White Paper

An Improved and Simplified Interface Protocol for DDR Memory

James Gay
Senior Member of the Technical Staff, Freescale Semiconductor, Inc.

Abstract

Multiple data rate memories are now common for dynamic (SDRAM) and non-volatile (Flash) type memories. Typically, these are implemented with a double data rate (DDR) protocol and interface that provides source synchronous data and clock and/or data strobe. But the protocol used for the transfer of data between the memory controller and memory is not symmetric with respect to the direction of data transfer. Write operations to memory from the controller allow easy acquisition of the data by the memory device, but this is not the case for read operations from the memory to the controller. Presented here is a proposal for improving this protocol that alleviates the implementation difficulties and logic in the memory controller design without impacting system performance and without significantly affecting memory device design or complexity.

Table of Contents

1 Introduction

The Double Data Rate (DDR) memory interface for synchronous dynamic random-access memory (SDRAM) has been around since the mid-1990s with standards being released starting in 2000 by the Joint Electron Device Engineering Council (JEDEC). Since that time there have been updates to the specifications and new SDRAM varieties such as DDR2, DDR3, DDR4, LPDDR, LPDDR2 and others but all using some form of double data rate transfer. The double data rate method of data transfer has also been applied to non-volatile (Flash) memories, particularly high performance Flash using a Quad SPI (QSPI) interface.

From a memory device standpoint the read and write data transfer DDR interface is simple and easy to implement in that data being transferred from the host memory controller to the memory during a write cycle provides a clock and / or data strobe (DQS) signal that is source synchronous and centered within the valid write data window. This allows for guaranteed setup and hold time about the clock or DQS edges to allow the memory device to easily capture or register the write data. However, during a read cycle data transfer the memory drives out new data changing with the DQS edges that are only valid between the data strobe edges. The data strobe for the read data transfer is source synchronous with the read data, but not aligned to be centered within the valid data window. Thus, the protocol used for the transfer of data between the host memory controller and memory device is not symmetric with respect to the direction of data transfer. Write operations to memory from the host controller allow easy acquisition of the data by the memory device. But this is not the case for read operations from the memory to the host controller. Figure 1 below shows the inconsistency in the protocol between read and write data transfer operation and is highlighted by the red oval areas in the figure that denote the valid data window with respect to the data strobe (DQS) edge placement. The figure shows a data strobe signal for write operations which is generally used by DDR SDRAM, but is typically not used for DDR QSPI Flash since the clock, which aligns with the data strobe, serves this purpose. Note in the figures the illustration of differential pair signals for clocks and data strobe signals. This is done for completeness, but may or may not be necessary for certain devices at certain frequencies of operation.

