

A Precise Measurement Platform for LPDDR4 Memories

Johannes Feldmann
j.feldmann@rptu.de
University of Kaiserslautern-Landau
Kaiserslautern, Germany

Lukas Steiner
lukas.steiner@rptu.de
University of Kaiserslautern-Landau
Kaiserslautern, Germany

Derek Christ
derek.christ@iese.fraunhofer.de
Fraunhofer IESE
Kaiserslautern, Germany

Thomas Psota
thomas.psota@iese.fraunhofer.de
Fraunhofer IESE
Kaiserslautern, Germany

Matthias Jung
m.jung@uni-wuerzburg.de
JMU Würzburg
Würzburg, Germany

Norbert Wehn
norbert.wehn@rptu.de
University of Kaiserslautern-Landau
Kaiserslautern, Germany

ABSTRACT

LPDDR4, the most widely used low-power DRAM standard in the industry, plays a crucial role in modern system design. Understanding the non-functional properties of these DRAMs such as power consumption and reliability is essential for accurate system design, as the values reported in vendor’s data sheets tend to be overly pessimistic compared to real-world performance. This research paper aims to characterize LPDDR4 memories of the three major vendors by conducting retention time and current measurements. To achieve these results, a custom measurement platform was developed, capable of precisely heating up the DRAM within a range of $\pm 1.0^\circ\text{C}$.

1 INTRODUCTION

Data-driven applications are increasingly becoming the focus of our information technology society. AI techniques disruptively change almost all areas of our society and economy. A common feature of all these applications is the enormous amount of data that needs to be captured, stored, and processed. As a result, external memory systems, particularly in System-on-Chip (SoC) or compute architectures, are gaining greater prominence. The state-of-the-art external memory systems are known as *Dynamic Random Access Memories* (DRAM), which come in various types (DDR_x, LPDDR_x, HBM_x, etc.). These different DRAM types vary significantly in their maximum bandwidths and latencies, storage capacity, reliability, and energy consumption. Of particular importance is energy consumption. Google recently demonstrated that more than 90 % of system energy is consumed by large machine learning models in memory. In *Augmented Reality Devices* for the Metaverse, according to Meta, memory can account for up to 80 % of energy consumption. Hence, selecting the appropriate DRAM memory based on applications and system context is of utmost importance. Moreover, in autonomous systems, AI-driven signal processing of numerous sensor data requires DRAM memory, where reliability takes precedence alongside power consumption. Therefore, it is crucial for system design to thoroughly characterize the reliability and power consumption of DRAM memories.

This paper presents an in-depth analysis of LPDDR4 DRAMs, which recently hold significant usage in the industry, including consumer products and safety-critical systems such as *Advanced Driving Assistant Systems* (ADAS) in automobiles. The primary focus of this analysis lies in assessing reliability, specifically retention errors, and power consumption. To accomplish this characterization,

we devised a customized measurement platform to meticulously examine the DRAMs from the three major vendors.

In summary, the paper makes the following new contributions:

- We present a precise retention error analysis that includes different temperature and data patterns, and compare the results of the measurements with the latest memory modules with previous measurements in literature.
- We present, to the best of our knowledge, for the first time, a current analysis of LPDDR4 DRAMs and compare the data with the vendor’s data sheets.
- We analyze the impact of the internal SEC-ECC on the reliability at different temperatures.
- From our results, we can conclude the internal array architectures of the different vendors.
- We present a sophisticated measurement platform that is capable to precisely heat up the DRAM and provide current measurements on all LPDDR4 voltage domains.

The paper is structured as follows: Section 2 presents the related work. The measurement platform is presented in 3, whereas Section 4 presents the experimental results. The paper is then finally concluded in Section 5.

2 RELATED WORK

There are several studies that analyze the retention behavior of DRAM chips. Kim and Lee presented in 2009 a detailed study on data retention times of nanoscaled DDR3 DRAMs [4]. Most of the following investigations, for instance [9, 10, 12] and [7], are based on these results. However, the authors of [3] showed with Wide I/O that even with high temperatures the majority of cells in this device can hold data much longer than 10,000s for a 0xFF data pattern. Later they confirmed in [2] for a similar DDR3 DRAM Device (same vendor) from 2009 that the presented numbers by [4] are way too pessimistic in the average case; DRAM cells can hold their value up to two orders of magnitude longer than assumed.

