



External Memory Interface Handbook Volume 2

Section II. Board Planning



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

EMI_PLAN_BOARD-3.0

Document last updated for Altera Complete Design Suite version:
Document publication date:

11.0
June 2011



[Subscribe](#)

© 2011 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Chapter 1. DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines

Leveling and Dynamic ODT	1-2
Read and Write Leveling	1-2
Calibrated Output Impedance and ODT	1-4
Dynamic ODT	1-4
Dynamic OCT in Stratix III and Stratix IV Devices	1-5
Dynamic OCT in Stratix V Devices	1-7
Board Termination for DDR2 SDRAM	1-7
External Parallel Termination	1-8
On-Chip Termination	1-9
Recommended Termination Schemes	1-10
Dynamic On-Chip Termination	1-12
FPGA Writing to Memory	1-13
FPGA Reading from Memory	1-16
On-Chip Termination (Non-Dynamic)	1-17
Class II External Parallel Termination	1-20
FPGA Writing to Memory	1-20
FPGA Reading from Memory	1-23
Class I External Parallel Termination	1-25
FPGA Writing to Memory	1-26
FPGA Reading from Memory	1-27
Class I Termination Using ODT	1-29
FPGA Writing to Memory	1-29
FPGA Reading from Memory	1-31
No-Parallel Termination	1-31
FPGA Writing to Memory	1-31
FPGA Reading from Memory	1-33
Board Termination for DDR3 SDRAM	1-36
Single-Rank DDR3 SDRAM Unbuffered DIMM	1-36
DQS, DQ, and DM for DDR3 SDRAM UDIMM	1-38
Memory Clocks for DDR3 SDRAM UDIMM	1-40
Commands and Addresses for DDR3 SDRAM UDIMM	1-42
Stratix III, Stratix IV, and Stratix V FPGAs	1-43
DQS, DQ, and DM for Stratix III, Stratix IV, and Stratix V FPGA	1-43
Memory Clocks for Stratix III, Stratix IV, and Stratix V FPGA	1-44
Commands and Addresses for Stratix III and Stratix IV FPGA	1-44
Multi-Rank DDR3 SDRAM Unbuffered DIMM	1-44
DDR3 SDRAM Registered DIMM	1-46
DDR3 SDRAM Components With Leveling	1-46
DDR3 SDRAM Components With or Without Leveling	1-46
Stratix III, Stratix IV, and Stratix V FPGAs	1-50
Drive Strength	1-51
How Strong is Strong Enough?	1-52
System Loading	1-53
Component Versus DIMM	1-53
FPGA Writing to Memory	1-53
FPGA Reading from Memory	1-55
Single- Versus Dual-Rank DIMM	1-56
Single DIMM Versus Multiple DIMMs	1-58

Design Layout Guidelines	1–58
Layout Guidelines for DDR2 SDRAM Interface	1–58
Layout Guidelines for DDR3 SDRAM Interface	1–61
Layout Guidelines for DDR3 SDRAM Wide Interface (>72 bits)	1–66
Fly-By Network Design for Clock, Command, and Address Signals	1–66

Chapter 2. Dual-DIMM DDR2 and DDR3 SDRAM Interface Termination, Drive Strength, Loading, and Board Layout Guidelines

DDR2 SDRAM	2–1
Stratix II High Speed Board	2–2
Overview of ODT Control	2–3
DIMM Configuration	2–5
Dual-DIMM Memory Interface with Slot 1 Populated	2–5
FPGA Writing to Memory	2–5
Write to Memory Using an ODT Setting of 150 Ω	2–5
Reading from Memory	2–7
Dual-DIMM with Slot 2 Populated	2–8
FPGA Writing to Memory	2–8
Write to Memory Using an ODT Setting of 150 Ω	2–9
Reading from Memory	2–10
Dual-DIMM Memory Interface with Both Slot 1 and Slot 2 Populated	2–11
FPGA Writing to Memory	2–12
Write to Memory in Slot 1 Using an ODT Setting of 75- Ω	2–12
Write to Memory in Slot 2 Using an ODT Setting of 75- Ω	2–14
Reading From Memory	2–15
FPGA OCT Features	2–20
Stratix III and Stratix IV Devices	2–20
Arria II GX Devices	2–20
Dual-DIMM DDR2 Clock, Address, and Command Termination and Topology	2–21
Address and Command Signals	2–21
Control Group Signals	2–22
Clock Group Signals	2–22
DDR3 SDRAM	2–23
Comparison of DDR3 and DDR2 DQ and DQS ODT Features and Topology	2–23
Dual-DIMM DDR3 Clock, Address, and Command Termination and Topology	2–24
Address and Command Signals	2–24
Control Group Signals	2–24
Clock Group Signals	2–24
Write to Memory in Slot 1 Using an ODT Setting of 75 Ω With One Slot Populated	2–25
Write to Memory in Slot 2 Using an ODT Setting of 75 Ω With One Slot Populated	2–26
Write to Memory in Slot 1 Using an ODT Setting of 150 Ω With Both Slots Populated	2–27
Write to Memory in Slot 2 Using an ODT Setting of 150 Ω With Both Slots Populated	2–28
Read from Memory in Slot 1 Using an ODT Setting of 150 Ω on Slot 2 with Both Slots Populated ..	2–29
Read From Memory in Slot 2 Using an ODT Setting of 150 Ω on Slot 1 With Both Slots Populated ..	2–30
Conclusion	2–30
References	2–31

Chapter 3. QDR II SRAM Interface Termination and Layout Guidelines

I/O Standards	3–1
QDR II SRAM Configurations	3–1
Signal Terminations	3–4
Output from the FPGA to the QDR II SRAM Component	3–5
Input to the FPGA from the QDR II SRAM Component	3–13

Termination Schemes	3-17
PCB Layout Guidelines	3-19

Chapter 4. RLDRAM II Interface Termination and Layout Guidelines

I/O Standards	4-1
RLDRAM II Configurations	4-1
Signal Terminations	4-3
Outputs from the FPGA to the RLDRAM II Component	4-5
Input to the FPGA from the RLDRAM II Component	4-9
Termination Schemes	4-10
PCB Layout Guidelines	4-11

Chapter 5. Power Estimation Methods for External Memory Interface Designs

Additional Information

Document Revision History	Info-1
How to Contact Altera	Info-1
Typographic Conventions	Info-2

This chapter provides guidelines on how to improve the signal integrity of your system and layout guidelines to help you successfully implement a DDR2 or DDR3 SDRAM interface on your system.


DDR3 SDRAM is the third generation of the DDR SDRAM family, and offers improved power, higher data bandwidth, and enhanced signal quality with multiple on-die termination (ODT) selection and output driver impedance control while maintaining partial backward compatibility with the existing DDR2 SDRAM standard.

This chapter focuses on the following key factors that affect signal quality at the receiver:

- Leveling and dynamic ODT
- Proper use of termination
- Output driver drive strength setting
- Loading at the receiver
- Layout guidelines

As memory interface performance increases, board designers must pay closer attention to the quality of the signal seen at the receiver because poorly transmitted signals can dramatically reduce the overall data-valid margin at the receiver. [Figure 1-1](#) shows the differences between an ideal and real signal seen by the receiver.

Figure 1-1. Ideal and Real Signal at the Receiver



In addition, this chapter compares various types of termination schemes, and their effects on the signal quality on the receiver. It also discusses the proper drive strength setting on the FPGA to optimize the signal integrity at the receiver, and the effects of different loading types, such as components versus DIMM configuration, on signal quality.

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

Leveling and Dynamic ODT

DDR3 SDRAM DIMMs, as specified by JEDEC, always use a fly-by topology for the address, command, and clock signals. This standard DDR3 SDRAM topology requires the use of Altera DDR3 SDRAM Controller with UniPHY or ALTMEMPHY with read and write leveling.

Altera recommends that for full DDR3 SDRAM compatibility when using discrete DDR3 SDRAM components, you should mimic the JEDEC DDR3 UDIMM fly-by topology on your custom printed circuit boards (PCB).




Arria II devices do not support DDR3 SDRAM with read or write leveling, so these devices do not support standard DDR3 SDRAM DIMMs or DDR3 SDRAM components using the standard DDR3 SDRAM fly-by address, command, and clock layout topology.

Read and Write Leveling

One major difference between DDR2 and DDR3 SDRAM is the use of leveling. To improve signal integrity and support higher frequency operations, the JEDEC committee defined a fly-by termination scheme used with clocks, and command and address bus signals. Fly-by topology reduces simultaneous switching noise (SSN) by deliberately causing flight-time skew between the data and strobes at every DRAM as the clock, address, and command signals traverse the DIMM (Figure 1-2).

Figure 1-2. DDR3 DIMM Fly-By Topology Requiring Write Leveling




The flight-time skew caused by the fly-by topology led the JEDEC committee to introduce the write leveling feature on the DDR3 SDRAMs; thus requiring controllers to compensate for this skew by adjusting the timing per byte lane.

During a write, DQS groups launch at separate times to coincide with a clock arriving at components on the DIMM, and must meet the timing parameter between the memory clock and DQS defined as t_{DQS} of $\pm 0.25 t_{CK}$.

During the read operation, the memory controller must compensate for the delays introduced by the fly-by topology. The Stratix III, Stratix IV, and Stratix V FPGAs have alignment and synchronization registers built in the I/O element (IOE) to properly capture the data.

Figure 1-3 shows two DQS groups returning from the DIMM for the same read command.

Figure 1-3. DDR3 DIMM Fly-By Topology Requiring Read Leveling



Calibrated Output Impedance and ODT

In DDR2 SDRAM, there are only two drive strength settings, full or reduced, which correspond to the output impedance of 18 Ω and 40 Ω , respectively. These output drive strength settings are static settings and are not calibrated; as a result, the output impedance varies as the voltage and temperature drifts.

The DDR3 SDRAM uses a programmable impedance output buffer. Currently, there are two drive strength settings, 34 Ω and 40 Ω . The 40- Ω drive strength setting is currently a reserved specification defined by JEDEC, but available on the DDR3 SDRAM, as offered by some memory vendors. Refer to the datasheet of the respective memory vendors for more information about the output impedance setting. You select the drive strength settings by programming the memory mode register defined by mode register 1 (MR1). To calibrate output driver impedance, an external precision resistor, RZQ, connects the ZQ pin and VSSQ. The value of this resistor must be $240\ \Omega \pm 1\%$.

If you are using a DDR3 SDRAM DIMM, RZQ is soldered on the DIMM so you do not need to layout your board to account for it. Output impedance is set during initialization. To calibrate output driver impedance after power-up, the DDR3 SDRAM needs a calibration command that is part of the initialization and reset procedure and is updated periodically when the controller issues a calibration command.

In addition to calibrated output impedance, the DDR3 SDRAM also supports calibrated parallel ODT through the same external precision resistor, RZQ, which is possible by using a merged output driver structure in the DDR3 SDRAM, which also helps to improve pin capacitance in the DQ and DQS pins. The ODT values supported in DDR3 SDRAM are 20 Ω , 30 Ω , 40 Ω , 60 Ω , and 120 Ω , assuming that RZQ is 240 Ω .

In DDR3 SDRAM, there are two commands related to the calibration of the output driver impedance and ODT. The controller often uses the first calibration command, ZQ CALIBRATION LONG (ZQCL), at initial power-up or when the DDR3 SDRAM is in a reset condition. This command calibrates the output driver impedance and ODT to the initial temperature and voltage condition, and compensates for any process variation due to manufacturing. If the controller issues the ZQCL command at initialization or reset, it takes 512 memory clock cycles to complete; otherwise, it requires 256 memory clock cycles to complete. The controller uses the second calibration command, ZQ CALIBRATION SHORT (ZQCS) during regular operation to track any variation in temperature or voltage. The ZQCS command takes 64 memory clock cycles to complete. Use the ZQCL command any time there is more impedance error than can be corrected with a ZQCS command.

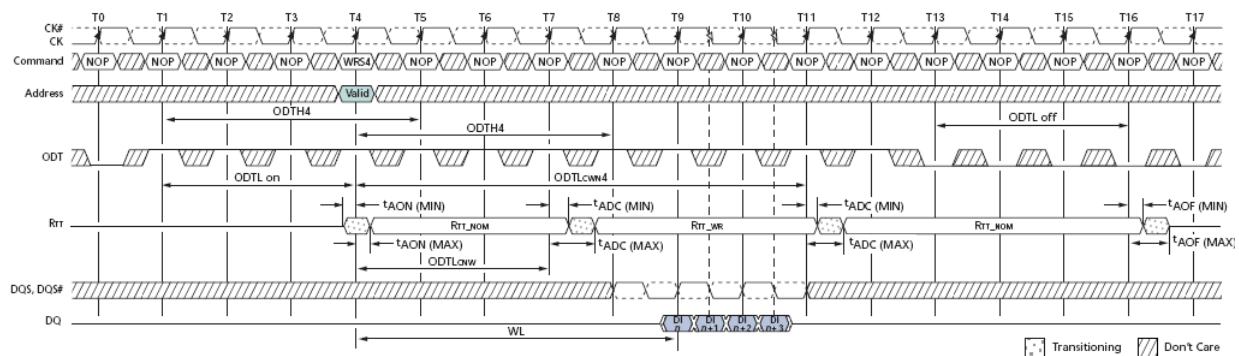
For more information about using ZQ Calibration in DDR3 SDRAM, refer to the application note by Micron, *TN-41-02 DDR3 ZQ Calibration*.

Dynamic ODT

Dynamic ODT is a new feature in DDR3 SDRAM, and not available in DDR2 SDRAM. Dynamic ODT can change the ODT setting without issuing a mode register set (MRS) command. When you enable dynamic ODT, and there is no write operation, the DDR3 SDRAM terminates to a termination setting of `RTT_NORM`; when there is a write operation, the DDR3 SDRAM terminates to a setting of `RTT_WR`. You can preset the values of `RTT_NORM` and `RTT_WR` by programming the mode registers, MR1 and MR2.

Figure 1-4 shows the behavior of ODT when you enable dynamic ODT.

Figure 1-4. Dynamic ODT: Behavior with ODT Asserted Before and After the Write (Note 1)



Note to Figure 1-4:

(1) Source: *TN-41-04 DDR3 Dynamic On-Die Termination*, Micron.

In the two-DIMM DDR3 SDRAM configuration, dynamic ODT helps reduce the jitter at the module being accessed, and minimizes reflections from any secondary modules.

For more information about using the dynamic ODT on DDR3 SDRAM, refer to the application note by Micron, *TN-41-04 DDR3 Dynamic On-Die Termination*.

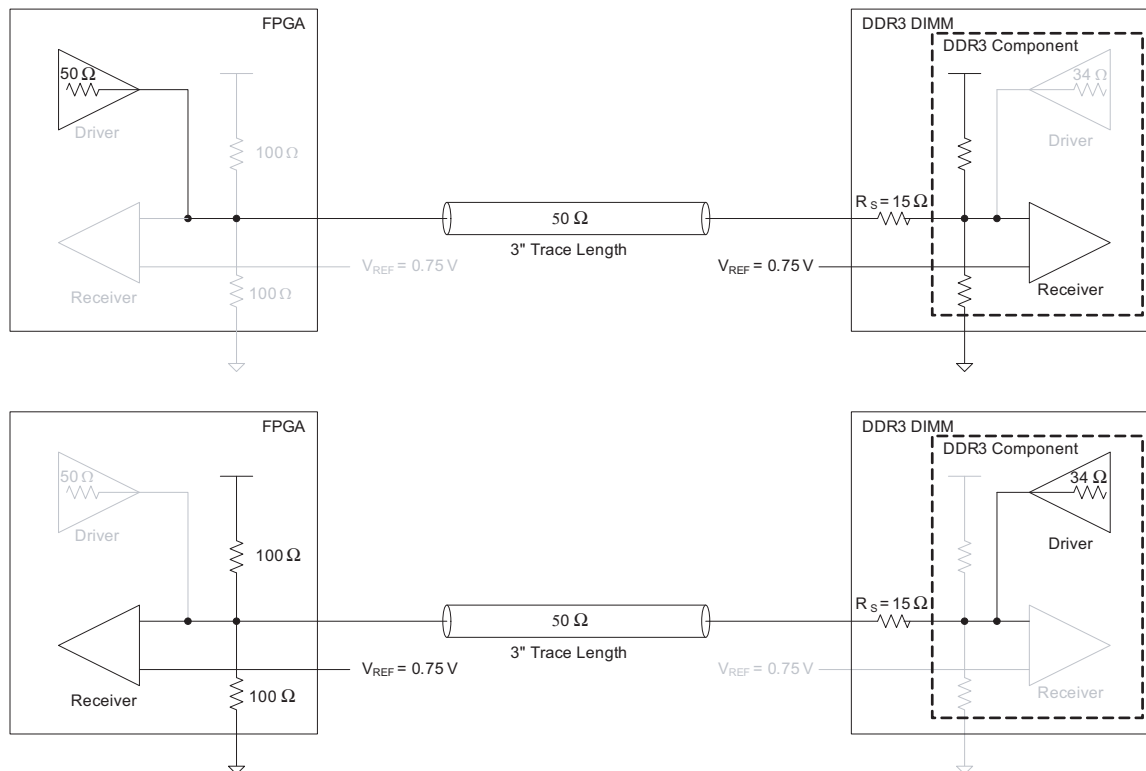
Dynamic OCT in Stratix III and Stratix IV Devices

Stratix III and Stratix IV devices support on-off dynamic series and parallel termination for a bi-directional I/O in all I/O banks. Dynamic OCT is a new feature in Stratix III and Stratix IV FPGA devices. You enable dynamic parallel termination only when the bidirectional I/O acts as a receiver and disable it when the bidirectional I/O acts as a driver. Similarly, you enable dynamic series termination only when the bidirectional I/O acts as a driver and disable it when the bidirectional I/O acts as a receiver. The default setting for dynamic OCT is series termination, to save power when the interface is idle—no active reads or writes.



Additionally, the dynamic control operation of the OCT is separate to the output enable signal for the buffer. Hence, UniPHY IP can only enable parallel OCT during read cycles, saving power when the interface is idle.

Figure 1–5. Dynamic OCT Between Stratix III and Stratix IV FPGA Devices



This feature is useful for terminating any high-performance bidirectional path because signal integrity is optimized depending on the direction of the data. In addition, dynamic OCT also eliminates the need for external termination resistors when used with memory devices that support ODT (such as DDR3 SDRAM), thus reducing cost and easing board layout.

However, dynamic OCT in Stratix III and Stratix IV FPGA devices is different from dynamic ODT in DDR3 SDRAM mentioned in previous sections and these features should not be assumed to be identical.



For detailed information about the dynamic OCT feature in the Stratix III FPGA, refer to the *Stratix III Device I/O Features* chapter in volume 1 of the *Stratix III Device Handbook*.

For detailed information about the dynamic OCT feature in the Stratix IV FPGA, refer to the *I/O Features in Stratix IV Devices* chapter in volume 1 of the *Stratix IV Device Handbook*.

Dynamic OCT in Stratix V Devices

Stratix V devices also support dynamic OCT feature and provide more flexibility. Stratix V OCT calibration uses one RZQ pin that exists in every OCT block. You can use any one of the following as a reference resistor on the RZQ pin to implement different OCT values:

- 240- Ω reference resistor—to implement R_S OCT of 34 Ω , 40 Ω , 48 Ω , 60 Ω , and 80 Ω ; and R_T OCT resistance of 20 Ω , 30 Ω , 40 Ω , and 120 Ω .
- 100 Ω reference resistor—to implement R_S OCT of 25 Ω and 50 Ω ; and R_T OCT resistance of 50 Ω .

For detailed information about the dynamic OCT feature in the Stratix V FPGA, refer to the *I/O Features in Stratix V Devices* chapter in volume 1 of the *Stratix V Device Handbook*.

Board Termination for DDR2 SDRAM

DDR2 adheres to the JEDEC standard of governing Stub-Series Terminated Logic (SSTL), JESD8-15a, which includes four different termination schemes.

Two commonly used termination schemes of SSTL are:

- Single parallel terminated output load with or without series resistors (Class I, as stated in JESD8-15a)
- Double parallel terminated output load with or without series resistors (Class II, as stated in JESD8-15a)


Depending on the type of signals you choose, you can use either termination scheme. Also, depending on your design's FPGA and SDRAM memory devices, you may choose external or internal termination schemes.

With the ever-increasing requirements to reduce system cost and simplify printed circuit board (PCB) layout design, you may choose not to have any parallel termination on the transmission line, and use point-to-point connections between the memory interface and the memory. In this case, you may take advantage of internal termination schemes such as on-chip termination (OCT) on the FPGA side and on-die termination (ODT) on the SDRAM side when it is offered on your chosen device.

External Parallel Termination


If you use external termination, you must study the locations of the termination resistors to determine which topology works best for your design. [Figure 1-6](#) and [Figure 1-7](#) illustrate the two most commonly used termination topologies: fly-by topology and non-fly-by topology, respectively.

Figure 1-6. Fly-By Placement of a Parallel Resistor



With fly-by topology ([Figure 1-6](#)), you place the parallel termination resistor after the receiver. This termination placement resolves the undesirable unterminated stub found in the non-fly-by topology. However, using this topology can be costly and complicate routing. The Stratix II Memory Board 2 uses the fly-by topology for the parallel terminating resistors placement. The Stratix II Memory Board 2 is a memory test board available only within Altera for the purpose of testing and validating Altera's memory interface.

Figure 1-7. Non-Fly-By Placement of a Parallel Resistor



With non-fly-by topology (Figure 1-7), the parallel termination resistor is placed between the driver and receiver (closest to the receiver). This termination placement is easier for board layout, but results in a short stub, which causes an unterminated transmission line between the terminating resistor and the receiver. The unterminated transmission line results in ringing and reflection at the receiver.

If you do not use external termination, DDR2 offers ODT and Altera® FPGAs have varying levels of OCT support. You should explore using ODT and OCT to decrease the board power consumption and reduce the required board real estate.

On-Chip Termination

OCT technology is offered on Arria II GX, Cyclone III, Stratix III, Stratix IV, and Stratix V devices. Table 1-1 summarizes the extent of OCT support for each device. This table provides information about SSTL-18 standards because SSTL-18 is the supported standard for DDR2 memory interface by Altera FPGAs.

On-chip series (R_S) termination is supported only on output and bidirectional buffers. The value of R_S with calibration is calibrated against a 25- Ω resistor for class II and 50- Ω resistor for class I connected to R_{UP} and R_{DN} pins and adjusted to $\pm 1\%$ of 25 Ω or 50 Ω . On-chip parallel (R_T) termination is supported only on inputs and bidirectional buffers. The value of R_T is calibrated against 100 Ω connected to the R_{UP} and R_{DN} pins. Calibration occurs at the end of device configuration. Dynamic OCT is supported only on bidirectional I/O buffers.

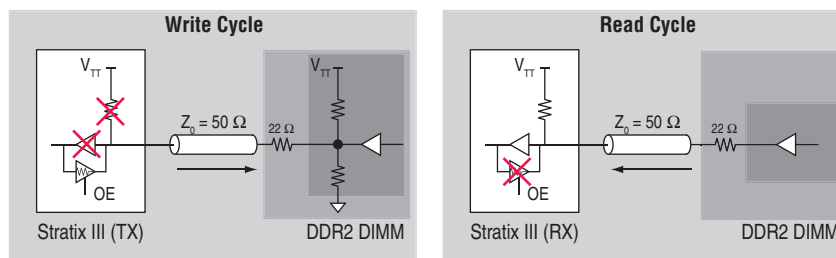
Table 1-1. On-Chip Termination Schemes

Termination Scheme	SSTL-18	FPGA Device							
		Arria II GX		Cyclone III and Cyclone IV		Stratix III and Stratix IV		Stratix V	
		Column I/O	Row I/O	Column I/O	Row I/O	Column I/O	Row I/O	Column I/O	Row I/O
On-Chip Series Termination without Calibration	Class I	50	50	50	50	50	50	50	—
	Class II	25	—	25	25	25	25	25	—
On-Chip Series Termination with Calibration	Class I	50	—	50	50	50	50	50	—
	Class II	—	—	25	25	25	25	25	—
On-Chip Parallel Termination with Calibration	Class I and Class II	—	—	—	—	50	50	50	—

The dynamic OCT scheme is only available in Stratix III, Stratix IV, and Stratix V FPGAs. The dynamic OCT scheme enables series termination (R_S) and parallel termination (R_T) to be dynamically turned on and off during the data transfer.

The series and parallel terminations are turned on or off depending on the read and write cycle of the interface. During the write cycle, the R_S is turned on and the R_T is turned off to match the line impedance. During the read cycle, the R_S is turned off and the R_T is turned on as the Stratix III FPGA implements the far-end termination of the bus (Figure 1-8).

Figure 1-8. Dynamic OCT for Memory Interfaces



Recommended Termination Schemes

Table 1-2 provides the recommended termination schemes for major DDR2 memory interface signals. Signals include data (DQ), data strobe (DQS/DQSn), data mask (DM), clocks (mem_clk/mem_clk_n), and address and command signals.

When interfacing with multiple DDR2 SDRAM components where the address, command, and memory clock pins are connected to more than one load, follow these steps:

1. Simulate the system to get the new slew-rate for these signals.
2. Use the derated tIS and tIH specifications from the DDR2 SDRAM datasheet based on the simulation results.
3. If timing deration causes your interface to fail timing requirements, consider signal duplication of these signals to lower their loading, and hence improve timing.



Altera uses Class I and Class II termination in this table to refer to drive strength, and not physical termination.



You must simulate your design for your system to ensure correct functionality.

Table 1-2. Termination Recommendations (Part 1 of 2) (Note 1)

Device Family	Signal Type	SSTL 18 IO Standard (2), (3), (4), (5), (6)	FPGA End Discrete Termination	Memory End Termination 1 Rank/DIMM	Memory I/O Standard
Arria II GX					
DDR2 component	DQ/DQS	Class I 8 mA	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	HALF (7)
	DM	Class I 8 mA	N/A	56 Ω parallel to V_{TT} discrete	N/A
	Address and command	Class I MAX	N/A		N/A
	Clock	Class I 8 mA	N/A	$\times 1 = 100 \Omega$ differential (10) $\times 2 = 200 \Omega$ differential (11)	N/A
DDR2 DIMM	DQ/DQS	Class I 8 mA	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	FULL (9)
	DM	Class I 8 mA	N/A	56 Ω parallel to V_{TT} discrete	N/A
	Address and command	Class I MAX	N/A		N/A
	Clock	Class I 8 mA	N/A	N/A = on DIMM	N/A
Cyclone III and Cyclone IV					
DDR2 component	DQ/DQS	Class I 12 mA	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	HALF (8)
	DM	Class I 12 mA	N/A	56 Ω parallel to V_{TT} discrete	N/A
	Address and command	Class I MAX	N/A		N/A
	Clock	Class I 12 mA	N/A	$\times 1 = 100 \Omega$ differential (10) $\times 2 = 200 \Omega$ differential (11)	N/A
DDR2 DIMM	DQ/DQS	Class I 12 mA	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	FULL (9)
	DM	Class I 12 mA	N/A	56 Ω parallel to V_{TT} discrete	N/A
	Address and command	Class I MAX	N/A		N/A
	Clock	Class I 12 mA	N/A	N/A = on DIMM	N/A

Table 1–2. Termination Recommendations (Part 2 of 2) (Note 1)

Device Family	Signal Type	SSTL 18 IO Standard (2), (3), (4), (5), (6)	FPGA End Discrete Termination	Memory End Termination 1 Rank/DIMM	Memory I/O Standard
Stratix III, Stratix IV, and Stratix V					
DDR2 component	DQ/DQS	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (7)
	DM	Class I R50 CAL	N/A	ODT75 (7)	N/A
	Address and command	Class I R50 CAL	N/A	56 Ω Parallel to V_{TT} discrete	N/A
	Clock	DIFF Class I R50 NO CAL	N/A	x1 = 100 Ω differential (10)	N/A
		DIFF Class I R50 NO CAL	N/A	X2 = 200 Ω differential (11)	N/A
	DQS DIFF recommended	DIFF Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (7)
	DQS SE (12)	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (7)
DDR2 DIMM	DQ/DQS	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	FULL (9)
	DM	Class I R50 CAL	N/A	ODT75 (7)	N/A
	Address and command	Class I MAX	N/A	56 Ω Parallel to V_{TT} discrete	N/A
	Clock	DIFF Class I R50 NO CAL	N/A	N/A = on DIMM	N/A
	DQS DIFF recommended	DIFF Class I R50/P50 DYN CAL	N/A	ODT75 (7)	FULL (9)
	DQS SE (12)	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	FULL (9)

Notes to Table 1–2:

- (1) N/A is not available.
- (2) R is series resistor.
- (3) P is parallel resistor.
- (4) DYN is dynamic OCT.
- (5) NO CAL is OCT without calibration.
- (6) CAL is OCT with calibration.
- (7) ODT75 vs. ODT50 on the memory has the effect of opening the eye more, with a limited increase in overshoot/undershoot.
- (8) HALF is reduced drive strength.
- (9) FULL is full drive strength.
- (10) x1 is a single-device load.
- (11) x2 is two-device load.
- (12) DQS SE is single-ended DQS.

Dynamic On-Chip Termination

The termination schemes are described in JEDEC standard JESD8-15a for SSTL 18 I/O. Dynamic OCT is available in Stratix III and Stratix IV. When the Stratix III FPGA (driver) is writing to the DDR2 SDRAM DIMM (receiver), series OCT is enabled dynamically to match the impedance of the transmission line. As a result, reflections are significantly reduced. Similarly, when the FPGA is reading from the DDR2 SDRAM DIMM, the parallel OCT is dynamically enabled.

- For information about setting the proper value for termination resistors, refer to the *Stratix III Device I/O Features* chapter in the *Stratix III Device Handbook* and the *I/O Features in Stratix IV Devices* chapter in the *Stratix IV Device Handbook*.

FPGA Writing to Memory

Figure 1-9 shows dynamic series OCT scheme when the FPGA is writing to the memory. The benefit of using dynamic series OCT is that when driver is driving the transmission line, it “sees” a matched transmission line with no external resistor termination.

Figure 1-9. Dynamic Series OCT Scheme with ODT on the Memory




Figure 1-10 and Figure 1-11 show the simulation and measurement results of a write to the DDR2 SDRAM DIMM. The system uses Class I termination with a 50- Ω series OCT measured at the DIMM with a full drive strength and a 75 Ω ODT at the DIMM. Both simulation and bench measurements are in 200 pS/div and 200 mV/div.

Figure 1-10. HyperLynx Simulation FPGA Writing to Memory




Figure 1-11. Board Measurement, FPGA Writing to Memory




Table 1-3 summarizes the comparison between the simulation and the board measurement of the signal seen at the DDR2 SDRAM DIMM.

Table 1-3. Signal Comparison When the FPGA is Writing to the Memory (Note 1)

	Eye Width (ns) (2)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Simulation	1.194	0.740	N/A	N/A
Board Measurement	1.08	0.7	N/A	N/A


Notes to Table 1-3:

- (1) N/A is not applicable.
- (2) The eye width is measured from $V_{IH}/V_{IL}(ac) = VREF \pm 250$ mV to $V_{IH}/V_{IL}(dc) = VREF \pm 125$ mV, where V_{IH} and V_{IL} are determined per the JEDEC specification for SSTL-18.

The data in Table 1-3 and Figure 1-10 and Figure 1-11 suggest that when the FPGA is writing to the memory, the bench measurements are closely matched with simulation measurements. They indicate that using the series dynamic on-chip termination scheme for your bidirectional I/Os maintains the integrity of the signal, while it removes the need for external termination.

Depending on the I/O standard, you should consider the four parameters listed in Table 1-3 when designing a memory interface. Although the simulation and board measurement appear to be similar, there are some discrepancies when the key parameters are measured. Although simulation does not fully model the duty cycle distortion of the I/O, crosstalk, or board power plane degradation, it provides a good indication on the performance of the board.

For memory interfaces, the eye width is important when determining if there is a sufficient window to correctly capture the data. Regarding the eye height, even though most memory interfaces use voltage-referenced I/O standards (in this case, SSTL-18), as long as there is sufficient eye opening below and above V_{IL} and V_{IH} , there should be enough margin to correctly capture the data. However, because effects such as crosstalk are not taken into account, it is critical to design a system to achieve the optimum eye height, because it impacts the overall margin of a system with a memory interface.

-  Refer to the memory vendors when determining the over- and undershoot. They typically specify a maximum limit on the input voltage to prevent reliability issues.

FPGA Reading from Memory

Figure 1-12 shows the dynamic parallel termination scheme when the FPGA is reading from memory. When the DDR2 SDRAM DIMM is driving the transmission line, the ringing and reflection is minimal because the FPGA-side termination 50- Ω pull-up resistor is matched with the transmission line. Figure 1-13 shows the simulation and measurement results of a read from DDR2 SDRAM DIMM. The system uses Class I termination with a 50- Ω calibrated parallel OCT measured at the FPGA end with a full drive strength and a 75- Ω ODT at the memory. Both simulation and bench measurements are in 200 pS/div and 200 mV/div.

