


This document discusses how optical interface technology embedded in an FPGA overcomes the reach, power, port density, cost, and circuit board complexity challenges associated with discrete copper interconnects. As data rates exceed 10 Gbps for various reach distances ranging from chip-to-chip, chip-to-module, rack-to-rack, and system-to-system interfaces, this technology enables designers to overcome those challenges and provides significant advantages over conventional discrete electrical or optical technologies.

Introduction

With today’s high-bandwidth, low-latency devices and associated applications—such as smart phones, tablets, HDTV, and 3DTV—computer and network system vendors endeavor to deliver systems that don’t significantly add to network or internet traffic congestion and latency. The sources of I/O data rate and density increases include servers, local area network (LAN) routers and switches, storage area network (SAN) switches and RAIDs, wide area network (WAN) optical switches, and transport systems. For example, the per lane data rate for the PCIe® interface has increased from 2.5 Gbps for Gen 1.0, to 8.0 Gbps for the current Gen 3.0, and is expected to increase to 16 Gbps by Gen 4.0. **Figure 1** shows the internet and IP traffic demand trend.

Figure 1. Internet bandwidth and IP Traffic Trend (CISCO VNI, 2010)




FPGAs play an important role in today's network, computer, data center, and communication ecosystems. Altera’s latest optical interface on FPGA overcomes the limits of copper interconnect by integrating the latest FPGA with a state of the art laser and photon detector at the FPGA package level. The optical FPGA interface provides reach-length, power, cost, density, and form-factor advantages that dramatically exceed conventional electrical signaling and interconnect and discrete electrical-optical signaling and interconnect capabilities.

Limits of Copper Interconnect


As they implement new systems, network and data center operators wish to avoid adding power and cost penalties compared to previous generation systems. There is a desire for power and data transfer associated with the merit of pJ/bit should progressively decrease with each generation. Figure 2 shows the data rate per lane trend from International Technology Roadmap for Semiconductors (ITRS, 2009 revision).

Figure 2. ITRS Projected High-Speed I/O Data Rate Trend



Designers widely use copper interconnect for chip-to-chip and chip-to-module interfaces over traces on a printed circuit board (PCB), in chip-to-chip over backplane, and in chip-to-chip over copper cable assemblies. At 10 Gbps, the reach distances are approximately 0.3 m for chip-to-chip and chip-to-module interfaces, 1 m for chip-to-chip over backplane, and 7 m for chip-to-chip over copper cable assemblies.

The challenge for copper based interconnect is that it does not scale with the data rate because of the frequency dependent loss. For example, in the widely used FR-4 copper trace material the loss is ~ 0.5-1.5 dB/in at 5 GHz (Nyquist for 10 Gbps rate), and the loss increases to ~ 2.0-3.0 dB/in at 12.5 GHz (Nyquist for 25 Gbps rate). Return loss and crosstalk can also increase with frequency. Figure 3 shows an example of insertion loss, return loss, and crosstalk for a PCI Express server channel. This example illustrates how the insertion loss and return loss increases with respect to the data rate. At 8 GHz (16 Gbps), the insertion loss is ~ -60 dB, which is far beyond the equalization dynamic range for a NRZ signaling (~40 dB). This example shows that the current technology cannot scale to 16 Gbps unless the channel length is reduced, which limits the application and flexibility. Furthermore, this method requires additional components, such as a repeater, which adds more power, cost, and complexity.

Figure 3. Insertion Loss, Return Loss, and Crosstalk for PCI Express Server Channel

In these copper-based systems, designers typically must compensate for insertion loss signal impairments, such as inter-symbol interference (ISI) or data-dependent jitter (DDJ), return-loss, and crosstalk. Designers adjust for these impairments by using various equalizers, such as a feed-forward equalizer (FFE), continuous time linear equalizer (CTLE), or decision feedback equalizer (DFE), implemented on the transmitter or receiver at the copper channel to ensure that the link performance (that is, bit error rate (BER) $< 10^{-12}$) is met. However, equalizers consume power and add penalties, especially the DFE. As the data rate increases, insertion loss, return loss, and crosstalk also increase and require even stronger equalizers (that is, more taps or larger DC/AC gains) to compensate for the resulting impairments, and to insure the same performance. This technique in turn, adds more power.

