



WP393 (v1.0) July 14, 2011

I/O and Memory Interfacing Features and Benefits in 7 Series Architecture

By: Matt Klein

The input and output (I/O) functionality in any FPGA needs to be designed to address a wide array of applications of varied performance and function. The I/O in the 7 series FPGAs and Zynq™ Embedded Processing Platforms (EPPs) from Xilinx have been optimized for a diverse set of requirements at both the physical and logical level, including high-speed memory, networking, video flat panel and sensor interfaces, high-speed ADC/DACs connectivity as well as legacy interfaces. Additionally, new hard blocks have been added specifically to benefit memory interfacing for high-speed DDR3 devices.

This white paper describes how the new I/O structures in the 7 series architecture support the range of performance and functionality challenges needed to address the broad range of application needs.

Introduction

Having to target a different I/O structure when migrating from one device to another can sometimes prove challenging for the user. 7 series devices use Xilinx's unified I/O architecture, providing identical functionality across all Artix™-7, Kintex™-7, and Virtex®-7 FPGAs, and Zynq EPPs. The unified architecture allows customers to reduce investment in developing and deploying products, and it simplifies migration of existing designs to 7 series FPGAs. It should be noted that when the Zynq EPP is referenced in this white paper, it is specific to the FPGA portion of the Zynq EPP.

I/O power efficiency is another breakthrough 7 series advantage, enabling better integration than in previous products. A range of customer-accessible and automatic power reduction features have been added to the I/O in the 7 series architecture. These features minimize the contribution of I/O power to total power consumption and fit into applications requiring tight power budgets.

The I/O in the 7 series FPGAs are designed to deliver the highest possible performance, including single-ended performance for DDR3 up to 1,866 Mb/s and differential LVDS up to 1,600 Mb/s. However, performance is not the only important attribute of the I/O. To support various applications, I/O needs to be flexible and offer a wide range of performance with different I/O standards, allowing applications to target FPGAs of the same family, between families of the 7 series FPGAs, and the Zynq EPPs.

Addressing I/O Challenges in 7 Series FPGAs and Zynq EPPs

The physical I/O capabilities and structures provide a range of I/O standards, terminations, and power-saving features. The number of I/O per I/O bank and how they are placed relative to clocking and new I/O resources is important as well as their arrangement within the FPGA die. Additionally, detailed I/O bound logical functions, such as input/output delays and serialization and deserialization functions, are key to allowing the broadest application support for the I/O. Lastly, the addition of new functional structures, such as the Phaser, I/O phase-locked loops (PLLs), and I/O FIFOs complete the interface feature, supporting the highest performance DDR3 and other memory interfaces. The basic I/O structure and new I/O related blocks are shown in [Figure 1](#).

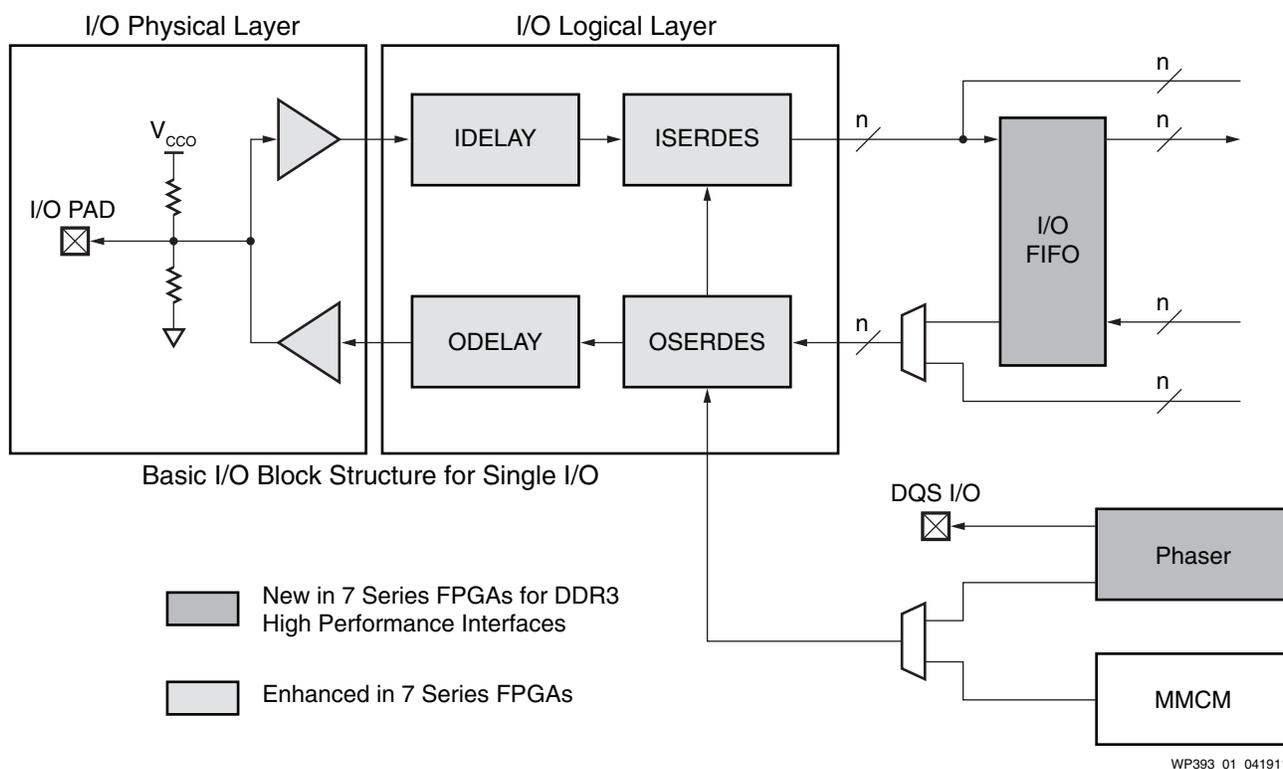


Figure 1: 7 Series FPGAs Basic I/O Blocks and Connectivity

I/O at the Physical Level

At the physical level, I/Os are required to support a range of drive voltages (or levels) and drive strengths and receive capabilities to interface to various I/O standards (e.g., compatible with PCI33/66 and fully compliant with LVCMOS, LVTTTL, LVDS, SSTL, and HSTL). I/Os also support various input and output termination features, which can be dynamically asserted and removed. Also, the I/O of the 7 series FPGAs and Zynq EPPs have features available that support dynamic, static, and user-controlled power reduction strategies.

The 7 series architecture has two types of I/O:

- High-performance I/O, arranged in banks called HP I/O banks
- High-range I/O, arranged in banks called HR I/O banks

Both I/O types are based on Virtex-6 architecture but have enhanced functionality and voltage range support. The two I/O types are bound to an entire I/O bank of 50 I/O. Some devices have all HP I/O banks, some have all HR I/O banks, and some have a combination of the two. All 50 I/O of a given bank might not be brought out to the balls of the FPGA depending on the part and package combination. The Artix-7, Kintex-7, Virtex-7 FPGAs, and the Zynq EPP product tables list how many of each type of I/O are present on a part/package combination basis:

<http://www.xilinx.com/technology/roadmap/7-series-fpgas.htm>

http://www.xilinx.com/publications/prod_mktg/zynq7000/Zynq-7000-hardware-designer-product-table.pdf

HP I/O and HP I/O Banks

The HP I/O banks are optimized for the highest performance applications to address DDR3 interfacing up to 1,866 Mb/s and other chip-to-chip interfaces, including LVDS interfaces at up to 1,600 Mb/s. The I/O in the HP I/O bank type are compatible with interface standards of up to 1.8V to work with the most demanding, highest performance interfaces.

The HR I/O banks are designed to support a wider range of I/O standards, with voltages up to 3.3V. The HR I/O banks are optimized for the broadest range of application coverage while still achieving DDR3 interface speeds up to 1,066 Mb/s. They also address LVDS interfaces at up to 1,055 Mb/s. The I/O in this bank type are compatible with modern and legacy interfaces.

