DesignCon 2018

Designing DC-Blocking Capacitor Transitions to Enable 56Gbps NRZ & 112Gbps PAM4

Scotty Neally, Samtec Teraspeed Consulting
Email: Scotty@Teraspeed.com

Scott McMorrow, Samtec Teraspeed Consulting
Email: Scott@Teraspeed.com

Abstract

DC blocking capacitors are required in almost all applications of high-speed SERDES design in order to level shift the differential signal to an operating point for optimum receiver performance and avoid DC ground loops. As we start the transition from 28G to 56G NRZ and 112G PAM4, it is crucial that the DC blocking capacitor present a high bandwidth, near reflectionless transition to the signal and maintain at least 32GHz of effective interconnect bandwidth. In this paper we discuss the steps necessary to generate a realistically detailed capacitor model for data transmission and optimize its layout for an electrically transparent design.

The genesis of this work is based on MLCC cross-sectional data provided through the use of scanning electron microscopy (SEM), modeling and simulation data derived from Ansys HFSS simulation software, and measurement data to be obtained with Agilent Vector Network Analyzers using advanced calibration methods.

In this paper we also present a more approachable first-plate capacitor model to be used in optimizing DC blocking capacitor transition design using 3D full wave solvers. We will explain in depth how to build a complex model of a multi-layered ceramic capacitor, optimize its transition region and build a correlation vehicle for testing. We will also exhibit the performance differences between the various capacitor models in the frequency and time-domain.

Authors Biography

Scotty Neally is an experienced signal integrity consultant with a background in high density PCB design and measurement, automation of signal integrity design flows, and currently focuses on system design for emerging technologies.

Scott McMorrow is an expert in high-performance design and signal integrity engineering, with a broad background in complex system design, interconnect modeling and measurement methodology, and professional training spanning over 25 years.

Introduction

Many electrical standards and interoperability agreements such as IEEE Std 802.3™ 2015, OIF-CEI 28G and PCIe Base/CEM require the link to meet minimum eye mask specifications at the receiver after applying equalization. While these standards often use time domain metrics to measure channel compliance, interconnect designs can be optimized more quickly and easily in frequency domain. To facilitate the design process of next generation high speed serial links, we must rely on frequency domain compliance curves as guidelines to properly weigh design decisions. While it is very difficult to predict link margin when considering the frequency domain alone, it is a widely accepted practice to optimize each interconnect region separately in the frequency domain to improve the overall system performance.

Design Relevance

DC blocking capacitors are essential to a variety of high speed electrical interfaces such as OIF-CEI 28G VSR, SR, MR, and LR channels. As the next generation of designs target data rates of 56G and above, it becomes increasingly important to characterize channel transitions accurately to ensure a high confidence of success. As such, some designers overlook the need for full wave modeling in component break out regions because of the difficulties they present to both procuring models and optimization in a standard flow. This becomes especially difficult for DC blocking capacitor simulation as component parasitics are hard to obtain and transition regions difficult to model accurately in planar solvers. For this reason, many chip and FPGA manufacturers give general guidelines for PCB layout and transition geometries to make design with their transceivers much simpler [1-2]. While this can expedite the design process, we feel that it is best to optimize design structures in-situ whenever possible to take into account 3D EM fields which planar solvers cannot.

In this study, we start with a shorted pad model to optimize a 0201 DC blocking capacitor transition in HFSS because of its computation speed. We then compare it to a first-plate capacitor model (referred to in plots as the “simple” model) and a full 3D MLCC model (referred to in plots as the “complex” model). Here we introduce the idea that a simplified first-plate capacitor model can be used to quickly optimize a DC blocking capacitor transition using full wave solvers and seeks to minimize computation time while maximizing performance similarities to a fully developed MLCC model. We would like to make the distinction that while we did not use this model as an upfront approach in designing our reference capacitor transition, it emerged as a byproduct to more quickly compare simulations of a full MLCC model and a shorted pad model.

The most common locations for placing DC blocking capacitors are near the connector transition regions where impedance discontinuities normally occur [3]. With this in mind, we have selected a differential microstrip configuration to model an outer layer breakout region which allows us to design a more ideal de-embedding fixture with minimal transitions. As a secondary finding of this paper, we will show how the performance of a given capacitor transition is affected by the substitution of other capacitor values in the same package family.

DC Blocking Transition Optimization

While this paper makes no endorsements of existing methods of transition optimization, our selected approach seeks to build, optimize, and characterize the proposed layout for relatively easy measurement correlation. We advise also consulting other bodies of work on this topic before starting a transition optimization so you may select the method that best fits your design needs [1-4]. Keeping in mind the goal of being able to isolate the effect of the capacitor transition alone, we started with

differential microstrip so we may be able to use de-embedding techniques to move the reference plane as close the DUT as possible (1mm). To begin, we created a differential trace geometry in an ECAD environment using a 0201 footprint based on dimensions given by AVX [5], a pre-determined stackup using I-TERA material, and a 15mil differential trace pitch based on the weave. We imported this geometry into HFSS, shorted the mounting pads together, and created two void openings on the nearest ground plane to encapsulate the shadow of each pair of shorted mounting pads. Finally, we parameterized these void openings in 1mil increments in both X and Y directions. See Figure 1 below.

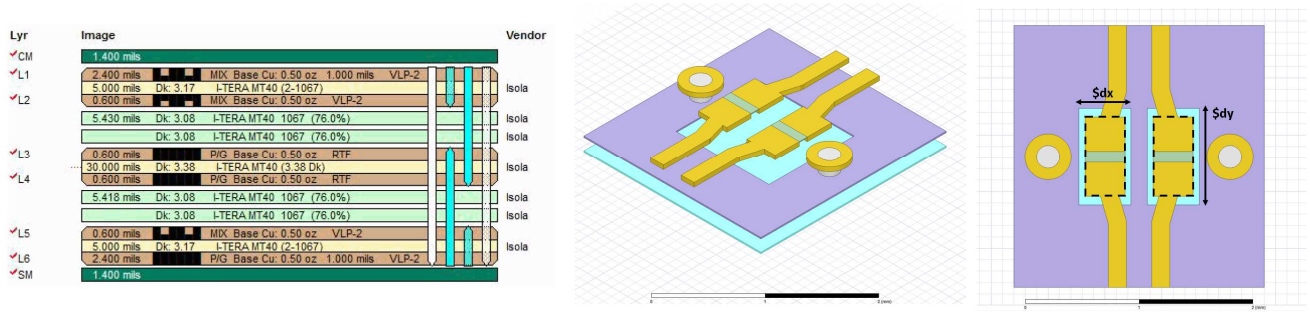

Figure 1 – Design Stackup and 220nF 0201 Shorted Pad Model with Ground Void Geometry Sweep

We chose to expedite the time required for the parametric solver by using wave port excitation and setting the solver error energy (deltaS) to 0.02. Once the parametric solve was complete, we narrowed down the results by finding candidate geometries which produced the largest frequency span which maintained at least -20dB of return loss. We further reduced this sample size to select the optimum geometry which had the best matched transition to 100ohms differential and produced the smallest settling time when viewed by TDR. The result of this analysis is shown below in Figure 2 and Figure 3 with the selected void geometry highlighted in neon green. The selected void openings are 22mils [0.5588mm] x 43.75mils [1.1112mm].

