UltraScale Brings Interlaken Onboard for High-Bandwidth Designs

Xilinx has integrated Interlaken IP on its UltraScale FPGAs to simplify interconnects in packet-processing systems.

by Martin Gilpatric
Transceiver Technical Marketing Manager
Xilinx, Inc.
martin.gilpatric@xilinx.com

Bandwidth is a funny thing. Ten years ago, most people barely understood the term "byte per second." But today, online video, smartphones and all of the trappings of our modern interconnected society have pushed bandwidth into the common consciousness. These applications have laid bare to the public why bandwidth is important and why they want more of it.

Unfortunately, that high-level, top-of-the-stack perspective obscures the technological advances required to supply those additional bits and bytes to our homes and to the devices in our pockets. All of that data requires that our infrastructure must keep up. Interconnects that used to rely on 10-Gbps optics have moved to 40 to 100 Gbps, and now we are looking at 400 Gbps in the very near future.

INTERLAKEN IN THE INDUSTRY TODAY

Through all of this, the fundamental packet-processing tasks and architectures have not changed drastically. Similarly, many of the packaging, power and thermal constraints facing 100-Gbps systems are the same faced by their predecessors. This means that every component of the system needs to operate at a much higher rate. More important, the interconnects between these devices need to be able to scale readily. Confounding this task is the multitude of suppliers for the FICs, NPUs, MACs and other ASSPs used in these systems. These devices use a variety of interconnect widths and rates to achieve the target throughput.

Now, Xilinx and the Interlaken protocol have stepped in to assist in the serial-interconnect bottlenecks between packet-processing functions (Figure 1). Interlaken arose from a joint effort between Cisco and Cortina, and was expressly designed to meet these challenges. Xilinx® FPGAs readily implement this standard via high-performance programmable logic and high-performance transceivers. Three things make Interlaken the optimal protocol for these situations: It is highly scalable, it was designed to work well with OTN and Ethernet systems, and it is an open protocol.

At the most fundamental level, Interlaken is designed to maximize

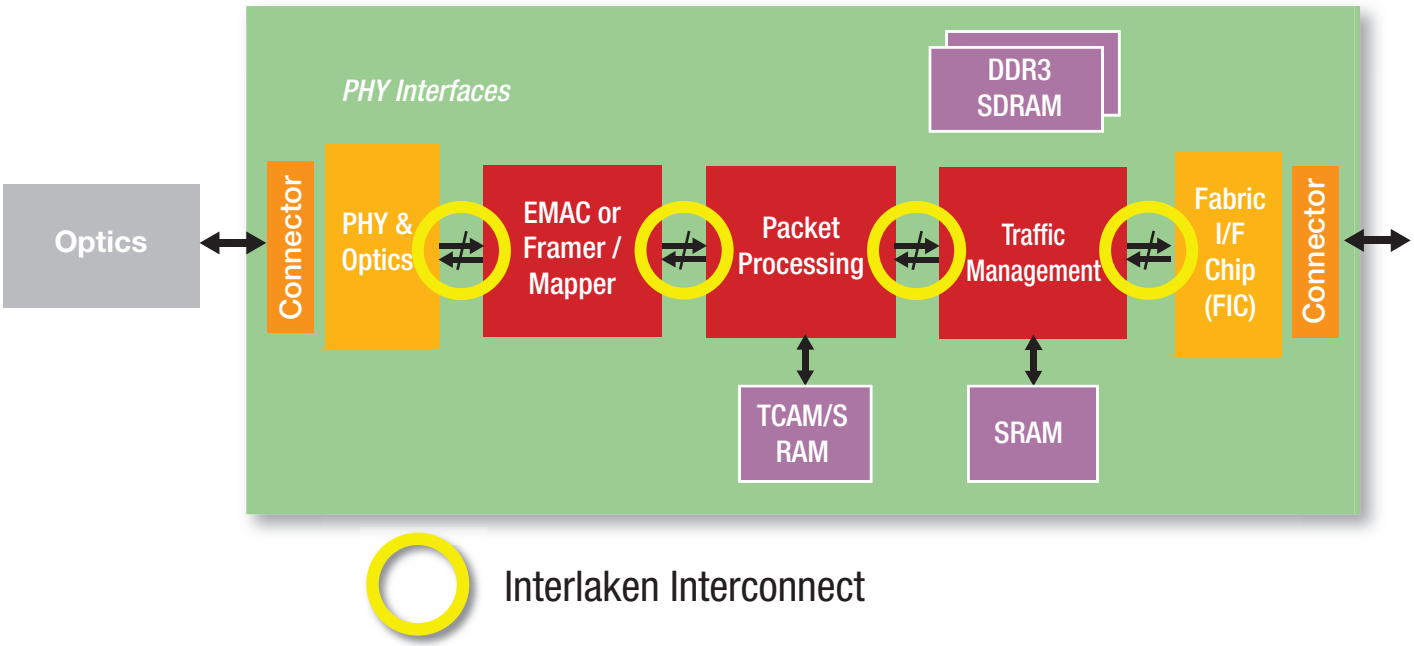

Figure 1 – Potential locations for Interlaken in packet processing

flexibility and scalability. There are no predefined line rates or lane widths mandated by the specification. This means, for example, that you can implement a 150-Gbps interface, enough to encapsulate a 100G Ethernet data path (with additional packet headers), in 12 lanes of 12.5 Gbps each or six lanes of 25 Gbps. In either case, the logical interface to Interlaken is maintained, simplifying the rest of the design. Furthermore, this flexibility makes it possible to adapt Interlaken for use in niche applications such as lookaside interfaces when working with TCAMs.

