Analyzing I/O Patterns for the Design of Energy-Efficient Image Servers

Jaemyoun Lee, Kyungtae Kang*
Dept. of Computer Science & Engineering, Hanyang University,
Ansan, Gyeonggi-do, Korea
Email: {jaemyoun, ktkang}@hanyang.ac.kr

Chang Song
Naver Labs
Bundang, Gyeonggi-do, Korea
Email: chang.song@navercorp.com

***Abstract*—Hard disks are one of the largest sources of power consumption in large-scale storage systems. The disk spin up/down technique has been shown to be an effective solution to this problem. Accordingly, the Open Compute Project (OCP) proposed a storage server for cold data, known as Cold Storage, to reduce power consumption using the spin-down technique in large-scale storage systems. With the aim of making effective use of Cold Storage, we characterize the power consumed by a hard disk in its various modes of operation. We then analyze the workload of an instant messaging service. These our contributions should provide guidelines for the implementation of a stable and energy-efficient distributed file system on a Cold Storage server, and to establish a spin-down policy that is power-proportional and promotes reliability.**

Keywords—cold data; power consumption; data center; instant messaging, green IT.

I. INTRODUCTION

The worldwide increase in the number of mobile devices running online applications, especially smart phones and tablets, has increased demand for large-scale storage systems [1]. Some service providers already have zettabyte data centers, where *zetta* is 1024 to the seventh power [2]. Billions of hard disks will be needed to cope with the explosive growth of data centers. However, these disks are one of the largest sources of power consumption in large-scale storage systems. With power consumption on the rise, concern is being raised over the growing implications in terms of energy bills, carbon emissions and the infrastructure requirements of data centers [3].

Current hard disks rotate faster, with higher storage density and lower error rates than previous designs, and there is little space to reduce their energy consumption. Several researchers have been working on means of replacing hard disks and optimizing power usage [4], [5], [6], resulting in development of Solid-State Disks (SSDs) comprising NAND flash memory. However, these disks unsatisfactory due to their higher cost and lower capacity compared to hard disks [4]. Thus hard disks remain useful, especially in conjunction with power-saving techniques that can be exploited by disk spin up/down solutions [3], [5].

The Open Compute Project (OCP) [7] has proposed a storage server for infrequently used data (i.e., cold data), which is called Cold Storage [8]. It uses the spin-down technique to reduce power consumption in large-scale storage systems. In a system of 30 hard disks, Cold Storage keeps only two are active at any time, and the other 28 are spun down. This has obvious potential for power saving, but the source code of this distributed file systems for power-aware, energy-efficient large-scale systems have not been published. Cold storage is an example of a system which is designed to match energy usage to I/O requirements; such power-proportional distributed file systems are an essential component of more energy-efficient data centers; however, the underlying energy consumption of a hard disk must be characterized before designing a file system. A big data workload is also investigated for the cold data stored in Cold Storage.

We present two estimations of extensive contributions that can be exploited as guidelines for designing power-proportional data centers: the trends in the power consumption of hard disks in various states, and the sophisticated investigation of a image sharing workload of a real instant messenger. The power consumed by a normal hard disk is significantly different from that consumed when the hard disk is spun down. Many researchers primarily focus on reducing the average power consumption; However, we focus on the peak power consumption in order to design a stable power-proportional hard-disk management system because burst requests are able to bring down entire data centers. Moreover, we have found that the spin-down technique is self-deceptive on SSDs. Because the spin-down technique increase the input/output (I/O) latency between a compute node and low-power storage systems, we intensively analyze a target workload and carefully design a proper hard disk management policy that improves the service quality. Our target workload is the instant messenger service LINE by LINE Corp. We recently discovered that the I/O request patterns to receive by storage servers from instant messenger services are entirely different from those of traditional instant messenger services because of smart phones and tablets. Whereas traditional instant messenger services have concentrated on text messages, modern instant messenger services have increasingly been dealing with text messages as well as photo messages. Our contributions can be of value as guidelines for implementing stable distributed file systems that optimize energy usage, and to establish a policy that is power-proportional and ensures reliability.

The remainder of this paper is organized as follows. In Section II we provide background on power-proportional storage servers and the spin-down technique, and give an example of a modern storage server workload. In Section III we describe the device that we developed to measure power consumption. In Section IV we outline the methodology used

*Corresponding author.

978-1-4799-7575-4/14/$31.00 ©2014 IEEE

in our experiments on hard-disk power consumption and in analyzing the resulting traces. In Section V we evaluate the experiments, and we conclude this paper in Section VI.

