

10G Ethernet MAC


Product Brief (HTK-10G-MAC-64)



The 10G dual-mode Ethernet MAC offers an IEEE802.3-2008 (802.3ae) compliant solution that meets the requirements for 10Gbps LAN in NIC (Network Interface Card) and Ethernet switching applications.

As shown in the figure below, the 10G MAC IP includes:

- 10G MAC with 10Gbps XGMII PCS side interface
- Statistics counter block (for RMON and MIB)
- MDIO and I2C cores for external PHY and optical module status/control



A complete reference design using an L2 (MAC level) packet generator/checker is also included to facilitate quick integration of the Ethernet MAC in a user design. A GUI application interacts with the reference design's hardware elements through a PCIe interface (a UART option is also available). A 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 64-bit data path operating at 156.25MHz.

Ethernet IP solution implements two user (application) side interfaces. The register access port can either be a 32-bit AXI4 interface or a 32-bit Avalon-MM interface. IP solution provides a highly flexible 10Gbps traffic port interface options. Depending upon the application layer, user can select an AXI-4 streaming bus or an Avalon Streaming bus to interface with the MAC block. In both modes, the MAC interface is a 64-bit bus operating at 156.25MHz.

On the Physical interface side, the MAC core implements a 64-bit SDR (Single Data Rate) XGMII interface for 10G which samples frame data at rising edge of the clock only. This interface is compatible with various vendor specific 10G PCS layer cores. Based upon the target design, user can select 10GBase-R, XAUI or RXAUI PCS for 10Gbps interface.

10G Ethernet MAC core 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.
- Dynamically configurable to support LAN and WAN (OC-192c/SDH SONET) rates.
- 64-bit SDR XGMII interface operating at 156.25 MHz.
- Deficit Idle Count (DIC) mechanism to ensure data rates of 10Gbps at the transmit interface.
- Dynamic Inter frame Gap calculation in WAN mode of operation with stretch ratio of 104.
- 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.
- Optional multicast address filtering with 64-bit HASH Filtering table providing imperfect filtering to reduce load on higher layers.
- CRC-32 generation and checking at high speed using Galois field multipliers and alternate polynomials.
- Optional forwarding of the CRC field to user application interface.
- Optional forwarding of received pause frames to the user application interface.
- Programmable frame maximum length providing support for any standard or proprietary frame length (e.g. 9K-Bytes Jumbo Frames).
-

10G Ethernet MAC

Product Brief (HTK-10G-MAC-64)



- 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.
- Implements a streaming 64-bit FIFO based application interface.
- Transmit and Receive FIFOs with configurable depths having a default depth of 1KB (128 64-bit words) each.
- 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.

See **Appendix B** for functional details of MAC core.

Licensing and Maintenance

- **True sign once licensing with NO yearly maintenance fees**
- Basic core licensing for a single vendor (either Xilinx or Altera) compiled (synthesized) binary
- Additional vendor license provided at only 50% cost of the base license. This allows for cost effective multi-vendor designs with identical user and control interfaces.
- 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)

Deliverables

- Compiled synthesizable binaries or encrypted RTL for the MAC core
- A **no-cost** 10GBase-R core for FPGAs with high speed transceivers ($\geq 10.3125\text{Gbps}$)
- 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 PCIe/UART 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
 - PCIe driver/API (source files, C) for Linux
 - UART and command interpreter blocks with the optional UART 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
- USA based technical support by developers

Contact and Sales Information

Phone: +1-301-528-2244

Email: info@mantaro.com

10G Ethernet MAC

Product Brief (HTK-10G-MAC-64)



MAC Resource Utilization

The MAC core contains two configurable depth FIFOs on both TX and RX user interfaces, for frame buffering. The following table provides the resource utilization for the 10GbE MAC with 1KB (128 64-bit words) FIFO depths.

