4. DDR2 and DDR3 SDRAM Board Design Guidelines

EMI_DG_004-5.0

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 4–1 shows the differences between an ideal and real signal seen by the receiver.

Figure 4–1. Ideal and Real Signal at the Receiver

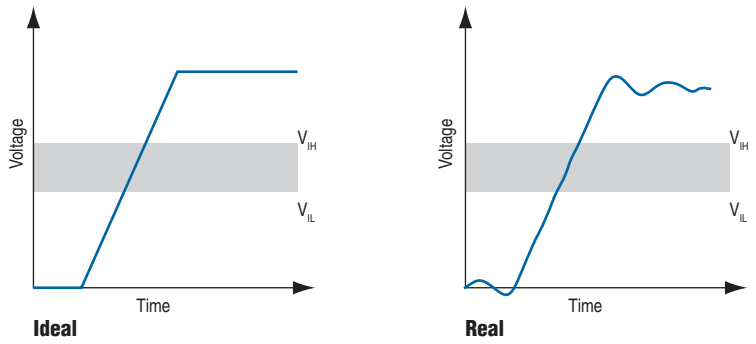

In addition, this chapter compares various types of termination schemes, and their effects on the signal quality on the receiver. It also discusses the proper drive strength setting on the FPGA to optimize the signal integrity at the receiver, and the effects of different loading types, such as components 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, Arria V, and Cyclone® V 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 4–2).

Figure 4–2. DDR3 DIMM Fly-By Topology Requiring Write Leveling

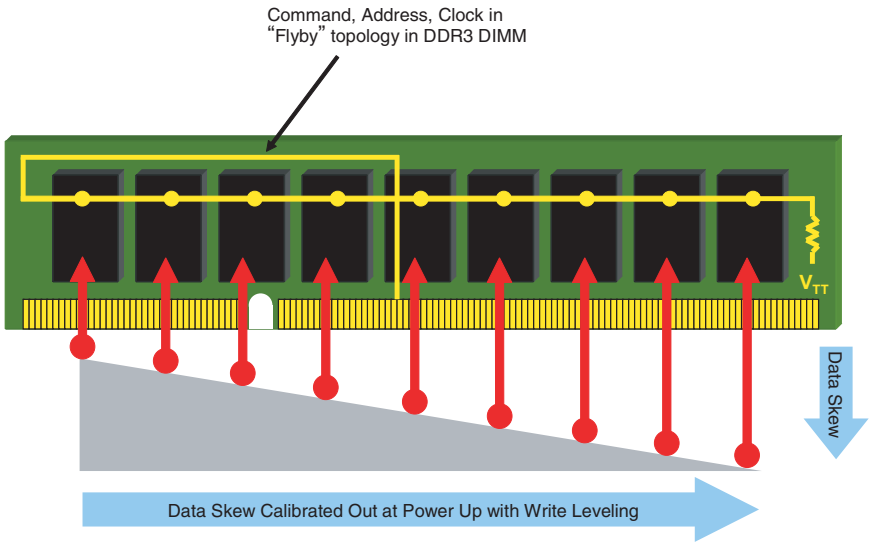

The flight-time skew caused by the fly-by topology led the JEDEC committee to introduce the write leveling feature on the DDR3 SDRAMs; thus requiring controllers to compensate for this skew by adjusting the timing per byte lane.

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_{DQSS} of ± 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.

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 Ω ± 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 4–3 shows the behavior of ODT when you enable dynamic ODT.

Figure 4–3. Dynamic ODT: Behavior with ODT Asserted Before and After the Write *(1)*

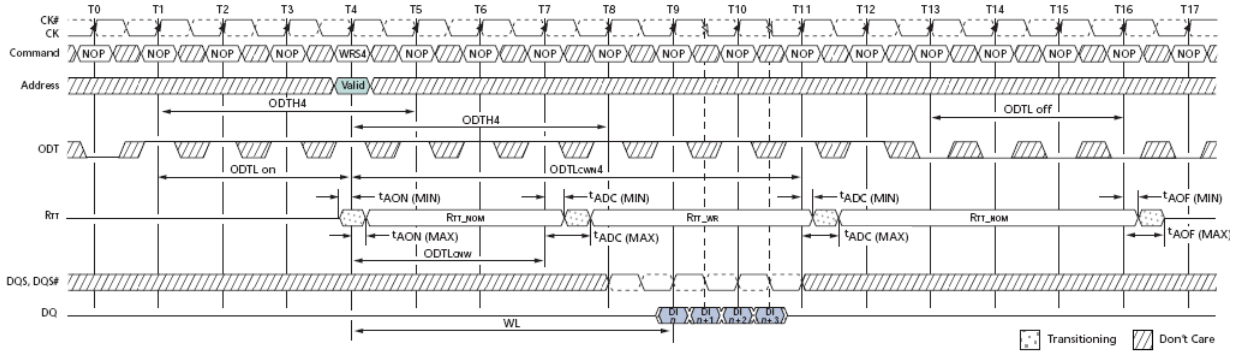

Note to Figure 4–3:

(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 bidirectional 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 is 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 4–4. Dynamic OCT Between Stratix III and Stratix IV FPGA Devices

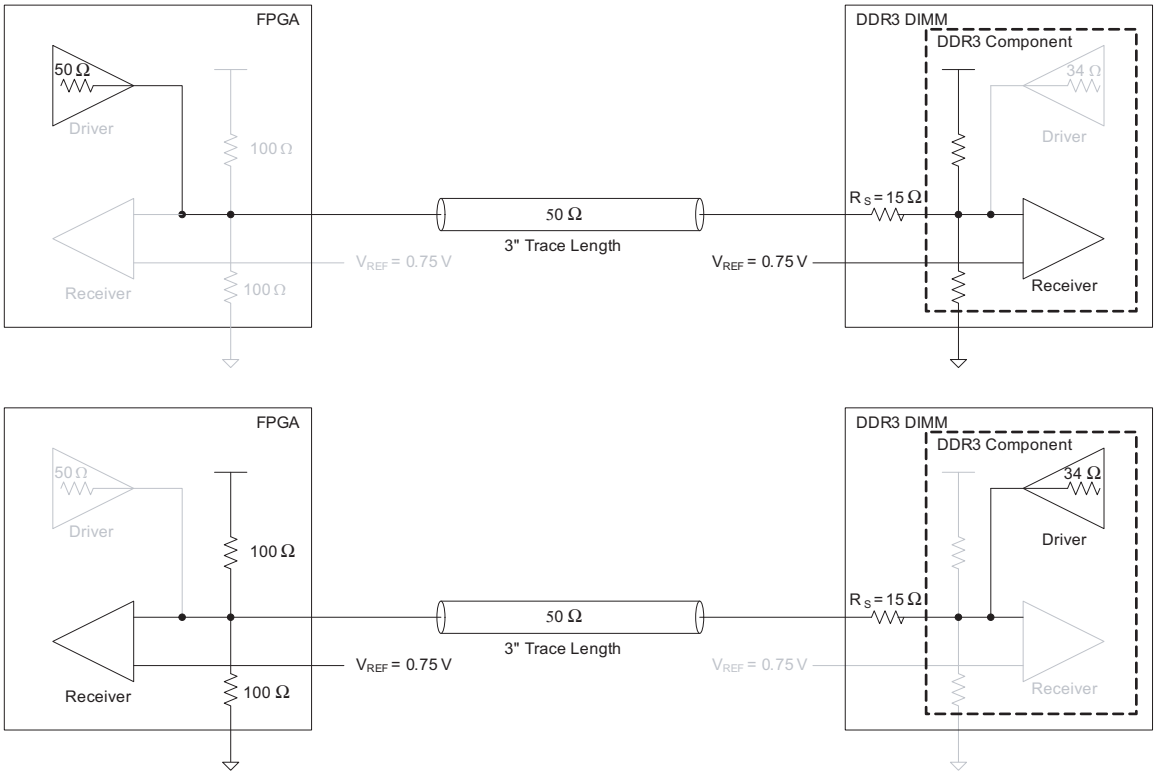

This feature is useful for terminating any high-performance bidirectional path because signal integrity is optimized depending on the direction of the data. In addition, dynamic OCT also eliminates the need for external termination resistors when used with memory devices that support ODT (such as DDR3 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 4–5 and Figure 4–6 illustrate the two most commonly used termination topologies: fly-by topology and non-fly-by topology, respectively.

Figure 4–5. Fly-By Placement of a Parallel Resistor

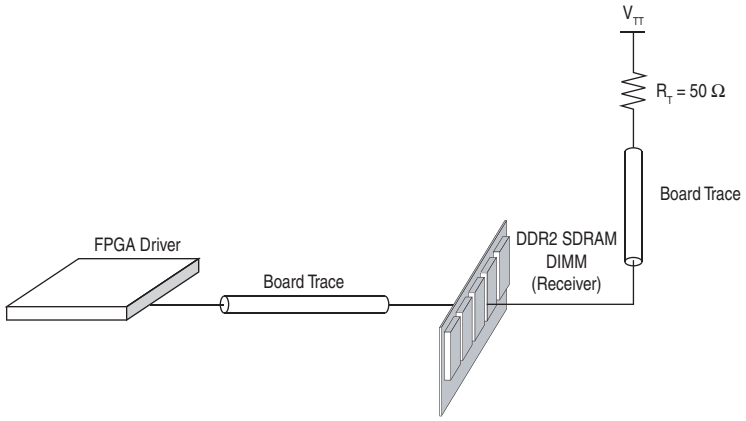

With fly-by topology (Figure 4–5), you place the parallel termination resistor after the receiver. This termination placement resolves the undesirable unterminated stub found in the non-fly-by topology. However, using this topology can be costly and complicate routing. The Stratix II Memory Board 2 uses 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 4–6. Non-Fly-By Placement of a Parallel Resistor

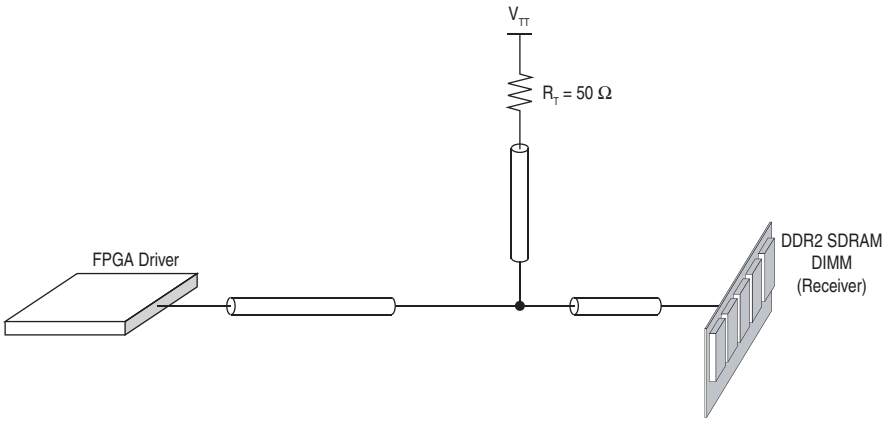

With non-fly-by topology (Figure 4–6), the parallel termination resistor is placed between the driver and receiver (closest to the receiver). This termination placement is easier for board layout, but results in a short stub, which causes an unterminated transmission line between the terminating resistor 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, Arria II GZ, Arria V, Cyclone III, Cyclone IV, Cyclone V, Stratix III, Stratix IV, and Stratix V devices. Table 4–1 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 ± 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 4–1. On-Chip Termination Schemes

Termination Scheme	SSTL-18	FPGA Device						
		Arria II GX	Arria II GZ	Arria V	Cyclone III and Cyclone IV	Cyclone V	Stratix III and Stratix IV	Stratix V (1)
		Column and Row I/O	Column and Row I/O	Column and Row I/O	Column and Row I/O	Column and Row I/O	Column and Row I/O	Column 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	25
On-Chip Series Termination with Calibration	Class I	50	50	50	50	50	50	50
	Class II	25	25	25	25	25	25	25
On-Chip Parallel Termination with Calibration	Class I and Class II	—	50	50	—	50	50	50

Note to Table 4–1:

(1) Row I/O is not available for external memory interfaces in Stratix V devices.

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 4–7).

Figure 4–7. Dynamic OCT for Memory Interfaces

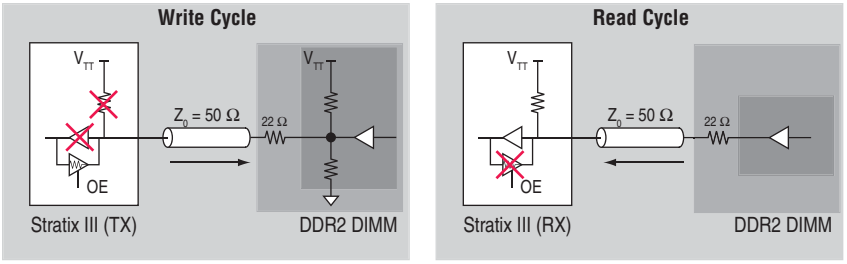

Recommended Termination Schemes

Table 4–2 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 4–2. Termination Recommendations (Part 1 of 3) (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	Class I R50 CAL	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	HALF (8)
	DQS DIFF (13)	DIFF Class R50 CAL	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	HALF (8)
	DQS SE (12)	Class I R50 CAL	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	HALF (8)
	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 CAL	N/A	×1 = 100 Ω differential (10) ×2 = 200 Ω differential (11)	N/A
DDR2 DIMM	DQ	Class I R50 CAL	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	FULL (9)
	DQS DIFF (13)	DIFF Class I R50 CAL	50 Ω Parallel to V_{TT} discrete	ODT75 (7)	FULL (9)
	DQS SE (12)	Class I R50 CAL	50 Ω Parallel to V_{TT} discrete	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 CAL	N/A	N/A = on DIMM	N/A
Arria V and Cyclone V					
DDR2 component	DQ	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	DQS DIFF (13)	DIFF Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	DQS SE (12)	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	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	×1 = 100 Ω differential (10) ×2 = 200 Ω differential (11)	N/A

Table 4–2. Termination Recommendations (Part 2 of 3) (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
DDR2 DIMM	DQ	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	FULL (9)
	DQS DIFF (13)	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)
	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
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	×1 = 100 Ω differential (10) ×2 = 200 Ω 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 I12 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
Arria II GZ, Stratix III, Stratix IV, and Stratix V					
DDR2 component	DQ	Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	DQS DIFF (13)	DIFF Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	DQS SE (12)	DIFF Class I R50/P50 DYN CAL	N/A	ODT75 (7)	HALF (8)
	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	x1 = 100 Ω differential (10) x2 = 200 Ω differential (11)	N/A

Table 4–2. Termination Recommendations (Part 3 of 3) *(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
DDR2 DIMM	DQ	Class I R50/P50 DYN CAL	N/A	ODT75 *(7)*	FULL *(9)*
	DQS DIFF *(13)*	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)*
	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

Notes to Table 4–2:

(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. For example, you can feed two out of nine devices on a single rank DIMM with a single clock pair.
(12) DQS SE is single-ended DQS.
(13) DQS DIFF is differential 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 4–8 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 4–8. Dynamic Series OCT Scheme with ODT on the Memory

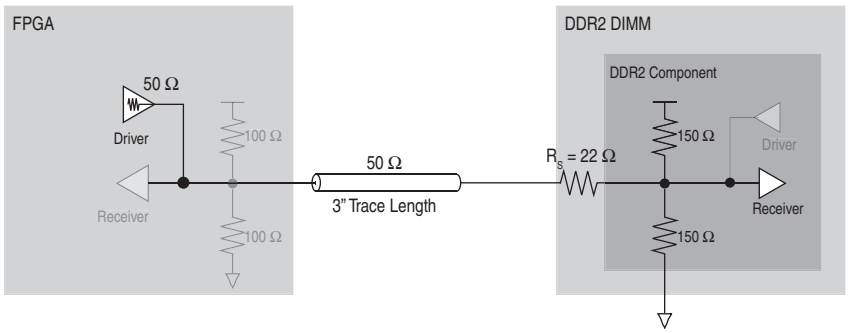

Figure 4–9 and Figure 4–10 show the simulation and measurement results of a write to the DDR2 SDRAM DIMM. The system uses Class I termination with a 50-Ω series OCT measured at the FPGA with a full drive strength and a 75 Ω ODT at the DIMM. Both simulation and bench measurements are in 200 pS/div and 200 mV/div.