The 7 series FPGAs contain both HR and HP I/O banks in several combinations:

- The Artix-7 devices only have 3.3V HR I/O banks.
- The Kintex-7 devices have both 3.3V HR I/O banks and 1.8V HP I/O banks; some devices have more HR banks than HP banks.
- The Virtex-7 FPGAs have both 3.3V HR and 1.8V HP I/O banks in some family members; however, there are more 1.8V HP I/O banks.

[Table 1](#) highlights the features supported in the HP and HR I/O banks. See the specific device family data sheet for details on the performance and other electrical requirements of the HP and HR I/O banks.

Table 1: Supported Features in the HR and HP I/O Banks

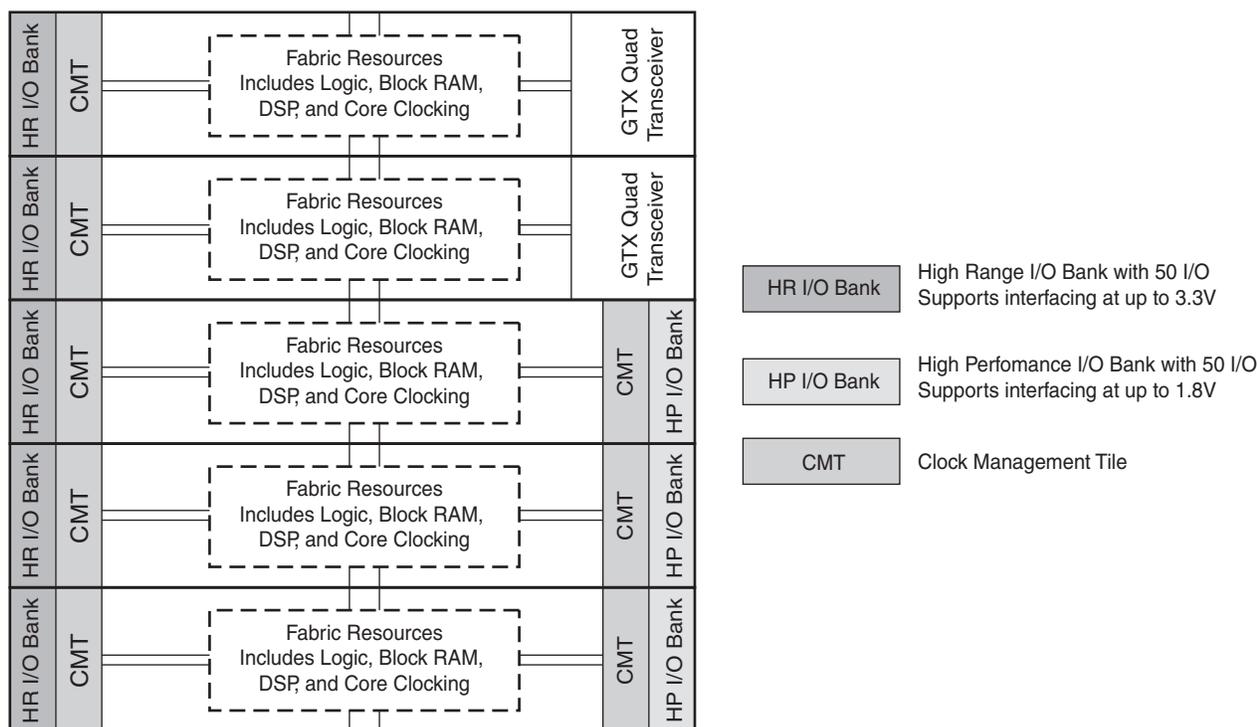
| Feature | HP I/O Banks | HR I/O Banks |
|--------------------------------------|--------------------|--------------|
| 3.3V I/O standards ⁽¹⁾ | N/A | Supported |
| 2.5V I/O standards ⁽¹⁾ | N/A ⁽²⁾ | Supported |
| 1.8V I/O standards ⁽¹⁾ | Supported | Supported |
| 1.5V I/O standards ⁽¹⁾ | Supported | Supported |
| 1.35V I/O standards ⁽¹⁾ | Supported | Supported |
| 1.2V I/O standards ⁽¹⁾ | Supported | Supported |
| VCCAUX_IO supply rail | Supported | N/A |
| Digitally-controlled impedance (DCI) | Supported | N/A |
| Untuned on-die termination | N/A | Supported |
| IDELAY | Supported | Supported |
| ODELAY | Supported | N/A |
| ISERDES | Supported | Supported |
| OSERDES | Supported | Supported |

Notes:

1. Not all I/O standards and drive strengths are supported in both the HP and HR I/O banks. See [UG471, 7 Series FPGAs SelectIO Resources User Guide](#) for the specific I/O standards that are available in the HP and HR I/O banks.
2. Although LVDS is generally considered a 2.5V I/O standard, it is supported in both the HR and HP I/O banks.

[Figure 2](#) shows the Kintex-7 XC7K160T FPGA layout with HR I/O banks, HP I/O banks, and clock management tiles (CMTs). CMTs have historically been used in

several generations of Xilinx FPGAs but have been enhanced in their functionality in relation to memory interfaces (see [CMTs and Phaser Blocks](#)).



WP393_02_032911

Figure 2: I/O Bank and CMT Layout in Kintex-7 FPGA

Power Reduction Features in the I/O

The I/O power for memory interfaces has three main components:

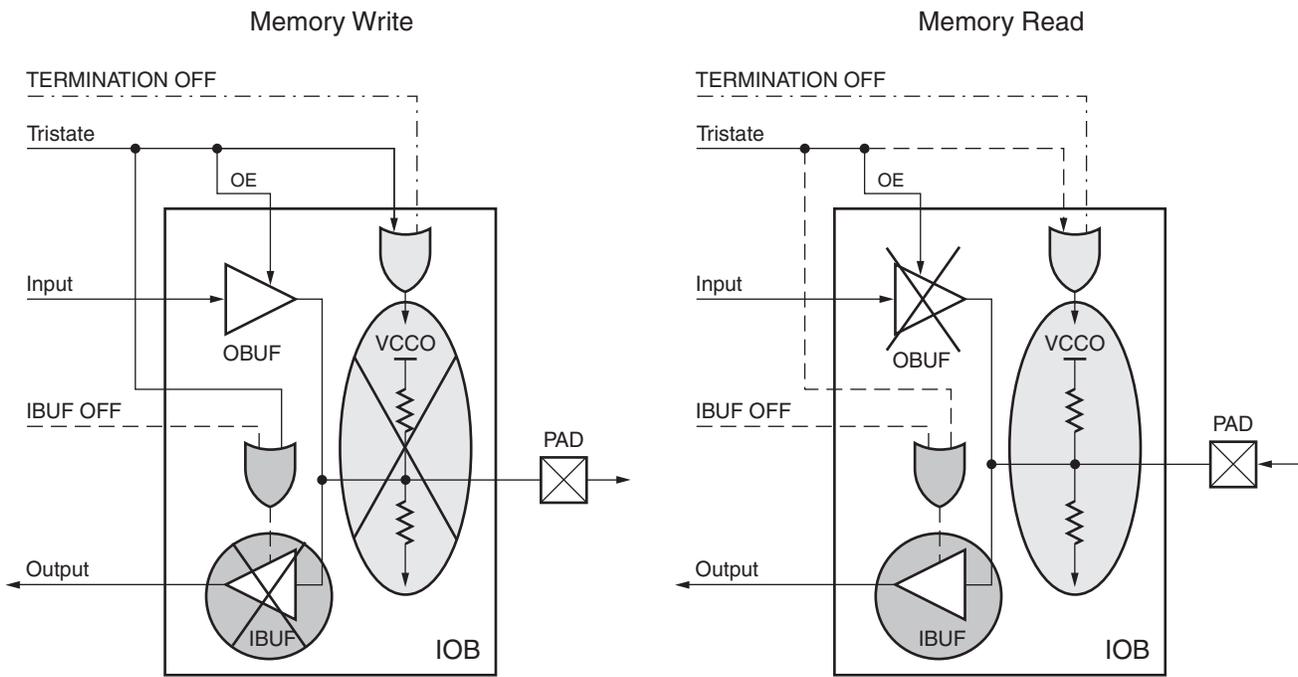
- DCI: Used for matching the impedance of the PCB trace
- Referenced-input receiver: Used to adjust the I/O voltage to the core voltage
- IDELAY: Used to sync the signal to a clock

To reduce power consumption in Virtex-6 FPGAs, the 3-state DCI automatically turned off the termination during memory writes, saving 50% termination power. Virtex-6 FPGAs also offer a referenced receiver in low power modes and IDELAY that save 70% and 50% respectively compared to the high-performance mode. These features save over 50% of the power consumed in the equivalent interface for Virtex-5 FPGAs.