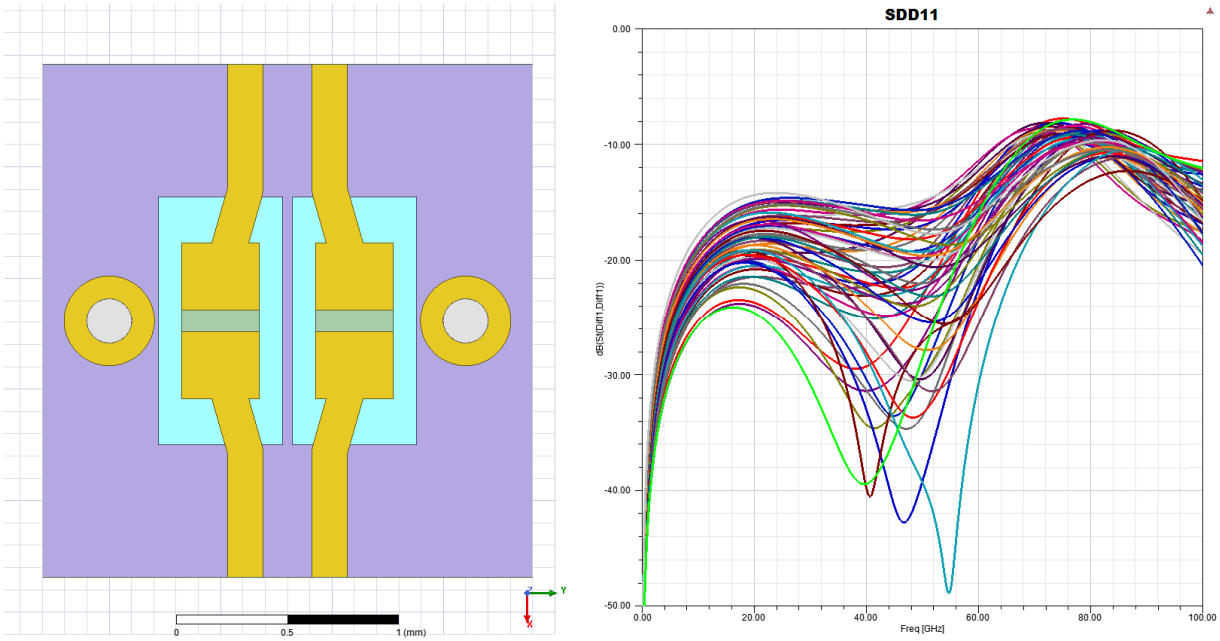

Figure 2 – Ground Void Optimization Layout and SDD11 Parametric Results

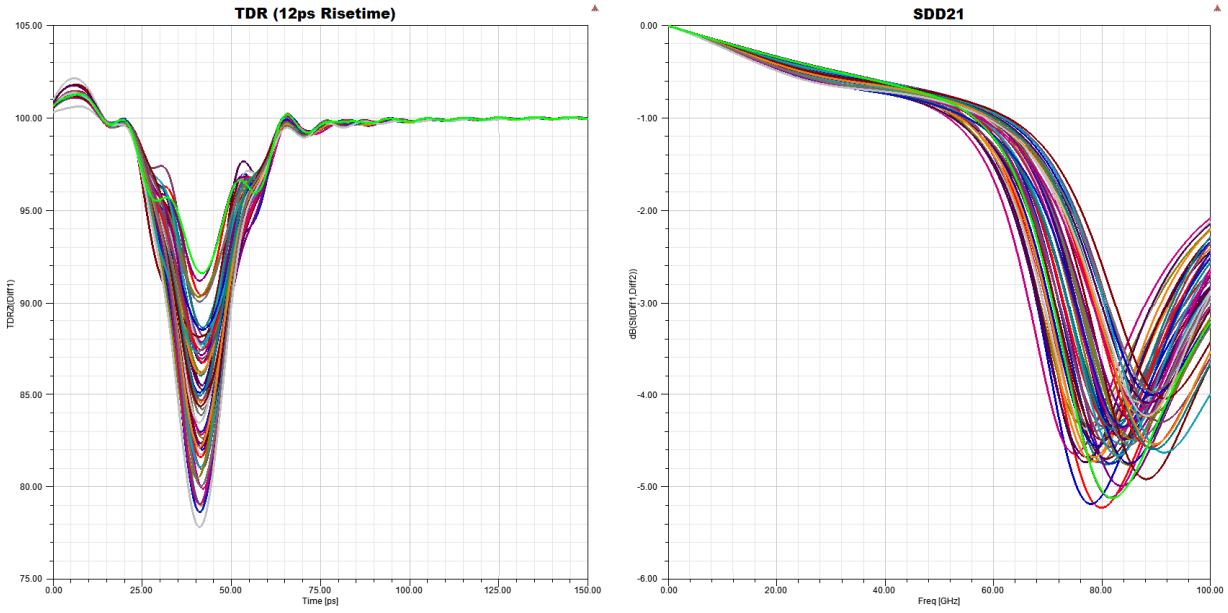

Figure 3 – Ground Void Optimization TDR and SDD21 Parametric Results

Correlation Test Vehicle

To enable our ultimate goal of correlating capacitor simulations to measurement results, we have developed a full 3D model of a multi-layer ceramic capacitor based on scanning electron microscope cross-sectional data. We have used the SEM to capture detailed images of 0201 package cross-sections across four different capacitor values in the TDK family (220nF, 150nF, 47nF and 22nF). The resulting SEM cross-sections are shown below in Figure 4.

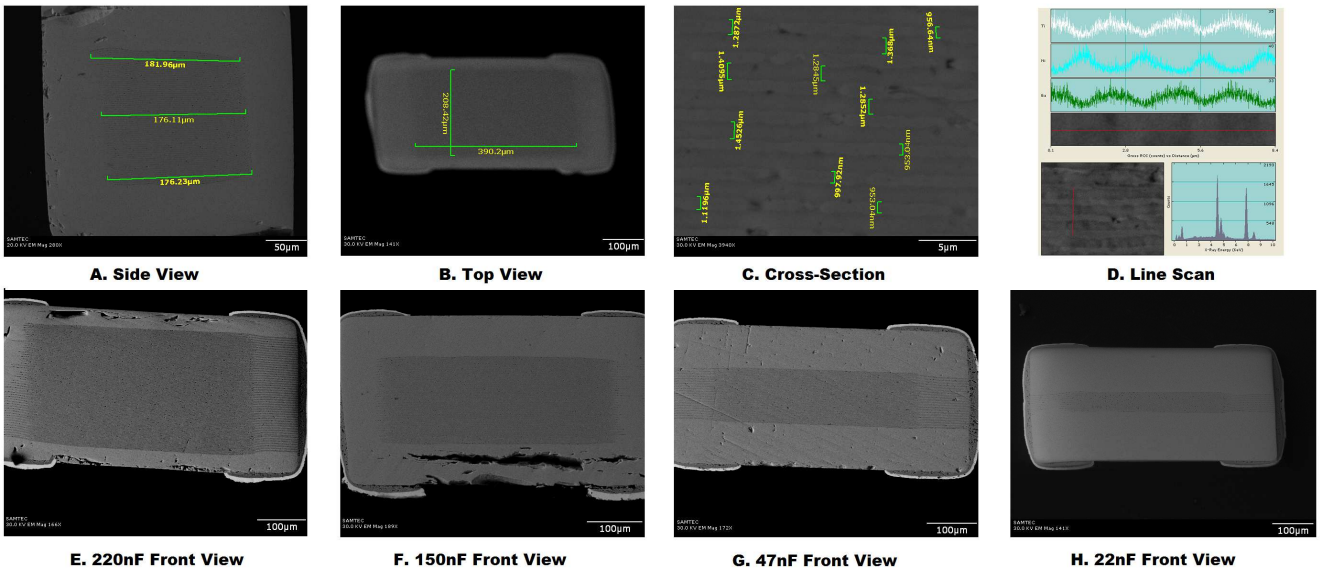

Figure 4 – 0201 Package SEM Cross-Sections and Material Line Scan

Complex MLCC Model

To complete the correlation vehicle, we have built a complex MLCC model of interdigitated parallel plates based on the cross-sectional data above and using MLCC modeling techniques referenced in a previous study [6]. When building this 3D model we have incorporated length variables to allow scalability to a variety of different manufacturing dimensions. Figure 5 below outlines the different model features and references their nominal dimensions for a 150nF (64 plate) 0201 TDK capacitor.

