
The Benefits of Altera's High-Speed DDR SDRAM Memory Interface Solution

Introduction

This white paper provides a general overview of a double data rate (DDR) SDRAM interface and discusses Altera's solution for implementing 400 megabits per second (Mbps) DDR interfaces using Stratix™ and Stratix GX FPGAs. Experimental results from hardware testing are also presented.

The Market Need for DDR SDRAMs

SDRAMs have traditionally been used in personal computers (PCs). Greatly increased performance of PCs has resulted in the need for faster, more efficient, larger and cheaper memories. To meet these needs, DDR SDRAMs were introduced as a cost-effective path for upgrading data bandwidth to memory. They have fast become the memory of choice in PC and server markets.

This drop in price has not gone un-noticed in the networking market. Ever increasing end-performance requirements in this market have resulted in greatly increased data bandwidths, which in turn resulted in the need for faster memories. Memory requirements for networking applications fall into several categories, with DRAMs being primarily used for packet buffer memory (where large amounts of packet memory are required to store entire packets while network processors process the packet headers).

General Overview of the DDR SDRAM Interface

Traditional single data rate (SDR) SDRAM architectures are performance-limited in their interfaces; therefore, DDR SDRAMs were introduced as an enhancement. While most of the addressing and command control interface is identical, the fundamental difference is in the data interface. The DDR SDRAM architecture employs a $2n$ -prefetch architecture, where the internal data bus is twice the width of the external data bus. A single read or write cycle involves a single $2n$ -bit wide, one-clock-cycle data transfer at the core, and two corresponding n -bit wide, one-half-clock-cycle data transfers at the I/O. Thus, this enables high-speed operation as the internal column accesses are half the frequency of the external data transfer rate.

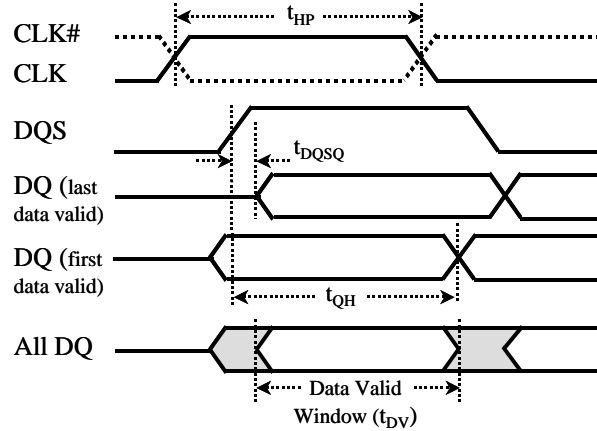
The data interface is designed to transfer two n -bit wide words per clock cycle. However, if the data transfers were based on a free-running system clock, the maximum frequency would be attained as soon as the total output access and flight time equaled the bit time. Additionally, in such a scheme, the data cannot track the clock with changes in temperature and loading, cutting down the effective data valid window and further limiting the maximum attainable frequency. To alleviate these limitations, DDR SDRAMs use a byte-wide, bidirectional data strobe (DQS) that is transmitted externally, along with data (DQ) for data capture. DQS is transmitted edge-aligned by the DDR SDRAM during reads, and center-aligned by the controller during writes to the memory. The DDR SDRAM utilizes on-chip delay-locked loops (DLLs) to clock out DQS and corresponding DQs, ensuring that they are well matched and that they track each other with changes in voltage and temperature.

DDR SDRAMs feature differential clock inputs (CK and CK#), helping to mitigate the effects of duty-cycle variation on the clock inputs. As with SDR SDRAMs, DDR SDRAMs also support the use of data mask (DM) signals to mask data bits during write cycles. All inputs and outputs are compliant with the JEDEC standard for SSTL-2. The following sections will provide a general overview of the read and write cycles, focusing on the DQS-DQ relationship.