II. Background

A. Cold Storage

The OCP [7] was founded by Facebook in 2011 with the objective of replicating the concepts underlying open source software in order to create an *open hardware* movement to build commodity systems for hyper-scale data centers. The OCP aims to share more efficient server and data center designs with the general information technology industry, and consequently has published specifications for various storage servers. It has also proposed a revised version of an OCP storage server known as Cold Storage in order to satisfy the requirements for cold data. Cold Storage is designed to improve the energy efficiency of data centers by exploiting the spin-down technique [8]. Cold Storage comprises 30 hard disks in two trays. A rack contains 16 Cold Storages; thus, one rack contains 480 hard disks. Only one of the 15 hard disks on a Cold Storage tray is able to spin up at a given time; the others spin down to conserve power. In other words, only two hard disks in the Cold Storage are in operation at any one time [8]. Thus, power is naturally conserved because most of the hard disks are spun down.

However, even though use of the spin-down technique can significantly reduce the power consumption of data centers, an appropriate methodology for I/O scheduling is needed because hard disks can usually spin down only a limited number of times. Furthermore, pathological workloads can completely negate the power-saving benefits of the spin-down technique, prematurely causing a disk to exceed its duty cycle rating, and significantly increasing the aggregate spin up latency [5]. Although the OCP has published hardware specifications for Cold Storage, its file systems specifications have not been published. Other open source codes for file systems are also elusive because file systems are obviously a good commercial item. Thus, established policies for file systems that consider the overheads of the spin-down technique are expected to play an increasingly important role in the future.

B. Spin-down technique

The spin-down technique, which puts a disk into low-power mode while it is idle, is used to reduce hard disk power consumption. In low-power mode, such as *standby*, the spindle motor is not spinning and the heads are parked; hence, the power consumption is reduced. Researchers have proposed several spin-down algorithms that are very efficient at reducing the hard disk power consumption [9], [10], [5], [11]. These algorithms are typically time-out driven, spinning down the disk if the time-out expires before a request occurs.

The spin-down technique in the Power Management feature [12] provides energy-saving hard disks, reducing the power consumption, and the hard disk state is changed from *active* to *idle*, *standby*, or *sleep*. In the *idle* state, the operations that can be performed are more restricted than in the *active* state. However, in the *idle* state, the spindle motor of the hard disk is still spinning, and the head is also on platters. Consequently, the amount of hard disk power conserved is very small. In the *standby* state, the spindle motor of the hard disk is spun down and the head is also parked. Since the spindle motor is not in operation, the hard disk is not able to access data. Naturally, only a few operations can be performed, but the power consumption is substantially reduced. Hard disks typically consume 510 times more energy while in *active* mode compared to *standby* mode [13]. The *sleep* state is similar to the *standby* state, but only a reset operation can be performed: hard or soft reset. Thus, in theory, the *sleep* state consumes the lest power of all modes.

Recently the *idle* state was combined with the *active* state to create the Advanced Power Management feature [14]. This feature allows the hard disk to automatically change its state to either *active* or *idle*. To enter the *standby* and the *sleep* states, a special command must first be manually input. This command is specified in [12]. The hard disk state is returned to *active* when read or write operations occur when the hard disk is in the spin-down state. The actual design and implementation of a concrete power management feature are left to the drive vendor's discretion.

C. Server I/O pattern

Albrecht et al. [15] conducted an experiment involving thousands of Google users and applications and services such as content indexing, advertisement serving, Gmail, and video processing, as well as smaller applications, such as MapReduce jobs owned by individual users. A large application may contain many component jobs. The workload characteristics and demands of the jobs in data centers are typically highly varied between users and jobs. As a result of the variation of mean read age over different jobs in Google's data centers, the mean read age of the bytes read over 15000 jobs is approximately 30 days, even though jobs access very young (one minute old) to very old (one year old) data. Another experiment has shown that 50% of the data stored by a particular user are less than one week old, but correspond to more than 90% of the read activity.

Parikh [16] discussed the necessity for Cold Storage in Facebook's data centers. He argued that 2.8 ZB (zetta bytes) of data were created in 2012, and that 40 ZB of data will exist by 2020 in the world. To store all of these data, data centers require billions of hard disks, with each at the current maximum capacity (4 TB). Since hard disks consume a significant amount of power, 153 million kilowatt hours of power were consumed by a single Facebook data center in 2012. This equates to as much power as in used in 13,000 homes [17]. However, most data such as photos are hot when they are created, but decrease in relevance over time, becoming warm. Eventually, the data change to cold and reads are hardly requested. More specifically, 82% of read traffic is serviced for only 8% of young photos in Facebook's data centers.

