Performance Analysis of SSD/HDD Hybrid Storage Manager[†]

Seongjin Lee, Seokhui Cho, Haesung Kim, Youjip Won
Dept. of Electronics and Computer Engineering
Hanyang University
Seoul, Korea
{insight, misost, matia513, yjwon}@hanyang.ac.kr

***Abstract*- The technology and cost barrier of the SSD motivated many research groups and companies to merge the SSD with low cost and large capacity storage device, Hard Disk Drive (HDD). Exploiting the SSD as a cache disk for the HDD not only reduces overall power consumption but also enhances responsiveness of a system. Four different cache management solutions are reviewed and validated advertised benchmark score. We also define a term, budget, to measure the maintenance cost of the cache managers. We show that HyperDuo is better that other device in terms of the budget.**

I. INTRODUCTION

Recently, due to dramatic price reduction of flash memory [1], NAND flash based Solid State Drive (SSD) are sought to be an alternative to reduce the I/O bottleneck. Although price reduction rate is increasing as bits per cell are increasing, the reliability of the SSD is still an issue because write endurance of a cell in NAND flash drops rapidly. While it would take some more years for consumers to fully enjoy benefits of SSD in lower price, manufacturers have found a new niche market. Exploiting the SSD as cache for Hard Disk Drive (HDD) has gained some popularity and many manufacturers are introducing various yet similar products.

There are many benefits in using SSD as a cache disk for the HDD. With attached SSD on the system, not only the life time of the SSD extends by differentiating direction of read and write operations, system-level performance increases by executing boot sequences and application data from the SSD which has very high read performance. In terms of power consumption of a system, it is better to operate on SSD rather than on HDD because when reads are directed to the SSD, it allows HDD to be in spin-down state longer and save about 500mW [2].

Since there are many types of SSD form-factors available in the market and satisfying different needs, it is hard to cover all types. Software solution may be enticing in terms of installation budget and have variety of interconnects choice from SATA to PCIe SSD form-factors; however, performance gap between hardware controller can be an issue.

There are at least four companies that provide methods to enhance the performance of HDD through cache. One of them is a software company called Nvelo [3], which was formed to develop a software solution for NVM cache management. Two of the companies, Marvell [4] and SilverStone [5] provide cache management solution through host controller. When Intel [6] announced Z68 architecture, Rapid Storage Technology and Smart Response Technology was a major step up to reduce the I/O bottleneck.

Main purpose of this article is to first review number of related products. In this article, we review four products: Dataplex from Nvelo, HyperDuo from Marvell, HDDBoost from SilverStone, and Smart Response Technology from Intel. Technical background and caching algorithms are proprietary of the manufacturer. It is not possible to fully understand mechanisms used but we collected as much information as possible to understand the product. Second purpose of this article is to run benchmark to verify the result with our own test environment. Although the manufactures provide benchmark results, the test results vary by the environmental settings. Finally, we defined a term for measuring the maintenance cost of the cache manager, called budget. By comparing the budget of each device, we show that HyperDuo is better that other device.

II. RELATED WORKS

Recent advent of technologies to enhance the reliability and capacity of NAND flash memory based SSD has created a new sub-segment in storage market by using the SSD as a cache disk of a HDD. Payer et al. proposed a method and a prototype to optimize the cost and the performance in a heterogeneous storage device [7]. They proposed Combo Drive that concatenates SSD with HDD which consists of I/O manager that communicates Combo Drive, Optimizer that moves files by the file types and access frequency of the files, and finally File Mover library which is implemented in user space to copy clusters of files in file systems to logical clusters of the disk.

While controller and firmware companies such as Marvell, LSI, and NVelo are introducing devices and software that supports heterogeneous storage configuration, Microsoft also have strived to reduce the I/O bottleneck in O/S level. One of recent technology Microsoft introduced is Superfetch [8]. To some extent, Superfetch can be both a friend and foe to the storage system. In essence, Superfetch is a learning algorithm that predicts next most likely application that a user will launch. In perspective of HDD, superfetch reads in most likely data to DRAM prior to application launch process. Reading in required data in application launch to DRAM enhances response time. But, when it is combined with heterogeneous disk configuration, Superfetch interferes with cache algorithm and deteriorates performance of cache considerably.