Figure 4–9. HyperLynx Simulation FPGA Writing to Memory

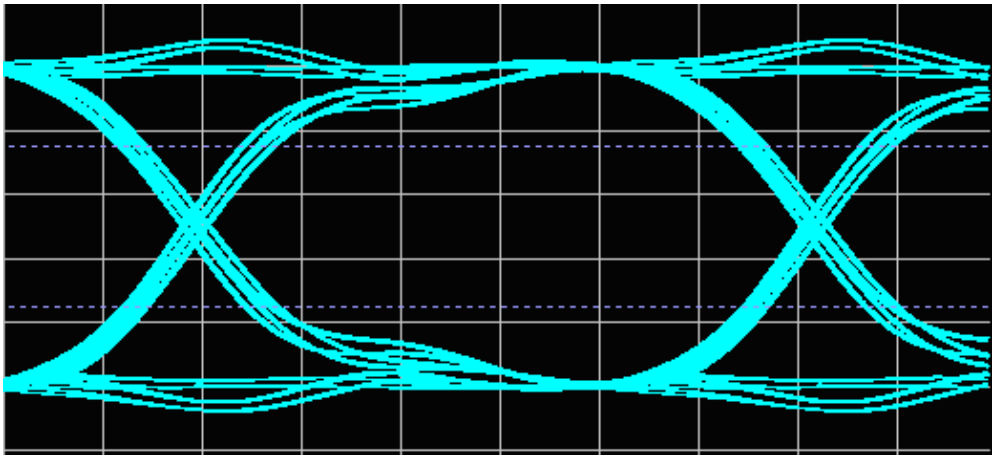

Figure 4–10. Board Measurement, FPGA Writing to Memory

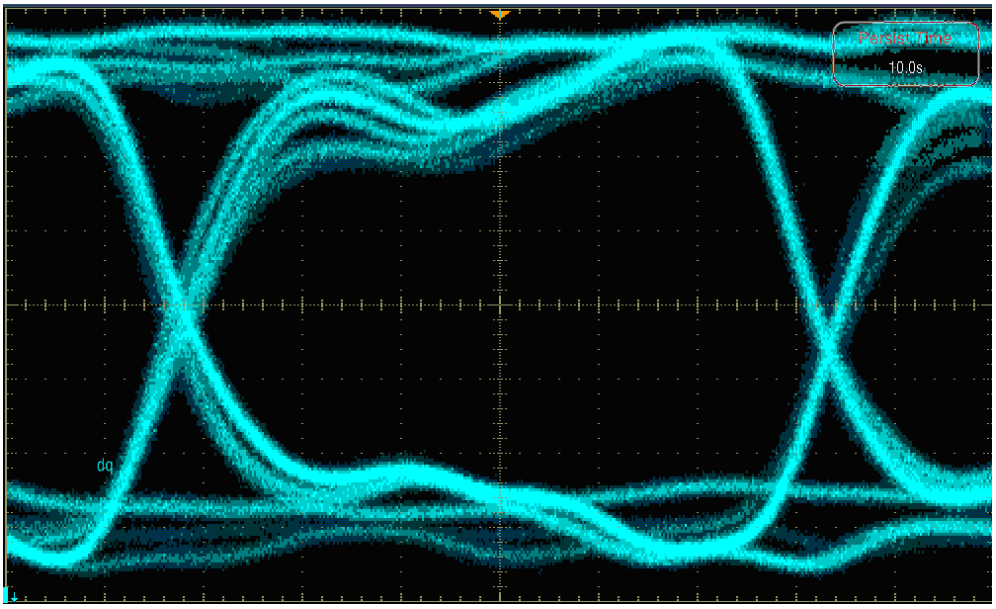

Table 4–3 lists the comparison between the simulation and the board measurement of the signal seen at the DDR2 SDRAM DIMM.

Table 4–3. Signal Comparison When the FPGA is Writing to the Memory *(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 4–3:

(1) N/A is not applicable.

(2) The eye width is measured from V_{IH}/V_{IL}(ac) = VREF ±250 mV to V_{IH}/V_{IL}(dc) = VREF ±125 mV, where V_{IH} and V_{IL} are determined per the JEDEC specification for SSTL-18.

The data in Table 4–3 and Figure 4–9 and Figure 4–10 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 4–3 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 VIL and VIH, 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 4–11 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 4–12 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 4–11. Dynamic Parallel OCT Scheme with Memory-Side Series Resistor

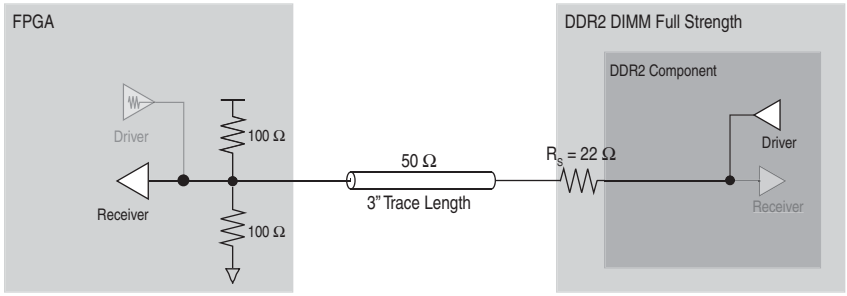

Figure 4–12. Hyperlynx Simulation and Board Measurement, FPGA Reading from Memory

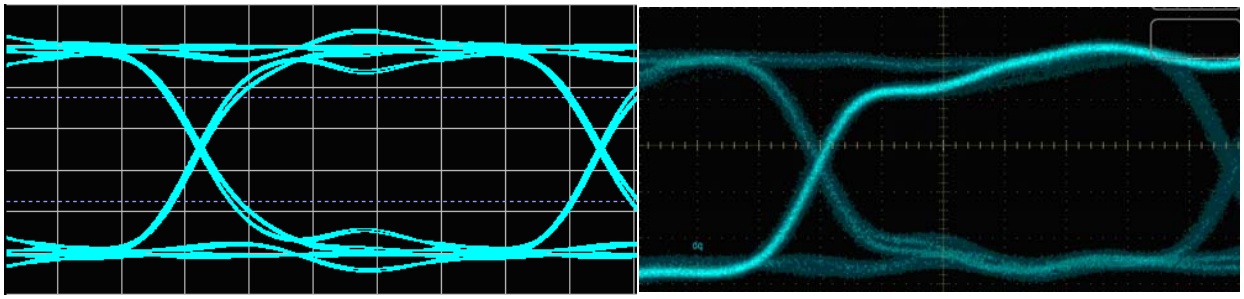

Table 4–4 lists the comparison between the simulation and the board measurement of the signal seen at the FPGA end.

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

	Eye Width (ns) (3)	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 4–4:

(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 ±250 mV to V_{IH}/V_{IL}(dc) = VREF ±125 mV, in which V_{IH} and V_{IL} are determined per the JEDEC specification for SSTL-18.

The data in Table 4–4 and Figure 4–13 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-ΩOCT 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 4–13 shows the Class I termination scheme using ODT when the 50-ΩOCT on the FPGA is turned on.

Figure 4–13. Class I Termination Using ODT with 50-ΩOCT

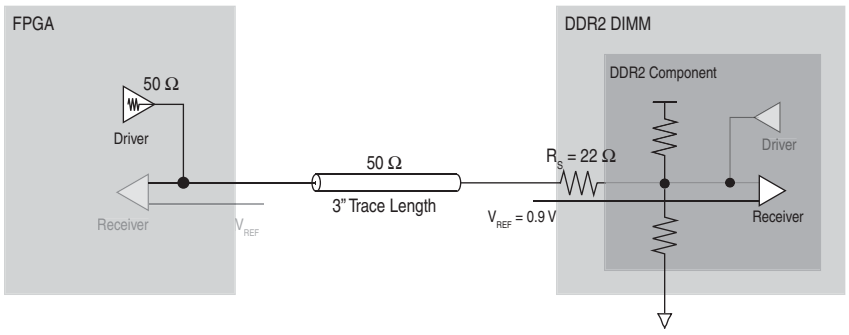

The resulting signal quality has a similar eye opening to the 8 mA drive strength setting (refer to "Drive Strength" on page 4–50) without any over- or undershoot. Figure 4–14 shows the simulation and measurement of the signal at the memory side (DDR2 SDRAM DIMM) with the drive strength setting of 50-Ω OCT in the FPGA.

Figure 4–14. HyperLynx Simulation and Measurement, FPGA Writing to Memory

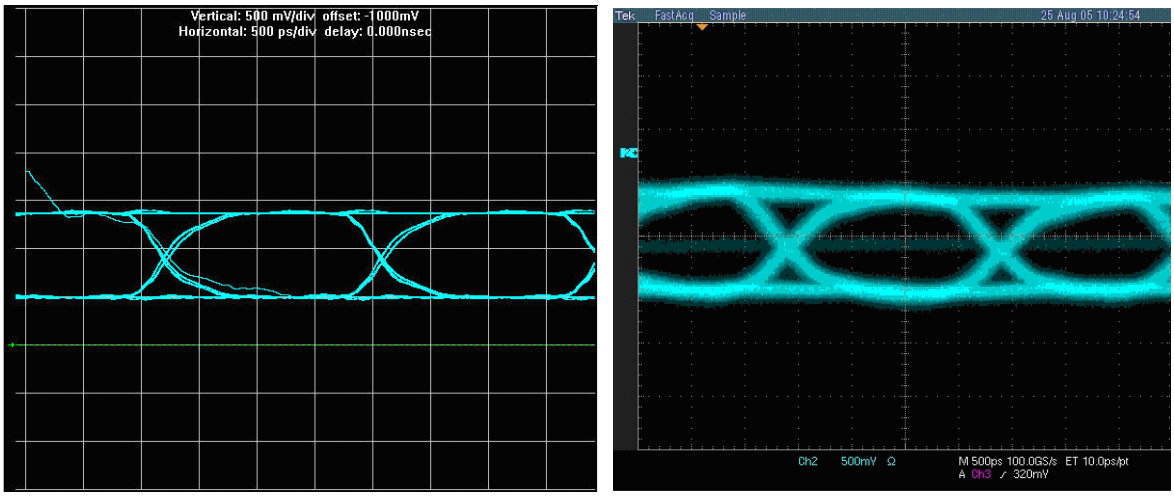

Table 4–5 lists data for the signal at the DDR2 SDRAM DIMM of a Class I scheme termination using ODT with a memory-side series resistor. The FPGA is writing to the memory with 50-Ω OCT.

Table 4–5. Simulation and Board Measurement Results for 50-Ω OCT and 8-mA Drive Strength Settings *(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 4–5:

(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-ΩOCT setting that you can use to improve the signal quality in a Class II terminated transmission line. Figure 4–15 shows the Class II termination scheme using ODT when the 25-ΩOCT on the FPGA is turned on.

Figure 4–15. Class II Termination Using ODT with 25-ΩOCT

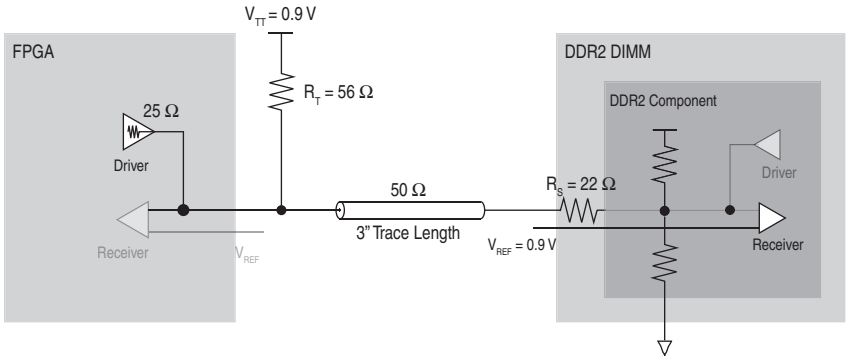

Figure 4–16 shows the simulation and measurement of the signal at the DDR2 SDRAM DIMM (receiver) with a drive strength setting of 25-ΩOCT in the FPGA.

Figure 4–16. HyperLynx Simulation and Measurement, FPGA Writing to Memory

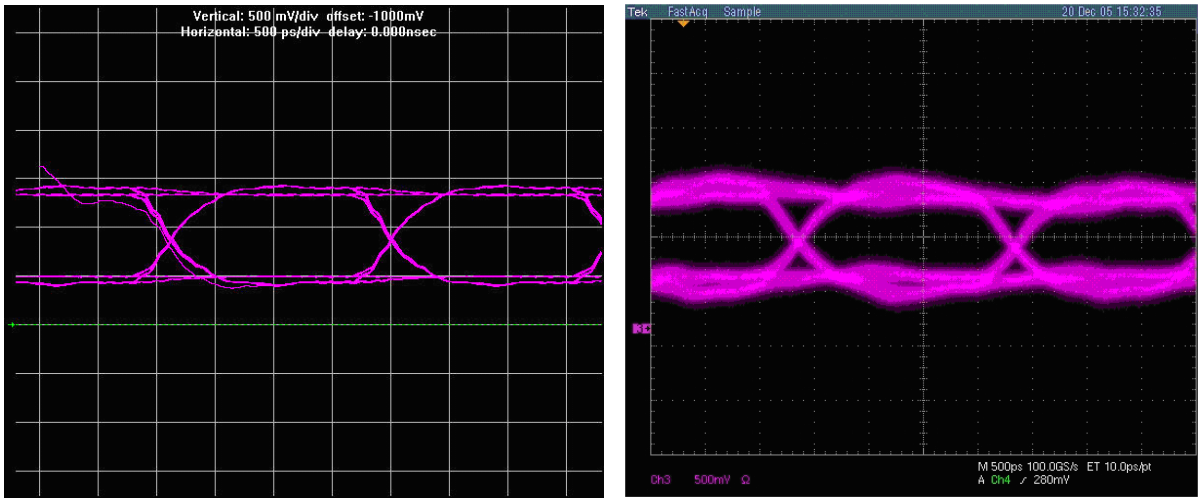

Table 4–6 lists the data for the signal at the DDR2 SDRAM DIMM of a Class II termination with a memory-side series resistor. The FPGA is writing to the memory with 25-Ω OCT.

Table 4–6. Simulation and Board Measurement Results for 25-Ω OCT and 16-mA Drive Strength Settings *(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
Board Measurement	1.47	0.51	N/A	N/A

Note to Table 4–6:

(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 4–17 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 4–17. Class-II Termination Scheme with Memory-Side Series Resistor

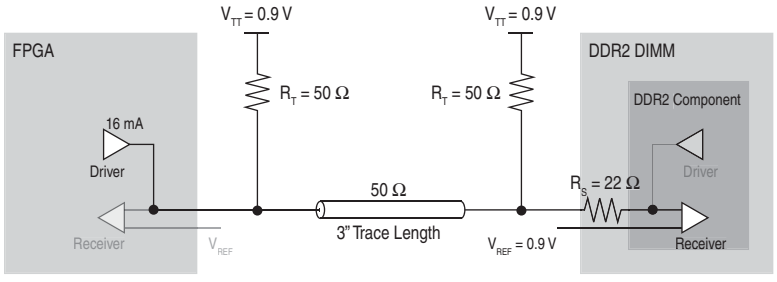

Figure 4–18 and Figure 4–19 show the simulation and measurement result of a write to the DDR2 SDRAM DIMM. The system uses Class II termination with a source-series resistor measured at the DIMM with a drive strength setting of 16 mA.

