Optimizing Hybrid Access Virtual Memory System using SCM/DRAM Unified Memory Management Unit

Yusuke Shirota
Corporate Research & Development Center
Toshiba Corporation
Kawasaki, Kanagawa, Japan
yusuke1.shirota@toshiba.co.jp

Shiyo Yoshimura
Corporate Research & Development Center
Toshiba Corporation
Kawasaki, Kanagawa, Japan
shiyo.yoshimura@toshiba.co.jp

Satoshi Shirai
Corporate Research & Development Center
Toshiba Corporation
Kawasaki, Kanagawa, Japan
satoshi.shirai@toshiba.co.jp

Tatsunori Kanai
Corporate Research & Development Center
Toshiba Corporation
Kawasaki, Kanagawa, Japan
tatsunori.kanai@toshiba.co.jp

ABSTRACT

In HPC systems, expectations for storage-class memory (SCM) are increasing in large-scale in-memory processing. While SCM can deliver higher capacity and lower standby power than DRAM, it is slower and the dynamic power is higher. Therefore, in order to realize high-speed, low-power and scalable main memory, it is necessary to build an SCM/DRAM unified memory, and dynamically optimize data placement between the two memories according to the memory access pattern. In this paper, we propose a new hybrid access type virtual memory method using TLB-extended unified memory management unit which enables collecting and extracting fine-grained memory access locality characteristics. We show that with the proposed method, Hybrid Access control, which is a memory hierarchy control that selectively uses Direct Access to bus attached byte-addressable SCM and low power Aggressive Paging using small DRAM as cache, can be made more accurate, and the efficiency of memory access can be significantly improved.

KEYWORDS

storage-class memory (SCM), virtual memory system, TLB

1 INTRODUCTION

Recently, in HPC systems, the emergence of new applications such as Big Data analytics, deep learning (AI) and CPS/IoT has increased the demand for in-memory processing of large-scale data. In the meanwhile, a new type of byte-addressable high-speed non-volatile memory called SCM [2,3,4] such as MRAM, ReRAM, PCM and *3D XPoint is expected to fill the large gap of access latency between DRAM and SSD/HDD in the memory hierarchy.

SC'19, November, 2019, Denver, Colorado USA

https://doi.org/10.1145/1234567890

Since SCM has larger capacity than DRAM but lower latency, it is necessary to build an SCM/DRAM unified main memory, and distribute data to SCM and DRAM according to the data access pattern. Therefore, in order to achieve high performance, memory hierarchical control methods have been proposed that adaptively use SCM and DRAM according to the memory access pattern[1, 6].

Hybrid Access [1], which is a memory hierarchical control method, is proposed for efficient unified memory management in SCM-aware virtual memory system. It is desirable that data with high locality be accessed on DRAM after the data is copied from SCM to DRAM at page granularity. On the other hand, data with low locality should be kept on memory bus attached byte-addressable SCM and accessed directly at cache line granularity. Thus Hybrid Access adaptively switches between the two modes, SCM-aware low power Aggressive Paging (AP) with small DRAM as cache and Direct Access (DA) to SCM, according to data access patterns. Similarly, *Intel® Optane DC Persistent Memory [6] has Memory Mode and App Direct Mode. DRAM acts as a cache for the most frequently-accessed data in Memory Mode, while applications can direct which type of data read or write is suitable for DRAM or SCM in App Direct Mode. However, since they switch the mode of the entire system, further optimization is difficult when data suitable for each mode is mixed simultaneously. Therefore, in this paper, we propose a compact memory management hardware unit for SCM/DRAM unified main memory that has a function to extract fine-grained memory access locality characteristics, and automatically optimize Hybrid Access enabled virtual memory system [1] using the new hardware.

2 VIRTUAL MEMORY SYSTEM USING UNIFIED MEMORY MANAGEMENT UNIT

2.1 describes a method to collect accurate information on fine-grained main memory access characteristics by adding a compact hardware to the processor's TLB. 2.2 describes a virtual memory method that performs Hybrid Access control that optimally arranges data on SCM and DRAM according to the memory access pattern extracted by the new memory management unit.

2.1 TLB-extended Unified Memory Management Unit

In order to apply an appropriate control method, that is AP or DA, for each data, it is necessary to collect fine-grained access characteristic information. Conventional TLBs often have a reference bit that records whether or not access is performed in page units, and thus it can only detect memory access at page granularity. Therefore, in this paper, we propose a memory management unit that extends the TLB to detect cache line unit access pattern of each page. Conventionally, TLB is a table that associates virtual page number with translated physical page number. Thus, we extend TLB, so that each entry of TLB also has access flag field. The access flag field holds the access flag corresponding to each cache line in page. Memory management unit accesses SCM or DRAM via memory bus if memory access request from processor core miss the cache. At this time, the access flag in the access flag field of the table entry corresponding to the virtual address of the memory to be accessed is set to 1. By referring to the access flag field, it is possible to detect access patterns in cache line units.

2.2 Hybrid Access type Virtual Memory System

When determining the optimum method for a certain page, the memory management unit determines the access pattern by referring to the access flag field of that page. If the number of 1s in the field is equal to or greater than the predetermined threshold, the locality of the page is considered high. Therefore, the page is copied from SCM to DRAM to exploit that locality (AP). On the other hand, when the number of 1s is smaller than the threshold value, the locality of access is considered low, and it is determined that it is optimal for the processor to directly access on SCM (DA).

