TECHNICAL NOTE

UPDATED JEDEC DDR2 SPECIFICATIONS

Introduction

DDR2 production ramps are accelerating, and systems containing initial DDR2-400 and DDR2-533 parts are starting to launch. Having been at the forefront of this technology wave since the beginning, Micron is committed to keeping our customers informed as specifications for this new architecture mature.

Technical Note TN-47-02, "DDR2 Offers New Features and Functionality," provides an excellent overview of the features included in DDR2 memory, including shortened page size for reduced activation power, burst lengths of four and eight for improved data bandwidth, and the addition of eight banks in densities of 1Gb and above. Designers who have not read this article are urged to do so, as it provides a foundation for understanding, along with important details about, new DDR2 functionality.

During the early implementation of the DDR2 architecture, several issues arose that required clarification or redefinition. In the last few months, JEDEC, together with DRAM suppliers, has resolved most of them, including the DDR2-400 and DDR2-533 speed grades. What follows is a discussion of how the initial JEDEC DDR2 requirements described in TN-47-02.

On-Die Termination

During the first stages of DDR2 development, JEDEC identified two values for on-die termination (ODT): 75 ohms and 150 ohms. However, some users requested a third option, 50 ohms. They were willing to trade the increase in power consumption the new specification required for the enhanced timing margin, however minimal it might be.

With a range of 40 ohms to 60 ohms, the 50-ohm ODT support is optional for all JEDEC-defined speed bins. The DDR2-400 and DDR2-533 design specifications do not require it; however, it will be mandated for several DDR2-667 and DDR2-800 platforms, in which the extended mode register's E2 and E6 bits select the 50-bit ODT when both are set to "1," as shown in Figure 1.

Figure 1: Extended Mode Register Control for On-Die Termination

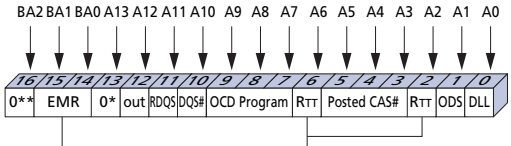

E6	E2	RTT (nominal)
0	0	RTT disabled
0	1	75 ohms
1	0	150 ohms
1	1	50 ohms (optional)

E15	E14	Mode Register Definition
0	0	Mode Register (MR)
0	1	Extended Mode Register (EMR)
1	0	Extended Mode Register (EMR2)
1	1	Extended Mode Register (EMR3)

**E13 (A13) is reserved for future use and must be programmed to "0".*

***E16 (BA2) is reserved for future use and must be programmed to "0".*

Delta VM

Delta VM is the percentage of voltage offset from the ODT's RTT voltage (pull-up resistor and pull-down resistor's voltages) with respect to VDDQ/2. The initial value of ±3.75 percent maximum deviation from VDDQ/2 was increased to ±6 percent to accommodate the third ODT value of 50 ohms. Figure 2 illustrates the range of VM offset allowed.

Figure 2: Delta VM Range

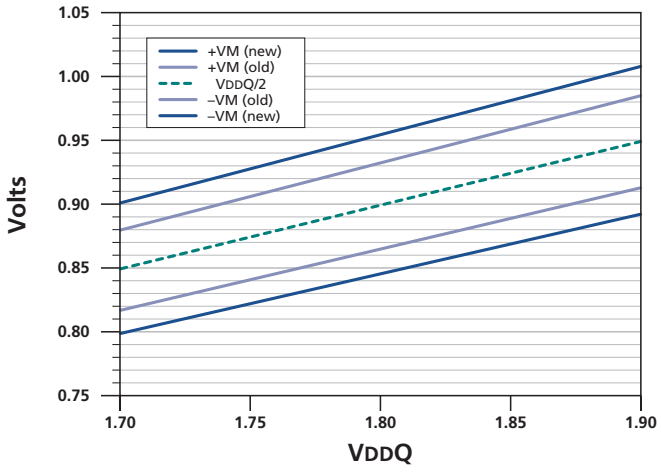

Off-Chip Driver Calibration

The addition of off-chip driver calibration (OCD) was originally intended to reduce the mismatch between the characteristics of the pull-up output driver and the pull-down output driver, thus enabling a tighter tolerance of the output buffer's impedance. In actual practice, however, it created a lot of confusion because users misinterpreted its usage. When system errors involving reference voltages, termination, and comparators are added to the equation, it is almost impossible to achieve the desired overall reduction in variation. In addition, some designers use OCD to move the output driver away from the original target of 18 ohms to compensate for improper termination schemes. In theory, this seems like a good idea, but again, actual practice proves otherwise. When the driver is moved from 18 ohms, the timing specifications become invalid and can no longer be guaranteed.

