Memory-Aware Management for Multi-Level Main Memory Complex using an Optimization of the Aging Paging Algorithm

Gal Oren
Ben-Gurion University of the Negev
Nuclear Research Center Negev
orenw@post.bgu.ac.il

Leonid Barenboim
The Open University of Israel
leonidb@openu.ac.il

Lior Amar
Parallel Machines
liorорama@gmail.com

Categories and Subject Descriptors D.4.2 [*Operating Systems*]: Storage Management, Storage Hierarchies.

Keywords Memory Hierarchies, Paging, Aging, SCM.

Abstract

Memory and storage are often assumed to be unsophisticated, flat resources, with simple properties, such as a constant access time. Over the years this assumption has been proven to be wrong, and understanding of the memory hierarchy could be useful in order to enhance the performance of an algorithm or a data structure [1]. For example, the Storage Class Memory (SCM) is a new technology which represents a new hybrid form of storage and memory with uniqe characteristics, meaning a memory which is non-volatile, cheap in per bit cost, has fast access times for both read and writes using cache line access, and is solid state. Also, the SCM is supposed to have different versions with different access speeds and volumes, meaning that it might be possible to add different SCM devices to the memory hierarchy as an extension of the RAM, and manage this enlarged main memory complex using special algorithms [2].

A combination between the SCM technology and a designated Memory Allocation Manager (MAM) that will allow the developer to manually control the different memories will be likely to achieve a new level of performance for memory-aware data structures. However, although the manual MAM seems to be the optimal approach for multi-level main memory complex management, this technique is still very far from being realistic because of several reasons, and the chances that it would be implemented in current codes using High Performance Computing (HPC) platforms is quite low. This premise means that the most reasonable way to introduce the SCM into any usable memory system would be by implementing an automated version of the MAM using the fundamentals of the paging algorithms, as used for a standard memory hierarchy. Our hypothesis is that achieving appropriate transferability between these new main memory complex levels may be possible using ideas of algorithms employed in current virtual memory systems, and that the memory-aware adaptation of those algorithms to a multi-level main memory complex is possible.

We investigated various paging algorithms, and found the ones that could be adapted successfully from a standard memory hierarchy to a hierarchy with multi-level main memory complex. We discovered that using a memory-aware adaptation of the Aging paging algorithm results in the best performances in terms of Hit / Miss ratio and access speed. Specifically, we show that this modification can improve the access speed of the main memory complex by about 75%, and that the new algorithm manages to achieve the same or better Hit / Miss ratio in almost all cases in comparison to the current alternatives.

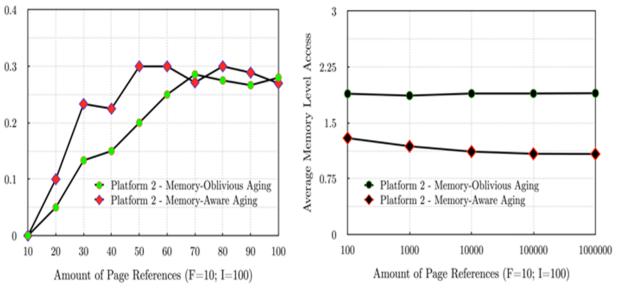

Figure 1 Memory-Oblivious vs. Memory-Aware Algorithm: Hit/Miss ratio (left), and Average Memory-Level (right).

References

[1] R. Dementiev, L. Kettner, and P. Sanders. STXXL: Standard Template Library for XXL Data Sets. In *Softw., Pract. Exper. 38.6*, pages 589–637, 2008.

[2] D. Hession, N. McKelvey, and K. Curran. Storage Class Memory. In *International J. E-Business Development 4.1*, 2013.

SYSTOR '16, June 06-08, 2016, Haifa, Israel.

ACM 978-1-4503-4381-7/16/06.
http://dx.doi.org/10.1145/2928275.2933276