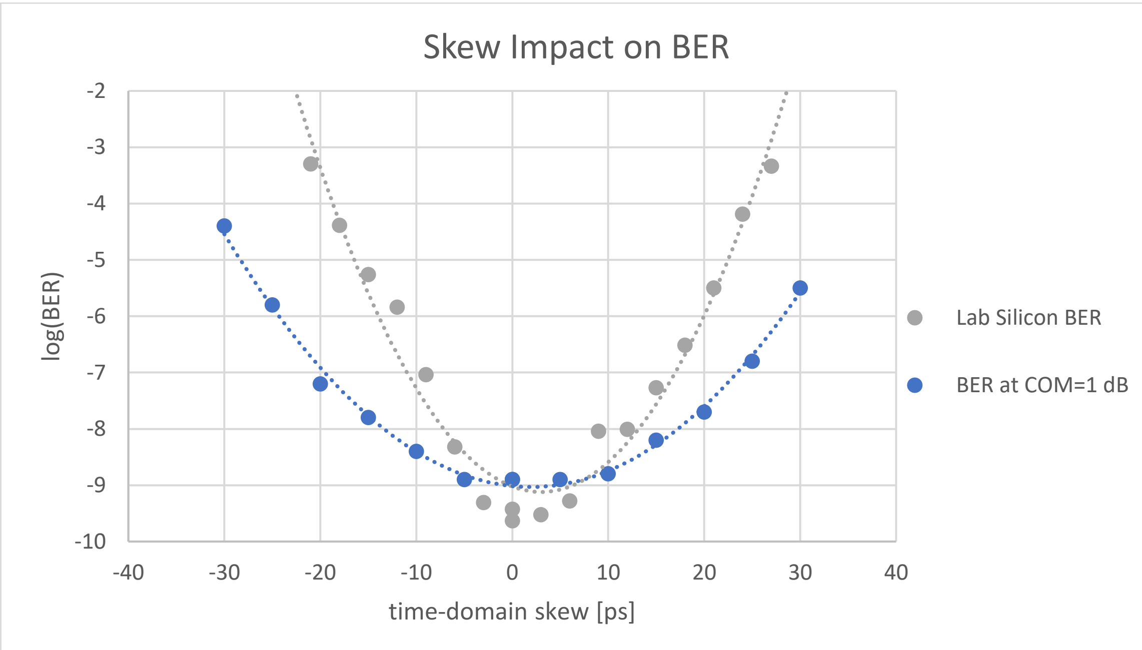

Figure 10 Measured and calculated BER for OPS structure

COM's relatively dismissive attitude towards skew stems from the fact that skew causes mode conversion, resulting in some non-zero common mode signal arriving at the receiver input. Typically a differential receiver will reject a large portion of this common mode signal, as characterized by its common mode rejection ratio (CMRR). The remaining common mode signal is then converted back to differential noise at the receiver decision circuit, reducing SNR. The implementation of COM in [9] does not account for CMRR or common-mode noise at the receiver input. Therefore, COM may be under estimating the noise term in its SNR calculations for interconnect with significant skew.

7. Skew FOM

In the course of high-speed channel design, it is desirable to trade-off several interconnect impairments with one another. However, frequency mask limitations do not enable such desired impairment trade-off analysis. The COM score is better suited for such a task and may be used as a channel design metric to trade-off the following signal integrity impairments: insertion loss, return loss, and crosstalk. However, the COM score's underestimate of skew impact on channel performance makes it a poorly suited metric to trade-off skew with other impairments. A new skew figure of merit (FOM) is defined for the purposes of a more complete interconnect impairment trade-off analysis.

Equation (3) defines FOM skew as an integrated frequency-domain skew metric weighted by a random bit stream's power spectral density and relevant TX and RX bandwidth filters, following the form of FOM_ILD [10] and ICN [11],

$$FOM_{skew} = sign\big(skew(f)\big) * \frac{df}{f_b} * 1e12 * \sum_{fmin}^{fb} \{|skew(f)| * W(f)\} \quad [ps] \tag{3}$$

where fb [symbols/second] and fmin [Hz] are the baud rate and the lower frequency integration limit respectively.

The weighting function W(f) is defined in Equation 4 and models the power spectral density (PSD) of the signal (NRZ or PAM4), and includes two Butterworth filters representing TX and RX bandwidth.

$$W(f) = sinc^2\left(\frac{f}{fb}\right) * \frac{1}{1+\left(\frac{f}{ft}\right)^4} * \frac{1}{1+\left(\frac{f}{fr}\right)^8} \quad for \quad f_{min} < f < f_b \tag{4}$$

Note that sinc refers to the normalized sinc function[2] while ft and fr [Hz] are respectively the TX and RX rise time Butterworth filters.

Figures 11 and 12 compare the time-domain skew to FOM_skew for OPS (frequency-independent) and a cable assemblies (frequency-dependent) respectively.

