

40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide

Last updated for Altera Complete Design Suite: 14.1

 [Subscribe](#)

 [Send Feedback](#)

UG-01088
2014.12.15

101 Innovation Drive
San Jose, CA 95134
www.altera.com

ALTERA®

Contents

| | |
|---|------------|
| About the 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function..... | 1-1 |
| 40- and 100-Gbps Ethernet MAC and PHY IP Core Supported Features..... | 1-3 |
| 40-100GbE IP Core Device Family and Speed Grade Support..... | 1-4 |
| Device Family Support..... | 1-4 |
| 40-100GbE IP Core Device Speed Grade Support..... | 1-5 |
| IP Core Verification..... | 1-6 |
| Simulation Environment..... | 1-7 |
| Hardware Testing..... | 1-7 |
| Performance and Resource Utilization..... | 1-7 |
| Resource Utilization for 40GbE IP Cores..... | 1-7 |
| Resource Utilization for 100GbE IP Cores..... | 1-12 |
| Release Information..... | 1-18 |
| | |
| Getting Started..... | 2-1 |
| Installing and Licensing IP Cores..... | 2-2 |
| OpenCore Plus IP Evaluation..... | 2-2 |
| Specifying the 40-100GbE IP Core Parameters and Options..... | 2-3 |
| IP Core Parameters..... | 2-3 |
| Files Generated for the 40-100GbE IP Core..... | 2-10 |
| Simulating the IP Core..... | 2-10 |
| Integrating Your IP Core in Your Design..... | 2-11 |
| Pin Assignments..... | 2-11 |
| External Transceiver Reconfiguration Controller..... | 2-11 |
| Placement Settings for the 40-100GbE IP Core..... | 2-14 |
| 40-100GbE IP Core Testbenches..... | 2-14 |
| Testbenches with Adapters..... | 2-15 |
| Testbenches without Adapters..... | 2-18 |
| Understanding the Testbench Behavior..... | 2-19 |
| Simulating the 40-100GbE IP Core With the Testbenches..... | 2-20 |
| Generating the 40-100GbE Testbench..... | 2-21 |
| Simulating with the Modelsim Simulator..... | 2-21 |
| Simulating with the NCSim Simulator..... | 2-21 |
| Simulating with the VCS Simulator..... | 2-21 |
| Testbench Output Example: 40GbE IP Core with Adapters..... | 2-21 |
| Testbench Output Example: 100GbE IP Core with Adapters..... | 2-23 |
| Compiling the Full Design and Programming the FPGA..... | 2-24 |
| Initializing the IP Core..... | 2-24 |
| | |
| Functional Description..... | 3-1 |
| High Level System Overview..... | 3-2 |
| 40-100GbE MAC and PHY Functional Description..... | 3-2 |

| | |
|---|-------|
| 40-100GbE IP Core TX Datapath..... | 3-3 |
| 40-100GbE IP Core TX Data Bus Interfaces..... | 3-6 |
| 40-100GbE IP Core RX Datapath..... | 3-20 |
| 40-100GbE IP Core RX Data Bus Interfaces..... | 3-25 |
| 40GbE Lower Rate 24.24 Gbps MAC and PHY..... | 3-32 |
| 100GbE CAUI-4 PHY..... | 3-32 |
| External Reconfiguration Controller..... | 3-32 |
| Congestion and Flow Control Using Pause Frames..... | 3-33 |
| Pause Control and Generation Interface..... | 3-35 |
| Pause Control Frame and Non-Pause Control Frame Filtering and Forwarding..... | 3-36 |
| 40-100GbE IP Core Modes of Operation | 3-37 |
| Link Fault Signaling Interface..... | 3-37 |
| Statistics Counters Interface..... | 3-39 |
| MAC – PHY XLGMII or CGMII Interface..... | 3-42 |
| Lane to Lane Deskew Interface..... | 3-43 |
| PCS Test Pattern Generation and Test Pattern Check..... | 3-44 |
| Transceiver PHY Serial Data Interface..... | 3-45 |
| 40GBASE-KR4 IP Core Variations..... | 3-46 |
| Control and Status Interface..... | 3-51 |
| Clocks..... | 3-51 |
| Resets..... | 3-54 |
| Signals..... | 3-55 |
| Signals of MAC and PHY Variations Without Adapters..... | 3-55 |
| Signals of MAC and PHY Variations With Adapters..... | 3-66 |
| Signals of 40-100GbE MAC-Only IP Core Variations..... | 3-68 |
| Signals of 40-100GbE PHY-Only IP Core Variations..... | 3-72 |
| Software Interface: Registers..... | 3-76 |
| 40-100GbE IP Core Registers..... | 3-79 |
| 40-100GbE Example Design Registers..... | 3-116 |
| Ethernet Glossary..... | 3-119 |

Debugging the 40GbE and 100GbE Link..... 4-1

40-100GbE IP Core Example Design..... A-1

Address Map Changes for the 40-100GbE IP Core v12.0 Release..... B-1

10GBASE-KR Registers..... C-1

10GBASE-KR PHY Register Definitions..... C-1

Additional Information..... D-1

40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide Revision History
..... D-1

How to Contact Altera..... D-9

Typographic Conventions..... D-9

About the 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function

1

2014.12.15

UG-01088



Subscribe



Send Feedback

The Altera® 40- and 100-Gbps Ethernet (40GbE and 100GbE) media access controller (MAC) and PHY MegaCore® functions implement the *IEEE 802.3ba 40G and 100G Ethernet Standard* with an option to support the *IEEE 802.3ap-2007 Backplane Ethernet Standard*. This product is included in the Altera MegaCore IP Library and available from the Quartus II IP Catalog.

This product provides support for Stratix IV, Arria V GZ, and Stratix V devices. For Arria 10 40- and 100-Gbps Ethernet support, please refer to the [Low Latency 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide](#).

Note: The full product name, 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function, is shortened to 40-100GbE IP core in this document. In addition, although multiple variations are available from the parameter editor, this document refers to this product as a single IP core, because all variations are configurable from the same parameter editor.


© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered



Figure 1-1: 40GbE and 100GbE MAC and PHY MegaCore Function

Main block, internal connections, and external block requirements.



As illustrated, on the MAC client side you can choose a wide, standard Avalon® Streaming (Avalon-ST) interface, or a narrower, custom streaming interface. Depending on the variant you choose, the MAC client side Avalon Streaming (Avalon-ST) interface is either 256 or 512 bits of data mapped to either four or ten 10.3125 Gbps transceiver PHY links, depending on data rate, or to four 25.78125 Gbps transceiver PHY links.

The 40GbE (XLAUI) interface has 4x10.3125 Gbps links. The 100GbE (CAUI) interface has 10x10.3125 Gbps links. Several additional options are available. For Arria V GZ, Stratix IV, and Stratix V devices, you can configure a lower-rate 40GbE option with 4x6.25 Gbps links. For Stratix V devices only, you can configure a 40GbE 40GBASE-KR4 variation to support Backplane Ethernet. For Stratix V GT devices only, you can configure a 100GbE CAUI-4 option, with 4x25.78125 Gbps links.

The FPGA serial transceivers are compliant with the IEEE 802.3ba standard XLAUI, CAUI, and CAUI-4 specifications. The IP core configures the transceivers to implement the relevant specification for your IP core variation. You can connect the transceiver interfaces directly to an external physical medium dependent (PMD) optical module or to another device.

You can configure and generate most configurations of the 40-100GbE IP core in transmit (TX) only, receive (RX) only, or duplex mode. The 100GbE CAUI-4 option and the 40GBASE-KR4 options are available in duplex mode only.

The IP core provides standard MAC and physical coding sublayer (PCS) functions with a variety of configuration and status registers. You can exclude the statistics registers. If you exclude these registers, you can monitor the statistics counter increment vectors that the IP core provides at the client side interface and maintain your own counters.

40- and 100-Gbps Ethernet MAC and PHY IP Core Supported Features

The 40- and 100-Gbps Ethernet MAC and PHY IP core offers the following features:

- Parameterizable through the IP Catalog available with the Quartus II software.
- Compliant with the *IEEE 802.3ba-2010 High Speed Ethernet Standard* available on the IEEE website (www.ieee.org).
- Soft PCS logic that interfaces seamlessly to Altera 10.3125 Gbps and 25.78125 Gbps serial transceivers.
- Standard XLAUI or CAUI external interface consisting of serial transceiver lanes operating at 10.3125 Gbps, or the CAUI-4 external interface consisting of four serial transceiver lanes operating at 25.78125 Gbps.
- Supports 40GBASE-R4, 100GBASE-R4, and 100GBASE-R10 PHY based on 64B/66B encoding with data striping and alignment markers to align data from multiple lanes.
- Supports 40GBASE-KR4 PHY and FEC option for interfacing to backplanes
- Supports Synchronous Ethernet (Sync-E)
 - Provides CDR recovered clock output signal to the device fabric.
 - Optionally accepts two separate input reference clocks for the transmit and receive transceiver paths.
- Supports a lower-rate 40GbE option at 24.24 Gbps (4 x 6.25 Gbps line rate).
- Ethernet MAC supports the 40GbE or 100GbE line rate with a flexible and configurable feature set.
- Avalon Memory-Mapped (Avalon-MM) management interface to access the IP core control and status registers.
- Avalon-ST data path interface connects to client logic with the start of frame in the most significant byte (MSB) when optional adapters are used. Interface has data width 256 or 512 bits depending on the data rate.
- Optional custom streaming data path interface with narrower bus width and a start frame possible on 64-bit word boundaries without the optional adapters. Interface has data width 128 or 320 bits depending on the data rate.
- MAC, PHY, or MAC and PHY options configurable at IP generation.
- TX only configuration options, RX only configuration options, and duplex configuration options; the 100GbE CAUI-4 option is available only in duplex mode.
- TX and RX CRC pass-through control.
- RX and TX preamble pass-through option for applications that require proprietary user management information transfer.
- TX automatic frame padding to meet the 64-byte minimum Ethernet frame length at the 40-100GbE Ethernet connection.
- Hardware and software reset control.
- TX MAC source address insertion control.
- One MAC address register for configurable RX destination address filtering.
- RX MAC padding removal control.
- Pause frame filtering control.
- Soft error detection on all internal RAMs for high reliability systems.
- RX FIFO in MAC provides cut-through or store-and-forward frame processing.
- Deficit idle counter (DIC) to maintain a 12-byte inter-packet gap (IPG) average.

- Programmable IPG fine adjustment for Ethernet repeater/bump-in-the-wire applications and traffic shaping.
- Ethernet flow control using the pause registers or pause interface.
- Programmable maximum receive frame length up to 9600 bytes (jumbo frame) in store-and-forward mode; there is no frame size limitation for cut-through mode.
- Promiscuous (transparent) and non-promiscuous (filtered) operation modes or received frame address filtering.
- Configurable received frame filtering with cyclic redundancy check (CRC), runt, or oversized frame error.
- Optional statistics counters.
- Additional testbench logic to demonstrate Ethernet IP core behavior and customize the interface.
- Statistics real-time output status signals vector.
- Fault signaling: detects and reports local fault and generates remote fault.

The 40-100GbE IP core can support full wire line speed with a 64-byte frame length and back-to-back or mixed length traffic, up to a programmable frame size greater than 9600 bytes, with no dropped packets.

For a detailed specification of the Ethernet protocol refer to the *IEEE 802.3ba-2010 High Speed Ethernet Standard*.

Related Information

[IEEE website](#)

The *IEEE 802.3ba-2010 High Speed Ethernet Standard* is available on the IEEE website.

40-100GbE IP Core Device Family and Speed Grade Support

The following sections list the device family and device speed grade support offered by the 40-100GbE IP core:

[Device Family Support](#) on page 1-4

[40-100GbE IP Core Device Speed Grade Support](#) on page 1-5

Device Family Support

Table 1-1: Altera IP Core Device Support Levels

| Device Support Level | Definition |
|----------------------|---|
| Preliminary | The IP core is verified with preliminary timing models for this device family. The IP core meets all functional requirements, but might still be undergoing timing analysis for the device family. It can be used in production designs with caution. |

| Device Support Level | Definition |
|----------------------|--|
| Final | The IP core is verified with final timing models for this device family. The IP core meets all functional and timing requirements for the device family and can be used in production designs. |

Table 1-2: 40-100GbE IP Core Device Family Support

Shows the level of support offered by the 40-100GbE IP core for each Altera device family.

| Device Family | Support |
|--------------------------------------|---------------|
| Arria V GZ | Preliminary |
| Stratix IV (GX and GT) | Final |
| Stratix V (GX, GT, and GS) | Final |
| Other device families ⁽¹⁾ | Not supported |

Related Information

[40-100GbE IP Core Device Speed Grade Support](#) on page 1-5

40-100GbE IP Core Device Speed Grade Support

Table 1-3: Slowest Supported Device Speed Grades

Lists the slowest supported device speed grades for different variations of the 40-100GbE IP core.

| MegaCore Function | Device Family | Supported Speed Grades |
|-------------------|-----------------|------------------------|
| 40GbE | Arria V GZ | I3, C3 |
| | Stratix IV (GT) | I2 |
| | Stratix V (GX) | I3, C3 |
| | Stratix V (GT) | I3, C2 |
| | Stratix V (GS) | I3, C3 |

⁽¹⁾ This product does not provide support for Arria 10 devices. For information about Arria 10 40-100GbE support, refer to the [Low Latency 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide](#).

| MegaCore Function | Device Family | Supported Speed Grades |
|----------------------------|-----------------|------------------------|
| 40GbE (24.24 Gbps option) | Arria V GZ | I3, C3 |
| | Stratix IV (GX) | I3, C3 |
| | Stratix IV (GT) | I3 |
| | Stratix V (GX) | I3, C3 |
| | Stratix V (GT) | I3, C2 |
| | Stratix V (GS) | I3, C3 |
| 40GbE (40GBASE-KR4 option) | Stratix V (GX) | I3, C3 |
| | Stratix V (GT) | I3, C2 |
| | Stratix V (GS) | I3, C3 |
| 100GbE | Arria V GZ | I3, C3 |
| | Stratix IV (GT) | I2 |
| | Stratix V (GX) | I3, C3 |
| | Stratix V (GT) | I3, C2 |
| | Stratix V (GS) | I3, C3 |
| 100GbE (CAUI-4 option) | Stratix V GT | C2 |

IP Core Verification

To ensure functional correctness of the 40-100GbE IP core, Altera performs extensive validation through both simulation and hardware testing. Before releasing a version of the 40- and 100-Gbps Ethernet MAC and PHY IP core, Altera runs comprehensive regression tests in the current version of the Quartus® II software.

Altera verifies that the current version of the Quartus II software compiles the previous version of each IP core. Any exceptions to this verification are reported in the *Altera IP Release Notes*. Altera does not verify compilation with IP core versions older than the previous release.

Related Information

[Altera IP Release Notes](#)

Simulation Environment

Altera performs the following tests on the 40-100GbE MAC and PHY IP core in the simulation environment using internal and third party standard bus functional models (BFM):

- Constrained random tests that cover randomized frame size and contents
- Randomized error injection tests that inject Frame Check Sequence (FCS) field errors, runt packets, and corrupt control characters, and then check for the proper response from the IP core
- Assertion based tests to confirm proper behavior of the IP core with respect to the specification
- Extensive coverage of our runtime configuration space and proper behavior in all possible modes of operation

Hardware Testing

Altera performs hardware testing of the key functions of the 40-100GbE MAC and PHY IP core. The Altera hardware tests of the 40-100GbE IP core also ensure reliable solution coverage for hardware related areas such as synchronization, and reset recovery. The IP core is tested with Stratix IV and Stratix V devices.

Performance and Resource Utilization

The following sections provide performance and resource utilization data for the 40GbE and 100GbE IP cores.

Resource Utilization for 40GbE IP Cores

Resource utilization changes if the statistics counters are configured in the IP core. You can specify whether to include or not include the statistics counters in the 40-100GbE parameter editor, but you cannot change the setting dynamically.

The 24.24 Gbps variations of the 40-100GbE IP core use the same resources as the standard 40GbE IP core variations. The 40GBASE-KR4 variations require more resources only for the PHY component.

Table 1-4: 40GbE IP Core FPGA Resource Utilization in Stratix V and Arria V GZ Devices

Lists the resources and expected performance for selected variations of the 40GbE IP cores in an Arria V GZ or Stratix V device. The results were obtained using the Quartus II software v13.1 for a Stratix V 5SGXEA7N2F45C2 device.

- Top-level modules are in bold.
- The numbers of ALMs and logic registers are rounded up to the nearest 100.
- The numbers of ALMs, before rounding, are the **ALMs needed** numbers from the Quartus II Fitter Report.

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|----------|
| | | | M20K |
| MAC&PHY with Avalon-ST client interface without statistics counters | 13600 | 23500 | 9 |
| MAC&PHY with Avalon-ST client interface and with statistics counters | 17700 | 30900 | 9 |
| MAC with Avalon-ST client interface without statistics counters | 7100 | 15000 | 9 |
| MAC with Avalon-ST client interface and with statistics counters | 11300 | 22300 | 9 |
| • alt_e40_adapter_rx:adapter_rx | 500 | 900 | 0 |
| • alt_e40_adapter_tx:adapter_tx | 300 | 700 | 0 |
| MAC with custom streaming client interface without statistics counters | 6200 | 13400 | 9 |
| MAC with custom streaming client interface and with statistics counters | 10400 | 20700 | 9 |

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|----------|
| | | | M20K |
| <ul style="list-style-type: none"> alt_e40_mac_rx:mac_rx | 3000 | 7000 | 9 |
| <ul style="list-style-type: none"> alt_e40_mac_tx:mac_tx | 2600 | 4800 | 0 |
| <ul style="list-style-type: none"> alt_e40_mac_csr:mac_csr without statistics counters | 700 | 2000 | 0 |
| <ul style="list-style-type: none"> alt_e40_mac_csr:mac_csr with statistics counters | 4600 | 8500 | 0 |
| PHY | 6800 | 8600 | 0 |
| <ul style="list-style-type: none"> alt_e40_phy_pcs:phy_pcs | 6200 | 8200 | 0 |
| <ul style="list-style-type: none"> alt_e40_pcs_rx:pcs_rx | 2800 | 3800 | 0 |
| <ul style="list-style-type: none"> alt_e40_pcs_tx:pcs_tx | 2900 | 3300 | 0 |
| <ul style="list-style-type: none"> alt_e40_phy_csr:phy_csr | 500 | 1100 | 0 |
| <ul style="list-style-type: none"> alt_e40_phy_pma:phy_pma | 200 | 400 | 0 |
| 40GBASE-KR4 PHY <ul style="list-style-type: none"> No auto-negotiation (AN) No link training (LT) Forward error correction (FEC) only Use M20K blocks for FEC buffer | 14800 | 16700 | 8 |

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|----------|
| | | | M20K |
| 40GBASE-KR4 PHY <ul style="list-style-type: none"> • AN • LT • FEC • Use M20K blocks for FEC buffer | 23800 | 24500 | 8 |
| 40GBASE-KR4 PHY <ul style="list-style-type: none"> • AN • LT • FEC • Do not use M20K blocks for FEC buffer | 31900 | 41600 | 0 |

Table 1-5: 40GbE IP Core FPGA Resource Utilization in Stratix IV Devices

Lists the resources and expected performance for selected variations of the 40GbE IP cores in a Stratix IV device. The results were obtained using the Quartus II software v13.1 for a Stratix IV EP4S100G5F45C2 device.

- Top-level modules are in bold.
- The numbers of ALMs and logic registers are rounded up to the nearest 100.

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|-----------|
| | | | M9K |
| MAC&PHY with Avalon-ST client interface without statistics counters | 18100 | 25000 | 20 |
| MAC&PHY with Avalon-ST client interface and with statistics counters | 22100 | 32100 | 20 |
| MAC with Avalon-ST client interface without statistics counters | 9700 | 15200 | 20 |

| Module | ALMs | Logic Registers | Memory |
|--|--------------|-----------------|-----------|
| | | | M9K |
| MAC with Avalon-ST client interface and with statistics counters | 13700 | 22200 | 20 |
| • alt_e40_adapter_rx:adapter_rx | 700 | 1000 | 0 |
| • alt_e40_adapter_tx:adapter_tx | 500 | 800 | 0 |
| MAC with custom streaming client interface without statistics counters | 8500 | 13400 | 20 |
| MAC with custom streaming client interface and with statistics counters | 12500 | 20400 | 20 |
| • alt_e40_mac_rx:mac_rx | 4300 | 7000 | 20 |
| • alt_e40_mac_tx:mac_tx | 3400 | 4800 | 0 |
| • alt_e40_mac_csr:mac_csr without statistics counters | 1400 | 1900 | 0 |
| • alt_e40_mac_csr:mac_csr with statistics counters | 5000 | 8300 | 0 |
| PHY | 8600 | 9900 | 0 |
| • alt_e40_phy_pcs:phy_pcs | 8100 | 9400 | 0 |
| • • alt_e40_pcs_rx:pcs_rx | 3700 | 4400 | 0 |

| Module | ALMs | Logic Registers | Memory |
|---|------|-----------------|--------|
| | | | M9K |
| <ul style="list-style-type: none"> alt_e40_pcs_tx:pcs_tx | 3600 | 3900 | 0 |
| <ul style="list-style-type: none"> alt_e40_phy_csr:phy_csr | 700 | 1100 | 0 |
| <ul style="list-style-type: none"> alt_e40_phy_pma_siv:pma | 600 | 500 | 0 |

Related Information**[Fitter Resources Reports in the Quartus II Help](#)**

Information about Quartus II resource utilization reporting, including **ALMs needed**.

Resource Utilization for 100GbE IP Cores

Resource utilization changes if the statistics counters are configured in the IP core. You can specify whether to include or not include the statistics counters in the 40-100GbE parameter editor, but you cannot change the setting dynamically.

Table 1-6: 100GbE IP Core FPGA Resource Utilization in Stratix V and Arria V GZ Devices

Lists the resources and expected performance for selected variations of the 100GbE IP cores in an Arria V GZ or Stratix V device. The results were obtained using the Quartus II software v13.1 for a Stratix V 5SGXEA7N2F45C2 device.

- Top-level modules are in bold.
- The numbers of ALMs and logic registers are rounded up to the nearest 100.
- The numbers of ALMs, before rounding, are the **ALMs needed** numbers from the Quartus II Fitter Report.

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|-----------|
| | | | M20K |
| MAC&PHY with Avalon-ST client interface without statistics counters | 45100 | 87700 | 28 |
| MAC&PHY with Avalon-ST client interface and with statistics counters | 49700 | 95500 | 28 |

| Module | ALMs | Logic Registers | Memory |
|--|--------------|-----------------|-----------|
| | | | M20K |
| MAC with Avalon-ST client interface without statistics counters | 21600 | 45200 | 28 |
| MAC with Avalon-ST client interface and with statistics counters | 26100 | 53000 | 28 |
| • alt_e100_adapter_rx:adapter_rx | 2700 | 6600 | 17 |
| • alt_e100_adapter_tx:adapter_tx | 2600 | 4900 | 0 |
| MAC with custom streaming client interface without statistics counters | 16200 | 33700 | 11 |
| MAC with custom streaming client interface and with statistics counters | 20700 | 41500 | 11 |
| • alt_e100_mac_rx:mac_rx | 6500 | 14900 | 11 |
| • alt_e100_mac_tx:mac_tx | 9200 | 17500 | 0 |
| • alt_e100_mac_csr:mac_csr without statistics counters | 700 | 2000 | 0 |
| • alt_e100_mac_csr:mac_csr with statistics counters | 4700 | 8500 | 0 |
| PHY | 23500 | 42500 | 0 |

| Module | ALMs | Logic Registers | Memory |
|------------------------------|-------|-----------------|--------|
| | | | M20K |
| • alt_e100_phy_pcs:phy_pcs | 23000 | 41700 | 0 |
| • • alt_e100_pcs_rx:pcs_rx | 13600 | 26300 | 0 |
| • • alt_e100_pcs_tx:pcs_tx | 8700 | 13700 | 0 |
| • • alt_e100_phy_csr:phy_csr | 700 | 1700 | 0 |
| • alt_e100_phy_pma_sv:pma | 500 | 800 | 0 |

Table 1-7: 100GbE IP Core FPGA Resource Utilization in Stratix IV Devices

Lists the resources and expected performance for selected variations of the 100GbE IP cores in a Stratix IV device. The results were obtained using the Quartus II software v13.1 for a Stratix IV EP4S100G5F45C2 device.

- Top-level modules are in bold.
- The numbers of ALMs and logic registers are rounded up to the nearest 100.

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|-----------|
| | | | M9K |
| MAC&PHY with Avalon-ST client interface without statistics counters | 60300 | 96000 | 29 |
| MAC&PHY with Avalon-ST client interface and with statistics counters | 65200 | 102400 | 29 |
| MAC with Avalon-ST client interface without statistics counters | 30700 | 48600 | 29 |

| Module | ALMs | Logic Registers | Memory |
|--|--------------|-----------------|-----------|
| | | | M9K |
| MAC with Avalon-ST client interface and with statistics counters | 35600 | 55000 | 29 |
| <ul style="list-style-type: none"> alt_e100_adapter_rx:adapter_rx | 4100 | 6300 | 17 |
| <ul style="list-style-type: none"> alt_e100_adapter_tx:adapter_tx | 3900 | 6400 | 0 |
| MAC with custom streaming client interface without statistics counters | 23300 | 35900 | 12 |
| MAC with custom streaming client interface and with statistics counters | 26900 | 42300 | 12 |
| <ul style="list-style-type: none"> alt_e100_mac_rx:mac_rx | 9500 | 15600 | 12 |
| <ul style="list-style-type: none"> alt_e100_mac_tx:mac_tx | 12600 | 18400 | 0 |
| <ul style="list-style-type: none"> alt_e100_mac_csr:mac_csr without statistics counters | 1200 | 1900 | 0 |
| <ul style="list-style-type: none"> alt_e100_mac_csr:mac_csr with statistics counters | 4900 | 8300 | 0 |
| PHY | 8600 | 9900 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_pcs:phy_pcs | 2900 | 46900 | 0 |
| <ul style="list-style-type: none"> alt_e100_pcs_rx:pcs_rx | 16700 | 28500 | 0 |

| Module | ALMs | Logic Registers | Memory |
|--|-------|-----------------|--------|
| | | | M9K |
| <ul style="list-style-type: none"> alt_e100_pcs_tx:pcs_tx | 11200 | 16600 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_csr:phy_csr | 1100 | 1700 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_pmasiv:pma | 600 | 500 | 0 |

In the standard 100GbE variations, as in the 40GbE variations, some resource utilization numbers decrease when statistics counters are not configured in the IP core. For example, compare the values for the MAC with custom streaming client interface on a Stratix IV device with statistics counters included or not included. When counters are included, the MAC requires 26600 ALMs, but when the counters are not included, the MAC requires 23000 ALMs. The difference is 3600 ALMs. In a Stratix V device, the difference is 2900 ALMs.

Related Information

[Fitter Resources Reports in the Quartus II Help](#)

Information about Quartus II resource utilization reporting, including **ALMs needed**.

Resource Utilization for 100GbE CAUI-4 IP Cores

Resource utilization changes if the statistics counters are configured in the IP core. You can specify whether to include or not include the statistics counters in the 40-100GbE parameter editor, but you cannot change the setting dynamically.

Table 1-8: 100GbE CAUI-4 IP Core FPGA Resource Utilization

Lists the resources and expected performance for selected variations of the 100GbE CAUI-4 IP core with statistics counters included or not included. The results were obtained using the Quartus II software v13.1 for a Stratix V 5SGTMC7K2F40C2 device.

- Top-level modules are in bold.
- The numbers of ALMs and logic registers are rounded up to the nearest 100.
- The numbers of ALMs, before rounding, are the **ALMs needed** numbers from the Quartus II Fitter Report.

| Module | ALMs | Logic Registers | Memory |
|--|--------------|-----------------|-----------|
| | | | M20K |
| MAC&PHY with Avalon-ST client interface without statistics counters | 50100 | 102700 | 28 |

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|-----------|
| | | | M20K |
| MAC&PHY with Avalon-ST client interface and with statistics counters | 54600 | 110100 | 28 |
| MAC with Avalon-ST client interface without statistics counters | 21500 | 45100 | 28 |
| MAC with Avalon-ST client interface and with statistics counters | 26100 | 52800 | 28 |
| • alt_e100_adapter_rx:adapter_rx | 2700 | 6500 | 17 |
| • alt_e100_adapter_tx:adapter_tx | 2600 | 4900 | 0 |
| MAC with custom streaming client interface without statistics counters | 16200 | 33700 | 11 |
| MAC with custom streaming client interface and with statistics counters | 20700 | 41300 | 11 |
| • alt_e100_mac_rx:mac_rx | 6500 | 14800 | 11 |
| • alt_e100_mac_tx:mac_tx | 9200 | 17400 | 0 |
| • alt_e100_mac_csr:mac_csr without statistics counters | 700 | 2000 | 0 |

| Module | ALMs | Logic Registers | Memory |
|---|--------------|-----------------|----------|
| | | | M20K |
| <ul style="list-style-type: none"> alt_e100_mac_csr:mac_csr with statistics counters | 4600 | 8500 | 0 |
| PHY | 28600 | 57400 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_pcs_cau4:phy_pcs | 27200 | 55000 | 0 |
| <ul style="list-style-type: none"> alt_e100_pcs_rx_cau4:pcs_rx | 18000 | 35700 | 0 |
| <ul style="list-style-type: none"> alt_e100_pcs_tx_cau4:pcs_tx | 8400 | 17500 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_csr_cau4:phy_csr | 700 | 1700 | 0 |
| <ul style="list-style-type: none"> alt_e100_phy_pma_sv_cau4:pma | 1400 | 2500 | 0 |

Related Information**[Fitter Resources Reports in the Quartus II Help](#)**

Information about Quartus II resource utilization reporting, including **ALMs needed**.

Release Information

Table 1-9: 40-100GbE IP Core Current Release Information

| Item | Description |
|--------------|-------------|
| Version | 14.1 |
| Release Date | 2014.12.15 |

| Item | Description |
|----------------|--|
| Ordering Codes | IP-40GEMAC IP-40GEPHY IP-100GEMAC IP-100GEPHY IP-40GEMACPHY IP-100GEMACPHY IP-40GBASEKR4PHY |
| Product ID | 40Gb Ethernet MAC: 00DF 40Gb Ethernet PHY: 00E0 100Gb Ethernet MAC: 00DD 100Gb Ethernet PHY: 00DE 40GBASE-KR4 with FEC: 0113 |
| Vendor ID | 6AF7 |

2014.12.15

UG-01088



Subscribe



Send Feedback

The following sections explain how to install, parameterize, simulate, and initialize the 40-100GbE IP core:

Installing and Licensing IP Cores on page 2-2

The 40-100GbE IP core is available with the Quartus II software in the Altera MegaCore IP Library.

Specifying the 40-100GbE IP Core Parameters and Options on page 2-3

The 40-100GbE IP core supports a standard customization and generation process from the Quartus II IP Catalog.. This IP core is not supported in Qsys.

IP Core Parameters on page 2-3

The 40-100GbE parameter editor provides the parameters you can set to configure the 40-100GbE IP core and simulation testbenches.

Files Generated for the 40-100GbE IP Core on page 2-10

The Quartus II software version 14.1 generates the following output for your 40-100GbE IP core.

Simulating the IP Core on page 2-10

Integrating Your IP Core in Your Design on page 2-11

40-100GbE IP Core Testbenches on page 2-14

Altera provides a testbench and an example design with most variations of the 40-100GbE IP core. The testbench is available for simulation of your IP core, and the example design targets a C2 speed grade device and can be run on hardware. You can run the testbench to observe the IP core behavior on the various interfaces in simulation.

Simulating the 40-100GbE IP Core With the Testbenches on page 2-20

Compiling the Full Design and Programming the FPGA on page 2-24

Initializing the IP Core on page 2-24

Related Information

Managing Quartus II Projects

Refer to the "Integrating IP Cores" section of this Quartus II Handbook chapter for more information about generating an Altera IP core and integrating it in your Quartus II project.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

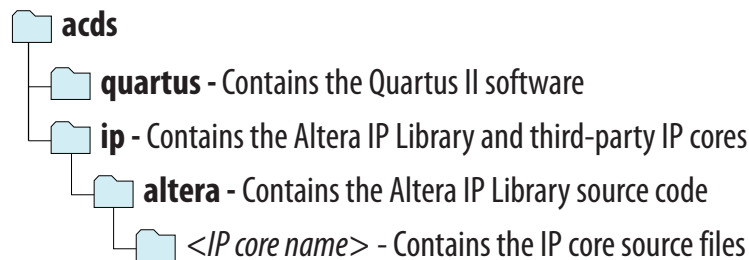
ISO
9001:2008
Registered



Installing and Licensing IP Cores

The Altera IP Library provides many useful IP core functions for production use without purchasing an additional license. You can evaluate any Altera IP core in simulation and compilation in the Quartus II software using the OpenCore® evaluation feature. Some Altera IP cores, such as MegaCore functions, require that you purchase a separate license for production use. You can use the OpenCore Plus feature to evaluate IP that requires purchase of an additional license until you are satisfied with the functionality and performance. After you purchase a license, visit the Self Service Licensing Center to obtain a license number for any Altera product.

Figure 2-1: IP Core Installation Path



Note: The default IP installation directory on Windows is `<drive>:\altera\<version number>`; on Linux it is `<home directory>/altera/ <version number>`.

Related Information

- [Altera Licensing Site](#)
- [Altera Software Installation and Licensing Manual](#)

OpenCore Plus IP Evaluation

Altera's free OpenCore Plus feature allows you to evaluate licensed MegaCore IP cores in simulation and hardware before purchase. You need only purchase a license for MegaCore IP cores if you decide to take your design to production. OpenCore Plus supports the following evaluations:

- Simulate the behavior of a licensed IP core in your system.
- Verify the functionality, size, and speed of the IP core quickly and easily.
- Generate time-limited device programming files for designs that include IP cores.
- Program a device with your IP core and verify your design in hardware

OpenCore Plus evaluation supports the following two operation modes:

- Untethered—run the design containing the licensed IP for a limited time.
- Tethered—run the design containing the licensed IP for a longer time or indefinitely. This requires a connection between your board and the host computer.

Note: All IP cores that use OpenCore Plus time out simultaneously when any IP core in the design times out.

Specifying the 40-100GbE IP Core Parameters and Options

The 40-100GbE parameter editor allows you to quickly configure your custom IP variation. Use the following steps to specify IP core options and parameters in the Quartus II software.

1. In the IP Catalog (**Tools > IP Catalog**), locate and double-click the name of the IP core to customize. The parameter editor appears.
2. Specify a top-level name for your custom IP variation. The parameter editor saves the IP variation settings in a file named `<your_ip>.qsys`. Click **OK**.
3. Specify the parameters and options for your IP variation in the parameter editor, including one or more of the following. Refer to your IP core user guide for information about specific IP core parameters.
 - Specify parameters defining the IP core functionality, port configurations, and device-specific features.
 - Specify options for processing the IP core files in other EDA tools.
4. Generate the IP core by following these steps:
 - a. Click **Generate**.
 - b. Optionally, to generate a simulation testbench or example project, follow the instructions in [Generating the 40-100GbE Testbench](#) on page 2-21.
5. Click **Finish**. The parameter editor adds the top-level `.qsys` file to the current project automatically. If you are prompted to manually add the `.qsys` file to the project, click **Project > Add/Remove Files in Project** to add the file.
6. After generating and instantiating your IP variation, make appropriate pin assignments to connect ports.

IP Core Parameters

The 40-100GbE parameter editor provides the parameters you can set to configure the 40-100GbE IP core and simulation testbenches.

The 40-100GbE parameter editor has two tabs, the **Main** tab and the **40GBASE-KR4** tab. The **40GBASE-KR4** tab in the 40-100GbE parameter editor is relevant only for certain variations that target a Stratix V device; for other variations, the parameters on the tab are unavailable.

Table 2-1: 40-100GbE Parameters: Main Tab

Describes the parameters for customizing the 40-100GbE IP core, on the Main tab of the 40-100GbE parameter editor.

| Parameter | Type | Range | Default Setting | Parameter Description |
|-------------------------------|--------|---|-----------------|----------------------------|
| General Design Options | | | | |
| Device family | String | <ul style="list-style-type: none"> • Stratix IV • Stratix V • Arria V GZ | Stratix V | Selects the device family. |

| Parameter | Type | Range | Default Setting | Parameter Description |
|---|---------|--|--|--|
| MAC configuration | String | <ul style="list-style-type: none"> 40 GbE 100 GbE | 100 GbE | Selects the MAC datapath width. |
| Core options | String | <ul style="list-style-type: none"> PHY MAC MAC & PHY | MAC & PHY | Selects the core components to generate. |
| PHY configuration ^{(2) (3) (4)} | Integer | 40 GbE: <ul style="list-style-type: none"> 24.24 Gbps (4x6.25) 40 Gbps (4x10) 100 GbE: <ul style="list-style-type: none"> 100 Gbps (10x10) CAUI-4 (4x25) | The default value depends on the MAC configuration value. 40 GbE: <ul style="list-style-type: none"> 40 Gbps (4x10) 100 GbE: <ul style="list-style-type: none"> 100 Gbps (10x10) | Selects the Ethernet speed and lane configuration. |
| MAC client interface ⁽⁵⁾ | String | <ul style="list-style-type: none"> Custom-ST interface Avalon-ST interface | Avalon-ST interface | Selects the Avalon-ST interface or the narrower, custom streaming client interface to the MAC. |
| Duplex mode ⁽⁶⁾ | Integer | <ul style="list-style-type: none"> RX TX Full Duplex | Full Duplex | Selects datapath mode to generate. |
| PHY Configuration Options | | | | |
| PHY PLL type ^{(2) (7) (8)} | String | <ul style="list-style-type: none"> ATX CMU | ATX | Configures the PHY PLL. |

⁽²⁾ This parameter is disabled in MAC-only operation.

⁽³⁾ The **PHY configuration** parameter is disabled when **MAC configuration** is set to 100GbE and **Device family** is not Stratix V. If the parameter is disabled, the IP core must always be set to the regular 10 Gbps PHY link option of 4 x 10.3125 or 10 x 10.3125.

⁽⁴⁾ For the **Device family** parameter, the CAUI-4 option requires the Stratix V GT device.

⁽⁵⁾ This parameter is disabled in PHY-only operation.

⁽⁶⁾ The **Duplex mode** parameter is disabled when **PHY configuration** is set to CAUI-4; CAUI-4 variations must always be set to the duplex configuration.

⁽⁷⁾ The **PHY PLL type** parameter is disabled when **PHY configuration** is set to CAUI-4; CAUI-4 variations must always be set to the ATX configuration.

⁽⁸⁾ The **PHY PLL type** parameter is disabled when the IP core targets a Stratix IV device; Stratix IV variations must always be set to the CMU configuration.

| Parameter | Type | Range | Default Setting | Parameter Description |
|---|--------------------|--|-----------------|---|
| PHY reference frequency ⁽²⁾ | Integer (encoding) | <p>The range and default settings depend on the PHY configuration.</p> <p>Despite its apparent availability in the parameter editor, CAUI-4 variations do not support the 322.265625 MHz clock frequency. For correct functioning of CAUI-4 variations, you must set this parameter to the value of 644.53125 MHz.</p> | | <p>Sets the expected incoming PHY <code>clk_ref</code> reference frequency. The input clock frequency must match the frequency you specify for this parameter.</p> <p>In Sync-E variations, the input clock frequencies for the <code>rx_ref_clk</code> and <code>tx_ref_clk</code> clocks must match the frequency you specify for this parameter, although the two clocks can be driven from different sources and need not be aligned with each other.</p> |

Advanced Design Options

| | | | | |
|---|---------|---|--|---|
| Status clock rate ⁽²⁾ | Float | <ul style="list-style-type: none"> Stratix IV: 37.5–50.0 Arria V GZ or Stratix V: 100.0–125.0 | <ul style="list-style-type: none"> Stratix IV: 37.5 Arria V GZ or Stratix V: 100.0 | Sets the clock rate of <code>clk_status</code> in MHz. |
| Statistics counters ⁽⁵⁾ | Boolean | <ul style="list-style-type: none"> True False | True | If turned on, the IP core includes built-in statistics counters. If turned off, the IP core is configured without statistics counters. |
| Enable SyncE support | Boolean | <ul style="list-style-type: none"> True False | False | Enables or disables a separate reference clock for the RX CDR block in the transceiver and exposes the RX recovered clock as an output signal. If this option is turned on (set to true), the TX PLL and the RX CDR in the transceiver have separate input reference clocks and the RX recovered clock is visible as an IP core output signal. If it is turned off, the two PLLs share one input reference clock and the RX recovered clock is not available as an output signal. |

Table 2-2: 40-100GbE Parameters: 40GBASE-KR4 Tab

Describes the parameters for customizing a 40GBASE-KR4 40-100GbE IP core, on the 40GBASE-KR4 tab of the 40-100GbE parameter editor. The parameters on this tab are available only if the following conditions hold:

- Your IP core targets a Stratix V device. You set the target device family for your Quartus II project or in the Quartus II software before you access the IP Catalog.
- You select the value of **40 GbE** for the **MAC configuration** parameter on the Main tab.
- You select a **Core option** value that includes a PHY component (**PHY** or **MAC & PHY**) on the Main tab.
- You select the value of **40 Gbps (4x10)** for the **PHY configuration** parameter on the Main tab.
- You select the value of **Full Duplex** for the **Duplex mode** parameter on the Main tab.

| Parameter | Type | Range | Default Setting | Parameter Description |
|---|--------------------|---|-----------------|---|
| KR4 General Options | | | | |
| Enable KR4 | Boolean | <ul style="list-style-type: none"> • True • False | False | If this parameter is turned on, the IP core is a 40GBASE-KR4 variation. If this parameter is turned off, the IP core is not a 40GBASE-KR4 variation, and the other parameters on this tab are not available. |
| Enable KR4 Reconfiguration | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core supports dynamic analog reconfiguration through a dedicated reconfiguration interface. If this parameter is turned off, the IP core cannot support auto-negotiation (AN) or link training (LT) modes, and the AN and LT parameters on this tab are not available. This parameter does not affect FEC availability. |
| Auto-Negotiation | | | | |
| Enable Auto-Negotiation | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core includes logic to implement auto-negotiation as defined in Clause 73 of <i>IEEE Std 802.3ap-2007</i> . If this parameter is turned off, the IP core does not include auto-negotiation logic and cannot perform auto-negotiation. Currently the IP core can only negotiate to KR4 mode. |
| Link fail inhibit time for 40Gb Ethernet | Integer (Unit: ms) | 500–510 ms | 504 ms | Specifies the time before <code>link_status</code> is set to FAIL or OK. A link fails if the time duration specified by this parameter expires before <code>link_status</code> is set to OK. For more information, refer to <i>Clause 73 Auto-Negotiation for Backplane Ethernet</i> in <i>IEEE Standard 802.3ap-2007</i> . The 40GBASE-KR4 IP core asserts the <code>lanes_deskewed</code> signal to indicate <code>link_status</code> is OK. |

| Parameter | Type | Range | Default Setting | Parameter Description |
|--------------------------------------|---------|--|-----------------|--|
| Auto-Negotiation Master | String | <ul style="list-style-type: none"> Lane 0 Lane 1 Lane 2 Lane 3 | Lane 0 | Selects the master channel for auto-negotiation. |
| Pause ability-C0 | Boolean | <ul style="list-style-type: none"> True False | True | If this parameter is turned on, the IP core supports symmetric pauses as defined in <i>Annex 28B</i> of Section 2 of <i>IEEE Std 802.3-2008</i> . |
| Pause ability-C1 | Boolean | <ul style="list-style-type: none"> True False | True | If this parameter is turned on, the IP core supports asymmetric pauses as defined in <i>Annex 28B</i> of Section 2 of <i>IEEE Std 802.3-2008</i> . |
| Link Training: PMA Parameters | | | | |
| VMAXRULE | Integer | 0-63 | 60 | Specifies the maximum V_{OD} . The default value, 60, represents 1200 mV. |
| VMINRULE | Integer | 0-63 | 9 | Specifies the minimum V_{OD} . The default value, 9, represents 165 mV. |
| VODMINRULE | Integer | 0-63 | 24 | Specifies the minimum V_{OD} for the first tap. The default value, 24, represents 440 mV. |
| VPOSTRULE | Integer | 0-31 | 31 | Specifies the maximum value that the internal algorithm for pre-emphasis will ever test in determining the optimum post-tap setting. |
| VPRERULE | Integer | 0-15 | 15 | Specifies the maximum value that the internal algorithm for pre-emphasis will ever test in determining the optimum pre-tap setting. |
| PREMAINVAL | Integer | 0-63 | 60 | Specifies the Preset V_{OD} value. This value is set by the Preset command of the link training protocol, defined in Clause 72.6.10.2.3.1 of <i>IEEE Std 802.3ap-2007</i> . |
| PREPOSTVAL | Integer | 0-31 | 0 | Specifies the preset Post-tap value. |
| PREPREVAL | Integer | 0-15 | 0 | Specifies the preset Pre-tap value. |
| INITMAINVAL | Integer | 0-63 | 52 | Specifies the initial V_{OD} value. This value is set by the Initialize command of the link training protocol, defined in Clause 72.6.10.2.3.2 of <i>IEEE Std 802.3ap-2007</i> . |
| INITPOSTVAL | Integer | 0-31 | 30 | Specifies the initial Post-tap value. |

| Parameter | Type | Range | Default Setting | Parameter Description |
|--|---------|---|-----------------|--|
| INITPREVAL | Integer | 0–15 | 5 | Specifies the initial Pre-tap value. |
| Link Training: General | | | | |
| Enable Link Training | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core includes the link training module, which configures the remote link partner TX PMD for the lowest Bit Error Rate (BER). LT is defined in Clause 72 of <i>IEEE Std 802.3ap-2007</i> . |
| Enable microprocessor interface | Boolean | <ul style="list-style-type: none"> • True • False | False | If this parameter is turned on, the IP core includes a dedicated interface through which you can control the link training. |
| Enable RX equalization | Boolean | <ul style="list-style-type: none"> • True • False | False | If this parameter is turned on, the IP core includes the RX part of the link training module. This part of the link training configures local receiver RX Continuous Linear Time Equalizer (CTLE) and Decision Feedback Equalizer (DFE) to achieve the lowest Bit Error Rate (BER) . |
| Maximum bit error count | Integer | $2^n - 1$ for n an integer in the range 4–12. | 4095 | <p>Specifies the maximum number of errors on a lane before the <code>Link Training Error</code> bit (40GBASE-KR4 register offset 0xD2, bit 4, 12, 20, or 28, depending on the lane) is set, indicating an unacceptable bit error rate.</p> <p>n is the width of the Bit Error Counter that is configured in the IP core. The value to which you set this parameter determines n, and thus the width of the Bit Error Counter. Because the default value of this parameter is 4095, the default width of the Bit Error Counter is 12 bits.</p> <p>You can use this parameter to tune PMA settings. For example, if you see no difference in error rates between two different sets of PMA settings, you can increase the width of the bit error counter to determine if a larger counter enables you to distinguish between PMA settings.</p> |
| Number of frames to send before sending actual data | Integer | <ul style="list-style-type: none"> • 127 • 255 | 127 | Specifies the number of additional training frames the local link partner delivers to ensure that the link partner can correctly detect the local receiver state. |
| FEC Options | | | | |

| Parameter | Type | Range | Default Setting | Parameter Description |
|---|---------|---|-----------------|--|
| Include FEC sublayer | Boolean | <ul style="list-style-type: none"> • True • False | False | If this parameter is turned on, the IP core includes logic to implement FEC. |
| Set FEC_ability bit on power up/reset | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core sets the FEC ability bit (40GBASE-KR4 register offset 0xB0, bit 16) on power up and reset. |
| Set FEC_Enable bit on power up/reset | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core sets the FEC enable bit (40GBASE-KR4 register offset 0xB0, bit 18) on power up and reset. If you turn on this parameter but do not turn on Set FEC_ability bit on power up/reset , this parameter has no effect: the IP core cannot specify the value of 1 for FEC Requested without specifying the value of 1 for FEC Ability. |
| Set FEC_Error_indication_ability bit on power up/reset | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core is programmed by default (40GBASE-KR4 register offset 0xB0, bit 17) to report decoding errors to the PCS. |
| Use M20K for FEC Buffer (if available) | Boolean | <ul style="list-style-type: none"> • True • False | True | If this parameter is turned on, the IP core is configured with a pipelined FEC buffer to support the Quartus II software in inferring M20K memory. Turning on this parameter potentially saves device resources. |

Table 2-3: 40-100GbE PHY Parameter Settings

Lists the PHY parameters that are configured automatically based on parameter values you select in the 40G/100G Ethernet parameter editor.

| Parameter | 40GbE Value 40GBASE-KR4 Value | 100GbE Value | 40GbE at 24.24 Gbps | 100GbE at CAUI-4 |
|---|----------------------------------|---------------------------------|-----------------------------|------------------|
| Lanes | 4 | 10 | 4 | 4 |
| Data rate per lane | 10312.5 Mbps | 10312.5 Mbps | 6250.0 Mbps | 25781.25 Mbps |
| Available PHY reference clock frequencies | 644.53125 MHz 322.265625 MHz | 644.53125 MHz 322.265625 MHz | 390.625 MHz 195.3125 MHz | 644.53125 MHz |

Related Information


- [Clocks](#) on page 3-51
The range and default settings for the **PHY reference frequency** parameter depend on the **PHY configuration** parameter value. The **PHY reference frequency** value is the required frequency of the transceiver reference clock or transceiver reference clocks.

