INFOTEH-JAHORINA Vol. 15, March 2016.

SSHD: Modeling and Performance Analysis

Borislav Đorđević, Valentina Timčenko, Slavica Boštjančič Rakas
Institute Mihailo Pupin, University of Belgrade
Belgrade, Serbia
bora@impcomputers.com, valentina.timcenko@pupin.rs, slavica.bostjancic@pupin.rs

Abstract— **In this paper we present specific characteristics of the new generation data discs - Solid State Hybrid Drives (SSHD), and propose an adequate SSHD mathematical model in the context of cache, which includes SSD and HDD behaviors. In a way to validate the model we further analyze the SSHD performances when comparing to the solid state drive (SSD) and hard discs (HDD). The analysis is based on the published set of test procedures, based on which are examined gained qualities, advantages and disadvantages of SSHD hybrid disc when compared to traditional and SSD drive.**

Ključne SSHD; model; performance; analysis; Windows

I. INTRODUCTION

The fast development of computer technologies demands for the increase of the performance of applied computer hardware. In addition to powerful microprocessors that handle large amounts of data per unit time, there is a need of making possible fast and accurate storage of large amounts of data. The traditional magnetic hard disk drive HDD is a type of secondary memory in which data is magnetically recorded in concentric tracks on the surface of the disk plate. Unfortunately, HDD cannot provide the required performances thus there is need of providing new storage options. One solution is to apply an electronic medium of mass storage, Solid State Drive (SSD), which is faster than HDD, has better fault tolerance, but has higher price and provides lower storage capabilities. SSDs are practically the upgrade devices to hard disks, which do not have movable parts, and are based solely on the use of electronic elements and assemblies. Unlike magnetic hard drives, SSDs do not rely on magnetic field for writing data, allowing lower power consumption, silent and faster access to data [1].

To overcome some of the disadvantages, but still keep with better performances than with HDD, a new generation of Solid State Hybrid Drives (SSHD) has taken the market.

SSD technology has limitations, as each transistor can be used limited number of times before it will no longer be operational. It uses the wear level mechanism that adjusts the consumption level in order to prevent premature consumption of the drive. But, as SSD cost-per-bit continues to decrease there is an increase of the interest in application of SSD for a wide set of environments [2, 3].

II. OBJECTIVE AND MOTIVATION

In this paper we analyze the performances of the hybrid version of HD and SSD drives, the SSHD drive. As SSHD are relatively novel drive format, our goal is to demonstrate what are their advantages and disadvantages, thus helping users in choosing the proper disk drive for their needs. We present test procedures encompassing a set of internal Windows 8 tests covering the performances of each subsystem within the used computer. Additionally, we present results obtained with the referent group of benchmark tools which allows operating system to estimate the efficiency of the installed hardware.

III. SOLID STATE HYBRID DRIVES

Solid State Hybrid Drive (SSHD) represents the combination of the traditional and solid state drive. Figure 1 shows the Seagate 500GB hybrid SSHD.

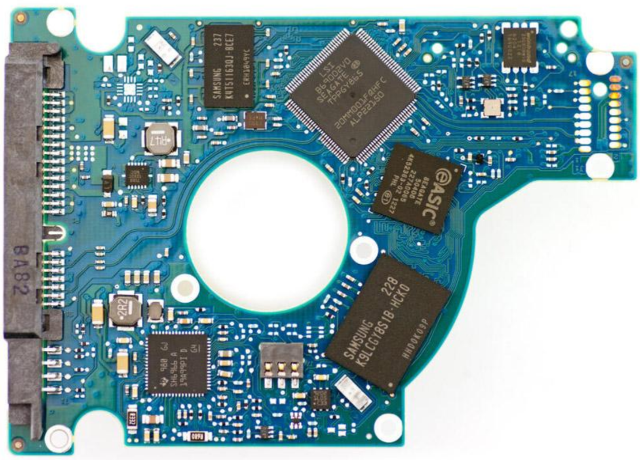

Figure1. Seagate 500GB hybrid SSHD [4].