[2] Defined as sin(πx)/(πx)

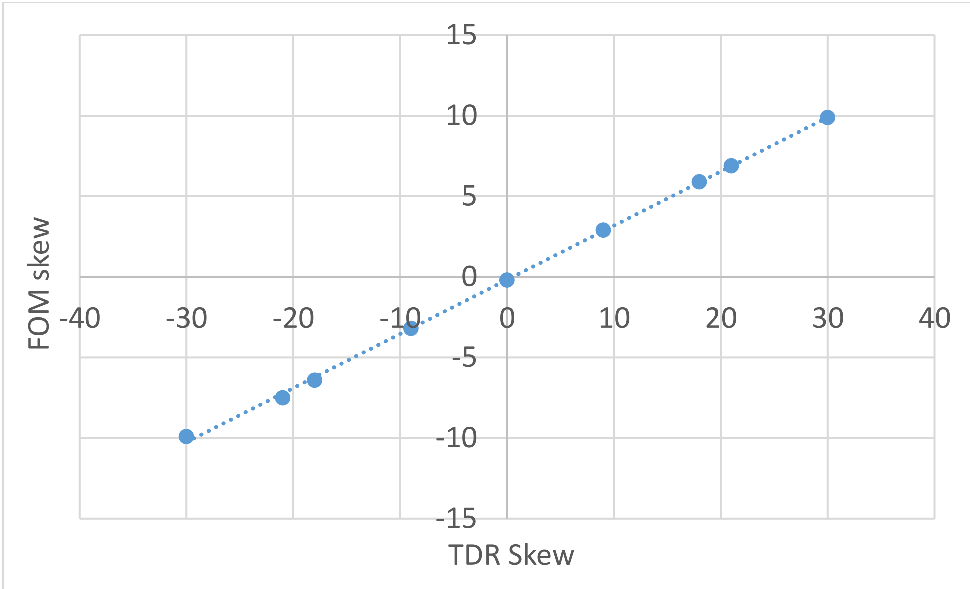

Figure 11 Time-domain skew vs FOM skew for OPS

Note that TDR-measured skew of the OPS structures correlates linearly with FOM skew in Figure 11. FOM skew can therefore be predicted from the TDR-measured skew if skew(f) exhibits frequency-independent behavior. By comparison, Figure 12 shows that TDR-measured skew of the cable assembly is quite uncorrelated to FOM skew; however, it is observed that the magnitude of FOM skew is reduced significantly as compared to TDR skew. This agrees with the earlier observation that frequency-dependent skew has less impact on channel BER than its TDR magnitude would suggest.

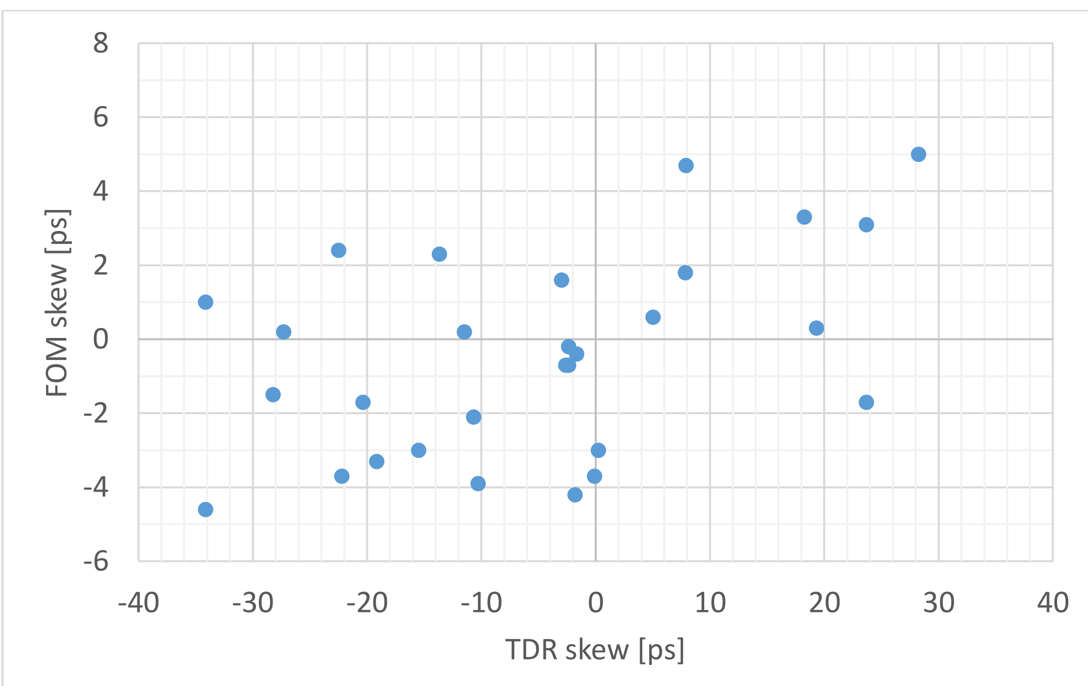

Figure 12 Time-domain skew for many measured cables

FOM skew continues to emphasize COM's short-comings for predicting the effects of skew, as shown in Figure 13:

