



Intel® Stratix® 10 External Memory Interfaces IP User Guide

Updated for Intel® Quartus® Prime Design Suite: **17.1**



Subscribe

Send Feedback

UG-S10EMI | 2017.11.06

Latest document on the web: [PDF](#) | [HTML](#)



Contents

1 Intel® Stratix® 10 EMIF IP Introduction.....	9
1.1 Intel Stratix® 10 EMIF IP Design Flow.....	10
1.2 Intel Stratix® 10 EMIF IP Design Checklist.....	11
2 Intel Stratix 10 EMIF IP Product Architecture.....	13
2.1 Intel Stratix 10 EMIF Architecture: Introduction.....	13
2.1.1 Intel Stratix 10 EMIF Architecture: I/O Subsystem.....	14
2.1.2 Intel Stratix 10 EMIF Architecture: I/O Column.....	15
2.1.3 Intel Stratix 10 EMIF Architecture: I/O SSM.....	15
2.1.4 Intel Stratix 10 EMIF Architecture: I/O Bank.....	16
2.1.5 Intel Stratix 10 EMIF Architecture: I/O Lane.....	17
2.1.6 Intel Stratix 10 EMIF Architecture: Input DQS Clock Tree.....	19
2.1.7 Intel Stratix 10 EMIF Architecture: PHY Clock Tree.....	20
2.1.8 Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks.....	20
2.1.9 Intel Stratix 10 EMIF Architecture: Clock Phase Alignment.....	21
2.2 Intel Stratix 10 EMIF Sequencer.....	22
2.2.1 Intel Stratix 10 EMIF DQS Tracking.....	23
2.3 Intel Stratix 10 EMIF Calibration.....	23
2.3.1 Intel Stratix 10 Calibration Stages	24
2.3.2 Intel Stratix 10 Calibration Stages Descriptions.....	24
2.3.3 Intel Stratix 10 Calibration Algorithms.....	25
2.3.4 Intel Stratix 10 Calibration Flowchart.....	28
2.4 Intel Stratix 10 EMIF IP Controller.....	29
2.4.1 Intel Stratix 10 Hard Memory Controller Rate Conversion Feature.....	29
2.5 Hardware Resource Sharing Among Multiple Intel Stratix 10 EMIFs.....	29
2.5.1 I/O SSM Sharing.....	29
2.5.2 I/O Bank Sharing.....	30
2.5.3 PLL Reference Clock Sharing.....	31
2.5.4 Core Clock Network Sharing.....	32
2.6 User-requested Reset in Intel Stratix 10 EMIF IP.....	32
2.7 Intel Stratix 10 EMIF for Hard Processor Subsystem.....	34
2.7.1 Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS.....	35
2.8 Intel Stratix 10 EMIF Ping Pong PHY.....	38
2.8.1 Intel Stratix 10 Ping Pong PHY Feature Description.....	38
2.8.2 Intel Stratix 10 Ping Pong PHY Architecture.....	39
2.8.3 Intel Stratix 10 Ping Pong PHY Limitations.....	41
2.8.4 Intel Stratix 10 Ping Pong PHY Calibration.....	43
2.8.5 Using the Ping Pong PHY.....	44
2.8.6 Ping Pong PHY Simulation Example Design.....	44
2.9 Intel Stratix 10 EMIF and SmartVID.....	44
3 Intel Stratix 10 EMIF IP End-User Signals.....	46
3.1 Interface and Signal Descriptions.....	46
3.1.1 Intel Stratix 10 EMIF IP Interfaces for DDR3.....	46
3.1.2 Intel Stratix 10 EMIF IP Interfaces for DDR4.....	57
3.1.3 Intel Stratix 10 EMIF IP Interfaces for LPDDR3.....	69
3.1.4 Intel Stratix 10 EMIF IP Interfaces for QDR II/II+/II+ Xtreme.....	78
3.1.5 Intel Stratix 10 EMIF IP Interfaces for QDR-IV.....	87



3.1.6 Intel Stratix 10 EMIF IP Interfaces for RLDRAM 3.....	97
3.2 AFI Signals.....	107
3.2.1 AFI Clock and Reset Signals.....	107
3.2.2 AFI Address and Command Signals.....	107
3.2.3 AFI Write Data Signals.....	108
3.2.4 AFI Read Data Signals.....	109
3.2.5 AFI Calibration Status Signals.....	110
3.2.6 AFI Tracking Management Signals.....	110
3.2.7 AFI Shadow Register Management Signals.....	111
3.3 AFI 4.0 Timing Diagrams.....	112
3.3.1 AFI Address and Command Timing Diagrams.....	112
3.3.2 AFI Write Sequence Timing Diagrams.....	115
3.3.3 AFI Read Sequence Timing Diagrams.....	123
3.3.4 AFI Calibration Status Timing Diagram.....	125
3.4 Intel Stratix 10 Memory Mapped Register (MMR) Tables.....	126
3.4.1 ctrlcfg0.....	127
3.4.2 ctrlcfg1.....	127
3.4.3 dramtiming0.....	129
3.4.4 caltiming0.....	129
3.4.5 caltiming1.....	129
3.4.6 caltiming2.....	130
3.4.7 caltiming3.....	130
3.4.8 caltiming4.....	130
3.4.9 caltiming9.....	131
3.4.10 dramaddrw.....	131
3.4.11 sideband0.....	131
3.4.12 sideband1.....	132
3.4.13 sideband2.....	132
3.4.14 sideband3.....	132
3.4.15 sideband4.....	132
3.4.16 sideband5.....	133
3.4.17 sideband6.....	133
3.4.18 sideband7.....	133
3.4.19 sideband8.....	134
3.4.20 sideband9.....	134
3.4.21 sideband10.....	134
3.4.22 sideband11.....	134
3.4.23 sideband12.....	135
3.4.24 sideband13.....	135
3.4.25 sideband14.....	136
3.4.26 dramsts.....	136
3.4.27 niosreserve0.....	136
3.4.28 niosreserve1.....	137
3.4.29 sideband16.....	137
3.4.30 ecc3: ECC Error and Interrupt Configuration.....	137
3.4.31 ecc4: Status and Error Information.....	138
3.4.32 ecc5: Address of Most Recent SBE/DBE.....	138
3.4.33 ecc6: Address of Most Recent Correction Command Dropped.....	139
3.4.34 ecc7: Extension for Address of Most Recent SBE/DBE.....	139
3.4.35 ecc8: Extension for Address of Most Recent Correction Command Dropped....	139



4 Intel Stratix 10 EMIF – Simulating Memory IP.....	140
4.1 Simulation Options.....	140
4.2 Simulation Walkthrough.....	141
4.2.1 Calibration Modes.....	142
4.2.2 Abstract PHY Simulation.....	143
4.2.3 Simulation Scripts.....	144
4.2.4 Functional Simulation with Verilog HDL.....	144
4.2.5 Functional Simulation with VHDL.....	144
4.2.6 Simulating the Example Design.....	145
5 Intel Stratix 10 EMIF IP for DDR3.....	148
5.1 Parameter Descriptions.....	148
5.1.1 Intel Stratix 10 EMIF IP DDR3 Parameters: General.....	148
5.1.2 Intel Stratix 10 EMIF IP DDR3 Parameters: Memory.....	150
5.1.3 Intel Stratix 10 EMIF IP DDR3 Parameters: Mem I/O.....	152
5.1.4 Intel Stratix 10 EMIF IP DDR3 Parameters: FPGA I/O.....	153
5.1.5 Intel Stratix 10 EMIF IP DDR3 Parameters: Mem Timing.....	154
5.1.6 Intel Stratix 10 EMIF IP DDR3 Parameters: Board.....	157
5.1.7 Intel Stratix 10 EMIF IP DDR3 Parameters: Controller.....	158
5.1.8 Intel Stratix 10 EMIF IP DDR3 Parameters: Diagnostics.....	161
5.1.9 Intel Stratix 10 EMIF IP DDR3 Parameters: Example Designs.....	163
5.2 Board Skew Equations.....	164
5.2.1 Equations for DDR3 Board Skew Parameters.....	164
5.3 Pin and Resource Planning.....	165
5.3.1 Interface Pins.....	166
5.3.2 FPGA Resources.....	170
5.3.3 Pin Guidelines for Intel Stratix 10 EMIF IP.....	171
5.4 DDR3 Board Design Guidelines.....	181
5.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices.....	182
5.4.2 Channel Signal Integrity Measurement.....	183
5.4.3 Layout Approach.....	187
5.4.4 Design Layout Guidelines.....	188
5.4.5 Package Deskew.....	198
6 Intel Stratix 10 EMIF IP for DDR4.....	202
6.1 Parameter Descriptions.....	202
6.1.1 Intel Stratix 10 EMIF IP DDR4 Parameters: General.....	202
6.1.2 Intel Stratix 10 EMIF IP DDR4 Parameters: Memory.....	204
6.1.3 Intel Stratix 10 EMIF IP DDR4 Parameters: Mem I/O.....	207
6.1.4 Intel Stratix 10 EMIF IP DDR4 Parameters: FPGA I/O.....	210
6.1.5 Intel Stratix 10 EMIF IP DDR4 Parameters: Mem Timing.....	211
6.1.6 Intel Stratix 10 EMIF IP DDR4 Parameters: Board.....	214
6.1.7 Intel Stratix 10 EMIF IP DDR4 Parameters: Controller.....	216
6.1.8 Intel Stratix 10 EMIF IP DDR4 Parameters: Diagnostics.....	218
6.1.9 Intel Stratix 10 EMIF IP DDR4 Parameters: Example Designs.....	221
6.2 Board Skew Equations.....	222
6.2.1 Equations for DDR4 Board Skew Parameters.....	222
6.3 Pin and Resource Planning.....	223
6.3.1 Interface Pins.....	224
6.3.2 FPGA Resources.....	228
6.3.3 Pin Guidelines for Intel Stratix 10 EMIF IP.....	229



6.3.4 Resource Sharing Guidelines (Multiple Interfaces).....	237
6.4 DDR4 Board Design Guidelines.....	238
6.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices.....	239
6.4.2 Channel Signal Integrity Measurement.....	240
6.4.3 Layout Approach.....	244
6.4.4 Design Layout Guidelines.....	245
6.4.5 Package Deskew.....	255
7 Intel Stratix 10 EMIF IP for QDR II/II+/II+ Xtreme.....	259
7.1 Parameter Descriptions.....	259
7.1.1 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: General.....	259
7.1.2 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: FPGA I/O.....	261
7.1.3 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Memory.....	263
7.1.4 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Mem Timing.....	263
7.1.5 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Board.....	264
7.1.6 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Controller.....	266
7.1.7 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Diagnostics.....	266
7.1.8 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Example Designs.....	269
7.2 Board Skew Equations.....	270
7.2.1 Equations for QDR II, QDR II+, and QDR II+ Xtreme Board Skew Parameters....	270
7.3 Pin and Resource Planning.....	271
7.3.1 Interface Pins.....	271
7.4 QDR II/II+/II+ Xtreme Board Design Guidelines.....	282
7.4.1 QDR II SRAM Configurations.....	282
7.4.2 General Layout Guidelines.....	284
7.4.3 QDR II Layout Guidelines.....	285
7.4.4 QDR II SRAM Layout Approach.....	286
7.4.5 Package Deskew.....	287
8 Intel Stratix 10 EMIF IP for QDR-IV.....	288
8.1 Parameter Descriptions.....	288
8.1.1 Intel Stratix 10 EMIF IP QDR-IV Parameters: General.....	288
8.1.2 Intel Stratix 10 EMIF IP QDR-IV Parameters: Memory.....	290
8.1.3 Intel Stratix 10 EMIF IP QDR-IV Parameters: Controller.....	291
8.1.4 Intel Stratix 10 EMIF IP QDR-IV Parameters: FPGA I/O.....	291
8.1.5 Intel Stratix 10 EMIF IP QDR-IV Parameters: Mem Timing.....	293
8.1.6 Intel Stratix 10 EMIF IP QDR-IV Parameters: Board.....	294
8.1.7 Intel Stratix 10 EMIF IP QDR-IV Parameters: Diagnostics.....	296
8.1.8 Intel Stratix 10 EMIF IP QDR-IV Parameters: Example Designs.....	298
8.2 Board Skew Equations.....	299
8.2.1 Equations for QDR-IV Board Skew Parameters.....	299
8.3 Pin and Resource Planning.....	300
8.3.1 Interface Pins.....	301
8.4 QDR-IV Board Design Guidelines.....	311
8.4.1 QDR-IV Layout Approach.....	312
8.4.2 General Layout Guidelines.....	313
8.4.3 QDR-IV Layout Guidelines.....	313
8.4.4 Package Deskew.....	315
9 Intel Stratix 10 EMIF IP for RLDRAM 3.....	316
9.1 Parameter Descriptions.....	316
9.1.1 Intel Stratix 10 EMIF IP RLDRAM 3 Parameters: General.....	316



9.1.2 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: Memory.....	318
9.1.3 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: FPGA I/O.....	319
9.1.4 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: Mem Timing.....	320
9.1.5 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: Board.....	321
9.1.6 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: Diagnostics.....	323
9.1.7 Intel Stratix 10 EMIF IP RLD RAM 3 Parameters: Example Designs.....	325
9.2 Board Skew Equations.....	327
9.2.1 Equations for RLD RAM 3 Board Skew Parameters.....	327
9.3 Pin and Resource Planning.....	328
9.3.1 Interface Pins.....	328
9.4 RLD RAM 3 Board Design Guidelines.....	339
9.4.1 RLD RAM 3 Configurations.....	339
9.4.2 General Layout Guidelines.....	341
9.4.3 RLD RAM 3 Layout Guidelines.....	342
9.4.4 Layout Approach.....	344
9.4.5 Package Deskew.....	345
10 Intel Stratix 10 EMIF IP for LPDDR3.....	346
10.1 Parameter Descriptions.....	346
10.1.1 Intel Stratix 10 EMIF IP LPDDR3 Parameters: General.....	346
10.1.2 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Memory.....	348
10.1.3 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Mem I/O.....	348
10.1.4 Intel Stratix 10 EMIF IP LPDDR3 Parameters: FPGA I/O.....	349
10.1.5 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Mem Timing.....	351
10.1.6 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Board.....	353
10.1.7 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Controller.....	355
10.1.8 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Diagnostics.....	357
10.1.9 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Example Designs.....	359
10.2 Board Skew Equations.....	361
10.2.1 Equations for LPDDR3 Board Skew Parameters.....	361
10.3 Pin and Resource Planning.....	362
10.3.1 Interface Pins.....	362
10.4 LPDDR3 Board Design Guidelines.....	371
10.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices.....	372
10.4.2 Channel Signal Integrity Measurement.....	373
10.4.3 Layout Approach.....	377
10.4.4 Design Layout Guidelines.....	378
10.4.5 Package Deskew.....	388
11 Intel Stratix 10 EMIF IP Timing Closure.....	392
11.1 Timing Closure	392
11.1.1 Timing Analysis.....	393
11.2 Timing Report DDR.....	396
11.3 Optimizing Timing.....	398
11.4 Early I/O Timing Estimation.....	400
11.4.1 Performing Early I/O Timing Analysis.....	401
12 Optimizing Controller Performance.....	402
12.1 Interface Standard.....	402
12.2 Bank Management Efficiency.....	403
12.3 Data Transfer.....	405
12.4 Improving Controller Efficiency.....	406



12.4.1 Auto-Precharge Commands.....	407
12.4.2 Latency.....	409
12.4.3 Calibration.....	411
12.4.4 Bank Interleaving.....	411
12.4.5 Command Queue Look-Ahead Depth.....	413
12.4.6 Additive Latency and Bank Interleaving.....	414
12.4.7 User-Controlled Refresh.....	416
12.4.8 Frequency of Operation.....	417
12.4.9 Series of Reads or Writes.....	418
12.4.10 Data Reordering.....	418
12.4.11 Starvation Control.....	419
12.4.12 Command Reordering.....	420
12.4.13 Bandwidth.....	421
13 Intel Stratix 10 EMIF IP Debugging.....	422
13.1 Interface Configuration Performance Issues.....	422
13.1.1 Interface Configuration Bottleneck and Efficiency Issues.....	423
13.2 Functional Issue Evaluation.....	423
13.2.1 Intel IP Memory Model.....	424
13.2.2 Vendor Memory Model.....	424
13.2.3 Transcript Window Messages.....	424
13.2.4 Modifying the Example Driver to Replicate the Failure.....	426
13.3 Timing Issue Characteristics.....	427
13.3.1 Evaluating FPGA Timing Issues.....	427
13.3.2 Evaluating External Memory Interface Timing Issues.....	428
13.4 Verifying Memory IP Using the Signal Tap II Logic Analyzer.....	429
13.4.1 Signals to Monitor with the Signal Tap II Logic Analyzer.....	430
13.5 Hardware Debugging Guidelines.....	431
13.5.1 Create a Simplified Design that Demonstrates the Same Issue.....	431
13.5.2 Measure Power Distribution Network.....	432
13.5.3 Measure Signal Integrity and Setup and Hold Margin.....	432
13.5.4 Vary Voltage.....	432
13.5.5 Operate at a Lower Speed.....	432
13.5.6 Determine Whether the Issue Exists in Previous Versions of Software.....	433
13.5.7 Determine Whether the Issue Exists in the Current Version of Software.....	433
13.5.8 Try A Different PCB.....	433
13.5.9 Try Other Configurations.....	434
13.5.10 Debugging Checklist.....	434
13.6 Categorizing Hardware Issues.....	435
13.6.1 Signal Integrity Issues.....	435
13.6.2 Hardware and Calibration Issues.....	437
13.7 Debugging Intel Stratix 10 EMIF IP.....	438
13.7.1 External Memory Interface Debug Toolkit.....	438
13.7.2 On-Chip Debug for Intel Stratix 10.....	439
13.7.3 Configuring Your EMIF IP for Use with the Debug Toolkit.....	439
13.7.4 Example Tcl Script for Running the EMIF Debug Toolkit.....	440
13.7.5 Using the EMIF Debug Toolkit with Intel Stratix 10 HPS Interfaces.....	441
13.7.6 Intel Stratix 10 EMIF Debugging Examples.....	441
13.8 User Interface.....	443
13.8.1 Communication.....	443
13.8.2 Setup and Use.....	444



13.9 On-Chip Debug Port for Intel Stratix 10 EMIF IP.....	447
13.9.1 EMIF On-Chip Debug Port.....	448
13.9.2 Access Protocol.....	448
13.9.3 On-Die Termination Calibration	449
13.9.4 Eye Diagram	449
13.10 Driver Margining for Intel Stratix 10 EMIF IP.....	450
13.10.1 Determining Margin.....	450
13.11 Traffic Generator 2.0.....	451
13.11.1 Configuring the Traffic Generator 2.0.....	451
13.11.2 Configurable Traffic Generator 2.0 Configuration Options.....	453
13.11.3 Running the Traffic Generator 2.0.....	458
13.11.4 Understanding the Custom Traffic Generator User Interface.....	459
13.11.5 Applying the Traffic Generator 2.0.....	463
13.11.6 Testing the EMIF Interface Using the Traffic Generator 2.0.....	466
13.11.7 The Traffic Generator 2.0 Report.....	470
13.12 Calibration Adjustment Delay Step Sizes for Intel Stratix 10 Devices.....	471
13.12.1 Addressing.....	471
13.12.2 Output and Strobe Enable Minimum and Maximum Phase Settings.....	473
14 Document Revision History for Intel Stratix 10 External Memory Interfaces IP User Guide.....	475



1 Intel® Stratix® 10 EMIF IP Introduction

Intel's fast, efficient, and low-latency external memory interface (EMIF) intellectual property (IP) cores easily interface with today's higher speed memory devices.

You can easily implement the EMIF IP core functions through the Intel® Quartus® Prime software. The Intel Quartus Prime software also provides external memory toolkits that help you test the implementation of the IP in the FPGA.

The EMIF IP provides the following components:

- A physical layer interface (PHY) which builds the data path and manages timing transfers between the FPGA and the memory device.
- A memory controller which implements all the memory commands and protocol-level requirements.

For information on the maximum speeds supported by the external memory interface IP, refer to the External Memory Interface Spec Estimator.

Intel Stratix® 10 EMIF IP Protocol and Feature Support

- Supports DDR4, DDR3, DDR3L, and LPDDR3 protocols with hard memory controller and hard PHY.
- Supports QDR-IV, QDR II + Xtreme, QDR II +, and QDR II using soft memory controller and hard PHY.
- Supports RLDRAM 3 using third-party soft controller.
- Supports UDIMM, RDIMM, LRDIMM and SODIMM memory devices.
- Supports 3D Stacked Die for DDR4 devices.
- Supports up to 4 physical ranks.
- Supports Ping Pong PHY mode, allowing two memory controllers to share command, address, and control pins.
- Supports error correction code (ECC) for both hard memory controller and soft memory controller.

Related Links

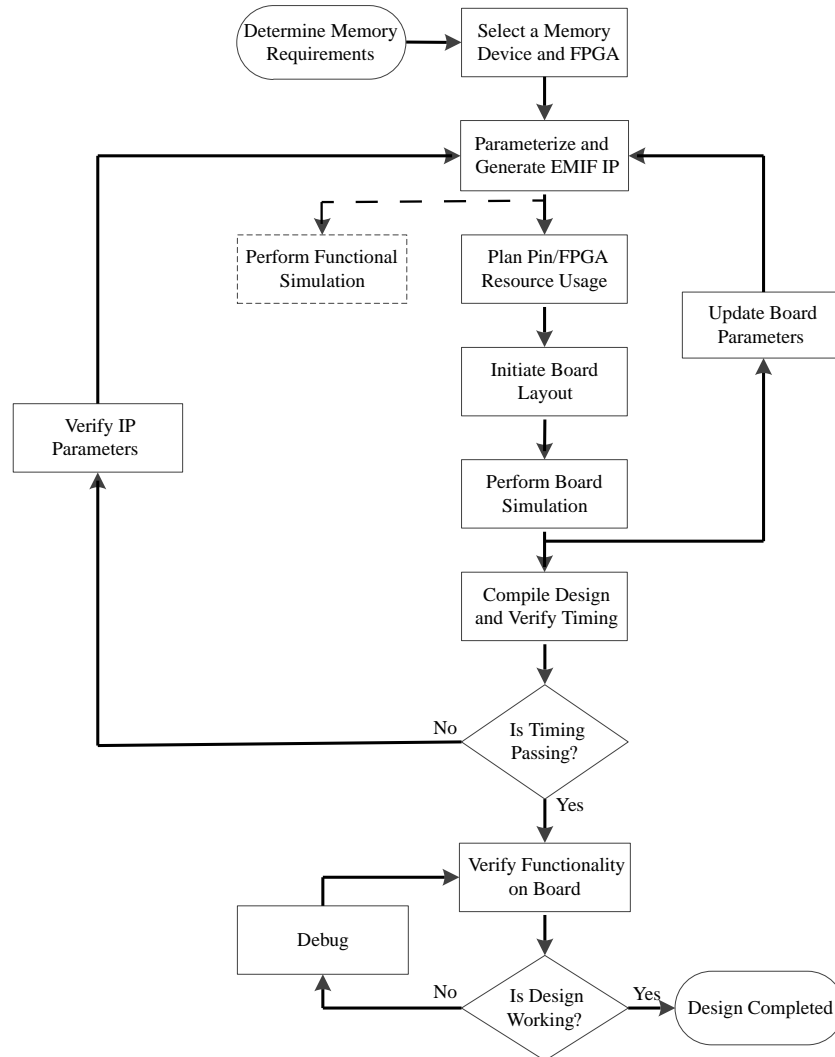
- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 General Purpose I/O User Guide](#)

1.1 Intel Stratix® 10 EMIF IP Design Flow

Intel recommends creating an example top-level file with the desired pin outs and all interface IPs instantiated. This enables the Intel Quartus Prime software to validate the design and resource allocation before PCB and schematic sign off.

The following figure shows the design flow to provide the fastest out-of-the-box experience with the EMIF IP.

Figure 1. EMIF IP Design Flow



Related Links

- [Introduction to Intel FPGA IP Cores](#)
- [Generating a Combined Simulator Setup Script](#)
- [Project Management Best Practices](#)



1.2 Intel Stratix® 10 EMIF IP Design Checklist

Refer to the following checklist as a quick reference for information about each step in the EMIF design flow.

Table 1. EMIF Design Checklist

Design Step	Description	Resources
Select an FPGA	Not all Intel FPGAs support all memory types and configurations. To help with the FPGA selection process, refer to these resources.	<ul style="list-style-type: none"> Intel FPGA Product Selector External Memory Interface Device Selector External Memory Interface Spec Estimator
Parameterize the IP	Correct IP parameterization is important for good EMIF IP operation. These resources define the memory parameters during IP generation.	<ul style="list-style-type: none"> DDR3 Parameter Descriptions DDR4 Parameter Descriptions QDR II/II+/II+ Xtreme Parameter Descriptions QDR-IV Parameter Descriptions RLDRAM 3 Parameter Descriptions LPDDR3 Parameter Descriptions
Generate initial IP and example design	After you have parameterized the EMIF IP, you can generate the IP, along with an optional example design. Refer to the Quick-Start Guide for a walkthrough of this process.	<ul style="list-style-type: none"> Design Example Quick Start Guide Design Example Description
Perform functional simulation	Simulation of the EMIF design helps to determine correct operation. These resources explain how to perform simulation and what differences exist between simulation and hardware implementation.	<ul style="list-style-type: none"> Design Example Quick Start Guide Simulating Memory IP
Make pin assignments	For guidance on pin placement, refer to these resources.	<ul style="list-style-type: none"> DDR3 Parameter Descriptions DDR4 Parameter Descriptions QDR II/II+/II+ Xtreme Parameter Descriptions QDR-IV Parameter Descriptions RLDRAM 3 Parameter Descriptions LPDDR3 Parameter Descriptions
Perform board simulation	Board simulation helps determine optimal settings for signal integrity, drive strength, as well as sufficient timing margins and eye openings. For guidance on board simulation, refer to these resources.	<ul style="list-style-type: none"> DDR3 Board Design Guidelines DDR4 Board Design Guidelines QDR II/II+/II+ Xtreme Board Design Guidelines QDR-IV Board Design Guidelines RLDRAM 3 Board Design Guidelines LPDDR3 Board Design Guidelines Board Skew Parameter Tool
Update board parameters in the IP	Board simulation is important to determine optimal settings for signal integrity, drive strength, and sufficient timing margins and eye openings. For guidance on board simulation refer to the mentioned resources.	<ul style="list-style-type: none"> DDR3 Board Design Guidelines DDR4 Board Design Guidelines QDR II/II+/II+ Xtreme Board Design Guidelines QDR-IV Board Design Guidelines RLDRAM 3 Board Design Guidelines LPDDR3 Board Design Guidelines Board Skew Parameter Tool

continued...



Design Step	Description	Resources
Verify timing closure	For information regarding compilation, system-level timing closure and timing reports refer to the Timing Closure section of this User Guide.	<ul style="list-style-type: none">• Timing Closure
Run the design on hardware	For instructions on how to program a FPGA refer to the Quick-Start Guide section of this User Guide.	<ul style="list-style-type: none">• Design Example Quick Start Guide
Debug issues with preceeding steps	Operational problems can generally be attributed to one of the following: interface configuration, pin/resource planning, signal integrity, or timing. These resources contain information on typical debug procedures and available tools to help diagnose hardware issues.	<ul style="list-style-type: none">• Debugging• External Memory Interfaces Support Center



2 Intel Stratix 10 EMIF IP Product Architecture

This chapter describes the Intel Stratix 10 product architecture.

2.1 Intel Stratix 10 EMIF Architecture: Introduction

The Intel Stratix 10 EMIF architecture contains many new hardware features designed to meet the high-speed requirements of emerging memory protocols, while consuming the smallest amount of core logic area and power.

The following are key hardware features of the Intel Stratix 10 EMIF architecture:

Hard Sequencer

The sequencer employs a hard Nios II processor, and can perform memory calibration for a wide range of protocols. You can share the sequencer among multiple memory interfaces of the same or different protocols.

Note: You cannot use the hard Nios II processor for any user applications after calibration is complete.

Hard PHY

The hard PHY in Intel Stratix 10 devices can interface with external memories running at speeds of up to 1.3 GHz. The PHY circuitry is hardened in the silicon, which simplifies the challenges of achieving timing closure and minimal power consumption.

Hard Memory Controller

The hard memory controller reduces latency and minimizes core logic consumption in the external memory interface. The hard memory controller supports the DDR3, DDR4, and LPDDR3 memory protocols.

PHY-Only Mode

Protocols that use a hard controller (DDR4, DDR3, LPDDR3, and RLDRAM 3), provide a PHY-only option, which generates only the PHY and sequencer, but not the controller. This PHY-only mode provides a mechanism by which to integrate your own custom soft controller.

High-Speed PHY Clock Tree

Dedicated high speed PHY clock networks clock the I/O buffers in Intel Stratix 10 EMIF IP. The PHY clock trees exhibit low jitter and low duty cycle distortion, maximizing the data valid window.

Automatic Clock Phase Alignment

Automatic clock phase alignment circuitry dynamically adjusts the clock phase of core clock networks to match the clock phase of the PHY clock networks. The clock phase alignment circuitry minimizes clock skew that can complicate timing closure in transfers between the FPGA core and the periphery.

Resource Sharing

The Intel Stratix 10 architecture simplifies resource sharing between memory interfaces. Resources such as the OCT calibration block, PLL reference clock pin, and core clock can be shared. The hard Nios processor in the I/O subsystem manager (I/O SSM) must be shared across all interfaces in a column.


Related Links

- [External Memory Interface Spec Estimator](#)
- [Introduction to Intel FPGA IP Cores](#)
- [Generating a Combined Simulator Setup Script](#)
- [Project Management Best Practices](#)

2.1.1 Intel Stratix 10 EMIF Architecture: I/O Subsystem

Depending on the Intel Stratix 10 device, the I/O subsystem consists of either two or three columns inside the core.

Figure 2. Stratix 10 I/O Subsystem



The I/O subsystem provides the following features:

- General-purpose I/O registers and I/O buffers
- On-chip termination control (OCT)
- I/O PLLs for external memory interfaces and user logic
- Low-voltage differential signaling (LVDS)
- External memory interface components, as follows:
 - Hard memory controller
 - Hard PHY
 - Hard Nios processor and calibration logic
 - DLL



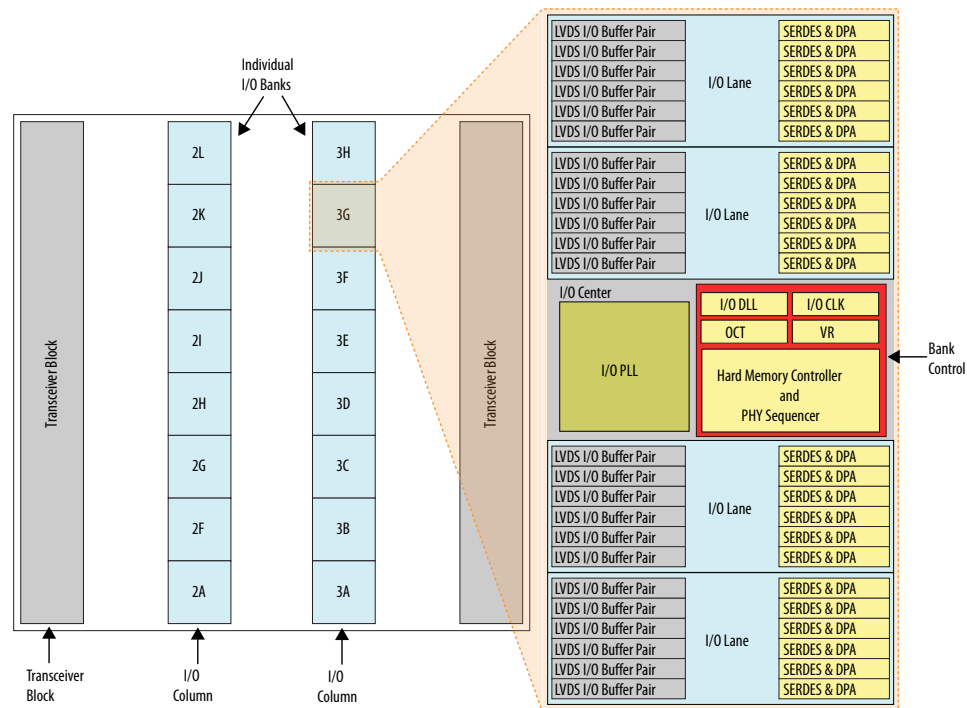
2.1.2 Intel Stratix 10 EMIF Architecture: I/O Column

Most Intel Stratix 10 devices have two I/O columns (some can have three), which contain the hardware related to external memory interfaces.

Each I/O column contains the following major parts:

- A hardened Nios processor with dedicated memory. This Nios block is referred to as the I/O SSM.
- Up to 13 I/O banks. Each I/O bank contains the hardware necessary for an external memory interface.

Figure 3. I/O Column



2.1.3 Intel Stratix 10 EMIF Architecture: I/O SSM

Each column includes one I/O subsystem manager (I/O SSM), which contains a hardened Nios II processor with dedicated memory. The I/O SSM is responsible for calibration of all the EMIFs in the column.

The I/O SSM includes dedicated memory which stores both the calibration algorithm and calibration run-time data. The hardened Nios II processor and the dedicated memory can be used only by an external memory interface, and cannot be employed for any other use. The I/O SSM can interface with soft logic, such as the debug toolkit, via an Avalon-MM bus.

The I/O SSM is clocked by an on-die oscillator, and therefore does not consume a PLL.

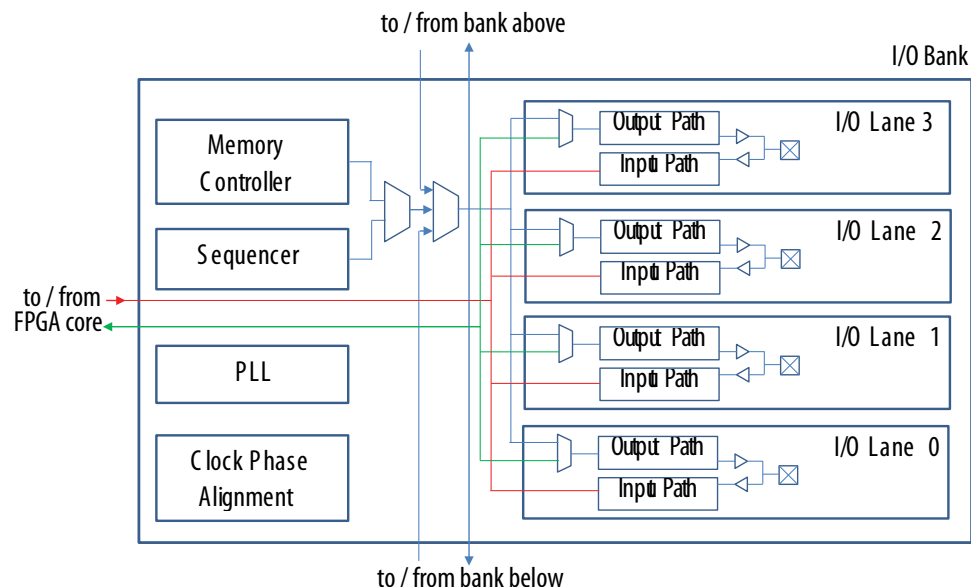
2.1.4 Intel Stratix 10 EMIF Architecture: I/O Bank

A single I/O bank contains all the hardware needed to build an external memory interface. Each I/O column contains up to 13 I/O banks; the exact number of banks depends on device size and pin package. You can make a wider interface by connecting multiple banks together.

Each I/O bank resides in an I/O column, and contains the following components:

- Hard memory controller
- Sequencer components
- PLL and PHY clock trees
- DLL
- Input DQS clock trees
- 48 pins, organized into four I/O lanes of 12 pins each

Figure 4. I/O Bank Architecture in Intel Stratix 10 Devices



I/O Bank Usage

The pins in an I/O bank can serve as address and command pins, data pins, or clock and strobe pins for an external memory interface. You can implement a narrow interface, such as a DDR3 or DDR4 x8 interface, with only a single I/O bank. A wider interface, such as x72 or x144, can be implemented by configuring multiple adjacent banks in a multi-bank interface. Any pins in a bank which are not used by the external memory interface remain available for use as general purpose I/O pins (of the same voltage standard).

Every I/O bank includes a hard memory controller which you can configure for DDR3, DDR4, or LPDDR3. In a multi-bank interface, only the controller of one bank is active; controllers in the remaining banks are turned off to conserve power.



To use a multi-bank Intel Stratix 10 EMIF interface, you must observe the following rules:

- Designate one bank as the address and command bank.
- The address and command bank must contain all the address and command pins.
- The locations of individual address and command pins within the address and command bank must adhere to the pin map defined in the pin table— regardless of whether you use the hard memory controller or not.
- If you do use the hard memory controller, the address and command bank contains the active hard controller.

All the I/O banks in a column are capable of functioning as the address and command bank. However, for minimal latency, you should select the center-most bank of the interface as the address and command bank.

2.1.5 Intel Stratix 10 EMIF Architecture: I/O Lane

An I/O bank contains 48 I/O pins, organized into four I/O lanes of 12 pins each.

Each I/O lane can implement one x8/x9 read capture group (DQS group), with two pins functioning as the read capture clock/strobe pair (DQS/DQS#), and up to 10 pins functioning as data pins (DQ and DM pins). To implement x18 and x36 groups, you can use multiple lanes within the same bank.

It is also possible to implement a pair of x4 groups in a lane. In this case, four pins function as clock/strobe pair, and 8 pins function as data pins. DM is not available for x4 groups. There must be an even number of x4 groups for each interface.

For x4 groups, DQS0 and DQS1 must be placed in the same I/O lane as a pair. Similarly, DQS2 and DQS3 must be paired. In general, DQS(x) and DQS(x+1) must be paired in the same I/O lane.

Table 2. Lanes Used Per Group

Group Size	Number of Lanes Used	Maximum Number of Data Pins per Group
x8 / x9	1	10
x18	2	22
x36	4	46
pair of x4	1	4 per group, 8 per lane

Figure 5. x4 Group

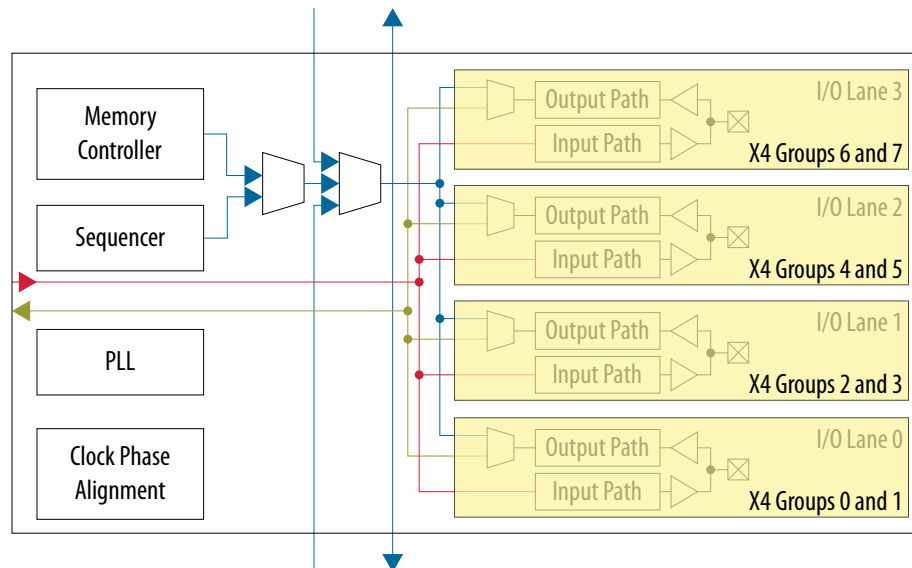


Figure 6. x8 Group

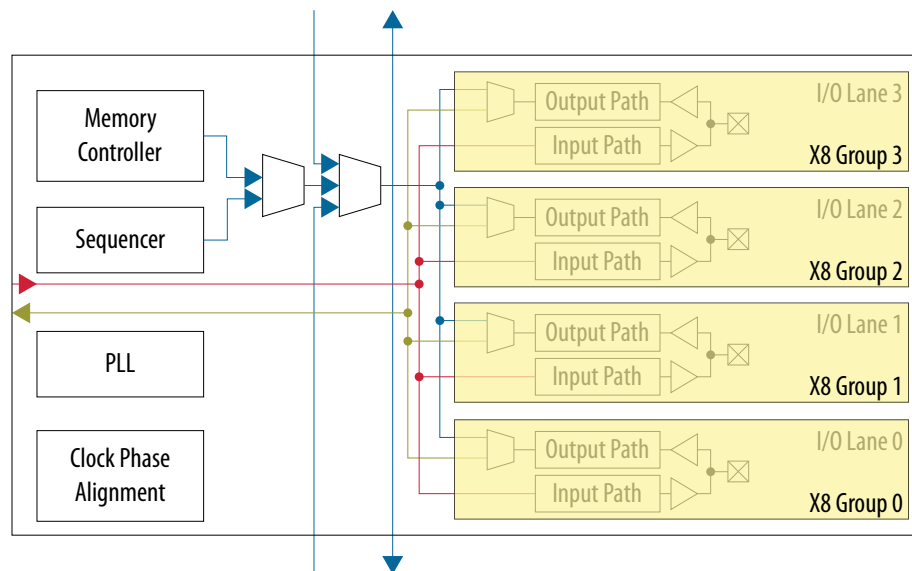


Figure 7. x18 Group

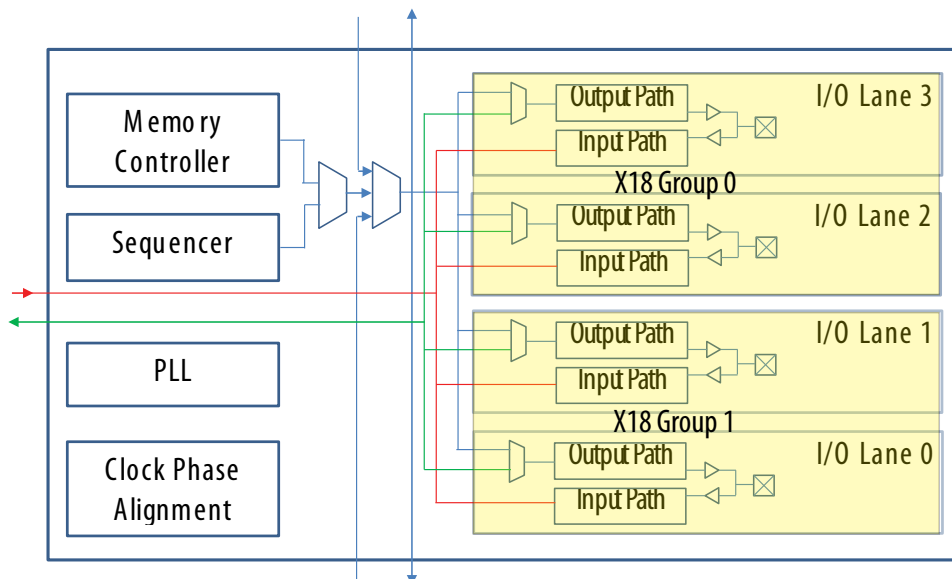
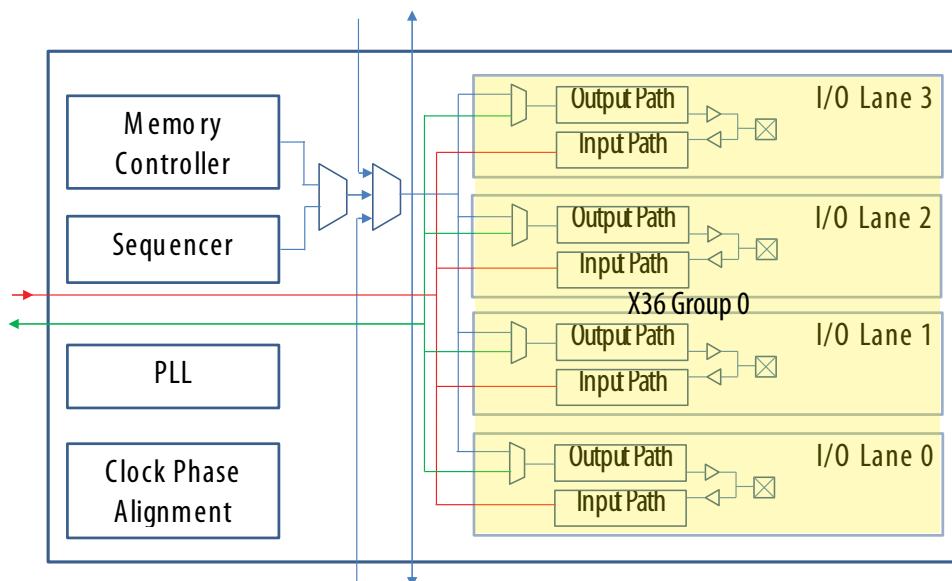


Figure 8. x36 Group



2.1.6 Intel Stratix 10 EMIF Architecture: Input DQS Clock Tree

The input DQS clock tree is a balanced clock network that distributes the read capture clock and strobe from the external memory device to the read capture registers inside the I/Os.

You can configure an input DQS clock tree in x4 mode, x8/x9 mode, x18 mode, or x36 mode.



Within every bank, only certain physical pins at specific locations can drive the input DQS clock trees. The pin locations that can drive the input DQS clock trees vary, depending on the size of the group.

Table 3. Pins Usable as Read Capture Clock / Strobe Pair

Group Size	Index of Lanes Spanned by Clock Tree	In-Bank Index of Pins Usable as Read Capture Clock / Strobe Pair	
		Positive Leg	Negative Leg
x4	0A	4	5
x4	0B	8	9
x4	1A	16	17
x4	1B	20	21
x4	2A	28	29
x4	2B	32	33
x4	3A	40	41
x4	3B	44	45
x8 / x9	0	4	5
x8 / x9	1	16	17
x8 / x9	2	28	29
x8 / x9	3	40	41
x18	0, 1	12	13
x18	2, 3	36	37
x36	0, 1, 2, 3	20	21

2.1.7 Intel Stratix 10 EMIF Architecture: PHY Clock Tree

Dedicated high-speed clock networks drive I/Os in Intel Stratix 10 EMIF. Each PHY clock network spans only one bank.

The relatively short span of the PHY clock trees results in low jitter and low duty-cycle distortion, maximizing the data valid window.

The PHY clock tree in Intel Stratix 10 devices can run as fast as 1.3 GHz. All Intel Stratix 10 external memory interfaces use the PHY clock trees.

2.1.8 Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks

Each I/O bank includes a PLL that can drive the PHY clock trees of that bank, through dedicated connections. In addition to supporting EMIF-specific functions, such PLLs can also serve as general-purpose PLLs for user logic.

Intel Stratix 10 external memory interfaces that span multiple banks use the PLL in each bank. The Intel Stratix 10 architecture allows for relatively short PHY clock networks, reducing jitter and duty-cycle distortion.

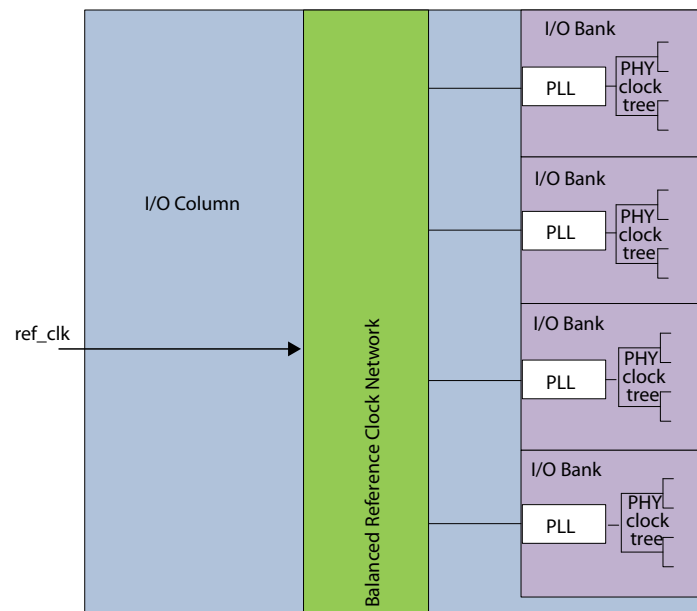


The following mechanisms ensure that the clock outputs of individual PLLs in a multi-bank interface remain in phase:

- A single PLL reference clock source feeds all PLLs. The reference clock signal reaches the PLLs by a balanced PLL reference clock tree. The Intel Quartus Prime software automatically configures the PLL reference clock tree so that it spans the correct number of banks.
- The EMIF IP sets the PLL M and N values appropriately to maintain synchronization among the clock dividers across the PLLs. This requirement restricts the legal PLL reference clock frequencies for a given memory interface frequency and clock rate. The Stratix 10 EMIF IP parameter editor automatically calculates and displays the set of legal PLL reference clock frequencies. If you plan to use an on-board oscillator, you must ensure that its frequency matches the PLL reference clock frequency that you select from the displayed list. The correct M and N values of the PLLs are set automatically based on the PLL reference clock frequency that you select.

Note: The PLL reference clock pin may be placed in the address and command I/O bank or in a data I/O bank, there is no implication on timing. However, for debug flexibility, it is recommended to place the PLL reference clock in the address and command I/O bank.

Figure 9. PLL Balanced Reference Clock Tree



Related Links

[Maximum Number of Interfaces](#) on page 169

2.1.9 Intel Stratix 10 EMIF Architecture: Clock Phase Alignment

In Intel Stratix 10 external memory interfaces, a global clock network clocks registers inside the FPGA core, and the PHY clock network clocks registers inside the FPGA periphery. Clock phase alignment circuitry employs negative feedback to dynamically adjust the phase of the core clock signal to match the phase of the PHY clock signal.

The clock phase alignment feature effectively eliminates the clock skew effect in all transfers between the core and the periphery, facilitating timing closure. All Stratix 10 external memory interfaces employ clock phase alignment circuitry.

Figure 10. Clock Phase Alignment Illustration

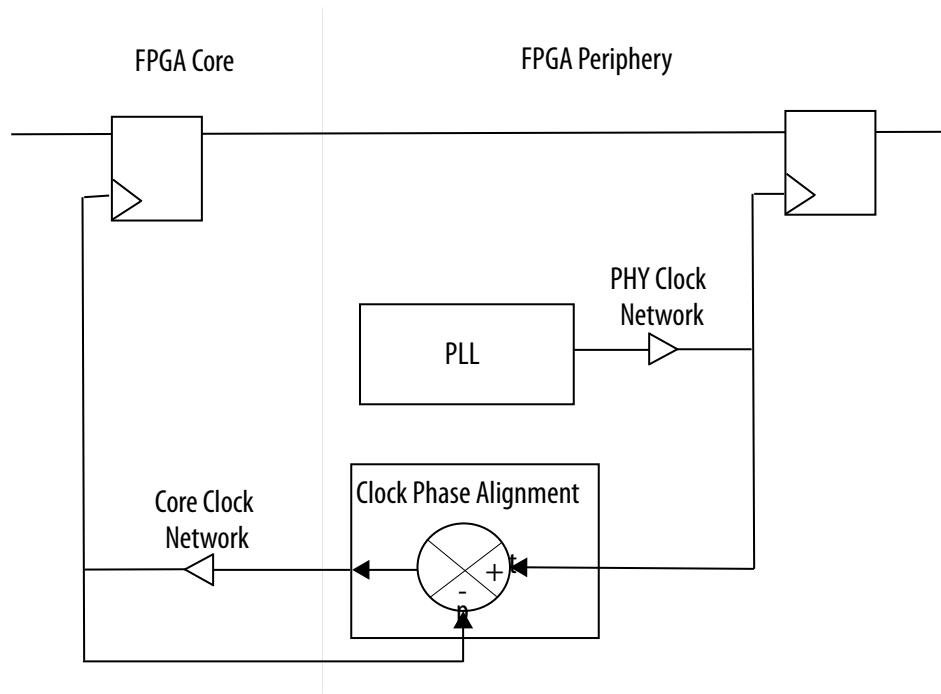
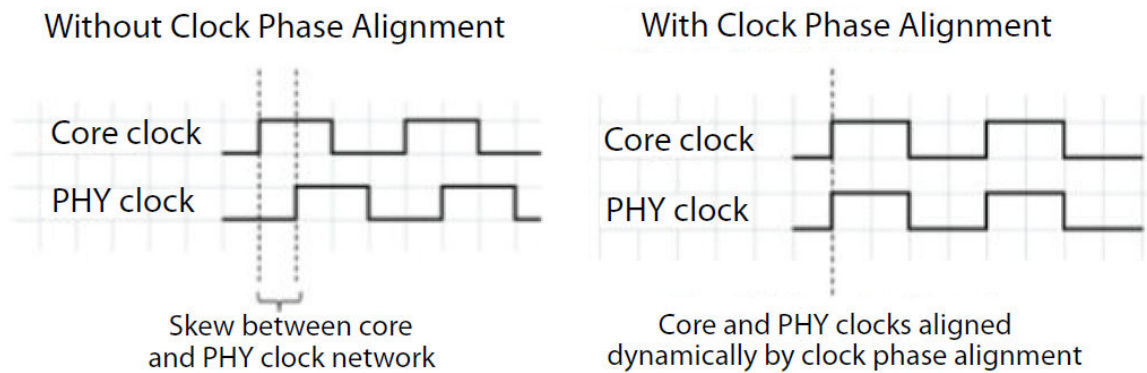


Figure 11. Effect of Clock Phase Alignment




2.2 Intel Stratix 10 EMIF Sequencer

The Intel Stratix 10 EMIF sequencer is fully hardened in silicon, with executable code to handle protocols and topologies. Hardened RAM contains the calibration algorithm.

The Intel Stratix 10 EMIF sequencer is responsible for the following operations:

- Initializes memory devices.
- Calibrates the external memory interface.
- Governs the hand-off of control to the memory controller.
- Handles recalibration requests and debug requests.
- Handles all supported protocols and configurations.

Figure 12. Intel Stratix 10 EMIF Sequencer Operation



2.2.1 Intel Stratix 10 EMIF DQS Tracking

DQS tracking is enabled for QDR II / II+ / QDR II+ Xtreme, RLDRAM 3, and LPDDR3 protocols. DQS tracking is not available for DDR3 and DDR4 protocols.

2.3 Intel Stratix 10 EMIF Calibration

The calibration process compensates for skews and delays in the external memory interface.

The following effects can be compensated for by the calibration process:

- Timing and electrical constraints, such as setup/hold time and V_{ref} variations.
- Circuit board and package factors, such as skew, fly-by effects, and manufacturing variations.
- Environmental uncertainties, such as variations in voltage and temperature.
- The demanding effects of small margins associated with high-speed operation.

Note: The calibration process is intended to maximize margins for robust EMIF operation; it cannot compensate for an inadequate PCB layout.

2.3.1 Intel Stratix 10 Calibration Stages

At a high level, the calibration routine consists of address and command calibration, read calibration, and write calibration.

The stages of calibration vary, depending on the protocol of the external memory interface.

Table 4. Calibration Stages by Protocol

Stage	DDR4	DDR3	LPDDR3	RLDRAM 3	QDR-IV	QDR II/II+
Address and command						
Leveling	Yes	Yes	—	—	—	—
Deskew	Yes	—	Yes	—	Yes	—
Read						
DQSen	Yes	Yes	Yes	Yes	Yes	Yes
Deskew	Yes	Yes	Yes	Yes	Yes	Yes
VREF-In	Yes	—	—	—	Yes	—
LFIFO	Yes	Yes	Yes	Yes	Yes	Yes
Write						
Leveling	Yes	Yes	Yes	Yes	Yes	—
Deskew	Yes	Yes	Yes	Yes	Yes	Yes
VREF-Out	Yes	—	—	—	—	—

2.3.2 Intel Stratix 10 Calibration Stages Descriptions

The various stages of calibration perform address and command calibration, read calibration, and write calibration.

Address and Command Calibration

The goal of address and command calibration is to delay address and command signals as necessary to optimize the address and command window. This stage is not available for all protocols, and cannot compensate for an inefficient board design.

Address and command calibration consists of the following parts:

- Leveling calibration— Centers the CS# signal and the entire address and command bus, relative to the CK clock. This operation is available for DDR3 and DDR4 interfaces only.
- Deskew calibration— Provides per-bit deskew for the address and command bus (except CS#), relative to the CK clock. This operation is available for DDR4 and QDR-IV interfaces only.



Read Calibration

Read calibration consists of the following parts:

- DQSen calibration— Calibrates the timing of the read capture clock gating and ungating, so that the PHY can gate and ungate the read clock at precisely the correct time—if too early or too late, data corruption can occur. The algorithm for this stage varies, depending on the memory protocol.
- Deskew calibration— Performs per-bit deskew of read data relative to the read strobe or clock.
- VREF-In calibration— Calibrates the VREF level at the FPGA.
- LFIFO calibration: Normalizes differences in read delays between groups due to fly-by, skews, and other variables and uncertainties.

Write Calibration

Write calibration consists of the following parts:

- Leveling calibration— Aligns the write strobe and clock to the memory clock, to compensate for skews, especially those associated with fly-by topology. The algorithm for this stage varies, depending on the memory protocol.
- Deskew calibration— Performs per-bit deskew of write data relative to the write strobe and clock.
- VREF-Out calibration— Calibrates the VREF level at the memory device.

2.3.3 Intel Stratix 10 Calibration Algorithms

The calibration algorithms sometimes vary, depending on the targeted memory protocol.

Address and Command Calibration

Address and command calibration consists of the following parts:



- Leveling calibration— (DDR3 and DDR4 only) Toggles the CS# and CAS# signals to send read commands while keeping other address and command signals constant. The algorithm monitors for incoming DQS signals, and if the DQS signal toggles, it indicates that the read commands have been accepted. The algorithm then repeats using different delay values, to find the optimal window.
- Deskew calibration— (DDR4, QDR-IV, and LPDDR3 only)
 - (DDR4) Uses the DDR4 address and command parity feature. The FPGA sends the address and command parity bit, and the DDR4 memory device responds with an alert signal if the parity bit is detected. The alert signal from the memory device tells the FPGA that the parity bit was received.

Deskew calibration requires use of the PAR/ALERT# pins, so you should not omit these pins from your design. One limitation of deskew calibration is that it cannot deskew ODT and CKE pins.
 - (QDR-IV) Uses the QDR-IV loopback mode. The FPGA sends address and command signals, and the memory device sends back the address and command signals which it captures, via the read data pins. The returned signals indicate to the FPGA what the memory device has captured. Deskew calibration can deskew all synchronous address and command signals.
 - (LPDDR3) Uses the LPDDR3 CA training mode. The FPGA sends signals onto the LPDDR3 CA bus, and the memory device sends back those signals that it captures, via the DQ pins. The returned signals indicate to the FPGA what the memory device has captured. Deskew calibration can deskew all signals on the CA bus. The remaining command signals (CS, CKE, and ODT) are calibrated based on the average of the deskewed CA bus.



Read Calibration

- DQSen calibration— (DDR3, DDR4, LPDDR3, RDRAMx and QDRx) DQSen calibration occurs before Read deskew, therefore only a single DQ bit is required to pass in order to achieve a successful read pass.
 - (DDR3, DDR4, and LPDDR3) The DQSen calibration algorithm searches the DQS preamble using a hardware state machine. The algorithm sends many back-to-back reads with a one clock cycle gap between. The hardware state machine searches for the DQS gap while sweeping DQSen delay values. The algorithm then increments the VFIFO value, and repeats the process until a pattern is found. The process is then repeated for all other read DQS groups.
 - (RDRAMx and QDRx) The DQSen calibration algorithm does not use a hardware state machine; rather, it calibrates cycle-level delays using software and subcycle delays using DQS tracking hardware. The algorithm requires good data in memory, and therefore relies on guaranteed writes. (Writing a burst of 0s to one location, and a burst of 1s to another; back-to-back reads from these two locations are used for read calibration.)

The algorithm enables DQS tracking to calibrate the phase component of DQS enable, and then issues a guaranteed write, followed by back-to-back reads. The algorithm sweeps DQSen values cycle by cycle until the read operation succeeds. The process is then repeated for all other read groups.
- Deskew calibration— Read deskew calibration is performed before write leveling, and must be performed at least twice: once before write calibration, using simple data patterns from guaranteed writes, and again after write calibration, using complex data patterns.

The deskew calibration algorithm performs a guaranteed write, and then sweeps `dqs_in` delay values from low to high, to find the right edge of the read window. The algorithm then sweeps `dq_in` delay values low to high, to find the left edge of the read window. Updated `dqs_in` and `dq_in` delay values are then applied to center the read window. The algorithm then repeats the process for all data pins.
- Vref-In calibration— Read Vref-In calibration begins by programming Vref-In with an arbitrary value. The algorithm then sweeps the Vref-In value from the starting value to both ends, and measures the read window for each value. The algorithm selects the Vref-In value which provides the maximum read window.
- LFIFO calibration— Read LFIFO calibration normalizes read delays between groups. The PHY must present all data to the controller as a single data bus. The LFIFO latency should be large enough for the slowest read data group, and large enough to allow proper synchronization across FIFOs.

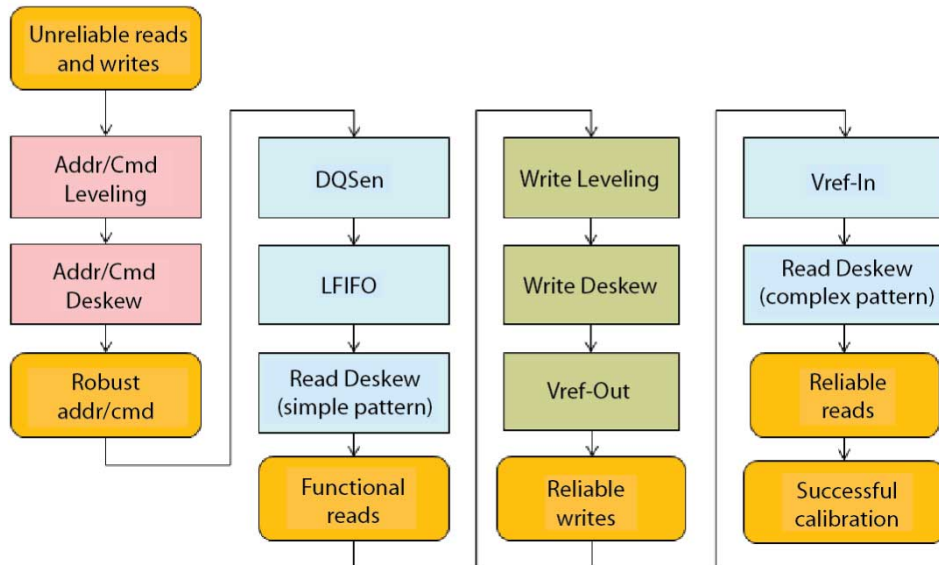
Write Calibration

- Leveling calibration— Write leveling calibration aligns the write strobe and clock to the memory clock, to compensate for skews. In general, leveling calibration tries a variety of delay values to determine the edges of the write window, and then selects an appropriate value to center the window. The details of the algorithm vary, depending on the memory protocol.
 - (DDR_x, LPDDR3) Write leveling occurs before write deskew, therefore only one successful DQ bit is required to register a pass. Write leveling staggers the DQ bus to ensure that at least one DQ bit falls within the valid write window.
 - (RLDRAM_x) Optimizes for the CK versus DK relationship.
 - (QDR-IV) Optimizes for the CK versus DK relationship. Is covered by address and command deskew using the loopback mode.
 - (QDR II/II+/Xtreme) The K clock is the only clock, therefore write leveling is not required.
- Deskew calibration— Performs per-bit deskew of write data relative to the write strobe and clock. Write deskew calibration does not change dqs_out delays; the write clock is aligned to the CK clock during write leveling.
- VREF-Out calibration— (DDR4) Calibrates the VREF level at the memory device. The VREF-Out calibration algorithm is similar to the VREF-In calibration algorithm.

2.3.4 Intel Stratix 10 Calibration Flowchart

The following flowchart illustrates the Intel Stratix 10 calibration flow.

Figure 13. Calibration Flowchart





2.4 Intel Stratix 10 EMIF IP Controller

2.4.1 Intel Stratix 10 Hard Memory Controller Rate Conversion Feature

The hard memory controller's rate conversion feature allows the hard memory controller and PHY to run at half-rate, even though user logic is configured to run at quarter-rate.

To facilitate timing closure, you may choose to clock your core user logic at quarter-rate, resulting in easier timing closure at the expense of increased area and latency. To improve efficiency and help reduce overall latency, you can run the hard memory controller and PHY at half rate.

The rate conversion feature converts traffic from the FPGA core to the hard memory controller from quarter-rate to half-rate, and traffic from the hard memory controller to the FPGA core from half-rate to quarter-rate. From the perspective of user logic inside the FPGA core, the effect is the same as if the hard memory controller were running at quarter-rate.

The rate conversion feature is enabled automatically during IP generation whenever all of the following conditions are met:

- The hard memory controller is in use.
- User logic runs at quarter-rate.
- The interface targets either an ES2 or production device.
- Running the hard memory controller at half-rate does not exceed the fMax specification of the hard memory controller and hard PHY.

When the rate conversion feature is enabled, you should see the following info message displayed in the IP generation GUI:

PHY and controller running at 2x the frequency of user logic for improved efficiency.

2.5 Hardware Resource Sharing Among Multiple Intel Stratix 10 EMIFs

Often, it is necessary or desirable to share certain hardware resources between interfaces.

2.5.1 I/O SSM Sharing

The I/O SSM contains a hard Nios-II processor and dedicated memory storing the calibration software code and data.

When a column contains multiple memory interfaces, the hard Nios-II processor calibrates each interface serially. Interfaces placed within the same I/O column always share the same I/O SSM. The Intel Quartus Prime Fitter handles I/O SSM sharing automatically.



2.5.2 I/O Bank Sharing

Data lanes from multiple compatible interfaces can share a physical I/O bank to achieve a more compact pin placement. To share an I/O bank, interfaces must use the same memory protocol, rate, frequency, I/O standard, and PLL reference clock signal.

Rules for Sharing I/O Banks

- A bank cannot serve as the address and command bank for more than one interface. This means that lanes which implement address and command pins for different interfaces cannot be allocated to the same physical bank.

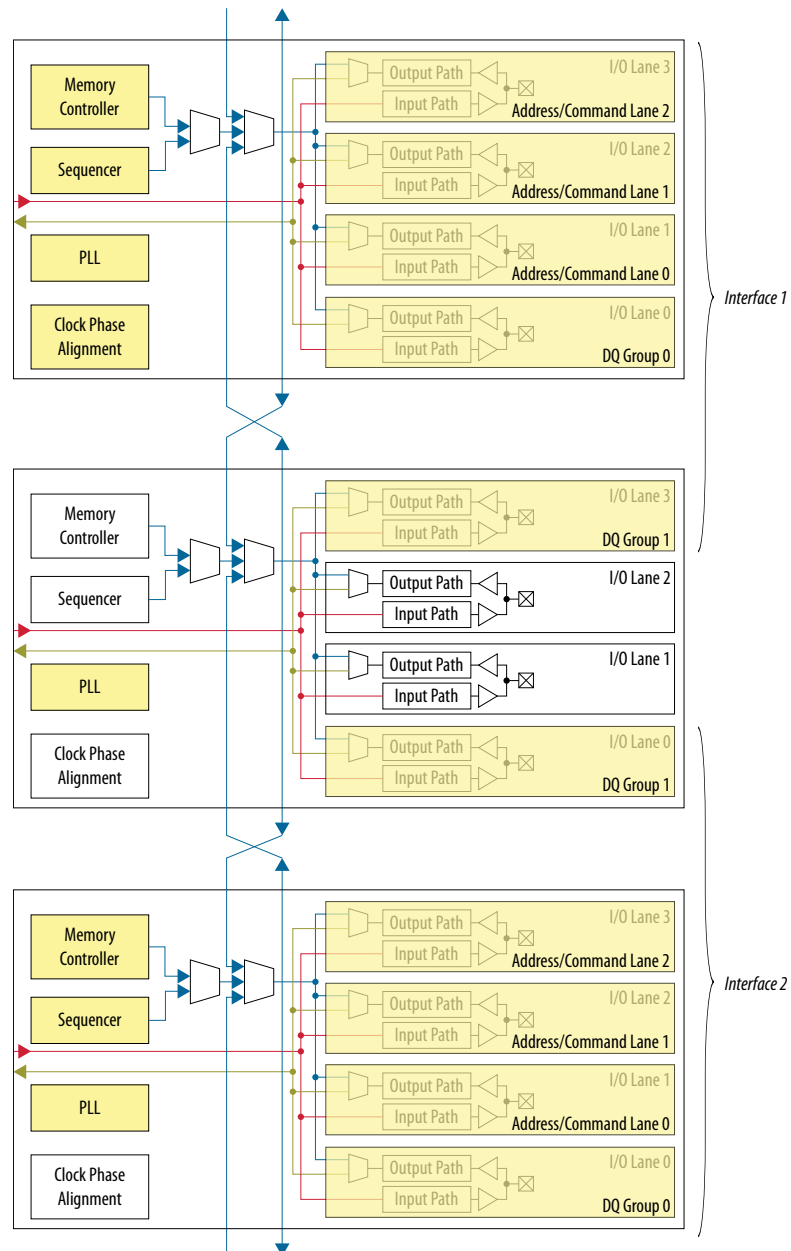
Note: An exception to the above rule exists when two interfaces are configured in a Ping-Pong PHY fashion. In such a configuration, two interfaces share the same set of address and command pins, effectively meaning that they share the same address and command tile.

- Pins within a lane cannot be shared by multiple memory interfaces.
- Pins that are not used by EMIF IP can serve as general-purpose I/Os of compatible voltage and termination settings.
- You can configure a bank as LVDS or as EMIF, but not both at the same time.
- Interfaces that share banks must reside at adjacent bank locations.

The following diagram illustrates two x16 interfaces sharing an I/O bank. The two interfaces share the same clock phase alignment block, so that one core clock signal can interact with both interfaces. Without sharing, the two interfaces would occupy a total of four physical banks instead of three.



Figure 14. I/O Bank Sharing



2.5.3 PLL Reference Clock Sharing

To implement PLL reference clock sharing, in your RTL code connect the PLL reference clock signal at your design's top-level to the PLL reference clock port of multiple interfaces.



To share a PLL reference clock, the following requirements must be met:

- Interfaces must expect a reference clock signal of the same frequency.
- Interfaces must be placed in the same column.
- Interfaces must be placed at adjacent bank locations.

2.5.4 Core Clock Network Sharing

It is often desirable or necessary for multiple memory interfaces to be accessible using a single clock domain in the FPGA core.

You might want to share core clock networks for the following reasons:

- To minimize the area and latency penalty associated with clock domain crossing.
- To minimize consumption of core clock networks.

Multiple memory interfaces can share the same core clock signals under the following conditions:

- The memory interfaces have the same protocol, rate, frequency, and PLL reference clock source.
- The interfaces reside in the same I/O column.
- The interfaces reside in adjacent bank locations.

For multiple memory interfaces to share core clocks, you must specify one of the interfaces as master and the remaining interfaces as slaves. Use the `Core clocks sharing` setting in the parameter editor to specify the master and slaves.

In your RTL, connect the `clks_sharing_master_out` signal from the master interface to the `clks_sharing_slave_in` signal of all the slave interfaces. Both the master and slave interfaces expose their own output clock ports in the RTL (e.g. `emif_usr_clk`, `afi_clk`), but the signals are equivalent, so it does not matter whether a clock port from a master or a slave is used.

Core clock sharing necessitates PLL reference clock sharing; therefore, only the master interface exposes an input port for the PLL reference clock. All slave interfaces use the same PLL reference clock signal.

2.6 User-requested Reset in Intel Stratix 10 EMIF IP

The following table summarizes information about the user-requested reset mechanism in the Intel Stratix 10 EMIF IP.

Table 5.

	Description
Reset-related signals	<code>local_reset_req</code> (input) <code>local_reset_done</code> (output)
When can user logic request a reset?	<code>local_reset_req</code> has effect only <code>local_reset_done</code> is high. After device power-on, the <code>local_reset_done</code> signal transitions high upon completion of the first calibration, whether the calibration is successful or not.
continued...	



	Description
Is user-requested reset a requirement?	A user-requested reset is optional. The I/O SSM automatically ensures that the memory interface begins from a known state as part of the device power-on sequence. A user-requested reset is necessarily only if the user logic must explicitly reset a memory interface after the device power-on sequence.
When does a user-requested reset actually happen?	A reset request is handled by the I/O SSM. If the I/O SSM receives a reset request from multiple interfaces within the same I/O column, it must serialize the reset sequence of the individual interfaces. You should avoid making assumptions on when the reset sequence will begin after a request is issued.
Timing requirement and triggering mechanism.	Reset request is sent by transitioning the <code>local_reset_req</code> signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. <code>local_reset_req</code> is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.
How long can an external memory interface be kept in reset?	It is not possible to keep an external memory interface in reset indefinitely. Asserting <code>local_reset_req</code> high continuously has no effect as a reset request is completed by a full 0->1->0 pulse.
Delaying initial calibration.	Initial calibration cannot be skipped. The <code>local_reset_done</code> signal is driven high only after initial calibration has completed.
Reset scope (within an external memory interface).	Only circuits that are required to restore EMIF to power-up state are reset. Excluded from the reset sequence are the IOSSM, the IOPLL(s), the DLL(s), and the CPA.
Reset scope (within an I/O column).	<code>local_reset_req</code> is a per-interface reset.

Method for Initiating a User-requested Reset

Step 1 - Precondition

Before asserting `local_reset_req`, user logic must ensure that the `local_reset_done` signal is high.

As part of the device power-on sequence, the `local_reset_done` signal automatically transitions to high upon the completion of the interface calibration sequence, regardless of whether calibration is successful or not.

Note: When targeting a group of interfaces that share the same core clocks, user logic must ensure that the `local_reset_done` signal of every interface is high.

Step 2 - Reset Request

After the pre-condition is satisfied, user logic can send a reset request by driving the `local_cal_req` signal from low to high and then low again (that is, by sending a pulse of 1).

- The 0-to-1 and 1-to-0 transitions need not happen in relation to any clock edges (that is, they can occur asynchronously); however, the pulse must meet a minimum pulse width of at least 2 EMIF core clock cycles. For example, if the `emif_usr_clk` has a period of 4ns, then the `local_reset_req` pulse must last at least 8ns (that is, two `emif_usr_clk` periods).
- The reset request is considered complete only after the 1-to-0 transition. The EMIF IP does not initiate the reset sequence when the `local_reset_req` is simply held high.
- Additional pulses to `local_reset_req` are ignored until the reset sequence is completed.

Optional - Detecting `local_reset_done` deassertion and assertion

If you want, you can monitor the status of the `local_reset_done` signal to to explicitly detect the status of the reset sequence.

- After the EMIF IP receives a reset request, it deasserts the `local_reset_done` signal. After initial power-up calibration, `local_reset_done` is de-asserted only in response to a user-requested reset. The reset sequence is imminent when `local_reset_done` has transitioned to low, although the exact timing depends on the current state of the I/O SSM. As part of the EMIF reset sequence, the core reset signal (`emif_usr_reset_n`, `afi_reset_n`) is driven low. Do not use a register reset by the core reset signal to sample `local_reset_done`.
- After the reset sequence has completed, `local_reset_done` is driven high again. `local_reset_done` being driven high indicates the completion of the reset sequence and the readiness to accept a new reset request; however, it does not imply that calibration was successful or that the hard memory controller is ready to accept requests. For these purposes, user logic must check signals such as `afi_cal_success`, `afi_cal_fail`, and `amm_ready`.

2.7 Intel Stratix 10 EMIF for Hard Processor Subsystem

The Intel Stratix 10 EMIF IP can enable the Intel Stratix 10 Hard Processor Subsystem (HPS) to access external DRAM memory devices.

To enable connectivity between the Intel Stratix 10 HPS and the Intel Stratix 10 EMIF IP, you must create and configure an instance of the Intel Stratix 10 External Memory Interface for HPS IP core, and use Platform Designer to connect it to the Intel Stratix 10 Hard Processor Subsystem instance in your system.



Supported Modes

The Intel Stratix 10 Hard Processor Subsystem is compatible with the following external memory configurations:

Table 6. Intel Stratix 10 Hard Processor Subsystem Compatibility

Protocol	DDR3, DDR4, LPDDR3
Maximum memory clock frequency	DDR3: 1.067 GHz DDR4: 1.333 GHz LPDDR3: 800 MHz
Configuration	Hard PHY with hard memory controller
Clock rate of PHY and hard memory controller	Half-rate
Data width (without ECC)	16-bit, 32-bit, 64-bit
Data width (with ECC)	24-bit, 40-bit, 72-bit
DQ width per group	x8
Maximum number of I/O lanes for address/command	3
Memory format	Discrete, UDIMM, SODIMM, RDIMM
Ranks / CS# width	Up to 2

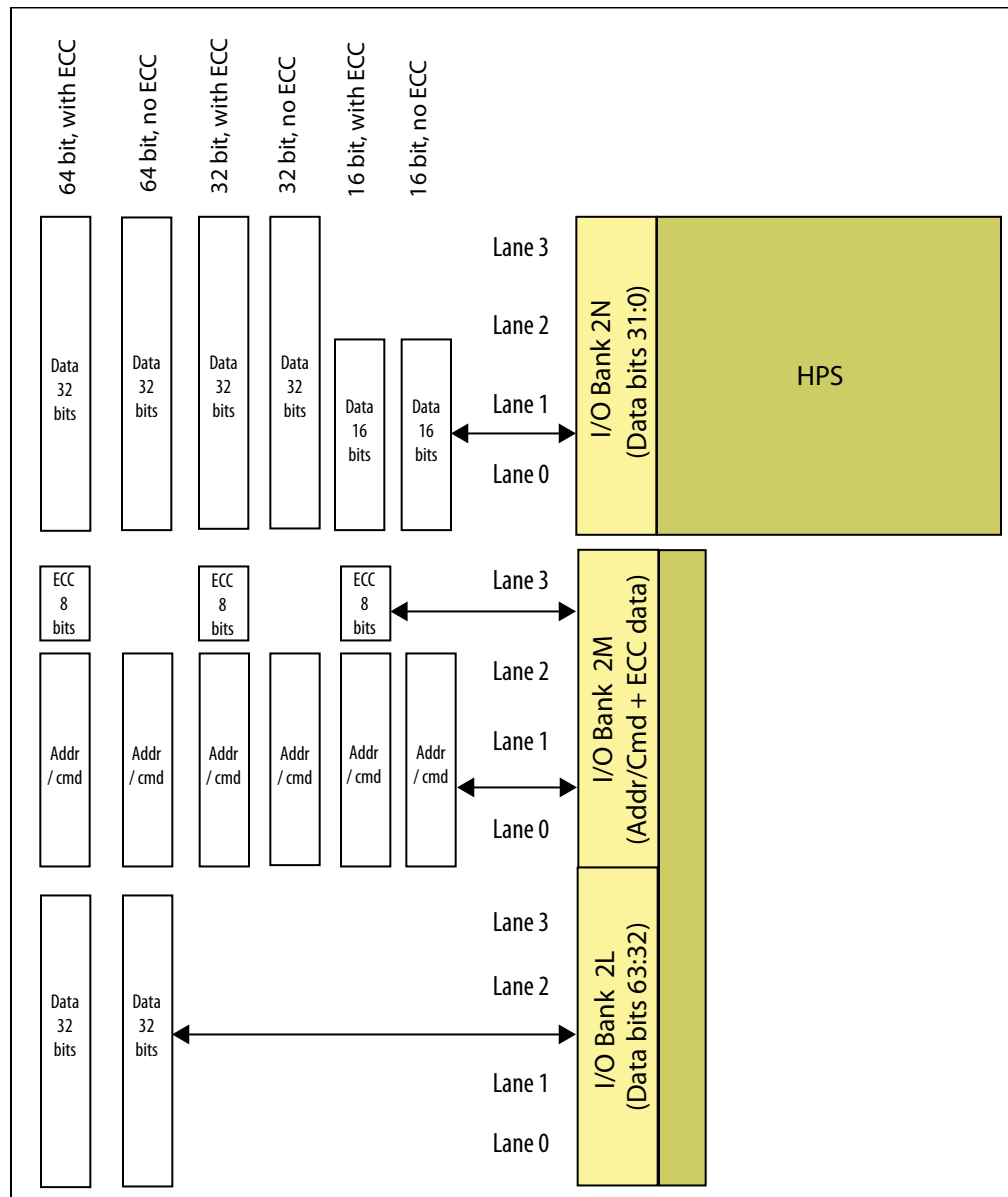
2.7.1 Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS

You can use only certain Intel Stratix 10 I/O banks to implement Intel Stratix 10 EMIF IP with the Intel Stratix 10 Hard Processor Subsystem.

The restrictions on I/O bank usage result from the Intel Stratix 10 HPS having hard-wired connections to the EMIF circuits in the I/O banks closest to the HPS. For any given EMIF configuration, the pin-out of the EMIF-to-HPS interface is fixed.

The following diagram illustrates the use of I/O banks and lanes for various EMIF-HPS data widths:

Figure 15. Intel Stratix 10 External Memory Interfaces I/O Bank and Lanes Usage



The HPS EMIF uses the closest located external memory interfaces I/O banks to connect to SDRAM. These banks include:

- Bank 2N—used for data I/Os (Data bits 31:0)
- Bank 2M—used for address, command and ECC data I/Os
- Bank 2L—used for data I/Os (Data bits 63:32)



If no HPS EMIF is used in a system, the entire HPS EMIF bank can be used as FPGA GPIO. If there is a HPS EMIF in a system, the unused HPS EMIF pins can be used as FPGA general I/O with restrictions:

- Bank 2M:
 - Lane 3 is used for SDRAM ECC data. Unused pins in lane 3 can be used as FPGA inputs only.
 - Lanes 2, 1, and 0 are used for SDRAM address and command. Unused pins in these lanes can be used as FPGA inputs or outputs.
- Bank 2N and Bank 2L :
 - Lanes 3, 2, 1, and 0 are used for data bits.
 - With 64-bit data widths, unused pins in these banks can be used as FPGA inputs only.
 - With 32-bit data widths, unused pins in Bank 2N can be used as FPGA inputs only. Unused pins for Bank 2L can be used as FPGA inputs or outputs.
 - With 16-bit data widths, Intel Quartus Prime assigns lane 0 and lane 1 as data lanes in bank 2N. Unused pins in lane 0 and lane 1 can be used as FPGA inputs only. The other two lanes are available to use as FPGA inputs or outputs.

By default, the Intel Stratix 10 External Memory Interface for HPS IP core together with the Intel Quartus Prime Fitter automatically implement the correct pin-out for HPS EMIF without you having to apply additional constraints. If you must modify the default pin-out for any reason, you must adhere to the following requirements, which are specific to HPS EMIF:

1. Within a single data lane (which implements a single x8 DQS group):
 - DQ pins must use pins at indices 1, 2, 3, 6, 7, 8, 9, 10. You may swap the locations between the DQ bits (that is, you may swap location of DQ[0] and DQ[3]) so long as the resulting pin-out uses pins at these indices only.
 - DM/DBI pin must use pin at index 11. There is no flexibility.
 - DQS/DQS# must use pins at index 4 and 5. There is no flexibility.
2. Assignment of data lanes must be as illustrated in the above figure. You are allowed to swap the locations of entire byte lanes (that is, you may swap locations of byte 0 and byte 3) so long as the resulting pin-out uses only the lanes permitted by your HPS EMIF configuration, as shown in the above figure.
3. You must not change placement of the address and command pins from the default.
4. You may place the `alert#` pin at any available pin location in either a data lane or an address and command lane.

To override the default generated pin assignments, comment out the relevant HPS_LOCATION assignments in the .qip file, and add your own location assignments (using `set_location_assignment`) in the .qsf file.

2.8 Intel Stratix 10 EMIF Ping Pong PHY


Ping Pong PHY allows two memory interfaces to share the address and command bus through time multiplexing. Compared to having two independent interfaces that allocate address and command lanes separately, Ping Pong PHY achieves the same throughput with fewer resources, by sharing the address and command lanes.

In Intel Stratix 10 EMIF, Ping Pong PHY supports both half-rate and quarter-rate interfaces for DDR3, and quarter-rate for DDR4.

2.8.1 Intel Stratix 10 Ping Pong PHY Feature Description


Conventionally, the address and command buses of a DDR3 or DDR4 half-rate or quarter-rate interface use 2T time—meaning that commands are issued for two full-rate clock cycles, as illustrated below.

Figure 16. 2T Command Timing



With the Ping Pong PHY, address and command signals from two independent controllers are multiplexed onto shared buses by delaying one of the controller outputs by one full-rate clock cycle. The result is 1T timing, with a new command being issued on each full-rate clock cycle. The following figure shows address and command timing for the Ping Pong PHY.

The command signals CS, ODT, and CKE have two signals (one for ping and one for pong); the other address and command signals are shared.

**Figure 17. 1T Command Timing Use by Ping Pong PHY**

2.8.2 Intel Stratix 10 Ping Pong PHY Architecture

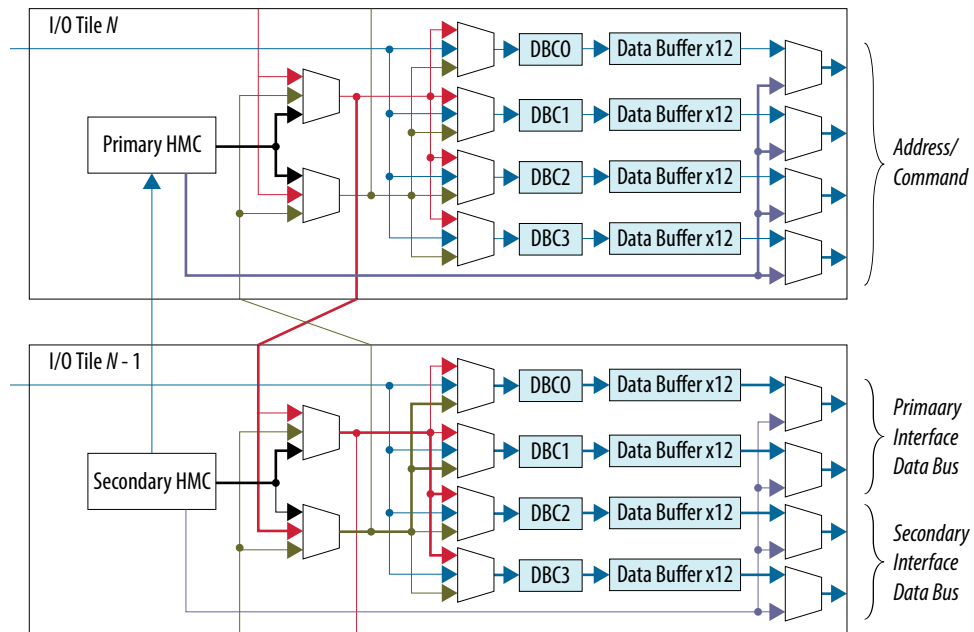
In Intel Stratix 10 EMIF, the Ping Pong PHY feature can be enabled only with the hard memory controller, where two hard memory controllers are instantiated—one for the primary interface and one for the secondary interface.

The hard memory controller I/O bank of the primary interface is used for address and command and is always adjacent and above the hard memory controller bank of the secondary interface. All four lanes of the primary hard memory controller bank are used for address and command.

The following example shows a 2x16 Ping Pong PHY bank-lane configuration. The upper bank (I/O bank N) is the address and command bank, which serves both the primary and secondary interfaces. The primary hard memory controller is linked to the secondary interface by the Ping Pong bus. The lower bank (I/O bank N-1) is the secondary interface bank, which carries the data buses for both primary and secondary interfaces. In the 2x16 case a total of four I/O banks are required for data, hence two banks in total are sufficient for the implementation.

The data for the primary interface is routed down to the top two lanes of the secondary I/O bank, and the data for the secondary interface is routed to the bottom two lanes of the secondary I/O bank.

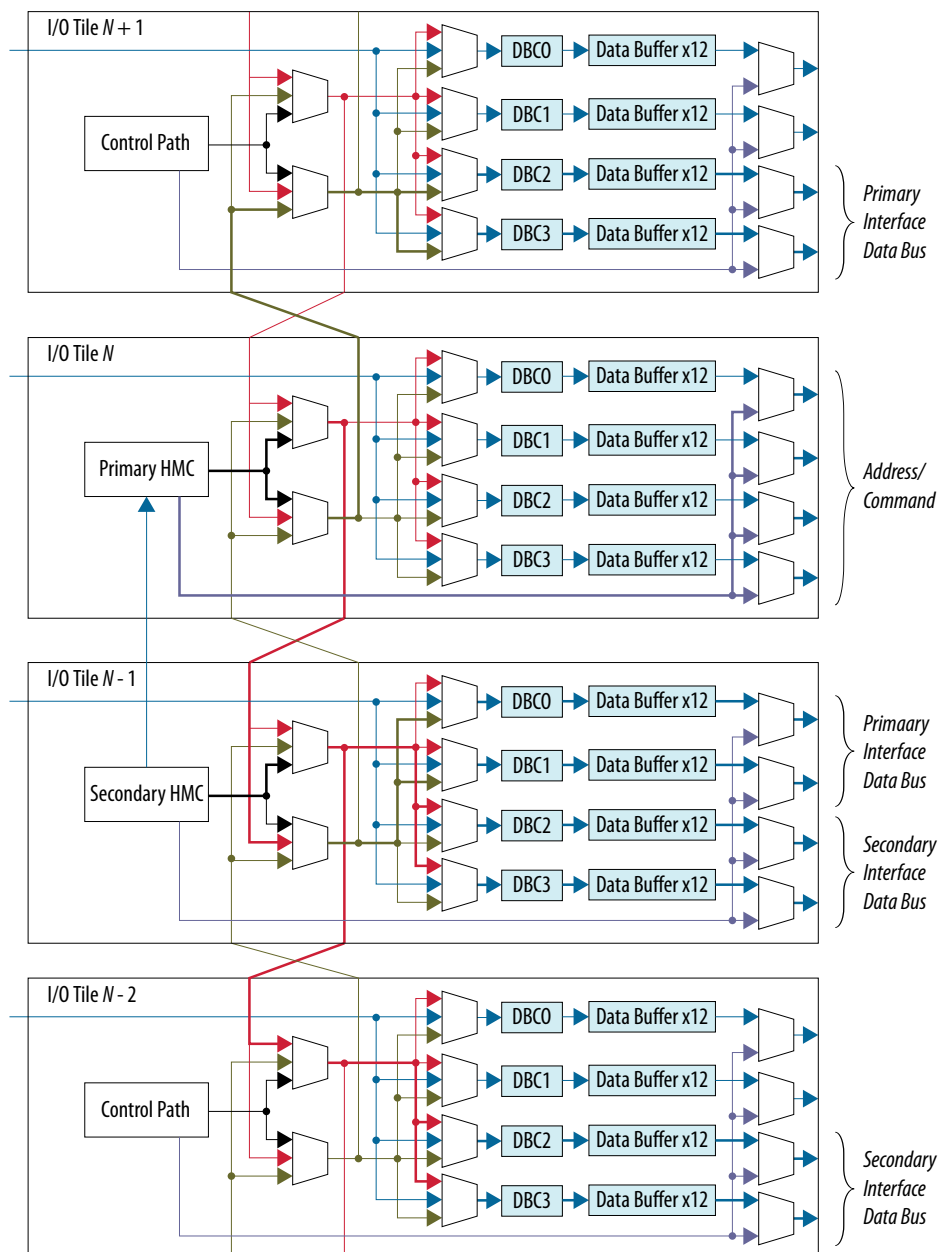
Figure 18. 2x16 Ping Pong PHY I/O Bank-Lane Configuration



A 2x32 interface can be implemented similarly, with the additional data lanes placed above and below the primary and secondary I/O banks, such that primary data lanes are placed above the primary bank and secondary data lanes are placed below the secondary bank.



Figure 19. 2x32 Ping Pong PHY I/O Bank-Lane Configuration.



2.8.3 Intel Stratix 10 Ping Pong PHY Limitations

Ping Pong PHY supports up to two ranks per memory interface. In addition, the maximum data width is x72, which is half the maximum width of x144 for a single interface.

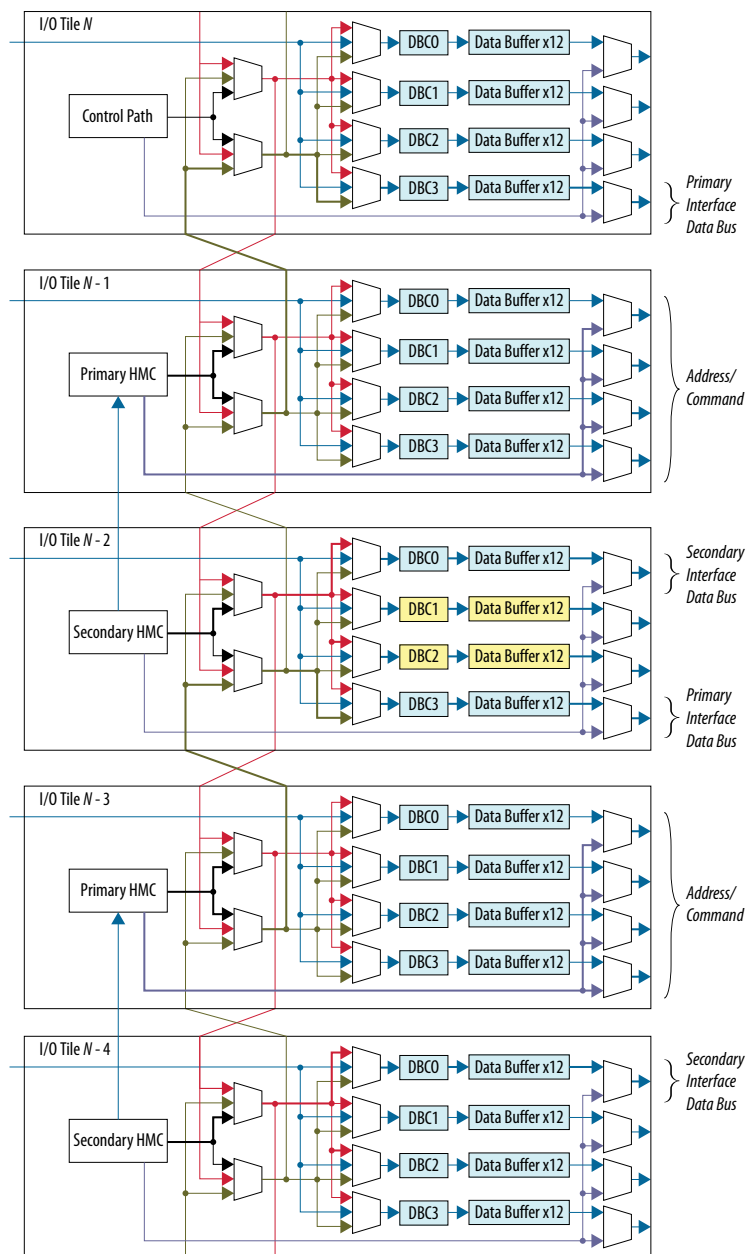


Ping Pong PHY uses all lanes of the address and command I/O bank as address and command. For information on the pin allocations of the DDR3 and DDR4 address and command I/O bank, refer to *DDR3 Scheme 1* and *DDR4 Scheme 3*, in *External Memory Interface Pin Information for Stratix 10 Devices*, on www.altera.com.

An additional limitation is that I/O lanes may be left unused when you instantiate multiple pairs of Ping Pong PHY interfaces. The following diagram shows two pairs of x8 Pin Pong controllers (a total of 4 interfaces). Lanes highlighted in yellow are not driven by any memory interfaces (unused lanes and pins can still serve as general purpose I/Os). Even with some I/O lanes left unused, the Ping Pong PHY approach is still beneficial in terms of resource usage, compared to independent interfaces. Memory widths of 24 bits and 40 bits have a similar situation, while 16 bit, 32 bit, and 64 bit memory widths do not suffer this limitation.



Figure 20. Two Pairs of x8 Pin-Pong PHY Controllers



2.8.4 Intel Stratix 10 Ping Pong PHY Calibration

A Ping Pong PHY interface is calibrated as a regular interface of double width.

Calibration of a Ping Pong PHY interface incorporates two sequencers, one on the primary hard memory controller I/O bank, and one on the secondary hard memory controller I/O bank. To ensure that the two sequencers issue instructions on the same memory clock cycle, the Nios II processor configures the sequencer on the primary hard memory controller to receive a token from the secondary interface, ignoring any

commands from the Avalon bus. Additional delays are programmed on the secondary interface to allow for the passing of the token from the sequencer on the secondary hard memory controller tile to the sequencer on the primary hard memory controller tile. During calibration, the Nios II processor assumes that commands are always issued from the sequencer on the primary hard memory controller I/O bank. After calibration, the Nios II processor adjusts the delays for use with the primary and secondary hard memory controllers.

