

Efficient Use of Low Cost SSDs for Cost Effective Solid State Caches

Yongseok Oh^{*1}, Eunjae Lee¹, Jongmoo Choi², Donghee Lee¹, and Sam H. Noh³

¹University of Seoul, Seoul, Korea, {ysoh, kckjn97, dhl_express}@uos.ac.kr

²Dankook University, Gyeonggi-do, Korea, choijm@dankook.ac.kr

³Hongik University, Seoul, Korea, http://next.hongik.ac.kr

1 Introduction

Conventional Hard Disk Drives (HDDs) and emerging Solid State Drives (SSDs) each has strengths and limitations in terms of latency, price, and lifetime. To mitigate the limitations and amplify their advantages, the Solid State Cache (SSCs), which use flash based SSD as a data cache, have been introduced in previous studies [4, 5]. These studies have focused only on improving the performance and endurance of the storage system with SSCs that comprises only of Single-Level Cell (SLC) or a Multi-Level Cell (MLC) flash memory chips. Our focus is on issues of Hybrid SSC that combines SLC, MLC, and TLC SSCs for performance and price efficiency.

Figure 1 shows the motivation of our work. The x-axis represents the capacity of SSCs per price (\$) and the y-axis represents the performance and endurance of SSCs. For clarity, we simply visualize the characteristics of each SSC without absolute values in the figure. In the figure, SLC-based SSC shows the best performance and write endurance at the expense of price. In contrast, MLC- and TLC-based SSCs are cheaper but they suffers from low performance and endurance. To meet the performance and price requirements, we consider hybrid approaches such as SLC+MLC, SLC+TLC, and MLC+TLC SSC combinations that are indicated by the shaded triangle in Figure 1. Specifically, we propose a management scheme of Hybrid Solid State Cache, namely *HySSC*, that integrates high performance SLC SSCs and low cost TLC SSCs.

2 Hybrid Solid State Cache

Let us now describe the design issues of *HySSC*. Figure 2 shows the overall structure of the storage system. The storage system consists of three layers, namely, the




Figure 1: Motivation of Hybrid Solid State Cache

HySSC Layer, the *SSC Layer*, and the *Disk Layer*. The *Disk Layer* has an HDD-based RAID system and provides large reliable storage service to upper layers. The *SSC Layer* consists of an SLC SSC and a TLC SSC. Typically, the TLC SSC has three orders of magnitude lower endurance than the SLC SSC while their read performance difference is relatively small. (Roughly, SLC has 100,000 P/E cycles, and TLC has 1,000 P/E cycles.) Due to their characteristics we decide to use TLC SSC as a read only cache.

Now, we describe the *HySSC Layer* that comprises three software modules, namely, *Page Replacer*, *Sequential I/O Detector*, and *Mapping Manager*. The *Sequential Detector* identifies consecutive I/O requests and makes them bypass the *SSC Layer* to prevent cache pollution. The *Page Replacer* uses two LRU lists to maintain the request order of data in SLC and TLC SSCs. If data not in the *SSC* is requested, the *Page Replacer* migrates the oldest data in the SLC SSC to TLC SSC or directly destages it to the *Disk Layer*. Specifically, if the oldest data is dirty, it is destaged to the *Disk layer*. Otherwise, the *Page Replacer* selectively removes the clean data from the SLC SSC or migrates it to the TLC SSC. We deliberately devised this selective migration to

*Student, Presenter

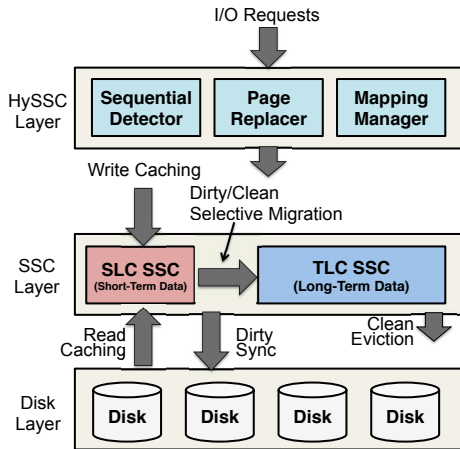


Figure 2: Architecture of Hybrid Solid State Cache

avoid frequent data updates that may hurt endurance of the TLC SSC. As a result, selective migration improves the endurance of the TLC SSC. To make room for the migrated data, the oldest data in the TLC SSC may be evicted. The Mapping Manager maintains a logical to physical address mapping table to quickly lookup the location of the requested data in the storage system.

