WP385 (v1.1) November 22, 2010

Industry's Highest Bandwidth FPGA Enables World's First Single-FPGA Solution for 400G Communications Line Cards

By: Greg Lara

To address the insatiable demand for more bandwidth, the telecommunications industry is accelerating development of Nx100G and 400G line cards for communication systems. These systems will leverage new optical interconnect standards to address power density and thermal management issues attendant with scaling the capacity of existing infrastructure. A critical component of the solution is 28 Gb/s electrical interfaces that increase port density and reduce power per bit.

Xilinx is responding to the demand for more bandwidth with two key developments. The first is high-fidelity 28 Gb/s transceiver technology that meets the stringent requirements of next-generation communication systems. The second is 28 nm Virtex®-7 HT FPGAs that integrate an unprecedented 16 x 28 Gb/s and 72 x 13.1 Gb/s transceivers with logic, memory, and I/O resources that enable the first silicon device (FPGA or otherwise) to support 400G line cards and the industry's largest single-FPGA solution for Nx100G line cards.

Scaling Network Bandwidth

Exploding Bandwidth Demand

The demand for bandwidth is growing at an astonishing rate. Cisco Systems forecasts IP traffic to grow to 64 exabytes (64 x 1,018 bytes) by 2014, driven largely by consumer internet traffic, especially the popularity of streaming/downloaded video (Figure 1).

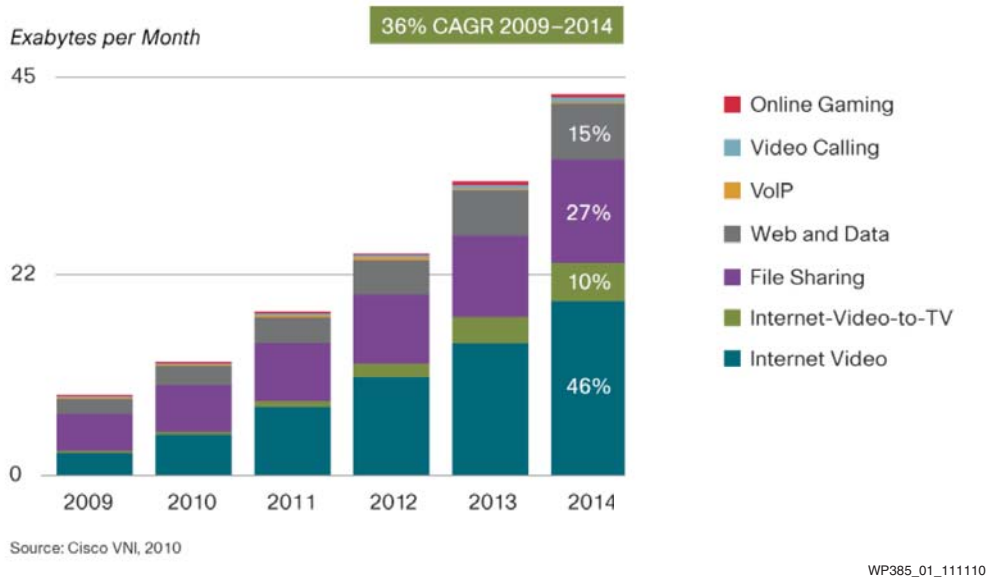

Figure 1: **Projected IP Bandwidth Demand**

In response to the demand for more bandwidth, the telecommunications industry roadmap calls for doubling system capacity every three years: 480 Gb/s systems employing 10 Gb/s Backplane Ethernet (10 GbE) and Optical channel Transport Unit 2 (OTU2) interfaces were introduced in 2008; terabit systems supporting 960 Gb/s with 40/100 GbE and OTU3/OTU4 protocols are now appearing; and the groundwork is being laid today for the multi-terabit (1.9 Tb/s) systems that are expected to appear in the 2014 time frame.

Power Density Management and Next Generation Optics

Examining the physical aspects of hardware design shows the challenges of scaling bandwidth within the constraints of existing infrastructure. The primary factor limiting the ability to achieve higher bandwidth is the increasing power density. This design challenge becomes apparent when considering the dimensions of the faceplate containing the optical modules that enable the optical fiber links (Figure 2). When using small form-factor pluggable (SFP+) optical modules with 15 mm pitch to connect 10G optical links, a single port blade can accommodate 48 fiber links. This arrangement provides 480 Gb/s of throughput, with 48W of power consumed by the optical modules.

