



Scaling 100G Wired Applications with Heterogeneous 3D FPGAs

By: Ehab Mohsen

To address the insatiable demand for bandwidth, the communications industry is accelerating development of Nx100G line cards for networking systems. In order for equipment manufacturers to scale infrastructure economically and effectively, they must leverage the latest optical interconnect technologies such as CFP2, and in the future CFP4, to increase bandwidth while lowering power and cost.

By working with network developers, Xilinx anticipated this need and developed transceiver-rich, high-performance, programmable devices comprised of heterogeneous silicon die. The technology supports the required 28 gigabits per second (Gb/s) channels for CFP2 optics and delivers optimal signal integrity due to its heterogeneous architecture. With high logic capacity and specific IP for communications applications, the latest Xilinx devices provide extensive levels of system integration to usher in the migration to next-generation optics.

The Demand for Bandwidth

Largely driven by streaming video, HD video, cloud computing, and mobile networking, the consumer market's relentless demand for network bandwidth compels the communications industry to double system capacity every three years. Service providers supporting the Internet's backbone must lead the migration to 100G and stay at the forefront of the latest technologies and standards.

Service providers not only demand more bandwidth but aim to reduce capital and operating expenses. For equipment manufacturers, this means rolling out solutions with leaps in performance, area efficiency, and cost effectiveness over previous generation products.

The Move to Next-Generation Optics for Nx100G

Most of today's network infrastructure is connected via optical fiber, hence the bandwidth and cost of optical interfaces are major development considerations. The type of optical modules that can be used depends on the architecture of the application's line cards. Three well known optical module standards include SFP+, CFP, and CFP2—each with varying throughput, cost per bit, power efficiency, and form factor:

- Simple form factor pluggable (SFP+) optical modules support 10G optical links and are currently shipping in high volume.
- Compact form factor pluggable (CFP) modules, also in production, support 100G optical links. Though they consume more power per bit than SFP+, integration to a single 100G fiber greatly reduces complexity and serviceability costs.
- The CFP2 optical module offers the same 100G bandwidth as a CFP, but in half the space, using half the power, and at a reduced cost.

Because of the 2X bandwidth-per-watt efficiency gained from CFP2 modules over CFP, the industry is eager to move to these optics. Without this technology, the cost of migrating to 100G is prohibitive for many service providers. The need for CFP2 is demonstrated in [Figure 1](#), showing a comparison of optical interfaces as they appear on the faceplate connector of a fixed-width line card. Because service providers postpone upgrading their chassis until economically feasible, network OEMs must strive to provide more capabilities within the same unit area and power envelope. The comparison in [Figure 1](#) shows that scaling bandwidth within existing infrastructure is driven by throughput per watt per unit area of optical ports.

When using SFP+ optical modules to connect 10G optical links, the top faceplate connector shown in [Figure 1](#) can accommodate 48 fiber links. The arrangement in this example provides 480 Gb/s of throughput in a 48W power envelope.

Comparatively, four CFP ports can be designed in the same footprint of 48 SFP+ modules. With each CFP accommodating a single 100G fiber link, this provides a total of 400 Gb/s of bandwidth. Though there is a slight increase in power, the integration reduces complexity and serviceability.

A CFP2 module, by contrast, provides the same 100G bandwidth of a CFP in half the width while consuming half the power per 100G port. In this example, within the same area, a module could accommodate eight CFP2 ports for an aggregate 800 Gb/s bandwidth within the same 60W power envelope. This is 33% higher bandwidth and power efficiency compared to SFP+ and double the efficiency provided by CFP modules.

