

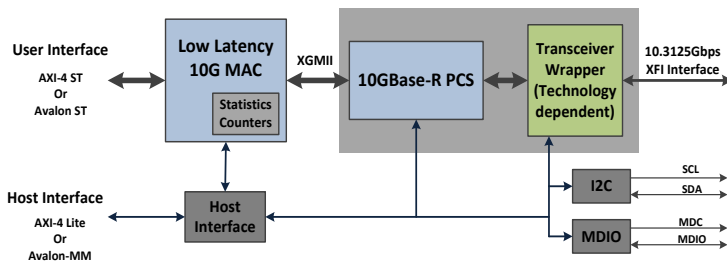
Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-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 extremely low latency solution is specifically targeted for demanding financial, high frequency trading and HPC applications. As shown in the figure below, the 10Gbps Ethernet IP includes:

- Low latency MAC; Tx = 50.0ns , Rx = 70.4ns; (32-bit user interface mode)
- Low latency PCS; Tx = 77.1ns , Rx = 121.3ns; (32-bit user interface mode)
- Flexible 10GBase-R PCS options with XFI interface for direct SFP+/XFP attachment
- 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 core is designed with 32-bit data path operating at 312.5MHz 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 is a 32-bit AXI4-Lite or Avalon-MM interface. Depending upon the application layer, user can select a 32-bit @312.5MHz or 64-bit @ 156.25MHz 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

MAC Core Features

- 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.
- Multiple user interface options for the MAC data path: AXI-4 or Avalon streaming with 32-bit data path at 312.5MHz or 64-bit data path at 156.25MHz
- PCS layer XGMII interface implemented as 64-bit (single data rate) SDR interface at 156.25MHz for direct interface to 10GBase-R, XAU1 and RXAU1 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.
- 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.

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



- 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 XGMII Loop-back.
- Statistics indicators for frame traffic as well as errors (alignment, CRC, length) and pause frames.
- Altera Avalon or Xilinx AXI4 interface compliant user (FIFO) interface.
- Transmit and Receive FIFOs with configurable depths having a default depth of 1KB/512B (128 64-bit/32-bit words) each, according to user interface bus width.
- 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.
- For devices without the support for the 10.3125Gbps transceivers, Ethernet IP solution can be configured to operate with vendor specific XAUI and RXAUI cores.
- Implements Inter Packet Gap (IPG) insertion/deletion for clock compensation while maintaining a minimum of 5 bytes IPG.
- Implements gear-box logic to convert 66-bit blocks to 40-bit for line side. The 40-bit interface operates at the transceiver reference clock. There is an option with Xilinx 7-series devices to use the gear-box logic available in the transceiver. In this case, the gear-box logic in PCS is bypassed and the 66-bit blocks are presented at line side operating 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.

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

Phone: +1-301-528-2244

Email: info@mantaro.com

PCS Core Features

- Implements 10GBase-R PCS core compliant with IEEE 802.3-2008 Specifications.
- Implements a 64-bit XGMII interface to operate at 156.25MHz 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.

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



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 LL Ethernet IP - Resource Usage for Xilinx Devices

<i>Device</i>	<i>User Interface (AXI4-ST)</i>	<i>Priority Flow Control (PFC)</i>	<i>Slice LUTs</i>	<i>Slice Registers</i>	<i>BRAMs</i>
UltraScale/ UltraScale+	32-Bit	No	5,365	5,822	18K = 2; 36K = 3
		Yes	5,567	6,344	18K = 2; 36K = 3
	64-Bit	No	5,570	6,079	18K = 0; 36K = 5
		Yes	5,773	6,601	18K = 0; 36K = 5
7-Series	32-Bit	No	5,523	5,822	18K = 2; 36K = 3
		Yes	5,728	6,344	18K = 2; 36K = 3
	64-Bit	No	5,705	6,079	18K = 0; 36K = 5
		Yes	5,926	6,601	18K = 0; 36K = 5
Virtex 6	32-Bit	No	6,509	5,781	18K = 2; 36K = 3
		Yes	6,708	6,303	18K = 2; 36K = 3
	64-Bit	No	6,696	6,104	18K = 0; 36K = 5
		Yes	6,902	6,626	18K = 0; 36K = 5
Note: <ul style="list-style-type: none">These utilization numbers include MAC and PCS Register files.Register based RMON statistics block adds additional 1948 LUTs and 1807 registers.					

