Power Variability in Contemporary DRAMs

Mark Gottscho, *Student Member, IEEE*, Abde Ali Kagalwalla, *Student Member, IEEE*, and Puneet Gupta, *Member, IEEE*

***Abstract*—Technology scaling has led to significant variability in chip performance and power consumption. In this work, we measured and analyzed the power variability in dynamic random access memories (DRAMs). We tested 22 double date rate third generation (DDR3) dual inline memory modules (DIMMs), and found that power usage in DRAMs depends on both operation type (write, read, and idle) as well as data, with write operations consuming more than reads, and 1s in the data generally costing more power than 0s. Temperature had little effect (1–3%) across the −50 °C to 50 °C range. Variations were up to 12.29% and 16.40% for idle power within a single model and for different models from the same vendor, respectively. In the scope of all tested 1 gigabyte (GB) modules, deviations were up to 21.84% in write power. Our ongoing work addresses memory management methods to leverage such power variations.**

***Index Terms*—Double data rate third generation (DDR3), dynamic random access memory (DRAM), power, variability.**

I. Introduction

MODERN digital integrated circuits (ICs) exhibit significant variability as a consequence of imperfections in the fabrication processes [1], [2], use patterns, aging, and the environment [3]. The typical approach of guardbanding for variability is expensive [4]. As a result, there is growing interest in software as well as hardware mechanisms that adapt to variations or compensate for them. Examples of explicit variation-awareness in the software stack include power management [5], embedded sensing [6], and video encoding [7].

To develop effective methods of addressing variations (especially in the software layers), it is important to understand the extent of variability in different components of computing systems and their dependence on the workload and environment. Though variability measurements through simple silicon test structures abound (e.g., [8], [9]), variability characterization of full components and systems have been scarce. Moreover, such measurements have been largely limited to processors (e.g., 14X variation in sleep power of embedded microprocessors [6] and 25% performance variation in an experimental 80-core Intel processor [10]). For a large class of applications, memory power is significant (e.g., 48% of total power in [11]) which has motivated several efforts to reduce dynamic random access memory (DRAM) power consumption (e.g., power-aware virtual memory systems [12]–[14]). These designs reduced power consumption of main memory, but they did not take into account hardware variability; instead they assumed all DRAMs to be equally power efficient.

A "variability-aware" operating system could exploit variability in DRAM, making physical allocation decisions in real time to reduce overall memory power consumption. Memory virtualization techniques such as in [15] could be extended to account for power variations in off-chip memories. To the best of our knowledge, [16] is the only study to present measured DIMM power variability; it explored running systems, including component-level sources such as CPUs and DDR2 DRAMs but primarily focused on vendor-dependent variations. The study in [17] included an investigation on operation and data dependence of memory power, but used SRAMs on an older 0.35 μm process node.

In our work, an Intel Atom-based testbed was constructed, running a modified version of Memtest86 [18] in order to control memory operations at a low level. We analyzed the write, read, and idle power consumption of several mainstream double data rate third generation (DDR3) dual inline memory modules (DIMMs), comprised of parts from several vendors and suppliers. The key contributions of this work are as follows:

- analysis of instance, vendor and temperature dependent power variability in contemporary DRAMs;
- characterization of power dependence on inputs and operation.

We begin with an overview of the test methodology in Section II, followed by an investigation of operation, data, and temperature dependencies in DRAM power as well as power variability in our set of memory modules in Section III. This work concludes with an overview of findings and suggestions for future work.

II. Test Methodology

A. Memory Equipment

Our DIMMs were comprised of several models from four vendors (see Table I), manufactured in 2010 and 2011 (the particular process technologies are unknown). For five of the DIMMs, we could not identify the DRAM suppliers. Most models were 1 GB[1] DDR3 modules, rated for 1066 MT/s (except for the Vendor 4 models, rated for 1800 MT/s) with a specified supply voltage of 1.5 V. We also included three 2 GB specimens from Vendor 1 to see if capacity had any effect on power consumption. The DIMMs are referred to henceforth by abbreviations such as V1S1M1 for Vendor 1, Supplier 1, Model 1.

[1]To avoid confusion in terminology, we refer to the gigabyte (GB) in the binary sense, i.e., 1 GB is 2^{30} bytes, not 10^9 bytes.

Manuscript received November 12, 2011; accepted March 14, 2012Date of publication April 03, 2012; date of current version June 15, 2012. Funding for this research was provided by the National Science Foundation (NSF) Variability Expedition under grant number CCF-1029030.This manuscript was recommended for publication by R. Kastner.

The authors are with the NanoCAD Laboratory in the Department of Electrical Engineering, University of California Los Angeles, Los Angeles, CA 90024 USA (e-mail: mgottscho@ucla.edu; abdeali@ucla.edu; puneet@ee.ucla.edu).

Digital Object Identifier 10.1109/LES.2012.2192414

TABLE I
DDR3 DIMM SELECTION

Category	Quantity
Vendors	4 (V1-V4)
Suppliers	3 known (S1-S3, SU)
Capacities	1 GB (V1-V4), 2 GB (V1 only)
Models	Up to 3 per vendor, 8 total (7 were 1 GB)
Total DIMMs	22 (19 were 1 GB)

TABLE II
TESTBED AND MEASUREMENT PARAMETERS

Parameter	Value
Testbed CPU	Intel Atom D525 @ 1.8 GHz
Number of CPU Cores Used	1
Cache Enabled	Yes
DIMM Capacities	1 GB, (2 GB)
DIMM Operating Clock Freq.	400 MHz
Effective DDR3 Rate	800 MT/s
DIMM Supply Voltage	1.5 V
Primary Ambient Temp.	30 °C
Secondary Ambient Temp.	-50, -30, -10, 10, 40, 50 °C
Primary Memory Test Software	Modified Memtest86 v3.5b [18]
Custom Test Routines	Seq. Write Pattern, Seq. Read, Idle
Digital Multimeter	Agilent 34411A
Sampling Frequency	10 ksamples/sec
Reading Accuracy	approx. 4.5 mW
Number of Samples Per Test	200000

B. Test Platform & Data Acquisition

The test platform utilized an Intel Atom D525 CPU running at 1.80 GHz, on a single core. Only one DIMM was installed at a time on the motherboard, and all other hardware was identical for all tests. No peripherals were attached to the system except for a keyboard, VGA monitor, and a USB flash drive containing the custom test routines. An Agilent 34411A digital multimeter sampled the voltage at 10 ksamples/s across a small 0.02 Ω resistor inserted on the V_{DD} line in between the DIMM and the motherboard slot, and this was used to derive the power consumption. Ambient temperature was regulated using a thermal chamber.

Because we required fine control over all memory I/Os, we developed custom modifications to Memtest86 v3.5b, which is typically used to diagnose memory faults [18]. The advantage of using Memtest86 as a foundation was the lack of any other processes or virtual memory, which granted us the flexibility to utilize memory at a low level.

We created a write function which wrote memory sequentially with a specified bit pattern, but never read it back. Similarly, a read function was created which only read memory sequentially without writing back. Each word location in memory could be initialized with an arbitrary pattern before executing the read test. The bit fade test, which was originally designed to detect bit errors over a period of DRAM inactivity, was modified to serve as an idle power test, with minimal memory usage.[2] For all tests, the cache was enabled to allow for maximum memory bus utilization. With the cache disabled, we observed dramatically lower data throughput and were unable to distinguish power differences between operations. As our intent was primarily to measure power variability between different modules, we used sequential access patterns to avoid the effects of caches and row buffers.

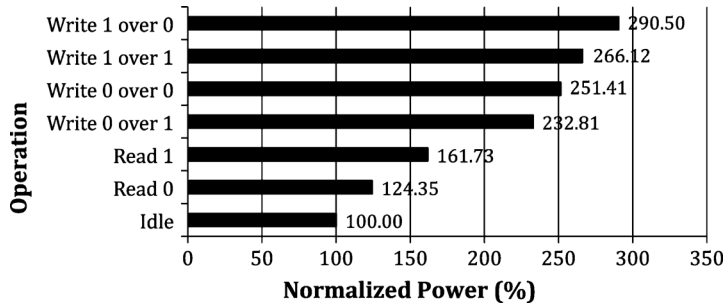

Fig. 1. Data and operation dependence of DIMM power.

Each test was sampled over a 20-s interval, during which several sequential passes over the entire memory were made. This allowed us to obtain the average power for each test over several iterations. Each reading had an estimated accuracy of 0.06 mV [19], which corresponds to approximately 4.5 mW assuming a constant supply voltage and resistor value. Table II summarizes the important test environment parameters. For further details on the testing methodology, see [20].

III. TEST RESULTS

A. Data Dependence of Power Consumption

Since DRAM power consumption is dependent on the type of operation as well as the data being read or written [21], we conducted experiments to find any such dependencies. Note that the background, pre-charge, and access power consumed in a DRAM should have no dependence on the data [22]. Note that this test is similar to one performed on SRAMs in [17].

Seven tests were performed on four DIMMs, each from a different vendor, at 30 °C to explore the basic data I/O combinations. The mean power for each test was calculated from the results of the four DIMMs. Fig. 1 depicts the results for each test with respect to the idle case ("Write 0 over 0" refers to continually writing only 0s to all of memory, whereas "Write 1 over 0" indicates that a memory full of 0s was overwritten sequentially by all 1s, and so on). Note that for the idle case, there was negligible data dependence, so we initialized memory to contain approximately equal amounts of 0s and 1s.

Interestingly, the power consumed in the operations was consistently ordered as seen in Fig. 1, with significant differences as a function of the data being read or written. There was also a large gap in power consumption between the reading and writing for all data inputs.

We presume that the difference between the write 0 over 0 case and the read 0 test is purely due to the DRAM I/O and peripheral circuitry, as the data in the cell array is identical. This would also apply to the write 1 over 1 case and its corresponding read 1 case. In both the read 1 and write 1 over 1 cases, more

[2]Although there are different "idle" DRAM states, they are not directly controllable through software; we did not distinguish between them.

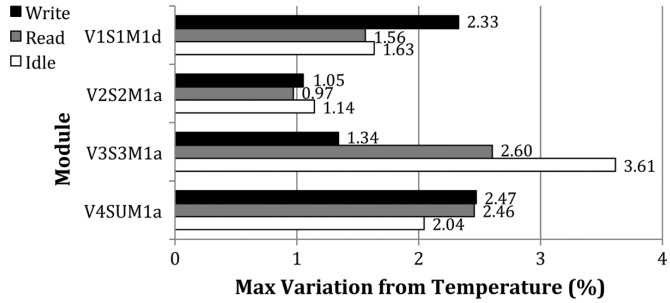

Fig. 2. Relative temperature effects on write, read, and idle DIMM power, −50 °C to 50 °C Range.

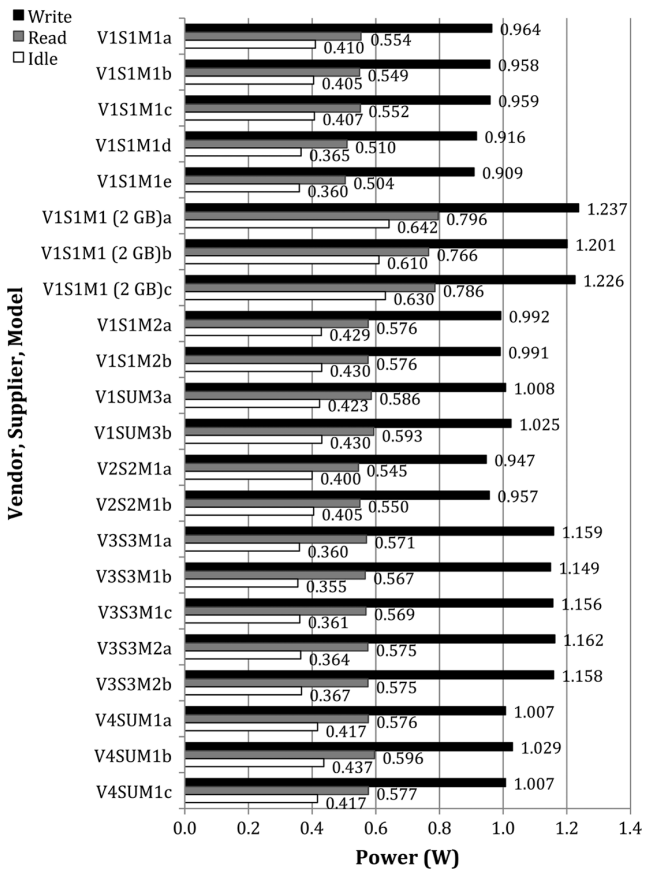

Fig. 3. Write, read, and idle power by DIMM, 30 °C.

power was consumed compared to the corresponding read 0 and write 0 over 0 cases. These deltas may be due to the restoration of cell values. Because a sense operation is destructive of cell data due to charge sharing [21], cells that originally contain 1s must be restored using additional supply current. In contrast, cells containing 0s need only be discharged.

Note that the write 0 over 1 test consumed *less* power than the write 0 over 0 test, whereas the write 1 over 0 case consumed *more* than the write 1 over 1 case. The write 1 over 0 case likely consumes the most power because the bit lines and cells must be fully charged from 0 to 1. In the write 0 over 1 case, it probably uses the least power because the bit lines and cells need only be discharged to 0. Further research and finer-grained measurement capabilities are required to fully explain these systematic dependencies. Nevertheless, these results indicate strong data and operation dependence in DRAM power consumption.

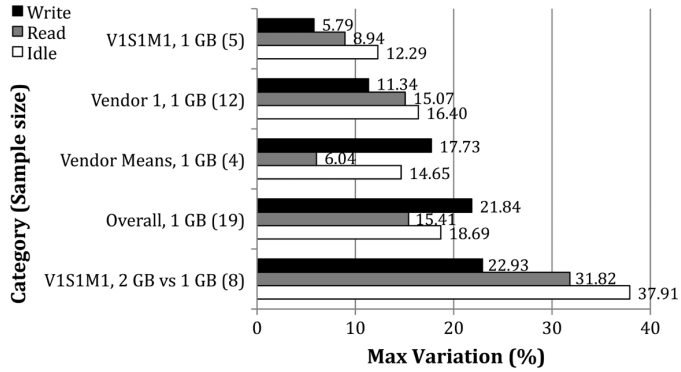

Fig. 4. Maximum variations in write, read, and idle power by DIMM category, 30 °C.

Because of the data dependencies in write and read operations, we decided to use memory addresses as the data for write and read in all subsequent tests, because over the entire address space, there are approximately equal quantities of 1s and 0s. Furthermore, memory addresses are common data in real applications. We verified that the average write and read power using addresses for data is approximately the same as the mean of the values for 1s and 0s as data.

B. Temperature Effects

To determine if temperature has any effect on memory power consumption, we tested four 1 GB modules, one from each vendor. Each DIMM was tested at ambient temperatures from −50 °C to 50 °C.[3] It is clear from Fig. 2 that temperature had a negligible effect on power consumption even across a large range. We speculate that this is partially due to the area and leakage-optimized DRAM architecture [1], but more substantially affected by modern refresh mechanisms. The use of rolling refreshes or conservative timings may consume significant dynamic power, overshadowing the temperature dependent components in the background power consumption. Since no DIMM exhibited more than 3.61% variation across a 100 °C range, all further tests were performed at an ambient temperature of 30 °C.

C. DIMM Power Variations

A plot of write, read, and idle power consumption for all 22 DIMMs at 30 °C is depicted in Fig. 3. The variability results are summarized in Fig. 4.

1) Variability Within DIMMs of the Same Model (1 GB): Consider a particular model,V1S1M1 in Fig. 3, of which we had the largest number (five) of specimens. While there was a maximum of 12.29% difference between the five DIMMs, there is a visible gap between the first group of three DIMMs and the second group of two (fourth and fifth in Fig. 3). This may be because the DIMMs come from two different production batches, resulting in lot-to-lot variability. The maximum deviations within the first group was only 1.34% for idle, and 1.47% within the second group. This suggests that the majority of the variation in V1S1M1 was between the two batches.

2) Variability Between Models of the Same Vendor/Supplier (1 GB): Now, consider all DIMMs from Vendor 1. We would

[3]Testing above an ambient temperature of 50 °C was not practical as it caused testbed hardware failure.

expect that there would be more variation in Vendor 1 overall than in V1S1M1 only, and this was confirmed in the data. The maximum variation observed in Vendor 1 (1 GB) was 16.40% for the idle case. This variability may be composed of batch variability or performance differences between models.

3) Variability Across Vendors (1 GB): In order to isolate variability as a function of vendors and to mitigate any effects of different sample sizes, we computed the mean powers for each vendor (1 GB). Vendor 3 consumed the most write power at 1.157 W. The variations for write, read, and idle power were 17.73%, 6.04%, and 14.65% respectively.

4) Overall Variability Amongst 1 GB DIMMs: As one may have expected, the variations across all DIMMs were significantly higher than within a model and among vendors, with the maximum variation occurring for write power at approximately 21.84%.

5) Effects of Capacity on Power Consumption: It is clear from Fig. 3 that the three 2 GB DIMMs of V1S1M1 consumed significantly more power than their 1 GB counterparts. This was expected, as there was bound to be higher idle power with twice as many DRAMs (in two ranks instead of one). Indeed, the maximum variation between the 2 GB and 1 GB versions was 37.91%, which occurred for idle power, whereas write power only differed by half as much. This is because background power is a smaller proportion of overall power when the DIMM is active.

IV. Conclusion and Future Work

We analyzed the power consumption of several mainstream DDR3 DIMMs from different vendors, suppliers, and models, and found several important trends. Firstly, we did not find any significant temperature dependence of power consumption. Manufacturing process induced variation (i.e., variation for the same model) was up to 12.29%. Among models from the same vendor, idle power generally varied the most (up to 16.40% among Vendor 1), followed by read and write power. However, a different trend was evident across vendors, with write power varying the most (up to 17.73%), followed by idle and read power. This pattern was dominant overall among all tested 1 GB DIMMs, where we observed up to 21.84% variations in write power. Lastly, we found that a 2 GB model consumed significantly more power than its matching 1 GB version, primarily due to its increased idle power (up to 37.91%). Data-dependence of power consumption was also very pronounced, with about 30% spread within read and 25% spread within write operations. These findings serve as a motivation for variability-aware software optimizations to reduce memory power consumption. In an arbitrary set of DIMMs, there can be considerable variation in power use, and an adaptable system can use this to its advantage. Because we observed negligible temperature dependence, we will not include it in our future models of DIMM power.

Our ongoing work aims to optimize for any variability present in a system's memory through a software-based approach. For example, a modified Linux memory management system could make variability-aware page allocation and migration decisions. We will consider different approaches toward this goal, such as a fully transparent decision mechanism, or semi-transparent methods that require compiler and/or application support.

Acknowledgment

The authors would like to thank L. Wanner of the UCLA Networked and Embedded Systems Laboratory (NESL) for kindly aiding us with instrumentation.

References

[1] J. M. Rabaey, A. Chandrakasan, and N. Borivoje, "Designing memory and array structures," in *Digital Integrated Circuits—A Design Perspective*, 2nd ed. Upper Saddle River, NJ: Pearson Education, Inc., 2003, pp. 623–717.

[2] S. Borkar, "Designing reliable systems from unreliable components: The challenges of transistor variability and degradation," *IEEE Micro,*, vol. 25, no. 6, pp. 10–16, 2005.

[3] K. Bernstein, D. J. Frank, A. E. Gattiker, W. Haensch, B. L. Ji, S. R. Nassif, E. J. Nowak, D. J. Pearson, and N. J. Rohrer, "High-performance CMOS variability in the 65-nm regime and beyond," *IBM J. of Res. Develop.*, vol. 50, no. 4.5, pp. 433–449, July 2006.

[4] K. Jeong, A. B. Kahng, and K. Samadi, "Impact of guardband reduction on design outcomes: A quantitative approach," *IEEE Trans. Semicond. Manuf.*, vol. 22, no. 4, pp. 552–565, Nov. 2009.

[5] K. Lahiri, A. Raghunathan, S. Dey, and D. Panigrahi, "Battery-driven system design: A new frontier in low power design," in *Proc. ASP-DAC, 15th Int.Conf. VLSI Design*, 2002, pp. 261–267.

[6] L. Wanner, R. Balani, S. Zahedi, C. Apte, P. Gupta, and M. Srivastava, "Variability-aware duty cycle scheduling in long running embedded sensing systems," in *Proc. Design, Automat. Test Eur. Conf. Exhibit. (DATE)*, Mar. 2011, pp. 1–6.

[7] A. Pant, P. Gupta, and M. van der Schaar, "AppAdapt: Opportunistic application adaptation in presence of hardware variation," *IEEE Trans. Very Large Scale Integr.(VLSI) Sys.*, vol. PP, no. 99, pp. 1–11, 2011.

[8] L. Cheng, P. Gupta, C. J. Spanos, K. Qian, and L. He, "Physically justifiable die-level modeling of spatial variation in view of systematic across wafer variability," *IEEE Trans. Comput.-Aided Design Integr. Circuits Sys.*, vol. 30, no. 3, pp. 388–401, Mar. 2011.

[9] L. Pang, K. Qian, C. J. Spanos, and B. Nikolic, "Measurement and analysis of variability in 45 nm Strained-Si CMOS technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 8, pp. 2233–2243, Aug. 2009.

[10] S. Dighe, S. Vangal, P. Aseron, S. Kumar, T. Jacob, K. Bowman, J. Howard, J. Tschanz, V. Erraguntla, and N. Borkar, "Within-die variation-aware dynamic-voltage-frequency scaling core mapping and thread hopping for an 80-core processor," in *Proc. IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers (ISSCC)*, 2010, pp. 174–175.

[11] K. Rajamani, C. Lefurgy, S. Ghiasi, J. C. Rubio, H. Hanson, and T. Keller, "Power management for computer systems and datacenters," in *Proc. Int. Symp. Low Power Electron. Design Tutorial*, 2008 [Online]. Available: http://www.islped.org/

[12] C. Lefurgy, K. Rajamani, F. Rawson, W. Felter, M. Kistler, and T. Keller, "Energy management for commercial servers," *Computer*, vol. 36, no. 12, pp. 39–48, 2003.

[13] V. Delaluz, A. Sivasubramaniam, M. Kandemir, N. Vijaykrishnan, and M. J. Irwin, "Scheduler-based DRAM energy management," in *Proc. 39th Annu. Design Automat. Conf.*, 2002, pp. 697–702.

[14] H. Huang, K. G. Shin, C. Lefurgy, and T. Keller, "Improving energy efficiency by making DRAM less randomly accessed," in *Proc. Int. Symp. Low Power Electron. Design*, 2005, pp. 393–398.

[15] L. A. D. Bathen, N. D. Dutt, D. Shin, and S. Lim, "SPMVisor: Dynamic scratchpad memory virtualization for secure, low power, and high performance distributed on-chip memories," in *Proc. IEEE 9th Int. Conf. HW/SW Codesign Syst. Synthesis (CODES+ISSS).* , Oct. 2011, pp. 79–88.

[16] H. Hanson, K. Rajamani, J. Rubio, S. Ghiasi, and F. Rawson, "Benchmarking for power and performance," in *SPEC Benchmark Workshop*, 2007 [Online]. Available: http://www.spec.org/workshops/2007/austin/papers/Benchmarking_Power_for_Performance.pdf

[17] J. Hezavei, N. Vijaykrishnan, and M. Irwin, "A comparative study of power efficient SRAM designs," in *Proc. 10th Great Lakes Symp. VLSI*, 2000, pp. 117–122.

[18] Memtest86 Test Algorithms [Online]. Available: http://memtest86.com/

[19] *"34410A/11A 6-1/2 Digit Multimeter User's Guide,"* Agilent , Santa Clara, CA, 2007 [Online]. Available: http://www.home.agilent.com/

[20] M. Gottscho, "Analyzing Power Variability of DDR3 Dual Inline Memory Modules," Aug. 2011 [Online]. Available: http://nanocad.ee.ucla.edu/pub/Main/Publications/UG2_paper.pdf

[21] K. Itoh, *VLSI Memory Chip Design*, 1st ed. New York: Springer, Apr. 2001.

[22] Calculating Memory System Power for DDR3 [Online]. Available: http://www.micron.com