ORTHOGONE Technologies Inc.

REFLEX CES Custom Embedded Systems

Ultra-Low Latency FPGA Solution for Electronic Trading and Networking Applications

White Paper – Xilinx UltraScale+ FPGA

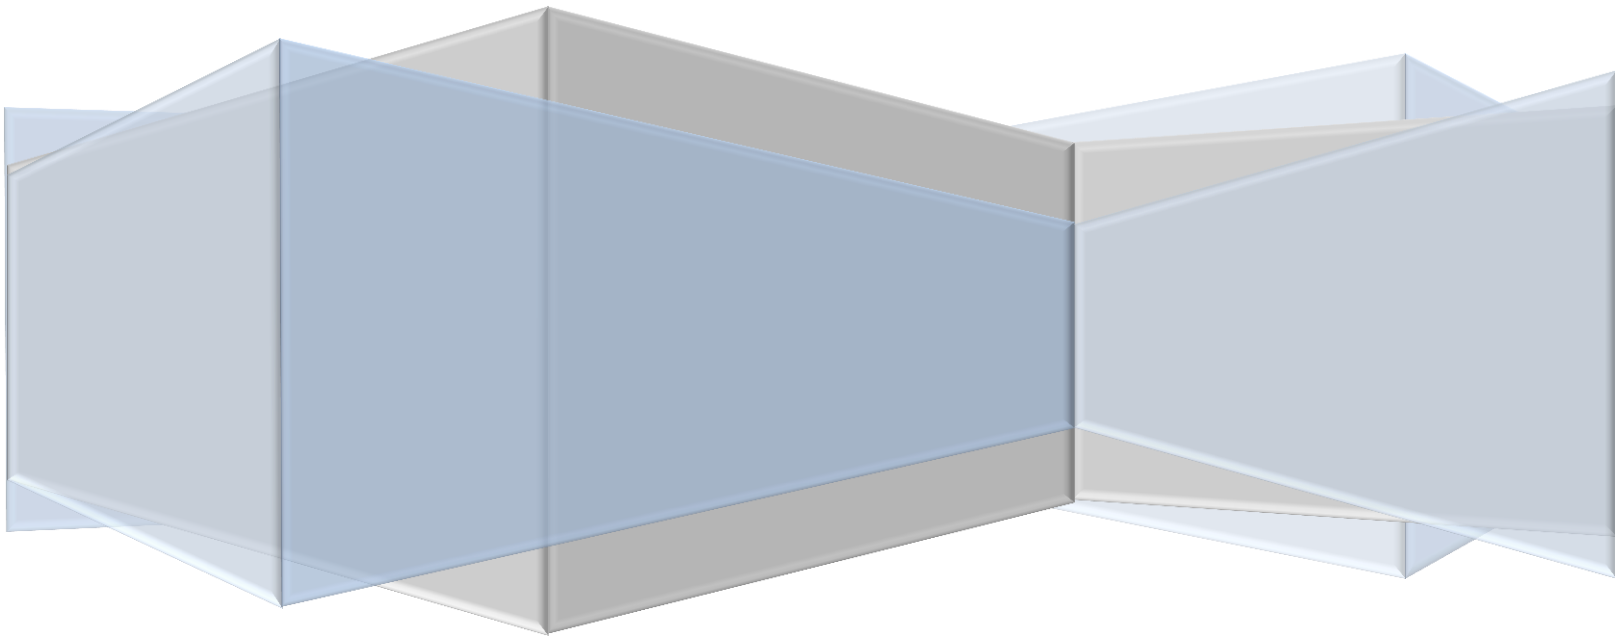

Introduction

REFLEX CES and **Orthogone Technologies** collaborated to offer an ultra-low latency networking solution for electronic trading. The solution is a low-profile PCIe Network Processing FPGA board based on Xilinx Virtex UltraScale+. The board has many salient features, including two QSFP28 ports with integrated ultra-low latency 10G MAC/PCS IP cores, 16GB of DDR4 memory, up to 576Mb of QDR2+ and a PCIe Gen. 3 x 16 interface. This makes it an ideal solution for FPGA accelerated Tick-to-Trade systems or market data feed handlers. Other applications such as High Performance Computing (HPC), networking functions and HW acceleration can also be implemented with this solution.

This white paper will present an overview of the developed solution and the comprehensive reference design that is offered to rapidly enable customers to build their trading infrastructure and networking components. Latency and network throughput performances are presented in detail and backed with lab measurements confirming the performances.