Regardless of the number of lanes the data is being transmitted across, the data is segmented and striped across those lanes. This approach has the advantage of lowering latency while working well with OTN streams or Ethernet packets. Within this segmentation, Interlaken allows for channelization, giving designers the flexibility to implement features like prioritization of packets while working within the protocol's existing feature set.

Interlaken's rich feature set and the fact that it was freely licensed have made for ready adoption. Also, including a variety of stakeholder companies in the original definition meant that Interlaken's feature set was ready to meet a wide selection of needs. A variety of ASSPs and IP quickly became available, easing both adoption and, of course, the use of Xilinx FPGAs.

XILINX AND INTERLAKEN: A PERFECT FIT

Xilinx has a reputation for making it easy for designers to adopt the latest high-performance standards. In the move to 100 Gbps and beyond, Xilinx devices have proven invaluable in leading packet-processing systems and test equipment. The availability of high-quality IP for 100G Ethernet MAC, OTU4 and Interlaken solutions complements Xilinx's high-speed FPGA fabric and world-class transceivers in providing flexible, powerful solutions for customers.

Figure 2 – With Virtex UltraScale's 32.5-Gbps GTY high-performance transceivers (demonstrated in this video) and integrated Interlaken blocks, customers can create higher-performance, lower-power and more tightly integrated networks without sacrificing flexibility.

Interlaken IP has been available for Xilinx FPGAs since the Virtex®-5 generation, originally from third-party vendors and then, after the acquisition of Sarance, through Xilinx itself. The UltraScale™ family of FPGAs is the fourth generation of Xilinx devices to offer Interlaken solutions, with one important advance.

In UltraScale FPGAs, Xilinx has integrated the Interlaken IP into the silicon itself. By making Interlaken a fixed feature, Xilinx frees the hands of designers by opening up more fabric logic and reducing the strain on timing that a soft implementation introduces. The integrated solution also saves both dynamic and static power without sacrificing flexibility. A single Interlaken block can implement as many as 12 lanes at any line rate up to 12.5 Gbps or up to six lanes all the way to 25 Gbps.

XILINX BUILDS ON INTERLAKEN'S ADVANTAGES

The integrated Interlaken IP is only one piece of what makes Xilinx UltraScale FPGAs so attractive. With a large number of feature-rich transceivers supporting a wide range of line rates, advanced high-speed programmable logic, integrated and soft IP for MAC solutions, and the capability to implement emerging standards, no other FPGA or ASSP can measure up.

The GTH and GTY transceivers found in the UltraScale FPGAs have a number of features that enable them to perform in a wide variety of situations. The GTH has line rates ranging from 500 Mbps to 16.3 Gbps, while the GTY can run all the way to 32.75 Gbps, ensuring that the FPGA can support any line rate required by a link partner (see video demo on YouTube, Figure 2). The equalization capabilities of these transceivers, including continuous-time linear equalization and decision-feedback equalization, allow users to connect to anything from another device on the board to optics or a device on the far side of a backplane.

To ease the bring-up of these high-speed interconnects, whether 100G Ethernet, OTU 4.4 or a wide Interlaken interface, all of the receiver equalization features are auto-adapting. This means that between the three taps of CTLE and 15 taps of DFE in the GTY, the millions of combinations that result are handled by the transceivers themselves. There's no need for the designer to hand-tune each link and maintain margin on those links over changes in process, voltage and temperature. With Interlaken supporting indefinitely wide interfaces, auto-adaptation keeps bring-up simple and robust.

Complementing the integrated Interlaken block is the UltraScale integrated 100G Ethernet MAC, or CMAC. With all of the same advantages as the integrated Interlaken block, the CMAC saves power, complexity and implementation time while providing a ready interface to today's 100G optics: CFP, CXP, CFP2, CFP4 and QSFP28. Connecting to these optics requires either a 10x10.3125-Gbps or 4x25.78-Gbps serial interface. The CMAC and the UltraScale transceivers support these two interfaces—GTH for 10.3125 Gbps and GTY for both line rates. The CMAC provides but one possible protocol to interface off a line card. Xilinx provides IP that can be implemented in fabric for OTN applications or an assortment of Ethernet standards: 1G, 10G, 40G, 100G and even the emerging 400G standard.

Bridging between ASSPs that leverage a multitude of Interlaken configurations and implementing MAC functions are two important uses for an FPGA. But there is a lot more potential to explore in FPGA design. Savvy designers have started to intelligently optimize their systems by having the FPGA handle functions that are costly to perform in their choice ASIC or ASSP. These functions can include varying levels of packet processing; the implementation of lookaside interfaces to a TCAM; or other features that set a product above the competition.

As technology always has, the world of high bandwidth continues its forward march. This means more throughput and new standards, two areas where Xilinx UltraScale devices are unbeatable. By integrating Interlaken into the UltraScale FPGAs, Xilinx enables existing 100G systems and all future systems to readily adopt new standards, enhance the features of existing architectures and better integrate systems across the board.