Read Cycle

During a read cycle, DQS and the corresponding DQ group are clocked out by the memory using the same internal clock. An on-chip DLL is used to edge-align DQ and DQS transitions to CK. Even though a DLL is used to clock out DQS and DQ, a finite mismatch occurs between DQS and its related byte-wide DQ group. This mismatch is represented by the t_{DQS0} parameter. Therefore, a data word is valid from the point of the latest switching signal in a group to the earliest switching signal (off the following DQS edge) in the group. Furthermore, variations in clock duty cycle eventually result in loss of the data valid window. The effective data valid window is depicted in Figure 1.

Figure 1 Data Valid Window at the Memory Output



The data valid window is further reduced by the time it takes the signals to arrive at the controller pins, primarily due to a combination of factors such as connector and board-induced skews. DQS can then be used for data capture at the controller after appropriately positioning the DQS in relation to the data valid window.

The amount of time that DQS must be delayed is governed by board-induced skew between the DQS and DQ group, the resulting data valid window at the controller, and the sampling window requirements at the controller capture registers. The minimum time that DQS must be delayed with respect to DQ should account for t_{DQSQ} and any board-induced skews. Similarly, the maximum amount of time that DQS should be delayed is calculated from t_{QH} and board-induced skew. Nominal delay would imply DQS is center-aligned with the DQ data valid window, as seen in Figure 2 and Table 1.

Figure 2 Optimal DQS Delay

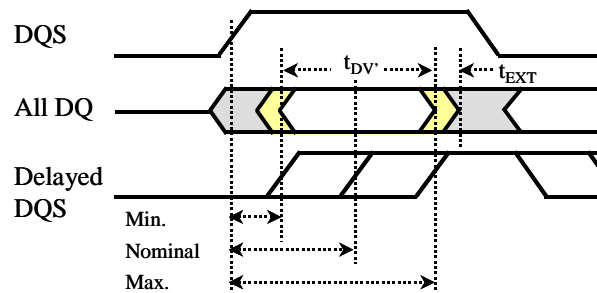


Table 1 Optimal DQS Delay Note (1)

Parameter	Symbol	DDR 333	DDR 400	Units
Board Trace Skew	t_{EXT}	0.05	0.05	ns
Data Valid Window at Controller	$t_{DV} = t_{DV} - 2 * t_{EXT}$	1.70	1.25	ns
Minimum DQS Delay	$t_{DQSQ} + t_{EXT}$	0.45	0.45	ns
Nominal DQS Delay	$(t_{DQSQ} + t_{EXT}) + (t_{DV}/2)$	1.30	1.08	ns
Maximum DQS Delay	$t_{QH} - t_{EXT}$	2.15	1.70	ns

Note to Table 1:

(1) t_{EXT} is based on estimated bus path mismatch of 50 ps. t_{EXT} does not account for additional board induced effects (e.g., crosstalk, ISI).

The controller designers can use several approaches to align the DQS to the center of the data valid window: board trace delay on DQS, on-chip trace delay on DQS, on-chip DLLs, or on-chip PLLs.

Board Trace Delay on DQS

Board trace delay on DQS is a more traditional approach used for aligning DQS and a related DQ group. While this can suffice for low-frequency systems and for relatively simple topologies, it proves to be a performance barrier for more sophisticated systems for the following reasons:

- To use the 400-Mbps case as an example, the nominal delay for DQS with respect to DQ is 1.08 ns. To achieve this delay, approximately 6 to 7 inches of trace length¹ must be added to the DQS line. Not only does this further complicate board layout, it also can result in increased board cost. Especially, when interfacing with DIMMs, routing the additional length required for each of the DQS signals can be difficult.
- The required delay and resulting trace length must be accurately pre-determined. This locks the interface to a specific frequency, and does not leave a controller designer with much flexibility. Any changes in interface frequency would require laying out the board again.
- Increased DQS trace length leads to a greater margin of error due to increased susceptibility to noise and skew injection, and delay variance with changing voltage and temperature.
- Increased trace length also results in higher load capacitance on the DQS line. Thus, rise/fall times are further compromised, limiting the maximum attainable frequency.