[†]This work is sponsored by IT R&D program MKE/KEIT. [No.10035202, Large Scale hyper-MLC SSD Technology Development], and partially supported by LG CST Gr. M&S LAB.

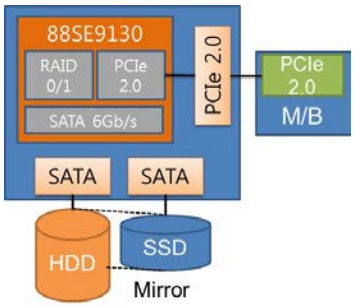

Figure 1 HyperDuo

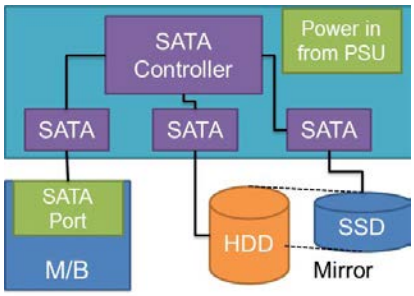

Figure 2 HDDBoost

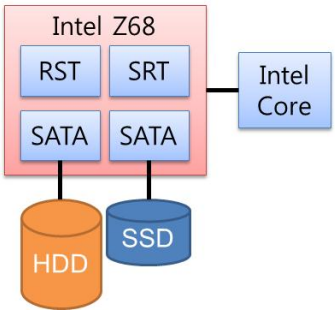

Figure 3 Smart Response Technology

III. Cache Managers

A. *Dataplex - Nvelo*

Little is known about Dataplex other than it is a software solution that combines the capacity and low cost of HDD with faster yet lower capacity SSD. There are three requirements in using Dataplex. First is that the cache device must have at least four times the system memory to produce right performance. Second, the SSD should be sufficiently faster than the HDD. Finally it requires intelligent and adaptive caching algorithm on the host system.

According to presentation given by Nvelo in Flash Memory Summit 2010 [9], intelligent algorithm of Dataplex learns user behavior and decides which data should reside in the SSD. Some of criteria in making decision on important data is frequency and recentness of the data. In order to make such decision on the data, block and file level location of the data must be visible to caching mechanism.

Advertise performance of Dataplex on Intel i7 CPU 920@ 2.67Ghz with Asustek P6X58D-E, 4GB 1066 Mhz DDR-DRAM, Intel X25-M G2 80GB SSD as a cache, and Hitachi 2.5'' 7200RPM 320GB HDD produces average of 18000 on PCMARK HDD Suite.

B. *HyperDuo - Marvell*

On January 2011, SATA embedded controller market leader Marvell introduced 6Gb/s SATA host controller which enables about 80 percent of the performance of a SSD at a lower cost [10]. Figure 1 shows the architecture of Marvel HyperDuo. Marvel 88SE9130 host controller supports two 6 Gb/s SATA peripheral interfaces and a one-lane 5.0 Gb/s PCIe host interface, which means that a HDD and a SSD is attached to the device with host controller with 6 Gb/s and the controller itself is connected to 5.0 Gb/s PCIe interface. This configuration confuses a bit at first because PCIe interface limits the bandwidth of transferring data to HDD or to SSD; however, current I/O bandwidth of HDD or SSD does not exceed over the 5.0 Gb/s.

Marvel 88SE9130 host controller implements a port multiplier to share a single 5.0 Gb/s PCIe lane and a hardware RAID 0/1 module to use SSD as a cache. The host controller provides two modes of operation which is shown in Figure 1. First mode is safe mode or mirror mode using RAID 1 scheme, and second mode is capacity mode. Safe mode duplicates first few GBytes of the HDD to SSD and manages hot files in SSD. According to Marvell product brief [11] states that Office, Media player, Adobe Creative Suites, iTunes, Internet browsers, and "hot'' accessed O/S-related files are pinned to SSD. Hot denotes data which are frequently accessed by O/S and pinned data denote sectors stored in the nonvolatile cache (NVCache). The document also states that MRU data structure periodically identifies hot and cold LBA ranges and pin the hot to SSD.

Advertised performance of Marvell is 25428 on PCMark benchmark [11], where Intel 80GB SSD is used as a cache disk, and Seagate 320GB is used as a system disk. The benchmark result is about five times faster than HDD alone and about 80% of SSD alone system.

