



Energy Consumption and Computation Models of Storage Systems

Wenlun Tong^(✉), Takumi Saito, and Makoto Takizawa

Hosei University, Tokyo, Japan
tongwenlun@gmail.com, takumi.saito.3j@stu.hosei.ac.jp,
makoto.takizawa@computer.org

Abstract. It is critical to reduce the electric energy consumption of information systems to realize green societies. Applications like database and web applications take usage of data in storages systems of servers. In this paper, we consider RAID storage systems which are composed of multiple drives like hard disk drives (HDDs) and solid state drives (SSDs). Types of RAID storage systems, RAID0, RAID10(1+0), and RAID5 are considered in this paper. The performance and reliability of RAID storage systems are so far studied by many researchers. The more number of storage drives are possibly in parallel accessed in the RAID storage systems, the more amount of electric energy is consumed while the higher reliability and availability are supported. The electric energy consumption of the RAID storage systems to read and write data is so far not discussed. In this paper, we measure the power consumption of RAID storage systems and time to read and write data in the storage systems in order to make a power consumption model of a storage system. We make clear how much energy each type of RAID storage system consumes to sequentially and randomly read and write data through experiment in this paper.

Keywords: Electric energy consumption · RAID storage systems · Hard Disk Drive (HDD) · Solid State Drive (SSD) · Power consumption model

1 Introduction

Information systems are getting more scalable like cloud computing systems [9] and IoT(Internet of Things) [5]. Here, huge amount of electric energy is consumed due to the scalability, e.g. millions of drives are interconnected in the IoT. It is critical to decrease electric energy consumed in information systems to reduce the carbon dioxide emission on the earth. The macro-level power consumption and computation models of a computer to perform application processes are proposed [8] [10] in order to design and implement energy-efficient information systems, models, and algorithms to select servers [8] [10] to perform application processes and make virtual machines migrate [7] in a cluster of servers are proposed in order

to reduce energy consumption. In this paper, we consider how much energy a storage system consumes to perform applications processes which read and write data in multiple drives of hard disk drives (HDDs) and solid state drives (SSDs).

The RAID (Redundant Arrays of Independent Disks) models of storage systems [6] are used to improve the performance and reliability of storage systems. Here, data units, i.e. blocks of each file f are distributed and replicated on multiple storage drives. Blocks in different drives are in parallel accessed to increase the performance. Multiple replicas of each block are stored in different drives to increase the reliability and availability. One replica of each block is properly operational on the other drive, even if one drive is faulty. The performance and reliability of RAID storage systems are so far studied by many researchers [6]. However, the energy consumption of the storage systems is not discussed. In this paper, we discuss how much electric energy each model of RAID storage systems consumes to read and write data. We measure the power consumption [W] of a RAID storage system and time to sequentially and randomly read and write data in the storages system through experiment.

In Sect. 2, we present a model of RAID storage system. In Sect. 3, we measure the electric energy consumption and time to read/write data in RAID storage systems.

2 System Model

A storage system S is composed of multiple storage drives SD_1, \dots, SD_d ($d \geq 1$) where files are stored and accessed. Each drive SD_i is a hard disk drive (HDD) or a solid state drive (SSD) ($i = 1, \dots, d$). We consider three types of RAID (Redundant Arrays of Independent Disks) [2] storage systems, RAID0, RAID10(0+1), and RAID5 in this paper.

A file f is a sequence of blocks b_1, \dots, b_m ($m \geq 1$). A block is a storage unit which is read and written in a read and write operation. That is, a block is a unit of read/write operation. A pair of blocks b_i and b_{i+1} ($1 \leq i < m$) are referred to as *consecutive* in the file f . If the file f is sequentially accessed, a block b_i is read before b_{i+1} . In the random access, a block b_i is directly access.

The RAID0 storage system supports striping Fig. 1. That is, the first block b_1 of a file f is stored in the first storage drive SD_1 and the second block b_2 is stored in the second drive SD_2 . Thus, each block b_i is stored in a drive $SD_{(i-1)\%d+1}$. Here, $x \% y$ stands for modulo of integer x by y . A pair of blocks b_i and b_j where $i \% d \neq j \% d$, i.e. which are in different drives, can be in parallel accessed. A pair of blocks b_i and b_j in different drives can be concurrently accessed. If one drive SD_i is faulty, blocks in the drive SD_i are lost since no block is replicated.

The RAID10 storage system supports reliability and availability by using the mirroring technologies in addition to the stripping i.e. parallel access Fig. 2. Suppose a storage system is composed of four drives ($d = 4$). Each block b_i is replicated in a pair of different drives. For example, two replicas of a block b_i are stored in different drives. The first block b_1 of a file f is stored in a pair of drives SD_1 and SD_2 . Next, the second block b_2 is stored in the drives SD_3 and SD_4 .