Figure 1-12. Dynamic Parallel OCT Scheme with Memory-Side Series Resistor




Figure 1-13. Hyperlynx Simulation and Board Measurement, FPGA Reading from Memory




Table 1–4 summarizes the comparison between the simulation and the board measurement of the signal seen at the FPGA end.

Table 1–4. Signal Comparison When the FPGA is Reading from the Memory (Note 1), (2)

	Eye Width (ns) <i>(3)</i>	Eye Height (V)	Overshoot (V)	Undershoot (V)
Simulation	1.206	0.740	N/A	N/A
Board Measurement	1.140	0.680	N/A	N/A

Notes to Table 1–4:


- (1) The drive strength on the memory DIMM is set to Full.
- (2) N/A is not applicable.
- (3) The eye width is measured from $V_{IH}/V_{IL}(ac) = VREF \pm 250$ mV to $V_{IH}/V_{IL}(dc) = VREF \pm 125$ mV, in which V_{IH} and V_{IL} are determined per the JEDEC specification for SSTL-18.

The data in [Table 1–4](#) and [Figure 1–13](#) suggest that bench measurements are closely matched with simulation measurements when the FPGA is reading from the memory. They indicate that using the parallel dynamic on-chip termination scheme in bidirectional I/Os maintains the integrity of the signal, while it removes the need for external termination.

On-Chip Termination (Non-Dynamic)


When you use the 50-Ω OCT feature in a Class I termination scheme using ODT with a memory-side series resistor, the output driver is tuned to 50 Ω, which matches the characteristic impedance of the transmission line. [Figure 1–14](#) shows the Class I termination scheme using ODT when the 50-Ω OCT on the FPGA is turned on.

Figure 1–14. Class I Termination Using ODT with 50-Ω OCT



The resulting signal quality has a similar eye opening to the 8 mA drive strength setting (refer to “[Drive Strength](#)” on page 1-51) without any over- or undershoot. [Figure 1-15](#) shows the simulation and measurement of the signal at the memory side (DDR2 SDRAM DIMM) with the drive strength setting of 50-Ω OCT in the FPGA.

Figure 1-15. HyperLynx Simulation and Measurement, FPGA Writing to Memory



[Table 1-5](#) shows data for the signal at the DDR2 SDRAM DIMM of a Class I scheme termination using ODT with a memory-side series resistor. The FPGA is writing to the memory with 50-Ω OCT.

Table 1-5. Simulation and Board Measurement Results for 50-Ω OCT and 8-mA Drive Strength Settings [\(Note 1\)](#)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
50-Ω OCT Drive Strength Setting				
Simulation	1.68	0.82	N/A	N/A
Board Measurement	1.30	0.70	N/A	N/A

Note to [Table 1-5](#):

(1) N/A is not applicable.

When you use the 50-Ω OCT setting on the FPGA, the signal quality for the Class I termination using ODT with a memory-side series resistor is further improved with lower over- and undershoot.

In addition to the 50- Ω OCT setting, Stratix II devices have a 25- Ω OCT setting that you can use to improve the signal quality in a Class II terminated transmission line. Figure 1-16 shows the Class II termination scheme using ODT when the 25- Ω OCT on the FPGA is turned on.

Figure 1-16. Class II Termination Using ODT with 25- Ω OCT




Figure 1-17 shows the simulation and measurement of the signal at the DDR2 SDRAM DIMM (receiver) with a drive strength setting of 25- Ω OCT in the FPGA.

Figure 1-17. HyperLynx Simulation and Measurement, FPGA Writing to Memory

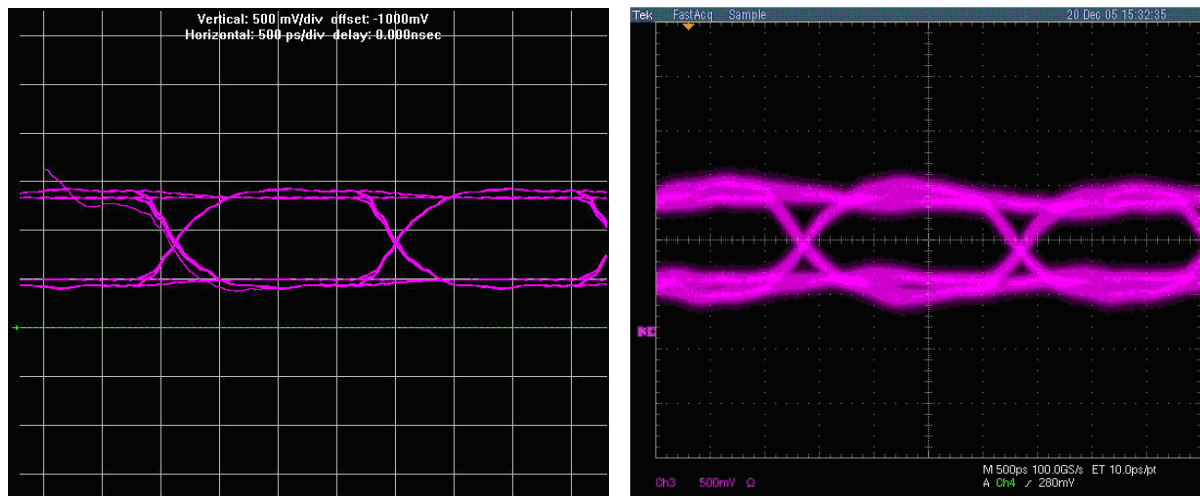


Table 1-6 shows the data for the signal at the DDR2 SDRAM DIMM of a Class II termination with a memory-side series resistor. The FPGA is writing to the memory with 25- Ω OCT.

Table 1-6. Simulation and Board Measurement Results for 25- Ω OCT and 16-mA Drive Strength Settings (Part 1 of 2) (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
25-Ω OCT Drive Strength Setting				
Simulation	1.70	0.81	N/A	N/A

Table 1-6. Simulation and Board Measurement Results for 25- Ω OCT and 16-mA Drive Strength Settings (Part 2 of 2) (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Board Measurement	1.47	0.51	N/A	N/A

Note to Table 1-6:

(1) N/A is not applicable.

This type of termination scheme is only used for bidirectional signals, such as data (DQ), data strobe (DQS), data mask (DM), and memory clocks (CK) found in DRAMs.

Class II External Parallel Termination

The double parallel (Class II) termination scheme is described in JEDEC standards JESD8-6 for HSTL I/O, JESD8-9b for SSTL-2 I/O, and JESD8-15a for SSTL-18 I/O. When the FPGA (driver) is writing to the DDR2 SDRAM DIMM (receiver), the transmission line is terminated at the DDR2 SDRAM DIMM. Similarly, when the FPGA is reading from the DDR2 SDRAM DIMM, the DDR2 SDRAM DIMM is now the driver and the transmission line is terminated at the FPGA (receiver). This type of termination scheme is typically used for bidirectional signals, such as data (DQ) and data strobe (DQS) signal found in DRAMs.

FPGA Writing to Memory

Figure 1-18 shows the Class II termination scheme when the FPGA is writing to the memory. The benefit of using Class II termination is that when either driver is driving the transmission line, it sees a matched transmission line because of the termination resistor at the receiver-end, thereby reducing ringing and reflection.



Figure 1-18. Class-II Termination Scheme with Memory-Side Series Resistor

Figure 1-19 and Figure 1-20 show the simulation and measurement result of a write to the DDR2 SDRAM DIMM. The system uses Class II termination with a source-series resistor measured at the DIMM with a drive strength setting of 16 mA.

Figure 1-19. HyperLynx Simulation, FPGA Writing to Memory



The simulation shows a clean signal with a good eye opening, but there is slight over- and undershoot of the 1.8-V signal specified by DDR2 SDRAM. The over- and undershoot can be attributed to either overdriving the transmission line using a higher than required drive strength setting on the driver or the over-termination on the receiver side by using an external resistor value that is higher than the characteristic impedance of the transmission line. As long as the over- and undershoot do not exceed the absolute maximum rating specification listed in the memory

vendor's DDR2 SDRAM data sheet, it does not result in any reliability issues. The simulation results are then correlated with actual board level measurements.

Figure 1-20 shows the measurement obtained from the Stratix II Memory Board 2. The FPGA is using a 16 mA drive strength to drive the DDR2 SDRAM DIMM on a Class II termination transmission line.

Figure 1-20. Board Measurement, FPGA Writing to Memory




Table 1-7 summarizes the comparison between the simulation and the board measurement of the signal seen at the DDR2 SDRAM DIMM.

Table 1-7. Signal Comparison When the FPGA is Writing to the Memory (Note 1)

	Eye Width (ns) (2)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Simulation	1.65	1.28	0.16	0.14
Board Measurement	1.35	0.83	0.16	0.18

Notes to Table 1-7:

- (1) The drive strength on the FPGA is set to 16 mA.
- (2) The eye width is measured from $V_{REF} \pm 125$ mV where V_{IH} and V_{IL} are determined per the JEDEC specification for SSTL-18.

A closer inspection of the simulation shows an ideal duty cycle of 50%–50%, while the board measurement shows that the duty cycle is non-ideal, around 53%–47%, resulting in the difference between the simulation and measured eye width. In addition, the board measurement is conducted on a 72-bit memory interface, but the simulation is performed on a single I/O.

FPGA Reading from Memory

Figure 1-21 shows the Class II termination scheme when the FPGA is reading from memory. When the DDR2 SDRAM DIMM is driving the transmission line, the ringing and reflection is minimal because of the matched FPGA-side termination pull-up resistor with the transmission line.

Figure 1-21. Class II Termination Scheme with Memory-Side Series Resistor




Figure 1-22 and Figure 1-23 show the simulation and measurement, respectively, of the signal at the FPGA side with the full drive strength setting on the DDR2 SDRAM DIMM. The simulation uses a Class II termination scheme with a source-series resistor transmission line. The FPGA is reading from the memory with a full drive strength setting on the DIMM.

Figure 1-22. HyperLynx Simulation, FPGA Reading from Memory




Figure 1-23. Board Measurement, FPGA Reading from Memory

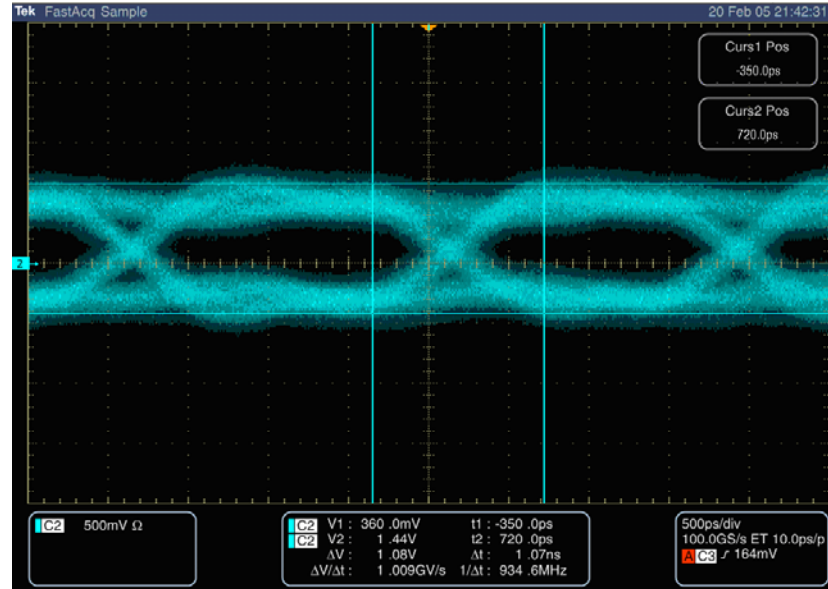


Table 1-8 summarizes the comparison between the simulation and board measurements of the signal seen by the FPGA when the FPGA is reading from memory (driver).

Table 1-8. Signal Comparison, FPGA is Reading from Memory (Note 1), (2)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Simulation	1.73	0.76	N/A	N/A
Board Measurement	1.28	0.43	N/A	N/A

Note to Table 1-8:

- (1) The drive strength on the DDR2 SDRAM DIMM is set to full strength.
- (2) N/A is not applicable.

Both simulation and measurement show a clean signal and a good eye opening without any over- and undershoot. However, the eye height when the FPGA is reading from the memory is smaller compared to the eye height when the FPGA is writing to the memory. The reduction in eye height is attributed to the voltage drop on the series resistor present on the DIMM. With the drive strength setting on the memory already set to full, you cannot increase the memory drive strength to improve the eye height. One option is to remove the series resistor on the DIMM when the FPGA is reading from memory (refer to the section “[Component Versus DIMM](#)” on page 1-53). Another option is to remove the external parallel resistor near the memory so that the memory driver sees less loading. For a DIMM configuration, the latter option is a better choice because the series resistors are part of the DIMM and you can easily turn on the ODT feature to use as the termination resistor when the FPGA is writing to the memory and turn off when the FPGA is reading from memory.

The results for the Class II termination scheme demonstrate that the scheme is ideal for bidirectional signals such as data strobe and data for DDR2 SDRAM memory. Terminations at the receiver eliminate reflections back to the driver and suppress any ringing at the receiver.

Class I External Parallel Termination

The single parallel (Class I) termination scheme refers to when the termination is located near the receiver side. Typically, this scheme is used for terminating unidirectional signals (such as clocks, address, and command signals) for DDR2 SDRAM.

However, because of board constraints, this form of termination scheme is sometimes used in bidirectional signals, such as data (DQ) and data strobe (DQS) signals. For bidirectional signals, you can place the termination on either the memory or the FPGA side. This section focuses only on the Class I termination scheme with memory-side termination. The memory-side termination ensures impedance matching when the signal reaches the receiver of the memory. However, when the FPGA is reading from the memory, there is no termination on the FPGA side, resulting in impedance mismatch. This section describes the signal quality of this termination scheme.

FPGA Writing to Memory

When the FPGA is writing to the memory (Figure 1-24), the transmission line is parallel-terminated at the memory side, resulting in minimal reflection on the receiver side because of the matched impedance seen by the transmission line. The benefit of this termination scheme is that only one external resistor is required. Alternatively, you can implement this termination scheme using an ODT resistor instead of an external resistor.

Refer to the section “Class I Termination Using ODT” on page 1-29 for more information about how an ODT resistor compares to an external termination resistor.

Figure 1-24. Class I Termination Scheme with Memory-Side Series Resistor




Figure 1-25 shows the simulation and measurement of the signal at the memory (DDR2 SDRAM DIMM) of Class I termination with a memory-side resistor. The FPGA writes to the memory with a 16 mA drive strength setting.

Figure 1-25. HyperLynx Simulation and Board Measurement, FPGA Writing to Memory

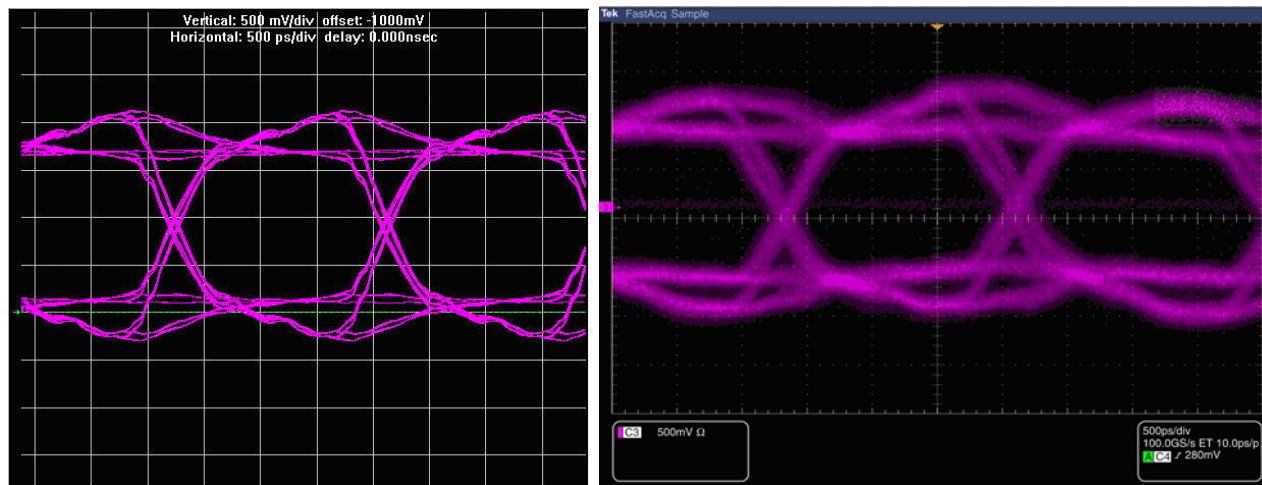


Table 1–9 summarizes the comparison of the signal at the DDR2 SDRAM DIMM of a Class I and Class II termination scheme using external resistors with memory-side series resistors. The FPGA (driver) writes to the memory (receiver).

Table 1–9. Signal Comparison When the FPGA is Writing to Memory (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Class I Termination Scheme With External Parallel Resistor				
Simulation	1.69	1.51	0.34	0.29
Board Measurement	1.25	1.08	0.41	0.34
Class II Termination Scheme With External Parallel Resistor				
Simulation	1.65	1.28	0.16	0.14
Board Measurement	1.35	0.83	0.16	0.18

Note to Table 1–9:

(1) The drive strength on the FPGA is set to 16 mA.

Table 1–9 shows the overall signal quality of a Class I termination scheme is comparable to the signal quality of a Class II termination scheme, except that the eye height of the Class I termination scheme is approximately 30% larger. The increase in eye height is due to the reduced loading “seen” by the driver, because the Class I termination scheme does not have an FPGA-side parallel termination resistor. However, increased eye height comes with a price: a 50% increase in the over- and undershoot of the signal using Class I versus Class II termination scheme. You can decrease the FPGA drive strength to compensate for the decreased loading seen by the driver to decrease the over- and undershoot.


Refer to the section “Drive Strength” on page 1-51 for more information about how drive strength affects the signal quality.

FPGA Reading from Memory

As described in the section “FPGA Writing to Memory” on page 1-26, in Class I termination, the termination is located near the receiver. However, if you use this termination scheme to terminate a bidirectional signal, the receiver can also be the driver. For example, in DDR2 SDRAM, the data signals are both receiver *and* driver.

Figure 1–26 shows a Class I termination scheme with a memory-side resistor. The FPGA reads from the memory.

Figure 1–26. Class I Termination Scheme with Memory-Side Series Resistor



When the FPGA reads from the memory (Figure 1-26), the transmission line is not terminated at the FPGA, resulting in an impedance mismatch, which then results in over- and undershoot. Figure 1-27 shows the simulation and measurement of the signal at the FPGA side (receiver) of a Class I termination. The FPGA reads from the memory with a full drive strength setting on the DDR2 SDRAM DIMM.

Figure 1-27. HyperLynx Simulation and Board Measurement, FPGA Reading from Memory




Table 1-10 summarizes the comparison of the signal “seen” at the FPGA of a Class I and Class II termination scheme using an external resistor with a memory-side series resistor. The FPGA (receiver) reads from the memory (driver).

Table 1-10. Signal Comparison When the FPGA is Reading From Memory (Note 1), (2)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Class I Termination Scheme with External Parallel Resistor				
Simulation	1.73	0.74	0.20	0.18
Board Measurement	1.24	0.58	0.09	0.14
Class II Termination Scheme with External Parallel Resistor				
Simulation	1.73	0.76	N/A	N/A
Board Measurement	1.28	0.43	N/A	N/A

Note to Table 1-10:

- (1) The drive strength on the DDR2 SDRAM DIMM is set to full strength.
- (2) N/A is not applicable.

When the FPGA reads from the memory using the Class I scheme, the signal quality is comparable to that of the Class II scheme, in terms of the eye height and width.

Table 1-10 shows the lack of termination at the receiver (FPGA) results in impedance mismatch, causing reflection and ringing that is not visible in the Class II termination scheme. As such, Altera recommends using the Class I termination scheme for unidirectional signals (such as command and address signals), between the FPGA and the memory.

Class I Termination Using ODT

Presently, ODT is becoming a common feature in memory, including SDRAMs, graphics DRAMs, and SRAMs. ODT helps reduce board termination cost and simplify board routing. This section describes the ODT feature of DDR2 SDRAM and the signal quality when the ODT feature is used.

FPGA Writing to Memory

DDR2 SDRAM has built-in ODT that eliminates the need for external termination resistors. To use the ODT feature of the memory, you must configure the memory to turn on the ODT feature during memory initialization. For DDR2 SDRAM, set the ODT feature by programming the extended mode register. In addition to programming the extended mode register during initialization of the DDR2 SDRAM, an ODT input pin on the DDR2 SDRAM must be driven high to activate the ODT.



Refer to the respective memory data sheet for additional information about setting the ODT feature and the timing requirements for driving the ODT pin in DDR2 SDRAM.

The ODT feature in DDR2 SDRAM is controlled dynamically—it is turned on while the FPGA is writing to the memory and turned off while the FPGA is reading from the memory. The ODT feature in DDR2 SDRAM has three settings: 50Ω, 75Ω, and 150Ω. If there are no external parallel termination resistors and the ODT feature is turned on, the termination scheme resembles the Class I termination described in “[Class I External Parallel Termination](#)” on page 1-25.

Figure 1-28 shows the termination scheme when the ODT on the DDR2 SDRAM is turned on.

Figure 1-28. Class I Termination Scheme Using ODT

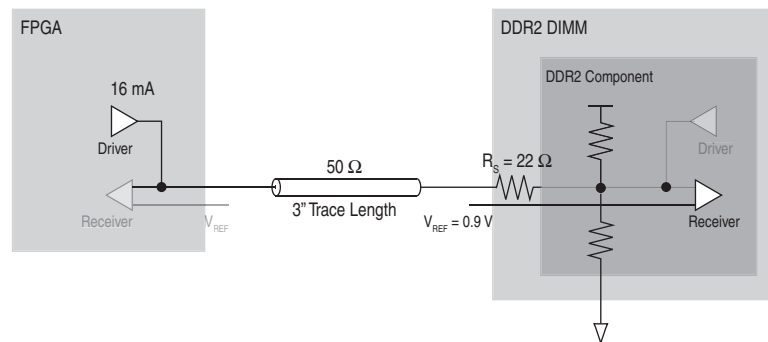


Figure 1-29 shows the simulation and measurement of the signal visible at the memory (receiver) using 50 Ω ODT with a memory-side series resistor transmission line. The FPGA writes to the memory with a 16 mA drive strength setting.

Figure 1-29. Simulation and Board Measurement, FPGA Writing to Memory




Table 1-11 summarizes the comparisons of the signal seen the DDR2 SDRAM DIMM of a Class I termination scheme using an external resistor and a Class I termination scheme using ODT with a memory-side series resistor. The FPGA (driver) writes to the memory (receiver).

Table 1-11. Signal Comparison When the FPGA is Writing to Memory (Note 1), (2)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Class I Termination Scheme with ODT				
Simulation	1.63	0.84	N/A	0.12
Board Measurement	1.51	0.76	0.05	0.15
Class I Termination Scheme with External Parallel Resistor				
Simulation	1.69	1.51	0.34	0.29
Board Measurement	1.25	1.08	0.41	0.34

Note to Table 1-11:

- (1) The drive strength on the FPGA is set to 16 mA.
- (2) N/A is not applicable.

When the ODT feature is enabled in the DDR2 SDRAM, the eye width is improved. There is some degradation to the eye height, but it is not significant. When ODT is enabled, the most significant improvement in signal quality is the reduction of the over- and undershoot, which helps mitigate any potential reliability issues on the memory devices.

Using memory ODT also eliminates the need for external resistors, which reduces board cost and simplifies board routing, allowing you to shrink your boards. Therefore, Altera recommends using the ODT feature on the DDR2 SDRAM memory.

FPGA Reading from Memory

Altera's Arria GX, Arria II GX, Cyclone series, and Stratix II series of devices are not equipped with ODT. When the DDR2 SDRAM ODT feature is turned off when the FPGA is reading from the memory, the termination scheme resembles the no-parallel termination scheme illustrated by [Figure 1-32 on page 1-33](#).

No-Parallel Termination

The no-parallel termination scheme is described in the JEDEC standards JESD8-6 for HSTL I/O, JESD8-9b for SSTL-2 I/O, and JESD8-15a for SSTL-18 I/O. Designers who attempt series-only termination schemes such as this often do so to eliminate the need for a V_{TT} power supply.

This is typically not recommended for any signals between an FPGA and DDR2 interface; however, information about this topic is included here as a reference point to clarify the challenges that may occur if you attempt to avoid parallel termination entirely.

FPGA Writing to Memory

[Figure 1-30](#) shows a no-parallel termination transmission line of the FPGA driving the memory. When the FPGA is driving the transmission line, the signals at the memory-side (DDR2 SDRAM DIMM) may suffer from signal degradation (for example, degradation in rise and fall time). This is due to impedance mismatch, because there is no parallel termination at the memory-side. Also, because of factors such as trace length and drive strength, the degradation seen at the receiver-end might be sufficient to result in a system failure. To understand the effects of each termination scheme on a system, perform system-level simulations before and after the board is designed.

Figure 1-30. No-Parallel Termination Scheme


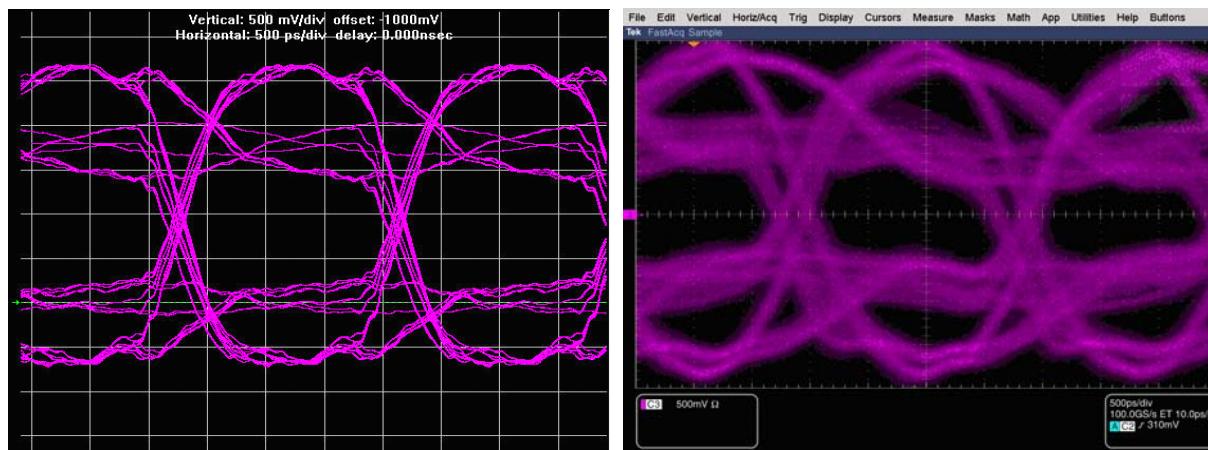


Figure 1–31 shows a HyperLynx simulation and measurement of the FPGA writing to the memory at 533 MHz with a no-parallel termination scheme using a 16 mA drive strength option. The measurement point is on the DDR2 SDRAM DIMM.

Figure 1–31. HyperLynx Simulation and Board Measurement, FPGA Writing to Memory



The simulated and measured signal shows that there is sufficient eye opening but also significant over- and undershoot of the 1.8-V signal specified by the DDR2 SDRAM. From the simulation and measurement, the overshoot is approximately 1 V higher than 1.8 V, and undershoot is approximately 0.8 V below ground. This over- and undershoot might result in a reliability issue, because it has exceeded the absolute maximum rating specification listed in the memory vendors' DDR2 SDRAM data sheet.

Table 1–12 summarizes the comparison of the signal visible at the DDR2 SDRAM DIMM of a no-parallel and a Class II termination scheme when the FPGA writes to the DDR2 SDRAM DIMM.

Table 1–12. Signal Comparison When the FPGA is Writing to Memory (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
No-Parallel Termination Scheme				
Simulation	1.66	1.10	0.90	0.80
Board Measurement	1.25	0.60	1.10	1.08
Class II Termination Scheme With External Parallel Resistor				
Simulation	1.65	1.28	0.16	0.14
Board Measurement	1.35	0.83	0.16	0.18

Note to Table 1–12:

(1) The drive strength on the FPGA is set to Class II 16 mA.

Although the appearance of the signal in a no-parallel termination scheme is not clean, when you take the key parameters into consideration, the eye width and height is comparable to that of a Class II termination scheme. The major disadvantage of using a no-parallel termination scheme is the over- and undershoot. There is no termination on the receiver, so there is an impedance mismatch when the signal arrives at the receiver, resulting in ringing and reflection. In addition, the 16-mA drive strength setting on the FPGA also results in overdriving the transmission line, causing the over- and undershoot. By reducing the drive strength setting, the over- and undershoot decreases and improves the signal quality “seen” by the receiver.

For more information about how drive strength affects the signal quality, refer to [“Drive Strength” on page 1-51](#).

FPGA Reading from Memory

In a no-parallel termination scheme ([Figure 1-32](#)), when the memory is driving the transmission line, the resistor, R_S , acts as a source termination resistor. The DDR2 SDRAM driver has two drive strength settings:

- Full strength, in which the output impedance is approximately 18Ω
- Reduced strength, in which the output impedance is approximately 40Ω

When the DDR2 SDRAM DIMM drives the transmission line, the combination of the $22\text{-}\Omega$ source-series resistor and the driver impedance should match that of the characteristic impedance of the transmission line. As such, there is less over- and undershoot of the signal visible at the receiver (FPGA).

Figure 1-32. No-Parallel Termination Scheme, FPGA Reading from Memory




Figure 1-33 shows the simulation and measurement of the signal visible at the FPGA (receiver) when the memory is driving the no-parallel termination transmission line with a memory-side series resistor.

Figure 1-33. HyperLynx Simulation and Board Measurement, FPGA Reading from Memory

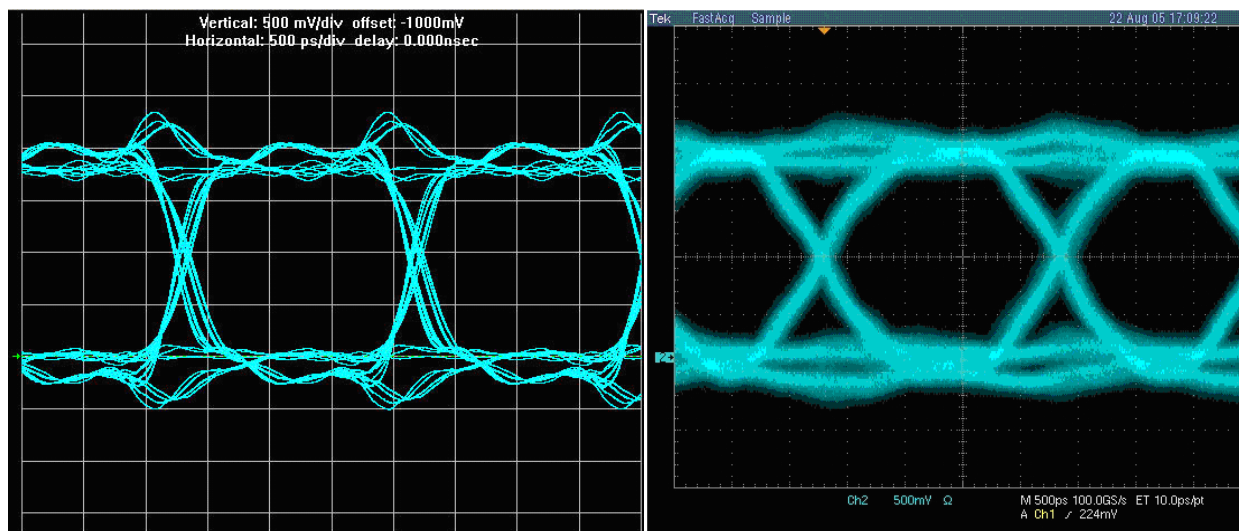


Table 1-13 summarizes the comparison of the signal seen on the FPGA with a no-parallel and a Class II termination scheme when the FPGA is reading from memory.

Table 1-13. Signal Comparison, FPGA Reading From Memory (Note 1), (2)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
No-Parallel Termination Scheme				
Simulation	1.82	1.57	0.51	0.51
Board Measurement	1.62	1.29	0.28	0.37
Class II Termination Scheme with External Parallel Resistor				
Simulation	1.73	0.76	N/A	N/A
Board Measurement	1.28	0.43	N/A	N/A

Note to Table 1-13:

- (1) The drive strength on the DDR2 SDRAM DIMM is set to full strength.
- (2) N/A is not applicable.

As in the section “FPGA Writing to Memory” on page 1-31, the eye width and height of the signal in a no-parallel termination scheme is comparable to a Class II termination scheme, but the disadvantage is the over- and undershoot. There is over- and undershoot because of the lack of termination on the transmission line, but the magnitude of the over- and undershoot is not as severe when compared to that described in “FPGA Writing to Memory” on page 1-31. This is attributed to the presence of the series resistor at the source (memory side), which dampens any reflection coming back to the driver and further reduces the effect of the reflection on the FPGA side.

