

UG0727
User Guide
PolarFire FPGA 10G Ethernet Solutions



Power Matters.™

Microsemi Corporate Headquarters

One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113

Outside the USA: +1 (949) 380-6100

Fax: +1 (949) 215-4996

Email: sales.support@microsemi.com

www.microsemi.com

© 2017 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

About Microsemi

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions, security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www.microsemi.com.

Contents

1	Revision History	1
1.1	Revision 3.0	1
1.2	Revision 2.0	1
1.3	Revision 1.0	1
2	10G Ethernet Overview	2
3	PolarFire FPGA Evaluation Kit Ethernet Support	3
4	Building Blocks for 10G Ethernet Solutions	4
4.1	Soft Processor IP	4
4.2	Ethernet MAC IP (Core10GMAC)	4
4.3	Ethernet Interface IP	5
4.4	Transmit PLL	5
4.5	IP Licensing	5
5	Implementing 10G Ethernet Solutions	6
5.1	10GBASE-KR Designs	6
5.1.1	Configuring IP and Transceiver Interface Using Libero SoC PolarFire	7
5.1.2	Clocking Requirements	8
5.2	10GBASE-R Designs	9
5.2.1	Configuring IP and Transceiver Lane Using Libero SoC PolarFire	9
5.2.2	Clocking Requirements	10
5.3	Firmware Support	10
6	Appendix: MAC Layers in the OSI Reference Model and Standard Ethernet Interfaces	11
6.1	MAC Layers	11
6.1.1	LLC	11
6.1.2	MAC Sublayer	11
6.1.3	MAC Control Sublayer	12
6.1.4	Reconciliation Sublayer	12
6.1.5	Physical Sublayers (PCS, PMA, and PMD)	12
6.2	Standard Ethernet MAC Interfaces	12
6.2.1	10GBASE-R	12
6.2.2	10GBASE-KR	13
7	Appendix: Ethernet Frame Format	14
7.1	Preamble	14
7.2	SFD	14
7.3	MAC Address Fields	14
7.4	VLAN Tag (for VLAN Frames Only)	15
7.5	Length/Type	15
7.6	Data	15
7.7	Pad	15
7.8	FCS	15
8	Appendix: Glossary	16

Figures

Figure 1	Sample Ethernet Application	2
Figure 2	PolarFire FPGA Evaluation Board Hardware Block Diagram	3
Figure 3	Sample Ethernet Application Using Microsemi IP Cores	4
Figure 4	Backplane 10GBASE-KR Ethernet Designs	6
Figure 5	Core10GMAC Configuration for 10GBASE-KR Designs	7
Figure 6	Transceiver Interface Configuration for 10GBASE-KR Designs	8
Figure 7	SFP-Based 10GBASE-R Designs	9
Figure 8	Core10GMAC Configuration for 10GBASE-R Designs	9
Figure 9	Transceiver Interface Configuration for 10GBASE-R Designs	10
Figure 10	IEEE Standard 802.3-2012 Ethernet Model	11
Figure 11	10GBASE-R System Diagram	12
Figure 12	10GBASE-KR System Diagram	13
Figure 13	Standard Ethernet Frame Format	14
Figure 14	Ethernet VLAN Frame Format	14

Tables

Table 1	License Information for Microsemi 10G Ethernet-Based IP	5
Table 2	Common Ethernet Terms	16

1 Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the current publication.

1.1 Revision 3.0

Throughout the document, Libero SoC PolarFire Configurator figures were updated to reflect the Libero SoC PolarFire v2.0 release.

1.2 Revision 2.0

The following is a summary of the changes made in revision 2.0 of this document.

- Throughout the document, Libero SoC PolarFire Configurator figures were updated to reflect the Libero SoC PolarFire v1.1 SP1 release.
- Clocking information was updated. For more information, see [10GBASE-KR Designs](#), page 6 and [10GBASE-R Designs](#), page 9.

1.3 Revision 1.0

Revision 1.0 was the first publication of this document.

2 10G Ethernet Overview

Ethernet is a family of networking interface standards used in systems and applications across multiple industries. Implementation of Ethernet solutions in FPGAs requires IP and design flows that reduce development time and utilize minimal device resources, thereby helping meet performance, power, and cost goals. Microsemi PolarFire® devices support Ethernet data transfer rates ranging from 10 Mbps to 10 Gbps on a single interface.


Microsemi PolarFire devices provide a complete range of solutions for implementing IEEE 802.3 standard-compliant Ethernet interfaces for chip-to-chip, board-to-board, and backplane interconnects. The high-speed serial interface and soft IP blocks available in PolarFire devices enable designers to build Ethernet solutions for use in embedded systems and systems connected over copper or optical cabling.

PolarFire FPGA 10G Ethernet support is compliant with the IEEE 802.3ae standard that supports data transfer rates of up to 10.3125 Gbps. Advantages offered by PolarFire FPGAs for building 10G Ethernet solutions include the use of low-power transceivers, low-power FPGA fabric, and SyncE-compliant jitter attenuation.

