The Feasibility of Utilizing Low-Power DRAM in Disaggregated Memory Systems

Ayaz Akram
yazakram@ucdavis.edu
Deptartment of Computer Science,
University of California, Davis

ABSTRACT

This study investigates low-power DRAM's viability in disaggregated memory systems. Through gem5 simulations and NAS Parallel Benchmark suite evaluations, we find that higher Compute Express Link (CXL) latencies reduce performance gaps between different DRAM devices, especially for LPDDR5. This insight informs the trade-offs between memory performance and system efficiency, impacting future disaggregated memory system design.

1 INTRODUCTION

Modern computing architectures grapple with memory constraints, and as a remedy, disaggregated memory systems have emerged as an exciting prospect [2, 5–8, 11]. These systems facilitate the sharing of memory resources across numerous nodes, thereby augmenting capacity and adaptability. Nevertheless, it is not clear how the memory device's properties can affect the overall system performance when the devices are used in a disaggregated setting. This study aims to evaluate the performance of diverse memory devices when employed as remote memory within disaggregated systems.

DRAM technologies have progressed significantly, introducing fresh technologies and architectural enhancements over the years. Presently, a wide spectrum of modern DRAM devices is available, each characterized by distinct attributes and performance characteristics. To unlock the full potential of these devices, it is imperative to comprehend their behavior and assess their suitability for high performance applications.

Our hypothesis is that the influence of a memory device's performance on overall system performance diminishes as these memory devices are integrated into a disaggregated environment. Therefore, low performance (higher latency and low bandwidth) but low power DRAM devices might be more feasible for future disaggregated memory systems. This study seeks to investigate this phenomenon through experiments.

2 EXPERIMENTS CONDUCTED

2.1 Methodology

We employ the gem5 simulator [3, 10] to conduct a series of experiments involving various DRAM devices. Our evaluation centers on workloads traditionally used to benchmark HPC systems, specifically the NAS Parallel Benchmark suite (NPB) [1]. This suite encompasses a variety of kernels and pseudo applications, serving as a long-standing tool for scrutinizing HPC systems. To expedite simulations, we focus on a limited execution interval of these benchmarks.

The details of the evaluated system are shown in Table 1. Our simulation models an 8-core CPU system featuring a two level cache hierarchy and a main memory outfitted with diverse DRAM devices. For our experiments, we select three distinct DRAM devices: DDR5_6400 (peak bandwidth: 51.2GB/s), DDR4_2400 (peak bandwidth: 19.2GB/s), and LPDDR5_6400 (12.8GB/s). The modeled DDR5 device comprises two individual channels, akin to real DDR5 devices, each with a peak bandwidth of 25.6GB/s. In addition, we incorporate various CXL [12] latencies, ranging from 50ns to 200ns.

Table 1: System Configuration Used for Experiments

Processors	
Number of cores	8
Frequency	5 GHz
Core type	Out of order, 8 wide
ROB entries/core	192
On-chip Caches	
Private L1 Inst.	32 KB
Private L1 Data	512 KB
Shared L2	8 MB
Main Memory	
Capacity	128GiB
Devices tested	DDR4_2400, LPDDR5_6400, DDR5_6400
Channels	1, 1, 2
Peak BW	19.2 GB/s, 12.8 GB/s, 51.2GB/s
Read/Write Buffer	64 entries each per channel
tRCD	14.16ns, 18ns, 14.375ns
tRAS	32ns, 42ns, 32ns
tRP	14.16ns, 18ns, 14.375ns
tCL	14.16ns, 21.25ns, 14.375ns
Tested CXL Attached Memory Latencies	
Round trip latency	0ns, 50ns, 100ns, 200ns

3 RESULTS

We present the outcomes of our experiments in Figure 1. This figure illustrates a comparison of execution times across various DRAM devices for diverse NPB applications under different CXL latencies. Upon analysis, we note that as CXL latency increases, distinctions in performance among DRAM devices diminish. Notably, this trend is particularly pronounced in the case of low-power DRAM, such as LPDDR5.

Table 2 outlines the normalized execution times of different DRAM devices relative to DDR5 for varying CXL latencies. For example, the difference in geometric mean execution time between DDR5 and LPDDR5 decreases from 88% to 23% as we transition from no CXL latency to a CXL latency of 200ns.