When the memory-side series resistor is removed (Figure 1-34), the memory driver impedance no longer matches the transmission line and there is no series resistor at the driver to dampen the reflection coming back from the unterminated FPGA side.

Figure 1-34. No-Parallel Termination Scheme, FPGA Reading from Memory




Figure 1-35 shows the simulation and measurement of the signal at the FPGA side in a no-parallel termination scheme with the full drive strength setting on the memory.

Figure 1-35. HyperLynx Simulation and Measurement, FPGA Reading from Memory

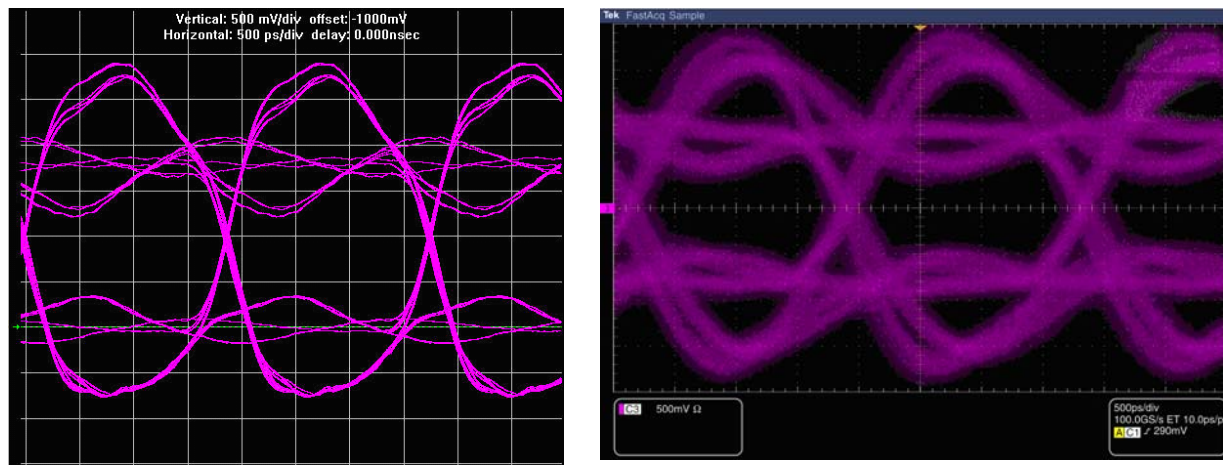


Table 1-14 summarizes the difference between no-parallel termination with and without memory-side series resistor when the memory (driver) writes to the FPGA (receiver).

Table 1-14. No-Parallel Termination with and without Memory-Side Series Resistor (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
Without Series Resistor				
Simulation	1.81	0.85	1.11	0.77
Board Measurement	1.51	0.92	0.96	0.99
With Series Resistor				
Simulation	1.82	1.57	0.51	0.51
Board Measurement	1.62	1.29	0.28	0.37

Note to Table 1-14:

(1) The drive strength on the memory is set to full drive strength.

Table 1-14 highlights the effect of the series resistor on the memory side with the dramatic increase in over- and undershoot and the decrease in the eye height. This result is similar to that described in “FPGA Writing to Memory” on page 1-31. In that simulation, there is a series resistor but it is located at the receiver side (memory-side), so it does not have the desired effect of reducing the drive strength of the driver and suppressing the reflection coming back from the unterminated receiver-end. As such, in a system without receiver-side termination, the series resistor on the driver helps reduce the drive strength of the driver and dampen the reflection coming back from the unterminated receiver-end.

Board Termination for DDR3 SDRAM

The following sections describe the correct way to terminate a DDR3 SDRAM interface together with Stratix III, Stratix IV, and Stratix V FPGA devices.

DDR3 DIMMs have terminations on all unidirectional signals, such as memory clocks, and addresses and commands; thus eliminating the need for them on the FPGA PCB. In addition, using the ODT feature on the DDR3 SDRAM and the dynamic OCT feature of Stratix III, Stratix IV, and Stratix V FPGA devices completely eliminates any external termination resistors; thus simplifying the layout for the DDR3 SDRAM interface when compared to that of the DDR2 SDRAM interface.

This section describes the termination for the following DDR3 SDRAM components:

- Single-Rank DDR3 SDRAM Unbuffered DIMM
- Multi-Rank DDR3 SDRAM Unbuffered DIMM
- DDR3 SDRAM Registered DIMM
- DDR3 SDRAM Components With Leveling



If you are using a DDR3 SDRAM without leveling interface, refer to the “Board Termination for DDR2 SDRAM” on page 1-7.

Single-Rank DDR3 SDRAM Unbuffered DIMM

The most common implementation of the DDR3 SDRAM interface is the unbuffered DIMM (UDIMM). You can find DDR3 SDRAM UDIMMs in many applications, especially in PC applications.

Table 1-15 lists the recommended termination and drive strength setting for UDIMM and Stratix III, Stratix IV, and Stratix V FPGA devices.



These settings are just recommendations for you to get started. Simulate with real board and try different settings to get the best SI.

Table 1-15. Drive Strength and ODT Setting Recommendations for Single-Rank UDIMM

Signal Type	SSTL 15 I/O Standard (1)	FPGA End On-Board Termination (2)	Memory End Termination for Write	Memory Driver Strength for Read
DQ	Class I R50C/G50C (3)	—	60 Ω ODT (4)	40 Ω (4)
DQS	Differential Class I R50C/G50C (3)	—	60 Ω ODT (4)	40 Ω (4)

Table 1–15. Drive Strength and ODT Setting Recommendations for Single-Rank UDIMM

Signal Type	SSTL 15 I/O Standard (1)	FPGA End On-Board Termination (2)	Memory End Termination for Write	Memory Driver Strength for Read
DM	Class I R50C (3)	—	60 Ω ODT (4)	40 Ω (4)
Address and Command	Class I with maximum drive strength	—	39 Ω on-board termination to V_{TT} (5)	
CK/CK#	Differential Class I R50C	—	On-board (5): 2.2 pF compensation cap before the first component; 36 Ω termination to V_{TT} for each arm (72 Ω differential); add 0.1 μ F just before V_{TT}	

Notes to Table 1–16:

- (1) UniPHY IP automatically implements these settings.
- (2) Altera recommends that you use dynamic on-chip termination (OCT) for Stratix III and Stratix IV device families.
- (3) R50C is series with calibration for write, G50C is parallel 50 with calibration for read.
- (4) You can specify these settings in the parameter editor.
- (5) For DIMM, these settings are already implemented on the DIMM card; for component topology, Altera recommends that you mimic termination scheme on the DIMM card on your board.

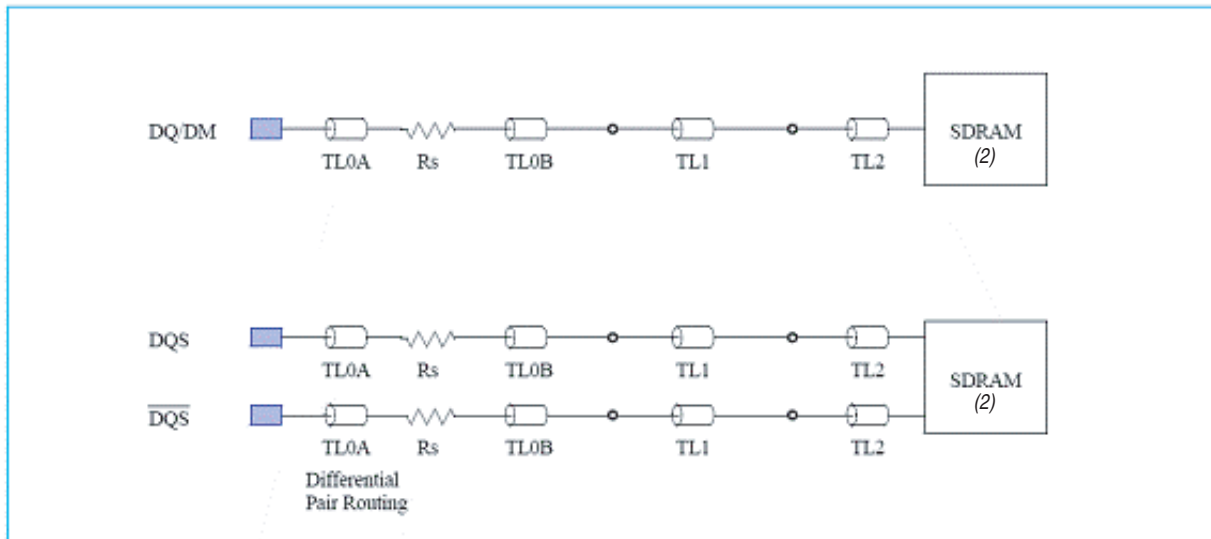
You can implement a DDR3 SDRAM UDIMM interface in several permutations, such as single DIMM or multiple DIMMs, using either single-ranked or dual-ranked UDIMMs. In addition to the UDIMM's form factor, these termination recommendations are also valid for small-outline (SO) DIMMs and MicroDIMMs.

DQS, DQ, and DM for DDR3 SDRAM UDIMM

On a single-ranked DIMM, DQS, and DQ signals are point-to-point signals.

Figure 1-36 shows the net structure for differential DQS and DQ signals. There is an external 15- Ω stub resistor, R_s , on each of the DQS and DQ signals soldered on the DIMM, which helps improve signal quality by dampening reflections from unused slots in a multi-DIMM configuration.

Figure 1-36. DQ and DQS Net Structure for 64-Bit DDR3 SDRAM UDIMM (Note 1)



Notes to Figure 1-36:

- (1) Source: *PC3-6400/PC3-8500/PC3-10600/PC3-12800 DDR3 SDRAM Unbuffered DIMM Design Specification*, July 2007, JEDEC Solid State Technology Association. For clarity of the signal connections in the illustration, the same SDRAM is drawn as two separate SDRAMs.

As mentioned in “Dynamic ODT” on page 1-4, DDR3 SDRAM supports calibrated ODT with different ODT value settings. If you do not enable dynamic ODT, there are three possible ODT settings available for RTT_NORM : 40 Ω , 60 Ω , and 120 Ω . If you enable dynamic ODT, the number of possible ODT settings available for RTT_NORM increases from three to five with the addition of 20 Ω and 30 Ω . Trace impedance on the DIMM and the recommended ODT setting is 60 Ω .

Figure 1-37 shows the simulated write-eye diagram at the DQ0 of a DDR3 SDRAM DIMM using the 60- Ω ODT setting, driven by a Stratix III or Stratix IV FPGA using a calibrated series 50- Ω OCT setting.

Figure 1-37. Simulated Write-Eye Diagram of a DDR3 SDRAM DIMM Using a 60- Ω ODT Setting





Figure 1-38 shows the measured write eye diagram using Altera's Stratix III or Stratix IV memory board.

Figure 1-38. Measured Write-Eye Diagram of a DDR3 SDRAM DIMM Using the 60- Ω ODT Setting

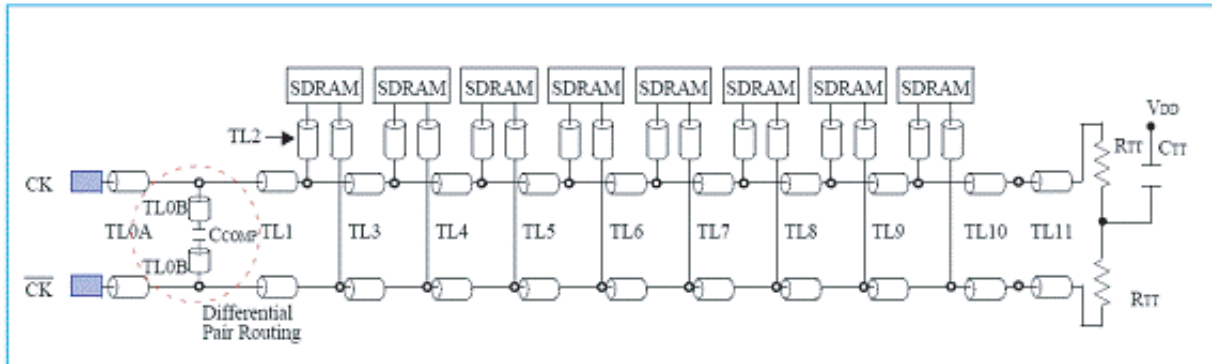


The measured eye diagram correlates well with the simulation. The faint line in the middle of the eye diagram is the effect of the refresh operation during a regular operation. Because these simulations and measurements are based on a narrow set of constraints, you must perform your own board-level simulation to ensure that the chosen ODT setting is right for your setup.

Memory Clocks for DDR3 SDRAM UDIMM

For the DDR3 SDRAM UDIMM, you do not need to place any termination on your board because the memory clocks are already terminated on the DIMM. Figure 1-39 shows the net structure for the memory clocks and the location of the termination resistors, R_{TT} . The value of R_{TT} is $36\ \Omega$, which results in an equivalent differential termination value of $72\ \Omega$. The DDR3 SDRAM DIMM also has a compensation capacitor, C_{COMP} of 2.2 pF, placed between the differential memory clocks to improve signal quality.

Figure 1-39. Clock Net Structure for a 64-Bit DDR3 SDRAM UDIMM (Note 1)



Note to Figure 1-39:

- (1) Source: *PC3-6400/PC3-8500/PC3-10600/PC3-12800 DDR3 SDRAM Unbuffered DIMM Design Specification*, July 2007, JEDEC Solid State Technology Association

From Figure 1-39, you can see that the DDR3 SDRAM clocks are routed in a fly-by topology, as mentioned in “Read and Write Leveling” on page 1-2, resulting in the need for write-and-read leveling. Figure 1-40 shows the HyperLynx simulation of the differential clock seen at the die of the first and last DDR3 SDRAM component on the UDIMM using the 50-Ω OCT setting on the output driver of the Stratix III or Stratix IV FPGA.

Figure 1-40. Differential Memory Clock of a DDR3 SDRAM DIMM at the First and Last Component on the DIMM





Figure 1-40 shows that the memory clock seen at the first DDR3 SDRAM component (the yellow signal) leads the memory clock seen at the last DDR3 SDRAM component (the green signal) by 1.3 ns, which is about 0.69 t_{CK} for a 533 MHz operation.

Commands and Addresses for DDR3 SDRAM UDIMM

Similar to memory clock signals, you do not need to place any termination on your board because the command and address signals are also terminated on the DIMM. **Figure 1-41** shows the net structure for the command and address signals, and the location of the termination resistor, R_{TT} , which has an R_{TT} value of $39\ \Omega$.

Figure 1-41. Command and Address Net Structure for a 64-Bit DDR3 SDRAM Unbuffered DIMM (Note 1)




Note to Figure 1-41:

- (1) Source: *PC3-6400/PC3-8500/PC3-10600/PC3-12800 DDR3 SDRAM Unbuffered DIMM Design Specification*, July 2007, JEDEC Solid State Technology Association

In **Figure 1-41**, observe that the DDR3 SDRAM command and address signals are routed in a fly-by topology, as mentioned in “**Read and Write Leveling**” on page 1-2, resulting in the need for write-and-read leveling.

Figure 1-42 shows the HyperLynx simulation of the command and address signal seen at the die of the first and last DDR3 SDRAM component on the UDIMM, using an OCT setting on the output driver of the Stratix III or Stratix IV FPGA.

Figure 1-42. Command and Address Eye Diagram of a DDR3 SDRAM DIMM at the First and Last DDR3 SDRAM Component at 533 MHz (Note 1)



Note to Figure 1-42:

- (1) The command and address simulation is performed using a bit period of 1.875 ns.

Figure 1-42 shows that the command and address signal seen at the first DDR3 SDRAM component (the green signal) leads the command and address signals seen at the last DDR3 SDRAM component (the red signal) by 1.2 ns, which is $0.64 t_{CK}$ for a 533-MHz operation.

Stratix III, Stratix IV, and Stratix V FPGAs

The following sections review the termination on the single-ranked single DDR3 SDRAM DIMM interface side and investigate the use of different termination features available in Stratix III, Stratix IV, and Stratix V FPGA devices to achieve optimum signal integrity for your DDR3 SDRAM interface.

DQS, DQ, and DM for Stratix III, Stratix IV, and Stratix V FPGA

As mentioned in “Dynamic OCT in Stratix III and Stratix IV Devices” on page 1-5, Stratix III, Stratix IV, and Stratix V FPGAs support the dynamic OCT feature, which switches from series termination to parallel termination depending on the mode of the I/O buffer. Because DQS and DQ are bidirectional signals, DQS and DQ can be both transmitters and receivers. “DQS, DQ, and DM for DDR3 SDRAM UDIMM” on page 1-38 describes the signal quality of DQ, DQS, and DM when the Stratix III, Stratix IV, or Stratix V FPGA device is the transmitter with the I/O buffer set to a 50- Ω series termination.

This section details the condition when the Stratix III, Stratix IV, or Stratix V device is the receiver, the Stratix III, Stratix IV, and Stratix V I/O buffer is set to a 50- Ω parallel termination, and the memory is the transmitter. DM is a unidirectional signal, so the DDR3 SDRAM component is always the receiver.



Refer to “DQS, DQ, and DM for DDR3 SDRAM UDIMM” on page 1-38 for receiver termination recommendations and transmitter output drive strength settings.

Figure 1-43 illustrates the DDR3 SDRAM interface when the Stratix III, Stratix IV, or Stratix V FPGA device is reading from the DDR3 SDRAM using a 50- Ω parallel OCT termination on the Stratix III, Stratix IV, or Stratix V FPGA device, and the DDR3 SDRAM driver output impedance is set to 34 Ω .

Figure 1-43. DDR3 SDRAM Component Driving the Stratix III, Stratix IV, and Stratix V FPGA Device with Parallel 50- Ω OCT Turned On





Figure 1-44 shows the simulation of a read from the DDR3 SDRAM DIMM with a 50- Ω parallel OCT setting on the Stratix III or Stratix IV FPGA device.

Figure 1-44. Read-Eye Diagram of a DDR3 SDRAM DIMM at the Stratix III and Stratix IV FPGA Using a Parallel 50- Ω OCT Setting



Use of the Stratix III, Stratix IV, or Stratix V parallel 50- Ω OCT feature matches receiver impedance with the transmission line characteristic impedance. This eliminates any reflection that causes ringing, and results in a clean eye diagram at the Stratix III, Stratix IV, or Stratix V FPGA.

Memory Clocks for Stratix III, Stratix IV, and Stratix V FPGA

Memory clocks are unidirectional signals. Refer to “[Memory Clocks for DDR3 SDRAM UDIMM](#)” on page 1-40 for receiver termination recommendations and transmitter output drive strength settings.

Commands and Addresses for Stratix III and Stratix IV FPGA

Commands and addresses are unidirectional signals. Refer to “[Commands and Addresses for DDR3 SDRAM UDIMM](#)” on page 1-42 for receiver termination recommendations and transmitter output drive strength settings.

Multi-Rank DDR3 SDRAM Unbuffered DIMM

You can implement a DDR3 SDRAM UDIMM interface in several permutations, such as single DIMM or multiple DIMMs, using either single-ranked or dual-ranked UDIMMs. In addition to the UDIMM’s form factor, these termination recommendations are also valid for small-outline (SO) DIMMs and MicroDIMMs.

Table 1-16 outlines the different permutations of a two-slot DDR3 SDRAM interface and the recommended ODT settings on both the memory and controller when writing to memory.

Table 1-16. DDR3 SDRAM ODT Matrix for Writes (Note 1) and (2)

Slot 1	Slot 2	Write To	Controller OCT (3)	Slot 1		Slot 2	
				Rank 1	Rank 2	Rank 1	Rank 2
DR	DR	Slot 1	Series 50 Ω	120 Ω (4)	ODT off	ODT off	40 Ω (4)
		Slot 2	Series 50 Ω	ODT off	40 Ω (4)	120 Ω (4)	ODT off
SR	SR	Slot 1	Series 50 Ω	120 Ω (4)	Unpopulated	40 Ω (4)	Unpopulated
		Slot 2	Series 50 Ω	40 Ω (4)	Unpopulated	120 Ω (4)	Unpopulated
DR	Empty	Slot 1	Series 50 Ω	120 Ω	ODT off	Unpopulated	Unpopulated
Empty	DR	Slot 2	Series 50 Ω	Unpopulated	Unpopulated	120 Ω	ODT off
SR	Empty	Slot 1	Series 50 Ω	120 Ω	Unpopulated	Unpopulated	Unpopulated
Empty	SR	Slot 2	Series 50 Ω	Unpopulated	Unpopulated	120 Ω	Unpopulated

Notes to Table 1-16:

- (1) SR: single-ranked DIMM; DR: dual-ranked DIMM.
- (2) These recommendations are taken from the *DDR3 ODT and Dynamic ODT* session of the JEDEC DDR3 2007 Conference, Oct 3-4, San Jose, CA.
- (3) The controller in this case is the FPGA.
- (4) Dynamic ODT is required. For example, the ODT of Slot 2 is set to the lower ODT value of 40 Ω when the memory controller is writing to Slot 1, resulting in termination and thus minimizing any reflection from Slot 2. Without dynamic ODT, Slot 2 will not be terminated.

Table 1-17 outlines the different permutations of a two-slot DDR3 SDRAM interface and the recommended ODT settings on both the memory and controller when reading from memory.

Table 1-17. DDR3 SDRAM ODT Matrix for Reads (Note 1) and (2)

Slot 1	Slot 2	Read From	Controller OCT (3)	Slot 1		Slot 2	
				Rank 1	Rank 2	Rank 1	Rank 2
DR	DR	Slot 1	Parallel 50 Ω	ODT off	ODT off	ODT off	40 Ω
		Slot 2	Parallel 50 Ω	ODT off	40 Ω	ODT off	ODT off
SR	SR	Slot 1	Parallel 50 Ω	ODT off	Unpopulated	40 Ω	Unpopulated
		Slot 2	Parallel 50 Ω	40 Ω	Unpopulated	ODT off	Unpopulated
DR	Empty	Slot 1	Parallel 50 Ω	ODT off	ODT off	Unpopulated	Unpopulated
Empty	DR	Slot 2	Parallel 50 Ω	Unpopulated	Unpopulated	ODT off	ODT off
SR	Empty	Slot 1	Parallel 50 Ω	ODT off	Unpopulated	Unpopulated	Unpopulated
Empty	SR	Slot 2	Parallel 50 Ω	Unpopulated	Unpopulated	ODT off	Unpopulated

Notes to Table 1-17:

- (1) SR: single-ranked DIMM; DR: dual-ranked DIMM.
- (2) These recommendations are taken from the *DDR3 ODT and Dynamic ODT* session of the JEDEC DDR3 2007 Conference, Oct 3-4, San Jose, CA.
- (3) The controller in this case is the FPGA. JEDEC typically recommends 60 Ω , but this value assumes that the typical motherboard trace impedance is 60 Ω and that the controller supports this termination. Altera recommends using a 50- Ω parallel OCT when reading from the memory.

DDR3 SDRAM Registered DIMM

The difference between a registered DIMM (RDIMM) and a UDIMM is that the clock, address, and command pins of the RDIMM are registered or buffered on the DIMM before they are distributed to the memory devices. For a controller, each clock, address, or command signal has only one load, which is the register or buffer. In a UDIMM, each controller pin must drive a fly-by wire with multiple loads.

You do not need to terminate the clock, address, and command signals on your board because these signals are terminated at the register. However, because of the register, these signals become point-to-point signals and have improved signal integrity making the drive strength requirements of the FPGA driver pins more relaxed. Similar to the signals in a UDIMM, the DQS, DQ, and DM signals on a RDIMM are not registered. To terminate these signals, refer to “[DQS, DQ, and DM for DDR3 SDRAM UDIMM](#)” on page 1-38.

DDR3 SDRAM Components With Leveling

This section discusses terminations used to achieve optimum performance for designing the DDR3 SDRAM interface using discrete DDR3 SDRAM components.

In addition to using DDR3 SDRAM DIMM to implement your DDR3 SDRAM interface, you can also use DDR3 SDRAM components. However, for applications that have limited board real estate, using DDR3 SDRAM components reduces the need for a DIMM connector and places components closer, resulting in denser layouts.

DDR3 SDRAM Components With or Without Leveling

The DDR3 SDRAM UDIMM is laid out to the JEDEC specification. The JEDEC specification is available from either the JEDEC Organization website (www.JEDEC.org) or from the memory vendors. However, when you are designing the DDR3 SDRAM interface using discrete SDRAM components, you may desire a layout scheme that is different than the DIMM specification. You have the following two options:

- Mimic the standard DDR3 SDRAM DIMM, using a fly-by topology for the memory clocks, address, and command signals. This option needs read and write leveling, so you must use the UniPHY IP with leveling.




For more information on this fly-by configuration, continue reading this chapter.

- Mimic a standard DDR2 SDRAM DIMM, using a balanced (symmetrical) tree-type topology for the memory clocks, address, and command signals. Using this topology results in unwanted stubs on the command, address, and clock, which degrades signal integrity and limits the performance of the DDR3 SDRAM interface.

DQS, DQ, and DM for DDR3 SDRAM Components

When you are laying out the DDR3 SDRAM interface using Stratix III, Stratix IV, or Stratix V devices, Altera recommends that you not include the 15- Ω stub series resistor that is on every DQS, DQ, and DM signal; unless your simulation shows that the absence of this resistor causes extra reflection. Although adding the the 15- Ω stub series resistor may help to maintain constant impedance in some cases, it also slightly reduces signal swing at the receiver. It is unlikely that by removing this resistor the waveform shows a noticeable reflection, but it is your responsibility to prove by simulating your board trace. Therefore, Altera recommends the DQS, DQ, and DM topology shown in Figure 1-45 when the Stratix III, Stratix IV, or Stratix V FPGA is writing to the DDR3 SDRAM.

Figure 1-45. Stratix III, Stratix IV, and Stratix V FPGA Writing to a DDR3 SDRAM Components




When you are using DDR3 SDRAM components, there are no DIMM connectors. This minimizes any impedance discontinuity, resulting in better signal integrity.

Memory Clocks for DDR3 SDRAM Components


When you use DDR3 SDRAM components, you must account for the compensation capacitor and differential termination resistor between the differential memory clocks of the DIMM. Figure 1-46 shows the HyperLynx simulation of the differential clock seen at the die of the first and last DDR3 SDRAM component using a fly-by topology on a board, without the 2.2 pF compensation capacitor using the 50- Ω OCT setting on the output driver of the Stratix III, Stratix IV, or Stratix V FPGA.

Figure 1-46. Differential Memory Clock of a DDR3 SDRAM Component without the Compensation Capacitor at the First and Last Component Using a Fly-by Topology on a Board



Without the compensation capacitor, the memory clocks (the yellow signal) at the first component have significant ringing, whereas, with the compensation capacitor the ringing is dampened. Similarly, the differential termination resistor needs to be included in the design. Depending on your board stackup and layout requirements, you choose your differential termination resistor value. Figure 1-47 shows the HyperLynx simulation of the differential clock seen at the die of the first and last DDR3 SDRAM component using a fly-by topology on a board, and terminated with $100\ \Omega$ instead of the $72\ \Omega$ used in the DIMM.

Figure 1-47. Differential Memory Clock of a DDR3 SDRAM DIMM Terminated with $100\ \Omega$ at the First and Last Component Using a Fly-by Topology on a Board



Terminating with $100\ \Omega$ instead of $72\ \Omega$ results in a slight reduction in peak-to-peak amplitude. To simplify your design, use the terminations outlined in the JEDEC specification for DDR3 SDRAM UDIMM as your guide and perform simulation to ensure that the DDR3 SDRAM UDIMM terminations provide you with optimum signal quality.

In addition to choosing the value of the differential termination, you must consider the trace length of the memory clocks. Altera's DDR3 UniPHY IP currently supports a flight-time skew of no more than $0.69\ t_{CK}$ in between the first and last memory component. If you use Altera's DDR3 UniPHY IP to create your DDR3 SDRAM interface, ensure that the flight-time skew of your memory clocks is not more than $0.69\ t_{CK}$. UniPHY IP also requires that the total skew combination of the clock fly-by skew and DQS skew is less than 1 clock cycle.




Refer to “[Layout Guidelines for DDR3 SDRAM Interface](#)” on page 1-61 for more information about layout guidelines for DDR3 SDRAM components.

Command and Address Signals for DDR3 SDRAM

As with memory clock signals, you must account for the termination resistor on the command and address signals when you use DDR3 SDRAM components. Choose your termination resistor value depending on your board stackup and layout requirements. Figure 1-48 shows the HyperLynx simulation of the command and address seen at the die of the first and last DDR3 SDRAM component using a fly-by topology on a board terminated with $60\ \Omega$ instead of the $39\ \Omega$ used in the DIMM.

Figure 1-48. Command and Address Eye Diagram of a DDR3 SDRAM Component Using Fly-by Topology on a Board at the First and Last DDR3 SDRAM Component at 533 MHz, Terminated with $60\ \Omega$



Terminating with $60\ \Omega$ instead of $39\ \Omega$ results in eye closure in the signal at the first component (the green signal), while there is no effect on the signal at the last component (the red signal). To simplify your design with discrete DDR3 SDRAM components, use the terminations outlined in the JEDEC specification for DDR3 SDRAM UDIMM as your guide, and perform simulation to ensure that the DDR3 SDRAM UDIMM terminations provide you with the optimum signal quality.

As with memory clocks, you must consider the trace length of the command and address signals so that they match the flight-time skew of the memory clocks.

Stratix III, Stratix IV, and Stratix V FPGAs

Stratix III, Stratix IV, or Stratix V FPGA termination settings for DIMM also applies to DDR3 SDRAM component interfaces.

Table 1-18 compares the effects of the series stub resistor on the eye diagram at the Stratix III or Stratix IV FPGA (receiver) when the Stratix III or Stratix IV FPGA is reading from the memory.

Table 1-18. Read-Eye Diagram with and without R_s Using 50- Ω Parallel OCT

ODT	Eye Height (V)	Eye Width (ps)	Overshoot (V)	Undershoot (V)
With R_s	0.70	685	—	—
Without R_s	0.73	724	—	—

Without the 15- Ω stub series resistor to dampen the signal, the signal at the receiver of the Stratix III or Stratix IV FPGA driven by the DDR3 SDRAM component is larger than the signal at the receiver of the Stratix III or Stratix IV FPGA driven by DDR3 SDRAM DIMM (Figure 1-43), and similar to the write-eye diagram in “DQS, DQ, and DM for DDR3 SDRAM Components” on page 1-47.

Drive Strength

Altera’s FPGA products offer numerous drive strength settings, allowing you to optimize your board designs to achieve the best signal quality. This section focuses on the most commonly used drive strength settings of 8 mA and 16 mA, as recommended by JEDEC for Class I and Class II termination schemes.



You are not restricted to using only these drive strength settings for your board designs. You should perform simulations using I/O models available from Altera and memory vendors to ensure that you use the proper drive strength setting to achieve optimum signal integrity.

How Strong is Strong Enough?

Figure 1-20 on page 1-22 shows a signal probed at the DDR2 SDRAM DIMM (receiver) of a far-end series-terminated transmission line when the FPGA writes to the DDR2 SDRAM DIMM using a drive strength setting of 16 mA. The resulting signal quality on the receiver shows excessive over- and undershoot. To reduce the over- and undershoot, you can reduce the drive strength setting on the FPGA from 16 mA to 8 mA. Figure 1-49 shows the simulation and measurement of the FPGA with a drive strength setting of 8 mA driving a no-parallel termination transmission line.

Figure 1-49. HyperLynx Simulation and Measurement, FPGA Writing to Memory




Table 1-19 compares the signals at the DDR2 SDRAM DIMM with no-parallel termination and memory-side series resistors when the FPGA is writing to the memory with 8-mA and 16-mA drive strength settings.

Table 1-19. Simulation and Board Measurement Results for 8 mA and 16 mA Drive Strength Settings

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)
8-mA Drive Strength Setting				
Simulation	1.48	1.71	0.24	0.35
Board Measurement	1.10	1.24	0.24	0.50
16-mA Drive Strength Setting				
Simulation	1.66	1.10	0.90	0.80
Board Measurements	1.25	0.60	1.10	1.08

With a lower strength drive setting, the overall signal quality is improved. The eye width is reduced, but the eye height is significantly larger with a lower drive strength and the over- and undershoot is reduced dramatically.

To improve the signal quality further, you should use 50- Ω on-chip series termination in place of an 8mA drive strength and 25- Ω on-chip series termination in place of a 16 mA drive strength. Refer to “On-Chip Termination (Non-Dynamic)” on page 1-17 for simulation and board measurements.

The drive strength setting is highly dependent on the termination scheme, so it is critical that you perform pre- and post-layout board-level simulations to determine the proper drive strength settings.

System Loading

You can use memory in a variety of forms, such as individual components or multiple DIMMs, resulting in different loading seen by the FPGA. This section describes the effect on signal quality when interfacing memories in component, dual rank, and dual DIMMs format.

Component Versus DIMM

When using discrete DDR2 SDRAM components, the additional loading from the DDR2 SDRAM DIMM connector is eliminated and the memory-side series resistor on the DDR2 SDRAM DIMM is no longer there. You must decide if the memory-side series resistor near the DDR2 SDRAM is required.