- [40-100GbE IP Core Testbenches](#) on page 2-14
Provides a list of IP core variations (parameter value choices) for which the Quartus II software can generate a testbench and example design if you turn on **Generate example design**.

Files Generated for the 40-100GbE IP Core

The Quartus II software version 14.1 generates the following output for your 40-100GbE IP core.

Figure 2-2: IP Core Generated Files



Notes:

1. If generated for your IP variation

Simulating the IP Core

You can simulate your 40GbE or 100GbE IP core variation with the functional simulation model and the testbench or example design generated with the IP core. The functional simulation model is a cycle-accurate model that allows for fast functional simulation of your IP core instance using industry-standard VHDL or Verilog HDL simulators. If your IP core variation does not generate a matching testbench, you can create your own testbench to exercise the IP core functional simulation model.

The functional simulation model and testbench files are generated in project subdirectories. These directories also include scripts to compile and run the example design.

Note: Use the simulation models only for simulation and not for synthesis or any other purposes. Using these models for synthesis creates a nonfunctional design.

In the top-level wrapper file for your simulation project, you can set the `FAST_SIMULATION` parameter to enable simulation optimization. Parameters are set through the IP core parameter editor. In general, you should not change them manually. The only exception is the `FAST_SIMULATION` parameter. You should set the `FAST_SIMULATION` parameter on the PHY blocks by adding the following line to the top-level wrapper file:

```
defparam <dut instance>.FAST_SIMULATION = 1;
```

Note: You can use the example testbench as a guide for setting the simulation parameters in your own simulation environment. This line is already present in the Altera-provided testbench that is generated with the IP core.

Related Information

- [Simulating the 40-100GbE IP Core With the Testbenches](#) on page 2-20
Instructions to simulate the 40GbE or 100GbE IP core with the IP core appropriate testbench you can generate.
- [40-100GbE IP Core Testbenches](#) on page 2-14
Altera provides a testbench and example design with most variations of the 40-100GbE IP core. The testbench is available for simulation of your IP core, and the example design can be run on hardware. This topic describes the testbench provided with the 40-100GbE IP core. For a complete list of models or libraries required to simulate your IP core, refer to the scripts provided with the testbench.
- [Simulating Altera Designs](#)
Chapter in volume 3 of the *Quartus II Handbook* that provides information about simulating Altera IP cores.

Integrating Your IP Core in Your Design

When you integrate your IP core instance in your design, you must pay attention to the following items:

[Pin Assignments](#) on page 2-11

[External Transceiver Reconfiguration Controller](#) on page 2-11

[Placement Settings for the 40-100GbE IP Core](#) on page 2-14

Pin Assignments

When you integrate your 40-100GbE IP core instance in your design, you must make appropriate pin assignments. You can create a virtual pin to avoid making specific pin assignments for top-level signals while you are simulating and not ready to map the design to hardware.

Related Information

[Quartus II Help](#)

For information about the Quartus II software, including virtual pins and the IP Catalog.

External Transceiver Reconfiguration Controller

40-100GbE IP cores that include the PHY component require an external reconfiguration controller to compile and to function correctly in hardware.

You can use the IP Catalog to generate an Altera transceiver reconfiguration controller.

- For Arria V GZ and Stratix V devices, select the Transceiver Reconfiguration Controller.
- For Stratix IV devices, select the `ALTGX_RECONFIG` transceiver reconfiguration block.

When you configure the Altera Transceiver Reconfiguration Controller, you must specify the number of reconfiguration interfaces. The number of reconfiguration interfaces required for the 40GbE and 100GbE IP cores depends on the IP core variation.

Table 2-4: Number of Reconfiguration Interfaces

Lists the number of reconfiguration interfaces you should specify for the Altera Transceiver Reconfiguration Controller for your Arria V GZ or Stratix V 40-100GbE IP core that includes a PHY component.

| PHY Configuration | RX Only | TX Only | Duplex |
|--|---------|---------|--------------------|
| Standard 40GbE and 40GBASE-KR4 (4x10.3125 lanes) | 4 | 8 | 8 |
| 100GbE (10x10.3125 lanes) | 10 | 20 | 20 |
| CAUI-4 (4x25.78125 lanes) | — | — | 4x3 ⁽⁹⁾ |

You can configure your reconfiguration controller with additional interfaces if your design connects with multiple transceiver IP cores. You can leave other options at the default settings or modify them for your preference.

You should connect the `reconfig_to_xcvr` and `reconfig_from_xcvr` ports of the 40-100GbE IP core to the corresponding ports of the reconfiguration controller.

The CAUI-4 variations have four reconfiguration channels, numbered consecutively from `reconfig_to_xcvr0` and `reconfig_from_xcvr0` to `reconfig_to_xcvr3` and `reconfig_from_xcvr3`. The CAUI-4 reconfiguration channels must be connected to the four reconfiguration controller groupings. The reconfiguration controller groupings include `ch0_2_from_xcvr`, `ch3_5_from_xcvr`, `ch6_8_from_xcvr`, and `ch9_11_from_xcvr`.

You must also connect the `mgmt_clk_clk` and `mgmt_rst_reset` ports of the Altera Transceiver Reconfiguration Controller. The `mgmt_clk_clk` port must have a clock setting in the range of 100–125MHz; this setting can be shared with the 40-100GbE IP core `clk_status` port. The `mgmt_rst_reset` port must be deasserted before, or deasserted simultaneously with, the 40-100GbE IP core `pma_arst_ST` port.

⁽⁹⁾ The CAUI-4 configuration requires 12 interfaces split into four groups of three; the interface grouping should be set to 3, 3, 3, 3.

Table 2-5: External Altera Transceiver Reconfiguration Controller Ports for Connection to 40-100GbE IP Core

| Signal Name | Direction | Description |
|---|-----------|---|
| reconfig_to_xcvr[559:0](40GbE) reconfig_to_xcvr[1399:0](100GbE) | Input | The 40-100GbE IP core reconfiguration controller to transceiver port in non-CAUI-4 configurations. Available only in the PHY and MAC & PHY configurations for Arria V GZ and Stratix V devices. |
| reconfig_from_xcvr[367:0](40GbE) reconfig_from_xcvr[919:0](100GbE) | Output | The 40-100GbE IP core reconfiguration controller from transceiver port in non-CAUI-4 configurations. Available only in the PHY and MAC & PHY configurations for Arria V GZ and Stratix V devices. |
| reconfig_to_xcvr0 | Input | The reconfiguration channel to CAUI-4 lane 0. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_to_xcvr1 | Input | The reconfiguration channel to CAUI-4 lane 1. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_to_xcvr2 | Input | The reconfiguration channel to CAUI-4 lane 2. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_to_xcvr3 | Input | The reconfiguration channel to CAUI-4 lane 3. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_from_xcvr0 | Output | The reconfiguration channel from CAUI-4 lane 0. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_from_xcvr1 | Output | The reconfiguration channel from CAUI-4 lane 1. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_from_xcvr2 | Output | The reconfiguration channel from CAUI-4 lane 2. Available only in the 100GbE CAUI-4 PHY configuration. |
| reconfig_from_xcvr3 | Output | The reconfiguration channel from CAUI-4 lane 3. Available only in the 100GbE CAUI-4 PHY configuration. |

Related Information

- [Altera Transceiver PHY IP Core User Guide](#)

For more information about the Altera Transceiver Reconfiguration Controller.

- [ALTGX_RECONFIG Megafunction User Guide for Stratix IV Devices](#)
Chapter in *volume 3: Transceiver Configuration Guide* of the *Stratix IV Device Handbook*. Describes the ALTGX_RECONFIG megafunction, which configures a transceiver reconfiguration block for a Stratix IV device.

Placement Settings for the 40-100GbE IP Core

The Quartus II software provides the options to specify design partitions and LogicLock™ regions for incremental compilation, to control placement on the device. To achieve timing closure for your design, you might need to provide floorplan guidelines using one or both of these features.

The appropriate floorplan is always design-specific, and depends on your full design.

Related Information

[Quartus II Handbook Volume 2: Design Implementation and Optimization](#)

Describes incremental compilation, design partitions, and LogicLock regions.

40-100GbE IP Core Testbenches

Altera provides a testbench and an example design with most variations of the 40-100GbE IP core. The testbench is available for simulation of your IP core, and the example design targets a C2 speed grade device and can be run on hardware. You can run the testbench to observe the IP core behavior on the various interfaces in simulation.

Altera offers testbenches for the following configurations:

- Non-40GBASE-KR4 IP core variations that have all of the following properties:
 - Includes both MAC and PHY components (**Core options** has the value of **MAC & PHY**)
 - Full duplex (**Duplex mode** has the value of **Full Duplex**)
- 40GBASE-KR4 IP core variations that have all of the following properties:
 - Includes both MAC and PHY components (**Core options** has the value of **MAC & PHY**)
 - With adapters (**MAC client interface** has the value of **Avalon-ST interface**)
 - Without Synchronous Ethernet support (**Enable SyncE support** is turned off)
 - Without the link training microprocessor interface (**Enable microprocessor interface** is turned off)
 - RX equalization enabled (**Enable RX equalization** is turned on)

When you generate your IP core and turn on **Generate example design**, the Quartus II software generates the testbench and example design for your variation. If your IP core variation does not meet the criteria for a testbench, the generation process does not create a testbench. Turning on **Generate example design** does not force the software to generate a testbench if none is defined for your variation.

MAC-only, PHY-only, TX-only, and RX-only IP core variations do not generate an example design and testbench. 40GBASE-KR4 IP core variations with the custom streaming interface, without RX equalization enabled, with Synchronous Ethernet support, or with the link training microprocessor interface, do not generate a testbench. (However, 40GBASE-KR4 IP core variations that conform to all the requirements with the exception of the requirement of adapters, do generate an example design that runs in hardware).

Conceptually, the testbenches for the 40-100GbE IP cores with adapters (IP cores with an Avalon-ST client interface) and the testbenches for the 40-100GbE IP cores without adapters (IP cores with the

custom streaming client interface) are identical, except for the bandwidth. The following sections first describe the testbenches that include adapters and then describe the testbenches without adapters.

You can simulate the testbench that you generate with your IP core variation. The testbench illustrates packet traffic, in addition to providing information regarding the transceiver PHY. The non-40GBASE-KR4 testbenches tie off the reconfiguration control interface for your IP core, and do not exercise transceiver reconfiguration. However, the 40GBASE-KR4 testbench exercises auto-negotiation and link training, in addition to generating and checking packet traffic.

Related Information

- [40-100GbE IP Core Example Design](#) on page 5-1
Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a C2 speed grade device.
- [Simulating the IP Core](#) on page 2-10
- [Simulating the 40-100GbE IP Core With the Testbenches](#) on page 2-20
Instructions to simulate the 40GbE or 100GbE IP core with the IP core appropriate testbench you can generate, including simulation parameters and supported simulators.

Testbenches with Adapters

Figure 2-3: 40-100GbE IP Core Testbenches with Adapters

Illustrates the top-level modules of the non-40GBASE-KR4 40GbE and 100GbE example testbenches that use adapters. In the file names, * denotes 40 for 40GbE IP cores and 100 for 100GbE IP cores.

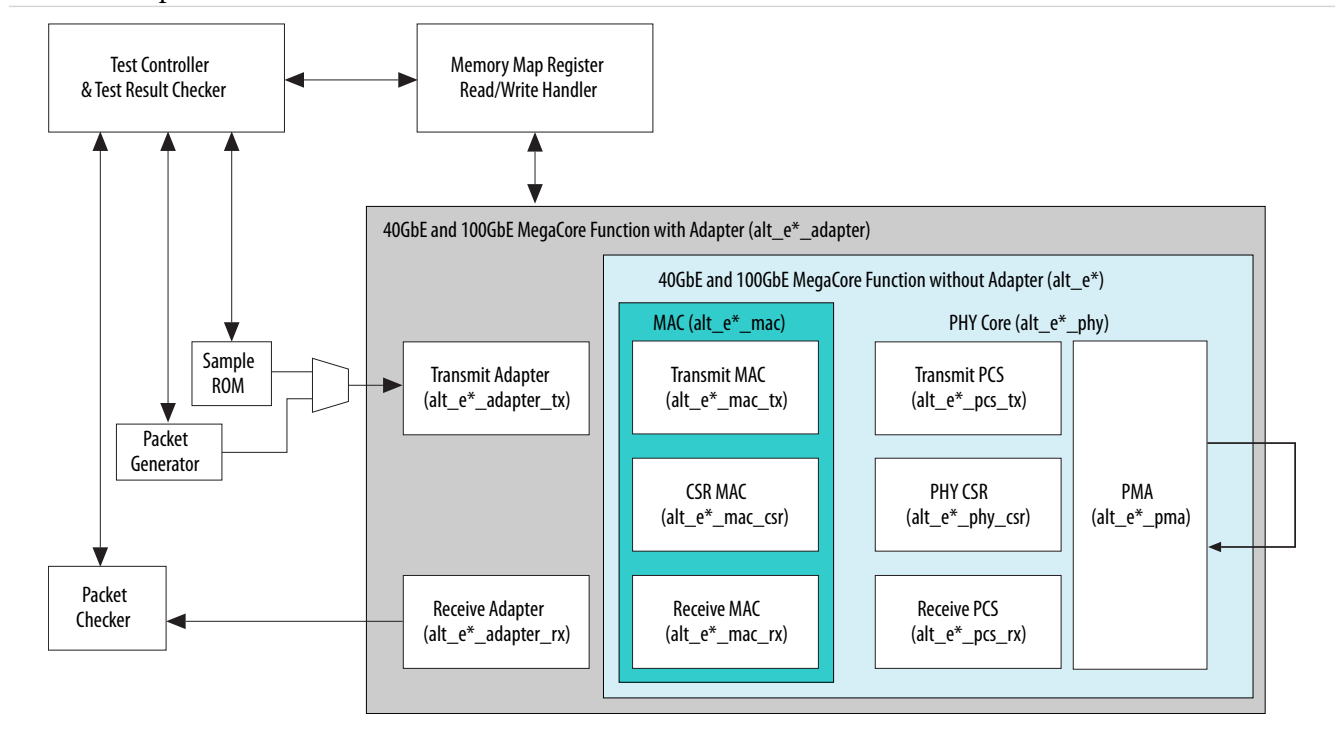


Figure 2-4: 40GBASE-KR4 40GbE IP Core Testbench with Adapters

Illustrates the top-level modules of the 40GBASE-KR4 example testbench that uses adapters. To support the simulation of auto-negotiation, the testbench uses two instances of the IP core instead of configuring the IP core in loopback mode.




Table 2-6: 40-100GbE IP Core Testbench with Adapters File Descriptions

Lists the key files that implement the example testbenches.


| File Names | Description |
|---|---|
| Testbench and Simulation Files | |
| alt_40gbe_tb.sv, alt_e40_avalon_kr4_tb.sv, alt_100gbe_tb.v | The testbench wrapper file. For non-KR4 variations, this file includes all of the testbench modules. |
| alt_e40_avalon_tb_packet_gen.v | The packet generator. This file is present only for 40GBASE-KR4 variations. |
| alt_e40_avalon_tb_packet_gen_sanity_check.v | The packet checker. This file is present only for 40GBASE-KR4 variations. |
| alt_e40_avalon_tb_sample_tx_rom.hex | The sample TX ROM. This file is present only for 40GBASE-KR4 variations. |
| alt_e40_avalon_tb_sample_tx_rom.v | Lists the contents of the sample TX ROM (alt_e40_avalon_tb_sample_tx_rom.hex). This file is present only for 40GBASE-KR4 variations. |
| Testbench Scripts | |
| run_vsim.do | The ModelSim script to run the testbench. |

| File Names | Description |
|---------------------------|--|
| <code>run_vcs.sh</code> | The Synopsys VCS script to run the testbench. |
| <code>run_ncsim.sh</code> | The Cadence NCSim script to run the testbench. |

Figure 2-5: Typical 40GbE Traffic on the Avalon-ST Interface Using the Four- to Two-Word Adapters

Shows typical traffic from the simulation testbench created using the `<instance_name>_example/alt_e40_e100/example_testbench/run_vsim.do` script in ModelSim.

Note: Client logic must maintain the `l4_tx_valid` signal asserted while asserting SOP, through the assertion of EOP. Client logic should not pull this signal low during a packet transmission.



The markers in the figure show the following sequence of events:

1. At marker 1, the application asserts `l4_tx_startofpacket`, indicating the beginning of a TX packet.
2. At marker 2, the application asserts `l4_tx_endofpacket`, indicating the end of the TX packet. The value on `l4_tx_empty[4:0]` indicates that the 2 least significant bytes of the last data cycle are empty.
3. At marker 3, the IP core asserts `l4_rx_startofpacket`, indicating the beginning of an RX packet. A second transfer has already started on the TX bus.
4. At marker 4, the 40GbE IP core deasserts `l4_rx_valid`, indicating that the IP core does not have new valid data to send to the client on `l4_rx_data[255:0]`. `l4_rx_data[255:0]` remains valid and unchanged for a second cycle.
5. At marker 5, the 40GbE IP core asserts `l4_rx_valid`, indicating that it has valid data to send to the client on `l4_rx_data[255:0]`.
6. At marker 6, the 40GbE IP core deasserts `l4_rx_valid`, indicating that it does not have new valid data to send to the client on `l4_rx_data[255:0]`. `l4_rx_data[255:0]` remains unchanged for a second cycle.

7. At marker 7, the 40GbE IP core asserts `l4_rx_valid`, indicating that the it has valid data to send to the client on `l4_rx_data[255:0]`.
8. At marker 8, the 40GbE IP core deasserts `l4_rx_valid`, indicating that the 40GbE IP core does not have new valid data to send to the client on `l4_rx_data[255:0]`. `l4_rx_data[255:0]` remains unchanged for a second cycle.
9. At marker 9, the IP core asserts `l4_rx_endofpacket`, indicating the end of the RX packet. `l4_rx_empty[4:0]` has a value of `0x1D`, indicating that 29 least significant bytes of the last cycle of the RX packet empty.

Note: The ready latency on the TX client interface with adapters is 0.

Related Information

[Avalon Interface Specifications](#)

For more information about the Avalon-ST protocol.

Testbenches without Adapters

Figure 2-6: 40-100GbE IP Core Testbench Without Adapters

Illustrates the top-level modules of the 40GbE and 100GbE example testbenches that do not use adapters. In the file names, * denotes 40 for 40GbE IP cores and 100 for 100GbE IP cores.

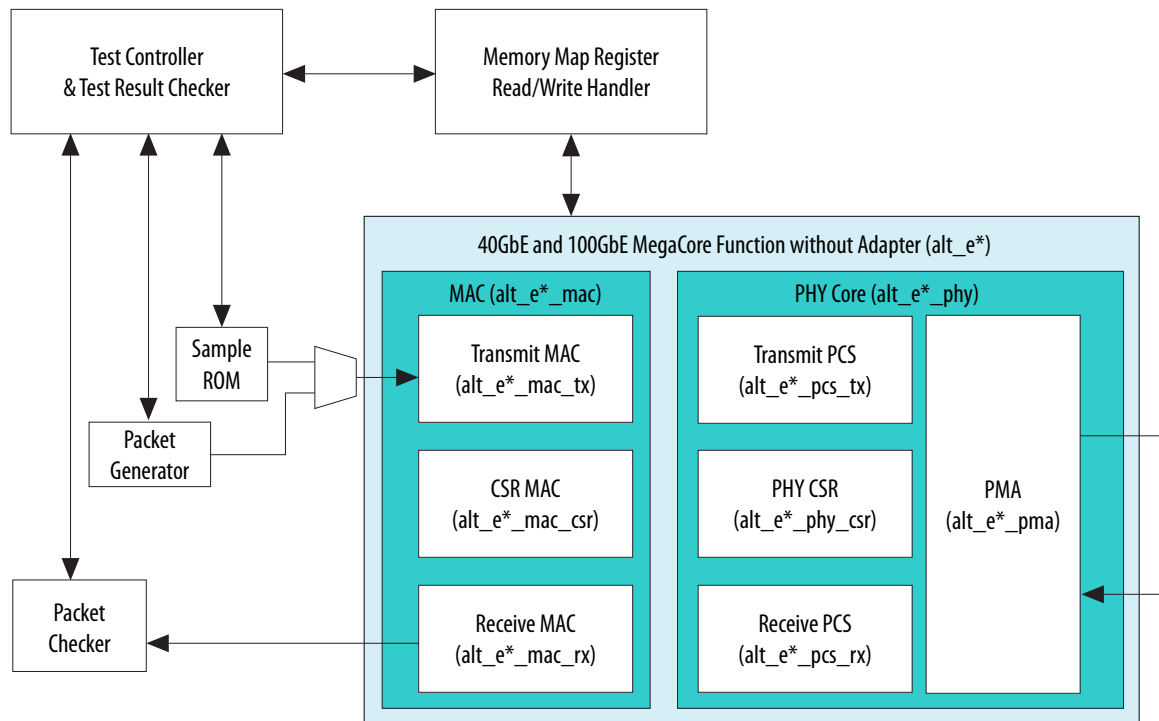


Table 2-7: 40-100GbE IP Core Testbench Without Adapters File Descriptions

Lists the key files that implement the example testbenches.

| File Names | Description |
|--------------------------------|-------------|
| Testbench and Simulation Files | |

| File Names | Description |
|---|--|
| alt_40gbe_tb.sv, alt_100gbe_tb.v | The testbench wrapper file. |
| Testbench Scripts | |
| run_vsim.do | The ModelSim script to run the testbench. |
| run_vcs.sh | The Synopsys VCS script to run the testbench. |
| run_ncsim.sh | The Cadence NCSim script to run the testbench. |

Understanding the Testbench Behavior

The non-40GBASE-KR4 testbenches send traffic through the IP core in loopback mode, exercising the transmit side and receive side of the IP core in the same data flow. These testbenches send traffic to allow the Ethernet lanes to lock, and then send packets to the transmit client data interface and check the data as it returns through the receive client data interface.

The 40GBASE-KR4 testbench sends traffic through the two IP cores in each direction, exercising the receive and transmit sides of both IP cores. This testbench exercises auto-negotiation and link training, and then sends and checks packets in data mode.

The 40-100GbE IP cores implement virtual lanes as defined in the *IEEE 802.3ba-2010 40G and 100G Ethernet Standard*. The 40GbE IP cores are fixed at four virtual lanes and each lane is sent over a 10 Gbps physical lane. The 100GbE IP cores are fixed at 20 virtual lanes; the 20 virtual lanes are typically bit-interleaved over ten 10-Gbps physical lanes. When the lanes arrive at the receiver the lane streams are in an undefined order. Each lane carries a periodic PCS-VLANE alignment tag to restore the original ordering. The simulation establishes a random permutation of the physical lanes that is used for the remainder of the simulation.

Within each virtual lane stream, the data is 64B/66B encoded. Each word has two framing bits which are always either 01 or 10, never 00 or 11. The RX logic uses this pattern to lock onto the correct word boundaries in each serial stream. The process is probabilistic due to false locks on the pseudo-random scrambled stream. To reduce hardware costs, the receiver does not test alignments in parallel; consequently, the process can be somewhat time-consuming in simulation.

In the 40GBASE-KR4 testbench, some register values are set to produce a shorter runtime. For example, timeout counters and the number of steps used in link training are set to smaller values than would be prudent in hardware. To override this behavior and use the normal settings in simulation, add the following line to your IP core variation top-level file or to the testbench top-level file,

alt_e40_avalon_kr4_tb.sv:

```
`define ALTERA_RESERVED_XCVR_FULL_KR_TIMERS
```

Both the word lock and the alignment marker lock implement hysteresis as defined in the *IEEE 802.3ba-2010 40G and 100G Ethernet Standard*. Multiple successes are required to acquire lock and multiple failures are required to lose lock. The “fully locked” messages in the simulation log indicate the point at which a physical lane has successfully identified the word boundary and virtual lane assignment.

In the event of a catastrophic error, the RX PCS automatically attempts to reacquire alignment. The MAC properly identifies errors in the datastream.

Simulating the 40-100GbE IP Core With the Testbenches

You can simulate the 40-100GbE IP core using the Altera-supported versions of the Mentor Graphics ModelSim® SE, Cadence NCSim, and Synopsys VCS simulators for the current version of the Quartus II software.

The example testbenches simulate packet traffic at the digital level. The testbenches do not require special SystemVerilog class libraries.

The top-level testbench file for non-40GBASE-KR4 variations consists of a simple packet generator and checker and one core in a loopback configuration. The packet generator skews and reorders its transmitter digital output to emulate actual transceiver behavior and optical cabling lane permutations.

The top-level testbench file for 40GBASE-KR4 variations consists of a symmetric arrangement with two IP cores and traffic between them. For each IP core there is a packet generator to send traffic on the TX side of the IP core and a packet checker to check the packets it receives from the other IP core. The two IP cores communicate with each other through their Ethernet link, in which the testbench injects random skew. The 40GBASE-KR4 testbench connects each IP core to a reconfiguration bundle, and exercises auto-negotiation, link training, and data mode.

The example testbenches contain the test files and run scripts for the ModelSim, Cadence, and Synopsys simulators. The run scripts use the file lists in the wrapper files. When you launch a simulation from the original directory, the relative filenames in the wrapper files allow the run script to locate the files correctly. You can access design files from any location if your directory structure matches the structure assumed in the run script path names.

The following examples provide directions for generating the testbench and running tests with the ModelSim, Cadence, and Synopsys simulators.

[Generating the 40-100GbE Testbench](#) on page 2-21

[Simulating with the Modelsim Simulator](#) on page 2-21

[Simulating with the NCSim Simulator](#) on page 2-21

[Simulating with the VCS Simulator](#) on page 2-21

[Testbench Output Example: 40GbE IP Core with Adapters](#) on page 2-21

[Testbench Output Example: 100GbE IP Core with Adapters](#) on page 2-23

Related Information

- [Simulating the IP Core](#) on page 2-10
- [40-100GbE IP Core Testbenches](#) on page 2-14

Altera provides a testbench and an example design with most variations of the 40-100GbE IP core. The testbench is available for simulation of your IP core, and the example design targets a C2 speed grade device and can be run on hardware. You can run the testbench to observe the IP core behavior on the various interfaces in simulation.

Generating the 40-100GbE Testbench

A single procedure generates both the testbench and the example project. To generate the testbench and example project:

1. Follow the steps in [Specifying the 40-100GbE IP Core Parameters and Options](#) to parameterize your IP core.
2. Generate the IP core by clicking **Generate**.

Note: When prompted at the start of generation, you must turn on **Generate example design**. Turning on **Generate example design** is the only process that generates a functional testbench and a functional example design.

When the IP core is generated in *<working directory>*, the testbench and example project are generated in *<working directory>/<IP core variation>/_example/alt_e40_e100*.

The directory with the testbench and example project has two subdirectories:

- **example**, which contains the example design project
- **example_testbench**, which contains the demonstration testbench

Simulating with the Modelsim Simulator

To run the simulation in the ModelSim simulation tool, follow these steps:

1. Change directory to the *<variation_name>* **_example/alt_e40_e100/example_testbench** directory.
2. In the command line, type: `vsim -c -do run_vsim.do`

The example testbench will run and pass.

Simulating with the NCSim Simulator

To run the simulation in the NCSim simulation tool, follow these steps:

1. Change directory to the *<variation_name>* **_example/alt_e40_e100/example_testbench** directory.
2. In the command line, type: `sh run_ncsim.sh`

The example testbench will run and pass.

Simulating with the VCS Simulator

To run the simulation in the VCS simulation tool, follow these steps:

1. Change directory to the *<variation_name>* **_example/alt_e40_e100/example_testbench** directory.
2. In the command line, type: `sh run_vcs.sh`

The example testbench will run and pass.

Testbench Output Example: 40GbE IP Core with Adapters

This section shows successful simulation using the 40GbE IP core with adapters testbench (**alt_40gbe_tb.sv**). The testbench connects the Ethernet TX lanes to the Ethernet RX lanes, so that the IP core is in an external loopback configuration. In simulation, the testbench resets the IP core and waits for lane alignment and deskew to complete successfully. The packet generator sends ten packets on the Ethernet TX lanes and the packet checker checks the packets when the IP core receives them on the Ethernet RX lanes.

The successful testbench run displays the following output:

```
#
# *****
# **      40g Ethernet Testbench
# **
# **
# ** Target Device:          Stratix V
# ** IP Configuration:      40 Gbe
# ** Variant Name:          abc
# ** Status Clock Rate:     100000 KHz
# ** Statistics Registers:  Enabled
# **
# ** This variant is MAC & PHY
# ** Interface:              Avalon-ST
# ** *****
# ** Resetting the IP Core...
# **
# **
# ** *****
# ** Waiting for alignment and deskew...
# **
# **
# ** Virtual lane locked:      None (lanes left:      4) |@@@|
# All lanes locked. Starting deskew at time          5528000
# ** Virtual lane locked:      0 (lanes left:         3) |@@\|
# ** Virtual lane locked:      1 (lanes left:         2) |@\|
# ** Virtual lane locked:      2 (lanes left:         1) |@\|
# ** Virtual lane locked:      3 (lanes left:         0) |/\|
# Deskew complete at time          6404800
# ** All virtual lanes locked and deskewed, ready for data |----|
# *****
# ** Starting TX traffic...
# **
# **
# ** Sending Packet          1...
# ** Sending Packet          2...
# ** Sending Packet          3...
# ** Sending Packet          4...
# ** Sending Packet          5...
# ** Sending Packet          6...
# ** Sending Packet          7...
# ** Sending Packet          8...
# ** Sending Packet          9...
# ** Sending Packet         10...
# ** Received Packet         1...
# ** Received Packet         2...
# ** Received Packet         3...
# ** Received Packet         4...
# ** Received Packet         5...
# ** Received Packet         6...
# ** Received Packet         7...
# ** Received Packet         8...
# ** Received Packet         9...
# ** Received Packet        10...
# **
# ** Testbench complete.
# **
# ** *****
# ** Note: $finish      : ./alt_40gbe_tb.v(181)
# ** Time: 7490400 ps  Iteration: 0  Instance: /alt_40gbe_tb
```

Testbench Output Example: 100GbE IP Core with Adapters

This section shows successful simulation using the 100GbE IP core with adapters testbench (**alt_100gbe_tb.sv**). The testbench connects the Ethernet TX lanes to the Ethernet RX lanes, so that the IP core is in an external loopback configuration. In simulation, the testbench resets the IP core and waits for lane alignment and deskew to complete successfully. The packet generator sends ten packets on the Ethernet TX lanes and the packet checker checks the packets when the IP core receives them on the Ethernet RX lanes.

The successful testbench run displays the following output:

```
# *****
# **      100g Ethernet Testbench
# **
# **
# ** Target Device:          Stratix IV
# ** IP Configuration:      100 Gbe
# ** Variant Name:          abc
# ** Status Clock Rate:     50000 KHz
# ** Statistics Registers:  Enabled
# **
# ** This variant is MAC & PHY
# ** Interface:              Avalon-ST
# *****
# ** Resetting the IP Core...
# **
# **
# *****
# ** Waiting for alignment and deskew...
# **
# **
# ** Virtual lane locked:    None (lanes left:      20) | @@@@|
# ** Virtual lane locked:    0 (lanes left:         19) | @@@@|
# ** Virtual lane locked:    3 (lanes left:         18) | @@@@|
# ** Virtual lane locked:    5 (lanes left:         17) | @@@@|
# ** Virtual lane locked:   11 (lanes left:         16) | @@@@|
# ** Virtual lane locked:    7 (lanes left:         15) | @@@@|
# ** Virtual lane locked:   12 (lanes left:         14) | @@@@|
# ** Virtual lane locked:   14 (lanes left:         13) | @@@@|
# ** Virtual lane locked:    4 (lanes left:         12) | @@@@|
# ** Virtual lane locked:    8 (lanes left:         11) | @@@@|
# ** Virtual lane locked:    2 (lanes left:         10) | @@@@|
# ** Virtual lane locked:   18 (lanes left:          9) | @@@@|
# ** Virtual lane locked:    1 (lanes left:          8) | @@@@|
# ** Virtual lane locked:   16 (lanes left:          7) | @@@@|
# ** Virtual lane locked:   17 (lanes left:          6) | @@@@|
# ** Virtual lane locked:   19 (lanes left:          5) | @@@@|
# ** Virtual lane locked:   15 (lanes left:          4) | @@@@|
# ** Virtual lane locked:   13 (lanes left:          3) | @@@@|
# ** Virtual lane locked:    9 (lanes left:          2) | @@@@|
# ** Virtual lane locked:   10 (lanes left:          1) | @@@@|
# All lanes locked. Starting deskew at time 29531200
# ** Virtual lane locked:    6 (lanes left:          0) | @@@@|
# Deskew complete at time 31243200
# ** All virtual lanes locked and deskewed, ready for data
# -----|
# *****
# ** Starting TX traffic...
# **
# **
# ** Sending Packet          1...
# ** Sending Packet          2...
# ** Sending Packet          3...
# ** Sending Packet          4...
```

```

# ** Sending Packet          5...
# ** Sending Packet          6...
# ** Sending Packet          7...
# ** Sending Packet          8...
# ** Sending Packet          9...
# ** Sending Packet         10...
# ** Received Packet         1...
# ** Received Packet         2...
# ** Received Packet         3...
# ** Received Packet         4...
# ** Received Packet         5...
# ** Received Packet         6...
# ** Received Packet         7...
# ** Received Packet         8...
# ** Received Packet         9...
# ** Received Packet        10...
# **
# ** Testbench complete.
# **
# *****
# ** Note: $finish      : ./alt_100gbe_tb.v(197)
#   Time: 32367200 ps  Iteration: 0  Instance: /alt_100gbe_tb

```

Compiling the Full Design and Programming the FPGA

You can use the **Start Compilation** command on the Processing menu in the Quartus II software to compile your design. After successfully compiling your design, program the targeted Altera device with the Programmer and verify the design in hardware.

Related Information

- [Quartus II Incremental Compilation for Hierarchical and Team-Based Design](#)
Information about compiling your design. Chapter in volume 1 of the *Quartus II Handbook*.
- [Quartus II Programmer](#)
Information about programming the device. Chapter in volume 3 of the *Quartus II Handbook*.

Initializing the IP Core

The testbench initializes the IP core. However, when you simulate or run your own design in hardware, you must implement the initialization steps yourself.

To initialize the 40-100GbE IP core in your own design, follow these steps:

1. Drive the clock ports.
2. Reset the IP core.
3. Clear the statistics counters by writing the value of 1 to bit 3 of the general control `MAC_CMD_config` register.

Related Information

- [Clocks](#) on page 3-51
In step 1, drive the clock ports as specified here.
- [Resets](#) on page 3-54
In step 2, reset the IP core according to these recommendations.

- **Control and Status Interface** on page 3-51
In step 3, write to the `MAC_CMD_config` register using this interface.
- **MAC Configuration and Filter Registers** on page 3-99
Information about the `MAC_CMD_config` register.

2014.12.15

UG-01088



Subscribe



Send Feedback

This chapter provides a detailed description of the 40-100GbE IP core. The chapter begins with a high-level overview of typical Ethernet systems and then provides detailed descriptions of the MAC, transmit (TX) and receive (RX) datapaths, signals, register descriptions, and an Ethernet glossary. This chapter includes the following sections:

High Level System Overview on page 3-2

40-100GbE MAC and PHY Functional Description on page 3-2

Signals on page 3-55

Software Interface: Registers on page 3-76

Ethernet Glossary on page 3-119

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.


ISO
9001:2008
Registered



High Level System Overview

Figure 3-1: 40GbE and 100GbE MAC and PHY MegaCore Function

Main block, internal connections, and external block requirements.



40-100GbE MAC and PHY Functional Description

The Altera 40-100GbE IP core implements the 40-100GbE Ethernet MAC in accordance with the *IEEE 802.3ba 2010 40G and 100G Ethernet Standard*. This IP core handles the frame encapsulation and flow of data between a client logic and Ethernet network via a 40-100GbE Ethernet PCS and PMA (PHY).

In the transmit direction, the MAC accepts client frames, and inserts inter-packet gap (IPG), preamble, start of frame delimiter (SFD), header, padding, and checksum bits before passing them to the PHY. The PHY encodes the MAC frame as required for reliable transmission over the media to the remote end.

In the receive direction, the PHY passes frames to the MAC. The MAC accepts frames from the PHY, performs checks, updates statistics counters, strips out the CRC, preamble, and SFD, and passes the rest of the frame to the client. In RX preamble pass-through mode, the MAC passes on the preamble and SFD to the client instead of stripping them out.

The MAC includes the following interfaces:

- Datapath client-interface—The following options are available:
 - 40GbE with adapters—Avalon-ST, 256 bits
 - 40GbE—Custom streaming, 128 bits
 - 100GbE with adapters—Avalon-ST, 512 bits
 - 100GbE—Custom streaming, 320 bits
- Datapath PHY side—The following options are available:
 - 40GbE—XLAUI
 - 100GbE—CAUI, CAUI-4
- Management interface—Avalon-MM host slave interface for MAC management. This interface has a data width of 32 bits and an address width of 16 bits.

The PHY includes the following interfaces:

- Datapath MAC—The following options are available:
 - 40GbE—XLAUI
 - 100GbE—CAUI, CAUI-4
- Datapath Ethernet interface—The following options are available:
 - 40GbE—Four 10.3125 Gbps serial links
 - 40GbE 24.24—Four 6.25 Gbps serial links
 - 100GbE—Ten 10.3125 Gbps serial links
 - 100GbE CAUI-4—Four 25.78125 Gbps serial links


40-100GbE IP Core TX Datapath

The TX MAC module receives the client payload data with the destination and source addresses and then adds, appends, or updates various header fields in accordance with the configuration specified. The MAC does not modify the destination address or the payload received from client. However, the TX MAC module adds a preamble (if the IP core is not programmed to receive the preamble from user logic), pads the payload to satisfy the minimum Ethernet frame payload of 46 bytes, and calculates the CRC over the entire MAC frame. (If padding is added, it is also included in CRC calculation). The TX MAC module can also modify the source address, and always inserts IDLE bytes to maintain an average IPG.

Figure 3-2: Typical Client Frame at the Transmit Interface

Illustrates the changes that the TX MAC makes to the client frame. This figure uses the following notational conventions:

- $\langle p \rangle$ = payload size = 0–1500 bytes, or 9600 bytes for jumbo frames.
- $\langle s \rangle$ = padding bytes = 0–46 bytes.
- $\langle l \rangle$ = number of IPG bytes



The following sections describe the functions that the TX module performs:

[Preamble, Start, and SFD Insertion](#) on page 3-4

[Address Insertion](#) on page 3-4

[Length/Type Field Processing](#) on page 3-5

[Frame Padding](#) on page 3-5

[Frame Check Sequence \(CRC-32\) Insertion](#) on page 3-5

[Inter-Packet Gap Generation and Insertion](#) on page 3-5

Preamble, Start, and SFD Insertion

In the TX datapath the MAC appends a one-byte Start, 6-byte preamble, and 1-byte SFD to the client frame. This MAC module also incorporates the functions of the reconciliation sublayer. The source of the 6-byte preamble and 1-byte SFD depends on whether you turn on the TX preamble pass-through feature by setting bit 1 of the `Preamble Pass-Through Configuration` register at offset 0x125.

If the TX preamble pass-through feature is turned on, the client provides the eight-byte preamble (Start byte, 6-byte preamble, and 1-byte SFD) on the data bus. However, the IP core overwrites the Start byte the client provides, and only passes on the 6-byte preamble and the SFD. The client is responsible for providing an appropriate SFD byte.

Related Information

[MAC Feature Configuration Registers](#) on page 3-105

Includes information about the `Preamble Pass-Through Configuration` register.

Address Insertion

The client provides the destination MAC address and the source address of the local MAC. However, if enabled by bit [31] of the `MADDR_CTRL` register at offset 0xC2, the source MAC address can be replaced by the source address contained in two, 32-bit MAC registers: `SRC_AD_LO` and `SRC_AD_HI`.

Related Information

[MAC Address Registers](#) on page 3-107

Includes information about the `MADDR_CTRL`, `SRC_AD_LO`, and `SRC_AD_HI` registers.

Length/Type Field Processing

This two-byte header represents either the length of the payload or the type of MAC frame. When the value of this field is equal to or greater than 1536 (0x600) it indicates a type field. Otherwise, this field provides the length of the payload data that ranges from 0–1500 bytes. The TX MAC does not modify this field before forwarding it to the network.

Frame Padding

When the length of client frame is less than 64 bytes (meaning the payload is less than 46 bytes), the TX MAC module inserts pad bytes (0x00) after the payload to create a frame length equal to the minimum size of 64 bytes.

Frame Check Sequence (CRC-32) Insertion

The TX MAC computes and inserts a CRC32 checksum in the transmitted MAC frame. The frame check sequence (FCS) field contains a 32-bit CRC value. The MAC computes the CRC32 over the frame bytes that include the source address, destination address, length, data, and pad. The CRC checksum computation excludes the preamble, SFD, and FCS. The encoding is defined by the following generating polynomial:

$$FCS(X) = X^{32} + X^{26} + X^{23} + X^{22} + X^{16} + X^{12} + X^{11} + X^{10} + X^8 + X^7 + X^5 + X^4 + X^2 + X + 1$$

CRC bits are transmitted with MSB (X32) first.

Independent user configuration register bits control FCS CRC insertion at runtime. Bit [0] of the `CRC_CONFIG` register enables and disables CRC insertion. By default, the CRC insertion feature is enabled.

Related Information

- [Order of Transmission](#) on page 3-16
Illustrations of the byte order and octet transmission order on the Avalon-ST client interface.
- [CRC Configuration Register](#) on page 3-105
Information about the `CRC_CONFIG` register.

Inter-Packet Gap Generation and Insertion

The TX MAC maintains the minimum inter-packet gap (IPG) between transmitted frames required by the IEEE 802.3 Ethernet standard. The standard requires an average minimum IPG of 96 bit times (or 12 byte times). The deficit idle counter maintains the average IPG of 12 bytes.

The MAC adjusts the IPG to compensate for Alignment Marker insertion by the PHY. You can program this adjustment using the `IPG_DEL_PERIOD` and `IPG_DEL_ENABLE` registers at offsets 0x126 and 0x127, respectively. By default, the adjustment removes one Idle byte for every 16384 bytes. This removal rate corresponds to the bandwidth used by the Alignment Marker that the PHY inserts in the outgoing Ethernet communication. You can modify the value in the `IPG_DEL_PERIOD` register to specify more or less frequent removal of Idle bytes from the sequence.

