How Many MLCs Should Impersonate SLCs to Optimize SSD Performance?

The 2nd International Symposium on Memory Systems (MEMSYS), USA, October 3-6, 2016.

Wei Wang[1], Wen Pan[1], Tao Xie[2], and Deng Zhou[1], San Diego State University, San Diego, CA 92182

Introduction

Typically, flash memory devices fall into two categories, SLC and MLC. An SLC flash memory cell can store only 1 bit data, whereas 2 or more bits can be stored in an MLC memory cell. An MLC flash cell can work in an SLC mode by selectively using its fast pages. We find that different workloads running on a fixed partitioning configuration exhibit distinctive partition utilizations, and different utilizations lead to various levels of performance. We develop a new FTL called BROMS that integrates the dynamic partitioning method, which optimizes the overall performance of a partitioned SSD via adjusting the capacities of its two partitions on the fly.

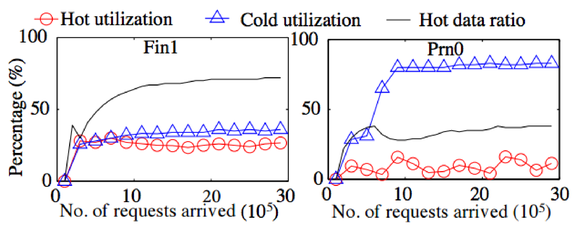

Fig. 1 Partition utilizations in Fin1 and Prn0.

Costs Analysis

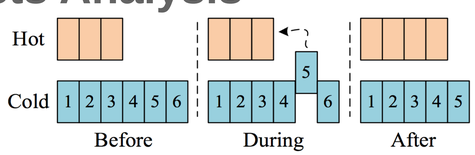

Fig. 2 An example of dynamic partitioning.

- A Dynamic Partitioning Method

$$\mu_c' = \frac{(1-\beta)\mu_h\mu_h'}{(1-\beta)\mu_h' - 2\beta\mu_h + 2\beta\mu_h'}$$

- Write Cost Model

$$C'_{overall} = \theta C'_{hot\_pw} + (1-\theta)C'_{cold\_pw} + \lambda\theta C'_m$$

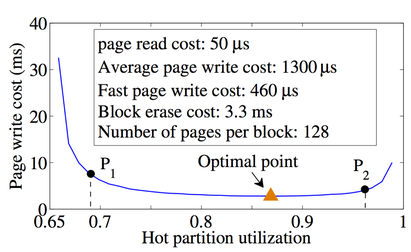

Fig. 3 An example of overall write cost.

- Optimal line is identified by this model

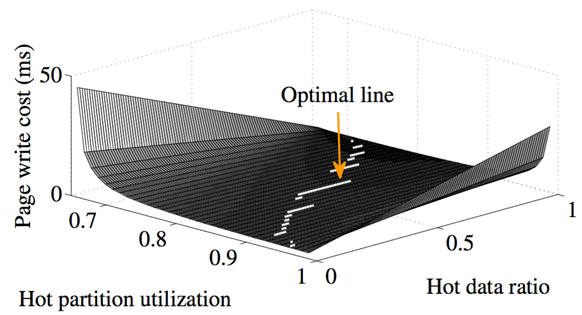

Fig. 4 Page write cost under different hot data ratios.

Model Verification

- We implement our scheme on an FPGA board to evaluate the accuracy of our model

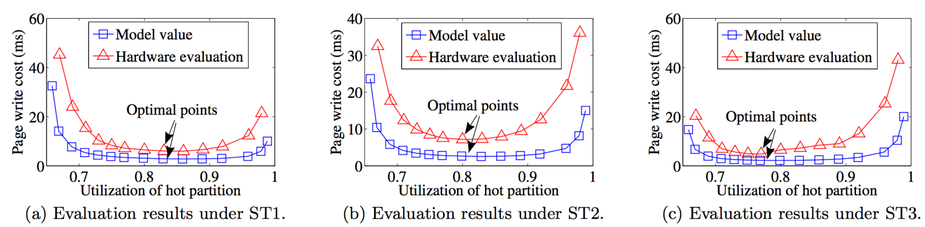

Fig. 5 Performance comparisons between model value and hardware evaluation results.

The BROMS FTL

- Architecture of BROMS.

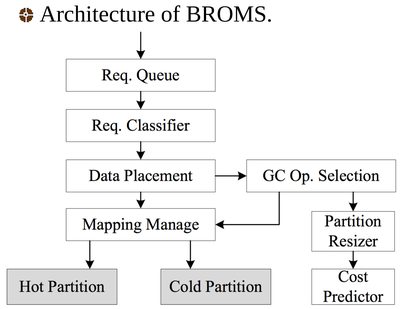

Fig. 6 Architecture of BROMS.

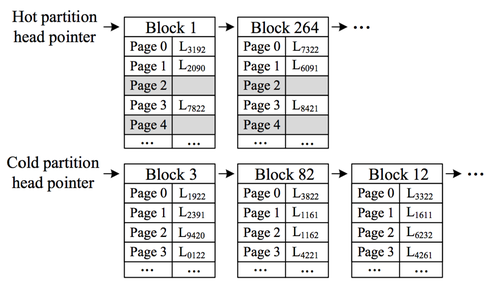

Fig. 7 An illustration of partition lists in BROMS.

EVALUATION OF BROMS

Table 1 Flash specifications [1].

Parameter	Value
Page size	8,192 + 44 bytes
Pages per block	256
Blocks per plane	2048
Planes per die	2
Chips per channel	1
Channel number	2
Read latency	75 µs
Write latency	1,300 µs
Erase latency	3.8 ms

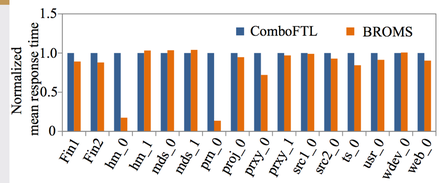

Fig. 8 Performance comparisons between the two FTLs.

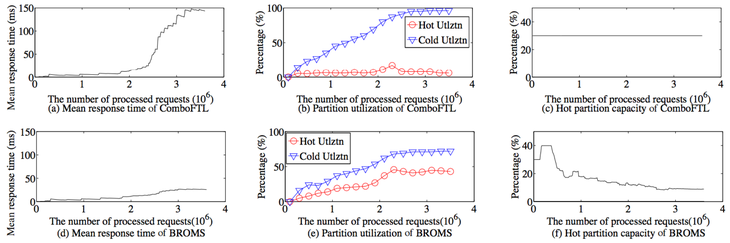

Fig. 9 A comparison between ComboFTL [2] and BROMS under the hm0 workload.

Conclusions

The conclusion is that a fixed partitioning normally leads to an inferior performance. To address this problem, we analyze the garbage collection cost in both partitions and then build a write cost model for an SSD. For each workload there always exists a best partition con- figuration

Acknowledgement

This work was supported by the US National Science Foundation under Grant CNS(CAREER)-0845105.

References

[1] Micron. Nand datasheet, MT29F64G08CEACA, 2014.
[2] S. Im and D. Shin. ComboFTL: Improving performance and lifespan of MLC flash memory using SLC flash buffer. Journal of Systems Architecture, 56(12):641 – 653, 2010.

1. Computational Science Research Center
2. Computer Science Department