

Technical Note

DDR2 Offers New Features and Functionality

Introduction

DDR2 SDRAM introduces features and functions that go beyond the DDR SDRAM specification and enable DDR2 to operate at data rates of 800MHz and above. The advancements of DDR2 are largely the result of changes in DRAM architecture and signaling, as well as additions to the mode register for lower power and improved command/data bandwidth. This article discusses the various the changes and new features and benefits of DDR2 technology.

DRAM Architecture Changes DDR2 Array Definition

Changes to the DRAM architecture include shortened page size for reduced activation power, burst lengths of four and eight for improved data bandwidth capability, and the addition of eight banks in 1Gb densities and above.

DDR2 Page Size

Page size is important to the power consumption of a DRAM device. Page size is defined as the minimum number of column locations that are accessed with a single ACTIVATE command. This is calculated as the number of column locations times the number of DQs on the DRAM. For example, the 512Mb, x8 DDR2 SDRAM has 1,024 column locations, so the page size is 1,024 columns times 8 DQs, which equals 8,192 bits. Dividing the 8,192 bits by a word length of 8 equals 1,024 bytes or 1KB.

Each time an ACTIVATE command is given, all bits in the page are read by the sense amplifiers and restored to the correct value. Because this process is a major contributor to the active power, a device with a shorter page size has a significantly lower operating current. The page size of the 512Mb DDR2 device is half that of DDR SDRAM. Therefore, the current consumed by the ACTIVATE commands is greatly reduced, resulting in improved IDD0, IDD1, and IDD7 specifications. (See Table 1 for a comparison of DDR2 and DDR page sizes and banks.)

Table 1: DRAM Architectures

Density		DDR	DDR2
256Mb	Page Size	1KB	1KB
	Banks	4	4
512Mb	Page Size	2KB	1KB (x4, x8), 2KB (x16)
	Banks	4	4
1Gb	Page Size	2KB	1KB (x4, x8), 2KB (x16)
	Banks	4	8
2Gb	Page Size	–	1KB (x4, x8), 2KB (x16)
	Banks	–	8

With the higher-density DDR2 SDRAM (512Mb, 1Gb, and 2Gb), the page size increases on the x16 device. The x4 and x8 devices continue to operate at lower activate currents with a 1KB page size, similar to the 256Mb DDR2 SDRAM. However, the x16 DDR2 page size increases to 2KB. While power consumption increases on the x16 DDR2 SDRAM, system power is typically no higher than that of the x8 device. A 64-bit bus requires eight x8 DRAM devices to support the bus width, while only four x16 DRAM devices are required. Thus, the activate current of a x16 component may be somewhat higher, but the overall system power consumption is still lower.

DDR2 Bank Accesses

Another change from DDR to DDR2 (as shown in Table 1) is the number of banks on the larger-density DDR2 devices. A typical four-bank access, which is applicable to the 256Mb and 512Mb DDR2 SDRAM, is shown in Figure 1. Bank accesses for DDR2 devices are similar to DDR, except ACTIVATE command spacing is 7.5ns for 1KB and 10ns for 2KB page sizes, due to the increased power consumption of the larger page size. When a bank is precharged, it may be activated again after t_{RC} is met, following the original ACTIVATE command.

The larger-density 1Gb and 2Gb DDR2 devices use eight banks, which increases flexibility in DRAM accesses. However, it also increases the power that must be delivered to the DRAM. To keep the power delivery from adversely affecting the functionality of the DRAM and the memory subsystem, an additional restriction (t_{FAW}) is applied to bank ACTIVATE commands, as shown in Figure 2. ACTIVATE-to-ACTIVATE spacing remains t_{RRD} , as in the four-bank case; however, there is an additional restriction within any t_{FAW} window where no more than four banks may be activated. This requires a short amount of extra time prior to activating the fifth bank.

Figure 1: ACTIVATE for Four-Bank DDR or DDR2 Devices

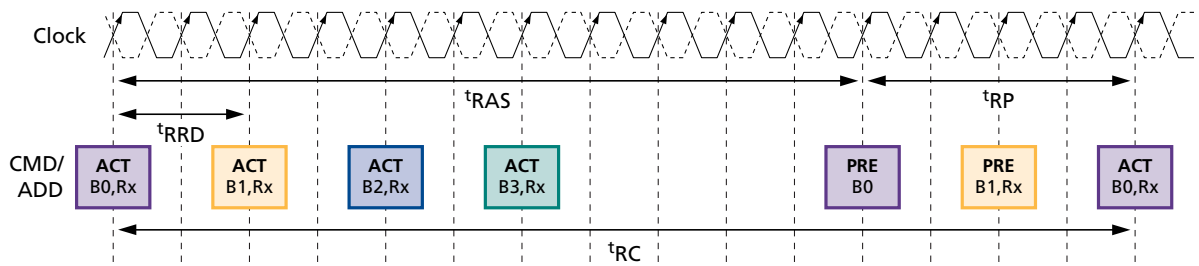
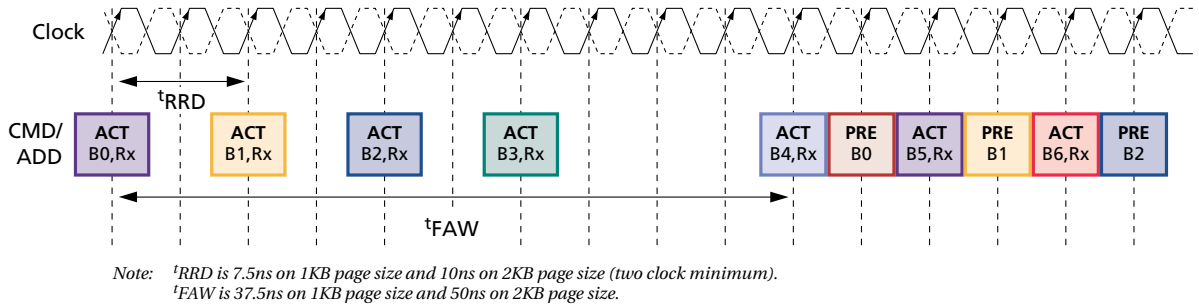


Figure 2: ACTIVATE for Eight-Bank DDR or DDR2 Devices



Signaling Features On-Die Termination

On-die termination (ODT) has been added to the DDR2 data signals to improve signal integrity in the system. Figure 3 shows a functional diagram of the on-die termination. The termination value of R_{TT} is the Thevenin equivalent of the resistors that terminate the DQ inputs to V_{SSQ} and V_{DDQ} . An ODT pin is added to the DRAM so the system can turn the termination on and off as needed. In a simple system with one DRAM load per DQ signal, the termination is turned on for WRITES and disabled for READS.