Figure 4–18. HyperLynx Simulation, FPGA Writing to Memory

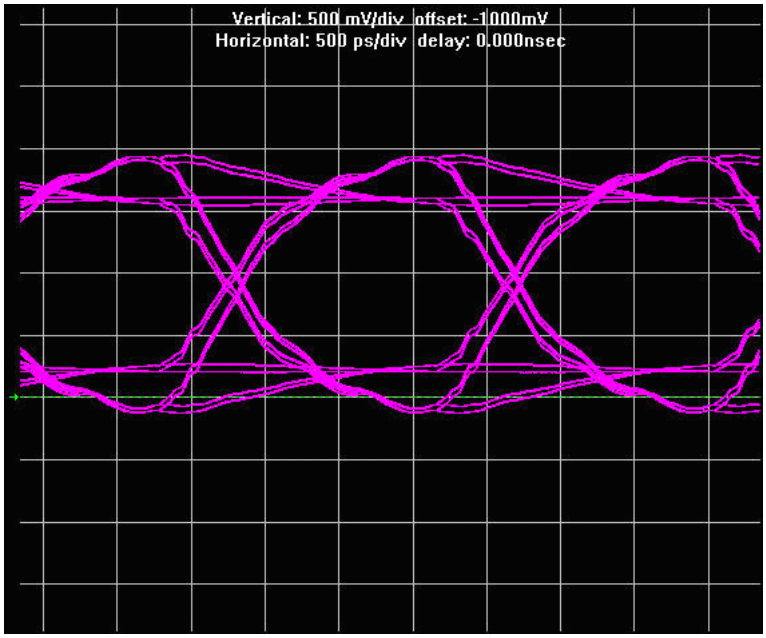

The simulation shows a clean signal with a good eye opening, but there is slight over- and undershoot of the 1.8-V signal specified by DDR2 SDRAM. The over- and undershoot can be attributed to either overdriving the transmission line using a higher than required drive strength setting on the driver or 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 4–19 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 4–19. Board Measurement, FPGA Writing to Memory

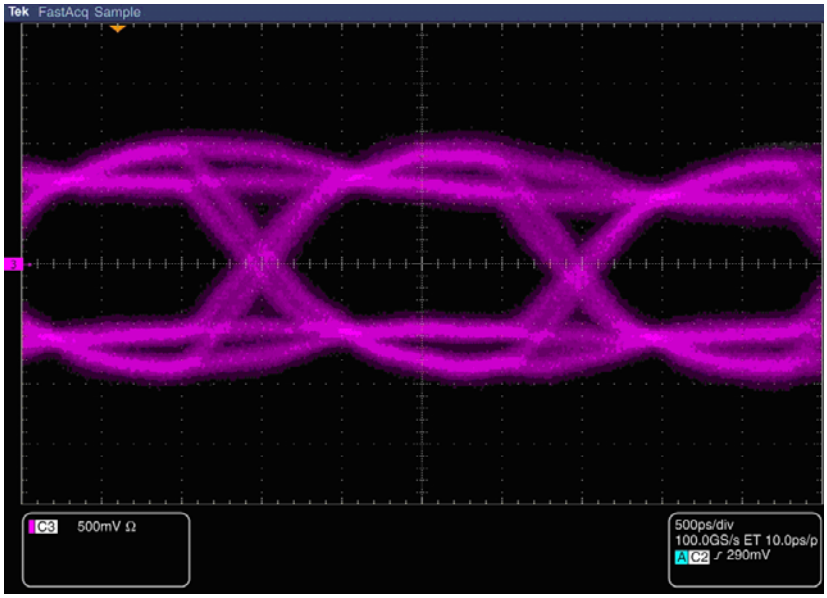

Table 4–7 lists the comparison between the simulation and the board measurement of the signal seen at the DDR2 SDRAM DIMM.

Table 4–7. Signal Comparison When the FPGA is Writing to the Memory (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 4–7:

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

(2) The eye width is measured from V_{REF} ± 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 4–20 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 4–20. Class II Termination Scheme with Memory-Side Series Resistor

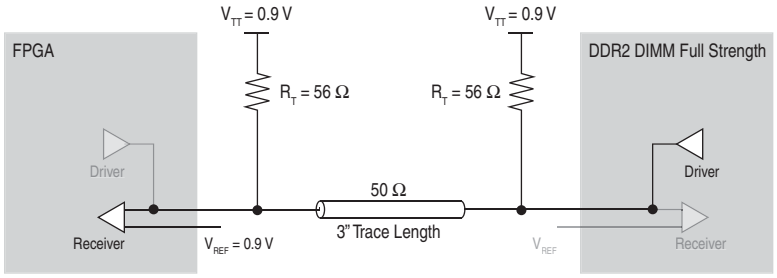

Figure 4–21 and Figure 4–22 show the simulation and measurement, respectively, of the signal at the FPGA side with the full drive strength setting on the DDR2 SDRAM DIMM. The simulation uses a Class II termination scheme with a source-series resistor transmission line. The FPGA is reading from the memory with a full drive strength setting on the DIMM.

Figure 4–21. HyperLynx Simulation, FPGA Reading from Memory

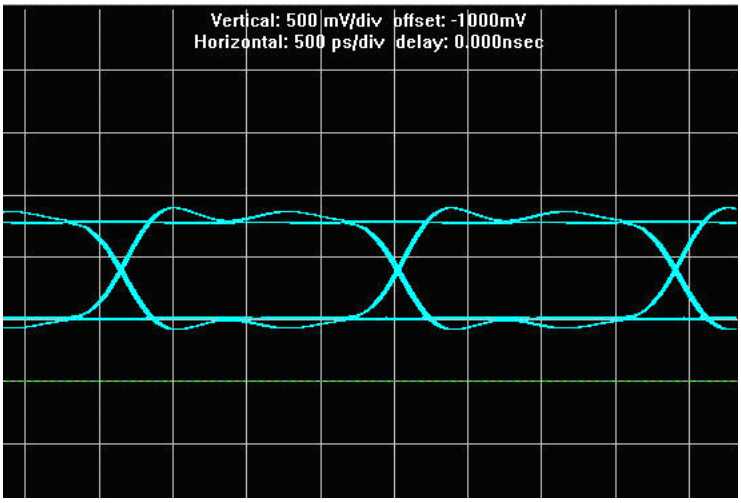

Figure 4–22. Board Measurement, FPGA Reading from Memory

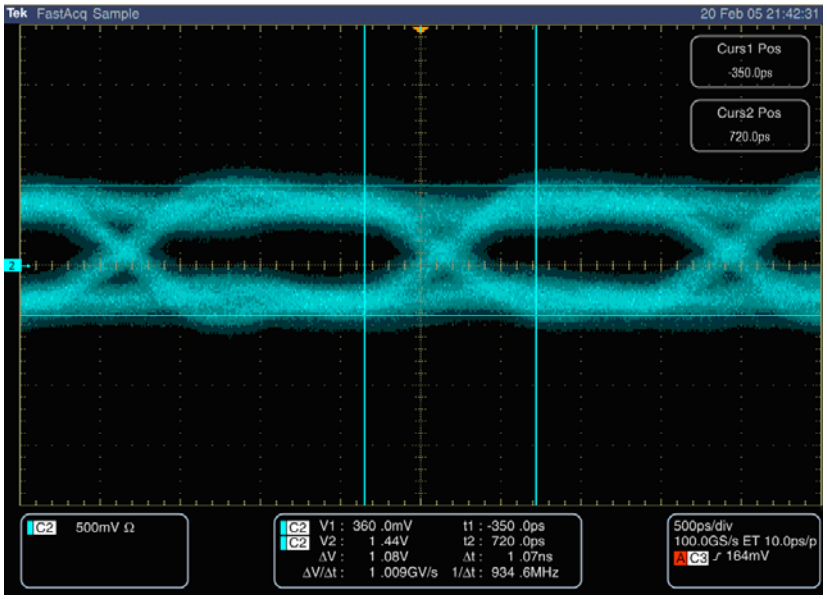

Table 4–8 lists the comparison between the simulation and board measurements of the signal seen by the FPGA when the FPGA is reading from memory (driver).

Table 4–8. Signal Comparison, FPGA is Reading from Memory (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

Notes to Table 4–8:

(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 4–52). 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 4–23), 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 4–28 for more information about how an ODT resistor compares to an external termination resistor.

Figure 4–23. Class I Termination Scheme with Memory-Side Series Resistor

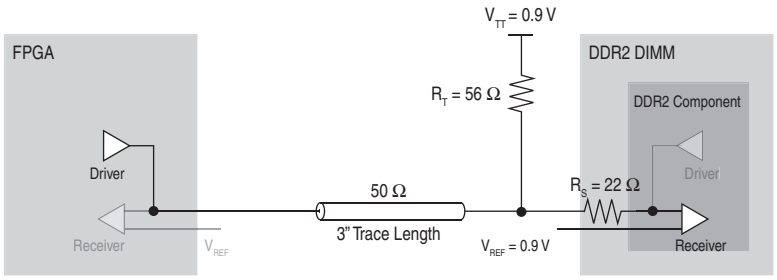

Figure 4–24 shows the simulation and measurement of the signal at the memory (DDR2 SDRAM DIMM) of Class I termination with a memory-side resistor. The FPGA writes to the memory with a 16 mA drive strength setting.

Figure 4–24. HyperLynx Simulation and Board Measurement, FPGA Writing to Memory

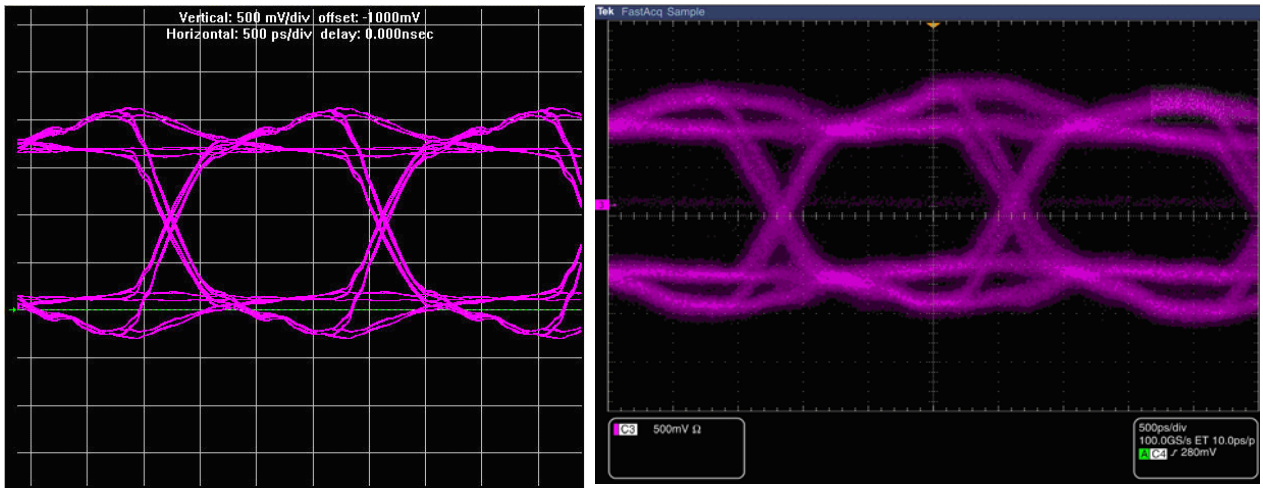

Table 4–9 lists the comparison of the signal at the DDR2 SDRAM DIMM of a Class I and Class II termination scheme using external resistors with memory-side series resistors. The FPGA (driver) writes to the memory (receiver).

Table 4–9. Signal Comparison When the FPGA is Writing to Memory *(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 4–9:

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

Table 4–9 lists 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.

For more information about how drive strength affects the signal quality, refer to "Drive Strength" on page 4–50.

FPGA Reading from Memory

As described in the section "FPGA Writing to Memory" on page 4–25, 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 4–25 shows a Class I termination scheme with a memory-side resistor. The FPGA reads from the memory.

Figure 4–25. Class I Termination Scheme with Memory-Side Series Resistor

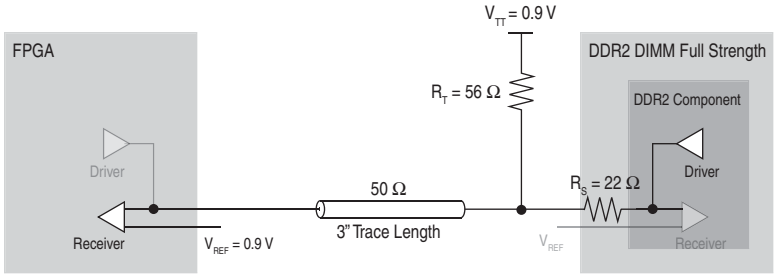

When the FPGA reads from the memory (Figure 4–25), the transmission line is not terminated at the FPGA, resulting in an impedance mismatch, which then results in over- and undershoot. Figure 4–26 shows the simulation and measurement of the signal at the FPGA side (receiver) of a Class I termination. The FPGA reads from the memory with a full drive strength setting on the DDR2 SDRAM DIMM.

Figure 4–26. HyperLynx Simulation and Board Measurement, FPGA Reading from Memory

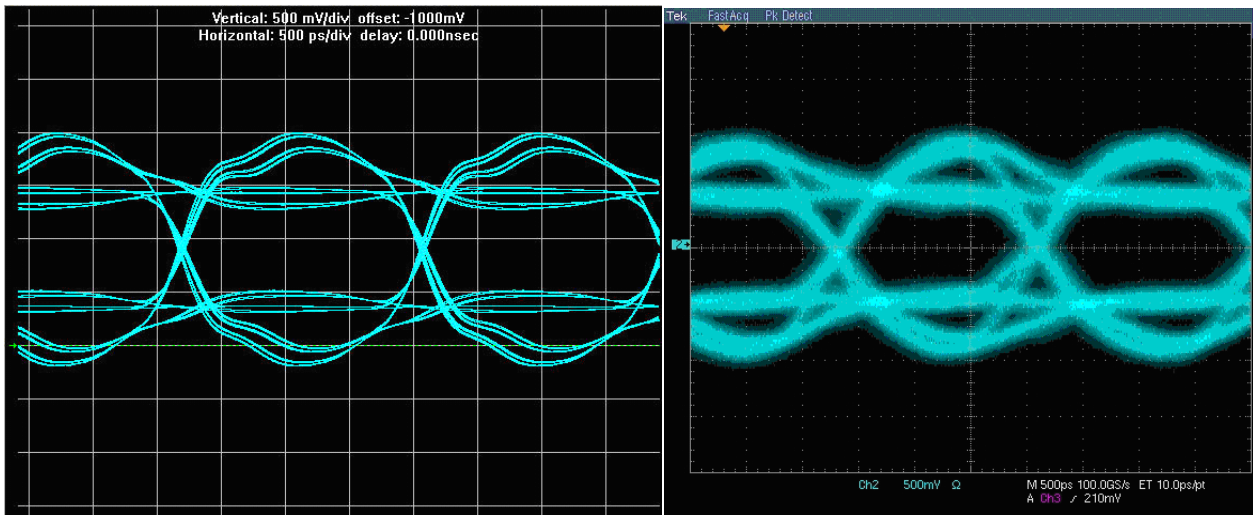

Table 4–10 lists the comparison of the signal "seen" at the FPGA of a Class I and Class II termination scheme using an external resistor with a memory-side series resistor. The FPGA (receiver) reads from the memory (driver).

