

Evolving to DDR3 technology

THE MOVE FROM DDR2 TO DDR3 REQUIRES NEW DESIGN TECHNIQUES TO IMPROVE SIGNAL INTEGRITY AND GAIN THE MAXIMUM BENEFIT FROM THE LATEST MEMORY-INTERFACE TECHNOLOGY.

Applications demanding higher system bandwidth and lower power, such as converged notebooks, desktop PCs, and servers, continue to drive the evolution of industry standards, including DDR3 (double-data-rate 3), as the JEDEC (Joint Electron Device Engineering Council) Solid State Technology Association defines it. The latest DDR3-memory standard, JEDEC JESD79-3A, supports these needs and the requirements of emerging dual-core- and multicore-processor systems. DDR3, the latest DDR-memory-interface technology, differs from the well-established DDR2 standard in several areas, including data rate, operating voltage, and logic (Table 1).

DDR2-memory-interface technology addresses and supports current system requirements. As memory requirements and

technology continue to evolve, however, the standards must adapt to enable approaches that deliver even more functions. DDR3 offers significant advantages over previous DDR generations. It supports data rates as high as 1600 Mbps per pin with an operating voltage of 1.5V, a 17% reduction from the previous generation of DDR2, which operates at 1.8V. DDR3's built-in power-conservation features, such as partial refresh, can be important in mobile-system applications in which battery power is not necessary just to refresh a portion of the DRAM not in active use. DDR3 also has a specification for an optional thermal sensor that could allow mobile-system engineers to save further power by providing minimum refresh cycles when the system is not in high-performance mode.

DDR3 uses eight internal banks compared with DDR2's four to further speed systems by allowing advance prefetch, which reduces access latency. This speed should become more apparent as the size of the DRAM increases in the future. The I/Os for DDR3 use the JEDEC standard SSTL (stub-series-terminated logic) 15, which employs 1.5V logic, whereas DDR2 uses

TABLE 1 DDR3 VERSUS DDR2		
	DDR2	DDR3/DDR3L
Rated speed (Mbps)	400 to 800	800 to 1600
Drain-to-drain voltage/supply voltage (V)	1.8±0.1	1.5±0.075 (DDR3)/ 1.35+0.1/-0.7 (DDR3L)
Internal banks	Four	Eight
Termination	Limited	All DQ signals
Topology	Conventional T	Fly-by
Driver	OCD calibration	Self-calibration with ZQ
Thermal	No	Optional




Figure 2 The DDR3 fly-by topology finds use in command/address and clock signals to improve signal integrity.

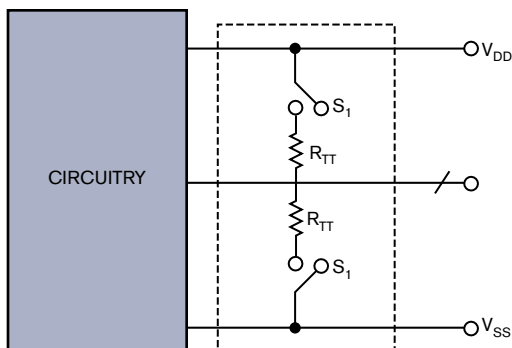


Figure 1 You can switch the on-die termination and select the effective value through a pullup/pulldown-resistor network.




Figure 3 With the optimal termination value, driver impedance, and output inversion enabled, the postregister DDR3 data-eye plot exhibits a wider opening with minimal ring-back.

JEDEC standard SSTL18, which uses 1.8V logic. The DDR3 architecture fully uses ODT (on-die termination), ZQ (zero-ohm) calibration, and a fly-by topology for improved signal integrity. With such a high demand for lower-power memory and cost-saving technologies, JEDEC is now defining a DDR3L (DDR3-low-voltage) node.

IMPROVING SI

Because DDR3 runs at higher memory speeds, the signal integrity of signals traveling through the memory module becomes more important. DDR3 uses a fly-by topology instead of the T branches that DDR2-module designs use. Thus, the address and control lines are in a single path chaining from one DRAM to another, whereas DDR2's T topology branches on the modules. Fly-by topology eliminates the mechanical line-balancing requirement and uses an automatic signal-time delay that the controller generates during memory-system training. Each DDR3 DRAM chip has an automatic leveling circuit for calibration and to memorize the calibration data.