In PolarFire devices, 10G Ethernet is implemented using the Core10GMAC soft IP media access control (MAC) core, which can be configured in 10GBASE-KR and 10GBASE-R modes. Core10GMAC supports standard Ethernet interfaces such as the 10 Gbps attachment unit interface (XAUI) and the 10 Gbps reduced attachment unit interface (RXAUI).

A typical Ethernet application, such as a switch or a router, requires an Ethernet MAC sublayer (commonly referred to as the MAC) that supports standard Ethernet interfaces, an Ethernet physical layer (PHY), and an SFP connector. The following illustration shows a sample Ethernet application.

Figure 1 • Sample Ethernet Application



For more information about the Ethernet MAC, including standard Ethernet interfaces, see [Appendix: MAC Layers in the OSI Reference Model and Standard Ethernet Interfaces](#), page 11.

3 PolarFire FPGA Evaluation Kit Ethernet Support


The PolarFire FPGA Evaluation Kit supports the 10GBASE-R standard (applicable to SFP applications), as well as the 10GBASE-KR standard (applicable to Ethernet backplane applications).

The PolarFire FPGA Evaluation Kit includes the following 10G Ethernet hardware components.

- SFP module supporting 1G and 10G Ethernet speeds that connects to a transceiver lane
- FPGA mezzanine card (FMC) high-pin count (HPC) connector
- SubMiniature version A (SMA) connectors

The following illustration shows the hardware of the PolarFire FPGA Evaluation Board. Highlighted in red are the hardware components used for implementing 10G Ethernet solutions.

Figure 2 • PolarFire FPGA Evaluation Board Hardware Block Diagram




4 Building Blocks for 10G Ethernet Solutions

Microsemi offers pre-designed and verified IP for all key markets and applications. A complete 10G Ethernet solution requires the following IP cores.

- Soft processor to implement control plane features, initialize the Ethernet MAC, and perform auto-negotiation
- Ethernet MAC to process Ethernet packets
- Ethernet interface to connect the Ethernet MAC with an SFP module
- Transceiver interface to send and receive serialized/deserialized data to and from the SFP module

The following illustration shows a sample Ethernet application developed using Microsemi IP cores.

Figure 3 • Sample Ethernet Application Using Microsemi IP Cores



For comprehensive information about all Microsemi IP, see www.microsemi.com/products/fpga-soc/designresources/ip-cores.

4.1 Soft Processor IP

CoreRISCV_AXI4: A 32-bit soft processor core such as CoreRISCV_AXI4 (or Cortex-M1) can be used to develop processor-based Ethernet solutions. The soft processor initializes the Ethernet MAC and runs real-time operating systems such as FreeRTOS. For more information about CoreRISCV_AXI4, see [CoreRISCV_AXI4 Handbook](#).

CoreABC: CoreABC is a simple, configurable, low gate-count controller primarily targeted at implementing Advanced Microcontroller Bus Architecture Advanced Peripheral Bus (AMBA APB)-based designs. In an Ethernet-based application, this core is used only for configuring the Ethernet MAC. For more information about CoreABC, see the [CoreABC Handbook](#).

4.2 Ethernet MAC IP (Core10GMAC)

Core10GMAC is a soft IP MAC core with built-in PCS that supports the 10GBASE-R and 10GBASE-KR Ethernet standards. It is compliant with the IEEE 802.3 standard, which contains PHY and MAC specifications for wired Ethernet.

Multiple Core10GMAC IP blocks can be used to develop Ethernet solutions in PolarFire devices. For more information, see [Core10GMAC Handbook](#).

4.3 Ethernet Interface IP

The PolarFire FPGA transceiver interface (PF_XCVR) provides the physical media attachment (PMA) for high-speed serial interfaces. The transceiver has a multi-lane architecture with each lane natively supporting serial data transfer rates ranging from 250 Mbps to 12.7 Gbps. For more information, see *UG0677: PolarFire FPGA Transceiver User Guide*.

4.4 Transmit PLL

The PolarFire FPGA transmit PLL (PF_TX_PLL) provides the high-speed bit clock for the PolarFire FPGA transceiver. When used with the transceiver, the transmit PLL supports jitter attenuation for loop-timing applications where recovered clocks are used as transmit reference clocks. The jitter attenuation feature controls low-speed wander, ensuring compliance with SyncE specifications. For more information, see *UG0677: PolarFire FPGA Transceiver User Guide*.

4.5 IP Licensing

The Libero® SoC PolarFire software provides free access to several Microsemi IP, but some IP require purchasing a separate license. Contact [Customer Service](#) for information about how to purchase licenses.

The following table lists license information for each Ethernet-based IP.

Table 1 • License Information for Microsemi 10G Ethernet-Based IP

IP Core	License Information
CoreRISCV_AXI4	Available with the Libero SoC PolarFire license
Cortex-M1	Available with clickthrough license ¹
CoreABC	Available with the Libero SoC PolarFire license
Core10GMAC	Must be purchased separately
PF_XCVR (transceiver interface)	Available with the Libero SoC PolarFire license

1. The IP is listed in the Libero SoC IP catalog but can be used only after signing an end-user license agreement (EULA).

5 Implementing 10G Ethernet Solutions

In PolarFire devices, 10G Ethernet solutions can be implemented using the Core10GMAC soft IP, which supports both the 10GBASE-R and 10GBASE-KR standards in both 32-bit and 64-bit data widths. The IP must be initialized and configured using a soft processor or a state machine hosted in the fabric.

