



Extreme Low Latency 10G Ethernet IP Solution


Product Brief (HTK-ELL10G-ETH-FPGA)

The 10Gbps Extreme Low Latency Ethernet IP solution offers a fully integrated IEEE802.3 compliant package for NIC (Network Interface Card) and Ethernet switching applications. This industry leading extreme low latency solution in an extremely small footprint is specifically targeted for demanding financial, high frequency trading, and high port count networking / HPC applications

- Round Trip Latency of 24.9ns + Device Specific Transceiver Latency

As shown in the figure below, the 10Gbps Ethernet IP solution includes:

- Extreme-Low latency Ethernet; Tx = 7.8ns (user data in to line preamble out); Tx = 14.0ns (user data in to line data out), Rx = 10.9ns (line data in to user data out) ALL latency numbers include full FCS generation and checking
- Technology dependent transceiver wrapper for Intel (Altera) and/or Xilinx FPGAs
- Statistics counter block (for RMON and MIB)
- MDIO and I2C cores for external module and optical module status/control



A complete reference design using a L2 (MAC level) packet generator/checker is also included to facilitate quick integration of the Ethernet IP in a user design. A GUI application interacts with the reference design's hardware elements through a UART interface (a PCIe option is also available). An application (with optional basic Linux PCIe driver/API) is also provided for memory mapped read/write access to the internal registers. See Appendix A for details.

MAC and PCS cores are designed to take advantage of high performance 14/16/20nm FPGA devices to achieve extreme low latency with a very small footprint.

As the PCS and transceiver wrapper is included with the Ethernet IP solution, the line side directly connects the 10.3125Gbps FPGA transceiver to the optical module (SFP+, XFP etc).

Ethernet IP solution implements two user (application) side interfaces. The register configuration and control port can be 32-bit AXI4-Lite or Avalon-MM interface. Depending upon the application layer, user can select a configurable bit width AXI-4 Streaming or Avalon Streaming bus to interface with the MAC block.

10Gbps Ethernet IP supports advanced features like per-priority pause frames (compliant with 802.3bd specifications) to enable Converged Enhanced Ethernet (CEE) applications like data center bridging that employ IEEE 802.1Qbb Priority Flow Control (PFC) to pause traffic based on the priority levels.

Features Overview

- Implements the full 802.3 specification with preamble/SFD generation, frame padding generation, CRC generation and checking on transmit and receive respectively.
- Implements 802.3bd specification with ability to generate and recognize PFC pause frames.
- Implements reconciliation sublayer functionality with start and terminate control characters alignment, error control character and fault sequence insertion and detection.
- Deficit Idle Count (DIC) mechanism to ensure data rates of 10Gbps at the transmit interface.
- Optional padding of frames if the size of frame is less than 64 bytes.
- Implements XON and XOFF Pause Frame (802.3 Annex 31A) generation and termination providing flow control.
- Pause frame generation additionally controllable by user application offering flexible traffic flow control.
- Support for VLAN tagged frames according to IEEE 802.1Q.
- Supports any type of Ethernet Frames such as SNAP / LLC, Ethernet II/DIX or IP traffic.
- Discards frames with mismatching destination address on receive (except Broadcast and Multicast frames).
- Programmable Promiscuous mode support to omit MAC destination address checking on receive path.
- Optional multicast address filtering with 64-bit Hash Filtering table providing imperfect filtering to reduce load on higher layers.
- High speed CRC-32 generation and checking.