Table 4–10. Signal Comparison When the FPGA is Reading From Memory (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

Notes to Table 4–10:

(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 4–10 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.

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

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 4–25.

Figure 4–27 shows the termination scheme when the ODT on the DDR2 SDRAM is turned on.

Figure 4–27. Class I Termination Scheme Using ODT

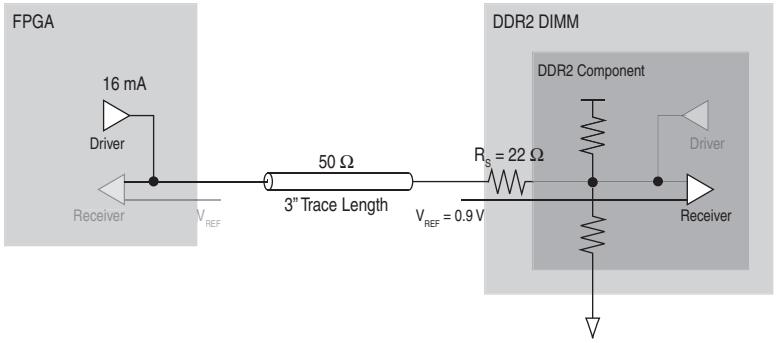

Figure 4–28 shows the simulation and measurement of the signal visible at the memory (receiver) using 50 ΩODT with a memory-side series resistor transmission line. The FPGA writes to the memory with a 16 mA drive strength setting.

Figure 4–28. Simulation and Board Measurement, FPGA Writing to Memory

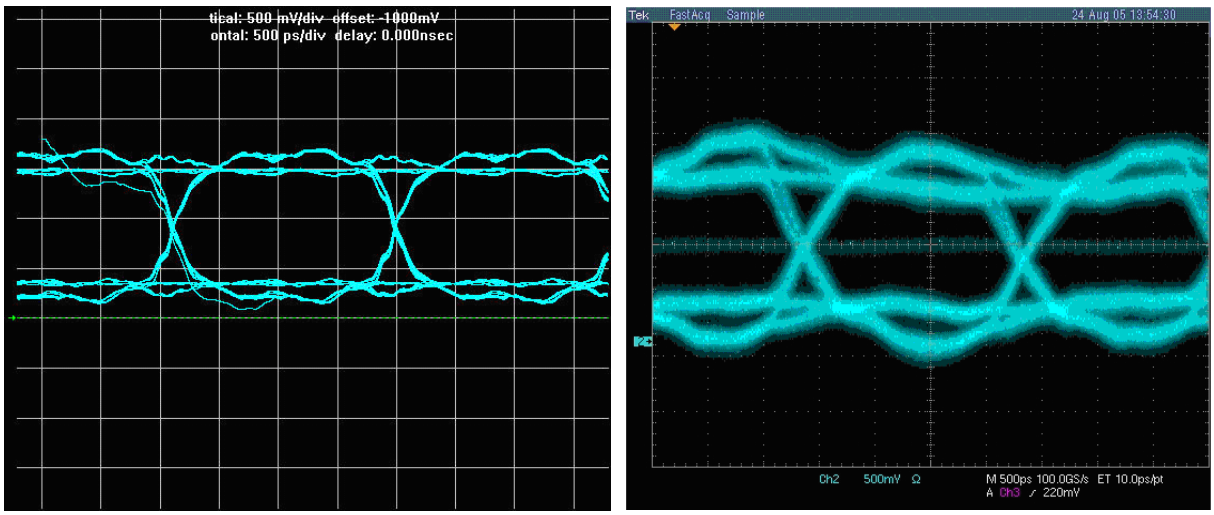

Table 4–11 lists the comparisons of the signal seen the DDR2 SDRAM DIMM of a Class I termination scheme using an external resistor and a Class I termination scheme using ODT with a memory-side series resistor. The FPGA (driver) writes to the memory (receiver).

Table 4–11. Signal Comparison When the FPGA is Writing to Memory (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

Notes to Table 4–11:

(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 parallel 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 4–31 on page 4–33.

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 4–29 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 4–29. No-Parallel Termination Scheme

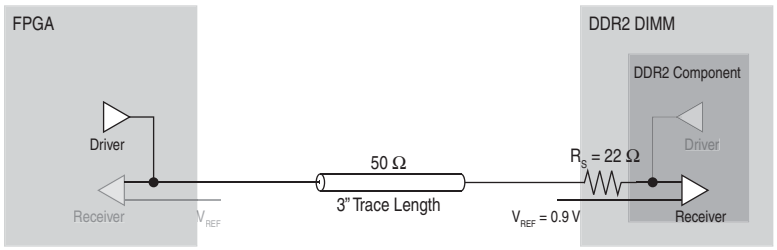

Figure 4–30 shows a HyperLynx simulation and measurement of the FPGA writing to the memory at 533 MHz with a no-parallel termination scheme using a 16 mA drive strength option. The measurement point is on the DDR2 SDRAM DIMM.

Figure 4–30. HyperLynx Simulation and Board Measurement, FPGA Writing to Memory

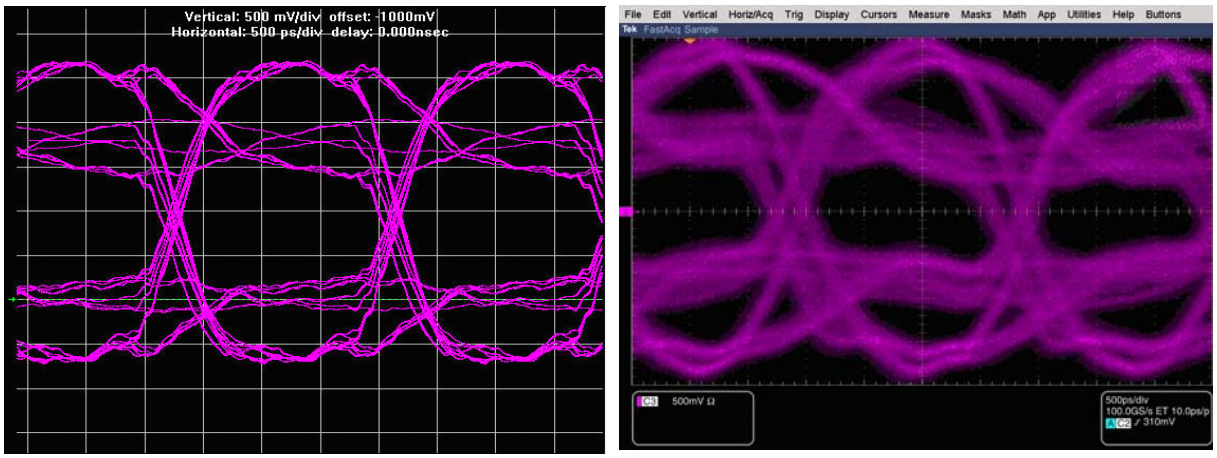

The simulated and measured signal shows that there is sufficient eye opening but also significant over- and undershoot of the 1.8-V signal specified by the DDR2 SDRAM. From the simulation and measurement, the overshoot is approximately 1 V higher than 1.8 V, and undershoot is approximately 0.8 V below 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 4–12 lists 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 4–12. Signal Comparison When the FPGA is Writing to Memory *(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 4–12:

(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 4–50.

FPGA Reading from Memory

In a no-parallel termination scheme (Figure 4–31), 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-Ω 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 4–31. No-Parallel Termination Scheme, FPGA Reading from Memory

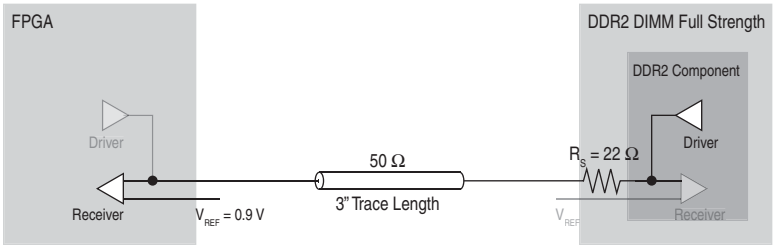

Figure 4–32 shows the simulation and measurement of the signal visible at the FPGA (receiver) when the memory is driving the no-parallel termination transmission line with a memory-side series resistor.

Figure 4–32. HyperLynx Simulation and Board Measurement, FPGA Reading from Memory

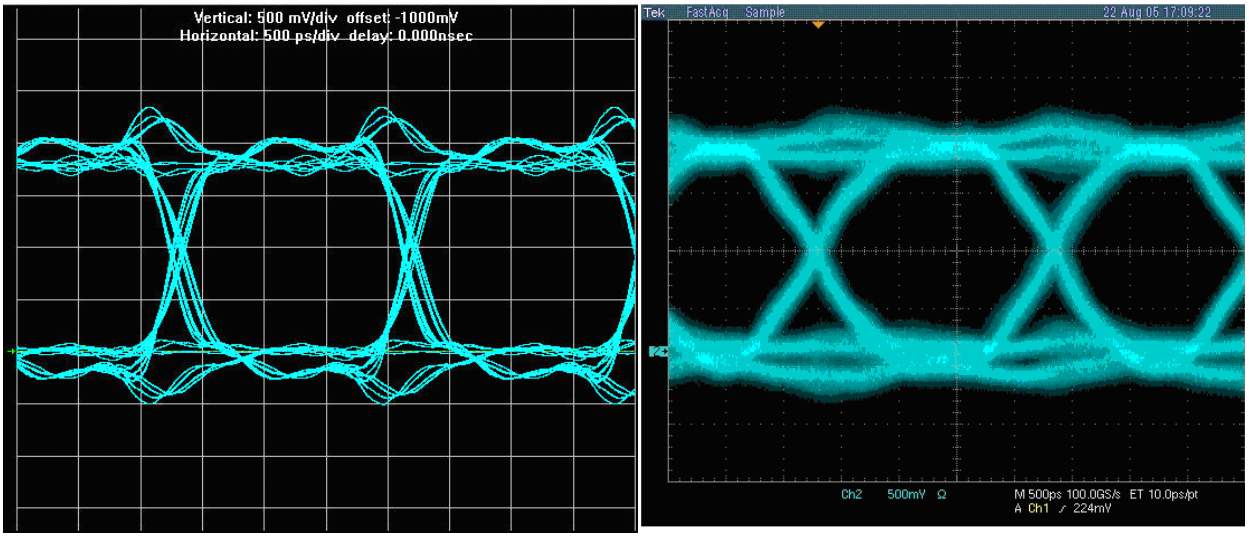

Table 4–13 lists the comparison of the signal seen on the FPGA with a no-parallel and a Class II termination scheme when the FPGA is reading from memory.

Table 4–13. Signal Comparison, FPGA Reading From Memory (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

Table 4–13. Signal Comparison, FPGA Reading From Memory (1), (2)

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

Notes to Table 4–13:

(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 4–31, 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 4–31. 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 4–33), 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 4–33. No-Parallel Termination Scheme, FPGA REading from Memory

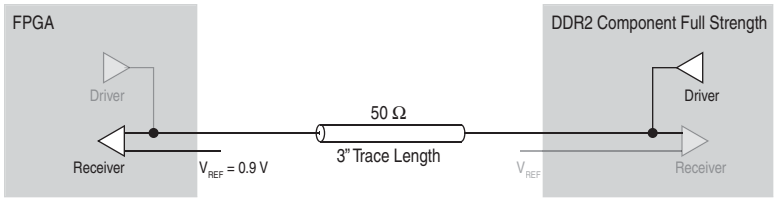

Figure 4–34 shows the simulation and measurement of the signal at the FPGA side in a no-parallel termination scheme with the full drive strength setting on the memory.

Figure 4–34. HyperLynx Simulation and Measurement, FPGA Reading from Memory

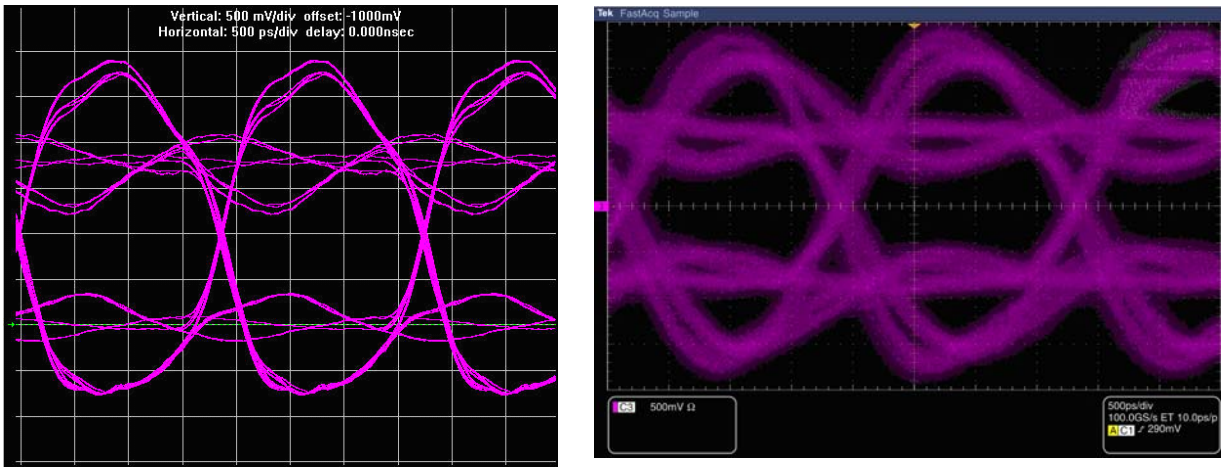

Table 4–14 lists the difference between no-parallel termination with and without memory-side series resistor when the memory (driver) writes to the FPGA (receiver).

Table 4–14. No-Parallel Termination with and without Memory-Side Series Resistor *(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 4–14:

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

Table 4–14 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 4–31. 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 4–7.

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 4–15 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 4–15. 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)
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 uF just before V_{TT} For more information, refer to Figure 4–38 on page 4–39.	

