

Ultra-Low Latency 25G Ethernet IP Solution



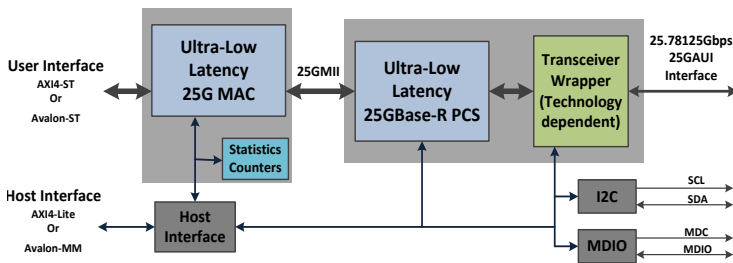
Product Brief (HTK-ULL25G-ETH-64-FPGA)

The 25Gbps 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.

25Gbps 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

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.
- 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.

Licensing and Maintenance

- ***NO*** yearly maintenance fees for upgrades and bug fixes
- Basic core licensing for a single vendor (either Xilinx or Altera) compiled (synthesized netlist) binary
- Other licensing options include:
 - Vendor and device family agnostic source code (Verilog) license
 - A **low cost, board locked** license for low budget prototyping (upgradeable to full license)

Contact and Sales Information

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

Ultra-Low Latency 25G Ethernet IP Solution

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



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	3,513	3,186	18K = 0; 36K = 0
		Yes	3,612	3,623	18K = 0; 36K = 0
7-Series	64-Bit	No	3,516	3,188	18K = 0; 36K = 0
		Yes	3,603	3,625	18K = 0; 36K = 0

Note:

- Support wrapper with MAC and PCS Registers adds additional 827 Slice LUTs and 893 Slice Registers.
- Register based RMON statistics block adds additional 1948 Slice LUTs and 1807 Slice Registers.
- RS-Frame Error Correction coming soon.

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	3,733	3,290	0
		Yes	3,810	3,731	0
Stratix V	64-Bit	No	3,741	3,297	0
		Yes	3,815	3,735	0

Note:

- Support wrapper with MAC and PCS Registers adds additional 670 Comb. ALUTs and 1010 Registers.
- Register based RMON statistics block adds additional 2003 Comb. ALUTs and 1808 registers.
- RS-Frame Error Correction coming soon.

Performance (Tx And Rx Latency)

The performance of the 25G 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 25G Ethernet MAC and the first bit of data output at PCS-Transceiver interface. These numbers will change with the change in programmable threshold levels used for reading the user interface FIFOs. For the latencies given here, the thresholds for both transmit and receive User FIFOs were set to minimum possible values for correct operation. Data path latency is also dependent upon the type of user interface FIFO used in the design. FIFO implementation can either be a SCFIFO (Single Clock FIFO, when the MAC and application clock are same) or can be a DCFIFO (Dual Clock FIFO, when the MAC and application clock are different). Following table lists the latencies for various user interface options.

<i>Technology</i>	<i>User Interface</i>	<i>MAC+PCS Latency (ns)</i>	
		<i>Tx</i>	<i>Rx</i>
Xilinx	64-bit (NO FIFO)	29.8	29.8
	64-bit (DCFIFO)	52.2	77.0
	128-bit (SCFIFO)	59.6	82.0
	128-bit (DCFIFO)	67.1	91.9
Altera	64-bit (NO FIFO)	29.8	29.8
	64-bit (DCFIFO)	52.2	77.0
	128-bit (SCFIFO)	59.6	82.0
	128-bit (DCFIFO)	67.1	91.9

Deliverables

- Compiled synthesizable binaries or encrypted RTL for the MAC and PCS cores
- Source code RTL (Verilog) for I2C, MDIO, RMON and Register-File blocks
- Self-checking behavioral models and test benches for simulation
- Constraint files and synthesis scripts for design compilation
- A complete UART/PCIe host interface based reference design with:
 - Top level wrapper (source files, Verilog) for user specific customizations
 - Source files (Verilog) for the PCIe application layer
 - Binaries for a basic L2 packet generator and checker
 - UART and command interpreter blocks with the UART host interface
 - PCIe driver/API (source files, C) for Linux with the optional PCIe host interface
 - GUI application (Linux only for PCIe, Linux and Windows for UART) for interfacing to the reference design
- Design guide(s) and user manuals