¹ This length estimation is based on an approximate delay of 160 ps/inch for a FR4 laminate Microstrip with a 50-Ω characteristic impedance.

On-Chip Trace Delay on DQS

DQS is routed to the capture registers through a line different from that of the DQ signals. Delay implemented on this dedicated DQS line either through increased length, or through use of a line with different characteristics, results in a delayed version of DQS reaching the registers. While this approach is similar in theory to the board trace delay approach, it is primarily used to fine-tune the delays achieved in the prior approach. Therefore, limitations of the board trace delay approach are not addressed. This approach continues to prevent higher performance.

On-Chip Delay Elements

This approach utilizes a number of delay elements connected in series to achieve a pre-determined delay. The delay and corresponding number of delay elements must be calculated based on frequency of operation and the right number of elements picked for each frequency bin. A designer can then implement varying design techniques to use a combination of coarse and fine delay to further fine-tune to the desired delay. However, delay elements are inherently susceptible to process, voltage, and temperature (PVT) variations, with variations seen up to $\pm 40\%$. This variation in delay increases the effective sampling window requirements of the controller, and does not scale with frequency. Thus, the limitation of this approach makes it useful for lower frequencies.

On-Chip DLLs & PLLs

On-chip DLLs and PLLs are used primarily in conjunction with on-chip delay elements to introduce delay onto the DQS lines. By using a reference clock that is at the desired interface frequency and basing the required delay as percent of that clock period, the DLLs or PLLs can then pick the right number of delay elements to achieve the desired delay. As such, the previously stated limitations of delay elements are addressed; the DLLs or PLLs help compensate for PVT variations, allowing for much higher frequencies to be attained. While this approach does involve the use of relatively complicated structures (DLLs and PLLs), the benefits of higher frequencies and system bandwidths offset the cost of implementation.

Write Cycle

During a write cycle, DQS must arrive at the memory pins center-aligned with data. Here, the controller designer can adopt a couple of different approaches: automatic 90° DQS phase offset at the controller outputs or use of board trace delays. The memory device then uses internally-matched routing for the DQS and data to allow for data capture with the DQS signals. The approach adopted by a controller designer must take into account effects on the read-cycle timing. For example, if board trace delays were used, limitations (that have been discussed in prior sections) with this approach would restrict the maximum attainable frequency.

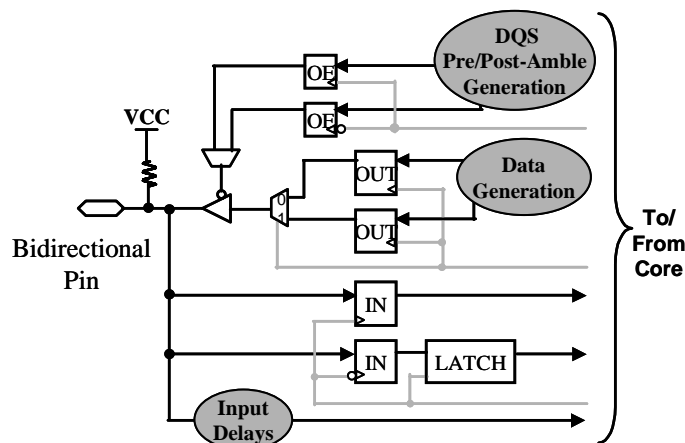
Memory Controller Implementation in Stratix & Stratix GX

This section will concentrate on some of the hardware features that have been implemented in Altera's Stratix and Stratix GX devices that address the system performance bottlenecks discussed in prior sections. In particular, features that enable FPGA-DDR SDRAM interface frequencies up to 400 Mbps are discussed.

