Extended Abstracts of the 2014 International Conference on Solid State Devices and Materials, Tsukuba, 2014, pp106-107

# Design Guidelines of All Storage Class Memory (SCM) SSD and Hybrid SCM/NAND Flash SSD to Balance Performance, Power, Endurance and Cost

Takahiro Onagi[1], Chao Sun[1,2] and Ken Takeuchi[1]
[1]Chuo University, 1-13-27 Kasuga, Bunkyo-ku, Tokyo, 112-9551 Japan, [2]Univ. of Tokyo, Phone: +81-3-3817-7374
E-mail: onagi@takeuchi-lab.org

**Abstract**

**The performance, energy consumption, endurance and cost are compared between the All storage-type storage class memory (S-SCM) solid-state drive (SSD) and Hybrid memory-type SCM (M-SCM)/NAND flash SSD. A wear leveling algorithm has been proposed for the SCM due to its limited endurance. With a 25% M-SCM/NAND capacity ratio, the All S-SCM SSD is faster if the S-SCM latency is <5μs. Moreover, assuming the cost ratio of M-SCM/NAND is 12, All S-SCM SSD is more cost-efficient when S-SCM/NAND cost ratio is < 4.**

## 1. Introduction

Storage class memories (SCMs) are the emerging non-volatile memories [1, 2]. Since the characteristics of each SCM device are different, memory-type SCM (M-SCM) such as STT-MRAM is named for the DRAM-like SCM with a high performance but low capacity. The NAND flash-like SCM such as ReRAM and PRAM is called storage-type SCM (S-SCM), which has a high capacity while low performance. Fig. 1 describes the switching mechanism of a super-lattice PRAM (SL-PRAM). It can be an ideal candidate for the future high density S-SCM because of the lower SET/RESET current than the GST-PRAM (Fig. 2) [3]. Since the perfect device of high-speed and low-cost does not exist, there is a tradeoff between the speed and cost. For example, the longer word-line and bit-line decrease the chip cost but its larger parasitic resistance and capacitance degrade the performance. For the Hybrid M-SCM/NAND flash solid-state drive (*Hybrid SSD*), its performance lies between M-SCM and NAND flash. More M-SCM capacity can accelerate the SSD speed but also introduces more cost. Therefore, in this paper, the design guidelines of the All S-SCM SSD and Hybrid SSD are provided considering the performance/cost tradeoff.

## 2. SSD Architecture and Algorithm

The architectures of the All S-SCM SSD and Hybrid SSD [4] are illustrated in Fig. 3. M-SCM is used as a write cache for the NAND flash to merge the frequently accessed data in the Hybrid SSD. S-SCM is adopted for the All S-SCM SSD due to its high capacity. The SSD controller manages the data storage. The algorithm tables are stored in a DRAM. In the Hybrid SSD, hot or fragmented data are buffered in the M-SCM while cold and sequential data are stored in the NAND flash, as shown in Fig. 4(a) [4]. When the M-SCM is almost full, cold and less fragmented data are evicted to the NAND flash. Since the SCM endurance is limited (~$10^9$), wear leveling is required. In the proposed wear leveling algorithm of Fig. 4(b), an overwrite count $WE_{\text{sector}}$ is calculated for each SCM sector (512 Bytes). When the maximum $WE_{\text{sector}}$ of a page surpasses a threshold $W_{\text{th}}$, the wear leveling procedure is triggered: the old data is read out, merged with the new data and written to a blank SCM page (same page size as the NAND flash), as illustrated in Fig. 5. After that, the old page is recycled as blank and the $W_{\text{th}}$ of this physical page is updated by adding an overwrite count $N_{\text{ow}}$. Fig. 6 demonstrates the developed trace-based simulation platform to evaluate the storage system. The system-level performance, energy consumption and endurance of the memory device are outputted as the evaluation results.

## 3. Results and Discussions

The specifications of the memory devices are listed in Table I. In the Hybrid SSD, the M-SCM capacity is varied from 5% to 25% of the SSD capacity. Workload of the financial server is used for evaluation [5]. Fig. 7 shows the sensitivity analyses of the parameter $N_{\text{ow}}$ for the All S-SCM SSD, whose endurance is the reciprocal of the write/erase (W/E) cycle of the mostly worn sector. As shown in Fig. 7(a), a small $N_{\text{ow}}$ causes degradation of the endurance, because excessively-frequent data migration is harmful to the SCM endurance. Contrarily, a large $N_{\text{ow}}$ leads to an insufficient wear leveling. As a result, write accesses among sectors are not balanced: some sectors are heavily overwritten while others not. From Fig. 7(b), frequent wear leveling procedures degrade the performance of the All SCM SSD. To find the optimum $N_{\text{ow}}$ for balancing performance and endurance, *Function(T,E)* is defined in Fig. 8, in which $T$ and $E$ are the regularized throughput and endurance. The cross point is considered as a balance point between the throughput and endurance. With this optimal $N_{\text{ow}}$, the speed of the All SCM SSD and Hybrid SSD are compared in Fig. 9. All S-SCM SSD with a <5 μs access latency is faster than the Hybrid SSD, when the M-SCM/NAND capacity ratio ($S_{\text{M-SCM/NAND}}$) is 25%. When $S_{\text{M-SCM/NAND}}$ is <10%, All S-SCM SSD is faster than the Hybrid SSD even S-SCM latency is >10 μs. Assuming the cost of each memory device, the total SSD cost is compared in Fig. 10. The Hybrid SSD with a 25% $S_{\text{M-SCM/NAND}}$ is more cost-effective if the S-SCM/NAND cost ratio ($R_{\text{S-SCM/NAND}}$) is >4.

