DesignCon 2019

Delivering 100 Gbps Solutions for Chip-to-Module and Direct Attach Copper (DAC) Cable Implementations

Bruce Champion, TE Connectivity
bruce.champion@te.com

Nathan Tracy, TE Connectivity
ntracy@te.com

Abstract

Efforts are underway to define 100 Gbps channels to support simpler, lower cost versions of 100G, 200G, and 400G Ethernet and also to enable 800G Ethernet links. This paper summarizes the design, build, and test of Chip-to-Module (C2M) and various Direct Attach Copper (DAC) configurations to give concrete evidence on what system designers can expect going forward with 100 Gbps channels. The paper includes channel impairment details that need to be addressed. It also presents methodologies required to support 100 Gbps next generation physical solutions for the channel, connectors, and cables.

Authors Biographies

Bruce Champion earned his undergraduate degree in Electric Engineering from Villanova University in 2007 and his Master of Science Degree from the Pennsylvania State University in 2016. For the past 10 years he was worked as a signal integrity engineer specializing in the analysis and design of high-speed components with a focus in I/O. Currently, he is a Staff Engineer at TE Connectivity where he focuses on high-frequency simulation, design, and test of multiple high-speed interconnect platforms.

Nathan Tracy is a Technologist on the system architecture team and manager of industry standards for TE Connectivity's Data and Devices business unit. Here he is responsible for driving standards activities as well as working with key customers to enable new system architectures. Nathan has more than 30 years of experience in technology development, marketing, and business development for TE in areas including RF/microwave and high speed signaling technology for the networking, telecom, wireless, automotive and defense markets. Nathan is active in numerous industry standards and associations. He currently serves the OIF as president and member of the board of directors, and is a regular attendee and contributor to the Ethernet Alliance, IEEE 802.3 and COBO. Additionally, he is active in a number of industry multi-source agreement (MSA) groups and forums where he has held leadership roles. Nathan earned his Bachelor of Science Electrical Engineering Technology degree from the University of Massachusetts, Dartmouth.

Introduction

In today's data and telecom networks, the current state of the art for electrical signaling in channels and interconnects is 28 Gbps NRZ for each differential pair. 56 Gbps electrical signaling definitions have recently been standardized and are starting to see limited deployment. There are numerous driving factors causing the industry to now begin the investigation of 112 Gbps signaling per differential pair. Throughout this paper, data rates are generically referred to as 28 Gbps, 56 Gbps, and 112 Gbps although the actual rates may differ depending on the specific industry group or standard considered. Figure 1 shows the historical and projected port shipment quantities by port speed from the year 2014 to 2025. The port data rate is based on all implementations of that speed. As an example, a 100G port could be implemented using 10 lanes of 10 Gbps, 4 lanes of 25 Gbps, 2 lanes of 50 Gbps, or 1 lane of 100 Gbps. The complexity and cost of a deployment always decreases as the number of lanes are reduced. As a result, next generation 800 Gbps port deployment will realize cost benefits if signaling at 112 Gbps per differential pair can be realized. This is due to higher aggregate data rates on a fewer number of lanes.

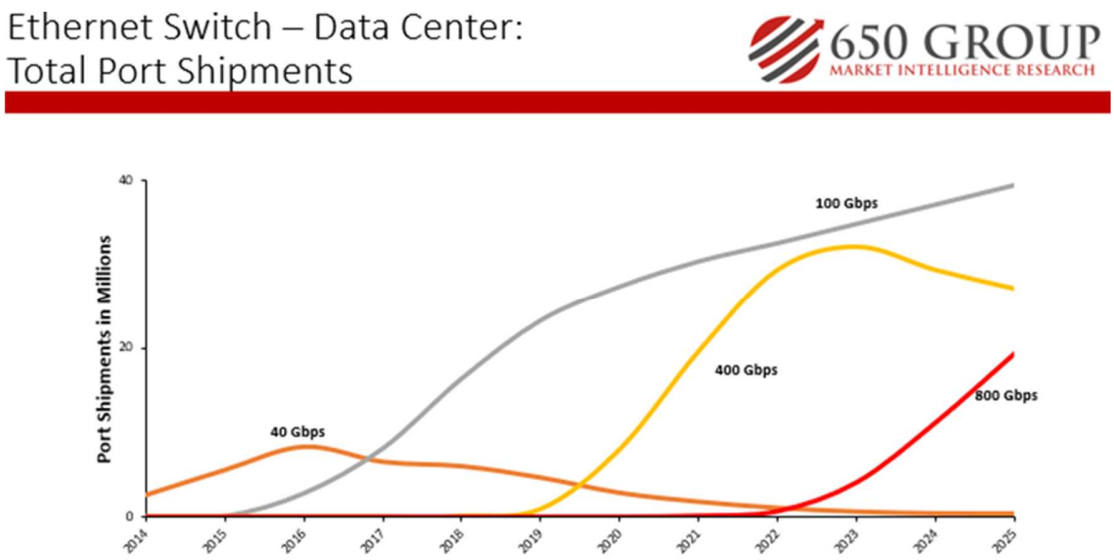

Figure 1. Projected Total Port Shipments (Chart used with the permission of 650 Group, LLC, Dec 2018)

For network switching silicon chips, the aggregate bandwidth of a switch is limited by the number of serializer/deserializer's (SerDes) that can fit on a chip and the data rate of each SerDes. By increasing the electrical data rate per lane on a switch, the aggregate switching bandwidth can be increased. Increased bandwidth enables the SerDes to keep up with data center and telecom network demand. Figure 2 shows the projected number of SerDes that are expected to be deployed at 112 Gbps.

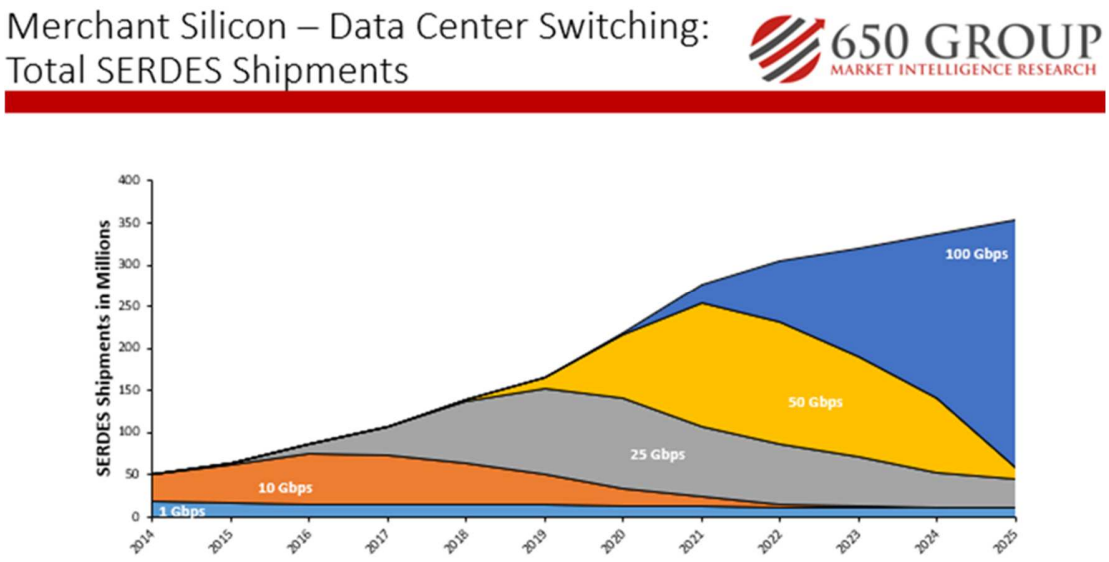

Figure 2. Projected SerDes Shipments (Chart used with the permission of 650 Group, LLC, Dec 2018)

The development of 112G electrical signaling is critical to the industry with the quantity of SerDes shipments projected. 112 Gbps electrical signaling, assuming the use of PAM-4 modulation, will require a 56 Gbaud symbol rate. At this data rate, insertion loss (IL), return loss (RL), and crosstalk (XT) are still vitally important; however, differential skew and mode conversion now have increased importance. The levels of these impairments will determine the quality and reach of the channel by reducing the signal that can be recovered by the receiver at the far end of the channel (and increasing the bit error rate).