I/O Elements

Figure 3 shows an I/O element (IOE) for Altera® Stratix and Stratix GX devices. Each IOE contains six registers and one latch. Two registers and a latch are used for input, two registers are used for output, and two registers are used for output-enable control. Input and output registers are independent of each other, enabling a bidirectional DDR I/O path to be implemented entirely in the IOE. When the input path is active, the output enable disables the tri-state buffer, which prevents data from being sent out on the output path.

Figure 3 DDR I/O Element




Dedicated DQS-DQ Groups

Stratix and Stratix GX devices feature dedicated DQS-DQ groups at the top and bottom of the die. When not interfacing with external memory, designers can use these pins as general-purpose I/O pins. However, when interfacing with external memories such as DDR SDRAMs, dedicated pins must be used for DQS. These in turn are associated with their respective DQ groups. DQS-DQ group ratios can be either 1:8 ($\times 8$ mode—one DQS per 8 DQ pins), or 1:16 ($\times 16$ mode—one DQS per 16 DQ pins), or 1:32 ($\times 32$ mode—one DQS per 32 DQ pins). When in $\times 8$ mode, up to 10 DQS-DQ groups can be implemented on either the top or bottom of the die, interfacing the FPGA to the 72-bit wide DIMMS. Figure 4 shows the layout for the DQS and DQ pins when in $\times 8$ mode.

The dedicated DQS pins then tie internally to a set of delay elements before being routed to the I/O input registers. The cumulative delay of the delay elements is controlled by a DQS phase shift circuitry, and is discussed further in the proceeding section. The dedicated DQS approach does take away some layout flexibility, but the benefits realized with this approach offset the limitations. Use of dedicated pins ensures tighter sampling window requirements at the controller pins. In addition, limiting the use of delay elements to a reduced number of I/O pins results in decreased die area consumption, leading to die cost savings.

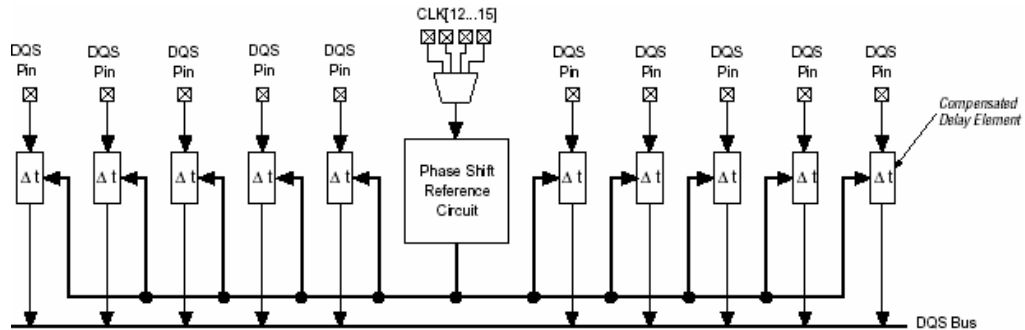
Figure 4 DQ-DQS Groups²



Dedicated DQS Phase Shift Circuitry

Use of dedicated DQS phase shift circuitry (see Figure 5) enables automatic, on-chip delay insertion on incoming DQS signals. This DQS phase shift circuitry uses a frequency reference to generate control signals for the delay elements on each of the dedicated DQS pins, allowing it to compensate for PVT variations.

Figure 5 Dedicated DQS Phase Shift Circuitry



² Altera Stratix Device Handbook, Chapter 8. Double Data Rate I/O Signaling in Stratix and Stratix GX Devices, July 2003, ver. 2.0

The approach of implementing a dedicated phase shift circuitry has several significant advantages:

- DQS can be treated as a 9th “DQ” (within an 8-bit DQ group) and thus can be laid out identical to DQ signals on a board. In addition, designers can treat each DQS-DQ group separately and routed accordingly (without forgetting that eventually all groups have to re-synchronize to a common system clock in the controller). This greatly simplifies board layout and reduces board costs.
- Any board-induced skew effects and board impedance variance with voltage and temperature will equally affect both DQS and DQ, ensuring a tighter timing relationship between DQS and DQ at the controller pins.
- As the on-chip phase-shift circuitry uses a frequency reference, the delay can be tuned to match a desired frequency without having to re-layout the board. Additionally, active PVT compensation ensures the tight DQS-DQ timing relationship is maintained all the way to the capture registers.

