


As various standard bodies finalize their 100G standards for transport, Ethernet, and optical interfaces, FPGAs play a vital role for early adopters of technology who want to design 100G production systems. Because of this increasing demand for more bandwidth, service providers are looking at emerging 40-GbE/100-GbE standards for their next-generation line card options. Altera® Stratix® V FPGAs solve the bandwidth problem by providing integrated 12.5-Gbps transceivers with hardened 100G PCS functions on the 28-nm technology node.

Introduction

Two qualities are important in a network: speed and reliability. Not only must the network be up all the time, it must also be fast. However, the load on networks has increased tremendously. Data is a minor component of what the network carries; voice, sound, and multimedia now form the major components.

According to the Cisco Visual Networking Index (VNI) Forecast, annual global IP traffic will reach two-thirds of a zettabyte (trillion gigabytes) by 2013. This number represents more than a fivefold increase over today's IP traffic. As [Figure 1](#) shows, video will account for 90 percent of the traffic growth in 2013.

Figure 1. Total Traffic Bandwidth Increase, 2008–2013



Source: Cisco VNI, June 2009

Driven by video, Internet traffic is growing at a phenomenal rate. The Beijing Olympics broadcast, viewed by over 3 million people, was estimated to generate 128 terabytes (Tbytes) of video traffic on the NBC network. On a typical day, YouTube generates over 1100 Tbytes of traffic to serve over 5 billion streams to over 91 million viewers.

Satisfying the Demand for High Bandwidth

Current service providers are continuously looking for technology advancements to keep up with demand, and seeking ways to optimize their network infrastructure. These service providers must continue to make a profit by reducing the cost per bit while simultaneously extending their service offerings. High-speed Ethernet, especially 100G Ethernet (100 GbE), offers the key solution. Today, many carriers already provide 10GbE links between their routers, and the adoption is at aggregation points at the transport level. As the demand for more bandwidth becomes increasingly prevalent, service providers are looking at emerging 40-GbE/100-GbE standards for their next-generation line card options. Many are contemplating a direct switch to 100-GbE, while others are evaluating the market dynamics by accessing the availability and economics of both 40-GbE and 100-GbE solutions to meet customer demands.

High-speed 100-GbE links are critical at the edge-router level as they enable operators to simultaneously achieve the least-cost bit transport and service for high-value, revenue-generating traffic, such as high-definition video, mobile LTE, and VPN content. The biggest challenge for such a line card is to deliver speed and quality of service (QoS) at the same time. Unfortunately, the existing infrastructure and routers are not yet geared up to efficiently support this type of traffic.