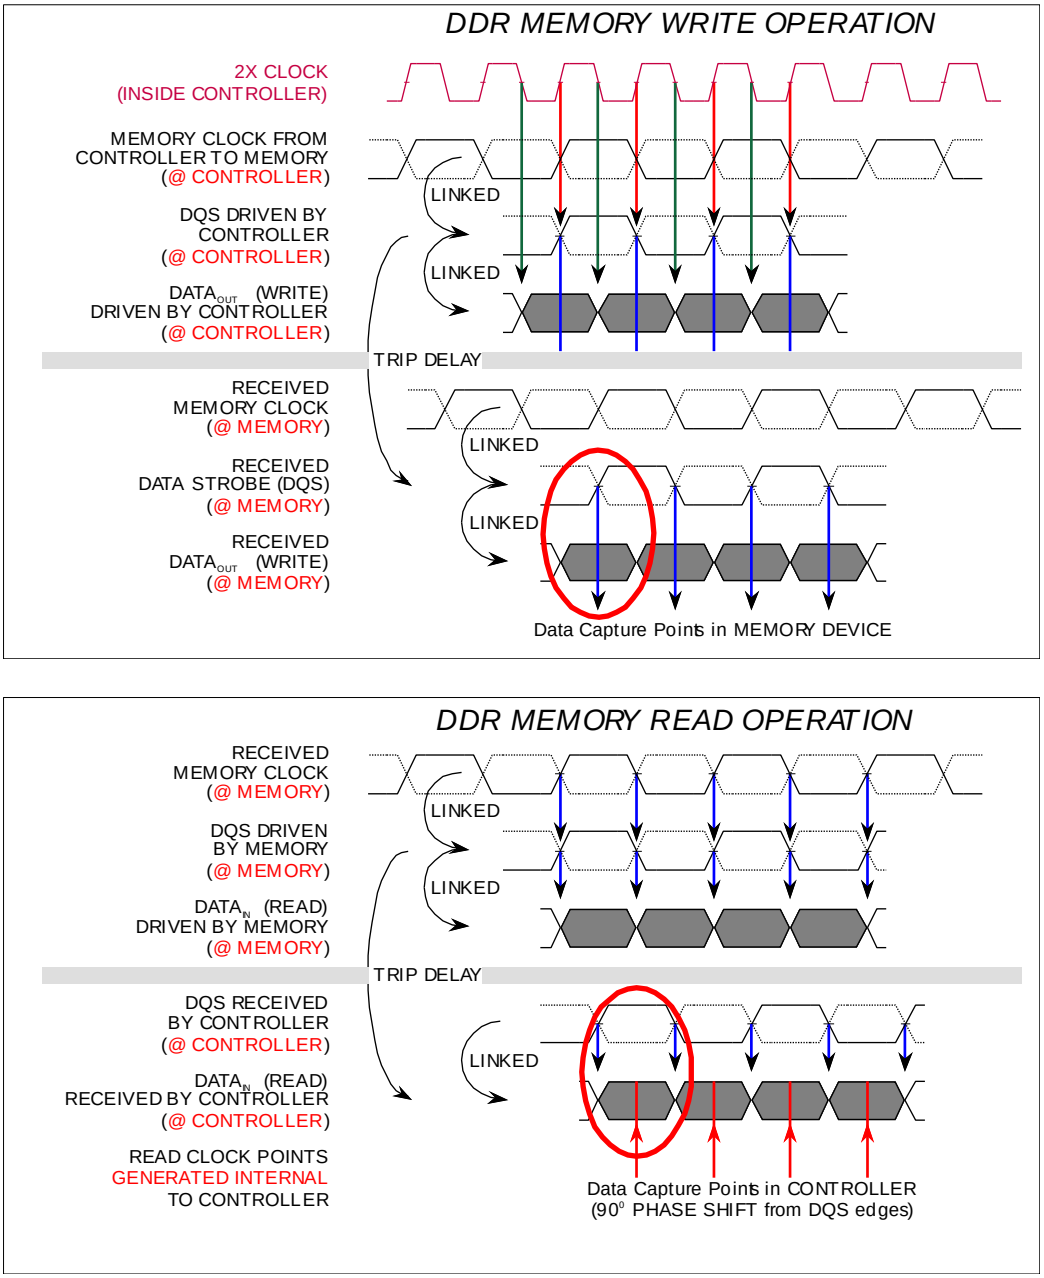

Figure 1. Inconsistent Protocol Between Write and Read Data Transfers

From the host memory controller standpoint, the non-center aligned data strobe for read data complicates logic in the host controller because the host controller must by some internal means find the proper point in time at which to register or capture read data within its valid data window. In addition, the data strobe is bursty in nature so capture point edges must be generated shortly after the data strobe starts toggling and must accommodate and adapt to fluctuations in voltage and temperature as the system operates in its environment.

Various methods have been devised to capture read data within its valid data window. Two basic methods are typically used to determine the proper time to capture valid data.

One method employs periodically running some sort of learning routine whereby known good data is written into the memory device and read back many times while adjusting the delay setting sample point.

For Flash memory protocol this may involve use of up-frequency multiplied clocks for over sampling data. Several sample points are taken with various delay settings tabulating those delay settings that result in correct data results and those that capture incorrect data. After determining the delay setting limits that provide correct data a suitable delay setting can be chosen typically centered between the first minimum and first maximum delay settings that resulted in incorrect data so that valid correct data can be captured for following read data transfers. Unfortunately, due to semiconductor processing variations and temperature and voltage variations (PVTs), a “one size fits all” setting may not be achievable. For example, even if the same host controller and memory types are used on two separate but identical Printed Circuit Boards (PCBs) the processing variation between the host controllers and/or between the memory device(s) and/or PCB manufacturing variations may result in different delay setting values between PCBs. Even on the same PCB with a fixed host controller and memory device(s), as the board heats up or cools down or if there are power supply variations, the chosen startup delay settings may prove problematic unless there is periodic recalibration of the delay settings. Of course the higher the clock frequency, the more critical that recalibration be periodically run.

A second, more sophisticated, method is to build an adaptive read clock recovery (RCR) system within the host memory controller that can adaptively adjust for process, voltage and temperature (PVT) variations. Since the host memory controller provides the clock to the memory device(s) and is generally derived from a 2x or higher frequency internal clock, it can use this clock or the actual memory clock along with some form of delay lock loop (DLL) to find the mid point of the memory clock period in terms of delay elements. When the bursty DQS is received from the memory device the previously determined number of delay elements is applied to the received DQS signal to delay the edges and position them where read data is valid for capture. The DLL feedback loop is adaptive in the sense that it adjusts with PVT and will either add or subtract the number of delay elements required to keep the delay centered at the mid point of the memory clock period. The number and delay granularity of each of the delay elements depend on the frequency ranges that must be supported by the host memory controller. If low frequency operation must be supported, more delay elements or elements with long delay may be needed to provide for enough delay while for very high frequency operation, very fast delay elements will be needed to have enough fine granularity. Thus the read clock recovery logic can be complex and difficult to design. In addition, testing a very fast host memory controller that uses delay elements having on the order of a few tens of picoseconds of delay in the RCR DLL is no trivial task especially using low cost testers.