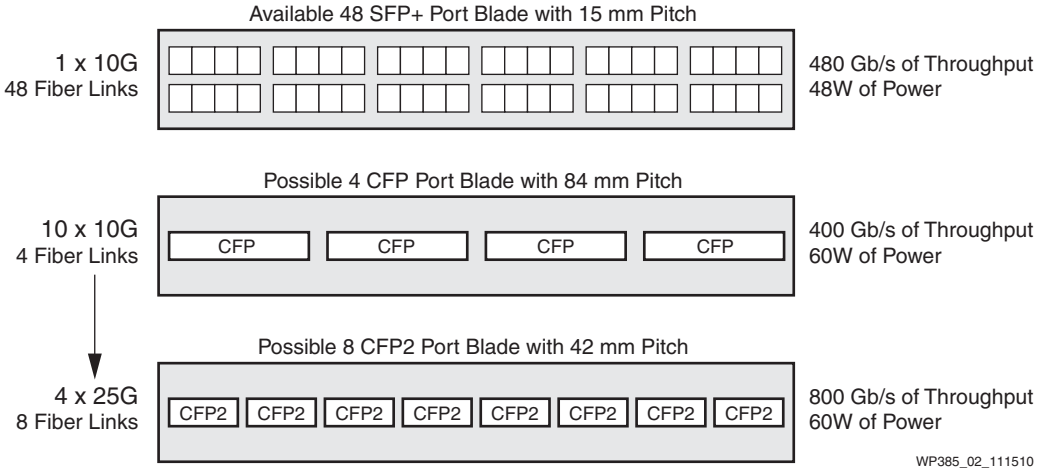

Figure 2: **Evolution of Port Blade Design**

An alternative approach uses CFP modules, which accommodate ten fiber links of 10G to provide an aggregate 100 Gb/s bandwidth. It is possible to design a blade with four CFP ports for a total of 400 Gb/s of bandwidth. Unfortunately, this approach provides lower bandwidth than the SFP+ implementation (400 Gb/s vs. 480 Gb/s) and consumes more power (60W vs. 48W). In the race to deliver higher bandwidth with lower power, the CFP approach does not achieve the goal.

The emerging CFP2 module specification addresses these drawbacks. The main difference between CFP2 and the earlier CFP module is the relocation of the "gearbox" function (which converts ten 10.3125G data streams into four 25G data streams) from within the optics to an external component, such as a discrete gearbox or the framer/MAC FPGA. The resulting optics form factor is smaller, and power consumption is lower. A single CFP2 module provides 100 Gb/s of bandwidth via four 25G fiber links in half the space of a CFP module (42 mm pitch vs. 84 mm pitch). Thus, the dimensions of a single blade can accommodate eight CFP2 ports for an aggregate 800 Gb/s bandwidth with the same 60W of power consumed by the CFP interface. This is 33% higher bandwidth/power efficiency compared to SFP+ and double the efficiency provided by CFP modules.

The Industry's Lowest-Jitter 28G Transceivers

Design Challenges and Solutions

The CEI-28G specification guiding the development of 28G networking imposes extremely tight jitter budgets (0.30 UI) on system designers building 28G chip-to-optics interfaces. Xilinx® GTZ 28 Gb/s transceivers are designed to deliver superior jitter performance. This begins with the phase-locked loop (PLL), which is based on an LC tank design. To ensure maximum design flexibility along with best jitter performance, Virtex-7 HT FPGAs employ a modular, multirate transceiver architecture that provides improved noise immunity. 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 Scalable SerDes Framer Interface (SFI-S) that limit acceptable skew to 500 ps.

Creating a single transceiver design that can support line rates spanning the range from 19.6 Gb/s to 28.05 Gb/s while meeting demanding jitter requirements is a

difficult design challenge. The Xilinx solution for achieving flexible multirate operation involves a unique design for the voltage-controlled oscillator (VCO) that drives the PLL.

Power system noise is another challenge that designers must conquer to ensure signal integrity. With hundreds of thousands of look-up tables and flip-flops driven by dozens of clock networks running at speeds up to 600 MHz and higher, the FPGA logic array creates a noisy environment that degrades transceiver jitter performance. To achieve the noise coupling reduction required to enable an FPGA with 16 x 28G and 72 x 13G transceivers, Xilinx employs a unique design approach to isolating the digital and analog circuits. This innovative noise isolation architecture delivers a 5–10 dB reduction in noise coupling compared to traditional approaches. As an additional benefit, the GTZ transceiver design eliminates the need for external reference resistors, thus minimizing bill-of-materials (BOM) and board design complexity.