As a result, OCD was found to be more of a liability than an asset in system production usage and will be de-emphasized as a JEDEC-approved DDR2 feature. Instead, it will be offered as an optional feature targeted to engineering debug only. Many DDR2 suppliers, including Micron, are following this strategy. Hence, this feature is no longer shown on Micron's marketing DDR2 data sheets.

Mode Register

To support high-speed DDR2-800 speeds sooner, an optional CAS latency of 6 clocks has been added to the mode register, as shown in Figure 3.

Figure 3: Mode Register Change

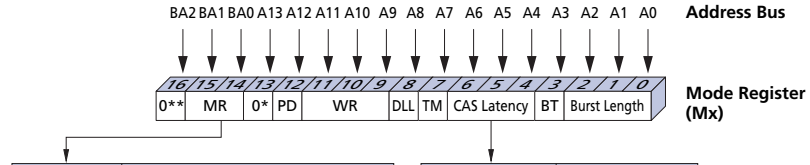

M15	M14	Mode Register Definition
0	0	Mode Register (MR)
0	1	Extended Mode Register (EMR)
1	0	Extended Mode Register (EMR2)
1	1	Extended Mode Register (EMR3)

*M13 (A13) is reserved for future use and must be programmed to "0".
**E16 (BA2) is reserved for future use and must be programmed to "0".

M6	M5	M4	CAS Latency
0	0	0	Reserved
0	0	1	Reserved
0	1	0	2 (optional)
0	1	1	3
1	0	0	4
1	0	1	5
1	1	0	6 (optional)
1	1	1	Reserved

Extended Mode Register

Supplementing the CAS latency of 6 clocks with an additive latency (AL) support of 5 clocks is also being considered. Doing so, however, would not be a straightforward change, so further investigation is required.

I/Os

The output slew rate specifications, which had been classified as "to be determined" (TBD), have been agreed on by JEDEC. The standard minimum slew rate allowed is 1.5 V/ns, and the maximum slew rate is set to 5.0 V/ns for DDR2-400, DDR2-533, and DDR2-667. Slew rate limits for DDR2-800 devices have not been established yet.

At the same time, capacitance minimum limits for the DDR2-400 and DDR2-533 devices were relaxed from 3.0pF to 2.5pF to provide backward compatibility with the newly defined DDR2-667 and DDR2-800 I/O capacitance limits of 2.5pF to 3.5pF.

Reduced I/O Drive

The JEDEC DDR2 specification has reserved EMR(1) bit1 for normal or weak output driver impedance control. The default state is "0" and provides "normal" or 100 percent of the full-drive target. JEDEC is evaluating what values to assign the "weak" drive curves as well as the "weak" state. The latter was assumed to be 60 percent of default. However, JEDEC is also considering 40 percent of default and 45 percent of default. The reduced I/O drive is intended to provide an impedance value that best supports point-to-point systems.

CKE Power-Down

A source of confusion with DDR2 was the definition of ^{t}CKE (MIN). It was addressed as follows:

1. The definition was changed to refer to the device "clock registrations" rather than the "pulse widths LOW."
2. The following note was added to the ^{t}CKE (MIN) definition:
   ^{t}CKE (MIN) of 3 clocks means CKE must be registered on 3 consecutive positive clock edges. CKE must remain at the valid input level the entire time it takes to achieve the 3 clocks of registration. Thus, after any CKE transition, CKE may not transition from its valid level during the time period of ^{t}IS + 2 × ^{t}CK + ^{t}IH
3. The ^{t}CKE (MIN) waveform in the timing diagram was improved to provide clarity and remove ambiguity, as shown in Figure 4.

