



CYPRESS

Interfacing QDR™-II SRAM with Virtex4™ Devices

Introduction

With the continuous demand for higher performance communications, networking, and digital signal processing (DSP) systems, memory devices are evolving to more closely match the needs of these applications. Specialized memory products that optimize memory bandwidth for a specific system architecture are successfully increasing overall performance in a variety of data processing systems. This application note describes a 250-MHz four-word burst QDR™-II SRAM interface implemented in a Virtex4™ XC4VP25 FF6688 –11 device. This implementation uses local clocking for data capture and an automatic calibration circuit to adjust the look-up table (LUT) delay on the local clocks. This circuit uses one digital clock manager (DCM) that can be shared between several memory interfaces running at the same frequency. The design has been implemented using Cypress CY7C1415AV18-250BZC memory devices. The aggregate throughput for this interface is 500 Mb/s per pin for a 36-bit write bus and a 36-bit read bus operating in DDR mode at 250 MHz.

This application note provides the following key information for the QDR II SRAM interface:

- Description of QDR-II functionality
- Description of the controller and interface in the FPGA for the QDR-II SRAM interface
- Details of the interface timing
- Overview of the implementation

QDR-II SRAM Discussion

Cypress, along with the other QDR consortium members, defined the QDR-II SRAM architecture for high-performance

communications systems. The QDR-II SRAM devices from these companies are pin-compatible.

QDR-II SRAM devices provide concurrent read and write operations and increased data throughput, allowing simultaneous access to the same address location. This innovative architecture outperforms other SRAM devices by up to four times in networking applications, where read and write operations are balanced.

QDR-II SRAM Functional Description

QDR-II SRAM can perform two data writes and two data reads per clock cycle. It uses one port for writing data (D) and another port for reading data (Q). These unidirectional data ports support simultaneous reads and writes and allow back-to-back transactions without the bus contention issues that may occur when using a single bidirectional data bus. Write and read operations share the address bus. QDR-II SRAM devices use three pairs of clocks: Input Clocks K and Kn (controlling the input signals), input clocks C and Cn (controlling the output data bus), and echo clocks CQ and CQn (source synchronous clocks).

QDR-II SRAM devices use either the 1.4V or 1.9V HSTL I/O standard. Write and read operations are burst-oriented and support burst lengths of two and four, so each read and write operation transfers either two or four data words. QDR-II SRAM performs write operations the same way as QDR SRAM. However, read operations in QDR-II SRAM output the data half a clock cycle later than QDRSRAM to improve the clock-to-output time (t_{CO}).

Figure 1 shows a block diagram of the QDR-II SRAM burst-of-4 architecture. The burst-of-2 SRAM architecture will be similar to this but with just two arrays instead of 4 memory arrays.

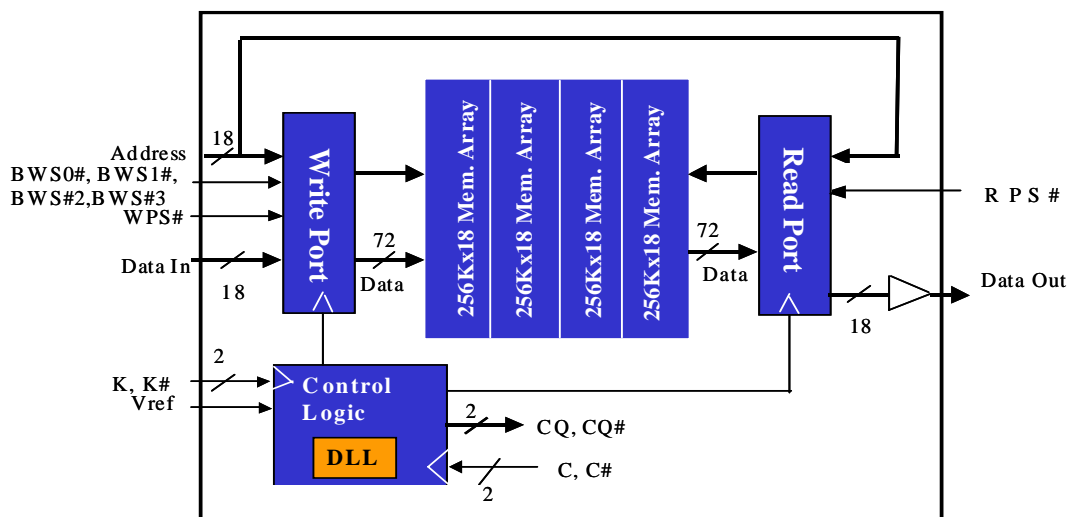


Figure 1. 18-Mbit QDR-II SRAM Burst-of-4 Architecture

QDR-II SRAM Functionality

Both burst-of-2 and burst-of-4 devices provide the same overall bandwidth at a given clock speed. This section describes the functionalities of burst-of-2 and burst-of-4 QDR-II SRAM devices.

Burst-of-2 QDR-II SRAM Devices

Burst-of-2 QDR-II SRAM devices support two-word data transfers on all write and read transactions, requiring a relatively simple controller implementation. This section outlines the basic burst-of-2 functionality for write-only, read-only, and combined read/write operations, assuming the K and Kn clocks are used for both reads and writes, and the C and Cn clocks are connected to VDD (Device is in Single Clock Mode).