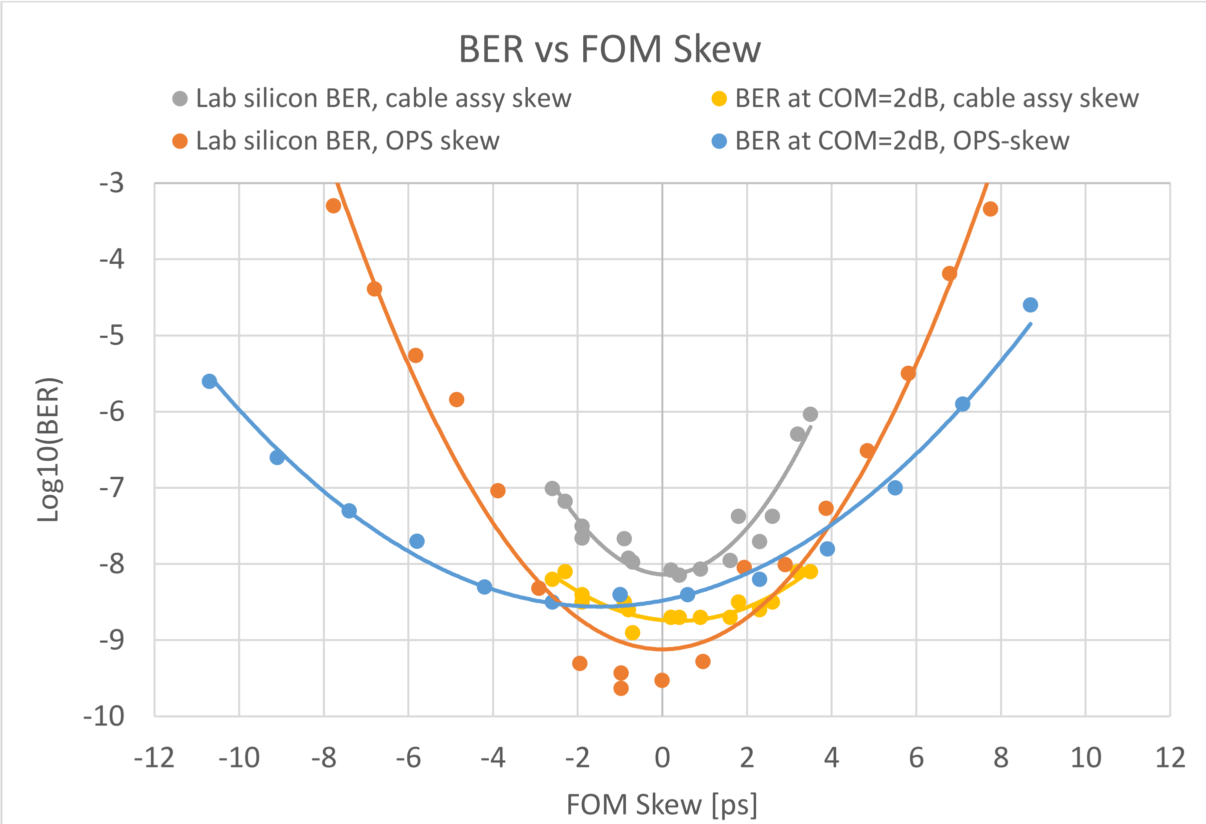

Figure 13 FOM skew for OPS (frequency-independent) and Cabled (frequency-dependent) channels.

Once curves such as those shown in Figure 13 are established for a given SerDes IP, trade-offs between channel impairments that COM models well and FOM skew can be made. It would be ideal if a universal metric could be included in COM independent of SerDes characteristics, but that will require more work by the industry to study and adopt.

In the next section both the TDR skew and the FOM skew on several additional cable assemblies are evaluated to see which one can predict actual silicon performance based on a passive characterization of the channel.

8. Validation of FOM

This section demonstrates FOM skew's ability to predict silicon BER's best and worst cases performances across a set of channels. The tested channels consist of a 2m OSFP cable assembly along with three additional QSFP28 cable assemblies (4m, 2.8m, and 1.8m). Measured silicon BER results with 50Gbps PAM4 signaling is compared to the FOM skew metric defined in Section 7, time domain skew , and the SCD21-SDD21 frequency domain mask.

The 2m OSFP cable assembly channel results are shown in Figure 14 and Figure 15. The two worst silicon BER lanes (TX7-RX7 and TX3-RX3) are predicted by all three metrics, though time domain skew shows a near three-way tie for the second worst lane.

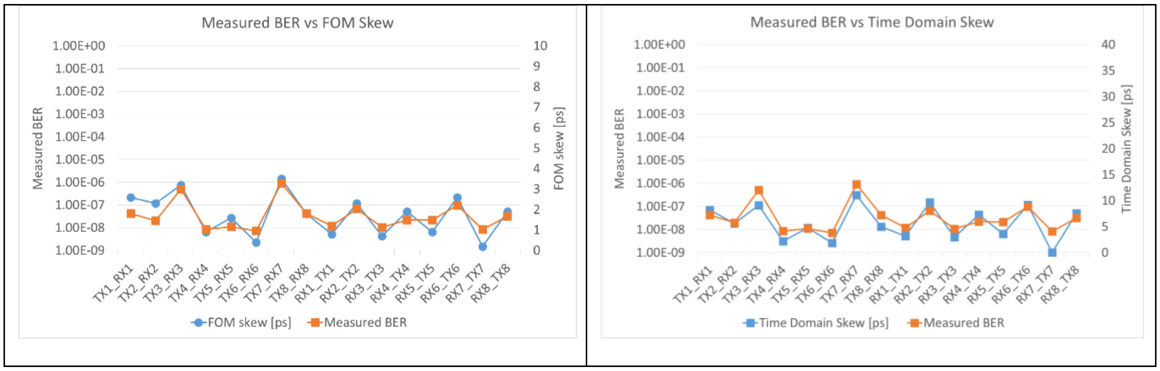

Figure 14 Time Domain skew and FOM skew vs Measured BER for a well behaved 2m OSFP cable