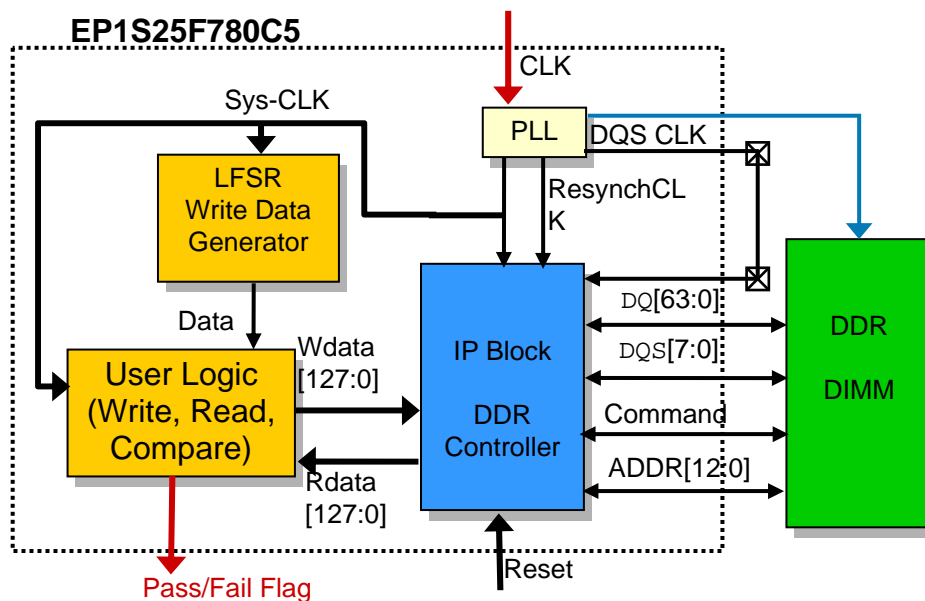
Stratix 400-Mbps DDR SDRAM Characterization

The Stratix DDR SDRAM memory interface solution was thoroughly tested on a hardware platform at 400 Mbps. Summary of the 400-Mbps Stratix DDR SDRAM interface test results, test setup, and FPGA design are presented in the following sections.

Design

Figure 6 shows the block diagram of the DDR SDRAM design used to test the Stratix DDR interface.

Figure 6 Block Diagram of DDR SDRAM Design



The design consists of DDR controller intellectual property (IP), linear feedback shift registers (LFSRs), and user logic blocks. The DDR controller IP block consists of the Altera DDR SDRAM Controller IP MegaCore[®] function. Customers can evaluate this controller by downloading the free version at www.altera.com. The LFSR write data generator block generates random data to be written to the memory. The user logic block verifies if the read data is the same as the write data. The match flag of the comparator is connected to an output pin and is monitored on the Lecroy scope. It goes high when both the inputs to the comparator are the same.

Test Setup

The test setup used to characterize the Stratix DDR SDRAM interface is described in the Table 2.


Table 2 Test Setup for Stratix DDR SDRAM Memory Interface Characterization

Parameter	Description
DDR Controller IP	Altera DDR SDRAM Controller IP
Memory Module	Micron MT16VDDT3264AG-403B, 200 MHz, 184-pin DIMM, 256 Mbytes
Data width	64-bit data bus
Stratix Device	EP1S25F780C5 (fastest speed grade device)
Data Transaction Sequences	Write, read, increment address, write, read, increment address... Write whole DIMM, read whole DIMM (read all, write all)
Termination	Off-chip termination
Burst Length	8 bytes
Data Pattern	PRBS8
Operating Conditions	Worst case, nominal, and best case conditions for temperature and voltage

Board Setup


The board setup used for the Stratix DDR SDRAM memory interface is shown in Figure 7. The Stratix EP1S25F780C5 device is soldered to the back of the motherboard. The daughter board contains four 184-pin DIMM sockets, and is connected to the motherboard via mictor connectors. Each DIMM consists of eight 32M × 8 bit DDR SDRAM memory modules.

