MASTER THESIS

JULY 2021

Modeling and assessment of complex interconnect features with focus on PAM-4 signal integrity

Student name   **Student number**
Marc Borrell Roig   s192377

DTU Supervisor

Tom Keinicke Johansen

External Supervisors

Claus Ek

Sune Graves Krohn

July 22, 2021

DTU Electrical Engineering

TABLE OF CONTENTS

Marc Borrell

ABSTRACT

As Internet demand continues to grow, the underlying technology struggles to keep up with the ever-increasing data rates. Existing 100G Ethernet is reaching its physical limits and superior solutions need to take over. PAM4 will be a key enabler capable of doubling the speeds of current NRZ solutions. However, it comes with its own implementation challenges, mainly from the signal integrity perspective.

This thesis evaluates the impact of complex channel characteristics on a 50 Gbps PAM4 transmission link using SI methods. The study is structured in three stages. First, several test boards are designed with multiple layout challenges (e.g., antipads, crosstalk, vias). Next, the effect of those designs on a PAM4 transmission are simulated using High-Speed Digital Design tools and evaluated in frequency and time domain. Finally, the boards are manufactured and measured with a VNA to validate the frequency results and an FPGA is used to validate the SERDES results. All these insights are condensed into layout guidelines.

The results show good agreement between simulations and measurements. In Part 1, the *Trace-to-Via transitions* and the *Pin-field Escape* show that a differential pair separation should be as close as possible to the via transition. The *Length equalization* and *Via types* results find the ideal case from various designs although the non-ideal solutions are also tolerable in most conditions. On the other hand, *Near-antipad* results proves that these do not cause significant issues. Finally, Part 2 validates that both the Ideal and the Tough *multi-feature interconnect* examples are adequate for a 50 Gbps PAM4 signal, even in crosstalk conditions. Therefore, the stackup used for this project is also validated.

Marc Borrell

ACKNOWLEDGEMENTS

NOMENCLATURE

ADS – Keysight Advanced Design System

AFR – Automatic fixture removal

ALA – Advanced link analyzer

BER – Bit error rate

BGA – Ball grid array

CEI – Common Electrical I/O

COM – Channel operating margin

DUT – Device under test

EM – Electromagnetic

EMI – Electromagnetic interference

FPGA – Field-programmable gate array

HSDD – High-speed digital design

IC – Integrated circuit

IEEE – Institute of Electrical and Electronics Engineers

ISI – Intersymbol interference

MR – Mid-range Reach

NIC – Network Interface Controller

NRZ – Non-return to zero

OIF – Optical Internetworking Forum

PAM4 – Pulse Amplitude Modulation 4-levels

PCB – Printed circuit board

PRBS – Pseudo-random Binary Sequence

QSFP – Quad Small Form-factor Pluggable

SERDES – Serializer de-serializer

SI – Signal Integrity

SNR – Signal to noise ratio

S-Parameters – Scattering parameters

TL – Transmission line

VNA – Vector Network Analyzer

VSR – Very Short Reach

Marc Borrell

LIST OF FIGURES

LIST OF TABLES

Marc Borrell

1 INTRODUCTION

1.1 BACKGROUND AND MOTIVATION

The thesis arises from the desire of accessing a specialized discipline known as High-speed digital design (HSDD). This is a field of electronics that focuses on high data throughput between devices across a physical interconnect. The need for faster digital communications is more relevant than ever in our society and the technologies that support it need to be developed further. HSDD is challenging because it requires deep knowledge in electromagnetics and expensive software and hardware tools. This project has been possible thanks to the support and collaboration between DTU as academic partner, and Napatech A/S as industry partner.

Napatech is a company leader in the SmartNIC's. They develop software for Internet datacenters and design their own boards. As a company, they continuously strive to achieve faster data rates to stay competitive, which puts a strain in the network channels across the board (see Figure 1). For this reason, the hardware department focuses on ensuring channel performance at a competitive cost.
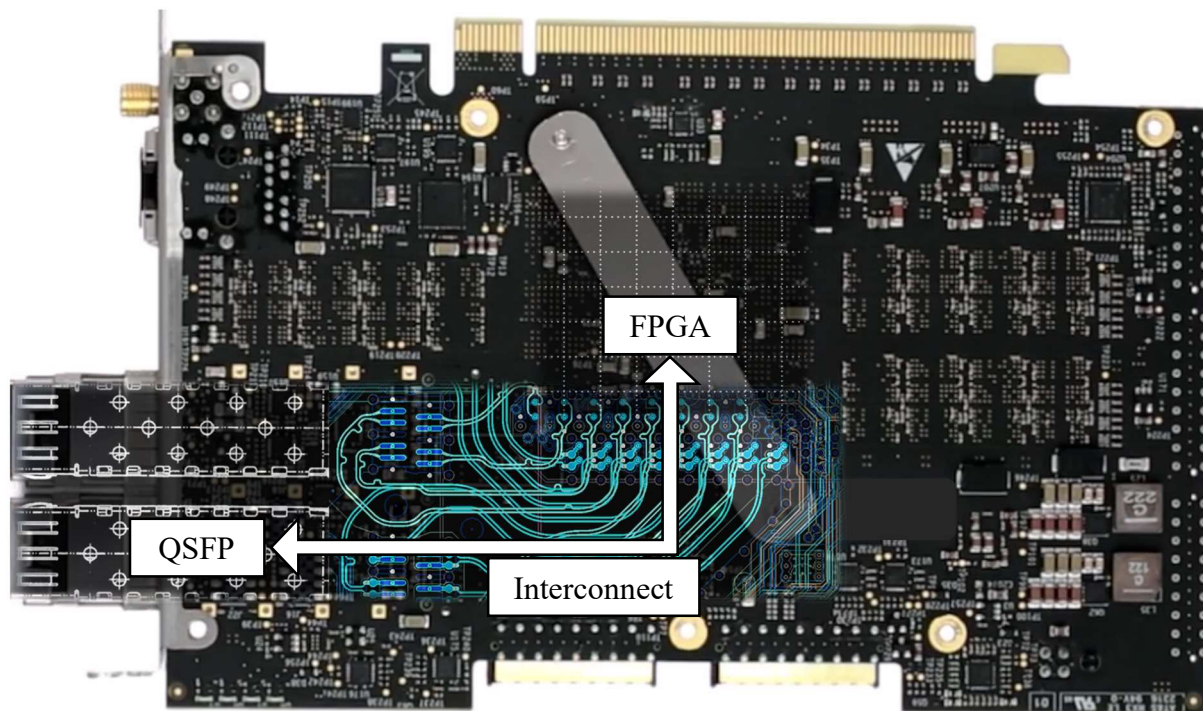
Figure 1: Interconnect location highlighted in light blue in a Napatech NIC

This thesis is inspired by the industry necessity to model and assess high data rate interconnects. PAM4 signals will be considered for this study as it is common in fast communication channels. Even though this modulation increases the complexity of the analysis, it will make it more relevant for current and future real-world applications.

1.2 OBJECTIVES AND RESEARCH QUESTIONS

The first objective is to model a broad set of interconnect discontinuities commonly used in PCB design (see Table 1). These interconnect characteristics are usually unavoidable as they

are imposed by the layout requirements, for example 90 degree turns, vias or other discontinuities. This part of the study sparks most of the research questions, such as what is the cause and effect of each discontinuity and what is the best way to avoid these effects. These discontinuities or features of study are compared against similar alternatives. In case there are no alternatives possible, they are compared against a standard to evaluate channel compliance.

Length matching, 90º turn, Pin-field Escape	Trace to Via transition, Vias	Antipad size, trace to pad transition	Reference plane discontinuity, Antipad size	Pin-field escape	Crosstalk

Table 1: Common interconnect discontinuities from older Napatech designs

In the second part, two demonstration channels are constructed. One consists of the best features that a designer could use in a channel representative of a real layout, while the other contains the worst features. The frequency and time results from these will help answer even more important questions among interconnect engineers, for example: Do the channels meet the standard requirements in each case? What is the most critical part of the interconnect?

Moreover, the overall objectives of the thesis are: 1) to understand the phenomenon behind interconnect discontinuities, 2) gain experience in design and modeling interconnects with HSDD tools while assessing each discontinuity. Finally, 3) verifying the simulation results with real measurements and correlating this assessment with the first point. All these studies are presented in a practical tone with useful real-world insights and conclusions.

1.3   RELATED WORK, ASSUMPTIONS AND LIMITATIONS

Fortunately, the study of interconnect features is an active field with plenty of resources to get inspiration from, for example journals, articles, books as well as other media formats (a small sample of these is found in the references). Be aware that each new board will contain different physical characteristics (e.g., the stackup or trace dimensions) that make each study unique. Therefore, available resources can serve as inspiration but not as design rules without a proper assessment.

For this reason, this project is limited to the materials and technologies used by Napatech as they offered to manufacture the test boards. The HSDD tools used also impose a few assumptions and limitations. There are two important limitations that need to be highlighted.

1. Keysight ADS FEM solver does not consider the etch of the traces so they need to be converted to a rectangular shape by averaging the top and bottom edge (see Figure 2)
2. Keysight ADS FEM solver does not consider surface roughness.

Figure 2: Trace section from trapezoidal to rectangular

Marc Borrell

1.4   CONTENT OF THE REPORT

THE FIRST PART contains individual features organized by topology which result in six categories:

- *Length matching*: Four sizes of bends are compared with two references
- *Via modeling*: Six designs with micro-vias and buried vias are validated
- *Trace-to-via transitions:* Six variations are studied to find the best design
- *Near-antipad discontinuity:* Three BGA sizes in five variations are analyzed
- *Pin-field escape*: Three escape methods with 90º turn are compared
- *Crosstalk*: two parameters studied. Spacing and coupling length.

All these are simulated and compared both in frequency and time domain. Frequency domain results are useful to characterize the performance as well as the cause of the discontinuity while the time domain results help evaluate the effect of each discontinuity. Finally, the results are concluded into layout guidelines. To ensure the simulations are accurate, the most relevant features are manufactured and verified individually in frequency domain with a 4-Port VNA.

THE SECOND PART contains two interconnects with features typically found in a NIC interconnect. One has the best performing or "Ideal" features from Part 1. The "Ideal" features can't always be used in a new layout; therefore a "Tough" channel is created with the most challenging counterparts that one could be forced to use, e.g., due to space limitations. These channels also allow to easily evaluate the limitations of the materials used in the stackup.

The two channels are characterized in frequency domain by cascading the S-Parameter models from Part 1 and validated using the masks specified in [1] and [2]. These channels are also manufactured and verified individually in frequency domain with a 4-Port VNA.

Finally, eye diagram verification is done with an Intel Stratix10TX PAM4 transmitter and receiver. To have accurate simulations of the SERDES channel, all discontinuities in the close-loop path (Stratix10TX board, coaxial cables and connectors) are also modeled and included.

1.5   TOOLS USED

Before commencing the core study, these are some of the most common tools used throughout the report. To keep the report concise, it includes only the tools used to explain the results therefore excluding other tools used such as the 3D solvers or the FPGA software. The only exception is the PAM4 explanation, which is necessary to understand the report.

1.5.1   Pulse Amplitude Modulation 4-levels (PAM4)

Up until recently, the most common differential signal modulation was non-return to zero (NRZ) but the telecommunications industry is constantly pushing for higher data rates. NRZ speeds can be increased by increasing the frequency of the signals until a certain point where the losses become too large. One option would be to compensate these losses with equalization but it is increasingly difficult.

Here is where PAM4 has an advantage over NRZ. As the name implies, there is 2 bits of data encoded in the amplitude of the signal instead of a single bit. For the same reason, to transmit the same amount of data, the Nyquist frequency is half of the NRZ (see Figure 3).
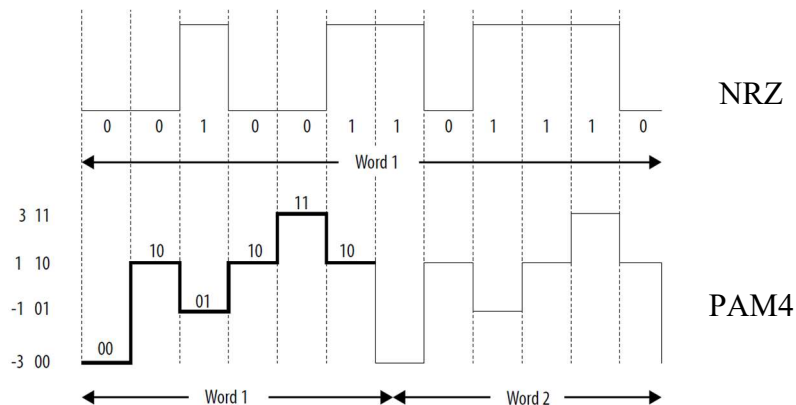
Figure 3: PAM4 vs NRZ signal modulation, [3]

The challenge that comes with PAM4 is that the signal-to-noise ratio (SNR) is reduced. Knowing that the modulation has three levels, the loss in SNR compared to an NRZ communication can be calculated with Equation 1. An in-depth look at PAM4 is found in [3].

$$SNR_{loss} = 20 \times \log_{10}\left(\frac{1}{3}\right) = -9.5 \, dB$$

Equation 1: SNR loss

Notice that the SNR loss increases the challenges from the signal integrity point of view. To verify that the signal is valid, Ethernet standards will be used as guidelines. It is important to use PAM4 modulation to verify the channel performance because it is a technology that is positioned as a key enabler of future high-speed communications.

1.5.2   Frequency domain tools

Since all these features are differential transmission lines, they are simulated with 4 single-ended nodal ports and converted to modal S-Parameters for the frequency domain study. The most relevant plots to look for in the frequency domain are:

- RETURN COEFFICIENT (SDD11): is related to the differential return loss and measures the differential response at port 1 when excited at port 1 in differential mode. It is useful for comparisons between features (the lower the better) and standard compliance with a mask.
- TRANSMISSION COEFFICIENT (SDD21): is related to differential insertion loss and measures the differential response at port 2 when excited at port 1 in differential mode. It typically shows a linear loss. It can also help spot issues, for example large dips can appear due to resonances.
- MODE CONVERSION COEFFICIENT (SCD21): is related to differential to common mode conversion and measures the common response at port 2 when excited at port 1 in differential mode. In practical terms, it can be thought as the susceptibility and generation of EM [4].

Marc Borrell

Figure 4 shows and example of the S-parameter & TDR template used to study and compare the features. The most relevant plots used throughout the report are highlighted in red. TDR was added to the template for convenience but belongs to the time domain analysis.
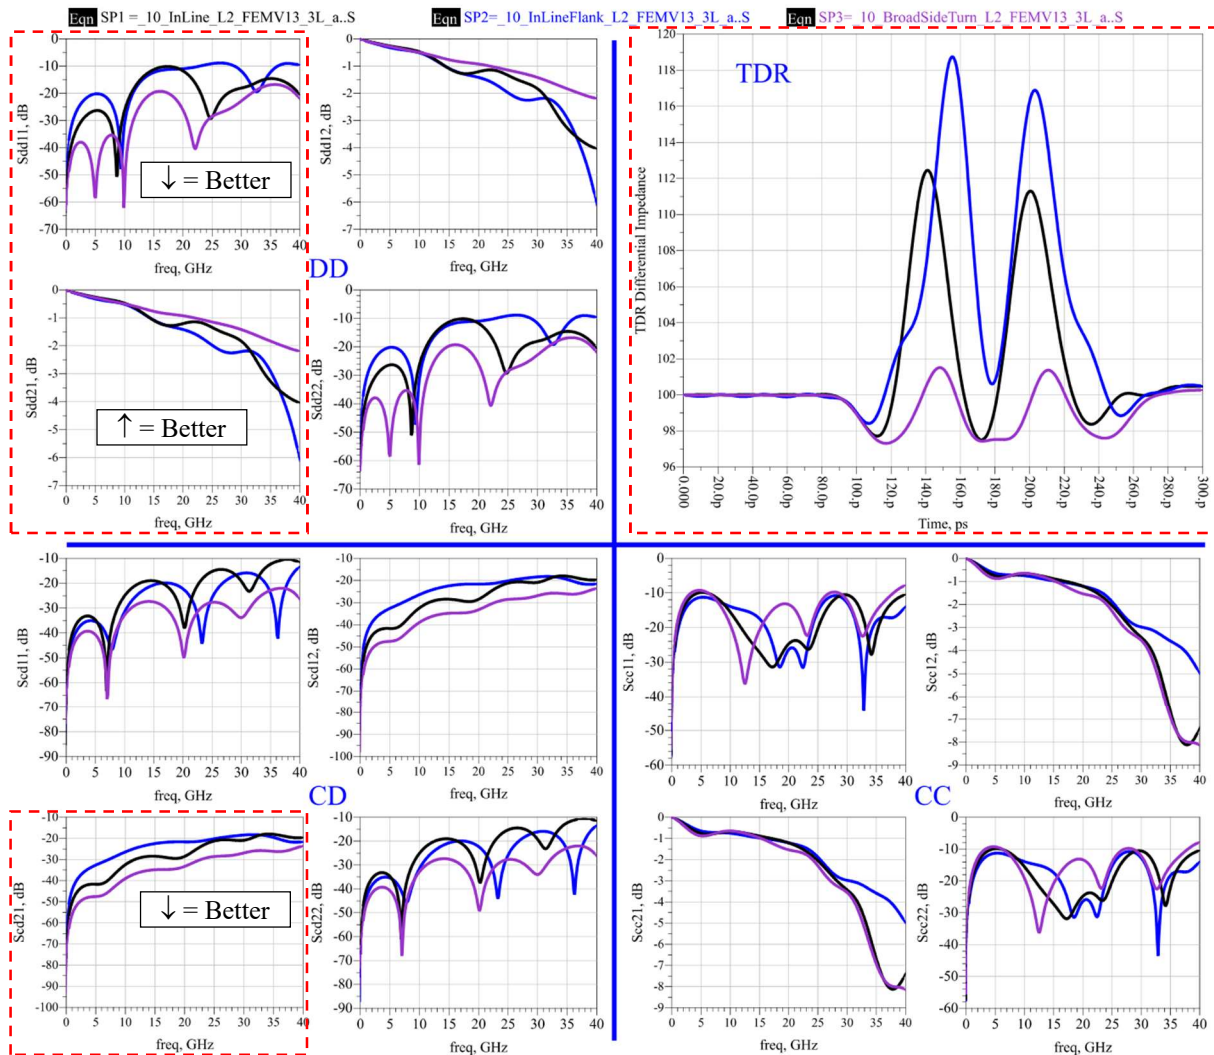
Figure 4: Complete S-Parameter matrix and TDR

1.5.3 Time domain tools

Starting with TDR, it provides a method for impedance monitoring of a transmission line. The TDR plot always starts at the input port with reference impedance (in this project the differential reference impedance is always 100 ohm) and then drops after 100 ps (because of port extension) to the simulation characteristic impedance of the transmission line (97 ohm).

Since the TDR is created from the frequency domain results (SDD11), the losses of the line will increase the impedance thus creating an upward ramp. Discontinuities will appear as spikes in the plot. To be specific, if the spike is above the characteristic impedance, the discontinuity is known to be inductive whereas down-spikes are capacitive. See the following example in Figure 5.

Marc Borrell
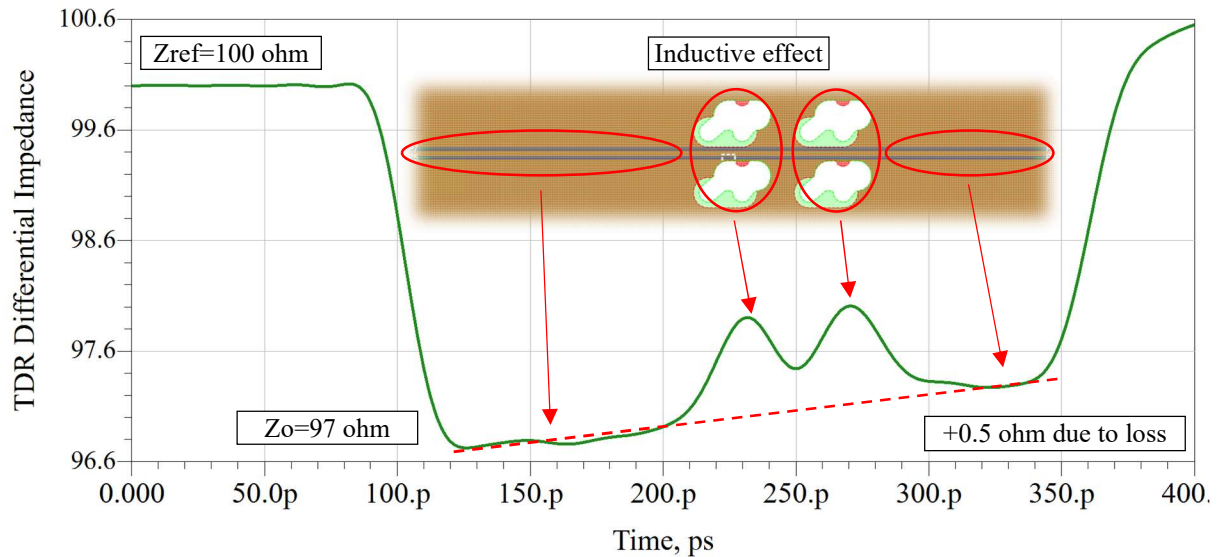
Figure 5: Simulated TDR example of an anti-pad feature

Also in the time domain, the most relevant plots for SERDES channels are Eye diagrams, BER contour plots and Channel Operating Margin (COM).

- EYE DENSITY PLOT is the superposition of multiple UI, slicing a PRBS and overlapping the rising and falling edges. Keysight ADS presents the eye diagram as a probability of a signal passing through that point in the plot. It is desirable to have open eyes and large enough that the receiver can distinguish between signal levels. The eye can be measured at the output of the channel or the receiver. The channel output shows the raw effect of the discontinuities on the signal. For example, if the channel is long, the eye will be closed due to losses. On the other hand, the receiver has two equalization tools (SR CTLE and SR FFE) that can compensate the loss and open the eye.

It is important to clarify how the eyes are measured. First one finds the time center of the middle eye at BER 10^{-3} (T_{mid}). Then the height of the eye is measured in volts as V_{upp}, V_{mid} and V_{low} at BER 10^{-6}. Next, the middle point of each voltage measured in the previous step is used as the level at which the width is measured (H_{upp}, H_{mid} and H_{low} in picoseconds). Finally, the eye amplitude (AV_{low}, AV_{mid}, AV_{upp}), is the height between decision levels (see Figure 6).

Once these measurements are acquired, additional calculations are required to be able to compare the measurements against the standard specifications in [1]. The requirements for OIF CEI 56G VSR PAM4 are:

- Minimum eye width at BER 10^{-6} (EW6) = 0.2 UI
- Minimum eye height at BER 10^{-6} (EH6) = 32mV
- Eye linearity (EL) = 0.85

Equation 2, Equation 3 and Equation 4 are used to calculate EW6, EH6 and EL where UI is 40 ps. This method and Figure 6 come from the Ethernet standard [1].

Marc Borrell

$$EW6 = \frac{\min\left(H_{low}, H_{mid}, H_{upp}\right)}{UI}$$

Equation 2

$$EH6 = \min\left(V_{low}, V_{mid}, V_{upp}\right)$$

Equation 3

$$EL = \frac{\min(AV_{low}, AV_{low}, AV_{low})}{\max(AV_{low}, AV_{low}, AV_{low})}$$
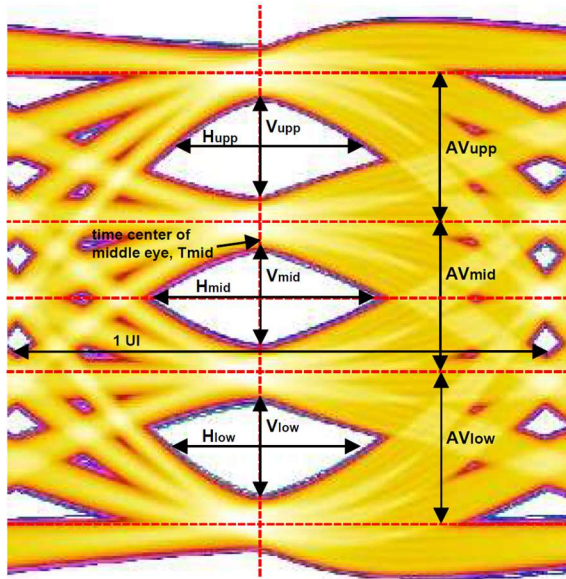
Equation 4

Figure 6: PAM4 eye example and measurements

- EYE CONTOUR PLOTS are a simplified version of the eye diagram. It displays the inner perimeter of the eye at a specific BER, usually BER 10^{-6} and allows easy comparison between multiple simulations or measurements. All the contour plots presented in Part 1 contain only the lower eye instead of the 3 eyes that PAM4 has. The intension is to show the most critical eye and remove confusion that the full plot might create.

- COM is a single metric that is computed in time domain and measures ISI, crosstalk and loss in a standardized transmitter and receiver configuration. COM analysis is configured in an Excel sheet specific to the channel characteristics which is provided by OIF [1]. This Excel sheet is used by a MATLAB script which can be run from most simulation tools (e.g ADS, Hyperlynx, Intel ALA). In this report, two standards are used depending on the length of the channel. OIF CEI 56G Very short reach (VSR) is used for individual features and open-loop while Mid-range reach MR. Only the MR standard uses COM to validate the channel performance.

1.5.4   Validation tools

Every part of this report is validated with real measurements. To do so, multiple test boards are manufactured with the same layout and stackup as used in the simulations. First, a Keysight 53GHz 4-Port VNA is used to characterize the features in frequency domain up to 40 GHz (limitation imposed by the maximum connector frequency). It is also used for TDR measurements with a 20 ps rise time (note the instrument uses the frequency domain data to produce the TDR. It is not a using a step). Then, an Intel Stratix10TX FPGA Evaluation Board is used to validate the SERDES simulations using the build-in PAM4 transceivers.

