A Novel Simulation Methodology for Accelerating Reliability Assessment of SSDs

Luyao Jiang
Department of Computer Science,
University of Virginia
lj5bp@virginia.edu

Sudhanva Gurumurthi
AMD Research, Advanced Micro Devices, Inc.
Department of Computer Science, University of Virginia
sudhanva.gurumurthi@amd.com

***Abstract*—Reliability is an important factor to consider when designing and deploying SSDs in storage systems. Both the endurance and the retention time of flash memory are affected by the history of low-level stress and recovery patterns in flash cells, which are determined by the workload characteristics, the time during which the workload utilizes the SSD, and the FTL algorithms. Accurately assessing SSD reliability requires simulating several years' of workload behavior, which is time-consuming. This paper presents a methodology that uses snapshot-based sampling and clustering techniques to help reduce the simulation time while maintaining high accuracy. The methodology leverages the key insight that most of the large changes in retention time occur early in the lifetime of the SSD, whereas most of the simulation time is spent in its later stages. This allows simulation acceleration to focus on the later stages without significant loss of accuracy. We show that our approach provides an average speed-up of 12X relative to detailed simulation with an error of 3.21% in the estimated mean and 6.42% in the estimated standard deviation of the retention times of the blocks in the SSD.**

***Keywords*-Storage, SSD, Flash, Reliability**

I. Introduction

Flash memory based solid-state drives (SSD) have gained tremendous popularity in recent years. SSDs are widely used in a variety of computing devices, from phones and tablets to desktops and servers. SSDs offer several advantages over hard disk drives (HDDs), including higher performance, lower power, improved acoustics, and ruggedness. Despite these positives, a major concern with SSDs is that the underlying flash memory technology has a limited lifetime, which affects both the number of writes that can be reliably done to flash, referred to as endurance and quantified in terms of program/erase (P/E) cycles, and the retention time of stored data. The retention time is the period of time during which data written to a flash memory cell can be read reliably [1]. The retention time decreases with the number of P/E cycles. These reliability concerns are paramount in data centers, where workloads are I/O-intense, data integrity requirements are high, and SSD replacement costs can be significant [2].

Accurate estimation of SSD reliability is important in data centers. Such accurate estimation cannot be obtained by merely counting the number of P/E cycles and instead requires capturing the pattern of stress and recovery to flash memory cells over time under the influence of a workload [1]. Also, while SSD performance can be evaluated using relatively short workload traces, accurate reliability assessments require capturing the stress-recovery patterns over an extended period of time, typically years, since endurance and retention time depends on the *history* of stresses and recovery over time. Existing reliability estimation approaches in the literature either merely count the number of P/E cycles [3][4], simulate the workload for only a very short interval of time, or use simplistic extrapolation of the results from a short simulation to a longer duration [5][1]. Because reliability is affected by the interplay of the workload behavior, the flash-translation layer (FTL) algorithms for page mapping, wear-leveling, garbage collection, and the distribution of stress and recovery events, any simplistic extrapolation of a short-duration simulation over a longer timescale is inherently error-prone. On the other hand, simulating a workload over a long time-scale is time-consuming. Having the means to perform reliability assessments of SSDs for different workloads with low turnaround time can help reduce the costs of capacity planning for deploying the storage systems.

Table I: Simulation time of enterprise workloads for 5 years

Workload	MSNFS	EXCHANGE	DAPPS	RADIUS
Simulation Time (hours)	185.93	127.51	86.52	49.31

Table I shows the simulation time of several data center workload traces [6] during a 5-year timescale in the DiskSim simulator [7] with the SSD extension [3]. These simulations were run on an 8-CPU quad-core 2.3 GHz Intel Xeon™ machine with 48 GB of RAM. We choose a 5-year timescale because it represents the typical hardware refresh interval in a data center [8]. We simulate 5 years worth of activity by repeatedly replaying the I/O traces, each of which capture one representative day's I/O traffic. As the table indicates, detailed simulation of years of workloads is very expensive in terms of time. In this paper, we present an acceleration framework to accurately measure SSD reliability by performing sampling over time on carefully trimmed workloads. Our methodology is generic and can be used for any workload or SSD architecture, thereby allowing for flexible design-space exploration studies.

In the rest of this paper, Section II describes the experimental setup and Section III characterizes the variation in the retention time during a multi-year timescale and motivates the sampling-based acceleration approach. Section IV describes our simulation acceleration methodology and Section V presents an evaluation of the accuracy and speedup attained using our approach. Section VI concludes the paper.