Figure 4: Basic Power-Down Entry and Exit

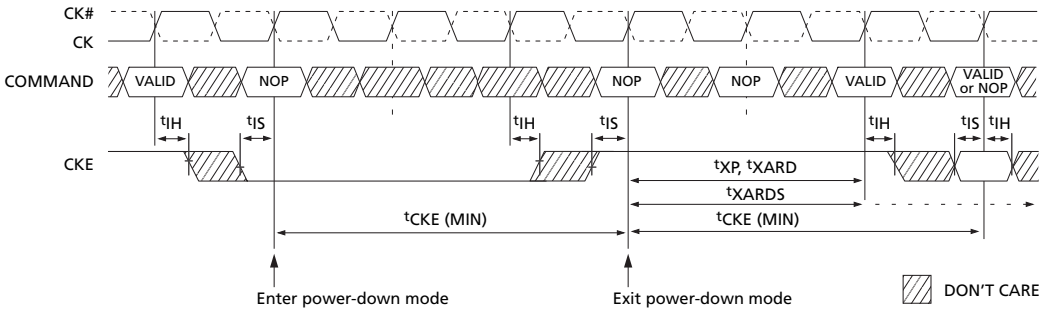

Command/Address Input Setup and Hold Timing

Previously, JEDEC specifications enabled symmetrical command/address setup and hold values by referencing the setup and hold timing specifications from the VREF level, represented by the blue reference points in Figure 5, to the positive clock (CK and CK# crossing). However, when these types of setup and hold specifications are derated for slow slew rates—those less than 1.0 V/ns—too many picoseconds (ps) are either lost or overmargined. This is because a large degree of uncertainty is introduced and is dependent on the VREF (dc) and VREF (ac) noise present.

The effects of the VREF noise become even more exaggerated at very slow slew rates, especially with nonlinear input signals, which may generate different timing results. This is dependent on which VREF was used as the reference point. The new methodology addresses these concerns by simplifying the measurement point outside the VREF region.

JEDEC approved a new methodology that references the command/address setup and hold timing from the appropriate logic level (the AC trip point for setups; the DC trip point for holds) to the positive clock's CK and CK# crossing, as indicated by the green reference points in Figure 5. Note that even though the specifications for the command/address setup and hold reference points and corresponding ^{t}IS/^{t}IH base parametric values have changed, they still represent the same timing relationship for a 1.0 V/ns linear input signal.

If this change in reference points is overlooked when using the new smaller values, the command/address setup and hold windows will lose almost 400ps of margin. This occurs if the new method's input timing values are used along with the old method's VREF reference points. For example, with a 1.0 V/ns input signal for DDR2-400, the previous setup and hold value was 600ps, while the new setup value is 350ps and the new hold value is 475ps. If the new setup and hold values are used correctly, they provide the same setup and hold times as the old method when the signals cross VREF. However, if the new setup and hold values are used with the old VREF method, the total setup and hold applied to the inputs is only 825ps, not 1,200ps. Thus, 375ps of timing requirement is lost (600ps + 600ps vs. 350ps + 475ps.)

JEDEC also recently added derating curves for command/address input signals with slew rates other than the specified base condition of 1.0 V/ns. Provided the input signal is monotonic, the new method does not require derating when the slew rate is less than 1.0 V/ns; but it does allow the designer to improve the "base" values. However, inputs with slew rates faster than 1.0 V/ns do require the "base" setup and hold values to be derated.