10G LL Ethernet IP - Resource Usage for Altera Devices

<i>Device</i>	<i>User Interface Width</i>	<i>Priority Flow Control (PFC)</i>	<i>COMB. ALUTs</i>	<i>Registers</i>	<i>Memory</i>
Arria 10	Avalon, 32-Bit	No	4,909	6,057	M20K = 8
		Yes	5,114	6,607	M20K = 8
	Avalon, 64-Bit	No	4,942	6,261	M20K = 10
		Yes	5,174	6,807	M20K = 10
Stratix V	Avalon, 32-Bit	No	4,906	6,035	M20K = 8
		Yes	5,111	6,586	M20K = 8
	Avalon, 64-Bit	No	4,939	6,243	M20K = 10
		Yes	5,171	6,766	M20K = 10
Stratix IV	Avalon, 32-Bit	No	4,939	5,656	M9K = 9
		Yes	5,135	6,176	M9K = 9
	Avalon, 64-Bit	No	4,977	5,839	M9K = 11
		Yes	5,204	6,360	M9K = 11
Note: <ul style="list-style-type: none">• <i>These utilization numbers include MAC and PCS Register files.</i>• <i>Register based RMON statistics block adds additional 2003 LUTs and 1808 registers.</i>					

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



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. 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	32-bit (SCFIFO)	125.1	191.7
	32-bit (DCFIFO)	132.9	199.9
	64-bit (SCFIFO)	129.9	202.9
	64-bit (DCFIFO)	137.6	233.6
Altera	32-bit (SCFIFO)	120.4	190.0
	32-bit (DCFIFO)	132.1	194.3
	64-bit (SCFIFO)	129.7	296.7
	64-bit (DCFIFO)	137.5	227.3

Deliverables

- Compiled synthesizable binaries or encrypted RTL for the MAC core
- 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 optional UART host interface
 - PCIe driver/API (source files, C) for Linux with the optional PCIe interface
 - GUI application (Linux only for PCIe, Linux and Windows for UART) for interfacing to the reference design
- Design guide(s) and user manuals

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



Ported/Validated Modules List

1. *HiTech Global HTG-K800, HTG-K816 and HTG-830; Xilinx Virtex Ultrascale and Kintex Ultrascale* FPGAs; Interface through FMC (HTG-FMC-X4SFP+) and Z-Ray (HTG-ZR-X3SFP+) SFP+ modules.
(<http://hitechglobal.com/Boards/Kintex-UltraScale.htm>)
(http://www.hitechglobal.com/Boards/Kintex_UltraScale_half-size_PCIe.htm)
(<http://hitechglobal.com/boards/Virtex-UltraScale-FPGA.htm>)
2. *HiTech Global HTG-K700; Xilinx Kintex-7* FPGA; Interface through FMC (HTG-FMC-X4SFP+) SFP+ module
(http://www.hitechglobal.com/Boards/Kintex-7_PCIe.htm)
3. *Xilinx VC707 (Xilinx Virtex-7 485) and Xilinx KC705 (Xilinx Kintex-7 325)* Evaluation modules
4. *HiTech Global HTG-616; Xilinx Virtex-6 HXT FPGA*, with integrated QSFP+ and SFP+ interfaces
(<http://hitechglobal.com/boards/x16pciexpress.htm>)
5. *HiTech Global HTG-510; Altera Stratix-V* FPGA, with integrated QSFP+ and SFP+ interfaces
(http://hitechglobal.com/Boards/Stratix-V_PCIExpress.htm)

