Technical Note

Designing for High-Density DDR2 Memory

Introduction

With densities ranging from 256Mb to 4Gb, DDR2 memory supports an extensive assortment of options for the system-level designer. Unlike the 4-bank-only technology of previous memory families, some configurations are now available with 8 banks. This abundance of options is important to the designer because of the different addressing and timing requirements for various devices.

To ensure maximum flexibility and a good migration path to higher-density devices, the system-level designer must understand the technical differences of the many DDR2 device options. This technical note focuses on configurtation, on the addressing schemes of each density, and on the subtle differences between the 4-bank and new 8-bank DDR2 devices.

Addressing Schemes and Why

When a memory device increases in density, either a bank, row, or column address must be added. Increasing a column address typically increases the activation current of the device. In some systems, the larger page size may provide optimal system-level performance by allowing more access (page hits) within a given bank.

Adding a single column address can substantially increase the activation current. See Figure 2 on page 3. When a DRAM row is activated (opened), it must transfer the contents of each individual cell within that row to separate sense amps. There are $2^{(n)}$ number of cells for each I/O within a row where *n* equals the number of column addresses.

For Example:

$$\textbf{\textit{Page size}} = \textit{Bits in the column address} \div 8$$

$$\textit{Bits in the column address} = 2^{(n)} \times \textit{Width of data path}$$

$$n = \textit{number of column addresses}$$

Using a x4 (DQ[3:0]) device with 11 addresses (A[9:0], A11)

$$\textbf{\textit{Page size}} = [2^{(11)} \times 4] \div 8 = 1024 = 1KB \quad \text{(EQ 1)}$$

If the column address is increased by one line, the page size for this same device doubles. A larger page size means each ACTIVE command must energize more of the array and additional sense amps.

Using a x4 (DQ[3:0]) device with 12 addresses (A[9:0], A11, A12)

$$\textbf{Page size} = \text{Bits in the column address} \div 8$$

$$\text{Bits in the column address} = 2^{(n)} \times \text{Width of data path}$$

$$n = \text{number of column addresses}$$

$$\textbf{Page size} = [2^{(12)} \times 4] \div 8 = 2048 = 2KB \quad \text{(EQ 2)}$$

Likewise, by increasing row addresses, the refresh overhead may change, or the device may have to execute multiple internal REFRESH cycles for each external AUTO REFRESH command. Also, the refresh power may rise due to the increased number of refreshes required, but the individual bank activation current would not increase for high-density designs (see Figure 1).

Which addressing method is best depends on the end-user application and complexity of the DRAM design—which may increase die size and cost (see Figure 2 on page 3). Other potential constraints include the number of device pins available, compatibility with other configurations, and PCB routing options. For DDR2, the Joint Electron Device Engineering Council (JEDEC) has defined the addressing requirements, as shown in Table 1 on page 4.