These results indicates that the continually increasing requirements for cold data stored on disks but almost never read, such as legal data or backups of third copies of data. Consequently, a tier system that separates data into hot, warm, and cold storage has been proposed. Furthermore, empirically, aged data is likely to become cold data.

Finally, Thereska et al. [18] discussed the I/O patterns of instant messenger services. Modern instant messenger ser-

Fig. 1. Implemented power measurement device

Fig. 2. Measurement results of a hard disk

vices have significantly different I/O patterns from those of traditional instant messengers. Concomitant with the increase in mobile devices with high quality cameras, the demand to exchange photos and videos in chat has increased. Therefore, the I/O trace data of the LINE service's photo servers must be analyzed to determine the workload characteristics of modern instant messenger services.

III. Hard Disk Power Consumption Measurement Device

Research related to the spin-down technique [19], [20], [5] is predominantly focused on average power consumption, with no regard for instantaneous power consumption. In the experiments conducted, however, we observed that the graph of hard disk power consumption first increases steeply then flattens out when that hard disk state changes from *standby* to *active*. The peak power consumption is five times greater than the average power consumption. Thus, the instantaneous power consumption has to be considered when designing a spin-down scheduler.

To effectively measure only hard disk power consumption, we developed a hard disk power consumption measurement device comprising an electric current shunt and power monitors. Since a 3.5 inch hard disk is supplied with both 12 V DC and 5 V DC, both voltages and currents have to be measured for precise hard disk power consumption. In addition, several hard disks have to be measured because our aim is to reduce the power consumption of storage servers. However, the devices currently on the market are inadequate for this work in terms of cost, effectiveness, and functionality. The measurement device that has eventually been developed is very durable because the experiments are conducted with real workloads over several days.

The INA226 by Texas Instruments [21] is used in the measurement device. The INA226 monitors both the shunt voltage drop and the bus supply voltage. The programmable calibration value, conversion times, and averaging, combined with an internal multiplier, enable direct readouts of the current in amperes and the power in watts. The INA226 detects current on buses that can vary from 0 V to 36 V, while the device obtains its power from a single 2.7 V to 5.5 V supply, typically drawing 330 µA of supply current. The INA226 is specified over the operating temperature range of −40 °C to 125 °C. The I^2C interface features 16 programmable addresses.

IV. Experiments

To execute the spinning down of the hard disks, we used Hdparm on Ubuntu 14.04, as shown in Table I. Hdparm is a command line utility for Linux that sets and views ATA hard disk drive hardware parameters. It can set parameters such as the drive caches, sleep mode, power management, acoustic management, and Direct Memory Access (DMA) settings [22]. We conducted experiments on both the *standby* and *sleep* states of the hard disk. However, we focused only on hard disks

TABLE I. Experimental conditions of spin-down technique on hard disk and SSD

Processor	Intel Core i3-2100 Processor (3 M Cache, 3.1 GHz)
RAM	DDR3 8 GB
Operating system	Linux 3.13.0 (Ubuntu 14.04)
Hard disk	Seagate Barracuda ST1000DM003
SSD	Samsung MZ-7PC256

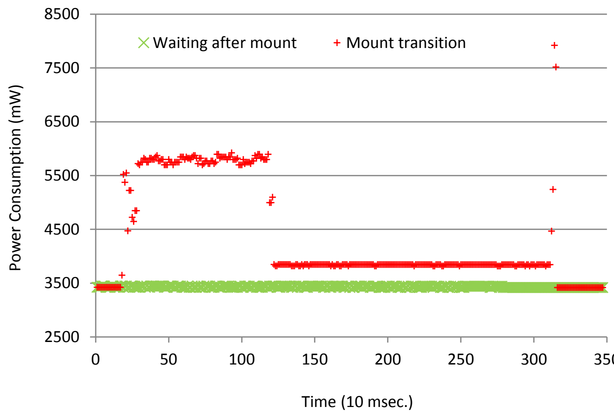

Fig. 3. Power consumption when hard disk is mounted (a green cross) and executing a mount command (a red plus). The mount command was carried out at 170 ms, and completed at 1250 ms.