Figure 7 Board Setup for Stratix DDR SDRAM Memory Interface



DDR Daughter Board with 4 DIMM Sockets

Stratix DDR SDRAM Mother Board



Flip Side of the Daughter Board

Micro Connectors

Experiment Results

Table 3 lists the f_{MAX} of the Stratix DDR SDRAM interface under different operating conditions f_{MAX} exceeds 400 Mbps across all operating conditions.

Table 3 f_{MAX} of Stratix DDR SDRAM Interface at Different Operating Conditions

Temperature	V_{CC}	Write-all-Read-all	Write-Read-Write-Read
		f_{MAX} (Mbps)	f_{MAX} (Mbps)
70° C	Nom – 5% (1)	410.8	413.2
	Nom (2)	413.6	419.4
	Nom + 5% (3)	419.2	426.4
25° C	Nom – 5%	425.2	430.8
	Nom	434.2	440.6
	Nom + 5%	437.4	443.8
0° C	Nom – 5%	433.2	427.6
	Nom	435	444.4
	Nom + 5%	438.8	450.2

Notes to Table 3:

- (1) Nom – 5%: $V_{CC} = 1.425V$, $V_{DD} = 2.55 V$
- (2) Nom: $V_{CC} = 1.5 V$, $V_{DD} = 2.65 V$
- (3) Nom + 5%: $V_{CC} = 1.575 V$, $V_{DD} = 2.75 V$

Summary of Test Results

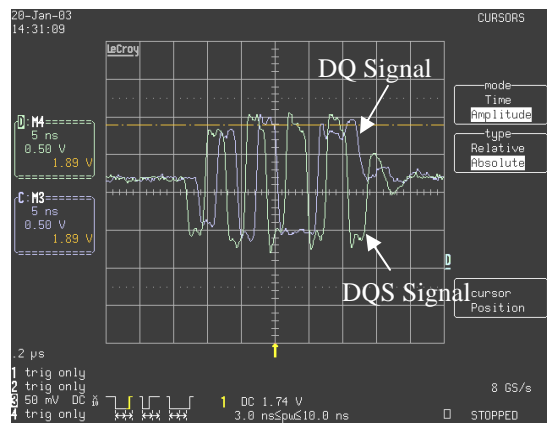
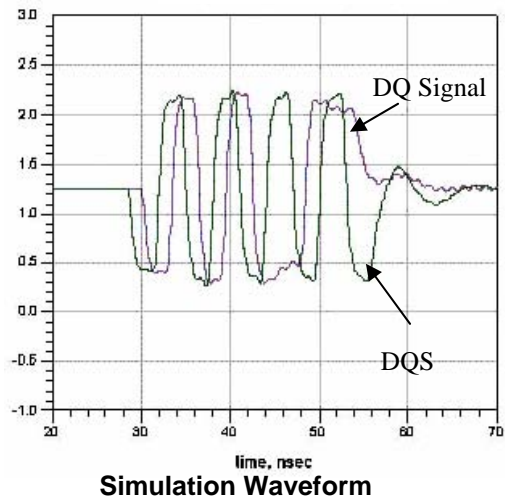
The following shows interoperability test results between the Stratix FPGA and Micron DDR SDRAM DIMM.

- System performance exceeded 400 Mbps under worst-case conditions (low voltage, high temperature, read-all write-all configuration, a noisy signal path that includes Mictor connectors, and a device socket)
- System performance under typical operating conditions is approximately 440 Mbps

Board-Level Simulation Correlation

SPICE simulations were performed on the Stratix DDR board setup and were correlated with measurements. Simulations and board captures correlate well. Figure 8 compares a simulation waveform with the measured waveform for the DQ and DQS in a write transaction. The scale on both the waveforms is the same.