The 10G Ethernet MAC 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 MAC - Resource Usage for Xilinx Devices

<i>Device</i>	<i>Speed Grade</i>	<i>Slices</i>	<i>Slice Registers (FFs)</i>	<i>Memory (Block RAMs)</i>
Virtex-4	-11	4374	4546	7 RAMB16
Virtex-5	-2	2782	4453	7 18K-BRAM
Virtex-6	-2	1773	4280	1 18K-BRAM 6 36K-BRAM

10G MAC - Resource Usage for Altera Devices

<i>Device</i>	<i>Speed Grade</i>	<i>Logic Cells</i>	<i>Registers (FFs)</i>	<i>Memory (Block RAMs)</i>
Cyclone-II	-6	7619 LEs	4286	7 M4K
Stratix-1	-5	7590 LEs	4268	7 M4K
Stratix-II	-5	4681 ALUTs	4346	7 M4K
Stratix-IV	-4	4097 ALUTs	4220	7 M9K

A. Reference Design Details

A.1 Overview

A 10G MAC reference design is included as part of the IP deliverable to facilitate quick integration and verification of the 10Gbps Ethernet on target platform.

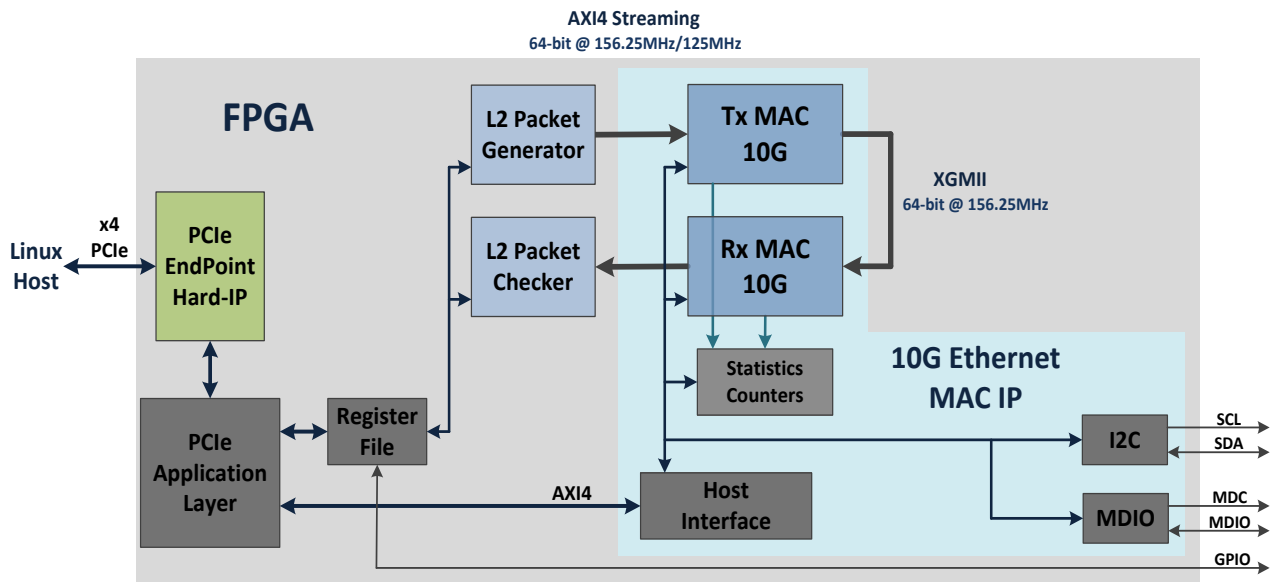
The PCIe based 10G Ethernet reference design can be seamlessly ported to various COTS PCIe from factor FPGA networking and evaluation modules (see section for the list of verified modules). This reference design can also be used on custom embedded design where the FPGA connects to the host processor via a PCIe interface. For the PCIe control interface, GUI application is hosted on a Linux platform.

For the embedded designs and evaluation modules that do not have PCIe control interface but do have UART debug interface (usually through a USB-to-UART converter chip), 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.

