Extending SSD Lifetime via Balancing Layer Endurance in 3D NAND Flash Memory

Siyi Huang†, Yajuan Du*†, Yi Fan†, Cheng Ji‡
† Wuhan University of Technology, Wuhan, China
‡ Nanjing University of Science and Technology, Nanjing, China
{huangsiyi, dyj, fy2685018622}@whut.edu.cn, cheng.ji@njust.edu.cn
* The corresponding author

***Abstract*—By stacking layers vertically, 3D flash memory enables continuous growth in capacity. In this paper, we study the layer variation in 3D flash blocks and find that bottom layer pages exhibit the lowest endurance, whereas middle layer pages demonstrate the highest endurance. The imbalanced endurance across different layers will diminish the overall SSD lifetime. To address this issue, we introduce a novel layer-aware write strategy, named LA-Write. It performs write-skip operations with layer-specific probabilities. The endurance of bottom layer pages with the highest probability would be noticeably improved, which could balance the layer endurance. Experiment results show that LA-Write can improve SSD lifetime by 29%.**

***Index Terms*—3D NAND flash memory, layer endurance variation, SSD lifetime**

I. Introduction

Flash vendors have been aggressively increasing the number of vertical layers to scale up the capacity of 3D flash memory [1]. However, its cell structure significantly varies along the z-axis, leading to significant layer process variation [2].

To grasp the impact of vertical organization on 3D NAND flash memory process variation, we conducted a process characterization experiment and observed the *Bit Error Rate* (BER) of pages at the bottom layers increased most rapidly when blocks underwent some *Program/Erase* (P/E) cycles, i.e. pages at the bottom layers exhibit lower endurance than those at the middle and top layers. Once the page with the worst endurance wears out, the whole block will be deemed a bad block. The *Flash Translation Layer* (FTL) stops using the bad block, causing under-utilization of other pages with better endurance. Therefore, unbalanced layer endurance within blocks negatively impacts the SSD lifetime.

LA-Write performs write-skip operations with probabilities on pages, which can improve the endurance of these pages by reducing their wear stress from write operations. In LA-Write, distinct probabilities are allocated to the bottom, top, and middle layers respectively, and layers with the lowest endurance (i.e., the bottom layers) are assigned the highest probabilities. Therefore, LA-Write performs most write-skip operations on bottom layers, significantly improving their endurance. Consequently, the distinct probabilities corresponding to different layers help balance endurance across the bottom, top, and middle layers. Experimental results show that LA-Write can improve SSD lifetime by 29% on average.

II. Motivation And Proposed Method

A. Motivation

In investigating BER variation across layers, a preliminary experiment was conducted on a 32-stack 3D MLC flash, each layer comprising upper and lower pages. The results show inter-layer error variation and intra-layer error similarity.

Inter-Layer Error Variation: In Figure 1, curves of various colors depict distinct P/E cycle counts. With increasing P/E cycles (e.g., up to 20K), the BER profile intensifies. Notably, the bottom layers (L2-L15) display markedly higher BER than other layers, while the top layers (L29-L32) exhibit a slightly elevated BER compared to the middle layers.

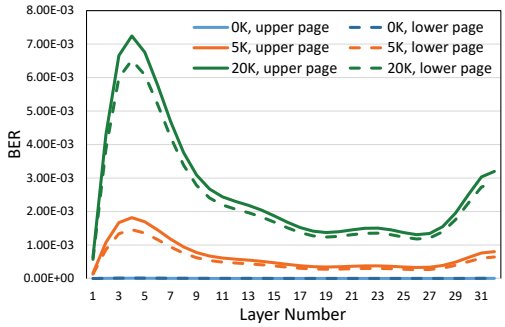

Fig. 1. BER variations in the flash block.

Intra-Layer Error Similarity: In Figure 1, solid and dashed lines represent upper and lower pages in flash layers. The BER of upper and lower pages within the same layer is nearly identical, suggesting uniform endurance across pages in the same layer.

These two observations lead to the conclusion that pages at the bottom layers experience earlier wear-out, indicating worse endurance. However, once any page wears out, the entire block is considered to be bad. Eventually, the SSD would die with no health blocks. Therefore, it is necessary to mitigate the wear of pages at the bottom layers.

B. The Proposed LA-Write Method

As shown in Figure 2, the FTL chooses an active page as the target write location for the incoming write. Traditional write strategies follow the selection of FTL and direct the write to the active page. However, LA-Write first determines

979-8-3503-4859-0/DATE24/© 2024 EDAA

which layer the active page belongs to, and then performs a write-skip operation with the corresponding probability of that layer on the page. The probabilities for the bottom, top, and middle layers are 20%, 10%, and 5%, respectively. If a write-skip operation is carried out on the page, the active page is switched to its next page. Consequently, the original active page remains unwritten, reducing its wear caused by writing. This process is iterated until an active page is found that does not undergo a write-skip operation. Finally, the incoming write is directed to this identified active page.

