

# UltraScale Architecture FPGAs Memory Interface Solutions v7.1

## *LogiCORE IP Product Guide*

Vivado Design Suite

PG150 June 24, 2015

# Table of Contents

## SECTION I: SUMMARY

### IP Facts

## SECTION II: DDR3/DDR4

### Chapter 1: Overview

Feature Summary . . . . .	12
Licensing and Ordering Information . . . . .	13

### Chapter 2: Product Specification

Standards . . . . .	15
Performance . . . . .	15
Resource Utilization . . . . .	15
Port Descriptions . . . . .	16

### Chapter 3: Core Architecture

Overview . . . . .	17
Memory Controller . . . . .	18
ECC . . . . .	23
PHY . . . . .	26

### Chapter 4: Designing with the Core

Clocking . . . . .	73
Resets . . . . .	77
PCB Guidelines for DDR3 . . . . .	77
PCB Guidelines for DDR4 . . . . .	77
Pin and Bank Rules . . . . .	77
Pin Mapping for x4 RDIMMs . . . . .	90
Protocol Description . . . . .	92
Performance . . . . .	145

## Chapter 5: Design Flow Steps

Customizing and Generating the Core .....	157
MIG I/O Planning .....	164
Constraining the Core .....	165
Simulation .....	166
Synthesis and Implementation .....	166

## Chapter 6: Example Design

Simulating the Example Design (Designs with Standard User Interface) .....	169
Project-Based Simulation .....	170
Non-Project-Based Simulation .....	178
Simulation Speed .....	178
Synplify Pro Black Box Testing .....	179
CLOCK_DEDICATED_ROUTE Constraints and BUFG Instantiation .....	181

## Chapter 7: Test Bench

Stimulus Pattern .....	183
Bus Utilization .....	184
Example Patterns .....	185
Simulating the Performance Traffic Generator .....	188

## SECTION III: QDR II+ SRAM

### Chapter 8: Overview

Feature Summary .....	192
Licensing and Ordering Information .....	192

### Chapter 9: Product Specification

Standards .....	194
Performance .....	194
Resource Utilization .....	194
Port Descriptions .....	195

### Chapter 10: Core Architecture

Overview .....	196
PHY .....	197

### Chapter 11: Designing with the Core

Clocking .....	202
----------------	-----

Resets .....	206
PCB Guidelines for QDR II+ SRAM. ....	206
Pin and Bank Rules. ....	206
Protocol Description .....	211

## Chapter 12: Design Flow Steps

Customizing and Generating the Core .....	216
MIG I/O Planning .....	218
Constraining the Core .....	218
Simulation .....	219
Synthesis and Implementation .....	219

## Chapter 13: Example Design

Simulating the Example Design (Designs with Standard User Interface). ....	222
Project-Based Simulation .....	223
Non-Project-Based Simulation .....	234
Simulation Speed .....	234
Synplify Black Box Testing .....	236
CLOCK_DEDICATED_ROUTE Constraints and BUFG Instantiation .....	237

## Chapter 14: Test Bench

# SECTION IV: RLDRAM 3

## Chapter 15: Overview

Feature Summary. ....	241
Licensing and Ordering Information. ....	242

## Chapter 16: Product Specification

Standards .....	243
Performance. ....	243
Resource Utilization. ....	243
Port Descriptions .....	244

## Chapter 17: Core Architecture

Overview .....	245
Memory Controller .....	247
User Interface Allocation. ....	247
PHY .....	247

## Chapter 18: Designing with the Core

Clocking.....	251
Resets .....	255
PCB Guidelines for RLD RAM 3.....	255
Pin and Bank Rules.....	255
Protocol Description .....	260

## Chapter 19: Design Flow Steps

Customizing and Generating the Core .....	269
MIG I/O Planning .....	271
Constraining the Core .....	271
Simulation .....	272
Synthesis and Implementation .....	272

## Chapter 20: Example Design

Simulating the Example Design (Designs with Standard User Interface).....	275
Project-Based Simulation .....	276
Non-Project-Based Simulation .....	284
Simulation Speed .....	284
CLOCK_DEDICATED_ROUTE Constraints and BUFG Instantiation .....	285

## Chapter 21: Test Bench

## SECTION V: TRAFFIC GENERATOR

### Chapter 22: Traffic Generator

Overview .....	289
Simple Traffic Generator .....	290
Advanced Traffic Generator .....	290

## SECTION VI: MULTIPLE IP CORES

### Chapter 23: Multiple IP Cores

Creating a Design with Multiple IP Cores.....	295
Sharing of a Bank .....	295
Sharing of Input Clock Source .....	296
XSDB and dbg_clk Changes .....	296
PBLOCK and MMCM Constraints .....	296

## SECTION VII: APPENDICES

### Appendix A: Migrating and Upgrading

### Appendix B: Debugging

Finding Help on Xilinx.com .....	299
Debug Tools .....	301
Hardware Debug .....	301

### Appendix C: Additional Resources and Legal Notices

Xilinx Resources .....	303
References .....	303
Revision History .....	304
Please Read: Important Legal Notices .....	313

# SECTION I: SUMMARY

## IP Facts

## Introduction

The Xilinx® UltraScale™ architecture FPGAs Memory Interface Solutions (MIS) core is a combined pre-engineered controller and physical layer (PHY) for interfacing UltraScale architecture FPGA user designs to DDR3 and DDR4 SDRAM, QDR II+ SRAM, and RLDRAM 3 devices.

This product guide provides information about using, customizing, and simulating a LogiCORE™ IP DDR3 or DDR4 SDRAM, QDR II+ SRAM, or a RLDRAM 3 interface core for UltraScale architecture FPGAs. It also describes the core architecture and provides details on customizing and interfacing to the core.

## Features

For feature information on the DDR3/DDR4 SDRAM, QDR II+ SRAM, and RLDRAM 3 interfaces, see the following sections:

- [Feature Summary in Chapter 1](#) for DDR3/DDR4 SDRAM
- [Feature Summary in Chapter 8](#) for QDR II+ SRAM
- [Feature Summary in Chapter 15](#) for RLDRAM 3

LogiCORE IP Facts Table	
Core Specifics	
Supported Device Family <sup>(1)</sup>	Virtex® and Kintex® UltraScale Families
Supported User Interfaces	User
Resources	See <a href="#">Table 2-1</a> , <a href="#">Table 2-2</a> , <a href="#">Table 9-1</a> , and <a href="#">Table 16-1</a> .
Provided with Core	
Design Files	RTL
Example Design	Verilog
Test Bench	Verilog
Constraints File	XDC
Simulation Model	Not Provided
Supported S/W Driver	N/A
Tested Design Flows <sup>(2)</sup>	
Design Entry	Vivado Design Suite
Simulation <sup>(3)</sup>	For supported simulators, see the <a href="#">Xilinx Design Tools: Release Notes Guide</a> .
Synthesis	Vivado Synthesis
Support	
Provided by Xilinx at the <a href="#">Xilinx Support web page</a> .	

### Notes:

1. For a complete listing of supported devices, see the Vivado IP catalog.
2. For the supported versions of the tools, see the [Xilinx Design Tools: Release Notes Guide](#).
3. Behavioral simulations are only supported and netlist (post-synthesis and post-implementation) simulations are not supported.



## SECTION II: DDR3/DDR4

Overview

Product Specification

Core Architecture

Designing with the Core

Design Flow Steps

Example Design

Test Bench

## Overview

The Xilinx® UltraScale™ architecture includes the DDR3/DDR4 SDRAM Memory Interface Solutions (MIS) cores. These MIS cores provide solutions for interfacing with these SDRAM memory types. Both a complete Memory Controller and a physical (PHY) layer only solution are supported. The UltraScale architecture for the DDR3/DDR4 cores are organized in the following high-level blocks.

- **Controller** – The controller accepts burst transactions from the user interface and generates transactions to and from the SDRAM. The controller takes care of the SDRAM timing parameters and refresh. It coalesces write and read transactions to reduce the number of dead cycles involved in turning the bus around. The controller also reorders commands to improve the utilization of the data bus to the SDRAM.
- **Physical Layer** – The physical layer provides a high-speed interface to the SDRAM. This layer includes the hard blocks inside the FPGA and the soft blocks calibration logic necessary to ensure optimal timing of the hard blocks interfacing to the SDRAM.

The new hard blocks in the UltraScale architecture allow interface rates of up to 2,400 Mb/s to be achieved. The application logic is responsible for all SDRAM transactions, timing, and refresh.

- These hard blocks include:
  - Data serialization and transmission
  - Data capture and deserialization
  - High-speed clock generation and synchronization
  - Coarse and fine delay elements per pin with voltage and temperature tracking
- The soft blocks include:
  - **Memory Initialization** – The calibration modules provide a JEDEC®-compliant initialization routine for the particular memory type. The delays in the initialization process can be bypassed to speed up simulation time, if desired.
  - **Calibration** – The calibration modules provide a complete method to set all delays in the hard blocks and soft IP to work with the memory interface. Each bit is individually trained and then combined to ensure optimal interface performance. Results of the calibration process are available through the Xilinx debug tools. After completion of calibration, the PHY layer presents raw interface to the SDRAM.

- **Application Interface** – The user interface layer provides a simple FIFO-like interface to the application. Data is buffered and read data is presented in request order.

The above user interface is layered on top of the native interface to the controller. The native interface is not accessible by the user application and has no buffering and presents return data to the user interface as it is received from the SDRAM which is not necessarily in the original request order. The user interface then buffers the read and write data and reorders the data as needed.

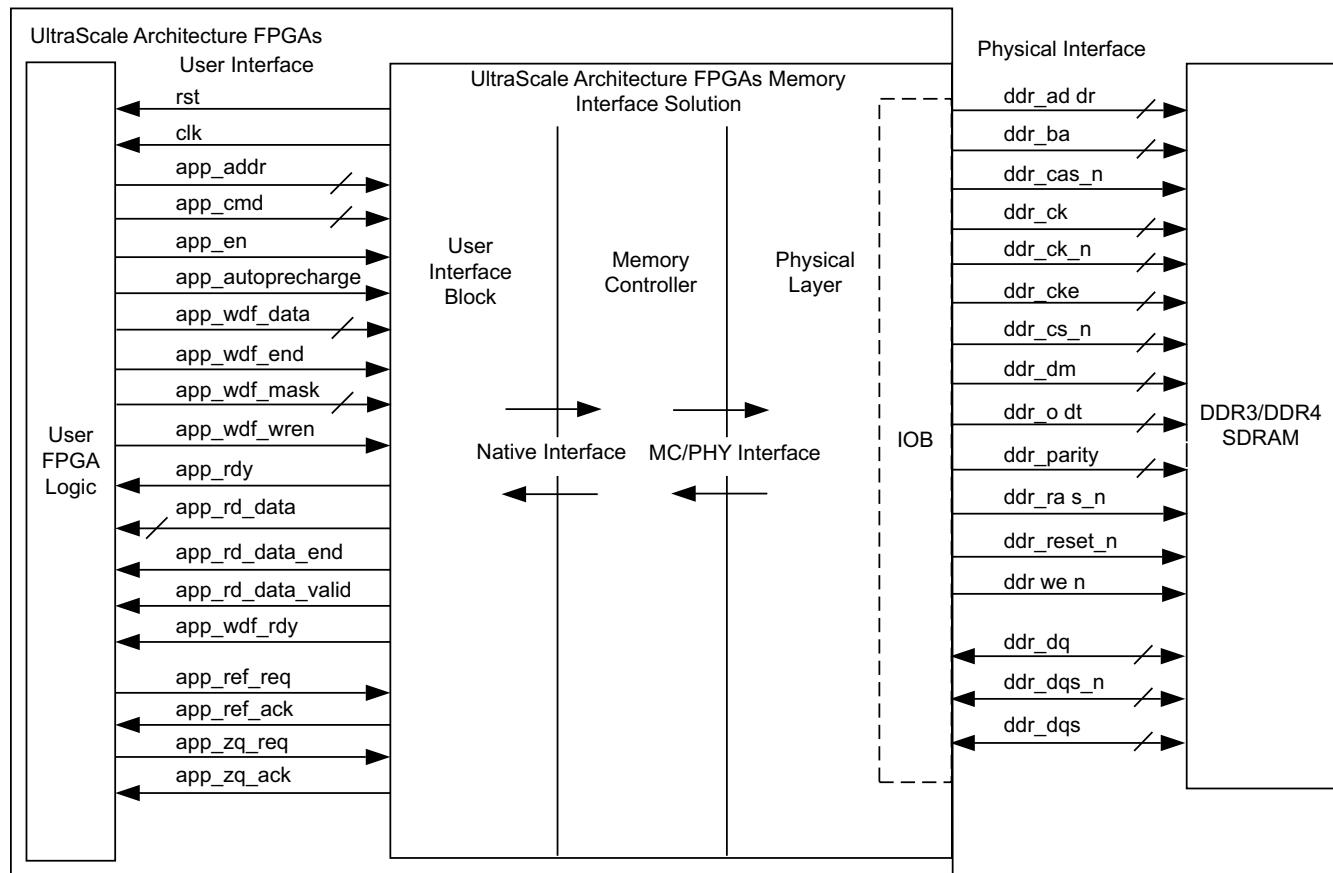


Figure 1-1: UltraScale Architecture FPGAs Memory Interface Solution

---

## Feature Summary

### DDR3 SDRAM

- Component support for interface width of 8 to 80 bits (RDIMM, UDIMM, and SODIMM support)
  - Maximum component limit is 9 and this restriction is valid for components only and not for DIMMs
- DDR3 (1.5V)
- 8 GB density device support
- 8-bank support
- x4 (x4 devices must be used in even multiples), x8, and x16 device support  
**Note:** x4 devices are not supported for AXI interface.
- 8:1 DQ:DQS ratio support for x8 and x16 devices
- 4:1 DQ:DQS ratio support for x4 devices
- 8-word burst support
- Support for 5 to 14 cycles of column-address strobe (CAS) latency (CL)
- On-die termination (ODT) support
- Support for 5 to 10 cycles of CAS write latency
- Write leveling support for DDR3 (fly-by routing topology required component designs)
- JEDEC®-compliant DDR3 initialization support
- Source code delivery in Verilog
- 4:1 memory to FPGA logic interface clock ratio
- Open, closed, and transaction based pre-charge controller policy
- Interface calibration and training information available through the Vivado hardware manager

### DDR4 SDRAM

- Component support for interface width of 8 to 80 bits (RDIMM, UDIMM, and SODIMM support)
  - Maximum component limit is 9 and this restriction is valid for components only and not for DIMMs
- 8 GB density device support

- x4 (x4 devices must be used in even multiples), x8, and x16 device support
  - **Note:** x4 devices are not supported for AXI interface.
  - 8:1 DQ:DQS ratio support for x8 and x16 devices
  - 4:1 DQ:DQS ratio support for x4 devices
  - 8-word burst support
  - Support for 9 to 24 cycles of column-address strobe (CAS) latency (CL)
  - ODT support
  - Support for 9 to 18 cycles of CAS write latency
  - Write leveling support for DDR4 (fly-by routing topology required component designs)
  - JEDEC-compliant DDR4 initialization support
  - Source code delivery in Verilog
  - 4:1 memory to FPGA logic interface clock ratio
  - Open, closed, and transaction based pre-charge controller policy
  - Interface calibration and training information available through the Vivado hardware manager
- 

## Licensing and Ordering Information

This Xilinx LogiCORE IP module is provided at no additional cost with the Xilinx Vivado Design Suite under the terms of the [Xilinx End User License](#). Information about this and other Xilinx LogiCORE IP modules is available at the [Xilinx Intellectual Property](#) page. For information about pricing and availability of other Xilinx LogiCORE IP modules and tools, contact your [local Xilinx sales representative](#).

## License Checkers

If the IP requires a license key, the key must be verified. The Vivado<sup>®</sup> design tools have several license checkpoints for gating licensed IP through the flow. If the license check succeeds, the IP can continue generation. Otherwise, generation halts with error. License checkpoints are enforced by the following tools:

- Vivado design tools: Vivado synthesis
- Vivado implementation
- write\_bitstream (Tcl command)



---

**IMPORTANT:** *IP license level is ignored at checkpoints. The test confirms a valid license exists. It does not check IP license level.*

---

## Product Specification

### Standards

This core supports DRAMs that are compliant to the JESD79-3F, *DDR3 SDRAM Standard* and JESD79-4, *DDR4 SDRAM Standard*, JEDEC® Solid State Technology Association [Ref 1].

For more information on UltraScale™ architecture documents, see [References, page 303](#).

### Performance

#### Maximum Frequencies

For more information on the maximum frequencies, see *Kintex UltraScale Architecture Data Sheet, DC and AC Switching Characteristics* (DS892) [Ref 2].

### Resource Utilization

#### Kintex UltraScale Devices

[Table 2-1](#) and [Table 2-2](#) provide approximate resource counts using Kintex® UltraScale devices.

**Table 2-1: Device Utilization – Kintex UltraScale FPGAs for DDR3**

Parameter Values	Device Resources						
Interface Width	FFs	LUTs	Memory LUTs	RAMB36E2/ RAMB18E2	BUFGs	PLLE3_ADV	MMCME3_ADV
72	13,423	11,940	1,114	25.5	5	3	1
32	8,099	7,305	622	25.5	5	2	1
16	6,020	5,621	426	25.5	5	1	1

Table 2-2: Device Utilization – Kintex UltraScale FPGAs for DDR4

Parameter Values	Device Resources						
Interface Width	FFs	LUTs	Memory LUTs	RAMB36E2/ RAMB18E2	BUFGs	PLLE3_ADV	MMCM3E3_ADV
72	13,535	11,905	1,105	25.5	5	3	1
32	8,010	7,161	622	25.5	5	2	1
16	5,864	5,363	426	25.5	5	1	1

Resources required for the UltraScale architecture FPGAs MIS core have been estimated for the Kintex UltraScale devices. These values were generated using Vivado® IP catalog. They are derived from post-synthesis reports, and might change during implementation.

## Port Descriptions

For a complete Memory Controller solution there are three port categories at the top-level of the memory interface core called the “user design.”

- The first category is the memory interface signals that directly interfaces with the SDRAM. These are defined by the JEDEC specification.
- The second category is the application interface signals. These are described in the [Protocol Description, page 92](#).
- The third category includes other signals necessary for proper operation of the core. These include the clocks, reset, and status signals from the core. The clocking and reset signals are described in their respective sections.

The active-High `init_calib_complete` signal indicates that the initialization and calibration are complete and that the interface is now ready to accept commands for the interface.

For a PHY layer only solution, the top-level application interface signals are replaced with the PHY interface. These signals are described in the [PHY Only Interface, page 122](#).

The signals that interface directly with the SDRAM and the clocking and reset signals are the same as for the Memory Controller solution.



## Core Architecture

This chapter describes the UltraScale™ device FPGAs Memory Interface Solutions core with an overview of the modules and interfaces.

### Overview

The UltraScale architecture FPGAs Memory Interface Solutions is shown in [Figure 3-1](#).




Figure 3-1: UltraScale Architecture FPGAs Memory Interface Solution Core

---

## Memory Controller

The Memory Controller (MC) is designed to take Read, Write, and Read-Modify-Write transactions from the user interface (UI) block and issues them to memory efficiently with low latency, meeting all DRAM protocol and timing requirements, while using minimal FPGA resources. The controller operates with a DRAM to system clock ratio of 4:1 and can issue one Activate, one CAS, and one Precharge command on each system clock cycle.

The controller supports an open page policy and can achieve very high efficiencies with workloads with a high degree of spatial locality. The controller also supports a closed page policy and the ability to reorder transactions to efficiently schedule workloads with address patterns that are more random. The controller also allows a degree of control over low-level functions with a UI control signal for AutoPrecharge on a per transaction basis as well as signals that can be used to determine when DRAM refresh commands are issued.

The key blocks of the controller command path include:

1. The Group FSMs that queue up transactions, check DRAM timing, and decide when to request Precharge, Activate, and CAS DRAM commands.
2. The "Safe" logic and arbitration units that reorder transactions between Group FSMs based on additional DRAM timing checks while also ensuring forward progress for all DRAM command requests.
3. The Final Arbiter that makes the final decision about which commands are issued to the PHY and feeds the result back to the previous stages.

There are also maintenance blocks that generate refresh and ZQCS commands as well as commands needed for VT tracking. Also, there is an optional block that implements a SECDED ECC for 72-bit wide data buses.

Figure 3-2 shows the Memory Controller block diagram.

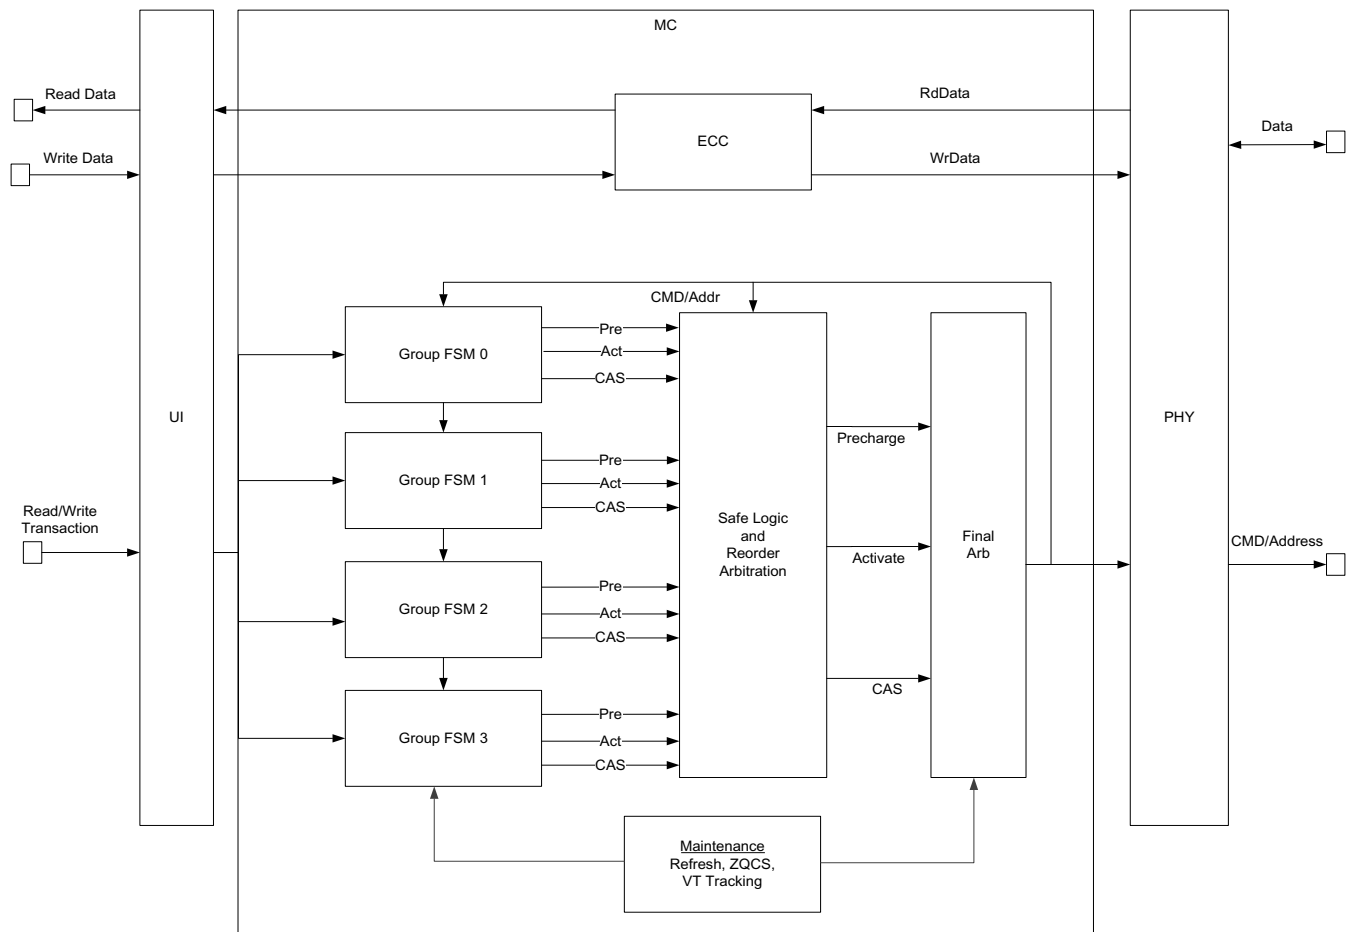


Figure 3-2: Memory Controller Block Diagram

## Native Interface

The UI block is connected to the Memory Controller by the native interface, and provides the controller with address decode and read/write data buffering. On writes, data is requested by the controller one cycle before it is needed by presenting the data buffer address on the native interface. This data is expected to be supplied by the UI block on the next cycle. Hence there is no buffering of any kind for data (except due to the barrel shifting to place the data on a particular DDR clock).

On reads, the data is offered by the MC on the cycle it is available. Read data, along with a buffer address is presented on the native interface as soon as it is ready. The data has to be accepted by the UI block.

Read and write transactions are mapped to an mcGroup instance based on bank group and bank address bits of the decoded address from the UI block. Although there are no groups in DDR3, the name group represents either a real group in DDR4 x4 and x8 devices (which serves four banks of that group). For DDR3, each mcGroup module would service two banks.

In the case of DDR4 x16 interface, the mcGroup represents 1-bit of group (there are only one group bit in x16) and 1-bit of bank, whereby the mcGroup serves two banks.

The total number of outstanding requests depends on the number of mcGroup instances, as well as the round trip delay from the controller to memory and back. When the controller issues an SDRAM CAS command to memory, an mcGroup instance becomes available to take a new request, while the previous CAS commands, read return data, or write data might still be in flight.

## Control and Datapaths

### *Control Path*

The control path starts at the mcGroup instances. The mapping of SDRAM group and bank addresses to mcGroup instance ensures that transactions to the same full address map to the same mcGroup instance. Because each mcGroup instance processes the transactions it receives in order, read-after-write and write-after-write address hazards are prevented.

### *Datapath*

Read and write data pass through the Memory Controller. If ECC is enabled, a SECDEC code word is generated on writes and checked on reads. For more information, see [ECC, page 23](#). The MC generates the requisite control signals to the mcRead and mcWrite modules telling them the timing of read and write data. The two modules acquire or provide the data as required at the right time.

## Read and Write Coalescing

The controller prioritizes reads over writes when reordering is enabled. If both read and write CAS commands are safe to issue on the SDRAM command bus, the controller selects only read CAS commands for arbitration. When a read CAS issues, write CAS commands are blocked for several SDRAM clocks specified by parameter tRTW. This extra time required for a write CAS to become safe after issuing a read CAS allows groups of reads to issue on the command bus without being interrupted by pending writes.

## Reordering

Requests that map to the same mcGroup are never reordered. Reordering between the mcGroup instances is controlled with the ORDERING parameter. When set to "NORM," reordering is enabled and the arbiter implements a round-robin priority plan, selecting in priority order among the mcGroups with a command that is safe to issue to the SDRAM.

The timing of when it is safe to issue a command to the SDRAM can vary on the target bank or bank group and its page status. This often contributes to reordering.

When the ORDERING parameter is set to "STRICT," all requests have their CAS commands issued in the order in which the requests were accepted at the native interface. STRICT ordering overrides all other controller mechanisms, such as the tendency to coalesce read requests, and can therefore degrade data bandwidth utilization in some workloads.

## Group Machines

In the Memory Controller, there are four group state machines. These state machines are allocated depending on technology (DDR3 or DDR4) and width (x4, x8, and x16). The following summarizes the allocation to each group machine. In this description, GM refers to the Group Machine (0 to 3), BG refers to group address, and BA refers to bank address. Note that group in the context of a group state machine denotes a notional group and does not necessarily refer to a real group (except in case of DDR4, part x4 and x8).

- DDR3, any part – Total of eight banks
  - GM 0: BA[2:1] == 2'b00; services banks 0 and 1
  - GM 1: BA[2:1] == 2'b01; services banks 2 and 3
  - GM 2: BA[2:1] == 2'b10; services banks 4 and 5
  - GM 3: BA[2:1] == 2'b11; services banks 6 and 7
- DDR4, x4 and x8 parts – Total of 16 banks
  - GM 0: services BG 0; four banks per group
  - GM 1: services BG 1; four banks per group
  - GM 2: services BG 2; four banks per group
  - GM 3: services BG 3; four banks per group
- DDR4, x16 parts – Total of eight banks
  - GM 0: services BG 0, BA[0] == 0; 2 banks per group
  - GM 1: services BG 0, BA[0] == 1; 2 banks per group
  - GM 2: services BG 1, BA[0] == 0; 2 banks per group
  - GM 3: services BG 1, BA[0] == 1; 2 banks per group

Figure 3-3 shows the Group FSM block diagram for one instance. There are two main sections to the Group FSM block, stage 1 and stage 2, each containing a FIFO and an FSM. Stage 1 interfaces to the UI, issues Precharge and Activate commands, and tracks the DRAM page status.

Stage 2 issues CAS commands and manages the RMW flow. There is also a set of DRAM timers for each rank and bank used by the FSMs to schedule DRAM commands at the earliest safe time. The Group FSM block is designed so that each instance queues up multiple transactions from the UI, interleaves DRAM commands from multiple transactions onto the DDR bus for efficiency, and executes CAS commands strictly in order.

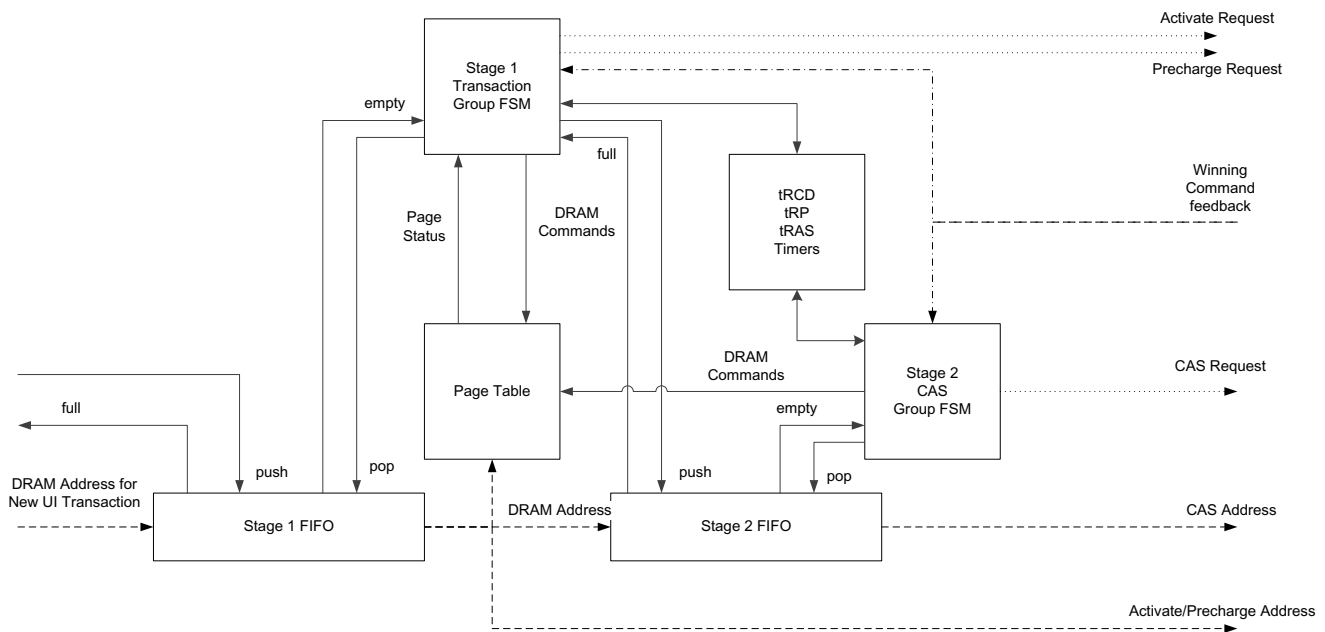


Figure 3-3: Group FSM Block Diagram

When a new transaction is accepted from the UI, it is pushed into the stage 1 transaction FIFO. The page status of the transaction at the head of the stage 1 FIFO is checked and provided to the stage 1 transaction FSM. The FSM decides if a Precharge or Activate command needs to be issued, and when it is safe to issue them based on the DRAM timers.

When the page is open and not already scheduled to be closed due to a pending RDA or WRA in the stage 2 FIFO, the transaction is transferred from the stage 1 FIFO to the stage 2 FIFO. At this point, the stage 1 FIFO is popped and the stage 1 FSM begins processing the next transaction. In parallel, the stage 2 FSM processes the CAS command phase of the transaction at the head of the stage 2 FIFO. The stage 2 FSM issues a CAS command request when it is safe based on the tRCD timers. The stage 2 FSM also issues both a read and write CAS request for RMW transactions.

## ECC

The MC supports an optional SECDED ECC scheme that detects and corrects read data errors with 1-bit error per DQ bus burst and detects all 2-bit errors per burst. The 2-bit errors are not corrected. Three or more bit errors per burst might or might not be detected, but are never corrected. Enabling ECC adds four DRAM clock cycles of latency to all reads, whether errors are detected/corrected or not.

A Read-Modify-Write (RMW) scheme is also implemented to support Partial Writes when ECC is enabled. Partial Writes have one or more user interface write data mask bits set High. Partial Writes with ECC disabled are handled by sending the data mask bits to the DRAM Data Mask (DM) pins, so the RMW flow is used only when ECC is enabled. When ECC is enabled, Partial Writes require their own command, `wr_bytes` or 0x3, so the MC knows when to use the RMW flow.

## Read-Modify-Write Flow

When a `wr_bytes` command is accepted at the user interface it is eventually assigned to a group state machine like other write or read transactions. The group machine breaks the Partial Write into a read phase and a write phase. The read phase performs the following:

1. First reads data from memory.
2. Checks for errors in the read data.
3. Corrects single bit errors.
4. Stores the result inside the Memory Controller.

Data from the read phase is not returned to the user interface. If errors are detected in the read data, an ECC error signal is asserted at the native interface. After read data is stored in the controller, the write phase begins as follows:

1. Write data is merged with the stored read data based on the write data mask bits.
2. New ECC check bits are generated for the merged data and check bits are written to memory.
3. Any multiple bit errors in the read phase results in the error being made undetectable in the write phase as new check bits are generated for the merged data. This is why the ECC error signal is generated on the read phase even though data is not returned to the user interface. This allows the system to know if an uncorrectable error has been turned into an undetectable error.

When the write phase completes, the group machine becomes available to process a new transaction. The RMW flow ties up a group machine for a longer time than a simple read or write, and therefore might impact performance.

## ECC Module

The ECC module is instantiated inside the DDR3/DDR4 Memory Controller. It is made up of five submodules as shown in Figure 3-4.

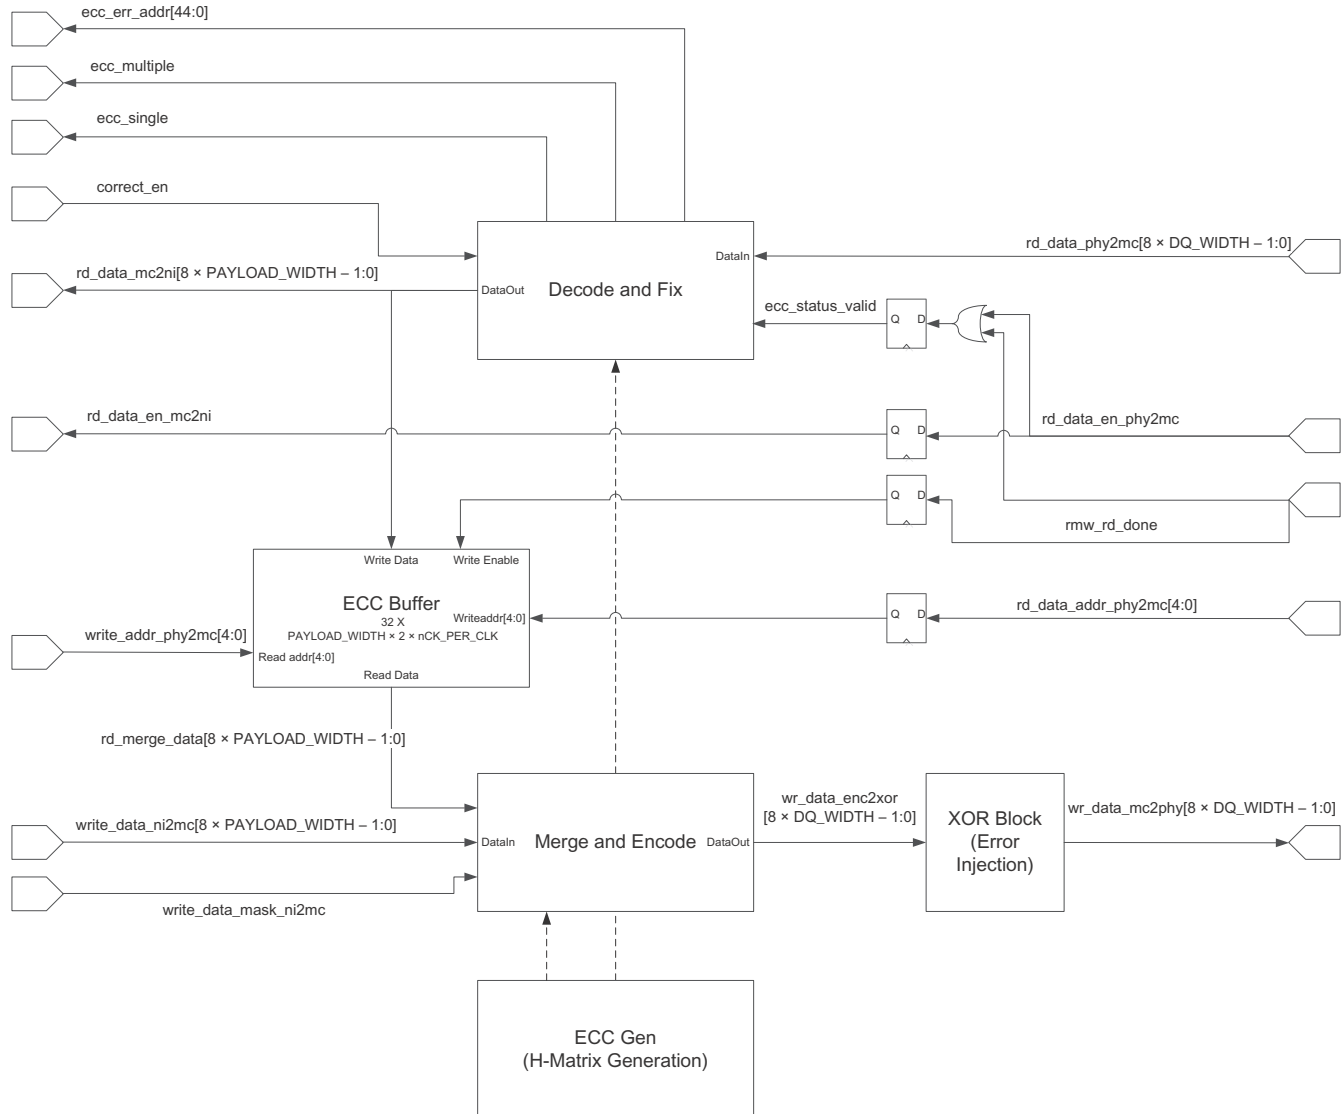


Figure 3-4: ECC Block Diagram

Read data and check bits from the PHY are sent to the Decode block, and on the next system clock cycle data and error indicators `ecc_single`/`ecc_multiple` are sent to the NI. `ecc_single` asserts when a correctable error is detected and the read data has been corrected. `ecc_multiple` asserts when an uncorrectable error is detected.

Read data is not modified by the ECC logic on an uncorrectable error. Error indicators are never asserted for “periodic reads,” which are read transactions generated by the controller only for the purposes of VT tracking and are not returned to the user interface or written back to memory in an RMW flow.



Write data is merged in the Encode block with read data stored in the ECC Buffer. The merge is controlled on a per byte basis by the write data mask signal. All writes use this flow, so full writes are required to have all data mask bits deasserted to prevent unintended merging. After the Merge stage, the Encode block generates check bits for the write data. The data and check bits are output from the Encode block with a one system clock cycle delay.

The ECC Gen block implements an algorithm that generates an H-matrix for ECC check bit generation and error checking/correction. The generated code depends only on the PAYLOAD\_WIDTH and DQ\_WIDTH parameters, where  $DQ\_WIDTH = PAYLOAD\_WIDTH + ECC\_WIDTH$ . Currently only  $DQ\_WIDTH = 72$  and  $ECC\_WIDTH = 8$  is supported.

## Error Address

Each time a read CAS command is issued, the full DRAM address is stored in a FIFO in the decode block. When read data is returned and checked for errors, the DRAM address is popped from the FIFO and `ecc_err_addr[44:0]` is returned on the same cycle as signals `ecc_single` and `ecc_multiple` for the purposes of error logging or debug. Table 3-1 is a common definition of this address for DDR3 and DDR4.

Table 3-1: ECC Error Address Definition

<b>ecc_err_addr [44:0]</b>	<b>44</b>	<b>43:42</b>	<b>41:40</b>	<b>39:24</b>	<b>23:22</b>	<b>21:18</b>	<b>17:8</b>	<b>7:6</b>	<b>5:4</b>	<b>3</b>	<b>2</b>	<b>1:0</b>
DDR4 (x4/x8)	RMW	RSVD	Row[17:0]	RSVD	RSVD	Col [9:0]	RSVD	Rank [1:0]	Group [1:0]			Bank [1:0]
DDR (x16)	RMW	RSVD	Row[17:0]	RSVD	RSVD	Col [9:0]	RSVD	Rank [1:0]	RSVD	Group [0]		Bank [1:0]
DDR3	RMW	RSVD		Row [15:0]	RSVD	Col[13:0]	RSVD	Rank [1:0]	RSVD		Bank[2:0]	

## Latency

When the parameter ECC is "ON," the ECC modules are instantiated and read and write data latency through the MC increases by one system clock cycle. When ECC is "OFF," the data buses just pass through the MC and all ECC logic should be optimized out.

## PHY

PHY is considered the low-level physical interface to an external DDR3 or DDR4 SDRAM device as well as all calibration logic for ensuring reliable operation of the physical interface itself. PHY generates the signal timing and sequencing required to interface to the memory device.

PHY contains the following features:

- Clock/address/control-generation logics
- Write and read datapaths
- Logic for initializing the SDRAM after power-up

In addition, PHY contains calibration logic to perform timing training of the read and write datapaths to account for system static and dynamic delays.

The PHY is included in the complete Memory Controller solution, but can also be implemented as a standalone PHY only block. A PHY only solution can be selected if you plan to implement a custom Memory Controller. For details about interfacing to the PHY only block, see the [PHY Only Interface, page 122](#).

## Overall PHY Architecture

The UltraScale architecture PHY is composed of dedicated blocks and soft calibration logic. The dedicated blocks are structured adjacent to one another with back-to-back interconnects to minimize the clock and datapath routing necessary to build high performance physical layers.

The Memory Controller and calibration logic communicate with this dedicated PHY in the slow frequency clock domain, which is either divided by four or divided by two. This depends on the DDR3 or DDR4 memory clock. A more detailed block diagram of the PHY design is shown in [Figure 3-5](#).

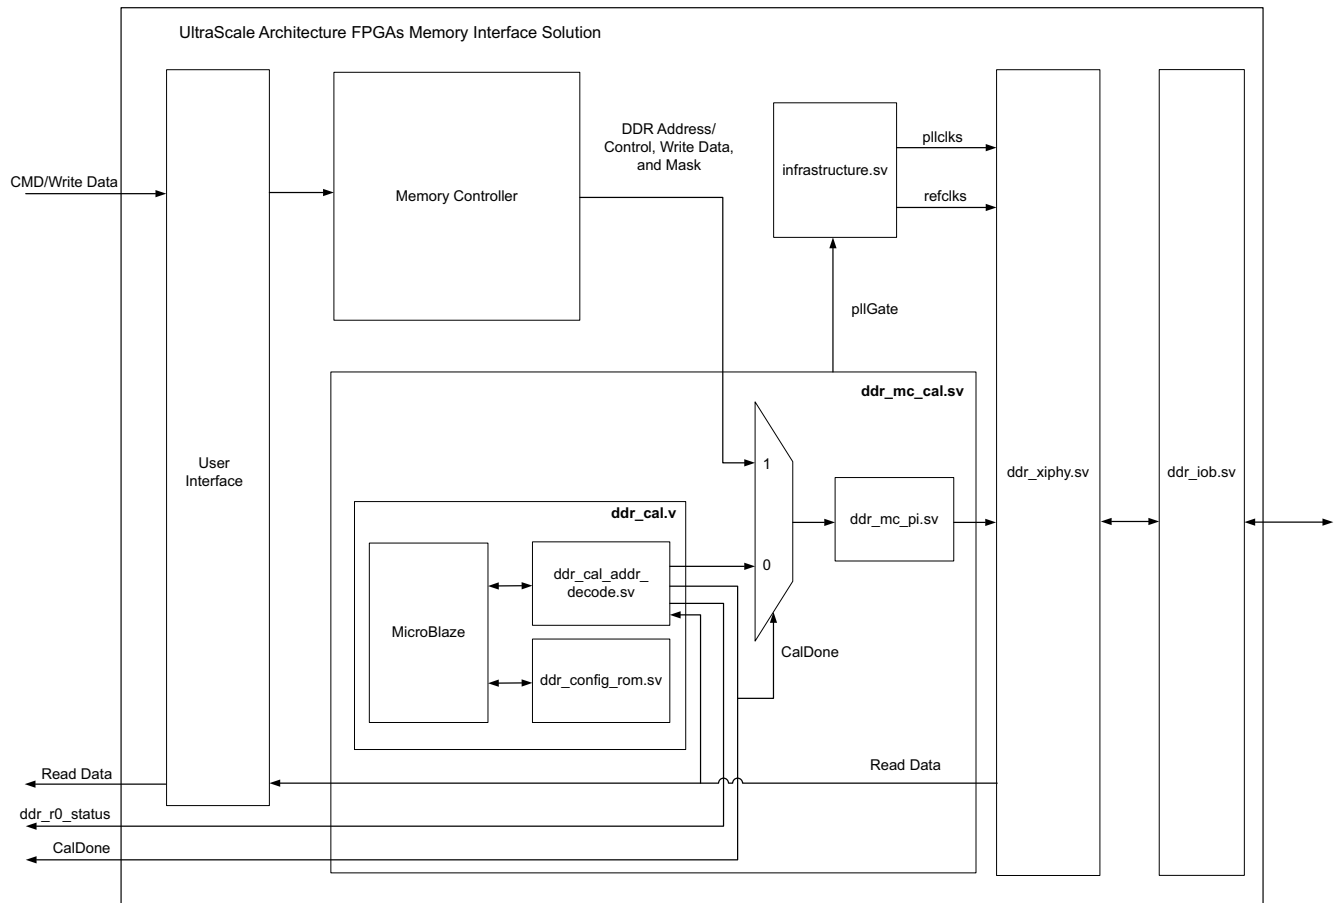


Figure 3-5: PHY Block Diagram

The Memory Controller is designed to separate out the command processing from the low-level PHY requirements to ensure a clean separation between the controller and physical layer. The command processing can be replaced with custom logic if desired, while the logic for interacting with the PHY stays the same and can still be used by the calibration logic.

Table 3-2: PHY Modules

Module Name	Description
ddr_mc_cal.v	Contains ddr_cal.v, ddr_mc_pi.v, and MUXes between the calibration and the Memory Controller.
ddr_cal.v	Contains the MicroBlaze processing system and associated logic.
ddr_mc_pi.v	Adjusts signal timing for the PHY for reads and writes.
ddr_cal_addr_decode.v	FPGA logic interface for the MicroBlaze processor.
ddr_config_rom.v	Configuration storage for calibration options.
microblaze	MicroBlaze processor
ddr_iob.v	Instantiates all byte IOB modules.
ddr_iob_byte.v	Generates the I/O buffers for all the signals in a given byte lane.

Table 3-2: PHY Modules (Cont'd)

Module Name	Description
ddr_debug_microblaze.sv	Simulation-only file to parse debug statements from software running in MicroBlaze to indicate status and calibration results to the log.
ddr_cal_cplx.sv	RTL state machine for complex pattern calibration.
ddr_cal_cplx_data.sv	Data patterns used for complex pattern calibration.
ddr_xiphy.sv	Top-level XIPHY module.
ddr_phy.sv	Top-level of the PHY, contains ddr_mc_cal.sv and ddr_xiphy.sv modules.

The PHY architecture encompasses all of the logic contained in `ddr_phy.sv`. The PHY contains wrappers around dedicated hard blocks to build up the memory interface from smaller components. A byte lane contains all of the clocks, resets, and datapaths for a given subset of I/O. Multiple byte lanes are grouped together, along with dedicated clocking resources, to make up a single bank memory interface. Each nibble in the PHY contains a Register Interface Unit (RIU), a dedicated integrated block in the XIPHY that provides an interface to the general interconnect logic for changing settings and delays for calibration. For more information on the hard silicon physical layer architecture, see the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4].

The memory initialization is executed in Verilog RTL. The calibration and training are implemented by an embedded MicroBlaze™ processor. The MicroBlaze Controller System (MCS) is configured with an I/O Module and a block RAM. The `ddr_cal_addr_decode.sv` module provides the interface for the processor to the rest of the system and implements helper logic. The `ddr_config_rom.sv` module stores settings that control the operation of initialization and calibration, providing run time options that can be adjusted without having to recompile the source code.

The address unit connects the MCS to the local register set and the PHY by performing address decode and control translation on the I/O module bus from spaces in the memory map and MUXing return data (`ddr_cal_addr_decode.sv`). In addition, it provides address translation (also known as “mapping”) from a logical conceptualization of the DRAM interface to the appropriate pinout-dependent location of the delay control in the PHY address space.

Although the calibration architecture presents a simple and organized address map for manipulating the delay elements for individual data, control and command bits, there is flexibility in how those I/O pins are placed. For a given I/O placement, the path to the FPGA logic is locked to a given pin. To enable a single binary software file to work with any memory interface pinout, a translation block converts the simplified RIU addressing into the pinout-specific RIU address for the target design (see Table 3-3). The specific address translation is written by MIG after a pinout is selected and cannot be modified. The code shows an example of the RTL structure that supports this.

```

Casez(io_address)// MicroBlaze I/O module address
// ... static address decoding skipped
//=====//
//=====DQ ODELAYS=====//
//=====//
//Byte0
28'h0004100: begin //c0_ddr4_dq[0] IO_L20P_T3L_N2_AD1P_44
    riu_addr_cal = 6'hD;
    riu_nibble = 'h6;
end
// ... additional dynamic addressing follows

```

In this example, DQ0 is pinned out on Bit[0] of nibble 0 (nibble 0 according to instantiation order). The RIU address for the ODELAY for Bit[0] is 0x0D. When DQ0 is addressed — indicated by address 0x000\_4100), this snippet of code is active. It enables nibble 0 (decoded to one-hot downstream) and forwards the address 0x0D to the RIU address bus.

The MicroBlaze I/O module interface is not always fast enough for implementing all of the functions required in calibration. A helper circuit implemented in `ddr_cal_addr_decode.sv` is required to obtain commands from the registers and translate at least a portion into single-cycle accuracy for submission to the PHY. In addition, it supports command repetition to enable back-to-back read transactions and read data comparison.

**Table 3-3: XIPHY RIU Addressing and Description**

RIU Address	Name	Description
0x00	NIBBLE_CTRL0	Nibble Control 0. Control for enabling DQS gate in the XIPHY, GT_STATUS for gate feedback, and clear gate which resets gate circuit.
0x01	NIBBLE_CTRL1	Nibble Control 1. TX_DATA_PHASE control for every bit in the nibble.
0x02	CALIB_CTRL	Calibration Control. XIPHY control and status for BISC.
0x03	Reserved	Reserved
0x04	Reserved	Reserved
0x05	BS_CTRL	Bit slice reset. Resets the ISERDES and IFIFOs in a given nibble.
0x06	Reserved	Reserved
0x07	PQTR	Rising edge delay for DQS.
0x08	NQTR	Falling edge delay for DQS.
0x09	Reserved	Reserved
0x0A	TRISTATE_ODELAY	Output delay for 3-state.
0x0B	ODELAY0	Output delay for bit slice 0.
0x0C	ODELAY1	Output delay for bit slice 1.
0x0D	ODELAY2	Output delay for bit slice 2.
0x0E	ODELAY3	Output delay for bit slice 3.
0x0F	ODELAY4	Output delay for bit slice 4.
0x10	ODELAY5	Output delay for bit slice 5.

**Table 3-3: XIPHY RIU Addressing and Description (Cont'd)**

RIU Address	Name	Description
0x11	ODELAY6	Output delay for bit slice 6.
0x12	IDELAY0	Input delay for bit slice 0.
0x13	IDELAY1	Input delay for bit slice 1.
0x14	IDELAY2	Input delay for bit slice 2.
0x15	IDELAY3	Input delay for bit slice 3.
0x16	IDELAY4	Input delay for bit slice 4.
0x17	IDELAY5	Input delay for bit slice 5.
0x18	IDELAY6	Input delay for bit slice 6.
0x19	PQTR Align	BISC edge alignment computation for rising edge DQS.
0x1A	NQTR Align	BISC edge alignment computation for falling edge DQS.
0x1B to 0x2B	Reserved	Reserved
0x2C	WL_DLY_RNK0	Write Level register for Rank 0. Coarse and fine delay, WL_TRAIN.
0x2D	WL_DLY_RNK1	Write Level register for Rank 1. Coarse and fine delay.
0x2E	WL_DLY_RNK2	Write Level register for Rank 2. Coarse and fine delay.
0x2F	WL_DLY_RNK3	Write Level register for Rank 3. Coarse and fine delay.
0x30	RL_DLY_RNK0	DQS Gate register for Rank 0. Coarse and fine delay.
0x31	RL_DLY_RNK1	DQS Gate register for Rank 1. Coarse and fine delay.
0x32	RL_DLY_RNK2	DQS Gate register for Rank 2. Coarse and fine delay.
0x33	RL_DLY_RNK3	DQS Gate register for Rank 3. Coarse and fine delay.
0x34 to 0x3F	Reserved	Reserved

## Memory Initialization and Calibration Sequence

After deassertion of the system reset, PHY performs some required internal calibration steps first.

1. The built-in self-check of the PHY (BISC) is run. BISC is used in the PHY to compute internal skews for use in voltage and temperature tracking after calibration is completed.
2. After BISC is completed, calibration logic performs the required power-on initialization sequence for the memory.
3. This is followed by several stages of timing calibration for the write and read datapaths.
4. After calibration is completed, PHY calculates internal offsets to be used in voltage and temperature tracking.
5. PHY indicates calibration is finished and the controller begins issuing commands to the memory.

Figure 3-6 shows the overall flow of memory initialization and the different stages of calibration.

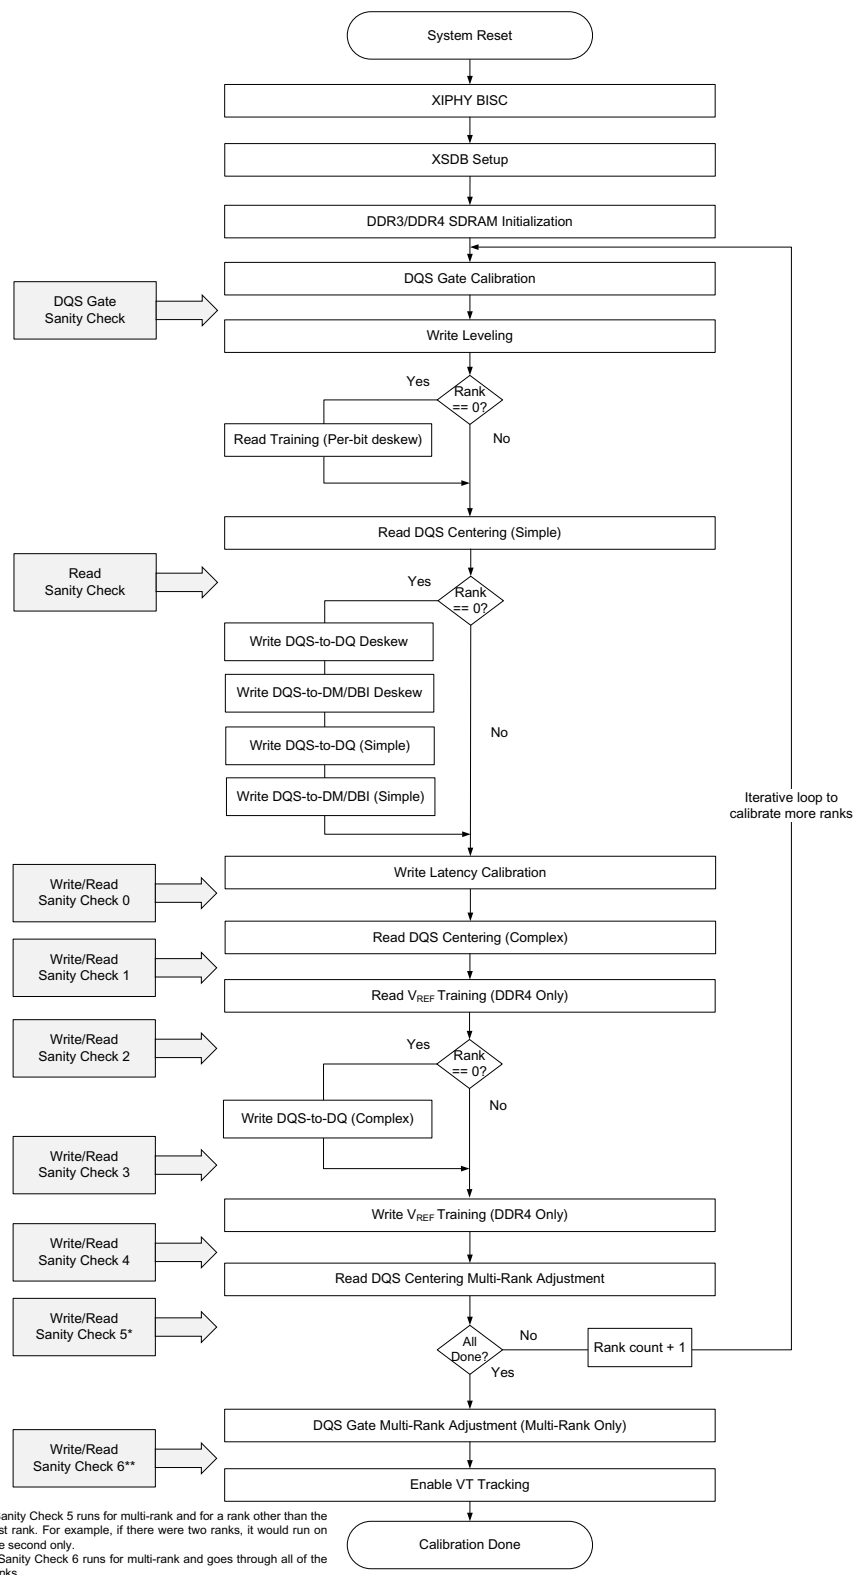


Figure 3-6: PHY Overall Initialization and Calibration Sequence

When simulating a design out of MIG, the calibration is set to be bypassed to enable you to generate traffic to and from the DRAM as quickly as possible. When running in hardware or simulating with calibration, enabled signals are provided to indicate what step of calibration is running or, if an error occurs, where an error occurred.

The first step in determining calibration status is to check the CalDone port. After the CalDone port is checked, the status bits should be checked to indicate the steps that were ran and completed. Calibration halts on the very first error encountered, so the status bits indicate which step of calibration was last run. The status and error signals can be checked through either connecting the Vivado analyzer signals to these ports or through the XSDB tool (also through Vivado).

The calibration status is provided through the XSDB port, which stores useful information regarding calibration for display in the Vivado IDE. The calibration status and error signals are also provided as ports to allow for debug or triggering. Table 3-4 lists the pre-calibration status signal description.

**Table 3-4: Pre-Calibration XSDB Status Signal Description**

XSDB Status Register	XSDB Bits[8:0]	Description	Pre-Calibration Step
DDR_PRE_CAL_STATUS	0	Done	MicroBlaze has started up
	1	Done	Reserved
	2	Done	Reserved
	3	Done	Reserved
	4	Done	XSDB Setup Complete
	5	–	Reserved
	6	–	Reserved
	7	–	Reserved
	8	–	Reserved



Table 3-5 lists the status signals in the port as well as how they relate to the core XSDB data.

Table 3-5: XSDB Status Signal Descriptions

XSDB Status Register	XSDB Bits[8:0]	Status Port Bits[40:0]	Description	Calibration Stage Name	Calibration Stage Number
DDR_CAL_STATUS_RANKx_0	0	0	Start	DQS Gate	1
	1	1	Done	–	–
	2	2	Start	Check for DQS gate	2
	3	3	Done	–	–
	4	4	Start	Write leveling	3
	5	5	Done	–	–
	6	6	Start	Read Per-bit Deskew	4
	7	7	Done	–	–
	8	8	Start	Reserved	5
DDR_CAL_STATUS_RANKx_1	0	9	Done	–	–
	1	10	Start	Read DQS Centering (Simple)	6
	2	11	Done	–	–
	3	12	Start	Read Sanity Check	7
	4	13	Done	–	–
	5	14	Start	Write DQS-to-DQ Deskew	8
	6	15	Done	–	–
	7	16	Start	Write DQS-to-DM Deskew	9
	8	17	Done	–	–
DDR_CAL_STATUS_RANKx_2	0	18	Start	Write DQS-to-DQ (Simple)	10
	1	19	Done	–	–
	2	20	Start	Write DQS-to-DM (Simple)	11
	3	21	Done	–	–
	4	22	Start	Reserved	12
	5	23	Done	–	–
	6	24	Start	Write Latency Calibration	13
	7	25	Done	–	–
	8	26	Start	Write/Read Sanity Check 0	14

Table 3-5: XSDB Status Signal Descriptions (Cont'd)

XSDB Status Register	XSDB Bits[8:0]	Status Port Bits[40:0]	Description	Calibration Stage Name	Calibration Stage Number
DDR_CAL_STATUS_RANKx_3	0	27	Done	–	–
	1	28	Start	Read DQS Centering (Complex)	15
	2	29	Done	–	–
	3	30	Start	Write/Read Sanity Check 1	16
	4	31	Done	–	–
	5	32	Start	Reserved	17
	6	33	Done	–	–
	7	34	Start	Write/Read Sanity Check 2	18
	8	35	Done	–	–
DDR_CAL_STATUS_RANKx_4	0	36	Start	Write DQS-to-DQ (Complex)	19
	1	37	Done	–	–
	2	38	Start	Write DQS-to-DM (Complex)	20
	3	39	Done	–	–
	4	40	Start	Write/Read Sanity Check 3	21
	5	41	Done	–	–
	6	42	Start	Reserved	22
	7	43	Done	–	–
	8	44	Start	Write/Read Sanity Check 4	23
DDR_CAL_STATUS_RANKx_5	0	45	Done	–	–
	1	46	Start	Read Level Multi-Rank Adjustment	24
	2	47	Done	–	–
	3	48	Start	Write/Read Sanity Check 5 (For More than 1 Rank)	25
	4	49	Done	–	–
	5	50	Start	Multi-Rank Adjustments and Checks	26
	6	51	Done	–	–
	7	52	Start	Write/Read Sanity Check 6 (All Ranks)	27
	8	53	Done	–	–

Table 3-6 lists the post-calibration XSDB status signal descriptions.

Table 3-6: Post-Calibration XSDB Status Signal Description

XSDB Status Register	XSDB Bits[8:0]	Description	Post-Calibration Step
DDR_POST_CAL_STATUS	0	Running	DQS Gate Tracking
	1	Idle	
	2	Fail	
	3	Running	Read Margin Check (Reserved)
	4	Running	Write Margin Check (Reserved)
	5	–	Reserved
	6	–	Reserved
	7	–	Reserved
	8	–	Reserved

Table 3-7 lists the error signals and a description of each error. To decode the error first look at the status to determine which calibration stage failed (the start bit would be asserted, the associated done bit deasserted) then look at the error code provided. The error asserts the first time an error is encountered.

Table 3-7: Error Signal Descriptions

STAGE_NAME	Stage	Code	DDR_CAL_ERROR_1	DDR_CAL_ERROR_0	Error
DQS Gate	1	0x1	Byte	RIU Nibble	Calibration uses the calculated latency from the MR register as a starting point and then backs off and begins sampling. If the sample occurs too late in the DQS burst and there are no taps left to decrement for the latency, then an error has occurred.
		0x2	Byte	RIU Nibble	Expected pattern was not found on GT_STATUS.
		0x3	Byte	RIU Nibble	CAS latency is too low. Calibration starts at a CAS latency (CL) – 3. For allowable CAS latencies, see <a href="#">EXTRA_CMD_DELAY Configuration Settings, page 140</a> .
		0x4	Byte	RIU Nibble	Pattern not found on GT_STATUS, all 0s were sampled. Expecting to sample the preamble.
		0x5	Byte	RIU Nibble	Pattern not found on GT_STATUS, all 1s were sampled. Expecting to sample the preamble.
		0x6	Byte	RIU Nibble	Could not find the 0->1 transition with fine taps in at least 1 tck (estimated) of fine taps.
DQS Gate Sanity Check	2	0xF	N/A	N/A	PHY fails to return same number of data bursts as expected

Table 3-7: Error Signal Descriptions (Cont'd)

STAGE_NAME	Stage	Code	DDR_CAL_ERROR_1	DDR_CAL_ERROR_0	Error
Write leveling	3	0x1	Byte	N/A	Cannot find stable 0.
		0x2	Byte	N/A	Cannot find stable 1.
		0x3	Byte	N/A	Cannot find the left edge of noise region with fine taps.
		0x4	Byte	N/A	Could not find the 0->1 transition with fine taps in at least 1 tck (estimated) of ODELAY taps.
Read per-bit Deskew	4	0x1	Nibble	Bit	No valid data found for a given bit in the nibble when running the deskew pattern.
		0xF	Nibble	Bit	Timeout error waiting for read data bursts to return.
Read DQS Centering	6	0x1	Nibble	Bit	No valid data found for a given bit in the nibble.
		0x2	Nibble	Bit	Could not find the left edge of the data valid window to determine window size. All samples returned valid data.
		0xF	Nibble	Bit	Timeout error waiting for read data to return.
Read Sanity Check	7	0x1	Nibble	0	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for read data to return.
Write DQS-to-DQ Deskew	8	0x1	Byte	Bit	DQS deskew error. No valid data found; therefore, ran out of taps during search.
		0x2	Byte	Bit	DQ deskew error. Failure point not found.
		0xF	Byte	Bit	Timeout error waiting for all read data bursts to return.
Write DQS-to-DM/DBI Deskew	9	0x1	Byte	Bit	DQS deskew error. No valid data found; therefore, ran out of taps during search.
		0x2	Byte	Bit	DM/DBI deskew error. Failure point not found.
		0xF	Byte	Bit	Timeout error waiting for all read data bursts to return.
Write DQS-to-DQ (Simple)	10	0x1	Byte	N/A	No valid data found; therefore, ran out of taps during search.
		0xF	Byte	N/A	Timeout error waiting for read data to return.
Write DQS-to-DM (Simple)	11	0x1	Byte	N/A	No valid data found; therefore, ran out of taps during search.
		0xF	Byte	N/A	Timeout error waiting for all read data bursts to return.

