

## The Evolution from DDR2 to DDR3 and its Impact on Signal Integrity

By John Nieto, Sr. Applications Engineer, Inphi Corporation

[www.Inphi-Corp.com](http://www.Inphi-Corp.com)

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 as defined by JEDEC. The latest DDR3 memory standard (JEDEC JESD79-3A) is designed specifically to support these needs and the requirements of emerging dual and multi-core processor systems. DDR3, the latest DDR memory interface technology, differs from the well-established DDR2 standard in several areas including data rate, operating voltage, logic and other key areas. This sidebar will examine these differences, review design techniques that improve signal integrity using DDR3, and examine products currently available to help designers exploit the benefits of DDR3 technology.

DDR2 memory interface technology addresses and supports current system needs and requirements; but as memory requirements and technology continue to evolve, the standards must adapt to enable solutions that deliver even more. The newest DDR memory interface technology, DDR3, offers significant advantages over previous DDR generations. DDR3 supports data rates up to 1600 Mbps per pin with an operating voltage of 1.5 volts, a 17% reduction from the previous generation of DDR2, which operates at 1.8 volts. DDR3's built-in power conservation features, like partial refresh, can be particularly important in mobile applications where battery power will no longer be needed 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 engineers to save further power by providing minimum refresh cycles when the system is not in high performance mode.

Official JEDEC Specifications		
	DDR2	DDR3/DDR3L
<b>Rated Speed</b>	400-800 Mbps	800-1600 Mbps
<b>Vdd/Vddq</b>	1.8V +/- 0.1V	1.5V +/- 0.075V (DDR3) 1.35V + 0.1V /- 0.7V (DDR3L)
<b>Internal</b>	4	8
<b>Termination</b>	Limited	All DQ signals
<b>Topology</b>	Conventional T	Fly-by
<b>Driver</b>	OCD Calibration	Self Calibration with ZQ
<b>Thermal</b>	No	Yes (Optional)

DDR3 uses more internal banks - 8 instead of the 4 used by DDR2 - to further speed up systems by allowing advance prefetch which reduces access latency. This should become more apparent as the size of the DRAM increases in the future.

The Input/Outputs (I/Os) for DDR3 are designed to use the JEDEC standard SSTL15, which is based on 1.5V logic, while DDR2 uses JEDEC standard SSTL18, based on 1.8V logic. The DDR3 architecture fully utilizes on-die termination (ODT), ZQ calibration, and a fly-topology for improved signal integrity. Finally, with such a high demand in lower power memory and cost-saving solutions, a low voltage DDR3 (DDR3L) node is currently being defined in JEDEC.

### Optimizing Signal Integrity

Since DDR3 is designed to run at higher memory speeds the signal integrity of signals traveling through the memory module is now more important. DDR3 uses a "fly-by" topology instead of the "T branches" seen on DDR2 module designs. This means the address and control lines are a single path chaining from one DRAM to another, where DDR2 uses a T topology that branches on the modules. Fly-by topology takes away the mechanical line balancing requirement and uses an automatic signal time delay

generated by the controller fixed at the memory system training. Each DDR3 DRAM chip has an automatic leveling circuit for calibration and to memorize the calibration data.

Several impedance calibration sequences are implemented for DDR3 that optimize signal integrity. The Long ZQ calibration is used after power-up and the Short ZQ calibration is used periodically during normal operation to compensate for voltage and temperature drift. These calibration sequences vastly improve the connectivity between the output driver of the SDRAM and the PCB trace. A ZQ pin on the SDRAM is connected 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 variation 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 utilize dynamic ODT for improved signaling, especially for higher speeds.