Related Information

[MAC Feature Configuration Registers](#) on page 3-105

Includes information about the `IPG_DEL_PERIOD` and `IPG_DEL_ENABLE` registers.

40-100GbE IP Core TX Data Bus Interfaces

This section describes the TX data bus at the user interface and includes the following topics:

[40-100GbE IP Core User Interface Data Bus](#) on page 3-6

[40-100GbE IP Core TX Data Bus with Adapters \(Avalon-ST Interface\)](#) on page 3-6

[40-100GbE IP Core TX Data Bus Without Adapters \(Custom Streaming Interface\)](#) on page 3-9

[Bus Quantization Effects With Adapters](#) on page 3-10

[User Interface to Ethernet Transmission](#) on page 3-11

40-100GbE IP Core User Interface Data Bus**Table 3-1: User Interface Width Depends on IP Core Variation**

The 40-100GbE IP core provides two different client interfaces: the Avalon-ST interface and a custom interface. The Avalon-ST interface requires adapters and the custom streaming interface does not require adapters.

| Client Interface | Data Bus Width (Bits) | |
|--|-----------------------|----------------|
| | 40GbE IP Core | 100GbE IP Core |
| Custom streaming interface (no adapters) | 128 | 320 |
| Avalon-ST interface (with adapters) | 256 | 512 |

40-100GbE IP Core TX Data Bus with Adapters (Avalon-ST Interface)

The 40-100GbE IP core TX datapath with adapters employs the Avalon-ST protocol. The Avalon-ST protocol is a synchronous point-to-point, unidirectional interface that connects the producer of a data stream (source) to a consumer of data (sink). The key properties of this interface include:

- Start of packet (SOP) and end of packet (EOP) signals delimit frame transfers.
- A valid signal qualifies signals from source to sink.
- The sink applies backpressure to the source by using the ready signal. The source typically responds to the deassertion of the ready signal from the sink by driving the same data until the sink can accept it. The `readyLatency` defines the relationship between assertion and deassertion of the ready signal, and cycles which are considered to be `ready` for data transfer. The `readyLatency` on the TX client interface is zero cycles.

Altera provides an Avalon-ST interface with adapters for both the 40GbE and 100GbE IP cores. The Avalon-ST interface requires that the start of packet (SOP) always be in the MSB, simplifying the interpretation and processing of incoming data. The TX adapter for the 100GbE IP core increases the client interface Avalon-ST bus width from 5 words (320 bits) to 8 words (512 bits). The TX adapter for the 40GbE IP core increases the client interface Avalon-ST bus width from 2 words (128 bits) to 4 words (256

bits). In both cases the client interfaces operate at a frequency above 315 MHz in the standard IP core variations, and at or above the frequency of 190.90 MHz in 24.24 Gbps variations.

The client acts as a source and the TX MAC acts as a sink in the transmit direction.

Figure 3-3: TX Client to MAC Interface with Adapters (Avalon-ST)

The Avalon-ST interface bus width varies with the IP core variation. In the figure, $\langle n \rangle = 4$ for the 40GbE IP core and $\langle n \rangle = 8$ for the 100GbE IP core. $\langle l \rangle$ is $\log_2(8 * \langle n \rangle)$.




Table 3-2: Signals of the TX Client Interface with Adapters

In the table, $\langle n \rangle = 4$ for the 40GbE IP core and $\langle n \rangle = 8$ for the 100GbE IP core. $\langle l \rangle$ is $\log_2(8 * \langle n \rangle)$. All interface signals are clocked by the `clk_txmac` clock.

| Signal Name | Direction | Description |
|--|-----------|---|
| $l\langle n \rangle_tx_data[\langle n \rangle * 64 - 1 : 0]$ | Input | TX data. If the preamble pass-through feature is enabled, data begins with the preamble. |
| $l\langle n \rangle_tx_empty[\langle l \rangle - 1 : 0]$ | Input | Indicates the number of empty bytes on $l\langle n \rangle_tx_data$ when $l\langle n \rangle_tx_endofpacket$ is asserted. |
| $l\langle n \rangle_tx_startofpacket$ | Input | When asserted, indicates the start of a packet. The packet starts on the MSB. |
| $l\langle n \rangle_tx_endofpacket$ | Input | When asserted, indicates the end of packet. |

| Signal Name | Direction | Description |
|----------------------------------|-----------|---|
| <code>l<n>_tx_ready</code> | Output | <p>When asserted, the MAC is ready to receive data. The <code>l<n>_tx_ready</code> signal acts as an acknowledge. The source drives <code>l<n>_tx_valid</code> and <code>l<n>_tx_data[l<n>*64-1:0]</code>, then waits for the sink to assert <code>l<n>_tx_ready</code>. The <code>readyLatency</code> is zero cycles, so that the IP core accepts valid data in the same cycle in which it asserts <code>l<n>_tx_ready</code>.</p> <p>The <code>tx_ready</code> signal indicates the MAC is ready to receive data in normal operational model. However, the <code>tx_ready</code> signal might not be an adequate indication following reset. To avoid sending packets before the Ethernet link is able to transmit them reliably, you should ensure that the application does not send packets on the TX client interface until after the <code>lanes_deskewed</code> signal is asserted.</p> |
| <code>l<n>_tx_valid</code> | Input | <p>When asserted <code>l<n>_tx_data</code> is valid. This signal must be continuously asserted between the assertions of <code>l<n>_tx_startofpacket</code> and <code>l<n>_tx_endofpacket</code> for the same packet.</p> |

Figure 3-4: Traffic on the TX and RX Client Interface for 40GbE IP Core Using the Four- to Two-Word Adapters

Shows typical traffic for the TX and RX Avalon-ST interface 40GbE IP core. This example shows a part of a ModelSim simulation of the parallel testbench provided with the IP core.

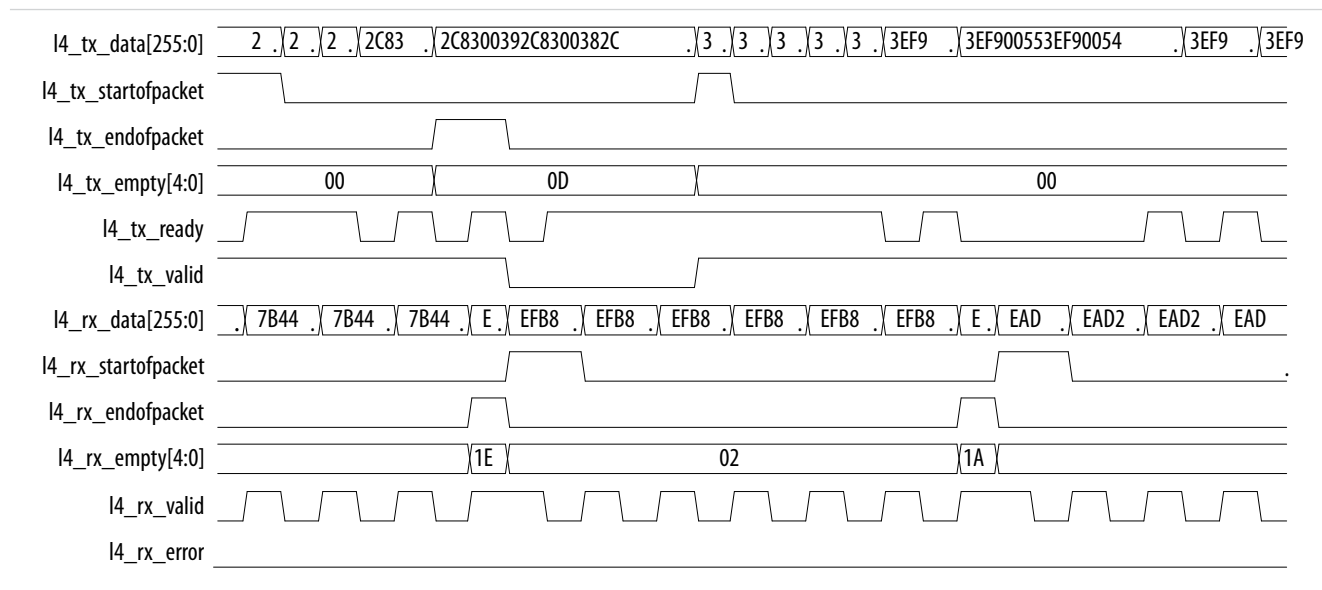



Figure 3-5: Traffic on the TX and RX Client Interface for 100GbE IP Core Using the Eight- to Five-Word Adapters

Shows typical traffic for the TX and RX Avalon-ST interface of the 100GbE IP core. This example shows a part of a ModelSim simulation of the parallel testbench provided with the IP core.



Related Information

- [Testbenches with Adapters](#) on page 2-15
Describes the 40GbE example in detail.
- [Avalon Interface Specifications](#)
For more information about the Avalon-ST interface.

40-100GbE IP Core TX Data Bus Without Adapters (Custom Streaming Interface)

When no adapters are used, the 40GbE custom interface bus width is 2 words (128 bits) and the 100GbE custom interface bus width is 5 words (320 bits). In both cases the client interfaces operate at a frequency above 315 MHz.

Figure 3-6: TX Client to MAC Interface Without Adapters

The custom streaming interface bus width varies with the IP core variation. In the figure, $\langle w \rangle = 2$ for the 40GbE IP core and $\langle w \rangle = 5$ for the 100GbE IP core.




Table 3-3: Signals of the TX Client Interface Without Adapters

In the table, $\langle w \rangle = 2$ for the 40GbE IP core and $\langle w \rangle = 5$ for the 100GbE IP core.

| Signal Name | Direction | Description |
|--|-----------|---|
| <code>din[$\langle w \rangle * 64 - 1 : 0$]</code> | Input | Data bytes to send in big-Endian mode. |
| <code>din_start[$\langle w \rangle - 1 : 0$]</code> | Input | Start of packet (SOP) location in the TX data bus. Only the most significant byte of each 64-bit word may be a start of packet. Bit 63 or 127 are possible for the 40GbE and bits 319, 255, 191, 127, or 63 are possible for 100 GbE. |
| <code>din_end_pos[$\langle w \rangle * 8 - 1 : 0$]</code> | Input | End of packet. Any byte may be the last byte in a packet. |
| <code>din_ack</code> | Output | Indicates that input data was accepted by the IP core. |
| <code>clk_txmac</code> | Input | TX MAC clock. The minimum clock frequency is 315 MHz for the circuit to function correctly. The <code>clk_txmac</code> and <code>clk_rxmac</code> which clocks the RX datapath are not related and their rates do not have to match. |

The IP core reads the bytes in big endian order. A packet may start in the most significant byte of any word. A packet may end on any byte.

To avoid sending packets before the IP core completes the reset sequence, you should ensure that the application does not send packets on the TX client interface until after the `lanes_deskewed` signal is asserted.

Related Information

- [40GbE IP Core Without Adapters](#) on page 3-12
Illustrates the 40GbE IP core TX client interface without adapters.
- [100GbE IP Core Without Adapters](#) on page 3-14
Illustrates the 100GbE IP core TX client interface without adapters.


Bus Quantization Effects With Adapters

The TX custom streaming interface allows a packet to start at any of two or five positions to maximize utilization of the link bandwidth. The TX Avalon-ST interface only allows start of packet (SOP) to be placed at the most significant position. If the SOP were restricted to the most significant position in the client logic data bus in the custom streaming interface, bus bandwidth would be reduced.


Figure 3-7: Reduced Bandwidth With Left-Aligned SOP Requirement

Illustrates the reduction of bandwidth that would be caused by left-aligning the SOP for the 100GbE IP core.

Example A



Example B



Example A shows the minimum-sized packet of eight words. Example B shows an 11-word packet which is the worst-case for bandwidth utilization. Assuming another packet is waiting for transmission, the effective ingress bandwidth is reduced by 20% and 26%, respectively. Running the MAC portion of the logic slightly faster than is required can mitigate this loss of bandwidth. Additional increases in the MAC frequency can provide further mitigation, although doing so makes timing closure more difficult. The wider data bus for the Avalon-ST interface also helps to compensate for the Avalon-ST left-aligned SOP requirement.

User Interface to Ethernet Transmission

The IP core reverses the bit stream for transmission per Ethernet requirements. The transmitter handles the insertion of the inter-packet gap, frame delimiters, and padding with zeros as necessary. The transmitter also handles FCS computation and insertion.

The Ethernet MAC and PHY transmit complete packets. After transmission begins, it must complete with no IDLE insertions. Between the end of one packet and the beginning of the next packet, the data input is not considered and the transmitter sends IDLE characters. An unbounded number of IDLE characters can be sent between packets.

Related Information

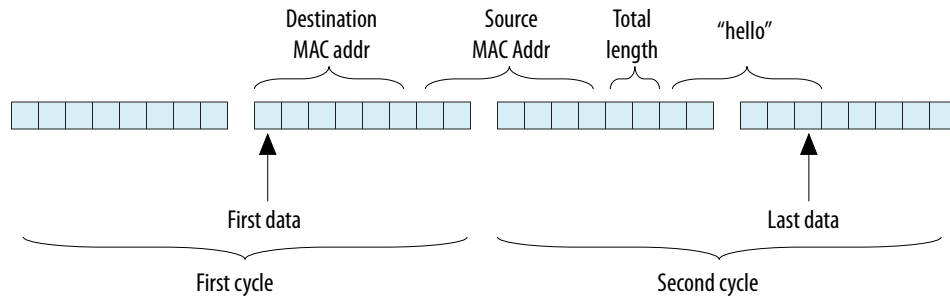
- [40GbE IP Core Without Adapters](#) on page 3-12
Illustrates the mapping from data on the 40GbE IP core TX client interface without adapters to the Ethernet packet fields.
- [100GbE IP Core Without Adapters](#) on page 3-14
Illustrates the mapping from data on the 100GbE IP core TX client interface without adapters to the Ethernet packet fields.

40GbE IP Core Without Adapters

The following figures illustrate the transmission of a short packet when preamble pass-through is turned off and when it is turned on.

Figure 3-8: Short Packet Example Without Preamble

Illustrates the transmission of a short packet when preamble pass-through is turned off.



Example 3-1: Bus Representation of a Short TX Packet Without Preamble

This example shows the Verilog HDL code that represents the simple packet illustrated in the previous figure. Note that bit `din_end[5]` in the second cycle, corresponding to the “Last data” in the figure, is asserted.

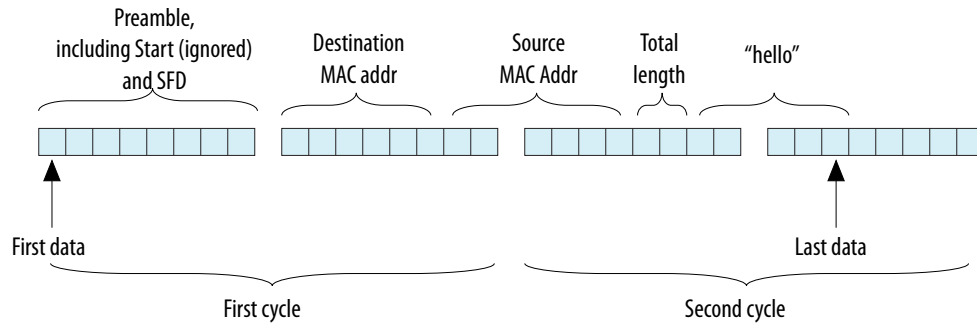
```
wire [6*8-1:0] dst_addr = 48'hffff ffff ffff;
wire [6*8-1:0] src_addr = 48'h0007 edff 1234;
wire [2*8-1:0] len = 16'd64;

\\First cycle:
din = {64'h0, dst_addr, src_addr[47:32]};
din_start = 2'b01;
din_end = 16'b00000000_00000000;

\\Second cycle:
din = {src_addr[31:0], len, "hello", 40'h0};
din_start = 2'b00;
din_end = 16'b00000000_00100000;
```

Figure 3-9: Short Packet Example With Preamble

Illustrates the transmission of a short packet when preamble pass-through is turned on. In this example, the preamble starts in the MSB; however, this need not be the case.

**Example 3-2: Bus Representation of a Short TX Packet With Preamble**

This example shows the Verilog HDL code that represents the simple packet illustrated in the previous figure. Note that bit `din_end[5]` in the second cycle, corresponding to the “Last data” in the figure, is asserted.


```
wire [8*8-1:0] preamble = 64'hXX55_5555_55D5;
wire [6*8-1:0] dst_addr = 48'hffff ffff ffff;
wire [6*8-1:0] src_addr = 48'h0007 edff 1234;
wire [2*8-1:0] len = 16'd64;

\\First cycle:
din = {preamble, dst_addr, src_addr[47:32]};
din_start = 2'b10;
din_end = 16'b00000000_00000000;

\\Second cycle:
din = {src_addr[31:0], len, "hello", 40'h0};
din_start = 2'b00;
din_end = 16'b00000000_00100000;
```

Figure 3-10: Sample 40GbE IP Core TX Bus Activity

Illustrates the deassertion of the `din_ack` signal. The data beginning with `0xe6e7` is not immediately accepted. The `din` bus must be held until `din_ack` returns to one. At this point normal data flow resumes.




100GbE IP Core Without Adapters

The following figures illustrate the transmission of a short packet when preamble pass-through is turned off and when it is turned on.

Figure 3-11: Short Packet Example Without Preamble

Illustrates the transmission of a short packet for the 100GbE IP core when preamble pass-through is turned off.



Example 3-3: Bus Representation of a Short TX Packet Without Preamble


This example shows the Verilog HDL code that represents the simple packet illustrated in the preceding figure. Note that bit `din_end[13]` corresponding to the "Last data" in the figure, is asserted.

```
wire [6*8-1:0] dst_addr = 48'hffff ffff ffff;
wire [6*8-1:0] src_addr = 48'h0007 edff 1234;
wire [2*8-1:0] len = 16'd64;

din = {64'h0, dst_addr, src_addr, len, "hello", 40'h0};
din_start = 5'b01000;
din_end = 40'b00000000_00000000_00000000_00100000_00000000;
```

Figure 3-12: Short TX Packet Example With Preamble

Illustrates the transmission of a short packet for the 100GbE IP core when preamble pass-through is turned on.



Example 3-4: Bus Representation of a Short TX Packet With Preamble

This example shows the Verilog HDL code that represents the simple packet illustrated in the preceding figure. Note that bit `din_end[5]` corresponding to the “Last data” in the figure, is asserted.


```

wire [8*8-1:0] preamble = 64'hXX55_5555_55D5;
wire [6*8-1:0] dst_addr = 48'hffff ffff ffff;
wire [6*8-1:0] src_addr = 48'h0007 edff 1234;
wire [2*8-1:0] len = 16'd64;

din = {64'h0, preamble, dst_addr, src_addr, len, "hello", 40'h0};
din_start = 5'b01000;
din_end = 40'b00000000_00000000_00000000_00000000_00100000;
    
```

Figure 3-13: Sample 100GbE IP Core TX Bus Activity

Illustrates the deassertion of the `din_ack` signal. The data beginning with `0x0202` is not immediately accepted. The `din` bus must be held until `din_ack` returns to one. At this point normal data flow resumes.



The TX logic supports packets of less than the usual length. However, no more than two start-of-packets can occur in the same clock cycle.

For example, `din_start` might be set to `5'b11000`, indicating the start of a new packet in two successive words. In this case, `din_end_pos` could equal `40'h0101000000`, indicating two packets of eight bytes. Each 8-byte packet is padded with zeros to create a 64-byte packet.

Order of Transmission

The IP core transmits bytes on the Ethernet link starting with the preamble and ending with the FCS in accordance with the IEEE 802.3 standard. Transmit frames the IP core receives on the client interface are big-endian. Frames the MAC sends to the PHY on the XGMII/CGMII between the MAC and the PHY are little-endian; the MAC TX transmits frames on this interface beginning with the least significant byte.

Figure 3-14: Byte Order on the Client Interface Lanes Without Preamble Pass-Through

Describes the byte order on the Avalon-ST interface when the preamble pass-through feature is turned off. Destination Address[40] is the broadcast/multicast bit (a type bit), and Destination Address[41] is a locally administered address bit.

| | Destination Address (DA) | | | | | | Source Address (SA) | | | | | | Type/ Length | | Data (D) | | |
|-------|--------------------------|---------|---------|---------|--------|-------|---------------------|---------|---------|---------|--------|-------|-----------------|-------|----------|-----|----------|
| Octet | 5 | 4 | 3 | 2 | 1 | 0 | 5 | 4 | 3 | 2 | 1 | 0 | 1 | 0 | 00 | ... | NN |
| Bit | [47:40] | [39:32] | [31:24] | [23:16] | [15:8] | [7:0] | [47:40] | [39:32] | [31:24] | [23:16] | [15:8] | [7:0] | [15:8] | [7:0] | MSB[7:0] | .. | LSB[7:0] |

For example, the destination MAC address includes the following six octets AC-DE-48-00-00-80. The first octet transmitted (octet 0 of the MAC address described in the 802.3 standard) is AC and the last octet transmitted (octet 7 of the MAC address) is 80. The first bit transmitted is the low-order bit of AC, a zero. The last bit transmitted is the high order bit of 80, a one.

The preceding table and the following figure show that in this example, 0xAC is driven on DA5 (DA[47:40]) and 0x80 is driven on DA0 (DA[7:0]).

Figure 3-15: Octet Transmission on the Avalon-ST Signals Without Preamble Pass-Through

Illustrates how the octets of the client frame are transferred over the TX datapath when preamble pass-through is turned off.




Figure 3-16: Byte Order on the Avalon-ST Interface Lanes With Preamble Pass-Through

Describes the byte order on the Avalon-ST interface when the preamble pass-through feature is turned on. Recall that the IP core overwrites the Start byte you provide on the client interface, but does not overwrite the SFD byte.

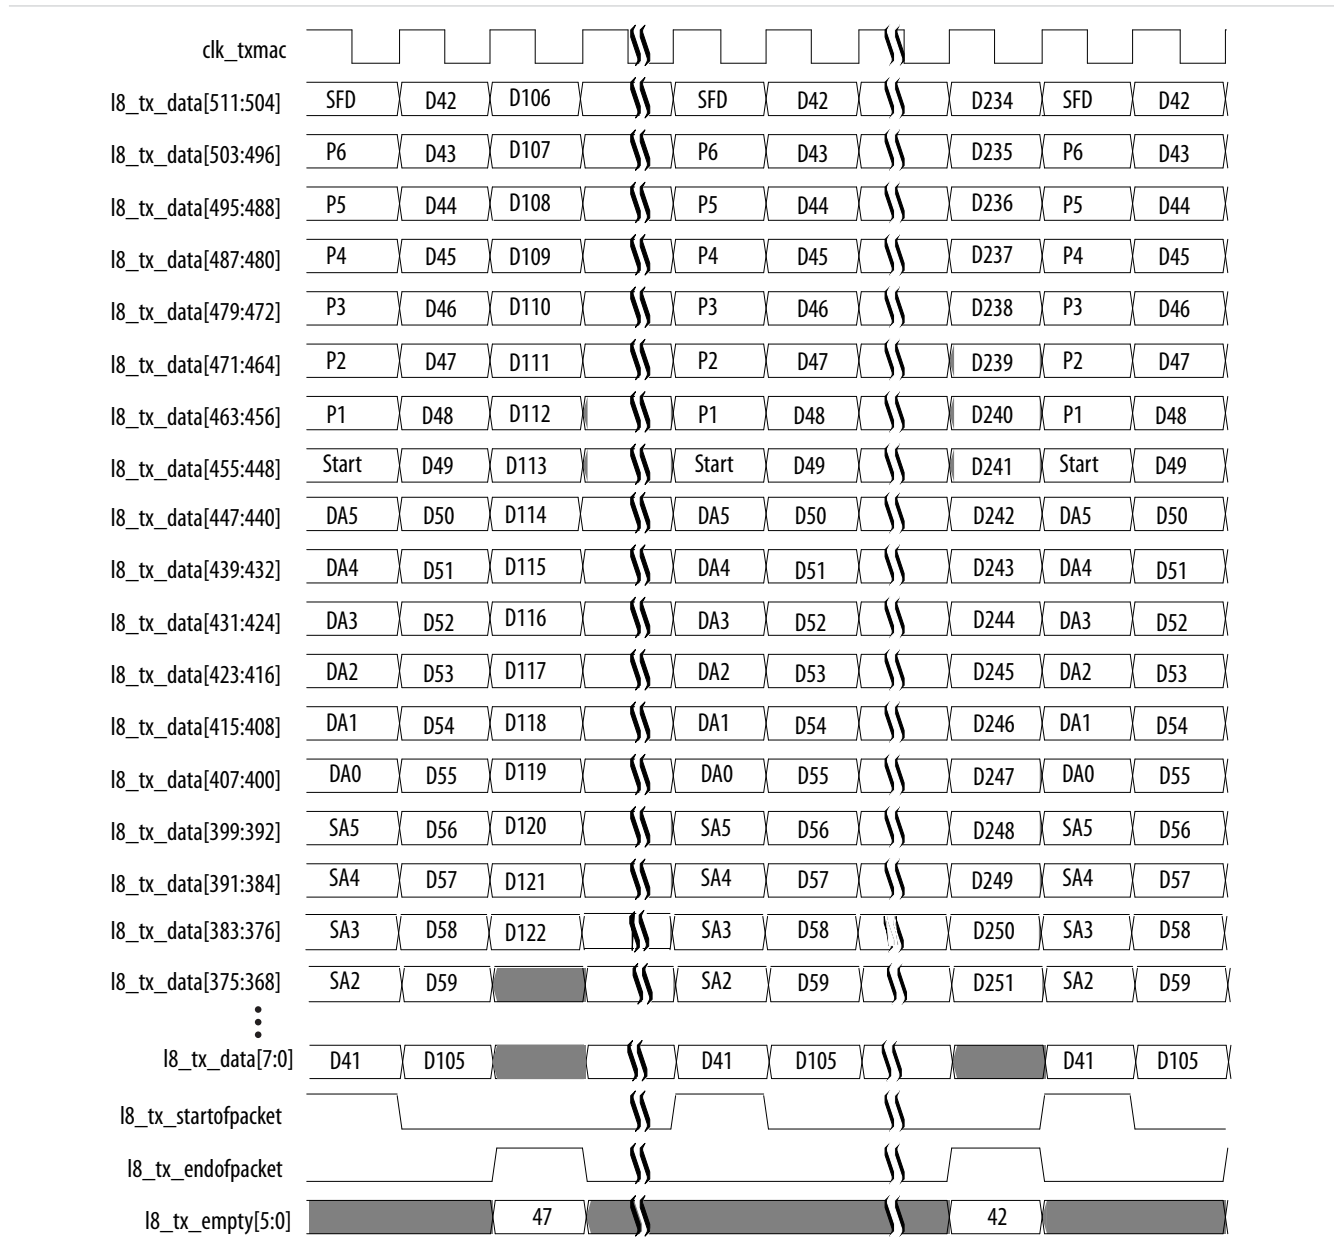
Destination Address[40] is the broadcast/multicast bit (a type bit), and Destination Address[41] is a locally administered address bit.

| | SFD | Preamble | | | | | | Start | Destination Address (DA) | | | | Source Address (SA) | | | | Type/Length | | Data (D) | | | | | | |
|-------|---------|----------|---------|---------|---------|---------|--------|-------|--------------------------|---------|---------|---------|---------------------|-------|---------|---------|-------------|---------|----------|-------|--------|-------|----------|-----|----------|
| Octet | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | 5 | 4 | 3 | 2 | 1 | 0 | 5 | 4 | 3 | 2 | 1 | 0 | 1 | 0 | 00 | ... | NN |
| Bit | [63:56] | [55:48] | [47:40] | [39:32] | [31:24] | [23:16] | [15:8] | [7:0] | [47:40] | [39:32] | [31:24] | [23:16] | [15:8] | [7:0] | [47:40] | [39:32] | [31:24] | [23:16] | [15:8] | [7:0] | [15:8] | [7:0] | MSB[7:0] | .. | LSB[7:0] |

Figure 3-17: Octet Transmission on the Avalon-ST Signals With Preamble Pass-Through

Illustrates how the octets of the client frame are transferred over the TX datapath when preamble pass-through is turned on. The eight preamble bytes precede the destination address bytes. The preamble bytes are reversed: the application must drive the Start byte on `l8_tx_data[455:448]` and the SFD byte on `l8_tx_data[511:504]`.

The destination address and source address bytes follow the preamble pass-through in the same order as in the case without preamble pass-through.



40-100GbE IP Core RX Datapath

The 40-100GbE RX MAC receives Ethernet frames from the PHY and forwards the payload with relevant header bytes to the client after performing some MAC functions on header bytes.

Figure 3-18: Flow of Frame Through the MAC RX Without Preamble Pass-Through

Illustrates the typical flow of frame through the MAC RX when the preamble pass-through feature is turned off. In this figure, $\langle p \rangle$ is payload size (0–1500 bytes), and $\langle s \rangle$ is the number of pad bytes (0–46 bytes).





Figure 3-19: Flow of Frame Through the MAC RX With Preamble Pass-Through Turned On

Illustrates the typical flow of frame through the MAC RX when the preamble pass-through feature is turned on. In this figure, $\langle p \rangle$ is payload size (0–1500 bytes), and $\langle s \rangle$ is the number of pad bytes (0–46 bytes).



The following sections describe the functions performed by the RX MAC:

[40-100GbE IP Core RX Filtering](#) on page 3-21

[40-100GbE IP Core Preamble Processing](#) on page 3-21

[40-100GbE IP Core FCS \(CRC-32\) Removal](#) on page 3-22

[40-100GbE IP Core CRC Checking](#) on page 3-22

[RX CRC Forwarding](#) on page 3-22

[RX Automatic Pad Removal Control](#) on page 3-22

[Address Checking](#) on page 3-24

[Inter-Packet Gap](#) on page 3-24

[Pause Ignore](#) on page 3-24

40-100GbE IP Core RX Filtering

The 40-100GbE IP core can operate in cut-through mode or in store and forward mode. In cut-through mode, the IP core does not buffer incoming Ethernet packets for filtering. It can filter out incoming runt packets, but cannot filter on any other criteria. The value in bit 0 of the `RX_FILTER_CTRL` register at offset 0x103 determines the mode, and the value in bit 3 determines whether the IP core filters runt packets.

When the IP core is in cut-through mode, it does not filter incoming Ethernet packets based on destination address. Therefore, when in cut-through mode, the IP core is in promiscuous receive mode. The Ethernet standard definition of promiscuous receive mode requires that the IP core accept all valid frames, regardless of destination address. The 40-100GbE IP core accepts or rejects invalid frames based on the filtering criteria that are turned on.

In store and forward mode, you can enable oversized-frame handling. When the maximum frame size is set to 9600 bytes, the IP core passes some of the frames between 9601-9644 bytes in size, and drops frames of 9645 bytes or more. For the 100GbE IP core, if the frame size is within 44 bytes over the specified maximum frame size, it may or may not be dropped, but oversized frames of over 44 bytes will always be dropped. For the 40GbE IP core, if the frame size is within 20 bytes over the specified maximum frame size, it may or may not be dropped, but oversized frames of over 20 bytes will always be dropped.

The 40-100GbE IP core supports the following filtering options:

- Destination address mismatch—refer to the descriptions of the `RX_FILTER_CTRL` register and the `MADDR_CTRL` register and the link below to the Address Checking topic.
- Runt frame—refer to the description of the `dout_runt_last_data` signal.
- Oversized frame—refer to the preceding description.
- Pause packet
- Control packet
- FCS error

Related Information

- [Address Checking](#) on page 3-24
Information about address filtering.
- [MAC Address Registers](#) on page 3-107
Information about the `MADDR_CTRL` register and related address-checking registers.
- [MAC Configuration and Filter Registers](#) on page 3-99
Additional information about the filtering options, including information about the `RX_FILTER_CTRL` register.
- [40-100GbE IP Core TX Data Bus Without Adapters \(Custom Streaming Interface\)](#) on page 3-9
Information about the `dout_runt_last_data` signal.
- [Pause Control Frame and Non-Pause Control Frame Filtering and Forwarding](#) on page 3-36
- [40-100GbE IP Core Modes of Operation](#) on page 3-37
Overview of filtering status.

40-100GbE IP Core Preamble Processing

The preamble sequence is Start, six preamble bytes, and SFD. If this sequence is incorrect the frame is ignored. The Start byte must be on receive lane 0 (most significant byte). The IP core uses the SFD byte (0xD5) to identify the last byte of the preamble. The MAC RX looks for the Start, six preamble bytes and SFD.

By default, the MAC RX removes all Start, SFD, preamble, and IPG bytes from accepted frames. However, if you turn on the RX preamble pass-through feature, by setting bit 0 of the `Preamble Pass-Through Configuration` register at offset 0x125, the MAC RX does not remove the eight-byte preamble sequence.

Related Information

[MAC Feature Configuration Registers](#) on page 3-105

Information about the `Preamble Pass-Through Configuration` register.

40-100GbE IP Core FCS (CRC-32) Removal

Independent user configuration register bits control FCS CRC removal at runtime. CRC removal supports both narrow and wide bus options. Bit 1 of the `CRC_CONFIG` register enables and disables CRC removal; by default, CRC removal is enabled.

In the user interface, the EOP signal (`l<n>_rx_endofpacket` or `dout_last_data`) indicates the end of CRC data if CRC is not removed. When CRC is removed, the EOP signal indicates the final byte of payload.

By default, the IP core asserts the FCS error signal (`l<n>_rx_fcs_error` or `dout_fcs_error`) and the EOP signal on the same clock cycle if the current frame has an FCS error. However, if the IP core is in RX automatic pad removal mode, the signals might not be asserted in the same clock cycle.

Related Information

[RX Automatic Pad Removal Control](#) on page 3-22

40-100GbE IP Core CRC Checking

The 32-bit CRC field is received in the order: X32, X30, . . . X1, and X0, where X32 is the most significant bit of the FCS field and occupies the least significant bit position in the first FCS byte.

If a CRC32 error is detected, the RX MAC marks the frame invalid by asserting the `dout_fcs_error` and `dout_fcs_valid` signals.

When operating in the cut-through or store and forward mode, with Avalon-ST or the custom streaming client interface, the FCS result is always preserved.

RX CRC Forwarding

The CRC-32 field is forwarded to the client interface after the final byte of data, if the CRC removal option is not enabled.

Related Information

[40-100GbE IP Core FCS \(CRC-32\) Removal](#) on page 3-22

RX Automatic Pad Removal Control

In the 40GbE and 100GbE MAC configurations, you can enable and disable RX automatic pad removal with a configuration register bit in run-time.

The following figures illustrate the normal format of received data at the MAC RX interface.

Figure 3-20: Flow of Frame Through the MAC RX Without Preamble Pass-Through

Illustrates the typical flow of frame through the MAC RX when the preamble pass-through feature is turned off. In this figure, $\langle p \rangle$ is payload size (0–1500 bytes), and $\langle s \rangle$ is the number of pad bytes (0–46 bytes).





Figure 3-21: Flow of Frame Through the MAC RX With Preamble Pass-Through Turned On

Illustrates the typical flow of frame through the MAC RX when the preamble pass-through feature is turned on. In this figure, $\langle p \rangle$ is payload size (0–1500 bytes), and $\langle s \rangle$ is the number of pad bytes (0–46 bytes).



In these figures, normal packet data, highlighted in gray, lasts until the end of the payload section. However, the IP core might pass along additional padding, marked with PAD, to ensure that the frame length is at least 64 bytes. The Ethernet standard requires padding insertion when the payload length is less than 46 bytes. EOP at the RX interface is normally marked after padding, but you can disable CRC removal to place the EOP at the end of the CRC block.

When enabled, RX automatic pad removal moves the EOP marker to the end of the payload as indicated in the length field, whether padding has been inserted or not. If the length is greater than 46 bytes, no padding has been inserted and this feature has no effect. When you enable RX automatic pad removal, the CRC is excluded from the EOP marker on packets that have stripped padding, and enabling CRC retention has no effect on padded packets.

Note: Signals ending in `*_fcs_valid` and `*_fcs_error` are not shifted along with the new EOP marker. Instead, they function as if pad removal were disabled. Do not rely on these signals when operating in RX automatic pad removal mode.

The `PAD_CONFIG` register controls RX automatic pad removal. By default, pad removal is disabled. Statistics counting is not affected by RX automatic pad removal; data reports as default, as if the padding were not removed.

Related Information

[MAC Feature Configuration Registers](#) on page 3-105

Information about the `PAD_CONFIG` register.

Address Checking

The RX MAC supports all three types of addresses:

- **Unicast**—Specifies a destination address is a unicast (individual) address. Bit 0 is 0.
- **Multicast**—Specifies a destination address is a multicast or group address. Bit 0 is 1.
- **Broadcast**—Specifies a broadcast address when all 48 bits in the destination address are all 1s, 48'hFFFF_FFFF_FFFF.

If destination address matching is enabled, IP core address checking compares the address to the address programmed in the destination address register, and accepts only the frames with a matching address. You must enable filtering to discard mismatched destination addresses.

To enable address checking, you must turn ensure your 40-100GbE IP core has the following values in the specified register fields:

- Bit 0 of the `RX_FILTER_CTRL` register at offset 0x103 has the value of 0.
- Bit 0 of the `MADDR_CTRL` register at offset 0x140 has the value of 1.
- Bit 30 of the `MADDR_CTRL` register has the value of 1.

The `MADDR_CTRL` fields allow you to turn off destination address checking but still enable the IP core to filter RX traffic based on other criteria.

If bit 0 of the `RX_FILTER_CTRL` register has the value of 1, the IP core is in promiscuous receive mode. In this mode, the IP core omits address checking and accepts all the Ethernet frames it receives, except possibly runt frames.

Related Information

- [MAC Address Registers](#) on page 3-107
Information about the `MADDR_CTRL` register and related address-checking registers. Describes the interactions between related register fields.
- [MAC Configuration and Filter Registers](#) on page 3-99
Additional information about the filtering options, including information about the `RX_FILTER_CTRL` register.

Inter-Packet Gap

The MAX RX removes all IPG octets received, and does not forward them to the client interface.

Pause Ignore

When the pause frame receive enable bits are not set, the IP core does not process incoming pause frames. In this case, the MAC TX traffic is not affected by the valid pause frames.

You can enable unicast or multicast pause receive by setting the appropriate bits of the pause registers.

Related Information

- [Congestion and Flow Control Using Pause Frames](#) on page 3-33
- [Pause Control and Generation Interface](#) on page 3-35

- [Pause Registers](#) on page 3-102

40-100GbE IP Core RX Data Bus Interfaces

This section describes the RX data bus at the user interface and includes the following topics:

[40-100GbE IP Core User Interface Data Bus](#) on page 3-6

[40-100GbE IP Core RX Data Bus with Adapters \(Avalon-ST Interface\)](#) on page 3-25

[40-100GbE IP Core RX Data Bus Without Adapters \(Custom Streaming Interface\)](#) on page 3-28

[100GbE IP Core RX Client Interface Examples](#) on page 3-30

[Error Conditions on the RX Datapath](#) on page 3-32

40-100GbE IP Core User Interface Data Bus

Table 3-4: User Interface Width Depends on IP Core Variation

The 40-100GbE IP core provides two different client interfaces: the Avalon-ST interface and a custom interface. The Avalon-ST interface requires adapters and the custom streaming interface does not require adapters.

| Client Interface | Data Bus Width (Bits) | |
|--|-----------------------|----------------|
| | 40GbE IP Core | 100GbE IP Core |
| Custom streaming interface (no adapters) | 128 | 320 |
| Avalon-ST interface (with adapters) | 256 | 512 |


40-100GbE IP Core RX Data Bus with Adapters (Avalon-ST Interface)

The adapter for the RX interface of the 100GbE IP core increases the bus width from 5 words (320 bits) to 8 words (512 bits). The adapter for the RX interface of the 40GbE IP core increases the bus width from 2 word (128 bits) to 4 words (256 bits). The Avalon-ST interface always locates the SOP at the MSB, simplifying the interpretation of incoming data.

The RX MAC acts as a source and the client acts as a sink in the receive direction.

Figure 3-22: RX MAC to Client Interface with Adapters

The Avalon-ST interface bus width varies with the IP core variation. In the figure, $\langle n \rangle = 4$ for the 40GbE IP core and $\langle n \rangle = 8$ for the 100GbE IP core. $\langle l \rangle$ is $\log_2(8 * \langle n \rangle)$.

**Table 3-5: Signals of the RX Client Interface with Adapters**

In the table, $\langle n \rangle = 4$ for the 40GbE IP core and $\langle n \rangle = 8$ for the 100GbE IP core. $\langle l \rangle$ is $\log_2(8 * \langle n \rangle)$.

| Name | Direction | Description |
|--|-----------|---|
| $l\langle n \rangle_rx_data[\langle n \rangle * 64 - 1 : 0]$ | Output | RX data. |
| $l\langle n \rangle_rx_empty[\langle l \rangle - 1 : 0]$ | Output | Indicates the number of empty bytes on $l\langle n \rangle_rx_data$ when $l\langle n \rangle_rx_endofpacket$ is asserted, starting from the least significant byte (LSB). |
| $l\langle n \rangle_rx_startofpacket$ | Output | When asserted, indicates the start of a packet. The packet starts on the MSB. |
| $l\langle n \rangle_rx_endofpacket$ | Output | When asserted, indicates the end of packet. |
| $l\langle n \rangle_rx_error$ | Output | When asserted indicates an error condition. |
| $l\langle n \rangle_rx_valid$ | Output | When asserted, indicates that RX data is valid. Only valid between the $l\langle n \rangle_rx_startofpacket$ and $l\langle n \rangle_rx_endofpacket$ signals. |
| $l\langle n \rangle_rx_fcs_valid$ | Output | When asserted, indicates that FCS is valid. |
| $l\langle n \rangle_rx_fcs_error$ | Output | When asserted, indicates an FCS error condition. |

Figure 3-23: Traffic on the TX and RX Client Interface for 40GbE IP Core Using the Four- to Two-Word Adapters

Shows typical traffic for the TX and RX Avalon-ST interface 40GbE IP core. This example shows a part of a ModelSim simulation of the parallel testbench provided with the IP core.

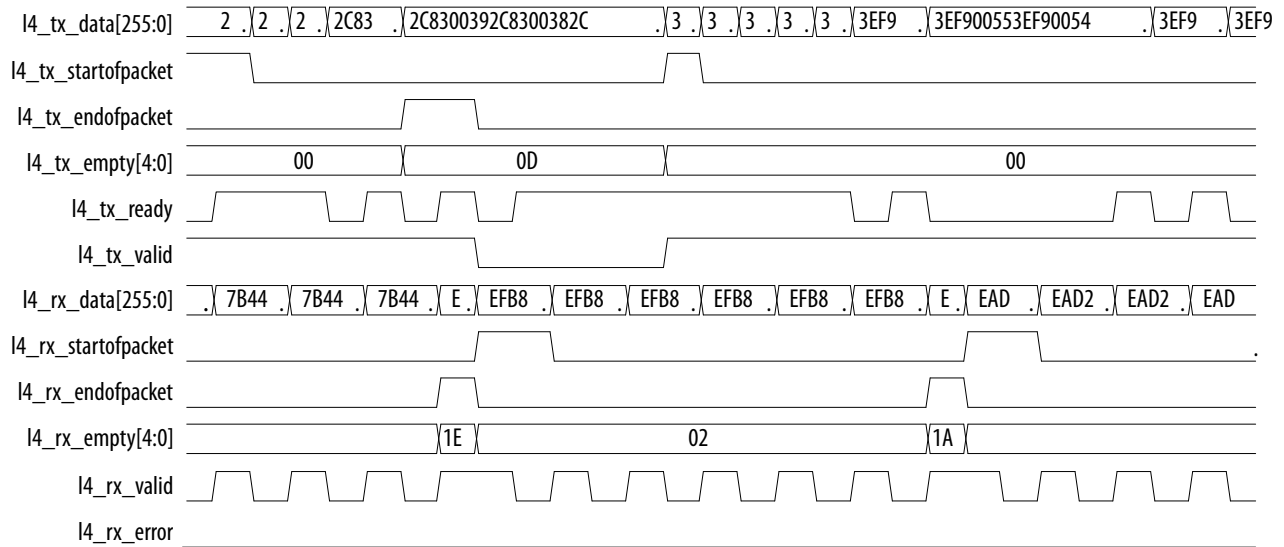



Figure 3-24: Traffic on the TX and RX Client Interface for 100GbE IP Core Using the Eight- to Five-Word Adapters

Shows typical traffic for the TX and RX Avalon-ST interface of the 100GbE IP core. This example shows a part of a ModelSim simulation of the parallel testbench provided with the IP core.



Related Information

[Avalon Interface Specifications](#)

For more information about the Avalon-ST interface.

40-100GbE IP Core RX Data Bus Without Adapters (Custom Streaming Interface)

The RX bus without adapters consists of five 8-byte words, or 320 bits, operating at a frequency above 315 MHz for the 100GbE IP core or two 8-byte words, or 128 bits, for the 40GbE IP core, nominally at 315 MHz. This bus drives data from the RX MAC to the RX client.

Figure 3-25: RX MAC to Client Interface Without Adapters

The custom streaming interface bus width varies with the IP core variation. In the figure, $\langle w \rangle = 2$ for the 40GbE IP core and $\langle w \rangle = 5$ for the 100GbE IP core.