Figure 2 shows the burst-of-2 timing diagram for reads and writes. The sizes of the Address, Data-In, and Data-Out buses depend on the memory device with which the FPGA interfaces. The BWSn signal is low for the entire cycle of Figure 2.

Write Cycle

On the rising edge of the K clock, the QDR-II SRAM device latches the control signals WPSn and BWSn and the lower data word on Data-In (D[A] at Cycle 1 of Figure 2). On the rising edge of the Kn clock, the QDR-II SRAM device latches the write address on B (Cycle 1 in Figure 2) and the upper data word (D[B + 1]) on Data-In, thus completing a write cycle.

Read Cycle

On the rising edge of the K clock, the QDR-II SRAM device latches the control signal RPSn and the read address A (Cycle 1 of Figure 2). After a one-and-a-half-clock-cycle latency, the rising edge of Kn (or Cn) clocks out the lower data word (Q[A]) of address A onto the Data-Out bus. The QDR-II SRAM device outputs the upper data word (Q[A + 1]) on the

next rising edge of the K (or C) signal, completing the read cycle.

Read/Write Cycle

Read and write operations occur during the same clock cycle on independent read and write datapaths along with the cycle-shared address bus. Performing concurrent reads and writes does not change the functionality of either transaction. If a read request occurs simultaneously with a write request at the same address, the data on Data-In is forwarded to Data-Out. Therefore, the normal Read latency is not required to access valid data.

Burst-of-4 QDR-II SRAM Devices

Burst-of-4 QDR-II SRAM devices support four-word data transfers on all writes and reads, which reduces address bus activity. However, the control circuitry needed to interface to burst-of-4 QDR-II SRAM devices is more complicated than control circuitry for burst-of-2 QDR-II SRAM devices. The following sections outline the basic burst-of-4 functionality for writes, reads, and read/write operations, assuming the K and Kn clocks are used for both read and write operations and the C and Cn clocks are connected to VDD (Device is in Single Clock Mode).

Figure 3 shows the burst-of-4 timing diagram for alternating reads and writes. The sizes of the Address, Data-In, and Data-Out buses depend on the memory device with which the FPGA interfaces. The BWSn signal is low for the entire cycle of Figure 3.

Write Cycle

The QDR-II SRAM device latches the control signals WPSn and BWSn and the write address A (A at Cycle 1 of Figure 3) on the rising edge of the K clock. On the following K clock rising edge, the QDR-II SRAM device latches the first data word (D[A]) on Data-In. On the next Kn clock rising edge, the second data word is latched (D[A + 1]). The third (D[A + 2])




Figure 2. QDRII SRAM Burst-of-2 Timing Diagram

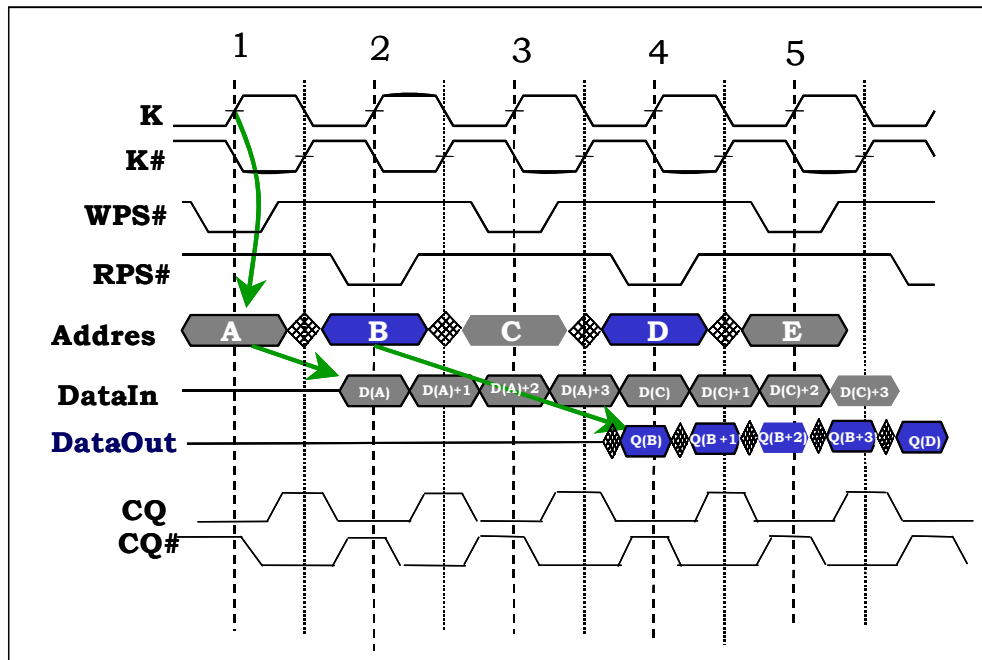


Figure 3. QDR-II SRAM Burst-of-4 Timing Diagram

and fourth ($D[A + 3]$) words are latched in on the subsequent K and Kn clock rising edges, respectively, completing a write cycle.