II. Experimental Setup

We carried out our simulations using Disksim[7] with the SSD extension module developed by Microsoft Research[3]. We modified Disksim to record statistics that impact reliability, such as the recovery time between successive stresses to flash, and augmented it with reliability models to calculate the retention time [1][8]. We simulate an enterprise-class 64GB 2-bit MLC SSD. The FTL of this SSD uses a greedy-based garbage collection approach with wear-leveling aware heuristics and cold data migration to distribute stress events evenly across all blocks [3]. The FTL uses a hybrid page-mapping policy that combines static and dynamic mapping mechanisms, with an allocation pool that operates at the granularity of a flash element. *Although we use Disksim to model a specific SSD in this paper, we believe our statistical acceleration methodology is generic enough to be applied to other FTLs, SSD organizations, and potentially other simulators.*

To the best of our knowledge, there are no recent publicly available workload traces that span multiple years of activity. *Cello* [9] is the longest trace that we know of; it spans one year. However, it is relatively old and its I/O activity might not be representative of modern data center workloads. The workloads we choose in this paper consist of four enterprise-class block I/O traces captured from modern Microsoft's data centers [6]. Because each trace spans only between 6 hours and 1 day of representative I/O activity, we repeatedly replay each I/O trace until the simulation time reaches 5 years, which aligns with the typical hardware refresh period in a data center. We believe that this approach is reasonable because prior research has shown that disk I/O traffic exhibits spatial and temporal self-similarity characteristics [10][6].

III. The Opportunity for Simulation Acceleration

The retention time of the flash memory cells is affected by the pattern of stresses and recovery periods over time. These patterns are governed by the workload, page-mapping, wear-leveling, and cleaning operations in the FTL. Disksim simulates all these characteristics. In the simulator, we track the retention times at the granularity of individual flash memory blocks. We simulate each of our workloads using the methodology described in Section II for a 5-year period and take snapshots of the entire system state, including all metrics related to retention time, every 15 days of simulated time. The histograms of the retention times over each 15-day period for one representative workload is shown in Figure 1. (The trends are similar for the other workloads). Each curve in the graph shows the histogram of the number of blocks in the SSD with a given retention time value at each 15-day interval. The x-axis in the graph is retention time. As noted in Section I, retention time decreases as P/E cycles increase, so the curves on the right represent retention time histograms early in the lifetime of an SSD.

The retention times of the blocks decrease rapidly early in the lifetime (denoted as *Phase I* in the figures) and then the retention times decrease at a much slower rate afterward (designated as *Phase II*). The *Phase I* and *Phase II* markers in Figure 1 are not the actual duration of these phases; rather, they show when the retention time distributions show a marked change in trends. The key reason for this trend is

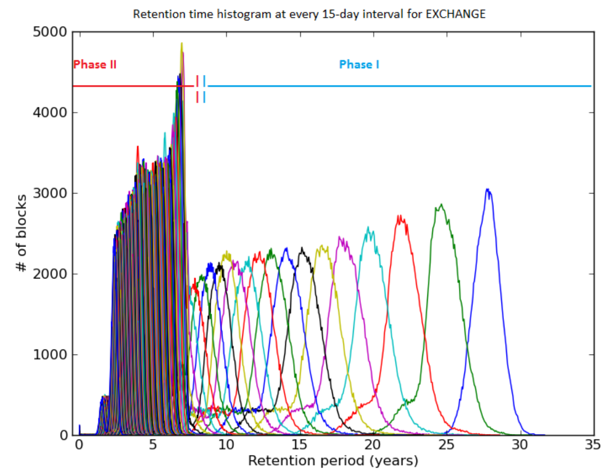

Figure 1: Retention time histograms over 5 years of simulated time. Curves on the right represent retention time histograms early in the lifetime of an SSD.