FPGA Writing to Memory

Figure 1-50 shows the Class II termination scheme without the memory-side series resistor when the FPGA is writing to the memory in the component format.

Figure 1-50. Class II Termination Scheme without Memory-Side Series Resistor




Figure 1-51 shows the simulation and measurement results of the signal seen at a DDR2 SDRAM component of a Class II termination scheme without the DIMM connector and the memory-side series resistor. The FPGA is writing to the memory with a 16-mA drive strength setting.

Figure 1-51. HyperLynx Simulation and Measurement of the Signal, FPGA Writing to Memory

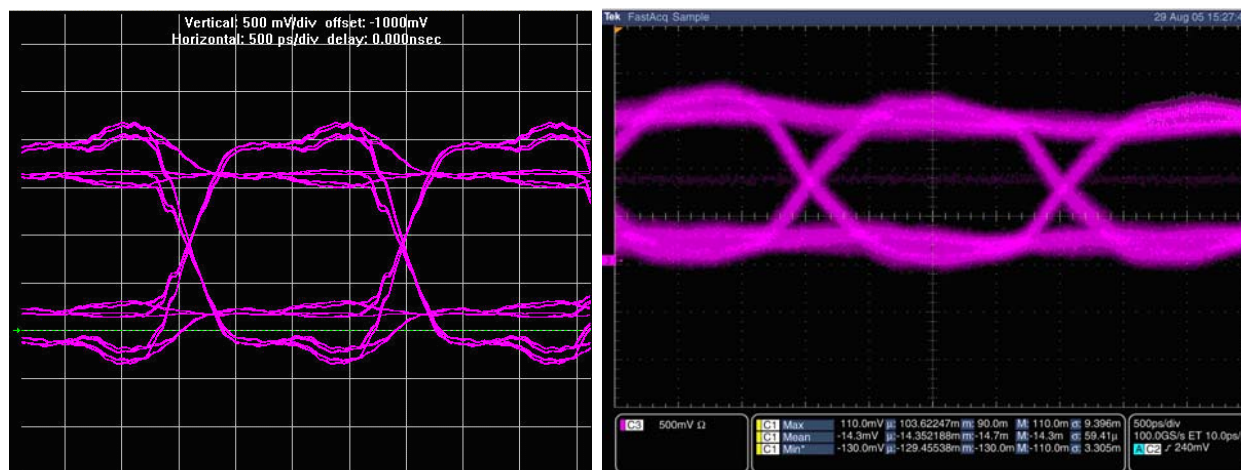


Table 1-20 compares the signal for a single rank DDR2 SDRAM DIMM and a single DDR2 SDRAM component in a Class II termination scheme when the FPGA is writing to the memory.

Table 1-20. Simulation and Board Measurement Results for Single Rank DDR2 SDRAM DIMM and Single DDR2 SDRAM Component (Note 1), (2)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Single DDR2 SDRAM Component						
Simulation	1.79	1.15	0.39	0.33	3.90	3.43
Measurement	1.43	0.96	0.10	0.13	1.43	1.43
Single Rank DDR2 SDRAM DIMM						
Simulation	1.65	0.86	N/A	N/A	1.71	1.95
Measurement	1.36	0.41	N/A	N/A	1.56	1.56

Note to Table 1-20:

- (1) The drive strength on the FPGA is set to Class II 16 mA.
- (2) N/A is not applicable.

The overall signal quality is comparable between the single rank DDR2 SDRAM DIMM and the single DDR2 SDRAM component, but the elimination of the DIMM connector and memory-side series resistor results in a more than 50% improvement in the eye height.

FPGA Reading from Memory

Figure 1-52 shows the Class II termination scheme without the memory-side series resistor when the FPGA is reading from memory. Without the memory-side series resistor, the memory driver has less loading to drive the Class II termination. Compare this result to the result of the DDR2 SDRAM DIMM described in “FPGA Reading from Memory” on page 1-33 where the memory-side series resistor is on the DIMM.

Figure 1-52. Class II Termination Scheme without Memory-Side Series Resistor




Figure 1-53 shows the simulation and measurement results of the signal seen at the FPGA. The FPGA reads from memory without the source-series resistor near the DDR2 SDRAM component on a Class II-terminated transmission line. The FPGA reads from memory with a full drive strength setting.

Figure 1-53. HyperLynx Simulation and Measurement, FPGA Reading from the DDR2 SDRAM Component

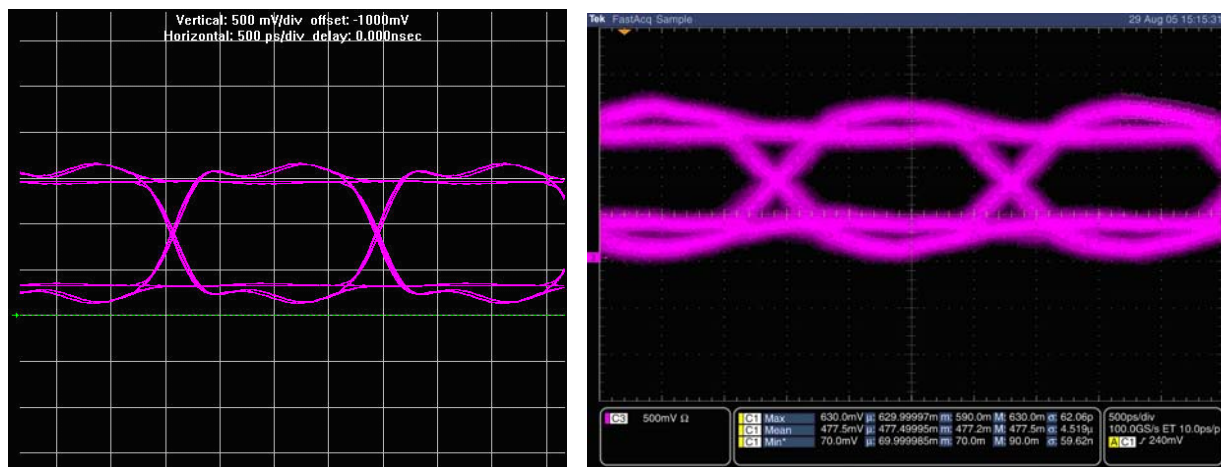


Table 1–21 compares the signal at a single rank DDR2 SDRAM DIMM and a single DDR2 SDRAM component of a Class II termination scheme. The FPGA is reading from memory with a full drive strength setting.

Table 1–21. Simulation and Board Measurement Results of Single Rank DDR2 SDRAM DIMM and DDR2 SDRAM Component (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Single DDR2 SDRAM Component						
Simulation	1.79	1.06	N/A	N/A	2.48	3.03
Measurement	1.36	0.63	0.13	0.00	1.79	1.14
Single Rank DDR2 SDRAM DIMM						
Simulation	1.73	0.76	N/A	N/A	1.71	1.95
Measurement	1.28	0.43	N/A	N/A	0.93	0.86

Note to Table 1–21:

(1) N/A is not applicable.

The effect of eliminating the DIMM connector and memory-side series resistor is evident in the improvement in the eye height.

Single- Versus Dual-Rank DIMM

DDR2 SDRAM DIMMs are available in either single- or dual-rank DIMM. Single-rank DIMMs are DIMMs with DDR2 SDRAM memory components on one side of the DIMM. Higher-density DIMMs are available as dual-rank, which has DDR2 SDRAM memory components on both sides of the DIMM. With the dual-rank DIMM configuration, the loading is twice that of a single-rank DIMM. Depending on the

board design, you must adjust the drive strength setting on the memory controller to account for this increase in loading. Figure 1-54 shows the simulation result of the signal seen at a dual rank DDR2 SDRAM DIMM. The simulation uses Class II termination with a memory-side series resistor transmission line. The FPGA uses a 16-mA drive strength setting.

Figure 1-54. HyperLynx Simulation with a 16-mA Drive Strength Setting on the FPGA




Table 1-22 compares the signals at a single- and dual-rank DDR2 SDRAM DIMM of a Class II and far-end source-series termination when the FPGA is writing to the memory with a 16-mA drive strength setting.

Table 1-22. Simulation Results of Single- and Dual-Rank DDR2 SDRAM DIMM (Note 1)

	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual Rank DDR2 SDRAM DIMM						
Simulation	1.34	1.27	0.12	0.12	0.99	0.94
Single Rank DDR2 SDRAM DIMM						
Simulation	1.65	1.27	0.10	0.10	1.71	1.95

Note to Table 1-22:

(1) The drive strength on the FPGA is set to Class II 16 mA.

In a dual-rank DDR2 SDRAM DIMM, the additional loading leads to a slower edge rate, which affects the eye width. The slower edge rate leads to the degradation of the setup and hold time required by the memory as well, which must be taken into consideration during the analysis of the timing for the interface. The overall signal quality remains comparable, but eye width is reduced in the dual-rank DIMM. This reduction in eye width leads to a smaller data capture window that must be taken into account when performing timing analysis for the memory interface.

Single DIMM Versus Multiple DIMMs

Some applications, such as packet buffering, require deeper memory, making a single DIMM interface insufficient. If you use a multiple DIMM configuration to increase memory depth, the memory controller is required to interface with multiple data strobes and the data lines instead of the point-to-point interface in a single DIMM configuration. This results in heavier loading on the interface, which can potentially impact the overall performance of the memory interface.



For detailed information about a multiple DIMM DDR2 SDRAM memory interface, refer to [Chapter 2, Dual-DIMM DDR2 and DDR3 SDRAM Interface Termination, Drive Strength, Loading, and Board Layout Guidelines](#).

Design Layout Guidelines

This section discusses general layout guidelines for designing your DDR2 and DDR3 SDRAM interfaces. These layout guidelines help you plan your board layout, but are not meant as strict rules that must be adhered to. Altera recommends that you perform your own board-level simulations to ensure that the layout you choose for your board allows you to achieve your desired performance.

These layout guidelines are for both ALTMEMPHY- and UniPHY-based IP designs, unless specified otherwise.



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



The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristic of the interface. They do not include any margin for crosstalk.

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

Layout Guidelines for DDR2 SDRAM Interface

[Table 1-23](#) summarizes DDR2 SDRAM layout guidelines.

 These layout guidelines also apply to DDR3 SDRAM without leveling interfaces.

Table 1–23. DDR2 SDRAM Layout Guidelines (Part 1 of 3) (Note 1)

Parameter	Guidelines
DIMMs	If you consider a normal DDR2 unbuffered, unregistered DIMM, essentially you are planning to perform the DIMM routing directly on your PCB. Therefore, each address and control pin routes from the FPGA (single pin) to all memory devices must be on the same side of the FPGA.
Impedance	<ul style="list-style-type: none"> ■ All signal planes must be 50-60-Ω, single-ended, ±10% ■ All signal planes must be 100Ω, differential ±10% ■ All unused via pads must be removed, because they cause unwanted capacitance
Decoupling Parameter	<ul style="list-style-type: none"> ■ Use 0.1 μF in 0402 size to minimize inductance ■ Make V_{TT} voltage decoupling close to pull-up resistors ■ Connect decoupling caps between V_{TT} and ground ■ Use a 0.1μF cap for every other V_{TT} pin and 0.01μF cap for every V_{DD} and V_{DDQ} pin
Power	<ul style="list-style-type: none"> ■ Route GND, 1.8 V as planes ■ Route V_{CCIO} for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation ■ Route V_{TT} as islands or 250-mil (6.35-mm) power traces ■ Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Altera recommend that signals from the same net group always be routed on the same layer.</p> <ul style="list-style-type: none"> ■ Use 45° angles (<i>not</i> 90° corners) ■ Avoid T-Junctions for critical nets or clocks ■ Avoid T-junctions greater than 250 mils (6.35 mm) ■ Disallow signals across split planes ■ Restrict routing other signals close to system reset signals ■ Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks ■ All data, address, and command signals must have matched length traces ± 50 ps (±0.250 inches or 6.35 mm) ■ All signals within a given Byte Lane Group should be matched length with maximum deviation of ±10 ps or approximately ±0.050 inches (1.27 mm) and routed in the same layer.
Clock Routing	<ul style="list-style-type: none"> ■ Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm) ■ These signals should maintain a 10-mil (0.254 mm) spacing from other nets ■ Clocks should maintain a length-matching between clock pairs of ±5 ps or approximately ±25 mils (0.635 mm) ■ Differential clocks should maintain a length-matching between P and N signals of ±2 ps or approximately ±10 mils (0.254 mm), routed in parallel ■ Space between different pairs should be at least three times the space between the differential pairs and must be routed differentially (5-mil trace, 10-15 mil space on centers) and equal to or up to 100 mils (2.54 mm) longer than signals in the Address/Command Group

Table 1–23. DDR2 SDRAM Layout Guidelines (Part 2 of 3) (Note 1)

Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> ■ Unbuffered address and command lines are more susceptible to cross-talk and are generally noisier than buffered address or command lines. Therefore, un-buffered address and command signals should be routed on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. ■ Do not route differential clock (CK) and clock enable (CKE) signals close to address signals.
External Memory Routing Rules	<ul style="list-style-type: none"> ■ Keep the distance from the pin on the DDR2 DIMM or component to the termination resistor pack (V_{TT}) to less than 500 mils for DQS [x] Data Groups. ■ Keep the distance from the pin on the DDR2 DIMM or component to the termination resistor pack (V_{TT}) to less than 1000 mils for the ADR_CMD_CTL Address Group. ■ Parallelism rules for the DQS [x] Data Groups are as follows: <ul style="list-style-type: none"> ■ 4 mils for parallel runs < 0.1 inch (approximately 1× spacing relative to plane distance) ■ 5 mils for parallel runs < 0.5 inch (approximately 1× spacing relative to plane distance) ■ 10 mils for parallel runs between 0.5 and 1.0 inches (approximately 2× spacing relative to plane distance) ■ 15 mils for parallel runs between 1.0 and 6.0 inch (approximately 3× spacing relative to plane distance) ■ Parallelism rules for the ADR_CMD_CTL group and CLOCKS group are as follows: <ul style="list-style-type: none"> ■ 4 mils for parallel runs < 0.1 inch (approximately 1× spacing relative to plane distance) ■ 10 mils for parallel runs < 0.5 inch (approximately 2× spacing relative to plane distance) ■ 15 mils for parallel runs between 0.5 and 1.0 inches (approximately 3× spacing relative to plane distance) ■ 20 mils for parallel runs between 1.0 and 6.0 inches (approximately 4× spacing relative to plane distance) ■ All signals are to maintain a 20-mil separation from other, non-related nets. ■ All signals must have a total length of < 6 inches.
Termination Rules	<ul style="list-style-type: none"> ■ When pull-up resistors are used, fly-by termination configuration is recommended. Fly-by helps reduce stub reflection issues. ■ Pull-ups should be within 0.5 to no more than 1 inch. ■ Pull up is typically 56 Ω. ■ If using resistor networks: <ul style="list-style-type: none"> ■ Do not share R-pack series resistors between address/command and data lines (DQ, DQS, and DM) to eliminate crosstalk within pack. ■ Series and pull up tolerances are 1–2%. ■ Series resistors are typically 10 to 20 Ω. ■ Address and control series resistor typically at the FPGA end of the link. ■ DM, DQS, DQ series resistor typically at the memory end of the link (or just before the first DIMM). ■ If termination resistor packs are used: <ul style="list-style-type: none"> ■ The distance to your memory device should be less than 750 mils. ■ The distance from your Altera's FPGA device should be less than 1250 mils.

Table 1–23. DDR2 SDRAM Layout Guidelines (Part 3 of 3) (Note 1)

Parameter	Guidelines
Quartus II Software Settings for Board Layout	<ul style="list-style-type: none"> ■ To perform timing analyses on board and I/O buffers, use third party simulation tool to simulate all timing information such as skew, ISI, crosstalk, and type the simulation result into the UniPHY board setting panel. ■ Do not use advanced I/O timing model (AIOT) or board trace model unless you do not have access to any third party tool. AIOT provides reasonable accuracy but tools like HyperLynx provides better result. In operations with higher frequency, it is crucial to properly simulate all signal integrity related uncertainties. ■ The Quartus II software does timing check to find how fast the controller issues a write command after a read command, which limits the maximum length of the DQ/DQS trace. Turn on the bus turnaround timing option and make sure the margin is positive before board fabrication. Functional failure happens if the margin is more than 0.

Notes to Table 1–23:

(1) For point-to-point and DIMM interface designs, refer to the Micron website, www.micron.com.

Layout Guidelines for DDR3 SDRAM Interface

Table 1–24 summarizes DDR3 SDRAM layout guidelines.



These layout guidelines are specifically for DDR3 UDIMMs and interfaces with discreet components using fly-by networks clocked at 1,066 MHz.

Table 1–24. DDR3 SDRAM UDIMM Layout Guidelines (Part 1 of 5) (Note 1)

Parameter	Guidelines
DIMMs	If you consider a normal DDR3 unbuffered, unregistered DIMM, essentially you are planning to perform the DIMM routing directly on your PCB. Therefore, each address and control pin routes from the FPGA (single pin) to all memory devices must be on the same side of the FPGA.
Impedance	<ul style="list-style-type: none"> ■ All signal planes must be 50 Ω, single-ended, $\pm 10\%$. ■ All signal planes must be 100 Ω, differential $\pm 10\%$. ■ All unused via pads must be removed, because they cause unwanted capacitance.
Decoupling Parameter	<ul style="list-style-type: none"> ■ Use 0.1 μF in 0402 size to minimize inductance. ■ Make V_{TT} voltage decoupling close to the DDR3 SDRAM components and pull-up resistors. ■ Connect decoupling caps between V_{TT} and V_{DD} using a 0.1 μF cap for every other V_{TT} pin. ■ Use a 0.1 μF cap and 0.01 μF cap for every V_{DDQ} pin.
Power	<ul style="list-style-type: none"> ■ Route GND, 1.5 V and 0.75 V as planes. ■ Route V_{CCIO} for memories in a single split plane with at least a 20-mil (0.020 inches, or 0.508 mm) gap of separation. ■ Route V_{TT} as islands or 250-mil (6.35-mm) power traces. ■ Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces.

Table 1–24. DDR3 SDRAM UDIMM Layout Guidelines (Part 2 of 5) (Note 1)

Parameter	Guidelines
Maximum Trace Length (2)	<ul style="list-style-type: none"> ■ Maximum trace length for all signals from FPGA to the first DIMM slot is 4.5 inches. ■ Maximum trace length for all signals from DIMM slot to DIMM slot is 0.425 inches. ■ When interface with multiple DDR3 SDRAM components, maximum trace length for address, command, control, and clock from FPGA to the first component must not be more than 7 inches. ■ Maximum trace length for DQ, DQS, DQS#, and DM from FPGA to the first component is 5 inches. ■ Even though there are no hard requirements for minimum trace length, you need to simulate the trace to ensure the signal integrity.
General Routing	<p>All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propagation variance, Altera recommend that you route signals from the same net group on the same layer.</p> <ul style="list-style-type: none"> ■ Use 45° angles (<i>not</i> 90° corners). ■ Disallow critical signals across split planes. ■ Route over appropriate V_{CC} and GND planes. ■ Keep signal routing layers close to GND and power planes. ■ Avoid routing memory signals closer than 0.025 inch (0.635 mm) to memory clocks.
Clock Routing	<ul style="list-style-type: none"> ■ Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm). The maximum length of the first SDRAM to the last SDRAM must not be more than 5 inches (127 mm) or $0.69 t_{CK}$ at 1.066 GHz ■ These signals should maintain the following spacings: <ul style="list-style-type: none"> ■ 10-mil (0.254 mm) spacing for parallel runs less than 0.5 inches or 2x trace-to-plane distance. ■ 15-mil spacing for parallel runs between 0.5 and 1 inches or 3x trace-to-plane distance. ■ 20-mil spacing for parallel runs between 1 and 6 inches or 4x trace-to-plane distance. ■ Clocks should maintain a length-matching between clock pairs of ± 5 ps or approximately ± 25 mils (0.635 mm). ■ Differential clocks should maintain a length-matching between positive (p) and negative (n) signals of ± 2 ps or approximately ± 10 mils (0.254 mm), routed in parallel. ■ Space between different pairs should be at least two times the trace width of the differential pair to minimize loss and maximize interconnect density. ■ Route differential clocks differentially (5-mil trace, 10-15 mil space on centers) and equal to length of signals in the Address/Command Group. ■ To avoid mismatched transmission line to via, Altera recommends that you use Ground Signal Signal Ground (GSSG) topology for your clock pattern—GND CLKP CKLN GND.

Table 1–24. DDR3 SDRAM UDIMM Layout Guidelines (Part 3 of 5) (Note 1)

Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> ■ Route address and command signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not be more than 5 inches (127 mm) or $0.69 t_{CK}$ at 1.066 GHz. For different DIMM configurations, check the appropriate JEDEC specifications. ■ UDIMMs are more susceptible to cross-talk and are generally noisier than buffered DIMMs. Therefore, route address and command signals of UDIMMs on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. Make sure that each net maintains the same consecutive order. ■ Do not route differential clock (CK) and clock enable (CKE) signals close to address signals. ■ Route all addresses and commands to match the clock signals to within ± 25 ps or approximately ± 125 mil (± 3.175 mm) to each discrete memory component. Refer to Figure 1–55.

Table 1–24. DDR3 SDRAM UDIMM Layout Guidelines (Part 4 of 5) (Note 1)

Parameter	Guidelines
External Memory Routing Rules	<ul style="list-style-type: none"> ■ Match in length all DQ, DQS, and DM signals within a given byte-lane group with a maximum deviation of ± 10 ps or approximately ± 50 mils (± 1.27 mm). ■ Ensure to route all DQ, DQS, and DM signals within a given byte-lane group on the same layer to avoid layer to layer transmission velocity differences, which otherwise increase the skew within the group. ■ For ALTMEMPHY-based interfaces, keep the maximum byte-lane group-to-byte group matched length deviation to ± 150 ps or ± 0.8 inches (± 20 mm). ■ Parallelism rules for address and command and clock signals are as follows: <ul style="list-style-type: none"> ■ 4 mils for parallel runs < 0.1 inch (approximately $1\times$ spacing relative to plane distance) ■ 10 mils for parallel runs < 0.5 inch (approximately $2\times$ spacing relative to plane distance) ■ 15 mils for parallel runs between 0.5 and 1.0 inches (approximately $3\times$ spacing relative to plane distance) ■ 20 mils for parallel runs between 1.0 and 6.0 inches (approximately $4\times$ spacing relative to plane distance) ■ Parallelism rules for all other signals are as follows: <ul style="list-style-type: none"> ■ 5 mils for parallel runs < 0.5 inch (approximately $1\times$ spacing relative to plane distance) ■ 10 mils for parallel runs between 0.5 and 1.0 inches (approximately $2\times$ spacing relative to plane distance) ■ 15 mils for parallel runs between 1.0 and 6.0 inch (approximately $3\times$ spacing relative to plane distance) ■ Do not use DDR3 deskew to correct for more than 20 ps of DQ group skew. The skew algorithm only removes the following possible uncertainties: <ul style="list-style-type: none"> ■ Minimum and maximum die IOE skew or delay mismatch ■ Minimum and maximum device package skew or mismatch ■ Board delay mismatch of 20 ps ■ Memory component DQ skew mismatch ■ Increasing any of these four parameters runs the risk of the deskew algorithm limiting, failing to correct for the total observed system skew. If the algorithm cannot compensate without limiting the correction, timing analysis shows reduced margins. ■ All the trace length matching requirements are from the FPGA package ball to DDR3 package ball, which means you have to take into account trace mismatching on different DIMM raw cards. ■ For UniPHY-based interfaces, the timing between the DQS and clock signals on each device calibrates dynamically when you enable leveling to meet t_{DQS}. To make sure the skew is not too large for the leveling circuit's capability, refer to Figure 1–56 and follow these rules: <ul style="list-style-type: none"> ■ Propagation delay of clock signal must not be shorter than propagation delay of DSQ signal at every device: $(CK_i - CK) - DQS_i > 0; 0 < i < \text{number of components} - 1$ ■ Total skew of CLK and DQS signal between groups is less than one clock cycle: $(CK_i - CK + DQS_i) \text{ max} - (CK_i - CK + DQS_i) \text{ min} < 1 \times t_{CK}$

Table 1–24. DDR3 SDRAM UDIMM Layout Guidelines (Part 5 of 5) (Note 1)

Parameter	Guidelines
Termination Rules	<ul style="list-style-type: none"> ■ When using DIMMs, you have no concerns about terminations on memory clocks, addresses, and commands. ■ If you are using components, use an external parallel termination of $40\ \Omega$ to V_{TT} at the end of the fly-by daisy chain topology on the addresses and commands. ■ For memory clocks, use an external parallel termination of $75\ \Omega$ differential at the end of the fly-by daisy chain topology on the memory clocks. Fly-by daisy chain topology helps reduce stub reflection issues. ■ Keep the length of the traces to the termination to within 0.5 inch (14 mm). ■ Use resistors with tolerances of 1 to 2%.
Quartus II Software Settings for Board Layout	<ul style="list-style-type: none"> ■ To perform timing analyses on board and I/O buffers, use third party simulation tool to simulate all timing information such as skew, ISI, crosstalk, and type the simulation result into the UniPHY board setting panel. ■ Do not use advanced I/O timing model (AIOT) or board trace model unless you do not have access to any third party tool. AIOT provides reasonable accuracy but tools like HyperLynx provides better result. In 1,066-MHz operation, it is crucial to properly simulate all signal integrity related uncertainties. ■ The Quartus II software does timing check to find how fast the controller issues a write command after a read command, which limits the maximum length of the DQ/DQS trace. Turn on the bus turnaround timing option and make sure the margin is positive before board fabrication. Functional failure happens if the margin is more than 0.

Notes to Table 1–23:

- (1) For point-to-point and DIMM interface designs, refer to the Micron website, www.micron.com.
- (2) For better efficiency, the UniPHY IP requires faster turnarounds from read commands to write.

Figure 1–55 shows the DDR3 SDRAM component routing guidelines for address and command signals.

Figure 1–55. DDR3 SDRAM Component Address and Command Routing Guidelines

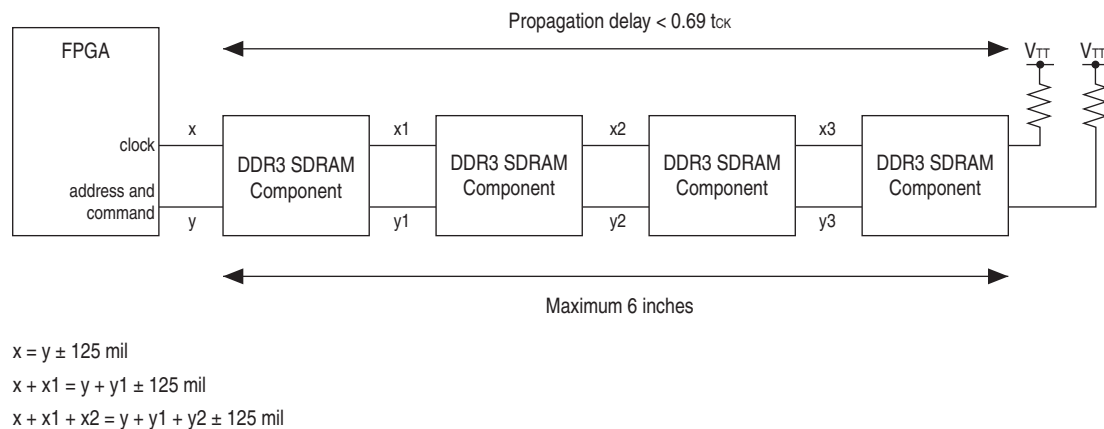
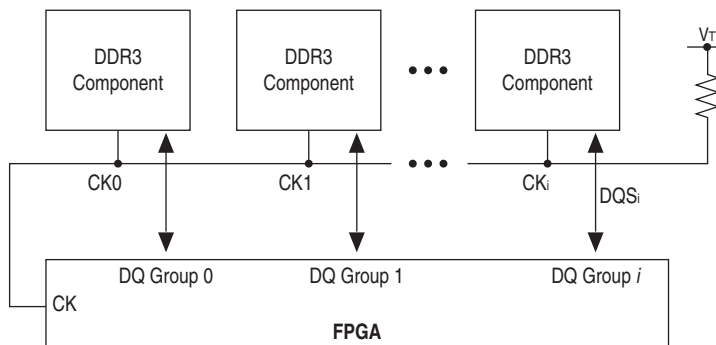


Figure 1-56 shows the delay requirements to align DQS and clock signals.

Figure 1-56. Delaying DQS Signal to Align DQS and Clock



(CK_i — CK) = Clock signal propagation delay to device *i*

DQS_i = DQ/DQS signals propagation delay to group *i*

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

This section discusses the different ways to lay out a wider DDR3 SDRAM interface to the FPGA. Choose the topology based on board trace simulation and the timing budget of your system.

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

When you lay out for a wider interface, all rules and constraints discussed in the previous sections still apply. The DQS, DQ, and DM signals are point-to-point, and all the same rules discussed in “Design Layout Guidelines” on page 1-58 apply.

The main challenge for the design of the fly-by network topology for the clock, command, and address signals is to avoid signal integrity issues, and to make sure you route the DQS, DQ, and DM signals with the chosen topology.

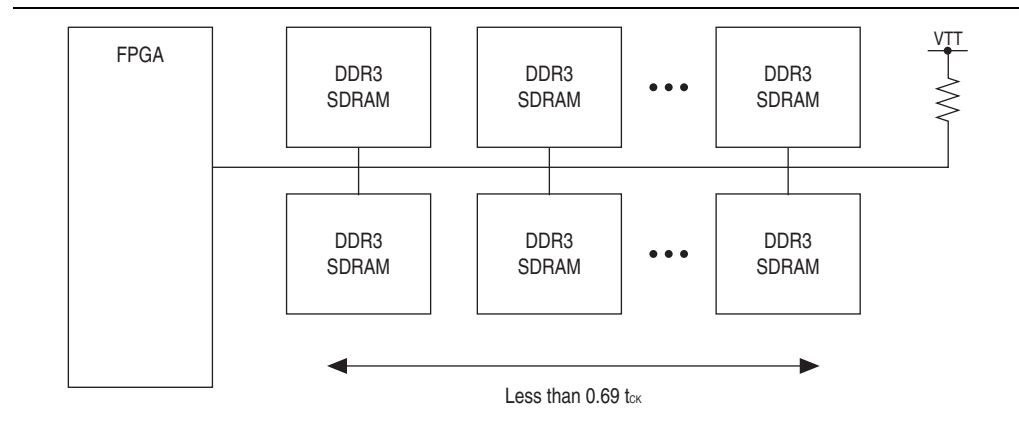
Fly-By Network Design for Clock, Command, and Address Signals

As described in “DDR3 SDRAM Components With Leveling” on page 1-46, the UniPHY IP requires the flight-time skew between the first DDR3 SDRAM component and the last DDR3 SDRAM component to be less than $0.69 t_{CK}$ for memory clocks. This constraint limits the number of components you can have for each fly-by network.

If you design with discrete components, you can choose to use one or more fly-by networks for the clock, command, and address signals.

Figure 1-57 shows an example of a single fly-by network topology.

Figure 1-57. Single Fly-By Network Topology

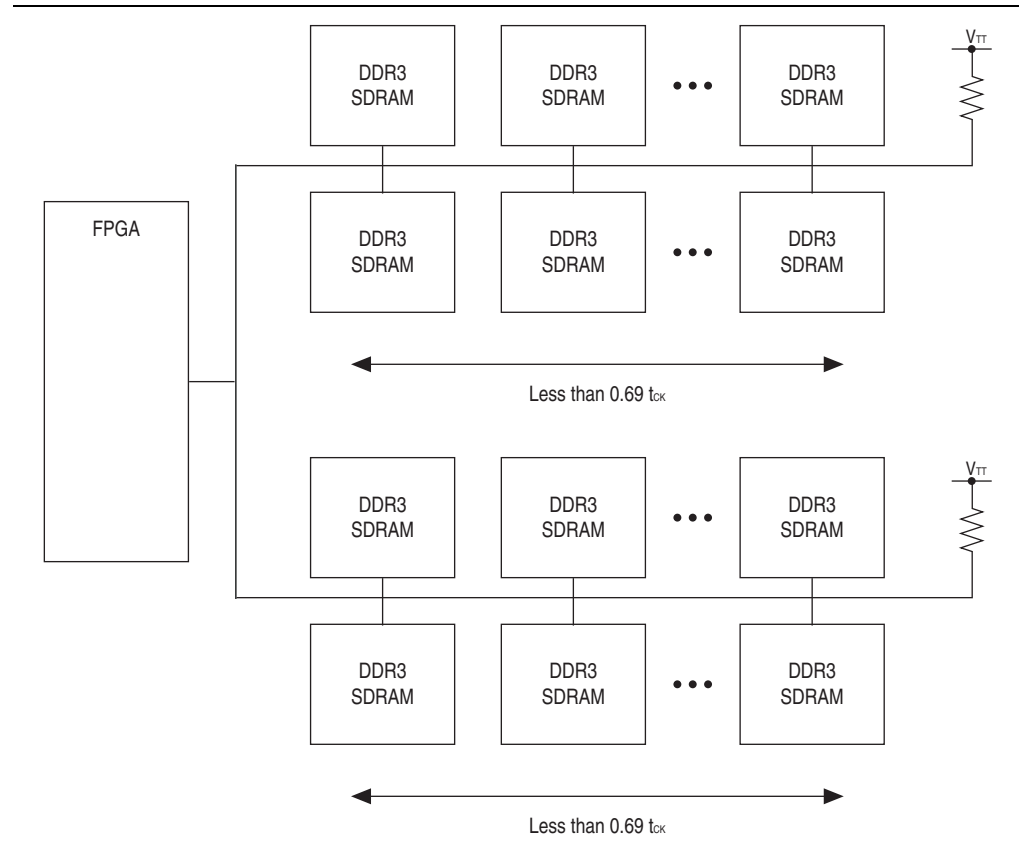


Every DDR3 SDRAM component connected to the signal is a small load that causes discontinuity and degrades the signal. When using a single fly-by network topology, to minimize signal distortion, follow these guidelines:

- Use ×16 device instead ×4 or ×8 to minimize the number of devices connected to the trace.
- Keep the stubs as short as possible.
- Even with added loads from additional components, keep the total trace length short; keep the distance between the FPGA and the first DDR3 SDRAM component less than 5 inches.
- Simulate clock signals to ensure a decent waveform.