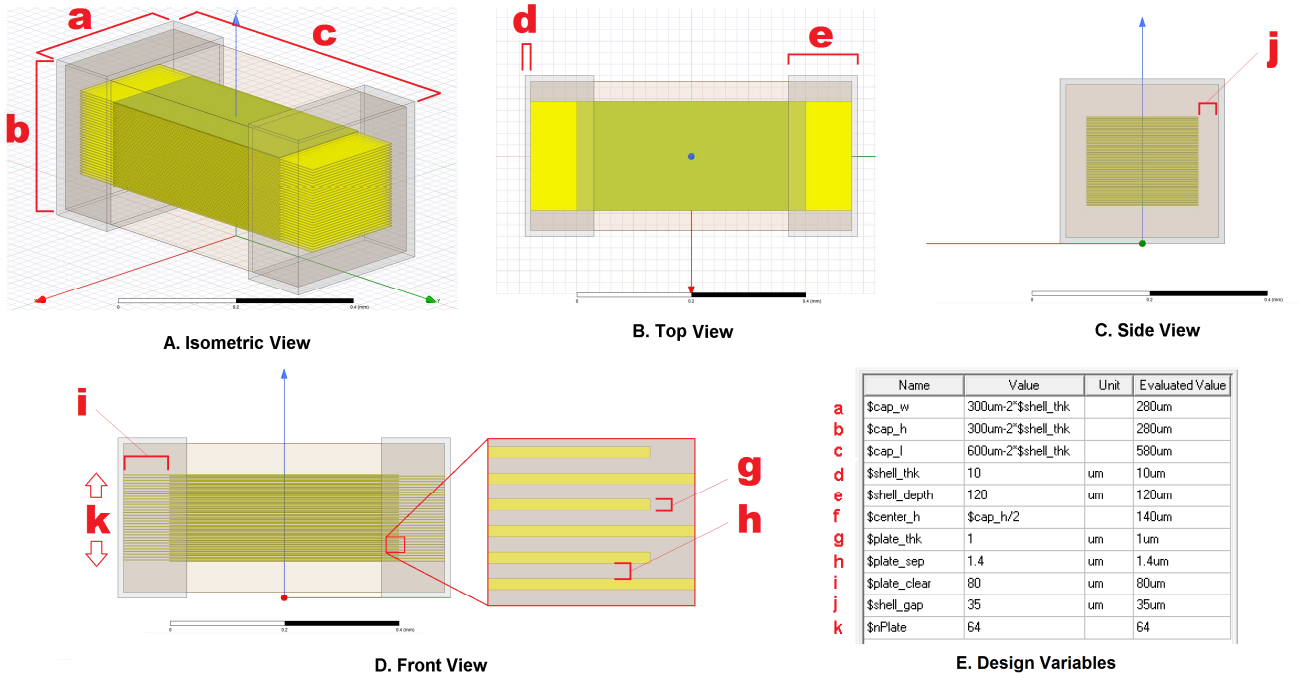

	Name	Value	Unit	Evaluated Value
a	$cap_w	300um-2*$shell_thk		280um
b	$cap_h	300um-2*$shell_thk		280um
c	$cap_l	600um-2*$shell_thk		580um
d	$shell_thk	10	um	10um
e	$shell_depth	120	um	120um
f	$center_h	$cap_h/2		140um
g	$plate_thk	1	um	1um
h	$plate_sep	1.4	um	1.4um
i	$plate_clear	80	um	80um
j	$shell_gap	35	um	35um
k	$nPlate	64		64

Figure 5 – HFSS 3D MLCC Model

Using the measurements derived from the SEM in Figure 4, we have back-calculated the dielectric constant required to achieve each capacitance value based on its observed number of plates and the parallel plate formula below with results seen in Table 1. It should be noted that each value of capacitor chosen for this study has a manufacturing tolerance of +/-10% with the exception of the 47nF capacitor (+/-20%) and can be shown that each value meets its nominal target within +/- 5%.

*Parallel Place Capacitance (Farads) = (E_o * E_r * n * A) / d*

Case Size	Capacitor Value (nF)	MFG, P/N	n Plates	Er	Estimated Capacitance (nF)	Tolerance (%)
201	220	TDK, C0603X7S0J224K030BC	90	5000	214.48	-2.51%
201	150	TDK, C0603X7S0J154K030BC	64	5000	151.82	1.21%
201	47	TDK, C0603X7S0J473M030BB	40	2500	46.99	-0.02%
201	22	TDK, C0603X7S0J223K030BB	20	2500	22.89	4.06%

Table 1 – Parallel Plate Calculation

Capacitor Shunt Configuration

To verify the accuracy of the complex MLCC model, we have simulated each value in shunt configuration and converted its reactance to equivalent capacitance in nanofarads to be compared with the manufacturer's nominal value. This formula is shown below. The results of this simulation can be seen in Figure 6 and Table 2 which shows that the models meet the nominal values within the specified tolerance for all but the 150nF capacitor case.

*Capacitance (Farads) = 1 / (2π * f * X_c)*

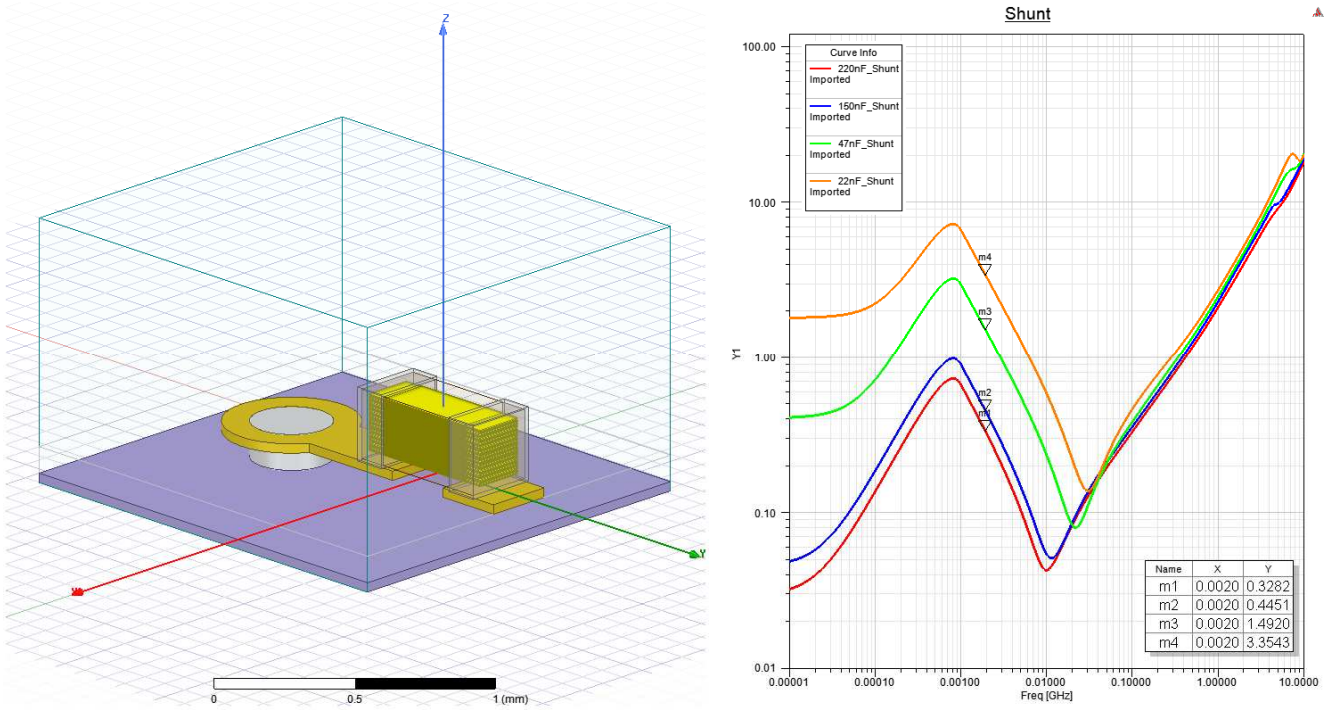

Figure 6 – Shunt Model & Z11 Results
220nF (red), 150nF (blue), 47nF (green), 22nF (orange)

