



WP427 (v1.0) December 20, 2012

Pluggable Optical Interfaces and Their Compatibility with Xilinx FPGAs

By: Martin Gilpatric

Electrical interconnects are limited in their practical drive distance based on the electrical loss characteristics of the channel. Moving a serial link from the electrical domain to the optical domain has a number of benefits that make it a spectacular solution in certain cases. Although optical interconnects suffer channel losses as well, the distances that can be spanned in the optical domain are orders of magnitude longer. The trade-offs presented by moving to the optical domain include additional components, additional manufacturing complexity, and a new set of variables that must be considered when designing a robust optical system.

This white paper is intended to help designers understand the available options in the pluggable optical module market, along with the trade-offs that must be considered when making design choices.

Introduction

Pluggable optical modules are used in systems to leverage the rapid, robust long-distance data transfer of which fiber optic networks are capable. Designers must ensure that systems integrating optical networks are robust. This requires knowledge of the *whole system* — the types of optical modules and the trade-offs associated with their use, the electrical and optical standards for which the optical module is designed, and the host transceiver features that are required to transmit and receive data through an optical system.

Optical Modules and Fiber

When a designer looks at implementing an optical interface, several choices need to be made based on the protocol and application:

- Types of optical modules.
Many form factors exist, varying in size, features, electrical and optical specifications, and lane width.
- Type of fiber.
The choice is between multi-mode and single-mode fibers, depending on the interface and the travel distance.
- Supported features.
The features that the module needs to support must be identified based on the electrical and optical specifications that need to be supported.

Types of Optical Modules

Optics have a long history of use in the communications sector. As a result, there is a wide variety of form factors and uses for optical modules. To simplify the discussion, this white paper focuses on the current set of optical modules utilized for 10G, 40G, and 100G Ethernet, though many of the topics translate directly to OTN or other standards.

These Ethernet rates can be supported by a variety of electrical lane widths (usually one, four, or ten lanes), which in turn determine the electrical line rates. For example, either four lanes of 3.125 Gb/s or a single lane of 10.3125 Gb/s can support 10G Ethernet; similarly, either ten lanes of 10.3125 Gb/s or four lanes 25.78125 Gb/s can support 100G Ethernet.

The current crop of high-end optical modules leverages a base rate of 10.3125 Gb/s, for 10G, 40G, and 100G Ethernet in one, four, or ten lanes. Popular optical modules are available that support each of these widths at this rate: SFP+ or XFP for a single lane, QSFP+ for four lanes, and CFP for ten lanes. Module type is constrained only by its form factor and pinout, supporting a variety of electrical and optical standards. This large field of choice adds to the complexity of selecting the correct parts for design into a system. See [Table 1](#).

Table 1: Example Ethernet Optics

Optical Form Factor	Re-timed, Limited, or Linear	Interface Width x Rate	Protocol	Electrical Standard	Supported 7 Series Transceivers
SFP+	Linear, Limiting	1 x 10.3125 Gb/s	10GBase-SR	SFF-8431	GTX, GTH
QSFP	Limiting	4 x 10.3125 Gb/s	40GBase-SR4	XLPII	GTX, GTH
CXP	Limiting	10 x 10.3125 Gb/s	100GBase-SR10	CPPI	GTX, GTH
CFP	Re-timed	10 x 10.3125 Gb/s	100GBase-LR4	CAUI	GTX, GTH
CFP2	Re-timed	4 x 25.78 Gb/s	100GBase-LR4	CAUI4	GTZ

Notes:

1. A selection of common optical form factors, their typical electrical interface type, common Ethernet protocol, common electrical standards, and the 7 series transceivers that support them. This is not an exhaustive list of optical interconnects supported by 7 series transceivers; contact your local sales representative for more information.

Impact of Optical Fiber on a System

One common differentiator between modules is the length and type of fiber a module can drive. SFP+ makes a good example for examining the options.

10G Ethernet defines a number of optical interfaces (SR, LR, LRM, and ER) along with other industry adopted standards (ZR and DWM) that specify the optical wavelength and the length and type of optical fiber that can be supported. Each element of this selection has a different impact on the link performance. For example, 10GBASE-SR uses an 850 nm wavelength and can support up to 300m of 50 μ m multi-mode fiber while 10GBASE-LR uses a 1,300 nm wavelength and can support 10 km of single mode fiber.