The primary subjects of this investigation are input/output (I/O) channels. These can include either short channels used to support pluggable optical modules, or they can consist of a full electrical I/O link with Direct Attach Copper (DAC) cables in the channel. These channels typically consist of a host printed circuit board (PCB) inside a switch enclosure with switch/ SerDes silicon mounted to the PCB. Differential traces on the PCB connect the SerDes to a connector receptacle that receives the pluggable optical module or DAC cable. These pluggable interfaces are usually based on industry standard interfaces that enable higher volumes, lower cost, and multivendor interoperability. One potential side effect of using industry standard interfaces is that those interfaces have dimensions and tolerances that create a low impedance point. These tolerances drive certain nominal dimensions within the interface and will create channel impairments that have an increasing effect at higher data rates. This study will outline how to improve the existing 56 Gbps mating interface for higher data rates. Additionally, this study will explore the effect of skew on channel reach and overall operating margin for channels operating at 112 Gbps.

Skew Impact on Insertion Loss

Skew is an important factor needing close attention as the industry approaches 50+ Gbaud symbol rates. Skew can adversely affect differential insertion loss (S_{dd21}) and mode conversion (S_{cd21}). Any degradation in S_{dd21} or S_{cd21} will affect channel performance. As a result, it is important to know how skew affects these parameters, and in turn, how this affects operating margin.

It has been shown that for a channel with weak intra-pair coupling the S_{dd21} can be estimated by:

$$S_{dd21} = |IL| cos\left(\frac{\Delta\theta}{2}\right) e^{j\theta_{21}} \quad (1)$$

In (1), $\Delta\theta$ is the phase difference between the two through parameter phases of the differential pair, $|IL| = |S_{21}| \cong |S_{43}|$, and $e^{j\theta_{21}}$is the phase component of S_{21}. Figure 3 compares the estimated results of (1) with actual results simulated in Keysight[1] Advanced Design System (ADS). In this case, since there is zero skew, $\Delta\theta = 0$.

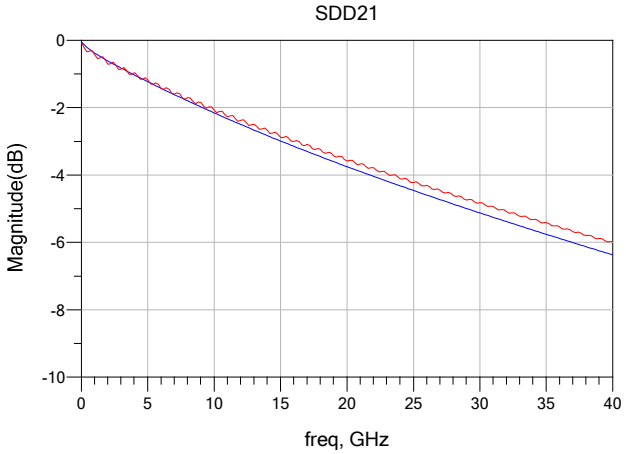

Figure 3. Estimated S_{dd21} (Red), Simulated (Blue)

This estimation does not provide an adequate result for this study. The difference in results is from the beginning assumption that the channel has weak intra-pair coupling. In looking at the traditional equation of S_{dd21} we cannot ignore the cross-coupling terms, $S_{41}$and S_{23}.

$$S_{dd21} = \frac{|S_{21}|e^{j\theta_{21}} + |S_{43}|e^{j\theta_{43}} - |S_{41}|e^{j\theta_{41}} - |S_{23}|e^{j\theta_{23}}}{2} \quad (2)$$

Incorporating the cross-coupling terms, the equation becomes

$$S_{dd21} = |IL| cos\left(\frac{\Delta\theta}{2}\right) e^{j\theta_{21}} - |X| cos\left(\frac{\Delta\theta_X}{2}\right) e^{j\theta_{41}} \quad (3)$$

In this case, $|X|=|S_{41}| \cong |S_{23}|$ and $\Delta\theta_X$ is the phase difference between the coupling terms S_{41} and S_{23}. Figure 4 shows a direct overlay when (3) is used to estimate S_{dd21} of the simulated trace. Again, since there is zero skew in the trace $\Delta\theta = \Delta\theta_X = 0$.

[1] Keysight is a trademark of Keysight Technologies Inc.

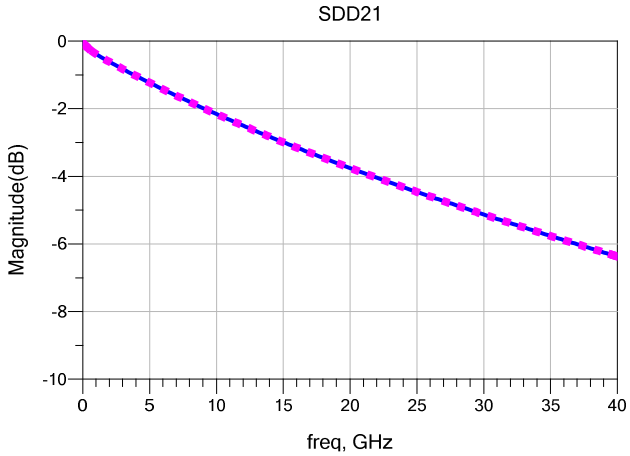

Figure 4. Estimated S_{dd21} (Pink Dash), Simulated (Blue)

S_{cd21} can be defined by:

$$S_{cd21} = \frac{|S_{21}|e^{j\theta_{21}} - |S_{43}|e^{j\theta_{43}} + |S_{41}|e^{j\theta_{41}} - |S_{23}|e^{j\theta_{23}}}{2} \tag{4}$$

Assuming $|IL| = |S_{21}| \cong |S_{43}|$ and $|X| = |S_{41}| \cong |S_{23}|$ (4) can be written as

$$S_{cd21} = \left|\frac{IL}{2}\right|\left(e^{j\theta_{21}} - e^{j\theta_{43}}\right) + \left|\frac{X}{2}\right|\left(e^{j\theta_{41}} - e^{j\theta_{23}}\right) \tag{5}$$

Using $\theta_{21} = \theta + \frac{\Delta\theta}{2}$, $\theta_{43} = \theta - \frac{\Delta\theta}{2}$, $\theta_{41} = \theta + \frac{\Delta\theta_X}{2}$, and $\theta_{23} = \theta - \frac{\Delta\theta_X}{2}$ (5) can become

$$S_{cd21} = |IL|\left[-jsin\left(\frac{\Delta\theta}{2}\right)\right]e^{j\theta_{21}} + |X|\left[-jsin\left(\frac{\Delta\theta_X}{2}\right)\right]e^{j\theta_{41}} \tag{6}$$

The main difference between (3) and (6) is the removal of the cosine term and the introduction of the sine term. We will refer to these terms as the S_{dd21} and S_{cd21} modifier values.

$$S_{dd21Mod} = cos\left(\frac{\Delta\theta}{2}\right) \tag{7}$$

$$S_{cd21Mod} = -jsin\left(\frac{\Delta\theta}{2}\right) \tag{8}$$

(7) and (8) are plotted in Figures 5 and 6 respectively for skew values of 3, 6, and 9 ps. Figure 7 and 8 show the estimated $S_{dd21}$and S_{cd21} of a differential stripline trace with 0, 3, 6, and 9 ps of skew. It should be noted that in Figure 8, the zero ps mode conversion plot is equal to zero. This is due to $\Delta\theta = 0$.

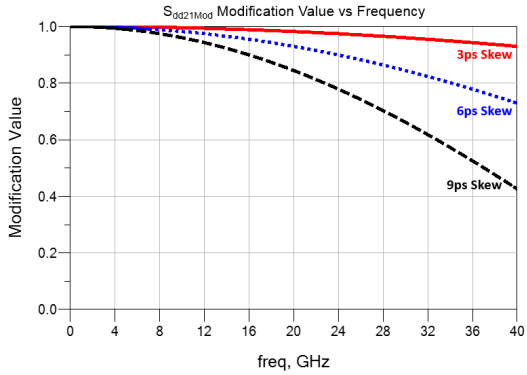

Figure 5. $S_{dd21Mod}$
3ps Skew (Red), 6ps (Blue), 9ps (Black)

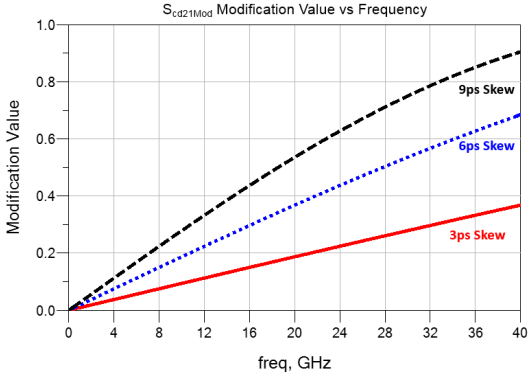

Figure 6. $S_{cd21Mod}$
3ps Skew (Red), 6ps (Blue), 9ps (Black)

