10 Gigabit Ethernet MAC

(with Atlantic User Application Interface)

PRODUCT BRIEF

1. INTRODUCTION

This document serves as an introduction to the 10 Gigabit Ethernet MAC core. The document outlines the various features supported by the 10 Gigabit Ethernet MAC and a brief functional description of the core. It also includes the implementation summary of the 10G EMAC on different devices.

1.1. Glossary

LAN	Local Area Network
WAN	Wide Area Network
MAC	Media Access Controller
VLAN	Virtual LAN
EMAC	Ethernet MAC
RS	Reconciliation Sublayer
PHY	Physical Interface
CRC	Cyclic Redundancy Check
FCS	Frame Check Sequence
XGMII	10 Gigabit Media Independent Interface
IFG	Inter-Frame Gap
DIC	Deficit Idle Count
MIB	Management Information Base
RMON	Remote Network Monitoring

2. 10G EMAC CORE FEATURES

The 10 Gigabit MAC core is designed to comply with the IEEE 802.3ae specifications and meets the requirements for both 10 Gbps LAN and 9.95328 Gbps OC-192c/SDH SONET WAN modes of operations. The core is designed to support different configuration modes controlled by the core's register file. The core also provides support for IEEE managed objects, IETF MIB and RMON for management applications.

The 10G core at the user application interface, implements a simple and flexible user application interface that is designed to transfer MAC frame data to/from the MAC core. On the Physical interface side, the MAC core implements a 64-bit SDR(Single Data Rate) XGMII interface which samples frame data at rising edge of the clock only. The 64-bit data from the MAC core can be directly connected to any specific DDR I/O structure which will provide 32-bits of data on both rising and falling edges of clock signal as specified in the IEEE XGMII specifications.

The various core features are outlined below

- Implements the full 802.3 specification with preamble/SFD generation, frame padding generation, CRC generation and checking on transmit and receive respectively.
- 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 10 Gbps 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.
- 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).
- 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 Atlantic interface compliant user (FIFO) 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.