In 7 series FPGAs, Xilinx is building on the Virtex-6 architecture and is fine-tuning each feature to save the most power. On the design side, lowering V_{CCAUX} from 2.5V to 1.8V saves 30% on all items powered by V_{CCAUX} , specifically the IDELAY and the input and output buffers.

In addition, a new feature for the 7 series FPGAs is the ability to dynamically disable the input buffer. This is an improvement and an extension of the dynamic 3-state DCI circuitry found in previous generations. The 3-state DCI feature is derived from the principle of turning off or disabling a function when it is not in use, which is exactly what happens to the termination when writing to the memory. However, in previous generations, the input buffer was still burning power during an output or memory write. In 7 series FPGAs, the input buffer can be disabled during a memory write

(output). This saves 50% of the power based on a 50% write/read balance. See Figure 3.



WP389_18_021011

Figure 3: Abstracted Input Buffer Disable and DCI Termination Disable

With these features, much of the power is saved during a memory write. However, in some cases, the memory interface is not reading or writing. These states are considered idle states and, in the past, both the DCI termination and input buffer would burn power in this state. In 7 series FPGAs, a new feature has been added to disable either the DCI or the input buffer, or both. This is a BUS IDLE state. Additionally, a number of other power saving features that are present in the I/O of the Virtex-6 FPGAs have been extended to the 7 series FPGAs. These and many other power saving features in the 7 series FPGAs are described in [WP389, Lowering Power at 28 nm with Xilinx 7 Series FPGAs](#).

I/O at the Logical Level

All inputs and outputs can be configured as either combinatorial or registered. Double data rate (DDR) is supported by all inputs and outputs. Any input and some outputs can be individually delayed by 0 to 31 sets of 78 ps or 52 ps, depending of the frequency of a reference clock. Such delays are implemented as programmable tapped delay lines called IDELAY and ODELAY. The tap delay resolution is varied by selecting an IDELAYCTRL reference clock from the range specified in the 7 series FPGAs and Zynq EPP data sheets at http://www.xilinx.com/support/documentation/7_series.htm. The number of delay steps can be set by configuration and can also be incremented or decremented while in use on a per I/O basis.

Every I/O block contains a programmable absolute delay primitive called IDELAYE2. The IDELAY can be connected to an ILOGICE2/ISERDESE2 or ILOGICE3/ISERDESE2 block. IDELAYE2 is a 31-tap, wraparound, delay primitive with a calibrated tap resolution. It can be applied to the combinatorial and registered

input paths. It can also be accessed directly in the FPGA logic. IDELAY allows incoming signals to be delayed on an individual basis.

Every HP I/O block contains a programmable absolute delay primitive called ODELAYE2. Note: This function is not available in the HR I/O blocks. The ODELAY can be connected to an OLOGICE2/OSERDESE2 block. ODELAY is a 31-tap, wraparound, delay primitive with a calibrated tap resolution. It can be applied to the combinatorial input path, registered input path, combinatorial output path, or registered output path. It can also be accessed directly in the FPGA logic. ODELAY allows outgoing signals to be delayed on an individual basis.

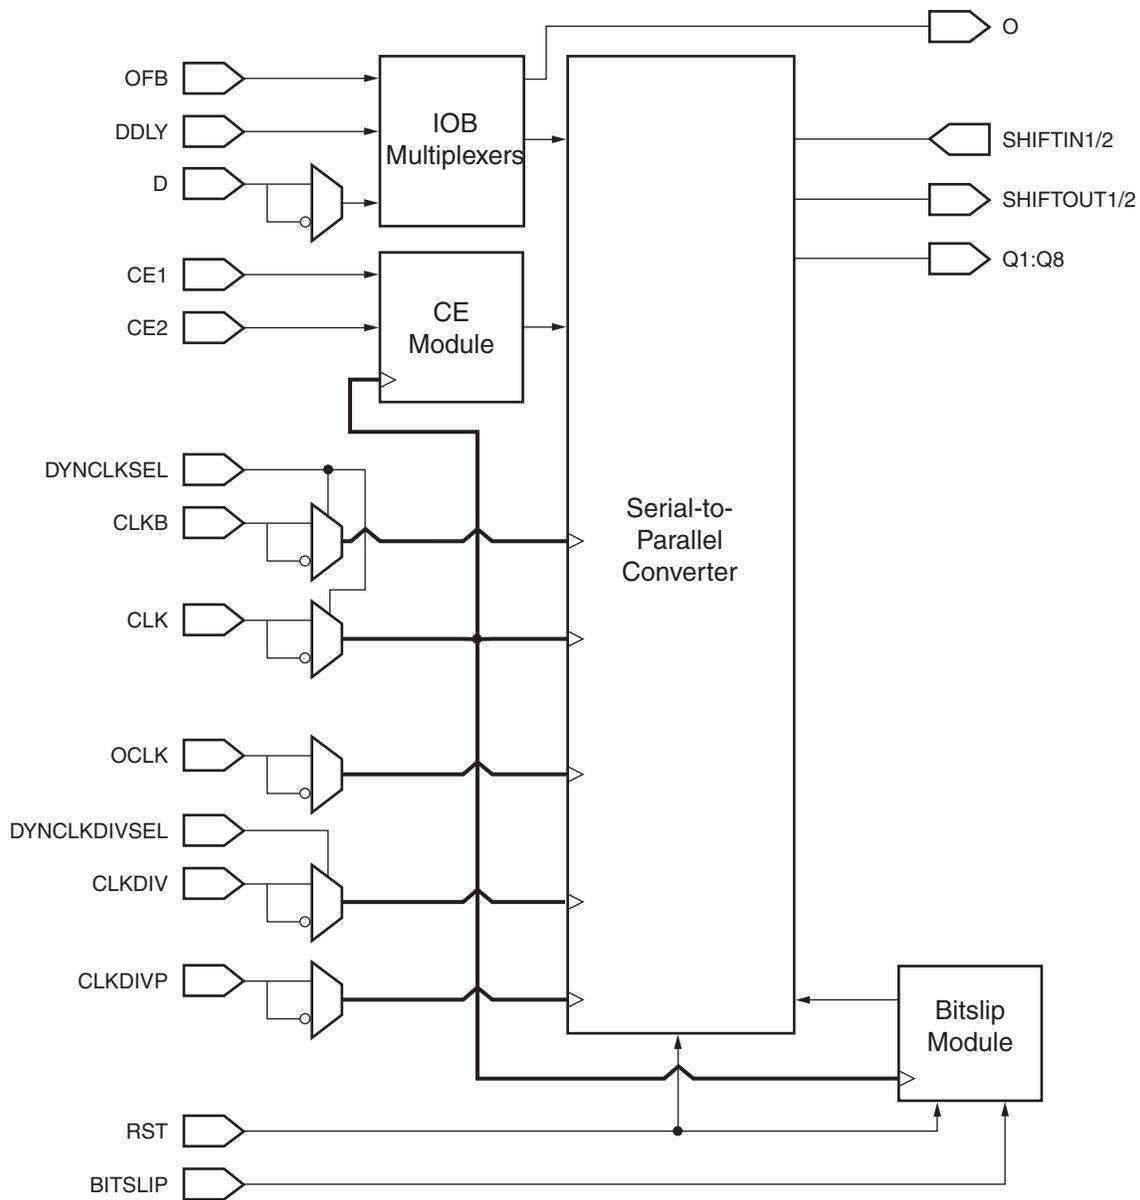
ISERDES and OSERDES

Many applications combine high-speed, bit-serial I/O with slower parallel operation inside the device. This requires a serializer and deserializer (SerDes) inside the I/O structure. Each I/O pin contains an 8-bit IOSERDES (called ISERDES and OSERDES in the 7 series FPGAs and Zynq EPPs) capable of performing serial-to-parallel or parallel-to-serial conversions with programmable widths of 2, 3, 4, 5, 6, 7, or 8 bits in single data rate (SDR) mode. In DDR mode, widths of 4, 6, and 8 bits are supported, and by cascading two IOSERDES from two adjacent pins (default from differential I/O), wider width conversions of 10 and 14 bits are supported.