40-GbE and 100-GbE IEEE 802.3ba Standards

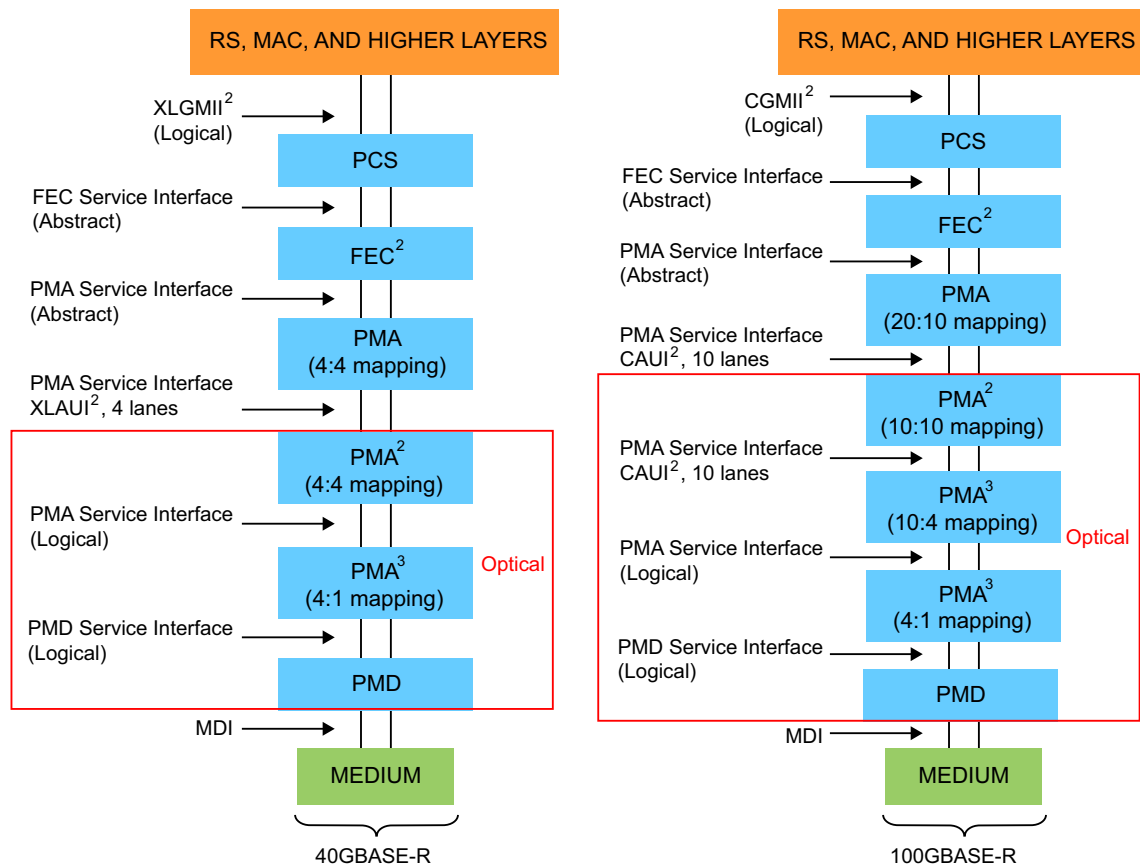
The IEEE 802.3ba High Speed Study Group (HSSG) was formed in late 2006 to study market needs and the definition of standards for the 100-GbE interface protocol. In 2007, an additional proposal to include a 40-Gbps rate was adopted. The initial draft from the HSSG was proposed on August 2008 and it provided the details of implementing this protocol. The project is targeted to be completed in the middle of 2010.

The purpose of the IEEE 802.3ba proposal was to extend the 10GBase-R Carrier Sense Multiple Access with Collision Detection (CSMA/CD) standards to be included for both the 40-Gbps and 100-Gbps protocols (shown in [Figure 2](#)) with the following objectives:

- Support full-duplex operations only
- Preserve the 802.3/Ethernet frame format utilizing the 802.3 MAC
- Preserve the frame size specification of the current 802.3 standard
- Support a bit error rate (BER) better than or equal to 10^{-12}
- Comply with the OTN for WAN applications
- Support a MAC data rate of 40 Gbps

- Provide physical-layer specifications, which support 40-Gbps operation over:
 - At least 10 km on SMF
 - At least 100 m on OM3 MMF
 - At least 10 m over a copper cable assembly
 - At least 1 m over a backplane
- Support a MAC data rate of 100 Gbps
- Provide physical-layer specifications, which support 100-Gbps operation over:
 - At least 40 km on SMF
 - At least 10 km on SMF
 - At least 100 m on OM3 MMF
 - At least 10 m over a copper cable assembly