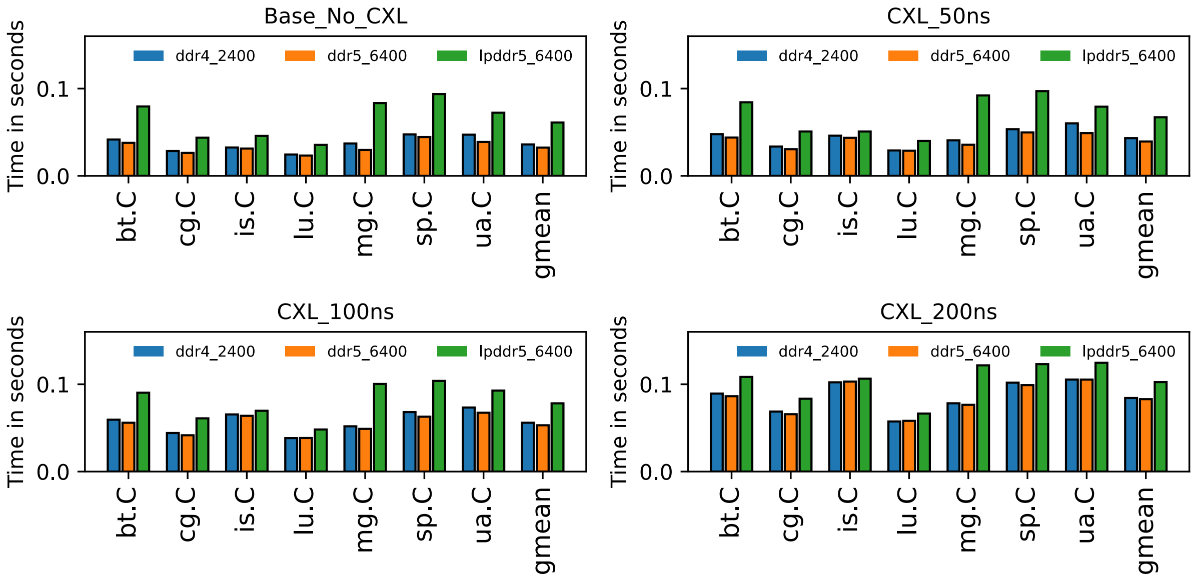

Figure 1: Execution time of different NAS Parallel Benchmarks for different DRAM types with different CXL latencies.

Table 2: Normalized Execution time of DRAM devices to DDR5 for different CXL latencies

DRAM	No CXL	50ns	100ns	200ns
DDR5_6400	1	1	1	1
DDR4_2400	1.11	1.09	1.05	1.01
LPDDR5_6400	1.88	1.70	1.47	1.23

3.1 Implications for Power Consumption in Disaggregated Systems

The surge in remote memory accesses within disaggregated memory systems has intensified concerns regarding power consumption. A substantial contributor to DRAM power consumption lies in the I/O interface responsible for transmitting data bits across the data bus, as highlighted by a study [9]. Notably, the power consumption of the system's I/O can far exceed that of on-chip I/O, with a notable difference of 10pJ/bit compared to 0.5pJ/bit [4]. The advantageous inclination of low-power DRAM devices towards disaggregated memory systems offers promising prospects for curbing overall power consumption costs by replacing more power-intensive DRAM counterparts like DDR5.

4 CONCLUSION

In conclusion, our study delved into the performance implications of employing diverse memory devices as remote memory within disaggregated systems. The findings underscore that the disparities in performance, particularly for low-power DRAM like LPDDR5, become less pronounced with heightened CXL latencies. This insight holds implications for the design and optimization of disaggregated memory systems, shedding light on the trade-offs between memory performance and system efficiency.

REFERENCES

[1] David H Bailey, Eric Barszcz, John T Barton, David S Browning, Robert L Carter, Leonardo Dagum, Rod A Fatoohi, Paul O Frederickson, Thomas A Lasinski, Rob S Schreiber, et al. 1991. The NAS Parallel Benchmarks. *The International Journal of Supercomputing Applic ations* 5, 3 (1991), 63–73.

[2] Daniel S Berger, Daniel Ernst, Huaicheng Li, Pantea Zardoshti, Monish Shah, Samir Rajadnya, Scott Lee, Lisa Hsu, Ishwar Agarwal, Mark D Hill, et al. 2023. Design Tradeoffs in CXL-Based Memory Pools for Public Cloud Platforms. *IEEE Micro* 43, 2 (2023), 30–38.

[3] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, et al. 2011. The gem5 Simulator. *ACM SIGARCH Computer Architecture News* 39, 2 (May 2011), 1–7.

[4] Allan Cantle. 2022. Redefining Computing Architecture Boundaries with Off Package Chiplets. In *HiPChips Chiplet Workshop in conjunction with the International Symposium on Computer Architecture.*

[5] Albert Cho, Anish Saxena, Moinuddin Qureshi, and Alexandros Daglis. 2023. A Case for CXL-Centric Server Processors. *arXiv preprint arXiv:2305.05033* (2023).

[6] Nan Ding, Samuel Williams, Hai Ah Nam, Taylor Groves, Muaaz Gul Awan, LeAnn Lindsey, Christopher Daley, Oguz Selvitopi, Leonid Oliker, and Nicholas Wright. 2022. Methodology for Evaluating the Potential of Disaggregated Memory Systems. In *2022 IEEE/ACM International Workshop on Resource Disaggregation in High-Performance Computing (REDIS).* IEEE, 1–11.

[7] Donghyun Gouk, Miryeong Kwon, Hanyeoreum Bae, Sangwon Lee, and Myoungsoo Jung. 2023. Memory pooling with cxl. *IEEE Micro* 43, 2 (2023), 48–57.

[8] Huaicheng Li, Daniel S Berger, Lisa Hsu, Daniel Ernst, Pantea Zardoshti, Stanko Novakovic, Monish Shah, Samir Rajadnya, Scott Lee, Ishwar Agarwal, et al. 2023. Pond: CXL-based memory pooling systems for cloud platforms. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2.* 574–587.

[9] Shang Li, Dhiraj Reddy, and Bruce Jacob. 2018. A performance & power comparison of modern high-speed dram architectures. In *Proceedings of the International Symposium on Memory Systems.* 341–353.

[10] Jason Lowe-Power, Abdul Mutaal Ahmad, Ayaz Akram, Mohammad Alian, Rico Amslinger, Matteo Andreozzi, Adrià Armejach, Nils Asmussen, Brad Beckmann, Srikant Bharadwaj, et al. 2020. The gem5 simulator: Version 20.0+. *arXiv preprint arXiv:2007.03152* (2020).

[11] George Michelogiannakis, Yehia Arafa, Brandon Cook, Liang Yuan Dai, Abdel Hameed Badawy, Madeleine Glick, Yuyang Wang, Keren Bergman, and John Shalf. 2023. Efficient Intra-Rack Resource Disaggregation for HPC Using Co-Packaged DWDM Photonics. *arXiv preprint arXiv:2301.03592* (2023).

[12] Stephen Van Doren. 2019. Hoti 2019: Compute express link. In *2019 IEEE Symposium on High-Performance Interconnects (HOTI).* IEEE, 18–18.