Read Cycle

The QDR-II SRAM device latches the control signal RPSn and the read address B (Cycle 2 in *Figure 3*) on the rising edge of the K clock. After a one-and-a-half-clock-cycle latency, the rising edge of Kn (or Cn) clocks out the first data word ($Q[B]$) of address B onto the Data-Out bus. The next rising edge of K (or C) clocks out the second data word ($Q[B + 1]$). The subsequent rising edges of Kn (or Cn) and K (or C) clock out the third ($Q[B + 2]$) and fourth ($Q[B + 3]$) words, respectively, completing a read cycle. Single-clock mode uses the K and Kn clocks for both reads and writes.

Read/Write Cycle

Read and write operations occur on subsequent clock cycles on the independent read and write datapaths along with the cycle-shared address bus. Performing concurrent reads and writes does not change the functionality of either transaction. If a read request occurs simultaneously with a write request at the same address, the data on Data-In is forwarded to Data-Out. Therefore, the normal Read latency is not required to access valid data.

QDR-II SRAM Interface Signals

This section provides a description of the clock, control, address, and data signals on a QDR-II SRAM device. *Table 1* shows the QDR-II SRAM interface pins and how to connect them to Virtex4.

Table 1. QDR-II SRAM Interface Pins

Pins	Description	Virtex-II Pin Utilization for QDR-II SRAM
D	Write Data	User I/O pin
Q	Read Data	User I/O pin
K	Write Clock	User I/O pin
Kn	Inverted Write Clock	User I/O pin
C	Read Clock	N/A
Cn	Inverted Read Clock	N/A
CQ	Echo Clock	User I/O pin
CQn	Inverted Echo Clock	User I/O pin
All Other	Address and Command	User I/O pin

When interfacing with one QDR-II SRAM device, Virtex4 uses a single-clock scheme where the QDR-II SRAM device's C or Cn port is tied to VDD (Single Clock Mode).

FPGA QDR-II Memory Interface Controller

This section provides a detailed description of the Controller provided

Interface Block Diagram

Figure 4 shows a block diagram of the QDR II SRAM design. The Controller is composed of three main elements.

- User Interface: Interface on the User side of the controller logic
- Read/Write State Machine
- Physical Interface: Interface to the actual memory device on the controller

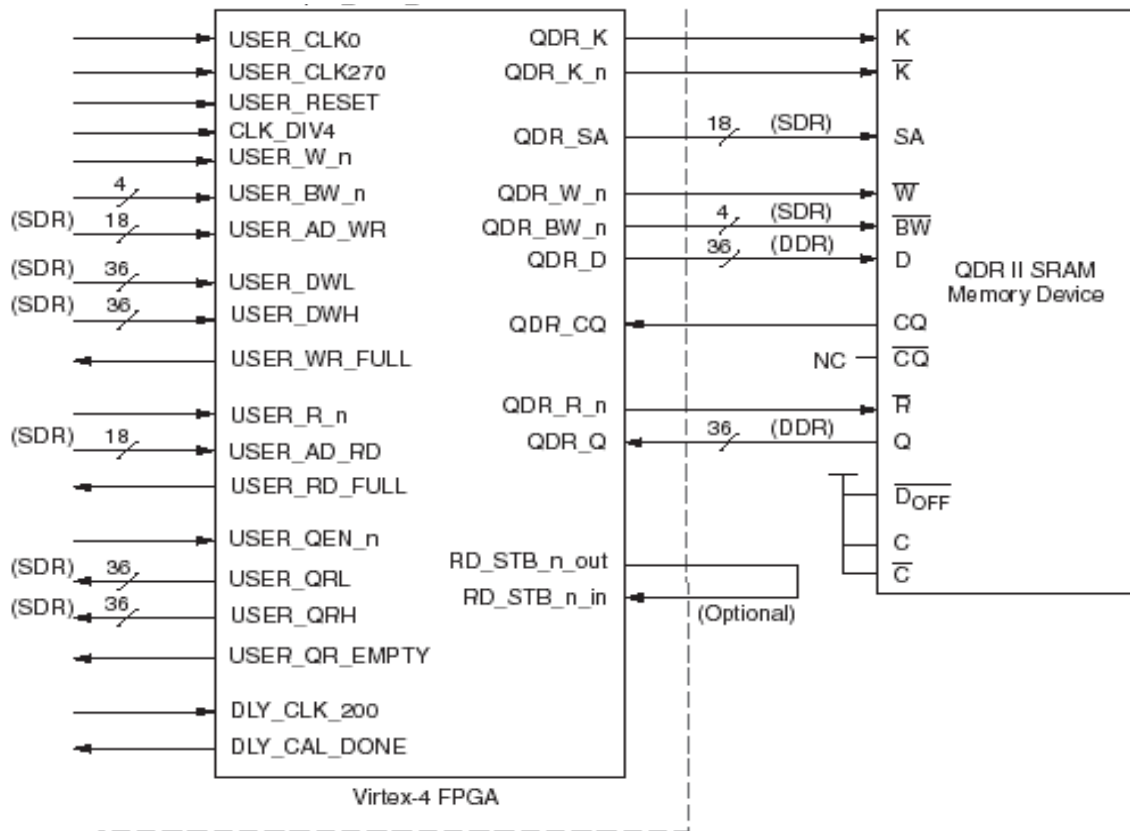


Figure 4. Top-Level Architecture Block Diagram of Reference Design

The signals from the user interface are labeled USER_* and are summarized below:

