Practical Implication of Analytical Models for SSD Write Amplification

Changhyun Park Seongjin Lee* Youjip Won Soohan Ahn**
Hanyang University, *Gyeongsang National University, **The University of Seoul
{pch1984|insight|yjwon}@hanyang.ac.kr, sahn@uos.ac.kr

ABSTRACT

A number of analytical models have been proposed to estimate the write amplification of the Flash storage to obtain the expected lifespan. This work is dedicated to examining the practical implication of the four existing analytical models for estimating the write amplification: Coupon Collector, Uniform Distribution, Expected Value and Markov model. Since the models assume uniform random workload in full utilization of an SSD to predict write amplification, they are not applicable in predicting write amplification in general workload. Moreover, the existing models have not been verified with the real SSD. In this work, we compare the write amplification of the models with that of a real SSD. When we use 0.147 as the overprovisioning factor of an SSD while running uniform random workload, the write amplification of Uniform Distribution, Expected Value, Markov model is 3.90, 4.08, and 4.08, respectively. However, write amplification of the real SSD shows 1.19, which is very different from that of the prediction models. Through experiment, we found that write amplification is closely related to the value of overprovisioning factor. To improve the accuracy of existing prediction models, we update the overprovisioning factor to take account of the ratio of a hot file and the utilization of the storage. We also find that by setting the overprovisioning factor to 1.15, we can obtain write amplification of 1.2 which is close to the write amplification of general workload in a real SSD.

Keywords

Solid State Drive, write amplification, overprovisioning

1. INTRODUCTIONS

The characteristics of SSDs such as low power and high performance makes the device attractive for many I/O intensive applications and workloads. However, the devices suffer from write amplification because of the intrinsic nature of the Flash memory. Write amplification pulls down the performance and reduces the life time of the devices. Thus, it is important to analyze the write amplification of an SSD for increasing its efficiency.

An append-only nature of the Flash-based storage entails that update of a Flash page invalidates the existing one. To accommodate the new updates, an SSD occasionally needs to consolidate the valid blocks and resets one or more blocks. A process to reclaim and consolidate such invalidated pages is called garbage collection (GC) [4, 5]. In the process, the garbage collection module first selects a block whose valid pages are to be migrated to a new location and which is to be erased. The selected block is a victim block. Selecting the best victim block is the crux of a successful garbage collection. A number of victim selection policies have been proposed, e.g. greedy [5], FIFO(LRU), d-choice[16], random. To effectively utilize all the Flash blocks in the storage, it is important that the number of erase/write (E/W) cycles are evenly distributed over all Flash blocks in the storage. Flash firmware levels the wear via selecting victims with smaller E/W cycle in garbage collection, *static wear-leveling* [4, 15, 6] or via occasionally switching the pages in two blocks whose difference in E/W cycle reaches a predefined threshold, *dynamic wear-leveling* [4, 7].

The write amplification describes undesirable behavior of an SSD caused by garbage collection and wear-leveling, which generates more volume of writes than the volume requested to the device. It is important to understand the effect of the write amplification of an SSD because the undesirable program operation performed on the SSD reduces the life time of the SSD. This has been the motivation for a number of works to model the write amplification of an SSD. In this work, we analyze and compare four models that are based on greedy garbage collection policy, then measure the accuracy of the models and compare them with the actual write amplification of the real SSD. Finally, we point out the limits of existing works and deduce implication, with which we address future works on the field.

*This work has been done while the author was Post-Doctoral Fellow at Hanyang university.

ICPE'17, April 22-26, 2017, L'Aquila, Italy
© 2017 ACM. ISBN 978-1-4503-4404-3/17/04...$15.00
DOI: http://dx.doi.org/10.1145/3030207.3030219

2. WRITE AMPLIFICATION

Let the number of pages in a block be N. In append-only storage, the number of IOs issued to the storage from the host and the actual number of IOs occurred in the storage does not coincide because there is the garbage collection and the wear-leveling operation in the Flash storage. The actual number of IO operations can be higher than user IO requests. We call it as Write Amplification, which, denoted by

Name	Description
A	Write Amplification
N	Given number of pages in a block
N_{vp} N_{ip}	Random variable representing the number of valid/invalid pages in a victim block
$\bar{N}_{vp}$ $\bar{N}_{ip}$	Expectations of N_{vp}/ N_{ip}
p_k	Probability that the victim blocks has k valid pages
T	Total number of physical blocks in an SSD
U	Number of available blocks to a user
ρ	Overprovisioning factor
s	Window size of windowed greedy GC
r	Number of free blocks to trigger greedy GC
p^*	Probability that a single page is invalidated
x^M	Minimum state number for Markov model

Table 1: Variables

A, and define as the average number of actual page programs for a single page program request from the host.