Notes to Table 4–16:

(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 4–35 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 4–35. DQ and DQS Net Structure for 64-Bit DDR3 SDRAM UDIMM *(1)*

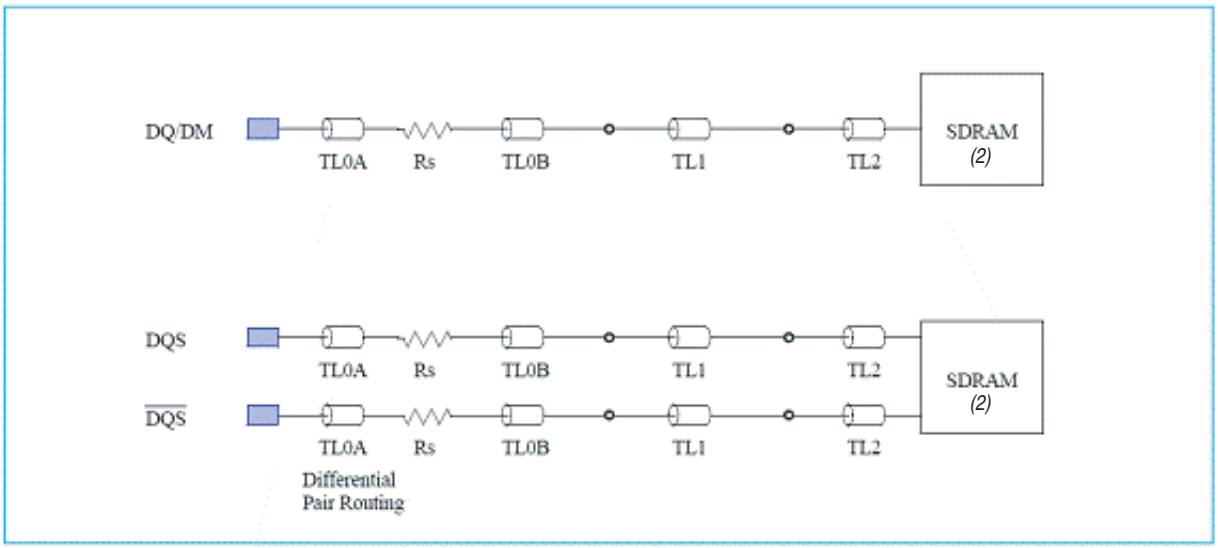

Note to Figure 4–35:

(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 4–5, 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 4–36 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 4–36. Simulated Write-Eye Diagram of a DDR3 SDRAM DIMM Using a 60-Ω ODT Setting

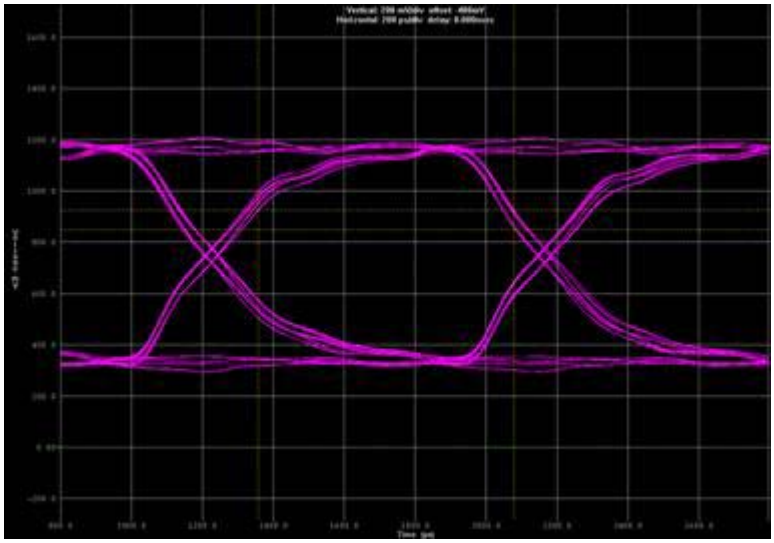

Figure 4–37 shows the measured write eye diagram using Altera's Stratix III or Stratix IV memory board.

Figure 4–37. Measured Write-Eye Diagram of a DDR3 SDRAM DIMM Using the 60-Ω ODT Setting

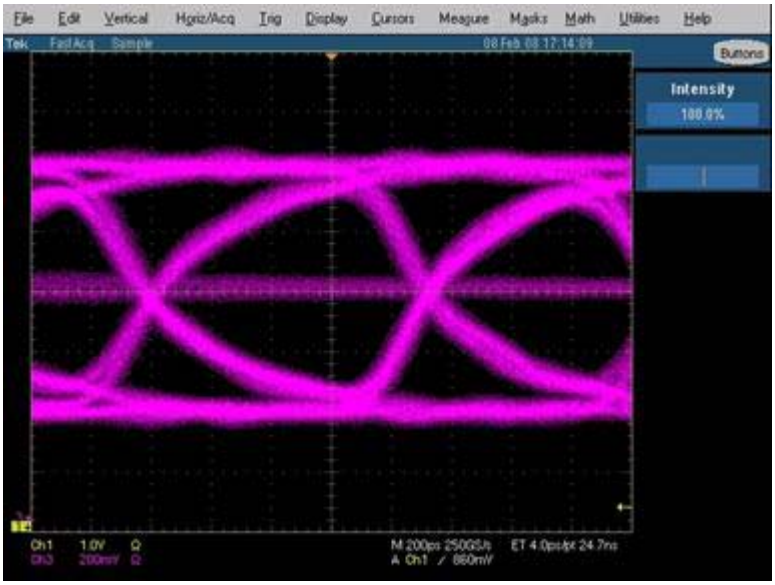

The measured eye diagram correlates well with the simulation. The faint line in the middle of the eye diagram is the effect of the refresh operation during a regular operation. Because these simulations and measurements are based on a narrow set of constraints, you must perform your own board-level simulation 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 4–38 shows the net structure for the memory clocks and the location of the termination resistors, R_{TT}. The value of R_{TT} is 36 Ω, which results in an equivalent differential termination value of 72 Ω. 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. The recommended center-tap-terminated (C_{TT}) value is 0.1 uF just before V_{TT}.

Figure 4–38. Clock Net Structure for a 64-Bit DDR3 SDRAM UDIMM *(1)*

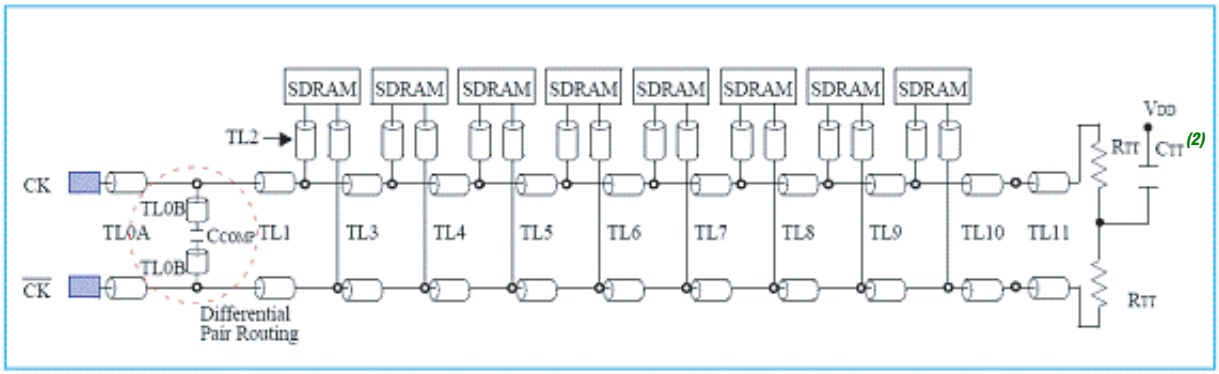

Note to Figure 4–38:

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

(2) The recommeded C_{TT} value is 0.1 uF just before V_{TT}..

From Figure 4–38, you can see that the DDR3 SDRAM clocks are routed in a fly-by topology, as mentioned in "Read and Write Leveling" on page 4–3, resulting in the need for write-and-read leveling. Figure 4–39 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 4–39. Differential Memory Clock of a DDR3 SDRAM DIMM at the First and Last Component on the DIMM

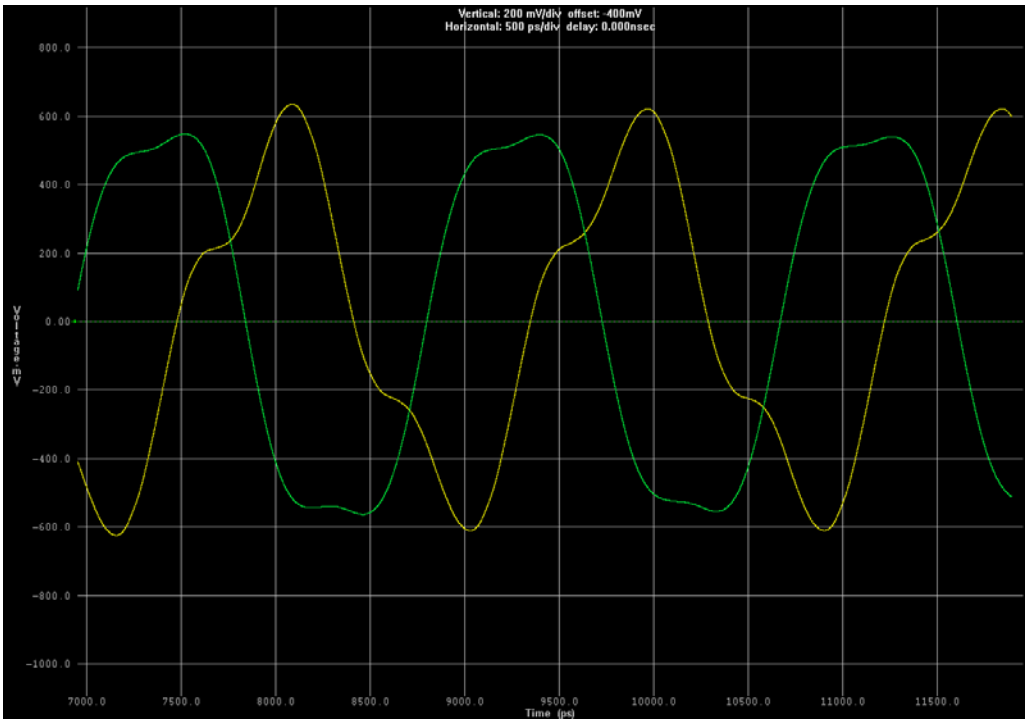

Figure 4–39 shows that the memory clock seen at the first DDR3 SDRAM component (the yellow signal) leads the memory clock seen at the last DDR3 SDRAM component (the green signal) by 1.3 ns, which is about 0.69 t_{CK} for a 533 MHz operation.

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 4–40 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 Ω

Figure 4–40. Command and Address Net Structure for a 64-Bit DDR3 SDRAM Unbuffered DIMM *(1)*

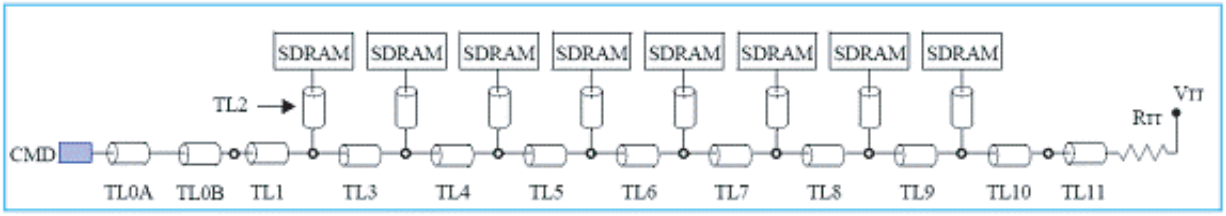

Note to Figure 4–40:

(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 4–40, 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 4–3, resulting in the need for write-and-read leveling.

Figure 4–41 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 4–41. Command and Address Eye Diagram of a DDR3 SDRAM DIMM at the First and Last DDR3 SDRAM Component at 533 MHz *(1)*

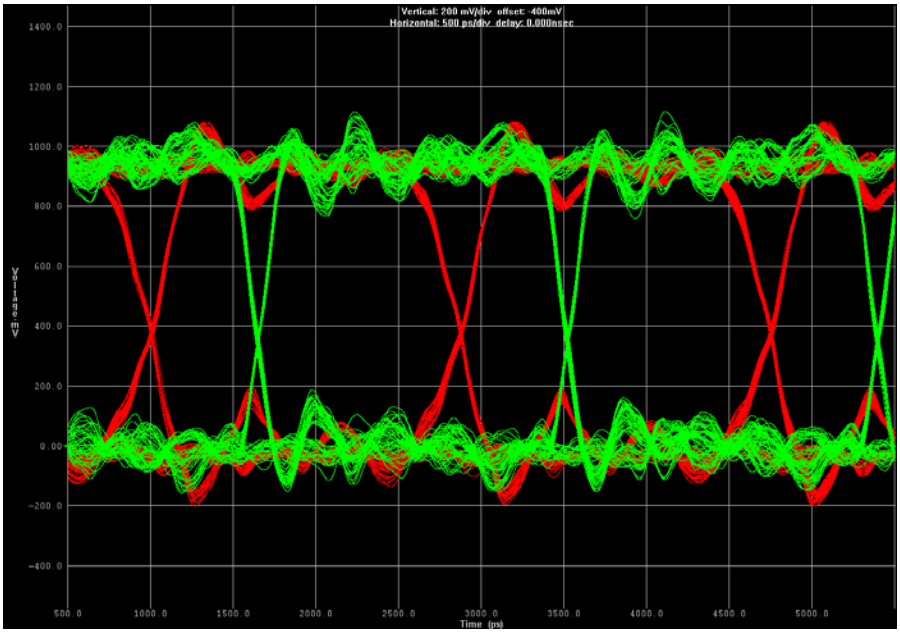

Note to Figure 4–41:

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

Figure 4–41 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 4–5, 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 4–37 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.

For receiver termination recommendations and transmitter output drive strength settings, refer to "DQS, DQ, and DM for DDR3 SDRAM UDIMM" on page 4–37.

Figure 4–42 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 4–42. DDR3 SDRAM Component Driving the Stratix III, Stratix IV, and Stratix V FPGA Device with Parallel 50-Ω OCT Turned On

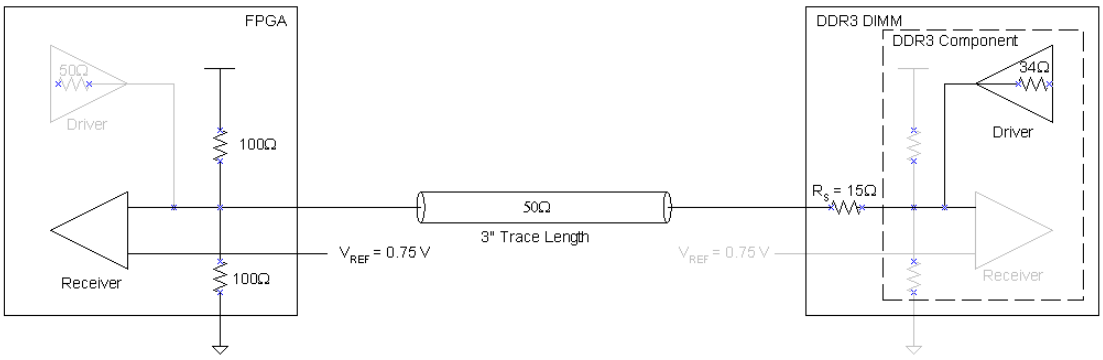

Figure 4–43 shows the simulation of a read from the DDR3 SDRAM DIMM with a 50-Ω parallel OCT setting on the Stratix III or Stratix IV FPGA device.

Figure 4–43. Read-Eye Diagram of a DDR3 SDRAM DIMM at the Stratix III and Stratix IV FPGA Using a Parallel 50-Ω OCT Setting

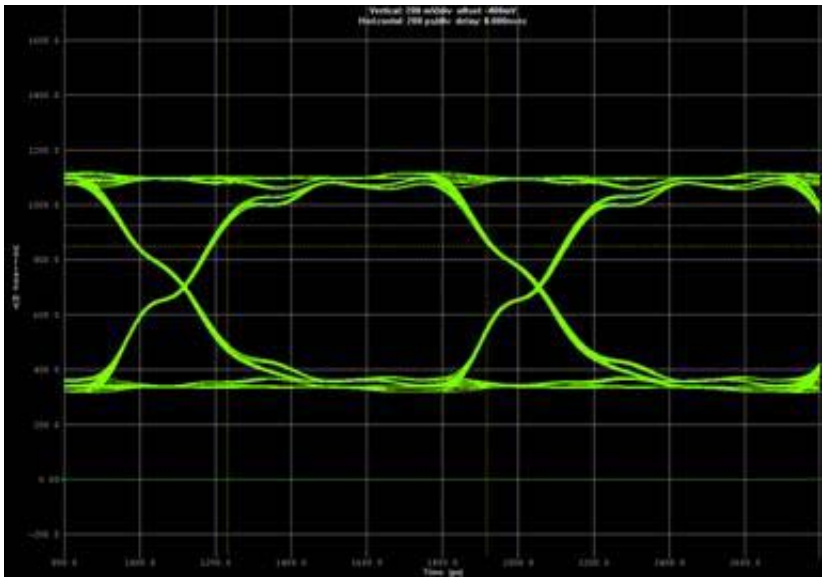

Use of the Stratix III, Stratix IV, or Stratix V parallel 50-Ω OCT feature matches receiver impedance with the transmission line characteristic impedance. This eliminates any reflection that causes ringing, and results in a clean eye diagram at the Stratix III, Stratix IV, or Stratix V FPGA.