A. Reference Design Details

A.1 Overview

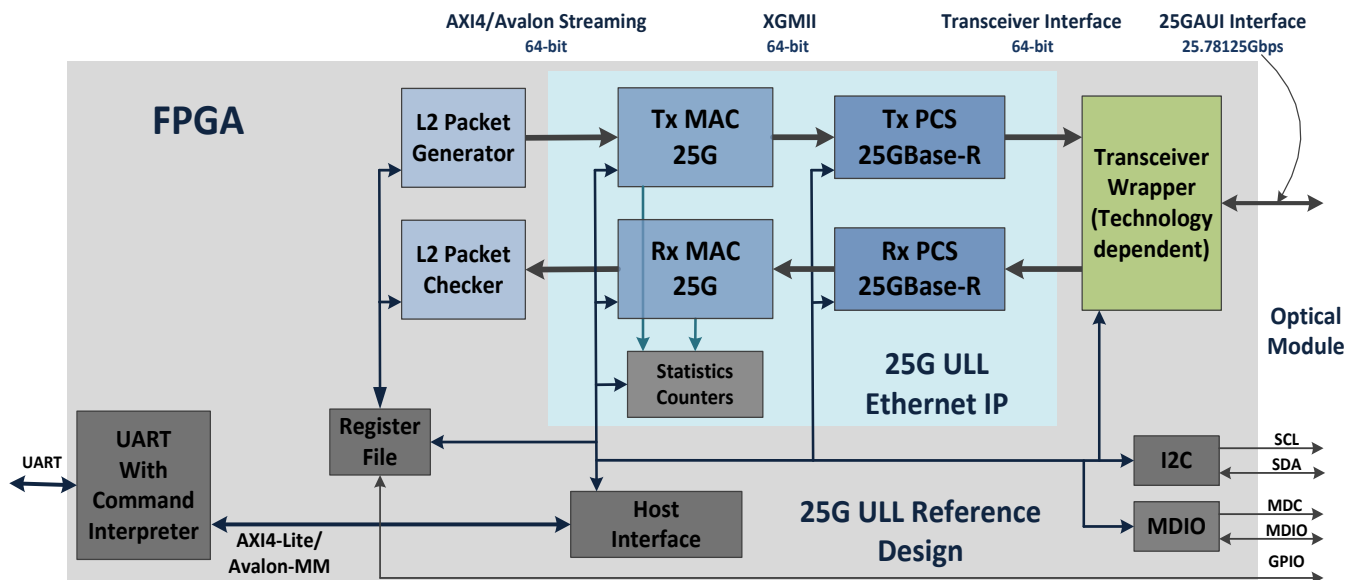
A 25Gbps reference design is included as part of the IP deliverable to facilitate quick L1 and L2 layer testing and verification of the 25Gbps Ethernet on target platform. The capability to run the L1 PRBS pattern and configure each transceiver independently can be for used for a fast module bring-up in the lab and can also be used for factory diagnostics.

The UART (normally through an onboard USB-to-UART converter chip) based 25G Ethernet reference design can be seamlessly ported to various COTS FPGA networking and evaluation modules (see section for the list of verified modules). A GUI application controls the register read/writes to the FPGA through a UART core with integrated command interpreter. Both Linux and Windows platforms are supported for the UART based interface control.

This reference design can also be used on custom embedded design where the FPGA connects to the host processor via a PCIe interface. For the PCIe control interface, GUI application is hosted on a Linux platform (as PCIe driver/API is provided for Linux OS only).

A.2 Functional Description

Following figure shows the connectivity and the elements of the 25Gbps Ethernet IP reference design. Usually the UART interface from the FPGA connects to an external (can be on the same module as well) USB-UART converter. A Linux or Windows host (through a USB port) running the GUI application is used to configure and control the 25G Ethernet. I2C and GPIO interfaces included in the reference design can be used to control any optical module on the target platform including the SFP28 and QSFP28 compliant modules.