The number of the valid pages in a victim block is denoted as N_{vp}, and the number of invalid pages in a victim block is denoted as $N_{ip} = N - N_{vp}$. The garbage collection copies N_{vp} pages to a free block. The Flash storage becomes to have N_{ip} free pages. Now, we introduce existing mathematical models on the write amplification in [11, 12, 3, 14, 9]. Write amplification is described in Eq. (1), where $\bar{N}_{vp}/\bar{N}_{ip}$ denotes the additional number of programs performed in copying the valid pages.

$$A = \frac{\bar{N}_{vp} + \bar{N}_{ip}}{\bar{N}_{ip}} = \frac{N}{\bar{N}_{ip}} = 1 + \frac{\bar{N}_{vp}}{\bar{N}_{ip}} \tag{1}$$

Here, $\bar{N}_{vp}/\bar{N}_{ip}$ is referred to as the write amplification factor [11, 12], denoted by A_f. Letting T and U represent the total number of available blocks in an SSD and the available blocks visible from the host respectively, the overprovisioning factor, ρ, is defined as follows [3, 14]:

$$\rho = \frac{T - U}{U} \tag{2}$$

3. WRITE AMPLIFICATION ANALYSIS MODELS

We examine four analytical models: Coupon Collector [11], Uniform Distribution [3], Expected Value [14], and Markov [9] model. These are named based upon how they estimate the number of valid pages in the victim block.

Note that we use ρ to unify overprovisioning factor and spare factor used in all four models. The summary of variables used in this paper are described in Table 1.

3.1 Coupon Collector Model

To describe and identify the work done by Hu et al. [11], we named the work as Coupon Collector model. This model uses windowed greedy garbage collection policy with window size of s. The garbage collection process selects a victim block that has the least number of valid pages within $0 \sim (s-1)$ blocks in the occupied block queue.

The basic idea of Coupon Collector model is simple and straight-forward. Let N_{vp}^{cc} and p_k^{cc} be the number of valid pages in the victim block and the probability that it has k valid pages, respectively. In Coupon Collector model, the expected number of valid pages, denoted by $\bar{N}_{vp}^{cc}$, is represented as $\bar{N}_{vp}^{cc} = \sum_{k=0}^{N} k p_k^{cc}$. For short, the write amplification in Coupon Collector model, denoted by A^{cc} can be computed as in Eq. 3.

$$A^{cc} = \frac{N}{N - \sum_{k=0}^{N} k p_k^{cc}} \tag{3}$$

The key ingredient of Coupon Collector model is p_k^{cc}. In computing the probability that the victim block contains k valid pages $p_{x_{cc}}^*$, Coupon Collector model needs to enumerate all combinations of the number of valid pages in the victim block and then find the cases where the largest number of valid pages in the blocks is k. The computational complexity to obtain A^{cc} is combinatorially increases with the number of pages in a block.

3.2 Uniform Distribution Model

Greedy Garbage Collection selects a block with the least number of valid pages as the victim block. Agarwal et al. [3] used Uniform Distribution to approximate the number of valid pages in a block and exploit them in finding the write amplification. We named the model proposed by Agarwal et al. [3] as Uniform Distribution model.

Let N_{vp}^{ud} be the number of valid pages in the victim block. By definition, the number of valid pages in the other blocks contain more than or equal to N_{vp}^{ud} number of valid pages. Let β_k be the probability that there are k valid pages in the other blocks given N_{vp}^{ud}. In Uniform distribution model, they assume that $\beta_{N_{vp}^{ud}} = \beta_{N_{vp}^{ud}+1} = \cdots = \beta_N$.

To find the probability of a valid page, it simply uses the ratio of number of pages in use and the total number of physical pages. It exploits the probability to find the distribution of valid pages, then it approximates the distribution with uniform distribution and predicts the write amplification.

$$A_{ud} = \frac{1}{2}\left(\frac{1+\rho}{\rho}\right) \tag{4}$$

As the Eq. (4) shows, the write amplification is solely dependent on ρ, which exhibits the simplest form among the four models but naturally it fails to captures the essence of the behavior because valid pages in the real SSDs do not follow uniform distribution.