Figure 1: 8K Refresh With 13-Row and 14-Row Addresses

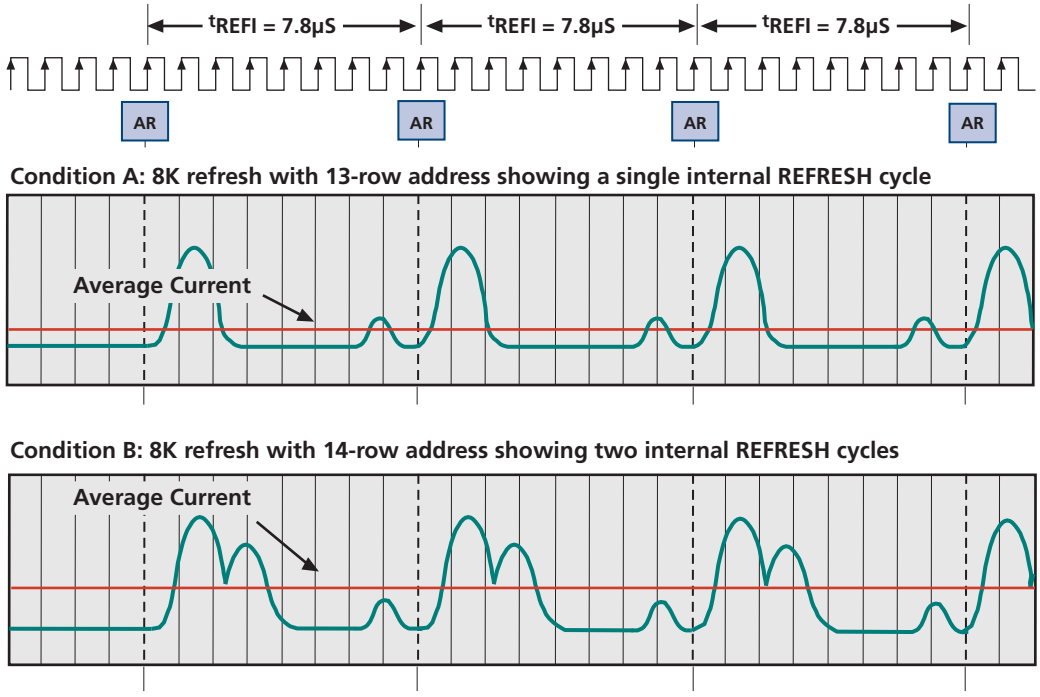

Figure 2: Typical DRAM Row/Column Array Within a Single Bank

Red shows circuitry for a *single column address*. Green shows circuitry for a single ACTIVE command (one active row for all DQ bits).