Methods used to capture valid read data can become complex and tricky to design. Thus, if there was a way to make the protocol between DDR write and read operations symmetric, system reliability could be improved and the host memory controller design complexity, silicon logic and area, power, test time, and, cost could all be reduced to benefit customers.

2 New DDR Center Aligned Read Strobe Protocol

Figure 2 below shows a means by which DDR write and read operations can be made symmetric with the addition of a 90^{o} phase shifted clock for the memory device to use to generate the data strobe centered within the valid read data window. Note that the data strobe edges, driven off the 90^{o} phase shifted clock, are now centered within the valid read data window.

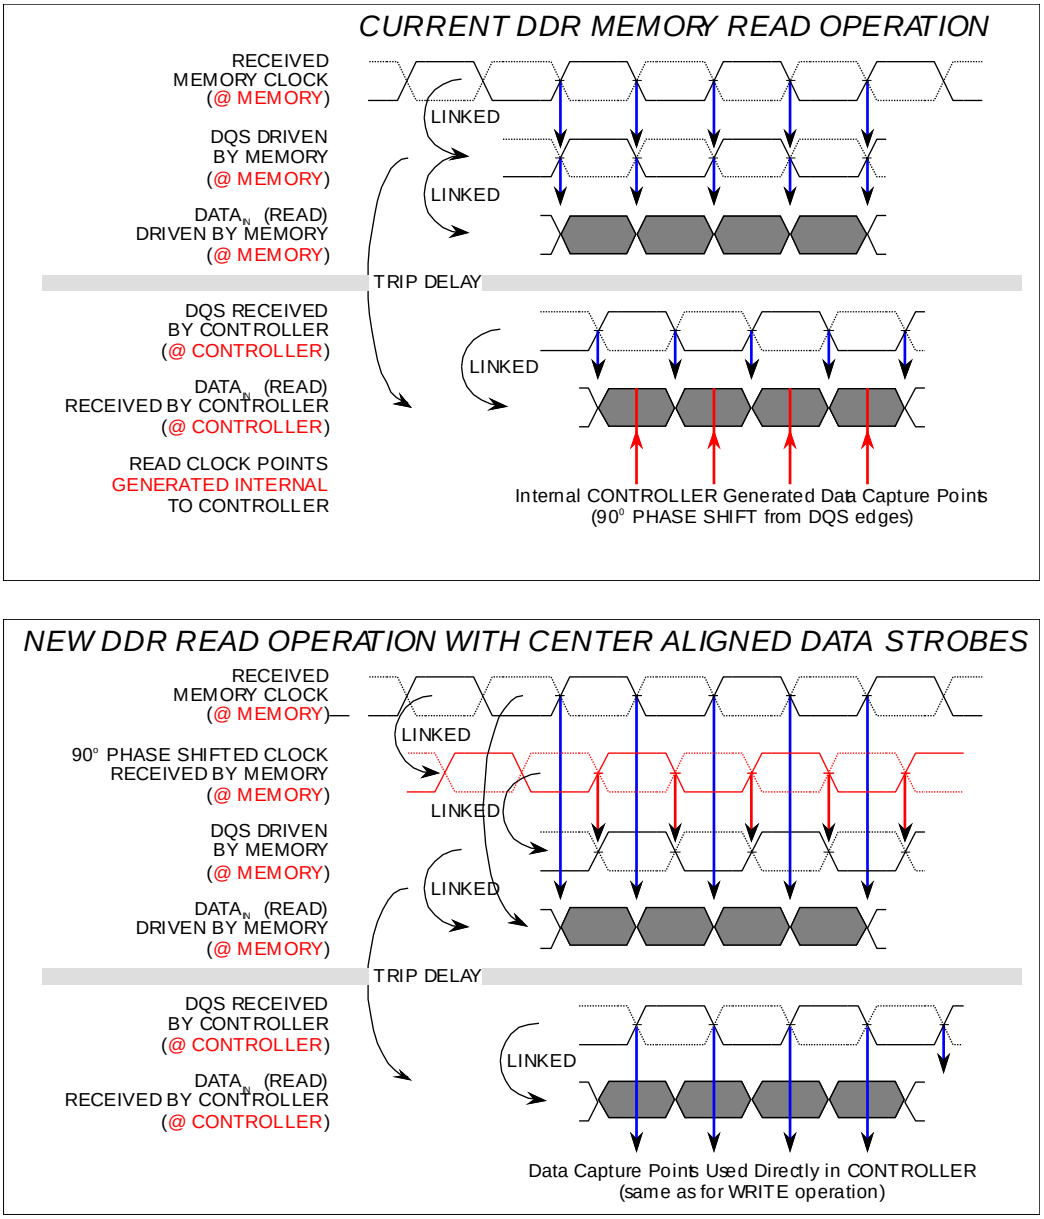

Figure 2. Legacy vs. New DDR Read Protocol With Center Aligned Data Strobes

It is quite easy for the host chip to generate 0^o and additional 90^o phase aligned memory clocks. The host chip that contains the memory controller generally has higher frequency clocks that are used to generate the traditional memory clock and can also be used to generate the 90^o phase shifted clock. A simple positive edge flop and negative edge flop that have equal clock-to-Q output times each configured as a divide by two function is an easy way to do this if a 2x clock is available. Some circuit like this, as shown in Figure 3, is probably already present in the host memory controller since it has to supply write data on the