Tuning lossy channels to achieve the desired bit error rate (BER) typically requires expensive diagnostic equipment. GTZ transceivers include a built-in "eye scan" function that automatically measures the height and width of the post-equalization data eye. The eye scan capability is a powerful diagnostic tool that enables system designers to perform jitter budget analysis on an active channel and optimize transceiver parameters to obtain the required signal integrity, all without the expense of specialized equipment.

28G Transceiver Performance

The eye diagram in Figure 3 demonstrates the low jitter and high signal quality achieved by the GTZ transceiver's designers. The GTZ transceiver presents an open eye without excessive over-equalization, while delivering 28 Gb/s PRBS31 data and while being driven by a 560 MHz reference clock.

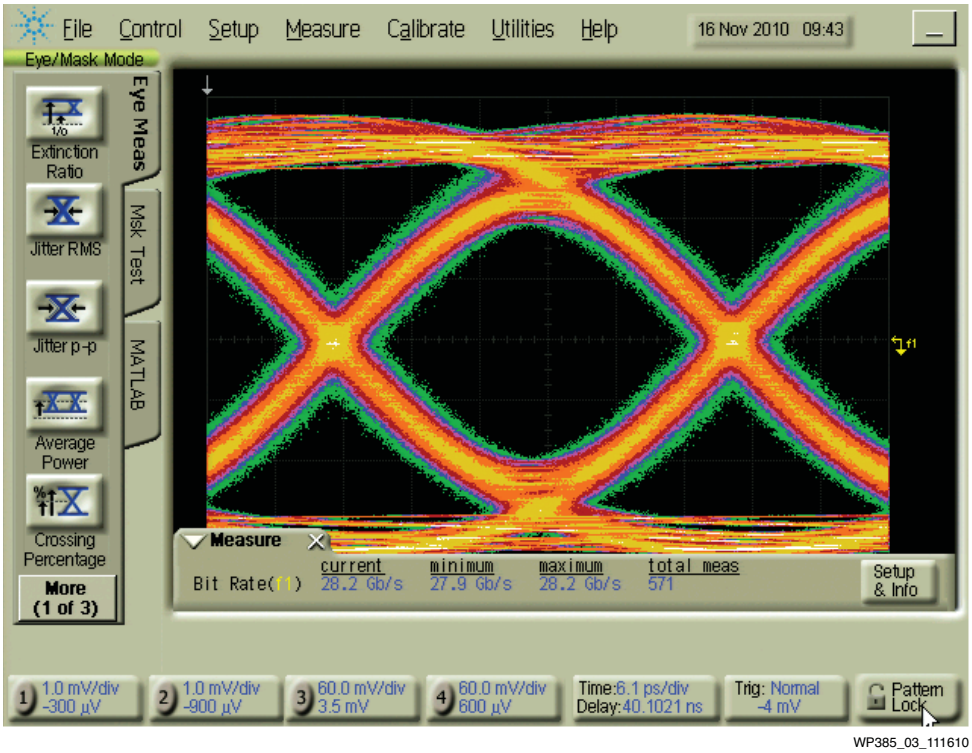

Figure 3: **28 Gb/s Eye Diagram of GTZ Transceiver Test Chip**

FPGAs for 400G and Nx100G Line Cards

Critical 28 Gb/s Electrical Interfaces

A look at a typical system architecture (Figure 4) reveals that serial I/O technology finds use in three capacities: the chip-to-optical-module interface on the network side; chip-to-chip communications on the board; and the chip-to-backplane interface that enables board-to-board communication. In next-generation systems, to achieve the goal of increasing faceplate bandwidth, the most urgent engineering pressure point is on the optical interface side. Achieving power-efficient and highest density 400G and Nx100G interfaces between the FPGA and the optical modules requires a move to 28 Gb/s serial I/O. The 28 Gb/s line rate allows for the coding overhead imposed by forward error correction (FEC) required for OTU4 long-haul transmission at 100 Gb/s data rates (via four bonded lanes). The G.709 standard defines a 7% overhead for FEC, resulting in a 28 Gb/s line rate.