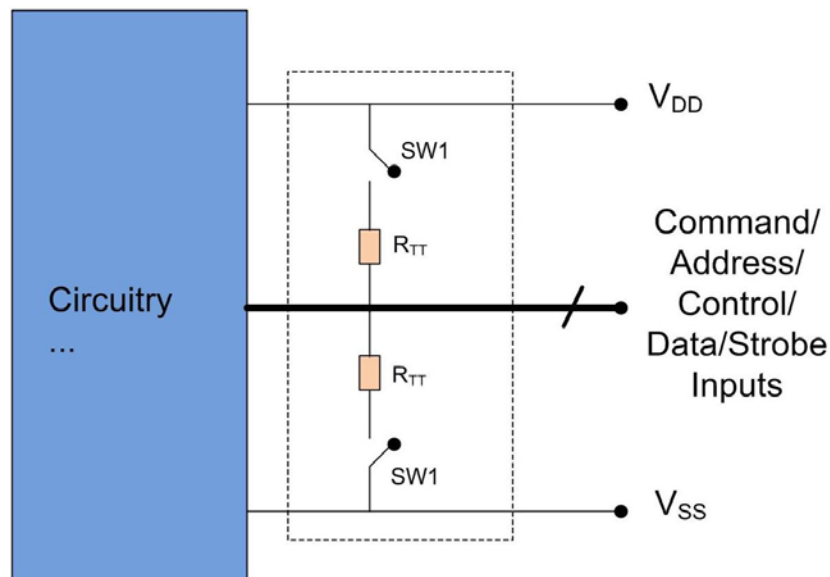


Figure 1. On Die Termination

At the DRAM, the ODT can be switched off and on and the effective value selected through a pullup and pulldown resistor network as shown in Figure 1. The dynamic nature of the circuit provides optimal command/address/control/data/strobe bus termination for better signal control and improved margins as opposed to using a static termination on the motherboard implemented in DDR2.

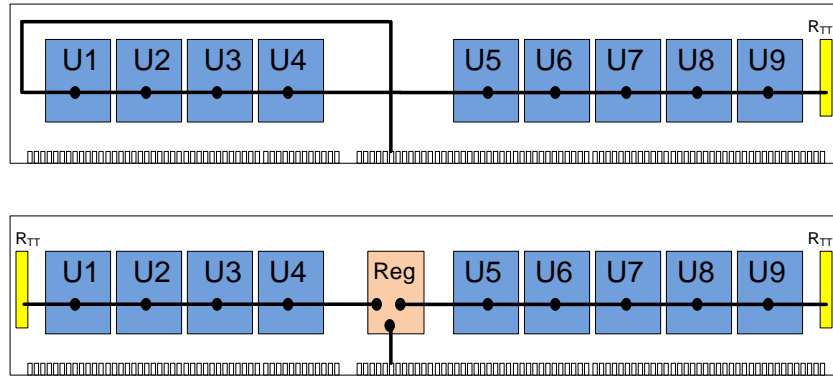


Figure 2. UDIMM and Registered DIMM fly-by topology

To further improve signal integrity, DDR3's fly-by-topology, as shown in Figure 2, is also utilized for command/address and clock signals. The signals are routed to the DRAMs in a linear fashion and to the edge of the card where the bus termination is located. For the registered DIMM, the command/address/control and clocks are buffered by an IC component. This helps reduce the number of stubs and stub lengths that normally would be present in DDR2's T-topology; however, this introduces flight time skew between the clocks and data/strobes at the DRAMs. The flight time skew, however, can be compensated from the controller side on the motherboard by performing a leveling technique for de-skew, 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 800Mbps per pin but can extend to 1066Mbps. In order to meet DDR3 bandwidths of up to 1600Mbps, an eight-word prefetch is used as opposed to the four-word prefetch used with DDR2. As a result, for every read or write operation, eight words within the core of the DRAM are accessed.

Registered DIMMS (RDIMMs) from previous generations, DDR and DDR2, exhibit excellent performance for systems that require higher bandwidths and better throughput efficiency. Inphi Corporation provides registers and PLLs for DDR2 registered DIMM and for DDR3 registered DIMM. Instead of two separate components, at least one register and a PLL are used in DDR2; 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

simultaneous switching noise. For power savings, the device includes output inversion, the ability to float the outputs, and completely power down the chips via input logic states. Another power savings technique can come from programming the output drivers' impedance. With enhanced termination techniques and impedance matching, the signal integrity of the data eye is vastly improved using an integrated device, which in turn increases the margin in the timing budget.

