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

Yusuke Shirota, Shiyo Yoshimura, Satoshi Shirai, Tatsunori Kanai
yusuke1.shirota@toshiba.co.jp, shiyo.yoshimura@toshiba.co.jp, satoshi.shirai@toshiba.co.jp, tatsunori.kanai@toshiba.co.jp
Toshiba Corporation, Kanagawa, Japan

1. Introduction

- New type of byte-addressable high-speed nonvolatile memory called SCM (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
- Memory hierarchical control methods have been proposed to use both SCM and DRAM
  - **Hybrid Access Method**: Aggressive Paging & Direct Access [1]
  - Intel Optane DC Persistent Memory: Memory Mode & App Direct Mode [2]
- To dynamically apply appropriate controls to each data, it is necessary to collect fine-grained memory access characteristic information
- However, conventional TLBs can only detect memory access at page granularity

2. Hybrid Access type Virtual Memory System using TLB-extended Unified Memory Management Unit

- Hybrid Access type Virtual Memory System
  - selectively uses **Direct Access** (DA) to bus attached SCM at cache line granularity and low power **Aggressive Paging** (AP) using small DRAM as cache
- TLB Extension
  - Each TLB entry has access flag field to detect cache line unit access pattern of each page
- Optimizing Hybrid Access
  - When determining the optimum method for a certain page,
    - # of 1s is equal to or greater than threshold → select Aggressive Paging
    - # of 1s is smaller than threshold → select Direct Access

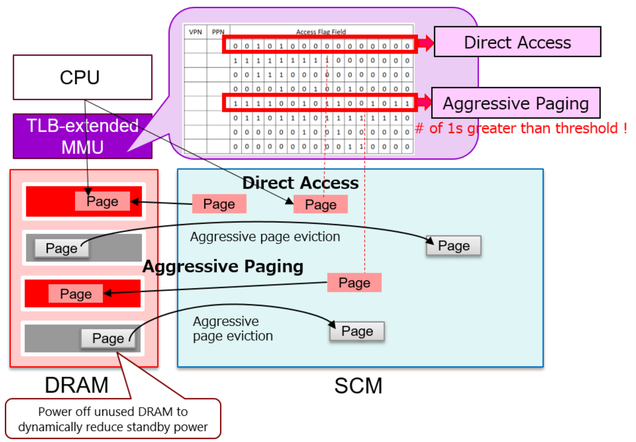

3. Evaluation

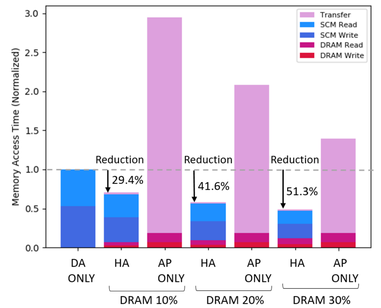

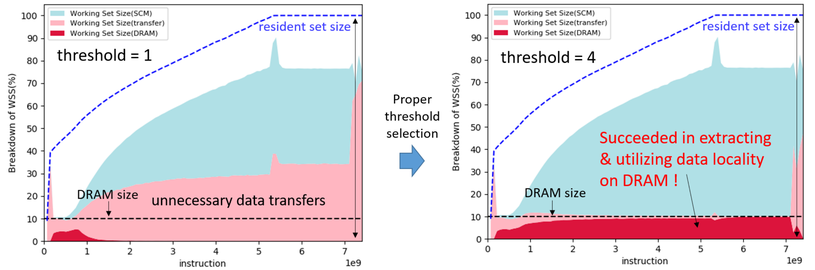

[Performance evaluation of Hybrid Access] [Working set size breakdown with different thresholds (canneal)]

Proper threshold selection is also important to avoid unnecessary data transfers between SCM and DRAM.

- Evaluated Hybrid Access (HA) on simulation platform using memory address traces of PARSEC benchmarks obtained on 2.60GHz *Intel® Xeon® processor E5-2690 v3
- Hybrid Access succeeded in reducing memory access time (**29.4%, 71.2% and 68.2%** respectively for canneal, facesim and raytrace) even with small DRAM (setting size ratio of DRAM to SCM to 1:9) compared to using SCM only
- Even for canneal where random access is the majority, pages adequate for AP were accurately extracted. As a result, memory access time was reduced by utilizing AP and DA separately for each page
- We assumed SCM with relatively low latency, but if latency is higher, the reduction rate will increase because the effect of using locality on DRAM will increase

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), 2019.
[2] Intel. https://www.intel.com/content/www/us/en/architecture-andtechnology/optane-dc-persistent-memory.html.

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