Simulated Reactance X (Ω)	Simulated Capacitance C (nF)	Capacitance Target C (nF)	Tolerance (%)
0.3282	242.47	220	9.3%
0.4451	178.79	150	16.1%
1.4920	53.34	47	11.9%
3.3543	23.72	22	7.3%

Table 2 – Shunt Model Capacitance Correlation Table

Measurement Test Fixture

We designed the test fixture below with slight modifications to our optimization in order to increase manufacturing yields by using one continuous rectangular opening instead of two to eliminate the copper bridge between them. Our design made use of a 450um pitch GSGGSG differential microprobe structure to measure 1x CAL, 2x CAL, VOIDed and SOLID ground plane structures across multiple test fixtures to incorporate capacitor values of 220nF, 150nF, 47nF, and 22nF respectively.

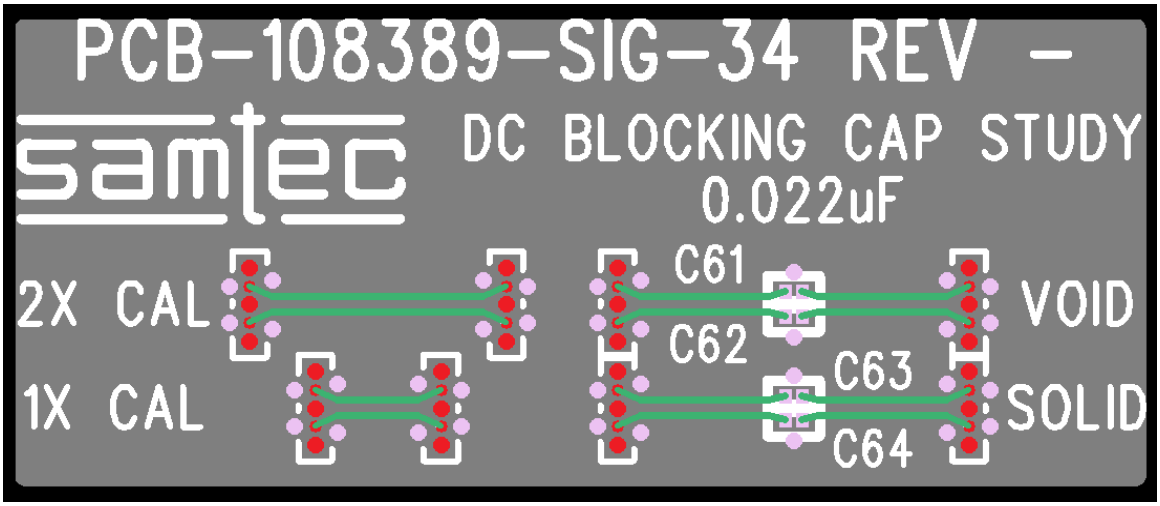

Figure 7 – Fixture Design

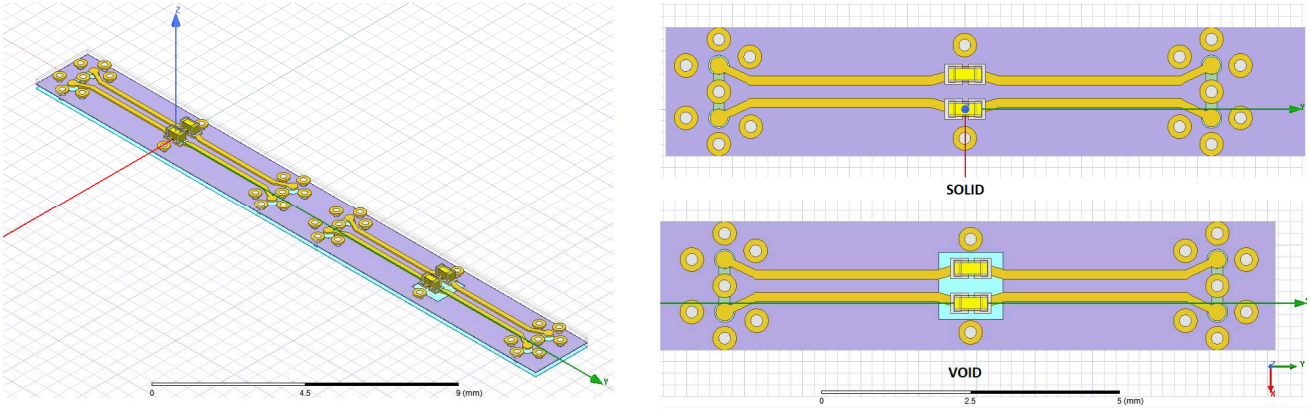

Figure 8 – DC Blocking Capacitor Fixture Simulation Model

The micro-probe launch is tuned to minimize the discontinuity when transitioning between the loosely-coupled and tightly-coupled regions of the probe launch into the DC blocking capacitor region. Below, we compare the simulation results of the test fixtures using different capacitor values and the effects of SOLID or VOIDed ground planes to later be compared with measurements. From the plots below we can conclude a few interesting results. First, each capacitor exhibits a primary resonance at a frequency inversely proportional to its shortest electrical path length as dictated by the lowest plate height and relative dielectric constant (see Table 1 for E_r values). Second, we see that the biggest impact of having the voided ground plane exists in frequencies greater than 20GHz.