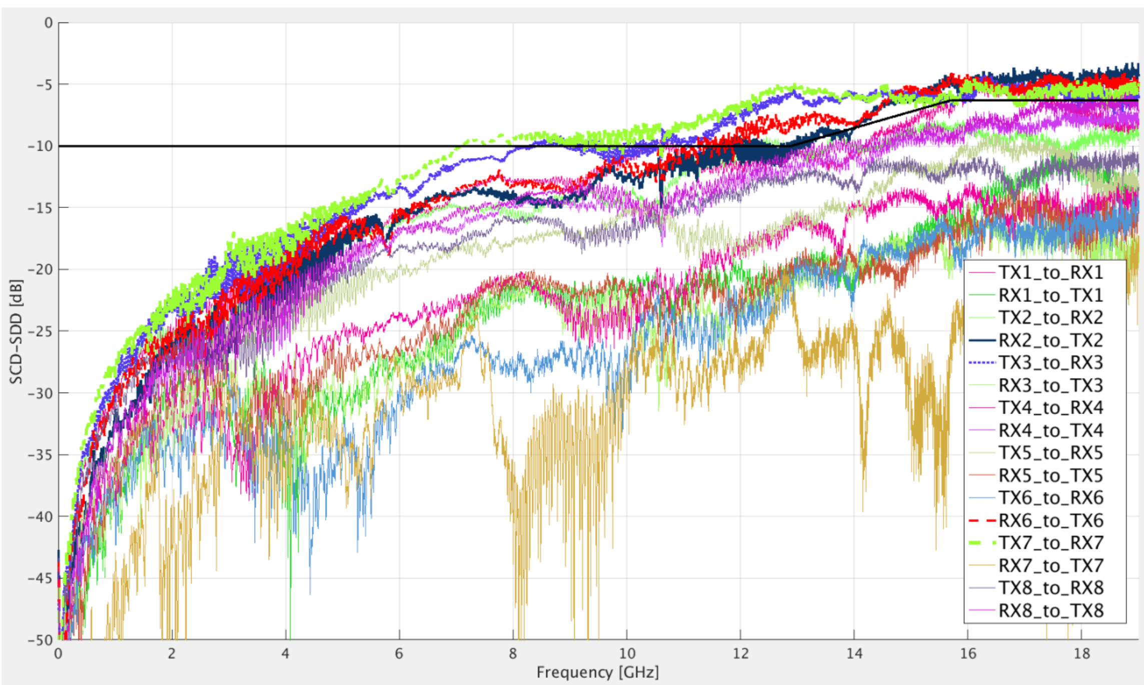

Figure 15 SCD21-SDD21 for a 2m OSFP cable

The 4m QSFP28 cable assembly channel results are shown in Figure 16 and Figure 17. The worst silicon BER lane (RX3-TX3) is predicted by FOM skew and SCD21-SDD21 compliance. Time domain skew does not track the silicon BER at all, incorrectly predicting seven other lanes to have worse BER than RX3-TX3.

RX4-TX4 violates the SCD21-SDD21 mask between 12GHz and 18GHz, but that infraction appears to be benign according to silicon BER, time domain skew, and FOM skew. SCD21-SDD21 indicates a false fail on RX4-TX4 because RX4-TX4 has the lowest silicon BER.