From an application point of view, the ISERDES has a special oversampling mode capable of asynchronous data recovery for applications like a 1.25 Gb/s LVDS I/O-based SGMII interface. Source-synchronous and system-synchronous interfaces are implemented using a range of the rich features in the ISERDES, OSERDES, and Input and Output Delay blocks. Specialized clock connectivity into the I/O has been added to the 7 series architecture to further enhance these functions.

Another function supported by the ISERDES in the 7 series FPGAs and Zynq EPPs is called Bitflip. The Bitflip function allows designers to reorder the sequence of the parallel data stream going into the FPGA fabric. This can be used for training source-synchronous interfaces that include a training pattern. The ISERDES also supports strobe-based memory interfaces through dedicated circuitry (including the OCLK input pin) to handle the strobe-to-FPGA clock domain crossover entirely within the ISERDES block. This allows for higher performance and a simplified implementation. In addition, the ISERDES supports modes for networking interfaces, DDR3 and other strobe-based memory interfaces, QDR memory interfaces as well as oversampling for asynchronous interfaces.

The exact implementation of the ISERDES in the 7 series FPGAs is ISERDESE2, which distinguishes it from older versions in previous generation Xilinx® FPGAs. See [Figure 4](#).



WP393_04_042111

Figure 4: ISERDES2 Block Diagram

The converse of the ISERDES2, the OSERDESE2 is a dedicated parallel-to-serial converter with specific clocking and logic resources designed to facilitate the implementation of high-speed source-synchronous interfaces.

Typically, the serialization and deserialization features are used for LVDS interfaces for SPI4.1 and SPI4.2. The wider serialization and deserializations, such as 10:1 and 14:1, are used for DACs, ADCs, and Flat Panel Display applications, also using LVDS, which takes advantage of cascading a master and slave ISERDES pair or OSERDES pair.

Every OSERDESE2 module includes a dedicated serializer for data and 3-state control. Both data and 3-state serializers can be configured in SDR and DDR mode. Data serialization can be up to 8:1 (10:1 and 14:1 if using OSERDESE2 width expansion). 3-state serialization can be up to 14:1. A dedicated DDR3 mode supports high-speed

memory applications. [Figure 5](#) shows a block diagram of the OSERDESE2, highlighting all the major components and features of the block.

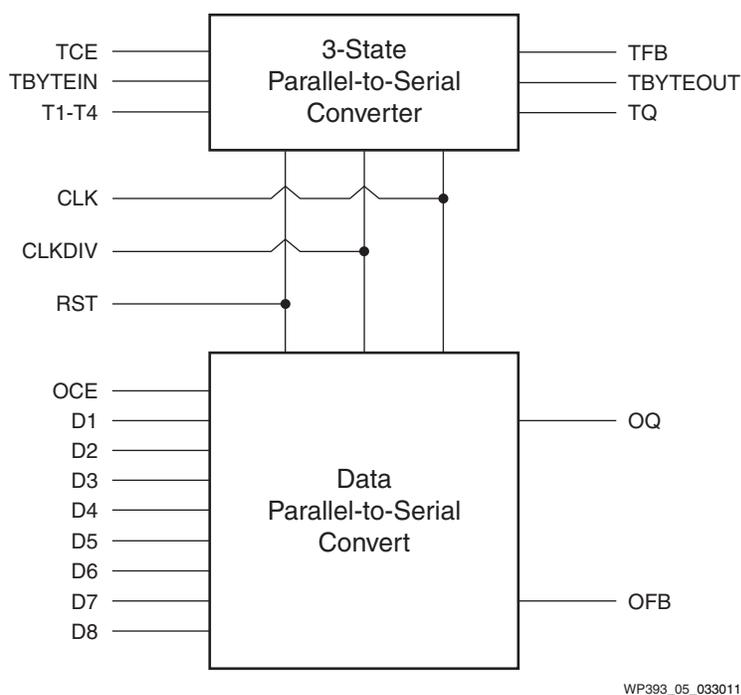


Figure 5: OSERDES2 Block Diagram

Detailed information on these blocks can be found in [UG471, 7 Series FPGAs SelectIO Resources User Guide](#).

Some of the many interfaces supported by the I/O in the HP I/O banks and HR I/O banks of the 7 series FPGAs and Zynq EPPs are shown in [Table 2](#).

Table 2: Example I/O Interfaces Enabled in 7 Series FPGAs

| Application | Single-Ended or Differential | HP I/O Banks | HR I/O Banks |
|---|-------------------------------|---|--|
| Memory Interfacing | Single-ended | Highest performance DDR2, DDR3, QDR2+, RLD RAM2, RLD RAM3, and LPDDR2 | High-performance DDR2, DDR3, QDR2+, RLD RAM2, RLD RAM3, and LPDDR2 |
| Networking Interfaces | Differential | LVDS for SPI4.2 and SGMII interfaces | LVDS for SPI4.2, but lower performance |
| | Single-ended | N/A | RGMII |
| Special High-Speed Interfaces | Differential | ASIC/ASSP up to 1,600 Mb/s | N/A |
| Consumer Video (e.g., flat panel and sensor interfaces) | Differential | Standard LVDS to 1,600 Mb/s | LVDS, sub-LVDS, TMDS, Mini-LVDS, and PPDS up to 1,055 Mb/s |
| DAC/ADC | Single-ended and Differential | High-performance DAC/ADC | General-purpose DAC/ADC |
| General-Purpose I/O | Single-ended | High-speed processor I/O and high-speed flash memory like Dual and Quad SPI | Legacy processors and flash memory using LVCMOS up to 3.3V |
| Legacy Flash Memory | Single-ended | N/A | NOR Flash and ZBT SRAM at 3.3V/2.5V |
| PCI33/66 | Single-ended | No | Yes |

High-Performance Memory Interfacing in 7 Series FPGAs and Zynq EPPs

The 7 series FPGAs and Zynq EPPs have several new blocks specifically designed for simplified interfacing to even the highest performance memory subsystems.

Design Challenges for Memory Interfaces

Several challenges exist with high-performance memory interfaces:

- Placing a capture clock edge in the middle of the data eye at the capture flip-flops (FFs) at 1,866 Mb/s
- Maintaining capture clock edges vs. the data eye across Process, Voltage, and Temperature (PVT)
- Synchronizing received I/O data with the rest of the FPGA system (i.e., transferring between the PHY capture domain and a fixed fabric clocking domain)
- Providing write clocks shifted 90° from the data clock in the memory interface

The first two challenges are especially difficult because the DQS clock used for each byte of memory is not continuous and has different timing for each byte as well as for each rank of a memory interface.

Many I/O interfaces require fast, jitter-free, phase-aligned clocks. In architectures prior to 7 series FPGAs, all clock resources must be shared between serial transceivers and memory interfaces. Resources like the Mixed Mode Clock Manager (MMCM) have features that are needed by the FPGA fabric but are not required when implementing a memory interface (e.g., DDR3). For example, the memory interface needs to respond to DQS edges that are not continuous and have different timing of a per-byte lane basis. Also, the PHY and FPGA fabric have arbitrary and variable phases of clocks. The 7 series FPGAs and Zynq EPPs have decoupled the PHY timing from the FPGA fabric timing for memory interfaces.