Solution Overview

The XpressVUP is a Low-Profile PCIe Network Processing FPGA Board, designed for HPC, Finance and Networking applications. The board includes a Virtex UltraScale+ XCVU9P-L2FLGB2104E FPGA and is available in two options (Ultra and Turbo), allowing the customers to select amongst various memory configurations. A one PPS (Pulse-per-second) input is available for application requiring inter-board time synchronization.

A photograph of the XpressVUP low-profile PCIe board is presented below.

Figure 1: REFLEX CES XpressVUP Low-Profile PCIe Board

A comprehensive documentation package and a Board Support Package (BSP) are provided to enable rapid deployment of the solution. Additional information is available on REFLEX CES website: https://www.reflexces.com/products-solutions/other-cots-boards/xilinx/xpressvup

10G MAC/PCS/PMA Reference Design Overview

Orthogone has developed an ultra-low latency 10G Ethernet MAC/PCS/PMA solution specifically designed for applications such as Electronic/High Frequency Trading. The core is designed to take advantage of the GTH/GTY transceiver low latency and the fast core fabric offered by the Xilinx US+ FPGA.

A multi-port 10G MAC reference design tailored to the XpressVUP board has been fully validated and characterized. A simple Verilog parameter allows the customers to select the number of 10G ports (1 – 4) to be instantiated per QSFP28. The reference design includes all the key components required to rapidly build low-latency Ethernet interfaces. The GTH/GTY transceivers, the clocks and PLL, the reset sequence, and the IP cores are automatically generated to rapidly enable a very flexible and reliable solution.

The reference design also includes test modules such as packet generator and analyzer, and CPU interface that are provided as clear Verilog code. All the necessary timing constraints and make files are provided to re-compile the reference design and ensure repeatable functionalities and performances.

A block diagram of the MAC/PCS/PMA solution is presented in the next figure (a 4-port configuration is illustrated). The modules highlighted in blue are generally provided as encrypted source code while the others are provided as clear Verilog.

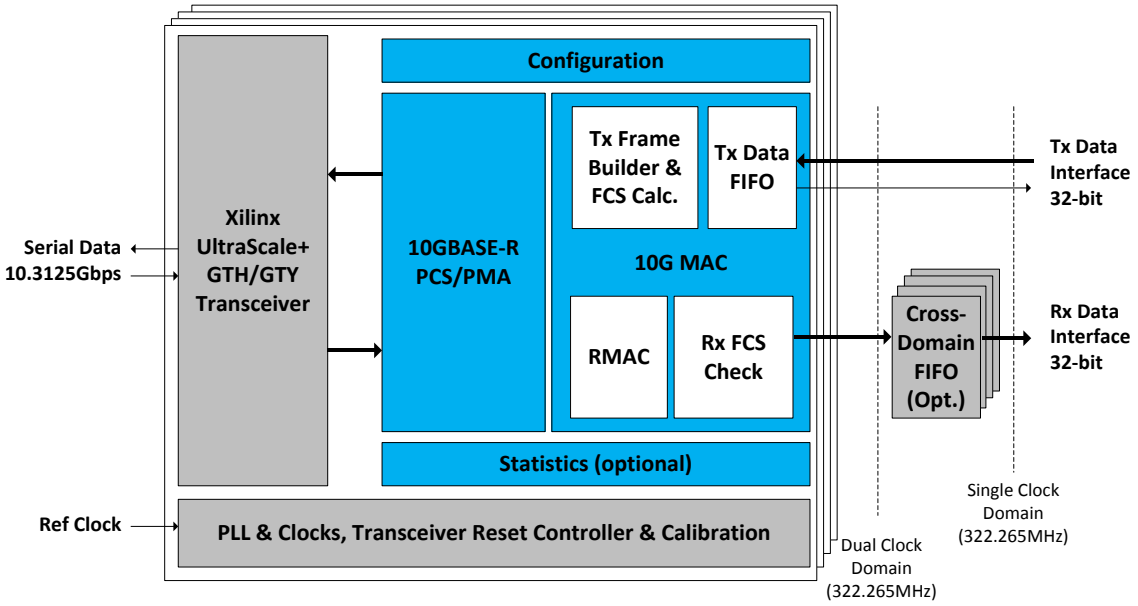

Figure 2: Orthogone Technologies Ultra-Low Latency 10G MAC/PCS/PMA Reference Design

The 10G MAC/PCS/PMA IP core can support full wire line speed with a 64-byte packet length. It also supports back-to-back or mixed length traffic, up to jumbo frame size, with no dropped packets.