2.8.5 Using the Ping Pong PHY

The following steps describe how to use the Ping Pong PHY for Intel Stratix 10 EMIF.

1. Configure a single memory interface according to your requirements.
2. Select **Instantiate two controllers sharing a Ping Pong PHY** on the **General** tab in the parameter editor.

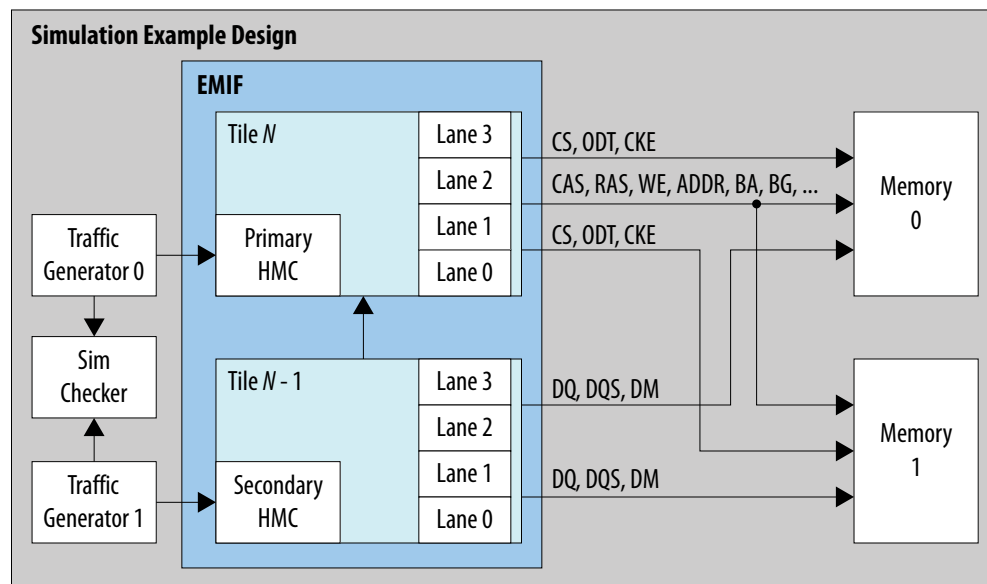
The Intel Quartus Prime software replicates the interface, resulting in two memory controllers and a shared PHY. The system configures the I/O bank-lane structure, without further input from you.

2.8.6 Ping Pong PHY Simulation Example Design

The following figure illustrates a top-level block diagram of a generated Ping Pong PHY simulation example design, using two I/O banks.

Functionally, the IP interfaces with user traffic separately, as it would with two independent memory interfaces. You can also generate synthesizable example designs, where the external memory interface IP interfaces with a traffic generator.

Figure 21. Ping Pong PHY Simulation Example Design



2.9 Intel Stratix 10 EMIF and SmartVID

Intel Stratix 10 EMIF IP can be used with the SmartVID voltage management system, to achieve reduced power consumption.



The SmartVID controller allows the FPGA to operate at a reduced V_{cc} , while maintaining performance. Because the SmartVID controller can adjust V_{cc} up or down in response to power requirements and temperature, it can have an impact on external memory interface performance. When used with the SmartVID controller, the EMIF IP implements a handshake protocol to ensure that EMIF calibration does not begin until after voltage adjustment has completed.

In extended speed grade devices, voltage adjustment occurs once when the FPGA is powered up, and no further voltage adjustments occur. The external memory calibration occurs after this initial voltage adjustment is completed. EMIF specifications are expected to be slightly lower in extended speed grade devices using SmartVID, than in devices not using SmartVID.

In industrial speed grade devices, voltage adjustment occurs at power up, and may also occur during operation, in response to temperature changes. External memory interface calibration does not occur until after the initial voltage adjustment at power up. However, the external memory interface is not recalibrated in response to subsequent voltage adjustments that occur during operation. As a result, EMIF specifications for industrial speed grade devices using SmartVID are expected to be lower than for extended speed grade devices.

Using Intel Stratix 10 EMIF IP with SmartVID

To employ Intel Stratix 10 EMIF IP with SmartVID, follow these steps:

1. Ensure that the Intel Quartus Prime project and Platform Designer system are configured to use VID components. This step exposes the `vid_cal_done_persist` interface on instantiated EMIF IP, which is required for communicating with the SmartVID controller.
2. Instantiate the SmartVID controller, using an I/O PLL IP core to drive the 125MHz `vid_clk` and the 25MHz `jtag_core_clk` inputs of the Smart VID controller.

Note: Do not connect the `emif_usr_clk` signal to either the `vid_clk` or `jtag_core_clk` inputs. Doing so would hold both the EMIF IP and the SmartVID controller in a perpetual reset condition.

3. Instantiate the Intel Stratix 10 EMIF IP.
4. Connect the `vid_cal_done_persist` signal from the EMIF IP with the `cal_done_persistent` signal on the SmartVID controller. This connection enables handshaking between the EMIF IP and the SmartVID controller, which allows the EMIF IP to delay memory calibration until after voltage levels are stabilized.

Note: The EMIF `vid_cal_done_persist` interface becomes available only when a VID-enabled device is selected.



3 Intel Stratix 10 EMIF IP End-User Signals

3.1 Interface and Signal Descriptions

The following sections describe each of the interfaces and their signals, by protocol, for the Intel Stratix 10 EMIF IP.

3.1.1 Intel Stratix 10 EMIF IP Interfaces for DDR3

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for DDR3.

Table 7. Interfaces for DDR3

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
afi_reset_n	Reset Output	AFI reset interface
afi_clk	Clock Output	AFI clock interface
<i>continued...</i>		

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

**ISO
9001:2008
Registered**



Interface Name	Interface Type	Description
afi_half_clk	Clock Output	AFI half-rate clock interface
afi	Conduit	Altera PHY Interface (AFI)
emif_usr_reset_n	Reset Output	User clock domain reset interface
emif_usr_clk	Clock Output	User clock interface
emif_usr_reset_n_sec	Reset Output	User clock domain reset interface (for the secondary interface in ping-pong configuration)
emif_usr_clk_sec	Clock Output	User clock interface (for the secondary interface in ping-pong configuration)
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
ctrl_amm	Avalon Memory-Mapped Slave	Controller Avalon Memory-Mapped interface
ctrl_ecc_user_interrupt	Conduit	Controller ECC user interrupt interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.1.1 global_reset_n for DDR3

Global reset interface

Table 8. Interface: global_reset_n

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.1.2 local_reset_req for DDR3

Local reset request interface

**Table 9. Interface: local_reset_req**

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.1.3 local_reset_status for DDR3

Local reset status interface

Table 10. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.

3.1.1.4 local_reset_req_out for DDR3

IF_LOCAL_RESET_REQ_OUT_DESC

Table 11. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.1.5 local_reset_status_in for DDR3

IF_LOCAL_RESET_STATUS_IN_DESC

Table 12. Interface: local_reset_status_in

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.1.6 pll_ref_clk for DDR3

PLL reference clock interface



Table 13. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input. Connect this interface to the clock output of the clock source that matches the PLL reference clock frequency value set in the parameter editor.

3.1.1.7 pll_ref_clk_out for DDR3

IF_PLL_REF_CLK_OUT_DESC

Table 14. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC

3.1.1.8 pll_locked for DDR3

IF_PLL_LOCKED_DESC

Table 15. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC

3.1.1.9 pll_extra_clk_0 for DDR3

IF_PLL_EXTRA_CLK_0_DESC

Table 16. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.1.10 pll_extra_clk_1 for DDR3

IF_PLL_EXTRA_CLK_1_DESC

Table 17. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.1.11 pll_extra_clk_2 for DDR3

IF_PLL_EXTRA_CLK_2_DESC



Table 18. Interface: pll_extra_clk_2

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.1.12 pll_extra_clk_3 for DDR3

IF_PLL_EXTRA_CLK_3_DESC

Table 19. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC

3.1.1.13 oct for DDR3

OCT interface

Table 20. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin

3.1.1.14 mem for DDR3

Interface between FPGA and external memory

Table 21. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_a	Output	Address
mem_ba	Output	Bank address
mem_cke	Output	Clock enable
mem_cs_n	Output	Chip select
mem_rm	Output	Rank multiplication for LRDIMM. Typically, mem_rm[0] and mem_rm[1] connect to CS2# and CS3# of the memory buffer of all LRDIMM slots.
mem_odt	Output	On-die termination
mem_ras_n	Output	RAS command
mem_cas_n	Output	CAS command
continued...		



Port Name	Direction	Description
mem_we_n	Output	WE command
mem_reset_n	Output	Asynchronous reset
mem_par	Output	Command and address parity. (memory format=RDIMM/LRDIMM)
mem_dm	Output	Write data mask
mem_dq	Bidirectional	Read/write data
mem_dqs	Bidirectional	Data strobe
mem_dqs_n	Bidirectional	Data strobe (negative leg)
mem_alert_n	Input	Alert flag

3.1.1.15 status for DDR3

PHY calibration status interface

Table 22. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration has failed.

3.1.1.16 afi_reset_n for DDR3

AFI reset interface

Table 23. Interface: afi_reset_n

Interface type: Reset Output

Port Name	Direction	Description
afi_reset_n	Output	Reset for the AFI clock domain. Asynchronous assertion and synchronous deassertion

3.1.1.17 afi_clk for DDR3

AFI clock interface

Table 24. Interface: afi_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_clk	Output	Clock for the Altera PHY Interface (AFI)

3.1.1.18 afi_half_clk for DDR3

AFI half-rate clock interface



Table 25. Interface: afi_half_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_half_clk	Output	Clock running at half the frequency of the AFI clock afi_clk

3.1.1.19 afi for DDR3

Altera PHY Interface (AFI)

Table 26. Interface: afi

Interface type: Conduit

Port Name	Direction	Description
afi_cal_success	Output	Signals calibration successful completion
afi_cal_fail	Output	Signals calibration failure
afi_cal_req	Input	When asserted, the interface is recalibrated
afi_rlat	Output	Latency in afi_clk cycles between read command and read data valid
afi_wlat	Output	Latency in afi_clk cycles between write command and write data valid
afi_addr	Input	Address
afi_ba	Input	Bank address
afi_cke	Input	Clock enable
afi_cs_n	Input	Chip select
afi_rm	Input	Rank multiplication for LRDIMM
afi_odt	Input	On-die termination
afi_ras_n	Input	RAS command
afi_cas_n	Input	CAS command
afi_we_n	Input	WE command
afi_rst_n	Input	Asynchronous reset
afi_par	Input	Command and address parity
afi_dm	Input	Write data mask
afi_dqs_burst	Input	Asserted by the controller to enable the output DQS signal
afi_wdata_valid	Input	Asserted by the controller to indicate that afi_wdata contains valid write data
afi_wdata	Input	Write data
afi_rdata_en_full	Input	Asserted by the controller to indicate the amount of relevant read data expected
afi_rdata	Output	Read data
<i>continued...</i>		



Port Name	Direction	Description
afi_rdata_valid	Output	Asserted by the PHY to indicate that afi_rdata contains valid read data
afi_rrank	Input	Asserted by the controller to indicate which rank is being read from, to control shadow register switching
afi_wrack	Input	Asserted by the controller to indicate which rank is being written to, to control shadow register switching

3.1.1.20 emif_usr_reset_n for DDR3

User clock domain reset interface

Table 27. Interface: emif_usr_reset_n

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion

3.1.1.21 emif_usr_clk for DDR3

User clock interface

Table 28. Interface: emif_usr_clk

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk	Output	User clock domain

3.1.1.22 emif_usr_reset_n_sec for DDR3

User clock domain reset interface (for the secondary interface in ping-pong configuration)

Table 29. Interface: emif_usr_reset_n_sec

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n_sec	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion. Intended for the secondary interface in a ping-pong configuration.

3.1.1.23 emif_usr_clk_sec for DDR3

User clock interface (for the secondary interface in ping-pong configuration)



Table 30. Interface: emif_usr_clk_sec

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk_sec	Output	User clock domain. Intended for the secondary interface in a ping-pong configuration.

3.1.1.24 cal_debug_reset_n for DDR3

User calibration debug clock domain reset interface

Table 31. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion

3.1.1.25 cal_debug_clk for DDR3

User calibration debug clock interface

Table 32. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.1.26 cal_debug_out_reset_n for DDR3

User calibration debug clock domain reset interface

Table 33. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion

3.1.1.27 cal_debug_out_clk for DDR3

User calibration debug clock interface

Table 34. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain



3.1.1.28 clks_sharing_master_out for DDR3

Core clocks sharing master interface

Table 35. Interface: clks_sharing_master_out

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.

3.1.1.29 clks_sharing_slave_in for DDR3

Core clocks sharing slave interface

Table 36. Interface: clks_sharing_slave_in

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.1.30 ctrl_amm for DDR3

Controller Avalon Memory-Mapped interface

Table 37. Interface: ctrl_amm

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
amm_ready	Output	Wait-request is asserted when controller is busy
amm_read	Input	Read request signal
amm_write	Input	Write request signal
amm_address	Input	Address for the read/write request
amm_readdata	Output	Read data
amm_writedata	Input	Write data
amm_burstcount	Input	Number of transfers in each read/write burst
amm_byteenable	Input	Byte-enable for write data
amm_beginbursttransfer	Input	Indicates when a burst is starting
amm_readdatavalid	Output	Indicates whether read data is valid

3.1.1.31 ctrl_ecc_user_interrupt for DDR3

Controller ECC user interrupt interface

**Table 38. Interface: ctrl_ecc_user_interrupt**

Interface type: Conduit

Port Name	Direction	Description
ctrl_ecc_user_interrupt	Output	Controller ECC user interrupt interface for connection to a custom control block that must be notified when either single-bit or double-bit ECC errors occur.

3.1.1.32 hps_emif for DDR3

Conduit between Hard Processor Subsystem and memory interface

Table 39. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface

3.1.1.33 cal_debug for DDR3

Calibration debug interface

Table 40. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.1.34 cal_debug_out for DDR3

Calibration debug interface

**Table 41. Interface: cal_debug_out**

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.1.35 generic_clk for DDR3

IF_GENERIC_CLK_DESC

Table 42. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC

3.1.1.36 generic_reset_n for DDR3

IF_GENERIC_RESET_DESC

Table 43. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.1.37 generic_conduit_reset_n for DDR3

IF_GENERIC_CONDUIT_RESET_DESC

Table 44. Interface: generic_conduit_reset_n

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.1.2 Intel Stratix 10 EMIF IP Interfaces for DDR4

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for DDR4.




Table 45. Interfaces for DDR4

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
afi_reset_n	Reset Output	AFI reset interface
afi_clk	Clock Output	AFI clock interface
afi_half_clk	Clock Output	AFI half-rate clock interface
afi	Conduit	Altera PHY Interface (AFI)
emif_usr_reset_n	Reset Output	User clock domain reset interface
emif_usr_clk	Clock Output	User clock interface
emif_usr_reset_n_sec	Reset Output	User clock domain reset interface (for the secondary interface in ping-pong configuration)
emif_usr_clk_sec	Clock Output	User clock interface (for the secondary interface in ping-pong configuration)
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
ctrl_amm	Avalon Memory-Mapped Slave	Controller Avalon Memory-Mapped interface
ctrl_ecc_user_interrupt	Conduit	Controller ECC user interrupt interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
<i>continued...</i>		



Interface Name	Interface Type	Description
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.2.1 global_reset_n for DDR4

Global reset interface

Table 46. Interface: global_reset_n

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.2.2 local_reset_req for DDR4

Local reset request interface

Table 47. Interface: local_reset_req

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.2.3 local_reset_status for DDR4

Local reset status interface

Table 48. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.



3.1.2.4 local_reset_req_out for DDR4

IF_LOCAL_RESET_REQ_OUT_DESC

Table 49. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.2.5 local_reset_status_in for DDR4

IF_LOCAL_RESET_STATUS_IN_DESC

Table 50. Interface: local_reset_status_in

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.2.6 pll_ref_clk for DDR4

PLL reference clock interface

Table 51. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input

3.1.2.7 pll_ref_clk_out for DDR4

IF_PLL_REF_CLK_OUT_DESC

Table 52. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC

3.1.2.8 pll_locked for DDR4

IF_PLL_LOCKED_DESC

Table 53. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC



3.1.2.9 pll_extra_clk_0 for DDR4

IF_PLL_EXTRA_CLK_0_DESC

Table 54. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.2.10 pll_extra_clk_1 for DDR4

IF_PLL_EXTRA_CLK_1_DESC

Table 55. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.2.11 pll_extra_clk_2 for DDR4

IF_PLL_EXTRA_CLK_2_DESC

Table 56. Interface: pll_extra_clk_2

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.2.12 pll_extra_clk_3 for DDR4

IF_PLL_EXTRA_CLK_3_DESC

Table 57. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC

3.1.2.13 oct for DDR4

OCT interface

Table 58. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin



3.1.2.14 mem for DDR4

Interface between FPGA and external memory

Table 59. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_a	Output	Address
mem_ba	Output	Bank address
mem_bg	Output	Bank group
mem_cke	Output	Clock enable
mem_cs_n	Output	Chip select
mem_odt	Output	On-die termination
mem_reset_n	Output	Asynchronous reset
mem_act_n	Output	Activation command
mem_par	Output	Command and address parity. (memory format=RDIMM/LRDIMM)
mem_dq	Bidirectional	Read/write data
mem_dbi_n	Bidirectional	Acts as either the data bus inversion pin, or the data mask pin, depending on configuration.
mem_dqs	Bidirectional	Data strobe
mem_dqs_n	Bidirectional	Data strobe (negative leg)
mem_alert_n	Input	Alert flag

3.1.2.15 status for DDR4

PHY calibration status interface

Table 60. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration failed

3.1.2.16 afi_reset_n for DDR4

AFI reset interface

**Table 61. Interface: afi_reset_n**

Interface type: Reset Output

Port Name	Direction	Description
afi_reset_n	Output	Reset for the AFI clock domain. Asynchronous assertion and synchronous deassertion

3.1.2.17 afi_clk for DDR4

AFI clock interface

Table 62. Interface: afi_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_clk	Output	Clock for the Altera PHY Interface (AFI)

3.1.2.18 afi_half_clk for DDR4

AFI half-rate clock interface

Table 63. Interface: afi_half_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_half_clk	Output	Clock running at half the frequency of the AFI clock afi_clk

3.1.2.19 afi for DDR4

Altera PHY Interface (AFI)

Table 64. Interface: afi

Interface type: Conduit

Port Name	Direction	Description
afi_cal_success	Output	Signals calibration successful completion
afi_cal_fail	Output	Signals calibration failure
afi_cal_req	Input	When asserted, the interface is recalibrated
afi_rlat	Output	Latency in afi_clk cycles between read command and read data valid
afi_wlat	Output	Latency in afi_clk cycles between write command and write data valid
afi_addr	Input	Address
afi_ba	Input	Bank address
afi_bg	Input	Bank group
afi_cke	Input	Clock enable
afi_cs_n	Input	Chip select
<i>continued...</i>		



Port Name	Direction	Description
afi_odt	Input	On-die termination
afi_rst_n	Input	Asynchronous reset
afi_act_n	Input	Activation command
afi_par	Input	Command and address parity
afi_dm_n	Input	Write data mask
afi_dqs_burst	Input	Asserted by the controller to enable the output DQS signal
afi_wdata_valid	Input	Asserted by the controller to indicate that afi_wdata contains valid write data
afi_wdata	Input	Write data
afi_rdata_en_full	Input	Asserted by the controller to indicate the amount of relevant read data expected
afi_rdata	Output	Read data
afi_rdata_valid	Output	Asserted by the PHY to indicate that afi_rdata contains valid read data
afi_rrank	Input	Asserted by the controller to indicate which rank is being read from, to control shadow register switching
afi_wrank	Input	Asserted by the controller to indicate which rank is being written to, to control shadow register switching

3.1.2.20 emif_usr_reset_n for DDR4

User clock domain reset interface

Table 65. Interface: emif_usr_reset_n

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion

3.1.2.21 emif_usr_clk for DDR4

User clock interface

Table 66. Interface: emif_usr_clk

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk	Output	User clock domain

3.1.2.22 emif_usr_reset_n_sec for DDR4

User clock domain reset interface (for the secondary interface in ping-pong configuration)



Table 67. Interface: emif_usr_reset_n_sec

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n_sec	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion. Intended for the secondary interface in a ping-pong configuration.

3.1.2.23 emif_usr_clk_sec for DDR4

User clock interface (for the secondary interface in ping-pong configuration)

Table 68. Interface: emif_usr_clk_sec

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk_sec	Output	User clock domain. Intended for the secondary interface in a ping-pong configuration.

3.1.2.24 cal_debug_reset_n for DDR4

User calibration debug clock domain reset interface

Table 69. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion

3.1.2.25 cal_debug_clk for DDR4

User calibration debug clock interface

Table 70. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.2.26 cal_debug_out_reset_n for DDR4

User calibration debug clock domain reset interface

Table 71. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion



3.1.2.27 cal_debug_out_clk for DDR4

User calibration debug clock interface

Table 72. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain

3.1.2.28 clks_sharing_master_out for DDR4

Core clocks sharing master interface

Table 73. Interface: clks_sharing_master_out

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.

3.1.2.29 clks_sharing_slave_in for DDR4

Core clocks sharing slave interface

Table 74. Interface: clks_sharing_slave_in

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.2.30 ctrl_amm for DDR4

Controller Avalon Memory-Mapped interface

Table 75. Interface: ctrl_amm

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
amm_ready	Output	Wait-request is asserted when controller is busy
amm_read	Input	Read request signal
amm_write	Input	Write request signal
amm_address	Input	Address for the read/write request
amm_readdata	Output	Read data
amm_writedata	Input	Write data
amm_burstcount	Input	Number of transfers in each read/write burst
<i>continued...</i>		



Port Name	Direction	Description
amm_byteenable	Input	Byte-enable for write data
amm_beginbursttransfer	Input	Indicates when a burst is starting
amm_readdatavalid	Output	Indicates whether read data is valid

3.1.2.31 ctrl_ecc_user_interrupt for DDR4

Controller ECC user interrupt interface

Table 76. Interface: ctrl_ecc_user_interrupt

Interface type: Conduit

Port Name	Direction	Description
ctrl_ecc_user_interrupt	Output	Controller ECC user interrupt signal to determine whether there is a bit error

3.1.2.32 hps_emif for DDR4

Conduit between Hard Processor Subsystem and memory interface

Table 77. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface

3.1.2.33 cal_debug for DDR4

Calibration debug interface

Table 78. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data
<i>continued...</i>		



Port Name	Direction	Description
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.2.34 cal_debug_out for DDR4

Calibration debug interface

Table 79. Interface: cal_debug_out

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.2.35 generic_clk for DDR4

IF_GENERIC_CLK_DESC

Table 80. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC

3.1.2.36 generic_reset_n for DDR4

IF_GENERIC_RESET_DESC

Table 81. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.2.37 generic_conduit_reset_n for DDR4

IF_GENERIC_CONDUIT_RESET_DESC

**Table 82. Interface: generic_conduit_reset_n**

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.1.3 Intel Stratix 10 EMIF IP Interfaces for LPDDR3

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for LPDDR3.

Table 83. Interfaces for LPDDR3

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
afi_reset_n	Reset Output	AFI reset interface
afi_clk	Clock Output	AFI clock interface
afi_half_clk	Clock Output	AFI half-rate clock interface
afi	Conduit	Altera PHY Interface (AFI)
emif_usr_reset_n	Reset Output	User clock domain reset interface
emif_usr_clk	Clock Output	User clock interface
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
<i>continued...</i>		



Interface Name	Interface Type	Description
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.3.1 global_reset_n for LPDDR3

Global reset interface

Table 84. Interface: global_reset_n

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.3.2 local_reset_req for LPDDR3

Local reset request interface

Table 85. Interface: local_reset_req

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.3.3 local_reset_status for LPDDR3

Local reset status interface



Table 86. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.

3.1.3.4 local_reset_req_out for LPDDR3

IF_LOCAL_RESET_REQ_OUT_DESC

Table 87. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.3.5 local_reset_status_in for LPDDR3

IF_LOCAL_RESET_STATUS_IN_DESC

Table 88. Interface: local_reset_status_in

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.3.6 pll_ref_clk for LPDDR3

PLL reference clock interface

Table 89. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input

3.1.3.7 pll_ref_clk_out for LPDDR3

IF_PLL_REF_CLK_OUT_DESC

Table 90. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC



3.1.3.8 pll_locked for LPDDR3

IF_PLL_LOCKED_DESC

Table 91. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC

3.1.3.9 pll_extra_clk_0 for LPDDR3

IF_PLL_EXTRA_CLK_0_DESC

Table 92. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.3.10 pll_extra_clk_1 for LPDDR3

IF_PLL_EXTRA_CLK_1_DESC

Table 93. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.3.11 pll_extra_clk_2 for LPDDR3

IF_PLL_EXTRA_CLK_2_DESC

Table 94. Interface: pll_extra_clk_2

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.3.12 pll_extra_clk_3 for LPDDR3

IF_PLL_EXTRA_CLK_3_DESC

Table 95. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC



3.1.3.13 oct for LPDDR3

OCT interface

Table 96. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin

3.1.3.14 mem for LPDDR3

Interface between FPGA and external memory

Table 97. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_a	Output	Address
mem_cke	Output	Clock enable
mem_cs_n	Output	Chip select
mem_odt	Output	On-die termination
mem_reset_n	Output	Asynchronous reset
mem_dm	Output	Write data mask
mem_dq	Bidirectional	Read/write data
mem_dqs	Bidirectional	Data strobe
mem_dqs_n	Bidirectional	Data strobe (negative leg)

3.1.3.15 status for LPDDR3

PHY calibration status interface

Table 98. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration failed

3.1.3.16 afi_reset_n for LPDDR3

AFI reset interface



Table 99. Interface: afi_reset_n

Interface type: Reset Output

Port Name	Direction	Description
afi_reset_n	Output	Reset for the AFI clock domain. Asynchronous assertion and synchronous deassertion

3.1.3.17 afi_clk for LPDDR3

AFI clock interface

Table 100. Interface: afi_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_clk	Output	Clock for the Altera PHY Interface (AFI)

3.1.3.18 afi_half_clk for LPDDR3

AFI half-rate clock interface

Table 101. Interface: afi_half_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_half_clk	Output	Clock running at half the frequency of the AFI clock afi_clk

3.1.3.19 afi for LPDDR3

Altera PHY Interface (AFI)

Table 102. Interface: afi

Interface type: Conduit

Port Name	Direction	Description
afi_cal_success	Output	Signals calibration successful completion
afi_cal_fail	Output	Signals calibration failure
afi_cal_req	Input	When asserted, the interface is recalibrated
afi_rlat	Output	Latency in afi_clk cycles between read command and read data valid
afi_wlat	Output	Latency in afi_clk cycles between write command and write data valid
afi_addr	Input	Address
afi_cke	Input	Clock enable
afi_cs_n	Input	Chip select
afi_odt	Input	On-die termination
afi_rst_n	Input	Asynchronous reset
continued...		



Port Name	Direction	Description
afi_dm	Input	Write data mask
afi_dqs_burst	Input	Asserted by the controller to enable the output DQS signal
afi_wdata_valid	Input	Asserted by the controller to indicate that afi_wdata contains valid write data
afi_wdata	Input	Write data
afi_rdata_en_full	Input	Asserted by the controller to indicate the amount of relevant read data expected
afi_rdata	Output	Read data
afi_rdata_valid	Output	Asserted by the PHY to indicate that afi_rdata contains valid read data
afi_rrank	Input	Asserted by the controller to indicate which rank is being read from, to control shadow register switching
afi_wrack	Input	Asserted by the controller to indicate which rank is being written to, to control shadow register switching

3.1.3.20 emif_usr_reset_n for LPDDR3

User clock domain reset interface

Table 103. Interface: emif_usr_reset_n

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion

3.1.3.21 emif_usr_clk for LPDDR3

User clock interface

Table 104. Interface: emif_usr_clk

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk	Output	User clock domain

3.1.3.22 cal_debug_reset_n for LPDDR3

User calibration debug clock domain reset interface

Table 105. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion



3.1.3.23 cal_debug_clk for LPDDR3

User calibration debug clock interface

Table 106. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.3.24 cal_debug_out_reset_n for LPDDR3

User calibration debug clock domain reset interface

Table 107. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion

3.1.3.25 cal_debug_out_clk for LPDDR3

User calibration debug clock interface

Table 108. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain

3.1.3.26 clks_sharing_master_out for LPDDR3

Core clocks sharing master interface

Table 109. Interface: clks_sharing_master_out

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.

3.1.3.27 clks_sharing_slave_in for LPDDR3

Core clocks sharing slave interface

**Table 110. Interface: clks_sharing_slave_in**

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.3.28 hps_emif for LPDDR3

Conduit between Hard Processor Subsystem and memory interface

Table 111. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface

3.1.3.29 cal_debug for LPDDR3

Calibration debug interface

Table 112. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.3.30 cal_debug_out for LPDDR3

Calibration debug interface



Table 113. Interface: cal_debug_out

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.3.31 generic_clk for LPDDR3

IF_GENERIC_CLK_DESC

Table 114. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC

3.1.3.32 generic_reset_n for LPDDR3

IF_GENERIC_RESET_DESC

Table 115. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.3.33 generic_conduit_reset_n for LPDDR3

IF_GENERIC_CONDUIT_RESET_DESC

Table 116. Interface: generic_conduit_reset_n

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.1.4 Intel Stratix 10 EMIF IP Interfaces for QDR II/II+/II+ Xtreme

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for QDR II/II+/II+ Xtreme.

**Table 117. Interfaces for QDR II/II+/II+ Xtreme**

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
emif_usr_reset_n	Reset Output	User clock domain reset interface
emif_usr_clk	Clock Output	User clock interface
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
ctrl_amm	Avalon Memory-Mapped Slave	Controller Avalon Memory-Mapped interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.4.1 global_reset_n for QDR II/II+/II+ Xtreme

Global reset interface

**Table 118. Interface: global_reset_n**

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.4.2 local_reset_req for QDR II/II+/II+ Xtreme

Local reset request interface

Table 119. Interface: local_reset_req

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.4.3 local_reset_status for QDR II/II+/II+ Xtreme

Local reset status interface

Table 120. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.

3.1.4.4 local_reset_req_out for QDR II/II+/II+ Xtreme

IF_LOCAL_RESET_REQ_OUT_DESC

Table 121. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.4.5 local_reset_status_in for QDR II/II+/II+ Xtreme

IF_LOCAL_RESET_STATUS_IN_DESC

**Table 122. Interface: local_reset_status_in**

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.4.6 pll_ref_clk for QDR II/II+/II+ Xtreme

PLL reference clock interface

Table 123. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input

3.1.4.7 pll_ref_clk_out for QDR II/II+/II+ Xtreme

IF_PLL_REF_CLK_OUT_DESC

Table 124. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC

3.1.4.8 pll_locked for QDR II/II+/II+ Xtreme

IF_PLL_LOCKED_DESC

Table 125. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC

3.1.4.9 pll_extra_clk_0 for QDR II/II+/II+ Xtreme

IF_PLL_EXTRA_CLK_0_DESC

Table 126. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.4.10 pll_extra_clk_1 for QDR II/II+/II+ Xtreme

IF_PLL_EXTRA_CLK_1_DESC



Table 127. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.4.11 pll_extra_clk_2 for QDR II/II+/II+ Xtreme

IF_PLL_EXTRA_CLK_2_DESC

Table 128. Interface: pll_extra_clk_2

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.4.12 pll_extra_clk_3 for QDR II/II+/II+ Xtreme

IF_PLL_EXTRA_CLK_3_DESC

Table 129. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC

3.1.4.13 oct for QDR II/II+/II+ Xtreme

OCT interface

Table 130. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin

3.1.4.14 mem for QDR II/II+/II+ Xtreme

Interface between FPGA and external memory

Table 131. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_k	Output	K clock
mem_k_n	Output	K clock (negative leg)
continued...		



Port Name	Direction	Description
mem_a	Output	Address
mem_reset_n	Output	Asynchronous reset
mem_wps_n	Output	Write port select
mem_rps_n	Output	Read port select
mem_doff_n	Output	DLL turn off
mem_bws_n	Output	Byte write select
mem_d	Output	Write data
mem_q	Input	Read data
mem_cq	Input	Echo clock
mem_cq_n	Input	Echo clock (negative leg)

3.1.4.15 status for QDR II/II+/II+ Xtreme

PHY calibration status interface

Table 132. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration failed

3.1.4.16 emif_usr_reset_n for QDR II/II+/II+ Xtreme

User clock domain reset interface

Table 133. Interface: emif_usr_reset_n

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion

3.1.4.17 emif_usr_clk for QDR II/II+/II+ Xtreme

User clock interface

Table 134. Interface: emif_usr_clk

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk	Output	User clock domain

3.1.4.18 cal_debug_reset_n for QDR II/II+/II+ Xtreme

User calibration debug clock domain reset interface



Table 135. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion

3.1.4.19 cal_debug_clk for QDR II/II+/II+ Xtreme

User calibration debug clock interface

Table 136. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.4.20 cal_debug_out_reset_n for QDR II/II+/II+ Xtreme

User calibration debug clock domain reset interface

Table 137. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion

3.1.4.21 cal_debug_out_clk for QDR II/II+/II+ Xtreme

User calibration debug clock interface

Table 138. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain

3.1.4.22 clks_sharing_master_out for QDR II/II+/II+ Xtreme

Core clocks sharing master interface

Table 139. Interface: clks_sharing_master_out

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.



3.1.4.23 clks_sharing_slave_in for QDR II/II+/II+ Xtreme

Core clocks sharing slave interface

Table 140. Interface: clks_sharing_slave_in

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.4.24 ctrl_amm for QDR II/II+/II+ Xtreme

Controller Avalon Memory-Mapped interface

Table 141. Interface: ctrl_amm

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
amm_ready	Output	Wait-request is asserted when controller is busy
amm_read	Input	Read request signal
amm_write	Input	Write request signal
amm_address	Input	Address for the read/write request
amm_readdata	Output	Read data
amm_writedata	Input	Write data
amm_burstcount	Input	Number of transfers in each read/write burst
amm_byteenable	Input	Byte-enable for write data
amm_beginbursttransfer	Input	Indicates when a burst is starting
amm_readdatavalid	Output	Indicates whether read data is valid

3.1.4.25 hps_emif for QDR II/II+/II+ Xtreme

Conduit between Hard Processor Subsystem and memory interface

Table 142. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface



3.1.4.26 cal_debug for QDR II/II+/II+ Xtreme

Calibration debug interface

Table 143. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.4.27 cal_debug_out for QDR II/II+/II+ Xtreme

Calibration debug interface

Table 144. Interface: cal_debug_out

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.4.28 generic_clk for QDR II/II+/II+ Xtreme

IF_GENERIC_CLK_DESC

Table 145. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC



3.1.4.29 generic_reset_n for QDR II/II+/II+ Xtreme

IF_GENERIC_RESET_DESC

Table 146. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.4.30 generic_conduit_reset_n for QDR II/II+/II+ Xtreme

IF_GENERIC_CONDUIT_RESET_DESC

Table 147. Interface: generic_conduit_reset_n

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.1.5 Intel Stratix 10 EMIF IP Interfaces for QDR-IV

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for QDR-IV.

Table 148. Interfaces for QDR-IV

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
<i>continued...</i>		



Interface Name	Interface Type	Description
afi_reset_n	Reset Output	AFI reset interface
afi_clk	Clock Output	AFI clock interface
afi_half_clk	Clock Output	AFI half-rate clock interface
afi	Conduit	Altera PHY Interface (AFI)
emif_usr_reset_n	Reset Output	User clock domain reset interface
emif_usr_clk	Clock Output	User clock interface
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
ctrl_amm	Avalon Memory-Mapped Slave	Controller Avalon Memory-Mapped interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.5.1 global_reset_n for QDR-IV

Global reset interface

Table 149. Interface: global_reset_n

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.5.2 local_reset_req for QDR-IV

Local reset request interface



Table 150. Interface: local_reset_req

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.5.3 local_reset_status for QDR-IV

Local reset status interface

Table 151. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.

3.1.5.4 local_reset_req_out for QDR-IV

IF_LOCAL_RESET_REQ_OUT_DESC

Table 152. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.5.5 local_reset_status_in for QDR-IV

IF_LOCAL_RESET_STATUS_IN_DESC

Table 153. Interface: local_reset_status_in

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.5.6 pll_ref_clk for QDR-IV

PLL reference clock interface



Table 154. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input

3.1.5.7 pll_ref_clk_out for QDR-IV

IF_PLL_REF_CLK_OUT_DESC

Table 155. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC

3.1.5.8 pll_locked for QDR-IV

IF_PLL_LOCKED_DESC

Table 156. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC

3.1.5.9 pll_extra_clk_0 for QDR-IV

IF_PLL_EXTRA_CLK_0_DESC

Table 157. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.5.10 pll_extra_clk_1 for QDR-IV

IF_PLL_EXTRA_CLK_1_DESC

Table 158. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.5.11 pll_extra_clk_2 for QDR-IV

IF_PLL_EXTRA_CLK_2_DESC

**Table 159. Interface: pll_extra_clk_2**

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.5.12 pll_extra_clk_3 for QDR-IV

IF_PLL_EXTRA_CLK_3_DESC

Table 160. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC

3.1.5.13 oct for QDR-IV

OCT interface

Table 161. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin

3.1.5.14 mem for QDR-IV

Interface between FPGA and external memory

Table 162. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_dka	Output	DK clock for port A
mem_dka_n	Output	DK clock for port A (negative leg)
mem_dkb	Output	DK clock for port B
mem_dkb_n	Output	DK clock for port B (negative leg)
mem_a	Output	Address
mem_reset_n	Output	Asynchronous reset
mem_lda_n	Output	Synchronous load for port A
mem_ldb_n	Output	Synchronous load for port B
mem_rwa_n	Output	Synchronous read/write for port A
<i>continued...</i>		



Port Name	Direction	Description
mem_rwb_n	Output	Synchronous read/write for port B
mem_lbk0_n	Output	Loopback mode
mem_lbk1_n	Output	Loopback mode
mem_cfg_n	Output	Configuration bit
mem_ap	Output	Address parity
mem_ainv	Output	Address inversion
mem_dqa	Bidirectional	Read/write data for port A
mem_dqb	Bidirectional	Read/write data for port B
mem_dinva	Bidirectional	Read/write data inversion for port A
mem_dinvb	Bidirectional	Read/write data inversion for port B
mem_qka	Input	Read data clock for port A
mem_qka_n	Input	Read data clock for port A (negative leg)
mem_qkb	Input	Read data clock for port B
mem_qkb_n	Input	Read data clock for port B (negative leg)
mem_pe_n	Input	Address parity error flag

3.1.5.15 status for QDR-IV

PHY calibration status interface

Table 163. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration failed

3.1.5.16 afi_reset_n for QDR-IV

AFI reset interface

Table 164. Interface: afi_reset_n

Interface type: Reset Output

Port Name	Direction	Description
afi_reset_n	Output	Reset for the AFI clock domain. Asynchronous assertion and synchronous deassertion

3.1.5.17 afi_clk for QDR-IV

AFI clock interface



Table 165. Interface: afi_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_clk	Output	Clock for the Altera PHY Interface (AFI)

3.1.5.18 afi_half_clk for QDR-IV

AFI half-rate clock interface

Table 166. Interface: afi_half_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_half_clk	Output	Clock running at half the frequency of the AFI clock afi_clk

3.1.5.19 afi for QDR-IV

Altera PHY Interface (AFI)

Table 167. Interface: afi

Interface type: Conduit

Port Name	Direction	Description
afi_id_n	Input	Synchronous load for port A and B
afi_rw_n	Input	Synchronous read/write for port A and B
afi_lbk0_n	Input	Loopback mode
afi_lbk1_n	Input	Loopback mode
afi_cfg_n	Input	Configuration bit
afi_ap	Input	Address parity
afi_ainv	Input	Address inversion
afi_rdata_dinv	Output	Data inversion for read data
afi_wdata_dinv	Input	Data inversion for write data
afi_pe_n	Output	Address parity error flag

3.1.5.20 emif_usr_reset_n for QDR-IV

User clock domain reset interface

Table 168. Interface: emif_usr_reset_n

Interface type: Reset Output

Port Name	Direction	Description
emif_usr_reset_n	Output	Reset for the user clock domain. Asynchronous assertion and synchronous deassertion



3.1.5.21 emif_usr_clk for QDR-IV

User clock interface

Table 169. Interface: emif_usr_clk

Interface type: Clock Output

Port Name	Direction	Description
emif_usr_clk	Output	User clock domain

3.1.5.22 cal_debug_reset_n for QDR-IV

User calibration debug clock domain reset interface

Table 170. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion

3.1.5.23 cal_debug_clk for QDR-IV

User calibration debug clock interface

Table 171. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.5.24 cal_debug_out_reset_n for QDR-IV

User calibration debug clock domain reset interface

Table 172. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion

3.1.5.25 cal_debug_out_clk for QDR-IV

User calibration debug clock interface



Table 173. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain

3.1.5.26 clks_sharing_master_out for QDR-IV

Core clocks sharing master interface

Table 174. Interface: clks_sharing_master_out

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.

3.1.5.27 clks_sharing_slave_in for QDR-IV

Core clocks sharing slave interface

Table 175. Interface: clks_sharing_slave_in

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.5.28 ctrl_amm for QDR-IV

Controller Avalon Memory-Mapped interface

Table 176. Interface: ctrl_amm

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
amm_ready	Output	Wait-request is asserted when controller is busy
amm_read	Input	Read request signal
amm_write	Input	Write request signal
amm_address	Input	Address for the read/write request
amm_readdata	Output	Read data
amm_writedata	Input	Write data
amm_burstcount	Input	Number of transfers in each read/write burst
amm_beginbursttransfer	Input	Indicates when a burst is starting
amm_readdatavalid	Output	Indicates whether read data is valid

3.1.5.29 hps_emif for QDR-IV

Conduit between Hard Processor Subsystem and memory interface



Table 177. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface

3.1.5.30 cal_debug for QDR-IV

Calibration debug interface

Table 178. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.5.31 cal_debug_out for QDR-IV

Calibration debug interface

Table 179. Interface: cal_debug_out

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
<i>continued...</i>		



Port Name	Direction	Description
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.5.32 generic_clk for QDR-IV

IF_GENERIC_CLK_DESC

Table 180. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC

3.1.5.33 generic_reset_n for QDR-IV

IF_GENERIC_RESET_DESC

Table 181. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.5.34 generic_conduit_reset_n for QDR-IV

IF_GENERIC_CONDUIT_RESET_DESC

Table 182. Interface: generic_conduit_reset_n

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.1.6 Intel Stratix 10 EMIF IP Interfaces for RLDRAM 3

The interfaces in the Intel Stratix 10 External Memory Interface IP each have signals that can be connected in Platform Designer (formerly known as Qsys). The following table lists the interfaces and corresponding interface types for RLDRAM 3.

Table 183. Interfaces for RLDRAM 3

Interface Name	Interface Type	Description
global_reset_n	Reset Input	Global reset interface
local_reset_req	Conduit	Local reset request interface
local_reset_status	Conduit	Local reset status interface
<i>continued...</i>		



Interface Name	Interface Type	Description
local_reset_req_out	Conduit	IF_LOCAL_RESET_REQ_OUT_DESC
local_reset_status_in	Conduit	IF_LOCAL_RESET_STATUS_IN_DESC
pll_ref_clk	Clock Input	PLL reference clock interface
pll_ref_clk_out	Clock Output	IF_PLL_REF_CLK_OUT_DESC
pll_locked	Conduit	IF_PLL_LOCKED_DESC
pll_extra_clk_0	Clock Output	IF_PLL_EXTRA_CLK_0_DESC
pll_extra_clk_1	Clock Output	IF_PLL_EXTRA_CLK_1_DESC
pll_extra_clk_2	Clock Output	IF_PLL_EXTRA_CLK_2_DESC
pll_extra_clk_3	Clock Output	IF_PLL_EXTRA_CLK_3_DESC
oct	Conduit	OCT interface
mem	Conduit	Interface between FPGA and external memory
status	Conduit	PHY calibration status interface
afi_reset_n	Reset Output	AFI reset interface
afi_clk	Clock Output	AFI clock interface
afi_half_clk	Clock Output	AFI half-rate clock interface
afi	Conduit	Altera PHY Interface (AFI)
cal_debug_reset_n	Reset Input	User calibration debug clock domain reset interface
cal_debug_clk	Clock Input	User calibration debug clock interface
cal_debug_out_reset_n	Reset Output	User calibration debug clock domain reset interface
cal_debug_out_clk	Clock Output	User calibration debug clock interface
clks_sharing_master_out	Conduit	Core clocks sharing master interface
clks_sharing_slave_in	Conduit	Core clocks sharing slave interface
hps_emif	Conduit	Conduit between Hard Processor Subsystem and memory interface
cal_debug	Avalon Memory-Mapped Slave	Calibration debug interface
cal_debug_out	Avalon Memory-Mapped Master	Calibration debug interface
generic_clk	Clock Input	IF_GENERIC_CLK_DESC
generic_reset_n	Reset Input	IF_GENERIC_RESET_DESC
generic_conduit_reset_n	Conduit	IF_GENERIC_CONDUIT_RESET_DESC

3.1.6.1 global_reset_n for RLDRAM 3

Global reset interface



Table 184. Interface: global_reset_n

Interface type: Reset Input

Port Name	Direction	Description
global_reset_n	Input	Asynchronous reset causes the memory interface to be reset and recalibrated. The global reset signal applies to all memory interfaces placed within an I/O column.

3.1.6.2 local_reset_req for RLDRAM 3

Local reset request interface

Table 185. Interface: local_reset_req

Interface type: Conduit

Port Name	Direction	Description
local_reset_req	Input	Signal from user logic to request the memory interface to be reset and recalibrated. Reset request is sent by transitioning the local_reset_req signal from low to high, then keeping the signal at the high state for a minimum of 2 EMIF core clock cycles, then transitioning the signal from high to low. local_reset_req is asynchronous in that there is no setup/hold timing to meet, but it must meet the minimum pulse width requirement of 2 EMIF core clock cycles.

3.1.6.3 local_reset_status for RLDRAM 3

Local reset status interface

Table 186. Interface: local_reset_status

Interface type: Conduit

Port Name	Direction	Description
local_reset_done	Output	Signal from memory interface to indicate whether it has completed a reset sequence, is currently out of reset, and is ready for a new reset request. When local_reset_done is low, the memory interface is in reset.

3.1.6.4 local_reset_req_out for RLDRAM 3

IF_LOCAL_RESET_REQ_OUT_DESC

Table 187. Interface: local_reset_req_out

Interface type: Conduit

Port Name	Direction	Description
local_reset_req_out	Output	PORT_LOCAL_RESET_REQ_OUT_DESC

3.1.6.5 local_reset_status_in for RLDRAM 3

IF_LOCAL_RESET_STATUS_IN_DESC



Table 188. Interface: local_reset_status_in

Interface type: Conduit

Port Name	Direction	Description
local_reset_done_in	Input	PORT_LOCAL_RESET_DONE_IN_DESC

3.1.6.6 pll_ref_clk for RLD RAM 3

PLL reference clock interface

Table 189. Interface: pll_ref_clk

Interface type: Clock Input

Port Name	Direction	Description
pll_ref_clk	Input	PLL reference clock input

3.1.6.7 pll_ref_clk_out for RLD RAM 3

IF_PLL_REF_CLK_OUT_DESC

Table 190. Interface: pll_ref_clk_out

Interface type: Clock Output

Port Name	Direction	Description
pll_ref_clk_out	Output	PORT_PLL_REF_CLK_OUT_DESC

3.1.6.8 pll_locked for RLD RAM 3

IF_PLL_LOCKED_DESC

Table 191. Interface: pll_locked

Interface type: Conduit

Port Name	Direction	Description
pll_locked	Output	PORT_PLL_LOCKED_DESC

3.1.6.9 pll_extra_clk_0 for RLD RAM 3

IF_PLL_EXTRA_CLK_0_DESC

Table 192. Interface: pll_extra_clk_0

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_0	Output	PORT_PLL_EXTRA_CLK_0_DESC

3.1.6.10 pll_extra_clk_1 for RLD RAM 3

IF_PLL_EXTRA_CLK_1_DESC



Table 193. Interface: pll_extra_clk_1

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_1	Output	PORT_PLL_EXTRA_CLK_1_DESC

3.1.6.11 pll_extra_clk_2 for RLD RAM 3

IF_PLL_EXTRA_CLK_2_DESC

Table 194. Interface: pll_extra_clk_2

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_2	Output	PORT_PLL_EXTRA_CLK_2_DESC

3.1.6.12 pll_extra_clk_3 for RLD RAM 3

IF_PLL_EXTRA_CLK_3_DESC

Table 195. Interface: pll_extra_clk_3

Interface type: Clock Output

Port Name	Direction	Description
pll_extra_clk_3	Output	PORT_PLL_EXTRA_CLK_3_DESC

3.1.6.13 oct for RLD RAM 3

OCT interface

Table 196. Interface: oct

Interface type: Conduit

Port Name	Direction	Description
oct_rzqin	Input	Calibrated On-Chip Termination (OCT) RZQ input pin

3.1.6.14 mem for RLD RAM 3

Interface between FPGA and external memory

Table 197. Interface: mem

Interface type: Conduit

Port Name	Direction	Description
mem_ck	Output	CK clock
mem_ck_n	Output	CK clock (negative leg)
mem_dk	Output	DK clock
mem_dk_n	Output	DK clock (negative leg)
<i>continued...</i>		



Port Name	Direction	Description
mem_a	Output	Address
mem_ba	Output	Bank address
mem_cs_n	Output	Chip select
mem_rm	Output	Rank multiplication for LRDIMM. Typically, mem_rm[0] and mem_rm[1] connect to CS2# and CS3# of the memory buffer of all LRDIMM slots.
mem_we_n	Output	WE command
mem_reset_n	Output	Asynchronous reset
mem_ref_n	Output	REF command
mem_dm	Output	Write data mask
mem_dq	Bidirectional	Read/write data
mem_qk	Input	Read data clock
mem_qk_n	Input	Read data clock (negative leg)

3.1.6.15 status for RLD RAM 3

PHY calibration status interface

Table 198. Interface: status

Interface type: Conduit

Port Name	Direction	Description
local_cal_success	Output	When high, indicates that PHY calibration was successful
local_cal_fail	Output	When high, indicates that PHY calibration failed

3.1.6.16 afi_reset_n for RLD RAM 3

AFI reset interface

Table 199. Interface: afi_reset_n

Interface type: Reset Output

Port Name	Direction	Description
afi_reset_n	Output	Reset for the AFI clock domain. Asynchronous assertion and synchronous deassertion

3.1.6.17 afi_clk for RLD RAM 3

AFI clock interface

Table 200. Interface: afi_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_clk	Output	Clock for the Altera PHY Interface (AFI)



3.1.6.18 afi_half_clk for RLDRAM 3

AFI half-rate clock interface

Table 201. Interface: afi_half_clk

Interface type: Clock Output

Port Name	Direction	Description
afi_half_clk	Output	Clock running at half the frequency of the AFI clock afi_clk

3.1.6.19 afi for RLDRAM 3

Altera PHY Interface (AFI)

Table 202. Interface: afi

Interface type: Conduit

Port Name	Direction	Description
afi_cal_success	Output	Signals calibration successful completion
afi_cal_fail	Output	Signals calibration failure
afi_cal_req	Input	When asserted, the interface is recalibrated
afi_rlat	Output	Latency in afi_clk cycles between read command and read data valid
afi_wlat	Output	Latency in afi_clk cycles between write command and write data valid
afi_addr	Input	Address
afi_ba	Input	Bank address
afi_cs_n	Input	Chip select
afi_we_n	Input	WE command
afi_rst_n	Input	Asynchronous reset
afi_ref_n	Input	REF command
afi_dm	Input	Write data mask
afi_wdata_valid	Input	Asserted by the controller to indicate that afi_wdata contains valid write data
afi_wdata	Input	Write data
afi_rdata_en_full	Input	Asserted by the controller to indicate the amount of relevant read data expected
afi_rdata	Output	Read data
afi_rdata_valid	Output	Asserted by the PHY to indicate that afi_rdata contains valid read data
afi_rrank	Input	Asserted by the controller to indicate which rank is being read from, to control shadow register switching
afi_wrank	Input	Asserted by the controller to indicate which rank is being written to, to control shadow register switching



3.1.6.20 cal_debug_reset_n for RLDRAM 3

User calibration debug clock domain reset interface

Table 203. Interface: cal_debug_reset_n

Interface type: Reset Input

Port Name	Direction	Description
cal_debug_reset_n	Input	Reset for the user clock connecting to the Avalon calibration debug bus. Asynchronous assertion and synchronous deassertion

3.1.6.21 cal_debug_clk for RLDRAM 3

User calibration debug clock interface

Table 204. Interface: cal_debug_clk

Interface type: Clock Input

Port Name	Direction	Description
cal_debug_clk	Input	User clock domain

3.1.6.22 cal_debug_out_reset_n for RLDRAM 3

User calibration debug clock domain reset interface

Table 205. Interface: cal_debug_out_reset_n

Interface type: Reset Output

Port Name	Direction	Description
cal_debug_out_reset_n	Output	Reset for the user clock connecting to the Avalon calibration debug_out bus. Asynchronous assertion and synchronous deassertion

3.1.6.23 cal_debug_out_clk for RLDRAM 3

User calibration debug clock interface

Table 206. Interface: cal_debug_out_clk

Interface type: Clock Output

Port Name	Direction	Description
cal_debug_out_clk	Output	User clock domain

3.1.6.24 clks_sharing_master_out for RLDRAM 3

Core clocks sharing master interface

**Table 207. Interface: clks_sharing_master_out**

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_master_out	Output	This port should fanout to all the core clocks sharing slaves.

3.1.6.25 clks_sharing_slave_in for RLDRAM 3

Core clocks sharing slave interface

Table 208. Interface: clks_sharing_slave_in

Interface type: Conduit

Port Name	Direction	Description
clks_sharing_slave_in	Input	This port should be connected to the core clocks sharing master.

3.1.6.26 hps_emif for RLDRAM 3

Conduit between Hard Processor Subsystem and memory interface

Table 209. Interface: hps_emif

Interface type: Conduit

Port Name	Direction	Description
hps_to_emif	Input	Signals coming from Hard Processor Subsystem to the memory interface
emif_to_hps	Output	Signals going to Hard Processor Subsystem from the memory interface
hps_to_emif_gp	Input	Signals coming from Hard Processor Subsystem GPIO to the memory interface
emif_to_hps_gp	Output	Signals going to Hard Processor Subsystem GPIO from the memory interface

3.1.6.27 cal_debug for RLDRAM 3

Calibration debug interface

Table 210. Interface: cal_debug

Interface type: Avalon Memory-Mapped Slave

Port Name	Direction	Description
cal_debug_waitrequest	Output	Wait-request is asserted when controller is busy
cal_debug_read	Input	Read request signal
cal_debug_write	Input	Write request signal
cal_debug_addr	Input	Address for the read/write request
cal_debug_read_data	Output	Read data

continued...



Port Name	Direction	Description
cal_debug_write_data	Input	Write data
cal_debug_byteenable	Input	Byte-enable for write data
cal_debug_read_data_valid	Output	Indicates whether read data is valid

3.1.6.28 cal_debug_out for RLDRAM 3

Calibration debug interface

Table 211. Interface: cal_debug_out

Interface type: Avalon Memory-Mapped Master

Port Name	Direction	Description
cal_debug_out_waitrequest	Input	Wait-request is asserted when controller is busy
cal_debug_out_read	Output	Read request signal
cal_debug_out_write	Output	Write request signal
cal_debug_out_addr	Output	Address for the read/write request
cal_debug_out_read_data	Input	Read data
cal_debug_out_write_data	Output	Write data
cal_debug_out_byteenable	Output	Byte-enable for write data
cal_debug_out_read_data_valid	Input	Indicates whether read data is valid

3.1.6.29 generic_clk for RLDRAM 3

IF_GENERIC_CLK_DESC

Table 212. Interface: generic_clk

Interface type: Clock Input

Port Name	Direction	Description
clk	Input	PORT_GENERIC_CLK_DESC

3.1.6.30 generic_reset_n for RLDRAM 3

IF_GENERIC_RESET_DESC

Table 213. Interface: generic_reset_n

Interface type: Reset Input

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_RESET_DESC

3.1.6.31 generic_conduit_reset_n for RLDRAM 3

IF_GENERIC_CONDUIT_RESET_DESC

**Table 214. Interface: generic_conduit_reset_n**

Interface type: Conduit

Port Name	Direction	Description
reset_n	Input	PORT_GENERIC_CONDUIT_RESET_DESC

3.2 AFI Signals

The following tables list Altera PHY interface (AFI) signals grouped according to their functions.

In each table, the **Direction** column denotes the direction of the signal relative to the PHY. For example, a signal defined as an output passes out of the PHY to the controller. The AFI specification does not include any bidirectional signals.

Not all signals are used for all protocols.

3.2.1 AFI Clock and Reset Signals

The AFI interface provides up to two clock signals and an asynchronous reset signal.

Table 215. Clock and Reset Signals

Signal Name	Direction	Width	Description
afi_clk	Output	1	Clock with which all data exchanged on the AFI bus is synchronized. In general, this clock is referred to as full-rate, half-rate, or quarter-rate, depending on the ratio between the frequency of this clock and the frequency of the memory device clock.
afi_half_clk	Output	1	Clock signal that runs at half the speed of the afi_clk. The controller uses this signal when the half-rate bridge feature is in use. This signal is optional.
afi_reset_n	Output	1	Asynchronous reset output signal. You must synchronize this signal to the clock domain in which you use it.

3.2.2 AFI Address and Command Signals

The address and command signals for AFI 4.0 encode read/write/configuration commands to send to the memory device. The address and command signals are single-data rate signals.

Table 216. Address and Command Signals

Signal Name	Direction	Width	Description
afi_addr	Input	AFI_ADDR_WIDTH	Address or CA bus (LPDDR3 only). ADDR_RATE_RATIO is 2 for LPDDR3 CA bus.
afi_bg	Input	AFI_BANKGROUP_WIDTH	Bank group (DDR4 only).
afi_ba	Input	AFI_BANKADDR_WIDTH	Bank address. (Not applicable for LPDDR3.)
<i>continued...</i>			



Signal Name	Direction	Width	Description
afi_cke	Input	AFI_CLK_EN_WIDTH	Clock enable.
afi_cs_n	Input	AFI_CS_WIDTH	Chip select signal. (The number of chip selects may not match the number of ranks; for example, RDIMMs and LRDIMMs require a minimum of 2 chip select signals for both single-rank and dual-rank configurations. Consult your memory device data sheet for information about chip select signal width.) (Matches the number of ranks for LPDDR3.)
afi_ras_n	Input	AFI_CONTROL_WIDTH	RAS# (for DDR3 memory devices.)
afi_we_n	Input	AFI_CONTROL_WIDTH	WE# (for DDR3 memory devices.)
afi_rw_n	Input	AFI_CONTROL_WIDTH * 2	RWA/B# (QDR-IV).
afi_cas_n	Input	AFI_CONTROL_WIDTH	CAS# (for DDR3 memory devices.)
afi_act_n	Input	AFI_CONTROL_WIDTH	ACT# (DDR4).
afi_rst_n	Input	AFI_CONTROL_WIDTH	RESET# (for DDR3 and DDR4 memory devices.)
afi_odt	Input	AFI_CLK_EN_WIDTH	On-die termination signal for DDR3, and LPDDR3 memory devices. (Do not confuse this memory device signal with the FPGA's internal on-chip termination signal.)
afi_par	Input	AFI_CS_WIDTH	Address and command parity input. (DDR4) Address parity input. (QDR-IV)
afi_ainv	Input	AFI_CONTROL_WIDTH	Address inversion. (QDR-IV)
afi_mem_clk_disable	Input	AFI_CLK_PAIR_COUNT	When this signal is asserted, mem_clk and mem_clk_n are disabled. This signal is used in low-power mode.
afi_wps_n	Output	AFI_CS_WIDTH	WPS (for QDR II/II+ memory devices.)
afi_rps_n	Output	AFI_CS_WIDTH	RPS (for QDR II/II+ memory devices.)

3.2.3 AFI Write Data Signals

Write Data Signals for AFI 4.0 control the data, data mask, and strobe signals passed to the memory device during write operations.

Table 217. Write Data Signals

Signal Name	Direction	Width	Description
afi_dqs_burst	Input	AFI_RATE_RATIO	Controls the enable on the strobe (DQS) pins for DDR3 and LPDDR3 memory devices. When this signal is asserted, mem_dqs and mem_dqsn are driven.
<i>continued...</i>			



Signal Name	Direction	Width	Description
			This signal must be asserted before <code>afi_wdata_valid</code> to implement the write preamble, and must be driven for the correct duration to generate a correctly timed <code>mem_dqs</code> signal.
<code>afi_wdata_valid</code>	Input	<code>AFI_RATE_RATIO</code>	Write data valid signal. This signal controls the output enable on the data and data mask pins.
<code>afi_wdata</code>	Input	<code>AFI_DQ_WIDTH</code>	Write data signal to send to the memory device at double-data rate. This signal controls the PHY's <code>mem_dq</code> output.
<code>afi_dm</code>	Input	<code>AFI_DM_WIDTH</code>	Data mask. This signal controls the PHY's <code>mem_dm</code> signal for DDR3 and LPDDR3 memory devices.) Also directly controls the PHY's <code>mem_dbi</code> signal for DDR4. The <code>mem_dm</code> and <code>mem_dbi</code> features share the same port on the memory device.
<code>afi_bws_n</code>	Input	<code>AFI_DM_WIDTH</code>	Data mask. This signal controls the PHY's <code>mem_bws_n</code> signal for QDR II/II+ memory devices.
<code>afi_dinv</code>	Input	<code>AFI_WRITE_DQS_WIDTH * 2</code>	Data inversion. It directly controls the PHY's <code>mem_dinva/b</code> signal for QDR-IV devices.

3.2.4 AFI Read Data Signals

Read Data Signals for AFI 4.0 control the data sent from the memory device during read operations.

Table 218. Read Data Signals

Signal Name	Direction	Width	Description
<code>afi_rdata_en_full</code>	Input	<code>AFI_RATE_RATIO</code>	Read data enable full. Indicates that the memory controller is currently performing a read operation. This signal is held high for the entire read burst. If this signal is aligned to even clock cycles, it is possible to use 1-bit even in half-rate mode (i.e., <code>AFI_RATE=2</code>).
<code>afi_rdata</code>	Output	<code>AFI_DQ_WIDTH</code>	Read data from the memory device. This data is considered valid only when <code>afi_rdata_valid</code> is asserted by the PHY.
<code>afi_rdata_valid</code>	Output	<code>AFI_RATE_RATIO</code>	Read data valid. When asserted, this signal indicates that the <code>afi_rdata</code> bus is valid. If this signal is aligned to even clock cycles, it is possible to use 1-bit even in half-rate mode (i.e., <code>AFI_RATE=2</code>).



3.2.5 AFI Calibration Status Signals

The PHY instantiates a sequencer which calibrates the memory interface with the memory device and some internal components such as read FIFOs and valid FIFOs. The sequencer reports the results of the calibration process to the controller through the Calibration Status Signals in the AFI interface.

Table 219. Calibration Status Signals

Signal Name	Direction	Width	Description
afi_cal_success	Output	1	Asserted to indicate that calibration has completed successfully.
afi_cal_fail	Output	1	Asserted to indicate that calibration has failed.
afi_cal_req	Input	1	Effectively a synchronous reset for the sequencer. When this signal is asserted, the sequencer returns to the reset state; when this signal is released, a new calibration sequence begins.
afi_wlat	Output	AFI_WLAT_WIDTH	The required write latency in afi_clk cycles, between address/command and write data being issued at the PHY/controller interface. The afi_wlat value can be different for different groups; each group's write latency can range from 0 to 63. If write latency is the same for all groups, only the lowest 6 bits are required.
afi_rlat (1)	Output	AFI_RLAT_WIDTH	The required read latency in afi_clk cycles between address/command and read data being returned to the PHY/controller interface. Values can range from 0 to 63.
Note to Table: 1. The afi_rlat signal is not supported for PHY-only designs. Instead, you can sample the afi_rdata_valid signal to determine when valid read data is available.			

3.2.6 AFI Tracking Management Signals

When tracking management is enabled, the sequencer can take control over the AFI 4.0 interface at given intervals, and issue commands to the memory device to track the internal DQS Enable signal alignment to the DQS signal returning from the memory device. The tracking management portion of the AFI 4.0 interface provides a means for the sequencer and the controller to exchange handshake signals.

**Table 220. Tracking Management Signals**

Signal Name	Direction	Width	Description
afi_ctl_refresh_done	Input	4	Handshaking signal from controller to tracking manager, indicating that a refresh has occurred and waiting for a response.
afi_seq_busy	Output	4	Handshaking signal from sequencer to controller, indicating when DQS tracking is in progress.
afi_ctl_long_idle	Input	4	Handshaking signal from controller to tracking manager, indicating that it has exited low power state without a periodic refresh, and waiting for response.

3.2.7 AFI Shadow Register Management Signals

Shadow registers are a feature that enables high-speed multi-rank support. Shadow registers allow the sequencer to calibrate each rank separately, and save the calibrated settings—such as deskew delay-chain configurations—of each rank in its own set of shadow registers.

During a rank-to-rank switch, the correct set of calibrated settings is restored just in time to optimize the data valid window. The PHY relies on additional AFI signals to control which set of shadow registers to activate.

Table 221. Shadow Register Management Signals

Signal Name	Direction	Width	Description
afi_wrnk	Input	AFI_WRANK_WIDTH	Signal from controller specifying which rank the write data is going to. The signal timing is identical to that of afi_dqs_burst. That is, afi_wrnk must be asserted at the same time and must last the same duration as the afi_dqs_burst signal.
afi_rrnk	Output	AFI_RRANK_WIDTH	Signal from controller specifying which rank is being read. The signal must be asserted at the same time as the afi_rdata_en signal when issuing a read command, but unlike afi_rdata_en, afi_rrnk is stateful. That is, once asserted, the signal value must remain unchanged until the controller issues a new read command to a different rank.

Both the `afi_wrank` and `afi_rrank` signals encode the rank being accessed using the one-hot scheme (e.g. in a quad-rank interface, 0001, 0010, 0100, 1000 refer to the 1st, 2nd, 3rd, 4th rank respectively). The ordering within the bus is the same as other AFI signals. Specifically the bus is ordered by time slots, for example:

```
Half-rate afi_w/rrank = {T1, T0}
```

```
Quarter-rate afi_w/rrank = {T3, T2, T1, T0}
```

Where T_x is a number of rank-bit words that one-hot encodes the rank being accessed at the y^{th} full-rate cycle.

Additional Requirements for Shadow Register Support

To ensure that the hardware has enough time to switch from one shadow register to another, the controller must satisfy the following minimum rank-to-rank-switch delays (t_{RTRS}):

- Two read commands going to different ranks must be separated by a minimum of 3 full-rate cycles (in addition to the burst length delay needed to avoid collision of data bursts).
- Two write commands going to different rank must be separated by a minimum of 4 full-rate cycles (in addition to the burst length delay needed to avoid collision of data bursts).

The FPGA device supports a maximum of 4 sets of shadow registers, each for an independent set of timings. More than 4 ranks are supported if those ranks have four or fewer sets of independent timing. For example, the rank multiplication mode of an LRDIMM allows more than one physical rank to share a set of timing data as a single logical rank. Therefore the device can support up to 4 logical ranks, though that means more than 4 physical ranks.

3.3 AFI 4.0 Timing Diagrams

3.3.1 AFI Address and Command Timing Diagrams


Depending on the ratio between the memory clock and the PHY clock, different numbers of bits must be provided per PHY clock on the AFI interface. The following figures illustrate the AFI address/command waveforms in full, half and quarter rate respectively.

The waveforms show how the AFI command phase corresponds to the memory command output. AFI command 0 corresponds to the first memory command slot, AFI command 1 corresponds to the second memory command slot, and so on.



Figure 22. AFI Address and Command Full-Rate

Memory Interface



AFI Interface





Figure 23. AFI Address and Command Half-Rate
Memory Interface



AFI Interface

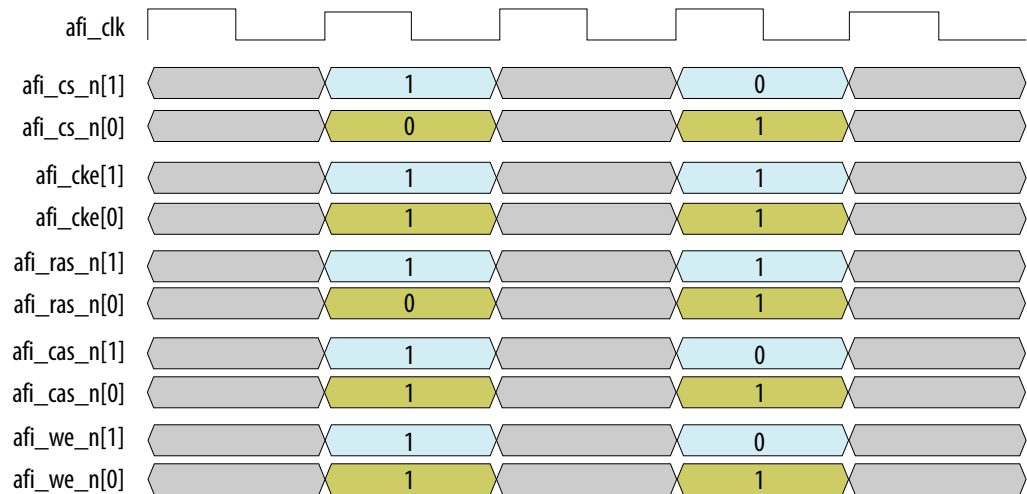






Figure 24. AFI Address and Command Quarter-Rate

Memory Interface



AFI Interface



3.3.2 AFI Write Sequence Timing Diagrams

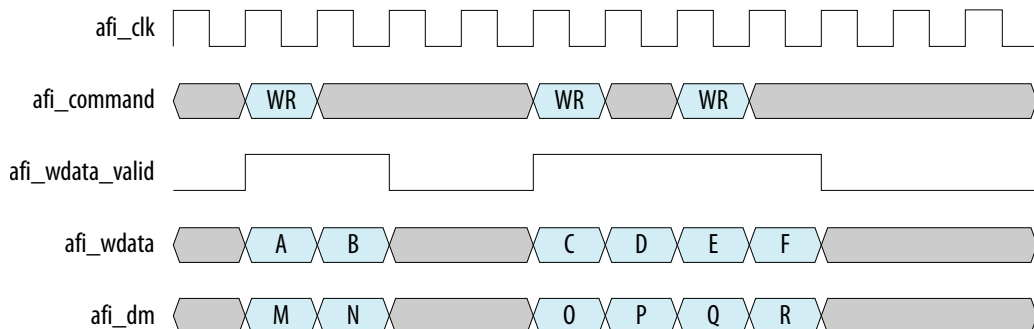
The following timing diagrams illustrate the relationships between the write command and corresponding write data and write enable signals, in full, half, and quarter rate.

For half rate and quarter rate, when the write command is sent on the first memory clock in a PHY clock (for example, `afi_cs_n[0] = 0`), that access is called *aligned access*; otherwise it is called *unaligned access*. You may use either aligned or unaligned access, or you may use both, but you must ensure that the distance

between the `write` command and the corresponding write data are constant on the AFI interface. For example, if a command is sent on the second memory clock in a PHY clock, the write data must also start at the second memory clock in a PHY clock.

Write sequences with `wlat=0`

Figure 25. AFI Write Data Full-Rate, `wlat=0`



The following diagrams illustrate both aligned and unaligned access. The first three write commands are aligned accesses where they were issued on LSB of `afi_command`. The fourth write command is unaligned access where it was issued on a different command slot. AFI signals must be shifted accordingly, based on the command slot.

Figure 26. AFI Write Data Half-Rate, `wlat=0`

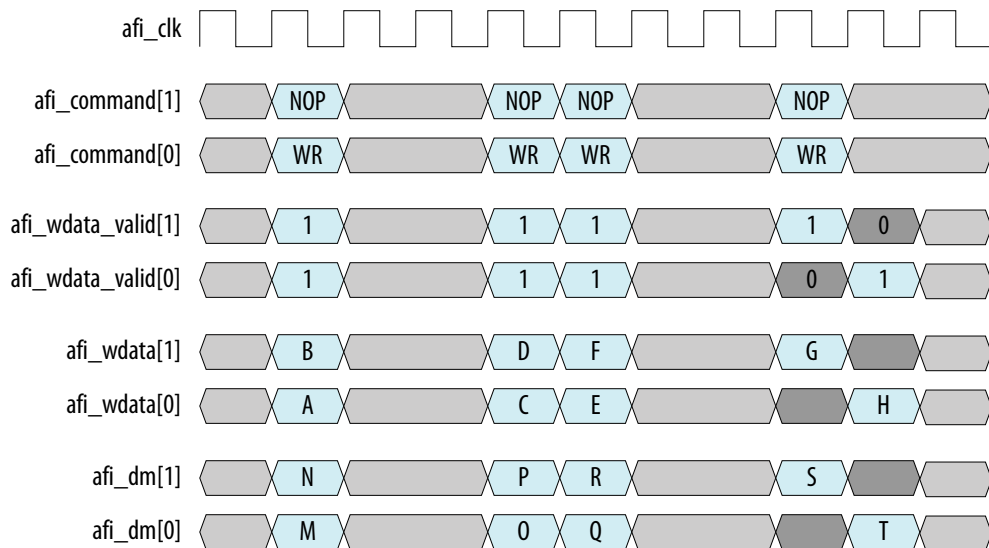
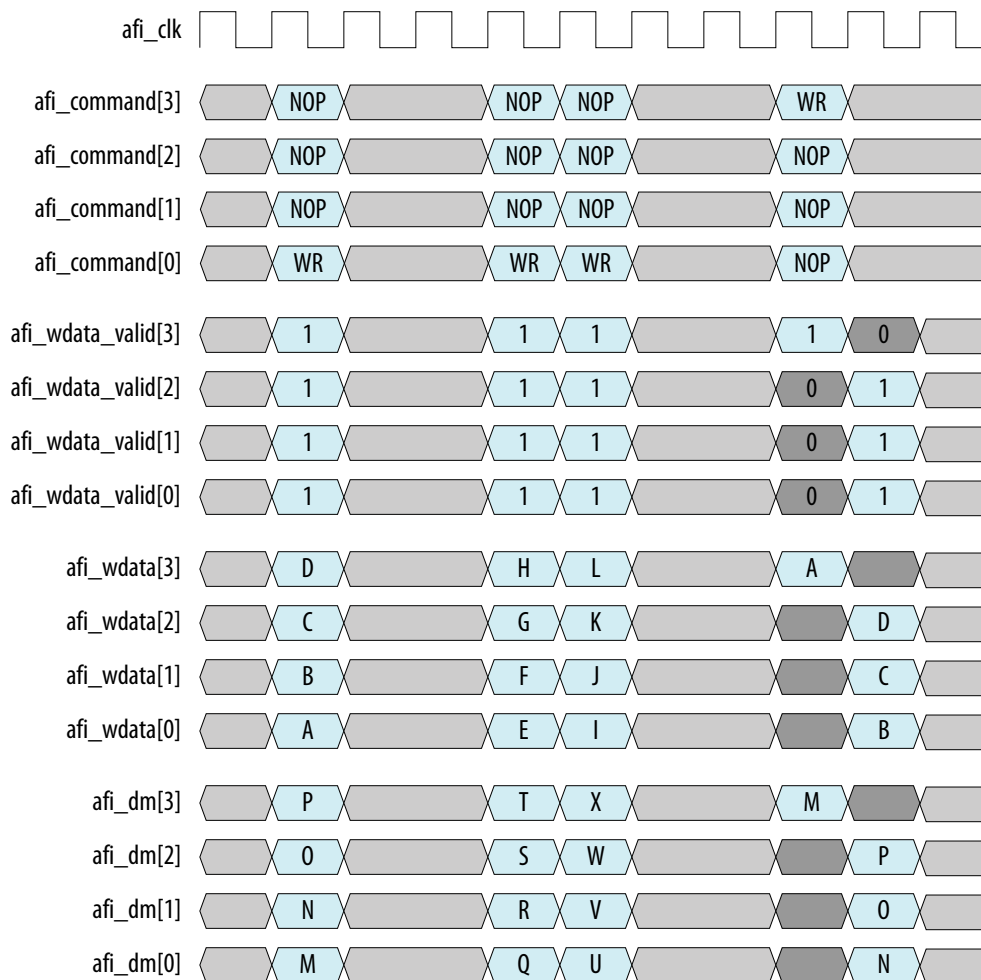




Figure 27. AFI Write Data Quarter-Rate, wlat=0



Write sequences with wlat=non-zero

The `afi_wlat` is a signal from the PHY. The controller must delay `afi_dqs_burst`, `afi_wdata_valid`, `afi_wdata` and `afi_dm` signals by a number of PHY clock cycles equal to `afi_wlat`, which is a static value determined by calibration before the PHY asserts `cal_success` to the controller. The following figures illustrate the cases when `wlat=1`. Note that `wlat` is in the number of PHY clocks and therefore `wlat=1` equals 1, 2, and 4 memory clocks delay, respectively, on full, half and quarter rate.

Figure 28. AFI Write Data Full-Rate, wlat=1

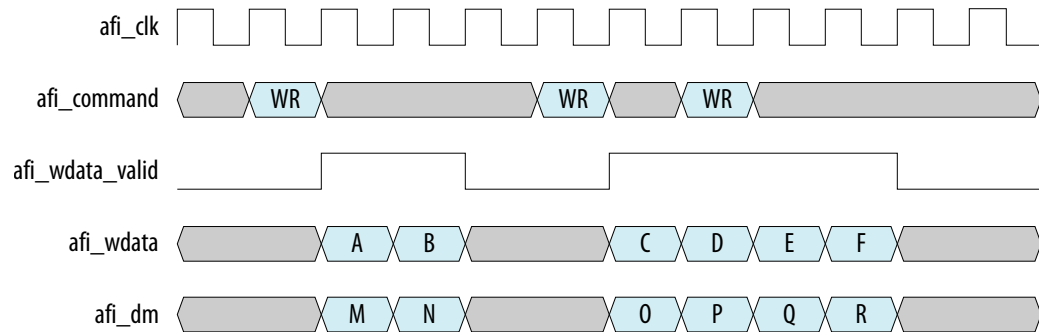


Figure 29. AFI Write Data Half-Rate, wlat=1

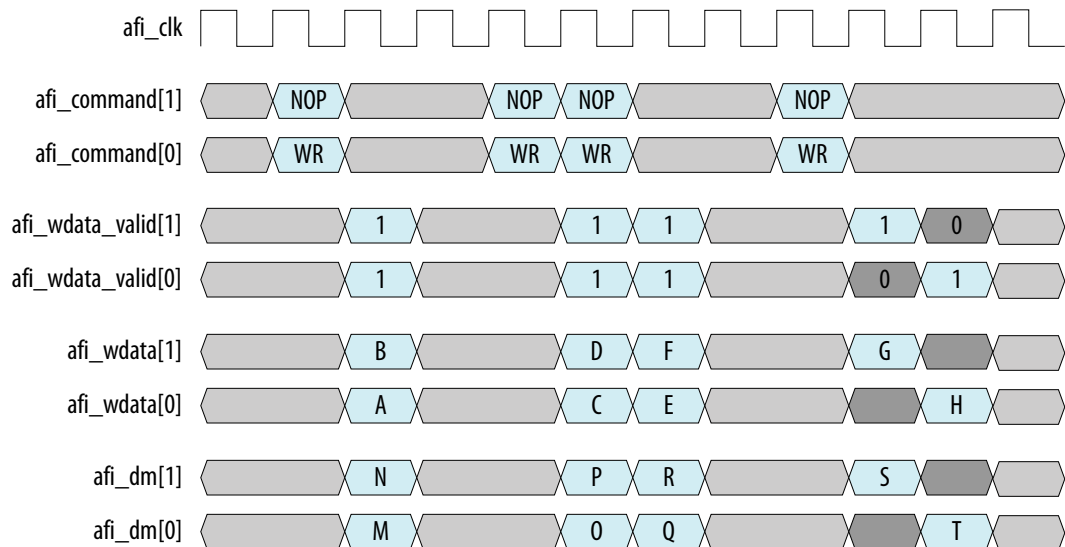
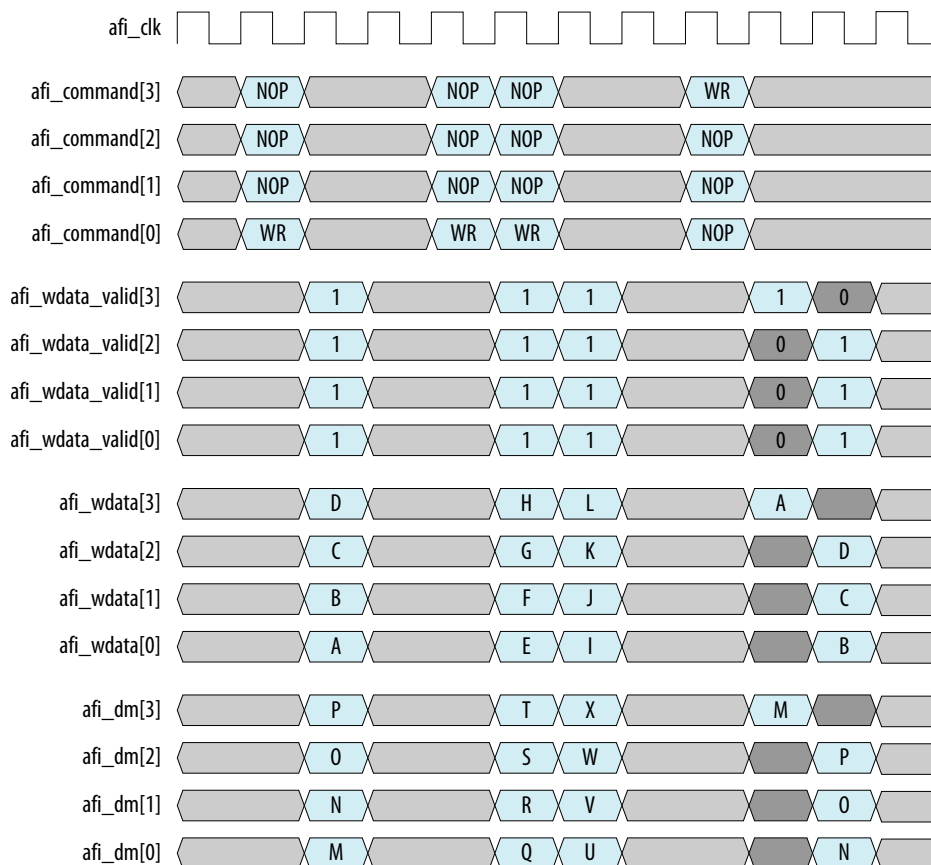




Figure 30. AFI Write Data Quarter-Rate, wlat=1



DQS burst

The `afi_dqs_burst` signal must be asserted one or two complete memory clock cycles earlier to generate DQS preamble. DQS preamble is equal to one-half and one-quarter AFI clock cycles in half and quarter rate, respectively.

A DQS preamble of two is required in DDR4, when the write preamble is set to two clock cycles.

The following diagrams illustrate how `afi_dqs_burst` must be asserted in full, half, and quarter-rate configurations.

Figure 31. AFI DQS Burst Full-Rate, wlat=1

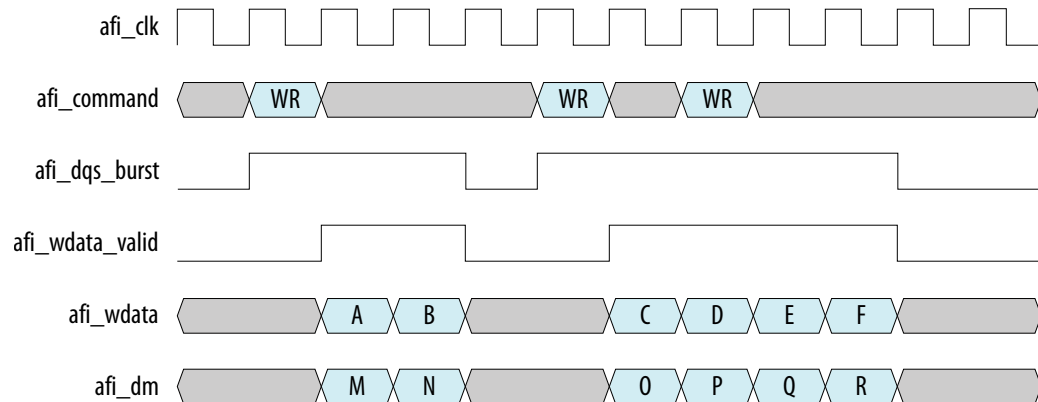


Figure 32. AFI DQS Burst Half-Rate, wlat=1

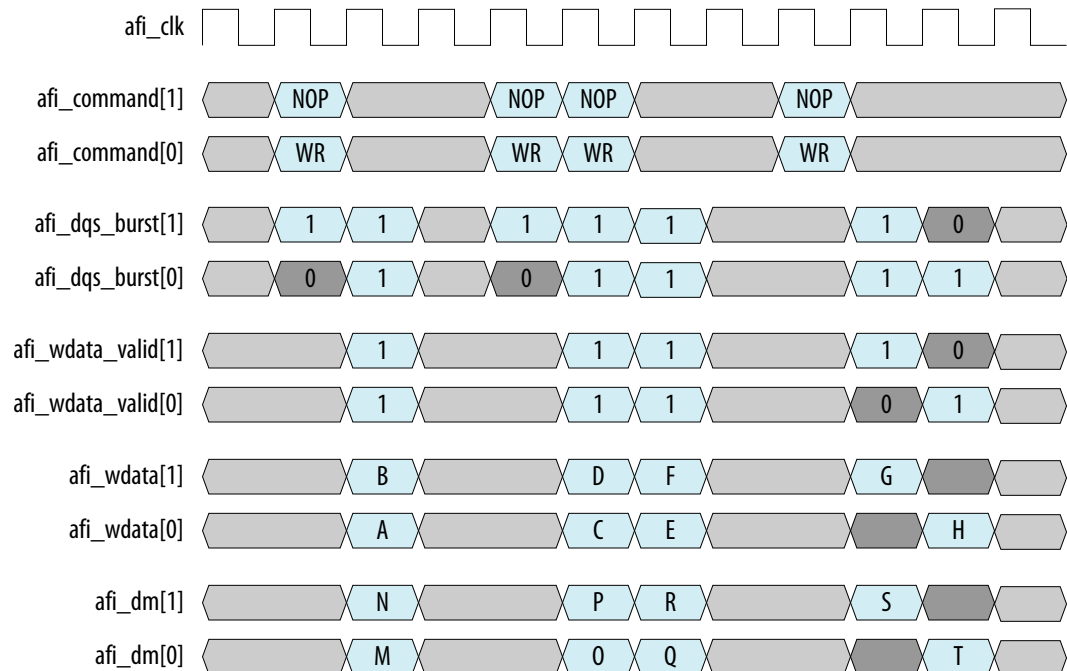
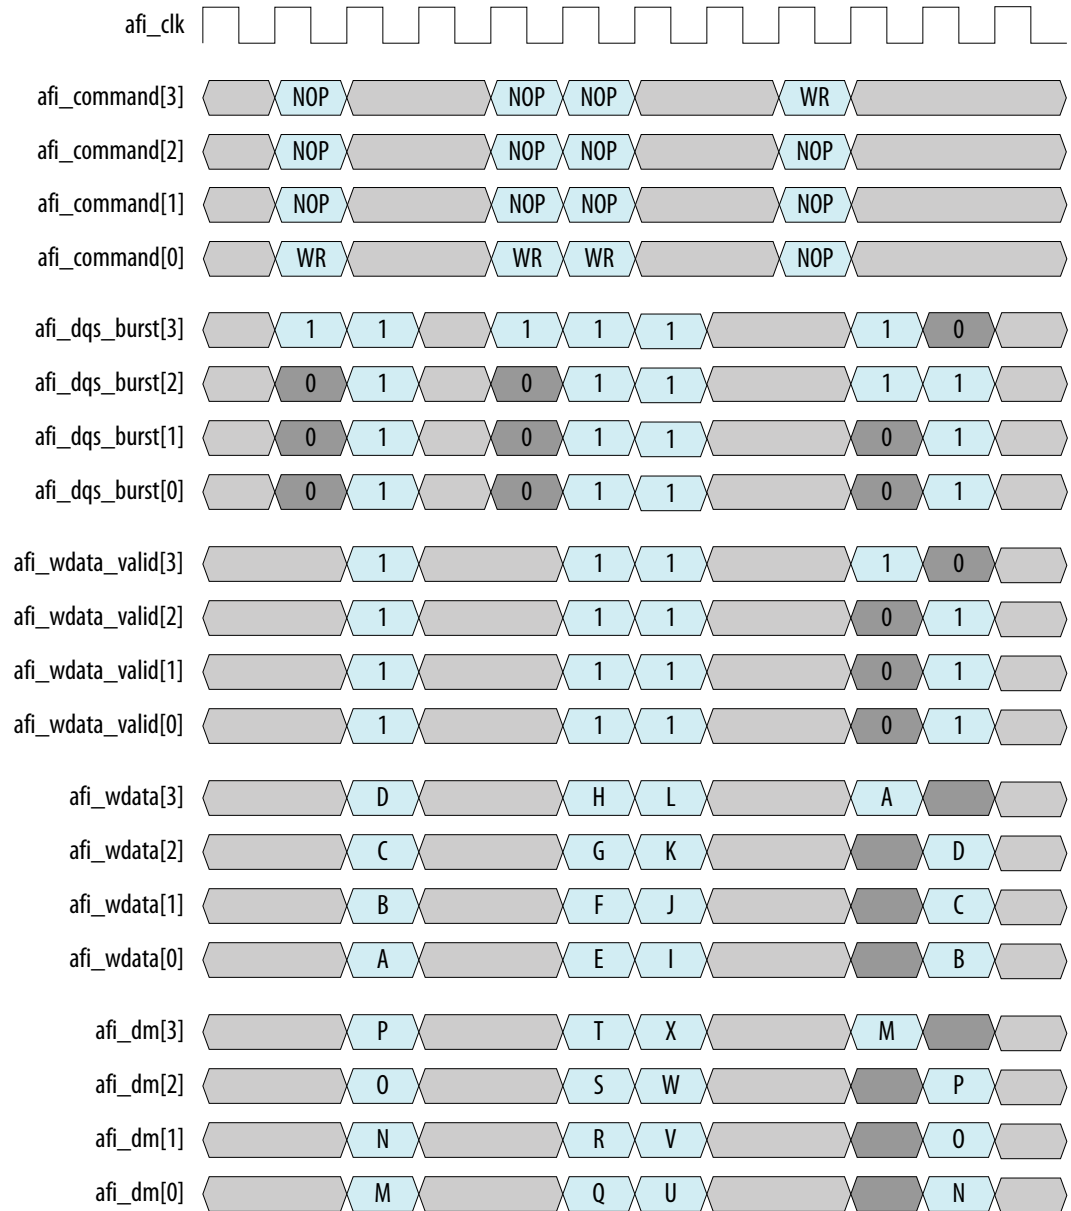




Figure 33. AFI DQS Burst Quarter-Rate, wlat=1



Write data sequence with DBI (DDR4 and QDRIV only)

The DDR4 write DBI feature is supported in the PHY, and when it is enabled, the PHY sends and receives the DBI signal without any controller involvement. The sequence is identical to non-DBI scenarios on the AFI interface.

Write data sequence with CRC (DDR4 only)

When the CRC feature of the PHY is enabled and used, the controller ensures at least one memory clock cycle between `write` commands, during which the PHY inserts the CRC data. Sending back to back `write` command would cause functional failure. The following figures show the legal sequences in CRC mode.

Entries marked as *0* and *RESERVE* must be observed by the controller; no information is allowed on those entries.

Figure 34. AFI Write Data with CRC Half-Rate, wlat=2

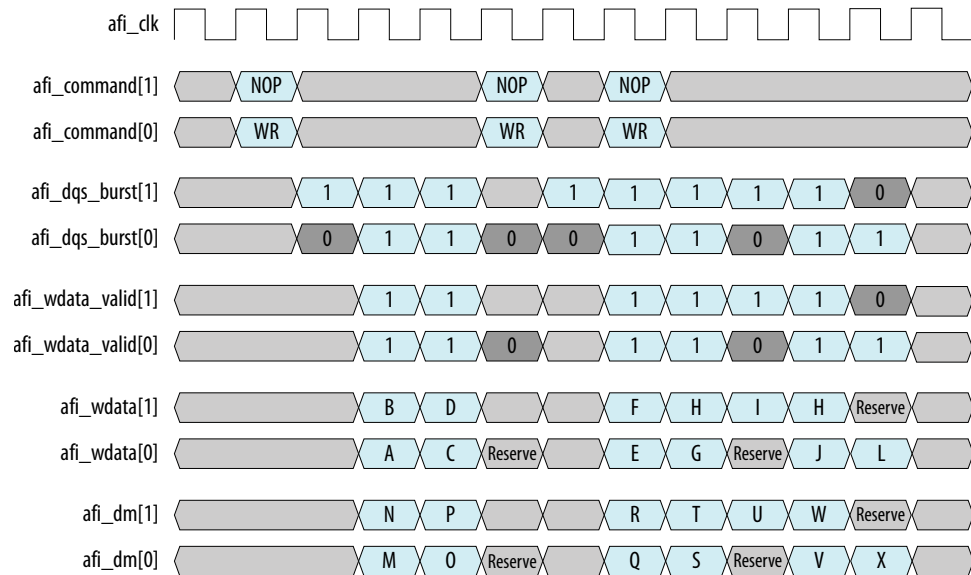
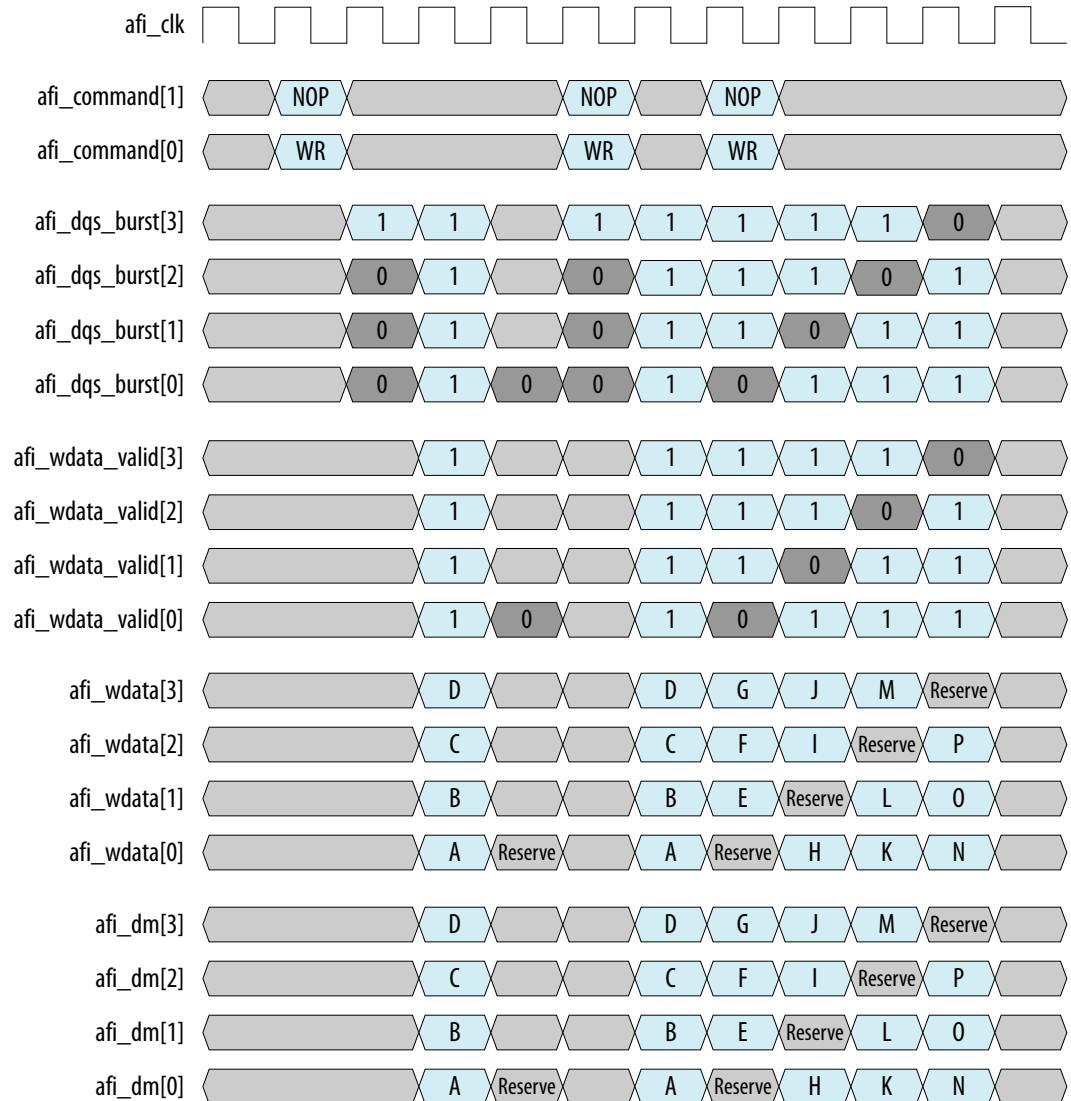




Figure 35. AFI Write Data with CRC Quarter-Rate, wlat=2



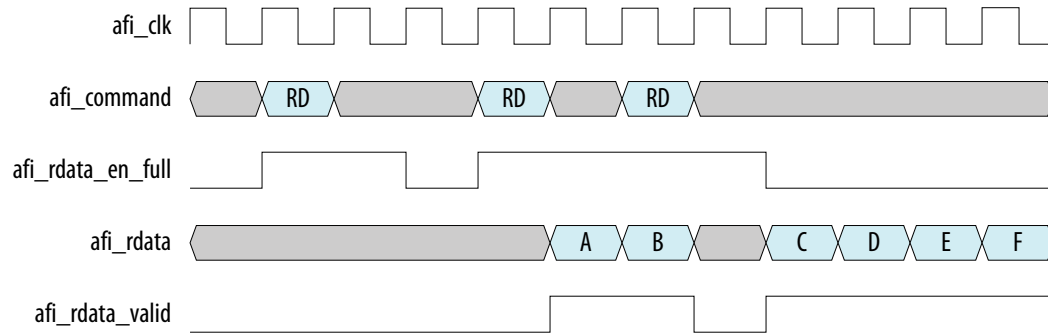
3.3.3 AFI Read Sequence Timing Diagrams

The following waveforms illustrate the AFI write data waveform in full, half, and quarter-rate, respectively.

