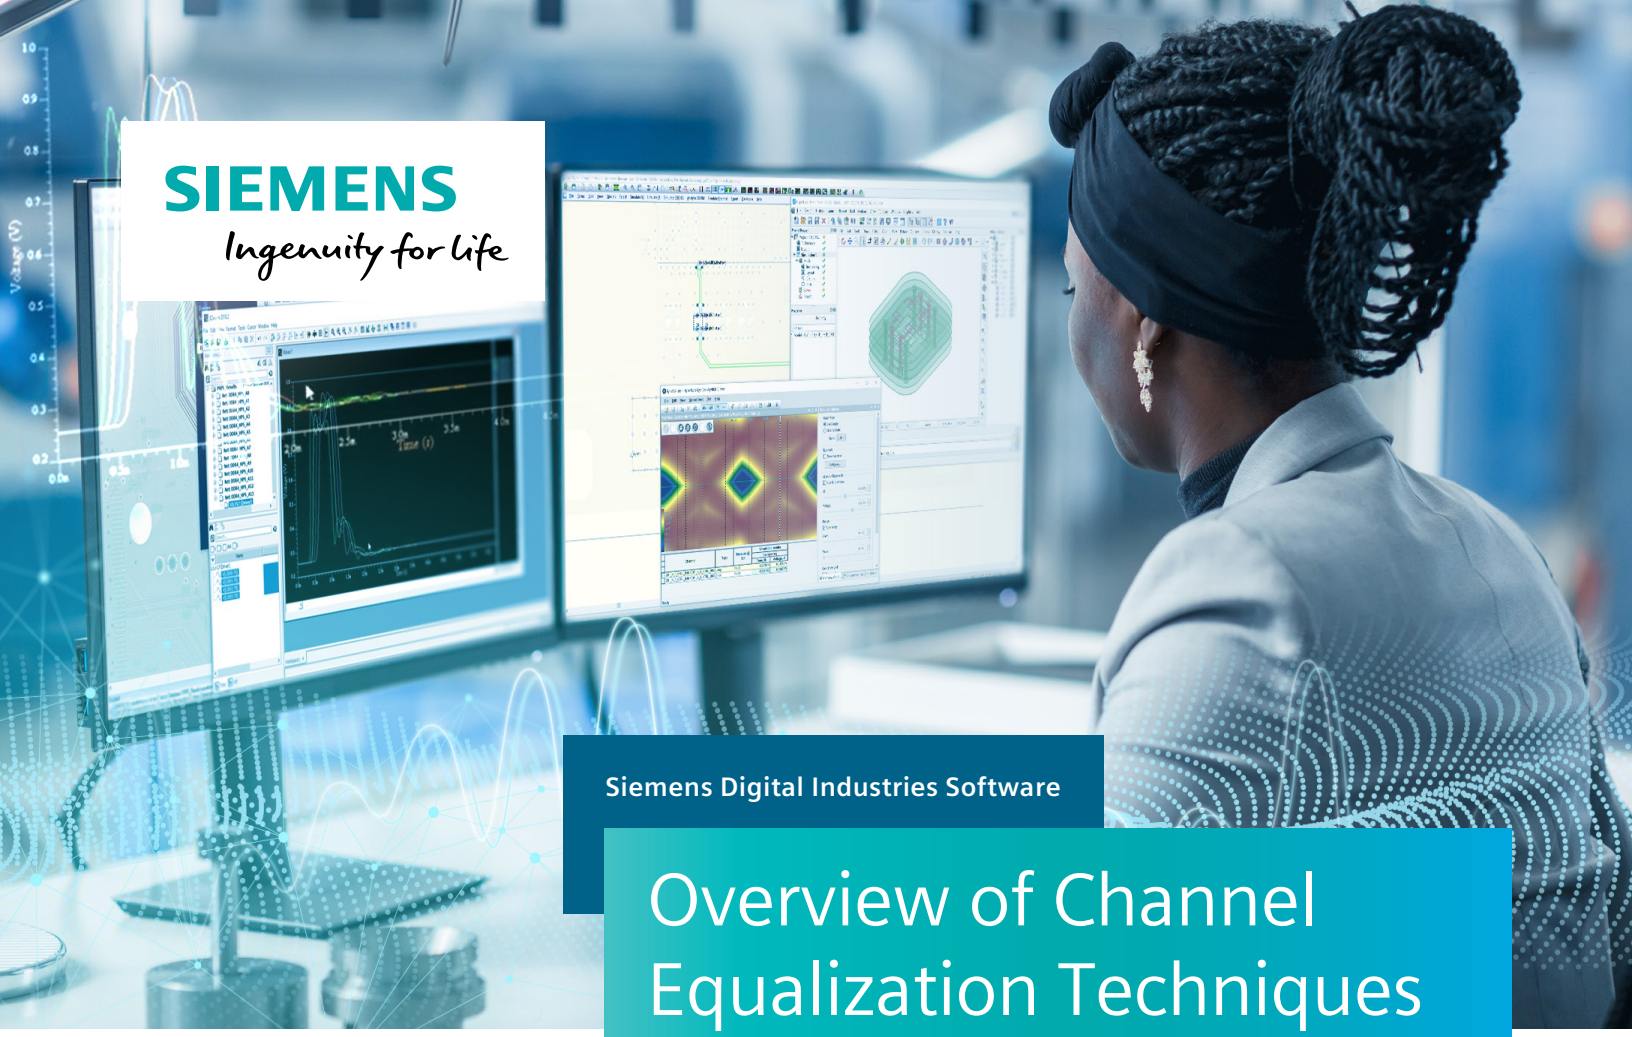

SIEMENS

Ingenuity for life

Siemens Digital Industries Software

# Overview of Channel Equalization Techniques for Serial Interfaces


## Executive summary

The newer industry standard SerDes protocols such as PCIe Gen6, USB4 and the 100G per-lane Ethernet and OIF/CEI standards offer an increasing challenge for PCB designers on multiple fronts. On the one hand, the speeds are approximately doubling for each generation. At the same time, the circuit board material used is often the same as the previous generations in order to keep costs down. In order to compensate for the increased loss at higher data rates, complex equalization techniques are employed.

In this paper, we will highlight some important aspects of the most popular interconnect specifications, with a focus on the reference equalizers.



Cristian Filip

# Contents

# Introduction

The maximum data rate at which a serial interface can operate is limited by various impairments such as reflections, losses and crosstalk. The frequency dependent attenuation of the channel has a low-pass filtering behavior causing "smearing" of a symbol into the adjacent symbols from a data stream. This phenomenon, which is the biggest source of noise for serial interfaces, degrades the performance of the link and is known as intersymbol interference (ISI). ISI is more pronounced when the signal baud rate is above the cut-off frequency of the low-pass filter.

Traditionally the performance of the channel has been evaluated using eye diagrams plotted at a specific bit error rate (BER). The effects of ISI can be identified on both the vertical eye opening, as a reduction in amplitude and on the horizontal eye opening as an ambiguity in the timing information. The ISI should be minimized or completely eliminated, in order to improve the voltage and timing margins of the SerDes links. Equalization is a signal conditioning technique that counters the effects of the channel, helps reduce ISI while increasing the signal to noise ratio and potentially increases the maximum achievable data rate.

Equalization can be applied at different locations within a link and it can be of different types, including:

- Transmitter emphasis
- Transmitter finite impulse response (FIR) equalization
- Receiver continuous time linear equalizer (CTLE)
- Decision feedback equalizer (DFE)
- Receiver Feed-forward equalizer (FFE)

Typical equalization architectures, commonly used in interfaces running at data rates up to 112Gbps are shown in Figure 1. For slower interfaces operating in the 2.5Gbps to 5Gbps range (i.e. PCIe Gen 1 and Gen2), transmitter equalization is enough to compensate for signal distortion. However, at and above 5Gbps receiver equalization is also needed, in addition to transmitter equalization. Receivers of the interfaces running at 5Gbps to 56Gbps usually contain a CTLE and/or a DFE block (i.e. PCIe Gen3, USB3.1 Gen2, 10GBASE-KR, 100GBSE-KR4, 50GBASE-KR, etc.). Attaining data rates above 56Gbps usually requires more complex receiver equalization which includes an FFE block besides CTLE and DFE. The new 112Gbps specifications such as CEI-112G-LR-PAM4 and CEI-112G-VSR-PAM4 that are in development at the time of writing this document are introducing these kind of receiver architectures.

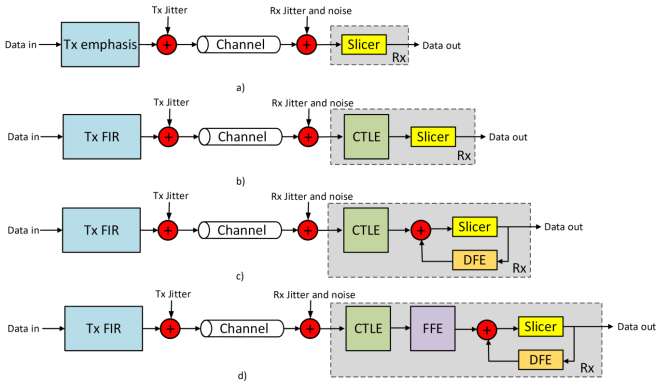


Figure 1 – Typical equalization architectures: Tx emphasis or FIR only a), Tx FIR and Rx CTLE b), Tx FIR and Rx CTLE + DFE c), and Tx FIR and Rx CTLE + FFE + DFE d)

# Transmitter emphasis

In frequency domain (FD), the effect of equalization can be explained as a technique that counters the disparity between low and high frequency losses of the low-pass filter. Ideally, the transfer function of the equalizer should be the inverse of the channel transfer function such that the combined responses result in a flat line from DC up to a frequency equal to half the maximum signaling rate, as depicted in Figure 2 [11].

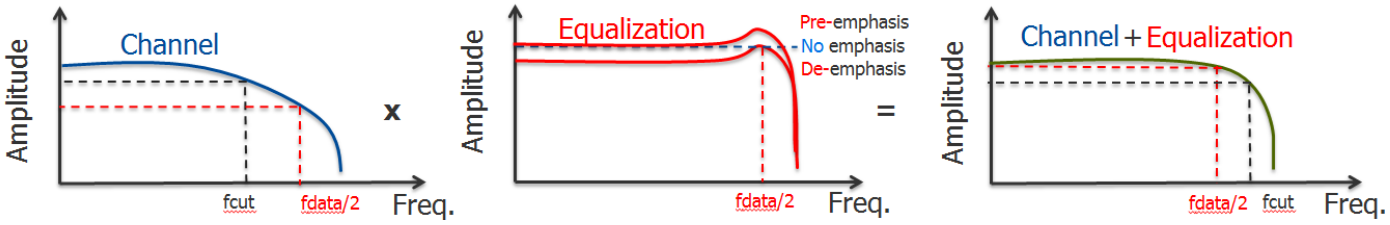


Figure 2 – The effect of equalization in frequency domain

Mathematically this can be expressed [1] as:

$$H_{eq}(f) = H_{Channel}^{-1}(f) \quad (1)$$

We use the pre-emphasis term if the channel loss is compensated by boosting the high frequency content of the signal and de-emphasis if the compensation is done by decreasing its low frequency content.

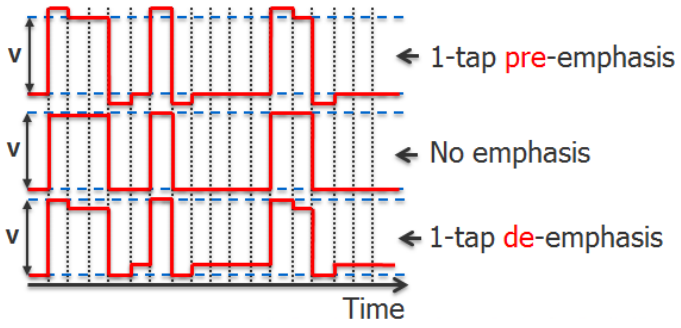


Figure 3 – Transmitter pre- and de-emphasis

In time domain, pre-emphasis or de-emphasis indicates whether the amplitude of the emphasized signal is larger or smaller than the amplitude of the non-emphasized signal. This terminology is exemplified in Figure 3 [11].

Considering that the peak-to-peak amplitude of the non-emphasized waveform is V, in the case of pre-emphasis, the peaking signal is added to the original waveform resulting in greater peak-to-peak signal amplitude. As opposed to that, in the case of de-emphasis, the peak-to-peak amplitude is maintained the same as that of the non-emphasized waveform (V) and the shoulder portion is suppressed from the original signal. Since the peaking duration is fixed and is triggered by the transitioning bits only, transmitter emphasis does not scale with the signal's data rate and it becomes less efficient when deviating from the targeted data rate [7].

# Transmitter FIR equalization

The most common transmitter equalization architectures are based on a digital FIR filter that is synchronized to the clock. A generic block diagram of a transmitter FIR equalizer is shown in Figure 4 below.

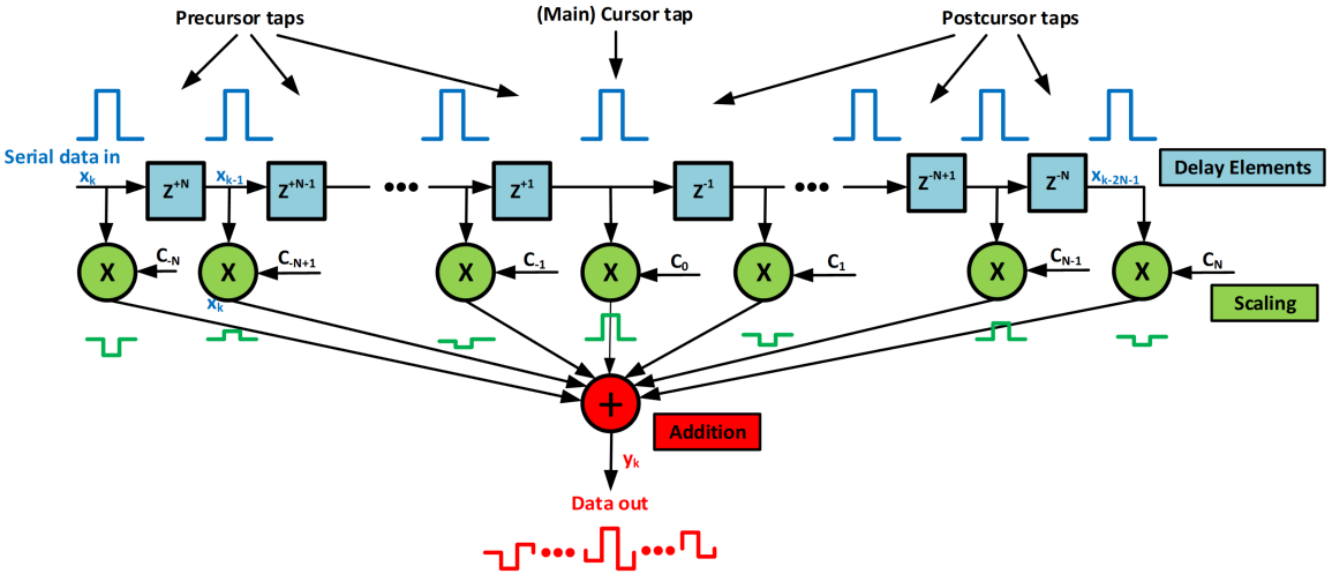


Figure 4 - Block diagram of a transmitter FIR equalizer

Samples of the incoming data stream flow through several delay elements, represented as blue squares in the block diagram. Those elements can be implemented in the hardware with flip-flops or shift registers and are usually referred to as the taps of the filter. The delay from one tap of the filter to the adjacent tap is known as tap spacing and usually is one unit interval (UI). In some cases fractional spacing of the taps is also allowed. The taps of the filter are scaled, at each stage, by a tap weight value or a filter tap coefficient, $C_i$ with "$_i$" being the tap index. The scaling is done with multipliers that are represented with green circles in the block diagram of the filter. The resulting signals are summed, in the red circle block, and sent to the data output.

Depending on the position of the taps relative to the tap $C_0$, which is called cursor or main cursor tap, the taps are identified either as precursor or postcursor taps. The precursors operate on an advanced signal and have negative indexes ($C_{-N}$ to $C_{-1}$), while the postcursors operate on a delayed signal and have positive indexes ($C_1$ to $C_N$). With this convention, the signal at the output of the filter can be computed using the following equation [5]:

$$y_{(k)} = \sum_{k=-N}^{N} x(k-n)c_n \quad (2)$$

In the above equation, $y_{(k)}$ is the output signal, $x_{(k)}$ is the input signal and k is the index of the tap coefficient into the bit stream or the sample number.

In most practical cases, the tap coefficients are constrained by the maximum voltage swing and power of the transmitter. If normalization to 1 is applied, they shall satisfy the following relation [5]:

$$\sum_i |ci| = 1 \quad (3)$$

This condition means that not all the coefficients need to be specified. Usually the precursor and post cursor taps are specified and the main cursor tap is implied.

The differential pulse response is convenient way to characterize the performance of a channel and how various impairments are degrading the signal quality. Similar to a TDR plot, a pulse response contains valuable information that can be used to debug failing channels. As the pulse propagates down the channel, it gets attenuated; consequently the received pulse will have smaller amplitude than the original pulse that was output by the transmitter. Moreover, due to dispersion (skin-effect, dielectric loss), the width of the pulse will be larger than one unit interval (1UI). The worst case noise margin degradation due to ISI can be calculated by adding the absolute value of the positive and negative terms. Finally the tail of the received pulse will contain ripples if any impedance discontinuities (layer transitions, connectors, traces over plane voids) are present in the channel.

A simulated plot of a differential pulse response observed at the receiver, after a sample channel, is shown in Figure 5. The dotted red lines represent the pre-and post-cursor ISI positive and negative terms, generated by channel loss and reflections, and the red dots are the sampling points that are uniformly spaced out and separated by the tap spacing T.

In frequency domain, the transfer function of the filter can be computed using the formula below [1]:

$$H_{(f)} = \sum_{k=-N_{pre}}^{N_{post}} c_k * e^{-j2\pi f(k-N_{pre})T} \quad (4)$$

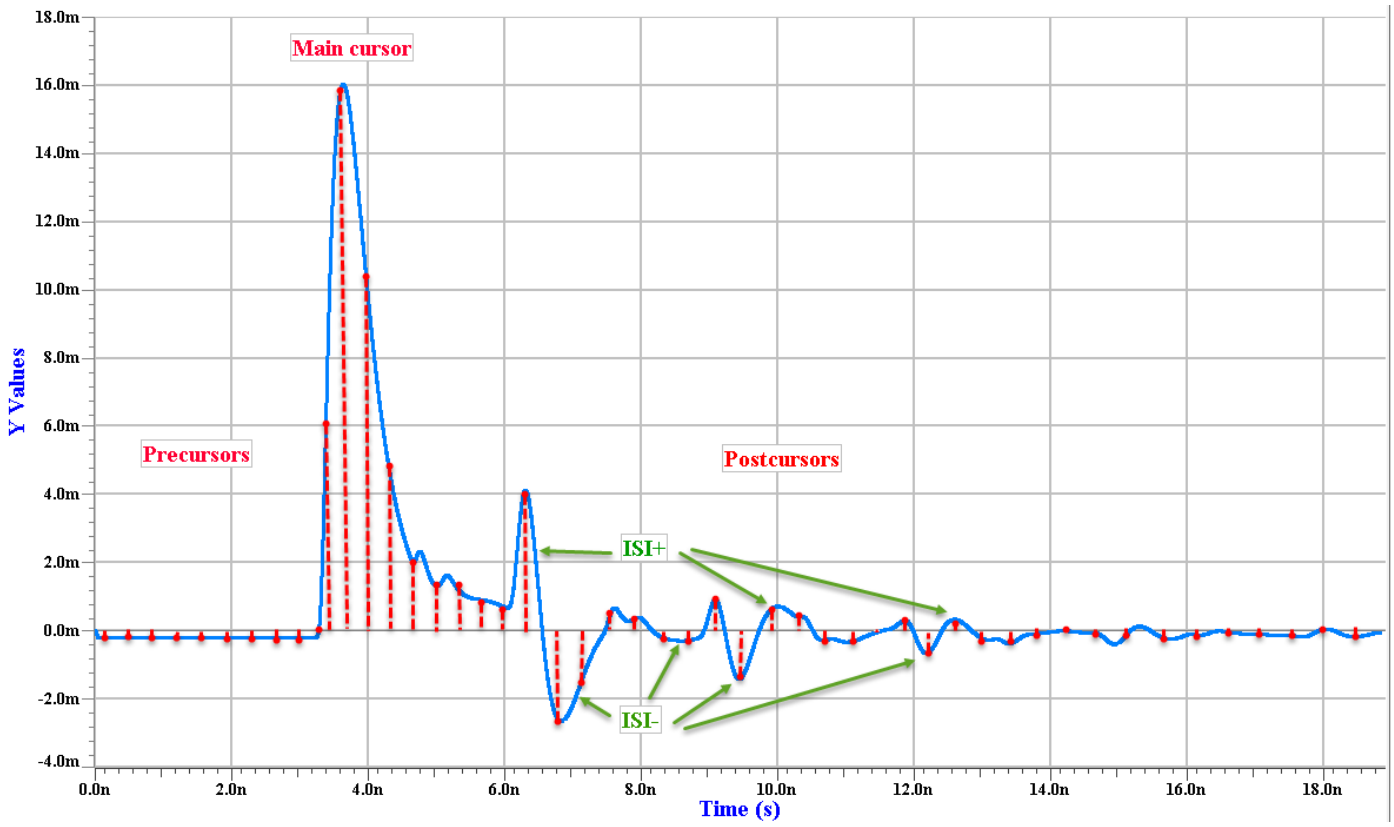


Figure 5 – Differential pulse response

A filter with a high enough number of taps can reduce or completely eliminate the system's ISI if the tap coefficients are carefully chosen. Consequently, taps optimization is one of the tasks that a simulation tool must perform during the channel analysis compliance process. We will explain how this optimization process works later in this document.

A high level block diagram of a possible hardware implementation using individual components is shown in Figure 6.

The analog signal is input to an analog-to-digital converter (ADC), where the data is sampled using the system's clock frequency. Each data sample is stored in the data memory block, while the tap coefficients are saved in the coefficient memory. The two are simultaneously fetched from the memory blocks and are loaded into the multiplier. The results of multiplications are accumulated to generate the output signal sample. Finally the data is converted back from digital to analog using a digital-to-analog converter (DAC). This FIR structure is implemented in the high-speed transceivers of the ASICs and FPGAs. It should be noted here that the resolution of the DAC is a limiting factor for the FIR tap coefficients resolution. Moreover the power consumption of ADC and DAC is an important concern.

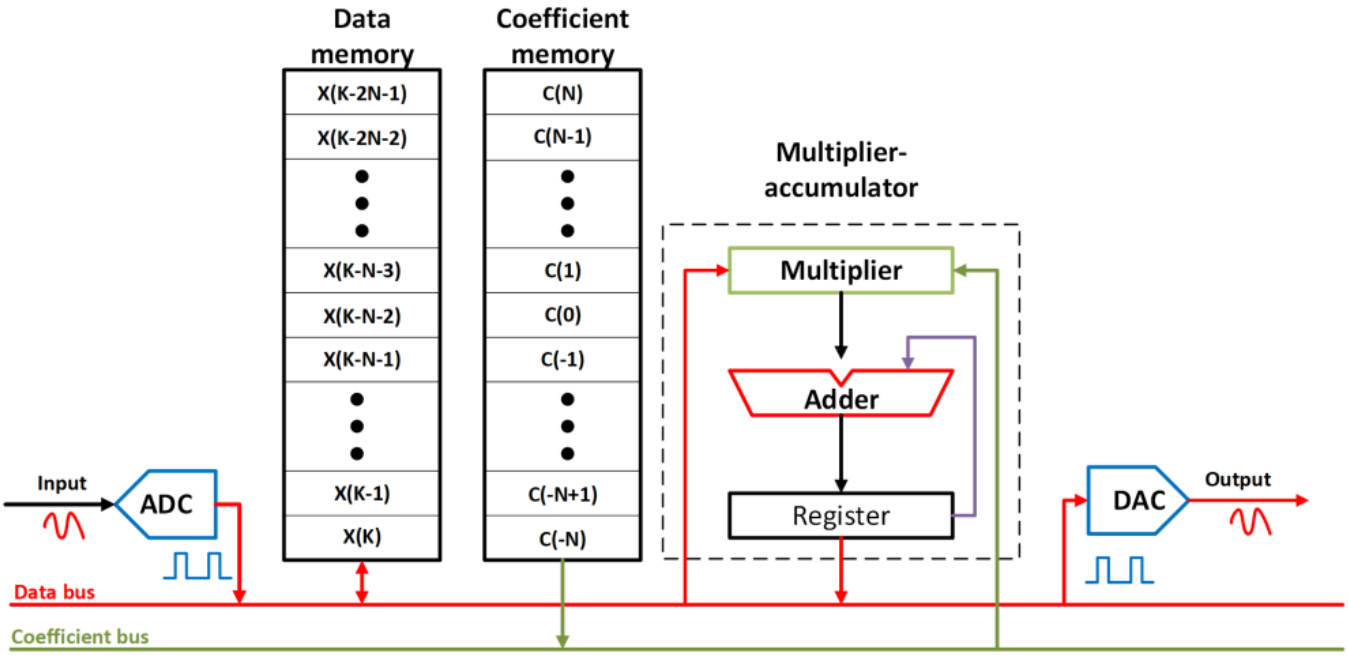


Figure 6 - Block diagram of a generic FIR digital filter implemented in hardware

# Examples of transmitter equalization architectures

Most of the SerDes protocol specifications define normative or informative reference implementations for transmitter equalization. Those reference architectures represent the minimum capability that a transmitter should have in order to be considered compliant to the specification requirements. More powerful implementations will produce better results at increased price and power consumption. From a system designer's perspective it is important to know what the reference architecture requirements are, as this is a critical aspect of the channel compliance analysis process.