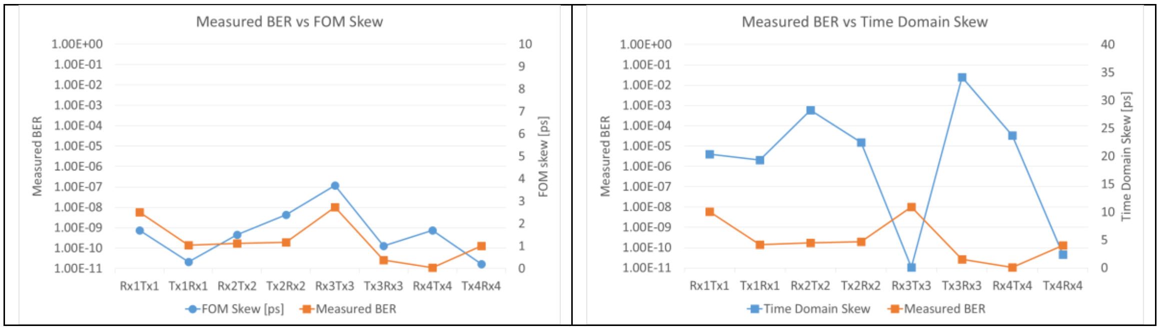

Figure 16 Time Domain skew and FOM skew vs Measured BER for a 4m QSFP28 cable

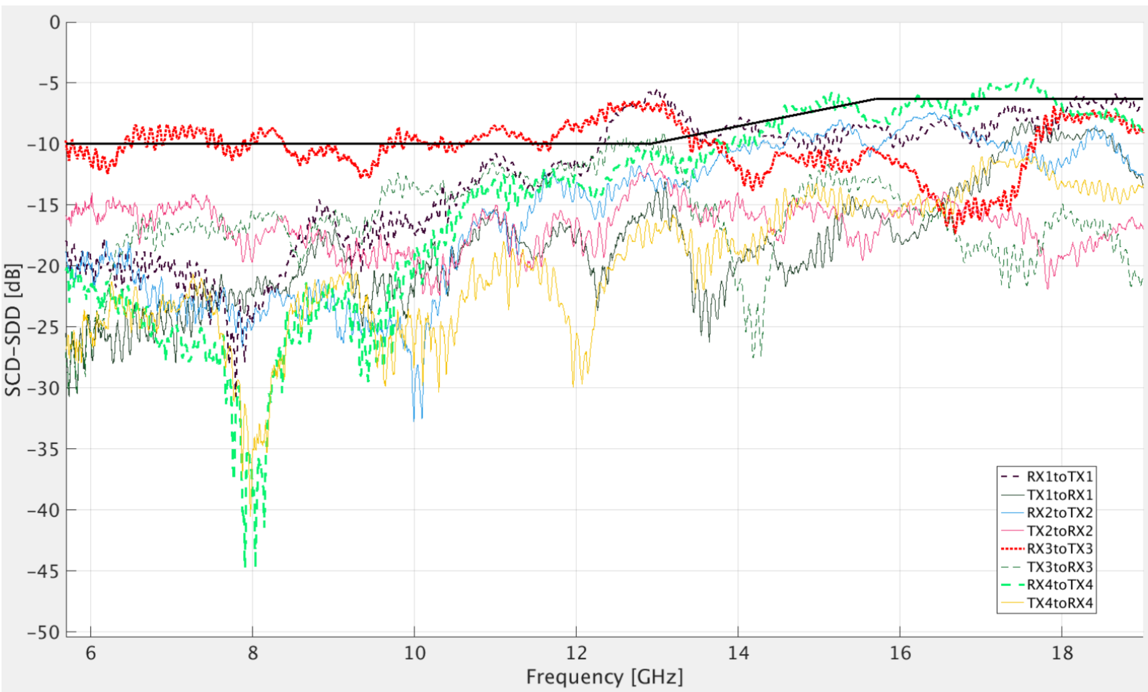

Figure 17 SCD21-SDD21 for a 4m QSFP28 cable

The 2.8m QSFP28 cable assembly channel results are shown in Figure 18 and Figure 19. The worst silicon BER lane (TX2-RX2) is again predicted by all three skew metrics, however, the second worst silicon BER lane (TX3-RX3) is not predicted by any of the skew metrics.

RX4-TX4 violates the SCD21-SDD21 mask around 12GHz, but that violation once more appears to be benign for silicon BER, time domain skew, and FOM skew. SCD21-SDD21 indicates a false fail on RX4-TX4 because RX4-TX4 has nearly the lowest silicon BER.