SSHD drives have necessary performance and speed, but they are expensive. This technology is considered as highly acceptable for data centers, Redundant Array of Independent Disk (RAID) architectures, etc [5]. The most important characteristics of SSHD are as follows:

A. Capacity

SSHD technology integrates a small amount of negative-AND (NAND) flash memory (4-8GB) inside the traditional HDDs. The result is a combination of the characteristics of a high-capacity hard disk of 500GB up to 4TB and high-speed delivery of commonly used data at approximately equal performances as ultra-fast SSD drive. They offer the same memory capacity as the HD with significantly higher operational speeds.

The point is that a small capacity of built-in SSD NAND is used as a form of temporary storage for frequently used data. Due to its much faster response to request for data, the operational speed is higher comparing to the HDD.

B. Hardware

As opposed to the HDD drive, the SSD has no movable parts, and depends solely on the electricity. It is faster, silent and operationally efficient. SSHD as a combination of these two technologies has taken the best from each of them. But, the frequent execution of clear and write data cycles can cause some SSHD shortcomings based on smaller capacity and possibilities for data manipulation. However, it overcomes the problem of HDD which slows the entire system due to time-consuming response of mechanical parts.

C. Software

Due to the limited capacity of NAND flash memory, SSHD can keep only a certain amount of data. Therefore, one of the most important parts of the SSHD disks is sophisticated software based on the algorithm that learns what are the most used data, thus keeping them in NAND. The adaptive memory adjusts to users operational routine and requirements, resulting in a faster time loading the operating system and more responsive to the execution of specific tasks.

IV. SSHD/SSD/HDD MATHEMATICAL MODELING

The mathematical model of SSHD encompasses three components: RAM disk, SSD disk and magnetic hard disk. RAM disk performes the caching for SSD disk, and is labeled as "L1 cache level" for the needs of this modeling. SSD disk is caching for the magnetic hard disk, and is marked as L2 cache level. It is presented in Figure 2.

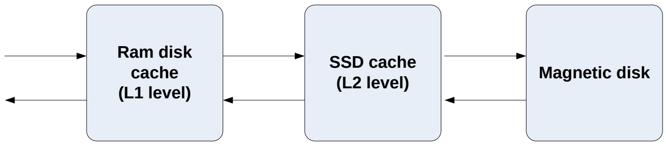

Figure2. SSHD mathematical model components.

Average SSHD access time, $T_{request}$, is modelled by the following equation:

$$T_{request} = hit\_rate_{L1} * hit\_time_{L1} + \\ + (1 - hit\_rate_{L1}) * miss\_penalty_{L1} \qquad (1)$$

Where the hit_rate_{L1} represents the probability of the successful hit in the RAM disk cache (L1), and hit_time_{L1} stands for the typical access time of RAM, which forms the RAM disk cache.

The $miss\_penalty_{L1}$ represents the measure of the L1 access to the SSD cache (L2). It is defined with the equation (2):

$$miss\_penalty_{L1} = hit\_rate_{L2} * hit\_time_{L2} + \\ + (1 - hit\_rate_{L2}) * miss\_penalty_{L2} \qquad (2)$$

where the hit_rate_{L2} represents the probability of the SSD disk cache hit (L2), the hit_time_{L1} is typical SSD access time, which strongly depends on the cycle type (read or write). The write cycle is more time/CPU consuming as it implies the delete operation over a huge block of data, which is preceded by the garbage collection procedure (also proceeded for the entire data block).

In this model, the parameter $miss\_penalty_{L2}$ represents the access time to the L2 cache of the magnet disk, and is calculated based on the equation (3):

$$miss\_penalty_{L2} = T_{Seek} + T_{rotational} + T_{media} \qquad (3)$$

where T_{Seek} stands for the time of positioning, $T_{rotational}$ defines the rotational delay time, and T_{Media} represents the time needed to access the magnetic media.

