ISSN No. 0976-5697

Volume 4, No. 3, March 2013 (Special Issue)

International Journal of Advanced Research in Computer Science

REVIEW ARTICLE

Available Online at www.ijarcs.info

Review of Data storage by Fusion Drive in MAC

Pranav Nerurkar
IT SFIT, Mumbai University
Mumbai, India
pranavn91@gmail.com

Vinay Gentyala
IT ARMIET, Mumbai University
Mumbai, India
vag.vin@gmail.com

Abstract: In the age of technology and performance big players are working arduously to compete with evolving technology. Apple has always tried to remain ahead when it comes to technology, sophistication, and performance. Apple Inc recently released a product which was often referred to as 'a game changer' in the data storage market. A product called 'fusion drive' was promoted which is a fast and efficient storage solution as claimed by Apple Inc. It was therefore necessary to investigate these claims made by Apple Inc, to find whether it delivers on all enumerations. In this paper, we performed a study and review of this technology and cite its pro's and con's. Also, we have laid out certain experimental results to demonstrate each function of the fusion drive along with metrics for measuring its performance. In this paper, we investigate this new technology along with its various aspects and testify if it is better than the current data storage methods. Also, we present a simple example which displays the working of this technology and the concept behind it.

Keywords: Fusion Drives, Caching, Data Storage, Emerging Trends.

I. INTRODUCTION

As we all know Apple has been involved in inventing new range of technology that has been successful evidently. Apple is known for its IPods, IMac, Macbook, IPad, IPhone and its propriety software.IMacs and Macbooks run on Mac OS whereas IPods, IPhone run on Apple OS. Mac OS which is an excellent piece of software with a Crystal UI. Mac OS is incredibly fast when comes to performance. And to enhance its performance Apple Inc. has come up with a data storage technology called 'Fusion Drive', developed specifically to boost its application performance and boot times. Fusion Drive combines a high-capacity hard drive with high-performance flash storage. Fusion Drive automatically and intelligently manages your data so that frequently used apps, documents, photos, and other files stay on the faster flash storage, while infrequently used items move to the hard drive.

During a technological event held October 23, 2012, Apple announced the Fusion Drive technology which combines a hard drive with 128GB of NAND flash storage (solid-state drive) and presents it as a single logical volume with the space of both drives combined. The operating system automatically manages the contents of the drive so that as the users works on the Mac his most frequently accessed data is identified and stored on the faster flash storage, while infrequently used items move to or stay on the hard drive. For example, if the user is a fan of photography and accesses his photos regularly then his images, photo editing software, apps used are stored in a flash storage drive for faster access. The technology also boasts of a intelligent software that creates a logical volume and speeds up performance of the computer by performing both caching for faster writes and auto tiering for faster reads.

Due to similarity in features of Apples Fusion Drive and current Hybrid drives which incorporate a large amount of NAND flash memory inside a hard disk like an extended cache, people felt this technology is not new. Also parallels were being drawn to Intel's Smart Response Technology which adds a 64GB dedicated SSD as a supporting cache for a slow hard disk.

A study however reveals that the Fusion Drive is neither similar to Intel's Technology nor the traditional hybrid drives. Apple's Fusion Drive does not appear to function like an SSD-backed disk cache, but rather seems more like a file-level implementation of a feature that has existed for some time in big enterprise disk arrays: automatic tiering.

II. RELATED WORK

Traditional HDD's prolonged for a long period of time with not much of a change. Technical specifications such has increasing the cache memory and the rotating speeds were upgraded which did not bring a major change in the performance. Although, SATA HDD's were much faster than IDE disks to boost performance hardware up gradation was a must where users would purchase Ram's and CPU's which indeed is a big investment yielding no significant results in performance. HDD's were just utilized as a storage device no one really cared that HDD's would make a difference in the performance. Apart from super computers majority of the computers were slow in terms of performance. Only the Apple computers were fast enough which are overpriced due to their superior hardware and software containment. Only the intellectual users cared about access times file transfer and boot speeds. These users would tweak their Operating system to boost performance but still no substantial results.

Then SSD's arrived, a data storage device that uses integrated circuit assemblies as memory to store data persistently. As of 2010, most SSDs use NAND-based flash memory, which retains data without power. For applications requiring fast access, but not necessarily data persistence after power loss, SSDs may be constructed from random-access memory (RAM). Such devices may employ separate power sources, such as batteries, to maintain data after power loss. SSDs had origins in the 1950s with two similar technologies: magnetic core memory and card capacitor read-only store (CCROS). These auxiliary memory units (as contemporaries called them) emerged during the era of vacuum-tube computers. But with the introduction of cheaper drum storage units their use ceased. Later, in the 1970s and 1980s, SSDs were implemented in semiconductor memory for early supercomputers of IBM, Amdahl and Cray, however, the prohibitively high price of the built-to-order SSDs made them quite seldom used.