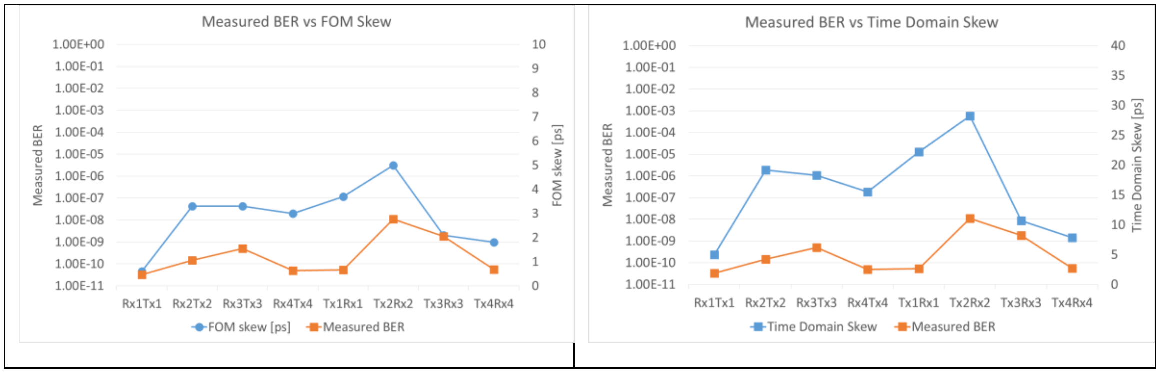

Figure 18 Time Domain skew and FOM skew vs Measured BER for a 2.8m QSFP28 cable

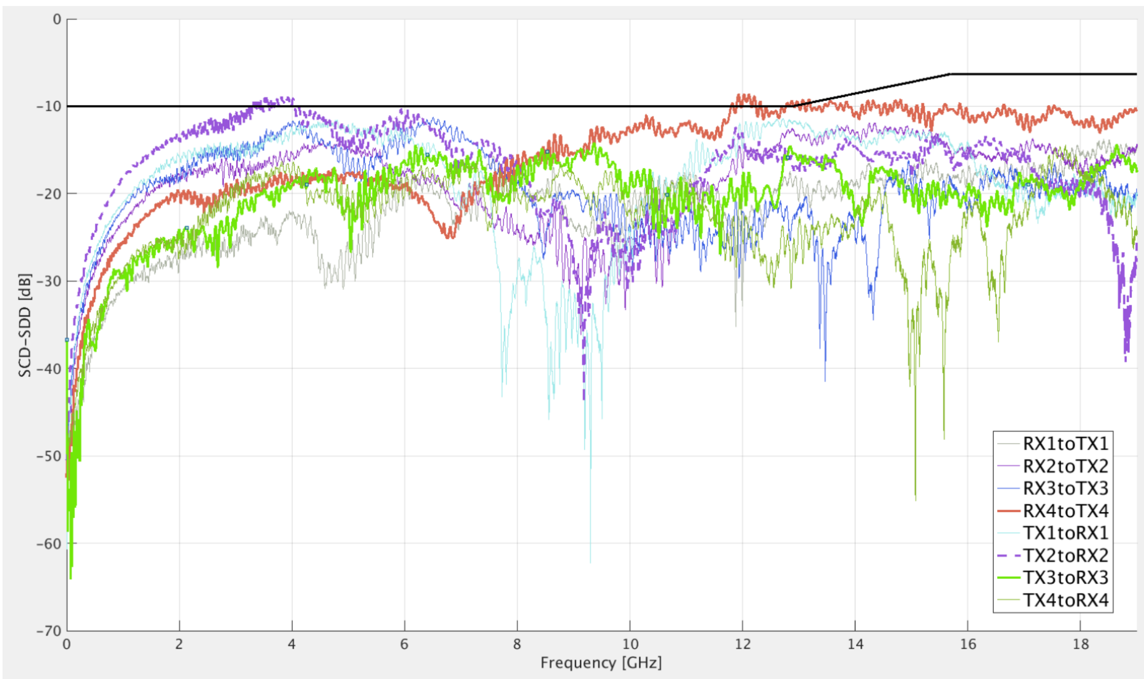

Figure 19 SCD21-SDD21 for a QSFP28 2.8m cable

The 1.8m QSFP28 cable assembly channel results are shown in Figure 20 and Figure 21. The worst silicon BER lane (RX2-TX2) is predicted by FOM skew and SCD21-SDD21 compliance. Time domain skew does not track the silicon BER at all, incorrectly predicting 3 other lanes to have worse BER than RX2-TX2.