The simplest transmitter equalization architecture is used in links running at relatively low speed by today's standards. For example, PCIe Gen 1 (2.5Gbps), PCIe Gen 2 (5Gbps), USB3.1 Gen 1 (5Gbps) and MIPI D-PHY 2.1 specifications define a two-tap FIR filter with fixed equalization levels. The equalizer structure and the transfer function for a reference USB3.1 Gen 1 compliant transmitter is shown in Figure 7.

The nominal de-emphasis value of this transmitter is -3.5dB with a +/-0.5dB tolerance. Figure 8 shows the overlapped unequalized and equalized pulse responses at the output of the transmitter. Those plots were generated by the HyperLynx SerDes Wizard with internally synthesized package models disabled.

The green waveform is the non-equalized pulse response and the red waveform is the de-emphasized pulse response. The amplitude of the non-equalized pulse is 1V while for the equalized waveform the measured amplitudes of the cursor taps are $C_0$ = 0.833V and $C_1$ = -166.7V respectively. Those values satisfy the constraints defined in equation (3). Note that the main cursor tap is not delayed, while the post cursor tap has 1UI delay. The amplitude of the equalized pulse response can be calculated as:

$$y_{(t)} = 0.833 * x_{(t)} - 0.1667 * x_{(t-T)}, \quad (5)$$

where $x_{(t-T)}$ is a copy of the original signal at the input of the FIR filter, but delayed one-bit time.

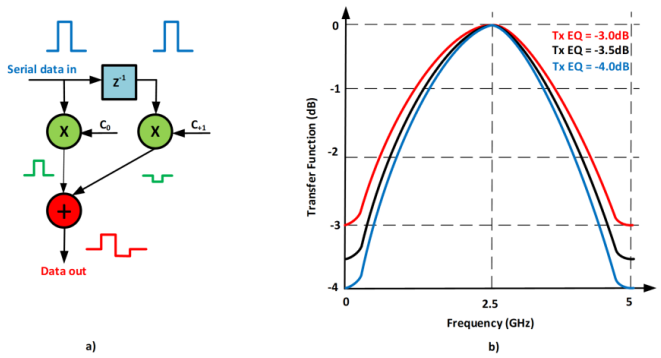


Figure 7 – USB3.1 Gen1 transmitter equalizer and transfer function: Equalizer structure (a) and Transfer function (b)

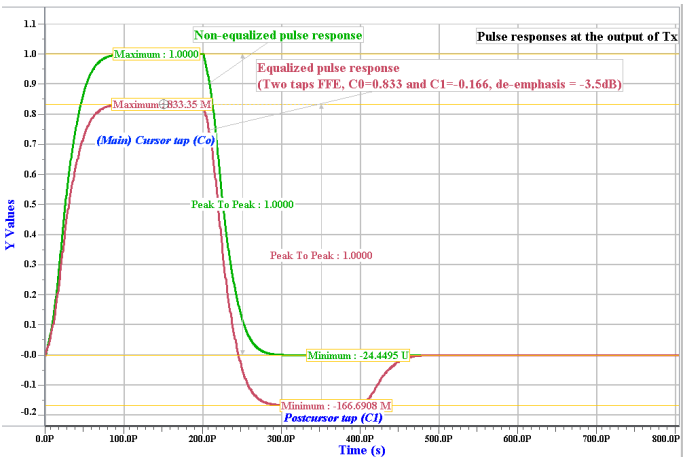


Figure 8 – Unequalized and equalized pulse responses at the output of a USB3.1 Gen1 compliant transmitter

Summary | Informative | Additional Info

## Informative

**Informative: Supplementary compliance based information**

### Optimized TX/RX Settings

Search

| # | Channel | Tx De-emphasis Min | Tx De-emphasis Max | FFE De-emphasis | FFE $C_0$ | FFE $C_1$ | CTLE DCGain | CTLE Pole1 Freq | CTLE Pole2 Freq |
|---|---|---|---|---|---|---|---|---|---|
| | Filter | Filter | Filter | Filter | Filter | Filter | Filter | Filter | Filter |
| 1 | Net001 / Net002 | 3dB | 4dB | -3.52183dB | 0.833333 | -0.166667 | NA | 1.95e+009 | 5e+009 |

Figure 9 - Optimized Tx FFE settings as shown in the HyperLynx SerDes Wizard's HTML report

The Tx tap weights are reported in the HTML report generated by the SerDes Wizard in HyperLynx as shown in Figure 9.

In Figure 10 we plotted the two pulse responses at the input of the receiver, after a 25 inch long microstrip channel with -9dB differential insertion loss at 2.5GHz (the Nyquist frequency for USB3.1 Gen1).

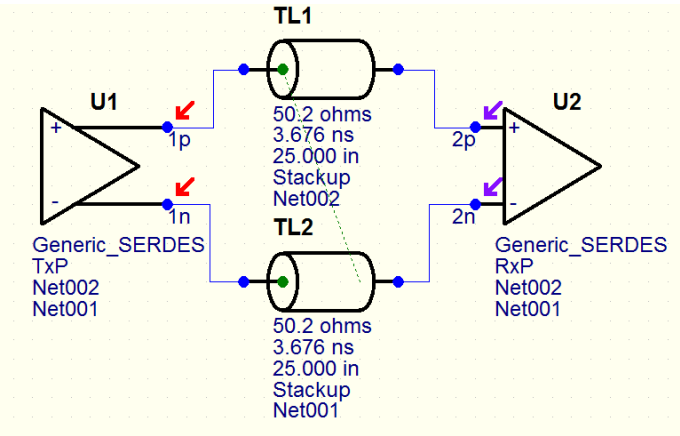


a)

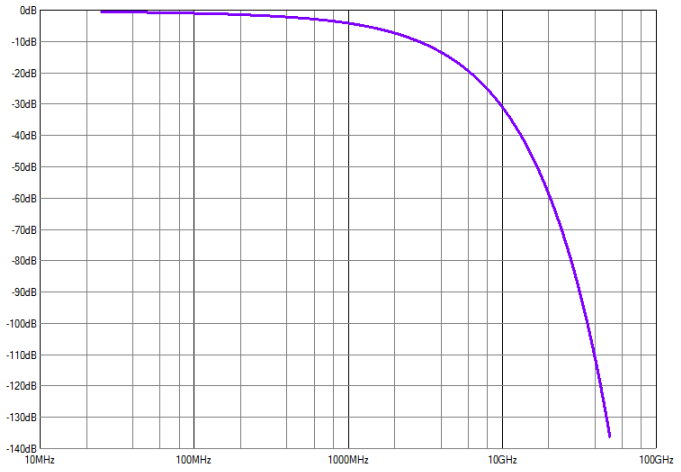


b)

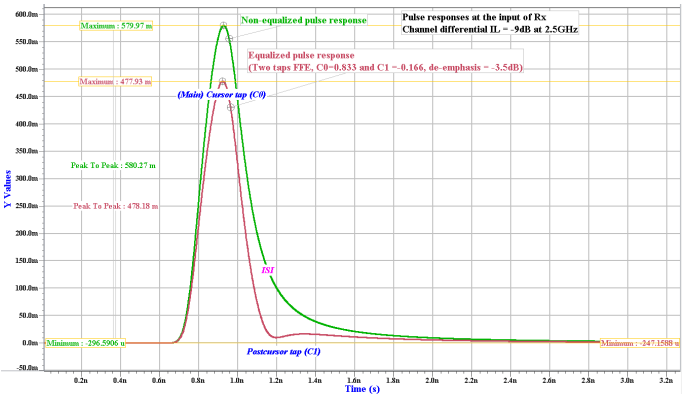


c)

Figure 10 - Pulse responses at the Rx input: schematic diagram (a), differential insertion loss plot (b) and unequalized/equalized pulse responses after the channel

It can be seen that the amplitude of the two pulse responses was attenuated by the channel, but the equalized pulse has lesser ISI compared to the non-equalized pulse. The resulting equalized and non-equalized eye density plots are shown in Figure 11.

Higher data rates require more complex transmitter equalization architectures with a larger number of taps that can compensate both the precursor and postcursor ISI. The PCIe Gen3 specification has introduced a 3-tap FIR filter as shown in Figure 12. For this filter, the tap spacing is 1UI, the main cursor is positive, while the precursor and the postcursor are negative. The next revisions of the PCIe standard, Gen 4 and Gen 5 preserved the same structure.

Along with this architecture, the specification has introduced new terminologies to define various voltage levels of the Tx waveforms that are shaped by Tx equalization. Those voltage levels and the derived ratios are depicted in Figure 13, where the top waveform represents the digital stimulus applied to the transmitter and the bottom waveforms are the emphasized signals (P and N) at the output of the transmitter. Those waveforms were generated from IBIS-AMI simulations using the SerDes Batch Wizard from HyperLynx SI/PI.

Here, the precursor ($V_c$) is called pre-shoot and the postcursor ($V_b$) is referred to as de-emphasis to retain the terminology from the previous revisions of the specification. When the equalization is turned off, both the precursor and the postcursor are zero; consequently the main cursor is one and the signal reaches the maximum amplitude $V_d$. When equalization is turned on, the voltage swing $V_a$ does not reach the maximum amplitude. With this naming convention, the specification defines the following equalization ratios measured in dB:

- De-emphasis $= 20 * \log_{10} \frac{Vd}{Vb}$ (dB) (6)
- Pre-shoot $= 20 * \log_{10} \frac{Vb}{Va}$ (dB) (7)
- Boost $= 20 * \log_{10} \frac{Vc}{Vb}$ (dB) (8)

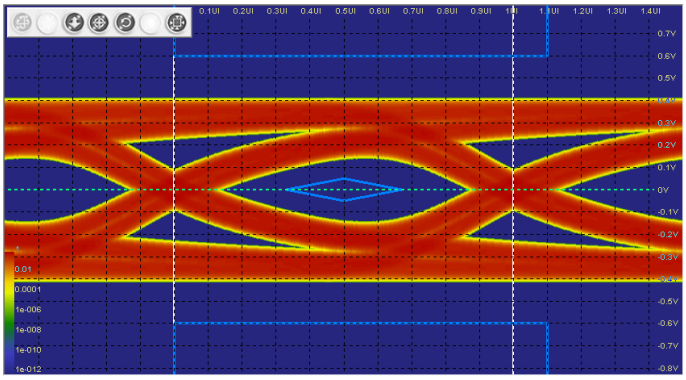

a) De-emphasis turned off – EH=270mV, EW=0.72UI

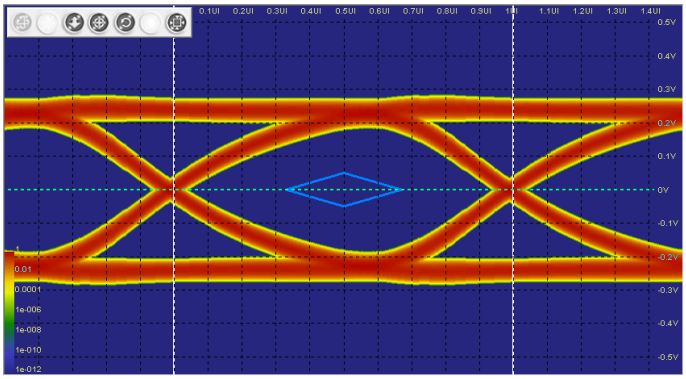

b) Deemphasis = -3.5dB, EH=350mV, EW=0.89UI

Figure 11 – Eye density plots at the input of Rx: de-emphasis off a), de-emphasis -3.5dB

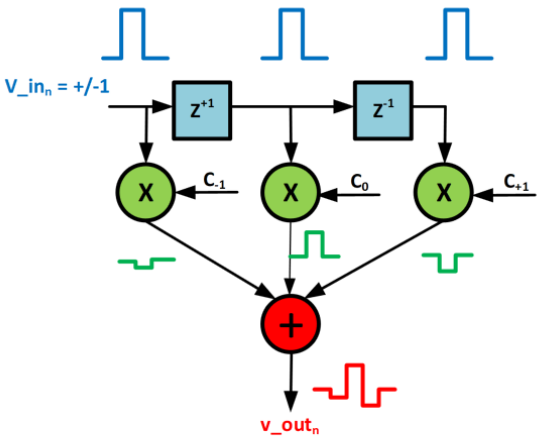


$V\_out_n = v\_in_{n-1}c_{n-1} + v\_in_n c_n + v\_in_{n+1}c_{n+1}$

$|c_{-1}|+|c_0|+|c_{+1}|=1,\ c_{+1} \leq 0$ and $c_{-1} \leq 0$

Figure 12 - PCIe Gen3, Gen4 and Gen5 Tx Equalization FIR

A simple relationship between voltages and tap coefficients can be established if we normalize the maximum voltage swing of the transmitter to 1. In this case $V_d = 1$ and the other voltages can be calculated using the following equations:

$$V_a = |C_0| - |C_{-1}| + |C_1| \quad (9)$$

$$V_b = |C_0| - |C_{-1}| - |C_1| \quad (10)$$

$$V_c = |C_0| + |C_{-1}| - |C_1| \quad (11)$$

The tap coefficients are defined as fractions of the maximum voltage swing of the transmitter and usual resolutions are 1/24 or 1/63. As a result the Tx equalization space is made of discrete values for de-emphasis, pre-shoot and boost. While the previous revision of the specification (PCIe Gen2) specified only two fixed selectable levels of Tx de-emphasis (-3.5dB and -6dB with +/-0.5 tolerance), the PCIe Gen3, 4 and 5 defines a set of values for tap coefficients, called presets. The 11 presets along with the corresponding de-emphasis, pre-shoot and voltage values are shown in Table 1.

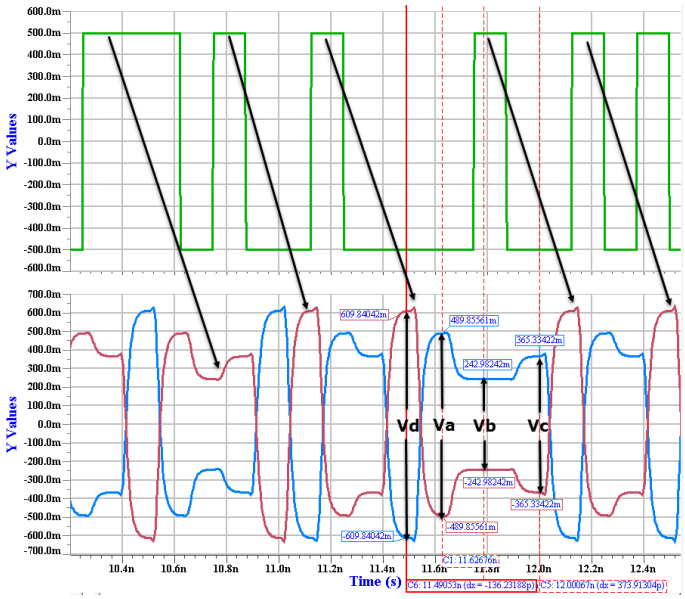


Figure 13 - Plots of the PCIe Gen3 Tx equalization with Preset 7 applied: digital waveforms top, analog waveforms bottom

| Preset # | C-1 | C0 | C1 | Va | Vb | Vc | De-emphasis (dB) | Pre-shoot (dB) |
|---|---|---|---|---|---|---|---|---|
| P0 | 0.000 | 0.750 | -0.250 | 1.000 | 0.500 | 0.500 | -6.0 +/-1.5 | 0.0 |
| P1 | 0.000 | 0.833 | -0.166 | 1.000 | 0.668 | 0.668 | -3.5 +/-1.0 | 0.0 |
| P2 | 0.000 | 0.800 | -0.200 | 1.000 | 0.600 | 0.600 | -4.4 +/-1.5 | 0.0 |
| P3 | 0.000 | 0.875 | -0.125 | 1.000 | 0.750 | 0.750 | -2.5 +/-1.0 | 0.0 |
| P4 | 0.000 | 1.000 | 0.000 | 1.000 | 1.000 | 1.000 | 0.0 | 0.0 |
| P5 | -0.100 | 0.900 | 0.000 | 0.800 | 0.800 | 1.000 | 0.0 | 1.9 +/-1.0 |
| P6 | -0.125 | 0.875 | 0.000 | 0.750 | 0.750 | 1.000 | 0.0 | 2.5 +/-1.0 |
| P7 | -0.100 | 0.700 | -0.200 | 0.800 | 0.400 | 0.600 | -6.0 +/-1.5 | 3.5 +/-1.0 |
| P8 | -0.125 | 0.750 | -0.125 | 0.750 | 0.500 | 0.750 | -3.5 +/-1.0 | 3.5 +/-1.0 |
| P9 | -0.166 | 0.833 | 0.000 | 0.668 | 0.668 | 1.000 | 0 | 3.5 +/-1.0 |
| P10 | 0.000 | | Note | | | | Note | 0.0 |

Table 1 – Tx Equalization Presets for PCIe Gen3, 4 and 5

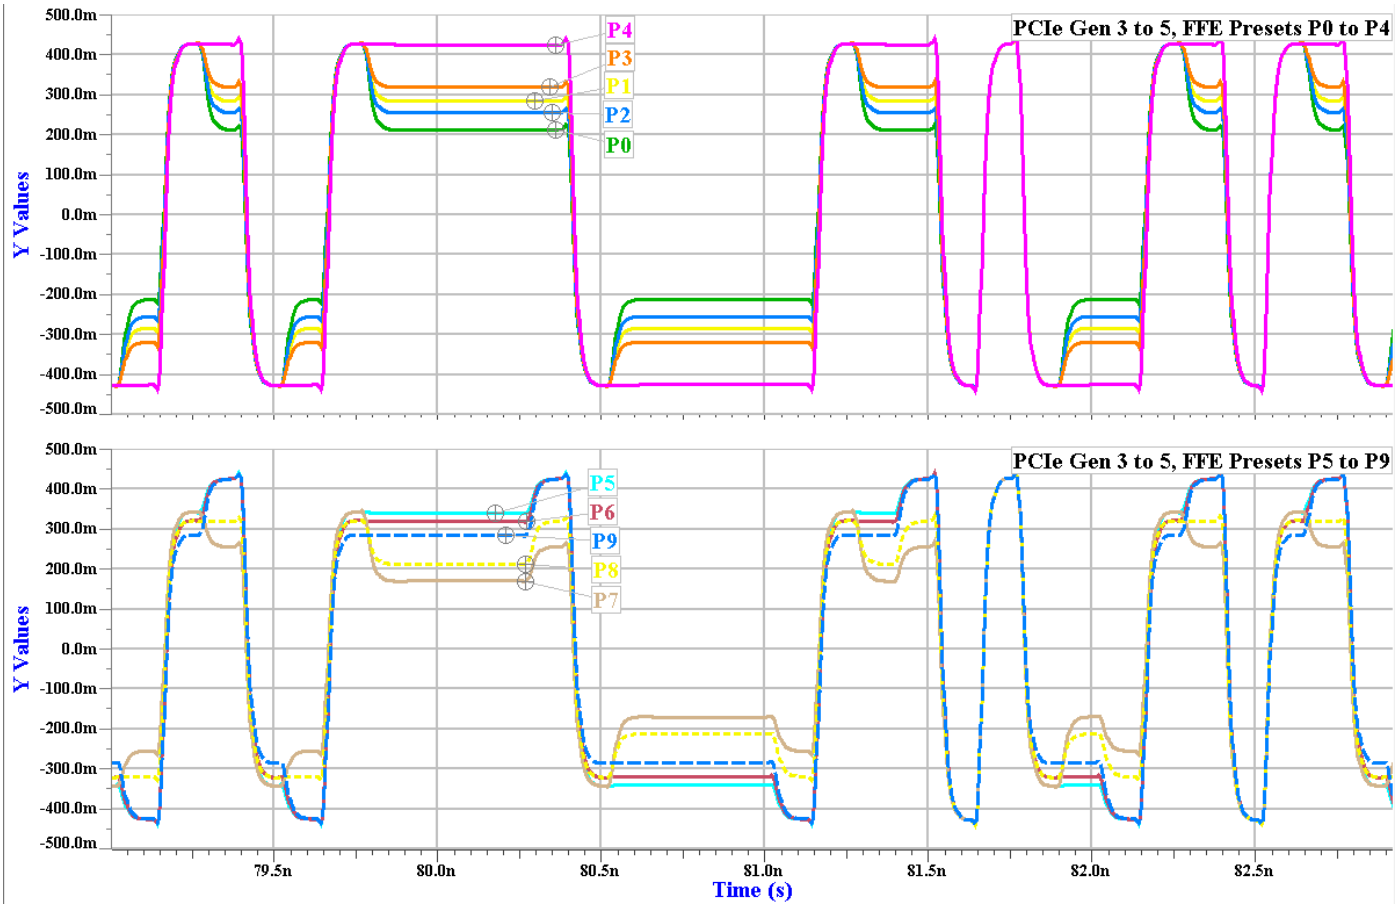


Figure 14 – The effects of the equalization presets on the waveforms plotted at the output of a PCIe Gen3, 4 or 5 compliant transmitter: P0 to P4 on the top waveforms, P5 to P9 on the bottom waveforms

The effects of the first ten presets on the signal at the output of the transmitter are shown in Figure 14.

Similarly, the effects of the first ten presets on the eye density plots at the output of a compliant PCIe Gen 3, Gen 4 or Gen 5 transmitter is depicted in Figure 15.