Fig. 1. RAID0 ($d = 4$).

Then, the third block b_3 is stored in the drives SD_1 and SD_2 . Thus, each block b_i is replicated in two different drives. Hence, the size of the storage to store a file f is double of the RAID0 since each block is replicated to two replicas. Even if one drive gets faulty, one replica of the block b_i is proper in the other drive.

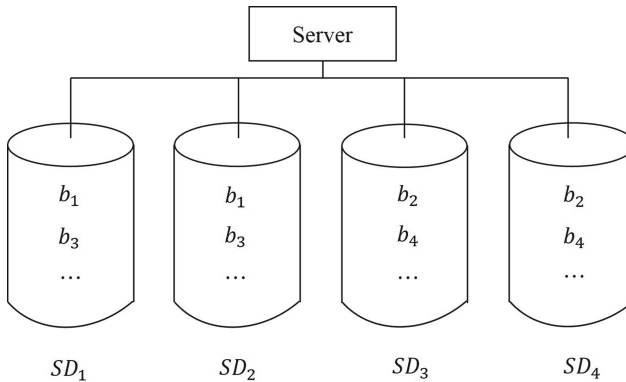


Fig. 2. RAID10 ($d = 4$).

In the RAID5 storage system, a subsequence of blocks b_1, \dots, b_m of a file f are divided to subsequences, each of which includes $(d - 1)$ consecutive blocks $b_i, b_{i+1}, \dots, b_{i+d-1}$. A parity block $pb_{i,i+d-1}$ is created for the subsequence of the blocks $b_i, b_{i+1}, \dots, b_{i+d-1}$ where $i \% (d - 1) = 0$. For example, a parity block $pb_{1,d-1}$ of a subsequence of blocks b_1, \dots, b_{d-1} is created by taking the exclusive or (xor) \oplus of the blocks, i.e. $pb_{1,d-1} = b_1 \oplus \dots \oplus b_{d-1}$. A subsequence $\langle b_1, \dots, b_{d-1}, pb_{1,d-1} \rangle$ of d blocks and the parity block $pb_{1,d-1}$ are stored in the drives $SD_1, \dots, SD_{d-1}, SD_d$, respectively. Then, the subsequence of blocks $b_d, \dots, b_{2d-2}, pb_{d,2d-1}, b_{2d-1}$

are stored in the drives $SD_1, SD_2, \dots, SD_{d-2}, SD_{d-1}, SD_d$, respectively. Let us consider a subsequence $\langle b_{kd+1}, b_{kd+2}, \dots, b_{(k+1)d} \rangle$ of d consecutive blocks of a file f ($k \geq 1$). A parity block $pb_{kd+1,(k+1)d}$ is created for the d blocks $b_{kd+1}, b_{kd+2}, \dots, b_{(k+1)d}$. The parity block $pb_{kd,(k+1)d-1}$ is the $((k+1)d - k\%d)$ the element of a subsequence $\langle b_{kd+1}, \dots, b_{(k+1)d-k\%d}, pb_{kd,(k+1)d-1}, b_{(k+1)d-k\%d+1}, \dots, b_{(k+1)d} \rangle$. Hence, the size of data stored in the RAID5 storage system is smaller than the RAID10 while larger than RAID0. Even if one block b_i is faulty, the faulty block b_i is recovered by taking the $xor \oplus$ of the other blocks and the parity block in the subsequence. For example, if the block b_5 is lost, the block b_5 is obtained as $b_5 = b_4 \oplus pb_{4,6} \oplus b_6$ (Fig. 3).




Fig. 3. RAID5 ($d = 4$).

Table 1 summarizes properties of each type of RAID storage systems, where a file of m (≥ 1) blocks b_1, \dots, b_m are stored in d ($d \geq 2$) drives SD_1, \dots, SD_d . In the RAID0 storage system, no replica of each block b_i is created, i.e. no redundancy. On the other hand, blocks in different drives can be concurrently accessed by applications. In the RAID10 storage system, two replicas of each block are stored in different drives. The RAID10 storage system supports more reliability than the RAID0. In the RAID 5 storage system, one parity block is created for $(d - 1)$ consecutive blocks. A subsequence of d blocks and the parity block is stored in different drives. Actually, $m \cdot d/(d-1)$ blocks are stored in d drives SD_1, \dots, SD_d . The size of data stored in the d drives of the RAID5 storage system are $d/(d-1)$ times larger than the RAID0 and $d/[2(d-1)]$ times smaller than the RAID10.