DDR3 implements several impedance-calibration sequences to improve signal integrity. It uses long-ZQ calibration after power-up and periodically uses short-ZQ calibration during normal operation to compensate for voltage and temperature drift. These calibration sequences vastly improve the connectivity between the output driver of the DRAM and the PCB (printed-circuit-board) trace. A ZQ pin on the DRAM connects to an external precision resistor that adjusts the output-driver impedance as well as the ODT to match the trace impedance, thus reducing impedance discontinuity and minimizing reflection on the signals. The use of external precision resistors reduces the effect of variations due to process, voltage, and temperature and maintains a tight tolerance for better-controlled impedance values. DDR2, on the other hand, employs on-chip resistors, which can exhibit larger variations. The system and DRAMs also use dynamic ODT for improved signaling, especially for higher speeds.

You can switch the ODT on and off at the DRAM and select




Figure 4 Normal operation inverts half of the address bus on the postregister to prevent simultaneous-switching-output noise.




Figure 5 During an MRS command, the timing through the register automatically switches to driving the data for three clock cycles to account for the simultaneous-switching-output effect.

the effective value through a network of pullup and pulldown resistors (Figure 1). The dynamic nature of the circuit provides optimal command-, address-, control-, data-, and strobe-bus termination for better signal control and improved margins than those that DDR2 achieves using static termination on the motherboard. To further improve signal integrity, DDR3 also uses its fly-by topology for command/address and clock signals (Figure 2). It routes the signals to the DRAMs linearly and to the edge of the card at its bus termination. For registered DIMMs (dual-inline-memory modules), an IC component buf-

fers the command, address, control, and clocks. This approach helps reduce the number of stubs and stub lengths that normally would be in DDR2's T topology; however, this approach also introduces flight-time skew between the clocks and data strobes at the DRAMs. You can compensate for this flight-time skew from the controller side on the motherboard by performing a leveling technique for deskew, which puts the DRAMs through a training sequence for tuning the DRAM clock.

The internal core speed of the DRAM basically remains unchanged in the transition from DDR2 to DDR3. DDR2 currently has a maximum bandwidth of 800 Mbps per pin but can extend to 1066 Mbps. To meet bandwidths as high as 1600 Mbps, DDR3 uses an 8-bit prefetch rather than the 4-bit prefetch that DDR2 uses. As a result, for every read or write operation, the technology accesses 8 bits in the DRAM core.

Registered DIMMs from previous generations, DDR and DDR2, exhibit excellent performance for systems that require higher bandwidths and better throughput efficiency. DDR2 employs at least one register and a PLL instead of two separate components; DDR3 employs a single monolithic-IC chip, which integrates the register and PLL. This integrated part features programmable-drive strength, input-bus termination, and output inversion, effectively reducing SSO (simultaneous-switching-output) noise. For power savings, the device includes output inversion, the ability to float the outputs, and the ability to power down the chips using input logic states. Another power-saving technique comes from programming the output drivers' impedance. Enhanced termination techniques and impedance matching vastly improve the signal integrity of the data eye using an integrated device, which in




Figure 7 You can use a flex circuit and a PCB interposer to stack BGA (a) or TSOP (b) DRAM.



Figure 6 The very-low-profile DIMMs are essentially the same as the standard DIMMs but at 60% of the height.

turn increases the margin in the timing budget.

Using the optimal termination value, driver impedance, and output inversion, the postregister data eye exhibits a wider opening with minimal ring-back (**Figure 3**). The preregister also exhibits similar data-eye quality if it receives sufficient setup-and-hold time from the controller to prevent any bit errors. The register-PLL component uses on-die termination to maximize throughput efficiency and achieve higher speeds without redesigning or doing any extensive modifications to the system.

PROGRAMMING CAPABILITY

The memory controller sends commands to the memory that force it into a mode of operation. The states of CS (chip-select), RAS (row-address select), CAS (column-address select), and WE (write enable) define the commands. The MRS (mode-register-set) command is one of the first in system initialization. It allows programming of the configuration registers in the DRAM for various functions, features, and modes, making the memory module flexible for applications. The configuration registers are at their default values at power-up and programmed accordingly for proper operation.

