Impact of Data Patterns and Interface Data Rates on DRAM Retention Reliability

Jarosław Warmbier, Paweł Pawłowski, Magdalena Szymkowiak
Institute of Automatic Control and Robotics
Poznan University of Technology
Poznań, Poland

EXTENDED ABSTRACT

Dynamic random-access memory (DRAM) remains the dominant technology for main memory in contemporary computing systems because it combines high density, high bandwidth, and low cost. At the same time, continued scaling of DRAM cells has reduced the physical area and stored charge of each bit cell, making data retention increasingly important for reliability, energy efficiency, and data-security-oriented applications. This paper investigates how scaling and increasing interface data rates influence the ability of modern DDR5 and LPDDR5 memories to retain stored information when refresh operations are disabled.

The work summarizes the architectural evolution of synchronous DRAM, with an emphasis on the transition toward higher bank counts, longer burst lengths, lower supply voltages, and substantially increased transfer rates in the DDR and LPDDR families. The fundamental retention problem is then related to the one-transistor, one-capacitor DRAM cell, in which data is stored as charge in the capacitor and must be periodically restored because of a leakage. Although JEDEC-compliant operation requires refreshing the whole memory every 32 ms for both DDR5 and LPDDR5, the actual retention margin of physical devices can be much longer and depends on temperature, data pattern, and neighboring-cell interactions.

Experiments were conducted on two x86-based platforms: one with a package-on-package LPDDR5 memory and another with a SO-DIMM DDR5 module. Both systems were operated under controlled thermal conditions at 85 °C, corresponding to a worst-case normal operating temperature. In the first experiment, memory regions were initialized with selected data patterns, memory refresh cycling was disabled for controlled time intervals, and data were subsequently read back after refresh was re-enabled. The tested patterns included all ones, all zeros, alternating bit patterns, a single-one-per-nibble pattern, and a pseudo-random LFSR pattern. In the second experiment, the mean time to first failure (MTFF) was measured for various operating data rates using repeated trials.

The results show a strong dependence of retention behavior on the written data pattern (Fig. 1). All-zero data remained stable in the evaluated time window, while all-one data produced the highest failure rate, indicating that logic zero acts as a default cell state in the tested devices. Alternating and pseudo-random patterns did not follow a simple independent-bit probability model, suggesting that retention failures are affected by neighboring-cell states. The cumulative failure distributions show that, for the all-one pattern, 50 % of failures occur earlier in LPDDR5 than in DDR5. The MTFF analysis further shows that DDR5 retains data longer than LPDDR5 in the evaluated platforms, while increasing data rate is associated with a negative trend in retention time for both technologies (Fig. 2).

The results show a negative trend in MTFF with increasing data rate, highlighting that further increases in interface bandwidth may come at the cost of reduced retention reliability. Furthermore, the results indicate that the ability to store data reliably strongly depends on the data pattern being stored.

The study also demonstrates that the industry-standard refresh interval of 32 ms may be overly conservative in light of the observed data retention times for both technologies.

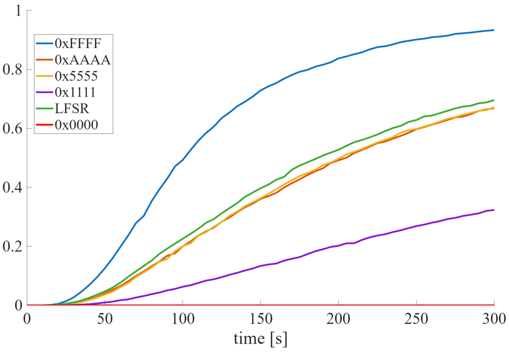

Fig. 1. Empirical cumulative relative number of failures for various data patterns in DDR5 memory

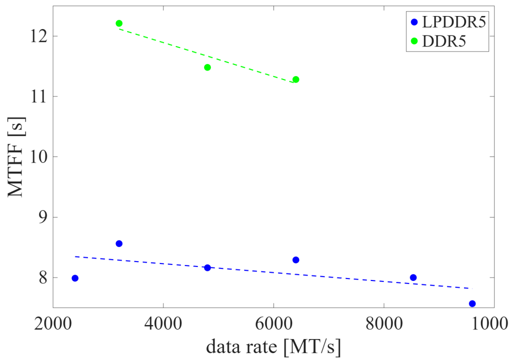

Fig. 2. Mean Time to First Failure (MTFF) for DDR5 and LPDDR5