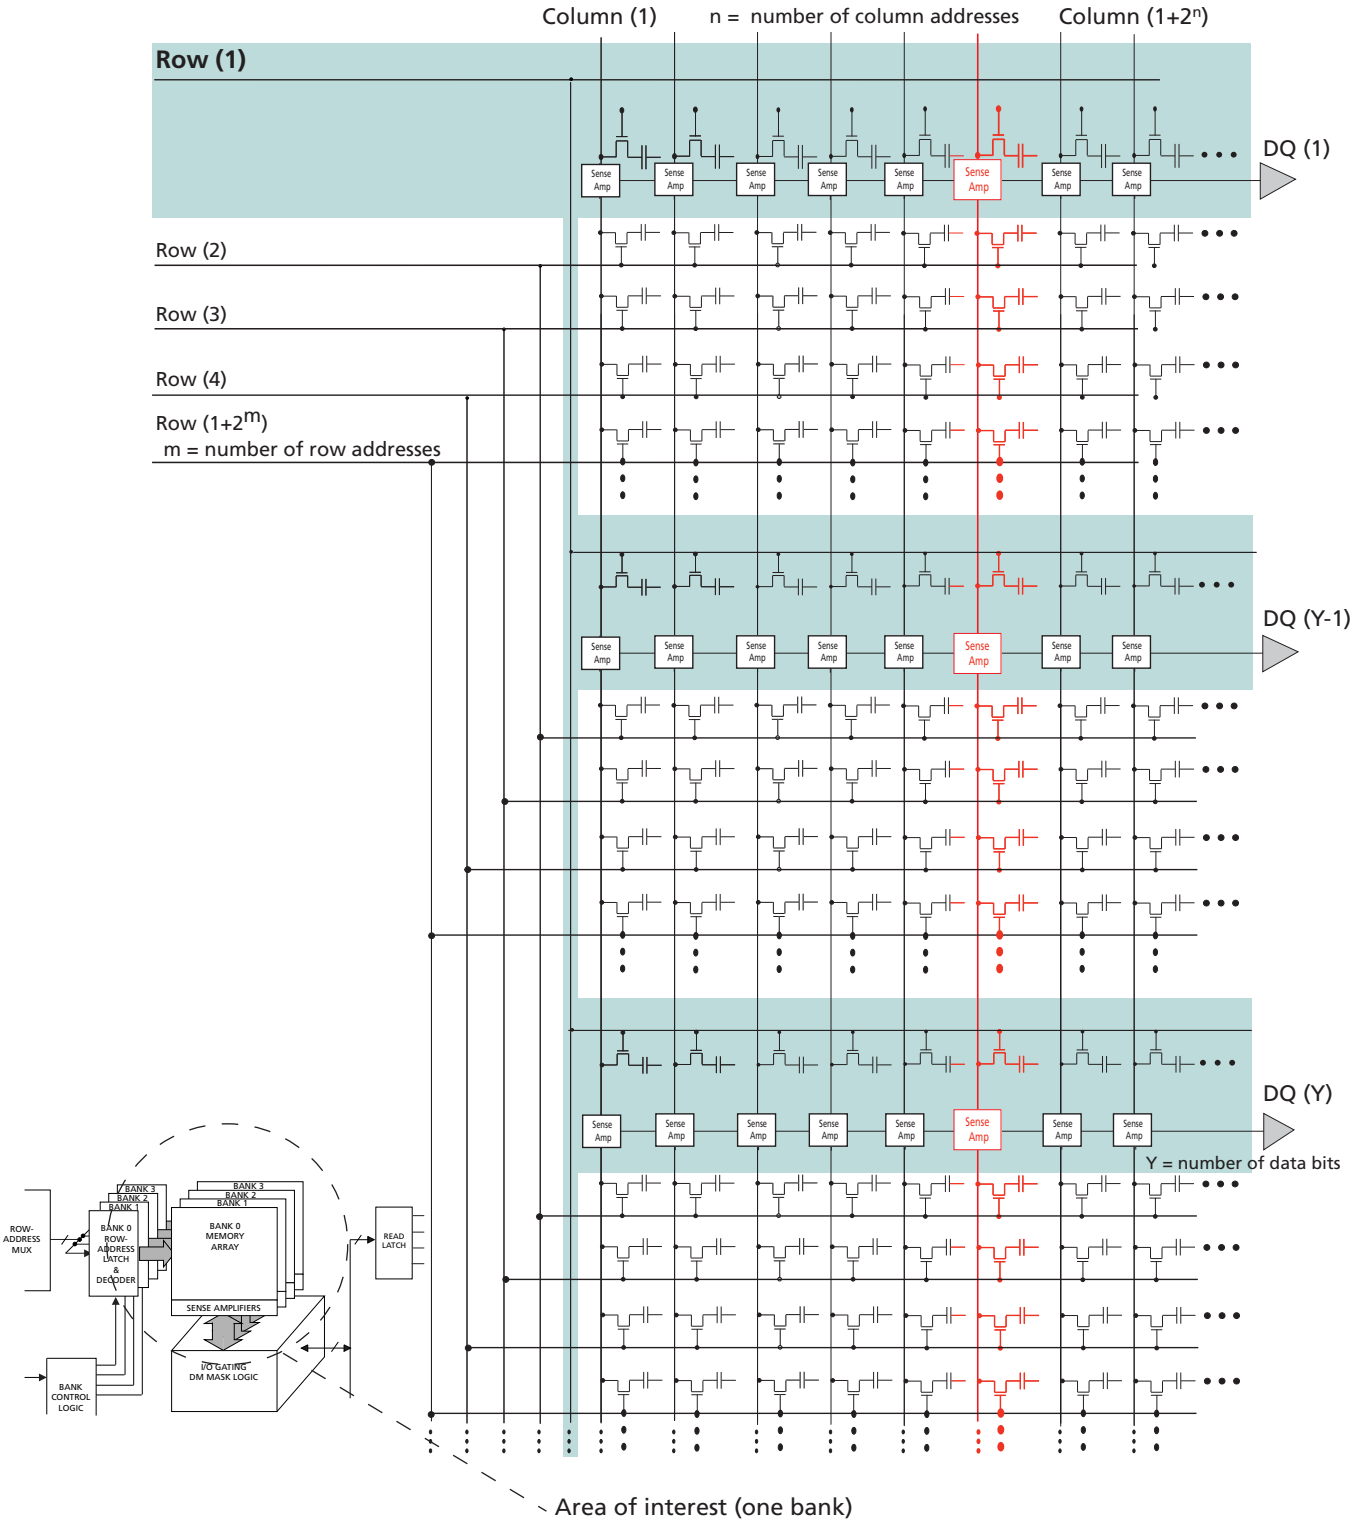

As shown in Table 1, the higher-density parts for DDR2 include only additional rows and/or bank addresses. By adding one more bank address, the DRAM can keep the same row and/or column addressing as the previous lower-density part, but the higher-density DRAM design goes from a 4-bank architecture to an 8-bank architecture.

Except for the x16 configuration, all DDR2 devices have a 1KB page size. For the x16 configuration, the page size is 2KB for all densities except the 256Mb device, which has a 1KB page size. Note that the page size is determined by the number of bits in the array divided by 8 (see Equation 1 on page 1).

Table 1: DDR2 Addressing 256Mb–4Gb

The 4Gb addressing is shown for completeness; currently Micron does not plan to support a 4Gb device.

Configuration	256Mb					
	64 Meg x 4		32 Meg x 8		16 Meg x 16	
Bank address	2	BA[1:0]	2	BA[1:0]	2	BA[1:0]
Row address	13	A[12:0]	13	A[12:0]	13	A[12:0]
Column address	11	A[9:0], A11	10	A[9:0]	9	A[8:0]

Configuration	512Mb					
	128 Meg x 4		64 Meg x 8		32 Meg x 16	
Bank address	2	BA[1:0]	2	BA[1:0]	2	BA[1:0]
Row address	14	A[13:0]	14	A[13:0]	13	A[12:0]
Column address	11	A[9:0], A11	10	A[9:0]	10	A[9:0]

Configuration	1Gb					
	256 Meg x 4		128 Meg x 8		64 Meg x 16	
Bank address	3	BA[2:0]	3	BA[2:0]	3	BA[2:0]
Row address	14	A[13:0]	14	A[13:0]	13	A[12:0]
Column address	11	A[9:0], A11	10	A[9:0]	10	A[9:0]

Configuration	2Gb					
	512 Meg x 4		256 Meg x 8		128 Meg x 16	
Bank address	3	BA[2:0]	3	BA[2:0]	3	BA[2:0]
Row address	15	A[14:0]	15	A[14:0]	14	A[13:0]
Column address	11	A[9:0], A11	10	A[9:0]	10	A[9:0]

Configuration	4Gb					