Table 3-7: Error Signal Descriptions (Cont'd)

STAGE_NAME	Stage	Code	DDR_CAL_ERROR_1	DDR_CAL_ERROR_0	Error
Write Latency Calibration	13	0x1	Byte	N/A	Could not find the data pattern within the allotted number of taps.
		0x2	Byte	N/A	Data pattern not found. Data late at the start, instead of "F0A55A96," found "00F0A55A."
		0x3	Byte	N/A	Data pattern not found. Data too early, not enough movement to find pattern. Found pattern of "A55A96FF," "5A96FFFF," or "96FFFFFF."
		0x4	Byte	N/A	Data pattern not found. Multiple reads to the same address resulted in a read mismatch.
		0xF	Byte	N/A	Timeout error waiting for read data to return.
Write Read Sanity Check	14	0x1	Nibble	0	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for read data to return.
Read_Leveling (Complex)	15	See Read DQS Centering error codes.			
Write Read Sanity Check	16	0x1	Nibble	N/A	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for all read data bursts to return.
Read V <sub>REF</sub> Training	17	0x1	Byte	N/A	No valid window found for any V <sub>REF</sub> value.
		0xF	Nibble	N/A	Timeout error waiting for read data to return.
Write Read Sanity Check	18	0x1	Nibble	0	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for read data to return.
Write DQS-to-DQ (Complex)	19	See Write DQS-to-DQ (Simple) error codes.			
Write Read Sanity Check	21	0x1	Nibble	N/A	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for all read data bursts to return.
Write V <sub>REF</sub> Training	22	0x1	Byte	N/A	No valid window found for any V <sub>REF</sub> value.
		0xF	Byte	N/A	Timeout error waiting for read data to return.
Write Read Sanity Check	23	0x1	Nibble	N/A	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for all read data bursts to return.
Write Read Sanity Check	25	0x1	Nibble	0	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for read data to return.

Table 3-7: Error Signal Descriptions (Cont'd)

STAGE_NAME	Stage	Code	DDR_CAL_ERROR_1	DDR_CAL_ERROR_0	Error
Multi-Rank Adjust and Checks	26	0x1	Byte	RIU Nibble	Could not find common setting across ranks for general interconnect read latency setting for given byte. Variance between ranks could not be compensated with coarse taps.
		0x2	Byte	RIU Nibble	DQS Gate skew between ranks for a given byte larger than 360°.
		0x3	Byte	RIU Nibble	Write skew between ranks for a given byte larger than 180°. Check Write Latency Coarse settings.
		0x4	Byte	N/A	Could not decrement coarse taps enough to limit coarse tap setting for all ranks.
		0x5	Byte	N/A	Violation of maximum read latency limit.
Write Read Sanity Check	27	0x1	Nibble	RIU Nibble	Read data comparison failure.
		0xF	N/A	N/A	Timeout error waiting for read data to return.
DQS Gate Tracking		0x1	Byte	Rank	Underflow of the coarse taps used for tracking.
		0x2	Byte	Rank	Overflow of the coarse taps used for tracking.

## DQS Gate

The XIPHY is used to capture read data from the DRAM by using the DQS strobe to clock in read data and transfer the data to an internal FIFO using that strobe. The first step in capturing data is to evaluate where that strobe is so the XIPHY can open the gate and allow the DQS to clock the data into the rest of the PHY.

The XIPHY uses an internal clock to sample the DQS during a read burst and provides a single binary value back called GT\_STATUS. This sample is used as part of a training algorithm to determine where the first rising edge of the DQS is in relation to the sampling clock.

Calibration logic issues individual read commands to the DRAM and asserts the `clb2phy_rd_en` signal to the XIPHY to open the gate which allows the sample of the DQS to occur. The `clb2phy_rd_en` signal has control over the timing of the gate opening on a DRAM-clock-cycle resolution (`DQS_GATE_READ_LATENCY_RANK#_BYTE#`). This signal is controlled on a per-byte basis in the PHY and is set in the `ddr_mc_pi` block for use by both calibration and the controller.

Calibration is responsible for determining the value used on a per-byte basis for use by the controller. The XIPHY provides for additional granularity in the time to open the gate through coarse and fine taps. Coarse taps offer 90° DRAM clock-cycle granularity (16 available) and each fine tap provides a 2.5 to 15 ps granularity for each tap (512 available). BISC provides the number of taps for 1/4 of a DDR clock cycle by taking

(BISC\_PQTR\_NIBBLE#-BISC\_ALIGN\_PQTR\_NIBBLE#) or (BISC\_NQTR\_NIBBLE#-BISC\_ALIGN\_NQTR\_NIBBLE#). These are used to estimate the per-tap resolution for a given nibble.

The search for the DQS begins with an estimate of when the DQS is expected back. The total latency for the read is a function of the delay through the PHY, PCB delay, and the configured latency of the DRAM (CAS latency, Additive latency, etc.). The search starts three DRAM clock cycles before the expected return of the DQS. The algorithm must start sampling before the first rising edge of the DQS, preferably in the preamble region. DDR3 and DDR4 have different preambles for the DQS as shown in Figure 3-7.




Figure 3-7: DDR3/DDR4 DQS Preamble

The specification for the DDR3 preamble is longer (3/4 of a DRAM clock cycle) and starts from the terminated 3-state while the DDR4 preamble is shorter (1/2 of a DRAM clock cycle) and starts from the rail terminated level. For DDR4, the preamble training mode is enabled during DQS gate calibration, so the DQS is driven low whenever the DQS is idle. This allows for the algorithm to look for the same sample pattern on the DQS for DDR3/DDR4 where the preamble is larger than half a clock cycle for both cases.

Given that DDR3 starts in the 3-state region before the burst, any accepted sample taken can either be a 0 or 1. To avoid this result, 20 samples (in hardware) are taken for each individual sample such that the probability of the 3-state region or noise in the sampling clock/strobe being mistaken for the actual DQS is low. This probability is given by the binomial probability shown in the binomial probability equation.


X = expected outcome

n= number of tries

P = probability of a single outcome

$$P(X = x) = \frac{n!}{x!(n-x)!} p^x (1-p)^{n-x}$$

When sampling in the 3-state region the result can be 0 or 1, so the probability of 20 samples all arriving at the same value is roughly  $9.5 \times 10^{-6}$ . Figure 3-8 shows an example of samples of a DQS burst with the expected sampling pattern to be found as the coarse taps are adjusted. The pattern is the expected level seen on the DQS over time as the sampling clock is adjusted in relation to the DQS.



**Figure 3-8: Example DQS Gate Samples Using Coarse Taps**

Each individual element of the pattern is 20 read bursts from the DRAM and samples from the XIPHY. The gate in the XIPHY is opened and a new sample is taken to indicate the level seen on the DQS. If each of the samples matches with the first sample taken, the value is accepted. If all samples are not the same value that value is marked as "X" in the pattern. The "X" in the pattern shown is to allow for jitter and DCD between the clocks, and to deal with uncertainty when dealing with clocks with an unknown alignment. Depending on how the clocks line up they can resolve to all 0s, all 1s, or a mix of values, and yet the DQS pattern can still be found properly.

The coarse taps in the XIPHY are incremented and the value recorded at each individual coarse tap location, looking for the full pattern "00X1X0X1X0." For the algorithm to incorrectly calculate the 3-state region as the actual DQS pattern, you would have to take 20 samples of all 0s at a given coarse tap, another 20 samples of all 0s at another, then 20 coarse taps of all 1s for the initial pattern ("00X1"). The probability of this occurring is  $8.67 \times 10^{-19}$ . This also only covers the initial scan and does not include the full pattern which scans over 10 coarse taps.

While the probability is fairly low, there is a chance of coupling or noise being mistaken as a DQS pattern. In this case, each sample is no longer random but a signal that can be fairly repeatable. To guard against mistaking the 3-state region in DDR3 systems with the actual DQS pulse, an extra step is taken to read data from the MPR register to validate the gate alignment. The read path is set up by BISC for capture of data, placing the capture clock roughly in the middle of the expected bit time back from the DRAM.

Because the algorithm is looking for a set pattern and does not know the exact alignment of the DQS with the clock used for sampling the data, there are four possible patterns, as shown in [Figure 3-9](#).






Figure 3-9: DQS Gate Calibration Possible Patterns

To speed up the pattern search, only the initial seven coarse taps are used to determine if the starting pattern is found. This eliminates the need to search additional coarse taps if the early samples do not match the expected result. If the result over the first coarse seven coarse taps is not one of the four shown in Figure 3-9, the following occurs:

- Coarse taps are reset to 0
- `clb2phy_rd_en` general interconnect control is adjusted to increase by one DRAM clock cycle
- Search starts again (this is the equivalent of starting at coarse tap four in Figure 3-9)

For DDR4, if the algorithm samples 1XX or 01X this means it started the sampling too late in relation to the DQS burst. The algorithm decreases the `clb2phy_rd_en` general interconnect control and try again. If the `clb2phy_rd_en` is at the low limit already it issues an error.

If all allowable values of `clb2phy_rd_en` for a given latency are checked and the expected pattern is still not found, the search begins again from the start but this time the sampling is offset by an estimated 45° using fine taps (half a coarse tap). This allows the sampling to occur at a different phase than the initial relationship. Each time through if the pattern is not found, the offset is reduced by half until all offset values have been exhausted.

Figure 3-10 shows an extreme case of DCD on the DQS that would result in the pattern not being found until an offset being applied using fine taps.




Figure 3-10: DQS Gate Calibration Fine Offset Example

After the pattern has been found, the final coarse tap (DQS\_GATE\_COARSE\_RANK#\_BYTE#) is set based on the alignment of the pattern previously checked (shown in Figure 3-9). The coarse tap is set to be the last 0 seen before the 1 (3 is used to indicate an unstable region, where multiple samples return 0 and 1) was found in the pattern shown in Figure 3-11. During this step, the final value of the coarse tap is set between five and nine.




Figure 3-11: DQS Gate Coarse Setting Before Fine Search

From this point the `clb2phy_rd_en` (`DQS_GATE_READ_LATENCY_RANK#_BYTE#`) is increased by 1 to position the gate in the final location before the start of the fine sweep. This is done to ensure the proper timing of the gate in relation to the full DQS burst during normal operation. Because this is sampling the strobe with another signal it can have jitter in relation to one another.

For example, when they are lined up taking multiple samples it might give you a different result each time as a new sample is taken. The fine search begins in an area where all samples returned a 0 so it is relatively stable, as shown in Figure 3-12. The fine taps are incremented until a non-zero value is returned (which indicates the left edge of the unstable region) and that value recorded as shown in Figure 3-14 (`DQS_GATE_FINE_LEFT_RANK#_BYTE#`).




Figure 3-12: DQS Gate Fine Adjustment, Sample a 0

The fine taps are then incremented until all samples taken return a 1, as shown in Figure 3-13. This is recorded as the right edge of the uncertain region as shown in Figure 3-14 (`DQS_GATE_FINE_RIGHT_RANK#_BYTE#`).





Figure 3-13: DQS Gate Fine Adjustment, Sample a 1



**Figure 3-14: DQS Gate Fine Adjustment, Uncertain Region**

The final fine tap is computed as the midpoint of the uncertain region,  $(\text{right} - \text{left})/2 + \text{left}$  (`DQS_GATE_FINE_CENTER_RANK#_BYTE#`). This ensures optimal placement of the gate in relation to the DQS. For simulation, speeding up a faster search is implemented for the fine tap adjustment. This is performed by using a binary search to jump the fine taps by larger values to quickly find the 0 to 1 transition.

For multi-rank systems, separate control exists in the XIPHY for each rank and every rank can be trained separately for coarse and fine taps. After calibration is complete, adjustments are made so that for each byte, the `clb2phy_rd_en` (`DQS_GATE_READ_LATENCY_RANK#_BYTE#`) value for a given byte matches across all ranks. The coarse taps are incremented/decremented accordingly to adjust the timing of the gate signal to match the timing found in calibration. If a common `clb2phy_rd_en` setting cannot be found for a given byte across all ranks, an error is asserted.

### DQS Gate Sanity Check

After completion of DQS gate calibration for all bytes in a given rank, read return timing is calculated and 10 read bursts with gaps between them are issued. Logic then checks that the FIFO is read 10 times. There is no data checking at this stage. This is just a basic functional check of the FIFO read port control logic, which is configured using the DQS gate calibration results. Read return timing is updated after DQS gate calibration for each rank. The final setting is determined by largest DQS gate delay out of all DQS lanes and all ranks.

## Write Leveling

The DDR3/DDR4 SDRAM memory modules use a fly-by topology on clocks, address, commands, and control signals to improve signal integrity. This topology causes a skew between DQS and CK at each memory device on the module. Write leveling is a feature in DDR3/DDR4 SDRAMs that allows the controller to adjust each write DQS phase independently with respect to the clock (CK) forwarded to the DDR3/DDR4 device to compensate for this skew and meet the `tDQSS` specification [Ref 1].

During write leveling, DQS is driven by the FPGA memory interface and DQ is driven by the DDR3/DDR4 SDRAM device to provide feedback. To start write leveling, an MRS command is sent to the DRAM to enable the feedback feature, while another MRS command is sent to disable write leveling at the end. Figure 3-15 shows the block diagram for the write leveling implementation.




Figure 3-15: Write Leveling Block Diagram

The XIPHY is set up for write leveling by setting various attributes in the RIU. WL\_TRAIN is set to decouple the DQS and DQ when driving out the DQS. This allows the XIPHY to capture the returning DQ from the DRAM. Because the DQ is returned without the returning DQS strobe for capture, the RX\_GATE is set to 0 in the XIPHY to disable DQS gate operation. While the write leveling algorithm acts on a single DQS at a time, all the XIPHY bytes are set up for write leveling to ensure there is no contention on the bus for the DQ.

DQS is delayed with ODELAY and coarse delay (WL\_DLY\_CRSE[12:9] applies to all bits in a nibble) provided in the RIU WL\_DLY\_RNKx register. The WL\_DLY\_FINE[8:0] location in the RIU is used to store the ODELAY value for write leveling for a given nibble (used by the XIPHY when switching ranks).

A DQS train of pulses is output by the FPGA to the DRAM to detect the relationship of CK and DQS at the DDR3/DDR4 memory device. DQS is delayed using the ODELAY and coarse taps in unit tap increments until a "0" to "1" transition is detected on the feedback DQ input. A single typical burst length of eight pattern is first put out on the DQS (four clock pulses), followed by a gap, and then 100 bursts length of eight patterns are sent to the DRAM (Figure 3-16).

The first part is to ensure the DRAM updates the feedback sample on the DQ being sent back, while the second provides a clock that is used by the XIPHY to clock into the XIPHY the level seen on the DQ. Sampling the DQ while driving the DQS helps to avoid ringing on the DQS at the end of a burst that can be mistaken as a clock edge by the DRAM.




Figure 3-16: Write Leveling DQS Bursts

To avoid false edge detection around the CK negative edge due to jitter, the DQS delays the entire window to find the large stable "0" and "1" region (Stable 0 or 1 indicates all samples taken return the same value). Check that you are to the left of this stable "1" region as the right side of this region is the CK negative edge being captured with the DQS, as shown in Figure 3-17.




Figure 3-17: Write Leveling Regions

Write leveling is performed in the following two steps:

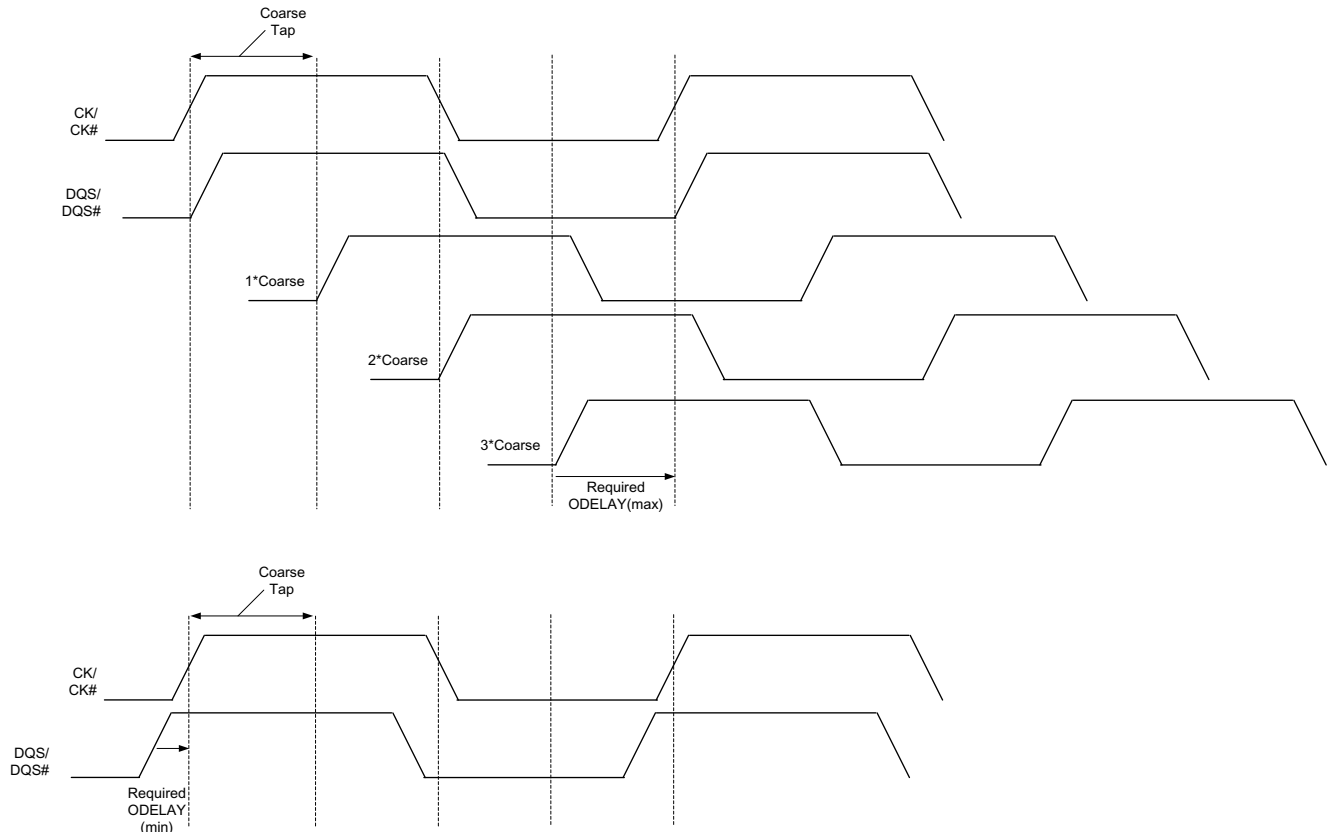
1. Find the transition from "0" to "1" using coarse taps and ODELAY taps (if needed).

During the first step, look for a static "0" to be returned from all samples taken. This means 64 samples were taken and it is certain the data is a "0." Record the coarse tap setting and keep incrementing the coarse tap.

- If the algorithm receives another stable "0," update the setting (WRLVL\_COARSE\_STABLE0\_RANK\_BYTE) and continue.

- If the algorithm receives a non-zero result (noise) or a stable "1" reading (WRLVL\_COARSE\_STABLE0\_RANK\_BYTE), the search has gone too far and the delay is backed up to the last coarse setting that gave a stable "0." This reference allows you to know the algorithm placed the coarse taps to the left of the transition desired.
  - If the algorithm never sees a transition from a stable "0" to the noise or stable "1" using the coarse taps, the ODELAY of the DQS is set to an offset value (first set at 45°, WRLVL\_ODELAY\_INITIAL\_OFFSET\_BYTE) and the coarse taps are checked again from "0." Check for the stable "0" to stable "1" transition (the algorithm might need to perform this if the noise region is close to 90° or there is a large amount of DCD).
  - If the transition is still not found, the offset is halved and the algorithm tries again. The final offset value used is stored at WRLVL\_ODELAY\_LAST\_OFFSET\_RANK\_BYTE. Because the algorithm is aligning the DQS with the nearest clock edge the coarse tap sweep is limited to five, which is 1.25 clock cycles. The final coarse setting is stored at WRLVL\_COARSE\_STABLE0\_RANK\_BYTE.
2. Find the center of the noise region around that transition from "0" to 1" using ODELAY taps.

The second step is to sweep with ODELAY taps and find both edges of the noise region (WRLVL\_ODELAY\_STABLE0\_RANK\_BYTE, WRLVL\_ODELAY\_STABLE1\_RANK\_BYTE while WRLVL\_ODELAY\_CENTER\_RANK\_BYTE holds the final value). The number of ODELAY taps used is determined by the initial alignment of the DQS and CK and the size of this noise region as shown in [Figure 3-18](#).



**Figure 3-18: Worst Case ODELAY Taps (Maximum and Minimum)**

After the final ODELAY setting is found, the value of ODELAY is loaded in the RIU in the WL\_DLY\_RNKx[8:0] register. This value is also loaded in the ODELAY register for the DQ and the DM to match the DQS. If any deskew has been performed on the DQS/DQ/DM when reaching this point (multi-rank systems), the deskew information is preserved and the offset is applied.

The lowest ODELAY value is stored at WRLVL\_ODELAY\_LOWEST\_COMMON\_BYTE, which is used to preserve the WRLVL element with the deskew portion of ODELAY for a given byte. During normal operation in a multi-rank system, the XIPHY is responsible for loading the ODELAY with the value stored for a given rank.

After write leveling, the MPR command is sent to the DRAM to disable the write leveling feature, the WL\_TRAIN is set back to the default "off" setting, and the DQS gate is turned back on to allow for capture of the DQ with the returning strobe DQS.



## Read DQS Deskew and Centering

After the gate has been trained and Write Leveling has completed, the next step is to ensure reliable capture of the read data with the DQS. This stage of Read Leveling is divided into two phases, Per-Bit Deskew and Read DQS Centering. Read DQS Centering utilizes the DDR3 and DDR4 Multi Purpose Register (MPR). The MPR contains a pattern that can be used to train the read DQS and DQ for read capture. While DDR4 allows for several patterns, DDR3 only has a single repeating pattern available.

To perform per-bit deskew, a non-repeating pattern is useful to deal with or diagnose cases of extreme skew between different bits in a byte. Because this is limited by the DDR3 MPR pattern, a long pattern is first written to the DRAM and then read back to perform per-bit deskew (only done on the first rank of a multi-rank system). When per-bit deskew is complete, the simple repeating pattern available through both DDR3 and DDR4 MPR is used to center the DQS in the DQ read eye.

The XIPHY provides separate delay elements (2.5 to 15 ps per tap, 512 total) for the DQS to clock the rising and falling edge DQ data (PQTR for rising edge, NQTR for falling edge) on a per-nibble basis (four DQ bits per PQTR/NQTR). This allows the algorithm to center the rising and falling edge DQS strobe independently to ensure more margin when dealing with DCD. The data captured in the PQTR clock domain is transferred to the NQTR clock domain before being sent to the read FIFO and to the general interconnect clock domain.

Due to this transfer of clock domains, the PQTR and NQTR clocks must be roughly 180° out of phase. This relationship between the PQTR/NQTR clock paths is set up as part of the BISC start-up routine, and thus calibration needs to maintain this relationship as part of the training (BISC\_ALIGN\_PQTR, BISC\_ALIGN\_NQTR, BISC\_PQTR, BISC\_NQTR).

### Read Per-bit Deskew

First, write 0x00 to address 0x000. Because the write latency calibration has not yet been performed, the address DQ is held for eight clock cycles before and after the expected write latency is expected. The DQS toggles extra time before/after is shown in [Figure 3-19](#). This ensures the data is written to the DRAM if the burst does not occur at the correct time the DRAM expects it.




Figure 3-19: Per-bit Deskew – Write 0x00 to Address 0x000

Next, write 0xFF to a different address to allow for back-to-back reads (Figure 3-20). For DDR3 address 0x008 is used, while for DDR4 address 0x000 and bank group 0x1 is used. At higher frequencies, DDR4 requires a change in the bank group to allow for back-to-back bursts of eight.




Figure 3-20: Per-bit Deskew – Write 0xFF to Other Address

After the data is written, back-to-back reads are issued to the DRAM to perform per-bit deskew (Figure 3-21).




Figure 3-21: Per-bit Deskew – Back-to-Back Reads (No Gaps)

Using this pattern each bit in a byte is left edge aligned with the DQS strobe (PQTR/NQTR). More than a bit time of skew can be seen and corrected as well.



**RECOMMENDED:** In general, a bit time of skew between bits is not ideal. Ensure the DDR3/DDR4 trace matching guidelines within DQS byte are met. See [PCB Guidelines for DDR3, page 77](#) and [PCB Guidelines for DDR4, page 77](#).

At the start of deskew, the PQTR/NQTR are decreased down together until one of them hits 0 (to preserve the initial relationship setup by BISC). Next, the data for a given bit is checked for the matching pattern. Only the rising edge data is checked for correctness. The falling edge comparison is thrown away to allow for extra delay on the PQTR/NQTR relative to the DQ.

While in the ideal case, the PQTR/NQTR are edge aligned with the DQ when the delays are set to 0. Due to extra delay in the PQTR/NQTR path, the NQTR might be pushed into the next burst transaction at higher frequencies and so it is excluded from the comparison (Figure 3-22 through Figure 3-23). More of the rising edge data of a given burst would need to be discarded to deal with more than a bit time of skew. If the last part of the burst was not excluded, the failure would cause the PQTR/NQTR to be pushed instead of the DQ IDELAY.




Figure 3-22: Per-bit Deskew – Delays Set to 0 (Ideal)




Figure 3-23: Per-bit Deskew – Delays Set to 0

If the pattern is found, the given IDELAY on that bit is incremented by 1, then checked again. If the pattern is not seen, the PQTR/NQTR are incremented by 1 and the data checked again. The algorithm checks for the passing and failing region for a given bit, adjusting either the PQTR/NQTR delays or the IDELAY for that bit.

To guard against noise in the uncertain region, the passing region is defined by a minimum window size (10), hence the passing region is not declared as found unless the PQTR/NQTR are incremented and a contiguous region of passing data is found for a given nibble. All of the bits are cycled through to push the PQTR/NQTR out to align with the latest bit in a given nibble. Figure 3-24 through Figure 3-26 show an example of the PQTR/NQTR and various bits being aligned during the deskew stage.




Figure 3-24: Per-bit Deskew – Initial Relationship Example

The algorithm takes the result of each bit at a time and decides based on the results of that bit only. The common PQTR/NQTR are delayed as needed to align with each bit, but is not decremented. This ensures it gets pushed out to the latest bit.


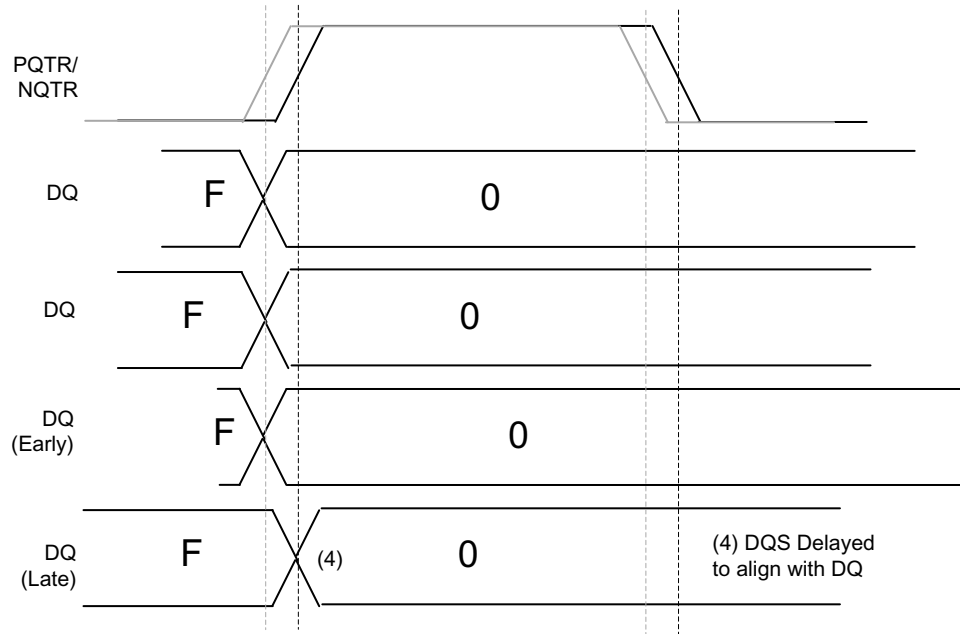
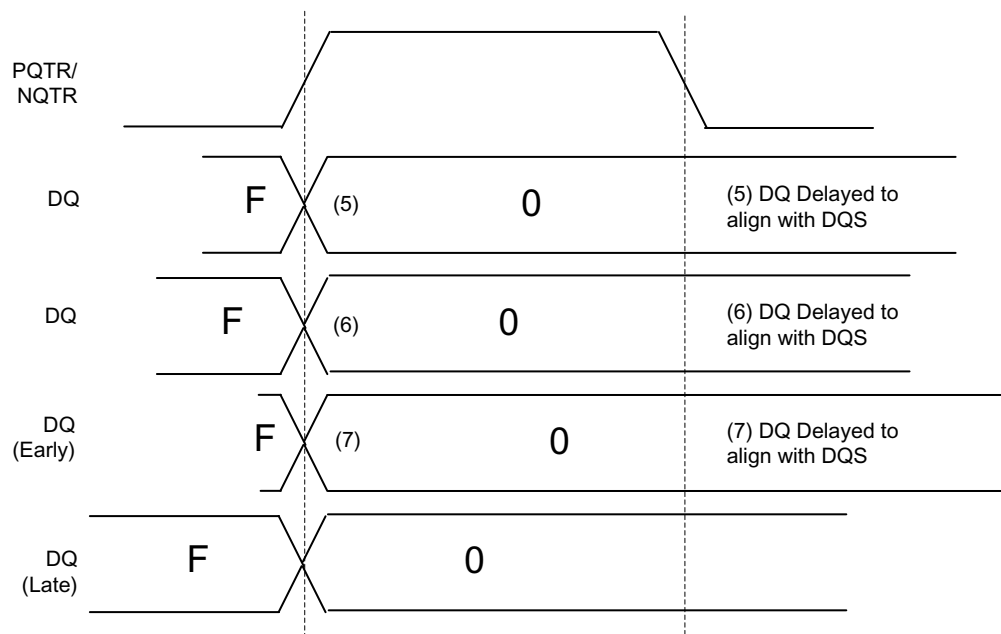


Figure 3-25: Per-bit Deskew – Early Bits Pushed Out



**Figure 3-26: Per-bit Deskew – PQTR/NQTR Delayed to Align with Late Bit**

When completed, the PQTR/NQTR are pushed out to align with the latest DQ bit (RDLVL\_DESKEW\_PQTR\_nibble, RDLVL\_DESKEW\_NQTR\_nibble), but DQ bits calibrated first might have been early as shown in the example. Accordingly, all bits are checked once again and aligned as needed (Figure 3-27).



**Figure 3-27: Per-bit Deskew – Push Early Bits as Needed to Align**

The final DQ IDELAY value from deskew is stored at RDLVL\_DESKEW\_IDELAY\_Byte\_Bit.

## Read DQS Centering

When the data is deskewed, the PQTR/NQTR delays need to be adjusted to center in the aggregate data valid window for a given nibble. The DRAM MPR register is used to provide the data pattern for centering. Therefore, the pattern changes each bit time and does not rely on being written into the DRAM first, eliminating some uncertainty. The simple clock pattern is used to allow for the same pattern checking for DDR3 and DDR4. Gaps in the reads to the DRAM are used to stress the initial centering to incorporate the effects of ISI on the first DQS pulse as shown in Figure 3-28.





Figure 3-28: Gap between MPR Reads

To properly account for jitter on the data and clock returned from the DRAM, multiple data samples are taken at a given tap value. 64 read bursts are used in hardware while five are used in simulation. More samples mean finding the best alignment in the data valid window.

Given that the PHY has two capture strobes PQTR/NQTR that need to be centered independently yet moved together, calibration needs to take special care to ensure the clocks stay in a certain phase relationship with one another.


The data and PQTR/NQTR delays start with the value found during deskew. Data is first delayed with IDELAY such that both the PQTR and NQTR clocks start out just to the left of the data valid window for all bits in a given nibble so the entire read window can be scanned with each clock (Figure 3-29, RDLVL\_IDELAY\_VALUE\_Rank\_Byte\_Bit). Scanning the window with the same delay element and computing the center with that delay element helps to minimize uncertainty in tap resolution that might arise from using different delay lines to find the edges of the read window.



**Figure 3-29: Delay DQ Thus PQTR and NQTR in Failing Region**

At the start of training, the PQTR/NQTR and data are roughly edge aligned, but because the pattern is different from the deskew step the edge might have changed a bit. Also, during deskew the aggregate edge for both PQTR/NQTR is found while you want to find a separate edge for each clock.

After making sure both PQTR/NQTR start outside the data valid region, the clocks are incremented to look for the passing region (Figure 3-30). Rising edge data is checked for PQTR while falling edge data is checked for NQTR, with a separate check being kept to indicate where the passing region/falling region is for each clock.



**Figure 3-30: PQTR and NQTR Delayed to Find Passing Region (Left Edge)**

When searching for the edge, a minimum window size of 10 is used to guarantee the noise region has been cleared and the true edge is found. The PQTR/NQTR delays are increased past the initial passing point until the minimum window size is found before the left edge is declared as found. If the minimum window is not located across the entire tap range for either clock, an error is asserted.

After the left edge is found (RDLVL\_PQTR\_LEFT\_Rank\_Nibble, RDLVL\_NQTR\_LEFT\_Rank\_Nibble), the right edge of the data valid window can be searched starting from the left edge + minimum window size. A minimum window size is not used when searching for the right edge, as the starting point already guarantees a minimum window size has been met.

Again, the PQTR/NQTR delays are incremented together and checked for error independently to keep track of the right edge of the window. Because the data from the PQTR domain is transferred into the NQTR clock domain in the XIPHY, the edge for NQTR is checked first, keeping track of the results for PQTR along the way (Figure 3-31).

When the NQTR edge is located, a flag is checked to see if the PQTR edge is found as well. If the PQTR edge was not found, the PQTR delay continues to search for the edge, while the NQTR delay stays at its right edge (RDLVL\_PQTR\_RIGHT\_Rank\_Nibble, RDLVL\_NQTR\_RIGHT\_Rank\_Nibble). For simulation, the right edge detection is sped up by having the delays adjusted by larger than one tap at a time.




Figure 3-31: PQTR and NQTR Delayed to Find Failing Region (Right Edge)

After both rising and falling edge windows are found, the final center point is calculated based on the left and right edges for each clock. The final delay for each clock (RDLVL\_PQTR\_CENTER\_Rank\_Nibble, RDLVL\_NQTR\_CENTER\_Rank\_Nibble) is computed by:

$$\text{left} + ((\text{right} - \text{left})/2).$$

For multi-rank systems deskew only runs on the first rank, while read DQS centering using the PQTR/NQTR runs on all ranks. After calibration is complete for all ranks, for a given DQ bit the IDELAY is set to the center of the range of values seen for all ranks (RDLVL\_IDELAY\_FINAL\_BYTE\_BIT). The PQTR/NQTR final value is also computed based on the range of values seen between all of the ranks (RDLVL\_PQTR\_CENTER\_FINAL\_NIBBLE, RDLVL\_NQTR\_CENTER\_FINAL\_NIBBLE).






---

**IMPORTANT:** For multi-rank systems, there must be overlap in the read window computation. Also, there is a limit in the allowed skew between ranks, see the [PCB Guidelines for DDR3 in Chapter 4](#) and [PCB Guidelines for DDR4 in Chapter 4](#).

---

## Read Sanity Check

After read DQS centering but before Write DQS-to-DQ, a check of the data is made to ensure the previous stage of calibration did not inadvertently leave the alignment of the read path in a bad spot. A single MPR read command is sent to the DRAM, and the data is checked against the expected data across all bytes before continuing.

## Write DQS-to-DQ

**Note:** The calibration step is only enabled for the first rank in a multi-rank system.

The DRAM requires the write DQS to be center-aligned with the DQ to ensure maximum write margin. Initially the write DQS is set to be roughly 90° out of phase with the DQ using the XIPHY TX\_DATA\_PHASE set for the DQS. The TX\_DATA\_PHASE is an optional per-bit adjustment that uses a fast internal XIPHY clock to generate a 90° offset between bits. The DQS and DQ ODELAY are used to fine tune the 90° phase alignment to ensure maximum margin at the DRAM.

A simple clock pattern of “10101010” is used initially because the write latency has not yet been determined. Due to fly-by routing on the PCB/DIMM module, the command to data timing is unknown until the next stage of calibration. Just as in read per-bit deskew when issuing a write to the DRAM, the DQS and DQ toggles for eight clock cycles before and after the expected write latency. This is used to ensure the data is written into the DRAM even if the command-to-write data relationship is still unknown. Write DQS-to-DQ is completed in two stages, per-bit deskew and DQS centering.

### Write DQS-to-DQ Per-bit Deskew

Initially all DQ bits have the same ODELAY setting based on the write leveling results, but the ODELAY for each bit might need to be adjusted to account for skew between bits.

Figure 3-32 shows an example of the initial timing relationship between a write DQS and DQ.




Figure 3-32: Initial Write DQS and DQ with Skew between Bits

1. Set TX\_DATA\_PHASE to 1 for DQ to add the 90° shift on the DQS relative to the DQ for a given byte (Figure 3-33). The data read back on some DQ bits are "10101010" while other DQ bits might be "01010101."


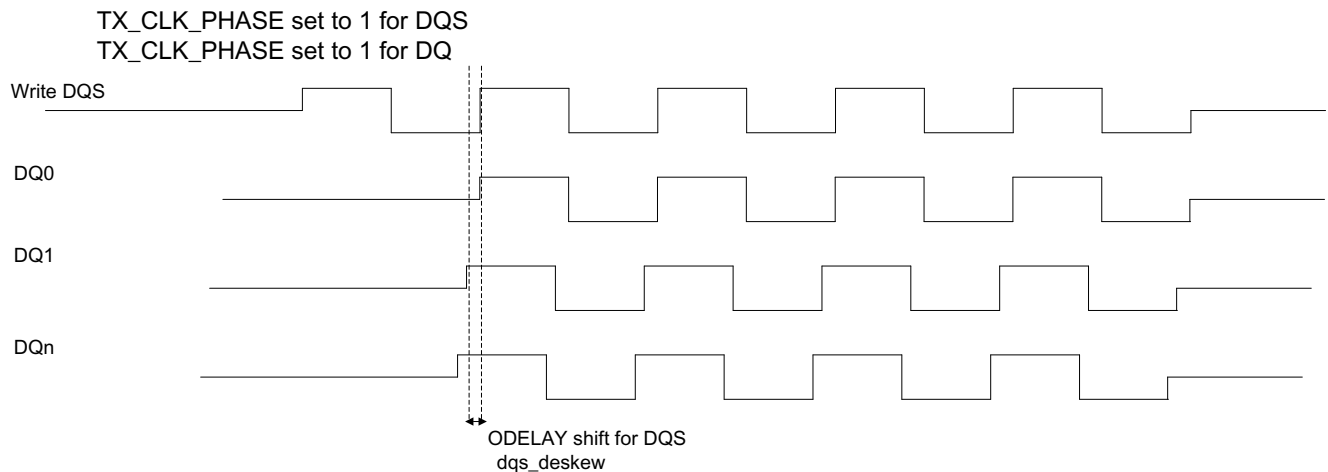


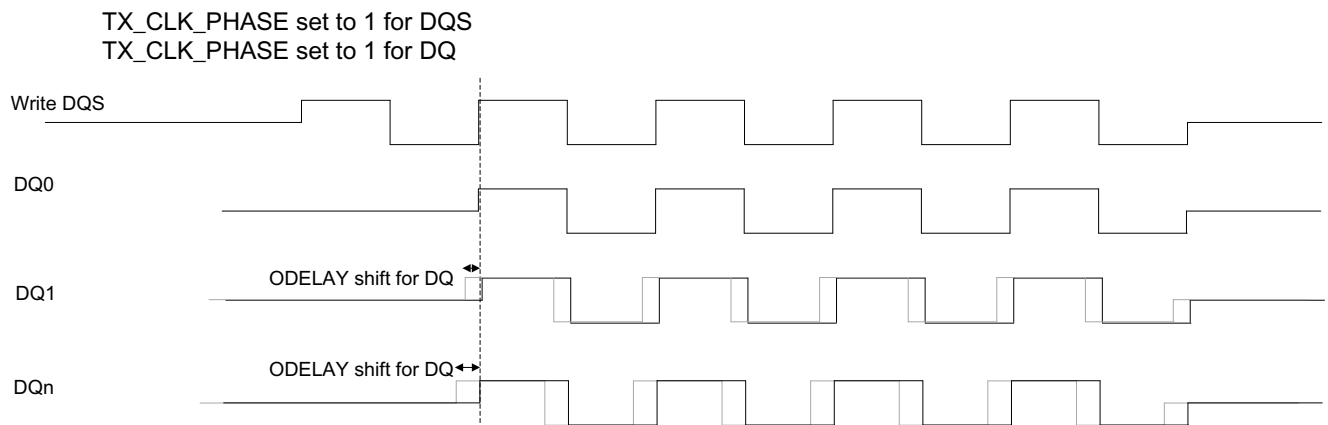
Figure 3-33: Add 90° Shift on DQ

2. If all the data for the byte does not match the expected data pattern, increment DQS ODELAY one tap at a time until the expected data pattern is found on all bits and save the delay as WRITE\_DQS\_TO\_DQ\_DESKEW\_DELAY\_Byte (Figure 3-34). As the DQS ODELAY is incremented, it moves away from the edge alignment with the CK. The deskew data is the inner edge of the data valid window for writes.



**Figure 3-34: Increment Write DQS ODELAY until All Bits Captured with Correct Pattern**

- Increment each DQ ODELAY until each bit fails to return the expected data pattern (the data is edge aligned with the write DQS, [Figure 3-35](#)).



**Figure 3-35: Per-bit Write Deskew**

- Return the DQ to the original position at the 0° shift using the TX\_DATA\_PHASE. Set DQS ODELAY back to starting value ([Figure 3-36](#)).




Figure 3-36: DQ Returned to Approximate 90° Offset with DQS

### Write DQS-to-DQ Centering

After per-bit write deskew, the next step is to determine the relative center of the DQS in the write data eye and compensate for any error in the TX\_DATA\_PHASE 90° offset.

1. Issue a set of write and read bursts with the data pattern "10101010" and check the read data. Just as in read write per-bit deskew when issuing a write to the DRAM, the DQS and DQ toggles for eight clock cycles before and after the expected write latency. This is used to ensure the data is written into the DRAM even if the command-to-write data relationship is still unknown.
2. Increment DQ ODELAY taps together until the read data pattern on all DQ bits changes from the expected data pattern "10101010." The amount of delay required to find the failing point is saved as WRITE\_DQS\_TO\_DQ\_PRE\_ADJUST\_MARGIN\_LEFT\_BYTE as shown in Figure 3-37.




Figure 3-37: Write DQS Centering – Left Edge

3. Return DQ ODELAY taps to their original value.

4. Find the right edge of the window by incrementing the DQS ODELAY taps until the data changes from the expected data pattern "10101010." The amount of delay required to find the failing point is saved as WRITE\_DQS\_TO\_DQ\_PRE\_ADJUST\_MARGIN\_RIGHT\_BYTE as shown in Figure 3-38.

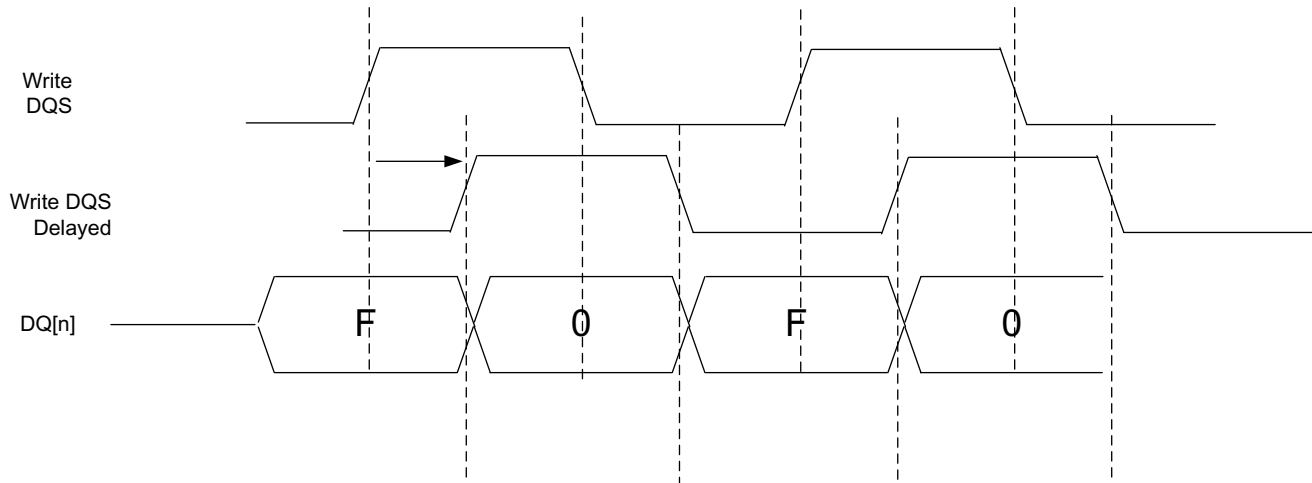


Figure 3-38: Write DQS Centering – Right Edge

5. Calculate the center tap location for the DQS ODELAY, based on skew, left and right edge.

$$\text{New DQS delay} = \text{deskew} - [(dly0 - dly1)/2]$$

Where dly0 is the original DQS delay + left margin and dly1 is the original DQS delay + right margin.

The final ODELAY tap setting for DQS is indicated by WRITE\_DQS\_TO\_DQ\_DQS\_ODELAY\_BYTE while the DQ is WRITE\_DQS\_TO\_DQ\_DQ\_ODELAY. The final computed left and right margin are WRITE\_DQS\_TO\_DQ\_MARGIN\_LEFT\_BYTE and WRITE\_DQS\_TO\_DQ\_MARGIN\_RIGHT\_BYTE.

## Write DQS-to-DM

**Note:** The calibration step is only enabled for the first rank in a multi-rank system.

During calibration the Data Mask (DM) signals are not used, they are deasserted during any writes before/after the required amount of time to ensure they have no impact on the pattern being written to the DRAM. If the DM signals are not used, this step of calibration is skipped.

Two patterns are used to calibrate the DM pin. The first pattern is written to the DRAM with the DM deasserted, ensuring the pattern is written to the DRAM properly. The second pattern overwrites the first pattern at the same address but with the DM asserted in a known position in the burst, as shown in Figure 3-39.

Because this stage takes place before Write Latency Calibration when issuing a write to the DRAM, the DQS and DQ/DM toggles for eight clock cycles before and after the expected write latency. This is used to ensure the data is written into the DRAM even though the command-to-write data relationship is still unknown.




Figure 3-39: DM Base Data Written

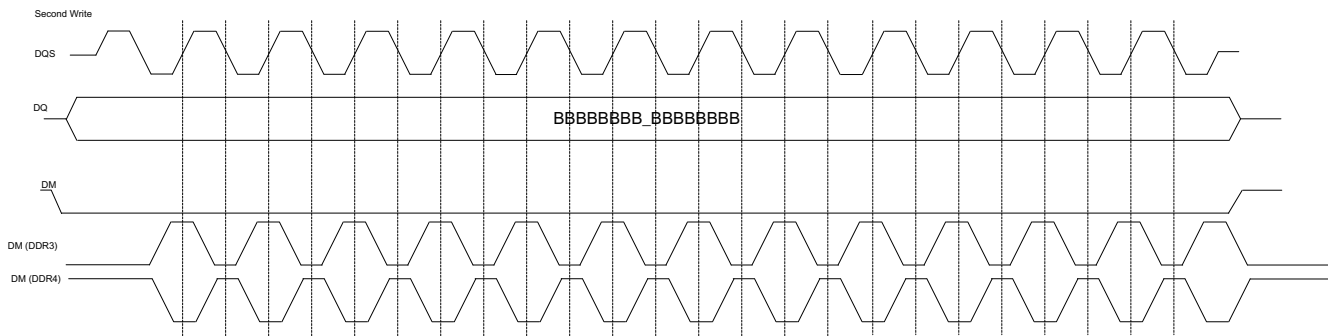


Figure 3-40: DM Asserted

The read back data for any given nibble is "5B5B\_5B5B," where the location of the "5" in the burst indicates where the DM is asserted. Because the data is constant during this step, the DQS-to-DQ alignment is not stressed. Only the DQS-to-DM is checked as the DQS and DM phase relationship is adjusted with each other.

### Write DQS-to-DM Per-Bit Deskew

This step is similar to Write DQS-to-DQ Per-Bit Deskew but involves the DM instead of the DQ bits. See [Write DQS-to-DQ, page 57](#) for an in-depth overview of the algorithm. The DQS ODELAY value used to edge align the DQS with the DM is stored as WRITE\_DQS\_TO\_DM\_DESKEW\_BYTE. The ODELAY value for the DM is stored as WRITE\_DQS\_TO\_DM\_DM\_ODELAY\_BYTE.

### Write DQS-to-DM Centering

This step is similar to Write DQS-to-DQ Centering but involves the DM instead of the DQ bits. See [Write DQS-to-DQ, page 57](#) for an in-depth overview of the algorithm. The tap value DM was set at to find the left edge is saved as WRITE\_DQS\_TO\_DM\_PRE\_ADJUST\_MARGIN\_LEFT\_BYTE. The tap value DQS was set at to find the right edge is saved as WRITE\_DQS\_TO\_DM\_PRE\_ADJUST\_MARGIN\_RIGHT\_BYTE.

The final DM margin is stored at `WRITE_DQS_TO_DM_MARGIN_LEFT_BYTE` and `WRITE_DQS_TO_DM_MARGIN_RIGHT_BYTE`.

Because the DQS ODELAY can only hold a single value, compute the aggregate smallest left/right margin between the DQ and DM. The DQS ODELAY value is set in the middle of this aggregate window. The final values of the DQS and DM can be found at `WRITE_DQS_ODELAY_FINAL` and `WRITE_DM_ODELAY_FINAL`.

## Write Latency Calibration

Write latency calibration is required to align the write DQS to the correct CK edge. During write leveling, the write DQS is aligned to the nearest rising edge of CK. However, this might not be the edge that captures the write command. Depending on the interface type (UDIMM, RDIMM, or component), the DQS could be up to three CK cycles later than, or aligned to the CK edge that captures the write command.

Write latency calibration makes use of the coarse tap in the `WL_DLY_RNK` of the XIPHY for adjusting the write latency on a per byte basis. Write leveling uses up a maximum of three coarse taps of the XIPHY delay to ensure each write DQS is aligned to the nearest clock edge. Memory Controller provides the write data 1TCK early to the PHY, which is then delayed by write leveling up to one memory clock cycle. This means for the zero PCB delay case of a typical simulation the data would be aligned at the DRAM without additional delay added from write calibration.

Write latency calibration can only account for early data, because in the case where the data arrives late at the DRAM there is no push back on the controller to provide the data earlier. With 16 XIPHY coarse taps available (each tap is 90°), four memory clock cycles of shift are available in the XIPHY with one memory clock used by write leveling. This leaves three memory clocks of delay available for write latency calibration.

Figure 3-41 shows the calibration flow to determine the setting required for each byte.

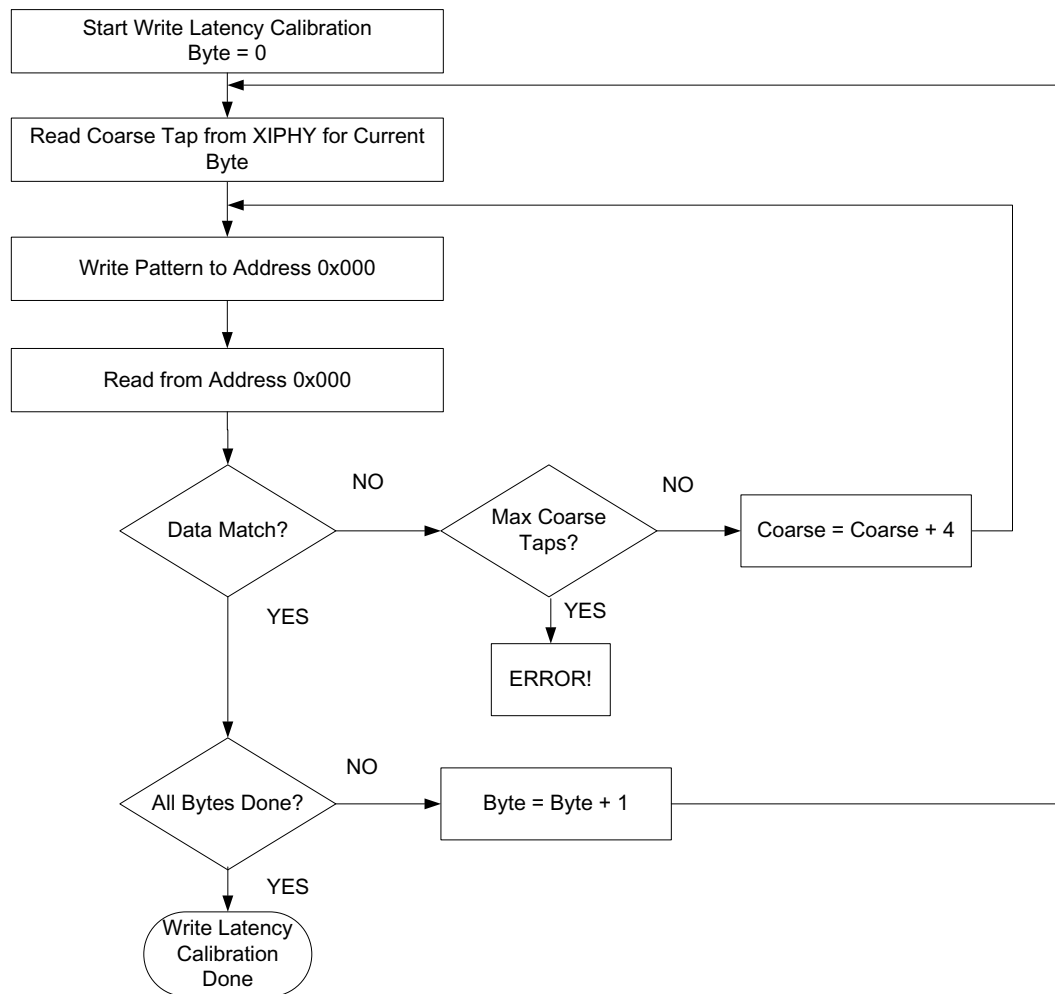


Figure 3-41: Initial Write DQS and DQ with Skew between Bits

The write DQS for the write command is extended for longer than required to ensure the DQS is toggling when the DRAM expects it to clock in the write data. A specific data pattern is used to check when the correct data pattern gets written into the DRAM, as shown in Figure 3-42.



In the example at the start of write latency calibration for the given byte, the target write latency falls in the middle of the data pattern. The returned data would be 55AA9966FFFFFFFF rather than the expected FF00AA5555AA9966. The write DQS and data are delayed using the XIPHY coarse delay and the operation is repeated, until the correct data pattern is found or there are no more coarse taps available. After the pattern is found, the amount of coarse delay required is indicated by WRITE\_LATENCY\_CALIBRATION\_COARSE\_Rank\_Byte.

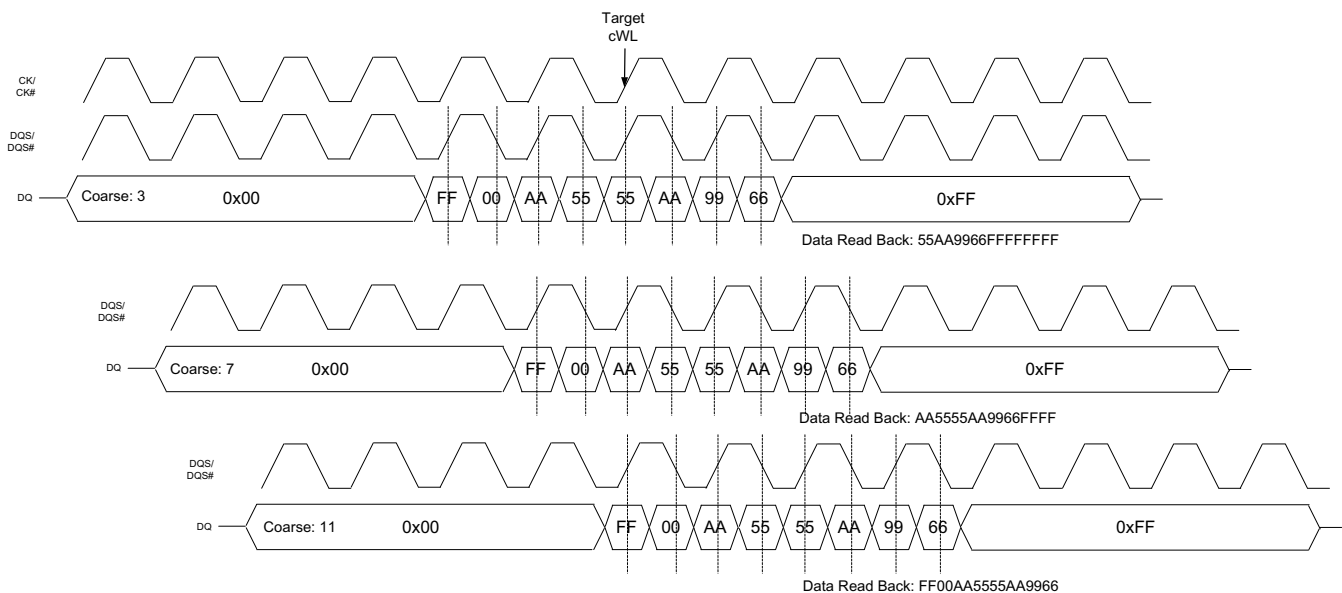


Figure 3-42: Write Latency Calibration Alignment Example

- If the data pattern is not found for a given byte, the data pattern found is checked to see if the data at the maximum delay available still arrives too early (indicating not enough adjustment was available in the XIPHY to align to the correct location) or if the first burst with no extra delay applied is already late (indicating at the start the data would need to be pulled back). The following data pattern is checked:
  - Expected pattern on a per-nibble basis: F0A55A96
  - Late Data Comparison: 00F0AA55A
  - Early Data Comparison: A55A96FF, 5A96FFFF, 96FFFFFF
- If neither of these cases holds true, an attempt is made to try to reclassify the error as either a write or a read failure. A single write burst is sent to the DRAM followed by 20 read bursts. The data from the first read burst is stored for comparison with the remaining 19 read bursts.
- If all the read data matches, the error is classified as a write failure.
- If the data does not match, it is marked as a read failure.

## Write/Read Sanity Check

After Write DQS-to-DQ, a check of the data is made to ensure the previous stage of calibration did not inadvertently leave the write or read path in a bad spot. A single write burst followed by a single read command to the same location is sent to the DRAM, and the data is checked against the expected data across all bytes before continuing. During this step, the expected data pattern as seen on a nibble is 937EC924.

## Read DQS Centering (Complex)

**Note:** Only enabled for data rates above 1,600 Mb/s.

Complex data patterns are used for advanced read DQS centering for memory systems to improve read timing margin. Long and complex data patterns on both the victim and aggressor DQ lanes impact the size and location of the data eye. The objective of the complex calibration step is to generate the worst case data eye on each DQ lane so that the DQS signal can be aligned, resulting in good setup/hold margin during normal operation with any work load.

There are two long data patterns stored in a block RAM, one for a victim DQ lane, and an aggressor pattern for all other DQ lanes. These patterns are used to generate write data, as well as expected data on reads for comparison and error logging. Each pattern consists of 157 8-bit chunks or BL8 bursts.

Each DQ lane of 1-byte takes a turn at being the victim. An RTL state machine automatically selects each DQ lane in turn, MUXing the victim or aggressor patterns to the appropriate DQ lanes, issues the read/write transactions, and records errors. The victim pattern is only walked across the DQ lanes of the selected byte to be calibrated, and all other DQ lanes carry the aggressor pattern, including all lanes in un-selected bytes if there is more than 1-byte lane.

Similar steps to those described in Read DQS Centering are performed, with the PQTR/NQTR starting out at the left edge of the simple window found previously. The complex pattern is written and read back. All bits in a nibble are checked to find the left edge of the window, incrementing the bits together as needed or the PQTR/NQTR to find the aggregate left edge. After the left and right edges are found, it steps through the entire data eye.

## Read $V_{REF}$ Calibration

**Note:** The calibration step is only enabled for the first rank in a multi-rank system.

In DDR4, the read  $V_{REF}$  calibration is enabled by default. Before the read  $V_{REF}$  calibration, the default read  $V_{REF}$  (component value is 20 and SODIMM, UDIMM, and RDIMM are 33) values are used in earlier stages of calibration.

During read  $V_{REF}$  calibration, the calibration logic looks for the read  $V_{REF}$  value with the maximum eye opening per data byte lane. In UltraScale architecture, each data byte lane can be programmed to 73 possible  $V_{REF}$  values. The  $V_{REF}$  search performs five coarse voltage searches at 18, 27, 36, 45, and 54.

Figure 3-43 shows five Coarse  $V_{REF}$  values checked and Coarse  $V_{REF2}$  has the biggest eye opening out of the five search voltages. The  $V_{REF}$  value with the maximum eye opening is chosen and used as the read  $V_{REF}$ . For DDR4 1866 and above, a fine voltage search with a range based on the coarse voltage search with  $\pm 4$  voltage value is done. For example, if the coarse voltage search lands on  $V_{REF}$  value of 27, the fine voltage search range is 23, 24, 25, 26, 27, 28, 29, 30, and 31.




Figure 3-43: Coarse Read  $V_{REF}$  Search

Figure 3-44 shows nine Fine  $V_{REF}$  values checked and the Fine  $V_{REF4}$  has the biggest eye opening out of the nine search voltages.




Figure 3-44: Fine Read  $V_{REF}$  Search for DDR4 1866 and Higher

At each read  $V_{REF}$  value, the read  $V_{REF}$  calibration takes these steps to check the eye opening. The DRAM MPR register is used to provide data pattern to search for an eye opening. The read  $V_{REF}$  calibration starts DQS tap values from the PQTR/NQTR delays from earlier calibration steps. PQTR/NQTR are incremented and traffic is sampled at each PQTR/NQTR tap value. The PQTR/NQTR increment and traffic sampling repeat until a mismatch is found for a given PQTR and NQTR tap value. The number of taps incremented for PQTR and NQTR until data mismatch defines the initial right margins (`pqtr_right` and `nqtr_right`). The sample count taken at each of the PQTR/NQTR tap value described above is low to reduce sampling time. This is indicated by the blue arrows in Figure 3-45.

To obtain a more accurate edge detection for the final `pqtr_right` and `nqtr_right`, the PQTR/NQTR tap values are decremented from `pqtr_right` and `nqtr_right` position with a high number of samples until there is data match. This is indicated by orange arrows in Figure 3-45.

DQS increment to find right edge




Figure 3-45: DQS Search for Right Margin

After the right margins are found, left margins, initial `pqtr_left` and `nqtr_left` are found similarly by decrementing PQTR/NQTR with a low sampling until data mismatch is discovered. Final `pqtr_left` and `nqtr_left` are found by incrementing PQTR/NQTR with a high sampling until data match is discovered. This is indicated in Figure 3-46.




Figure 3-46: DQS Search for Left Margin

Read  $V_{REF}$  is adjusted per byte basis. Minimum eye opening of two nibbles are used to determine the eye size for the byte. For each read  $V_{REF}$  value, the total eye width is determined by the following:

$$\text{MIN}(\text{pqtr\_right\_nibble0} + \text{pqtr\_left\_nibble0}, \text{nqtr\_right\_nibble0} + \text{nqtr\_left\_nibble0}, \text{pqtr\_right\_nibble1} + \text{pqtr\_left\_nibble1}, \text{nqtr\_right\_nibble1} + \text{nqtr\_left\_nibble1})$$

After optimal read  $V_{REF}$  value is determined by the voltage search described previously, DQS is re-centered at  $(\text{pqtr\_right} - \text{pqtr\_left})/2$  for PQTR and  $(\text{nqtr\_right} - \text{nqtr\_left})/2$  for NQTR. Complex data pattern is used for better centering performance.

## Write DQS-to-DQ Centering (Complex)

**Note:** The calibration step is only enabled for the first rank in a multi-rank system. Also, this is only enabled for data rates above 1,600 Mb/s.

For the same reasons as described in the [Read DQS Centering \(Complex\)](#), a complex data pattern is used on the write path to adjust the Write DQS-to-DQ alignment. The same steps as detailed in the [Write DQS-to-DQ Centering](#) are repeated just with a complex data pattern.

## Write $V_{REF}$ Calibration

In DDR4, the write  $V_{REF}$  calibration is enabled by default. Before the write  $V_{REF}$  calibration, the default write  $V_{REF}$  (value depends on  $\text{RTT\_NOM}$  setting in  $\text{MR1}[10:8]$ ) value is used in earlier stages of the calibration.

The write  $V_{REF}$  is similar to read  $V_{REF}$ . Each memory device can be programmed to 51  $V_{REF}$  values.  $V_{REF}$  search performs five coarse voltage search at write  $V_{REF}$  value at 12, 18, 24, 30, and 36.

[Figure 3-43](#) shows five Coarse  $V_{REF}$  values checked and Coarse  $V_{REF}2$  has the biggest eye opening out of the five search voltages. The  $V_{REF}$  value with the maximum eye opening is chosen and used as write  $V_{REF}$ . For DDR4 1866 and above, a fine voltage search with a range based on the coarse voltage search with  $\pm 4$  voltage value is done. For example, if coarse voltage search lands on  $V_{REF}$  value of 18, fine voltage search range is 14, 15, 16, 17, 18, 19, 20, 21, and 22.

Figure 3-44 shows nine Fine  $V_{REF}$  values checked and Fine  $V_{REF4}$  has the biggest eye opening out of the nine search voltages. At each write  $V_{REF}$  value, the write  $V_{REF}$  calibration takes these steps to check the eye opening. For the right margin, write  $V_{REF}$  operates similarly to the read  $V_{REF}$  calibration. The write  $V_{REF}$  calibration right margin searches for  $dqs\_right$  margin by incrementing the DQS ODELAY tap with a low sample count. Then, it follows by decrementing DQS ODELAY tap with a high sample count as shown in Figure 3-45.

Data compare is completed in the device width. For left margin, the write  $V_{REF}$  operates differently than read  $V_{REF}$  calibration. DQ\_ODELAY is incremented starting from the value from write DQS-to-DQ per bit deskew until data mismatch is found.

For each DQ\_ODELAY tap, low sample count is used to determine the initial `left_margin`. DQ\_ODELAY is then decremented with a high sample count to determine the final right margin as shown in Figure 3-47. Data compare is completed using DQ Bit[0] only.

For each  $V_{REF}$  value, the total eye width is determined by the following:

$$\text{MIN}(dqs\_right + dq\_left)$$




Figure 3-47: DQ Search for Left Margin

## Read Leveling Multi-Rank Adjustment

For multi-rank systems the read DQS centering algorithm is ran on each rank, but the final delay setting must be common for all ranks. The results of training each rank separately are stored in XSDB, but the final delay setting is a computed average of the training results across all ranks. The final PQTR/NQTR delay is indicated by `RDLVL_PQTR_CENTER_FINAL_NIBBLE/ RDLVL_NQTR_CENTER_FINAL_NIBBLE`, while the DQ IDelay is `RDLVL_IDELAY_FINAL_BYTE_BIT`.

## Multi-Rank Adjustments and Checks

### DQS Gate Multi-Rank Adjustment

During DQS gate calibration for multi-rank systems, each rank is allowed to calibrate independently given the algorithm as described in [DQS Gate, page 38](#). After all ranks have been calibrated, an adjustment is required before normal operation to ensure fast rank-to-rank switching. The general interconnect signal `clb2phy_rd_en` (indicated by `DQS_GATE_READ_LATENCY_RANK_BYTE` in XSDB) that controls the gate timing on a DRAM-clock-cycle resolution is adjusted here to be the same for a given byte across all ranks.

The coarse taps are adjusted so the timing of the gate opening stays the same for any given rank, where four coarse taps are equal to a single read latency adjustment in the general interconnect. During this step, the algorithm tries to find a common `clb2phy_rd_en` setting where across all ranks for a given byte the coarse setting would not overflow or underflow, starting with the lowest read latency setting found for the byte during calibration. If the lowest setting does not work for all ranks, the `clb2phy_rd_en` increments by one and the check is repeated. The fine tap setting is  $< 90^\circ$ , so it is not included in the adjustment.

If the check reaches the maximum `clb2phy_rd_en` setting initially found during calibration without finding a value that works between all ranks for a byte, an error is asserted ([Table 3-8](#), example #4). If after the adjustment is made and the coarse taps are larger than  $360^\circ$  (four coarse tap settings), a different error is asserted ([Table 3-8](#), example #5). For the error codes, see [Table 3-7, "Error Signal Descriptions," on page 35](#).