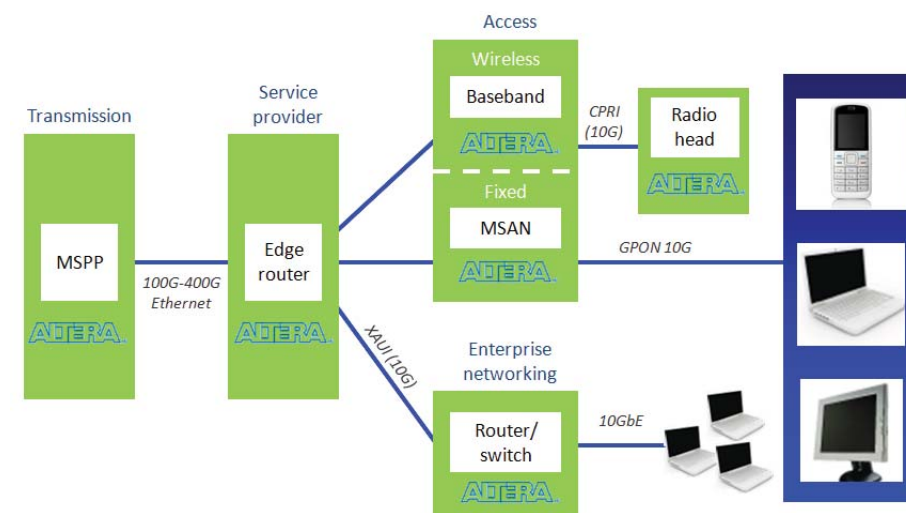
In today's interconnect ecosystem, designers use copper electrical signals mostly for reach distances up to 10 m, and fiber optic signal mostly for distances from 10 m and beyond, because of the characteristics and cost structure of copper electrical components and fiber optical components. Designers would like to apply optical fiber signaling to distances < 10 m to overcome the loss, signal integrity, and power challenges of copper electrical signaling. As the data rate increases, this goal is challenged by the discreteness of electric and optical components and associated cost and power.

Advantages of Optical FPGA Interfaces

In contrast with copper interfaces, optical fiber has virtually no loss. A multiple mode fiber (MMF) has a loss of ~ 3 dB/km and ~ 1 dB/km at 850-nm and 1300-nm wavelengths, respectively. A single mode fiber (SMF) has a loss of ~ 0.4 dB/km and 0.25 dB/km at 1300-nm and 1550-nm wavelengths respectively. MMF is less expensive due to its larger core (~ 50 micron) and has a bandwidth ~ 2 GHz km; while SMF is more expensive due to its smaller core (~ 9 micron) and has a bandwidth close 100 THz in practice. The laser that drives the optical signal over an MMF is commonly a light emitting diode (LED) or Vertical Cavity Surface Emitting Laser (VCSEL). The MMF is commonly used for reach distances of < 1 km, while SMF is used for reach distances of > 1 km to a few thousand km. At 10 Gbps, the reach distance for a MMF is ~ 300 m. Unlike the copper electrical link, power consumption and penalty of an optical link is relatively independent of reach length. Moreover, unlike an electrical signal, an optical signal is immune to electric-magnetic interference (EMI) and has no amplitude crosstalk, providing better signal integrity resilience. With the wavelength division multiplexing (WDM), multiple channels can be supported with the same optical fiber, enabling channel material savings.

FPGAs play an important role in today's network, computer, data center, and communication ecosystems. FPGAs offer critical reconfiguration and system on chip (SoC) capabilities for data processing and transport, and for providing computation, DSP, packet processing, frame processing, routing, switching, and bridging MAC/FEC capabilities. Figure 4 shows the I/O links and data processing of a network ecosystem for chip-to-chip, chip-to-module, and system-to-system provided by Altera's FPGA and advanced transceiver circuits.

Figure 4. Altera FPGA Crossing an Internet/Network Ecosystem



Altera Advanced Transceiver Technology

Altera's transceiver technology has evolved extensively over a decade, and provides electrical transmit and receive functionality with current data rates up to 28 Gbps on the 28-nm process node. These transceivers also support advanced clock generation, clock recovery, and equalization capabilities (that is, FFE, CTLE, and DFE). Figure 5 shows the FPGA transceiver advancement path with respect to the data rate and process node. This advanced transceiver technology integrates the FPGA with an optical laser and the receiver becomes seamless.