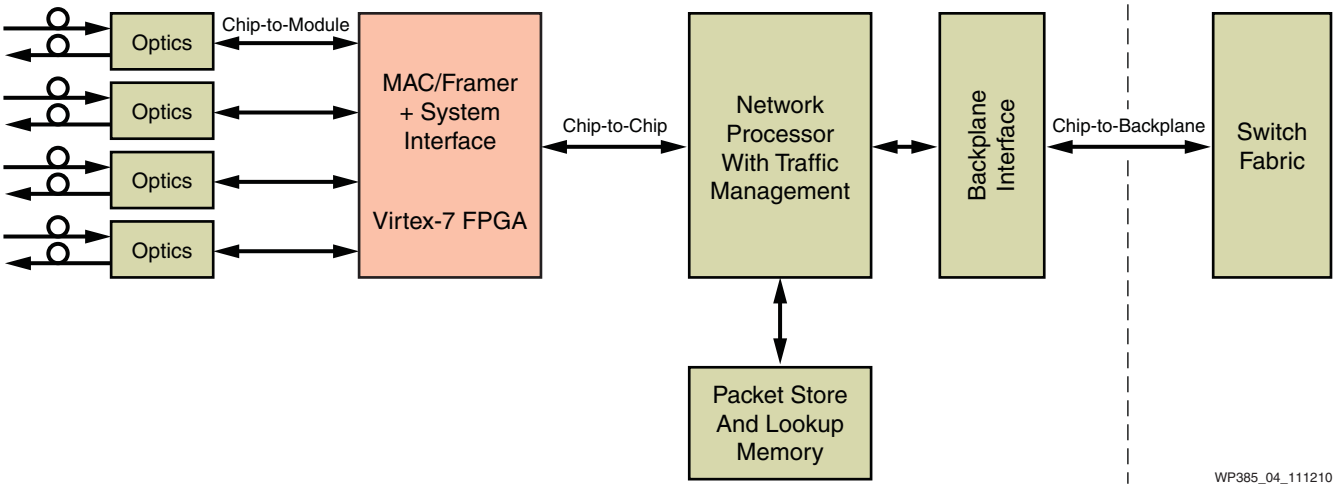

Figure 4: **FPGA Serial Interfaces in a Typical Line Card Architecture**

FPGAs play a critical role in telecommunication systems because their flexibility enables manufacturers to rapidly develop and deploy new systems that implement the latest networking standards, even as those standards are evolving. Virtex FPGA capabilities have continuously evolved to meet next-generation networking requirements by delivering greater capacity and performance, along with serial transceivers that support higher line rates and new protocols. To enable the industry's move to 25G optics, next-generation FPGAs must provide 28 Gb/s chip-to-module serial interfaces with support for advanced protocols, including 100GE, 400GE, OTU4, Common Packet Radio Interface (CPRI) 19.6, and 32G Fibre Channel.

Building the Industry's Highest Bandwidth FPGA

Virtex-7 FPGAs are fabricated on a 28 nm high-performance, low-power process developed to address the unique requirements of FPGAs. One distinction of the Virtex-7 family is unprecedented capacity, with up to 2,000,000 logic cells in a single FPGA. Because system performance is increasingly defined by bandwidth, advanced serial I/O is a critical component of Virtex-7 FPGA capabilities. The Virtex-7 family offers three types of serial transceivers:

- GTX transceivers enable 10.3125 Gb/s connectivity
- GTH transceivers provide 13.1 Gb/s line rates for 10G optical networking with up to 25% overhead for FEC or enhanced FEC (EFEC)
- GTZ transceivers deliver 28.05 Gb/s for next-generation optical networking

Virtex-7 HT devices expand the Virtex-7 family capabilities to enable 400G line cards and more (refer to DS180, *7 Series FPGAs Overview*). These new devices provide up to 16 x 28 Gb/s GTZ transceivers for optical module interfaces. This configuration satisfies a key requirement for 400G networking: the ability to monolithically terminate a 400 Gb/s data stream. Because each 400G packet is striped across all physical lanes in the 400G interface, recovering a complete packet requires terminating the entire stream. It is not possible to accomplish this using multiple 100G devices. Virtex-7 HT FPGAs also include up to 72 x 13.1 Gb/s GTH transceivers for chip-to-chip connections. This combination provides an unprecedented 2.8 Tb/s of bidirectional serial bandwidth in a single FPGA (Table 1).

