



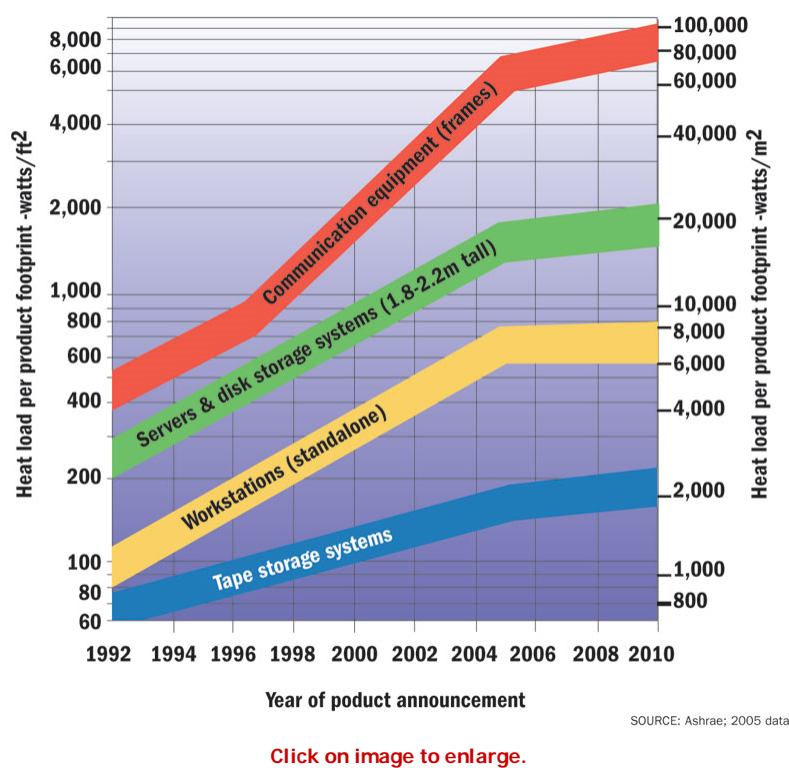
## Going green with DDR3 memory

Victor Cai  
 (01/26/2009 12:01 AM EST)  
 URL: <http://www.eetimes.eu/212902040>

Industry experts concede that about 82 percent of the \$500 million bill for a typical large data center is going toward the facility's mechanical and electrical infrastructure.

Industry experts concede that about 82 percent of the \$500 million bill for a typical large data center is going toward the facility's mechanical and electrical infrastructure.

For the data center server, one contributor to this higher power density is the memory subsystem. A typical 1U server requires a 600 to 1,000-W supply for its power. These high-end servers used in data centers support between 16 to 18 DRAM slots. On a DDR3 (double data rate 3) system, each single slot draws on average 9 W, using a typical 2-Gb module. The total memory subsystem power is around 144 to 162 W, which is nearly 25% of the available power for the system. Multiplying this by the power draw of 40,000 to 80,000 servers in a typical data center translates to between 5 to 13 MW (megawatts), just for the memory subsystem. Over 13,000 American homes can be powered using that same amount of power.



The memory community is pursuing three methods to improve the efficiency of the memory subsystem. These methods already have been tested in the lab environment and some forms of these methods will be rolling out in the commercial market in the near future.

The first method in power saving comes from the architecture of DDR3 memory subsystems. A key component of the DDR3 standard is the reduction of power. JEDEC, the leading developer of standards for the solid-state industry, is currently addressing two power-saving modes that will focus on how to reduce the power of registers and how to lower the power of the RDIMM (Registered Dual In-Line Memory Module) module. The first power-saving mode is the CKE (clock enable) power-down mode for the DDR3 register that requires the register device to shut off its outputs when both CKE input lines are low. When this mode is enabled by the memory controller, address outputs are tri-stated and the DDR3 register reduces its power consumption nearly 60 percent.

The second, power-saving mode is the S3 Power Management mode that allows the input clocks to float. When both inputs CK and CK# are being held LOW, the device stops operation and enters low-power static and standby operation. Measurement in the lab showed register power was reduced by 90 percent. The Clock Stopped power-down mode can only be utilized once the DRAM receives a self-refresh command prior to mode activation. In this state, the DRAM ignores all inputs except CKE. Hence, all register outputs except for two signals can be disabled further, lowering the overall power of the RDIMM module. However, the time required to wake up the memory subsystem from these two modes are different. While it only takes three clock cycles for the memory subsystem to start responding from a CKE power-down, only microseconds are needed coming out of the S3 Power Management mode. Memory-controller and system-management software are needed to support both of these functions to take full advantage of these power-saving features.

The second method is being explored by the memory communities to lower the supply voltage for the memory subsystem. This option explores migration of lower geometries in semiconductor processes, which allows reduction of the supply voltage.


New 1.35V-capable DDR3 DRAMs are being tested in limited quantities in laboratories. JEDEC and major memory manufacturers have been looking at the next generation of DDR3 DRAM that will operate both at 1.5 V and 1.35 V. Inphi Corp. is leading the effort to draft a 1.35-V register standard for the JEDEC organization. Inphi also has taken a leadership role in this effort with the first production-ready 1.35-V DDR3 register that is backward compatible to 1.5V. Preliminary testing on selected low-voltage DRAMs is already taking place at some OEM sites to run memory subsystems at 1.35 V. Early data shows approximately a 16 percent reduction in total power with little to no appreciable performance penalty. More testing is being done, and research is taking place to reduce operating voltages even further with future technologies.

The third and final method to conserve power is to reduce the output-swing voltage level for the driver. In the DDR3 RDIMM standard, command and address signals on the DIMM are terminated to Vtt, which is half of Vdd with 22-Ohm resistors. The minimum acceptable input level for DRAM is +/- 175 mV around Vtt. For reliable operation, the designer has provided three to four times that value even on lower density cards like Raw Card A. Reducing the output swing from 800 mV to just 400 mV can save around 20 mA per channel. With 41 channels per DIMM and 18 DIMMs per system, that is 22 W that potentially could be conserved per system and approximately 2 MW in a typical data center alone!

The energy challenges facing data centers need to be addressed to explore conservation solutions. Even with more efficient power-saving modes, reducing power supply voltages and lowering switching voltages, memory subsystems will still be one of the largest consumers of energy in a server system

---

*Victor Cai is an applications engineer with Inphi Corp. He holds B.S. and M.S. degrees in electrical engineering.*



All materials on this site [Copyright © 2009 TechInsights, a Division of United Business Media LLC.](#), [EETimes EU Copyright](#) All rights reserved.  
[Privacy Statement](#) | [Terms of Service](#)