

LogiCORE™ IP Ethernet 1000BASE-X PCS/PMA or SGMII v9.1

User Guide

UG155 March 24, 2008





Xilinx is disclosing this Specification to you solely for use in the development of designs to operate on Xilinx FPGAs. Except as stated herein, none of the Specification may be copied, reproduced, distributed, republished, downloaded, displayed, posted, or transmitted in any form or by any means including, but not limited to, electronic, mechanical, photocopying, recording, or otherwise, without the prior written consent of Xilinx. Any unauthorized use of this Specification may violate copyright laws, trademark laws, the laws of privacy and publicity, and communications regulations and statutes.

Xilinx does not assume any liability arising out of the application or use of the Specification; nor does Xilinx convey any license under its patents, copyrights, or any rights of others. You are responsible for obtaining any rights you may require for your use or implementation of the Specification. Xilinx reserves the right to make changes, at any time, to the Specification as deemed desirable in the sole discretion of Xilinx. Xilinx assumes no obligation to correct any errors contained herein or to advise you of any correction if such be made. Xilinx will not assume any liability for the accuracy or correctness of any engineering or technical support or assistance provided to you in connection with the Specification.

THE SPECIFICATION IS PROVIDED "AS IS" WITH ALL FAULTS, AND THE ENTIRE RISK AS TO ITS FUNCTION AND IMPLEMENTATION IS WITH YOU. YOU ACKNOWLEDGE AND AGREE THAT YOU HAVE NOT RELIED ON ANY ORAL OR WRITTEN INFORMATION OR ADVICE, WHETHER GIVEN BY XILINX, OR ITS AGENTS OR EMPLOYEES. XILINX MAKES NO OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, REGARDING THE SPECIFICATION, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NONINFRINGEMENT OF THIRD-PARTY RIGHTS.

IN NO EVENT WILL XILINX BE LIABLE FOR ANY CONSEQUENTIAL, INDIRECT, EXEMPLARY, SPECIAL, OR INCIDENTAL DAMAGES, INCLUDING ANY LOST DATA AND LOST PROFITS, ARISING FROM OR RELATING TO YOUR USE OF THE SPECIFICATION, EVEN IF YOU HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE TOTAL CUMULATIVE LIABILITY OF XILINX IN CONNECTION WITH YOUR USE OF THE SPECIFICATION, WHETHER IN CONTRACT OR TORT OR OTHERWISE, WILL IN NO EVENT EXCEED THE AMOUNT OF FEES PAID BY YOU TO XILINX HEREUNDER FOR USE OF THE SPECIFICATION. YOU ACKNOWLEDGE THAT THE FEES, IF ANY, REFLECT THE ALLOCATION OF RISK SET FORTH IN THIS AGREEMENT AND THAT XILINX WOULD NOT MAKE AVAILABLE THE SPECIFICATION TO YOU WITHOUT THESE LIMITATIONS OF LIABILITY.

The Specification is not designed or intended for use in the development of on-line control equipment in hazardous environments requiring fail-safe controls, such as in the operation of nuclear facilities, aircraft navigation or communications systems, air traffic control, life support, or weapons systems ("High-Risk Applications"). Xilinx specifically disclaims any express or implied warranties of fitness for such High-Risk Applications. You represent that use of the Specification in such High-Risk Applications is fully at your risk.

© 2004-2007 Xilinx, Inc. All rights reserved. XILINX, the Xilinx logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

Revision History

The following table shows the revision history for this document.

Date	Doc Version	Revision
09/30/04	1.0	Initial Xilinx release.
04/28/05	2.0	Updated to Xilinx tools 7.1i SP2, support for Virtex-4 Rocket IO.
01/18/06	3.0	Updated to Xilinx tools 8.1i SP1 for 7.0 release, added new chapter for dynamic switching.
07/13/06	4.0	Updated to core version 7.1; Xilinx tools 8.2i.
10/23/06	5.0	Updated to core version 8.0, support for Virtex-5 LXT and Spartan-3A families.
02/15/07	6.0	Updated to core version 8.1, Xilinx tools 9.1i.
08/08/07	7.0	Updated to core version 9.0, Xilinx tools 9.2i.
03/24/08	8.0	Updated to core version 9.1, Xilinx tools 10.1.

Table of Contents

Schedule of Figures	9
Schedule of Tables	13
Preface: About This Guide	
Guide Contents	17
Conventions	18
Typographical.....	18
Online Document.....	19
Chapter 1: Introduction	
About the Core	21
Designs Using RocketIO Transceivers.....	21
Recommended Design Experience	21
Additional Core Resources	21
Related Xilinx Ethernet Products and Services.....	22
Specifications.....	22
Technical Support	22
Feedback	22
Ethernet 1000BASE-X PCS/PMA or SGMII Core.....	22
Document.....	22
Chapter 2: Core Architecture	
System Overview	23
Ethernet 1000BASE-X PCS/PMA or SGMII Using A RocketIO Transceiver.....	23
Ethernet 1000BASE-X PCS/PMA or SGMII with Ten-Bit-Interface.....	25
Core Interfaces	26
Client Side Interface.....	31
Physical Side Interface.....	36
Chapter 3: Generating and Customizing the Core	
GUI Interface	39
Component Name.....	39
Select Standard.....	40
Core Functionality.....	40
SGMII/Dynamic Standard Switching Elastic Buffer Options.....	41
RocketIO Tile Configuration.....	43
Parameter Values in the XCO File	43
Output Generation	44
Chapter 4: Designing with the Core	
Design Overview	45
Design Guidelines	50
Generate the Core.....	50
Examine the Example Design Provided with the Core.....	50

Implement the Ethernet 1000BASE-X PCS/PMA or SGMII Core in Your Application	50
--	----

Chapter 5: Using the Client-side GMII Data Path

Designing with the Client-side GMII for the 1000BASE-X Standard	53
GMII Transmission	53
GMII Reception	54
status_vector[4:0] signals	56
Using the Virtex-II Pro RocketIO Transceiver CRC Functionality	57
Designing with Client-side GMII for the SGMII Standard	59
Overview	59
GMII Transmission	59
GMII Reception	60
Using the GMII as an Internal Connection	61
Implementing External GMII	61
GMII Transmitter Logic	61
GMII Receiver Logic	66

Chapter 6: The Ten-Bit Interface

Ten-Bit-Interface Logic	69
Transmitter Logic	69
Receiver Logic	70
Clock Sharing across Multiple Cores with TBI	77

Chapter 7: 1000BASE-X with RocketIO Transceivers

RocketIO Transceiver Logic	79
Virtex-II Pro Devices	79
Virtex-4 FX Devices	81
Virtex-5 LXT and SXT Devices	83
Virtex-5 FXT Devices	85
Clock Sharing Across Multiple Cores with RocketIO	87
Virtex-II Pro Devices	87
Virtex-4 FX Devices	88
Virtex-5 LXT and SXT Devices	90
Virtex-5 FXT Devices	92

Chapter 8: SGMII / Dynamic Standards Switching with RocketIO Transceivers

Receiver Elastic Buffer Implementations	95
Selecting the Buffer Implementation from the GUI	95
The Requirement for the FPGA Fabric Rx Elastic Buffer	96
The RocketIO Rx Elastic Buffer	97
RocketIO Logic using the RocketIO Rx Elastic Buffer	98
RocketIO Logic with the Fabric Rx Elastic Buffer	98
Virtex-II Pro Devices	98
Virtex-4 Devices for SGMII or Dynamic Standards Switching	101
Virtex-5 LXT or SXT Devices for SGMII or Dynamic Standards Switching	103
Virtex-5 FXT Devices for SGMII or Dynamic Standards Switching	105
Clock Sharing - Multiple Cores with RocketIO, Fabric Elastic Buffer	107
Virtex-II Pro Devices	107
Virtex-4 FX Devices	109

Virtex-5 LXT and SXT Devices	111
Virtex-5 FXT Devices	113

Chapter 9: Configuration and Status

MDIO Management Interface	115
MDIO Bus System	115
MDIO Transactions	116
MDIO Addressing	117
Connecting the MDIO to an Internally Integrated STA	118
Connecting the MDIO to an External STA	118
Management Registers	119
1000BASE-X Standard Using the Optional Auto-Negotiation	119
1000BASE-X Standard Without the Optional Auto-Negotiation	129
SGMII Standard Using the Optional Auto-Negotiation	135
SGMII Standard without the Optional Auto-Negotiation	145
Both 1000BASE-X and SGMII Standards	150
Optional Configuration Vector	151

Chapter 10: Auto-Negotiation

Overview of Operation	153
1000BASE-X Standard	153
SGMII Standard	155
Setting the Configurable Link Timer	156
1000BASE-X Standard	156
SGMII Standard	156
Simulating Auto-Negotiation	156
Using the Auto-Negotiation Interrupt	156

Chapter 11: Dynamic Switching of 1000BASE-X and SGMII Standards

Typical Application	157
Operation of the Core	158
Selecting the Power-On / Reset Standard	158
Switching the Standard Using MDIO	158
Auto-Negotiation State Machine	158
Setting the Auto-Negotiation Link Timer	158

Chapter 12: Constraining the Core

Required Constraints	161
Device, Package, and Speedgrade Selection	161
I/O Location Constraints	161
Placement Constraints	161
Virtex-II Pro RocketIO MGTs for 1000BASE-X Constraints	161
Virtex-II Pro RocketIO MGTs for SGMII or Dynamic Standards	
Switching Constraints	163
Virtex-4 RocketIO MGTs for 1000BASE-X Constraints	164
Virtex-4 RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints	166
Virtex-5 RocketIO GTP Transceivers for 1000BASE-X Constraints	166
Virtex-5 RocketIO GTP Transceivers for SGMII or Dynamic Standards	
Switching Constraints	167
Virtex-5 RocketIO GTX Transceivers for 1000BASE-X Constraints	167

Virtex-5 RocketIO GTX Transceivers for SGMII or Dynamic Standards	
Switching Constraints	168
Ten-Bit Interface Constraints	168
Constraints When Implementing an External GMII	172
Understanding Timing Reports for Setup/Hold Timing	176

Chapter 13: Interfacing to Other Cores

Integrating with the 1-Gigabit Ethernet MAC Core	179
Integration of the 1-Gigabit Ethernet MAC to 1000BASE-X PCS with TBI	179
Integration of the 1-Gigabit Ethernet MAC Using a RocketIO Transceiver	181
Integration of the 1-Gigabit Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality	185
Integrating with the Tri-Mode Ethernet MAC Core	185
Integration of the Tri-Mode Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality with TBI	185
Integration of the Tri-Mode Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality using RocketIO Transceivers	188

Chapter 14: Special Design Considerations

Power Management	197
Startup Sequencing	197
Loopback	197
Core with the TBI	197
Core with RocketIO Transceiver	198

Chapter 15: Implementing the Design

Pre-implementation Simulation	201
Using the Simulation Model	201
Synthesis	201
XST - VHDL	201
XST - Verilog	202
Implementation	202
Generating the Xilinx Netlist	202
Mapping the Design	202
Placing and Routing the Design	202
Static Timing Analysis	203
Generating a Bitstream	203
Post-Implementation Simulation	203
Generating a Simulation Model	203
Using the Model	203
Other Implementation Information	204

Appendix A: Core Verification, Compliance, and Interoperability

Verification	205
Simulation	205
Hardware Verification	205

Appendix B: Core Latency

Core Latency	207
Latency for 1000BASE-X PCS with TBI	207
Latency for 1000BASE-X PCS and PMA Using a RocketIO Transceiver	208
Latency for SGMII	208

Appendix C: Calculating the DCM Fixed Phase Shift Value

Requirement for DCM Phase Shifting	209
Finding the Ideal Phase Shift Value for Your System	209

Appendix D: 1000BASE-X State Machines

Introduction	211
Start of Frame Encoding	212
The Even Transmission Case	212
Reception of the Even Case	213
The Odd Transmission Case	213
Reception of the Odd Case	214
Preamble Shrinkage	215
End of Frame Encoding	215
The Even Transmission case	215
Reception of the Even Case	216
The Odd Transmission Case	216
Reception of the Odd Case	217

Appendix E: Rx Elastic Buffer Specifications

Introduction	219
Rx Elastic Buffers: Depths and Maximum Frame Sizes	219
RocketIO Rx Elastic Buffers	219
SGMII Fabric Rx Elastic Buffer	222
TBI Rx Elastic Buffer	223
Clock Correction	224
Maximum Frame Sizes for Sustained Frame Reception	226
Jumbo Frame Reception	226

Appendix F: Debugging Guide

General Checks	227
Problems with the MDIO	227
Problems with Data Reception or Transmission	227
Problems with Auto-Negotiation	228
Problems in Obtaining a Link (Auto-Negotiation Disabled)	228
Problems with a High Bit Error Rate	229
Symptoms	229
Debugging	229

Schedule of Figures

Chapter 2: Core Architecture

Figure 2-1: Functional Block Diagram Using RocketIO Transceiver	24
Figure 2-2: Functional Block Diagram with a Ten-Bit Interface	25
Figure 2-3: Component Pinout Using RocketIO Transceiver with PCS Management Registers	27
Figure 2-4: Component Pinout Using RocketIO Transceiver without PCS Management Registers	28
Figure 2-5: Component Pinout Using the Ten-Bit Interface with PCS Management Registers	29
Figure 2-6: Component Pinout Using Ten-Bit Interface without PCS Management Registers	30
Figure 2-7: Component Pinout with the Dynamic Switching Logic	31

Chapter 3: Generating and Customizing the Core

Figure 3-1: Core Customization Screen	39
Figure 3-2: 1000BASE-X Standard Options Screen	40
Figure 3-3: SGMII/Dynamic Standard Switching Options Screen	42
Figure 3-4: RocketIO Tile Configuration Screen	43

Chapter 4: Designing with the Core

Figure 4-1: 1000BASE-X Standard Using a RocketIO Transceiver	46
Figure 4-2: Example Design 1000BASE-X Standard Using TBI	47
Figure 4-3: Example Design Performing the SGMII Standard	48
Figure 4-4: Example Design Performing the SGMII Standard	49

Chapter 5: Using the Client-side GMII Data Path

Figure 5-1: GMII Normal Frame Transmission	53
Figure 5-2: GMII Error Propagation Within a Frame	54
Figure 5-3: GMII Normal Frame Reception	54
Figure 5-4: GMII Normal Frame Reception with Carrier Extension	55
Figure 5-5: GMII Frame Reception with Errors	55
Figure 5-6: False Carrier Indication	56
Figure 5-7: status_vector[4:2] timing	57
Figure 5-8: GMII Frame Transmission with RocketIO Transceiver CRC Logic Enabled	58
Figure 5-9: GMII Frame Reception with the RocketIO Transceiver CRC Logic Enabled	58
Figure 5-10: GMII Frame Transmission at 1 Gbps	59
Figure 5-11: GMII Data Transmission at 100 Mbps	59
Figure 5-12: GMII Frame Reception at 1 Gbps	60
Figure 5-13: GMII Data Reception at 100 Mbps	60
Figure 5-14: GMII Transmitter Logic	62
Figure 5-15: External GMII Transmitter Logic for Spartan-3, Spartan-3E and Spartan-3A Devices	63
Figure 5-16: External GMII Transmitter Logic for Virtex-4 Devices	64
Figure 5-17: External GMII Transmitter Logic for Virtex-5 Devices	65
Figure 5-18: External GMII Receiver Logic	67

Chapter 6: The Ten-Bit Interface

Figure 6-1: Ten-Bit Interface Transmitter Logic	70
Figure 6-2: Ten-Bit-Interface Receiver Logic	71
Figure 6-3: TBI Receiver Logic for Spartan-3, Spartan-3E, and Spartan-3A Devices	72
Figure 6-4: Ten-Bit Interface Receiver Logic - Virtex-4 Device (Example Design)	73
Figure 6-5: Alternate Ten-Bit Interface Receiver Logic for Virtex-4 Devices	74
Figure 6-6: Ten-Bit Interface Receiver Logic - Virtex-5 Device (Example Design)	75
Figure 6-7: Alternate Ten-Bit Interface Receiver Logic - Virtex-5 Devices	76
Figure 6-8: Clock Management, Multiple Core Instances with Ten-Bit Interface	77

Chapter 7: 1000BASE-X with RocketIO Transceivers

Figure 7-1: 1000BASE-X Connection to a Virtex-II Pro MGT	80
Figure 7-2: 1000BASE-X Connection to Virtex-4 MGT	82
Figure 7-3: 1000BASE-X Connection to Virtex-5 GTP Transceivers	84
Figure 7-4: 1000BASE-X Connection to Virtex-5 GTX Transceivers	86
Figure 7-5: Clock Management: Two Core Instances, Virtex-II Pro MGTs for 1000BASE-X	87
Figure 7-6: Clock Management - Multiple Core Instances, MGTs for 1000BASE-X	89
Figure 7-7: Clock Management - Multiple Core Instances, Virtex-5 RocketIO GTP Transceivers for 1000BASE-X	91
Figure 7-8: Clock Management - Multiple Core Instances, Virtex-5 RocketIO GTX Transceivers for 1000BASE-X	93

Chapter 8: SGMII / Dynamic Standards Switching with RocketIO Transceivers

Figure 8-1: SGMII Implementation using Separate Clock Sources	96
Figure 8-2: SGMII Implementation using Shared Clock Sources	97
Figure 8-3: SGMII Connection to a Virtex-II Pro RocketIO Transceiver	100
Figure 8-4: SGMII Connection to a Virtex-4 MGT	102
Figure 8-5: SGMII Connection to a Virtex-5 RocketIO GTP Transceiver	104
Figure 8-6: SGMII Connection to a Virtex-5 RocketIO GTX Transceiver	106
Figure 8-7: Clock Management with Multiple Core Instances with Virtex-II Pro RocketIO Transceivers for SGMII	108
Figure 8-8: Clock Management with Multiple Core Instances with Virtex-4 MGTs for SGMII ..	110
Figure 8-9: Clock Management with Multiple Core Instances with Virtex-5 GTP RocketIO Transceivers for SGMII	112
Figure 8-10: Clock Management with Multiple Core Instances with Virtex-5 GTX RocketIO Transceivers for SGMII	114

Chapter 9: Configuration and Status

Figure 9-1: A Typical MDIO-managed System	116
Figure 9-2: MDIO Write Transaction	117
Figure 9-3: MDIO Read Transaction	117
Figure 9-4: Creating an External MDIO Interface	119
Figure 9-5: Dynamic Switching (Register 17)	151

Chapter 10: Auto-Negotiation

Figure 10-1: 1000BASE-X Auto-Negotiation Overview	153
Figure 10-2: SGMII Auto-Negotiation	155

Chapter 11: Dynamic Switching of 1000BASE-X and SGMII Standards

<i>Figure 11-1: Typical Application for Dynamic Switching</i>	157
---	-----

Chapter 12: Constraining the Core

<i>Figure 12-1: Local Clock Place and Route for Top MGT</i>	164
<i>Figure 12-2: Input TBI timing</i>	170
<i>Figure 12-3: Input GMII timing</i>	174
<i>Figure 12-4: Timing Report Setup/Hold Illustration</i>	177

Chapter 13: Interfacing to Other Cores

<i>Figure 13-1: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS with TBI</i>	180
<i>Figure 13-2: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-II Pro MGT</i>	181
<i>Figure 13-3: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-4 MGT</i>	182
<i>Figure 13-4: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-5 GTP Transceiver</i>	183
<i>Figure 13-5: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-5 GTX Transceiver</i>	184
<i>Figure 13-6: Tri-Speed Ethernet MAC Extended to use an SGMII with TBI</i>	187
<i>Figure 13-7: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-II Pro</i>	189
<i>Figure 13-8: Tri-Speed Ethernet MAC Extended to Use an SGMII in Virtex-4</i>	191
<i>Figure 13-9: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-5 LXT/SXT</i>	193
<i>Figure 13-10: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-5 FXT</i>	195

Chapter 14: Special Design Considerations

<i>Figure 14-1: Loopback Implementation Using the TBI</i>	198
<i>Figure 14-2: Loopback Implementation When Using the Core with RocketIO Transceivers</i>	199

Appendix D: 1000BASE-X State Machines

<i>Figure D-1: 1000BASE-X Transmit State Machine Operation (Even Case)</i>	212
<i>Figure D-2: 1000BASE-X Reception State Machine Operation (Even Case)</i>	213
<i>Figure D-3: 1000BASE-X Transmit State Machine Operation (Odd Case)</i>	214
<i>Figure D-4: 1000BASE-X Reception State Machine Operation (Odd Case)</i>	214
<i>Figure D-5: 1000BASE-X Transmit State Machine Operation (Even Case)</i>	215
<i>Figure D-6: 1000BASE-X Reception State Machine Operation (Even Case)</i>	216
<i>Figure D-7: 1000BASE-X Transmit State Machine Operation (Even Case)</i>	217
<i>Figure D-8: 1000BASE-X Reception State Machine Operation (Odd Case)</i>	217

Appendix E: Rx Elastic Buffer Specifications

<i>Figure E-1: Elastic Buffer Sizes for all RocketIO Transceiver Families</i>	220
<i>Figure E-2: Elastic Buffer Size for all RocketIO families</i>	222
<i>Figure E-3: TBI Elastic Buffer Size for All Families</i>	223

Schedule of Tables

Chapter 2: Core Architecture

Table 2-1: GMII Interface Signal Pinout	32
Table 2-2: Other Common Signals	33
Table 2-3: Optional MDIO Interface Signal Pinout.	34
Table 2-4: Optional Configuration and Status Vectors	35
Table 2-5: Optional Auto-Negotiation Interface Signal Pinout.	35
Table 2-6: Optional Dynamic Standard Switching Signals	36
Table 2-7: Optional RocketIO Transceiver Interface Pinout	37
Table 2-8: Optional TBI Interface Signal Pinout	38

Chapter 3: Generating and Customizing the Core

Table 3-1: XCO File Values and Default Values.	44
--	----

Chapter 4: Designing with the Core

Table 4-1: Degree of Difficulty for Various Implementations	51
---	----

Chapter 9: Configuration and Status

Table 9-1: Abbreviations and Terms.	116
Table 9-2: MDIO Registers for 1000BASE-X with Auto-Negotiation.	119
Table 9-3: Control Register (Register 0)	120
Table 9-4: Status Register (Register 1)	122
Table 9-5: PHY Identifier (Registers 2 and 3)	124
Table 9-6: Auto-Negotiation Advertisement Register (Register 4)	124
Table 9-7: Auto-Negotiation Link Partner Ability Base Register (Register 5)	125
Table 9-8: Auto-Negotiation Expansion Register (Register 6)	126
Table 9-9: Auto-Negotiation Next Page Transmit (Register 7).	127
Table 9-10: Auto-Negotiation Next Page Receive (Register 8)	128
Table 9-11: Extended Status Register (Register 15)	128
Table 9-12: Vendor Specific Register: Auto-Negotiation Interrupt Control Register (Register 16)	129
Table 9-13: MDIO Registers for 1000BASE-X without Auto-Negotiation.	130
Table 9-14: Control Register (Register 0)	130
Table 9-15: Status Register (Register 1)	132
Table 9-16: PHY Identifier (Registers 2 and 3)	133
Table 9-17: Extended Status (Register 15)	134
Table 9-18: MDIO Registers for 1000BASE-X with Auto-Negotiation.	135
Table 9-19: SGMII Control (Register 0)	136
Table 9-20: SGMII Status (Register 1).	137

<i>Table 9-21: PHY Identifier (Registers 2 and 3)</i>	139
<i>Table 9-22: SGMII Auto-Negotiation Advertisement (Register 4)</i>	139
<i>Table 9-23: SGMII Auto-Negotiation Link Partner Ability Base (Register 5)</i>	140
<i>Table 9-24: SGMII Auto-Negotiation Expansion (Register 6)</i>	141
<i>Table 9-25: SGMII Auto-Negotiation Next Page Transmit (Register 7)</i>	141
<i>Table 9-26: SGMII Auto-Negotiation Next Page Receive (Register 8)</i>	142
<i>Table 9-27: SGMII Extended Status Register (Register 15)</i>	143
<i>Table 9-28: SGMII Auto-Negotiation Interrupt Control (Register 16)</i>	144
<i>Table 9-29: MDIO Registers for 1000BASE-X with Auto-Negotiation</i>	145
<i>Table 9-30: SGMII Control (Register 0)</i>	146
<i>Table 9-31: SGMII Status (Register 1)</i>	147
<i>Table 9-32: PHY Identifier (Registers 2 and 3)</i>	149
<i>Table 9-33: SGMII Auto-Negotiation Advertisement (Register 4)</i>	149
<i>Table 9-34: SGMII Extended Status Register (Register 15)</i>	150
<i>Table 9-35: Vendor-specific Register: Standard Selection Register (Register 17)</i>	151
<i>Table 9-36: Optional Configuration and Status Vectors</i>	152

Chapter 12: Constraining the Core

<i>Table 12-1: Input TBI Timing</i>	170
<i>Table 12-2: Input GMII Timing</i>	174

Appendix D: 1000BASE-X State Machines

<i>Table D-1: Defined Ordered Sets</i>	211
--	-----

Appendix E: Rx Elastic Buffer Specifications

<i>Table E-1: Maximum Frame Sizes: RocketIO Transceiver Rx Elastic Buffers (100ppm Clock Tolerance)</i>	220
<i>Table E-2: Maximum Frame Sizes: Fabric Rx Elastic Buffers (100ppm Clock Tolerance)</i>	223
<i>Table E-3: Maximum Frame Size: (Sustained Frame Reception) Capabilities of the Rx Elastic Buffers</i>	226

About This Guide

The *LogiCORE™ IP Ethernet 1000BASE-X PCS/PMA or SGMII User Guide* provides information about generating a Xilinx Ethernet 1000BASE-X PCS/PMA or SGMII core, customizing and simulating the core using the provided example design, and running the design files through implementation using the Xilinx tools.

Guide Contents

This guide contains the following information.

- [Preface, “About This Guide”](#) introduces the organization and purpose of this guide and defines the conventions used in this document.
- [Chapter 1, “Introduction”](#) describes the core and related information, including recommended design experience, additional documentation resources, technical support, and submitting feedback to Xilinx.
- [Chapter 2, “Core Architecture”](#) provides an overview of the core including all interfaces and major functional blocks.
- [Chapter 3, “Generating and Customizing the Core”](#) describes the Graphical User Interface (GUI) options used to generate and customize the core.
- [Chapter 4, “Designing with the Core”](#) provides general guidelines for creating designs with the core.
- [Chapter 5, “Using the Client-side GMII Data Path”](#) provides general guidelines for creating designs using client side GMII of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- [Chapter 6, “The Ten-Bit Interface”](#) provides general design guidelines when using the Ten-Bit Interface (TBI) as the Physical Side of the core.
- [Chapter 7, “1000BASE-X with RocketIO Transceivers”](#) provides general design guidelines when using the 1000BASE-X standard with the RocketIO™ transceiver as the physical side of the core.
- [Chapter 8, “SGMII / Dynamic Standards Switching with RocketIO Transceivers”](#) provides general design guidelines when using either the SGMII standard, or the Dynamic Switching option (between 1000BASE-X and SGMII standards). These options always use a RocketIO as the physical interface.
- [Chapter 9, “Configuration and Status”](#) provides general guidelines for configuring and monitoring the core, including a detailed description of the management registers present in the core.
- [Chapter 10, “Auto-Negotiation”](#) provides guidelines for Auto-Negotiation function of the core.

- [Chapter 11, “Dynamic Switching of 1000BASE-X and SGMII Standards”](#) provides general guidelines for using the core to perform dynamic standards switching between 1000BASE-X and SGMII.
- [Chapter 12, “Constraining the Core”](#) defines the constraint requirements of the core.
- [Chapter 13, “Interfacing to Other Cores”](#) describes additional design considerations associated with implementing the core with the 1-Gigabit Ethernet MAC and Tri-Mode Ethernet MAC cores.
- [Chapter 14, “Special Design Considerations”](#) describes additional design considerations associated with implementing the core.
- [Chapter 15, “Implementing the Design”](#) describes how to simulate and implement your design containing the core.
- [Appendix A, “Core Verification, Compliance, and Interoperability”](#) describes how the core was verified.
- [Appendix B, “Core Latency”](#) defines the latency of the core.
- [Appendix C, “Calculating the DCM Fixed Phase Shift Value”](#) instructs the user about how to calculate the system timing requirements when using DCMs with the core.
- [Appendix D, “1000BASE-X State Machines”](#) serves as a reference for the basic operation of the 1000BASE-X *IEEE 802.3* clause 36 transmitter and receiver state machines.
- [Appendix E, “Rx Elastic Buffer Specifications”](#) describes the depth of the Rx Elastic Buffers which are available with the core. The size of the buffer is related to the maximum frame size which the core can accommodate.
- [Appendix F, “Debugging Guide”](#) provides information for debugging the core within a system.

Conventions

This document uses the following conventions. An example illustrates each convention.

Typographical

The following typographical conventions are used in this document.

Convention	Meaning or Use	Example
Courier font	Messages, prompts, and program files that the system displays	speed grade: - 100
Courier bold	Literal commands you enter in a syntactical statement	ngdbuild <i>design_name</i>
<i>Italic font</i>	References to other manuals	See the <i>User Guide</i> for details.
	Emphasis in text	If a wire is drawn so that it overlaps the pin of a symbol, the two nets are <i>not</i> connected.
Dark Shading	Items that are not supported or reserved	This feature is not supported

Convention	Meaning or Use	Example
Square brackets []	An optional entry or parameter. However, in bus specifications, such as bus [7:0] , they are required.	ngdbuild [<i>option_name</i>] <i>design_name</i>
Braces { }	A list of items from which you must choose one or more	lowpwr = { on off }
Vertical bar	Separates items in a list of choices	lowpwr = { on off }
Vertical ellipsis .	Repetitive material that has been omitted	IOB #1: Name = QOUT' IOB #2: Name = CLKIN' . . .
Horizontal ellipsis ...	Repetitive material that has been omitted	allow block <i>block_name</i> <i>loc1 loc2 ... locn</i> ;
Notations	The prefix '0x' or the suffix 'h' indicate hexadecimal notation	A read of address 0x00112975 returned 45524943h.
	A '_n' means the signal is active low	usr_teof_n is active low.

Online Document

The following conventions are used in this document.

Convention	Meaning or Use	Example
Blue text	Cross-reference link to a location in the current document	See the section " Additional Resources " for details. See " Title Formats " in Chapter 1 for details.
Blue, underlined text	Hyperlink to a website (URL)	Go to www.xilinx.com for the latest speed files.

Introduction

The Ethernet 1000BASE-X PCS/PMA or SGMII core is a fully verified solution that supports Verilog HDL and VHDL. In addition, the example design provided with the core supports both Verilog and VHDL.

This chapter introduces the Ethernet 1000BASE-X PCS/PMA or SGMII core and provides related information, including recommended design experience, additional resources, technical support, and methods for submitting feedback to Xilinx.

About the Core

The Ethernet 1000BASE-X PCS/PMA or SGMII core is a Xilinx CORE Generator™ IP core, included in the latest IP Update on the Xilinx IP Center. For detailed information about the core, see the Ethernet 100BASE-X PCS/PMA [product page](#). For information about system requirements and licensing options, see Chapter 2, “Licensing the Core,” in the *Getting Started Guide*.

Designs Using RocketIO Transceivers

RocketIO transceivers are defined by device family in the following way:

- For Virtex-II Pro and Virtex-4 devices, RocketIO Multi-Gigabit Transceivers (MGT)
- For Virtex-5 LXT and SXT devices, RocketIO GTP transceivers; Virtex-5 FXT devices, RocketIO GTX transceivers

Recommended Design Experience

Although the Ethernet 1000BASE-X PCS/PMA or SGMII core is a fully-verified solution, the challenge associated with implementing a complete design varies depending on the configuration and functionality of the application. For best results, previous experience building high-performance, pipelined FPGA designs using Xilinx implementation software and User Constraint Files (UCF) is recommended.

Contact your local Xilinx representative for a closer review and estimation for your specific requirements.

Additional Core Resources

For detailed information and updates about the Ethernet 1000BASE-X PCS/PMA or SGMII core, see the following documents, located on the Xilinx Ethernet 100BASE-X PCS/PMA [product page](#).

- *Ethernet 1000BASE-X PCS/PMA or SGMII Data Sheet*
- *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide*

After generating the core, the following documents are available in the document directory:

- *Ethernet 1000BASE-X PCS/PMA or SGMII Release Notes*
- *Ethernet 1000BASE-X PCS/PMA or SGMII User Guide*

Related Xilinx Ethernet Products and Services

For information about all Xilinx Ethernet solutions, see www.xilinx.com/products/design_resources/conn_central/protocols/gigabit_ethernet.htm.

Specifications

- IEEE 802.3
- *Serial-GMII Specification* (CISCO SYSTEMS, ENG-46158)

Technical Support

To obtain technical support specific to the Ethernet 1000BASE-X PCS/PMA or SGMII core, visit www.support.xilinx.com/. Questions are routed to a team of engineers with expertise using the Ethernet 1000BASE-X PCS/PMA or SGMII core.

Xilinx provides technical support for use of this product as described in the *Ethernet 1000BASE-X PCS/PMA or SGMII User Guide* and the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide*. Xilinx cannot guarantee timing, functionality, or support of this product for designs that do not follow these guidelines.

Feedback

Xilinx welcomes comments and suggestions about the Ethernet 1000BASE-X PCS/PMA or SGMII core and the documentation supplied with the core.

Ethernet 1000BASE-X PCS/PMA or SGMII Core

For comments or suggestions about the core, please submit a WebCase from www.support.xilinx.com/. Be sure to include the following information:

- Product name
- Core version number
- Explanation of your comments

Document

For comments or suggestions about this document, please submit a WebCase from www.support.xilinx.com/. Be sure to include the following information:

- Document title
- Document number
- Page number(s) to which your comments refer
- Explanation of your comments

Core Architecture

This chapter describes the architecture of the Ethernet 1000BASE-X PCS/PMA or SGMII core, including all interfaces and major functional blocks.

System Overview

Ethernet 1000BASE-X PCS/PMA or SGMII Using A RocketIO Transceiver

The Ethernet 1000BASE-X PCS/PMA or SGMII core provides the functionality to implement the 1000BASE-X PCS and PMA sub-layers or used to provide a GMII to SGMII bridge when used with a RocketIO transceiver. RocketIO transceivers are defined in the following way:

- For Virtex-II Pro and Virtex-4 devices, RocketIO Multi-Gigabit Transceivers (MGT)
- For Virtex-5 LXT and SXT FPGAs, RocketIO GTP transceivers; Virtex-5 FXT FPGA, RocketIO GTX transceiver

The core interfaces to a RocketIO transceiver, providing some of the PCS layer functionality such as 8B/10B encoding/decoding, the PMA SERDES, and clock recovery. [Figure 2-1](#) illustrates the remaining PCS sublayer functionality, and also shows the major functional blocks of the core.

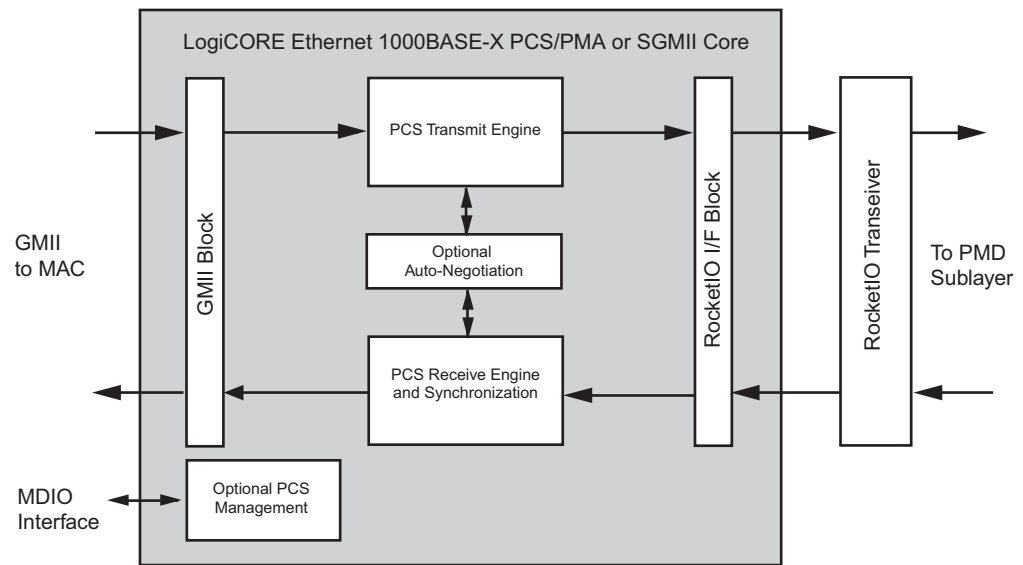


Figure 2-1: Functional Block Diagram Using RocketIO Transceiver

GMII Block

A client-side GMII is provided with the core, which can be used as an internal interface for connection to an embedded Media Access Controller (MAC) or other custom logic. Alternatively, the GMII may be routed to device IOBs to provide an external (off chip) GMII.

PCS Transmit Engine

The PCS transmit engine converts the GMII data octets into a sequence of ordered sets by implementing the state diagrams of *IEEE 802.3* (figures 36-5 and 36-6). See [Appendix D, “1000BASE-X State Machines.”](#)

PCS Receive Engine and Synchronization

The synchronization process implements the state diagram of *IEEE 802.3* (figure 36-9). The PCS receive engine converts the sequence of ordered sets to GMII data octets by implementing the state diagrams of *IEEE 802.3* (figures 36-7a and 36-7b). See [Appendix D, “1000BASE-X State Machines.”](#)

Optional Auto-Negotiation Block

IEEE 802.3 clause 37 describes the 1000BASE-X Auto-Negotiation function that allows a device to advertise the modes of operation that it supports to a device at the remote end of a link segment (link partner), and to detect corresponding operational modes that the link partner may be advertising.

Auto-Negotiation is controlled and monitored through the PCS Management Registers. See [Chapter 10, “Auto-Negotiation.”](#)

Optional PCS Management Registers

Configuration and status of the core, including access to and from the optional Auto-Negotiation function, uses the 1000BASE-X PCS Management Registers defined in *IEEE 802.3* clause 37. These registers are accessed through the serial Management Data Input/Output Interface (MDIO), defined in *IEEE 802.3* clause 22, as if it were an externally connected PHY.

The PCS Management Registers may be omitted from the core when the core is performing the 1000BASE-X standard. In this situation, configuration and status of the core is made possible with the use of an alternative configuration vector and a status signal.

When the core is performing the SGMII standard, the PCS Management Registers become mandatory and information in the registers takes on a different interpretation. For more information, see “Management Registers” in Chapter 9.

RocketIO Interface Block

The RocketIO Interface Block enables the core to connect to a Virtex-II Pro, Virtex-4, or Virtex-5 FPGA RocketIO transceiver.

Ethernet 1000BASE-X PCS/PMA or SGMII with Ten-Bit-Interface

The Ethernet 1000BASE-X PCS/PMA or SGMII core, when used with the Ten-Bit Interface (TBI), allows you to implement only the 1000BASE-X PCS sublayer.

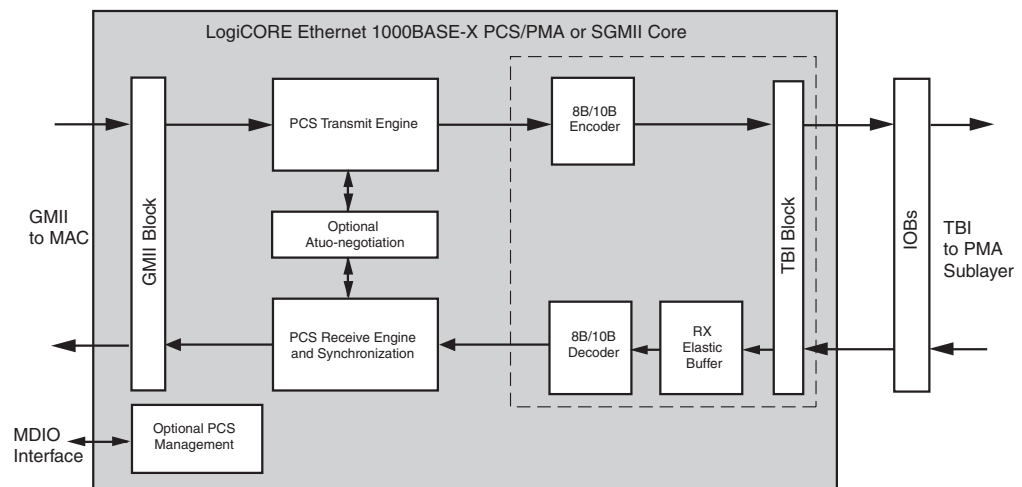


Figure 2-2: Functional Block Diagram with a Ten-Bit Interface

The optional TBI can be used in place of the RocketIO transceiver to provide a parallel interface for connection to an external PMA SERDES device. In this implementation, additional logic blocks are required to replace some of the RocketIO transceiver functionality. These are shown in the surrounded by the dotted line box in Figure 2-2 and are described in the following sections. The other blocks are described previously in this document.

8B/10B Encoder

8B10B encoding, as defined in *IEEE 802.3* (Tables 36-1a to 36-1e and Table 36-2), is implemented in a block SelectRAM™, configured as ROM, and used as a large look-up table.

8B/10B Decoder

8B10B decoding, as defined in *IEEE 802.3* (Table 36-1a to 36-1e and Table 36-2), is implemented in a block SelectRAM, configured as ROM, and used as a large look-up table.

Receiver Elastic Buffer

The Receiver Elastic Buffer enables the 10-bit parallel TBI data, received from the PMA sublayer synchronously to the TBI receiver clocks, to be transferred onto the core's internal 125 MHz clock domain. It is an asynchronous FIFO implemented in internal RAM. The Receiver Elastic Buffer attempts to maintain a constant occupancy by inserting or removing Idle sequences as necessary. This causes no corruption to the frames of data.

TBI Block

The core provides a TBI interface that should be routed to device IOBs to provide an off-chip TBI.

Core Interfaces

All ports of the core are internal connections in FPGA fabric. An HDL example design (delivered with the core) connects the core, where appropriate, to a RocketIO transceiver, and/or add IBUFs, OBUFs, and IOB flip-flops to the external signals of the GMII and TBI. IOBs are added to the remaining unconnected ports to take the example design through the Xilinx implementation software.

All clock management logic is placed in this example design, allowing you more flexibility in implementation (such as designs using multiple cores). This example design is provided in both VHDL and Verilog. For more information, see the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide*.

[Figure 2-3](#) shows the pinout for the Ethernet 1000BASE-X PCS/PMA or SGMII core using a RocketIO transceiver *with* the optional PCS Management Registers. The signals shown in the Auto-Negotiation box included only when the core includes the Auto-Negotiation

functionality. For more information, see [Chapter 3, “Generating and Customizing the Core.”](#)

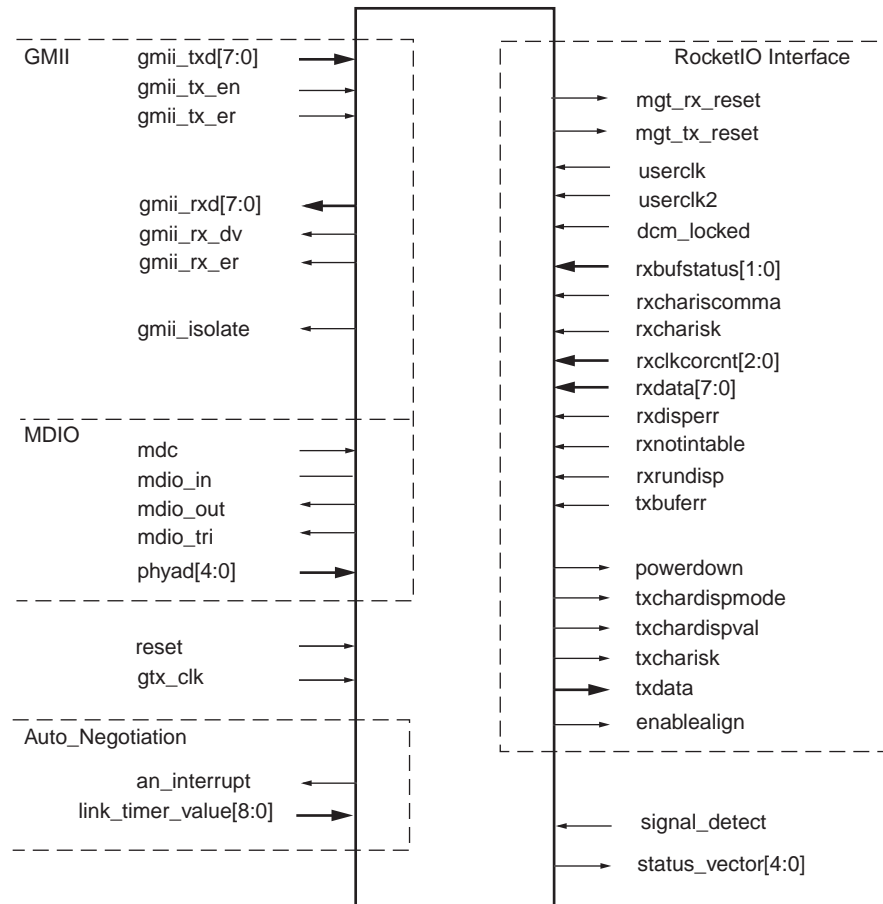


Figure 2-3: Component Pinout Using RocketIO Transceiver with PCS Management Registers

Figure 2-4 shows the pinout for the Ethernet 1000BASE-X PCS/PMA or SGMII core using a RocketIO transceiver *without* the optional PCS Management Registers

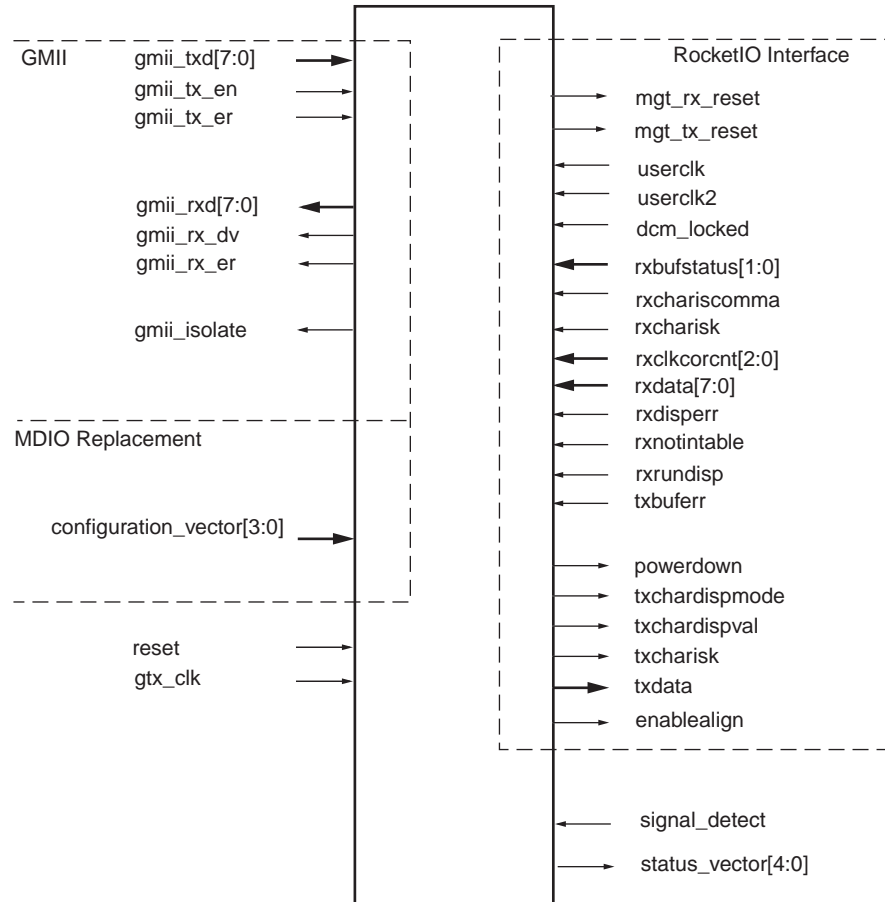


Figure 2-4: Component Pinout Using RocketIO Transceiver *without* PCS Management Registers

Figure 2-5 shows the pinout for the Ethernet 1000BASE-X PCS/PMA or SGMII core when using the TBI with optional PCS Management Registers. The signals shown in the Auto-Negotiation box are included only when the core includes the Auto-Negotiation functionality (see Chapter 3, “Generating and Customizing the Core”).

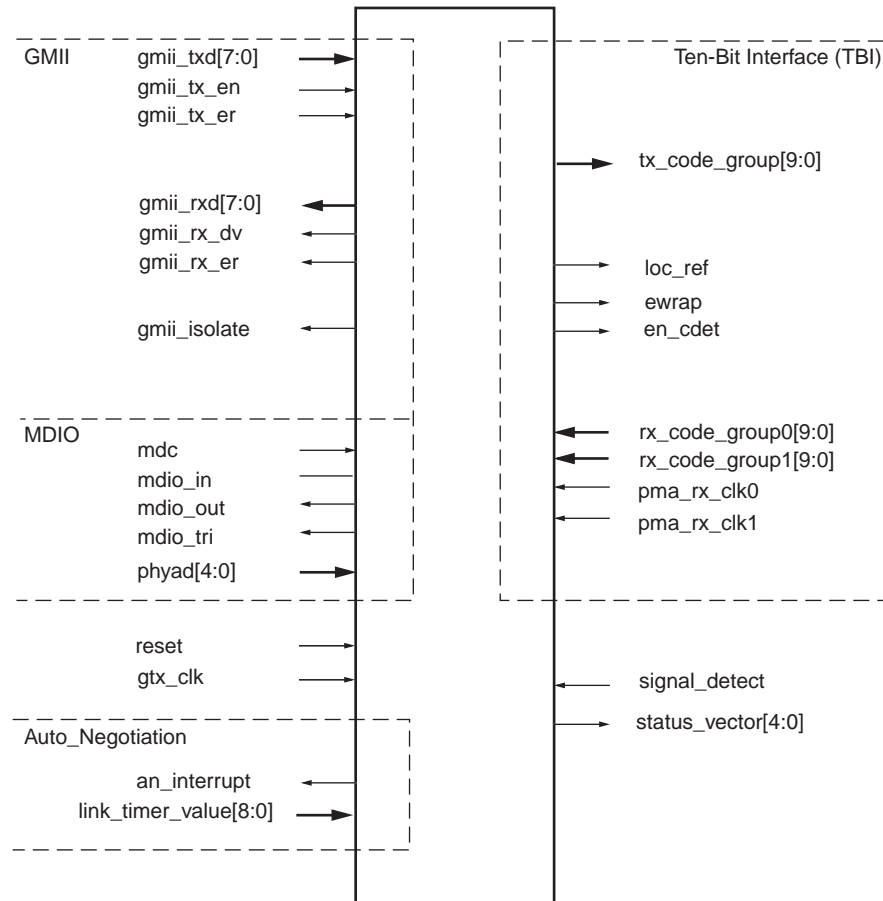


Figure 2-5: Component Pinout Using the Ten-Bit Interface with PCS Management Registers

Figure 2-6 shows the pinout for the Ethernet 1000BASE-X PCS/PMA or SGMII core when using a TBI without the optional PCS Management Registers.

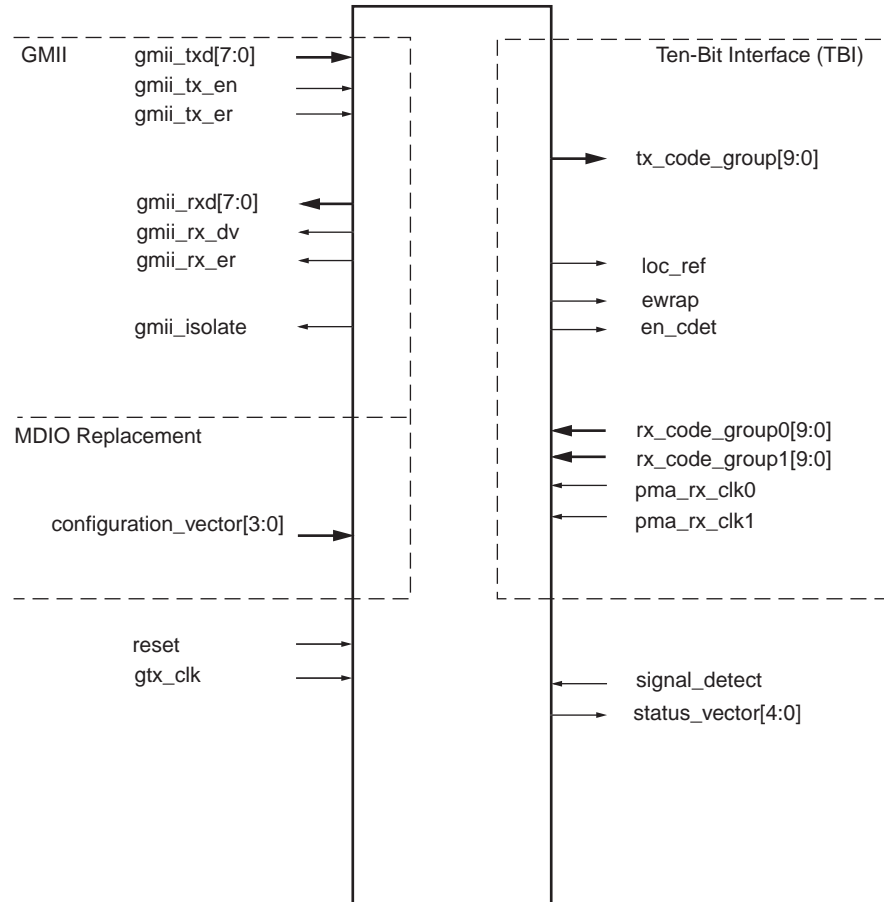


Figure 2-6: Component Pinout Using Ten-Bit Interface *without* PCS Management Registers

Figure 2-7 shows the pinout for the Ethernet 1000BASE-X PCS/PMA or SGMII core using the optional dynamic switching logic (between 1000BASE-X and SGMII standards). This mode is shown used with a RocketIO transceiver interface. For more information, see Chapter 11, “Dynamic Switching of 1000BASE-X and SGMII Standards.”

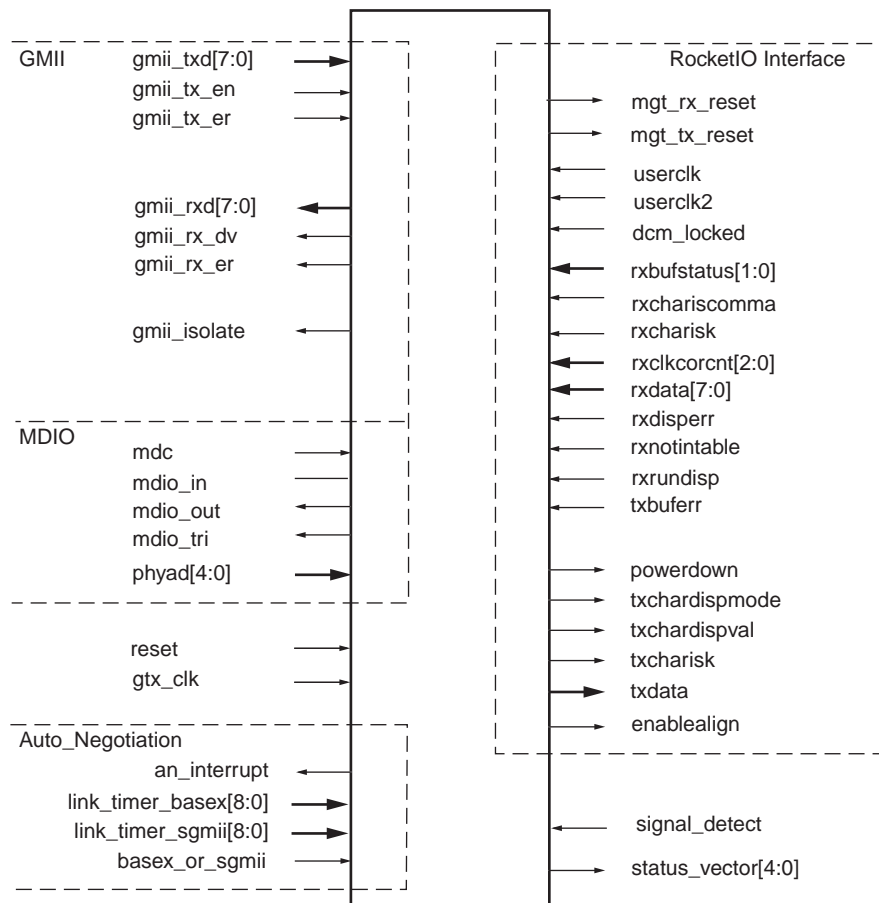


Figure 2-7: Component Pinout with the Dynamic Switching Logic

Client Side Interface

GMII Pinout

Table 2-1 describes the GMII-side interface signals of the core common to all parameterizations of the core. These are typically attached to an Ethernet MAC, either off-chip or internally integrated. The HDL example design delivered with the core connects these signals to IOBs to provide a place-and-route example.

For more information, see “Designing with the Client-side GMII for the 1000BASE-X Standard” in Chapter 5.

Table 2-1: GMII Interface Signal Pinout

Signal	Direction	Description
gmii_txd[7:0] ¹	Input	GMII Transmit data from MAC.
gmii_tx_en ¹	Input	GMII Transmit control signal from MAC.
gmii_tx_er ¹	Input	GMII Transmit control signal from MAC.
gmii_rxd[7:0] ²	Output	GMII Received data to MAC.
gmii_rx_dv ²	Output	GMII Received control signal to MAC.
gmii_rx_er ²	Output	GMII Received control signal to MAC.
gmii_isolate ²	Output	IOB Tri-state control for GMII Isolation. Only of use when implementing an External GMII as illustrated by the example design HDL.

1. When the Transmitter Elastic Buffer is present these signals are synchronous to gmii_tx_clk. When the Transmitter Elastic Buffer is omitted, see Note 2.
2. These signals are synchronous to the core's internal 125 MHz reference clock. This is userclk2 when the core is used with the RocketIO transceiver; gtx_clk when the core is used with TBI.

Common Signal Pinout

Table 2-2 describes the remaining signals common to all parameterizations of the core.

Table 2-2: Other Common Signals

Signal	Direction	Description
reset	Input	Asynchronous reset for the entire core. Active High. Clock domain is not applicable.
signal_detect	Input	Signal direct from PMD sublayer indicating the presence of light detected at the optical receiver. If set to '1,' indicates that the optical receiver has detected light. If set to '0,' this indicates the absence of light. If unused this signal should be set to '1' to enable correct operation the core. Clock domain is not applicable.
status_vector[4:0] ¹	Output	<p>Bit[0]: Link Status Indicates the status of the link.</p> <ul style="list-style-type: none"> • When high, the link is valid: synchronization of the link has been obtained and Auto-Negotiation (if present and enabled) has successfully completed. • When low, a valid link has not been established. Either link synchronization has failed or Auto-Negotiation (if present and enabled) has failed to complete. • When auto-negotiation is enabled this signal is identical to Status Register Bit 1.2: Link Status. • When auto-negotiation is disabled this signal is identical to status_vector Bit[1]. <p>Bit[1]: Link Synchronization Indicates the state of the synchronization state machine (IEEE802.3 figure 36-9) which is based on the reception of valid 8B10B code groups. This signal is similar to Bit[0] (Link Status), but is NOT qualified with Auto-Negotiation.</p> <ul style="list-style-type: none"> • When high, link synchronization has been obtained and in the synchronization state machine, sync_status = OK. • When low, synchronization has failed. <p>Bit[2]: RUDI(/C/) The core is receiving /C/ ordered sets (Auto-Negotiation Configuration sequences).</p> <p>Bit[3]: RUDI(/I/) The core is receiving /I/ ordered sets (Idles).</p> <p>Bit[4]: RUDI(INVALID) The core has received invalid data whilst receiving /C/ or /I/ ordered set. See "status_vector[4:0] signals" in Chapter 5 for more information.</p>

1. These signals are synchronous to the core's internal 125 MHz reference clock. This is userclk2 when the core is used with the RocketIO transceiver; this is gtx_clk when the core is used with TBI.