- BOARD CHARACTERIZATION: In order to achieve accurate results, calibration of the VNA is done with a passive calibration kit at the coax cables (see Figure 7). In addition, measurements are 1601 points with an average of 50 samples and a 2% smooth factor. To be able to compare the measurements with the simulation results, the connector transition must be de-embedded (this is from the coax cable to the differential stripline at the beginning of the feature). The VNA has an Automatic fixture removal (AFR)

Marc Borrell

which requires careful tuning of the de-embedded length (in this case 94.5 ps). A small piece of advice is that TDR measurements can be of great help during this "automatic" process. TDR is also captured for each feature and compared to its simulation counterpart. It is important to mention that the TDR from simulations is also extracted from the frequency domain data.

Figure 7: Keysight 53GHz 4-Port VNA connected to a test board

- SERDES MEASUREMENTS: The SERDES results, for example eye density plots, are acquired from a PRBS pattern of a PAM4 signal at 50 Gbps as specified by the thesis objectives. The Intel Stratix10TX is an FPGA with transceivers capable of achieving these speeds and measuring an eye density plot. A test board is connected in series between the transmitter and receiver of the FPGA (see Figure 8). The eye openings measured by the Stratix10TX depend, not only on the test board channel, but also on the connectors, coaxial cables and the channel on the FPGA board. Therefore, only the Closed-loop multi-feature channels discussed in Part 2 are compared against the Stratix10TX measurements because those consider the entire path.
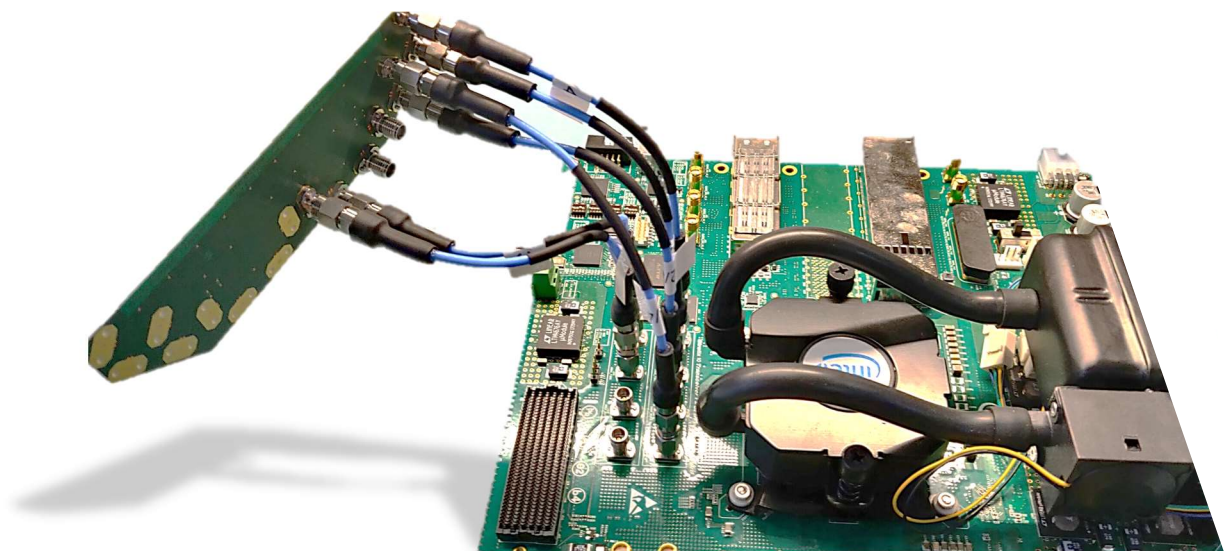
Figure 8: Intel Stratix10TX FPGA connected to a test board

Marc Borrell

2   PART 1: INDIVIDUAL FEATURES BY TOPOLOGY

The features selected for study aims to cover the most common discontinuities found in high-speed interconnects. The main source of inspiration comes from previous Napatech designs (see Table 1), as well as from IEEE papers which will be mentioned individually. The layout challenges for study have been created from scratch to facilitate comparison between similar cases and ease the conclusion process. Some of the features have been designed with parametric dimensions which unlocks advanced capabilities of the HSDD tools, for example these features can be swept, tuned or optimized.

2.1   LENGTH MATCHING DESIGN

Length matching is a required and unavoidable feature when a differential pair changes direction. The turn required to change direction makes the length of the internal trace shorter that the external and the effect of this miss-match is that one traveling wave will arrive later than its opposite pair at the receiver (see Figure 9). This difference is known as skew, and it is measured in time scale, e.g., in picoseconds. The consequence of skew is that the signal is converted from differential to common mode which are rejected by design in differential communications, i.e., part of the energy transmitted will be ignored by the receiver and will be radiated as EMI [4].
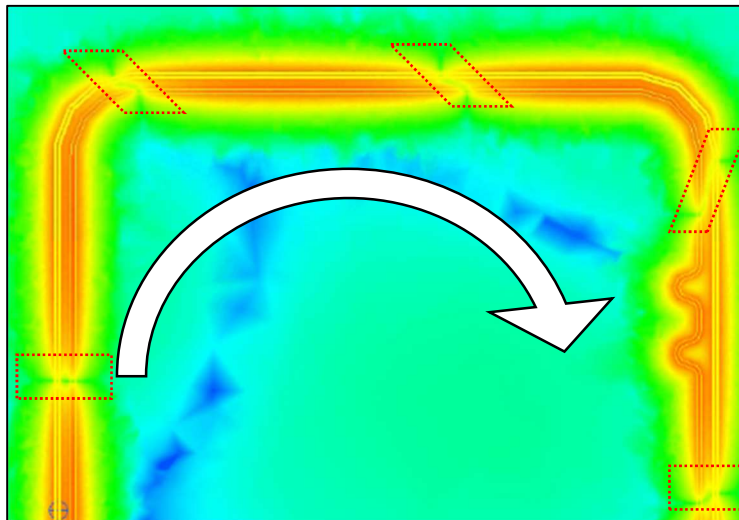
Figure 9: Length equalization feature "small" with E-field in log-scale

There are multiple options to match the length of a pair. One of the most common methods is creating a small bend or meander at the short trace. There are multiple bend shapes and sizes that can be used while having the same skew compensation. For example, the turn can be wider or smaller with one or multiple bends. Multiple options have been previously investigated such as Asymmetric dual bends [5], Novel length equalization structures [6] or Meander line optimization [7]. Based on [7], an extended set of bends is created for study: wide, big, slim and small (see Figure 10). These features compensate a 180º turn that causes the inner trace to be 1 mm shorter than the outer causing a 5.4 ps skew. To compensate this skew, different size bends are positioned 6 mm away from the turn at the inner trace.

Parameter	No-deskew	Wide	Big	Slim	Small (x2)	Straight
Width	0 mm	1 mm	0.75 mm	0.3 mm	0.37 mm	-
Hight	0 mm	0.74 mm	0.74 mm	0.62 mm	0.37mm	-
Skew	- 5.4 ps	0 ps	0 ps	0 ps	0 ps	0 ps
Length	24.14 mm	24.14 mm	24.14 mm	24.14 mm	24.14 mm	24.14 mm
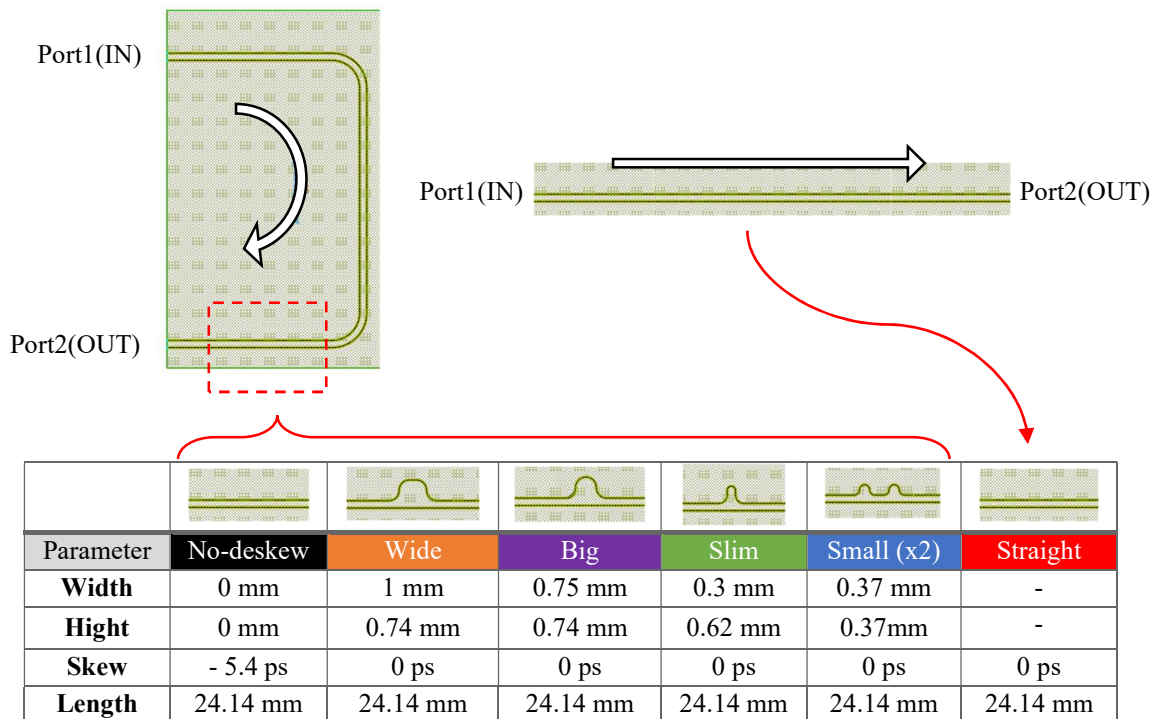
Figure 10: Length equalization features, layout and dimensions.
5 layouts are constructed with 2 90° turns with different deskew methods in the same position highlighted with a red rectangle. 1 additional layout does not have turns but has the same length as the other cases and serves as a baseline

2.1.1   Frequency domain analysis

After running a full 3D EM simulations, the modal S-Parameters show:

<u>Transmission coefficient, SDD21:</u> is the same in all cases but the **No-deskew**, which has a much higher loss at high frequencies (see Figure 12). For example, the expected loss of Straight feature is 3.2dB@40GHz but with 5.4 ps of uncorrected skew, the **No-deskew** has a loss of 5.2dB@40GHz (notice they are the same length). That is because skew converts the mode of propagation from differential to common and SDD21 measures only the differential signal out of Port2. Furthermore, when the frequency of the waveform increases, the waveform unit interval (UI) decreases and since the mode conversion can be expressed Equation 5, the effect becomes severe at higher frequencies.

$$\% \; mode \; conversion = \frac{skew}{UI}$$ *Equation 5: Percentage of mode conversion*
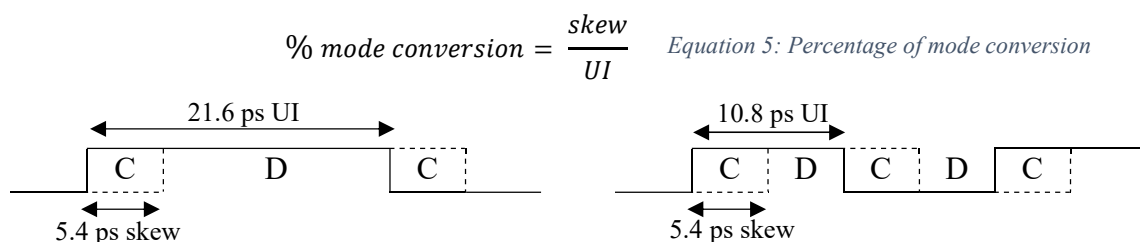
Figure 11: Illustration of mode conversion (C=common, D=differential)

<u>Mode conversion coefficient, SCD21:</u> Is the most relevant of the three because it shows the differential to common mode conversion which is expected when skew is not correctly compensated as discussed before. The results are ordered from best to worst mode conversion as follows: Straight, Big, Wide, Slim, Small, **No-deskew**.

Marc Borrell

It is expected that the **No-deskew** has the most mode conversion because it has 5.4 ps of uncorrected skew whereas Straight has the best performance as it does not have any turns or bumps. In fact, theoretically the Straight stripline should not have mode conversion but because the dielectric is not equally space around the trace and 3D solvers can only be accurate to a certain noise floor, the simulation finally converges to 59 dB@14GHz.

On the other hand, between the 4 different bump sizes, the best one in terms of mode conversion is the Big size. The results shows that there is a sweet spot where the size of the length equalization bump should be big but not too large (see Figure 12).

Return coefficient, SDD11: Straight and **No-deskew** are very similar and have the best return loss followed by Small and Slim and finally Wide and Big (see Figure 12). Notice that the difference between the best and the worst sizes is only 8 dB@14 GHz. The worst cases indicate that the signal is encountering a large discontinuity in the path and it is being reflected. The cause and effect of this is better explained with the time domain results (see Figure 13) discussed in the next paragraph.
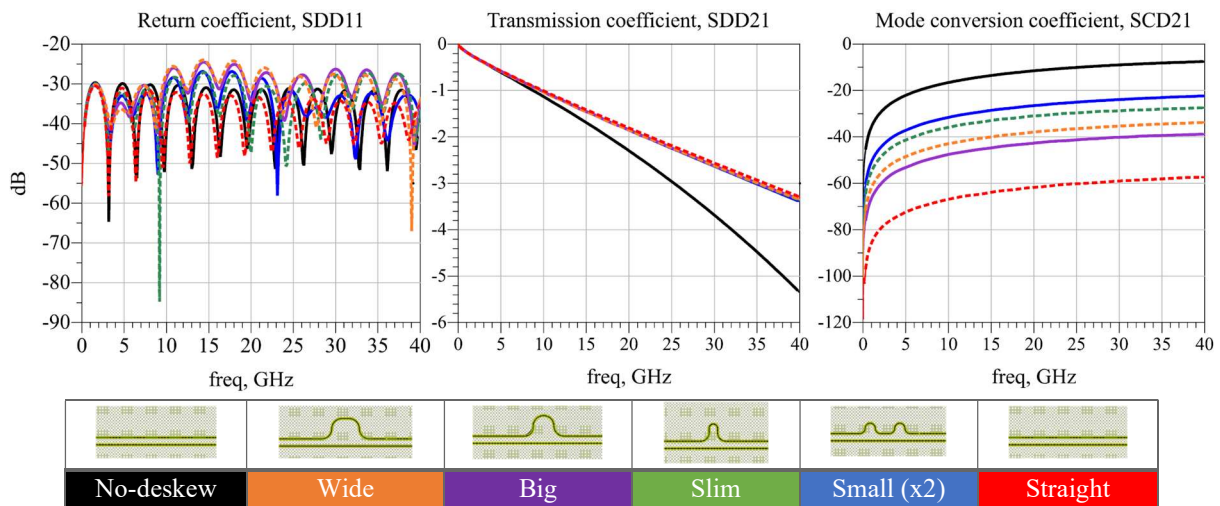
Figure 12: Length equalization Modal S-parameters

2.1.2   Time domain analysis

The time domain results help visualize the cause and effect of each discontinuity. Starting with the TDR plot at Port1, the smaller meanders (Slim and Small) are less inductive than the larger ones (Big and Wide) because they have smaller spikes. The **No-deskew** and Straight features do not have any spike because there is no discontinuity besides the turns which do not cause a significant impedance change.

Continuing with the time domain analysis, Figure 13 shows the contour plot (at BER 10^{-6}) of the lower eye measured at the channel, just before the receiver package. It is important to measure before, because the receiver modifies the waveform using equalization thus the discontinuity effects are hidden. Notice that the contours are rough because at BER 10^{-6} the eye is almost closed at the channel output (the largest eye height is only 6 mV) and there are only a few transitions captured at such low BER levels.

First, notice that the **No-deskew** feature has a very closed eye at 1.6 ps while the Straight feature has the largest eye at 5.4 ps. The 4 different bump sizes fall in between and the have

practically the same eye height of 11 mV. The width of the Big and Wide is the same at 4.6 ps followed by Small and Slim at 5 ps. One can conclude after the eye contour analysis, that skew needs to be compensated as it severely impacts the eye but the type of length equalization used is indifferent.
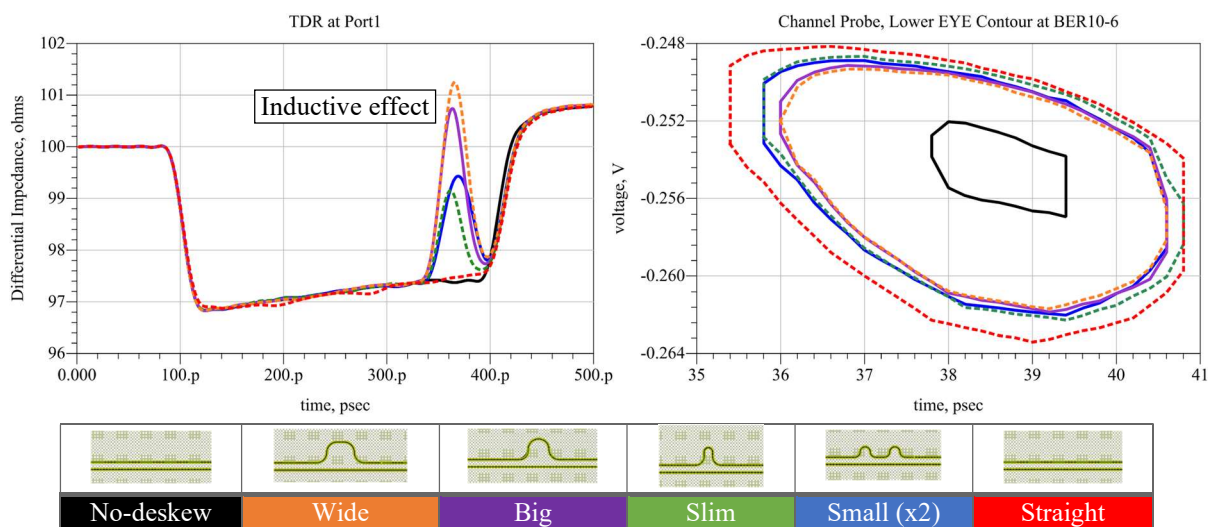
Figure 13: TDR at Port1 and Channel Lower EYE Contour at BER10-6

2.1.3   Conclusions

Then the question remains, which of the features tested is the best option? To answer this question let us look at Table 2 with a summary of the results discussed. The results are normalized, with 0% being the worst performing feature and 100% being the best. As mentioned previously, Mode conversion SCD21 and Eye contour are the most significant tools because they illustrate the effect of the discontinuities on the receiver (eye opening) and on the board (in from of EMI). Since the eye differences are very small, the eye opening could be ignored. It can be concluded that Big size is the best (highlighted in bold). Clearly there is room for improvement with better length equalization methods towards the Straight performance.
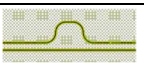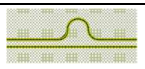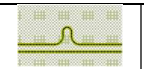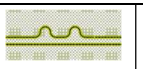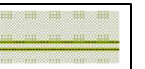
	No-deskew	Wide	Big	Slim	Small (x2)	Straight
SDD11	100%	2%	**0%**	41%	59%	100%
SDD21	0%	100%	**100%**	100%	100%	100%
SCD21	0%	58%	**69%**	42%	33%	100%
TDR	100%	0%	**13%**	55%	47%	100%
EYE	0%	66%	**68%**	81%	74%	100%

Table 2: Length equalization results summary

2.1.4   Validation

Commencing the frequency domain measurements, Figure 14 is commented by columns from left to right, always compering the top and bottom plots which correspond to simulations and measurements, respectively. Starting with the *Return coefficient*, the measurement results have resonances that were expected to be periodic and equally spaced as the simulation but because the de-embedding is not perfect, there are unknown resonances that distort the results. Nevertheless, the average return loss is around 25dB which is very close to the 30dB from simulation.

<div align="right">Marc Borrell</div>

Insertion loss (SDD21) is similar between simulation and measurement. It could be argued that the difference is because surface roughness is not considered in the simulations although since the channel studied is short, the effect is small. Notice that the impact of the skew in the **No-deskew** is very well characterized by the simulation.

Mode conversion also shows promising results. It is most obvious with the **No-deskew** case as it reaches the same mode conversion levels whereas all other features are less distinguished. Among the 4 different sizes of bumps, the Big size is the one that has less mode conversion. All other cases are practically the same with 10dB more than the Big.

Unfortunately, the Straight case was not manufactured. For this reason, the results presented are multiple straight references from other features that, after cascading the measured S-parameters, equal to the Straight case length.
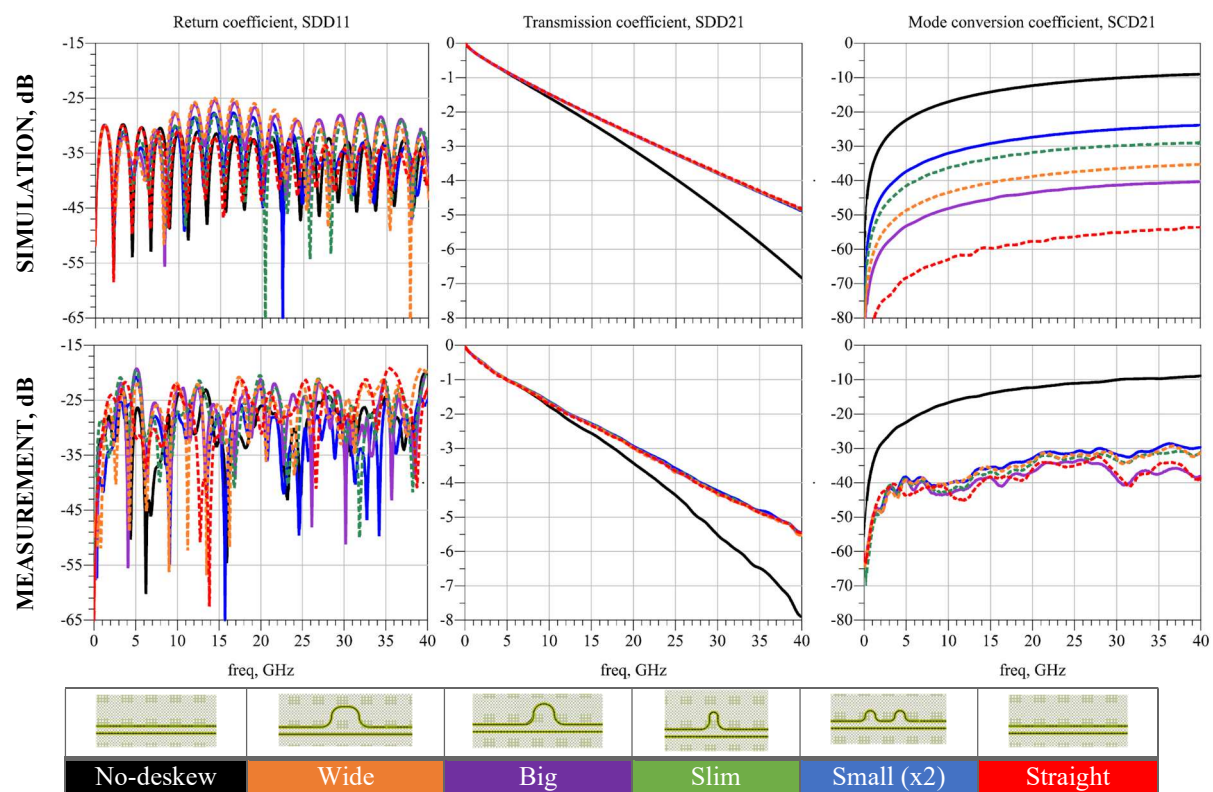
Figure 14: Length equalization S-parameters [Simulation -vs- Measurement]

On the other hand, time domain results confirm that each size has a distinct impact on the characteristic impedance of the line (see Figure 15). It is challenging to appreciate the magnitude of each feature because they do not have the same starting impedance. Also, the measurements do not have a Straight case because it was not manufactured.

