



WP419 (v1.0) March 27, 2012

Equalization for High-Speed Serial Interfaces in Xilinx 7 Series FPGA Transceivers

By: Harry Fu

The appetite for data is exploding, and the industry is being forced to transition from parallel to serial interface technology at constantly increasing speeds. Higher data rates produce higher bandwidths and present new demands and challenges in transceiver design.

Serial systems today require very low bit-error rates. As the serial signaling rate increases, channel-driven signal distortion increases with it, while the bit sampling time needed to process this distortion only grows shorter. The techniques used to compensate for signal distortion, therefore, become critical elements in the design.

This compensation is usually called *emphasis* in the transmit domain and *equalization* in the receive domain. The effective implementation of emphasis and equalization in Xilinx® 7 series FPGAs and Zynq™-7000 Extensible Processing Platform (EPP) is the subject of this white paper.

7 Series Transceiver Emphasis/Equalization Overview

The Xilinx 7 series FPGA and EPP portfolios are based on the concept of a scalable, optimized architecture, which allows all the FPGA elements to be scaled and optimized to the market and target application. This yields the flexibility required to port design investments to the “right fit” platform, based on current needs. The serial transceivers are no exception. Xilinx offers a portfolio of transceivers that run from 500 Mb/s up to 28.05 Gb/s. See [Figure 1](#).

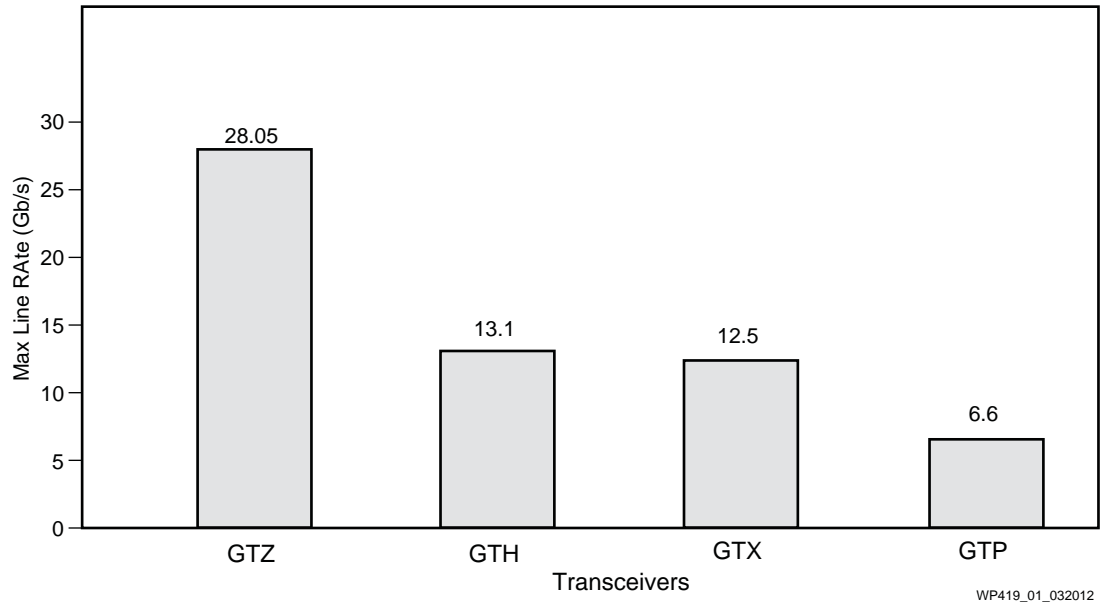
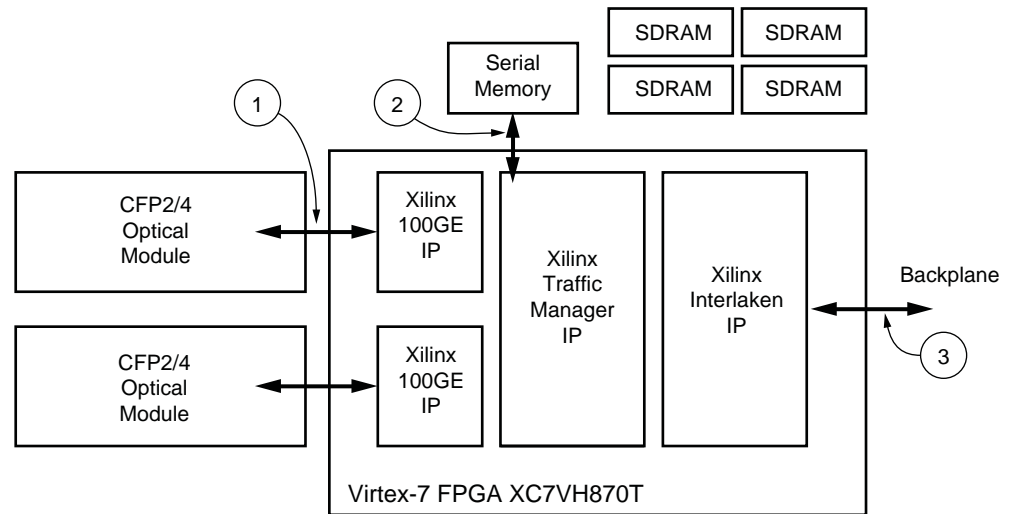


Figure 1: Xilinx 7 Series FPGA Transceiver Portfolio

This scalable approach enables cost-effective solutions for consumer applications, where transceivers are used at lower data rates, to the very powerful, sophisticated transceivers required by the telecommunications equipment used to move data around the world.

[Figure 2](#) depicts an example 2 x 100GE line card with traffic manager design showing three different uses of 7 series serial transceivers.



