DesignCon East 2003

How to Make Optimal Use of Signal Conditioning in 40 Gb/s Copper Interconnects

Jan De Geest, PhD
Jim Nadolny
Stefaan Sercu, PhD

FCI
Communications, Data, Consumer Division
FCI's-Hertogenbosch
Helftheuvelweg 11,
5222 AV 's-Hertogenbosch,
The Netherlands
tel: +31 73 6206 372
e-mail: jdegeest@fciconnect.com
sssercu@fciconnect.com

Abstract

The emergence of transceiver chips with active signal conditioning is pushing the limits for copper interconnects to higher data rates and higher interconnection lengths than anticipated before. By making use of such chips, transmitting 40 Gb/s (using 4 x 10 or 8 x 5 Gb/s/channel serial links in parallel) over cable and backplane links becomes feasible, not only over short distances but also over longer lengths. Signal conditioning techniques typically compensate for skin and/or dielectric losses in a system and as such they make the system budget requirements less severe. Some of these chips are adaptive, which means that the signal conditioning is automatically adapted to the channel and transmission becomes nearly channel independent. However, signal conditioning techniques do not solve all interconnection problems. They do not always compensate for e.g. resonances caused by via hole stubs, multiple reflections caused by impedance mismatches or crosstalk (in particular near end crosstalk). This paper focuses on the link design requirements that are necessary to make optimum use of available signal conditioning techniques. Different types of signal conditioning such as pre- and de-emphasis at transmitter side and passive and active equalization at receiver side are investigated. Interconnection length improvements that can be realized through the use of these devices are calculated. The system requirements with respect to relevant link parameters like crosstalk and via holes are studied. As a result of the study the interconnection link specifications needed for taking full benefit of signal conditioning are defined. Simulation and measurement results illustrate the conclusions of the paper.

Author Biographies

Jan De Geest 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 in 's-Hertogenbosch, The Netherlands. His work focuses on the modeling and simulation of high-speed interconnection links.

Mr. Nadolny has 19 years experience as a SI/EMI engineer with more than 20 publications. He has a NARTE certification and has worked for GE, Honeywell, AMP and FCI. Mr. Nadolny has a BSEE from the University of Connecticut and an MSEE from the University of New Mexico. Recent work has focused on simulation of high-speed backplane connectors and interconnects for the telecommunications industry. Mr. Nadolny has given numerous presentations related to high-speed data transmission (1-10 Gbps), shielded connector design and testing issues at DesignCon, IEEE EMC Symposia and at industry standard meetings.

Stefaan Sercu received the degree in electrical engineering from the University of Gent in 1992. From 1992 to 1998, he worked as research assistant at Department of Information Technology (INTEC) of the university of Gent. His research focused on the characterization and modeling of high-speed connectors and interconnections. In 1998 he joined FCI where he is responsible for high-speed signal integrity simulations and measurements. In 2001 he received the PhD degree in electrical engineering from the University of Gent

1. Introduction

The data rates and interconnection lengths that can be achieved by actual copper interconnection links are limited by the increasing losses in copper and dielectrics with increasing speed and length. High performance materials or more complex processing techniques can provide some benefit, however at a significant increase of cost. Optical fibers are also costly and cost-ineffective for short distances. Signal conditioning provides a more cost effective solution and is the enabling technology to achieve higher speed performance in copper links. Various signal conditioning techniques are now emerging, providing different levels of performance. Each signal conditioning technique may impose some interconnection link specification requirements to guarantee its benefit. If these requirements are not met the signal conditioning can become ineffective. Figure 1 shows differential eye pattern measurements on a 0.5 meter backpanel link at 10 Gb/s. Figure 1a shows the eye pattern when no signal conditioning is applied. Figure 1b and figure 1c show eye patterns when signal conditioning (in this case active receive equalization) is applied. In 1b the signal traces are routed on the bottom layer in the backplane. In 1c the signal traces are routed on the top layer, which results in a long via stub that completely eliminates the benefit of the signal conditioning. This paper will discuss the performance merits of the various signal conditioning solutions and elaborate on the design rules that need to be fulfilled to optimize the benefit from utilizing signal conditioning in differential interconnection links.

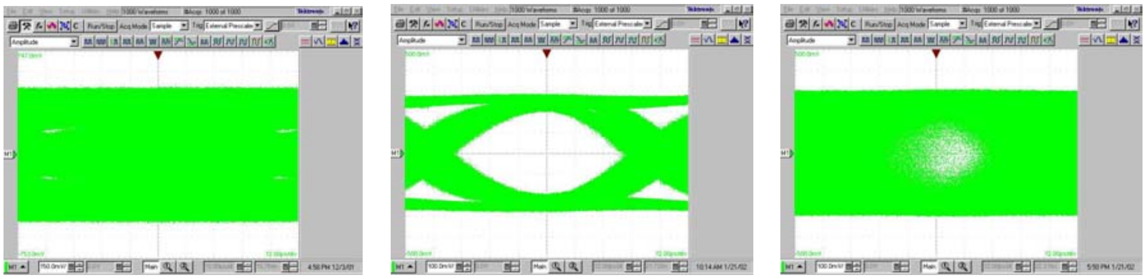
(a) (b) (c)

Figure 1: Eye pattern measurements of a 0.5 m backpanel link at 10 Gb/s: (a) without signal conditioning (b) with signal conditioning and traces routed on the bottom, (c) with signal conditioning and traces routed on the top.

2. Signal conditioning

There exist a number of different signal conditioning techniques: driver pre-emphasis, driver de-emphasis, passive receive equalization, active receive equalization or a combination of these (e.g. pre-emphasis + passive or active equalization). These techniques differ in power consumption, performance, thermal requirements and cost. The basic principle however is similar for each of these techniques. This section explains how the different signal conditioning techniques work and some advantages and disadvantages of each technique are highlighted.

2.1 How does signal conditioning work ?

As an example consider an interconnection link consisting of a driver, a 5 meter AWG 26 cable and a receiver (figure 2).

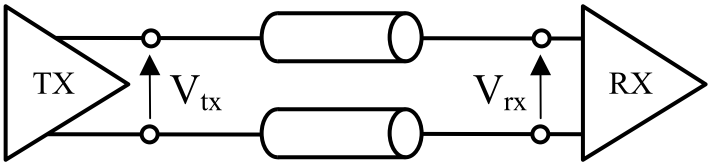

Figure 2: Schematic representation of a cable link.

Suppose the driver output impedance and the receiver input impedance are perfectly matched to the differential characteristic impedance of the cable. In that case the signal at the receiver (V_{rx}) is determined by the signal transmitted by the driver (V_{tx}) and by the differential attenuation of the cable (S_{12}): $V_{rx} = S_{12}.V_{tx}$ (in the frequency domain). The differential attenuation of the cable as a function of the frequency is shown in figure 3a. Figure 3b shows a short bitstream of nonreturn-to-zero (NRZ) data transmitted by the driver (at 10 Gb/s) and the corresponding received signal.

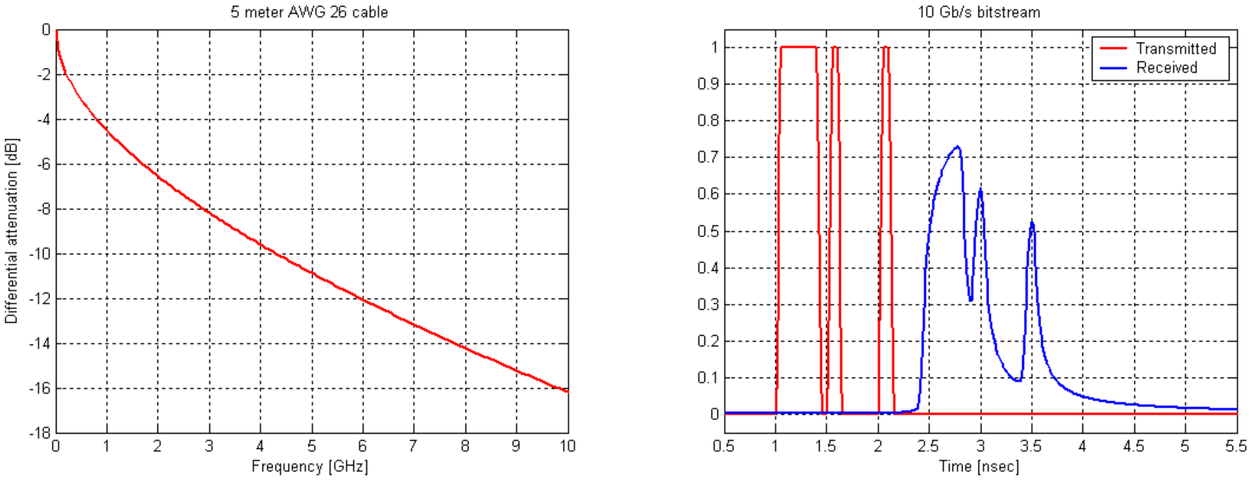

Figure 3: Attenuation of the cable as a function of the frequency (left) and transmitted and received bitstream (right).

The transmitted signal is attenuated by the cable and because of the dispersion in the attenuation (caused by the frequency dependency of the conductor and dielectric losses) the response to a single bit is several bitlengths long. This results in inter-symbolic interference (ISI), which implies that the voltage level of a bit at the receiver is not only determined by the value (0 or 1) of the transmitted bit, but also by the value of the previous bits. The level of a 0 to 1 transition following a number of consecutive 1's is different from the level of that same transition when this follows a number of consecutive 0's. In figure 4a the eye pattern of a 10 Gb/s PRBS signal is plotted. Because of the attenuation of the cable and the associated ISI the eye is nearly closed.

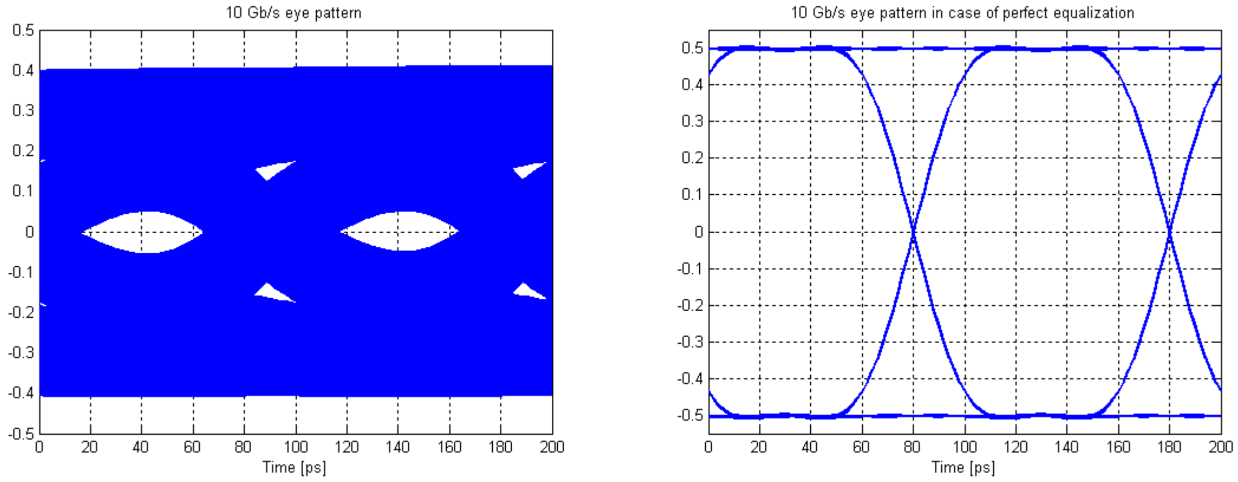

Figure 4: Eye pattern for a 10 Gb/s PRBS signal: without equalization (left) and with perfect equalization (right).