A.2 Functional Description

Following figure shows the connectivity and the elements of the 10G Ethernet MAC reference design. The physical side MAC interfaces (XGMII) is looped back from the Tx path to the Rx path in the reference design. End customer can connect appropriate (application dependent) PCS core to the MAC and perform extended simulation/verification with the same setup

A Linux host (embedded or standard PC) running a GUI application is used to configure and control the 10G Ethernet. I2C, MDIO 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 (requires PCS layer addition to the reference design).



GUI application uses the 10Gbps capable packet generator/checker inside the FPGA to generate and check MAC frames up to full line rate. The packet generator supports a basic rate control mechanism to control the packet/data rate on the interface. The generator can be configured for fixed size as well as pseudo random packet size packet transmission. An

10G Ethernet MAC

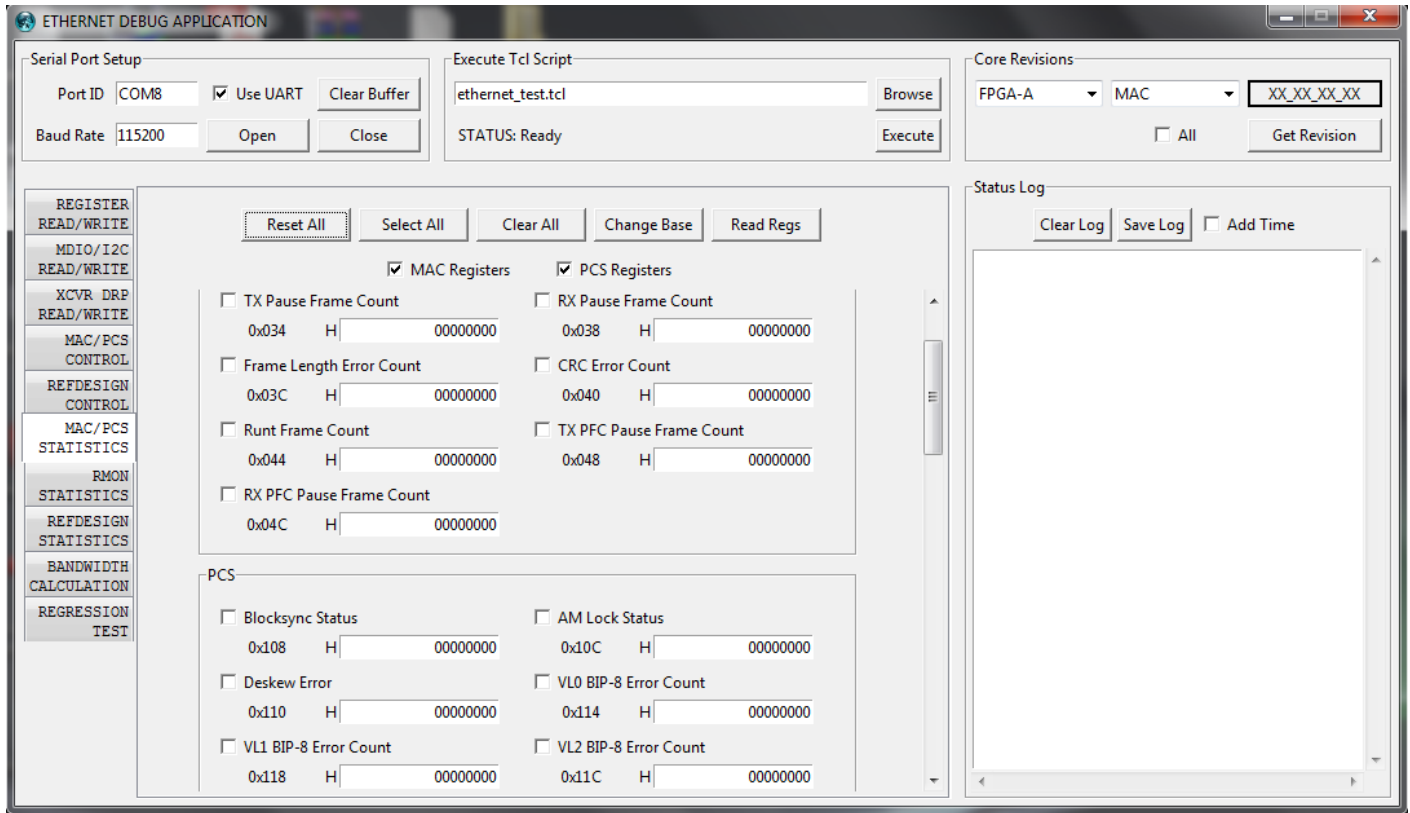
Product Brief (HTK-10G-MAC-64)



incrementing counter is used as payload for the MAC frames. The 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.

