



HAL
open science

Understanding Power Consumption Metric on Heterogeneous Memory Systems

Andrès Rubio Proaño, Kento Sato

► **To cite this version:**

Andrès Rubio Proaño, Kento Sato. Understanding Power Consumption Metric on Heterogeneous Memory Systems. 2023 IEEE 29th International Conference on Parallel and Distributed Systems (ICPADS), Dec 2023, Ocean Flower Island, France. pp.2859-2862, <10.1109/ICPADS60453.2023.00408>. <hal-05075654>

HAL Id: hal-05075654

<https://hal.science/hal-05075654v1>

Submitted on 20 May 2025


HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Understanding Power Consumption Metric on Heterogeneous Memory Systems

Andr es Rubio Proa o 
RIKEN R-CCS
Kobe, Japan
andres.rubioproano@riken.jp

Kento Sato 
RIKEN R-CCS
Kobe, Japan
kento.sato@riken.jp

Abstract—Contemporary memory systems contain a variety of memory types, each possessing distinct characteristics. This trend empowers applications to opt for memory types aligning with developer’s desired behavior. As a result, developers gain flexibility to tailor their applications to specific needs, factoring in attributes like latency, bandwidth, and power consumption. Our research centers on the aspect of power consumption within memory systems. We introduce an approach that equips developers with comprehensive insights into the power consumption of individual memory types. Additionally, we propose an ordered hierarchy of memory types. Through this methodology, developers can make informed decisions for efficient memory usage aligned with their unique requirements.

Index Terms—power consumption; nvm; heterogeneous memory

I. INTRODUCTION

High-performance computing (HPC) plays a crucial role in addressing a wide range of complex scientific challenges by utilizing advanced models and simulations.

Over the years, the architecture of supercomputers has undergone significant evolution to meet diverse computing requirements, including enhancing application performance and optimizing power utilization. A critical aspect of this evolution lies in the memory system, which has evolved into a heterogeneous structure with multiple levels or hierarchies, incorporating various technologies. In Figure 1, the memory-storage continuum illustrates the primary technologies associated with each hierarchy level.

Heterogeneity within memory systems is driven by the diverse nature of applications, each exhibiting affinities towards specific memory devices. By leveraging these different memory technologies, improvements in latency, bandwidth, and memory power consumption behaviors can be achieved.

Given the above, this paper presents a methodology that allows expanding the knowledge of the memory system in terms of power, offering the possibility of obtaining an ordering of memories given a heterogeneous memory system when using different applications.

The paper structure considers first a Section II that provides an overview of the pertinent memory devices and their integration into Heterogeneous Memory Systems(HMS). In Section IV, we elucidate the driving factors behind this study. Section V outlines our devised methodology for assessing application power consumption within HMS. Subsequently,

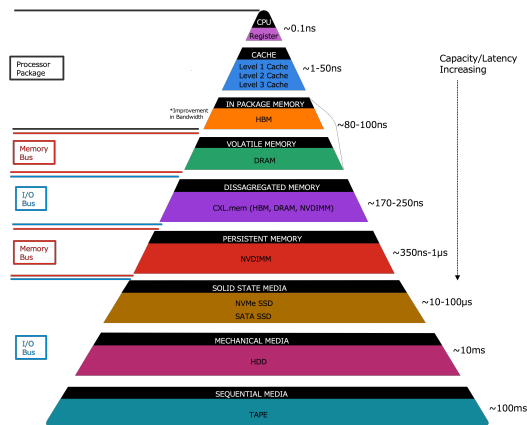


Fig. 1. Memory-Storage Continuum.

Section VI employs a suite of memory-bound applications to analyze memory power consumption within a heterogeneous memory system. The paper concludes by presenting future prospects and conclusions.

II. STATE OF THE ART

The acceleration in computation performance has been increasing enormously. However, memory systems have not been able to keep up with it.

Here we present some of the most relevant memory technologies.

A. HBM

Nowadays, High Bandwidth Memory (HBM) modules offers higher bandwidth and have become a part of other high-performance computing systems, including high-end GPUs and nodes equipped with A64FX processors.

B. NVM

NVM, is a type of memory that entered the market as a result of Intel’s decision to bridge the growing gap between DRAM and NAND SSD. Their flexibility enables them to co-exist with DRAM, fostering a heterogeneous environment that harmonizes the strengths of different memory technologies.

C. CXL.mem

CXL, which stands for Compute Express Link, instead of being a distinct form of memory, CXL is, in fact, an industry-endorsed Cache-Coherent Interconnect meticulously engineered to seamlessly integrate processors, memory, and accelerators. This standard is underpinned by an innovative protocol known as CXL.mem [1].