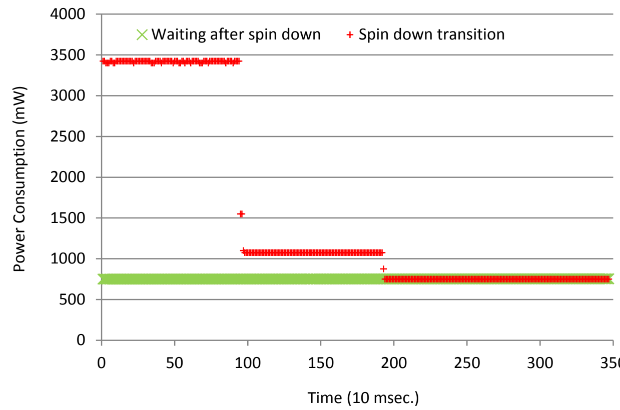

Fig. 4. Power consumption when hard disk is spun down (a green cross) and executing a spin-down command (a red plus) The spin-down command was carried out at 940 ms, and completed at 1960 ms.

with the *standby* state because the power consumption trends are virtually indistinguishable. In addition, we measured flash memory-based SSD power consumption. SSDs use integrated circuit assemblies as flash memory to store data without the spindle motor and the head used by hard disks. As a result, the SSD power consumption measurement shows a different pattern compared to hard disks.

The explosive growth of mobile devices such as smart phones and tablets that display text as well as images, has resulted in a change in the I/O pattern of instant messenger services in large-scale distributed storage servers. Therefore, analysis of the I/O pattern of current instant messenger services is needed. We used a real workload from the LINE instant messenger by LINE Corp. for an objective analysis. LINE users exchange text messages, graphics, video, and audio media; make free VoIP calls; and hold free audio and video conferences. LINE achieved more than 400 million registrants worldwide in 2014 [23]. We traced the I/O requests to LINE storage servers for images over seven days. LINE has 20,000 servers, including more than 10 Redis clusters and 10 HBase clusters. The logs in the trace data numbered in the billions. Accordingly, this workload is sufficient to simulate large-scale storage systems for instant messenger services. First we used R project to analyze this workload, but the ruinous data caused a memory overflow problem, even though the machine used contains two Xeon processors and 24 GB RAM. Consequently, we used MySQL, a widely used open source relational database management system appropriate for analyzing this workload. Its transaction processing speed is also correspondingly fast.

A. Measurement of hard disk power consumption

We measured the hard disk power consumption for various scenarios: waiting before mount, waiting after mount, the moment at which the unmount Linux command is executed, the moment at which the mount Linux command is executed, waiting after spinning down, the moment at which the hard disk is spinning down, and the moment at which the hard disk is spinning up. The measurements for the hard disk showed that it consumed the same amount of power when waiting after the mount as before the mount. It also consumes the same amount of power during execution of the mount and unmount Linux commands. Therefore, in Figure 3, we present only the results for the hard disk power consumption when waiting after the mount and the moment at which the mount Linux command is being executed. Note that there may be differences according to the particular hard disk manufacturers.

The hard disk consistently consumes 3425 mW while waiting for a command. It consumes 5699.3 mW for 1.08 s on average after receiving either the mount or the unmount command, and the power consumption is generally maintained at 3850 mW until the hard disk changes its state to *idle*. The time interval between the command and the state change is 1.8 s. At singularity, the peak power consumed by the hard disk is 2.3 times larger than the average for a fleeting moment when the state is changed. Figure 4 shows the results obtained when the hard disk is spinning down and the waiting power for the hard disk spun down. The spindle motor of the hard disk is able to stop within 1020 ms and then consumes 1079.5 mW on average. Thus, the hard disk spin-down scheduler does not have to consider spinning down all of the hard disks during a burst. After the hard disk has spun down, it consumes 700 mW, which is five times smaller than the average of the *active* hard disk. Thus, it is quite clear that the spin-down technique reduces the power consumed by large-scale storage servers. In addition, because the spin-down technique reduces the hard disk operating temperature, data center cooling systems, which are a primary target for energy efficiency improvements, can be scaled down. The spin-down technique can significantly reduce the cost of both establishing and operating data centers.

In the next experiment, we measured the hard disk power consumption when it was in a *standby* state or spinning up, as shown in Figure 5. The result represents the modern disk scheduling policy in a manner that reduces the electric power required to operate and have to focus on burst requests for spinning up. The operating time of the spinning up command is 4.3 s on average and is directly related to the overhead of the spin-down technique. In other words, when a hard disk that is in the *standby* state is requested by any command, the access latency is more than 4.3 s. Users have to wait longer than this time because of the network transfer time and the internal