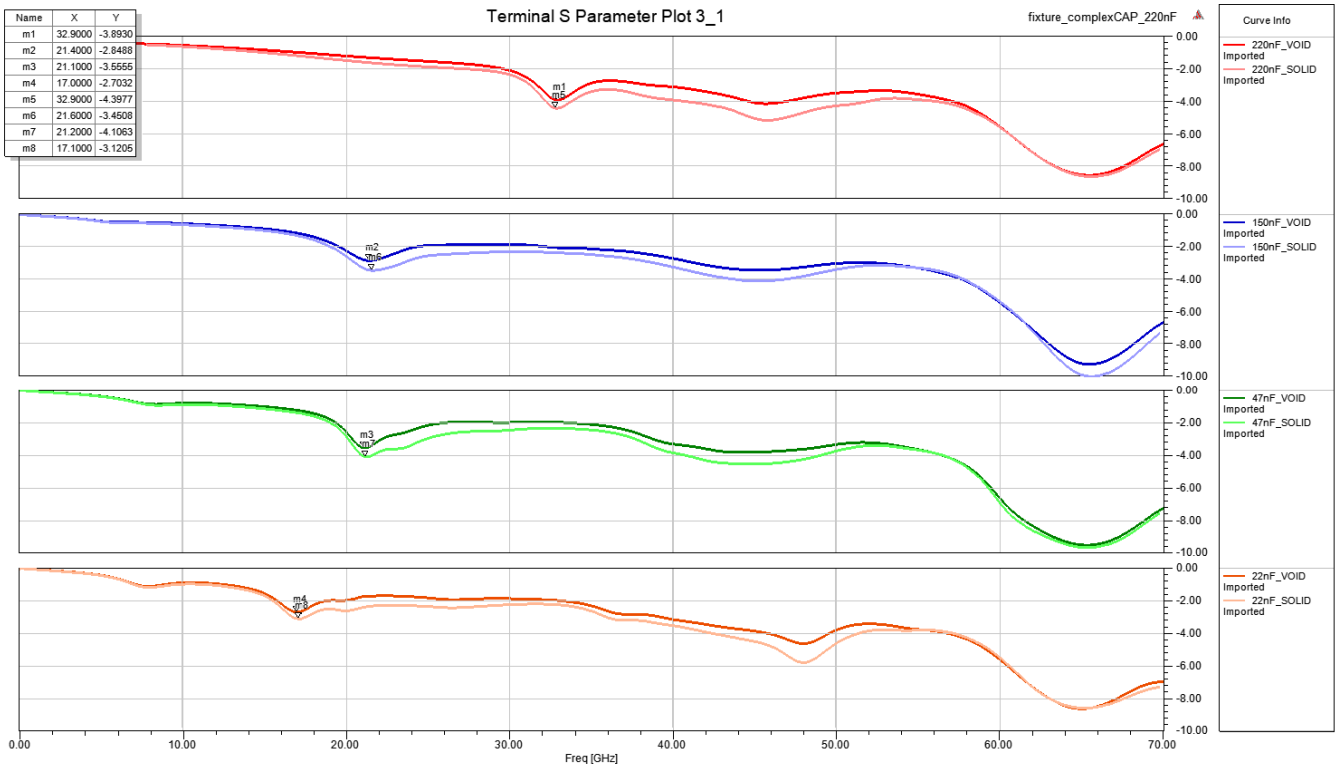

Figure 9 – HFSS Fixture Simulation of Capacitor Family Resonances
VOID gnd plane: 220nF (red), 150nF (blue), 47nF (green), 22nF (orange),
SOLID gnd plane: 220nF (light red), 150nF (light blue), 47nF (light green), 22nF (light orange)

Model Comparison Using First-Plate Capacitor Model

Based on the final void geometry used in this study, we compared the shorted pad, simple first-plate, and complex capacitor models to see the differences between them and how we might have optimized the structure more efficiently by using the first-plate model to approximate the complex capacitor model. The first-plate model incorporates only the lowest plate of the MLCC based on the least inductive path at high frequencies which contain the greatest current density [7]. The first-plate capacitor model provides many benefits in terms of simplicity and computation time in exchange for a slight penalty in accuracy. We believe this model has the potential to provide a good simplification for future work based on the results we show here. Based on the SEM measurements taken, we find the given 0201 capacitors with values of 220nF, 150nF, 47nF and 22nF have the first plate mounted at heights of 0.0527mm, 0.0839mm, 0.1127mm, and 0.1377mm respectively. Using the 220nF geometry as a starting point, we created a first-plate model of this capacitor by constructing two 10um nickel terminals shorted by a 1um thick tin plane at height 0.0527mm above the pad. We also included a dielectric slab that extends from the shorted plane down to a height of 10um above the surface of the mounting pad. The detailed geometry of this first-plate structure can be seen in the “SIMPLE_VOID” model in Figure 10 below. The use of the word “VOID” here is meant to differentiate between results which contain a voided ground geometry and those which have solid reference planes. In Figure 11, we include “COMPLEX_SOLID” to compare with the “COMPLEX_VOID” results to mark the performance difference seen when not using a voided ground plane. To achieve the results below, the solver was driven to an error energy (deltaS) of 0.0001.

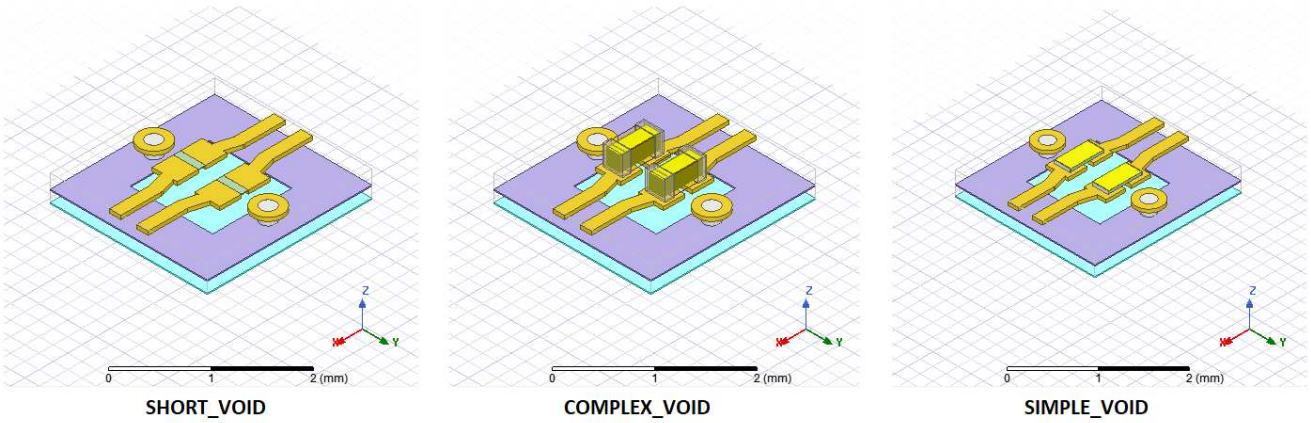

Figure 10 – 0201 220nF Model Comparison Layouts

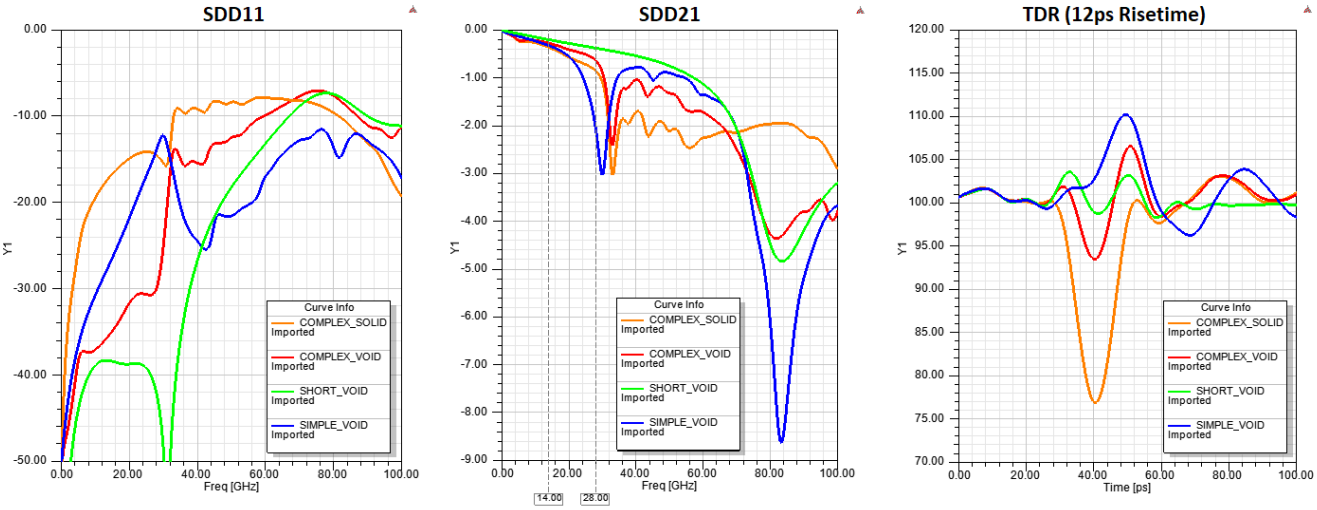

Figure 11 – 0201 220nF Model Comparison Results
COMPLEX_SOLID (orange), COMPLEX_VOID (red), SHORT_VOID (green), SIMPLE_VOID (blue)