Table 1: **Virtex-7 HT FPGA Competitive Comparison**

Capability	Virtex-7 HT FPGAs	Competing 28 nm FPGAs	Xilinx Advantage
28 Gb/s transceivers	16 x 28.05 Gb/s	4 x 28.0 Gb/s	4X
12+ Gb/s transceivers	72 x 13.1 Gb/s	32 x 12.5 Gb/s	2.25X
Aggregate serial bandwidth (bidirectional)	2.784 Tb/s	1.024 Tb/s	2.7X
Logic capacity	864K LC	622K LE	1.4X
Embedded RAM	64.8 Mb	50.0 Mb	1.3X

Multirate GTZ transceivers deliver the stringent jitter performance required for OTU4 and support line rates ranging from 19.6 Gb/s to 28.05 Gb/s to enable a number of next-generation protocols (Table 2).

Table 2: **Protocols Supported by Virtex-7 HT FPGA 28 Gb/s Transceivers**

Network	Electrical I/O Line Rate (Gb/s)	Bundled Line Rate (Gb/s)
100GE	25.78	4 x 25.78 = 103.125
400GE	25.78 (1)	16 x 25.78 = 412.48
OTN (OTU4)	27.95	4 x 27.95 = 111.81
CPRI 19.6 (Wireless)	19.6	614 Mb/s x 16 = 19.6
Fibre Channel (32GFC)	28.05	28.05

Notes:

1. Pending ratification by standards body.

Enabling 400G and Nx100G Applications with Virtex-7 HT FPGAs

A typical line card architecture (Figure 5) comprises an optical interface at one end and a backplane interface at the other. Various components perform electro-optical interfacing and functions such as framer/MAC and mapper, packet processing, and data management on the data passing through the card.

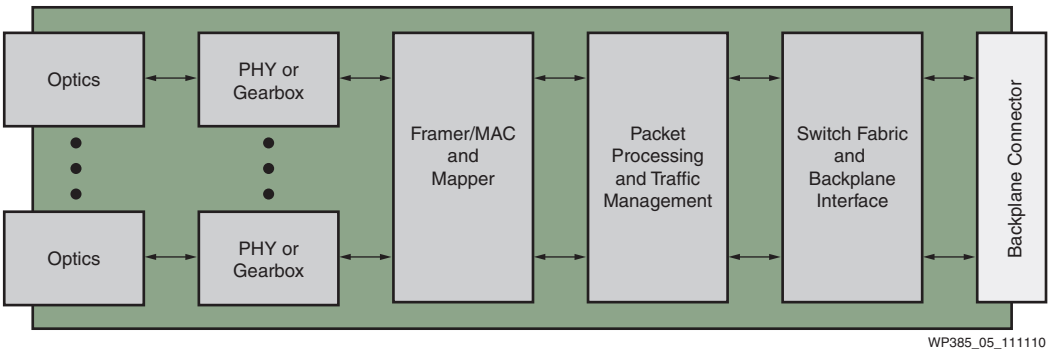

Figure 5: **Line Card Architecture**

To accelerate next-generation system deployment, designers can use Virtex-7 HT FPGAs with 28.05 Gb/s transceivers to create 400G and Nx100G line cards that leverage existing ASIC and ASSP components. 28.05 Gb/s GTZ transceivers enable interfaces to next-generation optical modules while 13.1 Gb/s GTH transceivers enable chip-to-chip interfaces with support for 10G data rates and the 25% overhead required for EFEC.

400G Line Card

Virtex-7 H870T FPGAs enable manufacturers to be first to market with 400GE line cards. These devices incorporate 16 GTZ transceivers for 400 Gb/s optical interfaces and up to 72 x 10.3125 Gb/s GTH transceivers for interfacing to as many as four ASICs/NPUs. With a capacity of 864,000 logic cells and 64.8 Mb of block RAM, these FPGAs enable sophisticated, high-performance in-line data processing (Figure 6 and Table 3).

Note: In Figure 6 to Figure 9, the red blocks at the top of each figure identify the portions of the line card architecture integrated into a single Virtex-7 FPGA.