3.3 Expected Value Model

Unlike the other models which are based on the valid pages in a victim block of greedy garbage collection to find the write amplification, the model proposed by Luojie et al. [14] is based on the number of invalid pages in a victim block. They use the invalid pages in a victim block as a random variable and finds the average that is the expected value of the random variable to predict the write amplification; thus, we name the model as Expected Value model.

This model assumes that all the pages in a victim block is independently invalidated with the probability of p^*. Then, the random variable N_{ip}^{ev} representing the number of invalid pages in the victim block is binomial distributed with parameters N and p^*. Then, its expectation $\bar{N}_{ip}^{ev}$ satisfies $\bar{N}_{ip}^{ev} = Np^*$. Hence, once we find p^*, we can compute $\bar{N}_{ip}^{ev}$.

We get Eq. (5) as the final equation for the write amplification, where W() denotes Lambert W function [8].

$$A_{ev} = \frac{-1-\rho}{-1-\rho-\mathrm{W}((-1-\rho)e^{-1-\rho})} \tag{5}$$

3.4 Markov Model

Desnoyers [9] further developed the earlier works and analyzed with various assumptions. Although it considers LRU cleaning and non-uniform workload, we only describe uniform workload distribution case to compare with the other models.

The model proposed by Desnoyers [9] predicts the write amplification by representing the distribution of valid pages in a block with Markov model. We used Markov model to represent the work.

The state in the model is defined as the number of valid pages in a block. It is assumed that the state i makes transition only to $i-1$ with probability of $\frac{i}{UN}$. Since this approach uses greedy garbage collection, under the assumption that x^M is the minimum value of the states of the blocks, the victim block has state x^M-1. Eq. (6) is the final equation to get the write amplification.

$$A^{mar} = \frac{N}{N-(x^M-1)} \tag{6}$$

$$x^M = \frac{1}{2} - \frac{N}{\rho+1}\mathrm{W}\left(-(1+\frac{1}{2N})(\rho+1)e^{-(1+\frac{1}{2N})(\rho+1)}\right) \tag{7}$$

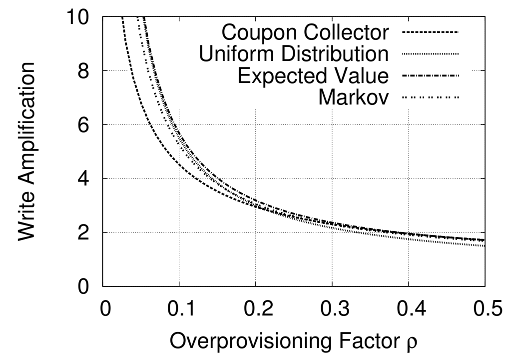

(a) N=64

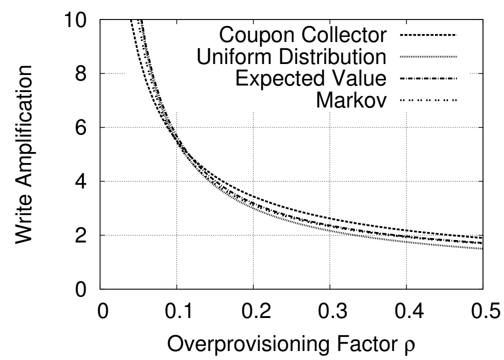

(b) N=128

Figure 1: Write Amplification Prediction with Respect to Overprovisioning Factor ρ

4. EXPERIMENTS

4.1 Experiment Setup

The environment used to compare the four models (Coupon Collector, Uniform Distribution, Expected Value, and Markov model) is described in Table 2. We first measure the write amplification using the four models, then measure the time complexity of running the model. We used MATLAB[1] running on Windows 7 64bit in Intel core i7-3820 machine with 24GB of main memory to compare the write amplification and the time complexity of analytical models.

