論文 / 著書情報
Article / Book Information

Title	A Power Saving Storage Method That Considers Individual Disk Rotation
Author	Satoshi Hikida, Hieu Hanh Le, Haruo Yokota
Journal/Book name	Database Systems for Advanced Applications Lecture Notes in Computer Science, Volume 7239/2012, , pp. 138-149
発行日 / Issue date	2012, 4
DOI	http://dx.doi.org/10.1007/978-3-642-29035-0_10
権利情報 / Copyright	The original publication is available at www.springerlink.com .
Note	このファイルは著者（最終）版です。 This file is author (final) version.

A Power Saving Storage Method That Considers Individual Disk Rotation

Satoshi Hikida, Hieu Hanh Le, and Haruo Yokota

Department of Computer Science, Tokyo Institute of Technology
2-12-1 Ookayama, Meguro-ku, Tokyo 152-8552, Japan
{hikida, hanhhlh}@de.cs.titech.ac.jp, yokota@cs.titech.ac.jp

Abstract. Reducing the power consumption of storage systems is now considered a major issue, alongside the maintenance of system reliability, availability, and performance. In this paper, we propose a method named as the Replica-assisted Power Saving Disk Array (RAPoSDA) to reduce the electrical consumption of storage systems. RAPoSDA utilizes a primary backup configuration to ensure system reliability and it dynamically controls the timing and targeting of disk access based on individual disk rotation states. We evaluated the effectiveness of RAPoSDA by developing a simulator that we used for comparing the performance and power consumption of RAPoSDA with Massive Arrays of Inactive Disks (MAID), which is a well-known power reduction disk array. The experimental results demonstrated that RAPoSDA provided superior power reduction and a shorter average response time compared with MAID.

Keywords: storage, power saving, performance, large scale, reliability

1 Introduction

Ongoing increases in the total electricity consumed by data centers present significant problems that must be solved to reduce the running cost of centers and to keep the global environment green. A governmental report estimated that the electricity consumption of US data centers in 2011 would be double that of 2006 [11].

Servers and storage systems are the two main data center components that consume electricity. They tend to grow at a numerically faster rate, because of the increasing requirements for processing load and the targets of processes. The amount of data stored in storage systems is increasing particularly rapidly, because of the explosive increase in the volume of data generated by the Internet. The IDC reports that the growing amount of digital data created and replicated will surpass 1.8 zettabytes in 2011 [7]. Many services are leveraged by cloud computing, which is becoming widespread in our daily life and business, including social network services, data sharing services, and movie sharing services. These services require large-scale storage, which is characterized by writing new data and reading recent data, while only a few read old data. In this paper, we focus on a power saving method for large-scale storage systems that involve time-skewed data access.

Many methods have been proposed to reduce the power consumption of storage systems. A typical approach is to spin down some of the hard disk drives in a storage

system, because rotation of the platters in a disk drive accounts for most of the electricity consumption by a storage system. Thus, the electricity consumption of a disk is decreased by spinning it down.

A well-known system that uses this approach is MAID (Massive Arrays of Inactive Disks) [3]. MAID keeps a small number of disk drives rotating at all times, which are used as a cache (known as cache disks), and this allows other disk drives to spin down (known as data disks). The MAID approach is effective for reducing power consumption when data access is time skewed, because numerous data disk drives store infrequently accessed old data and this allows disk rotation to be suspended. However, the original version of MAID did not consider system reliability. Thus, data stored on a disk are lost if one of data disks fails in the MAID system.

If we introduce a simple replica-based fault-tolerant configuration into MAID, synchronization between replicas during update operations would reduce the effect of saving power consumption because an increased number of write operations on disks elevates the number of spinning disks. In contrast, if we choose to improve timing when applying replica updates, we can reduce power consumption and maintain system reliability. In this paper, we propose a Replica-assisted Power Saving Disk Array (RAPoSDA) as a method for dynamically controlling the timing of disk write operations by considering the rotation status of disks storing replicas.

RAPoSDA employs a duplicated write cache memory to maintain reliable data and it waits for a suitable time to write to an appropriate disk, thereby avoiding unnecessary disk spin-ups. We assume that there is an independent power supply (UPS) for the duplicate cache memory to ensure tolerance of power shortages or a UPS failure. Data on data disks are replicated in RAPoSDA as a primary backup configuration to ensure system reliability. RAPoSDA also has cache disks that provide large read cache spaces, as found with MAID, but they do not need to be replicated. Cache disks are an optional feature in RAPoSDA if there is sufficient cache memory.