D. DRAM

A notable contrast to traditional DRAM DIMMs is the use of LPDDR (Low-Power Double Data Rate) memory in NVIDIA Grace CPUs [2]. However, there are still gaps in the memory-storage continuum. The need for these gaps depends on the requirements of applications to reach new levels of performance and power efficiency.

III. RELATED WORK

In earlier days, when memory systems exhibited homogeneity, there existed a sense of implicit assumptions regarding several factors, such as memory power. However, the landscape has transformed with the advent of heterogeneous memory systems.

Among these changes, several studies aimed to better handle diverse memory systems. For instance, in a paper referenced as [3], researchers explored ways to manage memory in heterogeneous systems.

In [4] we can observe techniques for power efficient computing in HPC systems by presenting measurement solutions that consider scalability, resolution and accuracy in power consumption measurements.

In [5], we can observe how authors give relevance to the memory power consumption metric. In fact, some old counters were used to understand power in memory; however, until the release of this paper, there was no way to differentiate the power consumption of each kind of memory.

IV. MOTIVATION

High-Performance Computing (HPC) systems have witnessed remarkable advancements in recent years, enabling researchers and scientists to tackle increasingly complex and data-intensive computational problems. A significant driver of this progress lies in the utilization of heterogeneous memory systems, which combine diverse memory technologies such as dynamic random access memory (DRAM), high bandwidth memory (HBM), and non-volatile memory (NVM).

However, with the proliferation of diverse memory devices, a new challenge arises: power consumption optimization. As the scale and complexity of scientific simulations and data-intensive applications continue to expand, the power consumed by the memory subsystem becomes a significant portion of the total power consumption in the system [6].

Given this, there is a critical need to provide a comprehensive ranking or ordering of memory devices based on their power consumption characteristics.

The power provisioning approach, as exemplified in [7], demands the optimal allocation of each component within

the High-Performance Computing (HPC) ecosystem, aligned with specific power requirements. Thus, the development of a pertinent strategy becomes pivotal, allowing us to establish a strategic ordering among memory devices. This ordering, in turn, contributes to effective power provisioning and allocation.

V. METHODOLOGY

In the present day, applications are capable of generating diverse workloads characterized by unique behaviors, which can respond differently within a given HMS. Additionally, the proliferation of various HMS configurations has introduced challenges for developers.

Here, we present the methodology employed for characterizing the power consumption metric within a HMS. We commence by detailing the process of configuring and exposing our memory system. Subsequently, we expound upon the selection of benchmarks and applications. Additionally, we introduce the tools to extract the requisite performance counters. Lastly, we use the information and undertake a comprehensive comparison of power performance variations across the various memory kinds.

A. Identifying Heterogeneous Memory Systems

The task of identifying memory system targets holds significant importance for both developers and applications. Particularly with the advent of HMS that incorporate multiple memory types, each characterized by distinct properties like bandwidth, latency, and power consumption.

A tool that greatly facilitates this task is `hwLoc` [8]. This tool streamlines the discovery of hardware topologies in HPC platforms.

It is crucial to highlight that NVM possesses the flexibility for two distinct applications: serving as storage or functioning as a memory device. In our study, our focus lies on utilizing NVM as a memory device to establish a HMS that incorporates both NVM and DRAM.

The configuration of our test machine features two sockets housing second-generation Intel Xeon Scalable processors and distinct types of memory local to each socket. This configuration results in a total of four NUMA nodes.

TABLE I
POWER CONSUMPTION OVER DIFFERENT APPLICATIONS IN A HMS
USING 18 THREADS AND A BIG PROBLEM SIZE

| Application.size | DRAM Power Consumption [Watts] | NVM Power Consumption [Watts] |
|------------------|--------------------------------|-------------------------------|
| BT.A | 86.91 | 74.73 |
| CG.W | 87.5 | 77.71 |
| FT.A | 99.65 | 72.91 |
| IS.B | 97.41 | 74.39 |
| MG.B | 90.36 | 75.21 |
| hpcg.400 | 115.15 | 78.16 |
| miniFE.400 | 103.14 | 78.18 |
| XSBench.large | 77.8 | 73.56 |

The characterization of our memory system encompasses various metrics, including bandwidth, latency, capacity, power performance, and others.

In our study, the power consumption metrics are primarily gathered through hardware counters during application profiling. Given the absence of dedicated memory-kind power consumption counters, we resort to employing the global memory power counter. To differentiate power consumption, we bind the entire process to each memory type, allowing us to work around this limitation.