The **afi_rdata_en_full** signal must be asserted for the entire read burst operation. The **afi_rdata_en** signal need only be asserted for the intended read data.

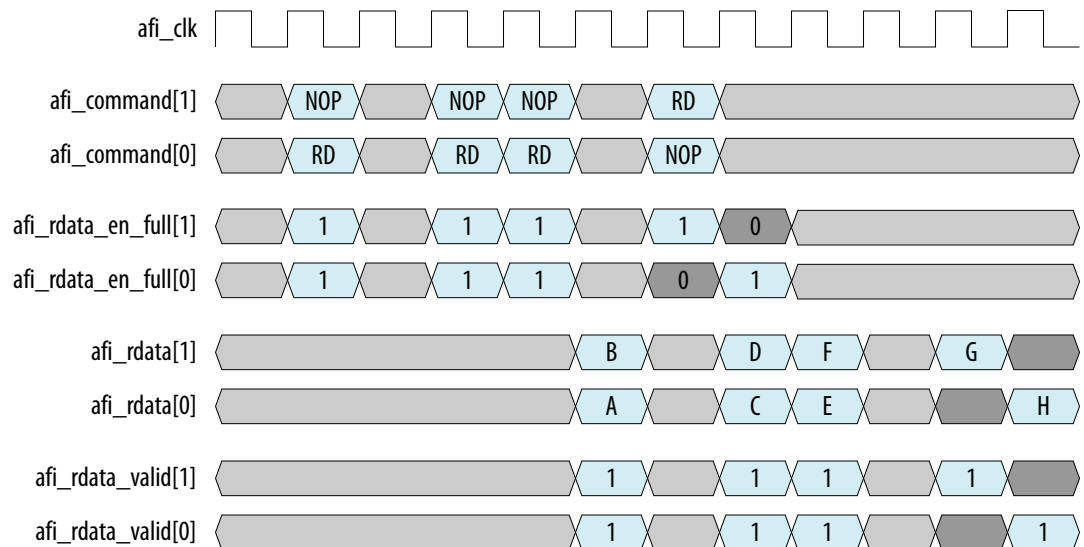
Aligned and unaligned access for read commands is similar to write commands; however, the `afi_rdata_en_full` signal must be sent on the same memory clock in a PHY clock as the read command. That is, if a read command is sent on the second memory clock in a PHY clock, `afi_rdata_en_full` must also be asserted, starting from the second memory clock in a PHY clock.

Figure 36. AFI Read Data Full-Rate



The following figure illustrates that the second and third reads require only the first and second half of data, respectively. The first three read commands are aligned accesses where they are issued on the LSB of `afi_command`. The fourth read command is unaligned access, where it is issued on a different command slot. AFI signals must be shifted accordingly, based on command slot.

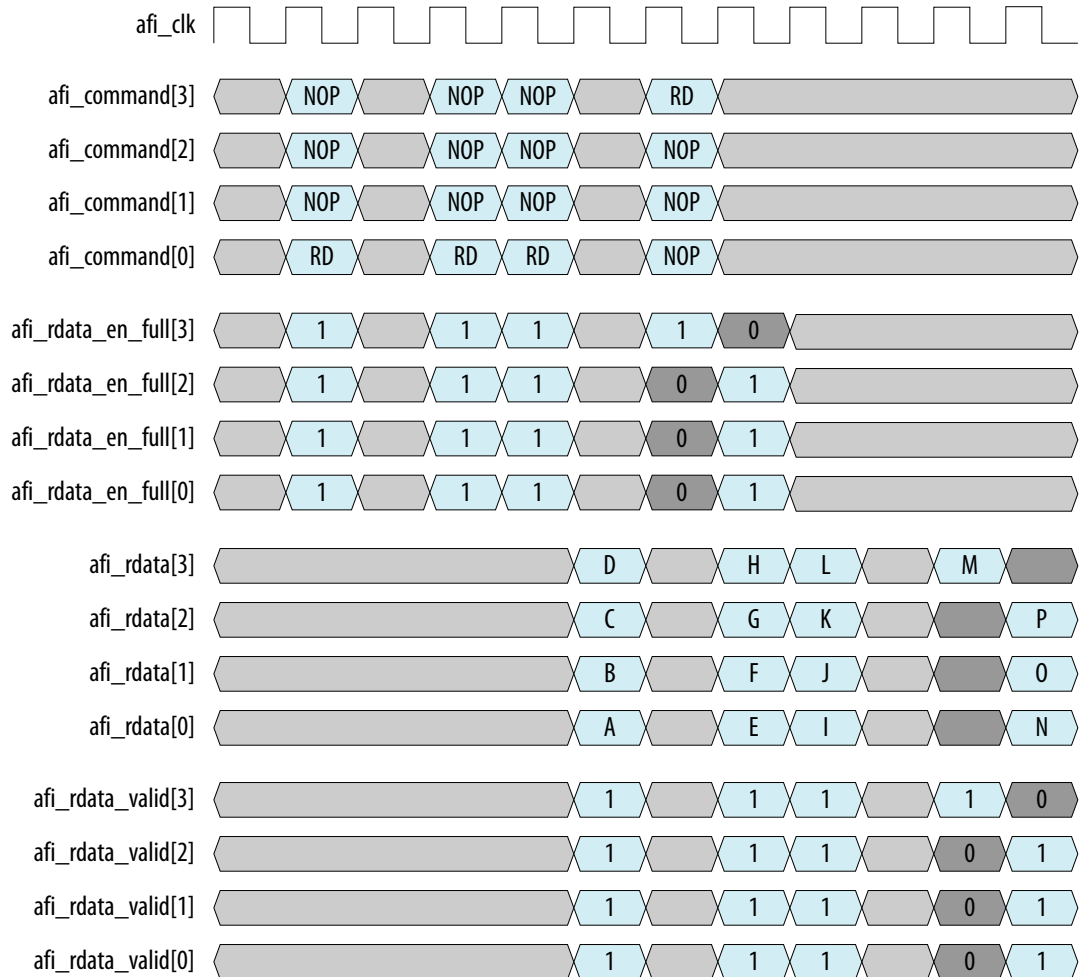
Figure 37. AFI Read Data Half-Rate



In the following figure, the first three read commands are aligned accesses where they are issued on the LSB of `afi_command`. The fourth read command is unaligned access, where it is issued on a different command slot. AFI signals must be shifted accordingly, based on command slot.



Figure 38. AFI Read Data Quarter-Rate

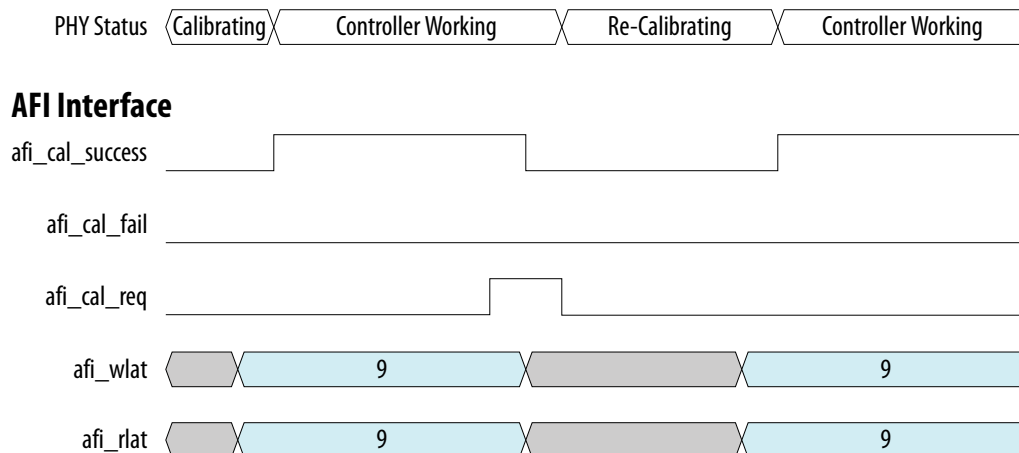


3.3.4 AFI Calibration Status Timing Diagram

The controller interacts with the PHY during calibration at power-up and at recalibration.

At power-up, the PHY holds `afi_cal_success` and `afi_cal_fail` 0 until calibration is done, when it asserts `afi_cal_success`, indicating to controller that the PHY is ready to use and `afi_wlat` and `afi_rlat` signals have valid values.

At recalibration, the controller asserts `afi_cal_req`, which triggers the same sequence as at power-up, and forces recalibration of the PHY.

Figure 39. Calibration


3.4 Intel Stratix 10 Memory Mapped Register (MMR) Tables

The address buses to read and write from the MMR registers are 10 bits wide, while the read and write data buses are configured to be 32 bits. The Bits Register Link column in the table below provides the mapping on the width of the data read within the 32-bit bus. The reads and writes are always performed using the 32-bit-wide bus.

Register Summary

Register	Address 32-bit Bus	Bits Register Link
ctrlcfg0	10	32
ctrlcfg1	11	32
dramtiming0	20	32
caltiming0	31	32
caltiming1	32	32
caltiming2	33	32
caltiming3	34	32
caltiming4	35	32
caltiming9	40	32
dramaddrw	42	32
sideband0	43	32
sideband1	44	32
sideband2	45	32
sideband3	46	32
sideband4	47	32
sideband5	48	32
continued...		



Register	Address 32-bit Bus	Bits Register Link
sideband6	49	32
sideband7	50	32
sideband8	51	32
sideband9	52	32
sideband10	53	32
sideband11	54	32
sideband12	55	32
sideband13	56	32
sideband14	57	32
dramsts	59	32
niosreserve0	68	32
niosreserve1	69	32
sideband16	79	32
ecc3	130	32
ecc4	144	32
ecc5	145	32
ecc6	146	32
ecc7	147	32
ecc8	148	32


Note: Addresses are in decimal format.

3.4.1 ctrlcfg0

address=10(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_mem_type	3	0	Indicates memory type. "0000" for DDR3 SDRAM, "0001" for DDR4 SDRAM, and "0010" for LPDDR3 SDRAM.	Read
cfg_dimm_type	6	4	Indicates dimm type.	Read
cfg_ac_pos	8	7	Indicates Command Address pin position.	Read
Reserved	31	9	Reserved.	Read

3.4.2 ctrlcfg1



address=11(32 bit)

Field	Bit High	Bit Low	Description	Access
Reserved	4	0	Reserved.	Read
cfg_addr_order	6	5	Indicates the order for address interleaving. This is related to mappings between Avalon-MM address and the SDRAM address. "00" - chip, row, bank(BG, BA), column; "01" - chip, bank(BG, BA), row, column; "10"-row, chip, bank(BG, BA), column.	Read
cfg_ctrl_enable_ecc	7	7	Enable the generation and checking of ECC.	Read
cfg_dbc0_enable_ecc	8	8	Enable the generation and checking of ECC.	Read
cfg_dbc1_enable_ecc	9	9	Enable the generation and checking of ECC.	Read
cfg_dbc2_enable_ecc	10	10	Enable the generation and checking of ECC.	Read
cfg_dbc3_enable_ecc	11	11	Enable the generation and checking of ECC.	Read
cfg_reorder_data	12	12	This bit controls whether the controller can reorder operations to optimize SDRAM bandwidth. It should generally be set to a one.	Read
cfg_ctrl_reorder_rdata	13	13	This bit controls whether the controller needs to reorder the read return data.	Read
cfg_dbc0_reorder_rdata	14	14	This bit controls whether the controller needs to reorder the read return data.	Read
cfg_dbc1_reorder_rdata	15	15	This bit controls whether the controller needs to reorder the read return data.	Read
cfg_dbc2_reorder_rdata	16	16	This bit controls whether the controller needs to reorder the read return data.	Read
cfg_dbc3_reorder_rdata	17	17	This bit controls whether the controller needs to reorder the read return data.	Read
cfg_reorder_read	18	18	This bit controls whether the controller can reorder read command.	Read
cfg_starve_limit	24	19	Specifies the number of DRAM burst transactions an individual transaction will allow to reorder ahead of it before its priority is raised in the memory controller.	Read
Reserved	25	25	Reserved.	Read
cfg_ctrl_enable_dm	26	26	Set to 1 to enable DRAM operation if DM pins are connected.	Read
cfg_dbc0_enable_dm	27	27	Set to 1 to enable DRAM operation if DM pins are connected.	Read
continued...				



Field	Bit High	Bit Low	Description	Access
cfg_dbc1_enable_dm	28	28	Set to 1 to enable DRAM operation if DM pins are connected.	Read
cfg_dbc2_enable_dm	29	29	Set to 1 to enable DRAM operation if DM pins are connected.	Read
cfg_dbc3_enable_dm	30	30	Set to 1 to enable DRAM operation if DM pins are connected.	Read

3.4.3 dramtiming0

address=20(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_tcl	6	0	Memory read latency.	Read
Reserved	31	7	Reserved.	Read

3.4.4 caltiming0

address=31(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_t_param_act_to_rdw	5	0	Activate to Read/Write command timing.	Read
cfg_t_param_act_to_pch	11	6	Active to precharge.	Read
cfg_t_param_act_to_act	17	12	Active to activate timing on same bank.	Read
cfg_t_param_act_to_act_diff_bank	23	18	Active to activate timing on different banks, for DDR4 same bank group.	Read
cfg_t_param_act_to_act_diff_bg	29	24	Active to activate timing on different bank groups, DDR4 only.	Read

3.4.5 caltiming1

address=32(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_t_param_rd_to_r	5	0	Read to read command timing on same bank.	Read
cfg_t_param_rd_to_r_diff_chip	11	6	Read to read command timing on different chips.	Read
<i>continued...</i>				



Field	Bit High	Bit Low	Description	Access
cfg_t_param_rd_to_r d_diff_bg	17	12	Read to read command timing on different chips.	Read
cfg_t_param_rd_to_ wr	23	18	Write to read command timing on same bank.	Read
cfg_t_param_rd_to_ wr_diff_chip	29	24	Read to write command timing on different chips	Read

3.4.6 caltiming2

address=33(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_t_param_rd_to_ wr_diff_bg	5	0	Read to write command timing on different bank groups.	Read
cfg_t_param_rd_to_ pch	11	6	Read to precharge command timing.	Read
cfg_t_param_rd_ap_ to_valid	17	12	Read command with autoprecharge to data valid timing.	Read
cfg_t_param_wr_to_ wr	23	18	Write to write command timing on same bank.	Read
cfg_t_param_wr_to_ wr_diff_chip	29	24	Write to write command timing on different chips.	Read

3.4.7 caltiming3

address=34(32 bit)

Field	Bit High	Bit Low	Description	Access
cfg_t_param_wr_to_ wr_diff_bg	5	0	Write to write command timing on different bank groups.	Read
cfg_t_param_wr_to_ rd	11	6	Write to read command timing.	Read
cfg_t_param_wr_to_ rd_diff_chip	17	12	Write to read command timing on different chips.	Read
cfg_t_param_wr_to_ rd_diff_bg	23	18	Write to read command timing on different bank groups.	Read
cfg_t_param_wr_to_ pch	29	24	Write to precharge command timing.	Read

3.4.8 caltiming4

**address=35(32 bit)**

Field	Bit High	Bit Low	Description	Access
cfg_t_param_wr_ap_to_valid	5	0	Write with autoprecharge to valid command timing.	Read
cfg_t_param_pch_to_valid	11	6	Precharge to valid command timing.	Read
cfg_t_param_pch_all_to_valid	17	12	Precharge all to banks being ready for bank activation command.	Read
cfg_t_param_arf_to_valid	25	18	Auto Refresh to valid DRAM command window.	Read
cfg_t_param_pdn_to_valid	31	26	Power down to valid bank command window.	Read

3.4.9 caltiming9**address=40(32 bit)**

Field	Bit High	Bit Low	Description	Access
cfg_t_param_4_act_to_act	7	0	The four-activate window timing parameter.	Read

3.4.10 dramaddrw**address=42(32 bit)**

Field	Bit High	Bit Low	Description	Access
cfg_col_addr_width	4	0	The number of column address bits for the memory devices in your memory interface.	Read
cfg_row_addr_width	9	5	The number of row address bits for the memory devices in your memory interface.	Read
cfg_bank_addr_width	13	10	The number of bank address bits for the memory devices in your memory interface.	Read
cfg_bank_group_addr_width	15	14	The number of bank group address bits for the memory devices in your memory interface.	Read
cfg_cs_addr_width	18	16	The number of chip select address bits for the memory devices in your memory interface.	Read

3.4.11 sideband0

**address=43(32 bit)**

Field	Bit High	Bit Low	Description	Access
mr_cmd_trigger	0	0	Mode Register Command Request. When asserted, indicates user request to execute mode register command. Controller clears bit to 0 when operation is completed. Register offset 37h and 38h must be properly configured before requesting Mode Register Command. Read offset 31h for Mode Register Command Status.	Read/Write

3.4.12 sideband1**address=44(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_refresh_req	3	0	Rank Refresh Request. When asserted, indicates a refresh request to the specific rank. Controller clears this bit to 0 when the refresh is executed.	Read/Write

3.4.13 sideband2**address=45(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_zqcal_long_req	0	0	Long ZQ calibration request. Asserting this bit sends a ZQ calibration command to the memory device. This is a self-clearing bit, the controller sets this bit back to 0 when the command is executed.	Read/Write

3.4.14 sideband3**address=46(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_zqcal_short_req	0	0	Short ZQ calibration request. Assertion of this bit sends the ZQ calibration command to the memory device. This is a self-clearing bit, the controller sets this bit back to 0 once the command is executed.	Read/Write

3.4.15 sideband4

**address=47(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_self_rfsh_req	3	0	Self-refresh request. When asserted, indicates a self-refresh request to DRAM. All 4 bits must be asserted or de-asserted at the same time. User clear to exit self refresh.	Read/Write

3.4.16 sideband5**address=48(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_dpd_mps_req	0	0	Deep Power Down/Maximum Power Saving request. Assertion of this bit invokes the deep power down/maximum power saving mode. You should poll for the acknowledge signal. When the acknowledge goes high, it indicates that the system has entered deep power down/maximum power saving mode. You may de-assert this bit to exit deep power down/maximum power saving mode, or keep this bit asserted to maintain deep power down/maximum power saving mode.	Read/Write

3.4.17 sideband6**address=49(32 bit)**

Field	Bit High	Bit Low	Description	Access
mr_cmd_ack	0	0	Register Command In Progress. When asserted, indicates Mode Register Command in progress.	Read

3.4.18 sideband7**address=50(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_refresh_ack	0	0	Refresh In Progress. Acknowledgement signal for refresh request. Indicates that refresh is in progress. Asserts when refresh request is sent out to PHY until $t_{RFC}/t_{param_arf_to_valid}$ is fulfilled.	Read



3.4.19 sideband8

address=51(32 bit)

Field	Bit High	Bit Low	Description	Access
mmr_zqcal_ack	0	0	ZQ Calibration in Progress. Acknowledgement signal for ZQ calibration request. When asserted, indicates that ZQ Calibration is in progress. Asserts when ZQ Calibration is sent to the PHY until the $tZQoper(t_param_zqcl_period) / tZQCS(t_param_zqcs_period)$ is fulfilled.	Read

3.4.20 sideband9

address=52(32 bit)

Field	Bit High	Bit Low	Description	Access
mmr_self_rfsh_ack	0	0	Self-refresh In Progress. Acknowledgement signal for the self-refresh request. A value of 1 indicates that memory is in self refresh mode.	Read

3.4.21 sideband10

address=53(32 bit)

Field	Bit High	Bit Low	Description	Access
mmr_dpd_mps_ack	0	0	Deep Power Down/Maximum Power Saving In Progress. Acknowledgement signal for the deep power down/maximum power saving request. A value of 1 indicates that the memory is in deep power down/maximum power saving mode.	Read

3.4.22 sideband11

address=54(32 bit)

Field	Bit High	Bit Low	Description	Access
mmr_auto_pd_ack	0	0	Auto Power Down In Progress. Acknowledgement signal for auto power down. A value of 1 indicates that the memory is in auto power down mode.	Read



3.4.23 sideband12

address=55(32 bit)

Field	Bit High	Bit Low	Description	Access
mr_cmd_type	2	0	Register command type. Indicates the type of register command.	Read/Write
			000 - Mode Register Set (DDR3, DDR4, and LPDDR3)	
			Others - Reserved	
mr_cmd_rank	6	3	Register command rank. Indicates the rank targeted by the register command.	Read/Write
			0001 - Chip select 0	
			0010 - Chip select 1	
			0011 - Chip select 0 and chip select 1	
			1111 - all chip selects	
			Mode Register Set - Any combination of chip selects.	

3.4.24 sideband13

address=56(32 bit)

Field	Bit High	Bit Low	Description	Access
mr_cmd_opcode	31	0	Register Command Opcode. Information used for register command.	Read/Write
			LPDDR3	
			[26:20] Reserved	
			[19:10] falling edge CA[9:0]	
			[9:4] rising edge CA[9:4]	
			[3:0] Reserved	
			MRW: [19:12] is OP[7:0], [11:4] is MA[7:0]	
			DDR4	
			[26:24] C2:C0	
			[23] ACT	
			[22:21] BG1:BG0	
			[20] Reserved	
			[19:18] BA1:BA0	

continued...



Field	Bit High	Bit Low	Description	Access
			[17] A17	
			[16] RAS#	
			[15] CAS#	
			[14] WE#	
			[13:0] A13:A0	
			MRS: [22:21] is BG1:BG0, [19:18] is BA1:BA0, [13:0] is Opcode[13:0]	
			DDR3	
			[26:21] Reserved	
			[20:18] BA2:BA0	
			[17] A14	
			[16] RAS#	
			[15] CAS#	
			[14] WE#	
			[13:0] A13:A0	
			MRS: [19:18] is BA1:BA0, [13:0] is Opcode[13:0]	

3.4.25 sideband14

address=57(32 bit)

Field	Bit High	Bit Low	Description	Access
mmr_refresh_cid	3	1	DDR4 3DS Chip ID Refresh. When asserted, indicates the logical rank chip ID for 3DS refresh. (This field is not applicable for DDR3 or LPDDR3.)	Read

3.4.26 dramsts

address=59(32 bit)

Field	Bit High	Bit Low	Description	Access
phy_cal_success	0	0	This bit is set to 1 if the PHY calibrates successfully.	Read
phy_cal_fail	1	1	This bit is set to 1 if the PHY does not calibrate successfully.	Read

3.4.27 niosreserve0

**address=68(32 bit)**

Field	Bit High	Bit Low	Description	Access
nios_reserve0	15	0	Indicates interface width.	Read

3.4.28 niosreserve1**address=69(32 bit)**

Field	Bit High	Bit Low	Description	Access
nios_reserve1	15	0	Indicates ACDS version.	Read

3.4.29 sideband16**address=79(32 bit)**

Field	Bit High	Bit Low	Description	Access
mmr_3ds_refresh_ack	31	0	DDR4 3DS Refresh Acknowledge. When asserted, indicates acknowledgement for the DDR4 3DS refresh.	Read
			[7:0] Refresh acknowledgement for logical rank [7:0] for physical rank 0.	
			[15:8] Refresh acknowledgement for logical rank [7:0] for physical rank 1.	
			[23:16] Refresh acknowledgement for logical rank [7:0] for physical rank 2.	
			[31:24] Refresh acknowledgement for logical rank [7:0] for physical rank 3.	

3.4.30 ecc3: ECC Error and Interrupt Configuration**address=130(32 bit)**

Field	Bit High	Bit Low	Description	Access
cfg_gen_sbe	0	0	A value of 1 enables the generate SBE feature. Generates a single bit error during the write process.	Read/Write
cfg_gen_dbe	1	1	A value of 1 enables the generate DBE feature. Generates a double bit error during the write process.	Read/Write
cfg_enable_intr	2	2	A value of 1 enables the interrupt feature. The interrupt signal notifies if an error condition occurs. The condition is configurable.	Read/Write
cfg_mask_sbe_intr	3	3	A value of 1 masks the interrupt signal when SBE occurs.	Read/Write
continued...				



Field	Bit High	Bit Low	Description	Access
cfg_mask_dbe_intr	4	4	A value of 1 masks the interrupt signal when DBE occurs.	Read/Write
cfg_mask_corr_dropped_intr	5	5	A value of 1 masks the interrupt signal when the auto correction command can't be scheduled, due to back-pressure (FIFO full).	Read/Write
cfg_mask_hmi_intr	6	6	A value of 1 masks the interrupt signal when the hard memory interface asserts an interrupt signal via the hmi_interrupt port.	Read/Write
cfg_clr_intr	7	7	Writing a value of 1 to this self-clearing bit clears the interrupt signal, error status, and address.	Read/Write
Reserved	31	8		Read

3.4.31 ecc4: Status and Error Information

address=144(32 bit)

Field	Bit High	Bit Low	Description	Access
sts_ecc_intr	0	0	Indicates the interrupt status; a value of 1 indicates an interrupt occurred.	Read
sts_sbe_error	1	1	Indicates the SBE status; a value of 1 indicates SBE occurred.	Read
sts_dbe_error	2	2	Indicates the DBE status; a value of 1 indicates DBE occurred.	Read
sts_corr_dropped	3	3	Indicates the status of correction command dropped; a value of 1 indicates correction command dropped.	Read
sts_sbe_count	7	4	Indicates the number of times SBE error has occurred. The counter will overflow.	Read
sts_dbe_count	11	8	Indicates the number of times DBE error has occurred. The counter will overflow.	Read
sts_corr_dropped_count	15	12	Indicates the number of times correction command has dropped. The counter will overflow.	Read
Reserved	31	16		Read

3.4.32 ecc5: Address of Most Recent SBE/DBE

**address=145(32 bit)**

Field	Bit High	Bit Low	Description	Access
sts_err_addr*	31	0	Address of the most recent single-bit error or double-bit error.	Read

3.4.33 ecc6: Address of Most Recent Correction Command Dropped**address=146(32 bit)**


Field	Bit High	Bit Low	Description	Access
sts_corr_dropped_addr	31	0	Address of the most recent correction command dropped.	Read

3.4.34 ecc7: Extension for Address of Most Recent SBE/DBE**address=147(32 bit)**

Field	Bit High	Bit Low	Description	Access
sts_err_addr_ext	2	0	Extension for address of the most recent single-bit error or double-bit error.	Read

3.4.35 ecc8: Extension for Address of Most Recent Correction Command Dropped**address=148(32 bit)**

Field	Bit High	Bit Low	Description	Access
sts_corr_dropped_addr_ext	2	0	Extension for address of the most recent correction command dropped.	Read



4 Intel Stratix 10 EMIF – Simulating Memory IP

To simulate your design you require the following components:

- A simulator—The simulator must be an Intel-supported VHDL or Verilog HDL simulator:
 - Aldec Riviera-Pro
 - Cadence NC Sim
 - Mentor Graphics* ModelSim
 - Synopsys* VCS/VCS-MX
- A design using Intel's External Memory Interface (EMIF) IP
- An example driver or traffic generator (to initiate read and write transactions)
- A testbench and a suitable memory simulation model

The Intel External Memory Interface IP is not compatible with the Platform Designer Testbench System. Instead, use the simulation example design from your generated IP to validate memory interface operation, or as a reference to create a full simulatable design, containing a memory interface, a memory model, and a traffic generator. For more information about the EMIF simulation design example, refer to the *Intel Stratix 10 EMIF IP Design Example User Guide*.

Memory Simulation Models

There are two types of memory simulation models that you can use:

- Intel-provided generic memory model

The Intel Quartus Prime software generates this model with the simulation example design. The model adheres to all the memory protocol specifications, and can be parameterized.

- Vendor-specific memory model

Vendor-specific memory models are simulation models for specific memory components from memory vendors such as Micron and Samsung. You can obtain these simulation models from the memory vendor's website.

Note: Intel does not provide support for vendor-specific memory models.

Related Links

[Modifying the Example Driver to Replicate the Failure](#) on page 426

4.1 Simulation Options

The following simulation options are available with the example testbench to improve simulation speed:

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

ISO
9001:2008
Registered



- Full calibration—Calibrates the same way as in hardware, and includes all phase sweeps, delay adjustments, and data centering.

Note: Intel Stratix 10 EMIF full calibration simulation will be available in a future release of the Intel Quartus Prime software.

- Skip calibration—Loads memory configuration settings and enters user mode, providing the fastest simulation time.

Note: For proper simulation of DQS Tracking, you must enable full calibration.

Both simulation options represent accurate controller efficiency and do not take into account board skew. This may cause a discrepancy in the simulated interface latency numbers. For more information regarding simulation assumptions and differences between RTL simulation and post-fit implementation, refer to the *Intel Stratix 10 EMIF IP Design Example User Guide*.

Table 222. Typical Simulation Times Using Intel Stratix 10 EMIF IP

Calibration Mode/Run Time ⁽¹⁾	Estimated Simulation Time	
	Small Interface (×8 Single Rank)	Large Interface (×72 Quad Rank)
Full <ul style="list-style-type: none"> • Full calibration • Includes all phase/delay sweeps and centering 	20 minutes	~ 1 day
Skip <ul style="list-style-type: none"> • Skip calibration • Preloads calculated settings 	10 minutes	25 minutes
Abstract PHY <ul style="list-style-type: none"> • Replace PHY and external memory model with a single abstract PHY model. • IMPORTANT: <i>External memory model is NOT used in this mode. No I/O switching occurs to the external memory model.</i> 	1 minute	5 minutes
Note to Table: 1. Uses one loop of driver test. One loop of driver is approximately 600 read or write requests, with burst length up to 64. 2. Simulation times shown in this table are approximate measurements made using Synopsys VCS. Simulation times can vary considerably, depending on the IP configuration, the simulator used, and the computer or server used.		

Related Links

[Simulation Walkthrough](#) on page 141

4.2 Simulation Walkthrough

Simulation is a good way to determine the latency of your system. However, the latency reflected in simulation may be different than the latency found on the board because functional simulation does not take into account board trace delays and different process, voltage, and temperature scenarios.

For a given design on a given board, the latency found may differ by one clock cycle (for full-rate designs) or two clock cycles (for half-rate designs) upon resetting the board. Different boards can also show different latencies even with the same design.

The Intel Stratix 10 EMIF IP supports functional simulation only. Functional simulation is supported at the RTL level after generating a post-fit functional simulation netlist. The post-fit netlist for designs that contain Intel Stratix 10 EMIF IP is a hybrid of the gate level (for FPGA core) and RTL level (for the external memory interface IP). You should validate the functional operation of your design using RTL simulation, and the timing of your design using timing analysis.

To perform functional simulation for an Intel Stratix 10 EMIF IP example design, locate the example design files in the `\<variation_name>_example_design` directory.

You can use the IP functional simulation model with any supported VHDL or Verilog HDL simulator.

After you have generated the memory IP, view the README.txt file located in the `\<variation_name>_example_design` directory for instructions on how to generate the simulation example design for Verilog HDL or VHDL. Simulation filesets for both Verilog HDL and VHDL are located in `\<variation_name>_example_design\sim`. The **README.txt** file also explains how to run simulation using the ModelSim* - Intel FPGA Edition. Simulation scripts for the Mentor Graphics, Cadence, Aldec, and Synopsys simulators are provided; however, detailed instructions on how to perform simulation using these third-party simulators are not provided.

Related Links

[Simulation Options](#) on page 140

4.2.1 Calibration Modes

Calibration occurs shortly after the memory device is initialized, to compensate for uncertainties in the hardware system, including silicon PVT variation, circuit board trace delays, and skewed arrival times. Such variations are usually not present in an RTL simulation environment, resulting in two simulatable calibration modes: Skip Calibration mode (which is the default), and Full Calibration mode.

Skip Calibration Mode

In Skip Calibration mode, the calibration processor assumes an ideal hardware environment, where PVT variations, board delays, and trace skews are all zero. Instead of running the actual calibration routine, the calibration processor calculates the expected arrival time of read data based on the memory latency values entered during EMIF IP generation, resulting in reduced simulation time. Skip calibration mode is recommended for use during system development, because it allows you to focus on interacting with the controller and optimizing your memory access patterns, thus facilitating rapid RTL development.

Full Calibration Mode

Full Calibration mode simulates every stage of the calibration algorithm immediately after memory device initialization. The calibration algorithm processes each data group sequentially and each pin in each group individually, causing simulation time to increase with the number of data pins in your interface. You can observe how the calibration algorithm compensates for various delays in the system by incorporating your own board delay model based on trace delays from your PCB design tools. Due to the large simulation overhead, Full Calibration simulation mode is not recommended for rapid development of IP cores.



VHDL Support

VHDL support for mixed-language simulators is implemented by generating the top-level wrapper for the core in VHDL, while all submodules are provided as clear text SystemVerilog files.

A set of precompiled device libraries is provided for use with the ModelSim - Intel FPGA Edition simulator, which is supplied with the Intel Quartus Prime software. Submodules normally provided as cleartext SystemVerilog files are encrypted using IEEE Verilog HDL encryption for ModelSim - Intel FPGA Edition.

4.2.2 Abstract PHY Simulation

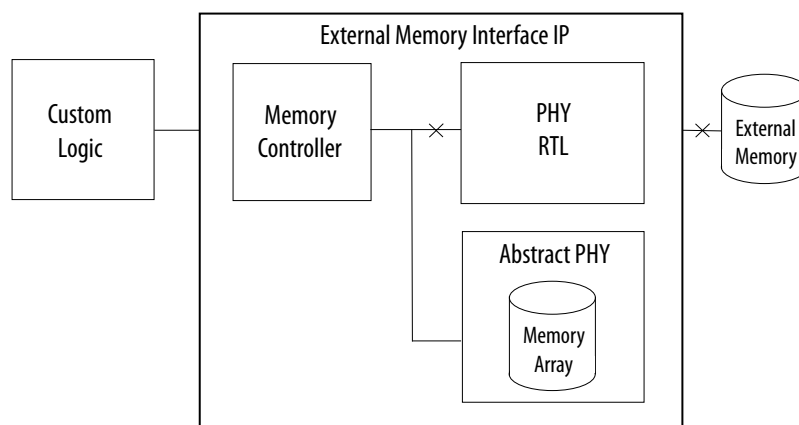
The abstract PHY is a simulation model of the EMIF PHY that can decrease simulation time by 3-10 times. The abstract PHY replaces the lane and the external memory model with a single model containing an internal memory array. No switching of the I/Os to the external memory model occurs when simulating with the abstract PHY.

Abstract PHY reduces simulation time by two mechanisms:

- The Nios processor has been disabled and is replaced by HDL forces that are applied at the beginning of simulation. The HDL forces are a minimum set of registers that configures the memory interface for simulation. The write and read latency values applied by the HDL forces are not representative of the post-calibration values applied to the memory interface running on hardware. However, as long as the customer logic is Avalon® and AFI-compliant, these values allow for successful RTL simulation.
- The abstract PHY eliminates the need for full-speed clocks and therefore simulation of the abstract PHY does not require full-speed clock simulation events.

To use the abstract PHY, enable **Simulation Options > Abstract PHY for fast simulation** on the **Diagnostic** tab during EMIF IP generation. When you enable Abstract PHY, the EMIF IP is configured as shown below. The PHY RTL and external memory model are disconnected from the data path and in their place is the abstract PHY containing an internal memory array.

Figure 40. Abstract PHY



Note: You cannot observe the external memory device signals when you are using Abstract PHY.

Note: Abstract PHY does not reflect accurate latency numbers.

4.2.3 Simulation Scripts

The Intel Quartus Prime software generates three simulation scripts during project generation for four different third party simulation tools—Cadence, Synopsys, Aldec, and Mentor Graphics.

The simulation scripts are located under the `<project_directory>\<variation_name>_sim` directory, in separate folders named after each supported simulator. The example designs also provide equivalent scripts after you run the `.tcl` script from the project located in the `<variation_name>_example_design\sim` directory.

4.2.4 Functional Simulation with Verilog HDL

Simulation scripts for the Synopsys, Cadence, Aldec, and Mentor Graphics simulators are provided for you to run the example design.

The simulation scripts are located in the following main folder locations:

Simulation scripts in the simulation folders are located as follows:

- `<variation_name>_example_design\sim\mentor\msim_setup.tcl`
- `<variation_name>_example_design\sim\synopsys\vcs\vcs_setup.sh`
- `<variation_name>_example_design\sim\synopsys\vcsmx\vcsmx_setup.sh`
- `<variation_name>_example_design\sim\aldec\rivierapro_setup.tcl`
- `<variation_name>_example_design\sim\cadence\ncsim_setup.sh`

Simulation scripts in the `<>_sim_folder` are located as follows:

- `<variation_name>_sim\mentor\msim_setup.tcl`
- `<variation_name>_sim\cadence\ncsim_setup.sh`
- `<variation_name>_sim\synopsys\vcs\vcs_setup.sh`
- `<variation_name>_sim\synopsys\vcsmx\vcsmx_setup.sh`
- `<variation_name>_sim\aldec\rivierapro_setup.tcl`

For more information about simulating Verilog HDL or VHDL designs using command lines, refer to the *Mentor Graphics ModelSim and QuestaSim Support* chapter in volume 3 of the Intel Quartus Prime Handbook.

Related Links

[ModelSim - Intel FPGA Edition, ModelSim, and QuestaSim Support](#)

4.2.5 Functional Simulation with VHDL

The EMIF VHDL fileset is provided for customers that wish to generate the top-level RTL instance of their EMIF IP cores in VHDL.



Prior to Intel Quartus Prime version 15.1, the VHDL fileset was comprised entirely of VHDL files. Beginning with Intel Quartus Prime version 15.1, only the top-level IP instance file is guaranteed to be written in VHDL; submodules can still be deployed as Verilog/SystemVerilog (encrypted or plaintext) files, or VHDL files. Note that the ModelSim - Intel FPGA Edition is no longer restricted to a single HDL language as of Intel Quartus Prime 15.1; however, some files may still be encrypted in order to be excluded from the maximum unencrypted module limit of this tool.

Because the VHDL fileset consists of both VHDL and Verilog files, you must follow certain mixed-language simulation guidelines. The general guideline for mixed-language simulation is that you must always link the Verilog files (whether encrypted or not) against the Verilog version of the libraries, and the VHDL files (whether SimGen-generated or pure VHDL) against the VHDL libraries.

Simulation scripts for the Synopsys, Cadence, Aldec, and Mentor Graphics simulators are provided for you to run the example design. These simulation scripts are located in the following main folder locations:

Simulation scripts in the simulation folders are located as follows:

- `<variation_name>_example_design\sim\mentor\msim_setup.tcl`
- `<variation_name>_example_design\sim\synopsys\vcsmx\vcsmx_setup.sh`
- `<variation_name>_example_design\sim\synopsys\vcs\vcs_setup.sh`
- `<variation_name>_example_design\sim\cadence\ncsim_setup.sh`
- `<variation_name>_example_design\sim\aldec\rivierapro_setup.tcl`

Simulation scripts in the `<>_sim_folder` are located as follows:

- `<variation_name>_sim\mentor\msim_setup.tcl`
- `<variation_name>_sim\cadence\ncsim_setup.sh`
- `<variation_name>_sim\synopsys\vcsmx\vcsmx_setup.sh`
- `<variation_name>_sim\aldec\rivierapro_setup.tcl`

For more information about simulating Verilog HDL or VHDL designs using command lines, refer to the *Mentor Graphics ModelSim and QuestaSim Support* chapter in volume 3 of the Intel Quartus Prime Handbook.

Related Links

[ModelSim - Intel FPGA Edition, ModelSim, and QuestaSim Support](#)

4.2.6 Simulating the Example Design

This topic describes how to simulate the example design in Cadence, Synopsys, Mentor Graphics, and Aldec simulators.



To simulate the example design in the Intel Quartus Prime software using the Cadence simulator, follow these steps:

1. At the Linux* shell command prompt, change directory to
`<name>_example_design\sim\cadence`
2. Run the simulation by typing the following command at the command prompt:

```
sh ncsim_setup.sh
```

To simulate the example design in the Intel Quartus Prime software using the Synopsys simulator, follow these steps:

1. At the Linux shell command prompt, change directory to
`<name>_example_design\sim\synopsys\vcsmx`
2. Run the simulation by typing the following command at the command prompt:

```
sh vcsmx_setup.sh
```

To simulate the example design in the Intel Quartus Prime software using the Mentor Graphics simulator, follow these steps:

1. At the Linux or Windows shell command prompt, change directory to
`<name>_example_design\sim\mentor`
2. Execute the **msim_setup.tcl** script that automatically compiles and runs the simulation by typing the following command at the Linux or Windows command prompt:

```
vsim -do msim_setup.tcl
```

or

Type the following command at the ModelSim command prompt:

```
do msim_setup.tcl
```


Note: Intel does not provide the `run.do` file for the example design with the EMIF interface.

To simulate the example design in the Intel Quartus Prime software using the Aldec simulator, follow these steps:

1. At the Linux or Windows shell command prompt, change directory to
`<name>_example_design\sim\aldec`
2. Execute the **rivierapro_setup.tcl** script that automatically compiles and runs the simulation by typing the following command at the Linux or Windows command prompt:
`vsim -do rivierapro.tcl`
3. To compile and elaborate the design after the script loads, type `ld_debug`.
4. Type `run -all` to run the simulation.

For more information about simulation, refer to the *Simulating Designs* chapter in volume 3 of the Intel Quartus Prime Handbook.

If your Intel Quartus Prime project appears to be configured correctly but the example testbench still fails, check the known issues on the Intel FPGA Knowledge Base before filing a service request.



Related Links

- [Simulating Intel FPGA Designs](#)
- [Intel FPGA Knowledge Base](#)



5 Intel Stratix 10 EMIF IP for DDR3

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for DDR3.

5.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

5.1.1 Intel Stratix 10 EMIF IP DDR3 Parameters: General

Table 223. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_DDR3_CONFIG_ENUM)
Instantiate two controllers sharing a Ping Pong PHY	Specifies the instantiation of two identical memory controllers that share an address/command bus through the use of Ping Pong PHY. This parameter is available only if you specify the Hard PHY and Hard Controller option. When this parameter is enabled, the IP exposes two independent Avalon interfaces to the user logic, and a single external memory interface with double width for the data bus and the CS#, CKE, ODT, and CK/CK# signals. (Identifier: PHY_DDR3_USER_PING_PONG_EN)

Table 224. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_DDR3_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_DDR3_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to
continued...	

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.



Display Name	Description
	use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_DDR3_USER_REF_CLK_FREQ_MHZ)
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_DDR3_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_DDR3_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_DDR3_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 225. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 226. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)



Table 227. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)

Table 228. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 229. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

5.1.2 Intel Stratix 10 EMIF IP DDR3 Parameters: Memory

Table 230. Group: Memory / Topology

Display Name	Description
Memory format	Specifies the format of the external memory device. The following formats are supported: Component - a Discrete memory device; UDIMM - Unregistered/Unbuffered DIMM where address/control, clock, and data are unbuffered; RDIMM - Registered DIMM where address/control and clock are buffered; LRDIMM - Load Reduction DIMM where address/control, clock, and data are buffered. LRDIMM reduces the load to increase memory speed and supports higher densities than RDIMM; SODIMM - Small Outline DIMM is similar to UDIMM but smaller in size and is typically used for systems with limited space. Some memory protocols may not be available in all formats. (Identifier: MEM_DDR3_FORMAT_ENUM)
DQ width	Specifies the total number of data pins in the interface. The maximum supported width is 144, or 72 in Ping Pong PHY mode. (Identifier: MEM_DDR3_DQ_WIDTH)
DQ pins per DQS group	Specifies the total number of DQ pins per DQS group. (Identifier: MEM_DDR3_DQ_PER_DQS)
Number of clocks	Specifies the number of CK/CK# clock pairs exposed by the memory interface. Usually more than 1 pair is required for RDIMM/LRDIMM formats. The value of this parameter depends on the memory device selected; <i>refer to the data sheet for your memory device</i> . (Identifier: MEM_DDR3_CK_WIDTH)
Number of chip selects	Specifies the total number of chip selects in the interface, up to a maximum of 4. This parameter applies to discrete components only . (Identifier: MEM_DDR3_DISCRETE_CS_WIDTH)
Number of DIMMs	Total number of DIMMs. (Identifier: MEM_DDR3_NUM_OF_DIMMS)
continued...	



Display Name	Description
Number of physical ranks per DIMM	Number of ranks per DIMM. For LRDIMM, this represents the number of physical ranks on the DIMM behind the memory buffer (Identifier: MEM_DDR3_RANKS_PER_DIMM)
Number of clock enables per DIMM	Number of clock enables (i.e. CKE) per DIMM. Only applicable to registered and load-reduced DIMMs. (Identifier: MEM_DDR3_CKE_PER_DIMM)
Row address width	Specifies the number of row address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of address pins needed for access to all available rows. (Identifier: MEM_DDR3_ROW_ADDR_WIDTH)
Column address width	Specifies the number of column address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of address pins needed for access to all available columns. (Identifier: MEM_DDR3_COL_ADDR_WIDTH)
Bank address width	Specifies the number of bank address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of bank address pins needed for access to all available banks. (Identifier: MEM_DDR3_BANK_ADDR_WIDTH)
Enable DM pins	Indicates whether the interface uses data mask (DM) pins. This feature allows specified portions of the data bus to be written to memory (not available in x4 mode). One DM pin exists per DQS group. (Identifier: MEM_DDR3_DM_EN)
Enable address mirroring for odd chip-selects	Enabling address mirroring for multi-CS discrete components. Typically used when components are arranged in a clamshell layout. (Identifier: MEM_DDR3_DISCRETE_MIRROR_ADDRESSING_EN)
Enable address mirroring for odd ranks	Enabling address mirroring for dual-rank or quad-rank DIMM. (Identifier: MEM_DDR3_MIRROR_ADDRESSING_EN)
ALERT# pin placement	Specifies placement for the mem_alert_n signal. If you select " I/O Lane with Address/Command Pins ", you can pick the I/O lane and pin index in the add/cmd bank with the subsequent drop down menus. If you select " I/O Lane with DQS Group ", you can specify the DQS group with which to place the mem_alert_n pin. If you select " Automatically select a location ", the IP automatically selects a pin for the mem_alert_n signal. If you select this option, no additional location constraints can be applied to the mem_alert_n pin, or a fitter error will result during compilation. For optimum signal integrity, you should choose " I/O Lane with Address/Command Pins ". For interfaces containing multiple memory devices, it is recommended to connect the ALERT# pins together to the ALERT# pin on the FPGA. (Identifier: MEM_DDR3_ALERT_N_PLACEMENT_ENUM)
DQS group of ALERT#	Select the DQS group with which the ALERT# pin is placed. (Identifier: MEM_DDR3_ALERT_N_DQS_GROUP)

Table 231. Group: Memory / Latency and Burst

Display Name	Description
Memory CAS latency setting	Specifies the number of clock cycles between the read command and the availability of the first bit of output data at the memory device. Overall read latency equals the additive latency (AL) + the CAS latency (CL). <i>Overall read latency depends on the memory device selected; refer to the datasheet for your device.</i> (Identifier: MEM_DDR3_TCL)
Memory write CAS latency setting	Specifies the number of clock cycles from the release of internal write to the latching of the first data in at the memory device. <i>This value depends on the memory device selected; refer to the datasheet for your device.</i> (Identifier: MEM_DDR3_WTCL)
Memory additive CAS latency setting	Determines the posted CAS additive latency of the memory device. Enable this feature to improve command and bus efficiency, and increase system bandwidth. (Identifier: MEM_DDR3_ATCL_ENUM)



Table 232. Group: Memory / Mode Register Settings

Display Name	Description
Hide advanced mode register settings	Show or hide advanced mode register settings. Changing advanced mode register settings to non-default values is strongly discouraged. (Identifier: MEM_DDR3_HIDE_ADV_MR_SETTINGS)
Burst Length	Specifies the DRAM burst length which determines how many consecutive addresses should be accessed for a given read/write command. (Identifier: MEM_DDR3_BL_ENUM)
Read Burst Type	Indicates whether accesses within a given burst are in sequential or interleaved order. Select sequential if you are using the Intel-provided memory controller. (Identifier: MEM_DDR3_BT_ENUM)
DLL precharge power down	Specifies whether the DLL in the memory device is off or on during precharge power-down (Identifier: MEM_DDR3_PD_ENUM)
Enable the DLL in memory device	Enable the DLL in memory device (Identifier: MEM_DDR3_DLL_EN)
Auto self-refresh method	Indicates whether to enable or disable auto self-refresh. Auto self-refresh allows the controller to issue self-refresh requests, rather than manually issuing self-refresh in order for memory to retain data. (Identifier: MEM_DDR3_ASR_ENUM)
Self-refresh temperature	Specifies the self-refresh temperature as "Normal" or "Extended" mode. <i>More information on Normal and Extended temperature modes can be found in the memory device datasheet.</i> (Identifier: MEM_DDR3_SRT_ENUM)
DDR3 RDIMM/LRDIMM control words	<i>Each 4-bit/8-bit setting can be obtained from the manufacturer's data sheet and should be entered in hexadecimal, starting with the 8-bit setting RCBx on the left and continuing to RC1x followed by the 4-bit setting RCOF and ending with RC00 on the right</i> (Identifier: MEM_DDR3_RDIMM_CONFIG)
DDR3 LRDIMM additional control words	<i>Each 4-bit setting can be obtained from the manufacturer's data sheet and should be entered in hexadecimal, starting with BC0F on the left and ending with BC00 on the right</i> (Identifier: MEM_DDR3_LRDIMM_EXTENDED_CONFIG)

5.1.3 Intel Stratix 10 EMIF IP DDR3 Parameters: Mem I/O

Table 233. Group: Mem I/O / Memory I/O Settings

Display Name	Description
Output drive strength setting	Specifies the output driver impedance setting at the memory device. <i>To obtain optimum signal integrity performance, select option based on board simulation results.</i> (Identifier: MEM_DDR3_DRV_STR_ENUM)
ODT Rtt nominal value	Determines the nominal on-die termination value applied to the DRAM. The termination is applied any time that ODT is asserted. If you specify a different value for RTT_WR, that value takes precedence over the values mentioned here. <i>For optimum signal integrity performance, select your option based on board simulation results.</i> (Identifier: MEM_DDR3_RTT_NOM_ENUM)
Dynamic ODT (Rtt_WR) value	Specifies the mode of the dynamic on-die termination (ODT) during writes to the memory device (used for multi-rank configurations). <i>For optimum signal integrity performance, select this option based on board simulation results.</i> (Identifier: MEM_DDR3_RTT_WR_ENUM)



Table 234. Group: Mem I/O / ODT Activation

Display Name	Description
Use Default ODT Assertion Tables	Enables the default ODT assertion pattern as determined from vendor guidelines. These settings are provided as a default only; <i>you should simulate your memory interface to determine the optimal ODT settings and assertion patterns.</i> (Identifier: MEM_DDR3_USE_DEFAULT_ODT)

5.1.4 Intel Stratix 10 EMIF IP DDR3 Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.

Table 235. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_DDR3_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_DDR3_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_DDR3_DEFAULT_IO)

Table 236. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR3_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR3_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_DDR3_USER_AC_SLEW_RATE_ENUM)



Table 237. Group: FPGA I/O / FPGA I/O Settings / Memory Clock

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR3_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR3_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_DDR3_USER_CK_SLEW_RATE_ENUM)

Table 238. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR3_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR3_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR3_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_DDR3_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_DDR3_USER_STARTING_VREFIN)

Table 239. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_DDR3_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_DDR3_USER_RZQ_IO_STD_ENUM)

5.1.5 Intel Stratix 10 EMIF IP DDR3 Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).



Table 240. Group: Mem Timing / Parameters dependent on Speed Bin

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_DDR3_SPEEDBIN_ENUM)
tIS (base)	tIS (base) refers to the setup time for the Address/Command/Control (A) bus to the rising edge of CK. (Identifier: MEM_DDR3_TIS_PS)
tIS (base) AC level	tIS (base) AC level refers to the voltage level which the address/command signal must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_DDR3_TIS_AC_MV)
tIH (base)	tIH (base) refers to the hold time for the Address/Command (A) bus after the rising edge of CK. Depending on what AC level the user has chosen for a design, the hold margin can vary (this variance will be automatically determined when the user chooses the " tIH (base) AC level "). (Identifier: MEM_DDR3_TIH_PS)
tIH (base) DC level	tIH (base) DC level refers to the voltage level which the address/command signal must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_DDR3_TIH_DC_MV)
tDS (base)	tDS(base) refers to the setup time for the Data(DQ) bus before the rising edge of the DQS strobe. (Identifier: MEM_DDR3_TDS_PS)
tDS (base) AC level	tDS (base) AC level refers to the voltage level which the data bus must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_DDR3_TDS_AC_MV)
tDH (base)	tDH (base) refers to the hold time for the Data (DQ) bus after the rising edge of CK. (Identifier: MEM_DDR3_TDH_PS)
tDH (base) DC level	tDH (base) DC level refers to the voltage level which the data bus must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_DDR3_TDH_DC_MV)
tDQSQ	tDQSQ describes the latest valid transition of the associated DQ pins for a READ . tDQSQ specifically refers to the DQS, DQS# to DQ skew. It is the length of time between the DQS, DQS# crossing to the last valid transition of the slowest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_DDR3_TDQSQ_PS)
tQH	tQH specifies the output hold time for the DQ in relation to DQS, DQS# . It is the length of time between the DQS, DQS# crossing to the earliest invalid transition of the fastest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_DDR3_TQH_CYC)
tDQSCK	tDQSCK describes the skew between the memory clock (CK) and the input data strobes (DQS) used for reads . It is the time between the rising data strobe edge (DQS, DQS#) relative to the rising CK edge. (Identifier: MEM_DDR3_TDQSCK_PS)
tDQSS	tDQSS describes the skew between the memory clock (CK) and the output data strobes used for writes . It is the time between the rising data strobe edge (DQS, DQS#) relative to the rising CK edge. (Identifier: MEM_DDR3_TDQSS_CYC)
tQSH	tQSH refers to the differential High Pulse Width, which is measured as a percentage of tCK. It is the time during which the DQS is high for a read . (Identifier: MEM_DDR3_TQSH_CYC)
continued...	



Display Name	Description
tDSH	tDSH specifies the write DQS hold time . This is the time difference between the rising CK edge and the falling edge of DQS, measured as a percentage of tCK. (Identifier: MEM_DDR3_TDSH_CYC)
tWLS	tWLS describes the write leveling setup time . It is measured from the rising edge of CK to the rising edge of DQS. (Identifier: MEM_DDR3_TWLS_PS)
tWLH	tWLH describes the write leveling hold time . It is measured from the rising edge of DQS to the rising edge of CK (Identifier: MEM_DDR3_TWLH_PS)
tDSS	tDSS describes the time between the falling edge of DQS to the rising edge of the next CK transition . (Identifier: MEM_DDR3_TDSS_CYC)
tINIT	tINIT describes the time duration of the memory initialization after a device power-up . After RESET_n is de-asserted, wait for another 500us until CKE becomes active. During this time, the DRAM starts internal initialization; this happens independently of external clocks. (Identifier: MEM_DDR3_TINIT_US)
tMRD	The mode register set command cycle time, tMRD is the minimum time period required between two MRS commands . (Identifier: MEM_DDR3_TMRD_CK_CYC)
tRAS	tRAS describes the activate to precharge duration . A row cannot be deactivated until the tRAS time has been met. Therefore tRAS determines how long the memory has to wait after a activate command before a precharge command can be issued to close the row. (Identifier: MEM_DDR3_TRAS_NS)
tRCD	tRCD, row command delay , describes the active to read/write time . It is the amount of delay between the activation of a row through the RAS command and the access to the data through the CAS command. (Identifier: MEM_DDR3_TRCD_NS)
tRP	tRP refers to the Precharge (PRE) command period . It describes how long it takes for the memory to disable access to a row by precharging and before it is ready to activate a different row. (Identifier: MEM_DDR3_TRP_NS)
tWR	tWR refers to the Write Recovery time . It specifies the amount of clock cycles needed to complete a write before a precharge command can be issued. (Identifier: MEM_DDR3_TWR_NS)

Table 241. Group: Mem Timing / Parameters dependent on Speed Bin, Operating Frequency, and Page Size

Display Name	Description
tRRD	tRRD refers to the Row Active to Row Active Delay . It is the minimum time interval (measured in memory clock cycles) between two activate commands to rows in different banks in the same rank (Identifier: MEM_DDR3_TRRD_CYC)
tFAW	tFAW refers to the four activate window time . It describes the period of time during which only four banks can be active. (Identifier: MEM_DDR3_TFAW_NS)
tWTR	tWTR or Write Timing Parameter describes the delay from start of internal write transaction to internal read command, for accesses to the same bank . The delay is measured from the first rising memory clock edge after the last write data is received to the rising memory clock edge when a read command is received. (Identifier: MEM_DDR3_TWTR_CYC)
tRTP	tRTP refers to the internal READ Command to PRECHARGE Command delay . It is the number of memory clock cycles that is needed between a read command and a precharge command to the same rank. (Identifier: MEM_DDR3_TRTP_CYC)



Table 242. Group: Mem Timing / Parameters dependent on Density and Temperature

Display Name	Description
tRFC	tRFC refers to the Refresh Cycle Time . It is the amount of delay after a refresh command before an activate command can be accepted by the memory. This parameter is dependent on the memory density and is necessary for proper hardware functionality. (Identifier: MEM_DDR3_TRFC_NS)
tREFI	tREFI refers to the average periodic refresh interval . It is the maximum amount of time the memory can tolerate in between each refresh command (Identifier: MEM_DDR3_TREFI_US)

5.1.6 Intel Stratix 10 EMIF IP DDR3 Parameters: Board

Table 243. Group: Board / Intersymbol Interference/Crosstalk

Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation using IBIS models and Hyperlynx*, and manually enter values based on your simulation results, instead of using the default values.</i> (Identifier: BOARD_DDR3_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to ISI and crosstalk effects. The number to be entered is the total loss of margin on both the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR3_USER_AC_ISI_NS)
Read DQS/DQS# ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQS/DQS# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR3_USER_RCLK_ISI_NS)
Read DQ ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQ signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR3_USER_RDATA_ISI_NS)
Write DQS/DQS# ISI/crosstalk	The reduction of the write data window due to ISI and crosstalk effects on the DQS/DQS# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR3_USER_WCLK_ISI_NS)
Write DQ ISI/crosstalk	The reduction of the write data window due to ISI and crosstalk effects on the DQ signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR3_USER_WDATA_ISI_NS)




Table 244. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (DQS group)	Enable this parameter if you are compensating for package skew on the DQ, DQS, and DM buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_DDR3_IS_SKEW_WITHIN_DQS_DESKEWED)
Maximum board skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. This value affects the read capture and write margins. (Identifier: BOARD_DDR3_BRD_SKEW_WITHIN_DQS_NS)
Maximum system skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_DDR3_PKG_BRD_SKEW_WITHIN_DQS_NS)
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_DDR3_IS_SKEW_WITHIN_AC_DESKEWED)
Maximum board skew within address/command bus	The largest skew between the address and command signals. Enter the board skew only; package skew is calculated automatically, based on the memory interface configuration, and added to this value. (Identifier: BOARD_DDR3_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	Maximum system skew within address/command bus refers to the largest skew between the address and command signals. (Identifier: BOARD_DDR3_PKG_BRD_SKEW_WITHIN_AC_NS)
Average delay difference between DQS and CK	The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS trace delay minus the CK trace delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. (Identifier: BOARD_DDR3_DQS_TO_CK_SKEW_NS)
Maximum delay difference between DIMMs/devices	The largest propagation delay on DQ signals between ranks (<i>applicable only when there is more than one rank</i>). For example: when you configure two ranks using one DIMM there is a short distance between the ranks for the same DQ pin; when you implement two ranks using two DIMMs the distance is larger. (Identifier: BOARD_DDR3_SKEW_BETWEEN_DIMMS_NS)
Maximum skew between DQS groups	The largest skew between DQS signals. (Identifier: BOARD_DDR3_SKEW_BETWEEN_DQS_NS)
Average delay difference between address/command and CK	The average delay difference between the address/command signals and the CK signal, calculated by averaging the longest and smallest address/command signal trace delay minus the maximum CK trace delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. (Identifier: BOARD_DDR3_AC_TO_CK_SKEW_NS)
Maximum CK delay to DIMM/device	The delay of the longest CK trace from the FPGA to any DIMM/device. (Identifier: BOARD_DDR3_MAX_CK_DELAY_NS)
Maximum DQS delay to DIMM/device	The delay of the longest DQS trace from the FPGA to any DIMM/device (Identifier: BOARD_DDR3_MAX_DQS_DELAY_NS)

5.1.7 Intel Stratix 10 EMIF IP DDR3 Parameters: Controller



Table 245. Group: Controller / Avalon Interface


Display Name	Description
Avalon Interface	Selects the Avalon Interface through which the controller interacts with user logic (Identifier: CTRL_DDR3_AVL_PROTOCOL_ENUM)

Table 246. Group: Controller / Low Power Mode

Display Name	Description
Enable Self-Refresh Control	Select this option to enable the self-refresh control on the controller top level. The control signal allows you to place the memory device into self-refresh mode, on a per chip-select basis. (Identifier: CTRL_DDR3_SELF_REFRESH_EN)
Enable Auto Power-Down	Enable this parameter to have the controller automatically place the memory device into power-down mode after a specified number of idle controller clock cycles. The idle wait time is configurable. All ranks must be idle to enter auto power-down. (Identifier: CTRL_DDR3_AUTO_POWER_DOWN_EN)
Auto Power-Down Cycles	Specifies the number of idle controller cycles after which the memory device is placed into power-down mode. You can configure the idle waiting time. The supported range for number of cycles is from 1 to 65534. (Identifier: CTRL_DDR3_AUTO_POWER_DOWN_CYCS)

Table 247. Group: Controller / Efficiency

Display Name	Description
Enable User Refresh Control	When enabled, user logic has complete control and is responsible for issuing adequate refresh commands to the memory devices, via the MMR interface. This feature provides increased control over worst-case read latency and enables you to issue refresh bursts during idle periods. (Identifier: CTRL_DDR3_USER_REFRESH_EN)
Enable Auto-Precharge Control	Select this parameter to enable the auto-precharge control on the controller top level. If you assert the auto-precharge control signal while requesting a read or write burst, you can specify whether the controller should close (auto-precharge) the currently open page at the end of the read or write burst, potentially making a future access to a different page of the same bank faster. (Identifier: CTRL_DDR3_AUTO_PRECHARGE_EN)
Address Ordering	Controls the mapping between Avalon addresses and memory device addresses. By changing the value of this parameter, you can change the mappings between the Avalon-MM address and the DRAM address. (Identifier: CTRL_DDR3_ADDR_ORDER_ENUM)
Enable Reordering	Enable this parameter to allow the controller to perform command and data reordering. Reordering can improve efficiency by reducing bus turnaround time and row/bank switching time. Data reordering allows the single-port memory controller to change the order of read and write commands to achieve highest efficiency. Command reordering allows the controller to issue bank management commands early based on incoming patterns, so that the desired row in memory is already open when the command reaches the memory interface. <i>For more information, refer to the Data Reordering topic in the EMIF Handbook.</i> (Identifier: CTRL_DDR3_REORDER_EN)
Starvation limit for each command	Specifies the number of commands that can be served before a waiting command is served. The controller employs a counter to ensure that all requests are served after a pre-defined interval -- this ensures that low priority requests are not ignored, when doing data reordering for
continued...	



Display Name	Description
	efficiency. The valid range for this parameter is from 1 to 63. <i>For more information, refer to the Starvation Control topic in the EMIF Handbook.</i> (Identifier: CTRL_DDR3_STARVE_LIMIT)
Enable Command Priority Control	Select this parameter to enable user-requested command priority control on the controller top level. This parameter instructs the controller to treat a read or write request as high-priority. The controller attempts to fill high-priority requests sooner, to reduce latency. Connect this interface to the conduit of your logic block that determines when the external memory interface IP treats the read or write request as a high-priority command. (Identifier: CTRL_DDR3_USER_PRIORITY_EN)

Table 248. Group: Controller / Configuration, Status and Error Handling

Display Name	Description
Enable Memory-Mapped Configuration and Status Register (MMR) Interface	Enable this parameter to change or read memory timing parameters, memory address size, mode register settings, controller status, and request sideband operations. (Identifier: CTRL_DDR3_MMR_EN)
Enable Error Detection and Correction Logic with ECC	Enables error-correction code (ECC) for single-bit error correction and double-bit error detection . Your memory interface must have a width of 16, 24, 40, or 72 bits to use ECC. <i>ECC is implemented as soft logic.</i> (Identifier: CTRL_DDR3_ECC_EN)
Enable Auto Error Correction to External Memory	Specifies that the controller automatically schedule and perform a write back to the external memory when a single-bit error is detected. Regardless of whether the option is enabled or disabled, the ECC feature always corrects single-bit errors before returning the read data to user logic. (Identifier: CTRL_DDR3_ECC_AUTO_CORRECTION_EN)

Table 249. Group: Controller / Data Bus Turnaround Time

Display Name	Description
Additional read-to-write turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read to a write within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_RD_TO_WR_SAME_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write to a read within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_WR_TO_RD_SAME_CHIP_DELTA_CYCS)
Additional read-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a read of another logical rank . This can resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_RD_TO_RD_DIFF_CHIP_DELTA_CYCS)
Additional read-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a write of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_RD_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a write of another logical rank . This can help resolve bus contention problems
continued...	



Display Name	Description
	specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_WR_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a read of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR3_WR_TO_RD_DIFF_CHIP_DELTA_CYCS)

5.1.8 Intel Stratix 10 EMIF IP DDR3 Parameters: Diagnostics

Table 250. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process. Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero. <i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations.</i> Abstract PHY is supported with skip calibration. (Identifier: DIAG_DDR3_SIM_CAL_MODE_ENUM)
Abstract phy for fast simulation	Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration . (Identifier: DIAG_DDR3_ABSTRACT_PHY)

Table 251. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic. If you set this parameter to "Disabled" , no debug features are enabled. If you set this parameter to "Export" , an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP core to it, or connect it to the cal_debug_out interface of another EMIF core. If you select "Add EMIF Debug Interface" , an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit. <i>Only one EMIF debug interface should be instantiated per I/O column.</i> You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first. (Identifier: DIAG_DDR3_EXPORT_SEQ_AVALON_SLAVE)
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies that the IP export an Avalon-MM master interface (cal_debug_out) which can connect to the cal_debug interface of other EMIF cores residing in the same I/O column. This parameter applies only <i>continued...</i>



Display Name	Description
	if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to the <i>Debugging Multiple EMIFs</i> wiki page for more information about debugging multiple EMIFs. (Identifier: DIAG_DDR3_EXPORT_SEQ_AVALON_MASTER)
Interface ID	Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled , the interface ID is unused. (Identifier: DIAG_DDR3_INTERFACE_ID)
Use Soft NIOS Processor for On-Chip Debug	Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port . <i>Only one interface in a column can activate this option.</i> (Identifier: DIAG_SOFT_NIOS_MODE)

Table 252. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the "Core clocks sharing" parameter in the "General" tab to "Master" or "Slave" . (Identifier: DIAG_DDR3_EX_DESIGN_NUM_OF_SLAVES)
Enable In-System-Sources-and-Probes	Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i> . This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_DDR3_EX_DESIGN_ISSP_EN)

Table 253. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_DDR3_USE_TG_AVL_2)
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_DDR3_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the <code>altera_emif_avl_tg_2_tb.sv</code> file. (Identifier: DIAG_DDR3_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_DDR3_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_DDR3_BYPASS_STRESS_STAGE)
Export Traffic Generator 2.0 configuration interface	Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)



Table 254. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_DDR3_EFFICIENCY_MONITOR)

Table 255. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

5.1.9 Intel Stratix 10 EMIF IP DDR3 Parameters: Example Designs

Table 256. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the <i>creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization.</i> After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_DDR3_SEL_DESIGN)

Table 257. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created.</i> Instead, the output directory will contain the ed_sim.qsys file which holds Qsys details of the simulation example design, and a make_sim_design.tcl file with other corresponding tcl files. You can run make_sim_design.tcl from a command line to generate the simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_DDR3_GEN_SIM)
Synthesis	Specifies that the 'Generate Example Design' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created.</i> Instead, the output directory will contain the ed_synth.qsys file which holds Qsys details of the synthesis example design, and a make_qii_design.tcl script with other corresponding tcl files. You can run make_qii_design.tcl from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_DDR3_GEN_SYNTH)



Table 258. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_DDR3_HDL_FORMAT)

Table 259. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select ' none ' from the ' Select board ' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_DDR3_TARGET_DEV_KIT)

5.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

5.2.1 Equations for DDR3 Board Skew Parameters

Table 260. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum CK delay to DIMM/device	The delay of the longest CK trace from the FPGA to any DIMM/device. $\max_r \left[\max_n (CK_{n,r} PathDelay) \right]$ <p>Where n is the number of memory clock and r is the number rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 pairs of memory clocks in each rank DIMM, the maximum CK delay is expressed by the following equation:</p> $\max(CK_1 PathDelayrank1, CK_2 PathDelayrank1, CK_1 PathDelayrank2, CK_2 PathDelayrank2)$
Maximum DQS delay to DIMM/device	The delay of the longest DQS trace from the FPGA to the DIMM/device. $\max_r \left[\max_n (DQS_{n,r} PathDelay) \right]$ <p>Where n is the number of DQS and r is the number of rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 DQS in each rank DIMM, the maximum DQS delay is expressed by the following equation:</p> $\max(DQS_1 PathDelayrank1, DQS_2 PathDelayrank1, DQS_1 PathDelayrank2, DQS_2 PathDelayrank2)$
Average delay difference between DQS and CK	The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS delay minus the CK delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the DQS signals for appropriate setup and hold margins.

continued...



Parameter	Description/Equation
	$\max_r \left[\max_{n, m} \left\{ \frac{(DQS_{m_r} Delay - CK_{n_r} Delay)}{2} + \min_r \left[\frac{\min}{n, m} \right] \left\{ (DQS_{m_r} Delay - CK_{n_r} Delay) \right\} \right\} \right]$ <p>Where n is the number of memory clock, m is the number of DQS, and r is the number of rank of DIMM/device.</p> <p>When using discrete components, the calculation differs slightly. Find the minimum and maximum values for (DQS-CK) over all groups and then divide by 2. Calculate the (DQS-CK) for each DQS group, by using the appropriate CLK for that group.</p> <p>For example, in a configuration with 5 x16 components, with each component having two DQS groups: To find the minimum and maximum, calculate the minimum and maximum of (DQS0 – CK0, DQS1 – CK0, DQS2 –CK1, DQS3 – CK1, and so forth) and then divide the result by 2.</p>
Maximum Board skew within DQS group	<p>The largest skew between all DQ and DM pins in a DQS group. Enter your board skew only. Package skew is calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins.</p> $\left[\max_{groups} \left[\max_g DQ_g - \min_g DQ_g \right] \right]$
Maximum skew between DQS groups	<p>The largest skew between DQS signals in different DQS groups.</p> $\left[\max_{groups} \left[DQS_g \right] \right] - \left[\min_{groups} \left[DQS_g \right] \right]$
Maximum system skew within address/command bus	<p>(MaxAC – MinAC)</p> <p>The largest skew between the address and command signals. Enter combined board and package skew. In the case of a component, find the maximum address/command and minimum address/command values across all component address signals.</p>
Average delay difference between address/command and CK	<p>A value equal to the average of the longest and smallest address/command signal delays, minus the delay of the CK signal. The value can be positive or negative.</p> <p>The average delay difference between the address/command and CK is expressed by the following equation:</p> $\frac{\sum_{n=1}^n \left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - CK_n PathDelay}{n}$ <p>where n is the number of memory clocks.</p>
Maximum delay difference between DIMMs/devices	<p>The largest propagation delay on DQ signals between ranks. For example, in a two-rank configuration where you place DIMMs in different slots there is also a propagation delay for DQ signals going to and coming back from the furthest DIMM compared to the nearest DIMM. This parameter is applicable only when there is more than one rank. Max_r { max_{n,m} [(DQ_{n_r} path delay- DQ_{n_r+1} path delay), (DQ_{m_r} path delay- DQ_{m_r+1} path delay)] }</p> <p>Where n is the number of DQ, m is the number of DQS and r is number of rank of DIMM/device .</p>

5.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.

5.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.


Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

5.3.1.1 Estimating Pin Requirements

You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:

1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.



You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)

5.3.1.2 DIMM Options

Unbuffered DIMMs (UDIMMs) require one set of chip-select ($CS\#$), on-die termination (ODT), clock-enable (CKE), and clock pair (CK/CK_n) for every physical rank on the DIMM. Registered DIMMs use only one pair of clocks. DDR3 registered DIMMs require a minimum of two chip-select signals, while DDR4 requires only one.

Compared to the unbuffered DIMMs (UDIMM), registered and load-reduced DIMMs (RDIMMs and LRDIMMs, respectively) use at least two chip-select signals $CS\#[1:0]$ in DDR3 and DDR4. Both RDIMMs and LRDIMMs require an additional parity signal for address, $RAS\#$, $CAS\#$, and $WE\#$ signals. A parity error signal is asserted by the module whenever a parity error is detected.

LRDIMMs expand on the operation of RDIMMs by buffering the DQ/DQS bus. Only one electrical load is presented to the controller regardless of the number of ranks, therefore only one clock enable (CKE) and ODT signal are required for LRDIMMs, regardless of the number of physical ranks. Because the number of physical ranks may exceed the number of physical chip-select signals, DDR3 LRDIMMs provide a feature known as rank multiplication, which aggregates two or four physical ranks into one larger logical rank. Refer to LRDIMM buffer documentation for details on rank multiplication.

The following table shows UDIMM and RDIMM pin options for DDR3.

Table 261. UDIMM and RDIMM Pin Options for DDR3

Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)
Data	72 bit $DQ[71:0] = \{CB[7:0], DQ[63:0]\}$	72 bit $DQ[71:0] = \{CB[7:0], DQ[63:0]\}$	72 bit $DQ[71:0] = \{CB[7:0], DQ[63:0]\}$	72 bit $DQ[71:0] = \{CB[7:0], DQ[63:0]\}$
Data Mask	$DM[8:0]$	$DM[8:0]$	$DM[8:0]$	$DM[8:0]$
Data Strobe	$DQS[8:0]$ and $DQS\#[8:0]$	$DQS[8:0]$ and $DQS\#[8:0]$	$DQS[8:0]$ and $DQS\#[8:0]$	$DQS[8:0]$ and $DQS\#[8:0]$
Address	$BA[2:0]$, $A[15:0]$ – 2 GB: $A[13:0]$ 4 GB: $A[14:0]$ 8 GB: $A[15:0]$	$BA[2:0]$, $A[15:0]$ – 2 GB: $A[13:0]$ 4 GB: $A[14:0]$ 8 GB: $A[15:0]$	$BA[2:0]$, $A[15:0]$ – 2 GB: $A[13:0]$ 4 GB: $A[14:0]$ 8 GB: $A[15:0]$	$BA[2:0]$, $A[15:0]$ – 2 GB: $A[13:0]$ 4 GB: $A[14:0]$ 8 GB: $A[15:0]$
Clock	$CK0/CK0\#$	$CK0/CK0\#$, $CK1/CK1\#$	$CK0/CK0\#$	$CK0/CK0\#$

continued...



Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)
Command	ODT, CS#, CKE, RAS#, CAS#, WE#	ODT[1:0], CS#[1:0], CKE[1:0], RAS#, CAS#, WE#	ODT, CS#[1:0], CKE, RAS#, CAS#, WE# ²	ODT[1:0], CS#[1:0], CKE[1:0], RAS#, CAS#, WE#
Parity	—	—	PAR, ALERT	PAR, ALERT
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#

The following table shows LRDIMM pin options for DDR3.

Table 262. LRDIMM Pin Options for DDR3

Pins	LRDIMM Pins (x4, 2R)	LRDIMM Pins (x4, 4R, RMF=1) ³	LRDIMM Pins (x4, 4R, RMF=2)	LRDIMM Pins (x4, 8R, RMF=2)	LRDIMM Pins (x4, 8R, RMF=4)	LRDIMM Pins (x8, 4R, RMF=1) ³	LRDIMM Pins (x8, 4R, RMF=2)
Data	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}
Data Mask	—	—	—	—	—	DM[8:0]	DM[8:0]
Data Strobe	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[8:0] and DQS#[8:0]	DQS[8:0] and DQS#[8:0]
Address	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]	BA[2:0], A[17:0] -16GB:A[15:0] 32GB:A[16:0] 64GB:A[17:0]	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]
Clock	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#
Command	ODT, CS[1:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[2:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[2:0]#, CKE, RAS#, CAS#, WE#
Parity	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#

Notes to Table:

- DM pins are not used for LRDIMMs that are constructed using x4 components.
- S#[2] is treated as A[16] (whose corresponding pins are labeled as CS#[2] or RM[0]) and S#[3] is treated as A[17] (whose corresponding pins are labeled as CS#[3] or RM[1]) for certain rank multiplication configuration.
- R = rank, RMF = rank multiplication factor.

The following table shows UDIMM, RDIMM, and LRDIMM pin options for DDR4.

**Table 263. UDIMM, RDIMM, and LRDIMM Pin Options for DDR4**

Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)	LRDIMM Pins (Dual Rank)	LRDIMM Pins (Quad Rank)
Data	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}
Data Mask	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	—	—
Data Strobe	x8: DQS[8:0] and DQS#[8:0]	x8: DQS[8:0] and DQS#[8:0]	x8: DQS[8:0] and DQS#[8:0] x4: DQS[17:0] and DQS#[17:0]	x8: DQS[8:0] and DQS#[8:0] x4: DQS[17:0] and DQS#[17:0]	x4: DQS[17:0] and DQS#[17:0]	x4: DQS[17:0] and DQS#[17:0]
Address	BA[1:0], BG[1:0], A[16:0] - 4GB: A[14:0] 8GB: A[15:0] 16GB: A[16:0] ⁽²⁾	BA[1:0], BG[1:0], A[16:0] - 8GB: A[14:0] 16GB: A[15:0] 32GB: A[16:0] ⁽²⁾	BA[1:0], BG[1:0], x8: A[16:0] - 4GB: A[14:0] 8GB: A[15:0] 16GB: A[16:0] ⁽²⁾ 32GB: A[17:0] ⁽³⁾	BA[1:0], BG[1:0], x8: A[16:0] x4: A[17:0] - 8GB: A[14:0] 16GB: A[15:0] 32GB: A[16:0] ⁽²⁾ 64GB: A[17:0] ⁽³⁾	BA[1:0], BG[1:0], A[17:0] - 16GB: A[15:0] 32GB: A[16:0] ⁽²⁾ 64GB: A[17:0] ⁽³⁾	BA[1:0], BG[1:0], A[17:0] - 32GB: A[15:0] 64GB: A[16:0] ⁽²⁾ 128GB: A[17:0] ⁽³⁾
Clock	CK0/CK0#	CK0/CK0#, CK1/CK1#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#
Command	ODT, CS#, CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT[1:0], CS#[1:0], CKE[1:0], ACT#, RAS#/ A16, CAS#/ A15, WE#/A14	ODT, CS#, CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT[1:0], CS#[1:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT, CS#[1:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT, CS#[3:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14
Parity	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#

Notes to Table:

1. DM/DBI pins are available only for DIMMs constructed using x8 or greater components.
2. This density requires 4Gb x4 or 2Gb x8 DRAM components.
3. This density requires 8Gb x4 DRAM components.
4. This table assumes a single slot configuration. The Intel Stratix 10 memory controller can support up to 4 ranks per channel. A single slot interface may have up to 4 ranks, and a dual slot interface may have up to 2 ranks per slot. In either case, the total number of ranks, calculated as the number of slots multiplied by the number of ranks per slot, must be less than or equal to 4.

5.3.1.3 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.



Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

5.3.2 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.

5.3.2.1 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

5.3.2.2 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:



- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.

5.3.3 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.

The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the `Bank Number` and `Index within I/O Bank` values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the `Bank Number` value identifies the I/O column, while the letter represents the I/O bank.
- The `Index within I/O Bank` value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its `Index within I/O Bank` number (if it is an even number), or by subtracting one from its `Index within I/O Bank` number (if it is an odd number).

For example, a physical pin with a `Bank Number` of 2M and `Index within I/O Bank` of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an `Index within I/O Bank` of 23 and `Bank Number` of 2M.

5.3.3.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.



If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version number/<synth/sim>/<variation_name>_altera_emif_arch_nd_version number_<unique ID>_readme.txt` file, w+ted with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (.qip), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLDRAMx , QDRx and LPDDR3 are not supported with HPS.

Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the `<variation_name>/altera_emif_arch_nd_<version>/<synth/sim>/<variation_name>_altera_emif_arch_nd_<version>_<unique ID>_readme.txt` file after you have generated your IP.



7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.

Note: a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.

b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).
10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:



- There must be an even number of x4 groups in an external memory interface.
- DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group $X+1$ must be in the same I/O lane, where X is an even number.
- When placing DQ pins in x4 mode, it is important to stay within a nibble when swapping pin locations. In other words, you may swap DQ pins within a given DQS group, but you may not swap pins across DQS groups. The following table illustrates an example, where DATA_A and DATA_B are swap groups, meaning that any pin in that index can move within that range of pins.

Index Within Lane	DQS x4 Locations
11	DATA_B[3:0]
10	DATA_B[3:0]
9	DQS_Bn
8	DQS_Bp
7	DATA_B[3:0]
6	DATA_B[3:0]
5	DQS_An
4	DQS_Ap
3	DATA_A[3:0]
2	DATA_A[3:0]
1	DATA_A[3:0]
0	DATA_A[3:0]

11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Extreme, and RLD RAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLD RAM 3, x36 device, $DQ[8:0]$ and $DQ[26:18]$ are referenced to $DK0/DK0\#$, and $DQ[17:9]$ and $DQ[35:27]$ are referenced to $DK1/DK1\#$.

12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.

13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.



Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLDRAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (`global_reset_n`) for each individual interface all come from the same input pin or signal.

I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.



CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)

The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the R_{ZQ} pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.



5.3.3.2 Command and Address Signals

Command and address signals in SDRAM devices are clocked into the memory device using the CK or CK# signal. These pins operate at single data rate (SDR) using only one clock edge. The number of address pins depends on the SDRAM device capacity. The address pins are multiplexed, so two clock cycles are required to send the row, column, and bank address.

For DDR3, the CS#, RAS#, CAS#, WE#, CKE, and ODT pins are SDRAM command and control pins. For DDR3 SDRAM, certain topologies such as RDIMM and LRDIMM include RESET#, PAR (1.5V LVC MOS I/O standard), and ALERT# (SSTL-15 I/O standard).

Although DDR4 operates in fundamentally the same way as other SDRAM, there are no longer dedicated pins for RAS#, CAS#, and WE#, as those are now shared with higher-order address pins. DDR4 still has CS#, CKE, ODT, and RESET# pins, similar to DDR3. DDR4 introduces some additional pins, including the ACT# (activate) pin and BG (bank group) pins. Depending on the memory format and the functions enabled, the following pins might also exist in DDR4: PAR (address command parity) pin and the ALERT# pin.

5.3.3.3 Clock Signals

DDR3 and DDR4 SDRAM devices use CK and CK# signals to clock the address and command signals into the memory. Furthermore, the memory uses these clock signals to generate the DQS signal during a read through the DLL inside the memory. The SDRAM data sheet specifies the following timings:

- t_{DQSCk} is the skew between the CK or CK# signals and the SDRAM-generated DQS signal
- t_{DSH} is the DQS falling edge from CK rising edge hold time
- t_{DSS} is the DQS falling edge from CK rising edge setup time
- t_{DQSS} is the positive DQS latching edge to CK rising edge

SDRAM have a write requirement (t_{DQSS}) that states the positive edge of the DQS signal on writes must be within $\pm 25\%$ ($\pm 90^\circ$) of the positive edge of the SDRAM clock input. Therefore, you should generate the CK and CK# signals using the DDR registers in the IOE to match with the DQS signal and reduce any variations across process, voltage, and temperature. The positive edge of the SDRAM clock, CK, is aligned with the DQS write to satisfy t_{DQSS} .

DDR3 SDRAM can use a daisy-chained control address command (CAC) topology, in which the memory clock must arrive at each chip at a different time. To compensate for the flight-time skew between devices when using the CAC topology, you should employ write leveling.

5.3.3.4 Data, Data Strobes, DM/DBI, and Optional ECC Signals

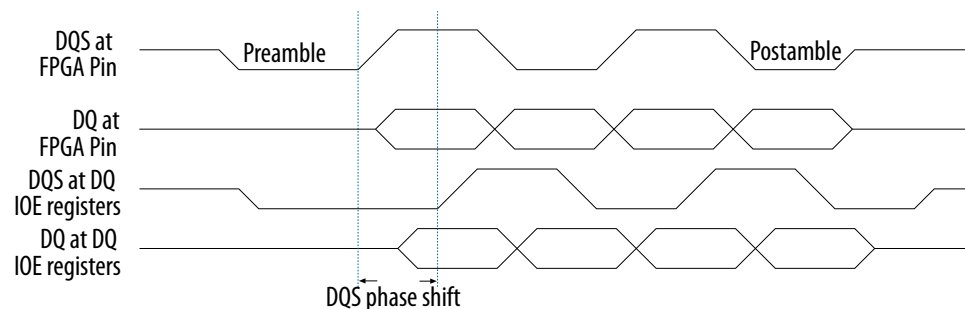
DDR3 and DDR4 SDRAM use bidirectional differential data strobes. Differential DQS operation enables improved system timing due to reduced crosstalk and less simultaneous switching noise on the strobe output drivers. The DQ pins are also bidirectional.

DQ pins in DDR3 and DDR4 SDRAM interfaces can operate in either $\times 4$ or $\times 8$ mode DQS groups, depending on your chosen memory device or DIMM, regardless of interface width. The $\times 4$ and $\times 8$ configurations use one pair of bidirectional data strobe

signals, DQS and DQSn, to capture input data. However, two pairs of data strobes, UDQS and UDQS# (upper byte) and LDQS and LDQS# (lower byte), are required by the $\times 16$ configuration devices. A group of DQ pins must remain associated with its respective DQS and DQSn pins.


The DQ signals are edge-aligned with the DQS signal during a read from the memory and are center-aligned with the DQS signal during a write to the memory. The memory controller shifts the DQ signals by -90 degrees during a write operation to center align the DQ and DQS signals. The PHY IP delays the DQS signal during a read, so that the DQ and DQS signals are center aligned at the capture register. Intel devices use a phase-locked loop (PLL) to center-align the DQS signal with respect to the DQ signals during writes and Intel devices use dedicated DQS phase-shift circuitry to shift the incoming DQS signal during reads. The following figure shows an example where the DQS signal is shifted by 90 degrees for a read from the DDR3 SDRAM.

Figure 41. Edge-aligned DQ and DQS Relationship During a DDR3 SDRAM Read in Burst-of-Four Mode



The following figure shows an example of the relationship between the data and data strobe during a burst-of-four write.

Figure 42. DQ and DQS Relationship During a DDR3 SDRAM Write in Burst-of-Four Mode



The memory device's setup (t_{DS}) and hold times (t_{DH}) for the DQ and DM pins during writes are relative to the edges of DQS write signals and not the CK or CK# clock. Setup and hold requirements are not necessarily balanced in DDR3 SDRAM.

The DQS signal is generated on the positive edge of the system clock to meet the t_{DQSS} requirement. DQ and DM signals use a clock shifted -90 degrees from the system clock, so that the DQS edges are centered on the DQ or DM signals when they arrive at the DDR3 SDRAM. The DQS, DQ, and DM board trace lengths need to be tightly matched (within 20 ps).

The SDRAM uses the DM pins during a write operation. Driving the DM pins low shows that the write is valid. The memory masks the DQ signals if the DM pins are driven high. To generate the DM signal, Intel recommends that you use the spare DQ pin within the same DQS group as the respective data, to minimize skew.



The DM signal's timing requirements at the SDRAM input are identical to those for DQ data. The DDR registers, clocked by the -90 degree shifted clock, create the DM signals.

DDR4 supports DM similarly to other SDRAM, except that in DDR4 DM is active LOW and bidirectional, because it supports Data Bus Inversion (DBI) through the same pin. DM is multiplexed with DBI by a Mode Register setting whereby only one function can be enabled at a time. DBI is an input/output identifying whether to store/output the true or inverted data. When enabled, if DBI is LOW, during a write operation the data is inverted and stored inside the DDR4 SDRAM; during a read operation, the data is inverted and output. The data is not inverted if DBI is HIGH. For Intel Stratix 10 interfaces, the DM (for DDR3) pins in each DQS group must be paired with a DQ pin for proper operation. DM/DBI (for DDR4) do not need to be paired with a DQ pin.

Some SDRAM modules support error correction coding (ECC) to allow the controller to detect and automatically correct error in data transmission. The 72-bit SDRAM modules contain eight extra data pins in addition to 64 data pins. The eight extra ECC pins should be connected to a single DQS or DQ group on the FPGA.

5.3.3.5 Resource Sharing Guidelines (Multiple Interfaces)

In Intel Cyclone 10Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.

PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.

Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `p11_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note:

You can place the `p11_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.

Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.

Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.

5.3.3.6 Ping-Pong PHY Implementation

The Ping Pong PHY feature instantiates two hard memory controllers—one for the primary interface and one for the secondary interface. The hard memory controller I/O bank of the primary interface is used for address and command and is always adjacent and above the hard memory controller I/O bank of the secondary interface. All four lanes of the primary hard memory controller I/O bank are used for address and command.

When you use Ping Pong PHY, the EMIF IP exposes two independent Avalon-MM interfaces to user logic; these interfaces correspond to the two hard memory controllers inside the interface. Each Avalon-MM interface has its own set of clock and reset signals. Refer to *Platform Designer Interfaces* for more information on the additional signals exposed by Ping Pong PHY interfaces.

For pin allocation information for Intel Stratix 10 devices, refer to *External Memory Interface Pin Information for Intel Stratix 10 Devices* on www.altera.com.



Additional Requirements for DDR3 and DDR4 Ping-Pong PHY Interfaces

If you are using Ping Pong PHY with a DDR3 or DDR4 external memory interface on an Intel Stratix 10 device, follow these guidelines:

- The address and command I/O bank must not contain any DQS group.
- I/O banks that are above the address and command I/O bank must contain only data pins of the primary interface—that is, the interface with the lower DQS group indices.
- The I/O bank immediately below the address and command I/O bank must contain at least one DQS group of the secondary interface—that is, the interface with the higher DQS group indices. This I/O bank can, but is not required to, contain DQS groups of the primary interface.
- I/O banks that are two or more banks below the address and command I/O bank must contain only data pins of the secondary interface.

Related Links

- [Pin-Out Files for Intel FPGA Devices](#)
- [Functional Description— Intel Stratix 10 EMIF IP](#)
- [External Memory Interface Pin Information for Intel Stratix 10 Devices](#)
- [Restrictions on I/O Bank Usage for Stratix 10 EMIF IP with HPS](#)

5.4 DDR3 Board Design Guidelines

The following topics provide guidelines for improving the signal integrity of your system and for successfully implementing a DDR3 SDRAM interface on your system.

The following areas are discussed:

- comparison of various types of termination schemes, and their effects on the signal quality on the receiver
- proper drive strength setting on the FPGA to optimize the signal integrity at the receiver
- effects of different loading types, such as components versus DIMM configuration, on signal quality


It is important to understand the trade-offs between different types of termination schemes, the effects of output drive strengths, and different loading types, so that you can swiftly navigate through the multiple combinations and choose the best possible settings for your designs.

The following key factors affect signal quality at the receiver:

- Leveling and dynamic ODT
- Proper use of termination
- Layout guidelines

As memory interface performance increases, board designers must pay closer attention to the quality of the signal seen at the receiver because poorly transmitted signals can dramatically reduce the overall data-valid margin at the receiver. The following figure shows the differences between an ideal and real signal seen by the receiver.

Figure 43. Ideal and Real Signal at the Receiver



Related Links

JEDEC.org

5.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices

The following topics describe considerations specific to DDR3 and DDR4 external memory interface protocols on Intel Stratix 10 devices.

5.4.1.1 Dynamic On-Chip Termination (OCT) in Intel Stratix 10 Devices

Depending upon the R_s (series) and R_t (parallel) OCT values that you want, you should choose appropriate values for the RZQ resistor and connect this resistor to the RZQ pin of the FPGA.

- Select a 240-ohm reference resistor to ground to implement R_s OCT values of 34-ohm, 40-ohm, 48-ohm, 60-ohm, and 80-ohm, and R_t OCT resistance values of 20-ohm, 30-ohm, 34-ohm, 40-ohm, 60-ohm, 80-ohm, 120-ohm and 240 ohm.
- Select a 100-ohm reference resistor to ground to implement R_s OCT values of 25-ohm and 50-ohm, and an R_t OCT resistance of 50-ohm.

Check the FPGA I/O tab of the parameter editor to determine the I/O standards and termination values supported for data, address and command, and memory clock signals.

Related Links

[Choosing Terminations on Intel Stratix 10 Devices](#) on page 182

5.4.1.2 Choosing Terminations on Intel Stratix 10 Devices

To determine optimal on-chip termination (OCT) and on-die termination (ODT) values for best signal integrity, you should simulate your memory interface in HyperLynx or a similar tool.

If the optimal OCT and ODT termination values as determined by simulation are not available in the list of available values in the parameter editor, select the closest available termination values for OCT and ODT.

For information about available ODT choices, refer to your memory vendor data sheet.



Related Links

[Dynamic On-Chip Termination \(OCT\) in Intel Stratix 10 Devices](#) on page 182

5.4.1.3 On-Chip Termination Recommendations for Intel Stratix 10 Devices

- Output mode (drive strength) for Address/Command/Clock and Data Signals: Depending upon the I/O standard that you have selected, you would have a range of selections expressed in terms of ohms or miliamps. A value of 34 to 40 ohms or 12 mA is a good starting point for output mode drive strength.
- Input mode (parallel termination) for Data and Data Strobe signals: A value of 40 or 60 ohms is a good starting point for FPGA side input termination.

5.4.2 Channel Signal Integrity Measurement

As external memory interface data rates increase, so does the importance of proper channel signal integrity measurement. By measuring the actual channel loss during the layout process and including that data in your parameterization, a realistic assessment of margins is achieved.

5.4.2.1 Importance of Accurate Channel Signal Integrity Information

Default values for channel loss (or eye reduction) can be used when calculating timing margins, however those default values may not accurately reflect the channel loss in your system. If the channel loss in your system is different than the default values, the calculated timing margins will vary accordingly.