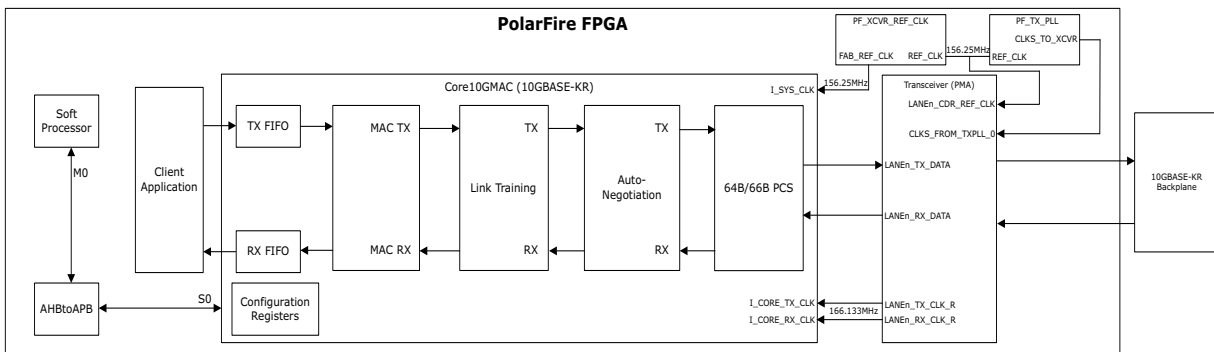
The following sections describe how the MAC interfaces with the PHY in PolarFire devices for various Ethernet interfaces. For more information about Ethernet interfaces, see [Appendix: MAC Layers in the OSI Reference Model and Standard Ethernet Interfaces](#), page 11.

Note: The designs illustrated in this section use the 64-bit data width for the MAC and the transceiver. If required, the data width can be changed to 32 bits in the Core10GMAC IP and the transceiver interface configurator. Updating the data width automatically updates the port names and the FPGA interface frequency.

5.1 10GBASE-KR Designs

For backplane 10GBASE-KR Ethernet designs, Core10GMAC is configured in 10GBASE-KR mode and connected to the transceiver, as shown in the following illustration.

Figure 4 • Backplane 10GBASE-KR Ethernet Designs

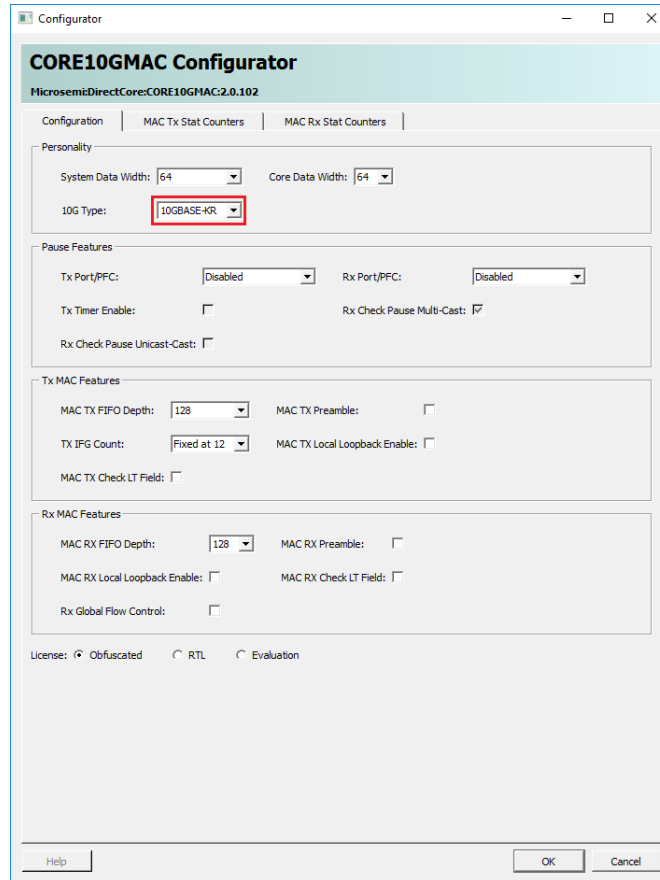


5.1.1 Configuring IP and Transceiver Interface Using Libero SoC PolarFire

To implement backplane 10GBASE-KR Ethernet designs, use the following settings in Libero SoC PolarFire.