Figure 2. Summary of 40GBASE-R and 100GBASE-R Sublayers

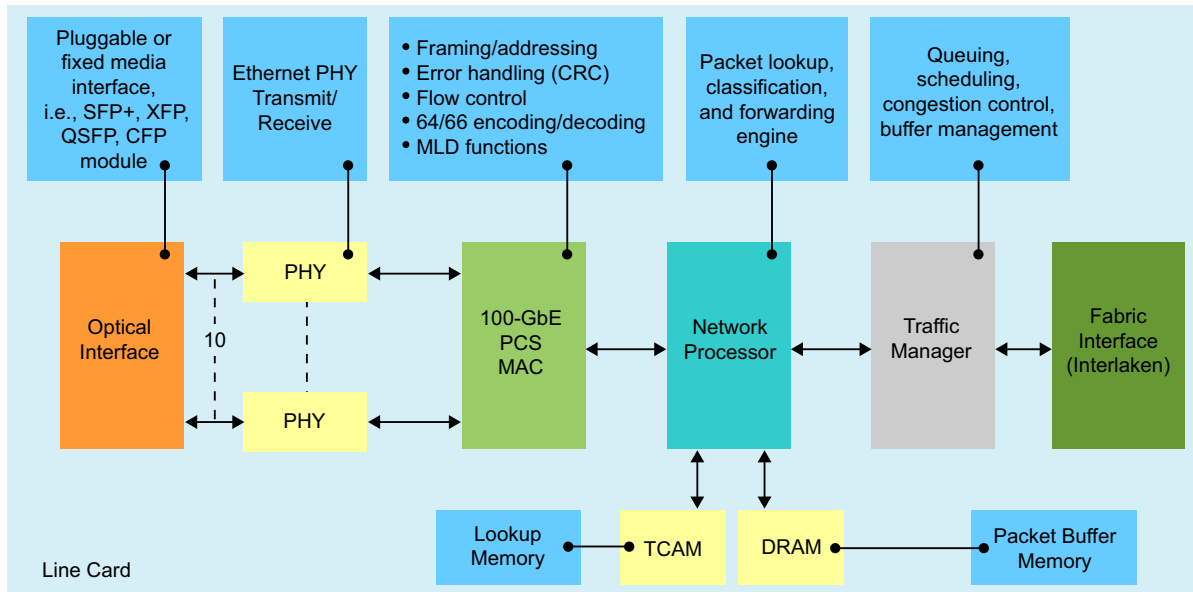


Notes:

- (1) Source: IEEE 802.3ba Section 82.1.3
- (2) Optional
- (3) Conditional based on PHY type

The FPGA market segment plays a vital role during the initial phases of implementing the 40-GbE and 100-GbE IEEE 802.3ba standards, which include the draft stage, evaluation of its merits, and finally prototyping it on the current platform. The basic block diagram of a 100-GbE line card is shown in Figure 3.

Figure 3. 100-GbE Line Card Components



A 100-GbE line card includes the following components:

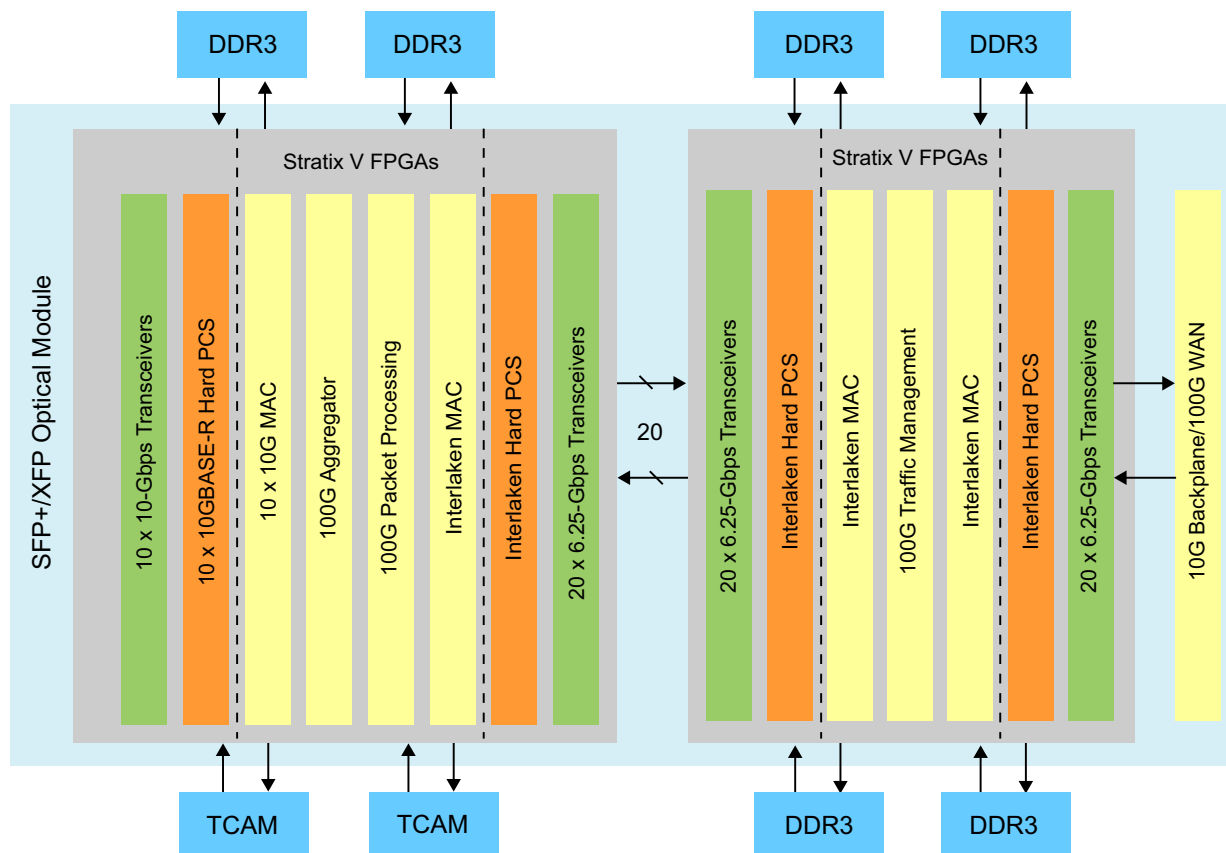
- **Optical interface**—The optical interface unit can consist of multiple SFP+ or XFP modules, or it can be driven by 100G traffic via CFP or QSFP modules.
- **PHY**—The PHY unit is the serializer/deserializer (SERDES) component of the line card. The PHY line rate and jitter specifications should be compliant with the optical interface.
- **MAC/PCS**—The MAC/PCS unit performs the gearbox, scrambling, and encoding functions based on the protocol. In the case of 40-GbE or 100-GbE implementations, there is a multilane distribution (MLD) function as per the IEEE 802.3ba specifications. In addition, flow control as well as error handling is performed by the MAC. In some cases, the received 10G data from the MAC unit is aggregated before it is passed over to the network processing unit (NPU).
- **NPU**—The key function of the NPU is to optimize the performance of packet processing in the evolving functional framework of the line card. Key functions include compression, classification/lookup, modification, and deep-packet inspection. The most common function of the NPU is to interface with a switch fabric device that performs complicated routing of the packets through the network.
- **Traffic manager**—The primary function of the traffic manager is to offer a large number of high-speed queues, optimize queue depths, and use sophisticated scheduling mechanisms to meet the QoS requirements of the application. Because NPUs are not designed with QoS in mind, they require excessive processing power and software optimization before they can function as efficiently as a dedicated traffic manager.