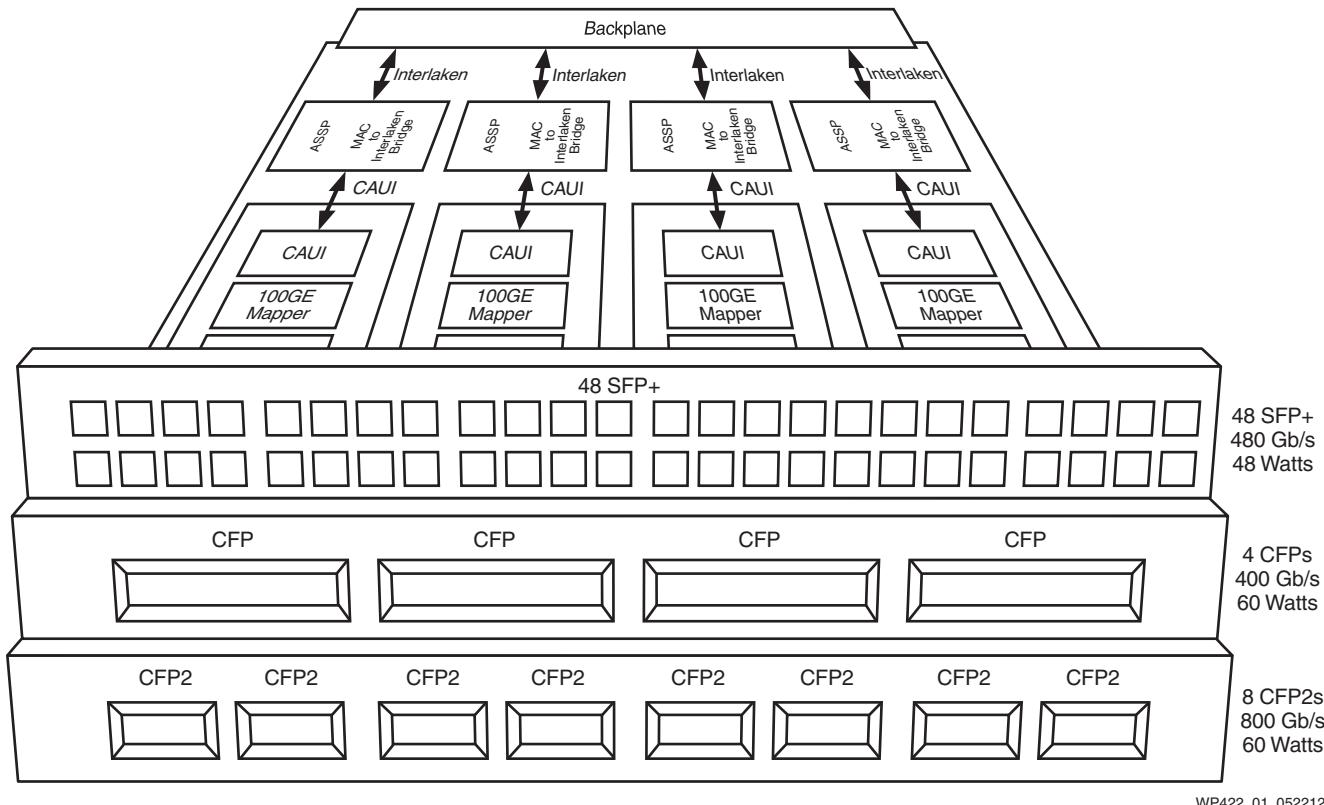


Figure 1: Throughput and Power for Line Cards and Face Plate Connectors of Fixed Width

The Challenge of Redesigning the Line Card for Nx100G

Migrating to CFP2 has its benefits, but the need for higher density front plates poses challenges on the line card itself. Effective integration is needed on the silicon side to support the incoming bandwidth so as not to nullify the power and cost efficiencies promised by a CFP2 transition.

A typical 100G transponder is shown in Figure 2, with an optical interface at one end and a backplane interface at the other. Typically, there is a forward error correction (FEC) block to minimize packet retransmission and framing and mapping functions to handle data transport. Transceiver interfaces such as CAUI are used for chip-to-chip communication, and Interlaken can be used for chip-to-chip or backplane communication.