Figure 5. Altera FPGA Transceiver Technology Advancement Path

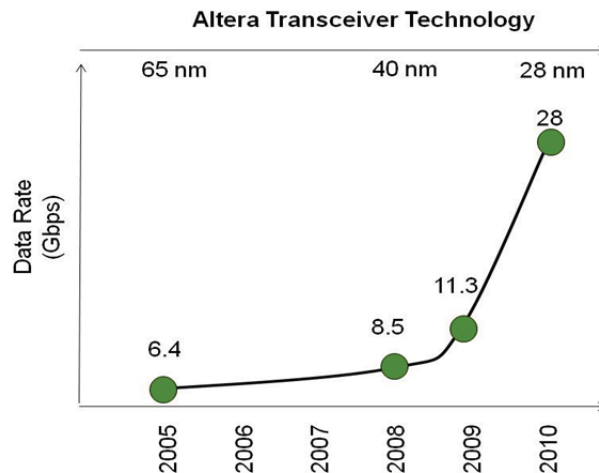



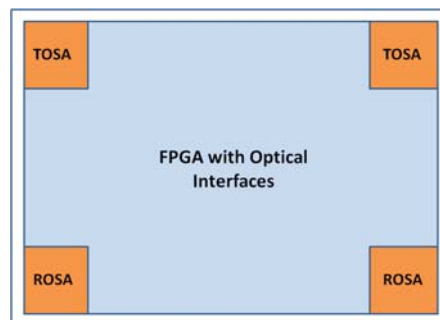
Figure 6 illustrates Altera's advanced transceiver capabilities and performance. On the transmitter (TX) side, jitter generation is extremely low and reaches ~300 fs or lower at 28 Gbps due to the use of advanced LC oscillator. TX FFE equalization can have up to four programmable taps. On the receiver (RX) side, there are multiple stages CTLE, with up to 12 dB DC gain and 20 dB AC gain, adjustable peaking frequency, and selectable automatic equalization coefficient finding. The RX DFE can have up to 5-taps that can compensate for most of the uncorrelated jitter and noise. Each tap weight is power optimized with its clock recovery (CR), based on a hybrid architecture that can lock to either the data or reference clock. These conditions produce excellent locking time/range and resilience of excessive jitter on the incoming data. The transceiver has a good run-length handling (~ a few hundred UI), bandwidth (~ a few hundred MHz), and jitter tolerance slope (better than -40 dB/decade). The transceiver has a built-in on-die instrumentation (ODI) that can measure BER contour and eye-diagram.

Figure 6. Altera Advanced FPGA Transceiver Capabilities and Performance

Altera's advanced ODI enables TX adaptive equalization, RX optimal sampling point finding and associated BER reduction. In addition, the transceivers support fast board and system bring up, debugging, non-intrusive link characterization, and life traffic and performance monitoring, all in a cost and power efficient way. A "closed" signal integrity "eye" can be easily "opened" by RX CTLE and/or DFE, and verified with ODI, as shown in Figure 6. At 28 nm, Altera's transceiver achieves single digit (~8 pJ/bit) power efficiency.

FPGA Optical Interface Application

With the advanced transceiver, integrating an Altera FPGA with optics becomes simple. Figure 7 shows a example of an FPGA with optical interfaces. The FPGA in Figure 7 is integrated with optics, such as a transmitter optical sub-assembly (TOSA) and receiver optical sub-assembly (ROSA), providing direct optical signal transmitting and receiving without the need for a discrete optical module.

Figure 7. FPGA with Optical Interface

Designers can replace chip-to-module link systems with an optical interface FPGA to capture significant power, resource, and cost reductions, as well as an increase in port density, as illustrated in Figure 8. This is possible because Altera's advanced and power efficient transceiver has best-in-class clock generation and clock recovery, signaling conditioning and equalization, and built-in test and measurement via ODI.