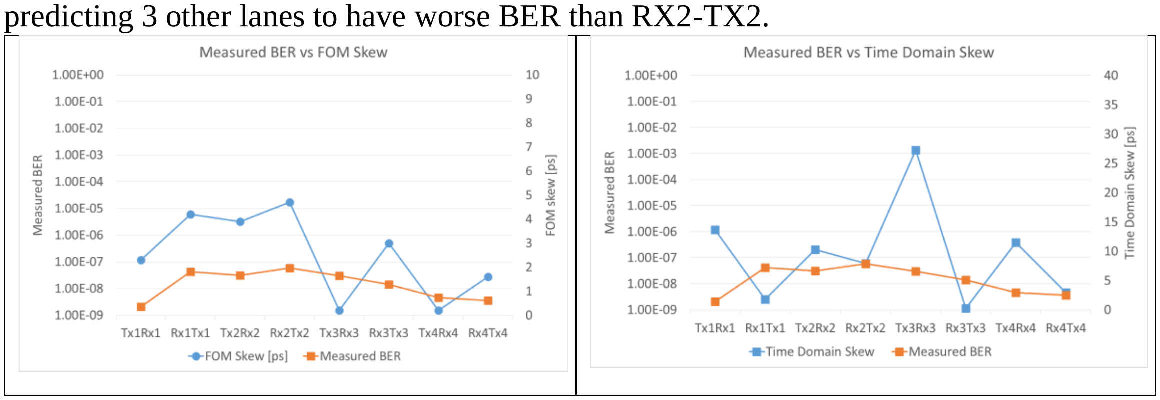

Figure 20 Time Domain skew and FOM skew vs Measured BER for a 1.8m QSFP28 cable

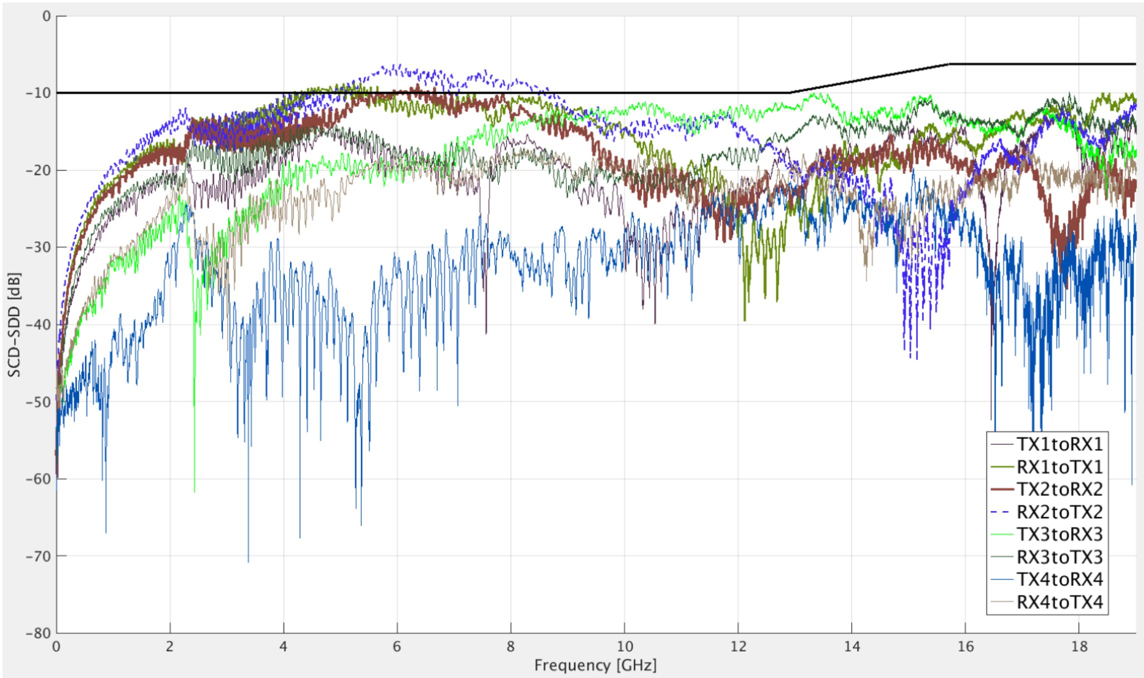

Figure 21 SCD21-SDD21 for a 1.8m QSFP28 cable

To summarize, of the four channels presented in this section, the FOM skew metric correctly predicts the worst case BER lane in every case. Time domain skew correctly predicts the worst BER lane in only 2 of 4 cases. SCD21-SDD21 correctly flags the worst BER lane in 4 of 4 cases by violations of the frequency domain mask, however it also flags false fails in 2 of 4 cases. Consequently, FOM skew is a good candidate metric to enable trade-off analysis with other quantitative channel pass/fail metrics such as BER or COM score.

9. Skew sensitivity for 50/100G SerDes IPs

Figure 1 shows the difference in skew sensitivity between NRZ and PAM4 modulation on a single 50Gbps SerDes design. Figure 22 compliments those results with skew's effects on 100Gbps PAM4 signaling. It is clear that skew will increasingly debilitate system performance at higher signaling rates. Although only a single representative result for each data rate is shown, all of these trends have been confirmed across multiple SerDes designs, emphasizing a pervasive industry challenge.

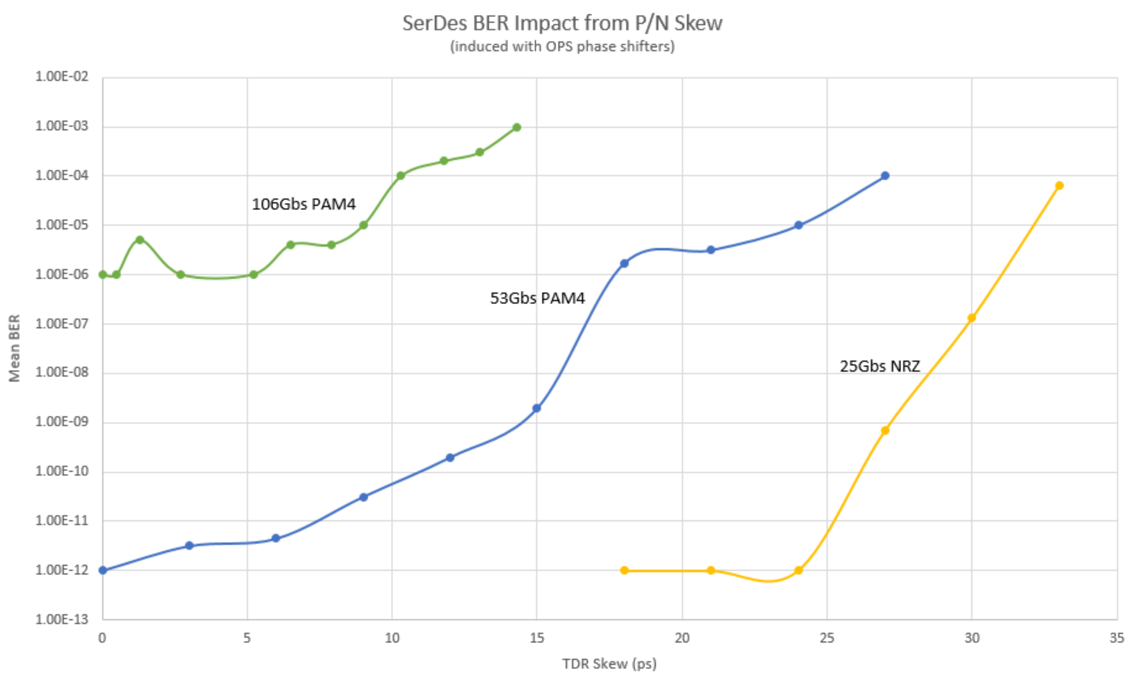

Figure 22 Skew sensitivity for 25Gb/s NRZ and 50/100Gb/s PAM4 receivers

10. Conclusions

The effects of skew on high-speed interfaces are increasingly devastating as the move to higher line rates and modulation levels becomes necessary to satisfy modern bandwidth demands. Typical NRZ interconnect channels of the past exhibited an impressive resilience to skew. Though PAM4 modulation is increasingly popular in mainstream interconnect standards, it is significantly more susceptible to skew and therefore requires more careful interconnect design and characterization.

Current methods for characterizing interconnect skew are demonstrated to be insufficient, particularly in recognizing and reconciling the differences between frequency-dependent and frequency-independent skew. In order to allow for more accurate interconnect design and performance analysis, a new metric for skew characterization is introduced. This method builds upon previously defined measures of frequency domain skew with the addition of a weighting function meant to properly model the PSD of the signal along with the bandwidths of the communication channel's transmitter and receiver.

This new FOM has been initially vetted with promising results in predicting PAM4 receiver behavior in the presence of various types of skew and should provide a more consistent metric for interconnect impairment trade-off analysis during 50 and 100 Gbps PAM4 interconnect design and characterization.