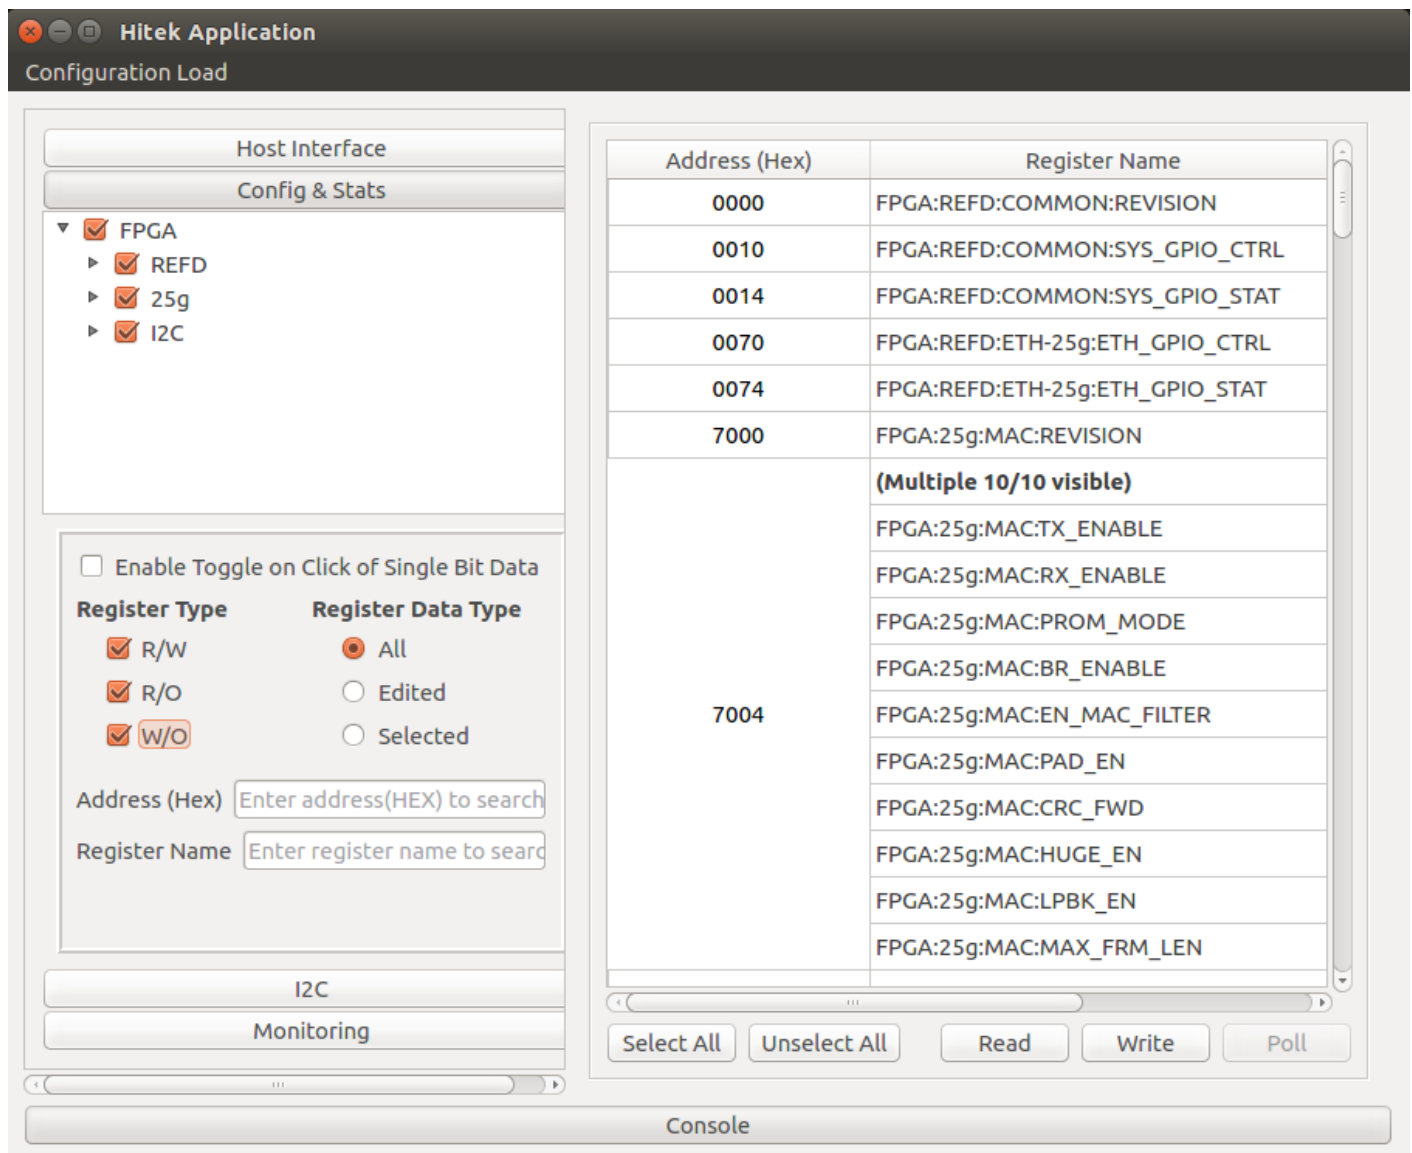
For L1 (physical layer verification and testing) GUI application provides an interface to independently control and configure 25.78125Gbps transceiver used for 25G Ethernet transport. User can configure the transceiver to run various PRBS pattern and configure various transceiver parameters like transmit voltage, transmit pre-emphasis, receive equalization and receive gain.