MDIO Management Interface Pinout (Optional)

Table 2-3 describes the optional MDIO interface signals of the core used to access the PCS Management Registers. These signals are typically connected to the MDIO port of a MAC device, either off-chip or to an internally integrated MAC core. For more information, see “Management Registers” in Chapter 9.

Table 2-3: Optional MDIO Interface Signal Pinout

Signal	Direction	Clock Domain	Description
mdc	Input	N/A	Management clock (<= 2.5 MHz).
mdio_in ¹	Input	mdc	Input data signal for communication with MDIO controller (for example, an Ethernet MAC). Tie high if unused.
mdio_out ¹	Output	mdc	Output data signal for communication with MDIO controller (for example, an Ethernet MAC).
mdio_tri ¹	Output	mdc	Tri-state control for MDIO signals; ‘0’ signals that the value on mdio_out should be asserted onto the MDIO interface.
phyad[4:0]	Input	N/A	Physical Address of the PCS Management register set. It is expected that this signal will be tied off to a logical value.

1. These signals can be connected to a Tri-state buffer to create a bidirectional mdio signal suitable for connection to an external MDIO controller (for example, an Ethernet MAC).

Configuration Vector (Optional)

Table 2-4 shows the alternative to the optional MDIO Management Interface, the configuration vector. See “Optional Configuration Vector” in Chapter 9.

Table 2-4: Optional Configuration and Status Vectors

Signal	Direction	Description
configuration_vector[3:0] ¹	Input	<p>Bit[0]: Reserved (currently unused)</p> <p>Bit[1]: Loopback Control</p> <ul style="list-style-type: none"> When the core with RocketIO transceiver is used, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to ewrap. When set to ‘1,’ this indicates to the external PMA module to enter loopback mode. <p>Bit[2]: Power Down</p> <ul style="list-style-type: none"> When the RocketIO transceiver is used (when set to ‘1’), the MGT is placed in a low power state. A reset must be applied to clear. With the TBI version this bit is unused. <p>Bit[3]: Isolate</p> <p>When set to ‘1,’ the GMII should be electrically isolated. When set to ‘0,’ normal operation is enabled.</p>

1. This signal is synchronous to the core’s internal 125 MHz reference clock. This is `userclk2` when the core is used with the RocketIO transceiver; this is `gtx_clk` when the core is used with TBI.

Auto-Negotiation Signal Pinout

Table 2-5 describes the signals present when the optional Auto-Negotiation functionality is present. For more information, see Chapter 10, “Auto-Negotiation.”

Table 2-5: Optional Auto-Negotiation Interface Signal Pinout

Signal	Direction	Description
link_timer_value[8:0] ¹	Input	<p>Used to configure the duration of the Auto-Negotiation Link Timer period. The duration of this timer is set to the binary number input into this port multiplied by 4096 clock periods of the 125 MHz reference clock (8 ns). It is expected that this signal will be tied off to a logical value.</p> <p>This port is replaced when using the dynamic switching mode.</p>
an_interrupt ¹	Output	<p>Active high interrupt to signal the completion of an Auto-Negotiation cycle. This interrupt can be enabled/disabled and cleared by writing to the appropriate PCS Management Register.</p>

1. These signals are synchronous to the core’s internal 125 MHz reference clock. This is `userclk2` when the core is used with the RocketIO transceiver; this is `gtx_clk` when the core is used with TBI.

Dynamic Switching Signal Pinout

Table 2-6 describes the signals present when the optional Dynamic Switching mode (between 1000BASE-X and SGMII standards) is selected. In this case, the MDIO (Table 2-3) and RocketIO transceiver (Table 2-7) interfaces are always present.

Table 2-6: Optional Dynamic Standard Switching Signals

Signal	Direction	Description
link_timer_basex[8:0] ¹	Input	Used to configure the duration of the Auto-Negotiation Link Timer period when performing the 1000BASE-X standard. The duration of this timer is set to the binary number input into this port multiplied by 4096 clock periods of the 125 MHz reference clock (8 ns). It is expected that this signal will be tied off to a logical value.
link_timer_sgmii[8:0] ¹	Input	Used to configure the duration of the Auto-Negotiation Link Timer period when performing the SGMII standard. The duration of this timer is set to the binary number input into this port multiplied by 4096 clock periods of the 125 MHz reference clock (8 ns). It is expected that this signal will be tied off to a logical value.
basex_or_sgmii ¹	Input	Used as the reset default to select the standard. It is expected that this signal will be tied off to a logical value. '0' signals that the core will come out of reset operating as 1000BASE-X. '1' signals that the core will come out of reset operating as SGMII. Note: The standard can be set following reset through the MDIO Management.

1. Clock domain is userclk2.

Physical Side Interface

1000BASE-X PCS with PMA Using RocketIO Transceiver Signal Pinout (Optional)

Table 2-7 describes the optional interface to the RocketIO transceiver. The core is connected to a RocketIO transceiver in the appropriate HDL example design delivered with the core. For more information, see:

- Chapter 7, "1000BASE-X with RocketIO Transceivers"
- Chapter 8, "SGMII / Dynamic Standards Switching with RocketIO Transceivers"

Table 2-7: Optional RocketIO Transceiver Interface Pinout

Signal	Direction	Description
mgt_rx_reset ¹	Output	Reset signal issued by the core to the RocketIO transceiver receiver path. Connect to RXRESET signal of RocketIO transceiver.
mgt_tx_reset ¹	Output	Reset signal issued by the core to the RocketIO transceiver transmitter path. Connect to TXRESET signal of RocketIO transceiver.
userclk	Input	Also connected to TXUSRCLK and RXUSRCLK of the RocketIO transceiver. Clock domain is not applicable.
userclk2	Input	Also connected to TXUSRCLK2 and RXUSRCLK2 of the RocketIO transceiver. Clock domain is not applicable.
dcm_locked	Input	A DCM may be used to derive userclk and userclk2. This is implemented in the HDL design example delivered with the core. The core will use this input to hold the RocketIO transceiver in reset until the DCM obtains lock. Clock domain is not applicable.
rxbufstatus[1:0] ¹	Input	Connect to RocketIO signal of the same name.
rxchariscomma ¹	Input	Connects to RocketIO signal of the same name.
rxcharisk ¹	Input	Connects to RocketIO signal of the same name.
rxclkcorcnt[2:0] ¹	Input	Connect to RocketIO signal of the same name.
rxdata[7:0] ¹	Input	Connect to RocketIO signal of the same name.
rxdisperr ¹	Input	Connects to RocketIO signal of the same name.
rxnotintable ¹	Input	Connects to RocketIO signal of the same name.
rxrundisp ¹	Input	Connects to RocketIO signal of the same name.
txbuferr ¹	Input	Connects to RocketIO signal of the same name.
powerdown ¹	Output	Connects to RocketIO signal of the same name.
txchardispmode ¹	Output	Connects to RocketIO signal of the same name.
txchardispval ¹	Output	Connects to RocketIO signal of the same name.
txcharisk ¹	Output	Connects to RocketIO signal of the same name.
txdata[7:0] ¹	Output	Connect to RocketIO signal of the same name.
enablealign ¹	Output	Allows the transceivers to serially realign to a comma character. Connects to ENMCOMMAALIGN and ENPCOMMAALIGN of the RocketIO.

1. When the core is used with a RocketIO transceiver, `userclk2` is used as the 125 MHz reference clock for the entire core.

1000BASE-X PCS with TBI Pinout

Table 2-8 describes the optional TBI signals, used as an alternative to the RocketIO receiver interface. The appropriate HDL example design delivered with the core connects these signals to IOBs to provide an external TBI suitable for connection to an off-chip PMA SERDES device. When the core is used with the TBI, gtx_clk is used as the 125 MHz reference clock for the entire core. For more information, see Chapter 6, “The Ten-Bit Interface.”

Table 2-8: Optional TBI Interface Signal Pinout

Signal	Direction	Clock Domain	Description
gtx_clk	Input	N/A	Clock signal at 125 MHz. Tolerance must be within <i>IEEE 802.3</i> specification.
tx_code_group[9:0]	Output	gtx_clk	10-bit parallel transmit data to PMA Sublayer (SERDES).
loc_ref	Output	N/A	Causes the PMA sublayer clock recovery unit to lock to pma_tx_clk. This signal is currently tied to Ground.
ewrap	Output	gtx_clk	When '1,' this indicates to the external PMA SERDES device to enter loopback mode. When '0,' this indicates normal operation
rx_code_group0[9:0]	Input	pma_rx_clk0	10-bit parallel received data from PMA Sublayer (SERDES). This is synchronous to pma_rx_clk0.
rx_code_group1[9:0]	Input	pma_rx_clk1	10-bit parallel received data from PMA Sublayer (SERDES). This is synchronous to pma_rx_clk1.
pma_rx_clk0	Input	N/A	Received clock signal from PMA Sublayer (SERDES) at 62.5 MHz.
pma_rx_clk1	Input	N/A	Received clock signal from PMA Sublayer (SERDES) at 62.5 MHz. This is 180 degrees out of phase with pma_rx_clk0.
en_cdet	Output	gtx_clk	Enables the PMA Sublayer to perform comma realignment. This is driven from the PCS Receive Engine during the <i>Loss-Of-Sync</i> state.

Generating and Customizing the Core

The Ethernet 1000BASE-X PCS/PMA or SGMII core is generated using the CORE Generator. This chapter describes the GUI options used to generate and customize the core.

GUI Interface

Figure 3-1 displays the Ethernet 1000BASE-X PCS/PMA or SGMII customization screen, used to set core parameters and options. For help starting and using CORE Generator on your system, see the documentation included with ISE™, including the *CORE Generator Guide*, at www.xilinx.com/support/software_manuals.htm.

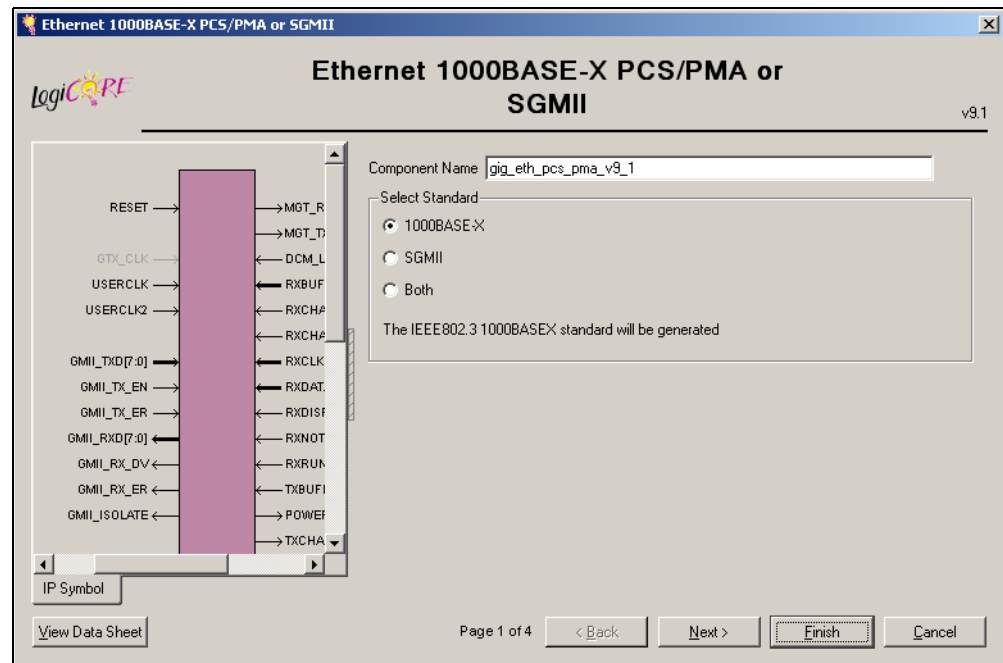


Figure 3-1: Core Customization Screen

Component Name

The component name is used as the base name of the output files generated for the core. Names must begin with a letter and must be composed from the following characters: a through z, 0 through 9 and “_.”

Select Standard

Select from the following standards for the core:

- **1000BASE-X.** 1000BASE-X Physical Coding Sublayer (PCS) functionality is designed to the IEEE 802.3 specification. Depending on the choice of physical interface, the functionality may be extended to include the 1000BASE-X Physical Medium Attachment (PMA) sublayer. Default setting.
- **SGMII.** Provides the functionality to provide a Gigabit Media Independent Interface (GMII) to Serial-GMII (SGMII) bridge, as defined in the Serial-GMII Specification (Cisco Systems, ENG-46158). SGMII may be used to replace GMII at a much lower pin count and for this reason often favored by PCB designers.
- **Both** (a combination of 1000BASE-X and SGMII). Combining the 1000BASE-X and SGMII standards lets you dynamically configure the core to switch between 1000BASE-X and SGMII standards. The core can be switched by writing through the MDIO Management Interface. For more information, see [Chapter 9, "Configuration and Status."](#)

Core Functionality

Figure 3-2 displays the Ethernet 1000BASE-X PCS/PMA or SGMII functionality screen.

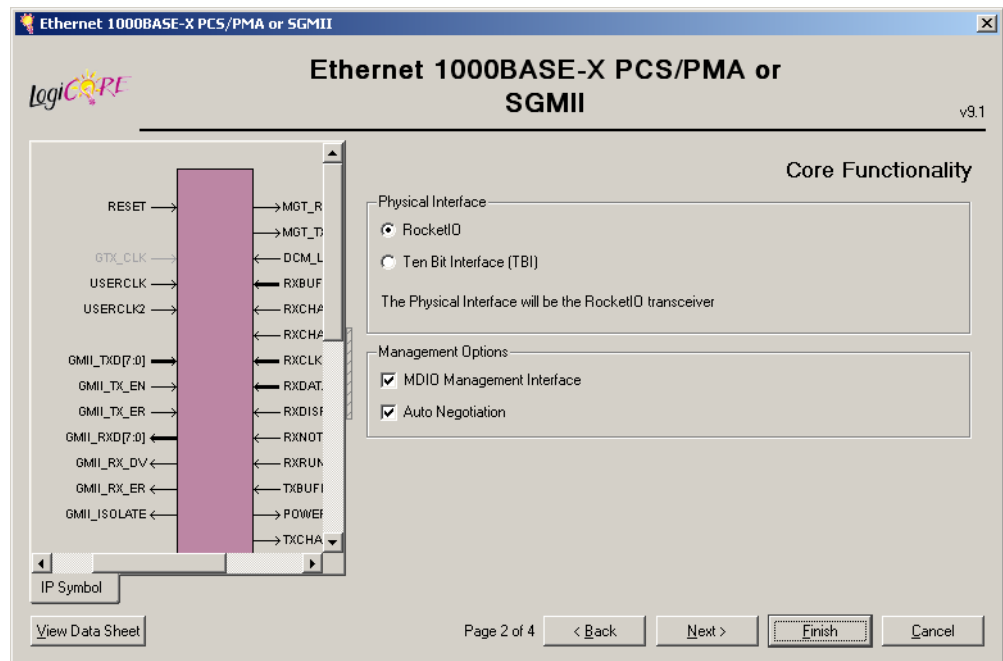


Figure 3-2: 1000BASE-X Standard Options Screen

Physical Interface

Depending on the target architecture, two physical interface options are available for the core.

- **RocketIO.** Uses a RocketIO transceiver specific to the selected device family to extend the 1000BASE-X functionality to include both PCS and PMA sub-layers. For this reason, it is available only for Virtex-II Pro, Virtex-4 FX, Virtex-5 LXT, Virtex-5 SXT, and Virtex-5 FXT devices. For additional information, see [“RocketIO Transceiver Logic” in Chapter 7](#).
- **Ten Bit Interface (TBI).** Available in all supported families and provides 1000BASE-X or SGMII functionality with a parallel TBI used to interface to an external SERDES. For more information, see [“Ten-Bit-Interface Logic” in Chapter 6](#). Default setting.

MDIO Management Interface

Select this option to include the MDIO Management Interface to access the PCS Configuration Registers. See [“MDIO Management Interface” in Chapter 9](#).

If this option is not selected, the core is generated with a replacement configuration vector. See [“Optional Configuration Vector” in Chapter 9](#). The Management Interface is selected by default.

Auto-Negotiation

Select this option to include Auto-Negotiation functionality with the core, available only if the core includes the optional Management Interface. For more information, see [Chapter 10, “Auto-Negotiation.”](#) The default is to include Auto-Negotiation.

RocketIO Transceiver CRC Logic

This option is visible in the GUI *only* when a Virtex-II Pro device family is selected, and then *only* when the RocketIO Interface is selected with the 1000BASE-X standard.

Select this option to use the built-in CRC functionality of the Virtex-II Pro RocketIO transceiver. See [Chapter 5, “Using the Virtex-II Pro RocketIO Transceiver CRC Functionality.”](#) This option is disabled (not displayed) by default.

SGMII/Dynamic Standard Switching Elastic Buffer Options

The SGMII/Dynamic Standard Switching Options screen, used to customize the Ethernet 1000BASE-X PCS/PMA or SGMII core, is *only* displayed if either SGMII or Both is selected in the Select Standard section of the initial customization screen, and *only* if RocketIO is selected as the Physical Standard.

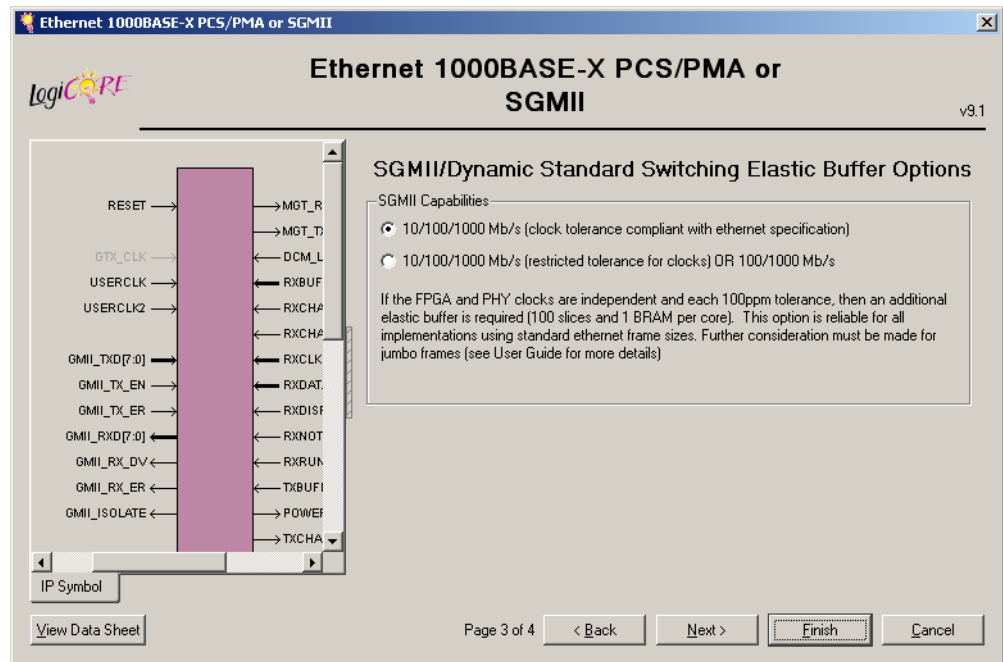


Figure 3-3: SGMII/Dynamic Standard Switching Options Screen

This screen lets you select the Receiver Elastic Buffer type to be used with the core. Before selecting this option, see [“Receiver Elastic Buffer Implementations”](#) in Chapter 8.

RocketIO Tile Configuration

The RocketIO Tile Configuration screen is only displayed if the RocketIO interface is used with the Virtex-4 or Virtex-5 device families.

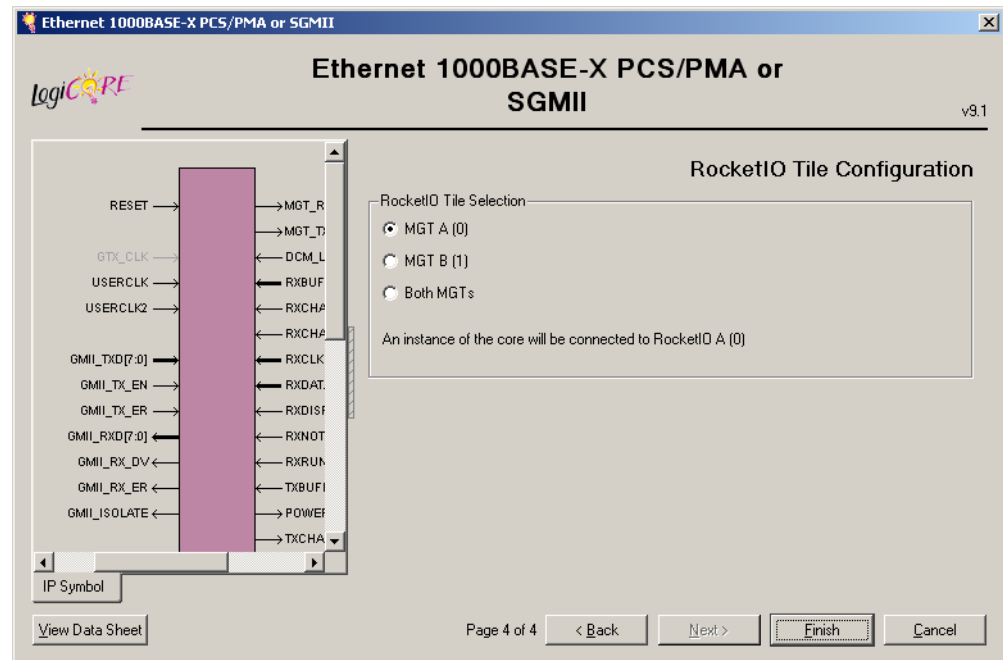


Figure 3-4: RocketIO Tile Configuration Screen

RocketIO transceivers for Virtex-4 FX and Virtex-5 device families are available in tiles, each tile consisting of a pair of transceivers. The RocketIO Tile Selection has no effect on the functionality of the core netlist, but determines the functionality of the example design delivered with the core.

Depending on the option selected, the example design instantiates a single core netlist and does one of the following:

- **MGT A (0)**. Connects to RocketIO transceiver A
- **MGT B (1)**. Connects to RocketIO transceiver B
- **Both MGTs**. Two instantiations of the core are created in the example design and connected to both RocketIO transceiver A and B.

Parameter Values in the XCO File

XCO file parameters are used to run the CORE Generator from the command line. XCO file parameter names and their values are similar to the names and values shown in the GUI, except that underscore characters (_) may be used instead of spaces. The text in an XCO file is not case sensitive.

Table 3-1 describes the XCO file parameters, values and summarizes the GUI defaults. The following is an example of the CSET parameters in an XCO file:

```
CSET component_name=gig_eth_pcs_pma_v9_1
CSET standard=1000BASEX
CSET physical_interface=TBI
CSET management_interface=true
CSET auto_negotiation=true
CSET mgt_crc_enabled=false
CSET sgmi_mode=10_100_1000
CSET rocketio_tile=A
```

Table 3-1: XCO File Values and Default Values

Parameter	XCO File Values	Default GUI Setting
component_name	ASCII text starting with a letter and based upon the following character set: a..z, 0..9 and _	gig_eth_pcs_pma_v9_1
standard	One of the following keywords: 1000BASEX, SGMII, Both	1000BASEX
physical_interface	One of the following keywords: TBI, RocketIO	TBI
management_interface	One of the following keywords: true, false	true
auto_negotiation	One of the following keywords: true, false	true
mgt_crc_enabled	One of the following keywords: true, false	false
sgmii_mode	One of the following keywords: 10_100_1000, 100_1000 <ul style="list-style-type: none"> • 10_100_1000 corresponds to “10/100/1000 Mbps (clock tolerance compliant with Ethernet specification)” • 100_1000 corresponds to “10/100/1000 Mbps (restricted tolerance for clocks) OR 100/1000 Mbps” 	10_100_1000
rocketio_tile	One of the following keywords: A, B, Both	A

Output Generation

The files output by the CORE Generator are placed in the CORE Generator project directory and include the following:

- The netlist file for the core
- Supporting CORE Generator files
- Release notes and documentation
- Subdirectories containing an HDL example design
- Scripts to run the core through the back-end tools and simulate the core using either Mentor Graphics® ModelSim®, Cadence® IUS, and Synopsys® simulators

See the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide* for a complete description of the CORE Generator output files, simulation requirements, and detailed information about the HDL example design.

Designing with the Core

This chapter provides information about creating your own designs using the Ethernet 1000BASE-X PCS/PMA or SGMII core. Design guidelines, as well as the variety of implementations presented, are based on the example design delivered with the core. See the *Xilinx Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide* for information about the example design delivered with the core.

Note that not all implementations require all of the design steps defined in this chapter. Carefully follow the provided logic design guidelines to ensure success.

Design Overview

An HDL example design built around the core is provided through the CORE Generator and allows for a demonstration of core functionality using either a simulation package or in hardware if placed on a suitable board. Four implementations of the core, based on the provided example design, are illustrated in the following sections.

- [“1000BASE-X Standard Using RocketIO Transceiver Example Design”](#)
- [“1000BASE-X Standard with TBI Example Design”](#)
- [“SGMII Standard Using a RocketIO Transceiver Example Design”](#)
- [“SGMII Standard with TBI Transceiver Example Design”](#)

1000BASE-X Standard Using RocketIO Transceiver Example Design

Figure 4-1 illustrates the example design in 1000BASE-X mode using the Virtex-II Pro or Virtex-4 MGT, Virtex-5 GTP or Virtex-5 GTX transceiver. As illustrated, the example is split between two hierarchical layers. The block level is designed so that it can be instantiated directly into your design and performs the following functions:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to a RocketIO transceiver

The top level of the example design creates a specific example that can be simulated, synthesized, implemented, and if required, placed on a suitable board and demonstrated in hardware. The top level of the example design performs the following functions:

- Instantiates the block level from HDL
- Derives the clock management logic for RocketIO and the core
- Implements an external GMII

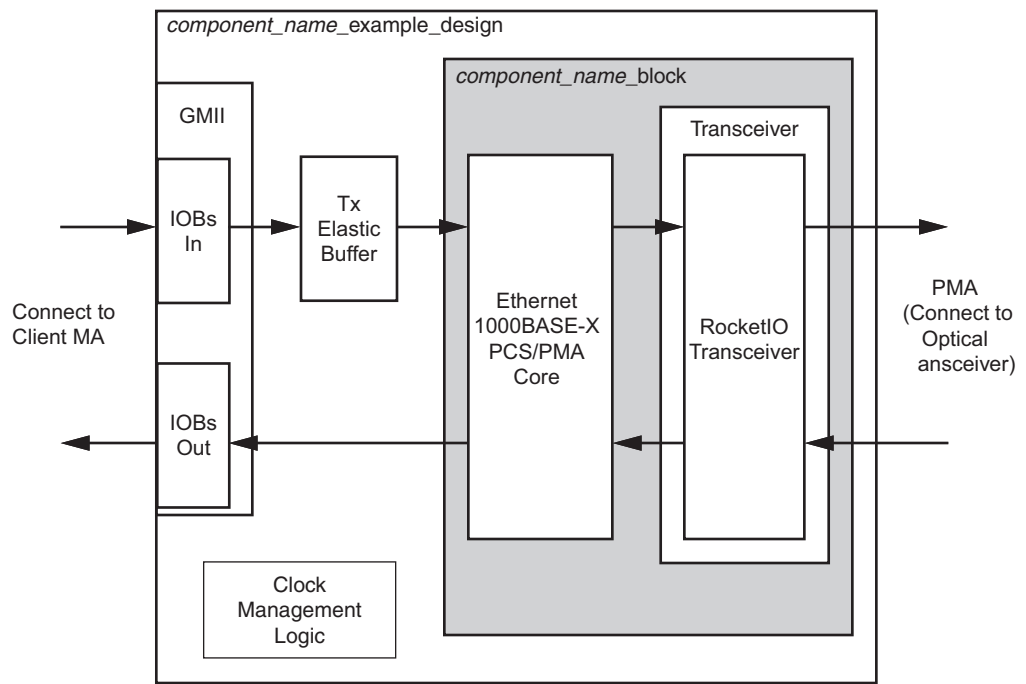


Figure 4-1: 1000BASE-X Standard Using a RocketIO Transceiver

1000BASE-X Standard with TBI Example Design

Figure 4-2 illustrates the example design in 1000BASE-X mode using a TBI. As illustrated, the example is split between two hierarchical layers. The block level is designed so that it can be instantiated directly into customer designs and performs the following functions:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to device IOBs, creating an external TBI. See Chapter 6, “The Ten-Bit Interface.”

The top level of the example design creates a specific example that can be simulated, synthesized, implemented, and if required, placed on a suitable board and demonstrated in hardware. The top level of the example design performs the following functions:

- Instantiates the block level from HDL
- Derives the clock management logic for the core
- Implements an external GMII

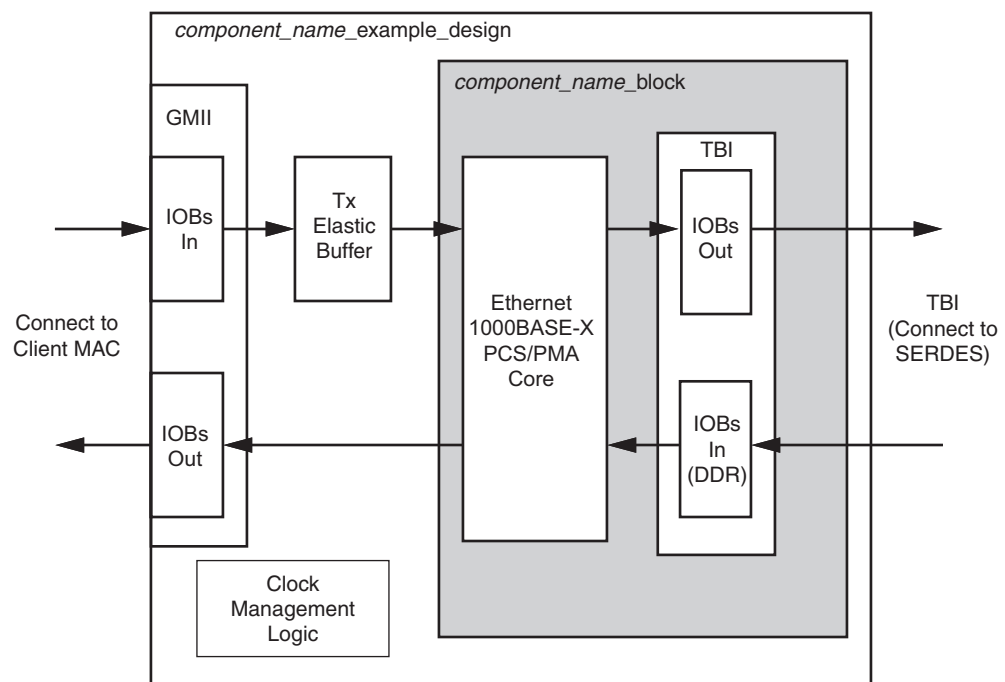


Figure 4-2: Example Design 1000BASE-X Standard Using TBI

SGMII Standard Using a RocketIO Transceiver Example Design

Figure 4-3 illustrates the example design in SGMII mode using the Virtex-II Pro or Virtex-4 MGT, Virtex-5 GTP or Virtex-5 GTX transceiver. This is also the example design created when the Dynamic Switching capability between SGMII and 1000BASE-X standards is present. As illustrated, the example is split between two hierarchical layers. The block level is designed so that it can be instantiated directly into customer designs and performs the following functions:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to a RocketIO transceiver
- Connects the client side GMII of the core to an SGMII Adaptation Module, which provides the functionality to operate at speeds of 1 Gbps, 100 Mbps and 10 Mbps

The top level of the example design creates a specific example which can be simulated, synthesized and implemented. The top level of the example design performs the following functions:

- Instantiates the block level from HDL
- Derives the clock management logic for RocketIO and the core
- Implements an external GMII-style interface

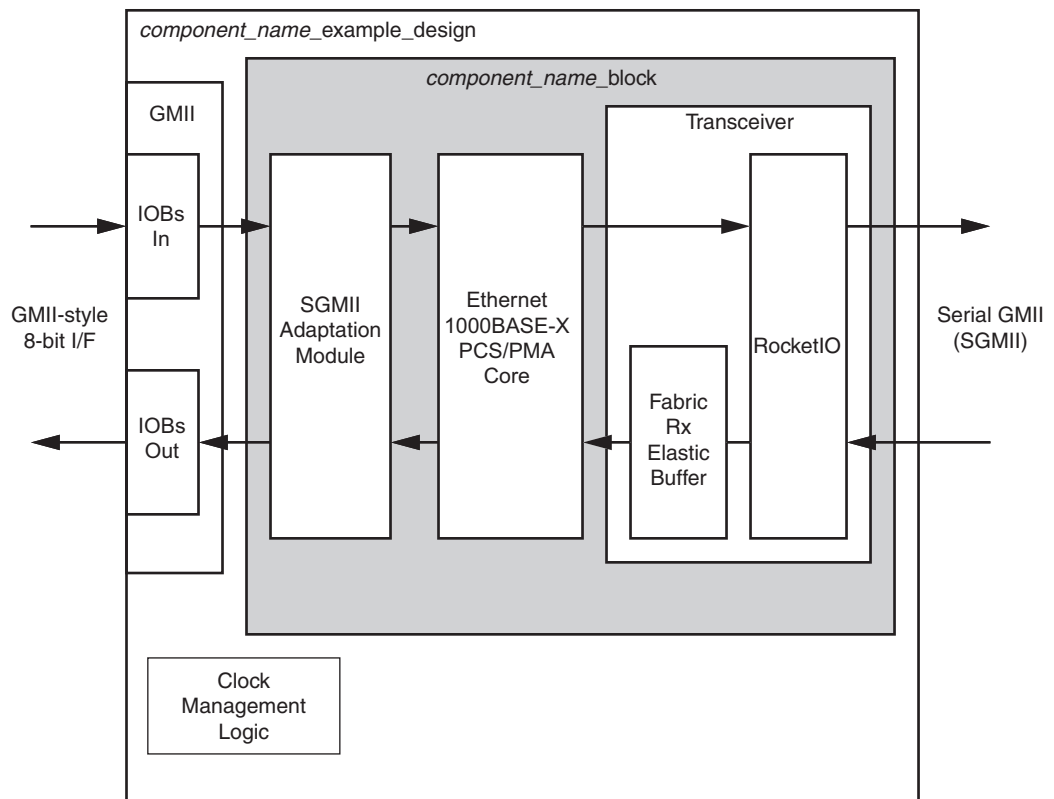


Figure 4-3: Example Design Performing the SGMII Standard

SGMII Standard with TBI Transceiver Example Design

Figure 4-3 illustrates the example design with the SGMII standard using a TBI. This is also the example design created when the Dynamic Switching capability between SGMII and 1000BASE-X standards is present. As illustrated, the example is split between two hierarchical layers. The block level is designed so that it can be instantiated directly into customer designs and performs the following functions:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to device IOBs, creating an external TBI. See Chapter 6, “The Ten-Bit Interface.”
- Connects the client side GMII of the core to an SGMII Adaptation Module, which provides the functionality to operate at speeds of 1 Gbps, 100 Mbps and 10 Mbps

The top level of the example design creates a specific example which can be simulated, synthesized and implemented. The top level of the example design performs the following functions:

- Instantiates the block level from HDL
- Derives the clock management logic for the core
- Implements an external GMII-style interface

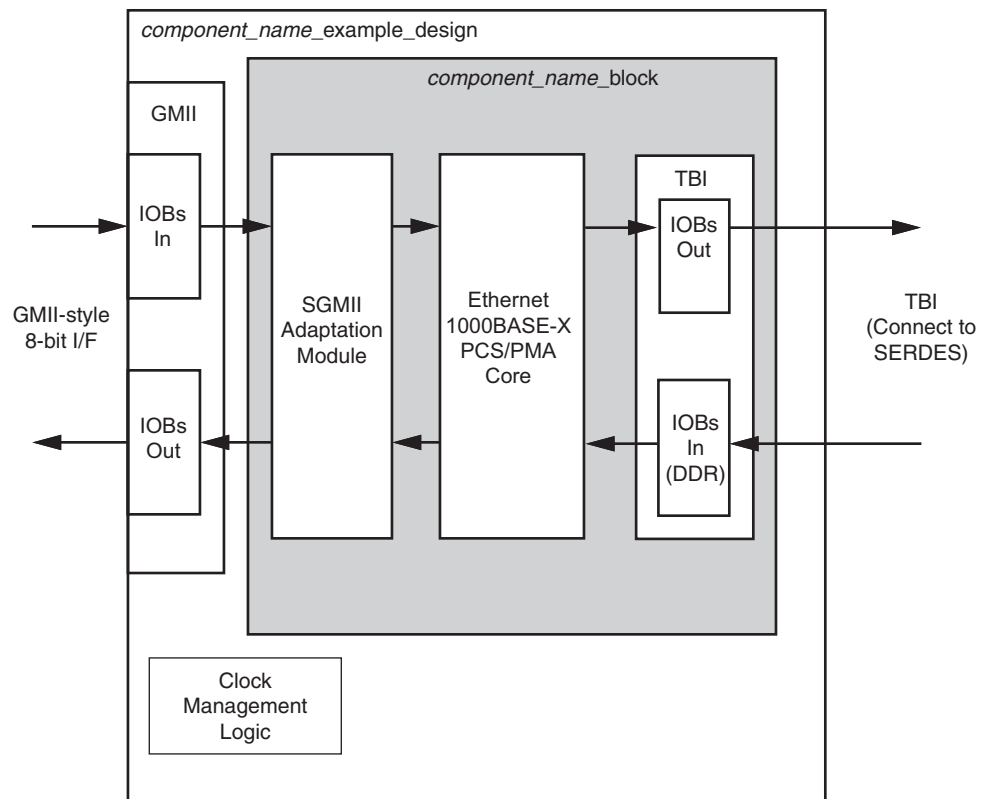


Figure 4-4: Example Design Performing the SGMII Standard

Design Guidelines

Generate the Core

Generate the core using the CORE Generator, as described in [Chapter 3, “Generating and Customizing the Core.”](#)

Examine the Example Design Provided with the Core

Before implementing the core in your application, examine the example design provided with the core to identify the steps that can be performed:

- Edit the HDL top level of the example design file to change the clocking scheme, add or remove IOBs as required, and replace the GMII IOB logic with user-specific application logic (for example, an Ethernet MAC).
- Synthesize the entire design.

The Xilinx Synthesis Tool (XST) script and Project file in the `/implement/vhdl` (or `/implement/verilog`) directory may be adapted to include any added user's HDL files.

- Run the `implement` script in the `/implement` directory to create a top-level netlist for the design.

The script may also run the Xilinx tools `map`, `par`, and `bitgen` to create a bitstream that can be downloaded to a Xilinx device. SimPrim-based simulation models for the entire design are also produced by the `implement` scripts.

- Simulate the entire design using the demonstration test bench provided in `/test/vhdl` (or `/test/verilog`) as a template.
- Download the bitstream to a target device.

Implement the Ethernet 1000BASE-X PCS/PMA or SGMII Core in Your Application

Before implementing your application, examine the example design delivered with the core for information about the following:

- Instantiating the core from HDL
- Connecting the physical-side interface of the core (RocketIO transceiver or TBI)
- Deriving the clock management logic

It is expected that the block level module from the example design will be instantiated directly into customer designs rather than the core netlist itself. The block level contains the core and a completed physical interface.

Write an HDL Application

After reviewing the example design delivered with the core, write an HDL application that uses single or multiple instances of the block level module for the Ethernet 1000BASE-X PCS/PMA or SGMII core. Client-side interfaces and operation of the core are described in [Chapter 5, “Using the Client-side GMII Data Path.”](#) See the following information for additional details:

- Using the Ethernet 1000BASE-X PCS/PMA or SGMII core in conjunction with the 1-Gigabit Ethernet MAC core in [“Integrating with the 1-Gigabit Ethernet MAC Core,”](#) page 179.
- Using the Ethernet 1000BASE-X PCS/PMA or SGMII core in conjunction with the Tri-Mode Ethernet MAC core in [“Integrating with the Tri-Mode Ethernet MAC Core,”](#) page 185.

Synthesize your Design

Synthesize your entire design using the desired synthesis tool. The Ethernet 1000BASE-X PCS/PMA or SGMII core is pre-synthesized and delivered as an NGC netlist—for this reason, it appears as a black box to synthesis tools.

Create a Bitstream

Run the Xilinx tools **map**, **par**, and **bitgen** to create a bitstream that can be downloaded to a Xilinx device. The UCF produced by the CORE Generator should be used as the basis for the user UCF and care must be taken to constrain the design correctly. See [Chapter 12, “Constraining the Core”](#) for more information.

Simulate and Download your Design

After creating a bitstream that can be downloaded to a Xilinx device, simulate the entire design and download it to the desired device.

Know the Degree of Difficulty

An Ethernet 1000BASE-X PCS/PMA or SGMII core is challenging to implement in any technology and as such, all Ethernet 1000BASE-X PCS/PMA or SGMII core applications require careful attention to system performance requirements. Pipelining, logic mapping, placement constraints, and logic duplication are all methods that help boost system performance.

Review [Table 4-1](#) to determine the relative level of difficulty associated with different designs. This relates to meeting the core’s required system clock frequency of 125 MHz.

Table 4-1: Degree of Difficulty for Various Implementations

Device Family	Difficulty
Virtex-II	Easy
Virtex-II Pro	Easy
Virtex-4	Easy
Virtex-5	Easy
Spartan™-3	Difficult

Table 4-1: Degree of Difficulty for Various Implementations

Device Family	Difficulty
Spartan-3E	Difficult
Spartan-3A	Difficult

Keep it Registered

To simplify timing and to increase system performance in an FPGA design, keep all inputs and outputs registered between the user application and the core. All inputs and outputs from the user application should come *from*, or connect *to*, a flip-flop. While registering signals may not be possible for all paths, it simplifies timing analysis and makes it easier for the Xilinx tools to place and route the design.

Recognize Timing Critical Signals

The UCF provided with the example design for the core identifies the critical signals and the timing constraints that should be applied. See [Chapter 12, “Constraining the Core”](#) for more information.

Use Supported Design Flows

The core is pre-synthesized and is delivered as an NGC netlist. The example implementation scripts provided currently uses ISE 10.1 as the synthesis tool for the HDL example design delivered with the core. Other synthesis tools may be used for the user application logic. The core will always be unknown to the synthesis tool and should appear as a black box. Post synthesis, only ISE 10.1i tools are supported.

Make Only Allowed Modifications

The Ethernet 1000BASE-X PCS/PMA or SGMII core should not be modified. Modifications may have adverse effects on system timing and protocol compliance. Supported user configurations of the Ethernet 1000BASE-X PCS/PMA or SGMII core can only be made by the selecting the options from within CORE Generator when the core is generated. See [Chapter 3, “Generating and Customizing the Core.”](#)

Using the Client-side GMII Data Path

This chapter provides general guidelines for creating designs using client-side GMII of the Ethernet 1000BASE-X PCS/PMA or SGMII core.

Designing with the Client-side GMII for the 1000BASE-X Standard

It is not within the scope of this document to define the Gigabit Media Independent Interface (GMII)— see clause 35 of the *IEEE 802.3* specification for information about the GMII. Timing diagrams and descriptions are provided only as an informational guide.

GMII Transmission

This section includes figures that illustrate GMII transmission. In these figures the clock is not labeled. The source of this clock signal varies, depending on the options selected when the core is generated. For more information on clocking, see Chapters 6, 7 and 8.

Normal Frame Transmission

Normal outbound frame transfer timing is illustrated in [Figure 5-1](#). This figure shows that an Ethernet frame is preceded by an 8-byte preamble field (inclusive of the Start of Frame Delimiter (SFD)), and completed with a 4-byte Frame Check Sequence (FCS) field. This frame is created by the MAC connected to the other end of the GMII. The PCS logic itself does not recognize the different fields within a frame and will treat any value placed on `gmi_i_txd[7:0]` within the `gmi_i_tx_en` assertion window as data.




Figure 5-1: GMII Normal Frame Transmission

Error Propagation

A corrupted frame transfer is illustrated in Figure 5-2. An error may be injected into the frame by asserting `gmii_tx_er` at any point during the `gmii_tx_en` assertion window. The core ensures that all errors are propagated through both transmit and receive paths so that the error is eventually detected by the link partner.




Figure 5-2: GMII Error Propagation Within a Frame

GMII Reception

This section includes figures that illustrate GMII reception. In these figures the clock is not labelled. The source of this clock signal will vary, depending on the options used when the core is generated. For more information on clocking, see Chapters 6, 7 and 8.

Normal Frame Reception

The timing of normal inbound frame transfer is illustrated in Figure 5-3. This shows that Ethernet frame reception is preceded by a preamble field. The IEEE 802.3 specification (see clause 35) allows for up to all of the seven preamble bytes that proceed the Start of Frame Delimiter (SFD) to be lost in the network. The SFD will always be present in well-formed frames.




Figure 5-3: GMII Normal Frame Reception

Normal Frame Reception with Extension Field

In accordance with the *IEEE 802.3*, clause 36, state machines for the 1000BASE-X PCS, `gmii_rx_er` may be driven high following reception of the end frame in conjunction with `gmii_rxd[7:0]` containing the hexadecimal value of 0x0F to signal carrier extension. This is illustrated in [Figure 5-4](#). See [Appendix D, “1000BASE-X State Machines”](#) for more information.

This is not an error condition and may occur even for full-duplex frames.




Figure 5-4: GMII Normal Frame Reception with Carrier Extension

Frame Reception with Errors

The signal `gmii_rx_er` when asserted within the assertion window signals that a frame was received with a detected error ([Figure 5-5](#)). In addition, a late error may also be detected during the Carrier Extension interval. This is indicated by `gmii_rxd[7:0]` containing the hexadecimal value 0x1F, also illustrated in [Figure 5-5](#).




Figure 5-5: GMII Frame Reception with Errors

False Carrier

Figure 5-6 illustrates the GMII signaling for a False Carrier condition. False Carrier is asserted by the core in response to certain error conditions, such as a frame with a corrupted start code, or for random noise.




Figure 5-6: False Carrier Indication

status_vector[4:0] signals

Bit[0]: Link Status

This signal indicates the status of the link. This information is duplicated in the optional PCS Management Registers, if present (“Status Register (Register 1),” bit 1.2). However, this would always serve a useful function as a Link Status LED.

When high, the link is valid: synchronization of the link has been obtained **and** Auto-Negotiation (if present and enabled) has completed.

When low, a valid link has not been established. Either link synchronization has failed or Auto-Negotiation (if present and enabled) has failed to complete.

Note: this bit is identical to the SYNC_ACQUIRED_STATUS port which was present in previous versions of the core.

Bit[1]: Link Synchronization

This signal indicates the state of the synchronization state machine (IEEE802.3 figure 36-9). This signal is similar to Bit[0] (Link Status), but is NOT qualified with Auto-Negotiation.

When high, link synchronization has been obtained.

When low, synchronization has failed.

Bits[4:2]: Code Group Reception Indicators

These signals indicate the reception of particular types of group, as defined below. [Figure 5-7](#) illustrates the timing of these signals, showing that they are aligned with the code groups themselves, as they appear on the output `gmii_rxd[7:0]` port.

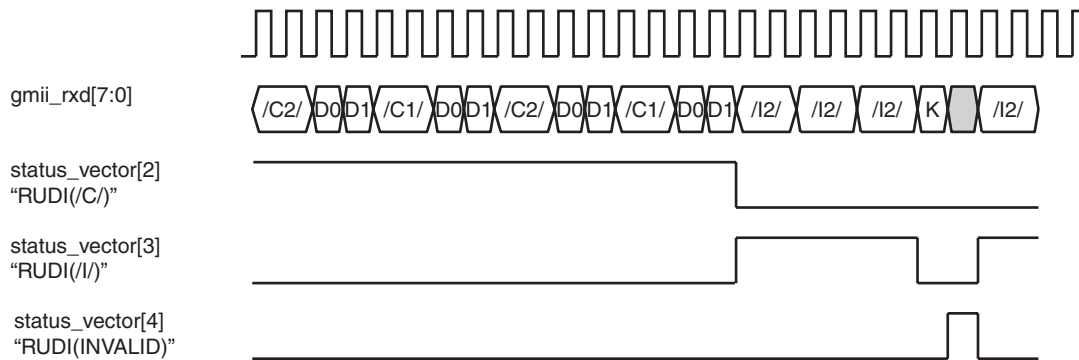


Figure 5-7: `status_vector[4:2]` timing

Bit[2]: RUDI(/C/)

The core is receiving `/C/` ordered sets (Auto-Negotiation Configuration sequences) as defined in IEEE802.3 clause 36.2.4.10.

Bit[3]: RUDI(/I/)

The core is receiving `/I/` ordered sets (Idles) as defined in IEEE802.3 clause 36.2.4.12.

Bit[4]: RUDI(INVALID)

The core has received invalid data whilst receiving `/C/` or `/I/` ordered set as defined in IEEE802.3 clause 36.2.5.1.6. This may be caused, for example, by bit errors occurring in any clock cycle of the `/C/` or `/I/` ordered set: [Figure 5-7](#) illustrates an error occurring in the second clock cycle of an `/I/` idle sequence.

Using the Virtex-II Pro RocketIO Transceiver CRC Functionality

When the core is generated with the Virtex-II Pro RocketIO transceiver, the CRC functionality of the RocketIO transceiver may be enabled. When the core is generated in this mode (see "[RocketIO Transceiver CRC Logic](#)" in [Chapter 3](#)), the core ensures that the `/K28.5/` characters are left-justified in the RocketIO transceiver internal two-byte data path. This is done by ensuring that the `/K28.5/` is strobed into the RocketIO transceiver on the rising edge of `TXUSRCLK2` only when `TXUSRCLK` is high. For more information, see the *Virtex-II Pro RocketIO Transceiver User Guide*.

Caution! Do not use the Virtex-II Pro RocketIO CRC functionality when using the SGMII standard.

GMII Transmission

The timing of normal outbound frame transfer with the RocketIO transceiver CRC functionality is illustrated in [Figure 5-8](#). This figure shows that an Ethernet frame can be completed by allowing the RocketIO transceiver to create the Frame Check Sequence field (FCS) using the in-built CRC logic. For this to be successful, four place-holder bytes must

be included in the frame supplied to the core. The RocketIO transceiver will replace these four bytes with the calculated CRC value.




Figure 5-8: GMII Frame Transmission with RocketIO Transceiver CRC Logic Enabled

GMII Reception

The timing of normal inbound frame transfer with RocketIO transceiver CRC functionality is illustrated in Figure 5-9. The RocketIO transceiver calculates the CRC value of the received frame and checks it against that contained in the frames FCS field. The RocketIO transceiver will assert RXCHECKINGCRC and RXCRCERR signals, as defined in the *Virtex-II Pro RocketIO Transceiver User Guide*. Figure 5-9 illustrates a frame received with a correct FCS field since RXCRCERR is not asserted.

Please note that RXCHECKINGCRC and RXCRCERR are obtained directly from the output of the RocketIO transceiver. The core receiver behavior is unchanged.




Figure 5-9: GMII Frame Reception with the RocketIO Transceiver CRC Logic Enabled

Designing with Client-side GMII for the SGMII Standard

Overview

When the core is generated for the SGMII standard, changes are made to the core that affect the PCS Management Registers and the Auto-Negotiation function (see “Select Standard” in Chapter 3). However, the data path through both transmitter and receiver sections of the core remains unchanged.

GMII Transmission

1 Gigabit per Second Frame Transmission

The timing of normal outbound frame transfer is illustrated in Figure 5-10. At 1 Gbps speed, the operation of the transmitter GMII signals remains identical to that described in “Designing with the Client-side GMII for the 1000BASE-X Standard.”




Figure 5-10: GMII Frame Transmission at 1 Gbps

100 Megabit per Second Frame Transmission

The operation of the core remains unchanged. It is the responsibility of the client logic (for example, an Ethernet MAC) to enter data at the correct rate. When operating at a speed of 100 Mbps, every byte of the MAC frame (from preamble field to the Frame Check Sequence field, inclusive) should each be repeated for 10 clock periods to achieve the desired bit rate, as illustrated in Figure 5-11. It is also the responsibility of the client logic to ensure that the interframe gap period is legal for the current speed of operation.




Figure 5-11: GMII Data Transmission at 100 Mbps

10 Megabit per Second Frame Transmission

The operation of the core remains unchanged. It is the responsibility of the client logic (for example, an Ethernet MAC), to enter data at the correct rate. When operating at a speed of 10 Mbps, every byte of the MAC frame (from destination address to the frame check sequence field, inclusive) should each be repeated for 100 clock periods to achieve the desired bit rate. It is also the responsibility of the client logic to ensure that the interframe gap period is legal for the current speed of operation.

GMII Reception

1 Gigabit per Second Frame Reception

The timing of normal inbound frame transfer is illustrated in Figure 5-12. At 1 Gbps speed, the operation of the receiver GMII signals remains identical to that described in “Designing with the Client-side GMII for the 1000BASE-X Standard” in Chapter 5.

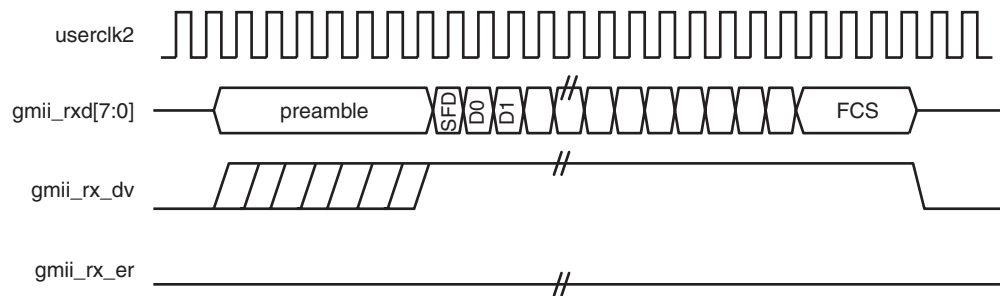


Figure 5-12: GMII Frame Reception at 1 Gbps

100 Megabit per Second Frame Reception

The operation of the core remains unchanged. When operating at a speed of 100 Mbps, every byte of the MAC frame (from destination address to the frame check sequence field, inclusive) is repeated for 10 clock periods to achieve the desired bit rate. See Figure 5-13. It is the responsibility of the client logic, for example an Ethernet MAC, to sample this data correctly.




Figure 5-13: GMII Data Reception at 100 Mbps

10 Megabit per Second Frame Reception

The operation of the core remains unchanged. When operating at a speed of 10 Mbps, every byte of the MAC frame (from destination address to the frame check sequence field, inclusive) is repeated for 100 clock periods to achieve the desired bit rate. It is the responsibility of the client logic (for example, an Ethernet MAC) to sample this data correctly.

Using the GMII as an Internal Connection

The client-side GMII of the core may be used to connect to an internally integrated Media Access Controller. For details, see [“Integrating with the 1-Gigabit Ethernet MAC Core,” page 179](#) and [“Integrating with the Tri-Mode Ethernet MAC Core,” page 185](#).

Implementing External GMII

GMII Transmitter Logic

When implementing an external GMII, the GMII transmitter signals will be synchronous to their own clock domain. The core must be used with a Transmitter Elastic Buffer to transfer these GMII transmitter signals onto the core's internal 125 MHz reference clock (`gtx_clk` when using the TBI; `userclk2` when using the RocketIO transceiver). A Transmitter Elastic Buffer is provided for the 1000BASE-X standard by the example design provided with the core. See the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide* for more information.

Virtex-II Pro and Virtex-II Devices

Figure 5-14 illustrates how to create an external GMII transmitter in a Virtex-II family device. The signal names and logic shown on the figure exactly match those delivered with the example design.

Figure 5-14 shows that the input transmitter signals are registered in device IOBs before presenting them to the FPGA fabric. This logic achieves the required setup and hold times.

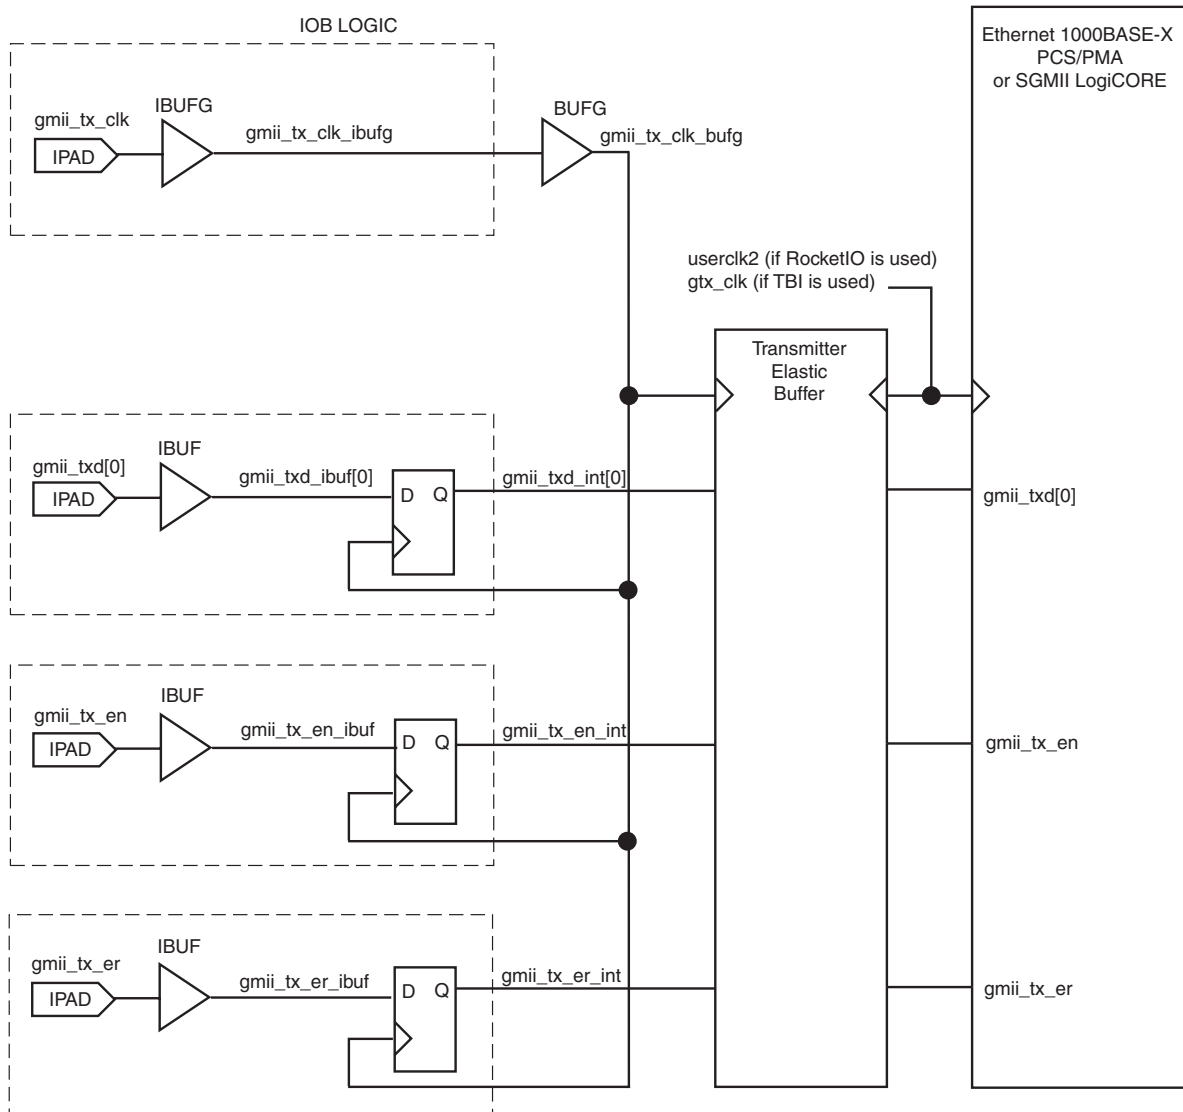


Figure 5-14: GMII Transmitter Logic

Spartan-3, Spartan-3E and Spartan-3A Devices

The logic described previously for Virtex-II and Virtex-II Pro devices does not meet the input setup and hold requirements for GMII with Spartan-3, Spartan-3E, and Spartan-3A devices. A DCM must be used on the `gmii_tx_clk` clock path, as illustrated in Figure 5-15. This is performed by the top-level example design delivered with the core (all signal names and logic match Figure 5-15). This DCM circuitry may optionally be used in other families.