equivalent of the positive edge of the internal 2x clock and the DQS signal on the equivalent of the negative edge of the 2x clock. If the higher frequency clock is not available inside the host chip it is usually a simple matter of tapping off of the host PLL ring oscillator used to generate the 0^o memory clock to generate a 90^o phase shifted clock for this new protocol. Again, for this type of host system, the

equivalent of the 0° and 90° clocks already have to be available to handle the current legacy DDR write protocol operations. Thus, the host controller generation of the new 90° clock is fairly trivial to produce as long as care is taken to manage equivalent delays for the two clocks.

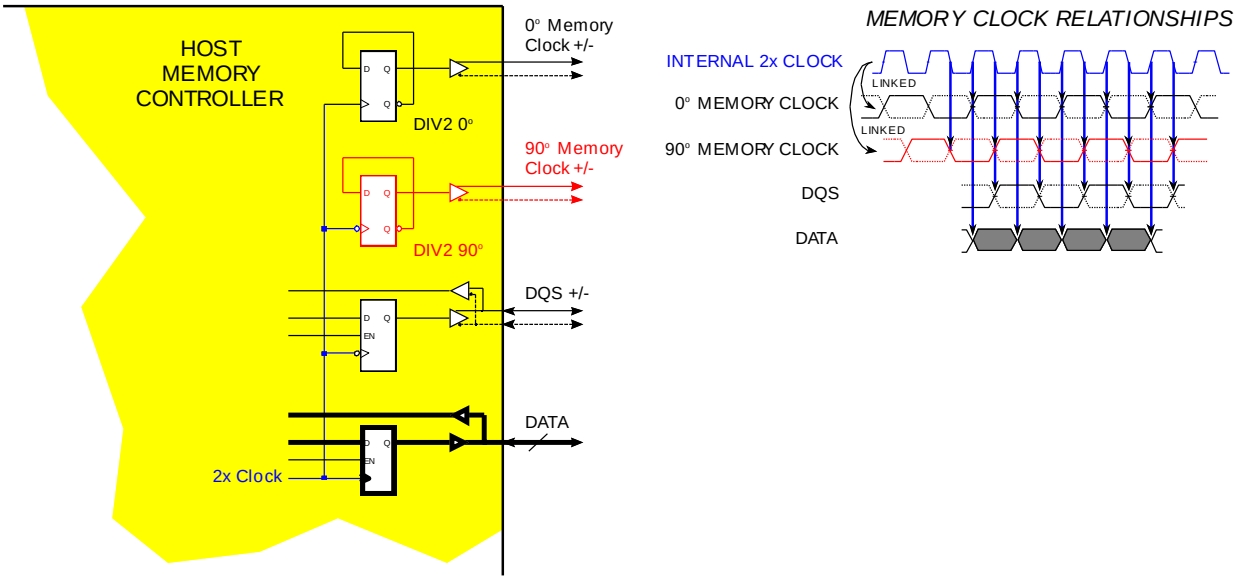

Figure 3. Host Controller Memory Clock Generation

For the memory device, the 0° memory clock provided by the host memory controller can continue to be used as the timing reference for the transfer of read data. With the addition of the 90° phase shifted clock provided by the host memory controller, the memory device can now use it as the timing reference for the assertion and negation of the DQS signal. Since the host controller is the source of the 0° memory clock *and* the 90° phase shifted clock, these clocks are temporally linked. Provided the memory device properly handles its internal delays driving data and DQS, the DQS and read data generated by the memory device are also temporally linked and source synchronous from the memory device back to the host memory controller.