All previously presented platforms for measuring DRAM reliability and power consumption have been DIMM-based (e.g., [2] for DDR3 or [11] for DDR4) and employ Peltier elements to heat up the DRAM. Other studies have utilized a thermal chamber to house the entire DRAM measurement system (e.g., [13]). They present a study on 368 LPDDR4 devices. However, the authors do not disclose any details about their measurement platform beside that they used a thermal chamber. The authors of [8] analyze the data retention behavior with respect to temperature, DPD, and VRT. However,

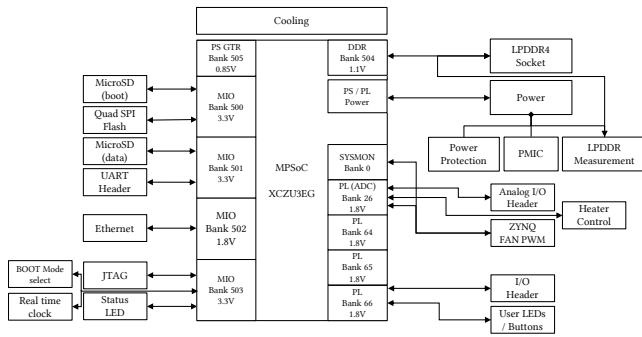


Figure 1: Block-Diagram of Measurement Platform

they measure the retention times only for several seconds (< 10 s) and use a thermal chamber.

Rahmati et al. [14] also use a thermal chamber, but moreover, they analyze a DRAM device from the 90’s. Therefore, their results have absolutely no relevance for nowadays DDR4/5 world. The authors of [1] measure only short retention times (< 3 s) and only one bank of one single DDR3 DRAM. Therefore, their results do not consider process variations. Moreover, they do not disclose details about their measurement setup, e.g. how they heat up the DRAM.

However, placing the entire measurement system within the thermal chamber may introduce undesired effects on the system itself, potentially influencing the results. Moreover, these platforms utilize the DIMM’s temperature sensor, which does not provide the actual chip temperature as it is not physically integrated into the DRAM chips. Therefore, it is important to directly heat up and measure the devices in order to obtain realistic results.

In [15] the authors measure the number of faulty pages with respect to ambient temperature and time. They heat up the DRAM chip manually by using a heat gun, which has the drawback that stabilizing the temperature over several hours is impossible.

3 MEASUREMENT PLATFORM

DIMM-based systems offer more convenient measurement, since off-the-shelf FPGA boards can be used. In such scenarios, currents can be measured using a specially designed adapter, as shown in [2, 11]. However, low-power DRAM systems like LPDDR4 do not utilize DIMMs; instead, the devices are directly soldered onto the *Printed Circuit Board* (PCB) of the computing system or even use *Package on Package* (PoP) technology. This presents a significant challenge, as it necessitates the design of an entire PCB, including FPGA and device socket, as a single adapter is insufficient for the measurement process. Therefore, we designed a custom measurement platform which allows a precise current and retention measurement while temperatures can be regulated with an accuracy of ± 1.0 °C. In the following we will describe the components of the platform, shown in Figure 1.

System on Chip (SoC). The measurement platform utilizes the Xilinx Zynq Ultrascale+ XCZU3EG-1SFVA625E, a *Multiprocessor System on Chip* (MPSoC), as its primary control unit. All essential tasks for conducting the measurements are executed on this platform. These tasks encompass heat regulation, generation of DRAM

data patterns, measurement of current and retention time, and the storage of acquired data.

Since the DRAM itself is the subject to test, it remains unavailable for use by the measurement software. The sole available main memory accessible to the software was a 256 kB on-chip SRAM memory. As a solution, we developed a microkernel tailored to operate on the four cores, optimized to fully reside within the platform’s on-chip memory. The purpose of the microkernel was to run diagnostics by filling the whole DRAM with certain specified patterns, while simultaneously controlling the output of the heating component. Each core was allocated for distinct functions within the system, encompassing tasks such as pattern writing and reading, heat control, and communication. Interaction with the platform was facilitated through the serial bus, enabling the exchange of JSON messages. These messages facilitated operations such as triggering new test runs, soliciting board diagnostics, and retrieving test results, all managed by external PC software written in Qt, as depicted in Figure 3.