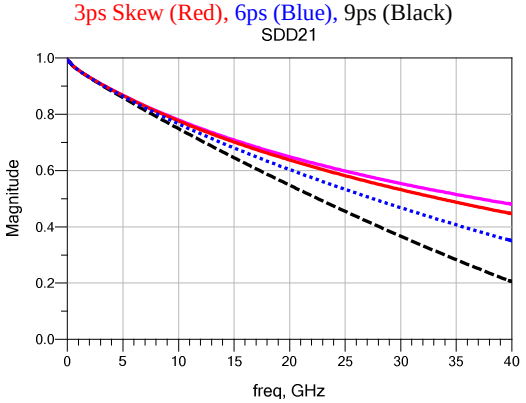

Figure 7. Stripline trace S_{dd21}
0ps Skew (Pink), 3ps (Red), 6ps (Blue), 9ps (Black)

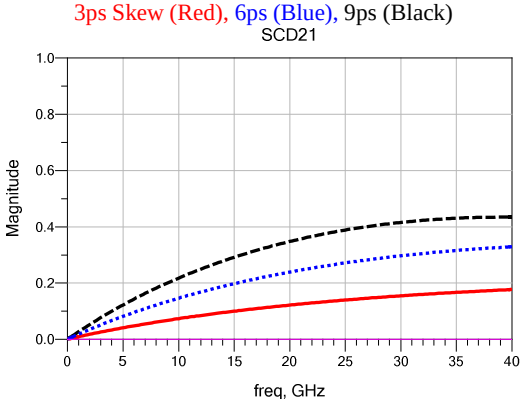

Figure 8. Stripline trace S_{cd21}
0ps Skew (Pink), 3ps (Red), 6ps (Blue), 9ps (Black)

S_{dd21} and S_{cd21} are related to one another in a lossless case such that the complex addition of the two terms follow:

$$S_{dd21} + S_{cd21} = 1 \tag{9}$$

This is seen for a lossless stripline trace. Figure 9 shows a lossless stripline trace with 0 ps of skew. The magnitude is 1 for S_{dd21} and 0 for S_{cd21} (since $\Delta\theta = 0$). As a result, (9) holds true. Figure 10, shows a lossless stripline trace with 9 ps of skew. The change in S_{dd21} and S_{cd21} follow (7) and (8) respectively. This change is attributed to added skew since $\Delta\theta \neq 0$. (9) still holds true.

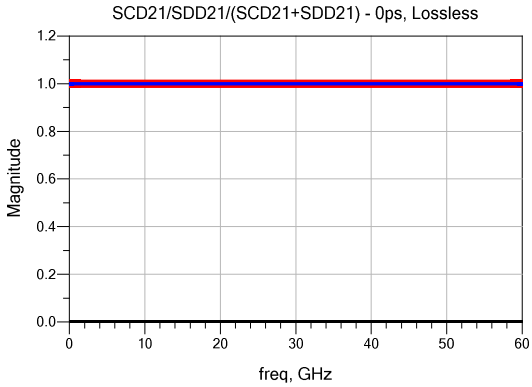

Figure 9. S_{dd21}(Red), S_{cd21}(Black), S_{dd21}+S_{cd21}(Blue)

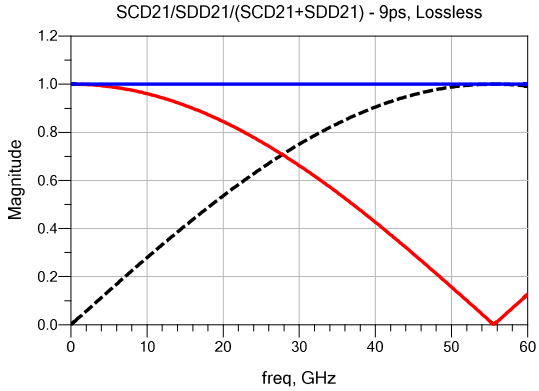

Figure 10. S_{dd21}(Red), S_{cd21}(Black), S_{dd21}+S_{cd21}(Blue)

When dielectric and conductive losses are introduced (9) no longer holds true as energy is now dissipated. Figure 11 shows a lossy stripline trace with 0 ps of skew. The magnitude is no longer 1 for S_{dd21} as a result of the dielectric and conductive losses. S_{cd21} still equals 0 (since $\Delta\theta = 0$). Figure 12 shows a lossy stripline trace with 9 ps of skew. The change in S_{dd21} and S_{cd21} is now attributed to (7) and (8) in addition to the dielectric and conductive losses.

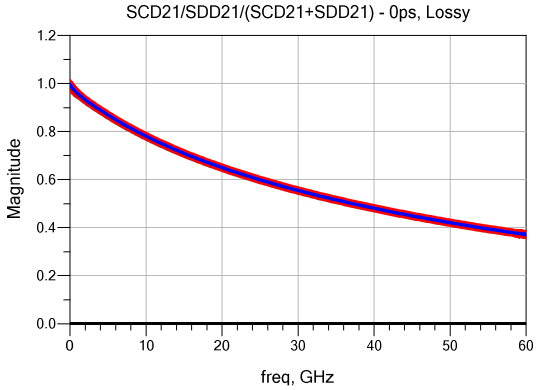

Figure 11 S_{dd21}(Red), S_{cd21}(Black), S_{dd21}+S_{cd21}(Blue)

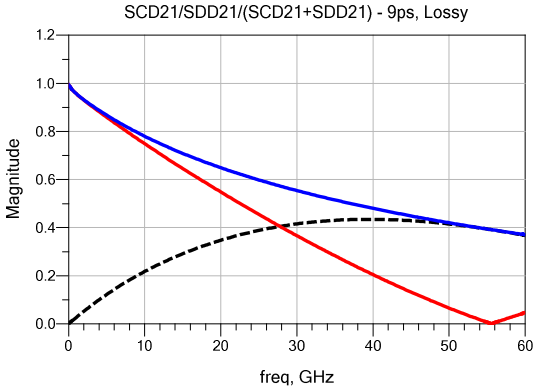

Figure 12. S_{dd21}(Red), S_{cd21}(Black), S_{dd21}+S_{cd21}(Blue)

As frequency increases, the impact of skew on a differential pair is more significant. For example, looking at the 9 ps skew case and at the Nyquist frequencies for 28G NRZ, 56G PAM-4, and 112G PAM4, $S_{dd21Mod}$ is not constant across those frequencies (14 and 28 GHz). Skew will have a larger impact on S_{dd21} in 112G channels as compared to 28 or 56G channels. Additionally, the impact on S_{dd21} from different skew values is not linear. For example, $S_{dd21Mod}$ at 26.56 GHz (Nyquist for 100G IEEE 802.3ck specification) for the 3, 6, and 9 ps skew values are 0.969, 0.877, and 0.730 respectively. Going from 3 to 6ps of skew yields a 9.5% change in $S_{dd21Mod}$. From 6 to 9ps of skew yields a 16.8% change in the $S_{dd21Mod}$. In looking at the Nyquist frequencies from a S_{cd21} standpoint, $S_{cd21Mod}$ isn't constant either. It can also be observed in Figure 8 that S_{cd21} is more sensitive to skew in the 0 to 30 GHz range. Therefore, if designers are looking to troubleshoot or monitor their design for skew, S_{cd21} should be a parameter taken into consideration as a low S_{cd21} value would mean differential insertion loss is maximized. To maximize reach of a channel, skew needs to be minimized such that differential energy is not converted into common energy because it manifests itself as additional differential insertion loss.

Improving Reflections for 112G Channels

One of the most problematic reflection sources in a C2M or DAC channel is the I/O connector-to-module mating interface where the transceiver or passive copper cable assembly meets the I/O connector inside the port. From an electrical standpoint, one of the most important interface dimensions is the mating pad width, which varies between form factors and is also a major source of unwanted reflections. Three common mating interfaces in multi-gigabit data transmission systems are from the QSFP28, QSFP-DD, and OSFP form factors. These form factors are driven by the SFF-8662, QSFP-DD MSA, and OSFP MSA respectively. QSFP28 was developed for 28 Gbps operation while QSFP-DD and OSFP were developed for 56 Gbps operation. The widths of the pads and pad pitch are shown in Figure 13.