The transmission over the cable can be improved by compensating for the attenuation. This can be done by adding an equalizer to the link, as is shown in figure 5. For perfect transmission the equalizer must be designed in such a way that its frequency response H(f) is the inverse of the attenuation of the cable: $H(f) = S_{12}^{-1}$ (suppose the equalizer input and output impedances are perfectly matched to the differential characteristic impedance of the cable and to the internal impedances of the driver and the receiver).

Figure 4b shows the eye pattern of a 10 Gb/s PRBS signal for the cable plus the equalizer in case of perfect equalization. Unlike this perfect equalizer circuit a practical equalizer only has a limited bandwidth in which it can generate the necessary amplification. Figure 6a shows the frequency response of the equalizer and of the cable plus the equalizer for different bandwidths. Here the bandwidth is defined as the maximum frequency up to which the equalizer can compensate for the attenuation. Above this frequency the equalizer frequency response goes down at 40dB/decade. Figure 6b shows the received 10 Gb/s bitstream for the different bandwidths. Down to a certain bandwidth the received signal is very close to the received signal in case of perfect equalization. Below this bandwidth the received signal deteriorates.
Instead of placing the equalizer between the cable and the receiver it can also be placed between the driver and the cable. In that case the driver and the equalizer can be combined into a pre-conditioned driver (figure 7). This driver transmits a pre-conditioned signal V_{pc} that is shaped in such a way that it compensates for the attenuation in the cable. In the frequency domain V_{pc} can be found by multiplying the regular driver signal V_{tx} with the equalizer frequency response: $V_{pc} = V_{tx}.H(f)$. By applying an inverse Fourier transform the time domain signal is obtained.

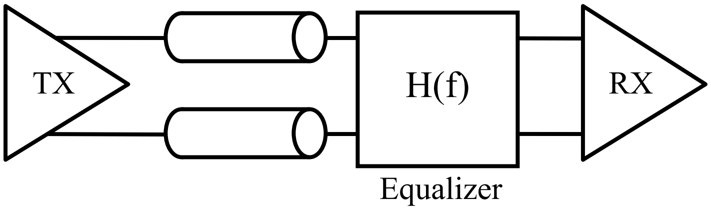

Figure 5: Equalizer added to the interconnection link.

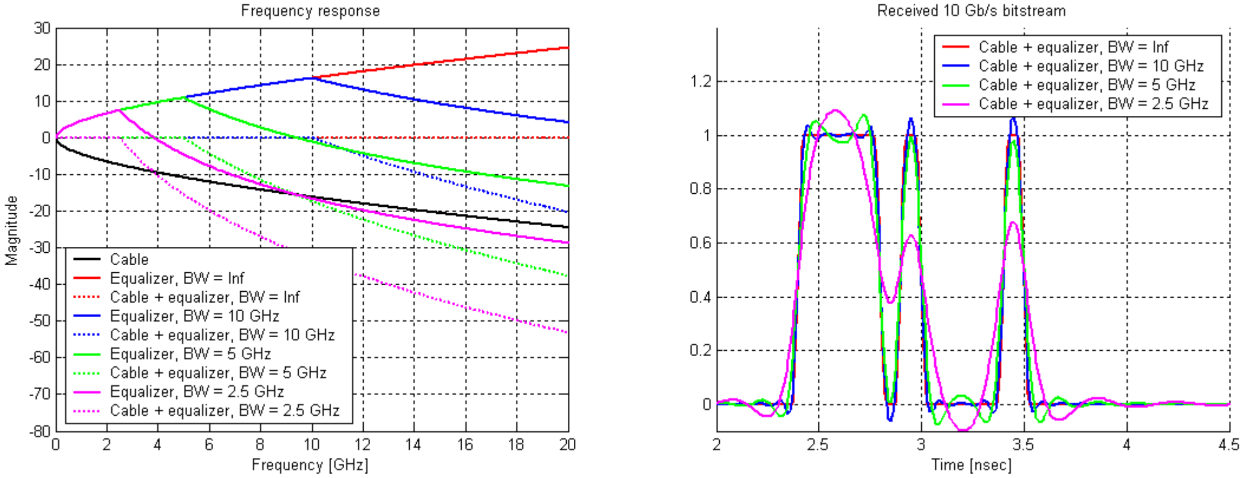

Figure 6: Frequency response (left) and received 10 Gb/s bitstream (right) for different equalizer bandwidths.

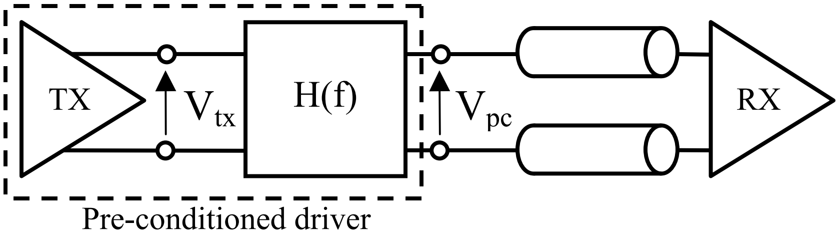

Figure 7: Driver and equalizer combined into a pre-conditioned driver.

Figure 8 shows the pre-conditioned transmitted signal in case the equalizer has a 5 GHz bandwidth. This demonstrates that receive equalization and driver pre-conditioning are two equivalent methods of applying signal conditioning to a link.

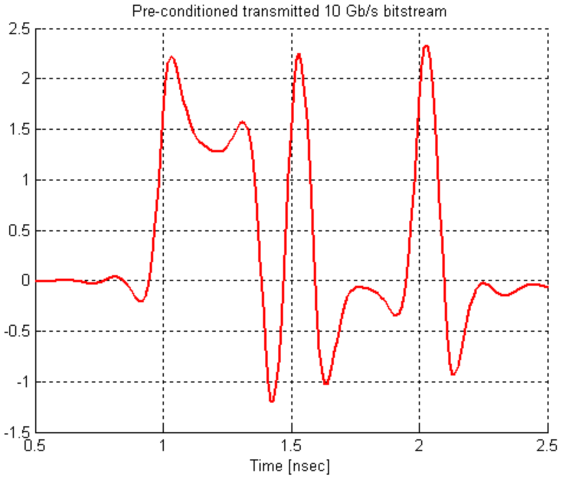

Figure 8: Pre-conditioned transmitted 10 Gb/s bitstream.

2.2 Driver pre-conditioning: pre-emphasis and de-emphasis

Figure 9a shows a short bitstream transmitted by a regular driver and by a pre-emphasis driver (at 10 Gb/s). A pre-emphasis driver boosts the level of bits following a transition. If there is no transition then the transmit level is the same for both drivers. The frequency response of an equivalent equalizer is shown in figure 9b. Boosting the level of bits following a transition results in an amplification of the power at the higher frequencies, which compensates for the higher losses in the cable at these higher frequencies. As a consequence the attenuation of the signal and the ISI are reduced, and the eye pattern improves (figure 10).

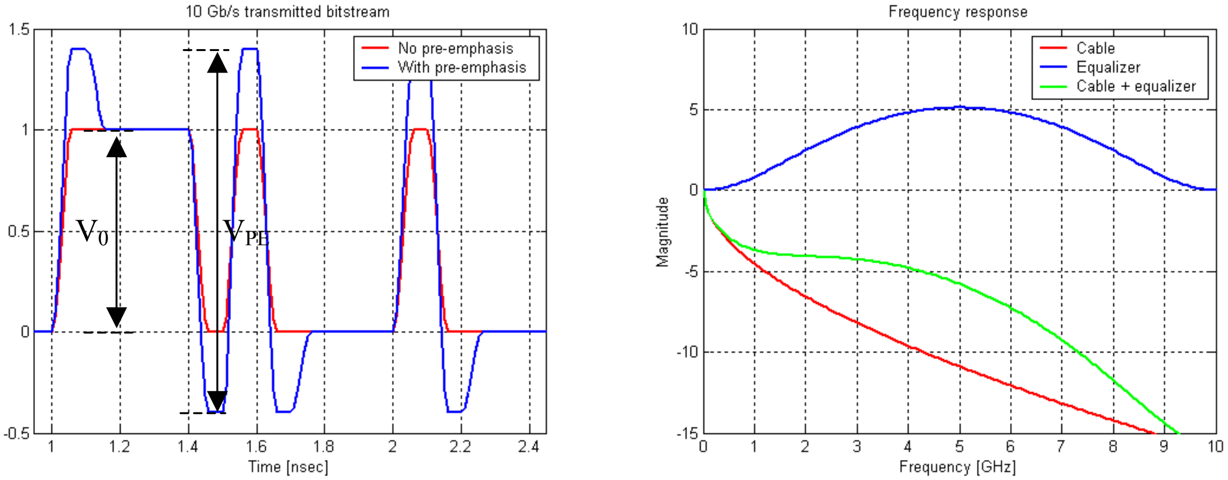

Figure 9: Short bitstream transmitted by a regular driver and by a pre-emphasis driver (left) and equivalent equalizer frequency response (right).

Evidently a pre-emphasis driver uses more power than a driver without pre-emphasis. However, since only the power at transitions (i.e. at the higher frequencies) is increased power consumption is less than what would be required for a regular driver with peak-to-peak voltage $V_{PE.}$ The power consumption and the efficiency of the equalization are a function of V_{PE}. Boosting results in large signal swings which are difficult to implement, consume power and often present EMI challenges. Therefore in practice the maximum amount of boost that can be achieved is limited [1,2].