The main contribution of this approach is the introduction of careful control over timing and targeting of disk access by replicas, which is achieved by dynamically checking the rotation status of individual disk drives. To ensure system reliability and reduce power consumption, RAID-based low power consumption storage systems [9, 8] and power proportional disk arrays using replicas [14, 13, 1] have also been proposed. However, those other methods do not check the rotation status of each disk drive.

A simulation was more appropriate than empirical experiments and disks when evaluating the performance and power consumption of a storage system with changing configurations, including various sizes of cache memory, disk capacity, number of disks, and particularly with large numbers of disks. Therefore, we developed a dedicated simulator to compare RAPoSDA with MAID and a simple disk array with no power saving mechanism. We also prepared synthetic workloads with skewed data access based on a Zipf distribution and different ratios of read/write requests. The experimental results using the simulator indicated that RAPoSDA provided superior power reduction and a shorter average response time compared with MAID.

The remainder of this paper is organized as follows. Section 2 describes the power consumption model of a disk drive and a general approach to power reduction for a storage system with disk drives. Section 3 describes the details of our proposed stor-

age system, RAPoSDA. Section 4 outlines the simulator we developed to evaluate the performance and power consumption of a storage system. A detailed discussion of our evaluation of RAPoSDA and MAID in terms of power saving and performance is reported in Section 5. Section 6 presents a number of related studies. Finally, Section 7 provides the conclusions of this paper.

2 Disk Drive Power Consumption Model

A disk drive is mainly composed of mechanical and control parts. The mechanical part has four components: platters, a spindle motor, read and write heads, and an actuator to move the heads over the platters. The controller is a component of the control part that performs read and write operations to comply with requests. Of these components, the spindle motor, actuator, and controller consume the most electricity in a disk drive, depending on its state.

A disk drive has three states, as shown in Table 1. It consumes the most power when it is in an active state, because the spindle motor, actuator, and controller are all working. It consumes the least in the standby state, because all components are not working, whereas only the spindle motor is working in the idle state. The standby state transfers to the active state after a read or write request. The idle state transfers to the standby state via spin-down, and vice versa via spin-up.

Compared with the active state, a disk drive consumes larger amounts of electricity when it spins up or down, especially at the beginning of spin-up. This means that frequent spin-ups and spin-downs to provide short standby states are not good for power saving. Thus it is important to keep unaccessed disk drives in a standby state for as long as possible. The duration of a standby state where a spin-up and spin-down requiring energy consumption does not exceed that of the idle state is known as the break-even time. A power saving is produced only if the duration of the standby state is longer than the break-even time.

Table 1. Disk-drive Status and Corresponding Power Consumption

State	I/O	RPM	Head location	Power
<i>Active</i>	In operation	max rotation	on disk	Large
<i>Idle</i>	No operation	max rotation	on disk	Middle
<i>Standby</i>	No operation	0	off disk	Small

3 RAPoSDA

One approach for reducing power consumption in a storage system with multiple hard disk drives is to divide the disk drives into two groups. Disk drives in one group are

Fig. 1. Configuration of RAPoSDA

mainly kept in the standby state for longer than the break-even time, while disk drives in the other group are kept in active or idle states by access localization.

MAID (Massive Arrays of Inactive Disks) [3] is a well-known system that takes this approach. MAID ensures that a small number of disk drive are always kept in idle or active state as cache disks, while the remaining large number of disks are kept in a standby state as data disks. All read and write requests are initially dealt with by the cache disks and a replacement algorithm is used to keep the hit ratio high. This approach is effective in reducing the power consumption of a storage system. However, the original MAID method had problems of reliability, which is a major requirement for data centers. There is no means of recovering data in the event of a data disk failure. Direct importation of replicas into MAID might make it more reliable, but it would have detrimental effects on power consumption because the increased disk access required for replicas might violate the break-even time restriction.

To solve this problem, we propose a method for controlling the timing and targeting of replica disk access by dynamically checking the rotation status of individual disk drives. This reduces the power consumption and maintains system reliability. To implement the method, we propose the RAPoSDA storage system configuration.