DRAM Connection and Socket. In contrast to DDR4, which is typically linked to the target platform through a DIMM inserted into the appropriate socket, LPDDR4 is commonly soldered directly onto the PCB. This soldered configuration poses challenges for a measurement platform, as replacing the DRAM devices could lead to damage to both the device and the PCB. To address this, our measurement platform features an Ironwood SBT-BGA200 socket, facilitating swift interchangeability of LPDDR4 devices, including those of varying dimensions. This socket-based approach obviates the need for soldering of both the socket and the device itself. All test devices are connected to the hard IP memory controller of the Xilinx Zynq Ultrascale+. The measurement platform achieves a maximum data rate of 2133 Mb/s/pin.

Current Measurement Circuit. The measurement platform features four independent current measurement channels to measure all three supply currents of the LPDDR4 (VDD_1 , VDD_2 , VDD_Q) and the $VCCO_{PSDDR}$ voltage domain of the Zynq Ultrascale+. Each current is measured using a shunt resistor, which is connected in series to the device. All resistance values were carefully selected to maximize the voltage drop across the shunt resistor while guaranteeing sufficient supply voltage to the device. The chosen values are listed in Table 1. The voltage drop across each resistor is first filtered using a low-pass filter with a cut-off frequency of 1.6 MHz and amplified with a factor of 200 using high-precision current sense amplifiers (Analog Devices MAX44284W). The amplified voltages are synchronously sampled by a 24-bit analog-to-digital converter (ADC) of type MAX11040K as well as the corresponding supply voltages itself using a cascaded device. These ADCs are connected to the Xilinx Zynq Ultrascale+, which processes the data. The measurement board provides also the option to measure the amplified voltages using an oscilloscope.

I/O Interfaces. The measurement platform offers a variety of I/O interfaces. A 1000BASE-T Ethernet connection provides remote access to the platform. A separate UART is used to print debug outputs which offers the opportunity to monitor the measurement procedure in detail. The UART can be accessed via USB using an FTDI UART-to-USB IC. Two micro-sd card slots are available to

Voltage domain	Shunt resistance
VDD_1	750 m Ω
VDD_2	27 m Ω
VDD_Q	120 m Ω
$VCCO_{PSDDR}$	50 m Ω

Table 1: Shunt resistor of current measurement circuit




Figure 2: PCB Layout

store all measured data. Additionally, there are several buttons, switches, leds, and I/O header which offer user communication as well as extension capabilities.

Heating Element. The heating element, which is attached to the LPDDR4 socket, consists of a CNC manufactured copper block which contains a 12 V heating cartridge as well as a temperature sensor. To isolate the heating element from the surrounding air, a 3D-printed cover, called the “Dome”, consisting of air-filled chambers is put over the heating element and the socket containing the DRAM, c.f. Figure 4. The 20 W heating cartridge can be controlled by the SoC using pulse width modulation. The current temperature is measured using a MCP9700A sensor, which is connected to the internal ADC of the Xilinx Zynq Ultrascale+.

4 EXPERIMENTAL RESULTS

With the presented platform of Section 3 we conducted several measurements with respect to retention time and power consumption.

4.1 Retention Analysis

The first set of experiments is an analysis of the data pattern dependencies that have been reported in previous studies, e.g., [11]. Therefore, we conducted experiments with LPDDR4 using $0x00$,