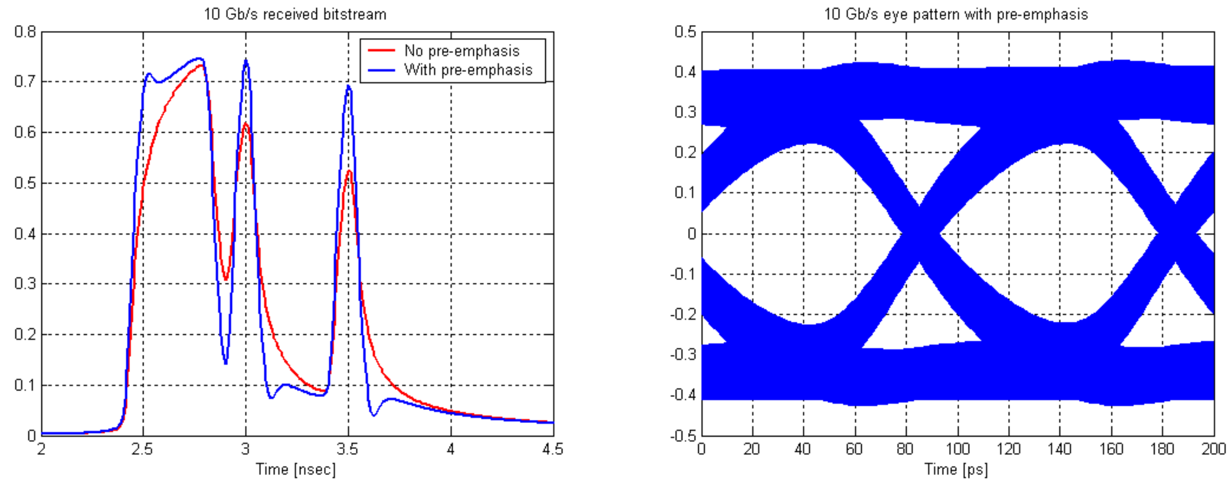

Figure 10: Received bitstream (left) and eye pattern for a 10 Gb/s PRBS signal (right) in case of pre-emphasis.

A pre-emphasis driver can be effectively implemented by a finite impulse response (FIR) equalization filter at the signal source (provided a synchronous clock is available). The FIR filter reduces ISI by inverting the channel low-pass characteristic. Figure 11 shows a two-tap digital FIR implementation (DATA is the data and CLK is the synchronous clock). The transfer function of the filter in the z-domain is $H(z) = (1+b)-bz^{-1}$. The output of the filter is $y(n) = (1+b)x(n) - bx(n-1)$. Note that in our case $b = (V_{PE}-V_0)/2$. The frequency response of the filter is $H(\omega) = (1+b)-b.\exp(-j\omega T)$, with T the pulse width of the signal (in our case T = 100 ps). The higher the boost given by the driver, the more the received eye pattern will be open. However, we can define an optimal boost as the amount of boost needed to make the tail of the pulse response as short as possible. Figure 12 shows the pulse response for different values of b. The optimal boost is achieved for b=0.4. Note that in figures 9 and 10, b = 0.4.

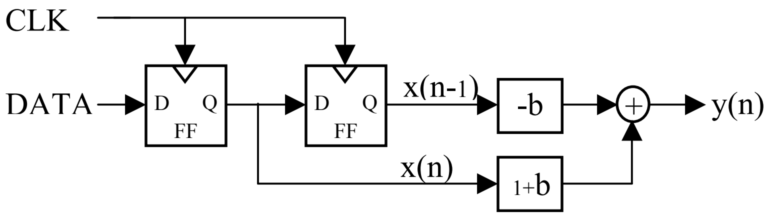

Figure 11: Two-tap digital FIR equalization filter.

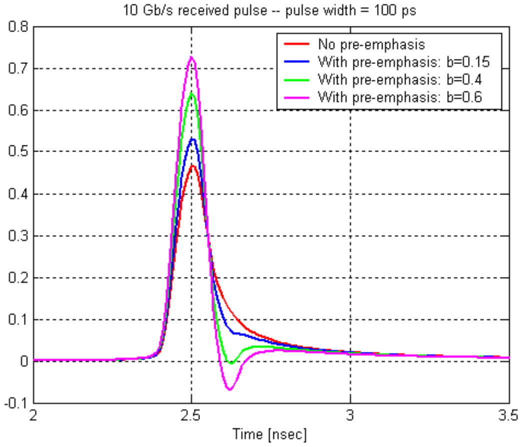

Figure 12: Pulse response for different boost levels.

To further reduce the tail of the pulse response FIR filters with more than two taps can be used to obtain more complex filter responses. Figure 13 shows the pulse response and the eye pattern for a 4-tap FIR filter implementation. The output of the filter is $y(n) = (1+a_1+a_2+a_3)x(n) - a_1x(n-1) - a_2x(n-2) - a_3x(n-3)$, with $a_1=0.4$, $a_2=0.09$ and $a_3=0.05$. There are pre-emphasis drivers on the market where the tap levels are set adaptively. This requires an additional feedback line from the receiver back to the driver or the ability to have bi-directional communication between driver and receiver. The receiver sends information back to the driver and this adapts its tap levels until the tail of the pulse response (and ISI) is minimized.

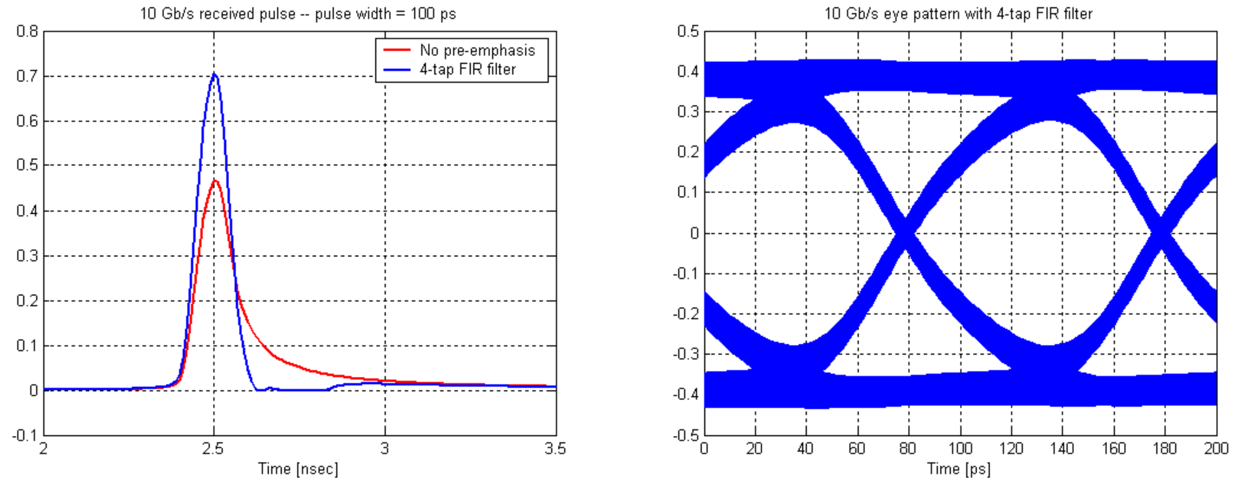

Figure 13: Pulse response (left) and the eye pattern at 10 Gb/s (right) for a pre-emphasis driver with a 4-tap FIR filter.

The problem with the higher power consumption and larger voltage swings associated with a pre-emphasis driver can be overcome by using de-emphasis. With a de-emphasis driver the peak-to-peak voltage is the same as with a regular driver, but the level of bits not following a transition is attenuated. This is equivalent to boosting the level of bits following a transition, except now the voltage swings during a transition are the same as with a regular driver and power consumption is even less than with a regular driver since the power at the lower frequencies is attenuated. Figure 14a shows a short bitstream transmitted by a regular driver and by a de-emphasis driver (at 10 Gb/s). The frequency response of an equivalent equalizer is shown in figure 14b.

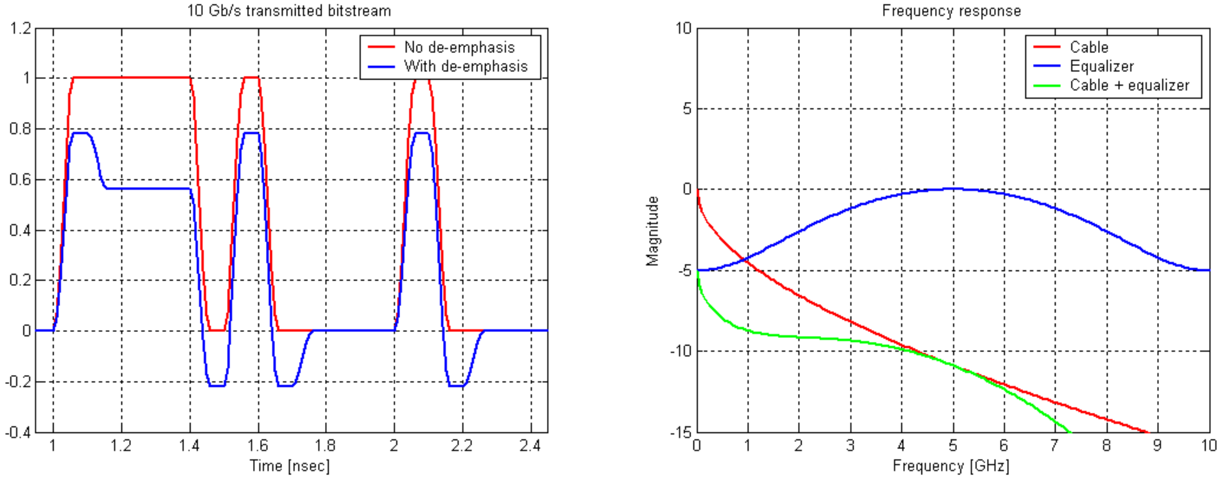

Figure 14: Short bitstream transmitted by a regular driver and by a de-emphasis driver (left) and equivalent equalizer frequency response (right).

Figure 15 shows the received bitstream and the eye pattern for a 10 Gb/s PRBS signal with de-emphasis. Although the amplitude of the received signal in case of the de-emphasis driver is lower than that in case of a regular driver, the eye pattern with the de-emphasis driver is more open than that with a regular driver. The de-emphasis driver reduces the dispersion in the attenuation of the cable by attenuating the lower frequencies. This reduces the ISI and therefore the eye pattern is improved.