3.1 RAPoSDA Configuration

In RAPoSDA, we adopt the chained declustering [6] method for data placement on data disks as a primary backup configuration that tolerates disk failures. When writing data onto data disks, RAPoSDA tries to select a disk that is currently rotating or one that has been longer in the standby state than the break-even time. This means that RAPoSDA has to maintain data elsewhere when waiting for the write timing. We use a write cache memory to temporarily maintain data. It is important to maintain reliability when the data are in the volatile cache memory, so we provide the cache memory with a primary backup configuration that corresponds to the data disks. As with MAID, RAPoSDA can use some disks as cache disks to provide larger read cache spaces. Figure 1 shows the configuration of RAPoSDA.

Fig. 2. Data Flow in RAPoSDA During Write Requests

Data Disks Chained declustering is a simple but effective strategy for data placement that provides good reliability and accessibility when the backup data is logically located in the disk next to the primary. We assume that a data disk spins down and moves to a standby state if the time without access is longer than a predefined threshold time.

Cache Disks When the size of cache memories is not sufficient for comparing the total capacity of a disk array, cache disks are effectively used for enlarging the read cache space. However, cache disks are optional in RAPoSDA if there is sufficient cache memory and access localization. The cache contents are copies, so they do not need to be replicated. In contrast to data disks, we assume that cache disks do not spin down. Therefore, a large number of cache disks is not appropriate for reducing power consumption.

Cache Memory RAPoSDA has two layers of cache memories that correspond to the primary and backup data disks, while one cache memory is shared by more than one disk drive. Each layer is connected to an individual power supply (UPS) to ensure the tolerance of power shortages or UPS failures.

3.2 Handling Write Requests

Write requests are processed in RAPoSDA as shown in Figure 2. Data are initially written into both the primary and backup layer of the cache memory. The written data are gathered in a corresponding buffer location in the cache memory of each individual disk that is responsible for storing the data. Buffered data are written onto their corresponding data disks when the amount of buffered data on the cache memory exceeds a predefined threshold. The rotation status of the data disk is investigated at this point. The data disk will spin up if it is in a standby state. When the disk enters an active state, all data on the location that corresponds to the disk in the cache memory are written to disks and then deleted from the cache memory. When a buffer overflow occurs in the primary layer corresponding to P_i , all data in the backup layer buffer for B_{i-1} are also transferred onto the data disk. However, if the buffer overflow occurs in B_i , the data

Fig. 3. Data Flow in RAPoSDA During Read Requests

in P_{i+1} are also transferred. Therefore, the amount of data in the primary layer of the cache memory is different from that in the backup layer.

This collective writing process reduces the frequency of spin-ups and spin-downs. Data are also copied to the cache disks to ensure quick responses for future read requests when there is time-skewed access. The remaining data disks with data still in the buffer stay in a standby state beyond the break-even time.

3.3 Handling Read Requests

As shown in Figure 3, read requests initially check the existence of data in the cache memory, followed by cache disks. The cache memory has primary and backup layers, and both layers are searched for data. If the target data are in the cache memory or cache disks, the data are returned without accessing the data disks. If the data do not exist in the cache memory or cache disk, the data are read from a data disk.

Chained declustering is the method used for primary backup data placement on data disks, so RAPoSDA selects an appropriate disk from the two disks that correspond to the primary and backup for data. This selection is also important for reducing power consumption while still maintaining the break-even restriction. We propose the following set of rules for selecting the disk.

- If only one data disk is active, select the one that is active.
- If both data disks are active, select the disk with the largest memory buffer capacity.
- If both data disks are in standby, select the disk with the longer standby duration period.

Data read from data disks are also copied to cache disks to ensure a rapid response in future read requests when there is time-skewed access.

4 Disk Array Simulator

A number of simulators, such as DiskSim [4], have been proposed for evaluating the performance of storage systems. However, many cannot measure power consumption, including DiskSim. Dempsey [15] used an expansion of DiskSim with a function for

Fig. 4. Simulator Configuration

measuring power consumption, but it needed to measure the power of the actual disk drives in advance and there was a limitation on the number of hard disks that could be used in a simulation. Thus, we developed a new simulator to evaluate the performance and power consumption of RAPoSDA and MAID.

The developed simulator is shown in Figure 4. It simulates behavior of each disk drive in the target storage system with a given workload, including the response time and power consumption, and it can flexibly change its configuration and workload.