The differences between multi-mode fiber and single-mode fiber are important. Single-mode fiber consists of a single strand of fiber that the data is transmitted across. Multi-mode fiber is composed of multiple strands of fiber bundled together, where light can pass across each strand. In either case, an effect known as dispersion can impact the fidelity of the transmit signal. As the signal passes through the fiber, the distribution of wavelengths of light that contain signal content are interacted with in slightly different ways, some wavelengths experiencing more delay or varying degrees of attenuation. The impact of this optical dispersion to the waveform is distinctly different from how copper impacts electrical signals. As a result, high levels of optical dispersion require special circuitry to compensate for, typically in the form of a feed-forward equalizer (FFE) in conjunction with the standard DFE and CTLE structures used for electrical channels.

While both single-mode and multi-mode fiber introduce some dispersion, the single-mode has the advantage that very little dispersion occurs per unit length compared to multi-mode fiber, leaving loss as the primary contributor to signal degradation at the far end of the link. The amount of dispersion introduced in multi-mode fiber varies depending on the laser wavelength. At 850 nm (used for 10GBASE-SR), very little optical dispersion is introduced over the maximum 300m of fiber. Comparatively, at 1,300 nm (used for 10GBASE-LRM), far more dispersion is introduced — enough that it needs to be compensated for after the signal is translated into the electrical domain.

Translating Optical Signals to the Electrical Domain

Plug-in modules act as translators between the electrical signaling used by the integrated circuits (ICs) and the optical signaling that goes over the fiber. As such, the electrical signaling developed by the received optical signal is as important to consider as the optical signaling itself. There are three types of optical modules, and each converts optical signals into electrical signals in a slightly different way.

- Linear modules convert the received optical power into a proportional electrical output. The advantage is that it is easier for an equalizer to compensate for dispersion or losses introduced by the optical channel, or even by the electrical channel on the far end of the link. The disadvantage of a linear system becomes apparent when only a small amount of optical power is received, and the launch amplitude of the electrical signal output by the module can be too low to sufficiently drive the trace between the module and the host IC.
- Limiting modules have an opposite set of trade-offs. A limiting amplifier in the electrical driver ensures that the electrical signal is always sent at full drive strength, eliminating concerns that low optical power limits the electrical drive strength. Equalizing for dispersion becomes more difficult, however, as the limiting amplifier introduces non-linearity between the optics and the host IC.
- Re-timed modules use clock and data recovery (CDR) circuits to clock the data within the optical module to ensure the highest signal integrity when transitioning between the electrical and optical domains. Data is clocked either as it is received from the Host IC in the transmit direction or immediately following the optical/electrical translation in the receive direction, clocked, and then re-transmitted into the other domain. These circuits contain robust equalization structures to ensure that the data being clocked and retransmitted in either direction is correct. By re-timing the data at the electrical/optical interface, any signal integrity degradation is removed, making the retransmitted data very clean.

Re-timed modules are typically more expensive and consume more power due to the additional circuitry required inside the module, but they provide the most robust solutions in applications where long or otherwise challenging channels need to be used.

Electrical Domain Interfacing

For a pluggable optical module to be useful, it needs a host IC that transmits and receives the data being passed through the optics. To make the connection between the module and the host IC, each partner needs to have some confidence that the other will comply to an agreed-upon electrical standard. Many electrical standards exist for serial communications protocols, and the links between optical modules and host ICs are no exception. SFF-8431 (published by the SFF Committee) and IEEE Std 802.3ba provide good examples of specifications that interface to each type of optics (i.e., linear, limiting, and re-timed).

Linear and Limiting Interfaces

SFF-8431 contains specifications for both linear and limiting modules. The transmitter specifications are identical between the two because the same desire exists for minimizing jitter injected into the optical domain. The host receiver specifications, however, change drastically. Without a limiting amplifier separating the optical domain from the electrical domain, all of the dispersion introduced by the optics is

passed into the electrical domain. While similar cross-talk specifications exist between the limiting and linear specifications, the linear specifications drastically change how jitter needs to be calculated and compensated for.

Both 10GBASE-LR and 10GBASE-LRM have their own specifications and highlight the difference between a low-dispersion, single-mode fiber implementation and a high-dispersion, multi-mode implementation. The concept of wave distortion penalty (WDP) is introduced to describe the impact that optical dispersion has on the resultant electrical signal. WDP is calculated for a number of different system eventualities, all of which need to be calibrated and tested against in order to claim compliance. For single-mode 10GBASE-LR systems, this is a single condition referred to as "Low WDP." 10GBASE-LRM has six separate WDP specifications that span both "High WDP" and "Low WDP" conditions.

The limiting standard has specifications that look very similar to other serial interconnect specifications like XAUI or PCI Express®. The transmitting host must meet strict jitter requirements to ensure that very little jitter is propagated into the optical domain. [Table 2](#) defines some of the specifications that need to be met.