3 Experiment

3.1 System Configuration

We measure the execution time [sec] to read and write data and the energy consumption [J] of each model of RAID storage system is first, RAID0, RAID10,

Table 1. Properties of RAID

RAID	Storage size	Redundancy	Number of blocks
0	m	0	m/d
10	$2m$	1	$2m/d$
5	$m \cdot (d/d-1)$	$d/(d-1)$	$m/(d-1)$

m = number of blocks. d = number of drives.

and RAID5. We consider a storage system “Yottamaster Y-Focus Series 4-Bay “ [4] by which the RAID0, RAID10, and RAID5 types can be used. Here, a storage system is composed of four storage drives SD_1, \dots, SD_4 ($d = 4$). For each drive SD_i , an HDD (Seagate BarraCuda,2 [TB]) and an SSD (Crucial MX500, 500 [GB]) can be installed to do the experiment.

The power consumption [W] of the storage systems are measured by using the UW meter [3]. The electric power [W] is supplied to the RAID storage system S through the UW meter as shown in Fig. 4. The power consumption [W] of the RAID storage system S can be measured every one hundred [millisecond] in the UW meter. The electric power [W] of the RAID storage system S measured by the UWmeter is transferred to a note PC by using the bluetooth communication.




Fig. 4. Experiment.

The storage system S is connected to a Windows PC as shown in Fig. 4. First, we measure the electric power of the storage system S where $k (\leq d)$ storage drives are used. Initially, the storage systems consumes the minimum power $minE = 3.618$ [W] where no drive is accessed. Figures 5 and 6 show the power consumption of the storage system S for the number $k (\leq d)$ of storage drives, HDDs and SSDs, respectively. Here a file f of 10 [GB] is copied to each of k drives. The storage system S of HDDs consumes 15.063, 16.207, 16.824, and 17.294 [W] for $k = 1, 2, 3,$ and 4 , respectively. The storage system S of SSDs consumes 5.491, 5.91, 6.063, and 6.142 [W] for $k = 1, 2, 3,$ and 4 , respectively.

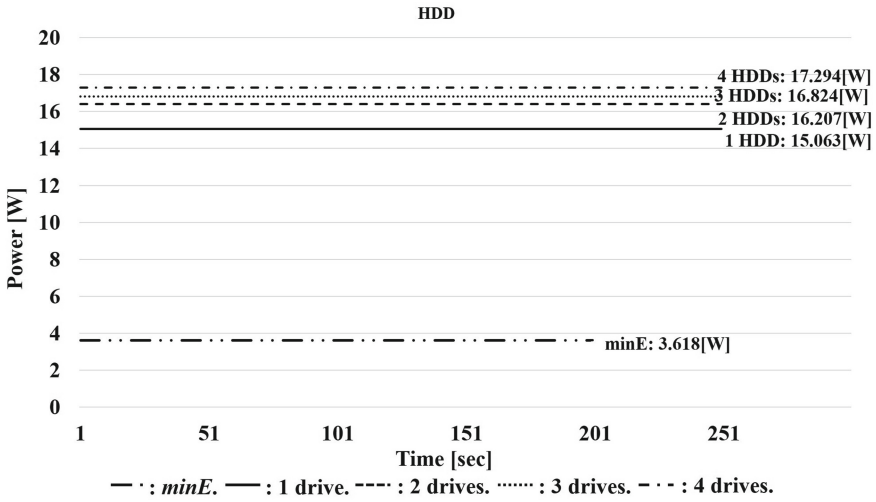


Fig. 5. Power consumption of HDDs.




Fig. 6. Power consumption of SSDs.

3.2 RAID for Sequential Access

First, the RAID type 0, 10, 5 or 5 is fixed in experiment. Then, a file f is sequentially written to the storage system by using a *copy* command from the PC. In turn, the file f is sequentially read by a *copy* command. Here, we measure the power consumption of storage system S with RAID0, RAID10, and RAID5 types.

Figures 7 and 8 show the power consumption of the RAID storage system S composed of HDDs and SSDs, respectively. Here, a file of 10 [GB] is written to and read from the RAID storage system. The power consumption of the storage system composed of HDDs is 16.99, 18.20, and 16.83 [W] for RAID0, RAID10, and RAID5 types, respectively. In the storage system composed of SSDs, the power consumption is 6.78, 6.93, and 6.84 [W] for RAID0, RAID10, and RAID5 types, respectively. The power consumption of the storage system S composed of the SSDs is about 60% smaller than the HDDs. The RAID10 storage system S consumes more power than the RAID0 and RAID5. In the HDDs, the RAID1 consumes more energy than the RAID0 and RAID5 and the RAID5 consumes the smaller power. On the other hand, the RAID0 consumes the smaller power while the RAID1 consumes the greatest power.

