DesignCon 2007

Memory Design Considerations when Migrating to DDR3 Interfaces from DDR2

Raj Mahajan, MemCore Inc.
raj@memcoreinc.com

Abstract

The emerging DDR3 memory standard will extend the performance range of DDR memories considerably, while maintaining some amount of backwards compatibility with the existing DDR2 memory standard. It is important to understand the similarities and differences between the DDR3 standard and the existing DDR2 standard in order to get the maximum benefit from the new standard while re-using as much as possible from any previous DDR2 memory interface design. This paper will provide the reader with a detailed understanding of the key design considerations when migrating to a DDR3 system interface from a DDR2 interface.

This paper will review the new DDR3 features and compare and contrast them to previous features available in the DDR2 specification. One of the biggest changes is the in Physical Layer (PHY) portion of the memory interface and these changes will be high-lighted and illustrated with an example design of a high performance processor interface. The areas where backwards compatibility should be maintained will also be illustrated with an example design, showing how simple changes can provide significant benefits in re-use and system flexibility.

Author(s) Biography

Raj Mahajan has more than 10 years of experience architecting, designing, and verifying memory access solutions for advanced ASICs for a variety of target markets. He started his career at Intel Corp, where he architected and designed advanced render cache controllers that shipped hundreds of millions of units in several generations of graphics-enabled PC chipsets. Following that he held a lead design position at 2Wire, Inc., a successful start-up addressing the residential broadband access market, where he led the integration and verification of their flagship SoC, which shipped first silicon. At Ingot Systems he led the architecture, design, and verification of MemCore's flagship memory controller. Raj holds a Bachelor of Computer Engineering from Georgia Institute of Technology.

Introduction

The emerging DDR3 memory standard will extend the performance range of DDR memories considerably, while maintaining some amount of backwards compatibility with the existing DDR2 memory standard. It is important to understand the similarities and differences between the DDR3 standard and the existing DDR2 standard in order to get the maximum benefit from the new standard while re-using as much as possible from any previous DDR2 memory interface design. This paper will provide the reader with a detailed understanding of the key design considerations when migrating to a DDR3 system interface from a DDR2 interface.

A Comparison of DDR2 and DDR3 Memory Standards

The DDR2 memory standard is being upgraded with the advent of the DDR3 standard. variety of memory devices available today provides the system architect with multiple options when selecting a memory. Before going into the detailed comparison of DDR2 and DDR3 let's review the key features of a typical DDR2 memory subsystem and the associated memory controller. This will serve as a baseline for the detailed comparison.

DDR2 Description

A typical DDR2 memory subsystem uses a DIMM to house multiple DDR2 memory devices. A typical DDR2 DIMM architecture is illustrated in Figure 1 below. The control and address signals come onto the DIMM and are routed to the memory devices in a T-branch topology. This architecture balances the delay to each memory device, but introduces additional skew due to the multiple stubs and the stub length.
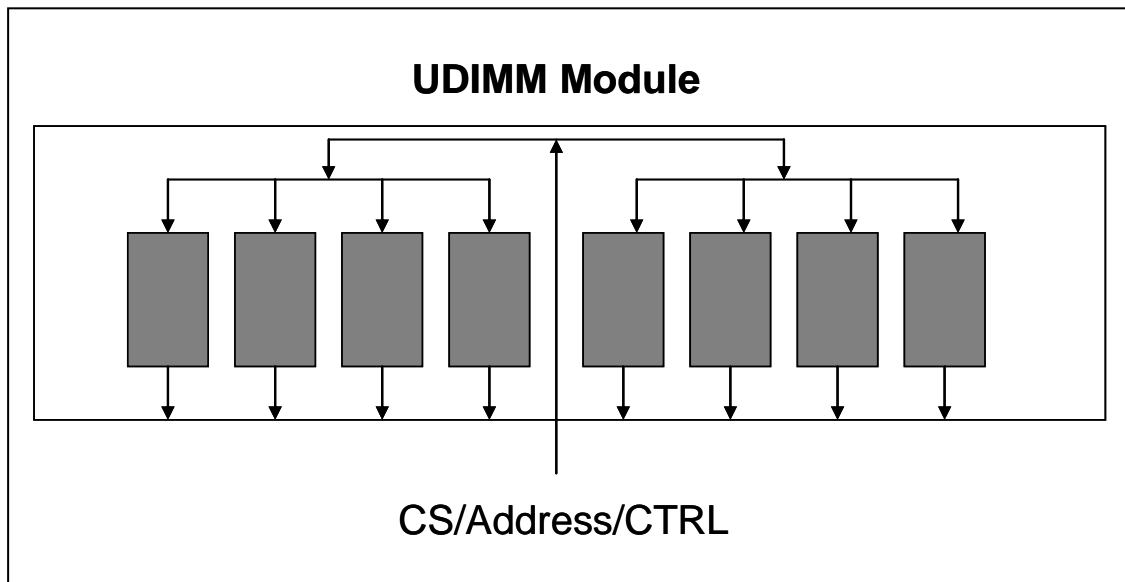