Environment	Prediction	Measurement
CPU	Intel core i7-3820	Intel core i5-2500k
Main memory	24GB	12GB
OS	Windows 7 64bit	Ubuntu 14.04 64bit
Storage	-	Samsung 843TN

Table 2: Environments to Compare with Models and with Real SSD

Parameter	Specification
Capacity	240 GByte
Overprovisioning Capacity	23.4 GByte
Page size(read/write unit)	8 KByte
Superblock size(erase unit)	256 MByte

Table 3: SSD Specification

We compare the write amplifications of the models with the actual one of a Samsung 843TN SSD. The environment is described in the right hand side of Table 2. We modified the firmware of the device to measure the volume of requested write operations and the actual volume of data programmed on the device. We used the measurements of the real SSD to calculate the write amplification. The specification of the SSD used in this paper is described in Table 3. NAND Flash memories of the SSD used in the paper are placed in multi-channel and multi-way configuration. Unlike other SSDs, the unit of erase operation is not a block, it operates on unit of superblock which is a set blocks. The superblock is a set of a block on each NAND Flash in all channels and ways. In such configuration, channels and ways do not have any effect on write amplification because it operates as if one-channel one-way configuration.

4.2 Write Amplification

Fig. 1 illustrates the predicted write amplification with respect to overprovisioning factor ρ of Coupon Collector, Uniform Distribution, Expected Value, and Markov Model. Fig. 1(a) and Fig. 1(b) shows the predicted write amplification when the number of pages in a block N is 64 and 128, respectively.

The result shows that the total number of pages in a block is not too important in the write amplification. Uniform distribution and Expected Value model does not include N in their final model and have the same write amplification. The write amplification in Markov model increased slightly as N increased. Only the write amplification of Coupon Collector model increased noticeably as N increased.

Table 4 shows the normalized predicted write amplifica-

[1]Version: 7.11.0.584(R2010b)

N	Coupon Collector	Uniform Distribution	Expected Value	Markov
64	94%	88%	100%	96%
128	105%	88%	100%	97%

Table 4: Normalized Prediction Difference against Expected Value Model

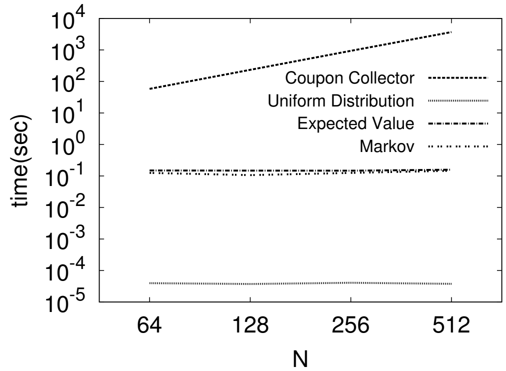

Figure 2: Time Complexity of Each Model

tion against Expected Value. It shows that Uniform distribution is off by 12% and Coupon Collector model differs by about $5 \sim 6\%$. Note the difference of Coupon Collector model between the $N = 64$ and $N = 128$ case, it is 6% lower when the total number of pages in a block is 64 and 5% larger when it has 128 pages. The result of Markov model is the closest to the result of Expected value model, it is about $3 \sim 4\%$ lower.

4.3 Time Complexity

Figure 2 shows the time spent to complete the calculation of the prediction models. We used MATLAB to measure the average time to calculate the write amplification at different overprovisioning factors. We equally divided the overprovisioning factor into 100 sections $\rho = 0.01, 0.02, ..., 0.99, 1$.

The time spent to measure the write amplification in Uniform distribution is about $4 \cdot 10^{-5}$ seconds. It shows that N has almost no effect on the time, except for Coupon Collector model. To compute the write amplification in Coupon Collector model, it takes 58 seconds, 233 seconds, 936 seconds, and 3690 seconds with $N = 64$, $N = 128$, $N = 256$, and $N = 512$, respectively. The time multiplies by factor of four every time the number of pages in a block doubles. It is because the model makes use of binomial distribution and factorial operation on N. As the number of pages in a block is increasing in SSDs, it seems Coupon Collector model is no longer feasible in practice. This is why we excluded Coupon Collector model from the comparion of this paper.

4.4 Comparison with Empirical Result

In order to have better understanding of each model and the effectiveness of the models in real life, we measured the write amplification of an SSD and compared it with the models. We format the device with EXT4 file system. We generate 8 KB write workload with direct IO to bypass the file system page cache and write immediately to the device. We also used uniform random distribution to touch all LBA addresses.

We set ρ as 0.28 which is the overprovisioning factor used in commercial enterprise SSDs [1, 2]. To set the overprovisioning factor on the SSD described in Table 3, we fill the device with cold data such that overprovisioning factor becomes 0.28. Note that, since models uses greedy garbage collection, the cold data are not selected as a victim block. We use this property to make overprovisioning factor to 0.28. We created cold data with the size of 102.1GB and 112.9GB of hot data. We generate 8KB uniform random writes on the hot data area. When the hot data is fully updated, we count it as one iteration, and repeated 15 iterations.