SSD's are still expensive than traditional SATA HDDs. At least the SSDs are into consumer market beating the traditional storage devices. SSD market has been growing since users realized that even SSDs being expensive they deliver performance. SSDs are now integrated into the new generation of laptops the "Ultra books" which are incredibly fast when it comes to performance since they started integrating SSDs with new generation of Intel processor technology. As of 2010, most SSDs use NAND-based flash memory, which retains data without power. For applications requiring fast access, but not necessarily data persistence after power loss, SSDs may be constructed from random-access memory (RAM).

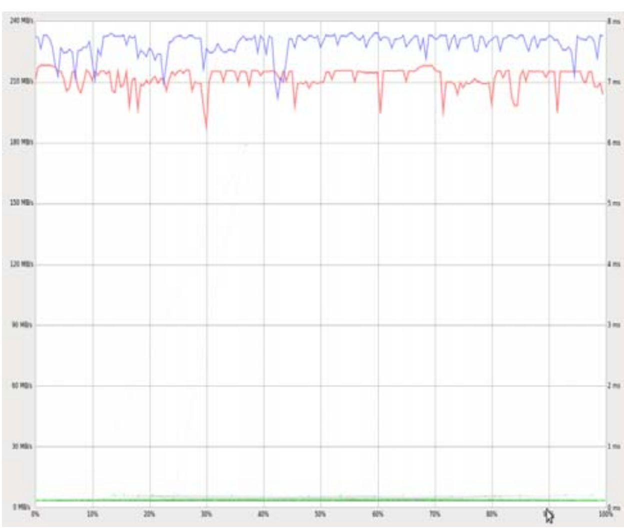

Figure: 1

SSD benchmark, showing about 230 MB/s reading speed, 210 MB/s writing speed and about 0.1 ms seek time, all independent from the accessed disk location.

Hybrid drives combine the features of SSDs and HDDs in the same unit, containing a large hard disk drive and an SSD cache to improve performance of frequently accessed data. These devices may offer near-SSD performance for many applications.

Traditional hard drives store their data in a linear, ordered manner. SSDs, however, constantly rearrange their data while keeping track of their locations for the purpose of wear leveling.

As such, the flash memory controller and its firmware play a critical role in maintaining data integrity. One major cause of data loss in SSDs is firmware bugs, which rarely cause problems in HDDs.

The following table shows a detailed overview of the advantages and disadvantages of both technologies.

Table: 1

Attribute or characteristic	**Solid-state drive**	**Hard disk drive**
Start-up time	Almost instantaneous; no mechanical components to prepare. May need a few milliseconds to come out of an automatic power-saving mode.	Disk spin-up may take several seconds. A system with many drives may need to stagger spin-up to limit peak power drawn, which is briefly high when an HDD is first started.
Random access time	About 0.1 ms - many times faster than HDDs because data is accessed directly from the flash memory.	Ranges from 2.9 (high end server drive) to 12 ms (laptop HDD) due to the need to move the heads and wait for the data to rotate under the read/write head.
Read latency time	Generally low because the data can be read directly from any location. In applications where hard disk seeks are the limiting factor, this results in faster boot and application launch times (see Amdahl's law).	Much higher than SSDs. Read time is different for every different seek, since the location of the data on the disk and the location of the read-head make a difference.
Data transfer rate	SSD technology can deliver rather consistent read/write speed, but when lots of individual smaller blocks are accessed, performance is reduced. In consumer products the maximum transfer rate typically ranges from about 100 MB/s to 600 MB/s, depending on the disk. Enterprise market offers devices with multi-gigabyte per second throughput.	Once the head is positioned, when reading or writing a continuous track, an enterprise HDD can transfer data at about 140 MB/s. In practice transfer speeds are many times lower due to constant seeking, as files are read from various locations or they are fragmented. Data transfer rate depends also upon rotational speed, which can range from 4,200 to 15,000 rpm.[72] and also upon the track (reading from the outer tracks is faster due higher absolute head velocity relative to the disk).