If your actual channel loss is greater than the default channel loss, and if you rely on default values, the available timing margins for the entire system will be lower than the values calculated during compilation. By relying on default values that do not accurately reflect your system, you may be led to believe that you have good timing margin, while in reality, your design may require changes to achieve good channel signal integrity.

5.4.2.2 Understanding Channel Signal Integrity Measurement

To measure channel signal integrity you need to measure the channel loss for various signals. For a particular signal or signal trace, channel loss is defined as loss of the eye width at $\pm V_{IH}(\text{ac and dc}) \pm V_{IL}(\text{ac and dc})$. V_{IH}/V_{IL} above or below V_{REF} is used to align with various requirements of the timing model for memory interfaces.

The example below shows a reference eye diagram where the channel loss on the setup- or leading-side of the eye is equal to the channel loss on the hold- or lagging-side of the eye; however, it does not necessarily have to be that way. Because Intel's calibrating PHY will calibrate to the center of the read and write eye, the Board Settings tab has parameters for the total extra channel loss for Write DQ and Read DQ. For address and command signals which are not-calibrated, the Board Settings tab allows you to enter setup- and hold-side channel losses that are not equal, allowing the Intel Quartus Prime software to place the clock statically within the center of the address and command eye.

Figure 44. Equal Setup and Hold-side Losses



5.4.2.3 How to Enter Calculated Channel Signal Integrity Values

You should enter calculated channel loss values in the **Channel Signal Integrity** section of the **Board** (or **Board Timing**) tab of the parameter editor.

For Intel Stratix 10 external memory interfaces, the default channel loss displayed in the parameter editor is based on the selected configuration (different values for single rank versus dual rank), and on internal Intel reference boards. You should replace the default value with the value that you calculate.

5.4.2.4 Guidelines for Calculating DDR3 Channel Signal Integrity

Address and Command ISI and Crosstalk

Simulate the address/command and control signals and capture eye at the DRAM pins, using the memory clock as the trigger for the memory interface's address/command and control signals. Measure the setup and hold channel losses at the voltage thresholds mentioned in the memory vendor's data sheet.


Address and command channel loss = Measured loss on the setup side + measured loss on the hold side.

$$V_{REF} = V_{DD}/2 = 0.75 \text{ mV for DDR3}$$

You should select the V_{IH} and V_{IL} voltage levels appropriately for the DDR3L memory device that you are using. Check with your memory vendor for the correct voltage levels, as the levels may vary for different speed grades of device.

The following figure illustrates a DDR3 example where $V_{IH(AC)}/V_{IL(AC)}$ is +/- 150 mV and $V_{IH(DC)}/V_{IL(DC)}$ is +/- 100 mV.

Figure 45.




Write DQ ISI and Crosstalk

Simulate the write DQ signals and capture eye at the DRAM pins, using DQ Strobe (DQS) as a trigger for the DQ signals of the memory interface simulation. Measure the setup and hold channel losses at the V_{IH} and V_{IL} mentioned in the memory vendor's data sheet. The following figure illustrates a DDR3 example where $V_{IH(AC)}/V_{IL(AC)}$ is ± 150 mV and $V_{IH(DC)}/V_{IL(DC)}$ is ± 100 mV.

Write Channel Loss = Measured Loss on the Setup side + Measured Loss on the Hold side

$$V_{REF} = V_{DD}/2 = 0.75 \text{ mV for DDR3}$$

Figure 46.



Read DQ ISI and Crosstalk


Simulate read DQ signals and capture eye at the FPGA die. Do not measure at the pin, because you might see unwanted reflections that could create a false representation of the eye opening at the input buffer of the FPGA. Use DQ Strobe (DQS) as a trigger for the DQ signals of your memory interface simulation. Measure the eye opening at ± 70 mV (V_{IH}/V_{IL}) with respect to V_{REF} .

Read Channel Loss = (UI) - (Eye opening at ± 70 mV with respect to V_{REF})

UI = Unit interval. For example, if you are running your interface at 800 Mhz, the effective data is 1600 Mbps, giving a unit interval of $1/1600 = 625$ ps

$V_{REF} = VDD/2 = 0.75$ mV for DDR3

Figure 47.




Write/Read DQS ISI and Crosstalk

Simulate the Write/Read DQS and capture eye, and measure the uncertainty at V_{REF} .

$V_{REF} = VDD/2 = 0.75$ mV for DDR3

Figure 48.

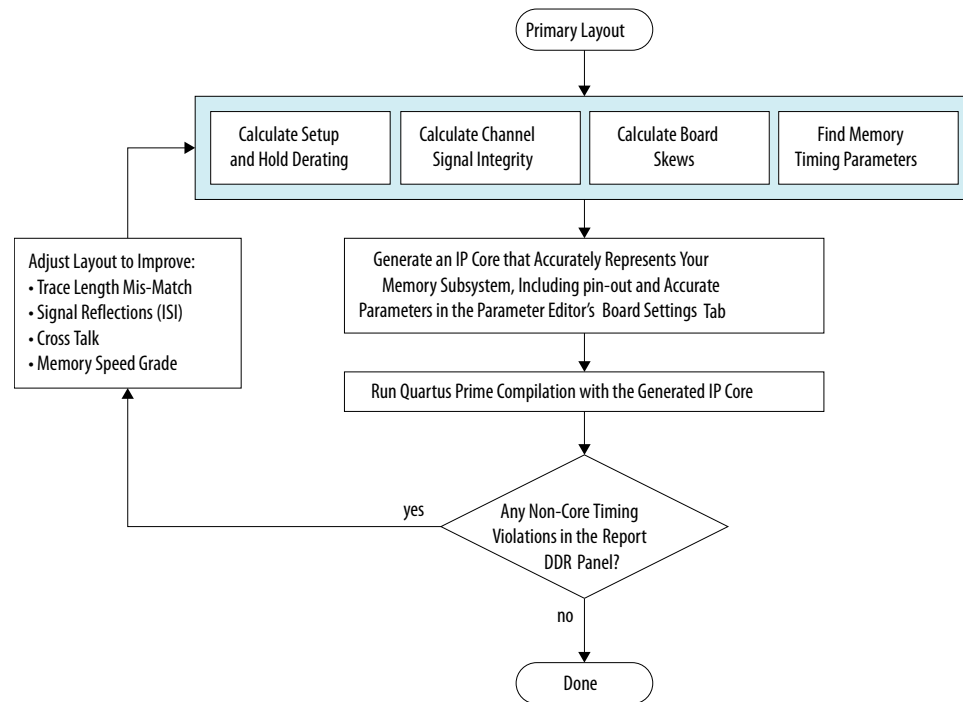


5.4.3 Layout Approach

For all practical purposes, you can regard the Timing Analyzer report on your memory interface as definitive for a given set of memory and board timing parameters.

You will find timing under **Report DDR** in the Timing Analyzer and on the **Timing Analysis** tab in the parameter editor.

The following flowchart illustrates the recommended process to follow during the board design phase, to determine timing margin and make iterative improvements to your design.



Board Skew

For information on calculating board skew parameters, refer to *Board Skew Equations*, in this chapter.

The Board Skew Parameter Tool is an interactive tool that can help you calculate board skew parameters if you know the absolute delay values for all the memory related traces.

Memory Timing Parameters

For information on the memory timing parameters to be entered into the parameter editor, refer to the datasheet for your external memory device.

Related Links

[Board Skew Parameter Tool](#)

5.4.4 Design Layout Guidelines

The general layout guidelines in the following topic apply to DDR3 and DDR4 SDRAM interfaces.

These guidelines will help you plan your board layout, but are not meant as strict rules that must be adhered to. Intel recommends that you perform your own board-level simulations to ensure that the layout you choose for your board allows you to achieve your desired performance.



For more information about how the memory manufacturers route these address and control signals on their DIMMs, refer to the Cadence PCB browser from the Cadence website, at www.cadence.com. The various JEDEC* example DIMM layouts are available from the JEDEC website, at www.jedec.org.

For assistance in calculating board skew parameters, refer to the board skew calculator tool, which is available at the Intel website.

Note:

1. The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristic of the interface. They do not include any margin for crosstalk.
2. To ensure reliable timing closure to and from the periphery of the device, signals to and from the periphery should be registered before any further logic is connected.

Intel recommends that you get accurate time base skew numbers for your design when you simulate the specific implementation.

Related Links

- [Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface \(>72 bits\)](#) on page 195
- JEDEC.org
- <https://www.cadence.com/>
- [Board Skew Parameter Tool](#)
- <https://www.mentor.com/>

5.4.4.1 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.

Table 264. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-Junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)


5.4.4.2 Layout Guidelines

The following table lists layout guidelines.

Unless otherwise specified, the guidelines in the following table apply to the following topologies:

- DIMM—UDIMM topology
- DIMM—RDIMM topology
- DIMM—LRDIMM topology
- Not all versions of the Intel Quartus Prime software support LRDIMM.
- Discrete components laid out in UDIMM topology
- Discrete components laid out in RDIMM topology

These guidelines are recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface.



For supported frequencies and topologies, refer to the *External Memory Interface Spec Estimator* <http://www.altera.com/technology/memory/estimator/mem-emif-index.html>.

For frequencies greater than 800 MHz, when you are calculating the delay associated with a trace, you must take the FPGA package delays into consideration.

Table 265. Layout Guidelines (1)

Parameter	Guidelines
Decoupling Parameter	<ul style="list-style-type: none"> Make VTT voltage decoupling close to the components and pull-up resistors. Connect decoupling caps between VTT and VDD using a 0.1μF cap for every other VTT pin. Use a 0.1 μF cap and 0.01 μF cap for every VDDQ pin.
Maximum Trace Length	<ul style="list-style-type: none"> Even though there are no hard requirements for minimum trace length, you need to simulate the trace to ensure the signal integrity. Shorter routes result in better timing. For DIMM topology only: <ul style="list-style-type: none"> Maximum trace length for all signals from FPGA to the first DIMM slot is 4.5 inches. Maximum trace length for all signals from DIMM slot to DIMM slot is 0.425 inches. For discrete components only: <ul style="list-style-type: none"> Maximum trace length for address, command, control, and clock from FPGA to the first component must not be more than 7 inches. Maximum trace length for DQ, DQS, DQS#, and DM from FPGA to the first component is 5 inches.
General Routing	<ul style="list-style-type: none"> Route over appropriate VCC and GND planes. Keep signal routing layers close to GND and power planes.
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always make sure that the signals related to memory interface are routed between appropriate GND or power layers. For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) for these traces. (Where H is the vertical distance to the closest return path for that particular trace.) For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) these traces. (Where H is the vertical distance to the closest return path for that particular trace.) For Clock traces: Maintain at least 5H spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm). Route clock signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not exceed 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specification. These signals should maintain the following spacings: <ul style="list-style-type: none"> Clocks should maintain a length-matching between clock pairs of ± 5 ps. Clocks should maintain a length-matching between positive (\bar{p}) and negative (\bar{n}) signals of ± 2 ps, routed in parallel. Space between different pairs should be at least two times the trace width of the differential pair to minimize loss and maximize interconnect density. To avoid mismatched transmission line to via, Intel recommends that you use Ground Signal Signal Ground (GSSG) topology for your clock pattern—GND CLKP CKLN GND. Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
continued...	



Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> Route address and command signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not be more than 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specifications. UDIMMs are more susceptible to cross-talk and are generally noisier than buffered DIMMs. Therefore, route address and command signals of UDIMMs on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. Do not route differential clock (CK) and clock enable (CKE) signals close to address signals. Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
DQ, DM, and DQS Routing Rules	<ul style="list-style-type: none"> All the trace length matching requirements are from the FPGA package ball to the SDRAM package ball, which means you must consider trace mismatching on different DIMM raw cards. Match in length all DQ, DQS, and DM signals within a given byte-lane group with a maximum deviation of ± 10 ps. Ensure to route all DQ, DQS, and DM signals within a given byte-lane group on the same layer to avoid layer to layer transmission velocity differences, which otherwise increase the skew within the group. Do not count on FPGAs to deskew for more than 20 ps of DQ group skew. The skew algorithm only removes the following possible uncertainties: <ul style="list-style-type: none"> Minimum and maximum die IOE skew or delay mismatch Minimum and maximum device package skew or mismatch Board delay mismatch of 20 ps Memory component DQ skew mismatch Increasing any of these four parameters runs the risk of the deskew algorithm limiting, failing to correct for the total observed system skew. If the algorithm cannot compensate without limiting the correction, timing analysis shows reduced margins. For memory interfaces with leveling, the timing between the DQS and clock signals on each device calibrates dynamically to meet tDQSS. To make sure the skew is not too large for the leveling circuit's capability, follow these rules: <ul style="list-style-type: none"> Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device: $(CK_i) - DQSi > 0$; $0 < i < \text{number of components} - 1$. For DIMMs, ensure that the CK trace is longer than the longest DQS trace at the DIMM connector. Total skew of CLK and DQS signal between groups is less than one clock cycle: $(CK_i + DQSi)_{\text{max}} - (CK_i + DQSi)_{\text{min}} < 1 \times tCK$ (If you are using a DIMM topology, your delay and skew must take into consideration values for the actual DIMM.)
continued...	



Parameter	Guidelines
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always ensure that the signals related to the memory interface are routed between appropriate GND or power layers. For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Clock traces: Maintain at least 5H spacing between two clock pairs or a clock pair and any other memory interface trace, where H is the vertical distance to the closest return path for that particular trace.
Intel Quartus Prime Software Settings for Board Layout	<ul style="list-style-type: none"> To perform timing analyses on board and I/O buffers, use a third-party simulation tool to simulate all timing information such as skew, ISI, crosstalk, and type the simulation result into the Board Settings tab in the parameter editor. Do not use advanced I/O timing model (AIOT) or board trace model unless you do not have access to any third party tool. AIOT provides reasonable accuracy but tools like HyperLynx provide better results.
Notes to Table: 1. For point-to-point and DIMM interface designs, refer to the Micron website, www.micron.com .	

Related Links


- [Package Deskew](#) on page 198
- [External Memory Interface Spec Estimator](#)
- www.micron.com

5.4.4.3 Length Matching Rules

The following topics provide guidance on length matching for different types of SDRAM signals.

Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. The following figure shows the component routing guidelines for address and command signals.

Figure 49. SDRAM Component Address and Command Routing Guidelines



If using discrete components:

$$\begin{aligned}
 x &= y \pm 20 \text{ ps} \\
 x + x1 &= y + y1 \pm 20 \text{ ps} \\
 x + x1 + x2 &= y + y1 + y2 \pm 20 \text{ ps} \\
 x + x1 + x2 + x3 &= y + y1 + y2 + y3 \pm 20 \text{ ps}
 \end{aligned}$$

If using a DIMM topology:

$$x = y \pm 20 \text{ ps}$$

The timing between the DQS and clock signals on each device calibrates dynamically to meet t_{DQSS} . The following figure shows the delay requirements to align DQS and clock signals. To ensure that the skew is not too large for the leveling circuit's capability, follow these rules:

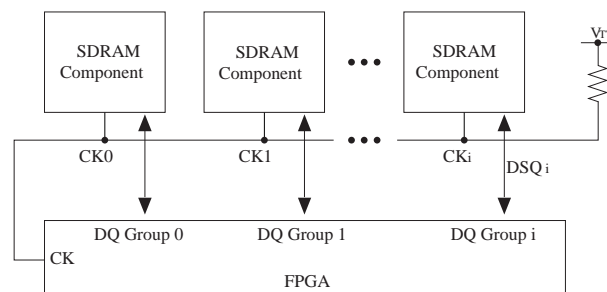
- Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device:

$$CK_i - DQSi > 0; 0 < i < \text{number of components} - 1$$

- Total skew of CLK and DQS signal between groups is less than one clock cycle:

$$(CK_i + DQSi)_{\max} - (CK_i + DQSi)_{\min} < 1 \times t_{CK}$$

Figure 50. Delaying DQS Signal to Align DQS and Clock



CK_i = Clock signal propagation delay to device i

$DQSi$ = DQ/DQS signals propagation delay to group i

Clk pair matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, match the trace lengths up to the DIMM connector. If you are using discrete components, match the lengths for all the memory components connected in the fly-by chain.

DQ group length matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, apply the DQ group trace matching rules described in the guideline table earlier up to the DIMM connector. If you are using discrete components, match the lengths up to the respective memory components.


When you are using DIMMs, it is assumed that lengths are tightly matched within the DIMM itself. You should check that appropriate traces are length-matched within the DIMM.

5.4.4.4 Spacing Guidelines

This topic provides recommendations for minimum spacing between board traces for various signal traces.

Spacing Guidelines for DQ, DQS, and DM Traces


Maintain a minimum of $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)






Spacing Guidelines for Address and Command and Control Traces

Maintain at least $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)



Spacing Guidelines for Clock Traces

Maintain at least $5H$ spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)



5.4.4.5 Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface (>72 bits)

The following topics discuss different ways to lay out a wider DDR3 or DDR4 SDRAM interface to the FPGA. Choose the topology based on board trace simulation and the timing budget of your system.

The EMIF IP supports up to a 144-bit wide DDR3 interface. You can use discrete components or DIMMs to implement a wide interface (any interface wider than 72 bits). Intel recommends using leveling when you implement a wide interface with DDR3 components.

When you lay out for a wider interface, all rules and constraints discussed in the previous sections still apply. The DQS, DQ, and DM signals are point-to-point, and all the same rules discussed in *Design Layout Guidelines* apply.

The main challenge for the design of the fly-by network topology for the clock, command, and address signals is to avoid signal integrity issues, and to make sure you route the DQS, DQ, and DM signals with the chosen topology.

Related Links

[Design Layout Guidelines](#) on page 188

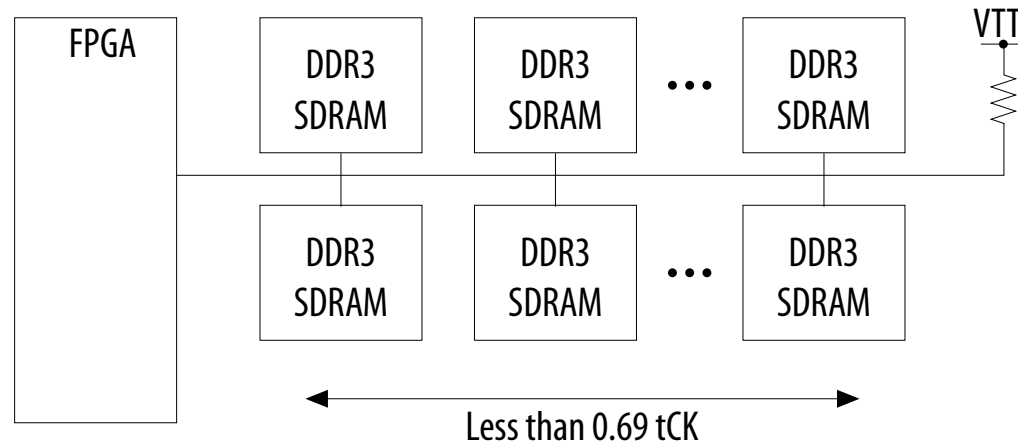
5.4.4.6 Fly-By Network Design for Clock, Command, and Address Signals

The EMIF IP requires the flight-time skew between the first SDRAM component and the last SDRAM component to be less than $0.69 t_{CK}$ for memory clocks. This constraint limits the number of components you can have for each fly-by network.

If you design with discrete components, you can choose to use one or more fly-by networks for the clock, command, and address signals.

The following figure shows an example of a single fly-by network topology.

Figure 51. Single Fly-By Network Topology



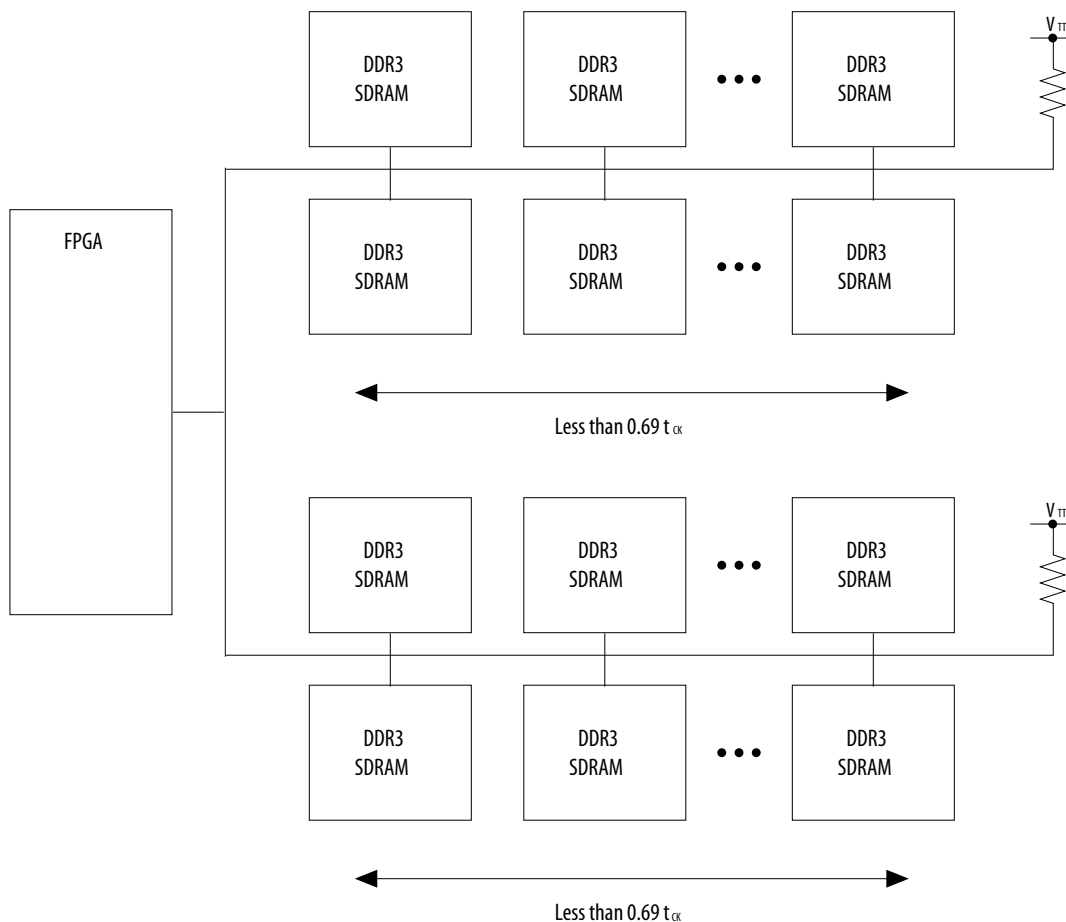
Every SDRAM component connected to the signal is a small load that causes discontinuity and degrades the signal. When using a single fly-by network topology, to minimize signal distortion, follow these guidelines:

- Use ×16 device instead ×4 or ×8 to minimize the number of devices connected to the trace.
- Keep the stubs as short as possible.
- Even with added loads from additional components, keep the total trace length short; keep the distance between the FPGA and the first SDRAM component less than 5 inches.
- Simulate clock signals to ensure a decent waveform.

The following figure shows an example of a double fly-by network topology. This topology is not rigid but you can use it as an alternative option. The advantage of using this topology is that you can have more SDRAM components in a system without violating the 0.69 tCK rule. However, as the signals branch out, the components still create discontinuity.




Figure 52. Double Fly-By Network Topology



You must perform simulations to find the location of the split, and the best impedance for the traces before and after the split.

The following figure shows a way to minimize the discontinuity effect. In this example, keep TL2 and TL3 matches in length. Keep TL1 longer than TL2 and TL3, so that it is easier to route all the signals during layout.

Figure 53. Minimizing Discontinuity Effect



You can also consider using a DIMM on each branch to replace the components. Because the trace impedance on the DIMM card is 40-ohm to 60-ohm, perform a board trace simulation to control the reflection to within the level your system can tolerate.

Using the fly-by daisy chain topology increases the complexity of the datapath and controller design to achieve leveling, but also greatly improves performance and eases board layout for SDRAM implementations.

You can also use the SDRAM components without leveling in a design if it may result in a more optimal solution, or use with devices that support the required electrical interface standard, but do not support the required read and write leveling functionality.

5.4.5 Package Deskew

Trace lengths inside the device package are not uniform for all package pins. The nonuniformity of package traces can affect system timing for high frequencies. A package deskew option is available in the Intel Quartus Prime software.

If you do not enable the package deskew option, the Intel Quartus Prime software uses the package delay numbers to adjust skews on the appropriate signals; you do not need to adjust for package delays on the board traces. If you do enable the package deskew option, the Intel Quartus Prime software does not use the package delay numbers for timing analysis, and you must deskew the package delays with the board traces for the appropriate signals for your design.

Related Links

[Layout Guidelines](#) on page 190

5.4.5.1 DQ/DQS/DM Deskew

To get the package delay information, follow these steps:



1. Select the **FPGA DQ/DQS Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

5.4.5.2 Address and Command Deskew

Deskew address and command delays as follows:

1. Select the **FPGA Address/Command Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

5.4.5.3 Package Deskew Recommendations for Intel Stratix 10 Devices

The following table shows package deskew recommendations for Intel Stratix 10 devices.

As operating frequencies increase, it becomes increasingly critical to perform package deskew. The frequencies listed in the table are the *minimum* frequencies for which you must perform package deskew.

If you plan to use a listed protocol at the specified frequency or higher, you must perform package deskew.

Protocol	Minimum Frequency (MHz) for Which to Perform Package Deskew		
	Single Rank	Dual Rank	Quad Rank
DDR4	933	800	667
DDR3	933	800	667
LPDDR3	667	533	Not required
QDR IV	933	Not applicable	Not applicable
RLDRAM 3	933	667	Not applicable
QDR II, II+, II+ Xtreme	Not required	Not applicable	Not applicable

5.4.5.4 Deskew Example

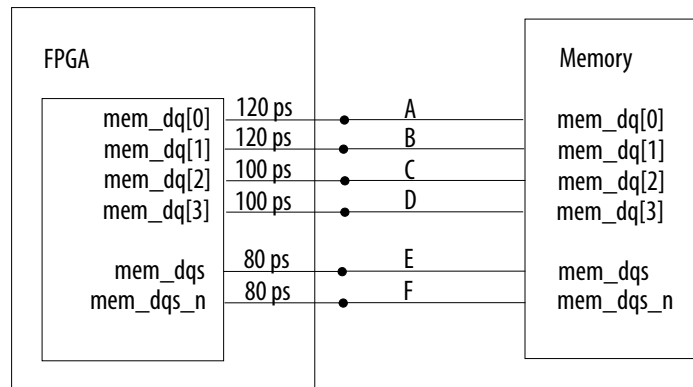
Consider an example where you want to deskew an interface with 4 DQ pins, 1 DQS pin, and 1 DQSn pin.

Let's assume an operating frequency of 667 MHz, and the package lengths for the pins reported in the **.pin** file as follows:

```
dq[0] = 120 ps
dq[1] = 120 ps
dq[2] = 100 ps
dq[3] = 100 ps
dqs = 80 ps
dqs_n = 80 ps
```

The following figure illustrates this example.

Figure 54. Deskew Example

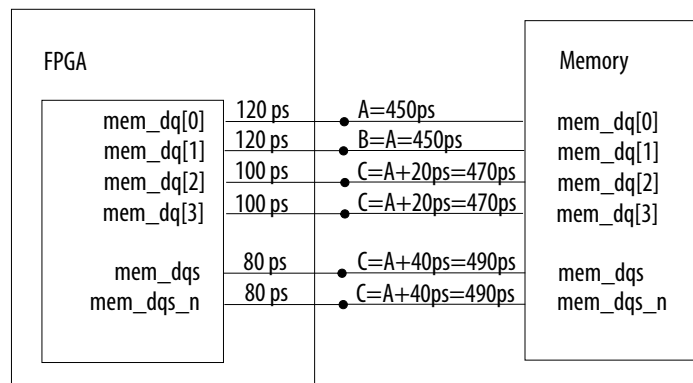


When you perform length matching for all the traces in the DQS group, you must take package delays into consideration. Because the package delays of traces A and B are 40 ps longer than the package delays of traces E and F, you would need to make the board traces for E and F 40 ps longer than the board traces for A and B.

A similar methodology would apply to traces C and D, which should be 20 ps longer than the lengths of traces A and B.

The following figure shows this scenario with the length of trace A at 450 ps.

Figure 55. Deskew Example with Trace Delay Calculations





When you enter the board skews into the Board Settings tab of the DDR3 parameter editor, you should calculate the board skew parameters as the sums of board delay and corresponding package delay. If a pin does not have a package delay (such as address and command pins), you should use the board delay only.

The example of the preceding figure shows an ideal case where board skews are perfectly matched. In reality, you should allow plus or minus 10 ps of skew mismatch within a DQS group (DQ/DQS/DM).



6 Intel Stratix 10 EMIF IP for DDR4

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for DDR4.

6.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

6.1.1 Intel Stratix 10 EMIF IP DDR4 Parameters: General

Table 266. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_DDR4_CONFIG_ENUM)
Instantiate two controllers sharing a Ping Pong PHY	Specifies the instantiation of two identical memory controllers that share an address/command bus through the use of Ping Pong PHY. This parameter is available only if you specify the Hard PHY and Hard Controller option. When this parameter is enabled, the IP exposes two independent Avalon interfaces to the user logic, and a single external memory interface with double width for the data bus and the CS#, CKE, ODT, and CK/CK# signals. (Identifier: PHY_DDR4_USER_PING_PONG_EN)

Table 267. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_DDR4_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_DDR4_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to
continued...	

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.



Display Name	Description
	use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_DDR4_USER_REF_CLK_FREQ_MHZ)
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_DDR4_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_DDR4_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_DDR4_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 268. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 269. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)



Table 270. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)

Table 271. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 272. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

6.1.2 Intel Stratix 10 EMIF IP DDR4 Parameters: Memory

Table 273. Group: Memory / Topology

Display Name	Description
Memory format	Specifies the format of the external memory device. The following formats are supported: Component - a Discrete memory device; UDIMM - Unregistered/Unbuffered DIMM where address/control, clock, and data are unbuffered; RDIMM - Registered DIMM where address/control and clock are buffered; LRDIMM - Load Reduction DIMM where address/control, clock, and data are buffered. LRDIMM reduces the load to increase memory speed and supports higher densities than RDIMM; SODIMM - Small Outline DIMM is similar to UDIMM but smaller in size and is typically used for systems with limited space. Some memory protocols may not be available in all formats. (Identifier: MEM_DDR4_FORMAT_ENUM)
DQ width	Specifies the total number of data pins in the interface. The maximum supported width is 144, or 72 in Ping Pong PHY mode. (Identifier: MEM_DDR4_DQ_WIDTH)
DQ pins per DQS group	Specifies the total number of DQ pins per DQS group. (Identifier: MEM_DDR4_DQ_PER_DQS)
Number of clocks	Specifies the number of CK/CK# clock pairs exposed by the memory interface. Usually more than 1 pair is required for RDIMM/LRDIMM formats. The value of this parameter depends on the memory device selected; <i>refer to the data sheet for your memory device</i> . (Identifier: MEM_DDR4_CK_WIDTH)
Number of chip selects	Specifies the total number of chip selects in the interface, up to a maximum of 4. This parameter applies to discrete components only . (Identifier: MEM_DDR4_DISCRETE_CS_WIDTH)
Number of DIMMs	Total number of DIMMs. (Identifier: MEM_DDR4_NUM_OF_DIMMS)
continued...	



Display Name	Description
Chip ID width	Specifies the number of chip ID pins. Only applicable to <i>registered and load-reduced DIMMs</i> that use 3DS/TSV memory devices. (Identifier: MEM_DDR4_CHIP_ID_WIDTH)
Number of physical ranks per DIMM	Number of ranks per DIMM. For LRDIMM, this represents the number of physical ranks on the DIMM behind the memory buffer (Identifier: MEM_DDR4_RANKS_PER_DIMM)
Number of clock enables per DIMM	Number of clock enables (i.e. CKE) per DIMM. Only applicable to registered and load-reduced DIMMs. (Identifier: MEM_DDR4_CKE_PER_DIMM)
Row address width	Specifies the number of row address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of address pins needed for access to all available rows. (Identifier: MEM_DDR4_ROW_ADDR_WIDTH)
Column address width	Specifies the number of column address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of address pins needed for access to all available columns. (Identifier: MEM_DDR4_COL_ADDR_WIDTH)
Bank address width	Specifies the number of bank address pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of bank address pins needed for access to all available banks. (Identifier: MEM_DDR4_BANK_ADDR_WIDTH)
Bank group width	Specifies the number of bank group pins. <i>Refer to the data sheet for your memory device.</i> The density of the selected memory device determines the number of bank group pins needed for access to all available bank groups. (Identifier: MEM_DDR4_BANK_GROUP_WIDTH)
Data mask	Indicates whether the interface uses data mask (DM) pins. This feature allows specified portions of the data bus to be written to memory (not available in x4 mode). One DM pin exists per DQS group. (Identifier: MEM_DDR4_DM_EN)
Write DBI	Indicates whether the interface uses write data bus inversion (DBI). This feature provides better signal integrity and write margin . This feature is unavailable if Data Mask is enabled or in x4 mode. (Identifier: MEM_DDR4_WRITE_DBI)
Read DBI	Specifies whether the interface uses read data bus inversion (DBI). Enable this feature for better signal integrity and read margin . This feature is not available in x4 configurations. (Identifier: MEM_DDR4_READ_DBI)
Enable address mirroring for odd chip-selects	Enabling address mirroring for multi-CS discrete components. Typically used when components are arranged in a clamshell layout. (Identifier: MEM_DDR4_DISCRETE_MIRROR_ADDRESSING_EN)
Enable address mirroring for odd ranks	Enabling address mirroring for dual-rank or quad-rank DIMM. (Identifier: MEM_DDR4_MIRROR_ADDRESSING_EN)
Enable ALERT#/PAR pins	Allows address/command calibration, which may provide better margins on the address/command bus. The <code>alert_n</code> signal is not accessible in the AFI or Avalon domains. This means there is no way to know whether a parity error has occurred during user mode. The parity pin is a dedicated pin in the address/command bank, but the <code>alert_n</code> pin can be placed in any bank that spans the memory interface. You should explicitly choose the location of the <code>alert_n</code> pin and place it in the address/command bank. (Identifier: MEM_DDR4_ALERT_PAR_EN)
ALERT# pin placement	Specifies placement for the <code>mem_alert_n</code> signal. If you select "I/O Lane with Address/Command Pins" , you can pick the I/O lane and pin index in the add/cmd bank with the subsequent drop down menus. If you select "I/O Lane with DQS Group" , you can specify the DQS group with which to place the <code>mem_alert_n</code> pin. If you select "Automatically select a location" , the IP automatically selects a pin for the <code>mem_alert_n</code> signal. If you select this option, no additional location constraints can be applied to

continued...



Display Name	Description
	the mem_alert_n pin, or a fitter error will result during compilation. For optimum signal integrity, you should choose " I/O Lane with Address/Command Pins ". For interfaces containing multiple memory devices, it is recommended to connect the ALERT# pins together to the ALERT# pin on the FPGA. (Identifier: MEM_DDR4_ALERT_N_PLACEMENT_ENUM)
DQS group of ALERT#	Select the DQS group with which the ALERT# pin is placed. (Identifier: MEM_DDR4_ALERT_N_DQS_GROUP)
Address/command I/O lane of ALERT#	Select the lane of the Address/Command I/O Tile where ALERT# pin is placed. (Identifier: MEM_DDR4_ALERT_N_AC_LANE)
Pin index of ALERT#	Select the pin of the Address/Command I/O Lane where ALERT# pin is placed. (Identifier: MEM_DDR4_ALERT_N_AC_PIN)

Table 274. Group: Memory / Latency and Burst

Display Name	Description
Memory CAS latency setting	Specifies the number of clock cycles between the read command and the availability of the first bit of output data at the memory device. Overall read latency equals the additive latency (AL) + the CAS latency (CL). Overall read latency depends on the memory device selected; <i>refer to the datasheet for your device</i> . (Identifier: MEM_DDR4_TCL)
Memory write CAS latency setting	Specifies the number of clock cycles from the release of internal write to the latching of the first data in at the memory device. This value depends on the memory device selected; <i>refer to the datasheet for your device</i> . (Identifier: MEM_DDR4_WTCL)
Memory additive CAS latency setting	Determines the posted CAS additive latency of the memory device. Enable this feature to improve command and bus efficiency, and increase system bandwidth . (Identifier: MEM_DDR4_ATCL_ENUM)

Table 275. Group: Memory / Mode Register Settings

Display Name	Description
Hide advanced mode register settings	Show or hide advanced mode register settings. Changing advanced mode register settings to non-default values is strongly discouraged. (Identifier: MEM_DDR4_HIDE_ADV_MR_SETTINGS)
Addr/CMD parity latency	Additional latency incurred by enabling address/command parity check after calibration. Select a value to enable address/command parity with the latency associated with the selected value. Select Disable to disable address/command parity. Address/command is enabled automatically and as-needed during calibration regardless of the value of this setting. (Identifier: MEM_DDR4_AC_PARITY_LATENCY)
Burst Length	Specifies the DRAM burst length which determines how many consecutive addresses should be accessed for a given read/write command. (Identifier: MEM_DDR4_BL_ENUM)
Read Burst Type	Indicates whether accesses within a given burst are in sequential or interleaved order. Select sequential if you are using the Intel-provided memory controller. (Identifier: MEM_DDR4_BT_ENUM)
Enable the DLL in memory device	Enable the DLL in memory device (Identifier: MEM_DDR4_DLL_EN)
Auto self-refresh method	Indicates whether to enable or disable auto self-refresh. Auto self-refresh allows the controller to issue self-refresh requests, rather than manually issuing self-refresh in order for memory to retain data. (Identifier: MEM_DDR4_ASR_ENUM)
Write CRC enable	Write CRC enable (Identifier: MEM_DDR4_WRITE_CRC)
continued...	



Display Name	Description
DDR4 geardown mode	Set DDR4 geardown mode for control signals at high frequency (Identifier: MEM_DDR4_GEARDOWN)
Per-DRAM addressability	Per-DRAM addressability enable (Identifier: MEM_DDR4_PER_DRAM_ADDR)
Temperature sensor readout	Temperature sensor readout enable (Identifier: MEM_DDR4_TEMP_SENSOR_READOUT)
Fine granularity refresh	Increased frequency of refresh in exchange for shorter refresh. Shorter tRFC and increased cycle time can produce higher bandwidth . (Identifier: MEM_DDR4_FINE_GRANULARITY_REFRESH)
MPR read format	Multipurpose register readout format (Identifier: MEM_DDR4_MPR_READ_FORMAT)
Maximum power down mode	Maximum power down mode (Identifier: MEM_DDR4_MAX_POWERDOWN)
Temperature controlled refresh range	Indicates temperature controlled refresh range where normal temperature mode covers 0C to 85C and extended mode covers 0C to 95C. (Identifier: MEM_DDR4_TEMP_CONTROLLED_RFSH_RANGE)
Temperature controlled refresh enable	Indicates whether to enable temperature controlled refresh, which allows the device to adjust the internal refresh period to be longer than tREFI of the normal temperature range by skipping external refresh commands. (Identifier: MEM_DDR4_TEMP_CONTROLLED_RFSH_ENA)
Internal VrefDQ monitor	Indicates whether to enable the internal VrefDQ monitor. (Identifier: MEM_DDR4_INTERNAL_VREFDQ_MONITOR)
CS to Addr/CMD Latency	CS to Addr/CMD Latency (CAL mode) for idle state DRAM receiver power reduction (Identifier: MEM_DDR4_CAL_MODE)
Self refresh abort	Self refresh abort for latency reduction. (Identifier: MEM_DDR4_SELF_RFSH_ABORT)
Read preamble training mode enable	Read preamble training mode enable (Identifier: MEM_DDR4_READ_PREAMBLE_TRAINING)
Read preamble	Number of read preamble cycles. This mode register setting determines the number of cycles DQS (read) will go low before starting to toggle. (Identifier: MEM_DDR4_READ_PREAMBLE)
Write preamble	Write preamble cycles. (Identifier: MEM_DDR4_WRITE_PREAMBLE)
ODT input buffer during powerdown mode	Indicates whether to enable on-die termination (ODT) input buffer during powerdown mode. (Identifier: MEM_DDR4_ODT_IN_POWERDOWN)
Addr/CMD persistent error	If set, Addr/CMD parity errors continue to be checked after a previous Addr/CMD parity error (Identifier: MEM_DDR4_AC_PERSISTENT_ERROR)

6.1.3 Intel Stratix 10 EMIF IP DDR4 Parameters: Mem I/O

Table 276. Group: Mem I/O / Memory I/O Settings

Display Name	Description
Output drive strength setting	Specifies the output driver impedance setting at the memory device. To obtain optimum signal integrity performance, select option based on board simulation results . (Identifier: MEM_DDR4_DRV_STR_ENUM)
Dynamic ODT (Rtt_WR) value	Specifies the mode of the dynamic on-die termination (ODT) during writes to the memory device (used for multi-rank configurations). For optimum signal integrity performance, select this option based on board simulation results . (Identifier: MEM_DDR4_RTT_WR_ENUM)

continued...



Display Name	Description
ODT Rtt nominal value	Determines the nominal on-die termination value applied to the DRAM. The termination is applied any time that ODT is asserted. If you specify a different value for <code>RTT_WR</code> , that value takes precedence over the values mentioned here. For optimum signal integrity performance, select your option based on board simulation results . (Identifier: <code>MEM_DDR4_RTT_NOM_ENUM</code>)
RTT PARK	If set, the value is applied when the DRAM is not being written AND ODT is not asserted HIGH. (Identifier: <code>MEM_DDR4_RTT_PARK</code>)
RCD CA Input Bus Termination	Specifies the input termination setting for the following pins of the registering clock driver: <code>DA0</code> . . <code>DA17</code> , <code>DBA0</code> . . <code>DBA1</code> , <code>DBG0</code> . . <code>DBG1</code> , <code>DACT_n</code> , <code>DC2</code> , <code>DPAR</code> . This parameter determines the value of bits <code>DA[1:0]</code> of control word <code>RC7x</code> of the registering clock driver. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_RCD_CA_IBT_ENUM</code>)
RCD DCS[3:0]_n Input Bus Termination	Specifies the input termination setting for the following pins of the registering clock driver: <code>DCS[3:0]_n</code> . This parameter determines the value of bits <code>DA[3:2]</code> of control word <code>RC7x</code> of the registering clock driver. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_RCD_CS_IBT_ENUM</code>)
RCD DCKE Input Bus Termination	Specifies the input termination setting for the following pins of the registering clock driver: <code>DCKE0</code> , <code>DCKE1</code> . This parameter determines the value of bits <code>DA[5:4]</code> of control word <code>RC7x</code> of the registering clock driver. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_RCD_CKE_IBT_ENUM</code>)
RCD DODT Input Bus Termination	Specifies the input termination setting for the following pins of the registering clock driver: <code>DODT0</code> , <code>DODT1</code> . This parameter determines the value of bits <code>DA[7:6]</code> of control word <code>RC7x</code> of the registering clock driver. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_RCD_ODT_IBT_ENUM</code>)
DB Host Interface DQ RTT_NOM	Specifies the <code>RTT_NOM</code> setting for the host interface of the data buffer. Only "RTT_NOM disabled" is supported. This parameter determines the value of the control word <code>BC00</code> of the data buffer. (Identifier: <code>MEM_DDR4_DB_RTT_NOM_ENUM</code>)
DB Host Interface DQ RTT_WR	Specifies the <code>RTT_WR</code> setting of the host interface of the data buffer. This parameter determines the value of the control word <code>BC01</code> of the data buffer. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_DB_RTT_WR_ENUM</code>)
DB Host Interface DQ RTT_PARK	Specifies the <code>RTT_PARK</code> setting for the host interface of the data buffer. This parameter determines the value of control word <code>BC02</code> of the data buffer. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_DB_RTT_PARK_ENUM</code>)
DB Host Interface DQ Driver	Specifies the driver impedance setting for the host interface of the data buffer. This parameter determines the value of the control word <code>BC03</code> of the data buffer. Perform board simulation to obtain the optimal value for this setting . (Identifier: <code>MEM_DDR4_DB_DQ_DRV_ENUM</code>)
Use recommended initial VrefDQ value	Specifies to use the recommended initial <code>VrefDQ</code> value. This value is used as a starting point and may change after calibration. (Identifier: <code>MEM_DDR4_DEFAULT_VREFOUT</code>)
VrefDQ training value	<code>VrefDQ</code> training value. (Identifier: <code>MEM_DDR4_USER_VREFDQ_TRAINING_VALUE</code>)
VrefDQ training range	<code>VrefDQ</code> training range. (Identifier: <code>MEM_DDR4_USER_VREFDQ_TRAINING_RANGE</code>)



Table 277. Group: Mem I/O / RDIMM/LRDIMM Serial Presence Detect (SPD) Data

Display Name	Description
SPD Byte 137 - RCD Drive Strength for Command/Address	Specifies the drive strength of the registering clock driver's control and command/address outputs to the DRAM. The value must come from Byte 137 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_137_RCD_CA_DRV)
SPD Byte 138 - RCD Drive Strength for CK	Specifies the drive strength of the registering clock driver's clock outputs to the DRAM. The value must come from Byte 138 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_138_RCD_CK_DRV)
SPD Byte 140 - DRAM VrefDQ for Package Rank 0	Specifies the VrefDQ setting for package rank 0 of an LRDIMM. The value must come from Byte 140 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_140_DRAM_VREFDQ_R0)
SPD Byte 141 - DRAM VrefDQ for Package Rank 1	Specifies the VrefDQ setting for package rank 1 of an LRDIMM. The value must come from Byte 141 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_141_DRAM_VREFDQ_R1)
SPD Byte 142 - DRAM VrefDQ for Package Rank 2	Specifies the VrefDQ setting for package rank 2 (if it exists) of an LRDIMM. The value must come from Byte 142 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_142_DRAM_VREFDQ_R2)
SPD Byte 143 - DRAM VrefDQ for Package Rank 3	Specifies the VrefDQ setting for package rank 3 (if it exists) of an LRDIMM. The value must come from Byte 143 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_143_DRAM_VREFDQ_R3)
SPD Byte 144 - DB VrefDQ for DRAM Interface	Specifies the VrefDQ setting of the data buffer's DRAM interface. The value must come from Byte 144 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_144_DB_VREFDQ)
SPD Byte 145-147 - DB MDQ Drive Strength and RTT	Specifies the drive strength of the MDQ pins of the data buffer's DRAM interface. The value must come from either Byte 145 (data rate = 1866), 146 (1866 data rate = 2400), or 147 (2400 data rate = 3200) of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_145_DB_MDQ_DRV)
SPD Byte 148 - DRAM Drive Strength	Specifies the drive strength of the DRAM. The value must come from Byte 148 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_148_DRAM_DRV)
SPD Byte 149-151 - DRAM ODT (RTT_WR and RTT_NOM)	Specifies the RTT_WR and RTT_NOM setting of the DRAM. The value must come from either Byte 149 (data rate = 1866), 150 (1866 data rate = 2400), or 151 (2400 data rate = 3200) of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_149_DRAM_RTT_WR_NOM)
SPD Byte 152-154 - DRAM ODT (RTT_PARK)	Specifies the RTT_PARK setting of the DRAM. The value must come from either Byte 152 (data rate = 1866), 153 (1866 data rate = 2400), or 154 (2400 data rate = 3200) of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_152_DRAM_RTT_PARK)
RCD and DB Manufacturer (LSB)	Specifies the LSB of the ID code of the registering clock driver and data buffer manufacturer. The value must come from Byte 133 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_133_RCD_DB_VENDOR_LSB)
RCD and DB Manufacturer (MSB)	Specifies the MSB of the ID code of the registering clock driver and data buffer manufacturer. The value must come from Byte 134 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_134_RCD_DB_VENDOR_MSB)
RCD Revision Number	Specifies the die revision of the registering clock driver. The value must come from Byte 135 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_135_RCD_REV)
DB Revision Number	Specifies the die revision of the data buffer. The value must come from Byte 139 of the SPD from the DIMM vendor. (Identifier: MEM_DDR4_SPD_139_DB_REV)



Table 278. Group: Mem I/O / ODT Activation

Display Name	Description
Use Default ODT Assertion Tables	Enables the default ODT assertion pattern as determined from vendor guidelines. These settings are provided as a default only; <i>you should simulate your memory interface to determine the optimal ODT settings and assertion patterns.</i> (Identifier: MEM_DDR4_USE_DEFAULT_ODT)

6.1.4 Intel Stratix 10 EMIF IP DDR4 Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.

Table 279. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_DDR4_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_DDR4_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_DDR4_DEFAULT_IO)

Table 280. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR4_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR4_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_DDR4_USER_AC_SLEW_RATE_ENUM)

**Table 281. Group: FPGA I/O / FPGA I/O Settings / Memory Clock**

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR4_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR4_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_DDR4_USER_CK_SLEW_RATE_ENUM)

Table 282. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_DDR4_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR4_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_DDR4_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_DDR4_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_DDR4_USER_STARTING_VREFIN)

Table 283. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_DDR4_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_DDR4_USER_RZQ_IO_STD_ENUM)

6.1.5 Intel Stratix 10 EMIF IP DDR4 Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).



Table 284. Group: Mem Timing / Parameters dependent on Speed Bin

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_DDR4_SPEEDBIN_ENUM)
tIS (base)	tIS (base) refers to the setup time for the Address/Command/Control (A) bus to the rising edge of CK. (Identifier: MEM_DDR4_TIS_PS)
tIS (base) AC level	tIS (base) AC level refers to the voltage level which the address/command signal must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_DDR4_TIS_AC_MV)
tIH (base)	tIH (base) refers to the hold time for the Address/Command (A) bus after the rising edge of CK. Depending on what AC level the user has chosen for a design, the hold margin can vary (this variance will be automatically determined when the user chooses the "tIH (base) AC level"). (Identifier: MEM_DDR4_TIH_PS)
tIH (base) DC level	tIH (base) DC level refers to the voltage level which the address/command signal must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_DDR4_TIH_DC_MV)
TdiVW_total	TdiVW_total describes the minimum horizontal width of the DQ eye opening required by the receiver (memory device/DIMM). It is measured in UI (1UI = half the memory clock period). (Identifier: MEM_DDR4_TDIVW_TOTAL_UI)
VdiVW_total	VdiVW_total describes the Rx Mask voltage , or the minimum vertical width of the DQ eye opening required by the receiver (memory device/DIMM). It is measured in mV. (Identifier: MEM_DDR4_VDIVW_TOTAL)
tDQSQ	tDQSQ describes the latest valid transition of the associated DQ pins for a READ . tDQSQ specifically refers to the DQS, DQS# to DQ skew. It is the length of time between the DQS, DQS# crossing to the last valid transition of the slowest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_DDR4_TDQSQ_UI)
tQH	tQH specifies the output hold time for the DQ in relation to DQS, DQS# . It is the length of time between the DQS, DQS# crossing to the earliest invalid transition of the fastest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_DDR4_TQH_UI)
tDVWp	Data valid window per device per pin (Identifier: MEM_DDR4_TDVWP_UI)
tDQSCK	tDQSCK describes the skew between the memory clock (CK) and the input data strobes (DQS) used for reads . It is the time between the rising data strobe edge (DQS, DQS#) relative to the rising CK edge. (Identifier: MEM_DDR4_TDQSCK_PS)
tDQSS	tDQSS describes the skew between the memory clock (CK) and the output data strobes used for writes . It is the time between the rising data strobe edge (DQS, DQS#) relative to the rising CK edge. (Identifier: MEM_DDR4_TDQSS_CYC)
tQSH	tQSH refers to the differential High Pulse Width, which is measured as a percentage of tCK. It is the time during which the DQS is high for a read . (Identifier: MEM_DDR4_TQSH_CYC)
tDSH	tDSH specifies the write DQS hold time . This is the time difference between the rising CK edge and the falling edge of DQS, measured as a percentage of tCK. (Identifier: MEM_DDR4_TDSH_CYC)
tDSS	tDSS describes the time between the falling edge of DQS to the rising edge of the next CK transition . (Identifier: MEM_DDR4_TDSS_CYC)
continued...	



Display Name	Description
tWLS	tWLS describes the write leveling setup time . It is measured from the rising edge of CK to the rising edge of DQS. (Identifier: MEM_DDR4_TWLS_CYC)
tWLH	tWLH describes the write leveling hold time . It is measured from the rising edge of DQS to the rising edge of CK. (Identifier: MEM_DDR4_TWLH_CYC)
tINIT	tINIT describes the time duration of the memory initialization after a device power-up . After RESET_n is de-asserted, wait for another 500us until CKE becomes active. During this time, the DRAM will start internal initialization; this will be done independently of external clocks. (Identifier: MEM_DDR4_TINIT_US)
tMRD	The mode register set command cycle time, tMRD is the minimum time period required between two MRS commands . (Identifier: MEM_DDR4_TMRD_CK_CYC)
tRAS	tRAS describes the activate to precharge duration . A row cannot be deactivated until the tRAS time has been met. Therefore tRAS determines how long the memory has to wait after a activate command before a precharge command can be issued to close the row. (Identifier: MEM_DDR4_TRAS_NS)
tRCD	tRCD, row command delay , describes the active to read/write time . It is the amount of delay between the activation of a row through the RAS command and the access to the data through the CAS command. (Identifier: MEM_DDR4_TRCD_NS)
tRP	tRP refers to the Precharge (PRE) command period . It describes how long it takes for the memory to disable access to a row by precharging and before it is ready to activate a different row. (Identifier: MEM_DDR4_TRP_NS)
tWR	tWR refers to the Write Recovery time . It specifies the amount of clock cycles needed to complete a write before a precharge command can be issued. (Identifier: MEM_DDR4_TWR_NS)

Table 285. Group: Mem Timing / Parameters dependent on Speed Bin, Operating Frequency, and Page Size

Display Name	Description
tRRD_S	tRRD_S refers to the Activate to Activate Command Period (short) . It is the minimum time interval between two activate commands to the different bank groups . For 3DS devices, this parameter is the same as tRRD_S_slr (i.e. tRRD_S within the same logical rank) in the memory data sheet. (Identifier: MEM_DDR4_TRRD_S_CYC)
tRRD_L	tRRD_L refers to the Activate to Activate Command Period (long) . It is the minimum time interval (measured in memory clock cycles) between two activate commands to the same bank group . For 3DS devices, this parameter is the same as tRRD_L_slr (i.e. tRRD_L within the same logical rank) in the memory data sheet. (Identifier: MEM_DDR4_TRRD_L_CYC)
tRRD_dlr	tRRD_dlr refers to the Activate to Activate Command Period to Different Logical Ranks . It is the minimum time interval (measured in memory clock cycles) between two activate commands to different logical ranks within a 3DS DDR4 device. (Identifier: MEM_DDR4_TRRD_DLR_CYC)
tFAW	tFAW refers to the four activate window time . It describes the period of time during which only four banks can be active. For 3DS devices, this parameter is the same as tFAW_slr (i.e. tFAW within the same logical rank) in the memory data sheet. (Identifier: MEM_DDR4_TFAW_NS)
<i>continued...</i>	



Display Name	Description
tFAW_dlr	tFAW_dlr refers to the four activate window to different logical ranks . It describes the period of time during which only four banks can be active across all logical ranks within a 3DS DDR4 device. (Identifier: MEM_DDR4_TFAW_DLR_CYC)
tCCD_S	tCCD_S refers to the CAS_n-to-CAS_n delay (short) . It is the minimum time interval between two read/write (CAS) commands to different bank groups . (Identifier: MEM_DDR4_TCCD_S_CYC)
tCCD_L	tCCD_L refers to the CAS_n-to-CAS_n delay (long) . It is the minimum time interval between two read/write (CAS) commands to the same bank group . (Identifier: MEM_DDR4_TCCD_L_CYC)
tWTR_S	tWTR_S or Write Timing Parameter refers to the Write to Read period for different bank groups . It describes the delay from start of internal write transaction to internal read command, for accesses to the different bank group. The delay is measured from the first rising memory clock edge after the last write data is received to the rising memory clock edge when a read command is received. (Identifier: MEM_DDR4_TWTR_S_CYC)
tWTR_L	tWTR_L or Write Timing Parameter refers to the Write to Read period for the same bank group . It describes the delay from start of internal write transaction to internal read command, for accesses to the same bank group. The delay is measured from the first rising memory clock edge after the last write data is received to the rising memory clock edge when a read command is received. (Identifier: MEM_DDR4_TWTR_L_CYC)

Table 286. Group: Mem Timing / Parameters dependent on Density and Temperature

Display Name	Description
tRFC	tRFC refers to the Refresh Cycle Time . It is the amount of delay after a refresh command before an activate command can be accepted by the memory. This parameter is dependent on the memory density and is necessary for proper hardware functionality. For 3DS devices, this parameter is the same as tRFC_slr (i.e. tRFC within the same logical rank) in the memory data sheet. (Identifier: MEM_DDR4_TRFC_NS)
tRFC_dlr	tRFC_dlr refers to the Refresh Cycle Time to different logical rank . It is the amount of delay after a refresh command to one logical rank before an activate command can be accepted by another logical rank within a 3DS DDR4 device. This parameter is dependent on the memory density and is necessary for proper hardware functionality. (Identifier: MEM_DDR4_TRFC_DLR_NS)
tREFI	tREFI refers to the average periodic refresh interval . It is the maximum amount of time the memory can tolerate in between each refresh command (Identifier: MEM_DDR4_TREFI_US)

6.1.6 Intel Stratix 10 EMIF IP DDR4 Parameters: Board

Table 287. Group: Board / Intersymbol Interference/Crosstalk


Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation</i>
continued...	



Display Name	Description
	using IBIS models and Hyperlynx*, and manually enter values based on your simulation results, instead of using the default values. (Identifier: BOARD_DDR4_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to ISI and crosstalk effects. The number to be entered is the total loss of margin on both the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR4_USER_AC_ISI_NS)
Read DQS/DQS# ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQS/DQS# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR4_USER_RCLK_ISI_NS)
Read DQ ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQ signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold side (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR4_USER_RDATA_ISI_NS)
Write DQS/DQS# ISI/crosstalk	The reduction of the write data window due to ISI and crosstalk effects on the DQS/DQS# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR4_USER_WCLK_ISI_NS)
Write DQ ISI/crosstalk	The reduction of the write data window due to ISI and crosstalk effects on the DQ signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_DDR4_USER_WDATA_ISI_NS)

Table 288. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (DQS group)	Enable this parameter if you are compensating for package skew on the DQ, DQS, and DM buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_DDR4_IS_SKEW_WITHIN_DQS_DESKEWED)
Maximum board skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. This value affects the read capture and write margins. (Identifier: BOARD_DDR4_BRD_SKEW_WITHIN_DQS_NS)
Maximum system skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_DDR4_PKG_BRD_SKEW_WITHIN_DQS_NS)
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_DDR4_IS_SKEW_WITHIN_AC_DESKEWED)
Maximum board skew within address/command bus	The largest skew between the address and command signals. (Identifier: BOARD_DDR4_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	The largest skew between the address and command signals. Enter combined board and package skew. (Identifier: BOARD_DDR4_PKG_BRD_SKEW_WITHIN_AC_NS)
continued...	



Display Name	Description
Average delay difference between DQS and CK	The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS trace delay minus the CK trace delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. (Identifier: BOARD_DDR4_DQS_TO_CK_SKEW_NS)
Maximum delay difference between DIMMs/devices	The largest propagation delay on DQ signals between ranks (<i>applicable only when there is more than one rank</i>). For example: when you configure two ranks using one DIMM there is a short distance between the ranks for the same DQ pin; when you implement two ranks using two DIMMs the distance is larger. (Identifier: BOARD_DDR4_SKEW_BETWEEN_DIMMS_NS)
Maximum skew between DQS groups	The largest skew between DQS signals. (Identifier: BOARD_DDR4_SKEW_BETWEEN_DQS_NS)
Average delay difference between address/command and CK	The average delay difference between the address/command signals and the CK signal, calculated by averaging the longest and smallest address/command signal trace delay minus the maximum CK trace delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. (Identifier: BOARD_DDR4_AC_TO_CK_SKEW_NS)
Maximum CK delay to DIMM/device	The delay of the longest CK trace from the FPGA to any DIMM/device. (Identifier: BOARD_DDR4_MAX_CK_DELAY_NS)
Maximum DQS delay to DIMM/device	The delay of the longest DQS trace from the FPGA to any DIMM/device (Identifier: BOARD_DDR4_MAX_DQS_DELAY_NS)

6.1.7 Intel Stratix 10 EMIF IP DDR4 Parameters: Controller

Table 289. Group: Controller / Avalon Interface

Display Name	Description
Avalon Interface	Selects the Avalon Interface through which the controller interacts with user logic (Identifier: CTRL_DDR4_AVL_PROTOCOL_ENUM)

Table 290. Group: Controller / Low Power Mode

Display Name	Description
Enable Self-Refresh Control	Select this option to enable the self-refresh control on the controller top level. The control signal allows you to place the memory device into self-refresh mode, on a per chip-select basis. (Identifier: CTRL_DDR4_SELF_REFRESH_EN)
Enable Auto Power-Down	Enable this parameter to have the controller automatically place the memory device into power-down mode after a specified number of idle controller clock cycles. The idle wait time is configurable. All ranks must be idle to enter auto power-down. (Identifier: CTRL_DDR4_AUTO_POWER_DOWN_EN)
Auto Power-Down Cycles	Specifies the number of idle controller cycles after which the memory device is placed into power-down mode. You can configure the idle waiting time. The supported range for number of cycles is from 1 to 65534. (Identifier: CTRL_DDR4_AUTO_POWER_DOWN_CYCS)




Table 291. Group: Controller / Efficiency

Display Name	Description
Enable User Refresh Control	When enabled, user logic has complete control and is responsible for issuing adequate refresh commands to the memory devices, via the MMR interface. This feature provides increased control over worst-case read latency and enables you to issue refresh bursts during idle periods. (Identifier: CTRL_DDR4_USER_REFRESH_EN)
Enable Auto-Precharge Control	Select this parameter to enable the auto-precharge control on the controller top level. If you assert the auto-precharge control signal while requesting a read or write burst, you can specify whether the controller should close (auto-precharge) the currently open page at the end of the read or write burst, potentially making a future access to a different page of the same bank faster. (Identifier: CTRL_DDR4_AUTO_PRECHARGE_EN)
Address Ordering	Controls the mapping between Avalon addresses and memory device addresses. By changing the value of this parameter, you can change the mappings between the Avalon-MM address and the DRAM address. (CS = chip select, CID = chip ID in 3DS/TSV devices, BG = bank group address, Bank = bank address, Row = row address, Col = column address) (Identifier: CTRL_DDR4_ADDR_ORDER_ENUM)
Enable Reordering	Enable this parameter to allow the controller to perform command and data reordering. Reordering can improve efficiency by reducing bus turnaround time and row/bank switching time. Data reordering allows the single-port memory controller to change the order of read and write commands to achieve highest efficiency. Command reordering allows the controller to issue bank management commands early based on incoming patterns, so that the desired row in memory is already open when the command reaches the memory interface. <i>For more information, refer to the Data Reordering topic in the EMIF Handbook.</i> (Identifier: CTRL_DDR4_REORDER_EN)
Starvation limit for each command	Specifies the number of commands that can be served before a waiting command is served. The controller employs a counter to ensure that all requests are served after a pre-defined interval -- this ensures that low priority requests are not ignored, when doing data reordering for efficiency. The valid range for this parameter is from 1 to 63. <i>For more information, refer to the Starvation Control topic in the EMIF Handbook.</i> (Identifier: CTRL_DDR4_STARVE_LIMIT)
Enable Command Priority Control	Select this parameter to enable user-requested command priority control on the controller top level. This parameter instructs the controller to treat a read or write request as high-priority. The controller attempts to fill high-priority requests sooner, to reduce latency. Connect this interface to the conduit of your logic block that determines when the external memory interface IP treats the read or write request as a high-priority command. (Identifier: CTRL_DDR4_USER_PRIORITY_EN)

Table 292. Group: Controller / Configuration, Status and Error Handling

Display Name	Description
Enable Memory-Mapped Configuration and Status Register (MMR) Interface	Enable this parameter to change or read memory timing parameters, memory address size, mode register settings, controller status, and request sideband operations. (Identifier: CTRL_DDR4_MMR_EN)
Enable Error Detection and Correction Logic with ECC	Enables error-correction code (ECC) for single-bit error correction and double-bit error detection. Your memory interface must have a width of 16, 24, 40, or 72 bits to use ECC. <i>ECC is implemented as soft logic.</i> (Identifier: CTRL_DDR4_ECC_EN)
Enable Auto Error Correction to External Memory	Specifies that the controller automatically schedule and perform a write back to the external memory when a single-bit error is detected. Regardless of whether the option is enabled or disabled, the ECC feature always corrects single-bit errors before returning the read data to user logic. (Identifier: CTRL_DDR4_ECC_AUTO_CORRECTION_EN)



Table 293. Group: Controller / Data Bus Turnaround Time

Display Name	Description
Additional read-to-write turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read to a write within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_RD_TO_WR_SAME_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write to a read within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_WR_TO_RD_SAME_CHIP_DELTA_CYCS)
Additional read-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a read of another logical rank . This can resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_RD_TO_RD_DIFF_CHIP_DELTA_CYCS)
Additional read-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a write of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_RD_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a write of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_WR_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a read of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_DDR4_WR_TO_RD_DIFF_CHIP_DELTA_CYCS)

6.1.8 Intel Stratix 10 EMIF IP DDR4 Parameters: Diagnostics

Table 294. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	<p>Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process.</p> <p>Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero.</p> <p><i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations.</i> Abstract PHY is supported with skip calibration.</p>
continued...	



Display Name	Description
	(Identifier: DIAG_DDR4_SIM_CAL_MODE_ENUM)
Abstract phy for fast simulation	Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration . (Identifier: DIAG_DDR4_ABSTRACT_PHY)

Table 295. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic. If you set this parameter to "Disabled" , no debug features are enabled. If you set this parameter to "Export" , an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP core to it, or connect it to the cal_debug_out interface of another EMIF core. If you select "Add EMIF Debug Interface" , an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit. <i>Only one EMIF debug interface should be instantiated per I/O column. You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first.</i> (Identifier: DIAG_DDR4_EXPORT_SEQ_AVALON_SLAVE)
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies that the IP export an Avalon-MM master interface (cal_debug_out) which can connect to the cal_debug interface of other EMIF cores residing in the same I/O column. This parameter applies only if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to the <i>Debugging Multiple EMIFs wiki page</i> for more information about debugging multiple EMIFs. (Identifier: DIAG_DDR4_EXPORT_SEQ_AVALON_MASTER)
Interface ID	Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled , the interface ID is unused. (Identifier: DIAG_DDR4_INTERFACE_ID)
Skip address/command leveling calibration	Specifies to skip the address/command leveling stage during calibration. Address/command leveling attempts to center the memory clock edge against CS# by adjusting delay elements inside the PHY, and then applying the same delay offset to the rest of the address and command pins. (Identifier: DIAG_DDR4_SKIP_CA_LEVEL)
Skip address/command deskew calibration	Specifies to skip the address/command deskew calibration stage. Address/command deskew performs per-bit deskew for the address and command pins. (Identifier: DIAG_DDR4_SKIP_CA_DESKEW)
Skip VREF calibration	Specifies to skip the VREF stage of calibration. Enable this parameter for debug purposes only ; generally, you should include the VREF calibration stage during normal operation. (Identifier: DIAG_DDR4_SKIP_VREF_CAL)
Use Soft NIOS Processor for On-Chip Debug	Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port . <i>Only one interface in a column can activate this option.</i> (Identifier: DIAG_SOFT_NIOS_MODE)



Table 296. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the "Core clocks sharing" parameter in the "General" tab to "Master" or "Slave". (Identifier: DIAG_DDR4_EX_DESIGN_NUM_OF_SLAVES)
Enable In-System-Sources-and-Probes	Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i> . This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_DDR4_EX_DESIGN_ISSP_EN)

Table 297. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_DDR4_USE_TG_AVL_2)
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_DDR4_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the <code>altera_emif_avl_tg_2_tb.sv</code> file. (Identifier: DIAG_DDR4_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_DDR4_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_DDR4_BYPASS_STRESS_STAGE)
Export Traffic Generator 2.0 configuration interface	Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)

Table 298. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_DDR4_EFFICIENCY_MONITOR)



Table 299. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

6.1.9 Intel Stratix 10 EMIF IP DDR4 Parameters: Example Designs

Table 300. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the <i>creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization</i> . After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_DDR4_SEL_DESIGN)

Table 301. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created</i> . Instead, the output directory will contain the ed_sim.qsys file which holds Qsys details of the simulation example design, and a make_sim_design.tcl file with other corresponding tcl files. You can run make_sim_design.tcl from a command line to generate the simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_DDR4_GEN_SIM)
Synthesis	Specifies that the 'Generate Example Design' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created</i> . Instead, the output directory will contain the ed_synth.qsys file which holds Qsys details of the synthesis example design, and a make_qii_design.tcl script with other corresponding tcl files. You can run make_qii_design.tcl from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_DDR4_GEN_SYNTH)

Table 302. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_DDR4_HDL_FORMAT)



Table 303. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select 'none' from the 'Select board' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_DDR4_TARGET_DEV_KIT)

6.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

6.2.1 Equations for DDR4 Board Skew Parameters

Table 304. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum CK delay to DIMM/device	<p>The delay of the longest CK trace from the FPGA to any DIMM/device.</p> $\max_r \left[\max_n (CK_{n_r} PathDelay) \right]$ <p>Where n is the number of memory clock and r is the number rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 pairs of memory clocks in each rank DIMM, the maximum CK delay is expressed by the following equation:</p> $\max(CK_1 PathDelayrank1, CK_2 PathDelayrank1, CK_1 PathDelayrank2, CK_2 PathDelayrank2)$
Maximum DQS delay to DIMM/device	<p>The delay of the longest DQS trace from the FPGA to the DIMM/device.</p> $\max_r \left[\max_n (DQS_{n_r} PathDelay) \right]$ <p>Where n is the number of DQS and r is the number of rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 DQS in each rank DIMM, the maximum DQS delay is expressed by the following equation:</p> $\max(DQS_1 PathDelayrank1, DQS_2 PathDelayrank1, DQS_1 PathDelayrank2, DQS_2 PathDelayrank2)$
Average delay difference between DQS and CK	<p>The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS delay minus the CK delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the DQS signals for appropriate setup and hold margins.</p> $\max_r \left[\frac{\max_{n,m} \left\{ (DQS_{m_r} Delay - CK_{n_r} Delay) \right\} + \min_{n,m} \left\{ (DQS_{m_r} Delay - CK_{n_r} Delay) \right\}}{2} \right]$ <p>Where n is the number of memory clock, m is the number of DQS, and r is the number of rank of DIMM/device.</p>

continued...



Parameter	Description/Equation
	<p>When using discrete components, the calculation differs slightly. Find the minimum and maximum values for (DQS-CK) over all groups and then divide by 2. Calculate the (DQS-CK) for each DQS group, by using the appropriate CLK for that group.</p> <p>For example, in a configuration with 5 x16 components, with each component having two DQS groups: To find the minimum and maximum, calculate the minimum and maximum of (DQS0 – CK0, DQS1 – CK0, DQS2 – CK1, DQS3 – CK1, and so forth) and then divide the result by 2.</p>
Maximum Board skew within DQS group	<p>The largest skew between all DQ and DM pins in a DQS group. Enter your board skew only. Package skew is calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins.</p> $\left[\begin{matrix} groups \\ Max_g \end{matrix} \left[\max DQ_g - \min DQ_g \right] \right]$
Maximum skew between DQS groups	<p>The largest skew between DQS signals in different DQS groups.</p> $\left[\begin{matrix} groups \\ Max_g \end{matrix} \left[DQS_g \right] \right] - \left[\begin{matrix} groups \\ Min_g \end{matrix} \left[DQS_g \right] \right]$
Maximum system skew within address/command bus	<p>(MaxAC – MinAC)</p> <p>The largest skew between the address and command signals. Enter combined board and package skew. In the case of a component, find the maximum address/command and minimum address/command values across all component address signals.</p>
Average delay difference between address/command and CK	<p>A value equal to the average of the longest and smallest address/command signal delays, minus the delay of the CK signal. The value can be positive or negative.</p> <p>The average delay difference between the address/command and CK is expressed by the following equation:</p> $\frac{\sum_{n=1}^n \left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - CK_n PathDelay}{n}$ <p>where n is the number of memory clocks.</p>
Maximum delay difference between DIMMs/devices	<p>The largest propagation delay on DQ signals between ranks. For example, in a two-rank configuration where you place DIMMs in different slots there is also a propagation delay for DQ signals going to and coming back from the furthest DIMM compared to the nearest DIMM. This parameter is applicable only when there is more than one rank. Max_r { max_{n,m} [(DQ_{n_r} path delay– DQ_{n_r+1} path delay), (DQ_{m_r} path delay– DQ_{m_r+1} path delay)] }</p> <p>Where n is the number of DQ, m is the number of DQS and r is number of rank of DIMM/device .</p>

6.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.

6.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.

Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

6.3.1.1 Estimating Pin Requirements


You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:

1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.

You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)



6.3.1.2 DIMM Options

Unbuffered DIMMs (UDIMMs) require one set of chip-select (CS#), on-die termination (ODT), clock-enable (CKE), and clock pair (CK/CKn) for every physical rank on the DIMM. Registered DIMMs use only one pair of clocks. DDR3 registered DIMMs require a minimum of two chip-select signals, while DDR4 requires only one.

Compared to the unbuffered DIMMs (UDIMM), registered and load-reduced DIMMs (RDIMMs and LRDIMMs, respectively) use at least two chip-select signals CS#[1:0] in DDR3 and DDR4. Both RDIMMs and LRDIMMs require an additional parity signal for address, RAS#, CAS#, and WE# signals. A parity error signal is asserted by the module whenever a parity error is detected.

LRDIMMs expand on the operation of RDIMMs by buffering the DQ/DQS bus. Only one electrical load is presented to the controller regardless of the number of ranks, therefore only one clock enable (CKE) and ODT signal are required for LRDIMMs, regardless of the number of physical ranks. Because the number of physical ranks may exceed the number of physical chip-select signals, DDR3 LRDIMMs provide a feature known as rank multiplication, which aggregates two or four physical ranks into one larger logical rank. Refer to LRDIMM buffer documentation for details on rank multiplication.

The following table shows UDIMM and RDIMM pin options for DDR3.

Table 305. UDIMM and RDIMM Pin Options for DDR3

Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)
Data	72 bit DQ[71:0] = {CB[7:0], DQ[63:0]}	72 bit DQ[71:0] = {CB[7:0], DQ[63:0]}	72 bit DQ[71:0] = {CB[7:0], DQ[63:0]}	72 bit DQ[71:0] = {CB[7:0], DQ[63:0]}
Data Mask	DM[8:0]	DM[8:0]	DM[8:0]	DM[8:0]
Data Strobe	DQS[8:0] and DQS#[8:0]	DQS[8:0] and DQS#[8:0]	DQS[8:0] and DQS#[8:0]	DQS[8:0] and DQS#[8:0]
Address	BA[2:0], A[15:0]– 2 GB: A[13:0] 4 GB: A[14:0] 8 GB: A[15:0]	BA[2:0], A[15:0]– 2 GB: A[13:0] 4 GB: A[14:0] 8 GB: A[15:0]	BA[2:0], A[15:0]– 2 GB: A[13:0] 4 GB: A[14:0] 8 GB: A[15:0]	BA[2:0], A[15:0]– 2 GB: A[13:0] 4 GB: A[14:0] 8 GB: A[15:0]
Clock	CK0/CK0#	CK0/CK0#, CK1/CK1#	CK0/CK0#	CK0/CK0#
Command	ODT, CS#, CKE, RAS#, CAS#, WE#	ODT[1:0], CS#[1:0], CKE[1:0], RAS#, CAS#, WE#	ODT, CS#[1:0], CKE, RAS#, CAS#, WE# ²	ODT[1:0], CS#[1:0], CKE[1:0], RAS#, CAS#, WE#
Parity	—	—	PAR, ALERT	PAR, ALERT
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#

The following table shows LRDIMM pin options for DDR3.




Table 306. LRDIMM Pin Options for DDR3

Pins	LRDIMM Pins (x4, 2R)	LRDIMM (x4, 4R, RMF=1) ³	LRDIMM Pins (x4, 4R, RMF=2)	LRDIMM Pins (x4, 8R, RMF=2)	LRDIMM Pins (x4, 8R, RMF=4)	LRDIMM (x8, 4R, RMF=1) ³	LRDIMM Pins (x8, 4R, RMF=2)
Data	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}	72 bit DQ [71:0]= {CB [7:0], DQ [63:0]}
Data Mask	—	—	—	—	—	DM[8:0]	DM[8:0]
Data Strobe	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[17:0] and DQS#[17:0]	DQS[8:0] and DQS#[8:0]	DQS[8:0] and DQS#[8:0]
Address	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]	BA[2:0], A[17:0] -16GB:A[15:0] 32GB:A[16:0] 64GB:A[17:0]	BA[2:0], A[15:0] -2GB:A[13:0] 4GB:A[14:0] 8GB:A[15:0]	BA[2:0], A[16:0] -4GB:A[14:0] 8GB:A[15:0] 16GB:A[16:0]
Clock	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#
Command	ODT, CS[1:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[2:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[3:0]#, CKE, RAS#, CAS#, WE#	ODT, CS[2:0]#, CKE, RAS#, CAS#, WE#
Parity	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT	PAR, ALERT
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#
Notes to Table: 1. DM pins are not used for LRDIMMs that are constructed using x4 components. 2. S#[2] is treated as A[16] (whose corresponding pins are labeled as CS#[2] or RM[0]) and S#[3] is treated as A[17] (whose corresponding pins are labeled as CS#[3] or RM[1]) for certain rank multiplication configuration. 3. R = rank, RMF = rank multiplication factor.							

The following table shows UDIMM, RDIMM, and LRDIMM pin options for DDR4.

Table 307. UDIMM, RDIMM, and LRDIMM Pin Options for DDR4

Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)	LRDIMM Pins (Dual Rank)	LRDIMM Pins (Quad Rank)
Data	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}	72 bit DQ[71:0]= {CB[7:0], DQ[63:0]}
Data Mask	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	DM#/DBI#[8:0] ⁽¹⁾	—	—
continued...						



Pins	UDIMM Pins (Single Rank)	UDIMM Pins (Dual Rank)	RDIMM Pins (Single Rank)	RDIMM Pins (Dual Rank)	LRDIMM Pins (Dual Rank)	LRDIMM Pins (Quad Rank)
Data Strobe	x8: DQS[8:0] and DQS#[8:0]	x8: DQS[8:0] and DQS#[8:0]	x8: DQS[8:0] and DQS#[8:0] x4: DQS[17:0] and DQS#[17:0]	x8: DQS[8:0] and DQS#[8:0] x4: DQS[17:0] and DQS#[17:0]	x4: DQS[17:0] and DQS#[17:0]	x4: DQS[17:0] and DQS#[17:0]
Address	BA[1:0], BG[1:0], A[16:0] - 4GB: A[14:0] 8GB: A[15:0] 16GB: A[16:0] (2)	BA[1:0], BG[1:0], A[16:0] - 8GB: A[14:0] 16GB: A[15:0] 32GB: A[16:0] (2)	BA[1:0], BG[1:0], x8: A[16:0] - 4GB: A[14:0] 8GB: A[15:0] 16GB: A[16:0] (2) 32GB: A[17:0] (3)	BA[1:0], BG[1:0],x8: A[16:0] x4: A[17:0] - 8GB: A[14:0] 16GB: A[15:0] 32GB: A[16:0] (2) 64GB: A[17:0] (3)	BA[1:0], BG[1:0], A[17:0] - 16GB: A[15:0] 32GB: A[16:0] (2) 64GB: A[17:0] (3)	BA[1:0], BG[1:0], A[17:0] - 32GB: A[15:0] 64GB: A[16:0] (2) 128GB: A[17:0] (3)
Clock	CK0/CK0#	CK0/CK0#, CK1/CK1#	CK0/CK0#	CK0/CK0#	CK0/CK0#	CK0/CK0#
Command	ODT, CS#, CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT[1:0], CS#[1:0], CKE[1:0], ACT#, RAS#/ A16, CAS#/ A15, WE#/A14	ODT, CS#, CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT[1:0], CS#[1:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT, CS#[1:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14	ODT, CS#[3:0], CKE, ACT#, RAS#/A16, CAS#/A15, WE#/A14
Parity	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#	PAR, ALERT#
Other Pins	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#	SA[2:0], SDA, SCL, EVENT#, RESET#
Notes to Table: 1. DM/DBI pins are available only for DIMMs constructed using x8 or greater components. 2. This density requires 4Gb x4 or 2Gb x8 DRAM components. 3. This density requires 8Gb x4 DRAM components. 4. This table assumes a single slot configuration. The Intel Stratix 10 memory controller can support up to 4 ranks per channel. A single slot interface may have up to 4 ranks, and a dual slot interface may have up to 2 ranks per slot. In either case, the total number of ranks, calculated as the number of slots multiplied by the number of ranks per slot, must be less than or equal to 4.						

6.3.1.3 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.

Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

6.3.2 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.

6.3.2.1 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

6.3.2.2 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:



- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.

6.3.3 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.

The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the `Bank Number` and `Index within I/O Bank` values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the `Bank Number` value identifies the I/O column, while the letter represents the I/O bank.
- The `Index within I/O Bank` value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its `Index within I/O Bank` number (if it is an even number), or by subtracting one from its `Index within I/O Bank` number (if it is an odd number).

For example, a physical pin with a `Bank Number` of 2M and `Index within I/O Bank` of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an `Index within I/O Bank` of 23 and `Bank Number` of 2M.

6.3.3.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.



If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version number/<synth/sim>/<variation_name>_altera_emif_arch_nd_version number_<unique ID>_readme.txt` file, w+t+ed with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (.qip), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLDRAMx, QDRx and LPDDR3 are not supported with HPS.

Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the `<variation_name>/altera_emif_arch_nd_<version>/<synth/sim>/<variation_name>_altera_emif_arch_nd_<version>_<unique ID>_readme.txt` file after you have generated your IP.



7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.

Note: a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.

b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).
10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:



- There must be an even number of x4 groups in an external memory interface.
- DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group $X+1$ must be in the same I/O lane, where X is an even number.
- When placing DQ pins in x4 mode, it is important to stay within a nibble when swapping pin locations. In other words, you may swap DQ pins within a given DQS group, but you may not swap pins across DQS groups. The following table illustrates an example, where DATA_A and DATA_B are swap groups, meaning that any pin in that index can move within that range of pins.

Index Within Lane	DQS x4 Locations
11	DATA_B[3:0]
10	DATA_B[3:0]
9	DQS_Bn
8	DQS_Bp
7	DATA_B[3:0]
6	DATA_B[3:0]
5	DQS_An
4	DQS_Ap
3	DATA_A[3:0]
2	DATA_A[3:0]
1	DATA_A[3:0]
0	DATA_A[3:0]


11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Extreme, and RLD RAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLD RAM 3, x36 device, $DQ[8:0]$ and $DQ[26:18]$ are referenced to $DK0/DK0\#$, and $DQ[17:9]$ and $DQ[35:27]$ are referenced to $DK1/DK1\#$.

12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.

13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.



Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLDRAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (`global_reset_n`) for each individual interface all come from the same input pin or signal.

I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.

CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)

The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the R_{ZQ} pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.



6.3.3.2 Command and Address Signals

Command and address signals in SDRAM devices are clocked into the memory device using the CK or CK# signal. These pins operate at single data rate (SDR) using only one clock edge. The number of address pins depends on the SDRAM device capacity. The address pins are multiplexed, so two clock cycles are required to send the row, column, and bank address.

For DDR3, the CS#, RAS#, CAS#, WE#, CKE, and ODT pins are SDRAM command and control pins. For DDR3 SDRAM, certain topologies such as RDIMM and LRDIMM include RESET#, PAR (1.5V LVC MOS I/O standard), and ALERT# (SSTL-15 I/O standard).

Although DDR4 operates in fundamentally the same way as other SDRAM, there are no longer dedicated pins for RAS#, CAS#, and WE#, as those are now shared with higher-order address pins. DDR4 still has CS#, CKE, ODT, and RESET# pins, similar to DDR3. DDR4 introduces some additional pins, including the ACT# (activate) pin and BG (bank group) pins. Depending on the memory format and the functions enabled, the following pins might also exist in DDR4: PAR (address command parity) pin and the ALERT# pin.

6.3.3.3 Clock Signals

DDR3 and DDR4 SDRAM devices use CK and CK# signals to clock the address and command signals into the memory. Furthermore, the memory uses these clock signals to generate the DQS signal during a read through the DLL inside the memory. The SDRAM data sheet specifies the following timings:

- t_{DQSCK} is the skew between the CK or CK# signals and the SDRAM-generated DQS signal
- t_{DSH} is the DQS falling edge from CK rising edge hold time
- t_{DSS} is the DQS falling edge from CK rising edge setup time
- t_{DQSS} is the positive DQS latching edge to CK rising edge

SDRAM have a write requirement (t_{DQSS}) that states the positive edge of the DQS signal on writes must be within $\pm 25\%$ ($\pm 90^\circ$) of the positive edge of the SDRAM clock input. Therefore, you should generate the CK and CK# signals using the DDR registers in the IOE to match with the DQS signal and reduce any variations across process, voltage, and temperature. The positive edge of the SDRAM clock, CK, is aligned with the DQS write to satisfy t_{DQSS} .

DDR3 SDRAM can use a daisy-chained control address command (CAC) topology, in which the memory clock must arrive at each chip at a different time. To compensate for the flight-time skew between devices when using the CAC topology, you should employ write leveling.

6.3.3.4 Data, Data Strobes, DM/DBI, and Optional ECC Signals


DDR3 and DDR4 SDRAM use bidirectional differential data strobes. Differential DQS operation enables improved system timing due to reduced crosstalk and less simultaneous switching noise on the strobe output drivers. The DQ pins are also bidirectional.

DQ pins in DDR3 and DDR4 SDRAM interfaces can operate in either $\times 4$ or $\times 8$ mode DQS groups, depending on your chosen memory device or DIMM, regardless of interface width. The $\times 4$ and $\times 8$ configurations use one pair of bidirectional data strobe

signals, DQS and DQSn, to capture input data. However, two pairs of data strobes, UDQS and UDQS# (upper byte) and LDQS and LDQS# (lower byte), are required by the $\times 16$ configuration devices. A group of DQ pins must remain associated with its respective DQS and DQSn pins.


The DQ signals are edge-aligned with the DQS signal during a read from the memory and are center-aligned with the DQS signal during a write to the memory. The memory controller shifts the DQ signals by -90 degrees during a write operation to center align the DQ and DQS signals. The PHY IP delays the DQS signal during a read, so that the DQ and DQS signals are center aligned at the capture register. Intel devices use a phase-locked loop (PLL) to center-align the DQS signal with respect to the DQ signals during writes and Intel devices use dedicated DQS phase-shift circuitry to shift the incoming DQS signal during reads. The following figure shows an example where the DQS signal is shifted by 90 degrees for a read from the DDR3 SDRAM.

Figure 56. Edge-aligned DQ and DQS Relationship During a DDR3 SDRAM Read in Burst-of-Four Mode



The following figure shows an example of the relationship between the data and data strobe during a burst-of-four write.

Figure 57. DQ and DQS Relationship During a DDR3 SDRAM Write in Burst-of-Four Mode



The memory device's setup (t_{DS}) and hold times (t_{DH}) for the DQ and DM pins during writes are relative to the edges of DQS write signals and not the CK or CK# clock. Setup and hold requirements are not necessarily balanced in DDR3 SDRAM.

The DQS signal is generated on the positive edge of the system clock to meet the t_{DQSS} requirement. DQ and DM signals use a clock shifted -90 degrees from the system clock, so that the DQS edges are centered on the DQ or DM signals when they arrive at the DDR3 SDRAM. The DQS, DQ, and DM board trace lengths need to be tightly matched (within 20 ps).

The SDRAM uses the DM pins during a write operation. Driving the DM pins low shows that the write is valid. The memory masks the DQ signals if the DM pins are driven high. To generate the DM signal, Intel recommends that you use the spare DQ pin within the same DQS group as the respective data, to minimize skew.



The DM signal's timing requirements at the SDRAM input are identical to those for DQ data. The DDR registers, clocked by the -90 degree shifted clock, create the DM signals.

DDR4 supports DM similarly to other SDRAM, except that in DDR4 DM is active LOW and bidirectional, because it supports Data Bus Inversion (DBI) through the same pin. DM is multiplexed with DBI by a Mode Register setting whereby only one function can be enabled at a time. DBI is an input/output identifying whether to store/output the true or inverted data. When enabled, if DBI is LOW, during a write operation the data is inverted and stored inside the DDR4 SDRAM; during a read operation, the data is inverted and output. The data is not inverted if DBI is HIGH. For Intel Stratix 10 interfaces, the DM (for DDR3) pins in each DQS group must be paired with a DQ pin for proper operation. DM/DBI (for DDR4) do not need to be paired with a DQ pin.

Some SDRAM modules support error correction coding (ECC) to allow the controller to detect and automatically correct error in data transmission. The 72-bit SDRAM modules contain eight extra data pins in addition to 64 data pins. The eight extra ECC pins should be connected to a single DQS or DQ group on the FPGA.

6.3.4 Resource Sharing Guidelines (Multiple Interfaces)

In Intel Cyclone 10/Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.

PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.

Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `p11_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note:

You can place the `p11_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.

Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.

Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.

6.4 DDR4 Board Design Guidelines

The following topics provide guidelines for improving the signal integrity of your system and for successfully implementing a DDR4 SDRAM interface on your system.

The following areas are discussed:

- comparison of various types of termination schemes, and their effects on the signal quality on the receiver
- proper drive strength setting on the FPGA to optimize the signal integrity at the receiver
- effects of different loading types, such as components versus DIMM configuration, on signal quality

It is important to understand the trade-offs between different types of termination schemes, the effects of output drive strengths, and different loading types, so that you can swiftly navigate through the multiple combinations and choose the best possible settings for your designs.


The following key factors affect signal quality at the receiver:

- Leveling and dynamic ODT
- Proper use of termination
- Layout guidelines



As memory interface performance increases, board designers must pay closer attention to the quality of the signal seen at the receiver because poorly transmitted signals can dramatically reduce the overall data-valid margin at the receiver. The following figure shows the differences between an ideal and real signal seen by the receiver.

Figure 58. Ideal and Real Signal at the Receiver



Related Links

JEDEC.org

6.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices

The following topics describe considerations specific to DDR3 and DDR4 external memory interface protocols on Intel Stratix 10 devices.

6.4.1.1 Dynamic On-Chip Termination (OCT) in Intel Stratix 10 Devices

Depending upon the R_s (series) and R_t (parallel) OCT values that you want, you should choose appropriate values for the RZQ resistor and connect this resistor to the RZQ pin of the FPGA.

- Select a 240-ohm reference resistor to ground to implement R_s OCT values of 34-ohm, 40-ohm, 48-ohm, 60-ohm, and 80-ohm, and R_t OCT resistance values of 20-ohm, 30-ohm, 34-ohm, 40-ohm, 60-ohm, 80-ohm, 120-ohm and 240 ohm.
- Select a 100-ohm reference resistor to ground to implement R_s OCT values of 25-ohm and 50-ohm, and an R_t OCT resistance of 50-ohm.

Check the FPGA I/O tab of the parameter editor to determine the I/O standards and termination values supported for data, address and command, and memory clock signals.

Related Links

[Choosing Terminations on Intel Stratix 10 Devices](#) on page 182

6.4.1.2 Dynamic On-Die Termination (ODT) in DDR4

In DDR4, in addition to the R_{tt_nom} and R_{tt_wr} values, which are applied during read and write respectively, a third option called R_{tt_park} is available. When R_{tt_park} is enabled, a selected termination value is set in the DRAM when ODT is driven low.



Rtt_nom and Rtt_wr work the same as in DDR3, which is described in *Dynamic ODT for DDR3*.

Refer to the DDR4 JEDEC specification or your memory vendor data sheet for details about available termination values and functional description for dynamic ODT in DDR4 devices.

For DDR4 LRDIMM, if SPD byte 152 calls for different values of Rtt_Park to be used for package ranks 0 and 1 versus package ranks 2 and 3, set the value to the larger of the two impedance settings.

6.4.1.3 Choosing Terminations on Intel Stratix 10 Devices

To determine optimal on-chip termination (OCT) and on-die termination (ODT) values for best signal integrity, you should simulate your memory interface in HyperLynx or a similar tool.

If the optimal OCT and ODT termination values as determined by simulation are not available in the list of available values in the parameter editor, select the closest available termination values for OCT and ODT.

For information about available ODT choices, refer to your memory vendor data sheet.

Related Links

[Dynamic On-Chip Termination \(OCT\) in Intel Stratix 10 Devices](#) on page 182

6.4.1.4 On-Chip Termination Recommendations for Intel Stratix 10 Devices

- Output mode (drive strength) for Address/Command/Clock and Data Signals: Depending upon the I/O standard that you have selected, you would have a range of selections expressed in terms of ohms or miliamps. A value of 34 to 40 ohms or 12 mA is a good starting point for output mode drive strength.
- Input mode (parallel termination) for Data and Data Strobe signals: A value of 40 or 60 ohms is a good starting point for FPGA side input termination.

6.4.2 Channel Signal Integrity Measurement

As external memory interface data rates increase, so does the importance of proper channel signal integrity measurement. By measuring the actual channel loss during the layout process and including that data in your parameterization, a realistic assessment of margins is achieved.

6.4.2.1 Importance of Accurate Channel Signal Integrity Information

Default values for channel loss (or eye reduction) can be used when calculating timing margins, however those default values may not accurately reflect the channel loss in your system. If the channel loss in your system is different than the default values, the calculated timing margins will vary accordingly.

If your actual channel loss is greater than the default channel loss, and if you rely on default values, the available timing margins for the entire system will be lower than the values calculated during compilation. By relying on default values that do not


accurately reflect your system, you may be lead to believe that you have good timing margin, while in reality, your design may require changes to achieve good channel signal integrity.

6.4.2.2 Understanding Channel Signal Integrity Measurement

To measure channel signal integrity you need to measure the channel loss for various signals. For a particular signal or signal trace, channel loss is defined as loss of the eye width at $\pm V_{IH}(\text{ac and dc}) \pm V_{IL}(\text{ac and dc})$. V_{IH}/V_{IL} above or below V_{REF} is used to align with various requirements of the timing model for memory interfaces.

The example below shows a reference eye diagram where the channel loss on the setup- or leading-side of the eye is equal to the channel loss on the hold- or lagging-side of the eye; however, it does not necessarily have to be that way. Because Intel's calibrating PHY will calibrate to the center of the read and write eye, the Board Settings tab has parameters for the total extra channel loss for Write DQ and Read DQ. For address and command signals which are not-calibrated, the Board Settings tab allows you to enter setup- and hold-side channel losses that are not equal, allowing the Intel Quartus Prime software to place the clock statically within the center of the address and command eye.

Figure 59. Equal Setup and Hold-side Losses



6.4.2.3 How to Enter Calculated Channel Signal Integrity Values

You should enter calculated channel loss values in the **Channel Signal Integrity** section of the **Board** (or **Board Timing**) tab of the parameter editor.

For Intel Stratix 10 external memory interfaces, the default channel loss displayed in the parameter editor is based on the selected configuration (different values for single rank versus dual rank), and on internal Intel reference boards. You should replace the default value with the value that you calculate.

6.4.2.4 Guidelines for Calculating DDR4 Channel Signal Integrity

Address and Command ISI and Crosstalk

Simulate the address/command and control signals and capture eye at the DRAM pins, using the memory clock as the trigger for the memory interface's address/command and control signals. Measure the setup and hold channel losses at the voltage thresholds mentioned in the memory vendor's data sheet.

Address and command channel loss = Measured loss on the setup side + measured loss on the hold side.


$V_{REF} = V_{DD}/2 = 0.75 \text{ mV}$ for address/command for DDR4.

You should select the V_{IH} and V_{IL} voltage levels appropriately for the DDR4 memory device that you are using. Check with your memory vendor for the correct voltage levels, as the levels may vary for different speed grades of device.

The following figure illustrates a DDR4-1200 example, where $V_{IH(AC)}/V_{IL(AC)}$ is +/- 100 mV and $V_{IH(DC)}/V_{IL(DC)}$ is +/- 75 mV.

Select the $V_{IH(AC)}$, $V_{IL(AC)}$, $V_{IH(DC)}$, and $V_{IL(DC)}$ for the speed grade of DDR4 memory device from the memory vendor's data sheet.

Figure 60.



Write DQ ISI and Crosstalk

Simulate the write DQ signals and capture eye at the DRAM pins, using DQ Strobe (DQS) as a trigger for the DQ signals of the memory interface simulation. Measure the setup and hold channel losses at the V_{IH} and V_{IL} mentioned in the memory vendor's data sheet

Write Channel Loss = Measured Loss on the Setup side + Measured Loss on the Hold side.

or

Write Channel Loss = UI - (Eye opening at V_{IH} or V_{IL}).