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 4–39 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 4–41 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 4–16 lists 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 4–16. DDR3 SDRAM ODT Matrix for Writes (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 4–16:

(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 4–17 lists 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 4–17. DDR3 SDRAM ODT Matrix for Reads (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 4–17:

(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 4–37.

DDR3 SDRAM Load-Reduced DIMM

RDIMM and LRDIMM differ in that DQ, DQS, and DM signals are registered or buffered in the LRDIMM. The LRDIMM buffer IC is a superset of the RDIMM buffer IC. The buffer IC isolates the memory interface signals from loading effects of the memory chip. Reduced electrical loading allows a system to operate at higher frequency and higher density.

If you want to use your DIMM socket for UDIMM and RDIMM/LRDIMM, you must create the necessary redundant connections on the board from the FPGA to the DIMM socket. For example, the number of chip select signals required for a single-rank UDIMM is one, but for single-rank RDIMM the number of chip selects required is two. RDIMM and LRDIMM have parity signals associated with the address and command bus which UDIMM does not have. Consult the DIMM manufacturer's data sheet for detailed information about the necessary pin connections for various DIMM topologies.

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 options needs read and write leveling, so you must use the UniPHY IP with leveling.

  For more information about 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 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 4–44 when the Stratix III, Stratix IV, or Stratix V FPGA is writing to the DDR3 SDRAM.

Figure 4–44. Stratix III, Stratix IV, and Stratix V FPGA Writing to a DDR3 SDRAM Components

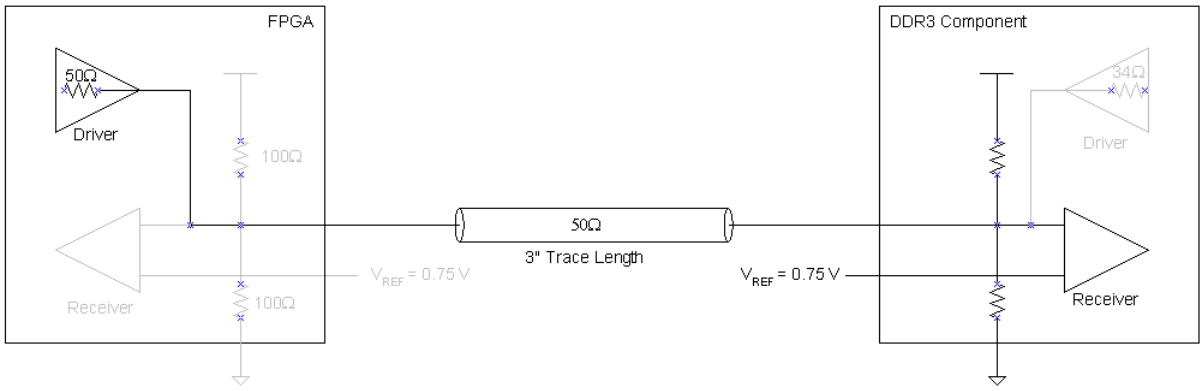

When you are using DDR3 SDRAM components, there are no DIMM connectors. This minimizes any impedance discontinuity, resulting in better signal integrity.

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 4–45 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-ΩOCT setting on the output driver of the Stratix III, Stratix IV, or Stratix V FPGA.

Figure 4–45. 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

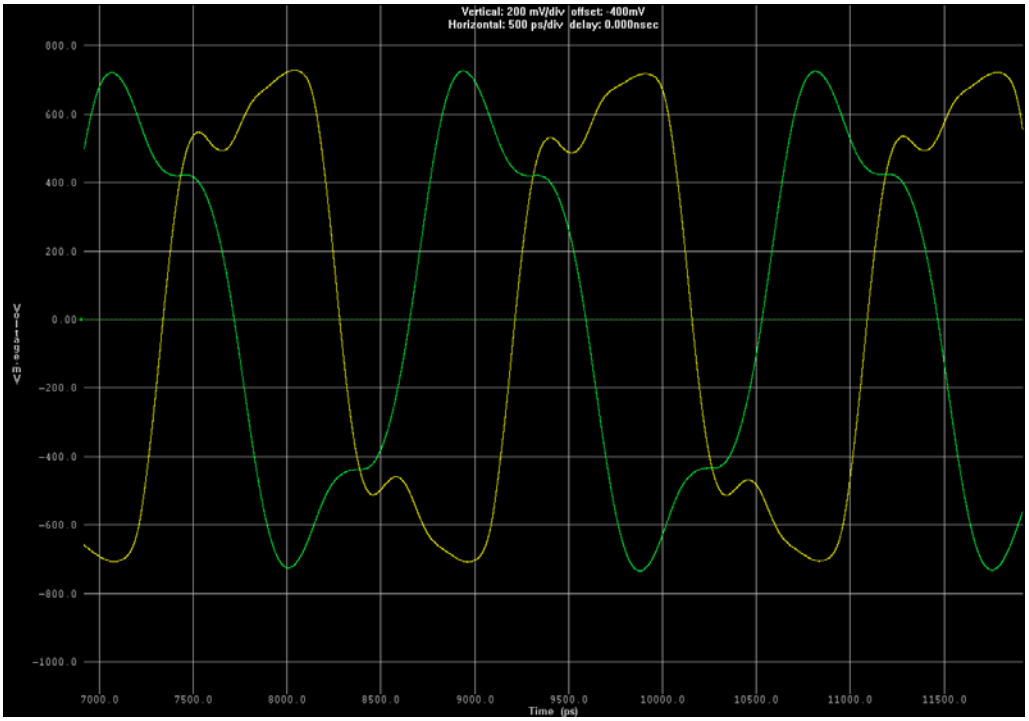

Without the compensation capacitor, the memory clocks (the yellow signal) at the first component have significant ringing, whereas, with the compensation capacitor the ringing is dampened. Similarly, the differential termination resistor needs to be included in the design. Depending on your board stackup and layout requirements, you choose your differential termination resistor value. Figure 4–46 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, and terminated with 100 Ω instead of the 72 Ω used in the DIMM.

Figure 4–46. Differential Memory Clock of a DDR3 SDRAM DIMM Terminated with 100 Ω at the First and Last Component Using a Fly-by Topology on a Board

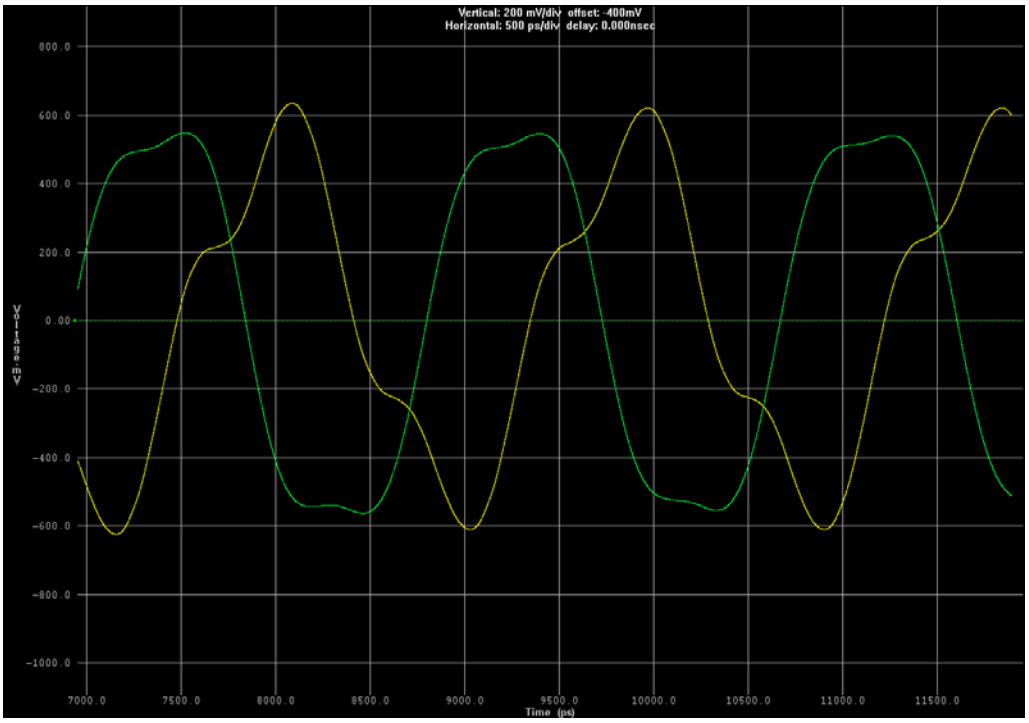

Terminating with 100 Ω instead of 72 Ω results in a slight reduction in peak-to-peak amplitude. To simplify your design, use the terminations outlined in the JEDEC specification for DDR3 SDRAM UDIMM as your guide and perform simulation to ensure that the DDR3 SDRAM UDIMM terminations provide you with 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 4–61 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 4–47 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 Ω instead of the 39 Ω used in the DIMM.

Figure 4–47. 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 Ω

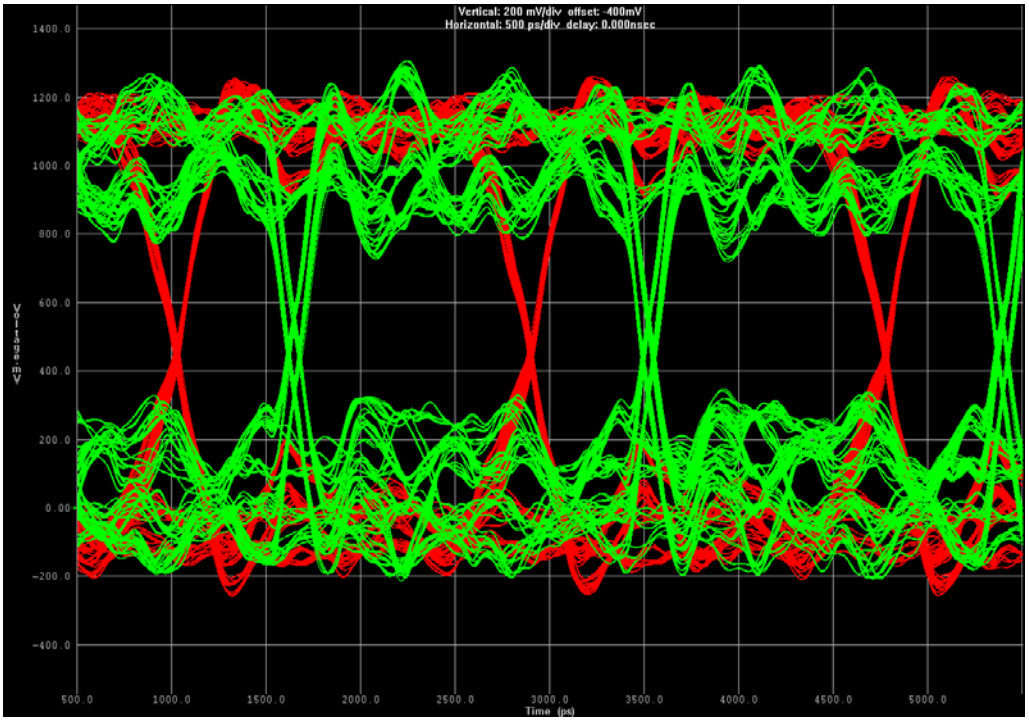

Terminating with 60 Ω instead of 39 Ω results in eye closure in the signal at the first component (the green signal), while there is no effect on the signal at the last component (the red signal). To simplify your design with discrete DDR3 SDRAM components, use the terminations outlined in the JEDEC 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 4–18 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 4–18. 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 4–42), and similar to the write-eye diagram in "DQS, DQ, and DM for DDR3 SDRAM Components" on page 4–46.

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 4–19 on page 4–22 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 4–48 shows the simulation and measurement of the FPGA with a drive strength setting of 8 mA driving a no-parallel termination transmission line.

Figure 4–48. HyperLynx Simulation and Measurement, FPGA Writing to Memory

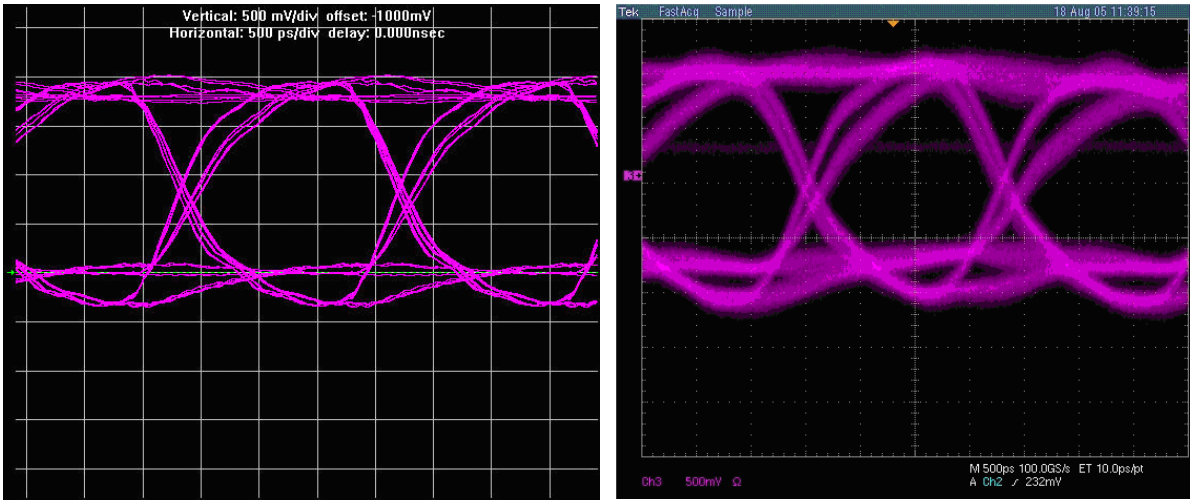

Table 4–19 compares the signals at the DDR2 SDRAM DIMM with no-parallel termination and memory-side series resistors when the FPGA is writing to the memory with 8-mA and 16-mA drive strength settings.

Table 4–19. 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 4–17 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 4–49 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 4–49. Class II Termination Scheme without Memory-Side Series Resistor

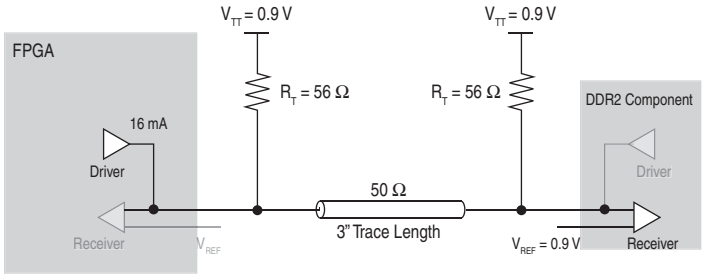

Figure 4–50 shows the simulation and measurement results of the signal seen at a DDR2 SDRAM component of a Class II termination scheme without the DIMM connector and the memory-side series resistor. The FPGA is writing to the memory with a 16-mA drive strength setting.