Phase-shifting may then be applied to the DCM to fine-tune the setup and hold times at the GMII IOB input flip-flops. The fixed phase shift is applied to the DCM with the example UCF for the example design. See “Constraints When Implementing an External GMII” in Chapter 12.

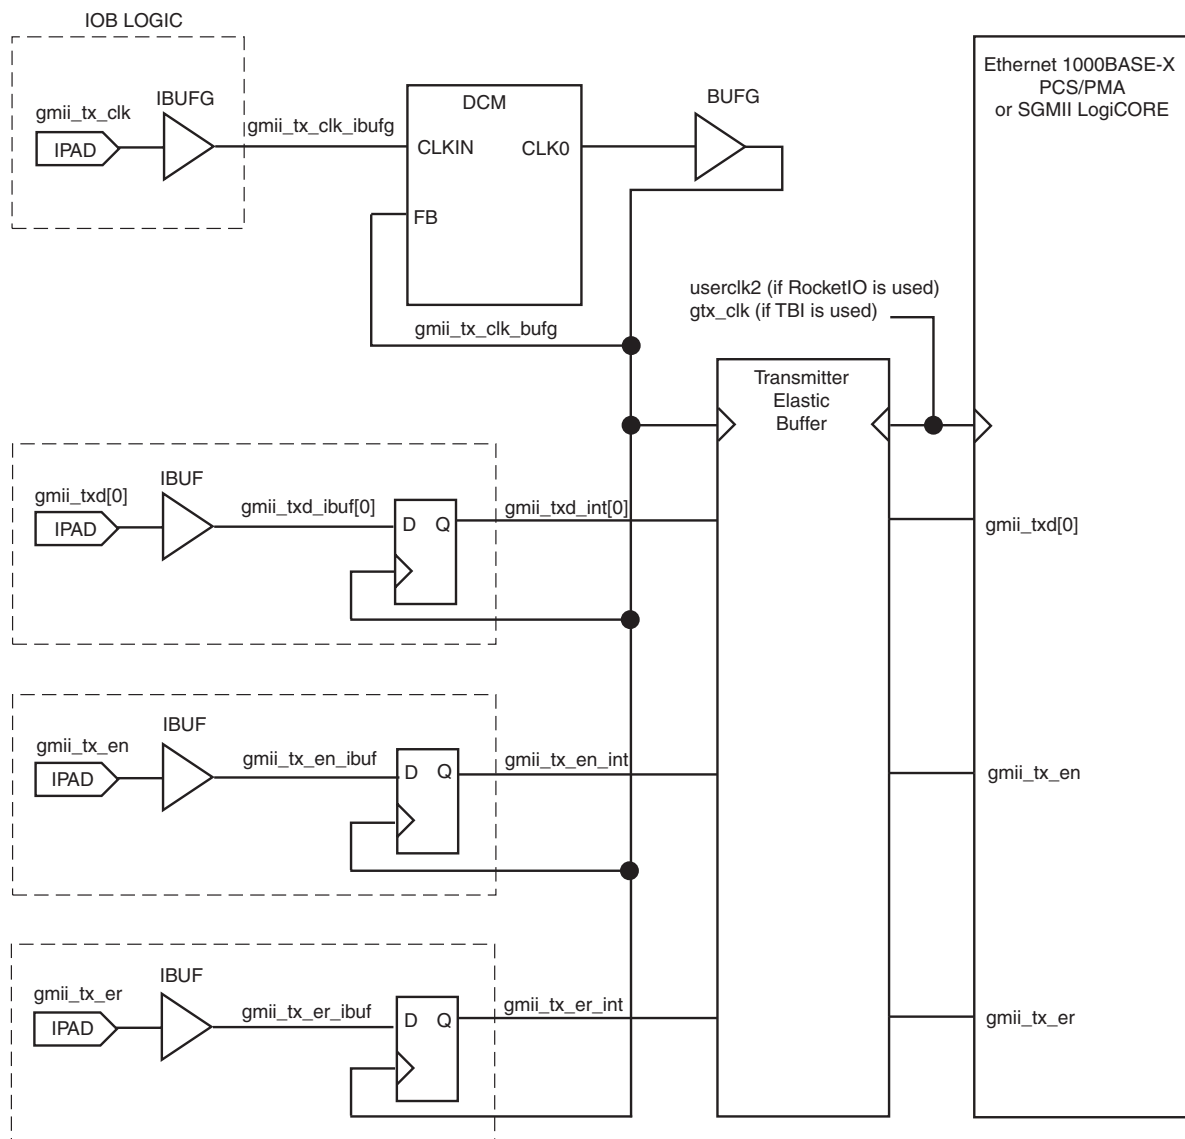


Figure 5-15: External GMII Transmitter Logic for Spartan-3, Spartan-3E and Spartan-3A Devices

Virtex-4 Devices

The logic described previously for Virtex-II and Virtex-II Pro devices does not meet the input setup and hold requirements for GMII with Virtex-4 devices. Two possible solutions are:

1. A DCM may be used on the `gmii_tx_clk` clock path for the Spartan-3 family, as illustrated in [Figure 5-15](#).
2. Input Delay Elements may be used on the GMII data path, as illustrated in [Figure 5-16](#).

The IODELAY elements can be adjusted to fine-tune the setup and hold times at the GMII IOB input flip-flops. The delay is applied to the IODELAY element using constraints in the UCF; these can be edited if desired. See [“Constraints When Implementing an External GMII”](#) in [Chapter 12](#) for more information.

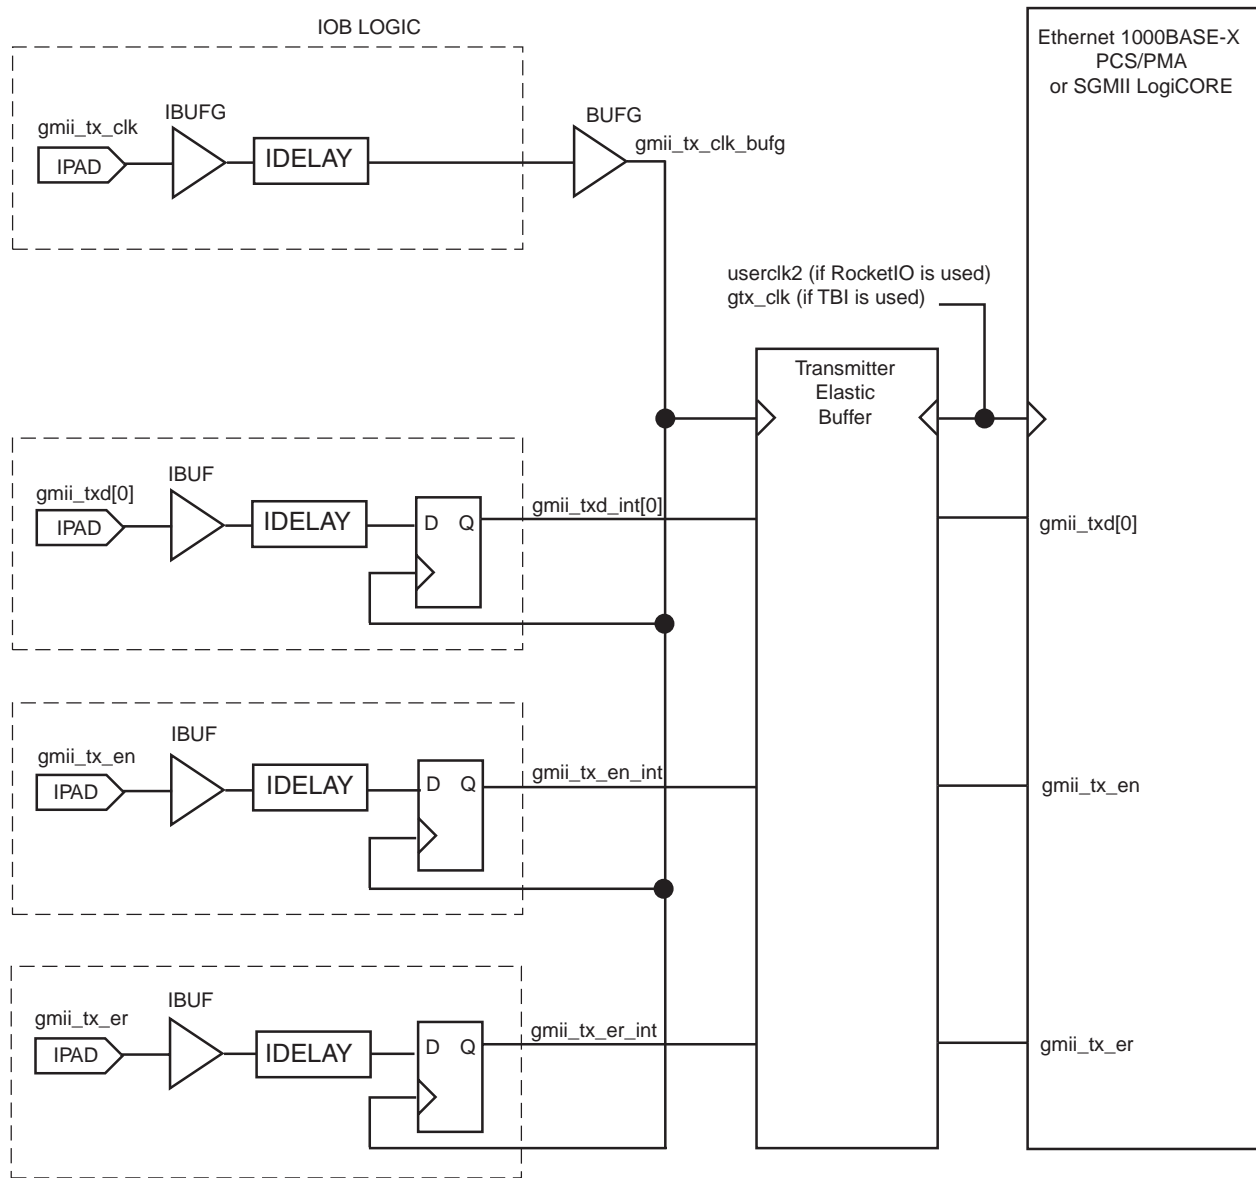


Figure 5-16: External GMII Transmitter Logic for Virtex-4 Devices

Virtex-5 Devices

Figure 5-17 illustrates how to create an external GMII transmitter in a Virtex-5 family device. The signal names and logic shown on the figure exactly match those delivered with the example design.

The IODELAY elements can be adjusted to fine-tune the setup and hold times at the GMII IOB input flip-flops. The delay is applied to the IODELAY element using constraints in the UCF; these can be edited if desired. See “Constraints When Implementing an External GMII” in Chapter 12 for more information.

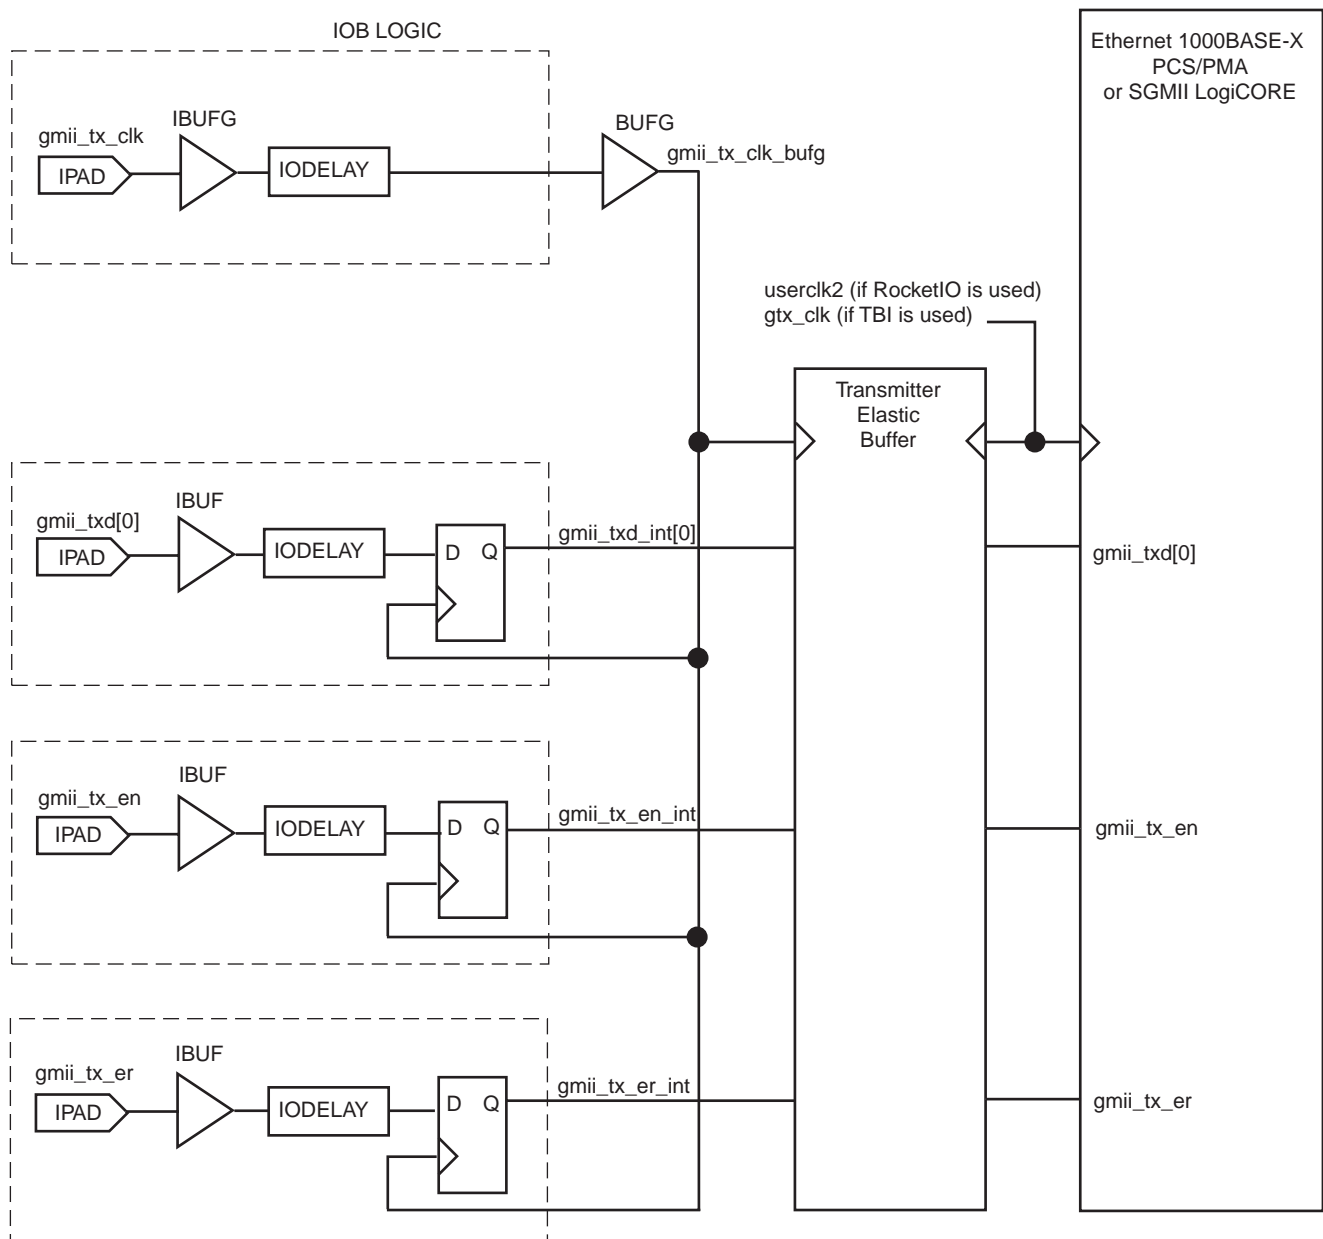


Figure 5-17: External GMII Transmitter Logic for Virtex-5 Devices

GMII Receiver Logic

Figure 5-18 illustrates an external GMII receiver created in a Virtex-II family device. The signal names and logic shown in the figure exactly match those delivered with the example design when the GMII is selected. If other families are selected, equivalent primitives and logic specific to that family is automatically used in the example design.

Figure 5-18 also shows that the output receiver signals are registered in device IOBs before driving them to the device pads. The logic required to forward the receiver GMII clock is also shown. This uses an IOB output Double-Data-Rate (DDR) register so that the clock signal produced incurs exactly the same delay as the data and control signals. This clock signal, `gmi_i_rx_clk`, is inverted so that the rising edge of `gmi_i_rx_clk` occurs in the center of the data valid window, which maximizes setup and hold times across the interface. All receiver logic is synchronous to a single clock domain.

The clock name varies depending on the CORE Generator options. When used with the RocketIO transceiver, the clock name is `userclk2`; when used with the TBI, the clock name is `gtx_clk`. For more information on clocking, see Chapters 6, 7 and 8.

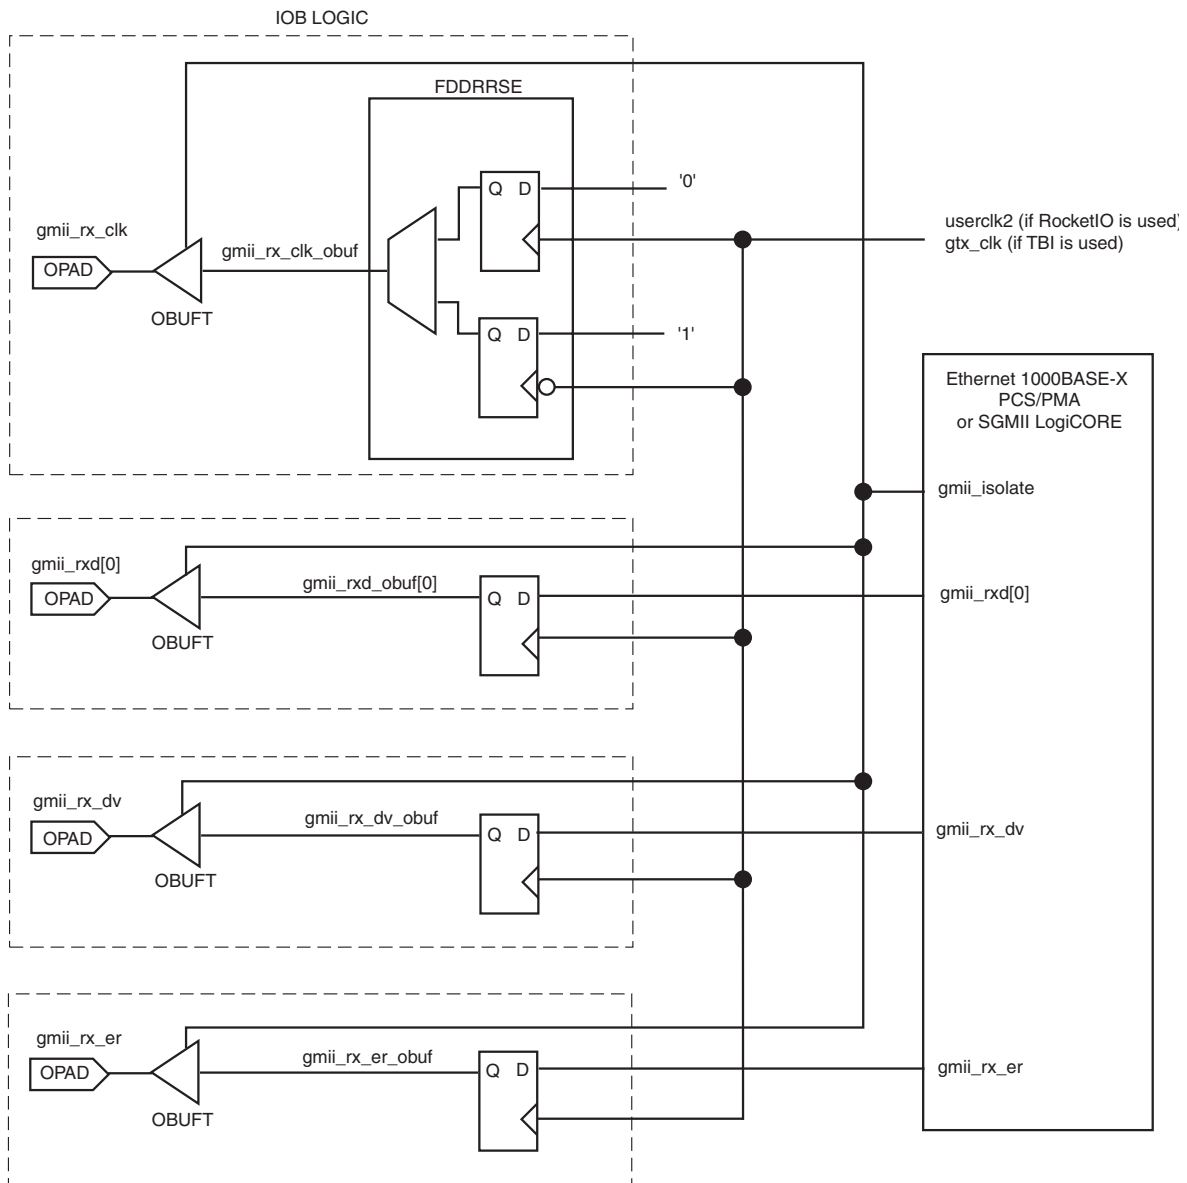


Figure 5-18: External GMII Receiver Logic

The Ten-Bit Interface

This chapter provides general guidelines for creating 1000BASE-X, SGMII or Dynamic Standards Switching designs using the Ten-Bit Interface (TBI). An explanation of the TBI logic in all supported device families is provided, as well as examples in which multiple instantiations of the core are required. Whenever possible, clock sharing should occur to save device resources.

Ten-Bit-Interface Logic

The example design delivered with the core is split between two hierarchical layers, as illustrated in [Figure 4-2](#). The block level is designed so that it can be instantiated directly into customer designs and provides the following functionality:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to device IOBs, creating an external TBI

The TBI logic implemented in the block level is illustrated in all the figures in this chapter.

Transmitter Logic

[Figure 6-1](#) illustrates the use of the physical transmitter interface of the core to create an external TBI in a Virtex-II family device. The signal names and logic shown exactly match those delivered with the example design when TBI is chosen. If other families are chosen, equivalent primitives and logic specific to that family will automatically be used in the example design.

[Figure 6-1](#) shows that the output transmitter data path signals are registered in device IOBs before driving them to the device pads. The logic required to forward the transmitter clock is also shown. The logic uses an IOB output Double-Data-Rate (DDR) register so that the clock signal produced incurs exactly the same delay as the data and control signals. This clock signal, `pma_tx_clk`, is inverted with respect to `gtx_clk` so that the rising edge of `pma_tx_clk` occurs in the center of the data valid window to maximize setup and hold times across the interface.

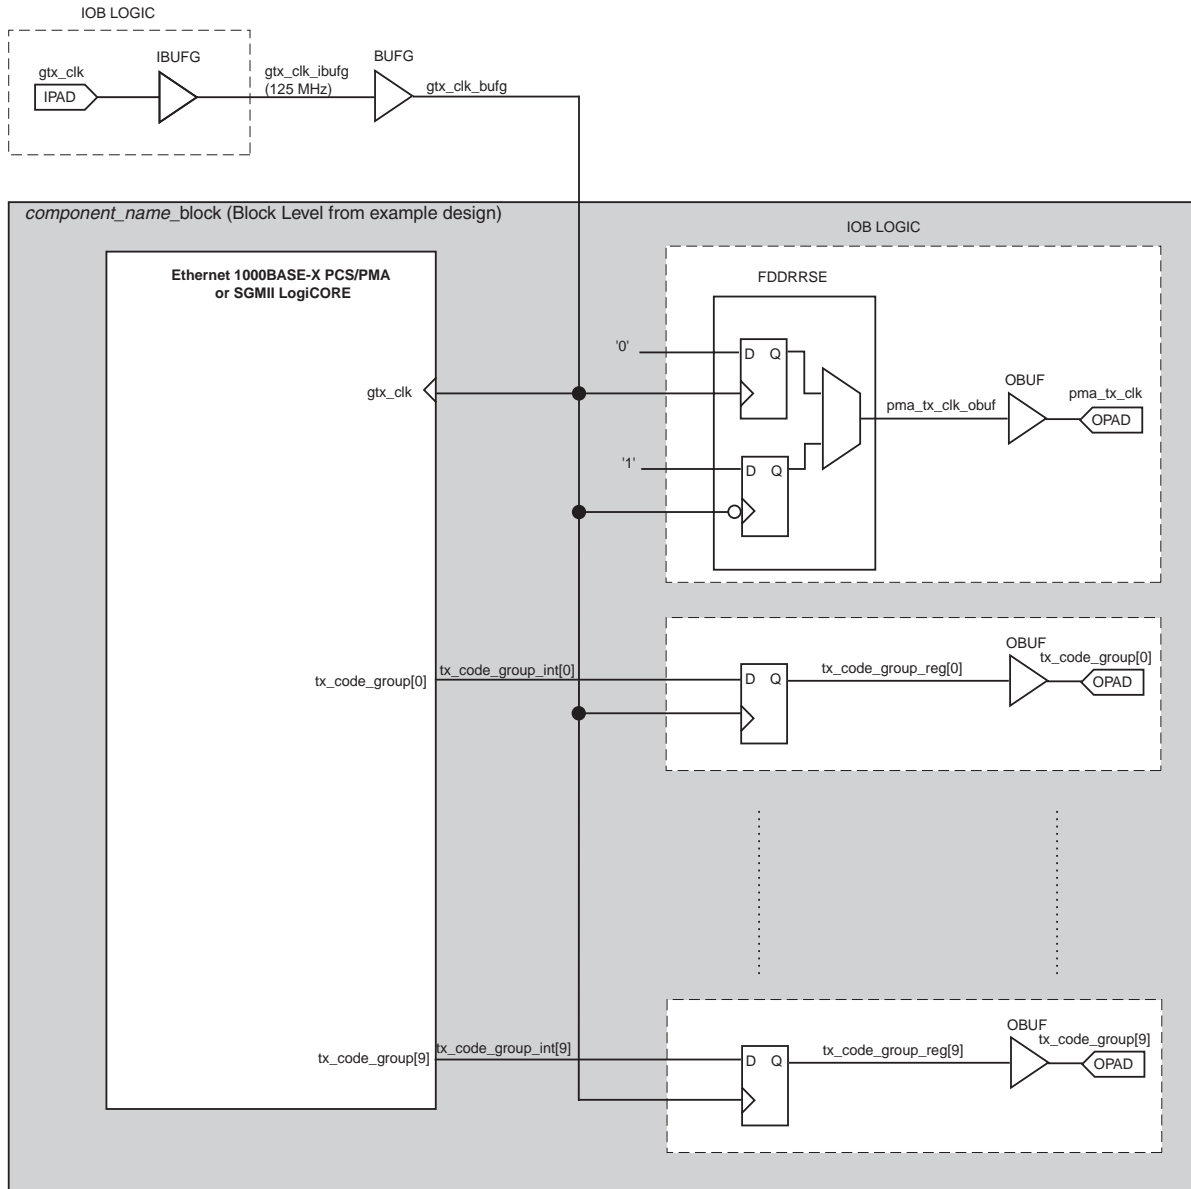


Figure 6-1: Ten-Bit Interface Transmitter Logic

Receiver Logic

Virtex-II and Virtex-II Pro Devices

Figure 6-2 illustrates an external receiver TBI in Virtex-II devices. The signal names and logic displayed precisely match those delivered with the example design when the TBI is chosen.

Figure 6-2 shows that the input receiver signals are registered in device IOB Double-Data Rate (DDR) input registers, alternatively on the rising edges of both `pma_rx_clk0_bufg` and `pma_rx_clk1_bufg` (`pma_rx_clk0` and `pma_rx_clk1` are 180 degrees out of phase with each other). This splits the input TBI data bus, `rx_code_group[9:0]`, up into two buses: `rx_code_group0_reg[9:0]` and `rx_code_group1_reg[9:0]`,

synchronous to pma_rx_clk0_bufg and pma_rx_clk1_bufg, respectively. These busses are then immediately registered inside the core on their respective clock.

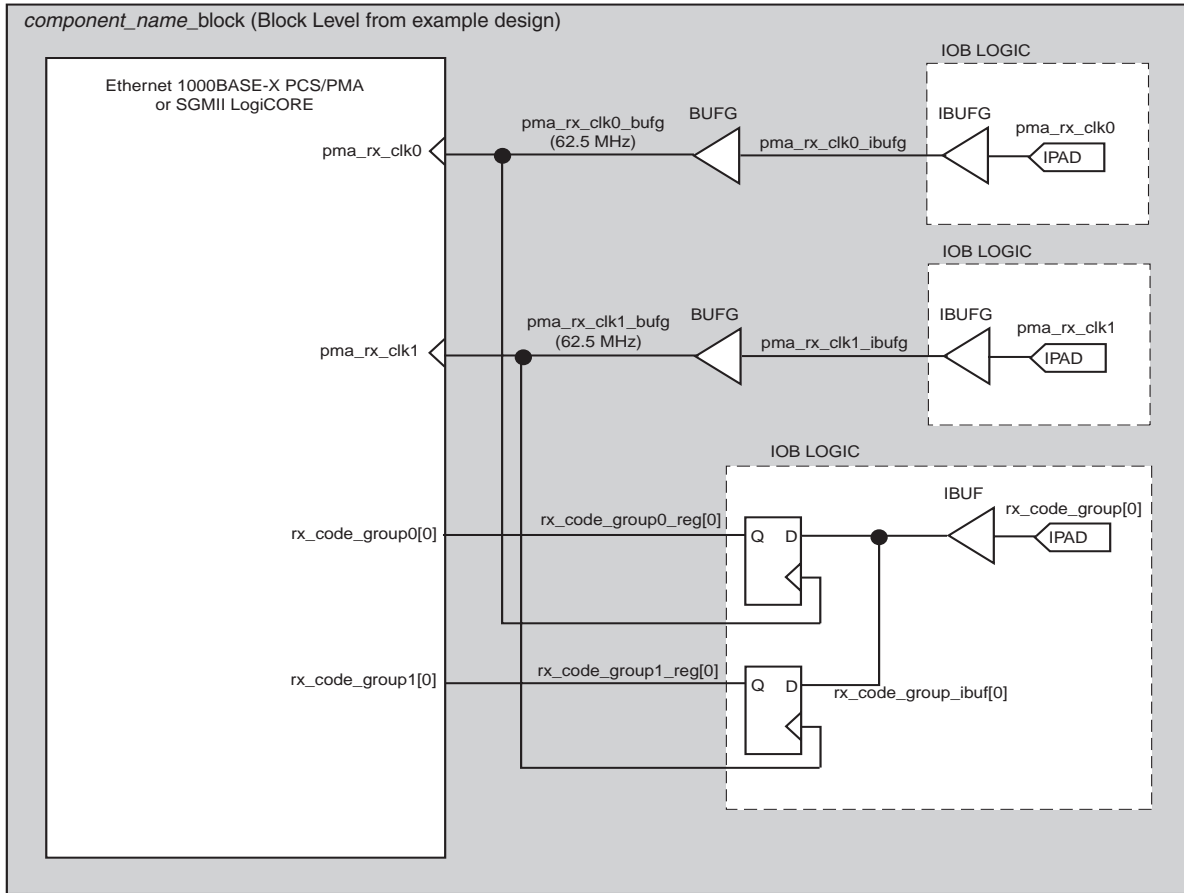


Figure 6-2: Ten-Bit-Interface Receiver Logic

Spartan-3, Spartan-3E and Spartan-3A Devices

The logic described previously for Virtex-II and Virtex-II Pro devices does not meet the input setup and hold requirements for TBI with Spartan-3, Spartan-3E and Spartan-3A devices. A DCM must be used on both the pma_rx_clk0 and pma_rx_clk1 clock paths (see Figure 6-3). This is performed by the example design delivered with the core (all signal names and logic match Figure 6-3).

Phase shifting may then be applied to the DCM to fine-tune the setup and hold times at the TBI IOB input flip-flops. Fixed phase shift is applied to the DCM using constraints in the example UCF for the example design. See “Constraints When Implementing an External GMII” in Chapter 12 for more information.

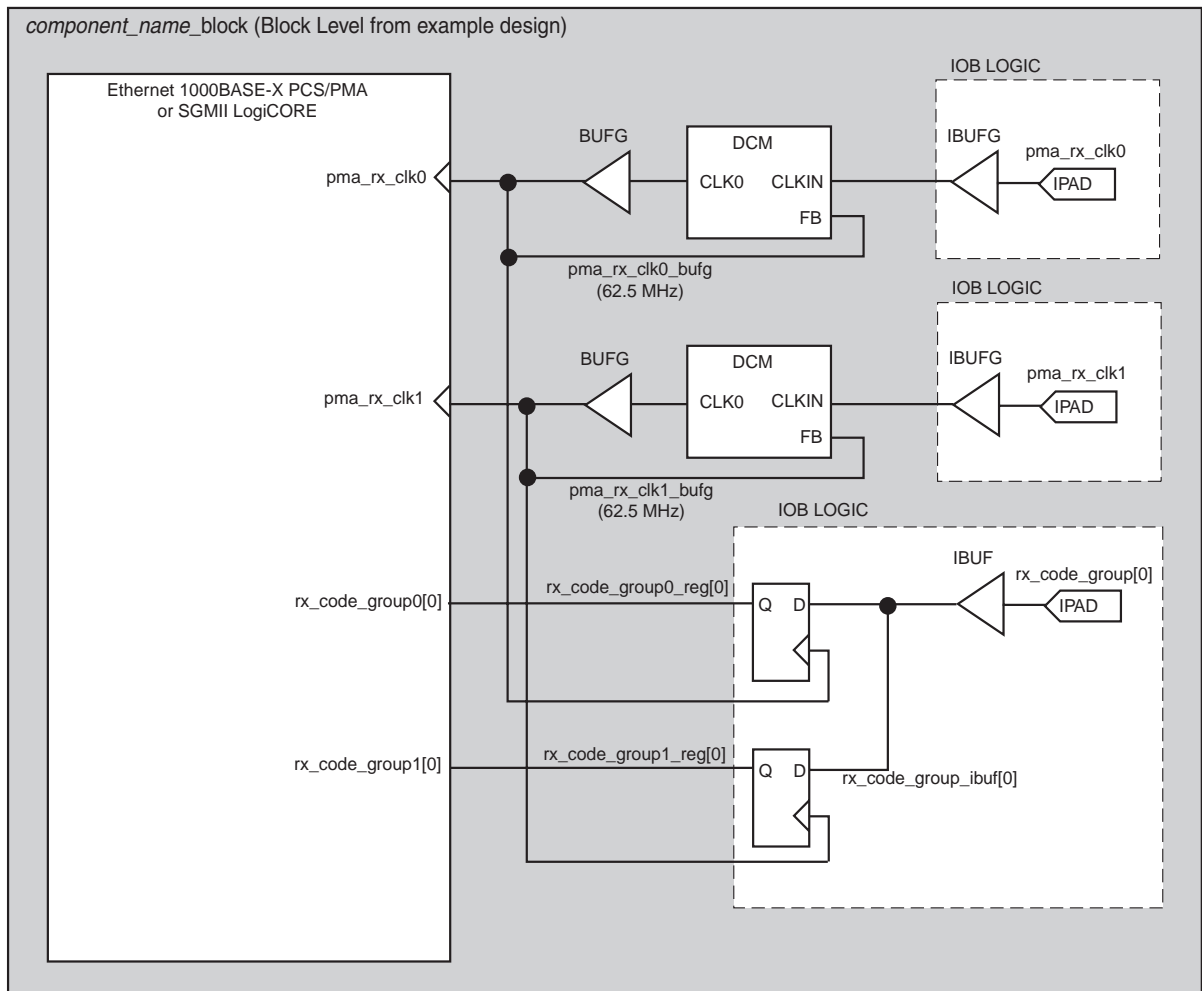


Figure 6-3: TBI Receiver Logic for Spartan-3, Spartan-3E, and Spartan-3A Devices

Virtex-4 Devices

Method 1

The Virtex-4 FPGA logic used by the example design delivered with the core is illustrated in Figure 6-4. This shows a Virtex-4 device IDDR primitive used with the `DDR_CLK_EDGE` attribute set to `SAME_EDGE` (see the *Virtex-4 FPGA User Guide*). This uses local inversion of `pma_rx_clk0` within the IOB logic to receive the `rx_code_group[9:0]` data bus on both the rising and falling edges of `pma_rx_clk0`. The `SAME_EDGE` attribute causes the IDDR to output both Q1 and Q2 data on the rising edge of `pma_rx_clk0`.

For this reason, `pma_rx_clk0` can be routed to both `pma_rx_clk0` and `pma_rx_clk1` clock inputs of the core as illustrated.

Caution! This logic relies on `pma_rx_clk0` and `pma_rx_clk1` being exactly 180 degrees out of phase with each other since the falling edge of `pma_rx_clk0` is used in place of `pma_rx_clk1`. See the data sheet for the attached SERDES to verify that this is the case.

The IDELAY elements can be adjusted to fine-tune the setup and hold times at the TBI IOB input flip-flops. The delay is applied to the IDELAY elements using constraints in the UCF; these can be edited if desired. See “Ten-Bit Interface Constraints” in Chapter 12 for more information.

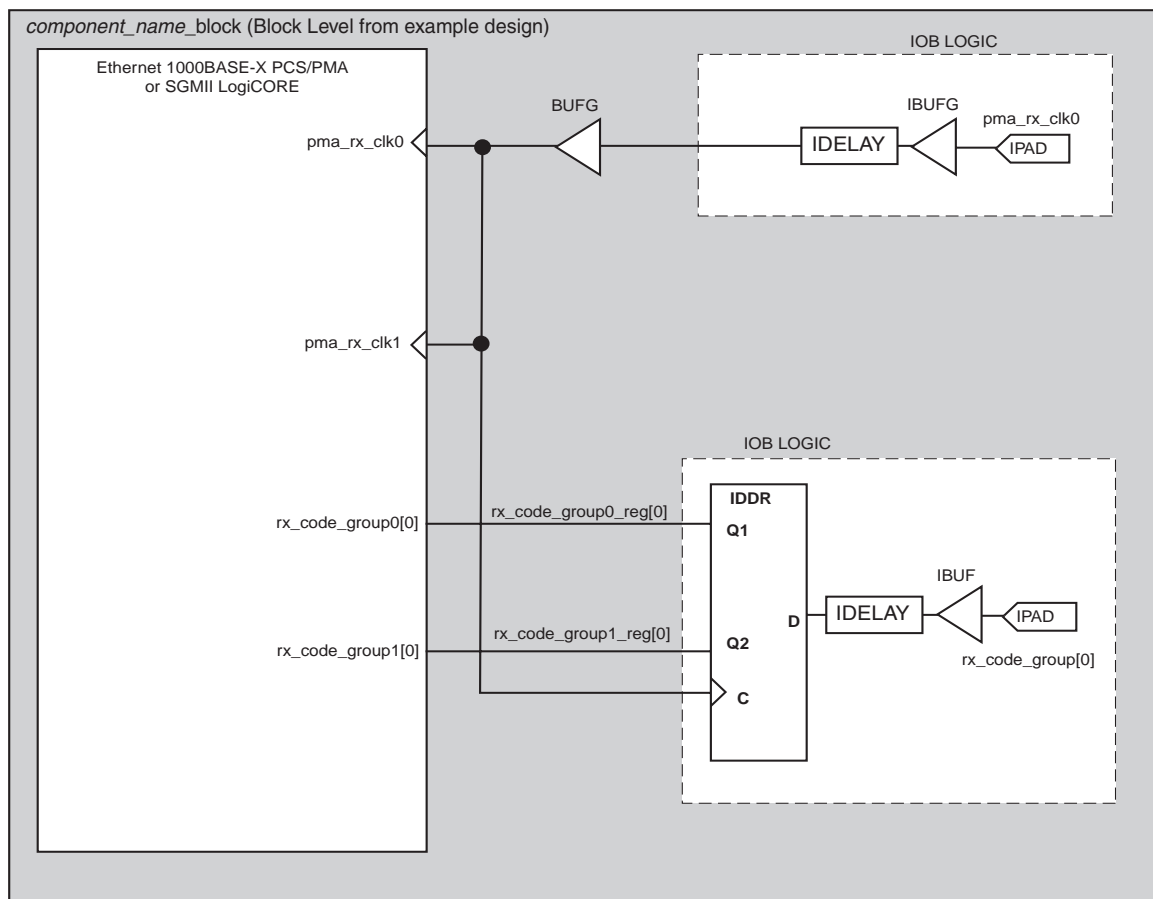


Figure 6-4: Ten-Bit Interface Receiver Logic - Virtex-4 Device (Example Design)

Method 2

This logic from method 1 relies on pma_rx_clk0 and pma_rx_clk1 being exactly 180 degrees out of phase with each other since the falling edge of pma_rx_clk0 is used in place of pma_rx_clk1. See the data sheet for the attached SERDES to verify that this is the case. If not, then the logic of Figure 6-5 illustrates an alternative implementation where both pma_rx_clk0 and pma_rx_clk1 are used as intended. Each bit of rx_code_group[9:0] must be routed to two separate device pads.

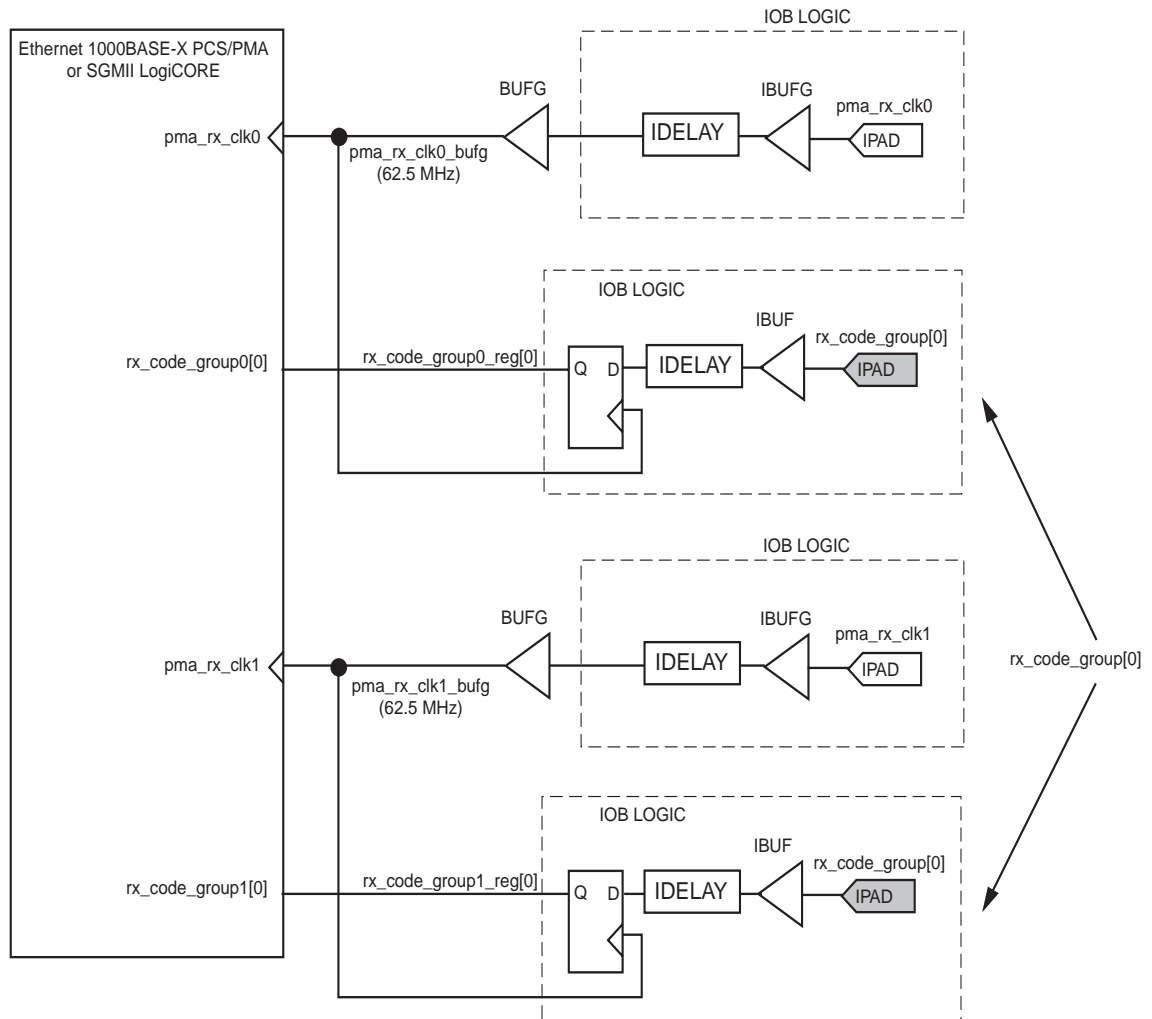


Figure 6-5: Alternate Ten-Bit Interface Receiver Logic for Virtex-4 Devices

Virtex-5 Devices

Method 1

The Virtex-5 FPGA logic used by the example design delivered with the core is illustrated in [Figure 6-6](#). This shows a Virtex-5 device IDDR primitive used with the `DDR_CLK_EDGE` attribute set to `SAME_EDGE` (see the *Virtex-5 FPGA User Guide*). This uses local inversion of `pma_rx_clk0` within the IOB logic to receive the `rx_code_group[9:0]` data bus on both the rising and falling edges of `pma_rx_clk0`. The `SAME_EDGE` attribute causes the IDDR to output both Q1 and Q2 data on the rising edge of `pma_rx_clk0`.

For this reason, `pma_rx_clk0` can be routed to both `pma_rx_clk0` and `pma_rx_clk1` clock inputs of the core as illustrated.

Caution! This logic relies on `pma_rx_clk0` and `pma_rx_clk1` being exactly 180 degrees out of phase with each other because the falling edge of `pma_rx_clk0` is used in place of `pma_rx_clk1`. See the data sheet for the attached SERDES to verify that this is the case.

The IODELAY elements can be adjusted to fine-tune the setup and hold times at the TBI IOB input flip-flops. The delay is applied to the IODELAY element using constraints in the UCF; these can be edited if desired. See [“Ten-Bit Interface Constraints”](#) in [Chapter 12](#) for more information.

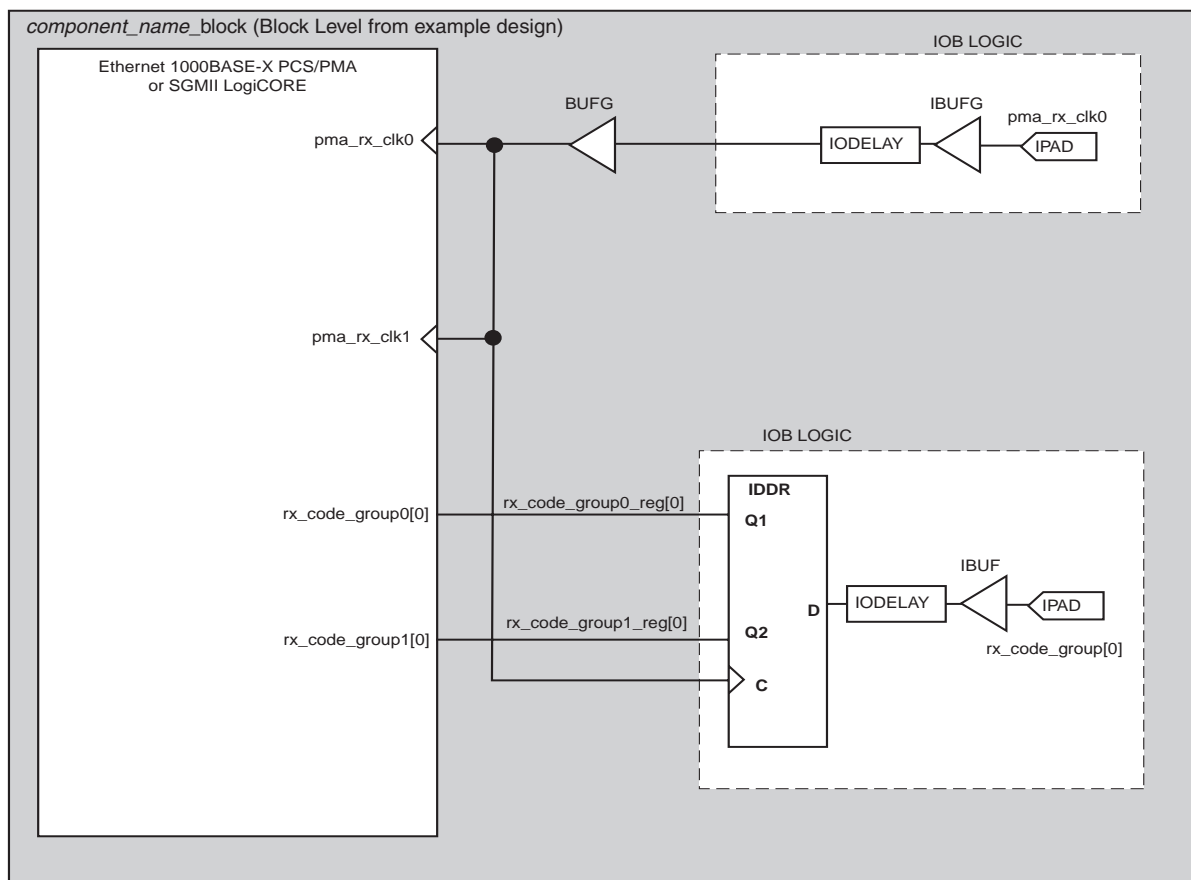


Figure 6-6: Ten-Bit Interface Receiver Logic - Virtex-5 Device (Example Design)

Method 2

This logic from method 1 relies on `pma_rx_clk0` and `pma_rx_clk1` being exactly 180 degrees out of phase with each other because the falling edge of `pma_rx_clk0` is used in place of `pma_rx_clk1`. See the data sheet for the attached SERDES to verify that this is the case. If not, the logic of [Figure 6-7](#) illustrates an alternate implementation where both `pma_rx_clk0` and `pma_rx_clk1` are used as intended. Each bit of `rx_code_group[9:0]` must be routed to two separate device pads. The `IODELAY` elements shown on [Figure 6-7](#) can be used to compensate for any bus skew that has resulted.

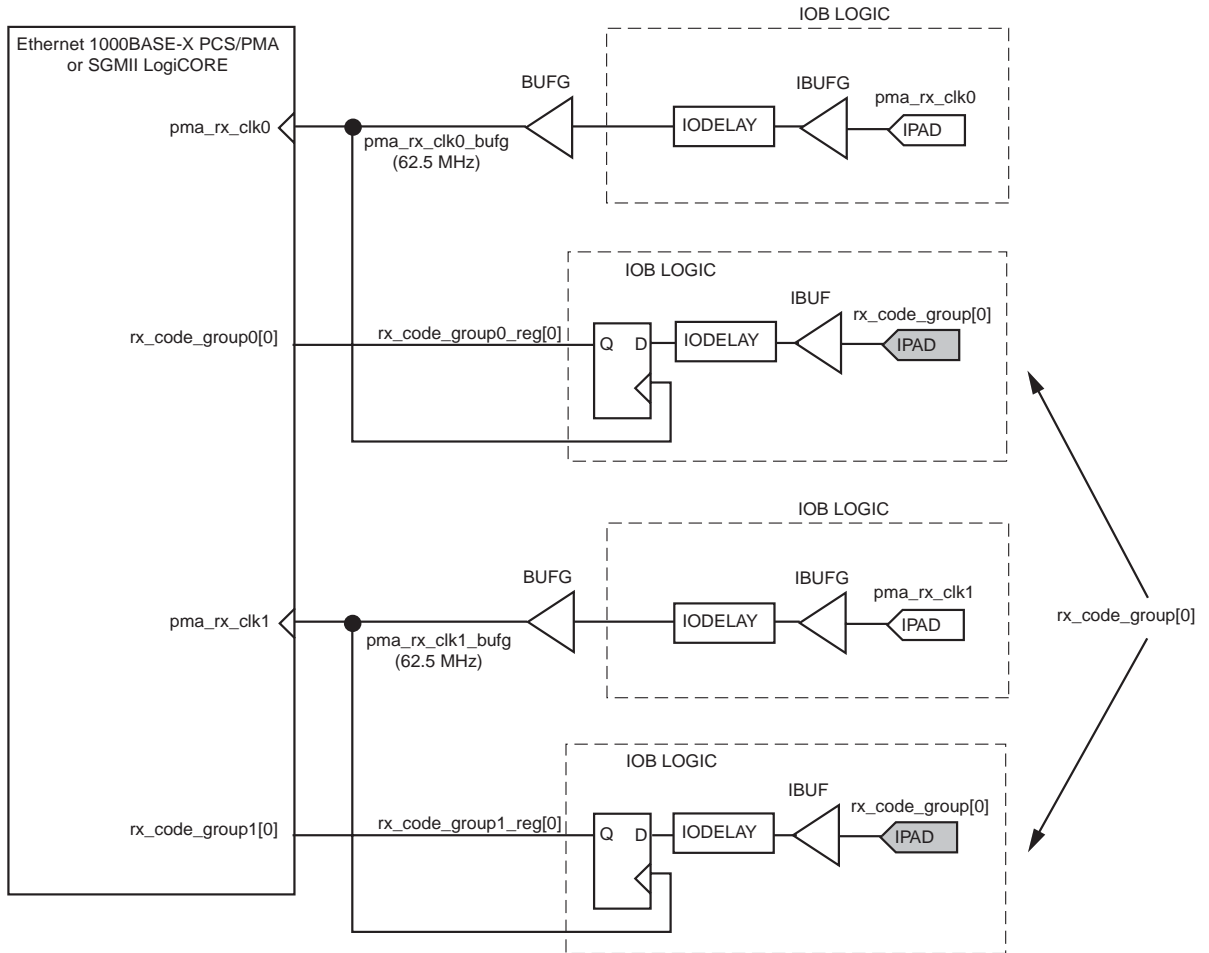


Figure 6-7: Alternate Ten-Bit Interface Receiver Logic - Virtex-5 Devices

Clock Sharing across Multiple Cores with TBI

Figure 6-8 illustrates sharing clock resources across multiple instantiations of the core when using the TBI. `gtx_clk` may be shared between multiple cores, resulting in a common clock domain across the device.

The receiver clocks `pma_rx_clk0` and `pma_rx_clk1` cannot be shared. Each core will be provided with its own versions of these clocks from its externally connected SERDES. Figure 6-8 illustrates the receiver clock logic used for the Virtex-II family. See “Receiver Logic,” page 70, for a description of the clock logic for other device families.

Figure 6-8 illustrates only two cores. However, more can be added using the same principle. This is done by instantiating the cores using the block level (from the example design) and sharing `gtx_clk` across all instantiations.

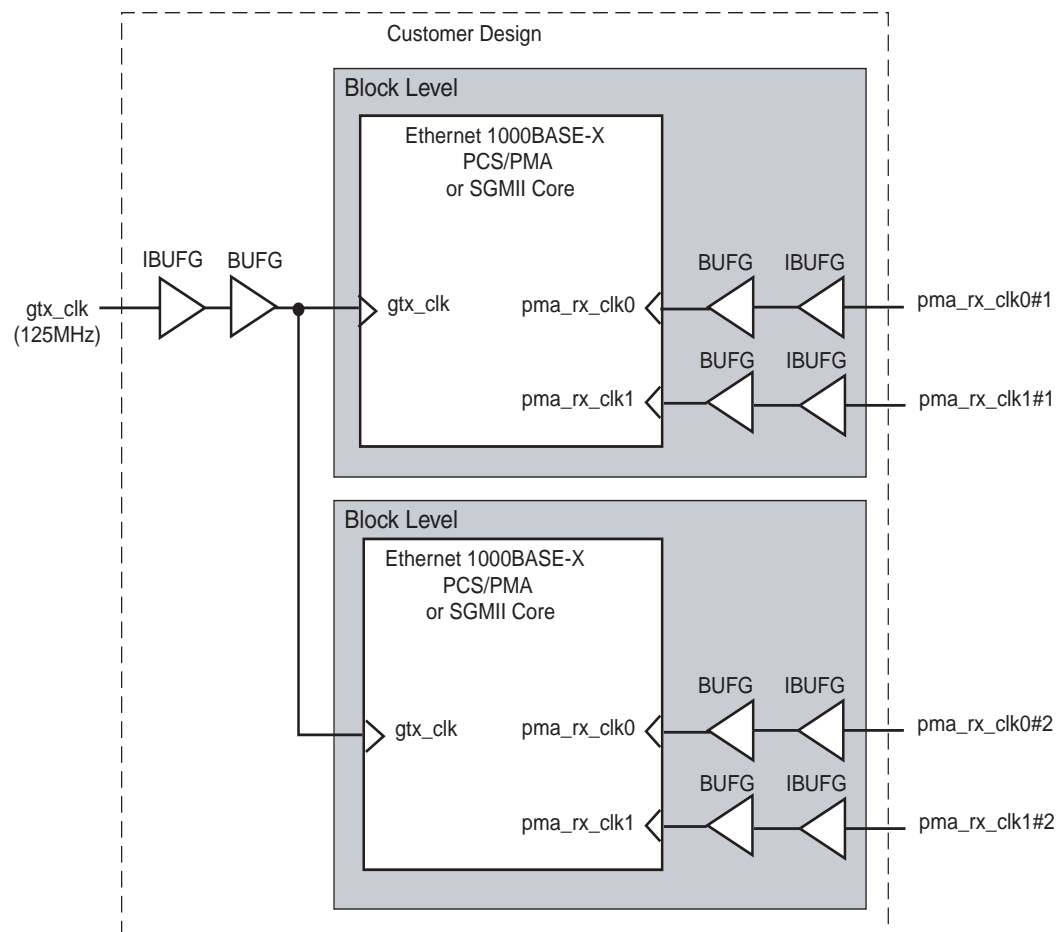


Figure 6-8: Clock Management, Multiple Core Instances with Ten-Bit Interface

1000BASE-X with RocketIO Transceivers

This chapter provides general guidelines for creating 1000BASE-X designs that use RocketIO transceivers for Virtex-II Pro, Virtex-4, and Virtex-5 devices. Information about RocketIO transceiver and core logic in all supported device families is provided, as well as information about designs requiring multiple instantiations of the core. Note that clock sharing should occur whenever possible to save device resources.

RocketIO Transceiver Logic

The example is split between two discrete hierarchical layers, as illustrated in [Figure 4-1](#). The block level is designed so that it can be instantiated directly into customer designs and provides the following functionality:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to a Virtex-II Pro, Virtex-4, or Virtex-5 RocketIO transceiver

The logic implemented in the block level is illustrated in all the figures in this chapter.

Virtex-II Pro Devices

The core is designed for seamless integration with the Virtex-II Pro RocketIO Multi-Gigabit Transceiver (MGT). [Figure 7-1](#) illustrates the connections and logic required between the core and the MGT—the signal names and logic in the figure precisely match those delivered with the example design when an MGT is used.

Some modifications can be made to the MGT. For example, REFCLK may be used instead of BREFCLK. See the *RocketIO Transceiver User Guide* (UG024) for details.

The placement of the flip-flop that connects to ENMCOMMAALIGN and ENPCOMMAALIGN is important (see [Figure 7-1](#)). For detailed information, see “[Virtex-II Pro RocketIO MGTs for 1000BASE-X Constraints](#),” and the *RocketIO Transceiver User Guide*.

Note: The `brefclk` differential pair applied to the MGT is of frequency 62.5 MHz.