Figure 5: Command/Address Setup and Holds

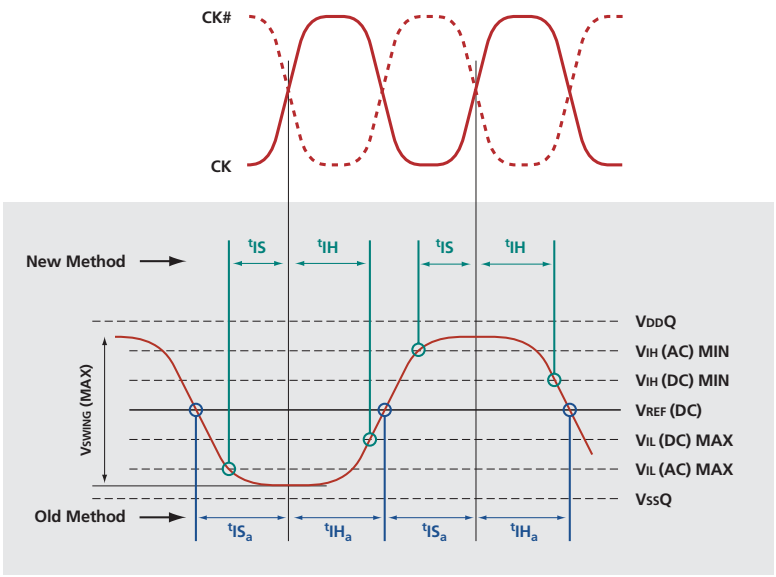

Data Setup and Hold Timing

Previously, JEDEC specifications also enabled symmetrical data input setup and hold values by referencing the setup and hold timing specifications from the VREF level to the positive clock (CK and CK# crossing), which are represented by the blue reference points in Figure 6. However, as with command/address input setup and hold timing, when these types of setup and hold specifications are derated for slow slew rates—again, less than 1.0 V/ns—too many picoseconds either get lost or overmargined. This is due to the large degree of uncertainty that is dependent on the VREF (dc) and VREF (ac) noise present. The VREF noise becomes even more exaggerated at very slow slew rates, especially with nonlinear input signals, which may generate different timing results. The new methodology addresses these concerns by simplifying the measurement point outside of the VREF region.

JEDEC approved a new methodology that references the data input setup and hold timing from the appropriate logic level (the AC trip point for setups; the DC trip point for holds) to the positive clock's CK and CK# crossing, as indicated by the green reference points in Figure 6. Note that even though the specifications for the data input setup and hold reference points and corresponding ^{t}DS/^{t}DH base parametric values have changed, they still represent the same timing relationship for a 1.0 V/ns linear input signal.

If this change in reference points is overlooked when using the new smaller values, the data input setup and hold windows will lose almost 400ps of margin. This occurs if the new method's input timing values are used with the old method's VREF reference points. For example, with a 1.0 V/ns input signal for DDR2-400, the previous setup and hold values were 400ps, while the new setup value is 150ps and the new hold is 275ps. If the new setup and hold values are used correctly, they provide the same setup and hold times as the old method when the signals cross VREF. However, if the new setup and hold values are used with the old VREF method, the total setup and hold applied to the inputs is only 425ps, not 800ps. Thus, 375ps of timing margin is lost (400ps + 400ps vs. 150ps + 275ps.)

JEDEC also recently added derating curves for data input signals with slew rates other than the specified base condition of 1.0 V/ns. Provided the input signal is monotonic, the new method does not require derating when the slew rate is less than 1.0 V/ns; but it does allow the designer to improve the "base" values. However, inputs with slew rates faster than 1.0 V/ns do require the "base" setup and hold values to be derated.