Figure 8. Chip-To-Module link Replaced by Single FPGA with Optical Interface

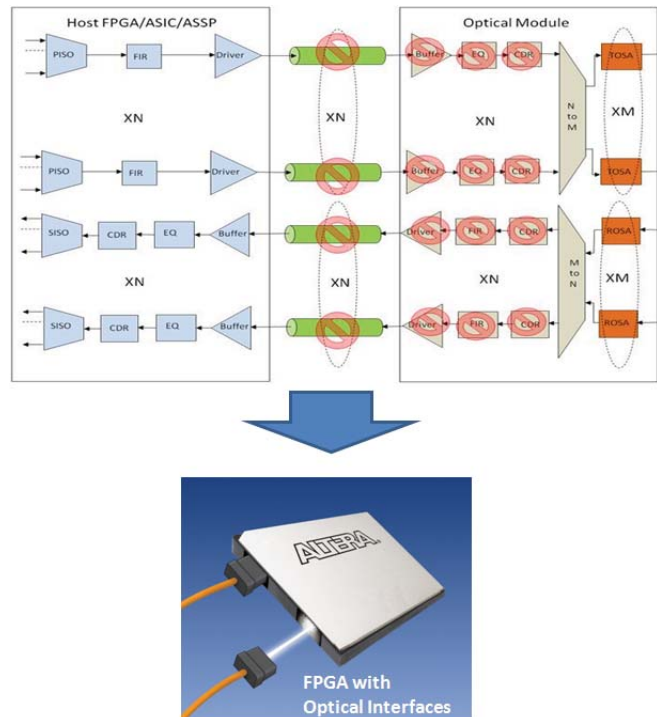



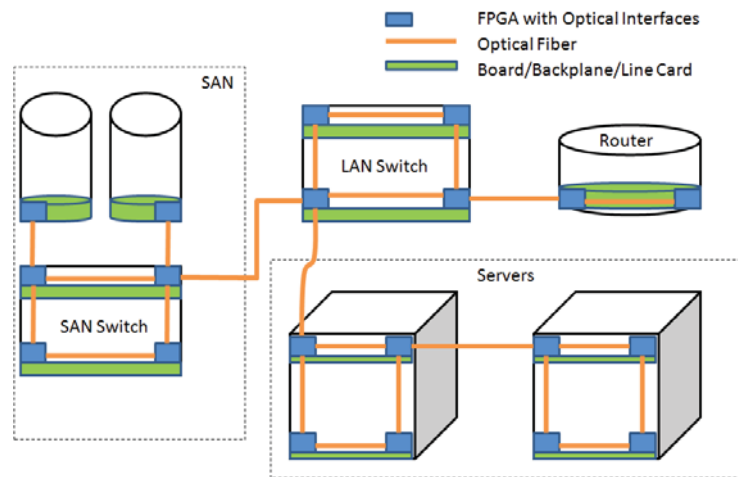
Figure 9 shows a generic application that includes an optical interface to the FPGA. Designers can use an FPGA with optical interface for an optical backplane, board-to-board, rack-to-rack, and system-to-system interconnect, with reach length up to 100 m at 10 Gbps line-rate. The maximum data rate for channel depends on the FPGA transceiver data rate, and is currently at 28 Gbps on the 28-nm process node.

Figure 9. Optical FPGA applications



A more detailed example for the use of an FPGA with optical interface in a data center (DC) is illustrated in Figure 10. This illustration shows the intranet board-backplane-line card, board-to-board, rack-to-rack, and system-to-system interconnects using this new FPGA as LAN switch, router, SAN switch and disk array, and server array. The FPGA with optical interface enables processing, as well as optical interconnects, for distances in the range of < 0.3 m to > 100 m, and is well suited for the entire DC interconnects. This provides significant power, density, and cost saving advantages compared with conventional technologies. These capabilities were not possible before this technology, especially for reach distance < 10 m.

Figure 10. FPGA with Optical Interfaces for a Data Center



Summary

Reach length, power, cost, board material, and circuit board complexity are key challenges for copper based, chip-to-chip interfaces. Similarly, cost, port density, power, form-factor, and board material are the primary challenges for host chip plus discrete optical module solutions currently used. As data rates approach 10 Gbps and higher, Altera's optical interface technology integrates the latest FPGA with state of the art laser and photon detector, overcoming challenges for both scenarios. An FPGA with optical interface can provide reach-length, power, cost, density, form-factor, and weight benefits compared with conventional electrical signaling and interconnect and electrical-optical signaling and interconnect. Designers can use the FPGA optical interface for backplane short reach chip-to-chip (~ or > 1 m), board-to-board, rack-to-rack, and system-to-system interconnect for a reach distance up to ~100 m.

It is easy to imagine that FPGAs with optical interface technology will be widely used in future datacom and telecom systems, data center, data transporting, and military network systems, as well as test and measurement, medical, and broadcasting systems due to reach-length, power, cost, density, and form-factor, weight, and EMI and crosstalk resilience benefits. This technology forever changes how future internet, network, data center, test and measurement, medical, and broadcast ecosystems will be designed and built.

Acknowledgements

- Dr. Mike Peng Li, Principal Architect, Product Engineering, Altera Corporation

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
April 2011	1.1	Corrected minor typos.
March 2011	1.0	Initial release.