Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability

Xiao Liu[1] David Roberts Rachata Ausavarungnirun[2] Onur Mutlu[3] Jishen Zhao[1]
[1]UC San Diego [2]King Mongkut's University of Technology North Bangkok [3]ETH Zürich

ABSTRACT

As memory capacity scales, traditional cache and memory hierarchy designs are facing increasingly difficult challenges in ensuring high reliability with low storage and performance cost. Recent developments in 3D die-stacked DRAM caches and nonvolatile memories (NVRAMs) introduce promising opportunities in tackling the reliability, performance, and capacity challenges, due to the diverse reliability characteristics of the technologies. However, simply replacing DRAM with NVRAM does not solve the reliability issues of the memory system, as conventional memory system designs maintain separate reliability schemes across caches and main memory. Our goal in this paper is to enable a reliable and high-performance memory hierarchy design, as memory capacity scales. To this end, we propose Binary Star, which coordinates the reliability schemes and consistent cache writeback between 3D-stacked DRAM last-level cache and NVRAM main memory to maintain the reliability of the cache and the memory hierarchy. Binary Star significantly reduces the performance and storage overhead of consistent cache writeback by coordinating it with NVRAM wear leveling. As a result, Binary Star is much more reliable and offers better performance than state-of-the-art memory systems with error correction. On a set of memory-intensive workloads, we show that Binary Star reduces memory failures in time (FIT) by 92.9% compared to state-of-the-art error correction schemes, while retaining 99% of the performance of a conventional DRAM design that provides no error correction.

CCS CONCEPTS

• **Computer systems organization → Reliability**; **Processors and memory architectures**; • **Hardware → Emerging technologies**.

KEYWORDS

reliability, nonvolatile memory, hybrid memory, memory systems

ACM Reference Format:

Liu, et al.. 2019. Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability. In *Proceedings of The 52nd Annual IEEE/ACM International Symposium on Microarchitecture, Columbus, OH, USA, October 12–16, 2019 (MICRO-52),* 14 pages.

MICRO-52, October 12–16, 2019, Columbus, OH, USA

ACM ISBN 978-1-4503-6938-1/19/10…$15.00
https://doi.org/10.1145/3352460.3358262

1 INTRODUCTION

Modern cloud servers adopt increasingly larger main memory and last-level caches (LLC) to accommodate the large working set of various in-memory computing [83, 85], big-data analytics [4, 73], deep learning [82], and server virtualization [4] applications. The memory capacity demand requires further process technology scaling of traditional DRAMs, e.g., to sub-20nm technology nodes [8, 53]. Yet, such aggressive technology scaling imposes substantial challenges in maintaining reliable and high-performance memory system operation due to two main reasons. First, resilience schemes to maintain memory reliability can impose high-performance overhead in future memory systems. Future sub-20nm DRAMs require stronger resilience techniques, such as in-DRAM error correction codes (IECC) [8, 34, 58, 62, 64, 67] and rank-level error correction codes (RECC) [46, 48, 75], compared to current commodity DRAMs. Figure 1 shows that such error correction code (ECC) schemes impose significant performance overhead to the memory hierarchy that employs sub-20nm DRAM as main memory. Second, even with strong resilience schemes, future memory systems do not appear to be as reliable as the state-of-the-art. In fact, due to increased bit error rates (BERs), even when we combine RECC and IECC, sub-20nm-DRAM-based memory systems with 3D-stacked DRAM LLC can have lower reliability than a 28nm-DRAM-based memory system that employs only RECC [8], as seen in Figure 1 (when comparing the triangle and the circle).

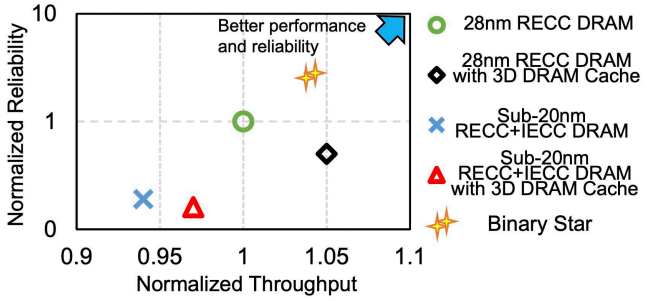

Figure 1: Reliability and throughput of various LLC and main memory hierarchy configurations normalized to 28nm DRAM with RECC. We define "reliability" as the reciprocal of device failures in time (FIT) [8, 79], i.e., 1/FIT. Reliability data is based on a recent study [8]. Throughput is measured using the benchmarks and system configurations described in Section 5. SRAM caches, which are less critical to scalability are kept constant across systems.

To address the DRAM technology scaling challenges, next-generation servers will adopt various new memory technologies. For example, Intel's next-generation Xeon servers and Lenovo's ThinkSystem SD650 servers support Optane DC PM [27, 29], which