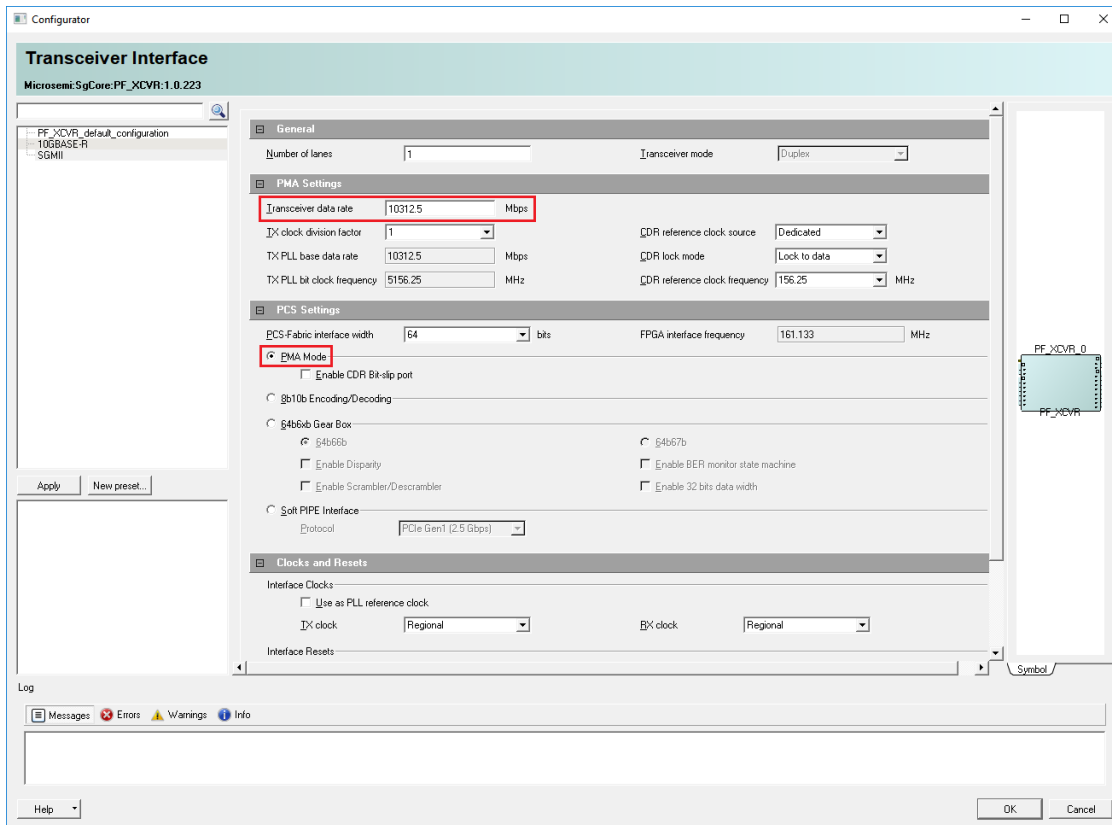
- Configure Core10GMAC in 10GBASE-KR mode, as shown in the following figure.

Figure 5 • Core10GMAC Configuration for 10GBASE-KR Designs



- Configure the transceiver in 10GBASE-KR mode, as shown in the following figure.

Figure 6 • Transceiver Interface Configuration for 10GBASE-KR Designs



When configuring the transceiver, the transceiver must be set to 10312.5 Mbps to match the 10GBASE-KR data transfer rate, as shown in the preceding figure.

The TX clock division factor allows the transceiver lane high-speed clock from the TX PLL to be divided, allowing the user to share a higher rate TX PLL (for data transfer rates of up to 6.4 Gbps only). The default value for this field is 1.

Because the Core10GMAC IP has built-in PCS functionality, PMA mode must be selected under PCS settings to allow the transceiver to transmit and receive serial data. PMA mode allows the transmission of the Manchester-encoded data for link-training.

5.1.2 Clocking Requirements

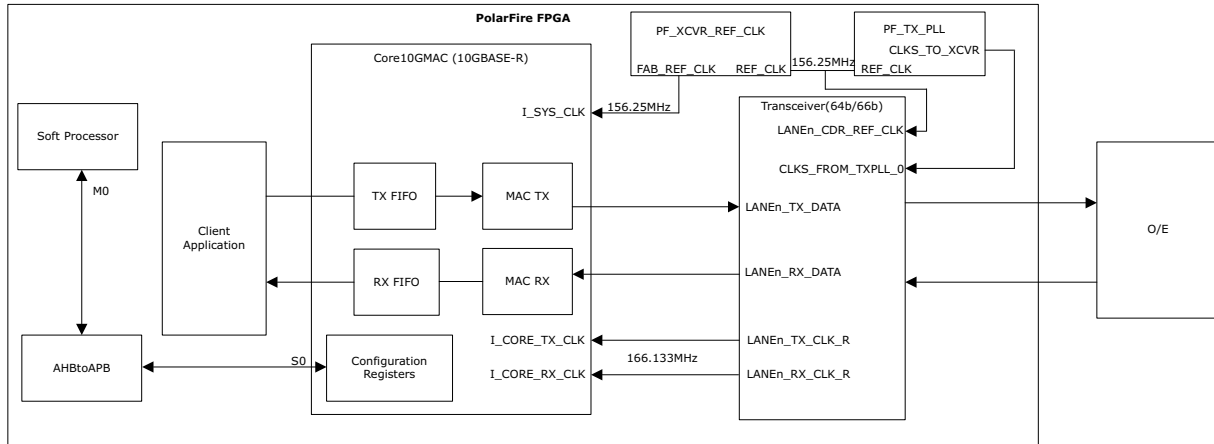
The following clocks are required for 10GBASE-KR designs.