This model integrally applies to SSHD, while equations (1) and (3) are valid for HDD, and (1) and (2) for SSD [6, 7].

V. SPECIFICATIONS OF THE TESTED DRIVES

For the needs of the evaluation of the proposed mathematical model, we are providing and examining the industrially performed tests and comparison [8]. They base test procedure on the comparison of the performances of 1TB SSHD disk, 1TB HDD and 240GB SSD. The tests are performed in the HD Tune and PC Mark 8 benchmarking software environments. Taking into consideration our model we have expected the following results: SSD performance should be significantly better when compared to SSHD and HDD. The main reason is the fact that it performs based on the equation (1). The SSHD should perform better that HDD, depending on the SSD cache efficiency (number of accomplished cache hits). The test environment characteristics are described in Table 1.

TABLE I. TEST ENVIRONMENT [8]

Motherboard	MSI Z87 MPower
Processor	Intel Core i7 4770K
Memory	32GB(4x8GB), DDR3 2166 CL11
SSHD, SDD, HDD	Seagate SSHD 1TB, ST1000LM014 Seagate SSD 600, 240GB, ST240HM000 Seagate Momentus Thin HDD 320GB
Driver	Catalyst 13.11 BetaV5
Operating System	Windows 8 Ultimate 64-bit SP1

The lack of the expensive disk environment has guided us to find a good reference to our assumptions, and provided review has very concretely answered to our assumptions [8]. Additionally, we present the test results for SSH Seagate 600 drive, in a way to compare performances with solid state disk representative.

Seagate ST1000LM014 is 1TB disk. It uses four magnetic heads and two two-side plates [9].

The Seagate 600 is SSD type drive with following characteristics: SATA 6 Gb/s, MLC, Sequential Read Rate (MB/s) >500 and Write Command Rate (MB/s) >400, 7mm in height [10].

Seagate Momentus Thin HDD 320GB represents a SATA 6Gb/s disk with 32MB cache memory capacity [11].

VI. PERFORMANCE ANALYSIS

The referent tests are presented by the sequential and random read and write [8]. The first set of tests is performed for obtaining the sequential read characteristics. In figure 3 we provide graphical representation of the obtained results. The obtained sequential read results show that SSD are still unchallenged disk option, but when making comparison of SSHD and HDD there is a slight difference and advantage for SSHD.

Actually, due to the large read data transfer, there is a small impact of L1 and L2 caches (equations (1) and (2)), and SSD (magnetic) components are dominant. Due to the absence of mechanical components, SSD considerably wins and SSHD and HDD, while there is just a slight difference between the SSHD and HDD.

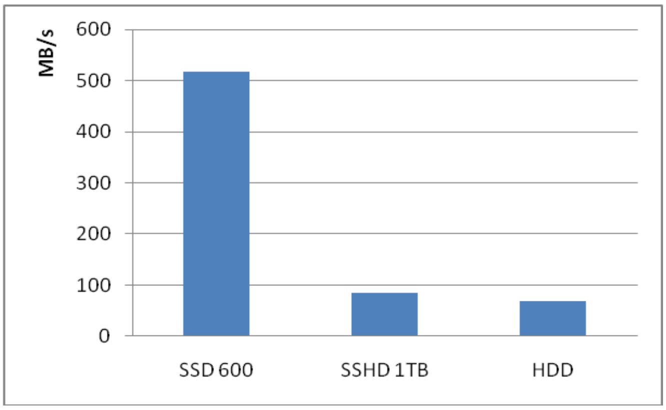

Figure 3. Sequential read [8].

In figure 4 we have presented the test results for sequential writing operation. The sequential write provides even better results for SSHD when compared to HDD, thus it is preferable and recommendable choice for users that cannot afford SSD but need performance improvement over the HDD.

When testing the sequential write performance, it is noticeable that due to large write data transfer there is a low impact of L1 cache. Still, the impact of the L2 cache is greater when compared to the obtained sequential read performances.