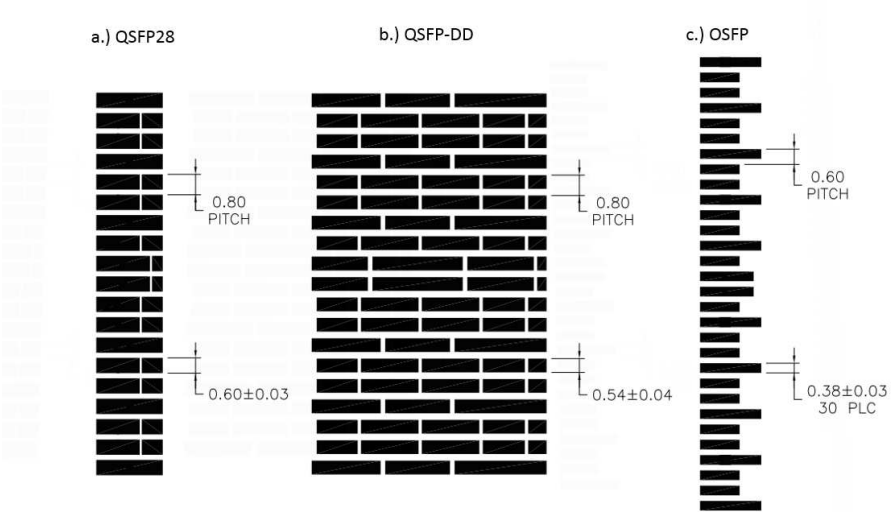

Figure 13. QSFP28, QSFP-DD, OSFP Mating Interfaces

QSFP28 has the highest pad width-to-pitch ratio at 0.75 while OSFP has the smallest pad width-to-pitch ratio at 0.63. These layouts inherently have their own characteristic impedance without considering the receptacle connector; that is, these pads act as microstrip traces in the PCB mating interface. These pads form the basic starting point for the mating interface impedance. Little, if anything, can be done in the receptacle connector to improve this impedance. Each of these configurations was set-up in a 2D field solver to calculate characteristic impedance. Geometries are shown if Figure 14 a-c.

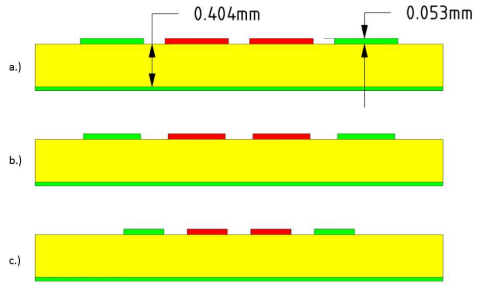

Figure 14. a.) QSFP28 b.) QSFP-DD c.) OSFP

The spacing between ground plane and the surface pads is fixed at .404mm. This distance allows the ground plane to be roughly in the center of a 1.0 mm thick PCB that includes pads on both the top and bottom surfaces. It is important to maximize this spacing to achieve optimal impedance while still allowing the ground plane to provide top-to-bottom noise isolation. Additionally, the top copper thickness is assumed to be 0.053mm which is a typical ½ oz copper with a 1mil plate-up. The dielectric constant chosen for this study is 3.8. The results of the 2D analysis are shown in Table 1.

Table 1. Characteristic Impedance of Mating Pads without Connector

Form Factor	Impedance (Ohms)
QSFP28	82
QSFP-DD	92
OSFP	101

QSFP28 has the worst impedance in the unmated PCB pad area while OSFP has the best impedance. The 10% reduction in pad size between QSFP-DD and QSFP28 yields a 10 Ohm

improvement in the mating interface area. As mentioned earlier these numbers do not consider the contribution of the receptacle connector which can have a large contribution to the mating interface impedance.

The largest impact on mating interface impedance in a receptacle connector comes from two geometries: the contact stub from the contact lead-in and the PCB pad stub from the amount of contact wipe. Historically these features have ensured reliable mechanical connection. A typical contact lead-in and contact wipe are shown in Figure 15.

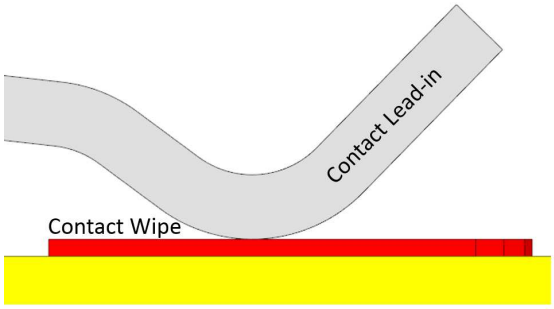

Figure 15. Typical Contact Lead-in and Contact Wipe

These two factors create a large stub in the mating interface, adding extra capacitance, degrading impedance, and increasing reflections. Each of the three form factors were measured on a Time Domain Reflectometer (TDR) to observe the difference between the results in Table 1 to the measured impedance of the mated connector. The TDR results of the mated connectors are shown in Figure 16.

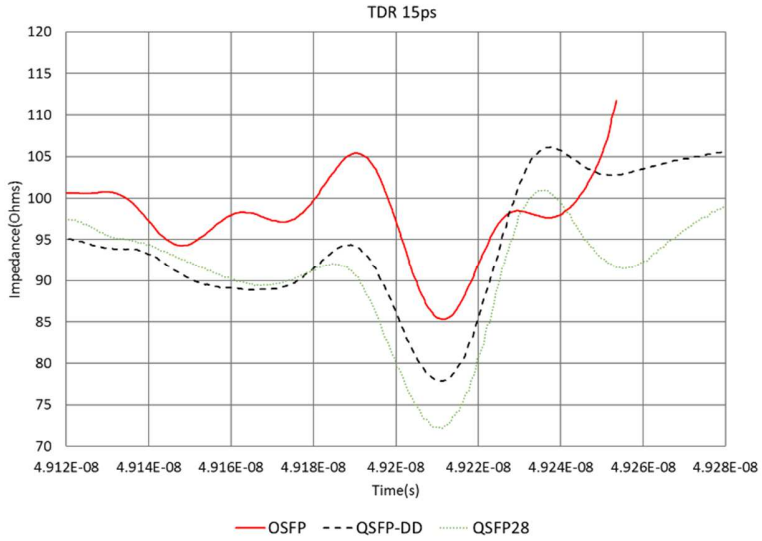

Figure 16. TDR results

Comparing Table 1 to Figure 16, the impedance in the mating zone is lowered by approximately 15 Ohms when the receptacle connector is introduced in the OSFP and QSFP-DD cases. The mating zone of QSFP28 was lowered by 10 Ohms.

To determine the benefits of improving the current OSFP mating zone, a 56G OSFP connector was modified. The modification involved reducing the size of the contact lead-in. This change led to an improvement of approximately 8 Ohms as shown in Figure 17. **For simplification, this**

modified 56G connector will be referred to as a 112G connector moving forward. This 112G connector was used to analyze a 112G C2M and DAC channel.

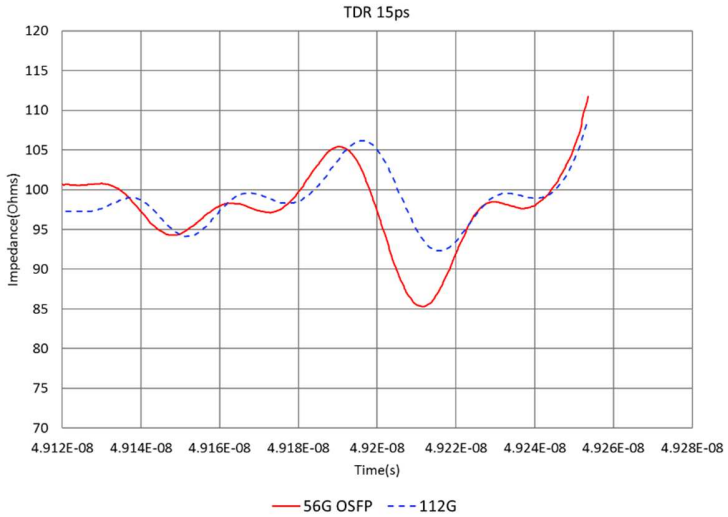

Figure 17. TDR results improved mating zone

Impact of Skew and Reflections for 112G DAC Channels

In the design of a DAC cable, both the mating interface impedance and skew are large contributors to performance. To ascertain their impact, 1 m/30 AWG and 1.5 m/28 AWG cable assemblies were manufactured and used as test vehicles. To show the effect of the mating zone on DAC performance these assemblies were measured using both an existing 56G connector and a 112G connector. To show the effect of skew, measurement data was brought into Keysight ADS and skew was injected as a phase delay. In each case Channel Operating Margin (COM) was analyzed to view the effect on channel performance. Figure 18 shows a depiction of the test set-up which includes approximately -2.55 dB of attenuation at 26.56 GHz on each side of the cable assembly from the test PCB. 26.56 GHz is the Nyquist frequency for the 100G IEEE 802.3ck specification.