WP419_02_031912

Figure 2: 2 x 100GE Traffic Manager: Example Transceiver Use Cases

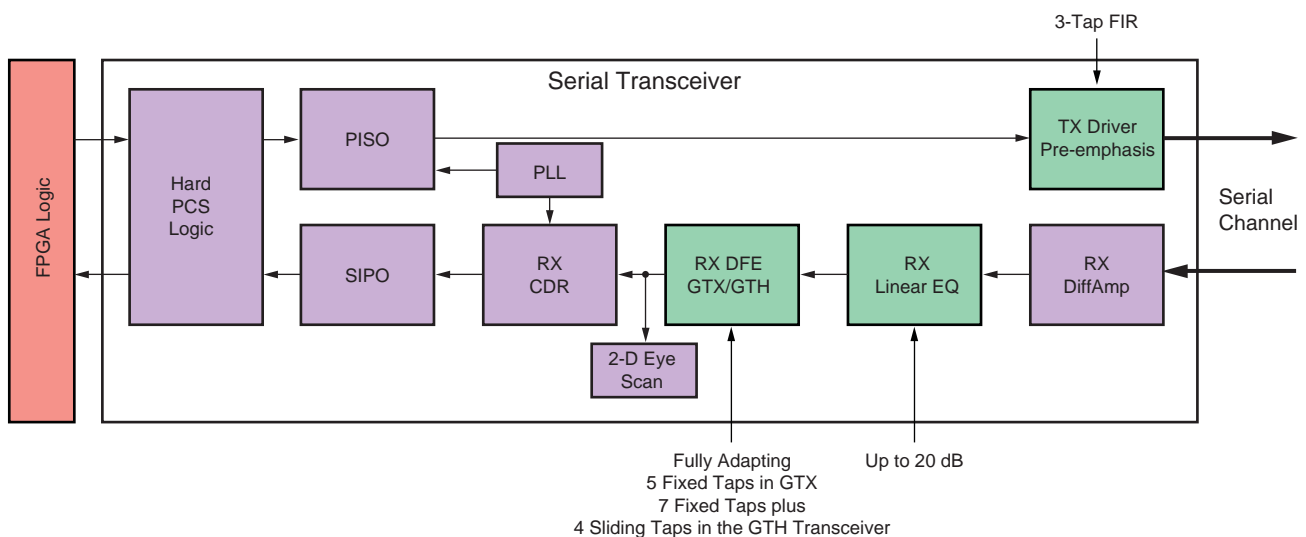
The three example use cases shown in Figure 2 are:

1. GTZ transceivers (8 x 25.78 Gb/s) connect directly to CFP2 optics via an OIF VSR channel (6 in.)
2. GTH transceivers (16 x 10.3125 Gb/s) connect to a serial memory (for storing descriptors used by the Traffic Manager IP) (6 in.)

Note: This case uses the special “low-power mode” of the GTH transceiver to optimize power for the short channel to the serial memory using TX emphasis and RX linear equalization.

3. GTH transceivers (24 x 12.5 Gb/s) connect over a mezzanine card connector, mid-plane, or backplane using decision feedback equalization (DFE) mode (40 in.)

To enable creation of a robust high-speed serial link system with these transceivers running at low power, equalization techniques are built into each transceiver, based on its targeted line rate and targeted applications. See Figure 3.



WP419_03_031912

Figure 3: 7 Series FPGA and EPP Transceiver Equalization Blocks

Transmitter emphasis is a commonly used technique to compensate for the channel losses between transceivers. All 7 series FPGA transceivers offer three-tap TX emphasis.

Receiver equalization varies with different transceivers:

- For systems running at ~6 Gb/s or below that utilize the GTP transceiver, the Continuous Time Linear Equalization-only (CTLE) structure has proven to be very efficient and has been used successfully over the past decade.
- For systems running at ~10 Gb/s – 15 Gb/s, the channel complexity can vary greatly, from simple chip-to-chip interfaces, to very complicated backplanes with long traces, blind vias, and connectors. In complex interface applications running at these line rates, signal distortions exist that make it impossible to guarantee adequate margins with a linear equalization process like CTLE. Therefore, in addition to CTLE, GTX/GTH transceivers offer DFE, a non-linear equalization technique and a popular option in the industry for 10G+ serial links.
- For systems running beyond 20 Gb/s, the GTZ transceiver is designed for short-reach applications directly connecting to CFP2 and CFP4 optics (OIF-CEI28-VSR), or for longer channels with re-timers. In this scenario, a CTLE-only receiver is sufficient.

Channel Conditions and Signal Distortion Components

Channel insertion loss and channel discontinuity are the most widely seen conditions that degrade signal quality. This section illustrates how each one impacts the signal.

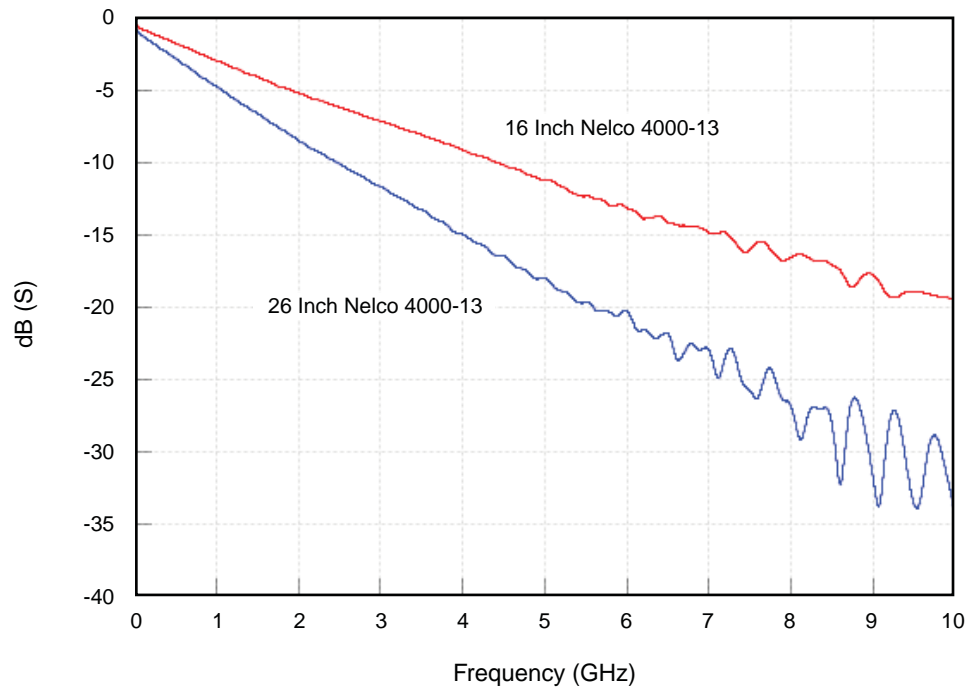
Channel Insertion Loss

When signals travel through a transmission line, such as a board trace or backplane, signal power is degraded and the high frequency components are attenuated more than the low frequency components. This impact is referred to as channel insertion loss. Channel conditions such as conductor loss, dielectric loss, and surface roughness can all result in channel insertion loss.