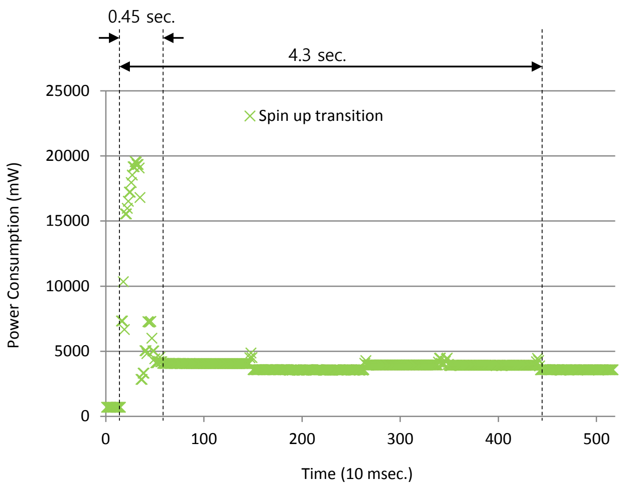

Fig. 5. Power consumption when hard disk is spinning up. The spin up command was carried out at 150 ms, and completed at 4450 ms.

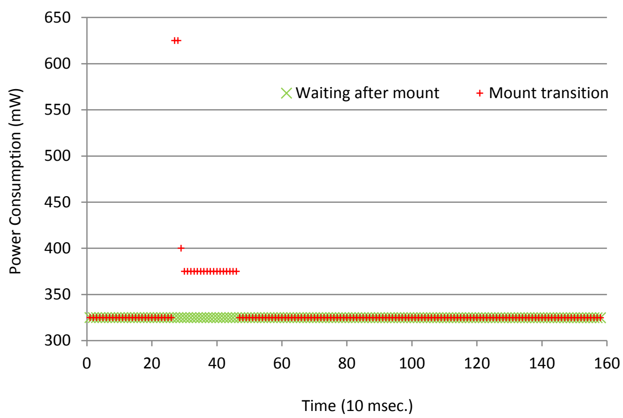

Fig. 6. Power consumption when SSD is mounted (a green cross) and executing a mount command (a red plus). The mount commands was carried out at 260 ms, and completed at 480 ms.

process time. This latency has repercussions for the overall service quality. Therefore, the hard disk operation scheduler must prudently choose which hard disk is spun down.

When a hard disk is spinning up, it consumes 5.8 times the average power for 0.45 s. This amount of power consumption is immense compared with those of other commands or other devices in a computer. If several hard disks were to spin up at the same time, the computer power supply would fail to meet with the demand and would be unable to deliver a stable power supply to the disks as well as the processors. Moreover, this results in accidents such as hardware failure and electrical shorts. Therefore, large-scale storage systems must take all reasonable precautions to protect against simultaneous spin up.

We also measured the SSD power consumption in order to compare it to that of hard disks. The SSD that was used consumes 10.5 times less power than the hard disk when waiting for any commands, as shown in Figure 6. This reduced amount is greater than when the hard disk is spun down. When the SSD is required to spin down, it carries out an action that

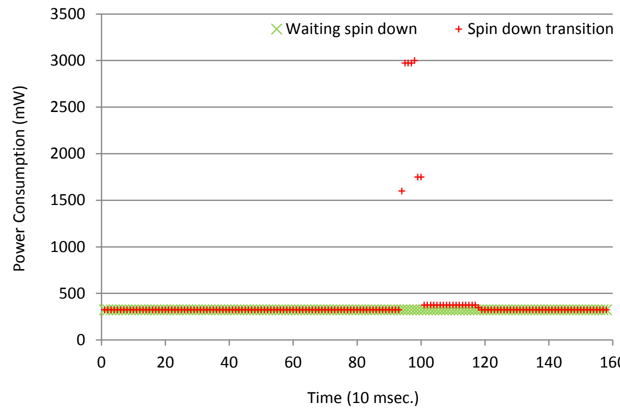

Fig. 7. Power consumption when SSD is spun down (a green cross) and executing a spin-down command (a red plus) The spin-down command was carried out at 930 ms, and completed at 1190 ms.

