

DesignCon 2010

Required Technologies for Transmission of 4 x 25 Gb/s Over a Copper Backplane

Jan De Geest, PhD, FCI

jan.degeest@fci.com

Stefaan Sercu, PhD, FCI

stefaan.sercu@fci.com

Abstract

In the quest to transmit higher data rates over copper backplanes, various standards organizations such as the IEEE 802.3ba committee are proposing extensions to existing specifications. One possible path to achieving transmission at 100 Gb/s would be to move to a 4 x 25 Gb/s implementation. New component requirements for the channel must be defined to accommodate these higher signal speeds. A new approach is presented whereby a theoretically perfect channel is first characterized to establish the feasibility of transmitting a data rate of 25 Gb/s over copper backplanes. Subsequently, models of real components are inserted into the channel, and the resultant performance degradation is observed. The analysis yields a set of requirements for key individual components including the PCB's, connectors, connector footprints and transceivers.

Authors biographies

Jan De Geest was born in Gent, Belgium on July 30, 1971. He received the degree in electrical engineering from the University of Ghent, Belgium in 1994 and the degree in supplementary studies in aerospace techniques from the University of Brussels, Belgium in 1995. From September 1995 to December 1999 he worked as a research assistant at the Department of Information Technology (INTEC) of the University of Ghent, where he received the PhD degree in electrical engineering in 2000. Since January 2000 he has been working for FCI ELX in 's-Hertogenbosch, The Netherlands. His work focuses on the design, modeling and optimization of high-speed connectors and interconnection links. He has authored and co-authored over 25 technical papers in international journals and international conference proceedings.

Stefaan Sercu was born in Ieper, Belgium, on February 6, 1969. He received the degree in electrical engineering from the University of Ghent in 1992. From 1992 to 1998, he worked as a research assistant at the Department of Information Technology (INTEC) of the University of Ghent. His research concentrated on the characterization and modeling of high-speed connectors and interconnections. Since 1998 he works for FCI Electronics in 's-Hertogenbosch, The Netherlands. In 2002 he received the PhD degree in electrical engineering from the University of Ghent. Currently he is leading the FCI Electronics Signal Integrity R&D team. He has authored and co-authored over 25 technical papers in international journals and international conference proceedings.

1 Introduction

In order to address the ever-increasing need for network bandwidth, the IEEE 802.3ba Gigabit Ethernet workgroup/taskforce is working to define the specifications for an aggregate data throughput of 40 Gb/s and 100 Gb/s. Both throughputs are realized by implementing a PHY with respectively 4 and 10 parallel lanes, each working at 10 Gb/s. One logical progression which will likely be considered by some standards organizations would be to move to a 4 x 25 Gb/s implementation which would reduce the density requirements and routing complexity. However, such an implementation would not be straightforward, as it would be fraught with many signal-integrity and other challenges that would need to be solved. Such challenges would include the development and selection of transceiver, PCB, connector, and board termination technologies.

This paper addresses the SI challenges inherent in such high-speed transmission and proposes requirements for each of the different components in a 25 Gb/s copper backplane link. A new strategy is proposed for deriving component requirements. Rather than using the traditional method of performing simulations with existing components to determine whether a link with these components would work, a structured top-down approach is proposed that makes use of theoretical models of ideal components in order to establish the feasibility of serial data transmission at 25 Gb/s. This methodology is then applied to determine the component requirements to achieve 4 x 25 Gb/s signal transmission over a copper backplane.

The analysis starts with a perfect link consisting of two ideal connectors with ideal footprints, a backplane and two component boards, all with no losses and no crosstalk. One by one, the perfect components are replaced by real (non-perfect) components including lossy PCB's, impedance mismatches such as those caused by footprints and transceiver parasitics, and crosstalk arising in footprints and connectors. The performance degradation caused by each of the incrementally-introduced imperfections is studied in combination with varying signal-conditioning schemes such as FFE and DFE. While still acknowledging that the link performance can be affected by interdependencies between the components, the analysis yields a set of performance requirements for each of the individual components with a focus on connector and connector footprint requirements. These requirements could be dependent on the level of signal conditioning that is applied in both transmitter and receiver.

2 How do we assess 25 Gb/s transmission ?

Currently no transceivers are commercially available that can transmit serial data at 25 Gb/s. So 25 Gb/s performance can only be verified through link simulations, not through actual system measurements. Different committees, like the IEEE 802.3ba and OIF CEI-25G-LR committees, are currently addressing the channel requirements and guidelines that may make it possible to transmit 25 Gb/s over a backplane link. These efforts are currently still in progress and no validated 25 Gb/s specification for copper backplane links is available yet.

Initial work we did on 25 Gb/s developments has been based on the IEEE 802.3ap spec, which we extrapolated to 25 Gb/s. Meanwhile the OIF initiative has been launched which

focuses on 25 Gb/s transmission over copper backplanes. In this paper we will use the parameters from the informative guidelines from the IEEE 802.3ap spec extrapolated to 25 Gb/s and from the OIF CEI-25G-LR spec as requirements for 25 Gb/s transmission. Besides using these informative guidelines we will also perform bit-error-rate (BER) simulations for which we assume a BER of 1e-15 at the receiver to be required for successful 25 Gb/s transmission.

2.1 IEEE 802.3ap spec extrapolated to 25 Gb/s

The IEEE 802.3ap 10GBASE-KR spec provides informative channel requirements for transmission of 10G Ethernet signals over up to 1 meter of enhanced FR4 PCB and 2 connectors with a BER of 1e-12 [1]. It defines high confidence acceptance regions for five parameters: Insertion Loss (IL), Attenuation (ATT), Insertion Loss Deviation (ILD), Return Loss (RL), and Insertion Loss to Crosstalk Ratio (ICR).

	IEEE 802.3ap	Extrapolation to 25 Gb/s
ATT	$ATT \geq A_{\min} = -20 \times \log_{10}(e) \times (2e-5 \times \sqrt{f} + 1.1e-10 \times f + 3.2e-20 \times (f)^2 - 1.2e-30 \times (f)^3)$ <i>between 1 and 6 GHz</i>	$ATT \geq A_{\min} = -20 \times \log_{10}(e) \times (2e-5 \times \sqrt{f/2.5} + 1.1e-10 \times f/2.5 + 3.2e-20 \times (f/2.5)^2 - 1.2e-30 \times (f/2.5)^3)$ <i>between 1 and 15 GHz</i>
IL	$ATT \geq A_{\min} - 0.8 - 2e-10 \times f$ <i>between 0.05 and 6 GHz</i> $ATT \geq A_{\min} - 0.8 - 2e-10 \times f - 1e-8 \times (f - 6e9)$ <i>between 6 and 15 GHz</i>	$ATT \geq A_{\min} - 0.8 - 2e-10 \times f$ <i>/2.5 between 0.05 and 15 GHz</i> $ATT \geq A_{\min} - 0.8 - 2e-10 \times f/2.5 - 1e-8 \times (f - 15e9)/2.5$ <i>between 15 and 37.5 GHz</i>
ILD	$ ILD \leq 1.0 + 0.5e-9 \times f$ <i>between 1 and 6 GHz</i>	$ ILD \leq 1.0 + 0.2e-9 \times f$ <i>between 1 and 15 GHz</i>
RL	$RL \geq 12$ <i>btwn 50 and 275 MHz</i> $RL \geq 12 - 6.75 \times \log_{10}(f/275e6)$ <i>btwn 275 MHz and 3 GHz</i> $RL \geq 5$ <i>btwn 3 and 10.3215 GHz</i>	$RL \geq 12$ <i>btwn 50 and 687.5 MHz</i> $RL \geq 12 - 6.75 \times \log_{10}(f/687.5e6)$ <i>btwn 687.5 MHz and 7.5 GHz</i> $RL \geq 5$ <i>btwn 7.5 and 25.80375 GHz</i>
ICR	$ICR \text{ fit} \geq 23.3 - 18.7 \times \log_{10}(f/5e9)$ <i>between 1 and 5.15625 GHz</i>	$ICR \text{ fit} \geq 23.3 - 15.15 \times \log_{10}(f/12.5e9)$ <i>between 1 and 12.890625 GHz</i>

Table 1: Extrapolation of 10 Gb/s IEEE 802.3ap limits to 25 Gb/s. Adjusted values are shown in bold blue font.

The IEEE 802.3ap spec is extrapolated to extend the frequency range of interest from the 6 or 15 GHz in the spec all the way up to frequencies more commensurate with the target data rate of 25 Gb/s. Extrapolation is done by scaling most of the high end frequencies in the spec by a factor of 2.5 (see table 1). Figure 1 shows the original and extrapolated limits for the IL, ATT and ILD. Figure 2 shows the original and extrapolated limits for the RL and ICR.