Figure 3: Functional Diagram for On-Die Termination

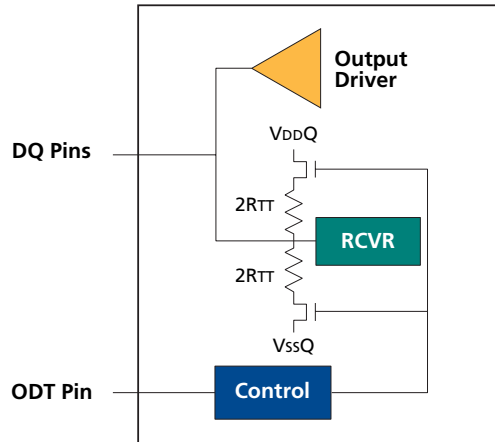
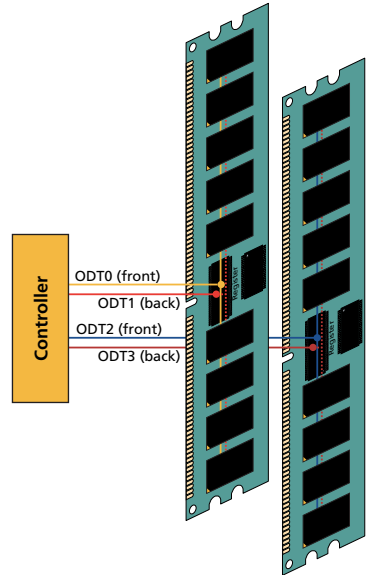


Figure 4 shows the ODT trace routing on a typical PC system with two memory slots. Each slot contains two ranks of memory (front and back). A separate ODT signal is routed to each rank of memory to provide flexibility in power management and the best system performance. The routing of ODT signals is similar to the routing of CS# signals.

Figure 4: Typical 2-Slot ODT Routing



ODT WRITES

Figure 5 shows a typical termination scheme for write data in a 533 Mbit/s, two-module system. If only one module is loaded in the system, a 150 ohm equivalent termination (300 ohm pull-up/300 ohm pull-down) must be enabled on the module during the WRITE. If two modules are loaded into the system, the termination on the module not receiving the write data should be enabled at 75 ohm equivalent termination resistance (150 ohm pull-up/150 ohm pull-down). This termination scheme provides optimal trade-offs between signal integrity and voltage swing.

The equivalent comparison for write data at 533 Mbit/s is shown in Figure 6. Neither bus has significant overshoot. However, the voltage margin on the SSTL bus is not sufficiently outside the DRAM input voltage requirements. Thus, when the input slew rate slows down in the two-slot system, the data eye that the DRAM must capture is significantly reduced. The ODT system shows increased voltage margin, and the data eye remains significantly larger, even under heavily loaded conditions.

Figure 5: Typical ODT Configuration for DRAM WRITES

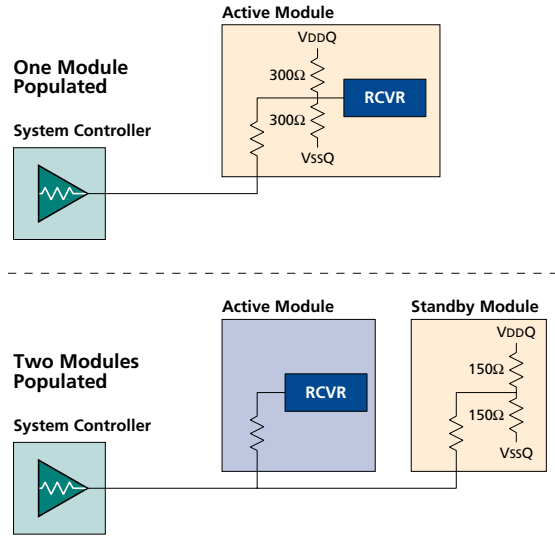
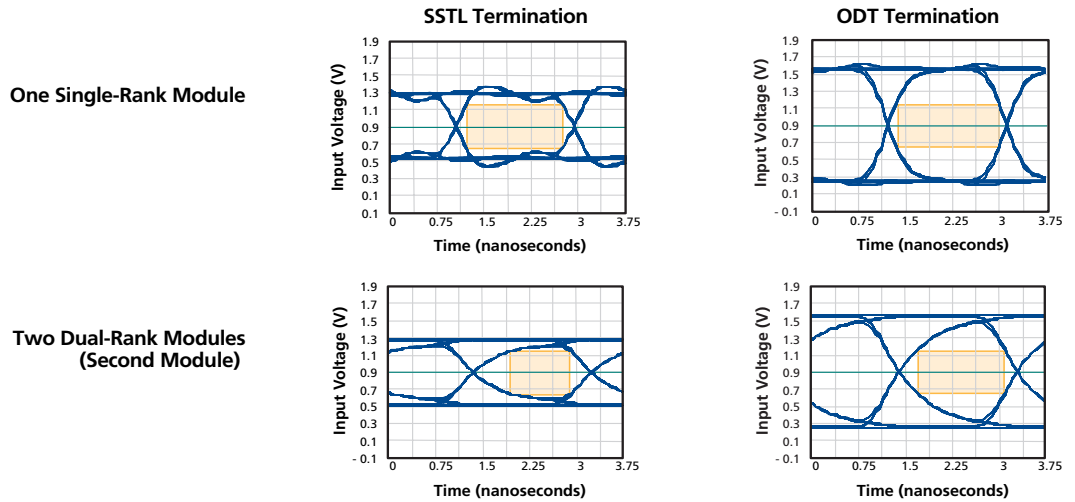


Figure 6: ODT vs. SSTL Termination Comparison for WRITES



ODT READS

Figure 7 shows a typical termination scheme for read data in a 533 Mbit/s, two-module system. For READs, on-die termination must be included on the controller. If only one module is loaded in the system, a 75 ohm termination should be enabled on the system controller. If two modules are loaded, the controller termination should be increased to 150 ohms and the unused module should terminate at 75 ohms. This configuration provides good voltage margin at the controller.

Figure 8 shows a comparison of a DDR2 DQ bus for a typical SSTL-terminated, two-module system using ODT termination versus a non-ODT system reading data at 533 MT/s. The most obvious difference between the two buses is the reduction in the overshoot on the data bus. The overshoot reduction significantly improves the voltage margin and reduces inter-symbol interference (ISI). The second major improvement using DDR2 ODT is the increase in the slew rate in the two-module case. The SSTL bus

starts “laying over” when more loads are added, which detracts from the timing margin in the system. Both of these improvements dramatically increase the size of the data eye that the system controller must capture to read data from the DRAM.

Overall, ODT implementation provides improved signal integrity, data eyes, and voltage margin for DDR2 to enable systems to attain higher data rates. At the same time, it is an option and can be disabled for systems that require a more traditional SSTL-type of bus.

