Stuck Errors in Bits and Blocks in GDDR6 Under High-energy Neutron Irradiation

Masakazu Yoshida Riku Iwamoto Masatoshi Itoh Masanori Hashimoto

2023 RADECS Data Workshop | 979-8-3503-8470-3/23/$31.00 ©2023 IEEE | DOI: 10.1109/RADECS59798.2023.10752846

***Abstract*—When DRAMs are irradiated with neutrons or protons, upsets occur in memory cells and logic circuits, resulting in single-bit errors and block errors, respectively. Among single-bit errors, those that occur repeatedly are called stuck bits and have been reported. We performed a neutron irradiation experiment on GDDR6 DRAM. Experimental results show that block errors occurred repeatedly at the same addresses in irradiated DRAMs named stuck block errors. We newly observed such stuck block errors in this work, while temporary block errors are reported in the literature. We also found that the stuck block errors disappeared after power cycling. Therefore, we consider that a possible reason is micro single event latch-up.**

***Index Terms*—Dynamic random access memory (DRAM), neutron irradiation, stuck bit error, stuck block error.**

I. Introduction

DRAMS are widely used in computing devices to store large data. For example, PCs use DRAM as the main memory and GPUs store image data in DRAM. In cloud servers and supercomputers, the reliability of DRAM has been a concern since the total amount of memory is tremendous. Also, in recent years, safety-critical applications, such as autonomous driving, are executed on GPUs. Therefore, DRAM reliability needs to be investigated every generation. For example, irradiation tests have been conducted on current DRAM standards, such as DDR3 [1], DDR4 [2], GDDR5 [3], and HBM2 [4].

As technology advances, the cell footprint and the size of transistors in DRAMs are decreasing. The soft error rate (SER) per bit tends to become lower according to technology scaling because the probability of a neutron having a nuclear reaction with silicon in a specific bit cell is smaller [5]. However, the error rate per chip has not changed much since total storage capacity has increased. Meanwhile, irradiating DRAM with neutrons or protons results in stuck bits that cause upsets repeatedly [1]. The increased leakage current due to displacement damage or total ionizing dose cause persistent errors in the memory cell. Furthermore, the recent literature (e.g., [2], [3]) points out that block errors (also known as cluster errors) occur in recent DRAMs, which originate from errors in peripheral logic circuits.

In this study, we performed a neutron irradiation experiment on GDDR6 memory to investigate stuck and temporary errors. In addition, we investigate the temporal variation of stuck bit errors. Also, the occurrence pattern of block errors is analyzed in time and space.

The rest of this paper is organized as follows. Section II reviews related work regarding neutron and proton irradiation to DRAM. Section III describes the experimental setup of neutron irradiation. Section IV shows the result of the experiment. Finally, Section V concludes the discussion.

This work is supported by JST CREST Grant Number JPMJCR19K5, and the Grant-in-Aid for Scientific Research (S) from Japan Society for the Promotion of Science (JSPS) under Grant JP19H05664.

Masakazu Yoshida and Masanori Hashimoto are with the Department of Informatics, Kyoto University.

Riku Iwamoto was with the Department of Information Systems Engineering, Osaka University.

Masatoshi Itoh is with Cyclotron and Radioisotope Center, Tohoku University.

II. Related Work

Soft error rate of DRAM is an important factor determining the reliability of computer systems, especially large cloud servers and supercomputers. Therefore, various studies have been conducted, including neutron and proton irradiation experiments. The following paragraphs introduce some of those works.

The literature describes that there are two main types of errors: transient errors and persistent errors, the former being single event upset (SEU) due to the particle-induced temporal noise and the latter being stuck bits due to increased leakage current. When a leakage current flows from a charged capacitor, the stored charge gradually decreases, and eventually the value stored in the memory is flipped. DRAM inherently suffers from the leakage current, and hence refresh operation is repeated performed to prevent such bit flipping. However, an excessive increase in leakage current causes bit flipping within the refresh interval, resulting in stuck bit errors.