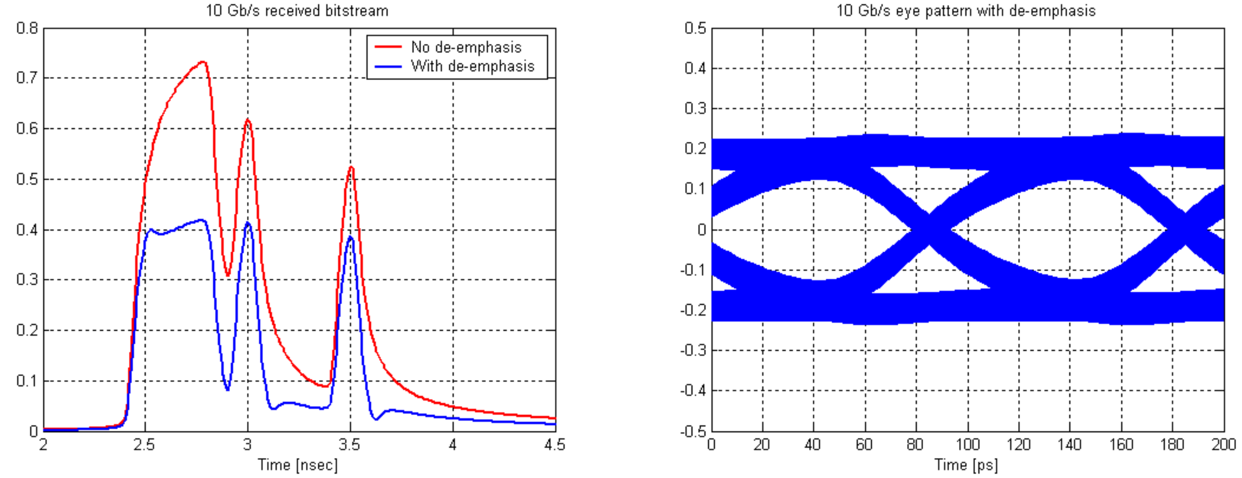

Figure 15: Received bitstream (left) and eye pattern for a 10 Gb/s PRBS signal (right) in case of de-emphasis.

A two-tap FIR filter implementation of a de-emphasis driver has a z-domain transfer function of the form $H(z) = (1-b)-bz^{-1}$. The output of the filter is $y(n) = (1-b)x(n) - bx(n-1)$. The optimal de-emphasis level can be defined as the de-emphasis level needed to make the tail of the pulse response as short as possible. This optimal pulse response is achieved for $b=0.22$. This corresponds to a power attenuation at DC of about 5.1 dB. Note that in figures 14 and 15, $b = 0.22$.

Once a receiver mask is defined the maximum interconnection length that can be achieved using pre-emphasis and de-emphasis can be calculated. Figure 16 shows the maximum interconnection length as a function of the pre-emphasis level in case of a AWG 26 cable and in case of a 100 Ω differential stripline on a FR4 backpanel ($\varepsilon_r = 4$ and $\mathrm{tg}\delta = 0.02$), for different bitrates and for two different receiver masks. The first mask (‘mask1’) is a rectangular mask with a width of 0.4 unit intervals (UI) and a height of 0.25 V. The second mask (‘mask2’) is a rectangular mask with a width of 0.25 UI and a height of 0.1 V.

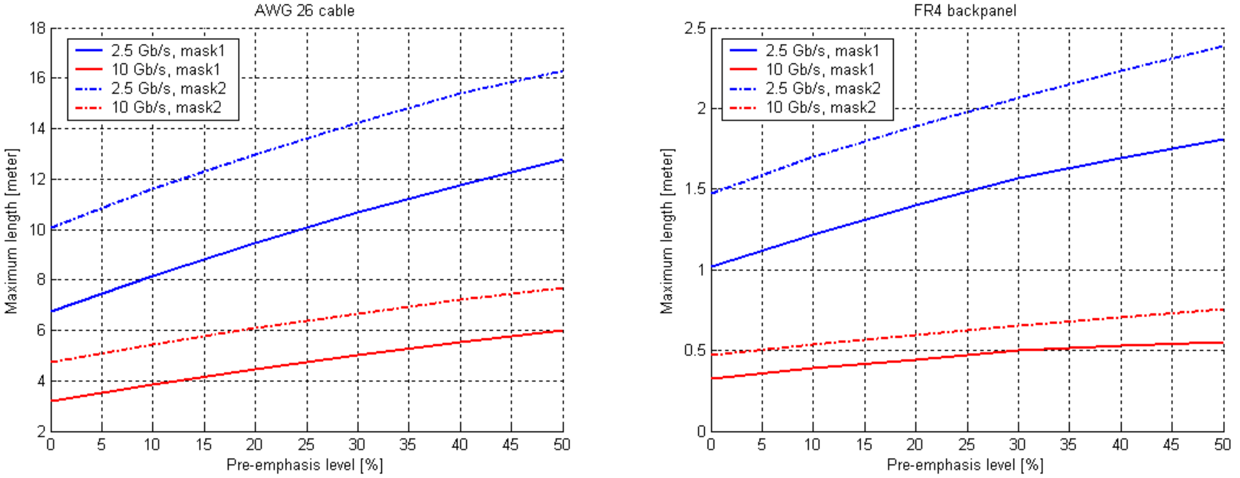

Figure 16: Maximum interconnection length as a function of the pre-emphasis level in case of a AWG 26 cable (left) and in case of a FR4 backpanel (right).

The traces on the backpanel have a width of 225 μm and a thickness of 18 μm. The distance between the traces is 220 μm. Figure 17 shows the maximum interconnection length as a function of the de-emphasis level for the same AWG 26 cable and FR4 backpanel, and for the same receiver mask definitions.

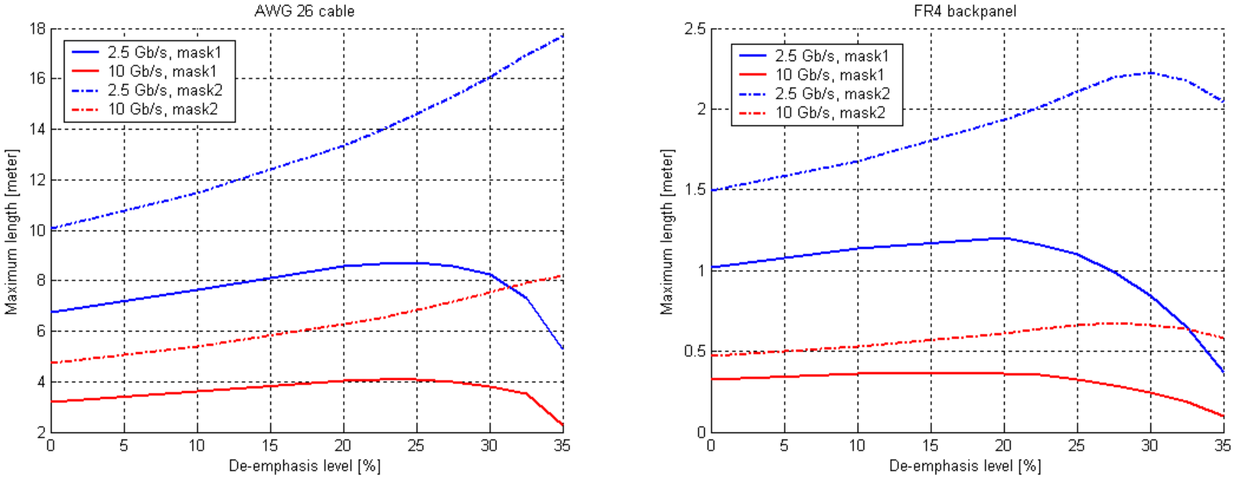

Figure 17: Maximum interconnection length as a function of the de-emphasis level in case of a AWG 26 cable (left) and in case of a FR4 backpanel (right).

Because of the additional losses that are induced de-emphasis is especially useful in applications with a large signal swing at the driver and a sensitive receiver [2]. This is also demonstrated in figure 17.

2.3 Passive and active receive equalization

Receive equalizers can be either passive or active circuits. Passive equalizers don’t suffer from additional power consumption and the associated thermal requirements. However, the frequency range over which equalization can be achieved is larger when using active equalizers. Active equalizers can amplify the power at higher frequencies, while passive circuits can only attenuate the power at the lower frequencies.

A passive equalization filter is made up of a number of lumped elements (resistors, inductors and capacitors). By using more elements more complex filter characteristics can be obtained and the bandwidth of the filter can be increased. Figure 18 shows a passive equalization filter consisting of a simple RC-filter in both signal lines and a matching impedance Z. The purpose of Z is to minimize the reflections caused by impedance mismatches between the cable and the equalizer. In the case where the equalizer is placed between the cable and the receiver this matching impedance can be omitted since reflected signals have to travel the length of the cable twice and are attenuated by the cable before they reach the equalizer. When the equalizer is placed between the driver and the cable the matching impedance is more important [5]. For the purpose of this paper we will ignore the matching impedance and assume there are no reflections at the interface between the cable and the equalizer.

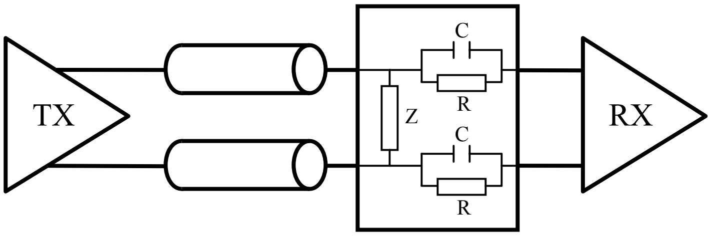

Figure 18: Passive equalization using a simple RC-filter in both signal lines.

Figure 19 shows the frequency response and the eye pattern for a 10 Gb/s PRBS signal in the case where R = 80 Ω and C = 5 pF. The dispersion in the attenuation of the cable including the passive equalization circuit is highly reduced from DC up to 1 GHz. This reduces the ISI and improves the eye pattern.

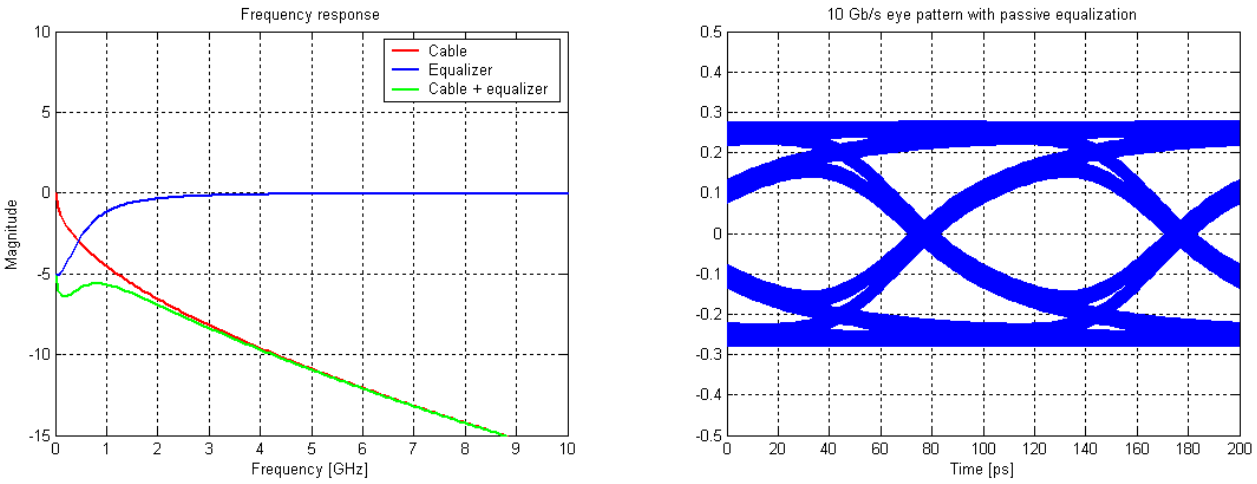

Figure 19: Frequency response (left) and eye pattern for a 10 Gb/s PRBS signal (right) in case of passive equalization.