Figure 7: Typical ODT Configuration for DRAM READS

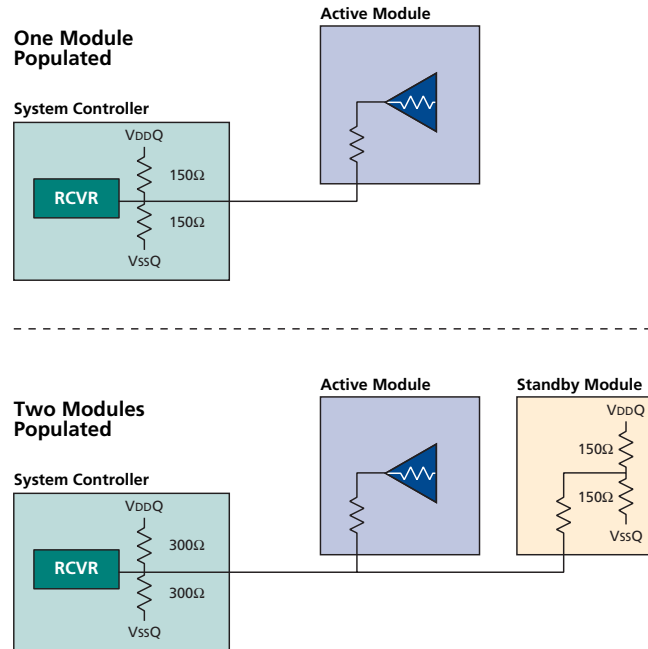
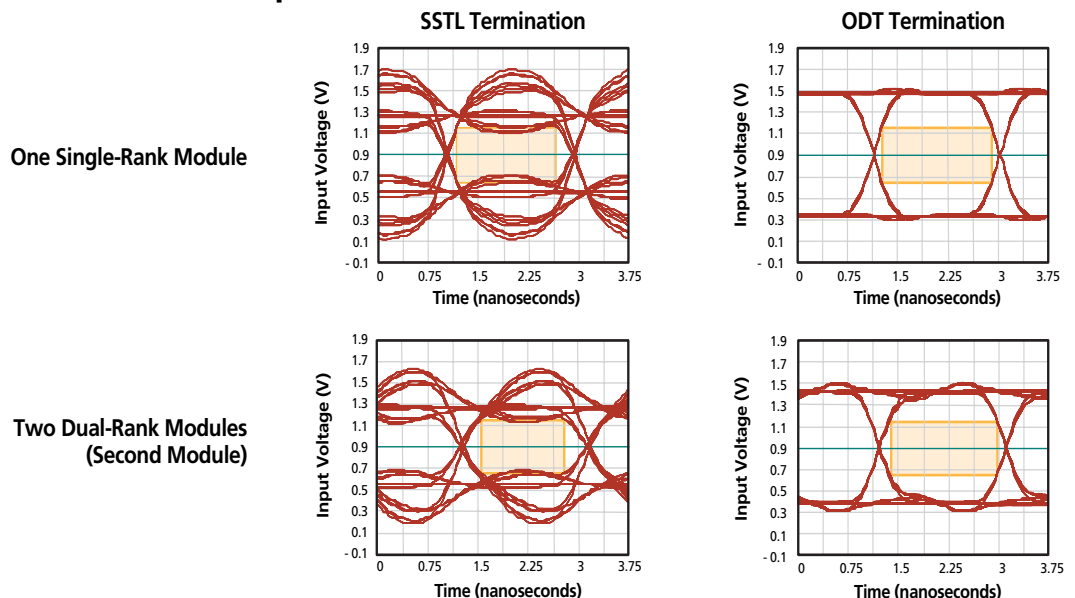


Figure 8: ODT vs. SSTL Termination Comparison for READS



Termination Values

As illustrated in Figures 5 and 7, on-die termination can have different values, depending on the system configuration. This is controlled via the extended mode register (EMR). (See Figure 9.) If extended mode bits 2 and 6 are both “0,” the on-die termination is disabled on the DRAM. Three additional options include selecting either an RTT equivalent value of 50 ohms (100 ohms pull-up and 100 ohms pull-down), or 75 ohms (150 ohms pull-up and 150 ohms pull-down) or an RTT equivalent of 150 ohms (300 ohms pull-up and 300 ohms pull-down). The proper termination value is selected to optimize signal integrity at minimum power, as shown in Figures 6 and 8.

Figure 9: Extended Mode Register Control for On-Die Termination

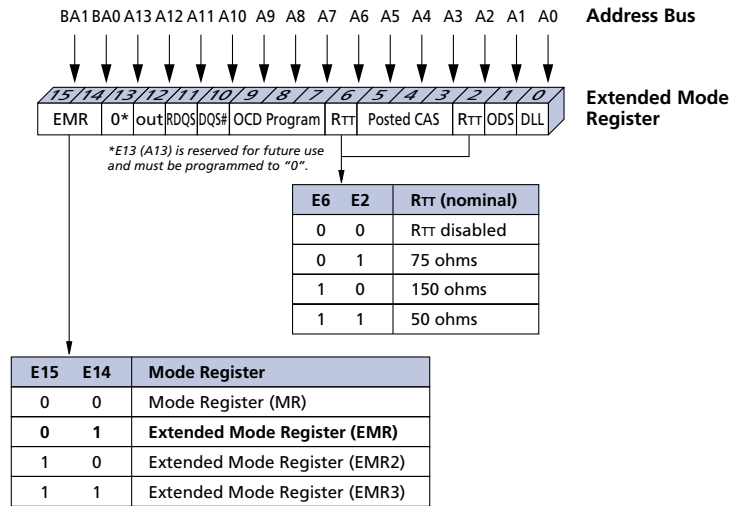
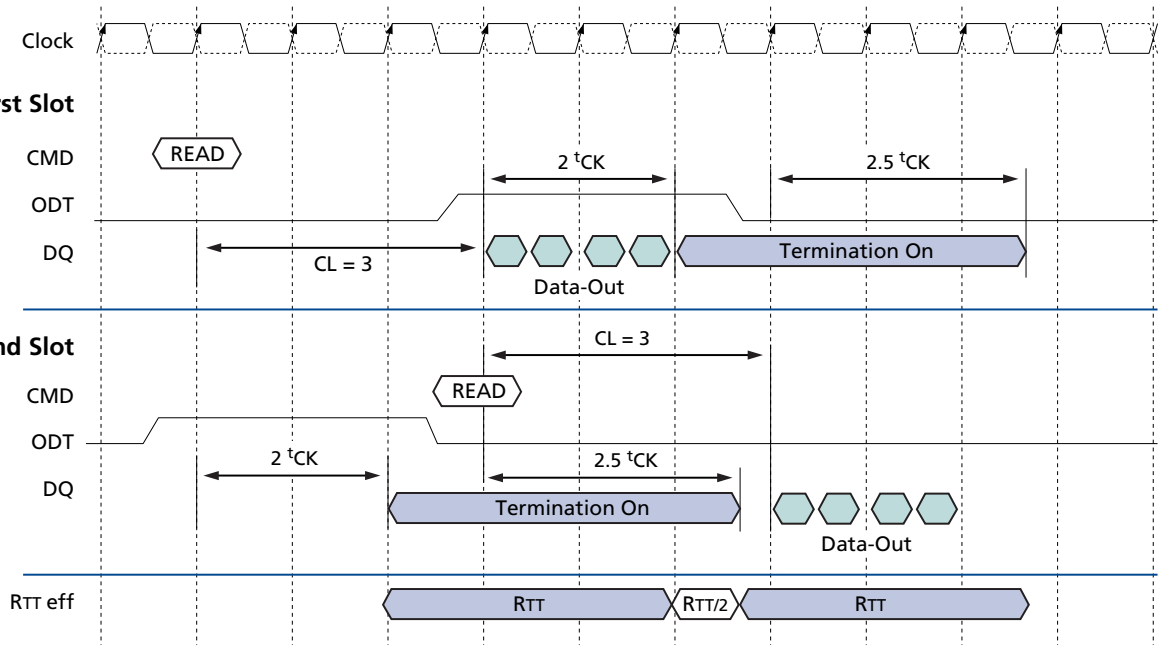