Figure 1: DDR2 Memory Module Architecture

A DDR2 memory controller is located on the chip driving the DIMM module. A typical DDR2 memory controller is show in the block diagram in Figure 2. The PHY is responsible for the physical interface between the DDR DRAM and the rest of the

system. Timing is controlled precisely to insure data is captured or presented in just the right relationship with the DRAM clocking signals. Data read from the DRAM is optionally corrected by the ECC block and then provided to the pending write and read modify write FIFO. If ECC is being used, the ECC check bits are computed prior to the write to memory by another optional ECC block in the write path.
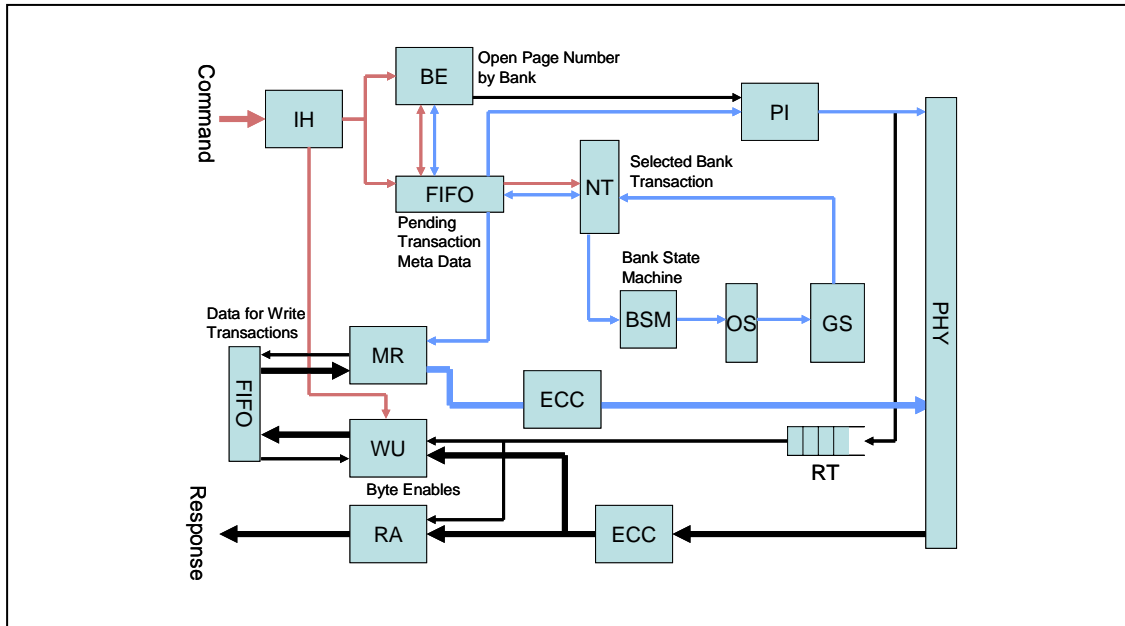

Figure 2: DDR2 Functional Block Diagram

The scheduler prioritizes the current list of commands and determines which command is the most urgent and issues that command to the DRAM. Data is read or written to the memory based on the schedulers computation of access priority. The scheduler constantly works towards the goal of maximizing overall system bandwidth while issuing all high priority command as quickly as possible.

Commands are optionally pipelined and added to the pending FIFO. If the command is most urgent (direct read) it bypasses the pending FIFO and is issued directly to the memory. Regular priority accesses make their way through the pending read FIFO or the read token FIFO for command completion.

DDR3 Description
The main thrust of the DDR3 memory standard is to increase bandwidth while making it relatively easy for the designer to take advantage of this bandwidth increase. Innovations in the physical layer (PHY) portion of the DDR3 interface support this increase in bandwidth. One of these innovations is the use of a Leveling technique that adjusts for the delay between DDR3 memories.