2.2 OIF CEI-25G-LR spec

The OIF CEI-25G-LR specification is the only existing specification presently under consideration that applies to a data rate of 25 Gb/s. It aims at driving 25 Gb/s over up to 68.6 cm (27") of PCB and up to 2 connectors [2]. It poses limits on five parameters:

Insertion Loss (IL), Attenuation (ATT), Insertion Loss Deviation (ILD), Return Loss (RL), and Integrated Crosstalk Noise (ICN). Figure 3 shows the limits for the IL, ATT and ILD. Figure 4 shows the limits for the RL and ICN.

Below approximately 12.5 GHz the extrapolated IEEE spec is more stringent than the OIF CEI-25G-LR spec for the IL. Above 12.5 GHz the OIF spec is more stringent. The extrapolated IEEE limit for the ILD only goes up to 15 GHz, while the OIF limit for the ILD goes up to 25 GHz. Below 15 GHz the extrapolated IEEE limit for the ILD is more stringent than the OIF spec. The RL limit in the extrapolated IEEE spec is more stringent than the RL limit in the OIF CEI-25G-LR spec.

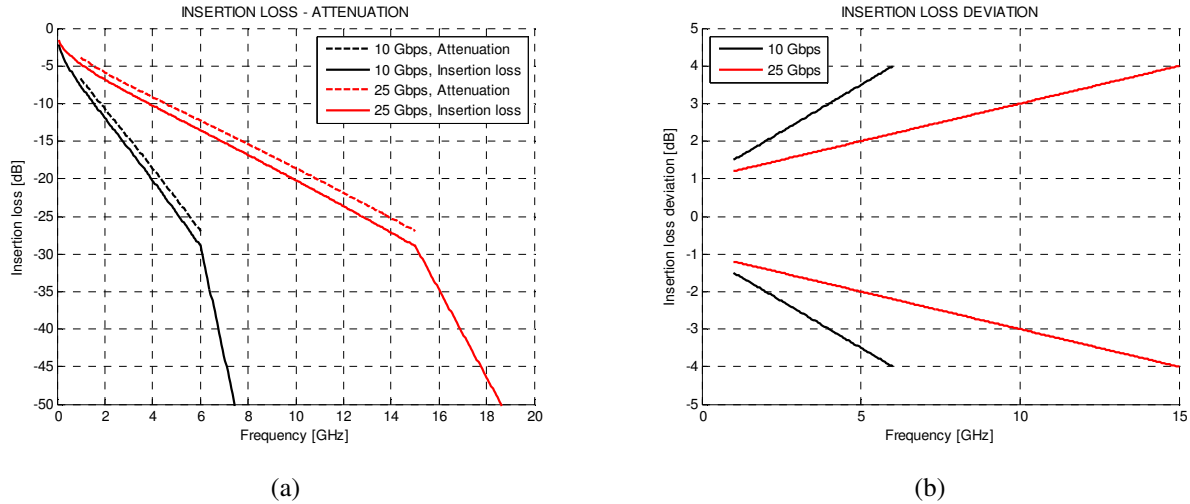


Figure 1: 10 Gb/s IEEE802.3ap spec and extrapolation to 25 Gb/s: IL/ATT (a) and ILD (b).

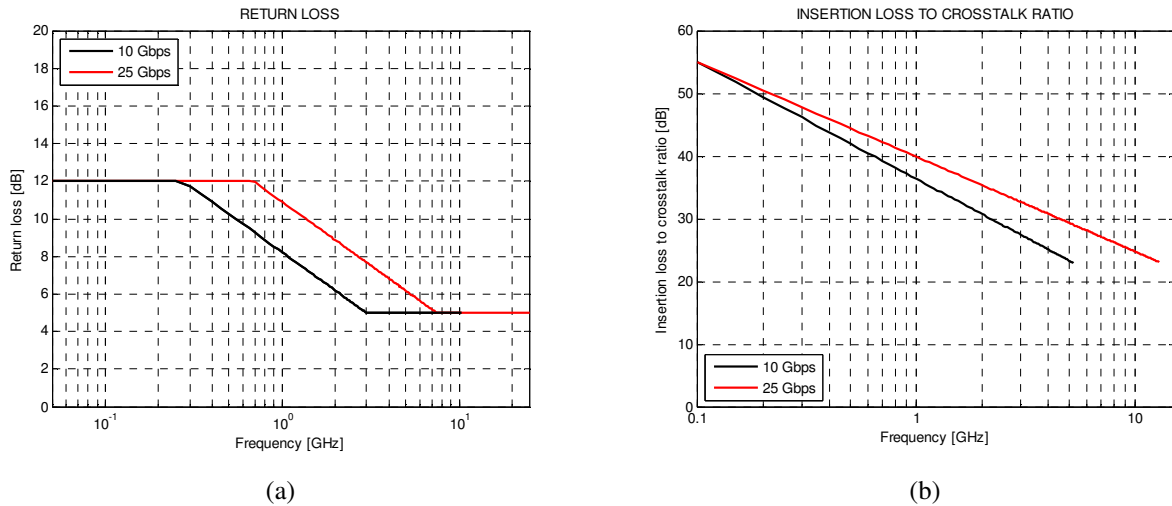


Figure 2: 10 Gb/s IEEE802.3ap spec and extrapolation to 25 Gb/s: RL (a) and ICR (b).

2.3 BER simulations

The extrapolated IEEE specification and the OIF CEI-25G-LR specification pose limits on the passive channel performance. Both specs are informative specs, implying that they provide high confidence regions where, if the passive performance figures (IL, ILD, RL and ICR or ICN) fall within these regions, there is a high confidence that the considered

link will be able to support 25 Gb/s serial data transmission. Some links that fail the passive channel specifications might still be able to support 25 Gb/s transmission.




Figure 3: OIF CEI-25G-LR spec: IL/ATT (a) and ILD (b).

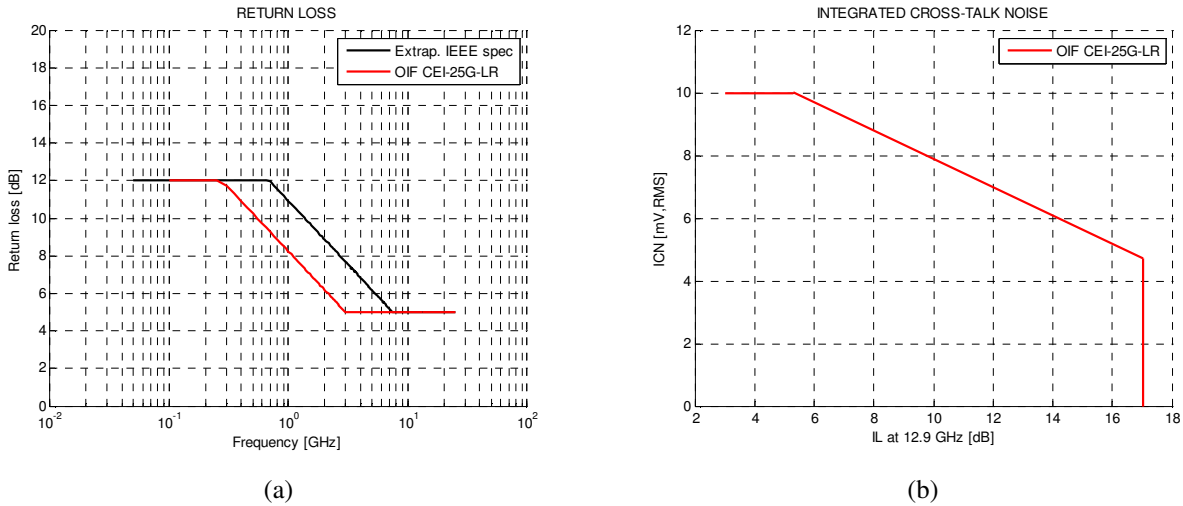


Figure 4: OIF CEI-25G-LR spec: RL (a) and ICN (b).

Ultimately BER measurements or simulations need to be performed as a final system verification to make sure a 25 Gb/s link really works. As there are no transceivers commercially available yet that can transmit at a data rate of 25 Gb/s, it is necessary to evaluate the performance of a link strictly through simulations rather than measurements. Therefore BER simulations will be run in combination with varying signal-conditioning schemes such as FFE and DFE and taking into account jitter and transceiver parasitics. The result of such a BER simulation is a BER eye (see figure 8a) and a bath-tub curve (BTC) (see figure 8b), which show the eye closure as a function of the BER.