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

For L2 testing, GUI application uses the 25Gbps packet generator/checker inside the FPGA to generate and check MAC frames up to full line rate. Packet generator supports a basic rate control mechanism to control the packet/data rate on the interface. Generator can be configured for fixed size as well as pseudo random packet size packet transmission. An incrementing counter is used as payload for the MAC frames. Checker on the receive side verifies the payload of receive MAC frames and reports error in the payload.

A comprehensive set of transmit and receive counters in the MAC core provide a detailed view of the packet statistics including various error types.

Following is a snapshot for the GUI application for the L2 packet test results screen.



The screenshot shows the 'Hitek Application' window with a 'Configuration Load' header. The interface is divided into several sections:

- Host Interface:** A tabbed interface with 'Config & Stats' selected.
- Config & Stats:** A tree view showing expanded options:
 - FPGA
 - REFD
 - 25g
 - I2C
- Options:**
 - Enable Toggle on Click of Single Bit Data
 - Register Type:**
 - R/W
 - R/O
 - W/O
 - Register Data Type:**
 - All
 - Edited
 - Selected
- Search:**
 - Address (Hex):
 - Register Name:
- Navigation:** Buttons for 'I2C' and 'Monitoring' are visible at the bottom of the left panel.
- Register List:** A table with columns 'Address (Hex)' and 'Register Name'.

Address (Hex)	Register Name
0000	FPGA:REFD:COMMON:REVISION
0010	FPGA:REFD:COMMON:SYS_GPIO_CTRL
0014	FPGA:REFD:COMMON:SYS_GPIO_STAT
0070	FPGA:REFD:ETH-25g:ETH_GPIO_CTRL
0074	FPGA:REFD:ETH-25g:ETH_GPIO_STAT
7000	FPGA:25g:MAC:REVISION
	(Multiple 10/10 visible)
	FPGA:25g:MAC:TX_ENABLE
	FPGA:25g:MAC:RX_ENABLE
	FPGA:25g:MAC:PROM_MODE
	FPGA:25g:MAC:BR_ENABLE
7004	FPGA:25g:MAC:EN_MAC_FILTER
	FPGA:25g:MAC:PAD_EN
	FPGA:25g:MAC:CRC_FWD
	FPGA:25g:MAC:HUGE_EN
	FPGA:25g:MAC:LPBK_EN
	FPGA:25g:MAC:MAX_FRM_LEN
- Actions:** Buttons for 'Select All', 'Unselect All', 'Read', 'Write', and 'Poll' are located below the register list.
- Console:** A console area is visible at the bottom of the application window.

A.3 Ported/Validated Modules List

1. *HiTech Global HTG-828 and HTG-9100 Modules: Xilinx Virtex UltraScale* FPGA with 25Gbps transceivers, Optical module with CFP4 and QSFP28 interfaces. Parallel MTP/MTO to quad SFP+ breakout cable used for 25Gbps interface.
(http://hitechglobal.com/Boards/Virtex_UltraScale_100G_NIC.htm)
(http://hitechglobal.com/Boards/QuadCFP4_QSFP28.htm)
2. *HiTech Global HTG-728 Module; Xilinx Virtex-7 H580T* FPGA with 25Gbps transceivers, Optical module with CFP2 and CFP4 interfaces. Parallel MTP/MTO to quad SFP+ breakout cable used for 25Gbps interface.
(<http://www.hitechglobal.com/Boards/x16PCIExpress-Gen3.htm>)