**Table 3-8: Examples of DQS Gate Multi-Rank Adjustment (2 Ranks)**

Example	Setting	Calibration		After Multi-Rank Adjustment		
		Rank 0	Rank 1	Rank 0	Rank 1	Result
#1	Read latency	14	15	15	15	Pass
	Coarse taps	8	6	4	6	
#2	Read latency	22	21	22	22	Pass
	Coarse taps	6	9	6	5	
#3	Read latency	10	15	N/A	N/A	Error
	Coarse taps	9	9	N/A	N/A	
#4	Read latency	10	11	10	10	Error
	Coarse taps	6	9	6	13	

For multi-rank systems, the coarse taps must be seven or less so additional delay is added using the general interconnect read latency to compensate for the coarse tap requirement.

## Write Latency Multi-Rank Check

The write latency is allowed to fall wherever it can in multi-rank systems, each rank is allowed to calibrate independently given the algorithms in Write Leveling and Write Latency Calibration. After all ranks have been calibrated and before it finishes, a check is made to ensure certain XIPHY requirements are met on the write path. The difference in write latency between the ranks is allowed to be 180° (or two XIPHY coarse taps).

## Enable VT Tracking

After the DQS gate multi-rank adjustment (if required), a signal is sent to the XIPHY to recalibrate internal delays to start voltage and temperature tracking. The XIPHY asserts a signal when complete, `phy2clb_phy_rdy_upp` for upper nibbles and `phy2clb_phy_rdy_low` for lower nibbles.

For multi-rank systems, when all nibbles are ready for normal operation there is a requirement of the XIPHY where two write-read bursts are required to be sent to the DRAM before starting normal traffic. A data pattern of F00FF00F is used for the first and 0FF00FF0 for the second. The data itself is not checked and is expected to fail.

## Write Read Sanity Check (Multi-Rank Only)

For multi-rank systems, a check of the data for each rank is made to ensure the previous stages of calibration did not inadvertently leave the write or read path in a bad spot. A single write burst followed by a single read command to the same location is sent to each DRAM rank. The data is checked against the expected data across all bytes before continuing.

During this step, the expected data pattern for each rank is shown in [Table 3-9](#).

**Table 3-9: Sanity Check Across All Ranks**

Rank	Expected Data Pattern for Single Burst as Seen on a Nibble
0	A1E04ED8
1	B1E04ED8
2	C1E04ED8
3	D1E04ED8

After all stages are completed across all ranks without any error, `calDone` gets asserted to indicate user traffic can begin. In XSDB, `DBG_END` contains 0x1 if calibration completes and 0x2 if there is a failure.



# Designing with the Core

This chapter includes guidelines and additional information to facilitate designing with the core.

---

## Clocking

The memory interface requires one MMCM, one TXPLL per I/O bank used by the memory interface and two BUFGCE\_DIVs. These clocking components are used to create the proper clock frequencies and phase shifts necessary for the proper operation of the memory interface.

There are two TXPLLs per bank. If a bank is shared by two memory interfaces, both TXPLLs in that bank are used.

**Note:** MIG generates the appropriate clocking structure and no modifications to the RTL are supported.

The MIG tool generates the appropriate clocking structure for the desired interface. This structure must not be modified. The allowed clock configuration is as follows:

- Differential reference clock source connected to GCIO
- GCIO to MMCM (located in center bank of memory interface)
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) driving FPGA logic and all TXPLLs
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) divide by two mode driving 1/2 rate FPGA logic
- Clocking pair of the interface must be in the same SLR of memory interface for the SSI technology devices

## Requirements

### ***GCIO***

- Must use a differential I/O standard
- Must be in the same I/O column as the memory interface
- Must be in the same SLR of memory interface for the SSI technology devices

### ***MMCM***

- MMCM is used to generate the FPGA logic system clock (1/4 of the memory clock)
- Must be located in the center bank of memory interface
- Must use internal feedback
- Input clock frequency divided by input divider must be  $\geq 70$  MHz ( $\text{CLKIN}_x / D \geq 70$  MHz)
- Must use integer multiply and output divide values

### ***BUFGCE\_DIVs and Clock Roots***

- One BUFGCE\_DIV is used to generate the system clock to FPGA logic and another BUFGCE\_DIV is used to divide the system clock by two.
- BUFGCE\_DIVs and clock roots must be located in center most bank of the memory interface.
  - For two bank systems, either bank can be used. MIG is always referred to the top-most selected bank in the Vivado Integrated Design Environment (IDE) as the center bank.
  - For four bank systems, either of the center banks can be chosen. MIG refers to the second bank from the top-most selected bank as the center bank.
  - Both the BUFGCE\_DIVs must be in the same bank.

### ***TXPLL***

- CLKOUTPHY from TXPLL drives XIPHY within its bank
- TXPLL must be set to use a CLKFBOUT phase shift of 90°
- TXPLL must be held in reset until the MMCM lock output goes High
- Must use internal feedback



For designs generated with System Clock configuration of **No Buffer**, MMCM must not be driven by another MMCM/PLL. Cascading clocking structures MMCM → BUFG → MMCM and PLL → BUFG → MMCM are not allowed.

If the MMCM is driven by the GCIO pin of the other bank, then the CLOCK\_DEDICATED\_ROUTE constraint with value "BACKBONE" must be set on the net that is driving MMCM or on the MMCM input. Setting up the CLOCK\_DEDICATED\_ROUTE constraint on the net is preferred. But when the same net is driving two MMCMs, the CLOCK\_DEDICATED\_ROUTE constraint must be managed by considering which MMCM needs the BACKBONE route.

In such cases, the CLOCK\_DEDICATED\_ROUTE constraint can be set on the MMCM input. To use the "BACKBONE" route, any clock buffer that exists in the same CMT tile as the GCIO must exist between the GCIO and MMCM input. The clock buffers that exist in the I/O CMT are BUFG, BUFGCE, BUFGCTRL, and BUFGCE\_DIV. So MIG instantiates BUFG between the GCIO and MMCM when the GCIO pins and MMCM are not in the same bank (see [Figure 4-1](#)).

If the GCIO pin and MMCM are allocated in different banks, MIG generates CLOCK\_DEDICATED\_ROUTE constraints with value as "BACKBONE." If the GCIO pin and MMCM are allocated in the same bank, there is no need to set any constraints on the MMCM input.

Similarly when designs are generated with System Clock Configuration as a **No Buffer** option, you must take care of the "BACKBONE" constraint and the BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV between GCIO and MMCM if GCIO pin and MMCM are allocated in different banks. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations. For more information on clocking, see the *UltraScale Architecture Clocking Resources User Guide* (UG572) [[Ref 3](#)].

**Note:** If two different GCIO pins are used for two MIG IP cores in the same bank, center bank of the memory interface is different for each IP. MIG generates MMCM LOC and CLOCK\_DEDICATED\_ROUTE constraints accordingly.

## Sharing of Input Clock Source (sys\_clk\_p)

If the same GCIO pin must be used for two IP cores, generate the two IP cores with System Clock Configuration option as **No Buffer**. Perform the following changes in the wrapper file in which both IPs are instantiated:

1. MIG generates a single-ended input for system clock pins, such as `sys_clk_i`. Connect the differential buffer output to the single-ended system clock inputs (`sys_clk_i`) of both the IP cores.
2. System clock pins must be allocated within the same I/O column of the memory interface pins allocated. Add the pin LOC constraints for system clock pins and clock constraints in your top-level XDC.

3. You must add a "BACKBONE" constraint on the net that is driving the MMCM or on the MMCM input if GCIO pin and MMCM are not allocated in the same bank. Apart from this, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM to use the "BACKBONE" route.

**Note:**

- The Ultrascale architecture includes an independent XIPHY power supply and TXPLL for each XIPHY. This results in clean, low jitter clocks for the memory system.
- Skew spanning across multiple BUFGs is not a concern because single point of contact exists between BUFG → TXPLL and the same BUFG → System Clock Logic.
- System input clock cannot span I/O columns because the longer the clock lines span, the more jitter is picked up.

---

## Resets

An asynchronous reset (`sys_rst`) input is provided. This is an active-High reset and the `sys_rst` must assert for a minimum pulse width of 5 ns. The `sys_rst` can be an internal or external pin.

---

## PCB Guidelines for DDR3

Strict adherence to all documented DDR3 PCB guidelines is required for successful operation. For more information on PCB guidelines, see the *UltraScale Architecture PCB Design and Pin Planning User Guide* (UG583) [Ref 5].

---

## PCB Guidelines for DDR4

Strict adherence to all documented DDR4 PCB guidelines is required for successful operation. For more information on PCB guidelines, see the *UltraScale Architecture PCB Design and Pin Planning User Guide* (UG583) [Ref 5].

---

## Pin and Bank Rules

### DDR3 Pin Rules

The rules are for single and multi-rank memory interfaces.

- Address/control means `cs_n`, `ras_n`, `cas_n`, `we_n`, `ba`, `ck`, `cke`, `a`, `parity` (valid for RDIMMs only), and `odt`. Multi-rank systems have one `cs_n`, `cke`, `odt`, and one `ck` pair per rank.
- Pins in a byte lane are numbered N0 to N12.
- Byte lanes in a bank are designed by T0, T1, T2, or T3. Nibbles within a byte lane are distinguished by a "U" or "L" designator added to the byte lane designator (T0, T1, T2, or T3). Thus they are T0L, T0U, T1L, T1U, T2L, T2U, T3L, and T3U.

**Note:** There are two PLLs per bank and a controller uses one PLL in every bank that is being used by the interface.

1. `dqs`, `dq`, and `dm` location.
  - a. Designs using x8 or x16 components – `dqs` must be located on a dedicated byte clock pair in the upper nibble designated with "U" (N6 and N7). `dq` associated with a `dqs` must be in same byte lane on any of the other pins except pins 1 and 12.
  - b. Designs using x4 components – `dqs` must be located on the dedicated `dqs` pair in the nibble (N0 and N1 in the lower nibble, N6 and N7 in the upper nibble). `dq`'s associated with a `dqs` must be in the same nibble on any of the other pins except pin N12 (upper nibble).
  - c. `dm` (if used) must be located on pin N0 in the byte lane with the corresponding `dqs`. When `dm` is disabled, pin N0 can be used for `dq` and pin N0 must not be used for address/control signal.

**Note:** `dm` is not supported with x4 devices.
2. The x4 components must be used in pairs. Odd numbers of x4 components are not permitted. Both the upper and lower nibbles of a data byte must be occupied by a x4 `dq/dqs` group.
3. Byte lanes with a `dqs` are considered to be data byte lanes. Pins N1 and N12 can be used for address/control in a data byte lane. If the data byte is in the same bank as the remaining address/control pins, see rule #4.
4. Address/control can be on any of the 13 pins in the address/control byte lanes. Address/control must be contained within the same bank.
5. One `vrp` pin per bank is used and a DCI is required for the interfaces. A `vrp` pin is required in I/O banks containing inputs as well as in output only banks. It is required in output only banks because address/control signals use SSTL15\_DCI/SSTL135\_DCI to enable usage of controlled output impedance. A DCI cascade is not permitted. All rules for the DCI in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4] must be followed.
6. `ck` pair(s) must be on any PN pair(s) in the Address/Control byte lanes.
7. `reset_n` can be on any pin as long as general interconnect timing is met and I/O standard can be accommodated for the chosen bank.
8. Banks can be shared between two controllers.

- a. Each byte lane is dedicated to a specific controller (except for `reset_n`).
  - b. Byte lanes from one controller cannot be placed inside the other. For example, with controllers A and B, "AABB" is allowed, while "ABAB" is not.
9. All I/O banks used by the memory interface must be in the same column.
  10. All I/O banks used by the memory interface must be in the same SLR of the column for the SSI technology devices.
  11. Maximum height of interface is five contiguous banks for 144-bit wide interface. The maximum supported interface is 80-bit wide.
  12. Bank skipping is not allowed.
  13. Input clock for the MMCM in the interface must come from a GCIO pair in the I/O column used for the memory interface. For more information, see [Clocking, page 73](#).
  14. There are dedicated  $V_{REF}$  pins (not included in the rules above). Either internal or external  $V_{REF}$  is permitted. If an external  $V_{REF}$  is not used, the  $V_{REF}$  pins must be pulled to ground by a resistor value specified in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4]. These pins must be connected appropriately for the standard in use.
  15. The interface must be contained within the same I/O bank type (High Range or High Performance). Mixing bank types is not permitted with the exceptions of the `reset_n` in step 7 and the input clock mentioned in step 12.

### DDR3 Pinout Examples



**IMPORTANT:** Due to the calibration stage, there is no need for `set_input_delay/`  
`set_output_delay` on the MIG. Ignore the unconstrained inputs and outputs for MIG and the signals which are calibrated.

Table 4-1 shows an example of a 16-bit DDR3 interface contained within one bank. This example is for a component interface using two x8 DDR3 components.

Table 4-1: 16-Bit DDR3 (x8/x16 Part) Interface Contained in One Bank

Bank	Signal Name	Byte Group	I/O Type
1	a0	T3U_12	–
1	a1	T3U_11	N
1	a2	T3U_10	P
1	a3	T3U_9	N
1	a4	T3U_8	P
1	a5	T3U_7	N
1	a6	T3U_6	P

**Table 4-1: 16-Bit DDR3 (x8/x16 Part) Interface Contained in One Bank (Cont'd)**

Bank	Signal Name	Byte Group	I/O Type
1	a7	T3L_5	N
1	a8	T3L_4	P
1	a9	T3L_3	N
1	a10	T3L_2	P
1	a11	T3L_1	N
1	a12	T3L_0	P
1	a13	T2U_12	–
1	a14	T2U_11	N
1	we_n	T2U_10	P
1	cas_n	T2U_9	N
1	ras_n	T2U_8	P
1	ck_n	T2U_7	N
1	ck_p	T2U_6	P
1	cs_n	T2L_5	N
1	ba0	T2L_4	P
1	ba1	T2L_3	N
1	ba2	T2L_2	P
1	sys_clk_n	T2L_1	N
1	sys_clk_p	T2L_0	P
1	cke	T1U_12	–
1	dq15	T1U_11	N
1	dq14	T1U_10	P
1	dq13	T1U_9	N
1	dq12	T1U_8	P
1	dqs1_n	T1U_7	N
1	dqs1_p	T1U_6	P
1	dq11	T1L_5	N
1	dq10	T1L_4	P
1	dq9	T1L_3	N
1	dq8	T1L_2	P
1	odt	T1L_1	N



**Table 4-1: 16-Bit DDR3 (x8/x16 Part) Interface Contained in One Bank (Cont'd)**

Bank	Signal Name	Byte Group	I/O Type
1	dm1	T1L_0	P
1	vrp	T0U_12	–
1	dq7	T0U_11	N
1	dq6	T0U_10	P
1	dq5	T0U_9	N
1	dq4	T0U_8	P
1	dqs0_n	T0U_7	N
1	dqs0_p	T0U_6	P
1	dq3	T0L_5	N
1	dq2	T0L_4	P
1	dq1	T0L_3	N
1	dq0	T0L_2	P
1	reset_n	T0L_1	N
1	dm0	T0L_0	P

Table 4-2 shows an example of a 16-bit DDR3 interface contained within one bank. This example is for a component interface using four x4 DDR3 components.

**Table 4-2: 16-Bit DDR3 Interface (x4 Part) Contained in One Bank**

Bank	Signal Name	Byte Group	I/O Type
1	a0	T3U_12	–
1	a1	T3U_11	N
1	a2	T3U_10	P
1	a3	T3U_9	N
1	a4	T3U_8	P
1	a5	T3U_7	N
1	a6	T3U_6	P
1	a7	T3L_5	N
1	a8	T3L_4	P
1	a9	T3L_3	N
1	a10	T3L_2	P
1	a11	T3L_1	N
1	a12	T3L_0	P

Table 4-2: 16-Bit DDR3 Interface (x4 Part) Contained in One Bank (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
1	a13	T2U_12	–
1	a14	T2U_11	N
1	we_n	T2U_10	P
1	cas_n	T2U_9	N
1	ras_n	T2U_8	P
1	ck_n	T2U_7	N
1	ck_p	T2U_6	P
1	cs_n	T2L_5	N
1	ba0	T2L_4	P
1	ba1	T2L_3	N
1	ba2	T2L_2	P
1	sys_clk_n	T2L_1	N
1	sys_clk_p	T2L_0	P
1	cke	T1U_12	–
1	dq15	T1U_11	N
1	dq14	T1U_10	P
1	dq13	T1U_9	N
1	dq12	T1U_8	P
1	dqs3_n	T1U_7	N
1	dqs3_p	T1U_6	P
1	dq11	T1L_5	N
1	dq10	T1L_4	P
1	dq9	T1L_3	N
1	dq8	T1L_2	P
1	dqs2_n	T1L_1	N
1	dqs2_p	T1L_0	P
1	vrp	T0U_12	–
1	dq7	T0U_11	N
1	dq6	T0U_10	P
1	dq5	T0U_9	N
1	dq4	T0U_8	P
1	dqs1_n	T0U_7	N
1	dqs1_p	T0U_6	P

Table 4-2: 16-Bit DDR3 Interface (x4 Part) Contained in One Bank (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
1	dq3	T0L_5	N
1	dq2	T0L_4	P
1	dq1	T0L_3	N
1	dq0	T0L_2	P
1	dqs0_n	T0L_1	N
1	dqs0_p	T0L_0	P

## DDR4 Pin Rules

The rules are for single and multi-rank memory interfaces.

- Address/control means `cs_n`, `ras_n`, `cas_n`, `we_n`, `ba`, `bg`, `ck`, `cke`, `a`, `odt`, `act_n`, and `parity` (valid for RDIMMs only). Multi-rank systems have one `cs_n`, `cke`, `odt`, and one `ck` pair per rank.
- Pins in a byte lane are numbered N0 to N12.
- Byte lanes in a bank are designed by T0, T1, T2, or T3. Nibbles within a byte lane are distinguished by a "U" or "L" designator added to the byte lane designator (T0, T1, T2, or T3). Thus they are T0L, T0U, T1L, T1U, T2L, T2U, T3L, and T3U.

**Note:** There are two PLLs per bank and a controller uses one PLL in every bank that is being used by the interface.

- `dqs`, `dq`, and `dm/dbi` location.
  - Designs using x8 or x16 components – `dqs` must be located on a dedicated byte clock pair in the upper nibble designated with "U" (N6 and N7). `dq` associated with a `dqs` must be in same byte lane on any of the other pins except pins N1 and N12.
  - Designs using x4 components – `dqs` must be located on a dedicated byte clock pair in the nibble (N0 and N1 in the lower nibble, N6 and N7 in the upper nibble). `dq` associated with a `dqs` must be in same nibble on any of the other pins except pin N12 (upper nibble). The lower nibble `dq` and upper nibble `dq` must be allocated in the same byte lane.

**Note:** The `dm/dbi` port is not supported in x4 DDR4 devices.

- `dm/dbi` must be on pin N0 in the byte lane with the associated `dqs`. Write and read `dbi` are required for per pin data rates above 2,133 Mb/s. Therefore, data mask functionality is not available above 2,133 Mb/s.
- The x16 components must have the `ldqs` connected to the even `dqs` and the `udqs` must be connected to the `ldqs + 1`. The first x16 component has `ldqs` connected to `dqs0` and `udqs` connected to `dqs1` in the XDC file. The second x16 component has `ldqs` connected to `dqs2` and `udqs` connected to `dqs3`. This pattern continues as needed for the interface. This does not restrict the physical location of the byte

lanes. The byte lanes associated with the  $dqs$  's might be moved as desired in the Vivado IDE to achieve optimal PCB routing.

2. The x4 components must be used in pairs. Odd numbers of x4 components are not permitted. Both the upper and lower nibbles of a data byte must be occupied by a x4  $dq/dqs$  group. Each byte lane containing two x4 nibbles must have sequential nibbles with the even nibble being the lower number. For example, a byte lane can have nibbles 0 and 1, or 2 and 3, but must not have 1 and 2. The ordering of the nibbles within a byte lane is not important.
3. Byte lanes with a  $dqs$  are considered to be data byte lanes. Pins N1 and N12 can be used for address/control in a data byte lane. If the data byte is in the same bank as the remaining address/control pins, see rule #4.
4. Address/control can be on any of the 13 pins in the address/control byte lanes. Address/control must be contained within the same bank.
5. One  $vrp$  pin per bank is used and a DCI is required for the interfaces. A  $vrp$  pin is required in I/O banks containing inputs as well as in output only banks. It is required in output only banks because address/control signals use SSTL12\_DCI to enable usage of controlled output impedance. A DCI cascade is not permitted. All rules for the DCI in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4] must be followed.
6.  $ck$  pair(s) must be on any PN pair(s) in the Address/Control byte lanes.
7.  $reset\_n$  can be on any pin as long as general interconnect timing is met and I/O standard can be accommodated for the chosen bank.
8. Banks can be shared between two controllers.
  - a. Each byte lane is dedicated to a specific controller (except for  $reset\_n$ ).
  - b. Byte lanes from one controller cannot be placed inside the other. For example, with controllers A and B, "AABB" is allowed, while "ABAB" is not.
9. All I/O banks used by the memory interface must be in the same column.
10. All I/O banks used by the memory interface must be in the same SLR of the column for the SSI technology devices.
11. Maximum height of interface is five contiguous banks for 144-bit wide interface. The maximum supported interface is 80-bit wide.
12. Bank skipping is not allowed.
13. Input clock for the MMCM in the interface must come from the a GCIO pair in the I/O column used for the memory interface. For more information, see [Clocking, page 73](#).
14. The dedicated  $V_{REF}$  pins in the banks used for DDR4 must be tied to ground with a resistor value specified in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4]. Internal  $V_{REF}$  is required for DDR4.

15. The interface must be contained within the same I/O bank type (High Performance). Mixing bank types is not permitted with the exceptions of the `reset_n` in step #7 and the input clock mentioned in step #13.
16. The `par` input for command and address parity, `alert_n` input/output, and the TEN input for Connectivity Test Mode are not supported by this interface. Consult the memory vendor for information on the proper connection for these pins when not used.



**IMPORTANT:** Component interfaces should be created with the same component for all components in the interface. x16 components have a different number of bank groups than the x8 components. For example, a 72-bit wide component interface should be created by using nine x8 components or five x16 components where half of one component is not used. Four x16 components and one x8 component is not permissible.

## DDR4 Pinout Examples



**IMPORTANT:** Due to the calibration stage, there is no need for `set_input_delay/`  
`set_output_delay` on the MIG. Ignore the unconstrained inputs and outputs for MIG and the signals which are calibrated.

Table 4-3 shows an example of a 32-bit DDR4 interface contained within two banks. This example is for a component interface using four x8 DDR4 components.

Table 4-3: 32-Bit DDR4 Interface Contained in Two Banks

Bank	Signal Name	Byte Group	I/O Type
<b>Bank 1</b>			
1	–	T3U_12	–
1	–	T3U_11	N
1	–	T3U_10	P
1	–	T3U_9	N
1	–	T3U_8	P
1	–	T3U_7	N
1	–	T3U_6	P
1	–	T3L_5	N
1	–	T3L_4	P
1	–	T3L_3	N
1	–	T3L_2	P
1	–	T3L_1	N
1	–	T3L_0	P

Table 4-3: 32-Bit DDR4 Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
1	–	T2U_12	–
1	–	T2U_11	N
1	–	T2U_10	P
1	–	T2U_9	N
1	–	T2U_8	P
1	–	T2U_7	N
1	–	T2U_6	P
1	–	T2L_5	N
1	–	T2L_4	P
1	–	T2L_3	N
1	–	T2L_2	P
1	–	T2L_1	N
1	–	T2L_0	P
1	reset_n	T1U_12	–
1	dq31	T1U_11	N
1	dq30	T1U_10	P
1	dq29	T1U_9	N
1	dq28	T1U_8	P
1	dqs3_c	T1U_7	N
1	dqs3_t	T1U_6	P
1	dq27	T1L_5	N
1	dq26	T1L_4	P
1	dq25	T1L_3	N
1	dq24	T1L_2	P
1	unused	T1L_1	N
1	dm3/dbi3	T1L_0	P
1	vrp	T0U_12	–
1	dq23	T0U_11	N
1	dq22	T0U_10	P
1	dq21	T0U_9	N
1	dq20	T0U_8	P

Table 4-3: 32-Bit DDR4 Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
1	dqs2_c	T0U_7	N
1	dqs2_t	T0U_6	P
1	dq19	T0L_5	N
1	dq18	T0L_4	P
1	dq17	T0L_3	N
1	dq16	T0L_2	P
1	–	T0L_1	N
1	dm2/dbi2	T0L_0	P
<b>Bank 2</b>			
2	a0	T3U_12	–
2	a1	T3U_11	N
2	a2	T3U_10	P
2	a3	T3U_9	N
2	a4	T3U_8	P
2	a5	T3U_7	N
2	a6	T3U_6	P
2	a7	T3L_5	N
2	a8	T3L_4	P
2	a9	T3L_3	N
2	a10	T3L_2	P
2	a11	T3L_1	N
2	a12	T3L_0	P
2	a13	T2U_12	–
2	we_n/a14	T2U_11	N
2	cas_n/a15	T2U_10	P
2	ras_n/a16	T2U_9	N
2	act_n	T2U_8	P
2	ck_c	T2U_7	N
2	ck_t	T2U_6	P
2	ba0	T2L_5	N
2	ba1	T2L_4	P
2	bg0	T2L_3	N

Table 4-3: 32-Bit DDR4 Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
2	bg1	T2L_2	P
2	sys_clk_n	T2L_1	N
2	sys_clk_p	T2L_0	P
2	cs_n	T1U_12	–
2	dq15	T1U_11	N
2	dq14	T1U_10	P
2	dq13	T1U_9	N
2	dq12	T1U_8	P
2	dqs1_c	T1U_7	N
2	dqs1_t	T1U_6	P
2	dq11	T1L_5	N
2	dq10	T1L_4	P
2	dq9	T1L_3	N
2	dq8	T1L_2	P
2	odt	T1L_1	N
2	dm1/dbi1	T1L_0	P
2	vrp	T0U_12	–
2	dq7	T0U_11	N
2	dq6	T0U_10	P
2	dq5	T0U_9	N
2	dq4	T0U_8	P
2	dqs0_c	T0U_7	N
2	dqs0_t	T0U_6	P
2	dq3	T0L_5	N
2	dq2	T0L_4	P
2	dq1	T0L_3	N
2	dq0	T0L_2	P
2	cke	T0L_1	N
2	dm0/dbi0	T0L_0	P



Table 4-4 shows an example of a 16-bit DDR4 interface contained within a single bank. This example is for a component interface using four x4 DDR4 components.

**Table 4-4: 16-Bit DDR4 Interface (x4 Part) Contained in One Bank**

Bank	Signal Name	Byte Group	I/O Type
1	a0	T3U_12	–
1	a1	T3U_11	N
1	a2	T3U_10	P
1	a3	T3U_9	N
1	a4	T3U_8	P
1	a5	T3U_7	N
1	a6	T3U_6	P
1	a7	T3L_5	N
1	a8	T3L_4	P
1	a9	T3L_3	N
1	a10	T3L_2	P
1	a11	T3L_1	N
1	a12	T3L_0	P
1	a13	T2U_12	–
1	we_n/a14	T2U_11	N
1	cas_n/a15	T2U_10	P
1	ras_n/a16	T2U_9	N
1	act_n	T2U_8	P
1	ck_c	T2U_7	N
1	ck_t	T2U_6	P
1	ba0	T2L_5	N
1	ba1	T2L_4	P
1	bg0	T2L_3	N
1	bg1	T2L_2	P
1	odt	T2L_1	N
1	cke	T2L_0	P
1	cs_n	T1U_12	–
1	dq15	T1U_11	N
1	dq14	T1U_10	P
1	dq13	T1U_9	N
1	dq12	T1U_8	P

**Table 4-4: 16-Bit DDR4 Interface (x4 Part) Contained in One Bank (Cont'd)**

Bank	Signal Name	Byte Group	I/O Type
1	dqs3_c	T1U_7	N
1	dqs3_t	T1U_6	P
1	dq11	T1L_5	N
1	dq10	T1L_4	P
1	dq9	T1L_3	N
1	dq8	T1L_2	P
1	dqs2_c	T1L_1	N
1	dqs2_t	T1L_0	P
1	vrp	T0U_12	–
1	dq7	T0U_11	N
1	dq6	T0U_10	P
1	dq5	T0U_9	N
1	dq4	T0U_8	P
1	dqs1_c	T0U_7	N
1	dqs1_t	T0U_6	P
1	dq3	T0L_5	N
1	dq2	T0L_4	P
1	dq1	T0L_3	N
1	dq0	T0L_2	P
1	dqs0_c	T0L_1	N
1	dqs0_t	T0L_0	P

**Note:** System clock pins (`sys_clk_p` and `sys_clk_n`) are allocated in different banks.

## Pin Mapping for x4 RDIMMs

Table 4-5 is an example showing the pin mapping for x4 DDR3 registered DIMMs between the memory data sheet and the XDC.

**Table 4-5: Pin Mapping for x4 DDR3 DIMMs**

Memory Data Sheet	MIG XDC
DQ[63:0]	DQ[63:0]
CB3 to CB0	DQ[67:64]
CB7 to CB4	DQ[71:68]

Table 4-5: Pin Mapping for x4 DDR3 DIMMs (Cont'd)

Memory Data Sheet	MIG XDC
DQS0, $\overline{\text{DQS0}}$	DQS[0], DQS_N[0]
DQS1, $\overline{\text{DQS1}}$	DQS[2], DQS_N[2]
DQS2, $\overline{\text{DQS2}}$	DQS[4], DQS_N[4]
DQS3, $\overline{\text{DQS3}}$	DQS[6], DQS_N[6]
DQS4, $\overline{\text{DQS4}}$	DQS[8], DQS_N[8]
DQS5, $\overline{\text{DQS5}}$	DQS[10], DQS_N[10]
DQS6, $\overline{\text{DQS6}}$	DQS[12], DQS_N[12]
DQS7, $\overline{\text{DQS7}}$	DQS[14], DQS_N[14]
DQS8, $\overline{\text{DQS8}}$	DQS[16], DQS_N[16]
DQS9, $\overline{\text{DQS9}}$	DQS[1], DQS_N[1]
DQS10, $\overline{\text{DQS10}}$	DQS[3], DQS_N[3]
DQS11, $\overline{\text{DQS11}}$	DQS[5], DQS_N[5]
DQS12, $\overline{\text{DQS12}}$	DQS[7], DQS_N[7]
DQS13, $\overline{\text{DQS13}}$	DQS[9], DQS_N[9]
DQS14, $\overline{\text{DQS14}}$	DQS[11], DQS_N[11]
DQS15, $\overline{\text{DQS15}}$	DQS[13], DQS_N[13]
DQS16, $\overline{\text{DQS16}}$	DQS[15], DQS_N[15]
DQS17, $\overline{\text{DQS17}}$	DQS[17], DQS_N[17]

Table 4-6 is an example showing the pin mapping for x4 DDR4 registered DIMMs between the memory data sheet and the XDC.

Table 4-6: Pin Mapping for x4 DDR4 DIMMs

Memory Data Sheet	MIG XDC
DQ[63:0]	DQ[63:0]
CB3 to CB0	DQ[67:64]
CB7 to CB4	DQ[71:68]
DQS0	DQS[0]
DQS1	DQS[2]
DQS2	DQS[4]
DQS3	DQS[6]
DQS4	DQS[8]
DQS5	DQS[10]
DQS6	DQS[12]
DQS7	DQS[14]
DQS8	DQS[16]

Table 4-6: Pin Mapping for x4 DDR4 DIMMs (Cont'd)

Memory Data Sheet	MIG XDC
DQS9	DQS[1]
DQS10	DQS[3]
DQS11	DQS[5]
DQS12	DQS[7]
DQS13	DQS[9]
DQS14	DQS[11]
DQS15	DQS[13]
DQS16	DQS[15]
DQS17	DQS[17]

## Protocol Description

This core has the following interfaces:

- [User Interface](#)
- [AXI4 Slave Interface](#)
- [PHY Only Interface](#)

## User Interface

The user interface signals are described in [Table 4-7](#) and connects to an FPGA user design to allow access to an external memory device. The user interface is layered on top of the native interface which is described earlier in the controller description.

Table 4-7: User Interface

Signal	Direction	Description
app_addr[ADDR_WIDTH – 1:0]	Input	This input indicates the address for the current request.
app_cmd[2:0]	Input	This input selects the command for the current request.
app_autoprecharge <sup>(1)</sup>	Input	This input instructs the controller to set the A10 autoprecharge bit on the DRAM CAS command for the current request.
app_en	Input	This is the active-High strobe for the app_addr[], app_cmd[2:0], app_sz, and app_hi_pri inputs.
app_rdy	Output	This output indicates that the user interface is ready to accept commands. If the signal is deasserted when app_en is enabled, the current app_cmd, app_autoprecharge, and app_addr must be retried until app_rdy is asserted.

Table 4-7: User Interface (Cont'd)

Signal	Direction	Description
app_hi_pri	Input	This input is reserved and should be tied to 0.
app_rd_data [APP_DATA_WIDTH – 1:0]	Output	This provides the output data from read commands.
app_rd_data_end	Output	This active-High output indicates that the current clock cycle is the last cycle of output data on app_rd_data[].
app_rd_data_valid	Output	This active-High output indicates that app_rd_data[] is valid.
app_sz	Input	This input is reserved and should be tied to 0.
app_wdf_data [APP_DATA_WIDTH – 1:0]	Input	This provides the data for write commands.
app_wdf_end	Input	This active-High input indicates that the current clock cycle is the last cycle of input data on app_wdf_data[].
app_wdf_mask [APP_MASK_WIDTH – 1:0]	Input	This provides the mask for app_wdf_data[].
app_wdf_rdy	Output	This output indicates that the write data FIFO is ready to receive data. Write data is accepted when app_wdf_rdy = 1'b1 and app_wdf_wren = 1'b1.
app_wdf_wren	Input	This is the active-High strobe for app_wdf_data[].
app_ref_req <sup>(2)</sup>	Input	User refresh request.
app_ref_ack <sup>(2)</sup>	Output	User refresh request completed.
app_zq_req <sup>(2)</sup>	Input	User ZQCS command request.
app_zq_ack <sup>(2)</sup>	Output	User ZQCS command request completed.
ui_clk	Output	This user interface clock must be one quarter of the DRAM clock.
init_calib_complete	Output	PHY asserts init_calib_complete when calibration is finished.
ui_clk_sync_rst	Output	This is the active-High user interface reset.
addn_ui_clkout1	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout2	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout3	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout4	Output	Additional clock outputs provided based on user requirement.
dbg_clk	Output	Debug Clock. Do not connect any signals to dbg_clk and keep the port open during instantiation.
sl_iport0	Input [36:0]	Input Port 0 (* KEEP = "true" *)
sl_oport0	Output [16:0]	Output Port 0 (* KEEP = "true" *)
c0_ddr4_app_correct_en_i	Input	DDR4 Correct Enable Input
app_raw_not_ecc	Input	Reserved for future use. Tie Low.

**Notes:**

1. This port appears when "Enable Precharge Input" option is enabled in the Vivado IDE.
2. These ports appear upon enabling "Enable User Refresh and ZQCS Input" option in the Vivado IDE.

### ***app\_addr[ADDR\_WIDTH – 1:0]***

This input indicates the address for the request currently being submitted to the user interface. The user interface aggregates all the address fields of the external SDRAM and presents a flat address space.

The MEM\_ADDR\_ORDER parameter determines how `app_addr` is mapped to the SDRAM address bus and chip select pins. This mapping can have a significant impact on memory bandwidth utilization. “ROW\_COLUMN\_BANK” is the recommended MEM\_ADDR\_ORDER setting. Table 4-8 through Table 4-11 show the “ROW\_COLUMN\_BANK” mapping for DDR3 and DDR4 with examples. Note that the three LSBs of `app_addr` map to the column address LSBs which correspond to SDRAM burst ordering.

The controller does not support burst ordering so these low order bits are ignored, making the effective minimum `app_addr` step size hex 8..

**Table 4-8: DDR3 “ROW\_COLUMN\_BANK” Mapping**

SDRAM	<code>app_addr</code> Mapping
Rank	$(\text{RANK} == 1) ? 1'b0 : \text{app\_addr}[\text{BANK\_WIDTH} + \text{COL\_WIDTH} + \text{ROW\_WIDTH} + : \text{RANK\_WIDTH}]$
Row	$\text{app\_addr}[\text{BANK\_WIDTH} + \text{COL\_WIDTH} + : \text{ROW\_WIDTH}]$
Column	$\text{app\_addr}[3 + \text{BANK\_WIDTH} + : \text{COL\_WIDTH} - 3], \text{app\_addr}[2:0]$
Bank	$\text{app\_addr}[3 + : \text{BANK\_WIDTH} - 1], \text{app\_addr}[2 + \text{BANK\_WIDTH} + : 1]$

**Table 4-9: DDR3 4 GB (512 MB x8) Single Rank Mapping Example**

SDRAM Bus	Row[15:0]	Column[9:0]	Bank[2:0]
<code>app_addr</code> Bits	28 through 13	12 through 6, and 2, 1, 0	4, 3, 5

**Table 4-10: DDR4 “ROW\_COLUMN\_BANK” Mapping**

SDRAM	<code>app_addr</code> Mapping
Rank	$(\text{RANK} == 1) ? 1'b0 : \text{app\_addr}[\text{BANK\_GROUP\_WIDTH} + \text{BANK\_WIDTH} + \text{COL\_WIDTH} + \text{ROW\_WIDTH} + : \text{RANK\_WIDTH}]$
Row	$\text{app\_addr}[\text{BANK\_GROUP\_WIDTH} + \text{BANK\_WIDTH} + \text{COL\_WIDTH} + : \text{ROW\_WIDTH}]$
Column	$\text{app\_addr}[3 + \text{BANK\_GROUP\_WIDTH} + \text{BANK\_WIDTH} + : \text{COL\_WIDTH} - 3], \text{app\_addr}[2:0]$
Bank	$\text{app\_addr}[3 + \text{BANK\_GROUP\_WIDTH} + : \text{BANK\_WIDTH}]$
Bank Group	$\text{app\_addr}[3 + : \text{BANK\_GROUP\_WIDTH}]$

**Table 4-11: DDR4 4 GB (512 MB x8) Single Rank Mapping Example**

SDRAM Bus	Row[14:0]	Column[9:0]	Bank[1:0]	Bank Group[1:0]
<code>app_addr</code> Bits	28 through 14	13 through 7, and 2, 1, 0	6, 5	4, 3

The "ROW\_COLUMN\_BANK" setting maps `app_addr[4:3]` to the DDR4 bank group bits or DDR3 bank bits used by the controller to interleave between its group FSMs. The lower order address bits equal to `app_addr[5]` and above map to the remaining SDRAM bank and column address bits. The highest order address bits map to the SDRAM row. This mapping is ideal for workloads that have address streams that increment linearly by a constant step size of hex 8 for long periods. With this configuration and workload, transactions sent to the user interface are evenly interleaved across the controller group FSMs, making the best use of the controller resources. In addition, this arrangement tends to generate hits to open pages in the SDRAM. The combination of group FSM interleaving and SDRAM page hits results in very high SDRAM data bus utilization.

Address streams other than the simple increment pattern tend to have lower SDRAM bus utilization. You can recover this performance loss by tuning the mapping of your design flat address space to the `app_addr` input port of the user interface. If you have knowledge of your address sequence, you can add logic to map your address bits with the highest toggle rate to the lowest `app_addr` bits, starting with `app_addr[3]` and working up from there.

For example, if you know that your workload address Bits[4:3] toggle much less than Bits[10:9], which toggle at the highest rate, you could add logic to swap these bits so that your address Bits[10:9] map to `app_addr[4:3]`. The result is an improvement in how the address stream interleaves across the controller group FSMs, resulting in better controller throughput and higher SDRAM data bus utilization.

### ***app\_cmd[2:0]***

This input specifies the command for the request currently being submitted to the user interface. The available commands are shown in Table 4-12. With ECC enabled, the `wr_bytes` operation is required for writes with any non-zero `app_wdf_mask` bits. The `wr_bytes` triggers a read-modify-write flow in the controller, which is needed only for writes with masked data in ECC mode.

Table 4-12: Commands for `app_cmd[2:0]`

Operation	<code>app_cmd[2:0]</code> Code
Write	000
Read	001
<code>wr_bytes</code>	011

### ***app\_autoprecharge***

This input specifies the state of the A10 `autoprecharge` bit for the DRAM CAS command for the request currently being submitted to the user interface. When this input is Low, the Memory Controller issues a DRAM RD or WR CAS command. When this input is High, the controller issues a DRAM RDA or WRA CAS command. This input provides per request control, but can also be tied off to configure the controller statically for open or closed page mode operation.

### ***app\_en***

This input strobes in a request. Apply the desired values to `app_addr[]`, `app_cmd[2:0]`, and `app_hi_pri`, and then assert `app_en` to submit the request to the user interface. This initiates a handshake that the user interface acknowledges by asserting `app_rdy`.

### ***app\_wdf\_data[APP\_DATA\_WIDTH – 1:0]***

This bus provides the data currently being written to the external memory.

### ***app\_wdf\_end***

This input indicates that the data on the `app_wdf_data[]` bus in the current cycle is the last data for the current request.

### ***app\_wdf\_mask[APP\_MASK\_WIDTH – 1:0]***

This bus indicates which bits of `app_wdf_data[]` are written to the external memory and which bits remain in their current state.

### ***app\_wdf\_wren***

This input indicates that the data on the `app_wdf_data[]` bus is valid.

### ***app\_rdy***

This output indicates whether the request currently being submitted to the user interface is accepted. If the user interface does not assert this signal after `app_en` is asserted, the current request must be retried. The `app_rdy` output is not asserted if:

- PHY/Memory initialization is not yet completed.
- All the controller Group FSMs are occupied (can be viewed as the command buffer being full).
  - A read is requested and the read buffer is full.
  - A write is requested and no write buffer pointers are available.
- A periodic read is being inserted.

### ***app\_rd\_data[APP\_DATA\_WIDTH – 1:0]***

This output contains the data read from the external memory.



### ***app\_rd\_data\_end***

This output indicates that the data on the `app_rd_data[]` bus in the current cycle is the last data for the current request.

### ***app\_rd\_data\_valid***

This output indicates that the data on the `app_rd_data[]` bus is valid.

### ***app\_wdf\_rdy***

This output indicates that the write data FIFO is ready to receive data. Write data is accepted when both `app_wdf_rdy` and `app_wdf_wren` are asserted.

### ***app\_ref\_req***

When asserted, this active-High input requests that the Memory Controller send a refresh command to the DRAM. It must be pulsed for a single cycle to make the request and then deasserted at least until the `app_ref_ack` signal is asserted to acknowledge the request and indicate that it has been sent.

### ***app\_ref\_ack***

When asserted, this active-High input acknowledges a refresh request and indicates that the command has been sent from the Memory Controller to the PHY.

### ***app\_zq\_req***

When asserted, this active-High input requests that the Memory Controller send a ZQ calibration command to the DRAM. It must be pulsed for a single cycle to make the request and then deasserted at least until the `app_zq_ack` signal is asserted to acknowledge the request and indicate that it has been sent.

### ***app\_zq\_ack***

When asserted, this active-High input acknowledges a ZQ calibration request and indicates that the command has been sent from the Memory Controller to the PHY.

### ***ui\_clk\_sync\_rst***

This is the reset from the user interface which is in synchronous with `ui_clk`.

## ***ui\_clk***

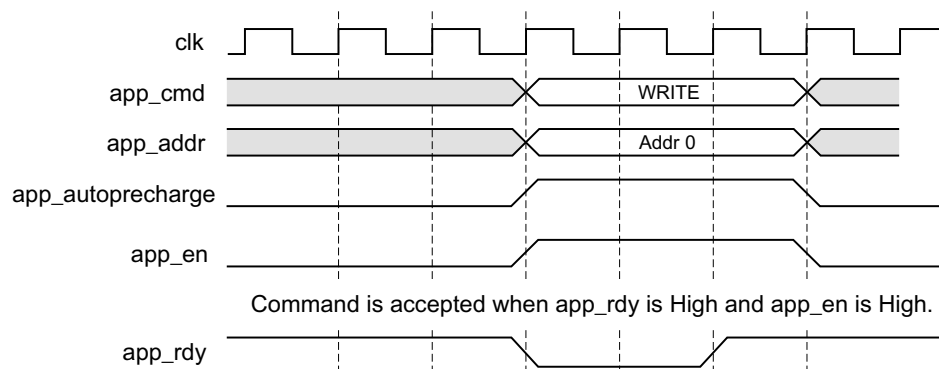
This is the output clock from the user interface. It must be a quarter the frequency of the clock going out to the external SDRAM, which depends on 4:1 mode selected in Vivado IDE.

## ***init\_calib\_complete***

PHY asserts `init_calib_complete` when calibration is finished. The application has no need to wait for `init_calib_complete` before sending commands to the Memory Controller.

## **Command Path**

When the user logic `app_en` signal is asserted and the `app_rdy` signal is asserted from the user interface, a command is accepted and written to the FIFO by the user interface. The command is ignored by the user interface whenever `app_rdy` is deasserted. The user logic needs to hold `app_en` High along with the valid command, autoprecharge, and address values until `app_rdy` is asserted as shown for the "write with autoprecharge" transaction in [Figure 4-2](#).



**Figure 4-2: User Interface Command Timing Diagram with `app_rdy` Asserted**

A non back-to-back write command can be issued as shown in [Figure 4-3](#). This figure depicts three scenarios for the `app_wdf_data`, `app_wdf_wren`, and `app_wdf_end` signals as follows:

1. Write data is presented along with the corresponding write command.
2. Write data is presented before the corresponding write command.
3. Write data is presented after the corresponding write command, but should not exceed the limitation of two clock cycles.

For write data that is output after the write command has been registered, as shown in Note 3 ([Figure 4-3](#)), the maximum delay is two clock cycles.


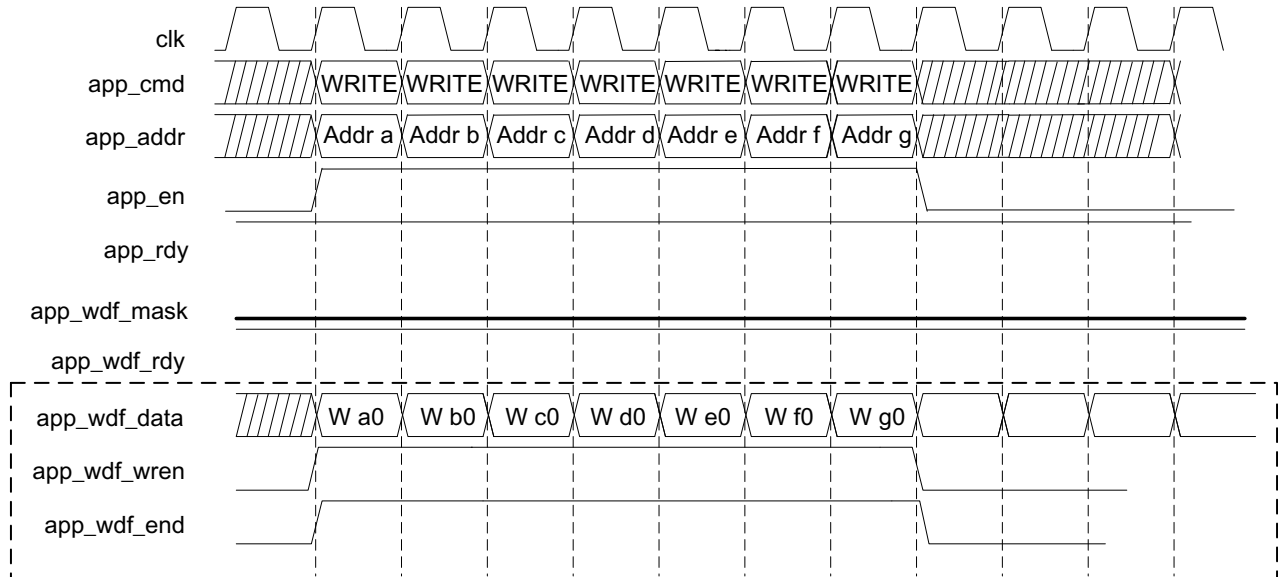


Figure 4-3: 4:1 Mode User Interface Write Timing Diagram (Memory Burst Type = BL8)

## Write Path

The write data is registered in the write FIFO when **app\_wdf\_wren** is asserted and **app\_wdf\_rdy** is High (Figure 4-4). If **app\_wdf\_rdy** is deasserted, the user logic needs to hold **app\_wdf\_wren** and **app\_wdf\_end** High along with the valid **app\_wdf\_data** value until **app\_wdf\_rdy** is asserted. The **app\_wdf\_mask** signal can be used to mask out the bytes to write to external memory.




**Figure 4-4: 4:1 Mode User Interface Back-to-Back Write Commands Timing Diagram (Memory Burst Type = BL8)**

The timing requirement for `app_wdf_data`, `app_wdf_wren`, and `app_wdf_end` relative to their associated write command is the same for back-to-back writes as it is for single writes, as shown in Figure 4-3.

The map of the application interface data to the DRAM output data can be explained with an example.

For a 4:1 Memory Controller to DRAM clock ratio with an 8-bit memory, at the application interface, if the 64-bit data driven is 0000\_0806\_0000\_0805 (Hex), the data at the DRAM interface is as shown in Figure 4-5. This is for a BL8 (Burst Length 8) transaction.



**Figure 4-5: Data at the DRAM Interface for 4:1 Mode**

The data values at different clock edges are as shown in [Table 4-13](#).

**Table 4-13: Data Values at Different Clock Edges**

Rise0	Fall0	Rise1	Fall1	Rise2	Fall2	Rise3	Fall3
05	08	00	00	06	08	00	00

[Table 4-14](#) shows a generalized representation of how DRAM DQ bus data is concatenated to form application interface data signals. `app_wdf_data` is shown in [Table 4-14](#), but the table applies equally to `app_rd_data`. Each byte of the DQ bus has eight bursts, Rise0 (burst 0) through Fall3 (burst 7) as shown previously in [Table 4-13](#), for a total of 64 data bits. When concatenated with Rise0 in the LSB position and Fall3 in the MSB position, a 64-bit chunk of the `app_wdf_data` signal is formed.

For example, the eight bursts of `ddr3_dq[7:0]` corresponds to DQ bus byte 0, and when concatenated as described here, they map to `app_wdf_data[63:0]`. To be clear on the concatenation order, `ddr3_dq[0]` from Rise0 (burst 0) maps to `app_wdf_data[0]`, and `ddr3_dq[7]` from Fall3 (burst 7) maps to `app_wdf_data[63]`. The table shows a second example, mapping DQ byte 1 to `app_wdf_data[127:64]`, as well as the formula for DQ byte N.

**Table 4-14: DRAM DQ Bus Data Map**

DQ Bus Byte	App Interface Signal	DDR Bus Signal at Each BL8 Burst Position				
		Fall3	...	Rise1	Fall0	Rise0
N	<code>app_wdf_data[(N + 1) × 64 – 1:N × 64]</code>	<code>ddr3_dq[(N + 1) × 8 – 1:N × 8]</code>	...	<code>ddr3_dq[(N + 1) × 8 – 1:N × 8]</code>	<code>ddr3_dq[(N + 1) × 8 – 1:N × 8]</code>	<code>ddr3_dq[(N + 1) × 8 – 1:N × 8]</code>
1	<code>app_wdf_data[127:64]</code>	<code>ddr3_dq[15:8]</code>	...	<code>ddr3_dq[15:8]</code>	<code>ddr3_dq[15:8]</code>	<code>ddr3_dq[15:8]</code>
0	<code>app_wdf_data[63:0]</code>	<code>ddr3_dq[7:0]</code>	...	<code>ddr3_dq[7:0]</code>	<code>ddr3_dq[7:0]</code>	<code>ddr3_dq[7:0]</code>

In a similar manner to the DQ bus mapping, the DM bus maps to `app_wdf_mask` by concatenating the DM bits in the same burst order. Example for the first two bytes of the DRAM bus are shown in [Table 4-15](#), and the formula for mapping DM for byte N is also given.

**Table 4-15: DRAM DM Bus Data Map**

DM Bus Byte	App Interface Signal	DDR Bus Signal at Each BL8 Burst Position				
		Fall3	...	Rise1	Fall0	Rise0
N	<code>app_wdf_mask[(N + 1) × 8 – 1:N × 8]</code>	<code>ddr3_dm[N]</code>	...	<code>ddr3_dm[N]</code>	<code>ddr3_dm[N]</code>	<code>ddr3_dm[N]</code>
1	<code>app_wdf_mask[15:0]</code>	<code>ddr3_dq[1]</code>	...	<code>ddr3_dm[1]</code>	<code>ddr3_dm[1]</code>	<code>ddr3_dm[1]</code>
0	<code>app_wdf_mask[7:0]</code>	<code>ddr3_dq[0]</code>	...	<code>ddr3_dm[0]</code>	<code>ddr3_dm[0]</code>	<code>ddr3_dm[0]</code>

## Read Path

The read data is returned by the user interface in the requested order and is valid when `app_rd_data_valid` is asserted (Figure 4-6 and Figure 4-7). The `app_rd_data_end` signal indicates the end of each read command burst and is not needed in user logic.

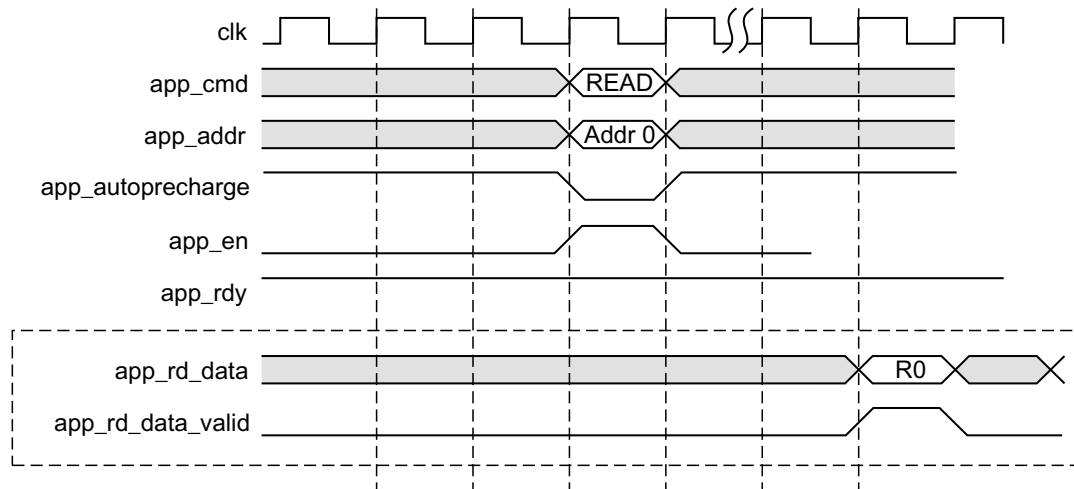


Figure 4-6: 4:1 Mode User Interface Read Timing Diagram (Memory Burst Type = BL8) #1

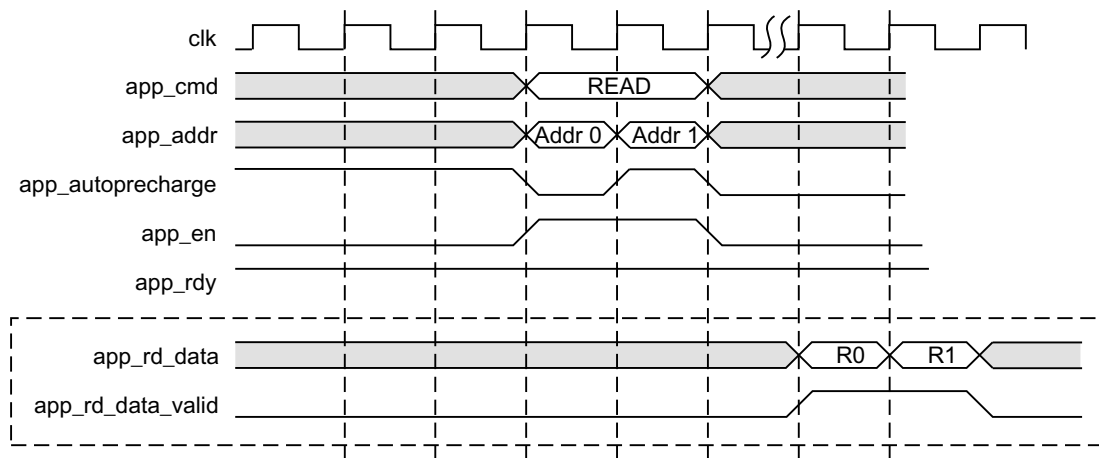


Figure 4-7: 4:1 Mode User Interface Read Timing Diagram (Memory Burst Type = BL8) #2

In Figure 4-7, the read data returned is always in the same order as the requests made on the address/control bus.

## Maintenance Commands

The UI can be configured by the Vivado IDE to enable two DRAM Refresh modes. The default mode configures the UI and the Memory Controller to automatically generate DRAM Refresh and ZQCS commands, meeting all DRAM protocol and timing requirements. The controller interrupts normal system traffic on a regular basis to issue these maintenance commands on the DRAM bus.

The User mode is enabled by checking the **Enable User Refresh and ZQCS Input** option in the Vivado IDE. In this mode, you are responsible for issuing Refresh and ZQCS commands at the rate required by the DRAM component specification after `init_calib_complete` asserts High. You use the `app_ref_req` and `app_zq_req` signals on the UI to request Refresh and ZQCS commands, and monitor `app_ref_ack` and `app_zq_ack` to know when the commands have completed. The controller manages all DRAM timing and protocol for these commands, other than the overall Refresh or ZQCS rate, just as it does for the default DRAM Refresh mode. These `request/ack` ports operate independently of the other UI command ports, like `app_cmd` and `app_en`.

The controller might not preserve the exact ordering of maintenance transactions presented to the UI on relative to regular read and write transactions. When you request a Refresh or ZQCS, the controller interrupts system traffic, just as in the default mode, and inserts the maintenance commands. To take the best advantage of this mode, you should request maintenance commands when the controller is idle or at least not very busy, keeping in mind that the DRAM Refresh rate and ZQCS rate requirements cannot be violated.

Figure 4-8 shows how the User mode ports are used and how they affect the DRAM command bus. This diagram shows the general idea about this mode of operation and is not timing accurate. Assuming the DRAM is idle with all banks closed, a short time after `app_ref_req` or `app_zq_req` are asserted High for one system clock cycle, the controller issues the requested commands on the DRAM command bus. The `app_ref_req` and `app_zq_req` can be asserted on the same cycle or different cycles, and they do not have to be asserted at the same rate. After a request signal is asserted High for one system clock, you must keep it deasserted until the acknowledge signal asserts.




Figure 4-8: User Mode Ports on DRAM Command Bus Timing Diagram

Figure 4-9 shows a case where the `app_en` is asserted and read transactions are presented continuously to the UI when the `app_ref_req` and `app_zq_req` are asserted. The controller interrupts the DRAM traffic following DRAM protocol and timing requirements, issues the Refresh and ZQCS, and then continues issuing the read transactions. Note that the `app_rdy` signal deasserts during this sequence. It is likely to deassert during a sequence like this since the controller command queue can easily fill up during `tRFC` or `tZQCS`. After the maintenance commands are issued and normal traffic resumes on the bus, the `app_rdy` signal asserts and new transactions are accepted again into the controller.




Figure 4-9: Read Transaction on User Interface Timing Diagram

Figure 4-9 shows the operation for a single rank. In a multi-rank system, a single refresh request generates a DRAM Refresh command to each rank, in series, staggered by  $tRFC/2$ . The Refresh commands are staggered since they are relatively high power consumption operations. A ZQCS command request generates a ZQCS command to all ranks in parallel.

## AXI4 Slave Interface

The AXI4 slave interface block maps AXI4 transactions to the UI to provide an industry-standard bus protocol interface to the Memory Controller. The AXI4 slave interface is optional in designs provided through the MIG tool. The RTL is consistent between both tools. For details on the AXI4 signaling protocol, see the ARM AMBA specifications [Ref 6].

The overall design is composed of separate blocks to handle each AXI channel, which allows for independent read and write transactions. Read and write commands to the UI rely on a simple round-robin arbiter to handle simultaneous requests. The address read/address write modules are responsible for chopping the AXI4 burst/wrap requests into smaller memory size burst lengths of either four or eight, and also conveying the smaller burst lengths to the read/write data modules so they can interact with the user interface.



If ECC is enabled, all write commands with any of the mask bits enabled are issued as read-modify-write operation.

Also if ECC is enabled, all write commands with none of the mask bits enabled are issued as write operation.

### AXI4 Slave Interface Parameters

Table 4-16 lists the AXI4 slave interface parameters.

Table 4-16: AXI4 Slave Interface Parameters

Parameter Name	Default Value	Allowable Values	Description
C_S_AXI_ADDR_WIDTH	32	32	This is the width of address read and address write signals. This value must be set to 32.
C_S_AXI_DATA_WIDTH	32	32, 64, 128, 256, 512	This is the width of data signals. Width of APP_DATA_WIDTH is recommended for better performance. Using a smaller width invokes an Upsizer, which would spend clocks in packing the data.
C_S_AXI_ID_WIDTH	4	1–16	This is the width of ID signals for every channel.
C_S_AXI_SUPPORTS_NARROW_BURST	1	0, 1	This parameter adds logic blocks to support narrow AXI transfers. It is required if any master connected to the Memory Controller issues narrow bursts. This parameter is automatically set if the AXI data width is smaller than the recommended value.
C_RD_WR_ARB_ALGORITHM	RD_PRI_REG	TDM, ROUND_ROBIN, RD_PRI_REG, RD_PRI_REG_STARVE_LIMIT, WRITE_PRIORITY_REG, WRITE_PRIORITY	This parameter indicates the Arbitration algorithm scheme. See <a href="#">Arbitration in AXI Shim</a> , page 108 for more information.

Table 4-16: AXI4 Slave Interface Parameters (Cont'd)

Parameter Name	Default Value	Allowable Values	Description
C_S_AXI_BASEADDR	–	Valid address	This parameter specifies the base address for the memory mapped slave interface. Address requests at this address map to rank 1, bank 0, row 0, column 0. The base/high address together define the accessible size of the memory. This accessible size must be a power of two. Additionally, the base/high address pair must be aligned to a multiple of the accessible size. The minimum accessible size is 4,096 bytes.
C_S_AXI_HIGHADDR	–	Valid address	This parameter specifies the high address for the memory mapped slave interface. Address requests received above this value wrap back to the base address. The base/high address together define the accessible size of the memory. This accessible size must be a power of two. Additionally, the base/high address pair must be aligned to a multiple of the accessible size. The minimum accessible size is 4,096 bytes.
C_S_AXI_PROTOCOL	AXI4	AXI4	This parameter specifies the AXI protocol.

### AXI4 Slave Interface Signals

Table 4-17 lists the AXI4 slave interface specific signal. `ui_clk` and `ui_clk_sync_rst` to the interface is provided from the Memory Controller. AXI interface is synchronous to `ui_clk`.

Table 4-17: AXI4 Slave Interface Signals

Name	Width	Direction	Active State	Description
<code>ui_clk</code>	1	Output		Output clock from the core to the interface.
<code>ui_clk_sync_rst</code>	1	Output	High	Output reset from the core to the interface.
<code>aresetn</code>	1	Input	Low	Input reset to the AXI Shim and it should be in synchronous with FPGA logic clock.
<code>s_axi_awid</code>	C_AXI_ID_WIDTH	Input		Write address ID

Table 4-17: AXI4 Slave Interface Signals (Cont'd)

Name	Width	Direction	Active State	Description
s_axi_awaddr	C_AXI_ADDR_WIDTH	Input		Write address
s_axi_awlen	8	Input		Burst length. The burst length gives the exact number of transfers in a burst.
s_axi_awsz	3	Input		Burst size. This signal indicates the size of each transfer in the burst.
s_axi_awburst	2	Input		Burst type
s_axi_awlock	1	Input		Lock type. (This is not used in the current implementation.)
s_axi_awcache	4	Input		Cache type. (This is not used in the current implementation.)
s_axi_awprot	3	Input		Protection type. (Not used in the current implementation.)
s_axi_awvalid	1	Input	High	Write address valid. This signal indicates that valid write address and control information are available.
s_axi_awready	1	Output	High	Write address ready. This signal indicates that the slave is ready to accept an address and associated control signals.
s_axi_wdata	C_AXI_DATA_WIDTH	Input		Write data
s_axi_wstrb	C_AXI_DATA_WIDTH/8	Input		Write strobes
s_axi_wlast	1	Input	High	Write last. This signal indicates the last transfer in a write burst.
s_axi_wvalid	1	Input	High	Write valid. This signal indicates that write data and strobe are available.
s_axi_wready	1	Output	High	Write ready
s_axi_bid	C_AXI_ID_WIDTH	Output		Response ID. The identification tag of the write response.
s_axi_bresp	2	Output		Write response. This signal indicates the status of the write response.
s_axi_bvalid	1	Output	High	Write response valid
s_axi_bready	1	Input	High	Response ready
s_axi_arid	C_AXI_ID_WIDTH	Input		Read address ID
s_axi_araddr	C_AXI_ADDR_WIDTH	Input		Read address
s_axi_arlen	8	Input		Read burst length
s_axi_arsz	3	Input		Read burst size
s_axi_arburst	2	Input		Read burst type
s_axi_arlock	1	Input		Lock type. (This is not used in the current implementation.)

Table 4-17: AXI4 Slave Interface Signals (Cont'd)

Name	Width	Direction	Active State	Description
s_axi_arcache	4	Input		Cache type. (This is not used in the current implementation.)
s_axi_arprot	3	Input		Protection type. (This is not used in the current implementation.)
s_axi_arvalid	1	Input	High	Read address valid
s_axi_arready	1	Output	High	Read address ready
s_axi_rid	C_AXI_ID_WIDTH	Output		Read ID tag
s_axi_rdata	C_AXI_DATA_WIDTH	Output		Read data
s_axi_rresp	2	Output		Read response
s_axi_rlast	1	Output		Read last
s_axi_rvalid	1	Output		Read valid
s_axi_rready	1	Input		Read ready
dbg_clk	1	Output		Debug Clock. Do not connect any signals to dbg_clk and keep the port open during instantiation.

### Arbitration in AXI Shim

The AXI4 protocol calls for independent read and write address channels. The Memory Controller has one address channel. The following arbitration options are available for arbitrating between the read and write address channels.

#### Time Division Multiplexing (TDM)

Equal priority is given to read and write address channels in this mode. The grant to the read and write address channels alternate every clock cycle. The read or write requests from the AXI master has no bearing on the grants. For example, the read requests are served in alternative clock cycles, even when there are no write requests. The slots are fixed and they are served in their respective slots only.

#### Round-Robin

Equal priority is given to read and write address channels in this mode. The grant to the read and write channels depends on the last served request granted from the AXI master. For example, if the last performed operation is write, then it gives precedence for read operation to be served over write operation. Similarly, if the last performed operation is read, then it gives precedence for write operation to be served over read operation. If both read and write channels requests at the same time when there are no pending requests, this scheme serves write channel ahead of read.

### Read Priority (RD\_PRI\_REG)

Read and write address channels are served with equal priority in this mode. The requests from the write address channel are processed when one of the following occurs:

- No pending requests from read address channel.
- Read starve limit of 256 is reached. It is only checked at the end of the burst.
- Read wait limit of 16 is reached.
- Write Quality of Service (QoS) is higher which is non-zero. It is only checked at the end of the burst.

The requests from the read address channel are processed in a similar method.

### Read Priority with Starve Limit (RD\_PRI\_REG\_STARVE\_LIMIT)

The read address channel is always given priority in this mode. The requests from the write address channel are processed when there are no pending requests from the read address channel or the starve limit for read is reached.

### Write Priority (WRITE\_PRIORITY, WRITE\_PRIORITY\_REG)

Write address channel is always given priority in this mode. The requests from the read address channel are processed when there are no pending requests from the write address channel. Arbitration outputs are registered in WRITE\_PRIORITY\_REG mode.