Consistent read performance	Read performance does not change based on where data is stored on an SSD	If data from different areas of the platter must be accessed, as with fragmented files, response times will be increased by the need to seek each fragment
Read/write performance symmetry	Less expensive SSDs typically have write speeds significantly lower than their read speeds. Higher performing SSDs have similar read and write speeds.	HDDs generally have slightly lower write speeds than their read speeds.

Hybrid drives combine the features of SSDs and HDDs in the same unit, containing a large hard disk drive and an SSD cache to improve performance of frequently accessed data. These devices may offer near-SSD performance for many applications.

III. MOTIVATION

Everyday new technologies are invented and more cutting edge features are designed and patented by Multi-National corporations in a bid to outwit each other. The main reason behind this competition is to increase the revenue of the Company. We feel that there is a need to study these technologies to ensure that they are really useful to the naïve users and will increase their experience as claimed by the manufacturing company. Hence we have compared this gadget with its competitors in the industry, to verify the tall claims made by Apple Inc. In announcing the Fusion Drive, Apple claimed that this new data storage solution was much better that then the methods used by its peers in the industry namely Intel and Kingston (see Figure 1). However, critics claim that this technology already existed in other architectures and is only called given a swanky name by Apple.

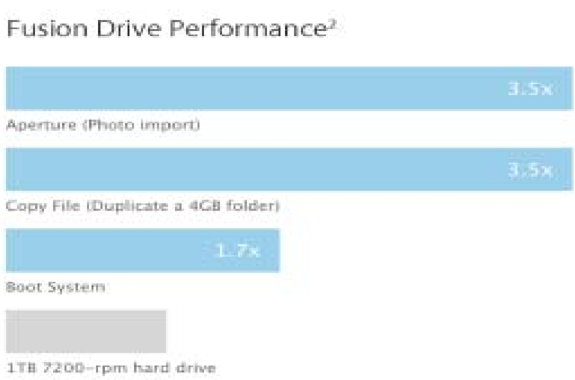

Figure 2: Fusion Drive performance as measured by apple

IV. METHODOLOGY

In OS X10.7 Apple introduced Core storage Logical volume manager. The core storage splits into two separate drives merged into a single volume. The volume manager handles the reallocation of files between the SSD and the HDD. Once the files are been placed in the SSD, file access becomes seemingly fast. Boot files are also placed in the SSD so that booting times are significantly faster. Core storage automatically puts the most used applications and boot (system files) files onto the SSD for performance boost.

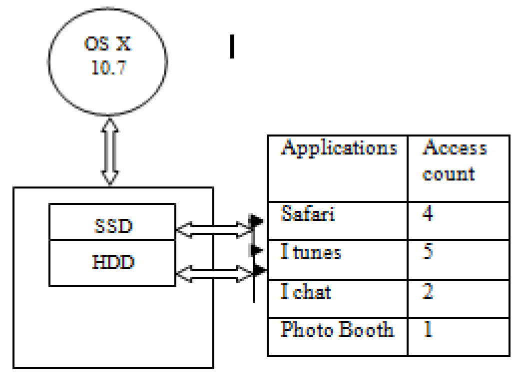

Figure: 3 Fusion Drive

The above figure (table) shows two most frequently used applications which are stored in SSD while least used applications are stored in HDD.

Most big disk arrays have different types of storage—some slow spinning disk, some faster spinning disk, and some solid state storage—and some have the ability to monitor what data is being accessed the most and can automatically move that data to a faster tier of disk as needed. These features typically operate at the block level, below the files, and can be done on large or small chunks of data, depending on what is frequently accessed and what isn't. Auto-tiering also includes the ability to take data that is no longer in demand, or no longer "hot" and demote it down off of fast disk and onto slower stuff. In this way, a file that doesn't get accessed very often might be stored on slow SATA disks, but if a hundred people need to open it repeatedly over a short period of time, it will get pulled up and kept on SSD until it's not needed anymore.

Based on Schiller's explanation, Fusion Drive sounds similar. In a caching solution, like Intel's, files live on the hard disk drive and are temporarily mirrored to the SSD cache as needed. In an enterprise auto-tiering situation, and with Fusion Drive, the data is actually moved from one tier to another, rather than only being temporarily cached there.

V. EXPERIMENTS

To prove the efficiency and of fusion drive in terms of speed that is promised by Apple Inc, it was necessary to carry out certain experiments on it.