- The USER_CLK270 is used to center-align the write path signals to the QDR_K and the QDR_K_n clocks. This is achieved by synchronizing the output registers for the write path signals to the USER_CLK270 clock. This clock signal is the same frequency as the USER_CLK0 but is 270 degrees out-of-phase with respect to the USER_CLK0. This ensures adequate set-up and hold margins for the memory device with respect to the incoming the K clocks.
- For memory writes, 36-bit data is presented on USER_DWL and USER_DWH in Single Data Mode (SDR). The corresponding write address is presented on USER_AD_WR. A write is indicated when the Write Command signal, USER_W_n, is asserted LOW. The data and address are only synchronously registered before the transfer to QDR-II SRAM to reduce the latency to a minimum value. A FIFO can be implemented on the transmit side as needed.
- For memory reads, the read address is presented on USER_AD_RD. A read is indicated when the Read Command signal, USER_R_n, is asserted Low. Resulting 2 x36-bit data is output on USER_QRL and USER_QRH in SDR Mode at the interface frequency of 250 MHz.
- USER_CLK is the clock input for the interface design.
- The DOFFbar, C and Cbar of the QDR-II device are tied HIGH to enable CQ echo clock feature of the QDRII device and is necessary for the operation of the reference design.
- USER_BW_n is the word enable input for the memory.
- USER_RESET is the reset input for the interface design. It has a synchronous mechanism in this design and is an input to the DCM reset.
- The optional RD_STB_n_out is an output of the reference design which is the duplicate of the QDR_R_n Read strobe signal routed over a length of board trace to match the delay to and from the memory device. This can be used as the data valid strobe synchronized with the values returning on the Read path (Q) from the memory. This strobe enters the RD_STB_n_in input to be used as a write enable signal to the Read data FIFOs inside the reference design.
- The USER_QEN_n is an active LOW signal used to indicate the completion of the Read operation.
- USER_WR_FULL indicated the status of the Write FIFOs.
- USER_WR_FULL indicated the status of the Write FIFOs. and the USER_QR_EMPTY indicates the status of the Read FIFOs.
- The CLK_DIV4 and the DLY_CAL_DONE are the signals used to control the IDELAY control blocks in the Xilinx Virtex4. Virtex-4 FPGA uses the Direct clocking method in which any IO can be delayed using the IDELAY instances by 75 ps. The DLY_CLK_200 is the reference clock used to clock these IDELAY instances and the CLK_DIV4 is used to synchronize these blocks with the QDRII design. More information on the IDELAY instances can be found in the documentation on Virtex-4 provided by Xilinx.

