

RealWear: Improving Performance and Lifetime of SSDs Using a NAND Aging Marker

Myungsuk Kim, Myoungjun Chun, Duwon Hong, Yoona Kim, Geonhee Cho, Dusol Lee, Jihong Kim

Department of Computer Science and Engineering, Seoul National University

{morssola75, mjchun, duwon.hong, yoonakim, ghcho, dslee, jihong}@davinci.snu.ac.kr

ABSTRACT

NAND flash memory has revolutionized how we manage data in modern digital systems, significant improvements are needed in flash-based storage systems to meet the requirements of emerging data-intensive applications. In this paper, we address the problem of NAND aging markers that represent the wearing degree of NAND cells. Since all flash operations are affected by the wearing status of NAND cells, an accurate NAND aging marker is critical to develop flash optimization techniques. From our evaluation study, we first show that the existing P/E cycle-based aging marker (PeWear) is inadequate to estimate the actual aging status of NAND blocks, thus losing opportunities for further optimizations. To overcome the limitations of PeWear, we propose a new NAND aging marker, RealWear, based on extensive characterization studies using real 3D TLC flash chips. By considering multiple variables that can affect the NAND cell wear, RealWear can accurately indicate the actual wear status of NAND blocks during run time. Using three case studies, we demonstrate that RealWear is effective in enhancing the lifetime and performance of a flash storage system. Our experimental results showed that RealWear can extend the lifetime of individual NAND blocks by 63% and can reduce the GC overhead by 21%. Furthermore, RealWear significantly mitigates read latency fluctuations, guaranteeing that the read latency can be bounded with at most 2 read retry operations.

Keywords

NAND aging marker, NAND cell wear, NAND flash memory, flash storage systems, storage reliability, storage performance

1. INTRODUCTION

NAND flash memory, which has been one of the key enablers of the modern digital revolution, played an important role in realizing various innovative digital products such as digital cameras, smart-phones and ultra-high-capacity SSDs. With on-going data-driven innovations, the role of NAND flash memory is expected to be even more critical in emerging new storage market areas (such as real-time analytics, machine learning, and self-driving cars) where a large

amount data should be collected and processed in a timely fashion. These new applications bring new challenges that NAND flash memory should guarantee high performance and lifetime even under the worst operating conditions.

Although there have been various flash memory solutions, they are ill-suited to meet these requirements except for the high-capacity requirement. For instance, the lifetime of the TLC flash memory, whose capacity is three times that of SLC flash memory, is only 1% of the SLC flash memory [1]. Furthermore, large variability in read latency makes it impossible to support *time-bounded* read operations that are essential for real-time systems. Unfortunately, as the flash memory technology further advances (e.g., QLC flash memory), both the performance and lifetime characteristics of the flash memory are expected to get even worse.

To meet the challenging new flash requirements of emerging data-intensive apps, therefore, significant improvements are needed across the key design abstractions of a flash storage system. In this paper, we address the key problem at the NAND flash layer. In particular, we focus on developing an accurate and practical NAND wear index (or marker) that represents the wearing degree of NAND cells. Since the inner working of all the flash operations are tightly connected to how much NAND cells were worn out, knowing the exact wear status of NAND cells is an essential prerequisite to develop efficient flash optimization techniques.

The most common wear indicator is to count the chronological age of a NAND cell based on the number of program/erase (P/E) cycles (i.e., the number of program and erase operations the NAND cell has experienced), in a similar fashion as the chronological age of a human being. Since a high electrical voltage ($> 20V$) is known to cause the wear of a NAND cell, the number of P/E cycles has been regarded as an effective and practical proxy indicating the wear status of NAND cells. Therefore, many flash optimization techniques have been designed using the number of P/E cycles as a NAND aging marker although its prediction accuracy as a NAND aging marker was not fully understood or its impact on the flash optimization techniques were not clear. The main goal of this paper is to investigate NAND aging markers in a comprehensive fashion so that highly efficient flash optimization techniques can be designed using an accurate NAND aging marker.

2. METHODOLOGY AND RESULTS

In order to develop a new NAND aging marker, we follow an overall procedure summarized in Figure 1. Before we

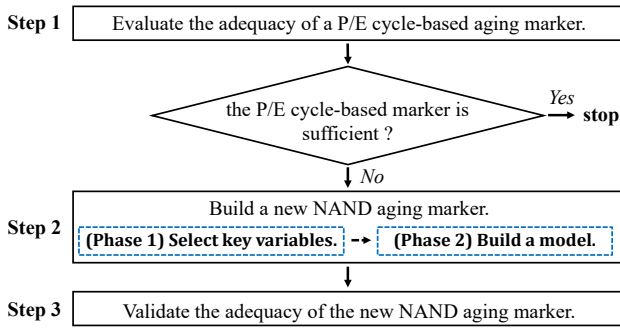


Figure 1: An overall procedure of our investigation study.