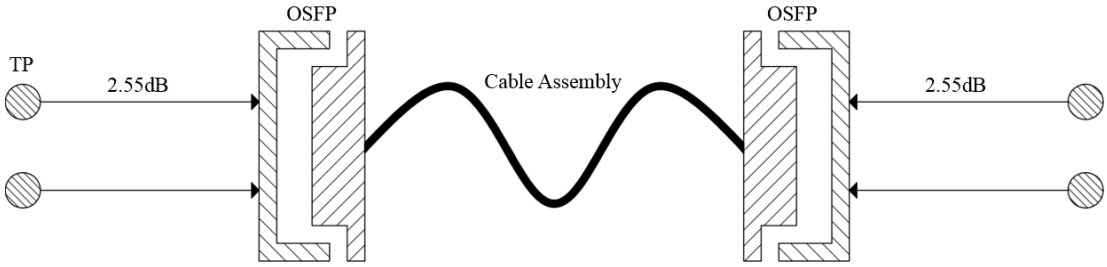

Figure 18. DAC Test Set-up

XT is also included in the analysis. For each of the two cable assemblies (1 m/30AWG, 1.5 m 28AWG), the Tx7 and Tx8 lanes were analyzed. In each case, 15 noise aggressors were used as shown in Figure 19.

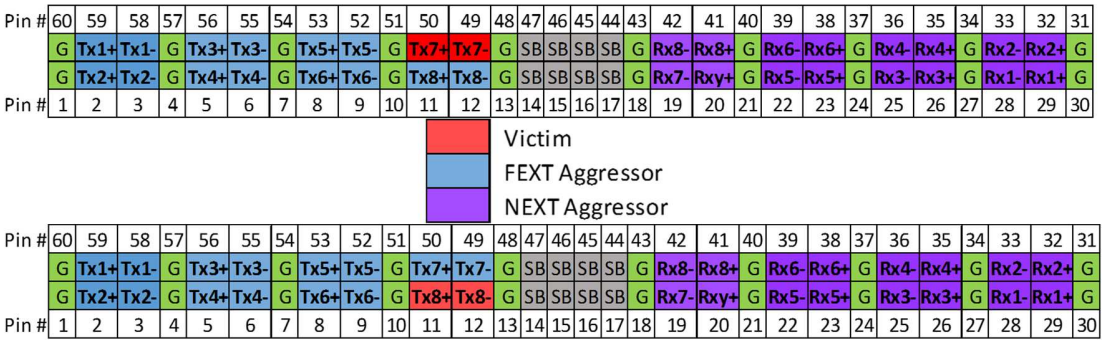

Figure 19. Crosstalk Aggressors for Tx7 and Tx8 Lanes

The S_{dd21} and XT results of the Tx7 and Tx8 lanes are shown for both cable assemblies (28 and 30AWG) each using the 56G connector and 112G connector in Figure 20.

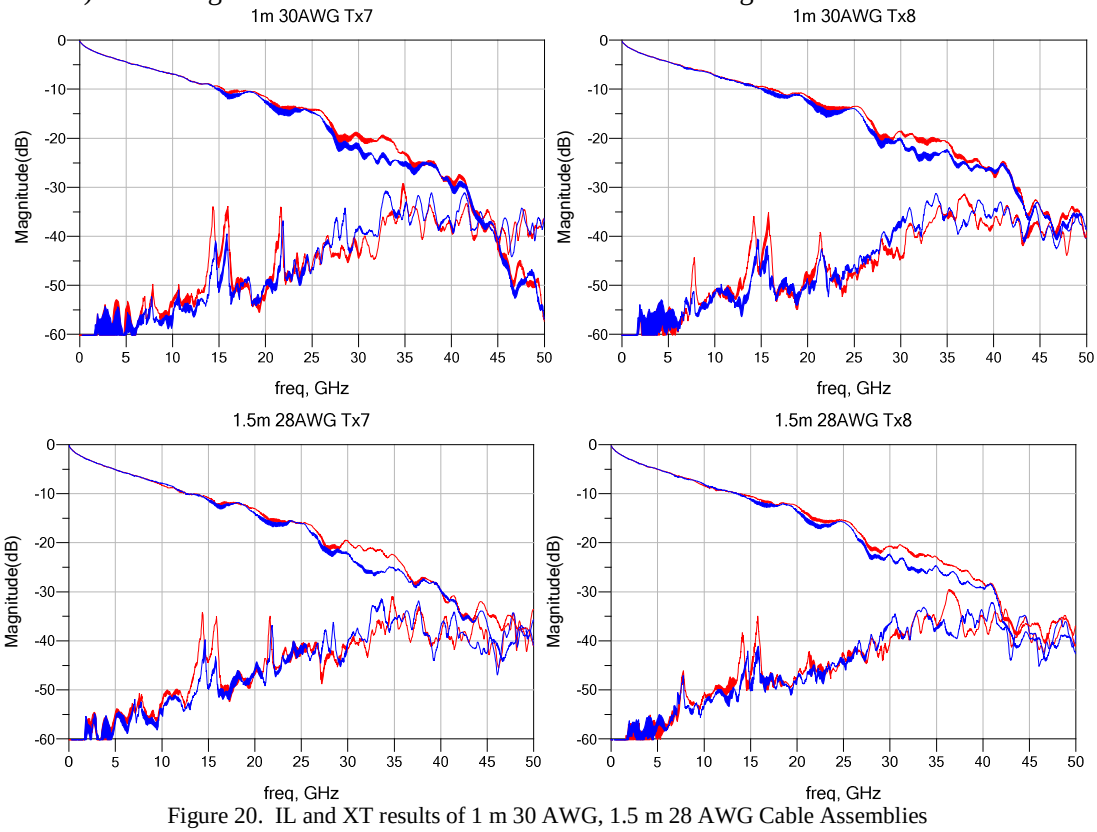

Figure 20. IL and XT results of 1 m 30 AWG, 1.5 m 28 AWG Cable Assemblies
56G Connector (Blue) 112G Connector (Red)

Skew for these channels is shown in Figure 21. The Tx7 lanes for the 28AWG and 30AWG with the 56G and 112G connector is shown in left plot and the same configurations are shown for the Tx8 lanes in the rightmost plot. It should be noted the skew from the test PCBs are also included in this result.

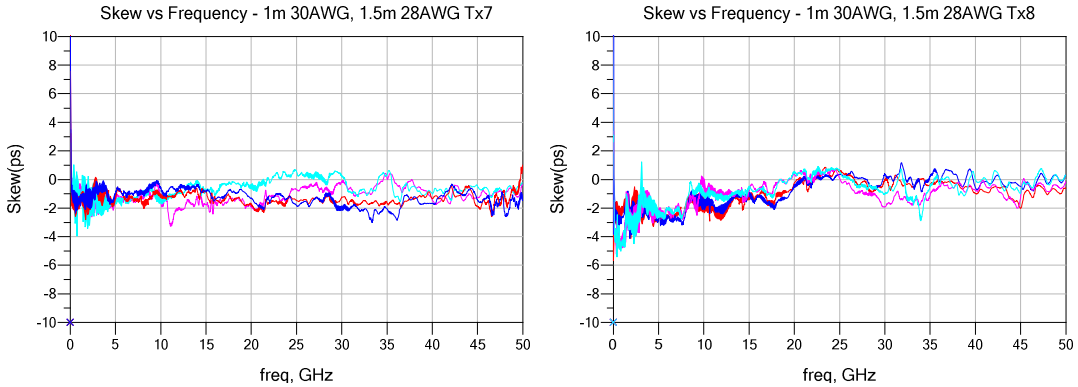

Figure 21. Skew vs Frequency, Tx7 Lanes (Left), Tx8 Lanes (Right)
30AWG/112G Conn(Red), 30AWG/56G Conn(Blue), 28AWG/112G Conn(Pink), 28AWG/56G Conn(Light Blue)

As mentioned earlier, skew was added to the measurement to observe the effect on S_{dd21}, S_{cd21}, and COM. To do this a phase delay was added to each of these lanes representing 3, 6, and 9ps of skew. The change in S_{dd21} and S_{cd21}because of the added skew for the Tx8 lane for both the 1.0m/30AWG and 1.5m/28AWG using the 112G connector is shown in Figure 22.