The simulator initially sets up the parameters for the workload and the configuration of the cache memories, cache disks, and data disks. Clients inside the simulator then generate requests at times assigned by the workload. The Storage Manager dispatches the requests to the cache memories, cache disks, and data disks, based on information from the Data Layout Manager. Logs of the operation status for each device are collected by the Log Collector and analyzed by the Analyzer after finishing the simulation.

5 Evaluation

We used the simulator described in the previous section to compare the performance and power consumption of RAPoSDA, MAID, and a Normal system that was composed of a simple disk array with no power saving mechanism. In this paper, we mainly focus on the effect of read/write ratio for their performance and power consumption. We also evaluated other aspects, including their scalability. The evaluation results on the scalability indicate the superiority of RAPoSDA to MAID. Because of the page limits, we will report the details of the scalability evaluation in the other chance.

The original MAID proposed by [3] had no cache memory. However, many practical storage systems have cache memories and the cache memories in RAPoSDA play a very important role in reducing power consumption. To ensure a fair comparison, we

Fig. 5. Parameters for the Hard Disk Drive Used in the Simulation

parameter	value
Capacity (TB)	2
Number of platters	5
RPM	7200
Disk cache size (MB)	32
Data transfer rate (MB/s)	134
Active power (Watt)	11.1
Idle power (Watt)	7.5
Standby power (Watt)	0.8
Spin-down energy (Joule)	35.0
Spin-up energy (Joule)	450.0
Spin-down time (sec)	0.7
Spin-up time (sec)	15.0

Fig. 6. Parameters of the Synthetic Workload

Workload parameter	value
Time	5 hours
read:write	7:3, 5:5, 3:7
Number of files	1,000,000 (32KB/file)
Amount of file size	64GB (Primary $\times$ Backup)
Number of requests	$\lambda \times 3600 \times \text{Time}$
Distribution of access	Zipf distribution
Request arrival distribution	Poisson process
Zipf factor	1.2
Mean arrival rate (λ)	25 (request/sec)

introduced cache memory into MAID and evaluated the effect of cache memory with the three storage systems.

To ensure reliability, we also modified MAID so it had two replicas on the data disks to match RAPoSDA. However, MAID had no replica in the cache memory and cache disks because the data in the cache was only a copy. To prevent data loss from the cache during failures, the cache memory and cache disk in MAID used the ‘write through’ protocol. To handle replicas in the data disks, MAID randomly selected disks for the replicas. Normal also used replicas and it provided a faster response than the other method. RAPoSDA determined the access disk based on the rotation state of individual disks.

The hard disk drive model used in our simulator was based on the Hitachi Deskstar 7K2000[12] produced by Hitachi Global Storage Technologies. Table 5 shows the parameters of the model. Furthermore, We prepared a synthetic workload for the evaluation. The workload parameters used in this experiment are listed in Table 6. In the workload, the access skew was based on Zipf distribution and the request arrival rate was based on a Poisson distribution. A workload was generated with three different types of read:write ratio, 7:3, 5:5, and 3:7.

In this evaluation, we set the simulation parameter as follows. The number of data disks was 64 with six cache disks and the total capacity of the cache memory was the number of data disks $\times$ 1/4 GB.

5.1 Power Consumption and Power Reduction Rate

Figure 7 shows the power consumption of Normal, MAID, and RAPoSDA, and the power reduction ratio of MAID and RAPoSDA when compared with Normal. This figure shows that Normal was the highest power consumer and that the power consumption of Normal increased with an increase in the write frequency. In contrast, MAID and RAPoSDA decreased power consumption with an increased write frequency. This

Fig. 7. Power Consumption of Each Storage System and the Power Reduction Ratios of RAPoSDA and MAID versus Normal

Fig. 8. Spin-up and Spin-down Counts with MAID and RAPoSDA

shows that MAID and RAPoSDA achieve significant power savings. The power reduction ratio of RAPoSDA was higher than MAID with all read/write ratios. This matched our assumption of time-skewed data access with high write requests.

The disk drive power consumption depended on the rotation state and the number of spin-ups or spin-downs. If spin-ups and spin-downs occurred frequently, the power consumption may exceed that with no disk spin-downs. Control of excessive spin-ups and spin-downs is very important for power saving in the storage system.