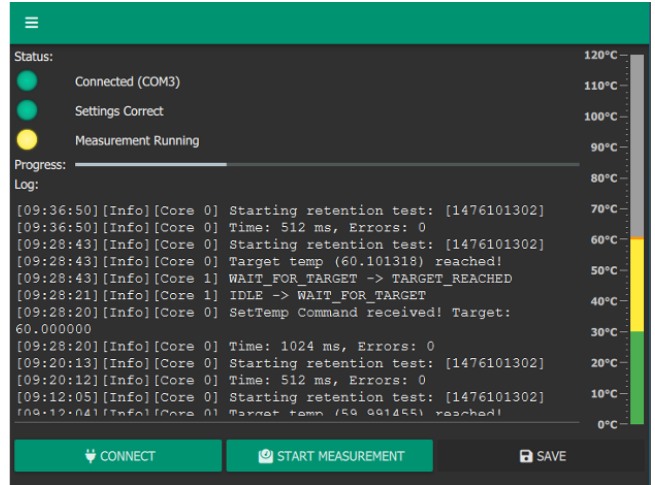


Figure 3: Control Software

$0x55$, $0xAA$, $0xFF$ and random data patterns. With these patterns it is possible to reverse engineer the internal DRAM architecture. Some DRAM vendors use only true-cells, other vendors are mixing true- and anti-cells in their DRAM architectures, where a true-cell stores the data value as it is and an anti-cell stores the inverse [8]. The reason for that is the architecture of the array and the primary sense amplifier. In order to analyze the data pattern dependencies, we analyzed one LPDDR4 device for each vendor. With these measurements it is also possible to rate the reliability, e.g., for *Approximate DRAM* scenarios.

For this purpose, the DRAM refresh is temporarily disabled, and the number of errors is measured within this time frame. It is important to highlight that all vendors meet the claimed reliability standards for typical refresh rates as required in the JEDEC standards (e.g., 64 ms).

For Vendor-A, we observe errors for all data patterns, as shown in Figure 5, indicating the usage of a mixed true- and anti-cell DRAM. In contrast, Vendor-B (Figure 6) and Vendor-C (Figure 7) do not exhibit errors for the $0x00$ data pattern, implying a true-cell DRAM. The data pattern dependency for Vendor-A is minimal, while the other two vendors show a much higher data pattern dependency. Notably, for Vendor-B and Vendor-C, the majority of errors occur with the $0xFF$ pattern, which aligns with the true-cell architecture. Conversely, Vendor-A experiences most errors with random data patterns.

To estimate the variance across multiple devices, we conducted measurements on five different devices from each vendor using a random pattern. The results are illustrated in Figure 8, where we present the maximum, minimum, and average errors. Among the vendors, Vendor-A exhibits the highest variance, while Vendor-C shows a medium variance, and Vendor-B has the lowest variance. A direct comparison between the vendors is depicted in Figure 9, whereas Figure 9 shows the same results in a linear form instead of logarithmic. Notably, at elevated temperatures, the performance disparity becomes evident, particularly with Vendor-A being considerably less reliable compared to Vendor-B and Vendor-C. Nonetheless, as mentioned before, all vendors meet the claimed reliability




Figure 4: Developed Measurement Platform

standards when the DRAM is used within the specification. Table 2 summarizes the qualitatively the measured results.

LPDDR4 is the first JEDEC standard that allows the integration of in-DRAM ECC into devices. For the tested devices the used ECC is a *Single Error Correction* (SEC) (136,128) shortened Hamming code [5, 6]. We analyzed the performance of this ECC for Vendor-A and compared it with previous DDR4 measurements for the same vendor from [11], as shown in Figure 11.

It is important to note that DDR4 uses an external ECC engine placed inside the memory controller, which covers all errors from the DRAM cell up to the memory controller, whereas LPDDR4’s ECC engine can only correct errors that appear within the memory array. Furthermore, it might be that for both devices different DRAM technologies are used. However, from an information-theoretical perspective, the (136,128) shortened Hamming code used in LPDDR4 devices has a higher code rate than the (72,64) shortened Hamming code of DDR4’s external ECC engine. Thus, its error correction capability is lower. This general trend can also be observed in our measurements.

Furthermore, we compared our results with prior art [13], as shown in Figure 12. They measured at 45 °C in a thermal chamber. Their values are closer to our 60 °C results than to our 30 °C. This might be because of the usage of a thermal chamber, which also extensively heats up the DRAM interface on the FPGA side, or they have used an earlier generation of LPDDR4 devices. Unfortunately, they do not disclose their measurement setup and the DRAMs they have analyzed.

4.2 Current Measurements

Using the measurement circuit presented in Section 3, we conducted current measurements for all three voltage domains, namely VDD_1 , VDD_2 , and VDD_Q . Figure 13 illustrates the measured currents, and the total current I_{total} is calculated as the sum of I_1 , I_2 , and I_Q . However, it was not possible to accurately measure the write current $IDD4W^*$, as the FPGA platform cannot send two write requests consecutively. Similarly, measuring the ACT-PRE current $IDD0^*$ in a JEDEC-compliant manner was not feasible, as it was not possible to execute ACT and PRE commands consecutively due to the absence of this feature in hard-IP memory controller of the FPGA. Instead, $IDD0^*$ reflects the sequence ACT-RD-PRE. Therefore, these two currents only provide approximate values, but they are still valid for comparing different vendors. For Vendor-C, no datasheet was provided. Notably, we observed that the values in the datasheet are overly pessimistic (especially for Vendor-B) compared to the measured results. These measurements can be invaluable to system designers in optimizing the power planning and ensuring a more optimistic design.