The impedance could be different for a combination of reasons. First, the TDR measurements captured with the VNA are a transformation of the frequency domain which is less accurate and introduces mathematical artefacts. In addition, the manufacturing process might create small geometric disturbances.
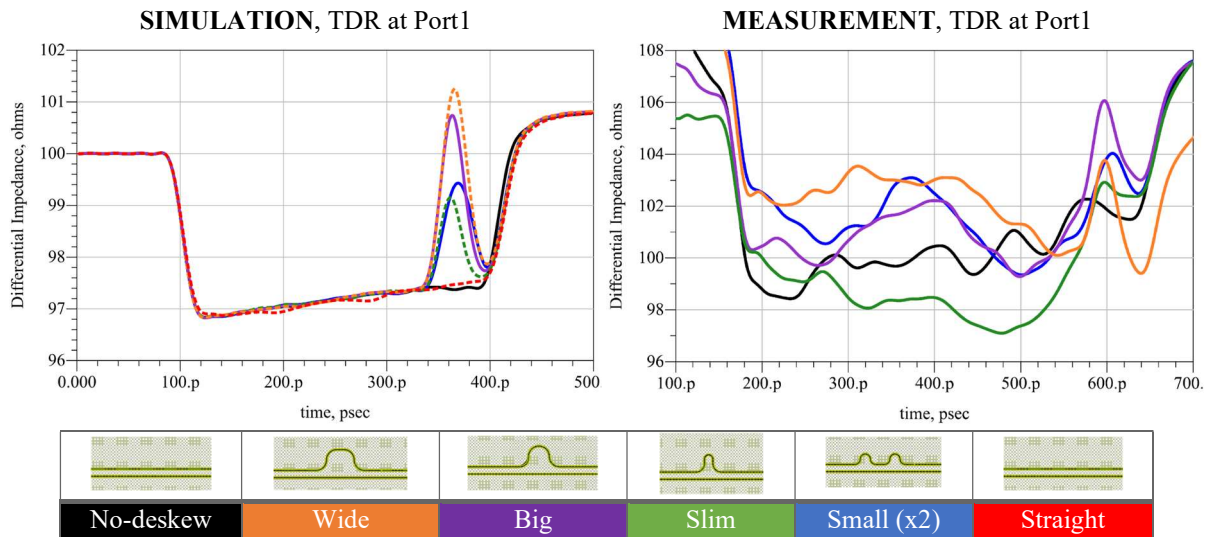
Figure 15: Length equalization, TDR at Port1 [Simulation -vs- Measurement]

If the ends of the spikes are matched together manually only by adjusting the height of each trace, the relative magnitude of the spikes can be compared more easily (see Figure 16). Notice that the Wide case has the largest inductance followed by the Big and the Small x2 cases. Even though Slim also shows a small spike, it is not as significant as expected from simulations. The **No-deskew** reference has large oscillations that make it unreliable.
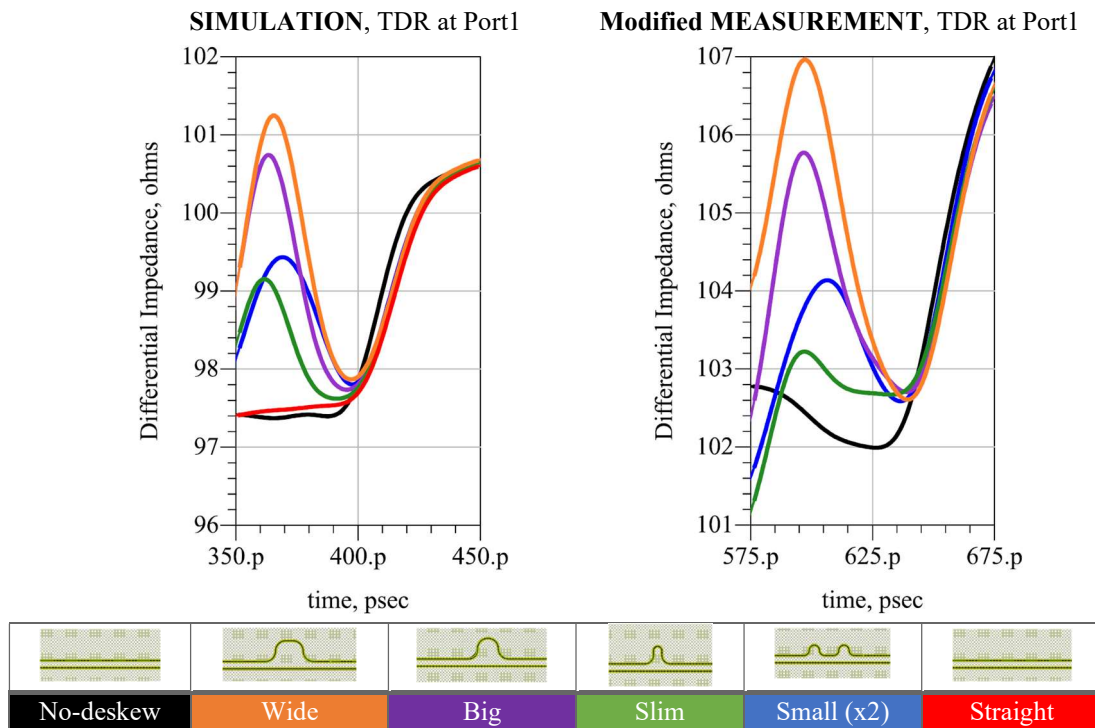
Figure 16: Manually adjusted measurements for comparison

Marc Borrell

2.2 VIA TYPES MODELING

Via transitions is yet another feature that is almost unavoidable when doing PCB layout. For example, a signal coming out from a device package needs to transition from the top layer to layer two so it can become a stripline. Then a basic question arises, is the via disturbance going to block the communication between devices?

Via modeling is a study specific to the stackup used, which depends on manufacturing capabilities and budget. This project uses the stackup provided by Napatech and has the following characteristics regarding via transitions. There are microvias from layer Top to layer 2 (L2) and from layer 2 to layer 3 (L3) as well as from layer Bottom to layer 13 (L13) and from layer 13 to layer 12 (L12). There is also available a buried via that connects L3 with L12. Through hole vias are also possible. This stackup configuration allows to transition to layer 2, 3 ,12 or 13 without stubs nor back-drilling (see Figure 17).
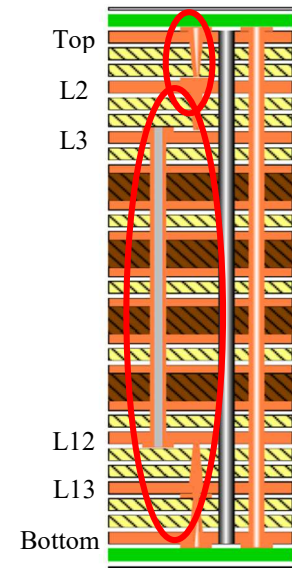
Figure 17: Napatech stackup

L2 and L13 are designed for 100 ohm edge coupled differential striplines. Since components are placed in the Top layer, the most critical via transitions happen between Top and L2 as well as Top to L13 (which is symmetrically equivalent to L2-Bottom). Considering the construction of the test board and its connectors, the connections to the test features is always done through L2.

Then the via structures studied are classified in two groups: "TOP-L2" (with cold color palette) and "L2-Bottom" (with warm color palette). For each, three distinct cases are created to test multiple scenarios. The first case is one commonly used by Napatech engineers with symmetric vias and open plane covers. Inspired by a paper that studied via pattern designs [8], the second case removes the ground vias from one side to test an asymmetric configuration which might be required in some layout situations. Finally, the third case closes the planes above and below the via structure (see Figure 18 and Figure 19).

| Top-L2 Reference | Top-L2 Asymmetric | Top-L2 Cover | Side section view |
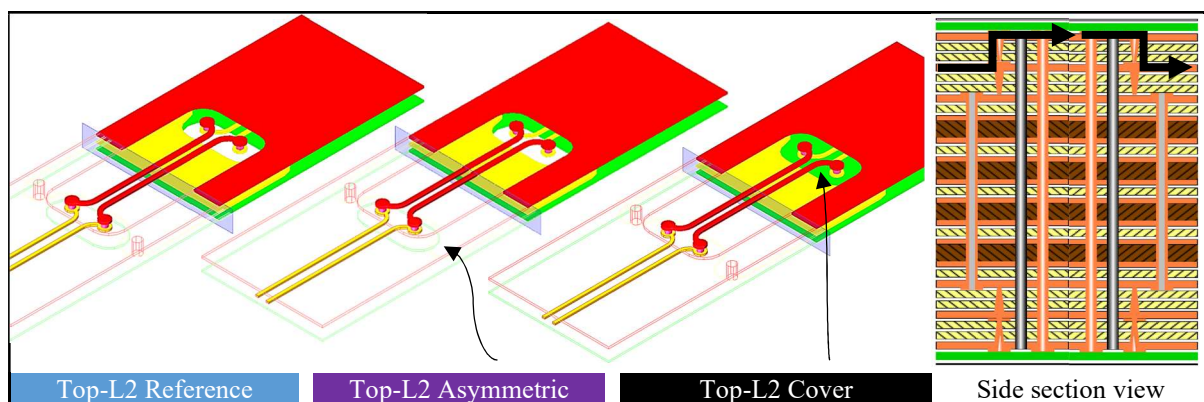
Figure 18: Top-L2 via features
(follow the arrow to see the difference introduced (asymmetry and cover) compared to the Reference case)

Marc Borrell
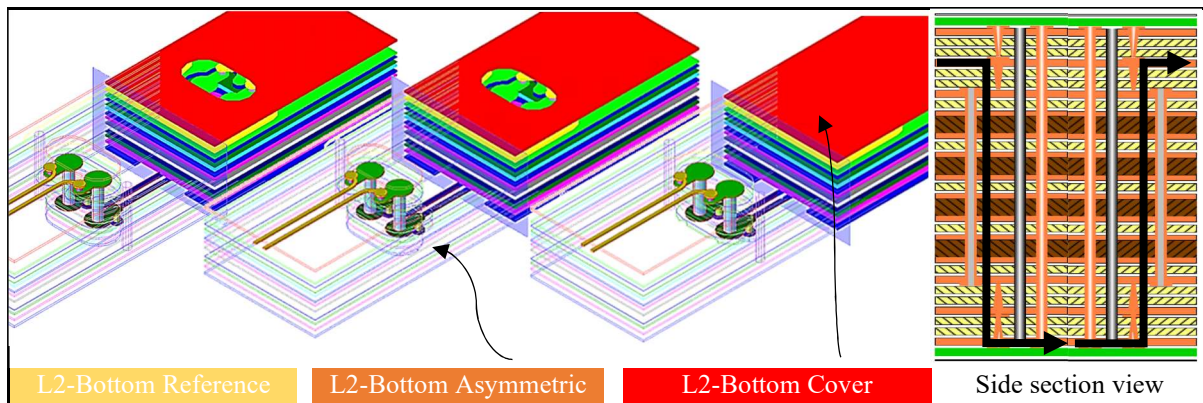
Figure 19: L2-Bottom via features
(follow the arrow to see the difference introduced (asymmetry and cover) compared to the Reference case)

2.2.1   Frequency domain analysis

After running the full 3D EM simulations, the modal S-Parameters show:

<u>Transmission coefficient, SDD21:</u> First of all, a microstrip is used to have a reference of the expected loss for these via transitions. In these cases, the expected loss for a microstrip is 1.1 dB@ 40 GHz both for Top-L2 (length 9.3 mm) and L2-Bottom (length 9.7 mm).

After simulating the six features with a 3D EM solver (see Figure 20), the first observation is that the second group (L2-Bottom) deviates significantly more from the reference insertion loss, almost 2 dB if the oscillations are omitted (see Figure 20). These oscillations are created by resonances of the capacitive and inductive parts of the vias. Group L2-Bottom has 2 resonance frequencies that when excited, severely block the signal, one at 14 GHz and one around 40 GHz. Group TOP-L2 has small resonances at 22 GHz and at higher frequencies than 40 GHz.

Notice that the effect of the cover planes in the "Cover" features has different effect for each group. In case of group 1, the capacitance kills the resonances which indicates that these were mostly inductive and are compensated by the added capacitive effect. On the other hand, group 2 is worst with the added capacitance indicating that the via transition was already capacitive. This is later verified in the time domain with the TDR analysis.

<u>Return coefficient, SDD11:</u> It shows four distinct patterns (see Figure 20):

1. Top-L2 Reference & Top-L2 Asymmetric
2. **Top-L2 Cover**
3. L2-Bottom Reference & L2-Bottom Asymmetric
4. L2-Bottom Cover

In all cases, the reference and asymmetric cases have the same resonant frequencies whereas the "Cover" cases have the center resonance frequencies shifted due to the capacitive effect mentioned before. The Top-L2 group has clearly fewer reflections and the overall best feature in terms of SDD11 is **Top-L2 Cover.**

<u>Mode conversion coefficient, SCD21:</u> Via transitions should be symmetrical along the signal path but asymmetric cases are also studied because it is not always possible to have symmetrical ground vias. This asymmetry in the return path creates differential to common

mode conversion. In particular, there is more mode conversion in the L2-Bottom Asymmetric case that in the other cases because the return path is longer which makes the effects more significant (see Figure 20).

It is important to mention that group Top-L2 is simulated using a simplified stackup with only 3 layers which reduces the mesh size. In comparison, group L2-Bottom has a denser mesh which improves the accuracy of SCD21 hence reducing the lower noise threshold (around 80 dB instead of the 50 dB for group Top-L2). The trade-off is the simulation time and memory used.

| Top-L2 Reference | Top-L2 Asymmetric | Top-L2 Cover |
| L2-Bottom Reference | L2-Bottom Asymmetric | L2-Bottom Cover |
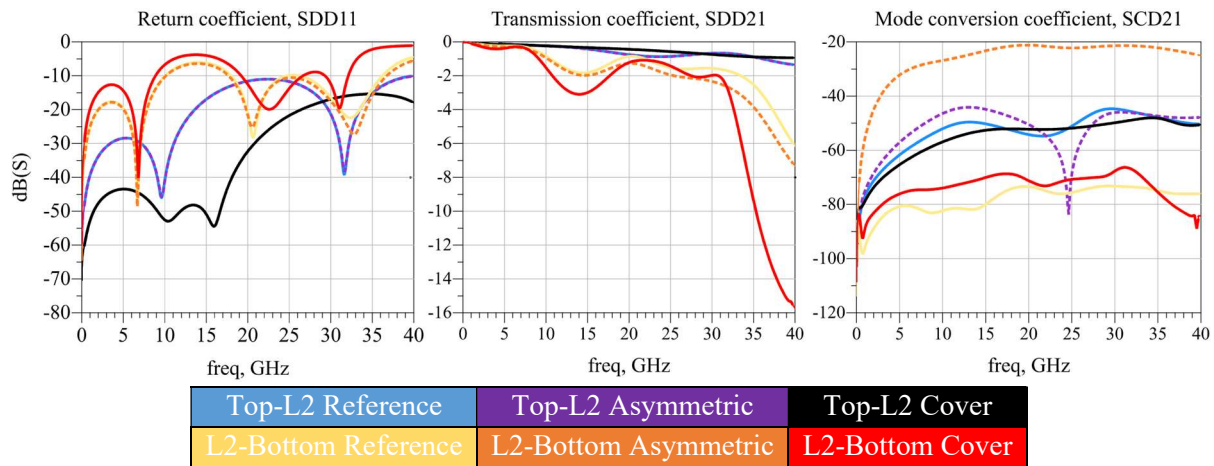
Figure 20: Via features Modal S-parameter Results

2.2.2   Time domain analysis

In the frequency domain study, there are some resonances in the via structure that are inductive or capacitive depending on the feature. The best tool to clarify the cause of the resonances is the TDR plot (see Figure 21). It shows that Top-L2 group are mostly inductive and the addition of covers helps reduce the spikes by introducing additional capacitance as concluded earlier.

It can be observed that the TDR of group L2-Bottom is longer because the via transition is 4 mm longer and the disturbances are larger which take more time to recover.

Finally, the contour plot (at BER 10^{-6}) of the lower eye measured at the channel is larger than the previously studied length equalization features (deskew). Comparing the maximum eye opening and the corresponding length of each type of feature (**deskew**: 6 mV, 5.4 ps, 24.14 mm // **vias**: 80 mV, 16 ps, 9.3 mm) it is clear that the losses of the deskew features, due to channel length, close the eye significantly at the channel output.

Focusing now on the via structures, the eye contour shows that L2-Bottom Reference is the best although Top-L2 Cover also performs adequately. In this Top-L2 group, the Reference and Asymmetric case are identical and have an elongated shape which is not desired as it can be a sign of ISI. By adding the Cover opens the eye height by 10mV. On the other hand, L2-Bottom group has significant differences between its variations. When using Asymmetric ground vias, the eye height is reduced by 10 mV with respect to the Reference. Also, if the Cover is added compared to the Reference, the eye closes severely.

Top-L2 Reference	Top-L2 Asymmetric	Top-L2 Cover
L2-Bottom Reference	L2-Bottom Asymmetric	L2-Bottom Cover
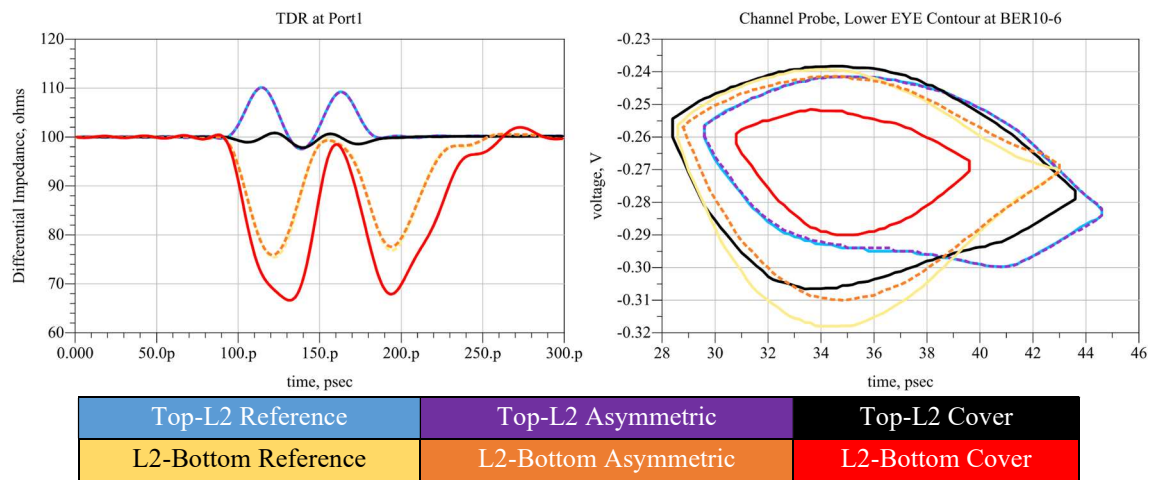
Figure 21: Via features - Time domain Results

2.2.3   Conclusions

To recap all results into a few layout guidelines, SCD21 shows that L2-Bottom Asymmetric ground vias can cause significant noise in EMI. Even though the eye contour shows good performance in most cases, the TDR shows that L2-Bottom cases are not desirable as they have large impedance discontinuities that could cause reflections. Yet, both Top-L2 and L2-Bottom transitions are viable for PAM4 interconnects, especially with **Top-L2 Cover** and L2-Bottom Reference (see Table 3 highlighted in bold). L2-Bottom would benefit from few optimizations that compensate the large capacitive and inductive effects. For example, to lower the capacitive effect, the pad size could be decreased and the antipad increased.

	Top-L2 Reference	Top-L2 Asymmetric	**Top-L2 Cover**	**L2-Bottom Reference**	L2-Bottom Asymmetric	L2-Bottom Cover
SDD11	52%	52%	**100%**	**19%**	24%	0%
SDD21	98%	98%	**100%**	**46%**	41%	0%
SCD21	45%	35%	**51%**	**100%**	**0%**	82%
TDR	62%	62%	**91%**	**30%**	31%	0%
EYE	81%	77%	**100%**	**90%**	78%	**0%**

Table 3: Via types, Results summary

2.2.4   Validation

The frequency domain measurements captured similar effects observed in simulations. Starting with the *Return coefficient* of the L2-Bottom features (warmer colors), results match well at lower frequencies and deviate significantly at higher frequencies. In particular, the center frequency of the highest resonance found in simulation appears between 30-35 GHz while in measurements appears at 40 GHz and higher.

On the other hand, the Top-L2 features (colder colors) have very similar main resonances but also contain unexpected smaller resonances. Especially in the **Top-L2 Cover** case, the 3D FEM solver has overestimated the capacitive effect of the cover plane.

In line with the return coefficient, the *Transmission coefficient* shows that the long via transitions from L2-Bottom, have not been characterized properly in simulation. For example, the severity of the resonance at 15 GHz is twice as large in reality, while the resonance found at 40 GHz is half of what is predicted in simulation. The small transitions from Top-L2 are much more accurate and the small resonances found in SDD11 barely have a difference between simulation and measurements.

Marc Borrell

Finally, *SCD21 measurements* confirm that L2-Bottom Asymmetric is the most susceptible to mode conversion. While the simulation suggests that L2-Bottom transition have better mode conversion than Top-L2, the measurements indicate the opposite. As explained before, simulations noise floor depends directly on the mesh used which could explain the difference. In any case, the difference is not enough to consider Top-L2 better in terms of mode conversion.
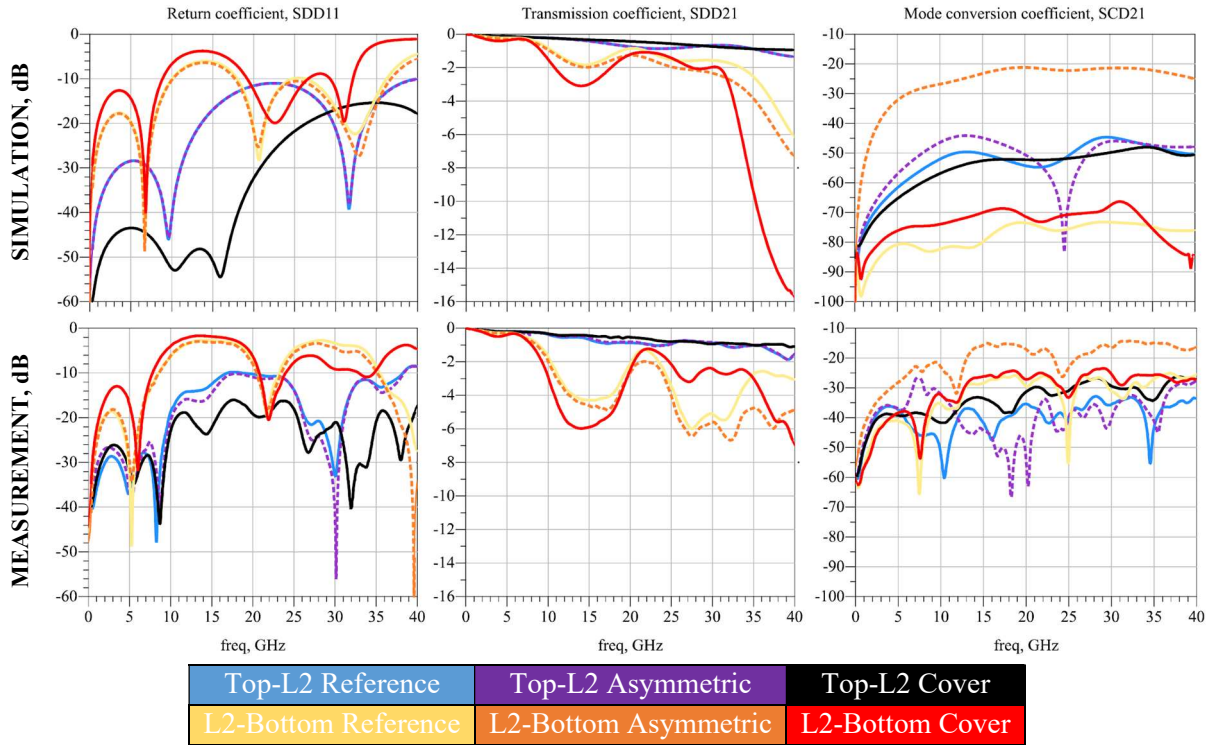
Figure 22: Via features S-parameters [Simulation -vs- Measurement]

Time domain measurements confirm that the transitions from Top-L2 are inductive and the L2-Bottom are very capacitive (see Figure 23). Among the Top-L2 features, the best is the **Top-L2 Cover** case because it transitions smoothly from the 110 ohm connector impedance to the 100 ohm impedance of the center microstrip. On the other hand, the L2-Bottom features do not reach the same microstrip impedance at the center yet the impedance drop at the via is similar to the simulation results.
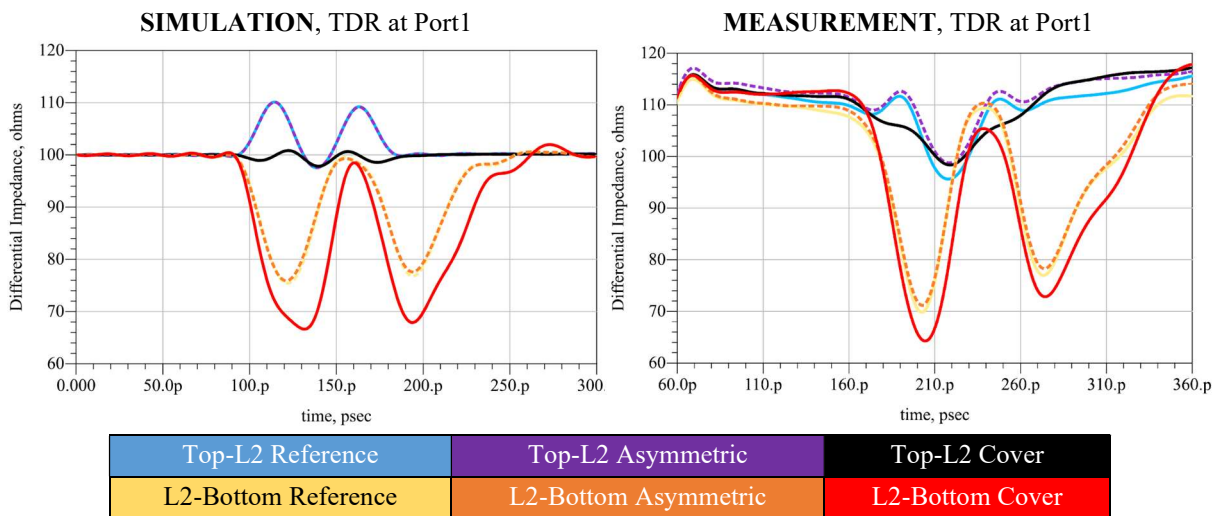
Figure 23: Via features, TDR at Port1 [Simulation -vs- Measurement]

Marc Borrell

2.3   TRACE TO VIA TRANSITION DESIGN

The transition from trace to via is used everywhere in layout design but it is usually not considered when modeling interconnect discontinuities. Nevertheless, this section elaborates a detailed study of 6 different possible transitions to clarify which is the best approach. The complete feature used in simulation and its variants is illustrated in Figure 24. The design of these is original and without previous knowledge that might hint what to expect. The simulations consider only the top 3 layers which is the minimum required to construct a stripline. As always, the input and output port of the features are required to be on layer 2. Then, the via transition to Top layer forces to use a microstrip before going back to L2. Observe that the type of transition from trace to via is repeated 4 times thus making the effect of the transition more noticeable.
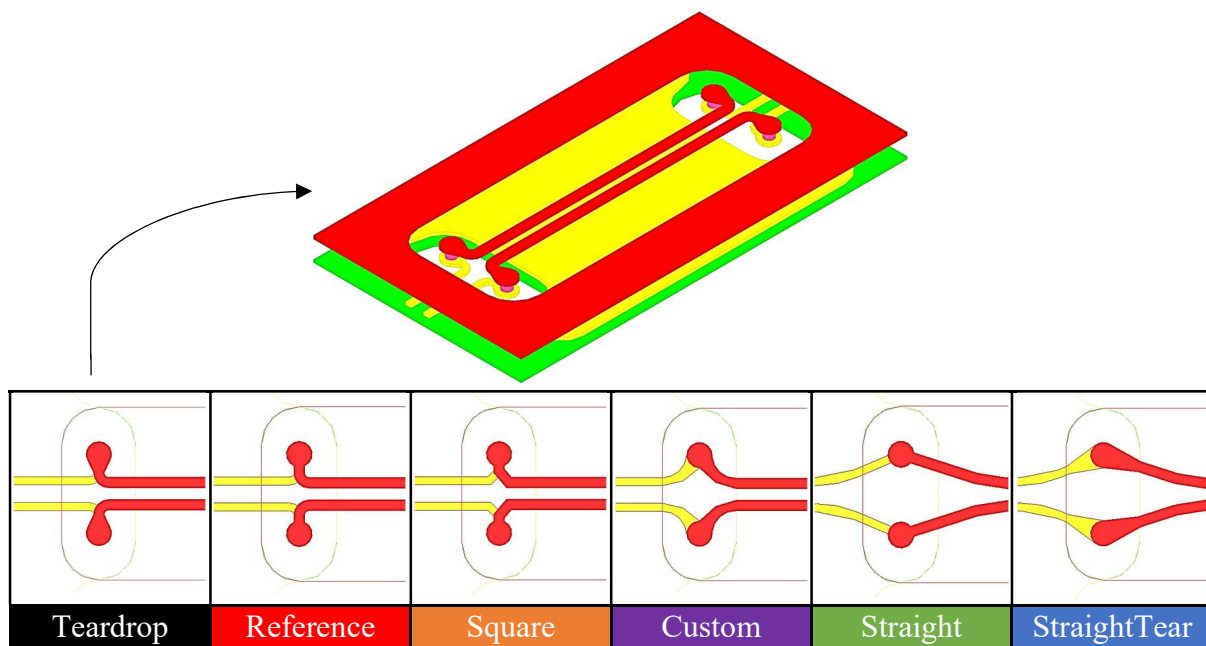
Figure 24: Trace to via features
(the top image illustrates the complete Teardrop feature whereas the bottom images display 6 variations of the same layout)

2.3.1   Frequency domain analysis

After running the full 3D EM simulations, the modal S-Parameters show:

Transmission coefficient, SDD21: Just as explained before in the Via types section, it is expected that the loss of 6.5 mm microstrip to be 0.9 dB@ 40 GHz. Then, the features that are closer to the ideal are the ones that have sharp corners: **Teardrop**, Reference and Square. In general, all cases appear to be resonant at 25 GHz and above 40 GHz. The extreme cases of resonance happen with the Straight and StraightTear. By adding the teardrop to the Straight feature decreases the losses at the resonance frequency. The TDR plot is required to further understand what the nature of these resonances are.