## ***AXI4-Lite Slave Control/Status Register Interface Block***

The AXI4-Lite Slave Control register block provides a processor accessible interface to the ECC memory option. The interface is available when ECC is enabled and the primary slave interface is AXI4. The block provides interrupts, interrupt enable, ECC status, ECC enable/disable, ECC correctable errors counter, first failing correctable/uncorrectable data, ECC, and address. Fault injection registers for software testing is provided when the ECC\_TEST\_FI\_XOR (C\_ECC\_TEST) parameter is "ON." The AXI4-Lite interface is fixed at 32 data bits and signaling follows the standard AMBA AXI4-Lite specifications [Ref 6].

The AXI4-Lite Control/Status register interface block is implemented in parallel to the AXI4 memory-mapped interface. The block monitors the output of the native interface to capture correctable (single bit) and uncorrectable (multiple bit) errors. When a correctable and/or uncorrectable error occurs, the interface also captures the byte address of the failure along with the failing data bits and ECC bits. Fault injection is provided by an XOR block placed in the write datapath after the ECC encoding has occurred. Only the first memory beat in a transaction can have errors inserted. For example, in a memory configuration with a data width of 72 and a mode register set to burst length 8, only the first 72 bits are corruptible through the fault injection interface. Interrupt generation based on either a correctable or uncorrectable error can be independently configured with the register interface.

## ECC Enable/Disable

The ECC\_ON\_OFF register enables/disables the ECC decode functionality. However, encoding is always enabled. The default value at start-up can be parameterized with C\_ECC\_ONOFF\_RESET\_VALUE. Assigning a value of 1 for the ECC\_ON\_OFF bit of this register results in the `correct_en` signal input into the `mem_intf` to be asserted. Writing a value of 0 to the ECC\_ON\_OFF bit of this register results in the `correct_en` signal to be deasserted. When `correct_en` is asserted, decoding is enabled, and the opposite is true when this signal is deasserted. ECC\_STATUS/ECC\_CE\_CNT are not updated when ECC\_ON\_OFF = 0. The FI\_D0, FI\_D1, FI\_D2, and FI\_D3 registers are not writable when ECC\_ON\_OFF = 0.

## Single Error and Double Error Reporting

Two vectored signals from the Memory Controller indicate an ECC error: `ecc_single` and `ecc_multiple`. The `ecc_single` signal indicates if there has been a correctable error and the `ecc_multiple` signal indicates if there has been an uncorrectable error. The widths of `ecc_multiple` and `ecc_single` are based on the C\_NCK\_PER\_CLK parameter. There can be between 0 and C\_NCK\_PER\_CLK × 2 errors per cycle with each data beat signaled by one of the vector bits. Multiple bits of the vector can be signaled per cycle indicating that multiple correctable errors or multiple uncorrectable errors have been detected. The `ecc_err_addr` signal (discussed in [Fault Collection](#)) is valid during the assertion of either `ecc_single` or `ecc_multiple`.

The ECC\_STATUS register sets the CE\_STATUS bit and/or UE\_STATUS bit for correctable error detection and uncorrectable error detection, respectively.




---

**CAUTION!** *Multiple bit error is a serious failure of memory because it is uncorrectable. In such cases, application cannot rely on contents of the memory. It is suggested to not perform any further transactions to memory.*

---

## Interrupt Generation

When interrupts are enabled with the CE\_EN\_IRQ and/or UE\_EN\_IRQ bits of the ECC\_EN\_IRQ register, and a correctable error or uncorrectable error occurs, the interrupt signal is asserted.

## Fault Collection

To aid the analysis of ECC errors, there are two banks of storage registers that collect information on the failing ECC decode. One bank of registers is for correctable errors, and another bank is for uncorrectable errors. The failing address, undecoded data, and ECC bits are saved into these register banks as CE\_FFA, CE\_FFD, and CE\_FFE for correctable errors. UE\_FFA, UE\_FFD, and UE\_FFE are for uncorrectable errors. The data in combination with the ECC bits can help determine which bit(s) have failed. CE\_FFA stores the address from the `ecc_err_addr` signal and converts it to a byte address. Upon error detection, the data is latched into the appropriate register. Only the first data beat with an error is stored.

When a correctable error occurs, there is also a counter that counts the number of correctable errors that have occurred. The counter can be read from the CE\_CNT register and is fixed as an 8-bit counter. It does not rollover when the maximum value is incremented.

### Fault Injection

The ECC Fault Injection registers, FI\_D and FI\_ECC, facilitates testing of the software drivers. When set, the ECC Fault Injection register XORs with the MIG DFI datapath to simulate errors in the memory. The DFI interface lies between the Memory Controller and the PHY. It is ideal for injection to occur here because this is after the encoding has been completed. There is only support to insert errors on the first data beat, therefore there are two to four FI\_D registers to accommodate this. During operation, after the error has been inserted into the datapath, the register clears itself.

### AXI4-Lite Slave Control/Status Register Interface Parameters

Table 4-18 lists the AXI4-Lite slave interface parameters.

Table 4-18: AXI4-Lite Slave Control/Status Register Parameters

Parameter Name	Default Value	Allowable Values	Description
C_S_AXI_CTRL_ADDR_WIDTH	32	32, 64	This is the width of the AXI4-Lite address buses.
C_S_AXI_CTRL_DATA_WIDTH	32	32	This is the width of the AXI4-Lite data buses.
C_ECC_ONOFF_RESET_VALUE	1	0, 1	Controls ECC on/off value at startup/reset.
C_S_AXI_CTRL_BASEADDR	–	Valid Address	This parameter specifies the base address for the AXI4-Lite slave interface.
C_S_AXI_CTRL_HIGHADDR	–	Valid Address	This parameter specifies the high address for the AXI4-Lite slave interface.
C_S_AXI_CTRL_PROTOCOL	AXI4LITE	AXI4LITE	AXI4-Lite protocol

### AXI4-Lite Slave Control/Status Register Interface Signals

Table 4-19 lists the AXI4 slave interface specific signals. Clock/reset to the interface is provided from the Memory Controller.

Table 4-19: List of New I/O Signals

Name	Width	Direction	Active State	Description
s_axi_ctrl_awaddr	C_S_AXI_CTRL_ADDR_WIDTH	Input		Write address
s_axi_ctrl_awvalid	1	Input	High	Write address valid. This signal indicates that valid write address and control information are available.

Table 4-19: List of New I/O Signals (Cont'd)

Name	Width	Direction	Active State	Description
s_axi_ctrl_awready	1	Output	High	Write address ready. This signal indicates that the slave is ready to accept an address and associated control signals.
s_axi_ctrl_wdata	C_S_AXI_CTRL_DATA_WIDTH	Input		Write data
s_axi_ctrl_wvalid	1	Input	High	Write valid. This signal indicates that write data and strobe are available.
s_axi_ctrl_wready	1	Output	High	Write ready
s_axi_ctrl_bvalid	1	Output	High	Write response valid
s_axi_ctrl_bready	1	Input	High	Response ready
s_axi_ctrl_araddr	C_S_AXI_CTRL_ADDR_WIDTH	Input		Read address
s_axi_ctrl_arvalid	1	Input	High	Read address valid
s_axi_ctrl_arready	1	Output	High	Read address
s_axi_ctrl_rdata	C_S_AXI_CTRL_DATA_WIDTH	Output		Read data
s_axi_ctrl_rvalid	1	Output		Read valid
s_axi_ctrl_rready	1	Input		Read ready
interrupt	1	Output	High	IP Global Interrupt signal

### AXI4-Lite Slave Control/Status Register Map

ECC register map is shown in Table 4-20. The register map is Little Endian. Write accesses to read-only or reserved values are ignored. Read accesses to write-only or reserved values return the value 0xDEADDEAD.

Table 4-20: ECC Control Register Map

Address Offset	Register Name	Access Type	Default Value	Description
0x00	ECC_STATUS	R/W	0x0	ECC Status Register
0x04	ECC_EN_IRQ	R/W	0x0	ECC Enable Interrupt Register
0x08	ECC_ON_OFF	R/W	0x0 or 0x1	ECC On/Off Register. If C_ECC_ONOFF_RESET_VALUE = 1, the default value is 0x1.
0x0C	CE_CNT	R/W	0x0	Correctable Error Count Register
<b>(0x10–0x9C) Reserved</b>				
0x100	CE_FFD[31:00]	R	0x0	Correctable Error First Failing Data Register
0x104	CE_FFD[63:32]	R	0x0	Correctable Error First Failing Data Register
0x108	CE_FFD[95:64] <sup>(1)</sup>	R	0x0	Correctable Error First Failing Data Register
0x10C	CE_FFD [127:96] <sup>(1)</sup>	R	0x0	Correctable Error First Failing Data Register



Table 4-20: ECC Control Register Map (Cont'd)

Address Offset	Register Name	Access Type	Default Value	Description
<b>(0x110–0x17C) Reserved</b>				
0x180	CE_FFE	R	0x0	Correctable Error First Failing ECC Register
<b>(0x184–0x1BC) Reserved</b>				
0x1C0	CE_FFA[31:0]	R	0x0	Correctable Error First Failing Address
0x1C4	CE_FFA[63:32] <sup>(2)</sup>	R	0x0	Correctable Error First Failing Address
<b>(0x1C8–0x1FC) Reserved</b>				
0x200	UE_FFD [31:00]	R	0x0	Uncorrectable Error First Failing Data Register
0x204	UE_FFD [63:32]	R	0x0	Uncorrectable Error First Failing Data Register
0x208	UE_FFD [95:64] <sup>(1)</sup>	R	0x0	Uncorrectable Error First Failing Data Register
0x20C	UE_FFD [127:96] <sup>(1)</sup>	R	0x0	Uncorrectable Error First Failing Data Register
<b>(0x210–0x27C) Reserved</b>				
0x280	UE_FFE	R	0x0	Uncorrectable Error First Failing ECC Register
<b>(0x284–0x2BC) Reserved</b>				
0x2C0	UE_FFA[31:0]	R	0x0	Uncorrectable Error First Failing Address
0x2C4	UE_FFA[63:32] <sup>(2)</sup>	R	0x0	Uncorrectable Error First Failing Address
<b>(0x2C8–0x2FC) Reserved</b>				
0x300	FI_D[31:0] <sup>(3)</sup>	W	0x0	Fault Inject Data Register
0x304	FI_D[63:32] <sup>(3)</sup>	W	0x0	Fault Inject Data Register
0x308	FI_D[95:64] <sup>(1)(3)</sup>	W	0x0	Fault Inject Data Register
0x30C	FI_D[127:96] <sup>(1)(3)</sup>	W	0x0	Fault Inject Data Register
<b>(0x340–0x37C) Reserved</b>				
0x380	FI_ECC <sup>(3)</sup>	W	0x0	Fault Inject ECC Register

**Notes:**

1. Data bits 64–127 are only enabled if the DQ width is 144 bits.
2. Reporting address bits 63–32 are only available if the address map is > 32 bits.
3. FI\_D\* and FI\_ECC\* are only enabled if ECC\_TEST parameter has been set to 1.

## AXI4-Lite Slave Control/Status Register Map Detailed Descriptions

### ECC\_STATUS

This register holds information on the occurrence of correctable and uncorrectable errors. The status bits are independently set to 1 for the first occurrence of each error type. The status bits are cleared by writing a 1 to the corresponding bit position; that is, the status bits can only be cleared to 0 and not set to 1 using a register write. The ECC Status register operates independently of the ECC Enable Interrupt register.

Table 4-21: ECC Status Register

Bits	Name	Core Access	Reset Value	Description
1	CE_STATUS	R/W	0	If 1, a correctable error has occurred. This bit is cleared when a 1 is written to this bit position.
0	UE_STATUS	R/W	0	If 1, an uncorrectable error has occurred. This bit is cleared when a 1 is written to this bit position.

### ECC\_EN\_IRQ

This register determines if the values of the CE\_STATUS and UE\_STATUS bits in the ECC Status register assert the Interrupt output signal (ECC\_Interrupt). If both CE\_EN\_IRQ and UE\_EN\_IRQ are set to 1 (enabled), the value of the Interrupt signal is the logical OR between the CE\_STATUS and UE\_STATUS bits.

Table 4-22: ECC Interrupt Enable Register

Bits	Name	Core Access	Reset Value	Description
1	CE_EN_IRQ	R/W	0	If 1, the value of the CE_STATUS bit of ECC Status register is propagated to the Interrupt signal. If 0, the value of the CE_STATUS bit of ECC Status register is not propagated to the Interrupt signal.
0	UE_EN_IRQ	R/W	0	If 1, the value of the UE_STATUS bit of ECC Status register is propagated to the Interrupt signal. If 0, the value of the UE_STATUS bit of ECC Status register is not propagated to the Interrupt signal.

### ECC\_ON\_OFF

The ECC On/Off Control register allows the application to enable or disable ECC checking. The design parameter, C\_ECC\_ONOFF\_RESET\_VALUE (default on) determines the reset value for the enable/disable setting of ECC. This facilitates start-up operations when ECC might or might not be initialized in the external memory. When disabled, ECC checking is disabled for read but ECC generation is active for write operations.

Table 4-23: ECC On/Off Control Register

Bits	Name	Core Access	Reset Value	Description
0	ECC_ON_OFF	R/W	Specified by design parameter, C_ECC_ONOFF_RESET_VALUE	If 0, ECC checking is disabled on read operations. (ECC generation is enabled on write operations when C_ECC = 1). If 1, ECC checking is enabled on read operations. All correctable and uncorrectable error conditions are captured and status is updated.

## CE\_CNT

This register counts the number of occurrences of correctable errors. It can be cleared or preset to any value using a register write. When the counter reaches its maximum value, it does not wrap around, but instead it stops incrementing and remains at the maximum value. The width of the counter is defined by the value of the C\_CE\_COUNTER\_WIDTH parameter. The value of the CE counter width is fixed to eight bits.

Table 4-24: Correctable Error Counter Register

Bits	Name	Core Access	Reset Value	Description
7:0	CE_CNT	R/W	0	Holds the number of correctable errors encountered.

## CE\_FFA[31:0]

This register stores the lower 32 bits of the decoded DRAM address (Bits[31:0]) of the first occurrence of an access with a correctable error. The address format is defined in [Table 3-1, page 25](#). When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the address of the next correctable error. Storing of the failing address is enabled after reset.

Table 4-25: Correctable Error First Failing Address [31:0] Register

Bits	Name	Core Access	Reset Value	Description
31:0	CE_FFA[31:0]	R	0	Address (Bits[31:0]) of the first occurrence of a correctable error.

## CE\_FFA[63:32]

This register stores the upper 32 bits of the decoded DRAM address (Bits[55:32]) of the first occurrence of an access with a correctable error. The address format is defined in [Table 3-1, page 25](#). In addition, the upper byte of this register stores the ecc\_single signal. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the address of the next correctable error. Storing of the failing address is enabled after reset.

Table 4-26: Correctable Error First Failing Address [63:32] Register

Bits	Name	Core Access	Reset Value	Description
31:24	CD_FFA[63:56]	R	0	ecc_single[7:0]. Indicates which bursts of the BL8 transaction associated with the logged address had a correctable error. Bit[24] corresponds to the first burst of the BL8 transfer.
23:0	CD_FFA[55:32]	R	0	Address (Bits[55:32]) of the first occurrence of a correctable error.

### CE\_FFD[31:0]

This register stores the (corrected) failing data (Bits[31:0]) of the first occurrence of an access with a correctable error. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next correctable error. Storing of the failing data is enabled after reset.

Table 4-27: Correctable Error First Failing Data [31:0] Register

Bits	Name	Core Access	Reset Value	Description
31:0	CE_FFD[31:0]	R	0	Data (Bits[31:0]) of the first occurrence of a correctable error.

### CE\_FFD[63:32]

This register stores the (corrected) failing data (Bits[63:32]) of the first occurrence of an access with a correctable error. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next correctable error. Storing of the failing data is enabled after reset.

Table 4-28: Correctable Error First Failing Data [63:32] Register

Bits	Name	Core Access	Reset Value	Description
31:0	CE_FFD[63:32]	R	0	Data (Bits[63:32]) of the first occurrence of a correctable error.

### CE\_FFD[95:64]

**Note:** This register is only used when DQ\_WIDTH == 144.

This register stores the (corrected) failing data (Bits[95:64]) of the first occurrence of an access with a correctable error. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next correctable error. Storing of the failing data is enabled after reset.

Table 4-29: Correctable Error First Failing Data [95:64] Register

Bits	Name	Core Access	Reset Value	Description
31:0	CE_FFD[95:64]	R	0	Data (Bits[95:64]) of the first occurrence of a correctable error.

### CE\_FFD[127:96]

**Note:** This register is only used when DQ\_WIDTH == 144.

This register stores the (corrected) failing data (Bits[127:96]) of the first occurrence of an access with a correctable error. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next correctable error. Storing of the failing data is enabled after reset.

Table 4-30: Correctable Error First Failing Data [127:96] Register

Bits	Name	Core Access	Reset Value	Description
31:0	CE_FFD [127:96]	R	0	Data (Bits[127:96]) of the first occurrence of a correctable error.

### CE\_FFE

This register stores the ECC bits of the first occurrence of an access with a correctable error. When the CE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the ECC of the next correctable error. Storing of the failing ECC is enabled after reset.

Table 4-31 describes the register bit usage when DQ\_WIDTH = 72.

Table 4-31: Correctable Error First Failing ECC Register for 72-Bit External Memory Width

Bits	Name	Core Access	Reset Value	Description
7:0	CE_FFE	R	0	ECC (Bits[7:0]) of the first occurrence of a correctable error.

Table 4-32 describes the register bit usage when DQ\_WIDTH = 144.

Table 4-32: Correctable Error First Failing ECC Register for 144-Bit External Memory Width

Bits	Name	Core Access	Reset Value	Description
15:0	CE_FFE	R	0	ECC (Bits[15:0]) of the first occurrence of a correctable error.

### UE\_FFA[31:0]

This register stores the decoded DRAM address (Bits[31:0]) of the first occurrence of an access with an uncorrectable error. The address format is defined in Table 3-1, page 25. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the address of the next uncorrectable error. Storing of the failing address is enabled after reset.

Table 4-33: Uncorrectable Error First Failing Address [31:0] Register

Bits	Name	Core Access	Reset Value	Description
31:0	UE_FFA [31:0]	R	0	Address (Bits[31:0]) of the first occurrence of an uncorrectable error.

### UE\_FFA[63:32]

This register stores the decoded address (Bits[55:32]) of the first occurrence of an access with an uncorrectable error. The address format is defined in Table 3-1, page 25. In addition, the upper byte of this register stores the ecc\_multiple signal. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the address of the next uncorrectable error. Storing of the failing address is enabled after reset.

Table 4-34: Uncorrectable Error First Failing Address [31:0] Register

Bits	Name	Core Access	Reset Value	Description
31:24	CD_FFA[63:56]	R	0	ecc_multiple[7:0]. Indicates which bursts of the BL8 transaction associated with the logged address had an uncorrectable error. Bit[24] corresponds to the first burst of the BL8 transfer.
23:0	CD_FFA[55:32]	R	0	Address (Bits[55:32]) of the first occurrence of a correctable error.

### UE\_FFD[31:0]

This register stores the (uncorrected) failing data (Bits[31:0]) of the first occurrence of an access with an uncorrectable error. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next uncorrectable error. Storing of the failing data is enabled after reset.

Table 4-35: Uncorrectable Error First Failing Data [31:0] Register

Bits	Name	Core Access	Reset Value	Description
31:0	UE_FFD[31:0]	R	0	Data (Bits[31:0]) of the first occurrence of an uncorrectable error.

### UE\_FFD[63:32]

This register stores the (uncorrected) failing data (Bits[63:32]) of the first occurrence of an access with an uncorrectable error. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next uncorrectable error. Storing of the failing data is enabled after reset.

Table 4-36: Uncorrectable Error First Failing Data [63:32] Register

Bits	Name	Core Access	Reset Value	Description
31:0	UE_FFD [63:32]	R	0	Data (Bits[63:32]) of the first occurrence of an uncorrectable error.

### UE\_FFD[95:64]

**Note:** This register is only used when the DQ\_WIDTH == 144.

This register stores the (uncorrected) failing data (Bits[95:64]) of the first occurrence of an access with an uncorrectable error. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next uncorrectable error. Storing of the failing data is enabled after reset.

Table 4-37: Uncorrectable Error First Failing Data [95:64] Register

Bits	Name	Core Access	Reset Value	Description
31:0	UE_FFD[95:64]	R	0	Data (Bits[95:64]) of the first occurrence of an uncorrectable error.

## UE\_FFD[127:96]

**Note:** This register is only used when the DQ\_WIDTH == 144.

This register stores the (uncorrected) failing data (Bits[127:96]) of the first occurrence of an access with an uncorrectable error. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the data of the next uncorrectable error. Storing of the failing data is enabled after reset.

Table 4-38: Uncorrectable Error First Failing Data [127:96] Register

Bits	Name	Core Access	Reset Value	Description
31:0	UE_FFD[127:96]	R	0	Data (Bits[127:96]) of the first occurrence of an uncorrectable error.

## UE\_FFE

This register stores the ECC bits of the first occurrence of an access with an uncorrectable error. When the UE\_STATUS bit in the ECC Status register is cleared, this register is re-enabled to store the ECC of the next uncorrectable error. Storing of the failing ECC is enabled after reset.

Table 4-39 describes the register bit usage when DQ\_WIDTH = 72.

Table 4-39: Uncorrectable Error First Failing ECC Register for 72-Bit External Memory Width

Bits	Name	Core Access	Reset Value	Description
7:0	UE_FFE	R	0	ECC (Bits[7:0]) of the first occurrence of an uncorrectable error.

Table 4-40 describes the register bit usage when DQ\_WIDTH = 144.

Table 4-40: Uncorrectable Error First Failing ECC Register for 144-Bit External Memory Width

Bits	Name	Core Access	Reset Value	Description
15:0	UE_FFE	R	0	ECC (Bits[15:0]) of the first occurrence of an uncorrectable error.

## FI\_D0

This register is used to inject errors in data (Bits[31:0]) written to memory and can be used to test the error correction and error signaling. The bits set in the register toggle the corresponding data bits (word 0 or Bits[31:0]) of the subsequent data written to the memory without affecting the ECC bits written. After the fault has been injected, the Fault Injection Data register is cleared automatically.

The register is only implemented if C\_ECC\_TEST = "ON" or ECC\_TEST\_FI\_XOR = "ON" and ECC = "ON" in a MIG design in the Vivado IP catalog.

Injecting faults should be performed in a critical region in software; that is, writing this register and the subsequent write to the memory must not be interrupted.

**Table 4-41: Fault Injection Data (Word 0) Register**

Bits	Name	Core Access	Reset Value	Description
31:0	FI_D0	W	0	Bit positions set to 1 toggle the corresponding Bits[31:0] of the next data word written to the memory. This register is automatically cleared after the fault has been injected.

Special consideration must be given across FI\_D0, FI\_D1, FI\_D2, and FI\_D3 such that only a single error condition is introduced.

### FI\_D1

This register is used to inject errors in data (Bits[63:32]) written to memory and can be used to test the error correction and error signaling. The bits set in the register toggle the corresponding data bits (word 1 or Bits[63:32]) of the subsequent data written to the memory without affecting the ECC bits written. After the fault has been injected, the Fault Injection Data register is cleared automatically.

This register is only implemented if C\_ECC\_TEST = "ON" or ECC\_TEST\_FI\_XOR = "ON" and ECC = "ON" in a MIG design in the Vivado IP catalog.

Injecting faults should be performed in a critical region in software; that is, writing this register and the subsequent write to the memory must not be interrupted.

**Table 4-42: Fault Injection Data (Word 1) Register**

Bits	Name	Core Access	Reset Value	Description
31:0	FI_D1	W	0	Bit positions set to 1 toggle the corresponding Bits[63:32] of the next data word written to the memory. This register is automatically cleared after the fault has been injected.



## FI\_D2

**Note:** This register is only used when DQ\_WIDTH =144.

This register is used to inject errors in data (Bits[95:64]) written to memory and can be used to test the error correction and error signaling. The bits set in the register toggle the corresponding data bits (word 2 or Bits[95:64]) of the subsequent data written to the memory without affecting the ECC bits written. After the fault has been injected, the Fault Injection Data register is cleared automatically.

This register is only implemented if C\_ECC\_TEST = "ON" or ECC\_TEST\_FI\_XOR = "ON" and ECC = "ON" in a MIG design in the Vivado IP catalog.

Injecting faults should be performed in a critical region in software; that is, writing this register and the subsequent write to the memory must not be interrupted.

**Table 4-43: Fault Injection Data (Word 2) Register**

Bits	Name	Core Access	Reset Value	Description
31:0	FI_D2	W	0	Bit positions set to 1 toggle the corresponding Bits[95:64] of the next data word written to the memory. This register is automatically cleared after the fault has been injected.

Special consideration must be given across FI\_D0, FI\_D1, FI\_D2, and FI\_D3 such that only a single error condition is introduced.

## FI\_D3

**Note:** This register is only used when DQ\_WIDTH =144.

This register is used to inject errors in data (Bits[127:96]) written to memory and can be used to test the error correction and error signaling. The bits set in the register toggle the corresponding data bits (word 3 or Bits[127:96]) of the subsequent data written to the memory without affecting the ECC bits written. After the fault has been injected, the Fault Injection Data register is cleared automatically.

The register is only implemented if C\_ECC\_TEST = "ON" or ECC\_TEST\_FI\_XOR = "ON" and ECC = "ON" in a MIG design in the Vivado IP catalog.

Injecting faults should be performed in a critical region in software; that is, writing this register and the subsequent write to the memory must not be interrupted.

**Table 4-44: Fault Injection Data (Word 3) Register**

Bits	Name	Core Access	Reset Value	Description
31:0	FI_D3	W	0	Bit positions set to 1 toggle the corresponding Bits[127:96] of the next data word written to the memory. The register is automatically cleared after the fault has been injected.

## FI\_ECC

This register is used to inject errors in the generated ECC written to the memory and can be used to test the error correction and error signaling. The bits set in the register toggle the corresponding ECC bits of the next data written to memory. After the fault has been injected, the Fault Injection ECC register is cleared automatically.

The register is only implemented if C\_ECC\_TEST = "ON" or ECC\_TEST\_FI\_XOR = "ON" and ECC = "ON" in a MIG design in the Vivado IP catalog.

Injecting faults should be performed in a critical region in software; that is, writing this register and the subsequent write to memory must not be interrupted.

Table 4-45 describes the register bit usage when DQ\_WIDTH = 72.

**Table 4-45: Fault Injection ECC Register for 72-Bit External Memory Width**

Bits	Name	Core Access	Reset Value	Description
7:0	FI_ECC	W	0	Bit positions set to 1 toggle the corresponding bit of the next ECC written to the memory. The register is automatically cleared after the fault has been injected.

Table 4-46 describes the register bit usage when DQ\_WIDTH = 144.

**Table 4-46: Fault Injection ECC Register for 144-Bit External Memory Width**

Bits	Name	Core Access	Reset Value	Description
15:0	FI_ECC	R	0	Bit positions set to 1 toggle the corresponding bit of the next ECC written to the memory. The register is automatically cleared after the fault has been injected.

## PHY Only Interface

This section describes the FPGA logic interface signals and key parameters of the DDR3 and DDR4 PHY. The goal is to implement a "PHY Only" solution that connects your own custom Memory Controller directly to the MIG generated PHY, instead of interfacing to the user interface or AXI Interface of a MIG generated Memory Controller. The PHY interface takes DRAM commands, like Activate, Precharge, Refresh, etc. at its input ports and issues them directly to the DRAM bus.

The PHY does not take in "memory transactions" like the user and AXI interfaces, which translate transactions into one or more DRAM commands that meet DRAM protocol and timing requirements. The PHY interface does no DRAM protocol or timing checking. When using a PHY Only option, you are responsible for meeting all DRAM protocol requirements and timing specifications of all DRAM components in the system.

The PHY runs at the system clock frequency, or 1/4 of the DRAM clock frequency. The PHY therefore accepts four DRAM commands per system clock and issues them serially on consecutive DRAM clock cycles on the DRAM bus. In other words, the PHY interface has four command slots: slots 0, 1, 2, and 3, which it accepts each system clock. The command in slot position 0 is issued on the DRAM bus first, and the command in slot 3 is issued last. The PHY does have limitations as to which slots can accept read and write CAS commands. For more information, see [CAS Command Timing Limitations, page 141](#). Except for CAS commands, each slot can accept arbitrary DRAM commands.

The PHY FPGA logic interface has an input port for each pin on a DDR3 or DDR4 bus. Each PHY command/address input port has a width that is eight times wider than its corresponding DRAM bus pin. For example, a DDR4 bus has one `act_n` pin, and the PHY has an 8-bit `mc_ACT_n` input port. Each pair of bits in the `mc_ACT_n` port corresponds to a "command slot." The two LSBs are slot0 and the two MSBs are slot3. The PHY address input port for a DDR4 design with 18 address pins is 144 bits wide, with each byte corresponding to the four command slots for one DDR4 address pin. There are two bits for each command slot in each input port of the PHY. This is due to the underlying design of the PHY and its support for double data rate data buses. But as the DRAM command/address bus is single data rate, you must always drive the two bits that correspond to a command slot to the same value. See the following interface tables for additional descriptions and examples in the timing diagrams that show how bytes and bits correspond to DRAM pins and command slots.

The PHY interface has read and write data ports with eight bits for each DRAM DQ pin. Each port bit represents one data bit on the DDR DRAM bus for a BL8 burst. Therefore one BL8 data burst for the entire DQ bus is transferred across the PHY interface on each system clock. The PHY only supports BL8 data transfers. The data format is the same as the user interface data format. For more information, see [PHY, page 26](#).

The PHY interface also has several control signals that you must drive and/or respond to when a read or write CAS command is issued. The control signals are used by the PHY to manage the transfer of read and write data between the PHY interface and the DRAM bus. See the following signal tables and timing diagrams.

Your custom Memory Controller must wait until the PHY output `calDone` is asserted before sending any DRAM commands to the PHY. The PHY initializes and trains the DRAM before asserting `calDone`. For more information on the PHY internal structures and training algorithms, see the [PHY, page 26](#). After `calDone` is asserted, the PHY is ready to accept any DRAM commands. The only required DRAM or PHY commands are related to VT tracking and DRAM refresh/ZQ. These requirements are detailed in [VT Tracking, page 143](#) and [Refresh and ZQ, page 144](#).

## ***PHY Interface Signals***

The PHY interface signals to the FPGA logic can be categorized into six groups:

- [Clocking and Reset](#)
- [Command and Address](#)
- [Write Data](#)
- [Read Data](#)
- [PHY Control](#)
- [Debug](#)

Clocking and Reset and Debug signals are described in other sections or documents. See the corresponding references. In this section, a description is given for each signal in the remaining four groups and timing diagrams show examples of the signals in use.

### **Clocking and Reset**

For more information on the clocking and reset, see the [Clocking, page 73](#) section.

### **Command and Address**

[Table 4-47](#) shows the command and address signals for a PHY only option.

Table 4-47: Command and Address

Signal	Direction	Description
mc_ACT_n[7:0]	Input	DRAM ACT_n command signal for four DRAM clock cycles. Bits[1:0] correspond to the first DRAM clock cycle, Bits[3:2] to the second, Bits[5:4] to the third, and Bits[8:7] to the fourth. For center alignment to the DRAM clock with 1N timing, both bits of a given bit pair should be asserted to the same value. See timing diagrams for examples. All of the command/address ports in this table follow the same eight bits per DRAM pin format. Active-Low. This signal is not used in DDR3 systems.
mc_ADR[ADDR_WIDTH × 8 – 1:0]	Input	DRAM address. Eight bits in the PHY interface for each address bit on the DRAM bus. Bits[7:0] corresponds to DRAM address Bit[0] on four DRAM clock cycles. Bits[15:8] corresponds to DRAM address Bit[1] on four DRAM clock cycles, etc. See the timing diagrams for examples. All of the multi-bit DRAM signals in this table follow the same format of 1-byte of the PHY interface port corresponding to four commands for one DRAM pin. Mixed active-Low and High depending on which type of DRAM command is being issued, but follows the DRAM pin active-High/Low behavior. The function of each byte of the mc_ADR port depends on whether the memory type is DDR4 or DDR3 and the particular DRAM command that is being issued. These functions match the DRAM address pin functions. For example, with DDR4 memory and the mc_ACT_n port bits asserted High, mc_ADR[135:112] have the function of RAS_n, CAS_n, and WE_n pins.
mc_RAS_n[7:0]	Input	DDR3 DRAM RAS_n pin. Not used in DDR4 systems.
mc_CAS_n[7:0]	Input	DDR3 DRAM CAS_n pin. Not used in DDR4 systems.
mc_WE_n[7:0]	Input	DDR3 DRAM WE_n pin. Not used in DDR4 systems.
mc_BA[BANK_WIDTH × 8 – 1:0]	Input	DRAM bank address. Eight bits for each DRAM bank address.
mc_BG[BANK_GROUP_WIDTH × 8 – 1:0]	Input	DRAM bank group address. Eight bits for each DRAM pin.
mc_CKE[CKE_WIDTH × 8 – 1:0]	Input	DRAM CKE. Eight bits for each DRAM pin.
mc_CS_n[CS_WIDTH × 8 – 1:0]	Input	DRAM CS_n. Eight bits for each DRAM pin. Active-Low.
mc_ODT[ODT_WIDTH × 8 – 1:0]	Input	DRAM ODT. Eight bits for each DRAM pin. Active-High.
mc_PAR[7:0]	Input	DRAM address parity. Eight bits for one DRAM parity pin. <b>Note:</b> This signal is valid for RDIMMs only.

Figure 4-10 shows the functional relationship between the PHY command/address input signals and a DDR4 command/address bus. The diagram shows an Activate command on system clock cycle N in the slot1 position. The `mc_ACT_n[3:2]` and `mc_CS_n[3:2]` are both asserted Low in cycle N, and all the other bits in cycle N are asserted High, generating an Activate in the slot1 position roughly two system clocks later and NOP/DESELECT commands on the other command slots.

On cycle N + 3, `mc_CS_n` and the `mc_ADR` bits corresponding to CAS/A15 are set to 0xFC. This asserts `mc_ADR[121:120]` and `mc_CS_n[1:0]` Low, and all other bits in cycle N + 3 High, generating a read command on slot0 and NOP/DESELECT commands on the other command slots two system clocks later. With the Activate and read command separated by three system clock cycles and taking into account the command slot position of both commands within their system clock cycle, expect the separation on the DDR4 bus to be 11 DRAM clocks, as shown in the DDR bus portion of Figure 4-10.

**Note:** Figure 4-10 shows the relative position of commands on the DDR bus based on the PHY input signals. Although the diagram shows some latency in going through the PHY to be somewhat realistic, this diagram does not represent the absolute command latency through the PHY to the DDR bus, or the system clock to DRAM clock phase alignment. The intention of this diagram is to show the concept of command slots at the PHY interface.




Figure 4-10: PHY Command/Address Input Signal with DDR4 Command/Address Bus

Figure 4-11 shows an example of using all four command slots in a single system clock. This example shows three commands to rank0, and one to rank1, in cycle N. BG and BA address pins are included in the diagram to spread the commands over different banks to not violate DRAM protocol. Table 4-48 lists the command in each command slot.

Table 4-48: Command Slots

Command Slot	0	1	2	3
DRAM Command	Read	Activate	Precharge	Refresh
Bank Group	0	1	2	0
Bank	0	3	1	0
Rank	0	0	0	1

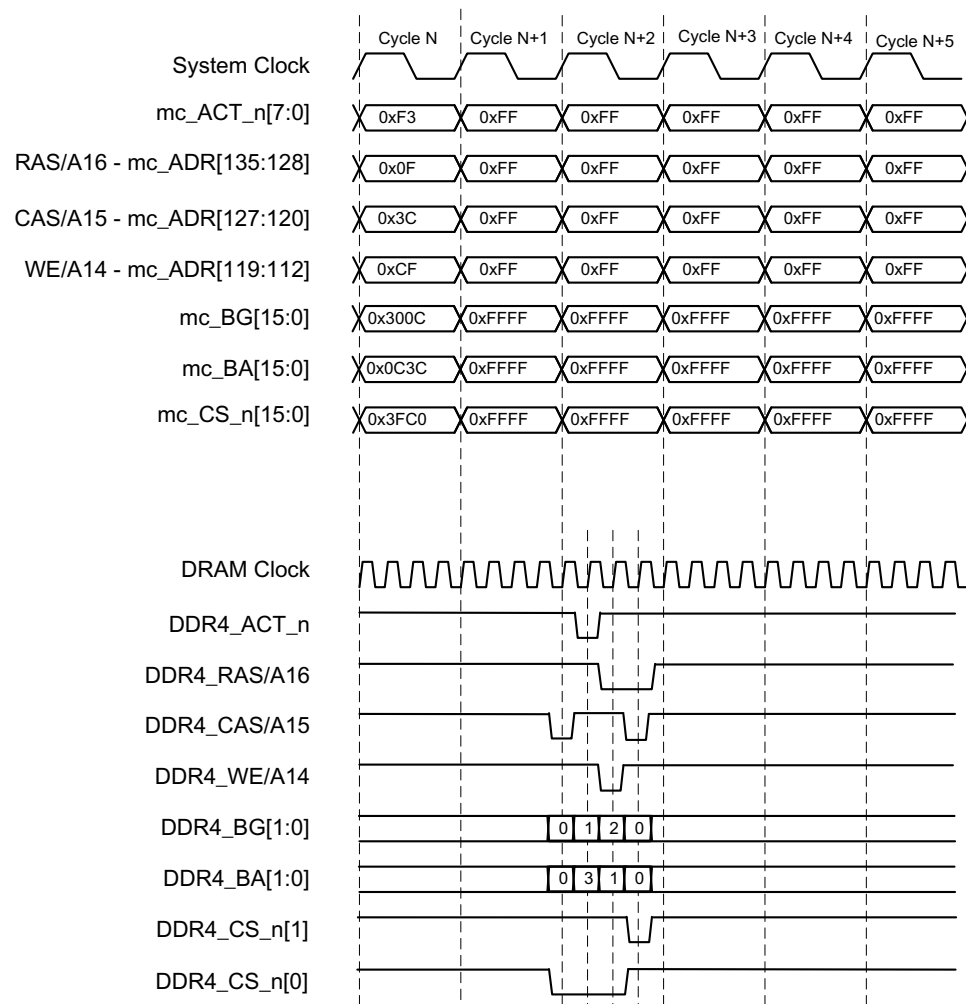


Figure 4-11: PHY Command/Address with All Four Command Slots

To understand how DRAM commands to different command slots are packed together, the following detailed example shows how to convert DRAM commands at the PHY interface to commands on the DRAM command/address bus. To convert PHY interface commands to DRAM commands, write out the PHY signal for one system clock in binary and reverse the bit order of each byte. You can also drop every other bit after the reversal because the bit pairs are required to have the same value. In the subsequent example, the `mc_BA[15:0]` signal has a cycle N value of 0x0C3C:

Hex	0x0C3C
Binary	16'b0000_1100_0011_1100
Reverse bits in each byte	16'b0011_0000_0011_1100

Take the upper eight bits for DRAM `BA[1]` and the lower eight bits for DRAM `BA[0]` and the expected pattern on the DRAM bus is:

BA[1]	00	11	00	00
	0	1	0	0
	Low	High	Low	Low
BA[0]	00	11	11	00
	0	1	1	0
	Low	High	High	Low

This matches the DRAM `BA[1:0]` signal values of 0, 3, 1, and 0 shown in the [Figure 4-11](#).

## Write Data

[Table 4-49](#) shows the write data signals for a PHY only option.

**Table 4-49: Write Data**

Signal	Direction	Description
<code>wrData[DQ_WIDTH × 8 – 1:0]</code>	Input	DRAM write data. Eight bits for each DQ lane on the DRAM bus. This port transfers data for an entire BL8 write on each system clock cycle. Write data must be provided to the PHY one cycle after the <code>wrDataEn</code> output signal asserts, or two cycles after if the ECC parameter is set to "ON." This protocol must be followed. There is no data buffering in the PHY.
<code>wrDataMask[DM_WIDTH × 8 – 1:0]</code>	Input	DRAM write DM/DBI port. One bit for each byte of the <code>wrData</code> port, corresponding to one bit for each byte of each burst of a BL8 transfer. <code>wrDataMask</code> is transferred on the same system clock cycle as <code>wrData</code> . Active-High.
<code>wrDataEn</code>	Output	Write data required. PHY asserts this port for one cycle for each write CAS command. Your design must provide <code>wrData</code> and <code>wrDataMask</code> at the PHY input ports on the cycle after <code>wrDataEn</code> asserts, or two cycles after if the ECC parameter is set to "ON."



Table 4-49: Write Data (Cont'd)

Signal	Direction	Description
wrDataAddr[DATA_BUF_ADDR_WIDTH – 1:0]	Output	Optional control signal. PHY stores and returns a data buffer address for each in-flight write CAS command. The wrDataAddr signal returns the stored addresses. It is only valid when the PHY asserts wrDataEn. You can use this signal to manage the process of sending write data into the PHY for a write CAS command, but this is completely optional.
tCWL[5:0]	Output	Optional control signal. This output indicates the CAS write latency used in the PHY.
dBufAdr[DATA_BUF_ADDR_WIDTH – 1:0]	Input	Reserved. Should be tied Low.

### Read Data

Table 4-50 shows the read data signals for a PHY only option.

Table 4-50: Read Data

Signal	Direction	Description
rdData[DQ_WIDTH × 8 – 1:0]	Output	DRAM read data. Eight bits for each DQ lane on the DRAM bus. This port transfers data for an entire BL8 read on each system clock cycle. rdData is only valid when the rdDataEn, per_rd_done, or rmw_rd_done is asserted. Your design must consume the read data when rdDataEn one of these “data valid” signals asserts. There is no data buffering in the PHY.
rdDataEn	Output	Read data valid. This signal asserts High to indicate that the rdData and rdDataAddr signals are valid. rdDataEn asserts High for one system clock cycle for each BL8 read, unless the read was tagged as a special type of read. See the optional per_rd_done and rmw_rd_done signals for details on special reads. rdData must be consumed when rdDataEn asserts or data is lost. Active-High.
rdDataAddr[DATA_BUF_ADDR_WIDTH – 1:0]	Output	Optional control signal. PHY stores and returns a data buffer address for each in-flight read CAS command. The rdDataAddr signal returns the stored addresses. It is only valid when the PHY asserts rdDataEn, per_rd_done, or rmw_rd_done. Your design can use this signal to manage the process of capturing and storing read data provided by the PHY, but this is completely optional.
per_rd_done	Output	Optional read data valid signal. This signal indicates that a special type of read has completed and its associated rdData and rdDataAddr signals are valid. When PHY input winInjTxn is asserted High at the same time as mcRdCAS, the read is tagged as a special type of read, and per_rd_done asserts instead of rdDataEn when data is returned.

Table 4-50: Read Data (Cont'd)

Signal	Direction	Description
rmw_rd_done	Output	Optional read data valid signal. This signal indicates that a special type of read has completed and its associated rdData and rdDataAddr signals are valid. When PHY input winRmw is asserted High at the same time as mcRdCAS, the read is tagged as a special type of read, and rmw_rd_done asserts instead of rdDataEn when data is returned.
rdDataEnd	Output	Unused. Tied High.

### PHY Control

Table 4-51 shows the PHY control signals for a PHY only option.

Table 4-51: PHY Control

Signal	Direction	Description
calDone	Output	Indication that the DRAM is powered up, initialized, and calibration is complete. This indicates that the PHY interface is available to send commands to the DRAM. Active-High.
mcRdCAS	Input	Read CAS command issued. This signal must be asserted for one system clock if and only if a read CAS command is asserted on one of the command slots at the PHY command/address input ports. Hold at 0x0 until calDone asserts. Active-High.
mcWrCAS	Input	Write CAS command issued. This signal must be asserted for one system clock if and only if a write CAS command is asserted on one of the command slots at the PHY command/address input ports. Hold at 0x0 until calDone asserts. Active-High.
winRank[1:0]	Input	Target rank for CAS commands. This value indicates which rank a CAS command is issued to. It must be valid when either mcRdCAS or mcWrCAS is asserted. The PHY passes the value from this input to the XIPHY to select the calibration results for the target rank of a CAS command in multi-rank systems. In a single rank system, this input port can be tied to 0x0.
mcCasSlot[1:0]	Input	CAS command slot select. The PHY only supports CAS commands on even command slots. mcCasSlot indicates which of these two possible command slots a read CAS or write CAS was issued on. mcCasSlot is used by the PHY to generate XIPHY control signals, like DQ output enables, that need DRAM clock cycle resolution relative to the command slot used for a CAS command. Valid values after calDone asserts are 0x0 and 0x2. Hold at 0x0 until calDone asserts. This signal must be valid if mcRdCAS or mcWrCAS is asserted. For more information, see the <a href="#">CAS Command Timing Limitations</a> , page 141.

Table 4-51: PHY Control (Cont'd)

Signal	Direction	Description
mcCasSlot2	Input	CAS slot 2 select. mcCasSlot2 serves a similar purpose as the mcCasSlot[1:0] signal, but mcCasSlot2 is used in timing critical logic in the PHY. Ideally mcCasSlot2 should be driven from separate flops from mcCasSlot[1:0] to allow synthesis/implementation to better optimize timing. mcCasSlot2 and mcCasSlot[1:0] must always be consistent if mcRdCAS or mcWrCAS is asserted. To be consistent, the following must be TRUE: mcCasSlot2==mcCasSlot[1]. Hold at 0x0 until calDone asserts. Active-High.
winInjTxn	Input	Optional read command type indication. When winInjTxn is asserted High on the same cycle as mcRdCAS, the read does not generate an assertion on rdDataEn when it completes. Instead, the per_rd_done signal asserts, indicating that a special type of read has completed and that its data is valid on the rdData output. In MIG controller designs, the winInjTxn/per_rd_done signals are used to track non-system read traffic by asserting winInjTxn only on read commands issued for the purpose of VT tracking.
winRmw	Input	Optional read command type indication. When winRmw is asserted High on the same cycle as mcRdCAS, the read does not generate an assertion on rdDataEn when it completes. Instead, the rmw_rd_done signal asserts, indicating that a special type of read has completed and that its data is valid on the rdData output. In MIG controller designs, the winRmw/rmw_rd_done signals are used to track reads issued as part of a read-modify-write flow. The MIG controller asserts winRmw only on read commands that are issued for the read phase of a RMW sequence.
winBuf[DATA_BUF_ADDR_WIDTH – 1:0]	Input	Optional control signal. When either mcRdCAS or mcWrCAS is asserted, PHY stores the value on the winBuf signal. The value is returned on rdDataAddr or wrDataAddr, depending on whether mcRdCAS or mcWrCAS was used to capture winBuf. In MIG controller designs, these signals are used to track the data buffer address used to source write data or sink read return data.
gt_data_ready	Input	Update VT Tracking. This signal triggers the PHY to read RIU registers in the XIPHY that measure how well the DQS Gate signal is aligned to the center of the read DQS preamble, and then adjust the alignment if needed. This signal must be asserted periodically to keep the DQS Gate aligned as voltage and temperature drift. For more information, see <a href="#">VT Tracking, page 143</a> . Hold at 0x0 until calDone asserts. Active-High.

Figure 4-12 shows a write command example. On cycle N, write command “A” is asserted on the PHY command/address inputs in the slot0 position. The mcWrCAS input is also asserted on cycle N, and a valid rank value is asserted on the winRank signal. In Figure 4-12, there is only one CS\_n pin, so the only valid winRank value is 0x0. The mcCasSlot[1:0] and mcCasSlot2 signals are valid on cycle N, and specify slot0.

Write command “B” is then asserted on cycle N + 1 in the slot2 position, with `mcWrCAS`, `winRank`, `mcCasSlot[1:0]`, and `mcCasSlot2` asserted to valid values as well. On cycle M, PHY asserts `wrDataEn` to indicate that `wrData` and `wrDataMask` values corresponding to command A need to be driven on cycle M + 1. On cycle M + 2, PHY asserts `wrDataEn` to indicate that `wrData` and `wrDataMask` values corresponding to command B need to be driven on cycle M + 3.

Figure 4-12 shows the data and mask widths assuming an 8-bit DDR4 DQ bus width. The delay between cycle N and cycle M is controlled by the PHY, based on the CWL and AL settings of the DRAM. `wrDataEn` also asserts on cycle M + 1 to indicate that `wrData` and `wrDataMask` values for command B are required on cycle M + 2. Although this example shows that `wrDataEn` is asserted on two consecutive system clock cycles, you should not assume this will always be the case, even if `mcWrCAS` is asserted on consecutive clock cycles as is shown here. There is no data buffering in the PHY and data is pulled into the PHY just in time. Depending on the CWL/AL settings and the command slot used, consecutive `mcWrCAS` assertions might not result in consecutive `wrDataEn` assertions.

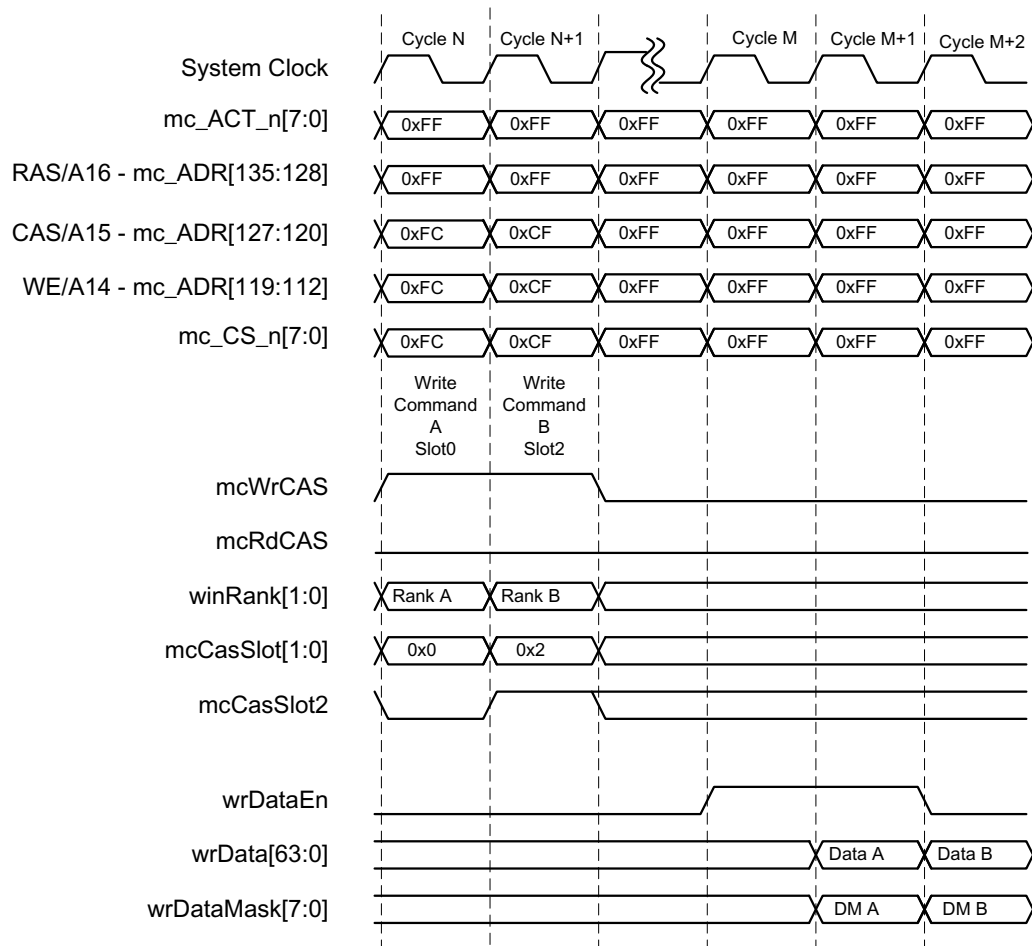


Figure 4-12: Write Command Example

Figure 4-13 shows a read command example. Read commands are issued on cycles N and N + 1 in slot positions 0 and 2, respectively. The `mcRdCAS`, `winRank`, `mcCasSlot`, and `mcCasSlot2` are asserted on these cycles as well. On cycles M + 1 and M + 2, PHY asserts `rdDataEn` and `rdData`.

**Note:** The separation between N and M + 1 is much larger than in the write example (Figure 4-12). In the read case, the separation is determined by the full round trip latency of command output, DRAM CL/AL, and data input through PHY.

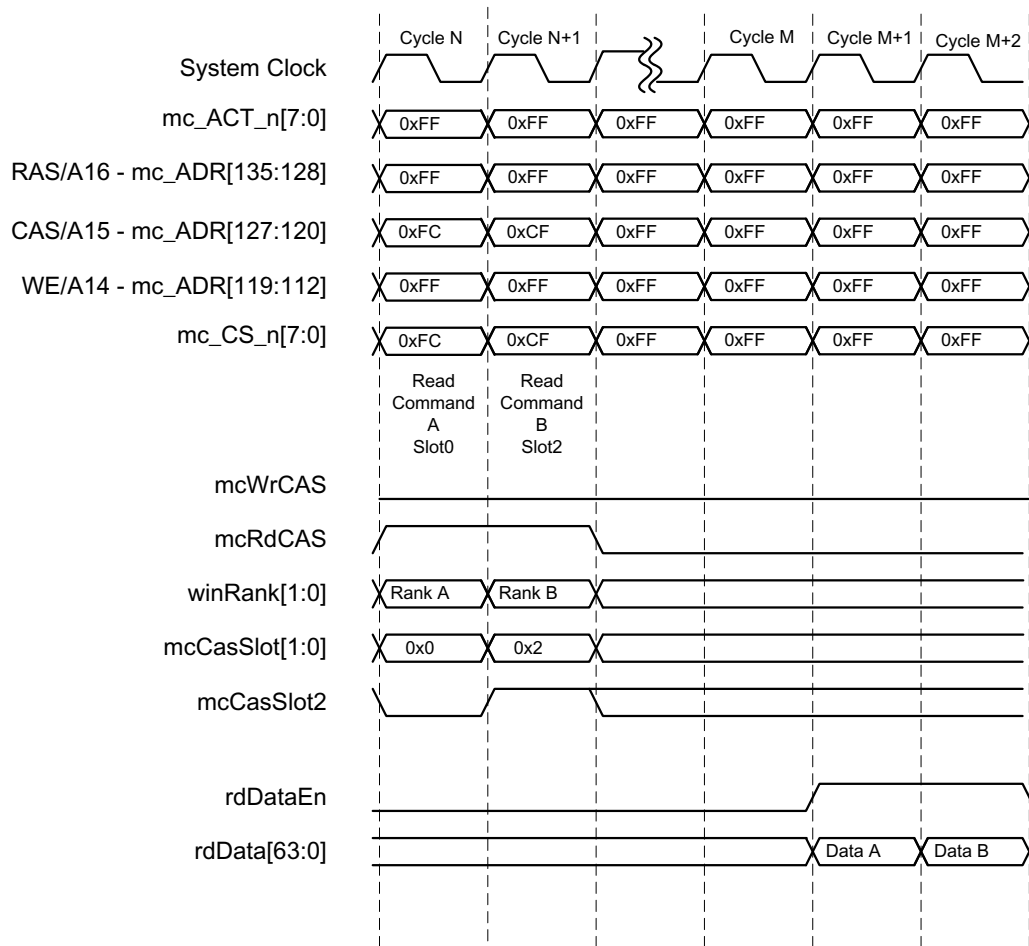


Figure 4-13: Read Command Example

## Debug

The debug signals are explained in Answer Record: [60305](#).

## PHY Only Parameters

All PHY parameters are configured by the MIG software. Table 4-52 describes the PHY parameters. These parameter values must not be modified in the MIG generated designs. The parameters are set during core generation. The core must be regenerated to change any parameter settings.

Table 4-52: PHY Only Parameters

Parameter Name	Default Value	Allowable Values	Description
ADDR_WIDTH	18	DDR4 18.. 17 DDR3 16.. 13	Number of DRAM Address pins
BANK_WIDTH	2	DDR4 2 DDR3 3	Number of DRAM Bank Address pins
BANK_GROUP_WIDTH	2	DDR4 2.. 1 DDR3 N/A	Number of DRAM Bank Group pins
CK_WIDTH	1	2.. 1	Number of DRAM Clock pins
CKE_WIDTH	1	2.. 1	Number of DRAM CKE pins
CS_WIDTH	1	2.. 1	Number of DRAM CS pins
ODT_WIDTH	1	4.. 1	Number of DRAM ODT pins
DRAM_TYPE	"DDR4"	"DDR4," "DDR3"	DRAM Technology
DQ_WIDTH	16	Minimum = 8 Must be multiple of 8	Number of DRAM DQ pins in the channel
DQS_WIDTH	2	Minimum = 1 x8 DRAM – 1 per DQ byte x4 DRAM – 1 per DQ nibble	Number of DRAM DQS pins in the channel
DM_WIDTH	2	Minimum = 0 x8 DRAM – 1 per DQ byte x4 DRAM – 0	Number of DRAM DM pins in the channel
DATA_BUF_ADDR_WIDTH	5	5	Number of data buffer address bits stored for a read or write transaction
ODTWR	0x8421	0xFFFF .. 0x0000	Reserved for future use
ODTWRDEL	8	Set to CWL	Reserved for future use
ODTWRDUR	6	7.. 6	Reserved for future use
ODTRD	0x0000	0xFFFF.. 0x0000	Reserved for future use
ODTRDDEL	11	Set to CL	Reserved for future use
ODTRDDUR	6	7.. 6	Reserved for future use
ODTWR0DEL ODTWR0DUR ODTRD0DEL ODTRD0DUR ODTNOP	N/A	N/A	Reserved for future use
MR0	0x630	Legal SDRAM configuration	DRAM MR0 setting
MR1	0x101	Legal SDRAM configuration	DRAM MR1 setting
MR2	0x10	Legal SDRAM configuration	DRAM MR2 setting

Table 4-52: PHY Only Parameters (Cont'd)

Parameter Name	Default Value	Allowable Values	Description
MR3	0x0	Legal SDRAM configuration	DRAM MR3 setting
MR4	0x0	Legal SDRAM configuration	DRAM MR4 setting. DDR4 only.
MR5	0x400	Legal SDRAM configuration	DRAM MR5 setting. DDR4 only.
MR6	0x800	Legal SDRAM configuration	DRAM MR6 setting. DDR4 only.
SLOT0_CONFIG	0x1	0x1 0x3 0x5 0xF	Reserved for future use
SLOT1_CONFIG	0x0	0x0 0x2 0xC 0xA	Reserved for future use
SLOT0_FUNC_CS	0x1	0x1 0x3 0x5 0xF	Memory bus CS_n pins used to send all DRAM commands including MRS to memory. Each bit of the parameter represents 1-bit of the CS_n bus, for example, the LSB indicates CS_n[0], and the MSB indicates CS_n[3]. For DIMMs this parameter specifies the CS_n pins connected to DIMM slot 0. <b>Note:</b> slot 0 used here should not be confused with the "command slot0" term used in the description of the PHY command/address interface.
SLOT1_FUNC_CS	0x0	0x0 0x2 0xC 0xA	See the SLOT0_FUNC_CS description. The only difference is that SLOT1_FUNC_CS specifies CS_n pins connected to DIMM slot 1.
REG_CTRL	OFF	"ON" "OFF"	Enable RDIMM RCD initialization and calibration
CA_MIRROR	OFF	"ON" "OFF"	Enable Address mirroring. This parameter is set to "ON" for the DIMMs that support address mirroring.
DDR4_REG_RC03	0x30	Legal RDIMM RCD configuration	RDIMM RCD control word 03
DDR4_REG_RC04	0x40	Legal RDIMM RCD configuration	RDIMM RCD control word 04
DDR4_REG_RC05	0x50	Legal RDIMM RCD configuration	RDIMM RCD control word 05
tCK	938	Minimum 833	DRAM clock period in ps

Table 4-52: PHY Only Parameters (Cont'd)

Parameter Name	Default Value	Allowable Values	Description
tXPR	72	Minimum 1. DRAM tXPR specification in system clocks	See JEDEC DDR SDRAM specification [Ref 1].
tMOD	6	Minimum 1. DRAM tMOD specification in system clocks	See JEDEC DDR SDRAM specification [Ref 1].
tMRD	2	Minimum 1. DRAM tMRD specification in system clocks	See JEDEC DDR SDRAM specification [Ref 1].
tZQINIT	256	Minimum 1. DRAM tZQINIT specification in system clocks	See JEDEC DDR SDRAM specification [Ref 1].
TCQ	100	100	Flop clock to Q in ps. For simulation purposes only.
EARLY_WR_DATA	OFF	OFF	Reserved for future use
EXTRA_CMD_DELAY	0	2.. 0	Added command latency in system clocks. Added command latency is required for some configurations. See details in CL/CWL section.
ECC	"OFF"	"OFF"	Enables early wrDataEn timing for MIG generated controllers when set to "ON." PHY only designs must set this to "OFF."
DM_DBI	"DM_NODBI"	"NONE" "DM_NODBI" "DM_DBIRD" "NODM_DBIWR" "NODM_DBIRD"	DDR4 DM/DBI configuration. For details, see Table 4-54.
USE_CS_PORT	1	0 = no CS_n pins 1 = CS_n pins used	Controls whether or not CS_n pins are connect to DRAM. If there are no CS_n pins the PHY initialization and training logic issues NOPs between DRAM commands. If there are no CS_n pins, The DRAM chip select pin (CS#) must be tied Low externally at the DRAM.
DRAM_WIDTH	8	16, 8, 4	DRAM component DQ width
RANKS	1	2.. 1	Number of ranks in the memory subsystem
nCK_PER_CLK	4	4	Number of DRAM clocks per system clock
C_FAMILY	"kintexu"	"kintexu" "virtexu"	Device information used by MicroBlaze controller in the PHY.



**Table 4-52: PHY Only Parameters (Cont'd)**

Parameter Name	Default Value	Allowable Values	Description
BYTES	4	Minimum 3	Number of XIPHY "bytes" used for data, command, and address
DBYTES	2	Minimum 1	Number of bytes in the DRAM DQ bus
IOBTYPE	{39'b001_001_001_001_001_001_001_001, 39'b001_001_001_001_001_001_001_001, 39'b000_011_011_011_011_011_011_011, 39'b001_011_011_011_011_011_011_011 }	3'b000 = Unused pin 3'b 001 = Single-ended output 3'b 010 = Single-ended input 3'b011 = Single-ended I/O 3'b100 = Unused pin 3'b 101 = Differential Output 3'b 110 = Differential Input 3'b 111 = Differential INOUT	IOB setting
PLL_WIDTH	1	MIG generated values	Number of PLLs
CLKOUTPHY_MODE	"VCO_2X"	VCO_2X	Determines the clock output frequency based on the VCO frequency for the BITSlice_CONTROL block
PLLCLK_SRC	0	0 = pll_clk0 1 = pll_clk1	XIPHY PLL clock source
DIV_MODE	0	0 = DIV4 1 = DIV2	XIPHY controller mode setting
DATA_WIDTH	8	8	XIPHY parallel input data width
CTRL_CLK	0x3	0 = Internal, local div_clk used 1 = External RIU clock used	Internal or external XIPHY clock for the RIU
INIT	{{(15 × BYTES){1'b1}}	1'b0 1'b1	3-state bitslice OSERDES initial value
RX_DATA_TYPE	{15'b000000_00_0000_00, 15'b000000_00_0000_00, 15'b011110_10_1110_01, 15'b011110_10_1110_01 }	2'b00 = None 2'b01 = DATA(DQ_EN) 2'b10 = CLOCK(DQS_EN) 2'b11 = DATA_AND_CLOCK	XIPHY bitslice setting

Table 4-52: PHY Only Parameters (Cont'd)

Parameter Name	Default Value	Allowable Values	Description
TX_OUTPUT_PHASE_90	{13'b111111111111, 13'b111111111111, 13'b0000011000010, 13'b1000011000010}	1'b0 = No offset 1'b1 = 90° offset applied	XIPHY setting to apply 90° offset on a given bitslice
RXTX_BITSLICE_EN	{13'b111101111111, 13'b111111111111, 13'b011110111111, 13'b111110111111}	1'b0 = No bitslice 1'b1 = Bitslice enabled	XIPHY setting to enable a bitslice
NATIVE_ODLAY_BYPASS	{{(13 × BYTES){1'b0}}	1'b0 = FALSE 1'b1 = TRUE (Bypass)	Bypass the ODELAY on output bitslices
EN_OTHER_PCLK	{BYTES{2'b01}}	1'b 0 = FALSE (not used) 1'b 1 = TRUE (used)	XIPHY setting to route capture clock from other bitslice
EN_OTHER_NCLK	{BYTES{2'b01}}	1'b 0 = FALSE (not used) 1'b 1 = TRUE (used)	XIPHY setting to route capture clock from other bitslice
RX_CLK_PHASE_P	{{(BYTES – DBYTES){2'b00}}, {DBYTES{2'b11}}}	2'b00 for Address/Control, 2'b11 for Data	XIPHY setting to shift the read clock DQS_P by 90° relative to the DQ
RX_CLK_PHASE_N	{{(BYTES – DBYTES){2'b00}}, {DBYTES{2'b11}}}	2'b00 for Address/Control, 2'b11 for Data	XIPHY setting to shift the read clock DQS_N by 90° relative to the DQ
TX_GATING	{{(BYTES – DBYTES){2'b00}}, {DBYTES{2'b11}}}	2'b00 for Address/Control, 2'b11 for Data	Write DQS gate setting for the XIPHY
RX_GATING	{{(BYTES – DBYTES){2'b00}}, {DBYTES{2'b11}}}	2'b00 for Address/Control, 2'b11 for Data	Read DQS gate setting for the XIPHY
EN_DYN_ODLY_MODE	{{(BYTES – DBYTES){2'b00}}, {DBYTES{2'b11}}}	2'b00 for Address/Control, 2'b11 for Data	Dynamic loading of the ODELAY by XIPHY
BANK_TYPE	"HP_IO"	"HP_IO" "HR_IO"	Indicates whether selected bank is HP or HR
SIM_MODE	"FULL"	"FULL", "BFM"	Flag to set if the XIPHY is used ("FULL") or the behavioral model for simulation speed up.
SELF_CALIBRATE	{{(2 × BYTES){1'b0}}	{{(2 × BYTES){1'b0}} for simulation, {{(2 × BYTES){1'b1}} for hardware	BISC self calibration

Table 4-52: PHY Only Parameters (Cont'd)

Parameter Name	Default Value	Allowable Values	Description
BYPASS_CAL	"FALSE"	"TRUE" for simulation, "FALSE" for hardware	Flag to turn calibration ON/OFF
CAL_WRLVL	"FULL"	"FULL"	Flag for calibration, write-leveling setting
CAL_DQS_GATE	"FULL"	"FULL"	Flag for calibration, DQS gate setting
CAL_RDLVL	"FULL"	"FULL"	Flag for calibration, read training setting
CAL_WR_DQS_DQ	"FULL"	"FULL"	Flag for calibration, write DQS-to-DQ setting
CAL_COMPLEX	"FULL"	"SKIP", "FULL"	Flag for calibration, complex pattern setting
CAL_RD_VREF	"SKIP"	"SKIP", "FULL"	Flag for calibration, read $V_{REF}$ setting
CAL_WR_VREF	"SKIP"	"SKIP", "FULL"	Flag for calibration, write $V_{REF}$ setting
CAL_JITTER	"FULL"	"FULL", "NONE"	Reserved for verification. Speed up calibration simulation. Must be set to "FULL" for all hardware test cases.
t200us	53305 decimal	0x3FFFF.. 1	Wait period after BISC complete to DRAM reset_n deassertion in system clocks
t500us	133263 decimal	0x3FFFF.. 1	Wait period after DRAM reset_n deassertion to CKE assertion in system clocks
SIM_MODE	BFM	"FULL", "BFM"	<ul style="list-style-type: none"> <li>FULL: Run example design simulations with XIPHY UNISIMs</li> <li>BFM: Run Fast Mode Simulations with XIPHY Bus Functional Model</li> </ul>

### EXTRA\_CMD\_DELAY Parameter

Depending on the number of ranks, ECC mode, and DRAM latency configuration, PHY must be programmed to add latency on the DRAM command address bus. This provides enough pipeline stages in the PHY programmable logic to close timing and to process `mcWrCAS`. Added command latency is generally needed at very low CWL in single rank configurations, or in multi-rank configurations. Enabling ECC might also require adding command latency, but this depends on whether your controller design (outside the PHY) depends on receiving the `wrDataEn` signal a system clock cycle early to allow for generating ECC check bits.