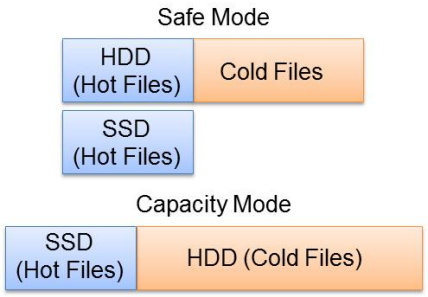

Figure 4 Two modes in HyperDuo

C. *HDDBoost - SiverStone*

HDDBoost is yet another device from Silverstone to provide solution to reduce the I/O bottleneck by installing SSD in between the I/O hierarchy. Figure 2 shows the architecture of HDDBoost. Configuration of HDDBoost is similar to that of HyperDuo except that the device uses SATA as communication interface. In essence HDDBoost is a two port RAID 1 controller that combines a HDD and a SSD. HDDBoost only provides mirroring mode. To solve the mismatch in capacity of two heterogeneous disks, first area of HDD is mirrored to the SSD. Capacity of mirrored region on HDD is decided by capacity of the SSD.

HDDBoost modified traditional RAID 1 scheme to work around the write endurance of the SSD. Initial installation of SSD in RAID 1 configuration duplicates entire storage area of SSD with mirrored area of the HDD. When the host requests a data from the storage, HDDBoost first checks if the requested data exists in SSD [12]. If the data resides in the SSD, the request data is serviced with the data in SSD, otherwise it fetches the data from the HDD. On the other hand when the host writes a new data to the storage, HDDBoost directs all of

the write operations to HDD including non-mirrored area; there are no write operations to the SSD from the host.

For example, when there is a write request to mirrored region on the HDD, HDDBoost writes the data to the HDD, and it mirrors the data back to SSD. Example of read and write operation on HDDBoost is shown in Figure 4. Initial state of data on the HDD is {1, 2, 3, 4, 5, 6, 7, 8} where each number represents a data and first 5 elements are mirrored on the SSD. If write request of {4, 5, 6, 7, 8, 9, 10} arrives, {4, 5} is directed to mirrored region of the HDD which is later updated to SSD and {6, 7, 8, 9, 10} is directed non-mirrored region of the HDD. If read operation on {4, 5, 6, 7, 8} is requested, HDDBoost checks if requested data is on SSD. Since {4, 5} is on the SSD, the data is serviced from the SSD, and {6, 7, 8} is serviced from the HDD.

For the record, Silverstone HDDBoost PCMark Vantage benchmark is 5860 with WD 2TB Caviar Green with X25-V 40GB SSD [13].

Figure 5 Read Write operation of HDDBoost

D. Smart Response Technology - Intel

Not only Intel Z68 chipset has best of two previous Intel chipsets P67 and H67, it also adds new capabilities to support SSD-like system responsiveness. The technology that enables the responsiveness is Smart Response Technology (SRT) under Rapid Storage Technology (RST). RST version 10.5 works with HDDs and SRT with SSDs. The abstract of SRT architecture is shown in Figure 3. SRT is directly embedded in the Z68 chipset, whereas other solutions have to pass through extra I/O hierarchy. A HDD and a SSD can be directly attached to the motherboard and two heterogeneous disks are integrated with RAID 0. An application has to be installed to enable the SRT feature.

The application provides two modes of operation as shown in Table 1. First mode of operation is Enhanced mode, which uses write-thru cache mechanism, and second mode of operation is Maximized mode, which uses write-back cache. Two modes in SRT shows that it not only cache read operations but also write operations. Performance difference between write-thru and write-back cache is clear, especially when the main disk is slower than the cache disk. Write-thru operation which is bounded by the performance of the HDD, reflects all writes to HDD and SSD simultaneously, whereas write-back cache holds the written data to the SSD, and synchronize the data to HDD at system idle time.

SRT is targeted on speeding up three categories, application loading, system boot times, and frequently accessed software [14]. Cache mechanism used in SRT takes individual blocks of data as unit to write on an SSD and distinguishes one-time accessed sequential data stream from application, user, and boot data which has higher value as a cache data. PCMark Vantage result of SRT is 16548 in the Enhanced mode with WD Raptor 10K RPM and Intel 311 SSD [15].