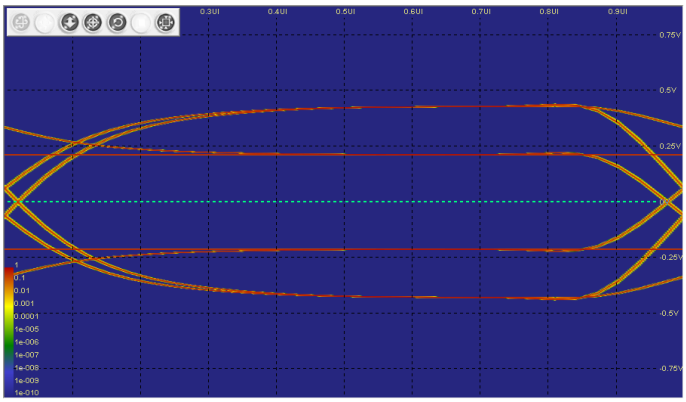

a) P0

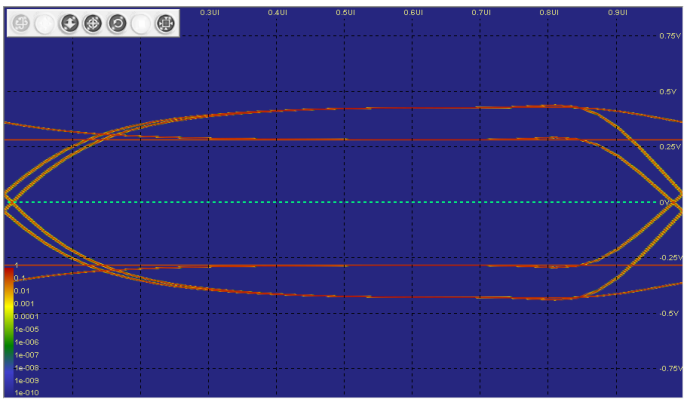

b) P1

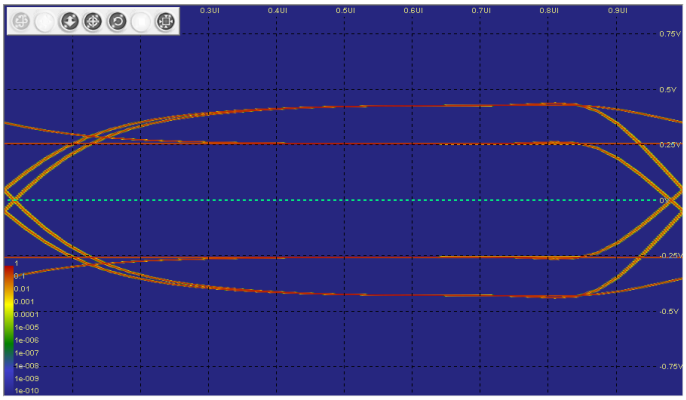

c) P2

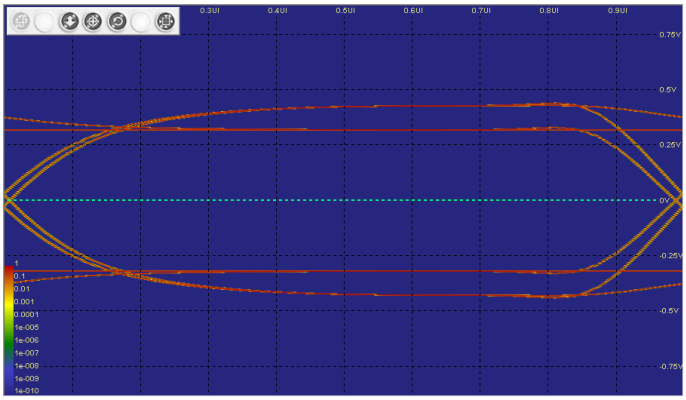

d) P3

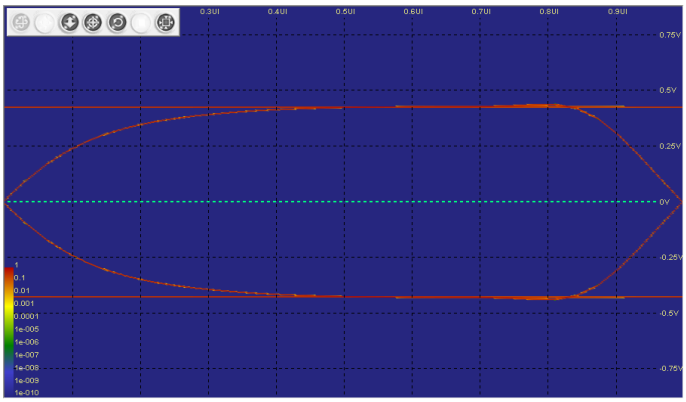

e) P4

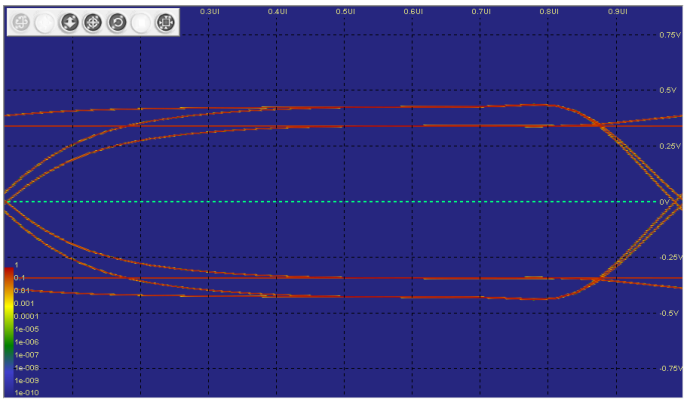

f) P5

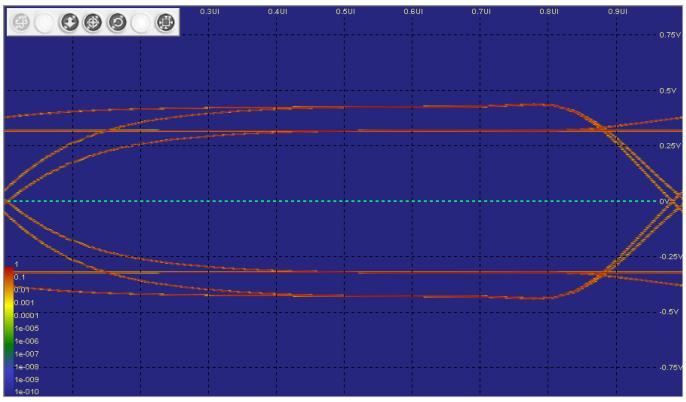

g) P6

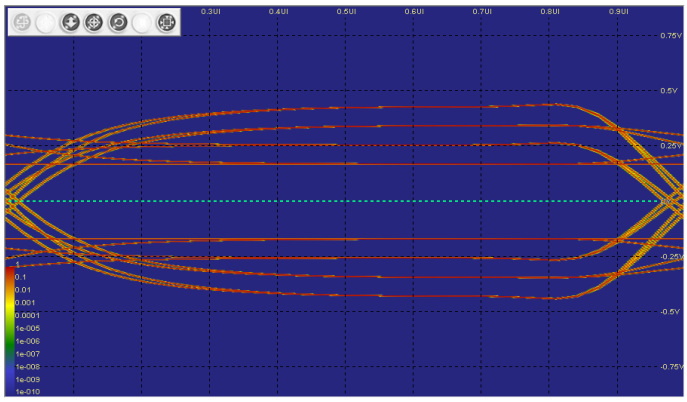

h) P7

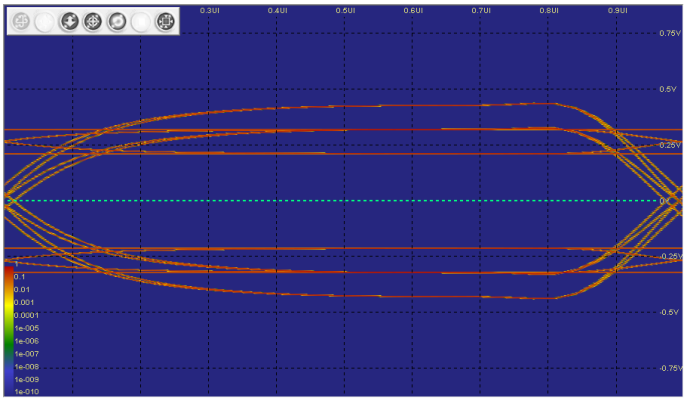

i) P8

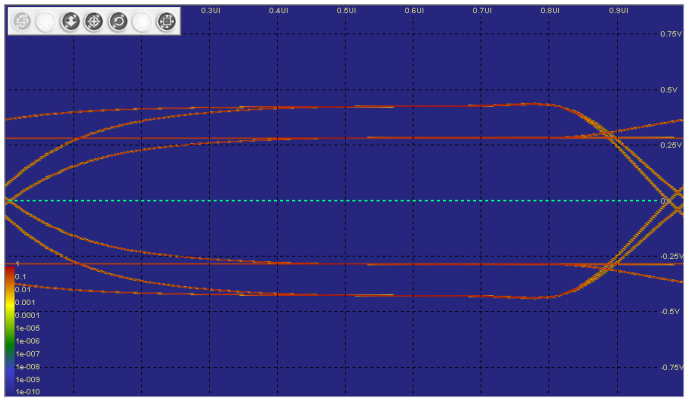

j) P9

Figure 15 - The effects of the equalization presets on the eye density plots at the output of a PCIe Gen3, 4 or 5 compliant transmitter: P0 a), P1 b), P2 c), P3 d), P4 e), P5 f), P6 g), P7 h), P8 i), and P9 j)

The note in the last row of the Table 1 specifies that for P10 the boost limits are not fixed. Moreover, the specification poses several constraints for the transmitter equalized output voltage that must meet the full swing, max reduced swing and min reduced swing limits. All those considerations lead to a triangular matrix [1], depicted in Table 2, which defines the equalization coefficient space.

Min Reduced Swing Limit

| PS DE | | $C_{-1}$ | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| BOOST | | 0/24 | 1/24 | 2/24 | 3/24 | 4/24 | 5/24 | 6/24 | 7/24 | 8/24 |
| $C_{-1}$ | 0/24 | 0.0 0.0 0.0 | 0.0 -0.8 0.8 | 0.0 -1.6 1.6 | 0.0 -2.5 2.5 | 0.0 -3.5 3.5 | 0.0 -4.7 4.7 | 0.0 -6.0 6.0 | 0.0 -7.6 7.6 | 0.0 -9.5 9.5 |
| | 1/24 | 0.8 0.0 0.8 | 0.8 -0.8 1.6 | 0.9 -1.7 2.5 | 1.0 -2.8 3.5 | 1.2 -3.9 4.7 | 1.3 -5.3 6.0 | 1.6 -6.8 7.6 | 1.9 -8.8 9.5 | |
| | 2/24 | 1.6 0.0 1.6 | 1.7 -0.9 2.5 | 1.9 -1.9 3.5 | 2.2 -3.1 4.7 | 2.5 -4.4 6.0 | 2.9 -6.0 7.6 | 3.5 -8.0 9.5 | | |
| | 3/24 | 2.5 0.0 2.5 | 2.8 -1.0 3.5 | 3.1 -2.2 4.7 | 3.5 -3.5 6.0 | 4.1 -5.1 7.6 | 4.9 -7.0 9.5 | | | |
| | 4/24 | 3.5 0.0 3.5 | 3.9 -1.2 4.7 | 4.4 -2.5 6.0 | 5.1 -4.1 7.6 | 6.0 -6.0 9.5 | | | | |
| | 5/24 | 4.7 0.0 4.7 | 5.3 -1.3 6.0 | 6.0 -2.9 7.6 | 7.0 -4.9 9.5 | | | | | |
| | 6/24 | 6.0 0.0 6.0 | 6.8 -1.6 7.6 | 8.0 -3.5 9.5 | | | | | | |

Full swing Limit or
Max reduced swing limit

Table 2 - Transmit Equalization Coefficient Space for PCIe Gen3, 4 and 5 [1]

| Channel differential IL (dB) | C-1 | C0 | C1 | Va | Vb | Vc | De-emphasis (dB) | Pre-shoot (dB) |
|---|---|---|---|---|---|---|---|---|
| <3.5 | 0.000 | 0.875 | -0.125 | 1.000 | 0.750 | 0.750 | -2.5 | 0.0 |
| ≥3.5 | -0.100 | 0.775 | -0.125 | 0.800 | 0.550 | 0.750 | -3.3 | 2.7 |

Table 3 – USB3.1 Gen2, informative transmitter equalization settings

The USB 3.1 Gen2 running at 10GT/s has borrowed the 3-tap FIR structure from PCIe Gen3 along with the pre-shoot and de-emphasis terminology and definitions. In addition to that, this protocol specification provides recommended (informative) tap coefficient values for channels with different differential insertion loss at the Nyquist frequency (5GHz) as shown in Table 3.

The effect of those fixed levels of equalization on the time domain waveforms and eye density plots at Rx, after the signal travels two channels, one with -1.6dB differential IL and another one with -5.2dB, is shown in Figure 16. The time domain waveforms at the output of Tx are shown overlapped with those at the input of Rx, for reference.

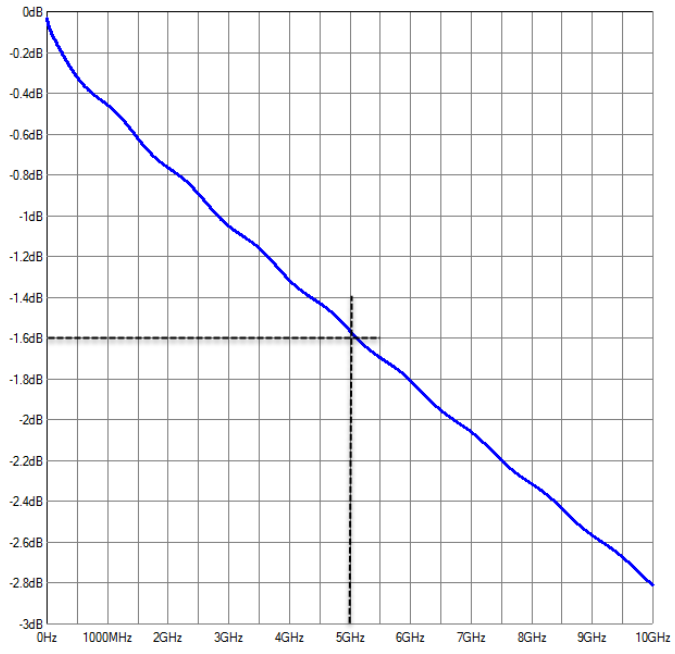


a)

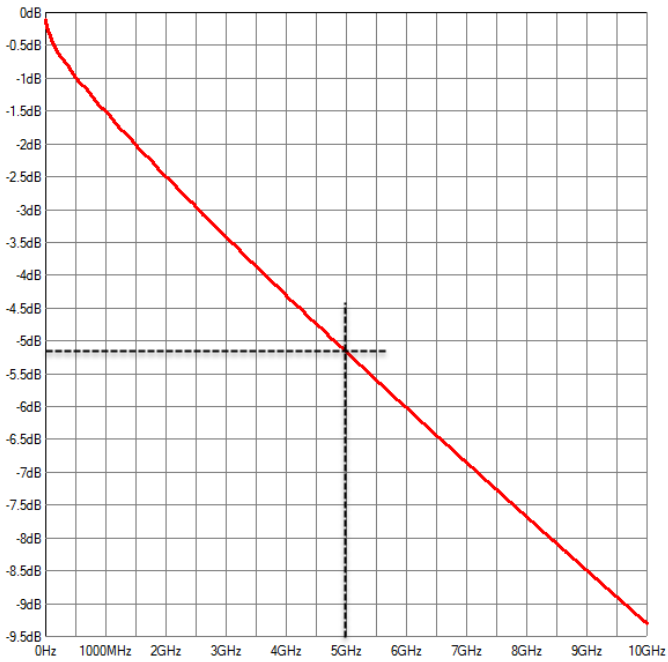


b)

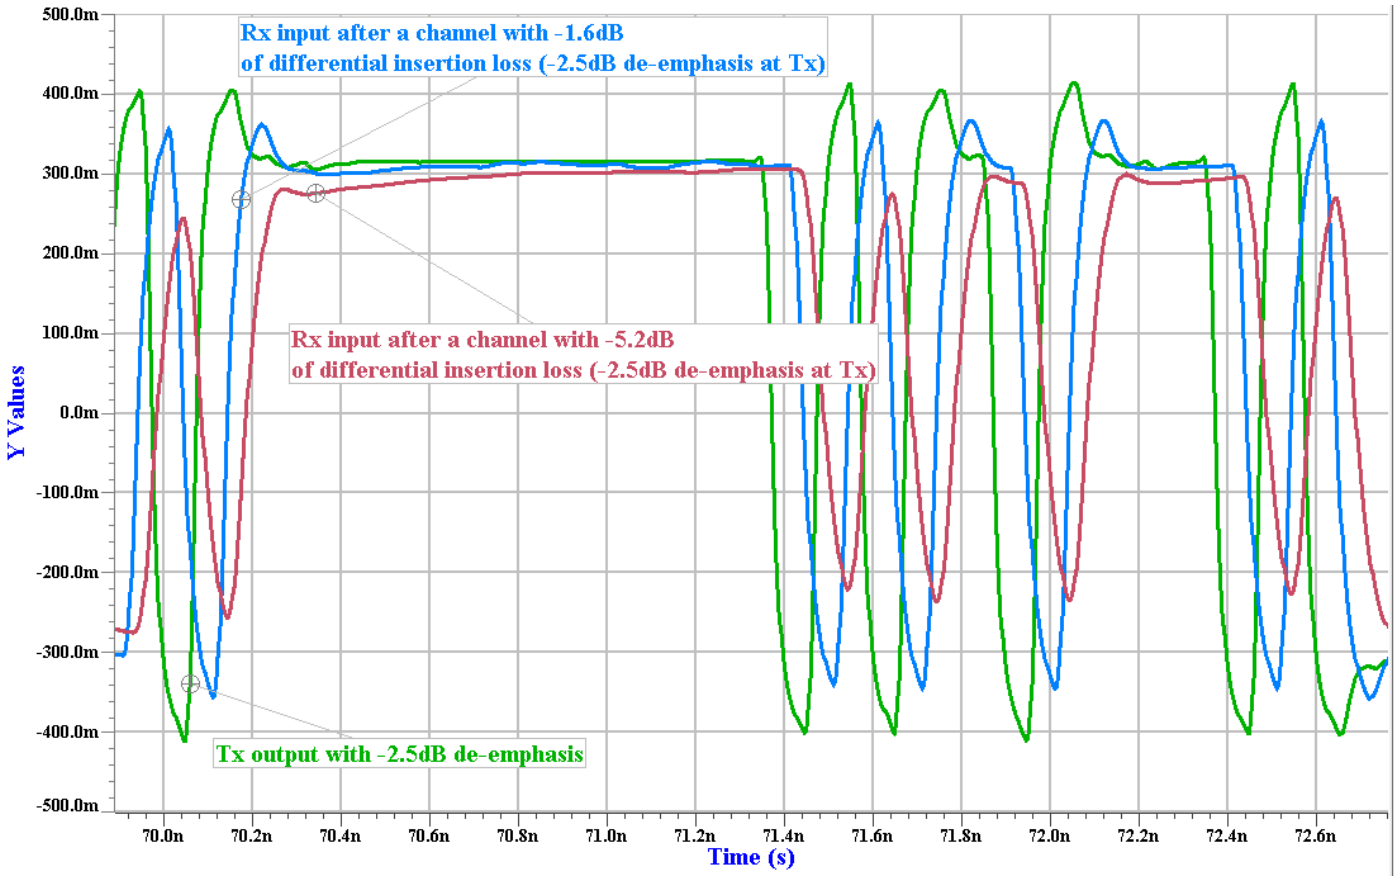
Rx input after a channel with -1.6dB
of differential insertion loss (-2.5dB de-emphasis at Tx)
Rx input after a channel with -5.2dB
of differential insertion loss (-2.5dB de-emphasis at Tx)
Tx output with -2.5dB de-emphasis
Y Values
Time (s)
500.0m
400.0m
300.0m
200.0m
100.0m
0.0m
-100.0m
-200.0m
-300.0m
-400.0m
-500.0m
70.0n
70.2n
70.4n
70.6n
70.8n
71.0n
71.2n
71.4n
71.6n
71.8n
72.0n
72.2n
72.4n
72.6n

c)

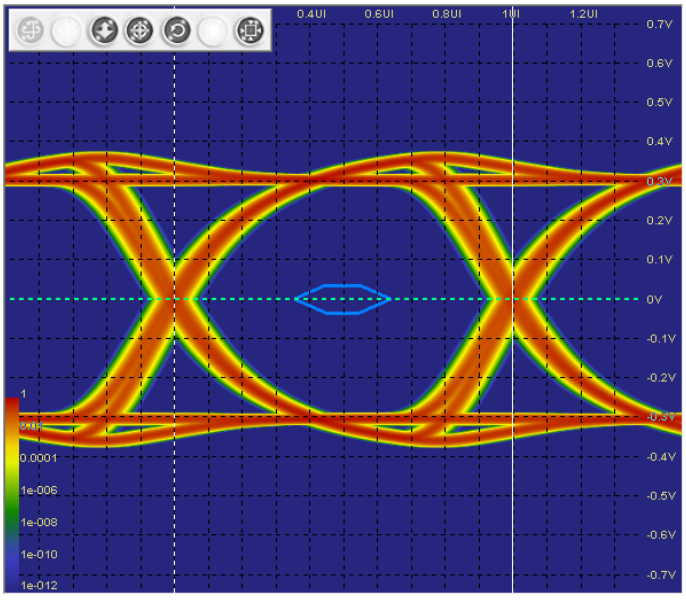
0.4UI
0.6UI
0.8UI
1UI
1.2UI
0.7V
0.6V
0.5V
0.4V
0.3V
0.2V
0.1V
0V
-0.1V
-0.2V
-0.3V
-0.4V
-0.5V
-0.6V
-0.7V
1
0.1
0.01
0.001
0.0001
1e-006
1e-008
1e-010
1e-012

d)

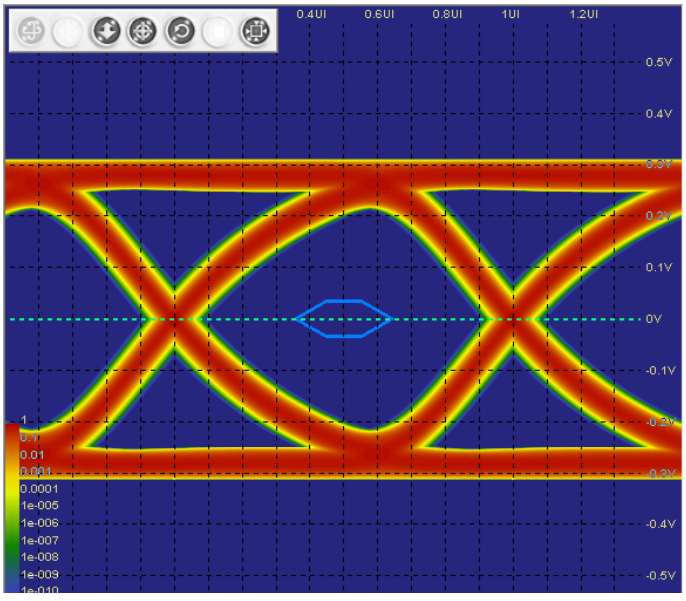
0.4UI
0.6UI
0.8UI
1UI
1.2UI
0.5V
0.4V
0.3V
0.2V
0.1V
0V
-0.1V
-0.2V
-0.3V
-0.4V
-0.5V
1
0.1
0.01
0.001
0.0001
1e-005
1e-006
1e-007
1e-008
1e-009
1e-010

e)

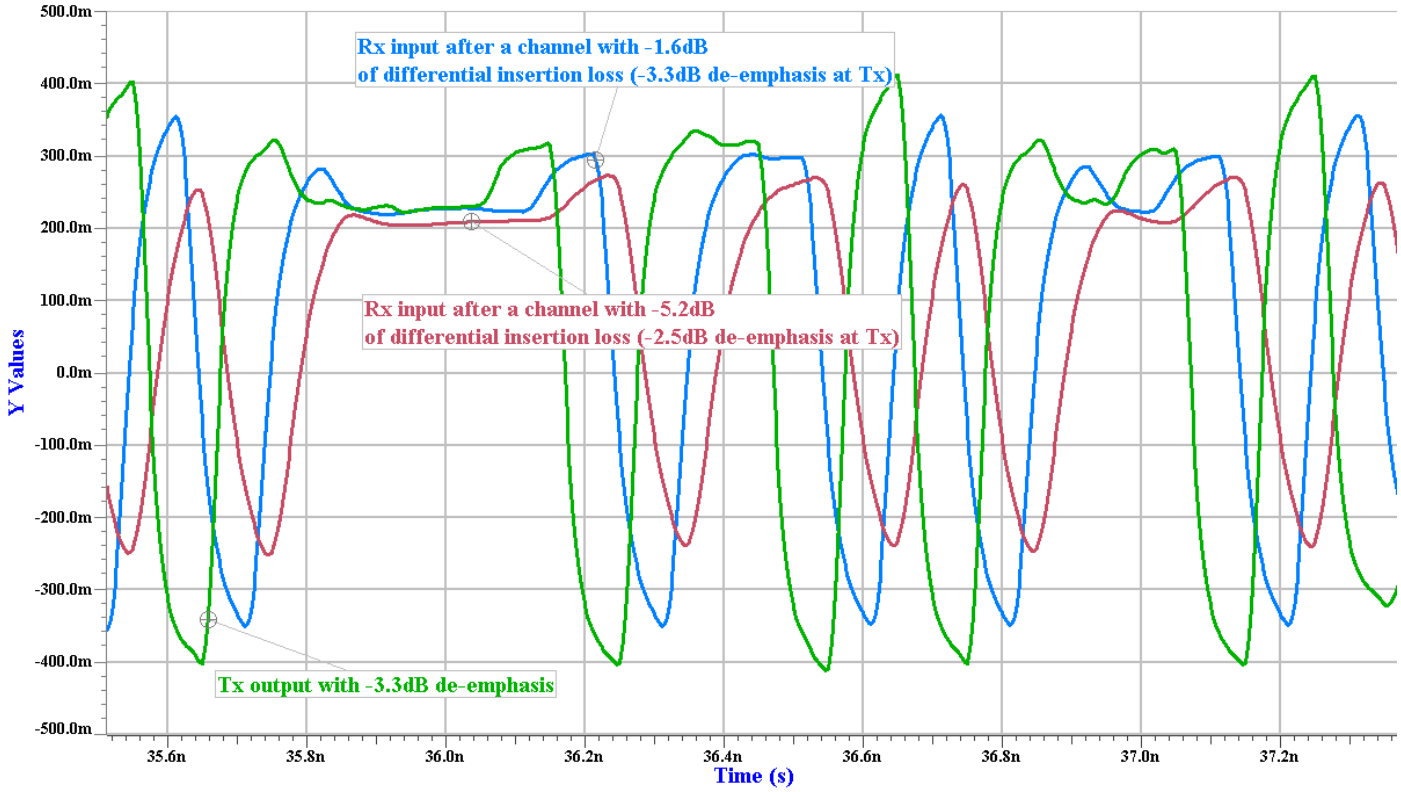


f)