C. Lim et al. reported that errors increase with longer retention times in DDR3, and stuck bits were recovered completely after annealing at 150 °C [1]. A. Rodriguez et al. reported that bit flips could occur once DRAM cells are irradiated, even when the beam is off. They also reported the relationship between retention time and stuck bits ratio [6]. M. Park et al. analyzed groups of flipped bits in DDR4. They reported block errors with the same row address and column addresses due to failure of logic functions [2]. Our previous work on GDDR5 reported that stuck bit errors could remain for months and that some block errors shift to different addresses in time [3]. M. B. Sullivan et al. investigated the number of simultaneous upset bits in a word and their occurrence frequency in HBM2. They also proposed two ECC schemes robust to multi bits errors [4].

In addition to the bit flipping mentioned earlier, the phenomenon of single event latch-up (SEL), which is an abnormal current state in a circuit and continues until the device is power cycled [7], has been investigated. In SRAMs, J. Tausch et al. reported experimental details of micro SEL occurrences and reported that the devices operating at voltages more than

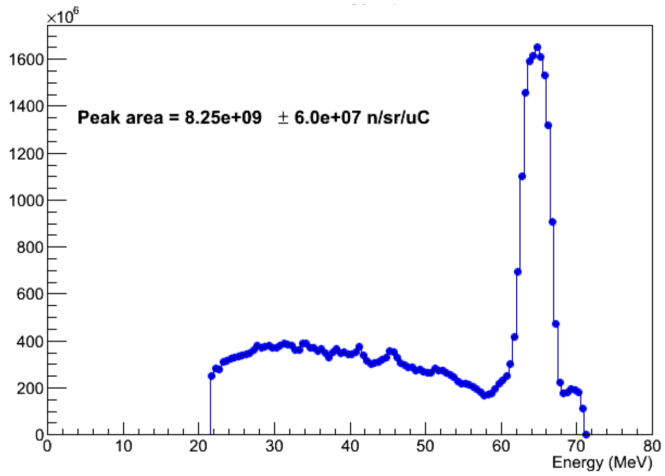

Fig. 1. A typical energy spectrum of the neutron beam at CYRIC [12].

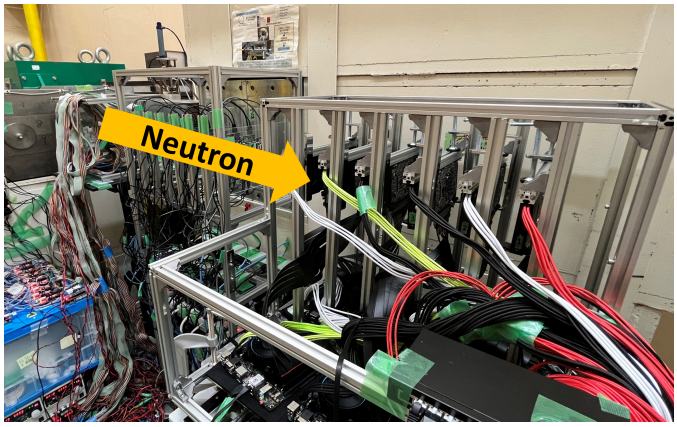

Fig. 2. Experiment setup for high-energy neutron irradiation to GPU at CYRIC.

1.1 V [8]. As for DRAMs, D. Kline et al. reported byte-level errors on DDR3 and hypothesized that each byte of every word belongs to the same chip, which experiences latch-up [9].

Thus, as described above, the reliability of GDDR6 under irradiation has not been reported, as far as the authors know. Also, the reported block errors on DRAMs are temporary; they disappear after a while.

III. Experimental Setup

We performed a quasi-monoenergetic neutron irradiation experiment at the Cyclotron and Radioisotope Center (CYRIC) at Tohoku University. Fig. 1 shows the spectrum of its neutron beam generated by 70-MeV protons. It has a flux peak at the energy near 70 MeV. The DUT is the 16 GiB GDDR6 onboard memory mounted on the NVIDIA RTX A4000 GPU card [10]. The RTX A4000 GPU has 48 streaming multiprocessors (SMs), and 6144 CUDA cores. The page size of 16 GiB GDDR6 memory can be either 2 KiB or 4 KiB, and the operating voltage can be either 1.25 V or 1.35 V [11]. As shown in Fig. 2, six GPUs were aligned with the beam, and one of them was used in this study. A 60-cm extension cable extends the PCIe between the GPU card and the host motherboard. Then, the motherboard is located off the beam track. A LAN was installed to control the experiment from outside the irradiation room. During irradiation, the GPU card and host PC can be hung or unresponsive. In such a case, the host was rebooted forcibly by a LAN-controlled rebooter, and the measurement was resumed. The fluence of neutrons is about $6.8 \times 10^{10}\,\mathrm{n/cm^2}$.