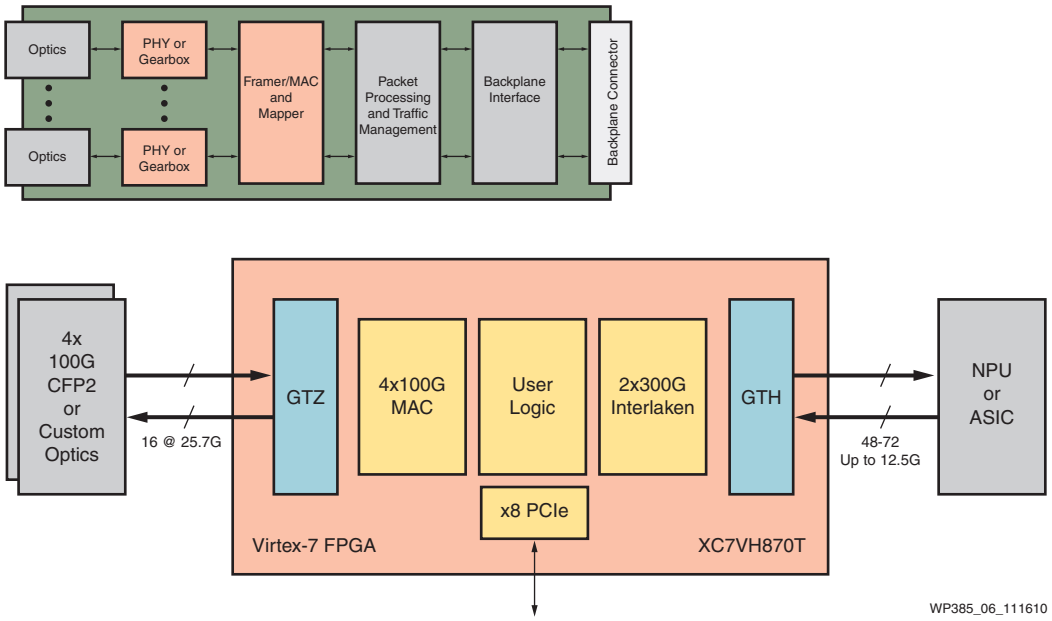

Figure 6: **400G Line Card Implementation**

Table 3: **400G Line Card Resource Utilization**

Resource	Number Required	Number Available
Logic Cells	As many as available	870k
28G Transceivers	16	8–16
13.1G Transceivers	52–72	72
Block RAM	40 Mb	64 Mb
SelectIO Interface	100	300–700

2x100G OTU4 Transponder/Line Card

Virtex-7 HT FPGAs are the only 28 nm FPGAs with resources that enable designers to integrate two 100G interfaces into a single FPGA for reduced board space, power, and cost. The devices accommodate the OTU4 Framer/MAC and mapper, along with a PHY "gearbox" function that converts ten lanes of 11.18 Gb/s data to four lanes of 28 Gb/s (Figure 7 and Table 4). The high-performance FPGA logic enables the addition of FEC for increased functionality. The Virtex-7 H580T FPGA is suitable for low-gain FEC like G.709. For implementing proprietary FEC, the Virtex-7 H870T FPGA delivers 40% more logic than competing 28G FPGAs. GTH transceivers support chip-to-chip connection with a variety of legacy ASIC interfaces, including 12.5G, 10.3125G, and 6.25G.