Due to the absence of mechanical components, SSD is considerably more efficient than SSHD and HDD, although it relies on block-erase overhead. Additionally, as a consequence of the equation (2) and positive impact of the SSD L2 cache, SSHD outperforms HDD.

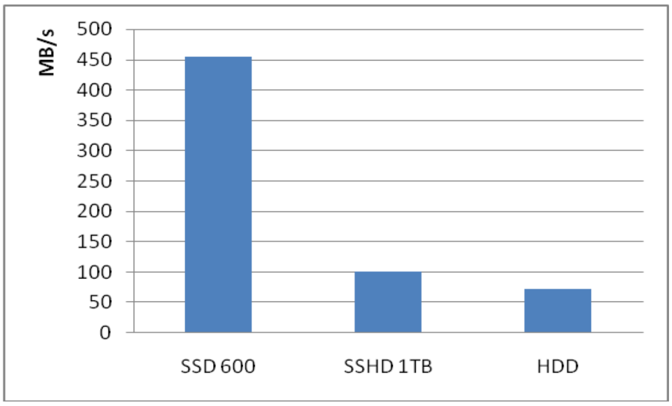

Figure 4. Sequential write [8].

In Table 2 and figures 5,6 we present the results obtained for random read and write performances. These tests have confirmed the SSD domination, but when comparing SSHD over the HDD, there are some interesting differences. HDD is slightly better with random reads, but the true difference comes with the random writes, as the SSHD is strongly overcoming the HDD option.

TABLE II. RANDOM READ AND WRITE [8]

Drive type	SSD 600	SSHD 1TB	HDD
Random Read [MB/s]	297,17	1,33	1,81
Random Write [MB/s]	229,2	21,3	1,78

Figure 5 and 6 represent the test results for random read and random write, respectively.

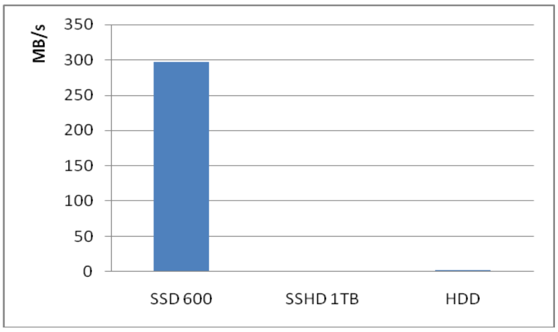

Figure 5. Random read [8].

In the case of the random testing it is noticeable the low impact of the L1 and L2 caches, which is in accordance to equations (1) and (2), while the dominant influence has an integrated media access technology, which is most efficient in the case of SSD, and drasticaly outperforms HDD and SSHD.

When analyzing the obtained performances of the HDD and SSHD, just in case of testing the random writes the L2 cache has demonstrated a significant impact, in a way that the L2 cache will absorb a large amount of small enrollments, thus HDD performs much better than SSHD. In the case of the random read, the impact of L2 cache is minor, to the point that we have detected that the HDD is again slightly faster.

We have noted that the random read and random write performances are very similar in the case of the HDD, but the real difference arises in the cases of SSD and SSHD. For SSD, random read test results significantly over perform random read for SSHD as a result of using features "block erase overhead" and "garbage collection". IN the case of the SSHD random write, we have an oposite case, as the random write performances are much faster than random reads, due to the L2 caching that practicaly absorbs the small data writes.

In general, for all random operations, due to used technology the SSD outperforms SSHD and HDD. The equation (3) is dominant in the case of the random workloads for SSHD and HDD, where the magnetic technology has great impact to the generated mechanical delays. On the other hand, the SSHD performs better than HDD, mostly due to the efficient SSD L2 cache (equation (2)) which increases random write operation performances.