Addressing the Challenges of Multiport 100-GbE Designs

The ever-growing need for more bandwidth is driving service providers to find quick routes to deploy 40-GbE and 100-GbE systems. Many service providers are adopting 40-GbE solutions first, because of the cost and complexity associated with 100-GbE optics. Others are jumping on the 100-GbE bandwagon immediately and are addressing the challenges of designing 100-GbE systems. Some of these challenges include designing ultra-high-bandwidth interfaces and integrating more functionality in existing systems while staying within existing power and cost budgets.

Altera's 28-nm Stratix V FPGA family is well-positioned to meet the performance and system bandwidth requirements of 100G transport and 100-GbE system designs. Stratix V FPGAs provide the highest density with integrated 12.5-Gbps and 28-Gbps transceivers along with hardened MLD and PCS functions supporting the IEEE 802.3ba specifications. In addition, Stratix V FPGAs deliver the high density and performance requirements for implementing the enhanced forward error correction (EFEC) functions, making them an ideal platform for OTN system designs for algorithm implementation, testing, and production. Figure 4 shows a typical example of how designers can use Stratix V FPGAs to develop their 100-GbE line cards.

Figure 4. Aggregated 100-GbE Packet Processing and Traffic Management Using Stratix V FPGAs



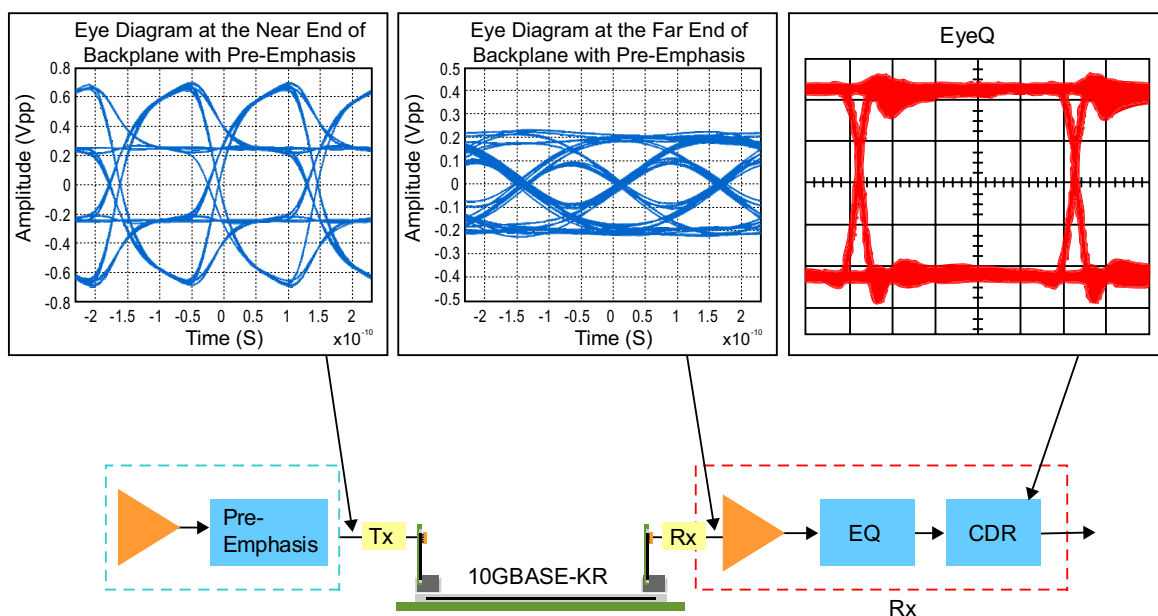
Ultra-High-Speed Serial Connectivity at the Lowest BER

100-GbE line cards require high bandwidth with the highest reliability. Stratix V FPGAs deliver best-in-class transceivers capable of directly driving backplanes and optical modules.

100G data traffic must be handled while maintaining the lowest BER. Advanced features in the transceivers simplify the PCB design and compensate for the inevitable board losses through adaptive linear and DFE equalization, multi-tap pre-emphasis, and EyeQ eye monitor. Additional features, such as on-die capacitance and on-package decoupling, enhance the Stratix V FPGA die and package, resulting in excellent signals, power integrity, and maximum user flexibility.

Stratix V FPGAs offer the EyeQ eye viewer (shown in Figure 5), an on-die instrumentation tool to minimize board bring-up and debug time. The EyeQ eye viewer allows users to fully reconstruct the eye diagram at the receiver side without interrupting the data path. Users can then leverage the dynamic reconfiguration of the transceivers to adjust the equalization settings and optimize the eye diagram for their applications.

Figure 5. 10-GbE Channel over 10GBASE-KR Backplane, Showing Stratix V Signal Conditioning at the Transmitter and Receiver Ends with the EyeQ Eye Viewer



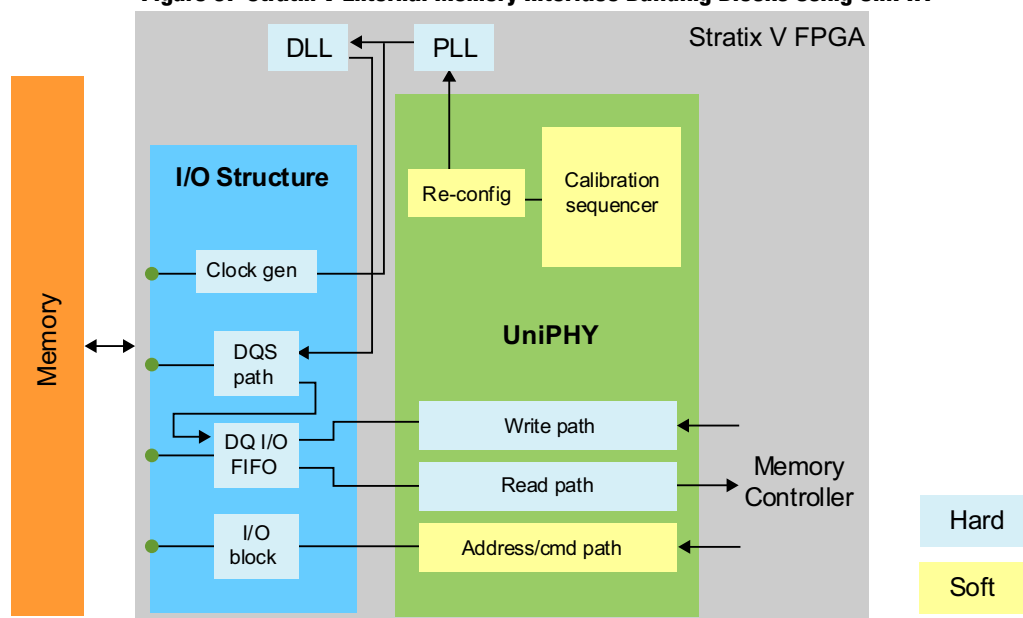
Stratix V FPGAs can handle up to 825-Gbps full-duplex serial bandwidth with up to 66 backplane-capable serial transceivers on one device. Data from various clients can be aggregated efficiently to form a 100G data pipe by leveraging the continuous operation range of each transceiver from 600 Mbps to 12.5 Gbps. Stratix V transceivers have an integrated electronic dispersion compensation (EDC) capability that enables them to drive various optical modules including SFP+, QSFP, and CFP without the need for any external PHY device.