5 CONCLUSION

This paper focuses on the analysis of LPDDR4 DRAMs, which are widely used in various applications, including consumer products and safety-critical systems like *Advanced Driving Assistant Systems* (ADAS) in cars. The study centers around assessing the reliability and power consumption of these DRAMs. To achieve this, we designed a customized measurement platform capable of

	Type	DPD	Variance	SEC	Most Errors
A	Mixed-Cell	very low	medium	yes	random
B	True-Cell	low	low	yes	0xFF
C	True-Cell	low	very low	yes	0xFF

Table 2: Qualitative Comparison of Vendors

accurately characterizing the DRAMs from the three major vendors. The paper makes several notable contributions, including a precise retention error analysis with different temperature and data patterns, a current analysis of LPDDR4 DRAMs, an evaluation of internal SEC-ECC impact on reliability at different temperatures, and insights into the internal array architectures of the vendors. We demonstrated that the datasheet values provided by the vendors were overly pessimistic. The study’s results offer valuable information to system designers for optimizing power planning and layout design.

ACKNOWLEDGEMENTS

We thank Martin Schultheis for the support. This paper was funded by the DFG and Fraunhofer within the MEMTOMOY project (NUMMER) and by the BMBF within the MEMTONOMY-2 project (GRANT Number).

REFERENCES

- [1] Chih-Sheng Hou, Jin-Fu Li, Chih-Yen Lo, Ding-Ming Kwai, Yung-Fa Chou, and Cheng-Wen Wu. 2013. An FPGA-based test platform for analyzing data retention time distribution of DRAMs. In *VLSI Design, Automation, and Test (VLSI-DAT), 2013 International Symposium on*. 1–4. <https://doi.org/10.1109/VLDI-DAT.2013.6533853>
- [2] Matthias Jung, Deepak Mathew, Carl Rheinländer, Christian Weis, and Norbert Wehn. 2017. A Platform to Analyze DDR3 DRAM’s Power and Retention Time. *IEEE Design & Test* 34, 4 (Aug 2017), 52–59. <https://doi.org/10.1109/MDAT.2017.2705144>
- [3] Matthias Jung, Deepak M. Mathew, Christian Weis, and Norbert Wehn. 2016. Efficient Reliability Management in SoCs - An Approximate DRAM Perspective. In *21st Asia and South Pacific Design Automation Conference (ASP-DAC)*.
- [4] Kinam Kim and Jooyoung Lee. 2009. A New Investigation of Data Retention Time in Truly Nanoscaled DRAMs. *Electron Device Letters, IEEE* 30, 8 (Aug 2009), 846–848. <https://doi.org/10.1109/LED.2009.2023248>
- [5] H. J. Kwon, E. Seo, C. Y. Lee, Y. H. Seo, G. H. Han, H. R. Kim, J. H. Lee, M. S. Jang, S. G. Do, S. H. Cho, J. K. Park, S. Y. Doo, J. B. Shin, S. H. Jung, H. J. Kim, I. H. Im, B. R. Cho, J. W. Lee, J. Y. Lee, K. H. Yu, H. K. Kim, C. H. Jeon, H. S. Park, S. S. Kim, S. H. Lee, J. W. Park, S. S. Lee, B. T. Lim, J. y. Park, Y. S. Park, H. J. Kwon, S. J. Bae, J. H. Choi, K. I. Park, S. J. Jang, and G. Y. Jin. 2017. 23.4 An extremely low-standby-power 3.73Gb/s/pin 2Gb LPDDR4 SDRAM for wearable devices. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. 394–395. <https://doi.org/10.1109/ISSCC.2017.7870427>