In the context of power-limited environments, the power consumption metric assumes vital importance. This is exemplified in [9], where NVM showcases higher power consumption efficiency compared to DRAM. Given our utilization of an HMS featuring three distinct memory types, the power performance ranking’s outcome remains less clear-cut.

In Table I, the disparities in power consumption across various applications are apparent (further elaborated in VI). These variations underline the potential for power savings when opting for NVM over DRAM.

B. Application

In the scope of our study, we deliberately opted for a variety of memory-bound applications. This selection includes two proxy applications from the Exascale Computing Project (ECP) - namely, miniFE and XSbench. We also incorporated the NASA Advanced Supercomputing (NAS) benchmark suite, encompassing applications such as Integer Sort (IS), Conjugate Gradient (CG), Multi-Grid (MG), Discrete 3D Fourier Transform (FT), and Block Tri-diagonal solver (BT). In addition, we subjected the applications to testing using the High Performance Conjugate Gradient (HPCG).

C. Profiling

In conducting our analysis, it is essential to take into account that the evaluations provide insights into power consumption across the complete memory system. This is due to the inherent limitation of tools in distinguishing between power consumption stemming from DRAM and that originating from NVM. However, by leveraging the capabilities of `hwloc`, we are able to allocate the entire process to various memory types, thereby facilitating a differentiation of power consumption among them.

To comprehensively understand the behavior of applications, particularly in terms of bandwidth, total power consumption, and memory system power consumption, we have employed performance counters managed by profilers. Our approach involves using two specific tools: The Intel Performance Counter Monitor (PCM) and Linux `perf`. In our study, we have utilized several PCM command-line utilities, including `pcm-numa`, `pcm-mem`, and `pcm-power` which retrieves information related to memory such as accesses throughput and power respectively [10].

The Linux `perf` tool enables the retrieval of main memory power consumption through the `MSR_DRAM_ENERGY_STATUS` register. It should be used with `hwloc - bind` to be capable to differentiate the power of each kind of memory.

This methodology empowers users to not only discern power consumption within memory but also to delineate this consumption based on memory types. The ability to differentiate between them becomes paramount a task achieved in our scenario through power consumption analysis.

VI. EVALUATION

This section provides a comprehensive overview of the evaluation of our strategy. Firstly, we will describe in detail how we exposed the memory system and the emulated kind of memory to increase heterogeneity in our evaluation. Then, we present an analysis of power consumption given an HMS.




Fig. 2. Emulation by binding application process to a remote memory target.

A. Emulating Heterogeneous Memory Systems

Our evaluation encompasses two memory types: DRAM and NVM. To enhance the diversity of our assessment, we have introduced a third memory type through emulation. Through a clever mechanism, we have successfully replicated this third type by leveraging one of CPU0’s remote access memories, such as CPU1’s local DRAM as indicated in Figure 2.

This strategic approach yields a memory module with distinct attributes compared to CPU0’s local memories.

B. HMS Power consumption analysis

In Figure 3 we present the power consumption of the Heterogeneous Memory System (HMS) using DRAM and NVM memories and thread counts. Our observations based on this figure include:

- In many cases, the power consumption of Remote DRAM is actually lower than that of Local DRAM. To reinforce this notion, it is crucial to reemphasize the concept of simulating memory through the utilization of the NUMA system.
- A key characteristic of an HPC application is its utilization of multithreaded executions. While this might result in heightened consumption within components like CPU cores, it is also evident that this phenomenon contributes to an escalation in memory power consumption. This effect arises from the concurrent access to memory facilitated by the multithreading nature of the application.
- There is a correlation of memory power consumption when increasing the parallelization level in all kinds of memory, i.e., there is a tendency for having more power consumption when increasing the number of threads.

Which memory should be chosen will depend on the application used and also on the memory system. We could say that the generated ordering depends on the memory system and the application.

- To achieve an equilibrium between power efficiency and performance, the clarity of the situation might not be immediate. It's important to recognize that the second rank in the hierarchy may not consistently correspond to a single memory type.