Figure 1-58 shows an example of a double fly-by network topology. This topology is not rigid but you can use it as an alternative option. The advantage of using this topology is that you can have more DDR3 SDRAM components in a system without violating the $0.69 t_{CK}$ rule. However, as the signals branch out, the components still create discontinuity.


Figure 1-58. Double Fly-By Network Topology



You need to carry out some simulations to find the location of the split, and the best impedance for the traces before and after the split.

Figure 1-59 shows a way to minimize the discontinuity effect. In this example, keep TL2 and TL3 matches in length. Keep TL1 longer than TL2 and TL3, so that it is easier to route all the signals during layout.

Figure 1-59. Minimizing Discontinuity Effect



You can also consider using a DIMM on each branch to replace the components. Because the trace impedance on the DIMM card is $40\ \Omega$ to $60\ \Omega$, perform a board trace simulation to control the reflection to within the level your system can tolerate.

By using the new features of the DDR3 SDRAM controller with UniPHY and the Stratix III, Stratix IV, or Stratix V devices, you simplify your design process. Using the fly-by daisy chain topology increases the complexity of the datapath and controller design to achieve leveling, but also greatly improves performance and eases board layout for DDR3 SDRAM.

You can also use the DDR3 SDRAM components without leveling in a design if it may result in a more optimal solution, or use with devices that support the required electrical interface standard, but do not support the required read and write leveling functionality.

This chapter describes guidelines for implementing dual unbuffered DIMM (UDIMM) DDR2 and DDR3 SDRAM interfaces. This chapter discusses the impact on signal integrity of the data signal with the following conditions in a dual-DIMM configuration:


- Populating just one slot versus populating both slots
- Populating slot 1 versus slot 2 when only one DIMM is used
- On-die termination (ODT) setting of 75 Ω versus an ODT setting of 150 Ω

 For detailed information about a single-DIMM DDR2 SDRAM interface, refer to [Chapter 1, DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines](#).

DDR2 SDRAM

This section describes guidelines for implementing a dual slot unbuffered DDR2 SDRAM interface, operating at up to 400-MHz and 800-Mbps data rates. [Figure 2-1](#) shows a typical DQS, DQ, and DM signal topology for a dual-DIMM interface configuration using the ODT feature of the DDR2 SDRAM components.

Figure 2-1. Dual-DIMM DDR2 SDRAM Interface Configuration *(Note 1)*



Note to Figure 2-1:

- (1) The parallel termination resistor $R_T = 54 \Omega$ to V_{TT} at the FPGA end of the line is optional for devices that support dynamic on-chip termination (OCT).


The simulations in this section use a Stratix II device-based board. Because of limitations of this FPGA device family, simulations are limited to 266 MHz and 533 Mbps so that comparison to actual hardware results can be directly made.

Stratix II High Speed Board

To properly study the dual-DIMM DDR2 SDRAM interface, the simulation and measurement setup evaluated in the following analysis features a Stratix II FPGA interfacing with two 267-MHz DDR2 SDRAM UDIMMs. This DDR2 SDRAM interface is built on the Stratix II High-Speed High-Density Board (Figure 2-2).

- For more information about the Stratix II High-Speed High-Density Board, contact your Altera representative.

Figure 2-2. Stratix II High-Speed Board with Dual-DIMM DDR2 SDRAM Interface

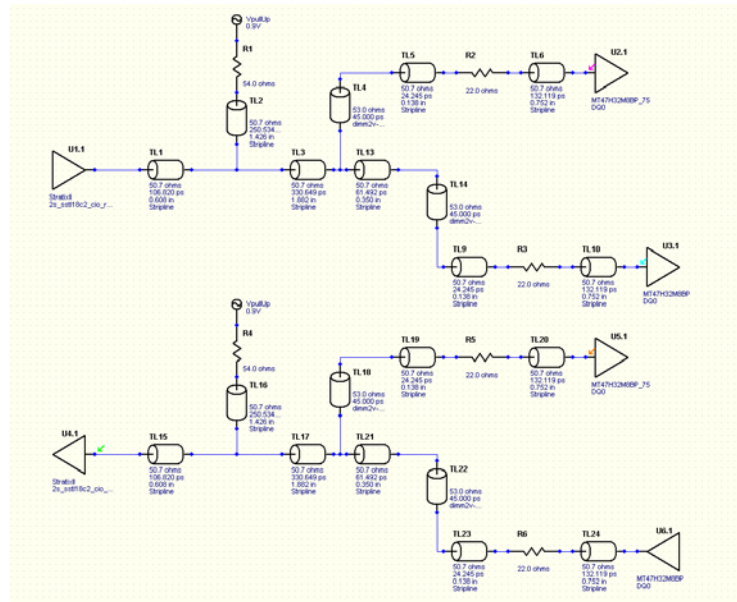


The Stratix II High-Speed Board uses a Stratix II 2S90F1508 device. For DQS, DQ, and DM signals, the board is designed without external parallel termination resistors near the DDR2 SDRAM DIMMs, to take advantage of the ODT feature of the DDR2 SDRAM components. Stratix II FPGA devices are not equipped with dynamic OCT, so external parallel termination resistors are used at the FPGA end of the line.

Stratix III and Stratix IV devices, which support dynamic OCT, do not require FPGA end parallel termination. Hence this discrete parallel termination is optional.

The DDR2 SDRAM DIMM contains a 22- Ω external series termination resistor for each data strobe and data line, so all the measurements and simulations need to account for the effect of these series termination resistors.

Figure 2-3. HyperLynx Setup for Simulating the Stratix II High Speed High Density with Dual-DIMM DDR2 SDRAM Interface



When there is only a single-DIMM on the board, the ODT control is relatively straightforward. During write to the memory, the ODT feature of the memory is turned on; during read from the memory, the ODT feature of the memory is turned off. However, when there are multiple DIMMs on the board, the ODT control becomes more complicated.

With a dual-DIMM interface on the system, the controller has different options for turning the memory ODT on or off during read or write. Table 2-1 shows the DDR2 SDRAM ODT control during write to the memory; Table 2-2 during read from the memory. These DDR2 SDRAM ODT controls are recommended by Samsung Electronics. The JEDEC DDR2 specification was updated to include optional support for $R_{TT}(\text{nominal}) = 50 \Omega$.



For more information about the DDR2 SDRAM ODT controls recommended by Samsung, refer to the *Samsung DDR2 Application Note: ODT (On Die Termination) Control*.

Table 2-1. DDR2 SDRAM ODT Control—Writes (Note 1)

Slot 1 (2)	Slot 2 (2)	Write To	FPGA	Module in Slot 1		Module in Slot 2	
				Rank 1	Rank 2	Rank 3	Rank 4
DR	DR	Slot 1	Series 50 Ω	Infinite	Infinite	75 or 50 Ω	Infinite
		Slot 2	Series 50 Ω	75 or 50 Ω	Infinite	Infinite	infinite
SR	SR	Slot 1	Series 50 Ω	Infinite	Unpopulated	75 or 50 Ω	Unpopulated
		Slot 2	Series 50 Ω	75 or 50 Ω	Unpopulated	Infinite	Unpopulated
DR	Empty	Slot 1	Series 50 Ω	150 Ω	Infinite	Unpopulated	Unpopulated
Empty	DR	Slot 2	Series 50 Ω	Unpopulated	Unpopulated	150 Ω	Infinite
SR	Empty	Slot 1	Series 50 Ω	150 Ω	Unpopulated	Unpopulated	Unpopulated
Empty	SR	Slot 2	Series 50 Ω	Unpopulated	Unpopulated	150 Ω	Unpopulated

Note to Table 2-1:

- (1) For DDR2 at 400 MHz and 533 Mbps = 75 Ω ; for DDR2 at 667 MHz and 800 Mbps = 50 Ω .
 (2) SR = single ranked; DR = dual ranked.

Table 2-2. DDR2 SDRAM ODT Control—Reads (Note 1)

Slot 1 (2)	Slot 2 (2)	Read From	FPGA	Module in Slot 1		Module in Slot 2	
				Rank 1	Rank 2	Rank 3	Rank 4
DR	DR	Slot 1	Parallel 50 Ω	Infinite	Infinite	75 or 50 Ω	Infinite
		Slot 2	Parallel 50 Ω	75 or 50 Ω	Infinite	Infinite	Infinite
SR	SR	Slot 1	Parallel 50 Ω	Infinite	Unpopulated	75 or 50 Ω	Unpopulated
		Slot 2	Parallel 50 Ω	75 or 50 Ω	Unpopulated	Infinite	Unpopulated
DR	Empty	Slot 1	Parallel 50 Ω	Infinite	Infinite	Unpopulated	Unpopulated
Empty	DR	Slot 2	Parallel 50 Ω	Unpopulated	Unpopulated	Infinite	Infinite
SR	Empty	Slot 1	Parallel 50 Ω	Infinite	Unpopulated	Unpopulated	Unpopulated
Empty	SR	Slot 2	Parallel 50 Ω	Unpopulated	Unpopulated	Infinite	Unpopulated

Note to Table 2-1:

- (1) For DDR2 at 400 MHz and 533 Mbps = 75 Ω ; for DDR2 at 667 MHz and 800 Mbps = 50 Ω .
 (2) SR = single ranked; DR = dual ranked.

A 54- Ω external parallel termination resistor is placed on all the data strobes and data lines near the Stratix II device on the Stratix II High Speed High Density Board. Although the characteristic impedance of the transmission is designed for 50 Ω , to account for any process variation, it is advisable to underterminate the termination seen at the receiver. This is why the termination resistors at the FPGA side use 54- Ω resistors.

DIMM Configuration

While populating both memory slots is common in a dual-DIMM memory system, there are some instances when only one slot is populated. For example, some systems are designed to have a certain amount of memory initially and as applications get more complex, the system can be easily upgraded to accommodate more memory by populating the second memory slot without re-designing the system. The following section discusses a dual-DIMM system where the dual-DIMM system only has one slot populated at one time and a dual-DIMM system where both slots are populated. ODT controls recommended by the memory vendors listed in Table 2-1 as well as other possible ODT settings will be evaluated for usefulness in an FPGA system.

Dual-DIMM Memory Interface with Slot 1 Populated

This section focuses on a dual-DIMM memory interface where slot 1 is populated and slot 2 is unpopulated. This section examines the impact on the signal quality due to an unpopulated DIMM slot and compares it to a single-DIMM memory interface.

FPGA Writing to Memory

In the DDR2 SDRAM, the ODT feature has two settings: 150 Ω and 75 Ω . In Table 2-1, the recommended ODT setting for a dual DIMM configuration with one slot occupied is 150 Ω .



On DDR2 SDRAM devices running at 333 MHz/667 Mbps and above, the ODT feature supports an additional setting of 50 Ω .



Refer to the respective memory decathlete for additional information about the ODT settings in DDR2 SDRAM devices.

Write to Memory Using an ODT Setting of 150 Ω

Figure 2-4 shows a double parallel termination scheme (Class II) using ODT on the memory with a memory-side series resistor when the FPGA is writing to the memory using a 25- Ω OCT drive strength setting on the FPGA.

Figure 2-4. Double Parallel Termination Scheme (Class II) Using ODT on DDR2 SDRAM DIMM with Memory-Side Series Resistor




Figure 2-5 shows a HyperLynx simulation and board measurement of a signal at the memory of a double parallel termination using ODT 150 Ω with a memory-side series resistor transmission line when the FPGA is writing to the memory with a 25- Ω OCT drive strength setting.

Figure 2-5. HyperLynx Simulation and Board Measurement of the Signal at the Memory in Slot 1 with Slot 2 Unpopulated

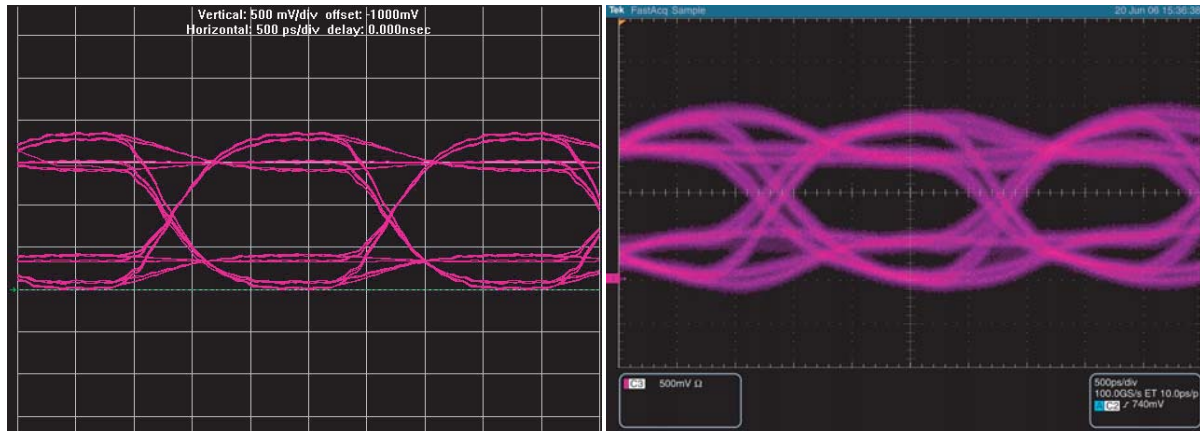


Table 2-3 summarizes the comparison between the simulation and board measurements of the signal at the memory of a single-DIMM and a dual-DIMM memory interface with slot 1 populated using a double parallel termination using an ODT setting of 150 Ω with a memory-side series resistor with a 25- Ω OCT strength setting on the FPGA.

Table 2-3. Comparison of Signal at the Memory of a Single-DIMM and a Dual-DIMM Interface with Slot 1 Populated
(Note 1)

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM memory interface with slot 1 populated						
Simulation	1.68	0.97	0.06	NA	2.08	1.96
Measurements	1.30	0.63	0.22	0.20	1.74	1.82
Single-DIMM						
Simulation	1.62	0.94	0.10	0.05	2.46	2.46
Measurements	1.34	0.77	0.04	0.13	1.56	1.39

Note to Table 2-3:

- (1) The simulation and board measurements of the single-DIMM DDR2 SDRAM interface are based on the Stratix II Memory Board 2. For more information about the single-DIMM DDR2 SDRAM interface, refer to [Chapter 1, DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines](#).

Table 2-3 indicates that there is not much difference between a single-DIMM memory interface or a dual-DIMM memory interface with slot 1 populated. The over and undershooting observed in both the simulations and board measurements can be attributed to the use of the ODT setting of 150 Ω on the memory resulting in over-termination at the receiver. In addition, there is no significant effect of the extra DIMM connector due to the unpopulated slot.

When the ODT setting is set to $75\ \Omega$, there is no difference in the eye width and height compared to the ODT setting of $150\ \Omega$. However, there is no overshoot and undershoot when the ODT setting is set to $75\ \Omega$, which is attributed to proper termination resulting in matched impedance seen by the DDR2 SDRAM devices.




For information about results obtained from using an ODT setting of $75\ \Omega$ refer to [page 2-25](#).

Reading from Memory

During read from the memory, the ODT feature is turned off. Thus, there is no difference between using an ODT setting of $150\ \Omega$ and $75\ \Omega$. As such, the termination scheme becomes a single parallel termination scheme (Class I) where there is an external resistor on the FPGA side and a series resistor on the memory side as shown in [Figure 2-6](#).

Figure 2-6. Single Parallel Termination Scheme (Class I) Using External Resistor and Memory-Side Series Resistor



[Figure 2-7](#) shows the simulation and board measurement of the signal at the FPGA of a single parallel termination using an external parallel resistor on the FPGA side with a memory-side series resistor with full drive strength setting on the memory.

Figure 2-7. HyperLynx Simulation and Board Measurement of the Signal at the FPGA When Reading From Slot 1 With Slot 2 Unpopulated

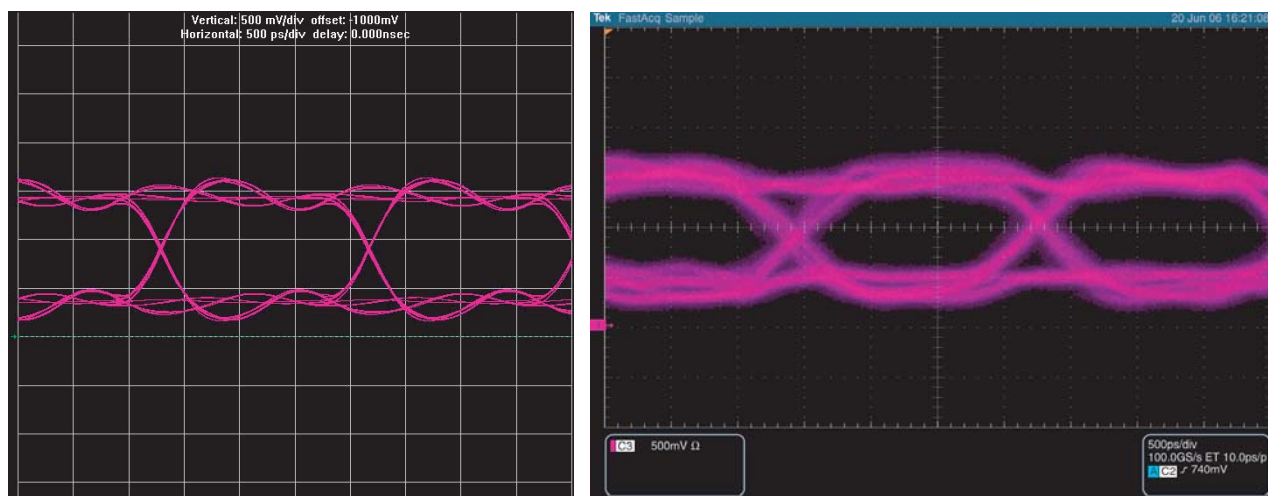


Table 2-4 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a single-DIMM and a dual-DIMM memory interface with a slot 1 populated memory interface using a single parallel termination using an external parallel resistor at the FPGA with a memory-side series resistor with full strength setting on the memory.

Table 2-4. Comparison of Signal at the FPGA of a Dual-DIMM Memory Interface with Slot 1 Populated (Note 1)

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM memory interface with slot 1 populated						
Simulation	1.76	0.80	NA	NA	2.29	2.29
Measurements	1.08	0.59	NA	NA	1.14	1.59
Single-DIMM¹						
Simulation	1.80	0.95	NA	NA	2.67	2.46
Measurements	1.03	0.58	NA	NA	1.10	1.30

Note to Table 2-4:

- (1) The simulation and board measurements of the single-DIMM DDR2 SDRAM interface are based on the Stratix II Memory Board 2. For more information about the single-DIMM DDR2 SDRAM interface, refer to [Chapter 1, DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines](#),

Table 2-4 demonstrates that there is not much difference between a single-DIMM memory interface or a dual-DIMM memory interface with only slot 1 populated. There is no significant effect of the extra DIMM connector due to the unpopulated slot.

Dual-DIMM with Slot 2 Populated

This section focuses on a dual-DIMM memory interface where slot 2 is populated and slot 1 is unpopulated. Specifically, this section discusses the impact of location of the DIMM on the signal quality.

FPGA Writing to Memory

The previous section focused on the dual-DIMM memory interface where slot 1 is populated resulting in the memory being located closer to the FPGA. When slot 2 is populated, the memory is located further away from the FPGA, resulting in additional trace length that potentially affects the signal quality seen by the memory. The next section explores if there are any differences between populating slot 1 and slot 2 of the dual-DIMM memory interface.

Write to Memory Using an ODT Setting of 150Ω

Figure 2-8 shows the double parallel termination scheme (Class II) using ODT on the memory with the memory-side series resistor when the FPGA is writing to the memory using a $25\text{-}\Omega$ OCT drive strength setting on the FPGA.

Figure 2-8. Double Parallel Termination Scheme (Class II) Using ODT on DDR2 SDRAM DIMM with Memory-side Series Resistor




Figure 2-9 shows the simulation and board measurement of the signal at the memory of a double parallel termination using an ODT setting of 150Ω with a memory-side series resistor transmission line when the FPGA is writing to the memory with a $25\text{-}\Omega$ OCT drive strength setting.

Figure 2-9. HyperLynx Simulation and Board Measurement of the Signal at the Memory in Slot 2 With Slot 1 Unpopulated

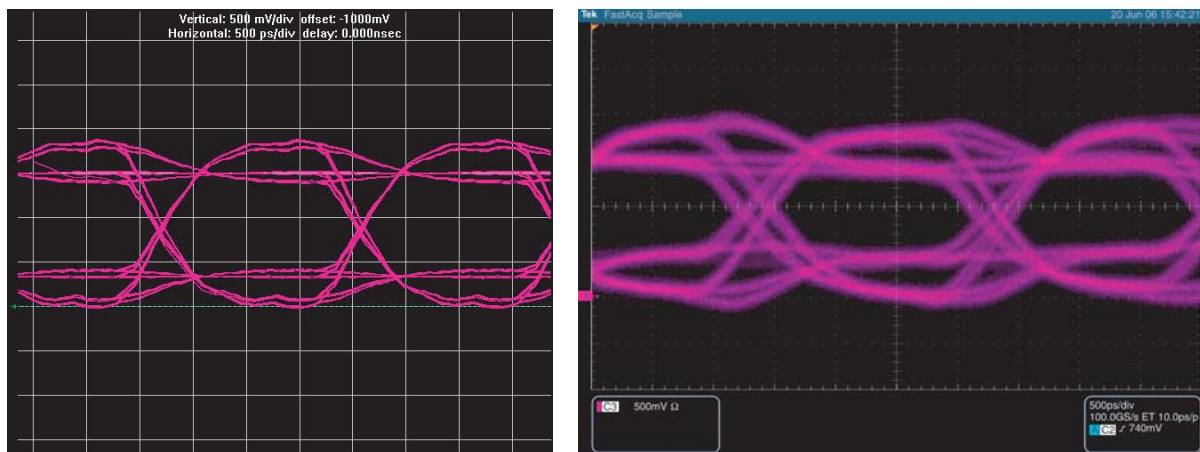


Table 2-5 summarizes the comparison between the simulation and board measurements of the signal seen at the DDR2 SDRAM DIMM of a dual-DIMM memory interface with either only slot 1 populated or only slot 2 populated using a double parallel termination using an ODT setting of $150\ \Omega$ with a memory-side series resistor with a $25\text{-}\Omega$ OCT strength setting on the FPGA.

Table 2-5. Comparison of Signal at the Memory of a Dual-DIMM Interface with Either Only Slot 1 Populated or Only Slot 2 Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM Memory Interface with Slot 2 Populated						
Simulation	1.69	0.94	0.07	0.02	1.96	2.08
Measurements	1.28	0.68	0.24	0.20	1.60	1.60
Dual-DIMM Memory Interface with Slot 1 Populated						
Simulation	1.68	0.97	0.06	NA	2.08	2.08
Measurements	1.30	0.63	0.22	0.20	1.74	1.82

Table 2-5 shows that there is not much difference between populating slot 1 or slot 2 in a dual-DIMM memory interface. The over and undershooting observed in both the simulations and board measurements can be attributed to the use of the ODT setting of $150\ \Omega$ on the memory, resulting in under-termination at the receiver.

When the ODT setting is set to $75\ \Omega$, there is no difference in the eye width and height compared to the ODT setting of $150\ \Omega$. However, there is no overshoot and undershoot when the ODT setting is set to $75\ \Omega$, which is attributed to proper termination resulting in matched impedance seen by the DDR2 SDRAM devices.



For detailed results for the ODT setting of $75\ \Omega$, refer to [page 2-26](#).

Reading from Memory

During read from memory, the ODT feature is turned off, thus there is no difference between using an ODT setting of $150\ \Omega$ and $75\ \Omega$. As such, the termination scheme becomes a single parallel termination scheme (Class I) where there is an external resistor on the FPGA side and a series resistor on the memory side, as shown in [Figure 2-10](#).

Figure 2-10. Single Parallel Termination Scheme (Class I) Using External Resistor and Memory-Side Series Resistor




Figure 2-11 shows the simulation and board measurement of the signal at the FPGA of a single parallel termination using an external parallel resistor on the FPGA side with a memory-side series resistor with full drive strength setting on the memory.

Figure 2-11. HyperLynx Simulation and Board Measurement of the Signal at the FPGA When Reading From Slot 2 With Slot 1 Unpopulated

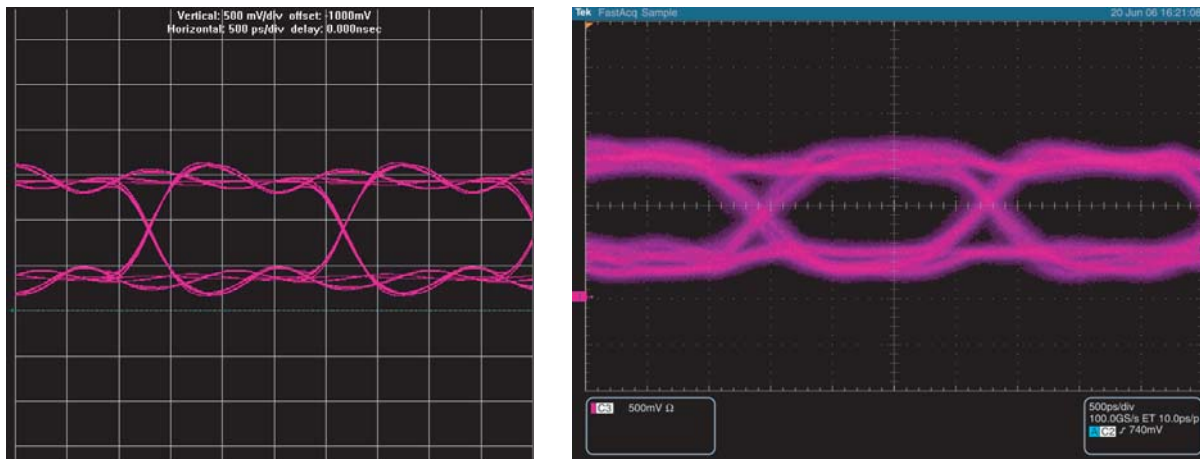


Table 2-6 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a dual-DIMM memory interface with either slot 1 or slot 2 populated using a single parallel termination using an external parallel resistor at the FPGA with a memory-side series resistor with full strength setting on the memory.

Table 2-6. Comparison of the Signal at the FPGA of a Dual-DIMM Memory Interface with Either Slot 1 or Slot 2 Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Slot 2 Populated						
Simulation	1.80	0.80	NA	NA	3.09	2.57
Measurements	1.17	0.66	NA	NA	1.25	1.54
Slot 1 Populated						
Simulation	1.80	0.95	NA	NA	2.67	2.46
Measurements	1.08	0.59	NA	NA	1.14	1.59

From Table 2-6, you can see the signal seen at the FPGA is similar whether the memory DIMM is located at either slot 1 or slot 2.

Dual-DIMM Memory Interface with Both Slot 1 and Slot 2 Populated

This section focuses on a dual-DIMM memory interface where both slot 1 and slot 2 are populated. As such, you can write to either the memory in slot 1 or the memory in slot 2.

FPGA Writing to Memory

In Table 2-1, the recommended ODT setting for a dual DIMM configuration with both slots occupied is $75\ \Omega$. Since there is an option for an ODT setting of $150\ \Omega$, this section explores the usage of the $150\ \Omega$ setting and compares the results to that of the recommended $75\ \Omega$.

Write to Memory in Slot 1 Using an ODT Setting of $75\ \Omega$

Figure 2-12 shows the double parallel termination scheme (Class II) using ODT on the memory with the memory-side series resistor when the FPGA is writing to the memory using a $25\text{-}\Omega$ OCT drive strength setting on the FPGA. In this scenario, the FPGA is writing to the memory in slot 1 and the ODT feature of the memory at slot 2 is turned on.

Figure 2-12. Double Parallel Termination Scheme (Class II) Using ODT on DDR2 SDRAM DIMM with a Memory-Side Series Resistor

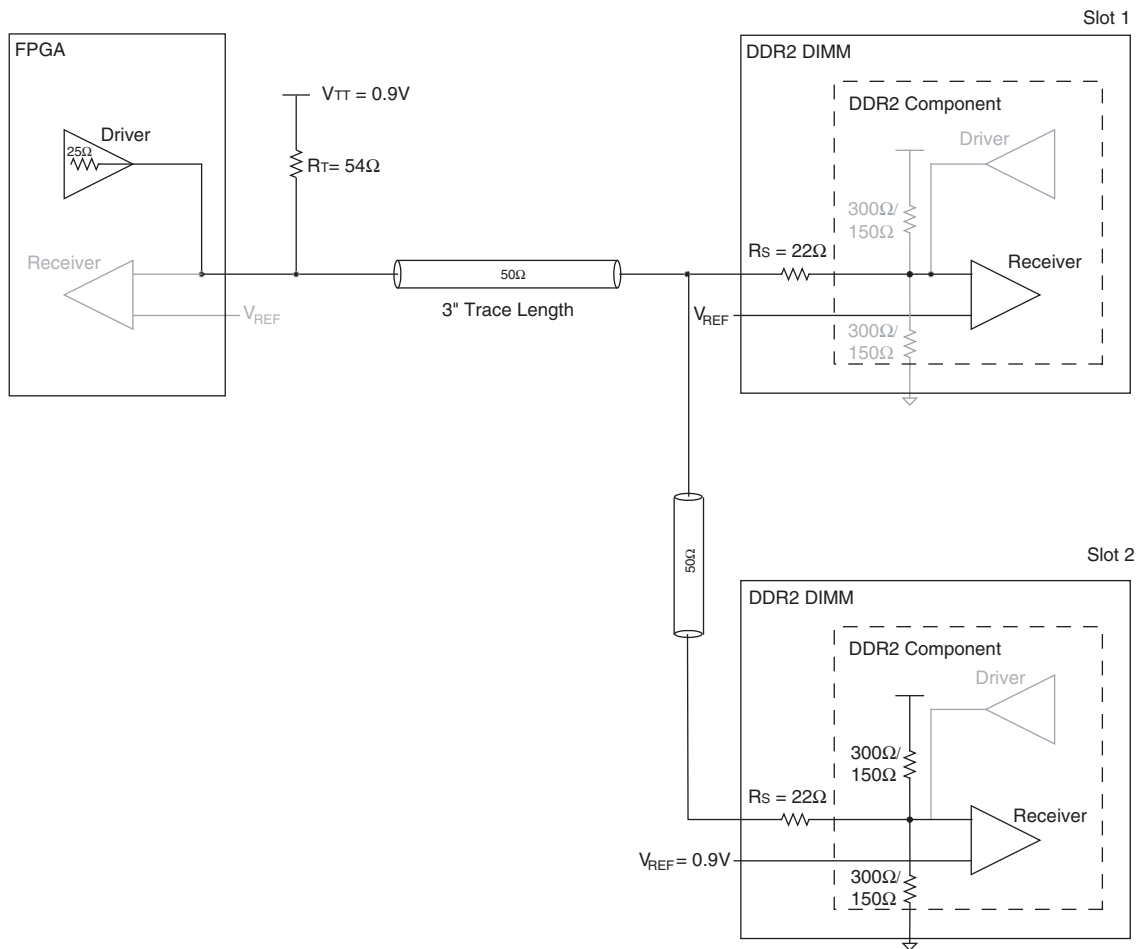


Figure 2-13 shows a HyperLynx simulation and board measurement of the signal at the memory in slot 1 of a double parallel termination using an ODT setting of $75\ \Omega$ with a memory-side series resistor transmission line when the FPGA is writing to the memory with a $25\text{-}\Omega$ OCT drive strength setting.