Table 1 Modes in Smart Response Technology

Mode	Performance	Dual Boot Safe	Data Caching
OFF	No Acceleration	N/A	N/A
Enhanced (Default)	Boot-time, Paging Run-time reads	Yes	Write-Through
Maximized	Boot-time, Paging Run-time reads Run-time write	No	Write-back

Table 2 Benchmark Environment

Device	Model
CPU	Intel i5-2500 3.3Ghz
RAM	DDR3 1333MHz 4G
M/B	Gigabyte Z68X-UD3H-B3
HDD	WD SD10EALX (1T, SATA3, 7200 RPM, 32M)
SSD	OCZ-Solid3 60G (SATA3, 20k IOPS)
O/S	Win7 32bit Home Premium Eng
Benchmark	PCMark Vantage 1.0.2
HyperDuo	Highpoint RocketRAID 62X (Safe Mode)
HDDBoost	Silverstone HDDBoost Rev-1.02
SRT	Rapid Storage Technology 10.6.0.1002 (Enhanced Mode)

III. ENVIRONMENT

Table 2 shows benchmark environment. We used motherboard with Z68 architecture supporting Smart Response Technology. In order to reduce the interference between different cache architectures, each cache devices are installed only when the corresponding device is being tested. Although the manual of the product states that the devices can be installed without having to reinstall the operating system, system disk is formatted and operating system is reinstalled when different device is attached to the system. We could not acquire Dataplex so we did not measure the performance of Dataplex.

Installing HDDBoost requires few more steps than installing other devices. At first only HDD is connected to HDDBoost controller. When the system is booted for the first time, BIOS setting is changed to support AHCI mode. After driver files are correctly installed then we attached SSD for caching.

Manufacturers use PCMark Vantage 1.0.2 [16] to measure the performance of a system or a disk and we use the same benchmark to compare the advertised performance with measured performance. PCMark Vantage provides 8 set of workloads. Workloads provided by the benchmark are as follows: 1 Windows Defender, 2 Game, 3Importing pictures to Windows Photo Gallery, 4 Start-up, 5 Video editing using Windows Movie Maker, 6 Windows Media Center, 7 Adding Music to Windows Media Player, and 8 Application loading. The number denotes the specific set of workload.

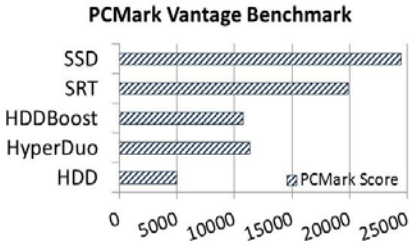

Figure 6 Benchmark Overall score

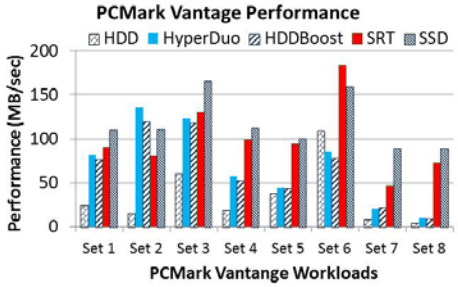

Figure 7 Benchmark Score for Each Set

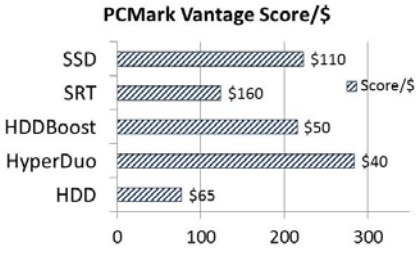

Figure 8 Benchmark Score/$

IV. Experiment

Figure 6 shows overall score of PCMark Vantage Benchmark and Figure 7 shows the score for each set of workloads. SSD and HDD is added in the benchmark test as well to show minimal and maximal expected score. On average HDD shows about 5000 and SSD shows about 25000. According to the overall benchmark score, SRT shows the best average performance of 20000, second HyperDuo with 11000, and HDDBoost with about 10000. SRT shows about 81% of performance of SSD and HyperDuo and HDDBoost shows about 46% and 44%, respectively. HyperDuo and HDDBoost shows better performance in set 2. For set 1, 5, 7, and 8, HyperDuo and HDDBoost shows similar performance. It is important to note that advertise performance differ greatly from the measured performance, which makes publicized benchmark scores unreliable.