3 Evaluation

To measure the performance of the storage system, we have developed a hybrid SSC simulator¹ based on the CMU DiskSim simulator [2] emulating HDDs and the MSR SSD extension [1] emulating SSDs. For performance evaluation, we use realistic workload traces Exchange and MSN that have been used in other studies [3, 4]. We set up the experiments to emulate pure SLC, MLC, and TLC SSCs with 32GB capacity and 10% over-provisioning space (OPS) as well as HySSC that comprises 16GB SLC SSC and 16GB TLC SSC with their aggregated OPS also set to 10%. The read and write operation times are 135us/350us for SLC, 175us/1.4ms for MLC, and 350us/2.5ms for TLC, respectively. The erase operation time is 1.5ms for SLC, 3.8ms for MLC, and 3ms for TLC, respectively. The page unit size is 4KB in both SLC and MLC chips and 8KB in TLC chips.

Figure 3 shows the response times of all SSC configurations. In the figure, the x -axis represents the name of the workloads and the y -axis denotes the mean response time of I/O requests in log scale. As we expected, the SLC SSC shows the best performance followed by the MLC SSC. The TLC SSC shows the worst performance due to slow operations of TLC flash memory chips. Moreover, the TLC SSC suffers from the small write problem because the page unit size of TLC flash memory chip is 8KB. Now we turn our attention to the performance of HySSC. In the figure, the performance of

¹The source is available at <http://embedded.uos.ac.kr/~hyssc>

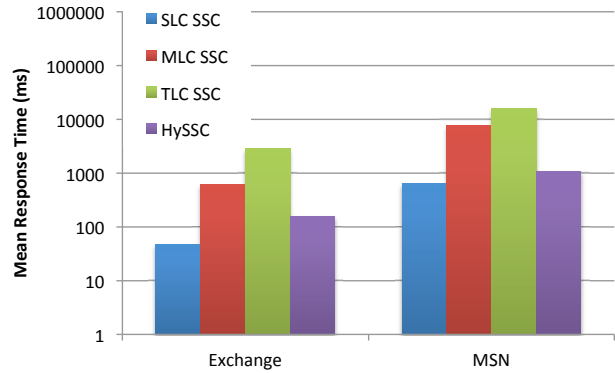


Figure 3: Measurement of Mean Response Time

HySSC is only slightly worse than that of SLC SSC. Despite that the HySSC comprises SLC and TLC chips, it outperforms MLC and TLC SSCs and its performance is almost on par with the SLC SSC. This experiment provides quantitative evidence that combining diverse flash memory chips in SSCs is a promising way to achieve both performance and price goals.

Acknowledgement This work was supported in part by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MEST) (No. 2012R1A2A2A01045733) and by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education, Science and Technology(2010-0025282)

References

- [1] AGRAWAL, N., PRABHAKARAN, V., WOBBER, T., DAVIS, J. D., MANASSE, M., AND PANIGRAHY, R. Design Tradeoffs for SSD Performance. In *Proc. of USENIX ATC* (2008), pp. 57–70.
- [2] BUCY, J. S., SCHINDLER, J., SCHLOSSER, S. W., AND GANGER, G. R. DiskSim 4.0. <http://www.pdl.cmu.edu/DiskSim/>.
- [3] NARAYANAN, D., THERESKA, E., DONNELLY, A., ELNIKETY, S., AND ROWSTRON, A. Migrating Server Storage to SSDs: Analysis of Tradeoffs. In *Proc. of EuroSys* (2009), pp. 145–158.
- [4] OH, Y., CHOI, J., LEE, D., AND NOH, S. H. Caching Less for Better Performance: Balancing Cache Size and Update Cost of Flash Memory Cache in Hybrid Storage Systems. In *Proceedings of FAST* (2012).
- [5] SAXENA, M., SWIFT, M. M., AND ZHANG, Y. FlashTier: A Lightweight, Consistent and Durable Storage Cache. In *Proceedings of EuroSys* (2012), pp. 267–280.