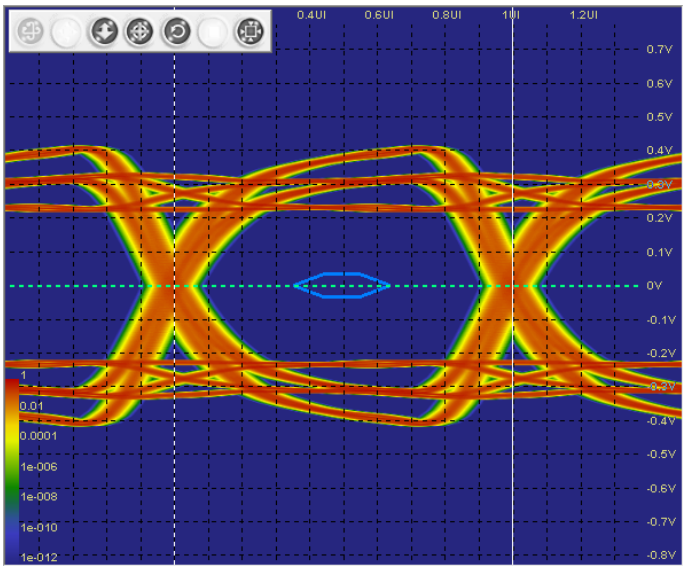


g)

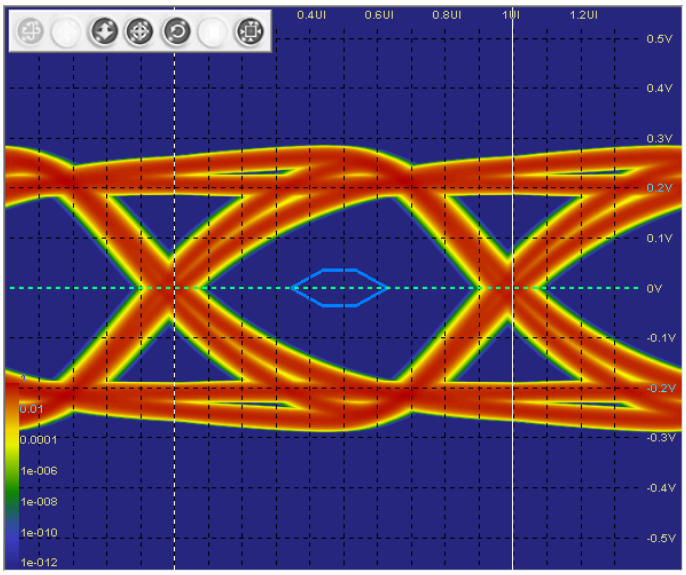


h)

Figure 16 - The effects of the USB3.1 Gen2 fixed equalization levels on the time domain waveforms and eye density plots after two channels with different differential IL losses: channel with -1.6dB differential IL @ 5GHz a), channel with -5.2dB differential IL @ 5GHz b), time domain waveforms at the output of Tx and input of Rx for first channel c), eye density plots at the input of Rx for first channel d) and e), time domain waveforms at the output of Tx and input of Rx for second channel f), eye density plots at the input of Rx for second channel g) and h).

The IEEE802.3ap specification defines three-tap FIR architecture for transmitter de-emphasis as well, which is similar to that described in the PCIe 3.0 to 5.0 specifications. The two 10GBASE-KR and 40GBASE-KR4 Physical Medium Dependent (PMD) introduced in this amendment are using this functional model, with programmable equalization. However the de-emphasis is slightly larger than that of PCIe 3.0, but the pre-shoot is lower. Moreover the constraints imposed on the coefficient limits are different than those for PCIe and are based on the values of the pre-/postcursor equalization ratios $R_{pre}$ / $R_{pst}$ and the restrictions placed on the minimum steady-state ($V_2$) or maximum peak voltage ($V_1$ and $V_3$). The two equalization ratios are constrained as well: -3.75dB ≤ $R_{pre}$ = $V_1/V_2$ ≤ +0.45dB and -12dB ≤ $R_{pst}$ = $V_3/V_2$ ≤ +0.45dB (including tolerance). In this case, the three coefficients are not required to be normalized to 1 ($|C_{-1}|+|C_0|+|C_{-1}|=1$).

All the PMDs introduced in the IEEE802.3bj, bm, by, and bs and supported in the HyperLynx SI/PI tool are using three-tap FIR architecture, that is modelled in the Channel Operating Model algorithm. The valid ranges and values for tap coefficients are specified, for each PMD, in a table containing the COM parameters for that particular PMD. Typical constraints are minimum value for the main cursor and minimum/maximum values and step size for the other cursors. Table 4 shows these types of constraints as they are defined for the 100GBASE-KR4 operating mode in the IEEE802.3bj specification.

| | Minimum | Maximum | Step size |
|---|---|---|---|
| Tx EQ precursor coefficient (c-1) | -0.18 | 0 | 0.02 |
| Tx EQ main cursor coefficient (c0) | 0.62 | - | |
| Tx EQ postcursor coefficient (c1) | -0.38 | 0 | 0.02 |

Table 4 – 100GBASE-KR4 Tx Equalization Constraints for Tap Coefficients

Newer specifications such as IEEE802.3cd have increased the complexity of Tx equalizers, moving from a 3-tap to 4-tap transversal filter as shown in Figure 17.

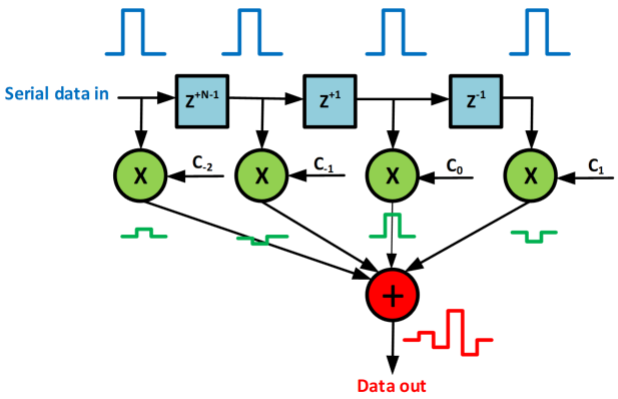


Figure 17 – Transmit equalizer functional model introduced in the IEEE802.3cd specification

Each of the six PMDs depicted in sections 136 and 137 of this amendment are using the 4-tap transmit equalizer functional model with two precursors and one postcursor:

- 50GBASE-CR, 100GBASE-CR2, and 200GBASE-CR4
- 50GBASE-KR, 100GBASE-KR2, and 200GBASE-KR4

The simulated equalized pulse responses for four different combinations of values for the four tap coefficients are plotted in Figure 18. The pulse responses are plotted at the input of a receiver after a channel with -16dB differential insertion loss at Nyquist frequency of 13.28125GHz. From this figure it can be seen that the amplitude of the equalized pulse responses is smaller than the amplitude of the non-equalized pulse, due to the channel attenuation of the signal. However, the width of the equalized pulse responses is smaller, closer to 1UI, which shows that the ISI is reduced. Moreover, the optimized pulse response is not the one with the highest amplitude, but the one that produces the best signal to noise ratio at Rx.

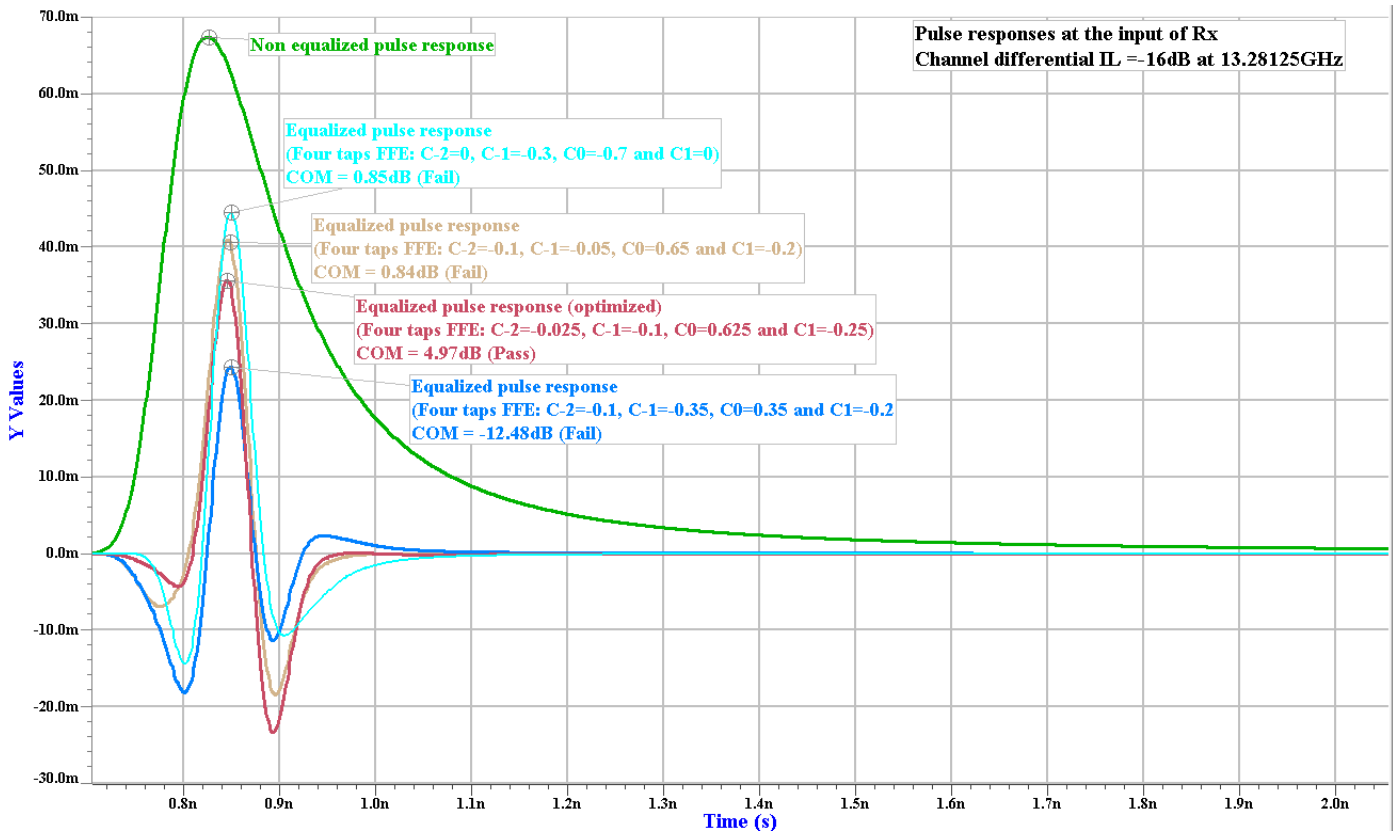


Figure 18 – Equalized and non-equalized pulse responses of a four taps FIR after a channel with -16dB differential insertion loss at Nyquist frequency

The eye diagrams generated out of the COM algorithm, corresponding to the above pulse responses, are plotted in Figure 19. In these experiments, one case (d) was intentionally chosen such that the constraint for the minimum cursor coefficient is violated ($c_{-1} \leq 0.6$). As a result, the COM value is negative and the corresponding eye density plot is completely closed.

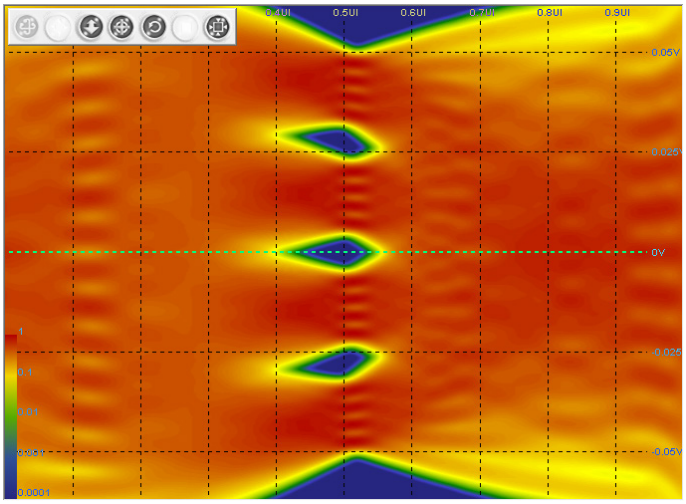


a) $C_{-2}$=0, $C_{-1}$=0.3, $C_0$=0.65, $C_{+1}$=0 /COM= 0.84dB

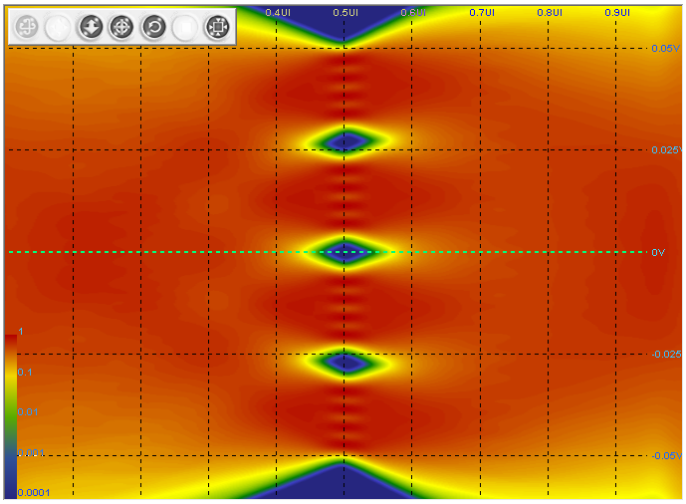


b) $C_{-2}$=-0.1, $C_{-1}$=-0.05, $C_0$=0.65, $C_{+1}$=-0.2/COM=0.84dB

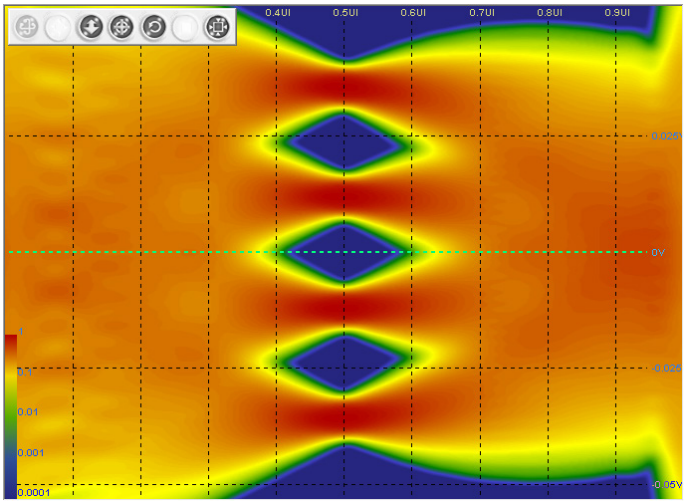


c) $C_{-2}$=-0,025, $C_{-1}$=-0.1, $C_0$=0.625, $C_{+1}$=-0.25/COM =4.97dB

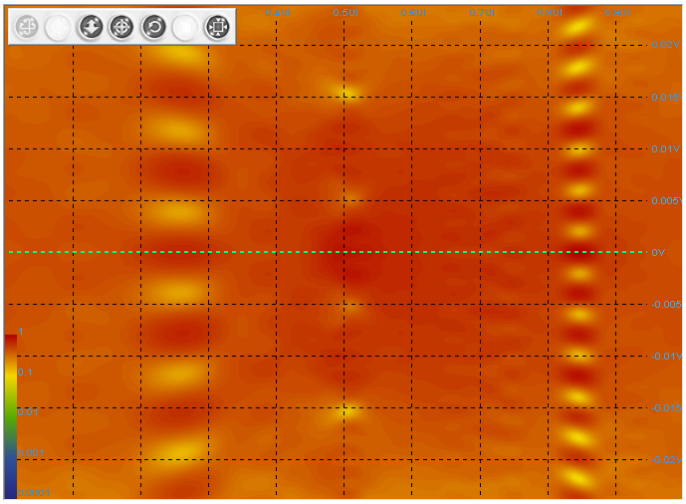


d) $C_{-2}$=-0.1, $C_{-1}$=-0.35, $C_0$=0.35, $C_{+1}$=-0.2/COM=-12.48dB

Figure 19 – Eye diagrams corresponding to four different sets of tap coefficients of a four taps FIR, after a channel with -16dB differential insertion loss at Nyquist frequency: $C_{-2}$=0, $C_{-1}$=0.3, $C_0$=0.65, $C_{+1}$=0 a), $C_{-2}$=-0.1, $C_{-1}$=-0.05, $C_0$=0.65, $C_{+1}$=-0.2 b), $C_{-2}$=-0,025, $C_{-1}$=-0.1, $C_0$=0.625, $C_{+1}$=-0.25 c), $C_{-2}$=-0.1, $C_{-1}$=-0.35, $C_0$=0.35, $C_{+1}$=-0.2 d)

# Receiver Continuous Time Linear Equalizer (CTLE)

The equalization architecture of the SerDes protocols operating at data rates above 5/8Gbps usually includes a Continuous Time Linear Equalizer (CTLE) block situated on the receiver side. As opposed to FIRs that are discrete devices, CTLEs are analog circuits that are operating continuously and have the characteristics of a high-pass filter. They can be implemented in the real hardware either using passive components or using active components. Instead of amplifying the high frequency components of the signal, passive CTLEs attenuate the low frequency components. Active CTLEs might provide some signal gain. Most of the CTLE architectures that are described in the SerDes protocol specifications are implemented with passive components integrated into the silicon. However the filter can be located anywhere in the channel and can even be built into cables or connectors. CTLE, in combination with other equalization blocks like DFE, can improve the quality of the received signal in a way that is not possible with DFE alone.

A typical transfer function of a CTLE with a single zero and two poles, is expressed in a mathematical format by equation (12) and it is depicted in Figure 20.

$$H_{ctf}(s) = k * \frac{(s+\omega_z)}{(s+\omega_{p1})(s+\omega_{p2})} \quad (12)$$

Where the complex frequency s = jω, k is the broadband gain of the CTLE, $\omega_z = 2\pi f_z$ is the zero angular frequency, $\omega_{P1} = 2\pi f_{P1}$ is the first pole angular frequency, $\omega_{P2} = 2\pi f_{P2}$ is the second pole angular frequency. The low frequency components from DC up to zero frequency will be attenuated by the value of the DC gain (Adc), which in this particular example is -12dB. The frequency components above the zero frequency will be amplified by the high frequency peak gain (Aac). The ratio of the amplified frequencies over the low frequencies is an important characteristic of the filter.

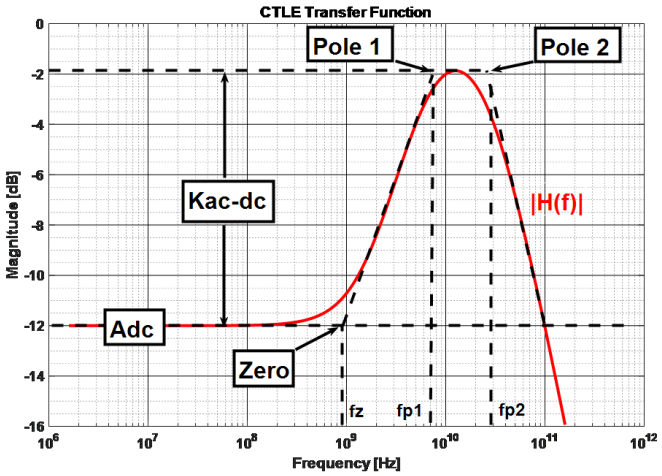


Figure 20 – Bode plot of a CTLE filter

The combined transfer functions of a printed circuit board interconnect (PCB) and the CTLE can be computed using equation (13) and is shown in Figure 21.

$$H(f) = H_{21PCB}(f)H_{ctf}(f) \quad (13)$$

where $H_{21PCB}(f)$ is the voltage transfer function of the terminated signal path.

From equation (12) and Figure 20 it can be seen that there are two controllable parameters that are affecting the performance of the CTLE: the peak gain and its location relative to the Nyquist frequency (half of the protocol data rate). In order to optimize the effect of the CTLE on the combined transfer function, the values of the two parameters shall be selected based on the channel characteristics, especially the differential insertion loss profile.

By applying the inverse Fourier transform to the above transfer functions we can convert from frequency-domain (FD) to time-domain (TD) and compute the impulse, step and pulse responses of the channel with and without CTLE effects. Figure 22 shows the effect of a passive CTLE on the pulse response of a channel. Since the gain of the CTLE at the Nyquist frequency is negative, in this particular case, the amplitude of the pulse shaped by the CTLE is lower, but the benefit of applying equalization consists on a narrower pulse reducing or completely eliminating the adjacent bit interference. The linear equalizer can suppress the long-tail ISI.

The effect of CTLE on a pseudo-random binary sequence (PRBS) pattern is depicted in Figure 23, where the digital waveform is shown in green color, the unequalized waveform in red and the equalized signal in blue. In the red waveform it can be seen the effect of the reduced amplitude and pulse-spreading for single bit transitions. Due to the limited channel bandwidth, the high frequency components of the signal are attenuated and the data stream loses sharpness. The transitions from one logic level to another take longer to complete causing the interference between adjacent bits, thus reducing the opening of the resulting eye diagrams and increasing the bit error rate (BER). In the blue waveforms, the amplitude of all transitions is equalized and the transitions are sharper; consequently the pulse spreading across bit boundaries is reduced. As a result, the equalized eye (c) has larger eye height (EH) and eye width (EW) than the unequalized eye (b).