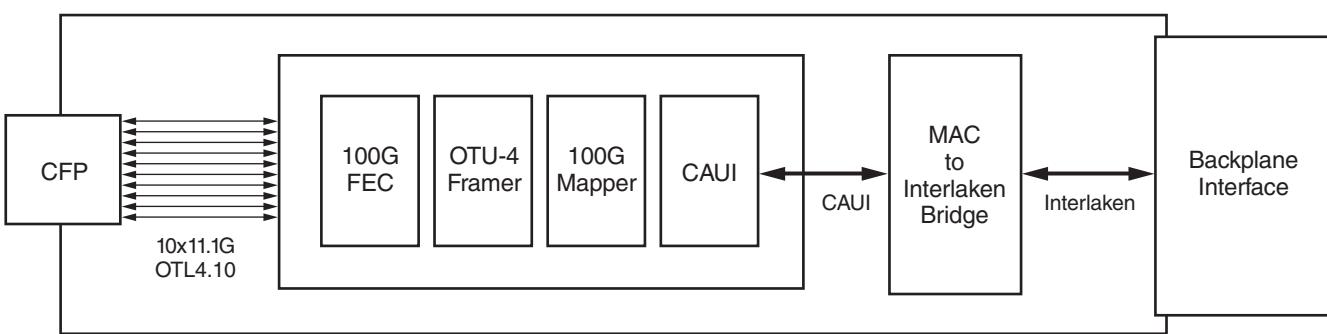
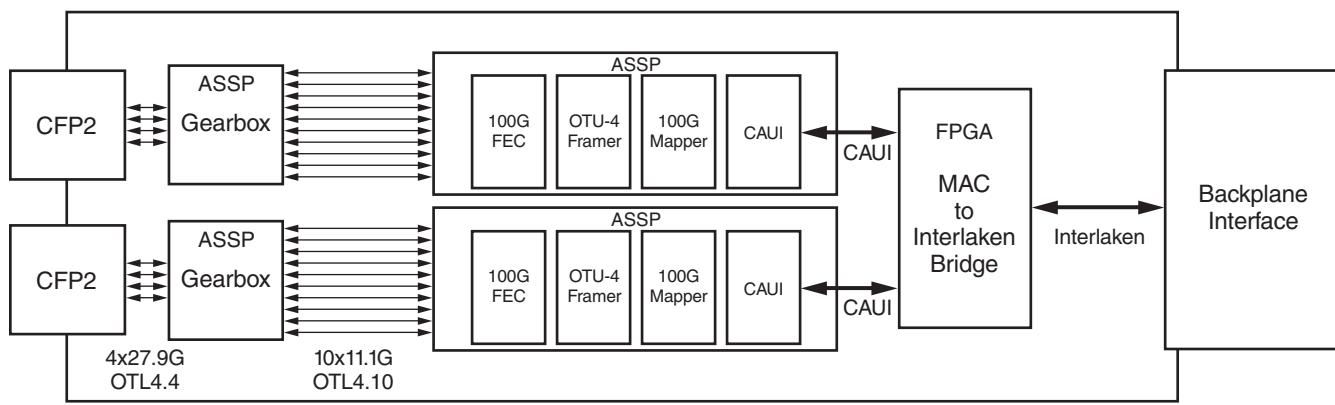


Figure 2: Generic 100G Transponder Line Card

To redesign the line card for 2X bandwidth, the interface to CPF2 must first be considered, given that it is based on 4x25G channels versus the 10x10G channels supported for CFP.

Through advanced multiplexing, a functional block known as a "gearbox" can convert a 100G interface comprised of 4x25G channels into 10x10G channels, allowing these modules to interface with existing silicon infrastructure. Consequently, the original devices (ASICs or ASSPs) that operate via 10x10G do not necessarily need to be replaced to support CFP2. The gearbox maps data between the ten and four serial lane interfaces, in both ingress and egress directions. It converts data streams of either four lanes of CAUI4 (4x 25.78G) or OTL4.4 (4x 27.95G) to CAUI (10x 10.3125G) or OTL4.10 (10x 11.18G).

Although the gearbox addresses optics connectivity, it still does not address the 2X bandwidth. If the CFP is replaced by two CFP2 modules, the system either has to support additional components of similar type within the same area or support completely new silicon to support 2x100G throughput. Migrating to new ASSP and NPU architectures can be prohibitive in terms of cost and schedule, and a new implementation using similar components has its own challenges. A re-design of the line card to support 2x100G using similar components and gearbox ASSPs is shown in [Figure 3](#). The increased number of components requires more area on the PCB. Even if such a layout is feasible, the increase in cost and power consumption can nullify the advantages of a CFP2 migration. Other alternatives are needed.