There are other ways to utilize the 0° memory clock and 90° phase shifted memory clock within the memory device to manage the timing of read data and DQS. For example, the 0° and 90° clocks can be combined to create a 2x memory clock inside the memory device for timing read data and the data strobe. This 2x memory clock internal to the memory device is similar to what might be used in the host memory controller when it has to provide data and a center aligned DQS relative to data windows for the memory device during write operations.

Formalizing this new protocol into a specification for a DDR QSPI Flash Memory Interface is shown in the following Figure 4. Note, only a Read Operation timing diagram is necessary since the Write Operation remains as it was.

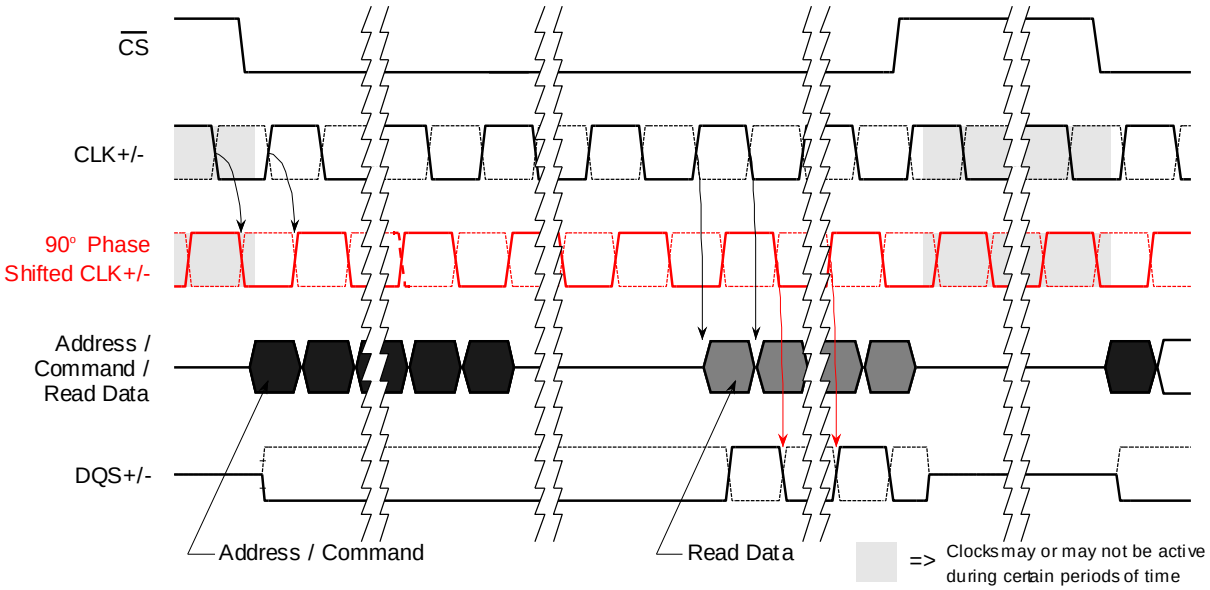

Figure 4. QSPI Read Operation Using Centered Read Strobe Protocol

The 90° phase shifted memory clock described in the above section and figures is easy to generate by the host memory controller, but other phase shift values could be used. Modifying the phase shift to less than 90° would shift DQS edges earlier in time with respect to the valid data window providing greater data hold time at the expense of data setup time for the receiving host memory controller. And conversely, adjusting the phase shift to greater than 90° would have the effect of shifting the DQS edges later in time with respect to the valid data window providing greater data setup time at the expense of data hold time. In certain systems having this ability to adjust the phase shift may be desirable for greater timing margins and higher frequencies of operation.

3 Other Benefits of the New Center Aligned DDR Read Strobe Protocol

There is another benefit for the new center aligned read strobe protocol for DDR SDRAM related to the elimination of preamble and postamble detection. This eliminates wasted clock cycles needed in the current legacy protocol for preamble and postamble time periods, increasing data transfer throughput and thereby improving system performance.

For example, consider the case from the existing read protocol in Figure 5 below showing two back-to-back read transfers.