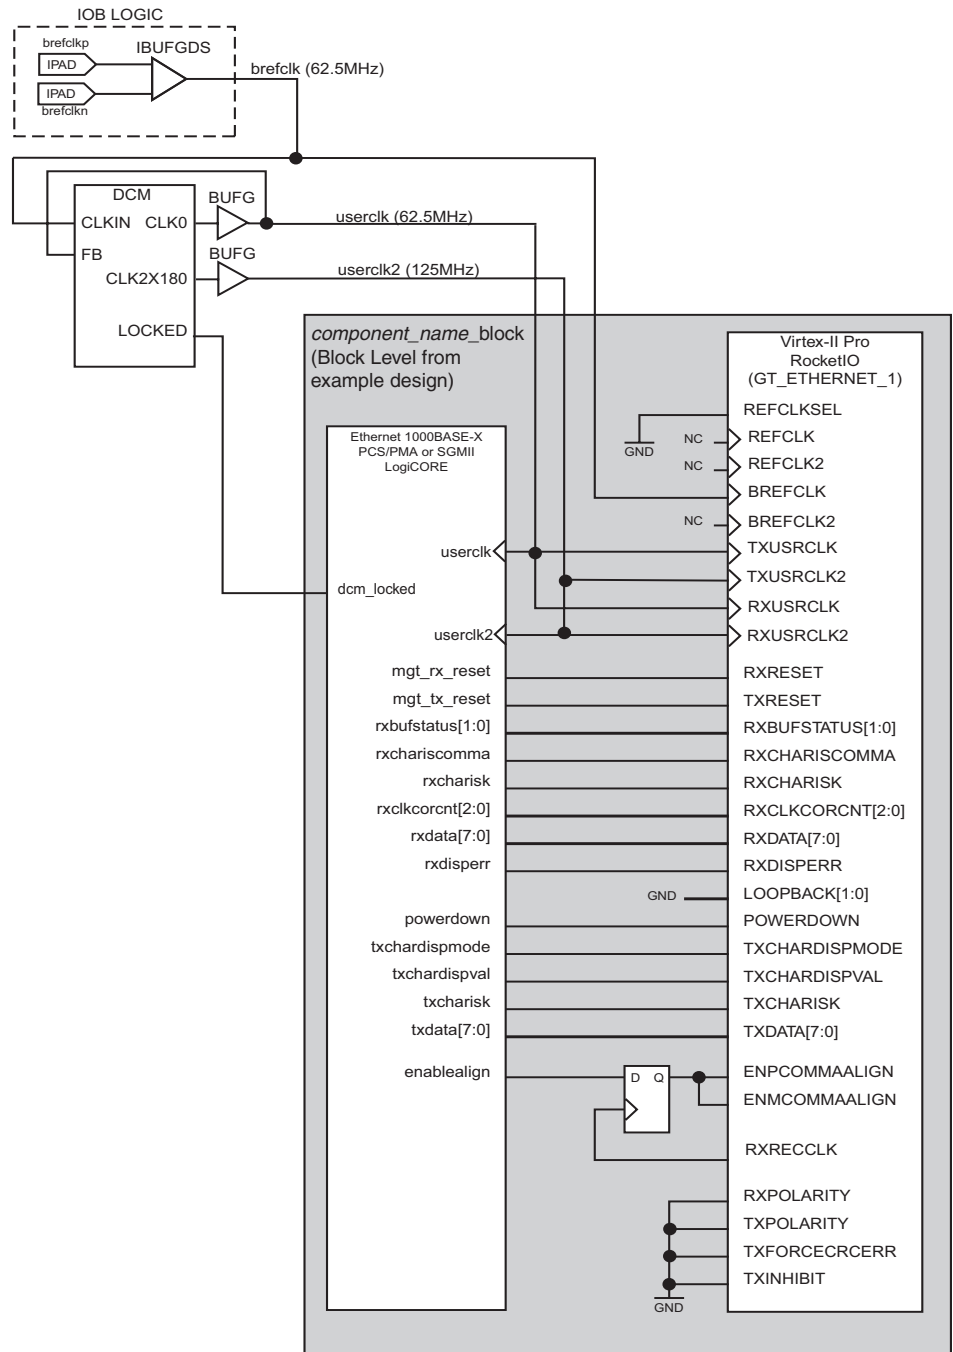


Figure 7-1: 100BASE-X Connection to a Virtex-II Pro MGT

Virtex-4 FX Devices

The core is designed to integrate with the Virtex-4 RocketIO MGT. [Figure 7-2](#) illustrates the connections and logic required between the core and MGT—the signal names and logic in the figure precisely match those delivered with the example design when an MGT is used.

Note: A small logic shim (included in the *block-level wrapper*) is required to convert between the port differences between the Virtex-II Pro and Virtex-4 RocketIO transceivers.

The MGT clock distribution in Virtex-4 devices is column-based and consists of multiple MGT tiles (each tile contains two MGTs). For this reason, the MGT wrapper delivered with the core always contains two MGT instantiations, even if only a single MGT is in use. [Figure 7-2](#) illustrates a single MGT tile for clarity.

A `GT11CLK_MGT` primitive is also instantiated to derive the reference clocks required by the MGT column-based tiles. See the *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide* (UG076) for information about layout and clock distribution.

The 250 MHz reference clock from the `GT11CLK_MGT` primitive is routed to the MGT, configured to internally synthesize a 125 MHz clock. This is output on the `TXOUTCLK1` port of the MGT and after placed onto global clock routing, can be used by all core logic. This clock is input back into the MGT on the user interface clock ports `rxusrclk2` and `txusrclk2`. With the attribute settings applied to the MGT from the example design, the `txusrclk` and `rxusrclk` ports are derived internally within the MGT using the internal clock dividers and do not need to be provided from the FPGA fabric.

The Virtex-4 FX MGTs require the inclusion of a calibration block in the fabric logic; the example design provided with the core instantiates calibration blocks as required. Calibration blocks require a clock source of between 25 to 50 MHz that is shared with the Dynamic Reconfiguration Port (DRP) of the MGT, which is named `dclk` in the example design. See Xilinx [Answer Record 22477](#) for more information.

[Figure 7-2](#) also illustrates the `TX_SIGNAL_DETECT` and `RX_SIGNAL_DETECT` ports of the calibration block, which should be driven to indicate whether or not dynamic data is being transmitted and received through the MGT (see [Virtex-4 Errata](#)). However, `RX_SIGNAL_DETECT` is connected to the `signal_detect` port of the example design. `signal_detect` is intended to be connected to the optical transceiver to indicate the presence of light. When light is detected, the optical transceiver provides dynamic data to the Rx ports of the MGT. When no light is detected, the calibration block switches the MGT into loopback to force dynamic data through the MGT receiver path.

Caution! `signal_detect` is an optional port in the *IEEE 802.3* specification. If this is not used, the `RX_SIGNAL_DETECT` port of the calibration block must be driven by an alternative method. Please refer to *XAPP732* for more information.




Figure 7-2: 1000BASE-X Connection to Virtex-4 MGT

Virtex-5 LXT and SXT Devices

The core is designed to integrate with the Virtex-5 RocketIO GTP transceiver. [Figure 7-3](#) illustrates the connections and logic required between the core and the GTP transceiver—the signal names and logic in the figure precisely match those delivered with the example design when a GTP transceiver is used.

Note: A small logic shim (included in the *block-level wrapper*) is required to convert between the port differences between the Virtex-II Pro and Virtex-5 GTP transceiver.

A GTP tile consists of a pair of transceivers. For this reason, the GTP transceiver wrapper delivered with the core always contains two GTP instantiations, even if only a single GTP transceiver tile is in use. [Figure 7-3](#) illustrates a single GTP transceiver tile.

The 125 MHz differential reference clock is routed directly to the GTP transceiver. The GTP transceiver is configured to output a version of this clock on the REFCLKOUT port and after placement onto global clock routing, can be used by all core logic. This clock is input back into the GTP transceiver on the user interface clock ports rxusrclk, rxusrclk2, txusrclk, and txusrclk2.

See also “[Virtex-5 RocketIO GTP Transceivers for 1000BASE-X Constraints](#),” page 166.

Virtex-5 RocketIO GTP Wizard

The two wrapper files immediately around the GTP transceiver pair, `rocketio_wrapper_gtp_tile` and `rocketio_wrapper_gtp` (see [Figure 7-3](#)), are generated from the *RocketIO GTP Wizard*. These files apply all the gigabit Ethernet attributes. Consequently, these files can be regenerated by customers and therefore be easily targeted at ES or Production silicon. Note that this core targets production silicon.

The CORE Generator log file (XCO file) which was created when the *RocketIO GTP Wizard* project was generated is available in the following location:

```
<project_directory>/<component_name>/example_design/transceiver/  
rocketio_wrapper_gtp.xco
```

This file can be used as an input to the CORE Generator to regenerate the RocketIO wrapper files. The XCO file itself contains a list of all of the GTP Wizard attributes which were used. For further information, please refer to the *Virtex-5 RocketIO GTP Wizard Getting Started Guide* (UG188) and the *CORE Generator Guide*, at www.xilinx.com/support/software_manuals.htm

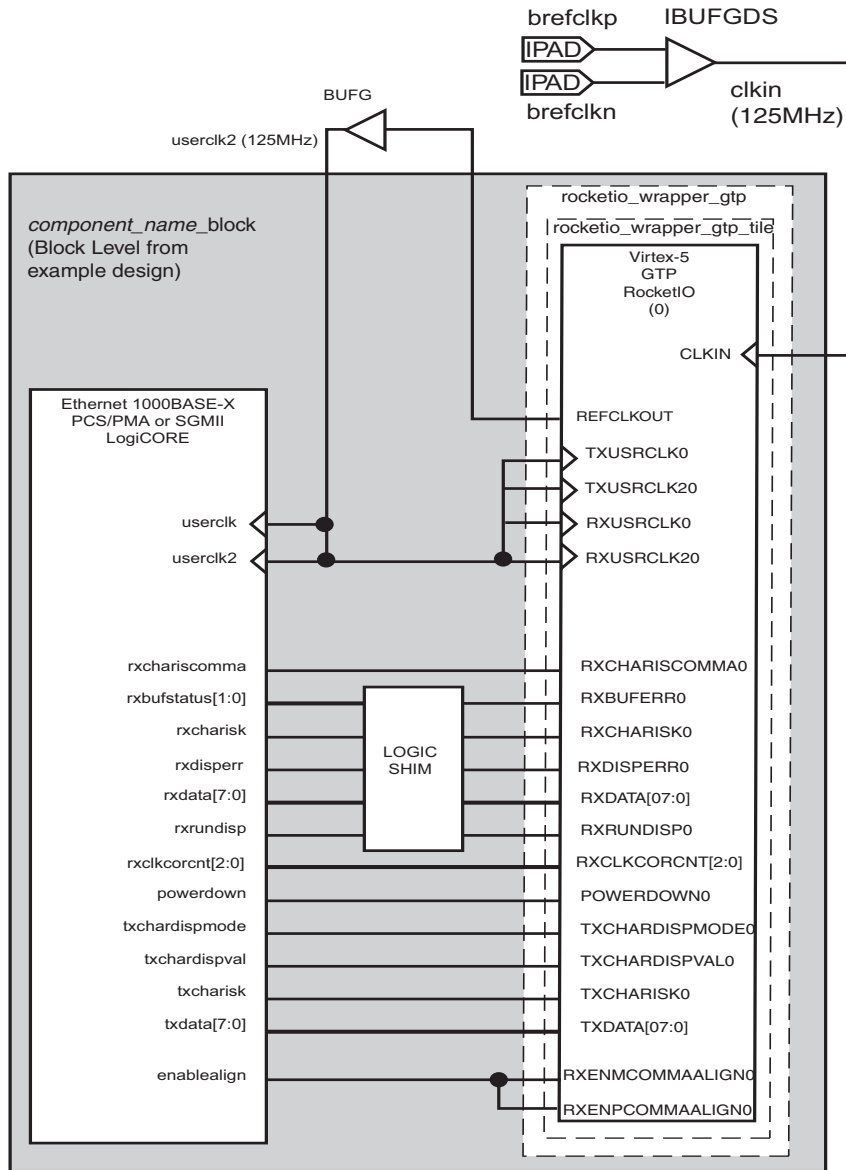


Figure 7-3: 100BASE-X Connection to Virtex-5 GTP Transceivers

Virtex-5 FXT Devices

The core is designed to integrate with the Virtex-5 RocketIO GTX transceiver. [Figure 7-4](#) illustrates the connections and logic required between the core and the GTX transceiver—the signal names and logic in the figure precisely match those delivered with the example design when a GTX transceiver is used.

Note: A small logic shim (included in the *block*-level wrapper) is required to convert between the port differences between the Virtex-II Pro and Virtex-5 GTX transceiver.

A GTX tile consists of a pair of transceivers. For this reason, the GTX transceiver wrapper delivered with the core always contains two GTX instantiations, even if only a single GTX transceiver tile is in use. [Figure 7-4](#) illustrates a single GTX transceiver tile.

The 125 MHz differential reference clock is routed directly to the GTX transceiver. The GTX transceiver is configured to output a version of this clock on the REFCLKOUT port: this is then routed to a DCM.

From the DCM, the CLK0 port (125MHz) is placed onto global clock routing and can be used as the 125MHz clock source for all core logic: this clock is also input back into the GTX transceiver on the user interface clock ports rxusrclk2 and txusrclk2.

From the DCM, the CLKDV port (62.5MHz) is placed onto global clock routing and is input back into the GTX transceiver on the user interface clock ports rxusrclk and txusrclk.

See also “[Virtex-5 RocketIO GTX Transceivers for 1000BASE-X Constraints](#),” page 167.

Virtex-5 RocketIO GTX Wizard

The two wrapper files immediately around the GTX transceiver pair, `rocketio_wrapper_gtx_tile` and `rocketio_wrapper_gtx` (see [Figure 7-4](#)), are generated from the *RocketIO GTX Wizard*. These files apply all the gigabit Ethernet attributes. Consequently, these files can be regenerated by customers and therefore be easily targeted at ES or Production silicon. Note that this core targets production silicon.

The CORE Generator log file (XCO file) which was created when the *RocketIO GTX Wizard* project was generated is available in the following location:

```
<project_directory>/<component_name>/example_design/transceiver/  
rocketio_wrapper_gtx.xco
```

This file can be used as an input to the CORE Generator to regenerate the RocketIO wrapper files. The XCO file itself contains a list of all of the GTX Wizard attributes which were used. For further information, please refer to the *Virtex-5 RocketIO GTX Wizard Getting Started Guide* and the *CORE Generator Guide*, at www.xilinx.com/support/software_manuals.htm




Figure 7-4: 100BASE-X Connection to Virtex-5 GTX Transceivers

Clock Sharing Across Multiple Cores with RocketIO

Virtex-II Pro Devices

Figure 7-5 illustrates sharing clock resources across two instantiations of the core on the same half of the device when using the core with the Virtex-II Pro MGT. Note that more can be added by instantiating the cores using the block level (from the example design) and continuing to share `userclk`, `userclk2`, and `brefclk` across all instantiations. For each core, `userclk` and `userclk2` must always be derived from the `brefclk` or `refclk` used by that core.

When using the fixed routing resources of `brefclk`, MGTs along the top edge of the device must use a separate `brefclk` routing resource to those along the bottom edge of the device. For more information, see the *Virtex-II Pro RocketIO Transceiver User Guide* (UG024). Each `brefclk` domain must use its own DCM to derive its version of `userclk` and `userclk2`.

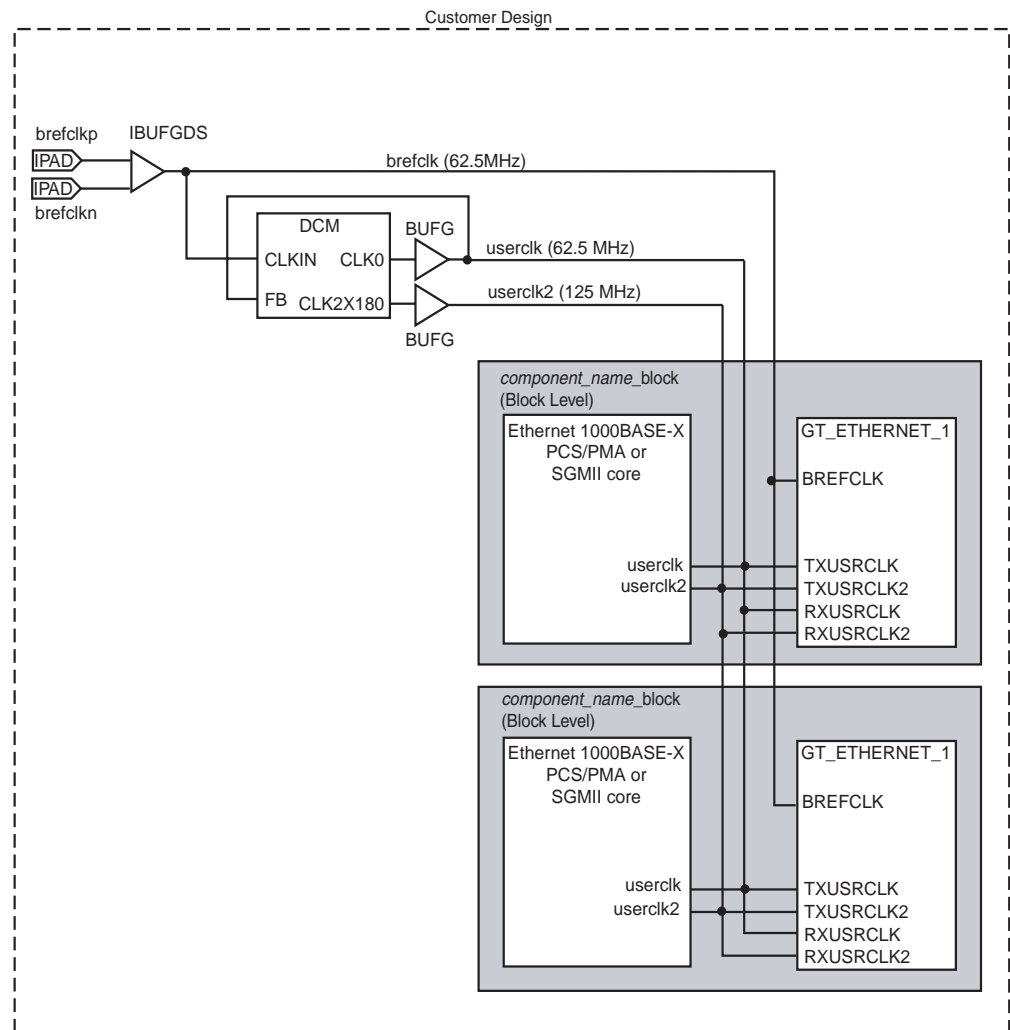


Figure 7-5: Clock Management: Two Core Instances, Virtex-II Pro MGTs for 100BASE-X

Virtex-4 FX Devices

Figure 7-6 illustrates sharing clock resources across multiple instantiations of the core when using MGTs. Note that the example design, when using the Virtex-4 family, can be generated to connect either a single instance of the core, or connect a pair of core instances to the transceiver pair present in an MGT tile. Figure 7-6 shows two instantiations of the block level, where each block contains a pair of cores, subsequently illustrating clock sharing between four cores in total.

More cores can be added by continuing to instantiate extra block-level modules. Share clocks only between the MGTs in a single column. For each column, use a single `brefclk_p` and `brefclk_n` differential clock pair and connect this to a `GT11CLK_MGT` primitive. The clock output from this primitive should be shared across all used RocketIO tiles in the column. See the *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide* (UG076) for more information.

To provide the 125 MHz clock for all core instances, select a `TXOUTCLK1` port from any MGT. This can be routed onto global clock routing using a BUFG as illustrated, and shared between all cores and MGTs in the column. Although not illustrated in Figure 7-6, `dc1k` (the clock used for the calibration blocks and for the Dynamic Reconfiguration Port (DRP) of the MGTs) can also be shared.

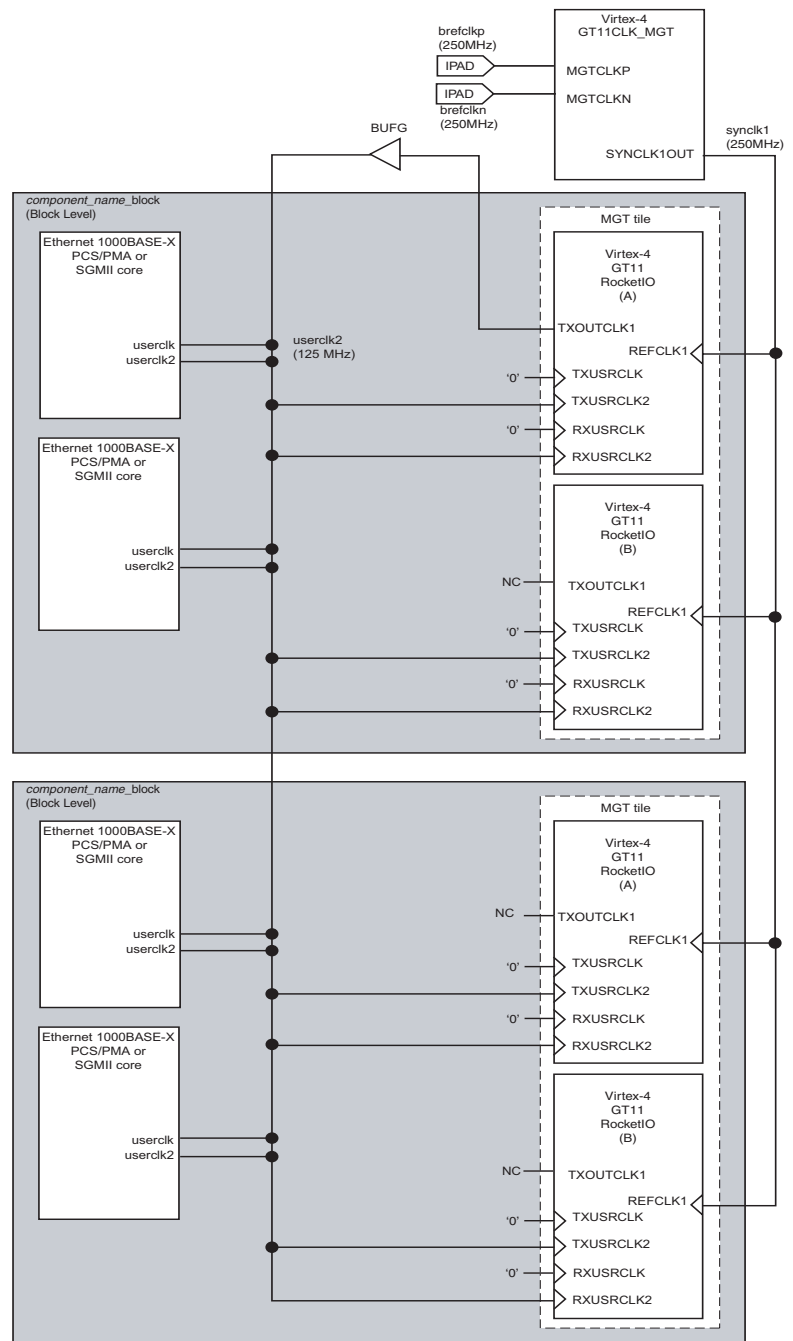


Figure 7-6: Clock Management - Multiple Core Instances, MGTs for 1000BASE-X

Virtex-5 LXT and SXT Devices

Figure 7-7 illustrates sharing clock resources across multiple instantiations of the core when using Virtex-5 RocketIO GTP transceivers.

The example design can be generated to connect either a single instance of the core or connect a pair of core instances to the transceiver pair present in a GTP tile. Figure 7-7 illustrates two instantiations of the block level, and each block level contains a pair of cores, consequently illustrating clock sharing between a total of four cores.

Additional cores can be added by continuing to instantiate extra block level modules. Share the `brefclk_p` and `brefclk_n` differential clock pair. See the *Virtex-5 RocketIO GTP Transceiver User Guide* (UG196) for more information.

To provide the 125 MHz clock for all core instances, select a REFCLKOUT port from any GTP transceiver. This can be routed onto global clock routing using a BUFG as illustrated and shared between all cores and GTP transceivers.

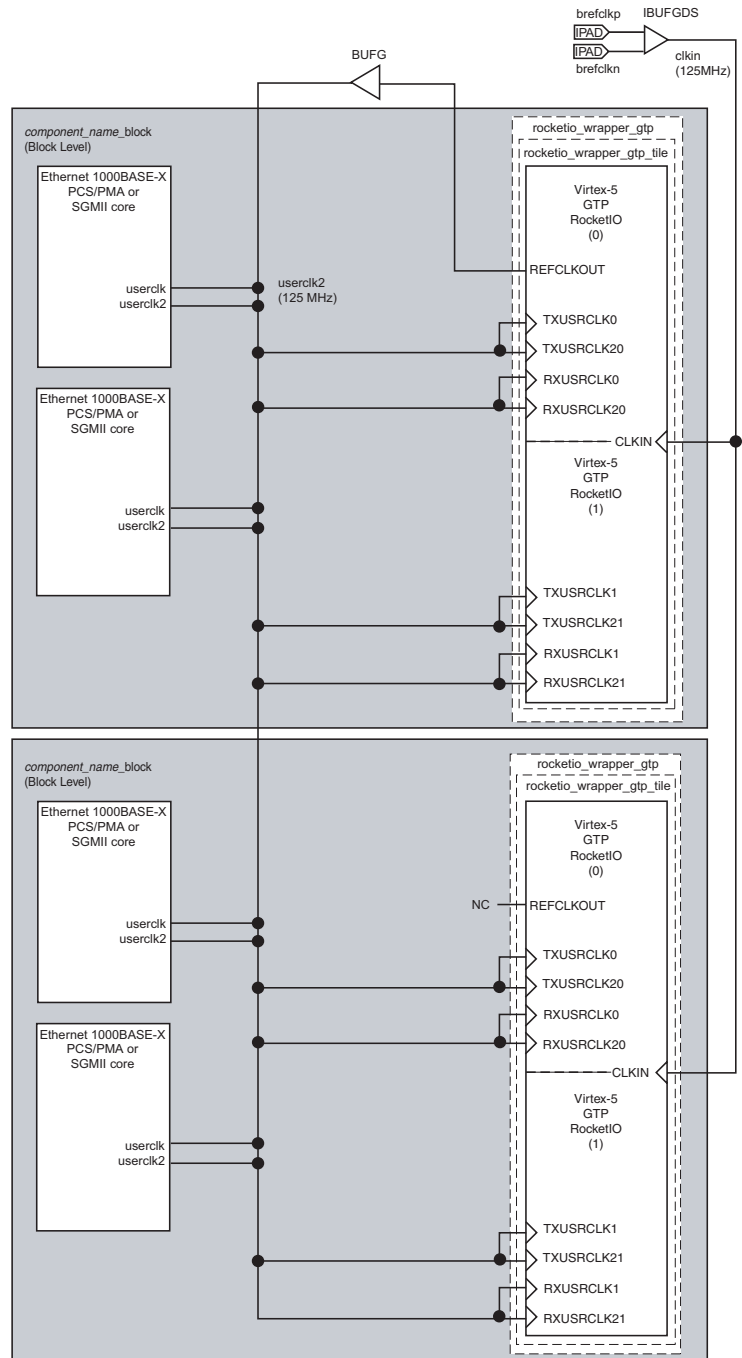


Figure 7-7: Clock Management - Multiple Core Instances, Virtex-5 RocketIO GTP Transceivers for 1000BASE-X

Virtex-5 FXT Devices

Figure 7-8 illustrates sharing clock resources across multiple instantiations of the core when using Virtex-5 RocketIO GTX transceivers.

The example design can be generated to connect either a single instance of the core or connect a pair of core instances to the transceiver pair present in a GTX tile. Figure 7-8 illustrates two instantiations of the block level, and each block level contains a pair of cores, consequently illustrating clock sharing between a total of four cores.

Additional cores can be added by continuing to instantiate extra block level modules. Share the `brefclk_p` and `brefclk_n` differential clock pair. See the *Virtex-5 RocketIO GTX Transceiver User Guide* for more information.

To provide the FPGA fabric clocks for all core instances, select a `REFCLKOUT` port from any GTX transceiver and route this to a single DCM. The `CLK0` (125MHz) and `CLKDV` (62.5MHz) outputs from this DCM, placed onto global clock routing using a `BUFGs`, can be shared across all core instances and GTX transceivers as illustrated.

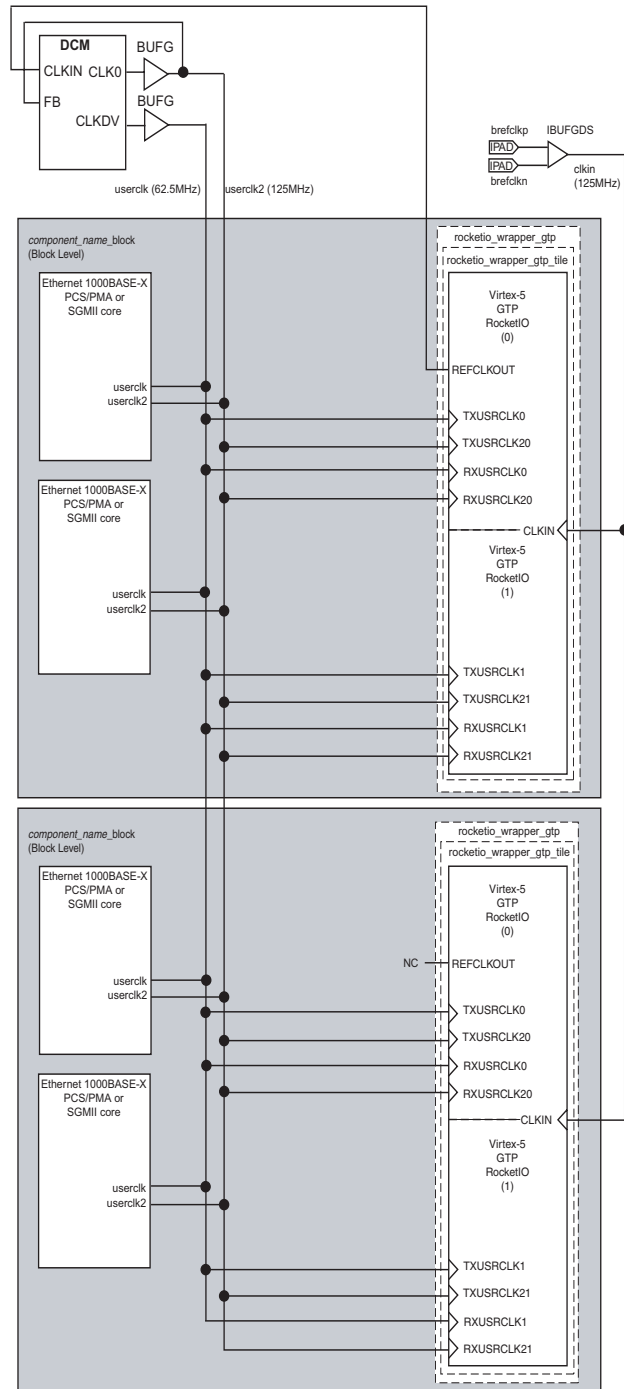


Figure 7-8: Clock Management - Multiple Core Instances, Virtex-5 RocketIO GTX Transceivers for 1000BASE-X

SGMII / Dynamic Standards Switching with RocketIO Transceivers

This chapter provides general guidelines for creating SGMII designs, and designs capable of switching between 1000BASE-X and SGMII standards (Dynamic Standards Switching), using a RocketIO transceiver. Throughout this chapter, any reference to SGMII also applies to the Dynamic Standards Switching implementation.

The chapter begins with an explanation of the two Receiver Elastic Buffer implementations: one implementation uses the buffer present in the RocketIO transceivers, and the other uses a larger buffer, implemented in the FPGA fabric.

After selecting the Rx Elastic Buffer implementation type, an explanation of the RocketIO transceiver and core logic in all supported device families is provided in the following sections:

- [“RocketIO Logic using the RocketIO Rx Elastic Buffer,” page 98](#)
- [“RocketIO Logic with the Fabric Rx Elastic Buffer,” page 98](#)

Instances where multiple instantiations of the core are required when using the fabric Receiver Elastic Buffer are then presented. Clock sharing should occur whenever possible to save device resources.

Receiver Elastic Buffer Implementations

Selecting the Buffer Implementation from the GUI

The GUI provides two SGMII Capability options:

- 10/100/1000 Mbps (clock tolerance compliant with Ethernet specification)
- 10/100/1000 Mbps (restricted tolerance for clocks) OR 100/1000 Mbps

The first option, 10/100/1000 Mbps (clock tolerance compliant with Ethernet specification) is the default and provides the implementation using the Receiver Elastic Buffer in FPGA fabric. This alternative Receiver Elastic Buffer uses a single block RAM to create a buffer twice as large as the one present in the RocketIO transceiver, for this reason consuming extra logic resources. However, this default mode is reliable for all implementations using standard Ethernet frame sizes. Further consideration must be made for jumbo frames.

The second option, 10/100/1000 Mbps (restricted tolerance for clocks) or 100/1000 Mbps, uses the receiver elastic buffer present in the RocketIOs. This is half the size and can potentially underflow or overflow during SGMII frame reception at 10 Mbps operation

(see the next section). However, there are logical implementations where this can be reliable and has the benefit of lower logic utilization.

The Requirement for the FPGA Fabric Rx Elastic Buffer

Figure 8-1 illustrates a simplified diagram of a common situation where the core, in SGMII mode, is interfaced to an external PHY device. Separate oscillator sources are used for the FPGA and the external PHY. The Ethernet specification uses clock sources with a tolerance of 100ppm. In Figure 8-1, the clock source for the PHY is slightly faster than the clock source to the FPGA. For this reason, during frame reception, the receiver elastic buffer (shown here as implemented in the RocketIO) starts to fill.

Following frame reception, in the interframe gap period, idles are removed from the received data stream to return the Rx Elastic Buffer to half-full occupancy. This is performed by the clock correction circuitry (see the RocketIO User Guide for the targeted device).

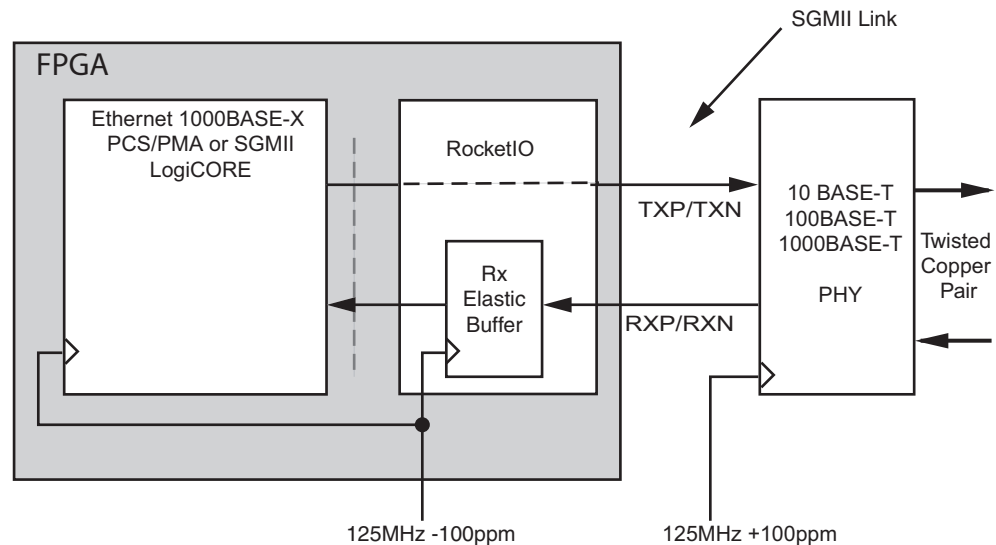


Figure 8-1: SGMII Implementation using Separate Clock Sources

Analysis

Assuming separate clock sources, each of tolerance 100 ppm, the maximum frequency difference between the two devices can be 200 ppm. It can be shown that this translates into a full clock period difference every 5000 clock periods.

Relating this to an Ethernet frame, there will be a single byte of difference every 5000 bytes of received frame data, and this will cause the Rx Elastic Buffer to either fill or empty by an occupancy of one.

The maximum Ethernet frame size (non-jumbo) is 1522 bytes for a VLAN frame.

- At 1 Gbps operation, this translates into 1522 clock cycles.
- At 100 Mbps operation, this translates into 15220 clock cycles (as each byte is repeated 10 times).
- At 10 Mbps operation, this translates into 152200 clock cycles (as each byte is repeated 100 times).

Considering the 10 Mbps case, we would need $152200/5000 = 31$ FIFO entries in the Elastic Buffer above and below the half way point to guarantee that the buffer will not under or overflow during frame reception. This assumes that frame reception begins when the buffer is exactly half full.

The size of the Rx Elastic Buffer in the RocketIOs is 64 entries. However, we cannot assume that the buffer is exactly half full at the start of frame reception. Additionally, the underflow and overflow thresholds are not exact (see [Appendix E, "Rx Elastic Buffer Specifications"](#) for more information).

To guarantee reliable SGMII operation at 10 Mbps (non-jumbo frames), the RocketIO Elastic Buffer must be bypassed and a larger buffer implemented in the FPGA fabric. The fabric buffer, provided by the example design, is twice the size of the RocketIO alternative. This has been proven to cope with standard (none jumbo) Ethernet frames at all three SGMII speeds.

[Appendix E, "Rx Elastic Buffer Specifications"](#) provides further information about all Rx Elastic Buffers used by the core. Information about the reception of jumbo frames is also provided.

The RocketIO Rx Elastic Buffer

The Elastic Buffer in the RocketIO can be used reliably when the following conditions are met:

- 10 Mbps operation is not required (for example, when connecting the core to the Xilinx 1-Gigabit Ethernet MAC to provide only 1 Gbps operation). Both 1 Gbps and 100 Mbps operation can be guaranteed.
- When the clocks are closely related (see the following section).

If there is any doubt, select the FPGA fabric Rx Elastic Buffer Implementation.

Closely Related Clock Sources

Case 1

Figure 8-2 illustrates a simplified diagram of a common situation where the core, in SGMII mode, is interfaced to an external PHY device. A common oscillator source is used for both the FPGA and the external PHY.

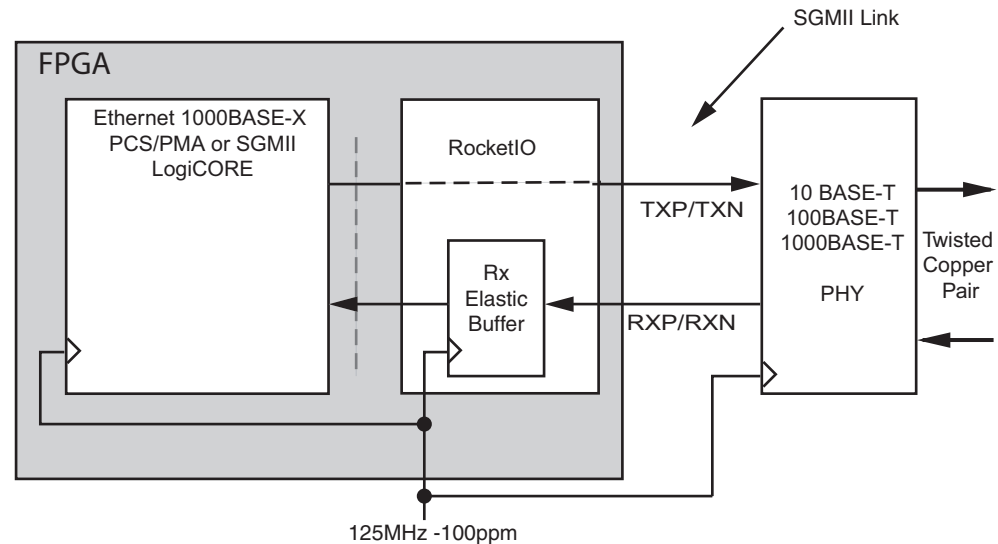


Figure 8-2: SGMII Implementation using Shared Clock Sources

If the PHY device sources the receiver SGMII stream synchronously from the shared oscillator (check PHY data sheet), the RocketIO will receive data at exactly the same rate as that used by the core. The receiver elastic buffer will neither empty nor fill, having the same frequency clock on either side.

In this situation, the receiver elastic buffer will not under or overflow, and the elastic buffer implementation in the RocketIO should be used to save logic resources.

Case 2

Consider again the case illustrated in Figure 8-1 with the following exception: assume that the clock sources used are both 50 ppm. Now the maximum frequency difference between the two devices is 100 ppm. It can be shown that this translates into a full clock period difference every 10000 clock periods, resulting in a requirement for 16 FIFO entries above and below the half-full point. This provides reliable operation with the RocketIO Rx Elastic Buffers. Again, however, check the PHY data sheet to ensure that the PHY device sources the receiver SGMII stream synchronously to its reference oscillator.

RocketIO Logic using the RocketIO Rx Elastic Buffer

When the RocketIO Rx Elastic Buffer implementation is selected, the connections between the core and the RocketIO as well as all clock circuitry in the system are identical to the 1000BASE-X implementation. For a detailed explanation, see Chapter 7, “1000BASE-X with RocketIO Transceivers.”

RocketIO Logic with the Fabric Rx Elastic Buffer

The example design delivered with the core is split between two hierarchical layers, as illustrated in [Figure 4-3](#). The block level is designed so to be instantiated directly into customer designs and provides the following functionality:

- Instantiates the core from HDL
- Connects the physical-side interface of the core to a Virtex-II Pro, Virtex-4 or Virtex-5 RocketIO transceiver via the fabric Rx Elastic Buffer

The logic implemented in the block level is illustrated in all figures throughout the remainder of this chapter.

Virtex-II Pro Devices

The core is designed for connection to a Virtex-II Pro MGT. The connections and logic required between the core and RocketIO transceiver are illustrated in [Figure 8-3](#)—the signal names and logic in the figure precisely match those delivered with the example design when an MGT transceiver is used.

Some modifications may be made to the MGT. For example, REFCLK may be used instead of BREFCLK. See the *Virtex-II Pro RocketIO Transceiver User Guide* (UG024) for details.

[Figure 8-3](#) shows that the Rx Elastic Buffer is implemented in the FPGA fabric between the MGT transceiver and the core. This replaces the Rx Elastic Buffer in the MGT (which is bypassed).

This alternative Receiver Elastic Buffer uses a single block RAM to create a buffer twice as large as the one present in the MGT. It is able to cope with larger frame sizes before clock tolerances accumulate and result in emptying or filling of the buffer. This is necessary to guarantee SGMII operation at 10 Mbps, where each frame size is effectively 100 times larger than the same frame would be at 1 Gbps because each byte is repeated 100 times (see “[Designing with Client-side GMII for the SGMII Standard](#),” page 59).

In bypassing the MGT Rx Elastic Buffer, data is clocked out of the MGT synchronously to `rxrecclk`. This must be placed on constrained local clock routing for reliable operation. See “[Virtex-II Pro RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints](#),” page 163 for constraint details. This methodology is also described in *XAPP763*.

Note: The `brefclk` differential pair applied to the MGT is of frequency 62.5 MHz.

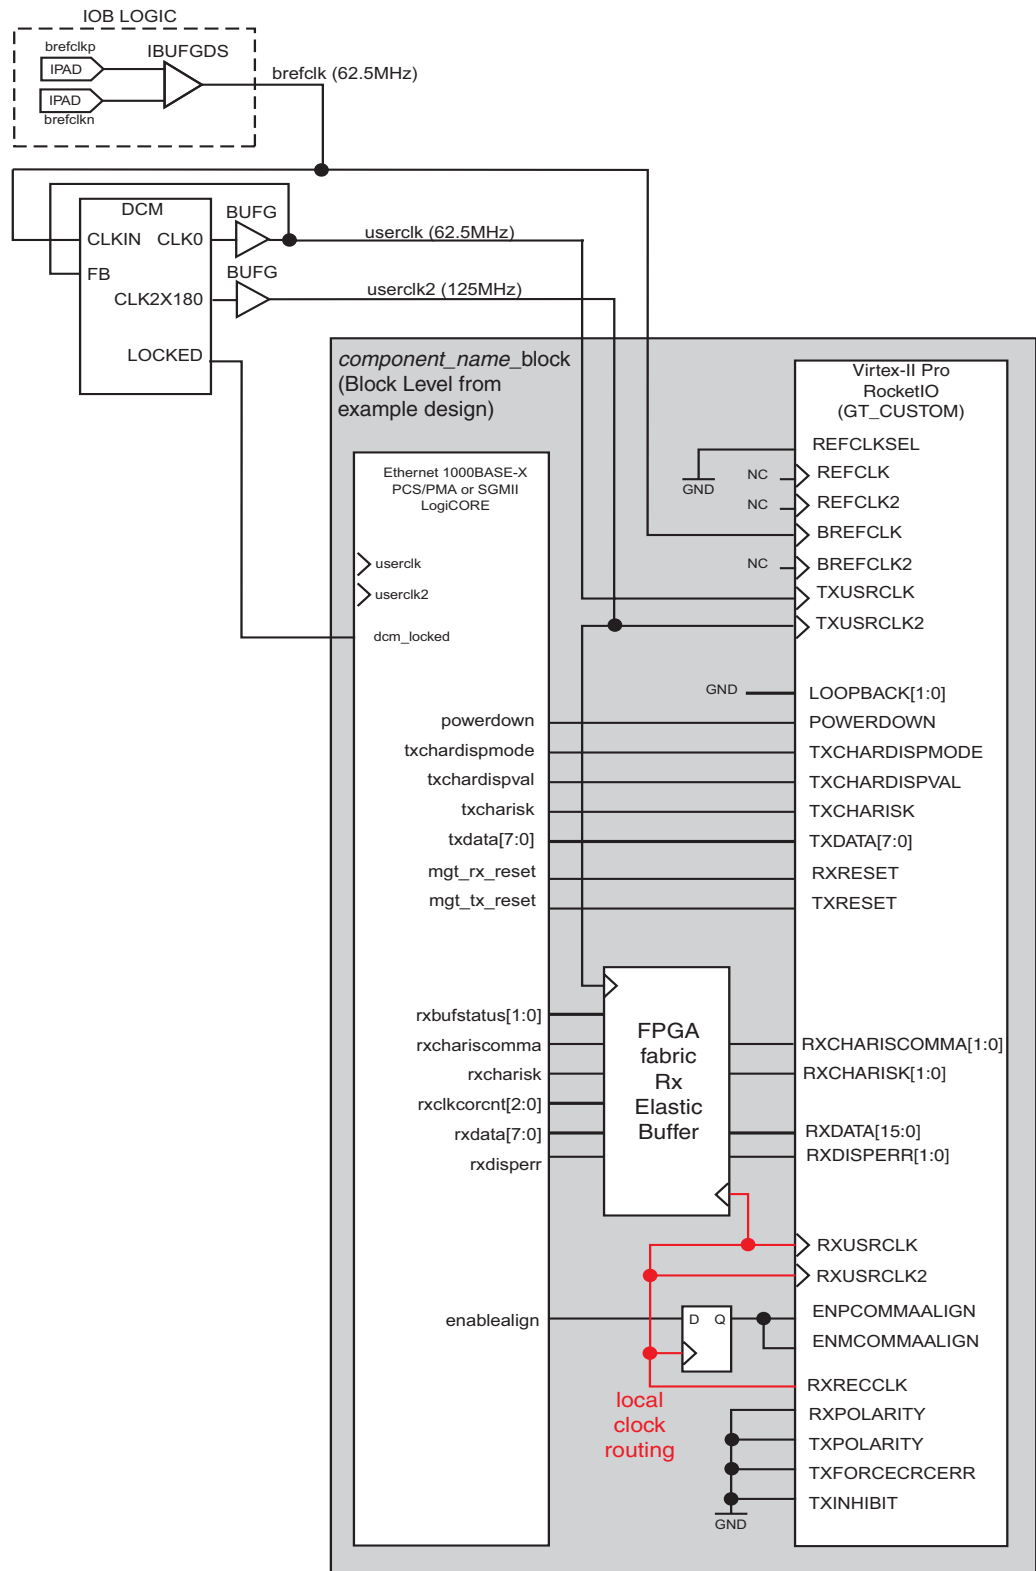


Figure 8-3: SGMII Connection to a Virtex-II Pro RocketIO Transceiver

Virtex-4 Devices for SGMII or Dynamic Standards Switching

The core is designed to integrate with the Virtex-4 MGT. The connections and logic required between the core and MGT transceiver are illustrated in [Figure 8-4](#)—the signal names and logic in the figure precisely match those delivered with the example design when an MGT transceiver is used.

Note: A small logic shim (included in the “block” level wrapper) is required to convert between the port differences between the Virtex-II Pro and Virtex-4 MGTs. This is not illustrated in [Figure 8-4](#).

The MGT clock distribution in Virtex-4 devices is column-based and consists of multiple MGT tiles (that contain two MGTs each). For this reason, the MGT transceiver wrapper delivered with the core always contains two MGT instantiations, even if only a single MGT is in use. [Figure 8-4](#) illustrates only a single MGT for clarity.

A `GT11CLK_MGT` primitive is also instantiated to derive the reference clocks required by the MGT column-based tiles. See the *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide* (UG076) for more information about layout and clock distribution.

The 250 MHz reference clock from the `GT11CLK_MGT` primitive is routed to the MGT, which is configured to internally synthesize a 125 MHz clock. This is output on the `TXOUTCLK1` port of the MGT and once placed onto global clock routing, can be used by all core logic. This clock is input back into the MGT on the user interface clock port `txusrclk2`. With the attribute settings applied to the MGT from the example design, the `txusrclk` port is derived internally within the MGT using the internal clock dividers and does not need to be provided from the FPGA fabric.

It can be seen from [Figure 8-4](#) that the Rx Elastic Buffer is implemented in the FPGA fabric between the MGT and the core. This replaces the Rx Elastic Buffer in the MGT (which is bypassed).

This alternative Receiver Elastic Buffer uses a single block RAM to create a buffer twice as large as the one present in the MGT. It is able to cope with larger frame sizes before clock tolerances accumulate and result in emptying or filling of the buffer. This is necessary to guarantee SGMII operation at 10 Mbps where each frame size is effectively 100 times larger than the same frame would be at 1 Gbps because each byte is repeated 100 times (see “[Designing with Client-side GMII for the SGMII Standard,](#)” page 59).

In bypassing the MGT Rx Elastic Buffer, data is clocked out of the MGT synchronously to `rxrecclk1`. This clock can be placed on a BUFR component and is used to synchronize the transfer of data between the MGT and the Elastic Buffer, as illustrated in [Figure 8-4](#). See also “[Virtex-4 RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints,](#)” page 166.

The MGT transceivers require a calibration block to be included in the fabric logic. The example design provided with the core instantiates calibration blocks as required. Calibration blocks require a clock source of between 25 to 50 MHz, which is shared with the Dynamic Reconfiguration Port (DRP) of the MGT, named `clk` in the example design. See Xilinx [Answer Record 22477](#) for more information.

[Figure 8-4](#) also illustrates the `TX_SIGNAL_DETECT` and `RX_SIGNAL_DETECT` ports of the calibration block, which should be driven to indicate whether or not dynamic data is being transmitted and received through the MGT (see [Virtex-4 Errata](#)). However, `RX_SIGNAL_DETECT` is connected to the `signal_detect` port of the example design. `signal_detect` is intended to indicate to the core that valid data is being received. When not asserted, the calibration block will switch the MGT into loopback to force dynamic data through the MGT receiver path.

Caution! The PHY connected via SGMII may always provide dynamic SGMII data (when powered up). If not, and if `signal_detect` is not present, the `RX_SIGNAL_DETECT` port of the calibration block must be driven by an alternative method. See *XAPP732* for more information.

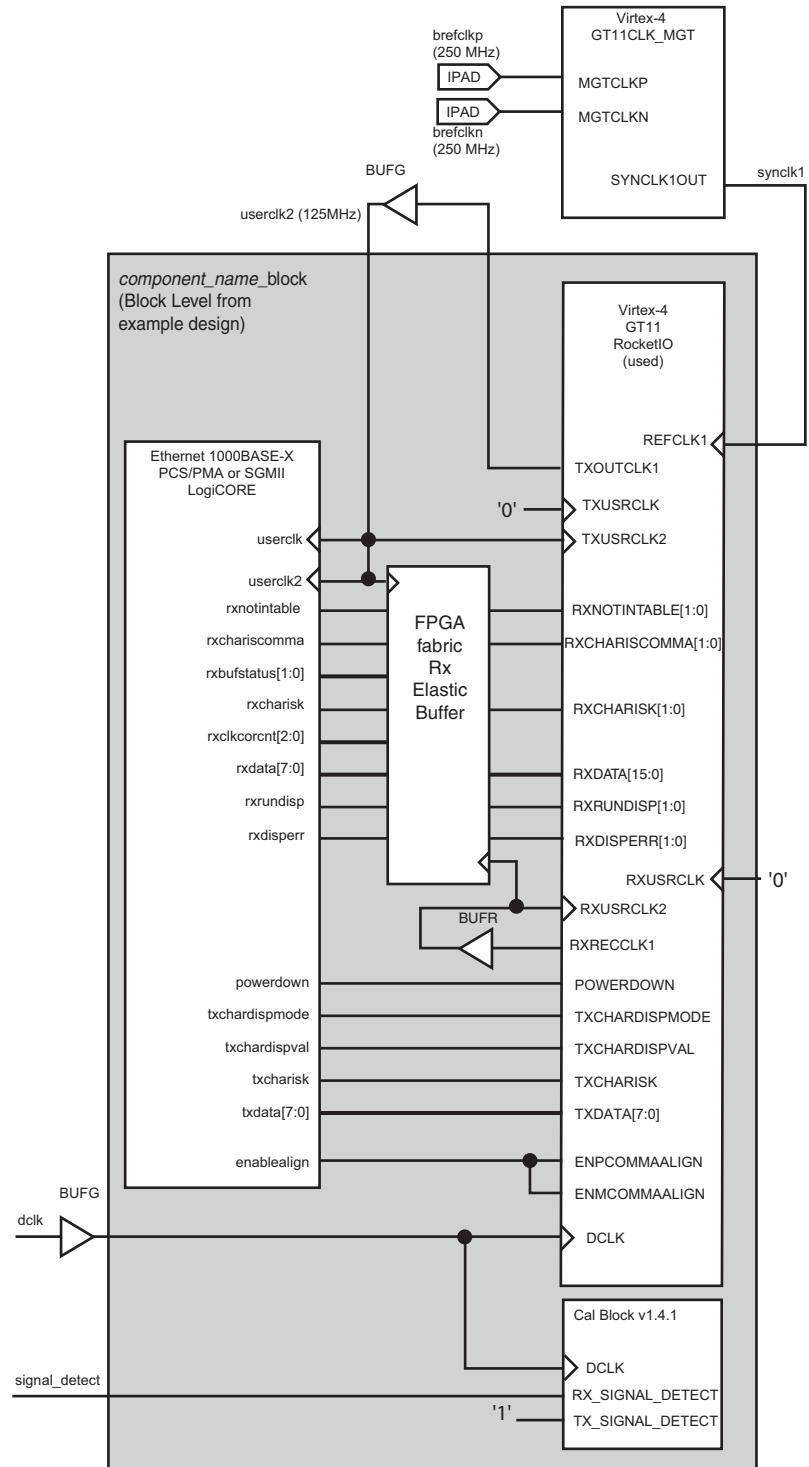


Figure 8-4: SGMII Connection to a Virtex-4 MGT

Virtex-5 LXT or SXT Devices for SGMII or Dynamic Standards Switching

The core is designed to integrate with the Virtex-5 RocketIO GTP transceiver. The connections and logic required between the core and GTP transceiver are illustrated in [Figure 8-5](#)—the signal names and logic in the figure precisely match those delivered with the example design when a GTP transceiver is used.

Note: A small logic shim (included in the block” level wrapper) is required to convert between the port differences between the Virtex-II Pro and Virtex-5 RocketIO GTP transceiver. This is not illustrated in [Figure 8-5](#).

A GTP tile consists of a pair of transceivers. For this reason, the GTP transceiver wrapper delivered with the core will always contain two GTP transceiver instantiations, even if only a single GTP is in use. [Figure 8-5](#) illustrates only a single GTP transceiver for clarity.

The 125 MHz differential reference clock is routed to the GTP transceiver, which is configured to output a version of this clock on the REFCLKOUT port, and once placed onto global clock routing can be used by all core logic. This clock is input back into the GTP transceiver on the user interface clock port txusrclk and txusrclk2.

It can be seen from [Figure 8-5](#) that the Rx Elastic Buffer is implemented in the FPGA fabric between the GTP transceiver and the core; this replaces the Rx Elastic Buffer in the GTP transceiver.

This alternative Receiver Elastic Buffer uses a single block RAM to create a buffer twice as large as the one present in the GTP transceiver. It is able to cope with larger frame sizes before clock tolerances accumulate and result in emptying or filling of the buffer. This is necessary to guarantee SGMII operation at 10 Mbps where each frame size is effectively 100 times larger than the same frame would be at 1 Gbps because each byte is repeated 100 times (see “[Designing with Client-side GMII for the SGMII Standard,](#)” page 59).

With this fabric Rx Elastic Buffer implementation, data is clocked out of the GTP transceiver synchronously to rxreclk0. This clock can be placed on a BUFR component and is used to synchronize the transfer of data between the GTP and the Elastic Buffer, as illustrated in [Figure 8-5](#). See also “[Virtex-5 RocketIO GTP Transceivers for SGMII or Dynamic Standards Switching Constraints,](#)” page 167.

Virtex-5 RocketIO GTP Wizard

The two wrapper files immediately around the GTP transceiver pair, `rocketio_wrapper_gtp_tile` and `rocketio_wrapper_gtp` (see [Figure 8-5](#)), are generated from the *RocketIO GTP Wizard*. These files apply all the gigabit Ethernet attributes. Consequently, these files can be regenerated by customers and therefore be easily targeted at ES or Production silicon. Note that this core targets production silicon.

The CORE Generator log file (XCO file) which was created when the *RocketIO GTP Wizard* project was generated is available in the following location:

```
<project_directory>/<component_name>/example_design/transceiver/  
rocketio_wrapper_gtp.xco
```

This file can be used as an input to the CORE Generator to regenerate the RocketIO wrapper files. The XCO file itself contains a list of all of the GTP Wizard attributes which were used. For further information, please refer to the *Virtex-5 RocketIO GTP Wizard Getting Started Guide* (UG188) and the *CORE Generator Guide*, at www.xilinx.com/support/software_manuals.htm

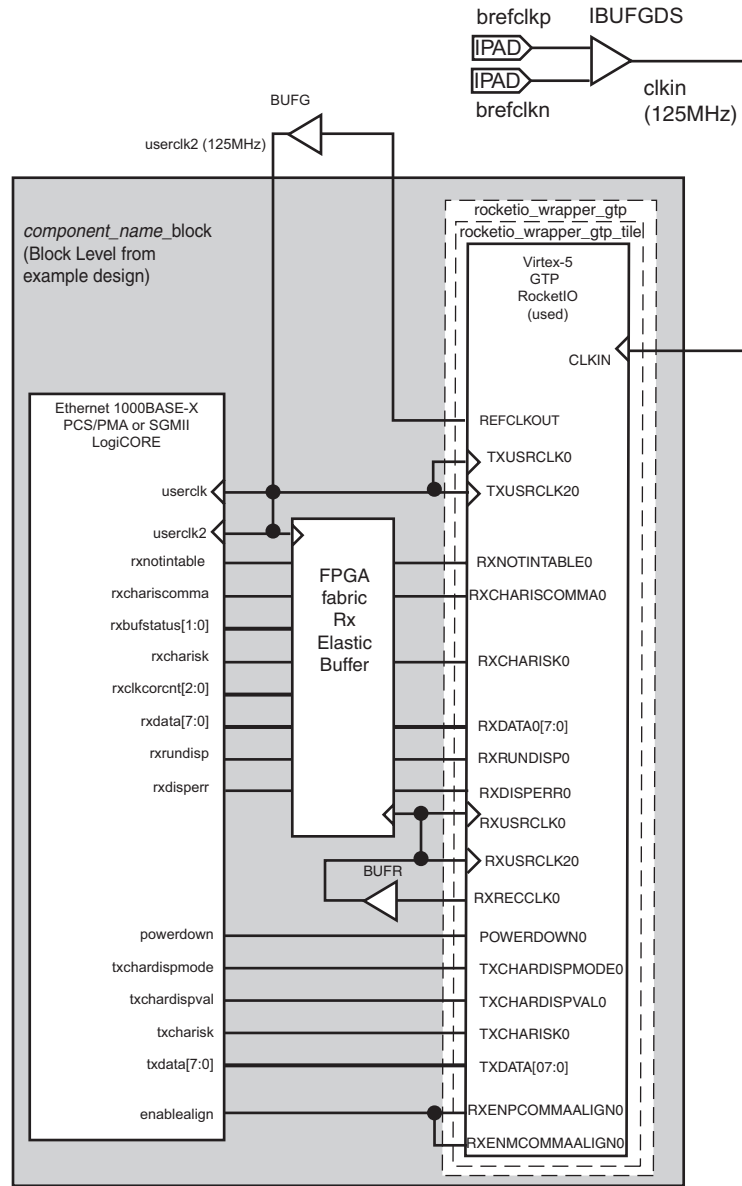


Figure 8-5: SGMII Connection to a Virtex-5 RocketIO GTP Transceiver

Virtex-5 FXT Devices for SGMII or Dynamic Standards Switching

The core is designed to integrate with the Virtex-5 RocketIO GTX transceiver. The connections and logic required between the core and GTX transceiver are illustrated in [Figure 8-6](#)—the signal names and logic in the figure precisely match those delivered with the example design when a GTX transceiver is used.