## 4. Conclusions

The All S-SCM SSD and Hybrid SSD are compared in the performance, power, endurance and cost. As summarized in Table II, compared with the Hybrid SSD, All S-SCM SSD is faster when the M-SCM/NAND capacity ratio of the Hybrid SSD ($S_{M\text{-}SCM/NAND}$) is <10%. When $S_{M\text{-}SCM/NAND}$ is 25%, S-SCM latency has to be <5 μs to make the All S-SCM SSD faster than the Hybrid SSD. The SSD cost comparisons are summarized in Table III, assuming the cost ratio of M-SCM/NAND is 12, if the S-SCM/NAND cost ratio is <4, All S-SCM SSD has lower cost than its hybrid counterpart with a 25% $S_{\text{M-SCM/NAND}}$.

**References** [1] R.F. Freitas et al., *IBM Journal of Research and Development, vol. 52, no. 4/5*(2008) 439-447. [2] G.W. Burr et al., *IBM Journal of Research and Development, vol. 52, no. 4/5*(2008) 449-464. [3] K. Johguchi at el., IRPS (2013) MY.5.1-MY.5.4 [4] C. Sun *et al.*, *TCAS-I* (2014) 382-392. [5] http://traces.cs.umass.edu/index.php/Storage/Storage

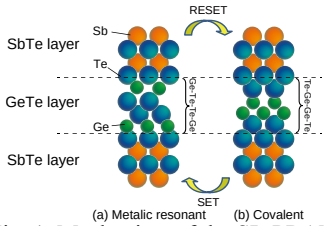


Fig. 1. Mechanism of the SL-PRAM [3].

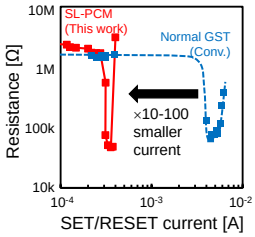


Fig. 2. SET/RESET current of the SL-PRAM and GST-PRAM [3].

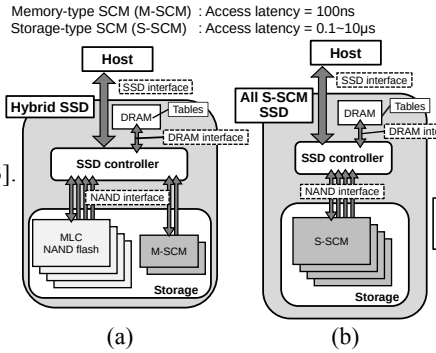


Fig. 3. (a) Hybrid M-SCM/NAND flash SSD architecture [4]. (b) All S-SCM SSD architecture.

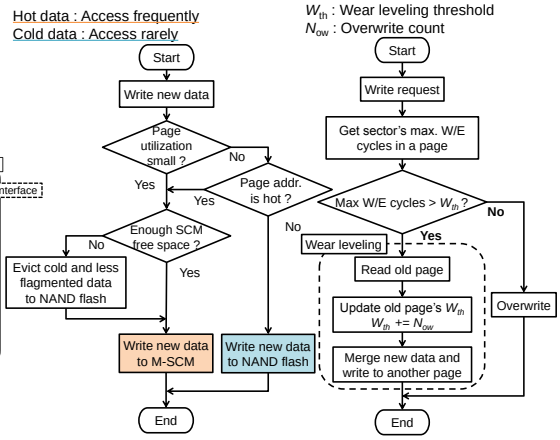


Fig. 4. (a) Conventional Hybrid M-SCM/NAND flash SSD write operation [4]. Fragmented or hot data are written to M-SCM. (b) Proposed SCM wear leveling algorithm.

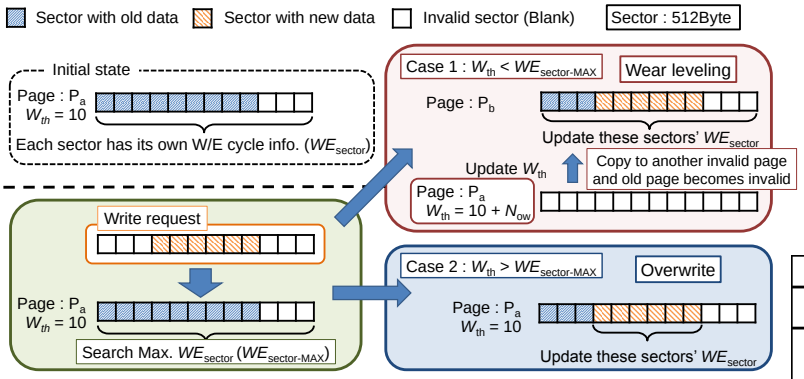


Fig. 5. An example of the wear leveling operation for SCM.