- [6] C. K. Lee, Y. J. Eom, J. H. Park, J. Lee, H. R. Kim, K. Kim, Y. Choi, H. J. Chang, J. Kim, J. M. Bang, S. Shin, H. Park, S. Park, Y. R. Choi, H. Lee, K. H. Jeon, J. Y. Lee, H. J. Ahn, K. H. Kim, J. S. Kim, S. Chang, H. R. Hwang, D. Kim, Y. H. Yoon, S. H. Hyun, J. Y. Park, Y. G. Song, Y. S. Park, H. J. Kwon, S. J. Bae, T. Y. Oh, I. D. Song, Y. C. Bae, J. H. Choi, K. I. Park, S. J. Jang, and G. Y. Jin. 2017. 23.2 A 5Gb/s/pin 8Gb LPDDR4X SDRAM with power-isolated LVSTL and split-die architecture with 2-die ZQ calibration scheme. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. 390–391. <https://doi.org/10.1109/ISSCC.2017.7870425>
- [7] Chung-Hsiang Lin, De-Yu Shen, Yi-Jung Chen, Chia-Lin Yang, and M. Wang. 2012. SECRET: Selective error correction for refresh energy reduction in DRAMs. In *IEEE 30th International Conference on Computer Design (ICCD)*. 67–74. <https://doi.org/10.1109/ICCD.2012.6378619>
- [8] Jamie Liu, Ben Jaiyen, Yoongu Kim, Chris Wilkerson, and Onur Mutlu. 2013. An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms. *SIGARCH Comput. Archit. News* 41, 3 (June 2013), 60–71. <https://doi.org/10.1145/2508148.2485928>
- [9] Jamie Liu, Ben Jaiyen, Richard Veras, and Onur Mutlu. 2012. RAIDR: Retention-Aware Intelligent DRAM Refresh. In *Proceedings of the 39th Annual International Symposium on Computer Architecture (Portland, Oregon) (ISCA ’12)*. IEEE Computer Society, Washington, DC, USA, 1–12. <http://dl.acm.org/citation.cfm?id=2337159.2337161>
- [10] Jan Lucas, Mauricio Alvarez-Mesa, Michael Andersch, and Ben Juurlink. 2014. Sparkk: Quality-Scalable Approximate Storage in DRAM. In *The Memory Forum*. <http://www.redaktion.tu-berlin.de/fileadmin/fg196/publication/sparkk2014.pdf>
- [11] Deepak M. Mathew, Martin Schultheis, Carl C. Rheinländer, Chirag Sudarshan, Matthias Jung, Christian Weis, and Norbert Wehn. 2018. An Analysis on Retention Error Behavior and Power Consumption of Recent DDR4 DRAMs. In *IEEE Conference Design, Automation and Test in Europe (DATE)*.
- [12] Prashant J. Nair, Dae-Hyun Kim, and Moinuddin K. Qureshi. 2013. ArchShield: Architectural Framework for Assisting DRAM Scaling by Tolerating High Error Rates. In *Proceedings of the 40th Annual International Symposium on Computer Architecture (Tel-Aviv, Israel) (ISCA ’13)*. ACM, New York, NY, USA, 72–83. <https://doi.org/10.1145/2485922.2485929>
- [13] Minesh Patel, Jeremie S. Kim, and Onur Mutlu. 2017. The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions. *SIGARCH Comput. Archit. News* 45, 2 (jun 2017), 255–268. <https://doi.org/10.1145/3140659.3080242>
- [14] Amir Rahmati, Matthew Hicks, Daniel Holcomb, and Kevin Fu. 2014. Refreshing thoughts on DRAM: Power saving vs. data integrity. In *Workshop on Approximate Computing Across the System Stack (WACAS)*.
- [15] R.K. Venkatesan, S. Herr, and E. Rotenberg. 2006. Retention-aware placement in DRAM (RAPID): software methods for quasi-non-volatile DRAM. In *The Twelfth International Symposium on High-Performance Computer Architecture, 2006*. 155–165. <https://doi.org/10.1109/HPCA.2006.1598122>




Figure 5: Data Pattern Analysis Vendor-A




Figure 6: Data Pattern Analysis Vendor-B




Figure 7: Data Pattern Analysis Vendor-C




Figure 8: Spread of Retention Errors for 5 Devices for each Vendor with Random Data Pattern




Figure 9: Comparison of Different Vendors with Random Data Pattern (log)




Figure 10: Comparison of Different Vendors with Random Data Pattern (linear)




Figure 11: Comparison of DDR4 SECEDED with LPDDR4 SEC




Figure 12: Comparison with Results of [13]




Figure 13: Currents for Different Vendors