By comparing the simulations of the three voided structures, we conclude that there are significant tradeoffs between the use of each model which influence both simulation time and resulting accuracy. The "SHORT_VOID" gives the most idealistic SDD21 response with no resonances in band, whereas the "SIMPLE_VOID" (first-plate) model more closely approximates the "COMPLEX_VOID" both in its first resonance magnitude near 30GHz and frequency-dependent loss. While the "SIMPLE_VOID" model exhibits similar characteristics to the complex MLCC model, its resonance band differs, resulting in a slightly more inductive TDR response. While not ideal, we believe the first-plate model better captures the complexity of the full MLCC capacitor model in the time and frequency domains and could be improved upon for better correlation in the future. It is interesting to note that there also exists a shallow 5.3GHz resonance in both of the complex capacitor simulations but not in either the shorted pad or the simple first-plate model. We believe this resonance is attributed to the length dimension of the overlapping plate geometries present in the complex MLCC model which exhibits a half-wave resonance given by the formula below (see this dimension in the shaded green area of Figure 5 – B. Top View).

$$f_{res}\ (GHz) = 1 / (2 * TD_{ns}) = 1 / (2* sqrt(5000) * 84.7 * 0.015748in / 1000) = 5.3012\ GHz$$

We also compare the single bit response (SBR) of the models above at 56Gbps using only feed-forward equalization (FFE) optimized for each capacitor case. It can be seen in Figure 12 below that the "SIMPLE_VOID" model over emphasizes the capacitor's contribution to ISI in terms of magnitude, whereas the "SHORT_VOID" model under predicts the contribution by limiting the duration of ISI as compared to the "COMPLEX_VOID" model.

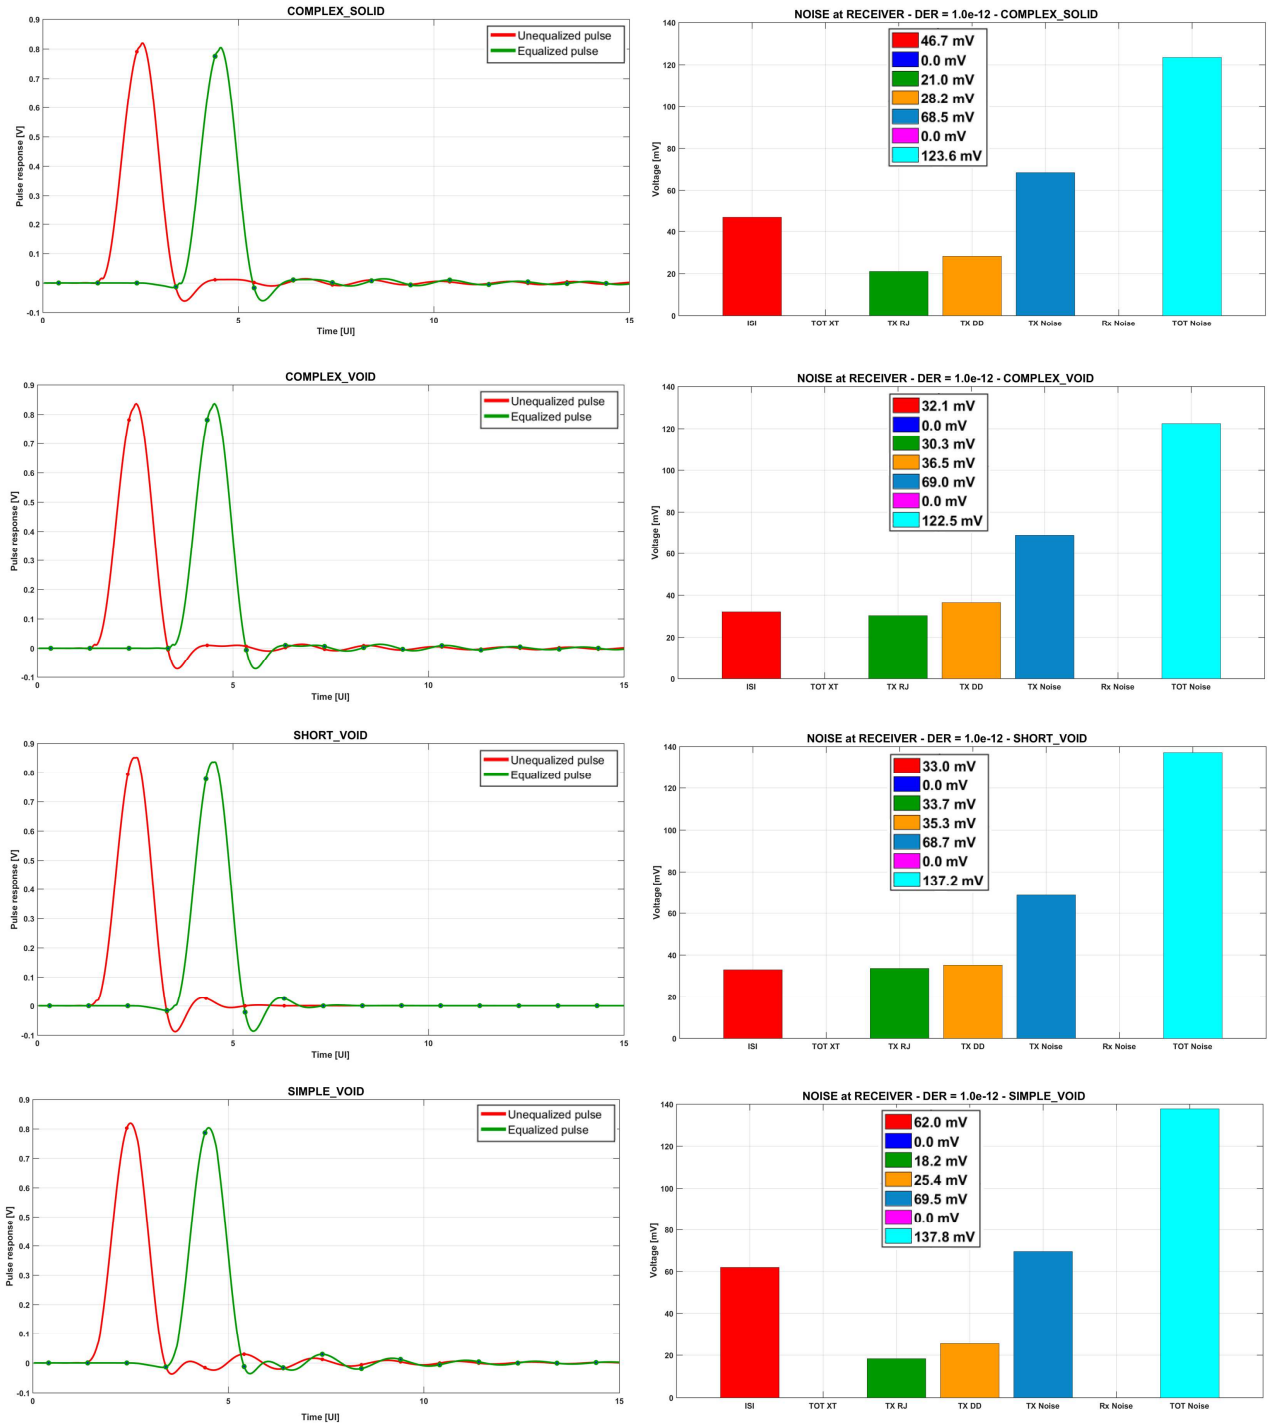

Figure 12 – Capacitor Model SBR and Noise Contribution Comparison

Full Channel Study

To study the impact of the different capacitor models in a full system design, we have simulated a single OIF-CEI-28G VSR chip to module channel utilizing the Samtec ARx6 connector in flyover configuration with no added crosstalk. The channel diagram shown below utilizes 1 inch of low-loss Meg6 stripline in the host board, 150mm of 100Ω differential twin-ax flyover cable, and 1 inch of Meg6 microstrip at the module. The compliance curves for this channel are shown below in Figure 13.