Channel loss is frequency-dependent and is usually quantified in the frequency domain at the Nyquist frequency of the data stream:

$$f_{Nyquist} = (\text{data rate}) / 2 \quad \text{Equation 1}$$

Figure 4 shows insertion loss from two sample traces made of Nelco 4000-13. One is 16 inches long, the other 26 inches long.



WP419_04_031912

Figure 4: Differential Insertion Loss of Sample Nelco 4000-13 Traces (16-Inch, 26-Inch)

In general, the insertion loss increases as the frequency increases. Given the same material and geometry, the insertion loss also increases as the trace length increases. Table 1 lists the insertion loss at 3 GHz ($f_{Nyquist}$ of 6 Gb/s data streams) and 5 GHz ($f_{Nyquist}$ of 10 Gb/s data streams) for the traces shown in Figure 4.

Table 1: Insertion Losses for Sample 16-Inch and 26-Inch Nelco Traces

	3 GHz	5 GHz
16 inch Nelco	~7 dB	~11.5 dB
26 inch Nelco	~12 dB	~18 dB

At a given data rate, the channel loss impact on the data stream can also be viewed in the time domain. Figure 5 shows the non-ideal pulse responses of the channels shown in Figure 4 at 10 Gb/s. Because of the insertion loss, one data symbol “lifts” the neighboring data symbols. The greater the insertion loss, the greater the lift — and, therefore, more symbols are lifted.

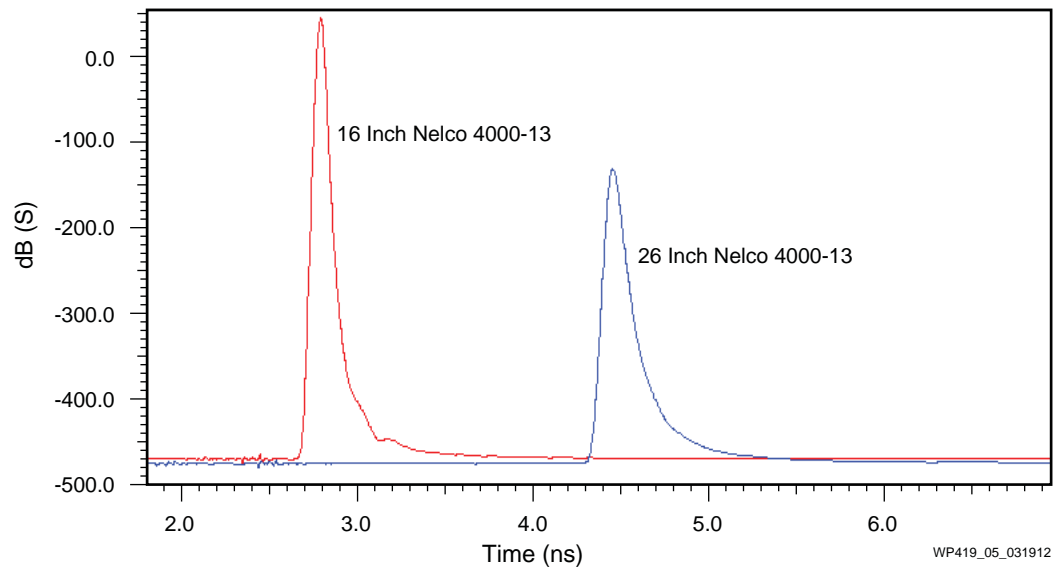


Figure 5: Non-ideal Pulse Responses of the Nelco Traces

The impact from neighboring symbols is usually referred to as “inter-symbol interference” (ISI). Figure 6 illustrates how ISI can change the data stream in a linear system. A received 01000 is changed by the pseudo-channel. In a linear system, a bit stream 011110 is a summation of UI-delayed 01000s. Because the pseudo-channel introduces ISI, the last '0' of the 011110 stream becomes a potential error bit.