Note: A small logic shim (included in the block” level wrapper) is required to convert between the port differences between the Virtex-II Pro and Virtex-5 RocketIO GTX transceiver. This is not illustrated in [Figure 8-6](#).

A GTX tile consists of a pair of transceivers. For this reason, the GTX transceiver wrapper delivered with the core will always contain two GTX transceiver instantiations, even if only a single GTX is in use. [Figure 8-6](#) illustrates only a single GTX transceiver for clarity.

The 125 MHz differential reference clock is routed directly to the GTX transceiver. The GTX transceiver is configured to output a version of this clock on the REFCLKOUT port: this is then routed to a DCM.

From the DCM, the CLK0 port (125MHz) is placed onto global clock routing and can be used as the 125MHz clock source for all core logic: this clock is also input back into the GTX transceiver on the user interface clock port txusrclk2.

From the DCM, the CLKDV port (62.5MHz) is placed onto global clock routing and is input back into the GTX transceiver on the user interface clock port txusrclk.

It can be seen from [Figure 8-6](#) that the Rx Elastic Buffer is implemented in the FPGA fabric between the GTX transceiver and the core; this replaces the Rx Elastic Buffer in the GTX transceiver.

This alternative Receiver Elastic Buffer uses a single block RAM to create a buffer twice as large as the one present in the GTX transceiver. It is able to cope with larger frame sizes before clock tolerances accumulate and result in emptying or filling of the buffer. This is necessary to guarantee SGMII operation at 10 Mbps where each frame size is effectively 100 times larger than the same frame would be at 1 Gbps because each byte is repeated 100 times (see “[Designing with Client-side GMII for the SGMII Standard,](#)” page 59).

With this fabric Rx Elastic Buffer implementation, data is clocked out of the GTX transceiver synchronously to rxrecclk0 (62.5MHz) on a 16-bit interface. This clock can be placed on a BUFR component and is used to synchronize the transfer of data between the GTX and the Elastic Buffer, as illustrated in [Figure 8-6](#). See also “[Virtex-5 RocketIO GTX Transceivers for SGMII or Dynamic Standards Switching Constraints,](#)” page 168.

Virtex-5 RocketIO GTX Wizard

The two wrapper files immediately around the GTX transceiver pair, `rocketio_wrapper_gtx_tile` and `rocketio_wrapper_gtx` (see [Figure 8-6](#)), are generated from the *RocketIO GTX Wizard*. These files apply all the gigabit Ethernet attributes. Consequently, these files can be regenerated by customers and therefore be easily targeted at ES or Production silicon. Note that this core targets production silicon.

The CORE Generator log file (XCO file) which was created when the *RocketIO GTX Wizard* project was generated is available in the following location:

```
<project_directory>/<component_name>/example_design/transceiver/  
rocketio_wrapper_gtx.xco
```

This file can be used as an input to the CORE Generator to regenerate the RocketIO wrapper files. The XCO file itself contains a list of all of the GTX Wizard attributes which were used. For further information, please refer to the *Virtex-5 RocketIO GTX Wizard*

Getting Started Guide and the CORE Generator Guide, at www.xilinx.com/support/software_manuals.htm




Figure 8-6: SGMII Connection to a Virtex-5 RocketIO GTX Transceiver

Clock Sharing - Multiple Cores with RocketIO, Fabric Elastic Buffer

Virtex-II Pro Devices

[Figure 8-7](#) illustrates sharing clock resources across multiple instantiations of the core on the same half of the device when using the core with the RocketIO MGT. [Figure 8-7](#) illustrates only two cores; however, more can be added by instantiating the cores using the block level (from the example design) and sharing `userclk`, `userclk2`, and `brefclk` across all instantiations. For each core, `userclk` and `userclk2` must always be derived from the `brefclk` or `refclk` used by that core.

Each MGT instantiated has its own independent clock domain synchronous to `RXRECCLK` which is placed on local clock routing. Each local clock domain must have area constraints added to place it in the region of the MGT. See “[Virtex-II Pro RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints](#),” page 163.

When using the fixed routing resources of `brefclk`, MGTs along the top edge of the device must use a separate `brefclk` routing resource to those along the bottom edge of

the device. For more information, see the *Virtex-II Pro RocketIO Transceiver User Guide*. Each `brefclk` domain must use its own DCM to derive its version of `userclk` and `userclk2`.

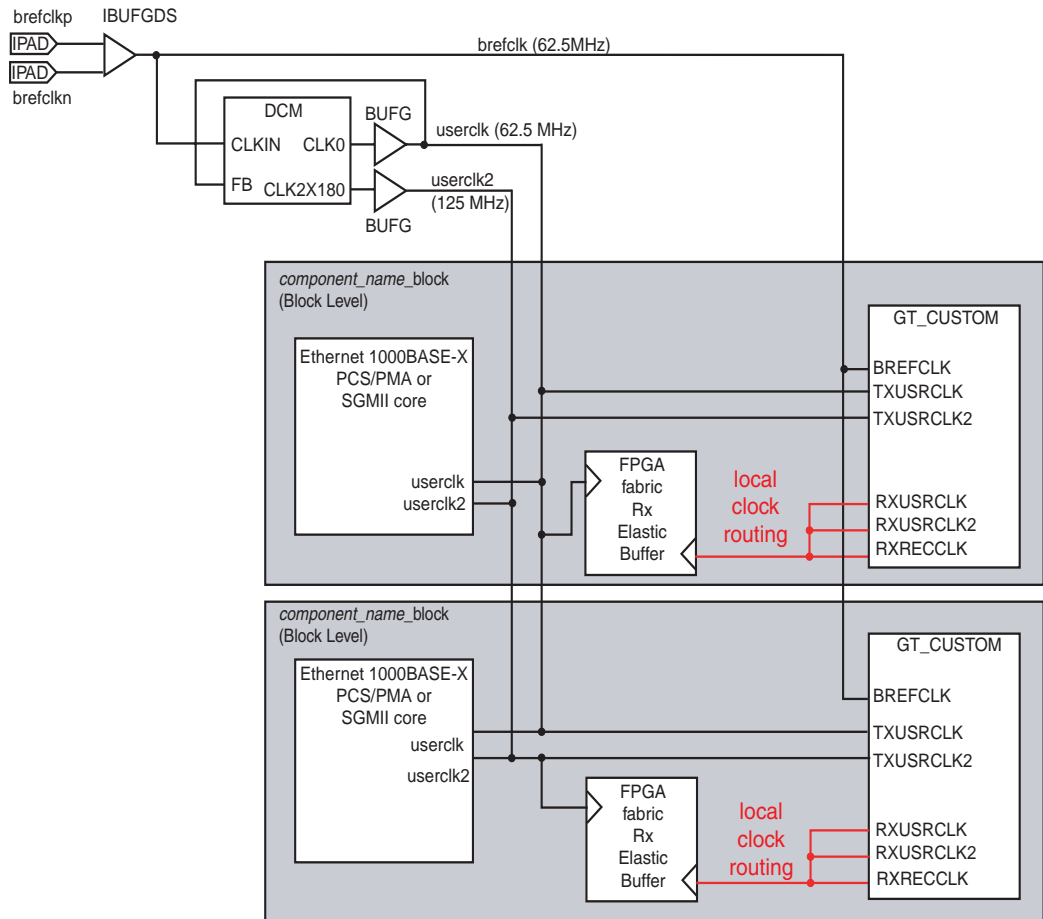


Figure 8-7: Clock Management with Multiple Core Instances with Virtex-II Pro RocketIO Transceivers for SGMII

Virtex-4 FX Devices

Figure 8-8 illustrates sharing clock resources across multiple instantiations of the core when using the Virtex-4 RocketIO MGT. Note that the example design, when using the Virtex-4 family, can be generated to connect either a single instance of the core, or connect a pair of core instances to the transceiver pair present in a MGT tile. Figure 8-8 illustrates two instantiations of the block level, and each block level contains a pair of cores, illustrating clock sharing between four cores.

More cores can be added by continuing to instantiate extra block level modules. Share clocks only between the MGTs in a single column. For each column, use a single `brefclk_p` and `brefclk_n` differential clock pair and connect this to a `GT11CLK_MGT` primitive. The clock output from this primitive should be shared across all used MGT tiles in the column. See the *Virtex-4 RocketIO Multi-Gigabit Transceiver User Guide* for more information.

To provide the 125 MHz clock for all core instances, select a `TXOUTCLK1` port from any MGT. This can be routed onto global clock routing using a `BUFG` as illustrated, and shared between all cores and MGTs in the column.

Each MGT and core pair instantiated has its own independent clock domain synchronous to `RXRECCLK1` which is placed on regional clock routing using a `BUFR`, as illustrated in Figure 8-8—these cannot be shared across multiple MGTs. Although not illustrated in Figure 8-8, `dc1k` (the clock used for the calibration blocks and for the Dynamic Reconfiguration Port (DRP) of the MGTs) can also be shared.

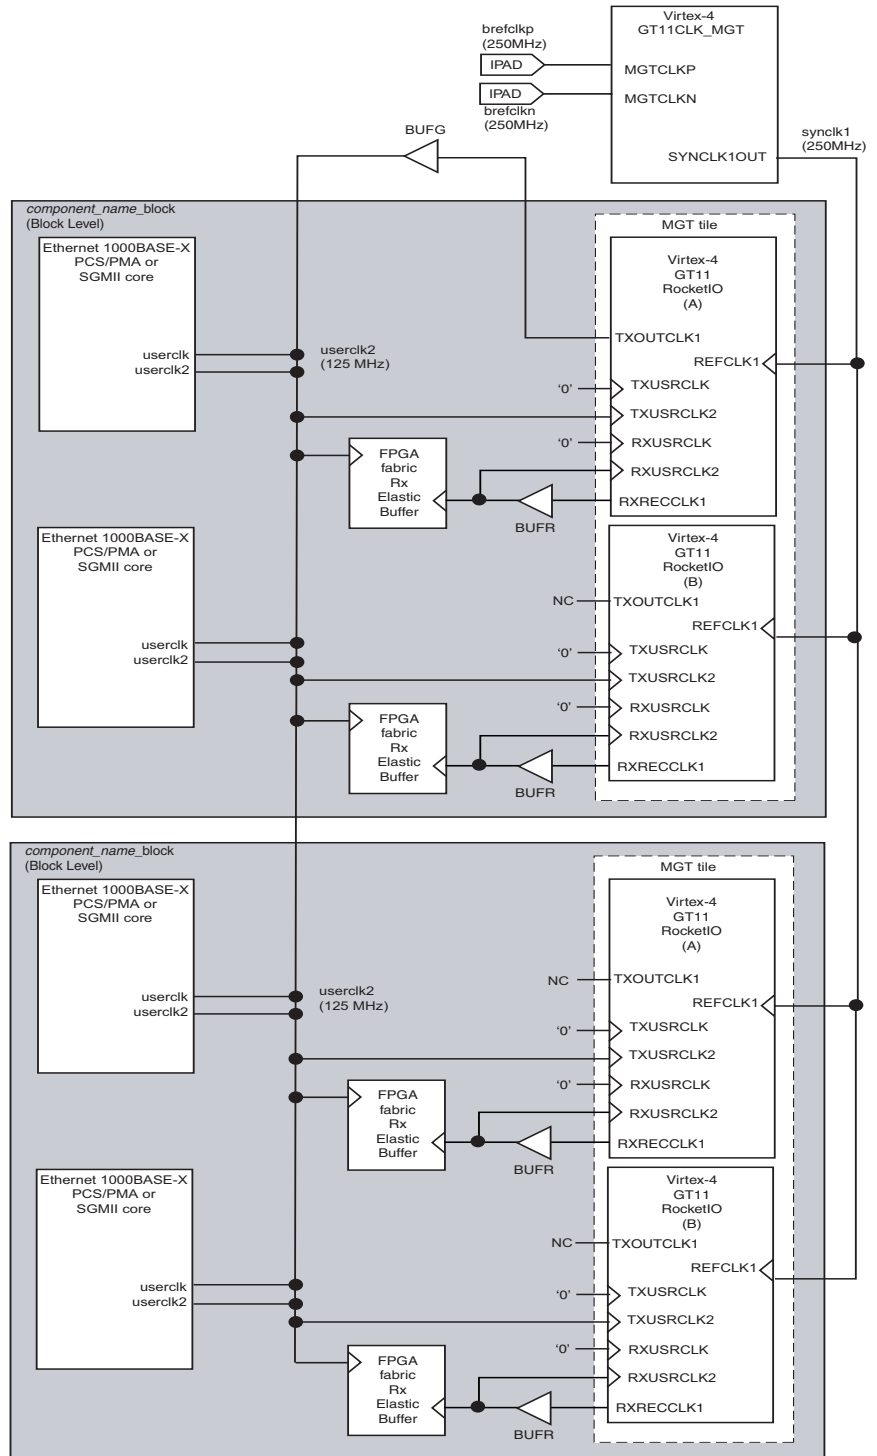


Figure 8-8: Clock Management with Multiple Core Instances with Virtex-4 MGTs for SGMII

Virtex-5 LXT and SXT Devices

Figure 8-9 illustrates sharing clock resources across multiple instantiations of the core when using the Virtex-5 RocketIO GTP transceiver. The example design can be generated to connect either a single instance of the core, or connect a pair of core instances to the transceiver pair present in a GTP transceiver tile. Figure 8-9 illustrates two instantiations of the block level, and each block level contains a pair of cores. Figure 8-9 illustrates clock sharing between four cores.

More cores can be added by instantiating extra block level modules. Share the `brefclk_p` and `brefclk_n` differential clock pairs. See the *Virtex-5 RocketIO GTP Transceiver User Guide* for more information.

To provide the 125 MHz clock for all core instances, select a REFCLKOUT port from any GTP transceiver. This can be routed onto global clock routing using a BUFG as illustrated and shared between all cores and GTP transceivers in the column.

Each GTP and core pair instantiated has its own independent clock domains synchronous to RXRECCLK0 and RXRECCLK1. These are placed on regional clock routing using a BUFR, as illustrated in Figure 8-9, and cannot be shared across multiple GTP transceivers.

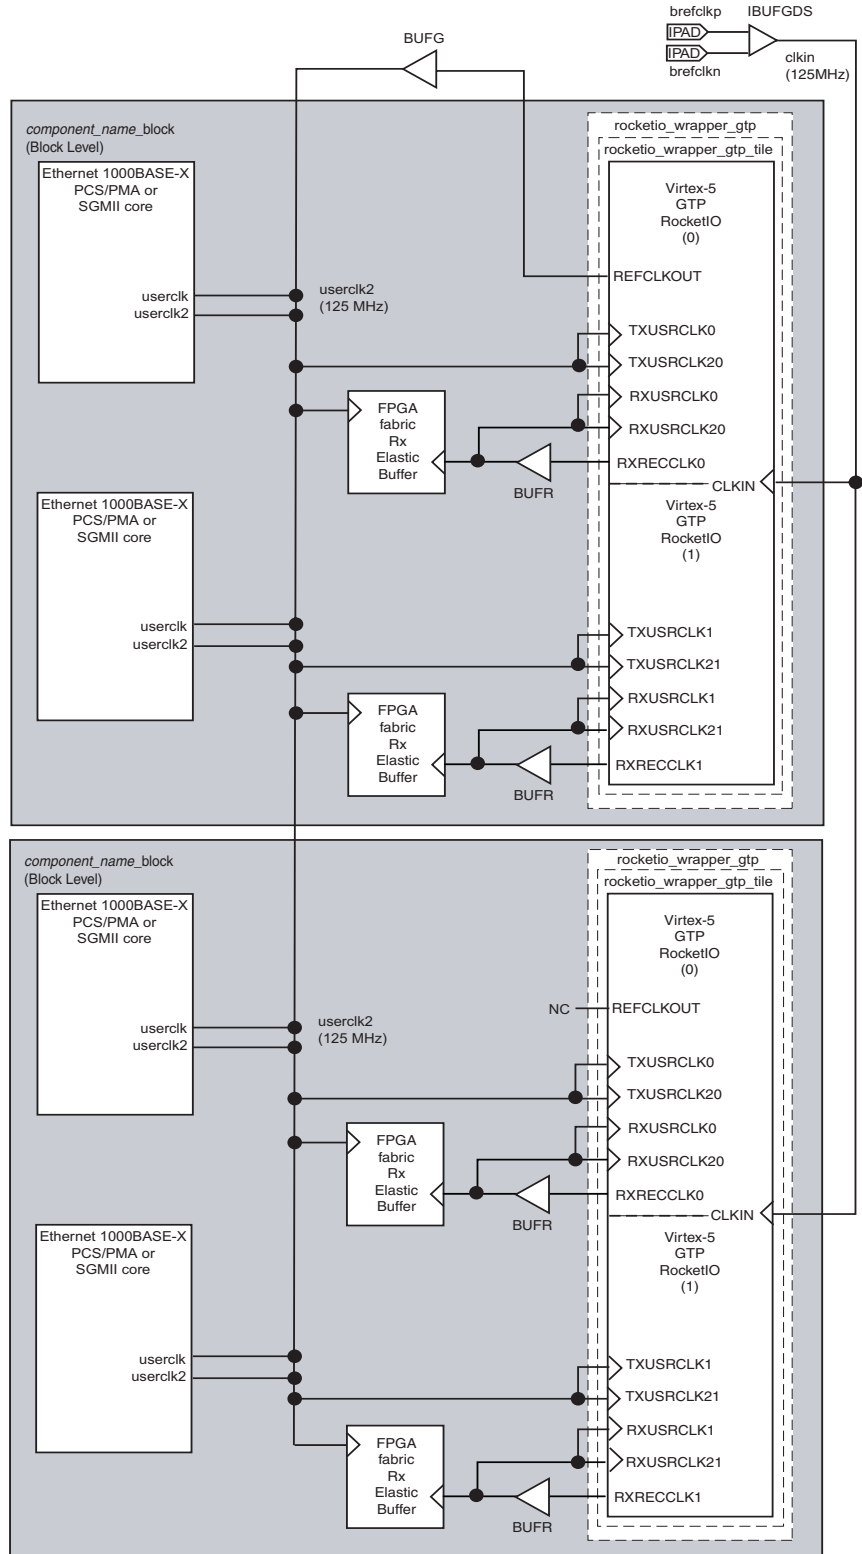


Figure 8-9: Clock Management with Multiple Core Instances with Virtex-5 GTP RocketIO Transceivers for SGMII

Virtex-5 FXT Devices

[Figure 8-9](#) illustrates sharing clock resources across multiple instantiations of the core when using the Virtex-5 RocketIO GTX transceiver. The example design can be generated to connect either a single instance of the core, or connect a pair of core instances to the transceiver pair present in a GTX transceiver tile. [Figure 8-9](#) illustrates two instantiations of the block level, and each block level contains a pair of cores. [Figure 8-9](#) illustrates clock sharing between four cores.

More cores can be added by instantiating extra block level modules. Share the `brefclk_p` and `brefclk_n` differential clock pairs. See the *Virtex-5 RocketIO GTX Transceiver User Guide* for more information.

To provide the FPGA fabric clocks for all core instances, select a `REFCLKOUT` port from any GTX transceiver and route this to a single DCM. The `CLK0` (125MHz) and `CLKDV` (62.5MHz) outputs from this DCM, placed onto global clock routing using a `BUFGs`, can be shared across all core instances and GTX transceivers as illustrated.

Each GTX and core pair instantiated has its own independent clock domains synchronous to `RXRECCLK0` and `RXRECCLK1`. These are placed on regional clock routing using a `BUFR`, as illustrated in [Figure 8-9](#), and cannot be shared across multiple GTX transceivers.

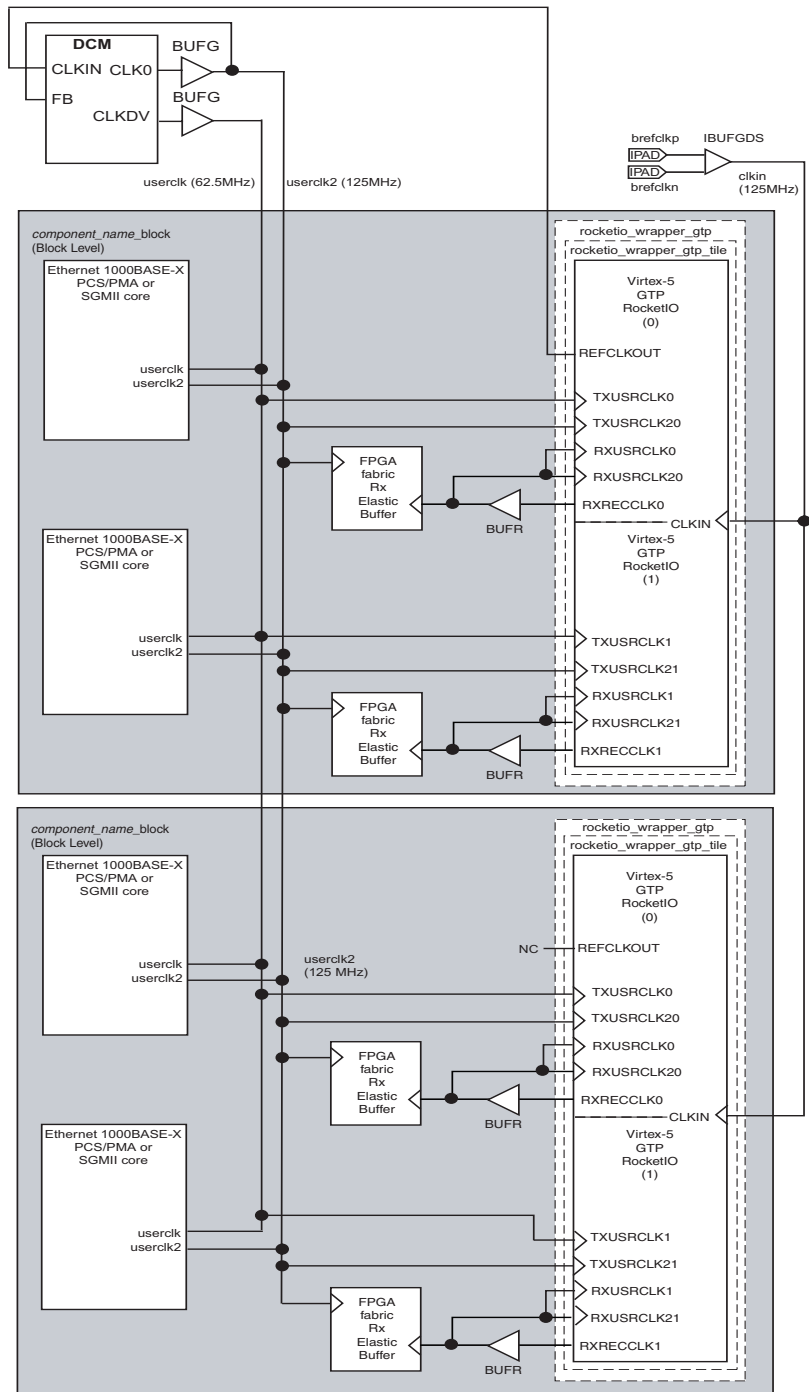


Figure 8-10: Clock Management with Multiple Core Instances with Virtex-5 GTX RocketIO Transceivers for SGMII

Configuration and Status

This chapter provides general guidelines for configuring and monitoring the Ethernet 1000BASE-X PCS/PMA or SGMII core, including a detailed description of the core management registers. It also describes Configuration Vector and status signals, an alternative to using the optional MDIO Management Interface.

MDIO Management Interface

When the optional MDIO Management Interface is selected, configuration and status of the core is achieved by the Management Registers accessed through the serial Management Data Input/Output Interface (MDIO). See “[MDIO Management Interface](#)” in [Chapter 3](#) for more information.

MDIO Bus System

The MDIO interface for 1 Gbps operation (and slower speeds) is defined in *IEEE 802.3*, clause 22. This two-wire interface consists of a clock (MDC) and a shared serial data line (MDIO). The maximum permitted frequency of MDC is set at 2.5 MHz.

[Figure 9-1](#) illustrates an example MDIO bus system.

An Ethernet MAC is shown as the MDIO bus master (the Station Management (STA) entity).

Two PHY devices are shown connected to the same bus, both of which are MDIO slaves (MDIO Managed Device (MMD) entities).

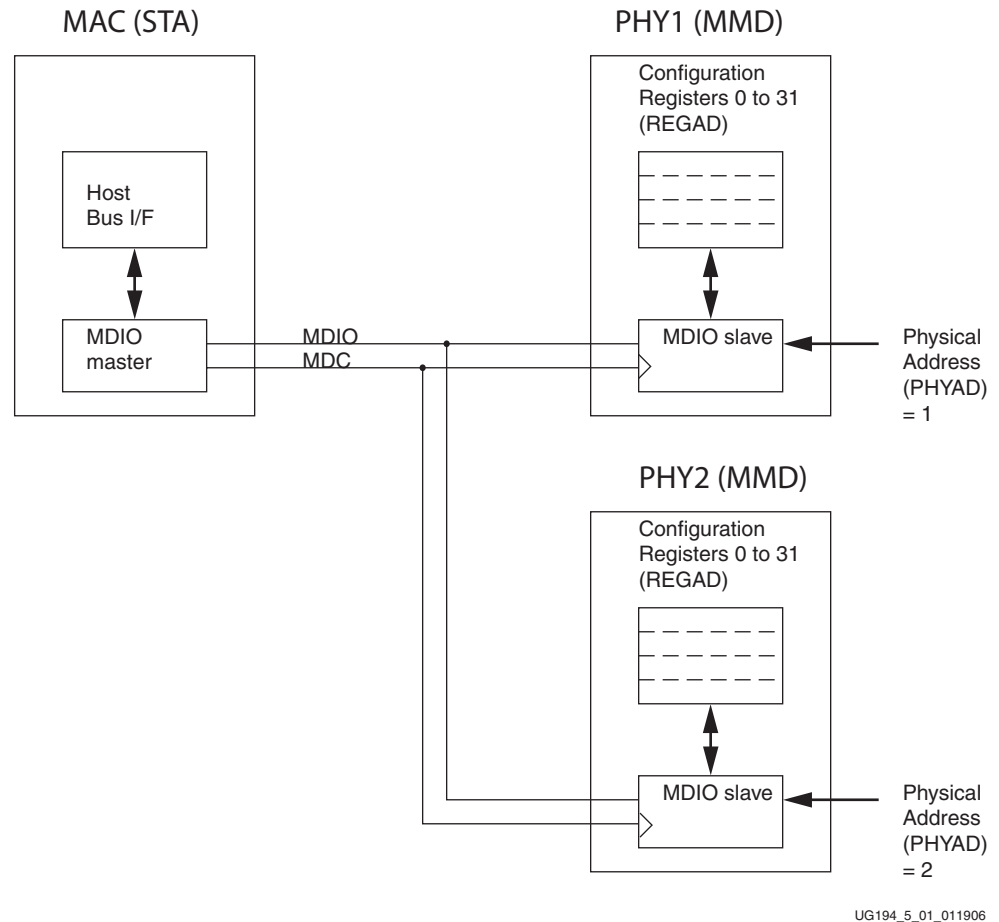


Figure 9-1: A Typical MDIO-managed System

The MDIO bus system is a standardized interface for accessing the configuration and status registers of Ethernet PHY devices. In the example illustrated, the Management Host Bus I/F of the Ethernet MAC is able to access the configuration and status registers of two PHY devices via the MDIO bus.

MDIO Transactions

All transactions, read or write, are initiated by the MDIO master. All MDIO slave devices, when addressed, must respond. MDIO transactions take the form of an MDIO frame, containing fields for transaction type, address and data. This MDIO frame is transferred across the MDIO wire synchronously to MDC. The abbreviations are used in this section are explained in Table 9-1.

Table 9-1: Abbreviations and Terms

Abbreviation	Term
PRE	Preamble
ST	Start of frame
OP	Operation code

Table 9-1: Abbreviations and Terms (Continued)

Abbreviation	Term
PHYAD	Physical address
REGAD	Register address
TA	Turnaround

Write Transaction

Figure 9-2 shows a write transaction across the MDIO, defined as OP="01." The addressed PHY device (with physical address PHYAD) takes the 16-bit word in the Data field and writes it to the register at REGAD.

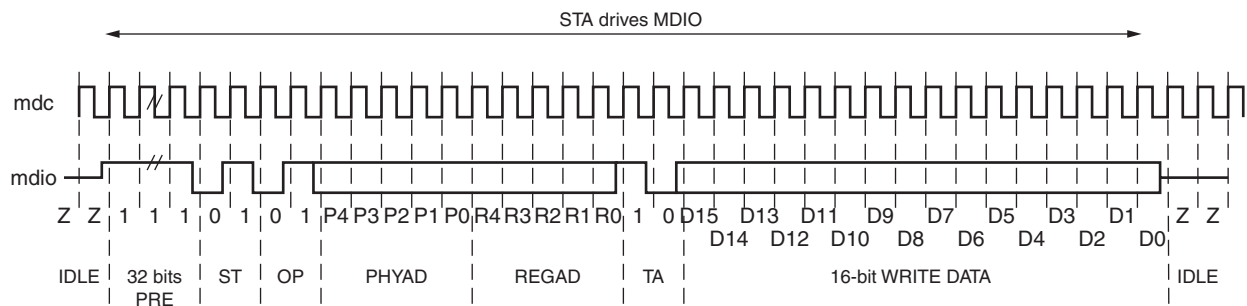


Figure 9-2: MDIO Write Transaction

Read Transaction

Figure 9-3 shows a read transaction, defined as OP="10." The addressed PHY device (with physical address PHYAD) takes control of the MDIO wire during the turnaround cycle and then returns the 16-bit word from the register at REGAD.

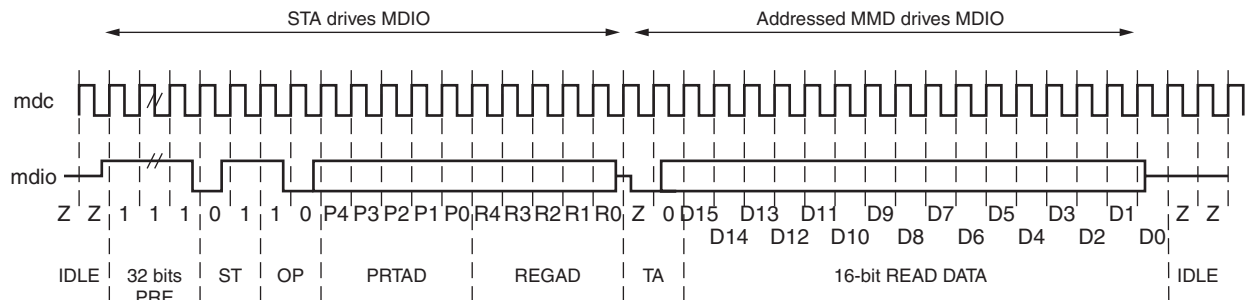


Figure 9-3: MDIO Read Transaction

MDIO Addressing

MDIO Addresses consists of two stages: Physical Address (PHYAD) and Register Address (REGAD).

Physical Address (PHYAD)

As shown in Figure 9-1, two PHY devices are attached to the MDIO bus. Each of these has a different physical address. To address the intended PHY, its physical address should be

known by the MDIO master (in this case an Ethernet MAC), and placed into the PHYAD field of the MDIO frame (see “[MDIO Transactions](#)”).

The PHYAD field for an MDIO frame is a 5-bit binary value capable of addressing 32 unique addresses. However, every MDIO slave must respond to physical address 0. This requirement dictates that the physical address for any particular PHY must not be set to 0 to avoid MDIO contention. Physical Addresses 1 through to 31 can be used to connect up to 31 PHY devices onto a single MDIO bus.

Physical Address 0 can be used to write a single command that is obeyed by all attached PHYs, such as a reset or power-down command.

Register Address (REGAD)

Having targeted a particular PHY using PHYAD, the individual configuration or status register within that particular PHY must now be addressed. This is achieved by placing the individual register address into the REGAD field of the MDIO frame (see “[MDIO Transactions](#)”).

The REGAD field for an MDIO frame is a 5-bit binary value capable of addressing 32 unique addresses. The first 16 of these (registers 0 to 15) are defined by the *IEEE 802.3*. The remaining 16 (registers 16 to 31) are reserved for PHY vendors own register definitions.

For details of the register map of PHY layer devices and a more extensive description of the operation of the MDIO Interface, see *IEEE 802.3-2002*.

Connecting the MDIO to an Internally Integrated STA

The MDIO ports of the Ethernet 1000BASE-X PCS/PMA or SGMII core can be connected to the MDIO ports of an internally integrated Station Management (STA) entity, such as the MDIO port of the 1-Gigabit Ethernet MAC core (see “[Integrating with the 1-Gigabit Ethernet MAC Core](#),” page 179) or the Tri-Mode Ethernet MAC core (see “[Integrating with the Tri-Mode Ethernet MAC Core](#),” page 185).

Connecting the MDIO to an External STA

[Figure 9-4](#) shows the MDIO ports of the Ethernet 1000BASE-X PCS/PMA or SGMII core connected to the MDIO of an external STA entity. In this situation, `mdio_in`, `mdio_out`, and `mdio_tri` must be connected to a Tri-State buffer to create a bidirectional wire, `mdio`. This Tri-State buffer can either be external to the FPGA, or internally integrated by using an IOB `IOBUF` component with an appropriate SelectIO™ standard suitable for the external PHY.

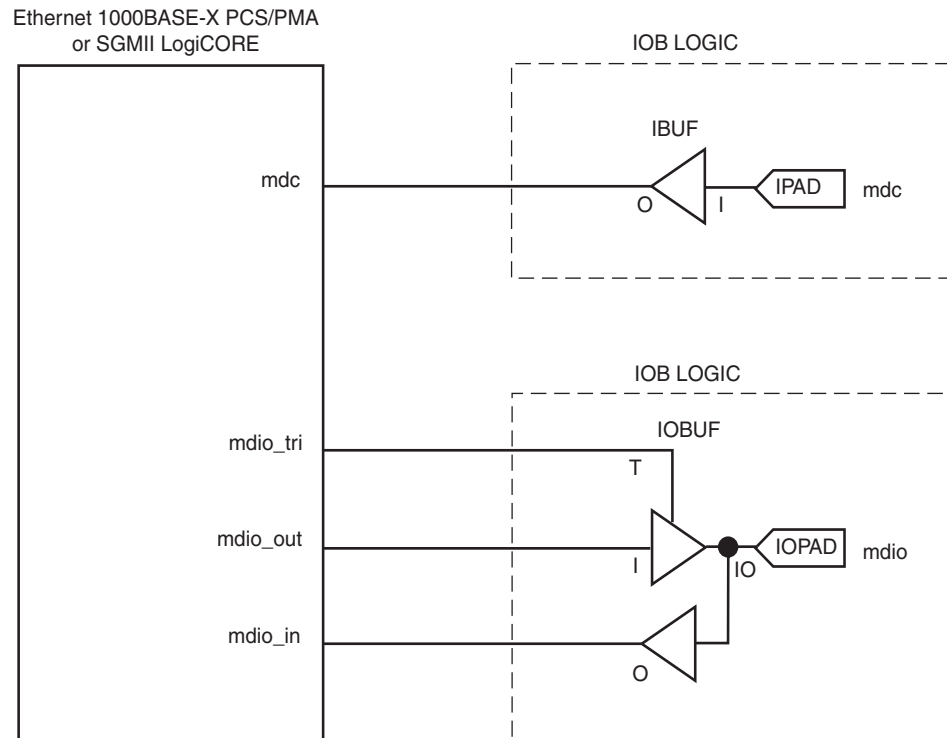


Figure 9-4: Creating an External MDIO Interface

Management Registers

The contents of the Management Registers can be accessed using the REGAD field of the MDIO frame. Contents will vary depending on the CORE Generator options, and are defined in the following sections in this guide.

- [1000BASE-X Standard Using the Optional Auto-Negotiation](#)
- [1000BASE-X Standard Without the Optional Auto-Negotiation](#)
- [SGMII Standard Using the Optional Auto-Negotiation](#)
- [SGMII Standard without the Optional Auto-Negotiation](#)
- [Both 1000BASE-X and SGMII Standards](#)

1000BASE-X Standard Using the Optional Auto-Negotiation

More information on the 1000BASE-X PCS Registers can be found in clause 37 and clause 22 of the *IEEE 802.3* specification. Registers at undefined addresses are read-only and return 0s.

Table 9-2: MDIO Registers for 1000BASE-X with Auto-Negotiation

Register Address	Register Name
0	Control Register
1	Status Register

Table 9-2: MDIO Registers for 1000BASE-X with Auto-Negotiation (Continued)

Register Address	Register Name
2,3	PHY Identifier
4	Auto-Negotiation Advertisement Register
5	Auto-Negotiation Link Partner Ability Base Register
6	Auto-Negotiation Expansion Register
7	Auto-Negotiation Next Page Transmit Register
8	Auto-Negotiation Next Page Receive Register
15	Extended Status Register
16	Vendor Specific: Auto-Negotiation Interrupt Control

Register 0: Control Register

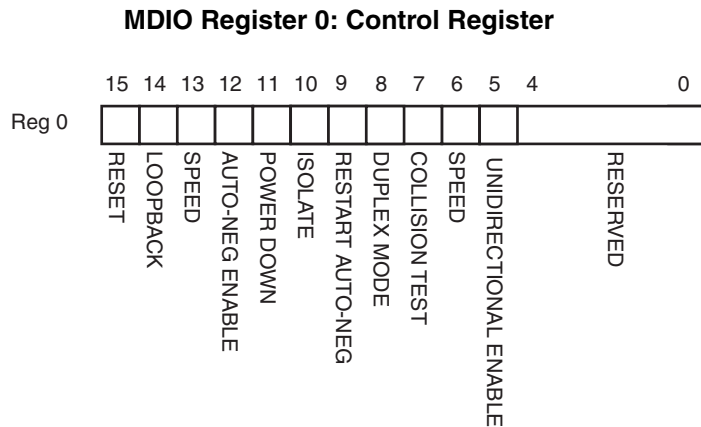


Table 9-3: Control Register (Register 0)

Bit(s)	Name	Description	Attributes	Default Value
0.15	Reset	1 = Core Reset 0 = Normal Operation	Read/write Self clearing	0
0.14	Loopback	1 = Enable Loopback Mode 0 = Disable Loopback Mode When used with a RocketIO transceiver, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to <i>ewrap</i> . When set to '1,' indicates to the external PMA module to enter loopback mode. See "Loopback," page 197.	Read/write	0

Table 9-3: Control Register (Register 0) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
0.13	Speed Selection (LSB)	Always returns a 0 for this bit. Together with bit 0.6, speed selection of 1000 Mbps is identified	Returns 0	0
0.12	Auto-Negotiation Enable	1 = Enable Auto-Negotiation Process 0 = Disable Auto-Negotiation Process	Read/write	1
0.11	Power Down	1 = Power down 0 = Normal operation With the PMA option, when set to '1' the RocketIO transceiver is placed in a low-power state. This bit requires a reset (see bit 0.15) to clear. With the TBI version this register bit has no effect.	Read/ write	0
0.10	Isolate	1 = Electrically Isolate PHY from GMII 0 = Normal operation	Read/write	1
0.9	Restart Auto-Negotiation	1 = Restart Auto-Negotiation Process 0 = Normal Operation	Read/write Self clearing	0
0.8	Duplex Mode	Always returns a '1' for this bit to signal Full-Duplex Mode.	Returns 1	1
0.7	Collision Test	Always returns a '0' for this bit to disable COL test.	Returns 0	0
0.6	Speed Selection (MSB)	Always returns a '1' for this bit. Together with bit 0.13, speed selection of 1000 Mbps is identified.	Returns 1	1
0.5	Unidirectional Enable	Enable transmit regardless of whether a valid link has been established.	Read/ write	0
0.4:0	Reserved	Always return 0s, writes ignored.	Returns 0s	00000

Register 1: Status Register

MDIO Register 1: Status Register

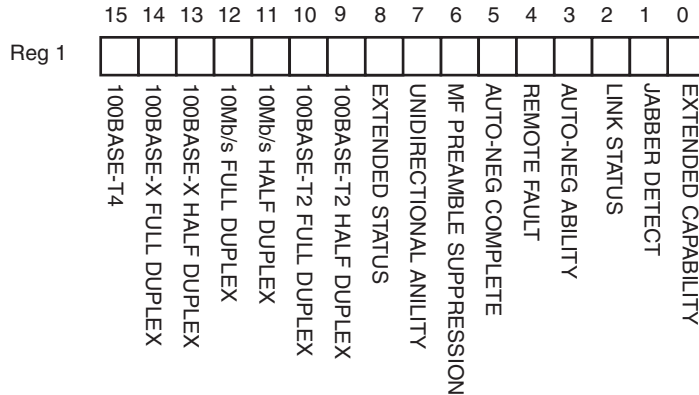


Table 9-4: Status Register (Register 1)

Bit(s)	Name	Description	Attributes	Default Value
1.15	100BASE-T4	Always returns a '0' as 100BASE-T4 is not supported.	Returns 0	0
1.14	100BASE-X Full Duplex	Always returns a '0' as 100BASE-X full duplex is not supported.	Returns 0	0
1.13	100BASE-X Half Duplex	Always returns a '0' as 100BASE-X half duplex is not supported.	Returns 0	0
1.12	10 Mbps Full Duplex	Always returns a '0' as 10 Mbps full duplex is not supported.	Returns 0	0
1.11	10 Mbps Half Duplex	Always returns a '0' as 10 Mbps half duplex is not supported	Returns 0	0
1.10	100BASE-T2 Full Duplex	Always returns a '0' as 100BASE-T2 full duplex is not supported.	Returns 0	0
1.9	100BASE-T2 Half Duplex	Always returns a '0' as 100BASE-T2 Half Duplex is not supported.	Returns 0	0
1.8	Extended Status	Always returns a '1' to indicate the presence of the Extended Register (Register 15).	Returns 1	1
1.7	Unidirectional Ability	Always returns a '1,' writes ignored	Returns 1	1
1.6	MF Preamble Suppression	Always returns a '1' to indicate that Management Frame Preamble Suppression is supported.	Returns 1	1
1.5	Auto- Negotiation Complete	1 = Auto-Negotiation process completed 0 = Auto-Negotiation process not completed	Read only	0

Table 9-4: Status Register (Register 1) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
1.4	Remote Fault	1 = Remote fault condition detected 0 = No remote fault condition detected	Read only Self-clearing on read	0
1.3	Auto-Negotiation Ability	Always returns a '1' for this bit to indicate that the PHY is capable of Auto-Negotiation.	Returns 1	1
1.2	Link Status ¹	1 = Link is up 0 = Link is down Latches '0' if Link Status goes down. Clears to current Link Status on read. See table note for Link Status behavior.	Read only Self-clearing on read	0
1.1	Jabber Detect	Always returns a '0' for this bit since Jabber Detect is not supported.	Returns 0	0
1.0	Extended Capability	Always returns a '0' for this bit since no extended register set is supported.	Returns 0	0

1. When high, the link is valid: synchronization of the link has been obtained and Auto-Negotiation (if present and enabled) has completed.

When low, a valid link has not been established. Either link synchronization has failed or Auto-Negotiation (if present and enabled) has failed to complete.

Regardless of whether Auto-Negotiation is enabled or disabled, there can be some delay to the deassertion of this signal following the loss of synchronization of a previously successful link. This is due to the Auto-Negotiation state machine which requires that synchronization is lost for an entire link timer duration before changing state. For more information, see the 802.3 specification.

Registers 2 and 3: PHY Identifiers

Registers 2 and 3: PHY Identifiers




Table 9-5: PHY Identifier (Registers 2 and 3)

Bit(s)	Name	Description	Attributes	Default Value
2.15:0	Organizationally Unique Identifier	Always return 0s	returns 0s	0000000000000000
3.15:10	Organizationally Unique Identifier	Always return 0s	returns 0s	000000
3.9:4	Manufacturer's model number	Always return 0s	returns 0s	000000
3.3:0	Revision Number	Always return 0s	returns 0s	0000

Register 4: Auto-Negotiation Advertisement

MDIO Register 4: Auto-Negotiation Advertisement




Table 9-6: Auto-Negotiation Advertisement Register (Register 4)

Bit(s)	Name	Description	Attributes	Default Value
4.15	Next Page	1 = Next Page functionality is advertised 0 = Next Page functionality is not advertised	read/write	0
4.14	Reserved	Always returns '0,' writes ignored	returns 0	0
4.13:12	Remote Fault	00 = No Error 01 = Offline 10 = Link Failure 11 = Auto-Negotiation Error	read/write self clearing to 00 after auto-negotiation	00
4.11:9	Reserved	Always return 0s, writes ignored	returns 0	0
4.8:7	Pause	00 = No PAUSE 01 = Symmetric PAUSE 10 = Asymmetric PAUSE towards link partner 11 = Both Symmetric PAUSE and Asymmetric PAUSE towards link partner	read/write	11

Table 9-6: Auto-Negotiation Advertisement Register (Register 4) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
4.6	Half Duplex	Always returns a '0' for this bit since Half Duplex Mode is not supported	returns 0	0
4.5	Full Duplex	1 = Full Duplex Mode is advertised 0 = Full Duplex Mode is not advertised	read/write	1
4.4:0	Reserved	Always return 0s , writes ignored	returns 0s	00000

Register 5: Auto-Negotiation Link Partner Base

MDIO Register 5: Auto-Negotiation Link Partner Base




Table 9-7: Auto-Negotiation Link Partner Ability Base Register (Register 5)

Bit(s)	Name	Description	Attributes	Default Value
5.15	Next Page	1 = Next Page functionality is supported 0 = Next Page functionality is not supported	read only	0
5.14	Acknowledge	Used by Auto-Negotiation function to indicate reception of a link partner's base or next page	read only	0
5.13:12	Remote Fault	00 = No Error 01 = Offline 10 = Link Failure 11 = Auto-Negotiation Error	read only	00
5.11:9	Reserved	Always return 0s	returns 0s	000
5.8:7	Pause	00 = No PAUSE 01 = Symmetric PAUSE 10 = Asymmetric PAUSE towards link partner 11 = Both Symmetric PAUSE and Asymmetric PAUSE supported	read only	00

Table 9-7: Auto-Negotiation Link Partner Ability Base Register (Register 5)

Bit(s)	Name	Description	Attributes	Default Value
5.6	Half Duplex	1 = Half Duplex Mode is supported 0 = Half Duplex Mode is not supported	read only	0
5.5	Full Duplex	1 = Full Duplex Mode is supported 0 = Full Duplex Mode is not supported	read only	0
5.4:0	Reserved	Always return 0s	returns 0s	00000

Register 6: Auto-Negotiation Expansion

MDIO Register 6: Auto-Negotiation Expansion




Table 9-8: Auto-Negotiation Expansion Register (Register 6)

Bit(s)	Name	Description	Attributes	Default Value
6.15:3	Reserved	Always returns 0s	returns 0s	0000000000000
6.2	Next Page Able	Always returns a '1' for this bit since the device is Next Page Able	returns 1	1
6.1	Page Received	1 = A new page has been received 0 = A new page has not been received	read only self clearing on read	0
6.0	Reserved	Always returns 0s	returns 0s	0000000

Register 7: Next Page Transmit

MDIO Register 7: Next Page Transmit




Table 9-9: Auto-Negotiation Next Page Transmit (Register 7)

Bit(s)	Name	Description	Attributes	Default Value
7.15	Next Page	1 = Additional Next Page(s) will follow 0 = Last page	read/ write	0
7.14	Reserved	Always returns '0'	returns 0	0
7.13	Message Page	1 = Message Page 0 = Unformatted Page	read/ write	1
7.12	Acknowledge 2	1 = Comply with message 0 = Cannot comply with message	read/ write	0
7.11	Toggle	Value toggles between subsequent Next Pages	read only	0
7.10:0	Message / Unformatted Code Field	Message Code Field or Unformatted Page Encoding as dictated by 7.13	read/ write	0000000001 (Null Message Code)

Register 8: Next Page Receive

MDIO Register 8: Next Page Receive




Table 9-10: Auto-Negotiation Next Page Receive (Register 8)

Bit(s)	Name	Description	Attributes	Default Value
8.15	Next Page	1 = Additional Next Page(s) will follow 0 = Last page	read only	0
8.14	Acknowledge	Used by Auto-Negotiation function to indicate reception of a link partner's base or next page	read only	0
8.13	Message Page	1 = Message Page 0 = Unformatted Page	read only	0

Table 9-10: Auto-Negotiation Next Page Receive (Register 8) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
8.12	Acknowledge 2	1 = Comply with message 0 = Cannot comply with message	read only	0
8.11	Toggle	Value toggles between subsequent Next Pages	read only	0
8.10:0	Message / Unformatted Code Field	Message Code Field or Unformatted Page Encoding as dictated by 8.13	read only	00000000000

Register 15: Extended Status

MDIO Register 15: Extended Status Register




Table 9-11: Extended Status Register (Register 15)

Bit(s)	Name	Description	Attributes	Default Value
15.15	1000BASE-X Full Duplex	Always returns a '1' for this bit since 1000BASE-X Full Duplex is supported	returns 1	1
15.14	1000BASE-X Half Duplex	Always returns a '0' for this bit since 1000BASE-X Half Duplex is not supported	returns 0	0
15.13	1000BASE-T Full Duplex	Always returns a '0' for this bit since 1000BASE-T Full Duplex is not supported	returns 0	0
15.12	1000BASE-T Half Duplex	Always returns a '0' for this bit since 1000BASE-T Half Duplex is not supported	returns 0	0
15:11:0	Reserved	Always return 0s	returns 0s	00000000000

Register 16: Vendor-Specific Auto-Negotiation Interrupt Control

MDIO Register 16: Vendor Specific Auto-Negotiation Interrupt Control




Table 9-12: Vendor Specific Register: Auto-Negotiation Interrupt Control Register (Register 16)

Bit(s)	Name	Description	Attributes	Default Value
16.15:2	Reserved	Always return 0s	returns 0s	00000000000000
16.1	Interrupt Status	1 = Interrupt is asserted 0 = Interrupt is not asserted If the interrupt is enabled, this bit is asserted on the completion of an Auto-Negotiation cycle; it is only cleared by writing '0' to this bit. If the Interrupt is disabled, the bit is set to '0.' NOTE: the an_interrupt port of the core is wired to this bit.	read/ write	0
16.0	Interrupt Enable	1 = Interrupt enabled 0 = Interrupt disabled	read/ write	1

1000BASE-X Standard Without the Optional Auto-Negotiation

It is not the intention of this document to fully describe the 1000BASE-X PCS Registers. See clauses 37 and 22 of the *IEEE 802.3* Specification for further information.

Registers at undefined addresses are read-only and return 0s.

Table 9-13: MDIO Registers for 1000BASE-X without Auto-Negotiation

Register Address	Register Name
0	Control Register
1	Status Register
2,3	PHY Identifier
15	Extended Status Register

Register 0: Control Register

MDIO Register 0: Control Register

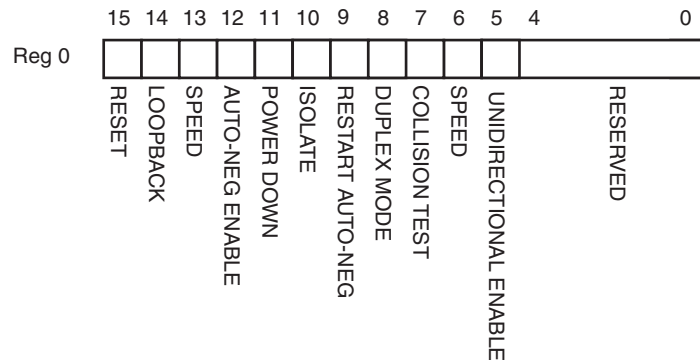


Table 9-14: Control Register (Register 0)

Bit(s)	Name	Description	Attributes	Default Value
0.15	Reset	1 = PCS/PMA reset 0 = Normal Operation	read/write self clearing	0
0.14	Loopback	1 = Enable Loopback Mode 0 = Disable Loopback Mode When used with a RocketIO transceiver, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to ewrap. When set to '1' indicates to the external PMA module to enter loopback mode. See "Loopback," page 197.	read/write	0
0.13	Speed Selection (LSB)	Always returns a 0 for this bit. Together with bit 0.6, speed selection of 1000 Mbps is identified.	returns 0	0
0.12	Auto-Negotiation Enable	Ignore this bit because Auto-Negotiation is not included.	read/ write	1
0.11	Power Down	1 = Power down 0 = Normal operation With the PMA option, when set to '1' the RocketIO transceiver is placed in a low-power state. This bit requires a reset (see bit 0.15) to clear. With the TBI version this register bit has no effect.	read/ write	0
0.10	Isolate	1 = Electrically Isolate PHY from GMII 0 = Normal operation	read/write	1

Table 9-14: Control Register (Register 0) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
0.9	Restart Auto-Negotiation	Ignore this bit because Auto-Negotiation is not included.	read/ write	0
0.8	Duplex Mode	Always returns a '1' for this bit to signal Full-Duplex Mode.	returns 1	1
0.7	Collision Test	Always returns a '0' for this bit to disable COL test.	returns 0	0
0.6	Speed Selection(MSB)	Always returns a '1' for this bit. Together with bit 0.13, speed selection of 1000 Mbps is identified	returns 1	1
0.5	Unidirectional Enable	Ignore this bit because Auto-Negotiation is not included.	read/ write	0
0.4:0	Reserved	Always return 0s , writes ignored.	returns 0s	00000

Register 1: Status Register

MDIO Register 1: Status Register

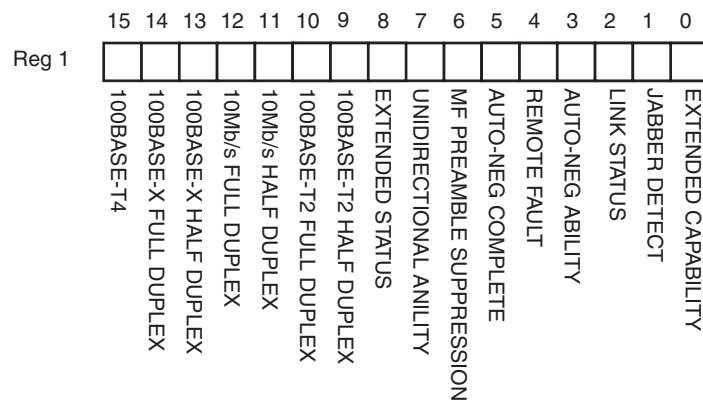


Table 9-15: Status Register (Register 1)

Bit(s)	Name	Description	Attributes	Default Value
1.15	100BASE-T4	Always returns a '0' for this bit since 100BASE-T4 is not supported	returns 0	0
1.14	100BASE-X Full Duplex	Always returns a '0' for this bit since 100BASE-X Full Duplex is not supported	returns 0	0
1.13	100BASE-X Half Duplex	Always returns a '0' for this bit since 100BASE-X Half Duplex is not supported	returns 0	0
1.12	10 Mbps Full Duplex	Always returns a '0' for this bit since 10 Mbps Full Duplex is not supported	returns 0	0
1.11	10 Mbps Half Duplex	Always returns a '0' for this bit since 10 Mbps Half Duplex is not supported	returns 0	0

Table 9-15: Status Register (Register 1) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
1.10	100BASE-T2 Full Duplex	Always returns a '0' for this bit since 100BASE-T2 Full Duplex is not supported	returns 0	0
1.9	100BASE-T2 Half Duplex	Always returns a '0' for this bit since 100BASE-T2 Half Duplex is not supported	returns 0	0
1.8	Extended Status	Always returns a '1' for this bit to indicate the presence of the Extended Register (Register 15)	returns 1	1
1.7	Unidirectional Ability	Always returns 1, writes ignored	returns 1	1
1.6	MF Preamble Suppression	Always returns a '1' for this bit to indicate that Management Frame Preamble Suppression is supported	returns 1	1
1.5	Auto-Negotiation Complete	Ignore this bit because Auto-Negotiation is not included.	returns 1	1
1.4	Remote Fault	Always returns a '0' for this bit because Auto-Negotiation is not included.	returns 0	0
1.3	Auto-Negotiation Ability	Ignore this bit because Auto-Negotiation is not included.	returns 0	0
1.2	Link Status ¹	1 = Link is up 0 = Link is down Latches '0' if Link Status goes down. Clears to current Link Status on read. See table note for Link Status behavior.	read only self clearing on read	0
1.1	Jabber Detect	Always returns a '0' for this bit since Jabber Detect is not supported	returns 0	0
1.0	Extended Capability	Always returns a '0' for this bit since no extended register set is supported	returns 0	0

1. When high, the link is valid: synchronization of the link has been obtained and Auto-Negotiation (if present and enabled) has completed.
When low, a valid link has not been established. Either link synchronization has failed or Auto-Negotiation (if present and enabled) has failed to complete.
Regardless of whether Auto-Negotiation is enabled or disabled, there can be some delay to the deassertion of this signal following the loss of synchronization of a previously successful link. This is due to the Auto-Negotiation state machine which requires that synchronization is lost for an entire link timer duration before changing state. For more information, see the 802.3 specification.

Registers 2 and 3: Phy Identifier

MDIO Registers 2 and 3: PHY Identifier




Table 9-16: PHY Identifier (Registers 2 and 3)

Bit(s)	Name	Description	Attributes	Default Value
2.15:0	Organizationally Unique Identifier	Always return 0s	returns 0s	0000000000000000
3.15:10	Organizationally Unique Identifier	Always return 0s	returns 0s	000000
3.9:4	Manufacturer's model number	Always return 0s	returns 0s	000000
3.3:0	Revision Number	Always return 0s	returns 0s	0000

Register 15: Extended Status

MDIO Register 15: Extended Status




Table 9-17: Extended Status (Register 15)

Bit(s)	Name	Description	Attributes	Default Value
15.15	1000BASE-X Full Duplex	Always returns a '1' since 1000BASE-X Full Duplex is supported	returns 1	1
15.14	1000BASE-X Half Duplex	Always returns a '0' since 1000BASE-X Half Duplex is not supported	returns 0	0
15.13	1000BASE-T Full Duplex	Always returns a '0' since 1000BASE-T Full Duplex is not supported	returns 0	0
15.12	1000BASE-T Half Duplex	Always returns a '0' since 1000BASE-T Half Duplex is not supported	returns 0	0
15:11:0	Reserved	Always return 0s	returns 0s	0000000000 00

SGMII Standard Using the Optional Auto-Negotiation

The registers provided for SGMII operation in this core are adaptations of those defined in *IEEE 802.3* clauses 37 and 22. In an SGMII implementation, two different types of links exist. They are the SGMII link between the MAC and PHY (SGMII link) and the link across the Ethernet Medium itself (Medium). See [Figure 10-2](#).

Information regarding the state of both of these links is contained within the following registers. Where applicable, the abbreviations *SGMII link* and *Medium* are used in the register descriptions. Registers at undefined addresses are read-only and return 0s.