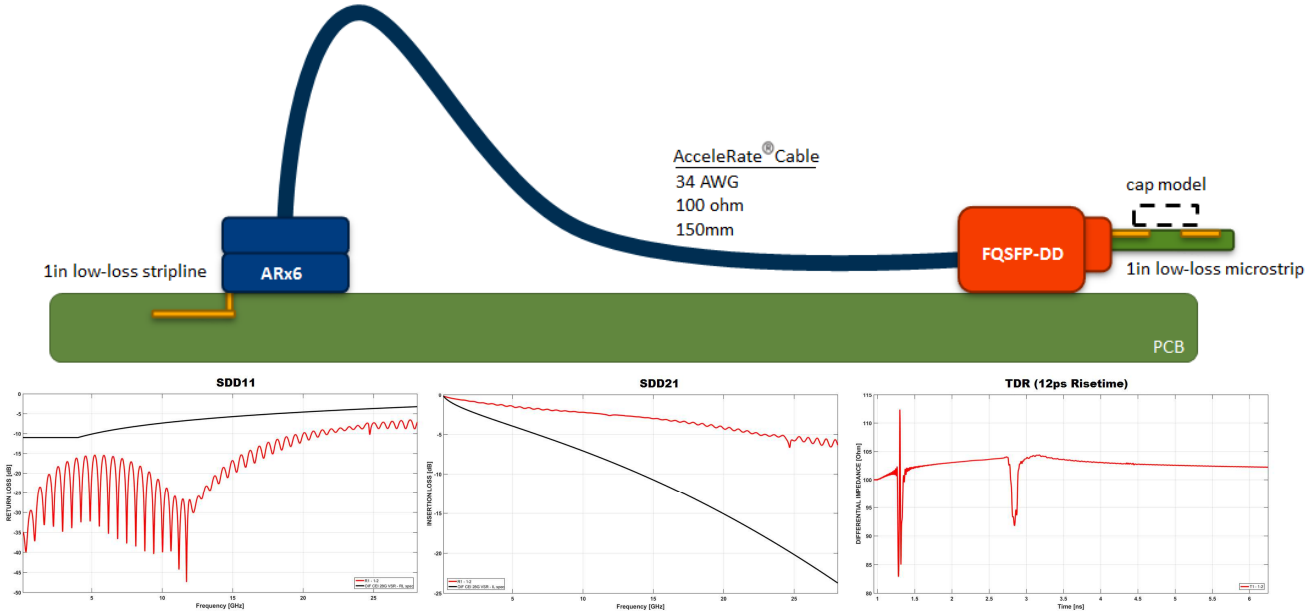

Figure 13 – Test Channel Diagram and OIF-CEI-28G VSR Compliance Curves

Since package and die parasitics present a major limitation to 56GHz operation, in this analysis we make the assumption that next generation designs will have minimal discontinuities to enable 56G NRZ and 112G PAM4. We have chosen the package and die values given below in Figure 14 for use with a low-loss packaging material of either ABF-GZ41 or ABF-GL102 (Df = 0.005) for both the host and module.

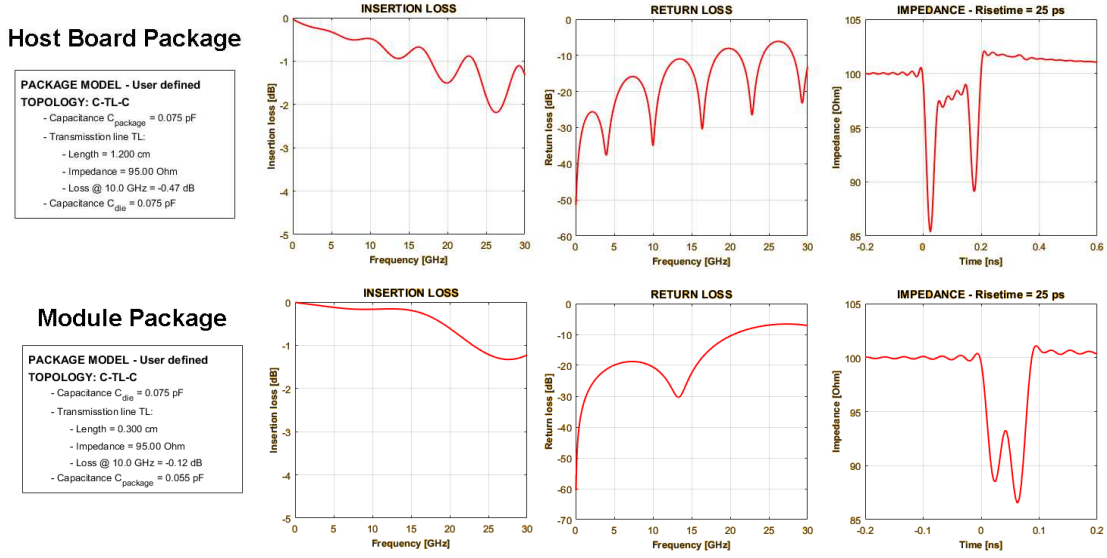

Figure 14 – Host and Module Package Parasitics

We have incorporated the different 220nF capacitor models into the VSR test channel and simulated at data rates of 28G NRZ, 56G NRZ, and 112G PAM4 using reasonable FFE and CTLE to get an understanding of how each configuration affects the channel performance under similar conditions. The equalization settings are shown below in Figure 15 and resulting channel performance in Figure 16.