3 Backplane link description

The interconnection link that is studied in this paper is a backplane link, consisting of two daughter cards connected through two right-angled backplane connectors to a backplane

(figure 5). The total interconnection link model is obtained by concatenating the S-parameter models of the individual components. In total there are 11 sub-models in the total channel model: the transmitter chip parasitics, the first daughter card traces, the connector footprint on the first daughter card, the first mated connector, the first connector footprint on the backplane, the backplane traces, the second connector footprint on the backplane, the second mated connector, the connector footprint on the second daughter card, the second daughter card traces and the receiver chip parasitics.




Figure 5: Backplane link.

	Backplane	Daughter cards
Trace width	6 mil	6 mil
Trace thickness	35 um	35 um
Trace separation	8 mil	8 mil
Routing width	20 mil	20 mil
Trace length	10 cm (4") to 60 cm (24")	10 cm (4")

Table 2: Backplane and daughter cards trace dimensions.

The daughter cards consist of 12 layers: 4 signal layers, 6 ground layers and 2 component layers. The build-up is shown in figure 6a. The thickness of all metal layers is 35 um. The thickness of the dielectric layers in the daughter cards is 220 um, except for the two outer layers and the middle dielectric layer which are 100 um thick. This results in a total thickness of the daughter cards of 2.4 mm. The backplane consists of 22 layers: 8 signal layers and 14 ground layers. The build-up is shown in figure 6b. The thickness of all metal layers is 35 um. The thickness of the dielectric layers in the backplane is 220 um, except for the three middle dielectric layers which are 550 um thick. This results in a total thickness of the backplane of 6.4 mm. Back-drilling is applied to the backplane and the daughter cards. The trace dimensions on the backplane and the daughter cards are given in table 2. The traces on the daughter cards are each 10 cm long. The trace length on the backplane varies between 10 and 60 cm. This results in a maximum total interconnection length of 80 cm (approximately 32"). The intention is to demonstrate successful 25 Gb/s transmission over a length of 27", which is in line with OIF CEI-25G-LR spec.




Figure 6: PCB build-ups: daughter cards (a) and backplane (b).

In the link simulations that follow later in this paper we will look at channels routed on different layers in the backplane and daughter cards. In a first channel, which will be referred to as “channel B” the traces are routed on the bottom signal layers in both the backplane and daughter cards, i.e. the layers highlighted in blue in figure 6. In a second channel, which will be referred to as “channel O” the traces are routed on the layers highlighted in red in figure 6.

4 Transceiver parameters

Driver and receiver parasitics are taken into account by adding a capacitance to ground at both ends of each trace. All traces are terminated by a 50 Ohms resistor to ground in parallel with a 0.15 pF chip capacitance. Figure 7 shows the input RL of the transmitter and the receiver compared with the OIF CEI-25G-LR spec. Using the above mentioned value for the chip capacitance the input RL meets the OIF CEI-25G-LR spec.

The driver transmits a NRZ bit stream with a bit rate of 25 Gb/s and a 10-90 % rise time of 15 ps. The driver has a peak-to-peak deterministic jitter of 6 ps (0.15 UI) and a random jitter with $\sigma = 0.30$ ps. Figure 8a shows the BER eye and figure 8b shows the BTC for a perfect lossless channel, but taking into account transceiver parasitics and driver jitter. The BTC has a width of about 0.7 UI for a BER of $1e-15$, which is in line with both the extrapolated IEEE spec and the OIF CEI-25G-LR spec.

Some kind of signal conditioning will be required to open the BER eyes at 25 Gb/s. In our link simulations the driver will use a baud-spaced pre-emphasis (PE) filter with up to 5 taps and the receiver will use a baud-spaced Decision-Feedback Equalizer (DFE) with up to 7 taps. For successful transmission we will assume a BER of $1e-15$ is required.




Figure 7: Transmitter and receiver input RL vs. OIF CEI-25G-LR spec..

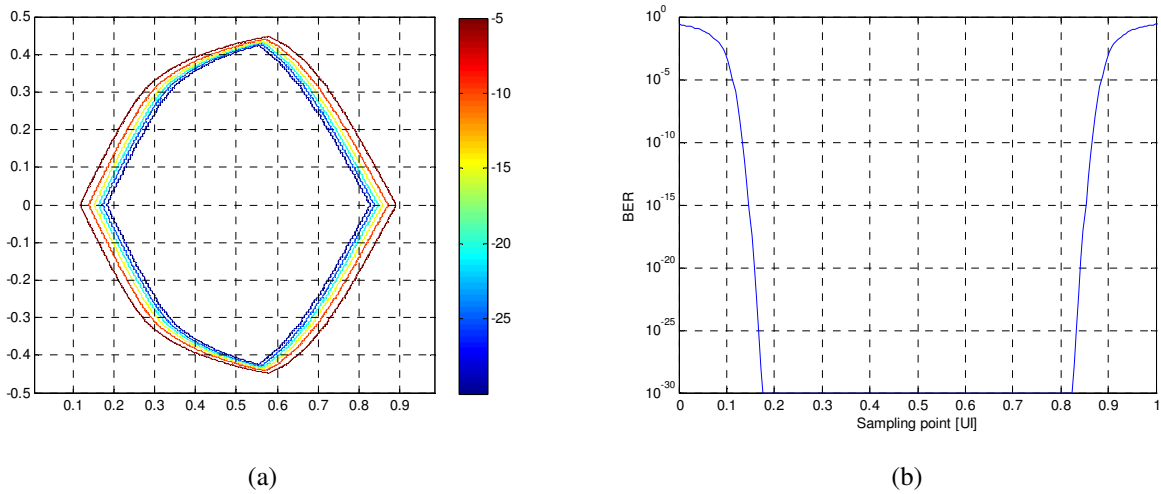


Figure 8: BER eye (a) and BTC (b) in case of a perfect lossless channel.

5 Board material

We start the analysis with a link consisting of only the backplane and the two component cards. We assume there is no crosstalk in the boards, and since the PCB's all have a 100 Ohm differential impedance there are no impedance mismatches either. The only losses in the channel are the PCB losses. PCB losses consist of dielectric losses and conductor losses. The dielectric losses are determined by the loss tangent (D_f) of the board material. The conductor losses are determined by the conductivity ($5.8e7$ S/m for copper) and the shape and dimensions of the metal traces. Different board imperfections which have a negligible effect at lower bit rates start to play an important role at bit rates of 25 Gb/s and higher. These board imperfections generate additional losses and these have to be taken into account. An important source of additional losses at high bit rates is the conductor surface roughness. This can be taken into account by using the well known Hammerstad and Jensen model (see section 5.1.1). Other effects like over-etching and tolerances on material parameters and board dimensions can be taken into account by

defining a certain ΔIL (which is a function of the frequency) that is added to the other board losses.

Using the PCB only model we can determine what material we should use in order to pass the extrapolated IEEE spec and the OIF CEI-25G-LR spec (for a given length). Furthermore, for a given material and a given length we can determine what type of signal conditioning is required to get a BER of $1e-15$ at the receiver.

5.1 Board imperfections

5.1.1 Surface roughness

Copper foils used in high-volume PCB's are not perfectly smooth, but have microscopically small dips and grooves on their surface. Often the foils are roughened to improve the adhesion between the foils and the dielectric layers. At high speeds the surface roughness can be of the same order as the skin depth, resulting in an increase in the conductor losses at a rate higher than the square root of the frequency.

The surface roughness is often quantified by a root mean square (RMS) value of the height of the surface discontinuities. This surface roughness RMS value depends on the type of copper foil. Typical surface roughness RMS values are given in table 3 [3]. In this paper we will assume a surface roughness RMS value of $0.81 \mu\text{m}$.

Copper foil type	RMS roughness [μm]
Microstrip	1.61
Stripline	0.81
Rolled copper (near-smooth)	0.32

Table 3: Surface roughness RMS values for different types of copper foils.

The losses caused by the surface roughness can be taken into account by defining a loss factor K_{SR} which is a function of the surface roughness RMS value:

$$K_{SR} = 1 + \frac{2}{\pi} \arctan\left[0.7\Delta^2 \sigma \omega \mu_0 \mu_r\right]$$

where Δ is the surface roughness RMS value, σ is the conductivity of the metal ($5.8e7$ S/m for copper), ω is the angular frequency, μ_0 is the permeability of free space ($4\pi \cdot 1e-7$ N/A²) and μ_r is the relative permeability of the conductor (1 for copper). The value of K_{SR} equals to 1 at DC, and it asymptotically approaches 2 at very high frequencies.

In a transmission line RLGC model the unit-length R-matrix is multiplied by this loss factor K_{SR} to take into account the surface roughness:

$$R(\omega) = R_{REF} K_{SR} (1 + j) \sqrt{\omega / \omega_{REF}}$$

The square root frequency dependency accounts for the skin effect.