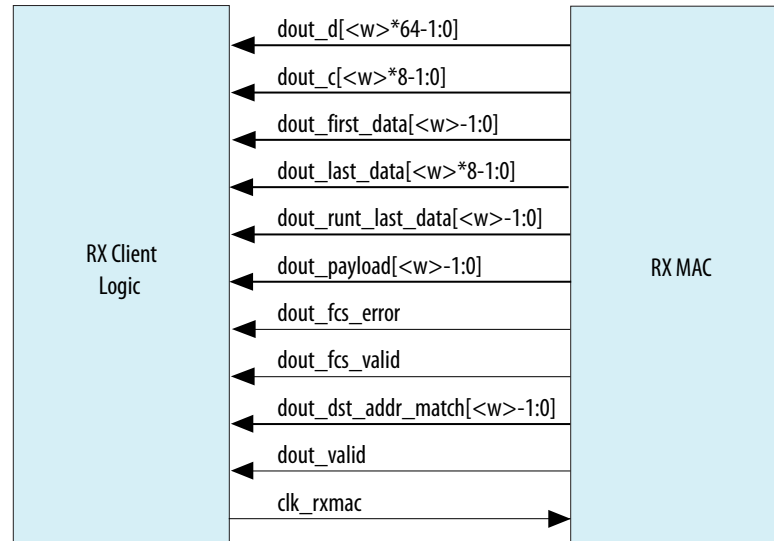


Table 3-6: Signals of the RX Client Interface Without Adapters

In the table, $\langle w \rangle = 2$ for the 40GbE IP core and $\langle w \rangle = 5$ for the 100GbE IP core.

| Signal Name | Direction | Description |
|--|-----------|--|
| <code>dout_d[<w>*64-1:0]</code> | Output | Received data and Idle bytes. In RX preamble pass-through mode, this bus also carries the preamble. |
| <code>dout_c[<w>*8-1:0]</code> | Output | Indicates control bytes on the data bus. Each bit of <code>dout_c</code> indicates whether the corresponding byte of <code>dout_d</code> is a control byte. A bit is asserted high if the corresponding byte on <code>dout_d</code> is an Idle byte or the Start byte, and has the value of zero if the corresponding byte is a data byte or, in preamble pass-through mode, a preamble or SFD byte. |
| <code>dout_first_data[<w>-1:0]</code> | Output | Indicates the first data word of a frame, in the current <code>clk_rxmac</code> cycle. In RX preamble pass-through mode, the first data word is the word that contains the preamble. When the RX preamble pass-through feature is turned off, the first data word is the first word of Ethernet data that follows the preamble. |
| <code>dout_last_data[<w>*8-1:0]</code> | Output | Indicates the final data byte of a frame, before the FCS, in the current <code>clk_rxmac</code> cycle. |

| Signal Name | Direction | Description |
|------------------------------|-----------|--|
| dout_runt_last_data[<w>-1:0] | Output | Indicates that the last_data (the final data byte of the frame) is the final data byte of a runt frame (< 64 bytes). If a frame is eight bytes or smaller, it is considered a decoding error and not a runt frame, and the IP core does not flag it with this signal. |
| dout_payload[<w>-1:0] | Output | Word contains packet data (including destination and source addresses) as opposed to only containing Idle bytes. CRC and padding bits are considered data for this signal. When preamble pass-through is turned on, the preamble is also considered data for this signal. |
| dout_fcs_error | Output | The current or most recent last_data byte is part of a frame with an incorrect FCS (CRC-32) value. By default, the IP core asserts dout_fcs_error in the same cycle as the dout_last_data signal. However, in RX automatic pad removal mode, the dout_fcs_error signal might lag the dout_last_data signal for the frame. |
| dout_fcs_valid | Output | The FCS error bit is valid. |
| dout_dst_addr_match[<w>-1:0] | Output | The first data word in a frame that matches the destination address in the DST_AD0_LO and DST_AD0_HI registers. However, if bit 30 of the MADDR_CTRL register has the value of 0, the address is always considered to match. Otherwise, if bit 0 of the MADDR_CTRL register has the value of 0, the address is always considered to not match. |
| dout_valid | Output | The dout_d bus contents are valid. This signal is occasionally deasserted due to clock crossing. |
| clk_rxmac | Input | RX MAC clock. The minimum clock frequency is 315 MHz. The clk_rxmac clock and the clk_txmac clock (which clocks the TX datapath) are not related and their rates do not have to match. |

The data bytes use 100 Gigabit Media Independent Interface (CGMII-like) encoding. For packet payload bytes, the `dout_c` bit is set to 0 and the `dout_d` byte is the packet data. You can use this information to transmit out-of-spec data such as customized preambles when implementing non-standard variants of the *IEEE 802.3ba-2010 100G Ethernet Standard*. If the additional customized data is not required, simply ignore all words which are marked with `dout_payload = 0` and discard the `dout_c` bus.

In RX preamble pass-through mode, `dout_c` has the value of 1 while the start byte of the preamble is presented on the RX interface, and `dout_c` has the value of 0 while the remainder of the preamble sequence (six-byte preamble plus SFD byte) is presented on the RX interface. While the preamble sequence is presented on the RX interface, `dout_payload` has the value of 1.


Related Information

- [RX Automatic Pad Removal Control](#) on page 3-22
Information about the automatic pad removal mode and its effect on the `dout_fcs_error` signal.
- [MAC Address Registers](#) on page 3-107
Additional information about the `MADDR_CTRL`, `DST_AD0_LO`, and `DST_AD0_HI` address registers and how various register settings interact to affect the value of the `dout_dst_addr_match` signal.

100GbE IP Core RX Client Interface Examples

Example on RX Client Interface Without Preamble Pass-Through

Figure 3-26: Typical Traffic on the RX Client Interface for 100GbE IP Core Without Adapters and with Preamble Pass-Through Turned Off



In the figure, `dout_last_data` is asserted in the second cycle, indicating the end of a packet. The `dout_d` signal returns to 0 (5'h1c = 5'b11100). The `dout_c` and `dout_d` busses are set to 1b'1 and 8'h07, respectively, to indicate idling. In the fourth cycle, `dout_first_data` is asserted and a short packet begins. This packet terminates successfully in the final cycle. Note that the first packet has a CRC error (`dout_fcs_error` = 1 and `fcs_valid` = 1).

The `dout_first_data` signal marks the first word of data. The first byte of data is always the most significant byte of the word. There are 5 legal starting positions for the 100GbE IP core and 2 legal starting positions for the 40GbE IP core. The `dout_last_data` signal marks the last data byte before the FCS (CRC). It can occur at any byte position.

The `dout_runt_last_data` signal indicates that the packet ending in this word is less than the minimum legal length of 64 bytes from first data to the last FCS byte inclusive. Runt frames of eight or fewer bytes cannot be legally represented in CGMII and trigger a decoding error rather than this flag.


The `dout_fcs_error` and `dout_fcs_valid` signals indicate the result of the CRC checking logic. `dout_fcs_valid` = 1 and `dout_fcs_error` = 1 indicate a corrupted frame. Note that the CRC checking network works correctly on runt frames of 40–63 bytes. Runt frames of less than 40 bytes may be incorrectly determined to have a proper CRC.

The `dout_payload` signal marks words that contain frame data payload. Words that contain data payload might also contain Idle or sequence control information, or preamble/frame delimiters. However, if the word contains any frame data, the `dout_payload` bit that corresponds to that word has the value of 1.

The `dout_valid` signal exists for clock rate compensation. This signal is almost always asserted. When `dout_valid` is deasserted all other `dout` signals should be ignored for one clock cycle.

Example on RX Client Interface With Preamble Pass-Through

Figure 3-27: Typical Traffic on the RX Client Interface for 100GbE IP Core Without Adapters and With Preamble Pass-Through



In the figure, the IP core is in preamble pass-through mode, and places the preamble sequence on the data bus `dout_d`. The `dout_first_data` signal marks the first word of data, including the full preamble sequence, beginning with the Start byte. The data in each packet begins with the Start byte (0x55), six-byte preamble (0x555555555555), and SFD (0xD5), and must align with one of the legal starting positions. The IP core asserts the `dout_c` signal high for the Start byte, and holds `dout_c` low for the remainder of the preamble, which is considered to be data. The `dout_c` signal is the only signal that distinguishes any part of the preamble sequence from data. The second and third packets in the example each begin in the first word of `dout_d`, with the preamble sequence. However, they could each legally begin at the start of any other word of `dout_d`; the start of each word is a legal starting position.

Error Conditions on the RX Datapath

The RX MAC indicates error conditions by asserting `l<n>_rx_error`. The following error conditions are defined:

- Received frame terminated early or with an error
- Received frame has a CRC error
- Error characters received from PHY
- Receive frame is too short (less than 64 bytes) or too long (longer than the maximum specified length)

40GbE Lower Rate 24.24 Gbps MAC and PHY

The 40GbE MAC and PHY IP core configured on certain device speed grades can run at 24.24 Gbps with a 4 x 6.25 Gbps external interface. To generate a 40GbE IP core that runs at the 24.24 Gbps rate, the Quartus II software generates the PHY with a data rate of 6250.0 Mbps, instead of the 10312.5 Mbps data rate for the normal 40GbE IP core variations.

Related Information

[40-100GbE IP Core Device Speed Grade Support](#) on page 1-5

Information about the device speed grades that support the 24.24 Gbps IP core variation.

100GbE CAUI-4 PHY

The 100GbE PHY IP core configured in a Stratix V GT device supports CAUI-4 PCS and PMA at 4 x 25.78125 Gbps. The PHY and CAUI PHY interfaces are interoperable with the 100GbE MAC.

The 100GbE CAUI-4 PHY is only available in duplex mode. The 100GbE CAUI-4 PHY, and all other Stratix V variants, require an external reconfiguration controller.

Related Information

- [40-100GbE IP Core Device Speed Grade Support](#) on page 1-5
Information about the device speed grades that support the 100GbE CAUI-4 Gbps IP core variation.
- [External Reconfiguration Controller](#) on page 3-32

External Reconfiguration Controller

40GbE and 100GbE IP cores that target a Stratix IV, Stratix V, or Arria V GZ device, and that include a PHY, require an external reconfiguration controller.

Altera recommends that you configure an Altera ALTGX_RECONFIG megafunction for your Stratix IV 40-100GbE IP core, and an Altera Transceiver Reconfiguration Controller for your Arria V GZ or Stratix V 40-100GbE IP core.

Related Information

- [External Transceiver Reconfiguration Controller](#) on page 2-11
Information about configuring and connecting the Altera Transceiver Reconfiguration Controller. Includes signal descriptions.
- [40-100GbE IP Core Example Design](#) on page 5-1
Refer to the Stratix V example designs for a guide to connecting Altera Transceiver Reconfiguration Controllers to your 40-100GbE IP core.

Congestion and Flow Control Using Pause Frames

The 40-100GbE IP core provides flow control to reduce congestion at the local or remote link partner. When either link partner experiences congestion, the respective transmit control sends pause frames. The pause frame instructs the remote transmitter to stop sending data for the duration that the congested receiver specified in an incoming XOFF frame.

When the IP core receives the XOFF pause control frame, if the following conditions all hold, the IP core stops transmitting frames for a period equal to the pause quanta of the incoming pause frame. While paused, the IP core does not transmit data but can transmit pause frames.

- The appropriate bit of the `RECEIVE_PAUSE_CONTROL` register has the value of 1.
- Address matching is positive.

The pause quanta can be configured in the pause quanta register of the device sending XOFF frames. If the pause frame is received in the middle of a frame transmission, the transmitter finishes sending the current frame and then suspends transmission for a period specified by the pause quanta. Data transmission resumes when a pause frame with quanta of zero is received or when the timer has expired. The pause quanta received overrides any counter currently stored. When more than one pause quanta is sent, the value of the pause is set to the last quanta received.

XOFF pause frames stop the remote transmitter. XON pause frames let the remote transmitter resume data transmission.

One pause quanta fraction is equivalent to 512 bit times, which equates to 512/64 (the width of the MAC data bus), or eight system clock cycles.

Figure 3-28: The XOFF and XON Pause Frames

| XOFF Frame | XON Frame |
|---|---|
| START[7:0] | START[7:0] |
| PREAMBLE[47:0] | PREAMBLE[47:0] |
| SFD[7:0] | SFD[7:0] |
| DESTINATION ADDRESS[47:0] = 0x010000C28001 ⁽¹⁰⁾ | DESTINATION ADDRESS[47:0] = 0x010000C28001 |
| SOURCE ADDRESS[47:0] | SOURCE ADDRESS[47:0] |
| TYPE[15:0] = 0x8808 | TYPE[15:0] = 0x8808 |
| OPCODE[15:0] = 0x001 | OPCODE[15:0] = 0x001 |
| PAUSE QUANTA[15:0] = 0xP1, 0xP2 ⁽¹¹⁾ | PAUSE QUANTA[15:0] = 0x0000 |
| PAD[335:0] | PAD[335:0] |
| CRC[31:0] | CRC[31:0] |

⁽¹⁰⁾ This is a multicast destination address. You can use a unicast address if unicast addresses are enabled in the pause register.

⁽¹¹⁾ The bytes P1 and P2 are filled with the value configured in the `pause_quanta` register.

Conditions Triggering XOFF Frame Transmission

The TX MAC transmits XOFF frames when one of the following conditions occurs:

- Client requests XOFF transmission—A client can explicitly request that XOFF frames be sent using the pause control interface signals. When `pause_insert_tx` is asserted and `pause_insert_time` is not zero, an XOFF frame is sent to the Ethernet network when the current frame transmission completes.
- Host (software) requests XOFF transmission—Setting the pause control register triggers a request that an XOFF frame be sent.

Both options are available in the 40-100GbE IP core with or without adapters.

Conditions Triggering XON Frame Transmission


The TX MAC transmits XON frames when one of the following conditions occurs:

- Client requests XON transmission—A client can explicitly request that XON frames be sent using the pause control interface signals. If `pause_insert_tx` is asserted and `pause_insert_time` is zero, an XON frame is sent to the Ethernet network when the current frame transmission completes.
- Host (software) requests XON transmission—Setting the pause control register triggers a request that an XON frame be sent.

Both options are available in the 40-100GbE IP core with or without adapters.

Pause Transmission Logic

Figure 3-29: Block Diagram of the Pause Transmission Logic



Pause Control and Generation Interface

The pause control interface implements flow control as specified by the *IEEE 802.3ba 2010 100G Ethernet Standard*. The pause logic, upon receiving a pause packet, temporarily stops packet transmission, and can pass the pause packets through as normal traffic or drop the pause control frames in the RX direction.

Table 3-7: Pause Control and Generation Signals

Describes the signals that implement pause control. You can also access the pause functionality using the pause registers for any variant of the 40-100GbE IP core.

| Signal Name | Direction | Description |
|--------------------------|-----------|--|
| pause_insert_tx | Input | Edge triggered signal which directs the IP core to insert a pause frame on the Ethernet link. |
| pause_insert_time [15:0] | Input | Specifies the duration of the pause in pause quanta. The pause control settings in the pause registers determine the duration of the pause quanta (pause quanta is equal to 512-bit time). |

| Signal Name | Direction | Description |
|------------------------------|-----------|---|
| pause_insert_mcast | Input | When asserted, specifies that the IP core should generate a pause packet with the well-known multicast address of 01-80-C2-00-00-01. If deasserted, the pause is generated with the specified MAC address (pause_insert_dst), which is typically a unicast address. |
| pause_insert_dst [47:0] | Input | Specifies the MAC address for a unicast pause. |
| pause_insert_src [47:0] | Input | Specifies source address of the pause packet. |
| pause_match_from_rx | Output | Asserted to indicate an RX pause signal match. Used only when RX configurations are instantiated. The IP core asserts this signal when it receives a pause request with an address match, to signal the TX MAC to throttle its transmissions on the Ethernet link. |
| pause_time_from_rx [15:0] | Output | Asserted for RX pause time. Used only when RX configurations are instantiated. |
| pause_match_to_tx | Input | Asserted to indicate a TX pause signal match. Used only when TX configurations are instantiated. |
| pause_time_to_tx [15:0] | Input | Asserted for TX pause time. Used only when TX configurations are instantiated. |

Related Information

- [Pause Control Frame and Non-Pause Control Frame Filtering and Forwarding](#) on page 3-36
Information about enabling and disabling the pause packets pass-through.
- [Pause Registers](#) on page 3-102
You can access the pause functionality using the pause registers for any 40-100GbE IP core variation that includes a MAC component. Values you program in the registers specify the pause quanta.

Pause Control Frame and Non-Pause Control Frame Filtering and Forwarding

The 40GbE and 100GbE MAC IP cores can pass the pause packets through as normal traffic or drop the pause control frames in the RX direction. You can enable and disable pass-through with the following configuration control bits:

- `RX_FILTER_CTRL` bit [4] enables and disables pause filtering.
- `RX_FILTER_CTRL` bit [5] enables and disables control filtering.

By default, pass-through is disabled.

The following rules define pause control frames filtering control:

1. The `RX_FILTER_CTRL` register contains options to filter different packets types, such as runt packets, FCS error packets, address mismatch packets, and so on, from the RX MAC. The `RX_FILTER_CTRL` register contains one bit to enable pause packet filtering and one bit to enable non-pause control

packet filtering. The reset state for both bits is 1, where filtering is enabled. The bits are gated by `RX_FILTER_CTRL` bit [0], which enables and disables all filtering.

2. If you have enabled pause packet filtering, the IP core drops packets that enter the RX MAC and match the length and type of 0x8808 with an opcode of 0x1 (pause packets) and does not process them or forward them to the client interface.
3. If you have enabled non-pause control packet filtering, the IP core drops packets that enter the RX MAC and match the length and type of 0x8808 with an opcode other than 0x1 (pause packets) and does not forward them to the client interface.
4. If you have disabled pause packet filtering, the RX MAC forwards pause packets to the client interface depending on their destination address. If destination address filtering is not enabled, you are forwarded all pause packets. If destination address filtering is enabled, you are only forwarded pause packets with a valid packet multicast address or a destination address matching the 40-100GbE IP core address.

Pause and control packet pass-through do not affect the pause functionality in the TX or RX MAC.

Related Information

[MAC Configuration and Filter Registers](#) on page 3-99
Information about the `RX_FILTER_CTRL` register.

40-100GbE IP Core Modes of Operation

This section explains the cut-through, store and forward, and promiscuous modes of the 40-100GbE IP core.

In the normal mode of operation, the 40-100GbE IP core MAC transmits and receives data through a PHY to and from a remote link partner Ethernet MAC. You can program IP core registers to control the way in which the IP core RX MAC operates.

You can program the RX MAC to selectively filter incoming Ethernet packets based on various criteria. For example, the RX MAC performs address filtering, various header checking, and control frame termination according to the IEEE 802.3 standard. You must enable filtering to discard mismatched destination addresses.

If you choose to accept all incoming Ethernet packets, and not filter on any criterion, except possibly to filter out runt packets, the IP core is configured in cut-through mode. If you filter based on any criterion other than runt packets, the IP core is configured in store and forward mode, in which it buffers the incoming packet for checking before processing in the MAC.

If the IP core is in cut-through mode, it meets the criteria for promiscuous receive mode, as defined in the Ethernet standard. This definition specifies that the Ethernet implementation accept all valid frames, regardless of destination address. In cut-through mode, the IP core accepts all Ethernet frames that are sufficiently well-formed to be identified. Runt frames are invalid frames, according to the Ethernet standard, and therefore their acceptance or rejection is immaterial to the criteria for promiscuous receive mode.

Related Information

[40-100GbE IP Core RX Filtering](#) on page 3-21


Link Fault Signaling Interface

The 40-100GbE IP core provides link fault signaling as defined in the *IEEE 802.3ba-2010 100G Ethernet Standard*. The 40GbE and 100GbE MAC include a Reconciliation Sublayer (RS) located between the

MAC and the XLGMII or CGMII to manage local and remote faults. Link fault signaling on the Ethernet link is disabled by default but can be enabled by the `Enable Link Fault Sequence` register. When enabled, the local RS TX logic can transmit remote fault sequences in case of a local fault and can transmit IDLE control words in case of a remote fault. An additional configuration register (`MAC/RS link fault sequence configuration`) is provided to select the type of information to be transmitted in case of a local or remote fault. Using the configuration bits, you can send remote fault sequence ordered sets, IDLE control words, or regular traffic in the case of a local or remote fault.

The RS RX logic sets `remote_fault_status` or `local_fault_status` to 1 when the RS RX block receives remote fault or local fault sequence ordered sets. When valid data is received in more than 127 columns, the RS RX logic resets the relevant fault status (`remote_fault_status` or `local_fault_status`) to 0.

Figure 3-30: Link Fault Signaling Example



The IEEE standard specifies RS monitoring of `RXC<7:0>` and `RXD<63:0>` for `Sequence ordered_sets`. For more information, refer to *Figure 81-9—Link Fault Signaling state diagram* and *Table 81-5—Sequence ordered_sets* in the *IEEE 802.3ba 2010 100G Ethernet Standard*. The variable `link_fault` is set to indicate the value of an RX `Sequence ordered_set` when four `fault_sequences` containing the same fault value are received with fault sequences separated by less than 128 columns and with no intervening `fault_sequences` of different fault values. The variable `link_fault` is set to OK following any interval of 128 columns not containing a remote fault or local fault `Sequence ordered_set`.

Table 3-8: Signals of the Link Fault Signaling Interface

| Signal Name | Direction | Description |
|-----------------------------------|-----------|--|
| <code>remote_fault_from_rx</code> | Output | Asserted when remote fault is detected in RX MAC. Available in RX-only variations. |
| <code>local_fault_from_rx</code> | Output | Asserted when local fault is detected in RX MAC. Available in RX-only variations. |
| <code>remote_fault_to_tx</code> | Input | Input to the TX MAC. Asserted when remote fault is detected. Visible in TX-only variations and used internally in duplex IP core variations. |
| <code>local_fault_to_tx</code> | Input | Input to the TX MAC. Asserted when local fault is detected. Visible in TX-only variations and used internally in duplex IP core variations. |
| <code>remote_fault_status</code> | Output | Asserted when remote fault is detected in RX MAC in a duplex IP core variation. In duplex IP core variations, <code>remote_fault_from_rx</code> is connected internally to <code>remote_fault_to_tx</code> , and this signal is available externally as <code>remote_fault_status</code> . |

| Signal Name | Direction | Description |
|--------------------|-----------|---|
| local_fault_status | Output | Asserted when local fault is detected in RX MAC in a duplex IP core variation. In duplex IP core variations, local_fault_from_rx is connected internally to local_fault_to_tx, and this signal is available externally as local_fault_status. |

Related Information

- [Link Fault Signaling Registers](#) on page 3-88
Information about the Enable Link Fault Sequence register and the MAC/RS link fault sequence configuration register.
- [IEEE website](#)
The *IEEE 802.3ba –2010 100G Ethernet Standard* is available on the IEEE website.

Statistics Counters Interface

The statistics counters module is a synthesis option that you select in the 40-100GbE parameter editor. However, the statistics status bit output vectors are provided whether you select the statistics counters module option or not.

The increment vectors are brought to the top level as output ports. If you configure the statistics counters module, the increment vectors also function as input ports to the control and status registers (CSR).

Table 3-9: Statistics Counters Increment Vectors

The TX statistics counter increment vectors are clocked by the `clk_txmac` clock, and the RX statistics counter increment vectors are clocked by the `clk_rxmac` clock.

| Name | Signal Direction | Description |
|---|------------------|---|
| TX Statistics Counter Increment Vectors | | |
| tx_inc_64 | Output | Asserted for one cycle when a 64-byte TX frame is transmitted. |
| tx_inc_127 | Output | Asserted for one cycle when a 65–127 byte TX frame is transmitted. |
| tx_inc_255 | Output | Asserted for one cycle when a 128–255 byte TX frame is transmitted. |
| tx_inc_511 | Output | Asserted for one cycle when a 256–511 byte TX frame is transmitted. |
| tx_inc_1023 | Output | Asserted for one cycle when a 512–1023 byte TX frame is transmitted. |
| tx_inc_1518 | Output | Asserted for one cycle when a 1024–1518 byte TX frame is transmitted. |
| tx_inc_max | Output | Asserted for one cycle when a maximum-size TX frame is transmitted. |
| tx_inc_over | Output | Asserted for one cycle when an oversized TX frame is transmitted. |

| Name | Signal Direction | Description |
|--|------------------|--|
| tx_inc_mcast_data_err | Output | Asserted for one cycle when an errored multicast TX frame, excluding control frames, is transmitted. |
| tx_inc_mcast_data_ok | Output | Asserted for one cycle when a valid multicast TX frame, excluding control frames, is transmitted. |
| tx_inc_bcast_data_err | Output | Asserted for one cycle when an errored broadcast TX frame, excluding control frames, is transmitted. |
| tx_inc_bcast_data_ok | Output | Asserted for one cycle when a valid broadcast TX frame, excluding control frames, is transmitted. |
| tx_inc_ucast_data_err | Output | Asserted for one cycle when an errored unicast TX frame, excluding control frames, is transmitted. |
| tx_inc_ucast_data_ok | Output | Asserted for one cycle when a valid unicast TX frame, excluding control frames, is transmitted. |
| tx_inc_mcast_ctrl | Output | Asserted for one cycle when a valid multicast TX frame is transmitted. |
| tx_inc_bcast_ctrl | Output | Asserted for one cycle when a valid broadcast TX frame is transmitted. |
| tx_inc_ucast_ctrl | Output | Asserted for one cycle when a valid unicast TX frame is transmitted. |
| tx_inc_pause | Output | Asserted for one cycle when a valid pause TX frame is transmitted. |
| tx_inc_fcs_err | Output | Asserted for one cycle when a TX packet with FCS errors is transmitted. |
| tx_inc_fragment | Output | Asserted for one cycle when a TX frame less than 64 bytes and reporting a CRC error is transmitted. |
| tx_inc_jabber | Output | Asserted for one cycle when an oversized TX frame reporting a CRC error is transmitted. |
| tx_inc_sizeok_fcserr | Output | Asserted for one cycle when a valid TX frame with FCS errors is transmitted. |
| RX Statistics Counter Increment Vectors | | |
| rx_inc_runt | Output | Asserted for one cycle when an RX runt packet is received. |
| rx_inc_64 | Output | Asserted for one cycle when a 64-byte RX frame is received. |

| Name | Signal Direction | Description |
|-----------------------|------------------|---|
| rx_inc_127 | Output | Asserted for one cycle when a 65–127 byte RX frame is received. |
| rx_inc_255 | Output | Asserted for one cycle when a 128–255 byte RX frame is received. |
| rx_inc_511 | Output | Asserted for one cycle when a 256–511 byte RX frame is received. |
| rx_inc_1023 | Output | Asserted for one cycle when a 512–1023 byte RX frame is received. |
| rx_inc_1518 | Output | Asserted for one cycle when a 1024–1518 byte RX frame is received. |
| rx_inc_max | Output | Asserted for one cycle when a maximum-size RX frame is received. |
| rx_inc_over | Output | Asserted for one cycle when an oversized RX frame is received. |
| rx_inc_mcast_data_err | Output | Asserted for one cycle when an errored multicast RX frame, excluding control frames, is received. |
| rx_inc_mcast_data_ok | Output | Asserted for one cycle when valid a multicast RX frame, excluding control frames, is received. |
| rx_inc_bcast_data_err | Output | Asserted for one cycle when an errored broadcast RX frame, excluding control frames, is received. |
| rx_inc_bcast_data_ok | Output | Asserted for one cycle when a valid broadcast RX frame, excluding control frames, is received. |
| rx_inc_ucast_data_err | Output | Asserted for one cycle when an errored unicast RX frame, excluding control frames, is received. |
| rx_inc_ucast_data_ok | Output | Asserted for one cycle when a valid unicast RX frame, excluding control frames, is received. |
| rx_inc_mcast_ctrl | Output | Asserted for one cycle when a valid multicast RX frame is received. |
| rx_inc_bcast_ctrl | Output | Asserted for one cycle when a valid broadcast RX frame is received. |
| rx_inc_ucast_ctrl | Output | Asserted for one cycle when a valid unicast RX frame is received. |
| rx_inc_pause | Output | Asserted for one cycle when valid RX pause frames are received. |

| Name | Signal Direction | Description |
|----------------------|------------------|--|
| rx_inc_fcs_err | Output | Asserted for one cycle when a RX packet with FCS errors is received. Assertion of this signal might be early or delayed compared to assertion of the <code>dout_fcs_error</code> signal on the RX custom streaming interface, because the IP core asserts <code>rx_inc_fcs_err</code> when the MAC sees the FCS error, and asserts <code>dout_fcs_error</code> when it presents the relevant frame on the custom streaming client interface. Depending on the filtering settings in the <code>RX_FILTER_CTRL</code> register, the frame might not appear at all on the client interface. |
| rx_inc_fragment | Output | Asserted for one cycle when a RX frame less than 64 bytes and reporting a CRC error is received. |
| rx_inc_jabber | Output | Asserted for one cycle when an oversized RX frame reporting a CRC error is received. |
| rx_inc_sizeok_fcserr | Output | Asserted for one cycle when a valid RX frame with FCS errors is received. |

Related Information

[Statistics Registers](#) on page 3-108

The statistics status bit output vectors are provided whether you select the statistics counters module option or not. The increment vectors are brought to the top level as output ports and function as input ports to the control and status registers (CSR).

MAC – PHY XLGMII or CGMII Interface

The PHY side of the MAC implements the XLGMII or CGMII protocol as defined by the IEEE 802.3ba standard. The standard XLGMII or CGMII implementation consists of 32 bit wide data bus. However, the Altera implementation uses a wider bus interface in connecting a MAC to the internal PHY. The width of this interface is 320 bits for the 100GbE IP core and 128 bits for the 40GbE IP core.

Table 3-10: XL/CGMII Permissible Encodings

Lists XL/CGMII permissible encodings. Memorizing a few of the XL/CGMII encodings greatly facilitates understanding of Ethernet waveforms. The XL/CGMII encodings are backwards compatible with older Ethernet and have convenient mnemonics. The `DATAPATH_OPTION` RTL parameter instantiates TX and RX backwards compatibility and is set by default.

| Control | Data | Description |
|---------|------|--|
| 0 | xx | Packet data, including preamble and FCS bytes. |
| 1 | 07 | Idle. |
| 1 | fb | Start of Frame (fb = frame begin). |

| Control | Data | Description |
|---------|------|--|
| 1 | fd | End of Frame (fd = frame done). |
| 1 | fe | XL/CGMII Error. Typically a bit error which switched a 66-bit block between data and control, or corrupt control information (fe = frame error). |

MAC to PHY Connection Interface

Table 3-11: MAC to PHY and PHY to MAC TX and RX Signals

The MAC-PHY connection interface is exposed in the 40-100GbE MAC-only and PHY-only IP core variations. In addition, the `tx_lanes_stable` output signal from the PHY component is available to provide status information to user logic in PHY-only IP core variations and in MAC and PHY IP core variations.

In the table, $\langle w \rangle = 2$ for the 40GbE IP core and $\langle w \rangle = 5$ for the 100GbE IP core.

| Signal Name | Direction | Description |
|--|----------------------------------|--|
| TX Connections (Data from MAC to PHY) | | |
| <code>tx_mii_d[$\langle w \rangle * 64 - 1 : 0$]</code> | Output from MAC and input to PHY | Media independent interface (MII) data TX connection, starting with the least significant bit (LSB). |
| <code>tx_mii_c[$\langle w \rangle * 8 - 1 : 0$]</code> | | MII control TX connection. |
| <code>tx_mii_valid</code> | | Asserted upon valid TX to MII connection. |
| <code>tx_mii_ready</code> | Output from PHY and input to MAC | Asserted when TX data bus is ready for MII connection. |
| <code>tx_lanes_stable</code> | | Asserted when TX lanes are stable and deskewed. This signal is available as an output status signal in MAC and PHY IP core variations as well as in PHY-only variations. |
| RX Connections (Data from PHY to MAC) | | |
| <code>rx_mii_d[$\langle w \rangle * 64 - 1 : 0$]</code> | Output from PHY and input to MAC | MII data RX connection, starting with the LSB and least significant word, outputting a 5-word data stream. |
| <code>rx_mii_c[$\langle w \rangle * 8 - 1 : 0$]</code> | | MII control RX connection. |
| <code>rx_mii_valid</code> | | Asserted when RX blocks are valid. |

Lane to Lane Deskew Interface

The lane to lane deskew signal is included in the 40-100GbE IP core with and without adapters. When both MAC and PHY options are selected, the lane to lane deskew input signal acts as an internal signal. The lane to lane deskew output signal from the PHY component is available to provide status information to user logic in both PHY-only and MAC and PHY IP core variations.

Table 3-12: Lane to Lane Deskew Interface Signals

| Signal Name | Direction | Description |
|----------------|-----------|--|
| lanes_deskewed | Input | Indicates lane to lane skew is corrected. Available as an input to the 40-100GbE MAC IP core only. |
| lanes_deskewed | Output | Indicates lane to lane skew is corrected. Available as an output from the 40-100GbE PHY IP core and the 40-100GbE MAC and PHY IP core. |


PCS Test Pattern Generation and Test Pattern Check

The PCS can generate a test pattern and detect a scrambled idle test pattern. PCS test-pattern mode is suitable for receiver tests and for certain transmitter tests.

When bit 0 of the `TEST_MODE` register at offset 0x019 has the value of 1, a scrambled idle pattern is enabled. In this mode, the scrambler generates a test pattern. The scrambler does not require seeding during test-pattern operation. The input to the scrambler is a control block (blocktype=0x1E). The TX PCS adds synchronous headers and alignment markers to the data stream, enabling the RX PCS to align and deskew the PCS lanes.

For information about the definition of idle test patterns, refer to *Figure 82-5—64B/66B block formats* illustrated in the *IEEE 802.3ba 2010 100G Ethernet Standard*.

Figure 3-31: PCS Test Pattern Generation



The scrambled idle test-pattern checker utilizes the block lock state diagram, the alignment marker state diagram, the PCS deskew state diagram, and the descrambler; these blocks operate as if in normal data reception. The bit error rate (BER) monitor state diagram is disabled during RX test-pattern mode. When

`align_status` is true and the scrambled idle RX test-pattern mode is active, the scrambled idle test-pattern checker observes the synchronous header and the output from the descrambler. When the synchronous header and the output of the descrambler is an idle pattern, a match is detected. When operating in scrambled idle test pattern, the test-pattern error counter counts blocks with a mismatch. Any mismatch indicates an error and shall increment the test-pattern error counter.

The test-pattern check uses the following register fields:

1. Bit 1 of the `TEST_MODE` register at offset 0x019 enables the RX test-pattern mode.
2. The `TEST_PATTERN_COUNTER` register at offset 0x01A, a 32-bit register that saturates, counts the number of mismatched blocks when the IP core is in test-pattern mode.
3. Bit 2 of the `TEST_MODE` register enables software to clear the test-pattern error counter.

Related Information

- [Test Mode Register](#) on page 3-88
Information about the `TEST_MODE` register.
- [Test Pattern Counter Register](#) on page 3-88
Information about the `TEST_PATTERN_COUNTER` register.
- [IEEE website](#)
The *IEEE 802.3ba-2010 100G Ethernet Standard* is available on the IEEE website.

Transceiver PHY Serial Data Interface

The core uses a 40-bit $\times n$ lane digital interface to send data to the TX high-speed serial I/O pins operating at 10.3125 Gbps in the standard 40GbE and 100GbE variations, at 6.25 Gbps in the 24.24 variations, and at 25.78125 Gbps in the CAUI-4 variations. The `rx_serial` and `tx_serial` ports connect to the 10.3125 Gbps, 6.25 Gbps, or 25.78125 Gbps pins. The protocol includes automatic reordering of serial lanes so that any ordering is acceptable. Virtual lanes 0 and 1 transmit data on `tx_serial[0]`.

PCS BER Monitor

The PCS implements bit error rate (BER) monitoring as specified by the *IEEE 802.3ba-2010 100G Ethernet Standard*. When the PCS deskews the data and aligns the lanes, the BER monitor checks the signal quality and asserts `hi_ber` if it detects excessive errors. When `align_status` is asserted and `hi_ber` is deasserted, the RX PCS continuously accepts blocks and generates RXD <63:0> and RXC <7:0> on the XLGMII or CGMII interface.

High BER occurs when 97 invalid 66-bit synchronous headers are detected for 100GbE within 500 μ s or detected for 40GbE within 1.25 ms. When fewer than 97 invalid 66-bit synchronous headers occur in the same window, the IP core exists the high BER state.

For more information, refer to *Figure 82-13—BER monitor state diagram* illustrated in the *IEEE 802.3ba-2010 100G Ethernet Standard*.

Related Information

[IEEE website](#)

The *IEEE 802.3ba-2010 100G Ethernet Standard* is available on the IEEE website.

40GBASE-KR4 IP Core Variations

The 40GBASE-KR4 IP core supports low-level control of analog transceiver properties for link training and auto-negotiation in the absence of a predetermined environment for the IP core. For example, an Ethernet IP core in a backplane may have to communicate with different link partners at different times. When it powers up, the environment parameters may be different than when it ran previously. The environment can also change dynamically, necessitating reset and renegotiation of the Ethernet link.

The 40-100GbE IP core 40GBASE-KR4 variations implement the *IEEE Backplane Ethernet Standard 802.3ap-2007*. The 40-100GbE IP core provides this reconfiguration functionality in Stratix V devices by configuring each physical Ethernet lane with an Altera Backplane Ethernet 10GBASE-KR PHY IP core if you turn on **Enable KR4** in the 40-100GbE parameter editor. The parameter is available in variations parameterized with these values:

- Device family: Stratix V
- MAC configuration: 40GbE
- Core option: "PHY" or "MAC & PHY"
- PHY configuration: 40Gbps (4 x 10)
- Duplex mode: Full Duplex

The PHY IP core includes the option to implement the following features:

- KR auto-negotiation provides a process to explore coordination with a link partner on a variety of different common features. The 40GBASE-KR4 variations of the 40-100GbE IP core can auto-negotiate only to a 40GBASE-KR4 configuration. Turn on the **Enable KR4 Reconfiguration** and **Enable Auto-Negotiation** parameters to configure support for auto-negotiation.
- Link training provides a process for the IP core to train the link to the data frequency of incoming data, while compensating for variations in process, voltage, and temperature. Turn on the **Enable KR4 Reconfiguration** and **Enable Link Training** parameters to configure support for TX link training.

To enable RX link training, you must also turn on the **Enable RX equalization** parameter. Two options are available for TX link training:

- A built-in TX adaptation algorithm.
- A microprocessor interface to support manual control of the link training process. Turn on the **Enable microprocessor interface** parameter to configure this support.
- After the link is up and running, forward error correction (FEC) provides an error detection and correction mechanism to enable noisy channels to achieve the Ethernet-mandated bit error rate (BER) of 10^{-12} . Turn on the **Include FEC sublayer** to configure support for FEC.

The 40GBASE-KR4 IP core variations include separate link training and FEC modules for each of the four Ethernet lanes, and a single auto-negotiation module. You specify the master lane for performing auto-negotiation in the parameter editor, and the IP core also provides register support to modify the selection dynamically.

The 40GBASE-KR4 IP core is designed to connect to a reconfiguration bundle, which includes the Altera Transceiver Reconfiguration Controller and logic to assist in reconfiguring the transceivers into the different modes of operation (AN, LT, and FEC data mode or non-FEC data mode). Altera provides the testbench and example design to assist you in integrating your 40GBASE-KR4 IP core into your complete design. The testbench and design example include the reconfiguration bundle. You can examine the reconfiguration bundle for an example of how to drive and connect the 40GBASE-KR4 IP core.

The 40GBASE-KR4 IP core variations provide two interfaces to control these processes.

Related Information**Altera Transceiver PHY IP Core User Guide**

The 40GBASE-KR4 variations of the 40-100GbE IP core use the Altera 10GBASE-KR PHY IP core. Information about this PHY IP core, including functional descriptions of the listed features, is available in the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide*. In this chapter, functional descriptions of the FEC, AN, and LT features are available in the "Forward Error Correction (Clause 74)", "Auto Negotiation (AN), Clause 73", and "Link Training (LT), Clause 72" sections, respectively.

40GBASE-KR4 Reconfiguration Interface

The 40GBASE-KR4 reconfiguration interface supports low-level control of analog transceiver properties for link training and auto-negotiation in the absence of a predetermined environment for the IP core.

Table 3-13: 40GBASE-KR4 Reconfiguration Interface Signals

Signals with a width of 4 x n are divided into fields of width n. Bits [n-1:0] refer to Lane 0, bits [2n-1:n] refer to Lane 1, bits [3n-1:2n] refer to Lane 2, and bits [4n-1:3n] refer to Lane 3. You can use these signals to dynamically change between auto-negotiation, link training, and normal data modes.

Note that the regular Stratix V dynamic reconfiguration interface, the `reconfig_from_xcvr`, `reconfig_to_xcvr`, and `reconfig_busy` signals, are also available in 40GBASE-KR4 IP core variations. The reconfiguration bundle in the example design includes the Altera Transceiver Reconfiguration Controller. For an example of how to coordinate dynamic transceiver reconfiguration using these two interfaces, the regular Stratix V transceiver reconfiguration interface and the 40GBASE-KR4 specific interface, refer to the example design reconfiguration bundle.

| Signal Name | Direction | Description |
|-------------------------------|-----------|---|
| <code>rc_busy[3:0]</code> | Input | When asserted, indicates that reconfiguration is in progress. |
| <code>lt_start_rc[3:0]</code> | Output | When asserted, starts the TX PMA equalization reconfiguration on the corresponding lane. This signal is present only if link training is enabled. |
| <code>main_rc[23:0]</code> | Output | The main TX equalization tap value, which is the same as V_{OD} . This signal is present only if link training is enabled. |
| <code>post_rc[19:0]</code> | Output | The post-cursor TX equalization tap value for the corresponding lane. This signal is present only if link training is enabled. |
| <code>pre_rc[15:0]</code> | Output | The pre-cursor TX equalization tap value for the corresponding lane. This signal is present only if link training is enabled. |

| Signal Name | Direction | Description |
|--------------------|-----------|---|
| tap_to_upd[11:0] | Output | Specifies the TX equalization tap to update to optimize signal quality. Each lane's field has the following valid values: <ul style="list-style-type: none"> 3'b100: main tap 3'b010: post-tap 3'b001: pre-tap This signal is present only if link training is enabled. |
| seq_start_rc[3:0] | Output | When a bit is asserted, starts PCS reconfiguration for the corresponding lane. |
| dfe_start_rc[3:0] | Output | When a bit is asserted, starts RX DFE equalization for the corresponding lane. This signal is present only if RX equalization is enabled. |
| dfe_mode[7:0] | Output | Specifies the DFE operation mode. Valid at the rising edge of the <code>def_start_rc</code> signal and held until the falling edge of the <code>rc_busy</code> signal. The following encodings are defined for each lane: <ul style="list-style-type: none"> 2'b00: Disable DFE 2'b01: DFE triggered mode (single shot) 2'b10 and 2'b11 are reserved. This signal is present only if RX equalization is enabled. |
| ctle_start_rc[3:0] | Output | When a bit is asserted, starts continuous time-linear equalization (CTLE) reconfiguration on the corresponding lane. This signal is present only if RX equalization is enabled. |
| ctle_rc[15:0] | Output | RX CTLE value. This signal is valid at the rising edge of the <code>ctle_start_rc</code> signal and held until the falling edge of the <code>rc_busy</code> signal. The valid range of values is 4'b0000–4'b1111. 4'b0000 indicates the RX CTLE is disabled and 4'b1111 indicates RX CTLE is at its maximum value. This signal is present only if RX equalization is enabled. |
| ctle_mode[7:0] | Output | Specifies the CTLE mode. This signal is valid at the rising edge of the <code>ctle_start_rc</code> signal and held until the falling edge of the <code>rc_busy</code> signal. The only valid value of this signal in the 40-100GbE IP core is 2'b0. This signal is present only if RX equalization is enabled. |

| Signal Name | Direction | Description |
|------------------|-----------|---|
| pcs_mode_rc[5:0] | Output | Specifies the PCS mode for reconfiguration. Has the following valid values: <ul style="list-style-type: none"> • b'000001: auto-negotiation mode • b'000010: link training mode • b'100000: data mode (normal operation) Other values are not valid for the 40GBASE-KR4 IP core. |
| en_lcl_rxeq[3:0] | Output | When a bit is asserted, it signals that an additional custom RX equalization is enabled for the corresponding lane. The bits are identical to the Link Trained status bits 0xD2 [0], [8], [16], and [24]. |
| rxeq_done[3:0] | Input | When asserted, indicates that custom RX equalization is complete. The PHY IP core ANDs each bit of this signal with rx_trained from the Training State Diagram for the corresponding lane. |
| reco_mif_done | Input | Reset signal the user logic asserts after completing MIF programming. |

Related Information**40-100GbE IP Core Example Design** on page 5-1

The 40GBASE-KR4 example design, specifically the reconfiguration bundle and its connections to the IP core, provide an example of how to coordinate use of the transceiver reconfiguration interface and the 40GBASE-KR4 specific reconfiguration interface to your 40GBASE-KR4 IP core.

40GBASE-KR4 Microprocessor Interface

The optional embedded processor interface signals allow you to use the embedded processor mode of Link Training. This mode overrides the TX adaptation algorithm and allows an embedded processor to initialize the link.