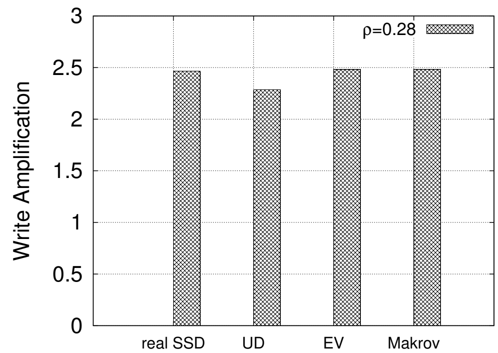

Figure 3: Write Amplification: SSD Vs. Models (N=32768)

Fig. 3 compares the write amplification of the SSD, Uniform Distribution (UD), Expected Value (EV), and Markov model, when overprovisioning factor ρ is set to 0.28 and the number of pages N. in a block is set to 32768. Coupon Collector model is not shown in the graph because it took too more than a day to compute an iteration.

The result shows that write amplification of the SSD is 2.466. Note that Expected Value and Markov model shows write amplification of 2.481. The prediction based on Expected Value and Markov model is about 0.6% lower than the actual write amplification in the SSD. Uniform distribution model shows difference of about 8%.

4.5 Write Amplification of a General Case

We performed experiments on an real SSD to compare the accuracies of the models. However, the assumption of each model on the workload is not realistic. The models assume that the size of a hot file and the utilization is as large as the size of the SSD. In this section, we run set of experiments to understand limits of the models and the effect of the utilization and the size of a file on write amplification.

We created EXT4 file system on the target SSD and created a cold and hot file. We vary the utilization from 60% to 90% of the file system partition with the size of 215 GB. We used uniform random, sequential, sequential+random workload on the hot file, and measured the write amplification of each workload. The ratio of sequential and random in the sequential+random workload is set to 1:1. We vary the size of the hot file from 10% to 40%, and the cold file occupies the rest of the file system up to specified utilization ratio.

The result of the experiment is shown in Fig. 4. It shows that the measured write amplification of number of workloads are less than 1.2 which is very different from the predicted values of the models. When the size of the hot file is small with utilization of 60% and 70% the write amplification is slightly over 1. In the case of 90% of utilization and

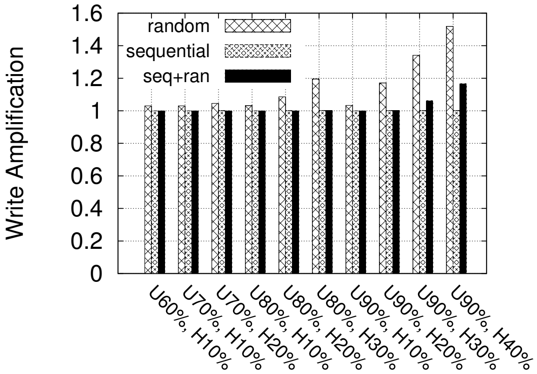

Figure 4: Write amplification by SSD utilization (U:utilization, H:hot file ratio, 10%=21.5GB)

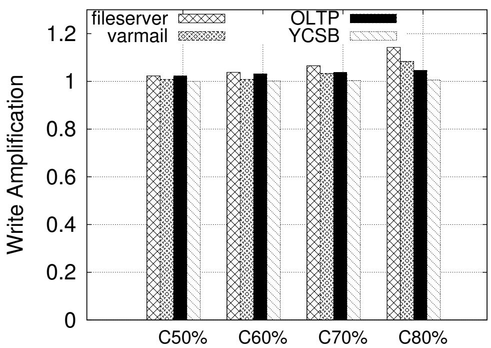

Figure 5: Write amplification of benchmarks (C:Cold file ratio, 10%=21.5GB)

the size of a hot file is set to 30% and 40% of the partition size, the write amplification of random workload is about 1.341 and 1.518, respectively.

Fig. 5 shows the write amplification of an SSD with respect to utilization. We use fileserver workload in Filebench, Varmail workload, OLTP with Sysbench, and Cassandra workload with YCSB. The configurations for the benchmarks are shown in Table 5. The result shows that the write amplification increases as the utilization increases, but the write amplification does not exceed 1.2.