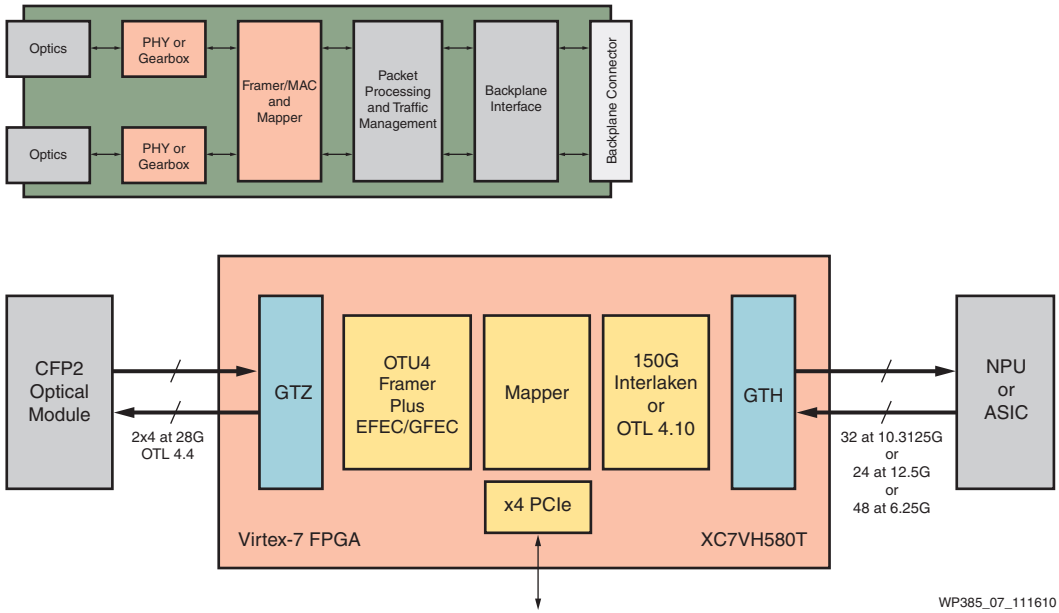

Figure 7: **2x100G OTU4 Transponder/Line Card Implementation**

Table 4: **2x100G OTU4 Transponder/Line Card Resource Utilization**

Resource	Number Required	Number Available
Logic Cells	300–600k	580–870k
28G Transceivers	8	8–16
13.1G Transceivers	34–52	48–72
Block RAM	20 Mb	43–64 Mb
SelectIO Interface	100	300–700

100G OTU4 Muxponder

Virtex-7 H580T FPGAs enable a flexible, single-FPGA multiplexing transponder implementation. Multirate GTH transceivers support line rates up to 13.1 Gb/s to carry the overhead required for FEC. Virtex-7 H580T devices enable up to 12 independent multi-protocol tributary interfaces, while the larger Virtex-7 H870T FPGAs offer a migration path to additional logic resources for implementing EFEC and up to 18 independent multi-protocol tributary interfaces (Figure 8 and Table 5).

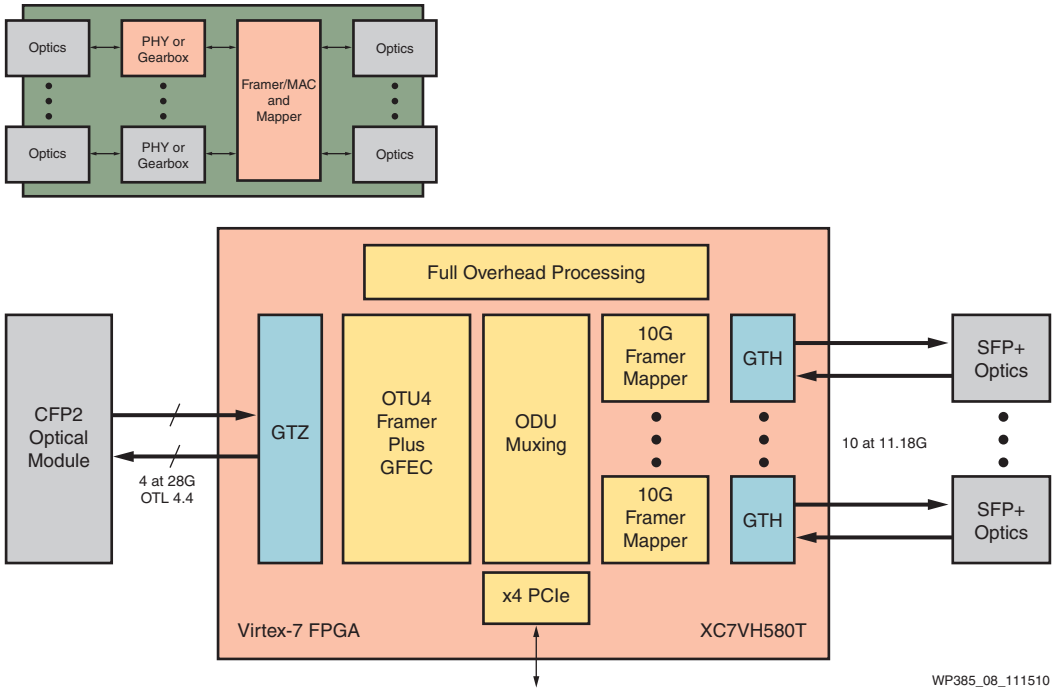

Figure 8: **100G OTU4 Muxponder Implementation**

Table 5: **100G OTU4 Muxponder Resource Utilization**

Resource	Number Required	Number Available
Logic Cells	400k–800k	580–870k
28G Transceivers	4	8–16
13.1G Transceivers	44	48–72
Block RAM	20 Mb	43–64 Mb
SelectIO Interface	100	300–700