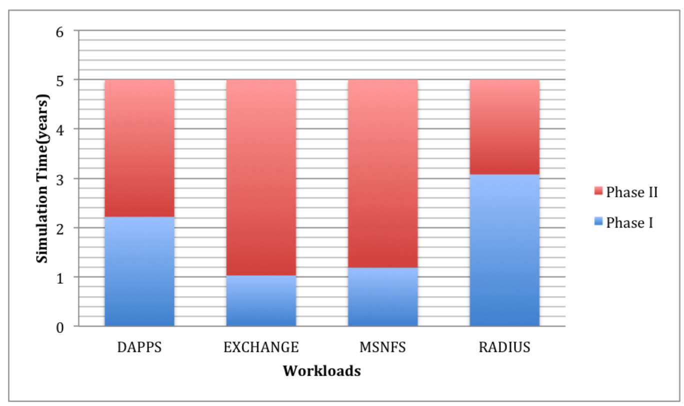

Figure 2: Ratio of simulation time between Phase I and Phase II of the usage period of the SSD.

that $\delta V_{th,s}$ has a power-law relationship with the number of P/E cycles [1]:

$$\delta V_{th,s} = \frac{(A.cycle^{0.62} + B.cycle^{0.3}).q}{C_{ox}} \quad (1)$$

where A and B are constants, $cycle$ is the number of P/E cycles, q is the charge of an electron, and C_{ox} is the oxide capacitance. Therefore, the rate of change of the retention time decreases with cycling. The ratio of the total simulated time in the two phases is shown in Figure 2.

The graph highlights two key trends that directly influence the ability to accelerate the simulation:

- It is important to simulate in detail the activity during Phase I. Otherwise, even small inaccuracies in the reliability estimation potentially can accumulate into large errors in the result.
- The bulk of the simulation time is spent in Phase II, where potential opportunity exists to skip regions of the simulation and extrapolate the behavior without resulting in large errors.

We use these two observations to develop a sampling-based approach to accelerate the simulation. This approach consists of two simulation modes: a detailed mode and a fast-forward functional mode in which the change in the reliability metrics are approximated. Detailed simulation is performed only during certain intervals, which we refer to as sampling units, whereas functional simulation is performed between the sampling units.

IV. Acceleration Methodology

As mentioned in Section III, our sampling-based simulation approach consists of two modes: a detailed mode and a fast-forward functional mode in which the change in the reliability metrics are approximated. To obtain high accuracy, sampling is performed only in Phase II of a simulation, where the changes in the retention time stabilize. To decrease simulation time further, we augment our systematic sampling-based approach with a snapshot-based clustering technique to reduce the size of the workload so that only a subset of requests are simulated during the detailed simulation mode. We will show that our trimmed workload is representative of the original workload in terms of its stress behavior and impact on SSD reliability.

The overall simulation acceleration framework is shown in Figure 3. The framework consists of three major components: simulator, workload trimmer and snapshots analyzer. The simulator, integrated with reliability models, performs simulation in the detailed and fast-forward modes. It takes workloads as inputs and periodically dumps snapshots to track the reliability-related characteristics. The simulation flow is also depicted in Figure 3. The simulation begins with detailed mode on the full workload (the dotted region) with no acceleration techniques applied. The stress behavior collector in the workload trimmer collects information about

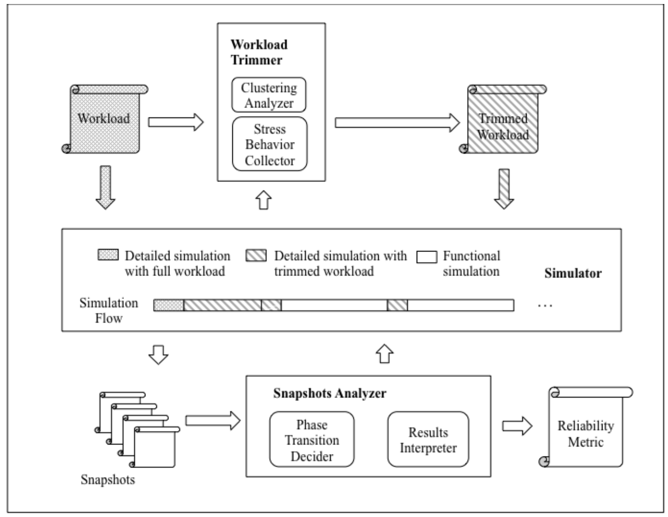

Figure 3: Overview of the acceleration framework

the stress patterns of the full workload as the simulation runs. As soon as sufficient information is collected, the clustering analyzer in the workload trimmer is triggered to perform a K-means clustering-based analysis on the information collected and reduce the size of the workload by selecting representative requests from the full workload. The rest of the simulation takes trimmed workload as inputs. During simulation, the simulator consults the phase-transition decider inside the snapshots analyzer periodically to see whether the simulation has entered Phase II and whether sampling can be performed. The phase-transition decider uses earth mover's distance (EMD) as a metric to quantify the distance between retention time histograms and an EMD threshold to decide whether the simulation has transitioned from Phase I to Phase II. The snapshots analyzer is also responsible for interpreting the snapshots to report the reliability metrics. Except for the beginning of the simulation (dotted region), the bulk of the simulation flow alternates between detailed simulation on the trimmed workload (striped region) and the fast-forward functional simulation (blank region), which speeds up the original simulation.