The available storage capacity after formatting the SSD with EXT4 is 215GB. We counted the number of blocks U in given capacity (215GB) to and we use U to calculate the overprovisioning factor ρ and get 0.147. When $\rho = 0.147$ is used in Expected Value and Markov Model, the predicted write amplification is 4.08 showing great difference compared to the measured write amplification. The two models are limited in predicting the write amplification of an SSD when the utilization and the workload pattern varies.

		file size	files	threads
Filebench	Fileserver	256KB	80K	50
	Varmail	16KB	720K	16
		records	requests	threads
Sysbench	OLTP	45M	1M	8
YCSB	Cassandra	6M	60M	10

Table 5: Workloads of benchmarks

$$\bar{\rho} = \frac{(1 - R_{util}) + \rho}{R_{hot}} \tag{8}$$

By replacing the ρ in existing models with $\bar{\rho}$, we can more accurately predict the write amplification of an SSD while varying the utilization (R_{util}) and the ratio of the hot file (R_{hot}). For example, when the utilization and the ratio of a hot file is 90% and 30%, respectively, then we get $\bar{\rho} = \frac{(1-0.9)+0.147}{0.3} \approx 0.823$. By applying $\bar{\rho} = 0.823$ on Expected Value and Markov model, we get write amplification of 1.349 which is very close to measured write amplification of 1.341.

Utilization	80%	90%	90%	90%
Hot ratio	30%	20%	30%	40%
$\bar{\rho}$	1.16	1.24	0.82	0.62
real SSD	1.19	1.17	1.34	1.52
EV	1.19	1.17	1.35	1.53
Markov	1.19	1.17	1.35	1.53

Table 6: write amplification by $\bar{\rho}$

Except for uniform random workloads, the measured write amplification of an SSD with general workloads is less than 1.2. We can set $\rho > 1.15$ to get the write amplification of 1.2 in Expected Value and Markov Model for general workloads.

5. RELATED WORK

As SSDs are becoming more prevalent, many are trying to analyze and analytically predict the behavior of SSDs. There are some works [11, 12] that exploits windowed greedy garbage collection algorithm with assumption the workload is uniformly or non-uniformly spread across the given storage space to find the write amplification. Agarwal et al. [3] and Luojie et al. [14] uses greedy garbage collection under uniform workload to analyze the write amplification and they also investigate the relationship between the write amplification and the size of an SSD. Desnoyer [9] also provides a model to find the write amplification in greedy and LRU (FIFO) garbage collection algorithm. His work takes a step further to analyze it under non-uniform workload to find the write amplification in real-world workload.

While most works are based on greedy garbage collection, Houdt [17] set out to analyze the write amplification with an SSD running a random garbage collection algorithm. He assumed d-Choices garbage collection and finds the relationship between d and write amplification. Li et al. [13] analyzes the effect of locality of workload in garbage collection algorithm. After analyzing the cost of cleaning and wear-leveling, they show that exploiting the locality increases the efficiency of the garbage collection.

6. CONCLUSION AND FUTURE WORK

As SSDs are being widely adopted by both users and storage system vendors for performance reasons, having thorough understanding of the write amplification of SSDs are becoming more important. It is mostly because the write amplification is closely related to the life time of the SSDs. A few have analyzed the write amplification, but most of them are based on simulation or an analytical model of the behavior. A problem in understanding or building on top of the previous works is that notations, basis, and assumptions are not the same to each other. In this work, we first analyze the existing works in depth to not only address the pros and cons of the works but also to re-evaluate the models with same standards and assumptions. Secondly, we point out that the time complexity of a model is as important as the accuracy of the model. For example, Coupon Collector model cannot be used in prediction because as the number of pages in a block increases with factor of two the time complexity increases with factor four. Finally, we evaluated the write amplification of a real SSD and compared it with analytical models, which previous works have failed to show. The result shows that Expected Value and Markov model show the most closest resemblance to the actual write amplification of an SSD with difference of only 0.6% in uniform random case. The accuracy of Uniform Distribution is much lower than the two, it is off by about 8%. The experiment shows that the write amplification of models varies greatly depending on the utilization and hot file ratio in the SSD. Use of models available in the field have to be considered carefully because they are not accurate in representing the write amplification of general use cases.