Figure 2-13. HyperLynx Simulation and Board Measurements of the Signal at the Memory in Slot 1 with Both Slots Populated

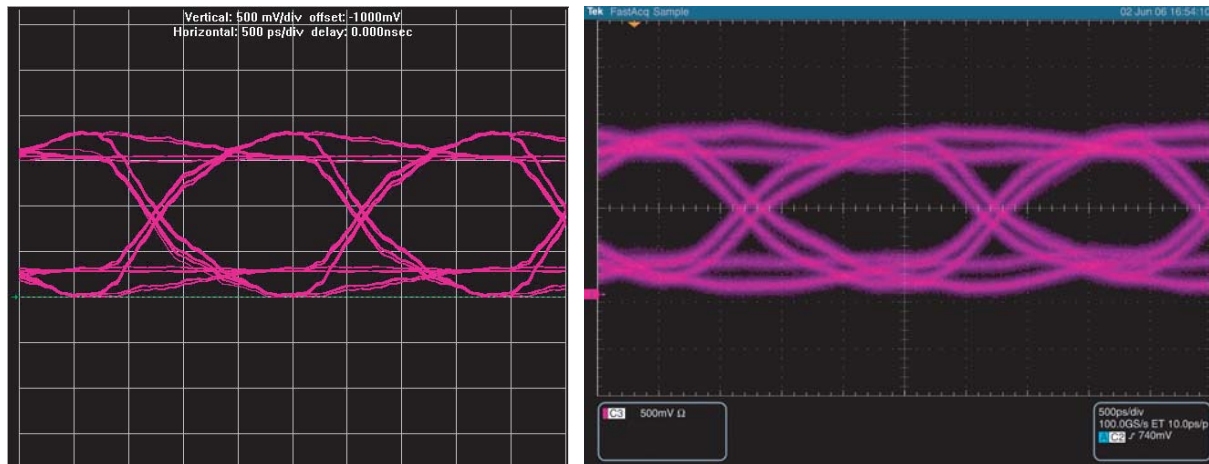


Table 2-7 summarizes the comparison of the signal at the memory of a dual-DIMM memory interface with one slot and with both slots populated using a double parallel termination using an ODT setting of $75\ \Omega$ with a memory-side series resistor with a $25\text{-}\Omega$ OCT strength setting on the FPGA.

Table 2-7. Comparison of the Signal at the Memory of a Dual-DIMM Interface With One Slot and With Both Slots Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM Interface with Both Slots Populated Writing to Slot 1						
Simulation	1.60	1.18	0.02	NA	1.71	1.71
Measurements	0.97	0.77	0.05	0.04	1.25	1.25
Dual-DIMM Interface with Slot 1 Populated						
Simulation	1.68	0.97	0.06	NA	2.08	2.08
Measurements	1.30	0.63	0.22	0.20	1.74	1.82

Table 2-7 shows that there is not much difference in the eye height between populating one slot or both slots. However, the additional loading due to the additional memory DIMM results in a slower edge rate, which results in smaller eye width and degrades the setup and hold time of the memory. This reduces the available data valid window.

When the ODT setting is set to $150\ \Omega$, there is no difference in the eye width and height compared to the ODT setting of $75\ \Omega$. However, there is some overshoot and undershoot when the ODT setting is set to $150\ \Omega$, which is attributed to under termination resulting in mismatched impedance seen by the DDR2 SDRAM devices.




For more information about the results obtained from using an ODT setting of $150\ \Omega$, refer to [page 2-27](#).

Write to Memory in Slot 2 Using an ODT Setting of $75\ \Omega$

In this scenario, the FPGA is writing to the memory in slot 2 and the ODT feature of the memory at slot 1 is turned on. [Figure 2-14](#) shows the HyperLynx simulation and board measurement of the signal at the memory in slot 1 of a double parallel termination using an ODT setting of $75\ \Omega$ with a memory-side series resistor transmission line when the FPGA is writing to the memory with a $25\text{-}\Omega$ OCT drive strength setting.

Figure 2-14. HyperLynx Simulation and Board Measurements of the Signal at the Memory in Slot 2 With Both Slots Populated



[Table 2-8](#) summarizes the comparison of the signal at the memory of a dual-DIMM memory interface with slot 1 populated using a double parallel termination using an ODT setting of $75\ \Omega$ with a memory-side series resistor with a $25\text{-}\Omega$ OCT strength setting on the FPGA.

Table 2-8. Comparison of the Signal at the Memory of a Dual-DIMM Interface With Both Slots Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rise Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM Interface with Both Slots Populated Writing to Slot 2						
Simulation	1.60	1.16	0.10	0.08	1.68	1.60
Measurements	1.10	0.85	0.16	0.19	1.11	1.25
Dual-DIMM Interface with Both Slots Populated Writing to Slot 1						
Simulation	1.60	1.18	0.02	NA	1.71	1.71
Measurements	1.30	0.77	0.05	0.04	1.25	1.25

From [Table 2-8](#), you can see that both simulations and board measurements demonstrate that the eye width is larger when writing to slot 1, which is due to better edge rate seen when writing to slot 1. The improvement on the eye when writing to slot 1 can be attributed to the location of the termination. When you are writing to slot 1, the ODT feature of slot 2 is turned on, resulting in a fly-by topology. When you are writing to slot 2, the ODT feature of slot 1 is turned on resulting in a non fly-by topology.

When the ODT setting is set to 150 Ω , there is no difference in the eye width and height compared to the ODT setting of 75 Ω . However, there is some overshoot and undershoot when the ODT setting is set to 150 Ω , which is attributed to under termination resulting in mismatched impedance seen by the DDR2 SDRAM devices.



For more information about the results obtained from using an ODT setting of 150 Ω , refer to [“Write to Memory in Slot 2 Using an ODT Setting of 150 \$\Omega\$ With Both Slots Populated”](#) on page 2-28.

Reading From Memory

In [Table 2-2](#), the recommended ODT setting for a dual-DIMM configuration with both slots occupied is to turn on the ODT feature using a setting of 75 Ω on the slot that is not read from. As there is an option for an ODT setting of 150 Ω , this section explores the usage of the 150 Ω setting and compares the results to that of the recommended 75 Ω .

Read From Memory in Slot 1 Using an ODT Setting of 75- Ω on Slot 2

Figure 2-15 shows the double parallel termination scheme (Class II) using ODT on the memory with the memory-side series resistor when the FPGA is reading from the memory using a full drive strength setting on the memory. In this scenario, the FPGA is reading from the memory in slot 1 and the ODT feature of the memory at slot 2 is turned on.

Figure 2-15. Double Parallel Termination Scheme (Class II) Using External Resistor and Memory-Side Series Resistor and ODT Feature Turned On

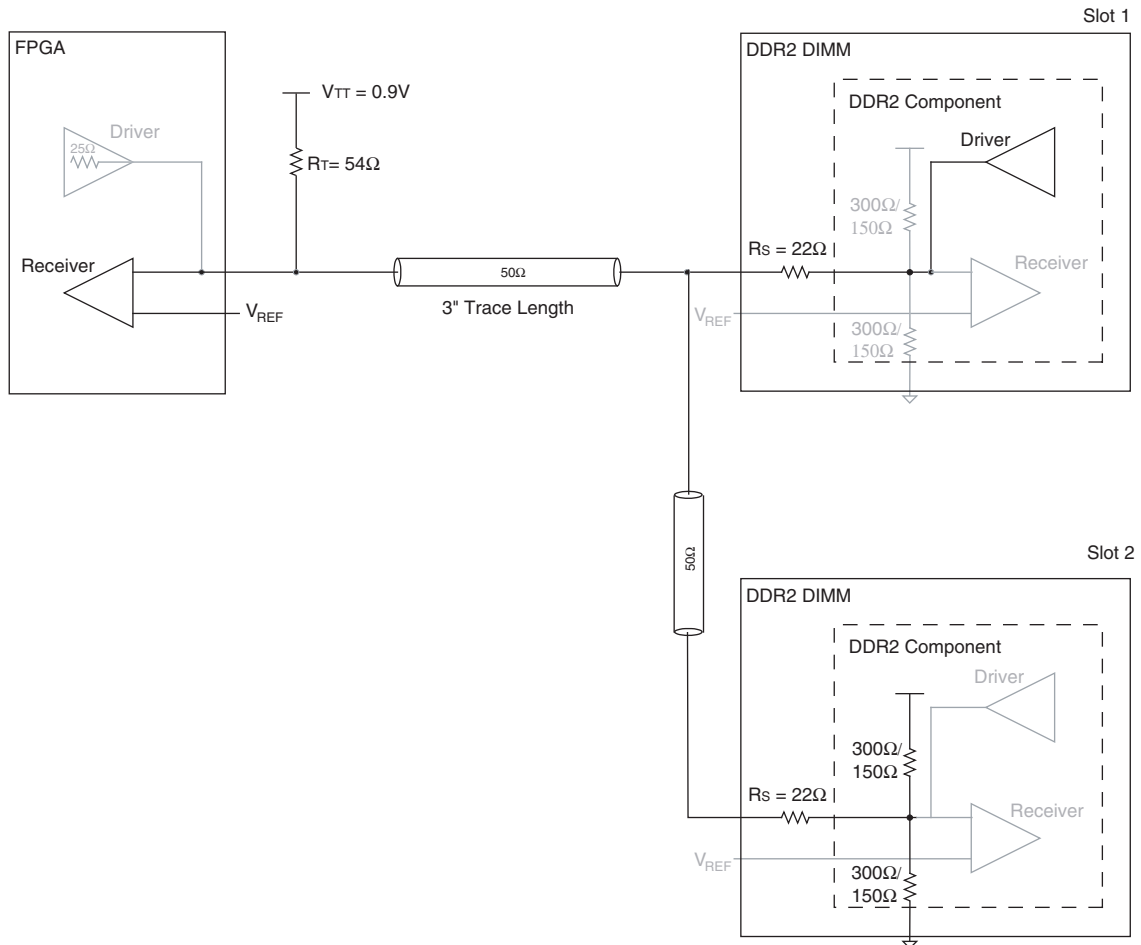
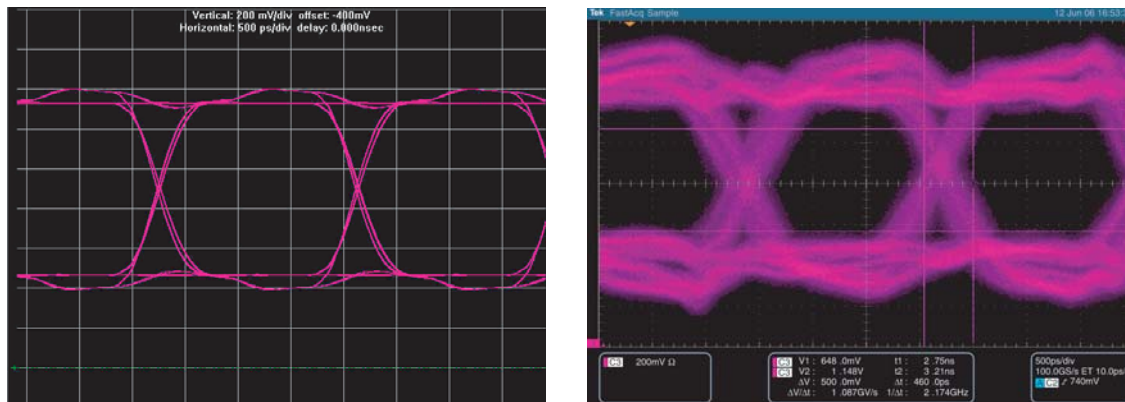


Figure 2-16 shows the simulation and board measurement of the signal at the FPGA when the FPGA is reading from the memory in slot 1 using a full drive strength setting on the memory.

Figure 2-16. HyperLynx Simulation and Board Measurement of the Signal at the FPGA When Reading From Slot 1 With Both Slots Populated (Note 1)



Note to Figure 2-16:

- (1) The vertical scale used for the simulation and measurement is set to 200 mV per division.

Table 2-9 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a dual-DIMM memory interface with both slots populated and a dual-DIMM memory interface with a slot 1 populated memory interface.

Table 2-9. Comparison of the Signal at the FPGA of a Dual-DIMM Interface Reading From Slot 1 With One Slot and With Both Slots Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM with One Slot Populated with an ODT Setting of 75-Ω on Slot 2						
Simulation	1.74	0.87	NA	NA	1.91	1.88
Measurements	0.86	0.58	NA	NA	1.11	1.09
Dual-DIMM with One Slot Populated in Slot 1 without ODT Setting						
Simulation	1.76	0.80	NA	NA	2.29	2.29
Measurements	1.08	0.59	NA	NA	1.14	1.59

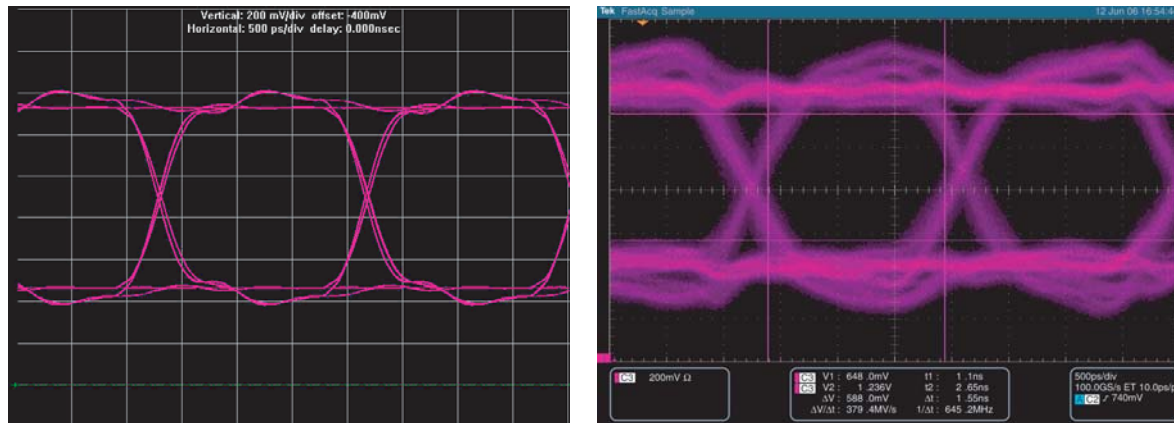
Table 2-9 shows that when both slots are populated, the additional loading due to the additional memory DIMM results in a slower edge rate, which results in a degradation in the eye width.



For more information about the results obtained from using an ODT setting of 150 Ω, refer to “Read from Memory in Slot 1 Using an ODT Setting of 150 Ω on Slot 2 with Both Slots Populated” on page 2-29.

Figure 2-18 shows the HyperLynx simulation and board measurement of the signal at the FPGA of a double parallel termination using an external parallel resistor on the FPGA side with a memory-side series resistor and an ODT setting of $75\ \Omega$ with a full drive strength setting on the memory.

Figure 2-18. HyperLynx Simulation and Board Measurements of the Signal at the FPGA When Reading From Slot 2 With Both Slots Populated (Note 1)



Notes to Figure 2-18:

- (1) The vertical scale used for the simulation and measurement is set to 200 mV per division.

Table 2-10 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a dual-DIMM memory interface with both slots populated and a dual-DIMM memory interface with a slot 1 populated memory interface.

Table 2-10. Comparison of the Signal at the FPGA of a Dual-DIMM Interface Reading From Slot 2 With One Slot and With Both Slots Populated

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
Dual-DIMM with Both Slots Populated with an ODT Setting of $75\text{-}\Omega$ Setting on Slot 1						
Simulation	1.70	0.81	NA	NA	1.72	1.99
Measurements	0.87	0.59	NA	NA	1.09	1.14
Dual-DIMM with One Slot Populated in Slot 2 without an ODT Setting						
Simulation	1.80	0.80	NA	NA	3.09	2.57
Measurements	1.17	0.66	NA	NA	1.25	1.54

Table 2-10 shows that when only one slot is populated in a dual-DIMM memory interface, the eye width is larger as compared to a dual-DIMM memory interface with both slots populated. This can be attributed to the loading from the DIMM located in slot 1.

When the ODT setting is set to $150\ \Omega$, there is no difference in the signal quality compared to the ODT setting of $75\ \Omega$.



For more information about the results obtained from using an ODT setting of $150\ \Omega$, refer to [“Read From Memory in Slot 2 Using an ODT Setting of \$150\ \Omega\$ on Slot 1 With Both Slots Populated”](#) on page 2-30.

FPGA OCT Features

Many FPGA devices offer OCT. Depending on the chosen device family, series (output), parallel (input) or dynamic (bidirectional) OCT may be supported.



For more information specific to your device family, refer to the respective I/O features chapter in the relevant device handbook.

Use series OCT in place of the near-end series terminator typically used in both Class I or Class II termination schemes that both DDR2 and DDR3 type interfaces use.

Use parallel OCT in place of the far-end parallel termination typically used in Class I termination schemes on unidirectional input only interfaces. For example, QDR-II type interfaces, when the FPGA is at the far end.

Use dynamic OCT in place of both the series and parallel termination at the FPGA end of the line. Typically use dynamic OCT for DQ and DQS signals in both DDR2 and DDR3 type interfaces. As the parallel termination is dynamically disabled during writes, the FPGA driver only ever drives into a Class I transmission line. When combined with dynamic ODT at the memory, a truly dynamic Class I termination scheme exists where both reads and writes are always fully Class I terminated in each direction. Hence, you can use a fully dynamic bidirectional Class I termination scheme instead of a static discretely terminated Class II topology, which saves power, printed circuit board (PCB) real estate, and component cost.

Stratix III and Stratix IV Devices

Stratix III and Stratix IV devices feature full dynamic OCT termination capability, Altera advise that you use this feature combined with the SDRAM ODT to simplify PCB layout and save power.

Arria II GX Devices

Arria II GX devices do not support dynamic OCT. Altera recommends that you use series OCT with SDRAM ODT. Use parallel discrete termination at the FPGA end of the line when necessary,




For more information, refer to [Chapter 1, DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines](#).

Dual-DIMM DDR2 Clock, Address, and Command Termination and Topology

The address and command signals on a DDR2 SDRAM interface are unidirectional signals that the FPGA memory controller drives to the DIMM slots. These signals are always Class-I terminated at the memory end of the line (Figure 2-19). Always place DDR2 SDRAM address and command Class-I termination after the last DIMM. The interface can have one or two DIMMs, but never more than two DIMMs total.

Figure 2-19. Multi DIMM DDR2 Address and Command Termination Topology



In Figure 2-19, observe the following points:

- Board trace A = 1.9 to 4.5 inches (48 to 115 mm)
- Board trace B = 0.425 inches (10.795 mm)
- Board trace C = 0.2 to 0.55 inches (5 to 13 mm)
- Total of board trace A + B + C = 2.5 to 5 inches (63 to 127 mm)
- $R_P = 36$ to 56Ω
- Length match all address and command signals to +250 mils (+5 mm) or ± 50 ps of memory clock length at the DIMM.

You may place a compensation capacitor directly before the first DIMM slot 1 to improve signal quality on the address and command signal group. If you fit a capacitor, Altera recommends a value of 24 pF.



For more information, refer to *Micron TN47-01*.

Address and Command Signals

The address and command group of signals: bank address, address, RAS#, CAS#, and WE#, operate a different toggle rate depending on whether you implement a full-rate or half-rate memory controller.

In full-rate designs, the address and command group of signals are 1T signals, which means that the signals can change every memory clock cycle. Address and command signals are also single data rate (SDR). Hence in a full-rate PHY design, the address and command signals operate at a maximum frequency of $0.5 \times$ the data rate. For example in a 266-MHz full rate design, the maximum address and command frequency is 133 MHz.

In half-rate designs the address and command group of signals are 2T signals, which means that the signals change only every two memory clock cycles. As the signals are also SDR, in a half-rate PHY design, the address and command signals operate at a maximum frequency of $0.25 \times$ the data rate. For example, in a 400-MHz half-rate design, the maximum address and command frequency is 100 MHz.

Control Group Signals

The control group of signals: chip select CS#, clock enable CKE, and ODT are always 1T regardless of whether you implement a full-rate or half-rate design. As the signals are also SDR, the control group signals operate at a maximum frequency of $0.5 \times$ the data rate. For example, in a 400-MHz design, the maximum control group frequency is 200 MHz.

Clock Group Signals

Depending on the specific form factor, DDR2 SDRAM DIMMs have two or three differential clock pairs, to ensure that the loading on the clock signals is not excessive. The clock signals are always terminated on the DIMMs and hence no termination is required on your PCB. Additionally, each DIMM slot is required to have its own dedicated set of clock signals. Hence clock signals are always point-to-point from the FPGA PHY to each individual DIMM slot. Individual memory clock signals should never be shared between two DIMM slots.


A typical two slot DDR2 DIMM design therefore has six differential memory clock pairs—three to the first DIMM and three to the second DIMM. All six memory clock pairs must be delay matched to each other to ± 25 mils (± 0.635 mm) and ± 10 mils (± 0.254 mm) for each CLK to CLK# signal.

You may place a compensation capacitor between each clock pair directly before the DIMM connector, to improve the clock slew rates. As FPGA devices have fully programmable drive strength and slew rate options, this capacitor is usually not required for FPGA design. However, Altera advise that you simulate your specific implementation to ascertain if this capacitor is required or not. If fitted the best value is typically 5 pF.

DDR3 SDRAM

This section details the system implementation of a dual slot unbuffered DDR3 SDRAM interface, operating at up to 400 MHz and 800 Mbps data rates. [Figure 2–20](#) shows a typical DQS, DQ, and DM, and address and command signal topology for a dual-DIMM interface configuration, using the ODT feature of the DDR3 SDRAM components combined with the dynamic OCT features available in Stratix III and Stratix IV devices.

Figure 2–20. Multi DIMM DDR3 DQS, DQ, and DM, and Address and Command Termination Topology



In [Figure 2–20](#), observe the following points:

- Board trace A = 1.9 to 4.5 inches (48 to 115 mm)
- Board trace B = 0.425 inches (10.795 mm)
- This topology to both DIMMs is accurate for DQS, DQ, and DM, and address and command signals
- This topology is not correct for CLK and CLK# and control group signals (CS#, CKE, and ODT), which are always point-to-point single rank only.

Comparison of DDR3 and DDR2 DQ and DQS ODT Features and Topology

DDR3 and DDR2 SDRAM systems are quite similar. The physical topology of the data group of signals may be considered nearly identical. The FPGA end (driver) I/O standard changes from SSTL18 for DDR2 to SSTL15 for DDR3, but all other OCT settings are identical. DDR3 offers enhanced ODT options for termination and drive-strength settings at the memory end of the line.



For more information, refer to the DDR3 SDRAM ODT matrix for writes and the DDR3 SDRAM ODT matrix for reads tables in [Chapter 1, DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines](#).

Dual-DIMM DDR3 Clock, Address, and Command Termination and Topology

One significant difference between DDR3 and DDR2 DIMM based interfaces is the address, command and clock signals. DDR3 uses a daisy chained based architecture when using JEDEC standard modules. The address, command, and clock signals are routed on each module in a daisy chain and feature a fly-by termination on the module. Impedance matching is required to make the dual-DIMM topology work effectively—40 to 50 Ω traces should be targeted on the main board.

Address and Command Signals

Two UDIMMs result in twice the effective load on the address and command signals, which reduces the slew rate and makes it more difficult to meet setup and hold timing (t_{IS} and t_{IH}). However, address and command signals operate at half the interface rate and are SDR. Hence a 400-Mbps data rate equates to an address and command fundamental frequency of 100 MHz.

Control Group Signals

The control group signals (chip Select CS#, clock enable CKE, and ODT) are only ever single rank. A dual-rank capable DDR3 DIMM slot has two copies of each signal, and a dual-DIMM slot interface has four copies of each signal. Hence the signal quality of these signals is identical to a single rank case. The control group of signals, are always 1T regardless of whether you implement a full-rate or half-rate design. As the signals are also SDR, the control group signals operate at a maximum frequency of $0.5 \times$ the data rate. For example, in a 400 MHz design, the maximum control group frequency is 200 MHz.

Clock Group Signals

Like the control group signals, the clock signals in DDR3 SDRAM are only ever single rank loaded. A dual-rank capable DDR3 DIMM slot has two copies of the signal, and a dual-slot interface has four copies of the `mem_clk` and `mem_clk_n` signals.



For more information about a DDR3 two-DIMM system design, refer to Micron *TN-41-08: DDR3 Design Guide for Two-DIMM Systems*.



The Altera DDR3 ALTMEMPHY megafunction does not support the 1T address and command topology referred to in this Micron Technical Note—only 2T implementations are supported.

Write to Memory in Slot 1 Using an ODT Setting of 75 Ω With One Slot Populated

Figure 2-21 shows the simulation and board measurement of the signal at the memory when the FPGA is writing to the memory with an ODT setting of 75 Ω and using a 25- Ω OCT drive strength setting on the FPGA.

Figure 2-21. HyperLynx Simulation and Board Measurement of the Signal at the Memory in Slot 1 With Slot 2 Unpopulated

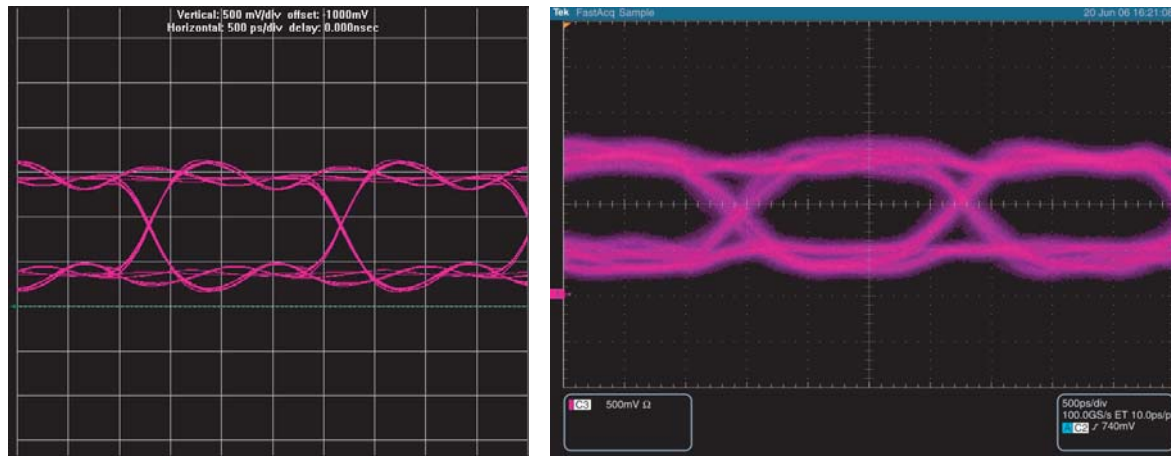


Table 2-11 summarizes the comparison between the simulation and board measurements of the signal seen at the DDR2 SDRAM of a dual-DIMM with slot 1 populated by a memory interface using a different ODT setting.

Table 2-11. Comparison of the Signal at the Memory of a Dual-DIMM Interface With Only Slot 1 Populated and a Different ODT Setting

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 75 Ω						
Simulation	1.68	0.91	NA	NA	1.88	1.88
Measurements	1.28	0.57	NA	NA	1.54	1.38
ODT Setting of 150 Ω						
Simulation	1.68	0.97	0.06	NA	2.67	2.13
Measurements	1.30	0.63	0.22	0.20	1.74	1.82

Write to Memory in Slot 2 Using an ODT Setting of 75 Ω With One Slot Populated

Figure 2-22 shows the simulation and measurements result of the signal seen at the memory when the FPGA is writing to the memory with an ODT setting of 75 Ω and using a 25- Ω OCT drive strength setting on the FPGA.

Figure 2-22. HyperLynx Simulation and Board Measurement of the Signal at the Memory in Slot 2 with Slot 1 Unpopulated

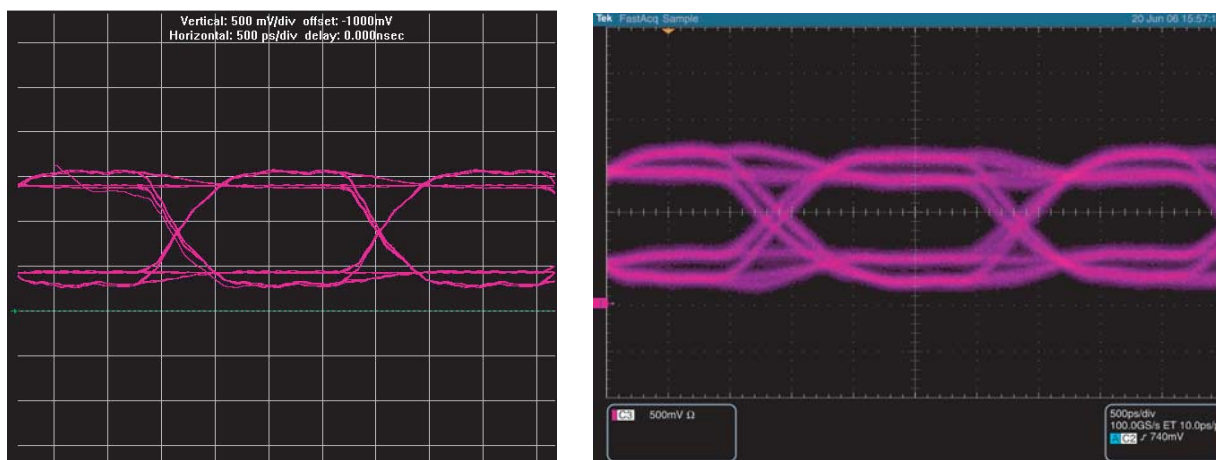


Table 2-12 summarizes the comparison of the signal at the memory of a dual-DIMM memory interface with either slot 1 or slot 2 populated using a double parallel termination using an ODT setting of 75 Ω with a memory-side series resistor with a 25- Ω OCT strength setting on the FPGA.

Table 2-12. Comparison of Signal at the Memory of a Dual-DIMM Interface With Only Slot 2 Populated and a Different ODT Setting

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 75 Ω						
Simulation	1.68	0.89	NA	NA	1.82	1.93
Measurements	1.29	0.59	NA	NA	1.60	1.29
ODT Setting of 150 Ω						
Simulation	1.69	0.94	0.07	0.02	1.88	2.29
Measurements	1.28	0.68	0.24	0.20	1.60	1.60

Write to Memory in Slot 1 Using an ODT Setting of 150 Ω With Both Slots Populated

Figure 2-23 shows the HyperLynx simulation and board measurement of the signal at the memory in slot 1 of a double parallel termination using an ODT setting of 150 Ω on Slot 2 with a memory-side series resistor transmission line when the FPGA is writing to the memory with a 25- Ω OCT drive strength setting.

Figure 2-23. HyperLynx Simulation and Board Measurement of the Signal at the Memory in Slot 1 With Both Slots Populated

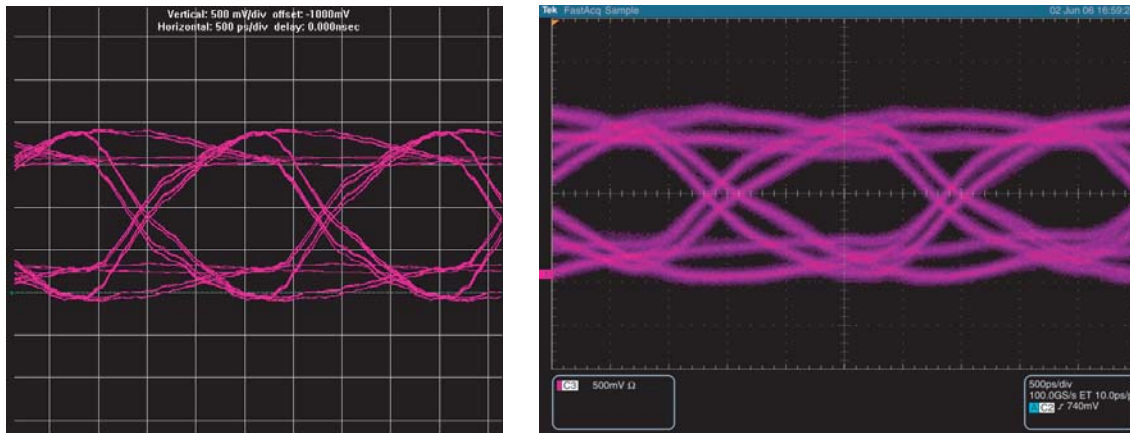


Table 2-13 summarizes the comparison between the simulation and board measurements of the signal seen at the memory in slot 1 of a dual-DIMM memory interface with both slots populated using a double parallel termination using a different ODT setting on Slot 2 with a memory-side series resistor with a 25- Ω OCT strength setting on the FPGA.

Table 2-13. Comparison of Signal at the Memory of a Dual-DIMM Interface with Both Slots Populated and a Different ODT Setting on Slot 2

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 150 Ω						
Simulation	1.60	1.18	0.02	NA	1.71	1.71
Measurements	0.89	0.78	0.13	0.17	1.19	1.32
ODT Setting of 75 Ω						
Simulation	1.60	1.18	0.02	NA	1.71	1.71
Measurements	0.97	0.77	0.05	0.04	1.25	1.25

Write to Memory in Slot 2 Using an ODT Setting of 150 Ω With Both Slots Populated

Figure 2-24 shows the HyperLynx simulation and board measurement of the signal at the memory in slot 2 of a double parallel termination using an ODT setting of 150 Ω on slot 1 with a memory-side series resistor transmission line when the FPGA is writing to the memory with a 25- Ω OCT drive strength setting.