Because it induces additional losses, like de-emphasis passive equalization is especially useful in applications with a large signal swing at the driver and a sensitive receiver. Passive equalization circuits can be built into cable connectors where the amount of equalization can be adapted to the cable assembly length. Figure 20 shows the maximum AWG 26 cable length as a function of the values of R and C for two different bitrates and for the two receiver mask specifications from the previous section.

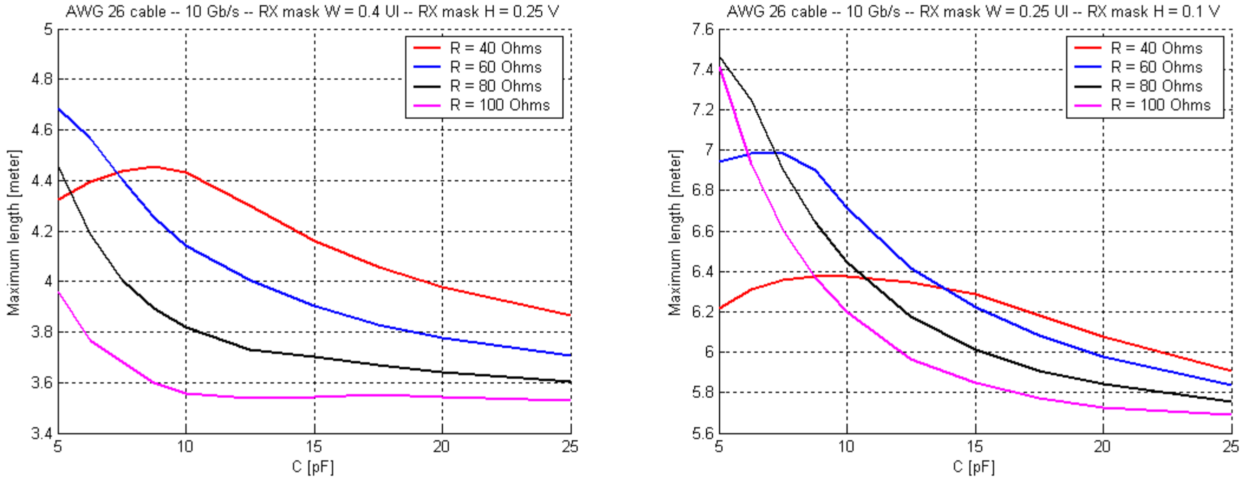

Figure 20: Maximum cable length at 10 Gb/s in case of passive equalization, as a function of R and C.

When maximum performance (i.e. the longest interconnection length) is needed, active receive equalizers must be used since these provide the greatest practical boost [2]. An adaptive active equalizer circuit improves transmission over an interconnection link by compensating for the attenuation. In case of cables and backplanes the attenuation is proportional to the length L: $S_{21} = \exp(L \cdot A(f))$. A(f) is a frequency dependent function which determines the shape of the attenuation and is defined by the geometry and the material properties of the cable or backpanel that is used. Cable losses are primarily determined by skin-effect losses in the conductors, which are proportional to the square root of the frequency. In a high-speed backplane the losses are primarily dielectric losses and these are proportional to the frequency.

The active equalizer must generate a frequency response H(f) which ideally is equal to the inverse of the link attenuation: $H(f) = S_{21}^{-1}$. This characteristic is approximated with the transfer function $H(f) = 1+\alpha G(f)$. Here G(f) is a high frequency boost stage and α is a gain between 0 and 1 which realizes the adaptive feature. Because of this adaptive feature the same equalization chip can be used for different interconnection lengths and the signal conditioning is less sensitive to construction tolerances. Cables

with different lengths can be plugged into a system and the equalizer will automatically adapt the amount of boost to keep the system working. Typically adaptive active equalizers are designed to operate over a range of interconnection lengths (and upto a certain maximum bitrate). The equalizer transfer function is optimized in case of full equalization, i.e. for $\alpha=1$ [4]. This corresponds to the maximum envisioned interconnection length. For shorter lengths α is automatically adapted through a feedback loop which samples the power spectrum of the received signal and adapts α in such a way that the typical $\sin^2(x)/x^2$ power spectrum characteristic of NRZ signals is being restored [2].
The high frequency boost is achieved by cascading a number of filter stages like the one shown in figure 21. The transfer function of such a filter stage is given by $V_{out}/V_{in} = Rg_m/(1+g_mZ/2)$ (with $g_m = qI/kT$). The impedance network Z provides the frequency-dependent characteristic. The locations of the poles and zeros of this network are optimized for best match to the inverse of the link attenuation. Instead of using analog filtering techniques, digital (DSP based) filters can be used as well.

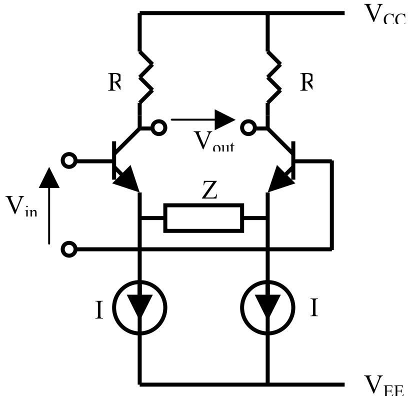

Figure 21: Active equalizer filter stage schematic.

Active receive equalizers often add additional timing jitter to the receiver output eye pattern. This is caused by improper equalization and by noise (e.g. due to ripples in the attenuation caused by multiple reflections). The amount of timing jitter that is being added can be reduced by adding a clock data recovery (CDR) circuit to the output of the equalizer. Active receive equalization is normally more difficult to implement at high speed than driver pre-conditioning. Digital receiver equalization using FIR filters requires high-resolution sampling AD convertors that run at high speed. Analog equalization requires very wide-bandwidth receiver circuits that run at the same high speed as the input data. To implement a transmit preshaping FIR filter the driver only requires adding the weighted values of the previous symbols to the present outgoing symbol value. This can be done by digital adders and multipliers to calculate the sum of the present symbol and the previous tap-weighted symbols and a high resolution DA convertor to gerenerate the analog signal. The need for complex logic can be removed by summing and modulating the signals in the analog domain [1].

3. Impact of interconnection link design parameters on signal conditioning performance

When signal conditioning is used in actual interconnection links it increases the available insertion loss budget. This implies that for a given bitrate longer interconnection lengths can be achieved or that more jitter and noise (e.g. crosstalk) can be tolerated for a given interconnection length. On the other hand the performance or efficiency of a given signal conditioning technique will depend on the design parameters of the link in which it is applied. Noise factors like crosstalk and discontinuities like via stubs and impedance mismatches will deteriorate the performance of the applied signal conditioning. Signal conditioning techniques do not always compensate for noise and discontinuities in a link, so careful link design is still necessary, even when signal conditioning is used.

To study the interaction between signal conditioning and link design parameters and to determine design rules for optimal signal conditioning performance an interconnection link modeling method is needed that takes into account the driver/receiver characteristics and the characteristics of possible signal conditioning circuits. In the past a measurement based link modeling technique using S-parameter models for each part of the interconnection link (cables, connectors, PCB's, via holes etc.) has been extensively discussed in amongst others [5-11]. These S-parameter models must be combined with driver/receiver models and models for active/passive equalizers. Available transceiver models from vendors are mostly SPICE models (usually these are encrypted to protect proprietary information) or behavioral models (like IBIS models).
In the following two sections the impact of the via stublength and of the connector crosstalk on the signal conditioning performance will be studied for the backpanel link shown in figure 22.

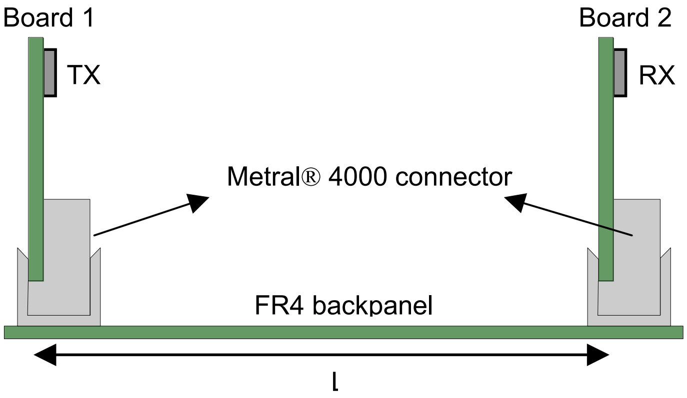

Figure 22: Metral® 4000 backpanel link.

This link consists of a 20 layer FR4 backpanel with two FR4 daughter cards connected through two FCI Metral® 4000 backplane connectors. The traces on the daughter cards and on the backplane were designed to have a 100 Ω differential impedance. The length of the traces on the daughter cards is 7.5 cm. The traces in the backpanel are routed on the top signal layer resulting in a via stublength of 5.32 mm. The backpanel length (L) can be varied. This backpanel link was part of a FCI high-speed demonstrator that was shown at DesignCon 2002. Full description details can be found in [7].

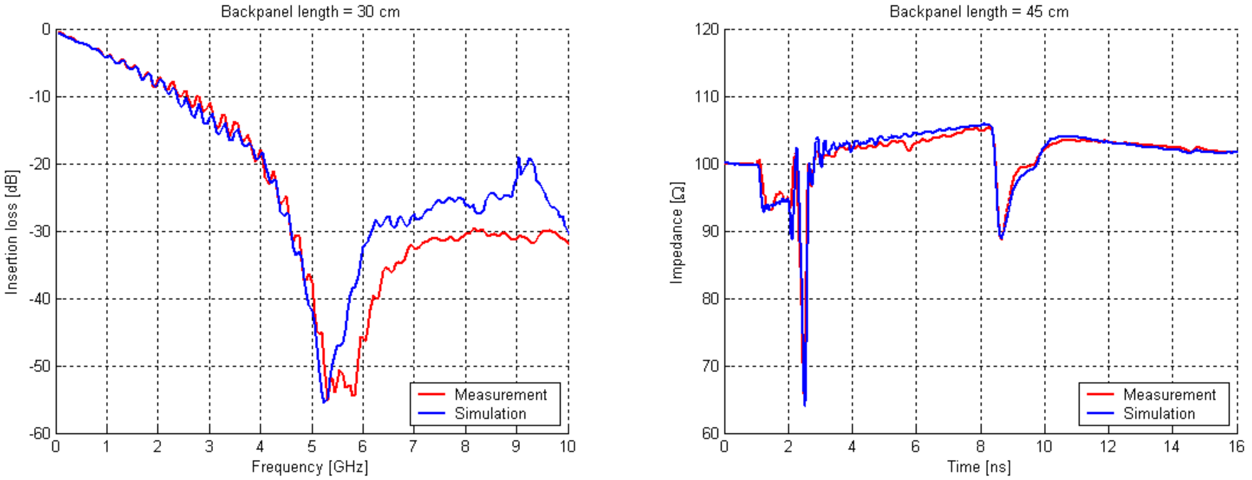

Figure 23: Measured vs. simulated performance of the Metral® 4000 backpanel: frequency domain results for a 30 cm backpanel trace (left) and time domain results for a 45 cm backpanel trace (right).