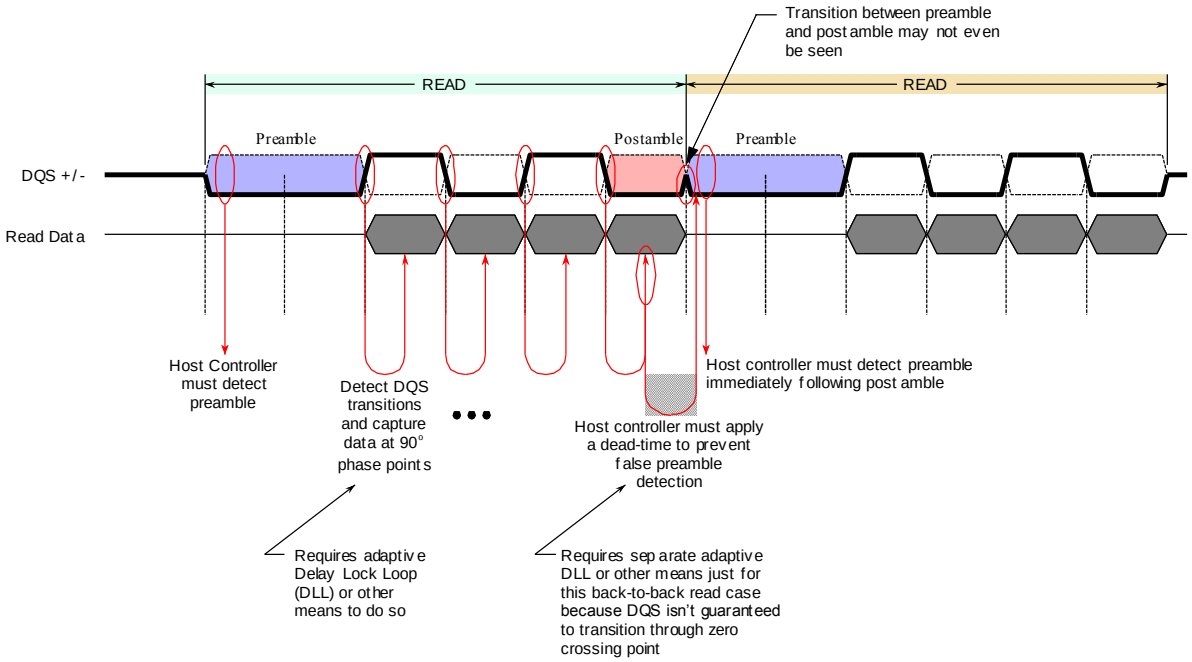

Figure 5. Back-to-back Read Operation Using Exiting DDR Read Protocol

Here the host memory controller must first detect the preamble to gate or allow upcoming DQS transitions to be seen as valid precursors of data transfers. Next it must use the DQS transitions as starting triggers to internally determine when to capture data somewhere hopefully near the center of the valid data windows via an adaptive DLL or other means. If there is a back-to-back read case as shown in Figure 5 where the postamble and preamble are abutted to each other, depending on host controller On Die Termination (ODT) settings, memory device drive settings, and bus RC parasitics, the tristating of DQS and reassertion of DQS by the memory device may not even be seen as a blip on the signal and is certainly not guaranteed to reach a valid logic state or zero-crossing point that would be easy to detect. Therefore, the host controller must also apply a time-out or dead-time between when the last expected data is received from a first read operation before another search for a new preamble from a second read operation is started. This involves the need for another adaptive DLL function or other method that is triggered not from a DQS transition, because there may not be one between back-to-back read transfers, but triggered from when the last data was captured. It's delay must be longer than the delay used from a DQS transition to the center of the data valid window to guarantee the search for a valid preamble does not start until the time point is beyond the period between the postamble and preamble abutment.

Preamble detection can become quite complex in IO buffer and control logic to guard against false preambles due to noise on external strobe signals. If the new center aligned read strobe protocol is used there's no need for preambles and postambles. Since the data strobe signal in this new protocol is centered with valid data, the host memory controller simply needs to capture valid data on DQS transitions. Figure 6 below repeats the previous timing example but using the new center aligned data strobe protocol.