During the beam irradiation, we repeatedly executed a CUDA program that repeats writes to and reads from DRAM 230 times with two write patterns of ALL0 and ALL1 alternately selected. All SMs are used to write and read in parallel, and if the written and read values do not match, the value and address are recorded. Since the only address visible to the programmer is a 47-bit virtual address, and the correspondence between the physical and virtual addresses is unknown, the relative address starting from the top address of the allocated memory area was recorded. Within 230 writes and reads in one CUDA program execution, the allocated physical memory is fixed, whereas it cannot be confirmed whether the allocated physical memory is the same in different runs of the error-checking CUDA program. The size of allocated memory was less than 16 GiB because not all memory space can be allocated to this CUDA program. In this experiment, it was 16,420,207,616 B for the first 2 hours and 8,010,103,808 B for the remaining time.

The observed errors are first classified into two categories. Errors occurring at consecutive addresses are classified as block errors, and errors occurring at isolated addresses are classified as single-bit errors. Orthogonal to this, errors occurring at the same address consecutively recorded in each error check, which consists of 230 times writing and reading, are classified as stuck errors and other errors as temporal errors. In summary, there are stuck block errors, stuck bit errors, temporary block errors, and temporary bit errors.

IV. Measurement Results

A. Single-bit errors

Regarding single-bit errors, none of the temporary single-bit errors were observed in ALL0 patterns, suggesting that the charged state corresponds to 1 in all DRAM cells of the DUT. This is similar to some of the literature (e.g., [1]). Thus, the rest of this section describes the results of the ALL1 pattern.

Fig. 3 shows the error rates of the stuck bits and the temporary bits as the function of time. Here, we need to mention that the GPU card had been used for irradiation experiments before this experiment. Therefore, the GPU card had already been affected at time 0. The neutron flux at each time point is also plotted using the right vertical axis. The duration in which the flux is zero means that the neutron beam was stopped. It can be seen that stuck bits increase when the neutron beam is irradiated and decrease when the neutron beam is stopped. For example, the time duration from 26 to 29 hours, the beam was mostly stopped. Before the beam stop, the stuck bit error rate is around 2.5×10^{-12}, while the error rate drops to around 1.5×10^{-12} during the beam stop. This result suggests that there is a recovery mechanism in each memory cell whose time constant is longer than an hour for stuck bit

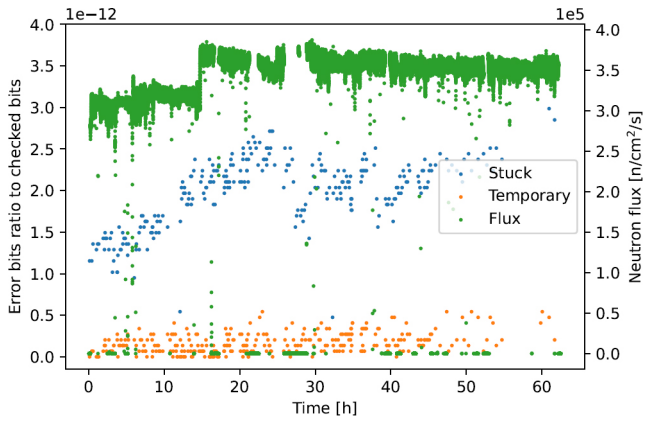

Fig. 3. Stuck and temporary single bit error ratio and neutron flux. The left vertical axis shows the ratio of bits in which each error occurred to all the checked bits. The right vertical axis shows the neutron flux.

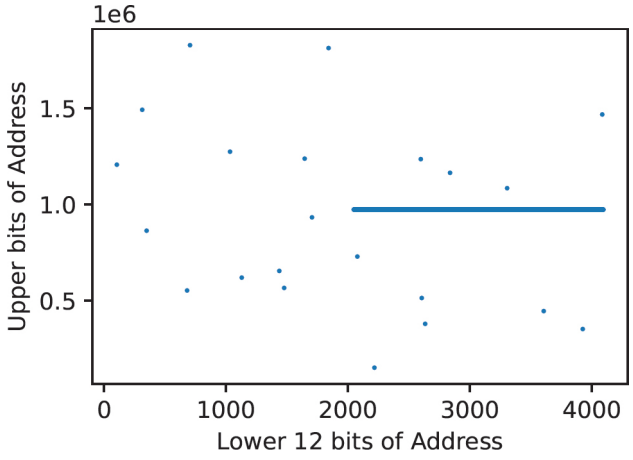

Fig. 4. Block errors. This shows the result of one run of error checking CUDA program.

errors. Meanwhile, the stuck bit error rate seems saturated in the latter half of the irradiation. A possible explanation might be the generation rate and recovery rate of stuck bits being balanced.