- Optional insertion of error control character in transmitted frame data.
- Programmable frame maximum length providing support for any standard or proprietary frame length (e.g. 9K-Bytes Jumbo Frames).
- Status signals available with each Frame on the user interface providing information such as frame length, VLAN frame type indication and error information.
- Optional padding termination on RX path for NIC applications or forwarding of unmodified data to the user interface.
- Statistics indicators for frame traffic as well as errors (alignment, CRC, length) and pause frames.
- Configurable bit width Avalon-ST or AXI4 Streaming user interface
- Implements statistics and event signals providing support for 802.3 basic and mandatory managed objects as well as IETF Management Information Database (MIB) package (RFC 2665) and Remote Network Monitoring (RMON) required in SNMP environments.
- Implements 10GBase-R PCS core compliant with IEEE 802.3-2008 Specifications.
- Implements 64b/66b encoding/decoding for transmit and receive PCS using 802.3-2008 specified control codes.
- Implements 10G scrambling/descrambling using 802.3-2008 specified polynomial $1 + x^{39} + x^{58}$.
- Implements 66-bit block synchronization state machine as specified in 802.3-2008 specifications.
- Implements Bit Error Rate (BER) monitor for monitoring excessive error ratio. In addition, the core implements various status and statistics required by the IEEE 802.3-2008 such as block synchronization status and test mode error counter.

Deliverables

- Encrypted MAC and PCS RTL for simulation and synthesis
- Encrypted L2 packet generator and checker RTL for simulation and synthesis
- Source code RTL (Verilog) for top level Ethernet wrappers to allow for user specific customizations.
- Technology specific transceiver wrappers for the selected device family
- Source code RTL (Verilog) for AXI4-Lite and Avalon-MM arbiters and address decoders
- Constraint files and synthesis scripts for design compilation
- Linux based APIs/tools to access core configuration and statistics registers
- Design guide(s) and user manual(s)

Links

<http://hiteksys.com/fpga-ip-cores/10g-extreme-low-latency-ethernet>

For sales or more information:



Hitek Systems LLC
 Phone: +1-301-528-8074
 Email: sales@hiteksys.com

Resource Utilization

The utilization summary of the 10G Ethernet solution is given in following tables. The utilization numbers are best in class as compared to other available 10G Ethernet cores with comparable feature set.

The Ethernet solution has been fully verified on different hardware platforms for Intel (Altera) / Xilinx FPGAs and has also been verified for interoperability with other 10G capable devices.

10G ELL Ethernet IP - Resource Usage for Intel (Altera) Devices

<i>Device</i>	<i>Priority Flow Control (PFC)</i>	<i>COMB. ALUTs</i>	<i>Registers</i>	<i>Memory M20K</i>
Arria 10	No	1,864	1,747	0
	Yes	2,017	2,025	0

Note:

- Support wrapper with MAC and PCS Registers adds additional 1,286 Comb. ALUTs and 1,453 Registers.
- Register based RMON statistics block adds additional 2003 Comb. ALUTs and 1808 registers.

10G ELL Ethernet IP - Resource Usage for Xilinx Devices

<i>Device</i>	<i>Priority Flow Control (PFC)</i>	<i>COMB. ALUTs</i>	<i>Registers</i>	<i>BRAMs</i>
Ultrascale	No	1,738	1,292	18K=0 ; 36K=0
	Yes	1,864	1,583	18K=0 ; 36K=0
Ultrascale+	No	1,742	1,292	18K=0 ; 36K=0
	Yes	1,853	1,583	18K=0 ; 36K=0

Note:

- Support wrapper with MAC and PCS Registers adds additional 908 Comb. ALUTs and 1,292 Registers.
- Register based RMON statistics block adds additional 2003 Comb. ALUTs and 1808 registers.

Performance (Tx And Rx Latency)

The performance of the 10G Ethernet solution is represented here in terms of individual latencies of transmit and receive paths, i.e. the time between the first bit of data input at 10G Ethernet MAC and the first bit of data output at PCS-Transceiver interface.

<i>Technology</i>	<i>User Interface</i>	<i>MAC+PCS Latency (ns)</i>	
		<i>Tx</i>	<i>Rx</i>
Xilinx/Intel (Altera)	Core with optimal user interface	14.0	10.9
	Core with nominal user interface	17.1	14.0

Latency Definitions

- The latency of transmit 10G-ELL core is defined as delay from first data bit in at the user interface to the first data bit out on the technology dependent transceiver interface. Some competitors provide latency values from first data bit in to preamble so this latency is also provided.
- The latency of receiver 10G-ELL core is defined as delay from first data bit in at the technology dependent transceiver interface to first data bit out at the MAC user interface.