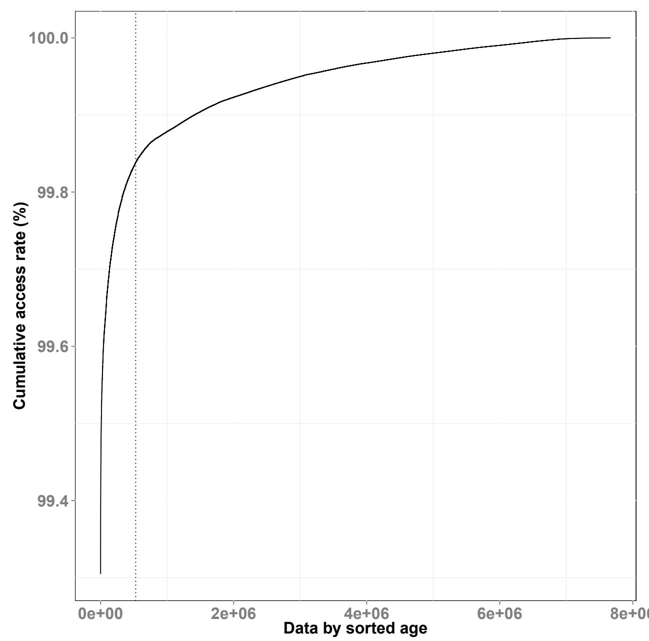

Fig. 8. Cumulative distribution function of the read and write operations sorted by the age of the data. The data on the left are younger than those on the right. The vertical dotted line means that 8% of the young data represents 99.85% of the traffic.

we are not able to measure because it depends on the drive vendor's discretion. Furthermore, the power consumed by the SSD after spinning down does not change, as shown in Figure 7. Thus, the spinning down on an SSD is a waste of electric power and is counterproductive because the amount of power consumed is five times larger than average. Accordingly, the SSD is permitted to receive a spin-down command because of compatibility, and the spin-down technique and the SSD negatively impact each other's performance.

B. Analysis of a real workload

We conducted experiments on a real workload using the trace data of I/O requests in order to image servers on a real

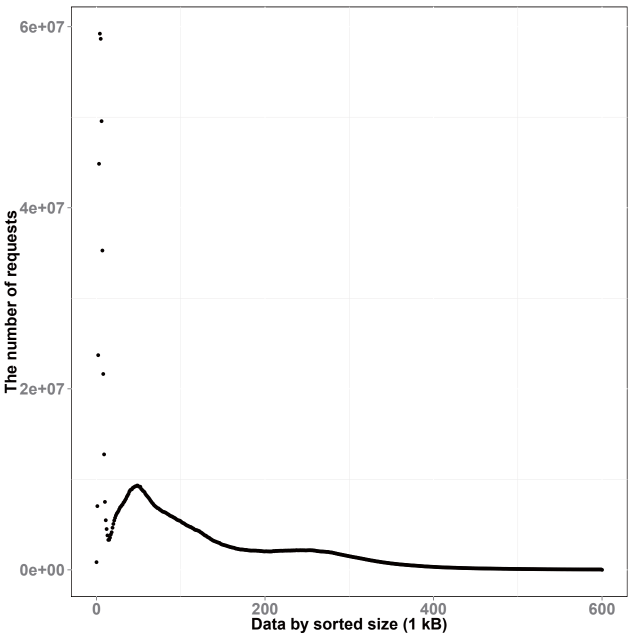

Fig. 9. I/O requests counting histogram sorted by data size. Only data sizes less than 600 KiB are represented because the data larger than 600 KiB are requested at most four times

instant messenger service called LINE. The trace data were obtained over seven days and included not only image files, but also the meta-data of file systems. Sophisticated analysis and scrupulous examination enabled us to determine the difference between a traditional instant messenger workload and that of a modern instant messenger.

First, we sorted the workload by file age and calculated the access rate in order to determine the I/O patterns, as shown in Figure 8. The horizontal axis is the number of files sorted by file age and represents just one of 100 files, because the amount of data is very large. According to the result, 8% of the young data represented 99.85% of the I/O traffic. In other words, when the data are young, they have a high probability of being read, and old data are never read or are unlikely to be read. These results convincingly demonstrate that most of the data stored in servers are cold data. Thus, if the cold data are migrated to power-proportional storage servers, an extraordinary amount of energy can be conserved with minimal performance degradation. The performance degradation only occurs when the data stored in a hard disk in the *standby* state is requested. However, large-scale storage systems have high-fidelity cache systems and instant messenger applications also have cache space. The cache technology helps power-proportional storage systems out of routing I/O requests to hard disks in the *standby* state. Note that the workload was traced with I/O requests occuring before entrance into cache systems.