The 7 series FPGAs and Zynq EPPs have new hard blocks, which mitigate these challenges. These blocks are the Phasers, the I/O FIFOs, and the I/O PLL, which are all contained within or are adjacent to the enhanced CMTs that encompass the traditional MMCM as in the Virtex-6 FPGAs.

Advantages over Other Methods

The new hard blocks (Phasers, I/O FIFOs, and I/O PLLs) in the 7 series FPGAs and Zynq EPPs provide and respond to a higher resolution of clock timing control. These blocks and the ISERDES and OSERDES also respond to higher input frequencies (up to 933 MHz for DDR3 at 1,866 Mb/s) and allow finer phase shift steps than older methods in Virtex-6 FPGAs.

The new hardware blocks used for high-performance memory interfaces contain PVT compensated delay lines based on a regulated bias. For memory reads, the delay line has a calibrated initial delay relative to a DQS input and phase detection for dynamic adjustment. Further, it can achieve a much longer delay length over two different stages and is much cleaner based on dedicated routing and a regulated bias.

Memory interfaces generated by the Memory Interface Generator (MIG) tool can consist of up to 72 bits, at which point, they span multiple I/O banks. Multiple banks of external memory and multiple clock regions (across multiple I/O banks) are supported through new clocking resources like direct connections to and from dedicated DQS I/O pins and the Phaser blocks for memory interfacing. They are also

supported by less dedicated BUFIO and BUFR (used in previous generation Virtex-6 FPGAs) and are connected to MMCMs for targeting all I/O clocking requirements.

Lastly, skew is much better controlled through direct connections and dedicated routing of the high speed clocking resources.

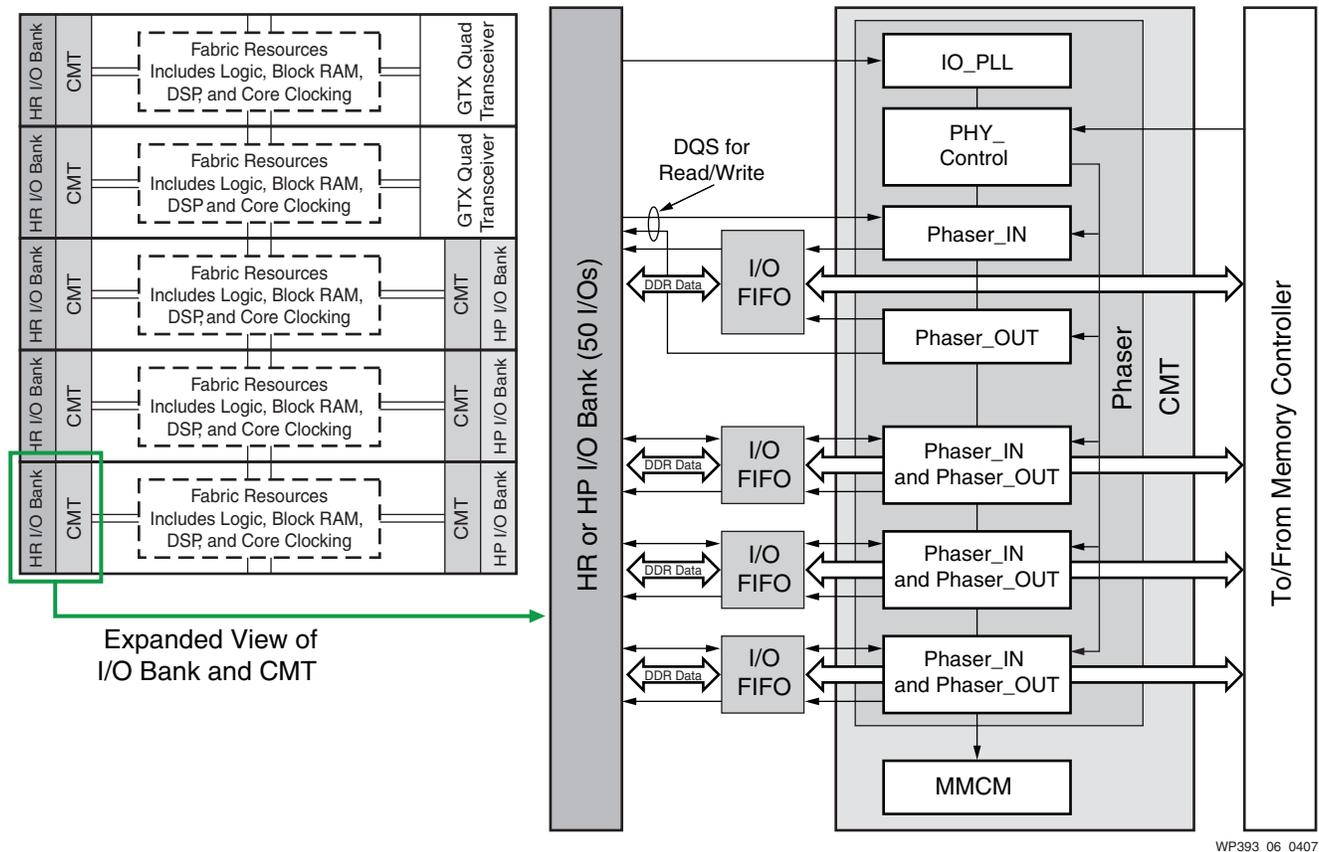
CMTs and Phaser Blocks

The CMTs in the 7 series architecture are dramatically enhanced. CMTs in previous generation Xilinx FPGAs only consisted of MMCMs, which are multi-function PLLs, and were placed throughout the core.

In 7 series FPGAs and Zynq EPPs, in addition to the MMCM, each CMT now contains the new Phaser block and a new additional PLL, especially for high performance memory interfacing. The Phaser block itself consists of several pieces, including a very high-performance I/O PLL, a Phaser Control block, four Phaser_IN blocks, and four Phaser_OUT blocks. The Phaser_IN blocks respond to DQS inputs and produce clocks for the input FIFOs and ISERDES blocks to synchronize memory reads to the DQS. The Phaser_OUT block produces memory write timing. It also provides clocks for the output FIFO and OSERDES blocks and takes care of keeping clock and data 90° phase shifted.

CMTs in the 7 series architecture been moved next to the I/O column; one exists on the left and on the right side of the FPGA per I/O bank. In previous generation FPGAs, the CMTs were placed throughout the core. Moving the CMTs physically closer to the I/O reduces lengths of clock paths and eases matching a phase detector reference path from a given byte lane DQS to a capture clock data path. The other major benefit in the 7 series FPGAs and Zynq EPPs is that a large DDR3 72-bit memory interface has upwards of 36 clock domains at the PHY level. The new blocks manage those domains cleanly and do not require or burden general-purpose clock resources, like BUFG, BUFR, or the MMCMs.

[Figure 6](#) shows a layout of a 7 series FPGAs with a pictorial expansion of the I/O bank, new I/O FIFOs, CMT, and new Phaser blocks. Each I/O bank also has four input FIFOs and output FIFOs, called IN_FIFO and OUT_FIFOs, which connect to up to 10 I/O in a byte lane and are adjacent to the I/O and fabric resources.



WP393_06_040711

Figure 6: Expanded View of the I/O Bank, I/O FIFOs, CMT, and Phaser Blocks

It should also be noted that non-memory interface structures can still be clocked in the traditional ways from the MMCMs, and both input and output clocks have similar I/O clock connectivity to previous Xilinx FPGAs like Virtex-6 FPGAs.