Table 9-18: MDIO Registers for 1000BASE-X with Auto-Negotiation

Register Address	Register Name
0	SGMII Control Register
1	SGMII Status Register
2,3	PHY Identifier
4	SGMII Auto-Negotiation Advertisement Register
5	SGMII Auto-Negotiation Link Partner Ability Base Register
6	SGMII Auto-Negotiation Expansion Register
7	SGMII Auto-Negotiation Next Page Transmit Register
8	SGMII Auto-Negotiation Next Page Receive Register
15	SGMII Extended Status Register
16	SGMII Vendor Specific: Auto-Negotiation Interrupt Control

Register 0: SGMII Control

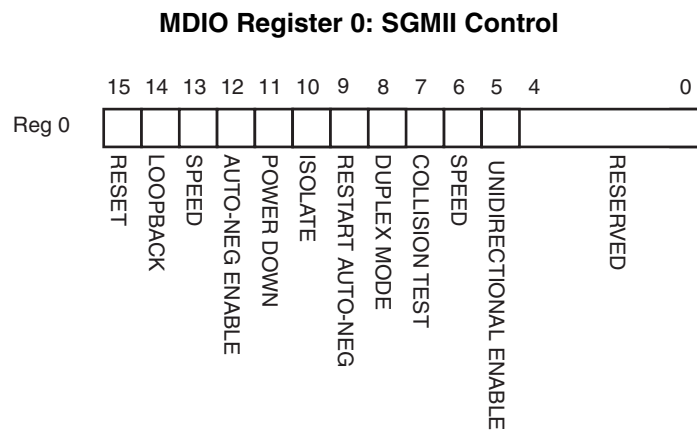


Table 9-19: SGMII Control (Register 0)

Bit(s)	Name	Description	Attributes	Default Value
0.15	Reset	1 = Core Reset 0 = Normal Operation	read/write self clearing	0
0.14	Loopback	1 = Enable Loopback Mode 0 = Disable Loopback Mode When used with a RocketIO transceiver, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to <i>ewrap</i> . When set to '1' indicates to the external PMA module to enter loopback mode. See "Loopback," page 197.	read/write	0
0.13	Speed Selection (LSB)	Always returns a '0' for this bit. Together with bit 0.6, speed selection of 1000 Mbps is identified	returns 0	0
0.12	Auto-Negotiation Enable	1 = Enable SGMII Auto-Negotiation Process 0 = Disable SGMII Auto-Negotiation Process	read/write	1
0.11	Power Down	1 = Power down 0 = Normal operation With the PMA option, when set to '1' the RocketIO transceiver is placed in a low-power state. This bit requires a reset (see bit 0.15) to clear. With the TBI version this register bit has no effect.	read/ write	0
0.10	Isolate	1 = Electrically Isolate SGMII logic from GMII 0 = Normal operation	read/write	1
0.9	Restart Auto-Negotiation	1 = Restart Auto-Negotiation Process across SGMII link 0 = Normal Operation	read/write self clearing	0
0.8	Duplex Mode	Always returns a '1' for this bit to signal Full-Duplex Mode	returns 1	1
0.7	Collision Test	Always returns a '0' for this bit to disable COL test	returns 0	0
0.6	Speed Selection (MSB)	Always returns a '1' for this bit. Together with bit 0.13, speed selection of 1000 Mbps is identified	returns 1	1

Table 9-19: SGMII Control (Register 0) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
0.5	Unidirectional Enable	Enable transmit regardless of whether a valid link has been established	read/ write	0
0.4:0	Reserved	Always return 0s , writes ignored	returns 0s	00000

Register 1: SGMII Status

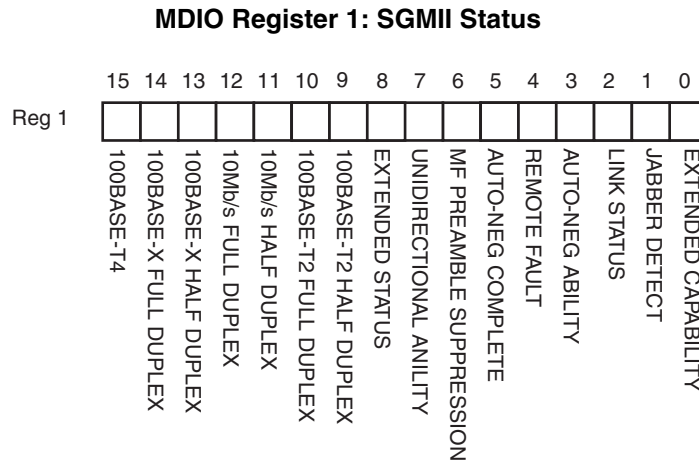


Table 9-20: SGMII Status (Register 1)

Bit(s)	Name	Description	Attributes	Default Value
1.15	100BASE-T4	Always returns a '0' for this bit because 100BASE-T4 is not supported	returns 0	0
1.14	100BASE-X Full Duplex	Always returns a '0' for this bit because 100BASE-X Full Duplex is not supported	returns 0	0
1.13	100BASE-X Half Duplex	Always returns a '0' for this bit because 100BASE-X Half Duplex is not supported	returns 0	0
1.12	10 Mbps Full Duplex	Always returns a '0' for this bit because 10 Mbps Full Duplex is not supported	returns 0	0
1.11	10 Mbps Half Duplex	Always returns a '0' for this bit because 10 Mbps Half Duplex is not supported	returns 0	0
1.10	100BASE-T2 Full Duplex	Always returns a '0' for this bit because 100BASE-T2 Full Duplex is not supported	returns 0	0
1.9	100BASE-T2 Half Duplex	Always returns a '0' for this bit because 100BASE-T2 Half Duplex is not supported	returns 0	0
1.8	Extended Status	Always returns a '1' for this bit to indicate the presence of the Extended Register (Register 15)	returns 1	1

Table 9-20: SGMII Status (Register 1) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
1.7	Unidirectional Ability	Always returns '1,' writes ignored	returns 1	1
1.6	MF Preamble Suppression	Always returns a '1' for this bit to indicate that Management Frame Preamble Suppression is supported	returns 1	1
1.5	Auto- Negotiation Complete	1 = Auto-Negotiation process completed across SGMII link 0 = Auto-Negotiation process not completed across SGMII link	read only	0
1.4	Remote Fault	1 = A fault on the Medium has been detected 0 = No fault of the Medium has been detected	read only self clearing on read	0
1.3	Auto- Negotiation Ability	Always returns a '1' for this bit to indicate that the SGMII core is capable of Auto-Negotiation	returns 1	1
1.2	SGMII Link Status ¹	1 = SGMII Link is up 0 = SGMII Link is down Latches '0' if SGMII Link Status goes down. Clears to current SGMII Link Status on read. See table note for SGMII Link Status behavior.	read only self clearing on read	0
1.1	Jabber Detect	Always returns a '0' for this bit since Jabber Detect is not supported	returns 0	0
1.0	Extended Capability	Always returns a '0' for this bit because no extended register set is supported	returns 0	0

1. When high, the link is valid: synchronization of the SGMII link has been obtained and SGMII Auto-Negotiation (if present and enabled) has completed.
When low, a valid link has not been established. Either SGMII link synchronization has failed or SGMII Auto-Negotiation (if present and enabled) has failed to complete.
Regardless of whether SGMII Auto-Negotiation is enabled or disabled, there can be some delay to the deassertion of this signal following the loss of synchronization of a previously successful SGMII link. This is due to the Auto-Negotiation state machine which requires that synchronization is lost for an entire link timer duration before changing state. For more information, see the 802.3 specification.

Registers 2 and 3: PHY Identifier

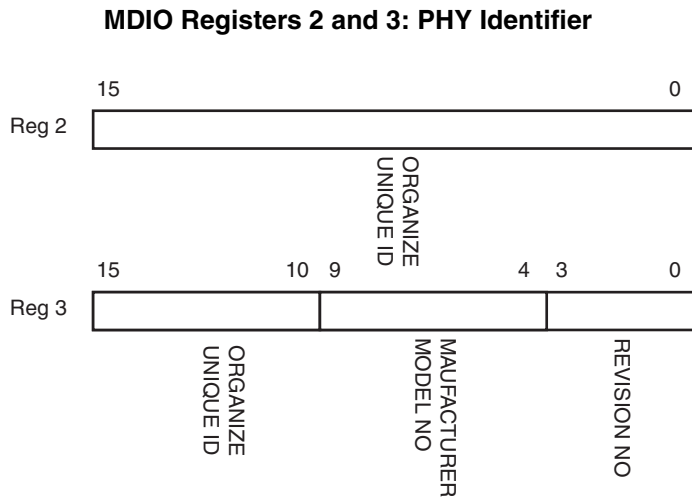


Table 9-21: PHY Identifier (Registers 2 and 3)

Bit(s)	Name	Description	Attributes	Default Value
2.15:0	Organizationally Unique Identifier	Always return 0s	returns 0s	0000000000000000
3.15:10	Organizationally Unique Identifier	Always return 0s	returns 0s	000000
3.9:4	Manufacturer's model number	Always return 0s	returns 0s	000000
3.3:0	Revision Number	Always return 0s	returns 0s	0000

Register 4: SGMII Auto-Negotiation Advertisement

MDIO Register 4: SGMII Auto-Negotiation Advertisement





Table 9-22: SGMII Auto-Negotiation Advertisement (Register 4)

Bit(s)	Name	Description	Attributes	Default Value
4.15:0	All bits	SGMII defined value sent from the MAC to the PHY	read only	0000000000000001

Register 5: SGMII Auto-Negotiation Link Partner Ability

MDIO Register 5: SGMII Auto-Negotiation Link Partner Ability



The Auto-Negotiation Ability Base Register (Register 5) contains information related to the status of the link between the PHY and its physical link partner across the Medium.

Table 9-23: SGMII Auto-Negotiation Link Partner Ability Base (Register 5)

Bit(s)	Name	Description	Attributes	Default Value
5.15	PHY Link Status	This refers to the link status of the PHY with its link partner across the Medium. 1 = Link Up 0 = Link Down	read only	1
5.14	Acknowledge	Used by Auto-Negotiation function to indicate reception of a link partner's base or next page	read only	0
5.13	Reserved	Always returns '0,' writes ignored	returns 0	0
5.12	Duplex Mode	1= Full Duplex 0 = Half Duplex	read only	0
5.11:10	Speed	11 = Reserved 10 = 1 Gbps 01 = 100 Mbps 00 = 10 Mbps	read only	00
5.9:1	Reserved	Always return 0s	returns 0s	000000000
5:0	Reserved	Always returns '1'	returns 1	1

Register 6: SGMII Auto-Negotiation Expansion

MDIO Register 6: SGMII Auto-Negotiation Expansion




Table 9-24: SGMII Auto-Negotiation Expansion (Register 6)

Bit(s)	Name	Description	Attributes	Default Value
6.15:3	Reserved	Always return 0s	returns 0s	0000000000000
6.2	Next Page Able	Always returns a '1' for this bit since the device is Next Page Able	returns 1	1
6.1	Page Received	1 = A new page has been received 0 = A new page has not been received	read only self clearing on read	0
6.0	Reserved	Always return 0s	returns 0s	0000000

Register 7: SGMII Auto-Negotiation Next Page Transmit

MDIO Register 7: SGMII Auto-Negotiation Next Page Transmit




Table 9-25: SGMII Auto-Negotiation Next Page Transmit (Register 7)

Bit(s)	Name	Description	Attributes	Default Value
7.15	Next Page	1 = Additional Next Page(s) will follow 0 = Last page	read/ write	0
7.14	Reserved	Always returns '0'	returns 0	0
7.13	Message Page	1 = Message Page 0 = Unformatted Page	read/ write	1

Table 9-25: SGMII Auto-Negotiation Next Page Transmit (Register 7)

Bit(s)	Name	Description	Attributes	Default Value
7.12	Acknowledge 2	1 = Comply with message 0 = Cannot comply with message	read/ write	0
7.11	Toggle	Value toggles between subsequent Next Pages	read only	0
7.10:0	Message / Unformatted Code Field	Message Code Field or Unformatted Page Encoding as dictated by 7.13	read/ write	0000000001 (Null Message Code)

Register 8: SGMII Next Page Receive

MDIO Register 8: SGMII Next Page Receive




Table 9-26: SGMII Auto-Negotiation Next Page Receive (Register 8)

Bit(s)	Name	Description	Attributes	Default Value
8.15	Next Page	1 = Additional Next Page(s) will follow 0 = Last page	read only	0
8.14	Acknowledge	Used by Auto-Negotiation function to indicate reception of a link partner's base or next page	read only	0
8.13	Message Page	1 = Message Page 0 = Unformatted Page	read only	0
8.12	Acknowledge 2	1 = Comply with message 0 = Cannot comply with message	read only	0
8.11	Toggle	Value toggles between subsequent Next Pages	read only	0
8.10:0	Message / Unformatted Code Field	Message Code Field or Unformatted Page Encoding as dictated by 8.13	read only	0000000000

Register 15: SGMII Extended Status

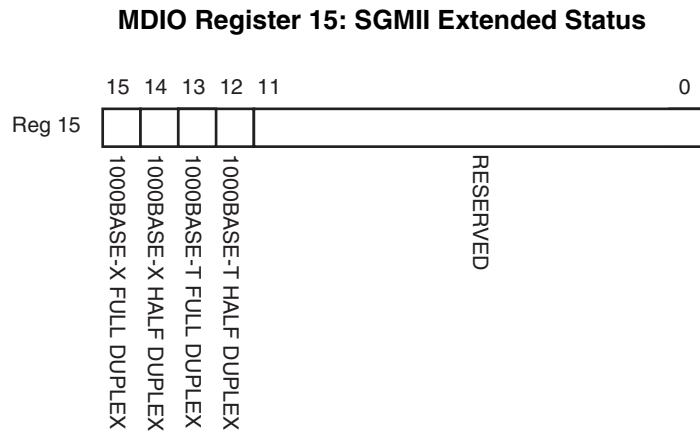


Table 9-27: SGMII Extended Status Register (Register 15)

Bit(s)	Name	Description	Attributes	Default Value
15.15	1000BASE-X Full Duplex	Always returns a '1' for this bit since 1000BASE-X Full Duplex is supported	returns 1	1
15.14	1000BASE-X Half Duplex	Always returns a '0' for this bit since 1000BASE-X Half Duplex is not supported	returns 0	0
15.13	1000BASE-T Full Duplex	Always returns a '0' for this bit since 1000BASE-T Full Duplex is not supported	returns 0	0
15.12	1000BASE-T Half Duplex	Always returns a '0' for this bit since 1000BASE-T Half Duplex is not supported	returns 0	0
15:11:0	Reserved	Always return 0s	returns 0s	000000000000

Register 16: SGMII Auto-Negotiation Interrupt Control

MDIO Register 16: SGMII Auto-Negotiation Interrupt Control




Table 9-28: SGMII Auto-Negotiation Interrupt Control (Register 16)

Bit(s)	Name	Description	Attributes	Default Value
16.15:2	Reserved	Always return 0s	returns 0s	00000000000000
16.1	Interrupt Status	1 = Interrupt is asserted 0 = Interrupt is not asserted If the interrupt is enabled, this bit is asserted on completion of an Auto-Negotiation cycle across the SGMII link; it is only cleared by writing '0' to this bit. If the Interrupt is disabled, the bit is set to '0.' NOTE: The an_interrupt port of the core is wired to this bit.	read/ write	0
16.0	Interrupt Enable	1 = Interrupt enabled 0 = Interrupt disabled	read/ write	1

Table 9-30: SGMII Control (Register 0)

Bit(s)	Name	Description	Attributes	Default Value
0.15	Reset	1 = Core Reset 0 = Normal Operation	read/write self clearing	0
0.14	Loopback	1 = Enable Loopback Mode 0 = Disable Loopback Mode When used with a RocketIO transceiver, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to <i>ewrap</i> . When set to '1' indicates to the external PMA module to enter loopback mode. See "Loopback," page 197.	read/write	0
0.13	Speed Selection (LSB)	Always returns a '0' for this bit. Together with bit 0.6, speed selection of 1000 Mbps is identified	returns 0	0
0.12	Auto-Negotiation Enable	1 = Enable SGMII Auto-Negotiation Process 0 = Disable SGMII Auto-Negotiation Process	read/write	1
0.11	Power Down	1 = Power down 0 = Normal operation With the PMA option, when set to '1' the RocketIO transceiver is placed in a low-power state. This bit requires a reset (see bit 0.15) to clear. With the TBI version this register bit has no effect.	read/ write	0
0.10	Isolate	1 = Electrically Isolate SGMII logic from GMII 0 = Normal operation	read/write	1
0.9	Restart Auto-Negotiation	1 = Restart Auto-Negotiation Process across SGMII link 0 = Normal Operation	read/write self clearing	0
0.8	Duplex Mode	Always returns a '1' for this bit to signal Full-Duplex Mode	returns 1	1
0.7	Collision Test	Always returns a '0' for this bit to disable COL test	returns 0	0
0.6	Speed Selection (MSB)	Always returns a '1' for this bit. Together with bit 0.13, speed selection of 1000 Mbps is identified	returns 1	1

Table 9-30: SGMII Control (Register 0) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
0.5	Unidirectional Enable	Enable transmit regardless of whether a valid link has been established	read/ write	0
0.4:0	Reserved	Always return 0s , writes ignored	returns 0s	00000

Register 1: SGMII Status




Table 9-31: SGMII Status (Register 1)

Bit(s)	Name	Description	Attributes	Default Value
1.15	100BASE-T4	Always returns a '0' for this bit because 100BASE-T4 is not supported	returns 0	0
1.14	100BASE-X Full Duplex	Always returns a '0' for this bit because 100BASE-X Full Duplex is not supported	returns 0	0
1.13	100BASE-X Half Duplex	Always returns a '0' for this bit because 100BASE-X Half Duplex is not supported	returns 0	0
1.12	10 Mbps Full Duplex	Always returns a '0' for this bit because 10 Mbps Full Duplex is not supported	returns 0	0
1.11	10 Mbps Half Duplex	Always returns a '0' for this bit because 10 Mbps Half Duplex is not supported	returns 0	0
1.10	100BASE-T2 Full Duplex	Always returns a '0' for this bit because 100BASE-T2 Full Duplex is not supported	returns 0	0
1.9	100BASE-T2 Half Duplex	Always returns a '0' for this bit because 100BASE-T2 Half Duplex is not supported	returns 0	0
1.8	Extended Status	Always returns a '1' for this bit to indicate the presence of the Extended Register (Register 15)	returns 1	1

Table 9-31: SGMII Status (Register 1) (Continued)

Bit(s)	Name	Description	Attributes	Default Value
1.7	Unidirectional Ability	Always returns '1,' writes ignored	returns 1	1
1.6	MF Preamble Suppression	Always returns a '1' for this bit to indicate that Management Frame Preamble Suppression is supported	returns 1	1
1.5	Auto- Negotiation Complete	Ignore this bit because Auto-Negotiation is not included.	returns 1	0
1.4	Remote Fault	Ignore this bit because Auto-Negotiation is not included	returns 0	0
1.3	Auto- Negotiation Ability	Ignore this bit because Auto-Negotiation is not included	returns 0	0
1.2	SGMII Link Status ¹	1 = SGMII Link is up 0 = SGMII Link is down Latches '0' if SGMII Link Status goes down. Clears to current SGMII Link Status on read. See table note for SGMII Link Status behavior.	read only self clearing on read	0
1.1	Jabber Detect	Always returns a '0' for this bit since Jabber Detect is not supported	returns 0	0
1.0	Extended Capability	Always returns a '0' for this bit because no extended register set is supported	returns 0	0

- When high, the link is valid: synchronization of the SGMII link has been obtained and SGMII Auto-Negotiation (if present and enabled) has completed.

When low, a valid link has not been established. Either SGMII link synchronization has failed or SGMII Auto-Negotiation (if present and enabled) has failed to complete.

Regardless of whether SGMII Auto-Negotiation is enabled or disabled, there can be some delay to the deassertion of this signal following the loss of synchronization of a previously successful SGMII link. This is due to the Auto-Negotiation state machine which requires that synchronization is lost for an entire link timer duration before changing state. For more information, see the 802.3 specification.

Registers 2 and 3: PHY Identifier




Table 9-32: PHY Identifier (Registers 2 and 3)

Bit(s)	Name	Description	Attributes	Default Value
2.15:0	Organizationally Unique Identifier	Always return 0s	returns 0s	0000000000000000
3.15:10	Organizationally Unique Identifier	Always return 0s	returns 0s	000000
3.9:4	Manufacturer’s model number	Always return 0s	returns 0s	000000
3.3:0	Revision Number	Always return 0s	returns 0s	0000

Register 4: SGMII Auto-Negotiation Advertisement

MDIO Register 4: SGMII Auto-Negotiation Advertisement




Table 9-33: SGMII Auto-Negotiation Advertisement (Register 4)

Bit(s)	Name	Description	Attributes	Default Value
4.15:0	All bits	Ignore this register because Auto-Negotiation is not included	read only	0000000000000001

Register 15: SGMII Extended Status

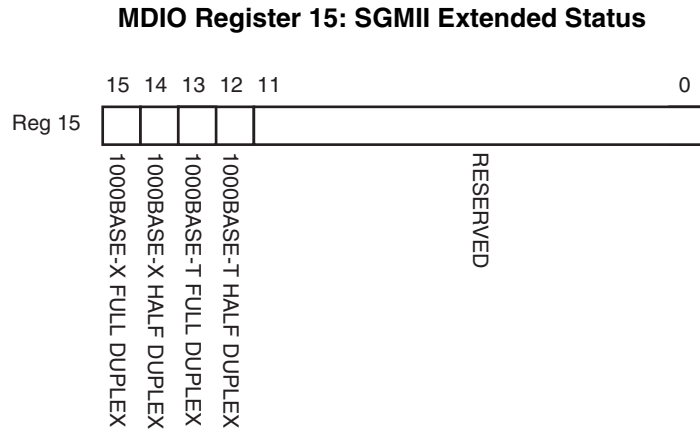


Table 9-34: SGMII Extended Status Register (Register 15)

Bit(s)	Name	Description	Attributes	Default Value
15.15	1000BASE-X Full Duplex	Always returns a '1' for this bit since 1000BASE-X Full Duplex is supported	returns 1	1
15.14	1000BASE-X Half Duplex	Always returns a '0' for this bit since 1000BASE-X Half Duplex is not supported	returns 0	0
15.13	1000BASE-T Full Duplex	Always returns a '0' for this bit since 1000BASE-T Full Duplex is not supported	returns 0	0
15.12	1000BASE-T Half Duplex	Always returns a '0' for this bit since 1000BASE-T Half Duplex is not supported	returns 0	0
15:11:0	Reserved	Always return 0s	returns 0s	000000000000

Both 1000BASE-X and SGMII Standards

Table 9-35 describes register 17, the vendor-specific Standard Selection Register. This register is only present when the core is generated with the capability to dynamically switch between 1000BASE-X and SGMII standards. See “Select Standard” in Chapter 3.

When this Register is configured to perform the 1000BASE-X standard, Registers 0 to 16 should be interpreted as per “1000BASE-X Standard Using the Optional Auto-Negotiation,” or “1000BASE-X Standard Without the Optional Auto-Negotiation.”

When this Register is configured to perform the SGMII standard, Registers 0 to 16 should be interpreted as per “SGMII Standard Using the Optional Auto-Negotiation,” or “SGMII Standard without the Optional Auto-Negotiation.” This register may be written to at any time. See Chapter 11, “Dynamic Switching of 1000BASE-X and SGMII Standards” for more information.

Register 17: Vendor-specific Standard Selection Register




Figure 9-5: Dynamic Switching (Register 17)

Table 9-35: Vendor-specific Register: Standard Selection Register (Register 17)

Bit(s)	Name	Description	Attributes	Default Value
17.15:1	Reserved	Always return 0s	Returns 0s	0000000000000000
16.0	Standard	0 = Core will perform the 1000BASE-X standard. Registers 0 to 16 will behave as per "1000BASE-X Standard Using the Optional Auto-Negotiation" 1 = Core will perform the SGMII standard. Registers 0 to 16 will behave as per "SGMII Standard Using the Optional Auto-Negotiation" .	read/write	Determined by the basex_or_sgmii port

Optional Configuration Vector

If ["MDIO Management Interface"](#) is omitted, relevant configuration signals are brought out of the core. These signals are bundled into the CONFIGURATION_VECTOR signal as defined in [Table 9-36](#).

These signals may be changed by the user application at any time. The *Clock Domain* heading denotes the clock domain the configuration signal is registered in before use by the core. It is not necessary to drive the signal from this clock domain.

Table 9-36: Optional Configuration and Status Vectors

Signal	Direction	Clock Domain	Description
configuration_vector [3:0]	Input	See note 1	Bit[0]: Reserved (currently unused) Bit[1]: Loopback Control <ul style="list-style-type: none"> When used with a RocketIO transceiver, the core is placed in internal loopback mode. With the TBI version, Bit 1 is connected to ewrap. When set to '1,' this indicates to the external PMA module to enter loopback mode. See "Loopback," page 197. Bit[2]: Power Down <ul style="list-style-type: none"> When a RocketIO transceiver is used, a setting of '1' places the RocketIO in a low-power state. A reset must be applied to clear. With the TBI version, this bit is unused. Bit[3]: Isolate <ul style="list-style-type: none"> When set to '1,' the GMII should be electrically isolated. When set to '0,' normal operation is enabled.

1. Signals are synchronous to the core's internal 125 MHz reference clock; this is `userclk2` when used with a RocketIO transceiver; `gtx_clk` when used with TBI.

Auto-Negotiation

This chapter provides general guidelines for using the Auto-Negotiation function of the Ethernet 1000BASE-X PCS/PMA or SGMII core. Auto-Negotiation is controlled and monitored through the PCS Management Registers and is only available when the optional MDIO Management Interface is present. For more information, see [Chapter 9, "Configuration and Status."](#)

Overview of Operation

For either standard, when considering Auto-Negotiation between two connected devices, it must be remembered that:

- Auto-Negotiation must be either enabled in **both** devices, or:
- Auto-Negotiation must be disabled in **both** devices.

1000BASE-X Standard

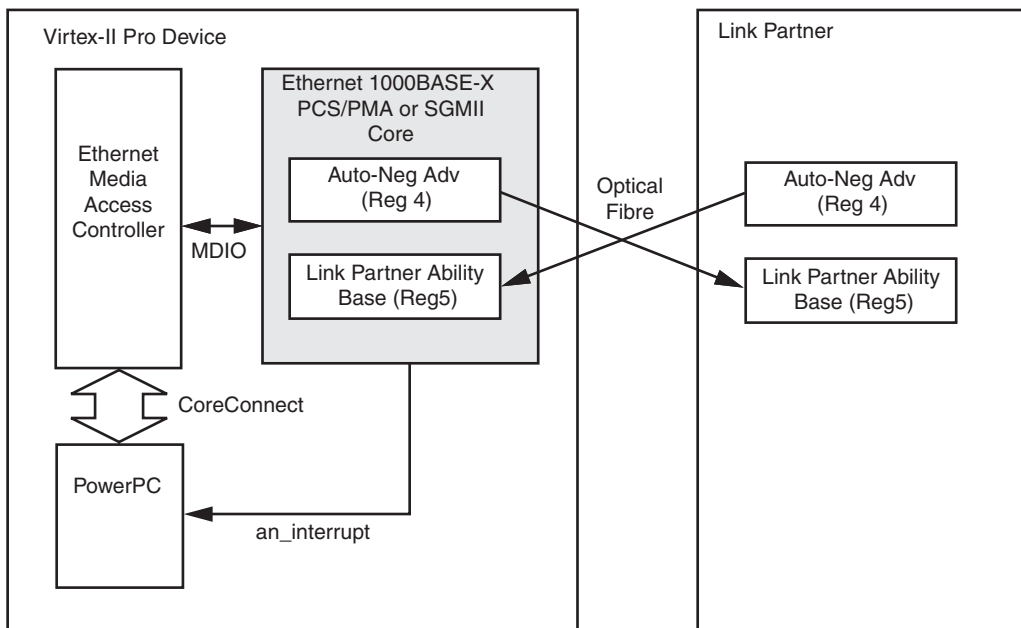


Figure 10-1: 1000BASE-X Auto-Negotiation Overview

IEEE 802.3 clause 37 describes the 1000BASE-X Auto-Negotiation function that allows a device to advertise the modes of operation that it supports to a device at the remote end of

a link segment (the link partner) and to detect corresponding operational modes that the link partner advertises. [Figure 10-1](#) illustrates the operation of 1000BASE-X Auto-Negotiation.

The following describes typical operation when Auto-Negotiation is enabled.

1. Auto-Negotiation starts automatically when any of the following conditions are met.
 - ◆ Power-up/reset
 - ◆ Upon loss of synchronization
 - ◆ The link partner initiates Auto-Negotiation
 - ◆ An Auto-Negotiation Restart is requested (See [“Control Register \(Register 0\),”](#) [page 120.](#))

2. During Auto-Negotiation, the contents of the Auto-Negotiation Advertisement Register are transferred to the link partner.

This register is writable through the MDIO, therefore enabling software control of the systems advertised abilities. See [“Auto-Negotiation Advertisement Register \(Register 4\),”](#) [page 124](#) for more information.

Information provided in this register includes:

- ◆ Fault Condition signaling
 - ◆ Duplex Mode
 - ◆ Flow Control capabilities for the attached MAC.
3. The advertised abilities of the Link Partner are simultaneously transferred into the Auto-Negotiation Link Partner Ability Base Register.

This register contains the same information as in the Auto-Negotiation Advertisement Register. See [“Auto-Negotiation Link Partner Ability Base Register \(Register 5\),”](#) [page 125](#) for more information.
 4. Under normal conditions, this completes the Auto-Negotiation information exchange.

It is now the responsibility of system management (for example, software running on an embedded PowerPC™ or MicroBlaze™) to complete the cycle. The results of the Auto-Negotiation should be read from Auto-Negotiation Link Partner Ability Base Register. Other networking components, such as an attached Ethernet MAC, should be configured accordingly. See [“Auto-Negotiation Link Partner Ability Base Register \(Register 5\)”](#) for more information.

There are two methods that a host processor uses to learn of the completion of an Auto-Negotiation cycle:

- ◆ Polling the Auto-Negotiation completion bit 1.5 in the Status Register (Register 1).
- ◆ Using the Auto-Negotiation interrupt port of the core (see [“Using the Auto-Negotiation Interrupt,”](#) [page 156.](#))

SGMII Standard

Figure 10-2 illustrates the operation of SGMII Auto-Negotiation. Additional information about SGMII Standard Auto-Negotiation is provided in the following sections.

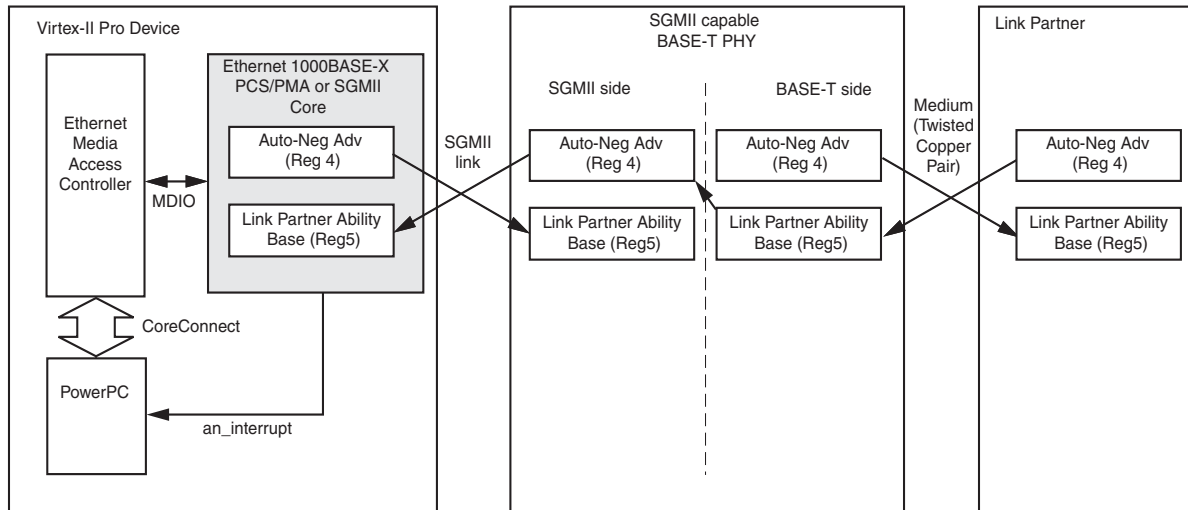


Figure 10-2: SGMII Auto-Negotiation

The SGMII capable PHY has two distinctive sides to Auto-Negotiation.

- The PHY performs Auto-Negotiation with its link partner using the relevant Auto-Negotiation standard for the chosen medium (BASE-T Auto-Negotiation is illustrated in Figure 10-2, using a twisted copper pair as its medium). This resolves the operational speed and duplex mode with the link partner.
- The PHY then passes the results of the Auto-Negotiation process with the link partner to the Ethernet 1000BASE-X PCS/PMA or SGMII core (in SGMII mode), by leveraging the 1000BASE-X Auto-Negotiation specification described in “1000BASE-X Auto-Negotiation Overview,” page 153. This transfers the results of the Link Partner Auto-Negotiation across the SGMII and is the only Auto-Negotiation observed by the core.

This SGMII Auto-Negotiation function, summarized previously, leverages the 1000BASE-X PCS/PMA Auto-Negotiation function but contains two differences.

- The duration of the Link Timer of the SGMII Auto-Negotiation is shrunk from 10 ms to 1.6 ms so that the entire Auto-Negotiation cycle is much faster. See “Setting the Configurable Link Timer,” page 156.
- The information exchanged is different and now contains speed resolution in addition to duplex mode. See “MDIO Register 5: SGMII Auto-Negotiation Link Partner Ability,” page 140.

There are no other differences and dealing with the results of Auto-Negotiation can be handled as described previously in “1000BASE-X Auto-Negotiation Overview.”

Setting the Configurable Link Timer

The optional Auto-Negotiation function has a Link Timer (`link_timer[8:0]`) port. This port sets the period of the Auto-Negotiation Link Timer. This port should be permanently tied to a logical binary value, and a binary value should be placed on this port. The duration of the timer is approximately equal to the binary value multiplied by 32.768 microseconds (4,96 clock periods of the 125 MHz clock provided to the core). See “[Auto-Negotiation Signal Pinout](#)” in Chapter 2.

Note: See Chapter 11, “[Dynamic Switching of 1000BASE-X and SGMII Standards](#)” for details of programming the Auto-Negotiation Link Timer when performing dynamic switching between 1000BASE-X and SGMII Standards.

The accuracy of this Link Timer is within the following range.

+0 to -32.768 microseconds

1000BASE-X Standard

The Link-Timer is defined as having a duration somewhere between 10 and 20 milliseconds. The example design delivered with the core sets the binary value as follows:

100111101 = 317 decimal

This corresponds to a duration of between 10.354 and 10.387 milliseconds.

SGMII Standard

The Link-Timer is defined as having a duration of 1.6 milliseconds. The example design delivered with the core sets the binary value to

000110010 = 50 decimal

This corresponds to a duration of between 1.606 and 1.638 milliseconds.

Simulating Auto-Negotiation

Auto-Negotiation requires a minimum of three link timer periods for completion. If simulating the Auto-Negotiation procedure, setting the `link_timer[8:0]` port to a low value will greatly reduce the simulation time required to complete Auto-Negotiation.

Using the Auto-Negotiation Interrupt

The Auto-Negotiation function has an `an_interrupt` port. This is designed to be used with common micro-processor bus architectures (for example, the CoreConnect bus interfacing to MicroBlaze or the Virtex-II Pro embedded IBM PowerPC). For more information, see “[Auto-Negotiation Signal Pinout](#)” in Chapter 2.

The operation of this port is enabled or disabled and cleared via the MDIO Register 16, the Vendor-specific Auto-Negotiation Interrupt Control Register.

- When disabled, this port is permanently tied to logic 0.
- When enabled, this port will be set to logic 1 following the completion of an Auto-Negotiation cycle. It will remain high until it is cleared by writing 0 to bit 16.1 (Interrupt Status bit) of the “[MDIO Register 16: Vendor Specific Auto-Negotiation Interrupt Control](#),” page 129.

Dynamic Switching of 1000BASE-X and SGMII Standards

This chapter provides general guidelines for using the core to perform dynamic standards switching between 1000BASE-X and SGMII. The core will only provide this capability if generated with the appropriate option, as described in [Chapter 3, "Generating and Customizing the Core."](#)

Typical Application

[Figure 11-1](#) illustrates a typical application for the Ethernet 1000BASE-X PCS/PMA or SGMII core with the ability to dynamically switch between 1000BASE-X and SGMII standards.

The FPGA is shown connected to an external, off-the-shelf PHY with the ability to perform both BASE-X and BASE-T standards.

- The core must operate in 1000BASE-X mode to use the optical fibre
- The core must operate in SGMII mode to provide BASE-T functionality and use the twisted copper pair.

The GMII of the Ethernet 1000BASE-X PCS/PMA or SGMII core is shown connected to an embedded Ethernet Media Access Controller (MAC), for example the Tri-Mode Ethernet MAC core from Xilinx.

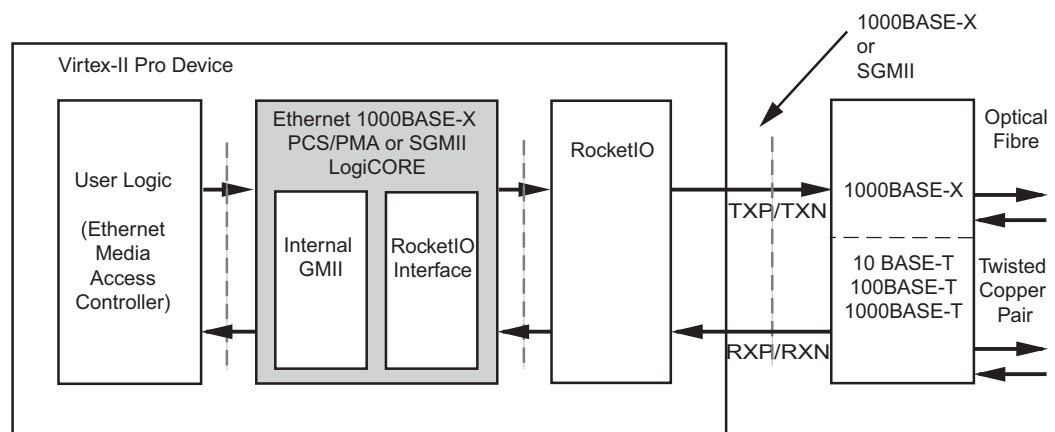


Figure 11-1: Typical Application for Dynamic Switching

Operation of the Core

Selecting the Power-On / Reset Standard

The external port of the core, `basex_or_sgmi i` (see “[Dynamic Switching Signal Pinout](#)” in [Chapter 2](#)), will select the default standard of the core as follows:

- Tie to logic ‘0’ in the core instantiation. The core powers-up and comes out of a reset cycle operating in the 1000BASE-X standard.
- Tie to logic ‘1’ in the core instantiation. The core powers-up and comes out of a reset cycle operating in the SGMII standard.

The `basex_or_sgmi i` port of the core could be dynamically driven. In this configuration, it is possible to drive a logical value onto the port, followed by a core reset cycle to switch the core to the desired standard. However, it is expected that the standard will be switched through the MDIO Management Registers.

Switching the Standard Using MDIO

The 1000BASE-X or SGMII standard of the core can be switched at any time by writing to the “[Vendor-specific Register: Standard Selection Register \(Register 17\)](#).” Following completion of this write, the MDIO Management Registers will immediately switch.

- Core set to 1000BASE-X standard. Management Registers 0 through 16 should be interpreted according to “[1000BASE-X Standard Using the Optional Auto-Negotiation](#),” [page 119](#).
- Core set to SGMII standard. Management Registers 0 through 16 should be interpreted according to “[SGMII Standard Using the Optional Auto-Negotiation](#),” [page 135](#).

Auto-Negotiation State Machine

- Core set to the 1000BASE-X standard. The Auto-Negotiation state machine operates as described in “[1000BASE-X Standard](#),” [page 156](#).
- Core set to perform the SGMII standard. The Auto-Negotiation state machine operates as described in “[SGMII Standard](#),” [page 156](#).
- Standard is switched during an Auto-Negotiation sequence. The Auto-Negotiation state machine will not immediately switch standards, but attempt to continue to completion at the original standard.
- Switching the standard using MDIO. This does not cause Auto-Negotiation to automatically restart. Xilinx recommends that after switching to a new standard using a MDIO write, immediately perform the following:
 - ◆ If you have switched to the 1000BASE-X standard, reprogram the Auto-Negotiation Advertisement Register (Register 4) to the desired settings.
 - ◆ For either standard, restart the Auto-Negotiation sequence by writing to bit 0.9 of the MDIO Control Register (Register 0).

Setting the Auto-Negotiation Link Timer

As described in “[Auto-Negotiation](#)” in [Chapter 10](#), the duration of the Auto-Negotiation Link Timer differs with the 1000BASE-X and the SGMII standards. To provide configurable link timer durations for both standards, the following ports are available. These ports

replace the `link_timer_value[8:0]` port that is used when the core is generated for a single standard.

- `link_timer_basex[8:0]` The value placed on this port is sampled at the beginning of the Auto-Negotiation cycle by the Link Timer when the core is set to perform the 1000BASE-X standard.
- `link_timer_sgmiic[8:0]` The value placed on this port is sampled at the beginning of the Auto-Negotiation cycle by the Link Timer when the core is set to perform the SGMII standard.

Both ports follow the same rules that are described in [“Setting the Configurable Link Timer,”](#) page 156.

Constraining the Core

This chapter defines the constraint requirements of the Ethernet 1000BASE-X PCS/PMA or SGMII core. An example UCF is provided with the HDL example design for the core to implement the constraints defined in this chapter.

See the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide* for a complete description of the CORE Generator output files and for details on the HDL example design.

Required Constraints

Device, Package, and Speedgrade Selection

The Ethernet 1000BASE-X PCS/PMA or SGMII core can be implemented in Virtex-II, Virtex-II Pro, Virtex-4, Virtex-5, Spartan-3, Spartan-3E, Spartan-3A, Spartan-3AN and Spartan-3 DSP devices. When selecting a device, be aware of the following considerations:

- Device must be large enough to accommodate the core
- Device must contain a sufficient number of IOBs
- -4 speed grade for Virtex-II, Spartan-3, Spartan-3E, Spartan-3A, Spartan-3AN and Spartan-3A DSP devices
- -5 speed grade for Virtex-II Pro FPGA
- -10 speed grade for Virtex-4 FPGA
- -1 speed grade for Virtex-5 FPGA
- The RocketIO transceiver is only supported in Virtex-II Pro, Virtex-4 FX, Virtex-5 LXT, Virtex-5 SXT, and Virtex-5 FXT FPGAs

I/O Location Constraints

No specific I/O location constraints required.

Placement Constraints

No specific placement constraints required.

Virtex-II Pro RocketIO MGTs for 1000BASE-X Constraints

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from the UCF are copied into the descriptions in the following sections to serve as examples. These should be studied in conjunction with

the HDL source code for the example design and with the information contained in [Chapter 7, "100BASE-X with RocketIO Transceivers."](#)

Clock Period Constraints

The clock provided to `userclk` must be constrained for a clock frequency of 62.5 MHz. The clock provided to `userclk2` must be constrained for a clock frequency of 125 MHz. The following UCF syntax shows the necessary constraints being applied to the example design.

```
#####
# PCS/PMA Clock period Constraints: please do not relax #
#####
NET "brefclk_ibufg" TNM_NET = "brefclk";
TIMESPEC "ts_brefclk" = PERIOD "brefclk" 16 ns HIGH 50 %;

NET "rocketio/rxrecclk" TNM_NET = "rxrecclk";
TIMESPEC "ts_rxrecclk" = PERIOD "rxrecclk" 16 ns;

NET "clk0" TNM_NET = "clk0";
TIMESPEC "ts_clk0" = PERIOD "clk0" "ts_brefclk";

NET "clk2x180" TNM_NET = "clk_tx";
TIMESPEC "ts_tx_clk" = PERIOD "clk_tx" "ts_brefclk"/2 PHASE + 4 ns HIGH
50 %;
```

Setting MGT Attributes

MGT attributes can be set by either of these methods:

- Directly from HDL source code during MGT instantiation (see the HDL source code for the example design)
- From the UCF

Attributes set from a UCF take priority. The UCF for the example design defines some user-modifiable attributes as illustrated in the following example. All attributes used in the example design UCF are based on the `GT_ETHERNET_1` defaults.

```
#####
# Rocket I/O constraints: #
# please refer to Rocket I/O documentation #
#####
INST "rocketio/mgt" TX_CRC_USE = FALSE;
INST "rocketio/mgt" RX_CRC_USE = FALSE;
INST "rocketio/mgt" REF_CLK_V_SEL = 1;
INST "rocketio/mgt" TERMINATION_IMP = 50;
INST "rocketio/mgt" TX_DIFF_CTRL = 500;
INST "rocketio/mgt" TX_PREEMPHASIS = 0;
```

MGT Transceiver Placement Constraints

The following UCF syntax illustrates the MGT transceiver placement for the example design. Special attention must be made to the placement of the SERDES alignment flip-flop as described in the *RocketIO Transceiver User Guide (Chapter 2, SERDES Alignment, Ports, and Attributes, ENPCOMMAALIGN, ENMCOMMAALIGN)*. This is the single flip-flop illustrated in [Figure 7-1](#).

```
#####
# Rocket I/O placement:                                     #
#####

# Place the Rocket I/O
INST "rocketio/mgt" LOC = "GT_X0Y1";

# Locate the SERDES alignment logic near the Rocket I/O.
# Please Refer to the Rocket I/O User Guide (Chapter 2,
# SERDES Alignment, Ports and Attributes, ENPCOMMAALIGN,
# ENMCOMMAALIGN).
# The following lock constraints are intended as an
# example of SERDES alignment logic placement in a
# XC2VP7 device when using GT_X0Y1. Please change the
# targeted slices appropriately for other combinations.
INST "rocketio/serdes_alignment" LOC = SLICE_X15Y72;
```

Virtex-II Pro RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints

All of the constraints documented in [“Virtex-II Pro RocketIO MGTs for 1000BASE-X Constraints”](#) apply. In addition, if the FPGA Fabric Rx Elastic Buffer is selected, area placement constraints are required to ensure that the correct local clock routing paths are used for `rxrecc1k`. This is described in *XAPP763* and in the remainder of this section.

With the MGT Rx Elastic Buffer bypassed, `rxrecc1k` clock is provided by the MGT to the FPGA fabric for the recovered receiver data signals leaving the transceiver. This data is then written into the replacement Rx Elastic Buffer implemented in the FPGA fabric. See [Chapter 8, “Virtex-II Pro Devices”](#) for more information about this logic.

For correct operation, `rxrecc1k` must be placed on specific clock routing in the vicinity of the MGT from which the clock signal originates. This is the MGT local clock route, a 5 x 12 Configurable Logic Block (CLB) array which is next to every MGT on the top of the device, or a 5 x 11 CLB array next to every MGT on the bottom of the device. Each array provides a minimum of 440 flip-flops plus two block SelectRAM™s; more than adequate for the fabric Rx Elastic Buffer requirements. A CLB array for the top of the device is illustrated in [Figure 12-1](#). This figure represents the view of this placement as seen in FPGA Editor.

The following UCF syntax shows an example of defining an `AREA_GROUP` for the `rxrecc1k` local clock route for all of the synchronous elements used in the example design. Because the block RAM is not included in the `AREA_GROUP`, a separate location constraint needs to be applied to the block RAM used.

```
#####
# Fabric Rx Elastic Buffer Placement:                       #
#####

# Constrain the slice area to be near the RocketIO
TIMEGRP "rxrecc1k" AREA_GROUP = "local_clk";
AREA_GROUP "local_clk" RANGE = SLICE_X6Y56:SLICE_X15Y79;

# Constrain the block RAM used for the fabric Rx Elastic
# Buffer to be near the RocketIO
INST "rocketio/clock_correction/dual_port_block_ram" LOC =
RAMB16_X1Y8;
```



Figure 12-1: Local Clock Place and Route for Top MGT

Virtex-4 RocketIO MGTs for 1000BASE-X Constraints

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from the UCF are copied into the following descriptions to serve as examples and should be studied in conjunction with the HDL source code for the example design. See also “Virtex-4 FX Devices” in Chapter 7.

Clock Period Constraints

The clock `txoutclk` is provided by the MGT for use in the FPGA fabric. It is connected to global clock routing to produce the `usrclk2` signal. This is the main 125 MHz clock used by all core logic and must be constrained.

DCLK is a clock with a frequency between 25 and 50 MHz, which must be provided to the Dynamic Reconfiguration Port and to the calibration block of the MGT. In the example design, this is constrained to 50 MHz.

The following UCF syntax shows these constraints being applied.

```

*****
# PCS/PMA Clock period Constraints: please do not relax      *
*****

NET "userclk2" TNM_NET = "userclk2";
TIMESPEC "TS_userclk2" = PERIOD "userclk2" 8 ns HIGH 50 %;

NET "dclk" TNM_NET = "dclk";
TIMESPEC "TS_dclk" = PERIOD "dclk" 20 ns HIGH 50 %;

```

Setting MGT Transceiver Attributes

The Virtex-4 MGT device has many attributes. These attributes are set directly from HDL source code for the transceiver wrapper file delivered with the example design. These are in the file `transceiver.vhd` (for VHDL design entry) or `transciever.v` (for Verilog design entry). See the *Ethernet 1000BASE-X PCS/PMA or SGMII Getting Started Guide* for a detailed description of the example design provided with the core.

This HDL transceiver wrapper file was initially created using Architecture Wizard. See the *Virtex-4 FPGA RocketIO Multi-Gigabit Transceiver User Guide* (UG076) for a description of available attributes.

MGT Placement Constraints

The following UCF syntax illustrates the MGT placement constraints for the example design. Because Virtex-4 MGTs are always available in pairs, two MGTs are always instantiated in the example design, even if one is inactive.

```

*****
# Example Rocket I/O placement      *
*****
# Lock down the REFCLK pins:
NET  brefclk_p   LOC = F26;
NET  brefclk_n   LOC = G26;

# Lock down the GT11 pair and GT11 clock module
INST "core_wrapper/rocketio/GT11_1000X_A" LOC = GT11_X0Y5;
INST "core_wrapper/rocketio/GT11_1000X_B" LOC = GT11_X0Y4;
INST "GT11CLK_MGT_INST" LOC = GT11CLK_X0Y3;

# Lock down the RocketIO pins:
NET  "rxp0" LOC = J26;
NET  "rxn0" LOC = K26;
NET  "txp0" LOC = M26;
NET  "txn0" LOC = N26;
NET  "rxp1" LOC = U26;
NET  "rxn1" LOC = V26;
NET  "txp1" LOC = P26;
NET  "txn1" LOC = R26;

```

Virtex-4 RocketIO MGTs for SGMII or Dynamic Standards Switching Constraints

All the constraints described in the section “[Virtex-4 RocketIO MGTs for 1000BASE-X Constraints](#).” In addition, if the FPGA Fabric Rx Elastic Buffer is selected, an extra clock period constraint of 16 ns is required for `rxrecclk1`.

With the MGT Rx Elastic Buffer bypassed, `rxrecclk1` is provided by the MGT to the FPGA fabric for the recovered receiver data signals leaving the transceiver. This data is then written into the replacement Rx Elastic Buffer implemented in the FPGA fabric. See “[Virtex-4 Devices for SGMII or Dynamic Standards Switching](#),” page 101.

The following UCF syntax shows the necessary constraint being applied to GT11 A.

```

#*****
# PCS/PMA Clock period Constraints for the GT11 A          *
# recovered clock: please do not relax                    *
#*****

NET "core_wrapper/rocketio/rxrecclk10" TNM_NET = "rxrecclk10";
TIMESPEC "ts_rxrecclk10" = PERIOD "rxrecclk10" 16 ns;

```

Virtex-5 RocketIO GTP Transceivers for 1000BASE-X Constraints

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from the UCF are copied into the following descriptions to serve as examples, and should be studied with the HDL source code for the example design. See also “[Virtex-5 LXT and SXT Devices](#)” in Chapter 7.

Clock Period Constraints

The `clkkin` clock is provided to the GTP transceiver. It is a high-quality reference clock with a frequency of 125 MHz and should be constrained.

The `refclkout` clock is provided by the GTP for use in the FPGA fabric, which is then connected to global clock routing to produce the `usrclk2` signal. This is the main 125 MHz clock used by all core logic and must be constrained.

The following UCF syntax shows these constraints being applied.

```

#*****
# PCS/PMA Clock period Constraints: please do not relax    *
#*****

NET "*clkkin" TNM_NET = "clkkin";
TIMESPEC "TS_clkkin" = PERIOD "clkkin" 8 ns HIGH 50 %;

NET "*refclkout" TNM_NET = "refclkout";
TIMESPEC "TS_refclkout" = PERIOD "refclkout" 8 ns HIGH 50 %;

```

Setting GTP Transceiver Attributes

The Virtex-5 GTP transceiver has many attributes that are set directly from HDL source code for the transceiver wrapper file delivered with the example design. These can be found in the `rocketio_wrapper_gtp_tile.vhd` file (for VHDL design entry) or the `rocketio_wrapper_gtp_tile.v` file (for Verilog design entry); these files were generated using the GTP Transceiver Wizard - to change the attributes, re-run the Wizard. See “[Virtex-5 RocketIO GTP Wizard](#)” in Chapter 7.

Virtex-5 RocketIO GTP Transceivers for SGMII or Dynamic Standards Switching Constraints

If the core is generated to use the GTP Rx Elastic Buffer, all of the constraints apply, as defined in “[Clock Period Constraints](#),” page 166. However, if the FPGA Fabric Rx Elastic Buffer is selected, an extra clock period constraint of 8 ns is required for `rxrecclk`: with the GTP Rx Elastic Buffer bypassed, `rxrecclk` is provided by the GTP transceiver to the FPGA fabric for the recovered receiver data signals leaving the transceiver. This data is then written into the replacement Rx Elastic Buffer implemented in the FPGA fabric. See “[Virtex-5 LXT or SXT Devices for SGMII or Dynamic Standards Switching](#),” page 103 for more information about this logic.

The following UCF syntax shows the necessary constraint being applied to the `rxrecclk` signal sourced from GTP 0.

```

#*****
# PCS/PMA Clock period Constraints for the GTP 0          *
# recovered clock: please do not relax                   *
#*****

NET "core_wrapper/rocketio/rxrecclk0" TNM_NET = "rxrecclk0";
TIMESPEC "ts_rxrecclk0" = PERIOD "rxrecclk0" 8 ns;

```

Setting GTP Transceiver Attributes

Additionally, if the FPGA Fabric Rx Elastic Buffer is selected, then the attributes of the Virtex-5 GTP transceiver which are set directly from HDL source code do differ from the standard case. These can be found in the `rocketio_wrapper_gtp_tile.vhd` file (for VHDL design entry) or the `rocketio_wrapper_gtp_tile.v` file (for Verilog design entry); these files were generated using the GTP RocketIO Wizard - to change the attributes, re-run the Wizard. See “[Virtex-5 RocketIO GTP Wizard](#)” in Chapter 8.

Virtex-5 RocketIO GTX Transceivers for 1000BASE-X Constraints

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from the UCF are copied into the following descriptions to serve as examples, and should be studied with the HDL source code for the example design. See also “[Virtex-5 FXT Devices](#)” in Chapter 7.

Clock Period Constraints

The `clk_in` clock is provided to the GTX transceiver. It is a high-quality reference clock with a frequency of 125 MHz and should be constrained.

The `refclkout` clock is provided by the GTX for use in the FPGA fabric—this is the main 125MHz clock reference source for the FPGA fabric and should be constrained. This is then connected to a DCM. The ports `CLK0` (125MHz) and `CLKDV` (62.5MHz) of this DCM are then placed onto global clock routing to produce the `usrclk2` and `usrclk` clock signals respectively. The Xilinx tools will trace the `refclkout` constraint through the DCM and automatically generate clock period constraints for the DCM output clocks. So constraints `usrclk2` and `usrclk` do not need to be manually applied.

The following UCF syntax shows these constraints being applied.

```

#*****
# PCS/PMA Clock period Constraints: please do not relax *
#*****

```

```
NET "*clkin" TNM_NET = "clkin";
TIMESPEC "TS_clkin" = PERIOD "clkin" 8 ns HIGH 50 %;

NET "*refclkout" TNM_NET = "refclkout";
TIMESPEC "TS_refclkout" = PERIOD "refclkout" 8 ns HIGH 50 %;
```

Setting GTX Transceiver Attributes

The Virtex-5 GTX transceiver has many attributes that are set directly from HDL source code for the transceiver wrapper file delivered with the example design. These can be found in the `rocketio_wrapper_gtx_tile.vhd` file (for VHDL design entry) or the `rocketio_wrapper_gtx_tile.v` file (for Verilog design entry): these files were generated using the GTX Transceiver Wizard - to change the attributes, re-run the Wizard. See “[Virtex-5 RocketIO GTX Wizard](#)” in [Chapter 7](#).

Virtex-5 RocketIO GTX Transceivers for SGMII or Dynamic Standards Switching Constraints

If the core is generated to use the GTX Rx Elastic Buffer, then all of the constraints documented in “[Clock Period Constraints](#),” [page 167](#), apply.

However, if the FPGA Fabric Rx Elastic Buffer is selected, then an extra clock period constraint of 16 ns is required for `rxrecclk`: with the GTX Rx Elastic Buffer bypassed, `rxrecclk` is provided by the GTX transceiver to the FPGA fabric for the recovered receiver data signals leaving the transceiver. This data is then written into the replacement Rx Elastic Buffer implemented in the FPGA fabric. See “[Virtex-5 FXT Devices for SGMII or Dynamic Standards Switching](#),” [page 105](#) for more information about this logic.

The following UCF syntax shows the necessary constraint being applied to the `rxrecclk` signal sourced from GTX 0.

```
*****
# PCS/PMA Clock period Constraints for the GTP/X 0          *
# recovered clock: please do not relax                      *
*****

NET "core_wrapper/rocketio/rxrecclk0" TNM_NET = "rxrecclk0";
TIMESPEC "ts_rxrecclk0" = PERIOD "rxrecclk0" 16 ns;
```

Setting GTX Transceiver Attributes

Additionally, if the FPGA Fabric Rx Elastic Buffer is selected, then the attributes of the Virtex-5 GTX transceiver which are set directly from HDL source code do differ from the standard case. These can be found in the `rocketio_wrapper_gtx_tile.vhd` file (for VHDL design entry) or the `rocketio_wrapper_gtx_tile.v` file (for Verilog design entry): these files were generated using the GTX RocketIO Wizard - to change the attributes, re-run the Wizard. See “[Virtex-5 RocketIO GTX Wizard](#)” in [Chapter 8](#).

Ten-Bit Interface Constraints

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from this UCF have been copied into the descriptions in this section to serve as examples, and should be studied with the HDL source code for the example design. See also [Chapter 6](#), “[The Ten-Bit Interface](#).”

Clock Period Constraints

The clocks provided to `pma_rx_clk0` and `pma_rx_clk1` must be constrained for a clock frequency of 62.5 MHz. The clock provided to `gtx_clk` must be constrained for a clock frequency of 125 MHz. The following UCF syntax shows the constraints being applied to the example design.

```
#####
# TBI Clock period Constraints: please do not relax      #
#####
NET "pma_rx_clk0" TNM_NET = "pma_rx_clk0";
TIMESPEC "ts_pma_rx_clk0" = PERIOD "pma_rx_clk0" 16000 ps HIGH 50 %;

NET "pma_rx_clk1" TNM_NET = "pma_rx_clk1";
TIMESPEC "ts_pma_rx_clk1" = PERIOD "pma_rx_clk1" 16000 ps HIGH 50 %;

NET "gtx_clk_bufg" TNM_NET = "clk_tx";
TIMESPEC "ts_tx_clk" = PERIOD "clk_tx" 8000 ps HIGH 50 %;
```

Period constraints should be applied to cover signals in to and out of the block memory based 8B/10B encoder and decoder.

```
# Constrain between flip-flops and the Block Memory for the 8B10B
encoder and decoder
INST "gig_eth_pcs_pma_core/BU2/U0/PCS_OUTPUT/DECODER/LOOK_UP_TABLE"
TNM = "codec8b10b";
INST "gig_eth_pcs_pma_core/BU2/U0/PCS_OUTPUT/ENCODER/LOOK_UP_TABLE"
TNM = "codec8b10b";
TIMESPEC "ts_ffs_to_codec8b10b" = FROM FFS TO "codec8b10b" 8000 ps;
TIMESPEC "ts_codec8b10b_to_ffs" = FROM "codec8b10b" TO FFS 8000 ps;
```

Ten-Bit Interface IOB Constraints

The following constraints target the flip-flops that are inferred in the top level HDL file for the example design. Constraints are set to ensure that these are placed in IOBs.

```
INST "tx_code_group_reg*"      IOB = true;
INST "ewrap_reg"               IOB = true;
INST "en_cdet_reg"             IOB = true;
INST "rx_code_group0_reg*"     IOB = true;
INST "rx_code_group1_reg*"     IOB = true;
```

Note: For Virtex-4 and Virtex-5 devices, the example design will directly instantiate IOB DDR components and the previous constraints are not included.