On the other hand, temporary single-bit errors continued to be observed even when the beam was stopped. Therefore, the temporary single-bit errors cannot be regarded as SEUs, and the true SEUs could be fewer than temporary single-bit errors. Meanwhile, the cross section calculated for temporary single-bit errors is $5.72\times10^{-11}\,\mathrm{cm}^2/\mathrm{Gb}$. According to [2], the cross section is about $7\times10^{-11}\,\mathrm{cm}^2/\mathrm{Gb}$ for DDR3 irradiated with atmospheric-like neutrons and $1\times10^{-10}\,\mathrm{cm}^2/\mathrm{Gb}$ for DDR4 irradiated with 480 MeV protons. The cross section of temporary single-bit errors in this study is smaller than these values.

B. Block errors

Fig. 4 shows the addresses of stuck bits. Each dot indicates an address where a stuck error is observed. The horizontal line indicates a stuck block error. Note that this figure is generated using the result of a single run of the error-checking CUDA program.

Fig. 5 shows all data and the frequency of stuck error occurrence at each address. Here, the frequency of 15, for example, means that the stuck errors are observed in 15 runs of the error-checking CUDA programs. It can be seen that errors have occurred at the same address several times, block errors have occurred at multiple address ranges, and the stuck block errors in red and green occurred more than ten times at the same address range.

Finally, Fig. 6 shows the temporal variation of the addresses where block errors occurred. We can see that once a block error occurs, it may continue for a while. We confirmed that the host PC was not rebooted within each time range in Fig. 6, and at the end of each time range, the host was hung, and the PC was rebooted. A possible reason is that DRAM cells in line or their sensing circuit are affected by micro SEL. Another reason could be register values in peripheral circuits being flipped.

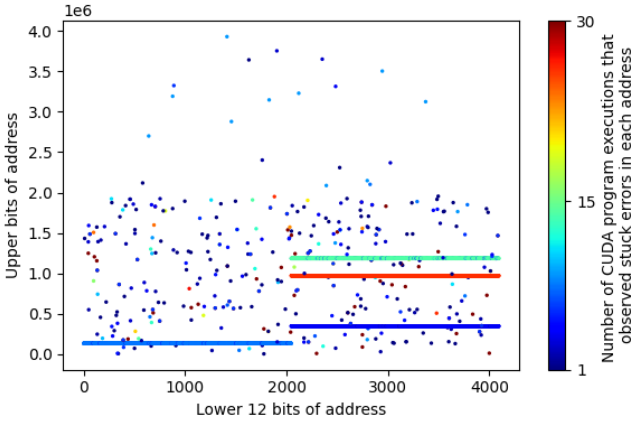

Fig. 5. The number of stuck block errors observed at each address during multiple runs of the error-checking program. Red and green lines indicate that stuck block errors have often occurred at the same address ranges. The errors that occurred more than 30 times are indicated by the same color as the one indicating 30 times.

V. Conclusion

We investigated the effect of high-energy neutron irradiation on GDDR6 DRAM. Similar to the results of previous studies, stuck bit errors, temporary block errors, and temporary bit errors were observed. In addition to this, we found stuck block errors occurred, and they disappeared after power cycling. A possible reason could be micro single event latch-up.

References

[1] C. Lim, H. S. Jeong, G. Bak, S. Baeg, S.-J. Wen, and R. Wong, "Stuck bits study in DDR3 SDRAMs using 45-MeV proton beam," *IEEE transactions on nuclear science*, vol. 62, no. 2, pp. 520–526, Apr. 2015.

[2] M. Park, S. Jeon, G. Bak, C. Lim, S. Baeg, S. Wen, R. Wong, and N. Yu, "Soft error study on DDR4 SDRAMs using a 480 MeV proton beam," in *2017 IEEE International Reliability Physics Symposium (IRPS)*, Apr. 2017, pp. SE–3.1–SE–3.6.