Table 2: Signal Characteristics at the Transmit End that Must Be Controlled

Specification	Description
Rise/Fall Time	Minimum value given to control cross-talk
Total Jitter	Total TIE jitter measurement for transmit data
Data-Dependent Jitter	Jitter associated with patterns within the transmit data
Data-Dependent Pulse Width Shrinkage	Maximum offset from a nominal UI that a single bit time can occupy due to patterns in transmit data
Uncorrelated Jitter	Maximum data that cannot be correlated to transmit data pattern
Eye Mask Coordinates	Definitions for an eye mask

Notes:

1. See SFF-8431 for a complete list of host transmitter requirements and their definitions.

Similarly, the receiving host needs to be able to receive data that has incurred some degradation from the far end electrical domain, the intervening optical channel, and the copper trace between the local optical module and the receiving host. These receive specifications (shown in [Table 3](#)) are similar to those of other serial protocols that have jitter tolerance, cross talk, and amplitude requirements.

Table 3: Signal Characteristics at the Receive End that Must Be Tolerated

Specification	Description
Crosstalk Source Rise/Fall Time and Amplitude	Defining the aggressors that need to be overcome to approximate a real system within the definition
99% Jitter	Bounding the inner 99% of the total jitter distribution
Total Jitter	Total TIE jitter measurement for stressed received data
Data Dependent Pulse Width Shrinkage	Maximum offset from a nominal UI that a single bit time can occupy due to patterns in transmit data
Eye Mask Coordinates	Definitions for an eye mask

Notes:

1. See SFF-8431 for complete list of host receiver requirements for supporting limiting modules and their definitions.

Re-Timed Interfaces

The last type, a re-timed module, is the most relaxed from the electrical point of view. The re-timer removes the vast majority of the negative signal effects before re-transmitting the data. The CAUI specification defined in IEEE Std. 802.3ba is targeted to such interfaces, namely those modules using a 100GBASE-LR4 interface where ten lanes of 10.3125 Gb/s electrical data get converted into four lanes of 25.78Gb/s optical data. The conversion between these two data rates is done in logic within the module, necessitating the re-timing of the data. The resulting system resembles a chip-to-chip electrical interface with short traces and a small, high-quality connector in the path. As a result, the CAUI specification defines an interconnect that has more relaxed requirements than either the SFP+ linear or limiting specifications. See [Table 4](#).

Table 4: Signal Characteristics for CAUI

Transmitter Specification	Description
Maximum Differential Output Voltage	Peak-to-peak between the P and N signal pins
Min/Max De-emphasis	Defining the ranges that de-emphasis (commonly known as pre-emphasis) can take
Total Jitter	Total TIE jitter measurement for transmit data
Deterministic Jitter	Deterministic jitter measurement of the transmit data
Eye Mask Coordinates	Definitions for an output eye mask
Receiver Specification	Description
Rise and Fall Time Tolerance	The RMS value defining the minimum rise and fall times that the receiver can tolerate
Input and Common-mode Input Return Loss	Return loss specifications for the input of the receiver
Minimum Total Input Jitter Tolerance	Calibrated total TIE jitter measurement of a stressed input signal the receiver can tolerate
Minimum Deterministic Input Tolerance	Calibrated deterministic jitter component of a stressed input signal that the receiver can tolerate
Eye Mask Coordinates	Definitions for an input eye mask

Notes:

1. See IEEE Std 802.3ba for a complete list of transmitter and receiver requirements for CAUI interfaces and their definitions.

Optical Compatibility of Xilinx 7 Series Transceivers

The purpose of a pluggable optical module is to translate the serial connectivity that cannot be supported by traditional electrical signaling. Since those same traditional electrical signals are required to interface to the optical module, it is worth considering how the performance of the host IC impacts the robustness of the interconnect. While a re-timed module has similar requirements to a standard chip-to-chip interface, linear

and limiting optics translate directly from the electrical to the optical domain and have special requirements as a result.


Transmit Jitter

In the transmit direction, from host to module, keeping jitter low is paramount. Any jitter present at the interface to the optics is converted, in some way, into the optical domain. If there is a limiting amplifier, the non-linearity makes it almost impossible for the far end receiver to equalize for any near-end board characteristics between the host and the optical module. The driver must therefore provide some sort of preemptive equalization for that portion of the channel. Xilinx has enabled the 7 series transceivers with two important features that help to provide the lowest possible jitter to the electrical-to-optical interface: low-jitter, high-performance PLLs, and three-tap equalizers in the transmit driver.

A large part of the random jitter in a system can be traced back to the PLL used to generate the signal that clocks out the serial data. By incorporating industry-leading, high-performance LC tank PLLs into the GTX, GTH, and GTZ 7 series transceivers, Xilinx has provided the best starting point for designers looking to interface to pluggable optical modules. Comparing data generated by a ring oscillator to an LC tank oscillator, an LC tank is capable of generating an order of magnitude less of random jitter.