Mode conversion coefficient, SCD21: Note that at low frequencies, the differential to common conversion is very noisy. This is a known limitation of the 3D solver when using modal wave ports and the results should be ignored below 500MHz. To be clear, this type of noise only appears when SCD21 has very low values and frequencies (below -80dB and <500MHz).

Marc Borrell

In general, these types of features should have very low mode conversion because there is no skew. This is corroborated with simulations because all cases are below 50 dB which is the noise floor of the simulation tool. It is for this reason that the results are less accurate. Nonetheless, it seems to be consistent about straight features (blue and green, see Figure 25) having a slight edge compared to the rest.

Return coefficient, SDD11: Finally, the return loss has similar conclusions as with insertion loss. The features with best performance are **Teardrop**, Reference and Square. Then the Custom feature follows by 5 dB and the straight features by almost 10 dB more. All features have the same number of resonances but with slightly different frequency points.
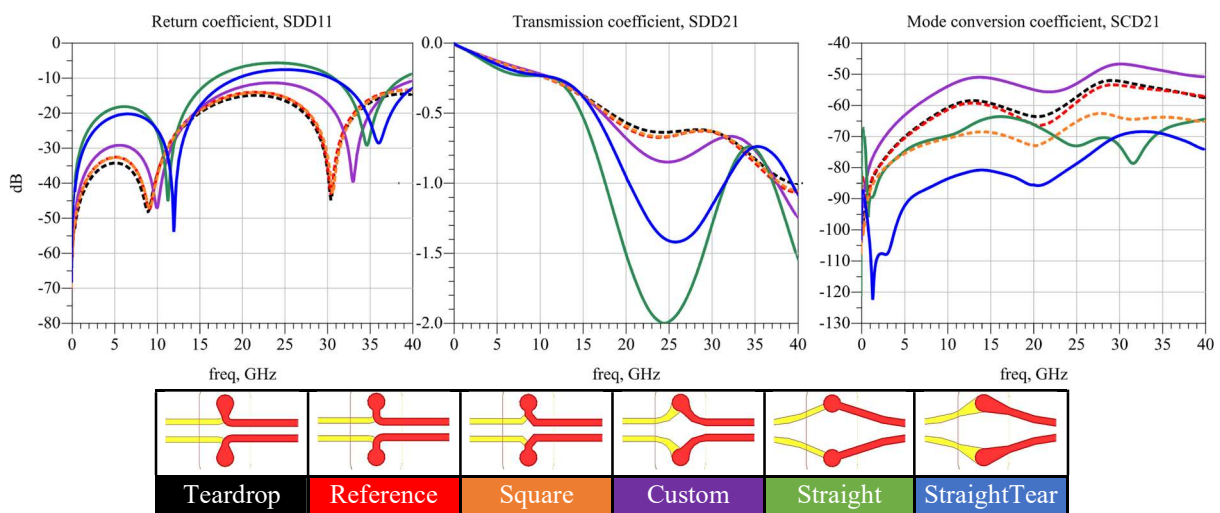
Figure 25: Trace to Via features, Frequency domain results

2.3.2   Time domain analysis

With the time domain results, it is possible to understand which are the causes and effects of these transitions. For example, the TDR plot highlights that Straight and StraightTear are very inductive (see Figure 26) and explains the large resonances that appeared in SDD21. It also corroborates the fact that by adding a teardrop which is mostly capacitive, the inductive effect of the Straight feature is reduced. It is also not a surprise to have a bit of inductance in **Teardrop**, Reference and Square features because SDD21 has some oscillations although compared with "Top-L2 Reference" from "Via Types" section, the TDR spike is reduced by 5 ohm. Then, it is possible to assume that most of the inductance comes from the via transition and that the Via to Trace feature is capacitive enough to reduce the inductance of the via.

On the other hand, the Eye contour shows the effects of each transition. The inductance on Straight and StraightTear causes the eye height to close significantly. This is a common characteristic with inductive features which tend to be more elongated whereas the capacitive features are more circular. **Teardrop**, Reference and Square have the largest opening and are identical. The Custom feature is just 8 mV worse.
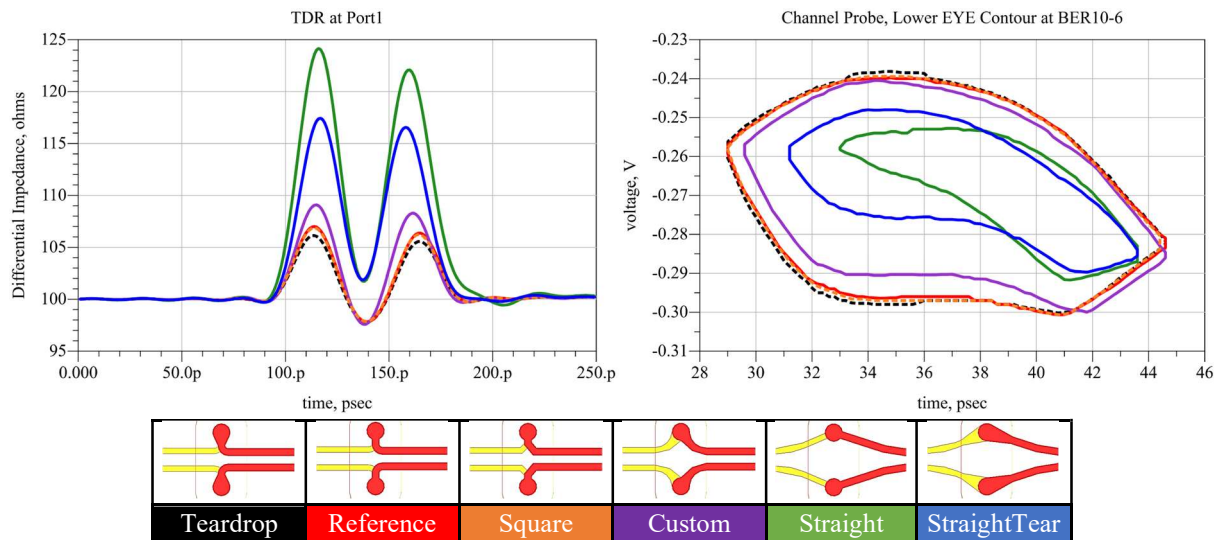
Figure 26: Trace to Via features, Time domain results

2.3.3   Conclusions

Even though via to trace transitions do not seem too critical, one could think that a smoother and straighter trace could ease the transition but instead turn out to be the opposite. The inductance added by breaking the differential pair early with the already inductive via transition ends up having an undesired effect to the eye contour.

Judging by Table 4 and especially the highlighted eye results, Straight and StraightTear transitions should not be used due to its small eye area. As expected, the best transitions to use are **Teardrop**, Reference and Square.

On the other hand, SCD21 cannot be a deciding factor as usual because the results are below the noise floor of the simulator and they strongly depend on mesh density used which can vary just by the number of edges a turn has. In other words, Square, Straight and StraightTear do not have as much edges as **Teardrop**, Reference or Custom meaning that the meshing on the first group is simpler which reduces the accuracy of the solution.

	Teardrop	Reference	Square	Custom	Straight	StraightTear
SDD11	**100%**	**92%**	**93%**	71%	0%	21%
SDD21	**100%**	**98%**	**99%**	90%	0%	47%
SCD21	44%	48%	**100%**	0%	59%	100%
TDR	69%	65%	65%	57%	0%	27%
EYE	**100%**	**97%**	**95%**	86%	0%	19%

Table 4: Trace to via, Results summary

2.3.4   Validation

Starting with the *Return coefficient*, simulation results match precisely with measurements (see Figure 27). Clearly the center frequencies of the two resonances observable are correctly found and the return loss has the same level between 10-20 dB. The small differences can be attributed to the de-embedding pre-process because it affects all results consistently.

Continuing with the *Transmission coefficient*, the loss measured is a bit larger than in simulations. This could be because the de-embedded fixture length is not perfect which would be consistent with the behavior observed in SDD11 but could also be due to the unknown

Marc Borrell

surface roughness losses. Nonetheless the results are very similar and correctly describe the performance of each trace to via transition.

Ending with the *Mode conversion*, at first glance simulations show very low levels of differential to common mode conversion. This happens because the simulation tool has much lower noise level that the VNA. As indicated before, the results at such low levels are often mathematical artefacts. Judging by the measurements, all cases have the same mode conversion but the **Teardrop** case is the best of all.
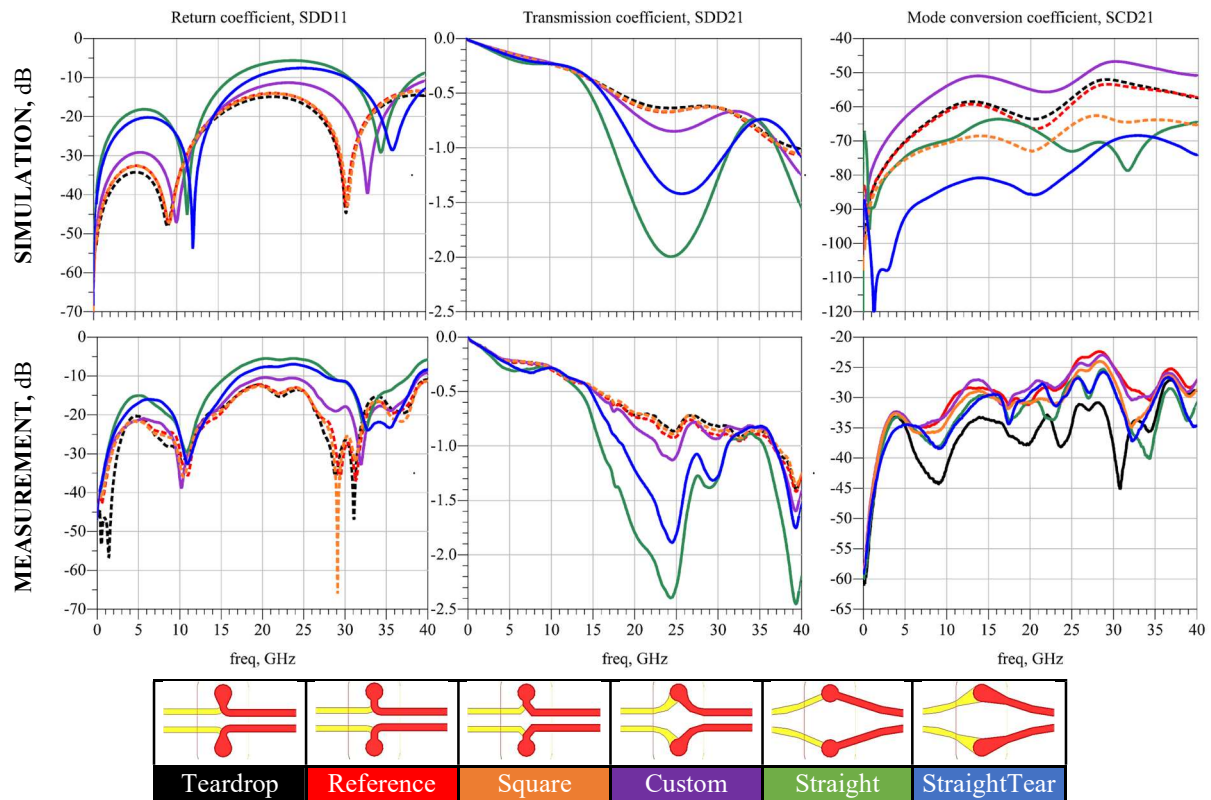
Figure 27: Trace to Via features S-parameters [Simulation -vs- Measurement]

Finally, the TDR plots show good agreement between measurements and simulations (see Figure 28). Notice that the straight cases are very inductive and they are not able to reach the characteristic impedance at the center microstrip.

Only the Straight is accurate at the maximum impedance reached, all the other cases have about 5 ohms more than what the simulations predicted. Although the reference is shifted, the relative magnitude is very similar to simulations and only the **Teardrop** case differs slightly possibly explained by its additional capacitance.
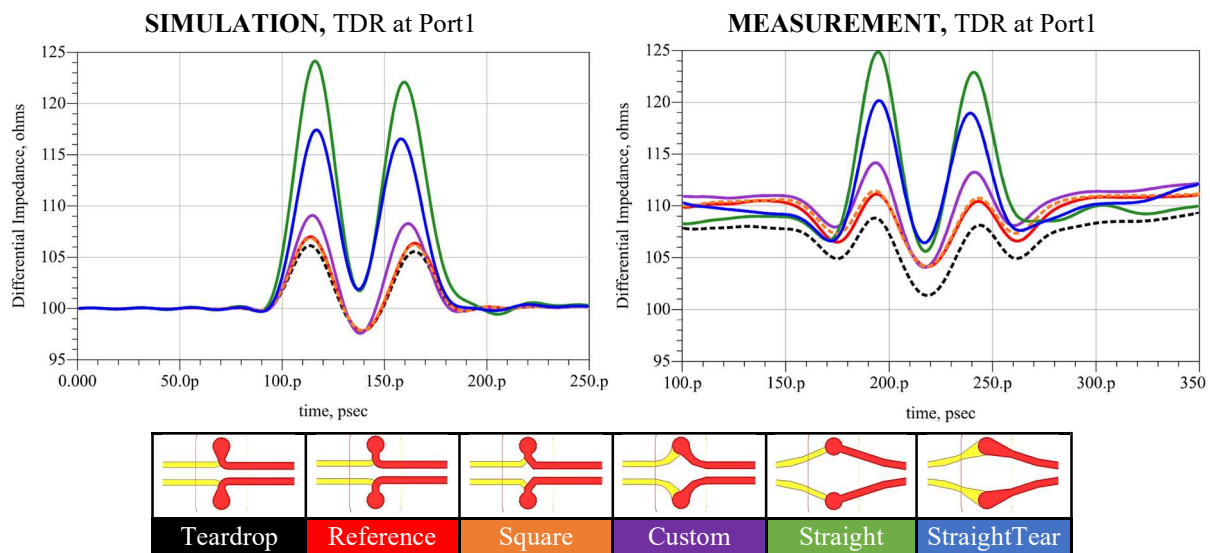
Figure 28: Trace to Via, TDR at Port1 [Simulation -vs- Measurement]

2.4   NEAR-ANTIPAD MODELING

As computing power continues to increase, so does the density of the packages. High density and high-performance packages usually take the form of Ball Grid Array (BGA). These typically have a small pitch between pins, for example 1 mm or 0.7 mm. It is easy to see that just considering the antipads, the number of discontinuities encountered by the traces when exiting the field is significant. This challenge has been previously studied in [9] and this report continues that research by adding new antipad shapes, configurations and analysis.

Multiple near anti-pad features have been designed to demonstrate the various effect that can occur when exiting a pin-field. There are three different grid pitches under test **0.7 mm**, 0.8 mm and the most common 1 mm. There is also an asymmetric antipad feature with 1 mm pitch. Finally, the 0.8 mm pitch antipad is duplicated to simulate the exit from further inside the package footprint. The antipad shape is inspired by the cutouts used in via transitions from L2-Bottom. The via itself is removed from the center of the antipad to avoid floating structures. The final test set is shown in Table 5.

Table 5: Antipad features overview

2.4.1   Frequency domain analysis

It is especially important to characterize these types of discontinuities with a 3D solver otherwise the plane cutouts will not be captured and results would be inaccurate. After running the full 3D EM simulations, the modal S-Parameters show:

Transmission coefficient, SDD21: It shows that none of the antipad features disturbs the insertion loss because all results are identical to the Reference which is just a homogeneous

stripline. The loss is almost linear because as pointed out in the assumptions, the solver does not consider surface roughness (see Figure 29).

Mode conversion coefficient, SCD21: Notice that almost all results fall below 50 dB which is considered to be the noise floor. The only feature that is slightly above 50 dB is the 1mm Asymmetric case. This result is consistent with the fact that asymmetric structures cause more differential to common mode conversion than symmetric ones. Still, the conversion is so low that this effect can be considered benign in most systems (see Figure 29).

Return coefficient, SDD11: Return loss is always expected to have oscillations, even in the ideal case because these come from the mismatch between the transmission line and the output port impedance, for example see the Reference. There are some features that defer heavily from the reference such as the **0.7mm** case. Clearly having plane cutouts near the trace creates additional reflections that are overlapped with the ideal case, thus creating the strange pattern shown in Figure 29. Using the TDR, one can interpret these reflections much better.

| 0.7 mm | 0.8 mm | 0.8 mm x2 | 1 mm Asymmetric | 1 mm | Reference |
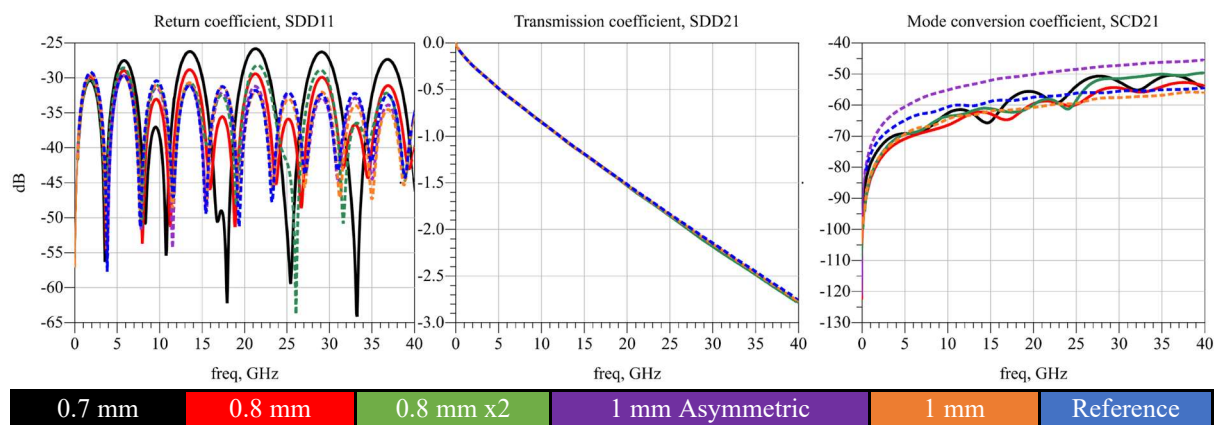
Figure 29: Antipad frequency domain results

2.4.2   Time domain analysis

Among the time domain results, TDR plot is the most significant because it converts the reflections observed in SDD11 into understandable impedance discontinuities that are the cause of those reflections. The Reference is the ideal case where only the losses on the line change the impedance. The 1mm Asymmetric and 1mm cases are almost ideal and there is no difference between the two.

On the other hand, 0.8mm and 0.8mm x2 are slightly inductive and the effect is longer in 0.8mm x2 as it is clearly appearing twice. The worst case is having **0.7mm** pitch because the impedance change is almost 3 times worse than with 0.8mm (see Figure 30).

The eye contour shows that all cases have an almost identical open eye. The small differences could come from the inductance which tends to elongate the eye but the change is minimal and probably not beneficial. It can be concluded that antipad discontinuities do not affect the eye.

Marc Borrell

| 0.7 mm | 0.8 mm | 0.8 mm x2 | 1 mm Asymmetric | 1 mm | Reference |
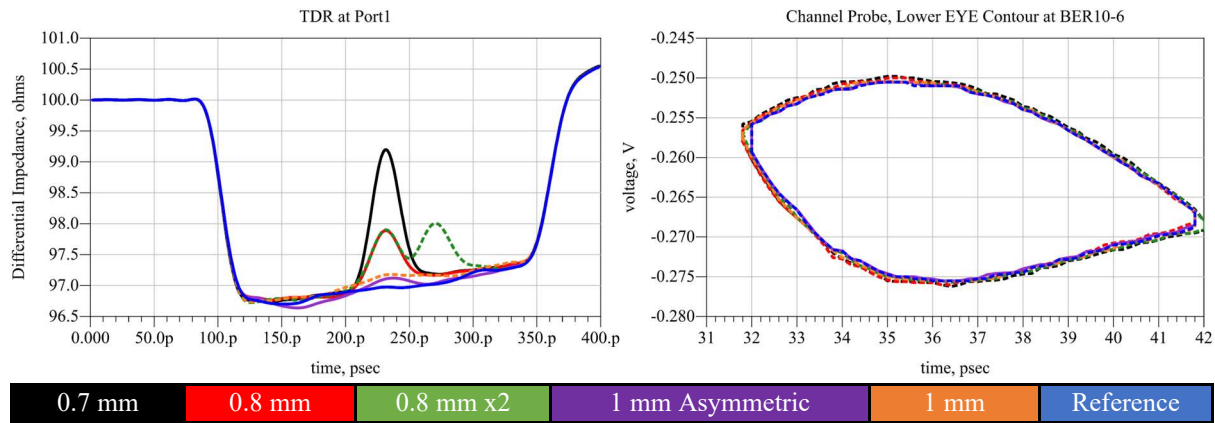
Figure 30: Antipad, Time domain results

2.4.3   Conclusions

Most of the results do not show significant difference between the cases studied. Based on the eye contour results, it is safe to say that any of the antipad configurations would be adequate for a layout. Of course, there is still and advantage when using wider pitches, such as 1 mm (highlighted in bold) because they produce less reflections as shown in SDD11 and TDR plots (see Table 6).

	0.7 mm	0.8 mm	0.8 mm x2	1 mm Asymmetric	1 mm	Reference
SDD11	0%	60%	10%	70%	**100%**	90%
SDD21	31%	62%	0%	69%	**23%**	100%
SCD21	72%	100%	79%	0%	**100%**	74%
TDR	0%	50%	50%	95%	**90%**	100%
EYE	100%	60%	48%	0%	**24%**	0%

Table 6: Antipad features results summary

2.4.4   Validation

Return coefficient measurements are difficult to interpret because the resonances are not periodic as they appear in simulation (see Figure 31). This inconsistency is attributed to small physical imperfections that cause reflections in the real board and to measurement limitations. This is confirmed when looking at the TDR response (see Figure 32). The same could be said about the level reached in simulation of 30 dB whereas measurement is worse with 25 dB.

Following with the *Transmission coefficient*, both behave linearly but the loss of the real case is 1 dB higher (see Figure 31). This could be explained with the same arguments as in the Trace to Via section where the de-embedded fixture length and the surface roughness effect are to blame. It is worth noticing that the measurements indicate that the 0.8 mm x2 feature is slightly worse than the rest which is understandable as it is the most severe case.

The *Mode conversion coefficient* measured differs significantly from simulations (see Figure 31). The simulations predicted that mode conversion would be very low (<50 dB) but the measurements show that the mode conversion is much higher at around 20 dB. This is far from the noise floor of the instrument observed in other measurements of around 40 dB. The difference could be because the medium is not homogeneous.