The DDR3 specification can support a fly-by architecture either on a memory module or on a board. In this architecture, illustrated in Figure 3 below, the signals from the memory controller are connected in series to each memory component. In this example, a

memory module, the signals from the DDR3 PHY come into the middle of the module and connect to each memory chip sequentially. This reduced the number of stubs and the stub lengths. Termination is placed just at the end of the signal. This improves the signal characteristics over the traditional DDR2 topology.
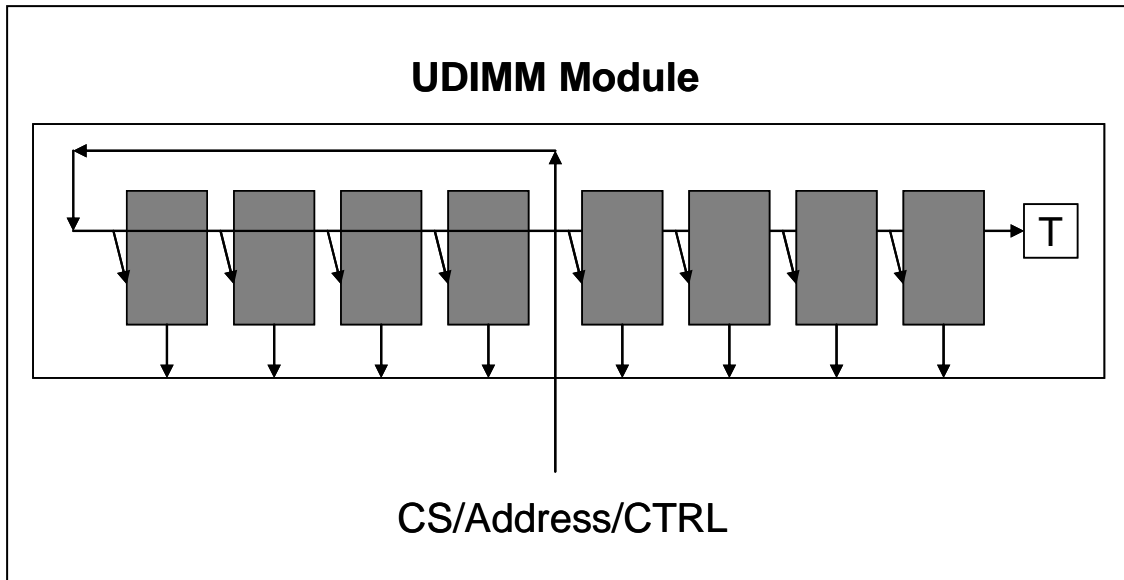

Figure 3: Fly-by Topology for DDR3 Unbuffered DIMM

The draw-back to this approach is that the delay from the PHY output signals to each memory is slightly different, depending on where the memory chip is in the sequence. This delay difference needs to be compensated for by the DDR3 PHY and uses the new Leveling feature required by the DDR3 specification. There is a different technique for both write and read leveling.

Write Leveling

During Write Leveling the memory controller needs to compensate for the additional flight time skew delay introduced by the fly-by topology with respect to strobe and clock. In particular, the tDQSS, tDSS and tDSH timing requirements would be very difficult to meet. These timing parameters can be met by using a programmable delay element on DQS with fine enough granularity so that the proper delay can be inserted to compensate for the additional skew delay. The figure below shows the needed timing relationship.

The source CK and DQS signals are delayed in getting to the destination as illustrated by arrow #1 and arrow #2 respectively. (This delay can be different for each memory component on the memory module and will be adjusted on a chip by chip basis and even on a byte basis if the chip has more than one byte of data. The diagram illustrates just one instance of a memory component). The memory controller repeatedly delays DQS, a step at a time, until a transition from a zero to a one is detected on the destination CK signal. This will re-align DQS and CL so that the destination data on the DQ bus can be captured reliably. Because all this is done automatically by the controller the board designer need