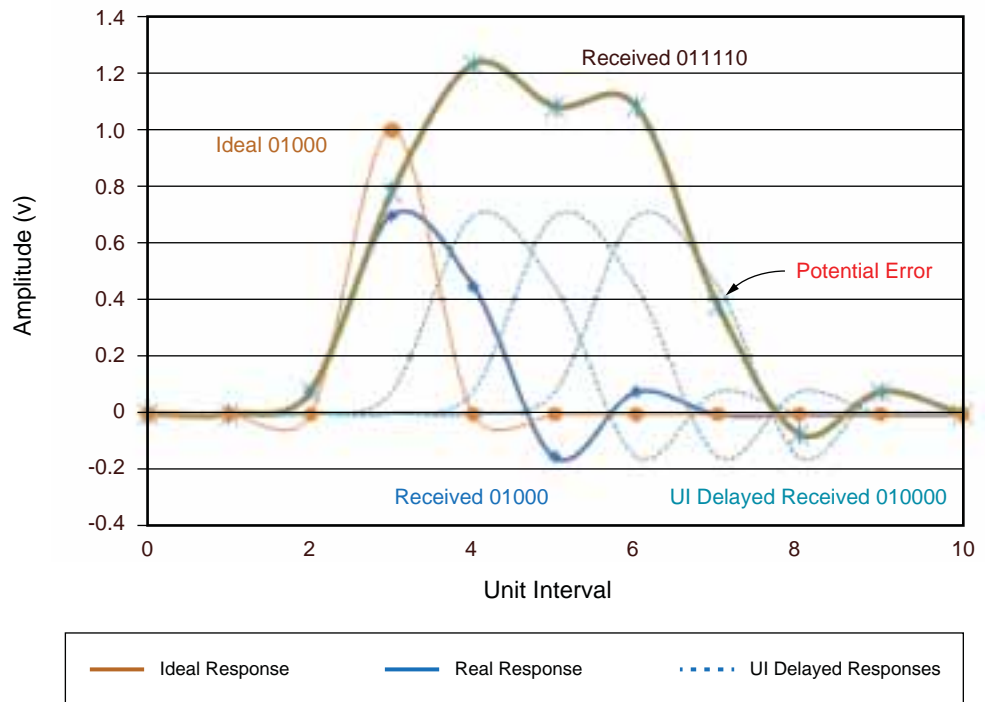


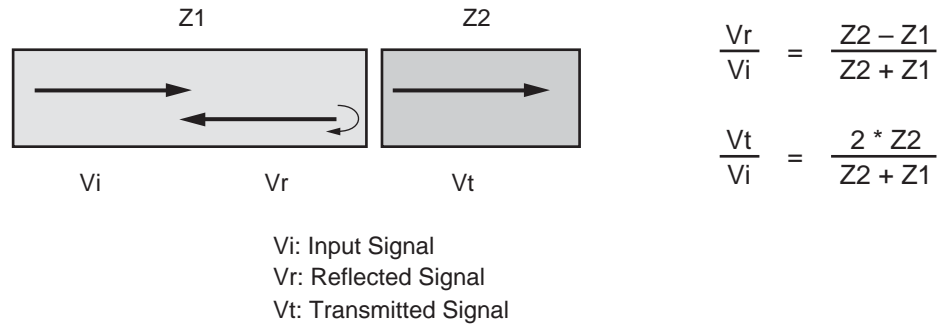
Figure 6: Sample ISI Impacts on Data Stream

Channel Discontinuity

Channel discontinuity happens when the characteristic impedance of the channel changes along the transmission line and results in signal reflection during the

transmission. Figure 7 illustrates how impedance mismatches cause reflections to occur and how the waveform splits across the interface.


Z1 and Z2 are the characteristic impedances on the two sides of the interface.



WP419_07_032012

Figure 7: Reflection Due to Impedance Mismatch

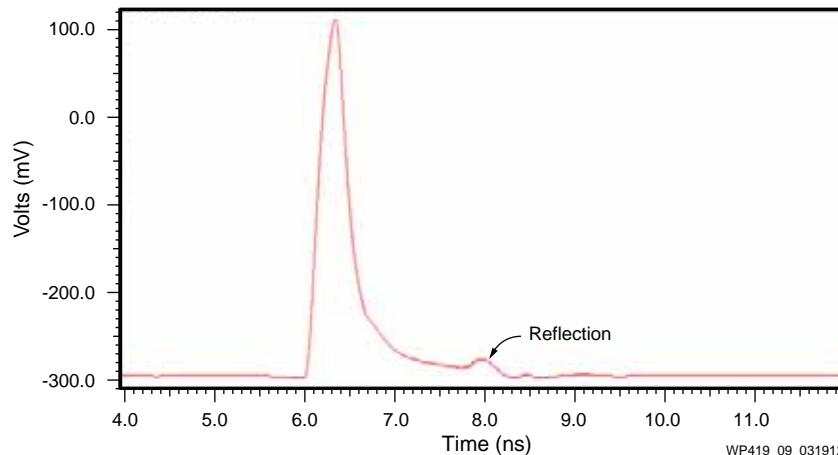
When multiple reflections occur, the transmitted signal is reshaped by the extra Vt. Figure 8 shows an extra Vt at the Z1-Z2 interface due to the second reflection.



WP419_08_031912

Figure 8: Extra Vt Due to Second Reflection

Signal reflection can result in signal amplification or signal destruction. These impacts can be observed in the pulse response. Figure 9 shows an example of a channel reflection appearing on the pulse response: i.e., a “bump” at ~8 ns and a “dip” following that. The distance between the bump to the main symbol depends on the length of the trace along which the reflected signal bounces back.



WP419_09_031912

Figure 9: Reflection on Pulse Response

Techniques Used to Combat Insertion Loss and Discontinuities

Here are a few techniques to combat insertion loss and discontinuities.

Transmitter End Techniques

TX emphasis has been used for many years as a simple solution to combat insertion loss. Because insertion loss causes high-frequency signal components to lose more energy than low-frequency components, TX emphasis either:

- boosts the high-frequency components
- suppresses the low-frequency components at the transmitter.

After channel losses degrade the high-frequency component of the TX emphasized signal, the high-frequency and low-frequency components of the signal should appear to be more balanced.

De-emphasis

De-emphasis suppresses the low-frequency components at the transmitter. When a 0-to-1 or 1-to-0 transition occurs, regular (nominal) voltage swing is applied to the symbols that have a new value after the transition. The symbols that keep the same value after the transition, however, use a reduced voltage swing. The amount of swing voltage reduction is referred to as the “tap weight” of the de-emphasis.

Figure 10 and Figure 11 show how TX de-emphasis can reshape a signal based on simulations with 7 series transceiver IBIS-AMI models.

Two GTX transmitters are configured to send a repeating 11110000 pattern.

Figure 10 shows the data stream without any de-emphasis. The symbols following the transition have a peak-to-peak amplitude of $\sim 0.28\text{V}$.




Figure 10: 11110000 Binary Pattern, No De-emphasis

Figure 11 shows the data stream with 6 dB of de-emphasis. The symbols following the transition now have a peak-to-peak amplitude of $\sim 0.14\text{V}$. The difference is because the de-emphasis tap weight applied to the data stream is 6 dB.


