



Ultra-Low Latency 25G/10G Ethernet IP


Product Brief (HTK-ULL25G-ETH-64-FPGA)

The 25G/10Gbps 64-bit Ethernet IP solution offers a fully integrated IEEE P802.3by compliant package for NIC (Network Interface Card) and Ethernet switching applications. This ultra-low latency solution is specifically targeted for demanding financial, high frequency trading and HPC applications.

- Round Trip Latency of 59.6ns + Device Specific Transceiver Latency

As shown in the figure below, the 25Gbps Ethernet IP includes:

- Ultra-Low latency MAC; Tx = 17.4ns , Rx = 17.4ns; (64-bit user interface mode)
- Ultra-Low latency PCS; Tx = 12.4ns , Rx = 12.4ns; (64-bit user interface mode)
- 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 64-bit data path to take advantage of high performance fabrics of the 28nm, 20nm and 16nm 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 25.78125Gbps FPGA transceiver to the optical module (SFP28-SR). 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 64-bit or 128-bit AXI-4 Streaming or Avalon Streaming bus to interface with the MAC block. 25G/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 (MAC Core Features)

- Implements the full P802.3by 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 25GMII interface implemented as 64-bit (single data rate) SDR.
- Deficit Idle Count (DIC) mechanism to ensure data rates of 25Gbps 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.

- Optional prevention of CRC appending in frame data by MAC to allow CRC to be pre-embedded in frame data by user application.
- Optional insertion of error control character in transmitted frame data.
- Optional forwarding of the CRC field to user application interface.
- 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.
- Optional internal 25GMII Loop-back with use of support wrapper containing 25GMII Loop-back Buffer.
- Statistics indicators for frame traffic as well as errors (alignment, CRC, length) and pause frames.
- Altera Avalon or Xilinx AXI4 interface compliant user interface with use of support wrapper containing Tx and Rx User Interface Buffers.
- 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.

Features (PCS Core Features)

- Implements 25GBase-R PCS core compliant with IEEE P802.3by Specifications.
- Implements a 64-bit 25GMII interface to operate at 25Gbps for 25G Ethernet.
- Implements 64b/66b encoding/decoding for transmit and receive PCS using 802.3-2015 specified control codes.
- Implements 25G scrambling/descrambling using 802.3-2008 specified polynomial $1 + x^{39} + x^{58}$.
- Ability to generate scrambled idle test pattern compliant with IEEE P802.3by Specifications.
- 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 25GMII Buffer.
- Implements gear-box logic to convert 66-bit blocks to 64-bit for line side. The 64-bit interface operates at the transceiver user clock.
- Implements Bit Error Rate (BER) monitor compliant with IEEE P802.3by Specifications. 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 25GMII remote loopback to loopback data received from Rx PCS back to Tx PCS with use of support wrapper containing 25GMII 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 reference design compilation
- Linux based APIs/tools to access core configuration and statistics registers
- Design guide(s) and user manual(s)

Product Ordering Code

(HTK-ULL25G-ETH-64-FPGA)

Links

<https://hiteksys.com/fpga-ip-cores>

<https://hiteksys.com/fpga-ip-cores/25g-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 25G Ethernet solution is given in following tables. The utilization numbers are best in class as compared to other available 25G 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 25G capable devices.

25G 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+	64-Bit	No	3,510	2,975	18K = 0; 36K = 0
		Yes	3,610	3,412	18K = 0; 36K = 0
7-Series	64-Bit	No	3,516	2977	18K = 0; 36K = 0
		Yes	3,604	3,414	18K = 0; 36K = 0

Note:

- Support wrapper with MAC and PCS Registers adds additional 681 Slice LUTs and 1065 Slice Registers.
- Register based RMON statistics block adds additional 1948 Slice LUTs and 1807 Slice Registers.
- RS-Forward Error Correction adds additional 37,553 Slice LUTs and 21,997 Slice Registers.

25G 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	64-Bit	No	3,683	3,114	0
		Yes	3,743	3,551	0
Stratix V	64-Bit	No	3,670	3,127	0
		Yes	3,742	3,565	0

Note:

- Support wrapper with MAC and PCS Registers adds additional 701 Comb. ALUTs and 1114 Registers.
- Register based RMON statistics block adds additional 2003 Comb. ALUTs and 1808 Registers.
- RS-Forward Error Correction adds additional 30,307 Comb. ALUTs and 27,555 Registers.