Marc Borrell

Finally, it can be argued that if measurements were accurate, the asymmetry present in 1 mm Asymmetric would increase the mode conversion over the other cases. Instead, measurement show that the worst mode conversion happens in 0.8 mm x2 followed by 1 mm Asymmetric, 1 mm, Reference and finally **0.7 mm** and 0.8 mm.
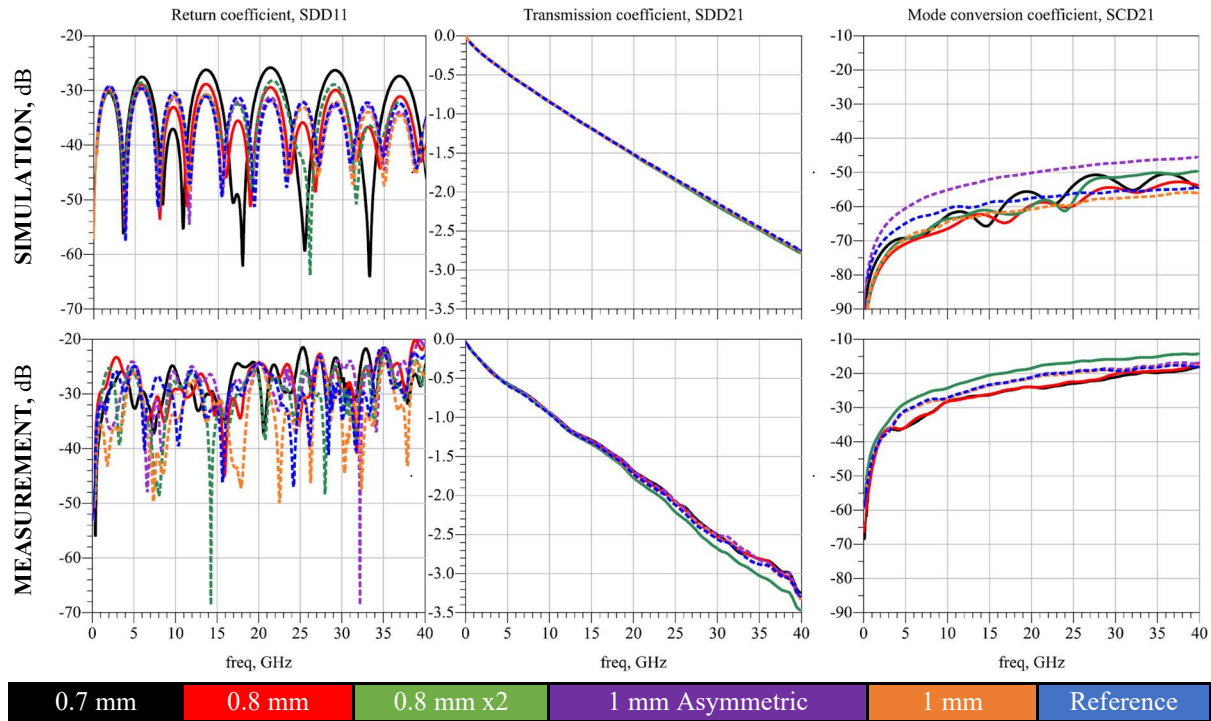
Figure 31: Antipad S-parameters [Simulation -vs- Measurement]

Recalling the simulation results, TDR was the most useful tool to analyze the effect of the antipad discontinuity because these show as an inductive spike. Unfortunately, the measurements do not have the same characteristic impedance neither it is constant (see Figure 32). Luckily three important observations can be extracted. First notice that the **0.7mm case** has a noticeable spike very similar in position and magnitude to the simulation. Case 0.8 mm x2 also has two visible bumps at 300 ps and 350 ps. Finally, only these two cases which are considered to be the worst are above the Reference. This could indicate that all other cases do not have a noticeable effect on the transmission line.
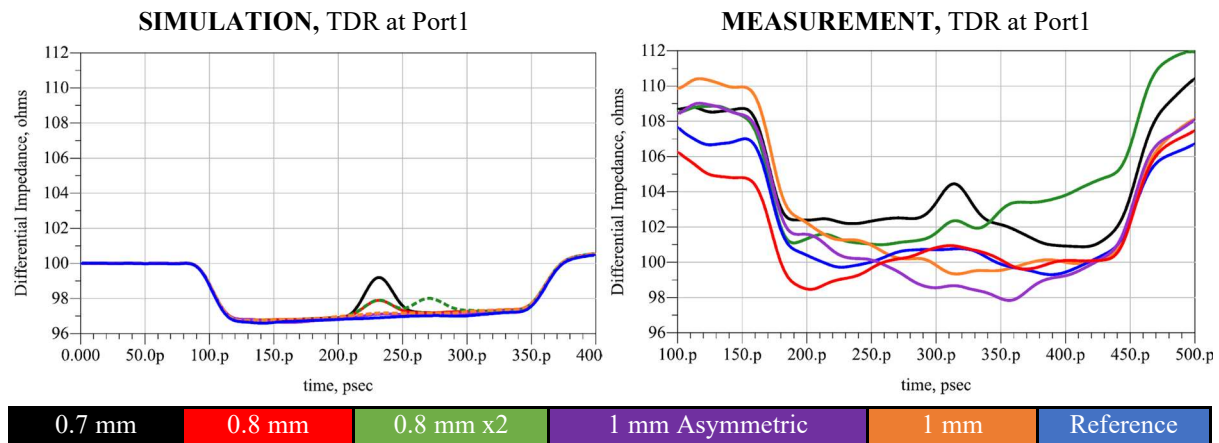
Figure 32: Antipad TDR at Port1 [Simulation -vs- Measurement]

Marc Borrell

2.5   PIN-FIELD ESCAPE DESIGN

Pin-field escape occurs at the pinout of an IC package usually in the form of BGA. One of the challenges found in pin-field escape design is tight spacing between antipads and traces which has already been studied in the previous section. Another common issue is finding the best method to route a stripline out of a pin-field when a 90-degree turn is forced.
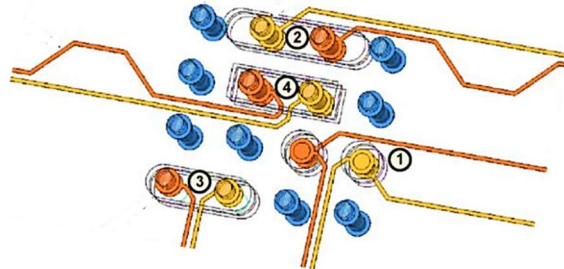
Figure 33: Pin-field escape examples from [10]
1) Inline, 2) Inline flank, 3) Broadside, 4) Broadside Turn

Inspired by previous Napatech designs as well as [10] (see Figure 33), three escape methods are tested: Broadside Turn, Inline Flank and **Inline**. In terms of layout, each one has its pros and cons. Inline Flank is very compact in width but it requires the most length equalization at the output (7.1ps), whereas Broadside Turn uses the same transition as studied in the trace to via transitions but introduces a 90-degree turn that needs less equalization (2.5ps). Finally, **Inline** does not require equalization of the output transmission line but it uses more space. Napatech's implementations of the **Inline** escape tries to re-unite the differential pair as soon as possible which ends up creating a coil-like shape (see Figure 34). It is worth mentioning that the "Broadside" case from Figure 33 is omitted because it has already been studied in the "Trace to Via" section and all other features have the disadvantage of having a 90 degree turn which would make the comparison unfair.

The features created for study are symmetric by design so that no equalization is required for any of the cases, thus isolating only the desired discontinuities for analysis. Since the signals need to come and exit through layer 2, the layout is created in such a way so that the same escape method is used in all transitions. Also, similar to the trace to via transitions, each feature includes 4 times the same escape method (see Figure 34). The via transitions both for signal and ground are spaced at 1mm pitch in a grid pattern to emulate a real pin-field escape scenario.

| Broadside Turn | Inline Flank | Inline |
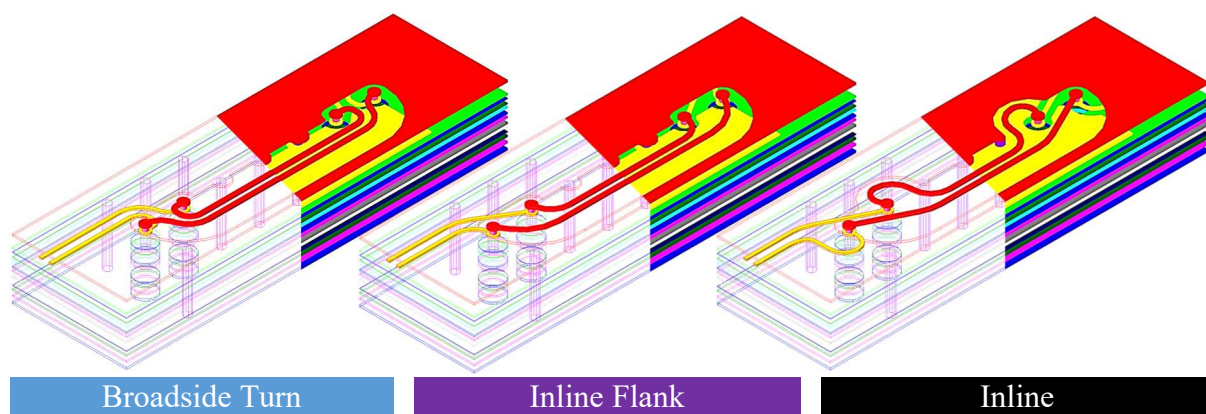
Figure 34: Pin-field escape features

Marc Borrell

2.5.1   Frequency domain analysis

After running the full 3D EM simulations, the modal S-Parameters show:

Transmission coefficient, SDD21: Out of the three features, Broadside Turn has almost ideal insertion loss as it does not have oscillations. Inline Flank has a small resonance at 16 GHz and above 40 GHz which creates almost doble the loss at 40 GHz. **Inline** has the worst SDD21 of all due to a large resonance above 40 GHz (see Figure 35). Even though the Nyquist frequency of a 50 Gbps PAM4 signal is 12.5 GHz, it is not desirable to have large resonances as it can filter out the high frequency spectrum and create ISI.

Mode conversion coefficient, SCD21: This time, the results of mode conversion are much more relevant than in other sections because the results are above 50 dB. Broadside Turn has the lowest common to differential mode conversion, followed by Inline Flank and **Inline** (see Figure 35). At first glance, this is not expected as **Inline** is the only feature that does not have any skew throughout the path. One reasonable explanation could be in the return path of the signal.

Return coefficient, SDD11: The return coefficient plot is not periodic, it has irregular oscillations that come from all the discontinuities in the path. Again, Broadside Turn has the best results followed by Inline Flank and **Inline** at 10 dB higher (see Figure 35).
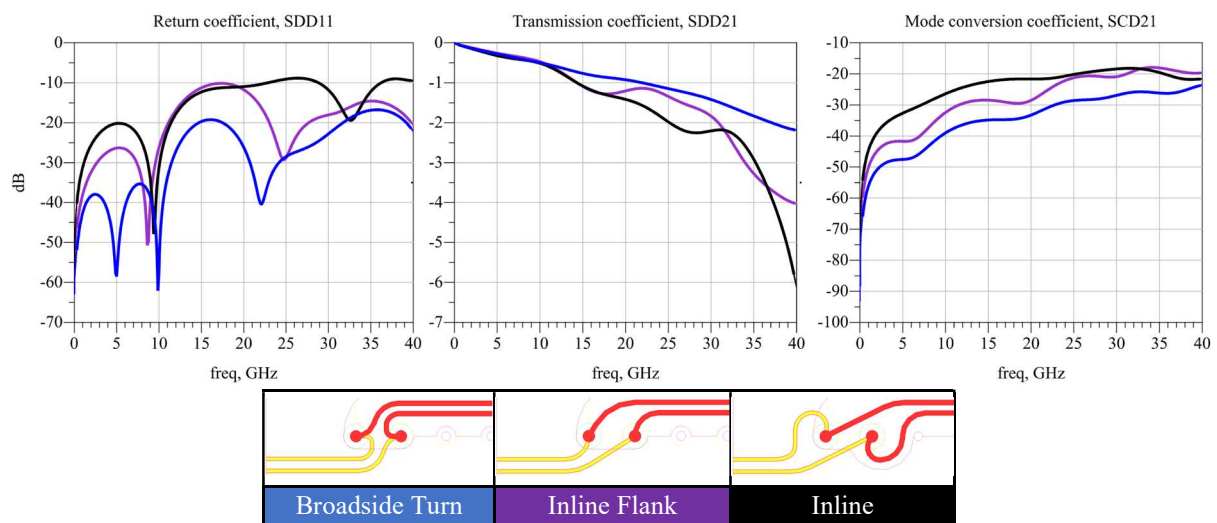
Figure 35: Pin-field escape frequency domain results

2.5.2   Time domain analysis

Starting the analysis with the TDR plot in Figure 36, the results explain the effects of the different geometries used. Broadside Turn and Inline Flank transitions are similar to previous features and the results correlate as expected. In the case of Broadside Turn, the transition is slightly inductive and most of this effect is coming from the via barrel. In comparison, Inline Flank is significantly more inductive because the differential pair is split earlier. Similarly, the **Inline** feature is very inductive but this time it is because the geometry creates a loop.
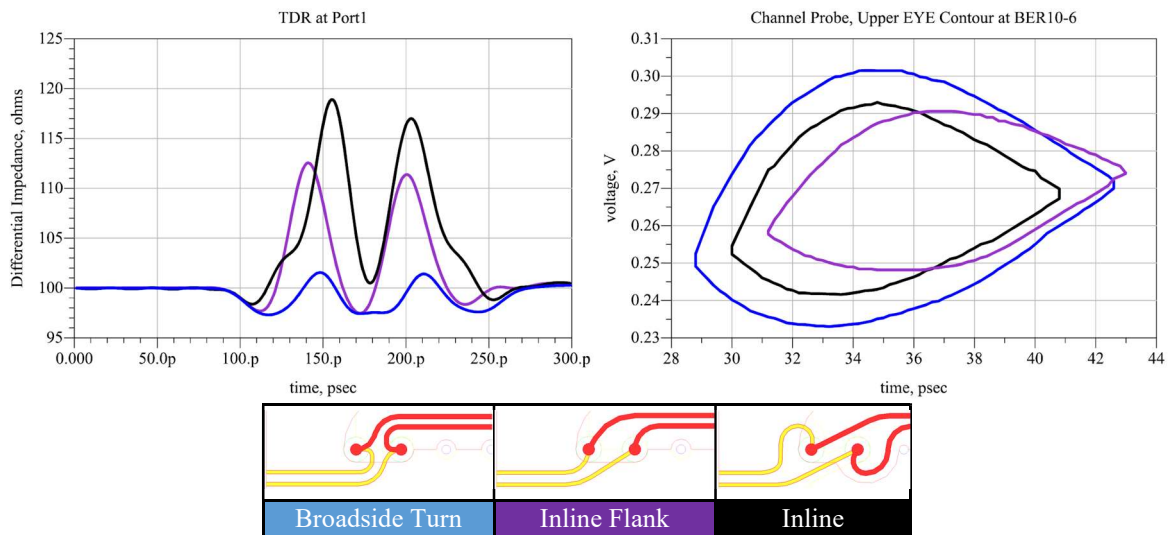
Figure 36: Pin-field escape, time domain results

Compared to previous eye contour plots, this one is much more difficult to evaluate. It is clear that Broadside Turn has the largest eye opening which indicates it is the best performer. Then, between **Inline** and Inline Flank one should not be misled by the eye opening without considering how the eye is measured following the 56G VSR PAM4 specification [1]. The centers of the eyes are found at around 38 ps (calculated at BER 10^{-3}). Then V_{upp} and H_{upp} are measured as explained in part 1.5.3. The final measurements shown in Table 7 point out that the second-best opening is Inline Flank followed by **Inline**.

	Broadside Turn	Inline Flank	Inline
T_{mid}, ps	38.2	38.8	38.0
V_{upp}, V	0.048	0.034	0.03
H_{upp}, ps	12.8	10.2	10.2
TDR, ohm	101.5	112.5	118.8
Height, (EH6), V	0.048	0.034	0.030
Width (EW6), %UI	0.320	0.255	0.255
Eye linearity (EL), %	0.990	0.995	0.981

Table 7: TDR and Eye Height, Width and linearity results

2.5.3   Conclusions

In conclusion, the pin-field escape features characterized have distinct results throughout all the measurements done (see Table 8). Out of the 3 cases, it is best to use Broadside Turn if possible because the lower inductance improves return and insertion loss. Mode conversion is also lower in this case and the eye opening is the largest. If not possible, use Inline Flank instead. Do not use **Inline** as it has large inductive effects that cause performance issues both in mode conversion and eye opening.

	Broadside Turn	Inline Flank	Inline
Mean SDD11, dB	30.6	23	17.6
Mean SDD21, dB	0.9	1.3	1.7
Mean SCD21, dB	37.3	31.4	26.9
TDR spike, ohm	101.5	112.5	118.8
Height, (EH6), V	0.048	0.034	0.030
Width (EW6), %UI	0.320	0.255	0.255
Eye linearity (EL), %	0.990	0.995	0.981

Table 8: Pin-field escape results summary

<div align="right">Marc Borrell</div>

2.5.4   Validation

Starting with the *Return coefficient* measurements, there is good agreement with the simulation predictions (see Figure 37). In particular, both Inline Flank and **Inline** have the best correlation with simulations as they reach similar return loss levels throughout the frequency range and the resonances appear at the same frequency points. Only Broadside Turn deviates from the simulations as it has more reflections. Notice that the Broadside Turn results are like the **Top-L2 Cover** results from Via transitions section. In both cases the simulation predicts a very low return loss but this is not captured properly by the measurements because of other reflections, possibly from the connector transitions.

Continuing with the *Transmission coefficient,* it also shows good correlation between simulations and measurements. For example, the loss is almost identical and the resonance appears at 40 GHz in the **Inline** as studied before. It even shows accurately that the Broadside Turn case has a gentle increase in loss at 30 GHz.

Finally, the *Mode conversion coefficient* also confirms that the characterization done for the Pin-field Escape features is very accurate in the frequency domain. The 3 cases are equally spaced which highlights that the best performing feature is the Broadside Turn followed by Inline Flank and **Inline**. Notice that at 30 GHz there is a sudden drop in mode conversion which could indicate that the channels has a resonance that is blocking the common mode.
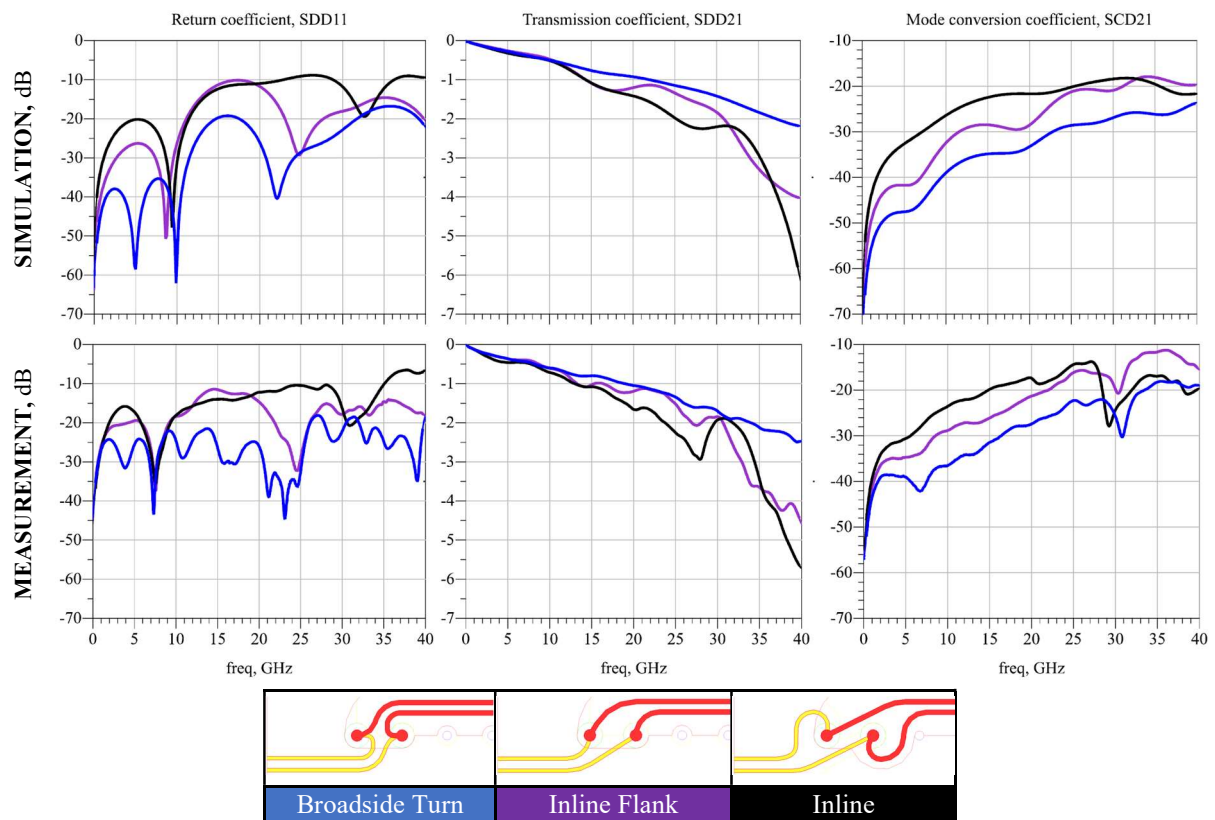
Figure 37: Pin-field escape S-parameters [Simulation -vs- Measurement]

On the other hand, the TDR measurements have a similar shape with some obvious differences (see Figure 38). Starting with the characteristic impedance of the center microstrip, in the measurement results it settles around 108 ohm while the simulations expected 98 ohm.

Considering this 10 ohm difference, the relative impedance change of the measurements is smaller than in the simulations. Despite these differences, the similarities are also easy to find. Clearly the **Inline** case is the most inductive and the Broadside Turn is the best because of the small impedance change.
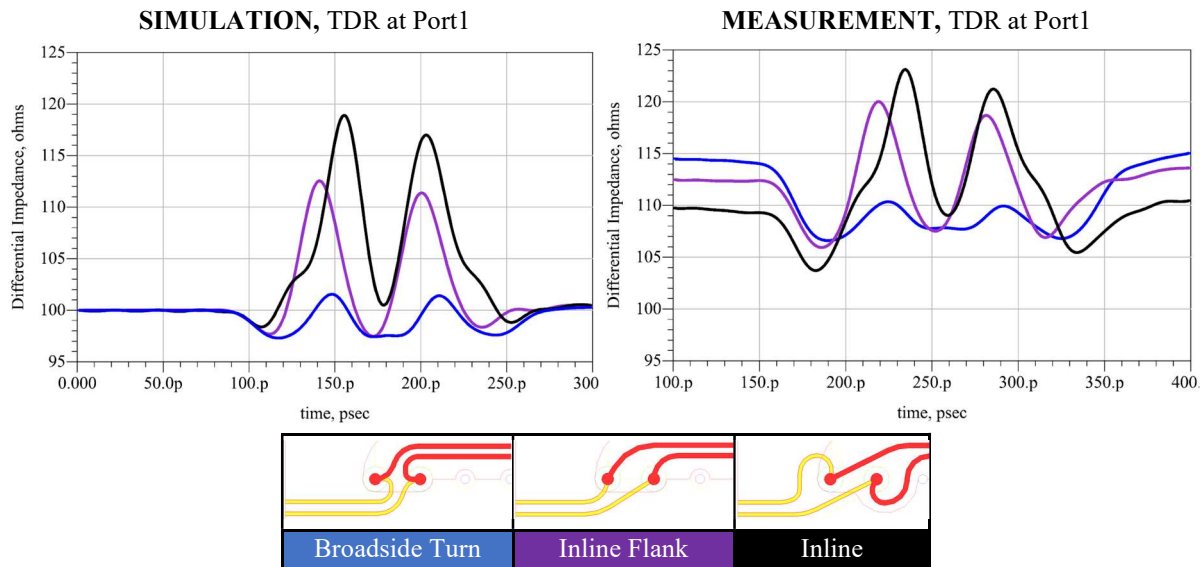
Figure 38: Pin-field escape TDR at Port1 [Simulation -vs- Measurement]

2.6 CROSSTALK

Crosstalk is a phenomenon that has been well studied and modeled but the magnitude of the effect is still dependent on stackup characteristics. The intent of this section is to quickly characterize crosstalk levels for the Napatech stackup and provide better understanding of this effect as preparation for Part 2.

With this in mind, using prebuild models in ADS is a convenient method for quantifying crosstalk because it can be solved quickly. When considering crosstalk scenarios, there is a victim and an aggressor that run side by side and the crosstalk effect varies with length and spacing between the two. The prebuild model has parametric length and spacing available for sweeping (see Figure 39).
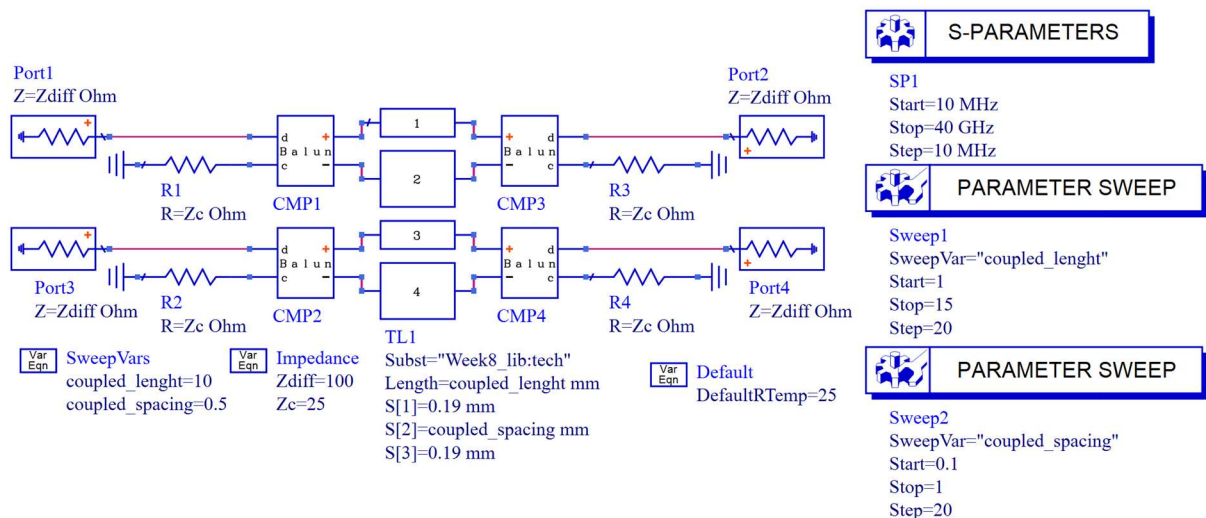
Figure 39: Crosstalk simulation using a prebuild model

Marc Borrell

In a board layout, crosstalk occurs when there is mutual capacitance and inductance between parallel transmission lines during signal transitions. Once crosstalk occurs at a certain point in the victim line, there are two signals created that flow away from the point of creation. One goes forward toward the end of the victim line. It is called forward crosstalk wave (Kf) and is the difference between the forward capacitive current and the inductive current. This signal might grow as it travels alongside the aggressor signal (for example in the case of a microstrip) or could be canceled if the mutual capacitance and inductance are equal (for example in the case of an ideal stripline).