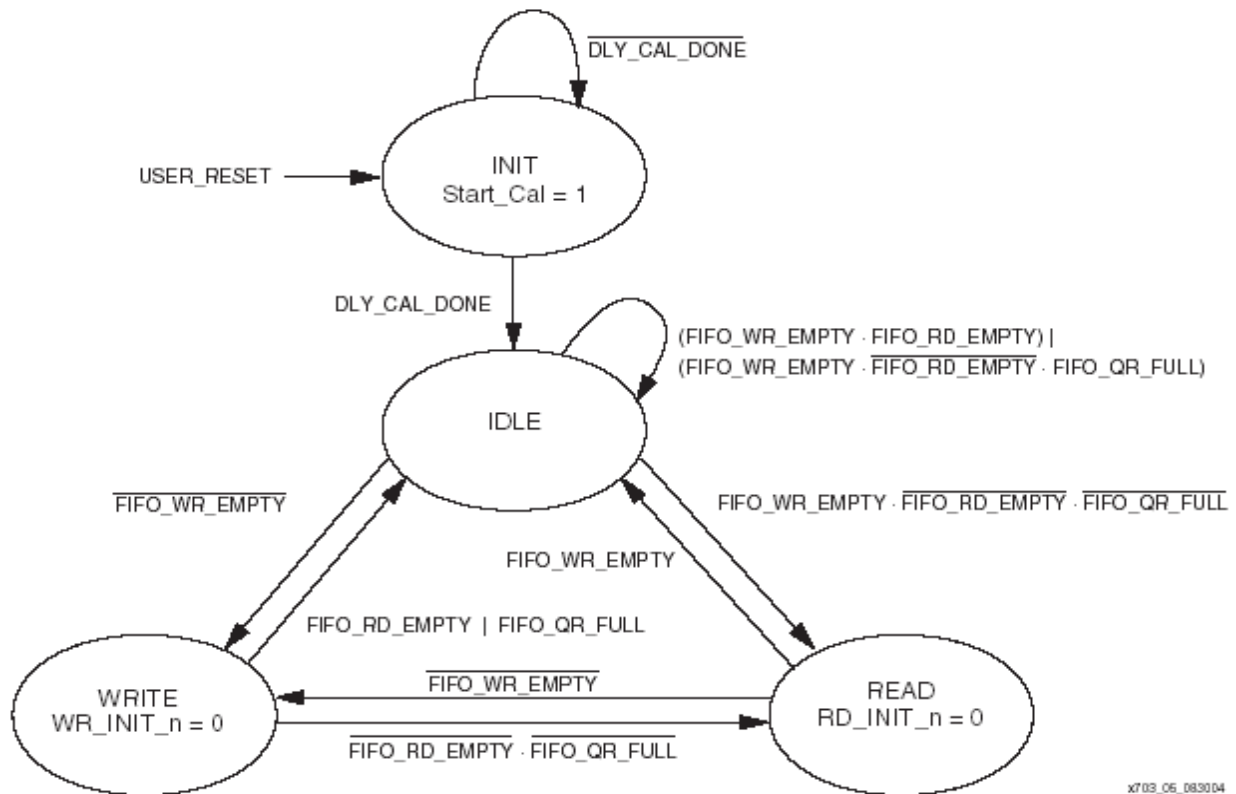


Figure 5. Read/Write Controller State Machine

x703_06_083004

The signals directed to memory are self-explanatory and in compliance with the QDR-II SRAM specification.

Read/Write Controller State Machine

Figure 5 shows the state machine of the controller for read and write accesses to the QDR-II SRAM. The state machine is responsible for coordinating the flow of data between the user interface and the physical interface. It initiates the Read/Write commands to the external memory device based on the requests stored in the user interface FIFOs.

Physical Interface

The physical interface

- transmits and receives all the signals to and from the memories.
- generates the actual I/O signaling and timing relationships for communication of Read/Write commands to the external memory device including the DDR signals.
- provides the necessary timing margins and I/O signaling standards required to meet the overall design performance specifications.

Design Specifications

Table 2 summarizes the specifications for the QDR-II reference design.

Table 2. QDR II Reference Design Specifications

Parameters	Specification/Details
Maximum Frequency	290 MHz (-11 part)
QDR-II Memory Operation	Burst of 4
Bus Width	36 bits
Signal Voltage	1.8V
Memory Device	CY7C1415AV18-250BZC
IO Standard	HSTL I (1.8V Signaling)
HDL Language support	Verilog/VHDL

Virtex-4 DDR I/O Block Implementation

The QDR-II design makes extensive use of the Input DDR (IDDR) and Output DDR (ODDR) primitives found in all Virtex-4 device I/O blocks. These built-in DDR register functions greatly simplify the task of generating the proper clock, address, data, and control signaling for communication to the QDR-II memory device. Both the IDDR and ODDR primitives have various modes of operation to determine how the captured or transmitted DDR data is presented to the FPGA fabric and I/O pins, respectively. More details on the IDDR and ODDR modes of operation are available in the Virtex-4 User Guide.

Clocking Scheme

The QDR-II SRAM standard requires sending data, address, and control signals center-aligned with respect to the forwarded synchronous clock. This requirement is achieved easily by using the clock-forwarding scheme to provide the clock to the FPGA. The clocking scheme in the QDR II design

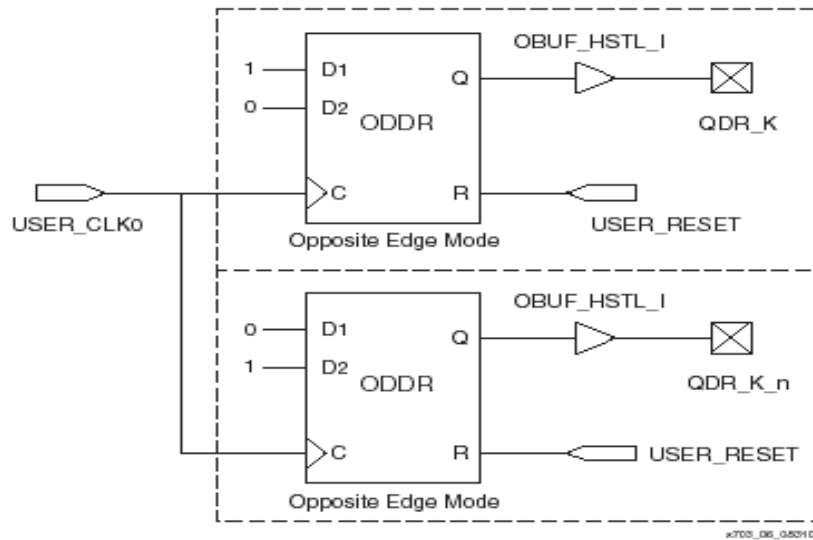


Figure 7. Clock Forwarding Scheme

uses the ODDR registers in the opposite-edge mode to generate the QDR_K and QDR_K_N clocks for the memory device.

This technique removes the clock-to-out parameter of the FPGA because the control, address, data, and clock paths are identical. *Figure 7* presents the current implementation of the clock-forwarding scheme.

The reference design uses the phase-shifted outputs of the DCM to clock the interface on the transmit side. This configuration gives the best jitter and skew characteristics. However, there can be a worst case of 280 ps of skew between the CLK0 output and any other main output of any DCM across the chip. The jitter between the two outputs (CLK0 and CLK270) is 50 ps (worst-case).