Figure 4–50. HyperLynx Simulation and Measurement of the Signal, FPGA Writing to Memory

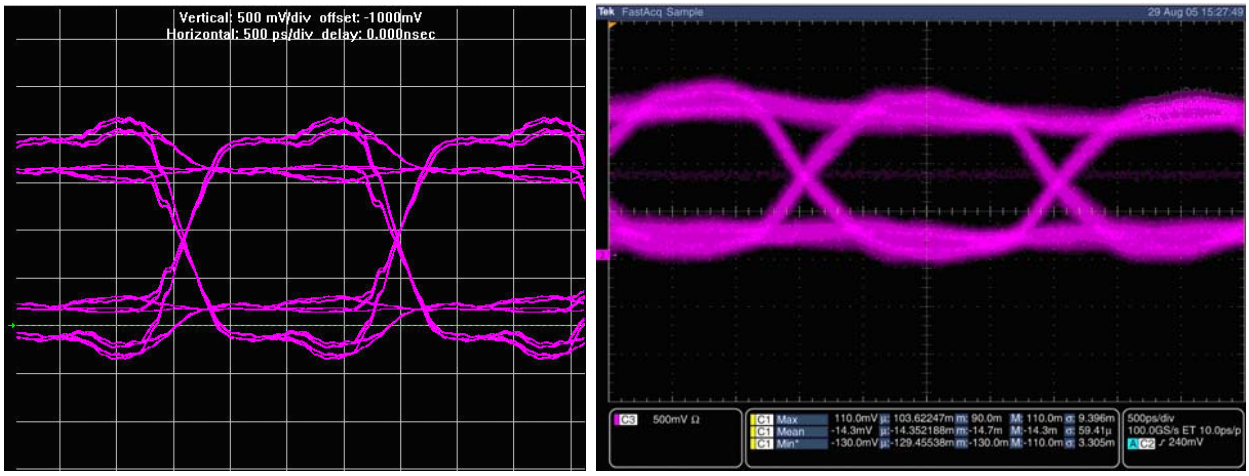

Table 4–20 compares the signal for a single rank DDR2 SDRAM DIMM and a single DDR2 SDRAM component in a Class II termination scheme when the FPGA is writing to the memory.

Table 4–20. Simulation and Board Measurement Results for Single Rank DDR2 SDRAM DIMM and Single DDR2 SDRAM Component *(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

Notes to Table 4–20:

(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 4–51 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 4–32 where the memory-side series resistor is on the DIMM.

Figure 4–51. Class II Termination Scheme without Memory-Side Series Resistor

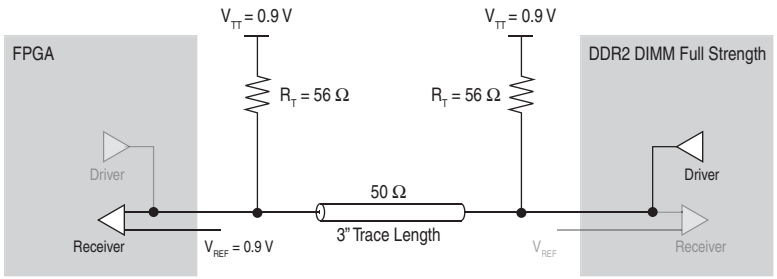

Figure 4–52 shows the simulation and measurement results of the signal seen at the FPGA. The FPGA reads from memory without the source-series resistor near the DDR2 SDRAM component on a Class II-terminated transmission line. The FPGA reads from memory with a full drive strength setting.

Figure 4–52. HyperLynx Simulation and Measurement, FPGA Reading from the DDR2 SDRAM Component

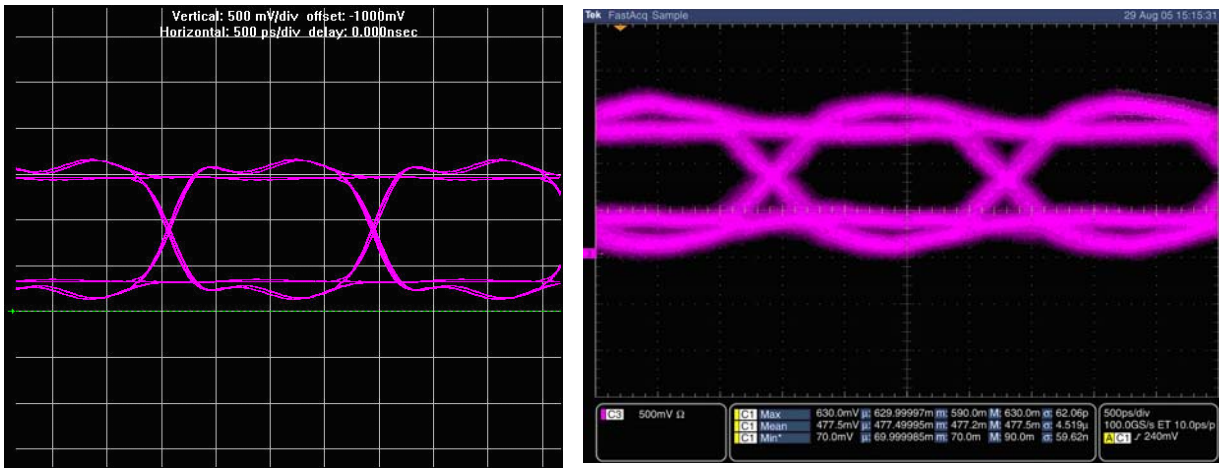

Table 4–21 compares the signal at a single rank DDR2 SDRAM DIMM and a single DDR2 SDRAM component of a Class II termination scheme. The FPGA is reading from memory with a full drive strength setting.

Table 4–21. Simulation and Board Measurement Results of Single Rank DDR2 SDRAM DIMM and DDR2 SDRAM Component *(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 4–21:

(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 4–53 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 4–53. HyperLynx Simulation with a 16-mA Drive Strength Setting on the FPGA

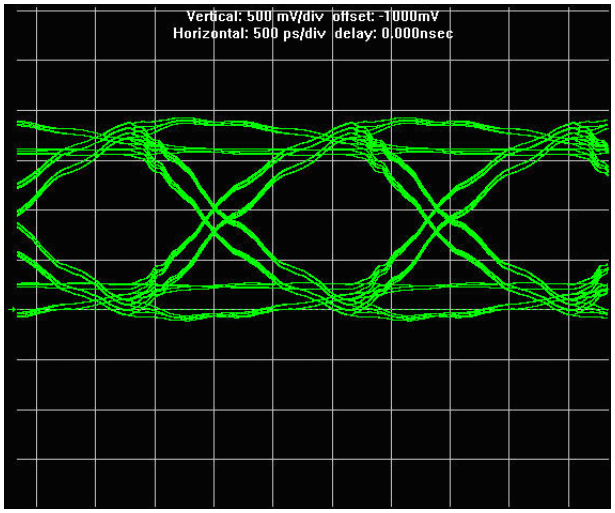

Table 4–22 compares the signals at a single- and dual-rank DDR2 SDRAM DIMM of a Class II and far-end source-series termination when the FPGA is writing to the memory with a 16-mA drive strength setting.

Table 4–22. Simulation Results of Single- and Dual-Rank DDR2 SDRAM DIMM *(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 4–22:

(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 the *Dual-DIMM DDR2 and DDR3 SDRAM Board Design Guidelines* chapter.

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.

For more information about board skew parameters, refer to Board Skews in the Implementing and Parameterizing Memory IP chapter. For assistance in calculating board skew parameters, refer to the Board Skew Parameters Tool, which is available through the solution at this location:
http://www.altera.com/support/kdb/solutions/rd10232012_771.html

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

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

- DIMM—UDIMM topology
- DIMM—RDIMM topology
- Discrete components laid out in UDIMM topology
- Discrete components laid out in RDIMM topology

Trace lengths for CLK and DQS should tightly match for each memory component. To match the trace lengths on the board, a balanced tree topology is recommended for clock and address and command signal routing. In addition to matching the trace lengths, you should ensure that DDR timing is passing in the Report DDR Timing report. For Stratix devices, this timing is shown as Write Leveling tDQSS timing. For Arria and Cyclone devices, this timing is shown as CK vs DQS timing

Table 4–23 lists device family topology support.

Table 4–23. Device Family Topology Support

Device	I/O Support
Arria II	Non-leveling
Arria V GX, Arria V GT	Non-leveling
Arria V GZ	Leveling
Cyclone V GX, Cyclone V GT	Non-leveling
Stratix III	Leveling
Stratix IV	Leveling
Stratix V	Leveling

Table 4–24 lists DDR2 SDRAM layout guidelines. These guidelines are Altera recommendations, and should not be considered as hard requirements. You should perform signal integrity simulation on all the traces to verify the signal integrity of the interface. You should extract the slew rate and propagation delay information, enter it into the IP and compile the design to ensure that timing requirements are met.

The following layout guidelines also apply to DDR3 SDRAM without leveling interfaces.

Table 4–24. DDR2 SDRAM Layout Guidelines (Part 1 of 4) (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	■ 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

Table 4–24. DDR2 SDRAM Layout Guidelines (Part 2 of 4) *(1)*

Parameter	Guidelines
Decoupling Parameter	■ 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 VDD and VDDQ pin
Power	■ 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	All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propogation variance, Altera recommend that signals from the same net group always be routed on the same layer. ■ Use 45° angles (*not* 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	■ A 4.7 KΩ resistor to ground is recommended for each Clock Enable signal. You can place the resistor at either the memory end or the FPGA end of the trace. ■ Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm) ■ These signals should maintain a10-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 the signals in the `Address/Command` Group or up to 100 mils (2.54 mm) longer than the signals in the `Address/Command` Group. ■ Trace lengths for CLK and DQS should closely match for each memory component. To match trace lengths on the board, a balanced tree topology is recommended for clock and address and command signal routing. For Stratix device families, ensure that Write Leveling tDQSS is passing in the DDR timing report; for Arria and Cyclone device families, verify that CK vs DQS timing is passing in the DDR timing report.
Address and Command Routing	■ 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.

Table 4–24. DDR2 SDRAM Layout Guidelines (Part 3 of 4) *(1)*

Parameter	Guidelines

DQ, DM, and DQS Routing Rules	■ 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:
	  ■ 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:
	  ■ 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.
	■ Trace lengths for CLK and DQS should closely match for each memory component. To match trace lengths on the board, a balanced tree topology is recommended for clock and address and command signal routing. For Stratix device families, ensure that Write Leveling tDQSS is passing in the DDR timing report; for Arria and Cyclone device families, verify that CK vs DQS timing is passing in the DDR timing report.

Termination Rules	■ 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:
	  ■ 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 20Ω
	  ■ 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:
	  ■ 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 4–24. DDR2 SDRAM Layout Guidelines (Part 4 of 4) *(1)*

Parameter	Guidelines
Quartus II Software Settings for Board Layout	■ 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. Check the turnaround timing in the Report DDR timing report and ensure the margin is positive before board fabrication. Functional failure happens if the margin is more than 0.

Note to Table 4–24:

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

Figure 4–54. Balanced Tree Topology

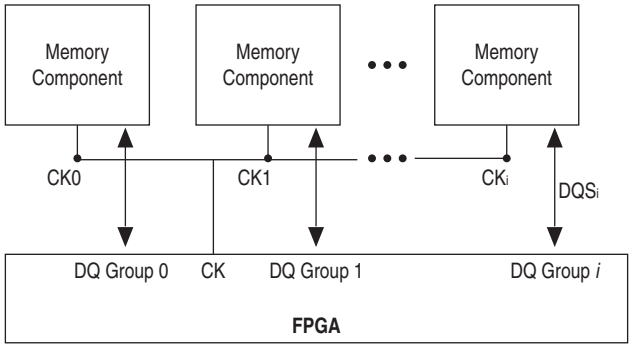

CK_i = Clock signal propagation delay to device *i*

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

Layout Guidelines for DDR3 SDRAM Interface

Table 4–25 lists DDR3 SDRAM layout guidelines.

Unless otherwise specified, the guidelines in Table 4–25 apply to the following topologies:

- DIMM—UDIMM topology
- DIMM—RDIMM topology
- DIMM—LRDIMM topology

☞ Not all versions of the Quartus II software support LRDIMM. Contact Altera if you require LRDIMM support in the Quartus II software.

- Discrete components laid out in UDIMM topology

- Discrete components laid out in RDIMM topology

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

Refer to the *External Memory Interface Spec Estimator* for all supported frequencies and topologies.

For frequencies greater than 533 MHz, when you are calculating the delay associated with a trace, you must take the FPGA package delays into consideration. For more information, refer to Package Deskew.

For device families that do not support write leveling, refer to Layout Guidelines for DDR2 SDRAM Interface.

Table 4–25. DDR3 SDRAM Layout Guidelines (Part 1 of 5) *(1)*

Parameter	Guidelines
Impedance	■ All signal planes must be 50 Ω, single-ended, ±10%. ■ All signal planes must be 100 Ω, differential ±10%. ■ All unused via pads must be removed, because they cause unwanted capacitance.
Decoupling Parameter	■ 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 0.1µF cap for every other V_{TT} pin. ■ Use a 0.1µF cap and 0.01µF cap for every V_{DDQ} pin.
Power	■ 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.

Maximum Trace Length *(2)*	■ Even though there are no hard requirements for minimum trace length, you need to simulate the trace to ensure the signal integrity. Shorter routes result in better timing.
	■ For DIMM topology only:
	  ■ Maximum trace length for all signals from FPGA to the first DIMM slot is 4.5 inches.
	  ■ Maximum trace length for all signals from DIMM slot to DIMM slot is 0.425 inches.
	■ For discrete components only:
	  ■ 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.

Table 4–25. DDR3 SDRAM Layout Guidelines (Part 2 of 5) *(1)*

Parameter	Guidelines
General Routing	All specified delay matching requirements include PCB trace delays, different layer propagation velocity variance, and crosstalk. To minimize PCB layer propogation variance, Altera recommend that you route signals from the same net group on the same layer. ■ Use 45° angles (*not* 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	■ A 4.7 KΩ resistor to ground is recommended for each Clock Enable signal. You can place the resistor at either the memory end or the FPGA end of the trace.
	■ Route clocks on inner layers with outer-layer run lengths held to under 500 mils (12.7 mm).
	■ Route clock signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not exceed 0.69 t_{CK}. For different DIMM configurations, check the appropriate JEDEC specification.
	■ These signals should maintain the following spacings:
	  ■ 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 3× trace-to-plane distance.
	  ■ 20-mil spacing for parallel runs between 1 and 6 inches or 4× trace-to-plane distance.
	■ Clocks should maintain a length-matching between clock pairs of ±5 ps or approximately ±25 mils (0.635 mm).
	■ 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.
	■ 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.
	■ Route all addresses and commands to match the clock signals to within ±20 ps or approximately ± 125 mil (± 3.175 mm) to each discrete memory component. Refer to Figure 4–55.