Table 3-14: 40GBASE-KR4 Microprocessor Interface Signals

Signals with a width of 4 x n are divided into fields of width n. Bits [n-1:0] refer to Lane 0, bits [2n-1:n] refer to Lane 1, bits [3n-1:2n] refer to Lane 2, and bits [4n-1:3n] refer to Lane 3. These signals are only available if you turn on **Enable microprocessor interface**.

| Signal Name | Direction | Description |
|------------------|-----------|--|
| upi_mode_en[3:0] | Input | When a bit is asserted, enables embedded processor mode on the corresponding lane. |

| Signal Name | Direction | Description |
|--|-----------|---|
| upi_adj[7:0] | Input | Selects the active tap for the corresponding lane. Each lane's field has the following valid values: <ul style="list-style-type: none"> • 2'b01: main tap • 2'b10: post-tap • 2'b11: pre-tap |
| upi_inc[3:0] | Input | When a bit is asserted, sends the increment command for the corresponding lane. |
| upi_dec[3:0] | Input | When a bit is asserted, sends the decrement command for the corresponding lane. |
| upi_pre[3:0] | Input | When a bit is asserted, sends the preset command for the corresponding lane. |
| upi_init[3:0] | Input | When a bit is asserted, sends the initialize command for the corresponding lane. |
| upi_st_bert[3:0] | Input | When a bit is asserted, starts the BER timer for the corresponding lane. |
| upi_train_err[3:0] | Input | When a bit is asserted, indicates a training error on the corresponding lane. |
| upi_lock_err[3:0] | Input | When a bit is asserted, indicates a training frame lock error on the corresponding lane. |
| upi_rx_trained[3:0] | Input | When a bit is asserted, the RX interface for the corresponding lane is trained. |
| upo_enable[3:0] | Output | When a bit is asserted, indicates that the IP core is ready to receive commands from the embedded processor for the corresponding lane. |
| upo_frame_lock[3:0] | Output | When a bit is asserted, indicates the receiver has achieved training frame lock on the corresponding lane. |
| upo_cm_done[3:0] | Output | When a bit is asserted, indicates the master state machine handshake for the corresponding lane is complete. |
| upo_bert_done[3:0] | Output | When a bit is asserted, indicates the BER timing for the corresponding lane is at its maximum count. |
| upo_ber_cnt[4*<bcw>-1:0] (width varies with <bcw> = BER counter width) | Output | Each four-bit field holds the current BER count for the corresponding lane. |

| Signal Name | Direction | Description |
|-------------------|-----------|---|
| upo_ber_max[3:0] | Output | When a bit is asserted, the BER counter for the corresponding lane has rolled over. |
| upo_coef_max[3:0] | Output | When a bit is asserted, indicates that the remote coefficients for the corresponding lane are at their maximum or minimum values. |

Control and Status Interface

The control and status interface provides an Avalon-MM interface to the control and status registers. The Avalon-MM interface implements a standard memory-mapped protocol. You can connect an embedded processor or JTAG Avalon master to this bus to access the control and status registers.

Table 3-15: Avalon-MM Control and Status Interface Signals

The `clk_status` clocks the signals on the 40-100GbE IP core control and status interface.

| Signal Name | Direction | Description |
|-------------------------|-----------|------------------------------|
| status_addr [15:0] | Input | Address for reads and writes |
| status_read | Input | Read command |
| status_write | Input | Write command |
| status_writedata [31:0] | Input | Data to be written |
| status_readdata [31:0] | Output | Read data |
| status_readdata_valid | Output | Read data is ready for use |

The status interface is designed to operate at a low frequencies, typically 50 MHz for Stratix IV devices and 100 MHz for Stratix V devices, so that control and status logic does not compete for resources with the surrounding high speed datapath.

Related Information

- [Software Interface: Registers](#) on page 3-76
- [40-100GbE IP Core Registers](#) on page 3-79
- [Avalon Interface Specifications](#)

For more information about the Avalon-MM protocol, including timing diagrams, refer to the *Avalon Memory-Mapped Interfaces* chapter.

Clocks

You must set the transceiver reference clock (`clk_ref`) frequency to a value that your IP core variation supports. For most variations, The 40-100GbE IP core supports `clk_ref` frequencies of 644.53125 MHz \pm 100 ppm and 322.265625 MHz \pm 100 ppm. The \pm 100ppm value is required for any clock source providing the transceiver reference clock. For CAUI-4 variations, you must set the frequency of `clk_ref`

to 644.53125 MHz \pm 100 ppm. For 24.24 Gbps variations, you must set the frequency of `clk_ref` either to 390.625 MHz \pm 100 ppm or to 195.3125 MHz \pm 100 ppm.

Sync-E IP core variations are duplex IP core variations for which you turn on **Enable SyncE support** in the parameter editor. These variations provide separate IP core input reference clock signals for the TX and RX transceiver PLLs, and provide the RX recovered clock as a top-level output signal.

The Synchronous Ethernet standard, described in the ITU-T G.8261, G.8262, and G.8264 recommendations, requires that the TX clock be filtered to maintain synchronization with the RX reference clock through a sequence of nodes. The expected usage is that user logic drives the `tx_ref_clk` signal with a filtered version of the RX recovered clock signal, to ensure the receive and transmit functions remain synchronized.

In a Sync-E IP core, the restrictions apply to each of the `rx_clk_ref` and `tx_clk_ref` input clocks.

The minimum clock frequency for the IP core is 315 MHz. The only exception is the 40GbE lower rate 24.24 Gbps MAC and PHY IP core, which has a minimum clock frequency of 190.90 MHz.

Table 3-16: Clock Inputs

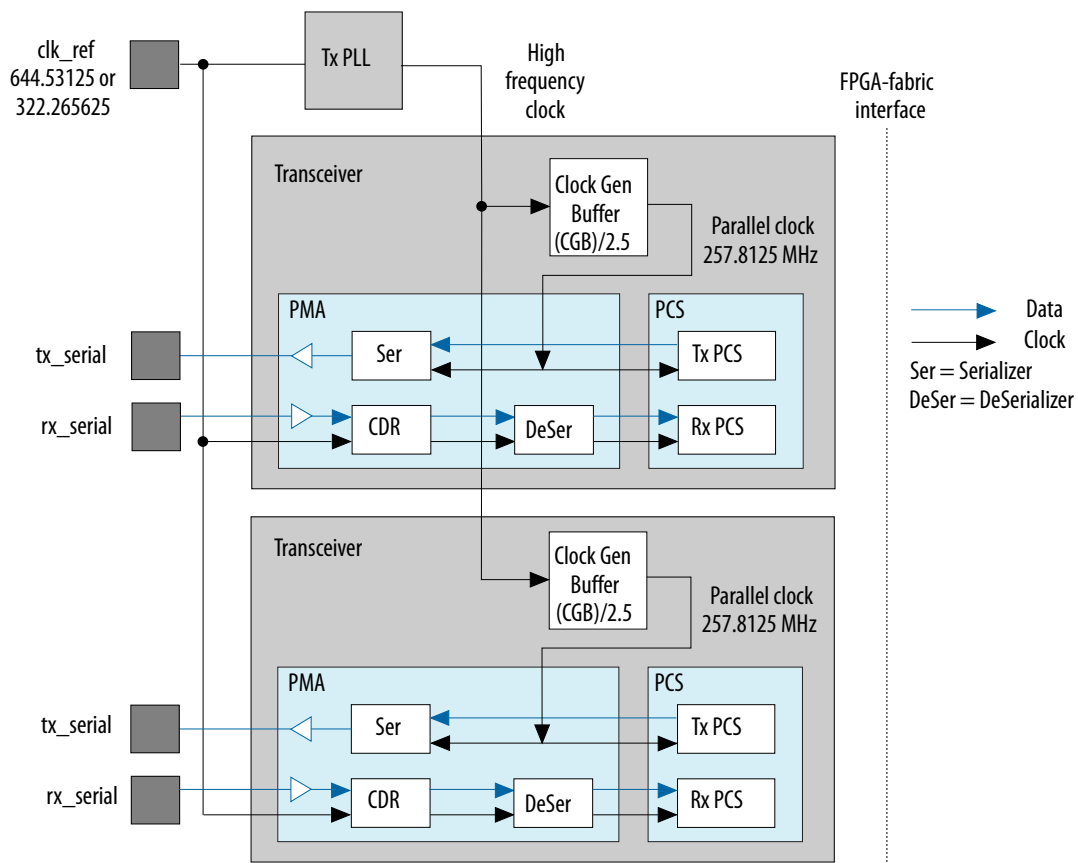
Describes the input clocks that you must provide.

| Signal Name | Description |
|-------------------------|---|
| <code>clk_status</code> | A clock for reconfiguration, offset cancellation, and housekeeping functions. This clock is also used for clocking the control and status interface. The clock quality and pin chosen are not critical. <code>clk_status</code> is expected to be a 37.5–50 MHz clock on Stratix IV devices and a 100–125 MHz clock on Stratix V devices. |
| <code>clk_ref</code> | <p><code>clk_ref</code> is the reference clock for the transceiver TX PLL and the RX CDR PLL. This input signal is not available in Sync-E variations.</p> <p>The frequency of this input clock must match the value you specify for PHY reference frequency in the IP core parameter editor.</p> <p>For the regular 40GbE and 100GbE IP core variations, this clock must have a frequency of 322.265625 or 644.53125 MHz with a \pm100 ppm accuracy per the <i>IEEE 802.3ba-2010 100G Ethernet Standard</i>.</p> <p>Despite its apparent availability in the 40-100GbE parameter editor, CAUI-4 variations do not support the 322 MHz clock frequency. For these variations, this clock must have a frequency of 644.53125 MHz with a \pm100 ppm accuracy.</p> <p>For 24.24 Gbps IP core variations, this clock must have a frequency of 390.625 or 195.3125 MHz with a \pm100 ppm accuracy.</p> <p>In addition, <code>clk_ref</code> must meet the jitter specification of the <i>IEEE 802.3ba-2010 100G Ethernet Standard</i>.</p> <p>The PLL and clock generation logic use this reference clock to derive the transceiver and PCS clocks. The input clock should be a high quality signal on the appropriate dedicated clock pin.</p> <p>The PCS clock frequency is 257.8125 MHz for standard variations, 201.416 MHz for CAUI-4 variations, and 156.25 MHz for 24.24 Gbps variations.</p> |

| Signal Name | Description |
|-------------|--|
| tx_clk_ref | <p>In Sync-E variations (IP core duplex variations with the Sync-E option enabled), this clock replaces <code>clk_ref</code> as the reference clock for the transceiver TX PLL.</p> <p>The frequency of this input clock must match the value you specify for PHY reference frequency in the IP core parameter editor.</p> |
| rx_clk_ref | <p>In Sync-E variations (IP core duplex variations with the Sync-E option enabled), this clock replaces <code>clk_ref</code> as the reference clock for the transceiver CDR PLL.</p> <p>The frequency of this input clock must match the value you specify for PHY reference frequency in the IP core parameter editor.</p> |
| clk_txmac | <p>The input TX clock for the IP core with or without adapters is <code>clk_txmac</code>. The recommended TX MAC clock frequency is 190.90 MHz for 24.24 Gbps variations, and 315 MHz for all other IP core variations.</p> |
| clk_rxmac | <p>The input RX clock for the IP core with or without adapters is <code>clk_rxmac</code>. The recommended TX MAC clock frequency is 190.90 MHz for 24.24 Gbps variations, and 315 MHz for all other IP core variations.</p> |

Figure 3-32: Clock Generation Circuitry

Provides a high-level view of the clock generation circuitry and clock distribution to the transceiver. In Sync-E variations, distinct clocks drive the TX PLL (`tx_clk_ref`) and the CDR block (`rx_clk_ref`), and the output clock from the CDR is brought out to the top level.



Resets

The 40-100GbE IP core provides the following two independent reset mechanisms:

- Asynchronous reset signals—A set of asynchronous reset signals you can assert to reset different parts of the IP core. Use this method to initialize your IP core.
- Reset registers—A set of register bits you can write to reset different parts of the IP core. This method is available for dynamic reset during operation.

Table 3-17: Asynchronous Reset Signals

The IP core provides five reset signals to allow independent reset control for all configurations. The MAC and PHY asynchronous reset signals are included in the 40-100GbE IP Core with adapters and without adapters.

| Signal Name | Direction | Description |
|-----------------------------|-----------|----------------------------------|
| <code>mac_rx_arst_ST</code> | Input | MAC RX asynchronous reset signal |
| <code>mac_tx_arst_ST</code> | Input | MAC TX asynchronous reset signal |

| Signal Name | Direction | Description |
|----------------|-----------|--------------------------------------|
| pcs_rx_arst_ST | Input | PHY PCS RX asynchronous reset signal |
| pcs_tx_arst_ST | Input | PHY PCS TX asynchronous reset signal |
| pma_arst_ST | Input | PHY PMA asynchronous reset signal |

Note: In any MAC and PHY variation, when you reset the TX MAC you must also reset the TX PCS to avoid transmitting corrupted packets. Therefore, Altera recommends that you reset the IP core with the following conditions:

- Reset the TX MAC and the TX PCS together (assert `pcs_tx_arst_ST` and `mac_tx_arst_ST` simultaneously).
- Release `pcs_tx_arst_ST` and `mac_tx_arst_ST` simultaneously or release `pcs_tx_arst_ST` after you release `mac_tx_arst_ST`.

Altera recommends that you release all parts of the 40-100GbE IP core from reset simultaneously.

Note: Each reset signal must be asserted for at least one `clk_status` cycle.

You should not release any reset signal until after you observe that the reference clock is stable. If the reference clock is generated from an fPLL, wait until after the fPLL locks.

Related Information

[MAC and PHY Reset Registers](#) on page 3-89

For information about the IP core reset registers, including the recommended reset sequence. Use the reset registers for dynamic reset during operation.

Signals

This section lists the external signals of the different 40-100GbE IP core variations.

[Signals of MAC and PHY Variations Without Adapters](#) on page 3-55

[Signals of MAC and PHY Variations With Adapters](#) on page 3-66

[Signals of 40-100GbE MAC-Only IP Core Variations](#) on page 3-68


[Signals of 40-100GbE PHY-Only IP Core Variations](#) on page 3-72

Signals of MAC and PHY Variations Without Adapters

The signals of the MAC and PHY variations without adapters are described in the following formats:

- The figure identifies the IP core interfaces and the presence or absence of various signals in MAC and PHY, PHY-only, and MAC-only IP core variations.
- The tables identify the different signals available in TX-only, RX-only, and duplex IP core variations.
- Links guide you to descriptions for the individual signals, by interface. The links are available only if you are viewing this topic in the context of the Functional Description chapter of the *40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide*.

Figure 3-33: Top-Level Signals of the 40-100GbE IP Core Without Adapters



The following tables list the signals for the 40-100GbE MAC and PHY IP core without adapters:

- Transmit side signals configured only in TX-only and duplex IP core variations
- Receive side signals configured only in RX-only and duplex IP core variations
- Signals configured in TX-only, RX-only, and duplex IP core variations

For links to the relevant interface information, including a description of each signal, refer to the Related Links below.

Table 3-18: 40-100GbE MAC and PHY IP Core Without Adapters: Transmit Side Signals

| Signal Name | Direction | Interface |
|---|-----------|--|
| mac_tx_arst_ST | Input | Resets |
| pcs_tx_arst_ST | Input | |
| tx_serial [3:0] (40GbE and CAUI-4) tx_serial [9:0] (standard 100GbE) | Output | Transceiver PHY serial data interface |
| tx_lanes_stable | Output | PHY status |
| clk_txmac | Input | Clocks TX client interface without adapters (custom streaming interface) |
| din[<w>*64-1:0] | Input | TX client interface without adapters (custom streaming interface) |
| din_start[<w>-1:0] | Input | |
| din_end_pos[<w>*8-1:0] | Input | |
| din_ack | Output | |
| pause_insert_tx | Input | Pause control and generation interface The _to_tx signals are not visible in duplex variations. |
| pause_insert_time[15:0] | Input | |
| pause_insert_mcast | Input | |
| pause_insert_dst[47:0] | Input | |
| pause_insert_src[47:0] | Input | |
| pause_match_to_tx | Input | |
| pause_time_to_tx[15:0] | Input | Link fault signaling interface These two signals are not visible in duplex variations. |
| remote_fault_to_tx | Input | |
| local_fault_to_tx | Input | |

| Signal Name | Direction | Interface |
|-----------------------|-----------|--------------------------------------|
| tx_inc_64 | Output | Statistics counter increment vectors |
| tx_inc_127 | Output | |
| tx_inc_255 | Output | |
| tx_inc_511 | Output | |
| tx_inc_1023 | Output | |
| tx_inc_1518 | Output | |
| tx_inc_max | Output | |
| tx_inc_over | Output | |
| tx_inc_mcast_data_err | Output | |
| tx_inc_mcast_data_ok | Output | |
| tx_inc_bcast_data_err | Output | |
| tx_inc_bcast_data_ok | Output | |
| tx_inc_ucast_data_err | Output | |
| tx_inc_ucast_data_ok | Output | |
| tx_inc_mcast_ctrl | Output | |
| tx_inc_bcast_ctrl | Output | |
| tx_inc_ucast_ctrl | Output | |
| tx_inc_pause | Output | |
| tx_inc_fcs_err | Output | |
| tx_inc_fragment | Output | |
| tx_inc_jabber | Output | |
| tx_inc_sizeok_fcserr | Output | |

Table 3-19: 40-100GbE MAC and PHY IP Core Without Adapters: Receive Side Signals

| Signal Name | Direction | Description |
|---|-----------|---|
| mac_rx_arst_ST | Input | Resets |
| pcs_rx_arst_ST | Input | |
| rx_serial [3:0] (40GbE and CAUI-4) rx_serial [9:0] (standard 100GbE) | Input | Transceiver PHY serial data interface |
| lanes_deskewed | Output | PHY status |
| clk_rxmac | Input | Clocks RX client interface without adapters (custom streaming interface) |
| dout_d[<w>*64-1:0] | Output | RX client interface without adapters (custom streaming interface) |
| dout_c[<w>*8-1:0] | Output | |
| dout_first_data[<w>-1:0] | Output | |
| dout_last_data[<w>*8-1:0] | Output | |
| dout_runt_last_data[<w>-1:0] | Output | |
| dout_payload[<w>-1:0] | Output | |
| dout_fcs_error | Output | |
| dout_fcs_valid | Output | |
| dout_dst_addr_match[<w>-1:0] | Output | |
| dout_valid | Output | |
| pause_match_from_rx | Output | |
| pause_time_from_rx[15:0] | Output | These two signals are not visible in duplex variations. |

| Signal Name | Direction | Description |
|----------------------|-----------|---|
| remote_fault_from_rx | Output | Link fault signaling interface |
| local_fault_from_rx | Output | These two signals are not visible in duplex variations. |

| Signal Name | Direction | Description |
|-----------------------|-----------|--------------------------------------|
| rx_inc_runt | Output | Statistics counter increment vectors |
| rx_inc_64 | Output | |
| rx_inc_127 | Output | |
| rx_inc_255 | Output | |
| rx_inc_511 | Output | |
| rx_inc_1023 | Output | |
| rx_inc_1518 | Output | |
| rx_inc_max | Output | |
| rx_inc_over | Output | |
| rx_inc_mcast_data_err | Output | |
| rx_inc_mcast_data_ok | Output | |
| rx_inc_bcast_data_err | Output | |
| rx_inc_bcast_data_ok | Output | |
| rx_inc_ucast_data_err | Output | |
| rx_inc_ucast_data_ok | Output | |
| rx_inc_mcast_ctrl | Output | |
| rx_inc_bcast_ctrl | Output | |
| rx_inc_ucast_ctrl | Output | |
| rx_inc_pause | Output | |
| rx_inc_fcs_err | Output | |
| rx_inc_fragment | Output | |
| rx_inc_jabber | Output | |
| rx_inc_sizeok_fcserr | Output | |

Table 3-20: 40-100GbE MAC and PHY IP Core Without Adapters: Common Signals

| Signal Name | Direction | Description |
|------------------------|-----------|---|
| pma_arst_ST | Input | Resets |
| clk_ref | Input | Clocks |
| tx_clk_ref | Input | Clocks |
| rx_clk_ref | Input | In Sync-E variations, these two clock signals replace clk_ref. All Sync-E variations are duplex IP cores. |
| rx_recovered_clk | Output | Clocks This signal is present only in Sync-E variations. |
| status_addr[15:0] | Input | Control and status interface |
| status_read | Input | |
| status_write | Input | |
| status_writedata[31:0] | Input | |
| status_readdata[31:0] | Output | |
| status_readdata_valid | Output | |
| clk_status | Input | Clocks |
| remote_fault_status | Output | Link fault signaling interface |
| local_fault_status | Output | These two signals are available only in duplex variations. |

| | | |
|---------------------|--------|---|
| reconfig_to_xcvr | Input | <p>External reconfiguration controller interface</p> <p>These signals are available in Arria V GZ and Stratix V devices only.</p> <p>The <code>_to_xcvr<n></code> and <code>_from_xcvr<n></code> signals are present only in CAUI-4 variations.</p> |
| reconfig_to_xcvr0 | Input | |
| reconfig_to_xcvr1 | Input | |
| reconfig_to_xcvr2 | Input | |
| reconfig_to_xcvr3 | Input | |
| reconfig_from_xcvr | Output | |
| reconfig_from_xcvr0 | Output | |
| reconfig_from_xcvr1 | Output | |
| reconfig_from_xcvr2 | Output | |
| reconfig_from_xcvr3 | Output | |

| | | |
|---|--------|--|
| upi_mode_en[3:0] | Input | 40GBASE-KR4 microprocessor interface. These signals are present only in 40GBASE-KR4 variations for which you turn on Enable microprocessor interface . All 40GBASE-KR4 variations are in Duplex mode. |
| upi_adj[7:0] | Input | |
| upi_inc[3:0] | Input | |
| upi_dec[3:0] | Input | |
| upi_pre[3:0] | Input | |
| upi_init[3:0] | Input | |
| upi_st_bert[3:0] | Input | |
| upi_train_err[3:0] | Input | |
| upi_lock_err[3:0] | Input | |
| upi_rx_trained[3:0] | Input | |
| upo_enable[3:0] | Output | |
| upo_frame_lock[3:0] | Output | |
| upo_cm_done[3:0] | Output | |
| upo_bert_done[3:0] | Output | |
| upo_ber_cnt[4*<bcw>-1:0] (width varies with <bcw> = BER counter width) | Output | |
| upo_ber_max[3:0] | Output | |
| upo_coef_max[3:0] | Output | |

| | | |
|---------------------|--------|---|
| rc_busy[3:0] | Input | |
| lt_start_rc[3:0] | Output | |
| main_rc[23:0] | Output | |
| post_rc[19:0] | Output | |
| pre_rc[15:0] | Output | 40GBASE-KR4 reconfiguration interface. These signals are present only in 40GBASE-KR4 variations for which you turn on Enable KR4 Reconfiguration . All 40GBASE-KR4 variations are in Duplex mode. The lt_start_rc, main_rc, post_rc, pre_rc, and tap_to_upd signals are present only if you turn on Enable Link Training . The dfe_start_rc, dfe_mode, ctile_start_rc, ctile_rc, and ctile_mode signals are present only if you turn on Enable RX equalization . |
| tap_to_upd[11:0] | Output | |
| seq_start_rc[3:0] | Output | |
| dfe_start_rc[3:0] | Output | |
| dfe_mode[7:0] | Output | |
| ctile_start_rc[3:0] | Output | |
| ctile_rc[15:0] | Output | |
| ctile_mode[7:0] | Output | |
| pcs_mode_rc[5:0] | Output | |
| en_lcl_rxeq[3:0] | Output | |
| rxeq_done[3:0] | Input | |
| reco_mif_done | Input | |

Related Information

- [Resets](#) on page 3-54
Overview of IP core reset access. Includes list of reset signals and recommended reset sequence.
- [Transceiver PHY Serial Data Interface](#) on page 3-45
- [MAC to PHY Connection Interface](#) on page 3-43
Describes the tx_lanes_stable PHY output signal.
- [Clocks](#) on page 3-51
Overview of IP core clocks. Includes list of clock signals and recommended and required frequencies.
- [40-100GbE IP Core TX Data Bus Without Adapters \(Custom Streaming Interface\)](#) on page 3-9
- [Pause Control and Generation Interface](#) on page 3-35
- [Link Fault Signaling Interface](#) on page 3-37
- [Statistics Counters Interface](#) on page 3-39
- [Statistics Registers](#) on page 3-108
- [Lane to Lane Deskew Interface](#) on page 3-43
Describes the lanes_deskewed PHY output signal.

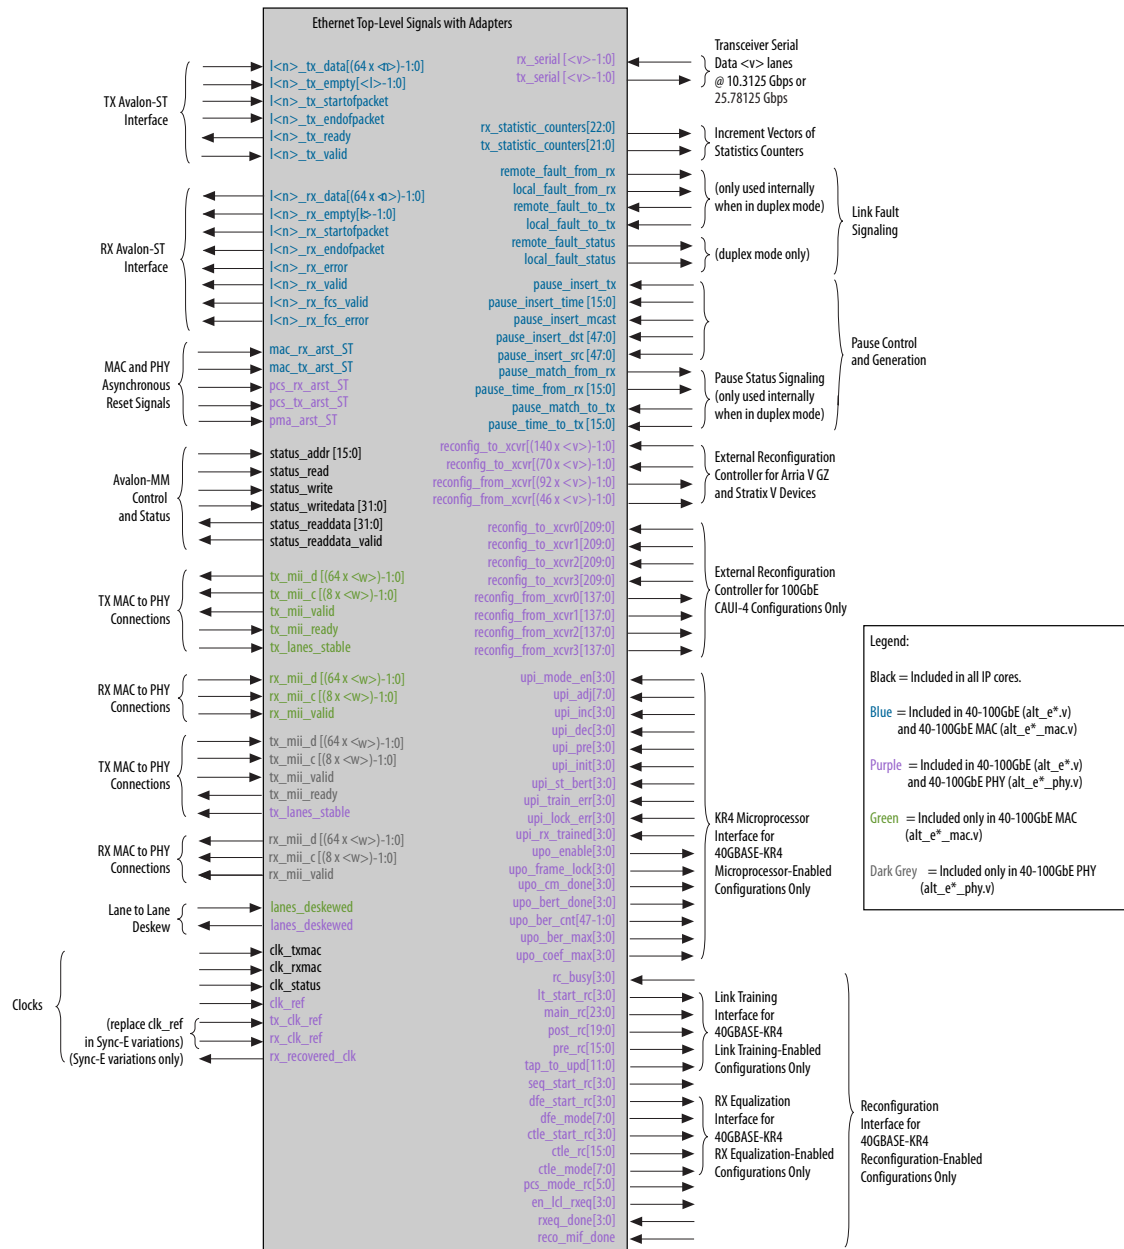
- [40-100GbE IP Core RX Data Bus Without Adapters \(Custom Streaming Interface\)](#) on page 3-28
- [Control and Status Interface](#) on page 3-51
- [External Reconfiguration Controller](#) on page 3-32
- [40GBASE-KR4 Microprocessor Interface](#) on page 3-49
- [40GBASE-KR4 Reconfiguration Interface](#) on page 3-47
- [40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide](#)
If you read this topic in the full user guide, links guide you to the descriptions of the individual signals, by interface.

Signals of MAC and PHY Variations With Adapters

Most signals of the 40-100GbE IP core with adapters are identical to the signals for the IP core without adapters. The IP core variations with adapters have the following differences:

- On the transmit side, in TX-only and duplex variations, the TX Avalon-ST client interface signals replace the TX custom streaming client interface signals.
- On the receive side, in RX-only and duplex variations, the RX Avalon-ST client interface signals replace the RX custom streaming client interface signals.

Figure 3-34: Top-Level Signals of the 40-100GbE IP Core with Adapters



Related Information

- **Signals of MAC and PHY Variations Without Adapters** on page 3-55
Signals in variations with adapters are the same as signals in variations without adapters on most interfaces. The only exceptions are the RX and TX client interfaces (completely different).
- **40-100GbE IP Core TX Data Bus with Adapters (Avalon-ST Interface)** on page 3-6
Describes the TX client interface in IP core variations with adapters.
- **40-100GbE IP Core RX Data Bus with Adapters (Avalon-ST Interface)** on page 3-25
Describes the RX client interface in IP core variations with adapters.

Signals of 40-100GbE MAC-Only IP Core Variations

40-100GbE MAC-only IP core variations are the variations that do not include a PHY. The signals in a MAC-only variation with and without adapters are described in the following formats:

- The figures identify the IP core interfaces and the presence or absence of various signals in MAC-only IP core variations.
- Text lists the basic differences between the MAC and PHY IP core variations and the MAC-only IP core variations.
- Text lists the MAC and PHY IP core variation signals absent in the MAC-only variations.
- A table lists the signals visible in the MAC-only variations that are not visible in the MAC and PHY variations.

40-100GbE MAC-only IP core variations are the variations that do not include a PHY. The signals in a MAC-only variation with and without adapters are shown in the following figures in black, blue, or green. Signals shown in purple and in dark gray are not available in the MAC-only variations.

Figure 3-35: Top-Level Signals of the 40-100GbE IP Core with Adapters

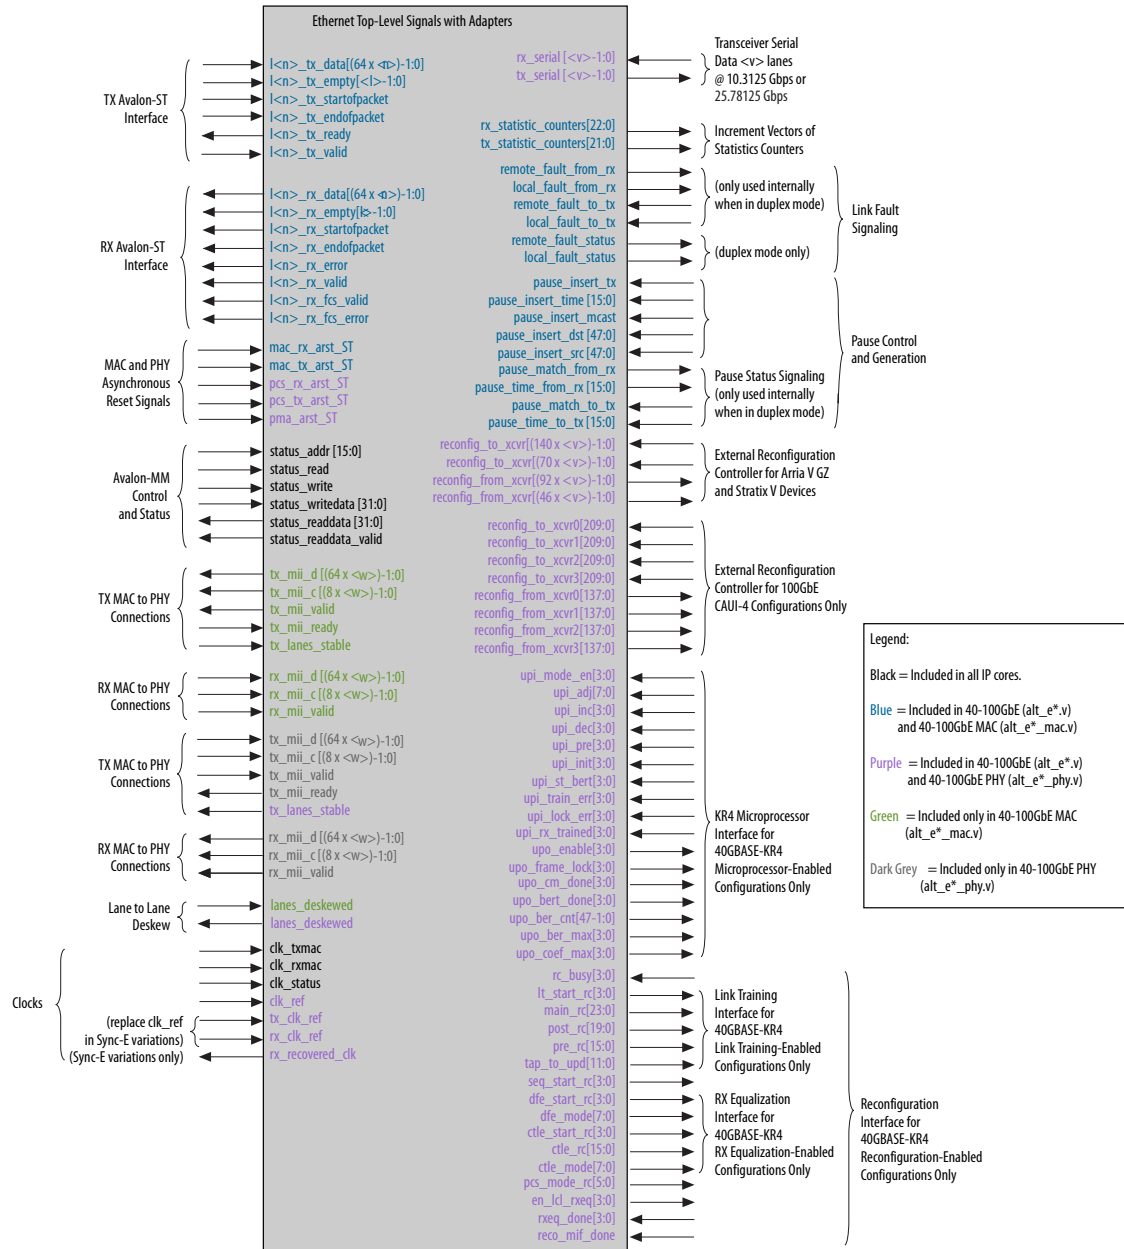



Figure 3-36: Top-Level Signals of the 40-100GbE IP Core Without Adapters



40-100GbE IP core variations that do not include a PHY have the following differences from the variations that include the PHY:

- The 40-100GbE PHY output signals are not available in the MAC-only variations.
- The MAC-only variations include an interface to connect to a PHY.

The following 40-100GbE MAC and PHY IP core signals are not available in MAC-only IP core variations:

- Clock signals:
 - `clk_ref` (relevant only for IP core variations with the Sync-E support option turned off)
 - `tx_clk_ref` (relevant only for IP core variations with the Sync-E support option turned on)
 - `rx_clk_ref` (relevant only for IP core variations with the Sync-E support option turned on)
 - `rx_recovered_clk` (relevant only for IP core variations with the Sync-E support option turned on)
- PCS and PMA reset signals:
 - `pcs_tx_arst_ST`
 - `pcs_rx_arst_ST`
 - `pma_arst_ST`
- Ethernet link signals `tx_serial` and `rx_serial`
- PHY output status signals `tx_lanes_stable` and `lanes_deskewed`. Note that input signals with these names are available, and should be supplied by the user logic that implements the PHY layer.
- Signals to connect to the external transceiver reconfiguration controller (relevant for Arria V GZ and Stratix V devices only).
- 40GBASE-KR4 microprocessor interface and reconfiguration interface signals.

Table 3-21: 40-100GbE MAC IP Core Signals Not Visible in 40-100GbE MAC and PHY IP Core

| Signal Name | Direction | Description |
|---|-----------|---------------------------------|
| Transmit Side Signals | | |
| <code>tx_mii_d[<w>*64-1:0]</code> | Output | MAC to PHY connection interface |
| <code>tx_mii_c[<w>*8-1:0]</code> | Output | |
| <code>tx_mii_valid</code> | Output | |
| <code>tx_mii_ready</code> | Input | |
| <code>tx_lanes_stable</code> | Input | |
| Receive Side Signals | | |
| <code>rx_mii_d[<w>*64-1:0]</code> | Input | PHY to MAC connection interface |
| <code>rx_mii_c[<w>*8-1:0]</code> | Input | |
| <code>rx_mii_valid</code> | Input | |
| <code>lanes_deskewed</code> | Input | Lane to lane deskew interface |

Note: When simulating the full design the `lanes_deskewed` input comes from the output of the RX PCS, indicating a fully locked status. To avoid confusion, when simulating the `alt_e40_mac` or `alt_e100_mac` wrapper as the top level, drive the `lanes_deskewed` input and the `tx_mii_ready` input to '1'.

Related Information

- [Lane to Lane Deskew Interface](#) on page 3-43
- [Signals of MAC and PHY Variations Without Adapters](#) on page 3-55
- [Signals of MAC and PHY Variations With Adapters](#) on page 3-66
- [MAC to PHY Connection Interface](#) on page 3-43

Signals of 40-100GbE PHY-Only IP Core Variations

40-100GbE PHY-only IP core variations are the variations that do not include a MAC. The signals in a PHY-only variation with and without adapters are described in the following formats:

- The figures identify the IP core interfaces and the presence or absence of various signals in PHY-only IP core variations.
- Text lists the MAC and PHY IP core variation signals present in the PHY-only variations.
- A table lists the signals visible in the PHY-only variations that are not visible in the MAC and PHY variations.

40-100GbE PHY-only IP core variations are the variations that do not include a MAC. The signals in a PHY-only variation with and without adapters are shown in the following figures in black, purple, or dark gray. Signals shown in blue or in green are not available in the PHY-only variations.

Figure 3-37: Top-Level Signals of the 40-100GbE IP Core with Adapters





Figure 3-38: Top-Level Signals of the 40-100GbE IP Core Without Adapters



The following 40-100GbE MAC and PHY IP core signals are available in PHY-only IP core variations:

- Clock signals:
 - `clk_ref` (relevant only for IP core variations with the Sync-E support option turned off)
 - `tx_clk_ref` (relevant only for IP core variations with the Sync-E support option turned on)
 - `rx_clk_ref` (relevant only for IP core variations with the Sync-E support option turned on)
 - `rx_recovered_clk` (relevant only for IP core variations with the Sync-E support option turned on)
 - `clk_rxmac`
 - `clk_txmac`
- PCS and PMA reset signals:
 - `pcs_tx_arst_ST`
 - `pcs_rx_arst_ST`
 - `pma_arst_ST`
- Control and status interface signals to access PHY component registers.
- Ethernet link signals `tx_serial` and `rx_serial`.
- PHY output status signals `tx_lanes_stable` and `lanes_deskewed`.
- Signals to connect to the external transceiver reconfiguration controller (relevant for Arria V GZ and Stratix V devices only).
- 40GBASE-KR4 microprocessor interface and reconfiguration interface signals.

The remaining 40-100GbE MAC and PHY IP core signals are associated with the MAC and are not available in PHY-only IP core variations.

Table 3-22: 40-100GbE PHY IP Core Signals Not Visible in 40-100GbE MAC and PHY IP Core

| Signal Name | Direction | Description |
|---|-----------|---------------------------------|
| Transmit Side Signals | | |
| <code>tx_mii_d[<w>*64-1:0]</code> | Input | MAC to PHY connection interface |
| <code>tx_mii_c[<w>*8 -1:0]</code> | Input | |
| <code>tx_mii_valid</code> | Input | |
| <code>tx_mii_ready</code> | Output | |
| Receive Side Signals | | |
| <code>rx_mii_d[<w>*64-1:0]</code> | Output | PHY to MAC connection interface |
| <code>rx_mii_c[<w>*8 -1:0]</code> | Output | |
| <code>rx_blocks_valid</code> | Output | |

Related Information

- [Resets](#) on page 3-54
Overview of IP core reset access. Includes list of reset signals and recommended reset sequence.
- [Transceiver PHY Serial Data Interface](#) on page 3-45

- [Lane to Lane Deskew Interface](#) on page 3-43
Describes the `lanes_deskewed` PHY output signal.
- [MAC to PHY Connection Interface](#) on page 3-43
- [External Reconfiguration Controller](#) on page 3-32

Software Interface: Registers

This section provides information about the memory-mapped registers. You access these registers using the IP core control and status interface. The registers use 32-bit addresses; they are not byte addressable.

Write operations to a read-only register field have no effect. Read operations that address a Reserved register return an unspecified constant. Write operations to Reserved registers have no effect. Accesses to registers that do not exist in your IP core variation have an unspecified result.

The following tables list the memory mapped registers for the 40-100GbE IP core and the memory mapped registers for the 40-100GbE IP core example design.

Table 3-23: 40-100GbE IP Core Register Map Overview

Lists the main ranges of the memory mapped registers for the 40-100GbE IP core.

| Word Offset | Register Category |
|-----------------|--|
| 0x0–0x3F | PCS registers |
| 0x40–0x7F | Low Latency PHY IP core registers |
| 0x80–0xFF | 40GBASE-KR4 registers |
| 0x100 and above | 40-100GbE IP core MAC and example design registers |

Table 3-24: 40-100GbE IP Core Address Map

Lists the memory mapped registers for the 40-100GbE IP core. Each register is 32 bits, and the addresses (word offsets) each address a full word.

| Word Offset | Register Description |
|-------------|---|
| 0x000–0x009 | Transceiver PHY control and status registers. Some of these registers are available only for Stratix IV devices. These registers are not the Altera transceiver PHY IP core registers. |
| 0x00A–0x00F | Reserved. |
| 0x010–0x014 | Lock status registers. |
| 0x015–0x016 | Bit error flag registers. |
| 0x017 | PCS hardware error register. |
| 0x018 | BER monitor register. |
| 0x019 | Test mode register. |

| Word Offset | Register Description |
|---|--|
| 0x01A | Test pattern counter register. |
| 0x01B | One of two link fault signaling registers. |
| 0x01C | Reserved. |
| 0x01D | PHY reset register. |
| 0x01E–0x01F | Reserved. |
| 0x020–0x023 | PCS-VLANE registers. |
| 0x024–0x02F | Reserved. |
| 0x030–0x032 | PRBS registers. |
| 0x033–0x03F | Reserved. |
| 0x040–0x07F (Stratix IV only) | Reserved. |
| 0x040–0x07F (Arria V GZ and Stratix V) | For all options except the CAUI-4 configuration: Maps to word addresses 0x040-0x07F in the Low Latency PHY IP core register map. For the CAUI-4 configuration only: Reserved. |
| 0x080–0x0FF | Reserved in non-40GBASE-KR4 IP core variations. |
| 0x080–0x0AF | Reserved in 40GBASE-KR4 IP core variations. |
| 0x0B0–0x0BD | 40GBASE-KR4 top-level and FEC registers. |
| 0x0BE–0x0BF | Reserved. |
| 0x0C0–0x0CC | 40GBASE-KR4 auto-negotiation registers. |
| 0x0CD–0x0CF | Reserved. |
| 0x0D0–0x0EB | 40GBASE-KR4 link training registers. |
| 0x0EC–0x0FF | Reserved. |
| 0x100–0x103 | MAC configuration and filter registers. |
| 0x104–0x10F | Reserved. |
| 0x110–0x117 | Pause registers. |

| Word Offset | Register Description |
|-------------|---|
| 0x118–0x11F | Reserved. |
| 0x120 | MAC hardware error register. |
| 0x121 | MAC reset register. |
| 0x122 | One of two link fault signaling registers. |
| 0x123 | CRC configuration register. |
| 0x124 | Padding configuration register. One of the MAC feature configuration registers. |
| 0x125 | Preamble pass-through configuration register. One of the MAC feature configuration registers. |
| 0x126–0x127 | IPG adjustment registers. Two of the MAC feature configuration registers. |
| 0x140–0x17F | MAC address registers. |
| 0x180–0x1FF | Reserved. |
| 0x200–0x229 | Transmit side statistics registers that are not TX packet statistics registers. |
| 0x22A–0x22D | TX Packet Statistics. |
| 0x22E–0x27F | Reserved. |
| 0x280–0x2A9 | Receive side statistics registers that are not RX packet statistics registers. |
| 0x2AA–0x2B8 | RX packet statistics registers. |
| 0x2B9–0x3FF | Reserved. |
| 0x400–0x423 | 40-100GbE example design registers. Refer to the following table. |
| 0x800-0x9FF | Altera's Low Latency PHY IP core registers, for use with the CAUI-4 configuration only. Offsets 0x800-0x9FF are for CAUI-4 lane 1, offsets 0xA00-0xBFF are for CAUI-4 lane 2, offsets 0xC00-0xDFF are for CAUI-4 lane 3, and offsets 0xE00-0xFFF are for CAUI-4 lane 4. Each lane maps to word addresses 0x000-0x1FF for one low latency PHY transceiver channel. |
| 0xA00-0xBFF | |
| 0xC00-0xDFF | |
| 0xE00-0xFFF | |

Table 3-25: 40-100GbE Example Design Registers

Lists the memory mapped registers for the 40-100GbE IP core example design.

| Word Offset | Register Category |
|-------------|-----------------------------------|
| 0x400–0x403 | PMD registers |
| 0x404–0xF0F | Reserved |
| 0x410–0x413 | MDIO registers |
| 0x414–0x41F | Reserved |
| 0x420–0x423 | 2-wire serial interface registers |

Related Information

- [Control and Status Interface](#) on page 3-51
- [Lock Status Registers](#) on page 3-84
- [Link Fault Signaling Registers](#) on page 3-88
- [MAC and PHY Reset Registers](#) on page 3-89
- [PCS-VLANE Registers](#) on page 3-91
- [PRBS Registers](#) on page 3-92
- [40GBASE-KR4 Registers](#) on page 3-92
- [MAC Configuration and Filter Registers](#) on page 3-99
- [MAC Hardware Error Register](#) on page 3-104
- [CRC Configuration Register](#) on page 3-105
- [PMD Registers](#) on page 3-116
- [MDIO Registers](#) on page 3-117
- [2-Wire Serial Interface Registers](#) on page 3-118
- [Altera Transceiver PHY IP Core User Guide](#)

The CAUI–4 variations of the 40-100GbE IP core use the Low Latency PHY IP core registers at internal offsets 0x000-0x1FF (at IP core register map offsets 0x800–0xFFFF), and the non-CAUI–4 Arria V GZ and Stratix V variations use the Low Latency PHY IP core registers at internal offsets 0x040–0x07F (at IP core register map offsets 0x040–0x07F). Information about this PHY IP core, including loopback configuration, is available in the *Low Latency PHY IP Core* chapter of the *Altera Transceiver PHY IP Core User Guide*.