Figure 10: Rank-to-Rank ODT Control – READs



On-Die Termination Control in Active Mode

After the proper value of termination is selected, the system controller must switch the termination on and off at the proper times. While the DRAM is in active or fast exit power-down mode, the termination turns on two clock cycles after the ODT signal is registered HIGH, as shown in Figure 10. For module 2, the ODT signal is latched HIGH on the same clock cycle as the READ command to module 1. Two clock cycles later, the on-die termination on the module 2 DRAM turns on. To schedule on-die termination, ODT should be transitioned HIGH three clocks before the data is on the bus to ensure it is on prior to the preamble.

To disable on-die termination, the ODT signal must be registered LOW by the clock. Then, two and a half cycles later, the termination turns off, as shown in Figure 10. The extra half cycle allows the termination to remain enabled until after the data has stopped driving on the bus. Thus, the ODT pin can be registered LOW two clock cycles prior to the end of the data burst on the bus.

On-Die Termination Control in Power-Down

Asynchronous operations are also associated with the ODT pin when the DRAM is in slow exit active power-down or precharge power-down mode. During this operation, the turn on and off time is an analog delay off the registration of the ODT signal. The timing isn't as critical since data can't be read from or written to a rank that is in power-down mode. Thus, the system only has to enable the termination to form a window around the data from the other ranks. See the DDR2 data sheet for additional information on specific ODT timings and restrictions.

On-Die Termination Control in Self Refresh

During self refresh, the DRAM disables on-die termination and ignores the ODT pin, regardless of the values programmed into the EMR. A memory subsystem that is in self refresh mode cannot read or write data, so disabling on-die termination does not affect signal quality, and it allows the DRAM to minimize power consumption.

Off-Chip Driver (OCD)

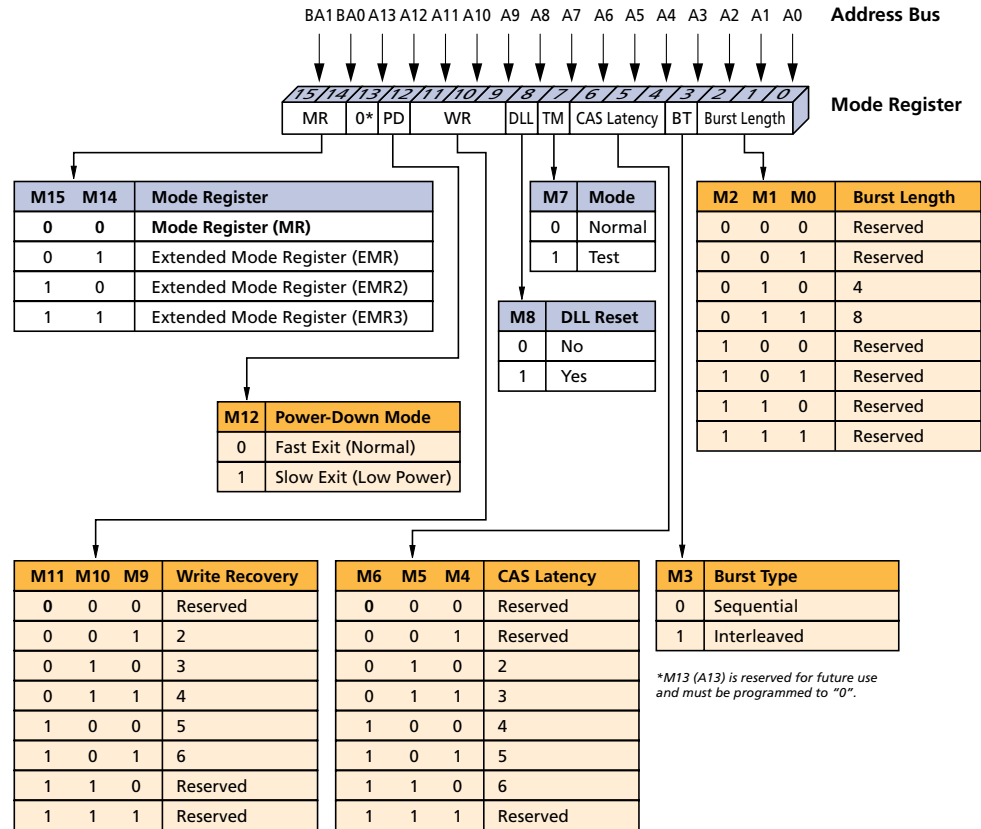
The Off-Chip Driver (OCD) function is no longer required on DDR2 devices and is not supported on Micron devices. During initialization, the OCD mode register should be enabled to its default setting. Please reference specific data sheets for the initialization sequence instructions.

Additional Mode Register Changes Mode Register (MR)

Some of the additional features and functionality that have been added to DDR2 require changes in the MR and the EMR.

Figure 11 highlights the changes made in the MR for DDR2. These include support for burst lengths of 4 and 8 only, a modification of the sequential interleave mode, and extension of the CAS latencies supported. In addition, a low-power definition for active power-down mode and a write recovery parameter have been added.

Figure 11: Mode Register Changes for DDR2



CAS Latency

DDR2 operates at higher clock frequencies than DDR SDRAM. Figure 13 illustrates the difference between a DDR READ at a 266 MHz data rate and a DDR2 READ at a 533 MHz data rate. The latency of each device is the same, approximately 15ns. For the DDR SDRAM, this is equivalent to two clocks; for DDR2, four clocks.

As clock rates increase, the number of latencies possible must also increase. Figure 12 shows the latencies supported for various clock rates and CAS latencies. The DRAM latency is calculated as $t_{CK} \times CAS \text{ latency supported}$. DDR2 is targeted at latencies of 12ns–20ns. If the DRAM latency is beyond this range, it is unlikely that a system will utilize that latency due to lost system performance. Likewise, it is unlikely the DRAM will meet latencies below 12ns due to DRAM core performance. To support DDR2's higher clock rates, the CAS latency programmability has been set at three to six clocks. Half clock latencies are not supported.

Figure 12: DDR2 READ Latencies

	READ Latency (clocks)				
	RL t _{CK}	3	4	5	6
DDR2-400	5ns	15ns	20ns	25ns	30ns
DDR2-533	3.75ns	11.25ns	15ns	18.75ns	22.5ns
DDR2-667	3ns	9ns	12ns	15ns	18ns
DDR2-800	2.5ns	7.5ns	10ns	12.5ns	15ns

Unlikely operation

Note: Refer to individual datasheets to determine which latencies are supported for each speed grade and part number combination.