To intelligently identify cold data, we conducted additional experiments. We sorted the data by file size and grouped the files with the same size into sets. We then counted the number of requests in the sets. The results are shown in Figure 9. The unit on the horizontal axis (file size) is 1 KiB because the histogram is too dispersed if the file size unit is one byte. The majority of requests occur in the 4 ˜ 6 KiB range. This size corresponds to the profile thumbnail image size in the instant messenger application. Hence, this set is usually for thumbnail images. The I/O patterns suggest that this set is only marginally dependent on file age, which means that old files in this set are sometimes requested by a read command. Therefore, it is better that files in this set are not separated as cold data, even though the files are old. File sizes in the range 40 ˜ 50 KiB are the second most requested. Most of the files in this set are for normal image messages. Although the actual size of the image file is larger than the specified range, the instant messenger application usually transmits the compressed image file to storage servers at a low resolution. Therefore, the size of the file stored in storage servers is less than that of the original image file. Normally the popular area is about 50 KiB. The I/O pattern in this area shows that the number of requests is maintained at a proper number. The proper number seems to be the number of people in a chat room. Thus, these files are best extracted as cold data. The files that are larger than 200 KiB are large media files such as video and voice messages. These files are also better to extract as cold data. Even though the exact file size are different among instant messenger applications, all modern instant messenger services provide thumbnail images, image messages, and video messages. Consequently, similar I/O patterns can be observed in any instant messenger services.

V. Evaluation

Our results demonstrate that only a very small amount of data needs to be stored in hot storage servers, because much of the data is cold. Compared with [16], the access rate of young data is extreme in an instant messenger service. Parikh [16] indicated that 8% of young data represents 82% of the traffic; however, the same 8 % of instant messenger workload represents 99.85% of the traffic. Upon further inspection, we determined that the instant messenger service has contains massive amounts of cold data because the persons in chat rooms of the instant messenger service are fixed. If someone sends an image to others, the number of images sent is equal to the number of others. In addition, an image that is sent once does not need to be sent again because of the application cache. The instant messenger server does not need to access an image after it is sent a few times. As a result, the access rate of files decreases for older files. For this reason, the instant messenger workload has extremely cold data. This feature only applies to modern instant messenger services. This I/O pattern is not shown in traditional instant messenger I/O patterns [18].

To reduce the power consumption of storage servers, the servers should be operated according to a special policy because the instant messenger workload is very distinctive. There are other papers that concentrate more generally on energy efficiency and low-level dynamic power management. However, this work focuses on the data age and the access rate on workloads without considering the number of requests. As described above, the number of requests is an important fact in the instant messenger workload and can be more than the data age. All of these facts suggest that the policies of power-proportional storage systems should be customized to consider the data age as well as the number of requests.

The proposed policy should also be considered for hard disk spin up scheduling. When hard disks are spinning up, they consume a lot of power, as described in Section 4. Although the timing of consumption for a significant quantity of electric power is instantaneous, it may causes damages to large-scale storage systems because the peak power is greater than five times that of the norm. Moreover, server racks need to have a high-capacity power supply in order to provide sufficiently stable electric power into storage servers. However, utilization of a high-capacity power supply is very low because high-capacity is not needed at other times. Therefore, the policy has to address the individual spin up of a hard disk. The time during which a hard disk consumes peak electric power is just $0.45\,\text{s}$, whereas a hard disk is spinning up for $4.3\,\text{s}$. Consequently, a simple sequential spin up of a hard disk causes substantial performance degradation. To avoid this degradation, the scheduler should receive feedback control with an electric power measurement. At present, passing the peak electric power, the scheduler can spin up other hard disks. These features make the scheduler ideally suited for managing hard disks in extremely large-scale storage systems where low-power, performance, and reliability are critically important.

This evaluation serves as a guideline to researchers in the field of energy-efficient large-scale storage systems, while providing a structured exposition and discussion of current hard disk low-power technology and modern instant messenger workload. Exact power consumption values and access rates of the workload vary on a case by case basis. However, the trends are the same under all circumstances. Therefore, the proposed guidelines are useful for power-proportional large-scale storage systems.

VI. Conclusions

Current data centers require energy-efficient and power-proportional storage servers, which in turn require energy-efficient storage devices. SSDs consume a small amount of power; however, they are more expensive and have less capacity than hard disks. Hence hard disks are still useful and their power requirement can be reduced by the spin-down technique, which is used in large-scale storage systems like Cold Storage. We measured hard-disk power consumption in a variety of situations, and analyzed the real workload of an instant messaging, with the aim of improving the effectiveness of Cold Storage and minimizing the overhead of the spin-down technique.

The measurement results indicate that a hard disk uses a lot of power to spin up. Our analysis of a modern instant messaging service indicates that huge amounts of cold data have to be stored; and the age of the data is important as well as the number of I/O requests. Our contributions can be used as guidelines to implement distributed file systems that are stable and optimizes energy usage, and to establish a policy that is power-proportional and promotes reliability.

Acknowledgment

This research would not have been possible without the help of Naver Labs and LINE Corp. The authors would like to thank Soonho Kwon for his advices and active support for this project. Juyoung Park and Kuyeon Lee also contributed ideas. We are grateful to Goun Choe for helping with experiments and analysis.