The `EXTRA_CMD_DELAY` parameter is used to add one or two system clock cycles of delay on the DRAM command/address path. The parameter does not delay the `mcWrCAS` or `mcRdCAS` signals. This gives the PHY more time from the assertion of `mcWrCAS` or `mcRdCAS` to generate XIPHY control signals. To the PHY, an `EXTRA_CMD_DELAY` setting of one or two is the same as having a higher CWL or AL setting.

Table 4-53 shows the required EXTRA\_CMD\_DELAY setting for various configurations of CWL, CL, and AL.

Table 4-53: EXTRA\_CMD\_DELAY Configuration Settings

DRAM Configuration			Required EXTRA_CMD_DELAY	
DRAM CAS Write Latency CWL	DRAM CAS Latency CL	DRAM Additive Latency MR1[4:3]	Single Rank without ECC	Single Rank with ECC or Multi-Rank
5	5	0	1	2
5	5	1	0	1
5	5	2	1	2
5	6	0	1	2
5	6	1	0	1
5	6	2	0	1
6	6	0	1	2
6	6	1	0	1
6	6	2	0	1
6	7	0	1	2
6	7	1	0	1
6	7	2	0	1
6	8	0	1	2
6	8	1	0	0
6	8	2	0	1
7	7	0	1	2
7	7	1	0	0
7	7	2	0	1
7	8	0	1	2
7	8	1	0	0
7	8	2	0	0
7	9	0	1	2
7	9	1	0	0
7	9	2	0	0
7	10	0	1	2
7	10	1	0	0
7	10	2	0	0
8	8	0	1	2
8	8	1	0	0
8	8	2	0	0
8	9	0	1	2

Table 4-53: EXTRA\_CMD\_DELAY Configuration Settings (Cont'd)

DRAM Configuration			Required EXTRA_CMD_DELAY	
DRAM CAS Write Latency CWL	DRAM CAS Latency CL	DRAM Additive Latency MR1[4:3]	Single Rank without ECC	Single Rank with ECC or Multi-Rank
8	9	1	0	0
8	9	2	0	0
8	10	0	1	2
8	10	1	0	0
8	10	2	0	0
8	11	0	1	2
8	11	1	0	0
8	11	2	0	0
9 to 12	X	0	0	1
9 to 12	X	1 or 2	0	0
≥13	X	0	0	0
≥13	X	1 or 2	0	0

### DM\_DBI Parameter

The PHY supports the DDR4 DBI function on the read path and write path. Table 4-54 show how read and write DBI can be enabled separately or in combination. When write DBI is enabled, Data Mask is disabled. The DM\_DBI parameter only configures the PHY and the MRS parameters must also be set to configure the DRAM for DM/DBI.

Table 4-54: DM\_DBI PHY Settings

DM_DBI Parameter Value	PHY Read DBI	PHY Write DBI	PHY Write Data Mask
None	Disabled	Disabled	Enabled
DM_NODBI	Disabled	Disabled	Enabled
DM_DBI RD	Enabled	Disabled	Enabled
NODM_DBI WR	Disabled	Enabled	Disabled
NODM_DBI RD	Enabled	Enabled	Disabled

### CAS Command Timing Limitations

The PHY only supports CAS commands on even command slots, that is, 0 and 2. This limitation is due to the complexity of the PHY logic driven by the PHY control inputs, like the `mcWrCAS` and `mcRdCAS` signals, not the actual DRAM command signals like `mc_ACT_n[7:0]`, which just pass through the PHY after `calDone` asserts. The PHY logic is complex because it generates XIPHY control signals based on the DRAM CWL and CL values with DRAM clock resolution, not just system clock resolution.

Supporting two different command slots for CAS commands adds a significant amount of logic on the XIPHY control paths. There are very few pipeline stages available to break up the logic due to protocol requirements of the XIPHY. CAS command support on all four slots would further increase the complexity and degrade timing.

### ***Minimum Write CAS Command Spacing***

The minimum Write CAS to Write CAS command spacing to different ranks is eight DRAM clocks. This is a PHY limitation. If you violate this timing, the PHY might not have enough time to switch its internal delay settings and drive Write DQ/DQS on the DDR bus with correct timing. The internal delay settings are determined during calibration, and it varies with system layout.

Following the memory system layout guidelines ensures that a spacing of eight DRAM clocks is sufficient for correct operation. Write to Write timing to the same rank is limited only by the DRAM specification and the command slot limitations for CAS commands discussed earlier.

### ***System Considerations for CAS Command Spacing***

System layout and timing uncertainties should be considered in how your custom controller sets minimum CAS command spacing. The controller must space the CAS commands so that there are no DRAM timing violations and no DQ/DQS bus drive fights. When a MIG generated memory controller is instantiated, the layout guidelines are considered and command spacing is adjusted accordingly for a worst case layout.

Consider Read to Write command spacing, the JEDEC<sup>®</sup> DRAM specification [Ref 1] shows the component requirement as:  $RL + BL/2 + 2 - WL$ . This formula only spaces the Read DQS post-amble and Write DQS preamble by one DRAM clock on an ideal bus with no timing skews. Any DQS flight time, write leveling uncertainty, jitter, etc. reduces this margin. When these timing errors add up to more than one DRAM clock, there is a drive fight at the FPGA DQS pins which likely corrupts the Read transaction. A MIG generated controller uses the following formula to delay Write CAS after a Read CAS to allow for a worst case timing budget for a system following the layout guidelines:  $RL + BL/2 + 4 - WL$ .

Read CAS to Read CAS commands to different ranks must also be spaced by your custom controller to avoid drive fights, particularly when reading first from a "far" rank and then from a "near" rank. A MIG generated controller spaces the Read CAS commands to different ranks by at least six DRAM clock cycles.

Write CAS to Read CAS to the same rank is defined by the JEDEC DRAM specification [Ref 1]. Your controller must follow this DRAM requirement, and it ensures that there is no possibility of drive fights for Write to Read to the same rank. Write CAS to Read CAS spacing to different ranks, however, must also be limited by your controller. This spacing is not defined by the JEDEC DRAM specification [Ref 1] directly.

Write to Read to different ranks can be spaced much closer together than Write to Read to the same rank, but factors to consider include write leveling uncertainty, jitter, and tDQSCK. A MIG generated controller spaces Write CAS to Read CAS to different ranks by at least six DRAM clocks.

### Additive Latency

The PHY supports DRAM additive latency. The only effect on the PHY interface due to enabling Additive Latency in the MRS parameters is in the timing of the `wrDataEn` signal after `mcWrCAS` assertion. The PHY takes the AL setting into account when scheduling `wrDataEn`. You can also find the `rdDataEn` asserts much later after `mcRdCAS` because the DRAM returns data much later. The AL setting also has an impact on whether or not the `EXTRA_CMD_DELAY` parameter needs to be set to a non-zero value.

### VT Tracking

The PHY requires read commands to be issued at a minimum rate to keep the read DQS gate signal aligned to the read DQS preamble after `calDone` is asserted. In addition, the `gt_data_ready` signal needs to be pulsed at regular intervals to instruct the PHY to update its read DQS training values in the RIU. Finally, the PHY requires periodic gaps in read traffic to allow the XIPHY to update its gate alignment circuits with the values the PHY programs into the RIU. Specifically, the PHY requires the following after `calDone` asserts:

1. At least one read command every 1  $\mu$ s. For a multi-rank system any rank is acceptable.
2. The `gt_data_ready` signal is asserted for one system clock cycle after the `rdDataEn` signal asserts at least once within each 1  $\mu$ s interval.
3. There is a three contiguous system clock cycle period with no read CAS commands asserted at the PHY interface every 1  $\mu$ s.

The PHY cannot interrupt traffic to meet these requirements. It is therefore your custom Memory Controller's responsibility to issue DRAM commands and assert the `gt_data_ready` input signal in a way that meets the above requirements.

Figure 4-14 shows two examples where the custom controller must interrupt normal traffic to meet the VT tracking requirements. The first example is a High read bandwidth workload with `mcRdCAS` asserted continuously for almost 1  $\mu$ s. The controller must stop issuing read commands for three contiguous system clocks once each 1  $\mu$ s period, and assert `gt_data_ready` once per period.

The second example is a High write bandwidth workload with `mcWrCAS` asserted continuously for almost 1  $\mu$ s. The controller must stop issuing writes, issue at least one read command, and then assert `gt_data_ready` once per 1  $\mu$ s period.




---

**IMPORTANT:** *The controller must not violate DRAM protocol or timing requirements during this process.*

---

**Note:** The VT tracking diagrams are not drawn to scale.




Figure 4-14: VT Tracking Diagrams

A workload that has a mix of read and write traffic in every 1  $\mu$ s interval might naturally meet the first and third VT tracking requirements listed above. In this case, the only extra step required is to assert the `gt_data_ready` signal every 1  $\mu$ s and regular traffic would not be interrupted at all. The custom controller, however, is responsible for ensuring all three requirements are met for all workloads. MIG generated controllers monitor the `mcRdCAS` and `mcWrCAS` signals and decide each 1  $\mu$ s period what actions, if any, need to be taken to meet the VT tracking requirements. Your custom controller can implement any scheme that meets the requirements described here.

## Refresh and ZQ

After `calDone` is asserted by the PHY, periodic DRAM refresh and ZQ calibration are the responsibility of your custom Memory Controller. Your controller must issue refresh and ZQ commands, meet DRAM refresh and ZQ interval requirements, while meeting all other DRAM protocol and timing requirements. For example, if a refresh is due and you have open pages in the DRAM, you must precharge the pages, wait `tRP`, and then issue a refresh command, etc. The PHY does not perform the precharge or any other part of this process for you.



## Performance

The efficiency of a memory system is affected by many factors including limitations due to the memory, such as cycle time (t<sub>RC</sub>) within a single bank, or Activate to Activate spacing to the same DDR4 bank group (t<sub>RRD\_L</sub>). When given multiple transactions to work on, the Memory Controller schedules commands to the DRAM in a way that attempts to minimize the impact of these DRAM timing requirements. But there are also limitations due to the Memory Controller architecture itself. This section explains the key controller limitations and options for obtaining the best performance out of the controller.

## Address Map

The `app_addr` to the DRAM address map is described in the [User Interface](#). Three mapping options are included:

- ROW\_COLUMN\_BANK
- ROW\_BANK\_COLUMN
- BANK\_ROW\_COLUMN

For a purely random address stream at the user interface, all three options would result in a similar efficiency. For a sequential `app_addr` address stream, or any workload that tends to have a small stride through the `app_addr` memory space, the ROW\_COLUMN\_BANK mapping generally provides a better overall efficiency. This is due to the Memory Controller architecture and the interleaving of transactions across the Group FSMs. The Group FSMs are described in the [Memory Controller, page 18](#). This controller architecture impact on efficiency should be considered even for situations where DRAM timing is not limiting efficiency. [Table 4-55](#) shows two mapping options for the 4 Gb (x8) DRAM components.

**Table 4-55: DDR3/DDR4 4 Gb (x8) DRAM Address Mapping Options**

DRAM Address	DDR3 4 Gb (x8)		DDR4 4 Gb (x8)	
	ROW_BANK_COLUMN	ROW_COLUMN_BANK	ROW_BANK_COLUMN	ROW_COLUMN_BANK
Row 15	28	28	–	–
Row 14	27	27	28	28
Row 13	26	26	27	27
Row 12	25	25	26	26
Row 11	24	24	25	25
Row 10	23	23	24	24
Row 9	22	22	23	23
Row 8	21	21	22	22
Row 7	20	20	21	21

Table 4-55: DDR3/DDR4 4 Gb (x8) DRAM Address Mapping Options (Cont'd)

DRAM Address	DDR3 4 Gb (x8)		DDR4 4 Gb (x8)	
	ROW_BANK_COLUMN	ROW_COLUMN_BANK	ROW_BANK_COLUMN	ROW_COLUMN_BANK
Row 6	19	19	20	20
Row 5	18	18	19	19
Row 4	17	17	18	18
Row 3	16	16	17	17
Row 2	15	15	16	16
Row 1	14	14	15	15
Row 0	13	13	14	14
Column 9	9	12	9	13
Column 8	8	11	8	12
Column 7	7	10	7	11
Column 6	6	9	6	10
Column 5	5	8	5	9
Column 4	4	7	4	8
Column 3	3	6	3	7
Column 2	2	2	2	2
Column 1	1	1	1	1
Column 0	0	0	0	0
Bank 2	12	4	–	–
Bank 1	11	3	11	6
Bank 0	10	5	10	5
Bank Group 1	–	–	13	4
Bank Group 0	–	–	12	3

**Note:** Highlighted bits are used to map addresses to Group FSMs in the controller.

From the DDR3 map, you might expect reasonable efficiency with the ROW\_BANK\_COLUMN option with a simple address increment pattern. The increment pattern would generate page hits to a single bank, which DDR3 could handle as a stream of back-to-back CAS commands resulting in high efficiency. But looking at the italic bank bits in Table 4-55 show that the address increment pattern also maps the long stream of page hits to the same controller Group FSM.

For example, [Table 4-56](#) shows how the first 12 `app_addr` addresses decode to the DRAM addresses and map to the Group FSMs for both mapping options. The `ROW_BANK_COLUMN` option only maps to the Group FSM 0 over this address range.

**Table 4-56: DDR3/DDR4 4 Gb (x8) `app_addr` Mapping Options**

app_addr	DDR3 4 Gb (x8) ROW_BANK_COLUMN				DDR3 4 Gb (x8) ROW_COLUMN_BANK			
	Row	Column	Bank	Group_FSM	Row	Column	Bank	Group_FSM
0x58	0x0	0x58	0x0	0	0x0	0x8	0x6	3
0x50	0x0	0x50	0x0	0	0x0	0x8	0x4	2
0x48	0x0	0x48	0x0	0	0x0	0x8	0x2	1
0x40	0x0	0x40	0x0	0	0x0	0x8	0x0	0
0x38	0x0	0x38	0x0	0	0x0	0x0	0x7	4
0x30	0x0	0x30	0x0	0	0x0	0x0	0x5	3
0x28	0x0	0x28	0x0	0	0x0	0x0	0x3	2
0x20	0x0	0x20	0x0	0	0x0	0x0	0x1	0
0x18	0x0	0x18	0x0	0	0x0	0x0	0x6	3
0x10	0x0	0x10	0x0	0	0x0	0x0	0x4	2
0x8	0x0	0x8	0x0	0	0x0	0x0	0x2	1
0x0	0x0	0x0	0x0	0	0x0	0x0	0x0	0

As mentioned in the [Memory Controller, page 18](#), a Group FSM can issue one CAS command every three system clock cycles, or every 12 DRAM clock cycles, even for page hits. Therefore with only a single Group FSM issuing page hit commands to the DRAM for long periods, the maximum efficiency is 33%.

[Table 4-56](#) shows that the `ROW_COLUMN_BANK` option maps these same 12 addresses evenly across all eight DRAM banks and all four controller Group FSMs. This generates eight “page empty” transactions which open up all eight DRAM banks, followed by page hits to the open banks.

With all four Group FSMs issuing page hits, the efficiency can hit 100%, for as long as the address increment pattern continues, or until a refresh interrupts the pattern, or there is bus dead time for a DQ bus turnaround, etc. [Figure 4-15](#) shows the Group FSM issue over a larger address range for the `ROW_BANK_COLUMN` option. Note that the first 2k addresses map to two DRAM banks, but only one Group FSM.




Figure 4-15: DDR3 4 Gb (x8) Address Map ROW\_BANK\_COLUMN Graph

The address map graph for the ROW\_COLUMN\_BANK option is shown in Figure 4-16. Note that the address range in this graph is only 64 bytes, not 8k bytes. This graph is showing the same information as in the Address Decode in Table 4-56. With an address pattern that tends to stride through memory in minimum sized steps, efficiency tends to be High with the ROW\_COLUMN\_BANK option.




Figure 4-16: **DDR3 4 Gb (x8) Address Map ROW\_COLUMN\_BANK Graph**

Note that the ROW\_COLUMN\_BANK option does not result in High bus efficiency for all strides through memory. Consider the case of a stride of 16 bytes. This maps to only two Group FSMs resulting in a maximum efficiency of 67%. A stride of 32 bytes maps to only one Group FSM and the maximum efficiency is the same as the ROW\_BANK\_COLUMN option, just 33%. For an address pattern with variable strides, but strides that tend to be < 1k in the app\_addr address space, the ROW\_COLUMN\_BANK option is much more likely to result in good efficiency.

The same Group FSM issue exists for DDR4. With an address increment pattern and the DDR4 ROW\_BANK\_COLUMN option, the first 4k transactions map to a single Group FSM, as well as mapping to banks within a single DRAM bank group. The DRAM would limit the address increment pattern efficiency due to the tCCD\_L timing restriction. The controller limitation in this case is even more restrictive, due to the single Group FSM. Again the efficiency would be limited to 33%.

With the ROW\_COLUMN\_BANK option, the address increment pattern interleaves across all the DRAM banks and bank groups and all of the Group FSMs over a small address range.

Figure 4-17 shows how the DDR4 4 Gb (x8) ROW\_COLUMN\_BANK address map for the first 128 bytes of `app_addr`. This graph shows how the addresses map evenly across all DRAM banks and bank groups, and all four controller Group FSMs.

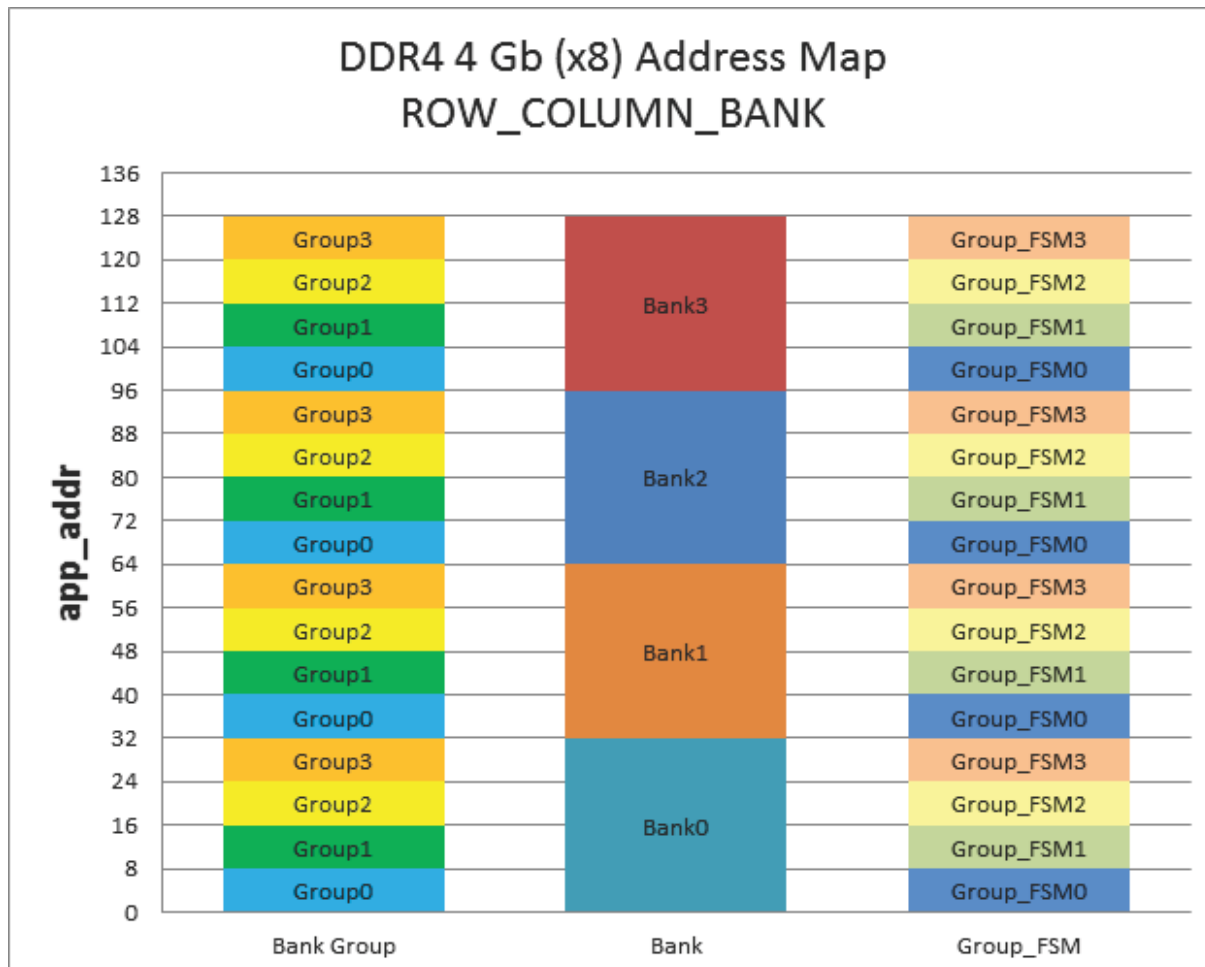


Figure 4-17: DDR4 4 Gb (x8) Address Map ROW\_COLUMN\_BANK Graph

When considering whether an address pattern at the user interface results in good DRAM efficiency, the mapping of the pattern to the controller Group FSMs is just as important as the mapping to the DRAM address. The `app_addr` bits that map `app_addr` addresses to the Group FSMs are shown in Table 4-57 for 4 Gb and 8 Gb components.

Table 4-57: DDR3/DDR4 Map Options for 4 Gb and 8 Gb

Memory Type	DDR3				DDR4		
Map Option	ROW_BANK_COLUMN			ROW_COLUMN_BANK	ROW_BANK_COLUMN		ROW_COLUMN_BANK
DRAM Component Width	x4	x8	x16	x4, x8, x16	x4, x8	x16	x4, x8, x16
Component Density	–	–	–	–	–	–	–

Table 4-57: DDR3/DDR4 Map Options for 4 Gb and 8 Gb (Cont'd)

Memory Type	DDR3				DDR4		
Map Option	ROW_BANK_COLUMN			ROW_COLUMN_BANK	ROW_BANK_COLUMN		ROW_COLUMN_BANK
4 Gb	13,12	12,11	12,11	4,3	13,12	12,10	4,3
8 Gb	14,13	13,12	12,11	4,3	13,12	12,10	4,3

Consider an example where you try to obtain good efficiency using only four DDR3 banks at a time. Assume you are using a 4 Gb (x8) with the ROW\_COLUMN\_BANK option and you decide to open a page in banks 0, 1, 2, and 3, and issue transactions to four column addresses in each bank. Using the address map from [Address Map](#), determine the `app_addr` pattern that decodes to this DRAM sequence. Applying the Group FSM map from [Table 4-57](#), determine how this `app_addr` pattern maps to the FSMs. The result is shown in [Table 4-58](#).

Table 4-58: Four Banks Sequence on DDR3 4 Gb (x8)

app_addr	Bank 0, 1, 2, 3 Sequence DDR3 4 Gb (x8) ROW_COLUMN_BANK			
	Row	Column	Bank	Group_FSM
0xE8	0x0	0x18	0x3	1
0xC8	0x0	0x18	0x2	1
0xE0	0x0	0x18	0x1	0
0xC0	0x0	0x18	0x0	0
0xA8	0x0	0x10	0x3	1
0x88	0x0	0x10	0x2	1
0xA0	0x0	0x10	0x1	0
0x80	0x0	0x10	0x0	0
0x68	0x0	0x8	0x3	1
0x48	0x0	0x8	0x2	1
0x60	0x0	0x8	0x1	0
0x40	0x0	0x8	0x0	0
0x28	0x0	0x0	0x3	1
0x08	0x0	0x0	0x2	1
0x20	0x0	0x0	0x1	0
0x00	0x0	0x0	0x0	0

The four bank pattern in [Table 4-58](#) works well from a DRAM point of view, but the controller only uses two of its four Group FSMs and the maximum efficiency is 67%. In practice it is even lower due to other timing restrictions like tRCD. A better bank pattern would be to open all the even banks and send four transactions to each as shown in [Table 4-59](#).

Table 4-59: Four Even Banks Sequence on DDR3 4 Gb (x8)

app_addr	Bank 0, 2, 4, 6 Sequence DDR3 4 Gb (x8) ROW_COLUMN_BANK			
	Row	Column	Bank	Group_FSM
0xD8	0x0	0x18	0x6	3
0xD0	0x0	0x18	0x4	2
0xC8	0x0	0x18	0x2	1
0xC0	0x0	0x18	0x0	0
0x98	0x0	0x10	0x6	3
0x90	0x0	0x10	0x4	2
0x88	0x0	0x10	0x2	1
0x80	0x0	0x10	0x0	0
0x58	0x0	0x8	0x6	3
0x50	0x0	0x8	0x4	2
0x48	0x0	0x8	0x2	1
0x40	0x0	0x8	0x0	0
0x18	0x0	0x0	0x6	3
0x10	0x0	0x0	0x4	2
0x08	0x0	0x0	0x2	1
0x00	0x0	0x0	0x0	0

The “even bank” pattern uses all of the Group FSMs and therefore has better efficiency than the previous pattern.

## Controller Head of Line Blocking and Look Ahead

As described in the [Memory Controller, page 18](#), each Group FSM has an associated transaction FIFO that is intended to improve efficiency by reducing “head of line blocking.” Head of line blocking occurs when one or more Group FSMs are fully occupied and cannot accept any new transactions for the moment, but the transaction presented to the user interface command port maps to one of the unavailable Group FSMs. This not only causes a delay in issuing new transactions to those busy FSMs, but to all the other FSMs as well, even if they are idle.

For good efficiency, you want to keep as many Group FSMs busy in parallel as you can. You could try changing the transaction presented to the user interface to one that maps to a different FSM, but you do not have visibility at the user interface as to which FSMs have space to take new transactions. The transaction FIFOs prevent this type of head of line blocking until a UI command maps to an FSM with a full FIFO.



A Group FSM FIFO structure can hold up to six transactions, depending on the page status of the target rank and bank. The FIFO structure is made up of two stages that also implement a “Look Ahead” function. New transactions are placed in the first FIFO stage and are operated on when they reach the head of the FIFO. Then depending on the transaction page status, the Group FSM either arbitrates to open the transaction page, or if the page is already open, the FSM pushes the page hit into the second FIFO stage. This scheme allows multiple page hits to be queued up while the FSM looks ahead into the logical FIFO structure for pages that need to be opened. Looking ahead into the queue allows an FSM to interleave DRAM commands for multiple transactions on the DDR bus. This helps to hide DRAM tRCD and tRP timing associated with opening and closing pages.

The following conceptual timing diagram shows the transaction flow from the UI to the DDR command bus, through the Group FSMs, for a series of transactions. The diagram is conceptual in that the latency from the UI to the DDR bus is not considered and not all DRAM timing requirements are met. Although not completely timing accurate, the diagram does follow DRAM protocol well enough to help explain the controller features under discussion.

Four transactions are presented at the UI, the first three mapping to the Group FSM0 and the fourth to FSM1. On system clock cycle 1, FSM0 accepts transaction 1 to Row 0, Column 0, and Bank 0 into its stage 1 FIFO and issues an Activate command.

On clock 2, transaction 1 is moved into the FSM0 stage 2 FIFO and transaction 2 is accepted into FSM0 stage 1 FIFO. On clock cycles 2 through 4, FSM0 is arbitrating to issue a CAS command for transaction 1, and an Activate command for transaction 2. FSM0 is looking ahead to schedule commands for transaction 2 even though transaction 1 is not complete. Note that the time when these DRAM commands win arbitration is determined by DRAM timing such as tRCD and controller pipeline delays, which explains why the commands are spaced on the DDR command bus as shown.

On cycle 3, transaction 3 is accepted into FSM0 stage 1 FIFO, but it is not processed until clock cycle 5 when it comes to the head of the stage 1 FIFO. Cycle 5 is where FSM0 begins looking ahead at transaction 3 while also arbitrating to issue the CAS command for transaction 2. Finally on cycle 4, transaction 4 is accepted into FSM1 stage 1 FIFO. If FSM0 did not have at least a three deep FIFO, transaction 4 would have been blocked until cycle 6.

Transaction Flow													
System Clock Cycle	1	2	3	4	5	6	7	8	9	10	11	12	13
UI Transaction Number	1	2	3	4	–	–	–	–	–	–	–	–	–
UI Transaction	R0, C0, B0	R0, C0, B1	R1, C0, B0	R0, C0, B2	–	–	–	–	–	–	–	–	–

Transaction Flow													
FSM0 FIFO Stage 2	–	R0, C0, B0	R0, C0, B0	R0, C0, B0	R0, C0, B1	R0, C0, B1	R0, B0, B1	–	–	R1, C0, B0	R1, C0, B0	R1, C0, B0	–
FSM0 FIFO Stage 1	R0, C0, B0	R0, C0, B1	R0, C0, B1 R1, C0, B0	R0, C0, B1 R1, C0, B0	R1, C0, B0	R1, C0, B0	R1, C0, B0	R1, C0, B0	R1, C0, B0	–	–	–	–
FSM1 FIFO Stage 2	–	–	–	–	–	R0, C0, B2	R0, C0, B2	R0, C0, B2	–	–	–	–	–
FSM1 FIFO Stage 1	–	–	–	R0, C0, B2	R0, C0, B2	–	–	–	–	–	–	–	–
DDR Command Bus	Act R0, B0	–	–	Act R0, B1	ACT R0, B2 CAS C0, B0	Pre B0	–	CAS C0, B1	Act R1, B0 CAS C0, B2	–	–	–	CAS C0, B0

This diagram does not show a high efficiency transaction pattern. There are no page hits and only two Group FSMs are involved. But the example does show how a single Group FSM interleaves DRAM commands for multiple transactions on the DDR bus and minimizes blocking of the UI, thereby improving efficiency.

## Autoprecharge

The Memory Controller defaults to a page open policy. It leaves banks open, even when there are no transactions pending. It only closes banks when a refresh is due, a page miss transaction is being processed, or when explicitly instructed to issue a transaction with a RDA or WRA CAS command. The `app_autoprecharge` port on the UI allows you to explicitly instruct the controller to issue a RDA or WRA command in the CAS command phase of processing a transaction, on a per transaction basis. You can use this signal to improve efficiency when you have knowledge of what transactions will be sent to the UI in the future.

The following diagram is a modified version of the “look ahead” example from the previous section. The page miss transaction that was previously presented to the UI in cycle 3 is now moved out to cycle 9. The controller can no longer “look ahead” and issues the Precharge to Bank 0 in cycle 6 because it does not know about the page miss until cycle 9. But if you know that transaction 1 in cycle 1 is the only transaction to Row 0 in Bank0, assert the `app_autoprecharge` port in cycle 1. Then, the CAS command for transaction 1 in cycle 5 is a RDA or WRA, and the transaction to Row 1, Bank 0 in cycle 9 is no longer a page miss.

The transaction in cycle 9 is only needed as an Activate command instead of a Precharge followed by an Activate tRP later.

Transaction Flow													
System Clock Cycle	1	2	3	4	5	6	7	8	9	10	11	12	13
UI Transaction Number	1	2	–	3	–	–	–	–	4	–	–	–	–
UI Transaction	R0, C0, B0 AutoPrecharge	R0, C0, B1	–	R0, C0, B2	–	–	–	–	R1, C0, B0	–	–	–	–
FSM0 FIFO Stage 2	–	R0, C0, B0	R0, C0, B0	R0, C0, B0	R0, C0, B1	R0, C0, B1	R0, B0, B1	–	–	R1, C0, B0	R1, C0, B0	R1, C0, B0	–
FSM0 FIFO Stage 1	R0, C0, B0	R0, C0, B1	R0, C0, B1	R0, C0, B1	–	–	–	–	R1, C0, B0	–	–	–	–
FSM1 FIFO Stage 2	–	–	–	–	–	R0, C0, B2	R0, C0, B2	R0, C0, B2	–	–	–	–	–
FSM1 FIFO Stage 1	–	–	–	R0, C0, B2	R0, C0, B2	–	–	–	–	–	–	–	–
DDR Command Bus	Act R0, B0	–	–	Act R0, B1	Act R0, B2 CAS-A C0, B0	–	–	CAS C0, B1	Act R1, B0 CAS C0, B2	–	–	–	CAS C0, B0

A general rule for improving efficiency is to assert `app_autoprecharge` on the last transaction to a page. An extreme example is an address pattern that never generates page hits. In this situation, it is best to assert `app_autoprecharge` on every transactions issued to the UI.

## User Refresh and ZQCS

The Memory Controller can be configured to automatically generate DRAM refresh and ZQCS maintenance commands to meet DRAM timing requirements. In this mode, the controller blocks the UI transactions on a regular basis to issue the maintenance commands, reducing efficiency.

If you have knowledge of the UI traffic pattern, you might be able to schedule DRAM maintenance commands with less impact on system efficiency. You can use the `app_ref` and `app_zq` ports at the UI to schedule these commands when the controller is configured for User Refresh and ZQCS. In this mode, the controller does not schedule the DRAM maintenance commands and only issues them based on the `app_ref` and `app_zq` ports. You are responsible for meeting all DRAM timing requirements for refresh and ZQCS.

Consider a case where the system needs to move a large amount of data into or out of the DRAM with the highest possible efficiency over a 50  $\mu$ s period. If the controller schedules the maintenance commands, this 50  $\mu$ s data burst would be interrupted multiple times for refresh, reducing efficiency roughly 4%. In User Refresh mode, however, you can decide to postpone refreshes during the 50  $\mu$ s burst and make them up later. The DRAM specification allows up to eight refreshes to be postponed, giving you flexibility to schedule refreshes over a  $9 \times t_{REFI}$  period, more than enough to cover the 50  $\mu$ s in this example.

While User Refresh and ZQCS enable you to optimize efficiency, their incorrect use can lead to DRAM timing violations and data loss in the DRAM. Use this mode only if you thoroughly understand DRAM refresh and ZQCS requirements as well as the operation of the `app_ref` and `app_zq` UI ports. The UI port operation is described in the [User Interface](#).

## Periodic Reads

The FPGA DDR PHY requires at least one DRAM RD or RDA command to be issued every 1  $\mu$ s. This requirement is described in the [User Interface](#). If this requirement is not met by the transaction pattern at the UI, the controller detects the lack of reads and injects a read transaction into Group FSM0. This injected read is issued to the DRAM following the normal mechanisms of the controller issuing transactions. The key difference is that no read data is returned to the UI. This is wasted DRAM bandwidth.

User interface patterns with long strings of write transactions are affected the most by the PHY periodic read requirement. Consider a pattern with a 50/50 read/write transaction ratio, but organized such that the pattern alternates between 2  $\mu$ s bursts of 100% page hit reads and 2  $\mu$ s bursts of 100% page hit writes. There is at least one injected read in the 2  $\mu$ s write burst, resulting in a loss of efficiency due to the read command and the turnaround time to switch the DRAM and DDR bus from writes to reads back to writes. This 2  $\mu$ s alternating burst pattern is slightly more efficient than alternating between reads and writes every 1  $\mu$ s. A 1  $\mu$ s or shorter alternating pattern would eliminate the need for the controller to inject reads, but there would still be more read-write turnarounds.

Bus turnarounds are expensive in terms of efficiency and should be avoided if possible. Long bursts of page hit writes, > 2  $\mu$ s in duration, are still the most efficient way to write to the DRAM, but the impact of one write-read-write turnaround each 1  $\mu$ s must be taken into account when calculating the maximum write efficiency.

# Design Flow Steps

This chapter describes customizing and generating the core, constraining the core, and the simulation, synthesis and implementation steps that are specific to this IP core. More detailed information about the standard Vivado® design flows and the Vivado IP integrator can be found in the following Vivado Design Suite user guides:

- *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#)
- *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#)
- *Vivado Design Suite User Guide: Getting Started* (UG910) [\[Ref 9\]](#)
- *Vivado Design Suite User Guide: Logic Simulation* (UG900) [\[Ref 10\]](#)

---

## Customizing and Generating the Core



**CAUTION!** The Windows operating system has a 260-character limit for path lengths, which can affect the Vivado tools. To avoid this issue, use the shortest possible names and directory locations when creating projects, defining IP or managed IP projects, and creating block designs.

---

This section includes information about using Xilinx® tools to customize and generate the core in the Vivado Design Suite.

If you are customizing and generating the core in the IP integrator, see the *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#) for detailed information. IP integrator might auto-compute certain configuration values when validating or generating the design. To check whether the values change, see the description of the parameter in this chapter. To view the parameter value, run the `validate_bd_design` command in the Tcl Console.

You can customize the IP for use in your design by specifying values for the various parameters associated with the IP core using the following steps:

1. Select the IP from the Vivado IP catalog.
2. Double-click the selected IP or select the Customize IP command from the toolbar or right-click menu.

For more information about generating the core in Vivado, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8] and the *Vivado Design Suite User Guide: Getting Started* (UG910) [Ref 9].

**Note:** Figures in this chapter are illustrations of the Vivado Integrated Design Environment (IDE). This layout might vary from the current version.

## Basic Tab

Figure 5-1 shows the **Basic** tab when you start up the MIG.

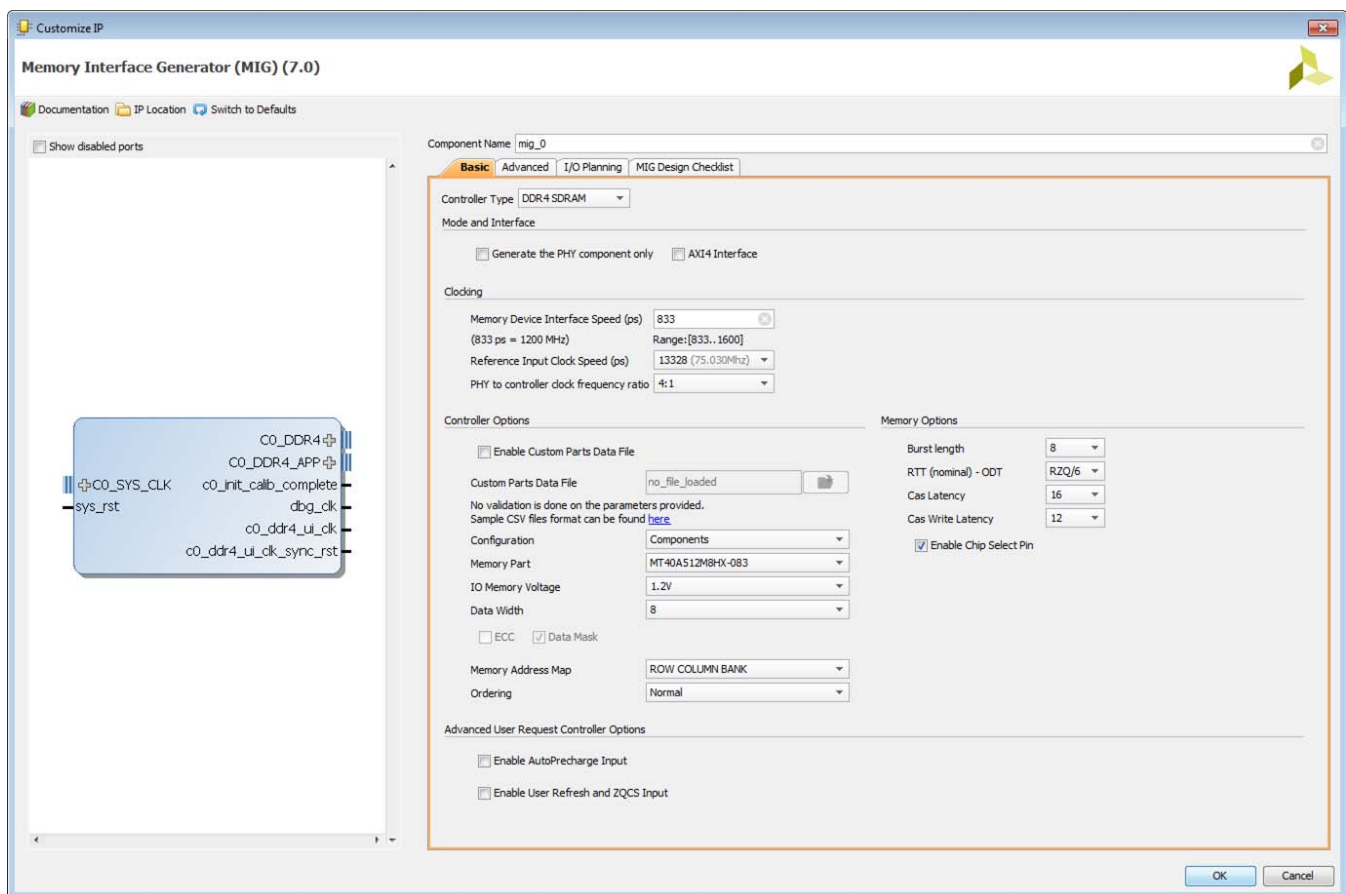


Figure 5-1: Vivado Customize IP Dialog Box – Basic Tab

For the Vivado IDE, all controllers (DDR3, DDR4, QDR II+, and RLDRAM 3) can be created and available for instantiation.

In IP integrator, only one controller instance can be created and only two kinds of controllers are available for instantiation:

- DDR3
- DDR4

1. After a controller is added in the pull-down menu, select the **Mode and Interface** for the controller. Select the **AXI4 Interface** or have the option to select the **Generate the PHY component only**.
2. Select the settings in the **Clocking, Controller Options, Memory Options, AXI Options**, and **Advanced User Request Controller Options**.

In **Clocking**, the **Memory Device Interface Speed** sets the speed of the interface. The speed entered drives the available **Reference Input Clock Speeds**. For more information on the clocking structure, see the [Clocking, page 73](#).

3. To use memory parts which are not available by default through the MIG GUI, you can create a custom parts CSV file, as specified in the AR: [63462](#). This CSV file has to be provided after enabling the **Custom Parts Data File** option. After selecting this option, you are able to see the custom memory parts along with the default memory parts.




---

**IMPORTANT:** *Data Mask (DM) option is always selected for AXI designs and is grayed out (you cannot select it). For AXI interfaces, Read Modify Write (RMW) is supported and for RMW to mask certain bytes of Data Mask bits should be present. Therefore, the DM is always enabled for AXI interface designs. This is the case for all data widths except 72-bit.*

*For 72-bit interfaces, ECC is enabled and DM is deselected and grayed out for 72-bit designs. If DM is enabled for 72-bit designs, computing ECC does is not compatible, so DM is disabled for 72-bit designs.*

---

## Advanced Tab

Figure 5-2 shows the next tab called **Advanced**. This displays the settings for **FPGA Options**, **Debug Signals for Controller**, **Simulation Options**, and **Clock Options** for the specific controller.

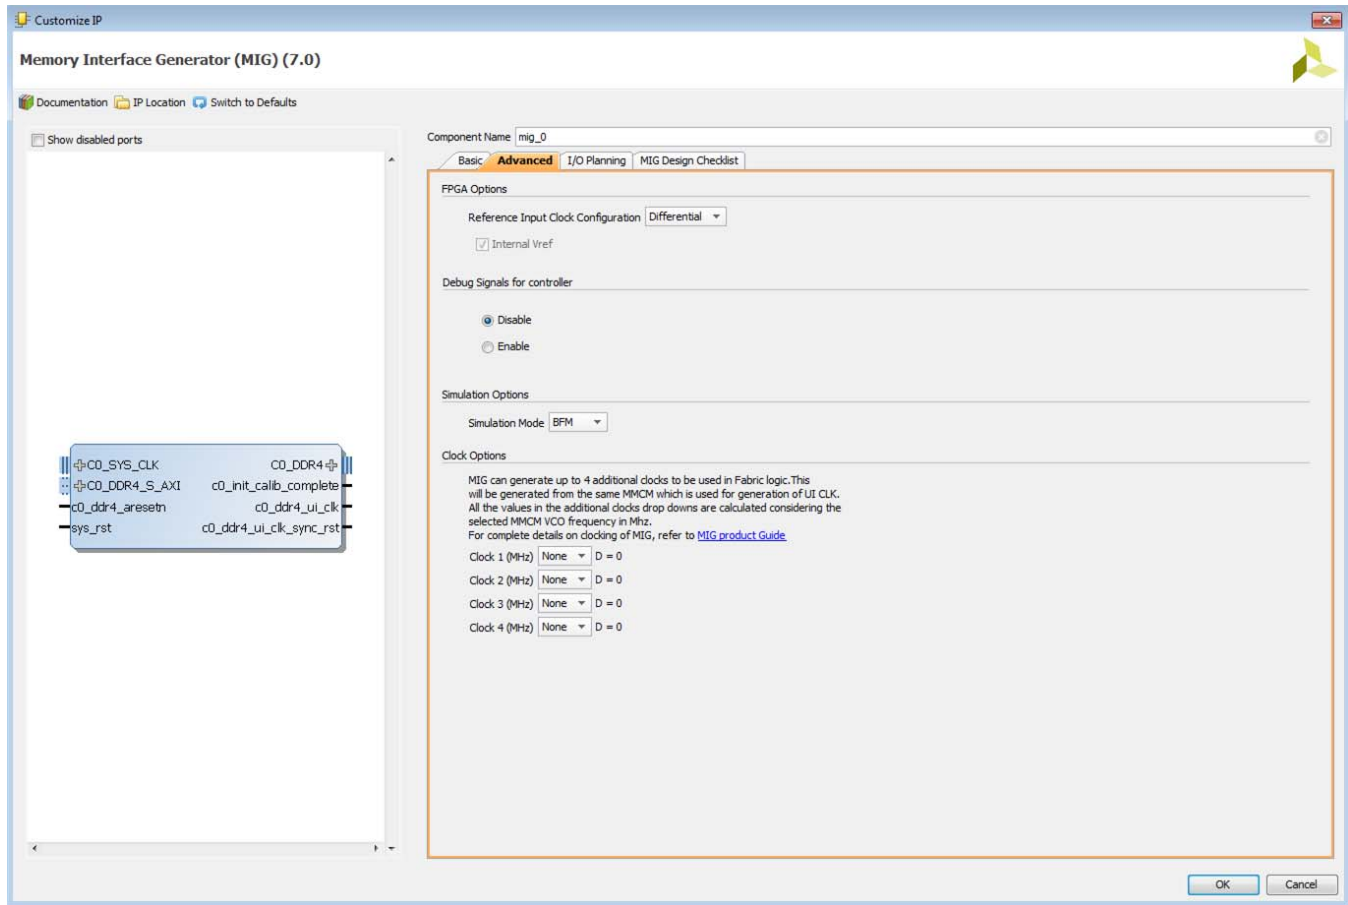


Figure 5-2: Vivado Customize IP Dialog Box – Advanced



**IMPORTANT:** All parameters shown in the controller options dialog box are limited selection options in this release.



## MIG I/O Planning Tab

Figure 5-3 shows the MIG **I/O Planning** tab and informs you that I/O planning has been taken away from the customization window.

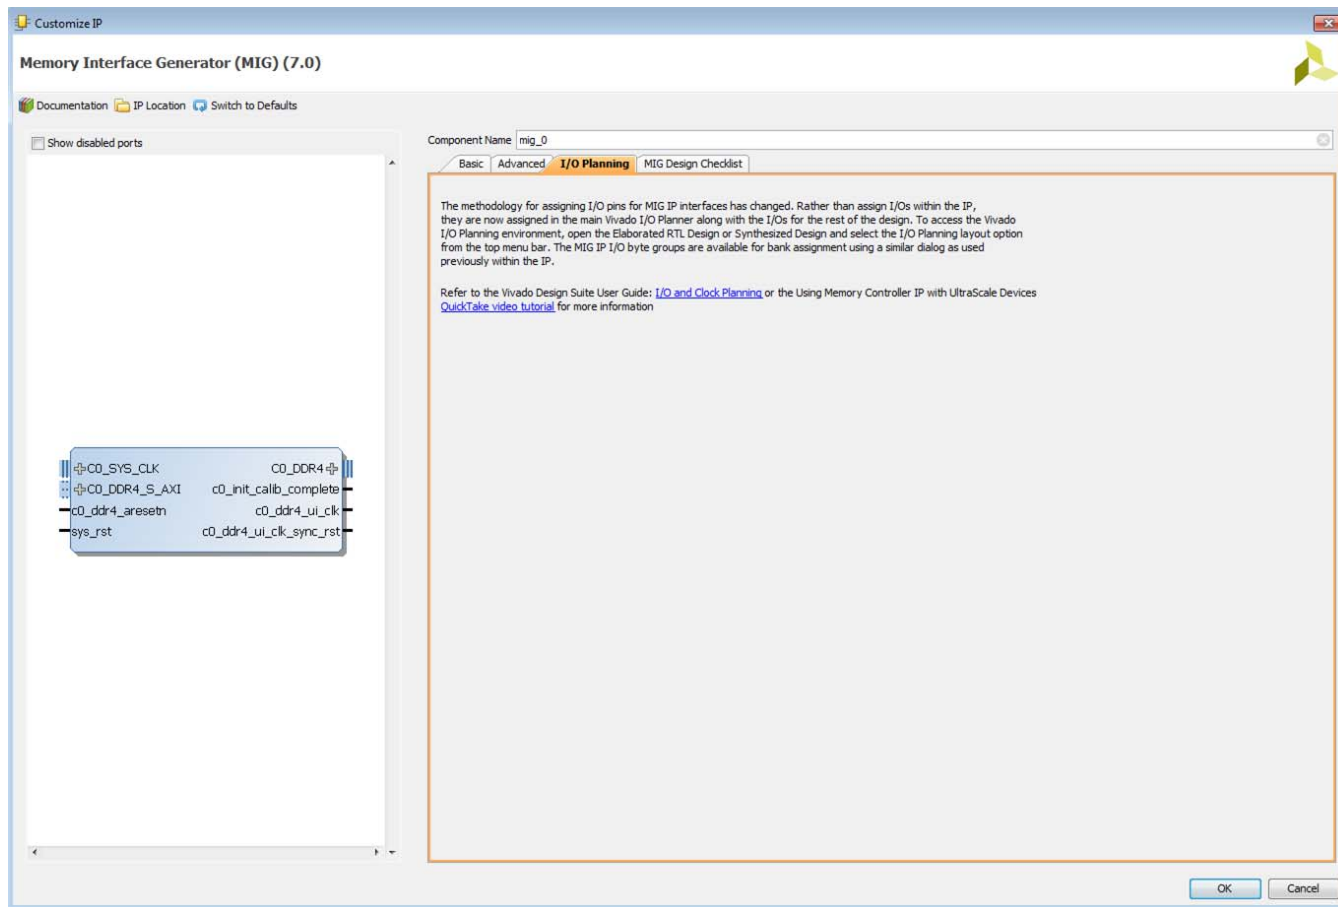


Figure 5-3: Vivado Customize IP Dialog Box – MIG I/O Planning Tab

For more information on the MIG I/O planning, see the [MIG I/O Planning](#).

## MIG Design Checklist Tab

Figure 5-4 shows the **MIG Design Checklist** usage information.

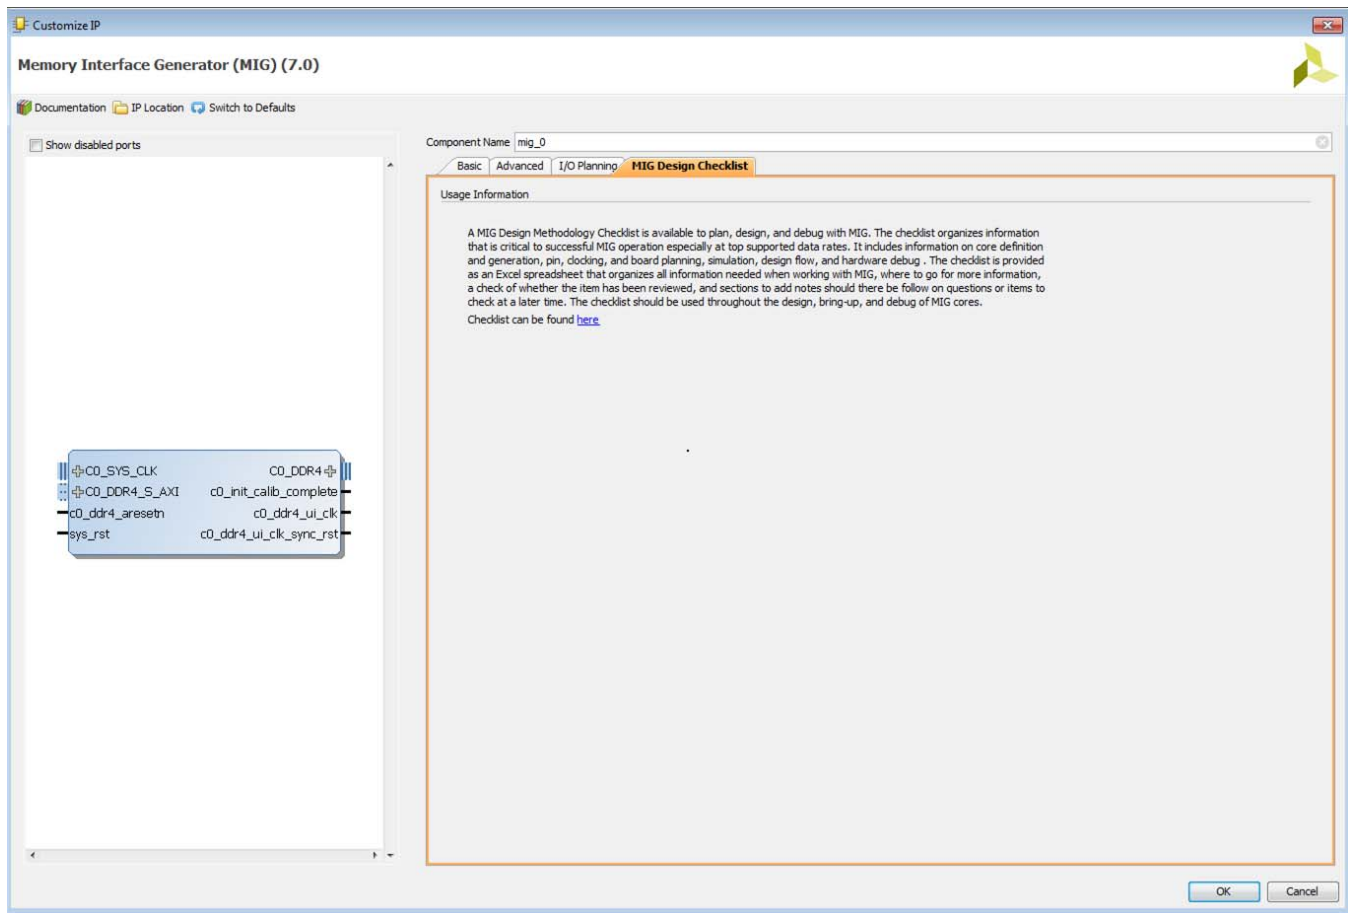


Figure 5-4: Vivado Customize IP Dialog Box – MIG Design Checklist Tab

## User Parameters

Table 5-1 shows the relationship between the GUI fields in the Vivado IDE and the User Parameters (which can be viewed in the Tcl Console).

Table 5-1: Vivado IDE Parameter to User Parameter Relationship

Vivado IDE Parameter/Value <sup>(1)</sup>	User Parameter/Value <sup>(1)</sup>	Default Value
System Clock Configuration	System_Clock	Differential
Internal V <sub>REF</sub>	Internal_Vref	TRUE
DCI Cascade	DCI_Cascade	FALSE
Debug Signal for Controller	Debug_Signal	Disable
Clock 1 (MHz)	ADDN_UI_CLKOUT1_FREQ_HZ	None
Clock 2 (MHz)	ADDN_UI_CLKOUT2_FREQ_HZ	None
Clock 3 (MHz)	ADDN_UI_CLKOUT3_FREQ_HZ	None

Table 5-1: Vivado IDE Parameter to User Parameter Relationship (Cont'd)

Vivado IDE Parameter/Value <sup>(1)</sup>	User Parameter/Value <sup>(1)</sup>	Default Value
Clock 4 (MHz)	ADDN_UI_CLKOUT4_FREQ_HZ	None
I/O Power Reduction	IOPowerReduction	OFF
Enable System Ports	Enable_SysPorts	TRUE
I/O Power Reduction	IO_Power_Reduction	FALSE
Default Bank Selections	Default_Bank_Selections	FALSE
Reference Clock	Reference_Clock	FALSE
Enable System Ports	Enable_SysPorts	TRUE
<b>DDR3</b>		
AXI4 Interface	C0.DDR3_AxiSelection	FALSE
Clock Period (ps)	C0.DDR3_TimePeriod	1,071
Input Clock Period (ps)	C0.DDR3_InputClockPeriod	13,947
General Interconnect to Memory Clock Ratio	C0.DDR3_PhyClockRatio	4:1
Data Width	C0.DDR3_AxiDataWidth	64
Arbitration Scheme	C0.DDR3_AxiArbitrationScheme	RD_PRI_REG
Address Width	C0.DDR3_AxiAddressWidth	27
AXI4 Narrow Burst	C0.DDR3_AxiNarrowBurst	FALSE
Configuration	C0.DDR3_MemoryType	Components
Memory Part	C0.DDR3_MemoryPart	MT41J128M16JT-093
Data Width	C0.DDR3_DataWidth	8
Data Mask	C0.DDR3_DataMask	TRUE
Burst Length	C0.DDR3_BurstLength	8
R <sub>TT</sub> (nominal)-ODT	C0.DDR3_OnDieTermination	RZQ/4
CAS Latency	C0.DDR3_CasLatency	11
CAS Write Latency	C0.DDR3_CasWriteLatency	9
Chip Select	C0.DDR3_ChipSelect	TRUE
Memory Address Map	C0.DDR3_Mem_Add_Map	ROW_COLUMN_BANK
Memory Voltage	C0.DDR3_MemoryVoltage	1.5
ECC	C0.DDR3_Ecc	FALSE
Ordering	C0.DDR3_Ordering	Normal
Burst Type	C0.DDR3_BurstType	Sequential
Output Driver Impedance Control	C0.DDR3_OutputDriverImpedanceControl	RZQ/7
AXI ID Width	C0.DDR3_AxiIDWidth	4
Capacity	C0.DDR3_Capacity	512
<b>DDR4</b>		
AXI4 Interface	C0.DDR4_AxiSelection	FALSE
Clock Period (ps)	C0.DDR4_TimePeriod	938
Input Clock Period (ps)	C0.DDR4_InputClockPeriod	104,045

Table 5-1: Vivado IDE Parameter to User Parameter Relationship (Cont'd)

Vivado IDE Parameter/Value <sup>(1)</sup>	User Parameter/Value <sup>(1)</sup>	Default Value
General Interconnect to Memory Clock Ratio	C0.DDR4_PhyClockRatio	4:1
Data Width	C0.DDR4_AxiDataWidth	64
Arbitration Scheme	C0.DDR4_AxiArbitrationScheme	RD_PRI_REG
Address Width	C0.DDR4_AxiAddressWidth	27
AXI4 Narrow Burst	C0.DDR4_AxiNarrowBurst	FALSE
Configuration	C0.DDR4_MemoryType	Components
Memory Part	C0.DDR4_MemoryPart	MT40A256M16HA-083
Data Width	C0.DDR4_DataWidth	8
Data Mask	C0.DDR4_DataMask	TRUE
Burst Length	C0.DDR4_BurstLength	8
R <sub>TT</sub> (nominal)-ODT	C0.DDR4_OnDieTermination	RZQ/6
CAS Latency	C0.DDR4_CasLatency	14
CAS Write Latency	C0.DDR4_CasWriteLatency	11
Chip Select	C0.DDR4_ChipSelect	TRUE
Memory Address Map	C0.DDR4_Mem_Add_Map	ROW_COLUMN_BANK
Memory Voltage	C0.DDR4_MemoryVoltage	1.2
ECC	C0.DDR4_Ecc	FALSE
Ordering	C0.DDR4_Ordering	Normal
Burst Type	C0.DDR4_BurstType	Sequential
Output Driver Impedance Control	C0.DDR4_OutputDriverImpedanceControl	RZQ/7
AXI ID Width	C0.DDR4_AxiIDWidth	4
Capacity	C0.DDR4_Capacity	512

**Notes:**

- Parameter values are listed in the table where the Vivado IDE parameter value differs from the user parameter value. Such values are shown in this table as indented below the associated parameter.

## Output Generation

For details, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

## MIG I/O Planning

MIG I/O pin planning is completed with the full design pin planning using the Vivado I/O Pin Planner. MIG I/O pins can be selected through several Vivado I/O Pin Planner features including assignments using I/O Ports view, Package view, or Memory Bank/Byte Planner. Pin assignments can additionally be made through importing an XDC or modifying the existing XDC file.

These options are available for all MIG designs and multiple MIG IP instances can be completed in one setting. To learn more about the available MIG pin planning options, see the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 12].

## Constraining the Core

This section contains information about constraining the core in the Vivado Design Suite.

### Required Constraints

For MIG Vivado IDE, you specify the pin location constraints. For more information on I/O standard and other constraints, see the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 12]. The location is chosen by the Vivado IDE according to the banks and byte lanes chosen for the design.

The I/O standard is chosen by the memory type selection and options in the Vivado IDE and by the pin type. A sample for `dq[0]` is shown here.

```
set_property PACKAGE_PIN AF20 [get_ports "c0_ddr4_dq[0]"]
set_property IOSTANDARD POD12_DCI [get_ports "c0_ddr4_dq[0]"]
```

Internal  $V_{REF}$  is always used for DDR4. Internal  $V_{REF}$  is optional for DDR3. A sample for DDR4 is shown here.

```
set_property INTERNAL_VREF 0.600 [get_iobanks 45]
```

**Note:** Internal  $V_{REF}$  is automatically generated by the tool and you do not need to specify it. The  $V_{REF}$  value listed in this constraint is not used with POD12 I/Os. The initial value is set to 0.84V. The calibration logic adjusts this voltage as needed for maximum interface performance.

The system clock must have the period set properly:

```
create_clock -name c0_sys_clk -period.938 [get_ports c0_sys_clk_p]
```

For HR banks, update the `output_impedance` of all the ports assigned to HR banks pins using the `reset_property` command. For more information, see AR: [63852](#).



**IMPORTANT:** Do not alter these constraints. If the pin locations need to be altered, rerun the MIG Vivado IDE to generate a new XDC file.

## Device, Package, and Speed Grade Selections

This section is not applicable for this IP core.

## Clock Frequencies

This section is not applicable for this IP core.

## Clock Management

For more information on clocking, see [Clocking, page 73](#).

## Clock Placement

This section is not applicable for this IP core.

## Banking

This section is not applicable for this IP core.

## Transceiver Placement

This section is not applicable for this IP core.

## I/O Standard and Placement

The MIG tool generates the appropriate I/O standards and placement based on the selections made in the Vivado IDE for the interface type and options.



---

**IMPORTANT:** *The `set_input_delay` and `set_output_delay` constraints are not needed on the external memory interface pins in this design due to the calibration process that automatically runs at start-up. Warnings seen during implementation for the pins can be ignored.*

---

---

## Simulation

For comprehensive information about Vivado simulation components, as well as information about using supported third-party tools, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [\[Ref 10\]](#). For more information on simulation, see [Chapter 6, Example Design](#) and [Chapter 7, Test Bench](#).

---

## Synthesis and Implementation

For details about synthesis and implementation, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#).

## Example Design

This chapter contains information about the example design provided in the Vivado® Design Suite.

Vivado supports Open IP Example Design flow. To create the example design using this flow, right-click the IP in the **Source Window**, as shown in [Figure 6-1](#) and select **Open IP Example Design**.




Figure 6-1: Open IP Example Design

This option creates a new Vivado project. Upon selecting the menu, a dialog box to enter the directory information for the new design project opens.

Select a directory, or use the defaults, and click **OK**. This launches a new Vivado with all of the example design files and a copy of the IP.

Figure 6-2 shows the example design with the PHY only option selected (controller module does not get generated).




Figure 6-2: Open IP Example Design with PHY Only Option Selected



Figure 6-3 shows the example design with the PHY only option not selected (controller module is generated).

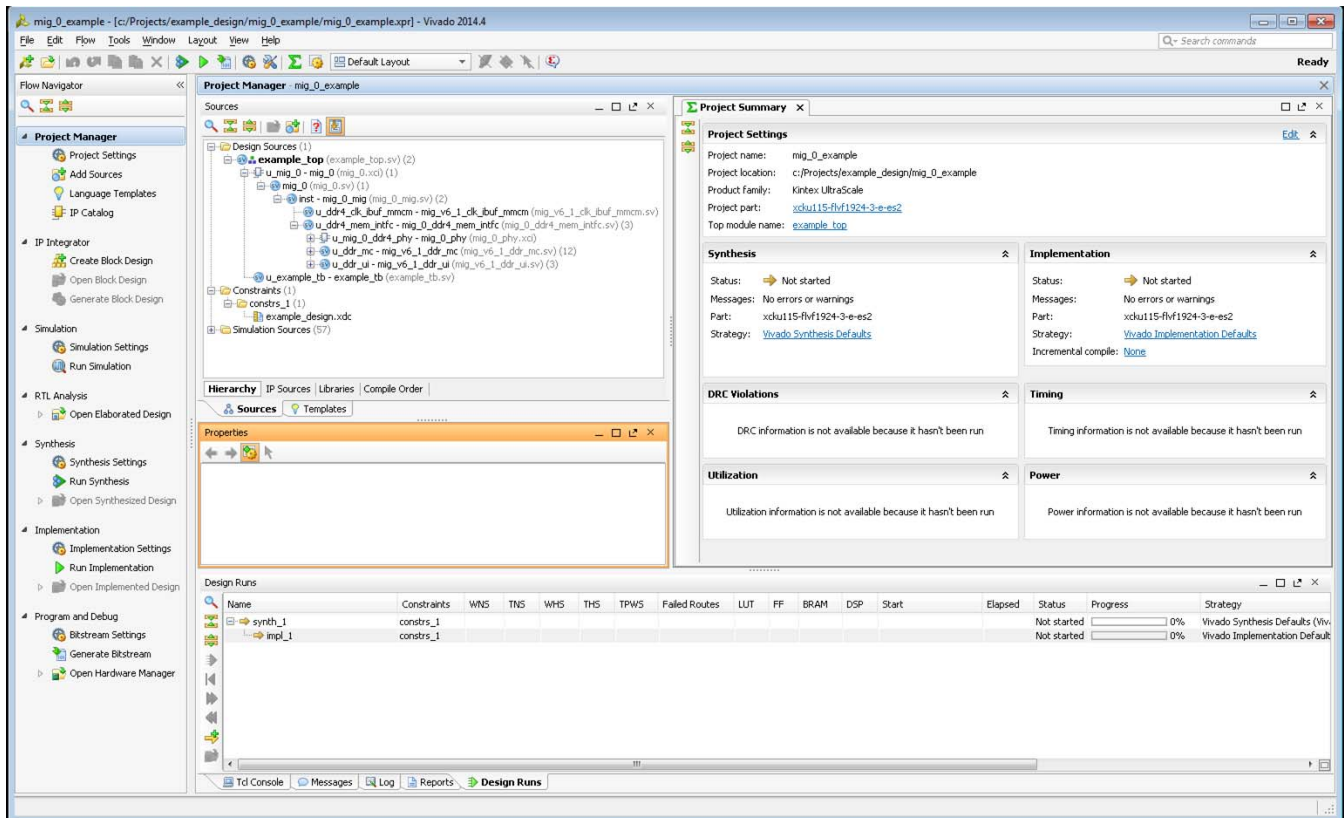


Figure 6-3: Open IP Example Design with PHY Only Option Not Selected

## Simulating the Example Design (Designs with Standard User Interface)

The example design provides a synthesizable test bench to generate a fixed simple data pattern. MIG generates the Simple Traffic Generator (STG) module as `example_tb` for native interface and `example_tb_phy` for PHY only interface. The STG native interface generates 100 writes and 100 reads. The STG PHY only interface generates 10 writes and 10 reads.

The example design can be simulated using one of the methods in the following sections.

## Project-Based Simulation

This method can be used to simulate the example design using the Vivado Integrated Design Environment (IDE). MIG delivers memory models for DDR3 and IEEE encrypted memory models for DDR4.