Figure 8 Correlation with Simulation



Stratix DDR SDRAM System-Level Solution

Altera provides a complete system-level solution for designing a DDR SDRAM memory interface with Stratix devices. See Table 4.

Table 4 Stratix DDR SDRAM Memory Interface Systems Solution

Collateral	Description
Intellectual Property	400-Mbps DDR SDRAM IP core available on the Altera web site at www.altera.com 333-Mbps DDR SDRAM IP core from Northwest Logic. Information available at www.altera.com 266-Mbps DDR SDRAM reference design from DCM Technologies
Demonstration Boards	400-Mbps DDR SDRAM with SODIMM Stratix PCI development kit 400-Mbps DDR SDRAM with DIMM Stratix GX development board 400-Mbps DDR SDRAM interface with DIMM (available through Altera FAEs)
Software Support	New features in Quartus® II software version 3.0 for DDR interface support The <code>ALTDQS</code> megafunction to enable the DQS circuitry Timing analysis to calculate round trip delay Timing analysis report will contain DQS read strobe to core register delay information Capability to perform min timing analysis
Timing Analysis	Stratix DDR SDRAM timing analysis (<i>DDR SDRAM Controller Megacore Function User Guide, Appendix B</i>) Timing analysis for resynchronization window (<i>Stratix Device Handbook, Chapter 8</i>) Round trip delay calculator (contact local Altera FAE)
Simulations	SPICE and IBIS models are available for customer simulation of their own board. IBIS models can be downloaded from the Altera web site at www.altera.com/support/software/download/ibis/ibs-ibis_index.html SPICE models – contact Altera for SPICE models
Documentation and Supporting Material	Contact your local FAE for documentation on the Stratix DDR memory interface.

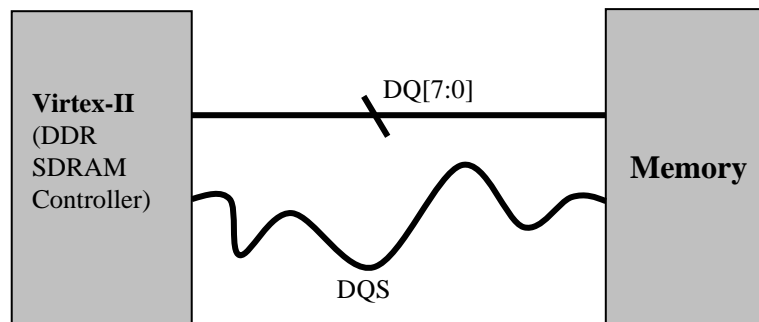
Competitive Advantages

The following section describes the competitive advantage Altera has over Xilinx regarding DDR SDRAM implementations.

Virtex-II 400 Mbps DDR SDRAM Implementation

Xilinx application note 253 (Synthesizable 400 Mb/s DDR SDRAM Controller) describes a Virtex-II DDR 32-bit memory interface at 400 Mbps. The 90° DQS phase shift is implemented as a combination of trace skew and on-chip routing delays. Depending on the performance³ of the system, the trace length of the DQS signal is required to be 4 to 12 inches longer than the DQ signal (see Figure 9).

Figure 9 *Virtex-II DDR SDRAM Interface Board Design*



The following timing equation is used to determine the DQ-DQS trace skew delay on the board:

$$t_{Local\ Routing} + t_{(DQ-DQS)\ Trace\ Skew} + t_{Other\ Skew} = \frac{T}{4} \text{ (for } 90^\circ \text{ Phase Shift)}$$

where T is the time period, $t_{Local\ Routing}$ is the on-chip routing delay of the DQS signal, $t_{(DQ-DQS)\ Trace\ Skew}$ is the trace skew between the DQ and DQS signals, and $t_{Other\ Skew}$ is miscellaneous skew on the chip. Detailed I/O and system timing analysis is required to determine the trace difference between the DQ and DQS signals. For the DDR SDRAM interface described in the Xilinx application note, the required on-board skew between DQ and DQS was calculated to be 750 ps. Assuming that a signal takes 166 ps to traverse one inch of trace on FR4 material, the trace length difference between the DQ and DQS signals is 4.5 inches at 400 Mbps.