Figure 6: Data Setup and Holds

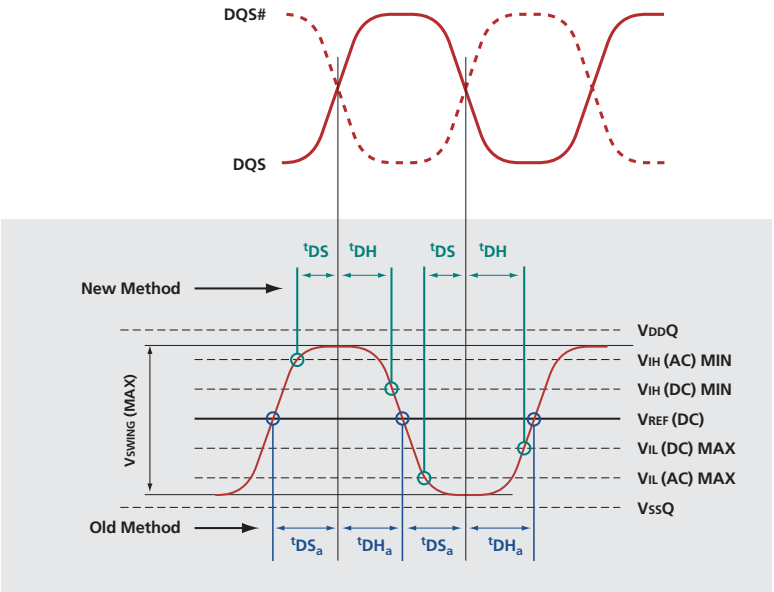

Self Refresh

An optional new feature known as high-temperature self refresh rate enable (HTSR) has been added to the extended mode register (EMR2) that allows the user to operate the DDR2 SDRAM at temperatures above the +85°C T_C limit. HTSR is enabled by setting bit7 of EMR2—which normally must be set to a "0"—to a "1," as shown in Figure 7. Because it has no upper limit standard, HTSR allows the DDR2 device to operate at the highest temperature limits defined by the part's manufacturer, even if they are higher than the +85°C T_C specification limit. In such cases, the DDR2 SDRAM will refresh itself internally at a faster rate to accommodate the higher T_C limit.

Micron does not currently support HTSR on its DDR2 SDRAM devices.

Figure 7: Extended Mode Register (2) With New High-Temperature, Self Refresh Option

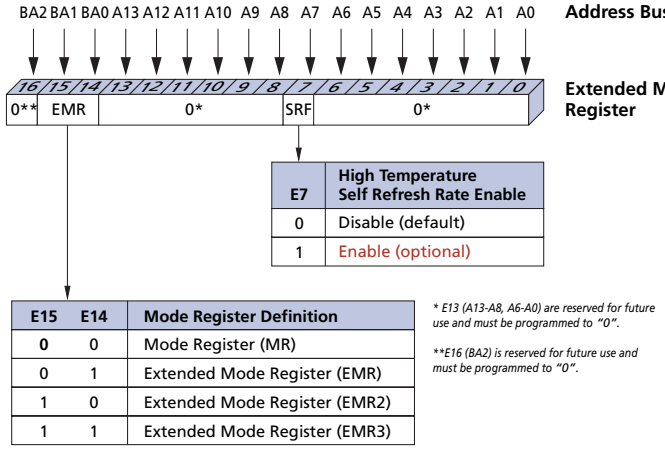

New Serial Presence-Detect Settings

Bytes 47–61 of the DDR2 module's serial presence-detect (SPD) feature were transferred from "reserved" status to "optional feature" status. Values were assigned to various thermal-related issues for special system support. The default value for all bytes is all "0s," which mirrors Micron's present support plans.

BYTES 47–61 define the following special options, if they are supported:

BYTE 47: Maximum case temperature delta
- Default is all "0s," i.e., +85°C T_C
- Any other value is in addition to +85°C

BYTE 48: Thermal resistance of device from case to ambient
- Default is all "0s," i.e., feature not supported

BYTE 49, bits1:0: If operation and self refresh are allowed to exceed +85°C
- Vendor defines value above +85°C
- Default is all "0s," i.e., feature not supported

BYTE 49, bit0: 0 = self refresh maximum T_C is +85°C; default
- Bit0: 1 = self refresh supports higher maximum T_C; optional
- Bit1: 0 = operating maximum T_C is +85°C; default
- Bit1: 1 = supports higher operating MAX T_C; optional

BYTES 50–61: Case temperature rises due to DRAM operations
- Includes BYTE 49, bits7:2
- From ambient to specified T_C; i.e. +85°C + BYTE 47
- DRAM operations such as IDD2N, IDD5, etc.
- Default is all "0s," i.e., feature not supported

Summary

After its standard evaluations and adjustments, JEDEC has finalized the specifications for DDR2 SDRAM components and modules. Consulting TN-47-02 "DDR2 Offers New Features and Functionality" as well as the updated information in this technical note, will provide designers with the most current DDR2 SDRAM standards and insights into the differences between DDR2 and DDR. The majority of changes or

clarifications in the JEDEC specifications are not significant in and of themselves; however, misunderstanding some of the changes could cause drastic errors in timing budgets.

Leading the Industry in DDR2 SDRAM Development

Micron Technology leads the field in the race to enable and support DDR2 SDRAM memory. In February 2004, Intel officially validated Micron's 533 MT/s speed, 1Gb DDR2 memory component, making Micron the first DRAM supplier to gain Intel® validation of DDR2 in all three core densities (256Mb, 512Mb, and 1Gb), in both speed bins (DDR2-400 and DDR2-533), and in all three configurations (x4, x8, and x16).

8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900
E-mail: prodmktg@micron.com, Internet: http://www.micron.com, Customer Comment Line: 800-932-4992