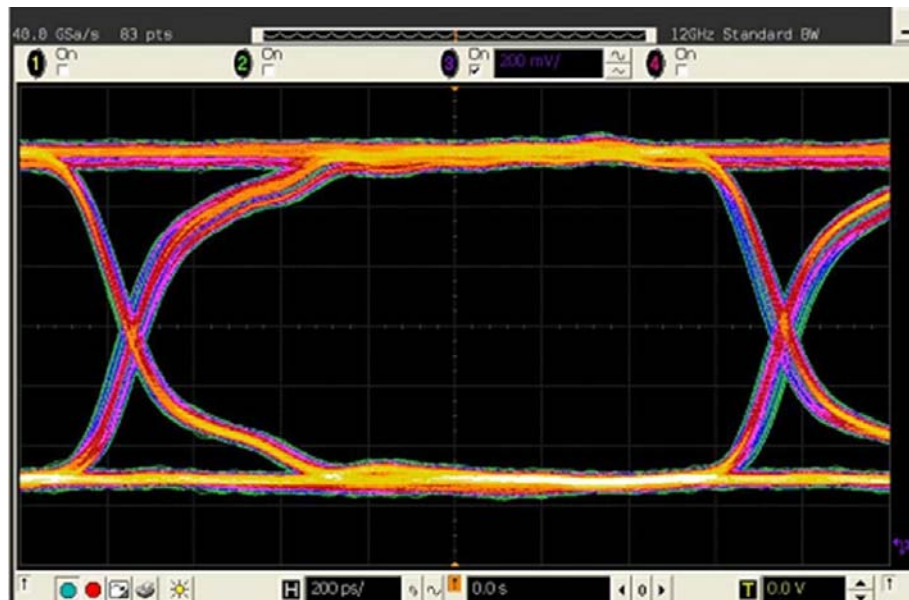


Figure 3. DDR3-1333 Data Eye plot

Using the optimal termination value, driver impedance, and output inversion enabled, the post-register data eye will exhibit a much wider opening with very minimal ringback as shown in Figure 3. Similar data eye quality is also expected at the pre-register, given enough setup and hold time from the controller in the system, in order to prevent any bit errors. This is a reason why the register PLL component utilizes on-die termination. Throughput efficiency is maximized and higher speeds easily achieved without having to redesign or do any extensive modifications to the system.

### **Programming Capability**

The memory controller sends out commands to the memory that will force it into a specific mode of operation. The commands are defined by the states of CS# (chip-select), RAS# (row address select), CAS# (column address select), and WE# (write enable). One of the first commands in a system boot up is the Mode Register Set (MRS)

command for initialization. This allows programming of the configuration registers in the DRAM for various functions, features, and modes, making the memory module flexible for different applications. The configuration registers are at its default values at power-up and programmed accordingly for proper operation.

The MRS command is present in both DDR2 and DDR3 DIMM applications but in DDR3 RDIMMs, the register IC chip (Register), which can also be programmed, is able to 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 (1T timing), half of the address bus on the post-register is intentionally inverted to prevent simultaneous switching output (SSO) noise. The timing diagram for normal operation is shown in Figure 4.