V_{REF} = Voltage level where the eye opening is highest.

$$V_{IH} = V_{REF} + (0.5 \times V_{diVW}).$$

$$V_{IL} = V_{REF} - (0.5 \times V_{diVW}).$$

Where V_{diVW} varies by frequency of operation; you can find the V_{diVW} value in your memory vendor's data sheet.

Figure 61.



Read DQ ISI and Crosstalk


Simulate read DQ signals and capture eye at the FPGA die. Do not measure at the pin, because you might see unwanted reflections that could create a false representation of the eye opening at the input buffer of the FPGA. Use DQ Strobe (DQS) as a trigger for the DQ signals of your memory interface simulation. Measure the eye opening at +/- 70 mV (V_{IH}/V_{IL}) with respect to V_{REF} .

Read Channel Loss = (UI) - (Eye opening at +/- 70 mV with respect to V_{REF} .)

UI = Unit interval. For example, if you are running your interface at 800 Mhz, the effective data is 1600 Mbps, giving a unit interval of $1/1600 = 625$ ps.

V_{REF} = Voltage level where the eye opening is highest.

Figure 62.




Write/Read DQS ISI and Crosstalk

Simulate write and read DQS and capture eye. Measure the uncertainty at V_{REF} .

V_{REF} = Voltage level where the eye opening is the highest.

Figure 63.



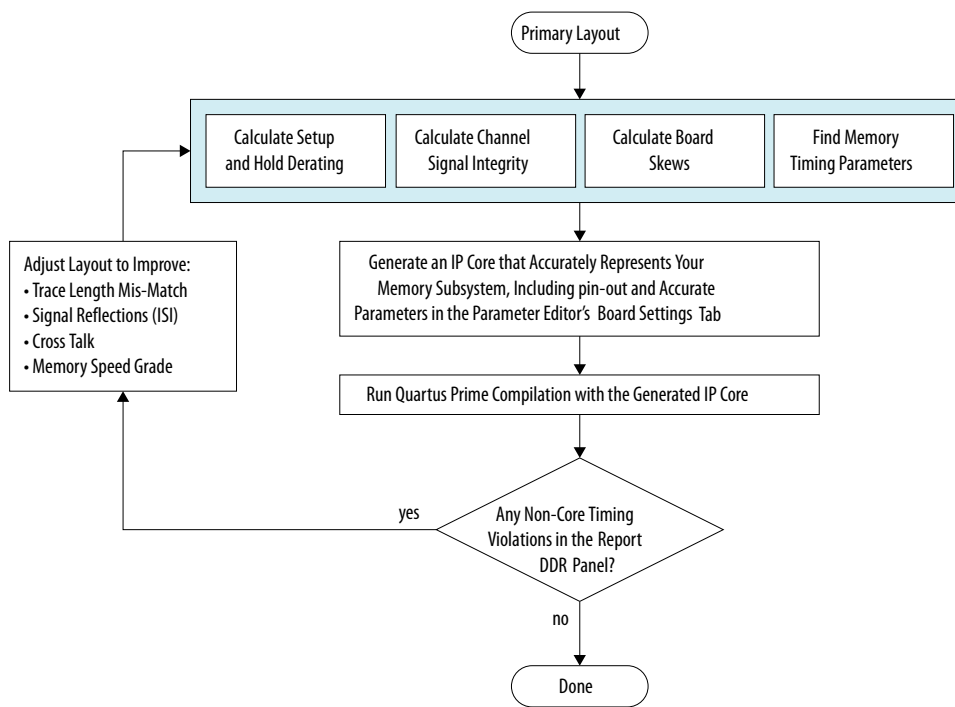
6.4.3 Layout Approach

For all practical purposes, you can regard the Timing Analyzer report on your memory interface as definitive for a given set of memory and board timing parameters.

You will find timing under **Report DDR** in the Timing Analyzer and on the **Timing Analysis** tab in the parameter editor.



The following flowchart illustrates the recommended process to follow during the board design phase, to determine timing margin and make iterative improvements to your design.



Board Skew

For information on calculating board skew parameters, refer to *Board Skew Equations*, in this chapter.

The Board Skew Parameter Tool is an interactive tool that can help you calculate board skew parameters if you know the absolute delay values for all the memory related traces.

Memory Timing Parameters

For information on the memory timing parameters to be entered into the parameter editor, refer to the datasheet for your external memory device.

Related Links

[Board Skew Parameter Tool](#)

6.4.4 Design Layout Guidelines

The general layout guidelines in the following topic apply to DDR3 and DDR4 SDRAM interfaces.

These guidelines will help you plan your board layout, but are not meant as strict rules that must be adhered to. Intel recommends that you perform your own board-level simulations to ensure that the layout you choose for your board allows you to achieve your desired performance.



For more information about how the memory manufacturers route these address and control signals on their DIMMs, refer to the Cadence PCB browser from the Cadence website, at www.cadence.com. The various JEDEC example DIMM layouts are available from the JEDEC website, at www.jedec.org.

For assistance in calculating board skew parameters, refer to the board skew calculator tool, which is available at the Intel website.

Note:

1. The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristic of the interface. They do not include any margin for crosstalk.
2. To ensure reliable timing closure to and from the periphery of the device, signals to and from the periphery should be registered before any further logic is connected.

Intel recommends that you get accurate time base skew numbers for your design when you simulate the specific implementation.

Related Links

- [Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface \(>72 bits\)](#) on page 195
- JEDEC.org
- <https://www.cadence.com/>
- [Board Skew Parameter Tool](#)
- <https://www.mentor.com/>

6.4.4.1 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.




Table 308. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)

6.4.4.2 Layout Guidelines

The following table lists layout guidelines.

Unless otherwise specified, the guidelines in the following table apply to the following topologies:

- DIMM—UDIMM topology
- DIMM—RDIMM topology
- DIMM—LRDIMM topology
- Not all versions of the Intel Quartus Prime software support LRDIMM.
- Discrete components laid out in UDIMM topology
- Discrete components laid out in RDIMM topology

These guidelines are recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface.



For supported frequencies and topologies, refer to the *External Memory Interface Spec Estimator* <http://www.altera.com/technology/memory/estimator/mem-emif-index.html>.

For frequencies greater than 800 MHz, when you are calculating the delay associated with a trace, you must take the FPGA package delays into consideration.

Table 309. Layout Guidelines ⁽¹⁾

Parameter	Guidelines
Decoupling Parameter	<ul style="list-style-type: none">• Make VTT voltage decoupling close to the components and pull-up resistors.• Connect decoupling caps between VTT and VDD using a 0.1μF cap for every other VTT pin.• Use a 0.1 μF cap and 0.01 μF cap for every VDDQ pin.
Maximum Trace Length	<ul style="list-style-type: none">• Even though there are no hard requirements for minimum trace length, you need to simulate the trace to ensure the signal integrity. Shorter routes result in better timing.• For DIMM topology only:<ul style="list-style-type: none">• Maximum trace length for all signals from FPGA to the first DIMM slot is 4.5 inches.• Maximum trace length for all signals from DIMM slot to DIMM slot is 0.425 inches.• For discrete components only:<ul style="list-style-type: none">• Maximum trace length for address, command, control, and clock from FPGA to the first component must not be more than 7 inches.• Maximum trace length for DQ, DQS, DQS#, and DM from FPGA to the first component is 5 inches.
General Routing	<ul style="list-style-type: none">• Route over appropriate VCC and GND planes.• Keep signal routing layers close to GND and power planes.
Spacing Guidelines	<ul style="list-style-type: none">• Avoid routing two signal layers next to each other. Always make sure that the signals related to memory interface are routed between appropriate GND or power layers.• For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) for these traces. (Where H is the vertical distance to the closest return path for that particular trace.)• For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) these traces. (Where H is the vertical distance to the closest return path for that particular trace.)• For Clock traces: Maintain at least 5H spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)
Clock Routing	<ul style="list-style-type: none">• Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm).• Route clock signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not exceed 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specification.• These signals should maintain the following spacings:<ul style="list-style-type: none">• Clocks should maintain a length-matching between clock pairs of ± 5 ps.• Clocks should maintain a length-matching between positive (p) and negative (n) signals of ± 2 ps, routed in parallel.• Space between different pairs should be at least two times the trace width of the differential pair to minimize loss and maximize interconnect density.• To avoid mismatched transmission line to via, Intel recommends that you use Ground Signal Signal Ground (GSSG) topology for your clock pattern—GND CLKP CKLN GND.• Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
continued...	



Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> Route address and command signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not be more than 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specifications. UDIMMs are more susceptible to cross-talk and are generally noisier than buffered DIMMs. Therefore, route address and command signals of UDIMMs on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. Do not route differential clock (CK) and clock enable (CKE) signals close to address signals. Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
DQ, DM, and DQS Routing Rules	<ul style="list-style-type: none"> All the trace length matching requirements are from the FPGA package ball to the SDRAM package ball, which means you must consider trace mismatching on different DIMM raw cards. Match in length all DQ, DQS, and DM signals within a given byte-lane group with a maximum deviation of ± 10 ps. Ensure to route all DQ, DQS, and DM signals within a given byte-lane group on the same layer to avoid layer to layer transmission velocity differences, which otherwise increase the skew within the group. Do not count on FPGAs to deskew for more than 20 ps of DQ group skew. The skew algorithm only removes the following possible uncertainties: <ul style="list-style-type: none"> Minimum and maximum die IOE skew or delay mismatch Minimum and maximum device package skew or mismatch Board delay mismatch of 20 ps Memory component DQ skew mismatch Increasing any of these four parameters runs the risk of the deskew algorithm limiting, failing to correct for the total observed system skew. If the algorithm cannot compensate without limiting the correction, timing analysis shows reduced margins. For memory interfaces with leveling, the timing between the DQS and clock signals on each device calibrates dynamically to meet tDQSS. To make sure the skew is not too large for the leveling circuit's capability, follow these rules: <ul style="list-style-type: none"> Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device: $(CK_i) - DQSi > 0$; $0 < i < \text{number of components} - 1$. For DIMMs, ensure that the CK trace is longer than the longest DQS trace at the DIMM connector. Total skew of CLK and DQS signal between groups is less than one clock cycle: $(CK_i + DQSi)_{\text{max}} - (CK_i + DQSi)_{\text{min}} < 1 \times tCK$ (If you are using a DIMM topology, your delay and skew must take into consideration values for the actual DIMM.)
continued...	

Parameter	Guidelines
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always ensure that the signals related to the memory interface are routed between appropriate GND or power layers. For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Clock traces: Maintain at least 5H spacing between two clock pairs or a clock pair and any other memory interface trace, where H is the vertical distance to the closest return path for that particular trace.
Intel Quartus Prime Software Settings for Board Layout	<ul style="list-style-type: none"> To perform timing analyses on board and I/O buffers, use a third-party simulation tool to simulate all timing information such as skew, ISI, crosstalk, and type the simulation result into the Board Settings tab in the parameter editor. Do not use advanced I/O timing model (AIOT) or board trace model unless you do not have access to any third party tool. AIOT provides reasonable accuracy but tools like HyperLynx provide better results.
Notes to Table: 1. For point-to-point and DIMM interface designs, refer to the Micron website, www.micron.com .	

Related Links

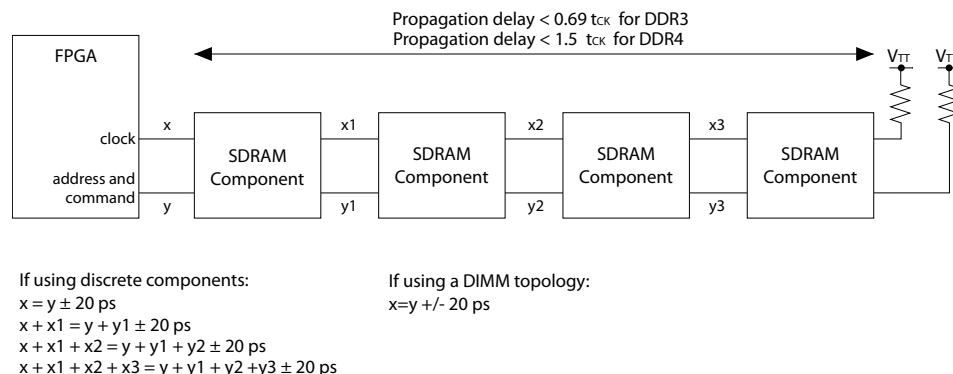
- [Package Deskew](#) on page 198
- [External Memory Interface Spec Estimator](#)
- www.micron.com

6.4.4.3 Length Matching Rules

The following topics provide guidance on length matching for different types of SDRAM signals.

Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. The following figure shows the component routing guidelines for address and command signals.

Figure 64. SDRAM Component Address and Command Routing Guidelines





The timing between the DQS and clock signals on each device calibrates dynamically to meet t_{DQSS} . The following figure shows the delay requirements to align DQS and clock signals. To ensure that the skew is not too large for the leveling circuit's capability, follow these rules:

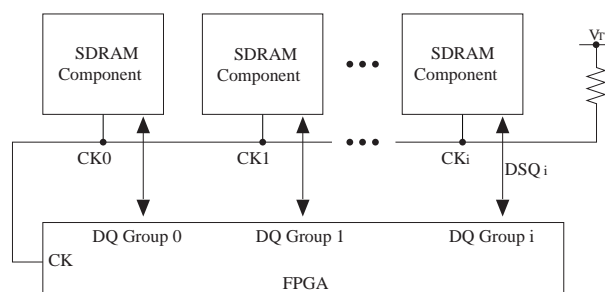
- Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device:

$$CK_i - DQSi > 0; 0 < i < \text{number of components} - 1$$

- Total skew of CLK and DQS signal between groups is less than one clock cycle:

$$(CK_i + DQSi)_{\max} - (CK_i + DQSi)_{\min} < 1 \times t_{CK}$$

Figure 65. Delaying DQS Signal to Align DQS and Clock



CK_i = Clock signal propagation delay to device i

$DQSi$ = DQ/DQS signals propagation delay to group i

Clk pair matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, match the trace lengths up to the DIMM connector. If you are using discrete components, match the lengths for all the memory components connected in the fly-by chain.

DQ group length matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, apply the DQ group trace matching rules described in the guideline table earlier up to the DIMM connector. If you are using discrete components, match the lengths up to the respective memory components.


When you are using DIMMs, it is assumed that lengths are tightly matched within the DIMM itself. You should check that appropriate traces are length-matched within the DIMM.

6.4.4.4 Spacing Guidelines

This topic provides recommendations for minimum spacing between board traces for various signal traces.


Spacing Guidelines for DQ, DQS, and DM Traces

Maintain a minimum of $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)




Spacing Guidelines for Address and Command and Control Traces

Maintain at least $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)



Spacing Guidelines for Clock Traces

Maintain at least $5H$ spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)



6.4.4.5 Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface (>72 bits)

The following topics discuss different ways to lay out a wider DDR3 or DDR4 SDRAM interface to the FPGA. Choose the topology based on board trace simulation and the timing budget of your system.

The EMIF IP supports up to a 144-bit wide DDR3 interface. You can use discrete components or DIMMs to implement a wide interface (any interface wider than 72 bits). Intel recommends using leveling when you implement a wide interface with DDR3 components.

When you lay out for a wider interface, all rules and constraints discussed in the previous sections still apply. The DQS, DQ, and DM signals are point-to-point, and all the same rules discussed in *Design Layout Guidelines* apply.

The main challenge for the design of the fly-by network topology for the clock, command, and address signals is to avoid signal integrity issues, and to make sure you route the DQS, DQ, and DM signals with the chosen topology.

Related Links

[Design Layout Guidelines](#) on page 188

6.4.4.6 Fly-By Network Design for Clock, Command, and Address Signals


The EMIF IP requires the flight-time skew between the first SDRAM component and the last SDRAM component to be less than 0.69 tCK for memory clocks. This constraint limits the number of components you can have for each fly-by network.

If you design with discrete components, you can choose to use one or more fly-by networks for the clock, command, and address signals.

The following figure shows an example of a single fly-by network topology.



Figure 66. Single Fly-By Network Topology

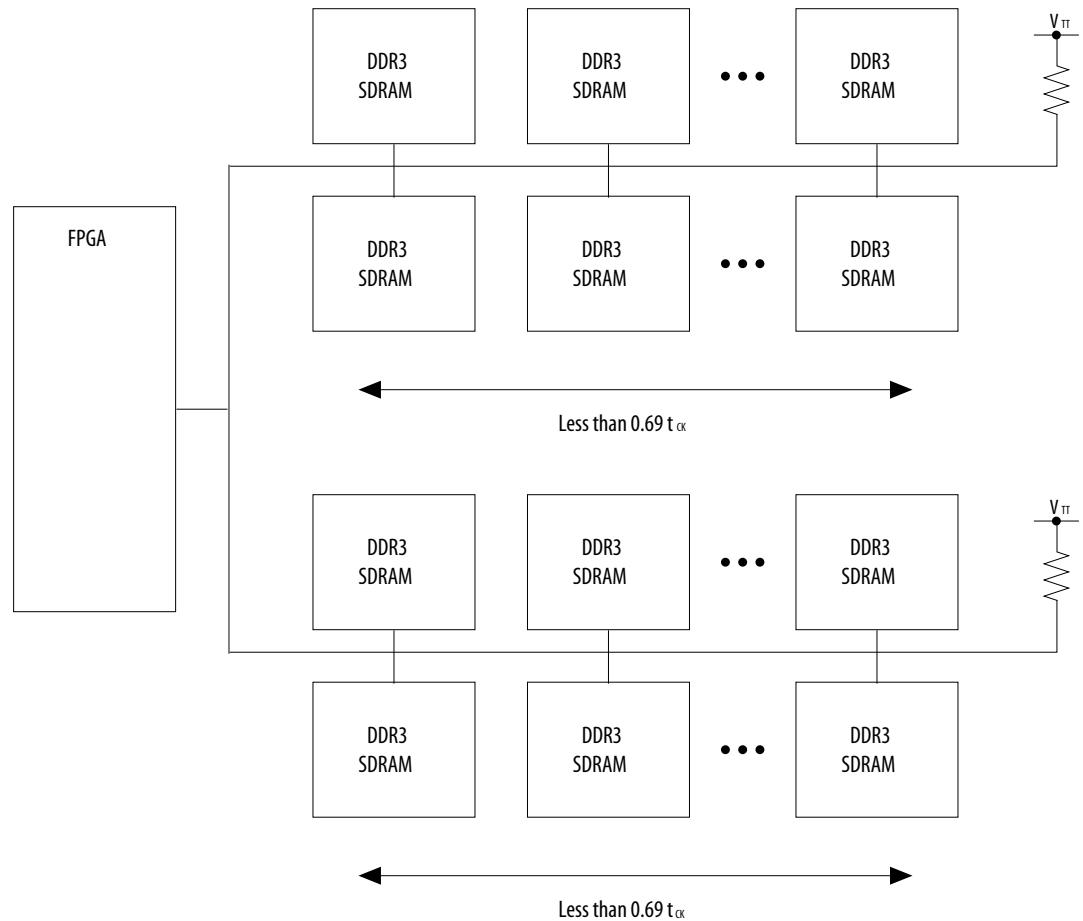


Every SDRAM component connected to the signal is a small load that causes discontinuity and degrades the signal. When using a single fly-by network topology, to minimize signal distortion, follow these guidelines:

- Use $\times 16$ device instead $\times 4$ or $\times 8$ to minimize the number of devices connected to the trace.
- Keep the stubs as short as possible.
- Even with added loads from additional components, keep the total trace length short; keep the distance between the FPGA and the first SDRAM component less than 5 inches.
- Simulate clock signals to ensure a decent waveform.

The following figure shows an example of a double fly-by network topology. This topology is not rigid but you can use it as an alternative option. The advantage of using this topology is that you can have more SDRAM components in a system without violating the 0.69 tCK rule. However, as the signals branch out, the components still create discontinuity.


Figure 67. Double Fly-By Network Topology



You must perform simulations to find the location of the split, and the best impedance for the traces before and after the split.

The following figure shows a way to minimize the discontinuity effect. In this example, keep TL2 and TL3 matches in length. Keep TL1 longer than TL2 and TL3, so that it is easier to route all the signals during layout.

Figure 68. Minimizing Discontinuity Effect



You can also consider using a DIMM on each branch to replace the components. Because the trace impedance on the DIMM card is 40-ohm to 60-ohm, perform a board trace simulation to control the reflection to within the level your system can tolerate.

Using the fly-by daisy chain topology increases the complexity of the datapath and controller design to achieve leveling, but also greatly improves performance and eases board layout for SDRAM implementations.

You can also use the SDRAM components without leveling in a design if it may result in a more optimal solution, or use with devices that support the required electrical interface standard, but do not support the required read and write leveling functionality.

6.4.5 Package Deskew

Trace lengths inside the device package are not uniform for all package pins. The nonuniformity of package traces can affect system timing for high frequencies. A package deskew option is available in the Intel Quartus Prime software.

If you do not enable the package deskew option, the Intel Quartus Prime software uses the package delay numbers to adjust skews on the appropriate signals; you do not need to adjust for package delays on the board traces. If you do enable the package deskew option, the Intel Quartus Prime software does not use the package delay numbers for timing analysis, and you must deskew the package delays with the board traces for the appropriate signals for your design.

Related Links

[Layout Guidelines](#) on page 190

6.4.5.1 DQ/DQS/DM Deskew

To get the package delay information, follow these steps:



1. Select the **FPGA DQ/DQS Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the <core_name>.pin file.

6.4.5.2 Address and Command Deskew

Deskew address and command delays as follows:

1. Select the **FPGA Address/Command Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the <core_name>.pin file.

6.4.5.3 Package Deskew Recommendations for Intel Stratix 10 Devices

The following table shows package deskew recommendations for Intel Stratix 10 devices.


As operating frequencies increase, it becomes increasingly critical to perform package deskew. The frequencies listed in the table are the *minimum* frequencies for which you must perform package deskew.

If you plan to use a listed protocol at the specified frequency or higher, you must perform package deskew.

Protocol	Minimum Frequency (MHz) for Which to Perform Package Deskew		
	Single Rank	Dual Rank	Quad Rank
DDR4	933	800	667
DDR3	933	800	667
LPDDR3	667	533	Not required
QDR IV	933	Not applicable	Not applicable
RLDRAM 3	933	667	Not applicable
QDR II, II+, II+ Xtreme	Not required	Not applicable	Not applicable

6.4.5.4 Deskew Example

Consider an example where you want to deskew an interface with 4 DQ pins, 1 DQS pin, and 1 DQSn pin.

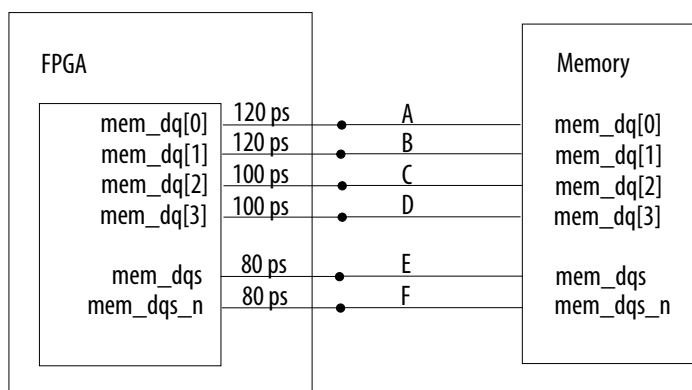


Let's assume an operating frequency of 667 MHz, and the package lengths for the pins reported in the **.pin** file as follows:

```
dq[0] = 120 ps
dq[1] = 120 ps
dq[2] = 100 ps
dq[3] = 100 ps
dqs = 80 ps
dqs_n = 80 ps
```

The following figure illustrates this example.

Figure 69. Deskew Example

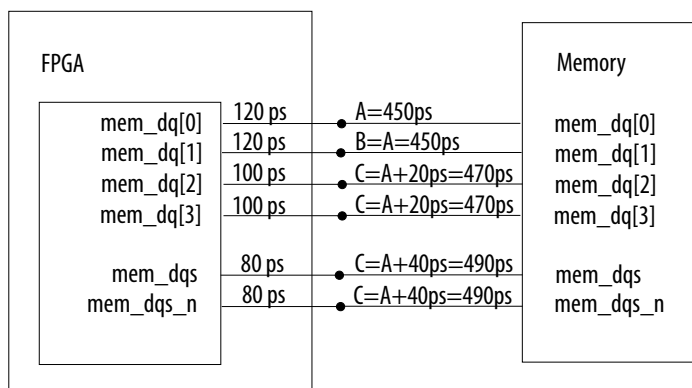


When you perform length matching for all the traces in the DQS group, you must take package delays into consideration. Because the package delays of traces A and B are 40 ps longer than the package delays of traces E and F, you would need to make the board traces for E and F 40 ps longer than the board traces for A and B.

A similar methodology would apply to traces C and D, which should be 20 ps longer than the lengths of traces A and B.

The following figure shows this scenario with the length of trace A at 450 ps.

Figure 70. Deskew Example with Trace Delay Calculations





When you enter the board skews into the Board Settings tab of the DDR3 parameter editor, you should calculate the board skew parameters as the sums of board delay and corresponding package delay. If a pin does not have a package delay (such as address and command pins), you should use the board delay only.

The example of the preceding figure shows an ideal case where board skews are perfectly matched. In reality, you should allow plus or minus 10 ps of skew mismatch within a DQS group (DQ/DQS/DM).



7 Intel Stratix 10 EMIF IP for QDR II/II+/II+ Xtreme

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for QDR II/II+/II+ Xtreme.

7.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

7.1.1 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: General

Table 310. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_QDR2_CONFIG_ENUM)

Table 311. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_QDR2_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_QDR2_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_QDR2_USER_REF_CLK_FREQ_MHZ)
<i>continued...</i>	



Display Name	Description
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_QDR2_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_QDR2_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_QDR2_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 312. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 313. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)

Table 314. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)



Table 315. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 316. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

7.1.2 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.

Table 317. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_QDR2_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_QDR2_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_QDR2_DEFAULT_IO)



Table 318. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR2_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR2_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_QDR2_USER_AC_SLEW_RATE_ENUM)

Table 319. Group: FPGA I/O / FPGA I/O Settings / Memory Clock

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR2_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR2_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_QDR2_USER_CK_SLEW_RATE_ENUM)

Table 320. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR2_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR2_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR2_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_QDR2_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_QDR2_USER_STARTING_VREFIN)



Table 321. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_QDR2_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_QDR2_USER_RZQ_IO_STD_ENUM)

7.1.3 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Memory

Table 322. Group: Memory / Topology

Display Name	Description
Data width per device	Number of D and Q pins per QDR II device. (Identifier: MEM_QDR2_DATA_PER_DEVICE)
Enable BWS# pins	Indicates whether the interface uses the BWS#(Byte Write Select) pins. If enabled, 1 BWS# pin for every 9 D pins will be added. (Identifier: MEM_QDR2_BWS_EN)
Enable width expansion	Indicates whether to combine two memory devices to double the data bus width. With two devices, the interface supports a width expansion configuration up to 72-bits. For width expansion configuration, the address and control signals are routed to 2 devices. (Identifier: MEM_QDR2_WIDTH_EXPANDED)
Address width	Number of address pins. (Identifier: MEM_QDR2_ADDR_WIDTH)
Burst length	Burst length of the memory device. (Identifier: MEM_QDR2_BL)

7.1.4 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).

Table 323. Group: Mem Timing

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_QDR2_SPEEDBIN_ENUM)
tRL	tRL refers to the QDR memory specific read latency . This parameter describes the length of time after a Read command has been registered on the rising edge of the Write Clock (K) at the QDR memory before the first piece of read data (Q) can be expected at the output of the memory. It is measured in Write Clock (K) cycles. The Read Latency is specific to a QDR memory device and cannot be modified to a different value. The Read Latency (tRL) can have the following values: 1.5, 2, 2.5 clk cycles. (Identifier: MEM_QDR2_TRL_CYC)
tSA	tSA refers to the setup time for the address and command bus (A) before the rising edge of the clock (K) . The address and command bus must be stable for at least tSA before the rising edge of K. (Identifier: MEM_QDR2_TSA_NS)
<i>continued...</i>	



Display Name	Description
tHA	tHA refers to the hold time after the rising edge of the clock (K) to the address and command control bus (A) . The address and command control bus must remain stable for at least tHA after the rising edge of K. (Identifier: MEM_QDR2_THA_NS)
tSD	tSD refers to the setup time for the data bus (D) before the rising edge of the clock (K) . The data bus must be stable for at least tSD before the rising edge of K. (Identifier: MEM_QDR2_TSD_NS)
tHD	tHD refers to the hold time after the rising edge of the clock (K) to the data bus (D) . The data bus must remain stable for at least tHD after the rising edge of K. (Identifier: MEM_QDR2_THD_NS)
tCQD	tCQD refers to the maximum time expected between an echo clock edge and valid data on the Read Data bus (Q). (Identifier: MEM_QDR2_TCQD_NS)
tCQDOH	tCQDOH refers to the minimum time expected between the echo clock (CQ or CQ#) edge and the last of the valid Read data (Q). (Identifier: MEM_QDR2_TCQDOH_NS)
Internal Jitter	QDRII internal jitter. (Identifier: MEM_QDR2_INTERNAL_JITTER_NS)
tCQH	tCQH describes the time period during which the echo clock (CQ, #CQ) is considered logically high. (Identifier: MEM_QDR2_TCQH_NS)
tCCQO	tCCQO describes the skew between the rising edge of the C clock to the rising edge of the echo clock (CQ) in QDRII memory devices. (Identifier: MEM_QDR2_TCCQO_NS)

7.1.5 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Board

Table 324. Group: Board / Intersymbol Interference/Crosstalk

Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation using IBIS models and Hyperlynx)*, and manually enter values based on your simulation results, instead of using the default values.</i> (Identifier: BOARD_QDR2_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to ISI and crosstalk effects. The number to be entered is the total loss of margin on both the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_QDR2_USER_AC_ISI_NS)
CQ/CQ# ISI/crosstalk	CQ/CQ# ISI/crosstalk describes the reduction of the read data window due to intersymbol interference and crosstalk effects on the CQ/CQ# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_QDR2_USER_RCLK_ISI_NS)
Read Q ISI/crosstalk	Read Q ISI/crosstalk describes the reduction of the read data window due to intersymbol interference and crosstalk effects on the CQ/CQ# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides
continued...	



Display Name	Description
	(measured loss on the setup side + measured loss on the hold side). Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_QDR2_USER_RDATA_ISI_NS)
K/K# ISI/crosstalk	K/K# ISI/crosstalk describes the reduction of the write data window due to intersymbol interference and crosstalk effects on the K/K# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_QDR2_USER_WCLK_ISI_NS)
Write D ISI/crosstalk	Write D ISI/crosstalk describes the reduction of the write data window due to intersymbol interference and crosstalk effects on the signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_QDR2_USER_WDATA_ISI_NS)

Table 325. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (Q group)	If you are compensating for package skew on the Q bus in the board layout (hence checking the box here), please include package skew in calculating the following board skew parameters . (Identifier: BOARD_QDR2_IS_SKEW_WITHIN_Q_DESKEWED)
Maximum board skew within Q group	This parameter describes the largest skew between all Q signals in a Q group. Q pins drive the data signals from the memory to the FPGA when the read operation is active. Users should enter their board skew only. Package skew will be calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins. (Identifier: BOARD_QDR2_BRD_SKEW_WITHIN_Q_NS)
Maximum system skew within Q group	The largest skew between all Q pins in a Q group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_QDR2_PKG_BRD_SKEW_WITHIN_Q_NS)
Package deskewed with board layout (D group)	If you are compensating for package skew on the D and BWS# signals in the board layout (hence checking the box here), please include package skew in calculating the following board skew parameters . (Identifier: BOARD_QDR2_IS_SKEW_WITHIN_D_DESKEWED)
Maximum board skew within D group	This parameter refers to the largest skew between all D and BWS# signals in a D group. D pins are used for driving data signals to the memory device during a write operation. BWS# pins are used as Byte Write Select signals to control which byte(s) are written to the memory during a write operation. Users should enter their board skew only. Package skew will be calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins. (Identifier: BOARD_QDR2_BRD_SKEW_WITHIN_D_NS)
Maximum system skew within D group	The largest skew between all D and BWS# pins in a D group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_QDR2_PKG_BRD_SKEW_WITHIN_D_NS)
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_QDR2_IS_SKEW_WITHIN_AC_DESKEWED)
continued...	



Display Name	Description
Maximum board skew within address/command bus	The largest skew between the address and command signals. Enter the board skew only; package skew is calculated automatically, based on the memory interface configuration, and added to this value. (Identifier: BOARD_QDR2_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	Maximum system skew within address/command bus refers to the largest skew between the address and command signals. (Identifier: BOARD_QDR2_PKG_BRD_SKEW_WITHIN_AC_NS)
Average delay difference between address/command and K	This parameter refers to the average delay difference between the Address/Command signals and the K signal, calculated by averaging the longest and smallest Address/Command trace delay minus the maximum K trace delay. Positive values represent address and command signals that are longer than K signals and negative values represent address and command signals that are shorter than K signals. (Identifier: BOARD_QDR2_AC_TO_K_SKEW_NS)
Maximum K delay to device	The maximum K delay to device refers to the delay of the longest K trace from the FPGA to any device (Identifier: BOARD_QDR2_MAX_K_DELAY_NS)

7.1.6 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Controller

Table 326. Group: Controller

Display Name	Description
Avalon Interface	Selects the Avalon Interface through which the controller interacts with user logic (Identifier: CTRL_QDR2_AVL_PROTOCOL_ENUM)
Maximum Avalon-MM burst length	Specifies the maximum burst length on the Avalon-MM bus. This will be used to configure the FIFOs to be able to manage the maximum data burst. More core logic will be required for an increase in FIFO length. (Identifier: CTRL_QDR2_AVL_MAX_BURST_COUNT)
Generate power-of-2 data bus widths for Qsys	If enabled, the Avalon data bus width is rounded down to the nearest power-of-2. The width of the symbols within the data bus is also rounded down to the nearest power-of-2. You should only enable this option if you know you will be connecting the memory interface to Qsys interconnect components that require the data bus and symbol width to be a power-of-2. If this option is enabled, you cannot utilize the full density of the memory device. For example, in x36 data width upon selecting this parameter, will define the Avalon data bus to 256-bit. This will ignore the upper 4-bit of data width. (Identifier: CTRL_QDR2_AVL_ENABLE_POWER_OF_TWO_BUS)

7.1.7 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Diagnostics

Table 327. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process.
continued...	



Display Name	Description
	<p>Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero.</p> <p><i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations.</i> Abstract PHY is supported with skip calibration. (Identifier: DIAG_QDR2_SIM_CAL_MODE_ENUM)</p>
Abstract phy for fast simulation	<p>Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration. (Identifier: DIAG_QDR2_ABSTRACT_PHY)</p>

Table 328. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	<p>Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic.</p> <p>If you set this parameter to "Disabled", no debug features are enabled. If you set this parameter to "Export", an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP core to it, or connect it to the cal_debug_out interface of another EMIF core. If you select "Add EMIF Debug Interface", an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit.</p> <p><i>Only one EMIF debug interface should be instantiated per I/O column.</i> You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first. (Identifier: DIAG_QDR2_EXPORT_SEQ_AVALON_SLAVE)</p>
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	<p>Specifies that the IP export an Avalon-MM master interface (cal_debug_out) which can connect to the cal_debug interface of other EMIF cores residing in the same I/O column. This parameter applies only if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to the <i>Debugging Multiple EMIFs wiki page</i> for more information about debugging multiple EMIFs. (Identifier: DIAG_QDR2_EXPORT_SEQ_AVALON_MASTER)</p>
Interface ID	<p>Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled, the interface ID is unused. (Identifier: DIAG_QDR2_INTERFACE_ID)</p>
Use Soft NIOS Processor for On-Chip Debug	<p>Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port. <i>Only one interface in a column can activate this option.</i> (Identifier: DIAG_SOFT_NIOS_MODE)</p>



Table 329. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the "Core clocks sharing" parameter in the "General" tab to "Master" or "Slave". (Identifier: DIAG_QDR2_EX_DESIGN_NUM_OF_SLAVES)
Enable In-System-Sources-and-Probes	Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i> . This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_QDR2_EX_DESIGN_ISSP_EN)

Table 330. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_QDR2_USE_TG_AVL_2)
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_QDR2_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the <code>altera_emif_avl_tg_2_tb.sv</code> file. (Identifier: DIAG_QDR2_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_QDR2_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_QDR2_BYPASS_STRESS_STAGE)
Export Traffic Generator 2.0 configuration interface	Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)

Table 331. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_QDR2_EFFICIENCY_MONITOR)



Table 332. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

7.1.8 Intel Stratix 10 EMIF IP QDR II/II+/II+ Xtreme Parameters: Example Designs

Table 333. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the <i>creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization</i> . After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_QDR2_SEL_DESIGN)

Table 334. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created.</i> Instead, the output directory will contain the <code>ed_sim.qsys</code> file which holds Qsys details of the simulation example design, and a <code>make_sim_design.tcl</code> file with other corresponding <code>tcl</code> files. You can run <code>make_sim_design.tcl</code> from a command line to generate the simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_QDR2_GEN_SIM)
Synthesis	Specifies that the 'Generate Example Design' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created.</i> Instead, the output directory will contain the <code>ed_synth.qsys</code> file which holds Qsys details of the synthesis example design, and a <code>make_qii_design.tcl</code> script with other corresponding <code>tcl</code> files. You can run <code>make_qii_design.tcl</code> from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_QDR2_GEN_SYNTH)

Table 335. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_QDR2_HDL_FORMAT)



Table 336. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select 'none' from the 'Select board' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_QDR2_TARGET_DEV_KIT)

7.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

7.2.1 Equations for QDR II, QDR II+, and QDR II+ Xtreme Board Skew Parameters

Table 337. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum system skew within address/command bus	$(MaxAC - MinAC)$ The largest skew between the address and command signals. Enter combined board and package skew.
Average delay difference between address/command and K	The average delay difference between the address and command signals and the K signal, calculated by averaging the longest and smallest Address/Command signal delay minus the K delay. Positive values represent address and command signals that are longer than K signals and negative values represent address and command signals that are shorter than K signals. The Quartus Prime software uses this skew to optimize the delay of the address and command signals to have appropriate setup and hold margins. $\sum_{n=1}^n \left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - K_n PathDelay$ <div style="text-align: center;">n</div> where n is the number of K clocks.
Maximum board skew within Q group	The largest skew between all Q pins in a Q group. Enter your board skew only. Package skew is calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins. $groups_{Max} \left[\max Q_g - \min Q_g \right]$ where g is the number of Q group.
Maximum board skew within D group	The largest skew between all D and BWS# pins in a D group. Enter your board skew only. Package skew is calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins. $groups_{Max} \left[\max D_g - \min D_g \right]$ where g is the number of D group.
Maximum K delay to device	$\max_n (K_n PathDelay)$ where n is the number of K clocks.



7.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.

7.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.

Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

7.3.1.1 Estimating Pin Requirements

You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:



1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.

You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)

7.3.1.2 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.

Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

7.3.1.3 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.



7.3.1.4 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

7.3.1.5 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:

- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.

7.3.1.6 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.

The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the Bank Number and Index within I/O Bank values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the Bank Number value identifies the I/O column, while the letter represents the I/O bank.
- The Index within I/O Bank value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its Index within I/O Bank number (if it is an even number), or by subtracting one from its Index within I/O Bank number (if it is an odd number).

For example, a physical pin with a Bank Number of 2M and Index within I/O Bank of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an Index within I/O Bank of 23 and Bank Number of 2M.

7.3.1.6.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.

If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version_number/<synth/sim>/<variation_name>_altera_emif_arch_nd_version_number_<unique ID>_readme.txt` file, w+ted with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (`.qip`), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLDRAMx , QDRx and LPDDR3 are not supported with HPS.



Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the `<variation_name>/altera_emif_arch_nd_<version>/<synth/sim>/<variation_name>_altera_emif_arch_nd_<version>_<unique ID>_readme.txt` file after you have generated your IP.

7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.



Note: a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.

- b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).

10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:

- There must be an even number of x4 groups in an external memory interface.
- DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group X+1 must be in the same I/O lane, where X is an even number.

11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Xtreme, and RLDRAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLDRAM 3, x36 device, DQ[8:0] and DQ[26:18] are referenced to DK0/DK0#, and DQ[17:9] and DQ[35:27] are referenced to DK1/DK1#.

12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.

13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.

Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLDRAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (global_reset_n) for each individual interface all come from the same input pin or signal.



I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.

CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)



The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the R_{ZQ} pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.

7.3.1.6.2 QDR II, QDR II+ and QDR II+ Xtreme SRAM Command Signals

QDR II, QDR II+ and QDR II+ Xtreme SRAM devices use the write port select (WPS#) signal to control write operations and the read port select (RPS#) signal to control read operations.

7.3.1.6.3 QDR II, QDR II+ and QDR II+ Xtreme SRAM Address Signals

QDR II, QDR II+ and QDR II+ Xtreme SRAM devices use one address bus (A) for both read and write accesses.

7.3.1.6.4 QDR II, QDR II+, and QDR II+ Xtreme SRAM Clock Signals

QDR II, QDR II+ and QDR II+ Xtreme SRAM devices have two pairs of clocks, listed below.

- Input clocks K and K#
- Echo clocks CQ and CQ#

In addition, QDR II devices have a third pair of input clocks, C and C#.



The positive input clock, κ , is the logical complement of the negative input clock, $\kappa\#$. Similarly, \mathcal{C} and \mathcal{CQ} are complements of $\mathcal{C}\#$ and $\mathcal{CQ}\#$, respectively. With these complementary clocks, the rising edges of each clock leg latch the DDR data.

The QDR II SRAM devices use the κ and $\kappa\#$ clocks for write access and the \mathcal{C} and $\mathcal{C}\#$ clocks for read accesses only when interfacing more than one QDR II SRAM device. Because the number of loads that the κ and $\kappa\#$ clocks drive affects the switching times of these outputs when a controller drives a single QDR II SRAM device, \mathcal{C} and $\mathcal{C}\#$ are unnecessary. This is because the propagation delays from the controller to the QDR II SRAM device and back are the same. Therefore, to reduce the number of loads on the clock traces, QDR II SRAM devices have a single-clock mode, and the κ and $\kappa\#$ clocks are used for both reads and writes. In this mode, the \mathcal{C} and $\mathcal{C}\#$ clocks are tied to the supply voltage (VDD). Intel FPGA external memory IP supports only single-clock mode.

For QDR II, QDR II+, or QDR II+ Xtreme SRAM devices, the rising edge of κ is used to capture synchronous inputs to the device and to drive out data through $Q[x:0]$, in similar fashion to QDR II SRAM devices in single clock mode. All accesses are initiated on the rising edge of κ .

\mathcal{CQ} and $\mathcal{CQ}\#$ are the source-synchronous output clocks from the QDR II, QDR II+, or QDR II+ Xtreme SRAM device that accompanies the read data.

The Intel device outputs the κ and $\kappa\#$ clocks, data, address, and command lines to the QDR II, QDR II+, or QDR II+ Xtreme SRAM device. For the controller to operate properly, the write data (\mathcal{D}), address (\mathcal{A}), and control signal trace lengths (and therefore the propagation times) should be equal to the κ and $\kappa\#$ clock trace lengths.

You can generate κ and $\kappa\#$ clocks using any of the PLL registers via the DDR registers. Because of strict skew requirements between κ and $\kappa\#$ signals, use adjacent pins to generate the clock pair. The propagation delays for κ and $\kappa\#$ from the FPGA to the QDR II, QDR II+, or QDR II+ Xtreme SRAM device are equal to the delays on the data and address (\mathcal{D} , \mathcal{A}) signals. Therefore, the signal skew effect on the write and read request operations is minimized by using identical DDR output circuits to generate clock and data inputs to the memory.

7.3.1.6.5 QDR II, QDR II+ and QDR II+ Xtreme SRAM Data, BWS, and QVLD Signals

QDR II, QDR II+ and QDR II+ Xtreme SRAM devices use two unidirectional data buses: one for writes (\mathcal{D}) and one for reads (\mathcal{Q}).

At the pin, the read data is edge-aligned with the \mathcal{CQ} and $\mathcal{CQ}\#$ clocks while the write data is center-aligned with the κ and $\kappa\#$ clocks (see the following figures).

The diagram illustrates the timing relationship between the FPGA Pin and the Capture Register for the Q and Q# signals. The signals shown are:

- CQ at FPGA Pin
- CQ# at FPGA Pin
- Q at FPGA Pin
- CQ at Capture Register
- CQ# at Capture Register
- Q at Capture Register

Two vertical dashed lines indicate the DQS phase shift, which is the time difference between the sampling clock (Q) at the FPGA Pin and the sampling clock (Q) at the Capture Register. The shift is shown as a time delay from the FPGA Pin to the Capture Register.

The timing diagram shows three signals: K at FPGA Pin, K# at FPGA Pin, and D at FPGA Pin. The K signal is a square wave that transitions from low to high at the first vertical dashed line and from high to low at the second vertical dashed line. The K# signal is the inverse of K, transitioning from high to low at the first vertical dashed line and from low to high at the second vertical dashed line. The D signal is a square wave that transitions from low to high at the first vertical dashed line and from high to low at the second vertical dashed line. The D signal transitions occur at the same time as the K and K# transitions.

QDR II+ and QDR II+ Xtreme SRAM devices also have a QVLD pin that indicates valid read data. The QVLD signal is edge-aligned with the echo clock and is asserted high for approximately half a clock cycle before data is output from memory.

Note: The Intel FPGA external memory interface IP does not use the QVLD signal.

In Intel Cyclone 10Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.



PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.

Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `pll_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note:

You can place the `pll_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.

Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.



Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.

7.4 QDR II/II+/II+ Xtreme Board Design Guidelines

The following topics provide guidelines for you to improve your system's signal integrity and layout guidelines to help successfully implement a QDR II, QDR II+, or QDR II+ Xtreme SRAM interface in your system.

Note: In the following topics, QDR II SRAM refers to QDR II, QDR II+, and QDR II+ Xtreme SRAM unless stated otherwise.

The following topics focus on the following key factors that affect signal integrity:

- I/O standards
- QDR II SRAM configurations
- Signal terminations
- Printed circuit board (PCB) layout guidelines

I/O Standards

QDR II SRAM interface signals use one of the following JEDEC I/O signalling standards:

- HSTL-15—provides the advantages of lower power and lower emissions.
- HSTL-18—provides increased noise immunity with slightly greater output voltage swings.

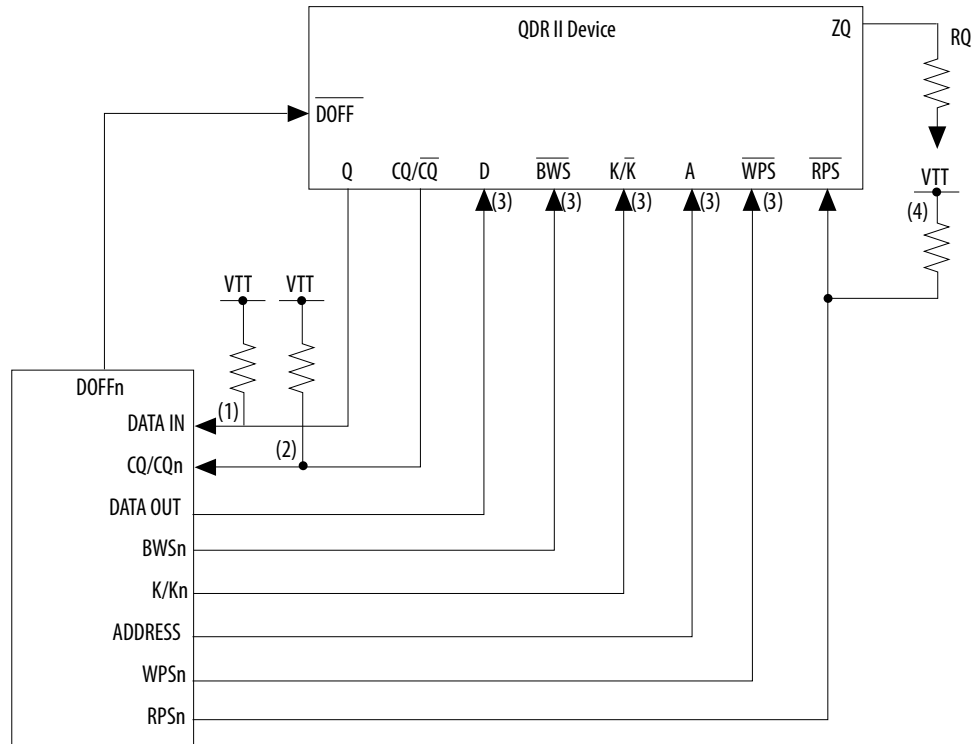
7.4.1 QDR II SRAM Configurations

The QDR II SRAM Controller for Intel Stratix 10 EMIF IP supports interfaces with a single device, and two devices in a width expansion configuration up to maximum width of 72 bits.

The following figure shows the main signal connections between the FPGA and a single QDR II SRAM component.

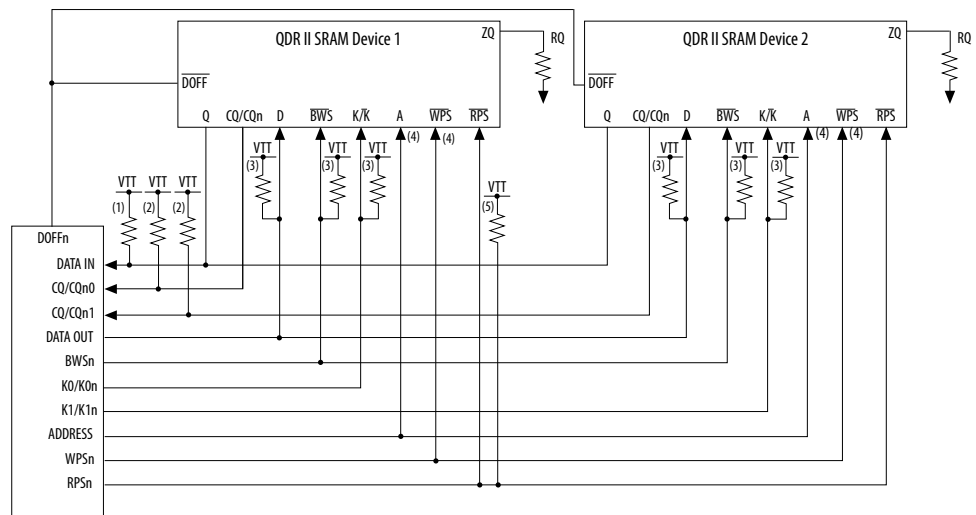


Figure 73. Configuration With A Single QDR II SRAM Component




The following figure shows the main signal connections between the FPGA and two QDR II SRAM components in a width expansion configuration.

Figure 74. Configuration With Two QDR II SRAM Components In A Width Expansion Configuration



The following figure shows the detailed balanced topology recommended for the address and command signals in the width expansion configuration.

Figure 75. External Parallel Termination for Balanced Topology



7.4.2 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.




Table 338. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-Junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)

7.4.3 QDR II Layout Guidelines

The following table summarizes QDR II and QDR II SRAM general routing layout guidelines.

- Note:**
- The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristics of your PCB implementation. They do not include any margin for crosstalk.
 - Intel recommends that you get accurate time base skew numbers when you simulate your specific implementation.
 - To reliably close timing to and from the periphery of the device, signals to and from the periphery should be registered before any further logic is connected.

Table 339. QDR II and QDR II+ SRAM Layout Guidelines

Parameter	Guidelines
General Routing	<ul style="list-style-type: none"> If signals of the same net group must be routed on different layers with the same impedance characteristic, you must simulate your worst case PCB trace tolerances to ascertain actual propagation delay differences. Typical later to later trace delay variations are of 15 ps/inch order. Avoid T-junctions greater than 150 ps.
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to under 150 ps. These signals should maintain a 10-mil (0.254 mm) spacing from other nets. Clocks should maintain a length-matching between clock pairs of ± 5 ps. Complementary clocks should maintain a length-matching between P and N signals of ± 2 ps. Keep the distance from the pin on the QDR II SRAM component to stub termination resistor (V_{TT}) to less than 50 ps for the K, K# clocks. Keep the distance from the pin on the QDR II SRAM component to fly-by termination resistor (V_{TT}) to less than 100 ps for the K, K# clocks. Keep the distance from the pin on the FPGA component to stub termination resistor (V_{TT}) to less than 50 ps for the echo clocks, CQ, CQ#, if they require an external discrete termination. Keep the distance from the pin on the FPGA component to fly-by termination resistor (V_{TT}) to less than 100 ps for the echo clocks, CQ, CQ#, if they require an external discrete termination.
External Memory Routing Rules	<ul style="list-style-type: none"> Keep the distance from the pin on the QDR II SRAM component to stub termination resistor (V_{TT}) to less than 50 ps for the write data, byte write select and address/command signal groups. Keep the distance from the pin on the QDR II SRAM component to fly-by termination resistor (V_{TT}) to less than 100 ps for the write data, byte write select and address/command signal groups. Keep the distance from the pin on the FPGA to stub termination resistor (V_{TT}) to less than 50 ps for the read data signal group. Keep the distance from the pin on the FPGA to fly-by termination resistor (V_{TT}) to less than 100 ps for the read data signal group. Parallelism rules for the QDR II SRAM data/address/command groups are as follows: <ul style="list-style-type: none"> 4 mils for parallel runs < 0.1 inch (approximately 1× spacing relative to plane distance). 5 mils for parallel runs < 0.5 inch (approximately 1× spacing relative to plane distance). 10 mils for parallel runs between 0.5 and 1.0 inches (approximately 2× spacing relative to plane distance). 15 mils for parallel runs between 1.0 and 6.0 inch (approximately 3× spacing relative to plane distance).
Maximum Trace Length	<ul style="list-style-type: none"> Keep the maximum trace length of all signals from the FPGA to the QDR II SRAM components to 6 inches.

Related Links

[Power Distribution Network](#)

7.4.4 QDR II SRAM Layout Approach

Using the layout guidelines in the above table, Intel recommends the following layout approach:



1. Route the $\kappa/\kappa\#$ clocks and set the clocks as the target trace propagation delays for the output signal group.
2. Route the write data output signal group (`write_data`, `byte_write_select`), ideally on the same layer as the $\kappa/\kappa\#$ clocks, to within ± 10 ps skew of the $\kappa/\kappa\#$ traces.
3. Route the address/control output signal group (`address`, `RPS`, `WPS`), ideally on the same layer as the $\kappa/\kappa\#$ clocks, to within ± 20 ps skew of the $\kappa/\kappa\#$ traces.
4. Route the $CQ/CQ\#$ clocks and set the clocks as the target trace propagation delays for the input signal group.
5. Route the read data output signal group (`read_data`), ideally on the same layer as the $CQ/CQ\#$ clocks, to within ± 10 ps skew of the $CQ/CQ\#$ traces.
6. The output and input groups do not need to have the same propagation delays, but they must have all the signals matched closely within the respective groups.

Note: Intel recommends that you create your project with a fully implemented external memory interface, and observe the interface timing margins to determine the actual margins for your design.

Although the recommendations in this section are based on simulations, you can apply the same general principles when determining the best termination scheme, drive strength setting, and loading style to any board designs. Even armed with this knowledge, it is still critical that you perform simulations, either using IBIS or HSPICE models, to determine the quality of signal integrity on your designs.

7.4.5 Package Deskew

You should follow Intel's package deskew guidance.

Related Links

[Package Deskew](#)



8 Intel Stratix 10 EMIF IP for QDR-IV

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for QDR-IV.

8.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

8.1.1 Intel Stratix 10 EMIF IP QDR-IV Parameters: General

Table 340. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_QDR4_CONFIG_ENUM)

Table 341. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_QDR4_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_QDR4_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_QDR4_USER_REF_CLK_FREQ_MHZ)
<i>continued...</i>	



Display Name	Description
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_QDR4_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_QDR4_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_QDR4_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 342. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 343. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)

Table 344. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)



Table 345. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 346. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

8.1.2 Intel Stratix 10 EMIF IP QDR-IV Parameters: Memory

Table 347. Group: Memory / Topology

Display Name	Description
DQ width per device	Specifies number of DQ pins per port per QDR IV device. Available widths for DQ are x18 and x36. (Identifier: MEM_QDR4_DQ_PER_PORT_PER_DEVICE)
Enable width expansion	Indicates whether to combine two memory devices to double the data bus width. With two devices, the interface supports a width expansion configuration up to 72-bits. For width expansion configuration, the address and control signals are routed to 2 devices. (Identifier: MEM_QDR4_WIDTH_EXPANDED)
Address width	Number of address pins. (Identifier: MEM_QDR4_ADDR_WIDTH)
Memory Type	The QDR-IV family includes two members: MEM_XP: QDR-IV Xtreme Performance (XP) with a Maximum Clock Frequency of 1066MHz MEM_HP: QDR-IV High Performance (HP) with a Maximum Clock Frequency of 667MHz. (Identifier: MEM_QDR4_MEM_TYPE_ENUM)

Table 348. Group: Memory / Configuration Register Settings

Display Name	Description
Address bus inversion	Enable address bus inversion. AINV are all active high at memory device. (Identifier: MEM_QDR4_ADDR_INV_ENA)
Data bus inversion	Enable data bus inversion for DQ pins. DINVA[1:0] and DINVB[1:0] are all active high. When set to 1, the corresponding bus is inverted at memory device. If the data inversion feature is programmed to be OFF, then the DINVA/DINVB output bits will always be driven to 0. (Identifier: MEM_QDR4_DATA_INV_ENA)
ODT (Clock)	Determines the configuration register setting that controls the clock ODT setting. (Identifier: MEM_QDR4_CK_ODT_MODE_ENUM)
ODT (Address/Command)	Determines the configuration register setting that controls the address/command ODT setting. (Identifier: MEM_QDR4_AC_ODT_MODE_ENUM)
continued...	



Display Name	Description
ODT (Data)	Determines the configuration register setting that controls the data ODT setting. (Identifier: MEM_QDR4_DATA_ODT_MODE_ENUM)
Output drive (pull-up)	Determines the configuration register setting that controls the pull-up output drive setting. (Identifier: MEM_QDR4_PU_OUTPUT_DRIVE_MODE_ENUM)
Output drive (pull-down)	Determines the configuration register setting that controls the pull-down output drive setting. (Identifier: MEM_QDR4_PD_OUTPUT_DRIVE_MODE_ENUM)

8.1.3 Intel Stratix 10 EMIF IP QDR-IV Parameters: Controller

Table 349. Group: Controller

Display Name	Description
Avalon Interface	Selects the Avalon Interface through which the controller interacts with user logic (Identifier: CTRL_QDR4_AVL_PROTOCOL_ENUM)
Maximum Avalon-MM burst length	Specifies the maximum burst length on the Avalon-MM bus. This will be used to configure the FIFOs to be able to manage the maximum data burst. More core logic will be required for an increase in FIFO length. (Identifier: CTRL_QDR4_AVL_MAX_BURST_COUNT)
Generate power-of-2 data bus widths for Qsys	If enabled, the Avalon data bus width is rounded down to the nearest power-of-2. The width of the symbols within the data bus is also rounded down to the nearest power-of-2. You should only enable this option if you know you will be connecting the memory interface to Qsys interconnect components that require the data bus and symbol width to be a power-of-2. If this option is enabled, you cannot utilize the full density of the memory device. For example, in x36 data width upon selecting this parameter, will define the Avalon data bus to 256-bit. This will ignore the upper 4-bit of data width. (Identifier: CTRL_QDR4_AVL_ENABLE_POWER_OF_TWO_BUS)
Additional read-after-write turnaround time	Specifies an additional number of idle memory cycles when switching the data bus (of a single port) from a write to a read. (Identifier: CTRL_QDR4_ADD_RAW_TURNAROUND_DELAY_CYC)
Additional write-after-read turnaround time	Specifies an additional number of idle memory cycles when switching the data bus (of a single port) from a read to a write. (Identifier: CTRL_QDR4_ADD_WAR_TURNAROUND_DELAY_CYC)

8.1.4 Intel Stratix 10 EMIF IP QDR-IV Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.




Table 350. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_QDR4_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_QDR4_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_QDR4_DEFAULT_IO)

Table 351. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR4_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR4_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_QDR4_USER_AC_SLEW_RATE_ENUM)

Table 352. Group: FPGA I/O / FPGA I/O Settings / Memory Clock

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR4_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR4_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_QDR4_USER_CK_SLEW_RATE_ENUM)



Table 353. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_QDR4_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR4_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_QDR4_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_QDR4_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_QDR4_USER_STARTING_VREFIN)

Table 354. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_QDR4_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_QDR4_USER_RZQ_IO_STD_ENUM)

8.1.5 Intel Stratix 10 EMIF IP QDR-IV Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).

Table 355. Group: Mem Timing

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_QDR4_SPEEDBIN_ENUM)
tISH	tISH provides the setup/hold window requirement for the entire data bus (DK or DINV) in all the data groups with respect to the DK clock . After deskew calibration, this parameter describes the intersection window for all the individual data bus signals setup/hold margins. (Identifier: MEM_QDR4_TISH_PS)
tQKQ_max	tQKQ_max describes the maximum skew between the read strobe (QK) clock edge to the data bus (DQ/DINV) edge. (Identifier: MEM_QDR4_TQKQ_MAX_PS)
tQH	tQH specifies the output hold time for the DQ/DINV in relation to QK . (Identifier: MEM_QDR4_TQH_CYC)
<i>continued...</i>	



Display Name	Description
tCKDK_max	tCKDK_max refers to the maximum skew from the memory clock (CK) to the write strobe (DK) . (Identifier: MEM_QDR4_TCKDK_MAX_PS)
tCKDK_min	tCKDK_min refers to the minimum skew from the memory clock (CK) to the write strobe (DK) . (Identifier: MEM_QDR4_TCKDK_MIN_PS)
tCKQK_max	tCKQK_max refers to the maximum skew from the memory clock (CK) to the read strobe (QK) . (Identifier: MEM_QDR4_TCKQK_MAX_PS)
tASH	tASH provides the setup/hold window requirement for the address bus in relation to the CK clock . Because the individual signals in the address bus may not be perfectly aligned with each other, this parameter describes the intersection window for all the individual address signals setup/hold margins. (Identifier: MEM_QDR4_TASH_PS)
tCSH	tCSH provides the setup/hold window requirement for the control bus (LD#, RW#) in relation to the CK clock . Because the individual signals in the control bus may not be perfectly aligned with each other, this parameter describes the intersection window for all the individual control signals setup/hold margins. (Identifier: MEM_QDR4_TCSH_PS)

8.1.6 Intel Stratix 10 EMIF IP QDR-IV Parameters: Board

Table 356. Group: Board / Intersymbol Interference/Crosstalk


Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation using IBIS models and Hyperlynx*, and manually enter values based on your simulation results, instead of using the default values.</i> (Identifier: BOARD_QDR4_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to ISI and crosstalk effects. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_QDR4_USER_AC_ISI_NS)
QK/QK# ISI/crosstalk	QK/QK# ISI/crosstalk describes the reduction of the read data window due to intersymbol interference and crosstalk effects on the QK/QK# signal when driven by the memory device during a read. The number to be entered in the Quartus Prime software is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_QDR4_USER_RCLK_ISI_NS)
Read DQ ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQ signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_QDR4_USER_RDATA_ISI_NS)
DK/DK# ISI/crosstalk	DK/DK# ISI/crosstalk describes the reduction of the write data window due to intersymbol interference and crosstalk effects on the DK/DK# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on
continued...	



Display Name	Description
	the setup side + measured loss on the hold side). Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_QDR4_USER_WCLK_ISI_NS)
Write DQ ISI/crosstalk	The reduction of the write data window due to intersymbol interference and crosstalk effects on the DQ signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side). Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_QDR4_USER_WDATA_ISI_NS)

Table 357. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (QK group)	If you are compensating for package skew on the QK bus in the board layout (hence checking the box here), please include package skew in calculating the following board skew parameters. (Identifier: BOARD_QDR4_IS_SKEW_WITHIN_QK_DESKEWED)
Maximum board skew within QK group	The largest skew between all DQ and DM pins in a QK group. Enter your board skew only. Package skew will be calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins. (Identifier: BOARD_QDR4_BRD_SKEW_WITHIN_QK_NS)
Maximum system skew within QK group	Maximum system skew within QK group refers to the largest skew between all DQ and DM pins in a QK group. This value can affect the read capture and write margins. (Identifier: BOARD_QDR4_PKG_BRD_SKEW_WITHIN_QK_NS)
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_QDR4_IS_SKEW_WITHIN_AC_DESKEWED)
Maximum board skew within address/command bus	The largest skew between the address and command signals. Enter the board skew only; package skew is calculated automatically, based on the memory interface configuration, and added to this value. (Identifier: BOARD_QDR4_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	Maximum system skew within address/command bus refers to the largest skew between the address and command signals. (Identifier: BOARD_QDR4_PKG_BRD_SKEW_WITHIN_AC_NS)
Average delay difference between DK and CK	This parameter describes the average delay difference between the DK signals and the CK signal, calculated by averaging the longest and smallest DK trace delay minus the CK trace delay. Positive values represent DK signals that are longer than CK signals and negative values represent DK signals that are shorter than CK signals. (Identifier: BOARD_QDR4_DK_TO_CK_SKEW_NS)
Maximum delay difference between devices	This parameter describes the largest propagation delay on the DQ signals between ranks. For example, in a two-rank configuration where devices are placed in series, there is an extra propagation delay for DQ signals going to and coming back from the furthest device compared to the nearest device. <i>This parameter is only applicable when there is more than one rank.</i> (Identifier: BOARD_QDR4_SKEW_BETWEEN_DIMMS_NS)
Maximum skew between DK groups	This parameter describes the largest skew between DK signals in different DK groups. (Identifier: BOARD_QDR4_SKEW_BETWEEN_DK_NS)
Average delay difference between address/command and CK	The average delay difference between the address/command signals and the CK signal, calculated by averaging the longest and smallest address/command signal trace delay minus the maximum CK trace delay. Positive values represent address and command signals that are longer than CK
continued...	



Display Name	Description
	signals and negative values represent address and command signals that are shorter than CK signals. (Identifier: BOARD_QDR4_AC_TO_CK_SKEW_NS)
Maximum CK delay to device	The maximum CK delay to device refers to the delay of the longest CK trace from the FPGA to any device. (Identifier: BOARD_QDR4_MAX_CK_DELAY_NS)
Maximum DK delay to device	The maximum DK delay to device refers to the delay of the longest DK trace from the FPGA to any device. (Identifier: BOARD_QDR4_MAX_DK_DELAY_NS)

8.1.7 Intel Stratix 10 EMIF IP QDR-IV Parameters: Diagnostics

Table 358. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	<p>Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process.</p> <p>Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero.</p> <p><i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations.</i> Abstract PHY is supported with skip calibration. (Identifier: DIAG_QDR4_SIM_CAL_MODE_ENUM)</p>
Abstract phy for fast simulation	<p>Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration. (Identifier: DIAG_QDR4_ABSTRACT_PHY)</p>

Table 359. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	<p>Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic.</p> <p>If you set this parameter to "Disabled", no debug features are enabled. If you set this parameter to "Export", an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP core to it, or connect it to the cal_debug_out interface of another EMIF core. If you select "Add EMIF Debug Interface", an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit.</p> <p><i>Only one EMIF debug interface should be instantiated per I/O column.</i> You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first. (Identifier: DIAG_QDR4_EXPORT_SEQ_AVALON_SLAVE)</p>
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	<p>Specifies that the IP export an Avalon-MM master interface (cal_debug_out) which can connect to the cal_debug interface of other EMIF cores residing in the same I/O column. This parameter applies only if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to</p>

continued...



Display Name	Description
	<i>the Debugging Multiple EMIFs wiki page for more information about debugging multiple EMIFs.</i> (Identifier: DIAG_QDR4_EXPORT_SEQ_AVALON_MASTER)
Interface ID	Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled , the interface ID is unused. (Identifier: DIAG_QDR4_INTERFACE_ID)
Skip VREF_in calibration	Specifies to skip the VREF stage of calibration. Enable this parameter for debug purposes only ; generally, you should include the VREF calibration stage during normal operation. (Identifier: DIAG_QDR4_SKIP_VREF_CAL)
Use Soft NIOS Processor for On-Chip Debug	Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port . <i>Only one interface in a column can activate this option.</i> (Identifier: DIAG_SOFT_NIOS_MODE)

Table 360. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the "Core clocks sharing" parameter in the "General" tab to "Master" or "Slave" . (Identifier: DIAG_QDR4_EX_DESIGN_NUM_OF_SLAVES)
Enable In-System-Sources-and-Probes	Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i> . This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_QDR4_EX_DESIGN_ISSP_EN)

Table 361. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_QDR4_USE_TG_AVL_2)
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_QDR4_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the <code>altera_emif_avl_tg_2_tb.sv</code> file. (Identifier: DIAG_QDR4_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_QDR4_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter,
continued...	



Display Name	Description
	the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_QDR4_BYPASS_STRESS_STAGE)
Export Traffic Generator 2.0 configuration interface	Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)

Table 362. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_QDR4_EFFICIENCY_MONITOR)

Table 363. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

8.1.8 Intel Stratix 10 EMIF IP QDR-IV Parameters: Example Designs

Table 364. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the <i>creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization.</i> After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_QDR4_SEL_DESIGN)

Table 365. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created.</i> Instead, the output directory will contain the ed_sim.qsys file which holds Qsys details of the simulation example design, and a make_sim_design.tcl file with other corresponding tcl files. You can run make_sim_design.tcl from a command line to generate the
continued...	



Display Name	Description
	simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_QDR4_GEN_SIM)
Synthesis	Specifies that the ' Generate Example Design ' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created.</i> Instead, the output directory will contain the ed_synth.qsys file which holds Qsys details of the synthesis example design, and a make_qii_design.tcl script with other corresponding tcl files. You can run make_qii_design.tcl from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_QDR4_GEN_SYNTH)

Table 366. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_QDR4_HDL_FORMAT)

Table 367. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select ' none ' from the ' Select board ' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_QDR4_TARGET_DEV_KIT)

8.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

8.2.1 Equations for QDR-IV Board Skew Parameters

Table 368. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum system skew within address/command bus	The largest skew between the address and command signals. Enter combined board and package skew.
continued...	



Parameter	Description/Equation
	$(MaxAC - MinAC)$
Average delay difference between address/command and CK	<p>The average delay difference between the address and command signals and the CK signal, calculated by averaging the longest and smallest Address/Command signal delay minus the CK delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the address and command signals to have appropriate setup and hold margins.</p> $\sum_{n=1}^n \left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - CK_n PathDelay$ <p>where n is the number of memory clocks.</p>
Maximum System skew within QK group	<p>The largest skew between all DQ and DM pins in a QK group. Enter combined board and package skew. This value affects the read capture and write margins.</p> $\max_n (\max DQ_n - \min DQ_n)$ <p>Where n includes both DQa and DQb</p>
Maximum CK delay to device	<p>The delay of the longest CK trace from the FPGA to any device.</p> $\max_n (CK_n PathDelay)$ <p>where n is the number of memory clocks.</p>
Maximum DK delay to device	<p>The delay of the longest DK trace from the FPGA to any device.</p> $\max_n (DK_n PathDelay)$ <p>where n is the number of DK.</p>
Average delay difference between DK and CK	<p>The average delay difference between the DK signals and the CK signal, calculated by averaging the longest and smallest DK delay minus the CK delay. Positive values represent DK signals that are longer than CK signals and negative values represent DK signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the DK signals to have appropriate setup and hold margins.</p> $\min_{n,m} \left(\frac{CK_n PathDelay - DK_m PathDelay + \max_{n,m} (CK_n PathDelay - DK_m PathDelay)}{2} \right)$ <p>CDO:/content/authoring/rto1474984235656.xml where n is the number of memory clocks and m is the number of DK.</p>
Maximum skew between DK groups	<p>The largest skew between DK signals in different DK groups.</p> $\max_n (\max DK_n - \min DK_n)$ <p>where n is the number of DK. Where n includes both DQa and DQb.</p>

8.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.



8.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.

Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

8.3.1.1 Estimating Pin Requirements

You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:

1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.

You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)

8.3.1.2 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.

Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

8.3.1.3 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.


8.3.1.4 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

8.3.1.5 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:



- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.

8.3.1.6 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.

The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the `Bank Number` and `Index within I/O Bank` values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the `Bank Number` value identifies the I/O column, while the letter represents the I/O bank.
- The `Index within I/O Bank` value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its `Index within I/O Bank` number (if it is an even number), or by subtracting one from its `Index within I/O Bank` number (if it is an odd number).

For example, a physical pin with a `Bank Number` of 2M and `Index within I/O Bank` of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an `Index within I/O Bank` of 23 and `Bank Number` of 2M.

8.3.1.6.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.



If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version number/<synth|sim>/<variation_name>_altera_emif_arch_nd_version number_<unique ID>_readme.txt` file, w+ted with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (.qip), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLDRAMx , QDRx and LPDDR3 are not supported with HPS.

Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the `<variation_name>/altera_emif_arch_nd_<version>/<synth|sim>/<variation_name>_altera_emif_arch_nd_<version>_<unique ID>_readme.txt` file after you have generated your IP.



7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.

Note: a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.

b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).
10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:
 - There must be an even number of x4 groups in an external memory interface.
 - DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group X+1 must be in the same I/O lane, where X is an even number.
11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Extreme, and RLDRAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLDRAM 3, x36 device, DQ[8:0] and DQ[26:18] are referenced to DK0/DK0#, and DQ[17:9] and DQ[35:27] are referenced to DK1/DK1#.
12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.
13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.

Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLDRAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (`global_reset_n`) for each individual interface all come from the same input pin or signal.

I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.



CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)

The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the R_{ZQ} pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.

8.3.1.6.2 QDR IV SRAM Commands and Addresses, AP, and AINV Signals

The CK and CK# signals clock the commands and addresses into the memory devices. There is one pair of CK and CK# pins per QDR IV SRAM device. These pins operate at double data rate using both rising and falling edge. The rising edge of CK latches the addresses for port A, while the falling edge of CK latches the addresses inputs for port B.

QDR IV SRAM devices have the ability to invert all address pins to reduce potential simultaneous switching noise. Such inversion is accomplished using the Address Inversion Pin for Address and Address Parity Inputs (AINV), which assumes an address parity of 0, and indicates whether the address bus and address parity are inverted.

The above features are available as **Option Control** under **Configuration Register Settings** in the parameter editor. The commands and addresses must meet the memory address and command setup (t_{AS} , t_{CS}) and hold (t_{AH} , t_{CH}) time requirements.

8.3.1.6.3 QDR IV SRAM Clock Signals

QDR IV SRAM devices have three pairs of differential clocks.

The three QDR IV differential clocks are as follows:

- Address and Command Input Clocks CK and CK#
- Data Input Clocks DKx and DKx#, where x can be A or B, referring to the respective ports
- Data Output Clocks, QKx and QKx#, where x can be A or B, referring to the respective ports

QDR IV SRAM devices have two independent bidirectional data ports, Port A and Port B, to support concurrent read/write transactions on both ports. These data ports are controlled by a common address port clocked by CK and CK# in double data rate. There is one pair of CK and CK# pins per QDR IV SRAM device.

DKx and DKx# samples the DQx inputs on both rising and falling edges. Similarly, QKx and QKx# samples the DQx outputs on both rising and falling edges.

QDR IV SRAM devices employ two sets of free running differential clocks to accompany the data. The DKx and DKx# clocks are the differential input data clocks used during writes. The QKx and QKx# clocks are the output data clocks used during reads. Each pair of DKx and DKx#, or QKx and QKx# clocks are associated with either 9 or 18 data bits.

The polarity of the QKB and QKB# pins in the Intel FPGA external memory interface IP was swapped with respect to the polarity of the differential input buffer on the FPGA. In other words, the QKB pins on the memory side must be connected to the negative pins of the input buffers on the FPGA side, and the QKB# pins on the memory side must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).



QDR IV SRAM devices are available in x18 and x36 bus width configurations. The exact clock-data relationships are as follows:

- For x18 data bus width configuration, there are 9 data bits associated with each pair of write and read clocks. So, there are two pairs of DKx and $DKx\#$ pins and two pairs of QKx or $QKx\#$ pins.
- For x36 data bus width configuration, there are 18 data bits associated with each pair of write and read clocks. So, there are two pairs of DKx and $DKx\#$ pins and two pairs of QKx or $QKx\#$ pins.

There are t_{CKDK} timing requirements for skew between CK and DKx or $CK\#$ and $DKx\#$. Similarly, there are t_{CKQK} timing requirements for skew between CK and QKx or $CK\#$ and $QKx\#$.

8.3.1.6.4 QDR IV SRAM Data, DINV, and QVLD Signals

The read data is edge-aligned with the QKA or $QKB\#$ clocks while the write data is center-aligned with the DKA and $DKB\#$ clocks.

QK is shifted by the DLL so that the clock edges can be used to clock in the DQ at the capture register.

Figure 76. Edge-Aligned DQ and QK Relationship During Read





Figure 77. Center-Aligned DQ and DK Relationship During Write



The polarity of the QKB and $QKB\#$ pins in the Intel FPGA external memory interface IP was swapped with respect to the polarity of the differential input buffer on the FPGA. In other words, the QKB pins on the memory side need to be connected to the negative pins of the input buffers on the FPGA side, and the $QKB\#$ pins on the memory side need to be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, `mem_qkb` is assigned to the negative buffer leg, and `mem_qkb_n` is assigned to the positive buffer leg).

The synchronous read/write input, $RWx\#$, is used in conjunction with the synchronous load input, $LDx\#$, to indicate a Read or Write Operation. For port A, these signals are sampled on the rising edge of CK clock, for port B, these signals are sampled on the falling edge of CK clock.

QDR IV SRAM devices have the ability to invert all data pins to reduce potential simultaneous switching noise, using the Data Inversion Pin for DQ Data Bus, $DINVx$. This pin indicates whether DQx pins are inverted or not.

To enable the data pin inversion feature, click **Configuration Register Settings > Option Control** in the parameter editor.

QDR IV SRAM devices also have a $QVLD$ pin which indicates valid read data. The $QVLD$ signal is edge-aligned with QKx or $QKx\#$ and is high approximately one-half clock cycle before data is output from the memory.

Note: The Intel ZFPGA external memory interface IP does not use the $QVLD$ signal.

8.3.1.7 Resource Sharing Guidelines (Multiple Interfaces)

In Intel Cyclone 10Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.

PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.


Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `pll_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note: You can place the `pll_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.



Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.

Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.

8.4 QDR-IV Board Design Guidelines

The following topics provide guidelines for you to improve your system's signal integrity and layout guidelines to help successfully implement a QDR-IV SRAM interface in your system.

The following topics focus on the following key factors that affect signal integrity:

- I/O standards
- QDR-IV SRAM configurations
- Signal terminations
- Printed circuit board (PCB) layout guidelines

I/O Standards

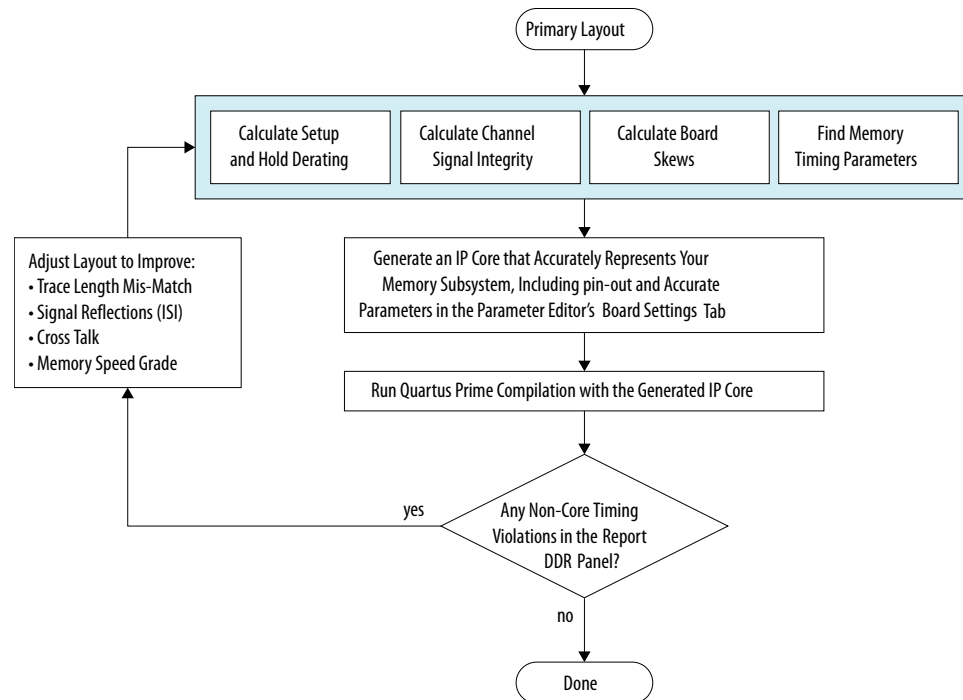
QDR-IV SRAM interface signals use one of the following JEDEC I/O signalling standards:

- HSTL-15—provides the advantages of lower power and lower emissions.
- HSTL-18—provides increased noise immunity with slightly greater output voltage swings.

8.4.1 QDR-IV Layout Approach

For all practical purposes, you can regard the Timing Analyzer report on your memory interface as definitive for a given set of memory and board timing parameters. You will find timing under Report DDR in Timing Analyzer and on the Timing Analysis tab in the parameter editor.

The following flowchart illustrates the recommended process to follow during the design phase, to determine timing margin and make iterative improvements to your design.



For more detailed simulation guidance, refer to the wiki: http://www.alterawiki.com/wiki/Arria_10_EMIF_Simulation_Guidance

Intersymbol Interference/Crosstalk

For information on intersymbol interference and crosstalk, refer to the wiki: http://www.alterawiki.com/wiki/Arria_10_EMIF_Simulation_Guidance

Board Skew

For information on calculating board skew parameters, refer to *Board Skew Equations*, in this chapter.

If you know the absolute delays for all the memory related traces, the interactive [Board Skew Parameter Tool](#) can help you calculate the necessary parameters.

Memory Timing Parameters

You can find the memory timing parameters to enter in the parameter editor, in your memory vendor's datasheet.



8.4.2 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.

Table 369. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-Junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)

8.4.3 QDR-IV Layout Guidelines

Observe the following layout guidelines for your QDR-IV interface.



Parameter	Guidelines
General Routing	<ul style="list-style-type: none"> If you must route signals of the same net group on different layers with the same impedance characteristic, simulate your worst case PCB trace tolerances to determine actual propagation delay differences. Typical layer-to-layer trace delay variations are on the order of 15 ps/inch. Avoid T-junctions greater than 150 ps. Match all signals within a given DQ group with a maximum skew of ± 10 ps and route on the same layer.
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to less than 150 ps. Clock signals should maintain a 10-mil (0.254 mm) spacing from other nets. Clocks should maintain a length-matching between clock pairs of ± 5 ps. Differential clocks should maintain a length-matching between P and N signals of ± 2 ps. Space between different clock pairs should be at least three times the space between the traces of a differential pair.
Address and Command Routing	<ul style="list-style-type: none"> - To minimize crosstalk, route address, bank address, and command signals on a different layer than the data signals. - Do not route the differential clock signals close to the address signals. - Keep the distance from the pin on the QDR-IV component to the stub termination resistor (VTT) to less than 50 ps for the address/command signal group. - Route the mem_ck (CK/CK#) clocks and set as the target trace propagation delays for the address/command signal group. Match the CK/CK# clock to within ± 50 ps of all the DK/DK# clocks for both ports. - Route the address/control signal group ideally on the same layer as the mem_ck (CK/CK#) clocks, to within ± 20 ps skew of the mem_ck (CK/CK#) traces.
Data Signals	<ul style="list-style-type: none"> For port B only: Swap the polarity of the QKB and QKB# signals with respect to the polarity of the differential buffer inputs on the FPGA. Connect the positive leg of the differential input buffer on the FPGA to QDR-IV QKB# (negative) pin and vice-versa. Note that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg). For each port, route the DK/DK# write clock and QK/QK# read clock associated with a DQ group on the same PCB layer. Match these clock pairs to within ± 5 ps. For each port, set the DK/DK# or QK/QK# clock as the target trace propagation delay for the associated data signals (DQ). For each port, route the data (DQ) signals for the DQ group ideally on the same layer as the associated QK/QK# and DK/DK# clocks to within ± 10 ps skew of the target clock.
Maximum Trace Length	<ul style="list-style-type: none"> Keep the maximum trace length of all signals from the FPGA to the QDR-IV components to 600 ps.
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always make sure that the signals related to memory interface are routed between appropriate GND or power layers. For Data and Data Strobe traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Clock (mem_CK) traces: Maintain at least 5H spacing between two clock pair or a clock pair and any other memory interface trace, where H is the vertical distance to the closest return path for that particular trace.
Trace Matching Guidance	<p>The following layout approach is recommended, based on the preceding guidelines:</p> <ol style="list-style-type: none"> For port B only: Swap the polarity of the QKB and QKB# signals with respect to the polarity of the differential buffer inputs on the FPGA. Connect the positive leg of the differential input buffer on the FPGA to QDR-IV QKB# (negative) pin and vice-versa. Note
continued...	



Parameter	Guidelines
	<p>that the port names at the top-level of the IP already reflect this swap (that is, <code>mem_qkb</code> is assigned to the negative buffer leg, and <code>mem_qkb_n</code> is assigned to the positive buffer leg).</p> <ol style="list-style-type: none"> For each port, set the <code>DK/DK#</code> or <code>QK/QK#</code> clock as the target trace propagation delay for the associated data signals (<code>DQ</code>). For each port, route the data (<code>DQ</code>) signals for the <code>DQ</code> group ideally on the same layer as the associated <code>QK/QK#</code> and <code>DK/DK#</code> clocks to within ± 10 ps skew of the target clock. Route the <code>mem_ck</code> (<code>CK/CK#</code>) clocks and set as the target trace propagation delays for the address/command signal group. Match the <code>CK/CK#</code> clock to within ± 50 ps of all the <code>DK/DK#</code> clocks for both ports. Route the <code>address/control</code> signal group ideally on the same layer as the <code>mem_ck</code> (<code>CK/CK#</code>) clocks, to within ± 10 ps skew of the <code>mem_ck</code> (<code>CK/CK#</code>) traces.

8.4.4 Package Deskew

You should follow Intel's package deskew guidance.

Related Links

[Package Deskew](#)



9 Intel Stratix 10 EMIF IP for RLD3 RAM 3

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for RLD3 RAM 3.

9.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

9.1.1 Intel Stratix 10 EMIF IP RLD3 RAM 3 Parameters: General

Table 370. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_RLD3_CONFIG_ENUM)

Table 371. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_RLD3_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_RLD3_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_RLD3_USER_REF_CLK_FREQ_MHZ)
<i>continued...</i>	



Display Name	Description
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_RLD3_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_RLD3_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_RLD3_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 372. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 373. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)

Table 374. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)



Table 375. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 376. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

9.1.2 Intel Stratix 10 EMIF IP RLD3 Parameters: Memory

Table 377. Group: Memory / Topology

Display Name	Description
DQ width per device	Specifies number of DQ pins per RLD3 device. Available widths for DQ are x18 and x36. (Identifier: MEM_RLD3_DQ_PER_DEVICE)
Enable DM pins	Indicates whether the interface uses the DM pins. If enabled, one DM pin per write data group is added. (Identifier: MEM_RLD3_DM_EN)
Enable width expansion	Indicates whether to combine two memory devices to double the data bus width. With two devices, the interface supports a width expansion configuration up to 72-bits. For width expansion configuration, the address and control signals are routed to 2 devices. (Identifier: MEM_RLD3_WIDTH_EXPANDED)
Enable depth expansion using twin die package	Indicates whether to combine two RLD3 devices to double the address space, resulting in more density. (Identifier: MEM_RLD3_DEPTH_EXPANDED)
Address width	Number of address pins. (Identifier: MEM_RLD3_ADDR_WIDTH)
Bank address width	Number of bank address pins (Identifier: MEM_RLD3_BANK_ADDR_WIDTH)

Table 378. Group: Memory / Mode Register Settings

Display Name	Description
tRC	Determines the mode register setting that controls the tRC(activate to activate timing parameter). Refer to the tRC table in the memory vendor data sheet. Set the tRC according to the memory speed grade and data latency. (Identifier: MEM_RLD3_T_RC_MODE_ENUM)
Data Latency	Determines the mode register setting that controls the data latency. Sets both READ and WRITE latency (RL and WL) . (Identifier: MEM_RLD3_DATA_LATENCY_MODE_ENUM)
Output drive	Determines the mode register setting that controls the output drive setting. (Identifier: MEM_RLD3_OUTPUT_DRIVE_MODE_ENUM)
ODT	Determines the mode register setting that controls the ODT setting. (Identifier: MEM_RLD3_ODT_MODE_ENUM)
continued...	



Display Name	Description
AREF protocol	Determines the mode register setting that controls the AREF protocol setting. The AUTO REFRESH (AREF) protocol is selected by setting mode register 1. There are two ways in which AREF commands can be issued to the RLD3, the memory controller can either issue bank address-controlled or multibank AREF commands. Multibank refresh protocol allows for the simultaneous refreshing of a row in up to four banks (Identifier: MEM_RLD3_AREF_PROTOCOL_ENUM)
Burst length	Determines the mode register setting that controls the burst length. (Identifier: MEM_RLD3_BL)
Write protocol	Determines the mode register setting that controls the write protocol setting. When multiple bank (dual bank or quad bank) is selected, identical data is written to multiple banks. (Identifier: MEM_RLD3_WRITE_PROTOCOL_ENUM)

9.1.3 Intel Stratix 10 EMIF IP RLD3 Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.

Table 379. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_RLD3_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_RLD3_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_RLD3_DEFAULT_IO)

Table 380. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_RLD3_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_RLD3_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_RLD3_USER_AC_SLEW_RATE_ENUM)



Table 381. Group: FPGA I/O / FPGA I/O Settings / Memory Clock

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_RLD3_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_RLD3_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_RLD3_USER_CK_SLEW_RATE_ENUM)

Table 382. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_RLD3_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_RLD3_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_RLD3_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_RLD3_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_RLD3_USER_STARTING_VREFIN)

Table 383. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_RLD3_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_RLD3_USER_RZQ_IO_STD_ENUM)

9.1.4 Intel Stratix 10 EMIF IP RLD3 IP Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).



Table 384. Group: Mem Timing

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_RLD3_SPEEDBIN_ENUM)
tDS (base)	tDS(base) refers to the setup time for the Data (DQ) bus before the rising edge of the DQS strobe. (Identifier: MEM_RLD3_TDS_PS)
tDS (base) AC level	tDS (base) AC level refers to the voltage level which the data bus must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_RLD3_TDS_AC_MV)
tDH (base)	tDH (base) refers to the hold time for the Data (DQ) bus after the rising edge of CK. (Identifier: MEM_RLD3_TDH_PS)
tDH (base) DC level	tDH (base) DC level refers to the voltage level which the data bus must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_RLD3_TDH_DC_MV)
tQKQ_max	tQKQ_max describes the maximum skew between the read strobe (QK) clock edge to the data bus (DQ/DINV) edge . (Identifier: MEM_RLD3_TQKQ_MAX_PS)
tQH	tQH specifies the output hold time for the DQ/DINV in relation to QK. (Identifier: MEM_RLD3_TQH_CYC)
tCKDK_max	tCKDK_max refers to the maximum skew from the memory clock (CK) to the write strobe (DK) . (Identifier: MEM_RLD3_TCKDK_MAX_CYC)
tCKDK_min	tCKDK_min refers to the minimum skew from the memory clock (CK) to the write strobe (DK) . (Identifier: MEM_RLD3_TCKDK_MIN_CYC)
tCKQK_max	tCKQK_max refers to the maximum skew from the memory clock (CK) to the read strobe (QK) . (Identifier: MEM_RLD3_TCKQK_MAX_PS)
tIS (base)	tIS (base) refers to the setup time for the Address/Command/Control (A) bus to the rising edge of CK. (Identifier: MEM_RLD3_TIS_PS)
tIS (base) AC level	tIS (base) AC level refers to the voltage level which the address/command signal must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_RLD3_TIS_AC_MV)
tIH (base)	tIH (base) refers to the hold time for the Address/Command (A) bus after the rising edge of CK. Depending on what AC level the user has chosen for a design, the hold margin can vary (this variance will be automatically determined when the user chooses the " tIH (base) AC level "). (Identifier: MEM_RLD3_TIH_PS)
tIH (base) DC level	tIH (base) DC level refers to the voltage level which the address/command signal must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_RLD3_TIH_DC_MV)

9.1.5 Intel Stratix 10 EMIF IP RLD3 Parameters: Board



Table 385. Group: Board / Intersymbol Interference/Crosstalk

Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation using IBIS models and Hyperlynx)*, and manually enter values based on your simulation results, instead of using the default values.</i> (Identifier: BOARD_RLD3_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to ISI and crosstalk effects. The number to be entered is the total loss of margin on both the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_RLD3_USER_AC_ISI_NS)
QK/QK# ISI/crosstalk	QK/QK# ISI/crosstalk describes the reduction of the read data window due to intersymbol interference and crosstalk effects on the QK/QK# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on both the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_RLD3_USER_RCLK_ISI_NS)
Read DQ ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQ signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold side (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_RLD3_USER_RDATA_ISI_NS)
DK/DK# ISI/crosstalk	DK/DK# ISI/crosstalk describes the reduction of the write data window due to intersymbol interference and crosstalk effects on the DK/DK# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold side (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_RLD3_USER_WCLK_ISI_NS)
Write DQ ISI/crosstalk	The reduction of the write data window due to ISI and crosstalk effects on the DQ signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold side (measured loss on the setup side + measured loss on the hold side) . Refer to the EMIF Simulation Guidance wiki page for additional information. (Identifier: BOARD_RLD3_USER_WDATA_ISI_NS)

Table 386. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (QK group)	If you are compensating for package skew on the QK bus in the board layout (hence checking the box here), please include package skew in calculating the following board skew parameters . (Identifier: BOARD_RLD3_IS_SKEW_WITHIN_QK_DESKEWED)
Maximum board skew within QK group	Maximum board skew within QK group refers to the largest skew between all DQ and DM pins in a QK group. This value can affect the read capture and write margins. (Identifier: BOARD_RLD3_BRD_SKEW_WITHIN_QK_NS)
Maximum system skew within QK group	The largest skew between all DQ and DM pins in a QK group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_RLD3_PKG_BRD_SKEW_WITHIN_QK_NS)
continued...	



Display Name	Description
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_RLD3_IS_SKEW_WITHIN_AC_DESKEWED)
Maximum board skew within address/command bus	The largest skew between the address and command signals. Enter the board skew only; package skew is calculated automatically, based on the memory interface configuration, and added to this value. (Identifier: BOARD_RLD3_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	Maximum system skew within address/command bus refers to the largest skew between the address and command signals. (Identifier: BOARD_RLD3_PKG_BRD_SKEW_WITHIN_AC_NS)
Average delay difference between DK and CK	This parameter describes the average delay difference between the DK signals and the CK signal, calculated by averaging the longest and smallest DK trace delay minus the CK trace delay. Positive values represent DK signals that are longer than CK signals and negative values represent DK signals that are shorter than CK signals. (Identifier: BOARD_RLD3_DK_TO_CK_SKEW_NS)
Maximum delay difference between devices	This parameter describes the largest propagation delay on the DQ signals between ranks. For example, in a two-rank configuration where devices are placed in series, there is an extra propagation delay for DQ signals going to and coming back from the furthest device compared to the nearest device. <i>This parameter is only applicable when there is more than one rank.</i> (Identifier: BOARD_RLD3_SKEW_BETWEEN_DIMMS_NS)
Maximum skew between DK groups	This parameter describes the largest skew between DK signals in different DK groups. (Identifier: BOARD_RLD3_SKEW_BETWEEN_DK_NS)
Average delay difference between address/command and CK	The average delay difference between the address/command signals and the CK signal, calculated by averaging the longest and smallest address/command signal trace delay minus the maximum CK trace delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. (Identifier: BOARD_RLD3_AC_TO_CK_SKEW_NS)
Maximum CK delay to device	The maximum CK delay to device refers to the delay of the longest CK trace from the FPGA to any device. (Identifier: BOARD_RLD3_MAX_CK_DELAY_NS)
Maximum DK delay to device	The maximum DK delay to device refers to the delay of the longest DK trace from the FPGA to any device. (Identifier: BOARD_RLD3_MAX_DK_DELAY_NS)

9.1.6 Intel Stratix 10 EMIF IP RLD3 Parameters: Diagnostics

Table 387. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process. Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero.
continued...	