There are other possibilities, such as multiplying the clock frequency using one DCM and using another DCM to generate a phase-shifted version. This latter possibility generates more jitter than the solution presented here.

When designing-in Virtex-4 devices, local clock inverters are not recommended when using the DCM. Use instead the CLK0 and CLK270 outputs of the DCM to generate the signals sent to the QDR-II SRAM.

Using Multiple QDR-II SRAMs on the Same Bus

There are three reasons for using multiple QDR-II SRAMs on the same bus:

1. To increase the density of the memory resource.
2. To increase the depth of the memory resource.
3. To divide the speed of the interface by using multiple devices to achieve a given bandwidth.

A design can require using QDR-II SRAMs for one or more of the above reasons. The increase in density increases the bandwidth of the memory resource when considering a fixed clock frequency.

This reference design interfaces to two external QDR-II SRAM devices, but it is possible to adapt it to interface to more devices with only slight modifications to the control

signal processing and the data bus width. QDR SRAM memory vendors have published technical papers on clocking strategies that cover both cases and give guidelines on clocking strategies.

Data Path

Transmit Side: Write Operations

At the memory device I/Os, data must be center-aligned with the clock. For this reason, the output registers for these signals are synchronized to the USER_CLK270 clock as discussed earlier. *Figure 8* demonstrates the use of USER_CLK270 and the ODDR registers to generate the DDR signaling required for the QDR_D Write data path. The ODDR register is configured in same-edge mode allowing both 36-bit data words (FIFO_DWL and FIFO_DWH) to be captured from the FPGA fabric on the same rising edge of USER_CLK270. The FIFO_DWL value is transmitted immediately after this rising edge onto the QDR_D Write data bus, while the FIFO_DWH value is subsequently transmitted out of the ODDR block on the next falling edge of USER_CLK270. This process repeats to generate the 4-word Write data burst. The Read/Write address, byte write enables, and Read/Write control strobes are generated in a similar manner using a single flip-flop within the I/O block to create SDR signals synchronized to USER_CLK270.




Figure 8. Write Data Path Implementation




Figure 9. Direct Clocking implementation for Read Path Data Capture

Receive Side: Read Operations

While Read data capture is fundamentally a more challenging operation than Write data transmission, the direct-clocking methodology of the Virtex-4 QDR-II reference design greatly simplifies this task.

As mentioned previously, each Virtex-4 input pin has a programmable delay element (IDELAY). *Figure 9* shows the use of the IDELAY primitives to implement the direct-clocking methodology of Read data capture.

As mentioned earlier, this methodology relies on the use of the CQ echo clock from the QDR-II memory device. This clock signal is used as a “training” signal to center align the QDR Read data with respect to the FPGA system clock USER_CLK0. The CQ clock enters the FPGA through an identical path to the QDR_Q data bus signals, an HSTL input buffer followed by the IDELAY block, followed by an IDDR register. All IDELAY blocks are configured to variable delay mode, this allows the tap delay setting to be dynamically adjusted. More information on the IDELAY Primitives can be obtained from the Xilinx QDR-II memory interface application note, a link to which is given in the references at the end of this document.

The delay calibration state machine monitors the state of the CQ clock input captured by the IDDR register on the rising edge of USER_CLK0. An edge detection algorithm finds the location of the rising and falling edges of the CQ clock by varying the IDELAY tap delay setting for this signal. Once the edges are found, the tap delay setting is adjusted to center the CQ clock edges around the rising edge of USER_CLK0. The CQ clock is edge-aligned to the QDR_Q data bus coming from the memory device. When the same tap delay setting is applied to this bus, the USER_CLK0 signal is automatically centered in the data-valid window of the incoming Read data words. In this manner, the Read data values can be captured directly into the FPGA system clock domain without using complex data recapture techniques or the advanced timing analysis typically required when crossing clock boundaries.

Figure 10 shows how the CQ clock and QDR_Q signals are delayed through the IDELAY blocks with identical tap settings to center align these signals to USER_CLK0. The qdr_cq_delay and qdr_q_delay signals represent the

waveforms at the inputs to the IDDR registers after passing through the IDELAY elements.

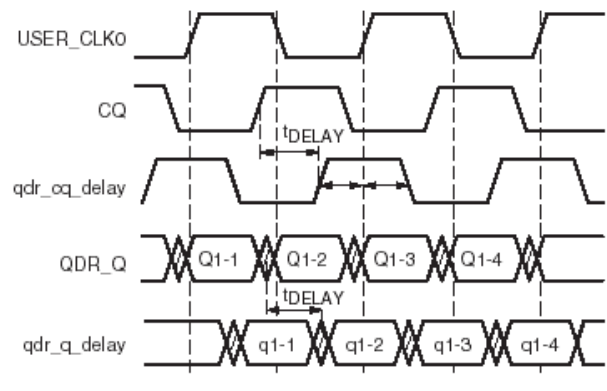


Figure 10. Alignment of QDR_Q inputs to USER_CLK0 using Tap delays

Timing Analysis

The Virtex-4 QDR-II reference design leverages the unique I/O and clocking features of the device to maximize performance and timing margins, while greatly reducing the need for detailed placement and pinout analysis.