develop a new NAND aging marker, to verify the adequacy of a P/E cycle-based aging marker, we evaluated how the wear status of NAND blocks and their P/E cycles are related. From extensive evaluation results using real 48-layer 3D TLC flash chips, we found that the P/E cycle-based aging marker is not a reliable wear indicator. Intuitively, the role of P/E cycle-based aging marker is similar to that of the chronological age of a human being. For example, individuals with the same chronological age could have widely different biological ages due to their genetic differences such as the length of telomere [2]. Even twins with the same genetic characteristics may have different biological ages because of differences in their life styles and living environments. Similarly, when two NAND blocks experience the same number of P/E cycles, their wearing degree could be significantly different. This difference in aging characteristics between NAND blocks is caused by various factors such as (1) process variations occurred while they are manufactured, (2) I/O workload variations in different NAND blocks, and (3) variations in their operating environment. Although these factors affect the wearing degree of the NAND block, the P/E cycle-based aging marker cannot reflect these factors at all. Therefore, there is a clear need for a more reliable NAND aging marker which can indicate the actual wear status of NAND blocks by considering key variation factors.

To derive an accurate NAND aging model that meets the above requirements, we have comprehensively investigated the effect of key factors affecting the NAND aging process using real 160 3D TLC flash chips. Based on our extensive experimental data, we developed a novel NAND aging marker, called *RealWear*, which can indicate the different aging status of an individual NAND block in a much more accurate fashion over the existing P/E cycle-based aging marker. We followed a general model construction approach which consists of two key phases: *variable selection* and *model building*. In the variable selection phase, we collected wear-related variables and check if they are sufficiently correlated with the wear status of the NAND block. In the model building phase, we constructed a NAND aging model for deriving *RealWear* using regression analysis of the selected variables.

We validated the adequacy of *RealWear* as a NAND aging predictor by comparing the predicted wear status of NAND blocks with that from actual measurements while varying operating conditions. We found that *RealWear* can effectively represent the different aging characteristics among NAND blocks and eventually prevent NAND blocks from being wasted. Unlike the P/E cycle-based aging marker, NAND blocks with the same *RealWear* values exhibited al-

most the same number of bit errors. In addition, *RealWear* can properly reflect the impact of I/O workload variation and temperature variation on the wear of NAND blocks.

To demonstrate the benefit of *RealWear* in flash storage systems, we present three case studies, *LongLive*, *FastCopy*, and *BoundedRead*. Our experimental results show that *LongLive* can extend the lifetime of individual NAND blocks by an average of 63% and *FastCopy* can improve the performance up to 21% by reducing GC overhead. *BoundedRead*, which reduces read latency fluctuations, achieves *time-bounded* reads for a flash storage system. Our evaluation shows that no read can experience more than two read retries under all possible operating conditions, thus significantly reducing the tail latency of the flash storage system.

3. CONTRIBUTION AND SUMMARY

This paper makes the following key contributions:

- We show that the existing P/E cycle-based NAND aging marker is not adequate to satisfy the new flash requirements for data-intensive emerging applications. Based on the comprehensive device-level measurements using real state-of-the-art 3D TLC NAND flash chips, we quantitatively show the key weak points of the P/E cycle-based aging marker.
- By using multiple variables, we present a new novel NAND aging marker, *RealWear*, which accurately represents the real wear status of NAND cells. To the best of our knowledge, *RealWear* is the most accurate NAND aging predictor among openly known NAND aging markers.
- We demonstrate that *RealWear* is effective in extending the lifetime and performance of a flash storage system. Compared with the P/E cycle-based aging marker, *LongLive* can extend the lifetime of a flash storage system and *FastCopy* can reduce the copy overhead of GC.
- We show that the worst-case flash read latency can be *bounded* using *RealWear*. To the best of our knowledge, *BoundedRead* is the first technique that supports the bounded read latency.

By considering multiple variables, *RealWear* outperforms the P/E cycle-based aging marker by distinguishing the wear status of NAND blocks in a more accurate fashion for a longer lifespan of a NAND block. Furthermore, *RealWear* can open new opportunities to further optimize the flash management techniques.

Acknowledgments

This work was supported by Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-IT1701-11. The ICT at Seoul National University provided research facilities for this study. (Corresponding Author: Jihong Kim)

4. REFERENCES

- [1] M. Kim, Y. Song, M. Jung, and J. Kim. SARO: A State-Aware Reliability Optimization Technique for High Density NAND Flash Memory. In *Proceedings of the ACM Great Lakes Symposium on VLSI (GLSVLSI)*, 2018.
- [2] J. Szostak and E. Blackburn. Cloning yeast telomeres on linear plasmid vectors. *Cell*, 29(1):245–255, 1982.