not worry about the details of the implementation- he or she just benefits from the additional margin created by the Write Leveling feature in the DDR3 memory controller.
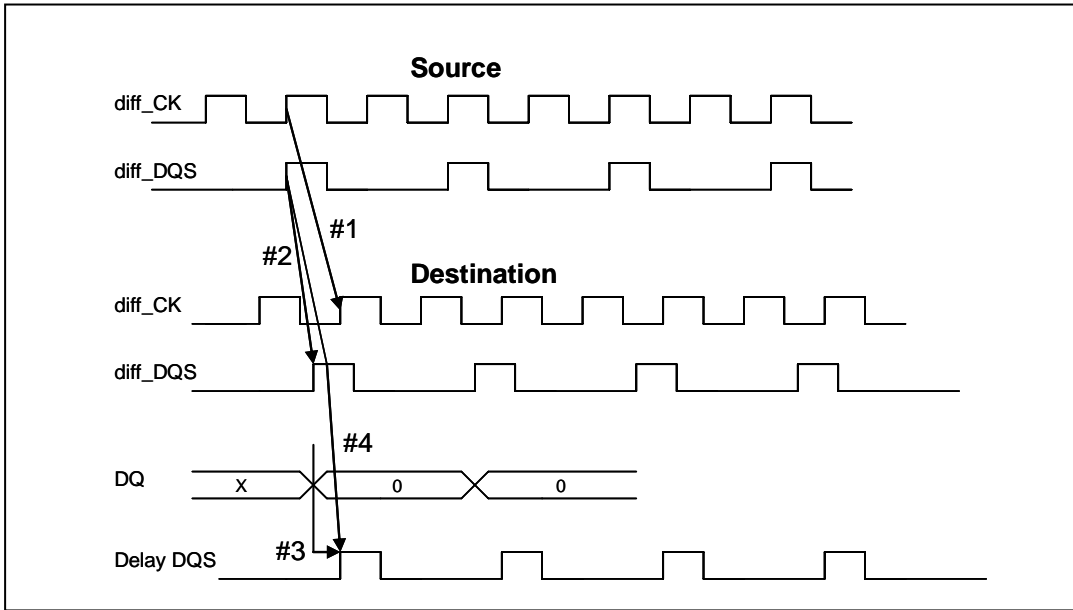

Figure 4. Timing Diagram for Write Leveling

Read Leveling

During Read Leveling the memory controller adjusts for the delays introduced by the fly-by memory topology that impact the read cycle. This is done via the addition of a special Multi-Purpose Register (MPR) in the DDR3 memory device. The MPR can be loaded with predefined data values via a special command from the memory controller. These data values can be used for system timing calibration by the memory controller.

As shown in Figure 5 below, the MPR can be selected, by setting a bit in another memory register (EMRS3, bit A2), to switch the source of data for memory read to come from the MPR, not the normal memory array. The MPR data is substituted for the DQ, DM, DQS and /DQS pads on the memory device. This feature allows the memory controller to calibrate the timing of the read path to adjust for any additional delays introduced by the DDR3 fly-by architecture.
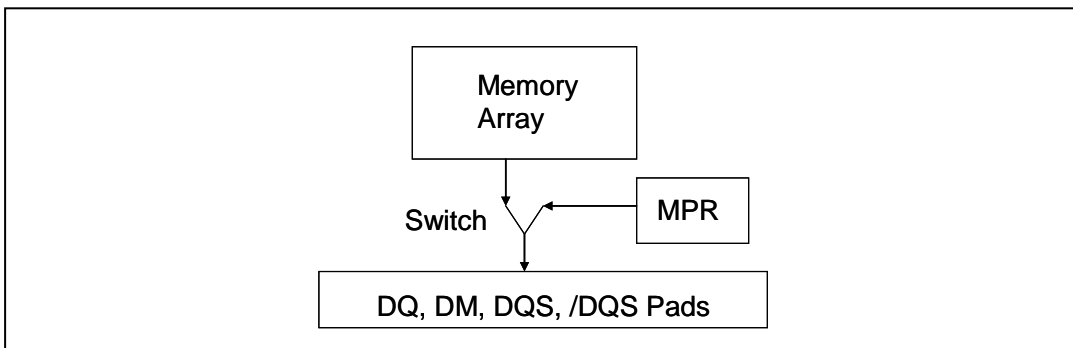

 Figure 5. Read Leveling Using MPR

Other DDR3 Features
DDR3 has additional features to improve performance and reliability. These include a Reset pin, an 8-bit pre-fetch, and ZQ calibration. A new Reset pin is used to clear all state information in the DDR3 memory device without the need to individually reset each control register or power down the device. This saves time and power when bringing the device to a known state. The 8-bit pre-fetch is used in conjunction with burst length of 4 or 8. This improves performance for sequential accesses. The new ZQ calibration feature allows the memory device to take a longer time for calibration at start-up and a smaller time during periodic calibration activities. The table below shows a feature by feature comparison of DDR, DDR2 and DDR3 memory devices.