The MRS command is present in both DDR2- and DDR3-DIMM applications. In DDR3-registered DIMMs, however, the register-IC chip, or register, which you can also program, can decode the command and switch to a different timing mode by driving the received data for three clock cycles instead of the conventional one. In normal operation, the system intentionally inverts half of the address bus on the postregister to prevent SSO noise (**Figure 4**).

An MRS command disables address inversion to allow correct access to the DRAMs, consequently increasing the propagation delay due to SSO noise. To account for SSO effects, the timing through the register automatically switches to driving the data for three clock cycles (**Figure 5**). This mode was not present in DDR2-registered-DIMM applications because the registers for DDR2 are strictly buffers without any programmability and decoding logic. You can program the register and the DRAMs for various functions, features, and modes. The JEDEC DDR3 register-specification document provides register mapping.

INCREASING MEMORY CAPACITY

With data centers needing more capacity, space is at a premium. Few options are available for maximizing space. One is the usage of more blade servers, which would require internal pe-




Figure 8 One version of die-stack-packaging technology has an interposer (a); another version does not (b).

ipherals to change to perhaps a smaller form factor, such as hard disks, memory, or I/O cards. JEDEC in 2005 established a VLP (very-low-profile) standard for smaller-form-factor DIMMs. As blade servers become more common, so will the VLP DIMMs, and the technology promises to become more of a presence as users begin to adopt DDR3-memory technology. The VLP DIMMs are essentially the same as the standard DIMMs but at 60% of the height (**Figure 6**). The height reduction improves airflow in the system, thus improving cooling. Other applications needing VLP DIMMs are ATCA (Advanced Telecommunications Com-

puting Architecture), Micro ATCA for the telecommunications market, and embedded single-board computers. Other small form factors for DIMMs are SO-DIMM (small-outline DIMM) and mini-DIMM.

Standard and VLP DIMMs have various densities and ranks. DIMMs are typically 64- or 72-bit-wide words. A *rank* is an identical arrangement of memory banks on a module. The following equation shows how to determine the DIMM's total memory density: DRAM density/8×bus width/DRAM width×no. of ranks=total density. The configuration and density of the component de-

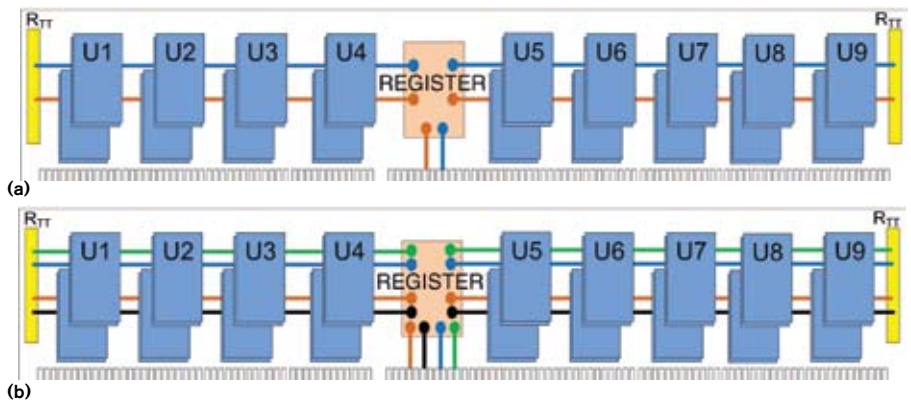


Figure 9 The dual-rank module requires only two chip-select signals (a), whereas quad-rank-registered DIMMs with stacked dual-die DRAM require four (b).

termine the number of ranks on a module. For example, two ranks of 512-Mbit DRAMs with 4-bit-wide data on a module with a 64-bit-wide bus would be $512 \text{ Mbits}/8 \times 64/4 \times 2 = 2$ Gbytes.

The most common ranks are single and double, but the total number of possible ranks is four. The quad-rank module is relatively new to DDR DIMMS, and until recently, memory controllers did not support it. Due to the ever-increasing demand for higher-density DIMMs, however, today's CPUs support quad-rank modules. In a relatively mature DDR2-memory market, JEDEC has defined and standardized only two quad-rank DIMMs, and two more are in development. In contrast, in an emerging DDR3-memory market, six quad-rank modules are in development.