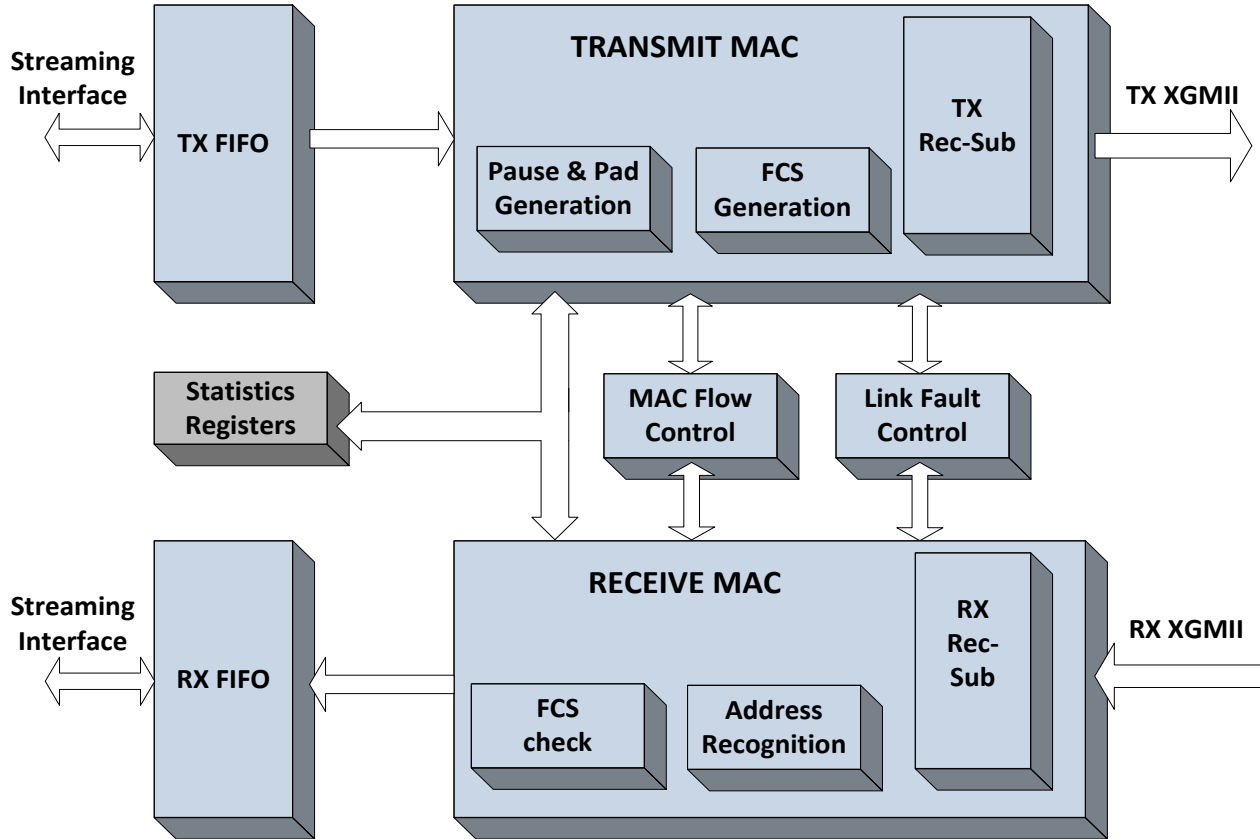


A.3 Ported/Validated Modules List

1. HTG-V6HXT-x16PCIE; Virtex-6 HXT FPGA, x16 PCIe Module with 2 QSFP+ and 2 SFP+ interfaces (http://www.mantaro.com/products/development_platforms/Xilinx_V6HXT_x16PCIE.htm)
2. HTG-V6HXT-x8PCIE; Virtex-6 HXT FPGA, x8 PCIe Module with 4 SFP+ interfaces (http://www.mantaro.com/products/development_platforms/Xilinx_V6HXT_x8PCIE.htm)
3. HTG-S4GT-PCIE; Virtex-6 HXT FPGA, x8 PCIe Module with 2 SFP+ interfaces (http://www.mantaro.com/products/development_platforms/Altera_S4G_PIE.htm)

B. MAC Functional Overview

The following figure shows the architecture for the 10G Ethernet MAC.



B.1 Reconciliation Sublayer (RS) Operation

The RS layer is responsible to map the data to/from the MAC sublayer to the XGMII interface. The RS layer provides a 64-bit (SDR) interface with data sampled at positive edge of the XGMII clock. This data interface can be connected directly to the DDR I/O structure which maps it to XGMII interface with 32-bits of data sampled at both rising and falling edge of clock signal. The 64-bit data is organized into eight 8-bit lanes with a control bit available for each lane.

B.2 MAC Sublayer Operation

The MAC sublayer is responsible to perform transmit and receive operations. The transmit MAC block transmits frames from a user application interface to the reconciliation sublayer, which then transmits these frames to the XGMII physical interface. The receive MAC block receives Ethernet frames from the reconciliation sublayer, validates the Ethernet frame and transfers this frame to the user application interface. The following description defines the various functions performed by transmit and receive Ethernet MAC engines.

Transmit Ethernet MAC

The transmit Ethernet MAC performs the following main functions:

- Accepts data including Destination Address, Source Address and length field from the MAC client.
- Appends preamble and SFD to the Ethernet frames.

10G Ethernet MAC

Product Brief (HTK-10G-MAC-64)



- Inserts PAD field for frames with length less than minimum frame length (64 bytes) and runt frames are not configured to pass through.
- Calculates and Appends proper FCS (CRC-32) value to outgoing frames and verifies full octet boundary alignment.