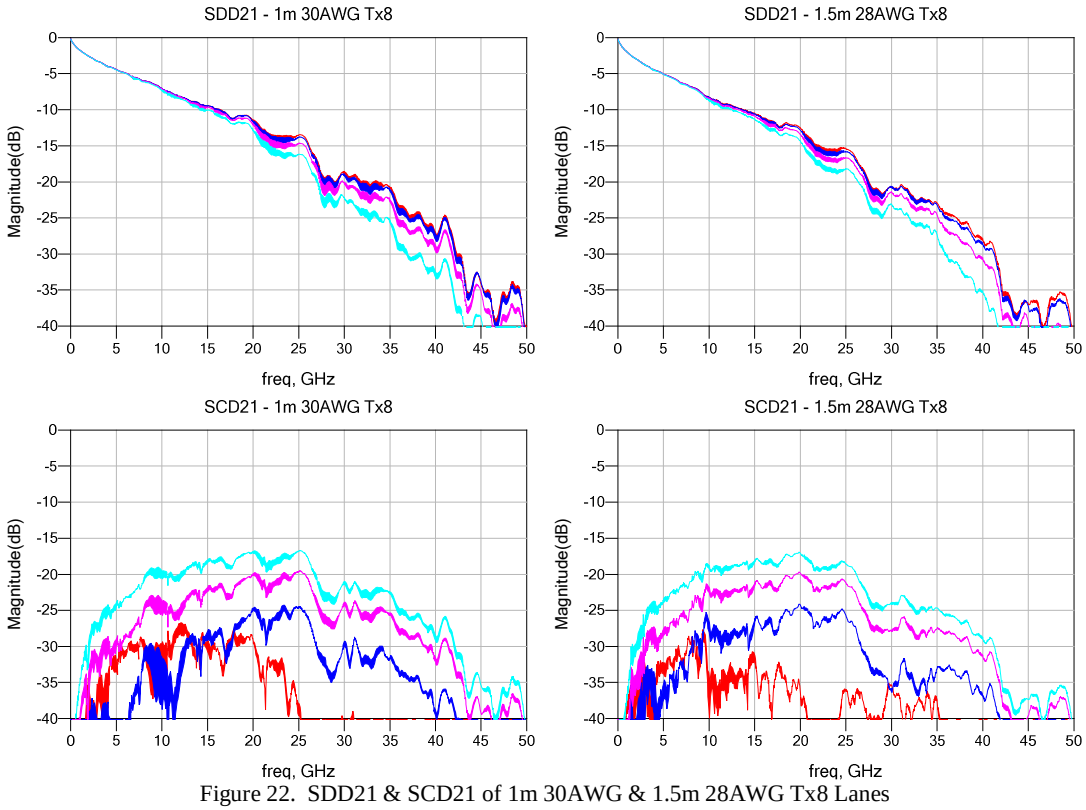

Figure 22. SDD21 & SCD21 of 1m 30AWG & 1.5m 28AWG Tx8 Lanes
0ps Added Skew (Red), 3ps Added Skew (Blue), 6ps (Pink), 9ps (Light Blue)

The changes in S_{dd21} and S_{cd21} match what was predicted earlier, that is, skew increased S_{dd21} and this degradation is not equal for all frequencies. Additionally, the increase in S_{dd21} is not proportional between 0 to 3ps, 3 to 6ps, or 6 to 9ps and the change in S_{cd21} is more dramatic than the change in S_{dd21}. Figure 23 shows S_{dd21}, S_{cd21}, and S_{dd21} + S_{cd21} for the Tx8 lane of the 30AWG assembly using the 112G connector. S_{cd21} is very low in magnitude and never surpasses S_{dd21}. S_{dd21} is nearly equivalent to S_{dd21} + S_{cd21}. Figure 24 shows S_{dd21}, S_{cd21}, and S_{dd21} + S_{cd21} for the same lane

with 9ps of injected skew. S_{dd21} + S_{cd21} between Figures 23 and 24 are equivalent; however, differential energy was converted to common energy causing $|S_{dd21}|$ to decrease in magnitude and $|S_{cd21}|$ to increase in magnitude. $|S_{cd21}|$ does intersect with $|S_{dd21}|$ in the same frequency range as predicted using (7) and (8) and is also higher in magnitude as predicted at the higher frequencies.

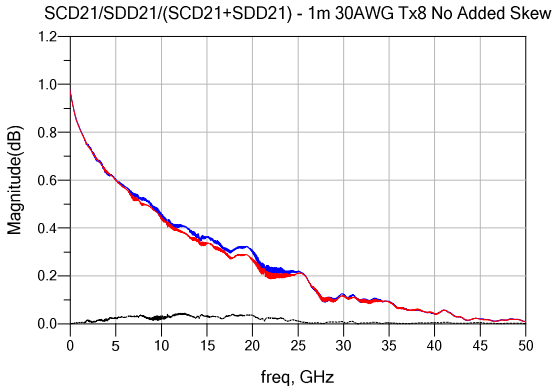

Figure 23 $|S_{dd21}|$(Red), $|S_{cd21}|$(Black), $|S_{dd21}|+|S_{cd21}|$(Blue)

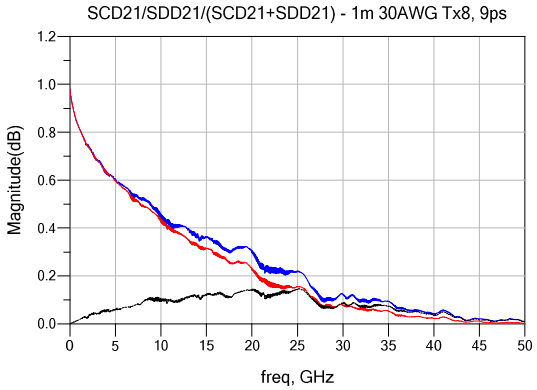

Figure 24 $|S_{dd21}|$(Red), $|S_{cd21}|$(Black), $|S_{dd21}|+|S_{cd21}|$(Blue)

Each of these channels were evaluated using the Version 2.51 COM script developed for the IEEE 802.3ck specification and is readily available to the public. The configuration settings used to calculate COM are also readily available as part of the 2.51 COM contribution for the IEEE 802.3ck working group. The results for COM Case 1 and 2 are shown in Figure 25. COM Case 1 corresponds to a shorter package length when compared to COM Case 2. The COM results include values for both the Tx7 and Tx8 lanes of the 28 and 30AWG assemblies. In addition, they include lanes with 3, 6, and 9ps of skew along with each of these skew values for the 56G and 112G connectors. In total, 32 different lanes were analyzed.