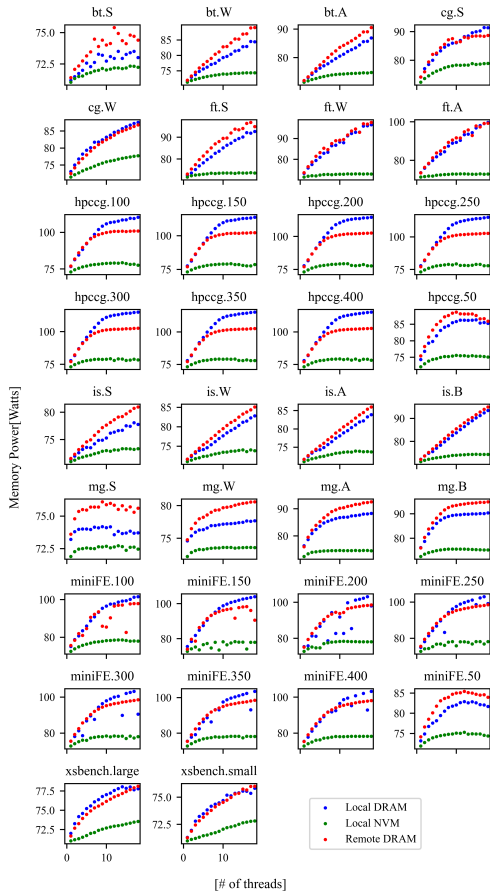


Fig. 3. Memory power consumption behavior taking into account the number of threads on executions bound to Local DRAM, Local NVM, and Remote DRAM.

VII. FUTURE WORK

The exploration of memory system behavior has revolved around determining a ranking that minimizes power consumption among memory types. We acknowledge that ultra-low power consumption can negatively impact application performance. Thus, a more favorable scenario could involve targeting a memory range with lower power consumption, potentially leading to enhanced application performance. This prompts us to delve into comprehending the intricate interplays among various memory metrics. The goal would be to identify optimal trade-offs that developers can leverage when seeking specific balances between these metrics.

VIII. CONCLUSION

This approach holds relevance not only for the specific HMS we have utilized, but also for other potential HMS configurations in principle. Of course, the strategy's extension would entail accommodating the intricacies that arise from accessing distinct memory types.

We have successfully derived an ordering among different memory types for a range of applications. Frequently, this ranking was not readily apparent, especially in cases where the memory type was unfamiliar. This methodology facilitates the creation of profiles that offer significant power savings, while also enabling the development of profiles that aim to strike a harmonious balance between power consumption and application performance.

Enhancing the user's comprehension of the memory system significantly promotes program portability across different memory architectures. Moreover, this serves to prevent the memory system from being underutilized or improperly employed.

REFERENCES

- [1] D. D. Sharma, R. Blankenship, and D. S. Berger, "An introduction to the compute express link (cxl) interconnect," 2023.
- [2] J. Evans, "Nvidia grace," in *2022 IEEE Hot Chips 34 Symposium (HCS)*. Los Alamitos, CA, USA: IEEE Computer Society, aug 2022, pp. 1–20. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/HCS55958.2022.9895599>
- [3] C. Cantalupo, V. Venkatesan, J. Hammond, K. Czurlyo, and S. D. Hammond, "memkind: An extensible heap memory manager for heterogeneous memory platforms and mixed memory policies." Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), Tech. Rep., 2015.
- [4] T. Ilsche, R. Schöne, J. Schuchart, D. Hackenberg, M. Simon, Y. Georgiou, and W. E. Nagel, "Power measurement techniques for energy-efficient computing: reconciling scalability, resolution, and accuracy," *SICS Software-Intensive Cyber-Physical Systems*, vol. 34, pp. 45–52, 2019.
- [5] H. David, E. Gorbatov, U. R. Hanebutte, R. Khanna, and C. Le, "Rapl: Memory power estimation and capping," in *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design*, ser. ISLPED '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 189–194. [Online]. Available: <https://doi.org/10.1145/1840845.1840883>
- [6] O. Sarood, A. Langer, L. Kalé, B. Rountree, and B. de Supinski, "Optimizing power allocation to cpu and memory subsystems in over-provisioned hpc systems," in *2013 IEEE International Conference on Cluster Computing (CLUSTER)*, 2013, pp. 1–8.
- [7] E. Arima, A. I. Comprés, and M. Schulz, "On the convergence of malleability and the hpc powerstack: Exploiting dynamism in over-provisioned and power-constrained hpc systems," in *International Conference on High Performance Computing*. Springer, 2022, pp. 206–217.
- [8] B. Goglin, "Exposing the locality of heterogeneous memory architectures to hpc applications," in *Proceedings of the Second International Symposium on Memory Systems*, ser. MEMSYS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 30–39. [Online]. Available: <https://doi.org/10.1145/2989081.2989115>
- [9] I. B. Peng, M. B. Gokhale, and E. W. Green, "System evaluation of the intel optane byte-addressable nvm," in *Proceedings of the International Symposium on Memory Systems*, ser. MEMSYS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 304–315. [Online]. Available: <https://doi.org/10.1145/3357526.3357568>
- [10] —, "System evaluation of the intel optane byte-addressable nvm," in *Proceedings of the International Symposium on Memory Systems*, 2019, pp. 304–315.