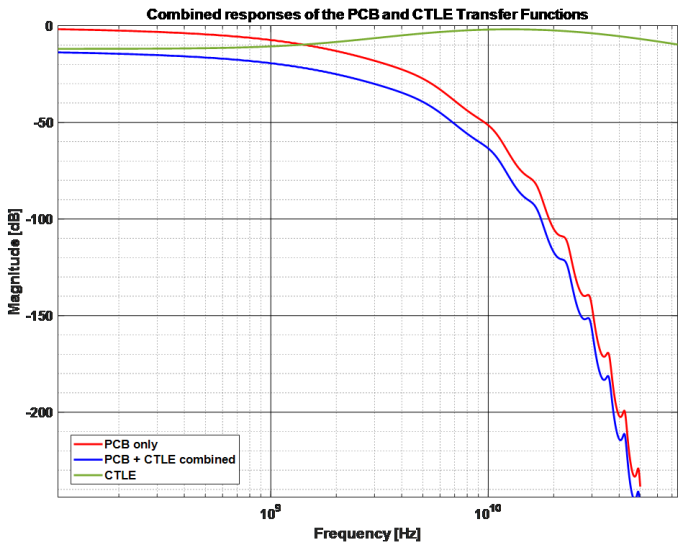


Figure 21 – The System Transfer Function with CTLE Effects Included

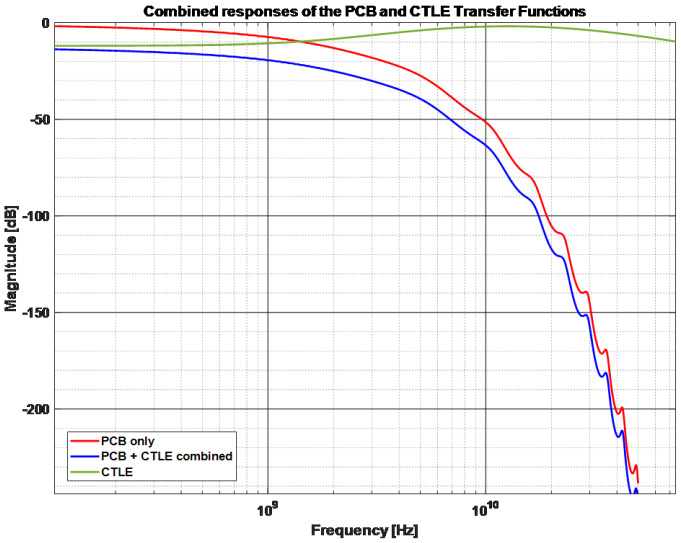


Figure 22 – Pulse responses of the channel with and without CTLE effects

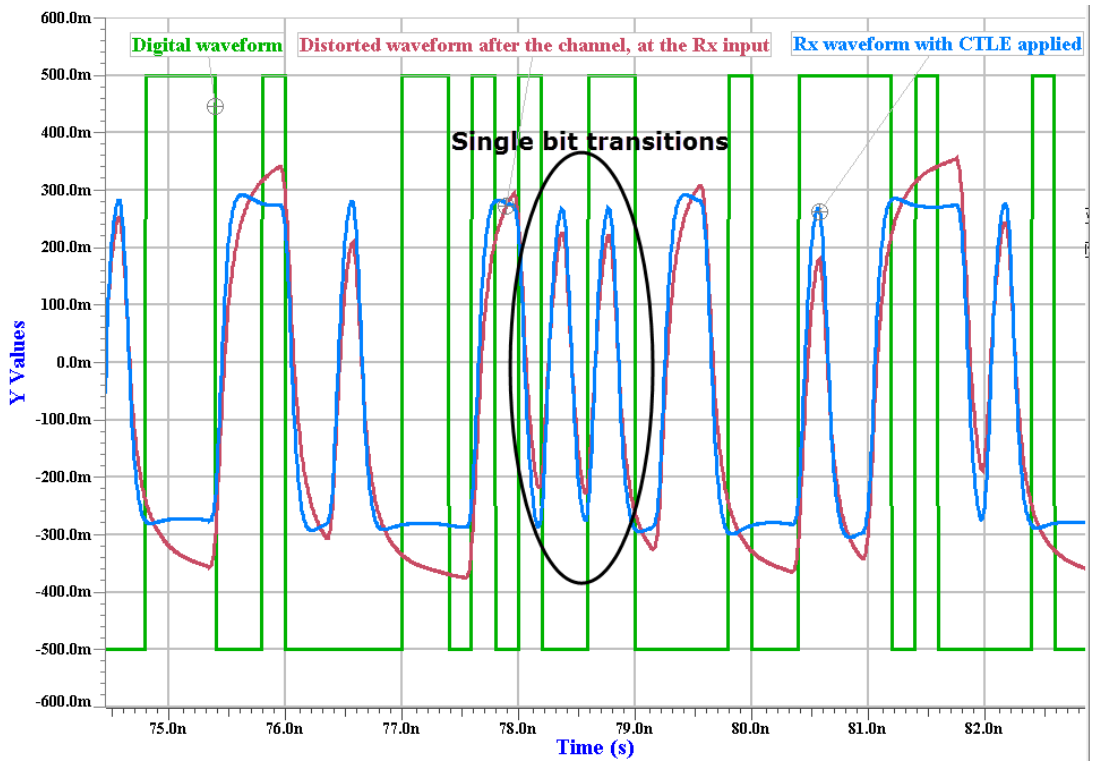


a)

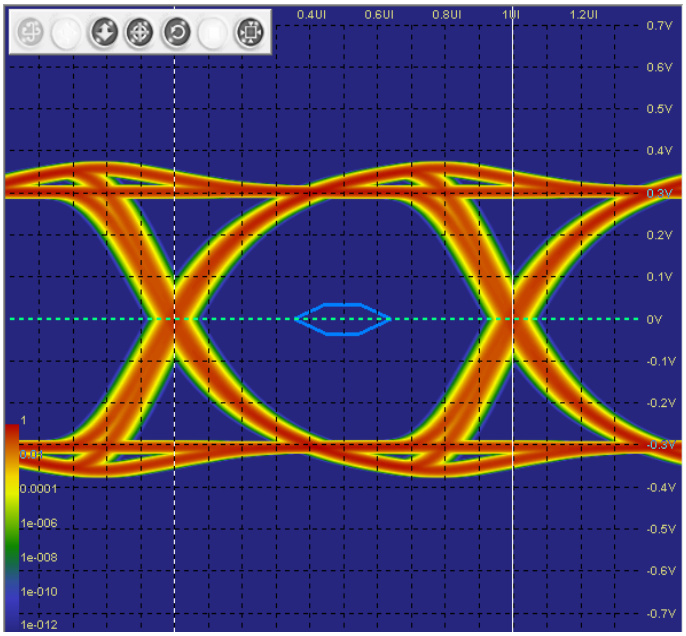


b)

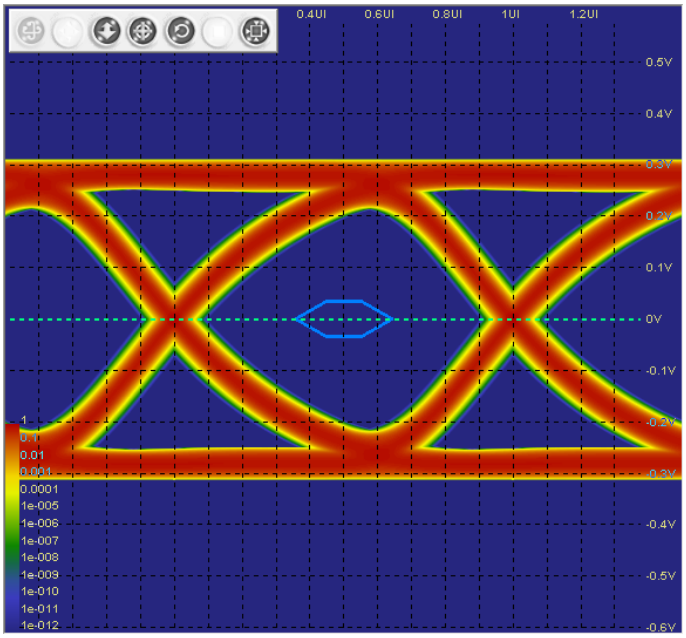


c)

Figure 23 – The effects of CTLE on the time-domain waveforms a), eye density plot without CTLE b), eye density plot with CTLE applied c)

A generalized CTLE transfer function can be expressed in the form of the following equation [4]:

$$H_{ctf(S)} = k * \frac{\prod_{n=1}^{m}(1+\frac{s}{fz_n})}{\prod_{n=1}^{k}(1+\frac{s}{fp_n})\prod_{n=1}^{y}(1+\frac{s}{fc_n})}, \quad (14)$$

Where, s = jω, k is the broadband gain of the CTLE, m is the number of adjustable zeros, k the number of adjustable poles, y the number of fixed poles, $f_{pn}$ $n^{th}$ adjustable pole's frequency, $f_{cn}$ $n^{th}$ fixed pole's frequency and $f_{zn}$ $n^{th}$ zero's frequency.

References [5], [10], [12], and [13] describe, in great detail, several possible hardware implementations for passive and active linear equalizers. A popular model for continuous-time equalizers is the split-path. The block diagram of a generic split-path equalizer is shown Figure 24.

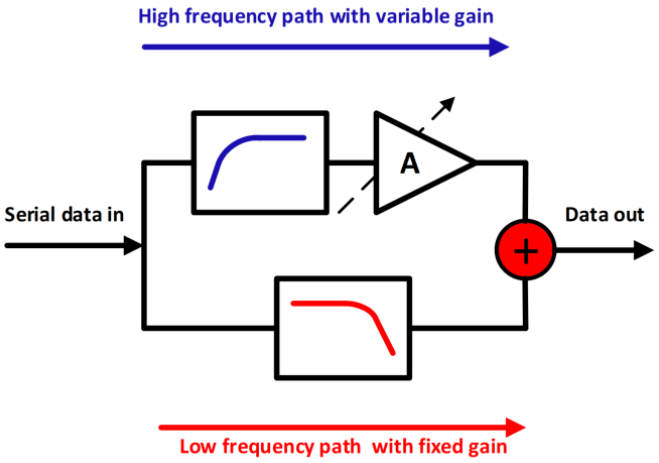


Figure 24 – Block diagram of a split-path CTLE

As depicted in the block diagram, the incoming serial data bits can take two different paths:

- A high-frequency path, which has an adaptive loop with variable gain;
- A low-frequency path with fixed gain or attenuation.

The two signals are then summed in the final stage to generate the equalized output.

A circuit implementation using passive components is depicted in Figure 25.

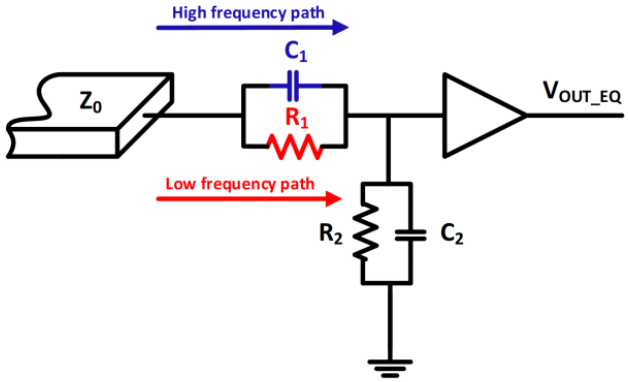


Figure 25 – Circuit implementation of a passive CTLE

In the passive circuit implementation, the capacitor C1 provides the high-frequency path for the signal, while the resistor $R_1$ attenuates the low-frequency components, with the overall result being the high frequency gain boosting. The transfer function for this circuit is given in equation [15]:

$$H_{ctf}(s) = \frac{R_2}{R_1+R_2} * \frac{1+R_1C_1s}{1+\frac{R_1R_2}{R_1+R_2}(C_1+C_2)s} \quad (15)$$

From the above equation, the pole and zero frequencies can be determined as:

$$\omega_p = \frac{1}{\frac{R_1R_2}{R_1+R_2}(C_1+C_2)} \quad (16)$$

$$\omega_z = \frac{1}{R_1R_2} \quad (17)$$

The two gains $A_{dc}$ and $A_{ac}$ are easily derived if we consider the capacitors as being open circuits at DC and low impedance elements at high frequency. With that consideration in mind, the two gains can be expressed as [10]:

$$A_{dc} = \frac{R_2}{R_1+R_2} \quad (18)$$

$$A_{ac} = \frac{C_1}{C_1+C_2} \quad (19)$$

The high frequency gain boosting factor (peaking) can be computed as the ratio of the two gains [12], [13]:

$$Peaking = \frac{A_{ac}}{A_{dc}} = \frac{\omega_p}{\omega_z} = \frac{R_2}{R_1+R_2}\frac{C_1}{C_1+C_2}$$

Similarly, a circuit implementation of a split-path model for an active CTLE is shown in Figure 26 [5].

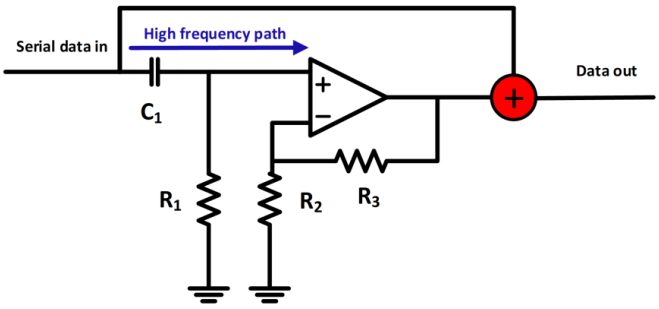


Figure 26 – Circuit implementation of an active CTLE

Once again, the high frequency content of the signal is extracted from the incoming data bits through a high-pass filter made off $C_1$ and $R_1$ and it is fed into the amplifier which has $R_2$ and $R_3$ in the feedback loop. The cutoff frequency of the high-pass filter is:

$$f_{cutoff} = 1 + \frac{1}{R_1 C_1} \tag{20}$$

And the voltage gain of the active filter is:

$$Gain = 1 + \frac{R_2}{R_3} \tag{21}$$

The output of the amplifier is summed with the incoming signal to generate the equalized response. The overall transfer function of the circuit is [5]:

$$H_{ctf}(f) = \frac{1+\frac{R_2}{R_3}}{1+\frac{1}{2\pi f R_1 C_1}}, \tag{22}$$

Regardless of its implementation, the CTLE with optimal equalization shall have the transfer function equal to the inverse transfer function of the channel.

# Examples of receiver CTLE architectures

In order to open the eye for long cables and channels, the USB3.1 Gen1 (5GT/s) specification has introduced a 1st order reference Rx CTLE with the transfer function described by the equation (23):

$$H_{(s)} = \frac{A_{dc}\omega_{p1}\omega_{p2}}{\omega_z} * \frac{(s+\omega_z)}{(s+\omega_{p1})(s+\omega_{p2})} \quad (23)$$

Note that in the above equation, the zero is explicitly defined. The plots of the Compliance equalization transfer functions for short and long channels along with corresponding parameters are shown in Figure 26.

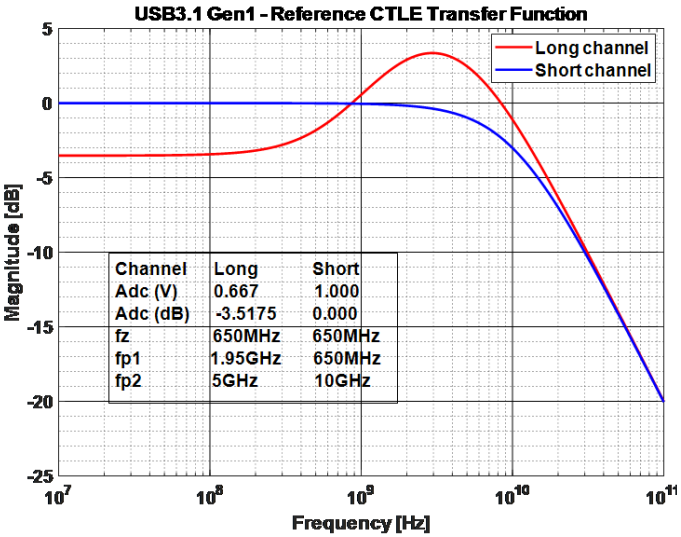


Figure 27 - USB3.1 Gen1 - Reference CTLE Transfer Function

From the above figure it can be seen that in the case of the long channel, the CTLE's transfer function has the characteristics of an active filter with positive gain around Nyquist frequency of 2.5GHz. As a result, the equalized pulse response after a channel with -9dB differential insertion loss, shown in red color in Figure 28, has larger amplitude than the unequalized pulse shown in green color. The pulse response with "short channel" CTLE is slightly attenuated due to the low pass filtering characteristic of the transfer function of a passive CTLE.

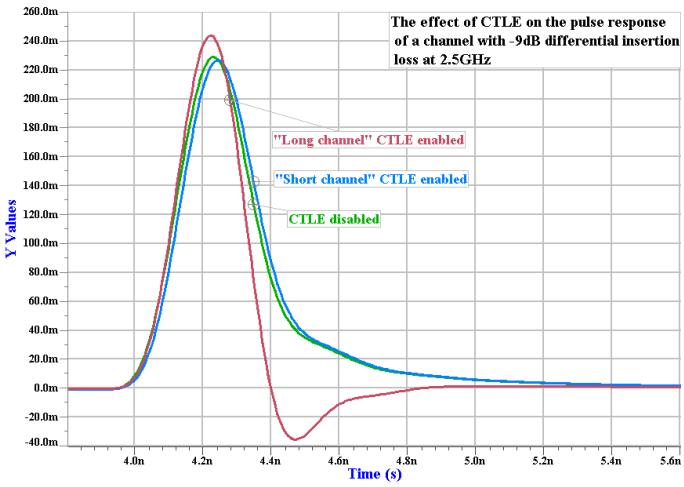


Figure 28 – The effect of USB Gen1 CTLEs on the pulse responses of a channel with -9dB differential insertion loss at Nyquist frequency of 2.5GHz

The corresponding eye density plots are depicted in Figure 29 below. As expected, the "long channel" CTLE generates an eye with the largest vertical and horizontal openings as in one side it amplifies the signal and on the other side it removes the ISI of the channel narrowing the pulse response of the channel.

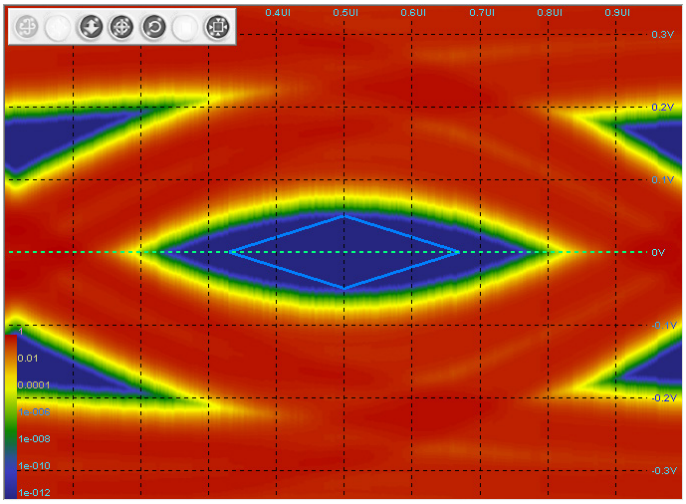


a) EH=114mV, EW=545UI

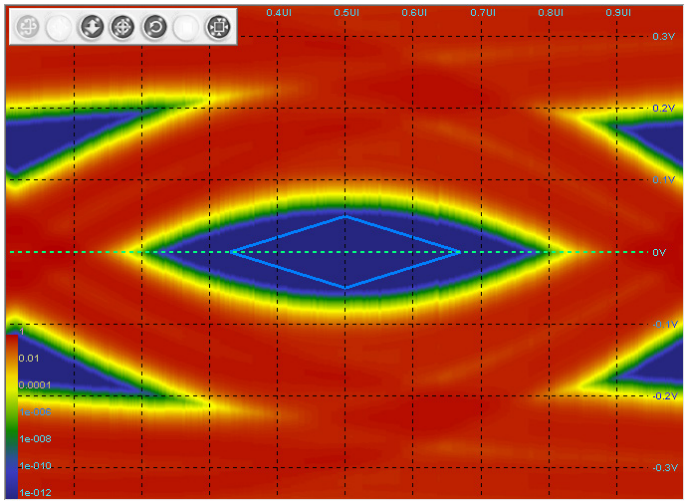


b) EH=108mV, EW=522UI

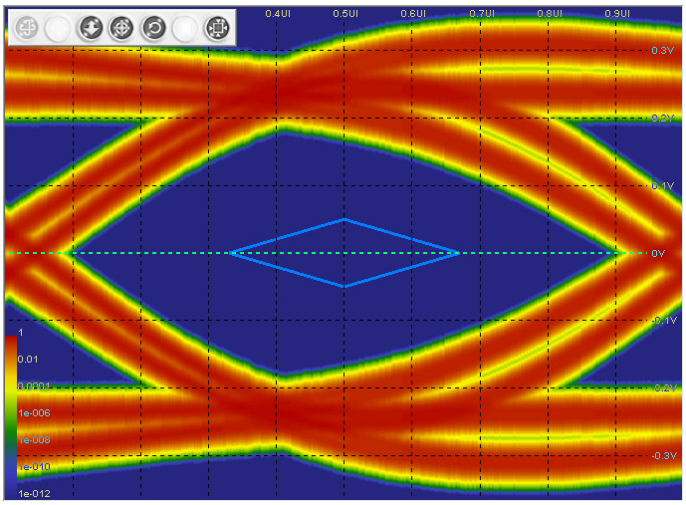


c) EH=345mV, EW=790UI

Figure 29 - The effect of USB Gen1 CTLEs on the eye density plots of a channel with -9dB differential insertion loss at Nyquist frequency of 2.5GHz: CTLE disabled a), "Short channel" CTLE enabled b), "Long channel" CTLE enabled c)

The second generation of the USB3.1 (10GT/s) has maintained similar 1st order CTLE architecture with transfer function defined by equation (24):

$$H_{(s)} = A_{ac}\omega_{p2} * \frac{s+\frac{A_{dc}}{A_{ac}}\omega_{p1}}{(s+\omega_{p1})(s+\omega_{p2})} \quad (24)$$

In the above equation, the zero frequency is indirectly defined by the location of pole 1 and the value of Adc, which can vary between 0 to -6dB, in 1dB steps:

$$\omega_z = A_{dc} * \omega_{p1} \quad (25)$$

The corresponding plot of the Compliance transfer function along with the associated parameters is depicted in Figure 30.

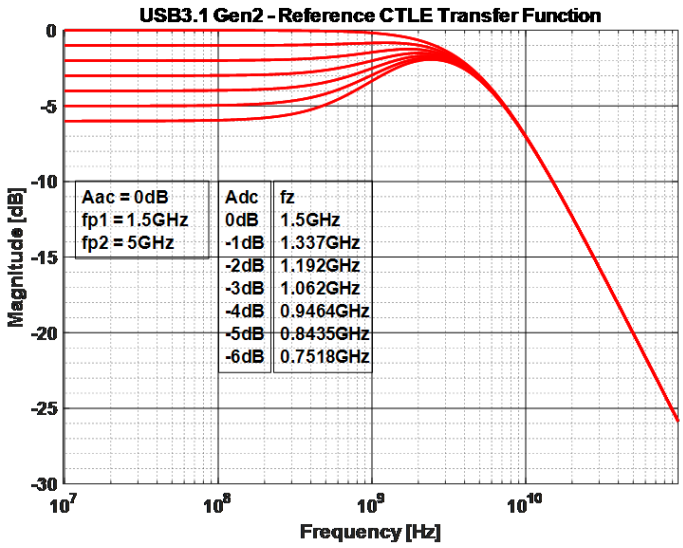


Figure 30 - USB3.1 Gen2 - Reference CTLE Transfer Function

The effect of the CTLE on the pulse responses of a channel with -6.5dB differential insertion loss at 5GHz, the Nyquist frequency of the USB3.1 Gen2 protocol, is shown in Figure 31. From this figure it can be noted that the unequalized pulse has the highest amplitude. The signal swing gets lower when the CTLE is turned on and continue to decrease proportionally with the Adc level: higher the attenuation at DC, lower the pulse amplitude. However, lower the pulse amplitude is, the narrower is its width, thus the associated ISI is reduced.

The corresponding eye density plots are depicted in Figure 32. As expected, the eye density plot with the highest vertical opening (EH) is the unequalized eye, but the eye with the largest horizontal opening (EW) is obtained when CTLE is turned on and Adc = -3dB. It is worth noting how the shape of the eye density plots is changing depending on the location of the zero on the transfer function and the Adc value.