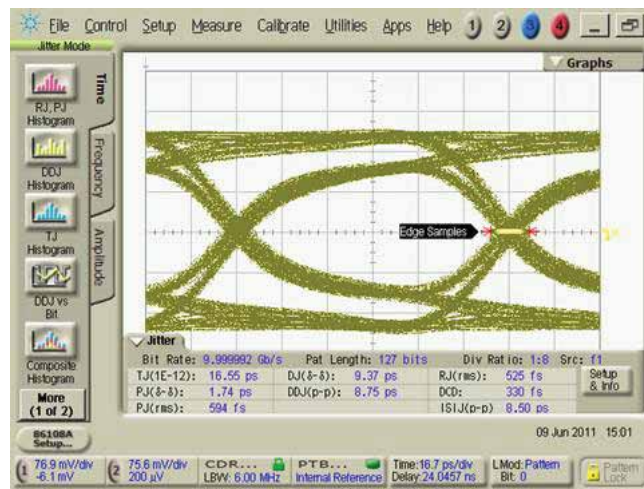


Figure 11: 11110000 Binary Pattern, 6 dB De-emphasis

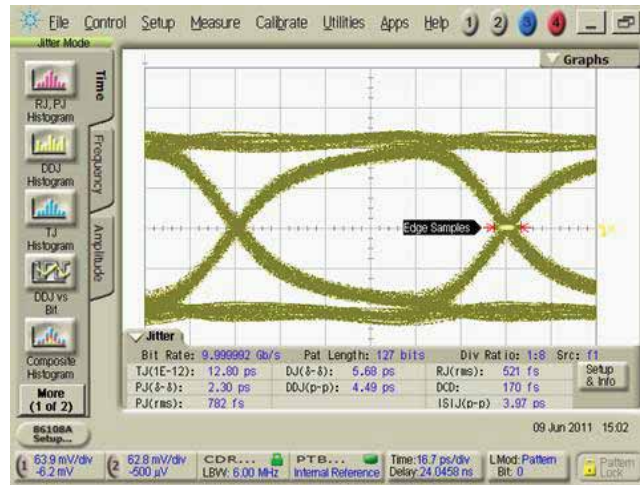
The de-emphasis implementation can reduce the ISI to ensure correct signal sampling at the receiver. Figure 12 and Figure 13 show scope shots of Kintex®-7 FPGA GTX transceiver eye diagrams and jitter decomposition at 10 Gb/s. The signal goes through the package, the characterization board trace, the connector, and the SMA cable. With the de-emphasis off, the eye shown in Figure 12 is distorted, and the peak-to-peak ISI jitter [ISI](p-p) reads 8.50 ps on the scope.



WP419_12_031912

Figure 12: GTX Transceiver 10 Gb/s Eye Diagram without De-emphasis

With ~2 dB post-tap de-emphasis, the eye in Figure 13 looks cleaner, and the ISI jitter decreases to 3.97 ps.

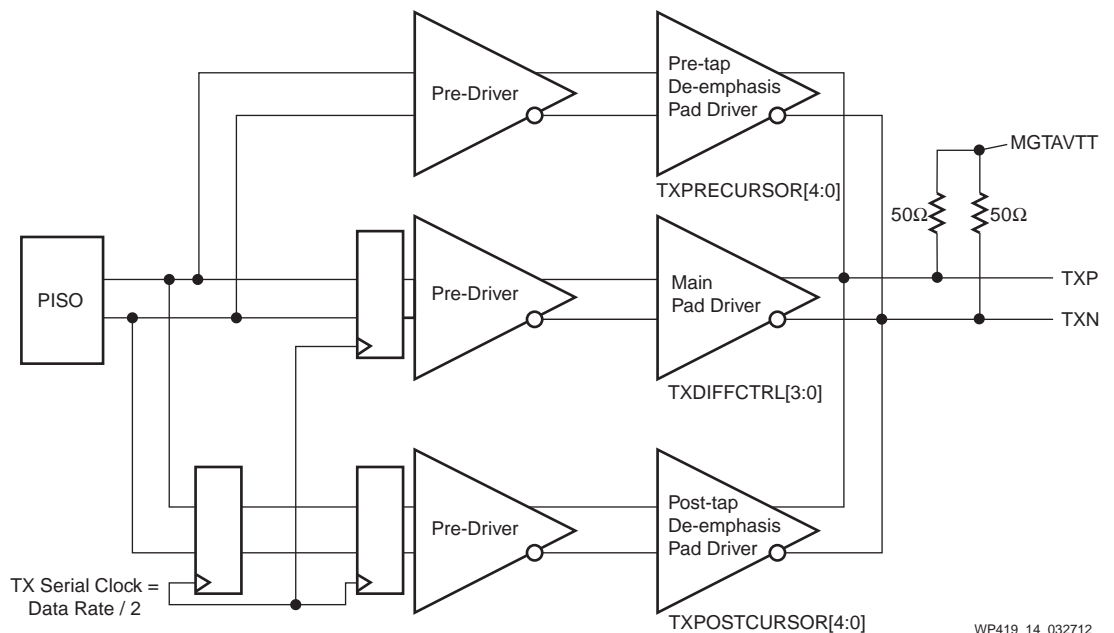


WP419_13_031912

Figure 13: GTX Transceiver 10 Gb/s Eye Diagram with 2 dB Post-tap De-emphasis

Three-Tap TX De-emphasis

In 7 series FPGAs and EPPs, all transceivers have a three-tap FIR to implement TX de-emphasis. The three taps include the pre-tap, main tap, and post-tap. Figure 14 shows the TX driver block diagram with three-tap de-emphasis.



WP419_14_032712

Figure 14: TX Driver Block Diagram with Three-Tap De-emphasis

Figure 10 and Figure 11 show the post-tap de-emphasis impact on the signal. The pre-tap de-emphasis is similar to post-tap. The difference is that the symbol that has the high swing is the one *before* the transition rather than after the transition. Figure 15 shows, a simulation example of GTX transmitter output, how a 6 dB pre-tap de-emphasis can reshape a signal repeating 11110000.




Figure 15: 11110000 Binary Pattern, 6 dB Pre-tap De-emphasis

Programmable TX De-emphasis

In 7 series FPGA transceivers, the tap weights are all programmable to meet different channel conditions. For example, the GTX/GTH/GTP transceivers have 32 settings for post-tap de-emphasis, up to 12.96 dB, and 21 settings for pre-tap de-emphasis, up to 6.02 dB.