The other primary component of the total jitter presented to the optical module is deterministic jitter. The deterministic jitter that can be compensated for by the transmit driver is due to the high-frequency channel losses introduced by the transmission medium itself. To provide a boost to the high frequency signal components and overcome these losses, pre- and post-emphasis circuits can be utilized. Xilinx includes pre- and post-emphasis circuits, both fully programmable, to enable designers to overcome channel losses and maximize jitter performance after the data has been translated into the optical domain.

Figure 1 shows eye diagrams of a 28 Gb/s signal with and without pre-emphasis driven over a short distance of PCB trace, where small adjustments in the pre- and post-emphasis values can have a large impact on the deterministic jitter presented to an optical module.



WP427_01_102512

Figure 1: 28 Gb/s Signal (a) without and (b) with Transmit Equalization

Receiver Compatibility

After data has been translated from the electrical domain and transmitted across the optical channel, it is translated back into the electrical domain, where it is received by the link partner. Depending on the type of module used in the receive path, there are different requirements that the electrical receiver must meet. First, varying amounts of dispersion and jitter introduced to the optical signal by the optical fiber have been transferred onto the electrical signal; this degradation must now be compensated for. Second, for any jitter that cannot be compensated for, the receiver needs to have a high inherent jitter tolerance to be able to receive valid and correct data.

Similar to the transmitter requiring the absolute minimum clock jitter to generate the lowest random jitter, the CDR circuitry in the receiver must start with the lowest jitter clock possible to maximize performance. The same LC tank PLLs that provide the transmitter with class-leading jitter performance in the 7 series GTX, GTH, and GTZ transceivers is used to drive the CDR circuitry to deliver class-leading receive jitter tolerance.

Dispersion or jitter from optical modules can come in a number of different forms. Electronic dispersion compensation (EDC) is commonly used to describe the features required by a receiver to compensate for interfacing with optical modules, but the implementation and extent of dispersion that needs to be compensated for varies widely between re-timed, limiting, and the many types of linear modules.

The most readily received form of electrical signaling comes from re-timed modules. When interfacing with a re-timed module, the same decision feedback equalizers (DFEs) and continuous time linear equalizers used as electrical interconnects provide the required compensation. These circuits are always important because they compensate for the board-level interference between any module and the receiver. The improvement in signal quality provided by re-timed modules enables the highest performing links. The Xilinx 28 Gb/s GTZ transceiver is designed to operate with the re-timed CFP2 optical modules. Furthermore, the GTX and GTH transceivers can interoperate with all of the major re-timed interfaces up to 12.5 Gb/s in the GTX transceiver and 13.1 Gb/s in the GTH transceiver.

When interfacing with limiting or linear modules, additional sources of interference need to be considered. Limiting modules have strict jitter specifications that must be met by both the module and receiver, which necessitates a receiver with strong native jitter tolerance.

Linear modules directly pass on any dispersion that might have been introduced by the optical cables. Linear modules vary widely in their application and, as such, vary widely in their requirements on the electrical receiver. For those interfaces with low dispersion, such as 10GBASE-SR or 10GBASE-LR, no additional circuitry beyond the DFE and CTLE present in the Xilinx GTX and GTH transceivers is required to compensate for the dispersion being passed directly from the optics into the electrical domain. 10GBASE-LRM differs from these standards by requiring the electrical receiver to compensate for large amounts of disparity under varying conditions. SFF-8431 defines an ideal receiver for calibrating characterization efforts, which consists of a combination of a feed-forward and decision-feedback equalizer, a very large circuit that consumes much more power than the standard DFE and CTLE circuits used for electronic compensation.

Conclusion

Optical interfaces are fast becoming a standard part of today's high-bandwidth landscape. Designers using optics are faced with a wide variety of choices of implementation rates, form factors, optical standards, electrical standards and the trade-offs between them. Knowing how the features of a given optical module impact the rest of the system, and how the host ICs on either side of the link can contribute to error free performance, is an advantage that cannot be ignored. From 10G SFP+ modules to 100G CFP2 modules, Xilinx is dedicated to supporting a wide variety of optical interfaces at today's 10Gb/s standards and beyond.

For more information on how to implement optical or other high-speed serial interfaces, go to: www.xilinx.com/products/technology/high-speed-serial/.

Additional Sources

1. SFF Committee SFF-8431, *Specification for Enhanced Small Form Factor Pluggable Module SFP+*, Revision 4.1, July 6th, 2009: <ftp://ftp.seagate.com/sff/SFF-8431.PDF>

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
12/20/12	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>.