The Ten-Bit Interface (TBI) is a 3.3 volt signal level interface. The 3.3 volt LVTTTL SelectIO standard is the default for Virtex-II devices. The following constraints may safely be added. The 3.3 volt LVTTTL SelectIO standard is not the default for other families. Use the following constraints and take into account the device IO Banking rules when fixing PADs.

```
INST "tx_code_group<?>"      IOSTANDARD = LVTTTL;
INST "pma_tx_clk"            IOSTANDARD = LVTTTL;

INST "rx_code_group<?>"      IOSTANDARD = LVTTTL;
INST "pma_rx_clk0"           IOSTANDARD = LVTTTL;

INST "loc_ref"               IOSTANDARD = LVTTTL;
INST "ewrap"                  IOSTANDARD = LVTTTL;
INST "en_cdet"                IOSTANDARD = LVTTTL;
```

In addition, the example design provides pad locking on the TBI for several families. This is included as a guideline only, and there are no specific I/O location constraints for this core.

TBI Input Setup/Hold Timing

Input TBI Timing Specification




Figure 12-2: Input TBI timing

Figure 12-2 and Table 12-1 illustrate the setup and hold time window for the input TBI signals. These specify the worst-case data valid window presented to the FPGA device pins. There is only a 2 ns data valid window of guaranteed data presented across the TBI input bus. This must be correctly sampled by the FPGA devices.

Table 12-1: Input TBI Timing

Symbol	Min	Max	Units
t_{SETUP}	2.00	-	ns
t_{HOLD}	0.00	-	ns

Virtex-II, and Virtex-II Pro Devices

Figure 6-2 illustrates the TBI input logic provided by the example design for the Virtex-II and Virtex-II Pro family. Although not illustrated, these families have input delay elements (always of a fixed delay). These are also automatically inserted by the Xilinx tools and are set to provide a zero-hold time. These input delays automatically meet input setup and hold timing on the TBI without any specific constraints.

Spartan-3, Spartan-3E, and Spartan-3A Devices

Figure 6-3, page 72 illustrates the TBI input logic provided by the example design for the Spartan-3 class family. DCMs are used on the pma_rx_clk0 and pma_rx_clk1 clock paths as illustrated. Phase-shifting is then applied to the DCMs to align the resultant clocks so that they correctly sample the 2 ns. TBI data valid window at the input DDR flip-flops.

The fixed phase shift is applied to the DCMs using the following UCF syntax.

```
INST "core_wrapper/tbi_rx_clk0_dcm" CLKOUT_PHASE_SHIFT = FIXED;
INST "core_wrapper/tbi_rx_clk0_dcm" PHASE_SHIFT = -10;
INST "core_wrapper/tbi_rx_clk0_dcm" DESKEW_ADJUST = 0;
```

```

INST "core_wrapper/tbi_rx_clk1_dcm" CLKOUT_PHASE_SHIFT = FIXED;
INST "core_wrapper/tbi_rx_clk1_dcm" PHASE_SHIFT = -10;
INST "core_wrapper/tbi_rx_clk1_dcm" DESKEW_ADJUST = 0;

```

The values of `PHASE_SHIFT` are preconfigured in the example designs to meet the setup and hold constraints for the example TBI pinout in the particular device. The setup/hold timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script).

For customers fixing their own pinout, the setup and hold figures reported in the TRCE report can be used to initially setup the approximate DCM phase shift values. [Appendix C, "Calculating the DCM Fixed Phase Shift Value"](#) describes a more accurate method for fixing the phase shift by using hardware measurement of a unique PCB design.

Virtex-4 Devices

[Figure 6-4, page 73](#) illustrates the TBI input logic provided by the example design for the Virtex-4 family. `IDELAY` elements are instantiated on the TBI data input path as illustrated: the number of tap delays is currently set to zero. This can be modified in the UCF file, if desired, to de-skew the bus for PCB routing.

A fixed tap delay is applied to delay the `pma_rx_clk0` clock so that it correctly samples the TBI data at the IOB IDDR register, thereby meeting TBI setup and hold timing.

The tap delays are applied using the following UCF syntax.

```

#-----
# To Adjust TBI Rx Input Setup/Hold Timing -
#-----
INST "core_wrapper/delay_pma_rx_clk" IOBDELAY_VALUE = "40";
INST "core_wrapper/tbi_rx_data_bus[9].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[8].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[7].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[6].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[5].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[4].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[3].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[2].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[1].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";
INST "core_wrapper/tbi_rx_data_bus[0].delay_tbi_rx_data"
IOBDELAY_VALUE = "0";

```

The value of `IOBDELAY_VALUE` for the `pma_rx_clk0` clock is preconfigured in the example designs to meet the setup and hold constraints for the example TBI pinout in the particular device. The setup/hold timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script). See ["Understanding Timing Reports for Setup/Hold Timing."](#)

Virtex-5 Devices

Figure 6-6, page 75 illustrates the TBI input logic provided by the example design for the Virtex-5 family. IODELAY elements are instantiated on the TBI data input path as illustrated: the number of tap delays is currently set to zero. This can be modified in the UCF file, if desired, to de-skew the bus for PCB routing. A fixed tap delay is applied to delay the `pma_rx_clk0` clock so that it correctly samples the TBI data at the IOB IDDR register, thereby meeting TBI setup and hold timing.

The tap delays are applied using the following UCF syntax.

```
#-----
# To Adjust TBI Rx Input Setup/Hold Timing
#-----
INST "core_wrapper/delay_pma_rx_clk" IDELAY_VALUE = "63";
INST "core_wrapper/tbi_rx_data_bus[9].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[8].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[7].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[6].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[5].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[4].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[3].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[2].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[1].delay_tbi_rx_data" IDELAY_VALUE
= "0";
INST "core_wrapper/tbi_rx_data_bus[0].delay_tbi_rx_data" IDELAY_VALUE
= "0";
```

The value of `IDELAY_VALUE` for the `pma_rx_clk0` clock is preconfigured in the example designs to meet the setup and hold constraints for the example TBI pinout in the particular device. The setup/hold timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script). See [“Understanding Timing Reports for Setup/Hold Timing.”](#)

Constraints When Implementing an External GMII

The constraints defined in this section are implemented in the UCF for the example designs delivered with the core. Sections from this UCF have been copied into the following examples, and should be studied in conjunction with the HDL source code for the example design. See also the section, [“Implementing External GMII,”](#) page 61.

Clock Period Constraints

When implementing an external GMII, the Transmitter Elastic Buffer delivered with the example design (or similar logic) must be used. The input transmitter GMII signals are then synchronous to their own clock domain (`gmi_i_tx_clk` is used in the example design). This clock must be constrained for a clock frequency of 125 MHz. The following UCF syntax shows the necessary constraints being applied to the example design.

```
#####
# GMI Clock period Constraints: please do not relax      #
#####
NET "gmii_tx_clk_bufg" TNM_NET = "gmii_tx_clk";
TIMESPEC "ts_gmii_tx_clk" = PERIOD "gmii_tx_clk" 8000 ps HIGH 50 %;
```

GMII IOB Constraints

The following constraints target the flip-flops that are inferred in the top level HDL file for the example design. Constraints are set to ensure that these are placed in IOBs.

```
#####
# GMI Transmitter Constraints: place flip-flops in IOB  #
#####
INST "gmii_txd*" IOB = true;
INST "gmii_tx_en" IOB = true;
INST "gmii_tx_er" IOB = true;

#####
# GMI Receiver Constraints: place flip-flops in IOB    #
#####
INST "gmii_rxd_obuf*" IOB = true;
INST "gmii_rx_dv_obuf" IOB = true;
INST "gmii_rx_er_obuf" IOB = true;
```

The GMII is a 3.3 volt signal level interface. The 3.3 volt LVTTTL SelectIO standard is the default for Virtex-II devices. The following constraints may be safely added. The 3.3 volt LVTTTL SelectIO standard is not the default for other families. Use the following constraints and take into account the device IO Banking rules when fixing PADs.

```
INST "gmii_txd<?>" IOSTANDARD = LVTTTL;
INST "gmii_tx_en" IOSTANDARD = LVTTTL;
INST "gmii_tx_er" IOSTANDARD = LVTTTL;

INST "gmii_rxd<?>" IOSTANDARD = LVTTTL;
INST "gmii_rx_dv" IOSTANDARD = LVTTTL;
INST "gmii_rx_er" IOSTANDARD = LVTTTL;

INST "gmii_tx_clk" IOSTANDARD = LVTTTL;
INST "gmii_rx_clk" IOSTANDARD = LVTTTL;
```

In addition, the example design provides pad locking on the GMII for several families. This is provided as a guideline only; there are no specific I/O location constraints for this core.

GMII Input Setup/Hold Timing

Input GMII timing specification

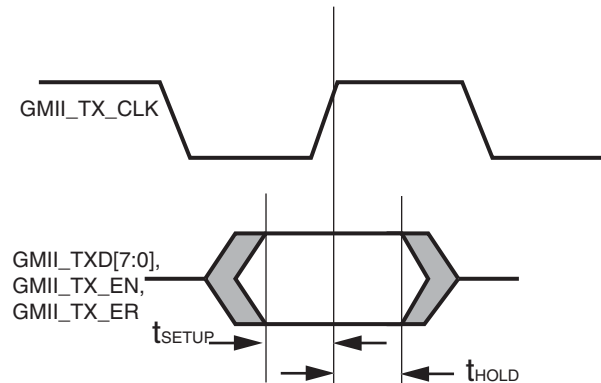


Figure 12-3: Input GMII timing

Figure 12-3 and Table 12-2 illustrate the setup and hold time window for the input GMII signals. These are the worst-case data valid window presented to the FPGA device pins.

Observe that there is, in total, a 2 ns data valid window of guaranteed data which is presented across the GMII input bus. This must be correctly sampled by the FPGA devices.

Table 12-2: Input GMII Timing

Symbol	Min	Max	Units
t_{SETUP}	2.00	-	ns
t_{HOLD}	0.00	-	ns

Virtex-II, and Virtex-II Pro devices

Figure 5-14 illustrates the GMII input logic which is provided by the example design for the Virtex-II and Virtex-II Pro family. Although not illustrated, these families have input delay elements (which are always of a fixed delay). These are also automatically inserted by the Xilinx tools and are set to provide a zero-hold time.

These input delays will automatically meet input setup and hold timing on the GMII without any specific constraints.

Spartan-3, Spartan-3E, and Spartan-3A devices

Figure 5-15 illustrates the GMII input logic which is provided by the example design for the Spartan-3 class family. A DCM must be used on the `gmii_tx_clk` clock path as illustrated. Phase-shifting is then applied to the DCM to align the resultant clock so that it will correctly sample the 2ns GMII data valid window at the input flip-flops.

The fixed phase shift is applied to the DCM using the following UCF syntax.

```
INST "gmii_tx_dcm" CLKOUT_PHASE_SHIFT = FIXED;
INST "gmii_tx_dcm" PHASE_SHIFT = -20;
INST "gmii_tx_dcm" DESKEW_ADJUST = 0;
```

The value of `PHASE_SHIFT` is preconfigured in the example designs to meet the setup and hold constraints for the example GMII pinout in the particular device. The setup/hold

timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script).

For customers fixing their own pinout, the setup and hold figures reported in the TRCE report can be used to initially setup the approximate DCM phase shift. [Appendix C, "Calculating the DCM Fixed Phase Shift Value"](#) describes a more accurate method for fixing the phase shift by using hardware measurement of a unique PCB design.

Virtex-4 devices

[Figure 5-16](#) illustrates the GMII input logic provided by the example design for the Virtex-4 family. IODELAY elements are instantiated on the GMII data input path as illustrated: the number of tap delays is currently set to zero. This can be modified in the UCF file, if desired, to de-skew the bus for PCB routing.

A fixed tap delay is applied to delay the `gmii_tx_clk` clock so that it correctly samples the GMII data at the IOB flip-flop, thereby meeting GMII setup and hold timing.

The tap delays are applied using the following UCF syntax.

```
#-----
# To Adjust GMII Tx Input Setup/Hold Timing           -
#-----
INST "delay_gmii_tx_en"  IOBDELAY_VALUE = "53";
INST "delay_gmii_tx_er"  IOBDELAY_VALUE = "53";

INST "gmii_data_bus[7].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[6].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[5].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[4].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[3].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[2].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[1].delay_gmii_txd"  IOBDELAY_VALUE = "53";
INST "gmii_data_bus[0].delay_gmii_txd"  IOBDELAY_VALUE = "53";
```

The value of `IOBDELAY_VALUE` for the `gmii_tx_clk` clock is preconfigured in the example designs to meet the setup and hold constraints for the example GMII pinout in the particular device. The setup/hold timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script). See ["Understanding Timing Reports for Setup/Hold Timing."](#)

Virtex-5 devices

[Figure 5-17](#) illustrates the GMII input logic provided by the example design for the Virtex-5 family. IODELAY elements are instantiated on the GMII data input path as illustrated: the number of tap delays is currently set to zero. This can be modified in the UCF file, if desired, to de-skew the bus for PCB routing.

A fixed tap delay is applied to delay the `gmii_tx_clk` clock so that it correctly samples the GMII data at the IOB flip-flop, thereby meeting GMII setup and hold timing.

The tap delays are applied using the following UCF syntax.

```
#-----
# To Adjust GMII Tx Input Setup/Hold Timing           -
#-----
INST "delay_gmii_tx_en"  IDELAY_VALUE = "33";
INST "delay_gmii_tx_er"  IDELAY_VALUE = "33";

INST "gmii_data_bus[7].delay_gmii_txd"  IDELAY_VALUE = "33";
```

```

INST "gmii_data_bus[6].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[5].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[4].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[3].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[2].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[1].delay_gmii_txd" IDELAY_VALUE = "33";
INST "gmii_data_bus[0].delay_gmii_txd" IDELAY_VALUE = "33";

```

The value of IDELAY_VALUE for the gmii_tx_clk clock is preconfigured in the example designs to meet the setup and hold constraints for the example GMII pinout in the particular device. The setup/hold timing which is achieved after place-and-route is reported in the datasheet section of the TRCE report (created by the implement script). See “Understanding Timing Reports for Setup/Hold Timing.”

Understanding Timing Reports for Setup/Hold Timing

Devices Other Than Virtex-4 or Virtex-5

Setup and Hold results for the TBI or GMII input busses for the following devices are defined in the Data Sheet Report section of the Timing Report: Virtex-II, Virtex-II Pro, Spartan-3, Spartan-3E, Spartan-3A, Spartan-3AN, and Spartan-3A DSP. The results are self-explanatory and show an obvious correlation and relationship to Figure 12-2 and Figure 12-3.

The following example shows the GMII report from a Virtex-II device. The implementation requires 1.531 ns of setup (this is less than the 2 ns required, to allow for slack). The implementation requires -0.125 ns of hold (this is less than the 0 ns required, to allow for slack).

Data Sheet report:

All values displayed in nanoseconds (ns)

Setup/Hold to clock gmii_tx_clk

Source	Setup to clk (edge)	Hold to clk (edge)	Internal Clock(s)	Clock Phase
gmii_tx_en	1.531 (R)	-0.141 (R)	gmii_tx_clk_bufg	0.000
gmii_tx_er	1.531 (R)	-0.141 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<0>	1.531 (R)	-0.141 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<1>	1.525 (R)	-0.135 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<2>	1.531 (R)	-0.141 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<3>	1.525 (R)	-0.135 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<4>	1.515 (R)	-0.125 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<5>	1.515 (R)	-0.125 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<6>	1.520 (R)	-0.130 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<7>	1.520 (R)	-0.130 (R)	gmii_tx_clk_bufg	0.000

Virtex-4 or Virtex-5 Devices

Setup and hold results for the TBI or GMII input busses can be found in the data sheet report section of the Timing Report. Note that initially, the results do not indicate an obvious relationship to Figure 12-2 and Figure 12-3. The following example shows the GMII report from a Virtex-4 device.

Data Sheet report:

All values displayed in nanoseconds (ns)

Setup/Hold to clock gmii_tx_clk

Source	Setup to clk (edge)	Hold to clk (edge)	Internal Clock(s)	Clock Phase
gmii_tx_en	-6.501 (R)	7.875 (R)	gmii_tx_clk_bufg	0.000
gmii_tx_er	-6.504 (R)	7.878 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<0>	-6.506 (R)	7.880 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<1>	-6.521 (R)	7.893 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<2>	-6.518 (R)	7.890 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<3>	-6.515 (R)	7.889 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<4>	-6.521 (R)	7.894 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<5>	-6.520 (R)	7.895 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<6>	-6.514 (R)	7.889 (R)	gmii_tx_clk_bufg	0.000
gmii_txd<7>	-6.513 (R)	7.889 (R)	gmii_tx_clk_bufg	0.000

The implementation requires -6.501 ns of setup. Figure 12-4 illustrates that this represents a figure of 1.499 ns relative to the following rising edge of the clock (because the IDELAY has acted to delay the clock by an entire period when measured from the input flip-flop). This is less than the 2 ns required, and so there is slack.

The implementation requires 7.893 ns of hold. Figure 12-4 illustrates that this represents a figure of -0.107 ns relative to the following rising edge of the clock (because the IDELAY has acted to delay the clock by an entire period when measured from the input flip-flop). This is less than the 0 ns required, and so there is slack.




Figure 12-4: Timing Report Setup/Hold Illustration

Interfacing to Other Cores

This chapter describes some additional design considerations associated with implementing the Ethernet 1000BASE-X PCS/PMA or SGMII core with other cores.

- 1-Gigabit Ethernet MAC
- Tri-Mode Ethernet MAC

Integrating with the 1-Gigabit Ethernet MAC Core

The 1000BASE-X PCS/PMA or SGMII core can be integrated in a single device with the 1-Gigabit Ethernet MAC core to extend the system functionality to include the MAC sublayer. This core supports full-duplex operation at 1 Gigabit per second.

A description of and instructions for obtaining the newest 1-Gigabit Ethernet MAC core are located on the 1-Gigabit Ethernet MAC product page:

www.xilinx.com/systemio/gmac/index.htm

Integration of the 1-Gigabit Ethernet MAC to 1000BASE-X PCS with TBI

Figure 13-1 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in 1000BASE-X mode with the parallel TBI) to the 1-Gigabit Ethernet MAC core.

Features of this configuration include:

- Direct internal connections are made between the GMII interfaces between the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected to that of the 1-Gigabit Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the embedded Receiver Elastic Buffer in the Ethernet 1000BASE-X PCS/PMA, the entire GMII is synchronous to a single clock domain. Therefore, `gtx_clk` is used as the 125 MHz reference clock for both cores, and the transmitter and receiver logic of the 1-Gigabit Ethernet MAC core operates in the same clock domain.

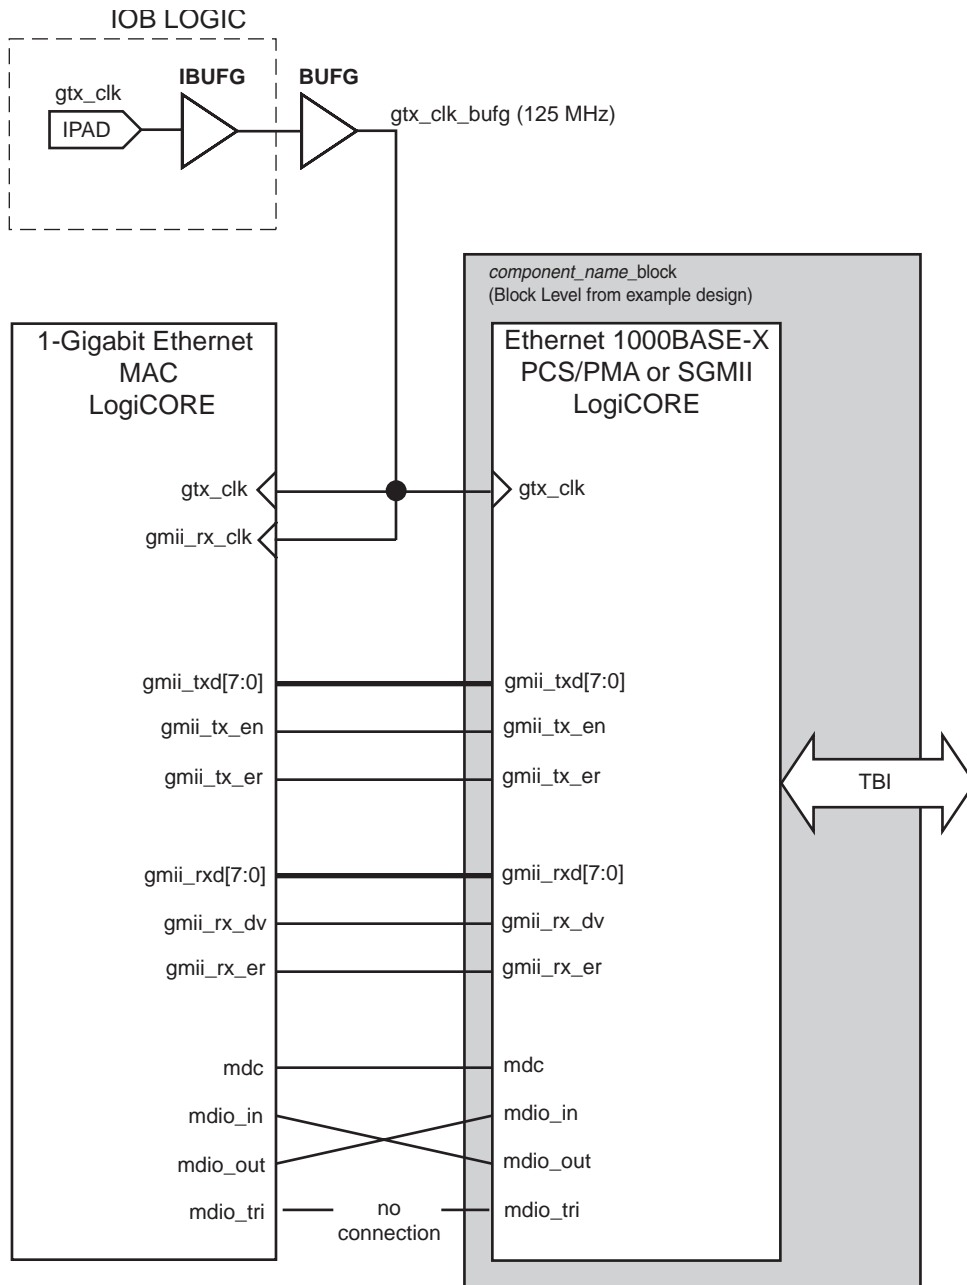


Figure 13-1: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS with TBI

Integration of the 1-Gigabit Ethernet MAC Using a RocketIO Transceiver

Virtex-II Pro Devices

Figure 13-2 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in 1000BASE-X mode) to the 1-Gigabit Ethernet MAC core.

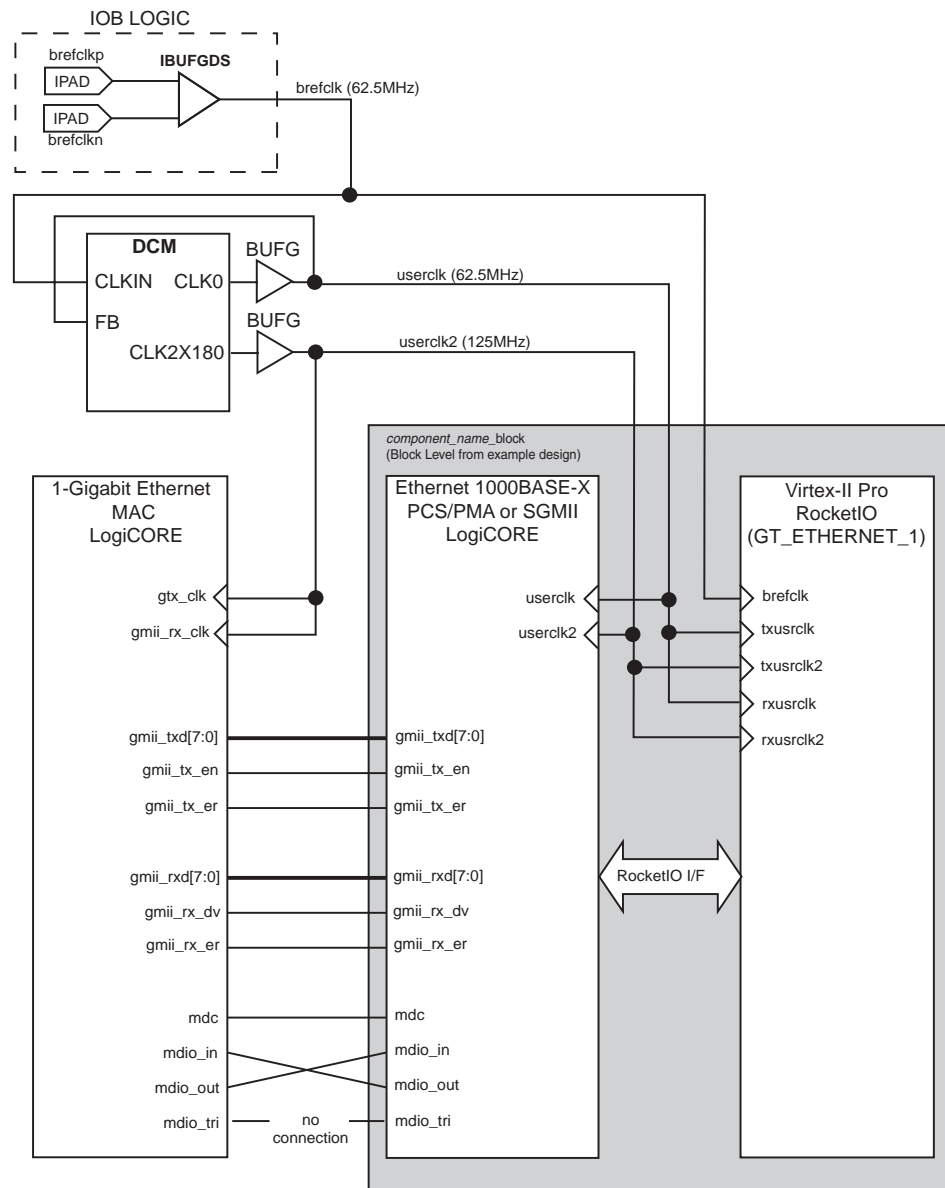


Figure 13-2: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-II Pro MGT

Features of this configuration include:

- Direct internal connections are made between the GMII interfaces between the two cores.

- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the 1-Gigabit Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the embedded Receiver Elastic Buffer in the MGT, the entire GMII is synchronous to a single-clock domain. Therefore, `userclk2` is used as the 125 MHz reference clock for both cores, and the transmitter and receiver logic of the 1-Gigabit Ethernet MAC core now operate in the same clock domain.

Virtex-4 Devices

Figure 13-2 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in 1000BASE-X mode) to the 1-Gigabit Ethernet MAC core.

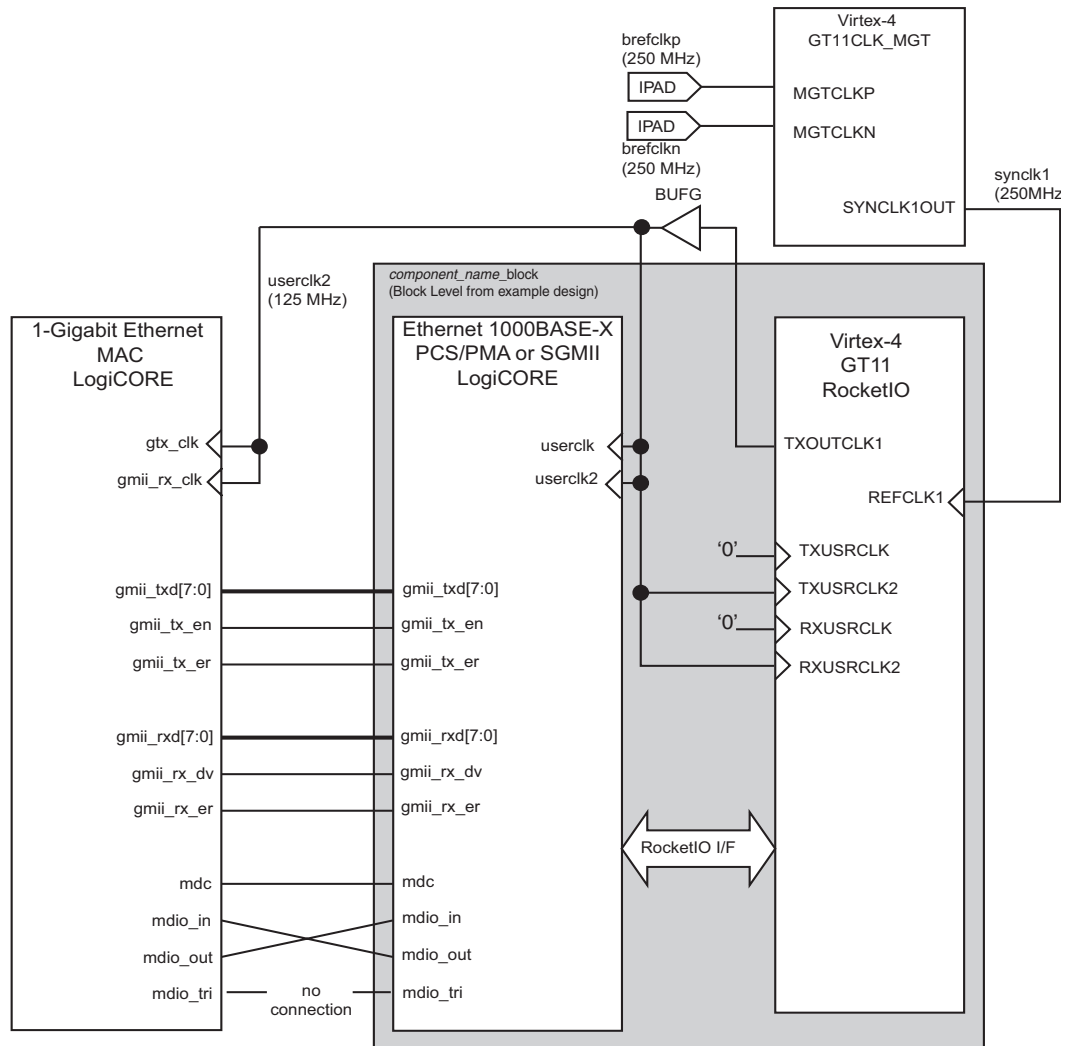


Figure 13-3: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-4 MGT

Features of this configuration include:

- Direct internal connections are made between the GMII interfaces between the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the 1-Gigabit Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the embedded Receiver Elastic Buffer in the MGT, the entire GMII is synchronous to a single clock domain. Therefore `userclk2` is used as the 125 MHz reference clock for both cores, and the transmitter and receiver logic of the 1-Gigabit Ethernet MAC core now operate in the same clock domain.

Virtex-5 LXT and SXT Devices

Figure 13-4 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in 1000BASE-X mode) to the 1-Gigabit Ethernet MAC core.

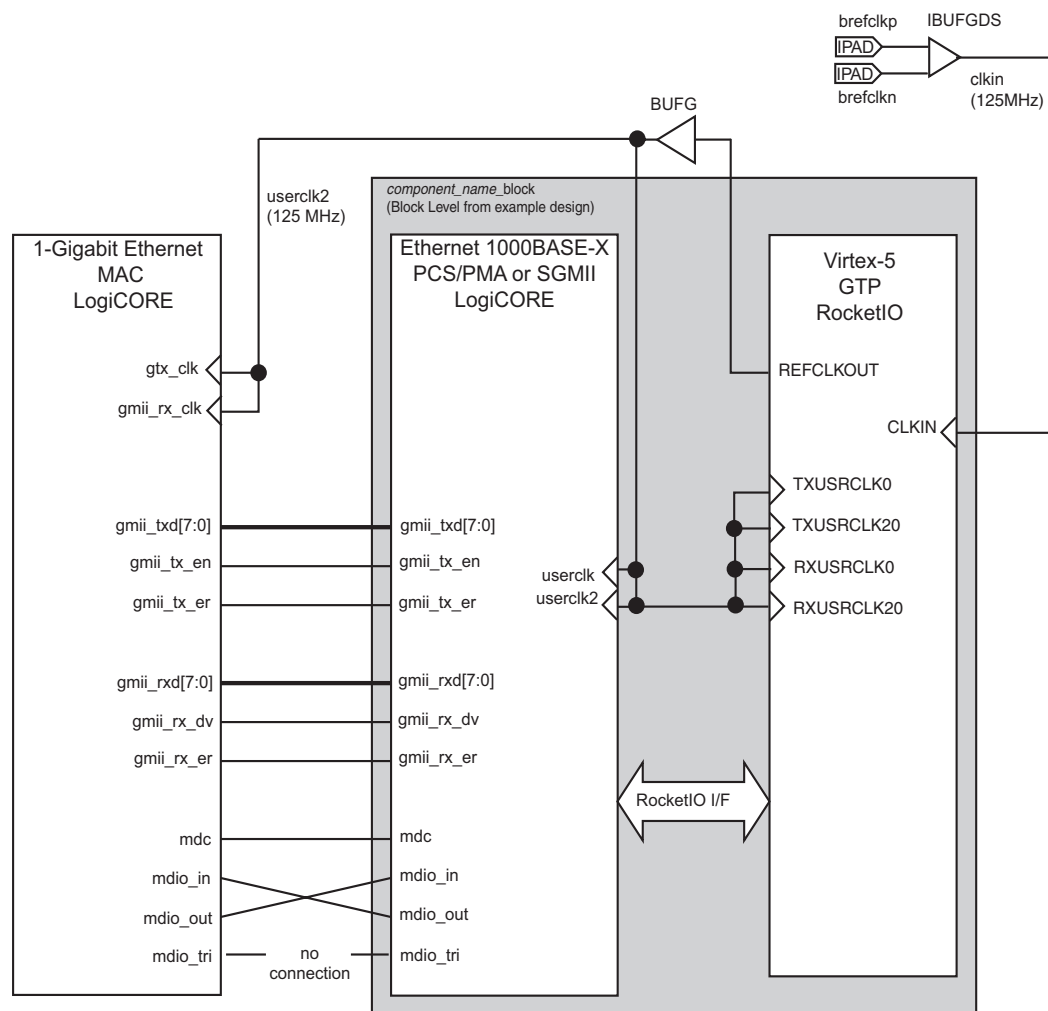


Figure 13-4: 1-Gigabit Ethernet MAC Extended to Include 1000BASE-X PCS and PMA Using a Virtex-5 GTP Transceiver

Features of this configuration include:

- Direct internal connections are made between the GMII interfaces between the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the 1-Gigabit Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the embedded Receiver Elastic Buffer in the GTX transceiver, the entire GMII is synchronous to a single clock domain. Therefore `userclk2` is used as the 125 MHz reference clock for both cores, and the transmitter and receiver logic of the 1-Gigabit Ethernet MAC core now operate in the same clock domain.

Integration of the 1-Gigabit Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality

The connections required to provide SGMII functionality are identical to the connections required for either “[Integration of the 1-Gigabit Ethernet MAC to 1000BASE-X PCS with TBI](#)” or “[Integration of the 1-Gigabit Ethernet MAC Using a RocketIO Transceiver](#)”, depending upon the chosen physical interface. The only difference is that the Ethernet 1000BASE-X PCS/PMA or SGMII core is generated with the SGMII or Dynamic switching option.

Note: When operating at 1 Gbps speed only, the Rx Elastic Buffer internal to the GTP transceiver should be used to save device resources. Additionally, when operating at 1 Gbps only, the SGMII Adaptation Module instantiated from within the block level of the example design is not required and can optionally be removed.

Integrating with the Tri-Mode Ethernet MAC Core

The 1000BASE-X PCS/PMA or SGMII core can be integrated in a single device with the Tri-Mode Ethernet MAC core to extend the system functionality to include the MAC sub-layer. This core provides support for operation at 10 Mbps, 100 Mbps, and 1 Gbps.

A description of the latest available IP update containing the Tri-Mode Ethernet MAC core and instructions can be found in the Tri-Mode Ethernet MAC product Web page:

www.xilinx.com/systemio/temac/index.htm

Caution! The Tri-Mode Ethernet MAC should always be configured for full-duplex operation when used with an SGMII. This constraint is due to the increased latency introduced by the SGMII logic. Without full-duplex operation, frame collisions could be undetected and the MAC response will not be timely.

Integration of the Tri-Mode Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality with TBI

Figure 13-7 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in SGMII mode with the TBI) to the Tri-Mode Ethernet MAC core. The following is a description of the functionality.

- The SGMII Adaptation module, provided in the example design for the Ethernet 1000BASE-X PCS/PMA or SGMII core when generated to the SGMII standard, can be used to interface the two cores.

- If both cores have been generated with the optional management interface, the MDIO port can be connected to that of the Tri-Speed Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the Receiver Elastic Buffer in the core, the entire GMII (transmitter and receiver paths) is synchronous to a single clock domain. Therefore, the `txcoreclk` and `rxcoreclk` inputs of the Tri-Speed Ethernet MAC core can always be driven from the same clock source.

Figure 13-7 illustrates the Tri-Mode Ethernet MAC core generated with the optional clock enable circuitry. This is the most efficient way to connect the two cores together in terms of clock resource usage and so is recommended. See the *Tri-Mode Ethernet MAC User Guide* for more information.

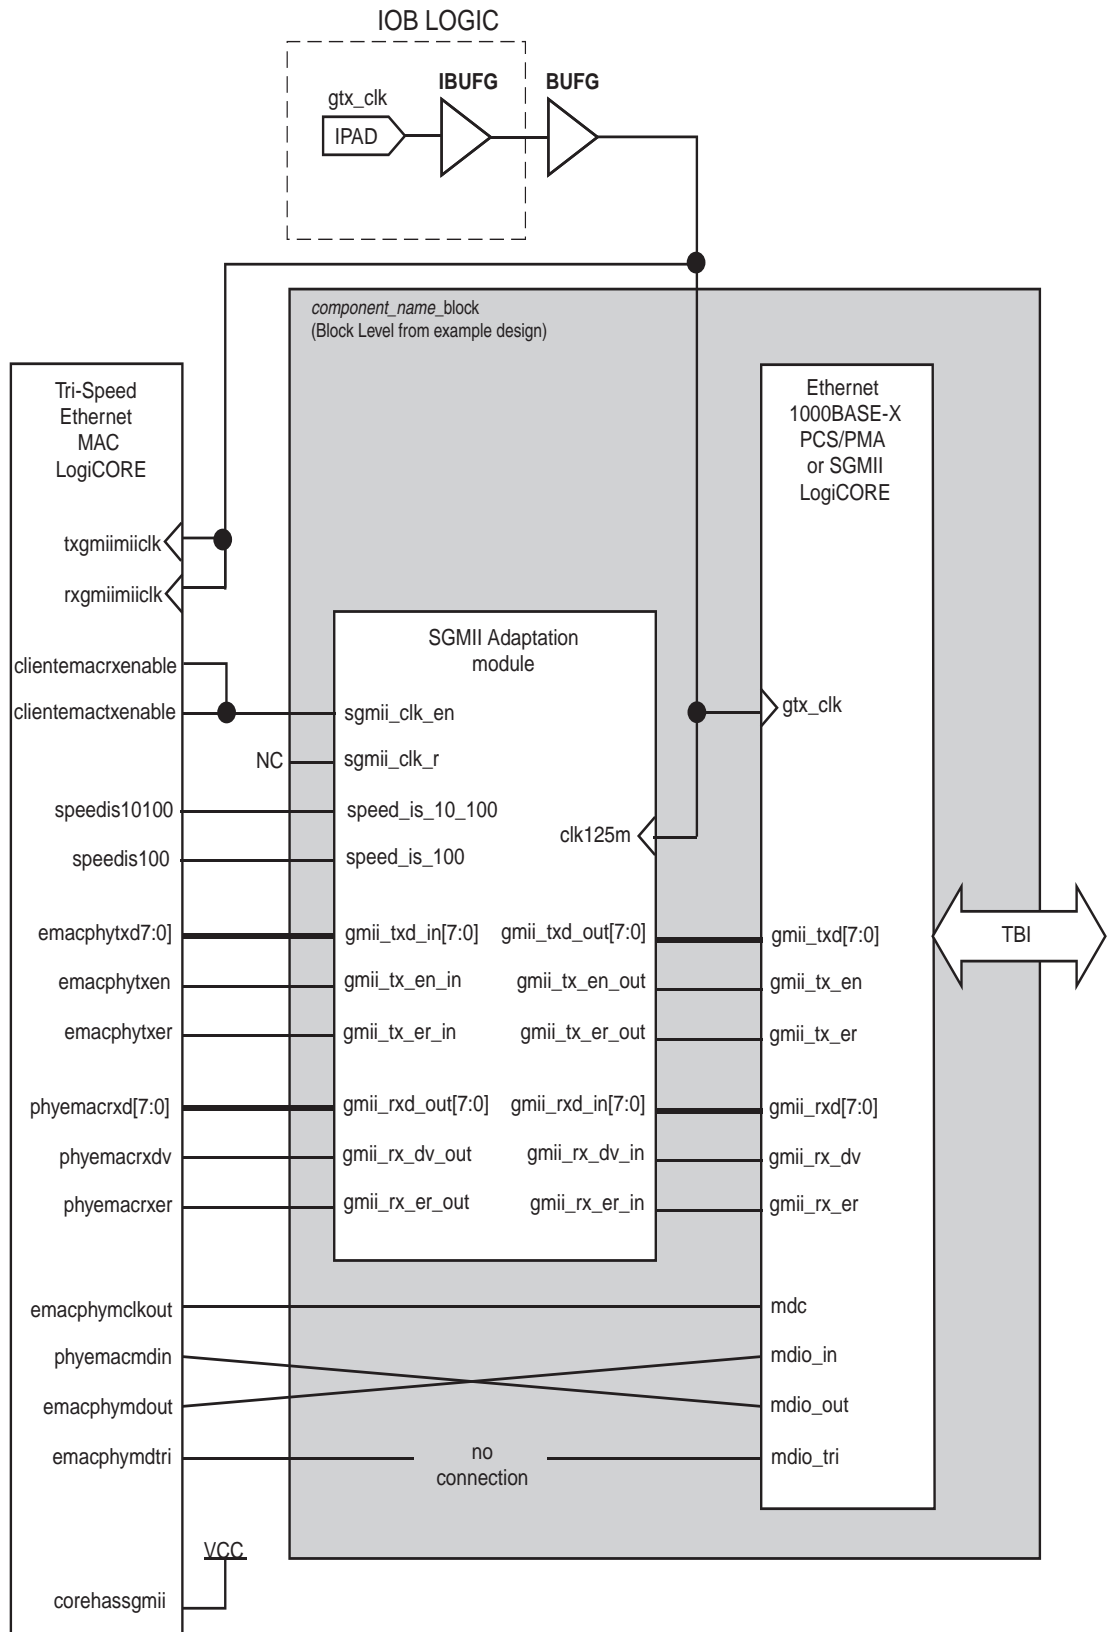


Figure 13-6: Tri-Speed Ethernet MAC Extended to use an SGMII with TBI

Integration of the Tri-Mode Ethernet MAC to Provide SGMII (or Dynamic Switching) Functionality using RocketIO Transceivers

Virtex-II Pro Devices

Figure 13-7 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in SGMII mode with the Virtex-II Pro MGT) to the Tri-Mode Ethernet MAC core. The following is a description of the functionality.

- The SGMII Adaptation module, provided in the example design for the Ethernet 1000BASE-X PCS/PMA or SGMII core when generated to the SGMII standard, can be used to interface the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected to that of the Tri-Speed Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the Receiver Elastic Buffer, the entire GMII (transmitter and receiver paths) is synchronous to a single clock domain. Therefore, the `txcoreclk` and `rxcoreclk` inputs of the Tri-Speed Ethernet MAC core can always be driven from the same clock source.

Figure 13-7 illustrates the Tri-Mode Ethernet MAC core generated with the optional clock enable circuitry. This is the most efficient way to connect the two cores together in terms of clock resource usage and so is recommended. See the *Tri-Mode Ethernet MAC User Guide* for more information.

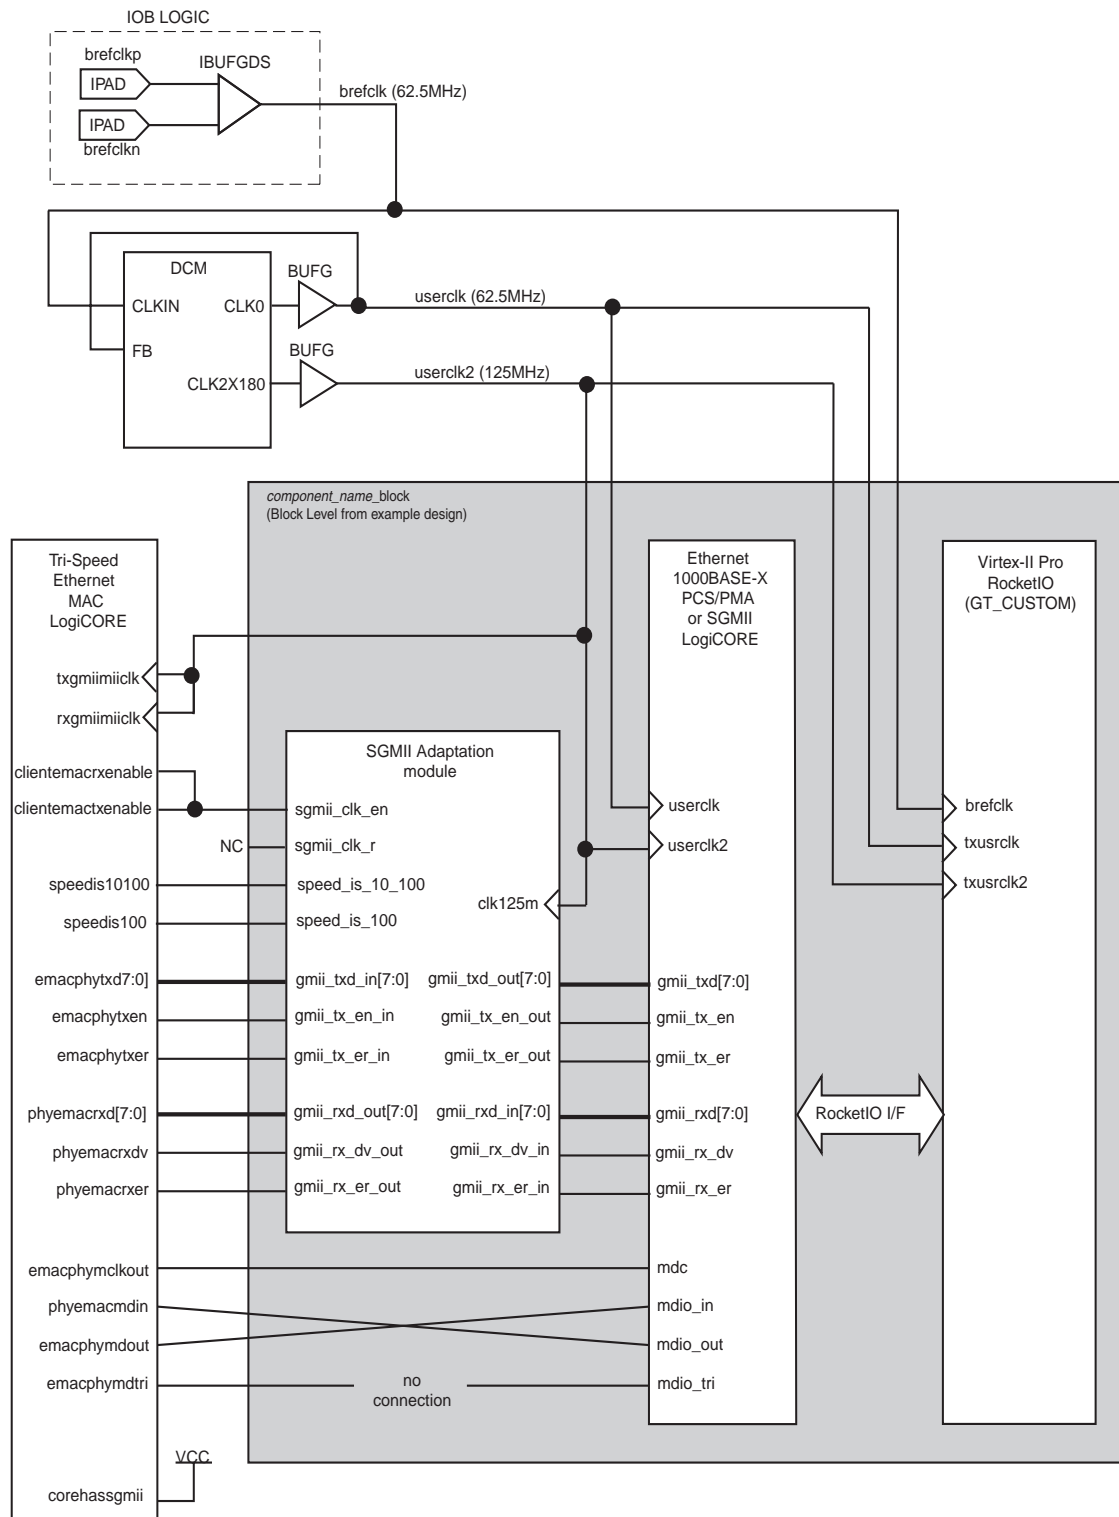


Figure 13-7: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-II Pro

Virtex-4 Devices

Figure 13-8 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in SGMII mode with the Virtex-4 MGT) to the Tri-Mode Ethernet MAC core.

The following conditions apply.

- The SGMII Adaptation module, as provided in the example design for the Ethernet 1000BASE-X PCS/PMA or SGMII core, when generated to the SGMII standard can be used to interface the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the Tri-Speed Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the Receiver Elastic Buffer, the entire GMII (transmitter and receiver paths) is synchronous to a single clock domain. Therefore the `txcoreclk` and `rxcoreclk` inputs of the Tri-Speed Ethernet MAC core can always be driven from the same clock source. The entire design is synchronous to the 125 MHz reference clock derived from the `CLK2X180` output of the DCM.

Figure 13-8 illustrates the Tri-Mode Ethernet MAC core generated with the optional clock enable circuitry. This is the most efficient way to connect the two cores together in terms of clock resource usage and so is recommended. See the *Tri-Mode Ethernet MAC User Guide* for more information.

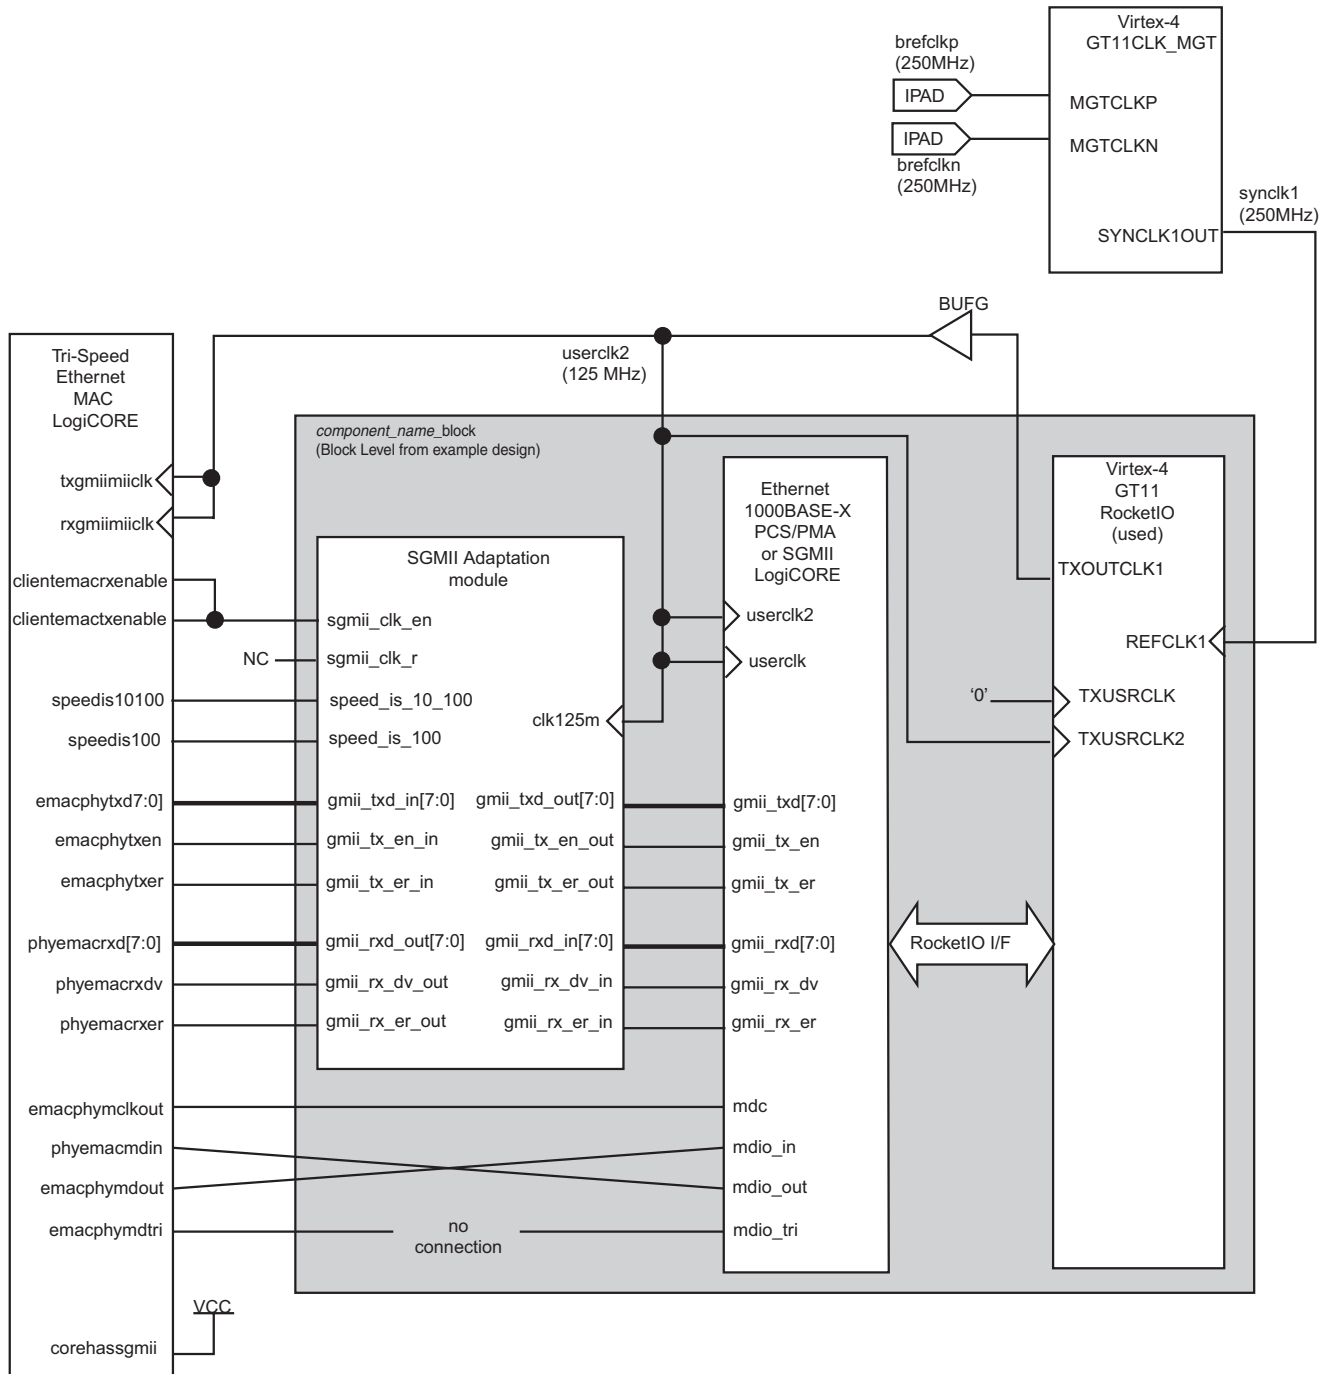


Figure 13-8: Tri-Speed Ethernet MAC Extended to Use an SGMII in Virtex-4

Virtex-5 LXT and SXT Devices

Figure 13-9 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in SGMII mode with the Virtex-5 GTP) to the Tri-Mode Ethernet MAC core.

The following conditions apply.

- The SGMII Adaptation module, as provided in the example design for the Ethernet 1000BASE-X PCS/PMA or SGMII core when generated to the SGMII standard, can be used to interface the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the Tri-Speed Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the Receiver Elastic Buffer, the entire GMII (transmitter and receiver paths) is synchronous to a single clock domain. Therefore the `txcoreclk` and `rxcoreclk` inputs of the Tri-Speed Ethernet MAC core can always be driven from the same clock source. The entire design is synchronous to the 125 MHz reference clock derived from the `CLK2X180` output of the DCM.

Figure 13-9 illustrates the Tri-Mode Ethernet MAC core generated with the optional clock enable circuitry. This is the most efficient way to connect the two cores together in terms of clock resource usage and so is recommended. See the *Tri-Mode Ethernet MAC User Guide* for more information.