The cores have been fully verified through extensive Universal Verification Methodology (UVM) based simulations and fully validated on hardware to comply with the IEEE 802.3 standards.

Thanks to their good timing margin performances, the cores are easy to integrate in the XpressVUP platform. Timing constraints and area placement guidelines are provided with the reference design.

The reference design has been developed using Xilinx Vivado version 2017.2.

Performance Simulation and Measurements

This section presents latency and throughput performances under various conditions.

In a dual clock domain configuration, the FPGA simulations demonstrated that the end-to-end round-trip latency is 67ns. Note that the latency performances remain constant regardless of the Ethernet throughput. The next figure presents the latency results obtained in simulations. The figure shows 67ns round-trip latency for a 64-byte packet.

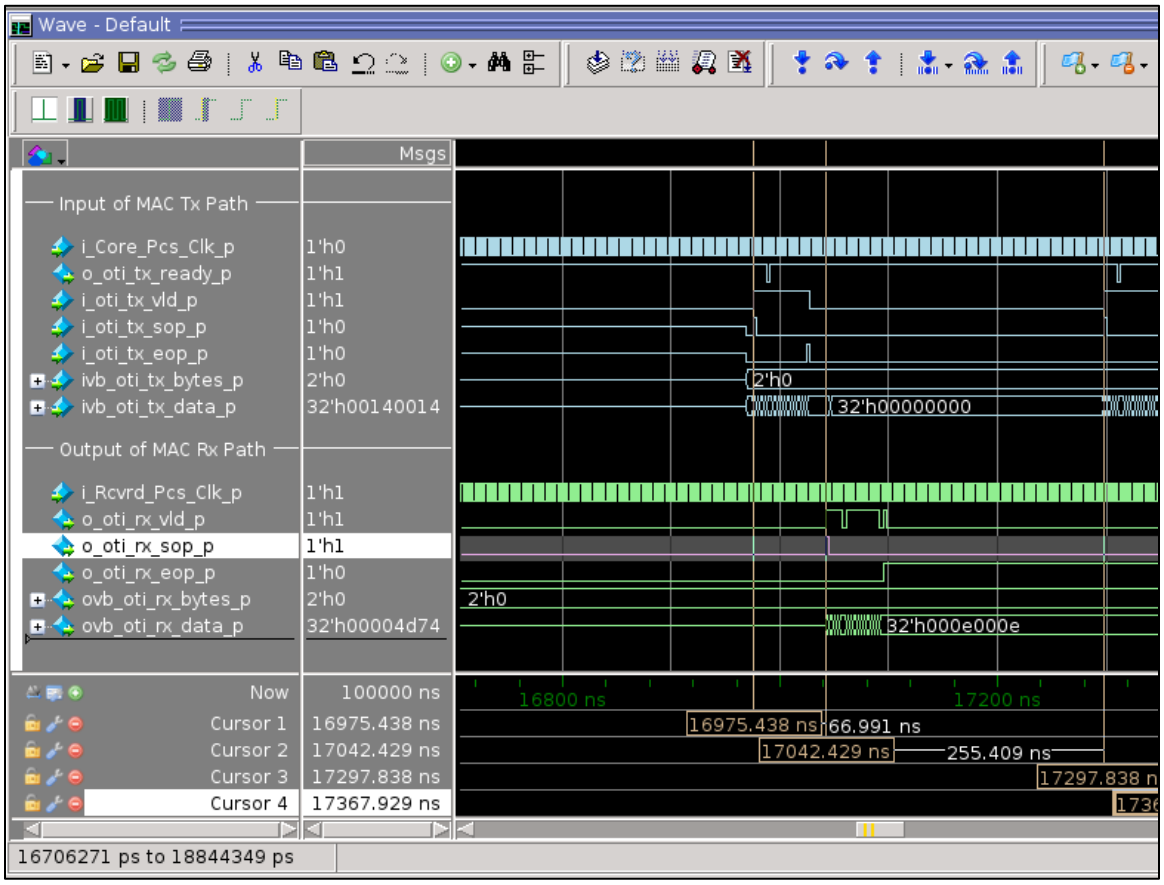

Figure 3: Latency Simulations (64B Frame)

It is important to understand that when a packet is presented to the TMAC it will be transmitted without any back pressure ensuring a deterministic ultra-low latency for any kind of traffic patterns. This is true as long as the TMAC is not driven with traffic rate exceeding 10-Gbps. The latency performance (dual clock domain configuration) is presented in the table below.