Figure 8 shows the number of disk spin-ups and spin-downs with MAID and RAPoSDA (Normal had no spin-downs). The graph shows that the trend for MAID was decreasing a number of spin-ups and spin-downs with a decreasing read frequency, while the trend for RAPoSDA was an increasing spin-up and spin-down count. From the perspective of the write cache size, MAID had a larger write cache than RAPoSDA because MAID used the cache disks for reads and writes, whereas RAPoSDA only used them for reads and this led to a smaller write cache memory compared with cache disks. This led to the possibility of a higher hit ratio of write requests for MAID and a smaller number of spin-ups and spin-downs with a frequent write workload.

However, the power reduction ratio shown in Figure 7 indicates that RAPoSDA gave a greater power reduction ratio compared with MAID. This shows that the group writing method of RAPoSDA provided a beneficial effect in maintaining the break-even time for individual data disks, even though the total counts of spin-ups and spin-downs with RAPoSDA exceeded those with MAID.

Fig. 9. Average Response Time with Normal, MAID, and RAPoSDA

5.2 Average Response Time

Figure 9 shows the average response time for each storage system. The graph shows that Normal had the fastest response time but Normal consumed the most amount of power, because the disks were always spinning. The average response time of RAPoSDA was faster than that of MAID. When the average response time of MAID changed from 2.302 sec to 3.308 sec, the corresponding values for RAPoSDA fell in the range 0.559 sec to 1.296 sec.

Based on the previous discussion of hit ratios for the write cache, it seemed possible that the average response time of MAID was shorter than that of RAPoSDA, because the cache hit was effective in reducing the response time. However, Figure 9 demonstrates that RAPoSDA was also superior to MAID in terms of the average response time. One of the main reasons for this was the cache protocol. MAID adopted a write through protocol for cache memories and cache disks to prevent data loss because of failures.

6 Related Studies

This section briefly reviews related systems other than MAID.

In DRPM [5], the power consumption of disk drives can be expressed as a function of the rotation speed (RPM). To achieve good performance and low power consumption, DRPM utilizes multispeed disk drives that can dynamically change the rotation speed of the disk depending on the system's workload. Similarly, Hibernator [16] exploited this concept to implement power saving and performance controlled by a RAID configuration of multispeed disk drives.

The reduction of power consumption by network servers, such as Web servers and Proxy servers, was investigated by Carrera [2]. According to that report, multispeed hard disk drives are necessary to reduce power consumption and maintain server response performance. However, dynamic changes in the frequency of disk drive rotations present many technical challenges. As a result, such hard disk drives are not currently popular in practical use.

PARAID [14] is a powerful power saving technique for targeting RAID-based storage [10]. The controller skews the access to a small number of disk drives. It then creates inaccessible disks and spins down these disks to reduce power consumption. GRAID [9] places an emphasis on ensuring reliability and power savings based on RAID10 disk arrays. Using information at the log disk that is added into the normal disk to store logs,

the system only needs to update the mirror disks periodically, so the system can spin down all the mirror disks to a low-power mode for most of the time and save energy. EERAIID [8] is a power saving method that is focused in the RAID controller layer. EERAIID reduces the power consumption using dynamic I/O scheduling and a cache management policy.

SIERRA[13] and RABBIT[1] are distributed storage systems that implement power proportionality through leveraging cluster-based data placement with data replication. These methods achieve proportional relation between power consumption and system performance by dividing all storing nodes into groups and controlling which groups are to be active to serve certain workloads.

7 Conclusions

Large-scale, high-performance, reliable, and low-power storage systems are required to construct better data centers for cloud computing. In this paper, we proposed RAPoSDA (Replica-assisted Power Saving Disk Array) for use with such storage systems. RAPoSDA carefully controls the timing and targeting of disk access by dynamically checking the rotation status of individual disk drives. We compared the performance and power saving effects of RAPoSDA with modified MAID in simulations with different ratios of read/write requests in the workloads. The original version of MAID had no cache memory and no replication mechanism, so we added them to MAID to make a fair comparison.

The experimental results showed that RAPoSDA and the modified MAID provide reduced power consumption compared with a simple disk array with no power saving mechanism. However, RAPoSDA was superior to the modified MAID. From the performance perspective, the simulation results showed that the average response time of RAPoSDA was shorter than that of the modified MAID. Thus, consideration of individual disk rotation is an effective method for reducing the power consumption of storage systems, while maintaining good performance.