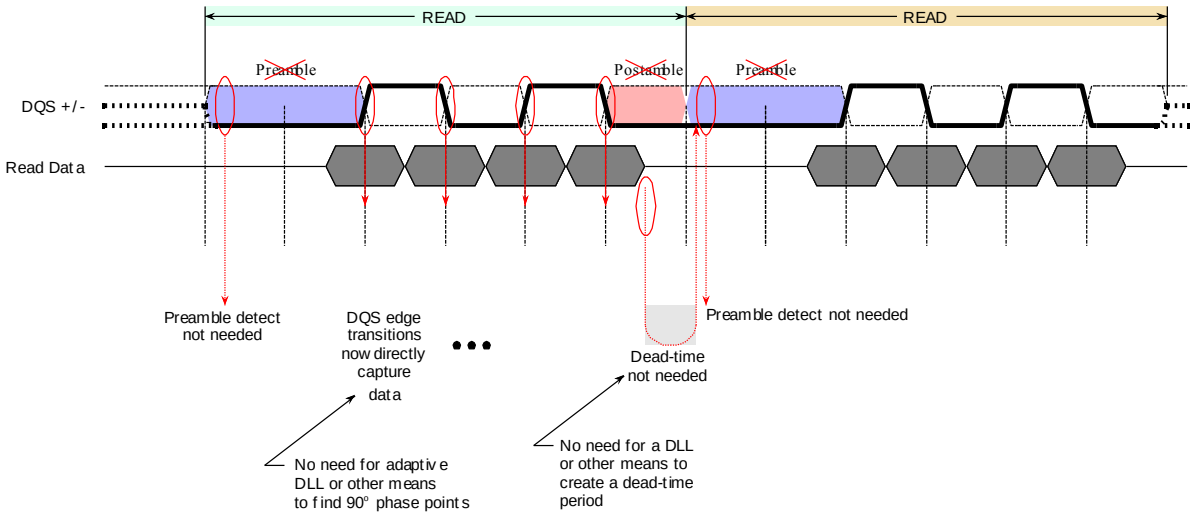

Figure 6. Back-to-back Read Operation Using New Center Aligned DDR Read Protocol

It is obvious from Figure 6 above there are now wasted time slots where there were preamble periods in the legacy protocol.

Figure 7 below shows how the current legacy protocol periods can be collapsed to save time and improve performance with greater data transfer throughput.

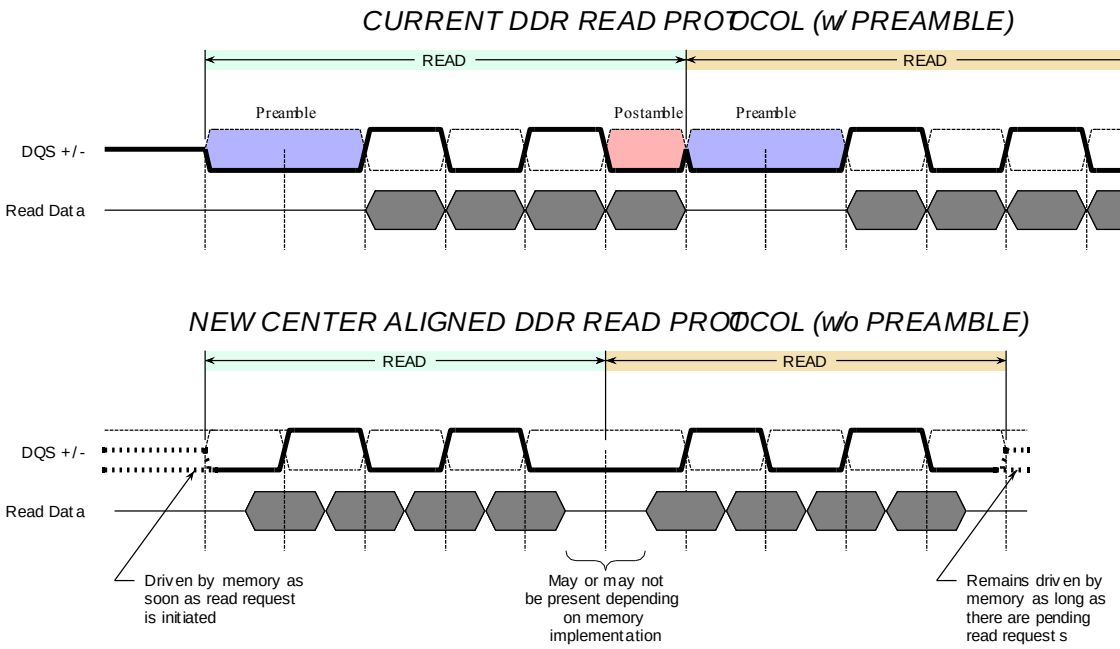

Figure 7. Legacy vs. New DDR Read Protocol That Removes Need For Preambles

The actual performance improvement gained with the new center aligned DQS protocol will depend on how often read bursts and single read transfers occur in a system's operation, but is nevertheless a significant improvement over the current legacy read protocol.

4 Perspectives on the New Center Aligned DDR Read Strobe Protocol

There are three primary perspectives from which to view the new center aligned read strobe protocol, customer, host controller design, and memory device design.

From a customer standpoint, the new protocol requires the customer route another clock with similar delay characteristics as the existing memory clock between the host controller and memory device(s). Doing this allows the customer to benefit from better system reliability and operating margins due to less complexity in the host memory controller. This translates into data error reduction at higher speeds, provides easier system debug, lowers power required in the host controller for operation of these functions, and perhaps most importantly, provides higher data transfer throughput. It is probably also the case that because of these simplified functions and symmetric protocol, higher frequency memory controllers and memory devices can be had yielding higher performance systems.