WP422_03_052412

Figure 3: Five ASSPs Required to Redesign Transponder to Support CFP2 and 2x100G

28G Enabled FPGAs as a Solution

FPGAs play a critical role in networking equipment because of their flexibility and ability to rapidly implement the latest networking standards, even as these standards evolve. FPGAs have also evolved to meet next-generation networking requirements by delivering greater capacity, performance, and features, along with more robust transceivers supporting higher line rates.

To interface to CFP2 modules, FPGAs must provide 25G–28G serial interfaces with support for advanced protocols and interface specifications. These include 100GE, OTU4, 400GE, CAUI, CAUI4, OTL4.4, SFI-S and other standards. A line card without 28G support simply cannot interface to CFP2 optics.

Transceiver support is only half the challenge for successful 28G design. Signal integrity is another consideration at this transmission rate. The CEI-28G specification

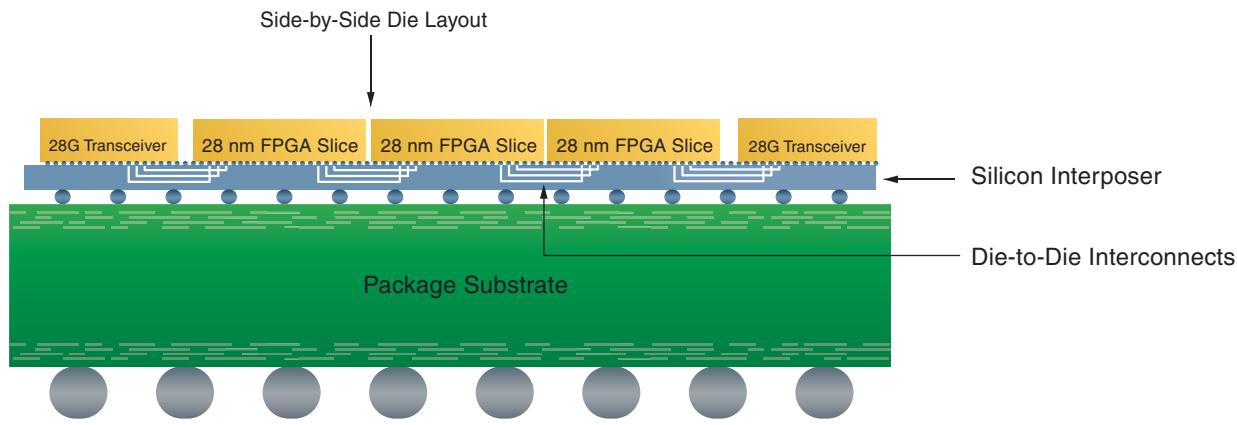
guiding the electrical specifications for 28G imposes very tight transmit jitter budgets (0.30 UI) on system designers and requires robust equalization techniques in the receiver to build 28G chip-to-optics interfaces.

An FPGA Tailored for Wired Communications Applications

Xilinx Virtex®-7HT FPGAs were designed to match these unique requirements, addressing the bandwidth needs, signal integrity challenges, and integration demands. As a single chip solution enabling Nx100G applications, the Virtex-7 HT FPGA ushers in the transition to CFP2 optical modules.

Based on Stacked Silicon Interconnect Technology

The Virtex-7 family is based on 3D Stacked Silicon Interconnect (SSI) technology, which combines enhanced FPGA die slices and a passive silicon interposer to create a three dimensional die stack. This interposer implements tens of thousands of die-to-die connections to provide ultra-high inter-die bandwidth with lower power consumption and one fifth the latency of standard I/Os. The Virtex-7 HT device shown in [Figure 4](#) ties together three Super Logic Regions (SLRs) fabricated on 28 nm. Next to these SLRs are separate 28G transceiver die. Because of its 3D nature, SSI technology has allowed the Virtex-7 family to outpace Moore's law in performance, capacity, and power efficiency. For more information, refer to [WP380](#), *Xilinx Stacked Silicon Interconnect Technology Delivers Breakthrough FPGA Capacity, Bandwidth, and Power Efficiency*.