The 40GBASE-KR4 variations of the 40-100GbE IP core use the 10GBASE-KR PHY IP core PHY registers at internal offsets 0xB0–0xFF (at IP core register map offsets 0xB0–0xFF). Information about this PHY IP core, including register descriptions, is available in the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide*.

40-100GbE IP Core Registers

The following sections describe the registers included in the 40-100GbE IP core.

[Transceiver PHY Control and Status Registers](#) on page 3-80

- [Lock Status Registers](#) on page 3-84
 - [Bit Error Flag Registers](#) on page 3-86
 - [PCS Hardware Error Register](#) on page 3-87
 - [BER Monitor Register](#) on page 3-87
 - [Test Mode Register](#) on page 3-88
 - [Test Pattern Counter Register](#) on page 3-88
 - [Link Fault Signaling Registers](#) on page 3-88
 - [MAC and PHY Reset Registers](#) on page 3-89
 - [PCS-VLANE Registers](#) on page 3-91
 - [PRBS Registers](#) on page 3-92
 - [40GBASE-KR4 Registers](#) on page 3-92
 - [MAC Configuration and Filter Registers](#) on page 3-99
 - [Pause Registers](#) on page 3-102
 - [MAC Hardware Error Register](#) on page 3-104
 - [CRC Configuration Register](#) on page 3-105
 - [MAC Feature Configuration Registers](#) on page 3-105
 - [MAC Address Registers](#) on page 3-107
 - [Statistics Registers](#) on page 3-108
- Related Information**
- [Control and Status Interface](#) on page 3-51

Transceiver PHY Control and Status Registers

The TX serial rate (PCS clock) is based on the input transceiver reference clock and should be precise and stable. The RX serial rate is recovered from the remote system. The RX serial clock typically shows some instability during lock acquisition.

In variations that target a Stratix IV device, the registers in the transceiver PHY provide dynamic access to the analog configuration capability on a per channel (pin) basis. You can also use these registers to place the transceivers in loopback mode for diagnostic or error injection testing. In loopback mode, the TX output connects to the corresponding RX channel.

Table 3-26: Scratch and Clock Registers for Applicable Devices

| Address | Name | Applicable Device(s) | Bit | Description | HW Reset Value | Access |
|-----------------------|-------------|---------------------------------------|--------|---|---|--------|
| 0x000 | PHY_VERSION | Arria V GZ, Stratix IV, and Stratix V | [31:0] | 40GbE/100GbE PHY IP core revision. | 0x00E01310 (40GbE) 0x00DE1310 (100GbE) | R |
| 0x001 | SCRATCH_PHY | Arria V GZ, Stratix IV, and Stratix V | [31:0] | Scratch register available for testing. | 0x00000000 | RW |
| 0x002 ⁽¹²⁾ | CLK_TXS | Arria V GZ, Stratix IV, and Stratix V | [19:0] | TX serial clock rate monitor in KHz. | 0x00000 | R |
| 0x003 ⁽¹²⁾ | CLK_RXS | Arria V GZ, Stratix IV, and Stratix V | [19:0] | RX serial clock rate monitor in KHz. | 0x00000 | R |
| 0x004 | CLK_TXC | Arria V GZ, Stratix IV, and Stratix V | [19:0] | TX core clock rate monitor in KHz. | 0x00000 | R |
| 0x005 | CLK_RXC | Arria V GZ, Stratix IV, and Stratix V | [19:0] | RX core clock rate monitor in KHz. | 0x00000 | R |
| 0x006 | — | Arria V GZ, Stratix IV, and Stratix V | - | Reserved. | — | — |
| 0x007 | — | Arria V GZ and Stratix V | - | Reserved. | — | — |

⁽¹²⁾ When the line rate is 10.3125 Gbps, the frequency of the serial clocks `clk_tx_serial` and `clk_rx_serial` is $10312.5/40 = 257.8125$ MHz.

| Address | Name | Applicable Device(s) | Bit | Description | HW Reset Value | Access |
|---------|----------|---------------------------------------|--------|--|----------------|--------|
| 0x007 | GX_CTRL1 | Stratix IV | [31] | When set, places the transceiver in Internal serial loopback, from TX to RX. | 1'b0 | RW |
| | | | [30] | When asserted, the transceiver channel's analog settings are read on the rising edge of <code>clk_status</code> . | 1'b0 | RW |
| | | | [29] | When asserted, the transceiver channel's analog settings are written on rising edge of <code>clk_status</code> . | 1'b0 | RW |
| | | | [28:4] | Specifies the analog settings to write. Refer to the following table for the register fields which are read on <code>GX_REPLY[24:0]</code> . Bits[28:4] of this register correspond to bits[24:0] of the <code>GX_REPLY</code> register. | 0x000000 | RW |
| | | | [3:0] | Specifies the logical channel to select [0-9]. | 0x0 | RW |
| 0x008 | GX_CTRL2 | Arria V GZ, Stratix IV, and Stratix V | [1] | Specifies bit error inject position 1 on rising edge. | 1'b0 | RW |
| | | | [0] | Specifies bit error inject position 0 on rising edge. | 1'b0 | RW |
| 0x009 | — | Arria V GZ and Stratix V | - | Reserved. | — | — |
| 0x009 | GX_REPLY | Stratix IV | [26] | Reserved. | 0 | R |
| | | | [25] | When asserted, indicates that read data is valid. | 1'b0 | R |
| | | | [24:0] | Contains the analog settings reported in the previous read of <code>GX_CTRL1[28:4]</code> . | 0x020080 | R |

Table 3-27: Stratix IV Transceiver Analog Settings Register GX_CTRL1—Offset 0x007—Bits [28:4]

Describes the fields in bits [28:4] of the analog settings register, GX_CTRL1. The default settings are a good starting point if you are connecting to a PMD module running at 10.3125 Gbps.

| Bits | Description | HW Reset Value | Access |
|---------|---|----------------|--------|
| [28:25] | RX equalization control. The equalizer uses a pass band filter. Specifying a low value passes low frequencies. Specifying a high value passes high frequencies. | 4'b0011 | RW |
| [24:22] | RX DC gain. Sets the equalization DC gain using one of the following settings: <ul style="list-style-type: none"> • 0: 0 dB • 1: 3 dB • 2: 6 dB • 3: 9 dB • 4: 12 dB • 5: 15 dB • 6: 18 dB • 7: 21 dB | 3'b000 | RW |
| [21:17] | Sets the pre-emphasis for TX tap 0. | 5'b10000 | RW |
| [16:12] | Sets the pre-emphasis for TX pre-emphasis tap 1. | 5'b00000 | RW |
| [11:7] | Sets the pre-emphasis for TX pre-emphasis tap 2. | 5'b10000 | RW |
| [6:4] | TX V _{OD} (amplitude). | 3'b010 | RW |

If you encounter signal integrity problems using the default settings in the table, the following procedure might be helpful:

1. Adjust V_{OD}. A value that is too high tends to cause interference with other lanes. You should select the lowest value that functions correctly.
2. Raise pre-emphasis tap 1 slightly. Raising tap 1 tends to help in the case of long trace length or multiple connectors causing signal loss.
3. Adjust the equalization control. This control is roughly analogous to a stereo equalizer. It emphasizes and de-emphasizes portions of the signal by frequency.
4. Repeat these steps, starting at step 1, making minor adjustments to all of the controls while monitoring the error rate. Note that the controls do interact.

Related Information

- [Clocks](#) on page 3-51
For information about required frequencies for the core clocks (`clk_txmac` and `clk_rxmac`) and the transceiver reference clock (`clk_ref`).
- [Altera Transceiver PHY IP Core User Guide](#)
For information about Stratix V and Arria V GZ transceiver PHY IP control and status registers.

Lock Status Registers

The following registers show the lock status of the high speed I/O and RX PCS. `RX_AGGREGATE[0]` aggregates the status of the individual RX PCS channels. When this bit is set to 1, the RX PCS is operating normally. When this bit is set to 0, the other information indicates the cause.

Table 3-28: Lock Status Registers

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--|----------------------|---|----------------|--------|
| 0x010 | IO_LOCKS For all options except the CAUI-4 configuration. | [31:25] (100 GbE) | Reserved. | 0x7f | R |
| | | [24:22] (100 GbE) | When asserted, indicates that the corresponding TX PLL is locked. | 0x0 | R |
| | | [21:12] (100 GbE) | Reserved. | 0x3ff | R |
| | | [11:2] (100 GbE) | When asserted, indicates that the corresponding RX CDR locked. The lowest bit corresponds to lane 0 and so forth. | 0x0 | R |
| | | [31:7] (40 GbE) | Reserved. | 0x1ffffff | R |
| | | [6] (40 GbE) | When asserted, indicates that the TX PLL is locked. | 1'b0 | R |
| | | [5:2] (40 GbE) | When asserted, indicates that the corresponding RX CDR is locked. Bit 2 corresponds to lane 1, bit 3 corresponds to lane 2, bit 4 corresponds to lane 3, and bit 5 corresponds to lane 4. | 0x0 | R |
| | | [1] | When asserted, indicates that the TX interface is ready. | 1'b0 | R |
| | | [0] | When asserted, indicates that the RX interface is ready. | 1'b0 | R |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--|---------|--|----------------|--------|
| 0x010 | IO_LOCKS For the CAUI-4 configuration only. | [31:26] | Reserved. | 0x3f | R |
| | | [25:22] | When asserted, indicates that the corresponding TX PLL is locked. | 0x0 | R |
| | | [21:6] | Reserved. | 0xffff | R |
| | | [5:2] | When asserted, indicates that the corresponding RX CDR is locked. The lowest bit corresponds to lane 0 and so forth. | 1'b0 | R |
| | | [1] | When asserted, indicates that the TX interface is ready. | 1'b0 | R |
| | | [0] | When asserted, indicates that the RX interface is ready. | 1'b0 | R |
| 0x011 | LOCKED_TIME | [31:0] | Counts the RX continuous up time in seconds. The counter will roll over in approximately 126 years. | 0x00000000 | R |
| 0x012 | WORD_LOCKS | [19:0] | When asserted, indicates that the physical channel has identified 66 bit block boundaries in the serial data stream. | 0x00000000 | R |
| 0x013 | AM_LOCKS | [19:0] | When asserted, indicates that the physical channel has identified virtual lane alignment markers in the data stream. | 0x00000000 | R |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--------------|-----|--|----------------|--------|
| 0x014 | RX_AGGREGATE | [4] | When asserted, indicates a change in PCS-VLANE permutation. This status bit clears on read. | 1'b0 | R |
| | | [3] | When asserted, indicates a change in lanes deskewed status. This status bit clears on read. | 1'b0 | R |
| | | [2] | When asserted, indicates a change in PCS-VLANE tag drop position. This status bit clears on read. | 1'b0 | R |
| | | [1] | When asserted, indicates that all lanes are locked and lane-to-lane deskew is complete so that the 40-100GbE IP core is in operating normally. | 1'b0 | R |
| | | [0] | When asserted, indicates that all lanes are word and alignment marker locked. | 1b'0 | R |

Bit Error Flag Registers

Bit errors occur naturally from time to time on high speed serial links. The higher level Ethernet protocol includes mechanisms to respond to errors and tally them appropriately. These lower level flags are useful for tracking errors in a physical link and computing error rates.

Table 3-29: Bit Error Flag Registers

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-------------|--------|--|----------------|--------|
| 0x015 | FRAMING_ERR | [31] | When asserted, indicates that a framing error has occurred on any lane. This status bit clears on read. | 1'b0 | R |
| | | [19:0] | When asserted, indicates that a framing error has occurred on the corresponding physical lane. This status bit clears on read. | 0x0000 0 | R |
| 0x016 | BIP_ERR | [19:0] | When asserted, indicates that a BIP (lane parity) error has occurred on the corresponding physical lane. This status bit clears on read. | 0x0000 0 | R |

PCS Hardware Error Register

Table 3-30: PCS Hardware Error Register

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|-----|---|----------------|--------|
| 0x017 | PCS_HW_ERR | [8] | When asserted, indicates a RX phase compensation error. | 1'b0 | R |
| | | [7] | When asserted, indicates a TX phase compensation error. | 1'b0 | R |
| | | [6] | When asserted, indicates a deskew failure suggesting a problem with the remote transmitter. | 1'b0 | R |
| | | [5] | When asserted, indicates a parity error in the RX lanes (RXL) section. | 1'b0 | R |
| | | [4] | When asserted, indicates a RX deskew FIFO buffer underflow error. | 1'b0 | R |
| | | [3] | When asserted, indicates a RX deskew FIFO buffer overflow error. | 1'b0 | R |
| | | [2] | When asserted, indicates a parity error in the TX lanes (TXL) section. | 1'b0 | R |
| | | [1] | When asserted, indicates a TX deskew FIFO buffer underflow error. | 1'b0 | R |
| | | [0] | When asserted, indicates a TX deskew FIFO buffer overflow error. | 1'b0 | R |

BER Monitor Register

Table 3-31: BER Monitor Register—Offset 0x018

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-------------|-----|---|----------------|--------|
| 0x018 | BER_MONITOR | [1] | This bit enables the BER monitor. | 1'b0 | RW |
| | | [0] | When asserted by the BER monitor block, this bit indicates the PCS is recording a high BER. | 1'b0 | R |

Test Mode Register

Table 3-32: Test Mode Register—Offset 0x019

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-----------|-----|--|----------------|--------|
| 0x019 | TEST_MODE | [2] | This bit clears the test-pattern counter (TEST_PATTERN_COUNTER register at offset 0x1A). | 1'b0 | RW |
| | | [1] | This bit enables RX test mode. | 1'b0 | RW |
| | | [0] | This bit enables TX test mode. | 1'b0 | RW |

Related Information

[PCS Test Pattern Generation and Test Pattern Check](#) on page 3-44

Describes test pattern usage.

Test Pattern Counter Register

Table 3-33: Test Pattern Counter Register—Offset 0x1A

Unlike other statistics counters, the RX TEST_PATTERN_COUNTER is 32 bits and saturates.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|----------------------|--------|---|----------------|--------|
| 0x1a | TEST_PATTERN_COUNTER | [31:0] | This register is the test pattern error counter. The counter saturates at 0xffffffff. | 32'b0 | R |

Related Information

[PCS Test Pattern Generation and Test Pattern Check](#) on page 3-44

Describes test pattern usage.

Link Fault Signaling Registers

Table 3-34: Link Fault Sequence Enable Register—Offset 0x01B

| Address | Name | Bit | Description | HW Value | Access |
|---------|----------------------------|-----|---|----------|--------|
| 0x1b | Enable Link Fault Sequence | [0] | When asserted, the PCS generates remote fault sequence if conditions are met. | 1'b0 | RW |

Table 3-35: Link Fault Signaling Configuration Register—Offset 0x122

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--|-------|---|----------------|--------|
| 0x122 | MAC/RS link fault sequence configuration | [5] | The remote fault status register. | 1'b0 | R |
| | | [4] | The local fault status register. | 1'b0 | R |
| | | [3:2] | The remote fault configuration register. Possible configurations include: <ul style="list-style-type: none"> 2'b01: sends idle frames when remote fault is received 2'b11: sends remote fault sequence when remote fault is received 2'bx0: sends normal traffic when remote fault is received | 2'b00 | RW |
| | | [1:0] | The local fault configuration register. Possible configurations include: <ul style="list-style-type: none"> 2'b01: sends idle frames when local fault is received 2'b11: sends remote fault sequence when local fault is received 2'bx0: sends normal traffic when local fault is received | 2'b00 | RW |

Related Information

[Link Fault Signaling Interface](#) on page 3-37

Describes how the IP core uses the register values.

MAC and PHY Reset Registers

The following registers control the 40-100GbE MAC and PHY resets. Writing a 1'b1 to any of the reset register fields initiates the corresponding reset sequence.

Note: Altera recommends that you reset the MAC and PHY using the following register reset sequence:

1. Write the value of 0x7 (all ones) to the PHY reset register at offset 0x1D.
2. Write the value of 0x3 (all ones) to the MAC reset register at offset 0x121.
3. Write the value of 0x0 (all zeros) to the MAC reset register at offset 0x121.
4. Write the value of 0x0 (all zeros) to the PHY reset register at offset 0x1D.

Table 3-36: MAC Reset Register

Writing a 1'b1 to any of the reset register fields initiates the corresponding reset sequence.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-----------|-----|----------------------------|----------------|--------|
| 0x121 | MAC Reset | [1] | The MAC TX reset register. | 1'b0 | RW |
| | | [0] | The MAC RX reset register. | 1'b0 | RW |

Table 3-37: PHY Reset Register

Writing a 1'b1 to any of the reset register fields initiates the corresponding reset sequence.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-----------|-----|----------------------------|----------------|--------|
| 0x01d | PHY reset | [2] | The PMA reset register. | 1'b0 | RW |
| | | [1] | The PCS TX reset register. | 1'b0 | RW |
| | | [0] | The PCS RX reset register. | 1'b0 | RW |

Related Information

[Resets](#) on page 3-54

PCS-VLANE Registers

Table 3-38: PCS-VLANE Registers

When the RX PCS is properly locked, the PCS-VLANE registers contain a permutation of the numbers [0–19] stored in 5-bit values. You may need this information to isolate a problem with the remote transmitter. Any problem with the permutation, for example two lanes with the same number, stops the reception of valid data.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|---------|------------------------------------|----------------|--------|
| 0x020 | PCS VLANE0 | [29:25] | Virtual index for physical lane 5 | 5h'00 | R |
| | | [24:20] | Virtual index for physical lane 4 | 5h'00 | R |
| | | [19:15] | Virtual index for physical lane 3 | 5h'00 | R |
| | | [14:10] | Virtual index for physical lane 2 | 5h'00 | R |
| | | [9:5] | Virtual index for physical lane 1 | 5h'00 | R |
| | | [4:0] | Virtual index for physical lane 0 | 5h'00 | R |
| 0x021 | PCS VLANE1 | [29:25] | Virtual index for physical lane 11 | 5h'00 | R |
| | | [24:20] | Virtual index for physical lane 10 | 5h'00 | R |
| | | [19:15] | Virtual index for physical lane 9 | 5h'00 | R |
| | | [14:10] | Virtual index for physical lane 8 | 5h'00 | R |
| | | [9:5] | Virtual index for physical lane 7 | 5h'00 | R |
| | | [4:0] | Virtual index for physical lane 6 | 5h'00 | R |
| 0x022 | PCS VLANE2 | [29:25] | Virtual index for physical lane 17 | 5h'00 | R |
| | | [24:20] | Virtual index for physical lane 16 | 5h'00 | R |
| | | [19:15] | Virtual index for physical lane 15 | 5h'00 | R |
| | | [14:10] | Virtual index for physical lane 14 | 5h'00 | R |
| | | [9:5] | Virtual index for physical lane 13 | 5h'00 | R |
| | | [4:0] | Virtual index for physical lane 12 | 5h'00 | R |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|-------|------------------------------------|----------------|--------|
| 0x023 | PCS_VLANE3 | [9:5] | Virtual index for physical lane 19 | 5h'00 | R |
| | | [4:0] | Virtual index for physical lane 18 | 5h'00 | R |

PRBS Registers

The PRBS feature operates on a per virtual lane basis. The PRBS streams are bit interleaved to form 10 Gbps lanes. PRBS transmissions are unframed.

Table 3-39: PRBS Registers

There is some deviation among manufacturers regarding the exact PRBS implementation on high rate Ethernet equipment. If the PRBS is working properly in loopback but consistently wrong with a remote device, Altera recommends that you confirm that the remote device is operating as described here.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--------------|--------|--|----------------|--------|
| 0x030 | PRBS_CTRL | [4] | When set to 1, selects PRBS-9 for the RX lane. When set to 0, selects PRBS-31 for the RX lane. | 1'b0 | RW |
| | | [3] | When set to 1, selects PRBS-9 for the TX lane. When set to 0, selects PRBS-31 for the TX lane. | 1'b0 | RW |
| | | [2] | When set to 1, enables the RX PRBS checker. | 1'b0 | RW |
| | | [1] | When set to 1, enables the TX PRBS transmitter. | 1'b0 | RW |
| 0x031 | PRBS_ERR_INJ | [19:0] | When set to 1, injects an error in the corresponding lane. This register is rising edge triggered. Write 0 to clear. | 20'b0 | RW |
| 0x032 | PRBS_EFLAGS | [19:0] | When set to 1, indicates a PRBS error in the corresponding PCS-VLANE. Clears on read. | 20'b0 | R |

40GBASE-KR4 Registers

Most 40GBASE-KR4 registers are 10GBASE-KR PHY registers of the 10GBASE-KR PHY IP core, documented in the [Altera Transceiver PHY IP Core User Guide](#). The register offsets are identical in the 40GBASE-KR4 variations of the 40-100GbE IP core. However, the 40GBASE-KR4 variations of the 40-100GbE IP core have additional 40GBASE-KR4 related registers and register fields.

For your convenience, the 40-100GbE IP core user guide includes an appendix with the 10GBASE-KR PHY register descriptions: [10GBASE-KR Registers](#).

Table 3-40: 40-100GbE IP Core 40GBASE-KR4 Registers and Register Fields Not in 10GBASE-KR PHY IP Core

Documents the differences between the 10GBASE-KR PHY register definitions and the 40GBASE-KR4 registers of the 40-100GbE IP core. All 10GBASE-KR PHY registers and register fields not listed in the table are available in the 40GBASE-KR4 variations of the 40-100GbE IP core.

The 10GBASE-KR PHY register listings are available in the [10GBASE-KR Registers](#) appendix and in the *10GBASE-KR PHY Register Definitions* section of the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide*. The information in these two sources should be identical. Refer to [10GBASE-KR Registers](#) for details.

Where the 10GBASE-KR PHY register definitions list 10GBASE-R, substitute 40GBASE-KR4 with auto-negotiation and link training both turned off, and where the 10GBASE-KR PHY register definitions list 10GBASE-KR (except in the description of 0xCB[24:0]), substitute 40GBASE-KR4. Where a register field description in the 10GBASE-KR PHY register definitions refers to link training or FEC in the single-lane 10GBASE-KR PHY IP core, substitute link training or FEC on Lane 0 of the 40GBASE-KR4 IP core variation. Where a register field description in the 10GBASE-KR PHY register definitions refers to the Transceiver Reconfiguration IP core, substitute the reconfiguration bundle.

To modify a field value in any 40GBASE-KR4 specific register, whether an underlying 10GBASE-KR PHY IP core register or one of the registers defined in this table, you must perform a read-modify-write operation to ensure you do not modify the values of any other fields in the register.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|---|---------|--|----------------|--------|
| 0x0B0 | Force Negotiate to FEC Mode | [19] | When set to 1, forces the IP core to use FEC mode regardless of the auto-negotiation result. You must write the value of 1 to 0xB0[0] (reset the sequencer) for this override to take effect. | 1'b0 | RW |
| | FEC Block Lock | [23:20] | FEC Block Lock for lanes [3:0]: bit [20] is FEC block lock for lane 0, bit [21] is FEC block lock for lane 1, bit [22] is FEC block lock for lane 2, and bit [23] is FEC block lock for lane 3. | 4'b0 | RO |
| 0x0B5 | Register 0xB2 refers to Lane 0. This register is the equivalent of register 0xB2 for Lane 1. (Refer to 10GBASE-KR PHY Register Definitions). | | | | RW |
| 0x0B6 | KR4 FEC Corrected Blocks, Lane 1 | [31:0] | Maintains count of corrected FEC blocks on Lane 1, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.1 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0B7 | KR4 FEC Uncorrected Blocks, Lane 1 | [31:0] | Maintains count of uncorrected (uncorrectable) FEC blocks on Lane 1, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.2 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0B8 | This register is the equivalent of register 0xB2 for Lane 2. (Refer to 10GBASE-KR PHY Register Definitions). | | | | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--|---------|---|----------------|--------|
| 0x0B9 | KR4 FEC Corrected Blocks, Lane 2 | [31:0] | Maintains count of corrected FEC blocks on Lane 2, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.1 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0BA | KR4 FEC Uncorrected Blocks, Lane 2 | [31:0] | Maintains count of uncorrected (uncorrectable) FEC blocks on Lane 2, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.2 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0BB | This register is the equivalent of register 0xB2 for Lane 3. (Refer to 10GBASE-KR PHY Register Definitions). | | | | RW |
| 0x0BC | KR4 FEC Corrected Blocks, Lane 3 | [31:0] | Maintains count of corrected FEC blocks on Lane 3, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.1 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0BD | KR4 FEC Uncorrected Blocks, Lane 3 | [31:0] | Maintains count of uncorrected FEC blocks on Lane 3, saturating (not rolling over) at $2^{32}-1$. Resets to 0 when read. Refer to <i>Clause 74.8.4.2 of IEEE Std 802.3ap-2007</i> . | 32'b0 | RO |
| 0x0C0 | Override AN Channel Enable | [6] | Overrides the auto-negotiation master channel that you set with the Auto-Negotiation Master parameter, setting the new master channel according to the value in register 0xCC[3:0]. While 0x0C0[6] has the value of 1, the channel encoded in 0xCC[3:0] is the master channel. While 0x0C0[6] has the value of 0, the master channel is the channel that you set with the Auto-Negotiation Master parameter. | 1'b0 | RW |
| 0x0CB | AN LP ADV FEC_F[1:0] | [26:25] | Received FEC ability bits. Bit [26] is FEC requested and bit [25] is FEC ability. FEC (F0:F1) is encoded in bits D46:D47 of the base Link Codeword in Clause 73 AN. F0 is FEC ability and F1 is FEC requested. Refer to <i>Clause 73.6.5 of IEEE Std 802.3ap-2007</i> . | 2'b0 | RO |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-------------------------------|-------|---|----------------|--------|
| 0x0CC | Override AN Channel Select | [3:0] | <p>If you set the value of the Override AN Channel Enable register field (0xC0[6]) to the value of 1, then while 0xC0[6] has the value of 1, the value in this register field (0xCC[3:0]) overrides the master channel you set with the Auto-Negotiation Master parameter.</p> <p>This register field has the following valid values:</p> <ul style="list-style-type: none"> 4'b0001: Selects Lane 0 4'b0010: Selects Lane 1 4'b0100: Selects Lane 2 4'b1000: Selects Lane 3 <p>All other values are invalid. The new master channel is encoded with one-hot encoding.</p> | 4'b0 | RW |
| 0x0D1 | Restart Link training, Lane 1 | [1] | <p>When set to 1, resets the 40GBASE-KR4 start-up protocol. When set to 0, continues normal operation. This bit self clears. Refer to the state variable <code>mr_restart_training</code> as defined in <i>Clause 72.6.10.3.1</i> and 10GBASE-KR PMD control register bit (1.150.0) in <i>IEEE Std 802.3ap-2007</i>.</p> <p>Register bit 0xD1[0] refers to Lane 0. This bit is the equivalent of register 0xD1[0] for Lane 1. (Refer to 10GBASE-KR PHY Register Definitions).</p> | 1'b0 | RW |
| | Restart Link training, Lane 2 | [2] | <p>This bit is the equivalent of register 0xD1[0] for Lane 2. (Refer to 10GBASE-KR PHY Register Definitions).</p> | 1'b0 | RW |
| | Restart Link training, Lane 3 | [3] | <p>This bit is the equivalent of register 0xD0[1] for Lane 3. (Refer to 10GBASE-KR PHY Register Definitions).</p> | 1'b0 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-----------------------------|-----|---|----------------|--------|
| 0x0D1 | Updated TX Coef new, Lane 1 | [5] | <p>When set to 1, indicates that new link partner coefficients are available to send. The LT logic starts sending the new values set in 0xD4[7:0] to the remote device. When set to 0, continues normal operation. This bit self clears.</p> <p>This override of normal operation can only occur if 0xD0[16] (Ovrld LP Coef enable) has the value of 1. If 0xD0[16] has the value of 0, this register field (0xD1[5]) has no effect.</p> <p>Register bit 0xD1[4] refers to Lane 0. This bit is the equivalent of register 0xD1[4] for Lane 1. (Refer to 10GBASE-KR PHY Register Definitions).</p> | 1'b0 | RW |
| | Updated TX Coef new, Lane 2 | [6] | This bit is the equivalent of register 0xD1[5] for Lane 2. | 1'b0 | RW |
| | Updated TX Coef new, Lane 3 | [7] | This bit is the equivalent of register 0xD1[5] for Lane 3. | 1'b0 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-----------------------------|------|---|----------------|--------|
| 0x0D1 | Updated RX Coef new, Lane 1 | [9] | <p>When set to 1, indicates that new local device coefficients are available for Lane 1. The LT logic changes the local TX equalizer coefficients as specified in 0xE1[23:16]. When set to 0, continues normal operation. This bit self clears.</p> <p>This override of normal operation can only occur if 0xD0[17] (Ovride Local RX Coef enable) has the value of 1. If 0xD0[17] has the value of 0, this register field (0xD1[9]) has no effect.</p> <p>Register bit 0xD1[8] refers to Lane 0. This bit is the equivalent of register 0xD1[8] for Lane 1. (Refer to 10GBASE-KR PHY Register Definitions).</p> | 1'b0 | RW |
| | Updated RX Coef new, Lane 2 | [10] | <p>When set to 1, indicates that new local device coefficients are available for Lane 2. The LT logic changes the local TX equalizer coefficients as specified in 0xE5[23:16]. When set to 0, continues normal operation. This bit self clears.</p> <p>This override of normal operation can only occur if 0xD0[17] (Ovride Local RX Coef enable) has the value of 1.</p> <p>This bit is the equivalent of register 0xD1[9] for Lane 2.</p> | 1'b0 | RW |
| | Updated RX Coef new, Lane 3 | [11] | <p>When set to 1, indicates that new local device coefficients are available for lane 3. The LT logic changes the local TX equalizer coefficients as specified in 0xE9[23:16]. When set to 0, continues normal operation. This bit self clears.</p> <p>This override of normal operation can only occur if 0xD0[17] (Ovride Local RX Coef enable) has the value of 1.</p> <p>This bit is the equivalent of register 0xD1[9] for Lane 3.</p> | 1'b0 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------|---------|--|----------------|--------|
| 0xD2 | | [15:8] | Register bits 0xD2[7:0] refer to Lane 0. These bits are the equivalent of 0xD2[7:0] for Lane 1. (Refer to 10GBASE-KR PHY Register Definitions). For Link Training Frame lock Error, Lane 1, if the tap settings specified by the fields of 0xE2 are the same as the initial parameter value, the frame lock error was unrecoverable. | | RO |
| | | [23:16] | These bits are the equivalent of 0xD2[7:0] for Lane 2. (Refer to 10GBASE-KR PHY Register Definitions). For Link Training Frame lock Error, Lane 2, if the tap settings specified by the fields of 0xE6 are the same as the initial parameter value, the frame lock error was unrecoverable. | | RO |
| | | [31:24] | These bits are the equivalent of 0xD2[7:0] for Lane 3. (Refer to 10GBASE-KR PHY Register Definitions). For Link Training Frame lock Error, Lane 3, if the tap settings specified by the fields of 0xEA are the same as the initial parameter value, the frame lock error was unrecoverable. | | RO |
| 0xE0 | | | Register 0xD3 refers to Lane 0. This register, register 0xE0, is the equivalent of register 0xD3 for Lane 1 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RW |
| 0xE1 | | | Register 0xD4 refers to Lane 0. This register, register 0xE1, is the equivalent of register 0xD4 for Lane 1 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RW |
| 0xE2 | | | Register 0xD5 refers to Lane 0. This register, register 0xE2, is the equivalent of register 0xD5 for Lane 1 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RO |
| 0xE3 | | | Register 0xD6 refers to Lane 0. This register, register 0xE3, is the equivalent of register 0xD6 for Lane 1 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RW |
| 0xE4 | | | This register is the equivalent of register 0xD3 for Lane 2 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RW |
| 0xE5 | | | This register is the equivalent of register 0xD4 for Lane 2 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | R / RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------|-----|---|----------------|--------|
| 0xE6 | | | This register is the equivalent of register 0xD5 for Lane 2 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RO |
| 0xE7 | | | This register is the equivalent of register 0xD6 for Lane 2 link training. (Refer to 10GBASE-KR PHY Register Definitions). | | RW |
| 0xE8 | | | This register is the equivalent of register 0xD3 for Lane 3 link training. | | RW |
| 0xE9 | | | This register is the equivalent of register 0xD4 for Lane 3 link training. | | R / RW |
| 0xEA | | | This register is the equivalent of register 0xD5 for Lane 3 link training. | | RO |
| 0xEB | | | This register is the equivalent of register 0xD6 for Lane 3 link training. | | RW |

Related Information

[Altera Transceiver PHY IP Core User Guide](#)

The 40GBASE-KR4 variations of the 40-100GbE IP core use the 10GBASE-KR PHY IP core PHY registers at internal offsets 0xB0–0xFF (at IP core register map offsets 0xB0–0xFF), in addition to the registers listed in this section. Information about this PHY IP core, including register descriptions, is available in the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide*.

MAC Configuration and Filter Registers

Table 3-41: General MAC Control Registers

Describes the MAC configuration and filter registers.

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-------------|--------|---|---|--------|
| 0x100 | MAC_VERSION | [31:0] | 40GbE/100GbE MAC IP core revision. | 0x00DF1310 (40GbE) 0x00DD1310 (100GbE) | RO |
| 0x101 | SCRATCH_MAC | [31:0] | Scratch register available for testing. | 0x00000000 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|--------------------|-------|---|----------------|--------|
| 0x102 | MAC_CMD_ config | [4] | When set to 1, the transmit CRC FIFO is cleared. Used for hardware diagnostics. Not required for normal operation. To clear this register bit, write a 0. | 1'b0 | RW |
| | | [3] | When set to 1, the statistics counters are reset. This register is self-clearing. | 1'b0 | RW |
| | | [2] | When set to 1, statistics collection is paused. The underlying counters continue to operate, but the readable values reflect a snapshot at the time the pause flag was activated. Write a 0 to release. | 1'b0 | RW |
| | | [1:0] | Allows you to override normal transmission for hardware diagnostic purposes: The following patterns are defined: <ul style="list-style-type: none"> • 2'b00 / 2'b10—normal operation (default) • 2'b01—Send a small repeating loop of random content frames • 2'b11—Send idles | 2'b00 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|----------------|--------|---|----------------|--------|
| 0x103 | RX_FILTER_CTRL | [17:8] | <p>Packet length limit in 16-byte words, for the 40GbE IP core. The IP core discards packets based on this length limit: frames of size within 20 bytes over the length limit may or may not be dropped, but oversized frames that are 20 or more bytes over the length limit are always dropped.</p> <p>The IP core records packets in the too long counter if they have length greater than the length specified in this register field.</p> | 0x280 | RW |
| | | [15:8] | <p>Packet length limit in 40-byte words, for the 100GbE IP core. The IP core discards packets based on this length limit: frames of size within 44 bytes over the length limit may or may not be dropped, but oversized frames that are 44 or more bytes over the length limit are always dropped.</p> <p>The IP core records packets in the too long counter if they have length greater than the length specified in this register field.</p> | 0xf0 | RW |
| | | [5] | When set to 1, non-pause control frames are removed; this filtering is only enabled when RX_FILTER_CTRL bit [0] is set to 0. | 1 | RW |
| | | [4] | When set to 1, pause frames are removed; this filtering is only enabled when RX_FILTER_CTRL bit [0] is set to 0. | 1 | RW |
| | | [3] | When set to 1, runt frames are removed regardless of the filtering enable bit 0x103[0]. The IP core identifies frames of length eight bytes or less as decoding errors rather than as runt frames. | 1b'1 | RW |
| | | [2] | When set to 1, the filter discard packets which do not target a matching destination address. | 1b'0 | RW |
| | | [1] | When set to 1, the filter discards packets with FCS errors. | 1b'1 | RW |
| | | [0] | When set to 0 enables filtering. When set to 1, accepts all traffic (promiscuous mode). However, when bit [3] is set, runts are removed regardless of the value of bit [0]. | 1b'0 | RW |

Related Information

- **40-100GbE IP Core RX Filtering** on page 3-21
Describes how the IP core interprets the packet length limits in `RX_FILTER_CTRL[17:8]` (for 100GbE IP cores) and `RX_FILTER_CTRL[15:8]` (for 40GbE IP cores).
- **Pause Control Frame and Non-Pause Control Frame Filtering and Forwarding** on page 3-36
For more information about the IP core behavior based on the `RX_FILTER_CTRL` register bits [0], [4], and [5].
- **40-100GbE IP Core Modes of Operation** on page 3-37
Overview of filtering status.

Pause Registers

The pause registers implement the pause functionality defined in the *IEEE 802.3ba-2010 100G Ethernet Standard*. You can program the pause registers to control the insertion and decoding of pause frames, to help reduce traffic in congested networks.

Alternatively, you can use the IP core pause signals.

Table 3-42: Pause Registers

| Addr | Name | Bit | Description | HW Reset Value | Access |
|-------|-------------------------|--------|--|------------------------------------|--------|
| 0x110 | RECEIVE_PAUSE_STATUS | [16] | When set to 1, indicates that a pause is in progress. | 1'b0 | RO |
| | | [15:0] | The time value for the pause. Reading this field locks the pause source address. | 0x140(40GbE) 0x01A0 (100GbE) | RO |
| 0x111 | RECEIVE_SOURCE_ADDR_LSB | [31:0] | Received pause source address lsb. | 0x00000000 | RO |
| 0x112 | RECEIVE_SOURCE_ADDR_MSB | [31:0] | Received pause source address msb. | 0x00000000 | RO |

| Addr | Name | Bit | Description | HW Reset Value | Access |
|-------|-----------------------|--------|--|----------------|--------|
| 0x113 | RECEIVE_PAUSE_CONTROL | [9] | When set to 1, enables unicast pause receive. | 1'b0 | RW |
| | | [8] | When set to 1, enables multicast pause receive. | 1'b0 | RW |
| | | [7:0] | Pause quantum time configuration as follows: <ul style="list-style-type: none"> 1 pause_quanta = 512 bit times For 40Gbps, 1 pause_quanta is 12.8 ns For 100Gbps, 1 pause_quanta is 5.12 ns pause_quantum_delta = 256 × Tclk / Tpause_quanta For example: 40 Gbps, clk =315 MHz, Tclk = 3.2ns, pause_quantum_delta = round(256*(3.2/12.8)) = 64. | 0x00000000 | RW |
| 0x114 | INSERT_PAUSE_CONTROL | [16] | When set to 1, sends a multicast pause request. When set to 0, sends a unicast pause request. | 1'b0 | RW |
| | | [15:0] | Specifies the pause time. A non-zero value specifies an XON. Zero specifies XOFF. | 0x0000 | RW |
| 0x115 | TX_PAUSE_DST_ADDR_LSB | [31:0] | Destination address LSB. | 0x00000000 | RW |
| 0x116 | TX_PAUSE_DST_ADDR_MSB | [31:0] | Destination address MSB. | 0x00000000 | RW |
| 0x117 | INSERT_PAUSE | [31:0] | Any write to this address triggers a pause packet insertion into the TX data stream. Other pause registers, described in this table, specify the properties of this pause packet. | 0x00000000 | W |

Related Information

[Pause Control and Generation Interface](#) on page 3-35

Describes the pause signals available in 40-100GbE IP core variations that include a MAC with a custom streaming client interface.

MAC Hardware Error Register

Table 3-43: MAC Hardware Error Register

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|-----|---|----------------|--------|
| 0x120 | MAC_HW_ERR | [6] | When asserted, indicates a parity error in the Drop on Error storage RAM section. | 1'b0 | R |
| | | [5] | When asserted, indicates a parity error in the DOE command FIFO buffer section. | 1'b0 | R |
| | | [4] | When asserted, indicates a DOE_COMMAND FIFO buffer overflow error. | 1'b0 | R |
| | | [3] | When asserted, indicates a parity error in the TX CRC read-ram (TXC) section. | 1'b0 | R |
| | | [2] | When asserted, indicates a parity error in the TX CRC write-ram (TXW) section. | 1'b0 | R |
| | | [1] | When asserted, indicates a parity error in the RX inspector (RXI) section. | 1'b0 | R |
| | | [0] | When asserted, indicates a TX CRC FIFO buffer overflow error. | 1'b0 | R |

CRC Configuration Register

Table 3-44: CRC Configuration Register

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|-----|--|----------------|--------|
| 0x123 | CRC_CONFIG | [1] | <p>The RX CRC configuration register. Possible configurations include:</p> <ul style="list-style-type: none"> 1'b0: removes RX CRC 1'b1: retains RX CRC <p>Turning on PAD_CONFIG (register 0x124 bit [0]) prevents this bit (CRC_CONFIG bit [1]) from having any effect on padded packets.</p> | 1'b0 | RW |
| | | [0] | <p>The TX CRC configuration register. Possible configurations include:</p> <ul style="list-style-type: none"> 1'b0: enables TX CRC insertion 1'b1: disables TX CRC insertion | 1'b0 | RW |

MAC Feature Configuration Registers

The MAC feature configuration registers control the following additional MAC features in the RX and TX datapaths:

- Padding configuration
- Preamble pass-through feature
- Inter-packet gap adjustment

Table 3-45: MAC Feature Configuration Registers

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|-----|--|----------------|--------|
| 0x124 | PAD_CONFIG | [0] | When set to 1, the IP core removes padding from frames on the RX datapath. When you set this bit to the value of 1, the IP core disables the CRC_CONFIG (register 0x123 bit [1]) from having any effect on padded packets. | 0 | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|-------------------------------------|--------|--|----------------|--------|
| 0x125 | Preamble Pass-Through Configuration | [1] | Enable TX preamble pass-through. This bit has the following valid values: <ul style="list-style-type: none"> 1'b0: Disable TX preamble pass-through. 1'b1: Enable TX preamble pass-through. | 1'b0 | RW |
| | | [0] | Enable RX preamble pass-through. This bit has the following valid values: <ul style="list-style-type: none"> 1'b0: Disable RX preamble pass-through. 1'b1: Enable RX preamble pass-through. | 1'b0 | RW |
| 0x126 | IPG_DEL_PERIOD | [11:0] | <p>If IPG_DEL_ENABLE[0] has the value of one, the IPG_DEL_PERIOD register controls the frequency of Idle byte deletion from the default inter-packet gap Idle sequence.</p> <p>If IPG_DEL_ENABLE[0] has the value of one, and IPG_DEL_PERIOD[11:0] has the value of N, the MAC deletes extra Idle bytes from the inter-packet gap in the ratio of one Idle byte for every N 8-byte blocks.</p> <p>At the default value of 2048, this deletion compensates for Alignment Marker insertion by the PHY (approximately 61 ppm). You can modify the value of this register field to increase bandwidth or adjust for clock compensation, as needed.</p> | 12'd20 48 | RW |
| 0x127 | IPG_DEL_ENABLE | [0] | <p>Enables the inter-packet gap adjustment feature. If this bit has the value of 1, the MAC deletes inter-packet gap Idle bytes at the frequency specified in the IPG_DEL_PERIOD register.</p> <p>If this bit has the value of 0, the MAC does not implement the inter-packet gap adjustment feature, but instead, generates the usual inter-packet gap, to an average IPG of 12 bytes.</p> | 1'b1 | RW |

Related Information

- [RX Automatic Pad Removal Control](#) on page 3-22
Describes the effect of PAD_CONFIG[0] on IP core behavior.
- [Preamble, Start, and SFD Insertion](#) on page 3-4
Information about the TX preamble pass-through feature and the requirements for the preamble the client application provides to the IP core.

- **40-100GbE IP Core Preamble Processing** on page 3-21
Describes the RX preamble pass-through feature and the effect of bit [0] of the `Preamble Pass-Through Configuration` register on IP core behavior.
- **Inter-Packet Gap Generation and Insertion** on page 3-5
Information about the effects of the `IPG_DEL_PERIOD` register and the `IPG_DEL_ENABLE` register on IP core behavior.

MAC Address Registers

All MAC addresses are stored in 2 words of 24 bits (3 bytes) in natural reading order. For example, an Altera Device ID of 00-07-ED-11-22-33 is stored with the more significant word set to 24'h0007ED and the least significant word set to 24'h112233.

The following table describes the MAC address registers. Additional information about the `MADDR_CTRL` register follows the table.