Figure 23 shows a comparison between measured and simulated data in both the time domain and the frequency domain for two different backpanel lengths. A very good correlation between measurements and simulations can be observed. This demonstrates the accuracy of the S-parameter models used to model the interconnection link.

3.1 Impact of the via stublength

Figure 24a shows the signal traces, ground planes and vias in the backpanel. On both sides of the backpanel there are two signal vias and two ground vias. The traces on the backpanel are routed on the top signal layer. This results in a via stublength of 5.32mm. This stublength (L_S) can be reduced by backdrilling of the signal vias. Figure 24b shows the differential insertion loss of the vias for different stublengths. A shorter stublength results in a lower insertion loss and a better transmission. The performance of the vias is calculated using a method described in [8-11].

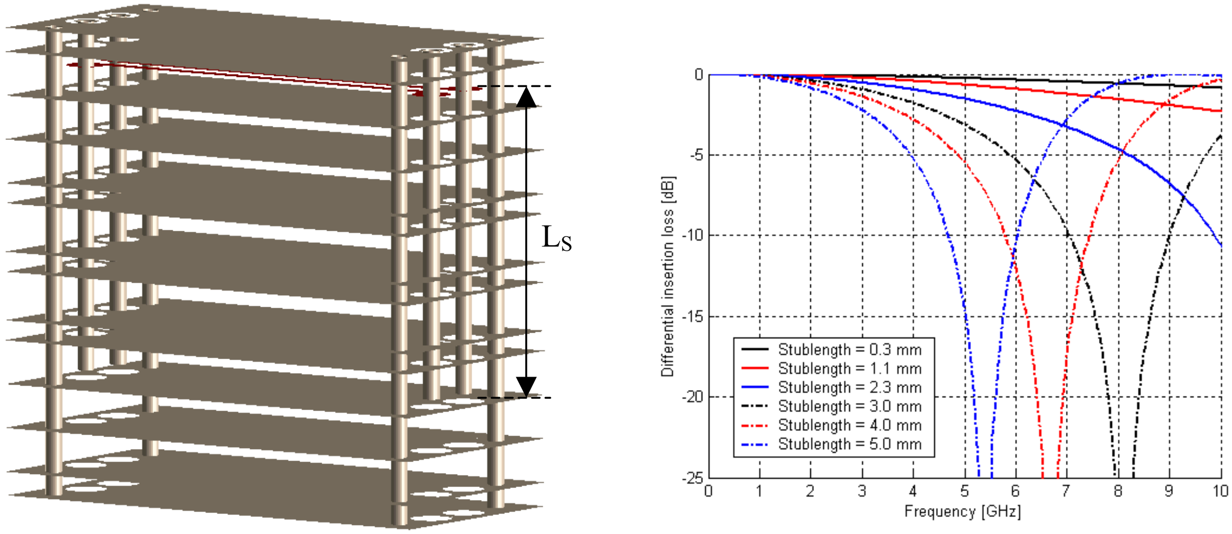

Figure 24: Signal traces, ground planes and vias in the backpanel (left) and differential insertion loss of the vias for different stublengths (right).

Figures 25 and 26 show the maximum backpanel trace length as a function of the via stublength for different pre-emphasis levels and for two different receiver mask specifications at 2.5 Gb/s and 10 Gb/s.

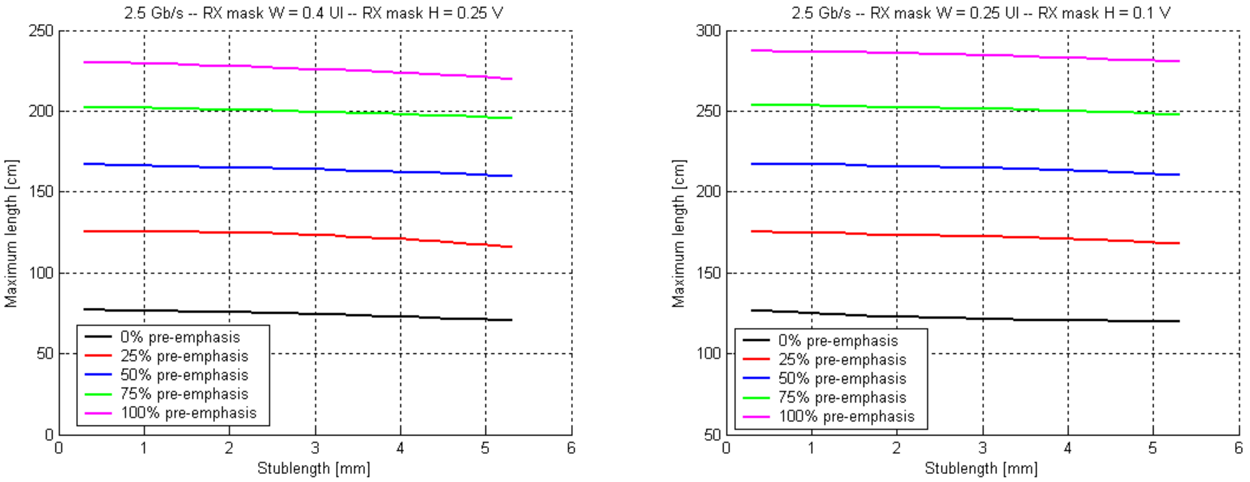

Figure 25: Maximum backpanel trace length vs. via stublength for different levels of pre-emphasis at 2.5 Gb/s.

At 2.5 Gb/s the via stublength has only limited impact on the achievable backpanel length. Decreasing the via stublength by 1 mm results in an increase of the backpanel length by 1.5 to 2 cm. By increasing the pre-emphasis level from 0% to 100% a backpanel length increase of roughly 1.5 m can be achieved, regardless of the via stublength. At 10 Gb/s the via stublength has a significant impact on the maximum backpanel length, even when pre-emphasis is used. By increasing the pre-emphasis level from 0% to

100% the maximum backpanel length increases by roughly 45 to 50 cm. Decreasing the via stublength by 1 mm results in an increase of the backpanel length by 4 to 6 cm. If the via stub is too long the maximum length becomes zero and transmission is no longer possible.

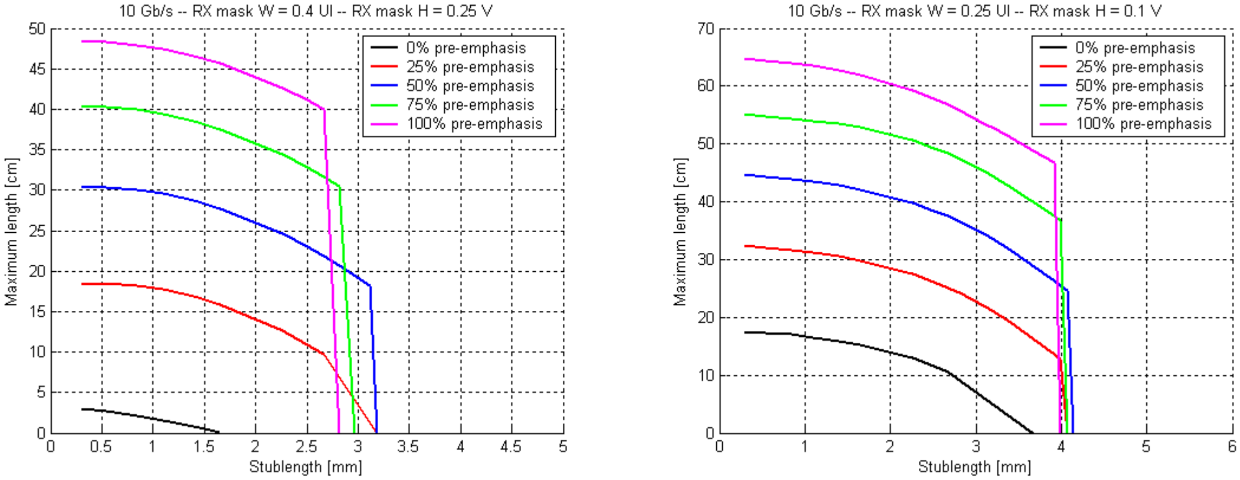

Figure 26: Maximum backpanel trace length vs. via stublength for different levels of pre-emphasis at 10 Gb/s.

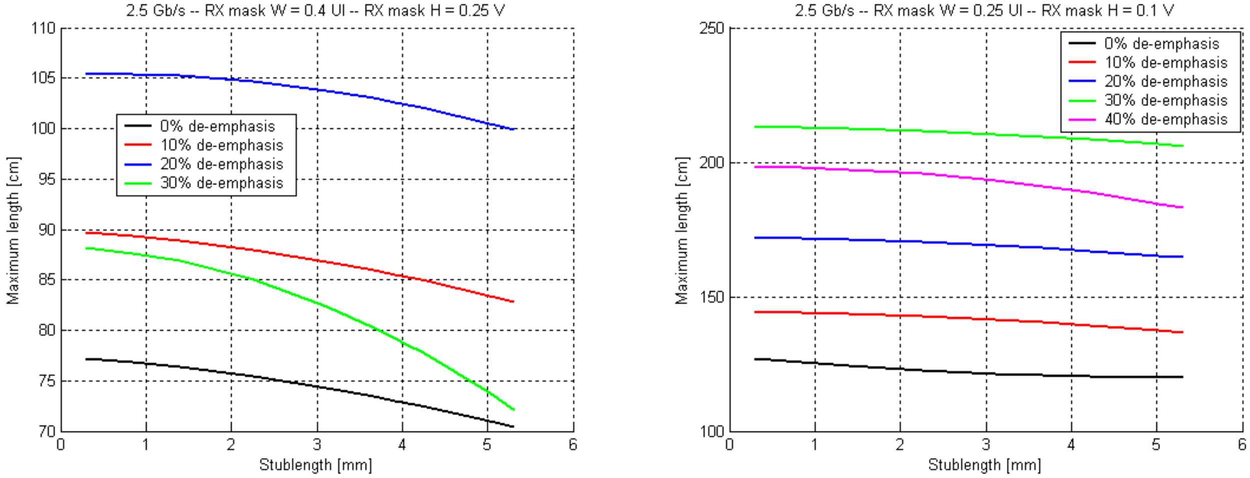

Figure 27: Maximum backpanel trace length vs. via stublength for different levels of de-emphasis at 2.5 Gb/s.