IO_PLL, Phaser_IN, and Phaser_OUT

The general features of the Phaser and its connectivity to I/O and clocking resources are very helpful because they compliment the byte clock structure of DDR memory interfaces and are better suited to byte-wide data channels. Dedicated clocking paths provide a smaller clock tree and reduced clock loading, and through the Phaser_IN and Phaser_OUT blocks, individual byte input and output DQS alignment and tracking are maintained. Because of their memory-specific functionality, the Phaser blocks are automatically configured and generally only supported by the Xilinx IP cores, like the MIG tool included in Xilinx ISE® Design Suite.

The new Phaser adds a per byte clock alignment, which provides fine phase-adjustment capability and 128 taps per clock period instead of fixed tap delays, which are not related to the period of memory clock.

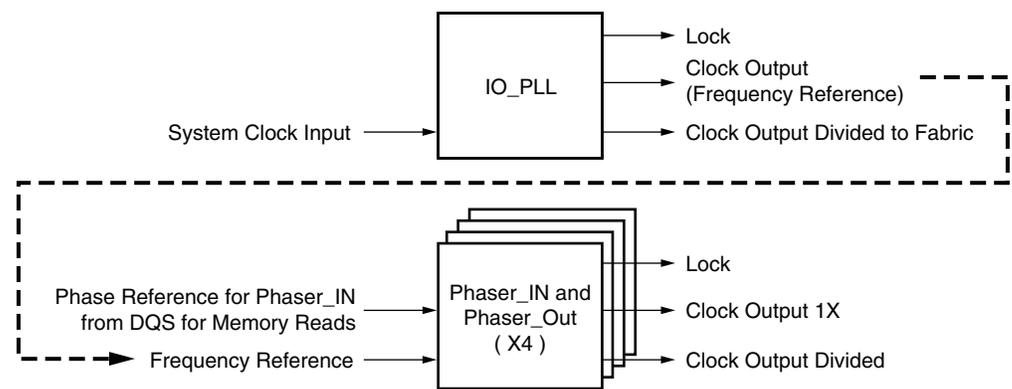
One of the new sections of the Phaser is called the IO_PLL. It is a simplified MMCM; it does not provide all of the dividing and multiplication or scaled output frequencies of the MMCM, but it does generate output frequencies up to 933 MHz for interfacing to DDR3 memory devices at 1,866 Mb/s. This block sources the CLK for the external memory devices and provides dedicated frequency references for the Phaser_IN and Phaser_OUT blocks and a divided down version for the fabric. If the IO_PLL is not being used for a given interface, it can be used for other purposes.

Another section of the Phaser is called the PHY_CONTROL block, which is not user accessible. The PHY_Control block receives PHY control words from the memory controller and initiates transfers of data within the memory PHY as requested by a PHY command word. The PHY_CONTROL block signals the start of read and write transactions to the Phaser_IN and Phaser_OUT blocks so that they know when to generate the clock and control signals needed to transfer data between the I/O SERDES and the I/O FIFOs. In addition, the advance notifications of the memory read transaction allows the Phaser_IN to prepare for the DQS edge detection and notifies the Phaser_IN of the memory rank that is being accessed.

The Phaser_IN block is user to dynamically lock to an incoming DQS clock during calibration and during memory reads. It additionally provides dedicated clock signals to the IN_FIFO and ISERDES for precision PHY timing control, serial data capture, and parallel data PHY-aligned timing to facilitate data capture during memory reads. The Phaser_IN block contains a PVT-stable circuit used for read data capture clocking for DDR3, which provides DQS phase detection and dynamic tracking of DQS position over PVT and shifting of the sampling clock into the middle of data eye.

The Phaser_OUT is used to provide alignment of outgoing the DQS/data byte group to memory CLK per byte compensated for the "flyby" CLK timing per byte group. It also provides a precise 90° phase shift of the DQS output clock relative to the output data for serial data for memory writes for the OSERDES as well as a divided down clock for the parallel side of the OSERDES and PHY side of the OUT_FIFO. As with the Phaser_IN, the Phaser_OUT also contains the same PVT stable circuit to maintain fixed timing of output clock and data relative to the CLK going to the external DDR memory.

Figure 7 shows a simple diagram of the IO_PLL, Phaser_IN, and Phaser_OUT block inputs and outputs.



WP393_07_033011

Figure 7: Simplified 7 Series FPGA IO_PLL, Phaser_IN, and Phase_OUT Connections

I/O Specific FIFOs (IN_FIFO and OUT_FIFO)

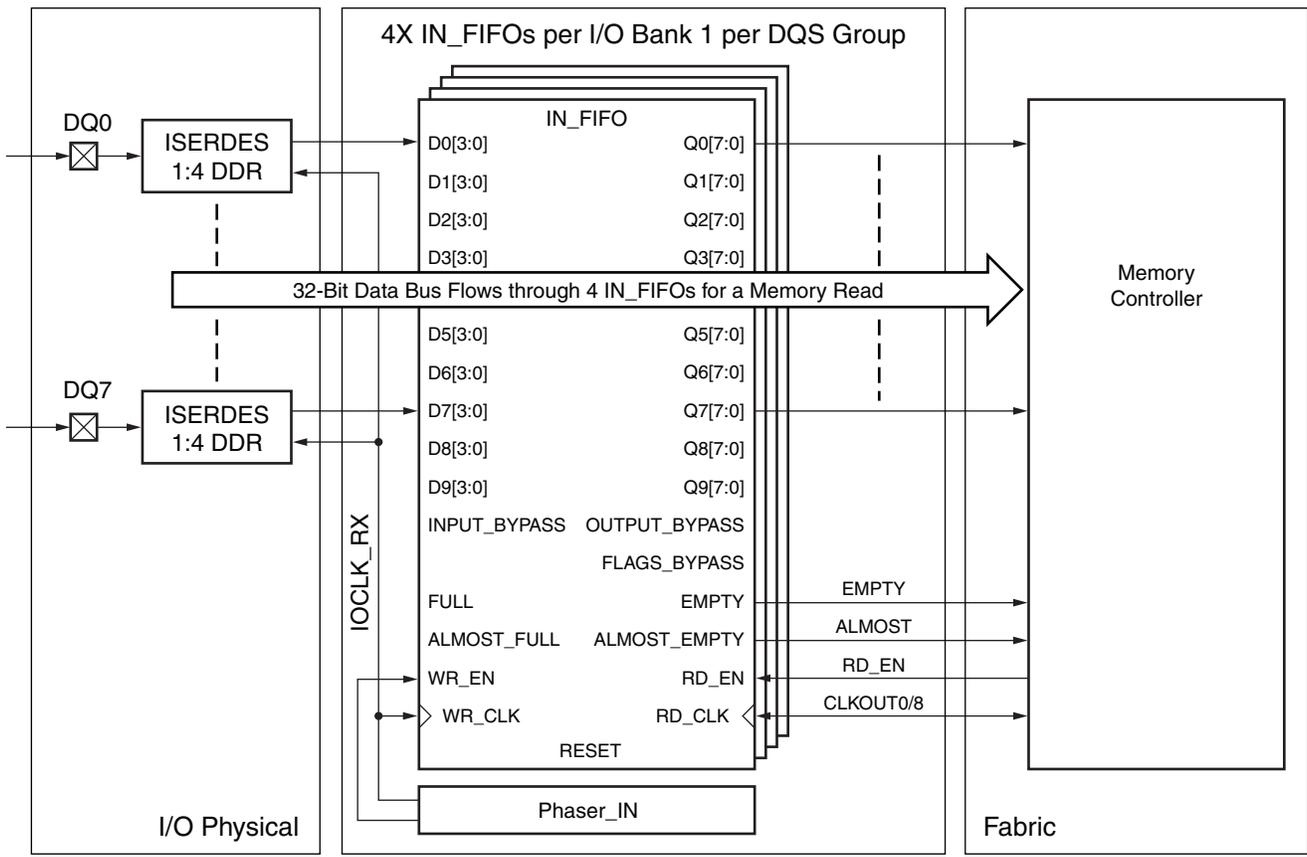
The I/O FIFOs solve system timing issues, which are normally caused by crossing PHY-side and fabric-side clock domains without a FIFO. There are four input and four output FIFOs, called IN_FIFOs and OUT_FIFOs. The IN_FIFOs and OUT_FIFOs are dedicated to up to 10 specific I/O, which correspond to one of four byte groups in a I/O bank used for memory interfacing.