3. FUNCTIONAL DESCRIPTION

3.1. Architectural Overview

The following figure shows the architecture for the 10 Gigabit Ethernet MAC

Figure 1: 10G EMAC High level block diagram

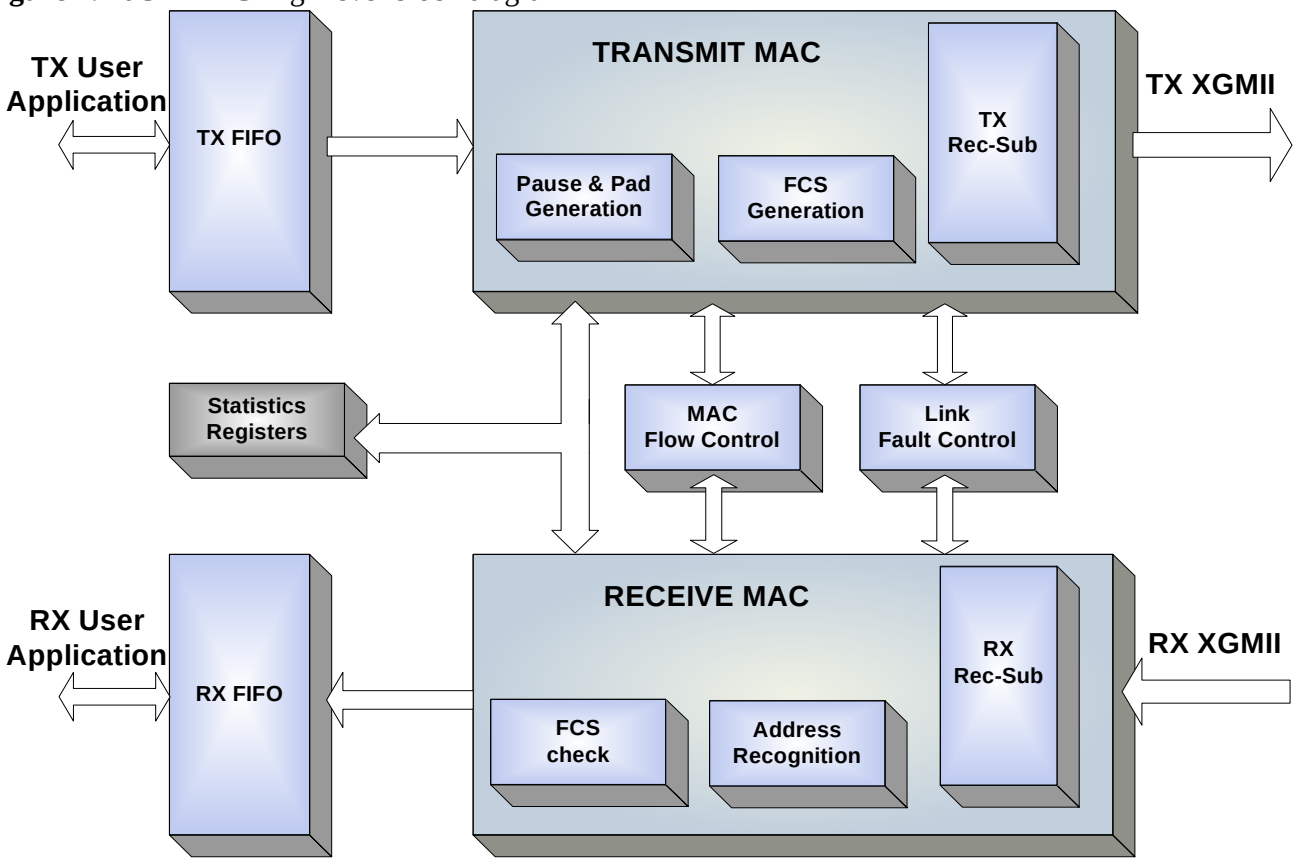

3.2. FUNCTIONAL OVERVIEW

3.2.1. 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.

3.2.2. 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 the transmit and receive Ethernet MAC engines.

3.2.2.1. 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.
- 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.

3.2.2.2. 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 in 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.

3.2.3. MAC Flow Control Operation

The MAC flow control block is responsible to maintain a proper flow of Ethernet frames through the 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.

4. MANAGEMENT INTERFACE

The 10G EMAC 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.

5. 10G EMAC IMPLEMENTATION SUMMARY

An estimate of the logic resources and memory utilization for the 10 Gigabit Ethernet MAC for different devices of Altera and Xilinx is shown in the following table. The MAC core contains two configurable depth FIFOs on both TX and RX user interfaces, for frame buffering. The following table gives the implementation for the 10GbE MAC with minimum verified FIFO depths that is (1KB = 128 64-bit words).

Device	***Speed Grade***	***Logic cells (LE,Slices)***	***Registers***	***Memory M4K, Block-RAM***	***Performance MHz***
ALTERA					
STRATIX	-5	7590 LEs	4268 FFs	7 M4K	163 MHz
CYCLONE II	-6	7619 LEs	4286 FFs	7 M4K	173 MHz
STRATIX II	-5	4681 ALUTs	4346 FFs	7 M4K	187 MHz
XILINX					
VIRTEX-4	-11	4374 Slices	4546 FFs	7 RAMB16	226 MHz
VIRTEX-5	-2	2782 Slices	4453 FFs	7 18K-BlockRAM	255 MHz

Table 1: 10GbE resource usage

6. REFERENCES

1. IEEE 802.3ae specification 2002
2. RFC2665, definition of Managed Objects for Ethernet Like Interface Type, www.ietf.org
3. RFC2863, The Interface Group MIB, www.ietf.org
4. RFC2819, Remote Network Monitoring (RMON) MIB, www.ietf.org