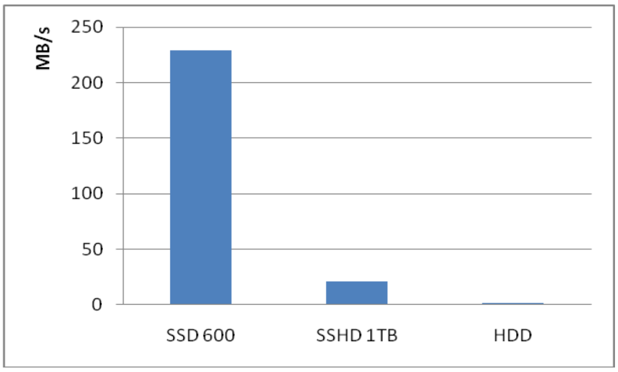

Figure 6. Random write [8].

The general conclusion is that SSD and SSHD over perform HDD mostly due to the huge beneficial caches, L1 and L2. This confirms our mathematical cache modeling, as complies perfectly with equations (1) and (2). Based on these equations, the random workloads are efficiently operated in SSD and SSHD environment, while the differences in sequential tests results are lower when comparing to the HHD environment. This result is due to the lower impact of the L1and L2 caches.

As it can be noticed, these test results have validated proposed mathematical cache model and proved our expectations and performance improvement when making usage transition from 2.5" HDD disks to SSHD models. The captured difference is not that obvious as it is shown when using fast SSD disks, but still it is noticeable. The mayor impact is on the random write operation, which is one of the most frequent scenarios when working in Windows environment.

In any case, we see the SSHD as a promising technology that will find the target users and allow better performances is some specific environments.

ACKNOWLEDGMENT

The work presented in this paper has partially been funded by the Ministry of Education, Science and Technological Development of Republic of Serbia: V. Timcenko and S. Bostjancic Rakas by grants TR-32037, TR -32025, and B. Djordjevic by grant III-43002.

REFERENCES

[1] H. H. Huang, S. Li, A. Szalay and A. Terzis "Performance Modeling and Analysis of Flash-based Storage Devices", 2011 IEEE 27th Symposium on Mass Storage Systems and Technologies (MSST), pp.

[2] S. Moon, J. Lee, Y. S. Kee"Introducing SSDs to the Hadoop MapReduce Framework", in proceeding of the 2014 IEEE International Conference on Cloud Computing, pp. 272-279, 2014

[3] B. Đorđević, V. Timčenko, S. Obradović, N. Čorni „Uticaj internog kontrolera na performanse SSD diskova“, YU INFO konferencija, Kopaonik, pp. 433 - 436, 2013.

[4] "Seagate 500GB hybrid SSHD", Seagate Technology LLC Online: http://www.seagate.com/www-content/product-content/momentus-fam/momentus-xt/en-gb/docs/laptop-sshd-family-ds1814-2-1406gb.pdf

[5] K. S. Krishnan, S. ArunKumar, "Data and Parity Storing Techniques:A Survey", International Journal of Engineering Sciences and Research Technology, pp. 338 - 343, 2014.

[6] P. Desnoyers, "Analytic Modeling of SSD Write Performance", ACM Transactions on Storage 10(2), 2012.

[7] N. Agrawal, V. Prabhakaran, T. Wobber, J. D. Davis, M. Manasse, R. Panigrahy "Design Tradeoffs for SSD Performance", Proceedings of the USENIX Technical Conference, June 2008.

[8] "SSHD – kombinacija tradicionalnog hard disk i Solid State Disk uređaja", BozzaBench

Online: http://www.bozzabench.com/Tests/TestView/tabid/95/id/455/

[9] "Laptop Thin SSHD", Seagate Technology LLC. Online: http://www.seagate.com/gb/en/internal-hard-drives/solid-state-hybrid/laptop-solid-state-hybrid-drive/

[10] 600 Pro SSD, Seagate Technology LLC

Online: http://www.seagate.com/gb/en/support/internal-hard-drives/solid-state-hybrid/600-pro-ssd/

[11] "Seagate Momentus Thin HDD", Seagate Technology LLC. Online: http://www.seagate.com/gb/en/support/internal-hard-drives/laptop-hard-drives/momentus-thin/