The Vivado simulator, QuestaSim, IES, and VCS tools are used for DDR3/DDR4 IP verification at each software release. The Vivado simulation tool is used for DDR3/DDR4 IP verification from 2015.1 Vivado software release. The following subsections describe steps to run a project-based simulation using each supported simulator tool.

### Project-Based Simulation Flow Using Vivado Simulator

1. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as **Vivado Simulator**.
  - a. Under the **Simulation** tab, set the `xsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in [Figure 6-4](#). For DDR3 simulation, set the `xsim.simulate.xsim.more_options` to `-testplusarg model_data+./`. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
3. Apply the settings and select **OK**.




Figure 6-4: Simulation with Vivado Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 6-5.




Figure 6-5: Run Behavioral Simulation

5. Vivado invokes Vivado simulator and simulations are run in the Vivado simulator tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using QuestaSim

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as **QuestaSim/ModelSim Simulator**.
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `modelsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 6-6. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected. For DDR3 simulation, set the `modelsim.simulate.vsim.more_options` to `+model_data+./.`
3. Apply the settings and select **OK**.




Figure 6-6: Simulation with QuestaSim

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 6-7.




Figure 6-7: Run Behavioral Simulation

5. Vivado invokes QuestaSim and simulations are run in the QuestaSim tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using IES

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as Incisive Enterprise Simulator (IES).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `ies.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 6-8. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected. For DDR3 simulation, set the `modelsim.simulate.vsim.more_options` to `+model_data+./`.
3. Apply the settings and select **OK**.




Figure 6-8: Simulation with IES Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 6-7.
5. Vivado invokes IES and simulations are run in the IES tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].



## Project-Based Simulation Flow Using VCS

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as Verilog Compiler Simulator (VCS).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `vcs.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in [Figure 6-9](#). The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected. For DDR3 simulation, set the `modelsim.simulate.vsim.more_options` to `+model_data+./.`
3. Apply the settings and select **OK**.






Figure 6-9: Simulation with VCS Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 6-7.
5. Vivado invokes VCS and simulations are run in the VCS tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation (UG900)* [Ref 10].

## Non-Project-Based Simulation



**IMPORTANT:** Xilinx® UNISIMS\_VER and SECUREIP library must be mapped into the simulator.

1. To run the simulation, go to this directory:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srscs/sim_1/imports/tb
```

If the MIG design is generated with the Component Name entered in the Vivado IDE as mig\_0, the simulation directory path is the following:

```
<project_dir>/example_project/mig_0_example/mig_0_example.srscs/  
sim_1/imports/tb
```

2. MIG delivers memory models for DDR3 and IEEE encrypted memory models for DDR4. Copy the memory models in the above directory for DDR4.
3. The QuestaSim, IES, and VCS simulation tools are used for verification of MIG IP at each software release.
4. Script files to run simulations with QuestaSim, IES, and VCS are generated in MIG generated output. See the `readme.txt` file located in the folder for running simulations. Other simulation tools can be used for MIG IP simulation but are not specifically verified by Xilinx.

## Simulation Speed

MIG provides a Vivado IDE option to reduce the simulation speed by selecting behavioral XIPHY model instead of UNISIM XIPHY model. Behavioral XIPHY model simulation is a default option for MIG designs. To select the simulation mode, click the **Advanced** tab and find the **Simulation Options** as shown in [Figure 6-10](#).




Figure 6-10: Advanced Tab – Simulation Options

The SIM\_MODE parameter in the RTL is given a different value based on the Vivado IDE selection.

- **SIM\_MODE = BFM** – If fast mode is selected in the Vivado IDE, the RTL parameter reflects this value for the SIM\_MODE parameter. This is the default option.
- **SIM\_MODE = FULL** – If FULL mode is selected in the Vivado IDE, XIPHY UNISIMs are selected and the parameter value in the RTL is FULL.

## Synplify Pro Black Box Testing

Using the Synopsys® Synplify Pro® black box testing for `example_design`, follow these steps to run black box synthesis with `synplify_pro` and implementation with Vivado.

1. Generate the UltraScale™ architecture MIG IP core with OOC flow to generate the `.dcp` file for implementation. The **Target Language** for the project can be selected as **verilog** or **VHDL**.
2. Create the example design for the MIG IP core using the information provided in the example design section and close the Vivado project.

3. Invoke the `synplify_pro` software which supports UltraScale FPGA and select the same UltraScale FPGA part selected at the time of generating the IP core.
4. Add the following files into `synplify_pro` project based on the **Target Language** selected at the time of invoking Vivado:

- a. For Verilog:

```
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/*stub.v
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>rtl/ip_top/
example_top.sv
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>tb/example_tb.sv
```

- b. For VHDL:

```
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/*stub.vhdl
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>rtl/ip_top/
example_top.sv
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>tb/example_tb.sv
```

5. Run `synplify_pro` synthesis to generate the `.edf` file. Then, close the `synplify_pro` project.
6. Open new Vivado project with Project Type as **Post-synthesis Project** and select the **Target Language** same as selected at the time of generating the IP core.
7. Add the `synplify_pro` generated `.edf` file to the Vivado project as **Design Source**.
8. Add the `.dcp` file generated in steps 1 and 2 to the Vivado project as **Design Source**. For example:

```
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/<Component_Name>.dcp
```

9. Add the `.xdc` file generated in step 2 to the Vivado project as **constraint**. For example:

```
<project_dir>/example_project/<Component_Name>example/
<Component_Name>_example.srcs/constrs_1/imports/par/example_design.xdc
```

10. Run implementation flow with the Vivado tool. For details about implementation, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

**Note:** Similar steps can be followed for the user design using appropriate `.dcp` and `.xdc` files.

---

## CLOCK\_DEDICATED\_ROUTE Constraints and BUFG Instantiation

If the GCIO pin and MMCM are not allocated in the same bank, the CLOCK\_DEDICATED\_ROUTE constraint must be set to BACKBONE. To use the BACKBONE route, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM input. MIG manages these constraints for designs generated with the **Reference Input Clock** option selected as **Differential** (at **Advanced > FPGA Options > Reference Input**). Also, MIG handles the IP and example design flows for all scenarios.

If the design is generated with the **Reference Input Clock** option selected as **No Buffer** (at **Advanced > FPGA Options > Reference Input**), the CLOCK\_DEDICATED\_ROUTE constraints and BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV instantiation based on GCIO and MMCM allocation needs to be handled manually for the IP flow. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations for the IP flow.

For an example design flow with **No Buffer** configurations, MIG generates the example design with differential buffer instantiation for system clock pins. MIG generates clock constraints in the `example_design.xdc`. It also generates a CLOCK\_DEDICATED\_ROUTE constraint as the "BACKBONE" and instantiates BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV between GCIO and MMCM input if the GCIO and MMCM are not in same bank to provide a complete solution. This is done for the example design flow as a reference when it is generated for the first time.

If in the example design, the I/O pins of the system clock pins are changed to some other pins with the I/O pin planner, the CLOCK\_DEDICATED\_ROUTE constraints and BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV instantiation needs to be managed manually. A DRC error is reported for the same.

## Test Bench

This chapter contains information about the test bench provided in the Vivado® Design Suite.

The intent of the performance test bench is for you to obtain an estimate on the efficiency for a given traffic pattern with the MIG controller. The test bench passes your supplied commands and address to the Memory Controller and measures the efficiency for the given pattern. The efficiency is measured by the occupancy of the `axi` bus. The primary use of the test bench is for efficiency measurements so no data integrity checks are performed. Static data is written into the memory during write transactions and the same data is always read back.

The stimulus to the Traffic Generator is provided through a `mig_v7_0_ddr4_stimulus.txt` file. The stimulus consists of command, address, and command repetition count. Each line in the stimulus file represents one stimulus (command repetition, address, and command). Multiple stimuli can be provided in a stimulus file and each stimulus is separated by the new line.

**Table 7-1: Modules for Performance Traffic Generator**

File Name	Description
<code>mig_v7_0_ddr4_traffic_generator.sv</code>	This file has the Traffic Generator code for sending out the traffic for DDR4 and also for the calculation of bus utilized.
<code>mig_v7_0_ddr4_stimulus.txt</code>	These files have the stimulus with Writes, Reads, and NOPs for DDR4 for the calculation of bus utilization.
<code>mig_v7_0_ddr3_traffic_generator.sv</code>	This file has the Traffic Generator code for sending out the traffic for DDR3 and also for the calculation of bus utilized.
<code>mig_v7_0_ddr3_stimulus.txt</code>	These files have the stimulus with Writes, Reads, and NOPs for DDR3 for the calculation of bus utilization.

## Stimulus Pattern

Each stimulus pattern is 48 bits and the format is described in [Table 7-2](#) and [Table 7-3](#).

**Table 7-2: Stimulus Command Pattern**

Command Repeat[47:40]	Address [39:4]	Command[3:0]
-----------------------	----------------	--------------

**Table 7-3: Stimulus Pattern Description**

Signal	Description
Command[3:0]	This corresponds to the WRITE/READ/NOP command that is sent to the user interface.
Address[35:0]	This corresponds to the address to the user interface.
Command Repeat[7:0]	This corresponds to the repetition count of the command. Up to 128 repetitions can be made for a command. In the burst length of eight mode, 128 transactions fill up the page in the memory.

## Command Encoding (Command[3:0])

**Table 7-4: Command Description**

Command	Code	Description
WRITE	0	This corresponds to the Write operation that needs to be performed.
READ	1	This corresponds to the Read operation that needs to be performed.
NOP	7	This corresponds to the idle situation for the bus.

## Address Encoding (Address[35:0])

Address is encoded in the stimulus as per [Figure 7-1](#) to [Figure 7-6](#). All the address fields need to be entered in the hexadecimal format. All the address fields are the width that is divisible by four to enter in the hexadecimal format. The test bench only sends the required bits of an address field to the Memory Controller.

For example, an eight bank configuration only bank Bits[2:0] is sent to the Memory Controller and the remaining bits are ignored. The extra bits for an address field are provided for you to enter the address in a hexadecimal format. You must confirm the value entered corresponds to the width of a given configuration.

**Table 7-5: Address Encoded**

Rank[3:0]	Bank[3:0]	Row[15:0]	Column[11:0]
-----------	-----------	-----------	--------------

- **Column Address (Column[11:0])** – Column Address in the stimulus is provided maximum of 12 bits, but you need to address this based on the column width parameter set in your design.

- **Row Address (Row[15:0])** – Row address in the stimulus is provided maximum of 16 bits, but you need to address this based on the row width parameter set in your design.
- **Bank Address (Bank[3:0])** – Bank address in the stimulus is provided maximum of four bits, but you need to address this based on the bank width parameter set in your design.

**Note:** For DDR4, use the 2-bit LSB for Bank Address and two bits of MSB for Bank Groups.

- **Rank Address (Rank[3:0])** – Rank address in the stimulus is provided maximum of four bits, but you need to address this based on the rank width parameter set in your design.

The address is assembled based on the top-level MEM\_ADDR\_ORDER parameter and sent to the user interface.

## Command Repeat (Command Repeat[7:0])

The command repetition count is the number of time the respective command is repeated at the User Interface. The address for each repetition is incremented by 8. The maximum repetition count is 128. The test bench does not check for the column boundary and it wraps around if the maximum column limit is reached during the increments. The 128 commands fill up the page. For any column address other than 0, the repetition count of 128 ends up crossing the column boundary and wrapping around to the start of the column address.

---

## Bus Utilization

The bus utilization is calculated at the User Interface taking total number of Reads and Writes into consideration and the following equation is used:

$$((rd\_command\_cnt + wr\_command\_cnt) \times (BURST\_LEN / 2) \times 100) \quad \text{Equation 7-1}$$

bw\_cumulative = -----

$$((end\_of\_stimulus - calib\_done) / tCK);$$

- BL8 takes four memory clock cycles.
- end\_of\_stimulus is the time when all the commands are done.
- calib\_done is the time when the calibration is done.



## Example Patterns

These examples are based on the MEM\_ADDR\_ORDER set to BANK\_ROW\_COLUMN.

### Single Read Pattern

**00\_0\_2\_000F\_00A\_1** – This pattern is a single read from 10<sup>th</sup> column, 15<sup>th</sup> row, and second bank.




Figure 7-1: Single Read Pattern

### Single Write Pattern

**00\_0\_1\_0040\_010\_0** – This pattern is a single write to the 32<sup>nd</sup> column, 128<sup>th</sup> row, and first bank.





Figure 7-2: Single Write Pattern

**00\_0\_2\_000F\_00A\_0** – This pattern is a single write to 10<sup>th</sup> column, 15<sup>th</sup> row, and second bank.

**00\_0\_2\_000F\_00A\_1** – This pattern is a single read from 10<sup>th</sup> column, 15<sup>th</sup> row, and second bank.



**0A\_0\_0\_0010\_000\_0** – This corresponds to 11 writes with address starting from 0 to 80 which can be seen in the column.



**0A\_0\_0\_0010\_000\_1** – This corresponds to 11 reads with address starting from 0 to 80 which can be seen in the column.




Figure 7-5: Multiple Reads with Same Address

## Page Wrap During Writes

**0A\_0\_2\_000F\_3F8\_0** – This corresponds to 11 writes with column address wrapped to the starting of the page after one write.




Figure 7-6: Page Wrap During Writes

## Simulating the Performance Traffic Generator

1. Map Xilinx® UNISIMS\_VER and SECUREIP library into the simulator.
2. MIG delivers memory models for DDR3 and memory models are not delivered for DDR4. Copy the memory models in the following directory for DDR4:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srscs/sim_1/imports/<Component_Name>/tb
```

3. Modify the mig\_v7\_0\_ddr4\_stimulus.txt for DDR4 and mig\_v7\_0\_ddr3\_stimulus.txt for DDR3 present in the following directory with the stimulus that wanted the percentage of bus utilization:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srscs/sim_1/imports/<Component_Name>/tb
```

4. Run the performance\_sim.do file.
5. For QuestaSim, run the following command vsim -do performance\_sim.do.
6. After the run in the tb directory, mig\_v7\_0\_ddr4\_band\_width\_cal.txt for DDR4 and mig\_v7\_0\_ddr3\_band\_width\_cal.txt for DDR3 are found with all of the output about bus percentage utilization metrics.

## SECTION III: QDR II+ SRAM

Overview

Product Specification

Core Architecture

Designing with the Core

Design Flow Steps

Example Design

Test Bench

## Overview

The Xilinx® UltraScale™ architecture includes the QDR II+ SRAM Memory Interface Solutions (MIS) core. This MIS core provides solutions for interfacing with the QDR II+ SRAM memory type.

The QDR II+ SRAM MIS core is a physical layer for interfacing Xilinx UltraScale FPGA user designs to the QDR II+ SRAM devices. QDR II+ SRAMs offer high-speed data transfers on separate read and write buses on the rising and falling edges of the clock. These memory devices are used in high-performance systems as temporary data storage, such as:

- Look-up tables in networking systems
- Packet buffers in network switches
- Cache memory in high-speed computing
- Data buffers in high-performance testers

The QDR II+ SRAM solutions core is a PHY that takes simple user commands, converts them to the QDR II+ protocol, and provides the converted commands to the memory. The design enables you to provide one read and one write request per cycle eliminating the need for a Memory Controller and the associated overhead, thereby reducing the latency through the core.

Figure 8-1 shows a high-level block diagram of the QDR II+ SRAM interface solution.

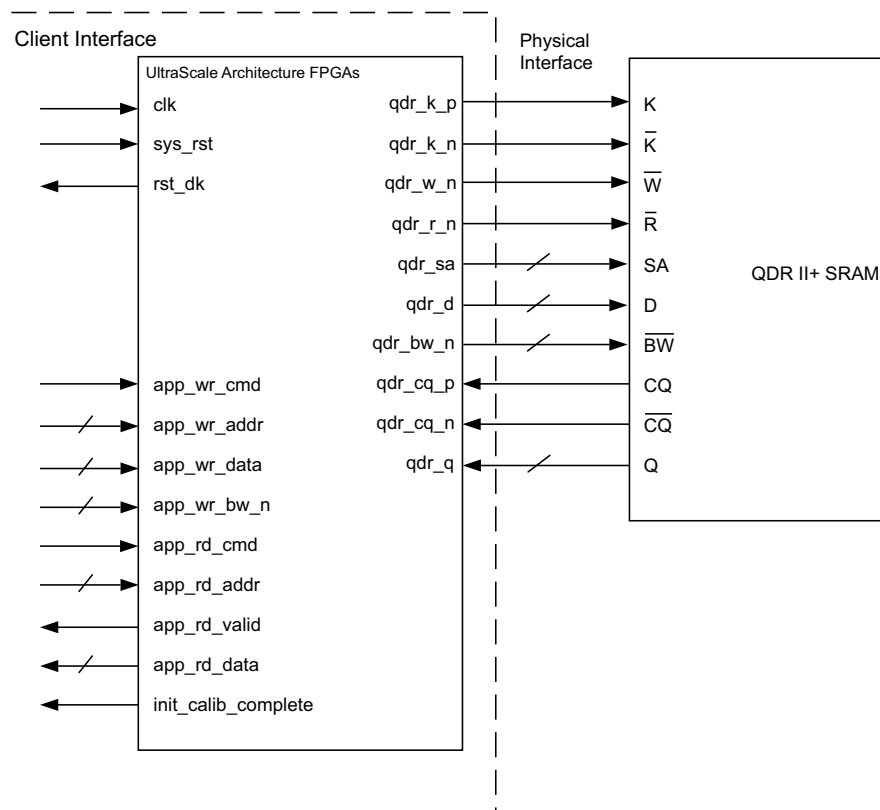


Figure 8-1: High-Level Block Diagram of QDR II+ Interface Solution

The physical layer includes the hard blocks inside the FPGA and the soft calibration logic necessary to ensure optimal timing of the hard blocks interfacing to the memory part.

These hard blocks include:

- Data serialization and transmission
- Data capture and deserialization
- High-speed clock generation and synchronization
- Coarse and fine delay elements per pin with voltage and temperature tracking

The soft blocks include:

- **Memory Initialization** – The calibration modules provide an initialization routine for the particular memory type. The delays in the initialization process can be bypassed to speed up simulation time if desired.

The QDR II+ memories do not require an elaborate initialization procedure. However, you must ensure that the `DoFF_n` signal is provided to the memory as required by the vendor. The QDR II+ SRAM interface design provided by the memory wizard drives the

`DoFF_n` signal from the FPGA. After the internal MMCM has locked, the `DoFF_n` signal is asserted High for 100  $\mu$ s without issuing any commands to the memory device.

For memory devices that require the `DoFF_n` signal to be terminated at the memory and not be driven from the FPGA, you must perform the required initialization procedure.

- **Calibration** – The calibration modules provide a complete method to set all delays in the hard blocks and soft IP to work with the memory interface. Each bit is individually trained and then combined to ensure optimal interface performance. Results of the calibration process is available through the Xilinx debug tools. After completion of calibration, the PHY layer presents raw interface to the memory part.

---

## Feature Summary

- Component support for interface widths up to 36 bits
- x18 and x36 memory device support
- 4-word and 2-word burst support
- Only HSTL\_I I/O standard support
- Cascaded data width support is available only for BL-4 designs
- Data rates up to 1,266 Mb/s for BL-4 designs
- Data rates up to 900 Mb/s for BL-2 designs
- Memory device support with 72 Mb density
- Support for 2.0 and 2.5 cycles of Read Latency
- Source code delivery in Verilog
- 2:1 memory to FPGA logic interface clock ratio
- Interface calibration and training information available through the Vivado hardware manager

---

## Licensing and Ordering Information

This Xilinx LogiCORE IP module is provided at no additional cost with the Xilinx Vivado Design Suite under the terms of the [Xilinx End User License](#). Information about this and other Xilinx LogiCORE IP modules is available at the [Xilinx Intellectual Property](#) page. For information about pricing and availability of other Xilinx LogiCORE IP modules and tools, contact your [local Xilinx sales representative](#).



## License Checkers

If the IP requires a license key, the key must be verified. The Vivado® design tools have several license checkpoints for gating licensed IP through the flow. If the license check succeeds, the IP can continue generation. Otherwise, generation halts with error. License checkpoints are enforced by the following tools:

- Vivado design tools: Vivado synthesis
- Vivado implementation
- write\_bitstream (Tcl command)



---

**IMPORTANT:** *IP license level is ignored at checkpoints. The test confirms a valid license exists. It does not check IP license level.*

---

## Product Specification

### Standards

This core complies with the QDR II+ SRAM standard defined by the QDR Consortium. For more information on UltraScale™ architecture documents, see [References](#), page 303.

### Performance

#### Maximum Frequencies

For more information on the maximum frequencies, see *Kintex UltraScale Architecture Data Sheet, DC and AC Switching Characteristics* (DS892) [\[Ref 2\]](#).

### Resource Utilization

#### Kintex UltraScale Devices

[Table 9-1](#) provides approximate resource counts on Kintex® UltraScale™ devices.

**Table 9-1: Device Utilization – Kintex UltraScale FPGAs**

Parameter Values	Device Resources						
Interface Width	FFs	LUTs	Memory LUTs	RAMB36E2/ RAMB18E2	BUFGs	PLLE3_ADV	MMCME3_ADV
36	6,741	4,456	106	16	4	3	1
18	4,271	3,083	106	16	4	2	1

Resources required for the UltraScale architecture FPGAs MIS core have been estimated for the Kintex UltraScale devices. These values were generated using Vivado® IP catalog. They are derived from post-synthesis reports, and might change during implementation.

---

## Port Descriptions

There are three port categories at the top-level of the memory interface core called the “user design.”

- The first category is the memory interface signals that directly interfaces with the memory part. These are defined by the QDR II+ SRAM specification.
- The second category is the application interface signals which is referred to as the “user interface.” This is described in the [Protocol Description, page 211](#).
- The third category includes other signals necessary for proper operation of the core. These include the clocks, reset, and status signals from the core. The clocking and reset signals are described in their respective sections.

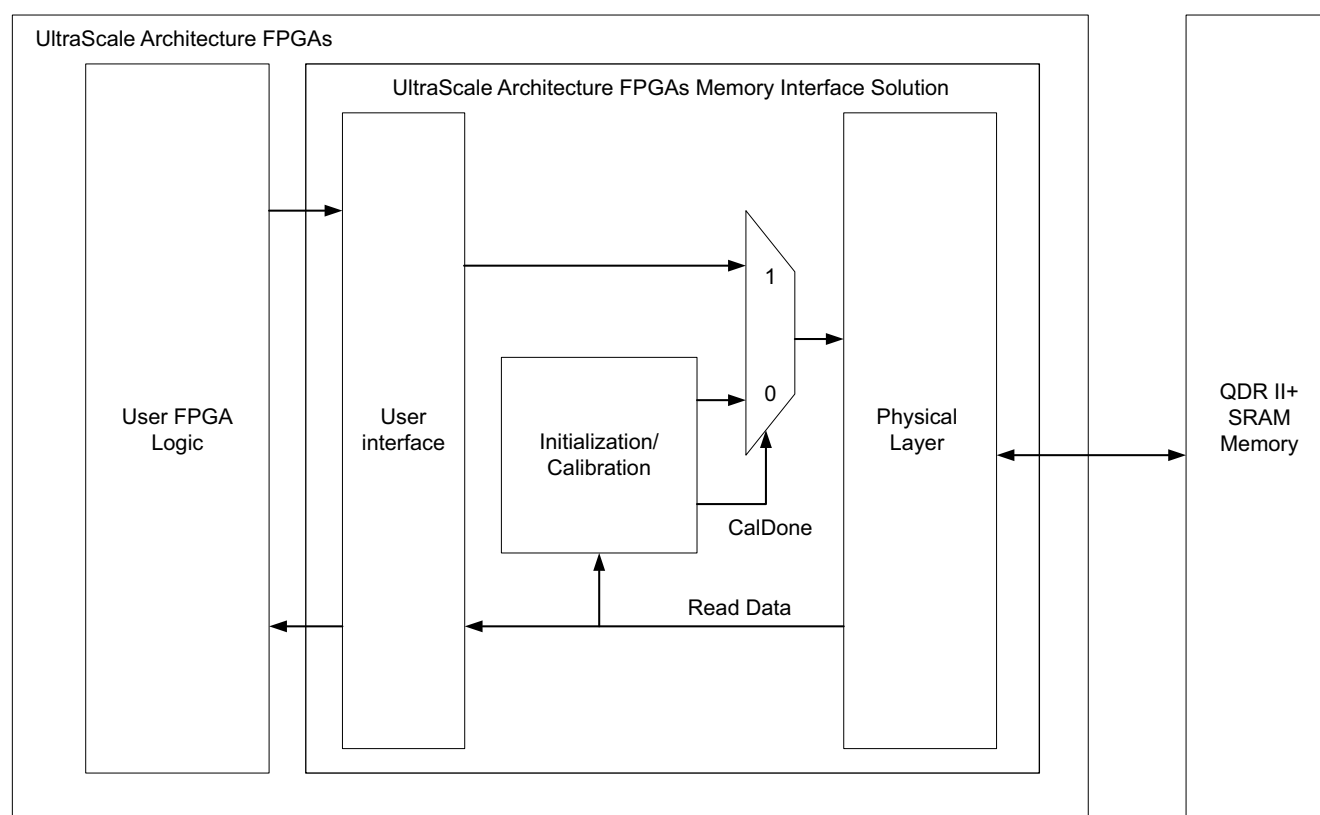
The active-High `init_calib_complete` signal indicates that the initialization and calibration are complete and that the interface is now ready to accept commands for the interface.

## Core Architecture

This chapter describes the UltraScale™ device FPGAs Memory Interface Solutions core with an overview of the modules and interfaces.

### Overview

The UltraScale architecture FPGAs Memory Interface Solutions is shown in [Figure 10-1](#).



**Figure 10-1: UltraScale Architecture FPGAs Memory Interface Solution Core**

The user interface uses a simple protocol based entirely on SDR signals to make read and write requests. For more details describing this protocol, see [User Interface in Chapter 11](#).

There is no requirement for the controller in QDR II+ SRAM protocol and thus, the Memory Controller contains only the physical interface. It takes commands from the user interface and adheres to the protocol requirements of the QDR II+ SRAM device. It is responsible to generate proper timing relationships and DDR signaling to communicate with the external memory device. For more details, see [Memory Interface in Chapter 11](#).

---

## PHY

PHY is considered the low-level physical interface to an external QDR II+ SRAM device. It contains the entire calibration logic for ensuring reliable operation of the physical interface itself. PHY generates the signal timing and sequencing required to interface to the memory device.

PHY contains the following features:

- Clock/address/control-generation logics
- Write and read datapaths
- Logic for initializing the SDRAM after power-up

In addition, PHY contains calibration logic to perform timing training of the read and write datapaths to account for system static and dynamic delays.

## Overall PHY Architecture

The UltraScale architecture PHY is composed of dedicated blocks and soft calibration logic. The dedicated blocks are structured adjacent to one another with back-to-back interconnects to minimize the clock and datapath routing necessary to build high performance physical layers.

The user interface and calibration logic communicate with this dedicated PHY in the slow frequency clock domain, which is divided by 2. A more detailed block diagram of the PHY design is shown in Figure 10-2.

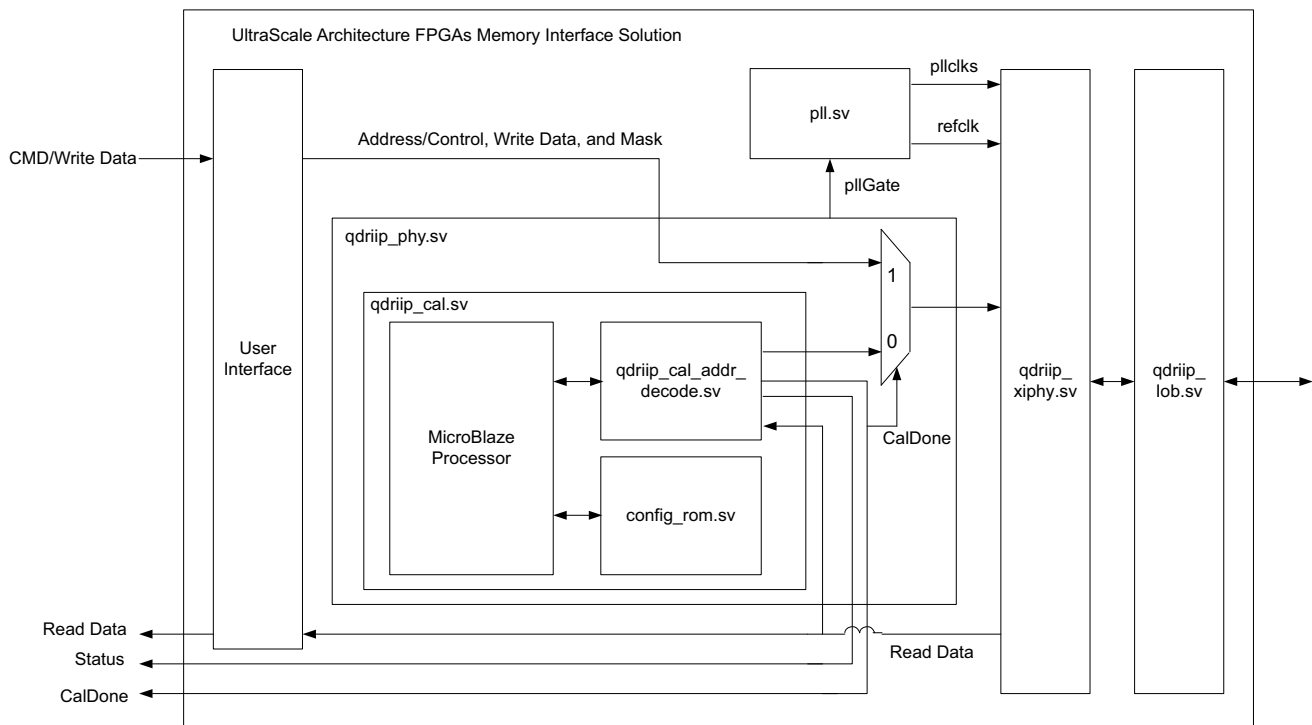


Figure 10-2: PHY Block Diagram

Table 10-1: PHY Modules

Module Name	Description
qdriip_phy.v	PHY top of QDR II+ design
qdriip_phycal.v	Contains the instances of XIPHY top and calibration top modules
qdriip_cal.v	Calibration top module
qdriip_cal_addr_decode.v	FPGA logic interface for the MicroBlaze processor
config_rom.v	Configuration storage for calibration options
debug_microblaze.v	MicroBlaze processor
qdriip_xiphy.v	Contains the XIPHY instance
qdriip_iob.v	Instantiates all byte IOB modules
qdriip_iob_byte.v	Generates the I/O buffers for all the signals in a given byte lane
qdriip_rd_bit_slip.v	Read bit slip

The PHY architecture encompasses all of the logic contained in `qdriip_xiphy.sv`. The PHY contains wrappers around dedicated hard blocks to build up the memory interface from smaller components. A byte lane contains all of the clocks, resets, and datapaths for a given subset of I/O. Multiple byte lanes are grouped together, along with dedicated clocking resources, to make up a single bank memory interface. For more information on the hard silicon physical layer architecture, see the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4].

The memory initialization and calibration are implemented in C programming on a small soft core processor. The MicroBlaze™ Controller System (MCS) is configured with an I/O Module and block RAM. The module `qdriip_cal_addr_decode.sv` module provides the interface for the processor to the rest of the system and implements helper logic. The `config_rom.sv` module stores settings that control the operation of initialization and calibration, providing run time options that can be adjusted without having to recompile the source code.

The address unit connects the MCS to the local register set and the PHY by performing address decode and control translation on the I/O module bus from spaces in the memory map and MUXing return data (`qdriip_cal_addr_decode.sv`). In addition, it provides address translation (also known as “mapping”) from a logical conceptualization of the DRAM interface to the appropriate pinout-dependent location of the delay control in the PHY address space.

Although the calibration architecture presents a simple and organized address map for manipulating the delay elements for individual data, control and command bits, there is flexibility in how those I/O pins are placed. For a given I/O placement, the path to the FPGA logic is locked to a given pin. To enable a single binary software file to work with any memory interface pinout, a translation block converts the simplified Register Interface Unit (RIU) addressing into the pinout-specific RIU address for the target design. The specific address translation is written by MIG after a pinout is selected. The code shows an example of the RTL structure that supports this.

```
Casez(io_address)// MicroBlaze I/O module address
// ... static address decoding skipped
//=====//
//=====DQ ODELAYS=====//
//=====//
//Byte0
28'h0004100: begin //dq2
    riu_addr_cal = /* MIG Generated */ 6'h0;
    riu_nibble = /* MIG Generated */ `h0;
end
// ... additional dynamic addressing follows
```

In this example, `DQ0` is pinned out on Bit[0] of nibble 0 (nibble 0 according to instantiation order). The RIU address for the ODELAY for Bit[0] is 0x0D. When `DQ0` is addressed — indicated by address 0x000\_4100), this snippet of code is active. It enables nibble 0 (decoded to one-hot downstream) and forwards the address 0x0D to the RIU address bus.

The MicroBlaze I/O interface operates at much slower frequency, which is not fast enough for implementing all the functions required in calibration. A helper circuit implemented in `qdriip_cal_adr_decode.sv` is required to obtain commands from the registers and translate at least a portion into single-cycle accuracy for submission to the PHY. In addition, it supports command repetition to enable back-to-back read transactions and read data comparison.

## Memory Initialization and Calibration Sequence

After deassertion of the system reset, PHY performs some required internal calibration steps first.

1. The built-in self-check (BISC) of the PHY is run. It is used to compensate the internal skews among the data bits and the strobe on the read path.
2. After BISC completion, the required steps for the power-on initialization of the memory part starts.
3. It requires several stages of calibration for tuning the write and read datapath skews as mentioned in [Figure 10-3](#).
4. After calibration is completed, PHY calculates internal offsets for the voltage and temperature tracking purpose by considering the taps used until the end of step 3.
5. When PHY indicates the calibration completion, the user interface command execution begins.



Figure 10-3 shows the overall flow of memory initialization and the different stages of calibration.

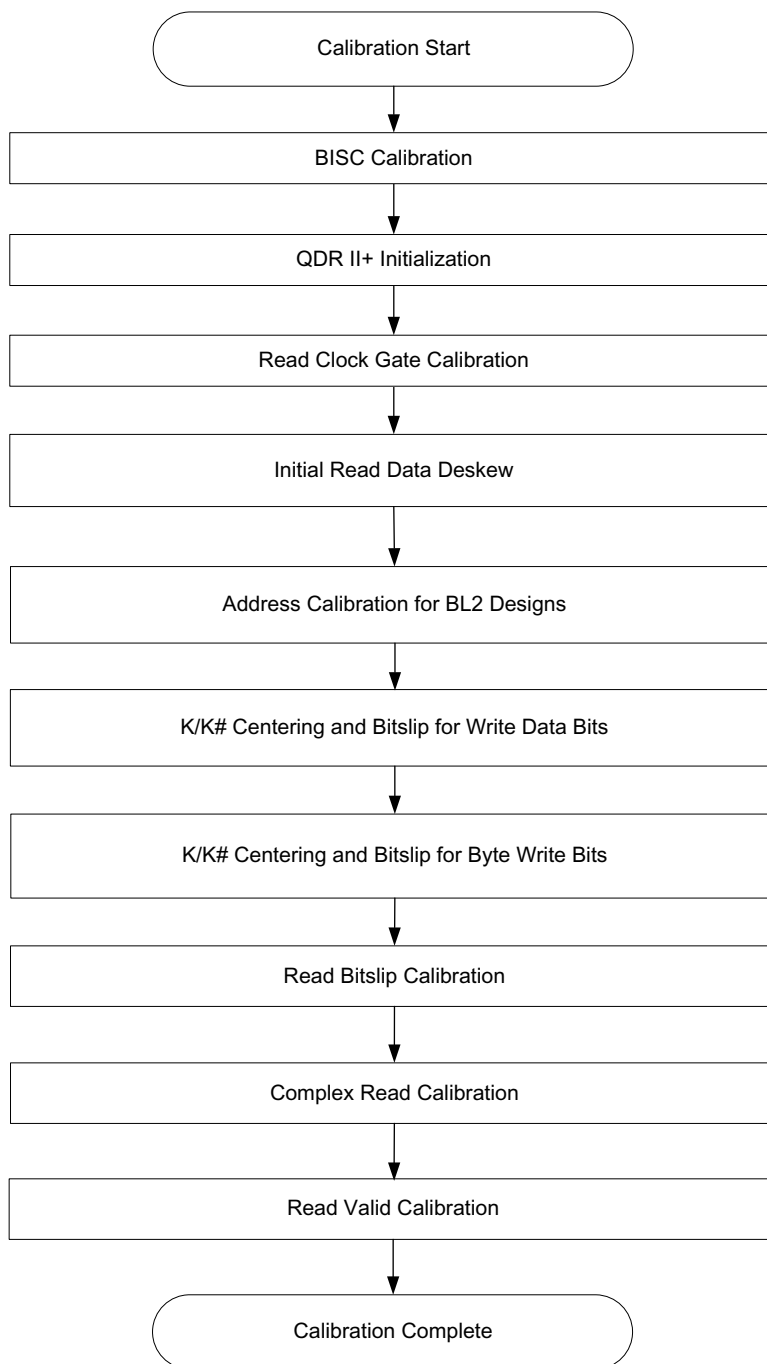


Figure 10-3: PHY Overall Initialization and Calibration Sequence

# Designing with the Core

This chapter includes guidelines and additional information to facilitate designing with the core.

---

## Clocking

The memory interface requires one MMCM, one TXPLL per I/O bank used by the memory interface and two BUFGCE\_DIVs. These clocking components are used to create the proper clock frequencies and phase shifts necessary for the proper operation of the memory interface.

There are two TXPLLs per bank. If a bank is shared by two memory interfaces, both TXPLLs in that bank are used.

**Note:** MIG generates the appropriate clocking structure and no modifications to the RTL are supported.

The MIG tool generates the appropriate clocking structure for the desired interface. This structure must not be modified. The allowed clock configuration is as follows:

- Differential reference clock source connected to GCIO
- GCIO to MMCM (located in center bank of memory interface)
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) driving FPGA logic and all TXPLLs
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) divide by two mode driving 1/2 rate FPGA logic
- Clocking pair of the interface must be in the same SLR of memory interface for the SSI technology devices

## Requirements

### ***GCIO***

- Must use a differential I/O standard
- Must be in the same I/O column as the memory interface
- Must be in the same SLR of memory interface for the SSI technology devices

### ***MMCM***

- MMCM is used to generate the FPGA logic system clock (1/2 of the memory clock)
- Must be located in the center bank of memory interface
- Must use internal feedback
- Input clock frequency divided by input divider must be  $\geq 70$  MHz ( $\text{CLKIN}_x / D \geq 70$  MHz)
- Must use integer multiply and output divide values

### ***BUFGCE\_DIVs and Clock Roots***

- One BUFGCE\_DIV is used to generate the system clock to FPGA logic and another BUFGCE\_DIV is used to divide the system clock by two.
- BUFGCE\_DIVs and clock roots must be located in center most bank of the memory interface.
  - For two bank systems, either bank can be used. MIG is always referred to the top-most selected bank in the Vivado Integrated Design Environment (IDE) as the center bank.
  - For four bank systems, either of the center banks can be chosen. MIG refers to the second bank from the top-most selected bank as the center bank.
  - Both the BUFGCE\_DIVs must be in the same bank.

### ***TXPLL***

- CLKOUTPHY from TXPLL drives XIPHY within its bank
- TXPLL must be set to use a CLKFBOUT phase shift of 90°
- TXPLL must be held in reset until the MMCM lock output goes High
- Must use internal feedback

Figure 11-1 shows an example of the clocking structure for a three bank memory interface. The GCIO drives the MMCM located at the center bank of the memory interface. MMCM drives both the BUFGCE\_DIVs located in the same bank. The BUFG CE\_DIV (which is used to generate system clock to FPGA logic) output drives the TXPLLs used in each bank of the interface.

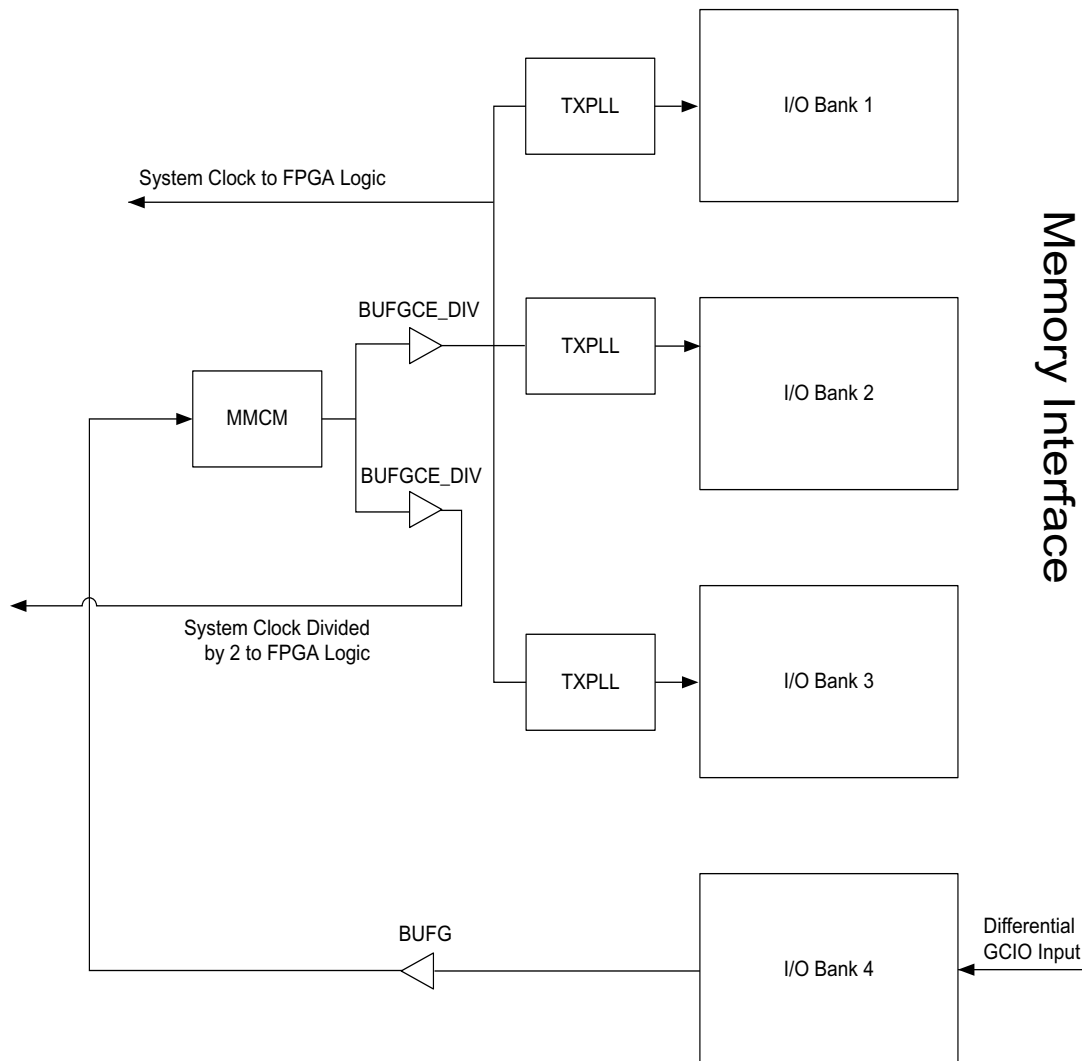


Figure 11-1: Clocking Structure for Three Bank Memory Interface

The MMCM is placed in the center bank of the memory interface.

- For two bank systems, MMCM is placed in a bank with the most number of bytes selected. If they both have the same number of bytes selected in two banks, then MMCM is placed in the top bank.
- For four bank systems, MMCM is placed in a second bank from the top.

For designs generated with System Clock configuration of **No Buffer**, MMCM must not be driven by another MMCM/PLL. Cascading clocking structures MMCM → BUFG → MMCM and PLL → BUFG → MMCM are not allowed.

If the MMCM is driven by the GCIO pin of the other bank, then the CLOCK\_DEDICATED\_ROUTE constraint with value "BACKBONE" must be set on the net that is driving MMCM or on the MMCM input. Setting up the CLOCK\_DEDICATED\_ROUTE constraint on the net is preferred. But when the same net is driving two MMCMs, the CLOCK\_DEDICATED\_ROUTE constraint must be managed by considering which MMCM needs the BACKBONE route.

In such cases, the CLOCK\_DEDICATED\_ROUTE constraint can be set on the MMCM input. To use the "BACKBONE" route, any clock buffer that exists in the same CMT tile as the GCIO must exist between the GCIO and MMCM input. The clock buffers that exist in the I/O CMT are BUFG, BUFGCE, BUFGCTRL, and BUFGCE\_DIV. So MIG instantiates BUFG between the GCIO and MMCM when the GCIO pins and MMCM are not in the same bank (see [Figure 11-1](#)).

If the GCIO pin and MMCM are allocated in different banks, MIG generates CLOCK\_DEDICATED\_ROUTE constraints with value as "BACKBONE." If the GCIO pin and MMCM are allocated in the same bank, there is no need to set any constraints on the MMCM input.

Similarly when designs are generated with System Clock Configuration as a **No Buffer** option, you must take care of the "BACKBONE" constraint and the BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV between GCIO and MMCM if GCIO pin and MMCM are allocated in different banks. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations. For more information on clocking, see the *UltraScale Architecture Clocking Resources User Guide* (UG572) [[Ref 3](#)].

**Note:** If two different GCIO pins are used for two MIG IP cores in the same bank, center bank of the memory interface is different for each IP. MIG generates MMCM LOC and CLOCK\_DEDICATED\_ROUTE constraints accordingly.

## Sharing of Input Clock Source (sys\_clk\_p)

If the same GCIO pin must be used for two IP cores, generate the two IP cores with System Clock Configuration option as **No Buffer**. Perform the following changes in the wrapper file in which both IPs are instantiated:

1. MIG generates a single-ended input for system clock pins, such as `sys_clk_i`. Connect the differential buffer output to the single-ended system clock inputs (`sys_clk_i`) of both the IP cores.
2. System clock pins must be allocated within the same I/O column of the memory interface pins allocated. Add the pin LOC constraints for system clock pins and clock constraints in your top-level XDC.

3. You must add a "BACKBONE" constraint on the net that is driving the MMCM or on the MMCM input if GCIO pin and MMCM are not allocated in the same bank. Apart from this, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM to use the "BACKBONE" route.

**Note:**

- The Ultrascale architecture includes an independent XIPHY power supply and TXPLL for each XIPHY. This results in clean, low jitter clocks for the memory system.
- Skew spanning across multiple BUFGs is not a concern because single point of contact exists between BUFG → TXPLL and the same BUFG → System Clock Logic.
- System input clock cannot span I/O columns because the longer the clock lines span, the more jitter is picked up.

---

## Resets

An asynchronous reset (`sys_rst`) input is provided. This is an active-High reset and the `sys_rst` must assert for a minimum pulse width of 5 ns. The `sys_rst` can be an internal or external pin.

---

## PCB Guidelines for QDR II+ SRAM

Strict adherence to all documented QDR II+ SRAM PCB guidelines is required for successful operation. For more information on PCB guidelines, see the *UltraScale Architecture PCB Design and Pin Planning User Guide* (UG583) [Ref 5].

---

## Pin and Bank Rules

### QDR II+ Pin Rules

This section describes the pin out rules for QDR II+ SRAM interface.

- Both HR and HP Banks are supported.
  - All signal groups that are write data, read data, address/control, and system clock interfaces must be selected in a single column.
  - All banks used must be adjacent. No skip banks allowed.
1. Write Data (D) and Byte Write (BW) Pins Allocation:

- a. The entire write data bus must be placed in a single bank regardless of the number of memory components.
  - b. Only one write data byte is allowed per byte lane.
  - c. All byte lanes that are used for the write data of a single component must be adjacent, no skip byte lanes are allowed.
  - d. One of the write data bytes of a memory component should be allocated in the center byte lanes (byte lanes 1 and 2).
  - e. Each byte write pin (BW) must be allocated in the corresponding write data byte lane.
2. Memory Clock (K/K#) Allocation:
- a. Memory Clock pair must be allocated in one of the byte lanes that are used for the write data of the corresponding memory component.
  - b. Memory clock should come from one of the center byte lanes (byte lanes 1 and 2).
  - c. K/K# can be allocated to any PN pair.
3. Read Data (Q) Allocation:
- a. The entire read data bus must be placed in a single bank irrespective of the number of memory components.
  - b. All byte lanes that are used for the read data of a single component must be adjacent, no skip byte lanes are allowed.
  - c. One of the read data bytes of a memory component should be allocated in the center byte lanes (byte lanes 1 and 2).
  - d. If a byte lane is used for read data, Bit[0] and Bit[6] must be used. Read clock (CQ or CQ#) gets the first priority and data (Q) is the next.
  - e. Read data buses of two components should not share a byte lane.
4. Read Clock (CQ/CQ#) Allocation:
- a. Read Clock pair must be allocated in one of the byte lanes that are used for the read data of the corresponding memory component.
  - b. CQ/CQ# pair must be allocated in a single byte lane.
  - c. CQ/CQ# must be allocated only in the center byte lanes (byte lanes 1 and 2) because other byte lanes cannot forward the clock out for read data capture.
  - d. CQ and CQ# must be allocated to either pin 0 or pin 6 of a byte lane. For example, if CQ is allocated to pin 0, CQ# should be allocated to pin 6 and vice versa.
5. For x36 and x18 component designs:
- a. All Read Data pins of a single component must not span more than three consecutive byte lanes and CQ/CQ# must always be allocated in center byte lane.

6. Address/Control (A/C) Pins Allocation:
  - a. All address/control (A/C) bits must be allocated in a single bank.
  - b. All A/C byte lanes should be contiguous and no skip byte lanes is allowed.
  - c. The address/control bank should be the same or adjacent to that of the write data bank.
  - d. There should not be any empty byte lane or read byte lane between A/C and write data byte lanes. This rule applies when A/C and write data share the same bank or allocated in adjacent banks.
  - e. Address/control pins should not share a byte lane with the write data as well as read data.
  - f. System clock pins (`sys_clk_p/sys_clk_n`) must be placed on any GCIO pin pair in the same column as that of the memory interface. For more information, see [Clocking, page 202](#).
7. All I/O banks used by the memory interface must be in the same SLR of the column for the SSI technology devices.
8. One `vrp` pin per bank is used and a DCI is required for the interfaces. A `vrp` pin is required in I/O banks containing inputs as well as output only banks. It is required in output only banks because address/control signals use `HSTL_I_DCI` to enable usage of controlled output impedance. A DCI cascade is not permitted. All rules for the DCI in the *UltraScale™ Device FPGAs SelectIO™ Resources User Guide* (UG571) [\[Ref 4\]](#) must be followed.

## QDR II+ Pinout Examples



**IMPORTANT:** Due to the calibration stage, there is no need for `set_input_delay/`  
`set_output_delay` on the MIG. Ignore the unconstrained inputs and outputs for MIG and the signals which are calibrated.

[Table 11-1](#) shows an example of an 18-bit QDR II+ SRAM interface contained within two banks.

**Table 11-1: 18-Bit QDR II+ Interface Contained in Two Banks**

Bank	Signal Name	Byte Group	I/O Type
1	–	T1U_12	–
1	<code>sys_clk_p</code>	T1U_11	N
1	<code>sys_clk_n</code>	T1U_10	P
1	–	T1U_9	N
1	<code>q17</code>	T1U_8	P
1	<code>q16</code>	T1U_7	N



Table 11-1: 18-Bit QDR II+ Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
1	cq_p	T1U_6	P
1	q15	T1L_5	N
1	q14	T1L_4	P
1	q13	T1L_3	N
1	q12	T1L_2	P
1	q11	T1L_1	N
1	cq_n	T1L_0	P
1	vrp	T0U_12	–
1	–	T0U_11	N
1	q10	T0U_10	P
1	q9	T0U_9	N
1	q8	T0U_8	P
1	q7	T0U_7	N
1	q6	T0U_6	P
1	q5	T0L_5	N
1	q4	T0L_4	P
1	q3	T0L_3	N
1	q2	T0L_2	P
1	q1	T0L_1	N
1	q0	T0L_0	P
0	–	T3U_12	–
0	–	T3U_11	N
0	–	T3U_10	P
0	d17	T3U_9	N
0	d16	T3U_8	P
0	d15	T3U_7	N
0	d14	T3U_6	P
0	d13	T3L_5	N
0	d12	T3L_4	P
0	d11	T3L_3	N
0	d10	T3L_2	P
0	bwsn1	T3L_1	N
0	d9	T3L_0	P

Table 11-1: 18-Bit QDR II+ Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
0	–	T2U_12	–
0	d8	T2U_11	N
0	d7	T2U_10	P
0	d6	T2U_9	N
0	d5	T2U_8	P
0	k_n	T2U_7	N
0	k_p	T2U_6	P
0	d4	T2L_5	N
0	d3	T2L_4	P
0	d2	T2L_3	N
0	d1	T2L_2	P
0	bwsn0	T2L_1	N
0	d0	T2L_0	P
0	doff	T1U_12	–
0	a21	T1U_11	N
0	a20	T1U_10	P
0	a19	T1U_9	N
0	a18	T1U_8	P
0	a17	T1U_7	N
0	a16	T1U_6	P
0	a15	T1L_5	N
0	a14	T1L_4	P
0	a13	T1L_3	N
0	a12	T1L_2	P
0	rpsn	T1L_1	N
0	a11	T1L_0	P
0	vrp	T0U_12	–
0	a10	T0U_11	N
0	a9	T0U_10	P
0	a8	T0U_9	N
0	a7	T0U_8	P
0	a6	T0U_7	N

Table 11-1: 18-Bit QDR II+ Interface Contained in Two Banks (Cont'd)

Bank	Signal Name	Byte Group	I/O Type
0	a5	T0U_6	P
0	a4	T0L_5	N
0	a3	T0L_4	P
0	a2	T0L_3	N
0	a1	T0L_2	P
0	wpsn	T0L_1	N
0	a0	T0L_0	P

## Protocol Description

This core has the following interfaces:

- [User Interface](#)
- [Memory Interface](#)

### User Interface

The user interface connects an FPGA user design to the QDR II+ SRAM solutions core to simplify interactions between you and the external memory device. The user interface provides a set of signals used to issue a read or write command to the memory device. These signals are summarized in [Table 11-2](#).

Table 11-2: User Interface Signals

Signal	Direction	Description
app_rd_addr0[ADDR_WIDTH – 1:0]	Input	Read Address. This bus provides the address to use for a read request. It is valid when app_rd_cmd0 is asserted.
app_rd_cmd0	Input	Read Command. This signal is used to issue a read request and indicates that the address on port0 is valid.
app_rd_data0[DBITS × BURST_LEN – 1:0]	Output	Read Data. This bus carries the data read back from the read command issued on app_rd_cmd0
app_rd_valid0	Output	Read Valid. This signal indicates that data read back from memory is now available on app_rd_data0 and should be sampled.
app_wr_addr0[ADDR_WIDTH – 1:0]	Input	Write Address. This bus provides the address for a write request. It is valid when app_wr_cmd0 is asserted.

Table 11-2: User Interface Signals (Cont'd)

Signal	Direction	Description
app_wr_bw_n0[(DBITS/9) × BURST_LEN – 1:0]	Input	Byte Writes. This bus provides the byte writes for a write request and indicates which bytes need to be written into the SRAM. It is valid when app_wr_cmd0 is asserted and is active-Low.
app_wr_cmd0	Input	Write Command. This signal is used to issue a write request and indicates that the corresponding sideband signals on write port0 are valid.
app_wr_data0[DBITS × BURST_LEN – 1:0]	Input	Write Data. This bus provides the data to use for a write request. It is valid when app_wr_cmd0 is asserted.
app_rd_addr1[ADDR_WIDTH – 1:0] <sup>(1)</sup>	Input	Read Address. This bus provides the address to use for a read request. It is valid when app_rd_cmd1 is asserted.
app_rd_cmd1 <sup>(1)</sup>	Input	Read Command. This signal is used to issue a read request and indicates that the address on port1 is valid.
app_rd_data1[DBITS × BURST_LEN – 1:0] <sup>(1)</sup>	Output	Read Data. This bus carries the data read back from the read command issued on app_rd_cmd1.
app_rd_valid1 <sup>(1)</sup>	Output	Read Valid. This signal indicates that data read back from memory is now available on app_rd_data1 and should be sampled.
app_wr_addr1[ADDR_WIDTH – 1:0] <sup>(1)</sup>	Input	Write Address. This bus provides the address for a write request. It is valid when app_wr_cmd1 is asserted.
app_wr_bw_n1[(DBITS/9) × BURST_LEN – 1:0] <sup>(1)</sup>	Input	Byte Writes. This bus provides the byte writes for a write request and indicates which bytes need to be written into the SRAM. It is valid when app_wr_cmd1 is asserted and is active-Low.
app_wr_cmd1 <sup>(1)</sup>	Input	Write Command. This signal is used to issue a write request and indicates that the corresponding sideband signals on write port1 are valid.
app_wr_data1[DBITS × BURST_LEN – 1:0] <sup>(1)</sup>	Input	Write Data. This bus provides the data to use for a write request. It is valid when app_wr_cmd1 is asserted.
clk	Output	User Interface clock.
rst_clk	Output	Reset signal synchronized by the User Interface clock.
Init_calib_complete	Output	Calibration Done. This signal indicates to the user design that read calibration is complete and the user can now initiate read and write requests from the client interface.
sys_rst	Input	Asynchronous system reset input.
sys_clk_p/n	Input	System clock to the Memory Controller.

Table 11-2: User Interface Signals (Cont'd)

Signal	Direction	Description
dbg_clk	Output	Debug Clock. Do not connect any signals to dbg_clk and keep the port open during instantiation.
dbg_bus	Output	Reserved. Do not connect any signals to dbg_bus and keep the port open during instantiation.

**Notes:**

1. These ports are available and valid only in BL2 configuration. For BL4 configuration, these ports are not available or if available, no need to be driven.

## Interfacing with the Core through the User Interface

Figure 11-2 shows the user interface protocol.

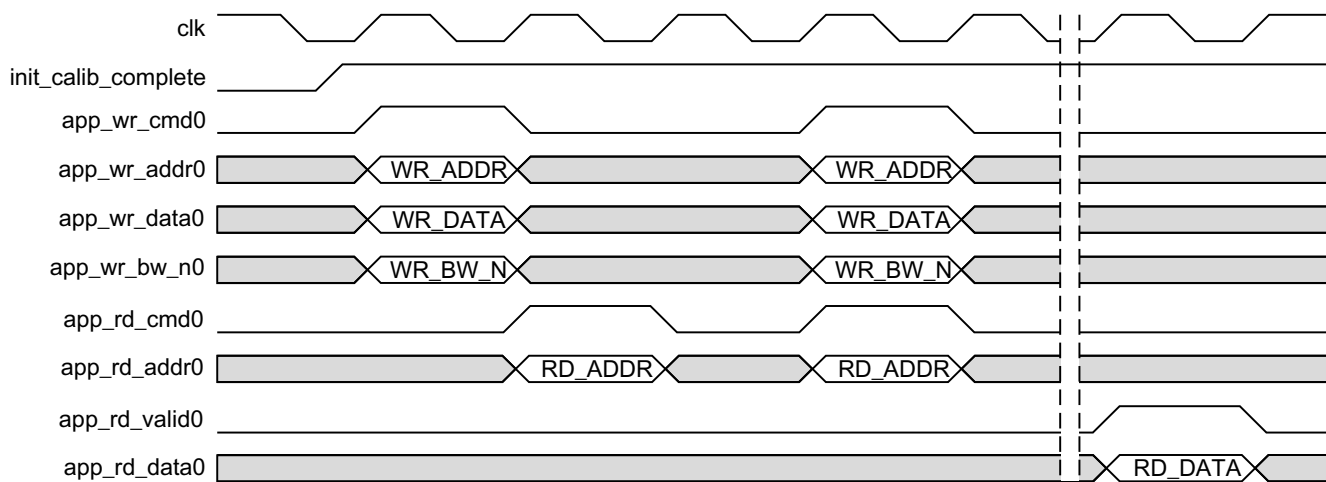


Figure 11-2: User Interface Write/Read Timing Diagram

Before any requests can be made, the `init_calib_complete` signal must be asserted High, as shown in Figure 11-2, no read or write requests can take place, and the assertion of `app_wr_cmd0` or `app_rd_cmd0` on the client interface is ignored. A write request is issued by asserting `app_wr_cmd0` as a single cycle pulse. At this time, the `app_wr_addr0`, `app_wr_data0`, and `app_wr_bw_n0` signals must be valid.

On the following cycle, a read request is issued by asserting `app_rd_cmd0` for a single cycle pulse. At this time, `app_rd_addr0` must be valid. After one cycle of idle time, a read and write request are both asserted on the same clock cycle. In this case, the read to the memory occurs first, followed by the write. The write and read commands can be applied in any order at the user interface, two examples are shown in the Figure 11-2.

Also, Figure 11-2 shows data returning from the memory device to the user design. The `app_rd_valid0` signal is asserted, indicating that `app_rd_data0` is now valid. This should be sampled on the same cycle that `app_rd_valid0` is asserted because the core does not buffer returning data.

In case of BL2, the same protocol should be followed on two independent ports: port-0 and port-1. [Figure 11-2](#) shows the user interface signals on port-0 only.

## Memory Interface

Memory interface is a connection from the FPGA memory solution to an external QDR II+ SRAM device. The I/O signals for this interface are defined in [Table 11-3](#). These signals can be directly connected to the corresponding signals on the memory device.

**Table 11-3: Memory Interface Signals**

Signal	Direction	Description
qdriip_cq_n	Input	QDR CQ#. This is the echo clock returned from the memory derived from qdr_k_n.
qdriip_cq_p	Input	QDR CQ. This is the echo clock returned from the memory derived from qdr_k_p.
qdriip_d	Output	QDR Data. This is the write data from the PHY to the QDR II+ memory device.
qdriip_dll_off_n	Output	QDR DLL Off. This signal turns off the DLL in the memory device.
qdriip_bw_n	Output	QDR Byte Write. This is the byte write signal from the PHY to the QDR II+ SRAM device.
qdriip_k_n	Output	QDR Clock K#. This is the inverted input clock to the memory device.
qdriip_k_p	Output	QDR Clock K. This is the input clock to the memory device.
qdriip_q	Input	QDR Data Q. This is the data returned from reads to memory.
qdriip_qvld	Input	QDR Q Valid. This signal indicates that the data on qdriip_q is valid. It is only present in QDR II+ SRAM devices.
qdriip_sa	Output	QDR Address. This is the address supplied for memory operations.
qdriip_w_n	Output	QDR Write. This is the write command to memory.
qdriip_r_n	Output	QDR Read. This is the read command to memory.

Figure 11-3 shows the timing diagram for the sample write and read operations at the memory interface of a BL4 QDR II+ SRAM device and Figure 11-4 is that of a BL2 device.




Figure 11-3: Interfacing with a Four-Word Burst Length Memory Device

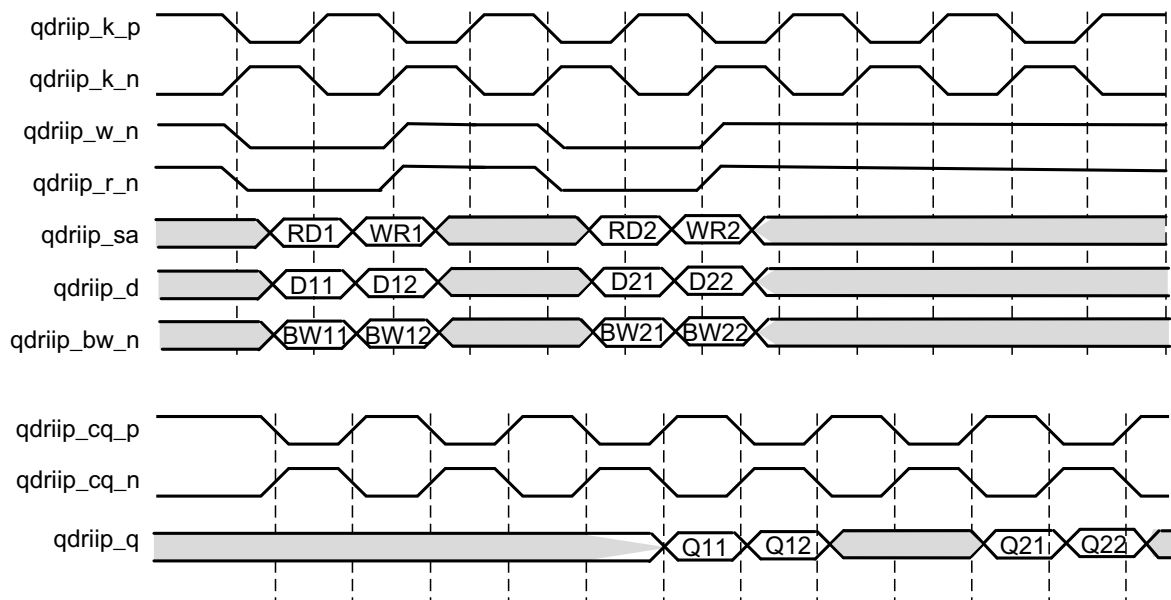


Figure 11-4: Interfacing with a Two-Word Burst Length Memory Device

# Design Flow Steps

This chapter describes customizing and generating the core, constraining the core, and the simulation, synthesis and implementation steps that are specific to this IP core. More detailed information about the standard Vivado® design flows and the Vivado IP integrator can be found in the following Vivado Design Suite user guides:

- *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#)
- *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#)
- *Vivado Design Suite User Guide: Getting Started* (UG910) [\[Ref 9\]](#)
- *Vivado Design Suite User Guide: Logic Simulation* (UG900) [\[Ref 10\]](#)

---

## Customizing and Generating the Core



**CAUTION!** *The Windows operating system has a 260-character limit for path lengths, which can affect the Vivado tools. To avoid this issue, use the shortest possible names and directory locations when creating projects, defining IP or managed IP projects, and creating block designs.*

---

This section includes information about using Xilinx® tools to customize and generate the core in the Vivado Design Suite.

If you are customizing and generating the core in the IP integrator, see the *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#) for detailed information. IP integrator might auto-compute certain configuration values when validating or generating the design. To check whether the values change, see the description of the parameter in this chapter. To view the parameter value, run the `validate_bd_design` command in the Tcl Console.

You can customize the IP for use in your design by specifying values for the various parameters associated with the IP core using the following steps:

1. Select the IP from the Vivado IP catalog.
2. Double-click the selected IP or select the Customize IP command from the toolbar or right-click menu.



For more information about generating the core in Vivado, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8] and the *Vivado Design Suite User Guide: Getting Started* (UG910) [Ref 9].

## User Parameters

Table 12-1 shows the relationship between the GUI fields in the Vivado IDE and the User Parameters (which can be viewed in the Tcl Console).

Table 12-1: Vivado IDE Parameter to User Parameter Relationship

Vivado IDE Parameter/Value <sup>(1)</sup>	User Parameter/Value <sup>(1)</sup>	Default Value
System Clock Configuration	System_Clock	Differential
Internal V <sub>REF</sub>	Internal_Vref	TRUE
DCI Cascade	DCI_Cascade	FALSE
Debug Signal for Controller	Debug_Signal	Disable
Clock 1 (MHz)	ADDN_UI_CLKOUT1_FREQ_HZ	None
Clock 2 (MHz)	ADDN_UI_CLKOUT2_FREQ_HZ	None
Clock 3 (MHz)	ADDN_UI_CLKOUT3_FREQ_HZ	None
Clock 4 (MHz)	ADDN_UI_CLKOUT4_FREQ_HZ	None
I/O Power Reduction	IOPowerReduction	OFF
Enable System Ports	Enable_SysPorts	TRUE
I/O Power Reduction	IO_Power_Reduction	FALSE
Default Bank Selections	Default_Bank_Selections	FALSE
Reference Clock	Reference_Clock	FALSE
Enable System Ports	Enable_SysPorts	TRUE
Clock Period (ps)	C0.QDRIIP_TimePeriod	1,819
Input Clock Period (ps)	C0.QDRIIP_InputClockPeriod	13,637
Configuration	C0.QDRIIP_MemoryType	Components
Memory Part	C0.QDRIIP_MemoryPart	CY7C2565XV18-633BZXC
Data Width	C0.QDRIIP_DataWidth	36
Burst Length	C0.QDRIIP_BurstLen	4
Memory Name	C0.QDRIIP_MemoryName	Main Memory

### Notes:

- Parameter values are listed in the table where the Vivado IDE parameter value differs from the user parameter value. Such values are shown in this table as indented below the associated parameter.

## Output Generation

For details, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

---

## MIG I/O Planning

For details on I/O planning, see [MIG I/O Planning, page 164](#).

---

## Constraining the Core

This section contains information about constraining the core in the Vivado Design Suite.

### Required Constraints

The MIG Vivado IDE generates the required constraints. A location constraint and an I/O standard constraint are added for each external pin in the design. The location is chosen by the Vivado IDE according to the banks and byte lanes chosen for the design.

The I/O standard is chosen by the memory type selection and options in the Vivado IDE and by the pin type. A sample for `qdriip_d[0]` is shown here.