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



A. Reference Design Details

A.1 Overview

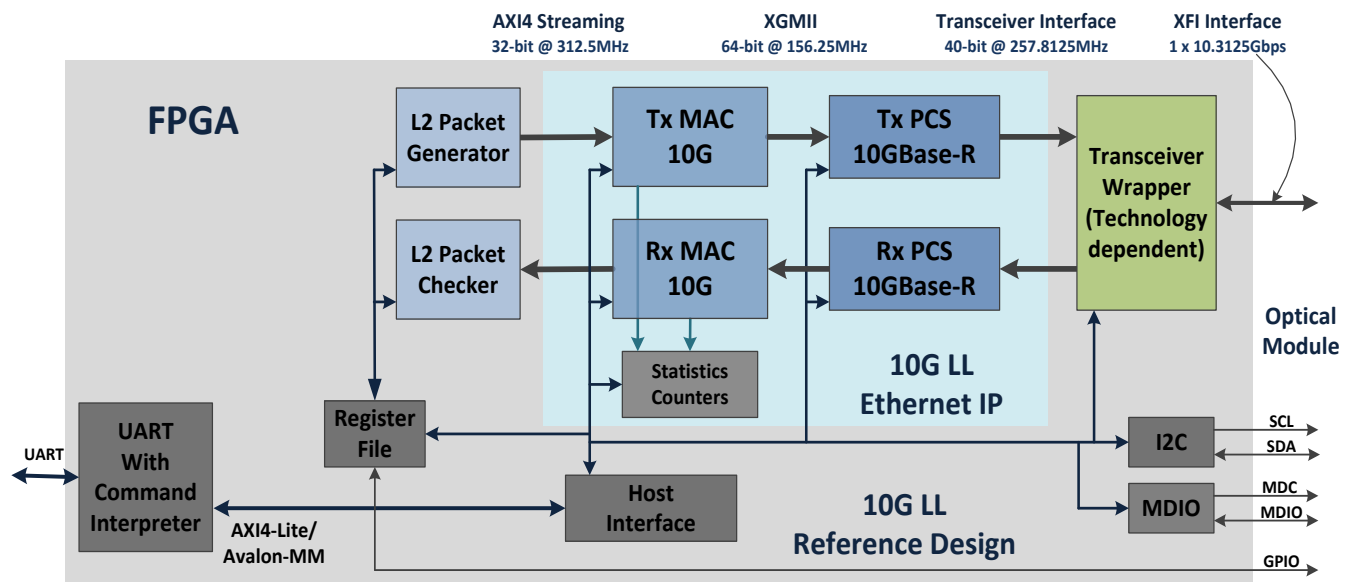
A 10Gbps reference design is included as part of the IP deliverable to facilitate quick L1 and L2 layer testing and verification of the 10Gbps Ethernet on target platform. The capability to run the L1 PRBS pattern and configure each transceiver independently can be for used for a fast module bring-up in the lab and can also be used for factory diagnostics.

The UART (normally through an onboard USB-to-UART converter chip) based 10G 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 10Gbps 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 host (embedded or standard PC) running a GUI application is used to configure and control the 10G Ethernet. I2C and GPIO interfaces included in the reference design can be used to control any optical module on the target platform including the XFP+ and XFP compliant modules.



For L1 (physical layer verification and testing) GUI application provides an interface to independently control and configure 10.3125Gbps transceiver used for 10G 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.

Low Latency 10G Ethernet IP Solution

Product Brief (HTK-LL10G-ETH-32-FPGA)



For L2 testing, GUI application uses the 10Gbps 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 "ETHERNET DEBUG APPLICATION" window. It has a sidebar on the left with a tree view containing items like "REGISTER READ/WRITE", "MDIO/I2C READ/WRITE", "XCVR DRP READ/WRITE", "MAC/PCS CONTROL", "REFDESIGN CONTROL", "MAC/PCS STATISTICS", "RMON STATISTICS", "REFDESIGN STATISTICS", "BANDWIDTH CALCULATION", and "REGRESSION TEST". The main area is divided into several sections. At the top left is a "Serial Port Setup" section with fields for "Port ID" (COM8), "Use UART" (checked), "Clear Buffer", "Baud Rate" (115200), "Open", and "Close" buttons. To its right is an "Execute Tcl Script" section with a text field containing "ethernet_test.tcl", a "Browse" button, "STATUS: Ready", and an "Execute" button. On the far right is a "Core Revisions" section with dropdowns for "FPGA-A", "MAC", and a text field with "XX_XX_XX_XX", along with "All" and "Get Revision" buttons. Below these are two large sections for "MAC Registers" and "PCS Registers". Each has checkboxes for "TX Pause Frame Count", "Frame Length Error Count", "Runt Frame Count", "RX PFC Pause Frame Count", "CRC Error Count", "TX PFC Pause Frame Count", "Blocksync Status", "Deskew Error", "VL1 BIP-8 Error Count", "AM Lock Status", "VL0 BIP-8 Error Count", and "VL2 BIP-8 Error Count". Each checkbox is followed by a hexadecimal address and a value field. At the bottom right is a "Status Log" section with "Clear Log", "Save Log", and "Add Time" buttons, and a large text area for the log output.