Figure 2-24. HyperLynx Simulation and Board Measurements of the Signal at the Memory in Slot 2 with Both Slots Populated

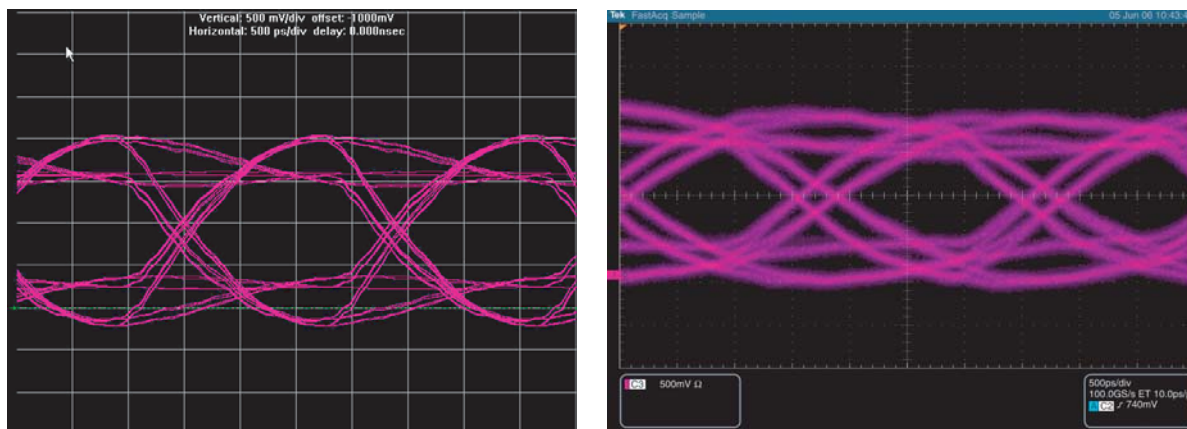


Table 2-14 summarizes the comparison between the simulation and board measurements of the signal seen at the memory of a dual-DIMM memory interface with both slots populated using a double parallel termination using a different ODT setting on Slot 1 with a memory-side series resistor with a 25- Ω OCT strength setting on the FPGA.


Table 2-14. Comparison of the Signal at the Memory of a Dual-DIMM Interface With Both Slots Populated and a Different ODT Setting on Slot 1

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 150 Ω						
Simulation	1.45	1.11	0.19	0.17	1.43	2.21
Measurements	0.71	0.81	0.12	0.20	0.93	1.00
ODT Setting of 75 Ω						
Simulation	1.60	1.16	0.10	0.08	1.68	1.60
Measurements	1.10	0.85	0.16	0.19	1.11	1.25

Read from Memory in Slot 1 Using an ODT Setting of 150 Ω on Slot 2 with Both Slots Populated

Figure 2-25 shows the HyperLynx simulation and board measurement of the signal at the FPGA of a double parallel termination using an external parallel resistor on the FPGA side with a memory-side series resistor and an ODT setting of 150 Ω with a full drive strength setting on the memory.

Figure 2-25. HyperLynx Simulation and Board Measurement of the Signal at the FPGA When Reading From Slot 1 With Both Slots Populated (Note 1)



Note to Figure 2-25:

- (1) The vertical scale used for the simulation and measurement is set to 200 mV per division.

Table 2-15 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a dual-DIMM memory interface with both slots populated using a different ODT setting on Slot 2.


Table 2-15. Comparison of Signal at the FPGA of a Dual-DIMM Interface With Both Slots Populated and a Different ODT Setting on Slot 2

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rise Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 150 Ω						
Simulation	1.68	0.77	NA	NA	1.88	1.88
Measurements	0.76	0.55	NA	NA	1.11	1.14
ODT Setting of 75 Ω						
Simulation	1.74	0.87	NA	NA	1.91	1.88
Measurements	0.86	0.59	NA	NA	1.11	1.09

Read From Memory in Slot 2 Using an ODT Setting of 150 Ω on Slot 1 With Both Slots Populated

Figure 2-26 shows the HyperLynx simulation board measurement of the signal seen at the FPGA of a double parallel termination using an external parallel resistor on the FPGA side with memory-side series resistor and an ODT setting of 150 Ω with a full drive strength setting on the memory.

Figure 2-26. HyperLynx Simulation Board Measurement of the Signal at the FPGA When Reading From Slot 2 With Both Slots Populated (Note 1)



Note to Figure 2-26:

- (1) The vertical scale used for the simulation and measurement is set to 200 mV per division.

Table 2-16 summarizes the comparison between the simulation and board measurements of the signal seen at the FPGA of a dual-DIMM memory interface with both slots populated using a different ODT setting on Slot 1.

Table 2-16. Comparison of Signal at the FPGA of a Dual-DIMM Interface With Both Slots Populated and a Different ODT Setting on Slot 1

Type	Eye Width (ns)	Eye Height (V)	Overshoot (V)	Undershoot (V)	Rising Edge Rate (V/ns)	Falling Edge Rate (V/ns)
ODT Setting of 150 Ω						
Simulation	1.70	0.74	NA	NA	1.91	1.64
Measurements	0.74	0.64	NA	NA	1.14	1.14
ODT Setting of 75 Ω						
Simulation	1.70	0.81	NA	NA	1.72	1.99
Measurements	0.87	0.59	NA	NA	1.09	1.14

Conclusion

This chapter looks at single- and dual-DIMM DDR2 and DDR3 SDRAM interfaces and makes recommendations on topology and termination, to ensure optimum design guidelines and best signal quality.

In the design of any dual-DIMM interface, you should follow the memory vendor recommendations on the optimum ODT setting, slot population, and operations to the DIMM locations. In addition, this chapter recommends the OCT settings to use at the FPGA, to ensure optimum configuration.

The simulations and experiments referenced throughout this chapter show that you can achieve good signal quality, if you follow the memory vendors recommended ODT settings. Although the DDR2 simulations and experimental results in this chapter are based on the Stratix II High Speed High Density Board, you can apply the general principles to any dual-DIMM design. The addition of dynamic OCT in Stratix III and Stratix IV devices has simplified the board design further by removing the need for the previously required FPGA end discrete parallel termination.

Even though this chapter covers several combinations of ODT and OCT termination, it is critical that as a board designer you perform system specific simulations to ensure good signal integrity in your dual-DIMM SDRAM designs.

References

- JEDEC Standard Publication JESD79-2, *DDR2 SDRAM Specification*, JEDEC Solid State Technology Association.
- JEDEC Standard Publication JESD8-15A, *Stub Series Termination Logic for 1.8 V (SSTL-18)*, JEDEC Solid State Technology Association.
- Micron TN-47-01: *DDR2 Design Guide for Two-DIMM Systems*.
- Micron TN-41-08: *DDR3 Design Guide for Two-DIMM Systems*
- Micron TN-47-17: *DDR2 SODIMM Optimized Address/Command Nets Introduction*.
- Samsung Electronics *Application Note: DDR2 ODT Control*
- *High-Speed Digital Design – A Handbook of Black Magic*, Howard Johnson and Martin Graham, Prentice Hall, 1993.
- *Circuits Interconnects, and Packaging for VLSI*, H.B. Bakoglu, Addison Wesley, 1990.
- *Signal Integrity – Simplified*, Eric Bogatin, Prentice Hall Modern Semiconductor Design Series, 2004.
- *Handbook of Digital Techniques for High-Speed Design*, Tom Granberg, Prentice Hall Modern Semiconductor Design Series, 2004.
- *Termination Placement in PCB Design How Much Does it Matter?*, Doug Brooks, UltraCAD Design Inc.
- *PC4300 DDR2 SDRAM Unbuffered DIMM Design Specification, Revision 0.5, Oct 30, 2003*.

This chapter provides guidelines for you to improve your system's signal integrity and layout guidelines to help successfully implement a QDR II or QDR II+ SRAM interface in your system.

The QDR II and QDR II+ SRAM Controller with UniPHY intellectual property (IP) enables you to implement QDR II and QDR II+ interfaces with Arria II GX, Stratix III, and Stratix IV devices.



In this chapter, QDR II SRAM refers to both QDR II and QDR II+ SRAM unless stated otherwise.

This chapter focuses on the following key factors that affect signal integrity:

- I/O standards
- QDR II SRAM configurations
- Signal terminations
- Printed circuit board (PCB) layout guidelines

I/O Standards

QDR II SRAM interface signals use one of the following JEDEC I/O signalling standards:

- HSTL-15—provides the advantages of lower power and lower emissions.
- HSTL-18—provides increased noise immunity with slightly greater output voltage swings.



To select the most appropriate standard for your interface, refer to the *Arria II GX Devices Data Sheet : Electrical Characteristics* chapter in the *Arria II Device Handbook*, *Stratix III Device Datasheet: DC and Switching Characteristics* chapter in the *Stratix III Device Handbook*, or the *Stratix IV Device Datasheet DC and Switching Characteristics* chapter in the *Stratix IV Device Handbook*.

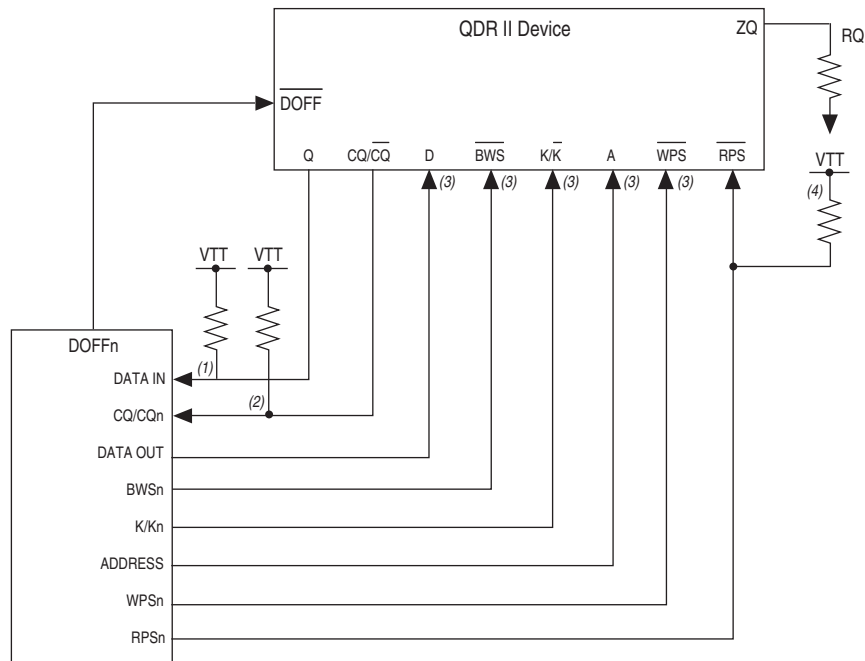
Altera QDR II SRAM Controller with UniPHY IP defaults to HSTL 1.5 V Class I outputs and HSTL 1.5 V inputs.

QDR II SRAM Configurations

The QDR II SRAM Controller with UniPHY IP supports interfaces with a single device, and two devices in a width expansion configuration up to maximum width of 72 bits.

Figure 3-1 shows the main signal connections between the FPGA and a single QDR II SRAM component.

Figure 3-1. Configuration With A Single QDR II SRAM Component

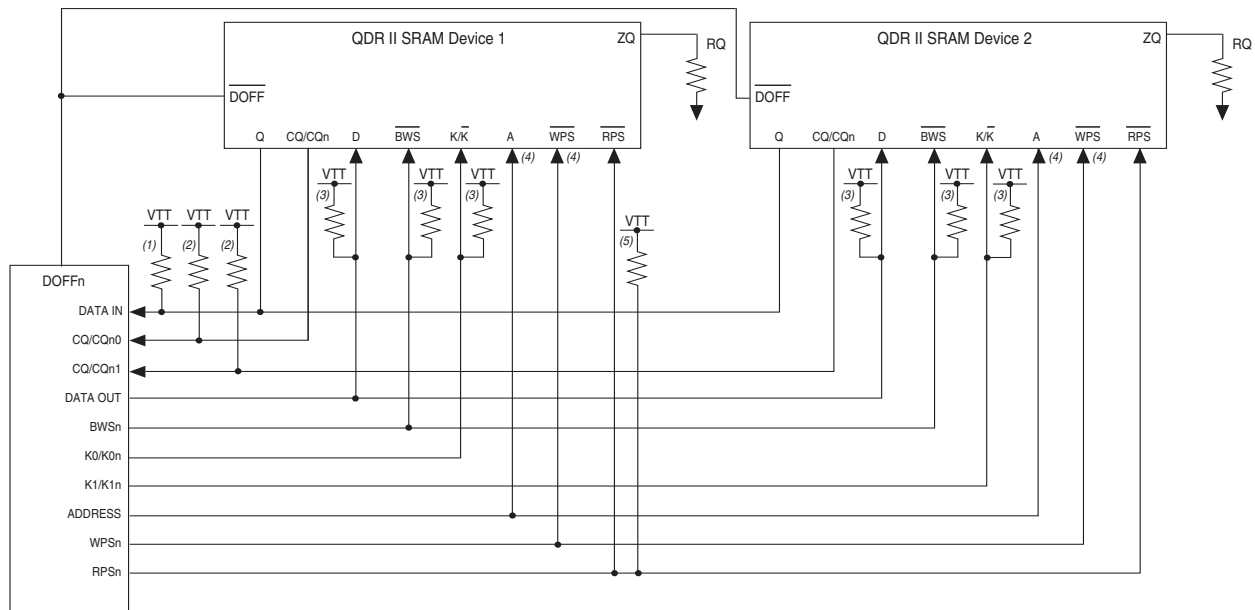


Notes to Figure 3-1:

- (1) Use external discrete termination only for data inputs targeting Arria II GX devices that do not support parallel OCT. For Stratix III and Stratix IV devices, use parallel OCT.
- (2) Use external discrete termination only for CQ/CQ# targeting Arria II GX devices, or for any device using ×36 emulated mode.
- (3) Use external discrete termination for this signal, as shown for RPS.
- (4) Use external discrete termination with fly-by placement to avoid stubs.

Figure 3-2 shows the main signal connections between the FPGA and two QDR II SRAM components in a width expansion configuration.

Figure 3-2. Configuration With Two QDR II SRAM Components In A Width Expansion Configuration

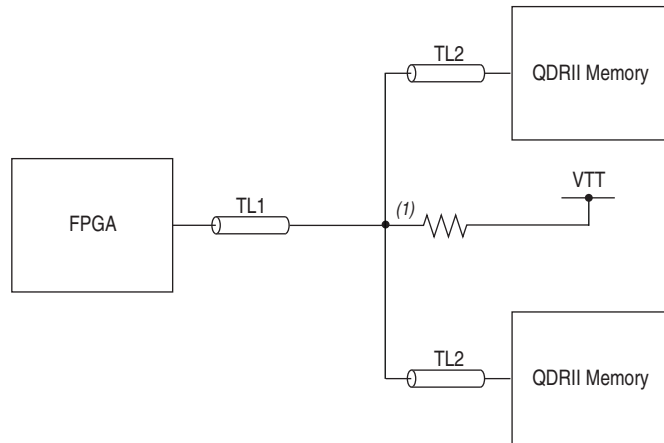


Notes to Figure 3-2:

- (1) Use external discrete termination only for data inputs targeting Arria II GX devices that do not support parallel OCT. For Stratix III and Stratix IV devices, use parallel OCT.
- (2) Use external discrete termination only for CQ/CQ# targeting Arria II GX devices, or for any device using ×36 emulated mode.
- (3) Use external discrete termination for data outputs, BWSn, and K/K# clocks with fly-by placement to avoid stubs.
- (4) Use external discrete termination for this signal, as shown for RPS.
- (5) Use external discrete termination at the trace split of the balanced T or Y topology.

Figure 3-3 shows the detailed balanced topology recommended for the address and command signals in the width expansion configuration.

Figure 3-3. External Parallel Termination for Balanced Topology



Note to Figure 3-3:

- (1) To minimize the reflections and parallel impedance discontinuity seen by the signal, place the trace split close to the QDR II SRAM memory components. Keep TL2 short so that the QDR II SRAM components appear as a lumped load.

Signal Terminations

Arria II GX, Stratix III and Stratix IV devices offer on-chip termination (OCT) technology.

Table 3-1 summarizes the extent of OCT support for each device.

Table 3-1. On-Chip Termination Schemes (Note 1)

Termination Scheme	HSTL-15 and HSTL-18	FPGA Device					
		Arria II GX		Arria II GZ, Stratix III, and Stratix IV		Stratix V	
		Column I/O	Row I/O	Column I/O	Row I/O	Column I/O	Row I/O
On-Chip Series Termination without Calibration	Class I	50	50	50	50	—	—
On-Chip Series Termination with Calibration	Class I	50	50	50	50	—	—
On-Chip Parallel Termination with Calibration	Class I	—	—	50	50	50	50

Note to Table 3-1:

- (1) This table provides information about HSTL-15 and HSTL-18 standards because these are the supported I/O standards for QDR II SRAM memory interfaces by Altera FPGAs.

On-chip series (R_S) termination is supported only on output and bidirectional buffers, while on-chip parallel (R_T) termination is supported only on input and bidirectional buffers. Because QDR II SRAM interfaces have unidirectional data paths, dynamic OCT is not required.

For Arria II GX, Stratix III and Stratix IV devices, the HSTL Class I I/O calibrated terminations are calibrated against $50\ \Omega$ 1% resistors connected to the R_{UP} and R_{DN} pins in an I/O bank with the same VCCIO as the QDR II SRAM interface. The calibration occurs at the end of the device configuration.

QDR II SRAM controllers have a ZQ pin which is connected via a resistor R_Q to ground. Typically the QDR II SRAM output signal impedance is $0.2 \times R_Q$. Refer to the QDR II SRAM device data sheet for more information.



For information about OCT, refer to the *I/O Features in Arria II GX Devices* chapter in the *Arria II GX Device Handbook*, *Stratix III Device I/O Features* chapter in the *Stratix III Device Handbook* and the *I/O Features in Stratix IV Devices* chapter in the *Stratix IV Device Handbook*.

The following section shows HyperLynx simulation eye diagrams to demonstrate signal termination options. Altera strongly recommends signal terminations to optimize signal integrity and timing margins, and to minimize unwanted emissions, reflections, and crosstalk.

All of the eye diagrams shown in this section are for a $50\ \Omega$ trace with a propagation delay of 720 ps which is approximately a 4-inch trace on a standard FR4 PCB. The signal I/O standard is HSTL-15.

For point-to-point signals, Altera recommends that you place a fly-by termination by terminating at the end of the transmission line after the receiver to avoid unterminated stubs. The guideline is to place the fly-by termination within 100 ps propagation delay of the receiver.

Although not recommended, you can place the termination before the receiver, which leaves an unterminated stub. The stub delay is critical because the stub between the termination and the receiver is effectively unterminated, causing additional ringing and reflections. Stub delays should be less than 50 ps.

The eye diagrams shown in this section show the best case achievable and do not take into account PCB vias, crosstalk and other degrading effects such as variations in the PCB structure due to manufacturing tolerances.



Simulate your design to ensure correct functionality.

Output from the FPGA to the QDR II SRAM Component

The following output signals are from the FPGA to the QDR II SRAM component:

- write data
- byte write select (BWS_n)
- address
- control (WPS_n and RPS_n)
- clocks, K and $K\#$

Altera recommends that you terminate the write clocks, K and $K\#$, with a single-ended fly-by $50\ \Omega$ parallel termination to V_{TT} . However, simulations show that you can consider a differential termination if the clock pair is well matched and routed differentially.

The HyperLynx simulation eye diagrams show simulation cases of write data and address signals with termination options. The QDR II SRAM write data is double data rate. The QDR II SRAM address is either double data rate (burst length of 2) or single data rate (burst length of 4).

Simulations show that lowering the drive strength does not make a significant difference to the eye diagrams. All eye diagrams are shown at the QDR II SRAM device receiver pin.

Figure 3-4 shows the fly-by terminated signal using Stratix IV Class I HSTL-15 with calibrated $50\ \Omega$ OCT output driver.

Figure 3-4. Write Data Simulation at 400 MHz with Fly-By $50\ \Omega$ Parallel Termination to V_{TT}




Figure 3-5 shows an unterminated signal using Stratix IV Class I HSTL-15 with a calibrated $50\ \Omega$ OCT output driver. This unterminated solution is not recommended.

Figure 3-5. Write Data Simulation at 400 MHz with No Far-End Termination




Figure 3-6 shows an unterminated signal at a lower frequency of 250 MHz using Arria II GX Class I HSTL-15 with calibrated 50 Ω OCT output driver. This unterminated solution may be passable for some systems, but is shown so that you can compare against the superior quality of the terminated signal in Figure 3-4.

Figure 3-6. Write Data Simulation at 250 MHz with No Far-End Termination




Figure 3-7 shows an unterminated signal at a frequency of 175 MHz with a point-to-point connection. QDR II SRAM interfaces using Stratix IV devices have a maximum supported frequency of 350 MHz. For QDR II SRAM with burst length of four interfaces, the address signals are effectively single data rate at 175 MHz. This unterminated solution is not recommended but can be considered. The FPGA output driver is Class I HSTL-15 with a calibrated 50 Ω OCT.

Figure 3-7. Address Simulation for QDR II SRAM Burst Length of 4 at 175 MHz with No Far-End Termination





Figure 3-8 shows a typical topology, which are used for two components in width expansion mode. Altera recommends that you match the stubs TL20 and TL22, but you can allow small differences allowed to achieve acceptable signal integrity.

Figure 3-8. Address for QDR II SRAM Burst Length of 2 in Width Expansion Mode Topology



The eye diagrams in Figure 3-9 and Figure 3-10 use the topology shown in Figure 3-8. The eye diagram in Figure 3-11 uses the topology shown in Figure 3-8 without the V_{TT} termination, R9 and TL21.

Figure 3-9 shows an address signal at a frequency of 400 MHz with parallel $50\ \Omega$ termination to V_{TT} for QDR II SRAM burst length of 2 width expansion using Stratix IV Class I HSTL-15 12 mA driver and fly-by $50\ \Omega$ parallel termination to V_{TT} .

Figure 3-9. Address Simulation Using Stratix IV Class I HSTL-15 12 mA Driver and Fly-by $50\ \Omega$ Parallel Termination to V_{TT}




Figure 3-10 shows an address signal at a frequency of 400 MHz with parallel $50\ \Omega$ termination to V_{TT} for QDR II SRAM burst length of 2 width expansion using Stratix IV Class I HSTL-15 with $50\ \Omega$ calibration driver and fly-by $50\ \Omega$ parallel termination to V_{TT} . The waveform eye is significantly improved compared to the maximum (12mA) drive strength case.

Figure 3-10. Address Simulation Using Stratix IV Class I HSTL-15 $50\ \Omega$ Calibration Driver and Fly-by $50\ \Omega$ Parallel Termination to V_{TT}





Figure 3-11 shows an unterminated address signal at a frequency of 400 MHz for QDR II SRAM burst length of 2 width expansion using Stratix IV Class I HSTL-15 with 50 Ω calibration driver. This unterminated address has small eye and is not recommended.

Figure 3-11. Address Simulation Using Stratix IV Class I HSTL-15 50 Ω Calibration Driver and No Termination



Input to the FPGA from the QDR II SRAM Component

The QDR II SRAM component drives the following input signals into the FPGA:

- read data
- echo clocks, CQ/CQ#

For point-to-point signals, Altera recommends that you use the FPGA parallel OCT wherever possible. For devices that do not support parallel OCT (Arria II GX), and for $\times 36$ emulated configuration CQ/CQ# termination, Altera recommends that you use a fly-by 50 Ω parallel termination to V_{TT} . Although not recommended, you can use parallel termination with a short stub of less than 50 ps propagation delay as an alternative option. The input echo clocks, CQ and CQ# must not use a differential termination.

The eye diagrams are shown at the FPGA receiver pin, and the QDR II SRAM output driver is Class I HSTL-15 using its ZQ calibration of 50 Ω . The QDR II SRAM read data is double data rate.

Figure 3-12 shows the ideal case of a fly-by terminated signal using 50 Ω calibrated parallel OCT with Stratix IV device.

Figure 3-12. Read Data Simulation at 400 MHz with 50 Ω Parallel OCT Termination




Figure 3-13 shows an external discrete component fly-by terminated signal at a lower frequency of 250 MHz using an Arria II GX device.

Figure 3-13. Read Data Simulation at 250 MHz with Fly-By Parallel 50 Ω Termination





Figure 3-14 shows an unterminated signal at a lower frequency of 250 MHz using an Arria II GX device. This unterminated solution is not recommended but is shown so that you can compare against the superior quality of the terminated signal in Figure 3-13.

Figure 3-14. Read Data Simulation at 250 MHz with No Far-End Termination



Termination Schemes

Table 3–2 and Table 3–3 provide the recommended termination schemes for major QDR II SRAM memory interface signals, which include write data (D), byte write select (BWS), read data (Q), clocks (K, K#, CQ, and CQ#), address and command (WPS and RPS).

Table 3–2. Termination Recommendations for Arria II GX Devices

Signal Type	HSTL 15/18 Standard (1), (2)	FPGA End Discrete Termination	Memory End Termination
K/K# Clocks	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
Write Data	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
BWS	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
Address (3), (4)	Class I Max Current	—	50 Ω Parallel to V_{TT}
WPS, RPS (3), (4)	Class I Max Current	—	50 Ω Parallel to V_{TT}
CQ/CQ#	Class I	50 Ω Parallel to V_{TT}	ZQ50
CQ/CQ# ×36 emulated (5)	Class I	50 Ω Parallel to V_{TT}	ZQ50
Read Data (Q)	Class I	50 Ω Parallel to V_{TT}	ZQ50
QVLD (6)	—	—	ZQ50

Notes to Table 3–2:

- (1) R is effective series output impedance.
- (2) CAL is calibrated OCT.
- (3) For width expansion configuration, the address and control signals are routed to 2 devices. Recommended termination is 50 Ω parallel to V_{TT} at the trace split of a balanced T or Y routing topology. For 400 MHz burst length 2 configurations where the address signals are double data rate, it is recommended to use a clamshell placement of the two QDR II SRAM components to achieve minimal stub delays and optimum signal integrity. Clamshell placement is when two devices overlay each other by being placed on opposite sides of the PCB.
- (4) The UniPHY default IP setting for this output is Max Current. A Class I 50 Ω output with calibration output is typically optimal in single load topologies.
- (5) For ×36 emulated mode, the recommended termination for the CQ/CQ# signals is a 50 Ω parallel termination to V_{TT} at the trace split, refer to Figure 3–15. Altera recommends that you use this termination when ×36 DQ/QS groups are not supported in the FPGA.
- (6) QVLD is not used in the QDR II or QDR II+ SRAM with UniPHY implementations.

Table 3–3. Termination Recommendations for Stratix III and Stratix IV Devices (Part 1 of 2)

Signal Type	HSTL 15/18 Standard (1), (2), (3)	FPGA End Discrete Termination	Memory End Termination
K/K# Clocks	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
Write Data	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
BWS	Class I R50 CAL	—	50 Ω Parallel to V_{TT}
Address (4), (5)	Class I Max Current	—	50 Ω Parallel to V_{TT}
WPS, RPS (4), (5)	Class I Max Current	—	50 Ω Parallel to V_{TT}
CQ/CQ#	Class I P50 CAL	—	ZQ50
CQ/CQ# ×36 emulated (6)	—	50 Ω Parallel to V_{TT}	ZQ50
Read Data (Q)	Class I P50 CAL	—	ZQ50

Table 3-3. Termination Recommendations for Stratix III and Stratix IV Devices (Part 2 of 2)

Signal Type	HSTL 15/18 Standard (1), (2), (3)	FPGA End Discrete Termination	Memory End Termination
QVLD (7)	Class I P50 CAL	—	ZQ50

Notes to Table 3-3:

- (1) R is effective series output impedance.
- (2) P is effective parallel input impedance.
- (3) CAL is calibrated OCT.
- (4) For width expansion configuration, the address and control signals are routed to 2 devices. Recommended termination is $50\ \Omega$ parallel to V_{TT} at the trace split of a balanced T or Y routing topology. For 400 MHz burst length 2 configurations where the address signals are double data rate, it is recommended to use a "clam shell" placement of the two QDR II SRAM components to achieve minimal stub delays and optimum signal integrity. "Clam shell" placement is when two devices overlay each other by being placed on opposite sides of the PCB.
- (5) The UniPHY default IP setting for this output is Max Current. A Class 1 $50\ \Omega$ output with calibration output is typically optimal in single load topologies.
- (6) For $\times 36$ emulated mode, the recommended termination for the $CQ/CQ\#$ signals is a $50\ \Omega$ parallel termination to V_{TT} at the trace split, refer to Figure 3-15. Altera recommends that you use this termination when $\times 36$ DQ/DQS groups are not supported in the FPGA.
- (7) QVLD is not used in the QDR II or QDR II+ SRAM Controller with UniPHY implementations.

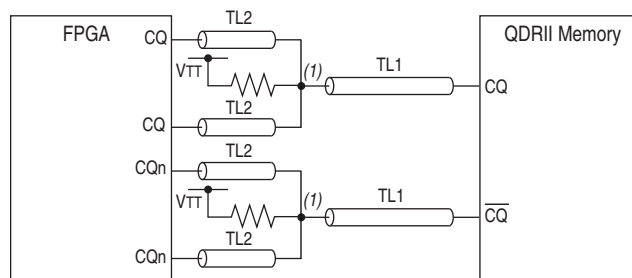


Altera recommends that you simulate your specific design for your system to ensure good signal integrity.



The QDR II SRAM Controller with UniPHY IP does not support the on-die termination (ODT) functionality featured in some QDR II SRAM components.

For a $\times 36$ QDR II SRAM interface that uses an emulated mode of two $\times 18$ DQS groups in the FPGA, there are two $CQ/CQ\#$ connections at the FPGA and a single $CQ/CQ\#$ output from the QDR II SRAM device. Altera recommends that you use a balanced T topology with the trace split close to the FPGA and a parallel termination at the split, as shown in Figure 3-15.

Figure 3-15. Emulated $\times 36$ Mode CQ/CQn Termination Topology**Note to Figure 3-15:**

- (1) To minimize the reflections and parallel impedance discontinuity seen by the signal, place the trace split close to the FPGA device. Keep TL2 short so that the FPGA inputs appear as a lumped load.



For more information about $\times 36$ emulated modes, refer to *Exceptions for $\times 36$ Emulated QDR II and QDR II+ SRAM Interfaces in Arria II GX, Stratix III, and Stratix IV Devices* in the *Device and Pin Planning* chapter in volume 2 of the *External Memory Interface Handbook*.

PCB Layout Guidelines

Table 3-4 summarizes QDR II and QDR II SRAM general routing layout guidelines.



The following layout guidelines include several +/- length based rules. These length based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristics of your PCB implementation. They do not include any margin for crosstalk.



Altera recommends that you get accurate time base skew numbers when you simulate your specific implementation.

Table 3-4. QDR II and QDR II+ SRAM Layout Guidelines (Part 1 of 2)

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> ■ All signal planes must be 50 Ω, single-ended, $\pm 10\%$. ■ All signal planes must be 100 Ω, differential $\pm 10\%$. ■ Remove all unused via pads, because they cause unwanted capacitance.
Decoupling Parameter	<ul style="list-style-type: none"> ■ Use 0.1 μF in 0402 size to minimize inductance. ■ Make V_{TT} voltage decoupling close to pull-up resistors. ■ Connect decoupling caps between V_{TT} and ground. ■ Use a 0.1 μF cap for every other V_{TT} pin. ■ Verify your capacitive decoupling using the Altera Power Distribution Network (PDN) Design tool.
Power	<ul style="list-style-type: none"> ■ Route GND, 1.5 V/1.8 V as planes. ■ Route V_{CCIO} for memories in a single split plane with at least a 20-mil (0.020 inches or 0.508 mm) gap of separation. ■ Route V_{TT} as islands or 250-mil (6.35-mm) power traces. ■ Route all oscillators and PLL power as islands or 100-mil (2.54-mm) power traces.
General Routing	<ul style="list-style-type: none"> ■ All specified delay matching requirements include PCB trace delays, different layer propagation, velocity variance, and crosstalk. To minimize PCB layer propagation variance, Altera recommends that signals from the same net group always be routed on the same layer. If signals of the same net group must be routed on different layers with the same impedance characteristic, you must simulate your worst case PCB trace tolerances to ascertain actual propagation delay differences. Typical later to later trace delay variations are of 15 ps/inch order. ■ Use 45° angles (not 90° corners). ■ Avoid T-Junctions for critical nets or clocks. ■ Avoid T-junctions greater than 150 ps (approximately 500 mils, 12.7 mm). ■ Disallow signals across split planes. ■ Restrict routing other signals close to system reset signals. ■ Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks.