Highest External Memory Interface Performance

100-GbE systems require various external memory types to support the data-intensive packet processing and traffic management functions. DDR3 memories are usually used for packet buffering, while QDR II+ memories are used whenever low latency is required. A third type is ternary content addressable memory (TCAM), mainly used for content-based processing. The performance of the external memory interface affects the overall system performance; therefore, 100G designers usually require the fastest external memory interfaces available. Stratix V FPGAs support 800-MHz DDR3 DIMM interfaces, 550-MHz QDR II+ interfaces, and next-generation serial-based TCAMs. The write and read paths are fully hardened in Stratix V FPGAs to guarantee timing closure at 800 MHz.

Stratix V FPGAs are supported by the new UniPHY, shown in Figure 6, in the Quartus® II design software. The UniPHY has enhanced features including a lower read latency and easier sharing of resources, as well as more DIMM and rank support. In addition, the UniPHY is available as an unencrypted cleartext with a Nios® embedded processor-based sequencer for easier debug, and is supported by flexible timing models for improved transparency of timing nodes and higher accuracy.

Figure 6. Stratix V External Memory Interface Building Blocks Using UniPHY



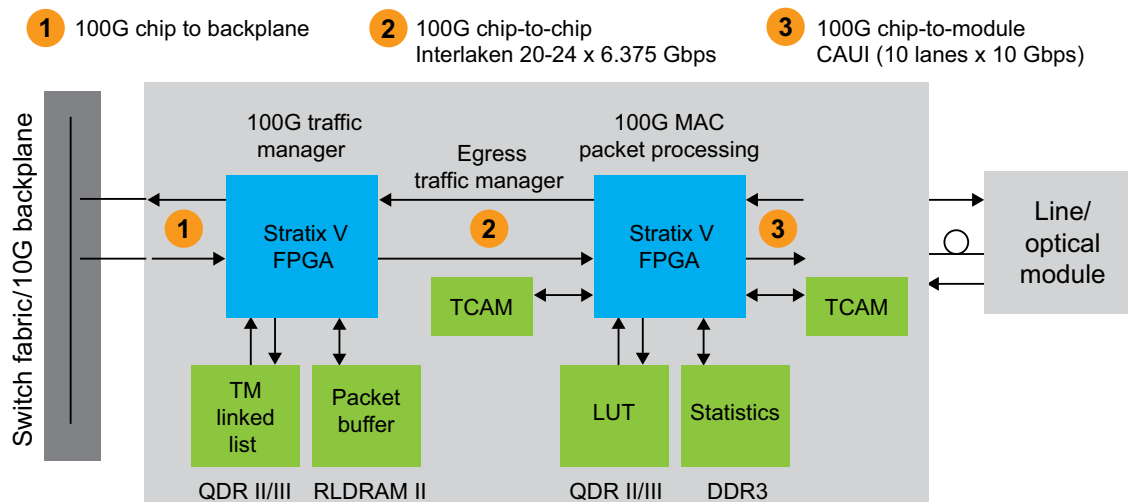
With the UniPHY and key innovations in hardware, Stratix V FPGAs deliver the high-bandwidth memory interfaces required by designers to implement an optimized solution for 100-GbE line cards.