Note: The maximum pre-tap boost is not as high as the maximum post-tap boost. The reason for this is that, in general, the pre-cursor degradation due to channel loss is small compared to post-cursor degradation.

Besides the weight, the polarity of the de-emphasis taps is also programmable. This gives more flexibility of the de-emphasis to fit different channel loss conditions.

Receiver End Techniques

At the receiver, there are also similar equalization techniques to boost the high-frequency components of the signals. Continuous Time Linear Equalization (CTLE) is one of the most popular linear equalization techniques at the receiver end of the line.

RX Continuous Time Linear Equalization (CTLE)

The concept of CTLE can be explained in the frequency domain. Figure 16 shows the conceptual frequency response of channel and CTLE and the result of the combination. There is more channel insertion loss for high-frequency signals than low-frequency signals. Thus, the link channel can be viewed as a low-pass filter. To compensate for the low-pass characteristics of the channel, a high-pass filter is added at the receiver to achieve balance between the high-frequency and low-frequency components of the data stream.




Figure 16: Frequency Domain Effects of CTLE

Selecting the Correct CTLE

When selecting a CTLE to compensate for channel loss, it is important to ensure that the peak of the CTLE frequency response is at the correct frequency and the correct gain (boost). Incorrect CTLE selection results in under-equalization or over-equalization, and thus, in suboptimal post-CTLE signal integrity. Figure 17 through Figure 19 show simulation examples of under-equalization, over-equalization, and correct equalization (respectively) using the GTX receiver and a sample 16-inch Nelco 4000-13 trace. Under-equalization does not open the eye enough in both the horizontal and vertical directions; over-equalization can result in high jitter. The eye with correct equalization shows low noise and low jitter.




Figure 17: Under-Equalized Eye Diagram




Figure 18: **Over-Equalized Eye Diagram**




Figure 19: **Correctly Equalized Eye Diagram**

Programmable CTLE

To support varying channel conditions, the correct CTLE selection is needed. All 7 series FPGAs and EPPs offer programmable CTLE in their receivers to address this concern. The thousands of possible frequency responses with different high-frequency peaking can address a multitude of channel variations.

CTLE Auto Adaptation for Ease of Use

With thousands of settings available, proper tuning of the CTLE is critically important. As illustrated in Figure 17 through Figure 19, under-equalization or over-equalization can result in suboptimal signal integrity.

To resolve this issue, 7 series FPGA transceivers offer an auto-adaptation feature. The adaptation removes the CTLE tuning burden from the user by automatically finding the optimal setting. Based on the incoming data streams, it tests different settings and selects one that provides the sufficient system margins.

Users can employ a one-time calibration for a given channel, or leave auto-adaptation on for run-time calibration. Using auto-adaptation adds less than a 5% power consumption increase to the transceiver. For applications that might experience voltage and temperature variations during operation, the run-time calibration can be very useful.

RX Decision Feedback Equalizer (DFE)

Linear equalization techniques such as RX CTLE have one major limitation: noise. When noise (such as reflections or crosstalk) is present on the channel, CTLE amplifies the high-frequency noise right along with the data.

Decision Feedback Equalizer (DFE) is a successful technique to mitigate ISI without amplifying the noise. It works by directly removing the ISI from a previous bit, allowing the following bits to be correctly sampled. Figure 20 shows the block diagram of a simple DFE implementation.