Figures 9 and 10 show the IL versus frequency of a 27" trace for materials with increasing dielectric losses. These results are for a surface roughness RMS value of 0.81

um. Clearly the surface roughness has a big impact on the trace losses. At 10 GHz the surface roughness adds about 5 dB to the trace losses, at 25 GHz almost 10 dB of losses are added by taking into account the surface roughness. A trace with surface roughness RMS value of 0.81 um in a material with a loss tangent of 0.005 has approximately the same performance as a perfectly smooth trace in a material with a loss tangent of 0.008.

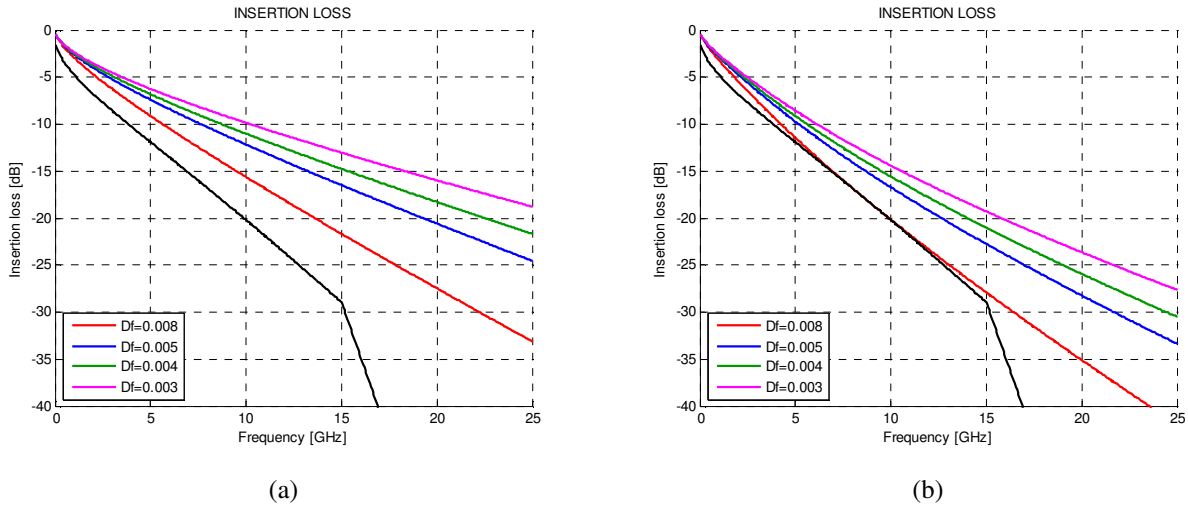


Figure 9: IL vs. frequency of a 27” trace on materials with increasing dielectric losses and compliance with extrapolated IEEE spec: without taking into account surface roughness (a) and with surface roughness included (b).




Figure 10: IL vs. frequency of a 27” trace on materials with increasing dielectric losses and compliance with CEI-25G-LR spec: without taking into account surface roughness (a) and with surface roughness included (b).

5.1.2 Over-etching and other tolerances on trace dimensions

Conductor traces are usually modeled having a rectangular cross-section. Because of over-etching, in real world applications the cross-section of a trace usually has a tapered or trapezoidal shape. Furthermore, because of tolerances on the dimensions the actual width and height of a trace after fabrication can differ from the nominal values. The tapering of the cross-section and the tolerances on the trace dimensions has two effects:

the impedance of the trace changes, and additional losses can be generated in the trace if the trace width and/or height are reduced.




Figure 11: Schematic representation of a copper trace with over-etching.

Figure 11 shows a schematic representation of a copper trace with over-etching. The trace width at the top of the trace ($W-2D$) is smaller than at the bottom (W). Figure 12 shows the insertion loss and the impedance of a differential pair as a function of over-etching and of tolerances on the nominal dimensions. The nominal trace width W and trace thickness H are 6 mil and 35 μm respectively. The trace separation is 8 mil, the trace length is 27". The board material has a dielectric constant of 3.4 and a loss tangent of 0.003.




Figure 12: Differential insertion loss (a) and differential impedance (b) of a 27" long trace as a function of over-etching and of dimensional tolerances.

At 10 GHz the insertion loss can increase with roughly 1 dB because of over-etching and tolerances on the trace dimensions. At 20 GHz the insertion loss can increase with roughly 2 dB. The differential impedance varies between 93.6 Ohms and 103.9 Ohms. The additional losses caused by the over-etching could easily be compensated by using a slightly larger width for the trace, although this might increase the crosstalk between adjacent traces.

5.2 Required board material for 25 Gb/s transmission

Using the PCB only model we can determine what material we should use in order to pass the extrapolated IEEE spec and the OIF CEI-25G-LR spec (for a given length). The connectors and footprints will also generate some additional losses, however the major contributors to the overall channel insertion loss are the backplane and the daughter

cards. The connectors and connector footprints will predominantly be responsible for the impedance mismatches and crosstalk. Figures 9 and 10 show the compliance with the IL limit in the extrapolated IEEE spec and in the OIF CEI-25G-LR spec, for different board materials (different loss tangents) and with and without surface roughness included.

As can be seen from figure 9b, a material with a loss tangent of 0.004-0.005 (e.g. Rogers 4350 material) is required to pass the extrapolated IEEE IL limit for a 27" trace with some margin for the additional losses caused by trace tolerances and for the connector and footprints that will be added later to the link model. A board material with a loss tangent of 0.003-0.004 (e.g. Megtron6) may be required to meet the OIF CEI-25G-LR IL limit with some margin (see figure 10).




Figure 13: BER eye opening vs. loss tangent for a 27" PCB only link: BTC width (a) and BER eye height (b) for BER = 1e-15.

For a given material and a given length we can determine what type of signal conditioning is required to get a BER of 1e-15 at the receiver. Figure 13 shows the BTC width and the BER eye height for a BER = 1e-15 as a function of the loss tangent for a 27" PCB only link. Note that the results shown in figure 13 include the effect of surface roughness, the multiple reflections caused by the transceiver parasitics and the driver jitter as defined in section 4. Without signal conditioning the BER eye is completely closed, even for very low loss materials. For low loss materials a 3-taps PE filter and a 5-taps PE filter have approximately the same performance. For higher loss materials (loss tangent of about 0.005 and higher) a 5-taps PE filter provides some additional BER eye opening compared to a 3-taps PE filter. A 7-taps DFE provides additional margin over a 5-taps DFE, regardless of the loss tangent.

Note that the results shown in figure 13 are for a PCB only model. The crosstalk and impedance mismatches generated by the connectors and their footprints are not yet included. These will be studied in sections 6 and 7.

6 Connector and footprint impedance

When the connectors and footprints are added to the channel model these introduce reflections caused by impedance mismatches and crosstalk. In this section we will investigate how much impedance mismatch can be tolerated in the connectors and the footprints to still pass the RL and ILD limits in the extrapolated IEEE spec. In section 7 we will investigate how much crosstalk can be tolerated in the connectors and in the footprints.

The connectors and footprints are modeled as short lossless transmission lines with varying differential impedance, and without crosstalk. The transmission lines representing the connectors have a delay of 180 ps, which corresponds to the delay of the longest pair in the connector used in section 8. The transmission lines representing the backplane and the daughter card footprints have delays of 20 ps and 13 ps respectively. The footprint impedance is swept between 60 and 160 Ohms, and the connector impedance is swept between 80 and 130 Ohms (for a 25 ps 10-90% rise time).

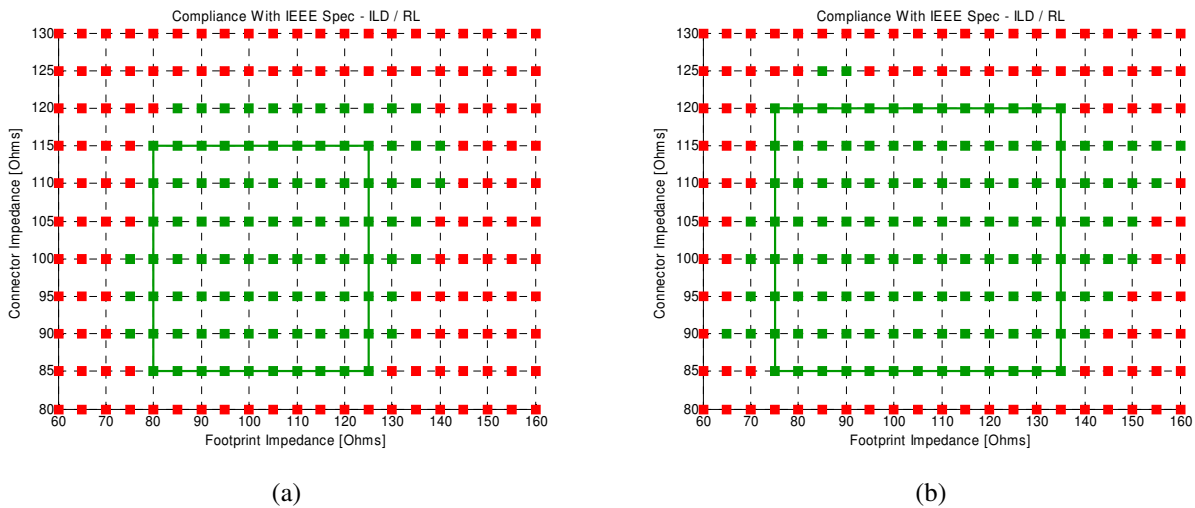


Figure 14: Compliance with extrapolated IEEE spec as a function of the connector and footprint impedance: for a 10 cm backplane (a) and for a 50 cm backplane (b).