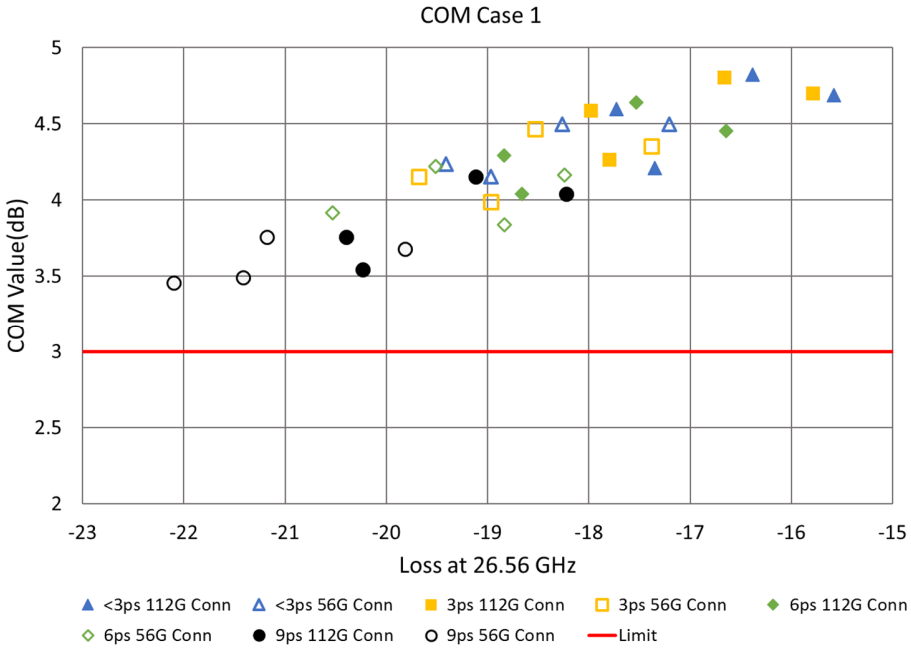

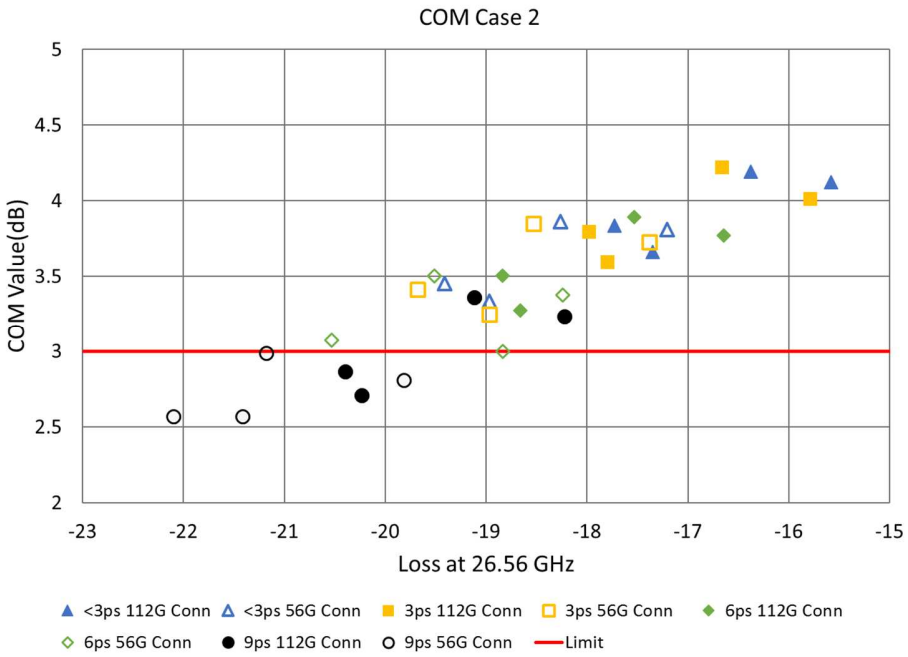

Figures 25. COM Case 1 & COM Case 2

All instances pass COM Case 1. This is the case for a shorter package length. There are a few instances which fail COM Case 2 utilizing a 30 mm package length. All instances of the 56G connector with 9 ps of skew fail this case. Only two instances of the 112G connector with 9 ps of skew fail COM Case 2. Although it cannot be determined from the plot, these instances are from the 1.5 m/28 AWG cable assembly which started out with a higher S_{dd21} than the 1 m/30 AWG assembly. Even with 9 ps of skew the 1 m/30 AWG assembly passes COM. The other cases close to the limit line for COM Case 2 are the lanes with 6 ps of skew using the 56G connector. There were no instances of the 112G connector and 6ps of skew failing COM. Therefore, in designing next generation 112G channels designers should look at improving the current mating zone of the 56G connectors in addition to controlling skew within the channel. Doing so will allow designers to achieve the maximum possible reach.

C2M Channel

To observe the effect of skew and an improved mating interface in a C2M channel, a set-up was designed, built, and tested as shown in Figure 26.

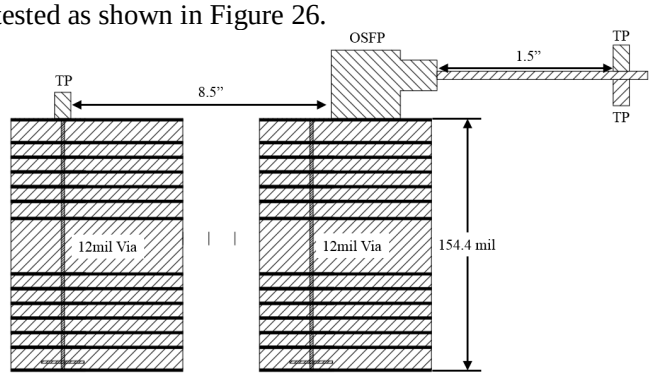

Figure 26. C2M Channel

The C2M uses MEGTRON[2] 7N HVLP material in both the host and module PCBs. The host uses 8.5" of 5.3-6.45-5.3 differential traces while the module uses 1.5" of 6.3mil single ended (SE) trace. An important note about this set-up is the use of 12mil vias for both the connector breakout and the test point. These vias are on the larger size for typical designers and provide a worst case for the breakout in regard to impedance and XT. Additionally, as in the DAC channel, the 56G and 112G connectors were analyzed. The 56G and 112G connector. For this set-up the Tx6 lane was analyzed using the 5 closest far-end crosstalk aggressors. The pinout is shown in Figure 27.

Pin #	60	59	58	57	56	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31
	G	Tx1+	Tx1-	G	Tx3+	Tx3-	G	Tx5+	Tx5-	G	Tx7+	Tx7-	G	SB	SB	SB	SB	G	Rx8-	Rx8+	G	Rx6-	Rx6+	G	Rx4-	Rx4+	G	Rx2-	Rx2+	G
	G	Tx2+	Tx2-	G	Tx4+	Tx4-	G	Tx6+	Tx6-	G	Tx8+	Tx8-	G	SB	SB	SB	SB	G	Rx7-	Rxy+	G	Rx5-	Rx5+	G	Rx3-	Rx3+	G	Rx1-	Rx1+	G
Pin #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30

Victim
FEXT Aggressor

Figure 27. Crosstalk Aggressors for Tx6 Lane

The S_{dd21} and crosstalk results of the Tx6 lane for the 56G and 112G connectors are shown in Figure 28. The skew for these lanes is shown Figure 29.

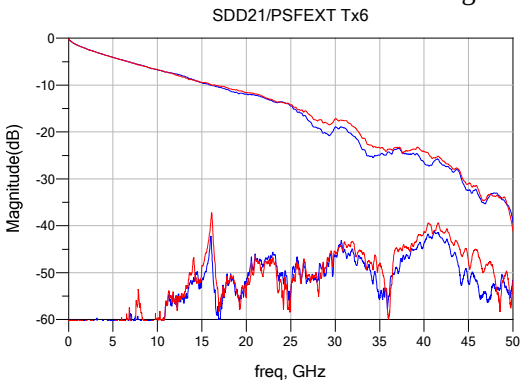

Figure 28. Sdd21 & XT Tx6
56G Connector (Blue), 112G Connector (Red)

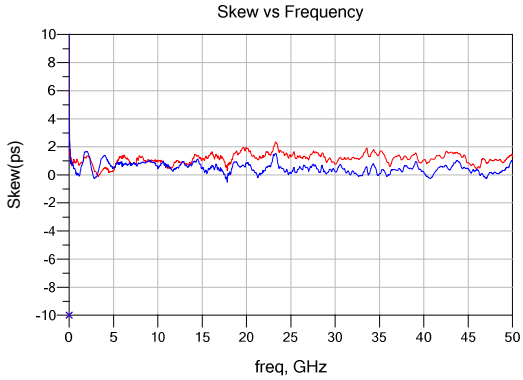

Figure 29. Skew Tx6
56G Connector (Blue), 112G Connector (Red)

The effect of skew on COM is also analyzed. To do this a phase delay was added to each of the channels representing 3, 6, and 9ps of skew. The change in S_{dd21} and S_{cd21} are result of the added skew for the Tx6 lane. Both the 56G and 112G connector are shown in Figure 30.