This research was supported in part by the MSIP (Ministry of Science, ICT and Future Planning), Korea and Naver Corp., under the IT/SW Creative Research Program supervised by the NIPA (National IT Industry Promotion Agency) (NIPA-2013-H0505-13-1003), and in part by the Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the MSIP (NRF-2013R1A1A1059188).

References

[1] B. Guenter, N. Jain, and C. Williams, "Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning," *Proc. 30th IEEE International Conference on Computer Communications*, 2011, pp. 1332–1340, 2011.

[2] S. Baboo and R. Kumar, "Next generation data warehouse design with big data for big analytics and better insights," *Global Journals Inc.*, vol. 13, no. 7, 2013, pp. 19–23, 2013.

[3] P. Llopis, J. G. Blas, F. Isaila, and J. Carretero, "Survey of energy-efficient and power-proportional storage systems," *The Computer Journal*, 2013; doi:10.1093/comjnl/bxt058.

[4] J. Lee, S. Han, J. Ryu, and K. Kang, "Strategy to reduce I/O access time of applications in virtual machines," *Proc. 2013 International Conference on IT Convergence and Security*, 2013, pp. 1–2.

[5] T. Bisson, S. Brandt, and D. D. E. Long, "A hybrid disk-aware spin-down algorithm with I/O subsystem support," *Proc. IEEE International Performance, Computing, and Communications Conference*, 2007, pp. 236–245.

[6] K. Bu, M. Wang, H. Nie, W. Huang, and B. Li, "The optimization of the hierarchical storage system based on the hybrid ssd technology," *Proc. 2nd International Conference on Intelligent System Design and Engineering Application*, 2012, pp. 6–7.

[7] "Open compute project"; www.opencompute.org.

[8] M. Yan, "Open vault storage hardware V0.5 ST-draco-abraxas-0.5," Open Compute Project, Jan. 2013.

[9] F. Douglis, P. Krishnan, and B. N. Bershad, "Adaptive disk spin-down policies for mobile computers," *Proc. 2nd Symposium on Mobile and Location-Independent Computing*, 1995, pp. 121–137.

[10] Y.-H. Lu and G. de Micheli, "Adaptive hard disk power management on personal computers," *Proc. 9th Great Lakes Symposium on VLSI*, 1999, pp. 50–53.

[11] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "DRPM: Dynamic speed control for power management in server class disks," *ACM Special Interest Group on Computer Architecture News*, vol. 31, no. 2, pp. 169–181, 2003.

[12] American National Std. T13/2015-D, *ATA/ATAPI Command Set - 2 (ACS-2)*, Western Digital Technologies Inc., 2011.

[13] J. Zedlewski, S. Sobti, N. Garg, F. Zheng, A. Krishnamurthy, and R. Wang, "Modeling hard-disk power consumption," *Proc. 2nd USENIX Conference on File and Storage Technologies*, 2003, pp. 217–230.

[14] *Power management on adaptec unified serial raid controllers and hgst deskstar hard drives*, HGST, a Western Digital company, 2008.

[15] C. Albrecht, A. Merchant, M. Stokely, M. Waliji, F. Labelle, N. Coehlo, X. Shi, and E. *Proc. USENIX Annual Technical Conference*, 2013, pp. 91–102.

[16] J. Parikh, "Cold Storage - Jay Parikh in Open Compute Summit IV"; http://new.livestream.com/accounts/2462150/events/1790124/videos/9502681.

[17] M. Rogoway, "Facebook: Power use in Prineville data center more than doubled last year"; www.oregonlive.com/siliconforest/index.ssf/2013/07/facebook_power_use_in_prinevil.html.

[18] E. Thereska, A. Donnelly, and D. Narayanan, "Sierra: Practical power-proportionality for data center storage," *Proc. 6th Conference on Computer Systems*, 2011, pp. 169–182.

[19] J. Chou, J. Kim, and D. Rotem, "Energy-aware scheduling in disk storage systems," *Proc. 31st International Conference on Distributed Computing Systems*, 2011, pp. 423–433.

[20] T. Bostoen, S. Mullender, and Y. Berbers, "Analysis of disk power management for data-center storage systems", *Proc. 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, 2012, pp. 1–10.

[21] Texas Instruments Incorporated, "High-Side Measurement, Bi-Directional Current/Power Monitor w/I2C Interface," Dallas, Texas, 2011

[22] "Sourceforge.net: hdparm - project web hosting - open source software"; http://hdparm.sourceforge.net.

[23] "LINE now has 400 million registered users! : LINE official blog"; http://officialblog.line.me/en/archives/1001168967.html.