WP422_04_052412

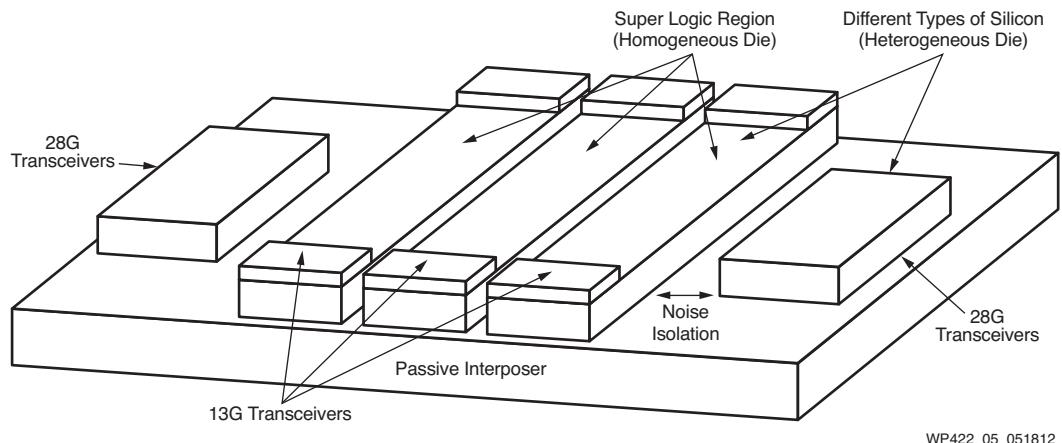
Figure 4: Xilinx Stacked Silicon Interconnect Technology (Side View)

Heterogeneous Silicon for Low Jitter and Noise Isolation

The Virtex-7 HT FPGA uses SSI technology with traditional FPGA SLR slices and 28 Gb/s transceiver slices to create the world's first heterogeneous device.

Xilinx employs a unique approach to isolate the digital logic from the analog transceiver circuit on the same interposer, as shown in [Figure 5](#)—in essence, placing heterogeneous die side-by-side to operate as one integrated device. If this were a monolithic device, the approach of competing solutions, the digital logic region would create a noisy environment that degrades transceiver performance. The electrical isolation of the digital and analog circuits in a heterogeneous device allows for low noise and jitter. The noise isolation architecture provides a considerable reduction in

noise coupling compared to traditional approaches. This simplifies the job of I/O specialists, accelerates 28G design closure, and reduces board cost.