3 EVALUATION

We implemented a simulation platform for evaluating Hybrid Access (HA) type virtual memory system using a TLB-extended memory management unit. Memory address traces of PARSEC benchmarks [5] obtained on 2.60GHz *Intel® Xeon® processor E5-2690 v3 were used for evaluation. In the simulation, the number of entries and the entry replacement algorithm are matched to that of the TLB of the above processor. Latency varies depending on the type of SCM, but in the evaluation, we assumed relatively fast SCM with read latency and write latency 4× and 8× that of DRAM, and 4 KB page transfer time of 1 us. Even when setting the size ratio of DRAM to SCM to 1:9 using small DRAM, memory access time was reduced by 29.4%, 71.2% and 68.2% respectively for canneal, facesim and raytrace by using HA method compared to using SCM only. Figure 1 depicts the breakdown of the working set size (WSS) for canneal in which the random access is dominant. Since canneal has a large percentage of randomly accessed pages, DA is properly selected for most pages. Moreover, some pages with high locality included in canneal is properly extracted, and for these pages, AP is selected, and the efficiency is improved by utilizing the locality on DRAM. Figure also shows that efficient memory access is being performed, without unnecessary data transfer between SCM and DRAM. Figure 2 depicts the breakdown of memory access time for canneal. Since small but frequently accessed data is successfully processed on DRAM with HA, it was shown that by concurrently utilizing AP and DA separately for each data, memory access time can be significantly reduced than the conventional mode switching of the two methods. Memory access time reduction rate will further increase if we assume SCM with higher latency.

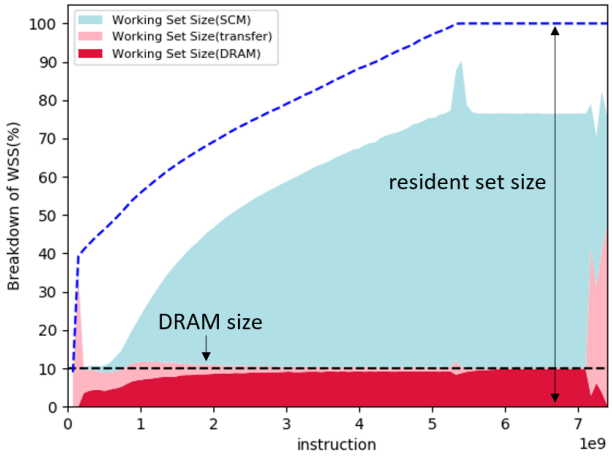

Figure 1: WSS breakdown with threshold set to 4 (canneal).

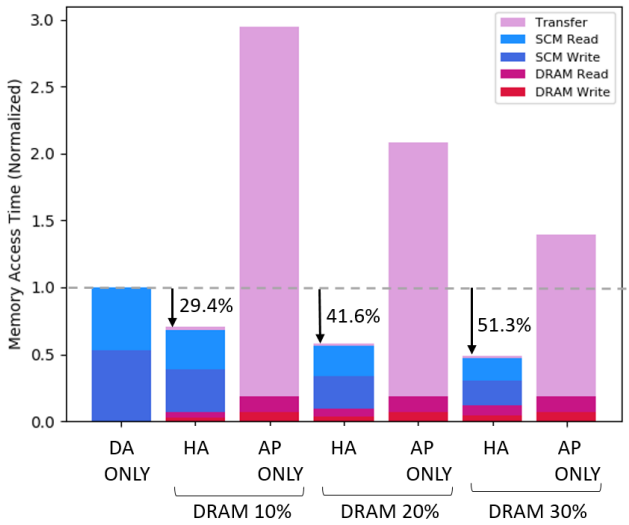

Figure 2: Performance evaluation of Hybrid Access method.

4 SUMMARY

This paper proposed a Hybrid Access type virtual memory system using TLB-extended memory management unit for SCM/DRAM unified memory. Our evaluation reveals that proposed method can automatically extract data for AP and DA, and further improves memory access efficiency. We will explore dynamic DRAM size adjustment and threshold setting suitable for SCM of various characteristics.

ACKNOWLEDGMENTS

This work was based on results obtained from a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

REFERENCES

[1] Y. Shirota, S. Shirai and T. Kanai, "Hybrid Access in Storage-class Memory-aware Low Power Virtual Memory System," 2019 IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS), pp. 1-3, 2019.

[2] R. F. Freitas and W. W. Wilcke, "Storage-class Memory: The Next Storage System Technology," IBM Journal of Research and Development, vol. 52, no. 4, pp. 439-447, 2008.

[3] G. Sun, J. Zhao, M. Poremba, C. Xu, Y. Xie, "Memory that never forgets: emerging nonvolatile memory and the implication for architecture design, " National Science Review, vol. 5, no. 4, pp. 2018, pp. 577-592, 2018.

[4] K. Hoya, K. Hatsuda, K. Tsuchida, Y. Watanabe, Y. Shirota and T. Kanai, "A perspective on NVRAM technology for future computing system," 2019 International Symposium on VLSI Design, Automation and Test (VLSI-DAT), pp. 1-2, 2019.

[5] C. Bienia, S. Kumar, J. P. Singh, K. Li, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," Princeton University Technical Report TR-811-08, 2008.

[6] Intel. https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html.