11. Appendix: Setup for skew experiments

The experiments to quantify BER impact due to Channel (OPS, or skewed cable) are shown below:

a. Experiments with OPS:

In a PVT controlled environment, channels with total insertion loss ranging from SR (short reach) to LR (long reach) were used along with OPS between the differential pair of the SerDes transmitter and receiver. For accurate performance comparison, all the experiments were run on the same SerDes lane for a fixed number of receiver equalizer tune repetitions. For each tune, errors were counted over a fixed dwell time. Channel skew is increased on one of the two OPS and error counter is restarted. After each change in OPS setting, the receiver is retuned. In order to remove the impact of skew outside the OPS, errors were counted in both direct and reverse connection. The average BER over the fixed number of tunes is used as the metric of comparison. VNA measurement of the entire channel is performed for each OPS setting and the s-parameters are post-processed to calculate skew.

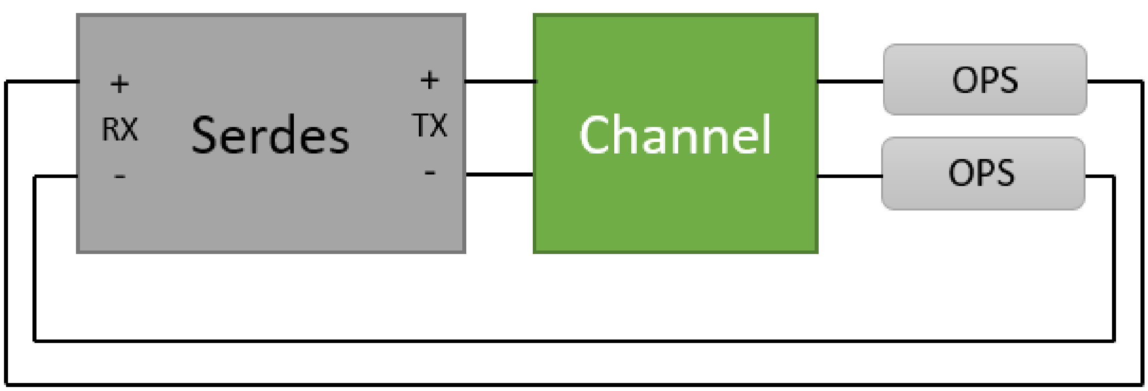

Figure 23 Setup with OPS.

b. Experiments with cable assembly:

A selected number of cables with inherent skew were used in this experiment. In a PVT controlled environment, the differential pair of each QSFP lane of the cable was connected between the differential pair of the SerDes transmitter and receiver using a breakout board or module compliance board (MCB) compatible with the connectors in the cable assembly. For accurate performance comparison of all the skewed cables, the same SerDes lane is used for all tests and a fixed number of tunes were performed on each lane of each cable taking one lane at a time and BER was measured for a fixed dwell time. In order to remove the impact of skew outside the channel setup {MCB+CABLE+MCB}, the receiver tuning and BER measurements were performed in both direct and reverse direction of each physical wire (QSFP lane). The average BER over the fixed number of tunes is used as the metric of comparison. VNA measurement of the entire channel is performed for each lane in the cable assembly and the s-parameters are post-processed to calculate skew.

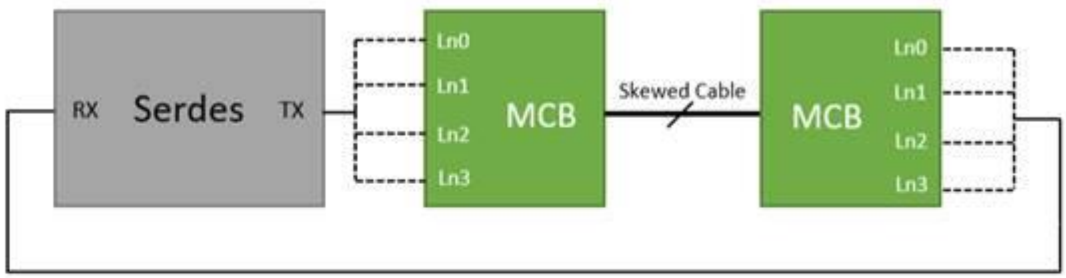

Figure 24 Setup with cable assembly.

12. References

[1] E. Bogatin, *Signal Integrity - Simplified*. Prentice Hall, 2003.

[2] E. Mayevskiy and J. Huffaker, “Limitations of the Intra-pair Skew Measurements in Gigabit Range Interconnects,” DesignCon 2016.

[3] D. E. Bockelman and W. R. Eisenstadt, "Combined differential and common-mode scattering parameters: theory and simulation," in IEEE Transactions on Microwave Theory and Techniques, vol. 43, no. 7, pp. 1530-1539, July 1995.

[4] H. Dsilva, S. Moon, A. Zhang, C. Kao, and B. Rothermel, “Mathematically De-mystifying Skew,” DesignCon 2017.

[5] IEEE 802.3 clause 92.10.5

[6] M. R. Buford, P. A. Levin, and T. J. Kazmierski, “Temporal skew and mode conversion management in differential pairs to 15GHz,” IEEE Electronics Letters, Vol. 44, No. 1, 2008.

[7] E. Kunz, J. Y. Choi, V. Kunda, L. Kocubinski, Y. Li, J. Miller, G. Blando, and I. Novak, “Sources and Compensation of Skew in Single-Ended and Differential Interconnects,” DesignCon 2014.

[8] S. B. Smith, S. S. Agili, and V. Balasubramanian, “Theory and Measurement of Unbalanced Differential-mode Transmission Lines,” DesignCon 2006.

[9] IEEE 802.3 clause 137.10

[10] IEEE 802.3 clause 93A.4

[11] IEEE 802.3 clause 92.11.3.6