[3] M. Hashimoto, Y. Zhang, and K. Ito, "Neutron-induced stuck error bits and their recovery in DRAMs on GPU cards," in *Proceedings of International Conference on Solid State Devices and Materials (SSDM)*, 2022.

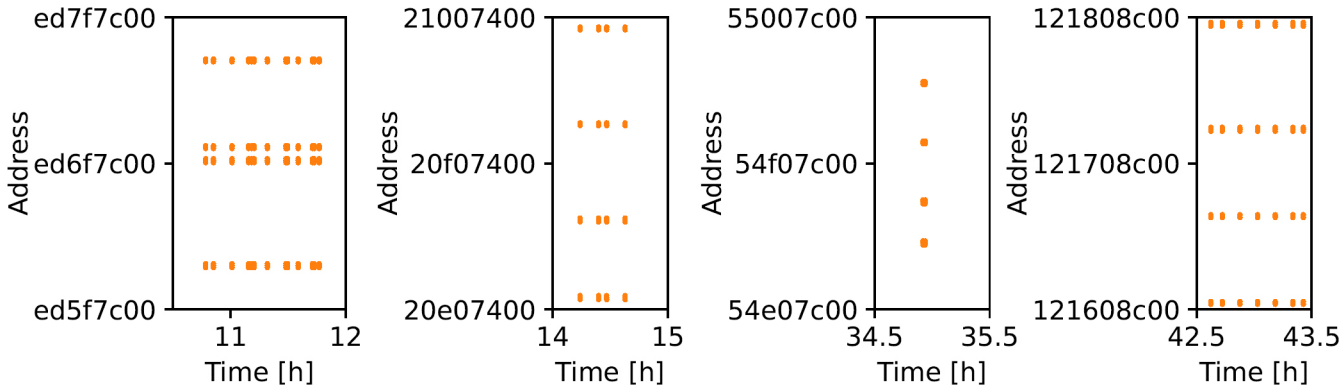

Fig. 6. Temporal and spatial distribution of stuck block errors.

[4] M. B. Sullivan, N. R. Saxena, M. O'Connor, D. Lee, P. Racunas, S. Hukerikar, T. Tsai, S. K. S. Hari, and S. W. Keckler, "Characterizing and mitigating soft errors in GPU DRAM," *IEEE Micro*, vol. 42, no. 4, pp. 69–77, Jul. 2022.

[5] R. C. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies," *IEEE Transactions on Device and Materials Reliability*, vol. 5, no. 3, pp. 305–316, Sep. 2005.

[6] A. Rodriguez, F. Wrobel, A. Samaras, F. Bezerra, B. Vandevelde, R. Ecoffet, A. Touboul, N. Chatry, L. Dilillo, and F. Saigné, "Proton-Induced Single-Event degradation in SDRAMs," *IEEE transactions on nuclear science*, vol. 63, no. 4, pp. 2115–2121, Aug. 2016.

[7] Jedec, "Measurement and reporting of alpha particles and terrestrial cosmic ray-induced soft errors in semiconductor devices : JESD89B," *JEDEC STANDARD, JEDEC Sold State Technology Association*, vol. 89, pp. 1–85, 2021.

[8] J. Tausch, D. Sleeter, D. Radaelli, and H. Puchner, "Neutron induced micro SEL events in COTS SRAM devices," in *2007 IEEE Radiation Effects Data Workshop*, 2007, pp. 185–188.

[9] D. Kline, S. Longofono, R. Melhem, and A. K. Jones, "Predicting single event effects in DRAM," in *2019 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, 2019, pp. 1–6.

[10] NVIDIA Corporation, "NVIDIA RTX A4000 graphics card," https://www.nvidia.com/en-us/design-visualization/rtx-a4000/, accessed: 2023-4-5.

[11] JEDEC Solid State Technology Association, "GRAPHICS DOUBLE DATA RATE (GDDR6) SGRAM STANDARD: JESD250C," 2021.

[12] Y. Sakemi, M. Itoh, and T. Wakui, "High intensity fast neutron beam facility at CYRIC," Tech. Rep. IAEA-TECDOC–1743(COMPANION CD-ROM), 2014.