```
set_property LOC AP25 [get_ports {c0_qdriip_d[0]}]
set_property IOSTANDARD HSTL_I [get_ports {c0_qdriip_d[0]}]
```

The system clock must have the period set properly:

```
create_clock -name c0_sys_clk -period 1.818 [get_ports c0_sys_clk_p]
```

### Device, Package, and Speed Grade Selections

This section is not applicable for this IP core.

### Clock Frequencies

This section is not applicable for this IP core.

### Clock Management

For more information on clocking, see [Clocking, page 202](#).

### Clock Placement

This section is not applicable for this IP core.

### Banking

This section is not applicable for this IP core.

## Transceiver Placement

This section is not applicable for this IP core.

## I/O Standard and Placement

The MIG tool generates the appropriate I/O standards and placement based on the selections made in the Vivado IDE for the interface type and options.



---

**IMPORTANT:** *The `set_input_delay` and `set_output_delay` constraints are not needed on the external memory interface pins in this design due to the calibration process that automatically runs at start-up. Warnings seen during implementation for the pins can be ignored.*

---

---

## Simulation

This section contains information about simulating the MIG generated IP. Vivado simulator, QuestaSim, IES, and VCS simulation tools are used for verification of the MIG IP at each software release. Vivado simulator is not supported yet. For more information on simulation, see [Chapter 13, Example Design](#) and [Chapter 14, Test Bench](#).

---


## Synthesis and Implementation

For details about synthesis and implementation, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#).

## Example Design

This chapter contains information about the example design provided in the Vivado® Design Suite.

Vivado supports Open IP Example Design flow. To create the example design using this flow, right-click the IP in the **Source Window**, as shown in [Figure 13-1](#) and select **Open IP Example Design**.



**Figure 13-1: Open IP Example Design**

This option creates a new Vivado project. Upon selecting the menu, a dialog box to enter the directory information for the new design project opens.

Select a directory, or use the defaults, and click **OK**. This launches a new Vivado with all of the example design files and a copy of the IP.

Figure 13-1 shows the example design with the PHY only option selected (controller module does not get generated).




Figure 13-2: Open IP Example Design with PHY Only Option Selected

Figure 13-3 shows the example design with the PHY only option not selected (controller module is generated).




Figure 13-3: Open IP Example Design with PHY Only Option Not Selected

## Simulating the Example Design (Designs with Standard User Interface)

The example design provides a synthesizable test bench to generate a fixed simple data pattern to the Memory Controller. This test bench consists of an IP wrapper and an `example_tb` that generates 10 writes and 10 reads. MIG does not deliver the QDR II+ memory models. The memory model required for the simulation must be downloaded from the memory vendor website.

The example design can be simulated using one of the methods in the following sections.

## Project-Based Simulation

This method can be used to simulate the example design using the Vivado Integrated Design Environment (IDE). MIG does not deliver the QDR II+ memory models. The memory model required for the simulation must be downloaded from the memory vendor website. The memory model file must be added in the example design using **Add simulation sources** option to run simulation.

The Vivado simulator, QuestaSim, IES, and VCS tools are used for QDR II+ IP verification at each software release. The Vivado simulation tool is used for QDR II+ IP verification from 2015.1 Vivado software release. The following subsections describe steps to run a project-based simulation using each supported simulator tool.

### Project-Based Simulation Flow Using Vivado Simulator

1. In the **Open IP Example Design** Vivado project, under **Add sources** option, select the **Add or create simulation sources** option, and click **Next** as shown in [Figure 13-4](#).



Figure 13-4: Add Source Option in Vivado

2. Add the memory model in the **Add or create simulation sources** page and click **Finish** as shown in Figure 13-9.




Figure 13-5: Add or Create Simulation Sources in Vivado

3. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
4. Select **Target simulator** as **Vivado Simulator**.
  - a. Under the **Simulation** tab, set the `xsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 13-10. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
5. Apply the settings and select **OK**.






Figure 13-6: Simulation with Vivado Simulator

6. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 13-7.




Figure 13-7: Run Behavioral Simulation

7. Vivado invokes Vivado simulator and simulations are run in the Vivado simulator tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using QuestaSim

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Add sources** option, select the **Add or create simulation sources** option, and click **Next** as shown in Figure 13-8.



Figure 13-8: Add Source Option in Vivado

2. Add the memory model in the **Add or create simulation sources** page and click **Finish** as shown in Figure 13-9.




Figure 13-9: Add or Create Simulation Sources in Vivado

3. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
4. Select **Target simulator** as **QuestaSim/ModelSim Simulator**.
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `modelsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 13-10. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
5. Apply the settings and select **OK**.




Figure 13-10: Simulation with QuestaSim

6. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 13-11.




Figure 13-11: Run Behavioral Simulation

7. Vivado invokes QuestaSim and simulations are run in the QuestaSim tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using IES

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Add sources option**, select the **Add or create simulation sources** option and click **Next** as shown in Figure 13-8.
2. Add the memory model in the **Add or create simulation sources** page and click **Finish** as shown in Figure 13-9.
3. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
4. Select **Target simulator** as Incisive Enterprise Simulator (IES).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `ies.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 13-12. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
5. Apply the settings and select **OK**.




Figure 13-12: Simulation with IES Simulator

6. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 13-11.
7. Vivado invokes IES and simulations are run in the IES tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using VCS

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Add sources option**, select the **Add or create simulation sources** option and click **Next** as shown in [Figure 13-8](#).
2. Add the memory model in the **Add or create simulation sources** page and click **Finish** as shown in [Figure 13-9](#).
3. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
4. Select **Target simulator** as Verilog Compiler Simulator (VCS).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `vcs.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in [Figure 13-13](#). The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
5. Apply the settings and select **OK**.






Figure 13-13: Simulation with VCS Simulator

6. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 13-11.
7. Vivado invokes VCS and simulations are run in the VCS tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation (UG900)* [Ref 10].

## Non-Project-Based Simulation



**IMPORTANT:** Xilinx® UNISIMS\_VER and SECUREIP library must be mapped into the simulator.

1. To run the simulation, go to this directory:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srscs/sim_1/imports/tb
```

If the MIG design is generated with the Component Name entered in the Vivado IDE as mig\_0, the simulation directory path is the following:

```
<project_dir>/example_project/mig_0_example/mig_0_example.srscs/  
sim_1/imports/tb
```

2. MIG does not deliver the QDR II+ memory models. The memory model required for the simulation must be downloaded from the memory vendor website. Copy the memory model into the tb folder. See the readme.txt file located in the folder for running simulations.
3. The QuestaSim, IES, and VCS simulation tools are used for verification of MIG IP at each software release.
4. Script files to run simulations with QuestaSim, IES, and VCS are generated in MIG generated output. Other simulation tools can be used for MIG IP simulation but are not specifically verified by Xilinx.

## Simulation Speed

MIG provides a Vivado IDE option to reduce the simulation speed by selecting behavioral XIPHY model instead of UNISIM XIPHY model. Behavioral XIPHY model simulation is a default option for MIG designs. To select the simulation mode, click the **Advanced** tab and find the **Simulation Options** as shown in [Figure 13-14](#).

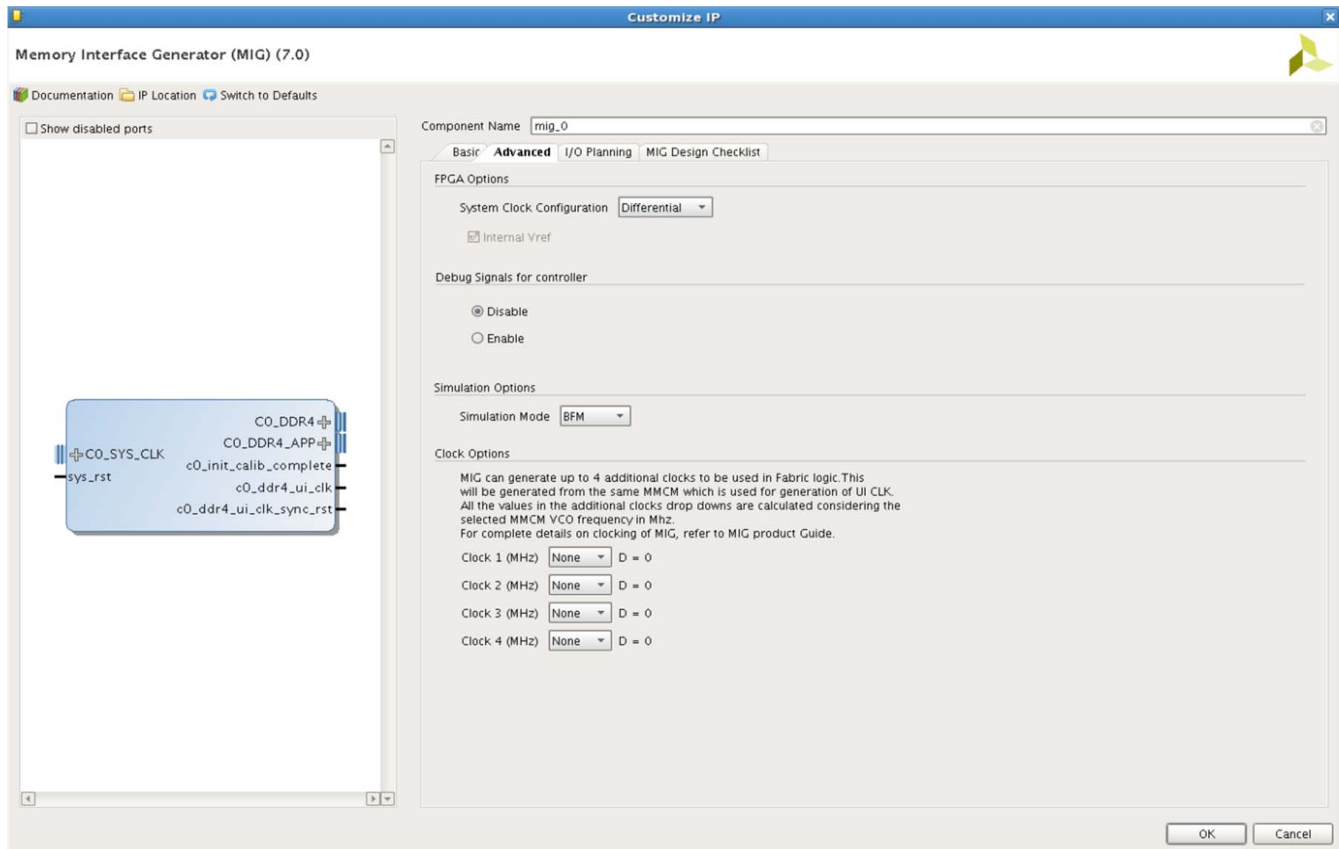


Figure 13-14: Advanced Tab – Simulation Options

The SIM\_MODE parameter in the RTL is given a different value based on the Vivado IDE selection.

- **SIM\_MODE = BFM** – If fast mode is selected in the Vivado IDE, the RTL parameter reflects this value for the SIM\_MODE parameter. This is the default option.
- **SIM\_MODE = FULL** – If FULL mode is selected in the Vivado IDE, XIPHY UNISIMs are selected and the parameter value in the RTL is FULL.



**IMPORTANT:** QDR II+ memory models from Cypress® Semiconductor need to be modified with the following two timing parameter values to run the simulations successfully:

```
`define tcqd #0
`define tcqdoh #0.15
```

## Synplify Black Box Testing

Using the Synopsys® Synplify Pro® black box testing for `example_design`, follow these steps to run black box synthesis with `synplify_pro` and implementation with Vivado.

1. Generate the UltraScale™ architecture MIG IP core with OOC flow to generate the `.dcp` file for implementation. The **Target Language** for the project can be selected as **verilog** or **VHDL**.
2. Create the example design for the MIG IP core using the information provided in the example design section and close the Vivado project.
3. Invoke the `synplify_pro` software which supports UltraScale FPGA and select the same UltraScale FPGA part selected at the time of generating the IP core.
4. Add the following files into `synplify_pro` project based on the **Target Language** selected at the time of invoking Vivado:

a. For Verilog:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/*stub.v  
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>/rtl/ip_top/  
example_top.sv  
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>/tb/example_tb.sv
```

b. For VHDL:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/*stub.vhdl  
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>/rtl/ip_top/  
example_top.sv  
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/imports/<Component_Name>/tb/example_tb.sv
```

5. Run `synplify_pro` synthesis to generate the `.edf` file. Then, close the `synplify_pro` project.
6. Open new Vivado project with Project Type as **Post-synthesis Project** and select the **Target Language** same as selected at the time of generating the IP core.
7. Add the `synplify_pro` generated `.edf` file to the Vivado project as **Design Source**.
8. Add the `.dcp` file generated in steps 1 and 2 to the Vivado project as **Design Source**. For example:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srcs/sources_1/ip/<Component_Name>/<Component_Name>.dcp
```

9. Add the `.xdc` file generated in step 2 to the Vivado project as **constraint**. For example:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srsrcs/constrs_1/imports/par/example_design.xdc
```

10. Run implementation flow with the Vivado tool. For details about implementation, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

**Note:** Similar steps can be followed for the user design using appropriate .dcp and .xdc files.

## CLOCK\_DEDICATED\_ROUTE Constraints and BUFG Instantiation

If the GCIO pin and MMCM are not allocated in the same bank, the CLOCK\_DEDICATED\_ROUTE constraint must be set to BACKBONE. To use the BACKBONE route, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM input. MIG manages these constraints for designs generated with the **Reference Input Clock** option selected as **Differential** (at **Advanced > FPGA Options > Reference Input**). Also, MIG handles the IP and example design flows for all scenarios.

If the design is generated with the **Reference Input Clock** option selected as **No Buffer** (at **Advanced > FPGA Options > Reference Input**), the CLOCK\_DEDICATED\_ROUTE constraints and BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV instantiation based on GCIO and MMCM allocation needs to be handled manually for the IP flow. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations for the IP flow.

For an example design flow with **No Buffer** configurations, MIG generates the example design with differential buffer instantiation for system clock pins. MIG generates clock constraints in the example\_design.xdc. It also generates a CLOCK\_DEDICATED\_ROUTE constraint as the "BACKBONE" and instantiates BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV between GCIO and MMCM input if the GCIO and MMCM are not in same bank to provide a complete solution. This is done for the example design flow as a reference when it is generated for the first time.

If in the example design, the I/O pins of the system clock pins are changed to some other pins with the I/O pin planner, the CLOCK\_DEDICATED\_ROUTE constraints and BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV instantiation needs to be managed manually. A DRC error is reported for the same.

# Test Bench

This chapter contains information about the test bench provided in the Vivado® Design Suite.

The Memory Controller is generated along with a simple test bench to verify the basic read and write operations. The stimulus contains 10 consecutive writes followed by 10 consecutive reads for data integrity check.

## SECTION IV: RLD RAM 3

Overview

Product Specification

Core Architecture

Designing with the Core

Design Flow Steps

Example Design

Test Bench

## Overview

The Xilinx® UltraScale™ architecture includes the RLD RAM 3 Memory Interface Solutions (MIS) core. This MIS core provides solutions for interfacing with these DRAM memory types. Both a complete Memory Controller and a physical (PHY) layer only solution are supported. The UltraScale architecture for the RLD RAM 3 cores is organized in the following high-level blocks.

- **Controller** – The controller accepts burst transactions from the User Interface and generates transactions to and from the RLD RAM 3. The controller takes care of the DRAM timing parameters and refresh.
- **Physical Layer** – The physical layer provides a high-speed interface to the DRAM. This layer includes the hard blocks inside the FPGA and the soft blocks calibration logic necessary to ensure optimal timing of the hard blocks interfacing to the DRAM.

The new hard blocks in the UltraScale architecture allow interface rates of up to 2,133 Mb/s to be achieved.

- These hard blocks include:
  - Data serialization and transmission
  - Data capture and deserialization
  - High-speed clock generation and synchronization
  - Fine delay elements per pin with voltage and temperature tracking
- The soft blocks include:
  - **Memory Initialization** – The calibration modules provide an initialization routine for RLD RAM 3. The delays in the initialization process are bypassed to speed up simulation time.
  - **Calibration** – The calibration modules provide a complete method to set all delays in the hard blocks and soft IP to work with the memory interface. Each bit is individually trained and then combined to ensure optimal interface performance. Results of the calibration process are available through the Xilinx debug tools. After completion of calibration, the PHY layer presents raw interface to the DRAM.
- **Application Interface** – The "User Interface" layer provides a simple FIFO interface to the application. Data is buffered and read data is presented in request order.



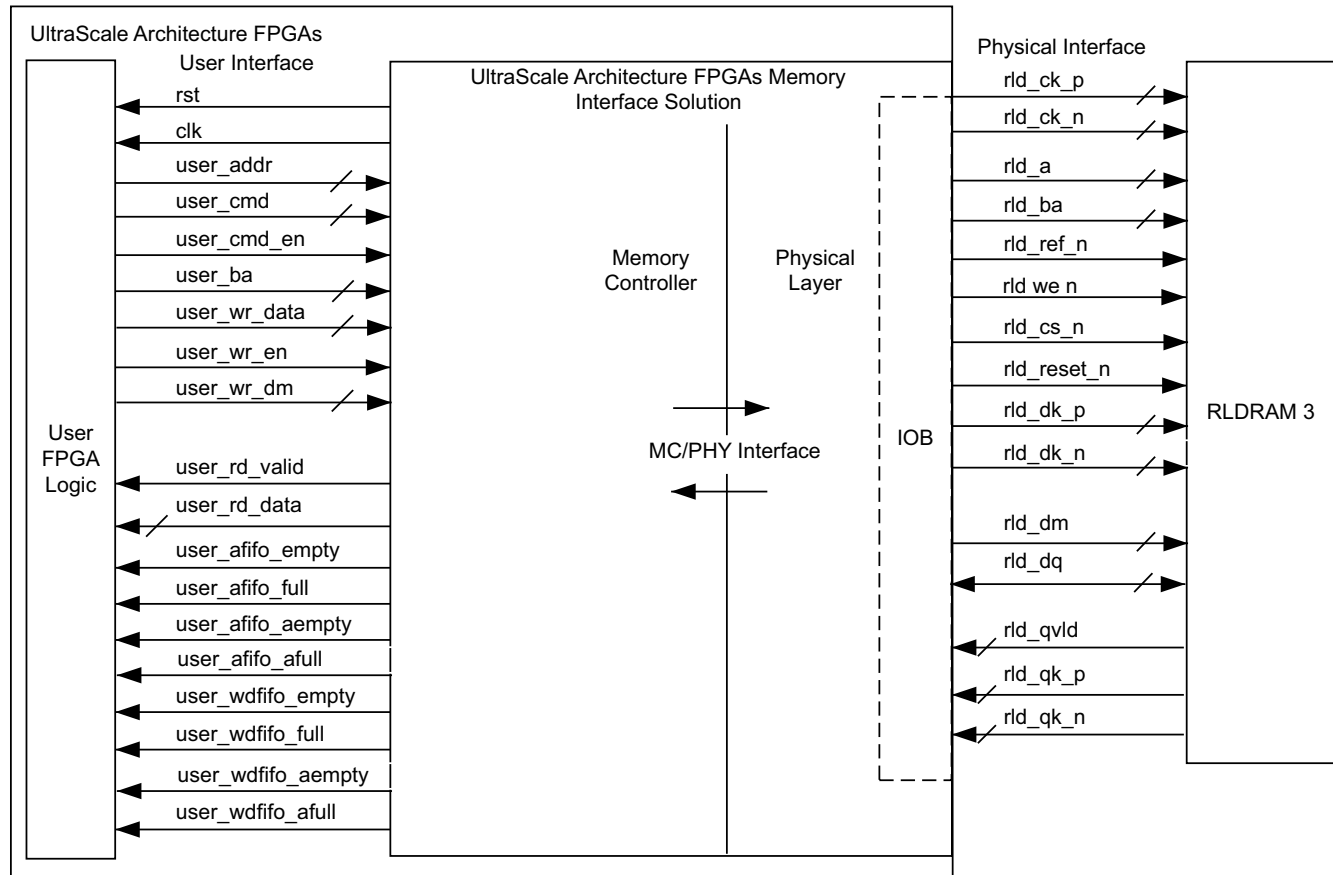


Figure 15-1: UltraScale Architecture FPGAs Memory Interface Solution

## Feature Summary

- Component support for interface widths of 18 and 36 bits

Table 15-1: Supported Configurations

Interface Width	Burst Length	Number of Device
36	BL2, BL4	1, 2
18	BL2, BL4, BL8	1, 2

- ODT support
- RLD RAM 3 initialization support
- Source code delivery in Verilog
- 4:1 memory to FPGA logic interface clock ratio
- Interface calibration and training information available through the Vivado hardware manager

---

## Licensing and Ordering Information

This Xilinx LogiCORE IP module is provided at no additional cost with the Xilinx Vivado Design Suite under the terms of the [Xilinx End User License](#). Information about this and other Xilinx LogiCORE IP modules is available at the [Xilinx Intellectual Property](#) page. For information about pricing and availability of other Xilinx LogiCORE IP modules and tools, contact your [local Xilinx sales representative](#).

### License Checkers

If the IP requires a license key, the key must be verified. The Vivado® design tools have several license checkpoints for gating licensed IP through the flow. If the license check succeeds, the IP can continue generation. Otherwise, generation halts with error. License checkpoints are enforced by the following tools:

- Vivado design tools: Vivado synthesis
- Vivado implementation
- write\_bitstream (Tcl command)



---

**IMPORTANT:** *IP license level is ignored at checkpoints. The test confirms a valid license exists. It does not check IP license level.*

---

## Product Specification

### Standards

For more information on UltraScale™ architecture documents, see [References, page 303](#).

### Performance

#### Maximum Frequencies

For more information on the maximum frequencies, see *Kintex UltraScale Architecture Data Sheet, DC and AC Switching Characteristics* (DS892) [\[Ref 2\]](#).

### Resource Utilization

#### Kintex UltraScale Devices

[Table 16-1](#) provides approximate resource counts on Kintex® UltraScale devices.

**Table 16-1: Device Utilization – Kintex UltraScale FPGAs**

Parameter Values	Device Resources						
Interface Width	FFs	LUTs	Memory LUTs	RAMB36E2/ RAMB18E2	BUFGs	PLLE3_ADV	MMCME3_ADV
36	5,311	4,799	463	29	6	2	1

Resources required for the UltraScale architecture FPGAs MIS core have been estimated for the Kintex UltraScale devices. These values were generated using Vivado® IP catalog. They are derived from post-synthesis reports, and might change during implementation.

---

## Port Descriptions

There are three port categories at the top-level of the memory interface core called the “user design.”

- The first category is the memory interface signals that directly interfaces with the RLD RAM. These are defined by the Micron<sup>®</sup> RLD RAM 3 specification.
- The second category is the application interface signals which is the “user interface.” These are described in the [Protocol Description, page 260](#).
- The third category includes other signals necessary for proper operation of the core. These include the clocks, reset, and status signals from the core. The clocking and reset signals are described in their respective sections.

The active-High `init_calib_complete` signal indicates that the initialization and calibration are complete and that the interface is now ready to accept commands for the interface.

## Core Architecture

This chapter describes the UltraScale™ device FPGAs Memory Interface Solutions core with an overview of the modules and interfaces.

### Overview

Figure 17-1 shows a high-level block diagram of the RLDRAM 3 MIS core. This figure shows both the internal FPGA connections to the user interface for initiating read and write commands, and the external interface to the memory device.

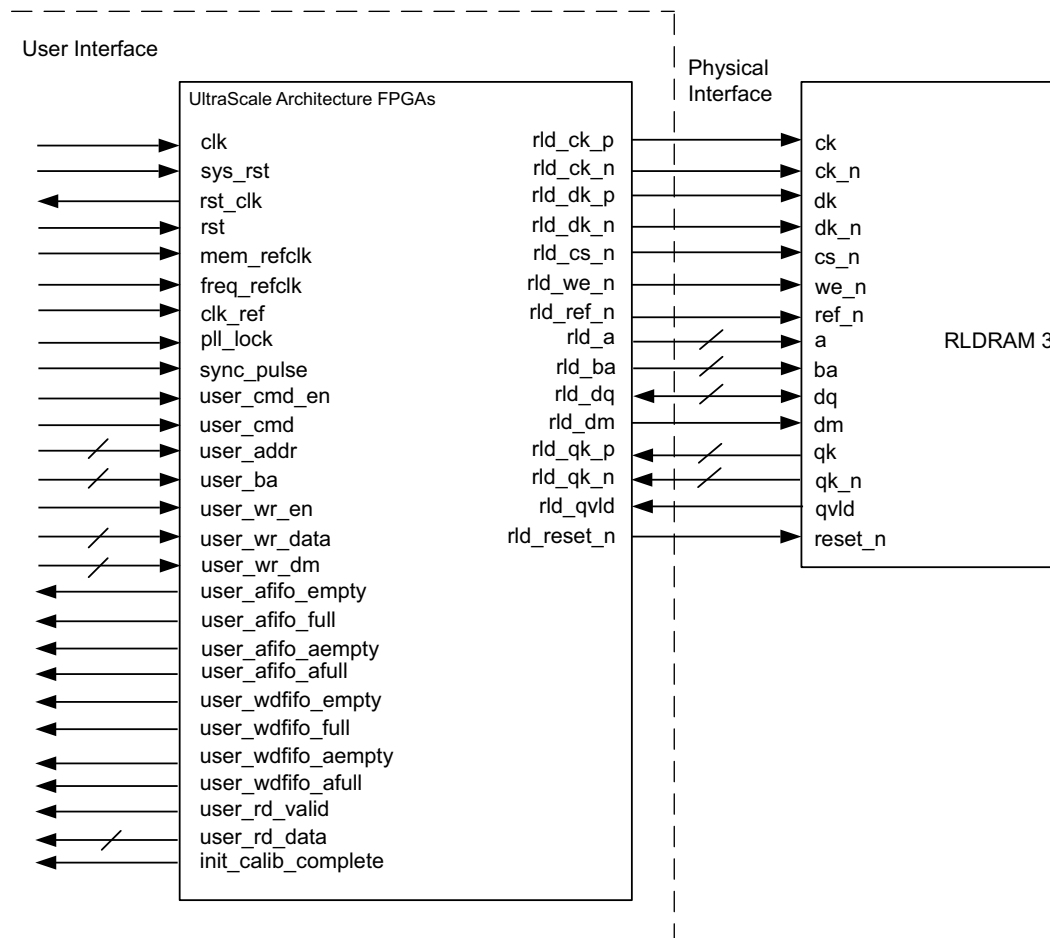


Figure 17-1: High-Level Block Diagram of RLDRAM 3 Interface Solution

Figure 17-2 shows the UltraScale architecture FPGAs Memory Interface Solutions diagram.

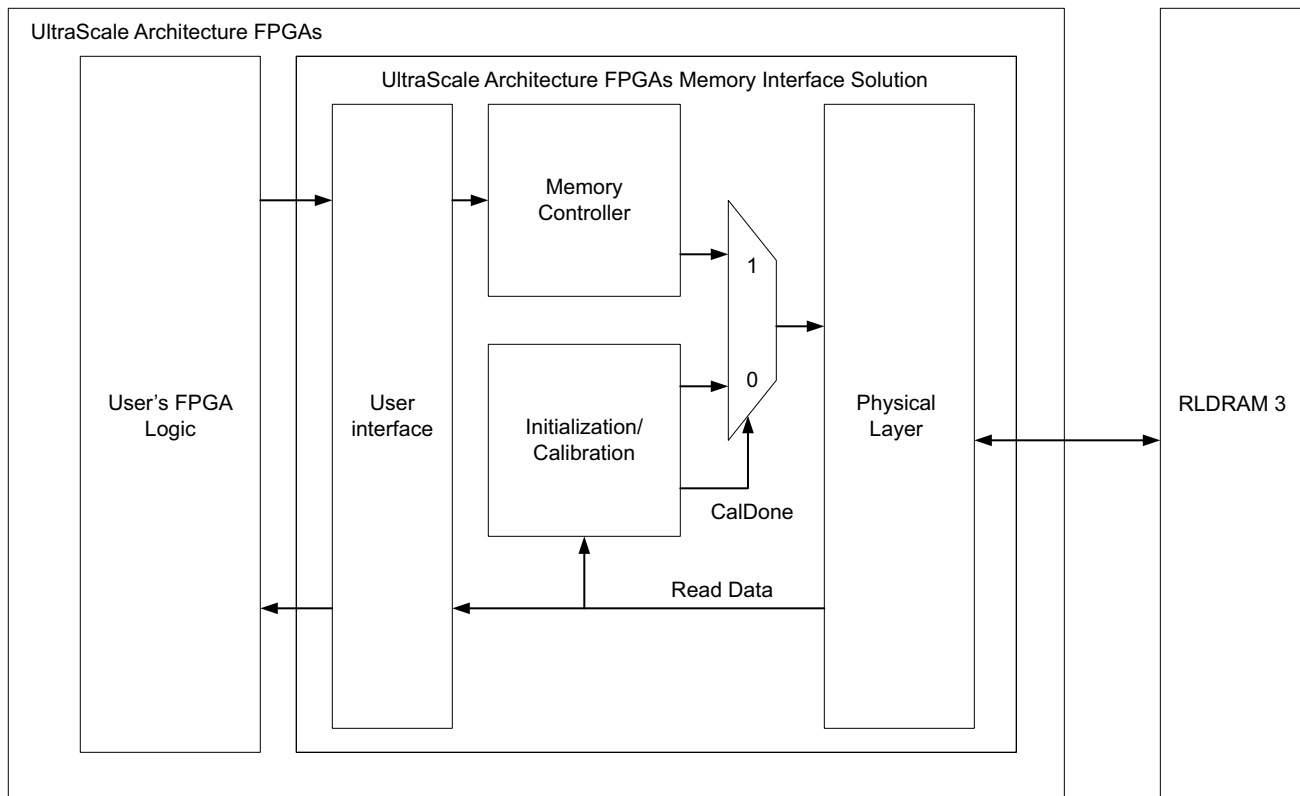


Figure 17-2: UltraScale Architecture FPGAs Memory Interface Solution Core

The user interface uses a simple protocol based entirely on SDR signals to make read and write requests. See [User Interface in Chapter 18](#) for more details describing this protocol.

The Memory Controller takes commands from the user interface and adheres to the protocol requirements of the RLDRAM 3 device. See [Memory Controller](#) for more details.

The physical interface generates the proper timing relationships and DDR signaling to communicate with the external memory device, while conforming to the RLDRAM 3 protocol and timing requirements. See [Physical Interface in Chapter 18](#) for more details.

---

## Memory Controller

The Memory Controller (MC) enforces the RLDRAM 3 access requirements and interfaces with the PHY. The controller processes commands in order, so the commands presented to the controller is the order in which they are presented to the memory device.

The MC first receives commands from the user interface and determines if the command can be processed immediately or needs to wait. When all requirements are met, the command is placed on the PHY interface. For a write command, the controller generates a signal for the user interface to provide the write data to the PHY. This signal is generated based on the memory configuration to ensure the proper command-to-data relationship. Auto-refresh commands are inserted into the command flow by the controller to meet the memory device refresh requirements.

The data bus is shared for read and write data in RLDRAM 3. Switching from read commands to write commands and vice versa introduces gaps in the command stream due to switching the bus. For better throughput, changes in the command bus should be minimized when possible.

CMD\_PER\_CLK is a top-level parameter used to determine how many memory commands are provided to the controller per FPGA logic clock cycle. It depends on nCK\_PER\_CLK and the burst length. For example if nCK\_PER\_CLK = 4, the CMD\_PER\_CLK is set to 1 for burst length = 8 and CMD\_PER\_CLK is set to 2 for burst length = 4 and CMD\_PER\_CLK is set to 4 for burst length = 2.

---

## User Interface Allocation

The address bits on `c0_rld3_user_addr` bus needs to be assigned as below.

---

## PHY

PHY is considered the low-level physical interface to an external RLDRAM 3 device as well as all calibration logic for ensuring reliable operation of the physical interface itself. PHY generates the signal timing and sequencing required to interface to the memory device.

PHY contains the following features:

- Clock/address/control-generation logics
- Write and read datapaths
- Logic for initializing the SDRAM after power-up

In addition, PHY contains calibration logic to perform timing training of the read and write datapaths to account for system static and dynamic delays.

## Overall PHY Architecture

The UltraScale architecture PHY is composed of dedicated blocks and soft calibration logic. The dedicated blocks are structured adjacent to one another with back-to-back interconnects to minimize the clock and datapath routing necessary to build high performance physical layers.

The Memory Controller and calibration logic communicate with this dedicated PHY in the slow frequency clock domain, which is divided by 4. A more detailed block diagram of the PHY design is shown in [Figure 17-3](#).

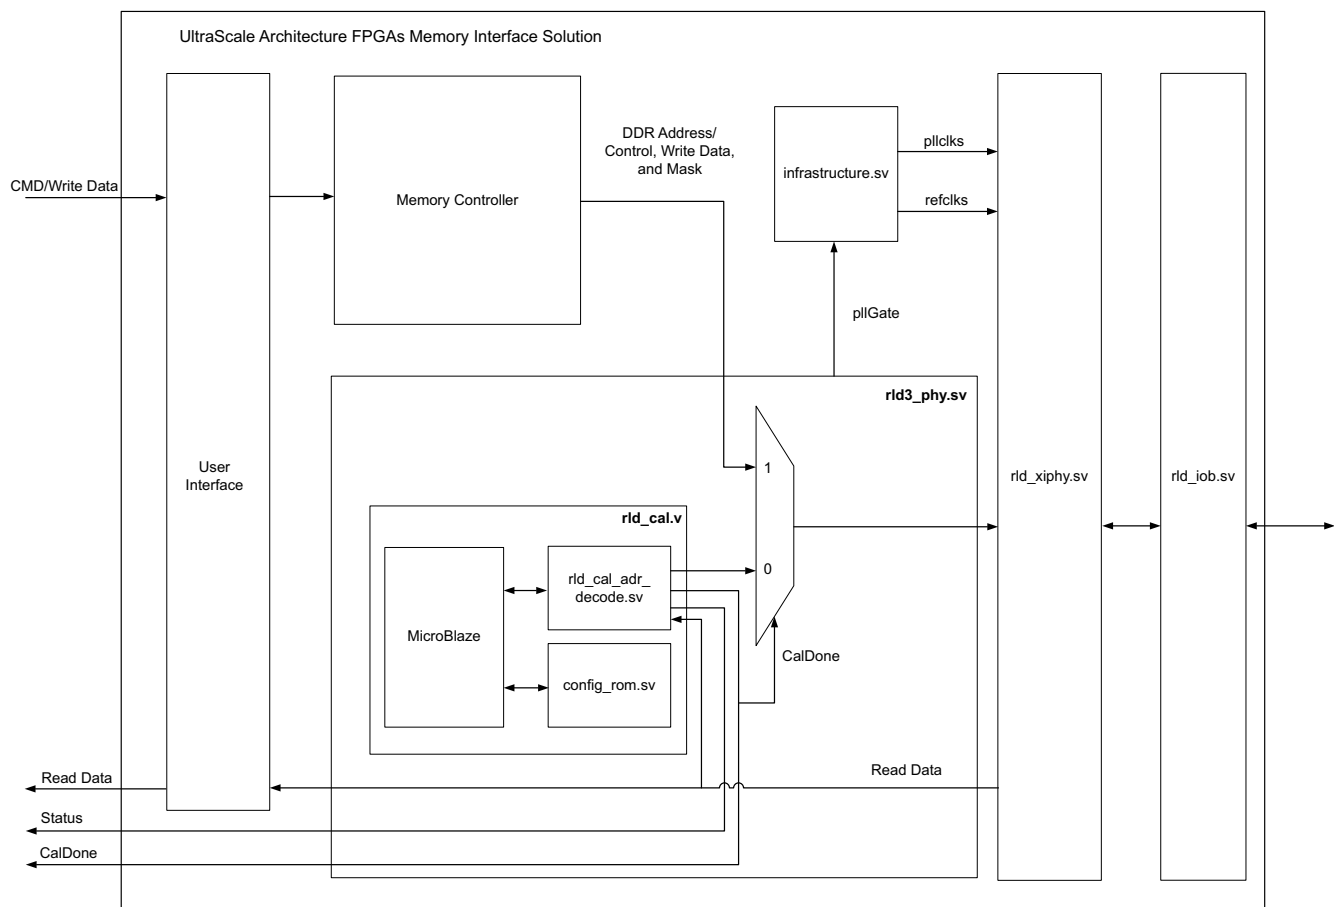


Figure 17-3: PHY Block Diagram

The MC is designed to separate out the command processing from the low-level PHY requirements to ensure a clean separation between the controller and physical layer. The command processing can be replaced with custom logic if desired, while the logic for interacting with the PHY stays the same and can still be used by the calibration logic.



Table 17-1: PHY Modules

Module Name	Description
rld3_phy.sv	Contains infrastructure (infrastructure.sv), rld_cal.sv, rld_xiphy.sv, and MUXes between the calibration and the Memory Controller.
rld_cal.sv	Contains the MicroBlaze processing system and associated logic.
rld_cal_adr_decode.sv	FPGA logic interface for the MicroBlaze processor.
config_rom.sv	Configuration storage for calibration options.
debug_microblaze.sv	MicroBlaze processor
rld_iob.sv	Instantiates all byte IOB modules
rld_iob_byte.sv	Generates the I/O buffers for all the signals in a given byte lane.
rld_addr_mux.sv	Address MUX
rld_rd_bit_slip.sv	Read bit slip
rld_wr_lat.sv	Write latency
rld_xiphy.sv	Top-level XIPHY module

The PHY architecture encompasses all of the logic contained in `rld_xiphy.sv`. The PHY contains wrappers around dedicated hard blocks to build up the memory interface from smaller components. A byte lane contains all of the clocks, resets, and datapaths for a given subset of I/O. Multiple byte lanes are grouped together, along with dedicated clocking resources, to make up a single bank memory interface. For more information on the hard silicon physical layer architecture, see the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4].

The memory initialization and calibration are implemented in C programming on a small soft core processor. The MicroBlaze™ Controller System (MCS) is configured with an I/O Module and block RAM. The `rld_cal_adr_decode.sv` module provides the interface for the processor to the rest of the system and implements helper logic. The `config_rom.sv` module stores settings that control the operation of initialization and calibration, providing run time options that can be adjusted without having to recompile the source code.

The MicroBlaze I/O module interface updates at a maximum rate of once every three clock cycles, which is not always fast enough for implementing all of the functions required in calibration. A helper circuit implemented in `rld_cal_adr_decode.sv` is required to obtain commands from the registers and translate at least a portion into single-cycle accuracy for submission to the PHY. In addition, it supports command repetition to enable back-to-back read transactions and read data comparison.

## Memory Initialization and Calibration Sequence

After deassertion of the system reset, calibration logic performs power-on initialization sequence for the memory. This is followed by several stages of timing calibration for the write and read datapaths. PHY indicates calibration is finished and controller begins issuing commands to the memory.

Figure 17-4 shows the overall flow of memory initialization and the different stages of calibration.

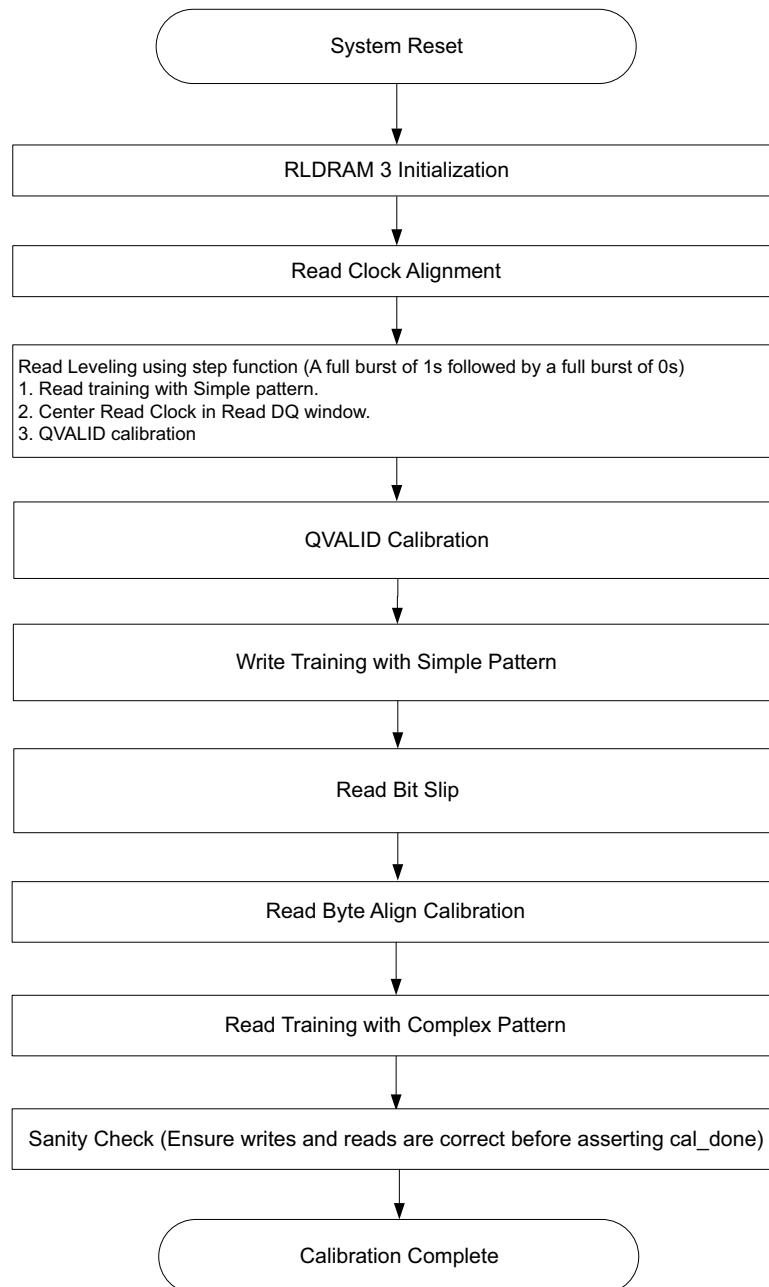


Figure 17-4: PHY Overall Initialization and Calibration Sequence

# Designing with the Core

This chapter includes guidelines and additional information to facilitate designing with the core.

---

## Clocking

The memory interface requires one MMCM, one TXPLL per I/O bank used by the memory interface and two BUFGCE\_DIVs. These clocking components are used to create the proper clock frequencies and phase shifts necessary for the proper operation of the memory interface.

There are two TXPLLs per bank. If a bank is shared by two memory interfaces, both TXPLLs in that bank are used.

**Note:** MIG generates the appropriate clocking structure and no modifications to the RTL are supported.

The MIG tool generates the appropriate clocking structure for the desired interface. This structure must not be modified. The allowed clock configuration is as follows:

- Differential reference clock source connected to GCIO
- GCIO to MMCM (located in center bank of memory interface)
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) driving FPGA logic and all TXPLLs
- MMCM to BUFGCE\_DIV (located at center bank of memory interface) divide by two mode driving 1/2 rate FPGA logic
- Clocking pair of the interface must be in the same SLR of memory interface for the SSI technology devices

## Requirements

### ***GCIO***

- Must use a differential I/O standard
- Must be in the same I/O column as the memory interface
- Must be in the same SLR of memory interface for the SSI technology devices

### ***MMCM***

- MMCM is used to generate the FPGA logic system clock (1/4 of the memory clock)
- Must be located in the center bank of memory interface
- Must use internal feedback
- Input clock frequency divided by input divider must be  $\geq 70$  MHz ( $\text{CLKIN}_x / D \geq 70$  MHz)
- Must use integer multiply and output divide values

### ***BUFGCE\_DIVs and Clock Roots***

- One BUFGCE\_DIV is used to generate the system clock to FPGA logic and another BUFGCE\_DIV is used to divide the system clock by two.
- BUFGCE\_DIVs and clock roots must be located in center most bank of the memory interface.
  - For two bank systems, either bank can be used. MIG is always referred to the top-most selected bank in the Vivado Integrated Design Environment (IDE) as the center bank.
  - For four bank systems, either of the center banks can be chosen. MIG refers to the second bank from the top-most selected bank as the center bank.
  - Both the BUFGCE\_DIVs must be in the same bank.

### ***TXPLL***

- CLKOUTPHY from TXPLL drives XIPHY within its bank
- TXPLL must be set to use a CLKFBOUT phase shift of 90°
- TXPLL must be held in reset until the MMCM lock output goes High
- Must use internal feedback

Figure 18-1 shows an example of the clocking structure for a three bank memory interface. The GCIO drives the MMCM located at the center bank of the memory interface. MMCM drives both the BUFGCE\_DIVs located in the same bank. The BUFG CE\_DIV (which is used to generate system clock to FPGA logic) output drives the TXPLLs used in each bank of the interface.




Figure 18-1: Clocking Structure for Three Bank Memory Interface

The MMCM is placed in the center bank of the memory interface.

- For two bank systems, MMCM is placed in a bank with the most number of bytes selected. If they both have the same number of bytes selected in two banks, then MMCM is placed in the top bank.
- For four bank systems, MMCM is placed in a second bank from the top.

For designs generated with System Clock configuration of **No Buffer**, MMCM must not be driven by another MMCM/PLL. Cascading clocking structures MMCM → BUFG → MMCM and PLL → BUFG → MMCM are not allowed.

If the MMCM is driven by the GCIO pin of the other bank, then the CLOCK\_DEDICATED\_ROUTE constraint with value "BACKBONE" must be set on the net that is driving MMCM or on the MMCM input. Setting up the CLOCK\_DEDICATED\_ROUTE constraint on the net is preferred. But when the same net is driving two MMCMs, the CLOCK\_DEDICATED\_ROUTE constraint must be managed by considering which MMCM needs the BACKBONE route.

In such cases, the CLOCK\_DEDICATED\_ROUTE constraint can be set on the MMCM input. To use the "BACKBONE" route, any clock buffer that exists in the same CMT tile as the GCIO must exist between the GCIO and MMCM input. The clock buffers that exist in the I/O CMT are BUFG, BUFGCE, BUFGCTRL, and BUFGCE\_DIV. So MIG instantiates BUFG between the GCIO and MMCM when the GCIO pins and MMCM are not in the same bank (see [Figure 18-1](#)).

If the GCIO pin and MMCM are allocated in different banks, MIG generates CLOCK\_DEDICATED\_ROUTE constraints with value as "BACKBONE." If the GCIO pin and MMCM are allocated in the same bank, there is no need to set any constraints on the MMCM input.

Similarly when designs are generated with System Clock Configuration as a **No Buffer** option, you must take care of the "BACKBONE" constraint and the BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV between GCIO and MMCM if GCIO pin and MMCM are allocated in different banks. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations. For more information on clocking, see the *UltraScale Architecture Clocking Resources User Guide* (UG572) [[Ref 3](#)].

**Note:** If two different GCIO pins are used for two MIG IP cores in the same bank, center bank of the memory interface is different for each IP. MIG generates MMCM LOC and CLOCK\_DEDICATED\_ROUTE constraints accordingly.

## Sharing of Input Clock Source (sys\_clk\_p)

If the same GCIO pin must be used for two IP cores, generate the two IP cores with System Clock Configuration option as **No Buffer**. Perform the following changes in the wrapper file in which both IPs are instantiated:

1. MIG generates a single-ended input for system clock pins, such as `sys_clk_i`. Connect the differential buffer output to the single-ended system clock inputs (`sys_clk_i`) of both the IP cores.
2. System clock pins must be allocated within the same I/O column of the memory interface pins allocated. Add the pin LOC constraints for system clock pins and clock constraints in your top-level XDC.

3. You must add a "BACKBONE" constraint on the net that is driving the MMCM or on the MMCM input if GCIO pin and MMCM are not allocated in the same bank. Apart from this, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM to use the "BACKBONE" route.

**Note:**

- The Ultrascale architecture includes an independent XIPHY power supply and TXPLL for each XIPHY. This results in clean, low jitter clocks for the memory system.
- Skew spanning across multiple BUFGs is not a concern because single point of contact exists between BUFG → TXPLL and the same BUFG → System Clock Logic.
- System input clock cannot span I/O columns because the longer the clock lines span, the more jitter is picked up.

---

## Resets

An asynchronous reset (`sys_rst`) input is provided. This is an active-High reset and the `sys_rst` must assert for a minimum pulse width of 5 ns. The `sys_rst` can be an internal or external pin.

---

## PCB Guidelines for RLDRAM 3

Strict adherence to all documented RLDRAM 3 PCB guidelines is required for successful operation. For more information on PCB guidelines, see the *UltraScale Architecture PCB Design and Pin Planning User Guide* (UG583) [Ref 5].

---

## Pin and Bank Rules

### RLDRAM 3 Pin Rules

The rules are for single rank memory interfaces.

- Address/control means `cs_n`, `ref_n`, `we_n`, `ba`, `ck`, `reset_n`, and `a`.
- All groups such as, Data, Address/Control, and System clock interfaces must be selected in a single column.
- Pins in a byte lane are numbered N0 to N12.

- Byte lanes in a bank are designed by T0, T1, T2, or T3. Nibbles within a byte lane are distinguished by a “U” or “L” designator added to the byte lane designator (T0, T1, T2, or T3). Thus they are T0L, T0U, T1L, T1U, T2L, T2U, T3L, and T3U.

**Note:** There are two PLLs per bank and a controller uses one PLL in every bank that is being used by the interface.

1. Read Clock ( $qk/qk\_n$ ), Write Clock ( $dk/dk\_n$ ),  $dq$ ,  $qvld$ , and  $dm$ .
  - a. Read Clock pairs ( $qkx\_p/n$ ) must be placed on N0 and N1 pins.  $dq$  associated with a  $qk/qk\_n$  pair must be in same byte lane on pins N2 to N11.
  - b. For the data mask off configurations, ensure that  $dm$  pin on the RLDRAM 3 device is grounded. When data mask is enabled, one  $dm$  pin is associated with nine bits in x18 devices or with 18 bits in x36 devices. It must be placed in its associated  $dq$  byte lanes as listed:
    - For x18 part,  $dm[0]$  must be allocated in  $dq[8:0]$  allocated byte group and  $dm[1]$  must be allocated in  $dq[17:9]$ .
    - For x36 part,  $dm[0]$  must be allocated in  $dq[8:0]$  or  $dq[26:18]$  allocated byte lane. Similarly  $dm[1]$  must be allocated in  $dq[17:9]$  or  $dq[35:27]$  allocated byte group.  $dq$  must be placed on one of the pins from N2 to N11 in the byte lane.
  - c.  $dk/dk\_n$  must be allocated to any P-N pair in the same byte lane as  $ck/ck\_n$  in the address/control bank.

**Note:** Pin 12 is not part of a pin pair and must not be used for differential clocks.

  - d.  $qvld$  (x18 device) or  $qvld0$  (x36 device) must be placed on one of the pins from N2 to N12 in the  $qk0$  or  $qk1$  data byte lane.  $qvld1$  (x36 device) must be placed on one of the pins from N2 to N12 in of the  $qk2$  or  $qk3$  data byte lane.
2. Byte lanes are configured as either data or address/control.
  - a. Pin N12 can be used for address/control in a data byte lane.
  - b. No data signals ( $qvalid$ ,  $dq$ ,  $dm$ ) can be placed in an address/control byte lane.
3. Address/control can be on any of the 13 pins in the address/control byte lanes. Address/control must be contained within the same bank.
4. One  $vrp$  pin per bank is used and a DCI is required for the interfaces. A  $vrp$  pin is required in I/O banks containing inputs as well as output only banks. It is required in output only banks because address/control signals use SSTL12\_DCI to enable usage of controlled output impedance. A DCI cascade is not permitted. All rules for the DCI in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4] must be followed.
5.  $ck$  must be on the PN pair in the Address/Control byte lane.
6.  $reset\_n$  can be on any pin as long as FPGA logic timing is met and I/O standard can be accommodated for the chosen bank (SSTL12).



7. Banks can be shared between two controllers.
  - a. Each byte lane is dedicated to a specific controller (except for `reset_n`).
  - b. Byte lanes from one controller cannot be placed inside the other. For example, with controllers A and B, "AABB" is allowed, while "ABAB" is not.
8. All I/O banks used by the memory interface must be in the same column.
9. All I/O banks used by the memory interface must be in the same SLR of the column for the SSI technology devices.
10. Maximum height of interface is three contiguous banks for 72-bit wide interface.
11. Bank skipping is not allowed.
12. The input clock for the MMCM in the interface must come from the a GCIO pair in the I/O column used for the memory interface. For more information, see [Clocking](#), page 251.
13. There are dedicated  $V_{REF}$  pins (not included in the rules above). If an external  $V_{REF}$  is not used, the  $V_{REF}$  pins must be pulled to ground by a resistor value specified in the *UltraScale™ Architecture FPGAs SelectIO™ Resources User Guide* (UG571) [Ref 4]. These pins must be connected appropriately for the standard in use.
14. The interface must be contained within the same I/O bank type (High Range or High Performance). Mixing bank types is not permitted with the exceptions of the `reset_n` in step 6 and the input clock mentioned in step 11.

### RLDRAM 3 Pinout Examples



**IMPORTANT:** Due to the calibration stage, there is no need for `set_input_delay/`  
`set_output_delay` on the MIG. Ignore the unconstrained inputs and outputs for MIG and the signals which are calibrated.

Table 18-1 shows an example of an 18-bit RLDRAM 3 interface contained within one bank. This example is for a component interface using one x18 RLDRAM3 component with Address Multiplexing.

Table 18-1: 18-Bit RLDRAM 3 Interface Contained in One Bank

Bank	Signal Name	Byte Group	I/O Type	Special Designation
1	qvld0	T3U_12	–	–
1	dq8	T3U_11	N	–
1	dq7	T3U_10	P	–
1	dq6	T3U_9	N	–
1	dq5	T3U_8	P	–
1	dq4	T3U_7	N	DBC-N

**Table 18-1: 18-Bit RLDRAM 3 Interface Contained in One Bank (Cont'd)**

Bank	Signal Name	Byte Group	I/O Type	Special Designation
1	dq3	T3U_6	P	DBC-P
1	dq2	T3L_5	N	–
1	dq1	T3L_4	P	–
1	dq0	T3L_3	N	–
1	dm0	T3L_2	P	–
1	qk0_n	T3L_1	N	DBC-N
1	qk0_p	T3L_0	P	DBC-P
1	reset_n	T2U_12	–	–
1	we#	T2U_11	N	–
1	a18	T2U_10	P	–
1	a17	T2U_9	N	–
1	a14	T2U_8	P	–
1	a13	T2U_7	N	QBC-N
1	a10	T2U_6	P	QBC-P
1	a9	T2L_5	N	–
1	a8	T2L_4	P	–
1	a5	T2L_3	N	–
1	a4	T2L_2	P	–
1	a3	T2L_1	N	QBC-N
1	a0	T2L_0	P	QBC-P
1	–	T1U_12	–	–
1	ba3	T1U_11	N	–
1	ba2	T1U_10	P	–
1	ba1	T1U_9	N	–
1	ba0	T1U_8	P	–
1	dk1_n	T1U_7	N	QBC-N
1	dk1_p	T1U_6	P	QBC-P
1	dk0_n	T1L_5	N	–
1	dk0_p	T1L_4	P	–
1	ck_n	T1L_3	N	–
1	ck_p	T1L_2	P	–
1	ref_n	T1L_1	N	QBC-N

Table 18-1: 18-Bit RLDRAM 3 Interface Contained in One Bank (Cont'd)

Bank	Signal Name	Byte Group	I/O Type	Special Designation
1	cs_n	T1L_0	P	QBC-P
1	vrp	T0U_12	–	–
1	dq17	T0U_11	N	–
1	dq16	T0U_10	P	–
1	dq15	T0U_9	N	–
1	dq14	T0U_8	P	–
1	dq13	T0U_7	N	DBC-N
1	dq12	T0U_6	P	DBC-P
1	dq11	T0L_5	N	–
1	dq10	T0L_4	P	–
1	dq9	T0L_3	N	–
1	dm1	T0L_2	P	–
1	qk1_n	T0L_1	N	DBC-N
1	qk1_p	T0L_0	P	DBC-P

---

## Protocol Description

This core has the following interfaces:

- [Memory Interface](#)
- [User Interface](#)
- [Physical Interface](#)

### Memory Interface

The RLDRAM 3 MIS core is customizable to support several configurations. The specific configuration is defined by Verilog parameters in the top-level of the core.

### User Interface

The user interface connects to an FPGA user design to the RLDRAM 3 MIS core to simplify interactions between the user design and the external memory device.

#### ***Command Request Signals***

The user interface provides a set of signals used to issue a read or write command to the memory device. These signals are summarized in [Table 18-2](#).

Table 18-2: User Interface Request Signals

Signal	Direction	Description
user_cmd_en	Input	Command Enable. This signal issues a read or write request and indicates that the corresponding command signals are valid.
user_cmd[2 × CMD_PER_CLK – 1:0]	Input	<p>Command. This signal issues a read, write, or NOP request. When user_cmd_en is asserted:</p> <p>2'b00 = Write Command</p> <p>2'b01 = Read Command</p> <p>2'b10 = NOP</p> <p>2'b11 = NOP</p> <p>The NOP command is useful when more than one command per clock cycle must be provided to the Memory Controller yet not all command slots are required in a given clock cycle. The Memory Controller acts on the other commands provided and ignore the NOP command. NOP is not supported when CMD_PER_CLK == 1. CMD_PER_CLK is a top-level parameter used to determine how many memory commands are provided to the controller per FPGA logic clock cycle, it depends on nCK_PER_CLK and the burst length (see <a href="#">Figure 18-2</a>)</p>
user_addr[CMD_PER_CLK × ADDR_WIDTH – 1:0]	Input	Command Address. This is the address to use for a command request. It is valid when user_cmd_en is asserted.
user_ba[CMD_PER_CLK × BANK_WIDTH – 1:0]	Input	Command Bank Address. This is the address to use for a write request. It is valid when user_cmd_en is asserted.
user_wr_en	Input	Write Data Enable. This signal issues the write data and data mask. It indicates that the corresponding user_wr_* signals are valid.
user_wr_data[2 × nCK_PER_CLK × DATA_WIDTH – 1:0]	Input	Write Data. This is the data to use for a write request and is composed of the rise and fall data concatenated together. It is valid when user_wr_en is asserted.
user_wr_dm[2 × nCK_PER_CLK × DM_WIDTH – 1:0]	Input	Write Data Mask. When active-High, the write data for a given selected device is masked and not written to the memory. It is valid when user_wr_en is asserted.
user_afifo_empty	Output	Address FIFO empty. If asserted, the command buffer is empty.
user_wdfifo_empty	Output	Write Data FIFO empty. If asserted, the write data buffer is empty.
user_afifo_full	Output	Address FIFO full. If asserted, the command buffer is full, and any writes to the FIFO are ignored until deasserted.

Table 18-2: User Interface Request Signals (Cont'd)

Signal	Direction	Description
user_wdfifo_full	Output	Write Data FIFO full. If asserted, the write data buffer is full, and any writes to the FIFO are ignored until deasserted.
user_afifo_aempty	Output	Address FIFO almost empty. If asserted, the command buffer is almost empty.
user_afifo_afull	Output	Address FIFO almost full. If asserted, the command buffer is almost full.
user_wdfifo_aempty	Output	Write Data FIFO almost empty. If asserted, the write data buffer is almost empty.
user_wdfifo_afull	Output	Write Data FIFO almost full. If asserted, the Write Data buffer is almost full.
user_rd_valid[nCK_PER_CLK – 1:0]	Output	Read Valid. This signal indicates that data read back from memory is available on user_rd_data and should be sampled.
user_rd_data[2 × nCK_PER_CLK × DATA_WIDTH – 1:0]	Output	Read Data. This is the data read back from the read command.
init_calib_complete	Output	Calibration Done. This signal indicates back to the user design that read calibration is complete and requests can now take place.
cx_rld3_ui_clk	Output	This User Interface clock should be one quarter of the RLD3 clock.
cx_rld3_ui_clk_sync_rst	Output	This is the active-High user interface reset.
cx_calib_error	Output	When asserted indicates error during calibration.
addn_ui_clkout1	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout2	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout3	Output	Additional clock outputs provided based on user requirement.
addn_ui_clkout4	Output	Additional clock outputs provided based on user requirement.
dbg_clk	Output	Debug Clock. Do not connect any signals to dbg_clk and keep the port open during instantiation.

### Interfacing with the Core through the User Interface

The width of certain user interface signals is dependent on the system clock frequency and the burst length. This allows the client to send multiple commands per FPGA logic clock cycle as might be required for certain configurations.

**Note:** Both write and read commands in the same `user_cmd` cycle is not allowed.

Figure 18-2 shows the `user_cmd` signal and how it is made up of multiple commands depending on the configuration.




Figure 18-2: Multiple Commands for `user_cmd` Signal

As shown in Figure 18-2, four command slots are present in a single user interface clock cycle for BL2. Similarly, two command slots are present in a single user interface clock cycle for BL4. These command slots are serviced sequentially and the return data for read commands are presented at the user interface in the same sequence. Note that the read data might not be available in the same slot as that of its read command. The slot of a read data is determined by the timing requirements of the controller and its command slot. One such example is mentioned in the following BL2 design configuration.

Assume that the following set of commands is presented at the user interface for a given user interface cycle.

Table 18-3: Command Set in User Interface Cycle

Slots	Commands
0	RD0
1	NOP
2	RD1
3	NOP

It is not guaranteed that the read data appears in {DATA0, NOP, DATA1, NOP} order. It might also appear in {NOP, DATA0, NOP, DATA1} or {NOP, NOP, DATA0, DATA1} etc. orders. In any case, the sequence of the commands are maintained.

### User Address Bit Allocation Based on RLD RAM 3 Configuration

The width of the address bus (not including bank address bits) at the user interface is always set in the multiple of 20 bits, which accounts for the maximum possible address width for RLD RAM 3 device. Depending on the RLD RAM 3 device configuration, the actual address width can be < 20 bits. Table 18-4 summarizes the address width for the various RLD RAM 3 configurations.

**Table 18-4: RLDram 3 Address Width**

Burst Length	Data Width	Address Width
2	18	20
2	36	19
4	18	18
4	36	18
8	18	18
8	36	Not supported by RLDram 3

The address bits at the user interface are concatenated based on the burst length as shown in [Figure 18-2](#). If the address width is < 20 bits, pad the unused bits with zero. An example for x36 burst length 4 configuration is shown here:

{00, (18-bit address), 00, (18-bit address)}



The user interface protocol for the RLDRAM 3 four-word burst architecture is shown in Figure 18-3.




Figure 18-3: RLDRAM 3 User Interface Protocol (Four-Word Burst Architecture)

Before any requests can be accepted, the `ui_clk_sync_rst` signal must be deasserted Low. After the `ui_clk_sync_rst` signal is deasserted, the user interface FIFOs can accept commands and data for storage. The `init_calib_complete` signal is asserted after the memory initialization procedure and PHY calibration are complete, and the core can begin to service client requests.

A command request is issued by asserting `user_cmd_en` as a single cycle pulse. At this time, the `user_cmd`, `user_addr`, and `user_ba` signals must be valid. To issue a read request, `user_cmd` is set to 2'b01, while for a write request, `user_cmd` is set to 2'b00. For a write request, the data is to be issued in the same cycle as the command by asserting the `user_wr_en` signal High and presenting valid data on `user_wr_data` and `user_wr_dm`. The user interface protocol for the RLDRAM 3 eight-word burst architecture is shown in Figure 18-4.




Figure 18-4: RLDRAM 3 User Interface Protocol (Eight-Word Burst Architecture)

When a read command is issued some time later (based on the configuration and latency of the system), the `user_rd_valid[0]` signal is asserted, indicating that `user_rd_data` is now valid, while `user_rd_valid[1]` is asserted indicating that `user_rd_data` is valid, as shown in Figure 18-5. The read data should be sampled on the same cycle that `user_rd_valid[0]` and `user_rd_valid[1]` are asserted because the core does not buffer returning data. This functionality can be added in, if desired.

The Memory Controller only puts commands on certain slots to the PHY such that the `user_rd_valid` signals are all asserted together and return the full width of data, but the extra `user_rd_valid` signals are provided in case of controller modifications.




Figure 18-5: User Interface Protocol Read Data

## Physical Interface

The physical interface is the connection from the FPGA MIS core to an external RLDRAM 3 device. The I/O signals for this interface are defined in Table 18-5. These signals can be directly connected to the corresponding signals on the RLDRAM 3 device.

Table 18-5: Physical Interface Signals

Signal	Direction	Description
rld_ck_p	Output	System Clock CK. This is the address/command clock to the memory device.
rld_ck_n	Output	System Clock CK#. This is the inverted system clock to the memory device.
rld_dk_p	Output	Write Clock DK. This is the write clock to the memory device.
rld_dk_n	Output	Write Clock DK#. This is the inverted write clock to the memory device.
rld_a	Output	Address. This is the address supplied for memory operations.
rld_ba	Output	Bank Address. This is the bank address supplied for memory operations.
rld_cs_n	Output	Chip Select CS#. This is the active-Low chip select control signal for the memory.
rld_we_n	Output	Write Enable WE#. This is the active-Low write enable control signal for the memory.
rld_ref_n	Output	Refresh REF#. This is the active-Low refresh control signal for the memory.
rld_dm	Output	Data Mask DM. This is the active-High mask signal, driven by the FPGA to mask data that a user does not want written to the memory during a write command.
rld_dq	Input/Output	Data DQ. This is a bidirectional data port, driven by the FPGA for writes and by the memory for reads.

Table 18-5: Physical Interface Signals (Cont'd)

Signal	Direction	Description
rld_qk_p	Input	Read Clock QK. This is the read clock returned from the memory edge aligned with read data on rld_dq. This clock (in conjunction with QK#) is used by the PHY to sample the read data on rld_dq.
rld_qk_n	Input	Read Clock QK#. This is the inverted read clock returned from the memory. This clock (in conjunction with QK) is used by the PHY to sample the read data on rld_dq.
rld_reset_n	Output	RLDRAM 3 reset pin. This is the active-Low reset to the RLDRAM 3 device.

# Design Flow Steps

This chapter describes customizing and generating the core, constraining the core, and the simulation, synthesis and implementation steps that are specific to this IP core. More detailed information about the standard Vivado® design flows and the Vivado IP integrator can be found in the following Vivado Design Suite user guides:

- *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#)
- *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#)
- *Vivado Design Suite User Guide: Getting Started* (UG910) [\[Ref 9\]](#)
- *Vivado Design Suite User Guide: Logic Simulation* (UG900) [\[Ref 10\]](#)

---

## Customizing and Generating the Core



**CAUTION!** *The Windows operating system has a 260-character limit for path lengths, which can affect the Vivado tools. To avoid this issue, use the shortest possible names and directory locations when creating projects, defining IP or managed IP projects, and creating block designs.*

---

This section includes information about using Xilinx® tools to customize and generate the core in the Vivado Design Suite.

If you are customizing and generating the core in the IP integrator, see the *Vivado Design Suite User Guide: Designing IP Subsystems using IP Integrator* (UG994) [\[Ref 7\]](#) for detailed information. IP integrator might auto-compute certain configuration values when validating or generating the design. To check whether the values change, see the description of the parameter in this chapter. To view the parameter value, run the `validate_bd_design` command in the Tcl Console.