This section presents an example timing analysis for the address/control paths, the Write data path, and the Read (or capture) data path.

Clocks, Address, and Control Signals

The FPGA generates all the clock and control signals for reads and writes to memory. The memory clocks typically are generated using a Double Data Rate (DDR) register. A Digital Clock Manager (DCM) generates the clock and its inverted version. Generating the clock this way has several advantages:

- The data, control, and clock signals all go through similar delay elements while exiting the FPGA.

- Clock duty cycle distortion is minimal when global clock nets are used for the clock and the 180° phase-shifted clock.

All address and control signals are registered and output at the IOB. The address and control signals are registered using a clock that is 270° shifted from the clock signal to the memory. Such an approach enables the address and control signals to have additional time margin before they are registered. The address and control signals easily meet the required timing.

Table 3 shows an example timing analysis for a DDR-1 interface implemented using an CY7C1415AV18-250BZC FPGA, – 11 speed grade.

Figure 11 illustrates the address and control signal timing margins. Because these signals are referenced to USER_CLK270, there is more trailing edge margin than leading edge margin with respect to the QDR_K clock edge. This allows use of fewer global clock buffers and still provides adequate margin on the leading edge.

Data Writes

The Write data path (QDR_D) is also synchronized to USER_CLK270. However, the Write data words are transmitted as DDR values, and therefore must have adequate

set-up and hold margins with respect to both the rising edge of QDR_K and the rising edge of QDR_K_n. Accordingly, the timing analysis for the Write data path shown in Table 4 incorporates the maximum duty cycle distortion of the memory clocks. This analysis is also for a 250-MHz QDR II memory device and a Virtex-4 device, -11 speed grade.




Figure 11. Address and Control Signal Timing Margins

Table 3. Address and Control Signal Timing Analysis

Parameter	Value (ps)	Leading-Edge Uncertainties	Trailing-Edge Uncertainties	Meaning
T _{CLOCK}	4000			Clock period (250-MHz clock frequency)
T _{CLOCK_SKEW_FPGA}	±50	50	50	Minimal skew since right/left sides are being used and the bits are close together. The number from TRACE for the specific clock area used to generate the address and control signals.
T _{PACKAGE_SKEW}	±30	30	30	Using the same bank reduces package skew
T _{SETUP}	500	500	0	Set-up time from memory data sheet
T _{HOLD}	500	0	500	Hold time from memory data sheet
T _{PCB_LAYOUT_SKEW}	±50	50	50	Skew between layout lines on the board. Because board skew varies per design, values listed here are examples
T _{PHASE_OFFSET_ERROR_DCM}	±140	140	140	Offset between different phases of the DCM outputs
T _{JITTER}	±50	50	50	Jitter component associated with the difference between USER_CLK0 and USER_CLK270.
Total Uncertainties	-	820	820	
Command Window	2360	820	3180	Worst-case window of 2360 ps

Table 4. Write Data Path Timing Analysis

Parameter	Value (ps)	Leading-Edge Uncertainties	Trailing-Edge Uncertainties	Meaning
T _{CLOCK}	4000			Clock period (250-MHz clock frequency)
T _{CLOCK_PHASE}	2000			Clock phase (50% of clock period)
T _{DCD}	200			Overall duty cycle distortion of clock to memory (5% of clock period)
T _{DATA_PERIOD}	1850			Total data period, T _{CLOCK_PHASE} – T _{DCD}
T _{CLOCK_SKEW}	±50	50	50	Minimal skew since right/left sides are being used and the bits are close together. The number from TRACE for the specific clock area used to generate the address and control signals.
T _{PACKAGE_SKEW}	±30	30	30	Skew due to package pins and board layout. This can be reduced further with tighter layout
T _{SETUP}	350	350	0	Set-up time from memory data sheet
T _{HOLD}	350	0	350	Hold time from memory data sheet
T _{PCB_LAYOUT_SKEW}	±50	50	50	Skew between layout lines on the board. Because board skew varies per design, values listed here are examples
T _{PHASE_OFFSET_ERROR}	±140	140	140	Offset error between different clocks from the same DCM
T _{JITTER}	±50	50	50	Jitter component associated with the difference between USER_CLK0 and USER_CLK270.
Total Uncertainties		670	670	Worst case for leading and trailing can never happen simultaneously
Window	460	670	1130	Total worst-case window is 460 ps

Figure 12 illustrates the Write data path timing margins. Only the analysis with respect to QDR_K is shown. The analysis with respect to QDR_K_n is identical.