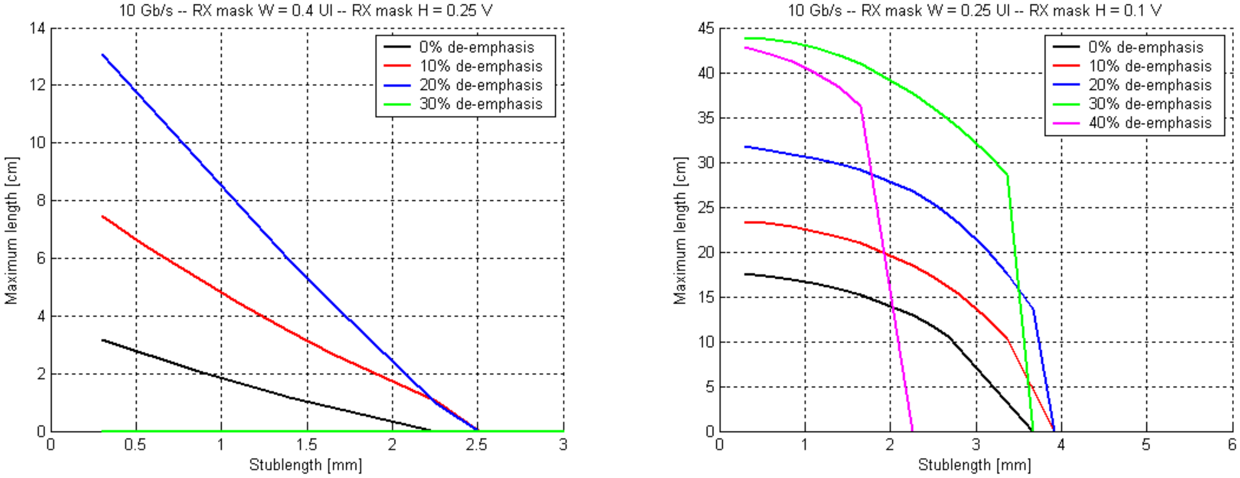

Figure 28: Maximum backpanel trace length vs. via stublength for different levels of de-emphasis at 2.5 Gb/s.

Figure 27 shows the maximum backpanel trace length as a function of the via stublength for different de-emphasis levels and for two different receiver mask specifications at 2.5 Gb/s. Figure 28 shows the maximum backpanel length as a function of the via stublength at 10 Gb/s. The maximum length increases with increasing de-emphasis level up to a certain optimum level. Beyond this optimum level the maximum length decreases again. The optimum de-emphasis level increases when a more sensitive receiver is used. The higher the de-emphasis level is, the more important the impact of the stub length on the maximum backpanel length becomes. At 10 Gb/s the maximum trace length is zero from a certain stublength on. The higher the de-emphasis level, the lower the stublength is at which transmission becomes impossible.

3.2 Impact of connector crosstalk

In this section the impact of the connector crosstalk on the maximum backpanel trace length is studied (we will assume the connector is the only component in the link that has crosstalk). This will be done for two different configurations: one where the crosstalk source is at the far end of the receiver and a second configuration where the crosstalk source is at the near end of the receiver (figure 29).

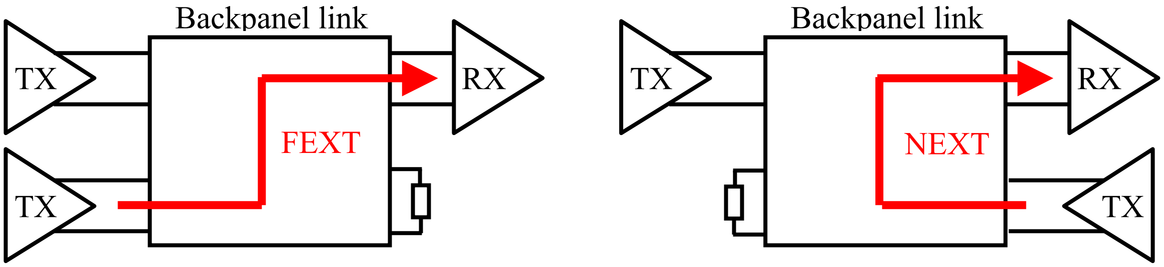

Figure 29: FEXT and NEXT configurations.

Figure 30 shows the Metral® 4000 NEXT and FEXT between columns for a step signal with a 35 ps risetime. The NEXT and FEXT peak values are 4% and 3% respectively. In this study the impact of the via hole stubs is minimized by putting the backpanel traces on the bottom signal layer (the associated via stublength is 0.5 mm).

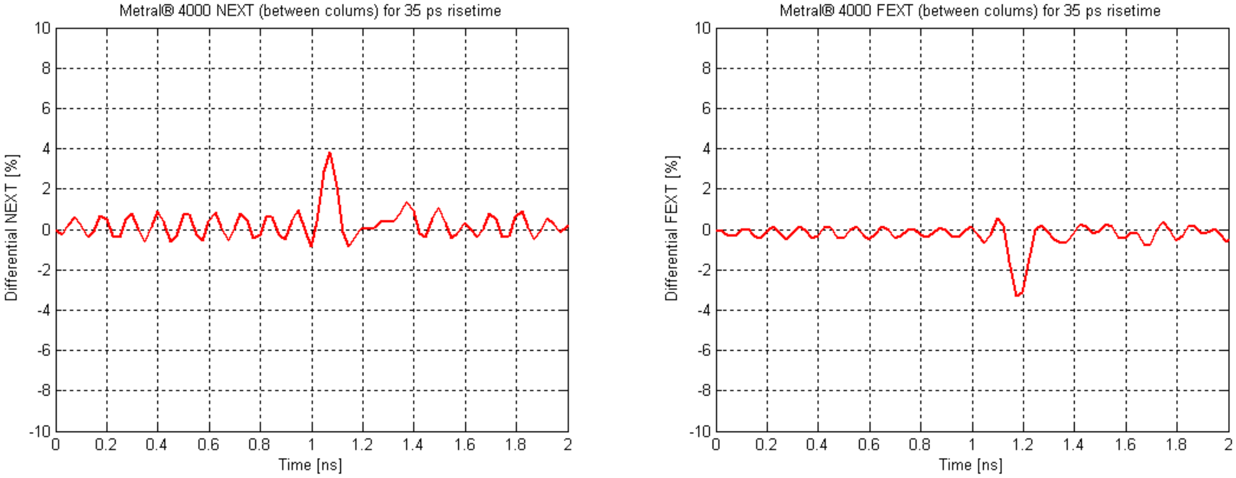

Figure 30: Metral® 4000 NEXT and FEXT between columns for a step signal with a 35 ps risetime.

Figure 31 shows the maximum backpanel trace length as a function of the connector NEXT for different pre-emphasis levels and for two different receiver mask specifications at 10 Gb/s. Figure 32 shows the maximum backpanel length as a function of the connector FEXT for different pre-emphasis levels at 10 Gb/s. Figure 33 shows the maximum backpanel trace length as a function of the connector NEXT for different de-emphasis at 10 Gb/s. Figure 34 shows the maximum backpanel length as a function of the connector FEXT for different de-emphasis levels at 10 Gb/s.

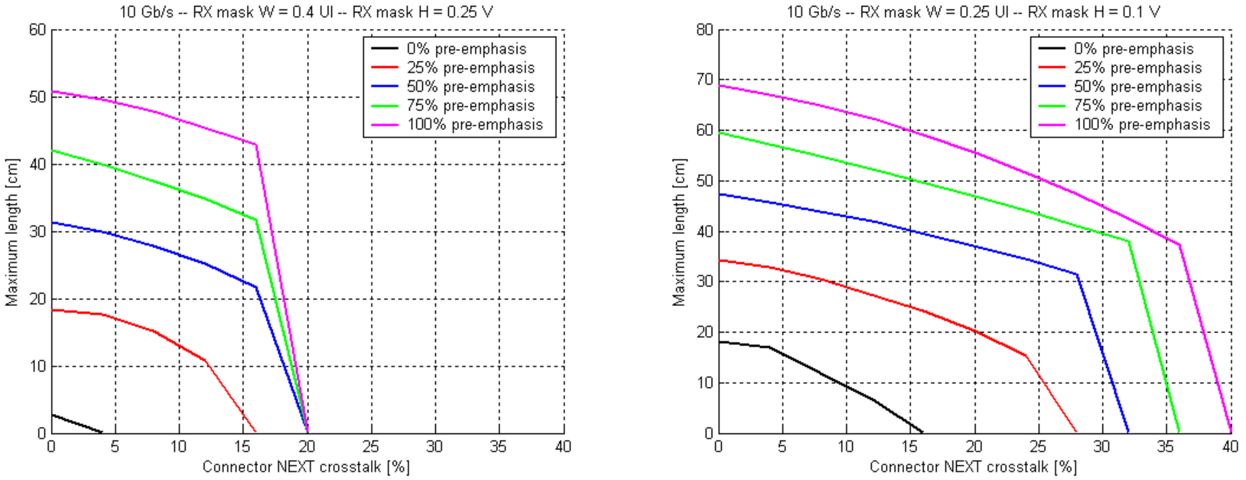

Figure 31: Maximum backpanel trace length vs. connector NEXT crosstalk for different levels of pre-emphasis at 10 Gb/s.

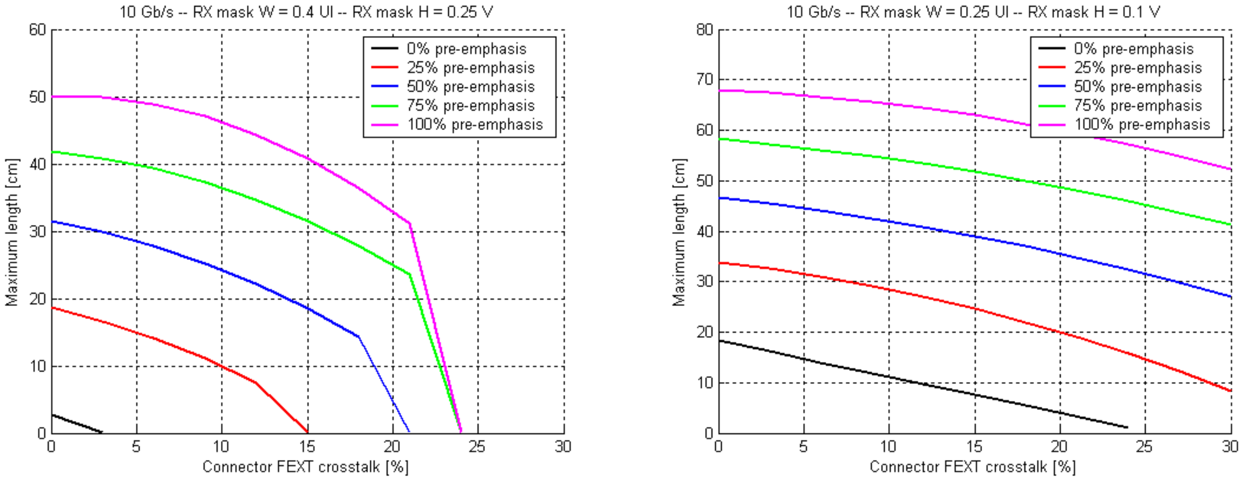

Figure 32: Maximum backpanel trace length vs. connector FEXT crosstalk for different levels of pre-emphasis at 10 Gb/s.