WRITE Latency

DDR2 also includes a modification to the WRITE latency specification. Figure 14 illustrates the differences between DDR and DDR2 SDRAM. DDR SDRAM has a WRITE latency of one clock, which means that approximately one clock after a WRITE command, the DRAM starts clocking in data. However, given the increase in clock rates, this time is very short, making it difficult for the DRAM to prepare to capture data. DDR2 enables the WRITE latency to track the CAS latency minus one clock.

Figure 14 shows a 533 MHz DDR2 device with CAS latency set to four clocks. The WRITE latency is set to $CAS \text{ latency} - 1 \text{ clock}$, which is three clocks.

The modification in WRITE latency allows more time for DRAM input buffers to support higher clock frequencies. Additionally, tracking WRITE latency with CAS latency simplifies system read-to-write timings.

Figure 13: DDR2 vs. DDR READ

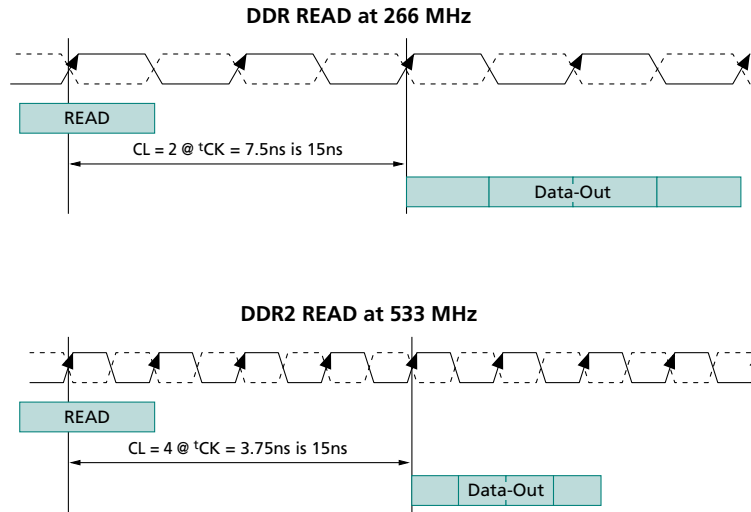
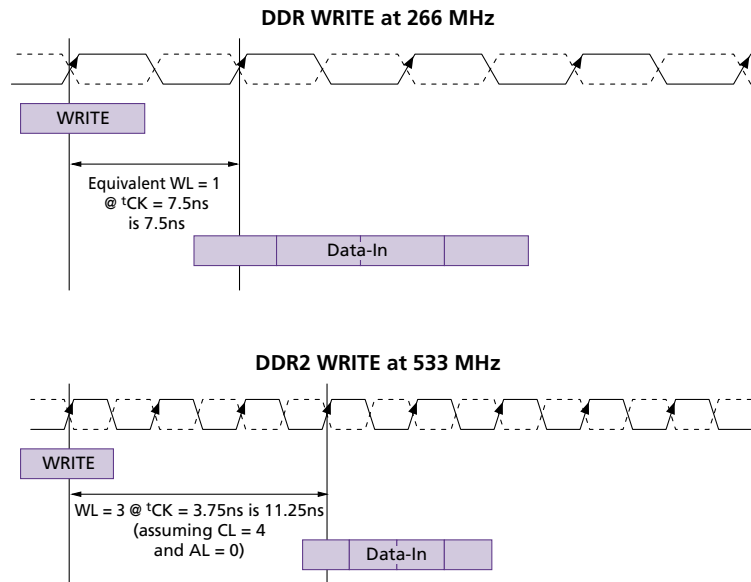


Figure 14: DDR2 vs. DDR WRITE



Sequential Interleave

To support the high data rates of DDR2, the DRAM must have the flexibility to prefetch four bits of data from the core instead of only two. Consequently, support for a burst length of two has been dropped. Prefetching four bits, however, makes it difficult to support the traditional sequential interleave mode for DRAM burst lengths of eight.

The sequential burst orders for DDR2 are shown in Figure 15. For burst lengths of four, there is no change from the DDR specification. However, for burst lengths of eight, a new burst type called sequential nibble is implemented. During sequential nibble operation, the DRAM breaks the eight-bit burst into two separate four-bit nibbles. The sequential

interleave is executed on the nibble with the starting column location, followed by a similar sequential interleave on the other nibble using the starting location that is equal to the lower bits of the starting column.

Using the nibble sequential burst, a system requiring sequential interleave can take advantage of the eight-bit burst and still use any starting column within the burst.

The interleave burst ordering for DDR2 is the same as DDR.

Figure 15: Sequential Burst Ordering

Burst Length	Starting Column (A2, A1, A0)	Order of Accesses Within a Burst Burst Type = Sequential	
		DDR	DDR2
4	0 0 0	0,1,2,3	0,1,2,3
	0 0 1	1,2,3,0	1,2,3,0
	0 1 0	2,3,0,1	2,3,0,1
	0 1 1	3,0,1,2	3,0,1,2
8	0 0 0	0,1,2,3,4,5,6,7	0,1,2,3,4,5,6,7
	0 0 1	1,2,3,4,5,6,7,0	1,2,3,0,5,6,7,4
	0 1 0	2,3,4,5,6,7,0,1	2,3,0,1,6,7,4,5
	0 1 1	3,4,5,6,7,0,1,2	3,0,1,2,7,4,5,6
	1 0 0	4,5,6,7,0,1,2,3	4,5,6,7,0,1,2,3
	1 0 1	5,6,7,0,1,2,3,4	5,6,7,4,1,2,3,0
	1 1 0	6,7,0,1,2,3,4,5	6,7,4,5,2,3,0,1
	1 1 1	7,0,1,2,3,4,5,6	7,4,5,6,3,0,1,2

Different

Active Power-Down Mode

Active power-down mode enables a system designer to optimize power and performance. In this mode, the transition timing between power-down exit and the next command is dependent upon the power-down mode bit.

If the power-down mode is set to fast exit mode, there is a two-clock delay from active power-down exit to any command (including READ). Alternatively, the power-down mode bit may be set to slow exit mode, which extends the power-down exit delay to six clocks, minus the additive latency to a READ command. Active power-down slow exit is the same as the precharge power-down time during all operations of the DRAM. The power-down mode bit is ignored in precharge power-down. (See Table 2.)

Write Recovery

Write recovery (WR) must be correctly entered into the MR to ensure proper DRAM operation. Write recovery is measured from the end of write data to the time when the internal auto precharge operation starts (see Figure 16). To correctly set the MR, WR is calculated as t_{WR}/t_{CK} , rounded up to the next integer.