Figure 14 shows the compliance with the ILD and RL limits in the extrapolated IEEE spec as a function of the connector impedance and the footprint impedance. Results are shown for a short (10 cm) and a long (50 cm) backplane trace. A loss tangent of 0.003 is used for all boards. A green square in figure 14 means that that particular combination of connector impedance and footprint impedance meets the extrapolated IEEE spec, a red square means the spec is not met. To be compliant with the extrapolated IEEE spec, for short tracks the footprint impedance must be between approximately 80 and 125 Ohms and the connector impedance must be between approximately 85 and 115 Ohms. For long tracks the footprint impedance and the connector impedance must be between approximately 75 and 135 Ohms and between approximately 85 and 120 Ohms respectively. Long traces can tolerate more impedance mismatch than short traces, because the multiple reflections are dampened by the losses in the trace. In the footprints

the impedance will be different depending on the routing layer. It can be beneficial to route the longest tracks on the layers which have the worst footprint impedance.

7 Connector and footprint crosstalk

The main crosstalk sources in the interconnection link are the connectors and the connector footprints. The crosstalk in the connector footprints includes crosstalk between the vias, crosstalk in the fan-out region, crosstalk between traces routed through a footprint and the footprint itself, and crosstalk between different signal layers through the ground plane antipads.

The total crosstalk at the receiver will be determined by a large number of parameters: the NEXT in the connector and in the footprints, the FEXT in the connector and in the footprints, the lengths of the victim and aggressor traces on the backplane and on the daughter cards, the number of channels, etc. It will also depend on the TX/RX configuration that is used (see section 8.3). Because of this large number of parameters it is difficult to define crosstalk requirements for individual components like connectors or connector footprints.




Figure 15: Model for deriving connector and connector footprint crosstalk requirements.

We can derive a rough estimate of the connector and connector footprint crosstalk requirements by using the simplified model shown in figure 15. In this model a number of assumptions are made. First of all we assume that all the crosstalk in the system is generated in the connectors and in the connector footprints and no crosstalk is generated in the boards. In the model each of the connectors together with its backplane and daughter card footprints is seen as a single block where crosstalk can occur. Furthermore we also assume that all NEXT channels can be combined into a single NEXT source and all FEXT channels can be combined into a single FEXT source. Using these assumptions and neglecting the multiple reflections caused by impedance mismatches in the system (which should be as low as possible anyhow for 25 Gb/s transmission) we can write the following expressions for the victim channel (the green channel in figure 15), the NEXT channel (the red channel) and the FEXT channel (the blue channel):

$$H_T \approx H_C^2 H^2 H_B$$

$$H_N \approx H_C^2 N (1 + H^2 H_B^2)$$

$$H_F \approx 2H_C^2 H F H_B$$

H_T , H_N , H_F , H_C and H_B are the transfer functions of the victim channel, the NEXT channel, the FEXT channel, the daughter cards and the backplane respectively. H , N and F are respectively the transmission, NEXT and FEXT transfer functions of the connector + connector footprints combination (as a function of frequency).

Using the above equations we can write the ratio between the total crosstalk at the receiver and the victim signal as:

$$ICR^{-1} \approx \frac{H_N + H_F}{H_T} \approx NH_B + \frac{N}{H^2 H_B} + 2 \frac{F}{H}$$

Based on the above equation and keeping in mind the assumptions that were made for deriving this equation we can note a number of things. First, the transfer function of the daughter cards H_C is not present in the equation for the ICR. This means that if the length of the daughter cards is changed the ICR should approximately be the same, which also implies that the crosstalk requirements for compliance with the extrapolated IEEE spec are independent of the daughter cards.

In an all FEXT configuration, i.e. for $N = 0$, the ICR is also independent of the transfer function of the backplane H_B , and therefore of the length of the backplane. The channel crosstalk requirements as set by the extrapolated IEEE spec can be directly translated to component crosstalk requirements. We can say that in an all FEXT configuration the fitted ICR of the connector + footprints combination should be 6 dB (i.e. a factor 2) better than the channel ICR limit (see table 1 or figure 2), so better than 61 dB at 0.1 GHz and better than 29.3 dB at 12.5 GHz.

In an all NEXT configuration, i.e. for $F = 0$, the situation is a bit more complex. For very short or lossless backplanes we can write $ICR^{-1} \approx 2N$. As the backplane length increases one component of the NEXT will be attenuated more than the other and will become less important. On the other hand the second component becomes more important because the victim channel is attenuated more when the backplane length increases. Short channels can tolerate more NEXT than longer channels. Regardless of the backplane length it is fair to say that in an all NEXT configuration the ICR should at least comply with the same limits as in an all FEXT configuration, i.e. the fitted ICR of the connector + footprints combination should be 6 dB (i.e. a factor 2) better than the channel ICR limit.

8 Link simulations

In this section the theoretical connector and footprint models that were used in sections 6 and 7 are replaced by actual models of FCI's AIRMAX WS™ connector and footprint. Link simulations for different TX/RX configurations are performed as a function of the total interconnection length. The results are compared against the extrapolated IEEE and OIF CEI-25G-LR specs. Finally, BER simulations are performed for different transceiver equalization settings. In these BER simulations the effects of surface roughness, the multiple reflections caused by the transceiver parasitics and the driver jitter are included. The simulations are performed using an FCI proprietary link simulation tool.

8.1 Connector performance

Figures 16 and 17 show the performance of the AIRMAX WS™ connector. Figure 16 shows the differential IL and the differential impedance at a 25 ps rise time. The connector has an insertion loss better than 0.5 dB up to 10 GHz and an insertion loss better than 1.5 dB up to approx. 22 GHz. The 3 dB bandwidth of the connector is approx. 25 GHz. The connector impedance varies between 87 and 109 Ohms and is well in line with the connector impedance requirement determined in section 6 (which is indicated by the black lines in figure 16). Figures 17a and 17b show the differential power-sum NEXT and FEXT respectively taking into account 11 crosstalk aggressors. Both the power-sum NEXT and FEXT are better than -35 dB up to approx. 12 GHz and better than -20 dB up to approx. 28 GHz.

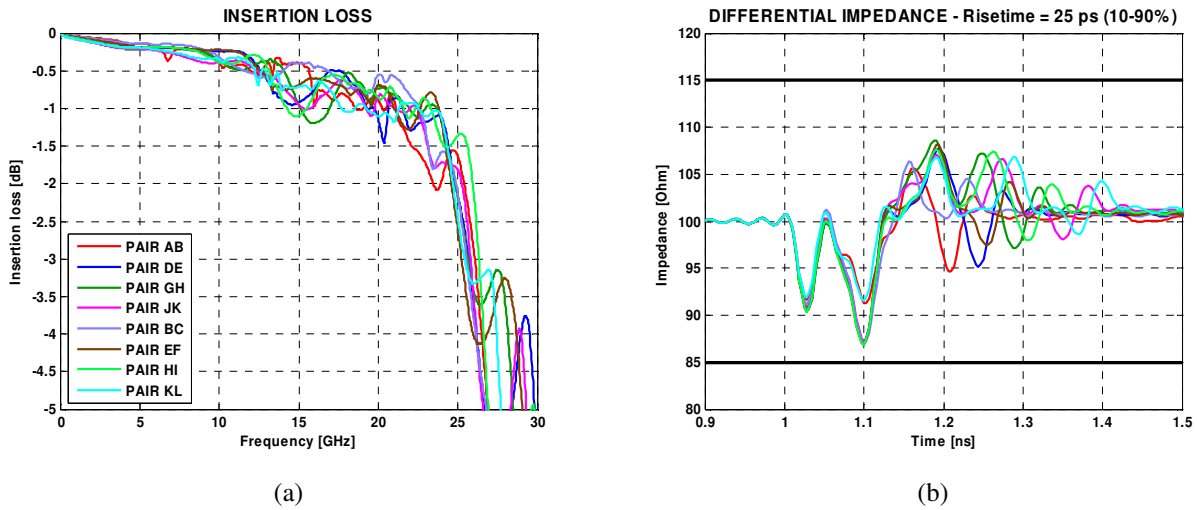


Figure 16: Connector performance: differential IL (a) and differential impedance (b).