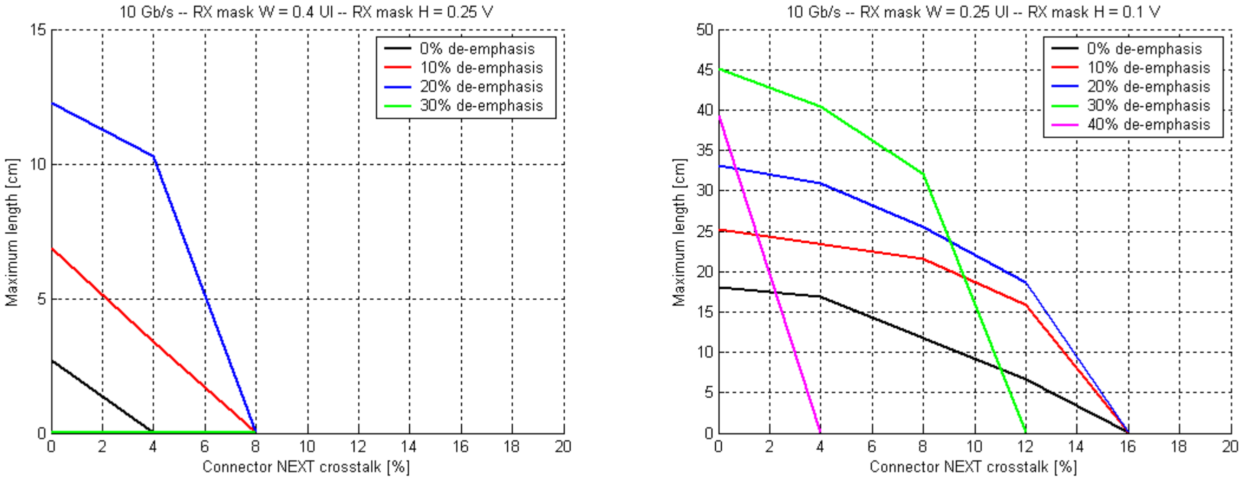

Figure 33: Maximum backpanel trace length vs. connector NEXT crosstalk for different levels of de-emphasis at 10 Gb/s.

Both the NEXT and the FEXT have a significant impact on the maximum backpanel length. If the crosstalk is too high data transmission at 10 Gb/s becomes impossible. In case of pre-emphasis the maximum backpanel length increases with increasing pre-emphasis level. In case of de-emphasis there is an optimal de-emphasis level. Beyond this optimum the maximum backpanel length decreases with increasing de-emphasis level. The higher the de-emphasis level is, the more important the impact of the crosstalk on the maximum length becomes. The NEXT has a bigger impact on the maximum length than the FEXT. This is most clear in case of de-emphasis. The reason for this is that the FEXT is attenuated by the backpanel and the daughter cards before it reaches the receiver, while the NEXT is only attenuated by the daughter cards. Therefore it is recommended not to connect transmitter and receiver signals adjacently in the connector.

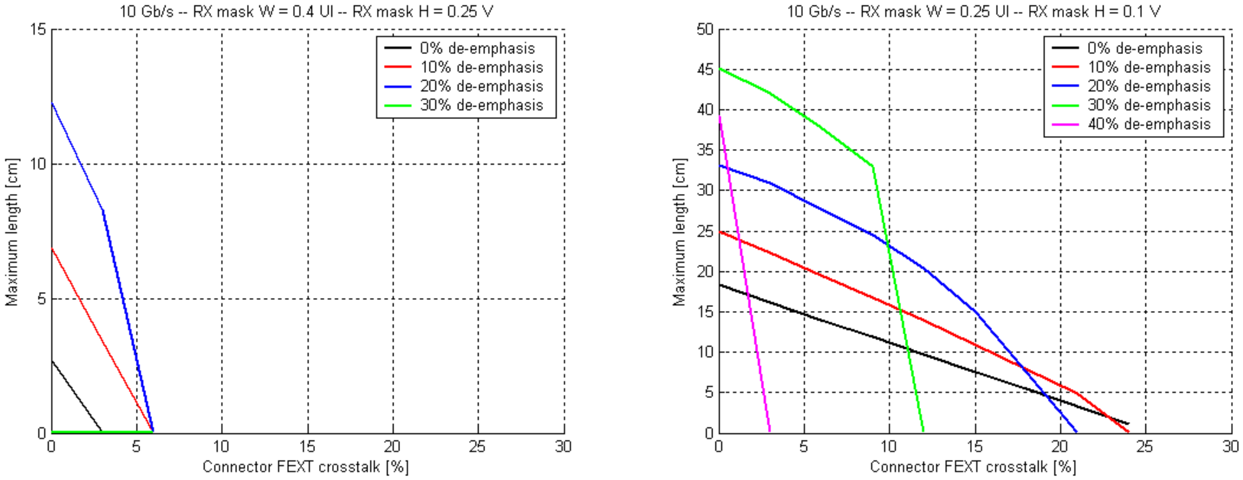

f

Figure 34: Maximum backpanel trace length vs. connector FEXT crosstalk for different levels of de-emphasis at 10 Gb/s.

3.3 Summary

Based on the results from the previous two sections a number of observations can be made:

- At 2.5 Gb/s backdrilling is not required. The via stublength only has a limited impact on the maximum backpanel length, whether driver pre-conditioning (pre-emphasis or de-emphasis) is used or not.
- At 10 Gb/s backdrilling is required. Pre-emphasis and de-emphasis cannot fully compensate for the discontinuity caused by the via stub. If the via stub is too long the benefit of using pre-emphasis or de-emphasis is lost.
- The maximum backpanel length increases with increasing pre-emphasis level. In case of pre-emphasis, the length loss caused by the via stub is comparable to the length loss when no signal conditioning is used.
- In case of de-emphasis the impact of the stublength increases with increasing de-emphasis level. For high de-emphasis levels the length loss caused by the via stub is significantly higher than the length loss when no signal conditioning is used.
- At 10 Gb/s the connector crosstalk has a significant impact on the performance improvement that can be achieved through pre-emphasis and de-emphasis. The NEXT has a bigger impact on the maximum length than the FEXT. This implies that it is best not to connect transmitter and receiver signals adjacently in the connector. This way the NEXT is minimized.
- If the crosstalk becomes too high the benefit of using pre-emphasis or de-emphasis is largely reduced.
- In case of pre-emphasis, the length loss caused by the crosstalk is comparable to the length loss when no signal conditioning is used.
- In case of de-emphasis the impact of the crosstalk increases with increasing de-emphasis level. For high de-emphasis levels the length loss caused by the crosstalk is significantly higher than the length loss when no signal conditioning is used.

4. Conclusion

Signal conditioning techniques provide a cost effective way of realizing higher speed performance in copper links. By making use of transceiver chips with active signal conditioning higher data rates and higher interconnection lengths can be achieved than anticipated before. However, signal conditioning techniques do not solve all interconnection problems. They do not always compensate for discontinuities caused by e.g. via hole stubs, multiple reflections or crosstalk.

In this paper different signal conditioning techniques (driver pre-conditioning and receive equalization) have been discussed and a number of advantages and disadvantages of these techniques have been highlighted. Interconnection length improvements that can be realized through the use of these devices have been calculated. The system requirements with respect to relevant link parameters like crosstalk and via holes have been studied. As a result of the study a number of design recommendations have been made to optimize the benefit of signal conditioning techniques.

5. References

[1] Ramin Farjad-Rad, Chih-Kong Ken Yang, Mark A. Horowitz and Thomas H. Lee, “A 0.4-um CMOS 10-Gb/s 4-PAM Pre-Emphasis Serial Link Transmitter”, IEEE Journal of Solid-State Circuits, vol. 34, no. 5, May 1999, pp. 580-585.

[2] “Equalizing Gigabit-per-Second Signals on Copper Media”, Application Note HFAN-06.0.1, MAXIM High-Frequency/Fiber Communications Group, March 2002.

[3] K. T. Oshiro, G. T. Uehara, A. K. Oki and B. Tang, “A 10-Gbps 83 mW GaAs HBT Equalizer/Detector for Coaxial Cable Channels”, Presented at the 1998 Custom Integrated Circuits Conference.

[4] M. H. Shakiba, “A 2.5 Gb/s Adaptive Cable Equalizer”, 1999 IEEE International Solid-State Circuits Conference, WP 23.3.

[5] J. Nadolny and S. Sercu, “Optimizing Copper Cable Assembly Design for Multi-Gigabit Serial Data Transmission”, DesignCon 2002, Santa Clara, California, Jan. 2002.

[6] J. De Geest, S. Sercu and J. Nadolny, "Electrical Modeling of High Speed Interconnection Links", Proc. of the 51st Electronics Components and Technologies Conference (ECTC), Orlando, Florida, May 2001.

[7] C. Clewell, J. De Geest, J. Nadolny, S. Sercu and C. Winings, "Design, Modeling and Measurement of Copper Links at Speeds to 10 Gb/s", TecForum HP-TF4, DesignCon, Santa Clara, California, Jan. 2002.

[8] E. Laermans, J. De Geest, D. De Zutter, F. Olyslager, S. Sercu and D. Morlion, "Modeling Differential Via Holes", Proc. of the 9th Topical Meeting on Electrical Performance of Electronic Packaging, Scottsdale, Arizona, Oct. 2000, pp. 127-130.

[9] E. Laermans, J. De Geest, D. De Zutter, F. Olyslager, S. Sercu and D. Morlion, "Modeling Differential Via Holes", IEEE Trans. on Advanced Packaging, Vol. 24, No 3, Aug. 2001, pp. 357-363.

[10] E. Laermans, J. De Geest, D. De Zutter, F. Olyslager, S. Sercu and D. Morlion, "Modeling Complex Via Hole Structures", Proc. of the 10th Topical Meeting on Electrical Performance of Electronic Packaging, Cambridge, Massachusetts, pp. 149-152, Oct. 2001.

[11] E. Laermans, J. De Geest, D. De Zutter, F. Olyslager, S. Sercu and D. Morlion, "Modeling Complex Via Hole Structures", Accepted for publication in IEEE Trans. on Advanced Packaging, Aug. 2002.