V. Evaluation of the Acceleration Framework

In this section, we evaluate the accuracy of our acceleration techniques and the speed-up achieved in the simulation time. We evaluate accuracy by comparing the histograms of the data retention times of the all the blocks in the SSD after a simulation period of 5 years for the accelerated variants to the base detailed simulation runs. The only exception is the EXCHANGE workload, which we simulate for only 3.5 years because this workload experiences block retention failures that exceed the capacity of the over-provisioning space of the SSD with the FTL algorithm we exploit. As mentioned earlier, we simulate the multi-year timescale by repeatedly playing back the workload trace. We call each such repetition a simulation round.

Figure 4 compares the retention time distribution across blocks after 5 years of simulated time among the detailed-simulation mode, sampling-simulation mode on full workloads, and sampling-simulation mode trimmed workloads. We compare these approaches to a simplistic extrapolation of retention time after the detailed simulation phase (the approach used in [1]). Our simulation framework generates estimates much closer to detailed simulation compared to the simplistic extrapolation in terms of the position, shape, and height of the retention time distribution. The histograms of the accelerated versions are very similar to the detailed simulation versions for DAPPS, EXCHANGE, and RADIUS. On average, our acceleration framework achieves a mean estimation error of 3.21% and a standard deviation estimation error of 6.42%.

Figure 5 shows the speed-up achieved using the sampling simulation mode on the full workload and trimmed workload with respect to the detailed simulation mode. Significant acceleration is achieved for all workloads except RADIUS. RADIUS is not write-intensive and does not incur as large a simulation time as the other workloads. With our sampling framework on the trimmed workload, an average of 12X speed-up is achieved.

VI. Conclusion

Accurate estimation of SSD reliability is important for data centers. However, accurately measuring the reliability of SSD requires capturing the impact of workload access patterns on flash memory over timescales that span several years, which requires long simulation times. Simplistic extrapolation of reliability from short simulations is inherently error-prone and can lead to incorrect conclusions. To address this problem, we present a framework to accelerate SSD reliability simulation while still providing high accuracy.

VII. Acknowlegements

We thank Sriram Sankar and Vidyabhushan Mohan for their valuable inputs. This research has been supported by NSF CAREER Award 0643925 and gifts from Google.

References

[1] V. Mohan, T. Siddiqua, S. Gurumurthi, and M. Stan, "How i learned to stop worrying and love flash endurance," in *Proceedings of the 2nd USENIX Conference on Hot Topics in Storage and File Systems*. USENIX Association, 2010, pp. 3–3.

[2] L. Barroso, "Warehouse-scale computing," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*. ACM, 2010, p. 2.