On the other hand, the backward crosstalk wave (Kb) signal travels in opposite direction towards the beginning of the victim line. It is defined by the sum of the backward capacitive current and the inductive current. The signal does not grow in size after one rise time but expands in width as long as the coupling lasts. Literature and in-depth formulation of this phenomenon is found in [11] and [12].

The magnitude of the crosstalk effect is measured in dB. In order to decide what is an acceptable level for a particular system, it is recommended to measure the eye opening of a channel with crosstalk and decide if it is acceptable based on the eye dimensions. In this way, it accounts for available budget of the interconnect. Following this approach, crosstalk is quantified in this section for a wide range of cases and two selected cases will be evaluated in Part 2 with eye diagrams.

2.6.1   Crosstalk analysis

The crosstalk analysis consists of a sweep of parameters from 0.1 mm to 1 mm of spacing between victim and aggressor and from 1 mm to 15 mm of coupled length. Each one with 20 steps resulting in 400 simulations total. The large quantity of simulations highlights the relevance of the quick simulations provided by the prebuild model. All this while it is still considering the specific stackup and differential pair dimensions.

A differential signal is introduced at port 1 of the aggressor (see Figure 39). Then, Near-end crosstalk (NEXT) is the differential signal coming out of port 3 of the victim and Far-end crosstalk (FEXT) is coming out of port 4. Also, all ports are terminated to minimize reflections. Each simulations runs a frequency sweep from 10 MHz to 40 GHz and finally the results are presented as the mean of all frequency points for each simulation. These 400 points are presented in a 3D surface plot that has *Spacing* and *Length* as the independent variables and the maximum value of each simulation as the dependent variable in dB.

The results show that FEXT is affected slightly by *coupling length* especially at large spacings (notice the blue region in Figure 41) and is strongly dependent on *spacing*. In comparison, *coupling length* has no difference on NEXT while being very susceptible to *spacing* (see Figure 40). Notice that NEXT grows almost linearly with the decreasing *spacing* whereas FEXT grows exponentially. Nevertheless, the FEXT effects are lower than NEXT. For example, take 2 mm spacing at any length and observe that NEXT reaches 50 dB while FEXT only reaches 80 dB.

Since the simulated transmission lines are striplines, theoretically, FEXT should be canceled but it is not for two reasons. One is that the dielectric is homogeneous but not symmetrical in the vertical axis because the stackup has a 10um difference between the top dielectric and the bottom dielectric. Secondly, FEXT is composed by the forward traveling wave and the reflected reverse traveling wave (the reflected NEXT).

From these observations it is recommended that when routing neighboring high-speed lines on a board, the signals on those lines should flow in the same direction and signals flowing in opposite directions should be reserved for different layers.
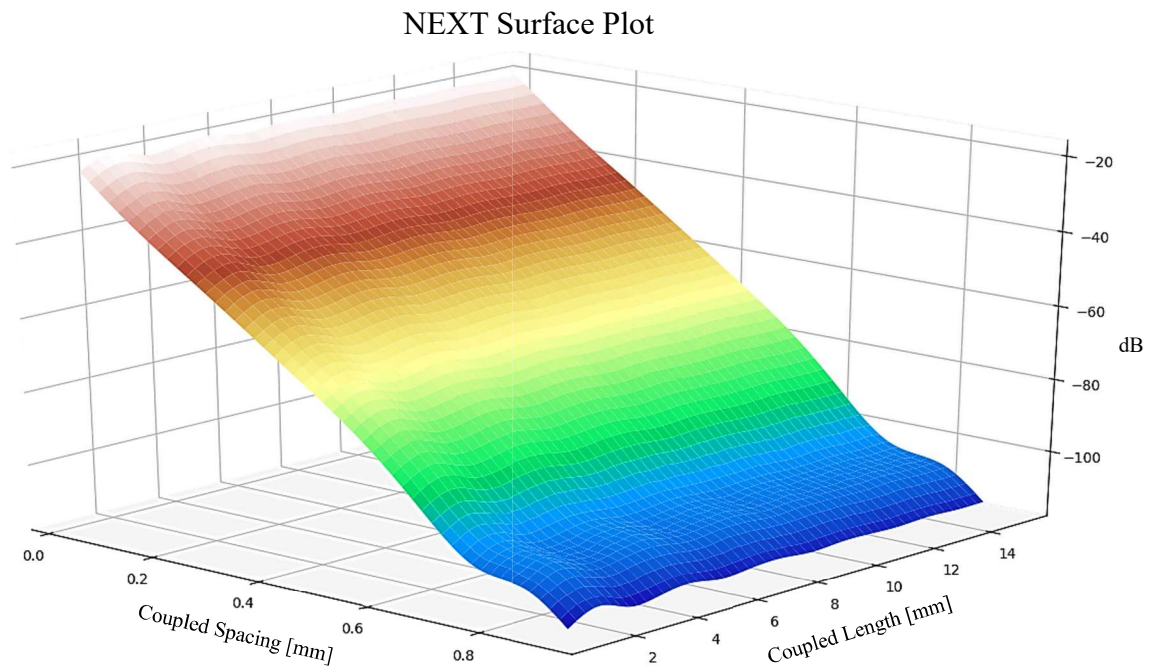
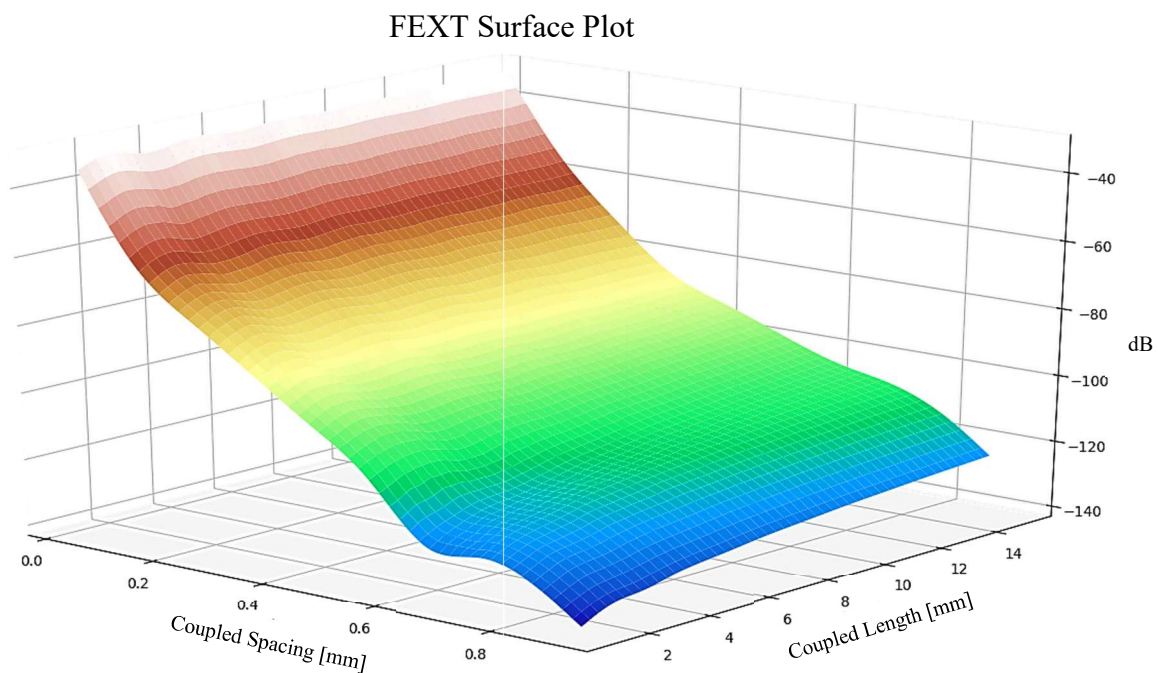

Figure 40: Near-end crosstalk NEXT magnitude

Figure 41: Far-end crosstalk FEXT magnitude

2.6.2 Validation

To validate the crosstalk simulations, 9 cases were manufactured with combinations of 0.1, 0.5 and 1 mm spacing and 5, 10 and 15 mm coupling length. This set had to be reduced to a minimum because of time constrains and only the two cases used in Part 2 were measured. The cases measured have 5 mm of coupling length and 0.1 and 1 mm of spacing.

The measurements show that crosstalk is stronger at 0.1 mm than at 1 mm. Also, NEXT has higher crosstalk levels than FEXT at the same spacing (see Table 9). In addition, there is a noticeable difference between simulations and measurements especially at 1 mm where the crosstalk effect is very weak. This is because the measurements are limited by the instrument sensitivity. The 10% and 18% discrepancy at 0.1 mm can be attributed to differences in physical characteristics between the simulated models and the manufactured test boards. Also notice that FEXT tends to have worst agreement because it depends on the dielectric characteristics as well as the reflections which are difficult to account for in simulations.

Since it is recommended to have the signals flowing in the same direction when having multiple high-speed lines on a board, FEXT will be the most relevant crosstalk for study. From these results, it is difficult to decide if the FEXT crosstalk levels (44dB@0.1mm and 99dB@1mm) are acceptable for a given interconnect. Instead, these values are used in Part 2 using eye diagrams. Those results show that 99 dB of crosstalk is acceptable while 44 dB closes the eye significantly, very close to the minimum dimensions acceptable (see Figure 49).

Spacing		Simulation	Measurement	Discrepancy
0.1 mm	NEXT	28 dB	31 dB	10%
	FEXT	52 dB	44 dB	18%
1 mm	NEXT	128 dB	96 dB	33%
	FEXT	160 dB	99 dB	62%

Table 9: Crosstalk at 5mm coupling length [Simulation -vs- measurement]

3 PART 2: MULTI-FEATURE CHANNELS

This second part pretends to validate if a typical interconnect would perform adequately using Napatech's stackup. To do so, this part has two sections:

FIRST SECTION, OPEN-LOOP: two distinct channels are created, named **Ideal** and Tough respectively. The **Ideal** uses some of the best characteristics one could choose when designing the layout of an interconnect. Nevertheless, as sometimes it is not possible to use ideal features because of geometric constrains, the engineer is forced to use less optimal solutions. That is what the Tough channel represents, and it is composed of some of the most challenging features encountered during this report. This section focuses on validating both channels in frequency domain. Notice that **Ideal** and Tough will be color-coded consistently throughout this Part 2.

SECOND SECTION, CLOSE-LOOP: focuses on the validation of the two channels using a Stratix10TX. To get good agreement between simulations and measurements, the complete signal path must be characterized, including connectors, coaxial cables and the Stratix10TX board traces. These will be referred as close loop simulations. The results are validated in the time domain.

3.1 OPEN-LOOP MULTI-FEATURE INTERCONNECTS

As introduced, two example channels are created to demonstrate the real capabilities and limitations of the stackup. These two demo channels have the same number of discontinuities in the same order and the equal length (see Figure 43) to ease the comparisons. The length of both is chosen to be like a VSR interconnect defined in OIF CEI VSR 56G [1] (100mm, see Figure 42), so that the S-parameter mask and eye mask can be used. There are three possible outcomes when comparing both channels:

1. <u>Both channels perform well</u>. Then the stackup has no limitations.

2. <u>Only the **Ideal** channel performs well</u>. Then new interconnect layouts need careful characterization when using this stackup.

3. <u>Both channels fail to perform</u>. Then the stackup needs better materials.
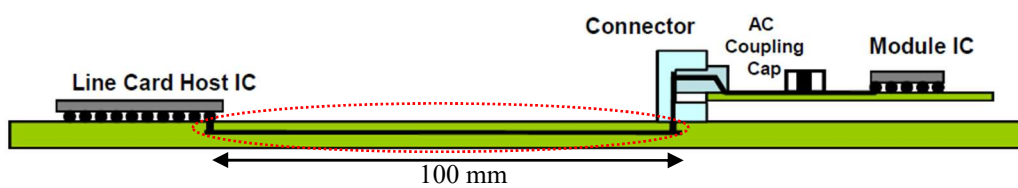

Figure 42: OIF CEI 56G VSR Interconnect diagram and length

The materials and dimensions used for the stackup cannot be specified because these are confidential. That is because the technology used in high-end boars is a core differentiator among the competition. Nevertheless, the conclusions presented still provide great insight on the channel performance when discontinuities, such as the ones previously studied, are encountered.

The two channels created have a large sample of the challenges discussed in Part 1. They try to reproduce a typical channel found on a NIC such as the one presented in Figure 1. On a real NIC, the channel would start at the FPGA package, then exit the pin-field usually next to an antipad from neighboring vias.

Then, the channel would find the most direct route to its destination, for example a transceiver. In doing so it typically does a 90º turn to align with the transceiver connector. Since these connectors are typically multilane, it might also be routed next to another high-speed channel. Finally, the channel requires a via to transition from the FPGA package to an inner layer and then exit at the connector side when using striplines.

The **Ideal** and the Tough channels presented in Figure 43 emulates all the discontinuities just described (see Table 10).

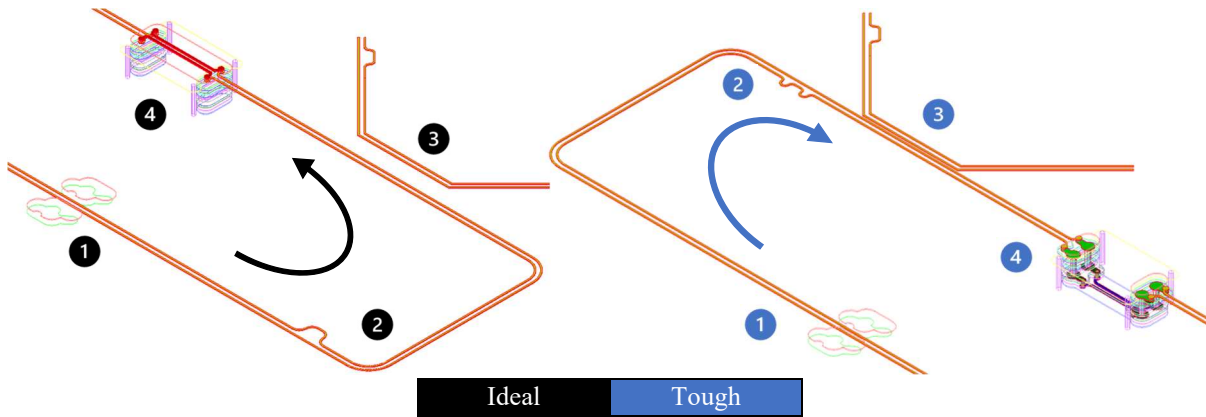Marc Borrell

Figure 43: Multi-feature interconnect overview

Discontinuity	IDEAL Channel	TOUGH Channel
1. Symmetrical Antipad	1 mm Pitch BGA	0.7 mm Pitch BGA
2. Length equalization	1 Big bump	2 Small bumps
3. 5 mm length Crosstalk	1 mm Spacing	0.1 mm Spacing
4. Via transitions	Top – L2	L2-Bottom
Total length	83 mm	

Table 10: Multi-feature interconnect comparison

3.1.1 Frequency domain analysis

Simulating the entire channel at once is not possible because the channels are 83mm long which would require too much memory. Instead, the S-parameter models created during Part 1 are cascaded to create a unified S-parameter model for each channel. As done previously, the most relevant S-parameter coefficient are discussed next. Nevertheless, in the following they also include the relevant mask in red (see Figure 44), to validate the channels using the OIF CEI 56G VSR standard [1].

Transmission coefficient, SDD21: Insertion loss is very good in both cases and, they pass the minimum defined by the standard with ample margin (see Figure 44). Both should have the same loss because they have the same length, but only the **Ideal** channel meets the expected loss. The Tough channel has small oscillations, the first at 15 GHz and the next above 40 GHz.

These are the same resonance frequencies found during the via type section for L2-Bottom case.

Mode conversion coefficient, SCD21: Mode conversion plot shows a clear difference between channels (see Figure 44). It is safe to assume that most of the mode conversion comes from the length equalization method. That is because all other discontinuities did not show large difference when it comes to mode conversion between the best and the worst features. Even though the Tough is the worst and could be improved, the amount of mode conversion is not problematic in most cases.

Return coefficient, SDD11: Return loss has more oscillations than usual because the length of the path is longer (see Figure 44). Both channels pass the maximum mask limit without issues. The only concern comes with the increased return loss at around 15 GHz for the Tough channel. This is due to the large capacitive effect that the via transition has.
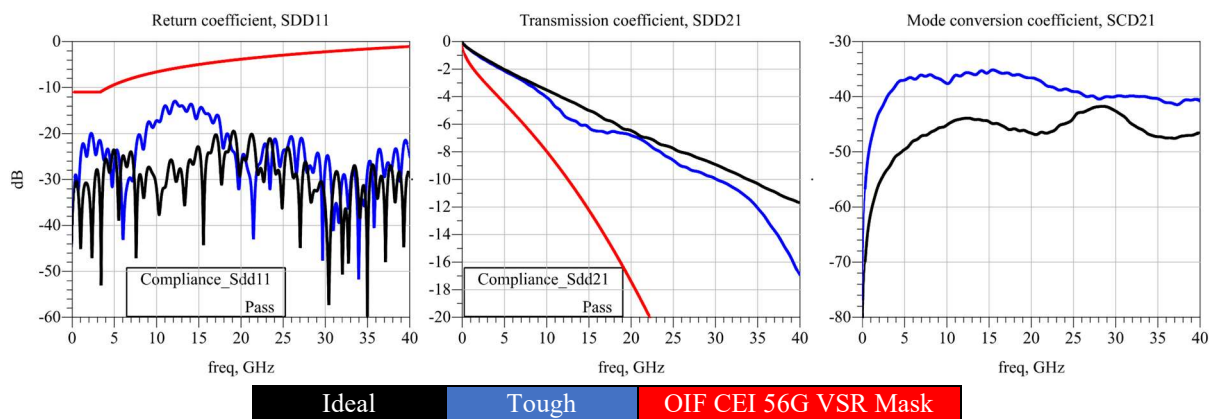

Figure 44: Open-loop Multi-feature interconnect, frequency domain results

3.1.2   Time domain analysis

Now that the TDR responses are well known from previous sections, it is easy to locate the effect of each feature in the complete channel. The first discontinuity (1) appears at 400 ps only in the Tough channel due to the near antipad, whereas the **Ideal** channel is not affected. The second discontinuity (2) takes place at 600 ps in the **Ideal** channel and at 800 ps in the Tough channel. Both spikes come from the inductive effect of the length equalization, which in the case of the Tough channel is unintentionally positioned before the 180-degree turn, whereas in the **Ideal** case it is positioned after. One could also have used port 2 for the TDR measurement on the Tough channel, but the performance is equivalent and the features would appear in reverse order in the plot.

The next disturbance comes from the Crosstalk aggressor trace (3). When the trace is at 0.1 mm from the Tough channel, the line becomes more capacitive thus lowering the impedance. The aggressor of the **Ideal** channel is too far away to cause such effect. Finally, both channels are affected by the via transition (4). The via transition in the **Ideal** channel is inductive while the transition in the Tough channel is very capacitive as observed in its own study.
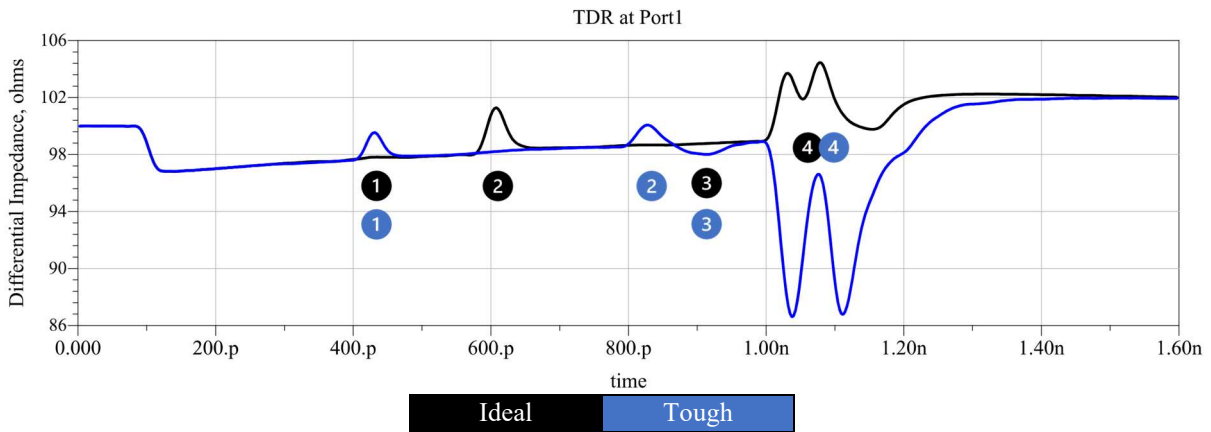
TDR at Port1

Figure 45: Open-loop TDR at Port 1, Differential impedance
(1) Symmetrical Antipad (2) Length equalization (3) Crosstalk5 mm length (4) Via transitions

The eye contour plots used previously do not provide enough information for the multi-feature channels and eye density plots are presented instead. The eye density plots are not measured at the channel but at the output of the receiver. These plots are useful because it also shows the dispersion of the signal transitions and the effect of the equalization. It is recommended to have Short Range (SR) CTLE and SR FFE on both channels because the loss of the channel is less than 15 dB@14 GHz [13].

Figure 46 shows the eye density plot for each channel of two UI. The eyes are open in both cases but the **Ideal** channel has slightly more height. The eyes are correctly equalized because the region with highest density (shown in red, see Figure 46) is at the center, very near 40 ps. Notice that the Tough channel has a lower maximum and minimum voltage of around 0.25V. Since they are very similar in shape and size, they will be compared more in detail in the conclusion of this section by using Figure 49 and Table 11.