	1024 Meg x 4		512 Meg x 8		256 Meg x 16	
Bank address	3	BA[2:0]	3	BA[2:0]	3	BA[2:0]
Row address	16	A[15:0]	16	A[15:0]	15	A[14:0]
Column address	11	A[9:0], A11	10	A[9:0]	10	A[9:0]

Address Locations

Extra address signals have been reserved for both the device packages and high-density DIMMs. Depending on the product, the ball or pin locations may vary. Table 2 identifies the ball locations for the FBGA packages and pin locations for each of the JEDEC-standard DIMMs. For lower-density products, these pins are considered RFU and are not utilized on the package or module. As such, routing higher-order address signals to these locations is enabled, which provides compatibility with the higher-density products.

Note that DDR2 supports many different package sizes and ball grid arrays. It is advisable to completely understand the complexity of layout options before starting a board design; refer to TN-47-08, "DDR2 Package Sizes and Layout Requirements," at http://www.micron.com/products/modules/ddr2sdram/technote.html for specific information.

Although they are not discussed in this document, some DDR2 devices might be utilized in a stacked environment. Most stacked DRAM will utilize the same footprint as the monilithic device, but will have additional balls to support the additional signals. For DDR2, these additional balls typically include chip selects, clock enable, and ODT signals. In some stacked solutions, the electrical array may be mapped to different balls, in which case a custom layout may be required.

Table 2: DDR2 High-Density Address Locations
Address pins are considered reserved for future use (RFU) for lower-density solutions.