These FIFOs bridge the variable phase clock domains of the PHY and the fixed phase clock domain of the fabric. These FIFOs also have bit width expansion for input data and bit width reduction for output data. These width ratios are 1:2 and 2:1 for IN_FIFO

for input data and OUT_FIFO for output data, which allow a lower clock frequency on the fabric side for design ease. Even at 1,866 Mb/s, the fabric clock is only 233 MHz.

Each I/O banks contains an IN_FIFO and an OUT_FIFO per byte group, meaning 4 of each per I/O bank. Additionally, the I/O FIFOs provide an 8-word depth and common FIFO flags (FULL, EMPTY, etc.).

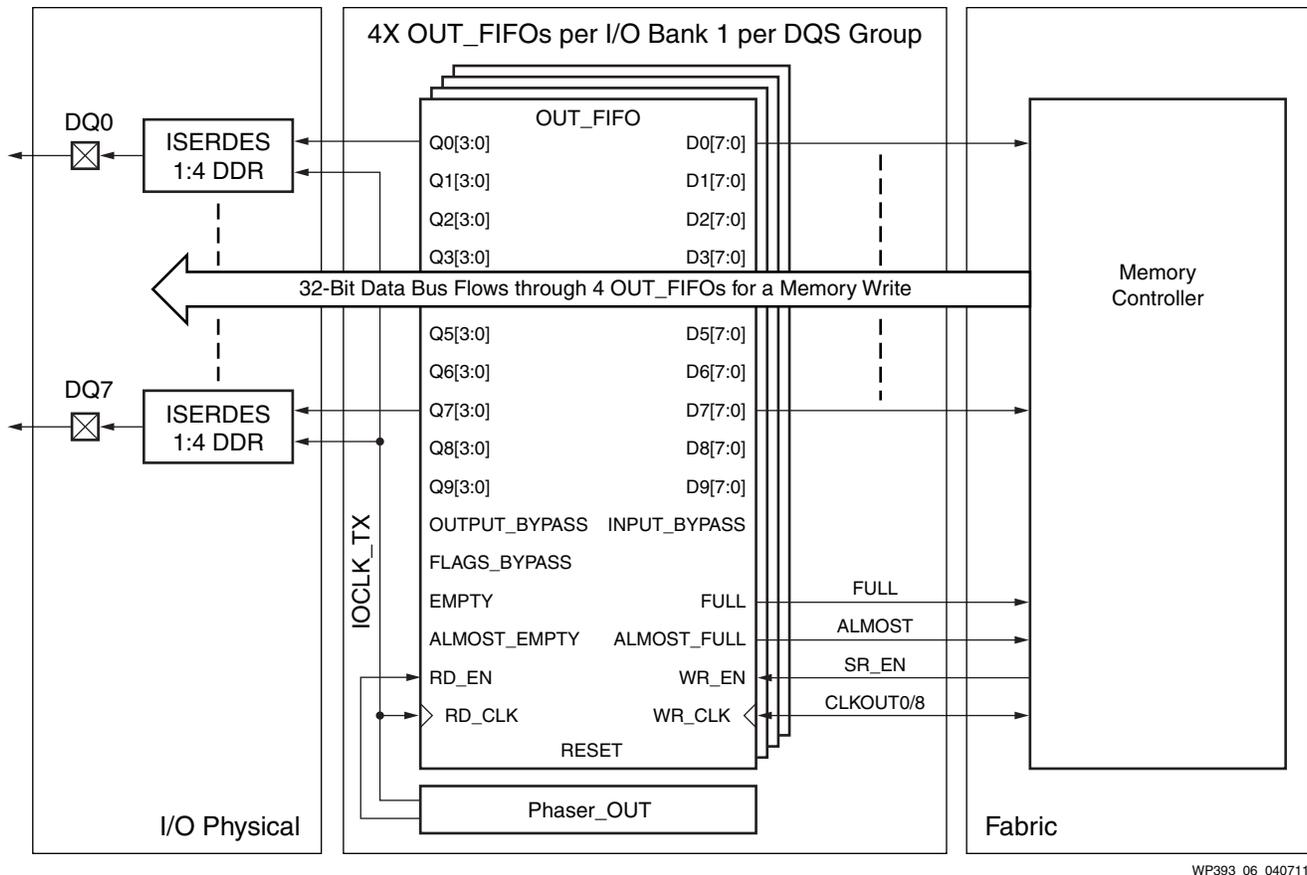
Both the IN_FIFO and OUT_FIFO have dedicated paths for their PHY (physical layer side) clocking, which come from the Phaser_IN block for the IN_FIFO and from the Phaser_OUT for OUT_FIFO. These same clocks go from the respective Phasers to the corresponding ISERDES and OSERDES in the same byte group. The basic data movement is shown in Figure 8 and Figure 9. Figure 8 shows the DDR3 input traffic flow during a memory read using the IN_FIFO and the decoupling of the I/O Physical or PHY domain from the fabric-based memory controller.



WP393_05_040711

Figure 8: DDR3 Input Traffic Flow during a Memory Read

Figure 9 shows the DDR3 output traffic flow during a memory write using the OUT_FIFO and decoupling fabric based memory controller from the I/O Physical or PHY domain.



WP393_06_040711

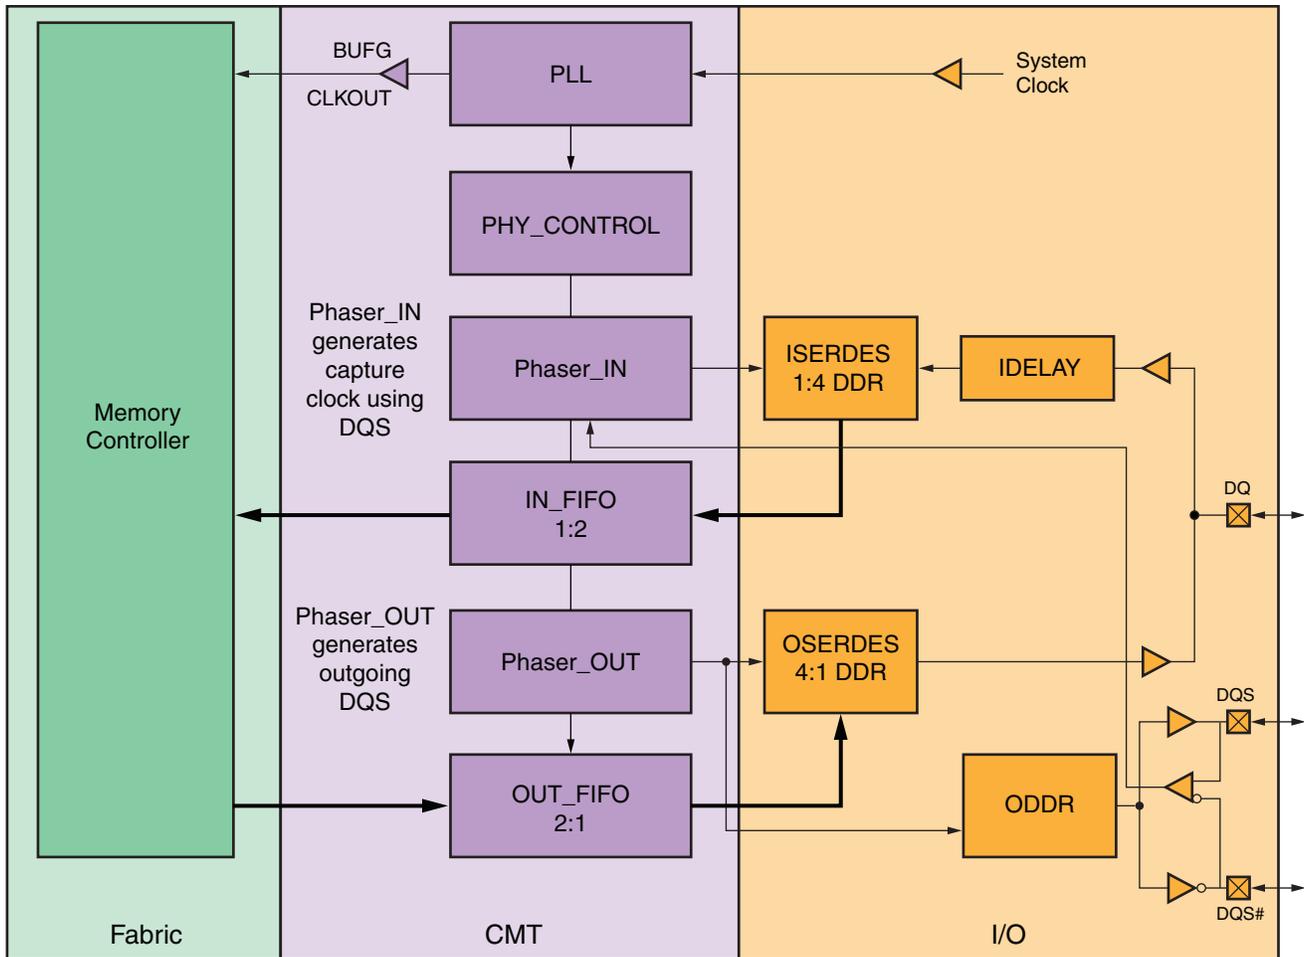
Figure 9: DDR3 Output Traffic Flow during a Memory Write