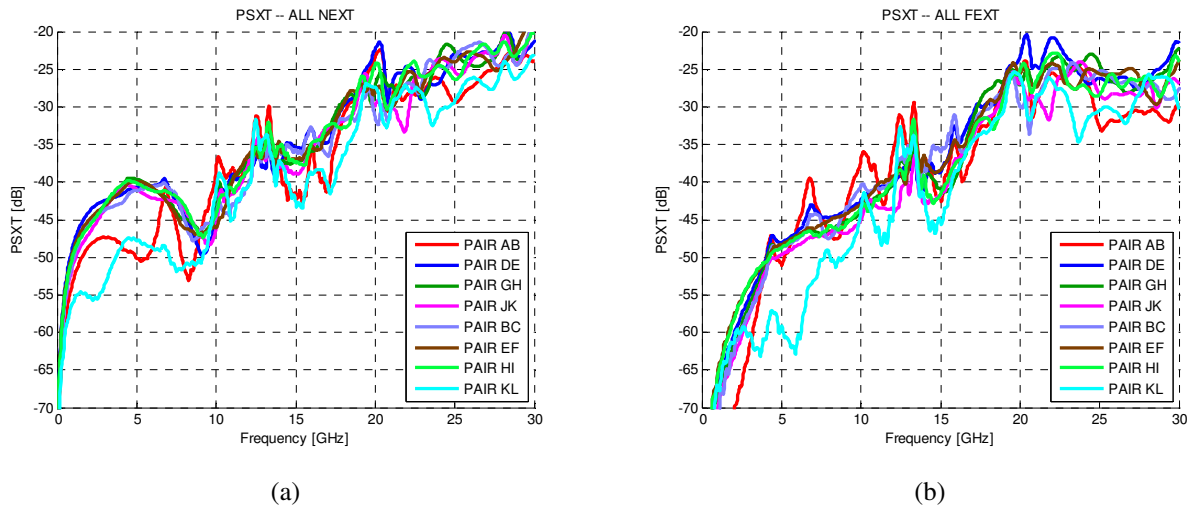


Figure 17: Connector performance: differential power-sum NEXT (a) and FEXT (b).

8.2 Connector footprint performance

Figures 18 and 19 show the differential impedance and the power-sum crosstalk (PSXT) of the backplane and daughter card footprints, for channel B and channel O respectively. The solid and dotted black lines on the impedance plots show the impedance requirement for long and short traces respectively, as determined in section 6. The impedance for channel B is better than the requirement for long traces, and the impedance for channel O is better than the requirement for short traces. For channel B the power-sum NEXT and FEXT in the backplane footprint are better than -35 dB up to approx. 15 GHz and better than -25 dB up to almost 30 GHz. The power-sum NEXT and FEXT in the daughter card footprint up to 30 GHz are better than -33dB and -35dB respectively. For channel O the PSXT in the footprint is better than -40 dB up to 20 GHz and better than -32 dB up to 30 GHz. The crosstalk for channel O is much smaller than for channel B because of the shorter via barrel length corresponding to traces routed in channel O.




Figure 18: Footprint performance for channel B: diff. impedance (a) and diff. PSXT (b).

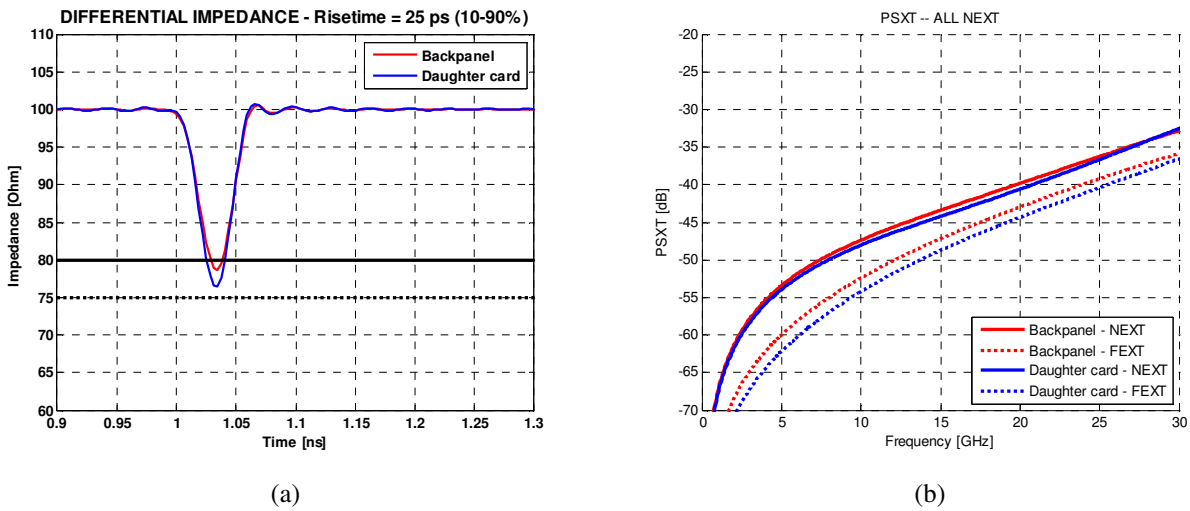


Figure 19: Footprint performance for channel O: diff. impedance (a) and diff. PSXT (b).

8.3 Different TX/RX configurations

The channel crosstalk performance will depend upon the TX/RX configuration that is used in the system. When assigning TX's and RX's to individual differential pairs, different options are available. Depending on where TX and RX pairs are assigned the crosstalk performance of a channel will be different. Figure 20 shows 4 possible TX/RX configurations for a system with 1 victim channel (the RX in the middle highlighted in red) and 11 aggressor channels. By selecting the most beneficial TX/RX configuration the total crosstalk on the victim channel at the receiver can be minimized.

Figures 21 to 24 show the compliance with the extrapolated IEEE spec of a channel with a total length of 27". The results are shown for channel B and channel O. It can first of all be noted that both channels comply with the IL and ILD spec. Also, the IL and ILD are independent of the TX/RX configuration. The ICR is different for different TX/RX configurations. For channel O the fitted ICR complies with the spec for all considered TX/RX configurations. For channel B none of the fitted ICR curves comply with the spec. TX/RX configuration (d) marginally fails the spec.




Figure 20: Different possibilities for TX/RX configurations.

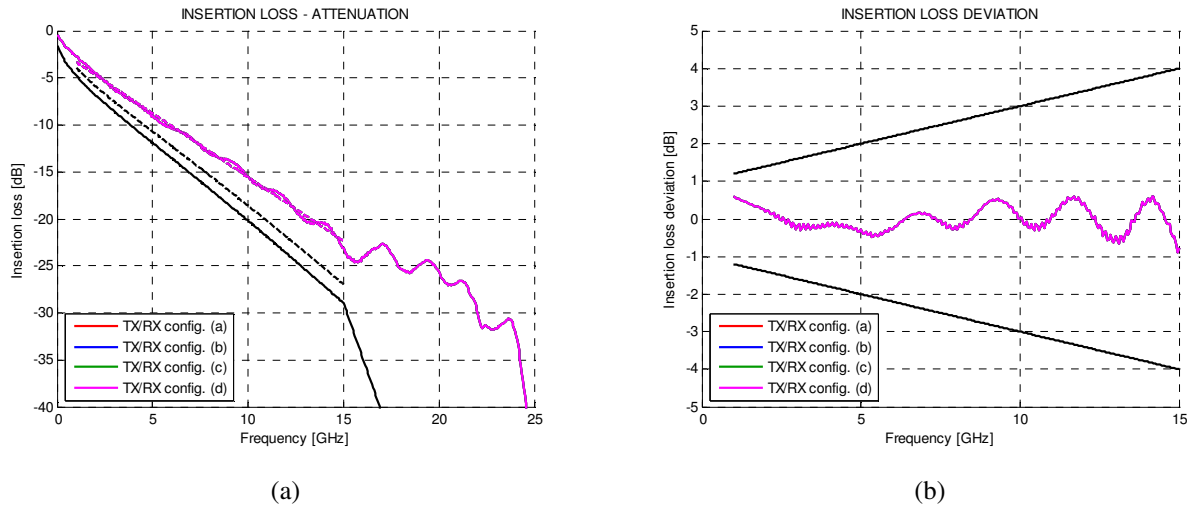
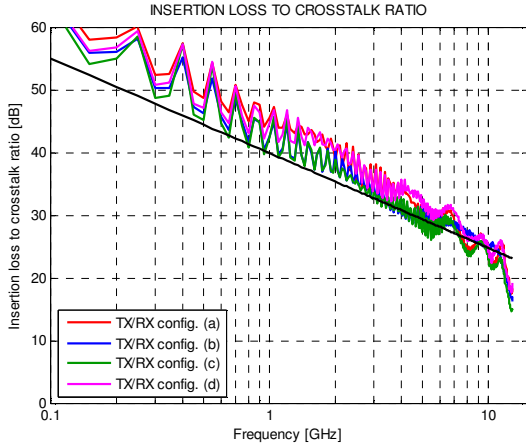
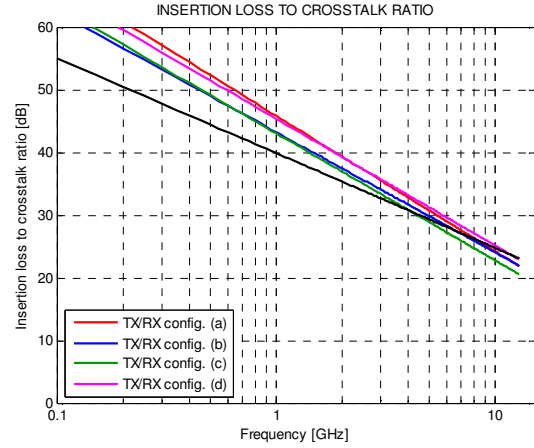