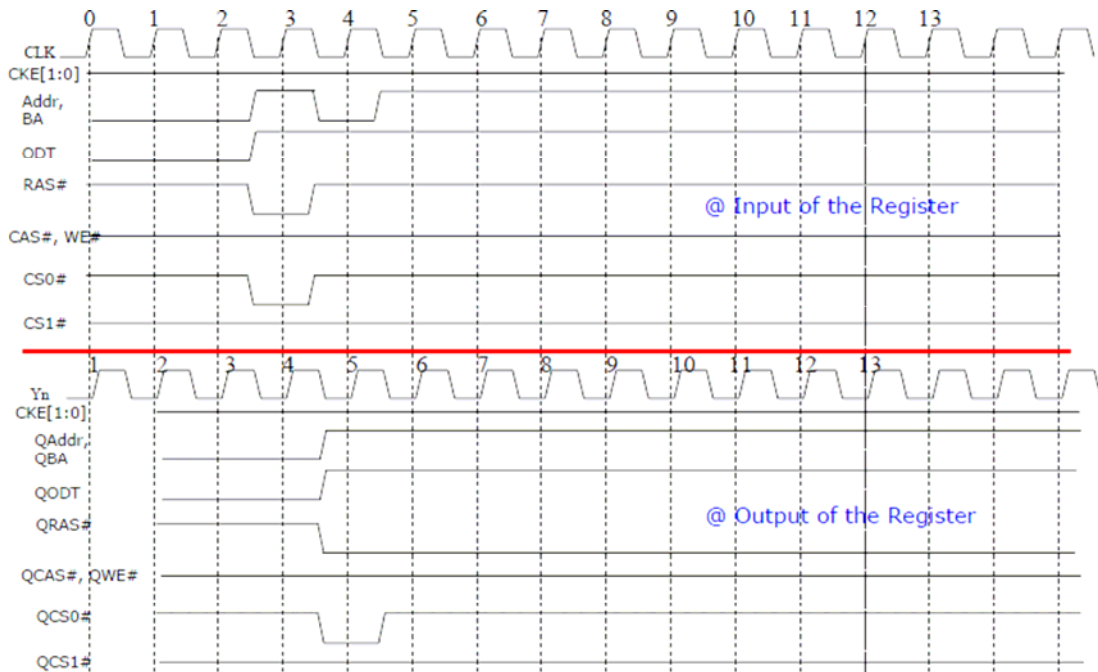


Figure 4. Normal Operation (1T timing) – Address Inversion Enabled

In an MRS command, the address inversion is disabled in order to allow correct access to the DRAMs, consequently increasing the propagation delay due to SSO noise. To account for SSO effect, the timing through the Register automatically switches to driving the data for three clock cycles as shown in Figure 5. This mode was not present in DDR2 RDIMM applications because the Registers for DDR2 are strictly buffers without any programmability and decoding logic.

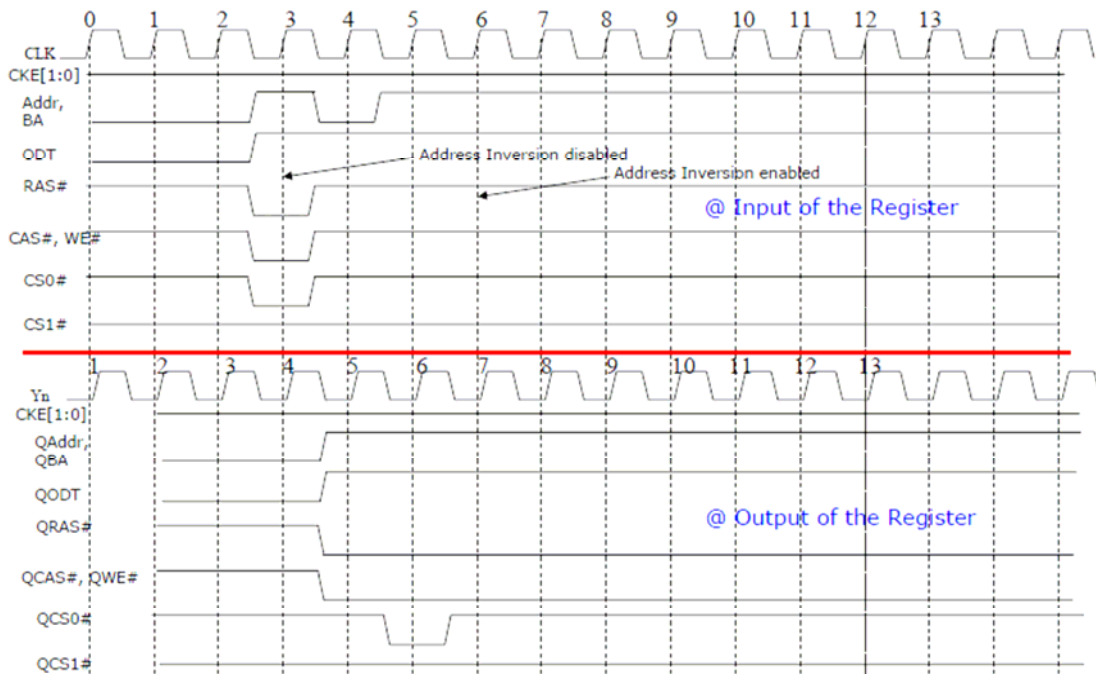


Figure 5. MRS Command (3T timing)

The Register is capable of being programmed for various functions, features, and modes along with the DRAMs. A register mapping is provided in the JEDEC DDR3 Register specification document.