Memory Interface Enabled by the New Phaser and I/O Subsystem

The new high-performance Phaser and I/O FIFO blocks are ideal hard blocks for creating high-performance DDR3 interfaces. The addition of the I/O FIFO decouples the physical timing from core timing and the CMT, which is tightly coupled to the I/O structure to meet the demanding performance needs for high-speed DDR3 interfaces. These dedicated clocking resources and their connections between I/O, I/O FIFOs, and the Phaser have taken 36 PHY domain clocks and locked down a high-performance behavior for them without burdening normal clock functionality in the rest of the FPGA.

These new structures are unified across Artix-7, Kintex-7, and Virtex-7 FPGAs, and Zynq EPPs. The unification of these blocks and structures makes IP development easier and faster, and allows the migration of IP across all Xilinx 7 series devices.

In actual use, the Phaser components, I/O FIFOs, IO_PLL, high-speed clock routing, and memory controller are automatically generated by the Memory Interface Generator (MIG) IP cores for various types of memory, bandwidths, and bit widths. [Figure 10](#) shows a basic block diagram of a DDR3 memory interface.



WP393_10_040711

Figure 10: DDR2/DDR3 Simplified Connections to DQ/DQS Interfaces, Phasers, I/O FIFOs, and Memory Controller

The proper hookups of the clock generation for DDR2/DDR3 memory devices, a memory read, and a memory write are automatically generated by the MIG tool (bundled with the Xilinx ISE Design Suite) for all supported memory configurations used in the 7 series FPGAs and Zynq EPPs.

Generating a CLK for the DDR Memory Device

The IO_PLL receives a System Clock into CCIO (clock capable I/O) in an I/O bank or from an MMCM and locks to it. The IO_PLL generates a clock output divided at 1/8 of the memory bit rate for fabric, feeding a BUFG at up to 233 MHz for DDR3 at 1,866 Mb/s (see Figure 7 for IO_PLL connections and Figure 10 for full DDR2/DDR3 I/O connections). The IO_PLL also generates a frequency reference for up to three I/O banks from its high-speed clock output at up to 933 MHz for DDR3 at 1,866 Mb/s and sources the PHASER_IN and PHASER_OUT blocks.

Reading from DDR Memory

When reading from DDR memory, the DQS Signal for one byte lane goes into the PHASEREF input of a Phaser_IN block for that lane at up to 933 MHz for DDR3 memory interfaces at 1,866 Mb/s (see [Figure 7](#) for IO_PLL connections and [Figure 10](#) for full DDR2/DDR3 I/O connections). The IO_PLL sources its clock output to the frequency reference input of the Phaser_IN blocks for each byte lane being used. When the input DQS is valid, first during initial calibration and then during memory reads, the Phaser_IN block syncs to the Phase Reference input (DQS) for that byte lane and stores alignment values (delay values). The Phaser_IN produces a DQS-aligned continuous high-speed output, which is a delayed version of the DQS input during the read. Just after the read, the high-speed output remains continuous even when the DQS becomes invalid. Additionally, a divided down version is produced. The extension of the non-continuous DQS is used to flush various flip-flops and complete parallel transfers from the ISERDES to the IN_FIFO. The eight or nine DQ0-7/8 inputs of a byte lane are captured through eight or nine sets of IDELAY and ISERDES blocks to drive the IN_FIFO for this byte lane. The Phaser_IN also provides write enables for the IN_FIFO capture. The process runs in parallel across up to four each of Phaser_IN, IN_FIFO, and groups of IDELAY and ISERDES blocks to allow up to 36 inputs to be read in a single I/O bank.

Writing to DDR Memory

For writing to DDR memory, the IO_PLL sources its clock output to the frequency reference input of the Phaser_OUT blocks for each byte lane being used in a given I/O bank. The Phaser_OUT performs a write leveling and "flyby" timing generation depending on byte position. Each of those delays are stored and tracked over PVT. The Phaser_OUT then generates a high-speed clock and divided clock with this timing for the serial data output to a memory bit and parallel data input to the OSERDES. A 90° shifted version of the high-speed Phaser_OUT signals go to the DQS outputs. Tri-state control for the DQS and data outputs for memory writes are also generated, and the Phaser_OUT block enables the OUT_FIFO via a read to transfer data from the OUT_FIFO to the OSERDES. The eight or nine DQ0-7/8 outputs of a byte lane are generated through eight or nine sets of OSERDES after being driven by the OUT_FIFO for this byte lane. The process runs in parallel across up to four Phaser_OUT, OUT_FIFO, and groups of OSERDES blocks to allow up to 36 inputs to be read in a single I/O bank.

For more information on memory interface trends and how Xilinx addresses them, go to [WP383](#), *Achieving High Performance DDR3 Data Rates in Virtex-7 and Kintex-7 FPGAs*.

Conclusion

The I/O functionality and features of the 7 series FPGAs are optimized for diverse requirements spanning memory interfacing, networking, and legacy interfacing as well as the applications shown in [Table 2, page 9](#).

The unified architecture of the 7 series devices is carried across the Artix-7, Kintex-7, Virtex-7 FPGAs and the FPGA portion of the Zynq EPPs. It is applied to the physical and logical I/O functions as well as the I/O enhancements and the new high-performance blocks of the Phasers, I/O FIFOs, and I/O PLLs for memory and interfacing and should not be underestimated. This unification (and where possible, the similarity to Virtex-6 FPGAs) reduces customer investment for developing and

deploying products across all 7 series FPGAs and Zynq EPPs, and also simplifies migration of designs.

Additionally, Xilinx's continued key focus on power as it pertains to the 7 series FPGAs and Zynq EPPs has resulted in exceptional power efficiency, not only for I/O but for the rest of the 7 series FPGAs, enabling more integration than ever before.

Lastly, the performance levels achievable in the 7 series FPGAs and Zynq EPPs through I/O enhancement and the addition of new very high-performance I/O functions (Phaser, I/O FIFOs, I/O PLLs, CMT) allows the 7 series FPGAs and Zynq EPPs to meet the most challenging performance levels for high-performance memory interfaces, networking, and other demanding applications.

For more information, go to:

<http://www.xilinx.com/technology/roadmap/7-series-fpgas.htm>.

Revision History

The following table shows the revision history for this document:

| Date | Version | Description of Revisions |
|----------|---------|--------------------------|
| 07/14/11 | 1.0 | Initial Xilinx release. |

Notice of Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials, or to advise you of any corrections or update. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.