Highest System Integration for Increased Functionality in a Smaller Line Card

A key challenge in designing 100-GbE line cards is determining the level of integration that can be achieved on a small form factor line card without exceeding the cost and power-consumption budget. Designing with Stratix V FPGAs allows developers to reach an unprecedented level of system integration on one FPGA while maintaining low cost and power consumption. Stratix V FPGAs offer a fully hardened PCS per transceiver supporting 3G, 6G, 10G, and Interlaken protocols. In addition, Stratix V FPGAs offer an Embedded HardCopy® Block, a programmable-gate array equivalent to 14.3 million ASIC gates or 1.19 million logic elements (LEs). This feature allows designers to harden standard or logic-intensive applications within the FPGA.

Using the Embedded HardCopy Blocks, designers can harden the full 100-GbE protocol, saving hundreds of LEs for implementing the remaining design functions. In this way, designers can fit 100-GbE logic functions, including traffic management and packet processing, in as few as two FPGAs, as shown in Figure 7.

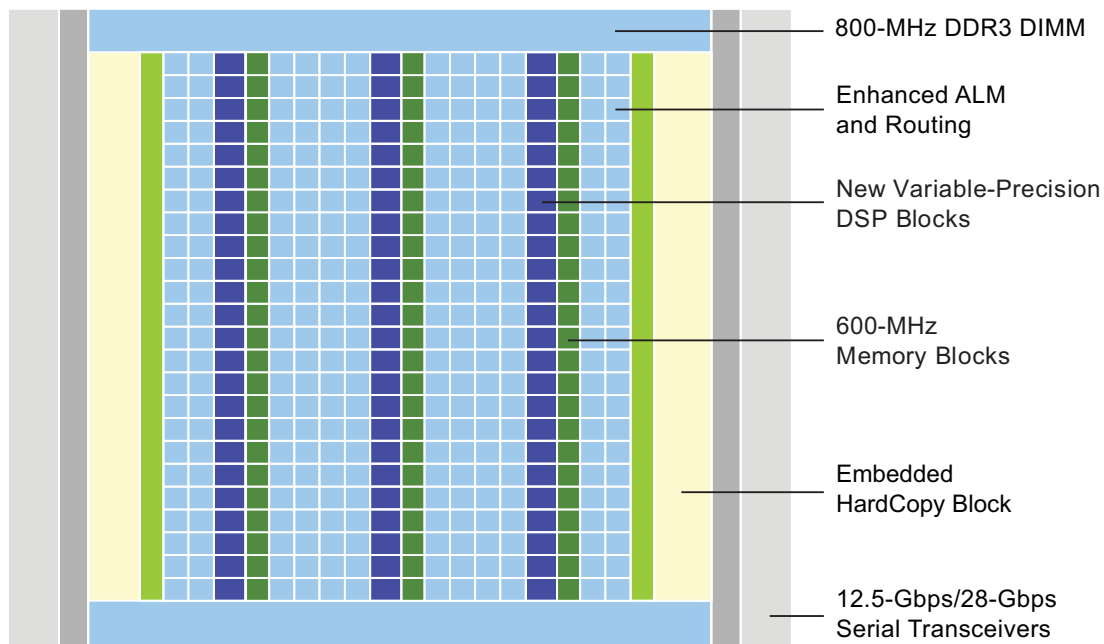
Figure 7. 100-GbE Line Card Using Two Stratix V FPGAs



There are up to six Embedded HardCopy Blocks on one Stratix V FPGA for a total of 1.19 million LEs. Stratix V FPGAs are customized using these Embedded HardCopy Blocks to deliver devices with PCI Express Gen1/Gen2/Gen3 hard IP or hardened 40-GbE/100-GbE PCS. For 100-GbE line card designs, developers can start with the standard offering from Altera, which hardens the 100-GbE PCS, and use the remaining Embedded HardCopy Blocks to harden other functions related to their application.

Highest System Performance at the Lowest Power

When protocols have not yet been standardized, designers often seek a flexible platform for their designs to allow them to easily adapt to any future changes in the protocol standards. Because of their flexibility, FPGAs are the ideal vehicles to implement protocols that are still in draft mode. Two key metrics for an FPGA to be used to implement high-bandwidth protocols such as 100-GbE are the I/O interface and the fabric performance, shown in Figure 8.