	DDR	DDR2	DDR3
Data Rate	200-400Mbps	400-800Mbps	800-1600Mbps
Interface	SSTL_2	SSTL_18	SSTL_15
Source Sync	Bi-directional DQS (Single ended default)	Bi-directional DQS (Single/Diff Option)	Bi-directional DQS (Differential default)
Burst Length	BL= 2, 4, 8 (2bit prefetch)	BL= 4, 8 (4bit prefetch)	BL= 4, 8 (8bit prefetch)
CL/tRCD/tRP	15ns each	15ns each	12ns each
Reset	No	No	Yes
ODT	No	Yes	Yes
Driver Calibration	No	Off-Chip	On-Chip with ZQ pin
Leveling	No	No	Yes

Table 1: DDR, DDR2 and DDR3 Feature Comparison

Planning For Migration- an Example Design

In order to explore how to prepare a DDR2 design for migration to a DDR3 design it will help to establish an example system. Let's assume that the system will require a DIMM interface for DDR2 and will want to use a similar type of memory module in the DDR3 implementation. Performance is increasingly important for many applications so the decision is to initially design the controller as a DDR2 design, but to allow future migration to DDR3. As much as possible, we want to make it easy to modify the board and the memory controller to migrate from the DDR2 implementation to DDR3.

Board Level Issues
One of the biggest issues when thinking of migrating from DDR2 to DDR3 is that the DIMM have different pin-outs and sizes. This means that it will be very difficult, at the board level to just plug in a new memory module. The best you can hope for is to take into account some of the other key characteristics and make it easy to spin the board for a DDR3 module. Probably the most important items to take care of at the board level will be DQS, the Reset pin.

DQS

In DDR3 DQS is specified as differential while in DDR2 it can be single ended or optionally differential. Clearly if the differential version is used in DDR2 it will make the transition to DDR3 easier. This may require additional pins in the memory controller, but if upward compatibility is important the extra pins will be worth it. Your DDR2 implementation will be more robust as well.

In DDR3 the DQS is sourced by each memory device in order to account for the additional delay from the fly-by topology. The number of DQS signals is thus larger in the DDR3 implementation that the DDR2 version. Again, if the additional pins are not a big issue it will help with the migration to DDR3 to implement the additional DQS signals in the DDR2 implementation.

Reset
The Reset pin present in DDR3 would be easy to add to DDR2. The pin wouldn't do anything in the DDR2 implementation, but including it would insure that the pin is available when its time to migrate to DDR3.

Memory Controller Issues
Other aspects of the DDR2 to DDR3 migration will require some impact to the memory controller. If the DDR2 memory controller is designed with some of these issues in mind it can simplify the process considerably. Some of the most important issues are the output drivers, DLLs for write launch, and Read Leveling.

Output Drivers
The DDR2 standard calls for 1.8V SSTL IOs. DDR3 calls for 1.5V SSTL IOs. It may be difficult to find an IO buffer that can support both standards. It might require a programmable IO, similar to those found on FPGAs, to support both standards. A change in IO buffers would require a spin of the chip driving the DDR3 memory, but perhaps a metal mask option could be used to make this change less expensive.

DLLs for Write Launch
Typical DDR2 memory controllers can get away with one DLL for several data outputs. In DDR3, due to the fly-by topology, it will be more usual to see a DDL for every 8-bits or so. This would require a larger number of DLL to be included in the DDR2 design in order to provide the resources required for the DDR3 migration. A digital DLL implementation can be very compact in die size and can minimize the overhead associated with the DDR3 requirement.

Read Leveling
Typical DDR2 memory controllers use an extra pair of IO pins to calibrate the controller read timing. These pins are used to help adjust the incoming data with respect to the strobe. Other controllers use a training sequence by writing and reading data from memory and adjusting the strobe to optimize the data capture point. In DDR3 the Read Leveling feature is used to do this and requires no additional pins. If the memory controller can be designed to include the Read Leveling feature, even if not use din DDR2, it would help with DDR3 migration.

Conclusion

DDR3 offers a substantial performance improvement over previous DDR2 memory systems. New DDR3 features, all transparently implemented in the memory controller, improve the signal integrity characteristics of DDR3 designs so that higher performance is achieved without an undue burden for the system designer. If proper consideration is given to any new DDR2 memory design, it can be a relatively easy upgrade to support DDR3 in the next generation design. This paper identified the key differences between DDR2 and DDR3 and illustrated some of the key issues that need to be addressed to easy migration to DDR3.