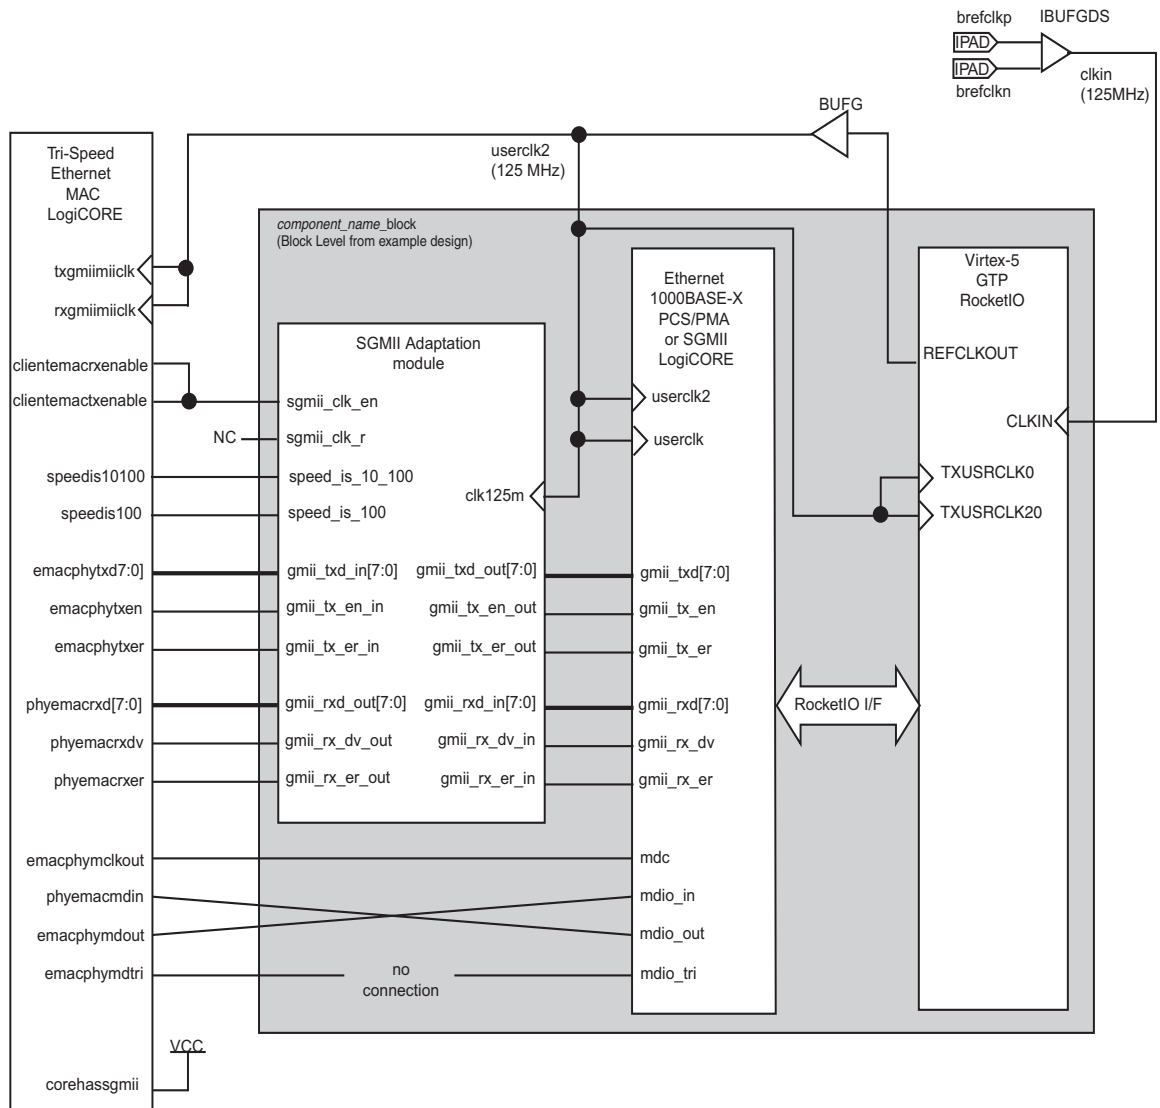


Figure 13-9: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-5 LXT/SXT

Virtex-5 FXT Devices

Figure 13-10 illustrates the connections and clock management logic required to interface the Ethernet 1000BASE-X PCS/PMA or SGMII core (when used in SGMII mode with the Virtex-5 GTX) to the Tri-Mode Ethernet MAC core.

The following conditions apply.

- The SGMII Adaptation module, as provided in the example design for the Ethernet 1000BASE-X PCS/PMA or SGMII core when generated to the SGMII standard, can be used to interface the two cores.
- If both cores have been generated with the optional management interface, the MDIO port can be connected up to that of the Tri-Speed Ethernet MAC core, allowing the MAC to access the embedded configuration and status registers of the Ethernet 1000BASE-X PCS/PMA or SGMII core.
- Due to the Receiver Elastic Buffer, the entire GMII (transmitter and receiver paths) is synchronous to a single clock domain. Therefore the `txcoreclk` and `rxcoreclk` inputs of the Tri-Speed Ethernet MAC core can always be driven from the same clock source. The entire design is synchronous to the 125 MHz reference clock derived from the `CLK2X180` output of the DCM.

Figure 13-10 illustrates the Tri-Mode Ethernet MAC core generated with the optional clock enable circuitry. This is the most efficient way to connect the two cores together in terms of clock resource usage and so is recommended. See the *Tri-Mode Ethernet MAC User Guide* for more information.

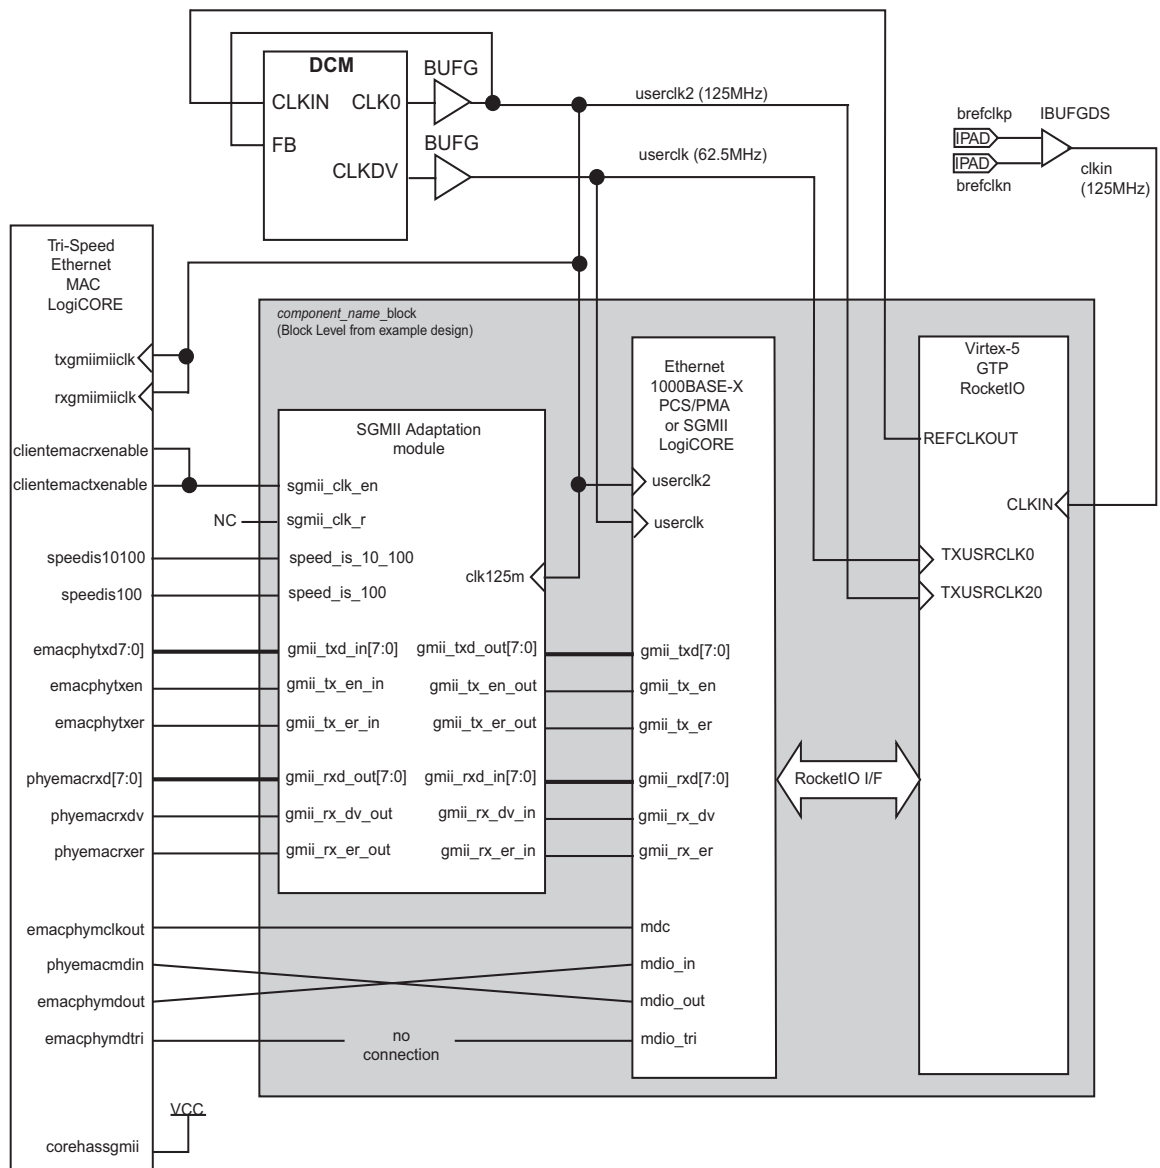


Figure 13-10: Tri-Speed Ethernet MAC Extended to use an SGMII in Virtex-5 FXT

Special Design Considerations

This chapter describes the unique design considerations associated with implementing the Ethernet 1000BASE-X PCS/PMA or SGMII core.

Power Management

No power management considerations are recommended for the Ethernet 1000BASE-X PCS/PMA or SGMII core when using it with the TBI. When using the Ethernet 1000BASE-X PCS/PMA or SGMII core with the Virtex-II Pro, the RocketIO transceiver may be placed in a low-power state in either of the following ways:

- Writing to the PCS Configuration Register 0 (if using the core with the optional Management Interface). The low-power state can only be removed by issuing the core with a reset. This reset can be achieved either by writing to the software reset bit in the PCS Configuration Register 0, or by driving the core reset port.
- Asserting the Power Down bit in the `configuration_vector` (if using the core without the optional Management Interface). The low-power state can only be removed by issuing the core with a reset by driving the reset port of the core.

Startup Sequencing

IEEE 802.3 clause 22.2.4.1.6 states that by default, a PHY should power-up in an isolate state (electrically isolated from the GMII).

- If you are using the core with the optional Management Interface, it is necessary to write to the PCS Configuration Register 0 to take the core out of the isolate state.
- If using the core without the optional Management interface, it is the responsibility of the client to ensure that the isolate input signal in the `configuration_vector` is asserted at power-on.

Loopback

This section details the implementation of the loopback feature. Loopback mode is enabled or disabled by either the [“MDIO Management Interface,”](#) page 115, or by the [“Optional Configuration Vector,”](#) page 151.

Core with the TBI

There is no physical loopback path in the core. Placing the core into loopback has the effect of asserting logic 1 on the `ewrap` signal of the TBI (see [“1000BASE-X PCS with TBI Pinout,”](#)

page 38). This instructs the attached PMA SERDES device to enter loopback mode as illustrated in Figure 14-1.

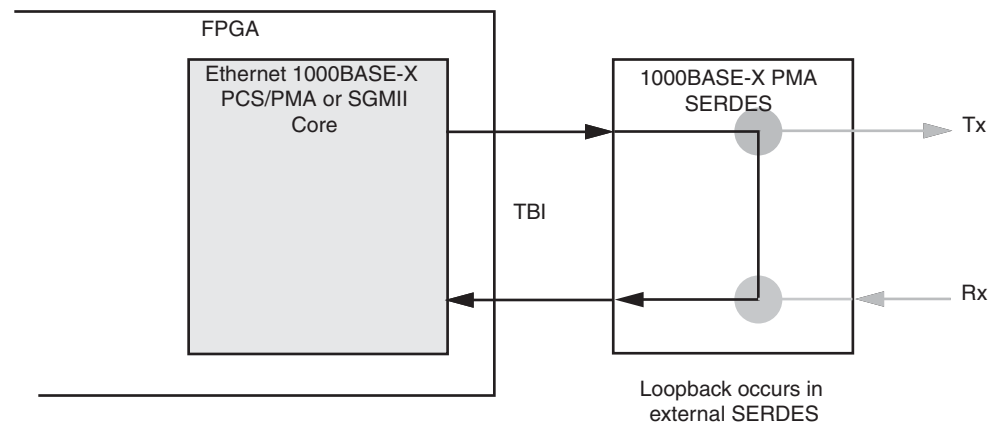


Figure 14-1: Loopback Implementation Using the TBI

Core with RocketIO Transceiver

The loopback path is implemented in the core as illustrated in Figure 14-2. When placed into loopback, the data is routed from the transmitter path to the receiver path at the last possible point in the core. This point is immediately before the RocketIO interface. When placed in loopback, the core creates a constant stream of Idle code groups that are transmitted through the MGT or GTP transceiver in accordance with the *IEEE 802.3* specification.

Earlier versions (before v5.0) of the core implemented loopback differently. The serial loopback feature of the RocketIO transceiver was used by driving the `LOOPBACK[1:0]` port of the RocketIO (MGT or GTP) transceiver. This is no longer the case, and the `loopback[1:0]` output port of the core is now permanently set to logic "00." However, for debugging purposes, the `LOOPBACK[1:0]` input port of the RocketIO transceiver may be directly driven by the user logic to place it in either parallel or serial loopback mode.

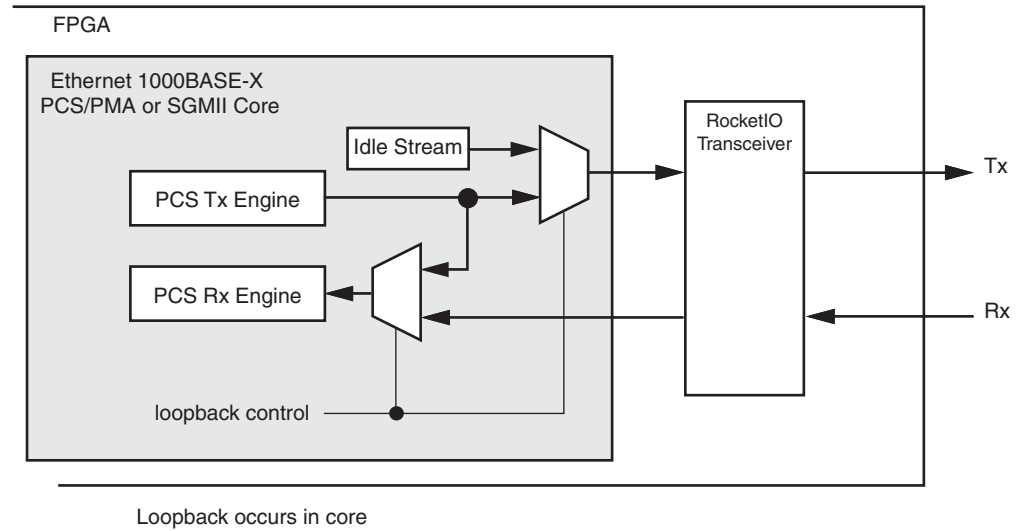


Figure 14-2: Loopback Implementation When Using the Core with RocketIO Transceivers

Implementing the Design

This chapter describes how to simulate and implement your design containing the Ethernet 1000BASE-X PCS/PMA or SGMII core.

Pre-implementation Simulation

A functional model of the Ethernet 1000BASE-X PCS/PMA or SGMII core netlist is generated by the CORE Generator to allow simulation of the core in the design-phase of the project.

Using the Simulation Model

For information about setting up your simulator to use the pre-implemented model, please consult the Xilinx *Synthesis and Verification Design Guide*, included in your Xilinx software installation.

The model is provided in the CORE Generator project directory.

VHDL Design Entry

```
<component_name>.vhd
```

Verilog Design Entry

```
<component_name>.v
```

This model can be compiled along with the user's code to simulate the overall system.

Synthesis

XST - VHDL

In the CORE Generator project directory, there is a `<component_name>.vho` file that is a component and instantiation template for the core. Use this to help instance the Ethernet 1000BASE-X PCS/PMA or SGMII core into your VHDL source.

After the entire design is complete, create the following:

- An XST project file `top_level_module_name.prj` listing all the user source code files
- An XST script file `top_level_module_name.scr` containing your required synthesis options.

To synthesize the design, run

```
$ xst -ifn top_level_module_name.scr
```

See the *XST User Guide* for more information on creating project and synthesis script files, and running the `xst` program.

XST - Verilog

There is a module declaration for the Ethernet 1000BASE-X PCS/PMA or SGMII core in the CORE Generator project directory:

```
<component_name>/implement/<component_name>.mod.v
```

Use this module to help instance the Ethernet 1000BASE-X PCS/PMA or SGMII core into your Verilog source.

After the entire design is complete, do the following:

- Generate an XST project file `top_level_module_name.prj` listing all user source code files.

Make sure to include the following as the first two files in the project list.

```
%XILINX%/verilog/src/ise/unisim_comp.v
```

and

```
<component_name>/implement/component_name.mod.v
```

- Generate an XST script file `top_level_module_name.scr` containing your required synthesis options.

To synthesize the design, run:

```
$ xst -ifn top_level_module_name.scr
```

See the *XST User Guide* for more information on creating project and synthesis script files, and running the `xst` program.

Implementation

Generating the Xilinx Netlist

To generate the Xilinx netlist, the `ngdbuild` tool is used to translate and merge the individual design netlists into a single design database—the NGD file. Also merged at this stage is the UCF for the design. An example of the `ngdbuild` command is:

```
$ ngdbuild -sd path_to_core_netlist -sd path_to_user_synth_results \  
-uc top_level_module_name.ucf top_level_module_name
```

Mapping the Design

To map the logic gates of the user design netlist into the CLBs and IOBs of the FPGA, run the `map` command. The `map` command writes out a physical design to an NCD file. An example of the `map` command is:

```
$ map -o top_level_module_name_map.ncd top_level_module_name.ngd \  
top_level_module_name.pcf
```

Placing and Routing the Design

The `par` command must be executed to place and route the user's design logic components (mapped physical logic cells) within an NCD file, in accordance with the

layout and timing requirements specified within the PCF file. The `par` command outputs the placed and routed physical design to an NCD file.

An example of the `par` command is:

```
$ par top_level_module_name_map.ncd top_level_module_name.ncd \  
    top_level_module_name.pcf
```

Static Timing Analysis

The `trce` command must be executed to evaluate timing closure on a design and create a Timing Report file (TWR) that is derived from static timing analysis of the Physical Design file (NCD). The analysis is typically based on constraints included in the optional PCF file.

An example of the `trce` command is:

```
$ trce -o top_level_module_name.twr top_level_module_name.ncd \  
    top_level_module_name.pcf
```

Generating a Bitstream

The `bitgen` command must be executed to create the configuration bitstream (BIT) file based on the contents of a physical implementation file (NCD). The BIT file defines the behavior of the programmed FPGA.

An example of the `bitgen` command is:

```
$ bitgen -w top_level_module_name.ncd
```

Post-Implementation Simulation

The purpose of post-implementation simulation is to verify that the design as implemented in the FPGA works as expected.

Generating a Simulation Model

To generate a chip-level simulation netlist for your design, the `netgen` command must be run.

VHDL

```
$ netgen -sim -ofmt vhdl -ngm top_level_module_name_map.ngm \  
    -tm netlist top_level_module_name.ncd \  
    top_level_module_name_postimp.vhd
```

Verilog

```
$ netgen -sim -ofmt verilog -ngm top_level_module_name_map.ngm \  
    -tm netlist top_level_module_name.ncd \  
    top_level_module_name_postimp.v
```

Using the Model

For information about setting up your simulator to use the pre-implemented model, please consult the Xilinx *Synthesis and Verification Design Guide*, included in your Xilinx software installation.

In addition, use the following guidelines to determine the simulator type required:

Virtex-5 Devices

Virtex-5 device designs incorporating a RocketIO transceiver require either a Verilog LRM-IEEE 1364-2005 encryption-compliant simulator or a SWIFT-compliant simulator.

- For a Verilog LRM-IEEE 1364-2005 encryption-compliant simulator, ModelSim v6.3c is currently supported.
- For a SWIFT-compliant simulator, Cadence IUS v6.1 and Synopsys Synopsys VCS 2006.06-SP1 are currently supported.

Virtex-4 and Virtex-II Pro Devices

Virtex-4 and Virtex-II Pro device designs incorporating a RocketIO transceiver require a SWIFT-compliant simulator. ModelSim, Cadence IUS, and Synopsys are currently supported using the versions defined above.

Other Implementation Information

For more information about using the Xilinx implementation tool flow, including command line switches and options, consult the software manuals provided with the Xilinx ISE software.

Core Verification, Compliance, and Interoperability

Verification

The Ethernet 1000BASE-X PCS/PMA or SGMII core has been verified with extensive simulation and hardware verification.

Simulation

A highly parameterizable transaction based test bench was used to test the core. Testing included the following:

- Register Access
- Loss of Synchronization
- Auto-Negotiation and error handling
- Frame Transmission and error handling
- Frame Reception and error handling
- Clock Compensation in the Elastic Buffers

Hardware Verification

The core has been tested in a variety of hardware test platforms at Xilinx to represent different parameterizations, including the following:

- The core with RocketIO transceiver and performing the 1000BASE-X standard was tested with the 1-Gigabit Ethernet MAC core from Xilinx.

This follows the architecture shown in [Figure 13-2](#). A test platform was built around these cores, including a back-end FIFO capable of performing a simple ping function, and a test pattern generator. Software running on the embedded PowerPC was used to provide access to all configuration and status registers. Version 3.0 of this core was taken to the University of New Hampshire Interoperability Lab (UNH IOL) where conformance and interoperability testing was performed.

- The core with RocketIO transceiver (all supported families) and performing the SGMII standard was tested with the Tri-speed Ethernet MAC core from Xilinx.

This was connected to an external PHY capable of performing 10BASE-T, 100BASE-T and 1000BASE-T. The system was tested at all three speeds, following the architecture shown in [Figure 13-7](#) and included the PowerPC based test platform.

Core Latency

Core Latency

The standalone core does not meet all the latency requirements specified in *IEEE 802.3* due to the latency of the Elastic Buffers in both TBI and RocketIO transceiver versions. However, the core may be used for backplane and other applications where strict adherence to the IEEE latency specification is not required.

Where strict adherence to the *IEEE 802.3* specification is required, the core may be used with an Ethernet MAC core that is within the IEEE specified latency for a MAC sublayer. For example, when the core is connected to the Xilinx 1-Gigabit Ethernet MAC core, the system as a whole is compliant with the overall *IEEE 802.3* latency specifications.

Latency for 1000BASE-X PCS with TBI

The following measurements are for the core only, and do not include any IOB registers or the Transmitter Elastic Buffer added in the example design.

Transmit Path Latency

As measured from a data octet input into `gmii_txd[7:0]` of the transmitter side GMII until that data appears on `tx_code_group[9:0]` on the TBI interface, the latency through the core in the transmit direction is 5 clock periods of `gtx_clk`.

Receive Path Latency

Measured from a data octet input into the core on `rx_code_group0[9:0]` or `rx_code_group1[9:0]` from the TBI interface (until that data appears on `gmii_rxd[7:0]` of the receiver side GMII), the latency through the core in the receive direction is equal to 16 clock periods of `gtx_clk`, plus an additional number of clock cycles equal to the current value of the Receiver Elastic Buffer.

The Receiver Elastic Buffer is 32 words deep. The nominal occupancy will be at half-full, thereby creating a nominal latency through the receiver side of the core equal to $16 + 16 = 32$ clock cycles of `gtx_clk`.

Latency for 1000BASE-X PCS and PMA Using a RocketIO Transceiver

These measurements are for the core only—they do not include the latency through the Virtex-II Pro or Virtex-4 MGT, Virtex-5 GTP transceiver, or the Transmitter Elastic Buffer added in the example design.

Transmit Path Latency

As measured from a data octet input into `gmi_i_txd[7:0]` of the transmitter side GMII (until that data appears on `txdata[7:0]` on the MGT interface), the latency through the core in the transmit direction is 4 clock periods of `userclk2`.

Receive Path Latency

As measured from a data octet input into the core on `rxdata[7:0]` from the MGT interface (until that data appears on `gmi_i_rxd[7:0]` of the receiver side GMII), the latency through the core in the receive direction is 6 clock periods of `userclk2`.

Latency for SGMII

When performing the SGMII standard, the core latency figures are identical to the Latency for 1000BASE-X PCS and PMA using the MGT. Again these figures do not include the latency through the MGT or any Elastic Buffers added in the example design.

Calculating the DCM Fixed Phase Shift Value

Requirement for DCM Phase Shifting

A DCM is used in the clock path to meet the input setup and hold requirements when using the core with a TBI (see [Chapter 6, “The Ten-Bit Interface”](#)) and with an external GMII implementation in Spartan-3, Spartan-3E, Spartan-3A, Spartan-3AN and Spartan-3A DSP devices (see [“Spartan-3, Spartan-3E and Spartan-3A Devices,”](#) page 63).

In these cases, a fixed phase shift offset is applied to the DCM to skew the clock. This will initiate a static alignment by using the clock DCM to shift the internal version of the clock so that its edges are centered on the data eye at the IOB DDR flip-flops. The ability to shift the internal clock in small increments is critical for sampling high-speed source synchronous signals such as TBI and GMII. For statically aligned systems, the DCM output clock phase offset (as set by the phase shift value) is a critical part of the system, as is the requirement that the PCB is designed with precise delay and impedance-matching for all the GMII/TBI data bus and control signals.

Determine the best DCM setting (phase shift) to ensure that the target system has the maximum system margin required to perform across voltage, temperature, and process (multiple chips) variations. Testing the system to determine the best DCM phase shift setting has the added advantage of providing a benchmark of the system margin based on the UI (unit interval or bit time).

System margin is defined as:

$$\text{System Margin (ps)} = \text{UI(ps)} * (\text{working phase shift range}/128)$$

Finding the Ideal Phase Shift Value for Your System

Xilinx cannot recommend a singular phase shift value that is effective across all hardware platforms. Xilinx does not recommend attempting to determine the phase shift setting empirically. In addition to the clock-to-data phase relationship, other factors such as package flight time (package skew) and clock routing delays (internal to the device) affect the clock-to-data relationship at the sample point (in the IOB) and are difficult to characterize.

Xilinx recommends extensive investigation of the phase shift setting during hardware integration and debugging. The phase shift settings provided in the example design UCF is a placeholder, and works successfully in back-annotated simulation of the example design.

Perform a complete sweep of phase shift settings during your initial system test. Use only positive (0 to 255) phase shift settings, and use a test range that covers a range of no less than 128, corresponding to a total 180 degrees of clock offset. This does not imply that 128

phase shift values must be tested; increments of 4 (52, 56, 60, etc.) correspond to roughly one DCM tap, and consequently provide an appropriate step size. It is not necessary to characterize areas outside the working phase shift range.

At the edge of the operating phase shift range, system behavior changes dramatically. In eight phase shift settings or fewer, the system can transition from no errors to exhibiting errors. Checking the operational edge at a step size of two (on more than one board) refines the typical operational phase shift range. Once the range is determined, choose the average of the high and low working phase shift values as the default. During the production test, Xilinx recommends that you re-examine the working range at corner case operating conditions to determine whether any adjustments to the final phase shift setting are needed.

You can use the FPGA Editor to generate the required test file set instead of resorting to multiple PAR runs. Performing the test on design files that differ only in phase shift setting prevents other variables from affecting the test results. FPGA Editor operations can even be scripted further, reducing the effort needed to perform this characterization.

1000BASE-X State Machines

This appendix is intended to serve as a reference for the basic operation of the 1000BASE-X *IEEE 802.3* clause 36 transmitter and receiver state machines.

Introduction

[Table D-1](#) illustrates the Ordered Sets defined in *IEEE 802.3* clause 36. These code group characters are inserted by the PCS Transmit Engine into the transmitted data stream, encapsulating the Ethernet frames indicated via the GMII transmit signals.

The PCS Receive Engine performs the opposite function; it uses the Ordered Sets to detect the Ethernet frames and from them creates the GMII receive signals.

Cross reference [Table D-1](#) with the remainder of this Appendix. See *IEEE 802.3* clause 36 for further information on these Orders Sets.

Table D-1: Defined Ordered Sets

Code	Ordered_Set	No. of Code-Groups	Encoding
/C/	Configuration		Alternating /C1/ and /C2/
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg ¹
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg ¹
/I/	IDLE		Correcting /I1/ Preserving /I2/
/I1/	IDLE_1	2	/K28.5/D5.6/
/I2/	IDLE_2	2	/K28.5/D16.2/
	Encapsulation		
/R/	Carrier_Extend	1	/K23.7/
/S/	Start_of_Packet	1	/K27.7/
/T/	End_of_Packet	1	/K29.7/
/V/	Error_Propagation	1	/K30.7/
1. Two data code-groups representing the Config_Reg value (contains Auto-Negotiation information)			

Start of Frame Encoding

The Even Transmission Case

Figure D-1 illustrates the translation of GMII encoding into the code-group stream performed by the PCS Transmit Engine. This stream is transmitted out of the core, either serially using the RocketIO transceiver or in parallel across the TBI.

It is important to note that the encoding of Idle periods /I2/ is constructed from a couple of code groups—the /K28.5/ character (considered the *even* position) and the /D16.2/ character (considered the *odd* position). In this example, the assertion of the `gmii_tx_en` signal of the GMII occurs in the even position. In response, the state machines insert a Start of Packet code group /S/ following the Idle (in the *even* position). This is inserted in place of the first byte of the frame preamble field.

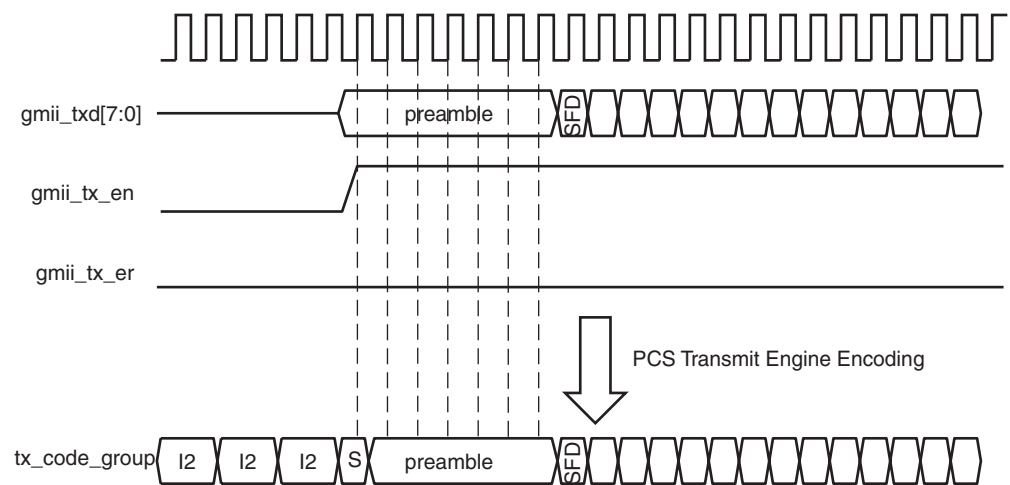


Figure D-1: 1000BASE-X Transmit State Machine Operation (Even Case)

Reception of the Even Case

Figure D-2 illustrates the reception of the in-bound code-group stream, received either serially using the RocketIO transceiver, or in parallel across the TBI, and translation of this code-group stream into the receiver GMII. This is performed by the PCS Receive Engine.

The Start of Packet code group /S/ is replaced with a preamble byte. This results in the restoration of the full preamble field.

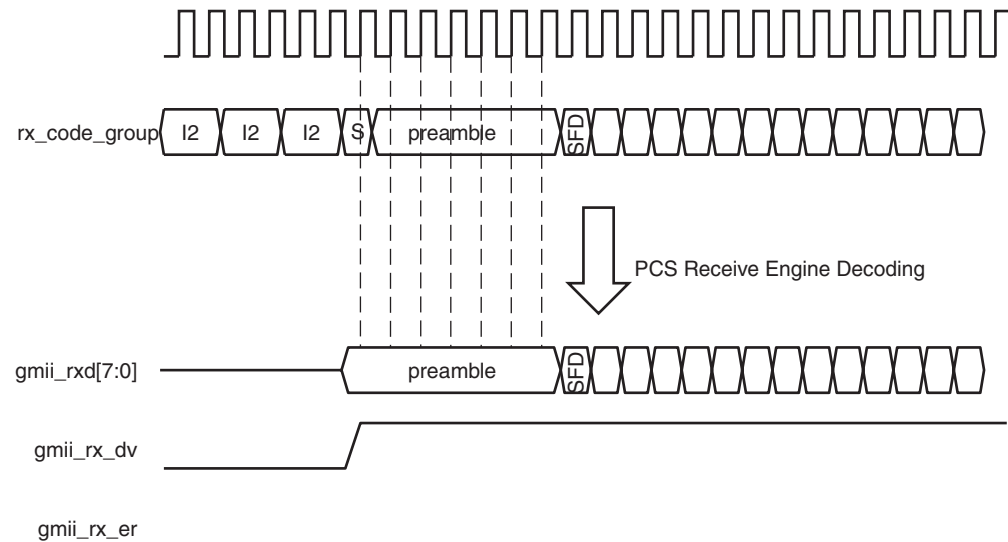


Figure D-2: 1000BASE-X Reception State Machine Operation (Even Case)

The Odd Transmission Case

Figure D-3 illustrates the translation of GMII encoding into the code-group stream performed by the PCS Transmit Engine; this stream is transmitted out of the core, either serially using the RocketIO transceiver, or in parallel across the TBI.

In this example, the assertion of the `gmii_tx_en` signal of the GMII occurs in the *odd* position; in response, the state machines are unable to immediately insert a Start-Of-Packet code group /S/ as the Idle character must first be completed. The Start of Packet code group /S/ is therefore inserted (in the *even* position) after completing the Idle. This results in the /D16.2/ character of the Idle /I2/ sequence being inserted in place of the first byte of the preamble field, and the Start-Of-Packet /S/ being inserted in place of the second byte of preamble as illustrated.




Figure D-3: 1000BASE-X Transmit State Machine Operation (Odd Case)

Reception of the Odd Case

Figure D-4 illustrates the reception of the in-bound code-group stream, received either serially using the RocketIO transceiver, or in parallel across the TBI, and translation of this code-group stream into the receiver GMII. This is performed by the PCS Receive Engine.

Note that the Start of Packet code group /S/ is again replaced with a preamble byte. However, the first preamble byte of the original transmit GMII (see Figure D-3) frame (which was replaced with the /D16.2/ character to complete the Idle sequence), has not been replaced. This has resulted in a single byte of preamble loss across the system.




Figure D-4: 1000BASE-X Reception State Machine Operation (Odd Case)

Preamble Shrinkage

As previously described, a single byte of preamble can be lost across the 1000BASE-X system (the actual loss occurs in the 1000BASE-X PCS transmitter state machine).

- There is no specific statement for this preamble loss in the *IEEE 802.3-2002* specification; the preamble loss falls out as a consequence of the state machines (see figures 36-5 and 36-6 in the *IEEE 802.3-2002* specification).
- *IEEE 802.3ah-2004* does, however, specifically state in clause 65.1.3.2.1:

“NOTE 1 – The 1000BASE-X PCS transmit function replaces the first octet of preamble with the /S/ code-group or it discards the first octet and replaces the second octet of preamble with the /S/ code-group. This decision is based upon the even or odd alignment of the PCS’s transmit state diagram (see Figure 36-5).”

End of Frame Encoding

The Even Transmission case

Figure D-5 illustrates the translation of GMII encoding into the code-group stream performed by the PCS Transmit Engine. This stream is transmitted out of the core, either serially using the RocketIO transceiver or in parallel across the TBI.

In response to the deassertion of `gmii_tx_en`, an End of Packet code group /T/ is immediately inserted. The even and odd alignment described in “[Start of Frame Encoding](#)” persists throughout the Ethernet frame. If the /T/ character occurs in the even position (the frame contained an even number of bytes starting from the /S/ character), then this is followed with a single Carrier Extend code group /R/. This allows the /K28.5/ character of the following Idle code group to be aligned to the even position.

Note: The first Idle to follow the frame termination sequence will be a /I1/ if the frame ended with positive running disparity or a /I2/ if the frame ended with negative running disparity. This is illustrated as the shaded code group.

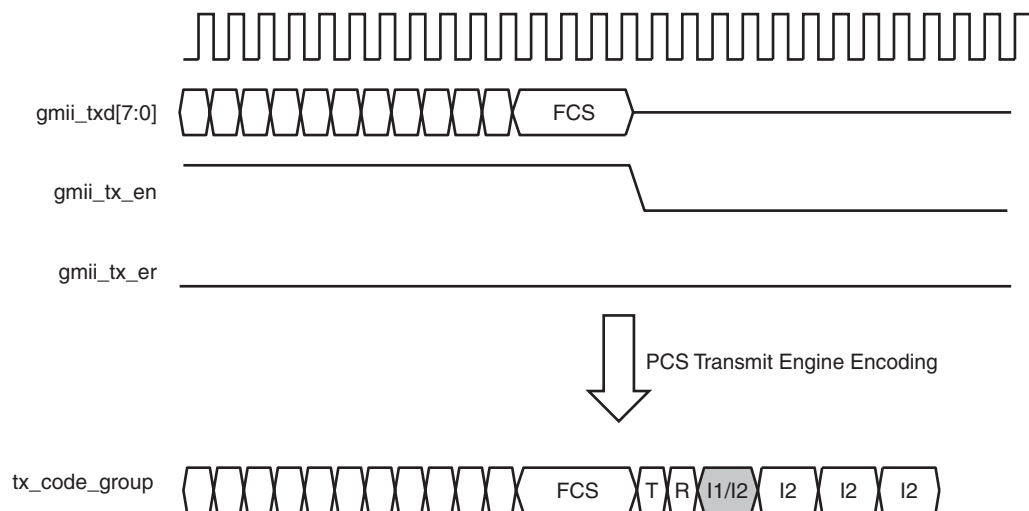


Figure D-5: 1000BASE-X Transmit State Machine Operation (Even Case)

Reception of the Even Case

Figure D-6 illustrates the reception of the in-bound code-group stream, received either serially using the RocketIO transceiver, or in parallel across the TBI, and translation of this code-group stream into the receiver GMII. This is performed by the PCS Receive Engine.




Figure D-6: 1000BASE-X Reception State Machine Operation (Even Case)

The Odd Transmission Case

Figure D-7 illustrates the translation of GMII encoding into the code-group stream performed by the PCS Transmit Engine; this stream is transmitted out of the core, either serially using the RocketIO transceiver, or in parallel across the TBI.

In response to the deassertion of `gmii_tx_en`, an End of Packet code group /T/ is immediately inserted. Note that the even and odd alignment described in “Start of Frame Encoding” persists throughout the Ethernet frame. If the /T/ character occurs in the odd position (the frame contained an odd number of bytes starting from the /S/ character), then this is followed with two Carrier Extend code groups /R/. This allows the /K28.5/ character of the following Idle code group to be aligned to the even position.

Note: The first Idle to follow the frame termination sequence will be a /I1/ if the frame ended with positive running disparity or a /I2/ if the frame ended with negative running disparity. This is illustrated as the shaded code group.

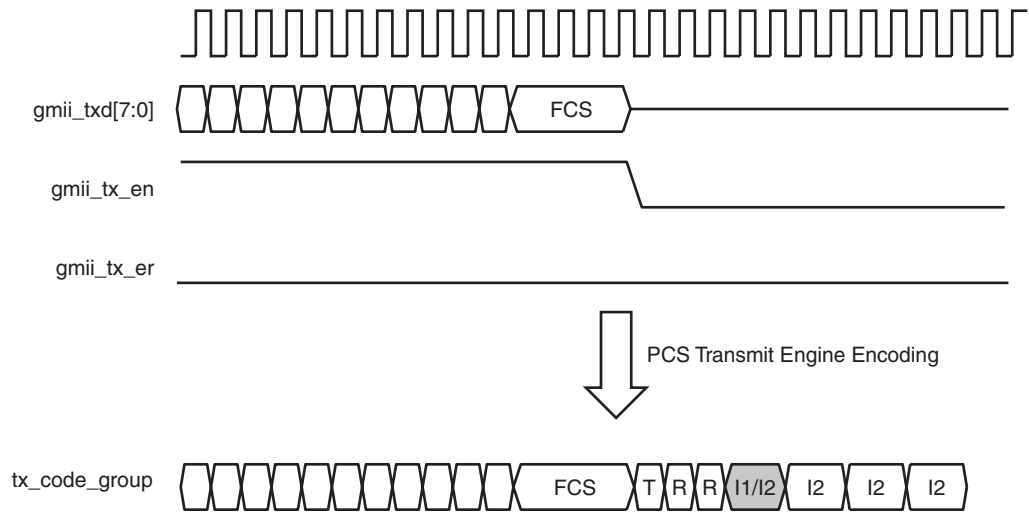


Figure D-7: 1000BASE-X Transmit State Machine Operation (Even Case)

Reception of the Odd Case

Figure D-8 illustrates the reception of the in-bound code-group stream, received either serially using the RocketIO transceiver, or in parallel across the TBI, and translation of this code-group stream into the receiver GMII. This is performed by the PCS Receive Engine.

Note that, as defined in IEEE 802.3 figure 36-7b, the combined /T/R/R/ sequence results in the GMII encoding of Frame Extension. This occurs even in full-duplex mode.




Figure D-8: 1000BASE-X Reception State Machine Operation (Odd Case)

Rx Elastic Buffer Specifications

This appendix is intended to serve as a reference for the Rx Elastic Buffer sizes used in the core, and the related maximum frame sizes that can be used without causing a buffer underflow or overflow error.

Throughout this appendix, all analyses are based on 100 ppm clock tolerances on both sides of the elastic buffer (200 ppm total difference). This corresponds to the Ethernet clock tolerance specification.

Introduction

The need for an Rx Elastic Buffer is illustrated in [“The Requirement for the FPGA Fabric Rx Elastic Buffer” in Chapter 8](#). The analysis included in this chapter shows that for standard Ethernet clock tolerances (100 ppm) there can be a maximum difference of one clock edge every 5000 clock periods of the nominal 125 MHz clock frequency.

This slight difference in clock frequency on either side of the buffer will accumulate and either start to fill or empty the Rx Elastic Buffer over time. The Rx Elastic buffer copes with this by performing clock correction during the interframe gaps by either inserting or removing Idle characters. The Rx Elastic Buffer will always attempt to restore the buffer occupancy to the half full level during an interframe gap. See [“Clock Correction,” page 224](#).

Rx Elastic Buffers: Depths and Maximum Frame Sizes

RocketIO Rx Elastic Buffers

[Figure E-1](#) illustrates the RocketIO transceiver Rx Elastic Buffer depths and thresholds in Virtex-II Pro, Virtex-4 FX, Virtex-5 LXT, SXT and FXT families. Each FIFO word corresponds to a single character of data (equivalent to a single byte of data following 8B10B decoding).




Figure E-1: Elastic Buffer Sizes for all RocketIO Transceiver Families

Virtex-II Pro and Virtex-5 Devices

Consider the Virtex-II Pro and Virtex-5 FPGA example, where the shaded area represents the usable buffer availability for the duration of frame reception.

- If the buffer is filling during frame reception, there are 52-34 = 18 FIFO locations available before the buffer reaches the overflow mark.
- If the buffer is emptying during reception, then there are 30-12 = 18 FIFO locations available before the buffer reaches the underflow mark.

This analysis assumes that the buffer is approximately at the half-full level at the start of the frame reception. As illustrated, there are two locations of uncertainty, above and below the exact half-full mark of 32, resulting from the clock correction decision, and is based across an asynchronous boundary.

Because there is a worst-case scenario of one clock edge difference every 5000 clock periods, the maximum number of clock cycles (bytes) that can exist in a single frame passing through the buffer before an error occurs is:

$$5000 \times 18 = 90000 \text{ bytes}$$

Table E-1 translates this into maximum frame size at different Ethernet speeds. At SGMII speeds lower than 1 Gbps, performance is diminished because bytes are repeated multiple times (see “Designing with Client-side GMII for the SGMII Standard” in Chapter 5).

Table E-1: Maximum Frame Sizes: RocketIO Transceiver Rx Elastic Buffers (100ppm Clock Tolerance)

Standard / Speed	Maximum Frame Size
1000BASE-X (1 Gbps only)	90000
SGMII (1 Gbps)	90000
SGMII (100 Mbps)	9000
SGMII (10 Mbps)	900

Virtex-4 FX

Consider the Virtex-4 FX case also illustrated in [Figure E-1](#). The thresholds are different to that of the Virtex-II Pro/Virtex-5 case, but the overall size of the buffer is the same. Instead of the half full point, there are configurable clock correction thresholds. During the interframe gap, clock correction will attempt to restore the occupancy to within these two thresholds.

However, by setting both `CLK_COR_MAX_LAT` and `CLK_COR_MIN_LAT` thresholds to the same value, symmetrically between overflow and underflow marks, it is possible to obtain the same figures as per the Virtex-II Pro/Virtex-5 analysis. For this reason, by adjusting the threshold attributes accordingly, [Table E-1](#) is also applicable.

SGMII Fabric Rx Elastic Buffer

Figure E-2 illustrates the alternative FPGA fabric Rx Elastic Buffer depth and thresholds in Virtex-II Pro, Virtex-4 FX and Virtex-5 LXT device families. Each FIFO word corresponds to a single character of data (equivalent to a single byte of data following 8B10B decoding). This buffer can optionally be used to replace the Rx Elastic Buffers of the RocketIO (see “Receiver Elastic Buffer Implementations” in Chapter 8).




Figure E-2: Elastic Buffer Size for all RocketIO families

The shaded area of Figure E-2 represents the usable buffer availability for the duration of frame reception.

- If the buffer is filling during frame reception, there are $122 - 66 = 56$ FIFO locations available before the buffer reaches the overflow mark.
- If the buffer is emptying during reception, then there are $62 - 6 = 56$ FIFO locations available before the buffer reaches the underflow mark.

Note that this analysis assumes the buffer is approximately at the half-full level at the start of the frame reception. As illustrated, there are two locations of uncertainty, above and below the exact half-full mark of 64. This is as a result of the clock correction decision, and is based across an asynchronous boundary.

Because there is a worst-case scenario of one clock edge difference every 5000 clock periods, the maximum number of clock cycles (bytes) that can exist in a single frame passing through the buffer before an error occurs is:

$$5000 \times 56 = 280000 \text{ bytes}$$

Table E-2 translates this into maximum frame size at different Ethernet speeds. At SGMII speeds lower than 1 Gbps, performance is diminished because bytes are repeated multiple times (see “Designing with Client-side GMII for the SGMII Standard” in Chapter 5).

Table E-2: Maximum Frame Sizes: Fabric Rx Elastic Buffers (100ppm Clock Tolerance)

Standard / Speed	Maximum Frame Size
1000BASE-X (1 Gbps only)	280000
SGMII (1 Gbps)	280000
SGMII (100 Mbps)	28000
SGMII (10 Mbps)	2800

TBI Rx Elastic Buffer

For SGMII / Dynamic Switching

The Rx Elastic Buffer used for the SGMII or Dynamic Standards Switching is identical to the method use in [“SGMII Fabric Rx Elastic Buffer.”](#)

For 1000BASE-X

[Figure E-3](#) illustrates the Rx Elastic Buffer depth and thresholds when using the Ten-Bit-Interface with the 1000BASE-X standard. This buffer is intentionally smaller than the equivalent buffer for SGMII/Dynamic Switching; because a larger size is not required, the buffer is kept smaller to save logic and keep latency low. Each FIFO word corresponds to a single character of data (equivalent to a single byte of data following 8B10B decoding).




Figure E-3: TBI Elastic Buffer Size for All Families

The shaded area of [Figure E-3](#) represents the usable buffer availability for the duration of frame reception.

- If the buffer is filling during frame reception, then there are $30 - 18 = 12$ FIFO locations available before the buffer reaches the overflow mark.
- If the buffer is emptying during reception, then there are $14 - 2 = 12$ FIFO locations available before the buffer reaches the underflow mark.

Note that this analysis assumes that the buffer is approximately at the half-full level at the start of the frame reception. As illustrated, there are two locations of uncertainty above and below the exact half-full mark of 16. This is as a result of the clock correction decision, and is based across an asynchronous boundary.

Since there is a worst-case scenario of 1 clock edge difference every 5000 clock periods, the maximum number of clock cycles (bytes) that can exist in a single frame passing through the buffer before an error occurs is:

$$5000 \times 12 = 60000 \text{ bytes}$$

This translates into a maximum frame size of 60000 bytes.

Clock Correction

The calculations in all previous sections assumes that the Rx Elastic Buffers are restored to approximately half occupancy at the start of each frame. This is achieved by the elastic buffer performing clock correction during the interframe gaps either by inserting or removing Idle characters as required.

- If the Rx Elastic Buffer is emptying during frame reception, there are no restrictions on the number of Idle characters that can be inserted due to clock correction. The occupancy will be restored to half full and the assumption holds true.
- If the Rx Elastic Buffer is filling during frame reception, Idle characters need to be removed. Restrictions that need to be considered are described in the following sections.

Idle Character Removal at 1Gbps (1000BASE-X and SGMII)

The minimum number of clock cycles that may be presented to an Ethernet receiver, according to the *IEEE 802.3* specification, is 64-bit times at any Ethernet speed. At 1 Gbps 1000BASE-X and SGMII, this corresponds to 8 bytes (8 clock cycles) of interframe gap. However, an interframe gap consists of a variety of code groups, namely /T/, /R/, /I1/ and /I2/ characters (please refer to [Appendix D, "1000BASE-X State Machines"](#)). Of these, only /I2/ can be used as clock correction characters.

In a minimum interframe gap at 1 Gbps, we can only assume that two /I2/ characters are available for removal. This corresponds to 4 bytes of data.

Looking at this from another perspective, 4 bytes of data will need to be removed in an elastic buffer (which is filling during frame reception) for a frame which is $5000 \times 4 = 20000$ bytes in length. So if the frame being received is 20000 bytes in length or shorter, at 1 Gbps, we can assume that the occupancy of the elastic buffer will always self correct to half full before the start of the subsequent frame.

For frames which are longer than 20000 bytes, the assumption that the elastic buffer will be restored to half full occupancy does not hold true. For example, for a long stream of 250000 byte frames, each separated by a minimum interframe gap, the Rx Elastic Buffer will eventually fill and overflow. This is despite the 250000 byte frame length being less than the maximum frame size calculated in the "[Rx Elastic Buffers: Depths and Maximum Frame Sizes](#)" section.

However, since the legal maximum frame size for Ethernet frames is 1522 bytes (for a VLAN frame), idle character removal restrictions are not usually an issue.

Idle Character Removal at 100 Mbps (SGMII)

At SGMII, 100 Mbps, each byte is repeated 10 times. This also applies to the interframe gap period. For this reason, the minimum of 8 bytes for the 1 Gbps case corresponds to a minimum of 80 bytes for the 100 Mbps case.

Additionally, the majority of characters in this 80-byte interframe-gap period are going to be the /I2/ clock correction characters. Because of the clock correction circuitry design, a minimum of 20 /I2/ code groups will be available for removal. This translates into 40 bytes, giving a maximum run size of $40 \times 5000 = 200000$ bytes. Because each byte at 100 Mbps is repeated ten times, this corresponds to an Ethernet frame size of 20000 bytes, the same size as the 1 Gbps case.

So in summary, at 100Mbps, for any frame size of 20000 bytes or less, it can still be assumed that the Elastic Buffer will return to half full occupancy before the start of the next frame. However, a frame size of 20000 is larger than can be received in the RocketIO Elastic Buffer (see "[Rx Elastic Buffers: Depths and Maximum Frame Sizes](#)"). Only the SGMII fabric Rx Elastic buffer is large enough.

Idle Character Removal at 10 Mbps (SGMII)

Using a similar argument to the 100 Mbps case, it can be shown that clock correction circuitry can also cope with a frame size up to 20000 bytes. However, this is larger than the maximum frame size for any Elastic Buffer provided with the core (see "[Rx Elastic Buffers: Depths and Maximum Frame Sizes](#)").

Maximum Frame Sizes for Sustained Frame Reception

Sustained frame reception refers to the maximum size of frames which can be continuously received when each frame is separated by a minimum interframe gap.

The size of frames that can be reliably received is dependent on the two considerations previously introduced in this appendix:

- The size of the Elastic Buffer, see [“Rx Elastic Buffers: Depths and Maximum Frame Sizes”](#)
- The number of clock correction characters present in a minimum interframe gap, (see [“Clock Correction”](#))

[Table E-3](#) summarizes the maximum frame sizes for sustained frame reception when used with the different Rx Elastic Buffers provided with the core. All frame sizes are provided in bytes.

Table E-3: Maximum Frame Size: (Sustained Frame Reception) Capabilities of the Rx Elastic Buffers

Ethernet Standard and Speed	Rx Elastic Buffer Type		
	TBI	RocketIO	SGMII Fabric Buffer
1000BASE-X (1 Gbps)	20000 (limited by clock correction)	20000 (limited by clock correction)	20000 (limited by clock correction)
SGMII 1 Gbps	20000 (limited by clock correction)	20000 (limited by clock correction)	20000 (limited by clock correction)
SGMII 100 Mbps	20000 (limited by clock correction)	9000 (limited by buffer size)	20000 (limited by clock correction)
SGMII 10 Mbps	2800 (limited by buffer size)	900 (limited by buffer size)	2800 (limited by buffer size)

Jumbo Frame Reception

A jumbo frame is an Ethernet frame which is deliberately larger than the maximum sized Ethernet frame allowed in the *IEEE 802.3* specification. The size of jumbo frames that can be reliably received is identical to the frame sizes defined in [“Maximum Frame Sizes for Sustained Frame Reception,”](#) page 226.

Debugging Guide

This appendix provides assistance for debugging the core within a system. For additional help, contact Xilinx by submitting a WebCase at support.xilinx.com/.

General Checks

- Ensure that all the timing constraints for the core were met during Place and Route.
- Does it work in timing simulation? If problems are seen in hardware but not in timing simulation, this could indicate a PCB issue.
- Ensure that all clock sources are clean. If using DCMs in the design, ensure that all DCMs have obtained lock by monitoring the LOCKED port.

Problems with the MDIO

- Ensure that the MDIO is driven properly. See “[MDIO Management Interface](#),” page 115 for detailed information about performing MDIO transactions.
- Check that the mdc clock is running and that the frequency is 2.5 MHz or less.
- Read from a configuration register that does not have all 0s as a default. If all 0s are read back, the read was unsuccessful. Check that the PHYAD field placed into the MDIO frame matches the value placed on the phyad[4:0] port of the core.

Problems with Data Reception or Transmission

When no data is being received or transmitted:

- Ensure that a valid link has been established between the core and its link partner, either by Auto-Negotiation or Manual Configuration: `status_vector[0]` and `status_vector[1]` should both be high. If no link has been established, see the topics discussed in the next section.
 - ◆ “[Problems with Auto-Negotiation](#)”
 - ◆ “[Problems in Obtaining a Link \(Auto-Negotiation Disabled\)](#)”

Note: Transmission through the core is not allowed unless a link has been established. This behavior can be overridden by setting the Unidirectional Enable bit.

- Ensure that the Isolate state has been disabled.

By default, the Isolate state is enabled after power-up. For an external GMII, the PHY will be electrically isolated from the GMII; for an internal GMII, it will behave as if it is isolated. This results in no data transfer across the GMII. See “[Startup Sequencing](#),” page 197 for more information.

If data is being transmitted and received between the core and its link partner, but with a high rate of packet loss, see [“Problems with a High Bit Error Rate.”](#)

Problems with Auto-Negotiation

Determine whether Auto-Negotiation has completed successfully by doing one of the following.

- Poll the Auto-Negotiation completion bit 1.5 in [“Status Register \(Register 1\)”](#)
- Use the Auto-Negotiation interrupt port of the core (see [“Using the Auto-Negotiation Interrupt,”](#) page 156)

If Auto-Negotiation is not completing:

1. Ensure that Auto-Negotiation is enabled in *both* the core and in the link partner (the device or test equipment connected to the core). Auto-Negotiation cannot complete successfully unless both devices are configured to perform Auto-Negotiation.

The Auto-Negotiation procedure requires that the Auto-Negotiation handshaking protocol between the core and its link partner, which lasts for several link timer periods, occur without a bit error. A detected bit error will cause Auto-Negotiation to go back to the beginning and restart. Therefore, a link with an exceptionally high bit error rate may not be capable of completing Auto-Negotiation, or may lead to a long Auto-Negotiation period caused by the numerous Auto-Negotiation restarts. If this appears to be the case, try the next step and see [“Problems with a High Bit Error Rate.”](#)

2. Try disabling Auto-Negotiation in both the core and the link partner and see if both devices report a valid link and are able to pass traffic. If they do, it proves that the core and link partner are otherwise configured correctly. If they do not pass traffic, see the next section, [“Problems in Obtaining a Link \(Auto-Negotiation Disabled\).”](#)

Problems in Obtaining a Link (Auto-Negotiation Disabled)

Determine whether the device has successfully obtained a link with its link partner by doing the following:

- Reading bit 1.2, Link Status, in [“Status Register \(Register 1\)”](#) when using the optional MDIO management interface (or look at `status_vector[1]`).
- Monitoring the state of `sync_acquired_status`. If this is logic ‘1,’ then synchronization, and therefore a link, has been established.

If the devices have failed to form a link then do the following:

- Ensure that Auto-Negotiation is disabled in *both* the core and in the link partner (the device or test equipment connected to the core).
 - Monitor the state of the `signal_detect` signal input to the core. This should either be:
 - ♦ connected to an optical module to detect the presence of light. Logic ‘1’ indicates that the optical module is correctly detecting light; logic ‘0’ indicates a fault. Therefore, ensure that this is driven with the correct polarity.
 - ♦ Signal must be tied to logic ‘1’ (if not connected to an optical module).
- Note:** When `signal_detect` is set to logic ‘0,’ this forces the receiver synchronization state machine of the core to remain in the loss of sync state.
- ♦ See the following section, [“Problems with a High Bit Error Rate.”](#)

RocketIO Transceiver Specific

When using a RocketIO transceiver, perform these additional checks:

- Ensure that the polarities of the TXN/TXP and RXN/RXP lines are not reversed. If they are, this can be easily fixed by using the TXPOLARITY and RXPOLARITY ports of the RocketIO.
- For Virtex-II Pro, ensure that the REF_CLK_V_SEL attribute matches the REFCLK or BREFCLK port that the clock source to which it is connected.
- Check that the RocketIO is not being held in reset by monitoring the mgt_tx_reset and mgt_rx_reset signals between the core and the RocketIO. If these are asserted then:
 - ◆ In Virtex-II Pro, this indicates that the DCM has not obtained lock.
 - ◆ In Virtex-4 and Virtex-5 this indicates that the PMA PLL circuitry in the RocketIO has not obtained lock.
- Monitor the RXBUFSTATUS [1] signal (Virtex-II Pro) or the RXBUFERR signal (Virtex-4 and Virtex-5) when Auto-Negotiation is disabled. If this is being asserted, the Elastic Buffer in the receiver path of the RocketIO is either under or overflowing. This indicates a clock correction problem caused by differences between the transmitting and receiving ends. Check all clock management circuitry and clock frequencies applied to the core and to the RocketIO.

Problems with a High Bit Error Rate

Symptoms

The severity of a high-bit error rate can vary and cause any of the following symptoms:

- Failure to complete Auto-Negotiation when Auto-Negotiation is enabled.
- Failure to obtain a link when Auto-Negotiation is disabled in both the core and the link partner.
- High proportion of lost packets when passed between two connected devices that are capable of obtaining a link through Auto-Negotiation or otherwise. This can usually be accurately measured if the Ethernet MAC attached to the core contains statistic counters.

Note: All bit errors detected by the 1000BASE-X PCS/PMA logic during frame reception will show up as Frame Check Sequence Errors in an attached Ethernet MAC.

Debugging

- Compare the problem across several devices or PCBs to ensure that the problem is not a one-off case.
- Try using an alternative link partner or test equipment and then compare results.
- Try putting the core into loopback (both by placing the core into internal loopback, and by looping back the optical cable) and compare the behavior. The core should always be capable of Auto-Negotiating with itself and looping back with itself from transmitter to receiver so direct comparisons can be made. If the core exhibits correct operation when placed into internal loopback, but not when loopback is performed via an optical cable, this may indicate a faulty optical module or a PCB problem.
- Try swapping the optical module on a misperforming device and repeat the tests.

RocketIO Transceiver Specific Checks

Perform these additional checks when using a RocketIO transceiver:

- Directly monitor the following ports of the RocketIO by attaching error counters to them, or by triggering on them using Chipscope or an external logic analyzer.

RXDISPERR

RXNOTINTABLE

These signals should not be asserted over the duration of a few seconds, minutes or even hours. If they are frequently asserted, it may indicate a problem with the RocketIO. Consult [Answer Record 19699](#) for debugging RocketIO issues.

- Place the RocketIO into parallel or serial loopback.
 - ◆ If the core exhibits correct operation in RocketIO serial loopback, but not when loopback is performed via an optical cable, it may indicate a faulty optical module.
 - ◆ If the core exhibits correct operation in RocketIO parallel loopback but not in serial loopback, this may indicate a RocketIO problem. See [Answer Record 19699](#) for details.
- A mild form of bit error rate may be solved by adjusting the transmitter TX_PREEMPHASIS, TX_DIFF_CTRL and TERMINATION_IMP attributes of the RocketIO.
- For Virtex-II Pro RocketIO only, check that the SERDES alignment logic is properly constrained to be placed near the RocketIO. See the *Virtex-II Pro RocketIO Transceiver User Guide* for more information. This constraint is not automatically adjusted for different RocketIO locations.

Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>