³ DDR interface performance between 200 Mbps and 400 Mbps

Issues associated with this implementation are discussed on page 4 in the *Board Trace Delay on DQS* and *On-chip Trace Delay on DQS* sections. Table 5 compares Stratix and Virtex-II DDR SDRAM solutions.

Table 5 Comparing Stratix & Virtex-II/PRO (Xilinx Application Note 253) DDR SDRAM Solutions

Feature	Stratix	Virtex-II/PRO
Ninety degree DQS phase shift	Implemented using dedicated circuitry in Stratix silicon	Implemented using board trace length.
Resource Usage (for 32-bit interface, 8:1 DQ/DQS ratio)	Logic Elements (LEs): 830, PLLs: 1	Slices: 873 (equivalent LEs = 1,746), DLLs: 3
Customer memory I/F upgrade-ability	You can change system performance of the memory I/F without having to re-spin the board.	System performance of the memory interface cannot be changed or a system cannot be upgraded without re-spinning the board
DIMM vs. surface mount (DIMM based systems are speed-limited by bus loading, line lengths, and stubs)	Stratix DDR SDRAM memory I/F was characterized at 400 Mbps with a 184-pin, non-buffered DIMM.	Xilinx application note 253: The surface mount memory device was used to test 400 Mbps memory I/F. Xilinx app note 608: 200 Mbps maximum performance with DIMM.
Data-width	64-bits with DIMM 400 Mbps.	32-bits with surface mount 400 Mbps. 64-bits with DIMM 200 Mbps.

The implementation described in Xilinx Application Note 253APP 253 is sensitive to changes in process, voltage and temperature (PVT) which effects the reliability of the memory interface. Xilinx application note 688 (XAPP688: Creating High-Speed Memory Interfaces with Virtex-II and Virtex-II Pro FPGAs) describes an alternate DDR implementation that only uses internal routing and look-up-table delay with a soft compensation circuit that tries to adjust for PVT changes.

The implementation is very intricate, there are several instances where precise delays and careful I/O and SLICE placement constraints are required. Since this is a soft solution (implemented using internal logic), the number of SLICES (or equivalent LEs) is significantly more than the Altera solution. In addition, this implementation is not backwards compatible with the previous solution as it requires significant board and RTL redesign.

Altera's Quartus software cleanly integrates the instantiation of the DDR SDRAM implementation by using the Megawizard. The Megawizard selects the appropriate I/O and configures the necessary constraints. This allows the design engineer to take advantage of design, development and characterization work that Altera has already completed.

Stratix Series FPGAs were developed with high-speed memories in mind so that future improvements to the solution would be limited to RTL changes and not board redesign.

Conclusion

DDR SDRAMs are quickly becoming the DRAM technology of choice for PCs and server applications, and are experiencing ever-increasing adoption in the networking market. At the same time, increasing system bandwidth requirements are pushing memory interface speeds. The burden of designing a robust interface eventually transfers to the memory system designer. While interface problems limit performance, designers do have solutions available to them to tackle these issues. Hardware approaches adopted in Altera's Stratix and Stratix GX FPGAs offer one such set of solutions, enabling FPGA-to-DDR SDRAM interface performance to run up to 400 Mbps with the dedicated DQS phase shift circuitry. Stratix and Stratix GX DDR interface solutions are hardware-verified to work at speeds up to 400 Mbps.



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
www.altera.com

Copyright © 2004 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries.* All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.