Table 3-46: MAC Address Registers

| Address | Name | Bit | Description | HW Reset Value | Access |
|---------|------------|--------|---|----------------|--------|
| 0x140 | MADDR_CTRL | [31] | When set to 1, the source address is Inserted in TX packets. | 1'b0 | RW |
| | | [30] | When set to 1, enables the IP core to check the destination address in RX packets. When set to 0, all RX packets are considered to have a matching destination address. If you enable destination address filtering, by setting bit 0 of the <code>RX_FILTER_CTRL</code> register (offset 0x103) to the value of 0, but this bit, bit 30 of the <code>MADDR_CTRL</code> register, is also set to the value of 0, the IP core considers all destination addresses to match. | 1'b0 | RW |
| | | [29:1] | Reserved. | 0x0000000 0 | |
| | | [0] | When set to 1, enables destination address 0x160-0x161 checking. If this bit has the value of 0 and bit 30 has the value of 1, the IP core considers all destination addresses to not match. | 1'b0 | RW |
| 0x141 | SRC_AD_LO | [23:0] | Source address (lower 24 bits). | 0x00FF1234 | RW |
| 0x142 | SRC_AD_HI | [23:0] | Source address (upper 24 bits). | 0x000007ED | RW |

| Address | Name | Bit | Description | HW Reset Value | Access |
|-------------|------------|--------|--|----------------|--------|
| 0x160 | DST_AD0_LO | [23:0] | Destination address 0 (lower 24 bits). | 0x00000000 | RW |
| 0x161 | DST_AD0_HI | [23:0] | Destination address 0 (upper 24 bits). | 0x00000000 | RW |
| 0x162-0x17f | RESERVED | [31:0] | Reserved. | 0x00000000 | 0 |

The MADDR_CTRL register bits 0 and 30 together with the RX_FILTER_CTRL register bit 0 (offset 0x103, Table 3–37 on page 3–59) control address checking as follows:

Table 3-47: Destination Address Checking Behavior

| RX_FILTER_CTRL[0] | MADDR_CTRL[30] | MADDR_CTRL[0] | Behavior |
|-------------------|----------------|---------------|--|
| 1 | — | — | Promiscuous receive mode: IP core accepts frames regardless of destination address. |
| 0 | 0 | — | IP core treats all packets as if they have a matching address. (However, the IP core may filter packets based on other criteria). |
| | 1 | 0 | IP core treats all packets as if they have a destination address that does not match (and therefore, accepts no packets). |
| | 1 | 1 | IP core performs destination address checking and filters out packets whose destination address does not match the address specified in the DST_AD0_LO and DST_AD0_HI registers. |

Related Information

[Address Checking](#) on page 3-24

Describes the register setting requirements to specify address checking.

Statistics Registers

The statistics registers count Ethernet traffic and errors. The 64-bit statistics registers are designed to roll over, to ensure timing closure on the FPGA. However, these registers should never roll over if the link is functioning properly. The statistics registers check the size of frames, which includes the following fields:

- Size of the destination address
- Size of the source address
- Size of the data
- Four bytes of CRC

The statistics counters module is a synthesis option. The statistics registers are counters that are implemented inside the CSR. When you turn on the **Statistics counters** parameter in the 40-100GbE parameter editor, the counters are implemented in the CSR. When you turn off the **Statistics counters**

parameter in the 40-100GbE parameter editor, the counters are not implemented in the CSR, and read access to the counters returns read data equal to 0.

To read the value of a statistics register without clearing it, write the value of 0x4 to the `MAC_CMD_config` register. To read the value of a statistics register and reset the counters, write the value of 0xC to the `MAC_CMD_config` register.

Table 3-48: Transmit Side Statistics Registers

The initial value after reset is 0 for all statistics registers.

| Address | Name- | Description | Access |
|---------|------------------------|--|--------|
| 0x200 | CNTR_TX_64B_LO | Number of 64-byte transmitted frames, including the CRC field but excluding the preamble and SFD bytes (lower 32 bits) | RO |
| 0x201 | CNTR_TX_64B_HI | Number of 64-byte transmitted frames (upper 32 bits), including the CRC field but excluding the preamble and SFD bytes (upper 32 bits) | RO |
| 0x202 | CNTR_TX_65to127B_LO | Number of transmitted frames between 65–127 bytes (lower 32 bits) | RO |
| 0x203 | CNTR_TX_65to127B_HI | Number of transmitted frames between 65–127 bytes (upper 32 bits) | RO |
| 0x204 | CNTR_TX_128to255B_LO | Number of transmitted frames between 128–255 bytes (lower 32 bits) | RO |
| 0x205 | CNTR_TX_128to255B_HI | Number of transmitted frames between 128–255 bytes (upper 32 bits) | RO |
| 0x206 | CNTR_TX_256to511B_LO | Number of transmitted frames between 256–511 bytes (lower 32 bits) | RO |
| 0x207 | CNTR_TX_256to511B_HI | Number of transmitted frames between 256–511 bytes (upper 32 bits) | RO |
| 0x208 | CNTR_TX_512to1023B_LO | Number of transmitted frames between 512–1023 bytes (lower 32 bits) | RO |
| 0x209 | CNTR_TX_512to1023B_HI | Number of transmitted frames between 512–1023 bytes (upper 32 bits) | RO |
| 0x20A | CNTR_TX_1024to1518B_LO | Number of transmitted frames between 1024–1518 bytes (lower 32 bits) | RO |
| 0x20B | CNTR_TX_1024to1518B_HI | Number of transmitted frames between 1024–1518 bytes (upper 32 bits) | RO |

| Address | Name- | Description | Access |
|---------|-------------------------------|--|--------|
| 0x20C | CNTR_TX_ 1519toMAXB_LO | Number of transmitted frames between 1519 and max size defined in 0x103 (lower 32 bits) | RO |
| 0x20D | CNTR_TX_ 1519toMAXB_HI | Number of transmitted frames between 1519 and max size defined in 0x103 (upper 32 bits) | RO |
| 0x20E | CNTR_TX_ OVERSIZE_LO | Number of oversized frames transmitted (lower 32 bits) | RO |
| 0x20F | CNTR_TX_ OVERSIZE_HI | Number of oversized frames transmitted (upper 32 bits) | RO |
| 0x210 | CNTR_TX_MCAST_ DATA_ERR_LO | Number of errored multicast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x211 | CNTR_TX_MCAST_ DATA_ERR_HI | Number of errored multicast frames transmitted, excluding control frames (upper 32 bits) | RO |
| 0x212 | CNTR_TX_MCAST_ DATA_OK_LO | Number of valid multicast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x213 | CNTR_TX_MCAST_ DATA_OK_HI | Number of valid multicast frames transmitted, excluding control frames (upper 32 bits) | RO |
| 0x214 | CNTR_TX_BCAST_ DATA_ERR_LO | Number of errored broadcast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x215 | CNTR_TX_BCAST_ DATA_ERR_HI | Number of errored broadcast frames transmitted, excluding control frames (upper 32 bits) | RO |
| 0x216 | CNTR_TX_BCAST_ DATA_OK_LO | Number of valid broadcast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x217 | CNTR_TX_BCAST_ DATA_OK_HI | Number of valid broadcast frames transmitted, excluding control frames (upper 32 bits) | RO |
| 0x218 | CNTR_TX_UCAST_ DATA_ERR_LO | Number of errored unicast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x219 | CNTR_TX_UCAST_ DATA_ERR_HI | Number of errored unicast frames transmitted, excluding control frames (upper 32 bits) | RO |
| 0x21A | CNTR_TX_UCAST_ DATA_OK_LO | Number of valid unicast frames transmitted, excluding control frames (lower 32 bits) | RO |
| 0x21B | CNTR_TX_UCAST_ DATA_OK_HI | Number of valid unicast frames transmitted, excluding control frames (upper 32 bits) | RO |

| Address | Name- | Description | Access |
|---------|-----------------------|--|--------|
| 0x21C | CNTR_TX_MCAST_CTRL_LO | Number of valid multicast frames transmitted (lower 32 bits) | RO |
| 0x21D | CNTR_TX_MCAST_CTRL_HI | Number of valid multicast frames transmitted (upper 32 bits) | RO |
| 0x21E | CNTR_TX_BCAST_CTRL_LO | Number of valid broadcast frames transmitted (lower 32 bits) | RO |
| 0x21F | CNTR_TX_BCAST_CTRL_HI | Number of valid broadcast frames transmitted (upper 32 bits) | RO |
| 0x220 | CNTR_TX_UCAST_CTRL_LO | Number of valid unicast frames transmitted (lower 32 bits) | RO |
| 0x221 | CNTR_TX_UCAST_CTRL_HI | Number of valid unicast frames transmitted (upper 32 bits) | RO |
| 0x222 | CNTR_TX_PAUSE_LO | Number of valid pause frames transmitted (lower 32 bits) | RO |
| 0x223 | CNTR_TX_PAUSE_HI | Number of valid pause frames transmitted (upper 32 bits) | RO |
| 0x224 | CNTR_TX_FRAGMENTS_LO | Number of transmitted frames of less than 64 bytes reporting a CRC error (lower 32 bits) | RO |
| 0x225 | CNTR_TX_FRAGMENTS_HI | Number of transmitted frames of less than 64 bytes reporting a CRC error (upper 32 bits) | RO |
| 0x226 | CNTR_TX_JABBERS_LO | Number of transmitted oversized frames reporting a CRC error (lower 32 bits) | RO |
| 0x227 | CNTR_TX_JABBERS_HI | Number of transmitted oversized frames reporting a CRC error (upper 32 bits) | RO |
| 0x228 | CNTR_TX_CRCERR_LO | Number of transmitted frames between 64 bytes and the value configured in 0x103 register reporting a CRC error (lower 32 bits) | RO |
| 0x229 | CNTR_TX_CRCERR_HI | Number of transmitted frames between 64 bytes and the value configured in 0x103 register reporting a CRC error (upper 32 bits) | RO |

TX Packet Statistics

| | | | |
|-------|---------------|--|----|
| 0x22A | CNTR_TX_ST_LO | Number of transmitted frame starts, excluding pause frames (lower 32 bits) | RO |
|-------|---------------|--|----|

| Address | Name- | Description | Access |
|---------|---------------|---|--------|
| 0x22B | CNTR_TX_ST_HI | Number of transmitted frame starts, excluding pause frames (upper 32 bits) | RO |
| 0x22C | CNTR_TX_DB_LO | Number of transmitted data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the MAC to the PHY that contain at least some data)(lower 32 bits) | RO |
| 0x22D | CNTR_TX_DB_HI | Number of transmitted data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the MAC to the PHY that contain at least some data) (upper 32 bits) | RO |

Table 3-49: Receive Side Statistics Registers

| Address | Name- | Description | Access |
|---------|-----------------------|---|--------|
| 0x280 | CNTR_RX_64B_LO | Number of 64-byte received frames (lower 32 bits), including the CRC field but excluding the preamble and SFD bytes | RO |
| 0x281 | CNTR_RX_64B_HI | Number of 64-byte received frames (upper 32 bits), including the CRC field but excluding the preamble and SFD bytes | RO |
| 0x282 | CNTR_RX_65to127B_LO | Number of received frames between 65–127 bytes (lower 32 bits) | RO |
| 0x283 | CNTR_RX_65to127B_HI | Number of received frames between 65–127 bytes (upper 32 bits) | RO |
| 0x284 | CNTR_RX_128to255B_LO | Number of received frames between 128 –255 bytes (lower 32 bits) | RO |
| 0x285 | CNTR_RX_128to255B_HI | Number of received frames between 128 –255 bytes (upper 32 bits) | RO |
| 0x286 | CNTR_RX_256to511B_LO | Number of received frames between 256 –511 bytes (lower 32 bits) | RO |
| 0x287 | CNTR_RX_256to511B_HI | Number of received frames between 256 –511 bytes (upper 32 bits) | RO |
| 0x288 | CNTR_RX_512to1023B_LO | Number of received frames between 512–1023 bytes (lower 32 bits) | RO |
| 0x289 | CNTR_RX_512to1023B_HI | Number of received frames between 512 –1023 bytes (upper 32 bits) | RO |

| Address | Name- | Description | Access |
|---------|-------------------------------|--|--------|
| 0x28A | CNTR_RX_ 1024to1518B_LO | Number of received frames between 1024–1518 bytes (lower 32 bits) | RO |
| 0x28B | CNTR_RX_ 1024to1518B_HI | Number of received frames between 1024–1518 bytes (upper 32 bits) | RO |
| 0x28C | CNTR_RX_ 1519toMAXB_LO | Number of received frames between 1519 and the maximum size defined in 0x103 (lower 32 bits) | RO |
| 0x18D | CNTR_RX_ 1519toMAXB_HI | Number of received frames between 1519 and the maximum size defined in 0x103 (upper 32 bits) | RO |
| 0x28E | CNTR_RX_ OVERSIZE_LO | Number of oversized frames received (lower 32 bits) | RO |
| 0x28F | CNTR_RX_ OVERSIZE_HI | Number of oversized frames received (upper 32 bits) | RO |
| 0x290 | CNTR_RX_MCAST_ DATA_ERR_LO | Number of errored multicast frames received, excluding control frames (lower 32 bits) | RO |
| 0x291 | CNTR_RX_MCAST_ DATA_ERR_HI | Number of errored multicast frames received, excluding control frames (upper 32 bits) | RO |
| 0x292 | CNTR_RX_MCAST_ DATA_OK_LO | Number of valid multicast frames received, excluding control frames (lower 32 bits) | RO |
| 0x293 | CNTR_RX_MCAST_ DATA_OK_HI | Number of valid multicast frames received, excluding control frames (upper 32 bits) | RO |
| 0x294 | CNTR_RX_BCAST_ DATA_ERR_LO | Number of errored broadcast frames received, excluding control frames (lower 32 bits) | RO |
| 0x295 | CNTR_RX_BCAST_ DATA_ERR_HI | Number of errored broadcast frames received, excluding control frames (upper 32 bits) | RO |
| 0x296 | CNTR_RX_BCAST_ DATA_OK_LO | Number of valid broadcast frames received, excluding control frames (lower 32 bits) | RO |
| 0x297 | CNTR_RX_BCAST_ DATA_OK_HI | Number of valid broadcast frames received, excluding control frames (upper 32 bits) | RO |
| 0x298 | CNTR_RX_UCAST_ DATA_ERR_LO | Number of errored unicast frames received, excluding control frames (lower 32 bits) | RO |
| 0x299 | CNTR_RX_UCAST_ DATA_ERR_HI | Number of errored unicast frames received, excluding control frames (upper 32 bits) | RO |

| Address | Name- | Description | Access |
|---------|--------------------------|--|--------|
| 0x29A | CNTR_RX_UCAST_DATA_OK_LO | Number of valid unicast frames received, excluding control frames (lower 32 bits) | RO |
| 0x29B | CNTR_RX_UCAST_DATA_OK_HI | Number of valid unicast frames received, excluding control frames (upper 32 bits) | RO |
| 0x29C | CNTR_RX_MCAST_CTRL_LO | Number of valid multicast frames received, excluding data frames (lower 32 bits) | RO |
| 0x29D | CNTR_RX_MCAST_CTRL_HI | Number of valid multicast frames received, excluding data frames (upper 32 bits) | RO |
| 0x29E | CNTR_RX_BCAST_CTRL_LO | Number of valid broadcast frames received, excluding data frames (lower 32 bits) | RO |
| 0x29F | CNTR_RX_BCAST_CTRL_HI | Number of valid broadcast frames received, excluding data frames (upper 32 bits) | RO |
| 0x2A0 | CNTR_RX_UCAST_CTRL_LO | Number of valid unicast frames received, excluding data frames (lower 32 bits) | RO |
| 0x2A1 | CNTR_RX_UCAST_CTRL_HI | Number of valid unicast frames received, excluding data frames (upper 32 bits) | RO |
| 0x2A2 | CNTR_RX_PAUSE_LO | Number of valid pause frames received (lower 32 bits) | RO |
| 0x2A3 | CNTR_RX_PAUSE_HI | Number of valid pause frames received (upper 32 bits) | RO |
| 0x2A4 | CNTR_RX_FRAGMENTS_LO | Number of received frames less than 64 bytes and reporting a CRC error (lower 32 bits) | RO |
| 0x2A5 | CNTR_RX_FRAGMENTS_HI | Number of received frames less than 64 bytes and reporting a CRC error (upper 32 bits) | RO |
| 0x2A6 | CNTR_RX_JABBERS_LO | Number of received oversized frames reporting a CRC error (lower 32 bits) | RO |
| 0x2A7 | CNTR_RX_JABBERS_HI | Number of received oversized frames reporting a CRC error (upper 32 bits) | RO |
| 0x2A8 | CNTR_RX_CRCERR_LO | Number of received frames between the length of 64 and the value configured in 0x103 register with CRC error (lower 32 bits) | RO |

| Address | Name- | Description | Access |
|----------------------|-------------------|--|--------|
| 0x2A9 | CNTR_RX_CRCERR_HI | Number of received frames between the length of 64 and the value configured in 0x103 register with CRC error (upper 32 bits) | RO |
| RX Packet Statistics | | | |
| 0x2AA | CNTR_RX_ST_LO | Number of received frame starts (lower 32 bits) | RO |
| 0x2AB | CNTR_RX_ST_HI | Number of received frame starts (upper 32 bits) | RO |
| 0x2AC | CNTR_RX_DB_LO | Number of received data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the PHY to the MAC that contain at least some data) (lower 32 bits) | RO |
| 0x2AD | CNTR_RX_DB_HI | Number of received data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the PHY to the MAC that contain at least some data) (upper 32 bits) | RO |
| 0x2AE | CNTR_RXF_ST_LO | Number of accepted frame starts (lower 32 bits) | RO |
| 0x2AF | CNTR_RXF_ST_HI | Number of accepted frame starts (upper 32 bits) | RO |
| 0x2B0 | CNTR_RXF_DB_LO | Number of accepted data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the PHY to the MAC that contain at least some data and that the MAC accepts) (lower 32 bits) | RO |
| 0x2B1 | CNTR_RXF_DB_HI | Number of accepted data blocks (number of full symbols (2 or 5 words, depending on the variation) that pass from the PHY to the MAC that contain at least some data and that the MAC accepts) (upper 32 bits) | RO |
| 0x2B2 | CNTR_RXF_LONG | Number of packets aborted due to excessive length. | RO |
| 0x2B3 | Reserved | | |
| 0x2B4 | CNTR_RX_EBLK | Number of received blocks with XGMII and CGMII errors. The number of received error blocks is equal to the RX symbol error counter defined by Operation Administration and Maintenance (OAM) in accordance with <i>IEEE standard 802.3ba</i> . | RO |
| 0x2B5 | Reserved | | |
| 0x2B6 | CNTR_RX_RUNT | Number of received runt packets | RO |

| Address | Name | Description | Access |
|---------|-------------|---|--------|
| 0x2B7 | Reserved | | |
| 0x2B8 | CNTR_RX_FCS | Number of received packets with FCS errors. This register maintains a count of the number of pulses on the <code>l<n>_rx_fcs_error</code> output signal in variations with a custom streaming client interface (variations without adapters). | RO |

Related Information

- [Statistics Counters Interface](#) on page 3-39
- [MAC Configuration and Filter Registers](#) on page 3-99
For more information about the `MAC_CMD_CONFIG` register.

40-100GbE Example Design Registers

The following sections describe the registers that are included in the 40-100GbE example design and are not a part of the 40-100GbE IP core. These registers are not available in the example designs for the 40GBASE-KR4 IP core variations.

[PMD Registers](#) on page 3-116

[MDIO Registers](#) on page 3-117

[2-Wire Serial Interface Registers](#) on page 3-118

Related Information

[40-100GbE IP Core Example Design](#) on page 5-1

Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a target device. The listed registers are part of the example design.

PMD Registers

Table 3-50: PMD Control and Status Registers

The physical medium dependent (PMD) control and status registers allow you to turn on and determine the status of the PMD device in the 40-100GbE IP core example design.

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|----------------|--------|---|----------------|--------|
| 0x400 | PMD_VERSION | [31:0] | PMD revision. The character string is "OPTs." | 0x4F505473 | R |
| 0x401 | SCRATCH_PMD | [31:0] | Scratch register available for testing. | 0x00000000 | RW |
| 0x402 | PMD_CMD_CONFIG | [0] | Writing a 1 enables the PMD. When you set this register to the value of 1, the IP core seeks RX lock. | 0x00000000 | RW |

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|------------|------|---|----------------|--------|
| 0x403 | CMD_STATUS | [6] | PMD global alarm. | 1b'0 | R |
| | | [5] | Programmable alarm 3. Defaults to module ready. | 1b'0 | R |
| | | [4] | Programmable alarm 2. Defaults to high power mode on. | 1b'0 | R |
| | | [3] | Programmable alarm 1. Defaults to RX CDR lock. | 1b'0 | R |
| | | [2] | RX side signal detected. | 1'b0 | R |
| | | [1] | When asserted, the PMD module is physically plugged in. When deasserted, it is offline. | 1b'1 | R |
| | | [0] | When 1, the PMD is resetting. When 0, the PMD is available. | 1'b1 | R |

Related Information

- [40-100GbE IP Core Example Design](#) on page 5-1
Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a target device. The listed registers are part of the example design.
- [CFP Multi-Source Agreement website](#)
For more information about PMD devices, refer to the *Next Gen PMD CFP MSA Baseline Specifications* available on the CFP Multi-Source Agreement website.

MDIO Registers

The management data input/output (MDIO) interface is a serial bus interface for the Ethernet. In the 40-100GbE IP core example design, client logic drives the MDIO module when an MDIO serial interface controls the external PMD or CFP device. This serial interface includes two wires, an MDC clock driven by the MAC and a bidirectional data line which can be driven by up to 31 PHY slave devices. The minimum clock frequency is 2.5 MHz. Refer to the specification of your PMD or CFP device for additional information.

Table 3-51: MDIO Registers

Describes the MDIO registers in the 40-100GbE IP core example design. These registers are typically used to communicate with an external optical module such as a PMD 100G optical module.

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|------------|--------|---------------------|----------------|--------|
| 0x410 | MDIO_WDATA | [15:0] | Data to be written. | 0x0000 | RW |

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|------------|---------|--|----------------|--------|
| 0x411 | MDIO_RDATA | [31] | Link is busy. | 1b'0 | R |
| | | [15:0] | Result of previous read. | 0x0000 | R |
| 0x412 | MDIO_ADDR | [14:10] | PHY address driven to external module. | 5'h00 | RW |
| | | [9:5] | PHY address to access. | 5'h00 | RW |
| | | [4:0] | DEV address. | 5'h01 | RW |
| 0x413 | MDIO_CMD | [3] | Address of the register to be written. | 1b'0 | RW |
| | | [2] | Write signal. | 1b'0 | RW |
| | | [1] | Address of the register to read. This address post increments. | 1b'0 | RW |
| | | [0] | Read signal. | 1b'0 | RW |

Related Information

[40-100GbE IP Core Example Design](#) on page 5-1

Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a target device. The listed registers are part of the example design.

2-Wire Serial Interface Registers

Several optical modules use the 2-wire serial interface protocol, which requires only limited resources. This interface uses two, bidirectional open-drain lines, a serial data line (SDA) and a serial clock (SCL). The master node drives the clock and addresses slaves. Data is sent most significant bit first. The slave node receives the clock and address. Unlike the MDIO protocol, the 2-wire serial interface can include multiple masters. In addition, the master and slave roles may be changed after a STOP message is sent. The interface has an eight-bit address space, so a maximum of 255 nodes can communicate. The interface runs at 400 KHz.

Table 3-52: 2-Wire Serial Interface Registers

Describes the 2-wire serial registers in the 40-100GbE IP core example design. Refer to the datasheet for your 2-wire device for more information.

Note: The 2-wire serial interface in the IP core example design does not support the specialized CRC-8 check or multibyte transaction modes.

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|-----------|-------|---------------------|----------------|--------|
| 0x420 | 2WS_WDATA | [7:0] | Data to be written. | 0x00 | RW |

| Address | Name | Bits | Description | HW Reset Value | Access |
|---------|-----------|--------|--|----------------|--------|
| 0x421 | 2WS_RDATA | [31] | When asserted, the link is busy. | 1b'0 | R |
| | | [30] | When asserted, indicates that the slave failed to acknowledge during the previous operation. | 1b'0 | R |
| | | [7:0] | Result of previous read. | 0x00 | R |
| 0x422 | 2WS_ADDR | [15:8] | Slave address. | 0xA0 | RW |
| | | [7:0] | Memory address. | 0x00 | RW |
| 0x423 | 2WS_CMD | [1] | When asserted, indicates a write command. | 0x0 | RW |
| | | [0] | When asserted, indicates a read. | 0x0 | RW |

Related Information

[40-100GbE IP Core Example Design](#) on page 5-1

Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a target device. The listed registers are part of the example design.

Ethernet Glossary

Table 3-53: Ethernet Glossary

Provides definitions for some terms associated with the Ethernet protocol.

| Term | Definition |
|-------|--|
| BIP | Bit Interleaved Parity. A diagonal parity field which is carried in the periodic alignment markers on each virtual lane, allowing isolation of a bit error to a physical channel. |
| CAUI | 100 gigabit attachment unit interface. (C is the symbol in Roman Numerals for 100). This is an electrical interface that which is based on a 10-lane interface with a bandwidth of 10 Gbps per lane. (In this implementation, the PMA multiplexes the 20 PCS lanes into 10 physical lanes. |
| CGMII | 100 gigabit media independent interface. (C is the symbol in Roman Numerals for 100). This is the byte-oriented interface protocol that connects the PCS and MAC. |
| DIC | Deficit Idle Counter. A rule for inserting and deleting idles as necessary to maintain the average IPG. The alternative is to always insert idles which could lead to reduced bandwidth. |

| Term | Definition |
|--------|--|
| FCS | Frame Check Sequence. A CRC-32 with bit reordering and inversion. |
| Frame | Ethernet formatted packet. A frame consists of a start delimiter byte, a 7 byte preamble, variable length data, 4-byte FCS, and an end delimiter byte. |
| IPG | Inter Packet Gap. Includes the end of frame delimiter and subsequent IDLE bytes up to, but not including the next start of frame delimiter. The protocol requires an average gap of 12 bytes. |
| MAC | Media Access Control. Formats a user packet stream into proper Ethernet frames for delivery to the PCS. The MAC generates the FCS and checks and maintains the IPG. |
| MII | Media Independent Interface. The byte-oriented protocol used by the PCS. Sometimes distinguished with roman numerals, XGMII (10), XLGMII (40), CGMII (100). |
| Octet | Byte. Note that Ethernet specifications primarily use least significant bit first ordering which is opposite from the default behavior of most contemporary CAD tools. |
| PCS | Physical Coding Sublayer. Presents the underlying hardware as a byte-oriented communication channel. |
| XLAUI | 40 gigabit attachment unit interface. (XL is the symbol in Roman Numerals for 40). This is an electrical interface that which is based on a 4-lane interface with a bandwidth of 10 Gbps per lane. |
| XLGMII | 40 gigabit media independent interface. (XL is the symbol in Roman Numerals for 40). This is the byte-oriented interface protocol that connects the PCS and MAC. |

2014.12.15

UG-01088



Subscribe



Send Feedback

If you are experiencing difficulties bringing up your 40-100GbE IP core link in hardware, Altera suggests that you begin debugging at the most basic level, with word lock. Then, consider higher level issues.

The following steps should help you identify and resolve common problems that occur when bringing up a 40-100GbE IP core link:

1. Establish word lock—The RX lanes should be able to achieve word lock even in the presence of extreme bit error rates. If unable to achieve word lock, check the transceiver clocking and data rate configuration. Check for cabling errors such as the reversal of the TX and RX lanes. Check the clock frequency monitors.
2. Establish the alignment marker lock—Virtual lane alignment marker lock requires a moderate quality transceiver connection. If the lock is completely absent, recheck the alignment marker period. If the lock is intermittent, recheck the transceiver physical connection and analog settings.
3. Establish lane integrity—When operating properly, the lanes should not experience bit errors at a rate greater than roughly one per hour per day. Bit errors within data packets are identified as FCS errors. Bit errors in control information including IDLE frames generally cause errors in XL/CGMII decoding. The bit interleaved parity (BIP) mechanism is a diagonal parity computation that enables tracing a protocol error back to a physical lane.
4. Verify packet traffic—The Ethernet protocol includes automatic lane reordering so the higher levels should follow the PCS. If the PCS is locked, but higher level traffic is corrupted, there may be a problem with the remote transmitter virtual lane tags.
5. Tuning—You can adjust analog parameters to minimize any remaining bit error rate. IDLE traffic is representative for analog purposes.

Related Information

- [Transceiver PHY Control and Status Registers](#) on page 3-80
For more information about the analog parameters for Stratix IV devices, refer to the description of the 40-100GbE IP core Stratix IV transceiver analog settings register GX_CTRL1 at offset 0x007.
- [Altera Transceiver PHY IP Core User Guide](#)
For information about the analog parameters for Arria V GZ devices and Stratix V devices.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered

40-100GbE IP Core Example Design



2014.12.15

UG-01088



Subscribe



Send Feedback

Altera provides an example design with the 40-100GbE IP core. This example design is ready for compilation and can be configured on a C2 speed grade device.

You can use the example design as an example for correct connection of your IP core to your design, or as a starter design you can customize for your own design requirements.

Separate figures illustrate the example design structure for 40GBASE-KR IP core variations and for the other IP core variations.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered

Figure A-1: High Level Block Diagram for the 40-100GbE Example Design

High level block diagram for non-40GBASE-KR4 example designs. You can generate a non-40GBASE-KR4 example design for any non-40GBASE-KR4, full duplex IP core variation that includes both MAC and PHY components.


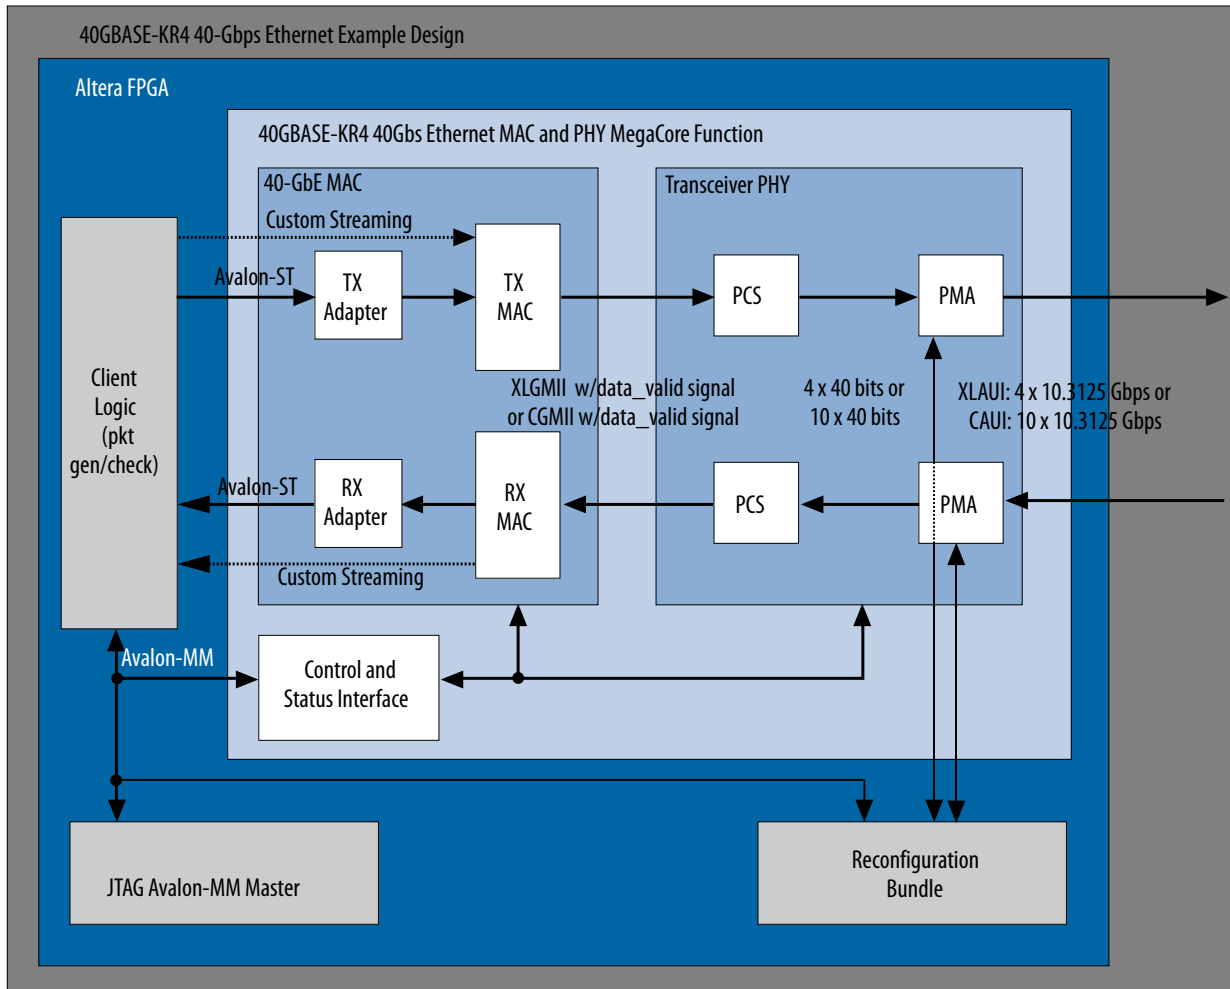


Figure A-2: High Level Block Diagram for the 40GBASE-KR4 Example Design

High level block diagram for 40GBASE-KR4 example designs. You can generate a 40GBASE-KR4 example design for any 40GBASE-KR4 IP core variation that includes both MAC and PHY components, does not have Synchronous Ethernet support, does not have the link training microporcoessor interface, and has RX equalization enabled. Variations that meet these criteria can be with or without adapters, unlike the testbench, which is available only for variations with adapters.



Altera's example design includes either the 40GbE or 100GbE IP core. Client logic connects to the TX and RX Ethernet MAC. The TX and RX logic in the MAC includes an optional adapter which is available for both the 40GbE and 100GbE IP cores. When you use the optional adapter, the start of packet (SOP) is always in the most significant word (64 bits) of the bus, simplifying the interpretation of incoming data. When the configuration includes adapters, the interface between the client logic and TX and RX FIFOs uses the Avalon-ST protocol. Without adapters, the interface between client logic and the TX and RX MAC is a custom streaming interface.

The bandwidth for the two versions is the same. However, the version without adapters achieves this bandwidth with a narrower bus, because it does not restrict the SOP to the most significant word of the bus. In the version without adapters, the SOP can be in the MSB of any word. The penalty for the less restrictive positioning of the SOP is a data stream which is more difficult to interpret.

The interface between the MAC and PHY modules of the IP core is XLGMII for the 40GbE IP core and CGMII for the 100GbE IP core. The interface between the PHY and external physical medium dependent (PMD) optical module or other device is XLAUI for the 40GbE IP core and CAUI for the 100GbE IP core, providing a bandwidth of either 4 or 10×10.3125 Gbps.

An Avalon-MM control and status (management) interface provides access to the MAC and PHY registers in the IP core and also controls the MDIO, 2-wire serial, and PMD controllers on the PCB (for non-40GBASE-KR4 variations) or the reconfiguration bundle (for 40GBASE-KR4 variations).

The example design is provided as a Quartus II project. The example design is crafted for you to configure on a C2 device in a specific Altera development kit. To use a different device or development kit, you must modify the project.

By default, the example design configures on one of the following Altera development kits, as appropriate for the IP core target device and variation:

- 100G Development Kit, Stratix IV GT Edition
- 100G Development Kit, Stratix V GX Edition
- Transceiver Signal Integrity Development Kit, Stratix V GT Edition
- Transceiver Signal Integrity Development Kit, Stratix V GX Edition (for 40GBASE-KR4 IP cores)

To set up and configure the example design on the device, follow these steps:

1. Follow the steps in Chapter 2, Getting Started to generate your IP core. Refer to the figures for information about the variations for which an example design can be generated.

Note: When prompted at the start of generation, you must turn on **Generate example design**.

2. In the Quartus II software, on the File menu, click **Open Project**.
3. Navigate to the example design project folder, select the Quartus Project File (**.qpf**), *<instance_name>_example/alt_e40_e100/example/<example_design_name>.qpf*, and click **Open**.
4. On the Processing menu, click **Start Compilation**.
5. Program the targeted Altera device with the Quartus II Programmer.

Related Information

- [40-100GbE Example Design Registers](#) on page 3-116
- [PMD Registers](#) on page 3-116
- [MDIO Registers](#) on page 3-117
- [2-Wire Serial Interface Registers](#) on page 3-118
- [40-100GbE IP Core File Set](#)
Illustrates the path to the example design Quartus II project file.
- [Avalon Interface Specifications](#)
For more information about the Avalon-MM and Avalon-ST protocols, including timing diagrams.
- [All Development Kits web page](#)
For information about the Altera development kits that the example designs target.
- [Volume 3 of the Quartus II Handbook](#)
For information about programming an Altera device, refer to the "Device Programming" section.

Address Map Changes for the 40-100GbE IP Core v12.0 Release



2014.12.15

UG-01088



Subscribe



Send Feedback

Table B-1: Address Map and Register Name Changes for the 40-100GbE IP Core v12.0 Release

For improved organization and clarity, version 12.0 of the 40- and 100-Gbps MAC and PHY MegaCore Function revised the address map and renamed a few registers. This table lists the address map changes in the 12.0 software release. These address map changes apply to Stratix IV and Stratix V devices.

| Current 40/100GbE Address | Current Name | Previous Address | Previous Name |
|---------------------------|--|------------------|---------------|
| 0x017 | PCS_HW_ERR | | HW_ERR |
| 0x018 | BER_MONITOR | | Reserved |
| 0x019 | TEST_MODE | | Reserved |
| 0x01a | TEST_PATTERN_COUNTER | | Reserved |
| 0x01b | Enable Link Fault sequence | | Reserved |
| 0x01d | PHY reset | | Reserved |
| 0x120 | MAC_HW_ERR | | Reserved |
| 0x121 | MAC Reset | | Reserved |
| 0x122 | MAC/RS link fault sequence configuration | | Reserved |
| 0x123 | CRC_CONFIG | | Reserved |
| 0x162-0x17f | Reserved | 0x160-0x017f | DST_ADR |

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered





2014.12.15

UG-01088



Subscribe



Send Feedback

This appendix duplicates the 10GBASE-KR PHY register listings from the [Altera Transceiver PHY IP Core User Guide](#). Altera provides this appendix as a convenience to make the full 40GBASE-KR4 register information available in the 40-100GbE IP core user guide. While Altera makes every attempt to keep the information in the appendix up-to-date, the most up-to-date information is always found in the [Altera Transceiver PHY IP Core User Guide](#), and the appendix is not guaranteed to be up-to-date at any particular time.

Most 40GBASE-KR4 registers are 10GBASE-KR PHY registers of the 10GBASE-KR PHY IP core, documented in the [Altera Transceiver PHY IP Core User Guide](#) and duplicated, with a potential time lag for updates, in this appendix. The register offsets are identical in the 40GBASE-KR4 variations of the 40-100GbE IP core. However, the 40GBASE-KR4 variations of the 40-100GbE IP core have additional 40GBASE-KR4 related registers and register fields.

[40GBASE-KR4 Registers](#) documents the differences between the 10GBASE-KR PHY register definitions in the *10GBASE-KR PHY Register Definitions* section of the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide* and the 40GBASE-KR4 registers of the 40-100GbE IP core. All 10GBASE-KR PHY registers and register fields not listed in [40GBASE-KR4 Registers](#) are available in the 40GBASE-KR4 variations of the 40-100GbE IP core.

Where the *Altera Transceiver PHY IP Core User Guide* and this appendix list 10GBASE-R, substitute 40GBASE-KR4 with auto-negotiation and link training both turned off, and where they list 10GBASE-KR (except in the description of 0xCB[24:0]), substitute 40GBASE-KR4. Where a register field description in the *Altera Transceiver PHY IP Core User Guide* and this appendix refers to link training or FEC in the single-lane 10GBASE-KR PHY IP core, substitute link training or FEC on Lane 0 of the 40GBASE-KR4 IP core variation. Where a register field description in the *Altera Transceiver PHY IP Core User Guide* and this appendix refers to the Transceiver Reconfiguration IP core, substitute the reconfiguration bundle.

Related Information

[Altera Transceiver PHY IP Core User Guide](#)

The 40GBASE-KR4 variations of the 40-100GbE IP core use the 10GBASE-KR PHY IP core PHY registers at internal offsets 0xB0–0xFF (at IP core register map offsets 0xB0–0xFF), in addition to the registers listed in [40GBASE-KR4 Registers](#). Information about this PHY IP core, including up-to-date register descriptions, is available in the *Backplane Ethernet 10GBASE-KR PHY IP Core with FEC Option* chapter of the *Altera Transceiver PHY IP Core User Guide*.

10GBASE-KR PHY Register Definitions

The Avalon-MM master interface signals provide access to the control and status registers.

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered

The following table specifies the control and status registers that you can access over the Avalon-MM PHY management interface. A single address space provides access to all registers.

Notes:

- Unless otherwise indicated, the default value of all registers is 0.
- Writing to reserved or undefined register addresses may have undefined side effects.
- To avoid any unspecified bits to be erroneously overwritten, you must perform read-modify-writes to change the register values.