### Increasing Memory Capacity

With data centers needing more capacity, space is becoming a premium. Trying to maximize space doesn't yield too many options. One option for maximizing space is the usage of more blade servers. This in turn would require internal peripherals to change to perhaps a smaller form factor such as hard disks, memory, I/O cards, etc. There is a new standard for smaller form factor DIMMs, namely, very low profile (VLP) DIMMs, that began in 2005. As blade servers become widely used, so will the VLP DIMMs, and promises to be more of a presence going into the DDR3 memory technology. The VLP DIMMs are essentially the same as the standard DIMMs but at 60% of the height. The height reduction improves airflow in the system, thus improving cooling. Other applications needing VLP DIMMs are Advanced Telecommunications Computing Architecture (ATCA), Micro ATCA for the telecommunications market and embedded Single Board Computer (SBC). Other small form factors for DIMMs are SO-DIMM and min-DIMM.



Figure 6. VLP DIMM vs. Standard DIMM

Within the standard and VLP DIMMs are different types of densities and ranks in memory modules. Memory modules are typically organized as either 64-or 72 bit wide words. A single “rank” is an identical arrangement of memory banks on a module. The total memory density of the module is determined by the depth and width. The number of ranks on a module is determined by the configuration and density of the component.

Eq.1)  $\text{SDRAM density} / 8 * \text{Bus width} / \text{SDRAM width} * \# \text{ of ranks} = \text{Total density of module}$

Examples: 2 ranks of 512Mb SDRAMs with 4 bit data width on a module with a 64 bit wide bus

$$512\text{Mb}/8 * 64/4 * 2 = 2\text{GB}$$

The most commonly used ranks are single and double or dual but the total number of possible ranks is four or quad. The quad-rank is relatively new to DDR memory modules and was not widely supported by memory controllers but due to the ever increasing demand for higher density memory modules, today’s CPUs are capable of supporting quad-rank modules. In a relatively mature DDR2 memory market, there are only two quad-rank DIMMs that have been defined and standardized through JEDEC and another two in development . In contrast, in an emerging DDR3 memory market, there are already six quad-rank modules in development.

The number of DRAM chips on a JEDEC standard memory module in a planarity configuration is limited to a maximum of 36. However, DRAM and module vendors have been able to develop several stacking technologies to further increase memory capacity

on any particular memory module. One stacking technology is the use of a PCB or flex circuit interposer to stack BGA or TSOP DRAM packages as shown in Figure 7.

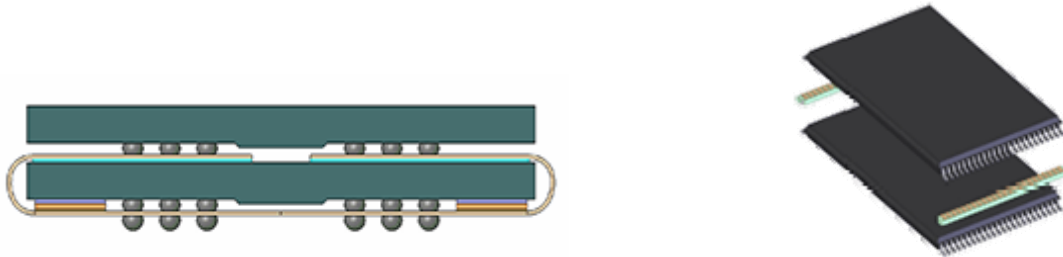


Figure 7. Stacked packaging using a flex circuit and PCB interposer

For improved manufacturing reliability, die stack packages is another alternative to stacking DRAMs. This technology enables up to two (dual) or four (quad) die in one chip package. A couple of die stack packaging technologies shown in Figure 8 where the use of a window chip scaled package with an interposer and both die facing opposite directions or wire bonding to 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 single memory module. However, due to electrical limitations, the maximum achievable is 72 for now.