	FBGA Packages				Standard DIMMs		
	92-Ball	84-Ball	68-Ball	60-Ball	SODIMM (200-Pin)	UDIMM (240-Pin)	RDIMM (240-Pin)
A13	V8	R8	R8	L8	116	196	196
A14	V3	R3	R3	L3	86	174	174
A15	V7	R7	R7	L7	84	173	173
BA2	P1	L1	L1	G1	85	54	54

Accessing Multiple Banks

To optimize the command bus and increase system performance, some systems will cycle through as many banks as possible. This could mean that several banks are open simultaneously for an extended period of time while READ or WRITE operations are randomly occurring from any open bank, or the system may be constantly interleaving between all banks as fast as possible (opening bank *x* while reading bank *y*, closing bank *w* and opening bank *z*, etc.). All of these scenarios are very strenuous and demanding on the internal power rails of the DRAM device.

For example, to open a bank, the controller sends the DRAM an ACTIVE command with a specific row address and bank address. This initiates the address decoders, determines which cells are to be transferred to the sense amplifiers for temporary storage, and provides access to the I/O circuitry. To close a bank, the PRECHARGE command is issued to initiate a write-back of data from the sense amps to each of the individual cells. These sequences of opening and closing individual banks can consume a sizeable surge of current. Traditionally, DDR DRAM has operated with a maximum of 4 internal banks. However, with DDR2 some higher-density devices will support 8 individual banks. For this reason, JEDEC has limited the number of banks that may be activated within a set period.

DDR2 devices support a new timing parameter called Four Active banks Window (^{t}FAW). This is the minimum amount of time that must pass before more than four ACTIVE commands may occur. It is acceptable to have more than 4 banks open simultaneously, but the additional ACTIVE command(s) must be spaced out past the ^{t}FAW(MIN) window. As shown in Figure 3, ^{t}RCD for the fourth opened bank is complete at T8. To satisfy ^{t}FAW(MIN), the fifth ACTIVE command cannot occur until T11.

Figure 3: ^{t}FAW Example in an 8-Bank Device

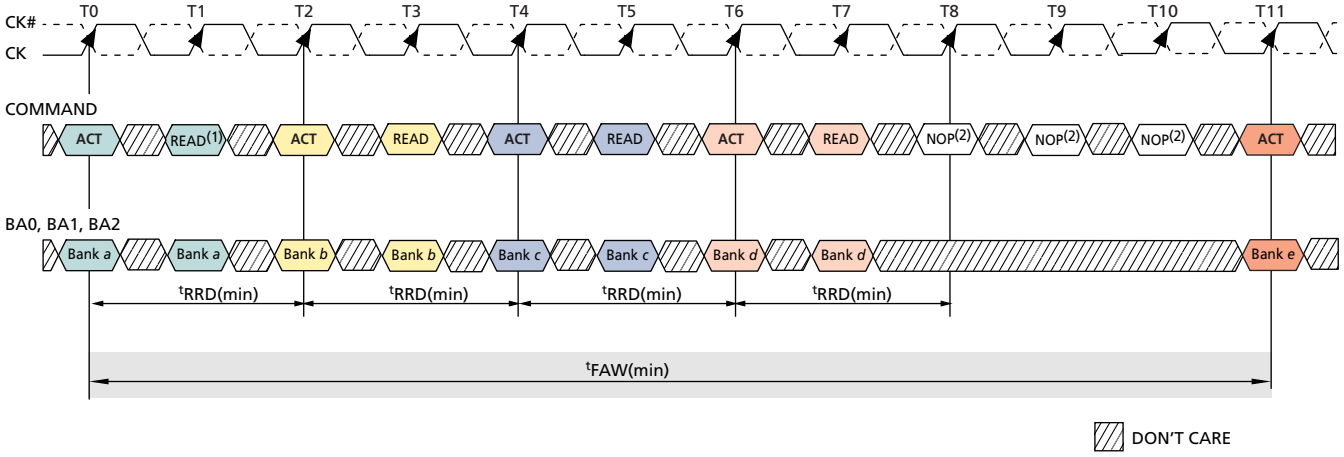

Notes:
1. Posted CL = 2.
2. NOP commands are shown for ease of illustration, but other commands may be valid at this time.
3. Example assumes ^{t}RRD = 2 clocks, ^{t}FAW = 11 clocks.