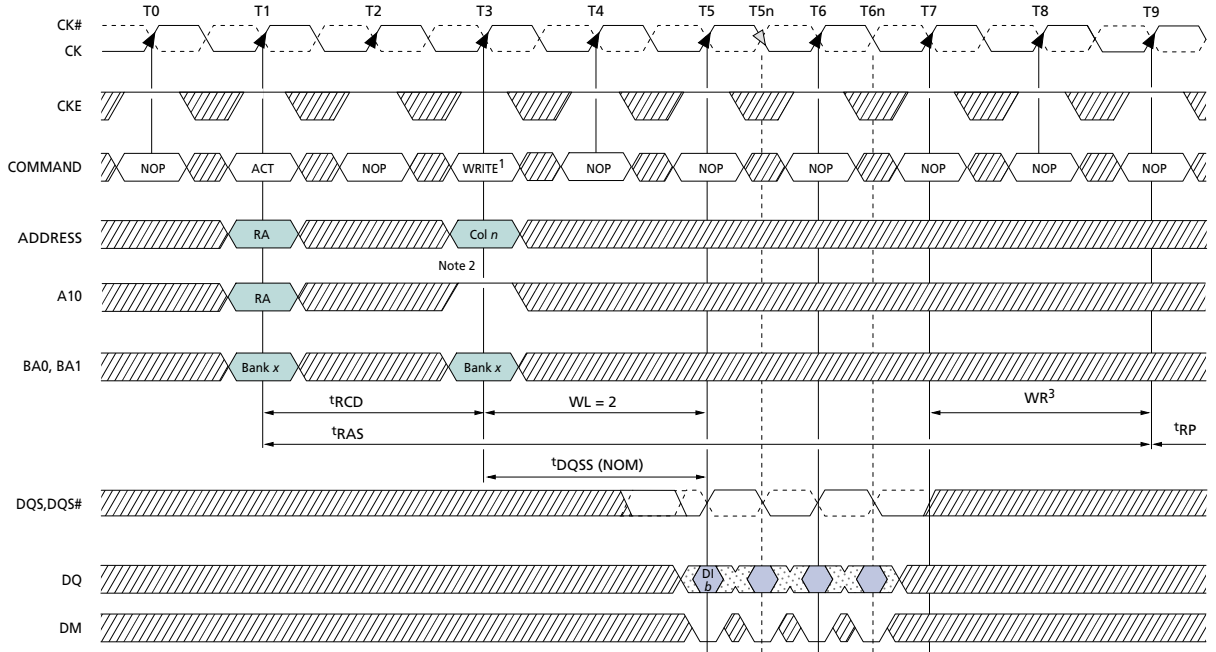
Extended Mode Register (EMR)

Additional enhancements have been made to the EMR, as shown in Figure 17. ODT was discussed previously, but other enhancements include additive latency, output disable, and DQ strobe configurability.

Table 2: DDR2 Active Power-Down Mode Options

	MR (bit 12)	PD Mode	Delay from CKE HIGH to Command
Active Power-Down	0	Fast Exit	DDR2-400, $t_{XARD} = 6\text{-AL}$ in clocks DDR2-533, $t_{XARD} = 6\text{-AL}$ in clocks DDR2-667, $t_{XARD} = 7\text{-AL}$ in clocks DDR2-800, $t_{XARD} = 8\text{-AL}$ in clocks
	1	Slow Exit	$t_{XARD} = 2$ clocks for all speeds
Precharge Power-Down		Ignored	2 clocks to any non-READ command <i>Additive latency – 6 clocks to READ</i> need to add 667 and 800 specs

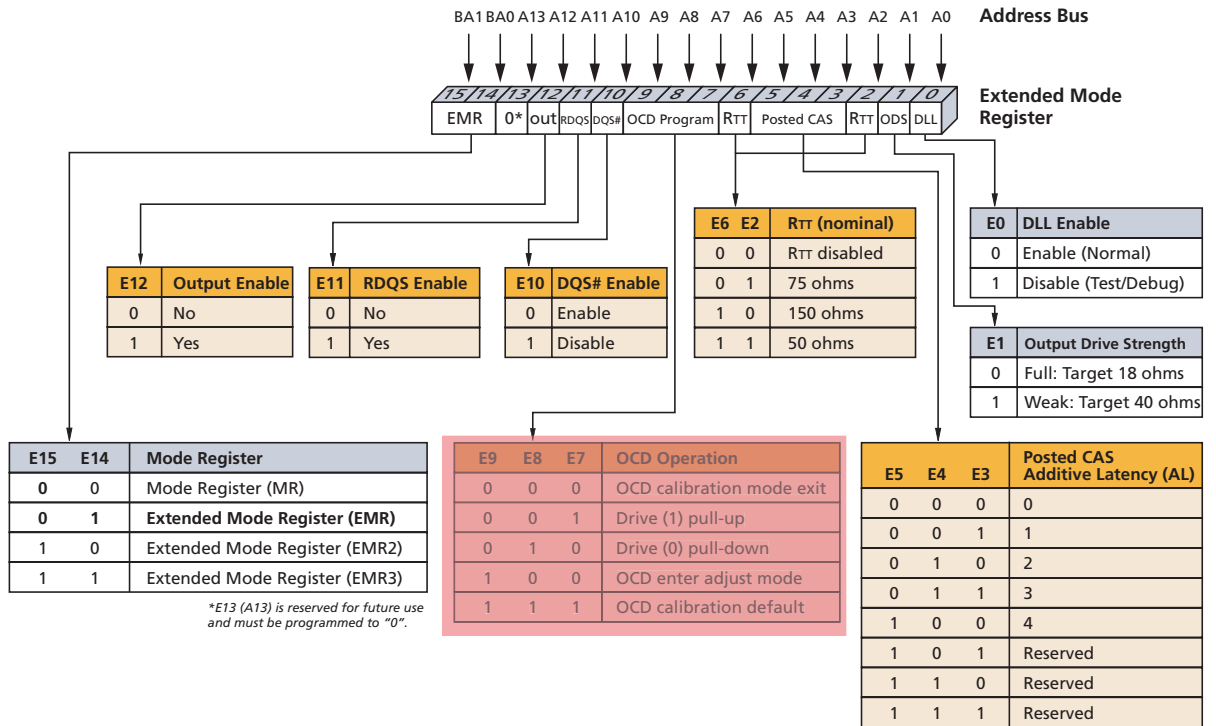
Figure 16: DDR2 Write-to-Precharge Timing



- NOTE:** 1. Burst length = 4, additive latency = 0, and WRITE latency = 2.
 2. Enable auto precharge.
 3. WR is programmed via EMR [11,10,9] and is calculated by dividing t_{WR} (ns) by t_{CK} and rounding up to the next integer value.

TRANSITIONING DATA
 DON'T CARE

Figure 17: Extended Mode Register



Additive Latency

Figure 18 shows an example of DDR2 bank interleave READs. During these accesses, a bank is opened with an ACTIVATE command. After t_{RCD} is met, a READ with auto precharge command is issued and data is read out of the DRAM. As seen in Figure 18, there are times when a command cannot be issued at the optimal location because there is a command slot conflict. For example, ACT (to bank 2) must be delayed by one clock because the RD_AP (to bank 0) is already in that spot. The result is a gap in the output data.

To improve system scheduling, DDR2 includes a mode register for additive latency. Using additive latency, the DRAM captures READ or WRITE commands but waits to execute them until the additive latency time expires.