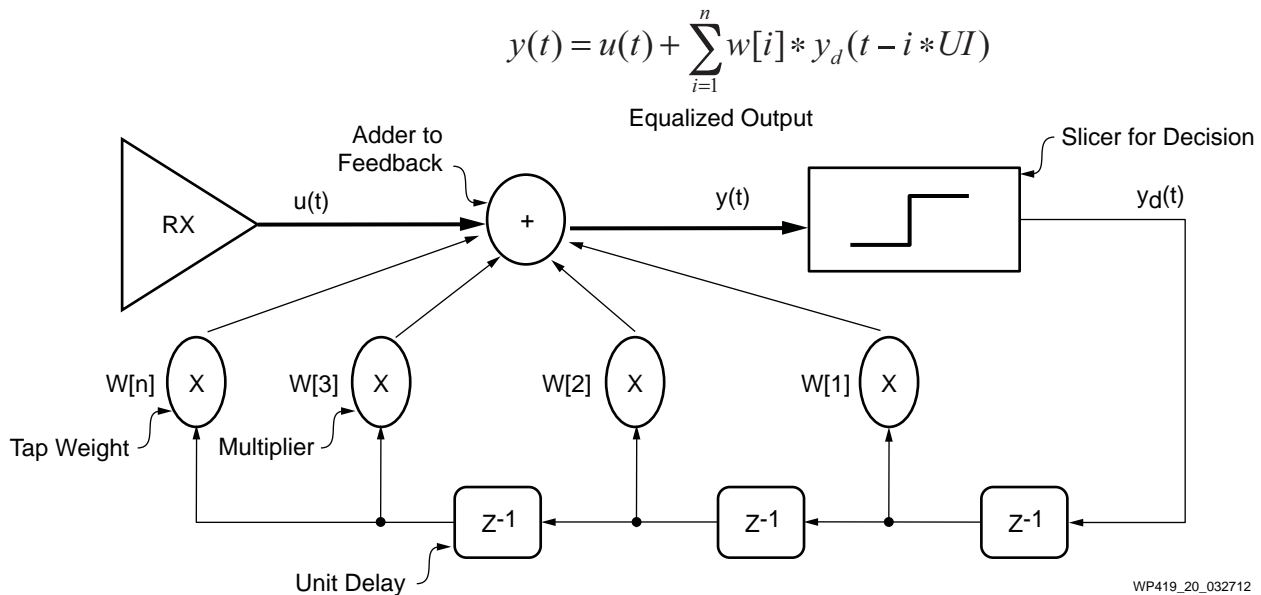



Figure 20: Simplified DFE Block Diagram

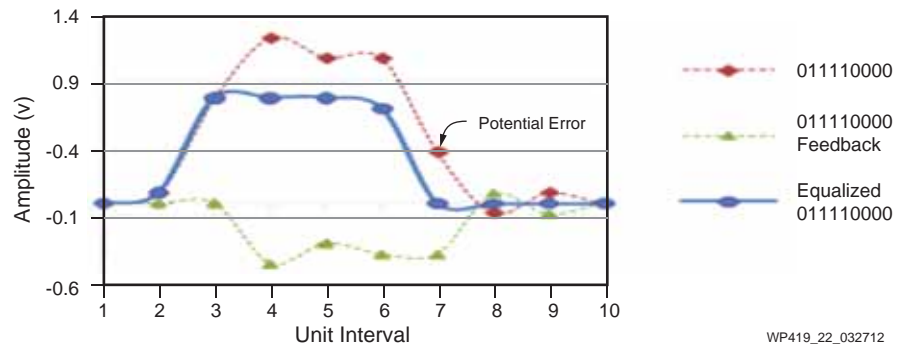
As shown in Figure 5, the ISI effect from this symbol applied to the following bits is determined by the pulse response, based on channel characteristics. DFE starts with a “decision slicer” to determine whether the current symbol is High or Low. The resulting symbol goes through unit delays and multiplies with the tap weights. The weighted delayed signals are added together to the input analog signals. If the tap weights are well selected to cancel the ISI at each following symbol, the result of this feedback loop is able to compensate for as many taps of ISI as the DFE has.

Figure 21 and Figure 22 illustrate two sample data streams equalized by the DFE.



WP419_21_032712

Figure 21: 010110000 Binary Pattern with DFE Equalization Applied



WP419_22_032712

Figure 22: 11110000 Binary Pattern with DFE Equalization Applied

The red curve in Figure 21 is the pre-equalized data stream of 010110000. Without any equalization, the second '0' could be sampled incorrectly due to the ISI. The green curve is the feedback stream added to the original data stream with well-selected DFE tap weights. The blue curve is the final data stream after applying the feedback stream. Compared to the red curve, the blue curve has much more margin and is likely to be sampled correctly.

Similarly, Figure 22 shows the DFE equalizing a data stream of 011110000. The first '0' after the '1's is a potential bit error. After DFE equalization, it has a large margin and is likely to be sampled correctly.

DFE Advantages in 7 Series FPGA GTX/GTH Transceivers

One way to judge the capability of a DFE is the number of taps it supports. The DFE compensation stops when it reaches the last tap of the feedback loop, so the more taps, the more ISI "tail" effect can be compensated for by the DFE. It is also notable that the reflections going through a trace with a length of more than a few inches shows up in the pulse response after a number of taps. For example, at 10 Gb/s, reflections through a 6-inch trace are likely to show up in the pulse response after 20 taps. However, more DFE taps can cost more silicon area as well as more power consumption.

DFE in GTX Transceivers

The DFE capability of GTX transceivers is the best offering in the FPGA industry. GTX transceivers support five fixed taps to cover high-frequency needs as well as short reflections, such as those that occur within packages and connectors. Competing devices have similar taps but have trouble delivering the full adaptation.

DFE in GTH Transceivers

GTH transceivers have the most DFE taps in the FPGA industry: seven fixed taps and four sliding taps that can reach up to 63 taps. The seven fixed taps can be used for high-frequency attenuation and short-distance reflections. The four sliding taps are designed to compensate for reflections on channels up to 18 inches long at 10 Gb/s, without sacrificing the power consumption of 63 fixed taps. This combination of fixed and floating taps enables great compensatory capability while keeping the power consumption low and costs in check.

Programmable Auto-adaptive DFE in GTX/GTH Transceivers

The tap weight selection is critical to the DFE functionality. If the tap weights are wrong, the feedback loop can have negative impact on the data stream, which can lead to wrong data sampling and then propagate to the following data streams. Variation of tap weights is necessary to meet different channel characteristics. Similar to TX emphasis, DFE tap weights in GTX/GTH transceivers are designed to be programmable.

Choosing the right setting out of thousands of DFE tap weights combined with many different possible CTLE settings is the major challenge for DFE designs. To make practical use of a DFE, an adaptation algorithm is required. Both GTX and GTH receivers support auto-adaptation when running under DFE + CTLE mode. Similar to CTLE adaptation, users can employ a one-time DFE calibration for a given channel, or leave auto-adaptive DFE on for constant run-time calibration.

Figure 23 illustrates DFE coverage in GTX and GTH transceivers.