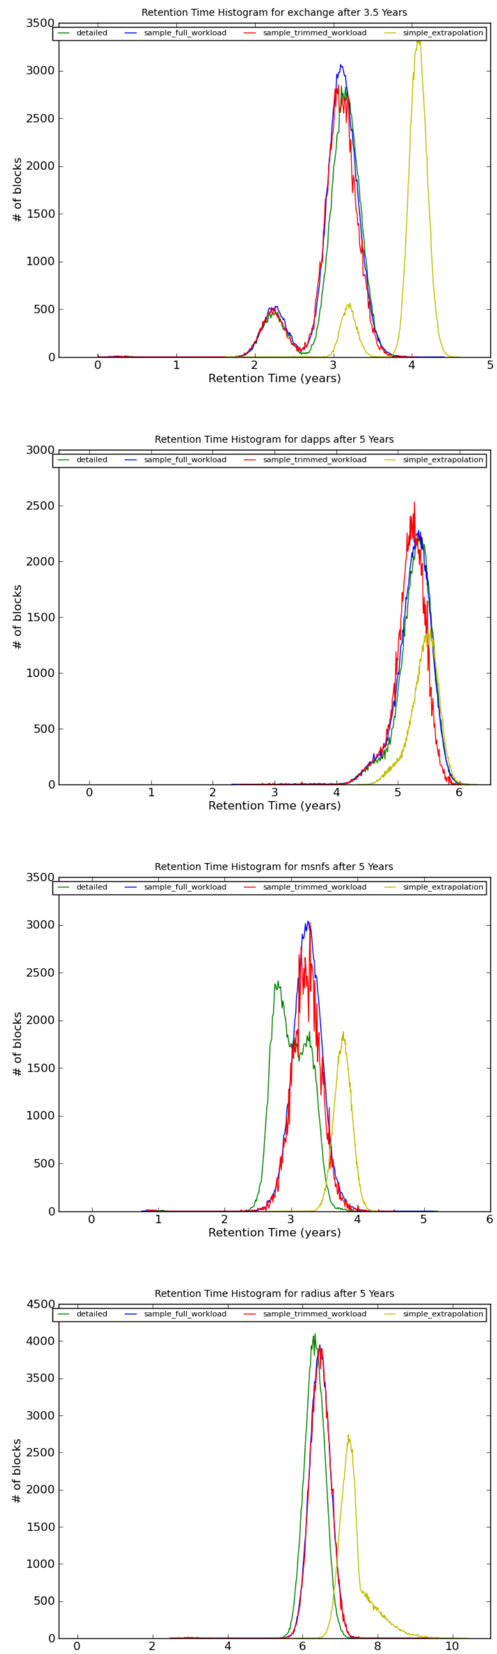

Figure 4: Comparison of retention time histograms between different acceleration modes.

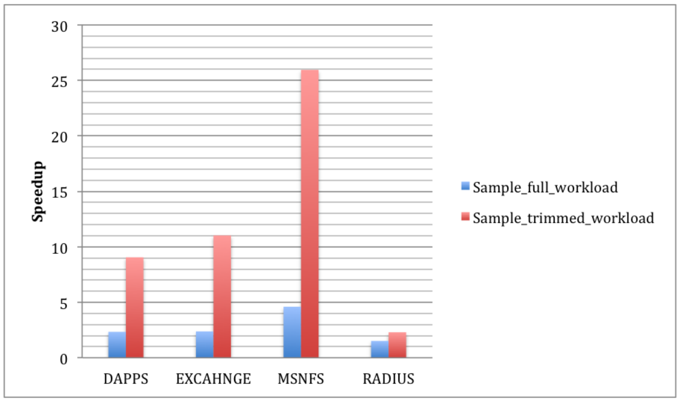

Figure 5: Speed-up achieved using different acceleration modes.

[3] N. Agrawal, V. Prabhakaran, T. Wobber, J. Davis, M. Manasse, and R. Panigrahy, "Design tradeoffs for ssd performance," in *USENIX 2008 Annual Technical Conference on Annual Technical Conference*, 2008, pp. 57–70.

[4] Y. Kim, B. Tauras, A. Gupta, and B. Urgaonkar, "Flashsim: a simulator for nand flash-based solid-state drives," in *Advances in System Simulation, 2009. SIMUL'09. First International Conference on*. IEEE, 2009, pp. 125–131.

[5] Y. Cai, G. Yalcin, O. Mutlu, E. Haratsch, A. Cristal, O. Unsal, and K. Mai, "Flash correct-and-refresh: retention-aware error management for increased flash memory lifetime."

[6] S. Kavalanekar, B. Worthington, Q. Zhang, and V. Sharda, "Characterization of storage workload traces from production windows servers," in *Workload Characterization, 2008. IISWC 2008. IEEE International Symposium on*. IEEE, 2008, pp. 119–128.

[7] J. Bucy, J. Schindler, S. Schlosser, and G. Ganger, "The disksim simulation environment version 4.0 reference manual (cmu-pdl-08-101)," *Parallel Data Laboratory*, p. 26, 2008.

[8] V. Mohan, S. Sankar, and S. Gurumurthi, "Refresh ssds: enabling high endurance, low cost flash in datacenters," University of Virginia, Tech. Rep. CS-2012-05, 2012.

[9] C. Ruemmler and J. Wilkes, *UNIX disk access patterns*. Hewlett-Packard Laboratories, 1992.

[10] M. Gomez and V. Santonja, "Analysis of self-similarity in i/o workload using structural modeling," in *Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 1999. Proceedings. 7th International Symposium on*. IEEE, 1999, pp. 234–242.