- **I_CORE_TX_CLOCK:** 322.266-MHz or 161.133-MHz transmit clock for 32-bit and 64-bit MAC configuration respectively, driven by the transceiver regional transmit clock
- **I_CORE_RX_CLOCK:** 322.266-MHz or 161.133-MHz receive clock for 32-bit and 64-bit MAC configuration respectively, driven by the transceiver regional receive clock
- **I_SYS_CLOCK:** 322.266-MHz or 161.133-MHz system clock driven by the clock that drives I_CORE_TX_CLK or I_CORE_RX_CLK, or a different clock with the same frequency as the clock that drives these clocks
- **LANEn_CDR_REF_CLK:** 156.25-MHz reference clock to lane CDR, driven by the transceiver reference clock or a fabric clock conditioning circuit (CCC)
- **CLKS_FROM_TXPLL_0:** Clocks from the transmit PLL bus interface port with the following underlying signals common to all lanes instantiated in the transceiver interface IP core:
 - TX_BIT_CLK_0
 - TX_PLL_LOCK_0
 - TX_PLL_REF_CLK_0

5.2 10GBASE-R Designs

For 10GBASE-R designs, Core10GMAC is configured in 10GBASE-R mode and connected to the transceiver, as shown in the following illustration.

Figure 7 • SFP-Based 10GBASE-R Designs

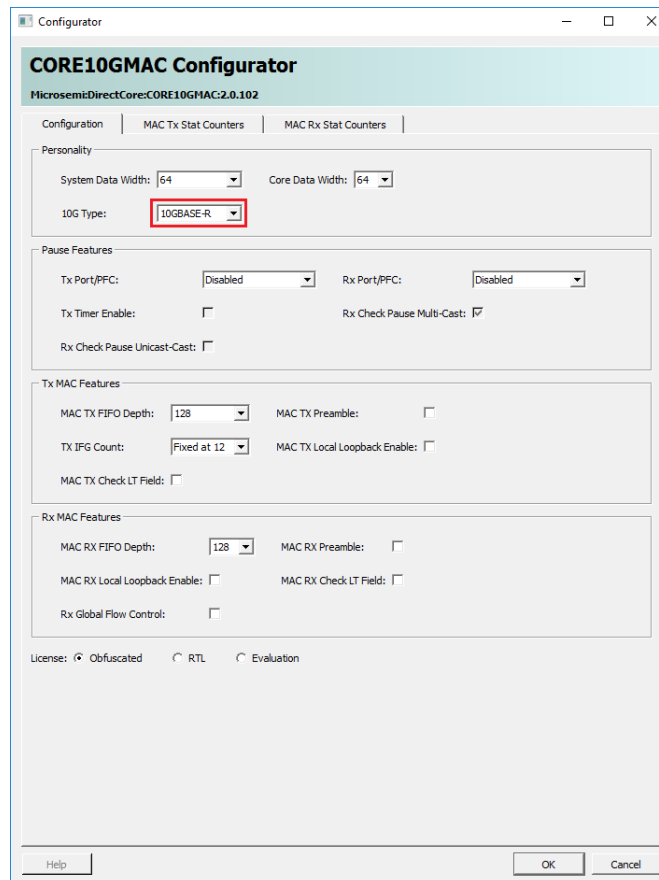


5.2.1 Configuring IP and Transceiver Lane Using Libero SoC PolarFire

To implement 10GBASE-R designs, use the following settings in Libero SOC PolarFire.

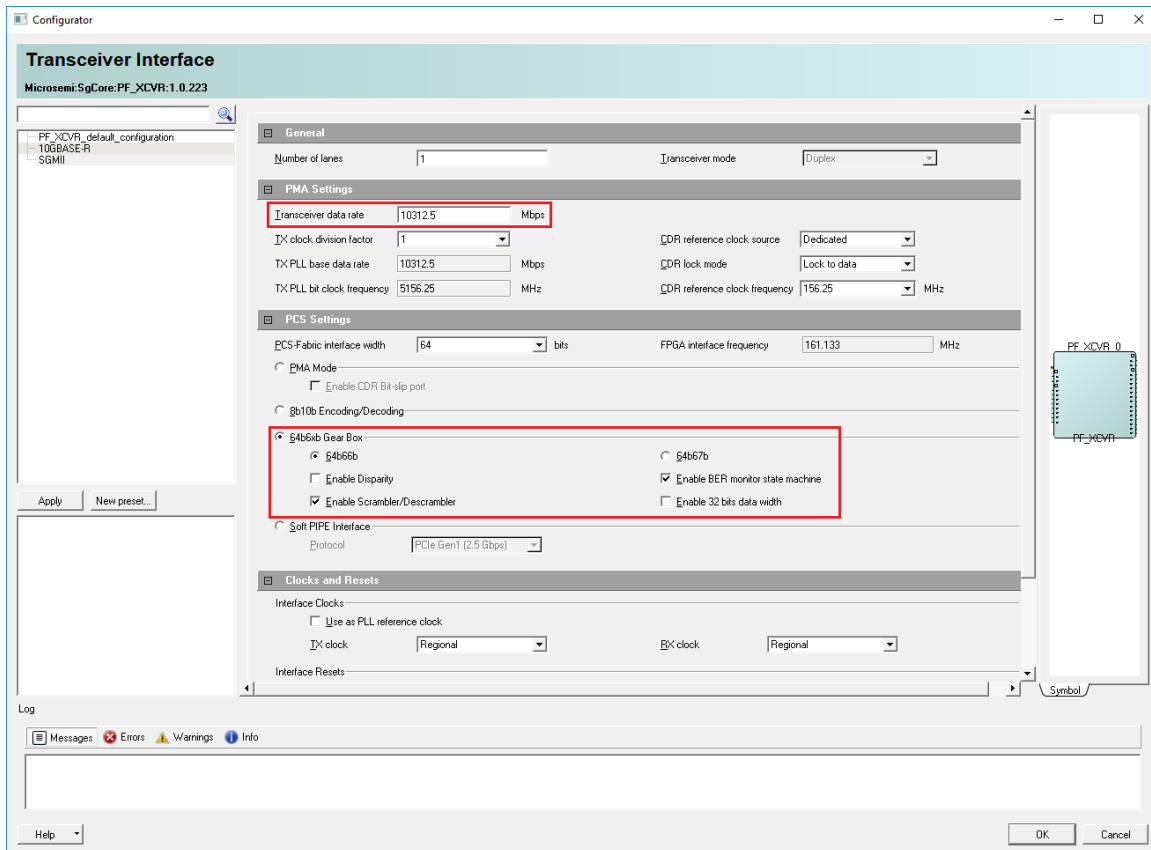
- Configure Core10GMAC in 10GBASE-R mode, as shown in the following figure.