In future work, we will consider distributed file systems such as Hadoop Distributed File System (HDFS) and Google File System (GFS) which accepts more than two replicas, as power saving storage systems. We also aim to develop an experimental system with actual disk drives and evaluate the system using actual workloads.

Acknowledgment

This work is partly supported by Grants-in-Aid for Scientific Research from Japan Science and Technology Agency (A) (#22240005).

References

1. Amur, H., Cipar, J., Gupta, V., Ganger, G.R., Kozuch, M.A., Schwan, K.: Robust and flexible power-proportional storage. In: Proceedings of the 1st ACM Symposium on Cloud Computing. pp. 217–228. SoCC '10, ACM, New York, NY, USA (2010)

2. Carrera, E.V., Pinheiro, E., Bianchini, R.: Conserving disk energy in network servers. In: ICS '03: Proceedings of the 17th Annual International Conference on Supercomputing. pp. 86–97. ACM, New York, NY, USA (2003)
3. Colarelli, D., Grunwald, D.: Massive arrays of idle disks for storage archives. In: Supercomputing '02: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing. pp. 1–11. IEEE Computer Society Press, Los Alamitos, CA, USA (2002)
4. Greg Ganger, et. al.: The DiskSim Simulation Environment (v4.0), <http://www.pdl.cmu.edu/DiskSim/>
5. Gurumurthi, S., Sivasubramaniam, A., Kandemir, M., Franke, H.: DRPM: Dynamic Speed Control for Power Management in Server Class Disks. International Symposium on Computer Architecture 0, 169 (2003)
6. Hsiao, H.I., DeWitt, D.J.: Chained Declustering: A New Availability Strategy for Multiprocessor Database Machines. In: Proceedings of the Sixth International Conference on Data Engineering. pp. 456–465. IEEE Computer Society, Washington, DC, USA (1990)
7. IDC: Extracting value from chaos (2011), <http://idcdocserv.com/1142>
8. Li, D., Wang, J.: EERAIID: energy efficient redundant and inexpensive disk array. In: EW 11: Proceedings of the 11th Workshop on ACM SIGOPS European Workshop. p. 29. ACM, New York, NY, USA (2004)
9. Mao, B., Feng, D., Wu, S., Zeng, L., Chen, J., Jiang, H.: GRAID: A Green RAID Storage Architecture with Improved Energy Efficiency and Reliability. In: MASCOTS. pp. 113–120 (2008)
10. Patterson, D.A., Gibson, G., Katz, R.H.: A case for redundant arrays of inexpensive disks (raid). SIGMOD Rec. 17(3), 109–116 (1988)
11. Program, U.E.P.A.E.S.: Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431 (2007), http://www.energystar.gov/ia/partners/prod_development/downloads/EPA.Datacenter.Report.Congress_Final11.pdf
12. Technologies, H.G.S.: Hard disk drive specification, hitachi deskstar 7k2000, [http://www.hitachigst.com/tech/techlib.nsf/techdocs/5F2DC3B35EA0311386257634000284AD/\\$file/USA7K2000_DS7K2000_OEMSpec.r1.4.pdf](http://www.hitachigst.com/tech/techlib.nsf/techdocs/5F2DC3B35EA0311386257634000284AD/$file/USA7K2000_DS7K2000_OEMSpec.r1.4.pdf)
13. Thereska, E., Donnelly, A., Narayanan, D.: Sierra: practical power-proportionality for data center storage. In: Proceedings of the Sixth Conference on Computer Systems. pp. 169–182. EuroSys '11, ACM, New York, NY, USA (2011)
14. Weddle, C., Oldham, M., Qian, J., Wang, A.I.A., Reiher, P., Kuenning, G.: PARAID: A gear-shifting power-aware RAID. Trans. Storage 3(3), 13 (2007)
15. Zedlewski, J., Sobti, S., Garg, N., Zheng, F., Krishnamurthy, A., Wang, R.: Modeling Hard-Disk Power Consumption. In: Proceedings of the 2nd USENIX Conference on File and Storage Technologies. pp. 217–230. USENIX Association, Berkeley, CA, USA (2003)
16. Zhu, Q., Chen, Z., Tan, L., Zhou, Y., Keeton, K., Wilkes, J.: Hibernator: helping disk arrays sleep through the winter. In: Proceedings of the Twentieth ACM Symposium on Operating Systems Principles. pp. 177–190. SOSP '05, ACM, New York, NY, USA (2005)