Page Scheduling on Hybrid Memory Systems with Machine Intelligence

Thaleia Dimitra Doudali*
Georgia Institute of Technology

Sergey Blagodurov
Advanced Micro Devices, Inc.

Abhinav Vishnu
Advanced Micro Devices, Inc.

Sudhanva Gurumurthi
Advanced Micro Devices, Inc.

Ada Gavrilovska
Georgia Institute of Technology

Abstract

Modern systems will feature heterogeneous technologies, in order to extend main memory capacity and service the increasing demand of Big Data analytics, in return for variability in the access latency and bandwidth of the memory components. In such hybrid memory systems, application performance enhancements are only viable through intelligent data management. To this extent, we propose a Page Scheduler with Machine Intelligence, that combines existing lightweight history-based data tiering methods with novel page scheduling, based on trained Recurrent Neural Networks. The Page Scheduler further makes a clever selection of the application pages that need machine intelligent based management. In this way, our solution can bridge 95% of the application performance gap that exists between the current state-of-the-art scheduling methodology and oracular data placement, while limiting the system's resource usage overheads, that are necessary for the learning process. Hence, this exploration lays the grounds for practical integration of Machine Intelligence into future systems.

Motivation. The current state-of-the-art [1] data management over hybrid memory systems, periodically migrates application pages such that the ones that are frequently accessed get prioritized for allocations on the memory component with the lowest access latency, until its capacity is full. However, this history-based approach provides limited performance enhancement opportunities, since it is depending on the application's data access pattern. Our experimental analysis over a broad variety of High Performance Computing benchmarks and mini-applications, shows that there can be up to **40%** performance gap between existing methodologies and an optimal solution with oracular prior knowledge of the access pattern.

Observations. We chose to deploy Recurrent Neural Networks, as their functionality best fits into the description of the Page Scheduler, who needs to predict future access pattern behavior. We identified that RNN-based predictions of per page access counts across the scheduling time intervals, can be highly accurate. More importantly, we observed that RNN-based scheduling of a cleverly selected subset of application pages can bring up to 90% of the possible performance improvements, while reducing up to 80% of the problem space that needs Machine Intelligence training. Finally, pages with easily predictable access pattern behavior can be optimally managed by existing lightweight history-based solutions.

Design. First, the Page Scheduler trains an individual Recurrent Neural Network for each of the cleverly selected pages, that will significantly boost performance, in order to learn how many accesses the page received during each scheduling interval. Then, he periodically aggregates the per page access counts (either RNN-based for the selected pages or history-based for the rest), orders pages in descending access count order and migrates the ordered pages in the fastest available memory component, until capacity is full.

Evaluation. We evaluate the Page Scheduler by simulating a hybrid memory system with DRAM and Non Volatile memory, using the runtime estimation model described in [1]. We show that our solution has the potential to bridge up to **95%** of the performance gap that exists between history-based and oracular page management. Also, we report low training and inference times for the RNN models. Our design aims to reduce the learning overhead, via cleverly selecting the number of pages, thus models, whose intelligent management will bring most of the possible application performance improvements.

Acknowledgment. AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

References

[1] M. R. Meswani, S. Blagodurov, D. Roberts, J. Slice, M. Ignatowski, and G. H. Loh. 2015. Heterogeneous memory architectures: A HW/SW approach for mixing die-stacked and off-package memories. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, Vol. 00. 126–136. https://doi.org/10.1109/HPCA.2015.7056027

*The only student author and presenter.