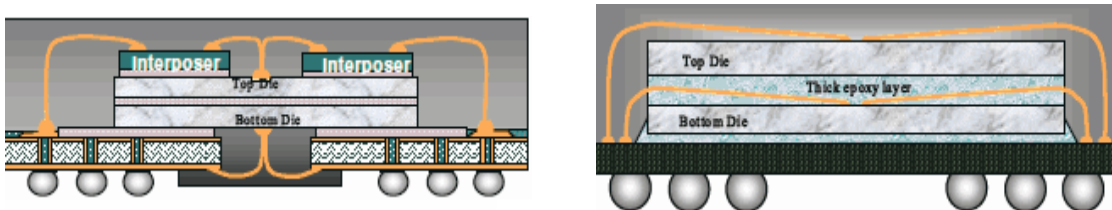


Figure 8. Die stack packaging

Figure 9 shows two registered DIMMs implementing dual-rank and quad rank with dual-die stacked DRAMs. The dual-rank module only requires two chip-select signals while quad-rank would require four, one for each rank.

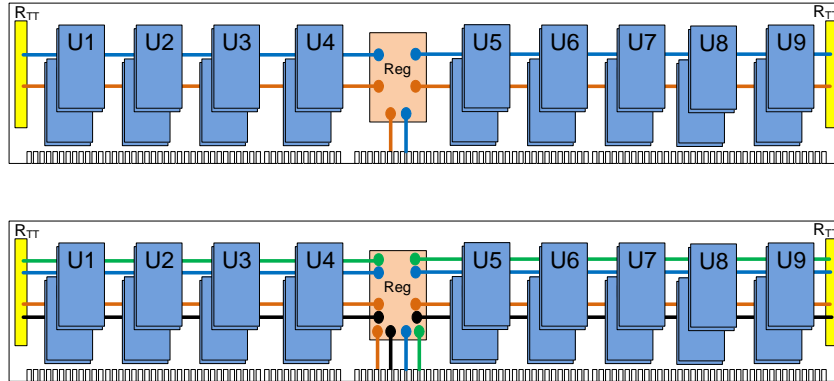


Figure 9. Dual-rank and Quad-rank RDIMM w/ stacked dual-die DRAM

The implementation of quad-die stacked DRAMs on modules will be commonplace heading into the DDR3 memory technology as the demand for higher density modules increases and smaller form factor modules such as VLP RDIMMs and SO-RDIMMs become the more popular.

### Low Power DDR Technology

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

The DDR market is currently addressing the need for low-power, cost-saving solutions. As previously mentioned, the DDR2 technology operates at 1.8 volts, which is about a 17% increase from DDR3, but there is a push for a lower voltage node down to 1.55 volts. A similar trend is beginning to develop in DDR3 where the original standard operating voltage is 1.5 volts but being pushed further down to a 1.35 volts 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 volts to 1.35 volts yields a 25% reduction in power per DIMM alone with some servers supporting up to 18 DIMMs. This is not only significant in the core of the memory and logic chips of the DIMMs but in the I/Os of both where the

majority of the power is consumed. As the operating voltage scales down so will the memory interface technology.