All the models lack practicality in real world because they are based on the assumption that workload follows uniform distribution and exploit a form of greedy garbage collection. Two of the models, Coupon Collector and Markov model, adopts the notion of hotness in to their models to consider the case with non-uniform traffic. If their non-uniform traffic approaches are to have any practical values, wear-leveling of the SSDs must be considered, but their models do not take the wear-leveling into account. Greedy garbage collection is not the only garbage collection algorithms, yet all the models are based on the primitive garbage collection algorithms. Since the write amplification is heavily dependent on the garbage collection policy, we need more in-depth studies on the write amplification with different garbage collection algorithms. Here are suggestions for future work. Since all of the given models are based on a single storage system, uniform workload distribution, and primitive garbage collection algorithm, this field of study asks for better model that explains parallelism of SSD, wear-leveling, sophisticated garbage collection policy, and real world workload. The study of write amplification of SSDs needs to broaden its focus on addressing the life span of SSDs as a function of workload, wear-leveling, and the write amplification.

7. ACKNOWLEDGEMENT

This work is supported by IT R&D program MKE/KEIT (No. 10041608, Embedded System Software for New-memory based Smart Device), the ICT R&D program of MSIP/IITP (R7117-16-0232, Development of extreme I/O storage technology for 32Gbps data services), the BK21 plus program through the National Research Foundation (NRF) funded by the Ministry of Education of Korea, and the Basic Science Research Program of the National Research Foundation of Korea (NRF-2015R1D1A1A09059663).

8. REFERENCES

[1] Understanding over-provisioning. Technical report, Kingston Technology, 2013.

[2] Over-provisioning: Maximize the lifetime and performance of your ssd with small effect to earn more. Technical report, Samsung Electronics Co., 2014.

[3] R. Agarwal and M. Marrow. A closed-form expression for write amplification in nand flash. In *Proceedings of GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, pages 1846–1850. IEEE, 2010.

[4] N. Agrawal, V. Prabhakaran, T. Wobber, J. D. Davis, M. S. Manasse, and R. Panigrahy. Design tradeoffs for ssd performance. In *Proceedings of USENIX Annual Technical Conference*, pages 57–70, 2008.

[5] W. Bux and I. Iliadis. Performance of greedy garbage collection in flash-based solid-state drives. *Performance Evaluation*, 67(11):1172–1186, 2010.

[6] L.-P. Chang. On efficient wear leveling for large-scale flash-memory storage systems. In *Proceedings of the 2007 ACM symposium on Applied computing*, pages 1126–1130. ACM, 2007.

[7] L.-P. Chang and T.-W. Kuo. An adaptive striping architecture for flash memory storage systems of embedded systems. In *Proceedings of 8th IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 187–196. IEEE, 2002.

[8] R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth. On the lambertw function. *Advances in Computational mathematics*, 5(1):329–359, 1996.

[9] P. Desnoyers. Analytic modeling of ssd write performance. In *Proceedings of the 5th Annual International Systems and Storage Conference*, page 12. ACM, 2012.

[10] P. Erdős. On a classical problem of probability theory. 1961.

[11] X.-Y. Hu, E. Eleftheriou, R. Haas, I. Iliadis, and R. Pletka. Write amplification analysis in flash-based solid state drives. In *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, page 10. ACM, 2009.

[12] X.-Y. Hu and R. Haas. The fundamental limit of flash random write performance: Understanding, analysis and performance modelling. Technical report, IBM Research Report, 2010/3/31, 2010.

[13] Y. Li, P. P. Lee, J. Lui, and Y. Xu. Impact of data locality on garbage collection in ssds: a general analytical study. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering*, pages 305–315. ACM, 2015.

[14] X. Luojie and B. M. Kurkoski. An improved analytic expression for write amplification in nand flash. In *Proceedings of Computing, Networking and Communications (ICNC), 2012 International Conference on*, pages 497–501. IEEE, 2012.

[15] D. Shmidt. Trueffs wear-leveling mechanism. *Technical Note (TN-DOC-017), M-Systems Flash Disk Pioneers, May*, 20(2002):1–4, 2002.

[16] B. Van Houdt. Analysis of the d-choices garbage collection algorithm with memory in flash-based ssds. In *Proceedings of the 7th International Conference on Performance Evaluation Methodologies and Tools*, pages 11–19. ICST, 2013.

[17] B. Van Houdt. A mean field model for a class of garbage collection algorithms in flash-based solid state drives. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*, pages 191–202. ACM, 2013.