Tx Latency [ns]	Rx Latency [ns]	Round-Trip Latency [ns]
32.3	34.6	66.9

Table 1: Pin-to-Pin Latency Performance Breakdown (Simulation)

The latency performances have also been validated on a REFLEX CES XpressVUP platform. In order to accurately measure the latency, an electrical loopback plug is inserted into the QSFP28 cage looping the Tx data back into the Rx port. A timestamp is inserted in every transmit packet and compared with the receive packet timestamp. Due to the timestamp resolution (3.1ns), the latency measurements will always be higher than the real latency (up to 3.1ns extra latency). It is also important to note that some small latency is introduced in the FPGA I/Os, board traces, and QSFP28 loopback module.

The wire to wire latency measurements are nominally 68.3ns to 71.3ns. Latency measurements are performed at full throughput with various packet sizes ranging between 64-bytes and 1522-bytes.

The IP core can sustain full throughput without losing any packets while maintaining ultra-low latency performances. Extensive long count tests have been successfully performed to ensure that the cores do not lose any packets under various traffic conditions.

Resources Utilization

The next table presents the resource utilization for one 10G port. The resource utilization for N 10G ports can be calculated by simply multiplying the 1 port resources by a factor of N.

The resources utilization includes the statistics counters.

	LUTs	Registers	RAMB18/36
One 10G MAC/PCS/PMA IP	3.31k	5.16k	0

Table 2: 10G MAC/PCS IP Resources Utilization

Performance Monitoring – Statistics

The 10G MAC/PCS/PMA core has an option to include statistics counters that can be used to provide network performance visibility. Over twenty statistics parameters are available for the MAC and PCS in each direction.

A complete list of statistics counters is available in the datasheet and user guide document. Contact Orthogone Technologies for additional information

Technical Support

Technical support for the IP core is provided for a one year period following the IP licensing date. Extended maintenance and technical support can optionally be purchased on an annual basis after the first year.

Design customization is also available upon request.

Board Pricing and IP Licensing Options

Contact REFLEX CES for XpressVUP boards pricing information.

The IP core is generally delivered as encrypted source code. Contact Orthogone Technologies to obtain information on licensing fees and discuss other licensing options. Evaluation licenses are available to simulate and test the cores before purchasing. Orthogone also offers additional Ethernet MAC/PCS IP cores supporting other date rates (e.g. 25G, 100G) and FEC options.

Conclusion

REFLEX CES and ORTHOGONE Technologies closely collaborated to develop a high-performance networking solution specifically designed for applications where latency and throughput performances are critical. This low-profile PCIe Network Processing FPGA card with integrated ultra-low latency 10G MAC enables customers to rapidly integrate their FPGA design into a mature and reliable solution that has been thoroughly tested and verified.

ABOUT REFLEX CES and ORTHOGONE TECHNOLOGIES

REFLEX CES

Since 2000, REFLEX CES designs and manufactures custom embedded systems and complex boards based on high-density FPGAs. The company is located in Evry, France and employs 100+ people with R&D and Manufacturing capabilities.

The company provides FPGA COTS boards for several markets, including the Finance market where Ultra Low Latency capability is a key element, and other markets like Networking.

Contact Information

Eric Penain, *Sales & Marketing – Chief Business Officer*
Email : epenain@reflexces.com
Web site : www.reflexces.com

ORTHOGONE TECHNOLOGIES

Orthogone offers a range of electronic design services, including FPGA/ASIC design and verification, HW design, PCB design, software development, and full turnkey electronic product development. The company was founded in 2007 and employs 40+ people at its headquarters in Montreal, Canada. Orthogone Technologies provides ultra-high-performance FPGA/ASIC IP cores used in data centers and networking & communication equipment. The company's IP (Intellectual Property) portfolio includes ultra-low latency and very low gate count Ethernet MAC and PCS solutions operating at 1, 10, 25, 40, and 100-Gbps. Developed by teams of expert FPGA/ASIC designers, Orthogone IP products undergo extensive verification and compliance testing.

Contact Information

Alexandre Raymond, *VP Product Development*
Email :araymond@orthogone.com
Web site : www.orthogone.com