Display Name	Description
	<i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations. Abstract PHY is supported with skip calibration.</i> (Identifier: DIAG_RLD3_SIM_CAL_MODE_ENUM)
Abstract phy for fast simulation	Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration . (Identifier: DIAG_RLD3_ABSTRACT_PHY)

Table 388. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic. If you set this parameter to " Disabled ", no debug features are enabled. If you set this parameter to " Export ", an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP core to it, or connect it to the cal_debug_out interface of another EMIF core. If you select " Add EMIF Debug Interface ", an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit. <i>Only one EMIF debug interface should be instantiated per I/O column. You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first.</i> (Identifier: DIAG_RLD3_EXPORT_SEQ_AVALON_SLAVE)
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies that the IP export an Avalon-MM master interface (cal_debug_out) which can connect to the cal_debug interface of other EMIF cores residing in the same I/O column. This parameter applies only if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to the <i>Debugging Multiple EMIFs</i> wiki page for more information about debugging multiple EMIFs. (Identifier: DIAG_RLD3_EXPORT_SEQ_AVALON_MASTER)
Interface ID	Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled , the interface ID is unused. (Identifier: DIAG_RLD3_INTERFACE_ID)
Use Soft NIOS Processor for On-Chip Debug	Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port . <i>Only one interface in a column can activate this option.</i> (Identifier: DIAG_SOFT_NIOS_MODE)

Table 389. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the " Core clocks sharing " parameter in the " General " tab to " Master " or " Slave ". (Identifier: DIAG_RLD3_EX_DESIGN_NUM_OF_SLAVES)
Enable In-System-Sources-and-Probes	Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i> . This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_RLD3_EX_DESIGN_ISSP_EN)



Table 390. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_RLD3_USE_TG_AVL_2)
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_RLD3_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the altera_emif_avl_tg_2_tb.sv file. (Identifier: DIAG_RLD3_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_RLD3_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_RLD3_BYPASS_STRESS_STAGE)
Export Traffic Generator 2.0 configuration interface	Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)

Table 391. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_RLD3_EFFICIENCY_MONITOR)

Table 392. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

9.1.7 Intel Stratix 10 EMIF IP RLD3 Parameters: Example Designs




Table 393. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization. After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_RLD3_SEL_DESIGN)

Table 394. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created.</i> Instead, the output directory will contain the ed_sim.qsys file which holds Qsys details of the simulation example design, and a make_sim_design.tcl file with other corresponding tcl files. You can run make_sim_design.tcl from a command line to generate the simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_RLD3_GEN_SIM)
Synthesis	Specifies that the 'Generate Example Design' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created.</i> Instead, the output directory will contain the ed_synth.qsys file which holds Qsys details of the synthesis example design, and a make_qii_design.tcl script with other corresponding tcl files. You can run make_qii_design.tcl from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_RLD3_GEN_SYNTH)

Table 395. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_RLD3_HDL_FORMAT)

Table 396. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select 'none' from the 'Select board' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_RLD3_TARGET_DEV_KIT)



9.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

9.2.1 Equations for RLD RAM 3 Board Skew Parameters

Table 397. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum CK delay to device	<p>The delay of the longest CK trace from the FPGA to any device.</p> $\max_n (CK_n PathDelay)$ <p>where n is the number of memory clocks. For example, the maximum CK delay for two pairs of memory clocks is expressed by the following equation:</p> $\max_2 (CK_1 PathDelay, CK_2 PathDelay)$
Maximum DK delay to device	<p>The delay of the longest DK trace from the FPGA to any device.</p> $\max_n (DK_n PathDelay)$ <p>where n is the number of DK. For example, the maximum DK delay for two DK is expressed by the following equation:</p> $\max_2 (DK_1 PathDelay, DK_2 PathDelay)$
Average delay difference between DK and CK	<p>The average delay difference between the DK signals and the CK signal, calculated by averaging the longest and smallest DK delay minus the CK delay. Positive values represent DK signals that are longer than CK signals and negative values represent DK signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the DK signals to have appropriate setup and hold margins.</p> $\frac{\max_{n, m} (CK_n PathDelay - DK_m PathDelay) + \min_{n, m} (CK_n PathDelay - DK_m PathDelay)}{2}$ <p>where n is the number of memory clocks and m is the number of DK.</p>
Maximum system skew within address/command bus	$(MaxAC - MinAC)$ <p>The largest skew between the address and command signals. Enter combined board and package skew.</p>
Average delay difference between address/command and CK	<p>The average delay difference between the address and command signals and the CK signal, calculated by averaging the longest and smallest Address/Command signal delay minus the CK delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the address and command signals to have appropriate setup and hold margins.</p> $\sum_{n=1}^n \left(\left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - CK_n PathDelay \right)$
Maximum board skew within QK group	<p>The largest skew between all DQ and DM pins in a QK group. Enter your board skew only. Package skew will be calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins.</p> $\max_n (\max DQ_n - \min DQ_n)$ <p>where n is the number of DQ.</p>
Maximum skew between DK groups	<p>The largest skew between DK signals in different DK groups.</p>

continued...



Parameter	Description/Equation
	$\max_n (\max DK_n - \min DK_n)$ <p>where n is the number of DQ.</p>

9.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.

9.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.

Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

9.3.1.1 Estimating Pin Requirements

You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:



1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.

You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)

9.3.1.2 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.

Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

9.3.1.3 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.

9.3.1.4 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

9.3.1.5 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:

- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.

9.3.1.6 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.



The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the Bank Number and Index within I/O Bank values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the Bank Number value identifies the I/O column, while the letter represents the I/O bank.
- The Index within I/O Bank value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its Index within I/O Bank number (if it is an even number), or by subtracting one from its Index within I/O Bank number (if it is an odd number).

For example, a physical pin with a Bank Number of 2M and Index within I/O Bank of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an Index within I/O Bank of 23 and Bank Number of 2M.

9.3.1.6.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.

If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version number/<synth/sim>/<variation_name>_altera_emif_arch_nd_version number_<unique ID>_readme.txt` file, w+ted with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (.qip), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLD RAMx , QDRx and LPDDR3 are not supported with HPS.



Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the

```
<variation_name>/altera_emif_arch_nd_<version>/<synth/  
sim>/
```

```
<variation_name>_altera_emif_arch_nd_<version>_<unique  
ID>_readme.txt file after you have generated your IP.
```

7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.



- Note:* a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.
- b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).

10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:
- There must be an even number of x4 groups in an external memory interface.
 - DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group X+1 must be in the same I/O lane, where X is an even number.

11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Extreme, and RLD RAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLD RAM 3, x36 device, DQ[8:0] and DQ[26:18] are referenced to DK0/DK0#, and DQ[17:9] and DQ[35:27] are referenced to DK1/DK1#.

12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.

13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.

Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLD RAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (global_reset_n) for each individual interface all come from the same input pin or signal.

I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.

CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)



The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the RZQ pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.

9.3.1.6.2 RLD RAM 3 Commands and Addresses

The CK and CK# signals clock the commands and addresses into the memory devices.

These pins operate at single data rate using only one clock edge. RLD RAM 3 supports both non-multiplexed and multiplexed addressing. Multiplexed addressing allows you to save a few user I/O pins while non-multiplexed addressing allows you to send the address signal within one clock cycle instead of two clock cycles. CS#, REF#, and WE# pins are input commands to the RLD RAM 3 device.

The commands and addresses must meet the memory address and command setup (t_{AS}, t_{CS}) and hold (t_{AH}, t_{CH}) time requirements.

Note: The RLD RAM 3 external memory interface IP does not support multiplexed addressing.

9.3.1.6.3 RLD RAM 3 Clock Signals

RLD RAM 3 devices use CK and CK# signals to clock the command and address bus in single data rate (SDR). There is one pair of CK and CK# pins per RLD RAM 3 device.

Instead of a strobe, RLD RAM 3 devices use two sets of free-running differential clocks to accompany the data. The DK and DK# clocks are the differential input data clocks used during writes while the QK or QK# clocks are the output data clocks used during

reads. Even though QK and QK# signals are not differential signals according to the RLD RAM 3 data sheet, Micron treats these signals as such for their testing and characterization. Each pair of DK and DK#, or QK and QK# clocks are associated with either 9 or 18 data bits.

The exact clock-data relationships are as follows:

- RLD RAM 3: For $\times 36$ data bus width configuration, there are 18 data bits associated with each pair of write clocks. There are 9 data bits associated with each pair of read clocks. So, there are two pairs of DK and DK# pins and four pairs of QK and QK# pins.
- RLD RAM 3: For $\times 18$ data bus width configuration, there are 9 data bits per one pair of write clocks and nine data bits per one pair of read clocks. So, there are two pairs of DK and DK# pins, and two pairs of QK and QK# pins
- RLD RAM 3: RLD RAM 3 does not have the $\times 9$ data bus width configuration.

There are t_{CKDK} timing requirements for skew between CK and DK or CK# and DK#.

For RLD RAM 3, because of the loads on these I/O pins, the maximum frequency you can achieve depends on the number of memory devices you are connecting to the Intel device. Perform SPICE or IBIS simulations to analyze the loading effects of the pin-pair on multiple RLD RAM 3 devices.

9.3.1.6.4 RLD RAM 3 Data, DM and QVLD Signals

The read data is edge-aligned with the QK or QK# clocks while the write data is center-aligned with the DK and DK# clocks (see the following figures). The memory controller shifts the DK and DK# signals to center align the DQ and DK or DK# signals during a write. It also shifts the QK signal during a read, so that the read data (DQ signals) and QK clock is center-aligned at the capture register.

Intel devices use dedicated DQS phase-shift circuitry to shift the incoming QK signal during reads and use a PLL to center-align the DK and DK# signals with respect to the DQ signals during writes.

Figure 78. Edge-aligned DQ and QK Relationship During RLD RAM 3 Read







Figure 79. Center-aligned DQ and DK Relationship During RLD RAM 3 Write



For RLD RAM 3, data mask (DM) pins are used only during a write. The memory controller drives the DM signal low when the write is valid and drives it high to mask the DQ signals.

For RLD RAM 3, there are two DM pins per memory device. DM0 is used to mask the lower byte for the x18 device and (DQ[8:0], DQ[26:18]) for the x36 device. DM1 is used to mask the upper byte for the x18 device and (DQ[17:9], DQ[35:27]) for the x36 device.

The DM timing requirements at the input to the memory device are identical to those for DQ data. The DDR registers, clocked by the write clock, create the DM signals. This reduces any skew between the DQ and DM signals.

The RLD RAM 3 device's setup time (t_{DS}) and hold (t_{DH}) time for the write DQ and DM pins are relative to the edges of the DK or DK# clocks. The DK and DK# signals are generated on the positive edge of system clock, so that the positive edge of CK or CK# is aligned with the positive edge of DK or DK# respectively to meet the tCKDK requirement. The DQ and DM signals are clocked using a shifted clock so that the edges of DK or DK# are center-aligned with respect to the DQ and DM signals when they arrive at the RLD RAM 3 device.

The clocks, data, and DM board trace lengths should be tightly matched to minimize the skew in the arrival time of these signals.

RLD RAM 3 devices also have a QVLD pin indicating valid read data. The QVLD signal is edge-aligned with QK or QK# and is high approximately half a clock cycle before data is output from the memory.

Note: The Intel FPGA external memory interface IP does not use the QVLD signal.

9.3.1.7 Resource Sharing Guidelines (Multiple Interfaces)

In Intel Cyclone 10/Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.

PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.

Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `pll_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note:

You can place the `pll_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.

Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.

Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.



9.4 RLD RAM 3 Board Design Guidelines

The following topics provide layout guidelines for you to improve your system's signal integrity and to successfully implement an RLD RAM 3 interface.

The following topics focus on the following key factors that affect signal integrity:

- I/O standards
- RLD RAM 3 configurations
- Signal terminations
- Printed circuit board (PCB) layout guidelines

I/O Standards

RLD RAM 3 interface signals use the following JEDEC I/O signalling standards: HSTL 1.2 V and SSTL-12.


The RLD RAM 3 IP defaults to HSTL 1.2 V Class I outputs and HSTL 1.2 V inputs.

9.4.1 RLD RAM 3 Configurations

The Intel Stratix 10 EMIF IP for RLD RAM 3 supports interfaces for CIO RLD RAM 3 with one or two devices. With two devices, the interface supports a width expansion configuration up to 72-bits. The termination and layout principles for SIO RLD RAM 3 interfaces are similar to CIO RLD RAM 3, except that SIO RLD RAM 3 interfaces have unidirectional data buses.

The following figure shows the main signal connections between the FPGA and a single CIO RLD RAM 3 component.

Figure 80. Configuration with a Single CIO RLD RAM 3 Component



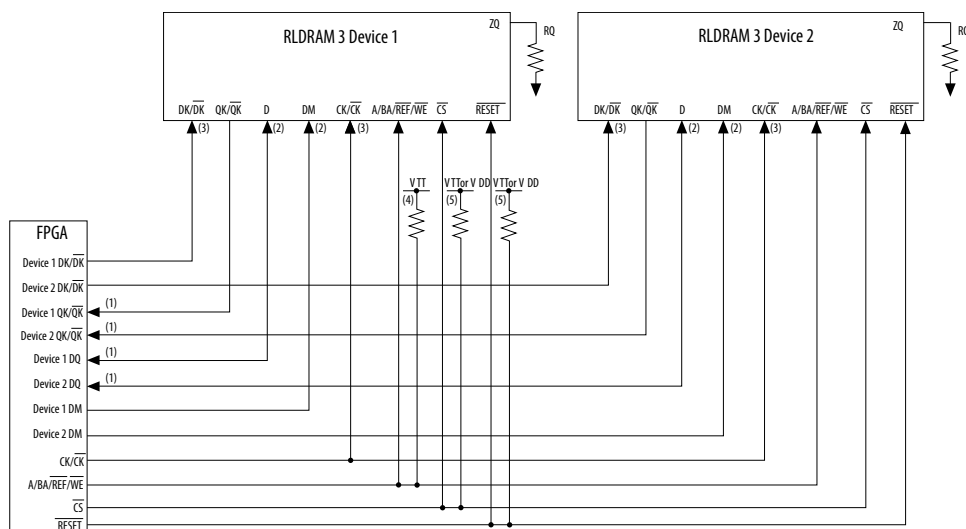
Notes to Figure:

1. Use external differential termination on CK/CK#.
2. Use FPGA parallel on-chip termination (OCT) for terminating QK/QK# and DQ on reads.
3. Use RLD RAM 3 component on-die termination (ODT) for terminating DQ, DM, and DK, DK# on writes.
4. Use external discrete termination with fly-by placement to avoid stubs.
5. Use external discrete termination for this signal, as shown for REF.
6. Use external discrete termination, as shown for REF, but you may require a pull-up resistor to VDD as an alternative option. Refer to the RLD RAM 3 device data sheet for more information about RLD RAM 3 power-up sequencing.

The following figure shows the main signal connections between the FPGA and two CIO RLD RAM 3 components in a width expansion configuration.



Figure 81. Configuration with Two CIO RLD RAM 3 Components in a Width Expansion Configuration



Notes to Figure:

1. Use FPGA parallel OCT for terminating QK/QK# and DQ on reads.
2. Use RLD RAM 3 component ODT for terminating DQ, DM, and DK on writes.
3. Use external dual 200 Ω differential termination.
4. Use external discrete termination at the trace split of the balanced T or Y topology.
5. Use external discrete termination at the trace split of the balanced T or Y topology, but you may require a pull-up resistor to VDD as an alternative option. Refer to the RLD RAM 3 device data sheet for more information about RLD RAM 3 power-up sequencing.

9.4.2 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.

Table 398. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-Junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)

9.4.3 RLD RAM 3 Layout Guidelines

The following table lists the RLD RAM 3 general routing layout guidelines. These guidelines apply to Intel Stratix 10 devices.



Table 399. RLD RAM 3 Layout Guidelines

Parameter	Guidelines
General Routing	<ul style="list-style-type: none"> If you must route signals of the same net group on different layers with the same impedance characteristic, simulate your worst case PCB trace tolerances to ascertain actual propagation delay differences. Typical layer to layer trace delay variations are of 15 ps/inch order. Avoid T-junctions greater than 150 ps. Match all signals within a given DQ group with a maximum skew of ± 10 ps and route on the same layer.
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to under 150 ps. These signals should maintain a 10-mil (0.254 mm) spacing from other nets. Clocks should maintain a length-matching between clock pairs of ± 5 ps. Differential clocks should maintain a length-matching between P and N signals of ± 2 ps. Space between different clock pairs should be at least three times the space between the traces of a differential pair.
Address and Command Routing	<ul style="list-style-type: none"> To minimize crosstalk, route address, bank address, and command signals on a different layer than the data and data mask signals. Do not route the differential clock signals close to the address signals. Keep the distance from the pin on the RLD RAM 3 component to the stub termination resistor (V_{TT}) to less than 50 ps for the address/command signal group. Keep the distance from the pin on the RLD RAM 3 component to the fly-by termination resistor (V_{TT}) to less than 100 ps for the address/command signal group.
External Memory Routing Rules	<ul style="list-style-type: none"> Apply the following parallelism rules for the RLD RAM 3 data/address/command groups: <ul style="list-style-type: none"> 4 mils for parallel runs < 0.1 inch (approximately 1\times spacing relative to plane distance). 5 mils for parallel runs < 0.5 inch (approximately 1\times spacing relative to plane distance). 10 mils for parallel runs between 0.5 and 1.0 inches (approximately 2\times spacing relative to plane distance). 15 mils for parallel runs between 1.0 and 3.3 inch (approximately 3\times spacing relative to plane distance).
Maximum Trace Length	<ul style="list-style-type: none"> Keep the maximum trace length of all signals from the FPGA to the RLD RAM 3 components to 600 ps.
Trace Matching Guidance	<p>The following layout approach is recommended, based on the preceding guidelines:</p> <ol style="list-style-type: none"> If the RLD RAM 3 interface has multiple DQ groups ($\times 18$ or $\times 36$ RLD RAM 3 component or width expansion configuration), match all the DK/DK# and QK, QK# clocks as tightly as possible to optimize the timing margins in your design. Route the DK/DK# write clock and QK/QK# read clock associated with a DQ group on the same PCB layer. Match these clock pairs to within ± 5 ps. Set the DK/DK# or QK/QK# clock as the target trace propagation delay for the associated data and data mask signals. Route the data and data mask signals for the DQ group ideally on the same layer as the associated QK/QK# and DK/DK# clocks to within ± 10 ps skew of the target clock. Route the CK/CK# clocks and set as the target trace propagation delays for the address/command signal group. Match the CK/CK# clock to within ± 50 ps of all the DK/DK# clocks. Route the address/control signal group (address, bank address, CS, WE, and REF) ideally on the same layer as the CK/CK# clocks, to within ± 20 ps skew of the CK/CK# traces.

continued...

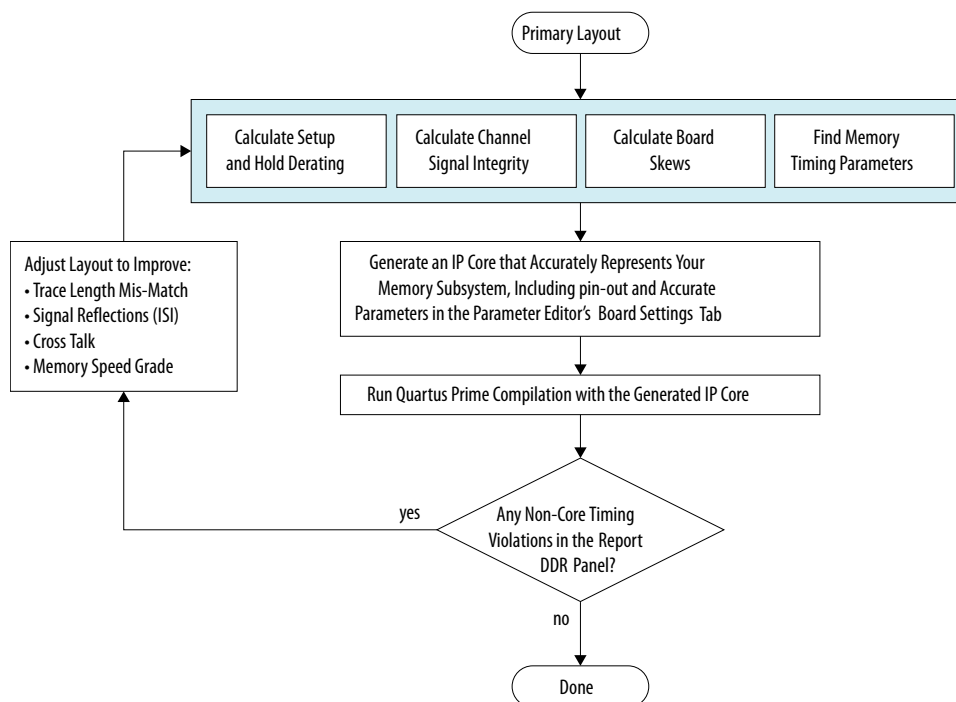
Parameter	Guidelines
	<p><i>Note:</i> It is important to match the delays of CK vs. DK, and CK vs. Addr-Cmd as much as possible.</p> <p>This layout approach provides a good starting point for a design requirement of the highest clock frequency supported for the RLD RAM 3 interface.</p>

9.4.4 Layout Approach

For all practical purposes, you can regard the Timing Analyzer report on your memory interface as definitive for a given set of memory and board timing parameters.

You will find timing under **Report DDR** in the Timing Analyzer and on the **Timing Analysis** tab in the parameter editor.

The following flowchart illustrates the recommended process to follow during the board design phase, to determine timing margin and make iterative improvements to your design.



Board Skew

For information on calculating board skew parameters, refer to *Board Skew Equations*, in this chapter.

The Board Skew Parameter Tool is an interactive tool that can help you calculate board skew parameters if you know the absolute delay values for all the memory related traces.

Memory Timing Parameters

For information on the memory timing parameters to be entered into the parameter editor, refer to the datasheet for your external memory device.



Related Links

[Board Skew Parameter Tool](#)

9.4.5 Package Deskew

You should follow Intel's package deskew guidance.

Related Links

[Package Deskew](#)



10 Intel Stratix 10 EMIF IP for LPDDR3

This chapter contains IP parameter descriptions, board skew equations, pin planning information, and board design guidance for Intel Stratix 10 external memory interfaces for LPDDR3.

10.1 Parameter Descriptions

The following topics describe the parameters available on each tab of the IP parameter editor, which you can use to configure your IP.

10.1.1 Intel Stratix 10 EMIF IP LPDDR3 Parameters: General

Table 400. Group: General / Interface

Display Name	Description
Configuration	Specifies the configuration of the memory interface. The available options depend on the protocol in use. Options include Hard PHY and Hard Controller , Hard PHY and Soft Controller , or Hard PHY only . If you select Hard PHY only , the AFI interface is exported to allow connection of a custom memory controller or third-party IP. (Identifier: PHY_LPDDR3_CONFIG_ENUM)

Table 401. Group: General / Clocks

Display Name	Description
Memory clock frequency	Specifies the operating frequency of the memory interface in MHz. If you change the memory frequency, you should update the memory latency parameters on the Memory tab and the memory timing parameters on the Mem Timing tab. (Identifier: PHY_LPDDR3_MEM_CLK_FREQ_MHZ)
Use recommended PLL reference clock frequency	Specifies that the PLL reference clock frequency is automatically calculated for best performance. <i>If you want to specify a different PLL reference clock frequency, uncheck the check box for this parameter.</i> (Identifier: PHY_LPDDR3_DEFAULT_REF_CLK_FREQ)
PLL reference clock frequency	This parameter tells the IP what PLL reference clock frequency the user will supply. Users must select a valid PLL reference clock frequency from the list. The values in the list can change when the memory interface frequency changes and/or the clock rate of user logic changes. It is recommended to use the fastest possible PLL reference clock frequency because it leads to better jitter performance. Selection is required only if the user does not check the "Use recommended PLL reference clock frequency" option. (Identifier: PHY_LPDDR3_USER_REF_CLK_FREQ_MHZ)
<i>continued...</i>	



Display Name	Description
PLL reference clock jitter	Specifies the peak-to-peak jitter on the PLL reference clock source. The clock source of the PLL reference clock must meet or exceed the following jitter requirements: 10ps peak to peak, or 1.42ps RMS at 1e-12 BER, 1.22ps at 1e-16 BER. (Identifier: PHY_LPDDR3_REF_CLK_JITTER_PS)
Clock rate of user logic	Specifies the relationship between the user logic clock frequency and the memory clock frequency. For example, if the memory clock sent from the FPGA to the memory device is toggling at 800MHz, a quarter-rate interface means that the user logic in the FPGA runs at 200MHz. (Identifier: PHY_LPDDR3_RATE_ENUM)
Core clocks sharing	<p>When a design contains multiple interfaces of the same protocol, rate, frequency, and PLL reference clock source, they can share a common set of core clock domains. By sharing core clock domains, they reduce clock network usage and avoid clock synchronization logic between the interfaces.</p> <p>To share core clocks, denote one of the interfaces as "Master", and the remaining interfaces as "Slave". In the RTL, connect the <code>clks_sharing_master_out</code> signal from the master interface to the <code>clks_sharing_slave_in</code> signal of all the slave interfaces.</p> <p>Both master and slave interfaces still expose their own output clock ports in the RTL (for example, <code>emif_usr_clk</code>, <code>afi_clk</code>), but the physical signals are equivalent, hence it does not matter whether a clock port from a master or a slave is used. <i>As the combined width of all interfaces sharing the same core clock increases, you may encounter timing closure difficulty for transfers between the FPGA core and the periphery.</i> (Identifier: PHY_LPDDR3_CORE_CLKS_SHARING_ENUM)</p>
Specify additional core clocks based on existing PLL	Displays additional parameters allowing you to create additional output clocks based on the existing PLL. This parameter provides an alternative clock-generation mechanism for when your design exhausts available PLL resources . The additional output clocks that you create can be fed into the core. Clock signals created with this parameter are synchronous to each other, but asynchronous to the memory interface core clock domains (such as <code>emif_usr_clk</code> or <code>afi_clk</code>). <i>You must follow proper clock-domain-crossing techniques when transferring data between clock domains.</i> (Identifier: PLL_ADD_EXTRA_CLKS)

Table 402. Group: General / Clocks / Additional Core Clocks

Display Name	Description
Number of additional core clocks	Specifies the number of additional output clocks to create from the PLL. (Identifier: PLL_USER_NUM_OF_EXTRA_CLKS)

Table 403. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_0

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_5)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_5)

Table 404. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_1

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_6)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_6)



Table 405. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_2

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_7)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_7)

Table 406. Group: General / Clocks / Additional Core Clocks / pll_extra_clk_3

Display Name	Description
Frequency	Specifies the frequency of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_FREQ_MHZ_GUI_8)
Phase shift	Specifies the phase shift of the core clock signal. (Identifier: PLL_EXTRA_CLK_ACTUAL_PHASE_PS_GUI_8)

10.1.2 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Memory

Table 407. Group: Memory / Topology

Display Name	Description
DQ width	Total number of DQ pins in the interface. (Identifier: MEM_LPDDR3_DQ_WIDTH)
Number of clocks	Number of CK/CK# clock pairs exposed by the memory interface. (Identifier: MEM_LPDDR3_CK_WIDTH)
Number of chip selects	Total number of chip selects in the interface. (Identifier: MEM_LPDDR3_DISCRETE_CS_WIDTH)
Row address width	The number of row address bits. (Identifier: MEM_LPDDR3_ROW_ADDR_WIDTH)
Column address width	The number of column address bits. (Identifier: MEM_LPDDR3_COL_ADDR_WIDTH)
Bank address width	The number of bank address bits. (Identifier: MEM_LPDDR3_BANK_ADDR_WIDTH)
Enable DM pins	Indicates whether interface uses data mask (DM) pins. This feature allows specified portions of the data bus to be written to memory (not available in x4 mode). One DM pin exists per DQS group. (Identifier: MEM_LPDDR3_DM_EN)

Table 408. Group: Memory / Latency and Burst

Display Name	Description
Data latency	Determines the mode register setting that controls the data latency. Sets both READ and WRITE latency (RL and WL). (Identifier: MEM_LPDDR3_DATA_LATENCY)
Burst length	Burst length of the memory device. (Identifier: MEM_LPDDR3_BL)

10.1.3 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Mem I/O



Table 409. Group: Mem I/O / Memory I/O Settings

Display Name	Description
Output drive strength setting	Specifies the output driver impedance setting at the memory device. (Identifier: MEM_LPDDR3_DRV_STR)
DQ ODT	The ODT setting for the DQ pins during writes. (Identifier: MEM_LPDDR3_DQODT)
Power down ODT	Turn on turn off ODT during power down. (Identifier: MEM_LPDDR3_PDODT)

Table 410. Group: Mem I/O / ODT Activation

Display Name	Description
Use Default ODT Assertion Tables	Enables the default ODT assertion pattern as determined from vendor guidelines. These settings are provided as a default only; you should simulate your memory interface to determine the optimal ODT settings and assertion patterns. (Identifier: MEM_LPDDR3_USE_DEFAULT_ODT)

10.1.4 Intel Stratix 10 EMIF IP LPDDR3 Parameters: FPGA I/O

You should use Hyperlynx* or similar simulators to determine the best settings for your board. Refer to the EMIF Simulation Guidance wiki page for additional information.

Table 411. Group: FPGA I/O / FPGA I/O Settings

Display Name	Description
Voltage	The voltage level for the I/O pins driving the signals between the memory device and the FPGA memory interface. (Identifier: PHY_LPDDR3_IO_VOLTAGE)
Periodic OCT re-calibration	Specifies that the system periodically recalibrate on-chip termination (OCT) to minimize variations in termination value caused by changing operating conditions (such as changes in temperature). By recalibrating OCT, I/O timing margins are improved. When enabled, this parameter causes the PHY to halt user traffic about every 0.5 seconds for about 1900 memory clock cycles, to perform OCT recalibration. Efficiency is reduced by about 1% when this option is enabled. (Identifier: PHY_LPDDR3_USER_PERIODIC_OCT_RECAL_ENUM)
Use default I/O settings	Specifies that a legal set of I/O settings are automatically selected. The default I/O settings are not necessarily optimized for a specific board. <i>To achieve optimal signal integrity, perform I/O simulations with IBIS models and enter the I/O settings manually, based on simulation results.</i> (Identifier: PHY_LPDDR3_DEFAULT_IO)



Table 412. Group: FPGA I/O / FPGA I/O Settings / Address/Command

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the address/command pins of the memory interface. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_LPDDR3_USER_AC_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_LPDDR3_USER_AC_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_LPDDR3_USER_AC_SLEW_RATE_ENUM)

Table 413. Group: FPGA I/O / FPGA I/O Settings / Memory Clock

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the memory clock pins. The selected I/O standard configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_LPDDR3_USER_CK_IO_STD_ENUM)
Output mode	This parameter allows you to change the current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_LPDDR3_USER_CK_MODE_ENUM)
Slew rate	Specifies the slew rate of the address/command output pins. The slew rate (or edge rate) describes how quickly the signal can transition, measured in voltage per unit time. <i>Perform board simulations to determine the slew rate that provides the best eye opening for the address and command signals.</i> (Identifier: PHY_LPDDR3_USER_CK_SLEW_RATE_ENUM)

Table 414. Group: FPGA I/O / FPGA I/O Settings / Data Bus

Display Name	Description
I/O standard	Specifies the I/O electrical standard for the data and data clock/strobe pins of the memory interface. The selected I/O standard option configures the circuit within the I/O buffer to match the industry standard. (Identifier: PHY_LPDDR3_USER_DATA_IO_STD_ENUM)
Output mode	This parameter allows you to change the output current drive strength or termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_LPDDR3_USER_DATA_OUT_MODE_ENUM)
Input mode	This parameter allows you to change the input termination settings for the selected I/O standard. <i>Perform board simulation with IBIS models to determine the best settings for your design.</i> (Identifier: PHY_LPDDR3_USER_DATA_IN_MODE_ENUM)
Use recommended initial Vrefin	Specifies that the initial Vrefin setting is calculated automatically, to a reasonable value based on termination settings. (Identifier: PHY_LPDDR3_USER_AUTO_STARTING_VREFIN_EN)
Initial Vrefin	Specifies the initial value for the reference voltage on the data pins(Vrefin) . This value is entered as a percentage of the supply voltage level on the I/O pins. The specified value serves as a starting point and may be overridden by calibration to provide better timing margins. If you choose to skip Vref calibration (Diagnostics tab) , this is the value that is used as the Vref for the interface. (Identifier: PHY_LPDDR3_USER_STARTING_VREFIN)

**Table 415. Group: FPGA I/O / FPGA I/O Settings / PHY Inputs**

Display Name	Description
PLL reference clock I/O standard	Specifies the I/O standard for the PLL reference clock of the memory interface. (Identifier: PHY_LPDDR3_USER_PLL_REF_CLK_IO_STD_ENUM)
RZQ I/O standard	Specifies the I/O standard for the RZQ pin used in the memory interface. (Identifier: PHY_LPDDR3_USER_RZQ_IO_STD_ENUM)

10.1.5 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Mem Timing

These parameters should be read from the table in the datasheet associated with the speed bin of the memory device (not necessarily the frequency at which the interface is running).

Table 416. Group: Mem Timing / Parameters dependent on Speed Bin

Display Name	Description
Speed bin	The speed grade of the memory device used. This parameter refers to the maximum rate at which the memory device is specified to run. (Identifier: MEM_LPDDR3_SPEEDBIN_ENUM)
tISCA (base)	Address and control setup to CK clock rise (Identifier: MEM_LPDDR3_TIS_PS)
tISCA (base) AC level	AC level of tIS (base) for derating purpose (Identifier: MEM_LPDDR3_TIS_AC_MV)
tIHCA (base)	Address and control hold after CK clock rise (Identifier: MEM_LPDDR3_TIH_PS)
tIHCA (base) DC level	DC level of tIH (base) for derating purpose (Identifier: MEM_LPDDR3_TIH_DC_MV)
tDS (base)	tDS(base) refers to the setup time for the Data (DQ) bus before the rising edge of the DQS strobe. (Identifier: MEM_LPDDR3_TDS_PS)
tDS (base) AC level	tDS (base) AC level refers to the voltage level which the data bus must cross and remain above during the setup margin window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire setup period. (Identifier: MEM_LPDDR3_TDS_AC_MV)
tDH (base)	tDH (base) refers to the hold time for the Data (DQ) bus after the rising edge of CK. (Identifier: MEM_LPDDR3_TDH_PS)
tDH (base) DC level	tDH (base) DC level refers to the voltage level which the data bus must not cross during the hold window . The signal is considered stable only if it remains above this voltage level (for a logic 1) or below this voltage level (for a logic 0) for the entire hold period. (Identifier: MEM_LPDDR3_TDH_DC_MV)
tDQSQ	tDQSQ describes the latest valid transition of the associated DQ pins for a READ . tDQSQ specifically refers to the DQS, DQS# to DQ skew. It is the length of time between the DQS, DQS# crossing to the last valid transition of the slowest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_LPDDR3_TDQSQ_PS)
tQH	tQH specifies the output hold time for the DQ in relation to DQS, DQS# . It is the length of time between the DQS, DQS# crossing to the earliest invalid transition of the fastest DQ pin in the DQ group associated with that DQS strobe. (Identifier: MEM_LPDDR3_TQH_CYC)
tDQSCKDL	Absolute difference between any two tDQSCK measurements (within a byte lane) within a contiguous sequence of bursts within a 32ms rolling window (Identifier: MEM_LPDDR3_TDQSCKDL)
<i>continued...</i>	



Display Name	Description
tDQSS (max)	First latching edge of DQS to associated clock edge (percentage of tCK) (Identifier: MEM_LPDDR3_TDQSS_CYC)
tQSH	tQSH refers to the differential High Pulse Width, which is measured as a percentage of tCK. It is the time during which the DQS is high for a read . (Identifier: MEM_LPDDR3_TQSH_CYC)
tDSH	tDSH specifies the write DQS hold time . This is the time difference between the rising CK edge and the falling edge of DQS, measured as a percentage of tCK. (Identifier: MEM_LPDDR3_TDSH_CYC)
tWLS	tWLS describes the write leveling setup time . It is measured from the rising edge of CK to the rising edge of DQS. (Identifier: MEM_LPDDR3_TWLS_PS)
tWLH	tWLH describes the write leveling hold time . It is measured from the rising edge of DQS to the rising edge of CK. (Identifier: MEM_LPDDR3_TWLH_PS)
tDSS	tDSS describes the time between the falling edge of DQS to the rising edge of the next CK transition . (Identifier: MEM_LPDDR3_TDSS_CYC)
tINIT	tINIT describes the time duration of the memory initialization after a device power-up . After RESET_n is de-asserted, wait for another 500us until CKE becomes active. During this time, the DRAM will start internal initialization; this will be done independently of external clocks. (Identifier: MEM_LPDDR3_TINIT_US)
tMRR	tMRR describes the minimum MODE REGISTER READ command period. (Identifier: MEM_LPDDR3_TMRR_CK_CYC)
tMRW	tMRW describes the minimum MODE REGISTER WRITE command period. (Identifier: MEM_LPDDR3_TMRW_CK_CYC)
tRAS	tRAS describes the activate to precharge duration . A row cannot be deactivated until the tRAS time has been met. Therefore tRAS determines how long the memory has to wait after a activate command before a precharge command can be issued to close the row. (Identifier: MEM_LPDDR3_TRAS_NS)
tRCD	tRCD, row command delay , describes the active to read/write time . It is the amount of delay between the activation of a row through the RAS command and the access to the data through the CAS command. (Identifier: MEM_LPDDR3_TRCD_NS)
tRPpb	Precharge command period (per bank) (Identifier: MEM_LPDDR3_TRP_NS)
tWR	tWR refers to the Write Recovery time . It specifies the amount of clock cycles needed to complete a write before a precharge command can be issued. (Identifier: MEM_LPDDR3_TWR_NS)


Table 417. Group: Mem Timing / Parameters dependent on Speed Bin, Operating Frequency, and Page Size

Display Name	Description
tRRD	tRRD refers to the Row Active to Row Active Delay . It is the minimum time interval (measured in memory clock cycles) between two activate commands to rows in different banks in the same rank (Identifier: MEM_LPDDR3_TRRD_CYC)
tFAW	tFAW refers to the four activate window time . It describes the period of time during which only four banks can be active. (Identifier: MEM_LPDDR3_TFAW_NS)
tWTR	tWTR or Write Timing Parameter describes the delay from start of internal write transaction to internal read command, for accesses to the same bank . The delay is measured from the first rising memory clock edge after the last write data is received to the rising memory clock edge when a read command is received. (Identifier: MEM_LPDDR3_TWTR_CYC)
tRTP	tRTP refers to the internal READ Command to PRECHARGE Command delay . It is the number of memory clock cycles that is needed between a read command and a precharge command to the same rank. (Identifier: MEM_LPDDR3_TRTP_CYC)

Table 418. Group: Mem Timing / Parameters dependent on Density and Temperature

Display Name	Description
tRFCab	Auto-refresh command interval (all banks) (Identifier: MEM_LPDDR3_TRFC_NS)
tREFI	tREFI refers to the average periodic refresh interval . It is the maximum amount of time the memory can tolerate in between each refresh command (Identifier: MEM_LPDDR3_TREFI_US)

10.1.6 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Board

Table 419. Group: Board / Intersymbol Interference/Crosstalk

Display Name	Description
Use default ISI/crosstalk values	You can enable this option to use default intersymbol interference and crosstalk values for your topology. Note that the default values are not optimized for your board. <i>For optimal signal integrity, it is recommended that you do not enable this parameter, but instead perform I/O simulation using IBIS models and Hyperlynx)*, and manually enter values based on your simulation results, instead of using the default values.</i> (Identifier: BOARD_LPDDR3_USE_DEFAULT_ISI_VALUES)
Address and command ISI/crosstalk	The address and command window reduction due to intersymbol interference and crosstalk effects. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_LPDDR3_USER_AC_ISI_NS)
Read DQS/DQS# ISI/crosstalk	The reduction of the read data window due to ISI and crosstalk effects on the DQS/DQS# signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_LPDDR3_USER_RCLK_ISI_NS)
continued...	



Display Name	Description
Read DQ ISI/crosstalk	The reduction of the read data window due to intersymbol interference and crosstalk effects on the DQ signal when driven by the memory device during a read. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_LPDDR3_USER_RDATA_ISI_NS)
Write DQS/DQS# ISI/crosstalk	The reduction of the write data window due to intersymbol interference and crosstalk effects on the DQS/DQS# signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_LPDDR3_USER_WCLK_ISI_NS)
Write DQ ISI/crosstalk	The reduction of the write data window due to intersymbol interference and crosstalk effects on the DQ signal when driven by the FPGA during a write. The number to be entered is the total loss of margin on the setup and hold sides (measured loss on the setup side + measured loss on the hold side) . Refer to the <i>EMIF Simulation Guidance wiki page</i> for additional information. (Identifier: BOARD_LPDDR3_USER_WDATA_ISI_NS)

Table 420. Group: Board / Board and Package Skews

Display Name	Description
Package deskewed with board layout (DQS group)	Enable this parameter if you are compensating for package skew on the DQ, DQS, and DM buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_LPDDR3_IS_SKEW_WITHIN_DQS_DESKEWED)
Maximum board skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. This value affects the read capture and write margins. (Identifier: BOARD_LPDDR3_BRD_SKEW_WITHIN_DQS_NS)
Maximum system skew within DQS group	The largest skew between all DQ and DM pins in a DQS group. Enter combined board and package skew. This value affects the read capture and write margins. (Identifier: BOARD_LPDDR3_PKG_BRD_SKEW_WITHIN_DQS_NS)
Package deskewed with board layout (address/command bus)	Enable this parameter if you are compensating for package skew on the address, command, control, and memory clock buses in the board layout. Include package skew in calculating the following board skew parameters. (Identifier: BOARD_LPDDR3_IS_SKEW_WITHIN_AC_DESKEWED)
Maximum board skew within address/command bus	The largest skew between the address and command signals. Enter the board skew only; package skew is calculated automatically, based on the memory interface configuration, and added to this value. (Identifier: BOARD_LPDDR3_BRD_SKEW_WITHIN_AC_NS)
Maximum system skew within address/command bus	Maximum system skew within address/command bus refers to the largest skew between the address and command signals. (Identifier: BOARD_LPDDR3_PKG_BRD_SKEW_WITHIN_AC_NS)
Average delay difference between DQS and CK	The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS trace delay minus the CK trace delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. (Identifier: BOARD_LPDDR3_DQS_TO_CK_SKEW_NS)
Maximum delay difference between devices	This parameter describes the largest propagation delay on the DQ signals between ranks. For example, in a two-rank configuration where devices are placed in series, there is an extra propagation delay for DQ signals going to and coming back from the furthest device compared to the nearest device. <i>This parameter is only applicable when there is more than one rank.</i>
continued...	



Display Name	Description
	(Identifier: BOARD_LPDDR3_SKEW_BETWEEN_DIMMS_NS)
Maximum skew between DQS groups	The largest skew between DQS signals. (Identifier: BOARD_LPDDR3_SKEW_BETWEEN_DQS_NS)
Average delay difference between address/command and CK	The average delay difference between the address/command signals and the CK signal, calculated by averaging the longest and smallest address/command signal trace delay minus the maximum CK trace delay. Positive values represent address and command signals that are longer than CK signals and negative values represent address and command signals that are shorter than CK signals. (Identifier: BOARD_LPDDR3_AC_TO_CK_SKEW_NS)
Maximum CK delay to device	The maximum CK delay to device refers to the delay of the longest CK trace from the FPGA to any device. (Identifier: BOARD_LPDDR3_MAX_CK_DELAY_NS)
Maximum DQS delay to device	The maximum DQS delay to device refers to the delay of the longest DQS trace from the FPGA to any device (Identifier: BOARD_LPDDR3_MAX_DQS_DELAY_NS)

10.1.7 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Controller

Table 421. Group: Controller / Avalon Interface

Display Name	Description
Avalon Interface	Selects the Avalon Interface through which the controller interacts with user logic (Identifier: CTRL_LPDDR3_AVL_PROTOCOL_ENUM)

Table 422. Group: Controller / Low Power Mode

Display Name	Description
Enable Self-Refresh Control	Select this option to enable the self-refresh control on the controller top level. The control signal allows you to place the memory device into self-refresh mode, on a per chip-select basis. (Identifier: CTRL_LPDDR3_SELF_REFRESH_EN)
Enable Auto Power-Down	Enable this parameter to have the controller automatically place the memory device into power-down mode after a specified number of idle controller clock cycles. The idle wait time is configurable. All ranks must be idle to enter auto power-down. (Identifier: CTRL_LPDDR3_AUTO_POWER_DOWN_EN)
Auto Power-Down Cycles	Specifies the number of idle controller cycles after which the memory device is placed into power-down mode. You can configure the idle waiting time. The supported range for number of cycles is from 1 to 65534. (Identifier: CTRL_LPDDR3_AUTO_POWER_DOWN_CYCS)

Table 423. Group: Controller / Efficiency

Display Name	Description
Enable User Refresh Control	When enabled, user logic has complete control and is responsible for issuing adequate refresh commands to the memory devices, via the MMR interface. This feature provides increased control over worst-case read latency and enables you to issue refresh bursts during idle periods. (Identifier: CTRL_LPDDR3_USER_REFRESH_EN)
Enable Auto-Precharge Control	Select this parameter to enable the auto-precharge control on the controller top level. If you assert the auto-precharge control signal while requesting a read or write burst, you can specify whether the controller should close

continued...



Display Name	Description
	(auto-precharge) the currently open page at the end of the read or write burst, potentially making a future access to a different page of the same bank faster. (Identifier: CTRL_LPDDR3_AUTO_PRECHARGE_EN)
Address Ordering	Controls the mapping between Avalon addresses and memory device addresses. By changing the value of this parameter, you can change the mappings between the Avalon-MM address and the DRAM address. (Identifier: CTRL_LPDDR3_ADDR_ORDER_ENUM)
Enable Reordering	Enable this parameter to allow the controller to perform command and data reordering. Reordering can improve efficiency by reducing bus turnaround time and row/bank switching time. Data reordering allows the single-port memory controller to change the order of read and write commands to achieve highest efficiency. Command reordering allows the controller to issue bank management commands early based on incoming patterns, so that the desired row in memory is already open when the command reaches the memory interface. <i>For more information, refer to the Data Reordering topic in the EMIF Handbook.</i> (Identifier: CTRL_LPDDR3_REORDER_EN)
Starvation limit for each command	Specifies the number of commands that can be served before a waiting command is served. The controller employs a counter to ensure that all requests are served after a pre-defined interval -- this ensures that low priority requests are not ignored, when doing data reordering for efficiency. The valid range for this parameter is from 1 to 63. <i>For more information, refer to the Starvation Control topic in the EMIF Handbook.</i> (Identifier: CTRL_LPDDR3_STARVE_LIMIT)
Enable Command Priority Control	Select this parameter to enable user-requested command priority control on the controller top level. This parameter instructs the controller to treat a read or write request as high-priority. The controller attempts to fill high-priority requests sooner, to reduce latency. Connect this interface to the conduit of your logic block that determines when the external memory interface IP treats the read or write request as a high-priority command. (Identifier: CTRL_LPDDR3_USER_PRIORITY_EN)

Table 424. Group: Controller / Configuration, Status and Error Handling

Display Name	Description
Enable Memory-Mapped Configuration and Status Register (MMR) Interface	Enable this parameter to change or read memory timing parameters, memory address size, mode register settings, controller status, and request sideband operations. (Identifier: CTRL_LPDDR3_MMR_EN)

Table 425. Group: Controller / Data Bus Turnaround Time

Display Name	Description
Additional read-to-write turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read to a write within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_RD_TO_WR_SAME_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (same rank)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write to a read within the same logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_WR_TO_RD_SAME_CHIP_DELTA_CYCS)
Additional read-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a read of another logical rank . This can resolve bus contention problems specific to
continued...	



Display Name	Description
	your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_RD_TO_RD_DIFF_CHIP_DELTA_CYCS)
Additional read-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a read of one logical rank to a write of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_RD_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-write turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a write of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_WR_TO_WR_DIFF_CHIP_DELTA_CYCS)
Additional write-to-read turnaround time (different ranks)	Specifies additional number of idle controller (not DRAM) cycles when switching the data bus from a write of one logical rank to a read of another logical rank . This can help resolve bus contention problems specific to your board topology. The value is added to the default which is calculated automatically. <i>Use the default setting unless you suspect a problem exists.</i> (Identifier: CTRL_LPDDR3_WR_TO_RD_DIFF_CHIP_DELTA_CYCS)


10.1.8 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Diagnostics

Table 426. Group: Diagnostics / Simulation Options

Display Name	Description
Calibration mode	Specifies whether to skip memory interface calibration during simulation, or to simulate the full calibration process. Simulating the full calibration process can take hours (or even days), depending on the width and depth of the memory interface. You can achieve much faster simulation times by skipping the calibration process, but that is only expected to work when the memory model is ideal and the interconnect delays are zero. <i>If you enable this parameter, the interface still performs some memory initialization before starting normal operations.</i> Abstract PHY is supported with skip calibration. (Identifier: DIAG_LPDDR3_SIM_CAL_MODE_ENUM)
Abstract phy for fast simulation	Specifies that the system use Abstract PHY for simulation. Abstract PHY replaces the PHY with a model for fast simulation and can reduce simulation time by 2-3 times. Abstract PHY is available for certain protocols and device families, and only when you select Skip Calibration . (Identifier: DIAG_LPDDR3_ABSTRACT_PHY)

Table 427. Group: Diagnostics / Calibration Debug Options

Display Name	Description
Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	Specifies the connectivity of an Avalon slave interface for use by the Quartus Prime EMIF Debug Toolkit or user core logic. If you set this parameter to " Disabled ", no debug features are enabled. If you set this parameter to " Export ", an Avalon slave interface named "cal_debug" is exported from the IP. To use this interface with the EMIF Debug Toolkit, you must instantiate and connect an EMIF debug interface IP
continued...	



Display Name	Description
	<p>core to it, or connect it to the <code>cal_debug_out</code> interface of another EMIF core. If you select "Add EMIF Debug Interface", an EMIF debug interface component containing a JTAG Avalon Master is connected to the debug port, allowing the core to be accessed by the EMIF Debug Toolkit.</p> <p><i>Only one EMIF debug interface should be instantiated per I/O column. You can chain additional EMIF or PHYLite cores to the first by enabling the "Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option for all cores in the chain, and selecting "Export" for the "Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port" option on all cores after the first.</i></p> <p>(Identifier: DIAG_LPDDR3_EXPORT_SEQ_AVALON_SLAVE)</p>
Use Soft NIOS Processor for On-Chip Debug	<p>Enables a soft Nios processor as a peripheral component to access the On-Chip Debug Port. <i>Only one interface in a column can activate this option.</i></p> <p>(Identifier: DIAG_SOFT_NIOS_MODE)</p>
Enable Daisy-Chaining for Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port	<p>Specifies that the IP export an Avalon-MM master interface (<code>cal_debug_out</code>) which can connect to the <code>cal_debug</code> interface of other EMIF cores residing in the same I/O column. This parameter applies only if the EMIF Debug Toolkit or On-Chip Debug Port is enabled. Refer to the <i>Debugging Multiple EMIFs</i> wiki page for more information about debugging multiple EMIFs. (Identifier: DIAG_LPDDR3_EXPORT_SEQ_AVALON_MASTER)</p>
Interface ID	<p>Identifies interfaces within the I/O column, for use by the EMIF Debug Toolkit and the On-Chip Debug Port. Interface IDs should be unique among EMIF cores within the same I/O column. If the Quartus Prime EMIF Debug Toolkit/On-Chip Debug Port parameter is set to Disabled, the interface ID is unused. (Identifier: DIAG_LPDDR3_INTERFACE_ID)</p>
Skip address/command leveling calibration	<p>Specifies to skip the address/command leveling stage during calibration. Address/command leveling attempts to center the memory clock edge against CS# by adjusting delay elements inside the PHY, and then applying the same delay offset to the rest of the address and command pins. (Identifier: DIAG_LPDDR3_SKIP_CA_LEVEL)</p>
Skip address/command deskew calibration	<p>Specifies to skip the address/command deskew calibration stage. Address/command deskew performs per-bit deskew for the address and command pins. (Identifier: DIAG_LPDDR3_SKIP_CA_DESKEW)</p>

Table 428. Group: Diagnostics / Example Design

Display Name	Description
Number of core clocks sharing slaves to instantiate in the example design	<p>Specifies the number of core clock sharing slaves to instantiate in the example design. This parameter applies only if you set the "Core clocks sharing" parameter in the "General" tab to "Master" or "Slave". (Identifier: DIAG_LPDDR3_EX_DESIGN_NUM_OF_SLAVES)</p>
Enable In-System-Sources-and-Probes	<p>Enables In-System-Sources-and-Probes in the example design for <i>common debug signals, such as calibration status or example traffic generator per-bit status</i>. This parameter must be enabled if you want to do driver margining. (Identifier: DIAG_LPDDR3_EX_DESIGN_ISSP_EN)</p>

Table 429. Group: Diagnostics / Traffic Generator

Display Name	Description
Use configurable Avalon traffic generator 2.0	<p>This option allows users to add the new configurable Avalon traffic generator to the example design. (Identifier: DIAG_LPDDR3_USE_TG_AVL_2)</p>
Export Traffic Generator 2.0 configuration interface	<p>Specifies that the IP export an Avalon-MM slave port for configuring the Traffic Generator. <i>This is required only if you are configuring the traffic generator through user logic and not through the EMIF Debug Toolkit.</i> (Identifier: DIAG_TG_AVL_2_EXPORT_CFG_INTERFACE)</p>
continued...	



Display Name	Description
Bypass the default traffic pattern	Specifies that the controller/interface bypass the traffic generator 2.0 default pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_LPDDR3_BYPASS_DEFAULT_PATTERN)
Bypass the user-configured traffic stage	Specifies that the controller/interface bypass the user-configured traffic generator's pattern after reset. If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. Configuration can be done by connecting to the traffic generator via the EMIF Debug Toolkit, or by using custom logic connected to the Avalon-MM configuration slave port on the traffic generator. Configuration can also be simulated using the example testbench provided in the altera_emif_avl_tg_2_tb.sv file. (Identifier: DIAG_LPDDR3_BYPASS_USER_STAGE)
Bypass the traffic generator repeated-writes/repeated-reads test pattern	Specifies that the controller/interface bypass the traffic generator's repeat test stage. <i>If you do not enable this parameter, every write and read is repeated several times.</i> (Identifier: DIAG_LPDDR3_BYPASS_REPEAT_STAGE)
Bypass the traffic generator stress pattern	Specifies that the controller/interface bypass the traffic generator's stress pattern stage. (Stress patterns are meant to create worst-case signal integrity patterns on the data pins.) If you do not enable this parameter, the traffic generator does not assert a pass or fail status until the generator is configured and signaled to start by its Avalon configuration interface. (Identifier: DIAG_LPDDR3_BYPASS_STRESS_STAGE)

Table 430. Group: Diagnostics / Performance

Display Name	Description
Enable Efficiency Monitor	Adds an Efficiency Monitor component to the Avalon-MM interface of the memory controller, allowing you to view efficiency statistics of the interface. <i>You can access the efficiency statistics using the EMIF Debug Toolkit.</i> (Identifier: DIAG_LPDDR3_EFFICIENCY_MONITOR)

Table 431. Group: Diagnostics / Miscellaneous

Display Name	Description
Use short Qsys interface names	Specifies the use of short interface names, for improved usability and consistency with other Qsys components. If this parameter is disabled, the names of Qsys interfaces exposed by the IP will include the type and direction of the interface. Long interface names are supported for backward-compatibility and will be removed in a future release. (Identifier: SHORT_QSYS_INTERFACE_NAMES)

10.1.9 Intel Stratix 10 EMIF IP LPDDR3 Parameters: Example Designs

Table 432. Group: Example Designs / Available Example Designs

Display Name	Description
Select design	Specifies the <i>creation of a full Quartus Prime project, instantiating an external memory interface and an example traffic generator, according to your parameterization.</i> After the design is created, you can specify the target device and pin location assignments, run a full compilation, verify timing closure, and test the interface on your board using the programming file created by the Quartus Prime assembler. The 'Generate Example Design' button lets you generate simulation or synthesis file sets. (Identifier: EX_DESIGN_GUI_LPDDR3_SEL_DESIGN)



Table 433. Group: Example Designs / Example Design Files

Display Name	Description
Simulation	Specifies that the 'Generate Example Design' button create all necessary file sets for simulation. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, simulation file sets are not created.</i> Instead, the output directory will contain the ed_sim.qsys file which holds Qsys details of the simulation example design, and a make_sim_design.tcl file with other corresponding tcl files. You can run make_sim_design.tcl from a command line to generate the simulation example design. The generated example designs for various simulators are stored in the /sim sub-directory . (Identifier: EX_DESIGN_GUI_LPDDR3_GEN_SIM)
Synthesis	Specifies that the 'Generate Example Design' button create all necessary file sets for synthesis. Expect a short additional delay as the file set is created. <i>If you do not enable this parameter, synthesis file sets are not created.</i> Instead, the output directory will contain the ed_synth.qsys file which holds Qsys details of the synthesis example design, and a make_qii_design.tcl script with other corresponding tcl files. You can run make_qii_design.tcl from a command line to generate the synthesis example design. The generated example design is stored in the /qii sub-directory . (Identifier: EX_DESIGN_GUI_LPDDR3_GEN_SYNTH)

Table 434. Group: Example Designs / Generated HDL Format

Display Name	Description
Simulation HDL format	This option lets you choose the format of HDL in which generated simulation files are created. (Identifier: EX_DESIGN_GUI_LPDDR3_HDL_FORMAT)

Table 435. Group: Example Designs / Target Development Kit

Display Name	Description
Select board	Specifies that <i>when you select a development kit with a memory module, the generated example design contains all settings and fixed pin assignments to run on the selected board. You must select a development kit preset to generate a working example design for the specified development kit.</i> Any IP settings not applied directly from a development kit preset will not have guaranteed results when testing the development kit. To exclude hardware support of the example design, select ' none ' from the ' Select board ' pull down menu. When you apply a development kit preset, all IP parameters are automatically set appropriately to match the selected preset. If you want to save your current settings, you should do so before you apply the preset. You can save your settings under a different name using File->Save as . (Identifier: EX_DESIGN_GUI_LPDDR3_TARGET_DEV_KIT)



10.2 Board Skew Equations

The following table presents the underlying equations for the board skew parameters.

10.2.1 Equations for LPDDR3 Board Skew Parameters

Table 436. Board Skew Parameter Equations

Parameter	Description/Equation
Maximum CK delay to DIMM/device	<p>The delay of the longest CK trace from the FPGA to any DIMM/device.</p> $\max_r \left[\max_n (CK_{n_r} PathDelay) \right]$ <p>Where n is the number of memory clock and r is the number rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 pairs of memory clocks in each rank DIMM, the maximum CK delay is expressed by the following equation:</p> $\max(CK_1 PathDelayrank1, CK_2 PathDelayrank1, CK_1 PathDelayrank2, CK_2 PathDelayrank2)$
Maximum DQS delay to DIMM/device	<p>The delay of the longest DQS trace from the FPGA to the DIMM/device.</p> $\max_r \left[\max_n (DQS_{n_r} PathDelay) \right]$ <p>Where n is the number of DQS and r is the number of rank of DIMM/device. For example in dual-rank DIMM implementation, if there are 2 DQS in each rank DIMM, the maximum DQS delay is expressed by the following equation:</p> $\max(DQS_1 PathDelayrank1, DQS_2 PathDelayrank1, DQS_1 PathDelayrank2, DQS_2 PathDelayrank2)$
Average delay difference between DQS and CK	<p>The average delay difference between the DQS signals and the CK signal, calculated by averaging the longest and smallest DQS delay minus the CK delay. Positive values represent DQS signals that are longer than CK signals and negative values represent DQS signals that are shorter than CK signals. The Quartus Prime software uses this skew to optimize the delay of the DQS signals for appropriate setup and hold margins.</p> $\max_r \left[\frac{\max_{n,m} \left\{ (DQS_{m_r} Delay - CK_{n_r} Delay) \right\} + \min_{n,m} \left\{ (DQS_{m_r} Delay - CK_{n_r} Delay) \right\}}{2} \right]$ <p>Where n is the number of memory clock, m is the number of DQS, and r is the number of rank of DIMM/device.</p> <p>When using discrete components, the calculation differs slightly. Find the minimum and maximum values for (DQS-CK) over all groups and then divide by 2. Calculate the (DQS-CK) for each DQS group, by using the appropriate CLK for that group.</p> <p>For example, in a configuration with 5 x16 components, with each component having two DQS groups: To find the minimum and maximum, calculate the minimum and maximum of (DQS0 – CK0, DQS1 – CK0, DQS2 – CK1, DQS3 – CK1, and so forth) and then divide the result by 2.</p>
Maximum Board skew within DQS group	<p>The largest skew between all DQ and DM pins in a DQS group. Enter your board skew only. Package skew is calculated automatically, based on the memory interface configuration, and added to this value. This value affects the read capture and write margins.</p> $\left[\max_{groups} \left[\max_{Max_g} DQ_g - \min_{Max_g} DQ_g \right] \right]$
Maximum skew between DQS groups	<p>The largest skew between DQS signals in different DQS groups.</p>

continued...



Parameter	Description/Equation
	$\left[\begin{matrix} groups \\ Max_g \end{matrix} \left[\begin{matrix} DQS_g \end{matrix} \right] \right] - \left[\begin{matrix} groups \\ Min_g \end{matrix} \left[\begin{matrix} DQS_g \end{matrix} \right] \right]$
Maximum system skew within address/command bus	<p>$(MaxAC - MinAC)$</p> <p>The largest skew between the address and command signals. Enter combined board and package skew. In the case of a component, find the maximum address/command and minimum address/command values across all component address signals.</p>
Average delay difference between address/command and CK	<p>A value equal to the average of the longest and smallest address/command signal delays, minus the delay of the CK signal. The value can be positive or negative.</p> <p>The average delay difference between the address/command and CK is expressed by the following equation:</p> $\sum_{n=1}^n \left(\frac{LongestACPathDelay + ShortestACPathDelay}{2} \right) - CK_n PathDelay$ <p style="text-align: center;">n</p> <p>where n is the number of memory clocks.</p>
Maximum delay difference between DIMMs/devices	<p>The largest propagation delay on DQ signals between ranks. For example, in a two-rank configuration where you place DIMMs in different slots there is also a propagation delay for DQ signals going to and coming back from the furthest DIMM compared to the nearest DIMM. This parameter is applicable only when there is more than one rank. $Max_r \{ \max_{n,m} [(DQn_r \text{ path delay} - DQn_{r+1} \text{ path delay}), (DQSm_r \text{ path delay} - DQSm_{r+1} \text{ path delay})] \}$</p> <p>Where n is the number of DQ, m is the number of DQS and r is number of rank of DIMM/device.</p>

10.3 Pin and Resource Planning

The following topics provide guidelines on pin placement for external memory interfaces.

Typically, all external memory interfaces require the following FPGA resources:

- Interface pins
- PLL and clock network
- Other FPGA resources—for example, core fabric logic, and on-chip termination (OCT) calibration blocks

Once all the requirements are known for your external memory interface, you can begin planning your system.

10.3.1 Interface Pins

Any I/O banks that do not support transceiver operations in Intel Stratix 10 devices support external memory interfaces. However, DQS (data strobe or data clock) and DQ (data) pins are listed in the device pin tables and are fixed at specific locations in the device. You must adhere to these pin locations to optimize routing, minimize skew, and maximize margins. Always check the pin table for the actual locations of the DQS and DQ pins.

Note: Maximum interface width varies from device to device depending on the number of I/O pins and DQS or DQ groups available. Achievable interface width also depends on the number of address and command pins that the design requires. To ensure adequate PLL, clock, and device routing resources are available, you should always test fit any IP in the Intel Quartus Prime software before PCB sign-off.



Intel devices do not limit the width of external memory interfaces beyond the following requirements:

- Maximum possible interface width in any particular device is limited by the number of DQS groups available.
- Sufficient clock networks are available to the interface PLL as required by the IP.
- Sufficient spare pins exist within the chosen bank or side of the device to include all other address and command, and clock pin placement requirements.
- The greater the number of banks, the greater the skew, hence Intel recommends that you always generate a test project of your desired configuration and confirm that it meets timing.

10.3.1.1 Estimating Pin Requirements

You should use the Intel Quartus Prime software for final pin fitting. However, you can estimate whether you have enough pins for your memory interface using the EMIF Device Selector on www.altera.com, or perform the following steps:

1. Determine how many read/write data pins are associated per data strobe or clock pair.
2. Calculate the number of other memory interface pins needed, including any other clocks (write clock or memory system clock), address, command, and RZQ. Refer to the External Memory Interface Pin Table to determine necessary Address/Command/Clock pins based on your desired configuration.
3. Calculate the total number of I/O banks required to implement the memory interface, given that an I/O bank supports up to 48 GPIO pins.

You should test the proposed pin-outs with the rest of your design in the Intel Quartus Prime software (with the correct I/O standard and OCT connections) before finalizing the pin-outs. There can be interactions between modules that are illegal in the Intel Quartus Prime software that you might not know about unless you compile the design and use the Intel Quartus Prime Pin Planner.

Related Links

[External Memory Interfaces Support Center](#)

10.3.1.2 Maximum Number of Interfaces

The maximum number of interfaces supported for a given memory protocol varies, depending on the FPGA in use.

Unless otherwise noted, the calculation for the maximum number of interfaces is based on independent interfaces where the address or command pins are not shared.

Note: You may need to share PLL clock outputs depending on your clock network usage.

For interface information for Intel Stratix 10, consult the EMIF Device Selector on www.altera.com.

Timing closure depends on device resource and routing utilization. For more information about timing closure, refer to the *Area and Timing Optimization Techniques* chapter in the *Intel Quartus Prime Handbook*.

Related Links

- [External Memory Interfaces Support Center](#)
- [Intel Stratix 10 EMIF Architecture: PLL Reference Clock Networks](#) on page 20
- [External Memory Interface Device Selector](#)
- [Intel Quartus Prime Pro Edition Handbook](#)

10.3.1.3 FPGA Resources

The Intel FPGA memory interface IP uses FPGA fabric, including registers and the Memory Block to implement the memory interface.

One OCT calibration block is used if you are using the FPGA OCT feature in the memory interface. The OCT calibration block uses a single pin (RZQ). You can select any of the available OCT calibration block as you do not need to place this block in the same bank or device side of your memory interface. The only requirement is that the I/O bank where you place the OCT calibration block uses the same VCCIO voltage as the memory interface. You can share multiple memory interfaces with the same OCT calibration block if the VCCIO voltage is the same.

10.3.1.4 OCT

If the memory interface uses any FPGA OCT calibrated series, parallel, or dynamic termination for any I/O in your design, you need a calibration block for the OCT circuitry. This calibration block is not required to be within the same bank or side of the device as the memory interface pins.

The R_{ZQ} pin in Intel Stratix 10 devices can be used as a general purpose I/O pin when it is not used to support OCT, provided the signal conforms to the bank voltage requirements.

10.3.1.5 PLL

When using PLL for external memory interfaces, you must consider the following guidelines:

- For the clock source, use the clock input pin specifically dedicated to the PLL that you want to use with your external memory interface. The input and output pins are only fully compensated when you use the dedicated PLL clock input pin. If the clock source for the PLL is not a dedicated clock input pin for the dedicated PLL, you would need an additional clock network to connect the clock source to the PLL block. Using additional clock network may increase clock jitter and degrade the timing margin.
- Pick a PLL and PLL input clock pin that are located on the same side of the device as the memory interface pins.
- Share the DLL and PLL static clocks for multiple memory interfaces provided the controllers are on the same or adjacent side of the device and run at the same memory clock frequency.
- If your design uses a dedicated PLL to only generate a DLL input reference clock, you must set the PLL mode to **No Compensation** in the Intel Quartus Prime software to minimize the jitter, or the software forces this setting automatically. The PLL does not generate other output, so it does not need to compensate for any clock path.



10.3.1.6 Pin Guidelines for Intel Stratix 10 EMIF IP

The Intel Stratix 10 device contains up to three I/O columns that can be used by external memory interfaces. The Intel Stratix 10 I/O subsystem resides in the I/O columns. Each column contains multiple I/O banks, each of which consists of four I/O lanes. An I/O lane is a group of twelve I/O ports.

The I/O column, I/O bank, I/O lane, adjacent I/O bank, and pairing pin for every physical I/O pin can be uniquely identified using the Bank Number and Index within I/O Bank values which are defined in each Intel Stratix 10 device pin-out file.

- The numeric component of the Bank Number value identifies the I/O column, while the letter represents the I/O bank.
- The Index within I/O Bank value falls within one of the following ranges: 0 to 11, 12 to 23, 24 to 35, or 36 to 47, and represents I/O lanes 1, 2, 3, and 4, respectively.
- The adjacent I/O bank is defined as the I/O bank with same column number but the letter is either before or after the respective I/O bank letter in the A-Z system.
- The pairing pin for an I/O pin is located in the same I/O bank. You can identify the pairing pin by adding one to its Index within I/O Bank number (if it is an even number), or by subtracting one from its Index within I/O Bank number (if it is an odd number).

For example, a physical pin with a Bank Number of 2M and Index within I/O Bank of 22, indicates that the pin resides in I/O lane 2, in I/O bank 2M, in column 2. The adjacent I/O banks are 2L and 2N. The pairing pin for this physical pin is the pin with an Index within I/O Bank of 23 and Bank Number of 2M.

10.3.1.6.1 General Guidelines

You should follow the recommended guidelines when performing pin placement for all external memory interface pins targeting Intel Stratix 10 devices, whether you are using the hard memory controller or your own solution.

If you are using the hard memory controller, you should employ the relative pin locations defined in the `<variation_name>/altera_emif_arch_nd_version number/<synth/sim>/<variation_name>_altera_emif_arch_nd_version number_<unique ID>_readme.txt` file, w+ted with your IP.

Note:

1. EMIF IP pin-out requirements for the Intel Stratix 10 Hard Processor Subsystem (HPS) are more restrictive than for a non-HPS memory interface. The HPS EMIF IP defines a fixed pin-out in the Intel Quartus Prime IP file (.qip), based on the IP configuration. When targeting Intel Stratix 10 HPS, you do not need to make location assignments for external memory interface pins. To obtain the HPS-specific external memory interface pin-out, compile the interface in the Intel Quartus Prime software. Alternatively, consult the device handbook or the device pin-out files. For information on how you can customize the HPS EMIF pin-out, refer to *Restrictions on I/O Bank Usage for Intel Stratix 10 EMIF IP with HPS*.
2. Ping Pong PHY, PHY only, RLDRAMx, QDRx and LPDDR3 are not supported with HPS.



Observe the following general guidelines when placing pins for your Intel Stratix 10 external memory interface:

1. Ensure that the pins of a single external memory interface reside within a single I/O column.
2. An external memory interface can occupy one or more banks in the same I/O column. When an interface must occupy multiple banks, ensure that those banks are adjacent to one another.
3. Any pin in the same bank that is not used by an external memory interface is available for use as a general purpose I/O of compatible voltage and termination settings.
4. All address and command pins and their associated clock pins (CK and CK#) must reside within a single bank. The bank containing the address and command pins is identified as the address and command bank.
5. To minimize latency, when the interface uses more than two banks, you must select the center bank of the interface as the address and command bank.
6. The address and command pins and their associated clock pins in the address and command bank must follow a fixed pin-out scheme, as defined in the *Intel Stratix 10 External Memory Interface Pin Information File*, which is available on www.altera.com.

You do not have to place every address and command pin manually. If you assign the location for one address and command pin, the Fitter automatically places the remaining address and command pins.

Note: The pin-out scheme is a hardware requirement that you must follow, and can vary according to the topology of the memory device. Some schemes require three lanes to implement address and command pins, while others require four lanes. To determine which scheme to follow, refer to the messages window during parameterization of your IP, or to the

```
<variation_name>/altera_emif_arch_nd_<version>/<synth/sim>/
```

```
<variation_name>_altera_emif_arch_nd_<version>_<unique ID>_readme.txt file after you have generated your IP.
```

7. An unused I/O lane in the address and command bank can serve to implement a data group, such as a x8 DQS group. The data group must be from the same controller as the address and command signals.
8. An I/O lane must not be used by both address and command pins and data pins.
9. Place read data groups according to the DQS grouping in the pin table and Pin Planner. Read data strobes (such as DQS and DQS#) or read clocks (such as CQ and CQ# / QK and QK#) must reside at physical pins capable of functioning as DQS/CQ and DQSn/CQn for a specific read data group size. You must place the associated read data pins (such as DQ and Q), within the same group.



- Note:* a. Unlike other device families, there is no need to swap CQ/CQ# pins in certain QDR II and QDR II+ latency configurations.
- b. QDR-IV requires that the polarity of all QKB/QKB# pins be swapped with respect to the polarity of the differential buffer inputs on the FPGA to ensure correct data capture on port B. All QKB pins on the memory device must be connected to the negative pins of the input buffers on the FPGA side, and all QKB# pins on the memory device must be connected to the positive pins of the input buffers on the FPGA side. Notice that the port names at the top-level of the IP already reflect this swap (that is, mem_qkb is assigned to the negative buffer leg, and mem_qkb_n is assigned to the positive buffer leg).

10. You can implement two x4 DQS groups with a single I/O lane. The pin table specifies which pins within an I/O lane can be used for the two pairs of DQS and DQS# signals. In addition, for x4 DQS groups you must observe the following rules:
 - There must be an even number of x4 groups in an external memory interface.
 - DQS group 0 and DQS group 1 must be placed in the same I/O lane. Similarly, DQS group 2 and group 3 must be in the same I/O lane. Generally, DQS group X and DQS group X+1 must be in the same I/O lane, where X is an even number.
11. You should place the write data groups according to the DQS grouping in the pin table and Pin Planner. Output-only data clocks for QDR II, QDR II+, and QDR II+ Extreme, and RLDRAM 3 protocols need not be placed on DQS/DQSn pins, but must be placed on a differential pin pair. They must be placed in the same I/O bank as the corresponding DQS group.

Note: For RLDRAM 3, x36 device, DQ[8:0] and DQ[26:18] are referenced to DK0/DK0#, and DQ[17:9] and DQ[35:27] are referenced to DK1/DK1#.
12. For protocols and topologies with bidirectional data pins where a write data group consists of multiple read data groups, you should place the data groups and their respective write and read clock in the same bank to improve I/O timing.

You do not need to specify the location of every data pin manually. If you assign the location for the read capture strobe/clock pin pairs, the Fitter will automatically place the remaining data pins.
13. Ensure that DM/BWS pins are paired with a write data pin by placing one in an I/O pin and another in the pairing pin for that I/O pin. It is recommended—though not required—that you follow the same rule for DBI pins, so that at a later date you have the freedom to repurpose the pin as DM.

Note:

1. x4 mode does not support DM/DBI, or Intel Stratix 10 EMIF IP for HPS.
2. If you are using an Intel Stratix 10 EMIF IP-based RLDRAM 3 external memory interface, you should ensure that all the pins in a DQS group (that is, DQ, DM, DK, and QK) are placed in the same I/O bank. This requirement facilitates timing closure and is necessary for successful compilation of your design.

Multiple Interfaces in the Same I/O Column

To place multiple interfaces in the same I/O column, you must ensure that the global reset signals (global_reset_n) for each individual interface all come from the same input pin or signal.

I/O Banks Selection

- For each memory interface, select consecutive I/O banks. (That is, select banks that contain the same column number and letter before or after the respective I/O bank letter.)
- A memory interface can only span across I/O banks in the same I/O column.
- The number of I/O banks that you require depends on the memory interface width.
- In some device packages, the number of I/O pins in some LVDS I/O banks is less than 48 pins.

Address/Command Pins Location

- All address/command pins for a controller must be in a single I/O bank.
- If your interface uses multiple I/O banks, the address/command pins must use the middle bank. If the number of banks used by the interface is even, any of the two middle I/O banks can be used for address/command pins.
- Address/command pins and data pins cannot share an I/O lane but can share an I/O bank.
- The address/command pin locations for the soft and hard memory controllers are predefined. In the *External Memory Interface Pin Information for Devices* spreadsheet, each index in the "Index within I/O bank" column denotes a dedicated address/command pin function for a given protocol. The index number of the pin specifies to which I/O lane the pin belongs:
 - I/O lane 0—Pins with index 0 to 11
 - I/O lane 1—Pins with index 12 to 23
 - I/O lane 2—Pins with index 24 to 35
 - I/O lane 3—Pins with index 36 to 47
- For memory topologies and protocols that require only three I/O lanes for the address/command pins, use I/O lanes 0, 1, and 2.
- Unused address/command pins in an I/O lane can be used as general-purpose I/O pins.

CK Pins Assignment

Assign the clock pin (CK pin) according to the number of I/O banks in an interface:

- If the number of I/O banks is odd, assign one CK pin to the middle I/O bank.
- If the number of I/O banks is even, assign the CK pin to either of the middle two I/O banks.

Although the Fitter can automatically select the required I/O banks, Intel recommends that you make the selection manually to reduce the pre-fit run time.

PLL Reference Clock Pin Placement

Place the PLL reference clock pin in the address/command bank. Other I/O banks may not have free pins that you can use as the PLL reference clock pin:

- If you are sharing the PLL reference clock pin between several interfaces, the I/O banks must be adjacent. (That is, the banks must contain the same column number and letter before or after the respective I/O bank letter.)



The Intel Stratix 10 external memory interface IP does not support PLL cascading.

RZQ Pin Placement

You may place the R_{ZQ} pin in any I/O bank in an I/O column with the correct V_{CCIO} and V_{CCPT} for the memory interface I/O standard in use. However, the recommended location is in the address/command I/O bank, for greater flexibility during debug if a narrower interface project is required for testing.

DQ and DQS Pins Assignment

Intel recommends that you assign the DQS pins to the remaining I/O lanes in the I/O banks as required:

- Constrain the DQ and DQS signals of the same DQS group to the same I/O lane.
- You cannot constrain DQ signals from two different DQS groups to the same I/O lane.

If you do not specify the DQS pins assignment, the Fitter selects the DQS pins automatically.

Sharing an I/O Bank Across Multiple Interfaces

If you are sharing an I/O bank across multiple external memory interfaces, follow these guidelines:

- The interfaces must use the same protocol, voltage, data rate, frequency, and PLL reference clock.
- You cannot use an I/O bank as the address/command bank for more than one interface. The memory controller and sequencer cannot be shared.
- You cannot share an I/O lane. There is only one DQS input per I/O lane, and an I/O lane can connect to only one memory controller.

10.3.1.6.2 LPDDR3 Clock Signal

CK and CK_n are differential clock inputs to the LPDDR3 interface. All the double data rate (DDR) inputs are sampled on both the positive and negative edges of the clock. Single data rate (SDR) inputs, CS_n and CKE , are sampled at the positive clock edge.

The clock is defined as the differential pair which consists of CK and CK_n . The positive clock edge is defined by the cross point of a rising CK and a falling CK_n . The negative clock edge is defined by the cross point of a falling CK and a rising CK_n .

The SDRAM data sheet specifies timing data for the following:

- t_{DSH} is the DQS falling edge hold time from CK .
- t_{DSS} is the DQS falling edge to the CK setup time.
- t_{DQSS} is the Write command to the first DQS latching transition.
- t_{DQSCK} is the DQS output access time from CK/CK_n .

10.3.1.6.3 LPDDR3 Command and Address Signal

All LPDDR3 devices use double data rate architecture on the command/address bus to reduce the number of input pins in the system. The 10-bit command/address bus contains command, address, and bank/row buffer information. Each command uses one clock cycle, during which command information is transferred on both the positive and negative edges of the clock.

10.3.1.6.4 LPDDR3 Data, Data Strobe, and DM Signals

LPDDR3 devices use bidirectional and differential data strobes.

Differential DQS operation enables improved system timing due to reduced crosstalk and less simultaneous switching noise on the strobe output drivers. The DQ pins are also bidirectional. DQS is edge-aligned with the read data and centered with the write data.

DM is the input mask for the write data signal. Input data is masked when DM is sampled high coincident with that input data during a write access.

10.3.1.6.5 Resource Sharing Guidelines (Multiple Interfaces)

In Intel Cyclone 10Intel Stratix 10 external memory interface IP, different external memory interfaces can share PLL reference clock pins, core clock networks, I/O banks, and hard Nios processors. Each I/O bank has DLL and PLL resources, therefore these do not need to be shared. The Intel Quartus Prime Fitter automatically merges DLL and PLL resources when a bank is shared by different external memory interfaces, and duplicates them for a multi-I/O-bank external memory interface.

PLL Reference Clock Pin

To conserve pin usage and enable core clock network and I/O bank sharing, you can share a PLL reference clock pin between multiple external memory interfaces; the interfaces must be of the same protocol, rate, and frequency. Sharing of a PLL reference clock pin also implies sharing of the reference clock network.

Observe the following guidelines for sharing the PLL reference clock pin:

1. To share a PLL reference clock pin, connect the same signal to the `pll_ref_clk` port of multiple external memory interfaces in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Place related external memory interfaces in adjacent I/O banks. If you leave an unused I/O bank between the I/O banks used by the external memory interfaces, that I/O bank cannot be used by any other external memory interface with a different PLL reference clock signal.

Note:

You can place the `pll_ref_clk` pin in the address and command I/O bank or in a data I/O bank, there is no impact on timing. However, for greatest flexibility during debug (such as when creating designs with narrower interfaces), the recommended placement is in the address and command I/O bank.

Core Clock Network

To access all external memory interfaces synchronously and to reduce global clock network usage, you may share the same core clock network with other external memory interfaces.



Observe the following guidelines for sharing the core clock network:

1. To share a core clock network, connect the `clks_sharing_master_out` of the master to the `clks_sharing_slave_in` of all slaves in the RTL code.
2. Place related external memory interfaces in the same I/O column.
3. Related external memory interface must have the same rate, memory clock frequency, and PLL reference clock.

I/O Bank

To reduce I/O bank utilization, you may share an I/O Bank with other external memory interfaces.

Observe the following guidelines for sharing an I/O Bank:

1. Related external memory interfaces must have the same protocol, rate, memory clock frequency, and PLL reference clock.
2. You cannot use a given I/O bank as the address and command bank for more than one external memory interface.
3. You cannot share an I/O lane between external memory interfaces, but an unused pin can serve as a general purpose I/O pin, of compatible voltage and termination standards.

Hard Nios Processor

All external memory interfaces residing in the same I/O column will share the same hard Nios processor. The shared hard Nios processor calibrates the external memory interfaces serially.

10.4 LPDDR3 Board Design Guidelines

The following topics provide guidelines for improving the signal integrity of your system and for successfully implementing an LPDDR3 interface on your system.

The following areas are discussed:

- comparison of various types of termination schemes, and their effects on the signal quality on the receiver
- proper drive strength setting on the FPGA to optimize the signal integrity at the receiver
- effects of different loading types, such as components versus DIMM configuration, on signal quality


It is important to understand the trade-offs between different types of termination schemes, the effects of output drive strengths, and different loading types, so that you can swiftly navigate through the multiple combinations and choose the best possible settings for your designs.

The following key factors affect signal quality at the receiver:

- Leveling and dynamic ODT
- Proper use of termination
- Layout guidelines

As memory interface performance increases, board designers must pay closer attention to the quality of the signal seen at the receiver because poorly transmitted signals can dramatically reduce the overall data-valid margin at the receiver. The following figure shows the differences between an ideal and real signal seen by the receiver.

Figure 82. Ideal and Real Signal at the Receiver



Related Links

JEDEC.org

10.4.1 Terminations for DDR3 and DDR4 with Intel Stratix 10 Devices

The following topics describe considerations specific to DDR3 and DDR4 external memory interface protocols on Intel Stratix 10 devices.

10.4.1.1 Dynamic On-Chip Termination (OCT) in Intel Stratix 10 Devices

Depending upon the R_s (series) and R_t (parallel) OCT values that you want, you should choose appropriate values for the RZQ resistor and connect this resistor to the RZQ pin of the FPGA.

- Select a 240-ohm reference resistor to ground to implement R_s OCT values of 34-ohm, 40-ohm, 48-ohm, 60-ohm, and 80-ohm, and R_t OCT resistance values of 20-ohm, 30-ohm, 34-ohm, 40-ohm, 60-ohm, 80-ohm, 120-ohm and 240 ohm.
- Select a 100-ohm reference resistor to ground to implement R_s OCT values of 25-ohm and 50-ohm, and an R_t OCT resistance of 50-ohm.

Check the FPGA I/O tab of the parameter editor to determine the I/O standards and termination values supported for data, address and command, and memory clock signals.

Related Links

[Choosing Terminations on Intel Stratix 10 Devices](#) on page 182

10.4.1.2 Choosing Terminations on Intel Stratix 10 Devices

To determine optimal on-chip termination (OCT) and on-die termination (ODT) values for best signal integrity, you should simulate your memory interface in HyperLynx or a similar tool.



If the optimal OCT and ODT termination values as determined by simulation are not available in the list of available values in the parameter editor, select the closest available termination values for OCT and ODT.

For information about available ODT choices, refer to your memory vendor data sheet.

Related Links

[Dynamic On-Chip Termination \(OCT\) in Intel Stratix 10 Devices](#) on page 182

10.4.1.3 On-Chip Termination Recommendations for Intel Stratix 10 Devices

- Output mode (drive strength) for Address/Command/Clock and Data Signals: Depending upon the I/O standard that you have selected, you would have a range of selections expressed in terms of ohms or miliamps. A value of 34 to 40 ohms or 12 mA is a good starting point for output mode drive strength.
- Input mode (parallel termination) for Data and Data Strobe signals: A value of 40 or 60 ohms is a good starting point for FPGA side input termination.

10.4.2 Channel Signal Integrity Measurement

As external memory interface data rates increase, so does the importance of proper channel signal integrity measurement. By measuring the actual channel loss during the layout process and including that data in your parameterization, a realistic assessment of margins is achieved.

10.4.2.1 Importance of Accurate Channel Signal Integrity Information

Default values for channel loss (or eye reduction) can be used when calculating timing margins, however those default values may not accurately reflect the channel loss in your system. If the channel loss in your system is different than the default values, the calculated timing margins will vary accordingly.

If your actual channel loss is greater than the default channel loss, and if you rely on default values, the available timing margins for the entire system will be lower than the values calculated during compilation. By relying on default values that do not accurately reflect your system, you may be lead to believe that you have good timing margin, while in reality, your design may require changes to achieve good channel signal integrity.


10.4.2.2 Understanding Channel Signal Integrity Measurement

To measure channel signal integrity you need to measure the channel loss for various signals. For a particular signal or signal trace, channel loss is defined as loss of the eye width at $\pm V_{IH}(\text{ac and dc}) \pm V_{IL}(\text{ac and dc})$. V_{IH}/V_{IL} above or below V_{REF} is used to align with various requirements of the timing model for memory interfaces.

The example below shows a reference eye diagram where the channel loss on the setup- or leading-side of the eye is equal to the channel loss on the hold- or lagging-side of the eye; however, it does not necessarily have to be that way. Because Intel's calibrating PHY will calibrate to the center of the read and write eye, the Board Settings tab has parameters for the total extra channel loss for Write DQ and Read DQ. For address and command signals which are not-calibrated, the Board Settings

tab allows you to enter setup- and hold-side channel losses that are not equal, allowing the Intel Quartus Prime software to place the clock statically within the center of the address and command eye.

Figure 83. Equal Setup and Hold-side Losses



10.4.2.3 How to Enter Calculated Channel Signal Integrity Values

You should enter calculated channel loss values in the **Channel Signal Integrity** section of the **Board** (or **Board Timing**) tab of the parameter editor.

For Intel Stratix 10 external memory interfaces, the default channel loss displayed in the parameter editor is based on the selected configuration (different values for single rank versus dual rank), and on internal Intel reference boards. You should replace the default value with the value that you calculate.

10.4.2.4 Guidelines for Calculating DDR3 Channel Signal Integrity

Address and Command ISI and Crosstalk

Simulate the address/command and control signals and capture eye at the DRAM pins, using the memory clock as the trigger for the memory interface's address/command and control signals. Measure the setup and hold channel losses at the voltage thresholds mentioned in the memory vendor's data sheet.


Address and command channel loss = Measured loss on the setup side + measured loss on the hold side.

$$V_{REF} = V_{DD}/2 = 0.75 \text{ mV for DDR3}$$

You should select the V_{IH} and V_{IL} voltage levels appropriately for the DDR3L memory device that you are using. Check with your memory vendor for the correct voltage levels, as the levels may vary for different speed grades of device.

The following figure illustrates a DDR3 example where $V_{IH(AC)}/V_{IL(AC)}$ is ± 150 mV and $V_{IH(DC)}/V_{IL(DC)}$ is ± 100 mV.

Figure 84.




Write DQ ISI and Crosstalk

Simulate the write DQ signals and capture eye at the DRAM pins, using DQ Strobe (DQS) as a trigger for the DQ signals of the memory interface simulation. Measure the setup and hold channel losses at the V_{IH} and V_{IL} mentioned in the memory vendor's data sheet. The following figure illustrates a DDR3 example where $V_{IH(AC)}/V_{IL(AC)}$ is ± 150 mV and $V_{IH(DC)}/V_{IL(DC)}$ is ± 100 mV.

Write Channel Loss = Measured Loss on the Setup side + Measured Loss on the Hold side

$$V_{REF} = V_{DD}/2 = 0.75 \text{ mV for DDR3}$$

Figure 85.



Read DQ ISI and Crosstalk


Simulate read DQ signals and capture eye at the FPGA die. Do not measure at the pin, because you might see unwanted reflections that could create a false representation of the eye opening at the input buffer of the FPGA. Use DQ Strobe (DQS) as a trigger for the DQ signals of your memory interface simulation. Measure the eye opening at ± 70 mV (V_{IH}/V_{IL}) with respect to V_{REF} .

Read Channel Loss = (UI) - (Eye opening at ± 70 mV with respect to V_{REF})

UI = Unit interval. For example, if you are running your interface at 800 Mhz, the effective data is 1600 Mbps, giving a unit interval of $1/1600 = 625$ ps

$V_{REF} = VDD/2 = 0.75$ mV for DDR3

Figure 86.




Write/Read DQS ISI and Crosstalk

Simulate the Write/Read DQS and capture eye, and measure the uncertainty at V_{REF} .

$V_{REF} = VDD/2 = 0.75$ mV for DDR3

Figure 87.

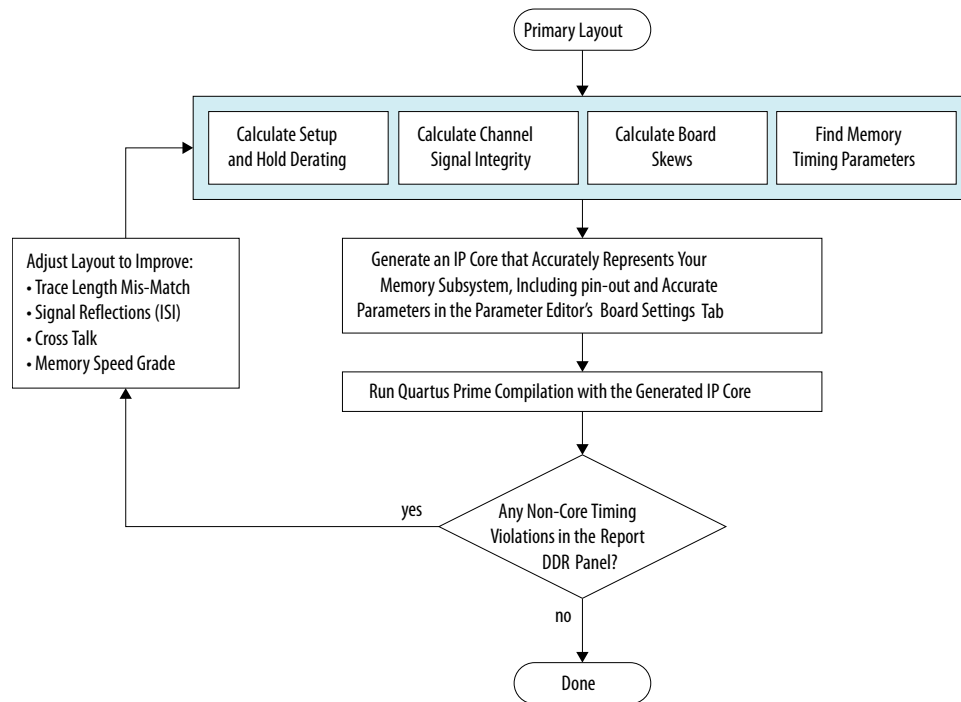


10.4.3 Layout Approach

For all practical purposes, you can regard the Timing Analyzer report on your memory interface as definitive for a given set of memory and board timing parameters.

You will find timing under **Report DDR** in the Timing Analyzer and on the **Timing Analysis** tab in the parameter editor.

The following flowchart illustrates the recommended process to follow during the board design phase, to determine timing margin and make iterative improvements to your design.



Board Skew

For information on calculating board skew parameters, refer to *Board Skew Equations*, in this chapter.

The Board Skew Parameter Tool is an interactive tool that can help you calculate board skew parameters if you know the absolute delay values for all the memory related traces.

Memory Timing Parameters

For information on the memory timing parameters to be entered into the parameter editor, refer to the datasheet for your external memory device.

Related Links

[Board Skew Parameter Tool](#)

10.4.4 Design Layout Guidelines

The general layout guidelines in the following topic apply to DDR3 and DDR4 SDRAM interfaces.

These guidelines will help you plan your board layout, but are not meant as strict rules that must be adhered to. Intel recommends that you perform your own board-level simulations to ensure that the layout you choose for your board allows you to achieve your desired performance.



For more information about how the memory manufacturers route these address and control signals on their DIMMs, refer to the Cadence PCB browser from the Cadence website, at www.cadence.com. The various JEDEC example DIMM layouts are available from the JEDEC website, at www.jedec.org.

For assistance in calculating board skew parameters, refer to the board skew calculator tool, which is available at the Intel website.

Note:

1. The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristic of the interface. They do not include any margin for crosstalk.
2. To ensure reliable timing closure to and from the periphery of the device, signals to and from the periphery should be registered before any further logic is connected.

Intel recommends that you get accurate time base skew numbers for your design when you simulate the specific implementation.

Related Links

- [Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface \(>72 bits\)](#) on page 195
- JEDEC.org
- <https://www.cadence.com/>
- [Board Skew Parameter Tool](#)
- <https://www.mentor.com/>

10.4.4.1 General Layout Guidelines

The following table lists general board design layout guidelines. These guidelines are Intel recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.

Table 437. General Layout Guidelines

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> All unused via pads must be removed, because they cause unwanted capacitance. Trace impedance plays an important role in the signal integrity. You must perform board level simulation to determine the best characteristic impedance for your PCB. For example, it is possible that for multi rank systems 40 ohms could yield better results than a traditional 50 ohm characteristic impedance.
Decoupling Parameter	<ul style="list-style-type: none"> Use 0.1 uF in 0402 size to minimize inductance Make VTT voltage decoupling close to termination resistors Connect decoupling caps between VTT and ground Use a 0.1 uF cap for every other VTT pin and 0.01 uF cap for every VDD and VDDQ pin Verify the capacitive decoupling using the Intel Power Distribution Network Design Tool
Power	<ul style="list-style-type: none"> Route GND and V_{CC} as planes Route VCCIO for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation Route VTT as islands or 250-mil (6.35-mm) power traces Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Intel recommends that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> Use 45° angles (<i>not</i> 90° corners) Avoid T-Junctions for critical nets or clocks Avoid T-junctions greater than 250 mils (6.35 mm) Disallow signals across split planes Restrict routing other signals close to system reset signals Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks

Related Links

[Power Distribution Network](#)

10.4.4.2 Layout Guidelines

The following table lists layout guidelines.

Unless otherwise specified, the guidelines in the following table apply to the following topologies:

- DIMM—UDIMM topology
- DIMM—RDIMM topology
- DIMM—LRDIMM topology
- Not all versions of the Intel Quartus Prime software support LRDIMM.
- Discrete components laid out in UDIMM topology
- Discrete components laid out in RDIMM topology

These guidelines are recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface.



For supported frequencies and topologies, refer to the *External Memory Interface Spec Estimator* <http://www.altera.com/technology/memory/estimator/mem-emif-index.html>.

For frequencies greater than 800 MHz, when you are calculating the delay associated with a trace, you must take the FPGA package delays into consideration.