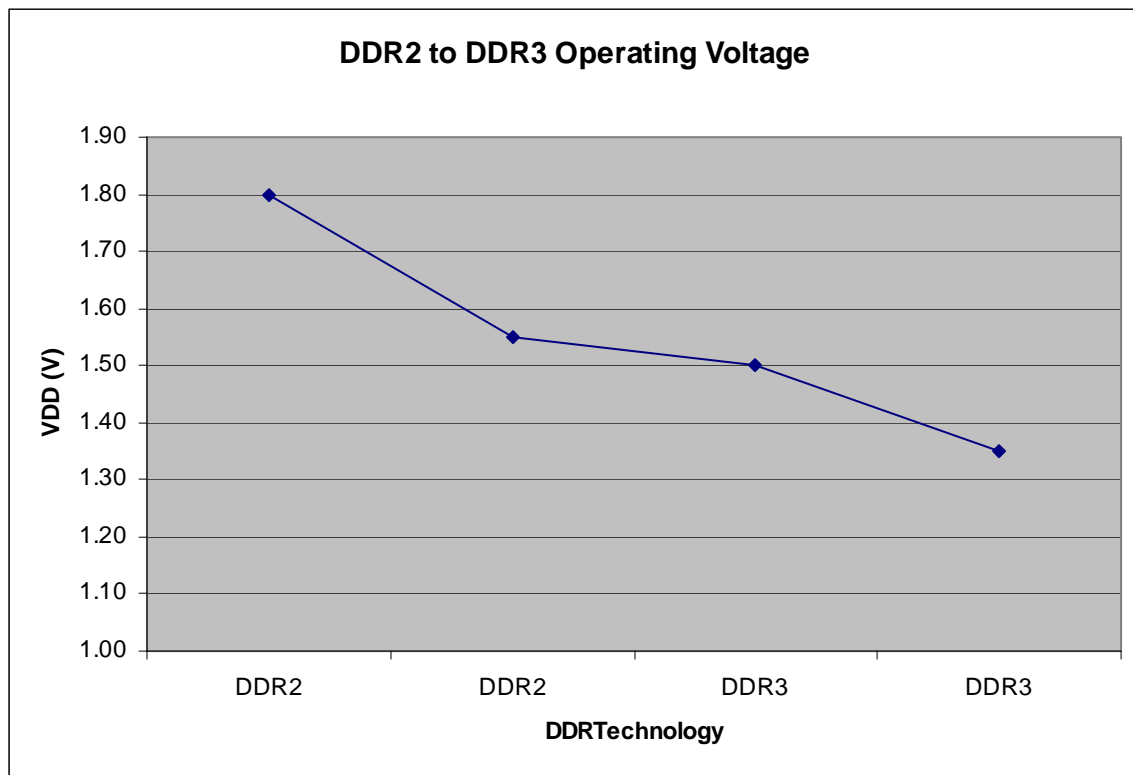


Figure 10. DDR2 to DDR3 Operating Voltage trend chart

Other areas of the memory interface technology are being studied for power reduction such as implementing special powerdown features and modes when the system is idling, more so in the server market.

### Conclusion

From mobile applications to notebooks to enterprise servers, there is a constant demand for higher bandwidth, lower power and improved throughput efficiency. All aspects of the DDR memory interface technology is being looked at in order to meet the needs and demands. The onus is on the DDR chip vendors to adapt and provide integrated circuit chips such as the DDR DRAM, register, and PLL with improved signal performance. Through JEDEC, new and innovative methods to improve signal integrity from one generation of DDR to the next have been defined in order to facilitate memory-driven application requirements. The DDR memory interface technology will be

continually pushed to its maximum limits and enhanced to meet industry needs. It is clear that DDR3 memory is addressing the need for higher density modules with lower power and efficient throughput.

## References

- [1] DDR2 SDRAM Specifications, JEDEC JESD79-2E, April 2008
- [2] DDR3 SDRAM Specifications, JEDEC JESD79-3B, April 2008
- [3] DDR3 Register Specifications, JEDEC Standard, October 2008
- [4] Liu Hao, Lim Beng Kuan, Desmond Y.R. Chong, Palei Win, Tan Hien Boon, "Stacked Die Technology Solutions for wCSP Memory Packages," EMPC, June 2005
- [5] Raj Mahajan, "Memory Design Considerations when Migrating to DDR3 Interfaces from DDR2," DesignCon 2007
- [6] [http://www.supertalent.com/datasheets/VLP\\_WHITEPAPER.pdf](http://www.supertalent.com/datasheets/VLP_WHITEPAPER.pdf)
- [7] <http://www.entorian.com/products/stakpak.html>