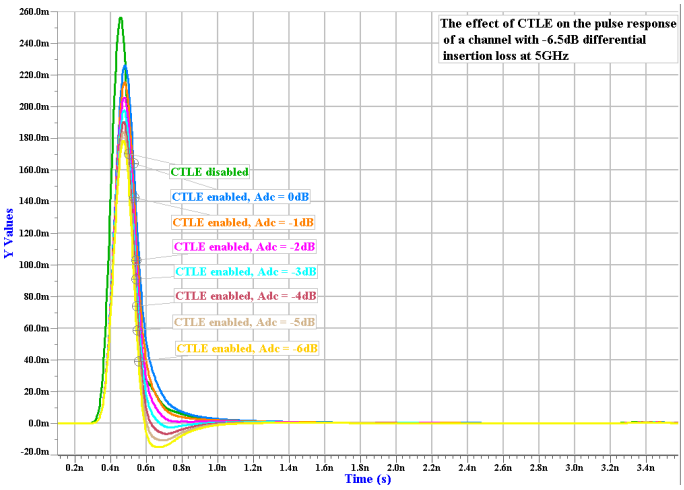


Figure 31 – The effect of USB Gen2 CTLEs on the pulse responses of a channel with -6.5dB differential insertion loss at Nyquist frequency of 5GHz

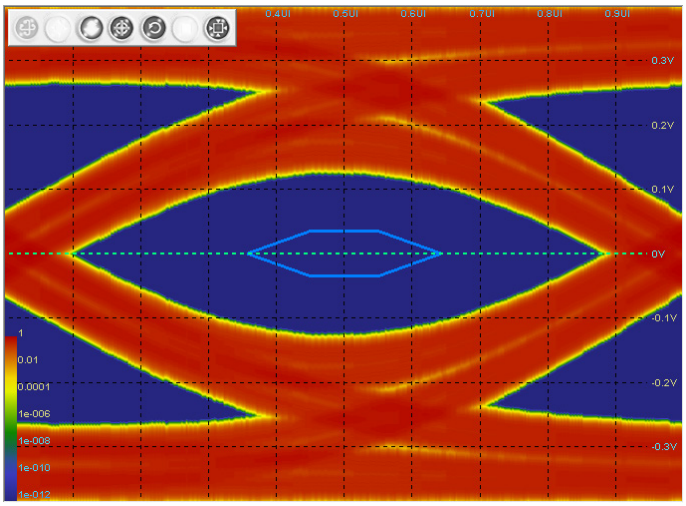


a) EH=248mV, EW=0.771UI

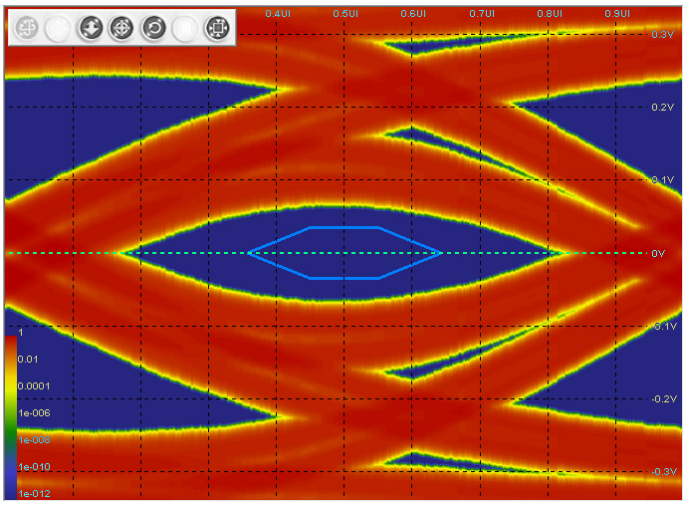


b) EH=128mV, EW=0.623UI

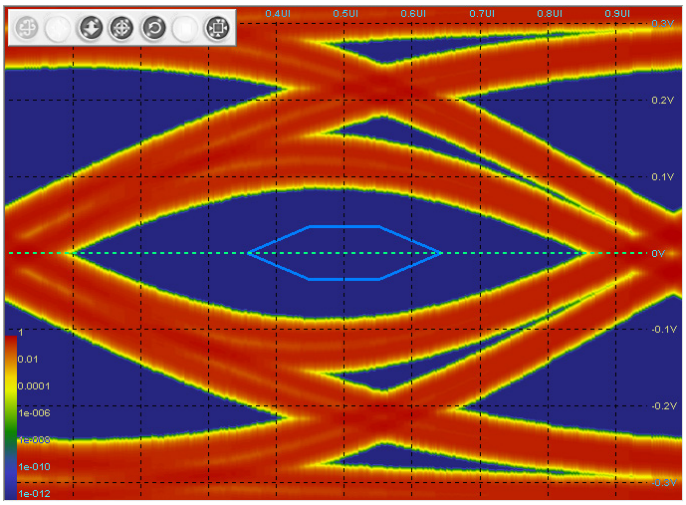


c) EH=167mV, EW=0.745UI

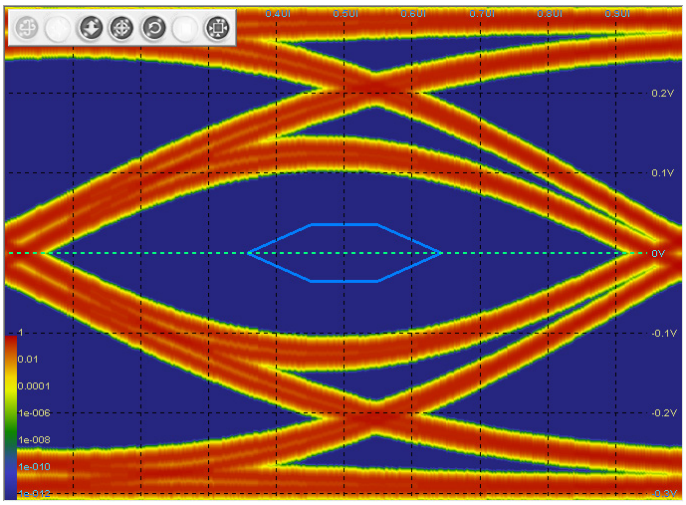


d) EH=200mV, EW=0.77UI

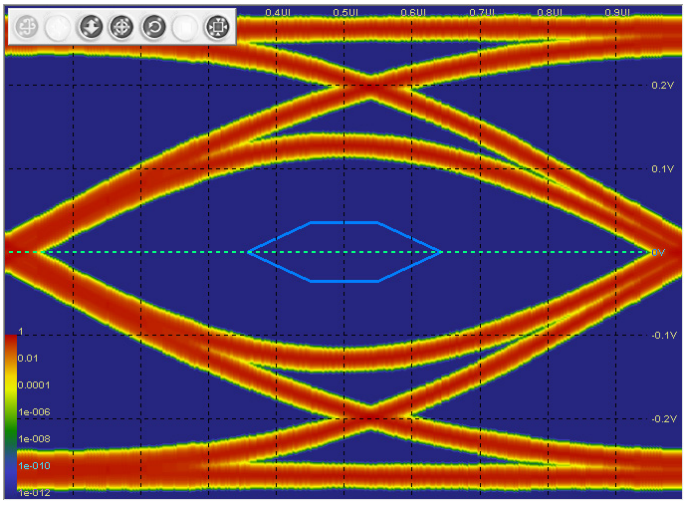


e) EH=224mV, EW=0.882UI

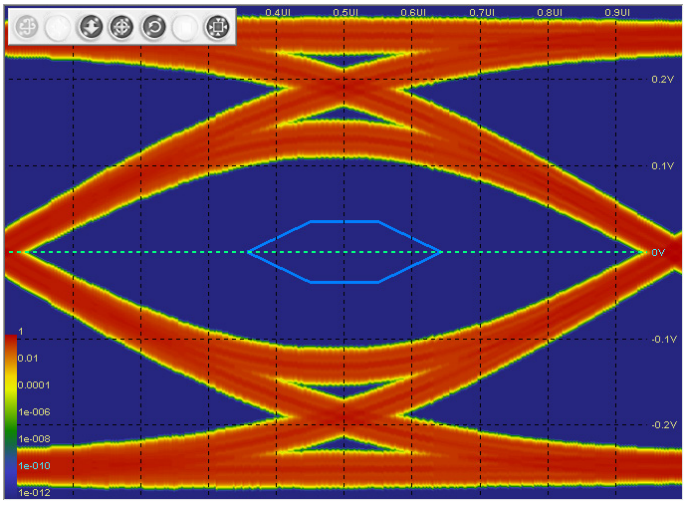


f) EH=216mV, EW=0.893UI

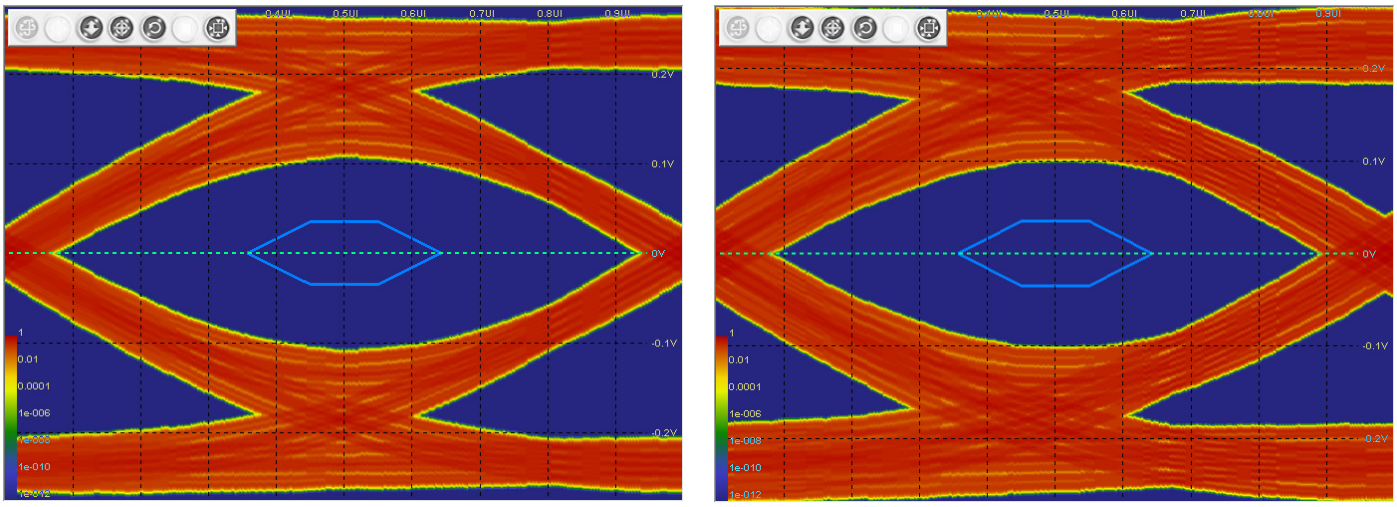

g) EH=212mV, EW=0.845UI

h) EH=197mV, EW=0.796UI

Figure 32 - The effect of USB Gen2 CTLEs on the eye density plots of a channel with -6.5dB differential insertion loss at Nyquist frequency of 2.5GHz: CTLE disabled a), CTLE enabled, Adc =0dB b), CTLE enabled, Adc =-1dB c), CTLE enabled, Adc =-2dB d), CTLE enabled, Adc =-3dB e), CTLE enabled, Adc =-4dB f), CTLE enabled, Adc =-5dB g), CTLE enabled, Adc =-6dB h)

Similarly to USB3.1, PCIe Gen3 (8GT/s) and Gen4 (16GT/s) are incorporating a 1st order CTLE with fixed low frequency (LF) and high frequency (HF) poles, and an adjustable Adc. The equation (26) that describes the transfer function is the same in the two cases; the Adc valid range and step is also the same, but the location of the HF poles is different (see Table 5).

$$H_{(s)} = \omega_{p2} * \frac{(s+\omega_{p1}*Adc)}{(s+\omega_{p1})(s+\omega_{p2})} \quad (26)$$

| | PCIe Gen 3 (8 GT/s) | PCIe Gen 4 (16 GT/s) |
|---|---|---|
| ωp1 (pole 1) | 2π*2 GHz | 2π*8 GHz |
| ωp2 (pole 2) | 2π*2 GHz | 2π*16 GHz |
| Adc | -6dB to -12dB, in steps of 1dB | -6dB to -12dB, in steps of 1dB |

Table 5 – Characteristics of the reference CTLEs for PCIe Gen3 and Gen4

The two loss curves for 8GT/s and 16GT/s behavioral CTLEs are shown in Figure 33 and Figure 34 respectively.

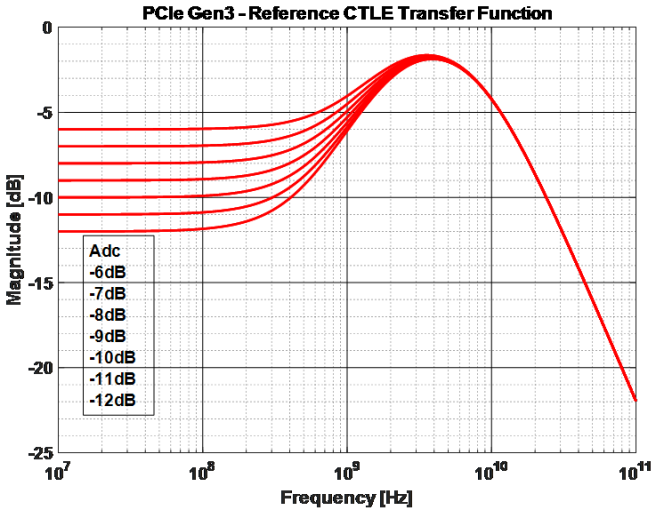


Figure 33 – PCIe Gen3 - Reference CTLE Transfer Function

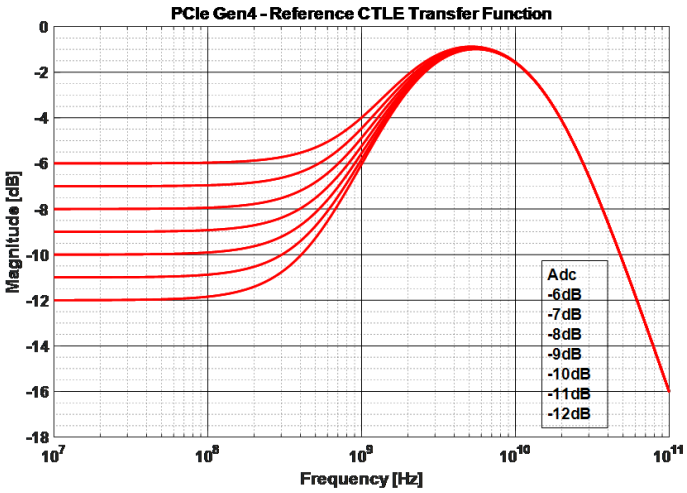


Figure 34 - PCIe Gen4 - Reference CTLE Transfer Function

PCIe Gen5 (32GT/s) has increased the complexity of the CTLE to a 2nd order transfer function as depicted in equation (27):

$$H_{(s)} = \frac{\omega_{p1}*\omega_{p3}*\omega_{p4}}{\omega_{z1}} * \frac{(s+\omega_{z1})(s+\omega_{p2}*Adc)}{(s+\omega_{p1})(s+\omega_{p2})(s+\omega_{p3})(s+\omega_{p4})} \quad (27)$$

As for previous speeds, the location of the poles is fixed and the DC gain is adjustable. The Adc valid range is from -5dB to -15dB and the step size is 1dB as depicted in Figure 35 and Table 6.

The IEEE802.3bj specification defines the transfer function of the CTLE for the supported PMDs, using a reduced form of the equation (11):

$$H_{ctf}(f) = \frac{10^{\frac{gDC}{20}} + j\frac{f}{f_z}}{(1+j\frac{f}{f_{p1}})(1+j\frac{f}{f_{p2}})}, \quad (28)$$

where gDC is the DC gain.

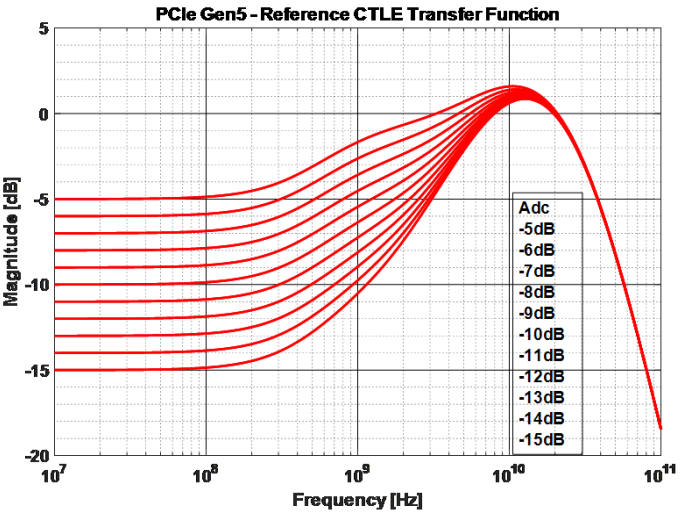


Figure 35- PCIe Gen5 - Reference CTLE Transfer Function

| | PCIe Gen 5 (32 GT/s) |
|---|---|
| Fz1 (zero 1) | 450 MHz |
| Fz2 (zero 2) | Mag(DC gain)*Fp2 |
| Fp1 (pole 1) | 1.65*Fz1 |
| Fp2 (pole 2) | 9.5 GHz |
| Fp3 (pole 3) | 28 GHz |
| Fp4 (pole 4) | 28 GHz |
| Adc | -5 to -15dB, in steps of 1dB |

Table 6- Characteristics of the reference CTLEs for PCIe Gen5

The 100GBASE-KR4 and 100GBASE-KP4 operating modes include CTLEs with transfer functions defined by the above equation. This equation is incorporated in the Channel Operating Margin (COM) computational algorithm. The DC gain minimum and maximum values as well as the step size are defined in tables that are specific for each operating mode. A typical range is 0 to -12dB in 1dB step size. The location of the poles and zeros is provided relative to the signaling rate $f_b$ value (e.g. $f_z = f_{p1} = \frac{f_b}{4}$ and $f_{p2} = f_b$). The transfer function of the100GBASE-KR4 operating mode is plotted in Figure 36. The combined responses of this CTLE and that of a PCB is shown in Figure 37, for each valid value of the DC gain.

The IEEE802.3bs specification is adding to equation (28) one low frequency (LF) pole and zero and a second DC gain ($g_{DC2}$). The new transfer function of the CTLE is depicted in equation (29) [3]:

$$H_{ctf}(f) = \frac{(10^{\frac{gDC}{20}} + j\frac{f}{f_z})(10^{\frac{gDC2}{20}} + j\frac{f}{f_{LF}})}{(1+j\frac{f}{f_{p1}})(1+j\frac{f}{f_{p2}})(1+j\frac{f}{f_{LF}})} \quad (29)$$

Ethernet 200GAUI-4, 400GAUI-8 chip-to-chip, 50GBASE-KR, 100GBASE-KR2, 200GBASE-KR4, 50GBASE-CR, 100GBASE-CR2, 200GBASE-CR4, and OIF CEI-56-MR-PAM4, and CEI-56G-LR-PAM4 all incorporate CTLEs with such transfer function. In some cases, the clause does not provide explicit values for the second DC gain and the low-frequency pole/zero, which implies that $g_{DC2}$ is set to 0 and $f_{LF}$ to 1. For exemplification, the loss plots for OIF CEI-56-LR-PAM4 protocol are shown in Figure 38. It should be noted that compared to the previously discussed protocols, this one has the finest granularity and range for the transfer functions, which means that the combined responses of the channel and CTLE can improve the behavior of a large variety of channels. Moreover, equation (29) was implemented in the latest revisions of the COM algorithm.

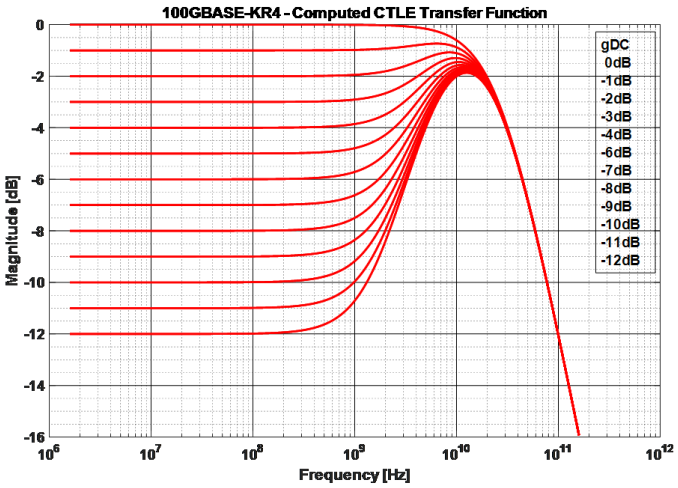


Figure 36 – IEEE802.3bj 100GBASE-KR4 Computed CTLE Transfer Function

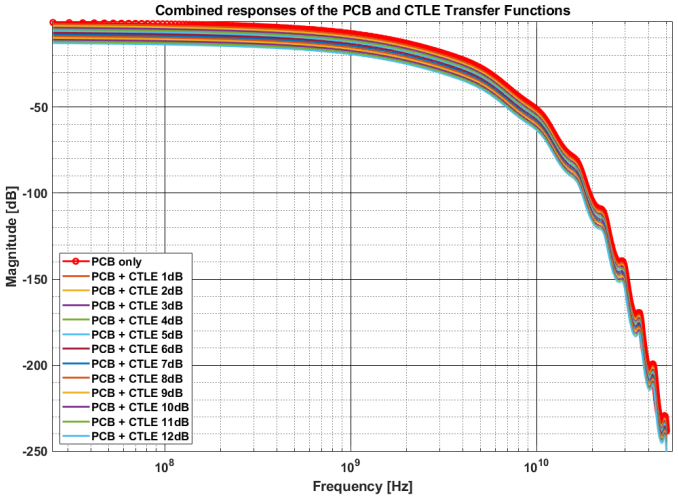


Figure 37 - IEEE802.3bj 100GBASE-KR4 combined response of PCB and CTLE Transfer Function

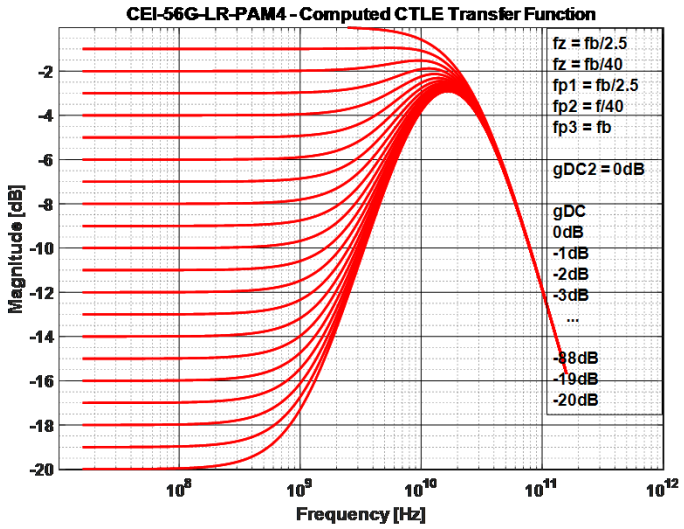


a) $g_{DC2}$ = 0dB, $g_{DC}$ = variable

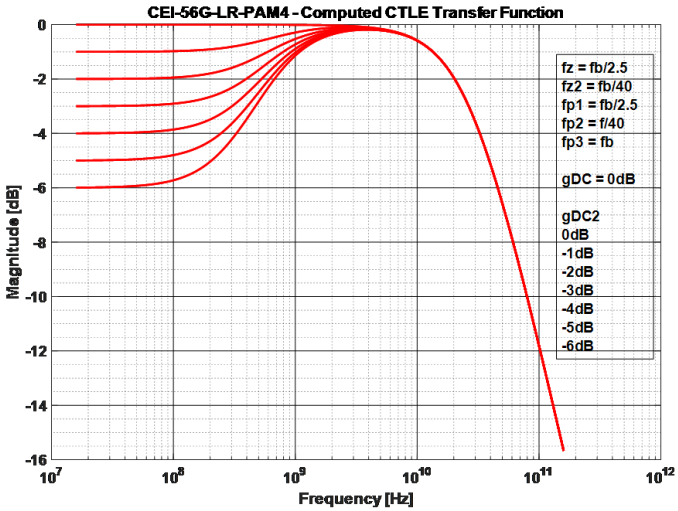


b) $g_{DC}$ = 0dB, $g_{DC2}$ = variable

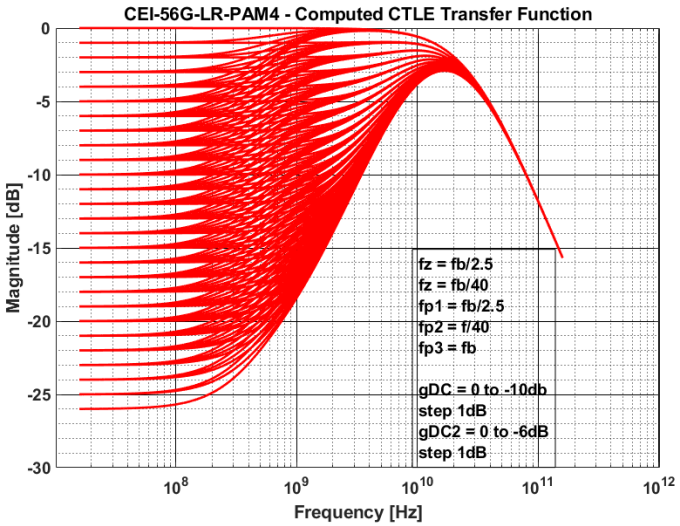


c) $g_{DC}$ = variable , $g_{DC2}$ = variable

Figure 38 – OIF CEI-56-LR-PAM4 Computed CTLE Transfer Function

The JESD204C has introduced three different CTLE implementations, one for each of the three types of reference receivers (class C-**S**hort, C-**M**edium, and C-**R**eflective). All of them are described by a 1st order CTLE model with adjustable DC gain and variable location of the zero and poles. The transfer functions are characterized by the general equation (30) [14] shown below:

$$H_{ctf}(f, setting) = g_{DC}(setting) * \frac{\prod_{i=1}^{np} -f_{pi}}{\prod_{i=1}^{nz} -fzi} * \frac{\prod_{i=1}^{nz}(2*\pi*(j*f-f_{zi}))}{\prod_{i=1}^{np}(2*\pi*(j*f-f_{pi}))} \quad (30)$$

where each setting represents a collection of predefined values for the DC gain, and locations for zero and poles [14]:

$$g_{DC}(setting) = 10^{\frac{g_{DC,dB}}{20}} [-] \quad (31)$$

$$Z = Z(settings) = (f_{z1}, f_{z2}, \dots, f_{zn}) \text{ [Hz]} \quad (32)$$

$$P = P(settings) = (f_{p1}, f_{p2}, \dots, f_{pn}) \text{ [Hz]} \quad (33)$$

The above equations are implemented in the JCOM algorithm and the associated loss plots are depicted in Figure 39.