Figure 18: DDR2 READs With no Additive Latency

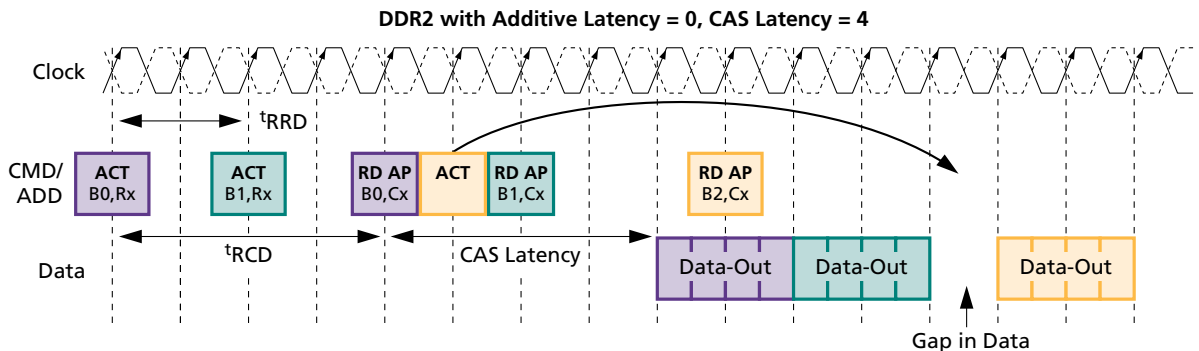


Figure 19 illustrates DRAM operation with two different additive latencies. The top figure shows an operation similar to DDR SDRAM using $AL = 0$ clocks. The bottom figure shows $AL = t_{RCD} - 1$ clock. In this configuration, the ACTIVATE and READ commands can be issued back-to-back. The DRAM holds the READ command for the additive latency time and then executes it on the same clock as the top diagram.

With additive latency set to $t_{RCD} - 1$ clock, a simplified read system timing can be attained. Figure 20 assumes $t_{RCD} = 4$ clocks. Thus, additive latency is set to three clocks. This allows the READ command to directly follow the ACTIVATE command. With a burst length of four, which requires two clocks to burst data, the two clocks for the back-to-back ACT and RD_AP commands match perfectly, enabling data to be continuously output from the DRAM without gaps. The simplicity of system bus scheduling is expected to be particularly popular in applications that have high page misses.

Figure 19: DDR2 With Additive Latency

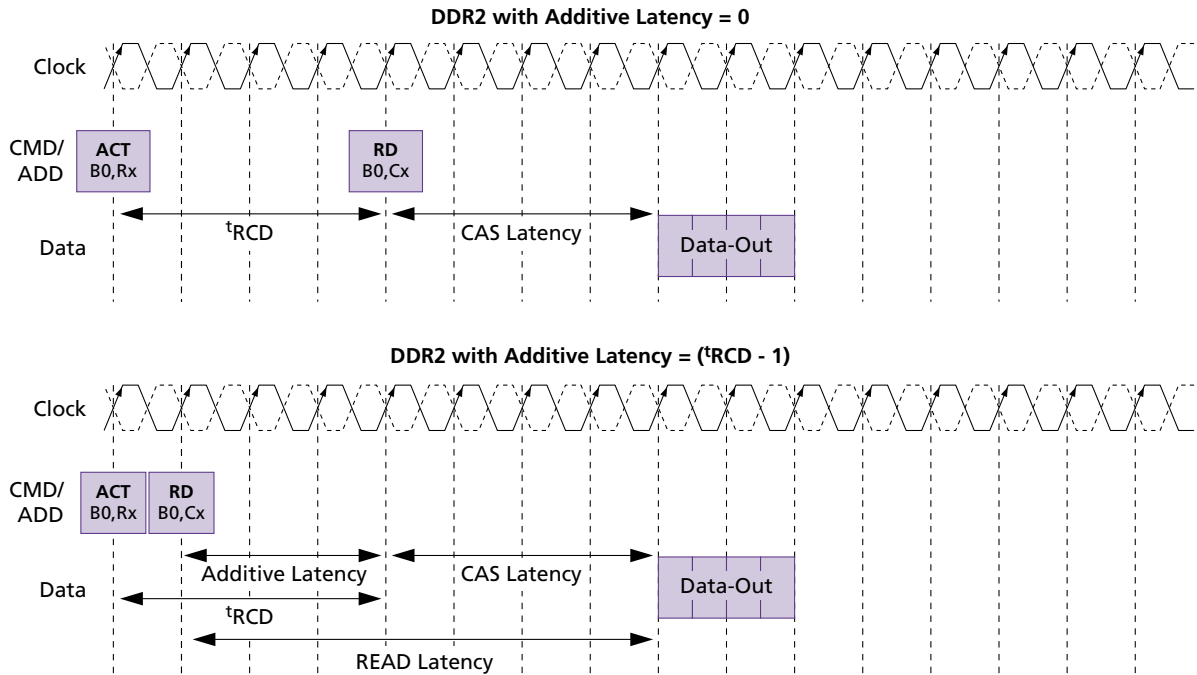


Figure 20: DDR2 READs Using Additive Latency

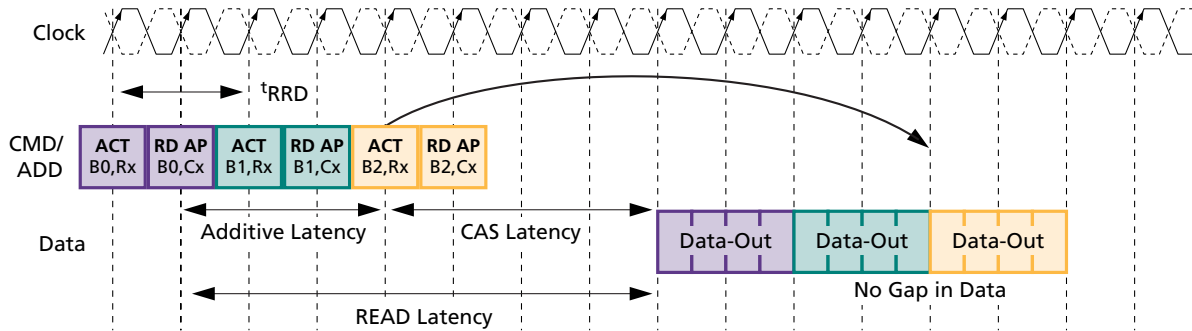
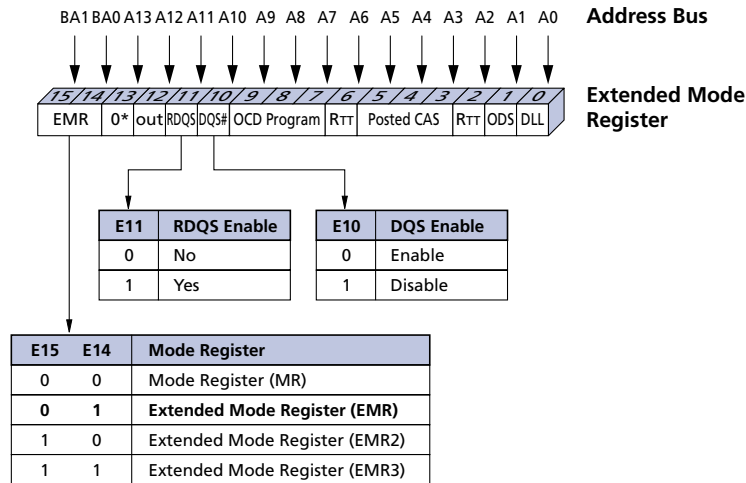


Figure 21: Extended Mode Register for DQS Configuration



DQS Configuration

The EMR also allows configuration of the DQ strobes for multiple system configurations. Figure 21 highlights the areas of the extended mode register that configure the strobes. Bit 10 switches the strobes between differential and single-ended mode, and bit 11 can be used to enable the redundant strobe on the x8 device.

