



Ultra-Low Latency 10G Ethernet IP Solution


Product Brief (HTK-ULL10G-ETH-32-FPGA)

The 10Gbps 32-bit Ethernet IP solution offers a fully integrated IEEE802.3-2015 compliant package for NIC (Network Interface Card) and Ethernet switching applications. This industry leading ultra-low latency solution is specifically targeted for demanding financial, high frequency trading and HPC applications.

- Round Trip Latency of 52.7ns + Device Specific Transceiver Latency

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

- Ultra-Low latency MAC; Tx = 12.4ns , Rx = 15.5ns; (32-bit user interface mode with FCS generation and checking)
- Ultra-Low latency PCS; Tx = 12.4ns , Rx = 12.4ns;
- Technology dependent transceiver wrapper for 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 with 32-bit data path to take advantage of high performance fabrics of the 28nm and 20nm FPGAs. This implementation approach also delivers industry's lowest latency and low area 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 32-bit or 64-bit AXI-4 Streaming or Avalon Streaming bus to interface with the MAC block.

10Gbps Ethernet IP supports advanced features like perpriority 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 (Mac Core)

- 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.
- PCS layer XGMII interface implemented as 32-bit (double data rate) DDR interface to operate at 10Gbps for direct interface to 10GBase-R, XAUI and RXAUI cores.
- 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 fully automated XON and XOFF Pause Frame (802.3 Annex 31A) generation and termination providing flow control without user application intervention. Non PFC Mode only with use of support wrapper containing Rx User Interface Buffer.
- Pause frame generation additionally controllable by user application offering flexible traffic flow control.
- Support for VLAN tagged frames according to IEEE 802.1Q.
- Support 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.
- 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.
- PCS layer XGMII interface implemented as 32-bit (double data rate) DDR interface to operate at 10Gbps for direct interface to 10GBase-R, XAU1 and RXUAI cores.
- 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 fully automated XON and XOFF Pause Frame (802.3 Annex 31A) generation and termination providing flow control without user application intervention. Non PFC Mode only with use of support wrapper containing Rx User Interface Buffer.
- Pause frame generation additionally controllable by user application offering flexible traffic flow control.
- Support for VLAN tagged frames according to IEEE 802.1Q.
- Support 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.

PCS Core Features

- Implements 10GBase-R PCS core compliant with IEEE 802.3-2008 Specifications.
- Implements a 32-bit XGMII interface to operate at 10Gbps for 10G Ethernet.
- 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.
- Automatic clock compensation without the need for Inter Packet Gap (IPG) insertion/deletion that is achieved with use of support wrapper containing Rx XGMII Buffer.
- Implements gear-box logic to connect to 32-bit transceivers for line side. The 32-bit interface operates at the transceiver reference clock.
- 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.

- Implements optional XGMII remote loopback to loopback data received from Rx PCS back to Tx PCS with use of support wrapper containing XGMII remote loop-back Buffer.

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-ultra-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 both Altera and Xilinx FPGAs and has also been verified for interoperability with other 10G capable devices.

10G ULL Ethernet IP - Resource Usage for Xilinx Devices

<i>Device</i>	<i>User Interface (AXI4)</i>	<i>Priority Flow Control (PFC)</i>	<i>Slice LUTs</i>	<i>Slice Registers</i>	<i>BRAMs</i>
UltraScale/ UltraScale+	32-Bit	No	2,437	1,704	18K = 0; 36K = 0
		Yes	2,534	2,129	18K = 0; 36K = 0
7-Series	32-Bit	No	2,455	1,703	18K = 0; 36K = 0
		Yes	2,565	2,128	18K = 0; 36K = 0
Note:					
<ul style="list-style-type: none"> Support wrapper with MAC and PCS Registers adds additional 529 Slice LUTs and 1276 Slice Registers. Register based RMON statistics block adds additional 1948 Slice LUTs and 1807 Slice Registers. 					

10G ULL Ethernet IP - Resource Usage for Altera Devices

<i>Device</i>	<i>User Interface (Avalon)</i>	<i>Priority Flow Control (PFC)</i>	<i>COMB. ALUTs</i>	<i>Registers</i>	<i>Memory M20K</i>
Arria 10	32-Bit	No	2,346	1,708	0
		Yes	2,415	2,135	0
Stratix V	32-Bit	No	2,351	1,703	0
		Yes	2,428	2,128	0
Note:					
<ul style="list-style-type: none"> Support wrapper with MAC and PCS Registers adds additional 596 Comb. ALUTs and 1093 Registers. Register based RMON statistics block adds additional 2003 Comb. ALUTs and 1808 registers. 					