As for the FFE, an important task for the simulation tool is to find the optimized CTLE settings. We will discuss this topic later in this document.

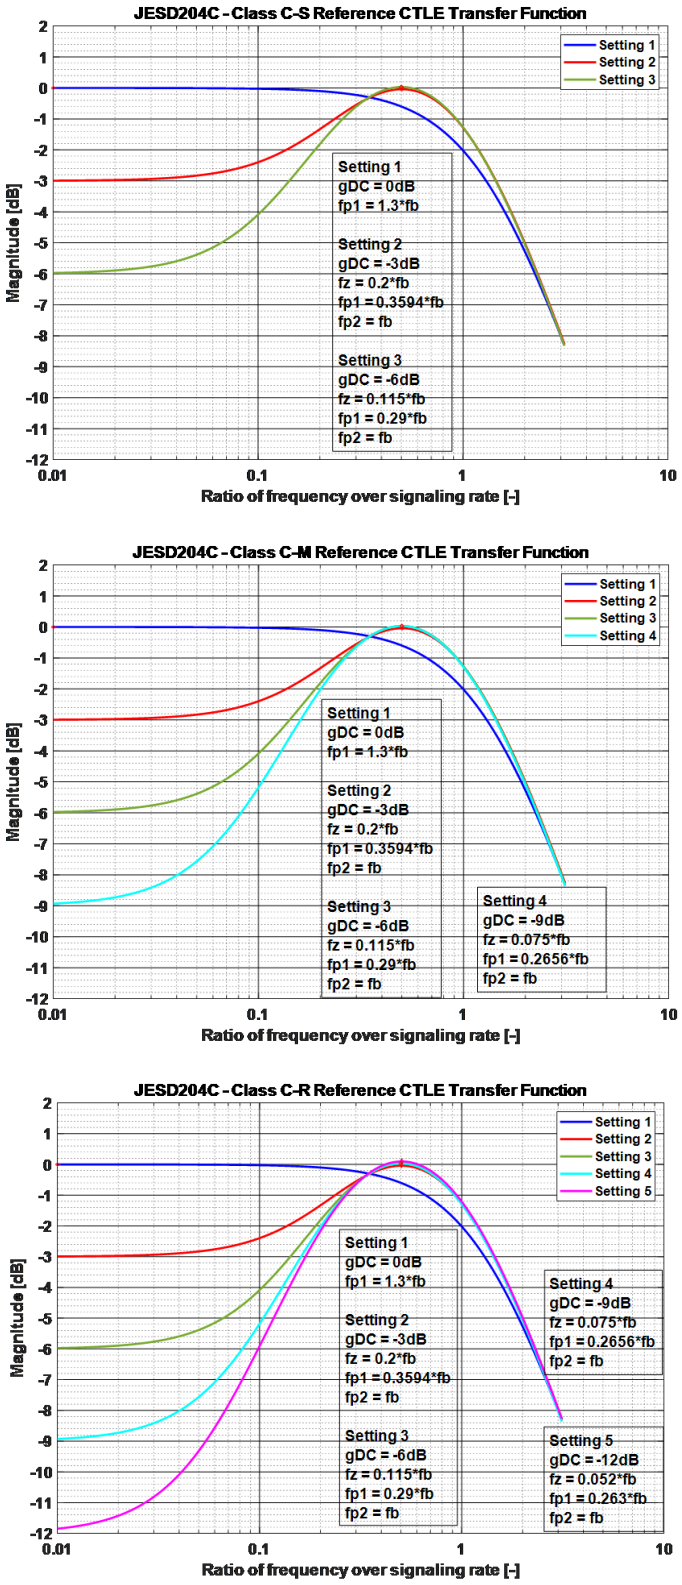


Figure 39 – JESD204C Computed CTLE Transfer Functions

# Decision Feedback Equalizer (DFE)

SerDes interfaces running at 8Gbps and above usually include a Decision Feedback Equalizer (DFE) in their equalization architectures. The DFE is a non-linear equalizer, typically placed after CTLE, and uses the previously detected symbols to estimate and cancel the postcursor's ISI from the input bit stream. As shown in Figure 40, at a high-level, the architecture of a behavioral DFE is very similar to that of a transmitter FIR. However, in this case, the filter is placed in a feedback loop together with a decision function (slicer). Similarly to transmitter FIR, the DFE requires clock to operate.

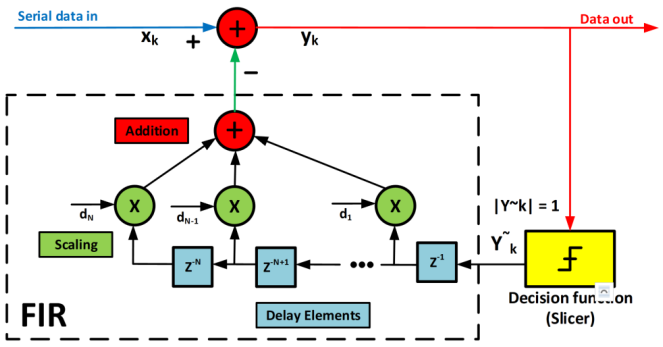


Figure 40 - Block diagram of a generic DFE with N taps

The decision function determines if the incoming signal is a "0" or a "1". The output of the slicer is fed into the FIR filter, where it gets multiplied by the filter's tap coefficients. The sum of the weighted and delayed signals is then subtracted from the input serial data. The signal at the output of the DFE can be computed using the following equation [14]:

$$y_k = x_k - \sum_{i=1}^{N} d_i * sgn(y_{k-i}) \quad (34)$$

Where $x_k$ is the DFE differential input voltage, $y_k$ is the DFE summer differential output voltage, $d_i$ are the feedback coefficients, k is the sample index in UI, $y_{\tilde{k}}$ is the decision function output voltage (with $|y_{\tilde{k}}|=1$) and N is the number of DFE taps.

The tap coefficients are calculated based on the ISI introduced by the lossy channel and based on the previous decision, therefore a DFE can only cancel postcursor ISI. If the tap coefficients are properly chosen, the filter can remove as many taps of ISI as the DFE has. However, if any of the previous decisions is not correct, it is fed back to the input of the FIR and can generate error propagation. The number of DFE taps and the values of their coefficients are affecting the length of the error burst.

Many hardware implementations are being described in the literature. Each of them has its own merits and is more suitable for a particular application. However from an educational perspective two simple architectures are of interest [9]. The first one is referred to as "direct" feedback architecture and is shown in Figure 41.

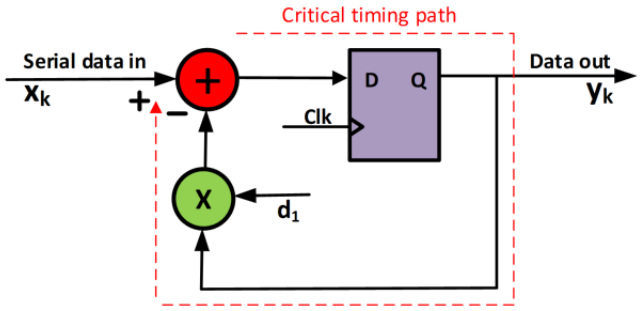


Figure 41 - Direct DFE architecture

The main element in this implementation is a latch with sample threshold controlled by an analog block. The threshold can move up or down depending on the value of the previously received bit. The decision process is illustrated in Figure 42 using a simulated waveform with a binary sequence of "001101111100" that is input to a lossy channel with a large amount of ISI.

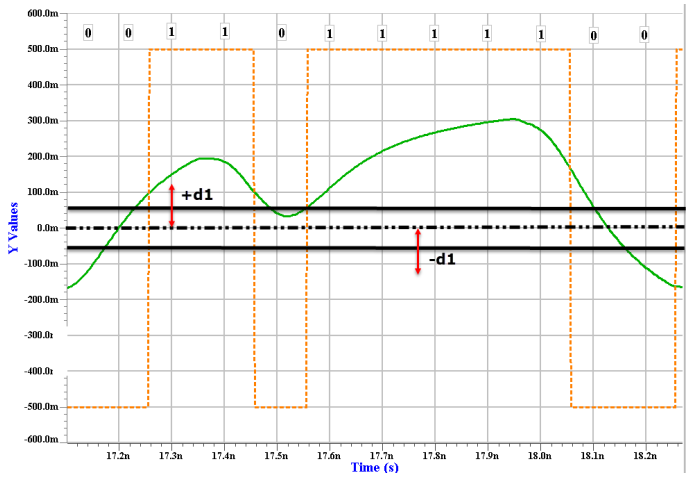


Figure 42 - Adaptive DFE threshold

The dotted orange line is the digital sequence at the output of the transmitter and the green waveform is the unequalized waveform at the input of the receiver. The skew between the two waveforms was intentionally removed and the waveforms are now edge aligned. A traditional latch with a fixed sampling threshold at zero, would interpret the isolated "0" in the middle of the bit stream as "1", since the received waveform does not cross the zero threshold. The received signal would be, in this case the sequence "001111111100".

As the bit preceding the erroneously received bit was "1", the analog function of the DFE moves the latch sample threshold up by the value of d1, where d1 is the value of the feedback coefficient. If d1 is properly chosen, the received signal will cross the new threshold and the isolated bit will be properly interpreted as "0". Similarly, if the previously received bit was "0", the latch sample threshold would be decreased by d1.

Although conceptually simple, this architecture suffers from the very narrow constraints imposed on the critical timing path. The feedback loop latency must be less than the bit period, which is a tough constraint for high data rate SerDes protocols.

The second architecture is called "speculative" (or "unrolled") DFE architecture and is illustrated in Figure 43.

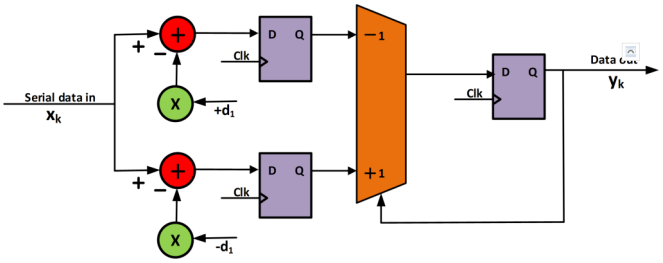


Figure 43 - Speculative DFE architecture

This architecture eliminates the critical timing constraints by duplicating the signal path. The path at the top uses the positive value of the tap coefficient, while the one at the bottom its negative value. Then, a multiplexor selects the correct output depending on the value of the previously received bit.

Analog implementations of a DFE block are also possible. The performance of two continuous-time DFE (CT-DFE) structures (1st and 2nd order) is discussed in [18] and shown in Figure 44.

In order for a CT-DFE to be efficient, its poles and gain values should be carefully chosen to emulate and remove the channel ISI. Higher the order CT-DFEs can improve more complex, bumpy channels.

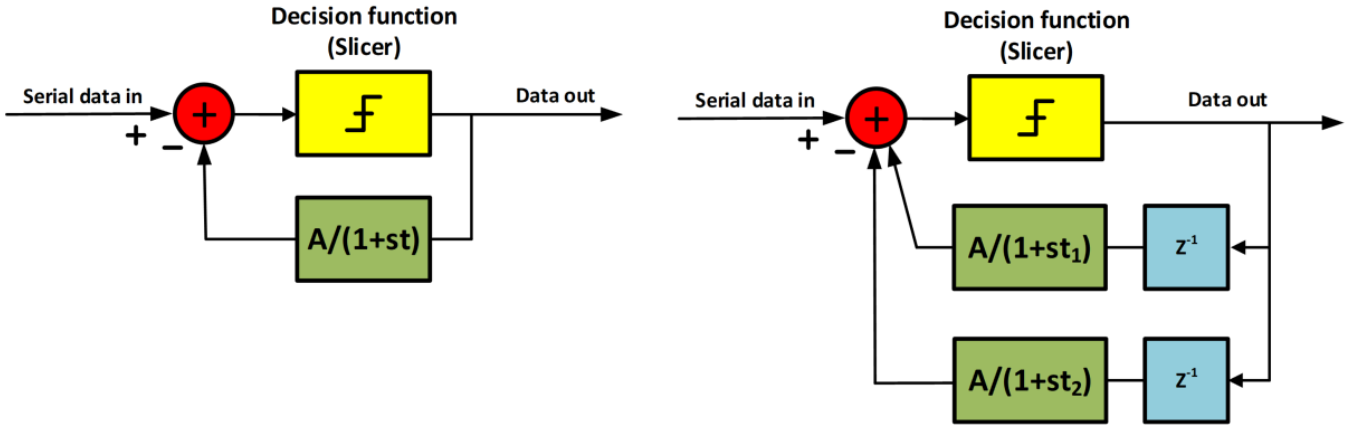


Figure 44 - Continuous-time DFE: 1st order a) and 2nd order b)

# Examples of Decision Feedback Equalizer architectures

Although there are many possible hardware implementations, the SerDes protocol specifications, do not go into these type of details. They provide only the high level requirements of a reference behavioral DFE and leave the implementation up to the silicon vendors. Most of the specifications define a minimum number of taps and a set of associated constraints (limits).

DFE architectures with 1 to 16 taps are frequent and the tap length typically increases with the data rate of the interface. The tap values are defined either as absolute values (usually mV) or relative to the main cursor value (see PCIe Gen 5). The tap limits can be defined in absolute value (mV) or normalized to the value of the equalized pulse response taken at the sampling location (the A_s parameter in the case of the COM algorithm). The effect of DFE is often modelled by subtracting the tap coefficient values from the input sampled signal and the error propagation is typically ignored.

The block diagram of a 1 tap DFE is shown in Figure 45. PCIe Gen3 and USB3.1 Gen2 protocols define such architecture in their specifications.

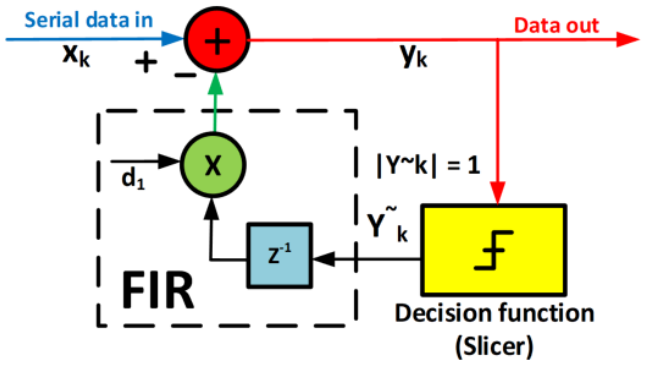


Figure 45 - Block diagram of a 1-tap DFE

The signal at the output of the DFE can be calculated using equation (35) [2], which is a particular case (N =1) of equation (34):

$$y_k = x_k - d_1 sgn(y\sim_{k-1}) \quad (35)$$

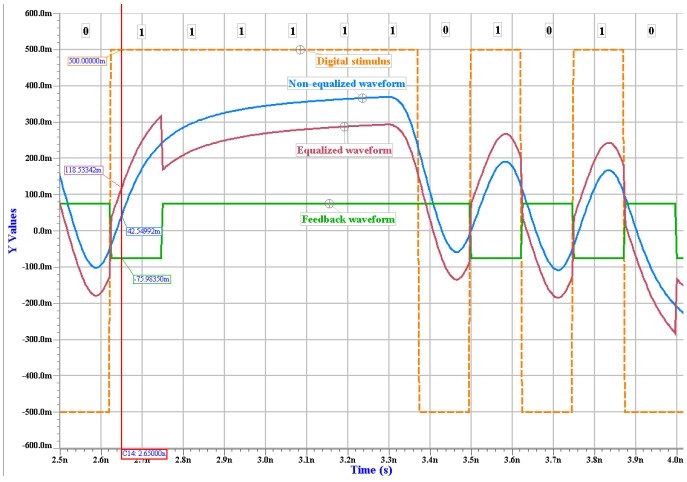


Figure 46 – 011111101010 binary pattern with 1 tap DFE applied (d1 =76mV)

The simulated waveforms of a data stream of 011111101010 are illustrated in Figure 46.

The data rate was chosen to be 8Gbps (PCIe Gen3), consequently the corresponding bit rate is 125ps. The orange dotted waveform represents the digital stimulus at the input of the transmitter. The blue colored waveform is the signal at the receiver, without any equalization applied to it. The green curve is the output of the FIR and represents the feedback signal that is summed up with the incoming signal to generate the equalized signal that is shown in red color. The skew between the signals was intentionally removed in this plot for simplification.

In this particular example d1 = 76 mV. As a result the feedback signal has two values: +76mV and -76mV. This is the signal at the output of the slicer (+/-1V), delayed by one UI and multiplied by d1. Since the first bits of the data stream are 0 and 1 respectively, the transitioning bit has lower margin due to the channel's ISI. As a result, its logic value of 1 might be wrongly interpreted. The DFE compensates for this impairment, by shifting the incoming signal up by the amplitude of the feedback signal, thus increasing the margins. The next five bits are non-transitioning bits with the same logic value of 1, the same as that of the second bit of the data stream. The slicer will toggle its output to -1V and the feedback signal will now subtract from the incoming signal, shifting it down (by 76mV). The output of the slicer will toggle again when it detects the next transition (from logic 1 to 0) and the process will continue until the last bit is received.

Similarly the block diagram of a 2 tap DFE, similar to the one used by PCIe Gen 4, is shown in Figure 47 and the simulated waveforms are illustrated in Figure 48.

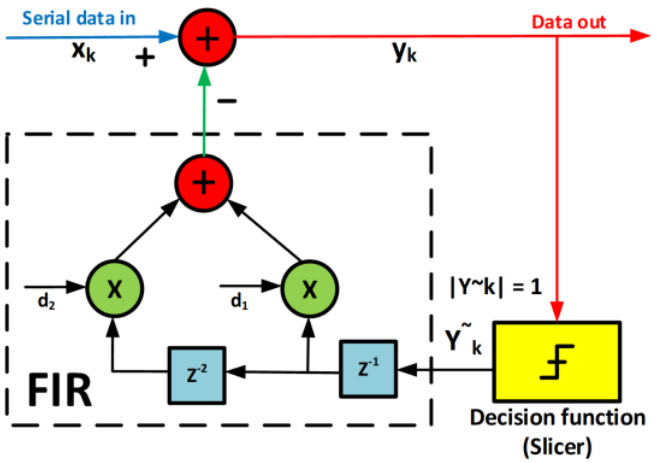


Figure 47 - Block diagram of a 2-tap DFE

The signal flow and the processing are the same as in the previous case, but now the equalized signal is calculated using the equation (36):

$$y_k = x_k - d_1 sgn(y\sim_{k-1}) - d_2 sgn(y\sim_{k-2}) \quad (36)$$

The feedback signal can have, in this case, multiple levels, corresponding to the possible combinations of the output of the slicer and the values of the tap coefficients d1 and d2.

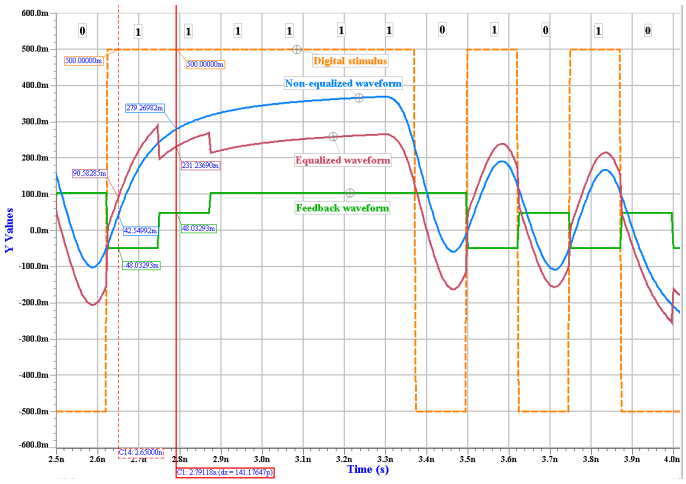


Figure 48 - 011111101010 binary pattern with 2 tap DFE applied (d1 =48mV, d2 =48mV)

In the case of a 15 tap DFE, the signal can be calculated using equation (37):

$$y_k = x_k - d_1 sgn(y\sim_{k-1}) \ldots - d_{14} sgn(y\sim_{k-14}) - d_{15} sgn(y\sim_{k-15}) \quad (37)$$

The simulated waveforms (at the same data rate) are shown in Figure 49.

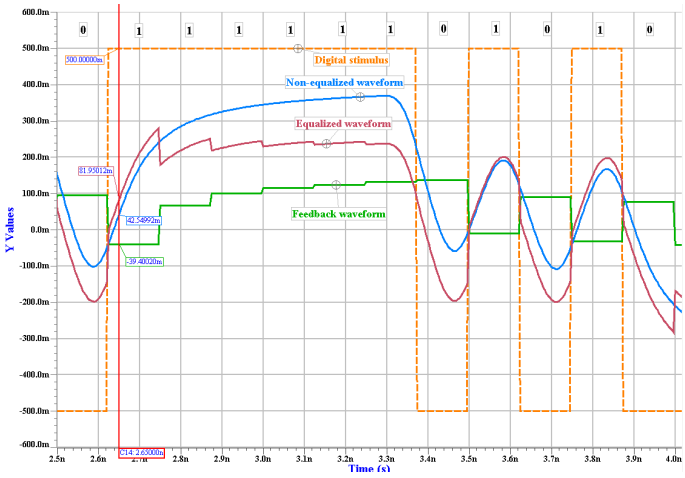


Figure 49 - 011111101010 binary pattern with 15 tap DFE applied

Since the feedback signal has many possible levels, the step size is finer and the equalized signal can better reproduce the shape of a digital waveform, thus increasing the margins and reducing the chance of an error to occur.

The effect of the number of DFE taps on the simulated eye density plots is shown in Figure 50.