Single/Differential Data Strobe

DDR2 supports single-ended strobes in the same manner as DDR SDRAM, as shown in Figure 22.

DDR2 SDRAM also supports differential strobes. If differential strobes are enabled, DQS operates the same as in single-ended mode with the addition of DQS# operating as the complement in differential signals. For READs the DRAM will output both signals, as shown in Figure 23. For WRITEs, the system must provide both input signals.

The use of differential strobes is system-dependent and may enable improved system timings due to reduced crosstalk and less simultaneous switching noise on the strobe output drivers.

Figure 22: Single-Ended Strobe

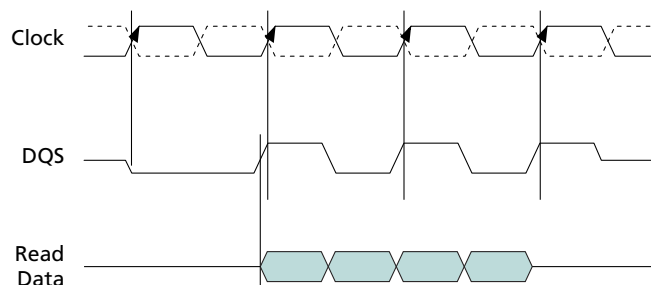
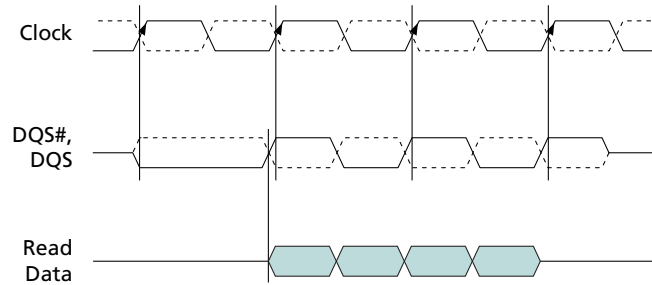


Figure 23: Differential Strobe



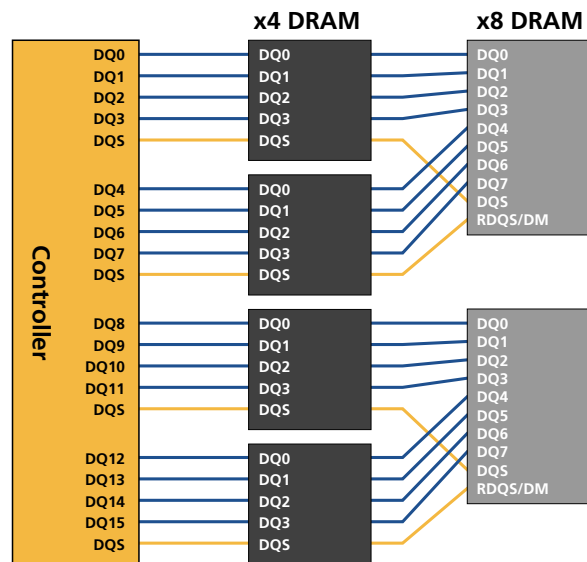
Redundant DQS (RDQS)

The RDQS option is only available on the x8 DRAM and is required when mixing x4 and x8 DRAM devices in the same system. An example of this is shown in Figure 24. Because the x8 DRAM has only one DQS input for the eight-bit byte and the x4 has a DQS for each four-bit nibble, the loading of the DQS lines is different when the devices are mixed in the same system.

RDQS reduces this effect. If devices are mixed in the system, the signal routing from the second x4 DRAM is routed to the RDQS pin on the x8 DRAM. During WRITES, the loading and termination on the RDQS is the same as the DQS signals, so all signaling is matched for the x4 DRAMs. For READS, the x8 DRAM drives the RDQS output the same way it does the DQS. Thus, the controller can use RDQS to capture read data on a nibble device and the x8 DRAM emulates the functionality of the x4.

Systems mixing x4 and x8 devices can utilize RDQS signals to improve timing. RDQS functionality is only an option on the x8 DRAM. Additionally, if RDQS is enabled, the data mask functionality of the x8 DRAM is disabled.

Figure 24: x4 and x8 Devices in Same System



Output Disable

For DDR SDRAM, it can be difficult to measure I_{DD} values in system environments. All I_{DD} values are specified with $I_{OUT} = 0mA$, which is challenging to accomplish in the system. Thus, DDR2 has an option to disable the output drivers from the EMR. If the output disable option is set, the last stage of the output driver is disabled. WRITES are completed normally and READs execute except for the final output stage. Thus, I_{DD} measurements within a system environment can be made.

Conclusion

DRAM architecture changes enable twice the bandwidth without increasing the demand on the DRAM core and while keeping power low. The evolutionary changes enable DDR2 to operate between 400 MHz, 533 MHz, 667 MHz, and 800 MHz. A summary of the functionality changes is shown in Table 3 on page 20.

Modifications to the DRAM architecture include shortened row lengths for reduced activation power, burst lengths of four and eight for improved data bandwidth capability, and the addition of eight banks in 1Gb densities and above. New signaling features include ODT and OCD. ODT provides improved signal quality due to better system termination on the data signals. Modifications were also made to the mode register and extended mode register, including CAS latency, additive latency, and programmable data strobes.

Table 3: Feature Overview

Feature/Option	DDR	DDR2
Data Transfer Rate	266, 333, 400 MHz	400, 533, 667, 800 MHz
Package	TSOP and FBGA	FBGA only
Operating Voltage	2.5V	1.8V
I/O Voltage	2.5V	1.8V
I/O Type	SSTL_2	SSTL_18
Densities	64Mb–1Gb	256Mb–2Gb
Internal Banks	4	4 and 8
Prefetch (MIN Write Burst)	2	4
CAS Latency (CL)	2, 2.5, 3 clocks	3, 4, 5, 6 clocks
Additive Latency (AL)	no	0, 1, 2, 3, 4 clocks
READ Latency	CL	AL + CL
WRITE Latency	fixed	<i>READ latency - 1 clock</i>
I/O Width	x4/ x8/ x16	x4/ x8/ x16
Output Calibration	none	none
Data Strobes	bidirectional strobe (single ended)	bidirectional strobe (single ended or differential) with RDQS
On-Die Termination	none	selectable
Burst Lengths	2, 4, 8	4, 8



8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900

prodmktg@micron.com www.micron.com Customer Comment Line: 800-932-4992

Micron, the M logo, and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners.