An Analysis on Retention Error Behavior and Power Consumption of Recent DDR4 DRAMs

Deepak M. Mathew, Martin Schultheis, Carl C. Rheinländer, Chirag Sudarshan, Christian Weis, Norbert Wehn University of Kaiserslautern, Germany Email: {deepak,weis,wehn}@eit.uni-kl.de

Abstract—DRAM technology is scaling aggressively that results in high leakage power, worse data retention time behavior, and large process variations. Due to these process variations, vendors provide large guard bands on various DRAM currents and timing specifications that are over pessimistic. Detailed knowledge on the DRAM retention behavior and currents for the average case allow to improve memory system performance and energy efficiency of specific applications by moving away from worst case behavior. In this paper, we present an advanced measurement platform to investigate off-the-shelf DDR4 DRAMs' retention behavior, and to precisely measure various DRAM currents (IDDs and IPPs) at a wide range of operating temperatures. Error Checking and Correction (ECC) schemes are popular in correcting randomly scattered single bit errors. Since retention failures also occur randomly, ECCs can be used to improve DRAM retention behavior. Therefore, for the first time, we show the influence of ECC on the retention behavior of recent DDR4 DRAMs, and how it varies across various DRAM architectures considering detailed structure of the DRAM (true-cell devices / mixed-cell devices).

I. INTRODUCTION

Its well known that more and more applications are memory centric. This puts Dynamic Random Access Memory (DRAM) in the focus to improve performance and energy efficiency of advanced computing systems. DRAM technology is scaling aggressively to fullfill the huge demands for data, which requires main memories in large capacity. As capacity increases, the time and energy overhead for refreshing the leaky DRAM cells are also increasing. It's a known fact that typically DRAM cells exhibit much longer retention times than specified [1], [2], [3], [4], and therefore the refresh interval can be reduced if the retention characteristics are known. There are previous investigations that make use of this typical DRAM cell retention behavior, especially for applications which have inherent error resilience: the so-called Approximate DRAM (ADRAM) [5], [6]. The main goal of those works were to reduce the energy and the performance penalty due to refresh. However, a very detailed knowledge of the DRAM retention behavior is crucial for these techniques to be effective.

Since the DRAM technology is approaching the end of scaling, it suffers more from various scaling related problems, such as higher leakage, lower charge retention time, larger process variations etc. [7]. Due to higher leakage and lower cell capacitance, the retention behavior of DRAMs varies for different technology nodes. Therefore, it is important to characterize the retention behavior of recent DRAMs in lower

Matthias Jung Fraunhofer Institute for Experimental Software Engineering (IESE) Kaiserslautern, Germany Email: matthias.jung@iese.fraunhofer.de

technology nodes at different temperatures before using them in the context of Approximate DRAMs.

As a result of another effect of scaling, which is larger process variations at lower technology nodes, DRAM vendors provide worst case currents in datasheets, which deviate largely from the nominal case. Various high level DRAM power estimation tools rely on those pessimistic datasheet values. Hence, to accurately estimate energy contribution of DRAMs for different applications, all DRAM currents have to be measured for the nominal case at different operating temperatures.

To the best of our knowledge, there exists no such measurement platform to obtain retention characteristics or currents for state-of-the-art DDR4 DRAMs. We present a novel platform to measure the DDR4 DRAMs' retention behavior and currents (at voltage domains VDD and VPP) for temperatures up to 95 °C. Using our advanced measurement platform, we first show the influence of temperature on the retention behavior of DDR4 DRAMs, and how the vendor specific internal architecture of DRAM affects the retention time behavior. Also, we compare DDR4 with DDR3 DRAMs' retention behavior to study the impacts of scaling.

ECC techniques are well known in correcting randomly scattered single bit errors. Since retention failures also occur randomly, ECCs can be used to improve DRAM retention behavior. However, there are not many investigations so far with the use of ECC techniques to mitigate the retention errors in off-the-shelf DDR4 memories. Therefore, we demonstrate the effectiveness of the most commonly used (72,64) Hamming Codes or Single Error Correction Double Error Detection (SECDED) ECC on DDR4 DRAMs' retention behavior. Our results show that benefits of ECC depends largely on the DRAM architecture (true-cell device / mixed-cell device). For mixed-cell devices, the failure probability can be reduced by up to 4 orders of magnitude using SECDED. It also verifies that simple Hamming codes are ineffective for longer refresh intervals and at higher temperatures. Finally, we provide the measurement results of various DDR4 DRAM currents (power consumption) and show how they vary with increasing temperature.

II. RELATED WORK

Most of the previous works on the DRAM retention profiling and the measurement platforms are listed in [3], [8]. In their work, they present *DRAMMeasure*, a low cost platform for measuring the retention errors and power consumption of DDR3 DRAMs. However, this platform cannot be used to measure DDR3 ECC-DIMMs or new DDR4 SO-DIMMs because they are not pin compatible and DDR4 features a much higher level of complexity, due to different voltage domains and higher frequencies. Most recently the authors of *REAPER* [9] presented a detailed study on 368 LPDDR4 devices. Since LPDDR4 devices are usually used in the *Package on Package* (PoP) method i.e. they are soldered on top of an *Multiprocessor System on Chip* (MPSoC) package, a sophisticated measurement setup would be required. However, the authors do not disclose any details about their measurement platform.

Various prior works related to refresh reduction are listed in [6], where they disable the refresh completely for applications which can tolerate certain degree of bit errors due to retention failures. Recently, Samsung [10] and Hynix [11] present new LPDDR4 DRAMs for IoT and wearable applications, which lower the refresh rates by a factor of $4\times$ in order to reduce power consumption. To avoid the occurance of retention errors, a on-die ECC technique is employed. Therefore, it is important to understand the effects of ECC on the DRAM's retention time [12].

In order to perform an analysis on retention error behavior and power consumption of state-of-the-art DDR4 DRAMs, we developed a custom experimental setup, which we will present in the following Section III.

III. EXPERIMENTAL SETUP



Figure 1: DDR4 Measurement Platform and Adapter Board

To analyze DDR4 DRAMs we developed a custom platform, shown in Figure 1, which is similar to our previous platform designed for DDR3 [8], but with various improvements. It is designed to measure power consumption, retention errors, and to heat up the DRAM devices of DDR4 SO-DIMM modules. The heating section consists of a mechanical setup, which is placed on the surface of the DRAM devices to heat them up within a range of 25°C to 95°C. The accuracy of the temperature control was determined to 2°C using thermal simulations.

To analyze the DDR4 DRAM currents of V_{DD} and V_{PP} voltage domains, we designed a JEDEC-conform adapter board for DDR4 SO-DIMMs, which is shown in Figure 1. Due to the DDR4 standard, this board can be used for ECC and non-ECC SO-DIMM modules. The power lines V_{DD} and V_{PP} are routed across $4 \text{ m}\Omega$ shunt resistors, whereas the data, address, and control lines are passed through. Due to precise

impedance-matched layout design, the adapter board works with DRAM clock frequencies greater than 1 GHz.

The voltages across the shunt resistors are amplified with current-sense amplifiers. High-precision 24-bit Analog to Digital Converters are synchronously sampling and converting these voltages into digital values. By using on-board converters, environmental impacts on the sensitive voltages are reduced significantly. We utilized calibrated precision measurement instruments by Keithley to ensure a current measurement accuracy of +0.5 mA/-0 mA.

The adapter board is designed to have the same dimensions as a regular DDR4 SO-DIMM. This way it can be directly plugged in to almost any kind of DDR4 host system. The test set-up for retention error measurements is similar to the one used for our previous DDR3 measurement platform [8].

IV. EXPERIMENTAL RESULTS

A. Retention Time Analysis

Using the platform described in Section III, we conducted retention measurements on off-the-shelf DDR4 DRAMs from two major vendors (*Vendor-A* and *Vendor-B*). Based on the internal architecture of DRAMs there exist different ways how information is stored:

- **True-Cell** bit storage, where a logical 1 is always stored at bitline high voltage (e.g. 1.1 V) and a logical 0 is stored at low voltage (0 V).
- Anti-Cell bit storage, where a logical 1 is stored as bitline low voltage (0 V) and a logical 0 is stored at high voltage (1.1 V) value – i.e. the bit value is stored inverted.
- **Mixed-Cells** a combination of both, true and anti-cells depending on the address of the accessed cell.

One of the objectives of our study was to evaluate the influence of those internal architectures of DRAM on its retention behavior. Therefore, for conducting measurements we chose two 4GB DDR4 SO-DIMMs with ECC capability: one with only true-cell DRAMs from *Vendor-A*, and the other with mixed-cell DRAMs from *Vendor-B*. Both SO-DIMMs consist of eight identical 4Gb DDR4 devices for storing data bits, and one additional device for storing ECC bits when ECC is enabled in the memory controller.

Figure 2 shows the retention behavior for *Vendor-A*, and Figure 3 depicts the retention behavior for *Vendor-B*. We plot retention times versus the cumulative cell failure probability obtained during each measurement step. Each device was tested with two data patterns: all 1s $(0 \times FF)$ and pseudorandom. The pseudo-random data pattern was generated in the memory controller using *Linear Feedback Shift Registers* (LFSRs). Each test was performed by first writing the complete DRAM with a specific data pattern, then waiting for a certain duration with refresh switched-off, and finally, reading back the written data from the entire DRAM. The read data is compared against the written data to count the number of retention failures.

Figure 2 shows the effect of data pattern on the retention behavior of a typical true-cell DRAM from *Vendor-A*. For shorter retention times (<1 s), and lower temperatures ($30 \degree$ C and $60 \degree$ C), the number of retention failures (cumulative failure





Figure 2: DRAM Retention Errors With and Without ECC for *Vendor-A* True-Cell DRAM

probability) are much higher for the random data pattern compared to the $0 \times FF$ data pattern. For random data pattern, retention failures start to occur even at 100 ms. This proves the very high *Data Pattern Dependency* (DPD) of the retention behavior of true-cell DRAMs. At very high temperatures (90 °C) and for longer retention times, the difference in the failure probability between $0 \times FF$ data pattern and the random data pattern decreases. This is because at higher temperatures and for longer retention times, there are already many bit flips from 1 to 0, for all 1's data pattern, such that the resulting data pattern is similar to the random data pattern.

Figure 3 shows the influence of data-pattern on the retention time behavior of a mixed cell DRAM. As opposite to the true cell devices, the cumulative failure probability with random data pattern is slightly reduced for the tail end (<1 s) of the retention curve compared to the $0 \times FF$ pattern. E.g., at 60 °C, bit flips start to occur from 500 ms for $0 \times FF$ pattern, while there are no bit flips till 1 s for random data pattern. This shows that mixed-cell DRAMs are more immune to retention failures due to DPDs.

Another goal of our investigation was to get a better insight on the influence of ECC on the retention behavior of various DRAM architectures. We used a (72,64) Hamming code for SEC-DED inside the DRAM controller. As shown in Figure 2, ECC is not much effective for true-cell DRAMs in correcting the retention failures due to random patterns. Using ECC, the number of failures can be reduced only by up to 1 order of magnitude in the tail of retention curve (<1 s). But, for mixed-

(b) Random Data Pattern

Figure 3: DRAM Retention Errors With and Without ECC for *Vendor-B* Mixed-Cell DRAM

cell devices, see Figure 3, ECC is much more effective in correcting the retention failures due to random data pattern. With ECC, failure probability in the lower end of retention curve (<1 s) can be reduced by 2 to 4 orders of magnitude even at (90 °C). Independent of DRAM architecture and data pattern, we also observe that typical (72,64) Hamming codes are not much effective in correcting the retention failures at higher temperatures (>60 °C) and for longer retention times (>10 s)

For ADRAMs, we mainly focus on the tail end of retention behavior (<1 s), and medium temperatures. Therefore mixedcell devices are better candidates for use in ADRAM applications.

We observe a strong bend in the retention distribution curve (Figure 2a) that is located between 400 ms to 1 s especially at $60 \,^{\circ}$ C). This conforms to the scaling trends predicted by [1] and observed by [8] for DDR3 DRAMs. This bend represents two failure probability distributions: one representing strong cells, and other one repesenting weak cells.

B. Current Consumption Measurement (IDDs)

Using our measurement platform, we measured various DRAM operational currrents at different temperatures for a DDR4 SO-DIMM from *Vendor-A*. All currents were measured at 1.2 GHz (DDR4-2400). There are two different voltage domains for DDR4 DRAMs: the VDD domain, which is 1.2 V, and the VPP domain, which is 2.5 V. The VPP domain supplies high voltage to Word Line Drivers during row activate



operation. Therefore, there are also different DRAM currents for both voltage domains: IDDs and IPPs.

Figure 4 shows a comparison of the measured IDD currents with the values given in the datasheet from *Vendor-A*. For majority of the currents, the datasheet values are much higher than the measured values. For example, for the Refresh current (IDD5B), the datasheet value is still 8% higher than measured value at 90°C. This is because vendors provide the currents measured for the worst case devices, which deviates largely from average case due to increased process variations.