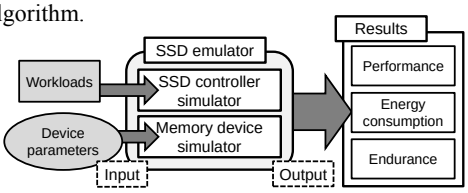


Fig. 6. Simulation platform.

Table I. MLC NAND flash and SCM specifications.

| | MLC NAND flash | M-SCM | S-SCM |
|---|---|---|---|
| Read latency | 85μs / page | 100ns / sector | 100ns ~ 10μs / sector |
| Write latency | Lower page : 400μs / page<br>Upper page : 2800μs / page | 100ns / sector | 100ns ~ 10μs / sector |
| Erase latency | 8500μs / block | N / A | |
| $V_{DD}$ (Core / I/O ) | 3.3V / 1.8V | 1.8V / 1.2V | |
| Access unit | Page ( 16KB ) | Sector ( 512B ) | |

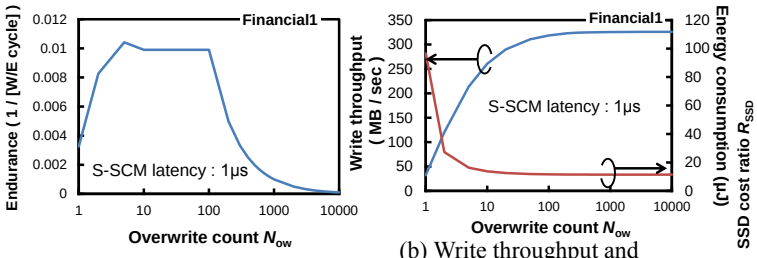


(a) Endurance

(b) Write throughput and Energy consumption

Fig. 7. Sensitivity analysis of the parameter $N_{ow}$ for the All S-SCM SSD.

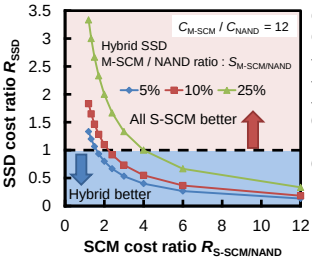


$C_{M-SCM}$ : Cost of M-SCM
$C_{S-SCM}$ : Cost of S-SCM
$C_{NAND}$ : Cost of NAND flash
$S_{M-SCM}$ : Capacity of M-SCM
$S_{S-SCM}$ : Capacity of S-SCM
$S_{NAND}$ : Capacity of NAND flash
Cost(Hybrid) = $S_{M-SCM}$×$C_{M-SCM}$ + $S_{NAND}$×$C_{NAND}$
Cost(SCM) = $S_{S-SCM}$×$C_{S-SCM}$
$R_{SSD}$ : Cost(Hybrid) / Cost(SCM)
$R_{S-SCM/NAND}$ : $C_{S-SCM}$ / $C_{NAND}$
Cost of DRAM is x16 of NAND flash

Fig. 10. Cost analysis of the All S-SCM SSD and the Hybrid M-SCM/NAND flash SSD.

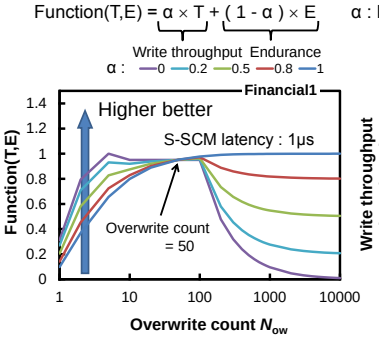


Fig. 8. The optimum $N_{ow}$ to balance throughput and endurance of the All S-SCM SSD.

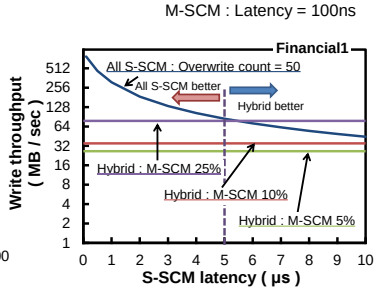


Fig. 9. Write throughput comparisons of the All S-SCM SSD and the Hybrid M-SCM/NAND SSD.

Table II. SSD throughput comparisons summary.

| Hybrid SSD M-SCM/NAND flash capacity ratio | S-SCM latency (μs) (Overwrite count $N_{ow}$ = 50) | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 5% | All S-SCM higher throughput | | | | | | | | | |
| 10% | All S-SCM higher throughput | | | | | | | | | |
| 25% | All S-SCM higher throughput | | | | | Hybrid higher throughput | | | | |

Table III. SSD cost comparisons summary.

| Hybrid SSD M-SCM/NAND flash capacity ratio | $R_{S-SCM/NAND}$ ($C_{M-SCM}$ / $C_{NAND}$ = 12) | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
| 5% | 1.71 | Hybrid low cost | | | | | | | | | | |
| 10% | 2.40 | Hybrid low cost | | | | | | | | | | |
| 25% | All S-SCM low cost | 4.00 | Hybrid low cost | | | | | | | | | |