Figure 46: Open-loop multi-feature channel, eye density plots WITHOUT crosstalk
Color bar indicates the number of signals in % crossing a particular point

The simulations of the two channels are repeated, but this time with crosstalk added through the adjacent aggressor line. The aggressor signal flows in the same direction, as it would typically do in a NIC. Recalling from Table 9, FEXT crosstalk is 44dB@0.1mm and 99dB@1mm. With this in mind, Figure 47 shows the effect that FEXT has on the eye opening of each channel. Clearly the **Ideal** channel with 1 mm spacing is completely open whereas the

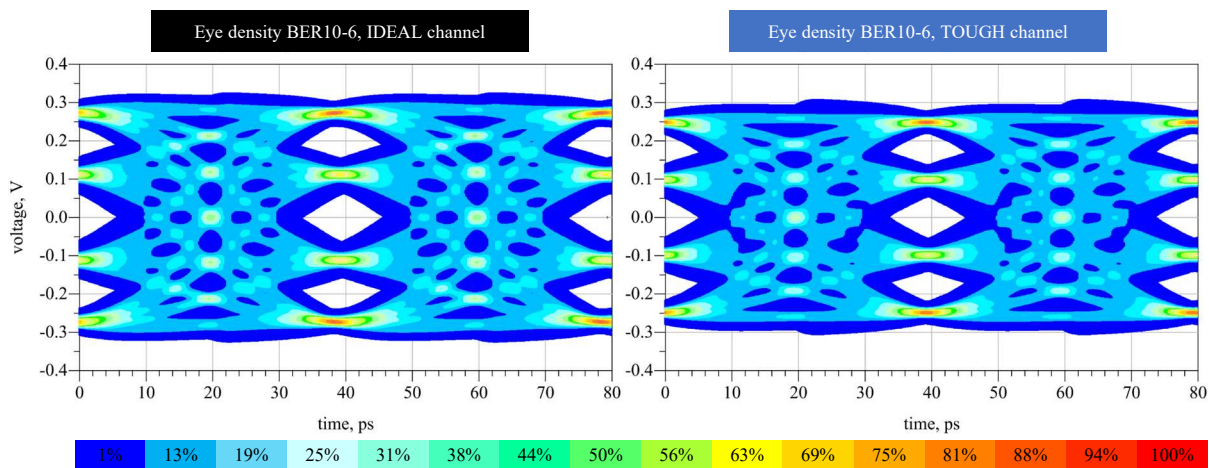channel with 0.1 mm of spacing between aggressor and victim has a much smaller eye opening.
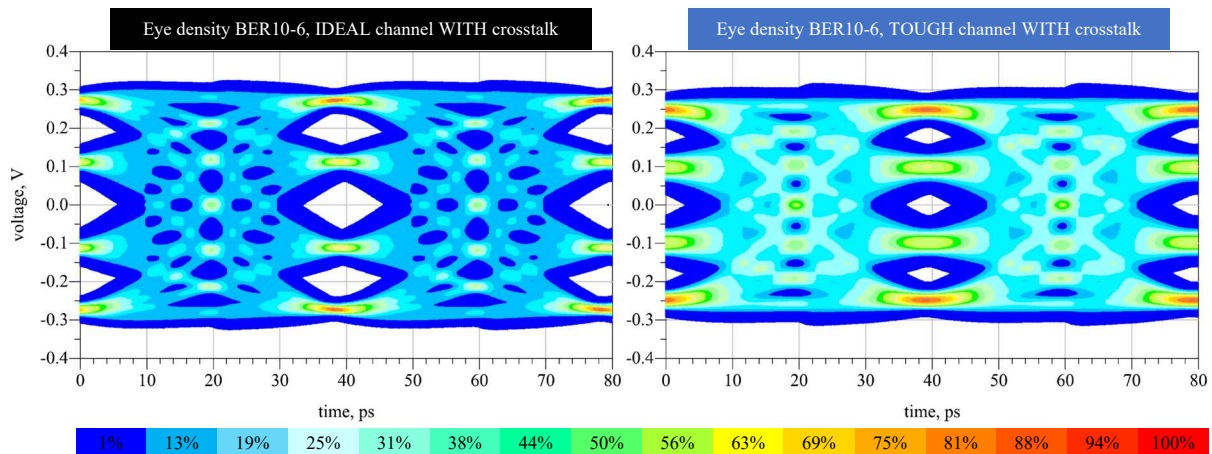

Figure 47: Open-loop multi-feature channel, eye density plots WITH crosstalk
Color bar indicates the number of signals (in %) crossing a particular point

With the intention of showing the effect of the equalization, the simulation with crosstalk in Figure 47 is repeated without SR CTLE and SR FFE enabled (see Figure 48). For this reason, the shape of the eye is less symmetrical and skewed to the left (the red region is very close to 30 ps). Moreover, the eye opening of the **Ideal** channel is still open while the Tough channel is significantly smaller (see Figure 48).
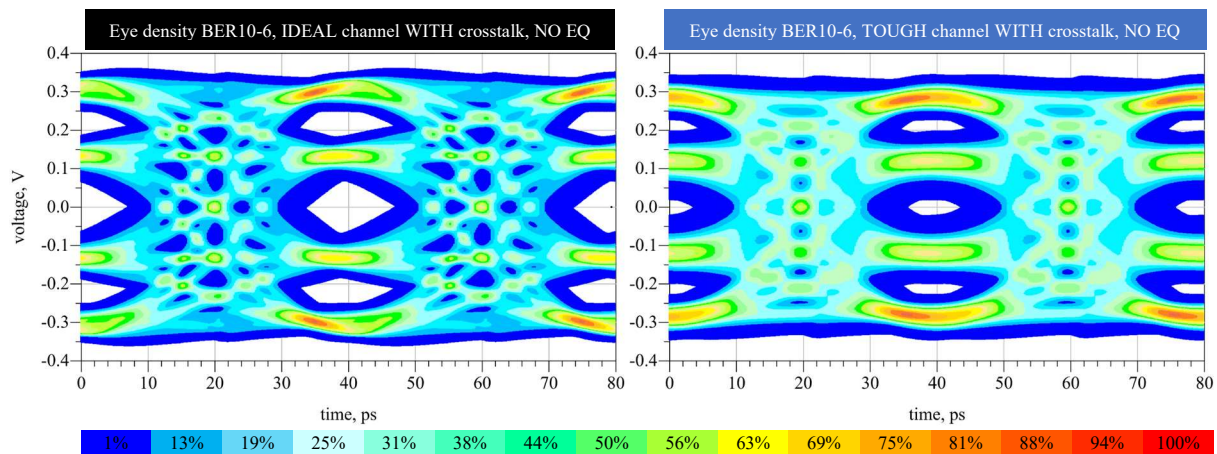
Figure 48: Open-loop with crosstalk without equalization, eye density plots
Color bar indicates the number of signals (in %) crossing a particular point

3.1.3   Conclusions

The **Ideal** channel is clearly superior in all measures, both in frequency and time domain. Nevertheless, still both channels have open eyes. To be able to compare the eye openings more accurately, a contour plot is presented in Figure 49. A comparison of the **Ideal** channel against the Tough channel is shown, with and without crosstalk. The figure illustrates that, even though the difference is small between the two channels without crosstalk, the opening of Tough eyes is narrower and have less height. This is even more noticeable when simulating with crosstalk where the **Ideal** channel is not affected by the 99 dB of FEXT crosstalk, whereas the Tough channel is severely closed by the 44dB of FEXT crosstalk.

Figure 49: Open-loop, Eye Contour plot, Crosstalk impact on Ideal -vs- Tough

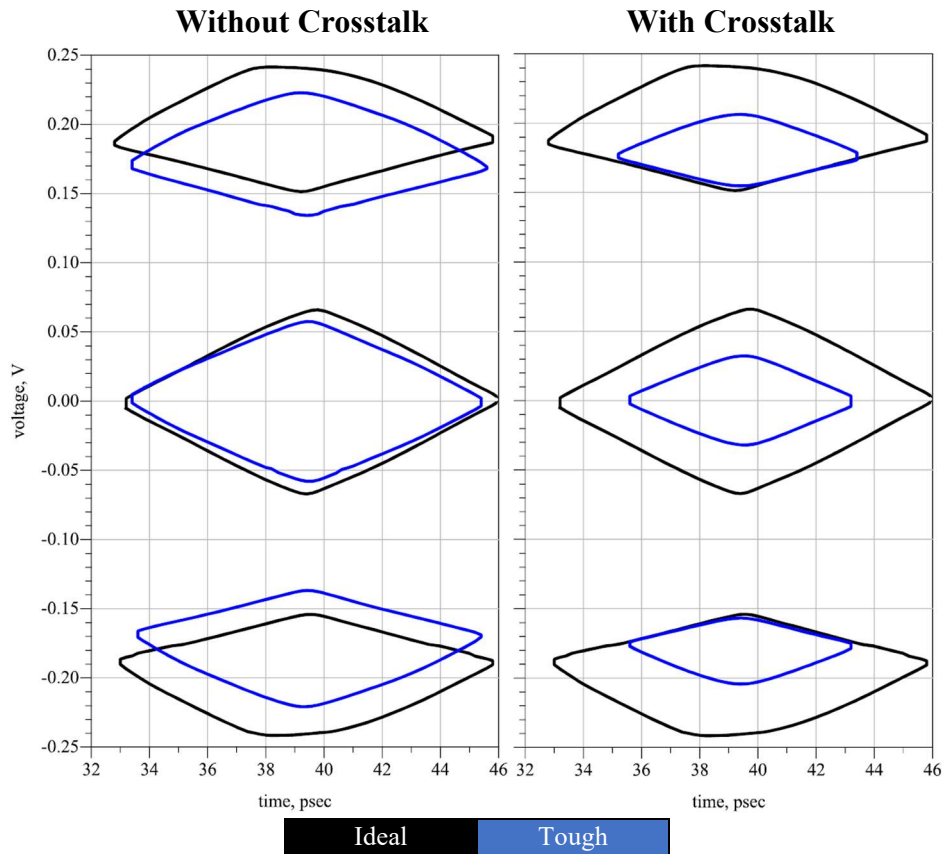It is also important to verify the channel performance using an adequate standard, in this case the OIF CEI 56G VSR PAM4 [1] or equivalently IEEE 802.3cd 50GAUI-1 C2M [2]. Table 11 has a compilation of the eye measurements which shows the raw values and the comparison of these to the standard mask. The standard defines that the minimum height must be 0.032, the width must be larger than 20% of the unit interval (UI = 40 ps) and the eye linearity must be more than 85%.

None of the scenarios pass eye linearity mask but they all pass the height and width mask. Since eye linearity is very close to the standard requirements, it is reasonable to assume that both channels perform adequately in both conditions. Be aware that the Tough channel has a smaller margin when there is crosstalk. Therefore, any trace longer than 5mm of coupling or tighter spacings could end up causing the channel to fail the standard.

	Ideal	Tough	Ideal, Crosstalk	Tough, Crosstalk
Height, (EH6), V	0.083	0.079	0.083	0.044
Width (EW6), %UI	30%	28%	30%	19%
Eye linearity (EL), %	72%	77%	72%	77%

* Results over 100% comply with the standard

EH6 / **0.032V**
EW6 / **20%**
EL / **85%**

	Ideal	Tough	Ideal, Crosstalk	Tough, Crosstalk
Eye Height Mask*	259%	247%	259%	138%
Eye Width Mask*	148%	140%	150%	95%
Eye linearity Mask*	85%	91%	85%	91%

Table 11: Eye opening summary, Ideal and Tough w/wo crosstalk

Marc Borrell

3.1.4   Validation

The frequency domain measurements have very good agreement with the simulations even though they are the most complex, as each channel has a combination of 4 different features (see Figure 50). Starting with *return coefficient*, both channels pass the SDD11 mask as the maximum levels stay under 10 dB. In particular, the Tough case has a challenging zone between 5 and 20 GHz, which comes from the L2-Bottom via transition that is accurately predicted by the simulation. Otherwise, both channels have very similar return loss.

Continuing with the *transmission coefficient*, the loss of the **Ideal** channel is 13 dB which is very close to the 12 dB predicted. The Tough channel has a resonance at 15 GHz, which is larger than expected, as well as it has a resonance at around 30 GHz that does not appear in simulation. These also come from the L2-Bottom via transition, as shown in Part 1.

To conclude the frequency domain with the *mode conversion* measurements, these are slightly noisier than the simulations but the trajectory is similar and there is a clear distinction between the **Ideal** and Tough channel performance.
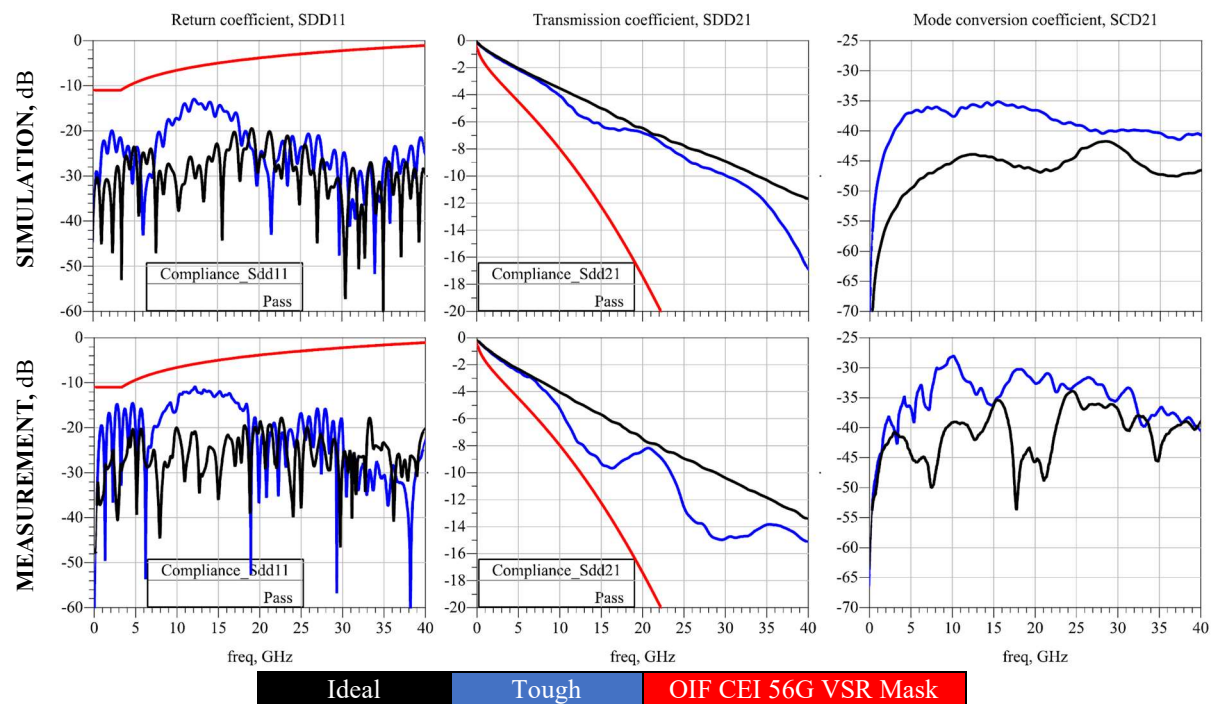

Figure 50: Open-loop, S-Parameters [Simulation -vs- Measurements]

As usually with the TDR measurements, the characteristic impedance of the straight striplines does not match the simulations but the discontinuities introduced are still noticeable (see the numbers highlighted in Figure 51). Starting the discontinuity analysis with the antipad (1), the **Ideal** channel does not have any disturbance. Instead, the Tough channel has a 2 ohm bump which matches the simulation. Then the **Ideal** channel has a larger disturbance than expected due to the length equalization feature (2). Further forward, the Tough channel also has a disturbance due to the length equalization. This is larger than the simulations because the impedance dropped during the 180º turn. The third point (3) shows a very small difference due to the crosstalk feature between the **Ideal** and the Tough case. Finally, the via transitions are

Marc Borrell

very distinct for each channel, as expected with simulations. The **Ideal** case is slightly less inductive than simulations, whereas the Tough channel is more capacitive than expected.
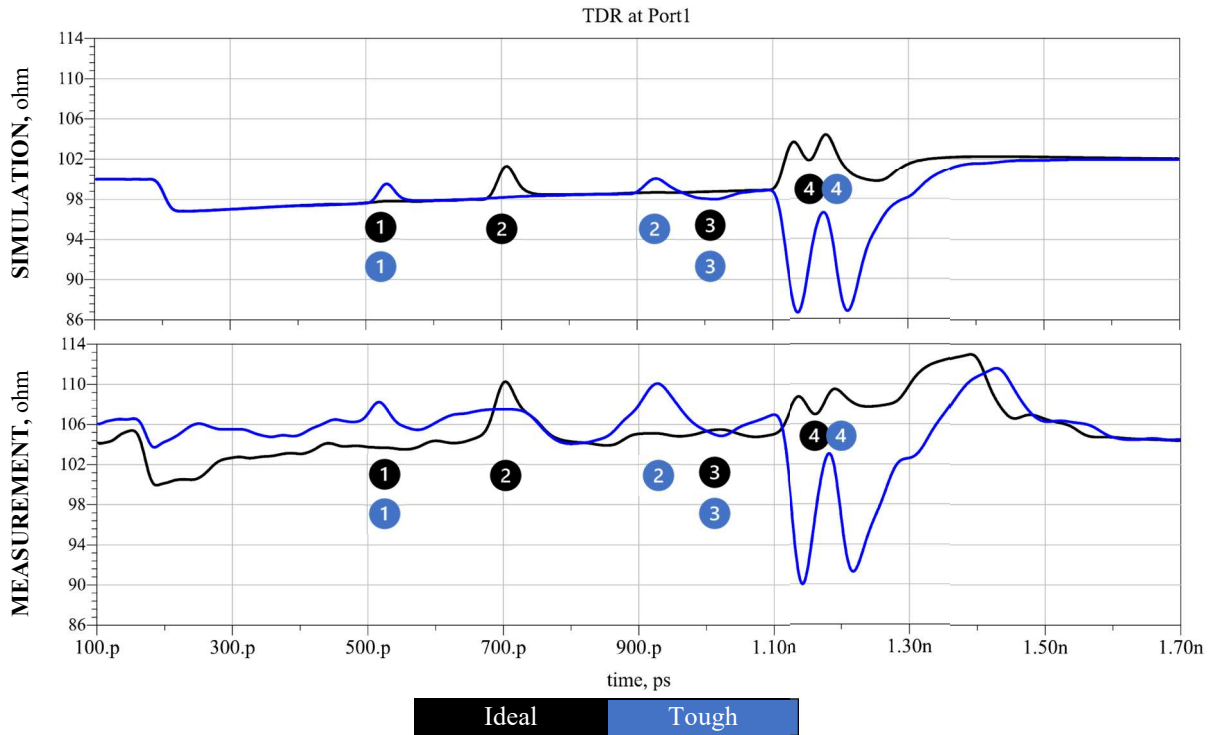

Figure 51: Open-loop TDR at Port 1[Simulation -vs- Measurement]
(1) Symmetrical Antipad (2) Length equalization (3) Crosstalk 5 mm length (4) Via transitions

3.2   CLOSE-LOOP MULTI-FEATURE INTERCONNECT

With the intention of correlating real eye measurements with eye simulations, the simulated channel must be identical to the real channel. The PAM4 signals used for the real measurement are transmitted and received from a Stratix10TX. The devices under test (DUT) are the **Ideal** and Tough channels. Therefore, the signal needs to flow through them, creating a close loop. Each individual component along the path has been characterized and cascaded together to create a two S-parameter models representing the close-loop for the Ideal channel **(Ideal CL)** and the close-loop for the Tough channel (Tough CL), see Figure 52.
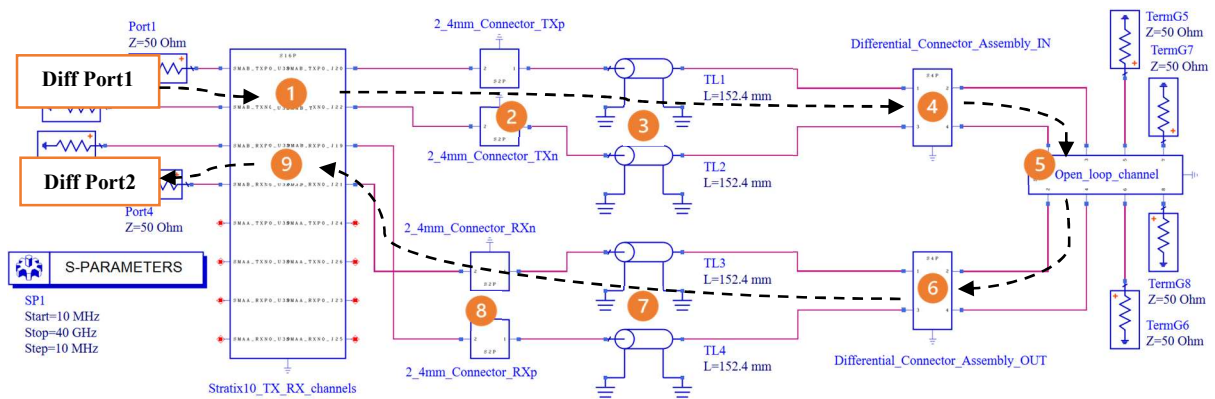

Figure 52: Close-loop multi-feature channel overview

(1) Stratix10 TX channel (2) 2.4mm Upright connector (3) 15 cm Coax cable (4) 2.92mm Connector to L2 differential stripline (5) Ideal & Tough channels (6) 2.92mm Connector to L2 differential stripline (7) 15 cm Coax cable (8) 2.4mm Upright connector (9) Stratix10 RX channel

Marc Borrell

3.2.1   Frequency domain analysis

In this close-loop section, the channel is extended well beyond the maximum length of a standard VSR interconnect. For this reason, the red VSR masks are only included for reference but are not relevant for interconnect validation. Instead, a validation for mid-range reach interconnects (MR, 50cm length) is done with COM and is discussed later. Nonetheless, the S-parameters are discussed to ensure characterization correctness of the components included in the path.

Transmission coefficient, SDD21: Its periodic oscillations are an indication of reflections and it is normal for long channels with many discontinuities. There are no signs of passivity or causality violations (usually in the form of non-periodic spikes or positive gain as explained in [14]). As a side note, both channels fail the VSR mask which is expected due to the longer path.

Mode conversion coefficient, SCD21: The mode conversion in both cases is very similar and have low conversion levels. Keep in mind that these results could differ significantly in the real setup where cables and fixtures might not be completely symmetrical in length thus creating skew.

Return coefficient, SDD11: Return loss has very low levels in both cases and even passes the VSR mask which is not required. It also has oscillations because of the long path and the reflections caused by the discontinuities.
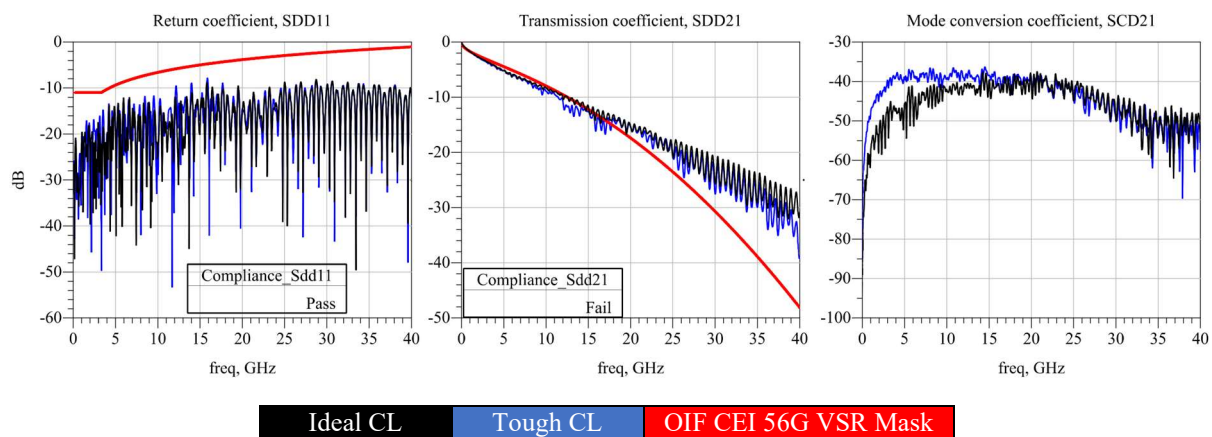

Figure 53: Close-loop Ideal and Tough channels, frequency domain results

3.2.2   Time domain analysis

The time domain analysis is the most important part of this section. It illustrates the cause of the discontinuities along the path and the effects it has in an eye diagram, both in simulation and reality.

The TDR plot at Port 1 show that the impedance of the channel always within 10% of the 100 ohm characteristic impedance. The largest close-loop discontinuities come from the connectors which have an inductive effect (see Figure 54, (2)(4)(6)(8)). Furthermore, the Stratix10TX paths and the coaxial cables are clearly distinguished by its long and uniform impedance (see Figure 54, (1)(3)(7)(9)). These do not have any visible discontinuity.
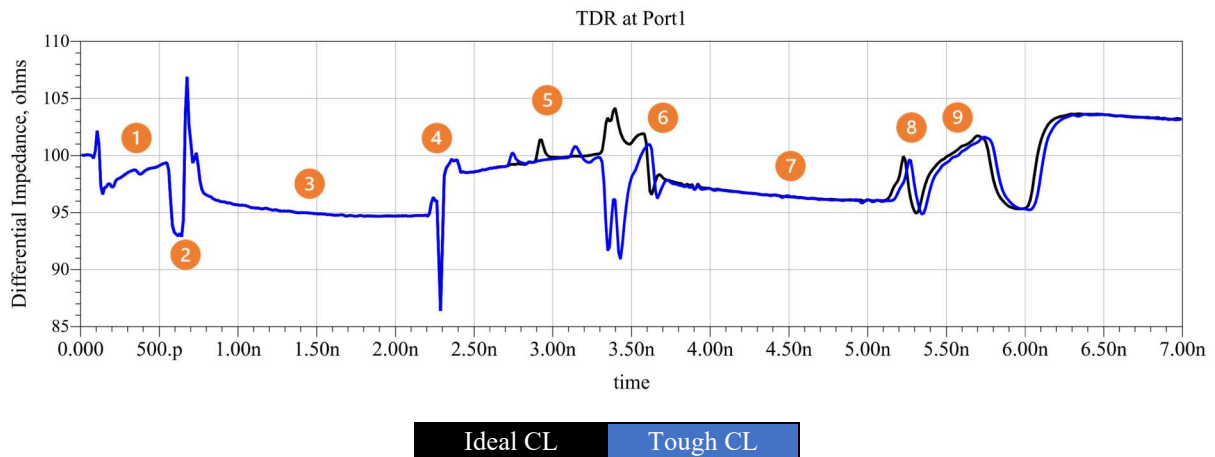
Figure 54: Close-loop multi-feature channels, TDR at Port 1, Differential impedance
(1) Stratix10 TX channel (2) 2.4mm Upright connector (3) 15 cm Coax cable (4) 2.92mm Connector to L2 differential stripline (5) Ideal & Tough channels (6) 2.92mm Connector to L2 differential stripline (7) 15 cm Coax cable (8) 2.4mm Upright connector (9) Stratix10 RX channel

When preparing the SERDES simulations for the eye density plots, one must consider if it is required to use SR CTLE and SR FFE. It is not recommended to use it with insertion loss higher than 15dB@14GHz [13]. Since the close-loop channels have 11dB of loss at 14GHz, it is still recommended to use equalization but it does not have a significant difference.