^{t}FAW applies for all devices, but for the 1Gb, 2Gb, and 4Gb (x16) configurations, the ^{t}FAW value is larger, as these are 8-bank devices with the larger 2KB page sizes. See Table 3 on page 7 for actual ^{t}FAW values.

Table 3: ^{t}FAW for All Densities and Speeds

(x4 and x8) Configurations	^{t}FAW				
	DDR2-400	DDR2-533	DDR2-667	DDR2-800	Units
256Mb (1KB page size)	37.5	37.5	37.5	35.0	ns
512Mb (1KB page size)	37.5	37.5	37.5	35.0	ns
1Gb (1KB page size)	37.5	37.5	37.5	35.0	ns
2Gb (1KB page size)	37.5	37.5	37.5	35.0	ns

(x16) Configuration	^{t}FAW				
	DDR2-400	DDR2-533	DDR2-667	DDR2-800	Units
256Mb (1KB page size)	37.5	37.5	37.5	35.0	ns
512Mb (2KB page size)	50.0	50.0	50.0	45.0	ns
1Gb (2KB page size)	50.0	50.0	50.0	45.0	ns
2Gb (2KB page size)	50.0	50.0	50.0	45.0	ns

Refresh Timing

When the DRAM is performing REFRESH cycles, it is idle and cannot perform a read or write sequence. Consequently, designers of high-throughput memory systems pay very close attention to the refresh rate of memory. As in previous DRAM technologies, the static refresh for DDR2 remains at 64ms. This means that every cell must be refreshed within the 64ms time limit, or data corruption may occur. Normally, most memory controllers utilize distributed REFRESH cycles. The distributed refresh rate (or the average amount of time between each AUTO REFRESH command) is determined by dividing the static refresh rate by the number of row addresses. So adding additional row addresses will typically increase the refresh rate (decrease the amount of time between distributed refreshes).

For example:

$$^{t}REFI = Static\ refresh\ time \div Number\ of\ rows \quad \text{(EQ 3)}$$

Using the standard equation above, the 512Mb DDR2 device should have a distributed refresh rate of 3.9µs (64ms/16k rows), but the 512Mb DDR2 device actually has a distributed refresh rate of 7.81µs. In fact, all DDR2 densities have an average periodic refresh interval (^{t}REFI) of 7.81µs (see Table 4).

Table 4: Key Refresh Parameters for DDR2 Devices

	Minimum Timing Required			
	256Mb	**512Mb**	**1Gb**	**2Gb**
Static refresh	64ms	64ms	64ms	64ms
Refresh interval ^{t}REFI (for ^{t}CASE = 85°C)	7.8µs	7.8µs	7.8µs	7.8µs
Refresh interval ^{t}REFI (for ^{t}CASE = 95°C)	3.9µs	3.9µs	3.9µs	3.9µs
Refresh time ^{t}RFC	75ns	105ns	127.5ns	195ns
Exit self-refresh to non-READ ^{t}XSNR	85ns	115ns	137.5ns	205ns

Notes:
1. ^{t}RFC is the time to complete one REFRESH cycle; only NOPs or DESELECT commands are allowed during this time.
2. ^{t}XSNR is the time immediately following an exit from self-refresh to any non-READ command. During this time, the only commands allowed are NOPs or DESELECT.
3. ^{t}REFI is the average time between distributed REFRESH commands.
4. Times in this table are for normal DRAM operating conditions and may increase for optional high temperature operation.
5. ^{t}CASE of 95°C is an optional feature and may not be supported on all designs.