A. *Automatic Tiering:*

Due to less details provided by Apples tech support manual of its underlying functionalities, it was difficult to understand how the automatic tiering was implemented in

CONFERENCE PAPER II International Conference on
"Advance Computing and Creating Entrepreneurs (ACCE2013)"
On 19-20 Feb 2013
Organized by
2nd SIG-WNs, Div IV & Udaipur Chapter , CSI , IEEE Computer Society Chapter India Council ,
IEEE Student Chapter Geetanjali Institute of Technical Studies, Udaipur, Rajasthan, India

fusion drive and who controls it. To find that out we performed an experiment:

A Fusion Drive consisted of a 128GB SSD connected to Mac's SATA bus and a 500GB HDD. Once that is done, using Core Storage we could create a single logical volume. Then using 'DISKUTIL' we create HFS+ volume inside it.

To check whether the experiment was successful we created 140GB of dummy files and directories on the volume using the "DD" command and the system automatically placed about 120GB of those on the SSD before storing the rest onto the HDD. After the files were all in place, we then triggered read activity on the volume, using the "DD" input file flag to constrain reads to the directories which had landed on the HDD.

'ISOSTAT' was used to monitor throughput of both the HDD and SSD at the device level. As soon as the reads stop, the file system is idle, SSD starts and a part of the writes are transferred to the SSD this can be observed by analyzing the write speeds. Now that the files are classified as frequently accessed they are transferred from HDD to SSD. This can be verified by again reading these files which lead to faster reads which means that the files now reside in SSD. Based on these findings, we can conclude that the main functionality either lies in the Core storage manager or the OS.

B. Block based or File based:

The next question was whether Fusion was block level or file level implementation of tiering. This was necessary to find out since operating at block level would have required it to move a document in its entire entirety from the HDD to SSD if it is found that even a small part of that document is frequently accessed. For that again using dd, we read the first megabyte of several 100MB files located on the HDD side of the Fusion Drive. After giving Fusion Drive some idle time to work, telling dd to read the entirety of the 100MB files, it generates significant IO on both SSD and HDD—the first megabyte of each file is coming off the SSD, and rest is coming off the HDD.

Thus, Fusion Drive is operating at the "sub-file" level, which means that it can overcome the disadvantages of file level implementations as their performance slows when large files are involved that have high level of change.

C. Results

The experiments prove that fusion drive is not a caching solution but a tiering solution. Instead of providing a copy of the data in the cache so that it can be accessed fast, it provides full transferring of data from the hard disk to the flash memory and this process is done when the CPU is idle. Hence it works in the background. This method is better than Caching technology because it avoids stealing CPU cycles for transferring so that the processor can do its housekeeping activities when the machine is idle.

VI. FUTURE SCOPE

Certain questions regarding fusion drive technology still remain unanswered. Regarding the nature of data movement – whether the sub files are promoted by being moved or are they simply copied.

VII. CONCLUSION

Though our experiments prove that fusion drive is a better technology than other caching solutions. It is not a new technology but a better implementation of tiering solution already present in big disk arrays. However, this concept was used only for large enterprises and Apple Inc. is the first company that chose to modify it and develop it as a solution for household users.

VIII. REFERENCES

[1] What is a Hybrid Drive? http://en.wikipedia.org/wiki/Hybrid_drive.

[2] Understanding Fusion Drive, "Apple's Latest Gadget". http://arstechnica.com/information-technology/2012/10/apple-fusion-drive-wait-what-how-does-this-work/.

[3] Fusion Drive Launch event October 23, 2012. http://www.anandtech.com/show/6406/understanding-apples-fusion-drive.

[4] What is a Fusion Drive and How does it work. "http://www.washingtonpost.com/business/technology/what-is-apples-fusion-drive-and-how-does-it-work/2012/10/25/9f86b330-1e08-11e2-ba31-3083ca97c314_story.html".

[5] Apple Inc. Technology event of October 23, 2012 http://www.apple.com/imac/performance/

[6] Mac mini (Late 2012) manual on fusion drives: https://support.apple.com/kb/HT5446?viewlocale=en_US&locale=en_US

[7] Creating a fusion drive: http://jollyjinx.tumblr.com/post/34638496292/fusion-drive-on-older-macs-yes-since-apple-has

[8] What is solid state drive? ttp://compreviews.about.com/od/storage/a/SSD.htm.

CONFERENCE PAPER
II International Conference on
"Advance Computing and Creating Entrepreneurs (ACCE2013)"
On 19-20 Feb 2013
Organized by
2nd SIG-WNs, Div IV & Udaipur Chapter , CSI , IEEE Computer Society Chapter India Council ,
IEEE Student Chapter Geetanjali Institute of Technical Studies, Udaipur, Rajasthan, India