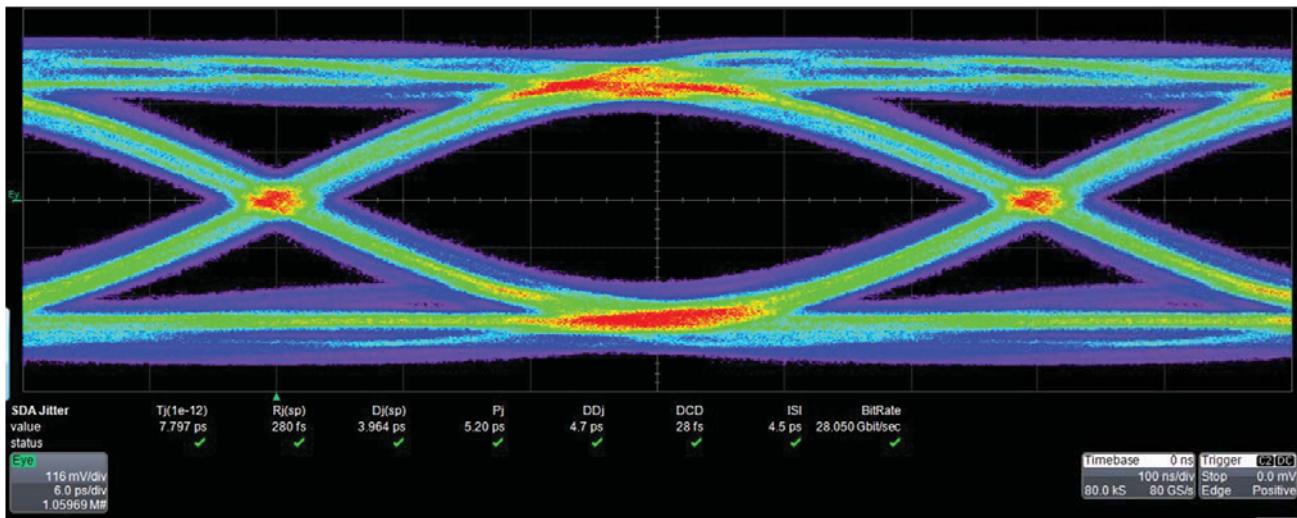
WP422_05_051812

Figure 5: Heterogeneous 3D FPGA Enables Low Jitter 28G Transceiver Design

In addition to noise isolation, the transceiver's jitter performance is improved with a narrowly tuned phase-locked loop (PLL) based on an LC tank design. Unique clocking, clock distribution, and PLL design minimizes jitter across multiple transceivers. Additional design features minimize lane-to-lane skew to support tough optical standards like the Scalable SerDes Framer Interface (SFI-S), which limits acceptable skew to 500 ps.

To compensate for channel loss and maintain signal integrity, Xilinx 28G transceivers employ a programmable main transmit driver, programmable transmit pre-emphasis, and an auto adapting continuous time linear equalizer (CTLE) in the receiver.

The eye diagram in Figure 6 demonstrates the low jitter and high signal quality of the 28G FPGA transceiver on the Virtex-7 XC7VH580T device. The 28G transceiver presents an open eye without excessive over-equalization.



WP422_06_052212

Figure 6: 28 Gb/s Eye Diagram of GTZ Transceiver on the Virtex-7 H580T FPGA

Highest Bandwidth FPGA with up to 88 Transceivers

The heterogeneous architecture also enables an unprecedented number of transceivers on Virtex-7 HT devices. The devices offer two types of serial transceivers:

- GTH transceivers operate up to 13.1 Gb/s and support optical, chip-to-chip, and backplane connectivity.
- GTZ transceivers deliver up to 28.05 Gb/s for 100G optical networking.

The Virtex-7 XC7VH870T device offers up to sixteen 28G transceivers—4X the competition—making it uniquely matched to interface to up to four CFP2 modules for 4x100G applications or 400 Gigabit Ethernet. With an additional seventy-two 13.1 Gb/s GTH transceivers on the same device, system designers have 6G or 10G Interlaken connectivity for high speed chip-to-chip connections. With up to 88 transceivers overall, the Virtex-7 HT device is the highest bandwidth FPGA available, providing 2.87 terabits per second (Tb/s) of bidirectional throughput.

Gearbox IP to Enable CFP2

As important as CFP2 connectivity is, the intellectual property (IP) cores needed for the line cards are equally critical. Xilinx provides gearbox IP that handles 4x25G to 10x10G conversion. It also supports a 10x10G-to-10x10G pass-through mode for designs that need to support both CFP and CFP2.

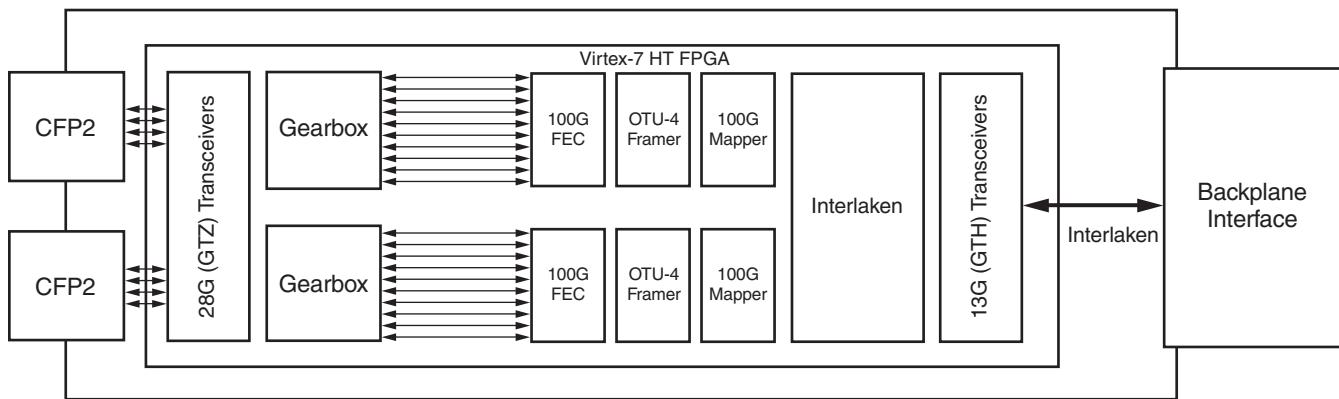
While a single Virtex-7 HT FPGA can provide up to 4x100G throughput via its gearbox IP connecting to up to four CFP2 ports, competing FPGA or ASSP solutions can provide only 100G throughput on a single device. As already shown in [Figure 3](#), if an ASSP approach were taken to upgrade the line card, separate gearbox chips would be needed, thereby increasing cost, power consumption, and board complexity.