Table 438. Layout Guidelines (1)

Parameter	Guidelines
Decoupling Parameter	<ul style="list-style-type: none"> Make VTT voltage decoupling close to the components and pull-up resistors. Connect decoupling caps between VTT and VDD using a 0.1 μF cap for every other VTT pin. Use a 0.1 μF cap and 0.01 μF cap for every VDDQ pin.
Maximum Trace Length	<ul style="list-style-type: none"> Even though there are no hard requirements for minimum trace length, you need to simulate the trace to ensure the signal integrity. Shorter routes result in better timing. For DIMM topology only: <ul style="list-style-type: none"> Maximum trace length for all signals from FPGA to the first DIMM slot is 4.5 inches. Maximum trace length for all signals from DIMM slot to DIMM slot is 0.425 inches. For discrete components only: <ul style="list-style-type: none"> Maximum trace length for address, command, control, and clock from FPGA to the first component must not be more than 7 inches. Maximum trace length for DQ, DQS, DQS#, and DM from FPGA to the first component is 5 inches.
General Routing	<ul style="list-style-type: none"> Route over appropriate VCC and GND planes. Keep signal routing layers close to GND and power planes.
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always make sure that the signals related to memory interface are routed between appropriate GND or power layers. For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) for these traces. (Where H is the vertical distance to the closest return path for that particular trace.) For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) these traces. (Where H is the vertical distance to the closest return path for that particular trace.) For Clock traces: Maintain at least 5H spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm). Route clock signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not exceed 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specification. These signals should maintain the following spacings: <ul style="list-style-type: none"> Clocks should maintain a length-matching between clock pairs of ± 5 ps. Clocks should maintain a length-matching between positive (\bar{p}) and negative (\bar{n}) signals of ± 2 ps, routed in parallel. Space between different pairs should be at least two times the trace width of the differential pair to minimize loss and maximize interconnect density. To avoid mismatched transmission line to via, Intel recommends that you use Ground Signal Signal Ground (GSSG) topology for your clock pattern—GND CLKP CKLN GND. Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
continued...	



Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> Route address and command signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not be more than 0.69 tCK for DDR3 and 1.5 tCK for DDR4. For different DIMM configurations, check the appropriate JEDEC specifications. UDIMMs are more susceptible to cross-talk and are generally noisier than buffered DIMMs. Therefore, route address and command signals of UDIMMs on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. Do not route differential clock (CK) and clock enable (CKE) signals close to address signals. Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. Refer to the following figure.
DQ, DM, and DQS Routing Rules	<ul style="list-style-type: none"> All the trace length matching requirements are from the FPGA package ball to the SDRAM package ball, which means you must consider trace mismatching on different DIMM raw cards. Match in length all DQ, DQS, and DM signals within a given byte-lane group with a maximum deviation of ± 10 ps. Ensure to route all DQ, DQS, and DM signals within a given byte-lane group on the same layer to avoid layer to layer transmission velocity differences, which otherwise increase the skew within the group. Do not count on FPGAs to deskew for more than 20 ps of DQ group skew. The skew algorithm only removes the following possible uncertainties: <ul style="list-style-type: none"> Minimum and maximum die IOE skew or delay mismatch Minimum and maximum device package skew or mismatch Board delay mismatch of 20 ps Memory component DQ skew mismatch Increasing any of these four parameters runs the risk of the deskew algorithm limiting, failing to correct for the total observed system skew. If the algorithm cannot compensate without limiting the correction, timing analysis shows reduced margins. For memory interfaces with leveling, the timing between the DQS and clock signals on each device calibrates dynamically to meet tDQSS. To make sure the skew is not too large for the leveling circuit's capability, follow these rules: <ul style="list-style-type: none"> Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device: $(CK_i) - DQSi > 0$; $0 < i < \text{number of components} - 1$. For DIMMs, ensure that the CK trace is longer than the longest DQS trace at the DIMM connector. Total skew of CLK and DQS signal between groups is less than one clock cycle: $(CK_i + DQSi)_{\text{max}} - (CK_i + DQSi)_{\text{min}} < 1 \times tCK$ (If you are using a DIMM topology, your delay and skew must take into consideration values for the actual DIMM.)
continued...	



Parameter	Guidelines
Spacing Guidelines	<ul style="list-style-type: none"> Avoid routing two signal layers next to each other. Always ensure that the signals related to the memory interface are routed between appropriate GND or power layers. For DQ/DQS/DM traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Address/Command/Control traces: Maintain at least 3H spacing between the edges (air-gap) of these traces, where H is the vertical distance to the closest return path for that particular trace. For Clock traces: Maintain at least 5H spacing between two clock pairs or a clock pair and any other memory interface trace, where H is the vertical distance to the closest return path for that particular trace.
Intel Quartus Prime Software Settings for Board Layout	<ul style="list-style-type: none"> To perform timing analyses on board and I/O buffers, use a third-party simulation tool to simulate all timing information such as skew, ISI, crosstalk, and type the simulation result into the Board Settings tab in the parameter editor. Do not use advanced I/O timing model (AIOT) or board trace model unless you do not have access to any third party tool. AIOT provides reasonable accuracy but tools like HyperLynx provide better results.
Notes to Table: 1. For point-to-point and DIMM interface designs, refer to the Micron website, www.micron.com .	

Related Links


- [Package Deskew](#) on page 198
- [External Memory Interface Spec Estimator](#)
- www.micron.com

10.4.4.3 Length Matching Rules

The following topics provide guidance on length matching for different types of SDRAM signals.

Route all addresses and commands to match the clock signals to within ± 20 ps to each discrete memory component. The following figure shows the component routing guidelines for address and command signals.

Figure 88. SDRAM Component Address and Command Routing Guidelines



If using discrete components:

$$\begin{aligned}
 x &= y \pm 20 \text{ ps} \\
 x + x1 &= y + y1 \pm 20 \text{ ps} \\
 x + x1 + x2 &= y + y1 + y2 \pm 20 \text{ ps} \\
 x + x1 + x2 + x3 &= y + y1 + y2 + y3 \pm 20 \text{ ps}
 \end{aligned}$$

If using a DIMM topology:

$$x = y \pm 20 \text{ ps}$$

The timing between the DQS and clock signals on each device calibrates dynamically to meet t_{DQSS} . The following figure shows the delay requirements to align DQS and clock signals. To ensure that the skew is not too large for the leveling circuit's capability, follow these rules:

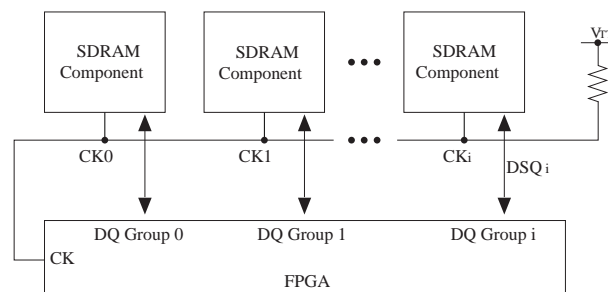
- Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device:

$$CK_i - DQSi > 0; 0 < i < \text{number of components} - 1$$

- Total skew of CLK and DQS signal between groups is less than one clock cycle:

$$(CK_i + DQSi)_{\max} - (CK_i + DQSi)_{\min} < 1 \times t_{CK}$$

Figure 89. Delaying DQS Signal to Align DQS and Clock



CK_i = Clock signal propagation delay to device i

$DQSi$ = DQ/DQS signals propagation delay to group i

Clk pair matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, match the trace lengths up to the DIMM connector. If you are using discrete components, match the lengths for all the memory components connected in the fly-by chain.

DQ group length matching—If you are using a DIMM (UDIMM, RDIMM, or LRDIMM) topology, apply the DQ group trace matching rules described in the guideline table earlier up to the DIMM connector. If you are using discrete components, match the lengths up to the respective memory components.


When you are using DIMMs, it is assumed that lengths are tightly matched within the DIMM itself. You should check that appropriate traces are length-matched within the DIMM.

10.4.4.4 Spacing Guidelines

This topic provides recommendations for minimum spacing between board traces for various signal traces.

Spacing Guidelines for DQ, DQS, and DM Traces


Maintain a minimum of $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)






Spacing Guidelines for Address and Command and Control Traces

Maintain at least $3H$ spacing between the edges (air-gap) of these traces. (Where H is the vertical distance to the closest return path for that particular trace.)



Spacing Guidelines for Clock Traces

Maintain at least $5H$ spacing between two clock pair or a clock pair and any other memory interface trace. (Where H is the vertical distance to the closest return path for that particular trace.)



10.4.4.5 Layout Guidelines for DDR3 and DDR4 SDRAM Wide Interface (>72 bits)

The following topics discuss different ways to lay out a wider DDR3 or DDR4 SDRAM interface to the FPGA. Choose the topology based on board trace simulation and the timing budget of your system.

The EMIF IP supports up to a 144-bit wide DDR3 interface. You can use discrete components or DIMMs to implement a wide interface (any interface wider than 72 bits). Intel recommends using leveling when you implement a wide interface with DDR3 components.

When you lay out for a wider interface, all rules and constraints discussed in the previous sections still apply. The DQS, DQ, and DM signals are point-to-point, and all the same rules discussed in *Design Layout Guidelines* apply.

The main challenge for the design of the fly-by network topology for the clock, command, and address signals is to avoid signal integrity issues, and to make sure you route the DQS, DQ, and DM signals with the chosen topology.

Related Links

[Design Layout Guidelines](#) on page 188

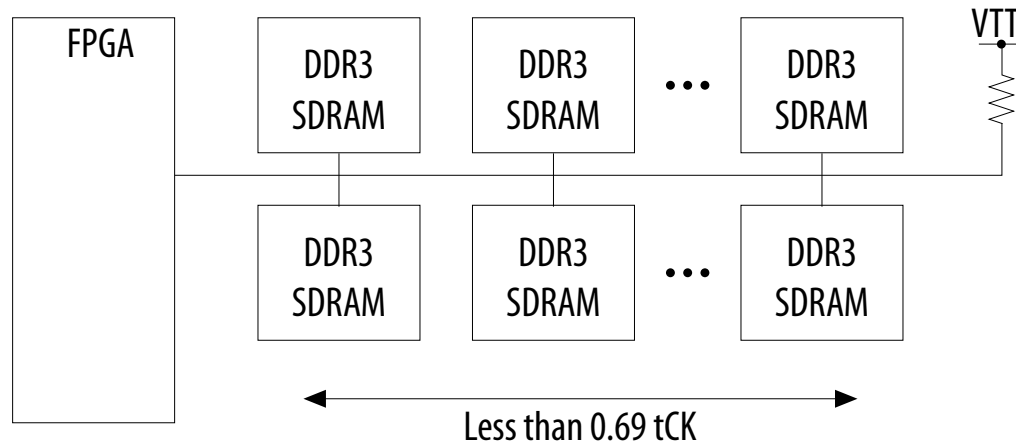
10.4.4.6 Fly-By Network Design for Clock, Command, and Address Signals

The EMIF IP requires the flight-time skew between the first SDRAM component and the last SDRAM component to be less than $0.69 t_{CK}$ for memory clocks. This constraint limits the number of components you can have for each fly-by network.

If you design with discrete components, you can choose to use one or more fly-by networks for the clock, command, and address signals.

The following figure shows an example of a single fly-by network topology.

Figure 90. Single Fly-By Network Topology



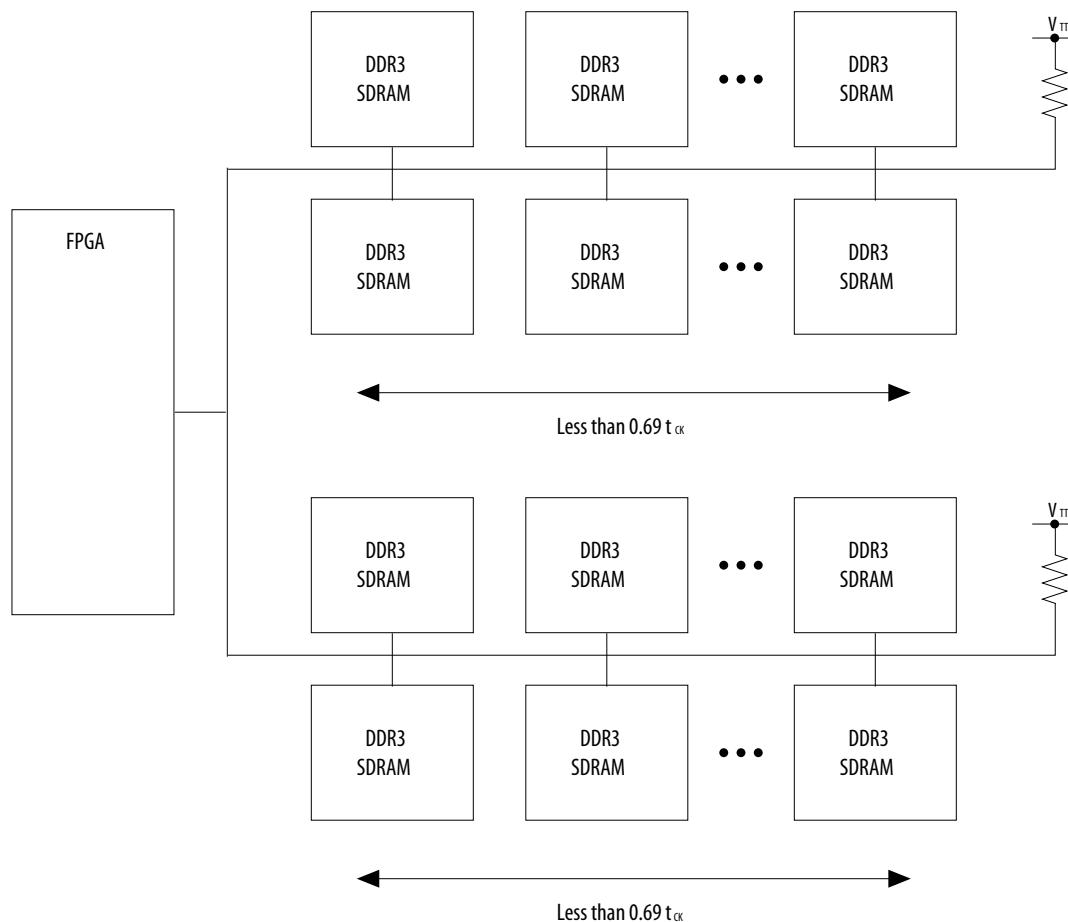
Every SDRAM component connected to the signal is a small load that causes discontinuity and degrades the signal. When using a single fly-by network topology, to minimize signal distortion, follow these guidelines:

- Use ×16 device instead ×4 or ×8 to minimize the number of devices connected to the trace.
- Keep the stubs as short as possible.
- Even with added loads from additional components, keep the total trace length short; keep the distance between the FPGA and the first SDRAM component less than 5 inches.
- Simulate clock signals to ensure a decent waveform.

The following figure shows an example of a double fly-by network topology. This topology is not rigid but you can use it as an alternative option. The advantage of using this topology is that you can have more SDRAM components in a system without violating the 0.69 tCK rule. However, as the signals branch out, the components still create discontinuity.




Figure 91. Double Fly-By Network Topology



You must perform simulations to find the location of the split, and the best impedance for the traces before and after the split.

The following figure shows a way to minimize the discontinuity effect. In this example, keep TL2 and TL3 matches in length. Keep TL1 longer than TL2 and TL3, so that it is easier to route all the signals during layout.

Figure 92. Minimizing Discontinuity Effect



You can also consider using a DIMM on each branch to replace the components. Because the trace impedance on the DIMM card is 40-ohm to 60-ohm, perform a board trace simulation to control the reflection to within the level your system can tolerate.

Using the fly-by daisy chain topology increases the complexity of the datapath and controller design to achieve leveling, but also greatly improves performance and eases board layout for SDRAM implementations.

You can also use the SDRAM components without leveling in a design if it may result in a more optimal solution, or use with devices that support the required electrical interface standard, but do not support the required read and write leveling functionality.

10.4.5 Package Deskew

Trace lengths inside the device package are not uniform for all package pins. The nonuniformity of package traces can affect system timing for high frequencies. A package deskew option is available in the Intel Quartus Prime software.

If you do not enable the package deskew option, the Intel Quartus Prime software uses the package delay numbers to adjust skews on the appropriate signals; you do not need to adjust for package delays on the board traces. If you do enable the package deskew option, the Intel Quartus Prime software does not use the package delay numbers for timing analysis, and you must deskew the package delays with the board traces for the appropriate signals for your design.

Related Links

[Layout Guidelines](#) on page 190

10.4.5.1 DQ/DQS/DM Deskew

To get the package delay information, follow these steps:



1. Select the **FPGA DQ/DQS Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

10.4.5.2 Address and Command Deskew

Deskew address and command delays as follows:

1. Select the **FPGA Address/Command Package Skews Deskewed on Board** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Compile your design.
5. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

10.4.5.3 Package Deskew Recommendations for Intel Stratix 10 Devices

The following table shows package deskew recommendations for Intel Stratix 10 devices.

As operating frequencies increase, it becomes increasingly critical to perform package deskew. The frequencies listed in the table are the *minimum* frequencies for which you must perform package deskew.

If you plan to use a listed protocol at the specified frequency or higher, you must perform package deskew.

Protocol	Minimum Frequency (MHz) for Which to Perform Package Deskew		
	Single Rank	Dual Rank	Quad Rank
DDR4	933	800	667
DDR3	933	800	667
LPDDR3	667	533	Not required
QDR IV	933	Not applicable	Not applicable
RLDRAM 3	933	667	Not applicable
QDR II, II+, II+ Xtreme	Not required	Not applicable	Not applicable

10.4.5.4 Deskew Example

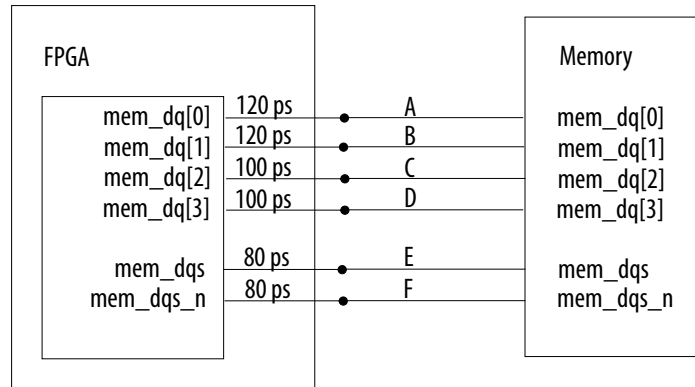
Consider an example where you want to deskew an interface with 4 DQ pins, 1 DQS pin, and 1 DQSn pin.

Let's assume an operating frequency of 667 MHz, and the package lengths for the pins reported in the **.pin** file as follows:

```
dq[0] = 120 ps
dq[1] = 120 ps
dq[2] = 100 ps
dq[3] = 100 ps
dqs = 80 ps
dqs_n = 80 ps
```

The following figure illustrates this example.

Figure 93. Deskew Example

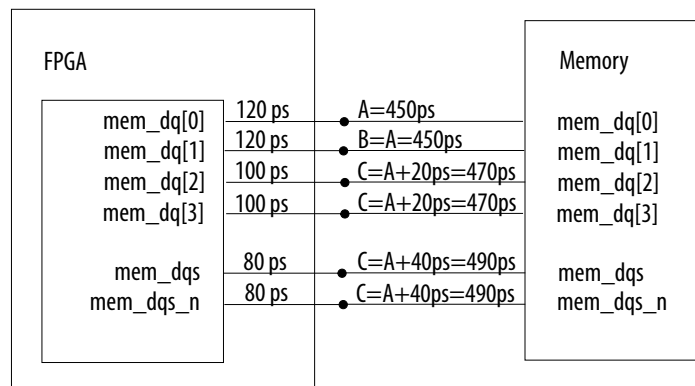


When you perform length matching for all the traces in the DQS group, you must take package delays into consideration. Because the package delays of traces A and B are 40 ps longer than the package delays of traces E and F, you would need to make the board traces for E and F 40 ps longer than the board traces for A and B.

A similar methodology would apply to traces C and D, which should be 20 ps longer than the lengths of traces A and B.

The following figure shows this scenario with the length of trace A at 450 ps.


Figure 94. Deskew Example with Trace Delay Calculations





When you enter the board skews into the Board Settings tab of the DDR3 parameter editor, you should calculate the board skew parameters as the sums of board delay and corresponding package delay. If a pin does not have a package delay (such as address and command pins), you should use the board delay only.

The example of the preceding figure shows an ideal case where board skews are perfectly matched. In reality, you should allow plus or minus 10 ps of skew mismatch within a DQS group (DQ/DQS/DM).



11 Intel Stratix 10 EMIF IP Timing Closure

This chapter describes timing analysis and optimization techniques that you can use to achieve timing closure.

11.1 Timing Closure

The following sections describe the timing analysis using the respective FPGA data sheet specifications and the user-specified memory data sheet parameters.

- Core to core (C2C) transfers have timing constraint created and are analyzed by the Timing Analyzer. Core timing does not include user logic timing within core or to and from EMIF block. The EMIF IP provides the constrained clock to the customer logic.
- Core to periphery (C2P) transfers have timing constraint created and are timing analyzed by the Timing Analyzer. Because of the increased number of C2P/P2C signals in 20nm families compared to previous families, more work is expected to ensure that these special timing arcs are properly modeled, both during timing analysis and compilation.
- Periphery to core (P2C) transfers have timing constraint created and are timing analyzed by the Timing Analyzer. Because of the increased number of C2P/P2C signals in 20nm families compared to previous families, more work is expected to ensure that these special timing arcs are properly modeled, both during timing analysis and compilation.
- Periphery to periphery (P2P) transfers are modeled entirely by a minimum pulse with violation on the hard block, and have no internal timing arc. P2P transfers are modeled only by a minimum pulse width violation on hardened block.

To account for the effects of calibration, the EMIF IP includes additional scripts that are part of the `<phy_variation_name>_report_timing.tcl` and `<phy_variation_name>_report_timing_core.tcl` files that determine the timing margin after calibration. These scripts use the setup and hold slacks of individual pins to emulate what is occurring during calibration to obtain timing margins that are representative of calibrated PHYs. The effects considered as part of the calibrated timing analysis include improvements in margin because of calibration, and quantization error and calibration uncertainty because of voltage and temperature changes after calibration.

Related Links

[Timing Analysis](#) on page 393



11.1.1.1 Timing Analysis

Timing analysis of Intel Stratix 10 EMIF IP is somewhat simpler than that of earlier device families, because Intel Stratix 10 devices have more hardened blocks and there are fewer soft logic registers to be analyzed, because most are user logic registers.

Your Intel Stratix 10 EMIF IP includes a Synopsys Design Constraints File (.sdc) which contains timing constraints specific to your IP. The .sdc file also contains Tool Command Language (.tcl) scripts which perform various timing analyses specific to memory interfaces.

Two timing analysis flows are available for Intel Stratix 10 EMIF IP:

- Early I/O Timing Analysis, which is a precompilation flow.
- Full Timing Analysis, which is a post-compilation flow.

Related Links

[Timing Closure](#) on page 392

11.1.1.1.1 PHY or Core

Timing analysis of the PHY or core path includes the path from the last set of registers in the core to the first set of registers in the periphery (C2P), path from the last set of registers in the periphery to the first of registers in the core (P2C) and ECC related path if it is enabled.

Core timing analysis excludes user logic timing to or from EMIF blocks. The EMIF IP provides a constrained clock (for example: ddr3_usr_clk) with which to clock customer logic; pll_afi_clk serves this purpose.

The PHY or core analyzes this path by calling the `report_timing` command in `<variation_name>_report_timing.tcl` and `<variation_name>_report_timing_core.tcl`.

Note: In version 14.1 and later, the *Spatial Pessimism Removal* slack values in the **Core to Periphery** and **Periphery to Core** tables are always equal to zero. This occurs because pessimism removal is integrated into the base timing analysis.

11.1.1.1.2 I/O Timing

I/O timing analysis includes analysis of read capture, write, address and command, DQS gating, and write leveling.

The Timing Analyzer provides a breakdown of the timing budgets which details margin loss due to transmitter, receiver, and channel. The Timing Analyzer displays the total margin in the last row of the timing report.

The I/O timing analysis described in the following topics is based on a 2 speed-grade device, interfacing with a DDR3 SDRAM UDIMM at 1066 MHz. A 1066 MHz DDR3 SDRAM UDIMM is used for the analysis.

11.1.1.1.2.1 Read Capture

Read capture timing analysis indicates the amount of slack on the DQ signals that are latched by the FPGA using the DQS strobe output of the memory device.

The Timing Analyzer analyzes read capture timing paths through conventional static timing analysis and further processing steps that account for memory calibration (which may include pessimism removal) and calibration uncertainties as shown in the following figure.

Figure 95. Read Capture Timing Analysis

ed_synth_ddr3w Read Capture		
	Operation	Margin
1	Ideal Timing Window	0.469
2	ISI	0.063
3	SSI	0.023
4	Slew Rate Derating	0.050
5	tDQSQ effect	0.075
6	tQH effect	0.112
7	Memory Calibration	-0.075
8	Jitter Effects	0.072
9	Duty Cycle Distortion	0.031
10	Setup/Hold Time	0.016
11	EOL	0.025
12	Calibration Uncertainty	0.036
13	Skew Effect	0.000
14	Final Read Margin	0.040

Channel Effects

Transmitter Effects (Memory)

Receiver Effects (FPGA)

11.1.1.2.2 Write

Write timing analysis indicates the amount of slack on the DQ signals that are latched by the memory device using the DQS strobe output from the FPGA device.

As with read capture, the Timing Analyzer analyzes write timing paths through conventional static timing analysis and further processing steps that account for memory calibration (which may include pessimism removal) and calibration uncertainties as shown in the following figure.

Figure 96. Write Timing Analysis

ed_synth_ddr3w Write		
	Operation	Margin
1	Ideal Timing Window	0.469
2	ISI	0.063
3	SSO	0.047
4	Slew Rate Derating	0.118
5	tDS effect	0.053
6	tDH effect	0.055
7	Memory Calibration	-0.043
8	Jitter Effects	0.039
9	Duty Cycle Distortion	0.016
10	EOL	0.013
11	Calibration Uncertainty	0.038
12	Skew Effect	0.000
13	Final Write Margin	0.071

Channel Effects

Receiver Effects (Memory)

Transmitter Effects (FPGA)

11.1.1.2.3 Address and Command

Address and command signals are single data rate signals latched by the memory device using the FPGA output clock; some are half-rate data signals, while others, such as the chip select, are full-rate signals.



The Timing Analyzer analyzes the address and command timing paths through conventional static timing analysis and further processing steps that account for memory pessimism removal (as shown in the following figure). Depending on the memory protocol in use, if address command calibration is performed, calibration uncertainty is subtracted from the timing window while PVT variation and skew effects are not subtracted, and vice versa

Figure 97. Address and Command Timing Analysis

ed_synth_ddr3w Address/Command			
	Operation	Margin	
1	Ideal Timing Window	0.937	
2	ISI	0.094	
3	SSO	0.023	
4	Slew Rate Derating	0.118	
5	tIS effect	0.060	
6	tIH effect	0.095	
7	Memory Calibration	-0.046	
8	Jitter Effects	0.083	
9	Duty Cycle Distortion	0.016	
10	EOL	0.013	
11	Calibration Uncertainty	0.041	
12	PVT variation	0.027	
13	Skew Effect	0.000	
14	Final CA Margin	0.442	

}

}

Channel Effects

Receiver Effects (Memory)

Transmitter Effects (FPGA)

11.1.1.2.4 DQS Gating / Postamble

Postamble timing is a setup period during which the DQS signal goes low after all the DQ data has been received from the memory device during a read operation. After postamble time, the DQS signal returns from a low-impedance to a high-impedance state to disable DQS and disallow any glitches from writing false data over valid data.

The Timing Analyzer analyzes the postamble timing path in DDRx memory protocols only through an equation which considers memory calibration, calibration uncertainty, and tracking uncertainties as shown in the following figure.

Figure 98. DQS Gating Timing Analysis

ed_synth_ddr3w DQS Gating			
	Operation	Margin	
1	Ideal Timing Window	0.750	
2	ISI	0.094	
3	SSI	0.012	
4	Slew Rate Derating	0.050	
5	tDQSK	0.360	
6	Memory Calibration	-0.144	
7	Jitter Effects	0.167	
8	Duty Cycle Distortion	0.002	
9	EOL	0.002	
10	Calibration Uncertainty	0.020	
11	Tracking Uncertainty	0.078	
12	Setup/Hold Time	0.016	
13	Skew Effect	0.000	
14	Final DQS Gating Margin	0.093	

}

}

Channel Effects

Transmitter Effects (Memory)

Receiver Effects (FPGA)

Intel® Stratix® 10 External Memory Interfaces IP User Guide
395

11.1.1.2.5 Write Leveling

In DDR3 SDRAM and DDR4 SDRAM interfaces, write leveling details the margin for the DQS strobe with respect to CK/CK# at the memory side.

The Timing Analyzer analyzes the write leveling timing path through an equation which considers memory calibration, calibration uncertainty and PVT variation as shown in the following figure.

Figure 99. Write Leveling Timing Analysis

ed_synth_ddr3w Write Levelling		
	Operation	Margin
1	— Ideal Timing Window	0.937
2	— ISI	0.031
3	— SSO	0.035
4	— tDQSS/tDSS/tDSH Effect	0.431
5	— Memory Calibration	-0.172
6	— tWLS/tWLH effect	0.000
7	— Jitter Effects	0.165
8	— Duty Cycle Distortion	0.016
9	— EOL	0.000
10	— Calibration Uncertainty	0.030
11	— PVT variation	0.862
12	— Skew Effect	0.000
13	Final Write Levelling Margin	0.401

Channel Effects

Receiver Effects (Memory)

Transmitter Effects (FPGA)

11.2 Timing Report DDR

The **Report DDR** task in the Timing Analyzer generates custom timing margin reports for all EMIF IP instances in your design. The Timing Analyzer generates this custom report by sourcing the wizard-generated `<variation_name>_report_timing.tcl` script.

This `<variation_name>_report_timing.tcl` script reports the following timing slacks on specific paths of the DDR SDRAM:

- Read capture
- Read resynchronization
- Mimic, address and command
- Core
- Core reset and removal
- Half-rate address and command
- DQS versus CK
- Write
- Write leveling (t_{DQSS})
- Write leveling (t_{DSS}/t_{DSH})
- DQS Gating (Postamble)



The `<variation_name>_report_timing.tcl` script checks basic design rules and assumptions; if violations are found, you receive critical warnings when the Timing Analyzer runs during compilation or when you run the **Report DDR** task.

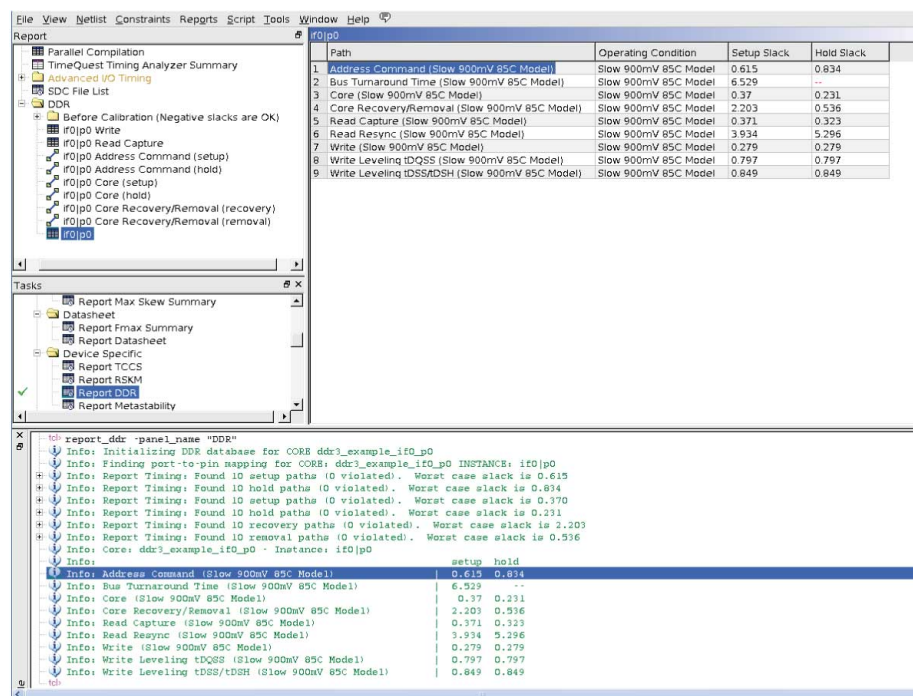
To generate a timing margin report, follow these steps:

1. Compile your design in the Intel Quartus Prime software.
 2. Launch the Timing Analyzer.
 3. Double-click **Report DDR** from the **Tasks** pane. This action automatically executes the **Create Timing Netlist**, **Read SDC File**, and **Update Timing Netlist** tasks for your project.
- The **.sdc** may not be applied correctly if the variation top-level file is the top-level file of the project. You must have the top-level file of the project instantiate the variation top-level file.

The **Report DDR** feature creates a new DDR folder in the Timing Analyzer **Report** pane.

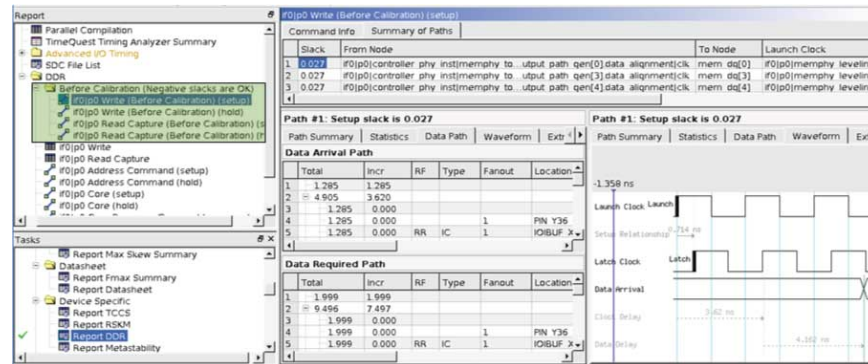
Expanding the DDR folder reveals the detailed timing information for each PHY timing path, in addition to an overall timing margin summary for the instance, as shown in the following figure.

Figure 100. Timing Margin Summary Window Generated by Report DDR Task



The following figure shows the timing analysis results calculated using FPGA timing model before adjustment in the **Before Calibration** panel.

Figure 101. Read and Write Before Calibration



The following two figures show the read capture and write margin summary window generated by the Report DDR Task for a DDR3 core. It first shows the timing results calculated using the FPGA timing model. The `<variation_name>_report_timing_core.tcl` then adjusts these numbers to account for effects that are not modeled by either the timing model or by Timing Analyzer.

Figure 102. Read Capture Margin Summary Window

if0 p0 Read Capture			
	Operation	Setup Slack	Hold Slack
1	After Calibration Read Capture	0.371	0.323
2	Before Calibration Read Capture	0.280	0.273
3	Memory Calibration	0.078	0.150
4	Deskew Read	0.157	0.044
5	Quantization error	-0.050	-0.050
6	Calibration uncertainty	-0.093	-0.093

Figure 103. Write Capture Margin Summary Window

if0 p0 Write			
	Operation	Setup Slack	Hold Slack
1	After Calibration Write	0.279	0.279
2	Before Calibration Write	0.027	0.026
3	Memory Calibration	0.135	0.113
4	Deskew Write and/or more clock pessimism removal	0.216	0.240
5	Quantization error	-0.050	-0.050
6	Calibration uncertainty	-0.050	-0.050

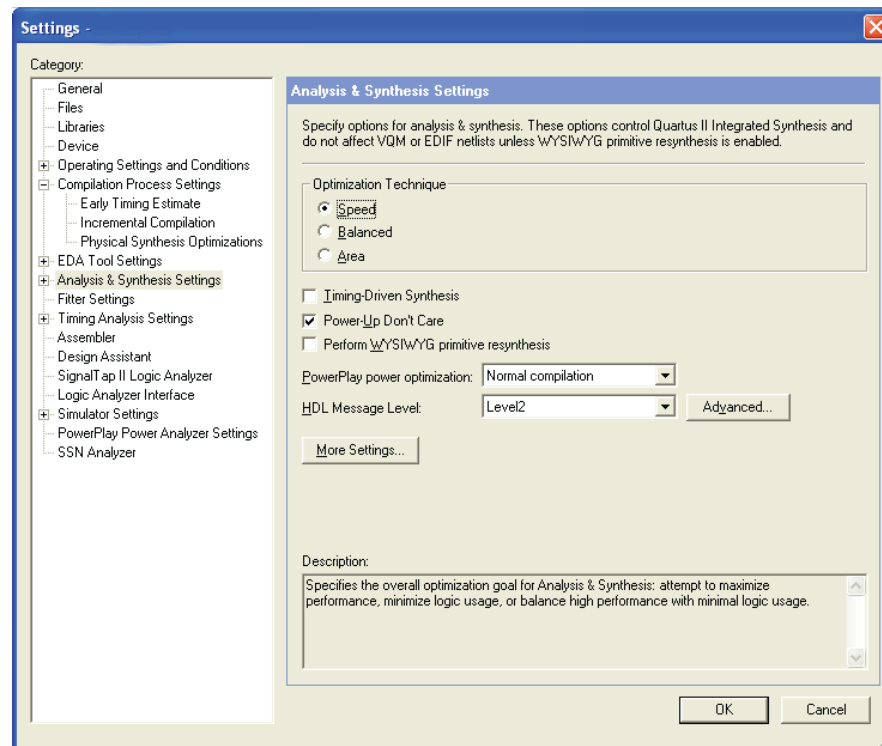
11.3 Optimizing Timing

For full-rate designs you may need to use some of the Intel Quartus Prime advanced features, to meet core timing, by following these steps:

1. On the Assignments menu click **Settings**. In the **Category** list, click **Analysis & Synthesis Settings**. For **Optimization Technique** select **Speed**.

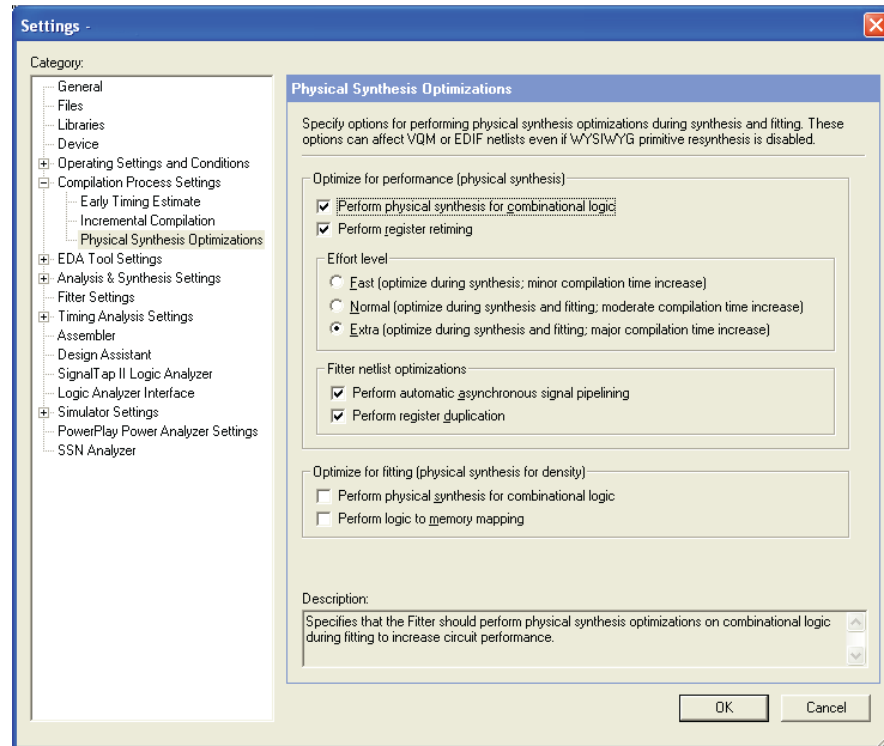


Figure 104. Optimization Technique



- Turn on **Perform physical synthesis for combinational logic**.
For more information about physical synthesis, refer to the *Netlist and Optimizations and Physical Synthesis* chapter in the *Intel Quartus Prime Software Handbook*.
- Turn on **Perform register retiming**
- Turn on **Perform automatic asynchronous signal pipelining**
- Turn on **Perform register duplication**
- You can initially select **Normal** for **Effort level**, then if the core timing is still not met, select **Extra**.

Figure 105. Physical Synthesis Optimizations



Related Links

[Netlist Optimizations and Physical Synthesis](#)

11.4 Early I/O Timing Estimation

Early I/O timing analysis allows you to run I/O timing analysis without first compiling your design. You can use early I/O timing analysis to quickly evaluate whether adequate timing margin exists on the I/O interface between the FPGA and external memory device.

Early I/O timing analysis performs the following analyses:

- Read analysis
- Write analysis
- Address and command analysis
- DQS gating analysis
- Write leveling analysis

Early I/O timing analysis takes into consideration the following factors:

- The timing parameters of the memory device
- The speed and topology of the memory interface
- The board timing and ISI characteristics
- The timing of the selected FPGA device



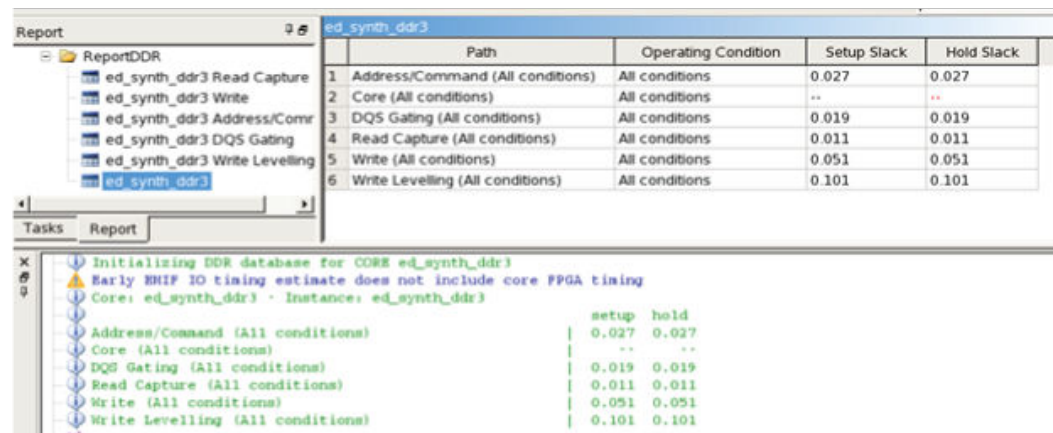
11.4.1 Performing Early I/O Timing Analysis

To perform early I/O timing analysis, follow these steps:


1. Instantiate an EMIF IP core.
 - a. On the **Memory Timing** tab, enter accurate memory parameters.
 - b. On the **Board Timing** tab, enter accurate values for Slew Rate, Intersymbol Interference, and Board and Package Skews.
2. After generating your IP core, create a Intel Quartus Prime project and select your device from the **Available devices** list.
3. To launch the Timing Analyzer, select **Timing Analyzer** from the **Tools** menu.
4. To run early I/O timing analysis:
 - a. Select **Run Tcl Script** from the **Script** menu.
 - b. Run **submodule\<variation_name>_report_io_timing.tcl**.

The following figure shows an early I/O timing analysis from the Timing Analyzer using a DDR3 example design.

Figure 106. Report DDR Timing Results



Report DDR details the read capture, write, address and command, DQS gating, and write leveling timing analyses, which are identical to those obtained after a full design compilation. Core FPGA timing paths are not included in early I/O timing analysis.



12 Optimizing Controller Performance

When designing an external memory interface, you should understand the ways available to increase the efficiency and bandwidth of the memory controller.

The following topics discuss factors that affect controller efficiency and ways to increase the efficiency of the controller.

Controller Efficiency

Controller efficiency varies depending on data transaction. The best way to determine the efficiency of the controller is to simulate the memory controller for your specific design.

Controller efficiency is expressed as:

Efficiency = number of active cycles of data transfer/total number of cycles

The total number of cycles includes the number of cycles required to issue commands or other requests.

Note: You calculate the number of active cycles of data transfer in terms of local clock cycles. For example, if the number of active cycles of data transfer is 2 memory clock cycles, you convert that to the local clock cycle which is 1.

The following cases are based on a high-performance controller design targeting an FPGA device with a CAS latency of 3, and burst length of 4 on the memory side (2 cycles of data transfer), with accessed bank and row in the memory device already open. The FPGA has a command latency of 9 cycles in half-rate mode. The `local_ready` signal is high.

- Case 1: The controller performs individual reads.
Efficiency = $1/(1 + \text{CAS} + \text{command latency}) = 1/(1+1.5+9) = 1/11.5 = 8.6\%$
- Case 2: The controller performs 4 back to back reads.
In this case, the number of data transfer active cycles is 8. The CAS latency is only counted once because the data coming back after the first read is continuous. Only the CAS latency for the first read has an impact on efficiency. The command latency is also counted once because the back to back read commands use the same bank and row.
Efficiency = $4/(4 + \text{CAS} + \text{command latency}) = 4/(4+1.5+9) = 1/14.5 = 27.5\%$

12.1 Interface Standard

Complying with certain interface standard specifications affects controller efficiency.



When interfacing the memory device to the memory controller, you must observe timing specifications and perform the following bank management operations:

- **Activate**

Before you issue any read (RD) or write (WR) commands to a bank within an SDRAM device, you must open a row in that bank using the activate (ACT) command. After you open a row, you can issue a read or write command to that row based on the t_{RCD} specification. Reading or writing to a closed row has negative impact on the efficiency as the controller has to first activate that row and then wait until t_{RCD} time to perform a read or write.

- **Precharge**

To open a different row in the same bank, you must issue a precharge command. The precharge command deactivates the open row in a particular bank or the open row in all banks. Switching a row has a negative impact on the efficiency as you must first precharge the open row, then activate the next row and wait t_{RCD} time to perform any read or write operation to the row.

- **Device CAS latency**

The higher the CAS latency, the less efficient an individual access. The memory device has its own read latency, which is about 12 ns to 20 ns regardless of the actual frequency of the operation. The higher the operating frequency, the longer the CAS latency is in number of cycles.

- **Refresh**

A refresh, in terms of cycles, consists of the precharge command and the waiting period for the auto refresh. Based on the memory data sheet, these components require the following values:

- $t_{RP} = 12$ ns, 3 clock cycles for a 200-MHz operation (5 ns period for 200 MHz)
- $t_{RFC} = 75$ ns, 15 clock cycles for a 200-MHz operation.

Based on this calculation, a refresh pauses read or write operations for 18 clock cycles. So, at 200 MHz, you lose 1.15% (18×5 ns/7.8 us) of the total efficiency.

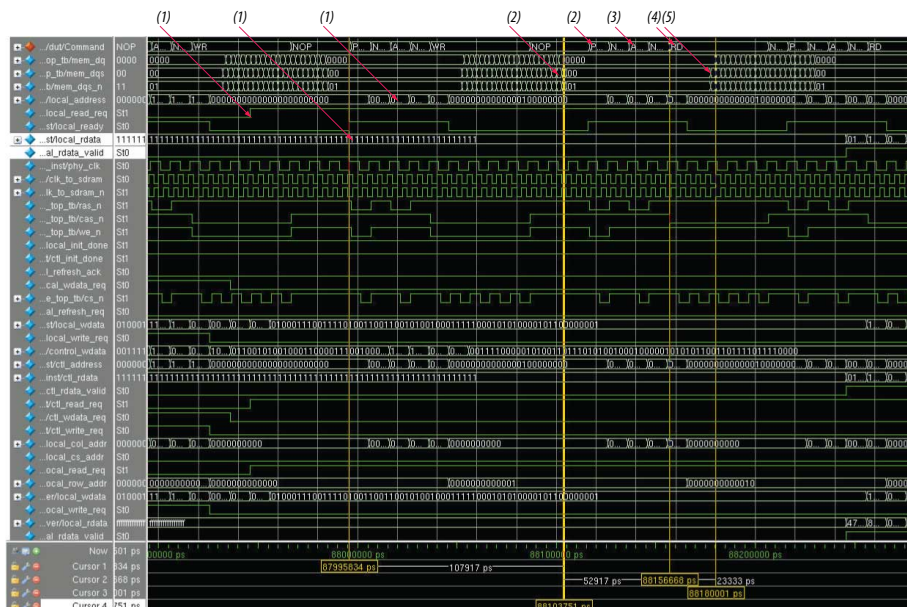
12.2 Bank Management Efficiency

The following figures show examples of how the bank management operations affect controller efficiency.

The first figure shows a read operation in which you have to change a row in a bank. This figure shows how CAS latency and precharge and activate commands affect efficiency.

The following figure illustrates a read-after-write operation. The controller changes the row address after the write-to-read from a different row.

Figure 107. Read Operation—Changing A Row in A Bank



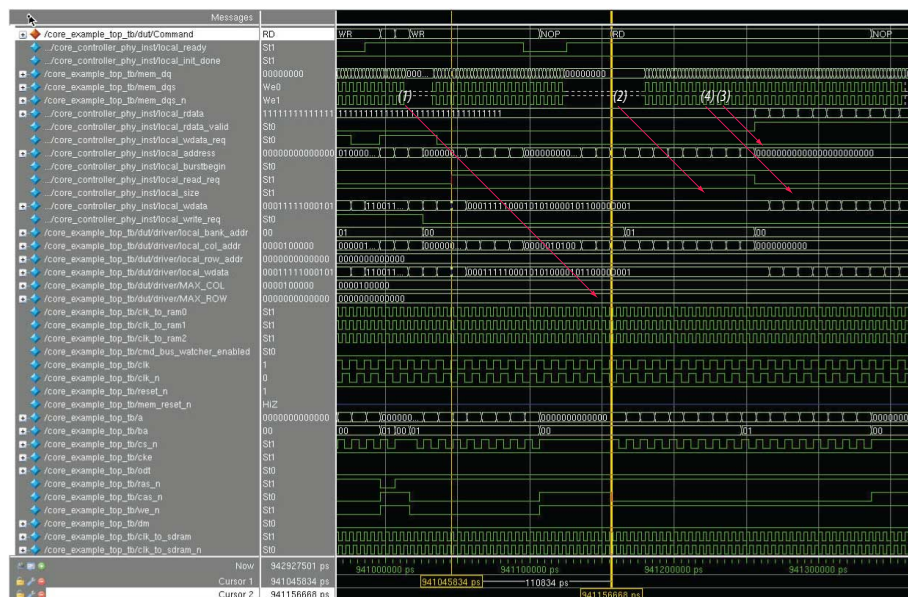
The following sequence of events describes the above figure:

1. The `local_read_req` signal goes high, and when the `local_ready` signal goes high, the controller accepts the read request along with the address.
2. After the memory receives the last write data, the row changes for read. Now you require a precharge command to close the row opened for write. The controller waits for t_{WR} time (3 memory clock cycles) to give the precharge command after the memory receives the last write data.
3. After the controller issues the precharge command, it must wait for t_{RP} time to issue an activate command to open a row.
4. After the controller gives the activate command to activate the row, it needs to wait t_{RCD} time to issue a read command.
5. After the memory receives the read command, it takes the memory some time to provide the data on the pin. This time is known as CAS latency, which is 3 memory clock cycles in this case.

For this particular case, you need approximately 17 local clock cycles to issue a read command to the memory. Because the row in the bank changes, the read operation takes a longer time, as the controller has to issue the precharge and activate commands first. You do not have to take into account t_{WTR} for this case because the precharge and activate operations already exceeded t_{WTR} time.

The following figure shows the case where you use the same the row and bank address when the controller switches from write to read. In this case, the read command latency is reduced.

Figure 108. Changing From Write to Read—Same Row and Bank Address



The following sequence of events describes the above figure:

1. The `local_read_req` signal goes high and the `local_ready` signal is high already. The controller accepts the read request along with the address.
2. When switching from write to read, the controller has to wait t_{WTR} time before it gives a read command to the memory.
3. The SDRAM device receives the read command.
4. After the SDRAM device receives the read command, it takes some time to give the data on the pin. This time is called CAS latency, which is 3 memory clock cycles in this case.

For the case illustrated in the second figure above, you need approximately 11 local clock cycles to issue a read command to the memory. Because the row in the bank remains the same, the controller does not have to issue the precharge and activate commands, which speeds up the read operation and in turn results in a better efficiency compared to the case in the first figure above.

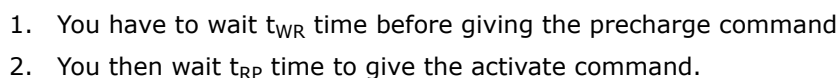
Similarly, if you do not switch between read and write often, the efficiency of your controller improves significantly.

12.3 Data Transfer

The following methods of data transfer reduce the efficiency of your controller:

- Performing individual read or write accesses is less efficient.
- Switching between read and write operation has a negative impact on the efficiency of the controller.
- Performing read or write operations from different rows within a bank or in a different bank—if the bank and a row you are accessing is not already open—also affects the efficiency of your controller.

Figure 109. Changing Row in the Same Bank



- Auto-Precharge Commands
- Additive Latency
- Bank Interleaving
- Command Queue Look-Ahead Depth
- Additive Latency and Bank Interleaving
- User-Controlled Refresh
- Frequency of Operation
- Burst Length
- Series of Reads or Writes

Intel® Stratix® 10 External Memory Interfaces IP User Guide
406



12.4.1 Auto-Precharge Commands

The auto-precharge read and write commands allow you to indicate to the memory device that a given read or write command is the last access to the currently opened row.

The memory device automatically closes or auto-precharges the page that is currently being accessed, so that the next access to the same bank is faster. The Auto-Precharge command is useful when you want to perform fast random memory accesses.

The Timer Bank Pool (TBP) block supports the dynamic page policy, where depending on user input on local autoprecharge input would keep a page open or close. In a closed-page policy, a page is always closed after it is accessed with auto-precharge command. When the data pattern consists of repeated reads or writes to addresses not within the same page, the optimal system achieves the maximum efficiency allowed by continuous page miss limited access. Efficiency losses are limited to those associated with activating and refreshing. An efficiency of 10-20% should be expected for this closed-page policy.

In an open-page policy, the page remains open after it is accessed for incoming commands. When the data pattern consists of repeated reads or writes to sequential addresses within the same page, the optimal system can achieve 100% efficiency for page-open transactions (ignoring the effects of periodic refreshes, which typically consume around 2-3% of total efficiency), with minimum latency for highest priority single transactions.

If you turn on **Enable Auto-Precharge Control**, you can instruct the controller to issue an autoprecharge read or write command. The next time you access that bank, the access will be faster because the controller does not have to precharge the bank before activating the row that you want to access.

The controller-derived autoprecharge logic evaluates the pending commands in the command buffer and determines the most efficient autoprecharge operation to perform. The autoprecharge logic can reorder commands if necessary. When all TBP are occupied due to tracking an open page, TBP uses a scheme called on-demand flush, where it stops tracking a page to create space for an incoming command.

The following figure compares auto-precharge with and without look-ahead support.

Figure 110. Comparison With and Without Look-ahead Auto-Precharge

Without Look-ahead Auto-Precharge			Look-ahead Auto-Precharge		
Cycle	Command	Data	Cycle	Command	Data
1	WRITE		1	WRITE with AP	
2	NOP	DATA0 (Burst 0, Burst 1)	2	NOP	DATA0 (Burst 0, Burst 1)
3	ACT	DATA0 (Burst 2, Burst 3)	3	ACT	DATA0 (Burst 2, Burst 3)
4	NOP	DATA0 (Burst 4, Burst 5)	4	NOP	DATA0 (Burst 4, Burst 5)
5	WRITE	DATA0 (Burst 6, Burst 7)	5	WRITE	DATA0 (Burst 6, Burst 7)
6	NOP	DATA1 (Burst 0, Burst 1)	6	NOP	DATA1 (Burst 0, Burst 1)
7	ACT	DATA1 (Burst 2, Burst 3)	7	ACT	DATA1 (Burst 2, Burst 3)
8	NOP	DATA1 (Burst 4, Burst 5)	8	NOP	DATA1 (Burst 4, Burst 5)
9	WRITE	DATA1 (Burst 6, Burst 7)	9	WRITE	DATA1 (Burst 6, Burst 7)
10	NOP	DATA2 (Burst 0, Burst 1)	10	NOP	DATA2 (Burst 0, Burst 1)
11	PCH	DATA2 (Burst 2, Burst 3)	11	ACT	DATA2 (Burst 2, Burst 3)
12	NOP	DATA2 (Burst 4, Burst 5)	12	NOP	DATA2 (Burst 4, Burst 5)
13	ACT	DATA2 (Burst 6, Burst 7)	13	WRITE	DATA2 (Burst 6, Burst 7)
14	NOP	Wasted Cycle	14	NOP	DATA3 (Burst 0, Burst 1)
15	WRITE	Wasted Cycle	15	NOP	DATA3 (Burst 2, Burst 3)
16	NOP	DATA3 (Burst 0, Burst 1)	16	NOP	DATA3 (Burst 4, Burst 5)
17	NOP	DATA3 (Burst 2, Burst 3)	17	NOP	DATA3 (Burst 6, Burst 7)
18	NOP	DATA3 (Burst 4, Burst 5)			
19	NOP	DATA3 (Burst 6, Burst 7)			

Command	Bank	Row	Condition
Write	Bank 0	Row 0	
Write	Bank 1	Row 0	Activate required
Write	Bank 2	Row 0	Activate required
Write	Bank 0	Row 1	Precharge required

Without using the look-ahead auto-precharge feature, the controller must precharge to close and then open the row before the write or read burst for every row change. When using the look-ahead precharge feature, the controller decides whether to do auto-precharge read/write by evaluating the incoming command; subsequent reads or writes to same bank/different row will require only an activate command.

As shown in the preceding figure, the controller performs an auto-precharge for the write command to bank 0 at cycle 1. The controller detects that the next write at cycle 13 is to a different row in bank 0, and hence saves 2 data cycles.

The following efficiency results apply to the above figure:

Table 439. Comparative Efficiencies With and Without Look-Ahead Auto-Precharge Feature

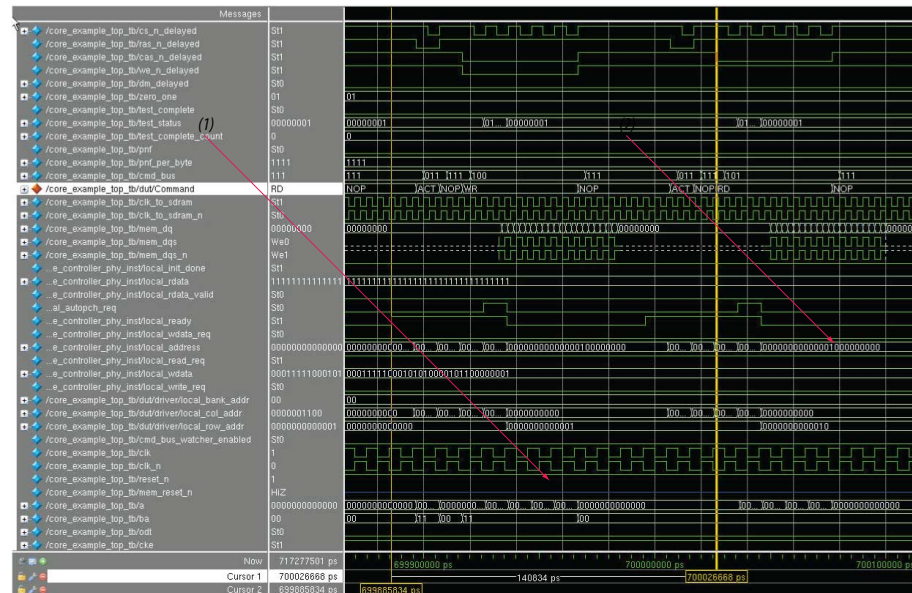
	Without Look-ahead Auto-precharge	With Look-ahead Auto-precharge
Active cycles of data transfer	16	16
Total number of cycles	19	17
Approximate efficiency	84%	94%

The look-ahead auto-precharge used increases efficiency by approximately 10%.

The following figure shows how you can improve controller efficiency using the auto-precharge command.



Figure 111. Improving Efficiency Using Auto-Precharge Command



The following sequence of events describes the above figure:

1. The controller accepts a read request from the local side as soon as the `local_ready` signal goes high.
2. The controller gives the activate command and then gives the read command. The read command latency is approximately 14 clock cycles for this case as compared to the similar case with no auto precharge which had approximately 17 clock cycles of latency (described in the "data Transfer" topic).

When using the auto-precharge option, note the following guidelines:

- Use the auto-precharge command if you know the controller is issuing the next read or write to a particular bank and a different row.
- Auto-precharge does not improve efficiency if you auto-precharge a row and immediately reopen it.

12.4.2 Latency

The following latency data applies to all memory protocols supported by the Intel Stratix 10 EMIF IP.

Table 440. Latency in Full-Rate Memory Clock Cycles

Rate ¹	Controller Address & Command	PHY Address & Command	Memory Read Latency ²	PHY Read Data Return	Controller Read Data Return	Round Trip	Round Trip Without Memory
Half:Write	12	2	3-23	—	—	—	—
Half:Read	8	2	3-23	6	8	27-47	24
Quarter:Write	14	2	3-23	—	—	—	—
continued...							

Rate ¹	Controller Address & Command	PHY Address & Command	Memory Read Latency ²	PHY Read Data Return	Controller Read Data Return	Round Trip	Round Trip Without Memory
Quarter: Read	10	2	3-23	6	14	35-55	32
Half: Write (ECC)	14	2	3-23	—	—	—	—
Half: Read (ECC)	12	2	3-23	6	8	31-51	28
Quarter: Write (ECC)	14	2	3-23	—	—	—	—
Quarter: Read (ECC)	12	2	3-23	6	14	37-57	34

1. User interface rate; the controller always operates in half rate.
2. Minimum and maximum read latency range for DDR3, DDR4, and LPDDR3.

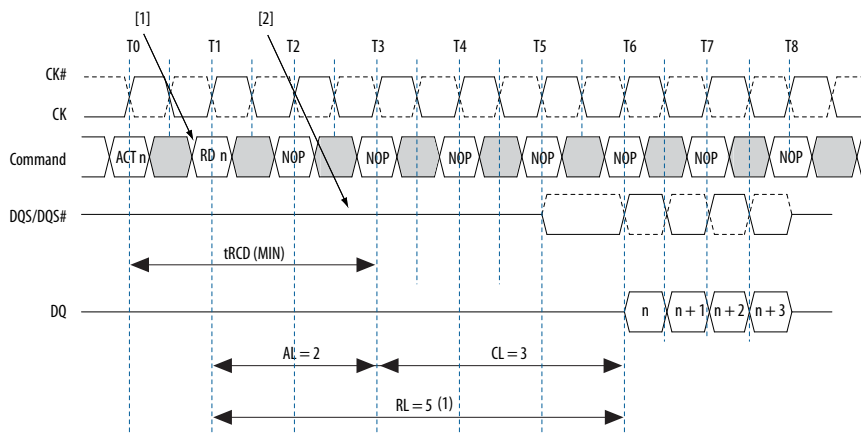
12.4.2.1 Additive Latency

Additive latency increases the efficiency of the command and data bus for sustainable bandwidths.

You may issue the commands externally but the device holds the commands internally for the duration of additive latency before executing, to improve the system scheduling. The delay helps to avoid collision on the command bus and gaps in data input or output bursts. Additive latency allows the controller to issue the row and column address commands—activate, and read or write—in consecutive clock cycles, so that the controller need not hold the column address for several (t_{RCD}) cycles. This gap between the activate and the read or write command can cause bubbles in the data stream.

The following figure shows an example of additive latency.

Figure 112. Additive Latency—Read





The following sequence of events describes the above figure:

1. The controller issues a read or write command before the t_{RCD} (MIN) requirement — additive latency less than or equal to t_{RCD} (MIN).
2. The controller holds the read or write command for the time defined by additive latency before issuing it internally to the SDRAM device.

Read latency = additive latency + CAS latency

Write latency = additive latency + CAS latency - t_{CK}

12.4.3 Calibration

The time needed for calibration varies, depending on many factors including the interface width, the number of ranks, frequency, board layout, and difficulty of calibration.

The following table lists approximate typical calibration times for various protocols and configurations.

Table 441. Intel Stratix 10 EMIF IP Approximate Calibration Times

Protocol	Rank and Frequency	Typical Calibration Time
DDR3, x64 UDIMM, DQS x8, DM on	1 rank, 933 MHz	102 ms
	1 rank, 800 MHz	106 ms
	2 rank, 933 MHz	198 ms
	2 rank, 800 MHz	206 ms
DDR4, x64 UDIMM, DQS x8, DBI on	1 rank, 1067 MHz	314 ms
	1 rank, 800 MHz	353 ms
	2 rank 1067 MHz	625 ms
	2 rank 800 MHz	727 ms
RLDRAM 3, x36	1200 MHz	2808 ms
	1067 MHz	2825 ms
	1200 MHz, with DM	2818 ms
	1067 MHz, with DM	2833 ms
QDR II, x36, BWS on	333 MHz	616 ms
	633 MHz	833 ms
QDR-IV, x36, BWS on	1067 MHz	1563 ms
	1067 MHz, with DBI	1556 ms

12.4.4 Bank Interleaving

You can use bank interleaving to sustain bus efficiency when the controller misses a page, and that page is in a different bank.

Note: Page size refers to the minimum number of column locations on any row that you access with a single activate command. For example: For a 512Mb x8 DDR3 SDRAM with 1024 column locations (column address A[9:0]), page size = 1024 columns x 8 = 8192 bits = 8192/8 bytes = 1024 bytes (1 KB)

Without interleaving, the controller sends the address to the SDRAM device, receives the data requested, and then waits for the SDRAM device to precharge and reactivate before initiating the next data transaction, thus wasting several clock cycles.

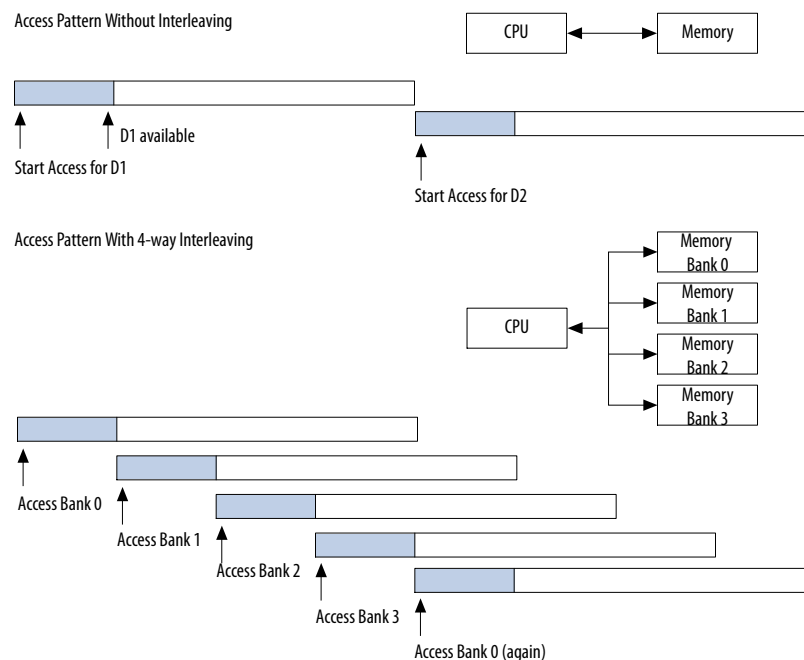
Interleaving allows banks of the SDRAM device to alternate their background operations and access cycles. One bank undergoes its precharge/activate cycle while another is being accessed. By alternating banks, the controller improves its performance by masking the precharge/activate time of each bank. If there are four banks in the system, the controller can ideally send one data request to each of the banks in consecutive clock cycles.

For example, in the first clock cycle, the CPU sends an address to Bank 0, and then sends the next address to Bank 1 in the second clock cycle, before sending the third and fourth addresses to Banks 2 and 3 in the third and fourth clock cycles respectively. The sequence is as follows:

1. Controller sends address 0 to Bank 0.
2. Controller sends address 1 to Bank 1 and receives data 0 from Bank 0.
3. Controller sends address 2 to Bank 2 and receives data 1 from Bank 1.
4. Controller sends address 3 to Bank 3 and receives data 2 from Bank 2.
5. Controller receives data 3 from Bank 3.

The following figure shows how you can use interleaving to increase bandwidth.

Figure 113. Using Interleaving to Increase Bandwidth





The controller supports three interleaving options:

Chip-Bank-Row-Col – This is a noninterleaved option. Select this option to improve efficiency with random traffic

Chip-Row-Bank-Col – This option uses bank interleaving without chip select interleaving. Select this option to improve efficiency with sequential traffic, by spreading smaller data structures across all banks in a chip.

Row-Chip-Bank-Col – This option uses bank interleaving with chip select interleaving. Select this option to improve efficiency with sequential traffic and multiple chip selects. This option allows smaller data structures to spread across multiple banks and chips.

Bank interleaving is a fixed pattern of data transactions, enabling best-case bandwidth and latency, and allowing for sufficient interleaved transactions between opening banks to completely hide t_{RC} . An optimal system can achieve 100% efficiency for bank interleave transactions with 8 banks. A system with less than 8 banks is unlikely to achieve 100%.

12.4.5 Command Queue Look-Ahead Depth

The command queue look-ahead depth value determines the number of read or write requests that the look-ahead bank management logic examines. The command queue look-ahead depth value also determines how many open pages the controller can track.

For example, if you set the command queue look-ahead depth value to 4, the controller can track 4 open pages. In a 4-bank interleaving case, the controller will receive repeated commands with addresses of bank A, bank B, bank C, and bank D. To receive the next set of commands, the controller issues a precharge command to exit the current page and then issues an activate command to track the new incoming page, leading to a drop in efficiency.

Figure 114. Simulation with Command Queue Look-ahead Depth of 4



With the command queue look-ahead depth set to 8, the controller can track 8 open pages and overall efficiency is much improved relative to a command queue look-ahead value of 4.

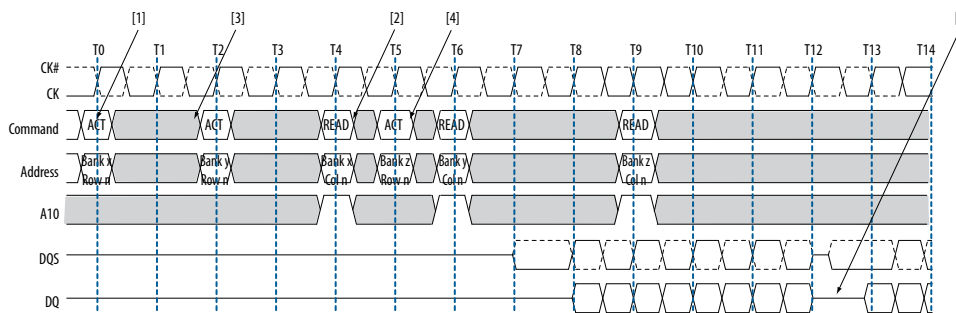
Note: User-selectable Command Queue Look-ahead depth is available only when using the soft memory controller. For the hard memory controller, the Command Queue Look-ahead depth value is hard-coded to 8.

12.4.6 Additive Latency and Bank Interleaving

Using additive latency together with bank interleaving increases the bandwidth of the controller.

The following figure shows an example of bank interleaving in a read operation without additive latency. The example uses bank interleave reads with CAS latency of 4, and burst length of 4.

Figure 116. Bank Interleaving—Without Additive Latency





The following sequence of events describes the above figure:


1. The controller issues an activate command to open the bank, which activates bank x and the row in it.
2. After t_{RCD} time, the controller issues a read with auto-precharge command to the specified bank.
3. Bank y receives an activate command after t_{RRD} time.
4. The controller cannot issue an activate command to bank z at its optimal location because it must wait for bank x to receive the read with auto-precharge command, thus delaying the activate command for one clock cycle.
5. The delay in activate command causes a gap in the output data from the memory device.

Note:

If you use additive latency of 1, the latency affects only read commands and not the timing for write commands.

The following figure shows an example of bank interleaving in a read operation with additive latency. The example uses bank interleave reads with additive latency of 3, CAS latency of 4, and burst length of 4. In this configuration, the controller issues back-to-back activate and read with auto-precharge commands.

Figure 117. Bank Interleaving—With Additive Latency



The following sequence of events describes the above figure:

1. The controller issues an activate command to bank x.
2. The controller issues a read with auto precharge command to bank x right after the activate command, before waiting for the t_{RCD} time.
3. The controller executes the read with auto-precharge command t_{RCD} time later on the rising edge T4.
4. 4 cycles of CAS latency later, the SDRAM device issues the data on the data bus.
5. For burst length of 4, you need 2 cycles for data transfer. With 2 clocks of giving activate and read with auto-precharge commands, you get a continuous flow of output data.

Compare the efficiency results in the two preceding figures:

- bank interleave reads with no additive latency, CAS latency of 4, and burst length of 4 (first figure),

Number of active cycles of data transfer = 6.

Total number of cycles = 15

Efficiency = 40%

- bank interleave reads with additive latency of 3, CAS latency of 4, and burst length of 4 (second figure),

Number of active cycles of data transfer = 6.

Total number of cycles = 14

Efficiency = approximately 43%

The interleaving reads used with additive latency increases efficiency by approximately 3%.

Note: Additive latency improves the efficiency of back-to-back interleaved reads or writes, but not individual random reads or writes.

12.4.7 User-Controlled Refresh

The requirement to periodically refresh memory contents is normally handled by the memory controller; however, the **User Controlled Refresh** option allows you to determine when memory refresh occurs.

With specific knowledge of traffic patterns, you can time the refresh operations so that they do not interrupt read or write operations, thus improving efficiency.

Note: If you enable the auto-precharge control, you must ensure that the average periodic refresh requirement is met, because the controller does not issue any refreshes until you instruct it to.

12.4.7.1 Back-to-Back User-Controlled Refresh Usage

The following diagram illustrates the user-controlled refresh for the hard memory controller (HMC), using the MMR interface.



Figure 118. User-Controlled Refresh via MMR Interface



To perform a user-controlled refresh in the hard memory controller using the MMR interface, follow these steps:

1. Write to the **cfg_user_rfsh_en** register (address=0x019) with the data 0x0000_0010 to enable user refresh.
2. Write to the **mmr_refresh_req** register (address=0x02c) with the data 0x0000_0001 to send a refresh request to rank 0.
 - Note:**
 - Each bit corresponds to one specific rank; for example, data 0x0000_0002 corresponds to rank 1.
 - You may program refreshes to more than one rank at a time.
3. Read from the **mmr_refresh_ack** register (address=0x032) until the readdatavalid signal is asserted and the read data is 1'b1, indicating that a refresh operation is in progress.
4. You can issue the next refresh request only after you see the the acknowledge signal asserted (at time 4).
5. Write to the **mmr_refresh_req** register (address=0x02c) with data 0x0000_0000 to disable the refresh request.
6. You can implement a timer to track tRFC before sending the next user-controlled refresh.

12.4.8 Frequency of Operation

Certain frequencies of operation give you the best possible latency based on the memory parameters. The memory parameters you specify through the parameter editor are converted to clock cycles and rounded up.

If you are using a memory device that has $t_{RCD} = 20$ ns and running the interface at 100 MHz, you get the following results:

- For full-rate implementation ($t_{CK} = 10$ ns):
 t_{RCD} convert to clock cycle = $20/10 = 2$.
- For half rate implementation ($t_{CK} = 20$ ns):
 t_{RCD} convert to clock cycle = $20/20 = 1$

This frequency and parameter combination is not easy to find because there are many memory parameters and frequencies for the memory device and the controller to run. Memory device parameters are optimal for the speed at which the device is designed to run, so you should run the device at that speed.

In most cases, the frequency and parameter combination is not optimal. If you are using a memory device that has $t_{RCD} = 20$ ns and running the interface at 133 MHz, you get the following results:

- For full-rate implementation ($t_{CK} = 7.5$ ns):
 t_{RCD} convert to clock cycle = $20/7.5 = 2.66$, rounded up to 3 clock cycles or 22.5 ns.
- For half rate implementation ($t_{CK} = 15$ ns):
 t_{RCD} convert to clock cycle = $20/15 = 1.33$, rounded up to 2 clock cycles or 30 ns.

There is no latency difference for this frequency and parameter combination.

12.4.9 Series of Reads or Writes

Performing a series of reads or writes from the same bank and row increases controller efficiency.

The case shown in the second figure in the "Bank Management Efficiency" topic demonstrates that a read performed from the same row takes only 14.5 clock cycles to transfer data, making the controller 27% efficient.

Do not perform random reads or random writes. When you perform reads and writes to random locations, the operations require row and bank changes. To change banks, the controller must precharge the previous bank and activate the row in the new bank. Even if you change the row in the same bank, the controller has to close the bank (precharge) and reopen it again just to open a new row (activate). Because of the precharge and activate commands, efficiency can decrease by as much as 3–15%, as the controller needs more time to issue a read or write.

If you must perform a random read or write, use additive latency and bank interleaving to increase efficiency.

Controller efficiency depends on the method of data transfer between the memory device and the FPGA, the memory standards specified by the memory device vendor, and the type of memory controller.

12.4.10 Data Reordering

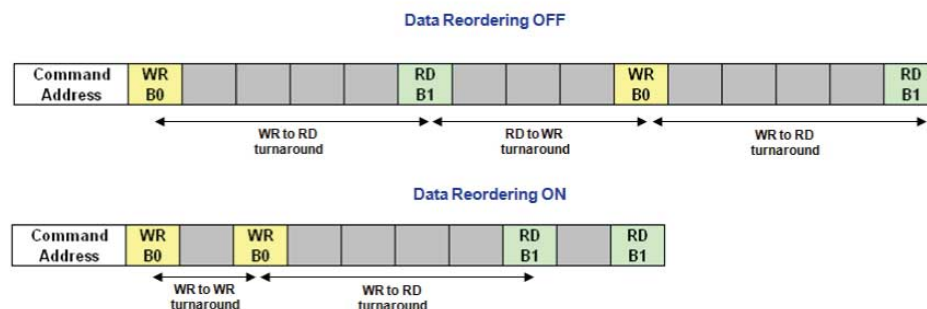
Data reordering and command reordering can both contribute towards achieving controller efficiency.



The Data Reordering feature allows the single-port memory controller to change the order of read and write commands to achieve highest efficiency. You can enable data reordering by turning on **Enable Reordering** on the **Controller Settings** tab of the parameter editor.

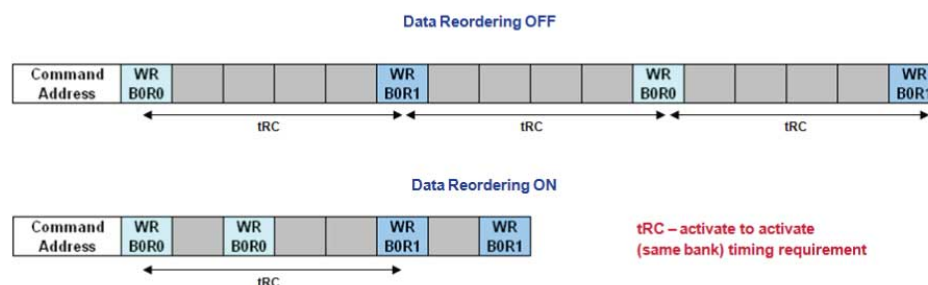
In the soft memory controller, inter-bank data reordering serves to minimize bus turnaround time by optimizing the ordering of read and write commands going to different banks; commands going to the same bank address are not reordered.

Figure 119. Data Reordering for Minimum Bus Turnaround



In the hard memory controller, inter-row data reordering serves to minimize t_{RC} by reordering commands going to different bank and row addresses; command going to the same bank and row address are not reordered. Inter-row data reordering inherits the minimum bus turnaround time benefit from inter-bank data reordering.

Figure 120. Data Reordering for Minimum t_{RC}



12.4.11 Starvation Control

The controller implements a starvation counter to ensure that lower-priority requests are not forgotten as higher-priority requests are reordered for efficiency.

In starvation control, a counter is incremented for every command served. You can set a starvation limit, to ensure that a waiting command is served immediately upon the starvation counter reaching the specified limit.

For example, if you set a starvation limit of 10, a lower-priority command will be treated as high priority and served immediately, after ten other commands are served before it.

12.4.12 Command Reordering

Data reordering and command reordering can both contribute towards achieving controller efficiency.

DDR protocols are naturally inefficient, because commands are fetched and processed sequentially. The DDRx command and DQ bus are not fully utilized as few potential cycles are wasted and degrading the efficiency

The command reordering feature, or look-ahead bank management feature, allows the controller to issue bank management commands early based on incoming patterns, so that when the command reaches the memory interface, the desired page in memory is already open.

The command cycles during the t_{RCD} period are idle and the bank-management commands are issued to next access banks. When the controller is serving the next command, the bank is already precharged. The command queue look-ahead depth is configurable from 1-16, to specify how many read or write requests the look-ahead bank management logic examines. With the look-ahead command queue, if consecutive write or read requests are to a sequential address with same row, same bank, and column incremental by 1, the controller merges the write or read requests at the memory transaction into a single burst.

Figure 121. Comparison With and Without Look-Ahead Bank Mangement Feature

Without Lookahead			With Lookahead		
Cycle	Command	Data	Cycle	Command	Data
1	ACT		1	ACT	
2	NOP		2		
3	NOP		3		
4	READ		4	READ	
5		DATA0 (Burst 0, Burst 1)	5	ACT	DATA0 (Burst 0, Burst 1)
6	NOP	DATA0 (Burst 2, Burst 3)	6	NOP	DATA0 (Burst 2, Burst 3)
7	ACT	DATA0 (Burst 4, Burst 5)	7	ACT	DATA0 (Burst 4, Burst 5)
8	NOP	DATA0 (Burst 6, Burst 7)	8	READ	DATA0 (Burst 6, Burst 7)
9	NOP	Wasted Cycle	9	NOP	DATA1 (Burst 0, Burst 1)
10	READ	Wasted Cycle	10	NOP	DATA1 (Burst 2, Burst 3)
11		DATA1 (Burst 0, Burst 1)	11	NOP	DATA1 (Burst 4, Burst 5)
12	NOP	DATA1 (Burst 2, Burst 3)	12	READ	DATA1 (Burst 6, Burst 7)
13	ACT	DATA1 (Burst 4, Burst 5)	13	NOP	DATA2 (Burst 0, Burst 1)
14	NOP	DATA1 (Burst 6, Burst 7)	14	NOP	DATA2 (Burst 2, Burst 3)
15	NOP	Wasted Cycle	15	NOP	DATA2 (Burst 4, Burst 5)
16	READ	Wasted Cycle	16	NOP	DATA2 (Burst 6, Burst 7)
17	NOP	DATA2 (Burst 0, Burst 1)			
18	NOP	DATA2 (Burst 2, Burst 3)			
19	NOP	DATA2 (Burst 4, Burst 5)			
20	NOP	DATA2 (Burst 6, Burst 7)			

Command	Address	Condition
Read	Bank 0	Activate required
Read	Bank 1	Precharge required
Read	Bank 2	Precharge required



Compare the following efficiency results for the above figure:

Table 442. Efficiency Results for Above Figure

	Without Look-ahead Bank Management	With Look-ahead Bank Management
Active cycles of data transfer	12	12
Total number of cycles	20	16
Approximate efficiency	60%	75%

In the above table, the use of look-ahead bank management increases efficiency by 15%. The bank look-ahead pattern verifies that the system is able to completely hide the bank precharge and activation for specific sequences in which the minimum number of page-open transactions are placed between transactions to closed pages to allow bank look-ahead to occur just in time for the closed pages. An optimal system would completely hide bank activation and precharge performance penalties for the bank look-ahead traffic pattern and achieve 100% efficiency, ignoring refresh.

12.4.13 Bandwidth

Bandwidth depends on the efficiency of the memory controller controlling the data transfer to and from the memory device.

You can express bandwidth as follows:

Bandwidth = data width (bits) × data transfer rate (1/s) × efficiency.


Data rate transfer (1/s) = 2 × frequency of operation (4 × for QDR SRAM interfaces).

The following example shows the bandwidth calculation for a 16-bit interface that has 70% efficiency and runs at 200 MHz frequency:

Bandwidth = 16 bits × 2 clock edges × 200 MHz × 70% = 4.48 Gbps.

DRAM typically has an efficiency of around 70%, but when you use the memory controller, efficiency can vary from 10 to 92%.

In QDR II+ or QDR II SRAM the IP implements two separate unidirectional write and read data buses, so the data transfer rate is four times the clock rate. The data transfer rate for a 400-MHz interface is 1,600 Mbps. The efficiency is the percentage of time the data bus is transferring data. It is dependent on the type of memory. For example, in a QDR II+ or QDR II SRAM interface with separate write and read ports, the efficiency is 100% when there is an equal number of read and write operations on these memory interfaces.



13 Intel Stratix 10 EMIF IP Debugging

This chapter discusses issues and strategies for debugging your external memory interface IP.

For support resources for external memory interface debugging, visit the External Memory Interfaces Support Center on www.altera.com.

Related Links

- [External Memory Interfaces Support Center](#)
- [Analyzing Timing of Memory IP](#)

13.1 Interface Configuration Performance Issues

There are a large number of interface combinations and configurations possible in an Intel design, therefore it is impractical for Intel to explicitly state the achievable f_{MAX} for every combination.

Intel seeks to provide guidance on typical performance, but this data is subject to memory component timing characteristics, interface widths, depths directly affecting timing deration requirements, and the achieved skew and timing numbers for a specific PCB.

FPGA timing issues should generally not be affected by interface loading or layout characteristics. In general, the Intel performance figures for any given device family and speed-grade combination should usually be achievable.

To resolve FPGA (PHY and PHY reset) timing issues, refer to the *Analyzing Timing of Memory IP* chapter.

Achievable interface timing (address and command, half-rate address and command, read and write capture) is directly affected by any layout issues (skew), loading issues (deration), signal integrity issues (crosstalk timing deration), and component speed grades (memory timing size and tolerance). Intel performance figures are typically stated for the default (single rank, unbuffered DIMM) case. Intel provides additional expected performance data where possible, but the f_{MAX} is not achievable in all configurations. Intel recommends that you optimize the following items whenever interface timing issues occur:

- Improve PCB layout tolerances
- Use a faster speed grade of memory component
- Ensure that the interface is fully and correctly terminated
- Reduce the loading (reduce the deration factor)



13.1.1 Interface Configuration Bottleneck and Efficiency Issues

Depending on the transaction types, efficiency issues can exist where the achieved data rate is lower than expected. Ideally, these issues should be assessed and resolved during the simulation stage because they are sometimes impossible to solve later without rearchitecting the product.

Any interface has a maximum theoretical data rate derived from the clock frequency, however, in practice this theoretical data rate can never be achieved continuously due to protocol overhead and bus turnaround times.

Simulate your desired configuration to ensure that you have specified a suitable external memory family and that your chosen controller configuration can achieve your required bandwidth.

Efficiency can be assessed in several different ways, and the primary requirement is an achievable continuous data rate. The local interface signals combined with the memory interface signals and a command decode trace should provide adequate visibility of the operation of the IP to understand whether your required data rate is sufficient and the cause of the efficiency issue.

To show if under ideal conditions the required data rate is possible in the chosen technology, follow these steps:

1. Use the memory vendors own testbench and your own transaction engine.
2. Use either your own driver, or modify the provided example driver, to replicate the transaction types typical of your system.
3. Simulate this performance using your chosen memory controller and decide if the achieved performance is still acceptable.

Observe the following points that may cause efficiency or bottleneck issues at this stage:

- Identify the memory controller rate (full, half, or quarter) and commands, which may take two or four times longer than necessary
- Determine whether the memory controller is starved for data by observing the appropriate request signals.
- Determine whether the memory controller processor transactions at a rate sufficient to meet throughput requirements by observing appropriate signals, including the local ready signal.

Intel has several versions and types of memory controller, and where possible you can evaluate different configurations based on the results of the first tests.

Consider using either a faster interface, or a different memory type to better align your data rate requirements to the IP available directly from Intel.

Intel also provides stand-alone PHY configurations so that you may develop custom controllers or use third-party controllers designed specifically for your requirements.

13.2 Functional Issue Evaluation

Functional issues occur at all frequencies (using the same conditions) and are not altered by speed grade, temperature, or PCB changes. You should use functional simulation to evaluate functional issues.



The Intel FPGA IP includes the option to autogenerate a testbench specific to your IP configuration, which provides an easy route to functional verification.

The following issues should be considered when trying to debug functional issues in a simulation environment.

13.2.1 Intel IP Memory Model

Intel memory IP autogenerates a generic simplified memory model that works in all cases. This simple read and write model is not designed or intended to verify all entered IP parameters or transaction requirements.

The Intel-generated memory model may be suitable to evaluate some limited functional issues, but it does not provide comprehensive functional simulation.

13.2.2 Vendor Memory Model

Contact the memory vendor directly, because many additional models are available from the vendor's support system.

When using memory vendor models, ensure that the model is correctly defined for the following characteristics:

- Speed grade
- Organization
- Memory allocation
- Maximum memory usage
- Number of ranks on a DIMM
- Buffering on the DIMM
- ECC

Note: Refer to the **readme.txt** file supplied with the memory vendor model, for more information about how to define this information for your configuration. Also refer to Transcript Window Messages, for more information.

Note: Intel does not provide support for vendor-specific memory models.

During simulation vendor models output a wealth of information regarding any device violations that may occur because of incorrectly parameterized IP.

13.2.3 Transcript Window Messages

When you are debugging a functional issue in simulation, vendor models typically provide much more detailed checks and feedback regarding the interface and their operational requirements than the Intel generic model.

In general, you should use a vendor-supplied model whenever one is available. Consider using second-source vendor models in preference to the Intel generic model.

Many issues can be traced to incorrectly configured IP for the specified memory components. Component data sheets usually contain settings information for several different speed grades of memory. Be aware data sheet specify parameters in fixed units of time, frequencies, or clock cycles.



The Intel generic memory model always matches the parameters specified in the IP, as it is generated using the same engine. Because vendor models are independent of the IP generation process, they offer a more robust IP parameterization check.

During simulation, review the transcript window messages and do not rely on the Simulation Passed message at the end of simulation. This message only indicates that the example driver successfully wrote and then read the correct data for a single test cycle.

Even if the interface functionally passes in simulation, the vendor model may report operational violations in the transcript window. These reported violations often specifically explain why an interface appears to pass in simulation, but fails in hardware.

Vendor models typically perform checks to ensure that the following types of parameters are correct:

- Burst length
- Burst order
- tMRD
- tMOD
- tRFC
- tREFPDEN
- tRP
- tRAS
- tRC
- tACTPDEN
- tWR
- tWRPDEN
- tRTP
- tRDPDEN
- tINIT
- tXPDLL
- tCKE
- tRRD
- tCCD
- tWTR
- tXPR
- PRECHARGE
- CAS length
- Drive strength
- AL
- tDQS
- CAS_WL

- Refresh
- Initialization
- tIH
- tIS
- tDH
- tDS

If a vendor model can verify all these parameters are compatible with your chosen component values and transactions, it provides a specific insight into hardware interface failures.

13.2.4 Modifying the Example Driver to Replicate the Failure

Often during debugging, you may discover that the example driver design works successfully, but that your custom logic encounters data errors.

When the example design works but your custom design doesn't, the underlying problem may be either of the following:

- Related to the way that the local interface transactions are occurring. You should probe and compare using the Signal Tap II analyzer.
- Related to the types or format of transactions on the external memory interface. You should try modifying the example design to replicate the problem.

Typical issues on the local interface side include:

- Incorrect local-address-to-memory-address translation causing the word order to be different than expected. Refer to *Burst Definition* in your memory vendor data sheet.
- Incorrect timing on the local interface. When your design requests a transaction, the local side must be ready to service that transaction as soon as it is accepted without any pause.
- For more information, refer to the *Avalon® Interface Specification*.

The default example driver performs only a limited set of transaction types, consequently potential bus contention or preamble and postamble issues can often be masked in its default operation. For successful debugging, isolate the custom logic transaction types that are causing the read and write failures and modify the example driver to demonstrate the same issue. Then, you can try to replicate the failure in RTL simulation with the modified driver.

For Intel Stratix 10 interfaces, you can enable the Traffic Generator 2.0 in the example design, allowing you to use the EMIF Debug Toolkit to configure different traffic pattern for debug purposes.

A problem that you can replicate in RTL simulation indicates a potential bug in the IP. You should recheck the IP parameters. A problem that you can not replicate in RTL simulation indicates a timing issue on the PCB. You can try to replicate the issue on an Intel development platform to rule out a board issue.

Note: Ensure that all PCB timing, loading, skew, and deration information is correctly defined in the Intel Quartus Prime software. The timing report is inaccurate if this initial data is not correct.



Functional simulation allows you to identify any issues with the configuration of either the memory controller or the PHY. You can then check the operation against both the memory vendor data sheet and the respective JEDEC specification. After you resolve functional issues, you can start testing hardware.

For more information about simulation, refer to the Simulation chapter.

Related Links

- [Avalon Interface Specifications](#)
- [Intel Stratix 10 EMIF – Simulating Memory IP](#) on page 140

13.3 Timing Issue Characteristics

The PHY and controller combinations autogenerate timing constraint files to ensure that the PHY and external interface are fully constrained and that timing is analyzed during compilation. However, timing issues can still occur. This topic discusses how to identify and resolve any timing issues that you may encounter.

Timing issues typically fall into two distinct categories:

- FPGA core timing reported issues
- External memory interface timing issues in a specific mode of operation or on a specific PCB

Timing Analyzer reports timing issues in two categories: core to core and core to IOE transfers. These timing issues include the PHY and PHY reset sections in the Timing Analyzer Report DDR subsection of timing analysis. External memory interface timing issues are specifically reported in the Timing Analyzer Report DDR subsection, excluding the PHY and PHY reset. The Report DDR PHY and PHY reset sections only include the PHY, and specifically exclude the controller, core, PHY-to-controller and local interface. Intel Quartus Prime timing issues should always be evaluated and corrected before proceeding to any hardware testing.

PCB timing issues are usually Intel Quartus Prime timing issues, which are not reported in the Intel Quartus Prime software, if incorrect or insufficient PCB topology and layout information is not supplied. PCB timing issues are typically characterized by calibration failure, or failures during user mode when the hardware is heated or cooled. Further PCB timing issues are typically hidden if the interface frequency is lowered.

13.3.1 Evaluating FPGA Timing Issues

Usually, you should not encounter timing issues with Intel-provided IP unless your design exceeds Intel's published performance range or you are using a device for which the Intel Quartus Prime software offers only preliminary timing model support. Nevertheless, timing issues can occur in the following circumstances:

- The **.sdc** files are incorrectly added to the Intel Quartus Prime project
- Intel Quartus Prime analysis and synthesis settings are not correct
- Intel Quartus Prime Fitter settings are not correct

For all of these issues, refer to the correct user guide for more information about recommended settings and follow these steps:

1. Ensure that the IP generated **.sdc** files are listed in the Intel Quartus Prime Timing Analyzer files to include in the project window.
2. Ensure that **Analysis and Synthesis Settings** are set to **Optimization Technique Speed**.
3. Ensure that **Fitter Settings** are set to **Fitter Effort Standard Fit**.
4. Use **Timing Analyzer Report Ignored Constraints**, to ensure that **.sdc** files are successfully applied.
5. Use **Timing Analyzer Report Unconstrained Paths**, to ensure that all critical paths are correctly constrained.

More complex timing problems can occur if any of the following conditions are true:

- The design includes multiple PHY or core projects
- Devices where the resources are heavily used
- The design includes wide, distributed, maximum performance interfaces in large die sizes

Any of the above conditions can lead to suboptimal placement results when the PHY or controller are distributed around the FPGA. To evaluate such issues, simplify the design to just the autogenerated example top-level file and determine if the core meets timing and you see a working interface. Failure implies that a more fundamental timing issue exists. If the standalone design passes core timing, evaluate how this placement and fit is different than your complete design.

Use LogicLock regions, or design partitions to better define the placement of your memory controllers. When you have your interface standalone placement, repeat for additional interfaces, combine, and finally add the rest of your design.

Additionally, use fitter seeds and increase the placement and router effort multiplier.

13.3.2 Evaluating External Memory Interface Timing Issues

External memory interface timing issues usually relate to the FPGA input and output characteristics, PCB timing, and the memory component characteristics.

The FPGA input and output characteristics are usually fixed values, because the IOE structure of the devices is fixed. Optimal PLL characteristics and clock routing characteristics do have an effect. Assuming the IP is correctly constrained with autogenerated assignments, and you follow implementation rules, the design should reach the stated performance figures.

Memory component characteristics are fixed for any given component or DIMM. Consider using faster components or DIMMs in marginal cases when PCB skew may be suboptimal, or your design includes multiple ranks when deration may cause read capture or write timing challenges. Using faster memory components often reduces the memory data output skew and uncertainty easing read capture, and lowering the memory's input setup and hold requirement, which eases write timing.



Increased PCB skew reduces margins on address, command, read capture and write timing. If you are narrowly failing timing on these paths, consider reducing the board skew (if possible), or using faster memory. Address and command timing typically requires you to manually balance the reported setup and hold values with the dedicated address and command phase in the IP.