Thanks to the probability settings, LA-Write performs more write-skip operations on the bottom layers, channeling more writes to the layers with strong endurance (i.e., the top and middle layers), thus mitigating wear on layers with poor endurance (i.e., the bottom layers), extending SSD lifetime.

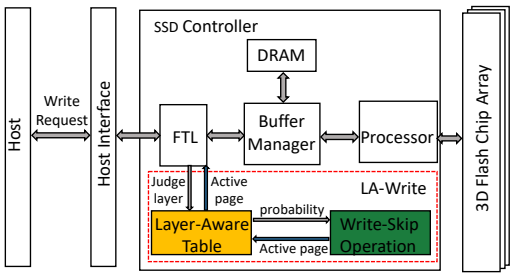

Fig. 2. The architecture overview of LA-Write.

III. Evaluation

A. Experimental Setup

We use Disksim with SSD extensions to simulate a 64GB 3D MLC flash. Five workloads are chosen from MSR. We implement three methods: baseline, BPM and LA-Write. BPM continuously utilizes pages with stronger endurance after the pages with poorer endurance wear out instead of classifying the entire block as a bad block.

B. Lifetime Improvement

In [3], [4], SSD lifetime is measured as the total amount of written data until the SSD failure. Subsequently, we compute the total amount of written data under Baseline, BPM, and LA-Write at the point of SSD failure and normalize it. As depicted in Figure 3, in comparison to Baseline, LA-Write demonstrates the highest average lifetime improvement, reaching 29%. In BPM, even after the page with the lowest endurance wears out, the entire block continues to be used until the number of worn-out pages reaches a certain threshold. As seen in Section II-A, pages at the bottom layers exhibit similar endurance, allowing for the rapid accumulation of several worn-out pages. Consequently, there is no noteworthy delay in the block's lifetime. In LA-Write, we reduce the fast wear of pages with worse endurance through write-skip operations, delaying the block failure, which leads to a significant increase in the total amount of written data.

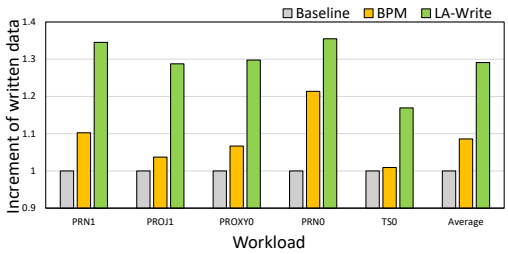

Fig. 3. Increment of written data in SSD.

C. Sensitivity Study

Considering the varying degrees of endurance differences observed among layers in Section II-A, we assigned specific probability proportions to the bottom, top, and middle layers. In Figure 4, LA-Write (20,10,5) indicates that the probabilities of the bottom, top, and middle layers are 20%, 10%, and 5%, respectively. Among these three probability configurations, LA-Write (20,10,5) exhibits an 11% increase in total amount of written data compared to LA-Write (20,15,10). However, in comparison to LA-Write (20,10,5), LA-Write (30,15,10) experiences a 1% reduction in total amount of written data. This is attributed to an increased probability of write-skip operations, leading to a reduction in the number of available pages for writing. Consequently, it is necessary to increase the frequency of garbage collection operations to meet the current workload, which in turn consumes the SSD lifetime.

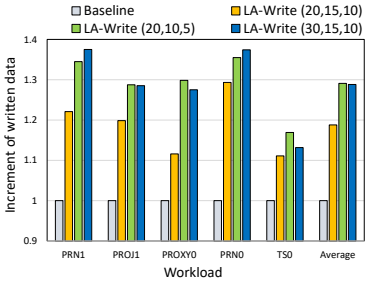

Fig. 4. The different probability groups in LA-Write

References

[1] C. Gao, L. Shi, C. Ji, Y. Di, K. Wu, C. J. Xue, and E. H.-M. Sha, "Exploiting parallelism for access conflict minimization in flash-based solid state drives," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 1, pp. 168–181, 2018.

[2] Y. Luo, S. Ghose, Y. Cai, E. F. Haratsch, and O. Mutlu, "Improving 3D NAND flash memory lifetime by tolerating early retention loss and process variation," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 2, no. 3, pp. 1–48, 2018.

[3] W. Debao, Q. Liyan, Z. Peng, and P. Xiyuan, "Bpm: A bad page management strategy for the lifetime extension of flash memory," in *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 57199. American Society of Mechanical Engineers, 2015, p. V009T07A010.

[4] D. Hong, M. Kim, G. Cho, D. Lee, and J. Kim, "GuardedErase: Extending SSD lifetimes by protecting weak wordlines," in *20th USENIX Conference on File and Storage Technologies (FAST 22)*, 2022, pp. 133–146.