Programmable Silicon Integration for a Single-Chip Solution

The other benefit of using an FPGA is flexibility when integrating IP. With Virtex-7 HT devices, designers can take integration to the next level by combining gearbox, Ethernet MAC, OTN transponder, OTN muxponder, Interlaken, differentiating IP, and chip-to-chip protocol interfaces such as CAUI, CAUI-4, SFI-S, and RXAUI—all on a single device.

A Comparison of Two Transponders

A 2x100G line card is shown in [Figure 7](#), where a Virtex-7 HT device is used to integrate two gearboxes, a MAC, and Interlaken bridge. The simplicity of this architecture is in stark contrast to the 5-chip alternative in [Figure 3](#). A designer can implement four of these Virtex-7 HT devices, producing an 8x100G system that can interface with eight CFP2 ports. An equivalent scaling of the ASSP implementation would require 20 devices—consuming excessive area, increasing PCB cost and power, and likely lengthening the project schedule.



WP422_07_052412

Figure 7: 2x100G Transponder Using a Single Virtex-7 FPGAs

Based on Xilinx estimates of pricing and power consumption of ASSPs advertised on the market today, a power-and-cost comparison of the two line card implementations is shown in [Table 1](#). The ASSP-based solution is comprised of five devices, consumes at least 40% additional power and costs 50% more than the FPGA implementation. Unaccounted for is the productivity and time-to-market gain from the simplified layout and integration of a single FPGA.

Table 1: Power/Cost of ASSP Implementation vs. FPGA-Based 8x100G Line Card

ASSP Implementation Costs (vs. Xilinx HT FPGAs)	
Power Consumption	40% or greater
BOM Cost	50% or greater

Logic resources and types of applications for Virtex-7 HT devices are shown in [Table 2](#). Due to its capacity and number of 28G transceivers, the Virtex-7 XC7VH870T is currently the only silicon on the market capable of implementing a 400 Gigabit Ethernet solution.

Table 2: Virtex-7 HT Family and Key Types of Applications

	XC7VH290T	XC7VH580T	XC7VH870T
Logic Cells	284,000	580,480	876,160
GTH Transceivers (13G)	24	48	72
GTZ Transceivers (28G)	8	8	16
Types of Applications	2x100G	2x100G	400G

Enabling CFP2 Connectivity and Beyond

The market need for higher-bandwidth networking line cards and next-generation optics is real. Xilinx is at the forefront of this movement with a heterogeneous architecture that provides the bandwidth and capacity for adopters of 100G, based on CFP2 optics. Without an FPGA solution of this caliber, the migration would not only be sub-optimal, but costly. By leveraging FPGAs, designers get a level of integration that is two-fold: optical connectivity at the system level and IP integration at the silicon level. By targeting Virtex-7 HT devices, designers achieve the greatest possible port density, protect themselves against evolving standards, and prepare themselves for optics even beyond CFP2.

Additional Information

To learn more about the Virtex-7 HT devices, go to the 28 Gb/s Serial Transceiver Technology webpage on xilinx.com:

<http://www.xilinx.com/technology/roadmap/28g-serial-transceiver-technology.htm>

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
05/29/12	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>.