You can customize the IP for use in your design by specifying values for the various parameters associated with the IP core using the following steps:

1. Select the IP from the Vivado IP catalog.
2. Double-click the selected IP or select the Customize IP command from the toolbar or right-click menu.

For more information about generating the core in Vivado, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8] and the *Vivado Design Suite User Guide: Getting Started* (UG910) [Ref 9].

## User Parameters

Table 19-1 shows the relationship between the GUI fields in the Vivado IDE and the User Parameters (which can be viewed in the Tcl Console).

Table 19-1: Vivado IDE Parameter to User Parameter Relationship

Vivado IDE Parameter/Value <sup>(1)</sup>	User Parameter/Value <sup>(1)</sup>	Default Value
System Clock Configuration	System_Clock	Differential
Internal V <sub>REF</sub>	Internal_Vref	TRUE
DCI Cascade	DCI_Cascade	FALSE
Debug Signal for Controller	Debug_Signal	Disable
Clock 1 (MHz)	ADDN_UI_CLKOUT1_FREQ_HZ	None
Clock 2 (MHz)	ADDN_UI_CLKOUT2_FREQ_HZ	None
Clock 3 (MHz)	ADDN_UI_CLKOUT3_FREQ_HZ	None
Clock 4 (MHz)	ADDN_UI_CLKOUT4_FREQ_HZ	None
I/O Power Reduction	IOPowerReduction	OFF
Enable System Ports	Enable_SysPorts	TRUE
I/O Power Reduction	IO_Power_Reduction	FALSE
Default Bank Selections	Default_Bank_Selections	FALSE
Reference Clock	Reference_Clock	FALSE
Enable System Ports	Enable_SysPorts	TRUE
Clock Period (ps)	C0.RLD3_TimePeriod	1,071
Input Clock Period (ps)	C0.RLD3_InputClockPeriod	13,947
General Interconnect to Memory Clock Ratio	C0.RLD3_PhyClockRatio	4:1
Configuration	C0.RLD3_MemoryType	Components
Memory Part	C0.RLD3_MemoryPart	MT44K16M36RB-093
Data Width	C0.RLD3_DataWidth	36
Data Mask	C0.RLD3_DataMask	TRUE
Burst Length	C0.RLD3_BurstLength	8
	C0.RLD3_MemoryVoltage	1.2

### Notes:

- Parameter values are listed in the table where the Vivado IDE parameter value differs from the user parameter value. Such values are shown in this table as indented below the associated parameter.

## Output Generation

For details, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

---

## MIG I/O Planning

For details on I/O planning, see [MIG I/O Planning, page 164](#).

---

## Constraining the Core

This section contains information about constraining the core in the Vivado Design Suite, if applicable.

### Required Constraints

Internal  $V_{REF}$  is not automatically set for RLDRAM 3 by the tool. It is required to be assigned to all Address/Control and DQ banks used by RLDRAM 3 through the pin planner tool or by the constraint file. A sample constraint for RLDRAM 3 is shown here:

```
set_property INTERNAL_VREF 0.600 [get_iobanks 45]
```

### Device, Package, and Speed Grade Selections

This section is not applicable for this IP core.

### Clock Frequencies

This section is not applicable for this IP core.

### Clock Management

For information on clocking, see [Clocking, page 251](#).

### Clock Placement

This section is not applicable for this IP core.

### Banking

This section is not applicable for this IP core.

### Transceiver Placement

This section is not applicable for this IP core.

## I/O Standard and Placement

The MIG tool generates the appropriate I/O standards and placement based on the selections made in the Vivado IDE for the interface type and options.



---

**IMPORTANT:** *The `set_input_delay` and `set_output_delay` constraints are not needed on the external memory interface pins in this design due to the calibration process that automatically runs at start-up. Warnings seen during implementation for the pins can be ignored.*

---

---

## Simulation

For comprehensive information about Vivado simulation components, as well as information about using supported third-party tools, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [\[Ref 10\]](#).

---

## Synthesis and Implementation

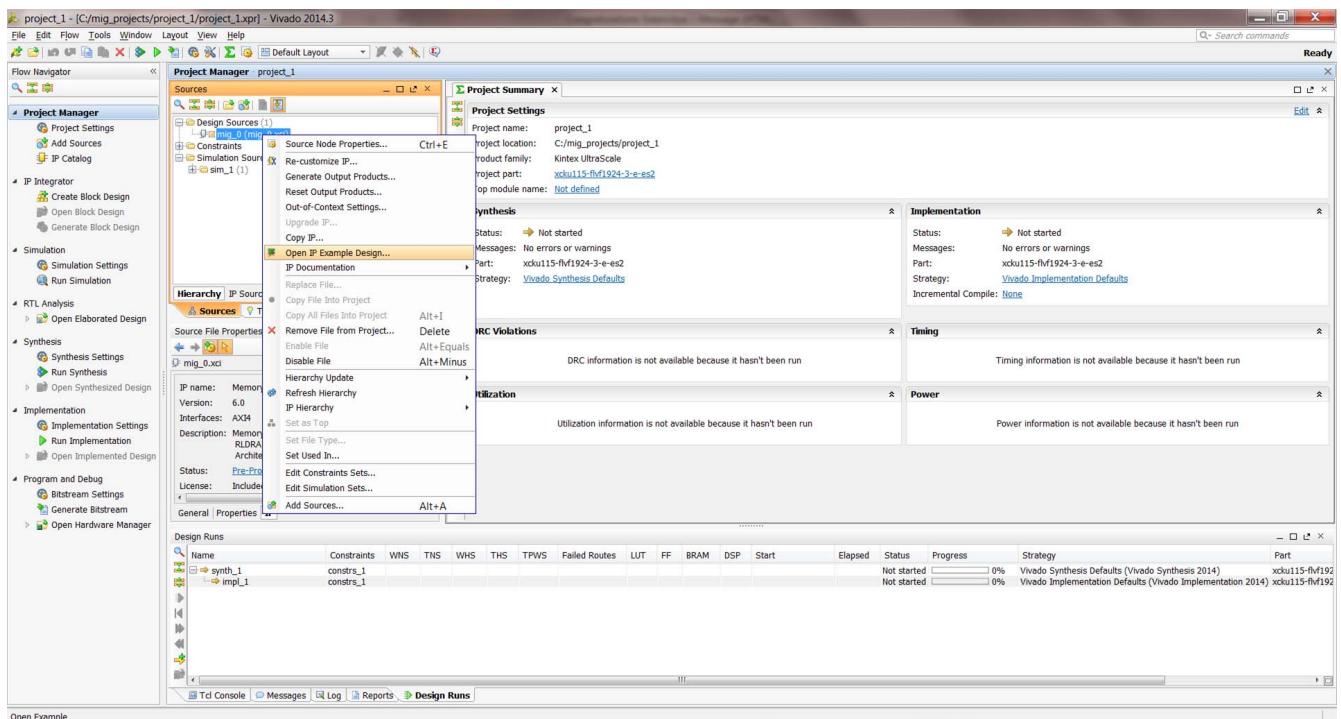
For details about synthesis and implementation, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 8\]](#).



## Example Design

This chapter contains information about the example design provided in the Vivado® Design Suite.

Vivado supports Open IP Example Design flow. To create the example design using this flow, right-click the IP in the **Source Window**, as shown in [Figure 20-1](#) and select **Open IP Example Design**.



**Figure 20-1: Open IP Example Design**

This option creates a new Vivado project. Upon selecting the menu, a dialog box to enter the directory information for the new design project opens.

Select a directory, or use the defaults, and click **OK**. This launches a new Vivado with all of the example design files and a copy of the IP.

Figure 20-1 shows the example design with the PHY only option selected (controller module does not get generated).




Figure 20-2: Open IP Example Design with PHY Only Option Selected

Figure 20-3 shows the example design with the PHY only option not selected (controller module is generated).

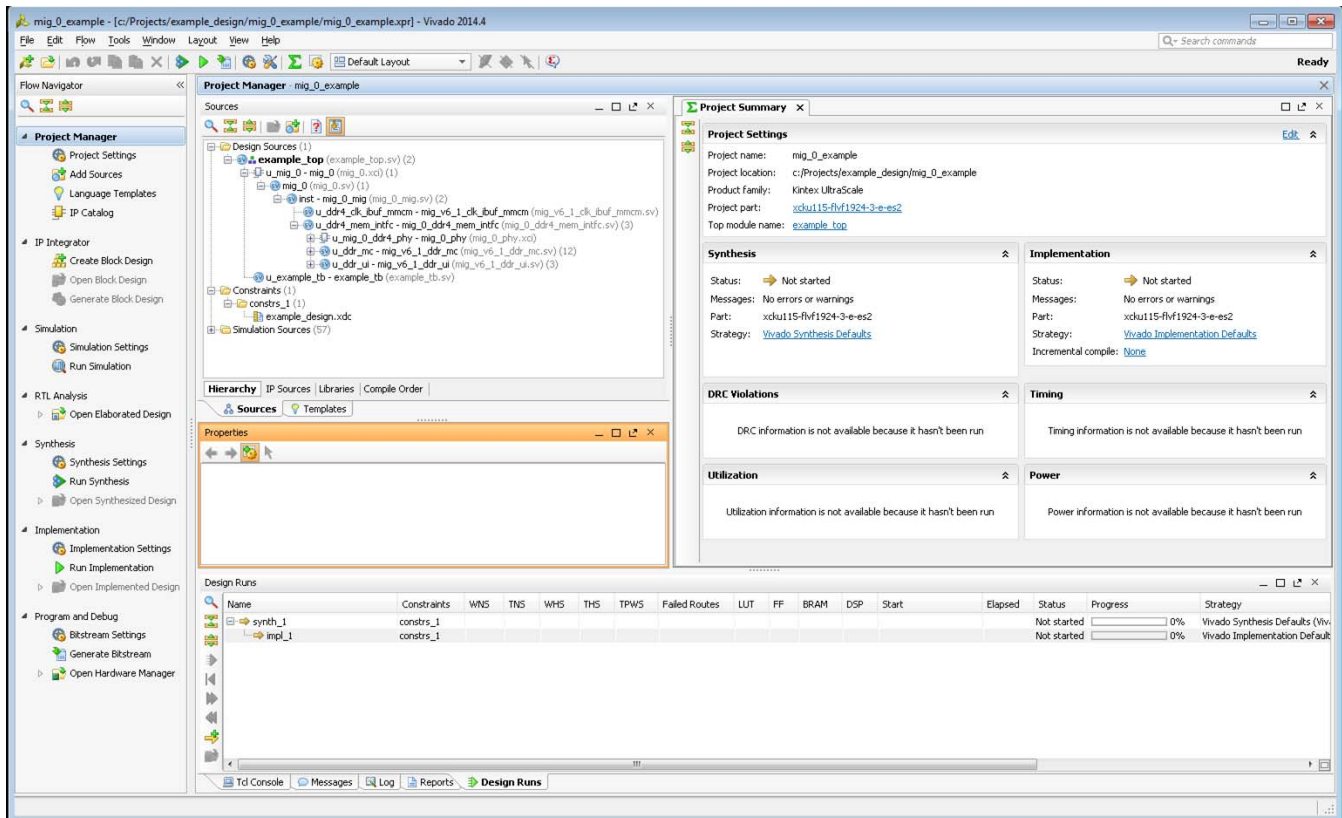


Figure 20-3: Open IP Example Design with PHY Only Option Not Selected

## Simulating the Example Design (Designs with Standard User Interface)

The example design provides a synthesizable test bench to generate a fixed simple data pattern to the Memory Controller. This test bench consists of an IP wrapper and an `example_tb` that generates 10 writes and 10 reads.

The example design can be simulated using one of the methods in the following sections.

---

## Project-Based Simulation

This method can be used to simulate the example design using the Vivado Integrated Design Environment (IDE). MIG delivers memory models for RLDRAM 3.

The Vivado simulator, QuestaSim, IES, and VCS tools are used for RLDRAM 3. IP verification at each software release. The Vivado simulation tool is used for RLDRAM 3. IP verification from 2015.1 Vivado software release. The following subsections describe steps to run a project-based simulation using each supported simulator tool.

### Project-Based Simulation Flow Using Vivado Simulator

1. In the **Open IP Example Design** Vivado project, under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as **Vivado Simulator**.
  - a. Under the **Simulation** tab, set the `xsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in [Figure 20-4](#). The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
3. Apply the settings and select **OK**.

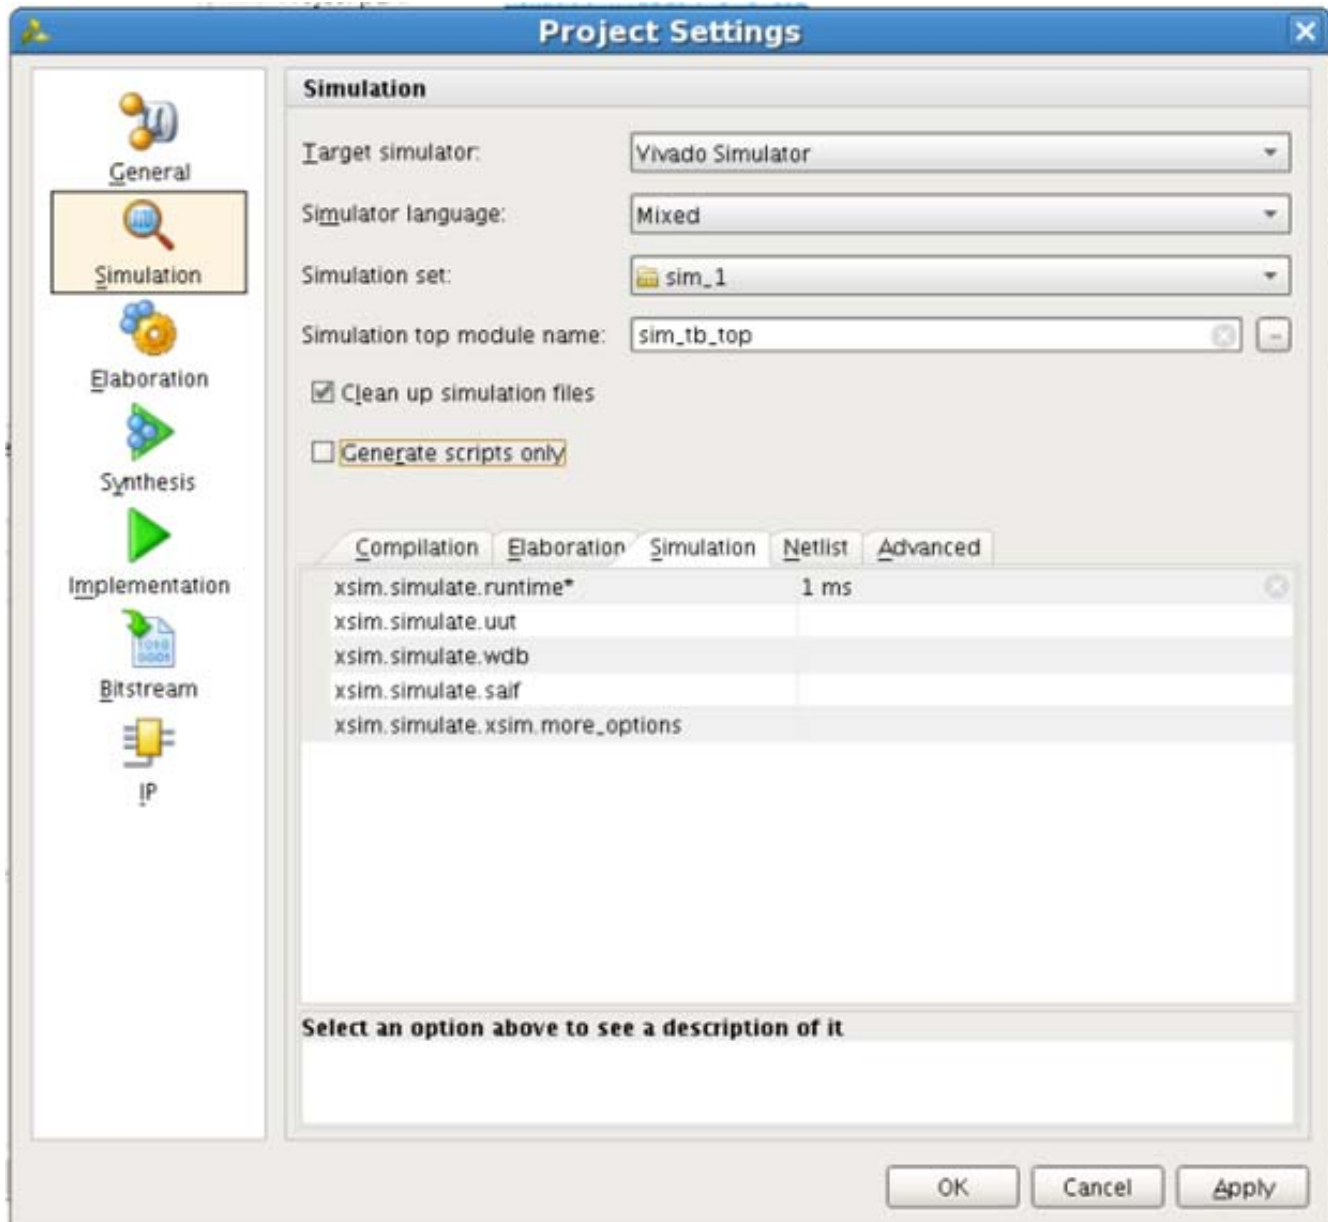


Figure 20-4: Simulation with Vivado Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 20-5.




Figure 20-5: Run Behavioral Simulation

5. Vivado invokes Vivado simulator and simulations are run in the Vivado simulator tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using QuestaSim

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as **QuestaSim/ModelSim Simulator**.
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `modelsim.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 20-6. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
3. Apply the settings and select **OK**.




Figure 20-6: Simulation with QuestaSim

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 20-7.






Figure 20-7: Run Behavioral Simulation

5. Vivado invokes QuestaSim and simulations are run in the QuestaSim tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using IES

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as Incisive Enterprise Simulator (IES).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `ies.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in Figure 20-8. The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
3. Apply the settings and select **OK**.






Figure 20-8: Simulation with IES Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 20-7.
5. Vivado invokes IES and simulations are run in the IES tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 10].

## Project-Based Simulation Flow Using VCS

1. Open a MIG example Vivado project (**Open IP Example Design...**), then under **Flow Navigator**, select **Simulation Settings**.
2. Select **Target simulator** as Verilog Compiler Simulator (VCS).
  - a. Browse to the compiled libraries location and set the path on **Compiled libraries location** option.
  - b. Under the **Simulation** tab, set the `vcs.simulate.runtime` to 1 ms (there are simulation RTL directives which stop the simulation after certain period of time, which is less than 1 ms) as shown in [Figure 20-9](#). The **Generate Scripts Only** option generates simulation scripts only. To run behavioral simulation, **Generate Scripts Only** option must be de-selected.
3. Apply the settings and select **OK**.

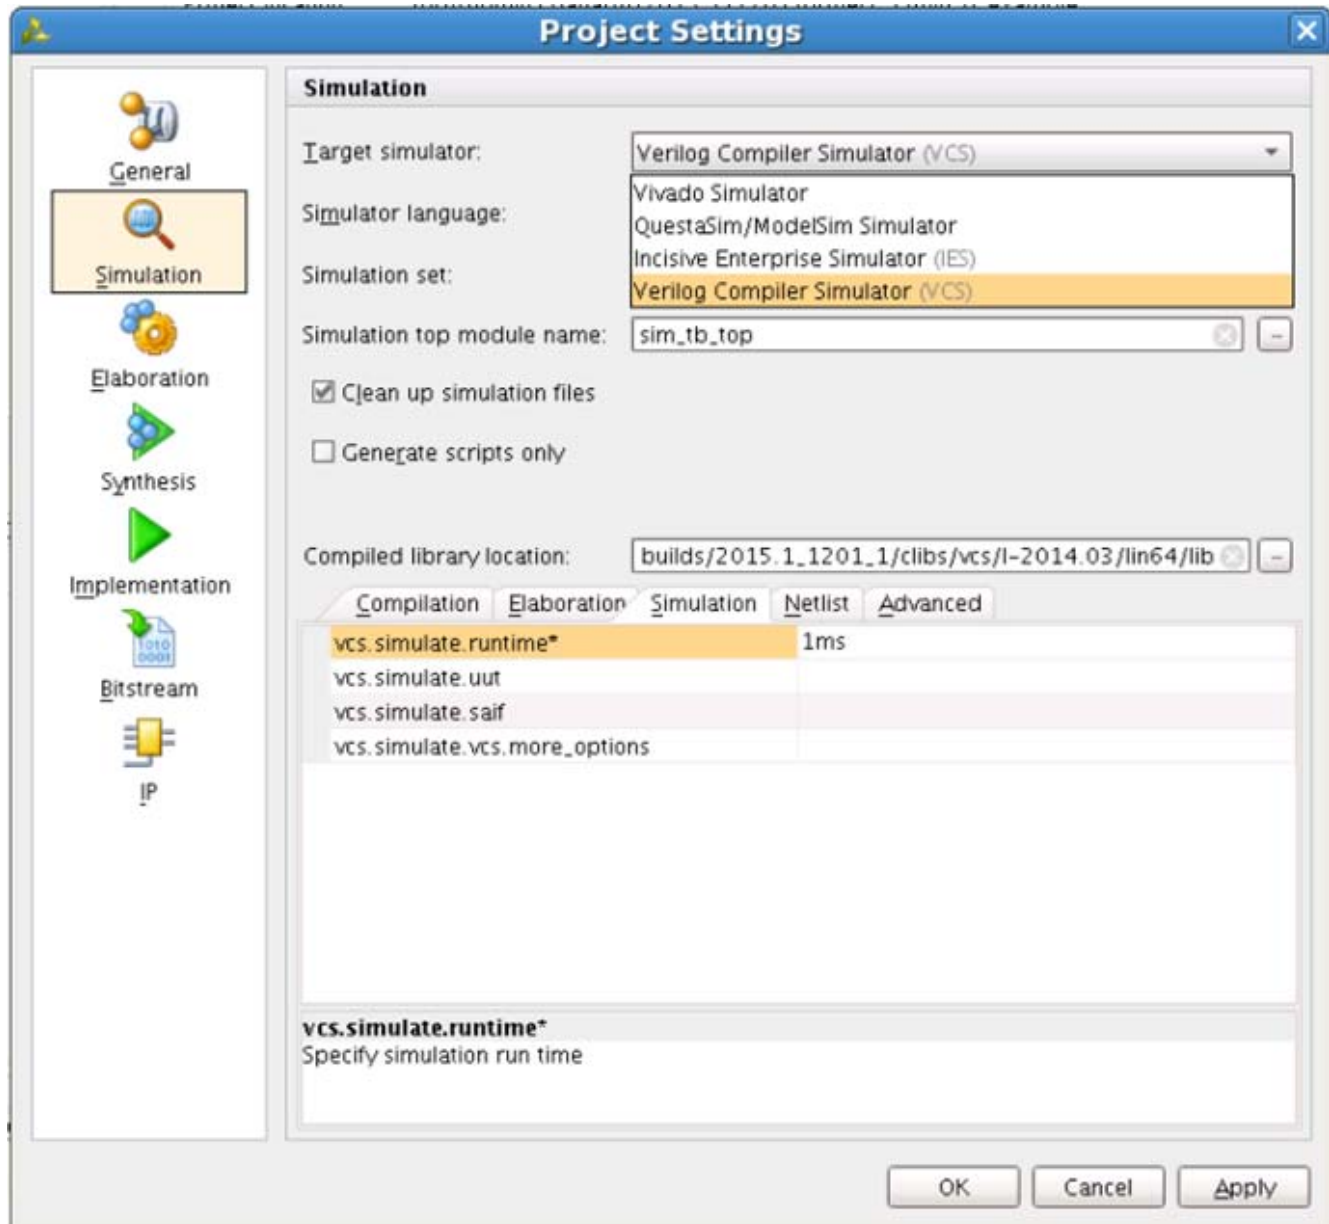


Figure 20-9: Simulation with VCS Simulator

4. In the **Flow Navigator** window, select **Run Simulation** and select **Run Behavioral Simulation** option as shown in Figure 20-7.
5. Vivado invokes VCS and simulations are run in the VCS tool. For more information, see the *Vivado Design Suite User Guide: Logic Simulation (UG900)* [Ref 10].

## Non-Project-Based Simulation



**IMPORTANT:** Xilinx® UNISIMS\_VER and SECUREIP library must be mapped into the simulator.

1. To run the simulation, go to this directory:

```
<project_dir>/example_project/<Component_Name>example/  
<Component_Name>_example.srscs/sim_1/imports/tb
```

If the MIG design is generated with the Component Name entered in the Vivado IDE as mig\_0, the simulation directory path is the following:

```
<project_dir>/example_project/mig_0_example/mig_0_example.srscs/  
sim_1/imports/tb
```

2. MIG delivers memory models for RLDRAM 3.
3. The QuestaSim, IES, and VCS simulation tools are used for verification of MIG IP at each software release.
4. Script files to run simulations with QuestaSim, IES, and VCS are generated in MIG generated output. See the `readme.txt` file located in the folder for running simulations. Other simulation tools can be used for MIG IP simulation but are not specifically verified by Xilinx.

## Simulation Speed

MIG provides a Vivado IDE option to reduce the simulation speed by selecting behavioral XIPHY model instead of UNISIM XIPHY model. Behavioral XIPHY model simulation is a default option for MIG designs. To select the simulation mode, click the **Advanced** tab and find the **Simulation Options** as shown in [Figure 20-10](#).

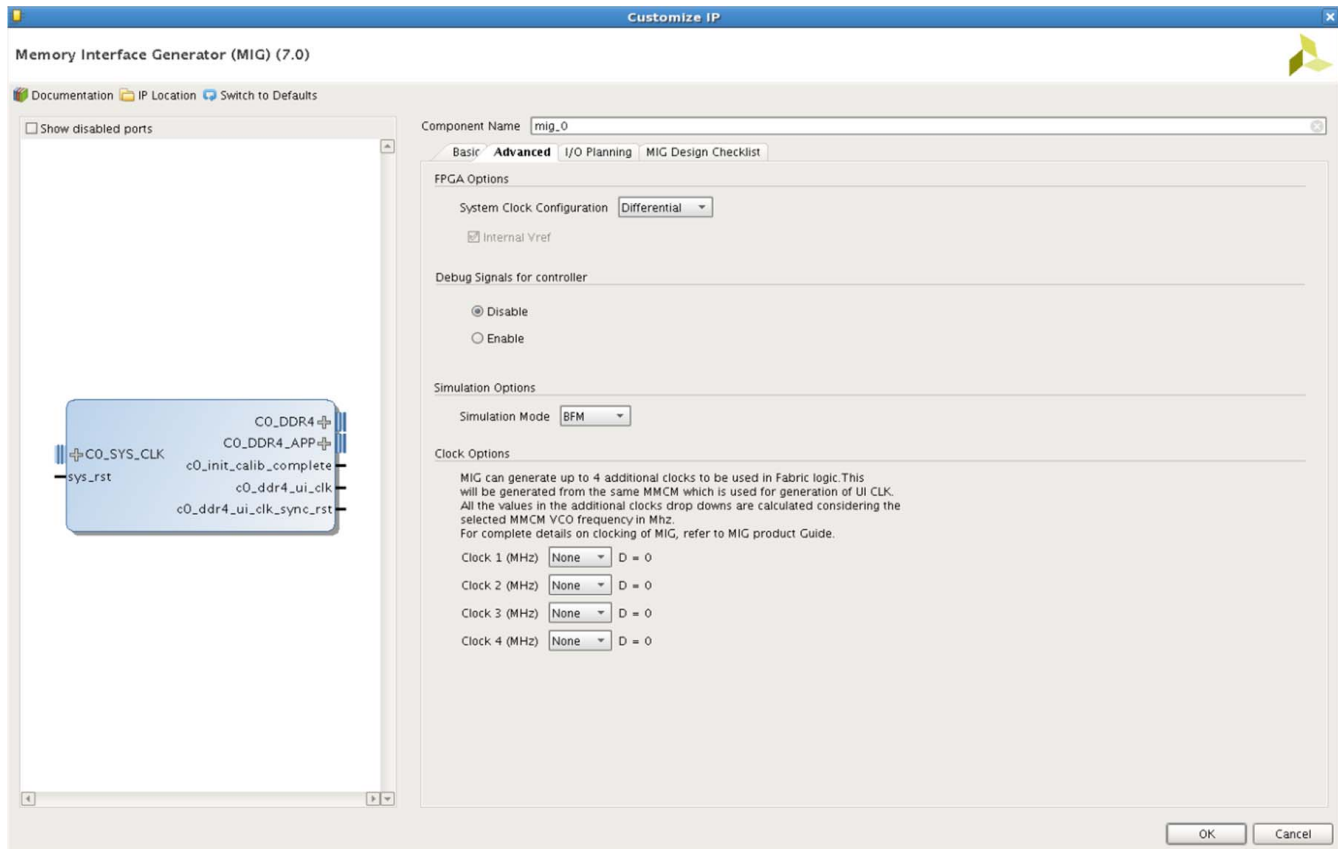


Figure 20-10: Advanced Tab – Simulation Options

The SIM\_MODE parameter in the RTL is given a different value based on the Vivado IDE selection.

- **SIM\_MODE = BFM** – If fast mode is selected in the Vivado IDE, the RTL parameter reflects this value for the SIM\_MODE parameter. This is the default option.
- **SIM\_MODE = FULL** – If FULL mode is selected in the Vivado IDE, XIPHY UNISIMs are selected and the parameter value in the RTL is FULL.

## CLOCK\_DEDICATED\_ROUTE Constraints and BUFG Instantiation

If the GCIO pin and MMCM are not allocated in the same bank, the CLOCK\_DEDICATED\_ROUTE constraint must be set to BACKBONE. To use the BACKBONE route, BUFG/BUFGCE/BUFGCTRL/BUFGCE\_DIV must be instantiated between GCIO and MMCM input. MIG manages these constraints for designs generated with the **Reference Input Clock** option selected as **Differential** (at **Advanced > FPGA Options > Reference Input**). Also, MIG handles the IP and example design flows for all scenarios.

If the design is generated with the **Reference Input Clock** option selected as **No Buffer** (at **Advanced > FPGA Options > Reference Input**), the `CLOCK_DEDICATED_ROUTE` constraints and `BUFG/BUFGCE/BUFGCTRL/BUFGCE_DIV` instantiation based on GCIO and MMCM allocation needs to be handled manually for the IP flow. MIG does not generate clock constraints in the XDC file for **No Buffer** configurations and you must take care of the clock constraints for **No Buffer** configurations for the IP flow.

For an example design flow with **No Buffer** configurations, MIG generates the example design with differential buffer instantiation for system clock pins. MIG generates clock constraints in the `example_design.xdc`. It also generates a `CLOCK_DEDICATED_ROUTE` constraint as the "BACKBONE" and instantiates `BUFG/BUFGCE/BUFGCTRL/BUFGCE_DIV` between GCIO and MMCM input if the GCIO and MMCM are not in same bank to provide a complete solution. This is done for the example design flow as a reference when it is generated for the first time.

If in the example design, the I/O pins of the system clock pins are changed to some other pins with the I/O pin planner, the `CLOCK_DEDICATED_ROUTE` constraints and `BUFG/BUFGCE/BUFGCTRL/BUFGCE_DIV` instantiation needs to be managed manually. A DRC error is reported for the same.

# Test Bench

This chapter contains information about the test bench provided in the Vivado® Design Suite.

The Memory Controller is generated along with a simple test bench to verify the basic read and write operations. The stimulus contains 10 consecutive writes followed by 10 consecutive reads for data integrity check.

# SECTION V: TRAFFIC GENERATOR

## Traffic Generator



## Chapter 22

# Traffic Generator

## Overview

This section describes the setup and behavior of the Traffic Generator. In the UltraScale™ architecture, Traffic Generator is instantiated in the example design (`example_top.sv`) to drive the memory design through the application interface (Figure 22-1).

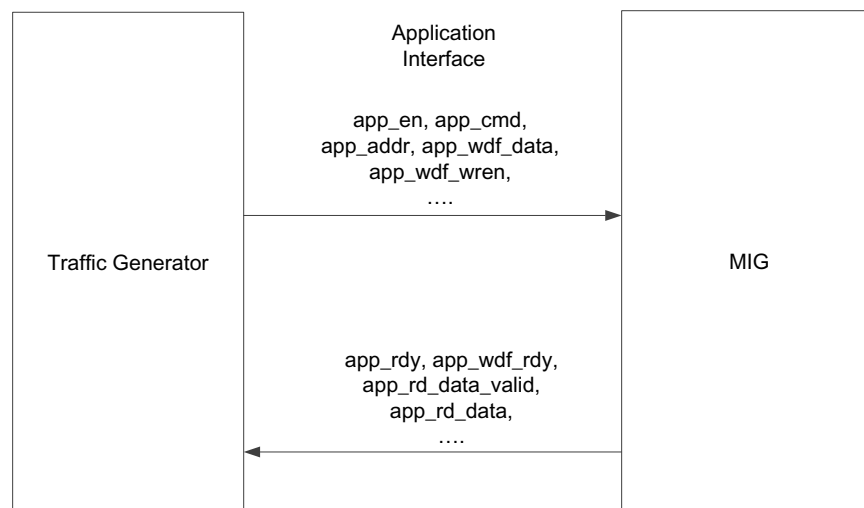


Figure 22-1: Traffic Generator and Application Interface

Two Traffic Generators are available to drive the memory design and they include:

- [Simple Traffic Generator](#)
- [Advanced Traffic Generator](#)

**Note:** The Advanced Traffic Generator is only available for the DDR3/DDR4 memory interfaces.

By default, Vivado® connects the memory design to the Simple Traffic Generator. You can choose to use the Advanced Traffic Generator by defining a switch "HW\_TG\_EN" in the `example_top.sv`. The Simple Traffic Generator is referred to as "STG" and the Advanced Traffic Generator is referred to as "ATG" for the remainder of this section.

## Simple Traffic Generator

MIG generates the STG modules as `example_tb` for native interface and `example_tb_phy` for PHY only interface. The STG native interface generates 100 writes and 100 reads. The STG PHY only interface generates 10 writes and 10 reads. Both address and data increment linearly. Data check is performed during reads. Data error is reported using the `compare_error` signal.

## Advanced Traffic Generator

The ATG is only supported for the user interface. When "HW\_TG\_EN" is defined, ATG is set to the default setting. To enable ATG (for both simulations and implementation), add "`define HW_TG_EN`" in the `example_top` module. The ATG default control connectivity in the example design created by Vivado is listed in [Table 22-1](#).

**Table 22-1: Default Traffic Generator Control Connection**

Signal	I/O	Width	Description
clk	I	1	Traffic Generator Clock
rst	I	1	Traffic Generator Reset
init_calib_complete	I	1	Calibration Complete
<b>General Control</b>			
vio_tg_start	I	1	Reserved signal. Tie to 1'b1.
vio_tg_rst	I	1	Reserved signal. Tie to 0.
vio_tg_restart	I	1	Reserved signal. Tie to 0.
vio_tg_pause	I	1	Reserved signal. Tie to 0.
vio_tg_err_chk_en	I	1	Reserved signal. Tie to 0.
vio_tg_err_clear	I	1	Reserved signal. Tie to 0.
vio_tg_err_clear_all	I	1	Reserved signal. Tie to 0.
vio_tg_err_continue	I	1	Reserved signal. Tie to 0.
<b>Instruction Table Programming</b>			
vio_tg_direct_instr_en	I	1	Reserved signal. Tie to 0.
vio_tg_instr_program_en	I	1	Reserved signal. Tie to 0.
vio_tg_instr_num	I	5	Reserved signal. Tie to 0.
vio_tg_instr_addr_mode	I	4	Reserved signal. Tie to 0.
vio_tg_instr_data_mode	I	4	Reserved signal. Tie to 0.
vio_tg_instr_rw_mode	I	4	Reserved signal. Tie to 0.
vio_tg_instr_rw_submode	I	2	Reserved signal. Tie to 0.

Table 22-1: Default Traffic Generator Control Connection (Cont'd)

Signal	I/O	Width	Description
vio_tg_instr_victim_mode	I	3	Reserved signal. Tie to 0.
vio_tg_instr_num_of_iter	I	32	Reserved signal. Tie to 0.
vio_tg_instr_m_nops_btw_n_burst_m	I	10	Reserved signal. Tie to 0.
vio_tg_instr_m_nops_btw_n_burst_n	I	32	Reserved signal. Tie to 0.
vio_tg_instr_nxt_instr	I	6	Reserved signal. Tie to 0.
<b>PRBS Data Seed Programming</b>			
vio_tg_seed_program_en	I	1	Reserved signal. Tie to 0.
vio_tg_seed_num	I	8	Reserved signal. Tie to 0.
vio_tg_seed_data	I	PRBS DATA WIDTH	Reserved signal. Tie to 0.
<b>Global Registers</b>			
vio_tg_glb_victim_bit	I	8	Reserved signal. Tie to 0.
vio_tg_glb_victim_aggr_delay	I	4	Reserved signal. Tie to 0.
vio_tg_glb_start_addr	I	APP_ADDR_WIDTH	Reserved signal. Tie to 0.
vio_tg_glb_qdriv_rw_submode	I	2	Reserved signal. Tie to 0.
<b>Traffic Generator Internal Signal</b>			
tg_qdriv_submode11_app_rd	I	1	Reserved signal. Tie to 0.

In default settings, the ATG performs memory writes follow by memory reads and data checks. Three types of patterns are generated sequentially:

1. PRBS23 data pattern
  - a. PRBS23 data pattern is used per data bit. Each data bit has a different default starting seed value.
  - b. Linear address pattern is used. Memory address space is walked through to cover full PRBS23 data pattern.
2. Hammer Zero pattern
  - a. Hammer Zero pattern is used for all data bits.
  - b. Linear address pattern is used. 1024 Traffic Generator commands are issued.
3. PRBS address pattern
  - a. PRBS23 data pattern is used per data bit. Each data bit has a different default starting seed value.
  - b. PRBS address pattern is used. 1024 Traffic Generator commands are issued.

The ATG repeats memory writes and reads on each of the two patterns infinitely. For simulations, ATG performs 1000 PRBS23 pattern followed by 1000 Hammer Zero pattern and 1000 PRBS address pattern.

You can check if there is a memory error in the Status register (`vio_tg_status_err_sticky_valid`) or if memory traffic stops (`vio_tg_status_watch_dog_hang`).

Upon the first memory error seen, the ATG logs the error address (`vio_tg_status_first_err_addr`) and bit mismatch (`vio_tg_status_first_err_bit`).

Table 22-2 shows the common Traffic Generator Status register output for debug.

**Table 22-2: Common Traffic Generator Status Register for Debug**

Signal	I/O	Width	Description
<b>Error Status Registers</b>			
<code>vio_tg_status_err_bit_valid</code>	O	1	Intermediate error detected. It is used as trigger to detect read error.
<code>vio_tg_status_err_bit</code>	O	APP_DATA_WIDTH	Intermediate error bit mismatch. Bitwise mismatch pattern.
<code>vio_tg_status_err_addr</code>	O	APP_ADDR_WIDTH	Intermediate error address. Address location of failed read.
<code>vio_tg_status_first_err_bit_valid</code>	O	1	If <code>vio_tg_err_chk_en</code> is set to 1, <code>first_err_bit_valid</code> is set to 1 when first mismatch error is encountered. This register is not overwritten until <code>vio_tg_err_clear</code> , <code>vio_tg_err_continue</code> , <code>vio_tg_restart</code> is triggered.
<code>vio_tg_status_first_err_bit</code>	O	APP_DATA_WIDTH	If <code>vio_tg_status_first_err_bit_valid</code> is set to 1, error mismatch bit pattern is stored in this register.
<code>vio_tg_status_first_err_addr</code>	O	APP_ADDR_WIDTH	If <code>vio_tg_status_first_err_bit_valid</code> is set to 1, error address is stored in this register.
<code>vio_tg_status_first_exp_bit_valid</code>	O	1	If <code>vio_tg_err_chk_en</code> is set to 1, this represents expected read data valid when first mismatch error is encountered.
<code>vio_tg_status_first_exp_bit</code>	O	APP_DATA_WIDTH	If <code>vio_tg_status_first_exp_bit_valid</code> is set to 1, expected read data is stored in this register.
<code>vio_tg_status_first_read_bit_valid</code>	O	1	If <code>vio_tg_err_chk_en</code> is set to 1, this represents read data valid when first mismatch error is encountered.
<code>vio_tg_status_first_read_bit</code>	O	APP_DATA_WIDTH	If <code>vio_tg_status_first_read_bit_valid</code> is set to 1, read data from memory is stored in this register.
<code>vio_tg_status_err_bit_sticky_valid</code>	O	1	Accumulated error mismatch valid over time. This register is reset by <code>vio_tg_err_clear</code> , <code>vio_tg_err_continue</code> , <code>vio_tg_restart</code> .
<code>vio_tg_status_err_bit_sticky</code>	O	APP_DATA_WIDTH	If <code>vio_tg_status_err_bit_sticky_valid</code> is set to 1, this represents accumulated error bit.

Table 22-2: Common Traffic Generator Status Register for Debug (Cont'd)

Signal	I/O	Width	Description
vio_tg_status_done	O	1	All traffic programmed completes. <b>Note:</b> If infinite loop is programmed, vio_tg_status_done does not assert.
vio_tg_status_watch_dog_hang	O	1	Watchdog hang. This register is set to 1 if there is no Read/Write command sent or no Read data return for a period of time (defined in tg_param.vh).

**Note:** See the corresponding section for application interface address/data format.

The ATG has a watchdog logic that checks if the ATG has any request sent to the application interface or the application interface has any read data return within N (parameter TG\_WATCH\_DOG\_MAX\_CNT) number of cycles. This provides information on whether memory traffic is running or stalled (because of reasons other than data mismatch).

# SECTION VI: MULTIPLE IP CORES

## Multiple IP Cores

## Chapter 23

# Multiple IP Cores

This chapter describes the specifications and pin rules for generating multiple IP cores.

---

## Creating a Design with Multiple IP Cores

The following steps must be followed to create a design with multiple IP cores:

1. Generate the target memory IP. If the design includes multiple instances of the same memory IP configuration, the IP only needs to be generated once. The same IP can be instantiated multiple times within the design.
  - If the IP shares the input `sys_clk`, select the **No Buffer** clocking option during IP generation. Memory IP that share `sys_clk` must be allocated in the same I/O column. For more information on Sharing of Input Clock Source, see the [Sharing of Input Clock Source](#) for a link of each controller section.
2. Create a wrapper file to instantiate the target memory IP cores.
3. Assign the pin locations for the Memory IP I/O signals. For more information on pin rules of the respective interface, see the [Sharing of a Bank](#) for a link of each controller section. Also, to learn more about the available MIG pin planning options, see the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [\[Ref 12\]](#).
4. Ensure the following specifications are followed.

---

## Sharing of a Bank

Pin rules of each controller must be followed during IP generation. For more information on pin rules of each interface, see the respective IP sections:

- [DDR3 Pin Rules in Chapter 4](#) and [DDR4 Pin Rules in Chapter 4](#)
- [QDR II+ Pin Rules in Chapter 11](#)
- [RLDRAM 3 Pin Rules in Chapter 18](#)

The same bank can be shared across multiple IP cores, but MIG allows sharing of banks across multiple IP cores if the rules for combining I/O standards in the same bank are followed.

For more information on the rules for combining I/O standards in the same bank, see the section "Rules for Combining I/O Standards in the Same Bank," in *UltraScale™ Architecture SelectIO™ Resources User Guide* (UG571) [Ref 4]. The DCI I/O banking rules are also captured in UG571.

---

## Sharing of Input Clock Source

One GCIO pin can be shared across multiple IP cores. There are certain rules that must be followed to share input clock source and you must perform a few manual changes in the wrapper files. For more information on Sharing of Input Clock Source, see the respective interfaces:

- [Sharing of Input Clock Source \(sys\\_clk\\_p\) in Chapter 4](#) (DDR3/DDR4)
- [Sharing of Input Clock Source \(sys\\_clk\\_p\) in Chapter 11](#) (QDR II+ SRAM)
- [Sharing of Input Clock Source \(sys\\_clk\\_p\) in Chapter 18](#) (RLDRAM 3)

---

## XSDB and dbg\_clk Changes

The `dbg_clk` port is an output from the MIG IP and it automatically connects to the `dbg_hub` logic by Vivado® during implementation. If multiple IP cores are instantiated in the same project, Vivado automatically connects the first IP `dbg_clk` to `dbg_hub`.

In the wrapper file in which multiple MIG IP cores are instantiated, do not connect any signal to `dbg_clk` and keep the port open during instantiation. Vivado takes care of the `dbg_clk` connection to the `dbg_hub`.

---

## PBLOCK and MMCM Constraints

To meet the timing designs, MIG generates the PBLOCK constraints. MIG takes care of the PBLOCK constraints for sharing of banks scenarios.

Similarly, MMCM must be allocated in the center bank of the memory I/Os selected banks. MIG generates the LOC constraints for MMCM such that there is no conflict if the same bank is shared across multiple IP cores.



## SECTION VII: APPENDICES

Migrating and Upgrading

Debugging

Additional Resources and Legal Notices

## Migrating and Upgrading

There are no port or parameter changes for upgrading the MIS core in the Vivado Design Suite at this time.

For general information on upgrading the MIG IP, see the “Upgrading IP” section in *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 8].

# Debugging

This appendix includes details about resources available on the Xilinx® Support website and debugging tools.



---

**TIP:** If the IP generation halts with an error, there might be a license issue. See [License Checkers in Chapter 1](#) for more details.

---

---

## Finding Help on Xilinx.com

To help in the design and debug process when using the MIS, the [Xilinx Support web page](#) contains key resources such as product documentation, release notes, answer records, information about known issues, and links for opening a Technical Support WebCase.

### Documentation

This product guide is the main document associated with the MIS. This guide, along with documentation related to all products that aid in the design process, can be found on the [Xilinx Support web page](#) or by using the Xilinx Documentation Navigator.

Download the Xilinx Documentation Navigator from the Design Tools tab on the [Downloads page](#). For more information about this tool and the features available, open the online help after installation.

### Solution Centers

See the [Xilinx Solution Centers](#) for support on devices, software tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips.

The Solution Center specific to the MIS core is located at [Xilinx MIG Solution Center](#).

### Answer Records

Answer Records include information about commonly encountered problems, helpful information on how to resolve these problems, and any known issues with a Xilinx product.

Answer Records are created and maintained daily ensuring that users have access to the most accurate information available.

Answer Records for this core can be located by using the Search Support box on the main [Xilinx support web page](#). To maximize your search results, use proper keywords such as:

- Product name
- Tool message(s)
- Summary of the issue encountered

A filter search is available after results are returned to further target the results.

### Master Answer Record for the MIS

AR: [58435](#)

## Contacting Technical Support

Xilinx provides technical support at [Xilinx support web page](#) for this LogiCORE™ IP product when used as described in the product documentation. Xilinx cannot guarantee timing, functionality, or support of product if implemented in devices that are not defined in the documentation, if customized beyond that allowed in the product documentation, or if changes are made to any section of the design labeled DO NOT MODIFY.

To contact Xilinx Technical Support:

1. Navigate to [Xilinx support web page](#).
2. Open a WebCase by selecting the [WebCase](#) link located under Additional Resources.

When opening a WebCase, include:

- Target FPGA including package and speed grade.
- All applicable Xilinx Design Tools and simulator software versions.
- Additional files based on the specific issue might also be required. See the relevant sections in this debug guide for guidelines about which file(s) to include with the WebCase.

**Note:** Access to WebCase is not available in all cases. Log in to the WebCase tool to see your specific support options.

---

## Debug Tools

There are many tools available to address MIG design issues. It is important to know which tools are useful for debugging various situations.

### Vivado Design Suite Debug Feature

The Vivado® Design Suite debug feature inserts logic analyzer and virtual I/O cores directly into your design. The debug feature also allows you to set trigger conditions to capture application and integrated block port signals in hardware. Captured signals can then be analyzed. This feature in the Vivado IDE is used for logic debugging and validation of a design running in Xilinx devices.

The Vivado logic analyzer is used with the logic debug IP cores, including:

- ILA 2.0 (and later versions)
- VIO 2.0 (and later versions)

See the *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [\[Ref 14\]](#).

---

## Hardware Debug

Hardware issues can range from link bring-up to problems seen after hours of testing. This section provides debug steps for common issues. Vivado Lab Edition is a valuable resource to use in hardware debug. The signal names mentioned in the following individual sections can be probed using Vivado Lab Edition for debugging the specific problems.

### General Checks

Ensure that all the timing constraints for the core were properly incorporated from the example design and that all constraints were met during implementation.

- If using MMCMs in the design, ensure that all MMCMs have obtained lock by monitoring the `locked` port.
- If your outputs go to 0, check your licensing.
- If you are experiencing issues with DDR3 or DDR4 interfaces, visit the Xilinx AR: [60305](#). Run the following Tcl Console commands in the Vivado when connected to the hardware:

```
refresh_hw_device [lindex [get_hw_devices] 0]

report_property [lindex [get_hw_migs] 0]
```

- For designs with more than one memory core, use the Tcl Console command `report_debug_core` to determine which index is the core of interest. The index number above is shown as 0.
- Copy all of the data reported and submit it as part of a WebCase. For more information on opening a WebCase, see the [Contacting Technical Support, page 300](#).

# Additional Resources and Legal Notices

---

## Xilinx Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see [Xilinx Support](#).

---

## References

These documents provide supplemental material useful with this product guide:

1. JESD79-3F, *DDR3 SDRAM Standard* and JESD79-4, *DDR4 SDRAM Standard*, JEDEC® Solid State Technology Association
2. *Kintex® UltraScale™ Architecture Data Sheet: DC and AC Switching Characteristics* ([DS892](#))
3. *UltraScale Architecture Clocking Resources User Guide* ([UG572](#))
4. *UltraScale Architecture SelectIO™ Resources User Guide* ([UG571](#))
5. *UltraScale Architecture PCB Design and Pin Planning User Guide* ([UG583](#))
6. [ARM® AMBA® Specifications](#)
7. *Vivado® Design Suite User Guide: Designing IP Subsystems using IP Integrator* ([UG994](#))
8. *Vivado Design Suite User Guide: Designing with IP* ([UG896](#))
9. *Vivado Design Suite User Guide: Getting Started* ([UG910](#))
10. *Vivado Design Suite User Guide: Logic Simulation* ([UG900](#))
11. *Vivado Design Suite User Guide: Implementation* ([UG904](#))
12. *Vivado Design Suite User Guide: I/O and Clock Planning* ([UG899](#))
13. *Vivado Design Suite User Guide: Release Notes, Installation, and Licensing* ([UG973](#))
14. *Vivado Design Suite User Guide: Programming and Debugging* ([UG908](#))

## Revision History

The following table shows the revision history for this document.

Date	Version	Revision
06/24/2015	7.1	<ul style="list-style-type: none"> <li>Updated all Resource Utilization sections.</li> <li>Added clocking reference in all Requirements sections.</li> <li>Updated description in all Resets section.</li> <li>Updated all Clocking sections.</li> <li>Updated all CLOCK_DEDICATED_ROUTE Constraints and BUFG Instantiation sections.</li> </ul>
		<b>DDR3/DDR4</b> <ul style="list-style-type: none"> <li>Added x4 devices are not supported in AXI note in Feature Summary section.</li> <li>Updated Fig. 3-6: PHY Overall Initialization and Calibration Sequence.</li> <li>Added Table 3-4: Pre-Calibration XSDB Status Signal Description.</li> <li>Updated Table 3-5: XSDB Status Signal Description</li> <li>Added Table 3-6: Post-Calibration XSDB Status Signal Description.</li> <li>Updated Read per-bit Deskew description in Table 3-6: Error Signal Descriptions.</li> <li>Updated description in Write DQS-to-DQ Centering section.</li> <li>Added Read DQS Centering (Complex) and Write DQS-to-DQ Centering (Complex) sections.</li> <li>Added Notes to Write DQS-to-DQ, Write DQS-to-DM, Write DQS-to-DQ Centering (Complex), Read <math>V_{REF}</math>, and Read DQS Centering (Complex).</li> <li>Added Read <math>V_{REF}</math> and Write <math>V_{REF}</math> Calibrations section.</li> <li>Updated letter b and c descriptions in DDR3 Pin Rules section.</li> <li>Updated AXI4-Lite Slave Control/Status Register Map Detailed Descriptions.</li> <li>Added description in Project-Based Simulation Flow Using Vivado Simulator section.</li> </ul>
		<b>QDR II+</b> <ul style="list-style-type: none"> <li>Added HSTL_I I/O standard support in Feature Summary.</li> <li>Added description to the Memory Initialization bullet in Overview section.</li> </ul> <b>RLDRAM 3</b> <ul style="list-style-type: none"> <li>Updated description in Required Constraints section.</li> <li>Updated Fig. 17-4: PHY Overall Initialization and Calibration Sequence.</li> <li>Updated description d. in RLDRAM 3 Pin Rules.</li> </ul> <b>Traffic Generator</b> <ul style="list-style-type: none"> <li>Updated Advanced Traffic Generator section.</li> </ul> <b>Debugging Appendix</b> <ul style="list-style-type: none"> <li>Added AR: 60305 in General Checks section.</li> </ul>



Date	Version	Revision
04/01/2015	7.0	<ul style="list-style-type: none"> <li>Updated Supported User Interface and added #3 footnote in IP Facts table.</li> <li>Updated Application Interface description in the Overview chapter.</li> <li>Updated descriptions and added BACKBONE description in all Clocking sections.</li> <li>Added sys_rst and dbg_clk references throughout book.</li> <li>Added Simulation Flow and Simulation Speed to all sections.</li> <li>Added Project-Based Simulation Flow Using Vivado Simulator to all sections.</li> <li>Added CLOCK_DEDICATED_ROUTE Constraints and BUFG Instantiation to all sections.</li> </ul> <p><b>DDR3/DDR4</b></p> <ul style="list-style-type: none"> <li>Updated Fig. 1-1: UltraScale Architecture-Based FPGAs Memory Interface Solution.</li> <li>Updated Feature Summary section.</li> <li>Updated Memory Controller section.</li> <li>Updated Group Machines section.</li> <li>Updated DQS section.</li> <li>Updated parameters in Write Leveling section.</li> <li>Updated and added Important note in Read DQS Centering section.</li> <li>Updated Read Leveling Multi-Rank Adjustment, Multi-Rank Adjustments and Checks, and added Write Latency Multi-Rank Check.</li> <li>Updated Write Per-bit Deskew section.</li> <li>Updated Write DQS-to-DM section.</li> <li>Updated Table 3-5: Error Signal Descriptions.</li> <li>Updated Table 3-6: Examples of DQS Gate Multi-Rank Adjustment (2 Ranks).</li> <li>Updated DDR3 and DDR4 Pin Rules sections.</li> </ul> <hr/> <ul style="list-style-type: none"> <li>Added Pin Mapping for x4 RDIMMs.</li> <li>Added app_ref_req, app_ref_ack, app_zq_req, and app_zq_ack in Table 4-7: User Interface.</li> <li>Updated Write Path section.</li> <li>Added Performance section.</li> <li>Added descriptions for app_ref_req, app_ref_ack, app_zq_req, and app_zq_ack.</li> <li>Added Maintenance Commands section.</li> <li>Updated Table 4-16: AXI4 Slave Interface Parameters.</li> <li>Added dbg_clk to Table 4-17: AXI4 Slave Interface Signals.</li> <li>Updated Time Division Multiplexing (TDM), Round-Robin, and Read Priority (RD_PRI_REG) sections.</li> </ul>

Date	Version	Revision
Continued		<ul style="list-style-type: none"> <li>• Updated Table 4-76: Read Data to Table 4-77: PHY Only Parameters.</li> <li>• Updated to 11 writes in Multiple Writes and Reads with Same Address to Page Wrap During Writes sections.</li> <li>• Added Minimum Write CAS Command Spacing and System Considerations for CAS Command Spacing sections.</li> <li>• Updated the Design Flow Steps chapter.</li> </ul> <p><b>QDR II+</b></p> <ul style="list-style-type: none"> <li>• Updated Feature Summary.</li> </ul> <p><b>RLDRAM 3</b></p> <ul style="list-style-type: none"> <li>• Added User Interface Allocation section.</li> <li>• Added User Address Bit Allocation Based on RLDRAM 3 Configuration section.</li> <li>• Added description to Interfacing with the Core through the User Interface section.</li> </ul> <p><b>Traffic Generator</b></p> <ul style="list-style-type: none"> <li>• Added Traffic Generator section.</li> </ul> <p><b>Multiple IP</b></p> <ul style="list-style-type: none"> <li>• Added Multiple IP section.</li> </ul> <p><b>Migrating and Upgrading Appendix</b></p> <ul style="list-style-type: none"> <li>• Added link to UG973 and description in Migrating and Upgrading chapter.</li> </ul> <p><b>Debugging Appendix</b></p> <ul style="list-style-type: none"> <li>• Added description in General Checks section.</li> </ul>

Date	Version	Revision
11/19/2014	6.1	<b>QDR II+</b> <ul style="list-style-type: none"> <li>Added interface calibration in Feature Summary section.</li> <li>Updated description #2 in Sharing of Input Clock Source (sys_clk_p) section.</li> <li>Added read data pins description and cross-ref to system clock pins description in QDR II+ Pin Rules section.</li> <li>Added vrp description in QDR II+ Pin Rules section.</li> <li>Updated User Parameters table.</li> </ul> Updated GUIs in Example Design chapter.
		<b>DDR3/DDR4</b> <ul style="list-style-type: none"> <li>Updated Fig. 1-1: UltraScale Architecture-Based FPGAs Memory Interface Solution.</li> <li>Added interface calibration in Feature Summary section.</li> <li>Updated RIU code in Overall PHY Architecture section.</li> <li>Updated description #2 in Sharing of Input Clock Source (sys_clk_p) section.</li> <li>Added ECC description in Datapath section and ECC section.</li> <li>Updated resetn, input clock description, and added x4 Part Contained in One Bank tables in DDR3 and DDR4 Pin Rules sections.</li> <li>Added app_raw_not_ecc in Table 4-5: User Interface.</li> <li>Updated descriptions in app_cmd[2:0] section.</li> <li>Updated Fig. 4-2 and Fig. 4-6 to Fig. 4-8.</li> <li>Added examples for DRAM clock in Write Path section.</li> <li>Added PHY Only section in Protocol Description.</li> <li>Updated R<sub>TT</sub> (nominal)-ODT default values in Table 5:1: Vivado IDE Parameter to User Parameter Relationship.</li> <li>Updated GUIs in Customizing and Generating the Core section.</li> <li>Updated User Parameters table.</li> <li>Updated GUIs in Example Design chapter.</li> </ul>
		<b>RLDRAM 3</b> <ul style="list-style-type: none"> <li>Added interface calibration in Feature Summary section.</li> <li>Updated Table 15-1: Supported Configurations and removed support for Read Latency in Feature Summary.</li> <li>Added CMD_PER_CLK description in Memory Controller section.</li> <li>Updated description #2 in Sharing of Input Clock Source (sys_clk_p) section.</li> <li>Updated input clock description in RLDRAM 3 Pin Rules section.</li> <li>Added note in Interfacing with the Core through the User Interface section.</li> <li>Updated Fig. 18-2: Multiple Commands for user_cmd Signal.</li> <li>Updated User Parameters table.</li> <li>Updated GUIs in Example Design chapter.</li> <li>Updated description in Simulating the Example Design (Designs with Standard User Interface) section.</li> </ul>

Date	Version	Revision
10/01/2014	6.0	<b>DDR3/DDR4</b> <ul style="list-style-type: none"> <li>• Updated Standards section.</li> <li>• Updated Feature Summary section.</li> <li>• Updated description in Memory Initialization and Calibration Sequence section.</li> <li>• Updated Overall PHY Architecture section.</li> <li>• Updated Fig. 3-4: PHY Overall Initialization and Calibration Sequence.</li> <li>• Added new calibration status descriptions in Memory Initialization and Calibration Sequence section.</li> <li>• Added DQS Gate, Write Leveling, Read Leveling, Read Sanity Check, Write DQS-to-DQ, Write Latency Calibration, Write/Read Sanity Check, Write DQS-to-DM, and Multi-Rank Adjustment sections.</li> <li>• Updated DDR3/DDR4 Pin Rules section.</li> <li>• Added AXI4 Slave Interface in Protocol Description section.</li> <li>• Added Multiple IP Cores and Sharing of Input Clock Source in Clocking section.</li> <li>• Removed Special Designation column in Table 4-1: 16-Bit Interface Contained in One Bank and Table 4-2: 32-Bit Interface Contained in Two Banks.</li> <li>• Added app_autoprecharge to Table 4-3: User Interface.</li> <li>• Added app_autoprecharge section.</li> <li>• Updated app_rdy section.</li> <li>• Updated ref_req and zq_req sections.</li> <li>• Updated Table 5-1: Vivado IDE Parameter to User Parameter Relationship.</li> <li>• Updated note description in Required Constraints section.</li> <li>• Updated description in Simulation section.</li> <li>• Updated GUIs in Example Design chapter.</li> </ul>

Date	Version	Revision
Continued		<b>QDR II+</b> <ul style="list-style-type: none"> <li>• Updated Feature Summary section.</li> <li>• Updated Table 9-1: Device Utilization – Kintex UltraScale FPGAs.</li> <li>• Updated Fig. 10-3: PHY Overall Initialization and Calibration Sequence.</li> <li>• Updated MicroBlaze description in Overall PHY Architecture section.</li> <li>• Updated Memory Initialization and Calibration Sequence section.</li> <li>• Updated Resets section.</li> <li>• Deleted Special Designation column in Table 11-1: 18-Bit QDR II+ Interface Contained in Two Banks.</li> <li>• Added Multiple IP Cores and Sharing of Input Clock Source in Clocking section.</li> <li>• Updated Protocol Description section.</li> <li>• Updated Simulation section.</li> <li>• Updated description in Simulating the Example Design (Designs with Standard User Interface) section.</li> <li>• Updated GUIs in Example Design chapter.</li> </ul>
		<b>RLDRAM 3</b> <ul style="list-style-type: none"> <li>• Added Configuration table in Feature Summary section.</li> <li>• Updated Memory Initialization bullet in Overview chapter.</li> <li>• Added description to burst support in Feature Summary section.</li> <li>• Updated Table 16-1: Device Utilization – Kintex UltraScale FPGAs.</li> <li>• Updated Memory Controller section.</li> <li>• Updated Overall PHY Architecture section.</li> <li>• Updated Memory Initialization and Calibration Sequence section.</li> <li>• Added Multiple IP Cores and Sharing of Input Clock Source in Clocking section.</li> <li>• Added data mask description to RLDRAM 3 Pin Rules section.</li> <li>• Updated GUIs in Example Design chapter.</li> </ul>
		<b>Appendix</b> Added Migrating Appendix.

Date	Version	Revision
06/04/2014	5.0	<ul style="list-style-type: none"> <li>Removed PCB sections and added link to UG583.</li> <li>Global replace BUFGCE to BUFGCE_DIV.</li> </ul> <b>DDR3/DDR4</b> <ul style="list-style-type: none"> <li>Updated CAS cycle description in DDR3 Feature Summary.</li> <li>Updated descriptions in Native Interface section.</li> <li>Updated Control Path section.</li> <li>Updated Read and Write Coalescing section.</li> <li>Updated Reordering section.</li> <li>Updated DDR4 x16 parts in Group Machines section.</li> <li>Updated Fig. 3-3: PHY Block Diagram.</li> <li>Updated Table 3-1: PHY Modules.</li> <li>Updated module names in Overall PHY Architecture section.</li> <li>Updated Fig. 3-4: PHY Overall Initialization and Calibration Sequence.</li> <li>Added description to Memory Initialization and Calibration Sequence section.</li> <li>Added SSI rule in Clocking section.</li> <li>Added SSI rule and updated Address and ck descriptions in DDR3/DDR4 Pin Rules sections.</li> <li>Added Important Note for calibration stage in DDR3/DDR4 Pinout Examples sections.</li> <li>Updated signal descriptions in Table 4-3: User Interface.</li> <li>Added new content in app_addr[ADDR_WIDTH – 1:0] section.</li> <li>Updated Write Path section.</li> <li>Updated Native Interface section.</li> <li>Added Important Note relating to Data Mask in Controller Options section.</li> <li>Added PHY Only section.</li> <li>Updated Fig. 5-1 to 5-8 in Customizing and Generating the Core section.</li> <li>Added User Parameters section in Design Flow Steps chapter.</li> <li>Updated I/O Standard and Placement section.</li> <li>Added Synplify Black Box Testing section in Example Design chapter.</li> </ul>

Date	Version	Revision
Continued		<b>QDR II+</b> <ul style="list-style-type: none"> <li>Updated Read Latency in Feature Summary section.</li> <li>Updated Fig. 10-2: PHY Block Diagram and Table 17-1: PHY Modules.</li> <li>Updated Table 11-2: User Interface.</li> <li>Added SSI rule in Clocking section.</li> <li>Added Important Note for calibration stage in QDR II+ Pinout Examples section.</li> <li>Added SSI rule in QDR II+ Pin Rules section.</li> <li>Updated I/O Standard and Placement section.</li> <li>Added User Parameters section in Design Flow Steps chapter.</li> <li>Updated the descriptions in Simulating the Example Design (Designs with Standard User Interface) section.</li> <li>Added Synplify Black Box Testing section in Example Design chapter.</li> </ul>
		<b>RLDRAM 3</b> <ul style="list-style-type: none"> <li>Added 18 bits in Feature Summary section.</li> <li>Updated Fig. 17-4: PHY Block Diagram.</li> <li>Updated module names in Table 17-1: PHY Modules.</li> <li>Updated module names in Overall PHY Architecture section.</li> <li>Added SSI rule in Clocking section.</li> <li>Updated c) and d) descriptions and added SSI rule in RLDRAM 3 Pin Rules section.</li> <li>Updated Table 18-2: User Interface Request Signals.</li> <li>Updated Fig. 18-2: Multiple Commands for user_cmd Signal.</li> <li>Added Important Note for calibration stage in RLDRAM 3 Pinout Examples section.</li> <li>Updated I/O Standard and Placement section.</li> <li>Added User Parameters section in Design Flow Steps chapter.</li> <li>Updated Test Bench chapter.</li> </ul>

Date	Version	Revision
04/02/2014	5.0	<ul style="list-style-type: none"> <li>Added Verilog Test Bench in IP Facts table.</li> </ul> <b>DDR3/DDR4</b> <ul style="list-style-type: none"> <li>Added Overview chapter.</li> <li>Updated component support to 80 bits in Feature Summary section.</li> <li>Updated DDR Device Utilization tables.</li> <li>Updated DDR Clocking section.</li> <li>Updated x4 DRAM to Four Component DRAM Configuration in Designing with the Core chapter.</li> <li>Updated Important note in PCB Guidelines for DDR3 and DDR4 Overview sections.</li> <li>Updated Important note in Reference Stack-Up for DDR3 and DDR4 sections.</li> <li>Updated trace length descriptions in DDR3 and DDR4 sections.</li> <li>Added <math>V_{TT}</math> Terminations guideline in Generic Routing Guideline for DDR3 and DDR4 sections.</li> <li>Removed Limitations section.</li> <li>Added <math>V_{REF}</math> note in Required Constraints section.</li> <li>Updated new figures in Design Flow Steps chapter.</li> <li>Added new descriptions in Example Design chapter.</li> <li>Added new description in Test Bench chapter.</li> </ul> <b>QDR II+ SRAM</b> <ul style="list-style-type: none"> <li>Added new QDR II+ section.</li> </ul> <b>RLDRAM 3</b> <ul style="list-style-type: none"> <li>Added Overview chapter.</li> <li>Added new Clocking section.</li> <li>Added new descriptions in Example Design chapter.</li> </ul> <b>Appendix</b> <ul style="list-style-type: none"> <li>Updated Debug Appendix.</li> </ul>
12/18/2013	4.2	Initial Xilinx release.



---

## Please Read: Important Legal Notices

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

© Copyright 2013–2015 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, UltraScale, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.