From a host controller viewpoint, the new protocol requires an additional pad(s) for the phase shifted clock with matching delay characteristics relative to the existing memory clock. But it can eliminate engineering resources required to migrate specialized custom hard macro PHY memory interfaces between technologies since this new protocol should afford automated tool synthesis of the logic to provide the necessary functions. The memory host controller no longer needs adaptive read clock recovery logic to generate an internal read clock to capture data, or alternately clock over-sampling with periodic software learning because the read data will be source synchronous with the data strobe and center aligned within data valid windows just as it is for write operations. PVT variations are mitigated in this manner. With this new protocol there is no need to count clocks to determine when to start looking for a preamble, DQS and data, i.e. handling differing memory response times are now trivial. It eliminates certain difficult cases related to detection of preamble and postamble time slots and eliminates the need for an additional DLL function required to handle these special cases. And most importantly, with no need for preambles, data transfer throughput and performance are increased. Custom specialized IO and control for data strobe preamble and postamble detection and allowance are also eliminated affording use of simple IO transceivers. All this speeds time-to-market, provides easier system debug, reduces design time and engineering costs (host die area especially related to on-die testability and test costs) improves margins and favors a competitive advantage.

From a memory vendor standpoint, the new protocol requires an additional pad(s) for the phase shifted clock and care within the memory device to keep the temporal linkage between the memory clock used to stage data and the phase shifted clock used to generate the data strobe. Most of the existing logic for these operations inside the memory device should be reusable by only changing the clock source for the logic staging out the data strobe. There is no need for a preamble or postamble which improves data transfer throughput and thereby also increases performance. Interestingly, the same memory die can be used for both the legacy protocol and the new protocol. That is, provided the design is done correctly if the phase shifted clock used to generate DQS is connected to the 0^{o} memory clock, DQS will revert to the legacy protocol. This saves costs by having one die that serves multiple purposes.

5 Devices Using This New Protocol Are Already Available

Freescale Semiconductor currently offers two devices that incorporate this new protocol for QSPI Flash memories.

The MAC57D5xx (offered in 4M, 3M, & 2M configurations; 4M part name is MAC57D54H) is currently sampling in the 208LQFP and 516MAPBGA with the 176LOQFP available in July 2015. Full AEC qualification is expected by the end of November 2015.

The MAC58Rxx has been sampling in the 529 MAPBGA since 2014 and is also planned to complete AEC qualification by the end of 2015.

Volume production for both of these products is expected in the latter half of 2016.

6 Summary and Conclusions

Introduced in this paper is an improved and simplified protocol for DDR memory that can be applied to Flash, SDRAM or other DDR-type interfaces. With this new protocol the direction of data transfer, read or write, is now consistent and symmetric with read data strobe edges center aligned to valid data windows maintaining source synchronicity across PVT variations as it is for write transfers to memory. Important beneficial side-effects are realized with the simplified protocol. When the host memory controller does not need to use an over sampling learning method or a read clock recovery system to properly capture read data in the memory controller because the data strobe is now directly linked to the center of the data valid window and source synchronous from the memory device, the detrimental effects of PVT variations do not require expensive and complex adaptive clock recovery logic to mitigate these variations.

Overall system complexity is greatly reduced, which translates into more reliable and robust systems. Because the new protocol eliminates the need for post- and preamble time slots, greater data transfer rates are achievable which translate into higher performance systems. Memory vendors also benefit because the same memory device die can be used for the legacy protocol and new protocol versions. Customers realize benefits for all of the above, especially reduced overall complexity, easier system debug, and more reliable, higher throughput and performance systems.

The new DDR memory interface using a centered aligned read strobe protocol that maintains consistency between read and write data operations presented here has many benefits for customers and designers alike. Gaining acceptance of this new protocol between memory vendors by creating compatible memory devices will benefit customers and avoid device fragmentation and market confusion.

7 References

JEDEC Standard JESD79C, *Double Data Rate (DDR) SDRAM Specification*

JEDEC Standard JESD79-2F, *DDR2 SDRAM STANDARD*

JEDEC Standard JESD79-3F, *DDR3 SDRAM STANDARD*

JEDEC Standard JESD79-4A, *DDR4 SDRAM STANDARD*

How to Reach Us

Home Page:
www.freescale.com

Web Support:
www.freescale.com/support

USA/Europe or Locations Not Listed:
Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or +1-480-768-2130
www.freescale.com/support

Europe, Middle East and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064, Japan
0120 191014
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com