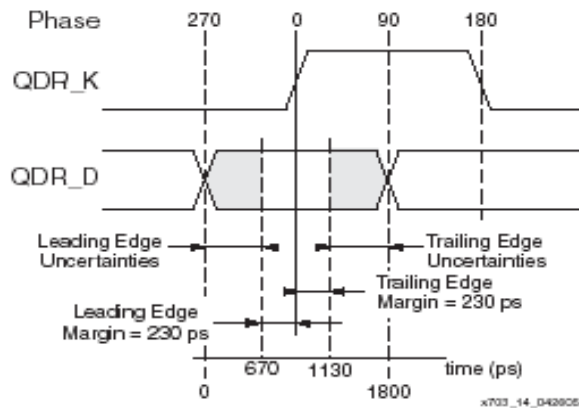


Figure 12. Write Data Path Timing Margins

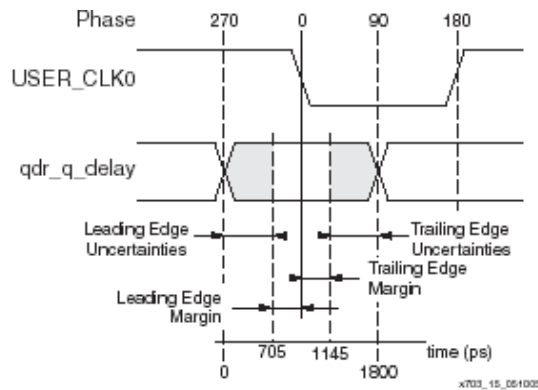
Data Captures

The Read data path (QDR_Q) values are captured directly into the USER_CLK0 clock domain using the previously described direct clocking techniques. Thus, the data capture timing analysis must be performed with respect to USER_CLK0 and consideration must be given to the IDELAY tap delay resolution. In addition, though the CQ echo clock from the memory is used only as a “training” signal for the edge detection algorithm that center-aligns the QDR_Q bus to USER_CLK0, the potential skew between the CQ clock and the QDR_Q bus must also be taken into account. Table 5 presents the timing analysis for Read data capture.

Table 5. Read Data Path Timing Analysis

Parameter	Value (ps)	Leading-Edge Uncertainties	Trailing-Edge Uncertainties	Meaning
T_{CLOCK}	4000			Clock period
$T_{\text{CLOCK_PHASE}}$	2000			Clock phase (50% of clock period)
$T_{\text{MEM_DCD}}$	200			Overall duty cycle distortion on receive clocks
$T_{\text{DATA_PERIOD}}$	1850			Total data period, $T_{\text{PHASE}} - T_{\text{MEM_DCD}}$
$T_{\text{CQ_TO_Q_SKEW}}$	± 300	300	300	CQ-to-data distortion from memory data sheet
$T_{\text{PACKAGE_SKEW}}$	± 50	50	50	This parameter depends on the exact package. Because the eight data bits are close together, skew is less than values listed
T_{SETUP}	100	100	0	Setup time from Virtex-4 data sheet for -11 part (T_{DICK})
T_{HOLD}	50	0	-50	Hold time from Virtex-4 data sheet for -11 part (T_{CKDI})
T_{JITTER}	± 100	100	100	Maximum USER_CLK0 jitter
$T_{\text{TAP_DELAY_RES}}$	± 75	75	75	TAP resolution of IDELAY element
$T_{\text{PCB_LAYOUT_SKEW}}$	± 50	50	50	Skew between data lines on the board
Total Uncertainties	-	705	655	
Valid Window	440	705	1145	Worst-case window is 440 ps

Figure 13 illustrates the Read data path timing margins.


Figure 13. Read Data Path Timing Margins

Conclusion

This application note describes the implementation and timing details of a four-word-burst QDR-II SRAM interface for Virtex-4 devices. The direct-clocking methodology utilized greatly simplifies the task of read data capture within the FPGA while providing a high-performance, robust, and scalable memory interface solution for current and next generation QDR II SRAM memory devices.

The total throughput for this interface is 500 Mb/s per pin for the 36-bit write bus and 36-bit read bus at 250 MHz.

Appendix A: Definition of Terms

The table below defines the mnemonics used in this application note.

Abbreviation	Description
CLB	Configurable Logic Block
DAR	Double Address Rate
DCM	Digital Clock Manager
DDR	Double Data Rate
DLL	Delay-Locked Loop
FDDR	Double Data Rate Flip-Flop
FIFO	First In, First Out
GND	Ground
HSTL	High-Speed Transceiver Logic
IOB	Input/Output Block
JEDEC	Joint Electron Device Engineering Council
JTAG	Joint Test Action Group
LUT	Look-Up Table
PCB	Printed Circuit Board
QDR SRAM	Quad Data Rate Synchronous Static Random Access Memory
SAR	Single Address Rate
SDR	Single Data Rate
SSO	Simultaneous Switching Outputs
TAP	Test Access Port

Appendix B: References

Most of the information provided in this document is obtained from the Xilinx application note “QDR-II SRAM Interface for Virtex-4 Devices”, the link to which is

<http://www.xilinx.com/bvdocs/appnotes/xapp703.pdf>

Other references include:

- QDR consortium website
— <http://www.qdrsram.com/>
- Xilinx Website
— http://www.xilinx.com/products/design_resources/mem_corner/resource/xaw_sram_qdr.htm

QDR RAMs and Quad Data Rate RAMs comprise a new family of products developed by Cypress Semiconductor, Renesas, IDT, NEC, and Samsung. Virtex-4 is a trademark of Xilinx, inc. All product and company names mentioned in this document are the trademarks of their respective holders.

AN5065 approved kkvtmp 8/10/05