To accomplish this, DDR2 DRAM has been designed to perform multiple internal REFRESH cycles for each external AUTO REFRESH command it receives. This slightly increases the time required to complete one REFRESH command cycle (^{t}RFC), but it keeps the refresh rate at a consistent value.

Figure 4 demonstrates the benefits of the DRAM's executing multiple internal refresh cycles for each single external refresh command. This example assumes ^{t}CASE = 85°C.

Figure 4: DDR2 Refresh Timing – Optimized for Maximum ^{t}REFI Time

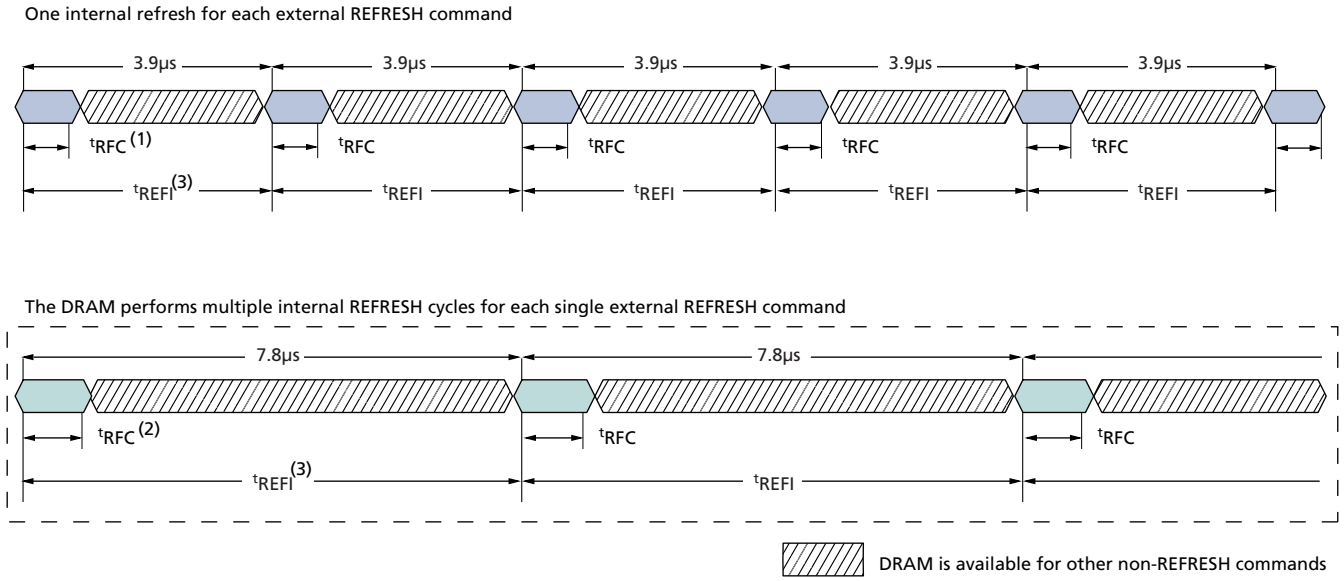

Notes:
1. DRAM performs one internal refresh for each external AR command (^{t}RFC time in ns).
2. DRAM performs two or more internal refreshes for each external AR command (^{t}RFC time is increased but is still in ns).
3. DRAM is available for any valid command after ^{t}RFC time (^{t}REFI time is in µm).

Conclusion

DDR2 promotes great design flexibility and a good migration path to higher-density memory, particularly if a few simple items are understood. Package sizes and ball arrays may vary, so be sure to review TN-47-08, "DDR2 Package Sizes and Layout Requirements," at http://www.micron.com/products/modules/ddr2sdram/technote.html before layout starts.

In addition, as the device density increases, either a row, column, and/or a bank address will be added. If a row is added, refresh timing is affected (for DDR2, ^{t}RFC time will be increased). When additional banks are added (some DDR2 devices support 8 banks), bank-to-bank timing will be affected, and a new timing parameter ^{t}FAW has been added to limit the number of ACTIVE commands within a specified time.

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