Figure 21: IL (a) and ILD (b) for different TX/RX configurations, $Df = 0.003$, channel B.

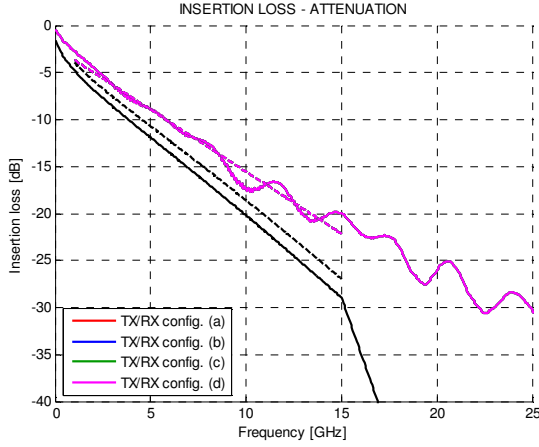


(a)

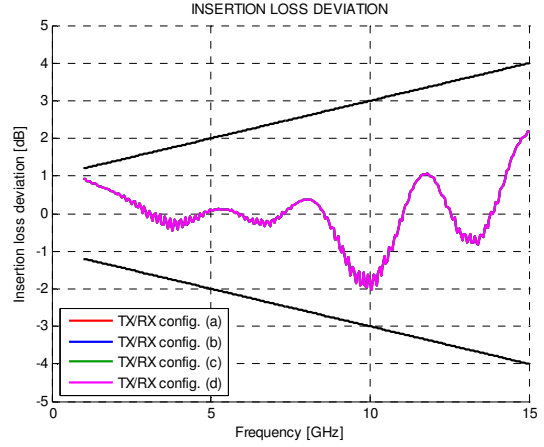


(b)

Figure 22: ICR (a) and fitted ICR (b) for different TX/RX configurations, $D_f = 0.003$, channel B.

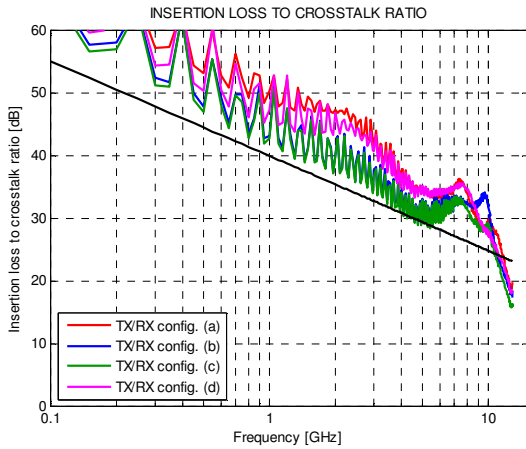


(a)




(b)

Figure 23: IL (a) and ILD (b) for different TX/RX configurations, $D_f = 0.003$, channel O.



(a)



(b)

Figure 24: ICR (a) and fitted ICR (b) for different TX/RX configurations, $D_f = 0.003$, channel O.

Figure 25 shows the PSXT of a connector with 2 footprints, one on the backplane and one on the daughter card, for different TX/RX configurations. Figure 25a shows the results for channel B and figure 25b shows the results for channel O. The black line in figures 25a and 25b represents 2 times the channel ICR limit (i.e. -61 dB at 0.1 GHz and -29.3 dB at 12.5 GHz).

As can be seen in figures 23 and 24, channel O passes the extrapolated IEEE guideline for all considered TX/RX configurations. The PSXT for all these TX/RX configurations is better than 2 times the channel ICR limit, as can be seen in figure 25b. For channel B the PSXT for TX/RX configurations (a) and (d) is better than 2 times the channel ICR limit, while for TX/RX configurations (b) and (c) this is not the case for some parts of the frequency range of interest (see figure 25a). In figure 22 it can be noticed that for channel B all TX/RX configurations marginally fail the ICR guideline, with TX/RX configurations (a) and (d) having the best performance.

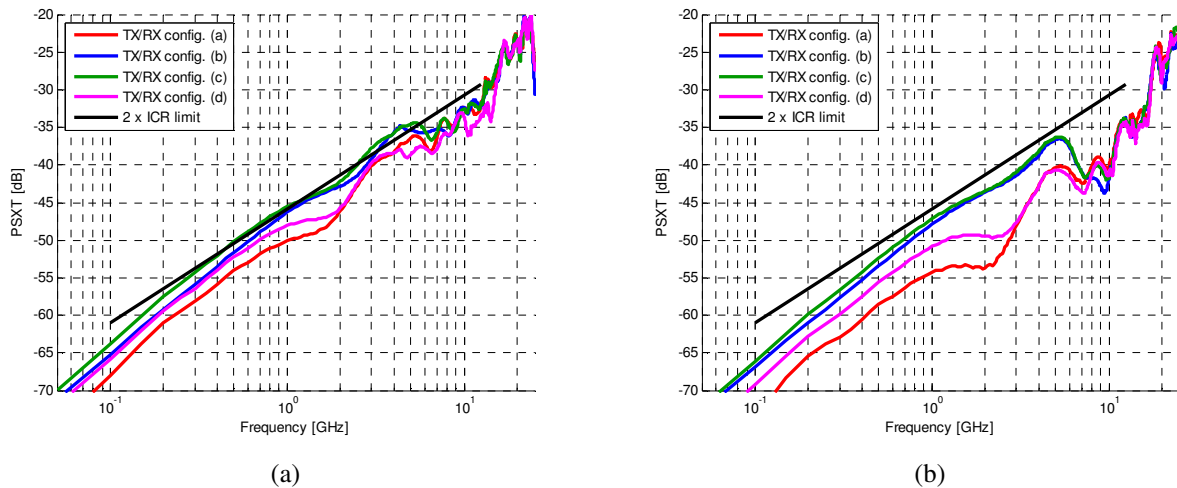
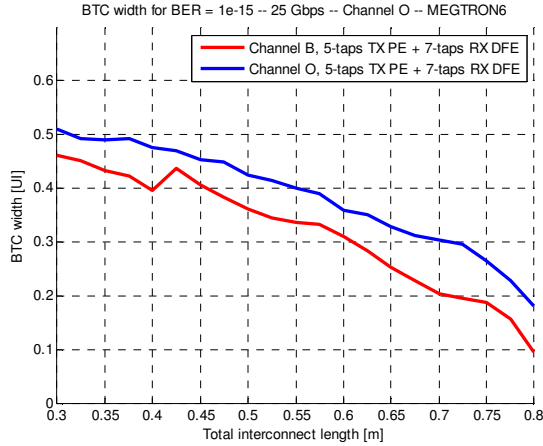


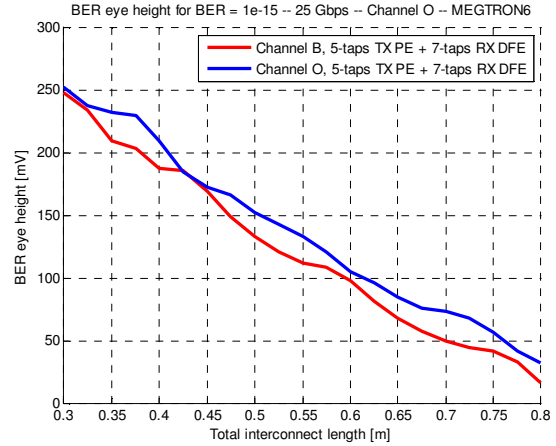
Figure 25: PSXT of a connector with 2 footprints for different TX/RX configurations: channel B (a) and channel O (b).