Figure 8. Stratix V FPGA I/O and Fabric Functional Blocks Enabling Higher System Performance

Faster I/O interfaces allow FPGAs to receive and transmit data at high rates and usually require high-bandwidth transceivers and high-performance memory interfaces. Stratix V FPGAs deliver the highest bandwidth transceivers in the FPGA industry with 12.5-Gbps and 28-Gbps transceivers. For memory interfaces, Stratix V FPGAs are targeted to support up to six 72-bit DDR3 DIMM interfaces, each running up to 800 MHz.

The second key metric to enable higher system performance is the FPGA fabric. Faster fabric allows designers to develop simpler designs and fit more logic in the FPGA. Slower FPGAs usually require designers to leverage more pipelining and parallelism to make up for the lack of performance, which leads to higher logic utilization, higher power and higher cost, thereby forcing designers to use higher density or multiple FPGAs.

Stratix V FPGAs are built on a high-performance process and include key innovations to address the high density and the high-system-performance requirements of 100-GbE designs. These innovations and enhancements include:

- 50% increase in memory-interface performance of up to 800 MHz made possible by the full hardening of the memory read and write paths
- Higher bandwidth transceivers with built-in advanced signal-conditioning circuitry, enabling the direct drive of 10GBASE-KR backplanes and lower power per gigabit
- Enhanced adaptive logic modules (ALMs) with four registers providing higher performance and easier timing closure for designs that are register rich and heavily pipelined
- Enhanced MultiTrack routing architecture with more routing resources, resulting in less routing congestion, higher logic utilization, and reduced compile times for tightly packed designs

- New 20-Kbit internal memory blocks enabling higher performance of up to 600 MHz in various memory modes, with built-in error correcting code (ECC) protection
- Enhanced distributed memory blocks (MLABs) with additional built-in registers delivering higher performance up to 600 MHz for optimized implementation of wide shallow FIFOs.

In addition, Stratix V FPGAs include process and architectural innovations to deliver the lowest power at the highest performance, with:

- High-performance, high-K metal gate 28-nm process technology optimized for low power, delivering leading-edge performance and power efficiency
- 0.85-V core supply voltage
- Programmable Power Technology to automatically reduce static power dissipation in the design while maintaining the high performance required for critical timing paths
- Clock gating technology to minimize dynamic power dissipation and prevent unnecessary switching noise

By offering a high-performance fabric with key enhancements for higher system performance, Stratix V FPGAs allow designers to stay within strict power and cost budgets. Efficient 100-GbE designs can be implemented in Stratix V FPGAs while avoiding the heavy pipeline and parallelism requirements of typical FPGAs.

Full Solution for 100-GbE Designs

Altera strives to offer designers the required resources to shorten their design's time to market. Altera and its partners offer a broad portfolio of off-the-shelf, configurable intellectual property (IP) cores optimized for Altera devices. These offerings include targeted reference designs, fully tested IP, development kits, and expert technical support. Stratix V FPGAs are supported by the industry-leading Quartus II design software featuring productivity-boosting tools such as power and performance-driven compilers, Advanced Blockset, and easy-to-use partial reconfiguration.

Conclusion

100G system designs are gaining popularity as bandwidth requirements continue to grow exponentially. Altera's 28-nm Stratix V FPGAs offer designers an ideal FPGA solution for 40-GbE/100-GbE systems by delivering the highest density FPGAs with integrated 12.5-Gbps transceivers and hardened industry-standard IP. Stratix V transceivers support the emerging 100-GbE standards and proprietary serial protocols with line rates up to 28 Gbps. Stratix V FPGAs are specifically built to meet the requirements of next-generation 100-GbE system designs, thereby offering system designers an accelerated time-to-market advantage with reduced risk compared to ASIC or ASSP solutions. By using Stratix V FPGAs, designers can implement current 100-GbE standards and easily adapt their designs to future enhancements.

Further Information

- Stratix V FPGAs: Built for Bandwidth:
www.altera.com/products/devices/stratix-fpgas/stratix-v/stxv-index.jsp
- Literature: Stratix V Devices:
www.altera.com/products/devices/stratix-fpgas/stratix-v/literature/stv-literature.jsp
- White paper: *Cisco Visual Networking Index: Forecast and Methodology, 2008-2013*:
www.cisco.com

Acknowledgements

- Rishi Chugh, Senior Product Marketing Manager, Altera Corporation
- Frank Yazbeck, Senior Technical Marketing Engineer, High-End FPGAs, Altera Corporation

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
July 2010	1.1	Minor text edits.
April 2010	1.0	Initial release.