Figure 5 shows a similar behavior for IPPs with respect to datasheet specifications. However, we observe for the IPP currents a decrease while increasing the temperature. As the DRAM vendors disclose no details of their technology, we assume this is due to frequency decrease of the oscillator circuitry driving the charge pumps.

V. CONCLUSION

In this paper, we presented a platform to analyze the retention behavior of DDR4 DRAMs at various temperatures, and to precisely measure various DRAM currents (IDDs and IPPs) at high frequencies (>1 GHz). Our experimental results reveal that retention behavior of DRAM is highly influenced by its internal architecture. Mixed-cell DRAMs (DRAMs with truecells and anti-cells) are less affected by the data stored in the neighboring cells, and therefore show less failure probability than true-cell only DRAMs when random data patterns are stored in the DRAM.

Our results applying ECC with random data pattern stored in the DRAM show that ECC is better suitable for Mixedcell DRAMs than true-cell only DRAMs, especially for retention times lower than 1 s. Hamming codes were able to reduce the retention failure probability by up to four orders of magnitude for Mixed-cell DRAMs. Therefore, we choose Mixed-cell DRAMs as the better candidate for applications in Approximate Computing/ADRAM.

Finally, we demonstrate the capability of our advanced measurement platform by conducting current measurements for different DDR4 command sequences defined by JEDEC. Knowledge of realistic DRAM currents will help in more accurate DRAM power estimations by various high level tools, which helps for a more optimistic power budget planning.

ACKNOWLEDGMENT

This work was supported by Carl-Zeiss Stiftung and by the German Research Foundation (DFG) as part of the priority program "Dependable Embedded Systems" (SPP 1500 - spp1500.itec.kit.edu) and by the DFG grants WE2442/10-1 and WE2442/9-3. Furthermore, this work was supported by the the Fraunhofer High Performance Center for Simulationand Software-based Innovation. The project OPRECOMP (http://oprecomp.eu) acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the European Union's Horizon 2020 research and innovation programme, under grant agreement No 732631.

REFERENCES

- K. Kim et al. A New Investigation of Data Retention Time in Truly Nanoscaled DRAMs. Electron Device Letters, IEEE, 30(8):846–848, Aug 2009.
 C. Weis, et al. Retention Time Measurements and Modelling of Bit
- [2] C. Weis, et al. Retention Time Measurements and Modelling of Bit Error Rates of WIDE I/O DRAM in MPSoCs. In Proceedings of the IEEE Conference on Design, Automation & Test in Europe (DATE). European Design and Automation Association, 2015.
- [3] M. Jung, et al. Efficient Reliability Management in SoCs An Approximate DRAM Perspective. In 21st Asia and South Pacific Design Automation Conference (ASP-DAC), 2016.
- [4] J. Liu, et al. RAIDR: Retention-Aware Intelligent DRAM Refresh. In Proceedings of the 39th Annual International Symposium on Computer Architecture, ISCA '12, pages 1–12, Washington, DC, USA, 2012. IEEE Computer Society.
- [5] J. Liu, et al. An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms. SIGARCH Comput. Archit. News, 41(3):60–71, June 2013.
- [6] M. Jung, et al. Omitting Refresh A Case Study for Commodity and Wide I/O DRAMs. In 1st International Symposium on Memory Systems (MEMSYS 2015), Washington, DC, USA, October 2015.
- [7] S. K. Park. Technology Scaling Challenge and Future Prospects of DRAM and NAND Flash Memory. In 2015 IEEE International Memory Workshop (IMW), 2015.
- [8] M. Jung, et al. A Platform to Analyze DDR3 DRAM's Power and Retention Time. IEEE Design & Test, 34(4):52–59, Aug 2017.
- [9] M. Patel, et al. The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions. In Proceedings of the 44th Annual International Symposium on Computer Architecture, 2017.
- [10] H. J. Kwon, et al. 23.4 An extremely low-standby-power 3.733Gb/s/pin 2Gb LPDDR4 SDRAM for wearable devices. In 2017 IEEE International Solid-State Circuits Conference (ISSCC), pages 394–395, Feb 2017.
- [11] C. K. Lee, et al. 23.2 A 5Gb/s/pin 8Gb LPDDR4X SDRAM with powerisolated LVSTL and split-die architecture with 2-die ZQ calibration scheme. In 2017 IEEE International Solid-State Circuits Conference (ISSCC), pages 390–391, Feb 2017.
- [12] S. Khan, et al. The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study. In The 2014 ACM International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '14, pages 519–532, New York, NY, USA, 2014. ACM.