Refer to the respective IP user guide for more information.

Multiple-slot multiple-rank UDIMM interfaces can place considerable loading on the FPGA driver. Typically a quad rank interface can have thirty-six loads. In multiple-rank configurations, Intel's stated maximum data rates are not likely to be achievable because of loading deration. Consider using different topologies, for example registered DIMMs, so that the loading is reduced.

Deration because of increased loading, or suboptimal layout may result in a lower than desired operating frequency meeting timing. You should close timing in the Timing Analyzer software using your expected loading and layout rules before committing to PCB fabrication.

Ensure that any design with an Intel PHY is correctly constrained and meets timing in the Timing Analyzer software. You must address any constraint or timing failures before testing hardware.

For more information about timing constraints, refer to the Timing Analysis chapter.

Related Links

[Analyzing Timing of Memory IP](#)

13.4 Verifying Memory IP Using the Signal Tap II Logic Analyzer

The Signal Tap II logic analyzer shows read and write activity in the system.

For more information about using the Signal Tap II logic analyzer, refer to the *Design Debugging Using the Signal Tap II Embedded Logic Analyzer* chapter in volume 3 of the *Intel Quartus Prime Handbook*

To add the Signal Tap II logic analyzer, follow these steps:

1. On the Tools menu click **Signal Tap II Logic Analyzer**.
2. In the **Signal Configuration** window next to the **Clock** box, click ... (Browse Node Finder).
3. Type the memory interface system clock (typically * phy_clk) in the **Named** box, for **Filter** select **Signal Tap II: presynthesis** and click **List**.
4. Select the memory interface clock that is exposed to the user logic.
5. Click **OK**.
6. Under Signal Configuration, specify the following settings:
 - For **Sample depth**, select **512**
 - For **RAM type**, select **Auto**
 - For **Trigger flow control**, select **Sequential**
 - For **Trigger position**, select **Center trigger position**
 - For **Trigger conditions**, select **1**

7. On the Edit menu, click **Add Nodes**.
8. Search for specific nodes that you want to monitor, and click **Add**.
Note: Signal Tap can probe only nodes that are exposed to FPGA core logic. Refer to pin descriptions for help in deciding which signals to monitor.
9. Decide which signal and event you want to trigger on, and set the corresponding trigger condition.
10. On the File menu, click **Save**, to save the Signal Tap II .stp file to your project.
Note: If you see the message **Do you want to enable Signal Tap II file "stp1.stp" for the current project**, click **Yes**.
11. After you add signals to the Signal Tap II logic analyzer, recompile your design by clicking **Start Compilation** on the **Processing** menu.
12. Following compilation, verify that Timing Analyzer timing analysis passes successfully.
13. Connect the development board to your computer.
14. On the Tools menu, click **Signal Tap II Logic Analyzer**.
15. Add the correct <project_name>.sof file to the SOF Manager:
 - a. Click ... to open the **Select Program Files** dialog box.
 - b. Select <your_project_name>.sof.
 - c. Click **Open**.
 - d. To download the file, click the **Program Device** button.
16. When the example design including Signal Tap II successfully downloads to your development board, click **Run Analysis** to run once, or click **Autorun Analysis** to run continuously.

Related Links

[Design Debugging with the Signal Tap Logic Analyzer](#)

13.4.1 Signals to Monitor with the Signal Tap II Logic Analyzer

This topic lists the memory controller signals you should consider analyzing for different memory interfaces. This list is not exhaustive, but is a starting point.

Monitor the following signals:

- amm_addr
- amm_rdata
- amm_rdata_valid
- amm_read_req
- amm_ready
- amm_wdata
- amm_write_req
- fail
- pass



- `afi_cal_fail`
- `afi_cal_success`
- `test_complete`
- `be_reg` (QDRII only)
- `pnf_per_bit`
- `rdata_reg`
- `rdata_valid_reg`
- `data_out`
- `data_in`
- `written_data_fifo|data_out`
- `usequencer|state *`
- `usequencer|phy_seq_rdata_valid`
- `usequencer|phy_seq_read_fifo_q`
- `usequencer|phy_read_increment_vfifo *`
- `usequencer|phy_read_latency_counter`
- `uread_datapath|afi_rdata_en`
- `uread_datapath|afi_rdata_valid`
- `uread_datapath|ddio_phy_dq`
- `qvld_wr_address *`
- `qvld_rd_address *`

13.5 Hardware Debugging Guidelines

Before debugging your design, confirm that it follows the recommended design flow. Refer to the *Design Flow* chapter in volume 1 of the *External Memory Interface Handbook*.

Always keep a record of tests, to avoid repeating the same tests later. To start debugging the design, perform the following initial steps.

Related Links

[Recommended Design Flow](#)

13.5.1 Create a Simplified Design that Demonstrates the Same Issue

To help debugging, create a simple design that replicates the problem.

A simple design should compile quickly and be easy to understand. The EMIF IP generates an example top-level file that is ideal for debugging. The example top-level file uses all the same parameters, pin-outs, and so on.

If you are using the Intel Stratix 10 example design and the Traffic Generator 2.0, you can configure the traffic pattern using the EMIF Debug Toolkit.



Related Links

[External Memory Interface Debug Toolkit](#)

13.5.2 Measure Power Distribution Network

Measure voltages of the various power supplies on their hardware development platform over a suitable time base and with a suitable trigger.

Ensure that you use an appropriate probe and grounding scheme. In addition, take the measurements directly on the pins or vias of the devices in question, and with the hardware operational.

13.5.3 Measure Signal Integrity and Setup and Hold Margin

Measure the signals on the PCB. When measuring any signal, consider the edge rate of the signal, not just its frequency. Modern FPGA devices have very fast edge rates, therefore you must use a suitable oscilloscope, probe, and grounding scheme when you measure the signals.

You can take measurements to capture the setup and hold time of key signal classes with respect to their clock or strobe. Ensure that the measured setup and hold margin is at least better than that reported in the Intel Quartus Prime software. A worse margin indicates a timing discrepancy somewhere in the project; however, this issue may not be the cause of your problem.

13.5.4 Vary Voltage

Vary the voltage of your system, if you suspect a marginality issue.

Increasing the voltage usually causes devices to operate faster and also usually provides increased noise margin.

13.5.5 Operate at a Lower Speed

Test the interface at a lower speed. If the interface works at a lower speed, the interface is correctly pinned out and functional.

If the interface fails at a lower speed, determine if the test is valid. Many high-speed memory components have a minimal operating frequency, or require subtly different configurations when operating at a lower speeds.

For example, DDR3 SDRAM typically requires modification to the following parameters if you want to operate the interface at lower speeds:

- t_{MRD}
- t_{WTR}
- CAS latency and CAS write latency



13.5.6 Determine Whether the Issue Exists in Previous Versions of Software

Hardware that works before an update to either the Intel Quartus Prime software or the memory IP indicates that the development platform is not the issue.

However, the previous generation IP may be less susceptible to a PCB issue, masking the issue.

13.5.7 Determine Whether the Issue Exists in the Current Version of Software

Designs are often tested using previous generations of Intel software or IP.

Projects may not be upgraded for various reasons:

- Multiple engineers are on the same project. To ensure compatibility, a common release of Intel software is used by all engineers for the duration of the product development. The design may be several releases behind the current Intel Quartus Prime software version.
- Many companies delay before adopting a new release of software so that they can first monitor Internet forums to get a feel for how successful other users say the software is.
- Many companies never use the latest version of any software, preferring to wait until the first service pack is released that fixes the primary issues.
- Some users may only have a license for the older version of the software and can only use that version until their company makes the financial decision to upgrade.
- The local interface specification from Intel FPGA IP to the customer's logic sometimes changes from software release to software release. If you have already spent resources designing interface logic, you may be reluctant to repeat this exercise. If a block of code is already signed off, you may be reluctant to modify it to upgrade to newer IP from Intel.

In all of the above scenarios, you must determine if the issue still exists in the latest version of the Intel software. Bug fixes and enhancements are added to the Intel FPGA IP every release. Depending on the nature of the bug or enhancement, it may not always be clearly documented in the release notes.

Finally, if the latest version of the software resolves the issue, it may be easier to debug the version of software that you are using.

13.5.8 Try A Different PCB

If you are using the same Intel FPGA IP on several different hardware platforms, determine whether the problem occurs on all platforms or just on one.

Multiple instances of the same PCB, or multiple instances of the same interface, on physically different hardware platforms may exhibit different behavior. You can determine if the configuration is fundamentally not working, or if some form of marginality is involved in the issue.



Issues are often reported on the alpha build of a development platform. These are produced in very limited numbers and often have received limited bare-board testing, or functional testing. These early boards are often more unreliable than production quality PCBs.

Additionally, if the IP is from a previous project to help save development resources, determine whether the specific IP configuration works on a previous platform.

13.5.9 Try Other Configurations

Designs are often quite large, using multiple blocks of IP in many different combinations. Determine whether any other configurations work correctly on the development platform.

The full project may have multiple external memory controllers in the same device, or may have configurations where only half the memory width or frequency is required. Find out what does and does not work to help the debugging of the issue.

13.5.10 Debugging Checklist

The following checklist is a good starting point when debugging an external memory interface.

Table 443. Checklist

Check	Item
<input type="checkbox"/>	Try a different fit.
<input type="checkbox"/>	Check IP parameters at the operating frequency (t_{MRD} , t_{WTR} for example).
<input type="checkbox"/>	Ensure you have constrained your design with proper timing deration and have closed timing.
<input type="checkbox"/>	Simulate the design. If it fails in simulation, it will fail in hardware.
<input type="checkbox"/>	Analyze timing.
<input type="checkbox"/>	Place and assign R_{ZQ} (OCT).
<input type="checkbox"/>	Measure the power distribution network (PDN).
<input type="checkbox"/>	Measure signal integrity.
<input type="checkbox"/>	Measure setup and hold timing.
<input type="checkbox"/>	Measure FPGA voltages.
<input type="checkbox"/>	Vary voltages.
<input type="checkbox"/>	Heat and cool the PCB.
<input type="checkbox"/>	Operate at a lower or higher frequency.
<input type="checkbox"/>	Check board timing and trace Information.
continued...	



Check	Item
<input type="checkbox"/>	Check LVDS and clock sources, I/O voltages and termination.
<input type="checkbox"/>	Check PLL clock source, specification, and jitter.
<input type="checkbox"/>	Retarget to a smaller interface width or a single bank.

13.6 Categorizing Hardware Issues

The following topics divide issues into categories. By determining which category (or categories) an issue belongs in, you may be able to better focus on the cause of the issue.

Hardware issues fall into three categories:

- Signal integrity issues
- Hardware and calibration issues
- Intermittent issues

13.6.1 Signal Integrity Issues

Many design issues, including some at the protocol layer, can be traced back to signal integrity problems. You should check circuit board construction, power systems, command, and data signaling to determine if they meet specifications.

If infrequent, random errors exist in the memory subsystem, product reliability suffers. Check the bare circuit board or PCB design file. Circuit board errors can cause poor signal integrity, signal loss, signal timing skew, and trace impedance mismatches. Differential traces with unbalanced lengths or signals that are routed too closely together can cause crosstalk.

13.6.1.1 Characteristics of Signal Integrity Issues

Signal integrity problems often appear when the performance of the hardware design is marginal.

The design may not always initialize and calibrate correctly, or may exhibit occasional bit errors in user mode. Severe signal integrity issues can result in total failure of an interface at certain data rates, and sporadic component failure because of electrical stress. PCB component variance and signal integrity issues often show up as failures on one PCB, but not on another identical board. Timing issues can have a similar characteristic. Multiple calibration windows or significant differences in the calibration results from one calibration to another can also indicate signal integrity issues.

13.6.1.2 Evaluating Signal Integrity Issues

Signal integrity problems can only really be evaluated in two ways:

- direct measurement using suitable test equipment like an oscilloscope and probe
- simulation using a tool like HyperLynx or Allegro PCB SI

Compare signals to the respective electrical specification. You should look for overshoot and undershoot, non-monotonicity, eye height and width, and crosstalk.

13.6.1.2.1 Skew

Ensure that all clocked signals, commands, addresses, and control signals arrive at the memory inputs at the same time.

Trace length variations cause data valid window variations between the signals, reducing margin. For example, DDR3-800 at 400 MHz has a data valid window that is smaller than 1,250 ps. Trace length skew or crosstalk can reduce this data valid window further, making it difficult to design a reliably operating memory interface. Ensure that the skew figure previously entered into the Intel FPGA IP matches that actually achieved on the PCB, otherwise Intel Quartus Prime timing analysis of the interface is accurate.

13.6.1.2.2 Crosstalk

Crosstalk is best evaluated early in the memory design phase.

Check the clock-to-data strobes, because they are bidirectional. Measure the crosstalk at both ends of the line. Check the data strobes to clock, because the clocks are unidirectional, these only need checking at the memory end of the line.

13.6.1.2.3 Power System

Some memory interfaces draw current in spikes from their power delivery system as SDRAMs are based on capacitive memory cells.

Rows are read and refreshed one at a time, which causes dynamic currents that can stress any power distribution network (PDN). The various power rails should be checked either at or as close as possible to the SDRAM power pins. Ideally, you should use a real-time oscilloscope set to fast glitch triggering to check the power rails.

13.6.1.2.4 Clock Signals

The clock signal quality is important for any external memory system.

Measurements include frequency, digital core design (DCD), high width, low width, amplitude, jitter, rise, and fall times.

13.6.1.2.5 Read Data Valid Window and Eye Diagram

The memory generates the read signals. Take measurements at the FPGA end of the line.

To ease read diagram capture, modify the example driver to mask writes or modify the PHY to include a signal that you can trigger on when performing reads.

13.6.1.2.6 Write Data Valid Window and Eye Diagram

The FPGA generates the write signals. Take measurements at the memory device end of the line.

To ease write diagram capture, modify the example driver to mask reads or modify the PHY export a signal that is asserted when performing writes.

13.6.1.2.7 OCT and ODT Usage

Modern external memory interface designs typically use OCT for the FPGA end of the line, and ODT for the memory component end of the line. If either the OCT or ODT are incorrectly configured or enabled, signal integrity problems occur.



If the design uses OCT, the R_{ZQ} pin must be placed correctly for the OCT to work. If you do not place the R_{ZQ} pin, the Intel Quartus Prime software allocates them automatically with the following warning:

```
Critical Warning(12677): No exact pin location assignment(s)
for 1 pins of 122 total pins. For the list of pins please refer to
the I/O Assignment Warnings table in the fitter report.
```

If you see these warnings, the R_{ZQ} pin may have been allocated to a pin that does not have the required external resistor present on the board. This allocation renders the OCT circuit faulty, resulting in unreliable calibration and or interface behavior. The pins with the required external resistor must be specified in the Intel Quartus Prime software.

For the FPGA, ensure that you perform the following:

- Connect the R_{ZQ} pin to the correct resistors and pull-down to ground in the schematic or PCB.
- Contain the R_{ZQ} pins within a bank of the device that is operating at the same VCCIO voltage as the interface that is terminated.
- Review the Fitter Pin-Out file for R_{ZQ} pins to ensure that they are on the correct pins, and that only the correct number of calibration blocks exists in your design.
- Check in the fitter report that the input, output, and bidirectional signals with calibrated OCT all have the termination control block applicable to the associated R_{ZQ} pins.

For the memory components, ensure that you perform the following:

- Connect the required resistor to the correct pin on each and every component, and ensure that it is pulled to the correct voltage.
- Place the required resistor close to the memory component.
- Correctly configure the IP to enable the desired termination at initialization time.
- Check that the speed grade of memory component supports the selected ODT setting.
- Check that the second source part that may have been fitted to the PCB, supports the same ODT settings as the original

13.6.2 Hardware and Calibration Issues

Hardware and calibration issues have the following definitions:

- Calibration issues result in calibration failure, which usually causes the `ctl_cal_fail` signal to be asserted.
- Hardware issues result in read and write failures, which usually causes the `pass not fail (pnf)` signal to be asserted.

Note: Ensure that functional, timing, and signal integrity issues are not the direct cause of your hardware issue, as functional, timing or signal integrity issues are usually the cause of any hardware issue.

13.6.2.1 Postamble Timing Issues and Margin

The postamble timing is set by the PHY during calibration.

You can diagnose postamble issues by viewing the `pnf_per_byte` signal from the example driver. Postamble timing issues mean only read data is corrupted during the last beat of any read request.

13.6.2.2 Intermittent Issue Evaluation

Intermittent issues are typically the hardest type of issue to debug—they appear randomly and are hard to replicate.

Errors that occur during run-time indicate a data-related issue, which you can identify by the following actions:

- Add the Signal Tap II logic analyzer and trigger on the post-trigger `pnf`
- Use a stress pattern of data or transactions, to increase the probability of the issue
- Heat up or cool down the system
- Run the system at a slightly faster frequency

If adding the Signal Tap II logic analyzer or modifying the project causes the issue to go away, the issue is likely to be placement or timing related.

Errors that occur at start-up indicate that the issue is related to calibration, which you can identify by the following actions:

- Modify the design to continually calibrate and reset in a loop until the error is observed
- Where possible, evaluate the calibration margin either from the debug toolkit or system console.
- Capture the calibration error stage or error code, and use this information with whatever specifically occurs at that stage of calibration to assist with your debugging of the issue.

Related Links

[External Memory Interface Debug Toolkit](#)

13.7 Debugging Intel Stratix 10 EMIF IP

You can debug hardware failures by connecting to the EMIF Debug Toolkit or by exporting an Avalon-MM slave port, from which you can access information gathered during calibration. You can also connect to this port to mask ranks and to request recalibration.

You can access the exported Avalon-MM port in two ways:

- Via the External Memory Interface Debug Toolkit
- Via On-Chip Debug (core logic on the FPGA)

13.7.1 External Memory Interface Debug Toolkit

The External Memory Interface Debug Toolkit provides access to data collected by the Nios II sequencer during memory calibration, and allows you to perform certain tasks.



The External Memory Interface Debug Toolkit provides access to data including the following:

- General interface information, such as protocol and interface width
- Calibration results per group, including pass/fail status, failure stage, and delay settings

You can also perform the following tasks:

- Mask ranks from calibration (you might do this to skip specific ranks)
- Request recalibration of the interface

13.7.2 On-Chip Debug for Intel Stratix 10

The On-Chip Debug feature allows user logic to access the same debug capabilities as the External Memory Interface Toolkit. You can use On-Chip Debug to monitor the calibration results of an external memory interface, without a connected computer.

To use On-Chip Debug, you need a C header file which is provided as part of the external memory interface IP. The C header file defines data structures that contain calibration data, and definitions of the commands that can be sent to the memory interface.

The On-Chip Debug feature accesses the data structures through the Avalon-MM port that is exposed by the EMIF IP when you turn on debugging features.

13.7.3 Configuring Your EMIF IP for Use with the Debug Toolkit

The Intel Stratix 10 EMIF Debug Interface IP core contains the access point through which the EMIF Debug Toolkit reads calibration data collected by the Nios II sequencer.

Connecting an EMIF IP Core to an Intel Stratix 10 EMIF Debug Interface

For the EMIF Debug Toolkit to access the calibration data for a Intel Stratix 10 EMIF IP core, you must connect one of the EMIF cores in each I/O column to a Intel Stratix 10 EMIF Debug Interface IP core. Subsequent EMIF IP cores in the same column must connect in a daisy chain to the first.

There are two ways that you can add the Intel Stratix 10 EMIF Debug Interface IP core to your design:

- When you generate your EMIF IP core, on the **Diagnostics** tab, select **Add EMIF Debug Interface** for the **EMIF Debug Toolkit/On-Chip Debug Port**; you do not have to separately instantiate a Intel Stratix 10 EMIF Debug Interface core. This method does not export an Avalon-MM slave port. You can use this method if you require only EMIF Debug Toolkit access to this I/O column; that is, if you do not require On-Chip Debug Port access, or PHYLite reconfiguration access.
- When you generate your EMIF IP core, on the **Diagnostics** tab, select **Export** for the **EMIF Debug Toolkit/On-Chip Debug Port**. Then, separately instantiate an Intel Stratix 10 EMIF Debug Interface core and connect its `to_ioaux` interface to the `cal_debug` interface on the EMIF IP core. This method is appropriate if you want to also have On-Chip Debug Port access to this I/O column, or PHYLite reconfiguration access.

For each of the above methods, you must assign a unique interface ID for each external memory interface in the I/O column, to identify that interface in the Debug Toolkit. You can assign an interface ID using the dropdown list that appears when you enable the **Debug Toolkit/On-Chip Debug Port** option.

Daisy-Chaining Additional EMIF IP Cores for Debugging

After you have connected a Intel Stratix 10 EMIF Debug Interface to one of the EMIF IP cores in an I/O column, you must then connect subsequent EMIF IP cores in that column in a daisy-chain manner. If you don't require debug capabilities for a particular EMIF IP core, you do not have to connect that core to the daisy chain.

To create a daisy chain of EMIF IP cores, follow these steps:

1. On the first EMIF IP core, select **Enable Daisy-Chaining for EMIF Debug Toolkit/On-Chip Debug Port** to create an Avalon-MM interface called `cal_debug_out`.
2. On the second EMIF IP core, select **Export** as the **EMIF Debug Toolkit/On-Chip Debug Port** mode, to export an Avalon-MM interface called `cal_debug`.
3. Connect the `cal_debug_out` interface of the first EMIF core to the `cal_debug` interface of the second EMIF core.
4. To connect more EMIF cores to the daisy chain, select the **Enable Daisy-Chaining for EMIF Debug Toolkit/On-Chip Debug Port** option on the second core, connect it to the next core using the **Export** option as described above. Repeat the process for subsequent EMIF cores.

If you place any PHYLite cores with dynamic reconfiguration enabled into the same I/O column as an EMIF IP core, you should instantiate and connect the PHYLite cores in a similar way. See the *Altera PHYLite for Memory Megafunction User Guide* for more information.

Related Links

[Intel FPGA PHYLite for Parallel Interfaces IP Core User Guide](#)

13.7.4 Example Tcl Script for Running the EMIF Debug Toolkit

If you want, you can run the EMIF Debug Toolkit using a Tcl script. The following example Tcl script is applicable to all device families.

The following example Tcl script opens a file, runs the debug toolkit, and writes the resulting calibration reports to a file.

You should adjust the variables in the script to match your design. You can then run the script using the command `quartus_sh -t example.tcl`.

```
# Modify the following variables for your project
set project "ed_synth.qpf"
# Index of the programming cable. Can be listed using "get_hardware_names"
set hardware_index 1
# Index of the device on the specified cable. Can be listed using
"get_device_names"
set device_index 1
# SOF file containing the EMIF to debug
set sof "ed_synth.sof"
# Connection ID of the EMIF debug interface. Can be listed using
"get_connections"
set connection_id 2
# Output file
```




```

set report "toolkit.rpt"

# The following code opens a project and writes its calibration reports to a
# file.
project_open $project
load_package ::quartus::external_memif_toolkit
initialize_connections
set hardware_name [lindex [get_hardware_names] $hardware_index]
set device_name [lindex [get_device_names -hardware_name $hardware_name]
$device_index]
link_project_to_device -device_name $device_name -hardware_name
$hardware_name -sof_file $sof
establish_connection -id $connection_id
create_connection_report -id $connection_id -report_type summary
create_connection_report -id $connection_id -report_type calib
write_connection_target_report -id $connection_id -file $report

```

13.7.5 Using the EMIF Debug Toolkit with Intel Stratix 10 HPS Interfaces

The External Memory Interface Debug Toolkit is not directly compatible with Intel Stratix 10 HPS interfaces.

To debug your Intel Stratix 10 HPS interface using the EMIF Debug Toolkit, you should create an identically parameterized, non-HPS version of your interface, and apply the EMIF Debug Toolkit to that interface. When you finish debugging this non-HPS interface, you can then apply any needed changes to your HPS interface, and continue your design development.

13.7.6 Intel Stratix 10 EMIF Debugging Examples


This topic provides examples of debugging a single external memory interface, and of adding additional EMIF instances to an I/O column.

Debugging a Single External Memory Interface

1. Under **EMIF Debug Toolkit/On-Chip Debug Port**, select **Add EMIF Debug Interface**.

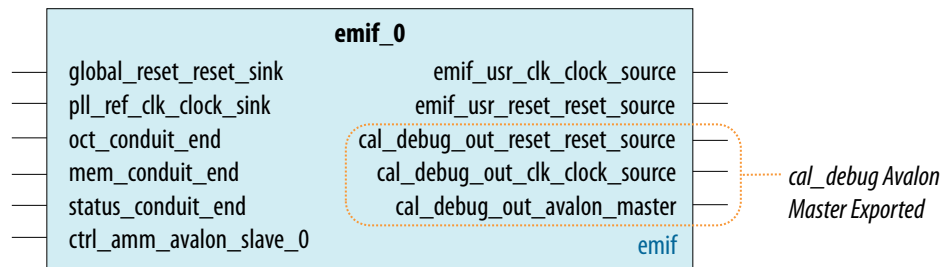
(If you want to use the On-Chip Debug Port instead of the EMIF Debug Toolkit, select **Export** instead.)

Figure 122. EMIF With Debug Interface Added (No Additional Ports)



2. If you want to connect additional EMIF or PHYLite components in this I/O column, select *Enable Daisy Chaining for EMIF Debug Toolkit/On-Chip Debug Port*.

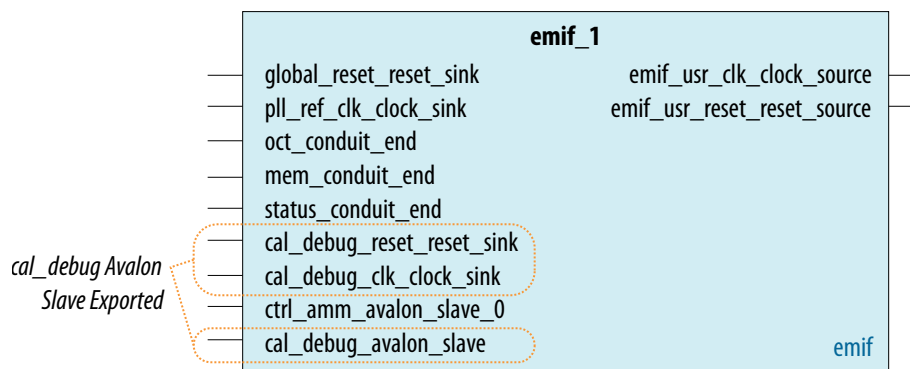
Figure 123. EMIF With cal_debug Avalon Master Exported



Adding Additional EMIF Instances to an I/O Column

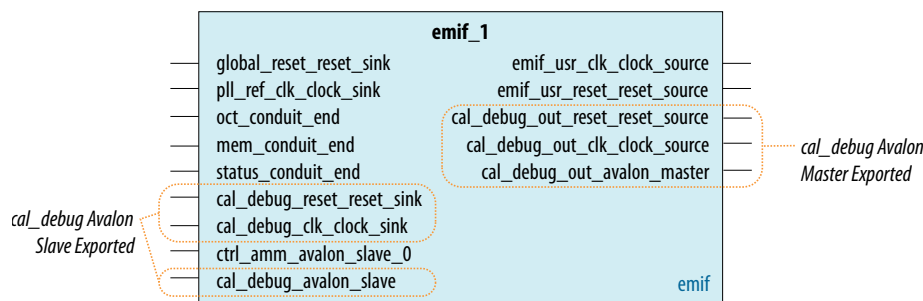
- Under **EMIF Debug Toolkit/On-Chip Debug Port**, select **Export**.

Figure 124. EMIF With cal_debug Avalon Slave Exported



- Specify a unique interface ID for this EMIF instance.
- If you want to connect additional EMIF or PHYLite components in this I/O column, select *Enable Daisy Chaining for EMIF Debug Toolkit/On-Chip Debug Port*.


Figure 125. EMIF With Both cal_debug Master and Slave Exported



- Connect the cal_debug Avalon Master, clock, and reset interfaces of the previous component to the cal_debug Avalon Slave, clock, and reset interfaces of this component.



Figure 126. EMIF Components Connected



13.8 User Interface

The EMIF toolkit provides a graphical user interface for communication with connections.

All functions provided in the toolkit are also available directly from the `quartus_sh` TCL shell, through the `external_memif_toolkit` TCL package. The availability of TCL support allows you to create scripts to run automatically from TCL. You can find information about specific TCL commands by running `help -pkg external_memif_toolkit` from the `quartus_sh` TCL shell.

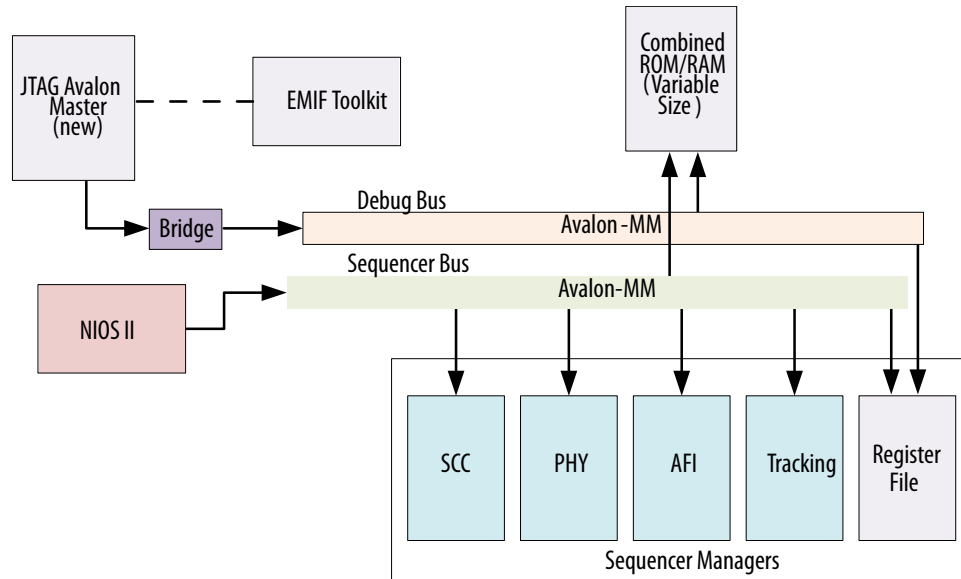
If you want, you can begin interacting with the toolkit through the GUI, and later automate your workflow by creating TCL scripts. The toolkit GUI records a history of the commands that you run. You can see the command history on the History tab in the toolkit GUI.

13.8.1 Communication

Communication between the EMIF Toolkit and external memory interface connections is achieved using a JTAG Avalon-MM master attached to the sequencer bus.

The following figure shows the structure of EMIF IP with JTAG Avalon-MM master attached to sequencer bus masters.

Figure 127. EMIF IP with JTAG Avalon-MM Master



13.8.2 Setup and Use

Before using the EMIF Toolkit, you should compile your design and program the target device with the resulting SRAM Object File (.sof). For designs compiled in the Intel Quartus Prime software, all debugging information resides in the .sof file.

You can run the toolkit using all your project files, or using only the Intel Quartus Prime Project File (.qpf), Intel Quartus Prime Settings File (.qsf), and .sof file.

After you have programmed the target device, you can run the EMIF Toolkit and open your project. You can then use the toolkit to create connections to the external memory interface.

13.8.2.1 General Workflow

To use the EMIF Toolkit, you must link your compiled project to a device, and create a communication channel to the connection that you want to examine.

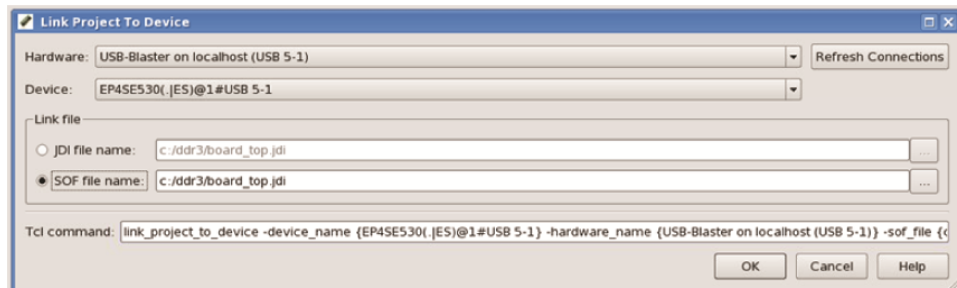
13.8.2.2 Linking the Project to a Device

1. To launch the toolkit, select External Memory Interface Toolkit from the Tools menu in the Intel Quartus Prime software.
2. After you have launched the toolkit, open your project and click the **Initialize connections** task in the **Tasks** window, to initialize a list of all known connections.
3. To link your project to a specific device on specific hardware, perform the following steps:



- a. Click the **Link Project to Device** task in the **Tasks** window.
- b. Select the desired hardware from the **Hardware** dropdown menu in the **Link Project to Device** dialog box.
- c. Select the desired device on the hardware from the **Device** dropdown menu in the **Link Project to Device** dialog box.
- d. Select **SOF** as the **Link file type**, verify that the **.sof** file is correct for your programmed device, and click **Ok**.

Figure 128. Link Project to Device Dialog Box



For designs compiled in the Intel Quartus Prime software, the SOF file contains a design hash to ensure the SOF file used to program the device matches the SOF file specified for linking to a project. If the hash does not match, an error message appears.

If the toolkit successfully verifies all connections, it then attempts to determine the connection type for each connection. Connections of a known type are listed in the Linked Connections report, and are available for the toolkit to use.

13.8.2.3 Establishing Communication to Connections

After you have completed linking the project, you can establish communication to the connections.

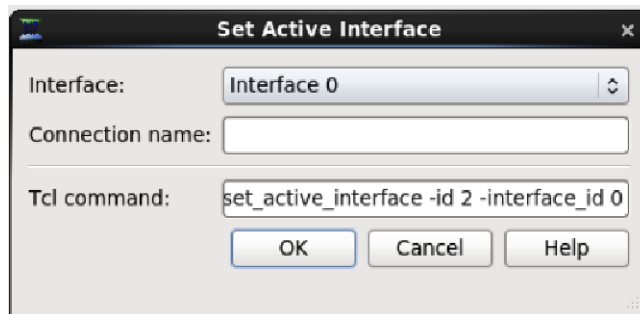
1. In the Tasks window,
 - Click **Establish Memory Interface Connection** to create a connection to the external memory interface.
 - Click **Establish Efficiency Monitor Connection** to create a connection to the efficiency monitor.
 - Click **Establish Traffic Generator Connection** to create a connection to the Traffic Generator 2.0.
2. To create a communication channel to a connection, select the desired connection from the displayed pulldown menu of connections, and click **Ok**. The toolkit establishes a communication channel to the connection, creates a report folder for the connection, and creates a folder of tasks for the connection.

Note: By default, the connection and the reports and tasks folders are named according to the hierarchy path of the connection. If you want, you can specify a different name for the connection and its folders.
3. You can run any of the tasks in the folder for the connection; any resulting reports appear in the reports folder for the connection.

13.8.2.4 Selecting an Active Interface

If you have more than one external memory interface in an I/O column, you can select one instance as the active interface for debugging.

1. To select one of multiple EMIF instances in an I/O column, use the **Set Active Interface** dialog box.



2. If you want to generate reports for the new active interface, you must first recalibrate the interface.

13.8.2.5 Reports

The toolkit can generate a variety of reports, including summary, calibration, and margining reports for external memory interface connections. To generate a supported type of report for a connection, you run the associated task in the tasks folder for that connection.

Summary Report

The Summary Report provides an overview of the memory interface; it consists of the following tables:

- Summary table. Provides a high-level summary of calibration results. This table lists details about the connection, IP version, IP protocol, and basic calibration results, including calibration failures. This table also lists the estimated average read and write data valid windows, and the calibrated read and write latencies.
- Interface Details table. Provides details about the parameterization of the memory IP. This table allows you to verify that the parameters in use match the actual memory device in use.
- Ranks Masked from Calibration tables (DDR3 only). Lists any ranks that were masked from calibration when calibration occurred. Masked ranks are ignored during calibration.



Calibration Report

The Calibration Report provides detailed information about the margins observed during calibration, and the settings applied to the memory interface during calibration; it consists of the following tables:

- Calibration Status Per Group table: Lists the pass/fail status per group.
- DQ Pin Margins Observed During Calibration table: Lists the DQ read/write margins and calibrated delay settings. These are the expected margins after calibration, based on calibration data patterns. This table also contains DM/DBI margins, if applicable.
- DQS Pin Margins Observed During Calibration table: Lists the DQS margins observed during calibration.
- FIFO Settings table: Lists the VFIFO and LFIFO settings made during calibration.
- Latency Observed During Calibration table: Lists the calibrated read/write latency.
- Address/Command Margins Observed During Calibration table: Lists the margins on calibrated A/C pins, for protocols that support Address/Command calibration.

13.9 On-Chip Debug Port for Intel Stratix 10 EMIF IP

The EMIF On-Chip Debug Port allows user logic to access the same calibration data used by the EMIF Toolkit, and allows user logic to send commands to the sequencer. You can use the EMIF On-Chip Debug Port to access calibration data for your design and to send commands to the sequencer just as the EMIF Toolkit would. The following information is available:


- Pass/fail status for each DQS group
- Read and write data valid windows for each group

In addition, user logic can request the following commands from the sequencer:

- Destructive recalibration of all groups
- Masking of groups and ranks
- Generation of per-DQ pin margining data as part of calibration

The user logic communicates through an Avalon-MM slave interface as shown below.

Figure 129. User Logic Access



Related Links

[Intel FPGA PHYLite for Parallel Interfaces IP Core User Guide](#)

13.9.1 EMIF On-Chip Debug Port

Access to on-chip debug is provided through software running on a Nios processor connected to the external memory interface.

If you enable the Use Soft Nios Processor for On-Chip Debug option, the system instantiates a soft Nios processor, and software files are provided as part of the EMIF IP.

Instructions on how to use the software are available in the following file: :
`<variation_name>/altera_emif_arch_nf_<version number>/<synth|sim>/<variation_name>_altera_emif_arch_nf_<version number>_<unique ID>_readme.txt.`

13.9.2 Access Protocol

The On-Chip Debug Port provides access to calibration data through an Avalon-MM slave interface. To send a command to the sequencer, user logic sends a command code to the command space in sequencer memory. The sequencer polls the command space for new commands after each group completes calibration, and continuously after overall calibration has completed.

The communication protocol to send commands from user logic to the sequencer uses a multistep handshake with a data structure as shown below, and an algorithm as shown in the figure which follows.

```
typedef struct_debug_data_struct {
...
// Command interaction
alt_u32 requested_command;
alt_u32 command_status;
alt_u32 command_parameters[COMMAND_PARAM_WORDS];...
}
```

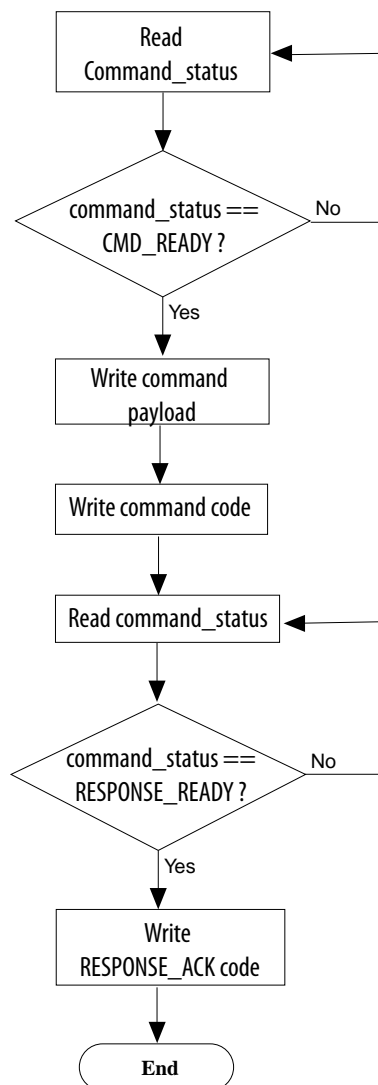
To send a command to the sequencer, user logic must first poll the `command_status` word for a value of `TCLDBG_TX_STATUS_CMD_READY`, which indicates that the sequencer is ready to accept commands. When the sequencer is ready to accept commands, user logic must write the command parameters into `command_parameters`, and then write the command code into `requested_command`.

The sequencer detects the command code and replaces `command_status` with `TCLDBG_TX_STATUS_CMD_EXE`, to indicate that it is processing the command. When the sequencer has finished running the command, it sets `command_status` to `TCLDBG_TX_STATUS_RESPONSE_READY` to indicate that the result of the command is available to be read. (If the sequencer rejects the requested command as illegal, it sets `command_status` to `TCLDBG_TX_STATUS_ILLEGAL_CMD`.)

User logic acknowledges completion of the command by writing `TCLDBG_CMD_RESPONSE_ACK` to `requested_command`. The sequencer responds by setting `command_status` back to `STATUS_CMD_READY`. (If an illegal command is received, it must be cleared using `CMD_RESPONSE_ACK`.)



Figure 130. Debugging Algorithm Flowchart



13.9.3 On-Die Termination Calibration

The **Calibrate Termination** feature lets you determine the optimal **On-Die Termination** and **Output Drive Strength** settings for your memory interface.

The **Calibrate Termination** function runs calibration with all available termination settings and selects the optimal settings based on the calibration margins.

The **Calibrate Termination** feature is available for DDR3, DDR4, and RLDRAM 3 protocols.

13.9.4 Eye Diagram

The **Generate Eye Diagram** feature allows you to create read and write eye diagrams for each pin in your memory interface.

The **Generate Eye Diagram** feature uses calibration data patterns to determine margins at each V_{ref} setting on both the FPGA pins and the memory device pins. A full calibration is done for each Vref setting. Other settings, such as DQ delay chains, will change for each calibration. At the end of a **Generate Eye Diagram** command, a default calibration is run to restore original behavior

The **Generate Eye Diagram** feature is available for DDR4 and QDR-IV protocols.

13.10 Driver Margining for Intel Stratix 10 EMIF IP

The Driver Margining feature lets you measure margins on your memory interface using a driver with arbitrary traffic patterns.

Margins measured with this feature may differ from margins measured during calibration, because of different traffic patterns. Driver margining is not available if ECC is enabled.

To use driver margining, ensure that the following signals on the driver are connected to In-System Sources/Probes:

- **Reset_n**: An active low reset signal
- **Pass**: A signal which indicates that the driver test has completed successfully. No further memory transactions must be sent after this signal is asserted.
- **Fail**: A signal which indicates that the driver test has failed. No further memory transactions must be sent after this signal is asserted.
- **PNF (Pass Not Fail)**: An array of signals that indicate the pass/fail status of individual bits of a data burst. The PNF should be arranged such that each bit index corresponds to $(\text{Bit of burst} * \text{DQ width}) + (\text{DQ pin})$. A 1 indicates pass, 0 indicates fail. If the PNF width exceeds the capacity of one In-System Probe, specify them in `PNF[1]` and `PNF[2]`; otherwise, leave them blank.

If you are using the example design with a single EMIF, the In-System Sources/Probes can be enabled by adding the following line to your .qsf file:

```
set_global_assignment -name VERILOG_MACRO
"ALTERA_EMIF_ENABLE_ISSP=1"
```

13.10.1 Determining Margin

The Driver Margining feature lets you measure margins on your EMIF IP interface using a driver with arbitrary traffic patterns.

The Driver Margining feature is available only for DDR3 and DDR4 interfaces, when ECC is not enabled.

1. Establish a connection to the desired interface and ensure that it has calibrated successfully.
2. Select **Driver Margining** from the **Commands** folder under the target interface connection.
3. Select the appropriate **In-System Sources/Probes** using the drop-down menus.
4. If required, set additional options in the **Advanced Options** section:



- Specify **Traffic Generator 2.0** to allow margining on a per-rank basis. Otherwise, margining is performed on all ranks together.
 - **Step size** specifies the granularity of the driver margining process. Larger step sizes allow faster margining but reduced accuracy. It is recommended to omit this setting.
 - **Adjust delays after margining** causes delay settings to be adjusted to the center of the window based on driver margining results.
 - The **Margin Read**, **Write**, **Write DM**, and **DBI** checkboxes allow you to control which settings are tested during driver margining. You can uncheck boxes to allow driver margining to complete more quickly.
5. Click **OK** to run the tests.
The toolkit measures margins for DQ read/write and DM. The process may take several minutes, depending on the margin size and the duration of the driver tests. The test results are available in the *Margin Report*.

13.11 Traffic Generator 2.0

The Traffic Generator 2.0 lets you emulate traffic to the external memory, and helps you test, debug, and understand the performance of your external memory interface on hardware in a standalone fashion, without having to incorporate your entire design.

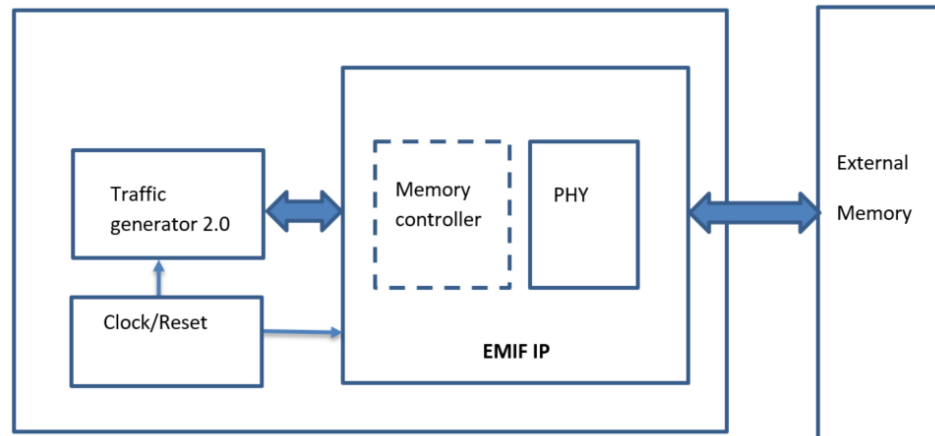
The Traffic Generator 2.0 lets you customize data patterns being written to the memory, address locations accessed in the memory, and the order of write and read transactions. You can use the traffic generator code with any FPGA architecture and memory protocol.

13.11.1 Configuring the Traffic Generator 2.0

The traffic generator replaces user logic to generate traffic to the external memory. You must incorporate the traffic generator design into the EMIF IP design during IP generation.

When you generate the example design in the parameter editor, the traffic generator module and EMIF IP are generated together. If you have an example design with the Traffic Generator 2.0 enabled, you can configure the traffic pattern using the EMIF Debug Toolkit.

Figure 131. Traffic Generator 2.0 Generated with EMIF IP in Example Design Mode



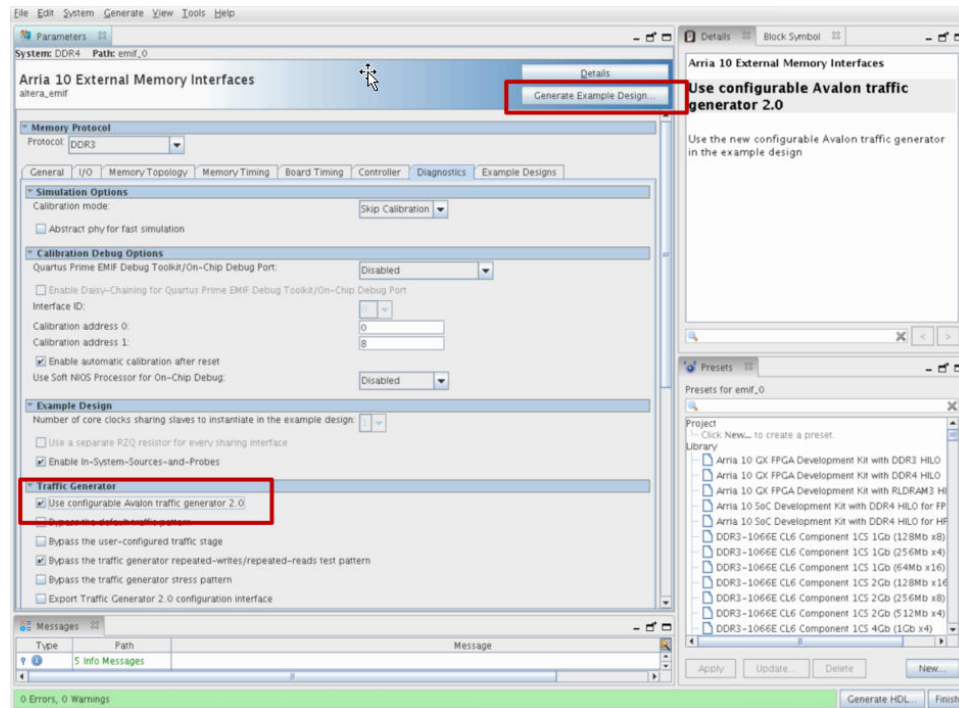
Generating the External Memory Interface

1. Select the FPGA and Memory parameters.
2. On the Diagnostics tab, configure the following parameters:
 - a. Select **Use Configurable Avalon Traffic Generator 2.0**.
 - b. Configure the desired traffic pattern, by specifying traffic patterns to be bypassed. The traffic pattern not bypassed is issued to the memory immediately after completion of calibration. You can choose to bypass any of the following traffic patterns:
 - **Bypass the default traffic pattern** Specifies not to use the default traffic patterns from the traffic generator. The default patterns include single read/write, byte-enabled read/write, and block read/write.
 - **Bypass the user-configured traffic stage.** Specifies to skip the stage that uses the user-defined test bench file to configure the traffic generator in simulation.
 - **Bypass the traffic generator repeated-writes/repeated-reads test pattern.** Bypasses the traffic generator's repeat test stage, which causes every write and read to be repeated several times.
 - **Bypass the traffic generator stress pattern.** Bypasses a test stage intended to stress-test signal integrity and memory interface calibration.
 - **Export Traffic Generator 2.0 configuration interface.** Instantiates a port for traffic generator configuration. Use this port if the traffic generator is to be configured by user logic.
3. Click **Generate Example Design** to generate the EMIF IP, including the Traffic Generator 2.0 design, with the traffic pattern that you have configured.

Note: If you click the **Generate HDL** option instead, the Traffic Generator 2.0 design is not included in the generated IP.



Figure 132. Enabling the Traffic Generator 2.0 in the Parameter Editor



13.11.2 Configurable Traffic Generator 2.0 Configuration Options

The EMIF Configurable Traffic Generator 2.0 offers a range of configuration options for fast debugging of your external memory interface. You configure the traffic generator by modifying address-mapped registers through the simulation test bench file or by creating a configuration test stage.

Configuration Syntax

The test bench includes an example test procedure.

The syntax for writing to a configuration register is as follows:

```
tg_send_cfg_write_<index>(<Register Name>, <Value to be written>);
```

The *index* value represents the index of the memory interface (which is usually 0, but can be 0/1 in ping-pong PHY mode).

The *register name* values are listed in the tables of configuration options, and can also be found in the `altera_emif_avl_tg_defs.sv` file, in the same directory as the test bench.

The final configuration command must be a write of any value to `TG_START`, which starts the traffic generator for the specified interface.

Configuration Options

Configuration options are divided into read/write, address, data, and data mask generation categories.

**Table 444. Configuration Options for Read/Write Generation**

Toolkit GUI Tab	Toolkit GUI Parameter	Register Name	Description
--	--	TG_START	Tells the system to perform a write to this register to initiate traffic generation test.
Loops	Loops	TG_LOOP_COUNT	Specifies the number of read/write loops to perform before completion of the test stage. A loop is a single iteration of writes and reads. Upon completion, the base address is either incremented (SEQ or RAND_SEQ) or replaced by a newly generated random address (RAND or RAND_SEQ).
	Writes per block	TG_WRITE_COUNT	Specifies the number of writes to be performed in each loop of the test.
	Reads per block	TG_READ_COUNT	Specifies the number of reads to be performed in each loop of the test. This register must have the same value as TG_WRITE_COUNT.
	--	TG_WRITE_REPEAT_COUNT	Specifies the number of times each write transaction is repeated.
	--	TG_READ_REPEAT_COUNT	Specifies the number of times each read transaction is repeated.
	Reads per block	TG_BURST_LENGTH	Configures the length of each write/read burst to memory.
--	--	TG_CLEAR_FIRST_FAIL	Clears the record of first failure occurrence. New failure information such as expected data, read data, and fail address, is written to the corresponding registers following the next failure.
--	--	TG_TEST_BYTEEN	Toggles the byte-enable (data mask enable) register within the traffic generator which allows the current test to use byte-enable signals.
--	--	TG_DATA_MODE	Specifies the source of data used for data signal generation during the test. Set to 0 to use pseudo-random data. Set to 1 to use user-specified values stored in the static data generators.
--	--	TG_BYTEEN_MODE	Specifies the source of data used for byte-enable signal generation during the test. Set to 0 to use pseudo-random data. Set to 1 to use user-specified values stored in the static data generators.

**Table 445. Configuration Options for Address Generation**

Toolkit GUI Tab	Toolkit GUI Parameter	Register Name	Description
Address	Start Address	TG_SEQ_START_ADDR_WR_L	Specifies the sequential start address (lower 32 bits).
	Start address	TG_SEQ_START_ADDR_WR_H	Specifies the sequential start address (upper 32 bits)
	Address mode	TG_ADDR_MODE_WR	Specifies how write addresses are generated by writing the value in parentheses to the register address. Values are: randomize the address for every write (0), increment sequentially from a specified address (1), increment sequentially from a random address (2), or perform one hot addressing (3).
	Number of sequential address	TG_RAND_SEQ_ADDRS_WR	Specifies the number of times to increment sequentially on the random base address before generating a new random write address.
	Return to start address	TG_RETURN_TO_START_ADDR	Return to start address in deterministic sequential address mode (if 1 is written to TG_ADDR_MODE_WR) after every loop of transactions.
	Rank Mask Mode	TG_RANK_MASK_EN	Specifies the rank masking mode by writing the value in parentheses to the register address. Values are: disable rank masking (0), maintain a static rank mask (1), cycle through rank masks incrementally (2).
	Bank Address Mask Mode	TG_BANK_MASK_EN	Specifies the bank masking mode by writing the value in parentheses to the register address. Values are: disable bank masking (0), maintain a static bank mask (1), cycle through bank masks incrementally (2), cycle through only three consecutive bank masks (3).
	Row Mask Mode	TG_ROW_MASK_EN	Specifies the mode for row masking by writing the value in parentheses to the register address. Values are: disable row masking (0), maintain a static row mask (1), cycle through row masks incrementally (2).
	Bank Group Mask Mode	TG_BG_MASK_EN	Specifies the mode for bank group masking by writing the value in parentheses to the register address. Values are: disable bank group masking (0), maintain a static bank group mask (1), cycle through bank group masks incrementally (2).
	Rank Mask Value	TG_RANK_MASK	Specifies the initial rank to be masked into the generated traffic generator address.
continued...			



Toolkit GUI Tab	Toolkit GUI Parameter	Register Name	Description
	Bank Address Mask Value	TG_BANK_MASK	Specifies the initial bank to be masked into the generated traffic generator address.
	Row Mask Value	TG_ROW_MASK	Specifies the initial row to be masked into the generated traffic generator address.
	Bank Group Mask Value	TG_BG_MASK	Specifies the initial bank group to be masked into the generated traffic generator address.
	Sequential Address Increment	TG_SEQ_ADDR_INCR	Specifies the increment to use when sequentially incrementing the address. This value is used by both deterministic sequential addressing and random sequential addressing. (Refer to TG_ADDR_MODE_WR)
	Start Address	TG_SEQ_START_ADDR_RD_L	Specifies the sequential start read address (lower 32 bits).
		TG_SEQ_START_ADDR_RD_H	Specifies the sequential start read address (upper 32 bits).
	Address Mode	TG_ADDR_MODE_RD	Similar to TG_ADDR_MODE_WR but for reads.
	Number of sequential address	TG_RAND_SEQ_ADDRS_RD	Specifies the number of times to increment the random sequential read address.

Table 446. Configuration Options for Data Generation

Toolkit GUI Tab	Toolkit GUI Parameter	Register Name	Description
Data	Seed/Fixed Pattern	TG_DATA_SEED	Specifies an initial value to the data generator corresponding to the index value.
	PRBS and Fixed Pattern radio buttons	TG_DATA_MODE	Specifies whether to treat the initial value of the data generator of corresponding index as a seed for generating pseudo-random data (value of 0) or to keep the initial value static (value of 1).

Table 447. Configuration Options for Data Mask Generation

Toolkit GUI Tab	Toolkit GUI Parameter	Register Name	Description
Data	Seed/Fixed Pattern	TG_BYTEEN_SEED	Specifies an initial value to the byte-enable generator corresponding to the index value.
	PRBS and Fixed Pattern radio buttons	TG_BYTEEN_MODE	Specifies whether to treat the initial value of the byte-enable generator of corresponding index as a seed and generate pseudo-random data (value of 0) or to keep the initial value static (value of 1).

13.11.2.1 Test Information

In the test bench file, register reads are encoded in a similar syntax to register writes.



The following example illustrates the syntax of a register read:

```
integer <Variable Name>;
tg_send_cfg_read_<index>(<Register Name>, <Variable Name>);
```

In hardware, you can probe the registers storing the test information (such as pnf per bit persist, first fail read address, first fail read data, and first fail expected data).

Table 448. Test Information Read-Accessible Through Register Addresses

Register Name	Description
TG_PASS	Returns a high value if the traffic generator passes at the end of all the test stages.
TG_FAIL	Returns a high value if the traffic generator fails at the end of all the test stages.
TG_FAIL_COUNT_L	Reports the failure count (lower 32 bits).
TG_FAIL_COUNT_H	Reports the failure count (upper 32 bits).
TG_FIRST_FAIL_ADDR_L	Reports the address of the first failure (lower 32 bits).
TG_FIRST_FAIL_ADDR_H	Reports the address of the first failure (upper 32 bits).
TG_FIRST_FAIL_IS_READ	First failure is Read Failure.
TG_FIRST_FAIL_IS_WRITE	First failure is Write Failure.
TG_VERSION	Reports the traffic generator version number.
TG_NUM_DATA_GEN	Reports the number of data generators.
TG_NUM_BYTEEN_GEN	Reports the number of byte-enable generators.
TG_RANK_ADDR_WIDTH	Reports the rank address width.
TG_BANK_ADDR_WIDTH	Reports the bank address width.
TG_ROW_ADDR_WIDTH	Reports the row address width.
TG_BANK_GROUP_WIDTH	Reports the bank group width.
TG_RDATA_WIDTH	Reports the width of all data and PNF signals within the traffic generator.
TG_DATA_PATTERN_LENGTH	Reports the length of the static pattern to be loaded into static per-pin data generators.
TG_BYTEEN_PATTERN_LENGTH	Reports the length of the static pattern to be loaded into static per-pin byte-enable generators.
TG_MIN_ADDR_INCR	Reports the minimum address increment permitted for sequential and random-sequential address generation.
TG_ERROR_REPORT	Reports error bits. Refer to <i>Error Report Register Bits</i> for details.
TG_PNF	Read the persistent PNF per bit as an array of 32-bit entries.
TG_FAIL_EXPECTED_DATA	Reports the first failure expected data. This is read as an array of 32-bit entries.
TG_FAIL_READ_DATA	Reports the first failure read data. This is read as an array of 32-bit entries.

The configuration error report register contains information on common misconfigurations of the traffic generator. The bit corresponding to a given configuration error is set high when that configuration error is detected. Ensure that all bits in this register are low, to avoid test failures or unexpected behavior due to improper configuration.

Table 449. Error Report Register Bits

Bit Index (LSB = 0x0)	Bit Name	Description of Error
0	ERR_MORE_READS_THAN_WRITES	You have requested more read transactions per loop than write transactions per loop.
1	ERR_BURSTLENGTH_GT_SEQ_ADDR_INCREMENT	The configured burst length is greater than the configured sequential address increment when address generation is in sequential or random sequential mode.
2	ERR_ADDR_DIVISIBLE_BY_GT_SEQ_ADDRESS_INCREMENT	The configured sequential address increment is less than the minimum required address increment when address generation is in sequential or random sequential mode.
3	ERR_SEQ_ADDR_INCREMENT_NOT_DIVISIBLE	The configured sequential address increment is not a multiple of the minimum required address increment when address generation is in sequential or random sequential mode.
4	ERR_READ_AND_WRITE_START_ADDRESSES_DIFFER	The configured start addresses for reads and writes are different when address generation is in sequential mode.
5	ERR_ADDR_MODES_DIFFERENT	The configured address modes for reads and writes are different.
6	ERR_NUMBER_OF_RANDOM_SEQ_ADDRESSES_DIFFERENT	The configured number of random sequential addresses for reads and writes are different when address generation is in random sequential mode.
7	ERR_REPEATS_SET_TO_ZERO	The number of read and/or write repeats is set to 0.
8	ERR_BOTH_BURST_AND_REPEAT_MODE_ACTIVE	The burst length is set greater than 1 and read/write requests are set greater than 1.
9-31	Reserved	--

13.11.3 Running the Traffic Generator 2.0

You can use the EMIF Debug Toolkit to configure the traffic generator infrastructure to send custom traffic patterns to the memory.

1. Launch the EMIF Debug Toolkit by selecting **Tools ► System Debugging Tools ► External Memory Interface Toolkit**.
2. After you launch the toolkit, you must establish the following connections, before running the custom traffic generator:



- Initialize Connections
 - Link Project to Device
 - Connections
 - Create Memory Interface Connection
 - Create Traffic Generator Connection
3. Launch the Traffic Generator by selecting **Traffic Generator > Settings > Run Custom Traffic Pattern**.

13.11.4 Understanding the Custom Traffic Generator User Interface

The Custom Traffic Generator interface lets you configure data patterns, the bytes to be enabled, the addressing mode, and the order in which traffic is organized.

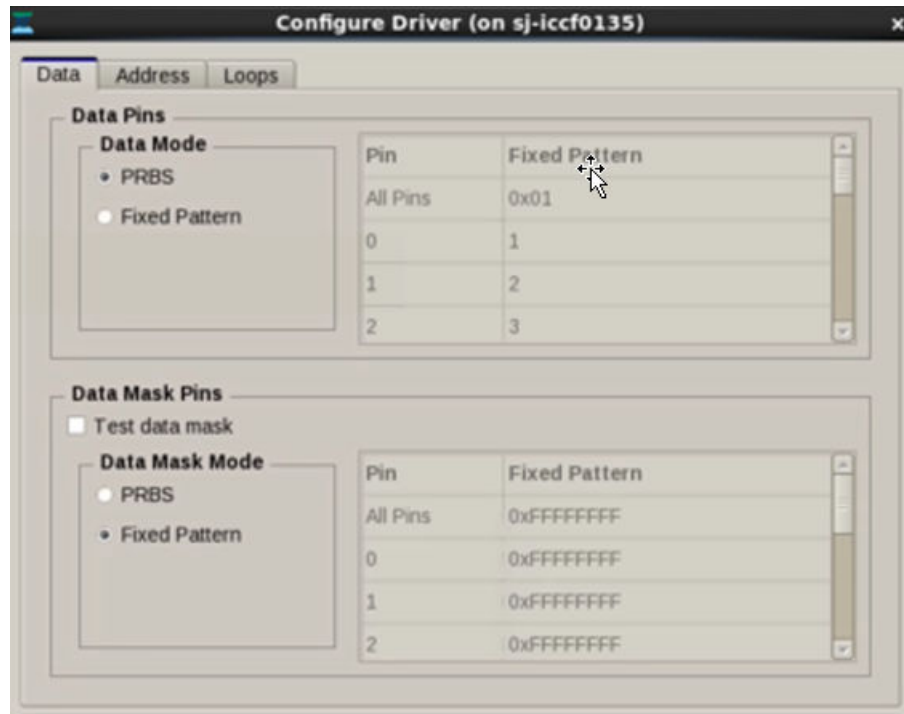
The interface has three tabs:

- **Data** tab
- **Address** tab
- **Loops** tab

Data Tab

The **Data** tab is divided into **Data Pins** and **Data Mask Pins** sections.

Figure 133. Data Tab



The **Data Pins** section helps with customizing the patterns selected for the data pins. You can choose between two options for **Data Mode**:

- **PRBS** The default write data to all the data pins.
- **Fixed Pattern** Lets you specify a pattern to be written to the memory.

Select the **All Pins** option when you want to write the same data pattern to all the data pins. If data must be individually assigned to the data pins, you must enter the data value for each individual pin. The width of the data entered is based on the AVL_TO_DQ_WIDTH_RATIO, which is based on the ratio of the memory clock to the user clock.

All data bytes are enabled by default; the **Data Mask Pins** section lets you disable any of the bytes if you want to. To disable data bytes individually, select **Test data mask**.

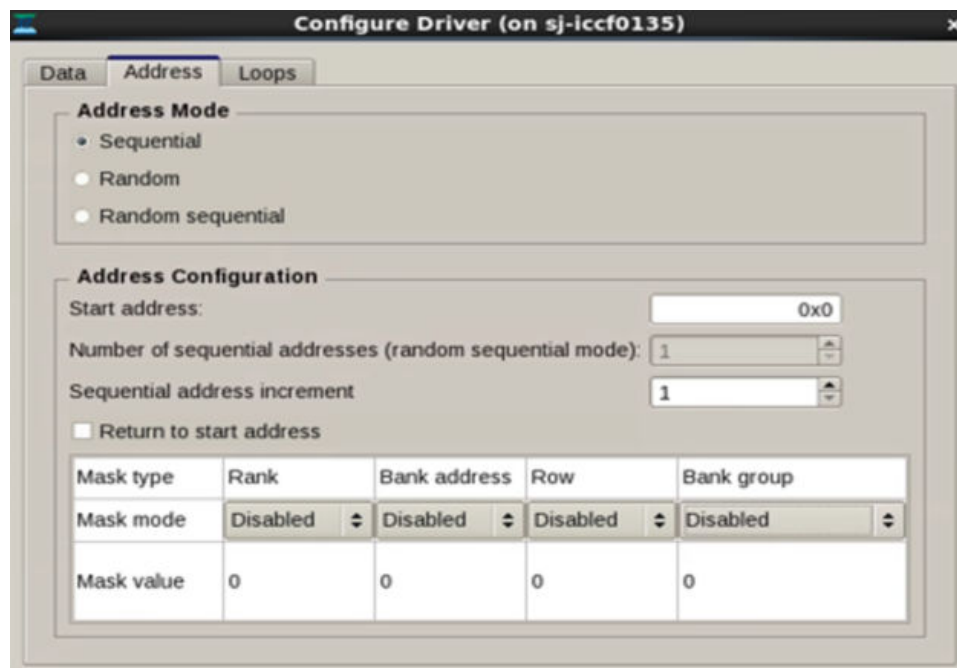
You can choose between two options for **Data Mode**:

- **PRBS** Specifies the PRBS pattern to enable or disable data bytes. A 1 denotes a data byte enabled, while a 0 denotes a data byte being masked or disabled.
- **Fixed Pattern** Lets you enable or disable individual bytes. You can apply byte enables to all pins or to individual bytes. A 1 denotes a data byte enabled, while a 0 denotes a data byte being masked or disabled.

Address Tab

The **Address** tab lets you configure sequential, random, or random sequential (where the initial start address is random, but sequential thereafter) addressing schemes. The **Address** tab is divided into **Address Mode** and **Address Configuration** sections.

Figure 134. Address Tab



Mask type	Rank	Bank address	Row	Bank group
Mask mode	Disabled	Disabled	Disabled	Disabled
Mask value	0	0	0	0



The **Address Mode** section lets you specify the pattern of addresses generated to access the memory. You can choose between three address modes :

- **Sequential** Each address is incremented by the **Sequential address increment** value that you specify. You also specify the **Start Address** from which the increments begin.
- **Random** Each address is generated randomly. (You set the number of random addresses on the **Loops** tab.)
- **Random Sequential** Each address is generated randomly, and then incremented sequentially. You specify the number of sequential increments in the **Number of sequential addresses** field.

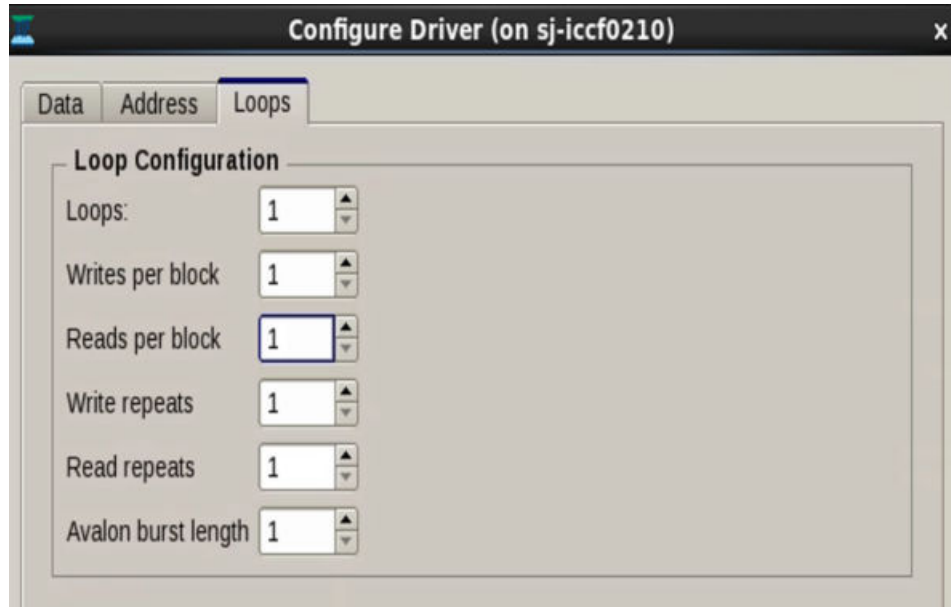
The **Address Configuration** section contains the settings with which you configure the address mode that you chose in the **Address Mode** section. The following settings are available:

- **Start address** Specifies the starting address for **Sequential Address Mode**. The maximum address value that can be reached is 1FF_FFFF. (The Traffic Generator 2.0 will accept higher values, but wraps back to 0 after the maximum value has been reached.) The **Start address** setting applies only to **Sequential Address Mode**.
- **Number of sequential addresses** Specifies the number of sequential addresses generated after the first random address generated. This setting applies only in **Random sequential** mode.
- **Sequential address increment** Specifies the size of increment between each address in the **Sequential** address mode and **Random sequential** address mode.
- **Return to start address** Specifies that the address value generated return back to the value entered in the **Start Address** field, after a block of transactions to the memory has completed. This setting applies only to **Sequential** address mode.
- **Address masking** Masking provides additional options for exploring certain specific address spaces in memory:
 - **Disabled** does not enable masking, and increments address based on the selected **Address Mode**.
 - **Fixed cycling** allows you to restrict the addressing to a specific row or a specific bank, which you can specify in the corresponding **Mask Value** field.

Loops Tab

The **Loops** tab lets you order the transactions to the memory as desired. A unit size of transactions to the memory is defined as a *block*; a *block* includes a set of write transaction(s) immediately followed by a set of read transaction(s).

Figure 135. Loops Tab



The screenshot shows a dialog box titled "Configure Driver (on sj-iccf0210)" with a close button (X) in the top right corner. The dialog has three tabs: "Data", "Address", and "Loops". The "Loops" tab is selected and active. Inside the "Loops" tab, there is a section titled "Loop Configuration" which contains six settings, each with a text input field and a spin button (up and down arrows):

- Loops: 1
- Writes per block: 1
- Reads per block: 1
- Write repeats: 1
- Read repeats: 1
- Avalon burst length: 1



The **Loops** tab provides the following configuration options:

- **Loops** Specifies the number of blocks of transactions to be sent to the memory. This option helps to extend the range of addresses that the controller can access. The address range is incremented as each loop is executed, unless you specify **Return to Start Address**, which causes each loop to begin from the same start address. The range of supported values for the **Loops** option is from 1 to 4095.
- **Writes per block** Specifies the size of the block (that is, the number of consecutive write operations that can be issued in a single block). The range of values for this option is as follows:
 - When address masking is disabled, the number of writes per block supported is 1 to 4094.
 - When address masking is enabled, the maximum number of writes issued inside a block is 255.
- **Reads per block** Specifies the number of consecutive read operations that can be issued in a single block, immediately following the consecutive writes issued. The number of reads per block should be identical to the number of writes per block, because data mismatches can occur when the two values are not identical. The range of values for this option is as follows:
 - When address masking is disabled, the number of reads per block supported is 1 to 4094.
 - When address masking is enabled, the maximum number of reads issued inside a block is 255.
- **Write repeats** Specifies the number of times each write command is issued in repetition to the same address. A maximum number of 255 repeat write transactions can be issued. The repeat writes are issued immediately after the first write command has been issued.
- **Read repeats** Specifies the number of times each read command is issued in repetition to the same address. A maximum number of 255 repeat read transactions can be issued. The repeat reads are issued immediately after the first read command has been issued.
- **Avalon burst length** Specifies the length of each Avalon burst. The value of this field should be less than the **Sequential address increment** specified on the **Address** tab. The number of write and read repeats default to 1 if the Avalon burst length is greater than 1.

13.11.5 Applying the Traffic Generator 2.0

You can apply the Traffic Generator 2.0 to run stress tests, debug your hardware platform for signal integrity problems, and to emulate actual memory transactions. This topic presents some common applications where you can benefit by using the Traffic Generator 2.0.

Testing Signal Integrity with PRBS Data Pattern

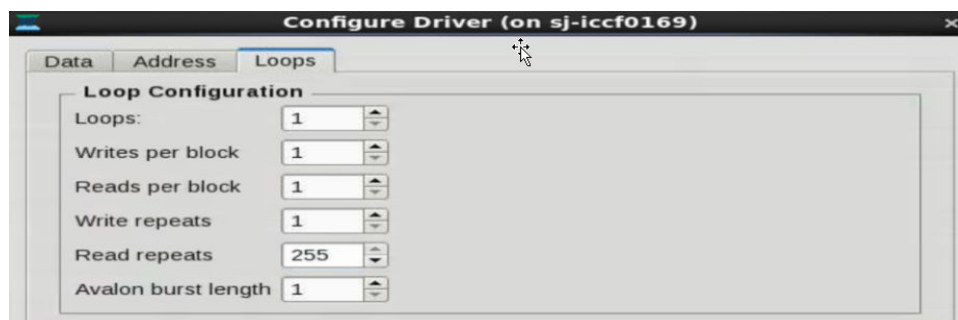
You can apply PRBS data to the data pins to help emulate an actual traffic pattern to the memory interface. The traffic generator uses a PRBS7 data pattern as the default traffic pattern on the data pins, and can support PRBS-15 and PRBS-31.

Debugging and Monitoring an Address for Reliable Data Capture

You can send a single write followed by multiple reads to a specific address to help debug and monitor a specific address for reliable data capture. You can do this with the following settings:

- Writes per block: 1
- Reads per block: 1
- Write repeats: 1
- Read repeats: 1 to 255

Figure 136. Configuring the Loops Tab for a Single Write Followed by Multiple Reads



If you specify a **Loops** value greater than 1, every block of write and multiple read transactions will follow the same pattern. If there is a specific address to which this transaction must be issued, you should specify that address in the **Start address** field on the **Address** tab, with the **Sequential** address mode selected.

Accessing Large Sections of Memory

The maximum number of unique addresses that can be written to in one block is 4094. Using the maximum **Loops** value of 4095, the address range that can be supported in one test is equal to the number of loops multiplied by the number of writes per block. Further address expansion can be achieved by changing the **Start address** value appropriately and reissuing the tests.

To continue addressing sections of the memory beyond the address range that can be specified in one set of toolkit configurations, you can incrementally access the next set of addresses in the memory by changing the **Start address** value.

For example, in a memory where row address width is 15, bank address width is 3 and column address width is 10, the total number of address locations to be accessed is: 2 (row address width) x (bank address width x 2 (column address width)). The maximum number of address locations that can be accessed is limited by the width of the internal address bus, which is 25 bits wide.

For the example described above, you must set the following values on the **Address** tab:

- Select the **Sequential** address mode.
- Set the **Start address** to 0x00.
- Ensure that you do not select **Return to start address**.
- Ensure that you disable address masking for rank, row, bank, and bank group.



Figure 137. Address Configuration to Access the First Set of Addresses

Mask type	Rank	Bank address	Row	Bank group
Mask mode	Disabled	Disabled	Disabled	Disabled
Mask value	0	0	0	0

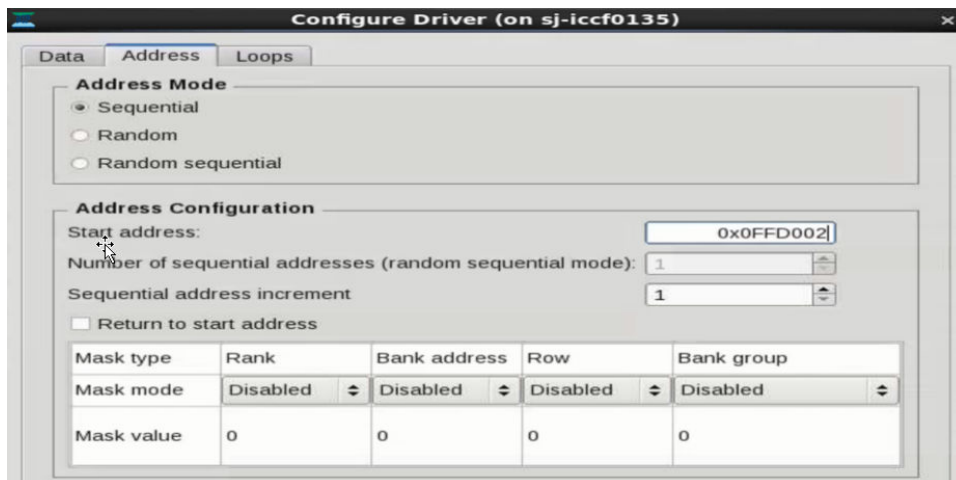
You must also set the following values on the **Loops** tab:

- Set **Loops** to the maximum value of 4095.
- Set **Writes per block** to the maximum value of 4094.
- Set **Reads per block** to the maximum value of 4094.
- Set **Write repeats** to 1.
- Set **Read repeats** to 1.

Figure 138. Loop Configuration to Access the First Set of Addresses

Each iteration can access a maximum of 4095 x 4094 locations (16,764,930 address locations i.e. Address ranging from 000_0000'h to FF_D001'h). To access the next 4095 x 4094 locations, the same settings as above must be repeated, except for the **Start address** value, which must be set to a hex value of 16,764,931 i.e. FF_D002. The same process can be repeated to further access memory locations inside the memory. The maximum value supported is 25'h 1FF_FFFF which is the equivalent of 33,554,432 locations inside the memory.

Figure 139. Address Configuration to Access the Second Set of Addresses



Mask type	Rank	Bank address	Row	Bank group
Mask mode	Disabled	Disabled	Disabled	Disabled
Mask value	0	0	0	0

13.11.6 Testing the EMIF Interface Using the Traffic Generator 2.0

The EMIF Configurable Traffic Generator 2.0 can assist in debugging and stress-testing your external memory interface.

Key Features

The traffic generator 2.0 has the following key features:

- Is a standalone, soft-logic device that resides in the FPGA core.
- Is independent of the FPGA architecture and the external memory protocol in use.
- Offers configuration options for the generation of reads and writes, addressing, data, and data mask.

13.11.6.1 Test Information

In the test bench file, register reads are encoded in a similar syntax to register writes.

The following example illustrates the syntax of a register read:

```
integer <Variable Name>;
tg_send_cfg_read_<index>(<Register Name>, <Variable Name>;
```

In hardware, you can probe the registers storing the test information (such as pnf per bit persist, first fail read address, first fail read data, and first fail expected data).

Table 450. Test Information Read-Accessible Through Register Addresses

Register Name	Description
TG_PASS	Returns a high value if the traffic generator passes at the end of all the test stages.
TG_FAIL	Returns a high value if the traffic generator fails at the end of all the test stages.
continued...	



Register Name	Description
TG_FAIL_COUNT_L	Reports the failure count (lower 32 bits).
TG_FAIL_COUNT_H	Reports the failure count (upper 32 bits).
TG_FIRST_FAIL_ADDR_L	Reports the address of the first failure (lower 32 bits).
TG_FIRST_FAIL_ADDR_H	Reports the address of the first failure (upper 32 bits).
TG_FIRST_FAIL_IS_READ	First failure is Read Failure.
TG_FIRST_FAIL_IS_WRITE	First failure is Write Failure.
TG_VERSION	Reports the traffic generator version number.
TG_NUM_DATA_GEN	Reports the number of data generators.
TG_NUM_BYTEEN_GEN	Reports the number of byte-enable generators.
TG_RANK_ADDR_WIDTH	Reports the rank address width.
TG_BANK_ADDR_WIDTH	Reports the bank address width.
TG_ROW_ADDR_WIDTH	Reports the row address width.
TG_BANK_GROUP_WIDTH	Reports the bank group width.
TG_RDATA_WIDTH	Reports the width of all data and PNF signals within the traffic generator.
TG_DATA_PATTERN_LENGTH	Reports the length of the static pattern to be loaded into static per-pin data generators.
TG_BYTEEN_PATTERN_LENGTH	Reports the length of the static pattern to be loaded into static per-pin byte-enable generators.
TG_MIN_ADDR_INCR	Reports the minimum address increment permitted for sequential and random-sequential address generation.
TG_ERROR_REPORT	Reports error bits. Refer to <i>Error Report Register Bits</i> for details.
TG_PNF	Read the persistent PNF per bit as an array of 32-bit entries.
TG_FAIL_EXPECTED_DATA	Reports the first failure expected data. This is read as an array of 32-bit entries.
TG_FAIL_READ_DATA	Reports the first failure read data. This is read as an array of 32-bit entries.

The configuration error report register contains information on common misconfigurations of the traffic generator. The bit corresponding to a given configuration error is set high when that configuration error is detected. Ensure that all bits in this register are low, to avoid test failures or unexpected behavior due to improper configuration.

Table 451. Error Report Register Bits

Bit Index (LSB = 0x0)	Bit Name	Description of Error
0	ERR_MORE_READS_THAN_WRITES	You have requested more read transactions per loop than write transactions per loop.
1	ERR_BURSTLENGTH_GT_SEQ_ADDR_INCREMENT	The configured burst length is greater than the configured sequential address increment when address generation is in sequential or random sequential mode.
2	ERR_ADDR_DIVISIBLE_BY_GT_SEQ_ADDRESS_INCREMENT	The configured sequential address increment is less than the minimum required address increment when address generation is in sequential or random sequential mode.
3	ERR_SEQ_ADDR_INCREMENT_NOT_DIVISIBLE	The configured sequential address increment is not a multiple of the minimum required address increment when address generation is in sequential or random sequential mode.
4	ERR_READ_AND_WRITE_START_ADDRESSES_DIFFER	The configured start addresses for reads and writes are different when address generation is in sequential mode.
5	ERR_ADDR_MODES_DIFFERENT	The configured address modes for reads and writes are different.
6	ERR_NUMBER_OF_RANDOM_SEQ_ADDRESSES_DIFFERENT	The configured number of random sequential addresses for reads and writes are different when address generation is in random sequential mode.
7	ERR_REPEATS_SET_TO_ZERO	The number of read and/or write repeats is set to 0.
8	ERR_BOTH_BURST_AND_REPEAT_MODE_ACTIVE	The burst length is set greater than 1 and read/write requests are set greater than 1.
9-31	Reserved	--

13.11.6.2 Performing Your Own Tests Using Traffic Generator 2.0

If you want, you can create your own configuration test stage for the EMIF Configurable Traffic Generator.

The general flow of a configuration test stage, including the default test stages, is as follows:

- Configure the number of loops to be completed by the traffic generator (TG_LOOP_COUNT).
- Configure the number of writes and reads to be complete per loop (TG_WRITE_COUNT and TG_READ_COUNT respectively).
- Choose the burst length of each write and read (TG_BURST_LENGTH).
- Select starting write address by writing to the lower and upper bits of the address register (TG_SEQ_START_ADDR_WR_L and TG_SEQ_START_ADDR_WR_H, respectively).



- Select write address generation mode (TG_ADDR_MODE_WR).
- Select starting read address by writing to the lower and upper bits of the address register (TG_SEQ_START_ADDR_RD_L and TG_SEQ_START_ADDR_RD_H, respectively).
- Select read address generation mode (TG_ADDR_MODE_RD).
- If applicable, select sequential address increment (TG_SEQ_ADDR_INCR).
- Write initial values/seeds to the data and byte-enable generators (TG_DATA_SEED and TG_BYTEEN_SEED).
- Select generation mode of the data and byte-enable generators (TG_DATA_MODE and TG_BYTEEN_MODE).
- Initiate test (TG_START).

Simulation

For a comprehensive example of how to write your own configuration test for simulation, refer to the test bench file, located at

```
<example_design_directory>/sim/ed_sim/  
altera_emif_tg_avl_2_<>/sim/altera_emif_avl_tg_2_tb.sv
```

To iterate over the data generators or byte-enable generators, you must read the number of data generators and number of byte-enable generators. These values are mapped to read-accessible registers TG_NUM_DATA_GEN and TG_NUM_BYTEEN_GEN, respectively. The following example illustrates how one would configure the data generators to continuously output the pattern *0x5a*, using the simulation test bench:

```
integer num_data_generators;  
...  
    tg_send_cfg_read_0(TG_NUM_DATA_GEN, num_data_generators);  
    tg_send_cfg_write_0(TG_DATA_MODE, 32'h1);  
    for (i = 0; i < num_data_generators; i = i + 1) begin  
        tg_send_cfg_write_0(TG_DATA_SEED + i, 32'h5A);  
    end  
end
```

Hardware

Configuration test stages in hardware must be inserted into the RTL, and will resemble the single read/write, byte-enable, and block read/write stages in the default test pattern. In most cases, you can modify one of the existing stages to create the desired custom test stage. The stages are linear, finite, state machines that write predetermined values to the configuration address-mapped registers. As always, the last state in configuration is a write to address *0x0* or TG_START. The state machine then waits for the traffic generator to return a signal signifying its completion of the test stage.

Refer to the aforementioned default test stages as examples of hardware test stages. The default test stages are contained within the following files:

- <example_design_directory>/qii/ed_synth/
altera_emif_tg_avl_2_<>/synth/altera_emif_avl_tg_2_rw_stage.sv
- <example_design_directory>/qii/ed_synth/
altera_emif_tg_avl_2_<>/synth/
altera_emif_avl_tg_2_byteenable_test_stage.sv

You can also configure the configurable traffic generator in real time using the EMIF Debug Toolkit. The configuration settings available in the Toolkit interface are detailed in the *Configurable Traffic Generator 2.0 Configuration Options* topic.

13.11.6.3 Signal Splitter Component

The signal splitter (`altera_emif_sig_splitter`) is an internal IP component which receives a single signal as its input and passes that signal directly to n outputs, where n is an interger value equal to or greater than 1. The signal splitter is useful because Platform Designer does not directly allow one-to-many connections for conduit interfaces.

The signal splitter contains no logic or memory elements. When you configure the signal splitter to have exactly one output port, it is functionally identical to a single wire, and can be replaced by one with no loss of performance.

The `rzq_splitter` is an instantiation of the signal splitter component specifically for the RZQ signal. The signal splitter facilitates the sharing of one RZQ signal among multiple memory interfaces in an EMIF example design.


13.11.7 The Traffic Generator 2.0 Report

The traffic generator report provides information about the configuration of the Traffic Generator 2.0 and the result of the most recent run of traffic.

Understanding the Traffic Generator 2.0 Report

The traffic generator report contains the following information:

- A **Pass flag** value of 1 indicates that the run completed with no errors.
- A **Fail flag** value of 1 indicates that the run encountered one or more errors.
- The **Failure Count** indicates the number of read transactions where data did not match the expected value.
- The **First failure** address indicates the address corresponding to the first data mismatch.
- The **Version** indicates the version number of the traffic generator.
- The **Number of data generators** indicates the number of data pins at the memory interface.
- The **Number of byte enable generators** indicates the number of byte enable and data mask pins at the memory interface.
- The **Rank Address**, **Bank address**, and **Bank group width** values indicate the number of bits in the Avalon address corresponding to each of those components.
- The **Data/Byte enable pattern length** indicates the number of bits in the fixed pattern used on each data/byte enable pin.
- The **PNF** (pass not fail) value indicates the persistent pass/fail status for each bit in the Avalon data. It is also presented on a per-memory-pin basis for each beat within a memory burst.
- **Fail Expected Data** is the data that was expected on the first failing transaction (if applicable).
- **Fail Read Data** is the data that was received on the first failing transaction (if applicable).



13.12 Calibration Adjustment Delay Step Sizes for Intel Stratix 10 Devices

Refer to the following tables for information on delay step sizes for calibration adjustment.

13.12.1 Addressing

Each reconfigurable feature of the interface has an associated memory address; however, this address is placement dependent. If Altera PHYLite for Parallel Interfaces IP cores and the Intel Stratix 10 External Memory Interfaces IP cores share the same I/O column, you must track the addresses of the interface lanes and the pins. Addressing is done at the 32-bit word boundary, where `avl_address[1:0]` are always 00.


Address Map

These points apply to the following table:

- `id[3:0]` refers to the **Interface ID** parameter.
- `lane_addr[7:0]` refers to the address of a given lane in an interface. The Fitter sets this address value. You can query this in the Parameter Table Lookup Operation Sequence as described in *Address Lookup* section of the *Intel PHYLite for Parallel Interfaces IP Core User Guide*.
- `pin[4:0]` refers to the physical location of the pin in a lane. You can use the Fitter to automatically determine a pin location or you can manually set the pin location through **.qsf** assignment. Refer to the Parameter Table Lookup Operation Sequence as described in *Address Lookup* section of the *Intel PHYLite for Parallel Interfaces IP Core User Guide* for more information.

Feature	Avalon Address R/W	Address CSR R	Control		Value
			Field	Range	
Pin Output Phase	{id[3:0], 3'h4, lane_addr[7:0], pin[4:0], 8'D0}	{id[3:0], 3'h4, lane_addr[7:0], pin[4:0], 8'E8}	Phase Value	12..0	Minimum Setting: Refer to Table 452 on page 474 Maximum Setting: Refer to Table 452 on page 474 Incremental Delay: 1/128th VCO clock period <i>Note:</i> The pin output phase switches from the CSR value to the Avalon value after the first Avalon write. It is only reset to the CSR value on a reset of the interface.
			Reserved ₁	31..13	—
Pin PVT Compensated Input Delay	{id[3:0], 3'h4, lane_addr[7:0]}	Not supported	Delay Value	8..0	Minimum Setting: 0 Maximum Setting: 511 VCO clock periods

continued...



Feature	Avalon Address R/W	Address CSR R	Control		Value
			Field	Range	
	4'hC, lgc_sel[1:0], pin_off[2:0], 4'h0} <ul style="list-style-type: none">lgc_sel[1:0] is:<ul style="list-style-type: none">2'b01 for DQ [5:0]2'b10 for DQ [11:6]pin_off[2:0] :<ul style="list-style-type: none">3'h0: DQ [0], DQ [6]3'h1: DQ [1], DQ [7]3'h2: DQ [2], DQ [8]3'h3: DQ [3], DQ [9]3'h4: DQ [4], DQ [10]3'h5: DQ [5], DQ [11]				Incremental Delay: 1/256th VCO clock period
			Reserved ₁	11..9	—
			Enable	12	0 = Delay value is 0. 1 = Select delay value from Avalon register
			Reserved ₁	31..13	—
Strobe PVT compensated input delay ²	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, lgc_sel[1:0], 3'h6, 4'h0} <ul style="list-style-type: none">lgc_sel[1:0] = 2'b01	Not supported	Delay Value	9..0	Minimum Setting: 0 Maximum Setting: 1023 VCO clock periods Incremental Delay: 1/256th VCO clock period
			Reserved ₁	11..10	—
			Enable	12	0 = Select delay value from CSR register. The CSR value is set through the Capture Strobe Phase Shift parameter during IP core instantiation. 1 = Select delay value from Avalon register
			Reserved ₁	31..13	—
Strobe enable phase ²	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, lgc_sel[1:0], 3'h7, 4'h0} <ul style="list-style-type: none">lgc_sel[1:0] = 2'b01	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h198}	Phase Value	12..0	Minimum Setting: Refer to Table 452 on page 474 Maximum Setting: Refer to Table 452 on page 474 Incremental Delay: 1/128th VCO clock period
			Reserved ₁	14..13	—
			Enable	15	0 = Select delay value from CSR register 1 = Select delay value from Avalon register
			Reserved ₁	31..16	—
continued...					



Feature	Avalon Address R/W	Address CSR R	Control		Value
			Field	Range	
Strobe enable delay ²	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h008}	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h1A8}	Delay Value	5..0	Minimum Setting: 0 external clock cycles Maximum Setting: 63 external memory clock cycles Incremental Delay: 1 external memory clock cycle
			Reserved ₁	14..6	—
			Enable	15	0 = Select delay value from CSR register 1 = Select delay value from Avalon register
			Reserved ₁	31..16	—
Read valid delay ²	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h00C}	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h1A4}	Delay Value	6..0	Minimum Setting: 0 external clock cycles Maximum Setting: 127 external memory clock cycles Incremental Delay: 1 external memory clock cycle
			Reserved ₁	14..7	—
			Enable	15	0 = Select delay value from CSR register 1 = Select delay value from Avalon register
			Reserved ₁	31..16	—
Internal VREF Code	{id[3:0], 3'h4, lane_addr[7:0], 4'hC, 9'h014}	Not supported	VREF Code	5..0	Refer to <i>Calibrated VREF Settings</i> in the <i>Intel PHYLite for Parallel Interfaces IP Core User Guide</i> .
			Reserved ₁	31..6 ⁹	—

1. Reserved bit ranges must be zero.

2. Modifying these values must be done on all lanes in a group.

Note: For more information about performing various clocking and delay calculations, depending on the interface frequency and rate, refer to [PHYLite_delay_calculations.xlsx](#).

13.12.2 Output and Strobe Enable Minimum and Maximum Phase Settings

When dynamically reconfiguring the interpolator phase settings, the values must be kept within the ranges below to ensure proper operation of the circuitry.



Table 452. Output and Strobe Enable Minimum and Maximum Phase Settings

VCO Multiplication Factor	Core Rate	Minimum Interpolator Phase			Maximum Interpolator Phase
		Output	Bidirectional	Bidirectional with OCT Enabled	
1	Full	0x080	0x100	0x100	0xA80
	Half	0x080	0x100	0x100	0xBC0
	Quarter	0x080	0x100	0x100	0xA00
2	Full	0x080	0x100	0x180	0x1400
	Half	0x080	0x100	0x180	0x1400
	Quarter	0x080	0x100	0x180	0x1400
4	Full	0x080	0x100	0x280	0x1FFF
	Half	0x080	0x100	0x280	0x1FFF
	Quarter	0x080	0x100	0x280	0x1FFF
8	Full	0x080	0x100	0x480	0x1FFF
	Half	0x080	0x100	0x480	0x1FFF
	Quarter	0x080	0x100	0x480	0x1FFF

For more information about performing various clocking and delay calculations, depending on the interface frequency and rate, refer to [PHYLite_delay_calculations.xlsx](#).



14 Document Revision History for Intel Stratix 10 External Memory Interfaces IP User Guide

Table 453. Document Revision History

Date	Version	Changes
November 2017	2017.11.06	<ul style="list-style-type: none"> Entire document extensively restructured and revised, consolidating relevant content from the <i>External Memory Interface Handbook</i>. Created <i>End-User Signals</i> chapter, comprising interface and signal descriptions, AFI signals and timing diagrams, and memory-mapped register (MMR) information. Created protocol-specific chapters consolidating parameter descriptions, board skew equations, pin planning information, and board design guidelines for each memory protocol. Created chapters for <i>Timing Closure</i>, <i>Optimizing Controller Performance</i>, and <i>Debugging</i>.
May 2017	2017.05.08	<ul style="list-style-type: none"> Updated the topics in the <i>I/O Column</i> section. Updated <i>DQ and DQS Pins Assignment</i> section with new pin information. Updated the <i>Placement Guidelines</i> section with more detailed description. Updated the <i>Resource Sharing Guidelines for Intel Stratix 10EMIF IP</i> section. Updated the <i>Parameterizing Intel Stratix 10 External Memory Interface IP</i> section. Updated the <i>Parameterizing Altera PHYLite for Parallel Interfaces IP Core</i> section. Added a topic about OCT in the <i>Altera PHYLite for Parallel Interfaces IP Core References</i> section. Added a note that you can only use the Report DDR function if you enable the dynamic reconfiguration feature. The dynamic reconfiguration feature is not available with the current version of the Altera PHYLite for Parallel Interfaces IP core.
October 2016	2016.10.31	<ul style="list-style-type: none"> Initial release.

Intel Corporation. All rights reserved. Intel, the Intel logo, Altera, Arria, Cyclone, Enpirion, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

ISO
9001:2008
Registered