A JEDEC-standard memory module in a planarity configuration has a maximum of 36 DRAM chips. DRAM and module vendors have developed several stacking technologies to further increase memory capacity on any memory module, however. One stacking technology is the use of a PCB or flex-circuit interposer to stack BGA (ball-grid-array) or TSOP (thin small-outline package) DRAM (**Figure 7**). For improved manufacturing reliability, die-stack packages are other alternatives to stacking DRAMs. This technology allows you to place two or four die in one chip package. A couple of die-stack-packaging technologies use a window chip-scaled package with an interposer (**Figure 8**). In this approach, both die face in opposite

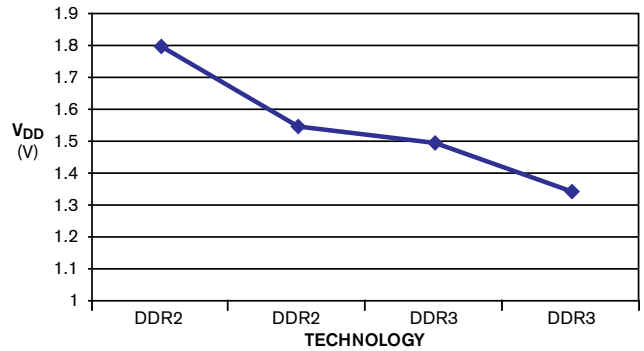


Figure 10 DDR3 uses 25% less power per DIMM than does DDR2.

directions, or you wire-bond both die in the same direction. In theory, a quad-die chip would maximize the number of DRAM die from 36 to 144 on a DIMM. However, due to electrical limitations, the maximum number achievable is now 72.

Two registered DIMMs implement dual- and quad-rank modules with dual-die stacked DRAMs (**Figure 9**). The dual-rank module requires only two chip-select signals, whereas the quad-rank module requires four, one for each rank. The implementation of quad-die-stacked DRAMs on modules will

be commonplace as higher-density modules and smaller-form-factor modules, such as VLP-registered DIMMs, SO-DIMMs, and mini-DIMMs, become more mainstream.

LOW-POWER DDR TECHNOLOGY

As data centers and server farms increase their capacity to meet higher bandwidth demand and increasing DDR-memory-interface data rates—and as companies grapple with rising energy costs—overall power consumption becomes more of a concern. The importance of energy saving is beginning to

take precedence. The DDR market is currently addressing the need for low-power, cost-saving approaches. The DDR2 technology operates at 1.8V, which is approximately a 17% increase from DDR3, but companies are pushing for reducing that node to 1.55V. A similar trend is beginning to develop in DDR3 in which the original standard operating voltage is

1.5V but is decreasing to the 1.35V operating-voltage node, which is another 10% reduction.

Figure 10 shows a linear downward trend of the operating voltage for the DDR technology. Going from 1.8 to 1.35V yields a 25% reduction in power per DIMM alone, with some servers supporting as many as 18 DIMMs. This reduction is significant not only in the core of the memory and logic chips of the DIMMs, but also in the I/Os of both, which consume most of the power. As the operating voltage scales down, so will the memory-interface technology. Manufacturers are studying other areas of memory-interface technology for power reduction. These areas include implementing special power-down features and modes when the system is idling.

From mobile-system applications to notebooks to enterprise servers, a constant demand exists for higher bandwidth, lower power, and improved throughput efficiency. The industry is examining all aspects of the DDR-memory-interface technology to meet the needs and demands. The onus is on the DDR-chip vendors to adapt and provide IC chips for DDR DRAMs, registers, and PLLs with improved signal performance. JEDEC has defined innovative methods to improve signal integrity from one generation of DDR to the next to facilitate memory-driven-application requirements. It is clear that DDR3 memory is addressing the need for higher-density modules with lower power and more efficient throughput. **EDN**

✚ For a list of references, go to www.edn.com/090528ms4317.

✚ Go to www.edn.com/ms4317 and click on Feedback Loop to post a comment on this article.

AUTHOR'S BIOGRAPHY

John Nieto is a senior applications engineer at Inphi Corp, where he is responsible for DDR2 and DDR3 registers and PLL support for DIMM-application simulations. He has a bachelor's degree in electrical engineering from the University of California—San Diego.