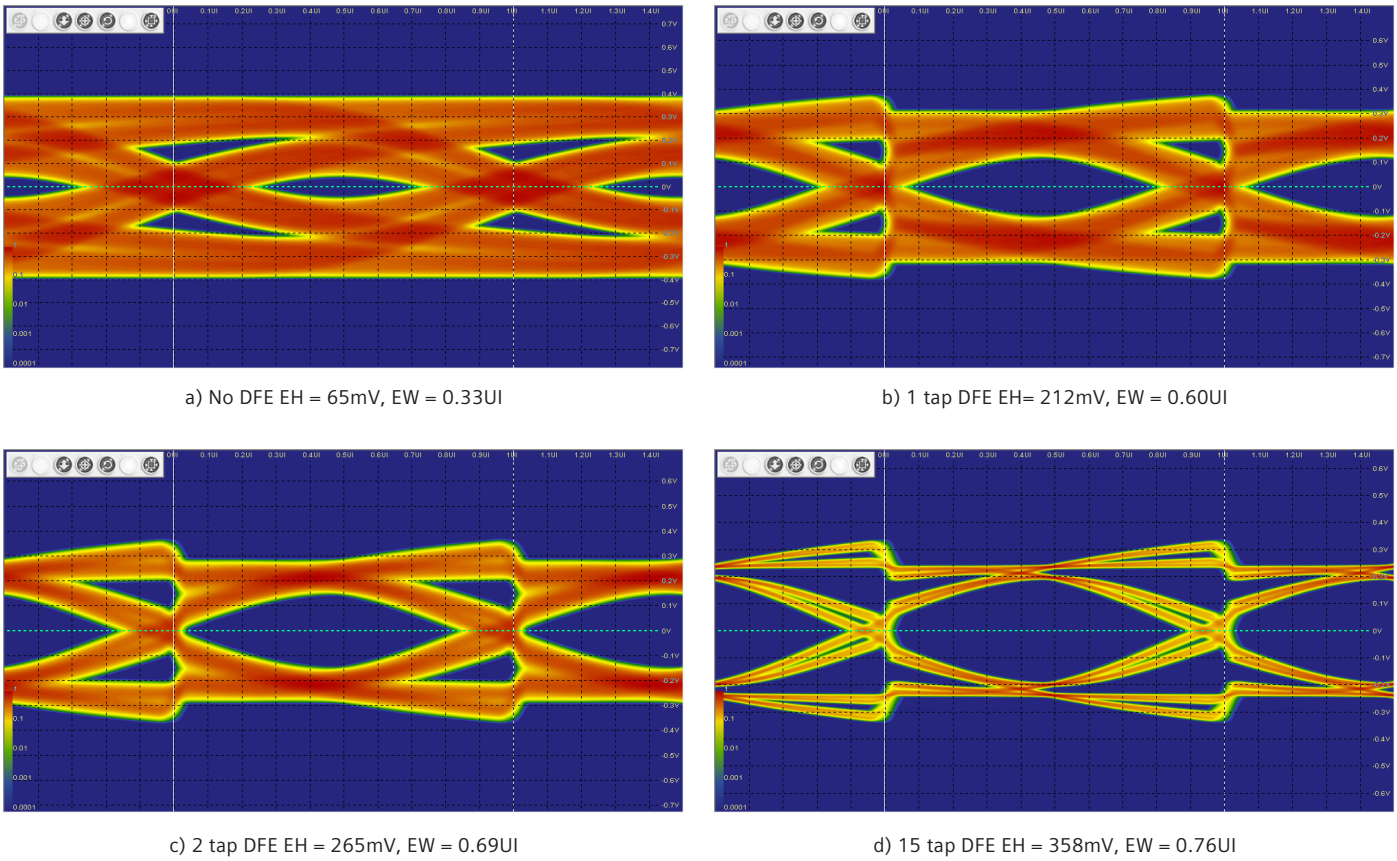

a) No DFE EH = 65mV, EW = 0.33UI

b) 1 tap DFE EH= 212mV, EW = 0.60UI

c) 2 tap DFE EH = 265mV, EW = 0.69UI

d) 15 tap DFE EH = 358mV, EW = 0.76UI

Figure 50 – Eye density plots at the output of a behavior DFE and after a channel with -13.5dB differential insertion loss at Nyquist frequency (4GHz): the DFE turned off a), 1 tap DFE b), 2 tap DFE c) and 15 tap DFE d)

We used here the same simulation deck and channel as the one from the experiments shown in Figure 10. The data rate was 8Gbps in each of the simulated cases. The only variable in those simulations was the number of the Rx DFE taps that were turned on. As it can be seen from the four eye density plots the higher the number of DFE taps is, the lower the signal to noise ratio is. However, in practical applications the maximum number of taps that can be implemented in silicon is limited by a number of factors such as timing requirements, power consumption, silicon area and cost.

One solution to this limitation is the Sliding Tap DFE (ST-DFE) [15]. As opposed to traditional DFEs, where the location of the taps is fixed, an ST-DFE has a number of taps with fixed location and a number of taps with variable location. The location of the fixed taps is from the first postcursor to the Nth postcursor, where N is the number of taps with fixed location. The location of the variable taps can be anywhere from $N+1^{th}$ to $M^{th}$ postcursor locations, where M depends on the technology (64 in [15]). The adaptation engine of a ST-DFE computes, besides the tap strengths, the optimized tap location for the variable taps.

As in the case of FFE and CTLE, it is important to understand how the DFE shapes the pulse responses of the channels in order to remove the effect of ISI. For example we provide in Figure 51 to Figure 54 the plots of non-equalized and equalized pulse responses of the same channel as they were shaped by the behavioral reference DFE of several protocols with different tap lengths and constraints.

As a general rule, the first DFE tap cancels the first post-cursor ISI by subtracting the positive value of the feedback signal, thus pulling the waveform down to bring the equalized sample near zero. The remaining taps can move the waveform either up or down, to remove the residual ISI. This situation would occur in the case of reflective channels with ripples on the tail of the pulse response. Since in those examples we used a long, lossy channel with smooth exponentially decaying ISI long-tail and no reflections at all, all the taps are subtracting from the waveforms. The shift on the voltage level occurs at locations between two consecutive sample points.

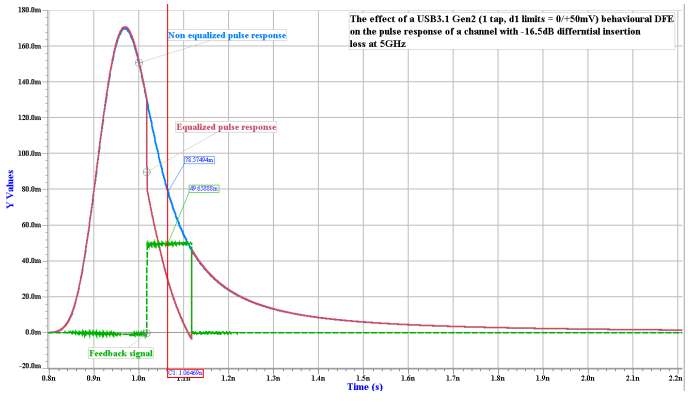


Figure 51 - The effect of a USB3.1 Gen2 (1 tap, d1 limits = 0/+50mV) behavioral DFE on the pulse response of a channel with -16.5dB differential insertion loss at Nyquist frequency of 5GHz

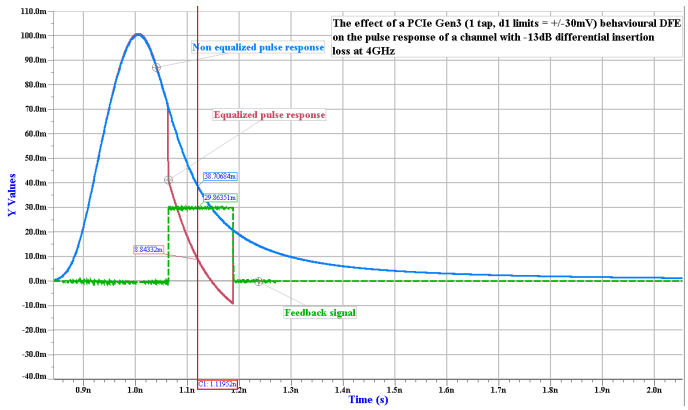


Figure 52 - The effect of a PCIe Gen3 (1 tap, d1 limits = +/-30mV) behavioral DFE on the pulse response of a channel with -13dB differential insertion loss at Nyquist frequency of 4GHz

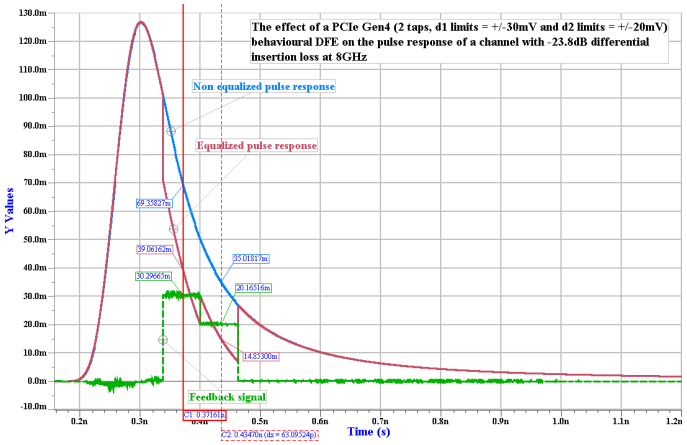


Figure 53 - The effect of a PCIe Gen4 (2 taps, d1 limits = +/-30mV and d2 limits = +/-20mV) behavioral DFE on the pulse response of a channel with -13dB differential insertion loss at Nyquist frequency of 4GHz

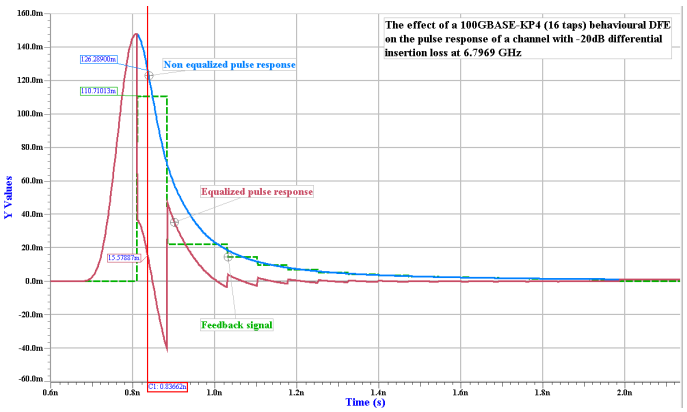


Figure 54 - The effect of a 100GBASE-KP4 (16 taps, max normalized d1 limit = 1 and max d2 to d16 limit = 0.2) behavioral DFE on the pulse response of a channel with -20dB differential insertion loss at Nyquist frequency of 6.7969 GHz

# Receiver Feed-forward Equalizer (FFE)

The most recent specifications, such as CEI-112G-LR-PAM4 and CEI-112G-VSR-PAM4, are introducing in their equalization architectures an FIR based Rx FFE [16], [17]. Actual hardware implementations can use digital or analog FIR filter. The conceptual block level diagram of a digital Rx FFE, shown in Figure 55, is similar to that of the Tx FIR equalization block except the addition of the analog to digital converter (ADC) that is needed at the input of the filter to deal with the analog nature of the signal. As for Tx FIR, the Rx FFE needs to be synchronized with the system's clock. Consequently for proper operation, either the clock needs to be recovered at Rx, or the operating frequency must be known.

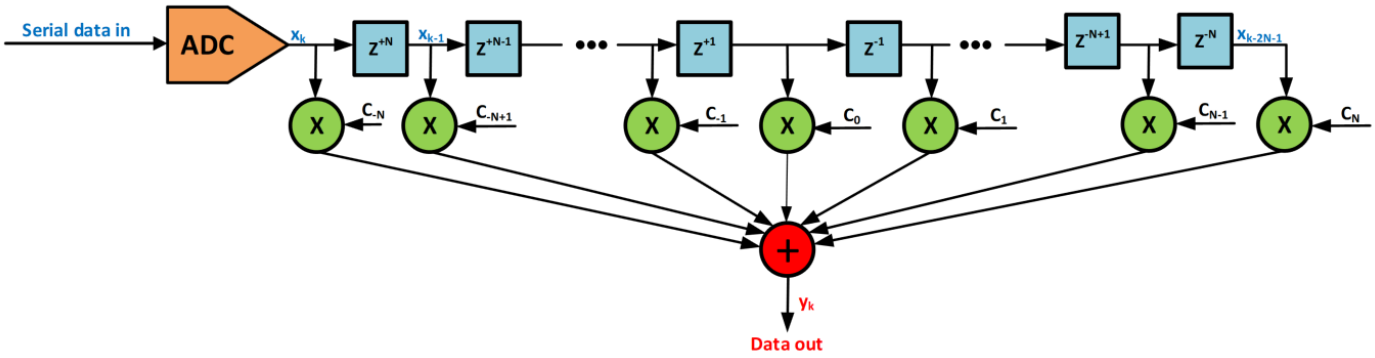


Figure 55- Block diagram of a receiver digital FFE equalizer

From a functional perspective the Tx FIR and Rx FFE are producing similar results as long as they are identically configured. The non-equalized and equalized pulse responses from COM simulations of a CEI-112G-LR-PAM4 reflective channel with differential IL of -34dB at 28GHz are shown in Figure 56.

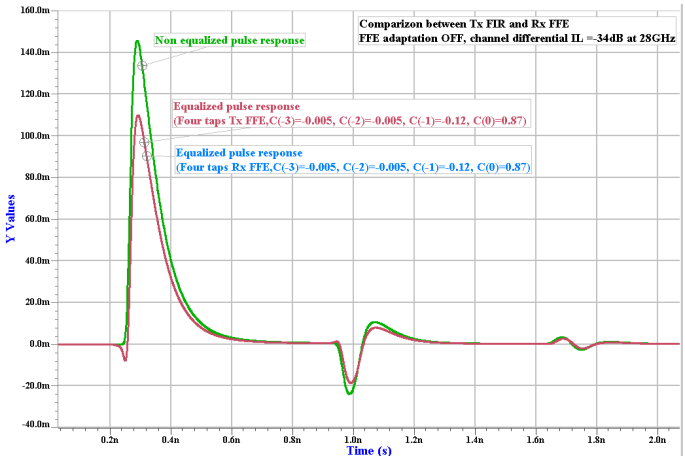


Figure 56 - Comparison between Tx FIR and Rx FFE: FFE adaption OFF, channel differential IL =-34dB at 28GHz

As can be seen from the above screenshot, the two equalized pulse responses perfectly overlap. In this experiment the adaptation of the tap coefficients was turned off in both cases and the tap coefficients were forced to fixed values. However, as we explained earlier in this document, the signal at the output of Tx has to meet some minimal magnitude (for compliant transmitters), which imposes additional constraints on the tap coefficients (|c(-2)|+ |c(-1)|+ |c(0)|+ |c(1)|)*2*T_Vf ≤ 1.2V for CEI-112G-LR-PAM4, where T_Vf = 0.4V to 0.6V is the steady-state voltage). The Rx FFE does not have such requirements; consequently the adaptation algorithms might generate different results as shown in Figure 57.

Compared to DFE, the Rx FFE has the advantage that it can compensate for both pre- and post-cursor ISI. When used together, the FFE is effective in precursor cancelation, helps reduce the number of required DFE taps and can eliminate the need for backchannel in the adaptation process.

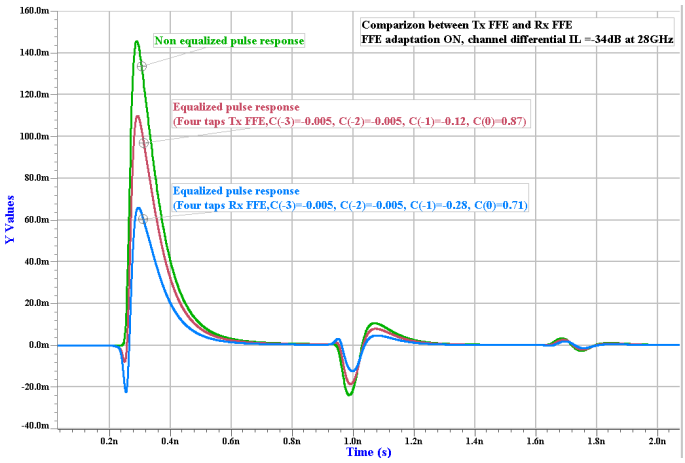


Figure 57 - Comparison between Tx FIR and Rx FFE: FFE adaption ON, channel differential IL =-34dB at 28GHz

# Performance of various equalization types

Higher data rates can be obtained using a combination of various equalization methods, modulation schemes and forward error correction (FEC) processing techniques. From an equalization perspective, combining multiple equalization circuits on the same link allows for achieving better results than relying on a single equalizer. Figure 58 shows overlapped pulse responses of a channel, with -33dB differential insertion loss at the Nyquist frequency, as they were reshaped by various equalization types and combinations of them. From this screenshot it can be seen that the combination of Tx FFE, Rx CTLE and Rx DFE is capable of restoring the transmitted pulse and almost completely removing the channel ISI.

Each of the previously described equalization techniques has its own merits and flaws. Table 7 contains a summary of the most important characteristics of the most commonly used equalization types.

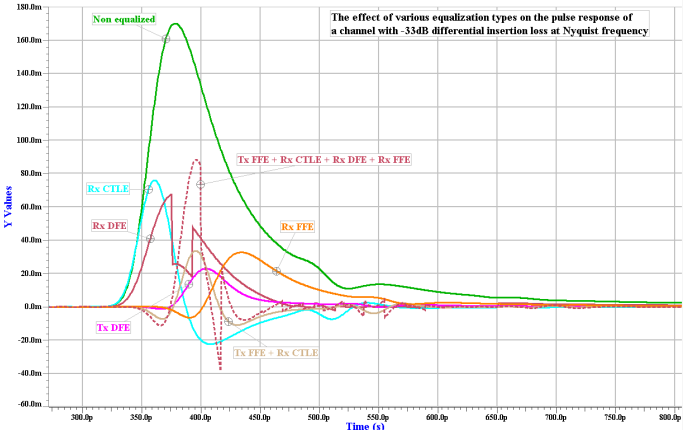


Figure 58 - The effect of various equalization types on the pulse response of a channel with -33dB differential insertion loss at Nyquist frequency

| Equalization Type | Advantages | Disadvantages |
|---|---|---|
| Tx FIR | Immune to noise enhancement, can cancel ISI in pre-cursor and beyond filter span. | Has limited output power, attenuates low-frequencies, and amplifies crosstalk; back channel is required for adaptation. |
| CTLE | Cost effective and simple to implement, active CTLE provides gain and equalization with low power and area overhead, long post-cursor reach, does not require backchannel for adaptation. | Equalization is limited to 1st order compensation, very sensitive to PVT and hard to tune, no precursor cancellation, amplifies high frequency noise and crosstalk, linearity can be challenging. |
| DFE | Noise and crosstalk neutral, does not require backchannel for adaptation. | Nonlinear, can only cancel postcursor ISI, error propagation, feedback loop latency with critical timing path, requires complex CDR design. |
| Rx FFE | Does not have limited output power as Tx FIR, does not require backchannel for adaptation. | Amplifies noise, precision, setting coefficients require adaptive algorithms. Tuning delays for analog implementations. |

Table 7 – The performance of various equalization types

# Adaptation algorithms for equalizers

The simulator needs to search through available equalization combinations to determine the best equalization settings that generate the maximum achievable margin for the link.

The search space can be defined by:

- Tx presets or coefficients space for FFE
- CTLE DC and AC gain
- CTLE poles location
- DFE coefficient magnitude

The most common equalization search routines are:

- Exhaustive
- Coordinate Descent
- Zero Forcing Solution (ZFS)
- Minimum Mean Square Error (MMSE)
- Least Mean Squares (LMS)
- Sign-Sign LMS
- Recursive Least Squares (RLS)

# Order of applying equalization during the optimization process

Since different equalization methods have their own advantages and disadvantages in the way they deal with various channel impairments, the order in which those equalization methods is used in the optimization process might lead to different results depending on the channel's characteristics. Several possible optimization schemes are:

- FIR -> CTLE -> DFE (recommended for heavy insertion loss channels)
- FIR -> CTLE + DFE (works better for large crosstalk noises)
- CTLE -> FIR -> DFE (good for strong impedance discontinuities)
- CTLE -> FIR + DFE (works well for moderate loss applications)
- CTLE -> FIR -> CTLE-> DFE (provides better performance in general)
- CTLE-> FIR + DFE-> CTLE + DFE (works well for moderate loss applications)
- CTLE + DFE-> FIR + DFE -> CTLE + DFE (works well for moderate loss applications)

# Figures of Merit (FOM)

Best equalization is determined based on a figure of merit (FOM). The most commonly used FOMs are:

- Eye height (EH)
- Eye width (EW)
- Sampler margin
- Eye area
- Signal to Noise Ratio (SNR)

# Summary and conclusions

From an IC design perspective, deciding what types of equalization circuits need to be incorporated in a transceiver's design has cost, power and silicon area implications. Often, a tradeoff between those considerations leads to the final decision. Incorporating complex, powerful transceivers in ICs simplifies the task of systems designers but increases the cost of the silicon. Cheap transceivers impose tough electrical and mechanical constraints on the PCB and package design. In most cases the SerDes protocol specifications are proposing budgets that balance the two.

From an EDA perspective, the large variety of compliance methods and ways to define equalization types along with their associated transfer functions and constraints leads to increased complexity of analysis tools. On one hand, the tool should be flexible enough to allow the user to simulate a large variety of protocols; on the other hand, the tool should be easy to use.

## References


1. PCI Express® Base Specification Revision 5.0 Version 1.0, PCI-SIG 02 April 2019
2. Universal Serial Bus 3.1 Specification, Revision 1.0, July 26, 2013
3. IEEE Std 802.3bj™-2014, IEEE Standard for Ethernet, Amendment 2: Physical Layer Specifications and Management Parameters for 100 Gb/s Operation Over Backplanes and Copper Cables
4. JESD204C, JEDEC Standard for Serial Interface for Data Converters, December 2017
5. Stephen H. Hall, Howard L. Heck, Advanced Signal Integrity for High-Speed Digital Designs
6. Emmanuel C. Ifeachor, Barrie W. Jervis, Digital Signal Processing , A practical Approach
7. Hsinho Wu, Masashi Shimanouchi, Mike Peng Li, Effective Link Equalization using FIR, CTLE, FFE, DFE, and FEC for Serial Links at 112 Gbps and Beyond, DesignCon 2018
8. Hossein Shakiba, Shayan Shahramian, Behzad Dehlaghi, David Cassan, Davide Tonietto, A Methododlogy for Performance Comparison of Center and Edge Sampling in Serial Links, DesignCon 2019
9. David Robert Stauffer, Jeanne Trinko Mechler, Michael A. Sorna, Kent Dramstad, Clarence Rosser Oglivie, Amanullah Mohammad, James Donald Rockrohr, High Speed SerDes Devices and Applications
10. Hanqiao Zhang, Steven Krooswyk, and Jeff Ou, High Speed Digital Design
11. Chulwoo Kim, Hyun-Woo Lee, Junyoung Song, High-Bandwidth Memory Interface
12. Sam Palermo, Texas A&M University, ECEN 689: High-Speed Links Circuits and Systems
13. Pavan Kumar Hanumolu, Gu-Yeon Wei and UN-KU Moon, Equalizers for High-Speed Serial Links, International Journal of High Speed Electronics and Systems, Vol 15, N0. 2 (2005) 429-458
14. Serial Attached SCSI-4 (SAS-4) Standard – Project T10/BSR INCITS 534, Revision 04, 15 January 2016
15. Joseph Aday, Hannah Lee, Juli Olenick, Hongtao Zhang, Geoff Zhang, Elimination of Highly Reflective Structures through Sliding Decision Feedback Equalization, DesignCon 2019
16. Optical Internetworking Forum – Clause 28: CEI-112G-LR-PAM4 Long Reach Interface, Implementation Agreement OIF-CEI-05.0, Draft March 24, 2019
17. Optical Internetworking Forum – Clause 23: CEI-112G-VSR-PAM4 Very Short Reach Interface, Implementation Agreement OIF-CEI-05.0, Draft April 15, 2019
18. Cathy Ye Liu, Joe Caroselli, Comparison of Signaling and Equalization Schemes in High Speed SerDes (10-25Gbps), DesignCon 2007
19. Brandon Gore, Richard Mellitz, An Exercise in Applying Channel Operating Margin (COM) for 10GBASE-KR Channel Design

**Siemens Digital Industries Software**

**Headquarters**
Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
+1 972 987 3000

**Americas**
Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
+1 314 264 8499

**Europe**
Stephenson House
Sir William Siemens Square
Frimley, Camberley
Surrey, GU16 8QD
+44 (0) 1276 413200

**Asia-Pacific**
Unit 901-902, 9/F
Tower B, Manulife Financial Centre
223-231 Wai Yip Street, Kwun Tong
Kowloon, Hong Kong
+852 2230 3333

## About Siemens Digital Industries Software

Siemens Digital Industries Software is driving transformation to enable a digital enterprise where engineering, manufacturing and electronics design meet tomorrow. Our solutions help companies of all sizes create and leverage digital twins that provide organizations with new insights, opportunities and levels of automation to drive innovation. For more information on Siemens Digital Industries Software products and services, visit siemens.com/software or follow us on LinkedIn, Twitter, Facebook and Instagram. Siemens Digital Industries Software – Where today meets tomorrow.

**siemens.com/eda**



84078-C1 07/21 D