Figure 8 • Core10GMAC Configuration for 10GBASE-R Designs



- Configure the transceiver interface in 10GBASE-R mode, as shown in the following figure.

Figure 9 • Transceiver Interface Configuration for 10GBASE-R Designs



When configuring the transceiver, the transceiver data rate must be set to 10312.5 Mbps to match the 10GBASE-R data transfer rate, as shown in the preceding figure.

The TX clock division factor allows the transceiver lane high-speed clock from the TX PLL to be divided. Thus, it allows the user to share a higher rate TX PLL and locally divide the clock (for data transfer rates up to 6.4 Gbps only). The default value for this field is 1.

The 64b6b gear box option must be selected under PCS settings to allow encoding and scrambling of Ethernet data.

5.2.2 Clocking Requirements

The clocking requirements for 10GBASE-R designs are the same as for 10GBASE-KR designs. For more information, see [Clocking Requirements](#), page 8.

5.3 Firmware Support

The Core10GMAC IP driver is distributed through the Microsemi SoC firmware catalog. This catalog provides access to the documentation for the driver, generates application projects from source files, and generate sample projects that illustrate how to use the driver.

The firmware catalog is available at:
www.microsemi.com/soc/products/software/firmwarecat/default.aspx.

The Core10GMAC driver supports auto-negotiation, link training, and MAC initialization.

For more information about the Core10GMAC driver, see *Core10GMAC Driver User Guide* (to be released).

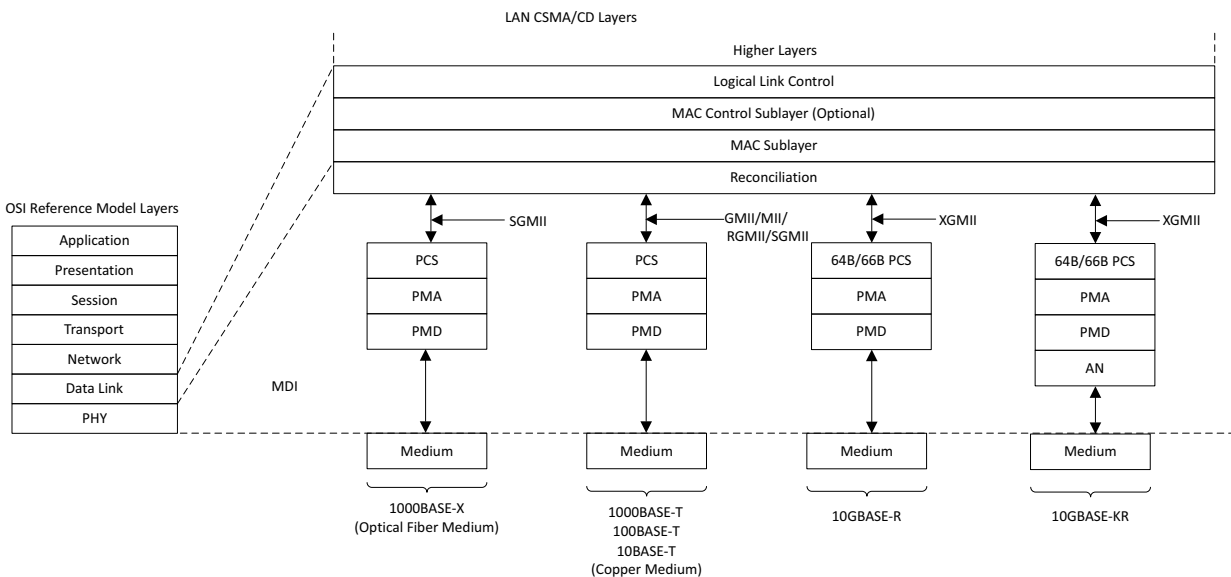
6 Appendix: MAC Layers in the OSI Reference Model and Standard Ethernet Interfaces

This section discusses how Ethernet MAC layers fit into the open systems interconnection (OSI) reference model and provides information about the standard interfaces used for Ethernet connections.