Table 4–25. DDR3 SDRAM Layout Guidelines (Part 3 of 5) *(1)*

Parameter	Guidelines
Address and Command Routing	■ Route address and command signals in a daisy chain topology from the first SDRAM to the last SDRAM. The maximum length of the first SDRAM to the last SDRAM must not be more than 0.69 t_{CK}. For different DIMM configurations, check the appropriate JEDEC specifications. ■ UDIMMs are more susceptible to cross-talk and are generally noisier than buffered DIMMs. Therefore, route address and command signals of UDIMMs on a different layer than data signals (DQ) and data mask signals (DM) and with greater spacing. ■ Do not route differential clock (CK) and clock enable (CKE) signals close to address signals. ■ Route all addresses and commands to match the clock signals to within ±20 ps or approximately ± 125 mil (± 3.175 mm) to each discrete memory component. Refer to Figure 4–55. ■ Parallelism rules for address and command and clock signals are as follows: ■ 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)

Table 4–25. DDR3 SDRAM Layout Guidelines (Part 4 of 5) *(1)*

Parameter	Guidelines

DQ, DM, and DQS Routing Rules	■ 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.
	■ 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.
	■ Parallelism rules for all signals (other than address and command) are as follows:
	  ■ 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)
	■ 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:
	  ■ Minimum and maximum die IOE skew or delay mismatch
	  ■ Minimum and maximum device package skew or mismatch
	  ■ Board delay mismatch of 20 ps
	  ■ Memory component DQ skew mismatch
	■ Increasing any of these four parameters runs the risk of the deskew algorithm limiting, failing to correct for the total observed system skew. If the algorithm cannot compensate without limiting the correction, timing analysis shows reduced margins.
	■ For ALTMEMPHY-based interfaces, keep the maximum byte-lane group-to-byte group matched length deviation to ± 150 ps or ± 0.8 inches (± 20 mm).
	■ For UniPHY-based interfaces, the timing between the DQS and clock signals on each device calibrates dynamically to meet t_{DQSS}. To make sure the skew is not too large for the leveling circuit's capability, refer to Figure 4–56 and follow these rules:
	  ■ Propagation delay of clock signal must not be shorter than propagation delay of DQS signal at every device:
	  $(CK_i) - DQS_i > 0$; $0 < i <$ number of components – 1
	  ■ Total skew of CLK and DQS signal between groups is less than one clock cycle:
	  $(CK_i + DQS_i)$ max – $(CK_i + DQS_i)$ min < 1 × t_{CK}
	
	  (If you are using a DIMM topology, your delay and skew must take into consideration values for the actual DIMM.)

Table 4–25. DDR3 SDRAM Layout Guidelines (Part 5 of 5) (1)

Parameter	Guidelines
Termination Rules	■ 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 Ω 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 Ω 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. ■ If you include a compensation capacitor at the first memory load, it may improve the waveform signal integrity. ■ 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	■ 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 provide better results.

Notes to Table 4–24:

(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.

Length Matching Rules

This section provides additional guidance on length matching for different types of DDR3 signals.

Route all addresses and commands to match the clock signals to within ±20 ps or approximately ± 125 mil (± 3.175 mm) to each discrete memory component. Figure 4–55 shows the DDR3 SDRAM component routing guidelines for address and command signals.

Figure 4–55. DDR3 SDRAM Component Address and Command Routing Guidelines

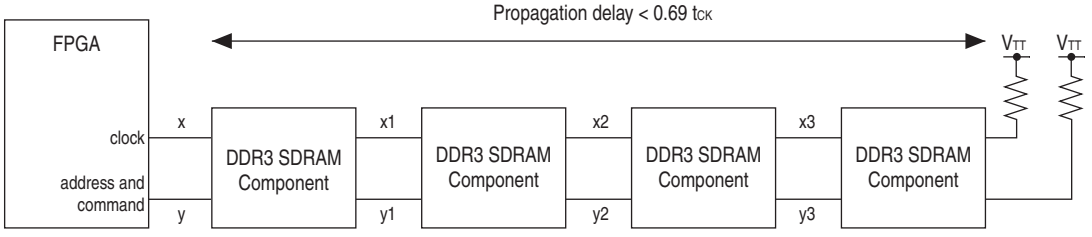

If using discrete components:
x = y ± 125 mil
x + x1 = y + y1 ± 125 mil
x + x1 + x2 = y + y1 + y2 ± 125 mil
x + x1 + x2 + x3 = y + y1 + y2 +y3 ± 125 mil

If using a DIMM topology:
x=y +/- 125 mil

The timing between the DQS and clock signals on each device calibrates dynamically to meet t_{DQSS}. Figure 4–56 shows the delay requirements to align DQS and clock signals. To ensure that the skew is not too large for the leveling circuit's capability, follow these rules:

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

```
CKi  – DQSi > 0; 0 < i < number of components – 1
```

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

```
(CKi + DQSi) max – (CKi + DQSi) min < 1 × tCK
```

Figure 4–56. Delaying DQS Signal to Align DQS and Clock

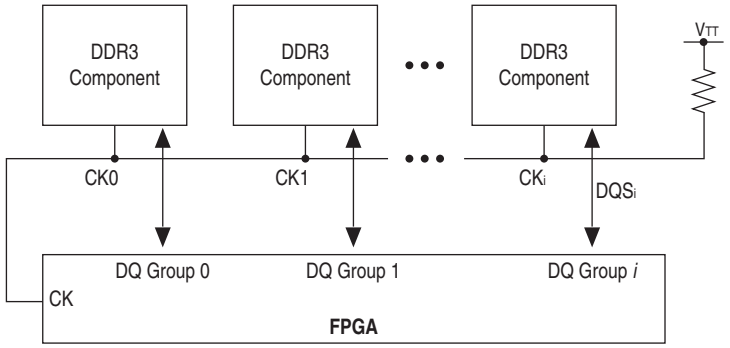

CK_i = Clock signal propagation delay to device *i*

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

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

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

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

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 4–57 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 4–45, 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 4–57 shows an example of a single fly-by network topology.

Figure 4–57. Single Fly-By Network Topology

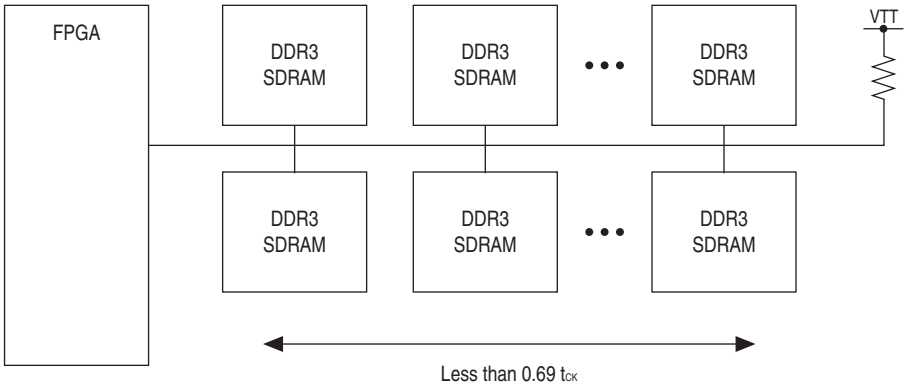

Every DDR3 SDRAM component connected to the signal is a small load that causes discontinuity and degrades the signal. When using a single fly-by network topology, to minimize signal distortion, follow these guidelines:

- 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 4–58 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 4–58. Double Fly-By Network Topology

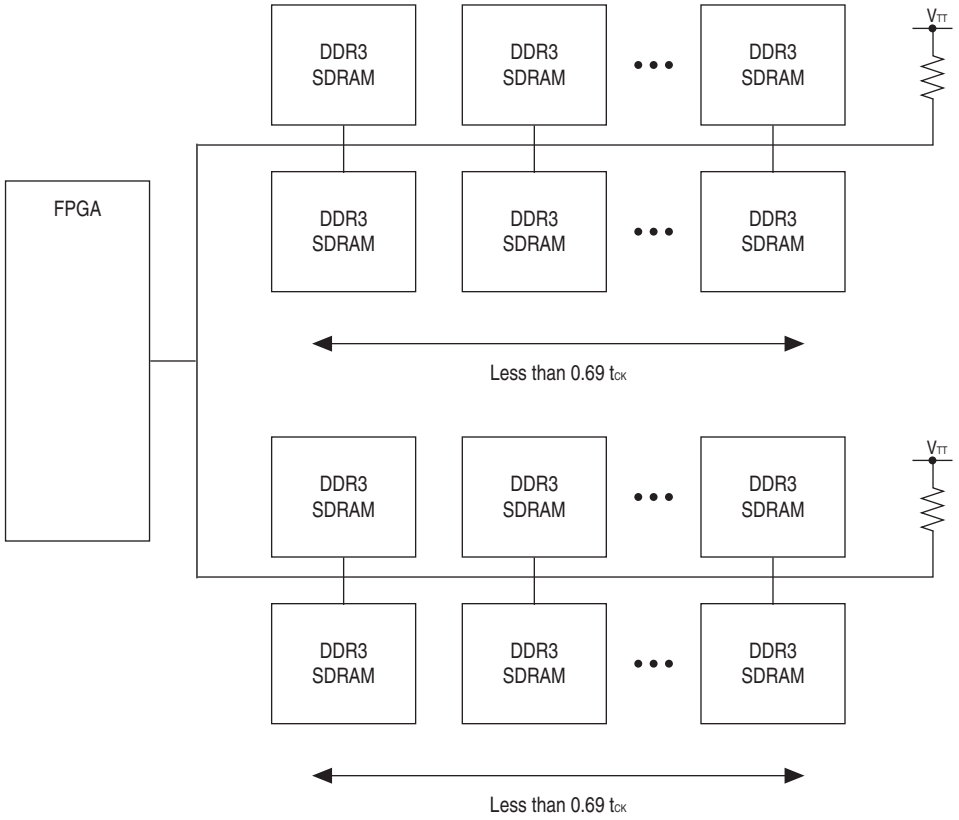

You need to carry out some simulations to find the location of the split, and the best impedance for the traces before and after the split.

Figure 4–59 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 4–59. Minimizing Discontinuity Effect

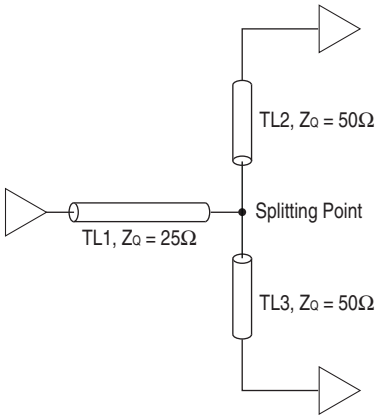

You can also consider using a DIMM on each branch to replace the components. Because the trade impedance on the DIMM card is 40 Ω to 60 Ω, perform a board trace simulation to control the reflection to within the level your system can tolerate.

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.

Package Deskew

Trace lengths inside the device package are not uniform for all package pins. The nonuniformity of package traces can affect system timing for high frequencies. In the Quartus II software version 12.0 and later, a package deskew option is available.

If you do not enable the package deskew option, the Quartus II software uses the package delay numbers to adjust skews on the appropriate signals; you do not need to adjust for package delays on the board traces. If you do enable the package deskew option, the Quartus II software does not use the package delay numbers for timing analysis, and you must deskew the package delays with the board traces for the appropriate signals for your design.

DQ/DQS/DM Deskew

For Stratix V DDR3 UniPHY designs running at frequencies equal to or greater than 533 MHz, you must take into account the package delays of all data pins (DQ/DQS/DM) when you calculate the total trace delay. To get the package delay information, follow these steps:

1. Select the **DQ/DQS Package Deskew** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Run **Analysis and Synthesis** in the Quartus II software.
5. Run the `<core_name>_p0_pin_assignment.tcl` script.
6. Compile your design.
7. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

Address and Command Deskew

For frequencies of 1 GHz or higher, you should deskew address and command delays as follows:

1. Select the **Address/Command Package Deskew** checkbox on the **Board Settings** tab of the parameter editor.
2. Generate your IP.
3. Instantiate your IP in the project.
4. Run **Analysis and Synthesis** in the Quartus II software.
5. Run the `<core_name>_p0_pin_assignment.tcl` script.
6. Compile your design.
7. Refer to the **All Package Pins** compilation report, or find the pin delays displayed in the `<core_name>.pin` file.

Deskew Example

Consider an example where you want to deskew an interface with 4 DQ pins, 1 DQS pin, and 1 DQSn pin. Let's assume an operating frequency of 667 MHz, and the package lengths for the pins reported in the **.pin** file as as follows:

```
dq[0] = 120 ps
dq[1] = 120 ps
dq[2] = 100 ps
dq[3] = 100 ps
dqs   = 80 ps
dqs_n = 80 ps
```

Figure 4–60 illustrates this example.

Figure 4–60. Deskew Example

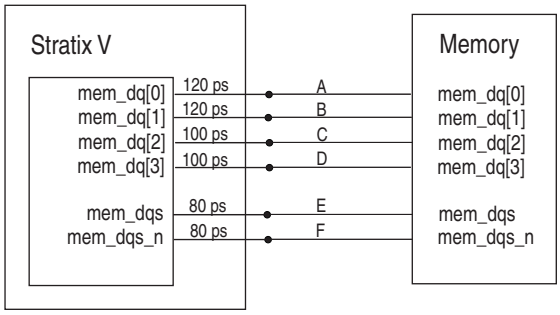

When you perform length matching for all the traces in the DQS group, you must take package delays into consideration. Because the package delays of traces A and B are 40 ps longer than the package delays of traces E and F, you would need to make the board traces for E and F 40 ps longer than the board traces for A and B.

A similar methodology would apply to traces C and D, which should be 20 ps longer than the lengths of traces A and B.

Figure 4–61 shows this scenario with the length of trace A at 450 ps.

Figure 4–61. Deskew Example With Trace Delay Calculations

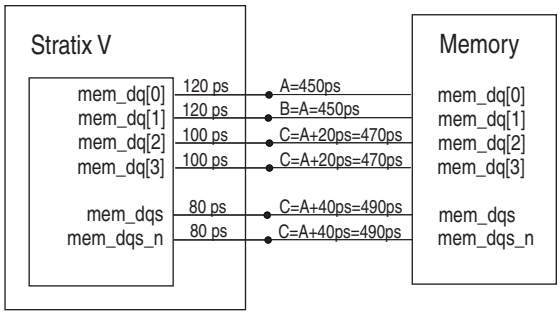

When you enter the board skews into the Board Settings tab of the DDR3 parameter editor, you should calculate the board skew parameters as the sums of board delay and corresponding package delay. If a pin does not have a package delay (such as address and command pins), you should use the board delay only.

The example of Figure 4–61 shows an ideal case where board skews are perfectly matched. In reality, you should allow plus or minus 2 ps of skew mismatch within a DQS group (DQ/DQS/DM).

Package Migration

It is important to note that package delays can be different for the same pin in different packages. If you want to use multiple migratable packages in your system, you should compensate for package skew as described below:

Assume two migratable packages, device A and device B, and that you want to compensate for the board trace lengths for device A. Follow these steps:

1. Compile your design for device A, with the Package Skew option enabled.
2. Note the skews in the `<core_name>.pin` file for device A. Deskew these package skews with board trace lengths as described in the preceding examples.
3. Recompile your design for device A.
4. Compile your design for device B, with the Package Skew option enabled.
5. Note the skews in the `<core_name>.pin` file for device B. Calculate the maximum absolute package skew difference between the two packages.
6. Take the value calculated in step 4, and add to it the board deskew uncertainty (which, by default, is 20 ps). Enter the resulting value as the **Maximum skew within a DQS group** on the board settings tab.
7. Recompile the design for device B.
8. Verify that timing is correct for both device A and device B.

Document Revision History

Table 4–26 lists the revision history for this document.

Table 4–26. Document Revision History

Date	Version	Changes
November 2012	5.0	■ Updated Layout Guidelines for DDR2 SDRAM Interface and Layout Guidelines for DDR3 SDRAM Interface. ■ Added LRDIMM support. ■ Added Package Deskew section.
June 2012	4.1	Added Feedback icon.
November 2011	4.0	Added Arria V and Cyclone V information.
June 2011	3.0	■ 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	Added *DDR3 SDRAM Interface Termination, Drive Strength, Loading, and Board Layout Guidelines* chapter with Stratix V information.
July 2010	2.0	Updated Arria II GX information.
April 2010	1.0	Initial release.