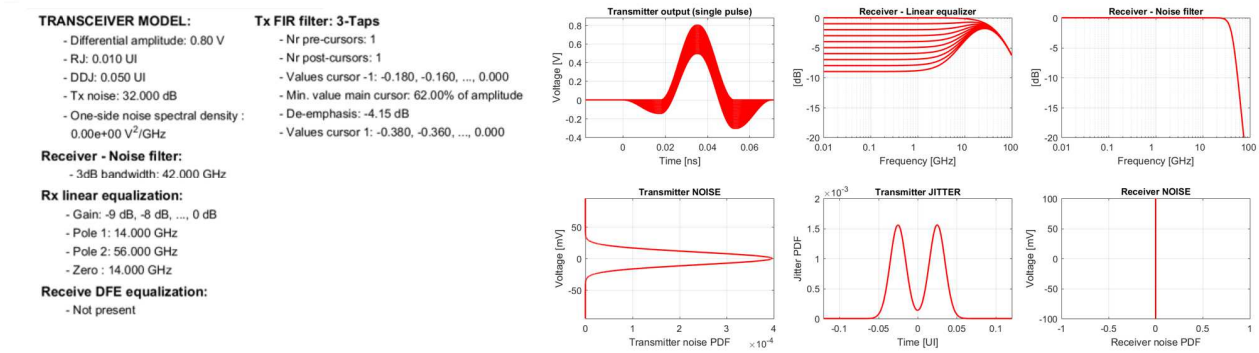

Figure 15 –Equalization Settings

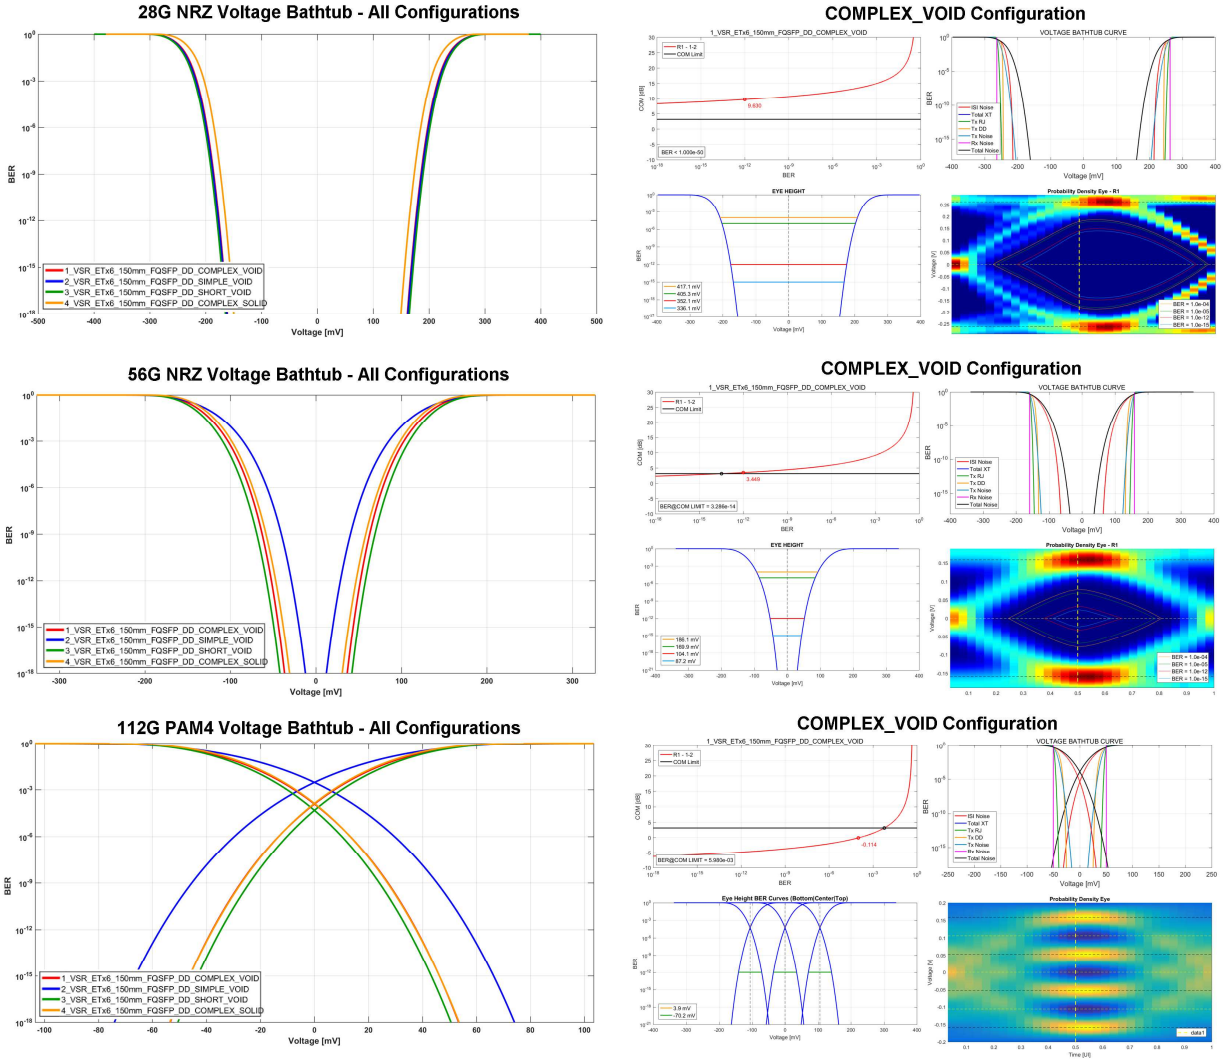

Figure 16 – Channel BER for 28G, 56G NRZ & 112G PAM4

As can be seen from the above voltage bathtub curves for all configurations, the "SHORT_VOID" model (green curve) gives a slightly more optimistic estimation of the BER at all data rates as compared to the "COMPLEX_VOID" model (red). Whereas, the "SIMPLE_VOID" model has a relatively good prediction of voltage at BER at 28G NRZ, but is very poor for 56G NRZ and 112G PAM4 due to the wideband resonance near 28GHz. The two markers at 14GHz and 28GHz in the SDD21 plot of Figure 11 highlight this difference in insertion loss between the "SIMPLE_VOID" and other models at these two frequencies. Due to this shortcoming, it is imperative we improve upon this model before use in future simulations. Lastly, to highlight the results of the 112G PAM4 simulation, it is likely that both FEC and DFE will be required to achieve an acceptable BER in this VSR channel.

Conclusion

The importance of good transition design for high speed serial links cannot be overstated. Below 10 Gigabit data rates, DC blocking capacitors and mounting parasitics are electrically small enough to be modeled as simple lumped element circuits. Starting with 25Gbps transmission, the signal bandwidth needs to be considered when generating channel models based on scattering parameters derived from physical attributes. Even greater care and modeling precision will be required as data rates transcend 100 Gbps. Capturing these detailed interactions will be especially relevant for future designs of PAM4 as it is more susceptible to common channel impairments at the same symbol rate than NRZ. These include a degraded SNR due to amplitude reduction, a higher susceptibility to transmitter variance, and less tolerance to ISI, insertion loss deviation, and cross-talk noise when operating in equivalent channels [8].

Preparation for designing at bit times approaching 18ps will require engineers to employ system level management of connectors, components, breakouts, and PCBs to minimize loss, deviation, intra-pair skew, inter-symbol interference, and cross-talk to maintain recoverable margins. While there exist alternative approaches to the design of board-level DC blocking, such as on-package or on-die implementations, each have significant tradeoffs to cost and performance that are not easily compared at the system level [9]. By generating accessible model data and continuing to improve and correlate our first-plate structure against a fully complex MLCC model and measurement data, we hope to make these time, cost, performance, and reliability tradeoffs more available to the designers of next generation systems.

Special Thanks

Special thanks to Dave Scopelliti, Brandon Gore, Rich Mellitz and John Abbott in recognition for their contributions to this paper.

References

[1] Altera AN-672, "Transceiver Link Design Guidelines for High-Gbps Data Rate Transmission," Altera Corp., CA.

[2] Altera AN-530 "Optimizing Impedance Discontinuity Caused by Surface Mount Pads for High-Speed Channel Designs," Altera Corp., CA.

[3] Weifeng Shu, Chunfei Ye, Dan Liu, Xiaoning Ye, Enrique Lopez, Xinjun Zhang, "DC Blocking Capacitor Design and Optimization for High Speed Signalling," IEEE Xplore.

[4] Mallikarjun Vasa, Arun Chada Reddy, Bhyrav Mutnury, Sanjay Kumar, Vasanth RD, “High Speed Interconnect Optimization, APEMC 2015.

[5] AVX Catalog, “Surface Mount Ceramic Capacitor Products”

[6] Scott McMorrow, Fabrizio Zanella, Steve Weir, “Using Full Wave Solvers for Practical Analysis of Capacitor Mounting Structures,” DesignCon, 2008.

[7] Istvan Novak, “A Black-Box Frequency Dependent Model of Capacitors for Frequency Domain Simulations,” DesignCon East, 2005.

[8] Hongtao Zhang, Brandon Jiao, Yu Liao, and Geoff Zhang, “PAM4 Signaling for 56G Serial Link Applications − A Tutorial,” DesignCon, 2016.

[9] Xin Wu, Casey Morrison, Bhavesh Patel, Raghav Nallan Chakravarth, Peerouz Amlesh, “SI and EMI Impact of AC Coupling Capacitors on 25Gbps+ Systems,” DesignCon. 2013