6.1 MAC Layers

The OSI reference model is a standard framework for data communication between networked systems. The following illustration shows the relationship between the OSI reference model and the Ethernet MAC as defined in the IEEE 802.3-2012 standard. It also illustrates where the supported physical interfaces (PCS, PMA, and PMD) fit into the architecture. The MAC and MAC control sublayers shown are handled by the Ethernet MAC.

Figure 10 • IEEE Standard 802.3-2012 Ethernet Model



The data link and physical layers in the OSI model are explained below.

6.1.1 LLC

The logical layer control (LLC) sublayer acts as an interface between the MAC and the network layer. It provides frame synchronization, flow control, and error checking mechanisms. It also offers multiplexing mechanisms to allow several network protocols to exist simultaneously in a multi-point network and share the same network medium for transmitting and receiving Ethernet packets.

6.1.2 MAC Sublayer

The MAC sublayer, referred to as the MAC, is defined in IEEE 802.3-2012, clauses 2, 3, and 4. The MAC is the second sublayer of the data link layer in the seven-layer OSI model. It provides addressing and channel access control mechanisms that allow terminals and network nodes in a multiple access network using a shared medium, such as an Ethernet network, to communicate with each other. The MAC is responsible for the transmission of data packets to and from the network-interface card. It is independent of the physical layer and can connect to any type of physical layer device.

6.1.3 MAC Control Sublayer

The MAC control sublayer, defined in IEEE 802.3-2012, clause 31 provides real-time flow control manipulation for the MAC. MAC and MAC control sublayer functions are performed by the Ethernet MAC in all modes of operation.

6.1.4 Reconciliation Sublayer

The reconciliation sublayer maps the signals between the physical medium interface and the MAC layer.

6.1.5 Physical Sublayers (PCS, PMA, and PMD)

The Ethernet physical layer consists of the following three sublayers.

- Physical coding sublayer (PCS)—performs auto-negotiation and coding operations such as 8b/10b encoding
- Physical medium attachment sublayer (PMA)—performs data serialization and clock recovery necessary to move serial data in and out of the device
- Physical medium-dependent sublayer (PMD)—hosts the transceiver that receives and transmits data through the physical medium


6.2 Standard Ethernet MAC Interfaces

This section describes standard interfaces used to connect the Ethernet MAC to the PHY and the signals associated with each interface.

6.2.1 10GBASE-R

10GBASE-R is a serial-encoded PCS that supports 32- and 64-bit data transmission over fiber-optic media. It allows Ethernet framing at 10.3125 Gbps.


Figure 11 • 10GBASE-R System Diagram



6.2.2 10GBASE-KR

10GBASE-KR is a serial-encoded PCS that supports 32-bit data transmission over copper media. It allows Ethernet framing at 10.3125 Gbps (similar to 10GBASE-R). The link training block and the optional auto-negotiation (AN) feature available in the 10GBASE-KR standard distinguishes it from the 10GBASE-R standard.


Figure 12 • 10GBASE-KR System Diagram



7 Appendix: Ethernet Frame Format


Ethernet data is encapsulated in frames consisting of a preamble, a start-of-frame delimiter (SFD), and the actual frame starting from the destination address and ending with the frame check sequence (FCS) field. The bytes within each field in an Ethernet frame are transmitted from left to right, going from the least significant bit to the most significant bit. A typical Ethernet frame's data length is between 0 bytes to 1500 bytes. Frames with data lengths greater than 1500 bytes are called jumbo Ethernet frames. The actual Ethernet frame begins after the SFD. The following illustration shows the format of a standard Ethernet frame.

Figure 13 • Standard Ethernet Frame Format



The Ethernet MAC also accepts virtual local area network (VLAN) frames. VLANs are specified in IEEE 802.1Q. A virtual, bridged LAN logically groups the network devices that share the same physical network. This way, the network traffic in a VLAN group is only visible to those devices that are members of that network group. For VLAN-type frames, the Ethernet MAC accepts four bytes more than a standard Ethernet frame. The following illustration shows the format of an Ethernet VLAN frame.

Figure 14 • Ethernet VLAN Frame Format



The following sections describe the individual fields in an Ethernet frame.

7.1 Preamble

The preamble field synchronizes receiver clocks within a network and contains seven bytes, with the pattern 0x55, transmitted from left to right. During transmission on the physical interface, this field is automatically inserted by the Ethernet MAC. On reception, it is stripped from the incoming frame before the data is passed to the MAC client. The MAC can receive Ethernet frames even if the preamble does not exist, as long as a valid SFD is available.

7.2 SFD

The SFD field marks the start of the Ethernet frame and must follow the pattern 0xD5. For transmission, this field is automatically inserted by the Ethernet MAC. On reception, it is stripped from the incoming frame before the data is passed to the MAC client.