100G “Smart Gearbox”

Xilinx 28 Gb/s FPGAs enable manufacturers to leverage legacy hardware to rapidly deploy new 100G networking products. In this application, an existing CFP-based 100GE and OTU4 line card is upgraded with the higher-density, lower-power CFP2 optics by using a Virtex-7 H290T FPGA to create an interface between an existing 10G-based ASIC and the new 28G optics. At costs equivalent to commodity gearbox devices, Virtex-7 HT FPGAs provide additional flexibility for bridging proprietary ASIC interfaces and custom optics as well as implementing standards-based interfaces. In addition, system manufacturers can add differentiation to proven ASICs and ASSPs by implementing custom in-line data processing functionality in the FPGA logic array (Figure 9 and Table 6).

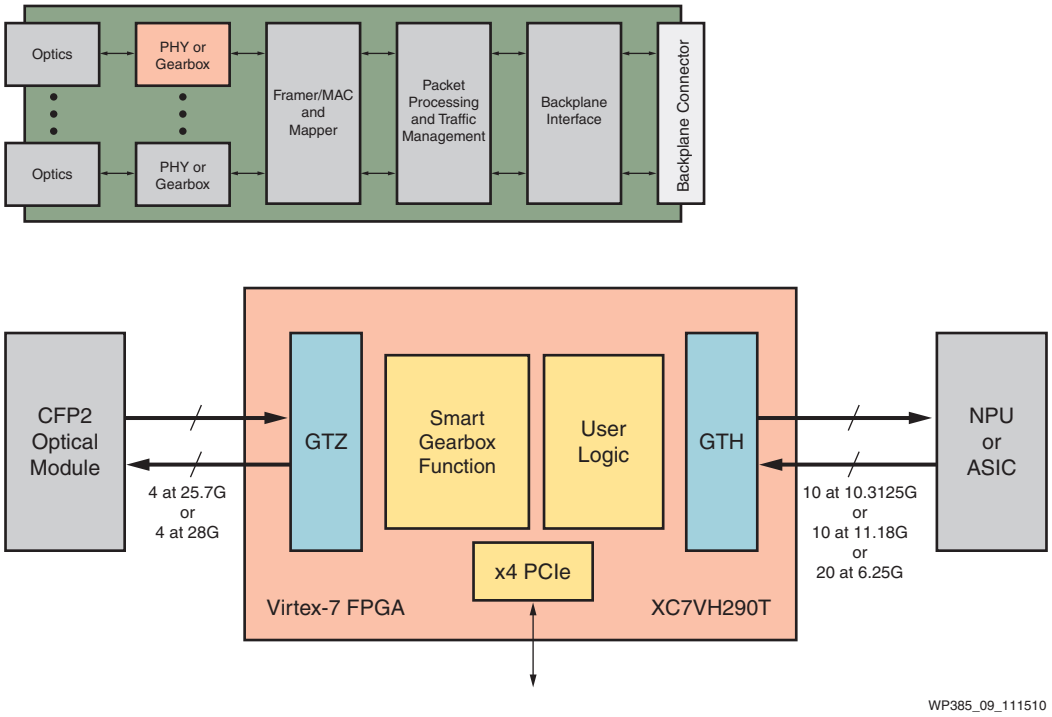

Figure 9: **100G “Smart Gearbox” Implementation**

Table 6: **100G “Smart Gearbox” Resource Utilization**

Resource	Number Required	Number Available
Logic Cells	200k	290k
28G Transceivers	4	4
13.1G Transceivers	10–24	24
Block RAM	10 Mb	20 Mb
SelectIO Interface	100	300

Conclusion

The market demand for higher bandwidth networking line cards is real. New technology is required to tackle the cost and power challenges of higher density 100G interfaces. Xilinx is at the forefront with new low-jitter 28 Gb/s transceivers that support the adoption of new high-density 100GE and OTU4 optics. Virtex-7 HT FPGAs combine the highest bandwidth serial interfaces with the highest density logic arrays to enable the industry's first solution for 400G and Nx100G networking.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
11/17/10	1.0	Initial Xilinx release.
11/22/10	1.1	Updated Figure 3, Figure 6, and Figure 7.

Notice of Disclaimer