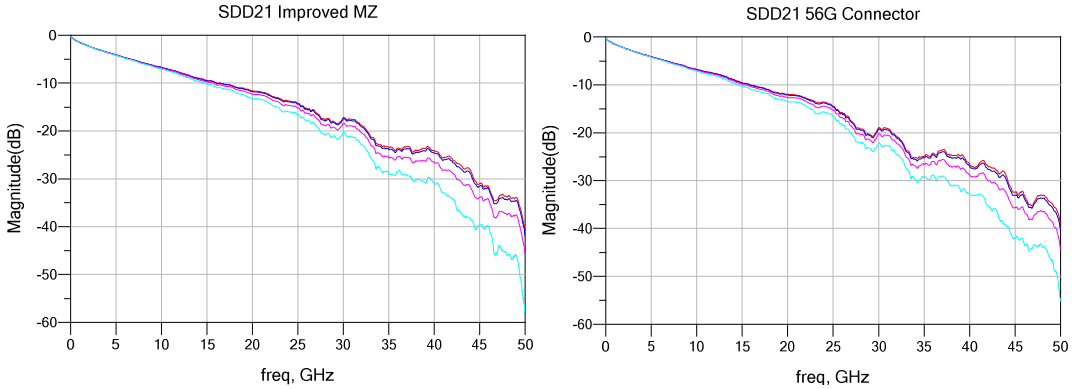

[2] MEGTRON is a trademark of PANASONIC Corporation

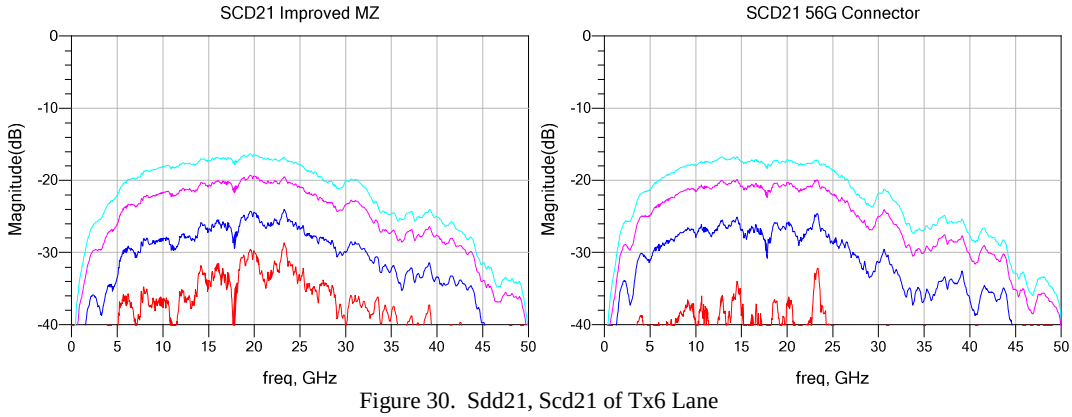

Figure 30. Sdd21, Scd21 of Tx6 Lane
0ps Added Skew (Red), 3ps Added Skew (Blue), 6ps (Pink), 9ps (Light Blue)

These 8 different channels were evaluated using the Version 2.41 COM script developed for the IEEE 802.3ck specification and is readily available to the public. This script is different than the 2.51 COM script used to analyze the DAC channel. The configuration settings used to calculate COM are also readily available as part of the 2.41 COM contribution for IEEE 802.3ck working group. One difference between what was posted on the public website to these settings is the addition of a DFE tap. The results for COM Case 1 and 2 are shown in Figure 31.

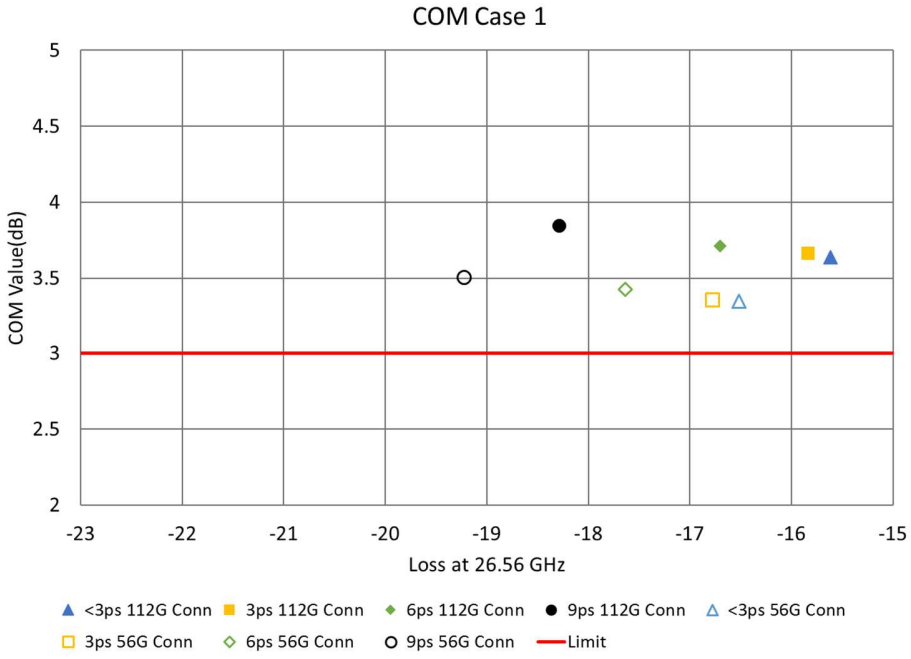

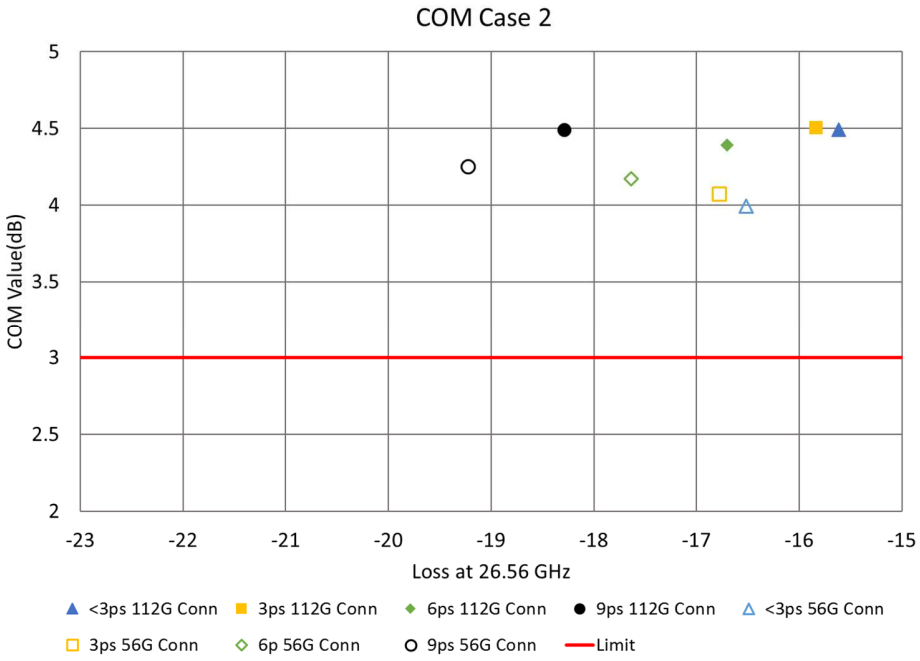

Figure 31. COM Case 1 & COM Case 2

All instances pass both COM cases. Unlike the DAC channel, the C2M channel is not as sensitive to S_{dd21}. In fact, the COM numbers improved with increased S_{dd21}. Skew does not appear to have as much of an impact on COM in the C2M channel as compared to a DAC channel. The 112G connector averaged a 3.72 dB COM value for COM Case 1 while the 56G connector averaged 3.41 dB.

Conclusion

Maximizing reach for 112G channels will be difficult. Two main drivers in extending reach are skew and mating zone reflections. A relationship between skew, differential insertion loss, and mode conversion was determined. By minimizing skew, less differential energy is converted to common mode energy, and differential throughput can be maximized. It was shown that for any degree of skew one can relate the change in differential insertion loss to the change in mode conversion.

Both a DAC and C2M channel were designed, built, and tested up to the frequencies of interest for 112G channels. Good margin was observed from a COM perspective with these channels. Channel impairments were also added to the channels themselves to observe performance impact. It was only when skew was added to the DAC channel that failures started to occur because of the increased differential insertion loss. If mating zone reflections can be minimized, a channel can be more tolerant of skew from a differential insertion loss standpoint. Minimizing both skew and mating zone reflections are strong drivers in being able to maximize channel reach for 112G.

References

[1] S. Farrahi, V. Kunda, Y. Li, X. Zhang, G. Blando and I. Novak, "Does Skew Really Degrade SERDES Peformance," in *DesignCon*, Santa Clara, 2015.

[2] R. Mellitz, "COM 2.41 with 100GEL Update," 2018. [Online]. Available: http://www.ieee802.org/3/ck/public/adhoc/aug15_18/mellitz_3ck_adhoc_01_081518.pdf.

[3] R. Mellitz, "COM 2.51 with rxFFE updates," 2018. [Online]. Available: http://www.ieee802.org/3/ck/public/adhoc/oct03_18/mellitz_3ck_adhoc_01_100318.pdf.

i