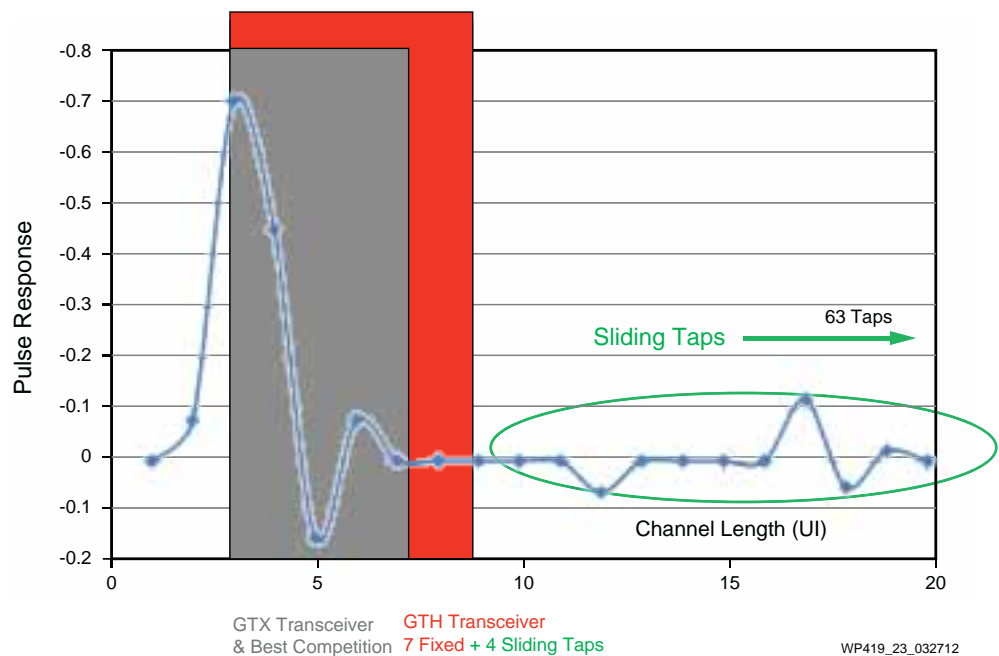


Figure 23: DFE Taps in GTX and GTH Transceivers

Robust Performance of TX Emphasis, RX CTLE, and RX DFE in GTX/GTH Transceivers

The DFE of GTX and GTH transceivers is designed to work together with TX Emphasis and RX CTLE, which are well suited to mitigate some or most of the insertion loss. The TX emphasis can boost up to 12.96 dB, and the CTLE in DFE mode is enhanced to boost up to 20 dB between 4 GHz and 5 GHz.

DFE is the final stage to compensate for the rest of the impairments in the channel before the incoming signal reaches Clock Data Recovery (CDR).

To avoid an “improper equalization” situation that can apply the wrong frequency boost, amplify noise, or directly introduce bit errors, it is critical to find the right combination of TX emphasis, RX CTLE, and RX DFE. The adaptation feature at the receiver makes this tuning process easy. Refer to [UG476, 7 Series FPGAs GTX Transceivers User Guide](#) for more details on how to use these equalization features.

The strategy of employing a combination of TX Emphasis, CTLE, and DFE has worked well for backplanes, optics, plug-in cards, or even simple chip-to-chip applications. [Figure 24](#) shows an internal eye diagram created after the DFE, measured on a Kintex-7 FPGA GTX receiver in a backplane system running at 10.3125 Gb/s. This eye diagram was produced by using the on-chip Xilinx 2-D Eye-Scan feature and the ChipScope™ toolset included in the ISE® Design Suite. The backplane system has a total differential insertion loss of more than 28 dB at 5 GHz. The plotted eye shows the equalization combination has well compensated the channel conditions and resulted in plenty of margin for the backplane system.

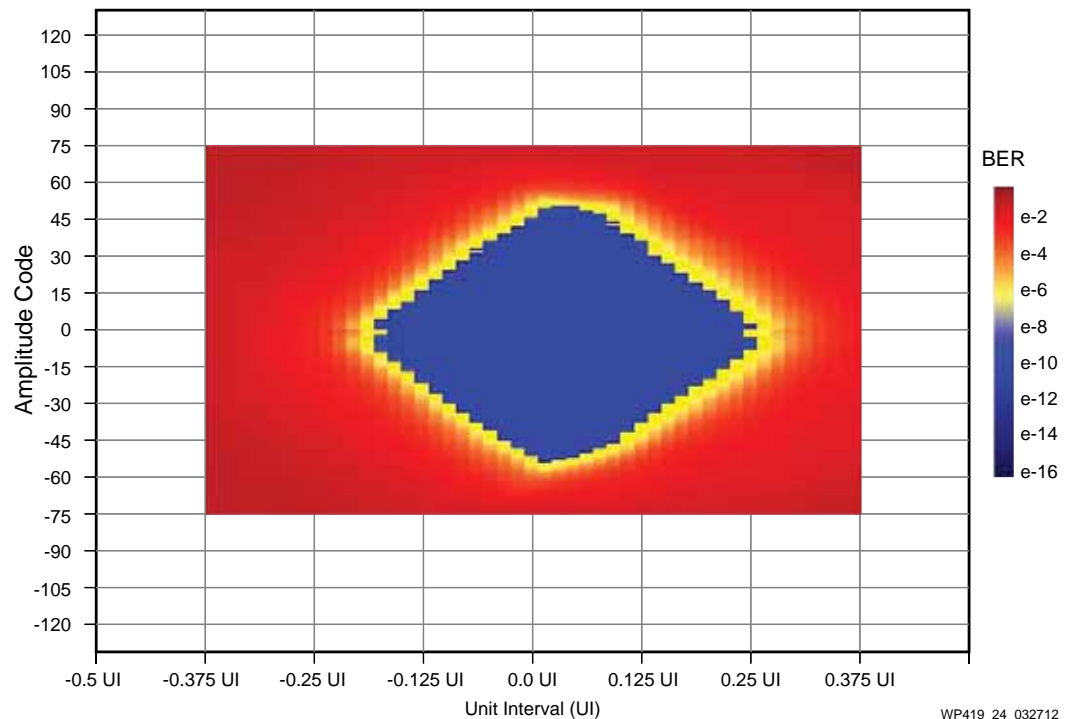


Figure 24: 2-D Eye Scan at the Kintex-7 FPGA Receiver through a Backplane System

Summary

7 series FPGAs offer a variety of transceivers to scale and optimize across different needs and applications. Each transceiver runs at different target line rates and has appropriate equalization techniques for the line rates and channels expected for those devices. The combinations of the on-chip equalization features such as TX emphasis, RX CTLE, and RX DFE (available in the GTX and GTH transceivers) offer the best equalization capability for even the most demanding backplane and optical applications in the FPGA industry.

Table 2 summarizes the equalization modes in each 7 series transceiver.

Table 2: Equalization Modes in Each 7 Series Transceiver

Transceiver	TX Emphasis	RX CTLE Only	RX CTLE + DFE
GTP	3-tap	>9 dB up to 6.6 Gb/s	No
GTX	3-tap	>9 dB up to 10.3125 Gb/s	Up to 20 dB CTLE and DFE (5 fixed-tap, auto-adaptive), up to 12.5 Gb/s
GTH	3-tap	>9 dB up to 10.3125 Gb/s	Up to 20 dB CTLE and DFE (7 fixed-tap + 4 sliding-tap up to 63, auto-adaptive), up to 13.1 Gb/s
GTZ	3-tap	>9dB	No

More information is available in the transceiver user guide for each transceiver. Also go to: <http://xilinx.com/products/technology/transceivers/index.htm>.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
03/27/12	1.0	Initial Xilinx release.

Notice of Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.