7.3 MAC Address Fields

The MAC address is a unique identifier assigned to PHY interfaces to allow devices to communicate over the network. A MAC address can either be a source address or a destination address, depending on whether it is transmitting the frame or receiving it.

- Destination address—the MAC address of the frame's intended recipient on the network. It is the first field in an Ethernet frame that is transmitted and received between stations.

- Source address—the MAC address of the frame's initiator on the network. It is the second field in an Ethernet frame.

The least significant bit of a MAC address determines if a MAC address is an individual (unicast) address, a group (multicast) address, or a broadcast address.

- An individual address, also known as a unicast address, is specific to a station (device) on the network. For this address type, the destination address ends with 0.
- A group address, also known as a multicast address, is used to group logically-related stations. For this address type, the destination address ends with 1.
- A broadcast address is a multicast address used to group all stations on the LAN. For this address type, the destination address field has all 1s.

The Ethernet MAC supports transmission and reception of unicast, multicast, and broadcast packets. During transmission, the bit representing the individual or group (multicast) address is the first bit to appear in the address field of an Ethernet frame.

7.4 VLAN Tag (for VLAN Frames Only)

A VLAN tag field consists of the VLAN ID inserted into a packet header to identify the VLAN the packet belongs to. Based on the VLAN ID, switches determine the port or interface to which the broadcast packet needs to be sent.

7.5 Length/Type

The value of this field determines if it is interpreted as a length field or a type field as defined by IEEE 802.3-2012. The MAC interprets a value of 1500 bytes or less is interpreted by the MAC as a length field and a value of 1536 bytes or more as a type field. A length field represents the number of bytes in the following data field. This value excludes any bytes inserted in the pad field following the data field. A length/type field value of 0x8100 indicates a VLAN frame, and a value of 0x8808 indicates a PAUSE MAC control frame.

During transmission, the Ethernet MAC does not process the length/type field. On reception, if this field is a length field, the Ethernet MAC's receive engine interprets this value and removes any padding that may be displayed in the pad field.

If the field is a length field and length/type checking is enabled in the Ethernet MAC, the MAC compares the length against the actual data field length and flags an error if a mismatch occurs. If the field is a type field, the Ethernet MAC ignores the value and simply passes it along with the packet data with no further processing. The length/type field is retained in the receive packet data.

7.6 Data

For a normal frame, the data field can vary from 0 to 1,500 bytes in length for a normal frame. The Ethernet MAC can handle jumbo frames of any length. This field is provided in the packet data for transmissions and is retained in the receive packet data.

7.7 Pad

The pad field ensures that the Ethernet frame is at least 64 bytes in length. This is the minimum length required for successful CSMA/CD operation. The field can vary from 0 to 46 bytes in length.

7.8 FCS

A frame check sequence (FCS) is an error-detecting code that can optionally be added to an Ethernet frame. The value of the FCS field is calculated using the data in the destination address, source address, length/type, data, and pad fields through a 32-bit cyclic redundancy check (CRC) algorithm. For transmission, this field can be either inserted automatically by the Ethernet MAC or supplied by the client. On reception, the incoming FCS value is verified for every frame. If an incorrect FCS value is received, the Ethernet MAC indicates to the client that it has received a bad frame. The FCS field can either be passed on to the client or dropped by the Ethernet MAC based on whether the FCS feature is enabled in the MAC.

8 Appendix: Glossary

This section defines common terms associated with Ethernet architecture used in this document.

Note: This section does not define MAC sublayers, standard Ethernet interfaces, Ethernet frame fields, and 10GBASE-R/10GBASE-KR standards, which are described in previous sections.

Table 2 • Common Ethernet Terms

Term	Definition
Physical layer (PHY)	The PHY is the physical layer of the MAC, which, when instantiated, connects a link layer device (often called a MAC) to a physical medium such as an optical fiber or a copper cable.
SFP	A small form factor pluggable (SFP) connector, commonly known as an SFP, is a compact, hot-pluggable transceiver (transmitter and receiver in a single package) used for carrying data over optical or copper wires.
Transceiver	A transceiver is a pair of functional blocks that converts serial data to parallel data and vice versa. The primary use of a transceiver is to provide data transmission over a single/differential line, thereby minimizing the number of I/O pins and interconnects required for the design.
Txp, Txn/Rxp, Rxn	Txp, txn and rxp, rxn are pairs of differential signals used to connect the transceiver to the SFP connectors.
Auto-negotiation	Auto-negotiation is an Ethernet procedure that allows two link partners to exchange capability parameters (speed, duplex mode, and flow control) and then choose the highest data transmission speed supported by both the devices. In the OSI model, the auto-negotiation capability resides in the PHY.
Link training	Link training is a process that enables communication between a transmitter and a receiver by controlling the transmit and receive signal integrity settings. Electrical characteristics and bit rate of the link are established during link training.
Differential Manchester encoding Scheme	Differential Manchester encoding Scheme is a mechanism used in auto-negotiation. It is a line code where data and clock signals are combined to form a two-level self-synchronizing data stream.
O/E	An optical-to-electrical converter (commonly known as an O/E) is a device that converts optical signals to electrical signals. The SFP module included in the PolarFire FPGA Evaluation Kit functions as an O/E in 10GBASE-R designs.