Table 3–4. QDR II and QDR II+ SRAM Layout Guidelines (Part 2 of 2)

Parameter	Guidelines
Clock Routing	<ul style="list-style-type: none"> Route clocks on inner layers with outer-layer run lengths held to under 150 ps (approximately 500 mils, 12.7 mm). These signals should maintain a 10-mil (0.254 mm) spacing from other nets. Clocks should maintain a length-matching between clock pairs of ± 5 ps or approximately ± 25 mils (0.635 mm). Complementary clocks should maintain a length-matching between P and N signals of ± 2 ps or approximately ± 10 mils (0.254 mm). Keep the distance from the pin on the QDR II SRAM component to stub termination resistor (V_{TT}) to less than 50 ps (approximately 250 mils, 6.35 mm) for the K, K# clocks. Keep the distance from the pin on the QDR II SRAM component to fly-by termination resistor (V_{TT}) to less than 100 ps (approximately 500 mils, 12.7 mm) for the K, K# clocks. Keep the distance from the pin on the FPGA component to stub termination resistor (V_{TT}) to less than 50 ps (approximately 250 mils, 6.35 mm) for the echo clocks, CQ, CQ#, if they require an external discrete termination. Keep the distance from the pin on the FPGA component to fly-by termination resistor (V_{TT}) to less than 100 ps (approximately 500 mils, 12.7 mm) for the echo clocks, CQ, CQ#, if they require an external discrete termination.
External Memory Routing Rules	<ul style="list-style-type: none"> Keep the distance from the pin on the QDR II SRAM component to stub termination resistor (V_{TT}) to less than 50 ps (approximately 250 mils, 6.35 mm) for the write data, byte write select and address/command signal groups. Keep the distance from the pin on the QDR II SRAM component to fly-by termination resistor (V_{TT}) to less than 100 ps (approximately 500 mils, 12.7 mm) for the write data, byte write select and address/command signal groups. Keep the distance from the pin on the FPGA (Arria II GX) to stub termination resistor (V_{TT}) to less than 50 ps (approximately 250 mils, 6.35 mm) for the read data signal group. Keep the distance from the pin on the FPGA (Arria II GX) to fly-by termination resistor (V_{TT}) to less than 100 ps (approximately 500 mils, 12.7 mm) for the read data signal group. Parallelism rules for the QDR II SRAM data/address/command groups are as follows: <ul style="list-style-type: none"> 4 mils for parallel runs < 0.1 inch (approximately 1\times spacing relative to plane distance). 5 mils for parallel runs < 0.5 inch (approximately 1\times spacing relative to plane distance). 10 mils for parallel runs between 0.5 and 1.0 inches (approximately 2\times spacing relative to plane distance). 15 mils for parallel runs between 1.0 and 6.0 inch (approximately 3\times spacing relative to plane distance).
Maximum Trace Length	<ul style="list-style-type: none"> Keep the maximum trace length of all signals from the FPGA to the QDR II SRAM components to 6 inches.

Using the layout guidelines in [Table 3–4](#), Altera recommends the following layout approach:

1. Route the K/K# clocks and set the clocks as the target trace propagation delays for the output signal group.
2. Route the write data output signal group (write data, byte write select), ideally on the same layer as the K/K# clocks, to within ± 10 ps skew of the K/K# traces.

3. Route the address/control output signal group (address, RPS, WPS), ideally on the same layer as the $\kappa/\kappa\#$ clocks, to within ± 20 ps skew of the $\kappa/\kappa\#$ traces.
4. Route the CQ/CQ# clocks and set the clocks as the target trace propagation delays for the input signal group.
5. Route the read data output signal group (read data), ideally on the same layer as the CQ/CQ# clocks, to within ± 10 ps skew of the CQ/CQ# traces.
6. The output and input groups do not need to have the same propagation delays, but they must have all the signals matched closely within the respective groups.

Table 3-5 and Table 3-6 show the typical margins for QDR II and QDR II+ SRAM interfaces, with the assumption that there is zero skew between the signal groups.

Table 3-5. Typical Worst Case Margins for QDR II SRAM Interfaces of Burst Length 2

Device	Speed Grade	Frequency (MHz)	Typical Margin Address/Command (ps)	Typical Margin Write Data (ps)	Typical Margin Read Data (ps)
Arria II GX	I5	250	± 240	± 80	± 170
Arria II GX ×36 emulated	I5	200	± 480	± 340	± 460
Stratix IV	—	350	—	—	—
Stratix IV ×36 emulated	C2	300	± 320	± 170	± 340

Table 3-6. Typical Worst Case Margins for QDR II+ SRAM Interfaces of Burst Length 4

Device	Speed Grade	Frequency (MHz)	Typical Margin Address/Command (ps) (1)	Typical Margin Write Data (ps)	Typical Margin Read Data (ps)
Arria II GX	I5	250	± 810	± 150	± 130
Arria II GX ×36 emulated	I5	200	± 1260	± 410	± 420
Stratix IV	C2	400	± 550	± 10	± 80
Stratix IV ×36 emulated	C2	300	± 860	± 180	± 300

Note to Table 3-6:

(1) The QDR II+ SRAM burst length of 4 designs have greater margins on the address signals because they are single data rate.

Other devices and speed grades typically show higher margins than the ones in Table 3-5 and Table 3-6.



Altera recommends that you create your project with a fully implemented QDR II or QDR II+ SRAM Controller with UniPHY interface, and observe the interface timing margins to determine the actual margins for your design.

Although the recommendations in this chapter are based on simulations, you can apply the same general principles when determining the best termination scheme, drive strength setting, and loading style to any board designs. Even armed with this knowledge, it is still critical that you perform simulations, either using IBIS or HSPICE models, to determine the quality of signal integrity on your designs.

This chapter provides guidelines for you to improve your system's signal integrity and layout guidelines to help successfully implement an RLD RAM II interface in your system.

The RLD RAM II Controller with UniPHY intellectual property (IP) enables you to implement Common I/O (CIO) RLD RAM II interfaces with Stratix III, Stratix IV, and Stratix V devices. You can implement Separate I/O (SIO) RLD RAM II interfaces with the ALTDQ_DQS or ALTDQ_DQS2 megafunctions.

This chapter focuses on the following key factors that affect signal integrity:

- I/O standards
- RLD RAM II configurations
- Signal terminations
- Printed circuit board (PCB) layout guidelines

I/O Standards

RLD RAM II interface signals use one of the following JEDEC I/O signalling standards:

- HSTL-15—provides the advantages of lower power and lower emissions.
- HSTL-18—provides increased noise immunity with slightly greater output voltage swings.



To select the most appropriate standard for your interface, refer to the *Device Datasheet for Arria II Devices* chapter in the *Arria II Device Handbook*, the *Stratix III Device Datasheet: DC and Switching Characteristics* chapter in the *Stratix III Device Handbook*, the *DC and Switching Characteristics for Stratix IV Devices* chapter in the *Stratix IV Device Handbook*, or the *DC and Switching Characteristics for Stratix V Devices* chapter in the *Stratix V Device Handbook*.


The RLD RAM II Controller with UniPHY IP defaults to HSTL 1.8 V Class I outputs and HSTL 1.8 V inputs.

RLD RAM II Configurations

The RLD RAM II Controller with UniPHY IP supports interfaces for CIO RLD RAM II with a single device, and two devices in a width expansion configuration up to maximum width of 72 bits. This chapter focusses the layout and guidelines for CIO RLD RAM II interfaces. However, the termination and layout principles for SIO RLD RAM II interfaces are similar to CIO RLD RAM II, except that SIO RLD RAM II interfaces have unidirectional data buses.

Figure 4-1 shows the main signal connections between the FPGA and a single CIO RLD RAM II component.

Figure 4-1. Configuration With A Single CIO RLD RAM II Component

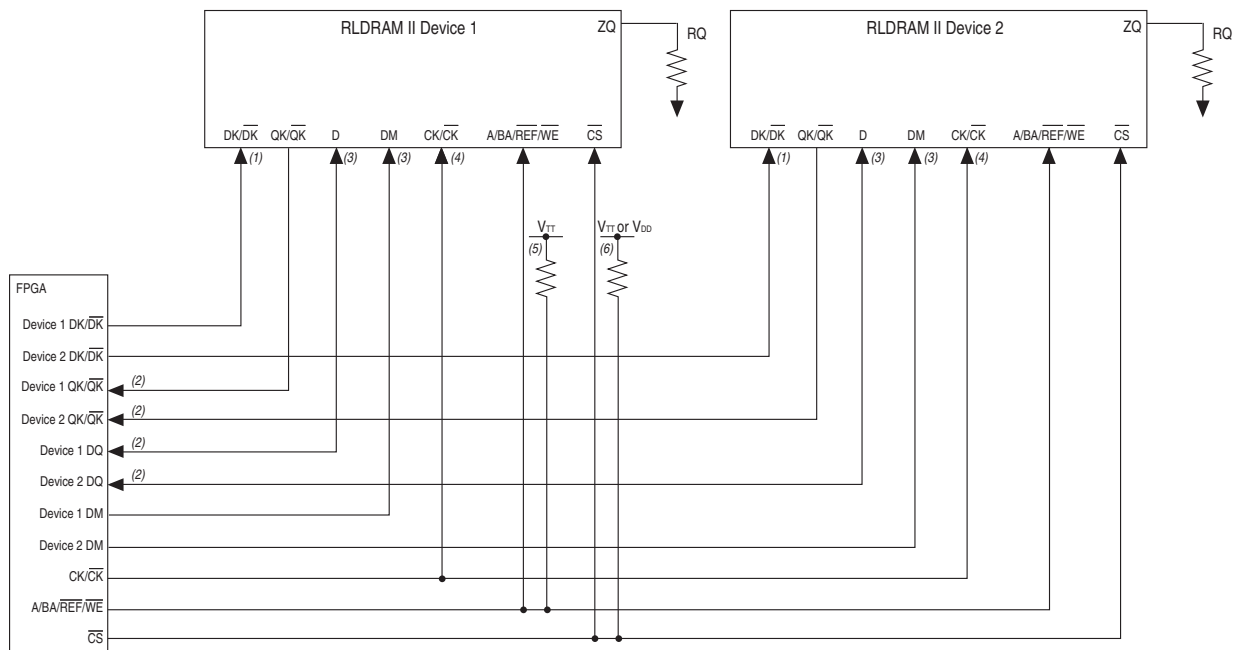


Notes to Figure 4-1:

- (1) Use external differential termination on DK/DK# and CK/CK#.
- (2) Use FPGA parallel on-chip termination (OCT) for terminating QK/QK# and DQ on reads.
- (3) Use RLD RAM II component on-die termination (ODT) for terminating DQ and DM on writes.
- (4) Use external discrete termination with fly-by placement to avoid stubs.
- (5) Use external discrete termination for this signal, as shown for REF.
- (6) Use external discrete termination, as shown for REF, but you may require a pull-up resistor to V_{DD} as an alternative option. Refer to the RLD RAM II device data sheet for more information about RLD RAM II power-up sequencing.

Figure 4-2 shows the main signal connections between the FPGA and two CIO RLDram II components in a width expansion configuration.

Figure 4-2. Configuration With Two CIO RLDram II Components In A Width Expansion Configuration



Notes to Figure 4-2:

- (1) Use external differential termination on DK/DK#.
- (2) Use FPGA parallel OCT for terminating QK/QK# and DQ on reads.
- (3) Use RLDram II component ODT for terminating DQ and DM on writes.
- (4) Use external dual 200 Ω differential termination.
- (5) Use external discrete termination at the trace split of the balanced T or Y topology.
- (6) Use external discrete termination at the trace split of the balanced T or Y topology, but you may require a pull-up resistor to V_{DD} as an alternative option. Refer to the RLDram II device data sheet for more information about RLDram II power-up sequencing.

Signal Terminations

Stratix III, Stratix IV, and Stratix V devices offer OCT technology.

Table 4-1 summarizes the extent of OCT support for each device.

Table 4-1. On-Chip Termination Schemes (Part 1 of 2)

Termination Scheme	HSTL-15 and HSTL-18	FPGA Device	
		Arria II GZ, Stratix III, and Stratix IV	Stratix V
		Row/Column I/O	Row/Column I/O
On-Chip Series Termination without Calibration	Class I	50	50
On-Chip Series Termination with Calibration	Class I	50	50 (1)

Table 4-1. On-Chip Termination Schemes (Part 2 of 2)

Termination Scheme	HSTL-15 and HSTL-18	FPGA Device	
		Arria II GZ, Stratix III, and Stratix IV	Stratix V
		Row/Column I/O	Row/Column I/O
On-Chip Parallel Termination with Calibration	Class I	50	50 (1)

Note to Table 4-1:

- (1) Although 50 Ω is the recommended option, Stratix V devices offer a wider range of calibrated termination impedances.

On-chip series (R_S) termination supports output buffers, and bidirectional buffers only when they are driving output signals. On-chip parallel (R_T) termination supports input buffers, and bidirectional buffers only when they are input signals. RLD RAM II CIO interfaces have bidirectional data paths. The UniPHY IP uses dynamic OCT on the datapath, which switches between series OCT for memory writes and parallel OCT for memory reads.

For Arria II GZ, Stratix III, and Stratix IV devices, the HSTL Class I I/O calibrated terminations are calibrated against 50 Ω 1% resistors connected to the R_{UP} and R_{DN} pins in an I/O bank with the same V_{CCIO} as the RLD RAM II interface. For Stratix V devices, the HSTL Class I I/O calibrated terminations are calibrated against 100 Ω or 240 Ω 1% resistors connected to the R_{ZQ} pins in an I/O bank with the same V_{CCIO} as the RLD RAM II interface.

The calibration occurs at the end of the device configuration.

RLD RAM II memory components have a ZQ pin which connects through a resistor RQ to ground. Typically the RLD RAM II output signal impedance is $0.2 \times RQ$. Refer to the RLD RAM II device data sheet for more information.



For information about OCT, refer to the *I/O Features in Arria II Devices* chapter in the *Arria II Device Handbook*, *Stratix III Device I/O Features* chapter in the *Stratix III Device Handbook*, the *I/O Features in Stratix IV Devices* chapter in the *Stratix IV Device Handbook*, or the *I/O Features in Stratix V Devices* chapter in the *Stratix V Device Handbook*.

The following section shows HyperLynx simulation eye diagrams to demonstrate signal termination options. Altera strongly recommends signal terminations to optimize signal integrity and timing margins, and to minimize unwanted emissions, reflections, and crosstalk.

All of the eye diagrams shown in this section are for a 50 Ω trace with a propagation delay of 600 ps which is approximately a 3.3-inch trace on a standard FR4 PCB. The signal I/O standard is HSTL-18.

The eye diagrams shown in this section show the best case achievable and do not take into account PCB vias, crosstalk and other degrading effects such as variations in the PCB structure due to manufacturing tolerances.



Simulate your design to ensure correct functionality.

Outputs from the FPGA to the RLD RAM II Component

The following output signals are from the FPGA to the RLD RAM II component:

- write data (DQ on the bidirectional data signals for CIO RLD RAM II)
- data mask (DM)
- address, bank address
- command (CS, WE, and REF)
- clocks (CK/CK# and DK/DK#)

For point-to-point single-ended signals requiring external termination, Altera recommends that you place a fly-by termination by terminating at the end of the transmission line after the receiver to avoid unterminated stubs. The guideline is to place the fly-by termination within 100 ps propagation delay of the receiver.

Although not recommended, you can place the termination before the receiver, which leaves an unterminated stub. The stub delay is critical because the stub between the termination and the receiver is effectively unterminated, causing additional ringing and reflections. Stub delays should be less than 50 ps.

Altera recommends that the differential clocks, CK, CK# and DK, DK#, use a differential termination at the end of the trace of the RLD RAM II component. Alternatively, you can terminate each clock output with a parallel termination to V_{TT} .

The HyperLynx simulation eye diagrams show simulation cases of write data, address, and chip-select signals with termination options. All eye diagrams are shown at the connection to the receiver device die.

Figure 4-3 shows the double data rate write data using a Stratix IV Class I HSTL-18 with calibrated 50 Ω OCT output driver and the nominal RLD RAM II ODT of 150 Ω .

Figure 4-3. Write Data Simulation at 400 MHz with RLD RAM II ODT




Figure 4-4 shows an address signal at a frequency of 200 MHz using Stratix IV Class I HSTL-18 with a calibrated 50 Ω OCT driver and a 100 ps fly-by 50 Ω parallel termination to V_{TT} .

Figure 4-4. Address Simulation Using Stratix IV Class I HSTL-18 50 Ω Calibration Driver and Fly-by 50 Ω Parallel Termination




Figure 4-5 shows an address signal at a frequency of 200 MHz using Stratix IV Class I HSTL-18 12 mA driver and a 50 ps stub 50 Ω parallel termination to V_{TT} .

Figure 4-5. Address Simulation Using Stratix IV Class I HSTL-18 50 Ω Calibration Driver and Stub 50 Ω Parallel Termination to V_{TT}





Figure 4-6 shows the chip-select signal at a frequency of 200 MHz using Stratix IV Class I HSTL-18 with a calibrated 50 Ω driver and a 10 K pull-up resistor to V_{DD} . The RLD RAM II power sequencing may require the chip selects to have a pull-up resistor. Refer to the RLD RAM II data sheet for further details.

Figure 4-6. Chip-Select Simulation Using Stratix IV Class I HSTL-18 50 Ω Calibration Driver and 10 K Pull-up Resistor to V_{DD}



For the RLD RAM II width expansion configuration for address and command, use the same principles recommended for “QDR II SRAM Interface Termination and Layout Guidelines” on page 3-1.

For external parallel termination recommended for a balanced T topology, refer to Figure 3-3 on page 3-4, and for HyperLynx simulation diagrams of the width expansion topology for address and command signals, refer to Figure 3-8 through Figure 3-11 on page 3-13.

Input to the FPGA from the RLD RAM II Component

The RLD RAM II component drives the following input signals into the FPGA:


- read data (DQ on the bidirectional data signals for CIO RLD RAM II)
- read clocks (QK/QK#)

Altera recommends that you use the FPGA parallel OCT to terminate the data on reads and read clocks.

The eye diagrams are shown at the FPGA die pin, and the RLD RAM II output driver is Class I HSTL-18 using its ZQ calibration of $50\ \Omega$. The RLD RAM II read data is double data rate.

Figure 4-7 shows the ideal case of a fly-by terminated signal using $50\ \Omega$ calibrated parallel OCT for a Stratix IV device.

Figure 4-7. Read Data Simulation at 400 MHz with $50\ \Omega$ Parallel OCT Termination



Termination Schemes

Table 4-2 shows the recommended termination schemes for major CIO RLD RAM II memory interface signals, which include data (DQ), data mask (DM), clocks (CK, CK#, DK, DK#, QK, and QK#), address, bank address, and command (WE#, REF#, and CS#).

Table 4-2. Termination Recommendations for Stratix III, Stratix IV, and Stratix V Devices (Part 1 of 2)

Signal Type	HSTL 15/18 Standard (1), (2), (3)	Memory End Termination
DK/DK# Clocks	Class I R50 CAL	$100\ \Omega$ Differential
QK/QK# Clocks	Class I P50 CAL	ZQ50
Data (Write)	Class I R50 CAL	ODT
Data (Read)	Class I P50 CAL	ZQ50
Data Mask	Class I R50 CAL	ODT

Table 4–2. Termination Recommendations for Stratix III, Stratix IV, and Stratix V Devices (Part 2 of 2)

Signal Type	HSTL 15/18 Standard (1), (2), (3)	Memory End Termination
CK/CK# Clocks	Class I R50 CAL	×1 = 100 Ω Differential (8) ×2 = 200 Ω Differential (9)
Address/Bank Address (4), (5)	Class I Max Current	50 Ω Parallel to V_{TT}
Command (WE, REF) (4), (5)	Class I Max Current	50 Ω Parallel to V_{TT}
Command (CS) (4), (5), (6)	Class I Max Current	50 Ω Parallel to V_{TT} or Pull-up to V_{DD}
QVLD (7)	Class I P50 CAL	ZQ50

Notes to Table 4–2:

- (1) R is effective series output impedance.
- (2) P is effective parallel input impedance.
- (3) CAL is calibrated OCT.
- (4) For width expansion configuration, the address and control signals are routed to 2 devices. Recommended termination is 50 Ω parallel to V_{TT} at the trace split of a balanced T or Y routing topology. Use a clamshell placement of the two RLD RAM II components to achieve minimal stub delays and optimum signal integrity. Clamshell placement is when two devices overlay each other by being placed on opposite sides of the PCB.
- (5) The UniPHY default IP setting for this output is Max Current. A Class I 50 Ω output with calibration output is typically optimal in single load topologies.
- (6) Altera recommends that you use a 50 Ω parallel termination to V_{TT} if your design meets the power sequencing requirements of the RLD RAM II component. Refer to the RLD RAM II data sheet for further information.
- (7) QVLD is not used in the RLD RAM II Controller with UniPHY implementations.
- (8) ×1 is a single-device load.
- (9) ×2 is a double-device load. An alternative option is to use a 100 Ω differential termination at the trace split.



Altera recommends that you simulate your specific design for your system to ensure good signal integrity.

PCB Layout Guidelines

Table 4–3 summarizes RLD RAM II general routing layout guidelines.



The following layout guidelines include several +/- length based rules. These length-based guidelines are for first order timing approximations if you cannot simulate the actual delay characteristics of your PCB implementation. They do not include any margin for crosstalk.



Altera recommends that you get accurate time base skew numbers when you simulate your specific implementation.

Table 4-3. RLD RAM II Layout Guidelines (Part 1 of 2)

Parameter	Guidelines
Impedance	<ul style="list-style-type: none"> ■ All signal planes must be 50 Ω, single-ended, $\pm 10\%$. ■ All signal planes must be 100 Ω, differential $\pm 10\%$. ■ Remove all unused via pads, because they cause unwanted capacitance.
Decoupling Parameter	<ul style="list-style-type: none"> ■ Use 0.1 μF in 0402 size to minimize inductance. ■ Make V_{TT} voltage decoupling close to pull-up resistors. ■ Connect decoupling caps between V_{TT} and ground. ■ Use a 0.1 μF cap for every other V_{TT} pin. ■ Verify your capacitive decoupling using the Altera Power Distribution Network (PDN) Design tool.
Power	<ul style="list-style-type: none"> ■ Route GND, 1.5 V/1.8 V as planes. ■ Route V_{CCIO} for memories in a single split plane with at least a 20-mil (0.020 inches or 0.508 mm) gap of separation. ■ Route V_{TT} as islands or 250-mil (6.35-mm) power traces. ■ Route oscillators and PLL power as islands or 100-mil (2.54-mm) power traces.
General Routing	<ul style="list-style-type: none"> ■ All specified delay matching requirements include PCB trace delays, different layer propagation, velocity variance, and crosstalk. To minimize PCB layer propagation variance, Altera recommends that signals from the same net group always be routed on the same layer. If you must route signals of the same net group on different layers with the same impedance characteristic, simulate your worst case PCB trace tolerances to ascertain actual propagation delay differences. Typical layer to layer trace delay variations are of 15 ps/inch order. ■ Use 45° angles (not 90° corners). ■ Avoid T-Junctions for critical nets or clocks. ■ Avoid T-junctions greater than 150 ps (approximately 500 mils, 12.7 mm). ■ Disallow signals across split planes. ■ Restrict routing other signals close to system reset signals. ■ Avoid routing memory signals closer than 0.025 inch (0.635 mm) to PCI or system clocks. ■ Match all signals within a given DQ group with a maximum skew of ± 10 ps or approximately ± 50 mils (0.254 mm) and route on the same layer.
Clock Routing	<ul style="list-style-type: none"> ■ Route clocks on inner layers with outer-layer run lengths held to under 150 ps (approximately 500 mils, 12.7 mm). ■ These signals should maintain a 10-mil (0.254 mm) spacing from other nets. ■ Clocks should maintain a length-matching between clock pairs of ± 5 ps or approximately ± 25 mils (0.635 mm). ■ Differential clocks should maintain a length-matching between P and N signals of ± 2 ps or approximately ± 10 mils (0.254 mm). ■ Space between different clock pairs should be at least three times the space between the traces of a differential pair.

Table 4-3. RLD RAM II Layout Guidelines (Part 2 of 2)

Parameter	Guidelines
Address and Command Routing	<ul style="list-style-type: none"> ■ To minimize crosstalk, route address, bank address, and command signals on a different layer than the data and data mask signals. ■ Do not route the differential clock signals close to the address signals. ■ Keep the distance from the pin on the RLD RAM II component to the stub termination resistor (V_{TT}) to less than 50 ps (approximately 250 mils, 6.35 mm) for the address/command signal group. ■ Keep the distance from the pin on the RLD RAM II component to the fly-by termination resistor (V_{TT}) to less than 100 ps (approximately 500 mils, 12.7 mm) for the address/command signal group.
External Memory Routing Rules	<ul style="list-style-type: none"> ■ Apply the following parallelism rules for the RLD RAM II data/address/command groups: <ul style="list-style-type: none"> ■ 4 mils for parallel runs < 0.1 inch (approximately 1× spacing relative to plane distance). ■ 5 mils for parallel runs < 0.5 inch (approximately 1× spacing relative to plane distance). ■ 10 mils for parallel runs between 0.5 and 1.0 inches (approximately 2× spacing relative to plane distance). ■ 15 mils for parallel runs between 1.0 and 3.3 inch (approximately 3× spacing relative to plane distance).
Maximum Trace Length	<ul style="list-style-type: none"> ■ Keep the maximum trace length of all signals from the FPGA to the RLD RAM II components to 600 ps (approximately 3,300 mils, 83.3 mm).

Using the layout guidelines in [Table 4-3](#), Altera recommends the following layout approach:

1. If the RLD RAM II interface has multiple DQ groups (×18 or ×36 RLD RAM II component or width expansion configuration), match all the $DK/DK\#$ and $QK/QK\#$ clocks as tightly as possible to optimize the timing margins in your design.
2. Route the $DK/DK\#$ write clock and $QK/QK\#$ read clock associated with a DQ group on the same PCB layer. Match these clock pairs to within ±5 ps.
3. Set the $DK/DK\#$ or $QK/QK\#$ clock as the target trace propagation delay for the associated data and data mask signals.
4. Route the data and data mask signals for the DQ group ideally on the same layer as the associated $QK/QK\#$ and $DK/DK\#$ clocks to within ±10 ps skew of the target clock.
5. Route the $CK/CK\#$ clocks and set as the target trace propagation delays for the address/command signal group. Match the $CK/CK\#$ clock to within ±50 ps of all the $DK/DK\#$ clocks.
6. Route the address/control signal group (address, bank address, CS, WE, and REF) ideally on the same layer as the $CK/CK\#$ clocks, to within ±20 ps skew of the $CK/CK\#$ traces.

This layout approach provides a good starting point for a design requirement of the highest clock frequency supported for the RLD RAM II interface.

Table 4-4 shows the typical margins for RLD RAM II interfaces, with the assumption that there is zero skew between the signal groups.

Table 4-4. Typical Worst Case Margins for CIO RLD RAM II Interfaces

Device	Speed Grade	Frequency (MHz)	Typical Margin Address/Command (ps)	Typical Margin Write Data (ps)	Typical Margin Read Data (ps)
Stratix IV	—	400	—	—	—
Stratix V	—	400	—	—	—

Other devices and speed grades typically show higher margins than those in Table 4-4.



Altera recommends that you create your project in the Quartus® II software with a fully implemented RLD RAM II Controller with UniPHY interface, and observe the interface timing margins to determine the actual margins for your design.

Although the recommendations in this chapter are based on simulations, you can apply the same general principles when determining the best termination scheme, drive strength setting, and loading style to any board designs. Even armed with this knowledge, it is still critical that you perform simulations, either using IBIS or HSPICE models, to determine the quality of signal integrity on your designs.

Table 5–1 shows the Altera-supported power estimation methods for external memory interfaces.

Table 5–1. Power Estimation Methods for External Memory Interfaces

Method	Vector Source	ALTMEMPHY Support	UniPHY Support	Accuracy	Estimation Time (1)
Early power estimator (EPE)	Not applicable	✓	✓	<div>Lowest</div> <div>↓</div> <div>Highest</div>	Fastest
Vector-less PowerPlay power analysis (PPPA)	Not applicable	✓	✓		<div>↓</div> <div>Slowest</div>
Vector-based PPPA	RTL simulation	✓	✓		
	Zero-delay simulation (2)	✓	✓		
	Timing simulation	(2)	(2)		

Note to Table 5–1:

- (1) To decrease the estimation time, you can skip power estimation during calibration. Power consumption during calibration is typically equivalent to power consumption during user mode.
- (2) Power analysis using timing simulation vectors is not supported.

When using Altera IP, you can use the zero-delay simulation method to analyze the power required for the external memory interface. Zero-delay simulation is as accurate as timing simulation for 95% designs (designs with no glitching). For a design with glitching, power may be under estimated.



For more information about zero-delay simulation, refer to the *Power Estimation and Analysis* section in the *Quartus II Handbook*.



The size of the vector file (.vcd) generated by zero-delay simulation of an Altera DDR3 SDRAM High-Performance Controller Example Design is 400 GB. The .vcd includes calibration and user mode activities. When vector generation of calibration phase is skipped, the vector size decreases to 1 GB.

To perform vector-based PPPA using zero-delay simulation, follow these steps:

1. Perform design compilation in the Quartus II software to generate your design's Netlist *<project_name>.vo*.



The *<project_name>.vo* is generated in the last stage of a compile EDA Netlist Writer.

2. In *<project_name>.vo*, search for the include statement for *<project_name>.sdo*, comment the statement out, and save the file.

3. Create a simulation script containing device model files and libraries and design specific files:
 - Netlist file for the design, `<project_name>.vo`
 - RTL or netlist file for the memory device
 - Testbench RTL file
4. Compile all the files.
5. Invoke simulator with commands to generate `.vcd` files.
6. Generate `.vcd` files for the parts of the design that contribute the most to power dissipation.
7. Run simulation
8. Use the generated `.vcd` files in PPPA tool as the signal activity input file.
9. Run PPPA



For more information about estimating power, refer to the *Power Estimation and Analysis* section in the *Quartus II Handbook*

This chapter provides additional information about the document and Altera.

Document Revision History

The following table shows the revision history for this document.

Date	Version	Changes
June 2011	3.0	<ul style="list-style-type: none"> Merged DDR2 and DDR3 chapters to DDR2 and DDR3 SDRAM Interface Termination and Layout Guidelines and updated with leveling information. Added Stratix V information.
December 2010	2.1	<ul style="list-style-type: none"> Added new chapter: RLDRAM II Interface Termination and Layout Guidelines. Added DDR3 SDRAM Interface Termination, Drive Strength, Loading, and Board Layout Guidelines chapter with Stratix V information.
July 2010	2.0	<ul style="list-style-type: none"> Added new chapter: QDR II SRAM Interface Termination and Layout Guidelines. Updated Arria II GX information.
April 2010	1.0	Initial release.

How to Contact Altera

To locate the most up-to-date information about Altera products, refer to the following table.








Contact (1)	Contact Method	Address
Technical support	Website	www.altera.com/support
Technical training	Website	www.altera.com/training
	Email	custrain@altera.com
Product literature	Website	www.altera.com/literature
Non-technical support (General) (Software Licensing)	Email	nacomp@altera.com
	Email	authorization@altera.com

Note to Table:

(1) You can also contact your local Altera sales office or sales representative.

Typographic Conventions

The following table shows the typographic conventions this document uses.

Visual Cue	Meaning
Bold Type with Initial Capital Letters	Indicate command names, dialog box titles, dialog box options, and other GUI labels. For example, Save As dialog box. For GUI elements, capitalization matches the GUI.
bold type	Indicates directory names, project names, disk drive names, file names, file name extensions, software utility names, and GUI labels. For example, \qdesigns directory, D: drive, and chiptrip.gdf file.
<i>Italic Type with Initial Capital Letters</i>	Indicate document titles. For example, <i>Stratix IV Design Guidelines</i> .
<i>italic type</i>	Indicates variables. For example, $n + 1$. Variable names are enclosed in angle brackets (< >). For example, <file name> and <project name>.pof file.
Initial Capital Letters	Indicate keyboard keys and menu names. For example, the Delete key and the Options menu.
“Subheading Title”	Quotation marks indicate references to sections within a document and titles of Quartus II Help topics. For example, “Typographic Conventions.”
Courier type	Indicates signal, port, register, bit, block, and primitive names. For example, data1, tdi, and input. The suffix n denotes an active-low signal. For example, resetn. Indicates command line commands and anything that must be typed exactly as it appears. For example, c:\qdesigns\tutorial\chiptrip.gdf. Also indicates sections of an actual file, such as a Report File, references to parts of files (for example, the AHDL keyword SUBDESIGN), and logic function names (for example, TRI).
	An angled arrow instructs you to press the Enter key.
1., 2., 3., and a., b., c., and so on	Numbered steps indicate a list of items when the sequence of the items is important, such as the steps listed in a procedure.
■ ■ ■	Bullets indicate a list of items when the sequence of the items is not important.
	The hand points to information that requires special attention.
	A question mark directs you to a software help system with related information.
	The feet direct you to another document or website with related information.
	A caution calls attention to a condition or possible situation that can damage or destroy the product or your work.
	A warning calls attention to a condition or possible situation that can cause you injury.
	The envelope links to the Email Subscription Management Center page of the Altera website, where you can sign up to receive update notifications for Altera documents.