- Delays transmission of frame data for specified inter-frame gap period.
- Controls Inter-frame gap timing for both LAN and WAN modes of operation.
- Generates preamble and SFD field before frame transmission.
- Manages local device flow control by generating PAUSE control frames.
- Manages Remote device congestion by transiting to HALT state for a specified time quanta.

Receive Ethernet MAC

The receive Ethernet MAC performs the following main functions:

- Receives a frame from the RS sub layer via a 64-bit data bus.
- Presents to the MAC client sublayer frames that are either frames with group address or directly addressed to the local station (Address recognition).
- Filters Multi-cast frames using hash filtering algorithm.
- Discards all frames not addressed to the receiving station when promiscuous mode is disabled.
- Accepts all frames destined to the EMAC if promiscuous mode is enabled.
- Checks incoming frames for transmission errors by way of FCS and verifies octet boundary alignment.
- Discards received transmissions that are less than a minimum length (64 bytes).
- Truncates frames with length greater than maximum frame length when Jumbo frames are not allowed to pass through.
- Optionally forwards pause frames to user application

B.3 MAC Flow Control Operation

The MAC flow control block is responsible to maintain a proper flow of Ethernet frames through transmit and receive engines. It performs the following main functions:

- Prevents the receive EMAC FIFO congestion by sending pause control frames.
- Prevents the remote device congestion by responding to pause frames and going into idle state for specified number of slot times.

Automatic flow control is only available during the non-PFC (legacy single priority flow control) mode of operation. For PFC mode, user layer is responsible for managing the flow control operation.

B.4 Management Interface

The 10G MAC core provides a set of signals which can be used to implement the statistics required in IEEE 802.3 basic, mandatory and recommended Management information packages. In addition the MAC core provides signals to generate the applicable objects of the Management Information Database (MIB, MIB II) according to IETF RFC2665.