8.4 Link simulation results

Figures 26 and 27 show the results of link simulations for different board materials. The TX/RX configuration from figure 20(d) is used in all simulations. The BTC width and BER eye height corresponding to a BER of 1e-15 are shown as a function of the total interconnection length. Simulations have been performed for channel B and channel O. In these simulations the effects of surface roughness, the multiple reflections caused by the transceiver parasitics and the driver jitter are included. Signal conditioning is required to open the eye at 25 Gb/s. The transmitter uses a 5-taps PE filter with 2 pre-cursor taps and 2 post-cursor taps. The receiver uses a DFE with 7 baud-spaced taps.

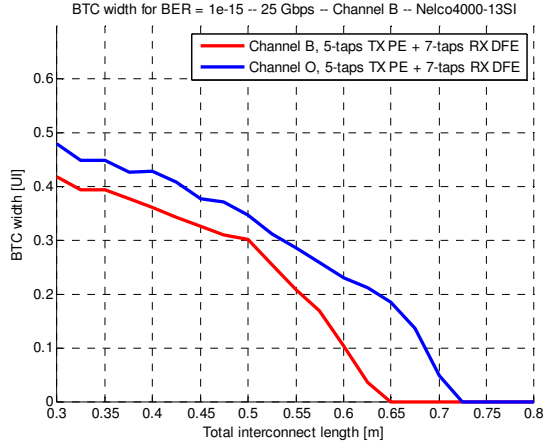


(a)

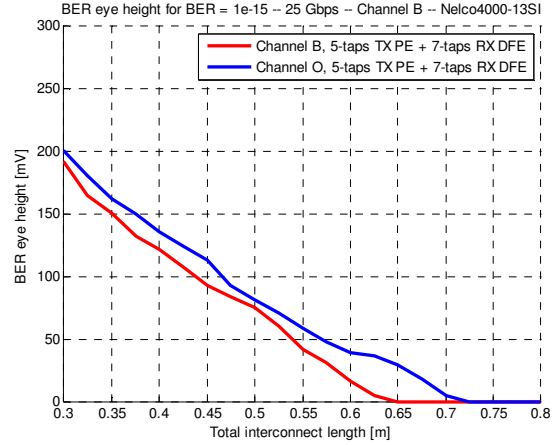


(b)

Figure 26: BER simulations, $Df = 0.003$, TX/RX configuration (d).



(a)



(b)

Figure 27: BER simulations, $Df = 0.008$, TX/RX configuration (d).

For a targeted total interconnection length of 70 cm, a low loss material (loss tangent of 0.003, like e.g. Megtron6) is required to provide sufficient eye opening. A higher loss material (loss tangent of 0.008, like e.g. Nelco4000-13SI) does not provide sufficient eye opening for a total interconnection length of 70 cm. However, these materials may be used in applications with shorter maximum lengths.

Table 4 summarizes the BER eye opening for a 5-taps TX PE filter in combination with a 7-taps DFE. The maximum total interconnection length that can be achieved is determined by the required receiver sensitivity. Assuming a eye opening metric requiring a 0.3 UI horizontal eye opening (12 ps) a total interconnection length of 50 cm can be achieved for both channel B and channel O, and using Nelco4000-13SI material. For a total interconnection length of 70 cm Megtron6 material is required.

	Total length = 0.5 m		Total length = 0.7 m	
	Eye width [UI]	Eye height [mV]	Eye width [UI]	Eye height [mV]
Channel B, Megtron6	0.36	133	0.20	49
Channel B, Nelco4000-13SI	0.30	75	0	0
Channel O, Megtron6	0.43	152	0.30	73
Channel O, Nelco4000-13SI	0.35	81	0.05	4

Table 4: Summary of BER eye opening for a 5-taps TX PE + 7-taps RX DFE.

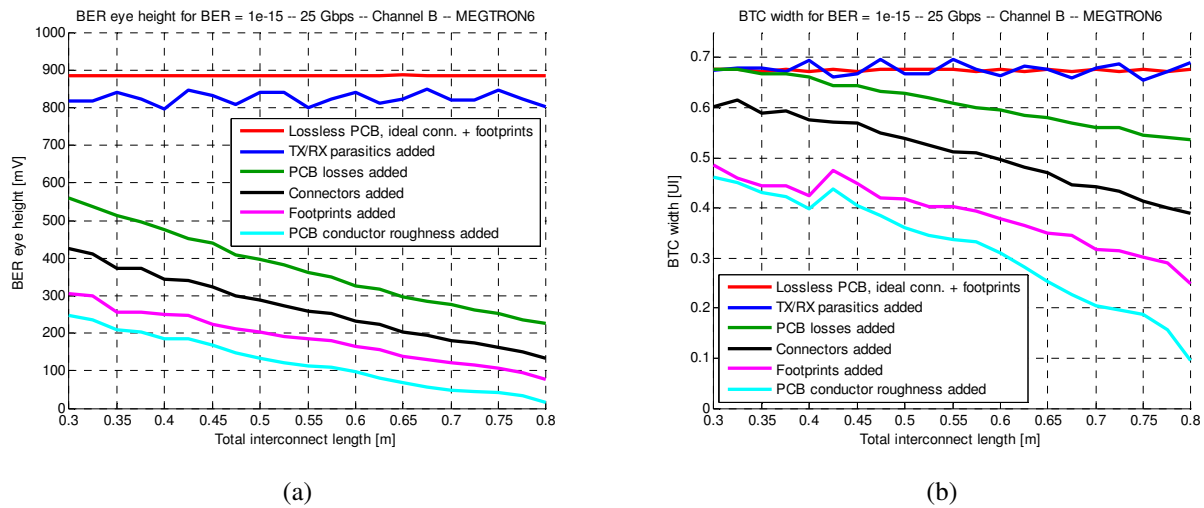


Figure 28: BER simulations, $D_f = 0.003$, TX/RX configuration (d), channel B: BER eye height (a) and BER eye width (b).

Figure 28 shows how the eye closes as different components and effects are added to the channel. The channel starts as an ideal channel with lossless PCB's, ideal connectors and ideal connector footprints and without transceiver parasitics. Step by step the transceiver parasitics, the PCB losses, the connectors, the connector footprints and the surface roughness are included in the channel model, and the eye closure caused by each component can be observed (the eye opening is shown for a 5-taps PE filter and a 7-taps DFE).

9 Conclusions

In this paper we addressed some of the SI challenges inherent to 25 Gb/s serial data transmission and we proposed requirements for different components in a 25 Gb/s copper backplane link. A top-down approach was used, starting with a perfect link consisting of two ideal connectors with ideal footprints, a backplane and two component boards, and then replacing the ideal components by real (non-perfect) components. The performance

degradation caused by each of the incrementally-introduced imperfections was studied in combination with FFE and DFE. Different 25 Gb/s specifications were compared and requirements for board material, connector and footprint impedance mismatches and crosstalk were studied. BER link simulations were performed in combination with different levels of signal conditioning.

Note that in our simulations only NRZ signaling with basic equalization schemes like baud-spaced TX pre-emphasis and RX DFE are used. For the design of 25 Gb/s SerDes more advanced signal conditioning and processing techniques are being considered, like multi-level signaling (PAM-4), duo-binary signaling, crosstalk cancellation, forward error control and others [4],[5]. These could improve the BER even further; however these were not part of the scope of this paper. Further study is required to assess the benefits of these techniques.

10 References

- [1] IEEE Std 802.3ap™-2007, Annex 69B.
- [2] CEI Implementation Agreement Draft X.X, Document OIF2008.161.07, October 14th 2009.
- [3] T. Liang, S. Hall, H. Heck and G. Brist, “A Practical Method for Modeling PCB Transmission Lines with Conductor Surface Roughness and Wideband Dielectric Properties”, Proc. of the IEEE MTT-S Symposium, San Francisco, CA, June 2006, pp. 1780-1783.
- [4] C. Zhong, C.Y. Liu, W. Jin, A. Malipatil, G. Tang and F.Y. Zhong, “Design Considerations in High Speed SerDes (25 Gbps)”, DesignCon 2008.
- [5] D.G. Kam, M.B. Ritter, T.J. Beukema, J.F. Bulzacchelli, P.K. Pepeljugoski, Y.H. Kwark, L. Shan, X. Gu, C.W. Baks, R.A. John, G. Hougham, C. Schuster, R. Rimolo-Donadio and B. Wu, “Is 25 Gb/s On-Board Signaling Viable?”, IEEE Trans. On Advanced Packaging, Vol. 32, No. 2, May 2009, pp. 328-344.