Comparing the **Ideal CL** channel and the Tough CL channel without crosstalk (see Figure 55), the two eye diagrams are practically identical. The **Ideal CL** channel has slightly better performance because of the dark blue zones at 20 ps and 60 ps which indicate that the transition lines are less disperse. In the section's conclusion, these slight differences are evaluated with concrete values. However, one can already see that both channels have the eyes significantly smaller than when not closing the loop. Keep in mind that the most relevant results in terms of absolute eye dimensions are the open-loop channels because they have the length and characteristics typically found in a NIC. The close-loop channels are meant for eye validation with real measurements captured with the Stratix10TX.
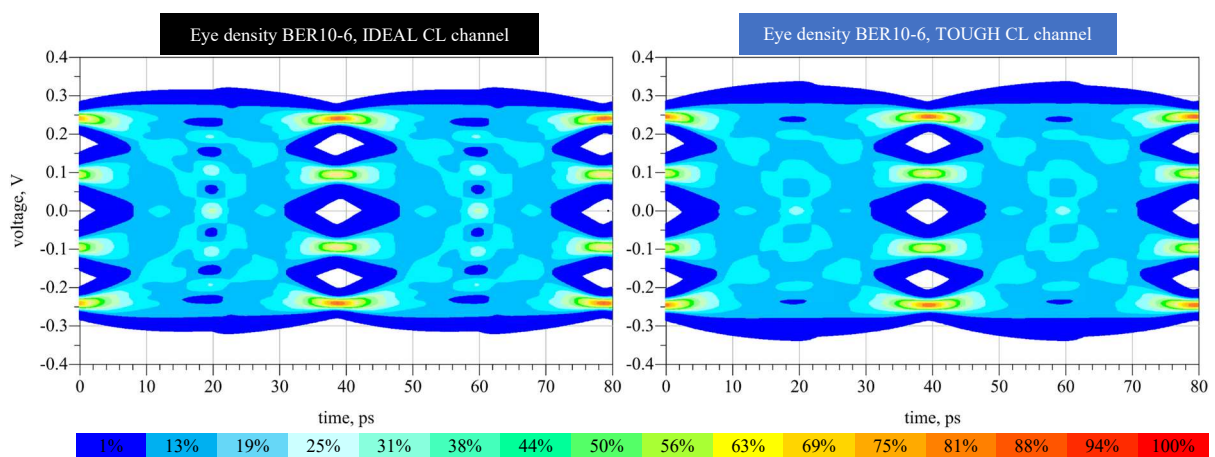

Figure 55: Close-loop Multi-feature channel without crosstalk, eye density plots
Color bar indicates the number of signals (in %) crossing a particular point

To complete the time domain analysis, Figure 56 illustrates the effects of crosstalk on the close-loop channels. The **Ideal CL** channel has the same opening compared with the previous simulation without crosstalk. Finally, the Tough CL channel has the eye almost completely

closed due to crosstalk. This is understandable, as the open-loop Tough channel with crosstalk had a small eye and could barely pass the VSR eye mask. Notice that when the eye is closed, the rising and falling edges of the signals are less defined and the probability of finding a transition at each point is higher. Even though the transitions are sparce, the voltage levels are the same, which indicates that the eye is mostly closed due to noise in the time axis. This is known as jitter and could be caused by crosstalk.
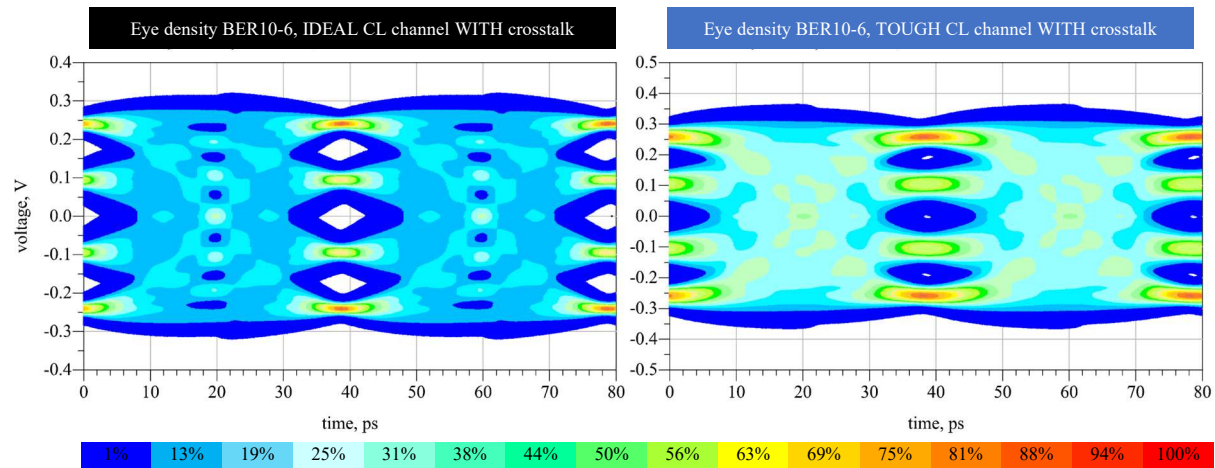

Figure 56: Close-loop Multi-feature channel with crosstalk, eye density plots
Color bar indicates the number of signals (in %) crossing a particular point

3.2.3   Conclusions

To be more precise with the comparisons, the eye dimensions of the close-loop cases are compared in Figure 57 and presented in Table 12. As mentioned before, the eye opening in the **Ideal CL** is the same with and without crosstalk whereas the Tough CL channel closes significantly.
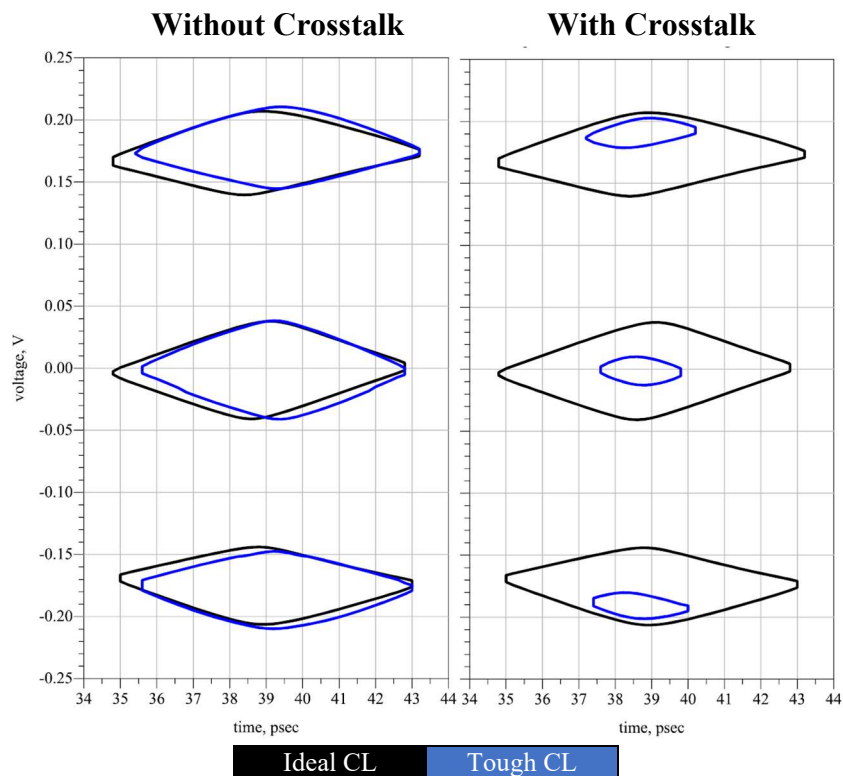

Figure 57: Eye Contour plot, Crosstalk impact on Ideal CL -vs- Tough CL

Marc Borrell

Similarly, the Table 12 shows that the Tough CL channel with crosstalk is almost closed. Otherwise, **Ideal CL** and Tough CL are almost identical. Notice that the VSR mask table is not included because it is not relevant for the close-loop case as it is considered a MR channel.

	Ideal CL	Tough CL	Ideal CL Crosstalk	Tough CL Crosstalk
Height, (EH6), V	0.058	0.058	0.058	0.015
Width (EW6), %UI	20%	18%	20%	7%
Eye linearity (EL), %	78%	77%	78%	75%

Table 12: Close-loop multi-feature channels, eye dimensions summary

3.2.4   Validation

The validation of the close-loop multi-feature channels is completely different from all other validation techniques used previously. Those are done with a VNA whereas this is done with an Intel Stratix10TX FPGA. The reason for this is to validate the channels in time domain with a SERDES setup. Then, the transmitter and receiver of the FPGA are used to create a signal that flows through the channel. The output of these measurements is an eye density plot.

Unfortunately, the Intel Quartus tool does not provide the results in international units and instead uses steps from -255 to 255 for the Y axis and from 53 to 148 for the X axis. There is no indication on what these correlates to. Therefore, from this point onward it is assumed that the Y axis voltage range from -0.5 V to 0.5 V and that the X axis is 1 UI (see Figure 58).
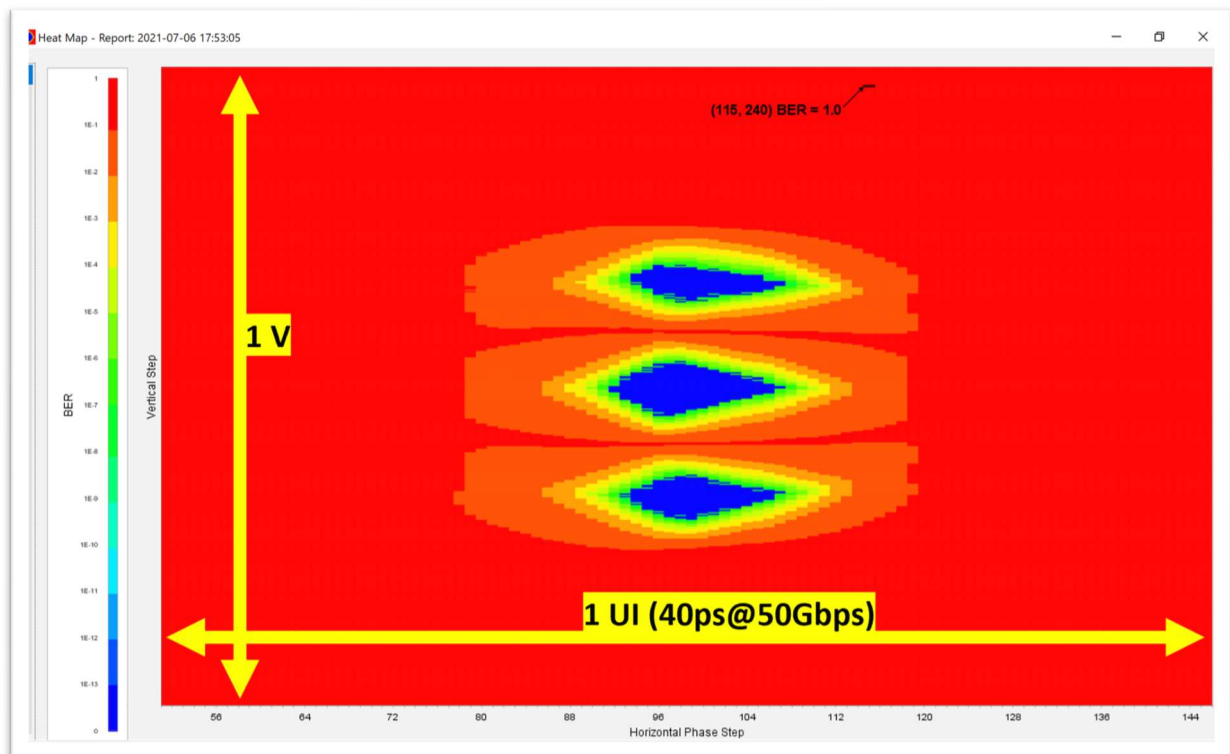

Figure 58: Intel Quartus Eye diagram, Assumptions in yellow

Once the fundamental basis for comparison have been established, the contour of the eyes is extracted manually by drawing a polygon overlaid on top of the eye diagram presented in Figure 58. This polygon is then imported into ADS for comparison (see Figure 59). Since there

has been several assumptions and manipulations, the units of the figure have been removed to bring the attention to one-to-one comparisons.

Now to the actual validation, Figure 59 shows the **Ideal CL** channel in the first two columns and the Tough CL channel in the last two. All cases have an overlay in red of the measured eye with the Stratix10TX. The measurements show in general good agreement with the simulations. Yet, the **Ideal CL** simulations do not capture properly the decrease in width. This could be due to the large deviation shown in Table 9 between measured FEXT of 99dB and simulated FEXT of 160 dB (discrepancy of 62%). Furthermore, this hypothesis also holds with the Tough CL channel results. This channel captures much better the crosstalk effect possibly because the measured discrepancy in Table 9 is also much lower at only 18%.
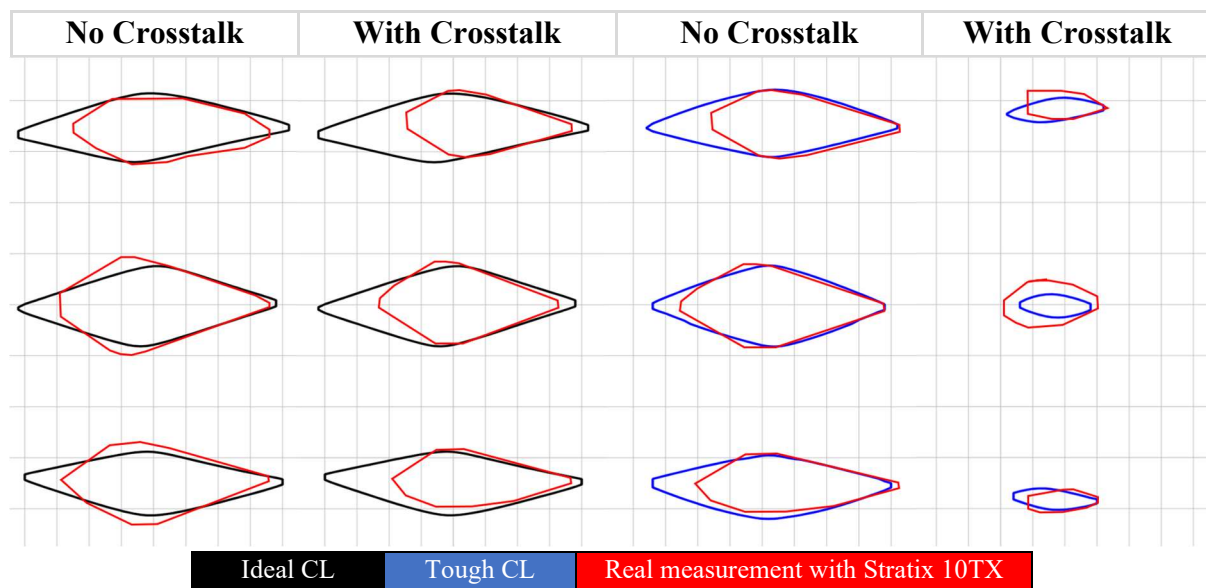

Figure 59: Eye contour plots, Close-loop [Simulation -vs- measurement]

3.2.5   Channel Operating Margin (COM)

The close-loop channel falls outside the VSR standard because of length. Therefore, it is considered as a Mid-range reach (MR). The OIF CEI 56G MR [1] introduces the Channel Operating Margin (COM) as a validation tool. Table 13 shows the COM results for the **Ideal** and Tough channels in the 4 scenarios discussed previously, including open-loop, close-loop and with or without crosstalk. The open-loop multi-feature channels from the previous section have been included for comparison purposes. However, these cannot be verified with COM because the VSR standard does not use it.

The results show that the close-loop cases without crosstalk are very similar and pass the standard as they are above 3dB. With crosstalk, the **Ideal CL** channel is barely affected and still passes the MR standard. On the other hand, the Tough CL channel completely fails with a score of 0dB. It is worth mentioning that having a 0.1mm spacing is not realistic because it is very small so this dramatic effect is to be expected.
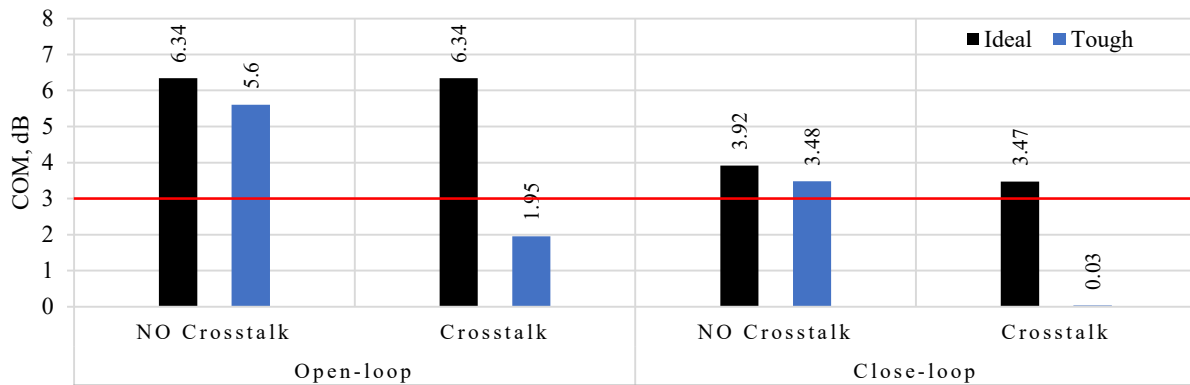
Marc Borrell

Table 13: Ideal and Tough channel COM results in all conditions

4 FIRST INSIGHT INTO FUTURE WORK

There are multiple ideas that sparked during this thesis which would be interesting to work on.

The first one is called "Distance to cause". When doing length equalization, it is not clear how far away can the compensation be from the turn. For example, in some cases the 90º turn is introduced inside a BGA pin-field but there is no space to create a bump for equalization. The layout engineer is then forced to position the bump further along the path. That creates a region where the modes are not compensated.

A preliminary test has been initiated (see Figure 60). It consists of a parametric ADS layout which has a bump that can be moved to the right by specifying a distance in millimeters. This parameter can be swept in an ADS schematic.
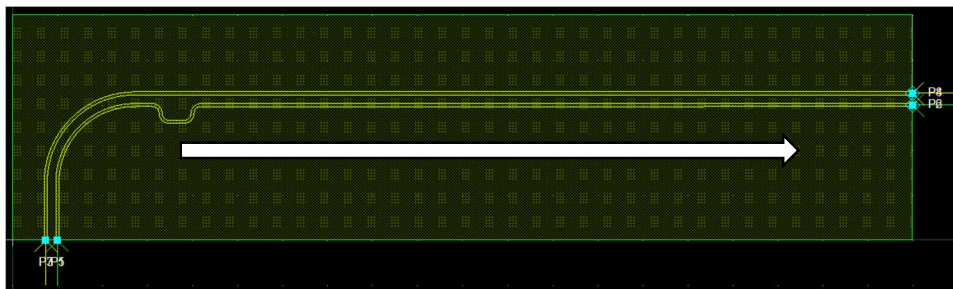

Figure 60: Distance to cause, ADS layout overview

Initial results show that there is no change in the difference to common mode conversion when sweeping between from 0 to 10 mm in 0.5 mm steps (see Figure 61). The plot on the left shows the mode conversion at each distance. The plot on the right calculates the mean of the SCD21 for each distance to see if there is a trend hidden in the left plot. Notice that the variation is very small of only 1 dB which is probably caused by simulation noise. The conclusion is that it makes no difference how far way the turn is compensated. Further investigation is required at other distances. Also, EMI radiation could be included as an evaluation tool.
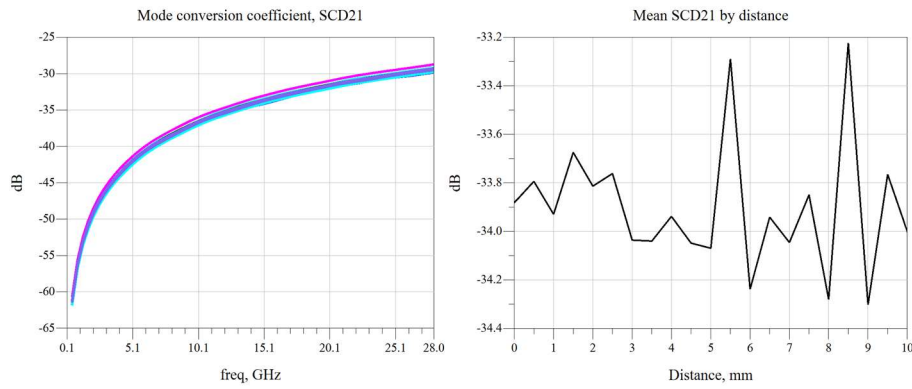
Marc Borrell

Figure 61: Distance to cause, SCD21 and Mean SCD21

Finally, another interesting challenge is the optimization of L2-Bottom via transition. This could help improve the large capacitive effect discovered during the via analysis. Initial tests have been very successful. It has been achieved reducing the pad size in L3 and L12 from 0.55 mm in diameter to 0.26 mm and increasing the antipad size in all layers by 0.34 mm. These configurations have been found by trial and error and there is margin for even more improvement with optimization tools. The frequency domain results show that the large resonance in the insertion loss is eliminated (see Figure 62).
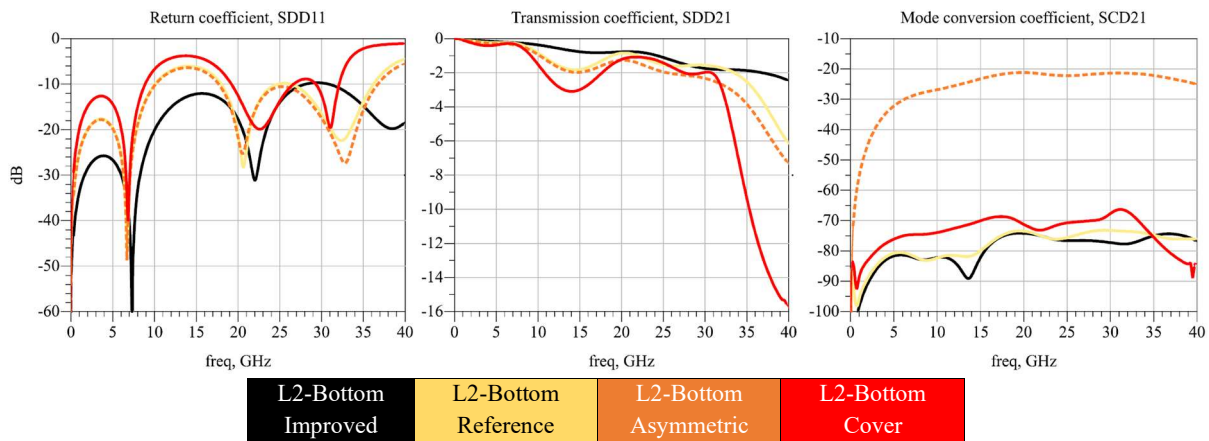

Figure 62: L2-Bottom via improvements, S-parameters

Similarly, the TDR plot illustrates that the transition capacitance has been reduced by 15 ohms (see Figure 63). To give validity to these results, they should be manufactured and measured like all the other features in the report.
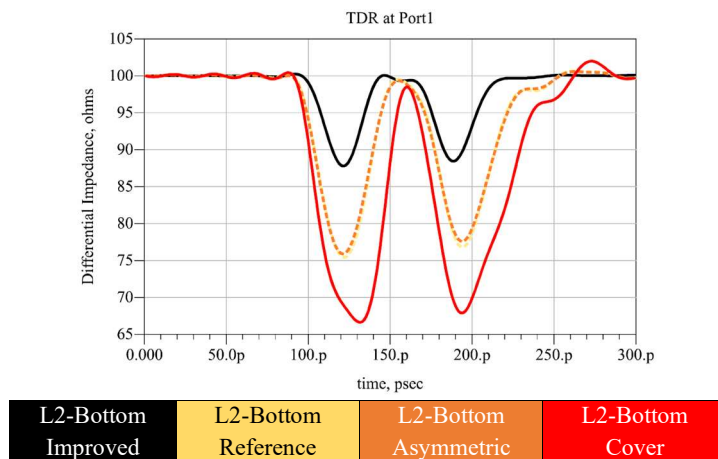

Figure 63: L2-Bottom Via improvements, TDR at Port1

Marc Borrell

5 CONCLUSION

THE FIRST PART

Length equalization: It can be concluded that there is an optimum size, which is the "Big" case, that minimizes the differential to common mode conversion. The causes of signal integrity issues in deskew methods comes from the inductive nature of the bend. Fortunately, it does not cause noticeable issues to the performance of the channel. For this reason, if it is not possible to use the optimum size other sizes can be considered.

Via types: The frequency domain characterization and TDR show that micro-vias are predominantly inductive whereas buried vias are capacitive. The SERDES simulations show that the long transitions from L2-Bottom have a strong impact on a PAM4 signal and, when chained with other challenges such as crosstalk, they can completely break the transmission.

Trace-to-via transitions: The main conclusion extracted from this study is that the differential pair should split up as close to the pad as possible. Otherwise, the via transition becomes even more inductive and degrades the signal eye opening.

Near-antipad: Even though these discontinuities are noticeable in the TDR results, the rest of the analysis indicates that they do not create any signal integrity issues even with the most challenging case of 0.7 mm pitch.

Pin-field escape: The Broadside Turn is recommended for similar reasons as with trace-to-via transitions. The results show that the differential pair should be split very close to the via. This design benefits with less mode conversion and better performance overall.

Crosstalk: Results show that NEXT crosstalk is more severe than FEXT and it is recommended to only have signals flowing in the same direction in the same layer.

THE SECOND PART

Open-loop multi-feature interconnect: This section shows that the Ideal and the Tough channel are adequate channels for a 50Gbps PAM4 signal even in crosstalk conditions. Indirectly, the stackup used is also validated by the VSR standard.

Close-loop multi-feature interconnect: The simulations and real measurements are consistent with the fact that the Tough channel with crosstalk would not perform correctly. This is also confirmed with a COM analysis.

6 REFERENCES

[1] OIF, "OIF-CEI-04.0," December 29, 2017.

[2] IEEE, "Std 802.3™, IEEE Standard for Ethernet," 2018.

[3] Intel, "AN 835: PAM4 Signaling," 2019.

[4] Y. Kayano, "Identifying dominant factor of imbalance component and EM radiation from differential-paired lines with serpentine equi-length," 2014.

[5] J. Lou, "Intra-pair Length Matching by Asymmetric Dual Bend to Reduce Mode Conversion," in *IEEE*, 2020.

[6] R. J. Sánchez-Mesa, "A Novel High-Performance Length Matching Element for High-Speed Interconnect Differential Channels," in *IEEE*, 2018.

[7] S. Jin, "Meander Line Optimization for High-Speed Differential Transmission Lines," in *IEEE*, 2020.

[8] C. Ye, "Via pattern design and optimization for differential signaling 25Gbps and above," in *IEEE*, 2016.

[9] J. Tang, "A Comprehensive Signal Integrity Study of Differential Pairs Routed within a PCB Via Field," in *IEEE*, 2018.

[10] M. J. Degerstrom, "Automating Pin Field Modeling for SerDes Channel Simulations," in *IEEE*, 2011.

[11] S. H. Hall, "Crosstalk," in *Advanced Signal Integrity for High-Speed Digital Design*, 2009, p. 145.

[12] S. McMorrow, "Trace Design for Crosstalk Reduction," Samtec, 2020. [Online]. Available: https://vimeo.com/432923374.

[13] Intel, *Stratix10TX E-tile IBIS-AMI Userguide V2.0,* 2019.

[14] S. H. Hall, "Properties of Physical S-Parameters," in *Advanced Signal Integrity for High-Speed Digital Designs*, 2009, p. 406.