Table C-1: 10GBASE-KR Register Definitions

| Word Addr | Bit | R/W | Name | Description |
|-----------|-----|-----|--------------------------------|---|
| 0xB0 | 0 | RW | Reset SEQ | When set to 1, resets the 10GBASE-KR sequencer, initiates a PCS reconfiguration, and may restart Auto-Negotiation, Link Training or both if AN and LT are enabled (10GBASE-KR mode). SEQ Force Mode[2:0] forces these modes. This reset self clears. |
| | 1 | RW | Disable AN Timer | Auto-Negotiation disable timer. If disabled (Disable AN Timer = 1), AN may get stuck and require software support to remove the ABILITY_DETECT capability if the link partner does not include this feature. In addition, software may have to take the link out of loopback mode if the link is stuck in the ACKNOWLEDGE_DETECT state. To enable this timer set Disable AN Timer = 0. |
| | 2 | RW | Disable LF Timer | When set to 1, disables the Link Fault timer. When set to 0, the Link Fault timer is enabled. |
| | 6:4 | RW | SEQ Force Mode[2:0] | Forces the sequencer to a specific protocol. Must write the Reset SEQ bit to 1 for the Force to take effect. The following encodings are defined: <ul style="list-style-type: none"> • 3'b000: No force • 3'b001: GigE • 3'b010: Reserved • 3'b011: Reserved • 3'b100: 10GBASE-R • 3'b101: 10GBASE-KR • Others: Reserved |
| | 16 | RW | Assert KR FEC Ability | When set to 1, indicates that the FEC ability is supported. This bit defaults to 1 if the Set FEC_ability bit on power up/reset bit is on. For more information, refer to the FEC variable FEC_Enable as defined in <i>Clause 74.8.2</i> and 10GBASE-KR PMD control register bit (1.171.0) IEEE 802.3ap-2007. |
| | 17 | RW | Enable KR FEC Error Indication | When set to 1, the FEC module indicates errors to the 10G PCS. For more information, refer to the KR FEC variable FEC_enable_Error_to_PCS and 10GBASE-KR PMD control register bit (1.171.1) as defined in <i>Clause 74.8.3 of IEEE 302.3ap-2007</i> . |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|------|--|--|
| | 18 | RW | Assert KR FEC Request | When set to 1, indicates that the core is requesting the FEC ability. When this bit changes, you must assert the <code>Reset_SEQ</code> bit (0xB0[0]) to renegotiate with the new value. |
| 0xB1 | 0 | R | SEQ Link Ready | When asserted, the sequencer is indicating that the link is ready. |
| | 1 | R | SEQ AN timeout | When asserted, the sequencer has had an Auto-Negotiation timeout. This bit is latched and is reset when the sequencer restarts Auto-Negotiation. |
| | 2 | R | SEQ LT timeout | When set, indicates that the Sequencer has had a timeout. |
| | 13:8 | R | SEQ Reconfig Mode[5:0] | Specifies the Sequencer mode for PCS reconfiguration. The following modes are defined: <ul style="list-style-type: none"> Bit 8, mode[0]: AN mode Bit 9, mode[1]: LT Mode Bit 10, mode[2]: 10G data mode Bit 11, mode[3]: Gige data mode Bit 12, mode[4]: Reserved for XAUI Bit13, mode[5]: 10G FEC mode |
| | 16 | R | KR FEC Ability | Indicates whether or not the 10GBASE-KR PHY supports FEC. For more information, refer to the FEC variable <code>FEC_Enable</code> as defined in <i>Clause 74.8.2</i> and 10GBASE-KR PMD control register bit (1.171.0) IEEE 802.3ap-2007. |
| | 17 | R | Enable KR FEC Error Indication Ability | When set to 1, indicates that the 10GBASE-KR PHY is capable of reporting FEC decoding errors to the PCS. For more information, refer to the KR FEC variable <code>FEC_enable_Error_to_PCS</code> and 10GBASE-KR PMD control register bit (1.171.1) as defined in <i>Clause 74.8.3 of IEEE 302.3ap-2007</i> . |
| 0xB2 | 0 | RW | FEC TX trans error | When asserted, indicates that the error insertion feature in the FEC Transcoder is enabled. |
| | 1 | RW | FEC TX burst error | When asserted, indicates that the error insertion feature in the FEC Encoder is enabled. |
| | 5:2 | RW | FEC TX burst length | Specifies the length of the error burst. Values 1-16 are available. |
| | 10:6 | | Reserved | |
| | 11 | RWSC | FEC TX Error Insert | Writing a 1 inserts 1 error pulse into the TX FEC depending on the Transcoder and Burst error settings. Software clears this register. |
| | 31:15 | RWSC | Reserved | |
| 0xB3 | 31:0 | RSC | FEC Corrected Blocks | Counts the number of corrected FEC blocks. Resets to 0 when read. Otherwise, it holds at the maximum count and does not roll over. Refer to <i>Clause 74.8.4.1 of IEEE 802.3ap-2000</i> for details. |

| Word Addr | Bit | R/W | Name | Description |
|-----------|------|-----|---------------------------|--|
| 0xB4 | 31:0 | RSC | FEC Uncorrected Blocks | Counts the number of uncorrectable FEC blocks. Resets to 0 when read. Otherwise, it holds at the maximum count and does not roll over. Refer to <i>Clause 74.8.4.1 of IEEE 802.3ap-2000</i> for details. |
| 0xC0 | 0 | RW | AN enable | When set to 1, enables Auto-Negotiation function. The default value is 1. For additional information, refer to bit 7.0.12 in Clause 73.8 Management Register Requirements, of <i>IEEE 802.3ap-2007</i> . |
| | 1 | RW | AN base pages ctrl | When set to 1, the user base pages are enabled. You can send any arbitrary data via the user base page low/high bits. When set to 0, the user base pages are disabled and the state machine generates the base pages to send. |
| | 2 | RW | AN next pages ctrl | When set to 1, the user next pages are enabled. You can send any arbitrary data via the user next page low/high bits. When set to 0, the user next pages are disabled. The state machine generates the null message to send as next pages. |
| | 3 | R | Local device remote fault | When set to 1, the local device signals Remote Faults in the Auto-Negotiation pages. When set to 0 a fault has not occurred. |
| | 4 | RW | Force TX nonce value | When set to 1, forces the TX none value to support some UNH-IOL testing modes. When set to 0, operates normally. |
| | 5 | RW | Override AN | When set to 1, the override settings defined by the AN_TECH, AN_FEC and AN_PAUSE registers take effect. |
| 0xC1 | 0 | RW | Reset AN | When set to 1, resets all the 10GBASE-KR Auto-Negotiation state machines. This bit is self-clearing. |
| | 4 | RW | Restart AN TX SM | When set to 1, restarts the 10GBASE-KR TX state machine. This bit self clears. This bit is active only when the TX state machine is in the AN state. For more information, refer to bit 7.0.9 in Clause 73.8 Management Register Requirements of <i>IEEE 802.3ap-2007</i> . |
| | 8 | RW | AN Next Page | When asserted, new next page info is ready to send. The data is in the XNP TX registers. When 0, the TX interface sends null pages. This bit self clears. Next Page (NP) is encoded in bit D15 of Link Codeword. For more information, refer to Clause 73.6.9 and bit 7.16.15 of Clause 45.2.7.6 of <i>IEEE 802.3ap-2007</i> . |
| 0xC2 | 1 | RO | AN page received | When set to 1, a page has been received. When 0, a page has not been received. The current value clears when the register is read. For more information, refer to bit 7.1.6 in Clause 73.8 of <i>IEEE 802.3ap-2007</i> . |
| | 2 | RO | AN Complete | When asserted, Auto-Negotiation has completed. When 0, Auto-Negotiation is in progress. For more information, refer to bit 7.1.5 in Clause 73.8 of <i>IEEE 802.3ap-2007</i> . |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|-----------------------|---|
| | 3 | RO | AN ADV Remote Fault | When set to 1, fault information has been sent to the link partner. When 0, a fault has not occurred. The current value clears when the register is read. Remote Fault (RF) is encoded in bit D13 of the base Link Codeword. For more information, refer to Clause 73.6.7 of and bit 7.16.13 of <i>IEEE 802.3ap-2007</i> . |
| | 4 | RO | AN RX SM Idle | When set to 1, the Auto-Negotiation state machine is in the idle state. Incoming data is not Clause 73 compatible. When 0, the Auto-Negotiation is in progress. |
| | 5 | RO | AN Ability | When set to 1, the transceiver PHY is able to perform Auto-Negotiation. When set to 0, the transceiver PHY is not able to perform Auto-Negotiation. If your variant includes Auto-Negotiation, this bit is tied to 1. For more information, refer to bits 7.1.3 and 7.48.0 of Clause 45 of <i>IEEE 802.3ap-2007</i> . |
| | 6 | RO | AN Status | When set to 1, link is up. When 0, the link is down. The current value clears when the register is read. For more information, refer to bit 7.1.2 of Clause 45 of <i>IEEE 802.3ap-2007</i> . |
| | 7 | RO | LP AN Ability | When set to 1, the link partner is able to perform Auto-Negotiation. When 0, the link partner is not able to perform Auto-Negotiation. For more information, refer to bit 7.1.0 of Clause 45 of <i>IEEE 802.3ap-2007</i> . |
| | 8 | RO | Enable FEC | When asserted, indicates that auto-negotiation is complete and that communicate includes FEC. For more information refer to <i>Clause 7.48.4</i> . |
| | 9 | RO | Seq AN Failure | When set to 1, a sequencer Auto-Negotiation failure has been detected. When set to 0, a Auto-Negotiation failure has not been detected. |
| | 17:12 | RO | KR AN Link Ready[5:0] | Provides a one-hot encoding of an_receive_idle = true and link status for the supported link as described in Clause 73.10.1. The following encodings are defined: <ul style="list-style-type: none"> • 6'b000000: 1000BASE-KX • 6'b000001: Reserved • 6'b000100: 10GBASE-KR • 6'b001000: Reserved • 6'b010000: Reserved • 6'b100000: Reserved |
| 0xC3 | 15:0 | RW | User base page low | The Auto-Negotiation TX state machine uses these bits if the AN base pages ctrl bit is set. The following bits are defined: |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|------------------------|---|
| | | | | <ul style="list-style-type: none"> [4:0]: Selector [9:5]: Echoed nonce which are set by the state machine [12:10]: Pause bits [13]: Remote Fault bit [14]: ACK which is controlled by the SM [15]: Next page bit <p>Bit 49, the PRBS bit, is generated by the Auto-Negotiation TX state machine.</p> |
| | 21:16 | RW | Override AN_TECH[5:0] | <p>Specifies an AN_TECH value to override. The following encodings are defined:</p> <ul style="list-style-type: none"> [16]: AN_TECH[0] = 1000Base-KX [17]: AN_TECH[1] = XAUI [18]: AN_TECH[2] = 10GBASE-KR [19]: AN_TECH[3] = 40G [20]: AN_TECH[4] = CR-4 [21]: AN_TECH[5] = 100G <p>You must write 0xC0[5] to 1'b1 for these overrides to take effect.</p> |
| | 25:24 | RW | Override AN_FEC[1:0] | <p>Specifies an AN_FEC value to override. The following encodings are defined:</p> <ul style="list-style-type: none"> [24]: AN_FEC [0] = Capability [25]: AN_FEC [1] = Request <p>You must write 0xC0[5] to 1'b1 for these overrides to take effect.</p> |
| | 30:28 | RW | Override AN_PAUSE[2:0] | <p>Specifies an AN_PAUSE value to override. The following encodings are defined:</p> <ul style="list-style-type: none"> [28]: AN_PAUSE[0] = Pause Ability [29]: AN_PAUSE[1] = Asymmetric Direction [30]: AN_PAUSE[2] = Reserved <p>Need to set 0xC0 bit-5 to take effect.</p> |
| 0xC4 | 31:0 | RW | User base page high | <p>The Auto-Negotiation TX state machine uses these bits if the Auto-Negotiation base pages ctrl bit is set. The following bits are defined:</p> <ul style="list-style-type: none"> [4:0]: Correspond to bits 20:16 which are TX nonce bits. [29:5]: Correspond to page bit 45:21 which are the technology ability. <p>Bit 49, the PRBS bit, is generated by the Auto-Negotiation TX state machine.</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|------|-----|---------------------|--|
| 0xC5 | 15:0 | RW | User Next page low | <p>The Auto-Negotiation TX state machine uses these bits if the Auto-Negotiation next pages ctrl bit is set. The following bits are defined:</p> <ul style="list-style-type: none"> [11]: Toggle bit [12]: ACK2 bit [13]: Message Page (MP) bit [14]: ACK controlled by the state machine [15]: Next page bit <p>For more information, refer to Clause 73.7.7.1 Next Page encodings of <i>IEEE 802.3ap-2007</i>. Bit 49, the PRBS bit, is generated by the Auto-Negotiation TX state machine.</p> |
| 0xC6 | 31:0 | RW | User Next page high | <p>The Auto-Negotiation TX state machine uses these bits if the Auto-Negotiation next pages ctrl bit is set. Bits [31:0] correspond to page bits [47:16]. Bit 49, the PRBS bit, is generated by the Auto-Negotiation TX state machine.</p> |
| 0xC7 | 15:0 | RO | LP base page low | <p>The AN RX state machine received these bits from the link partner. The following bits are defined:</p> <ul style="list-style-type: none"> [4:0] Selector [9:5] Echoed Nonce which are set by the state machine [12:10] Pause bits [12]: ACK2 bit [13]: RF bit [14]: ACK controlled by the state machine [15]: Next page bit |
| 0xC8 | 31:0 | RO | LP base page high | <p>The AN RX state machine received these bits from the link partner. The following bits are defined:</p> <ul style="list-style-type: none"> [31:30]: Reserved [29:5]: Correspond to page bits [45:21] which are the technology ability [4:0]: Correspond to bits [20:16] which are TX Nonce bits |
| 0xC9 | 15:0 | RO | LP Next page low | <p>The AN RX state machine receives these bits from the link partner. The following bits are defined:</p> <ul style="list-style-type: none"> [15]: Next page bit [14]: ACK which is controlled by the state machine [13]: MP bit [12] ACK2 bit [11] Toggle bit <p>For more information, refer to Clause 73.7.7.1 Next Page encodings of <i>IEEE 802.3ap-2007</i>.</p> |
| 0xCA | 31:0 | RO | LP Next page high | <p>The AN RX state machine receives these bits from the link partner. Bits [31:0] correspond to page bits [47:16]</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|--------------------------------|--|
| 0xCB | 24:0 | RO | AN_LP_ADV_Tech_A[24:0] | <p>Received technology ability field bits of Clause 73 Auto-Negotiation. The 10GBASE-KR PHY supports A0 and A2. The following protocols are defined:</p> <ul style="list-style-type: none"> • A0 1000BASE-KX • A1 10GBASE-KX4 • A2 10GBASE-KR • A3 40GBASE-KR4 • A4 40GBASE-CR4 • A5 100GBASE-CR10 • A24:6 are reserved <p>For more information, refer to Clause 73.6.4 and AN LP base page ability registers (7.19-7.21) of Clause 45 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 26:25 | RO | AN_LP_ADV_FEC_F[1:0] | <p>Received FEC ability bits. FEC [F0:F1] is encoded in bits D46:D47 of the base Link Codeword as described in Clause 73 AN, 73.6.5. Bit[26] corresponding to F1 is the request bit. Bit[25] corresponding to F0 is the FEC ability bit.</p> |
| | 27 | RO | AN_LP_ADV_Remote_Fault | <p>Received Remote Fault (RF) ability bits. RF is encoded in bit D13 of the base link codeword in Clause 73 AN. For more information, refer to Clause 73.6.7 and bits AN LP base page ability register AN LP base page ability registers (7.19-7.21) of Clause 45 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 30:28 | RO | AN_LP_ADV_Pause_Ability_C[2:0] | <p>Received pause ability bits. Pause (C0:C1) is encoded in bits D11:D10 of the base link codeword in Clause 73 AN as follows:</p> <ul style="list-style-type: none"> • C0 is the same as PAUSE as defined in Annex 28B • C1 is the same as ASM_DIR as defined in Annex 28B • C2 is reserved <p>For more information, refer to bits AN LP base page ability registers (7.19-7.21) of Clause 45 of <i>IEEE 802.3ap-2007</i>.</p> |
| 0xD0 | 0 | RW | Link_Training_enable | <p>When 1, enables the 10GBASE-KR start-up protocol. When 0, disables the 10GBASE-KR start-up protocol. The default value is 1. For more information, refer to Clause 72.6.10.3.1 and 10GBASE-KR PMD control register bit (1.150.1) of <i>IEEE 802.3ap-2007</i>.</p> |
| | 1 | RW | dis_max_wait_tmr | <p>When set to 1, disables the LT max_wait_timer . Used for characterization mode when setting much longer BER timer values.</p> |
| | 2 | RW | quick_mode | <p>When set to 1, only the init and preset values are used to calculate the best BER.</p> |
| | 3 | RW | pass_one | <p>When set to 1, the BER algorithm considers more than the first local minimum when searching for the lowest BER. The default value is 1.</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|---------|--|--|
| | 7:4 | RW | main_step_cnt [3:0] | Specifies the number of equalization steps for each main tap update. There are about 20 settings for the internal algorithm to test. The valid range is 1-15. The default value is 4'b0010. |
| | 11:8 | RW | prpo_step_cnt [3:0] | Specifies the number of equalization steps for each pre- and post- tap update. From 16-31 steps are possible. The default value is 4'b0001. |
| | 14:12 | RW | equal_cnt [2:0] | Adds hysteresis to the error count to avoid local minimums. The default value is 3'b010. The following encodings are defined: <ul style="list-style-type: none"> • 3'b000: 0 • 3'b001: 1 • 3'b010: 2 • 3'b100: 4 • 3'b101: 8 • 3'b110: 16 |
| | 15 | RW | disable initialize PMA on max_wait_timeout | When set to 1, does not initialize the PMA V_{OD} , pretap, posttap values upon entry into the Training_Failure state as defined in <i>Figure 72-5 of Clause 72.6.10.4.3 of IEEE 802.3ap-2007</i> . This failure occurs when the max_wait_timer_done timeout is reached setting the Link Training failure bit (0xD2[3]). Used during UNH-IOL testing. When set to 0, initializes the PMA values upon entry into Training_Failure state. |
| | 16 | RW | Ovride LP Coef enable | When set to 1, overrides the link partner's equalization coefficients; software changes the update commands sent to the link partner TX equalizer coefficients. When set to 0, uses the Link Training logic to determine the link partner coefficients. Used with 0xD1 bit-4 and 0xD4 bits[7:0]. |
| | 17 | RW | Ovride Local RX Coef enable | When set to 1, overrides the local device equalization coefficients generation protocol. When set, the software changes the local TX equalizer coefficients. When set to 0, uses the update command received from the link partner to determine local device coefficients. Used with 0xD1 bit-8 and 0xD4 bits[23:16]. The default value is 1. |
| | 19:18 | RM W | Reserved | You should not modify these bits. To update this register, first read the value of this register then change only the value for bits that are not reserved. |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|-----------------------|--|
| | 22:20 | RW | rx_ctle_mode | RX CTLE mode in the Link Training algorithm. The default value is 3'b000. The following encodings are defined: <ul style="list-style-type: none"> 3'b000: CTLE tuning in link training is disabled. Retains user set value of CTLE. 3'b001: Reserved. 3'b010: Reserved. 3'b011: CTLE tuning in link training is enabled. 3'b100 to 3'b111: reserved. |
| | 23 | RW | vod_up | When set to 1, V _{OD} is trained to high values. The default is set to 0 to save power and reduce crosstalk on the link. |
| | 26:24 | RW | rx_dfe_mode | RX DFE mode in the link training algorithm. The default value is 3'b000. The following bits are defined: <ul style="list-style-type: none"> 3'b000: DFE adaptation in link training is disabled 3'b001: Reserved 3'b010: DFE is triggered at the end of link training 3'b011: DFE is triggered at the end of V_{OD}, Post tap and Pre-tap training 3'b100 to 3'b111: Reserved |
| | 28 | RW | max_mode | When set to 1, link training operates in maximum TX equalization mode. Modifies the link training algorithm to settle on the max pretap and max V _{OD} if the BER counter reaches the maximum for all values. Link training settles on the max_post_step for the posttap value. |
| | 31:29 | RW | max_post_step | Number of TX posttap steps from the initialization state when in max_mode. |
| 0xD1 | 0 | RW | Restart Link training | When set to 1, resets the 10GBASE-KR start-up protocol. When set to 0, continues normal operation. This bit self clears. For more information, refer to the state variable mr_restart_training as defined in Clause 72.6.10.3.1 and 10GBASE-KR PMD control register bit (1.150.0) <i>IEEE 802.3ap-2007</i> . |
| | 4 | RW | Updated TX Coef new | When set to 1, there are new link partner coefficients available to send. The LT logic starts sending the new values set in 0xD4 bits[7:0] to the remote device. When set to 0, continues normal operation. This bit self clears. Must enable this override in 0xD0 bit16. |
| | 8 | RW | Updated RX coef new | When set to 1, new local device coefficients are available. The LT logic changes the local TX equalizer coefficients as specified in 0xD4 bits[23:16]. When set to 0, continues normal operation. This bit self clears. Must enable the override in 0xD0 bit17. |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-----|-----|--|--|
| 0xD2 | 0 | RO | Link Trained - Receiver status | When set to 1, the receiver is trained and is ready to receive data. When set to 0, receiver training is in progress. For more information, refer to the state variable rx_trained as defined in Clause 72.6.10.3.1 and bit 10GBASE-KR PMD control register bit 10GBASE_KR PMD status register bit (1.151.0) of <i>IEEE 802.3ap-2007</i> . |
| | 1 | RO | Link Training Frame lock | When set to 1, the training frame delineation has been detected. When set to 0, the training frame delineation has not been detected. For more information, refer to the state variable frame_lock as defined in Clause 72.6.10.3.1 and 10GBASE_KR PMD status register bit 10GBASE_KR PMD status register bit (1.151.1) of <i>IEEE 802.3ap-2007</i> . |
| | 2 | RO | Link Training Start-up protocol status | When set to 1, the start-up protocol is in progress. When set to 0, start-up protocol has completed. For more information, refer to the state training as defined in Clause 72.6.10.3.1 and 10GBASE_KR PMD status register bit (1.151.2) of <i>IEEE 802.3ap-2007</i> . |
| | 3 | RO | Link Training failure | When set to 1, a training failure has been detected. When set to 0, a training failure has not been detected. For more information, refer to the state variable training_failure as defined in Clause 72.6.10.3.1 and bit 10GBASE_KR PMD status register bit (1.151.3) of <i>IEEE 802.3ap-2007</i> . |
| | 4 | RO | Link Training Error | When set to 1, excessive errors occurred during Link Training. When set to 0, the BER is acceptable. |
| | 5 | RO | Link Training Frame lock Error | When set to 1, indicates a frame lock was lost during Link Training. If the tap settings specified by the fields of 0xD5 are the same as the initial parameter value, the frame lock error was unrecoverable. |
| | 6 | RO | CTLE Frame Lock Loss | When set to 1, indicates that fram lock was lost at some point during CTLE link training. |
| | 7 | RO | CTLE Tuning Error | When set to 1, indicates that CTLE did not achieve best results because the BER counter reached the maximum value for each step of CTLE tuning. |
| 0xD3 | 9:0 | RW | ber_time_frames | Specifies the number of training frames to examine for bit errors on the link for each step of the equalization settings. Used only when ber_time_k_frames is 0. The following values are defined: <ul style="list-style-type: none"> • A value of 2 is about 10³ bytes • A value of 20 is about 10⁴ bytes • A value of 200 is about 10⁵ bytes The default value for simulation is 2'b11. The default value for hardware is 0. |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|----------|----------------------------|---|
| | 19:10 | RW | ber_time_k_frames | <p>Specifies the number of thousands of training frames to examine for bit errors on the link for each step of the equalization settings. Set <i>ber_time_m_frames</i> = 0 for time/bits to match the following values:</p> <ul style="list-style-type: none"> • A value of 3 is about 10^7 bits = about 1.3 ms • A value of 25 is about 10^8 bits = about 11 ms • A value of 250 is about 10^9 bits = about 110 ms <p>The default value for simulation is 0. The default value for hardware is 0x15.</p> |
| | 29:20 | RW | ber_time_m_frames | <p>Specifies the number of millions of training frames to examine for bit errors on the link for each step of the equalization settings. Set <i>ber_time_k_frames</i> = 4'd1000 = 0x3E8 for time/bits to match the following values:</p> <ul style="list-style-type: none"> • A value of 3 is about 10^{10} bits = about 1.3 seconds • A value of 25 is about 10^{11} bits = about 11 seconds • A value of 250 is about 10^{12} bits = about 110 seconds |
| 0xD4 | 5:0 | RO or RW | LD coefficient update[5:0] | <p>Reflects the contents of the first 16-bit word of the training frame sent from the local device control channel. Normally, the bits in this register are read-only; however, when you override training by setting the Override Coef enable control bit, these bits become writeable. The following fields are defined:</p> <ul style="list-style-type: none"> • [5:4]: Coefficient (+1) update <ul style="list-style-type: none"> • 2'b11: Reserved • 2'b01: Increment • 2'b10: Decrement • 2'b00: Hold • [3:2]: Coefficient (0) update (same encoding as [5:4]) • [1:0]: Coefficient (-1) update (same encoding as [5:4]) <p>For more information, refer to bit 10G BASE-KR LD coefficient update register bits (1.154.5:0) in Clause 45.2.1.80.3 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 6 | RO or RW | LD Initialize Coefficients | <p>When set to 1, requests the link partner coefficients be set to configure the TX equalizer to its INITIALIZE state. When set to 0, continues normal operation. For more information, refer to 10G BASE-KR LD coefficient update register bits (1.154.12) in Clause 45.2.1.80.3 and Clause 72.6.10.2.3.2 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 7 | RO or RW | LD Preset Coefficients | <p>When set to 1, requests the link partner coefficients be set to a state where equalization is turned off. When set to 0 the link operates normally. For more information, refer to bit 10G BASE-KR LD coefficient update register bit (1.154.13) in Clause 45.2.1.80.3 and Clause 72.6.10.2.3.2 of <i>IEEE 802.3ap-2007</i>.</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|----------|---|--|
| | 13:8 | RO | LD coefficient status[5:0] | <p>Status report register for the contents of the second, 16-bit word of the training frame most recently sent from the local device control channel. The following fields are defined:</p> <ul style="list-style-type: none"> [5:4]: Coefficient (post-tap) <ul style="list-style-type: none"> 2'b11: Maximum 2'b01: Minimum 2'b10: Updated 2'b00: Not updated [3:2]: Coefficient (0) (same encoding as [5:4]) [1:0]: Coefficient (pre-tap) (same encoding as [5:4]) <p>For more information, refer to bit 10G BASE-KR LD status report register bit (1.155.5:0) in Clause 45.2.1.81 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 14 | RO | Link Training ready - LD Receiver ready | <p>When set to 1, the local device receiver has determined that training is complete and is prepared to receive data. When set to 0, the local device receiver is requesting that training continue. Values for the receiver ready bit are defined in Clause 72.6.10.2.4.4. For more information refer to For more information, refer to bit 10G BASE-KR LD status report register bit (1.155.15) in Clause 45.2.1.81 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 21:16 | RO or RW | LP coefficient update[5:0] | <p>Reflects the contents of the first 16-bit word of the training frame most recently received from the control channel.</p> <p>Normally the bits in this register are read only; however, when training is disabled by setting low the KR Training enable control bit, these bits become writeable. The following fields are defined:</p> <ul style="list-style-type: none"> [5: 4]: Coefficient (+1) update <ul style="list-style-type: none"> 2'b11: Reserved 2'b01: Increment 2'b10: Decrement 2'b00: Hold [3:2]: Coefficient (0) update (same encoding as [5:4]) [1:0]: Coefficient (-1) update (same encoding as [5:4]) <p>For more information, refer to bit 10G BASE-KR LP coefficient update register bits (1.152.5:0) in Clause 45.2.1.78.3 of <i>IEEE 802.3ap-2007</i>.</p> |
| | 22 | RO or RW | LP Initialize Coefficients | <p>When set to 1, the local device transmit equalizer coefficients are set to the INITIALIZE state. When set to 0, normal operation continues. The function and values of the initialize bit are defined in Clause 72.6.10.2.3.2. For more information, refer to bit 10G BASE-KR LP coefficient update register bits (1.152.12) in Clause 45.2.1.78.3 of <i>IEEE 802.3ap-2007</i>.</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|----------|----------------------------|---|
| | 23 | RO or RW | LP Preset Coefficients | When set to 1, The local device TX coefficients are set to a state where equalization is turned off. Preset coefficients are used. When set to 0, the local device operates normally. The function and values of the preset bit is defined in 72.6.10.2.3.1. The function and values of the initialize bit are defined in Clause 72.6.10.2.3.2. For more information, refer to bit 10G BASE-KR LP coefficient update register bits (1.152.13) in Clause 45.2.1.78.3 of IEEE 802.3ap-2007. |
| | 29:24 | RO | LP coefficient status[5:0] | Status report register reflects the contents of the second, 16-bit word of the training frame most recently received from the control channel. The following fields are defined: <ul style="list-style-type: none"> [5:4]: Coefficient (+1) <ul style="list-style-type: none"> 2'b11: Maximum 2'b01: Minimum 2'b10: Updated 2'b00: Not updated [3:2]: Coefficient (0) (same encoding as [5:4]) n [1:0]: Coefficient (-1) (same encoding as [5:4]) For more information, refer to bit 10G BASE-KR LP status report register bits (1.153.5:0) in Clause 45.2.1.79 of IEEE 802.3ap-2007. |
| | 30 | RO | LP Receiver ready | When set to 1, the link partner receiver has determined that training is complete and is prepared to receive data. When set to 0, the link partner receiver is requesting that training continue. <p>Values for the receiver ready bit are defined in Clause 72.6.10.2.4.4. For more information, refer to bit 10G BASE-KR LP status report register bits (1.153.15) in Clause 45.2.1.79 of IEEE 802.3ap-2007.</p> |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|----------------------------|---|
| 0xD5 | 5:0 | R | LT V _{OD} setting | Stores the most recent V _{OD} setting that LT specified using the Transceiver Reconfiguration Controller IP core. It reflects Link Partner commands to fine-tune the V _{OD} . |
| | 12:8 | R | LT Post-tap setting | Stores the most recent post-tap setting that LT specified using the Transceiver Reconfiguration Controller IP core. It reflects Link Partner commands to fine-tune the TX pre-emphasis taps. |
| | 19:16 | R | LT Pre-tap setting | Stores the most recent pre-tap setting that LT specified using the Transceiver Reconfiguration Controller IP core. It reflects Link Partner commands to fine-tune the TX pre-emphasis taps. |
| | 23:20 | R | RXEQ CTLE Setting | Stores the most recent CTLE setting sent to the Transceiver Reconfiguration IP Core during RX Equalization. |
| | 25:24 | R | RXEQ CTLE Mode | Stores the most recent CTLE mode that CTLE specified using the Transceiver Reconfiguration IP Core during RX Equalization. |
| | 27:26 | R | RXEQ DFE Mode | Stores the most recent DFE setting sent to the Transceiver Reconfiguration IP Core during RX Equalization. |
| 0xD6 | 5:0 | RW | LT VODMAX ovrđ | Override value for the VMAXRULE parameter. When enabled, this value substitutes for the VMAXRULE to allow channel-by-channel override of the device settings. This only effects the local device TX output for the channel specified. This value must be greater than the INITMAINVAL parameter for proper operation. Note this will also override the PREMAINVAL parameter value. |
| | 6 | RW | LT VODMAX ovrđ Enable | When set to 1, enables the override value for the VMAXRULE parameter stored in the LT VODMAX ovrđ register field. |
| | 13:8 | RW | LT VODMin ovrđ | Override value for the VODMINRULE parameter. When enabled, this value substitutes for the VMINRULE to allow channel-by-channel override of the device settings. This override only effects the local device TX output for this channel. The value to be substituted must be less than the INITMAINVAL parameter and greater than the VMINRULE parameter for proper operation. |
| | 14 | RW | LT VODMin ovrđ Enable | When set to 1, enables the override value for the VODMINRULE parameter stored in the LT VODMin ovrđ register field. |

| Word Addr | Bit | R/W | Name | Description |
|-----------|-------|-----|----------------------|---|
| | 20:16 | RW | LT VPOST ovrd | <p>Override value for the VPOSTRULE parameter. When enabled, this value substitutes for the VPOSTRULE to allow channel-by-channel override of the device settings. This override only effects the local device TX output for this channel.</p> <p>The value to be substituted must be greater than the INITPOSTVAL parameter for proper operation.</p> |
| | 21 | RW | LT VPOST ovrd Enable | When set to 1, enables the override value for the VPOSTRULE parameter stored in the LT VPOST ovrd register field. |
| | 27:24 | RW | LT VPre ovrd | <p>Override value for the VPRERULE parameter. When enabled, this value substitutes for the VPOSTRULE to allow channel-by-channel override of the device settings. This override only effects the local device TX output for this channel.</p> <p>The value greater than the INITPREVAL parameter for proper operation.</p> |
| | 28 | RW | LT VPre ovrd Enable | When set to 1, enables the override value for the VPRERULE parameter stored in the LT VPre ovrd register field. |

2014.12.15

UG-01088



Subscribe



Send Feedback

40- and 100-Gbps Ethernet MAC and PHY MegaCore Function User Guide Revision History

Table D-1: Document Revision History

Summarizes the new features and changes in the user guide for the 40- and 100-Gbps Ethernet MAC and PHY MegaCore function.

| Date | ACDS Version | Changes |
|------------|--------------|--|
| 2014.12.15 | v14.1 | <ul style="list-style-type: none"> • Updated release-specific information for the software release v14.1. • Updated for new Quartus II IP Catalog, which replaces the MegaWizard Plug-In Manager starting in the Quartus II software v14.0. Changes are located primarily in Getting Started. Reordered the chapter to accommodate the new descriptions. • Removed OpenCore Plus exception for 40GBASE-KR4 feature. OpenCore Plus evaluation is available for this feature in the IP core version 14.0 and later. • Moved licensing information and the description of the OpenCore Plus evaluation feature to Getting Started. • Corrected description of IPG_DEL_PERIOD register in MAC Feature Configuration Registers. • Corrected statements that the pause signals are available only in 40-100GbE IP core variations without adapters. Pause signals have been available in IP core variations with and without adapters since the IP core v12.0. Changes are located in Pause Control and Generation Interface, in Conditions Triggering XOFF Frame Transmission, in Conditions Triggering XON Frame Transmission, in Signals of MAC and PHY Variations With Adapters, and in Pause Registers. • Clarified that while the IP core is paused it can still transmit pause frames, in Congestion and Flow Control Using Pause Frames. • Clarified reset requirements in Resets. |

© 2014 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

ISO
9001:2008
Registered



| Date | ACDS Version | Changes |
|------------|-----------------------|---|
| | | <ul style="list-style-type: none"> Added link to Low Latency 40-100GbE IP core user guide in About the 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function and in Device Family Support. Added new 10GBASE-KR Registers appendix to provide the 10GBASE-KR PHY register information that forms the basis for the 40GBASE-KR4 Registers descriptions. This information is correct at the time of publication but may lag updates to the source document, the Altera Transceiver PHY IP Core User Guide. Removed mention of register 0x0C2 bit [8] from 40GBASE-KR4 Registers on page 3-92 because this register bit is now described in the 10GBASE-KR PHY register listing. Clarified that the <code>tx_ready</code> signal might an inadequate indication that the IP core is ready to accept and handle data on the TX client interface with adapters. Following reset, the application should not send data on this interface until after the <code>lanes_deskewed</code> signal is asserted, indicating the IP core can safely transmit packets on the Ethernet link. In the case of the TX client interface without adapters, the application should also wait until after the <code>lanes_deskewed</code> signal is asserted. Notice appears in 40-100GbE IP Core TX Data Bus with Adapters (Avalon-ST Interface) and in 40-100GbE IP Core TX Data Bus Without Adapters (Custom Streaming Interface). Corrected lists of testbench files in Testbenches with Adapters and in Testbenches without Adapters. The new lists are correct for the IP core v13.1 and later. Corrected assorted typos. |
| March 2014 | v13.1 (2014.03.04) | <ul style="list-style-type: none"> Corrected device support list: <ul style="list-style-type: none"> To include support for Stratix V GS devices. To specify Final support for Stratix V devices. Included "40-100GbE IP Core TX Data Bus with Adapters (Avalon-ST Interface)" section inadvertently omitted from previous version of user guide. In "Frame Check Sequence (CRC-32) Insertion" section, corrected <code>CRC_CONFIG</code> register bit that controls TX CRC insertion. Removed inconsistent statement in "40-100GbE IP core CRC Checking" section that incorrectly indicated that CRC aligns with EOP. In "Statistics Counters Interface" section, fixed descriptions of the size-based frame counting signals. In "Order of Transmission" section, improved figures showing byte order on the Avalon-ST interface with and without preamble pass-through enabled. In "About the 40- and 100-Gbps Ethernet MAC and PHY MegaCore Function" chapter, removed mention of 4 and 20 virtual lanes, and replaced with information correct for the Avalon-ST client-side interface. The virtual lanes are relevant only on the Ethernet link. In "OpenCore Plus Time-Out Behavior" section, removed mention of nonexistent <code>local_ready</code> signal. |

| Date | ACDS Version | Changes |
|---------------|--------------------|--|
| | | <ul style="list-style-type: none"> In "Simulating the 40-100GbE IP Core with the Testbenches" section, removed option to include or exclude waveform generation in the simulation scripts. The scripts do include this option in pre-13.1 releases. Corrected topics to include links to both internal (other sections) and external (other documents and websites). Previously only one or the other appeared in each section in the document. Fixed assorted typos. |
| November 2013 | v13.1 (2013.11.04) | <ul style="list-style-type: none"> Added new 40GBASE-KR4 support with FEC option, including new parameters, signals, registers, testbench, and example design. Added new parameter for Synchronous Ethernet support option to separate the reference input clocks for the RX CDR PLL and the TX PLL and make the RX recovered clock visible. Added the following new signals: <ul style="list-style-type: none"> Link fault signals visible in duplex variations: <code>remote_fault_status</code> and <code>local_fault_status</code>. PHY status signals visible in MAC+PHY variations: <code>tx_lanes_stable</code> and <code>lanes_deskewed</code>. New clock signals for SyncE variations: <code>rx_clk_ref</code>, <code>tx_clk_ref</code>, and <code>rx_recovered_clk</code>. New 40GBASE-KR4 signals for analog reconfiguration. Updated testbench descriptions to describe the new streamlined testbenches and the new 40GBASE-KR4 example design and testbench. Updated descriptions of testbenches and example designs to clarify the user no longer needs to configure the DUT with a specific name and clock rate. Updated resource utilization numbers. Corrected module names in 40GbE IP core resource utilization tables. |
| July 2013 | 1.3(v13.0) | <ul style="list-style-type: none"> Corrected signal widths and descriptions in Figure 3-2 on page 3-2 and Figure 3-3 on page 3-3. Corrected Avalon-ST client interface signal widths in Figure 3-5 on page 3-8, Table 3-2 on page 3-8, Figure 3-20 on page 3-23, and Table 3-6 on page 3-23. Renamed and reorganized "Signals" on page 3-40. |

| Date | ACDS Version | Changes |
|-----------|-----------------------------|--|
| June 2013 | 1.3(v13.0 software release) | <ul style="list-style-type: none"> • Updated for use with version 13.0 of the Quartus II software and the MegaWizard Plug-In Manager. • Added preamble pass-through option. This change affects various sections in “TX Datapath” on page 3–5 and “RX Datapath” on page 3–20, and includes the addition of a new <code>Preamble Pass-Through Configuration</code> register (offset 0x125), described in “MAC Feature Configuration Registers” on page 3–62. • Added transmitter average inter-packet gap (IPG) adjustment option. This change affects “Inter-Packet Gap Generation and Insertion” on page 3–7 and includes the addition of two new registers (<code>IPG_DEL_PERIOD</code> at offset 0x126 and <code>IPG_DEL_ENABLE</code> at offset 0x127), described in “R**MAC Feature Configuration Registers” on page 3–62. • Added new section “MAC Feature Configuration Registers” on page 3–62 for the new registers. Moved the description of the <code>CRC_CONFIG</code> register (offset 0x123) to this new section. • Reorganized Chapter 2, Getting Started to remove non-IP core specific information. • Moved instructions for IP core initialization from “Software Interface: Registers” to new section “Initializing the IP Core” on page 2–30. • Updated resource utilization numbers in “Performance and Resource Utilization” on page 1–5. • Clarified device speed grades per device family variant, in “Device Speed Grade Support” on page 1–4. • Clarified definitions of cut-through, store and forward, and promiscuous receive modes, in “40-100GbE IP Core Modes of Operation” on page 3–34. • Clarified destination address checking controls in “Address Checking” on page 3–22 and in “MAC Address Registers” on page 3–64. • Removed Appendix B, Address Map Changes for 12.1, and moved information to Document Revision History entry for Quartus II software v12.1 release. The update added registers but did not change register names or offsets of existing registers. • Fixed descriptions of allowed transceiver reference clock frequencies and PCS clock frequencies: <ul style="list-style-type: none"> • Clarified that 644 MHz is not an allowed frequency for CAUI–4 variations, despite its presence in the parameter editor as an apparently allowed value for the PHY Reference Frequency parameter. • Provided correct PCS clock frequency for CAUI–4 variations. • Provided correct frequencies for 24.24 Gbps variations. • Modified “100GbE IP Core without Adapters” on page 3–14 to clarify that only two SOPs can occur on the TX custom streaming client interface during the same clock cycle. • Added instructions for using the example design in Appendix A, 13.0 Example Design. • Improved descriptions of various signals and fixed typos. |

| Date | ACDS Version | Changes |
|---------------|--------------|---|
| December 2012 | 1.2 | <ul style="list-style-type: none"> • Updated Slowest Supported Device Speed Grades table on page 1–4: <ul style="list-style-type: none"> • Supported speed grades for the Arria V GZ device updated to I3L, C3. • Supported speed grades for the Stratix V device updated to I3, C3. |
| November 2012 | 1.1 | <ul style="list-style-type: none"> • Updated for use with version 12.1 of the Quartus II software and the MegaWizard Plug-In Manager. • Updated address map. • Updated device family support, including the addition of the Arria V GZ device. • Updated interfaces for 40-100GbE IP cores with adapters and without adapters to include the external reconfiguration controller. • Assignment of the <code>FAST_SIMULATION</code> parameter has been updated. For more information, refer to “MegaWizard Plug-In Manager Flow” on page 2–2. • Updated RTL hierarchy and directory structure. • Clocking revisions: <ul style="list-style-type: none"> • <code>clk_din</code> has been replaced by <code>clk_txmac</code> • <code>clk_dout</code> has been replaced by <code>clk_rxmac</code> |

| Date | ACDS Version | Changes |
|------|--------------|---|
| | | <ul style="list-style-type: none"> • If you are transitioning from an earlier version of the IP core, you must complete the following steps: <ul style="list-style-type: none"> • Generate the 12.1 release from the MegaWizard Plug-In Manager into a new project directory; this will generate a .qip file to include in the Quartus II software project, along with the required Verilog HDL files. For more information, refer to “Getting Started” on page 2–1. • Quartus II software assignments from previous releases that reference the internal IP hierarchy (such as logic lock regions) must be updated for changes in the internal hierarchy. Hierarchy changes include the following: <ul style="list-style-type: none"> • The top-level Verilog in the synthesis file set has one level below it named <i><name of your IP instance>_inst</i>. All further instances begin under the <i><name of your IP instance>_inst</i> directory. • The <code>pcs_tx</code>, <code>pcs_rx</code>, and <code>phy_csr</code> instances now reside under phy/phy_pcs • The <code>pcs_tx</code>, <code>pcs_rx</code>, and <code>phy_csr</code> instances now reside under phy/phy_pcs • Nodes for the Stratix V device PMA now reside under phy/pma/pma_bridge • Generate example designs from the 12.1 release of the MegaWizard Plug-In Manager for a complete set of new Quartus II software assignments • The reconfiguration controller must be instantiated and connected to the IP core. For more information, refer to “External Reconfiguration Controller” on page 3–26 and “12.1 Example Design” on page A–1 • Note the <code>GXB_0PPM_CORECLK</code> and <code>GXB_0PPM_CORE_CLOCK</code> Quartus II software settings are no longer required. • Feature additions: <ul style="list-style-type: none"> • 40GbE Lower Rate 24.24 Gbps MAC and PHY • 100GbE CAUI–4 PHY • RX Automatic Pad Removal Control • Pause Control Frames Filtering Control • Updated or added signals: <ul style="list-style-type: none"> • Top-level output and input high-speed serial lines from the transceivers • External reconfiguration controller • TX MAC to PHY connections • RX MAC to PHY connections • Added registers: <ul style="list-style-type: none"> • PAD_CONFIG at offset 0x124 • Low Latency PHY IP Core registers for CAUI–4 at offsets 0x800–0x9FF, 0xA00–0xBFF, 0xC00–0xDFF, and 0xE00–0xFFF. |

| | | |
|------------------|------------|---|
| <p>June 2012</p> | <p>1.0</p> | <ul style="list-style-type: none"> • Updated for use with version 12.0 of the Quartus II software. • Updated address map. • Updated device family support. • Updated interfaces for 40-100GbE IP cores with adapters and without adapters. Additional interfaces include: <ul style="list-style-type: none"> • MAC and PHY asynchronous resets. • MAC to PHY connections. • Lane-to-lane deskew. • Statistics counters increment vectors. • Link fault signaling, including remote fault and local fault. • Updated RTL hierarchy, directory structure, and wrapper reorganization. • Feature additions: <ul style="list-style-type: none"> • Controllable FCS (CRC) insertion and removal. • Cut-through mode runt removal. • PCS BER monitor. • PCS test pattern generation and check. • Reduced RX destination MAC address checking from 16 addresses to 1 address. • Preserved FCS result. • Statistics counters implemented as a synthesis option. • Updated or added software registers: <ul style="list-style-type: none"> • Test pattern counter. • Link fault signaling. • CRC configuration. • MAC hardware error. • MAC and PHY resets. • PCS hardware error. • BER monitor. • Test mode. • MAC address. • Statistics counters: roll-overs and increment vectors. • Clocking revisions: <ul style="list-style-type: none"> • <code>clk50</code> has been removed and replaced by <code>clk_status</code>. • <code>clk_din/clk_txmac</code> and <code>clk_dout/clk_rxmac</code> have been added as TX and RX input clocks. • Additional parameters: <code>STATS_CNTRS_OPTION</code> and <code>FAST_SIMULATION</code>. • Updated testbenches and simulation examples. • Updated reset signals and reset bits. • Additional information regarding oversized frames. |
|------------------|------------|---|

History

| | | |
|----------------|--------------|--|
| November 2011 | Early Access | <ul style="list-style-type: none"> • Corrected the following issues in the MegaCore function: <ul style="list-style-type: none"> • Corrected sequence ordered set encoding in PCS. • Corrected error control block encoding in PCS. • Corrected pause logic to accomodate multiple pause requests. • Timing performance improvements: <ul style="list-style-type: none"> • Optimized RTL for better timing performance. • Clocking revisions: <ul style="list-style-type: none"> • <code>clk_status</code> and <code>clk_csr</code> now support 100 MHz operation in Stratix V devices for PHY IP calibration. • Added Stratix V resource utilization information. • Updated definition <code>1 <n> _rx_ready</code> to include fact that the RX MAC can only be backpressured for a limited number of cycles; consequently, the application should be able to accept a continuous data stream. |
| September 2011 | Early Access | <ul style="list-style-type: none"> • Feature additions: <ul style="list-style-type: none"> • Added optional adapters that guarantee the start of packet is always in lane 0. • Added full statistics counters module. • Modified and improved register map. • Provided separate product ID and ordering code for 40GbE and 100GbE MAC and PHY. • Corrected the following issues in the MegaCore function: <ul style="list-style-type: none"> • Data is no longer reversed in the adapters. • The multicast address is used for the multicast pause frames. • The pause state machine loads new pause times as required. • Short frames on the TX datapath are converted correctly from 8 words to 5 words. • Updated resource utilization numbers. • User guide enhancements: <ul style="list-style-type: none"> • Rewrote and added new sections and drawings. • Added default values of registers after reset. • Combined 40GbE and 100GbE in one user guide. • Corrected description of <code>PMD_CMD_CONFIG</code> bit. Writing a 1 enables the PMD. • Corrected description of <code>PMD_CMD_CONFIG</code> bit. Writing a 1 enables the PMD. • Corrected description of <code>RX_AGGREGATE</code> register. Definitions for bit[0] and bit[1] were reversed. • Corrected description of <code>pause_quanta</code>. • Corrected descriptions of Figure 3–10 on page 3–11 through Figure 3–12 on page 3–12. • Removed <code>din_in_packet</code> from Figure 3–2 on page 3–3. |
| July 2010 | Early Access | Initial early access release. |

How to Contact Altera

Table D-2: How to Contact Altera

To locate the most up-to-date information about Altera products, refer to this table. You can also contact your local Altera sales office or sales representative.

| Contact | Contact Method | Address |
|--|----------------|--|
| Technical support | Website | www.altera.com/support |
| Technical training | Website | www.altera.com/training |
| | Email | custrain@altera.com |
| Product literature | Website | www.altera.com/literature |
| Nontechnical support: general | Email | nacomp@altera.com |
| Nontechnical support: software licensing | Email | authorization@altera.com |

Related Information

- www.altera.com/support
- www.altera.com/training
- custrain@altera.com
- www.altera.com/literature
- nacomp@altera.com
- authorization@altera.com

Typographic Conventions

Table D-3: Typographic Conventions

Lists the typographic conventions this document uses.

| Visual Cue | Meaning |
|---|--|
| Bold Type with Initial Capital Letters | Indicate command names, dialog box titles, dialog box options, and other GUI labels. For example, Save As dialog box. For GUI elements, capitalization matches the GUI. |
| bold type | Indicates directory names, project names, disk drive names, file names, file name extensions, software utility names, and GUI labels. For example, \ qdesigns directory, D: drive, and chiptrip.gdf file. |

| Visual Cue | Meaning |
|---|--|
| <i>Italic Type with Initial Capital Letters</i> | Indicate document titles. For example, <i>Stratix V Design Guidelines</i> . |
| <i>italic type</i> | Indicates variables. For example, $n + 1$. Variable names are enclosed in angle brackets (< >). For example, <file name> and <project name> . pdf file. |
| Initial Capital Letters | Indicate keyboard keys and menu names. For example, the Delete key and the Options menu. |
| “Subheading Title” | Quotation marks indicate references to sections in a document and titles of Quartus II Help topics. For example, “Typographic Conventions.” |
| Courier type | Indicates signal, port, register, bit, block, and primitive names. For example, data1, tdi, and input. The suffix n denotes an active-low signal. For example, resetn. Indicates command line commands and anything that must be typed exactly as it appears. For example, c:\qdesigns\tutorial\chiptrip.gdf. Also indicates sections of an actual file, such as a Report File, references to parts of files (for example, the AHDL keyword SUBDESIGN), and logic function names (for example, TRI). |
| 1., 2., 3., and a., b., c., and so on | Numbered steps indicate a list of items when the sequence of the items is important, such as the steps listed in a procedure. |
| • | Bullets indicate a list of items when the sequence of the items is not important. |

The **Subscribe** button links to the Email Subscription Management Center page of the Altera website, where you can sign up to receive update notifications for Altera documents.

The **Feedback** icon allows you to submit feedback to Altera about the document. Methods for collecting feedback vary as appropriate for each document.

Related Information

[Email Subscription Management Center](#)