PCMark Vantage score alone is a good measure to investigate the performance of different cache managers; however, cost of benefit is too important to be disregarded. Figure 8 shows PCMark Vantage Score per Dollar and price of the device excluding the price of the HDD or the SSD. Performance per Dollar of HDD (77) and SSD (223) is also shown to represent the minimal and maximal of expected performance per Dollar. Among all devices, HyperDuo shows best performance per Dollar and SRT comes in third place because of the price of the product.

Price of a cache manager alone can be somewhat misleading because the device alone can never improve the system performance. We set a budget as Eq. (1), which is benchmark score over total cost of using cache manager. C_{Cache} and C_{SSD} denote cost of the cache manager, SSD, respectively.

$$Budget = \frac{PCMark\ Vantage\ Score}{C_{Cache} + C_{SSD}} \quad (1)$$

Budget for HyperDuo, SRT, and HDDBoost is 76, 74, and 67, respectively. In terms of maintenance cost of the cache manager, HyperDuo holds a dominant position over SRT and HDDBoost.

V. Conclsion

Exploiting high performance SSD as cache storage for low performance HDD is becoming widely accepted solution to enhance performance of computing system. In this paper, we reviewed four different cache management solutions and compared advertised benchmark score with measured score. We found that the two scores do not conform to each other. We also define a term, budget, to measure the maintenance cost of the cache managers. We show that HyperDuo is better that other device in terms of the budget.

References

[1] D. Reinsel and J. Janukowicz, "Datacenter SSDs: Solid footing for growth," ed: IDC Corporation, Tech. Rep, 2008.

[2] J. C. Dave B. Anderson, Walter Fry, Rob Larsen, "Flash Cache Form Factors & HDD Economics," in *Flash Memory Summit 2010*, Santa Clara, CA, 2010.

[3] NVelo. (2010). Available: http://www.nvelo.com/

[4] Marvell. Available: http://www.marvell.com/

[5] SilverStone. (2003). Available: http://www.silverstonetek.com/

[6] Intel. Available: http://www.intel.com/content/www/us/en/homepage.html

[7] H. Payer, M. Sanvido, Z. Z. Bandic, and C. M. Kirsch, "Combo drive: Optimizing cost and performance in a heterogeneous storage device," in *WISH 2009 The First Workshop on Integrating Solid-state Memory into the Storage Hierarchy* Hamilton Crowne Plaza, Washington DC, USA 2009, pp. 1-8.

[8] E. Horvitz, "Machine learning, reasoning, and intelligence in daily life: Directions and challenges," in *Proceedings of Artificial Intelligence Techniques for Ambient Intelligence*, Hyderabad, India, 2007.

[9] K. Silver, "Realizing SSD Value through Caching Software," in *Flash Memory Summit 2010*, Santa Clara, CA, 2010.

[10] 04-Jan-2011, Marvell Introduces Advanced 6Gb/s SATA Controllers Powered By HyperDuo Technology. Available: http://www.prnewswire.com/news-releases/marvell-introduces-advanced-6gbs-sata-controllers-powered-by-hyperduo-technology-112860709.html

[11] Marvell. (2011, Marvell HyperDuo for 6Gb/s SATA Controllers: Automated SSD/HDD Tiering: 80% SSD Performance at 1/3 the Cost. Available: http://www.marvell.com/storage/system-solutions/assets/Marvell-HyperDuo-Product-Brief.pdf

[12] SilverStone. HDDBOOST Manual. Available: http://www.silverstonetek.com.tw/downloads/Manual/storage/Multi-HDDBOOST-Manual.pdf

[13] N. Edwards. (2010) Reviews of the Lates Hardware and Software: Silverstone HDDBoost. *Maximum PC*. p.86.

[14] Intel. Technical Brief: Intel Smart Response Technology. Available: www.intel.com/products/desktop/chipsets/ec-z68/z68-overview.htm

[15] M. Chiappetta. (2011, Intel Z68 Express with Smart Response Technology. Available: http://hothardware.com/printarticle.aspx?articleid=1676

[16] Futuremark. *PCMARK Vantage*. Available: http://www.futuremark.com/benchmarks/pcmarkvantage/introduction/