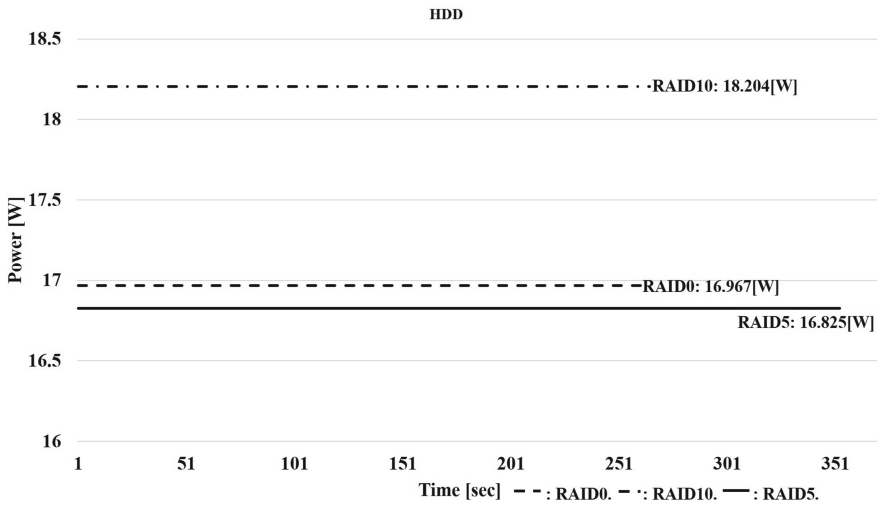


Fig. 7. Power consumption (HDDs) in sequential access.

Figures 9 and 10 show the execution time [sec] for data size [GB] of a file f to write data in the RAID storage system S . The execution time of the RAID storage system S composed of HDDs to write a file f of 10 [GB] is 261, 265, and 353 [sec] for RAID0, RAID10, and RAID5 types, respectively. The execution time of RAID0 is the fastest and RAID5 is the slowest in the HDDs. For the RAID storage system S composed of the SSDs, the execution time to write a file f of 10 [GB] is 259, 257, and 256 [sec] for the RAID0, RAID10, and RAID5, respectively. The execution time of the RAID5 storage system S is the fastest and the RAID0 is the slowest.




Fig. 8. Power consumption (SSDs) in sequential access.

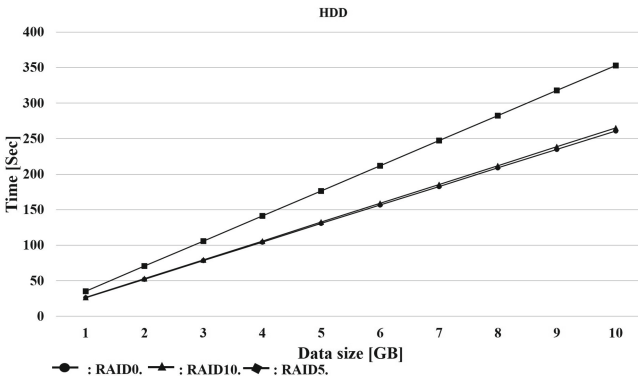


Fig. 9. Execution time (HDDs) in sequential access.

3.3 RAID for Random Access

Next, data in a RAID storage system S is randomly accessed. By using HDtune [1], data of size 512 [B] to 1 [MB] is randomly read in the file f of 10 [GB]. The average size of the data is 4 [KB]. Figures 11 and 12 show the power consumption [W] of the storage system S composed of HDDs and SSDs, respectively. The power consumption of the storage system composed of HDDs is 19.742, 17.05, and 19.318 [W] for RAID0, RAID10, and RAID5 types, respectively. The power consumption of the storage system composed of SSDs is 6.919, 5.933, and 7.047 [W] for RAID0, RAID10, RAID5 types, respectively. The power consumption of the storage system S composed of SSDs is about 65% smaller than the HDDs. The RAID0 storage system S of HDDs consumes more power than the RAID10




Fig. 10. Execution time (SSDs) in sequential access.




Fig. 11. Power consumption (HDDs) in random access.

and RAID5, and the RAID 10 consumes the smallest power consumption in RAID0, RAID10, and RAID5. Compare with the sequential access of RAID storage system composed of HDDs, the power consumption of the sequential access of RAID storage system is smaller than random access of RAID storage system of RAID0, RAID10, and RAID5 types. In the storage system composed of SSDs, the power consumption of random access is bigger than the sequential access.

Figures 13 and 14 show the execution time [sec] of the storage system S of HDDs and SSDs, respectively. The execution time in the random access of the RAID storage system S of HDDs to read a file f of 10 [GB] is 92, 87, and 92 [sec] for RAID0, RAID10, and RAID5 types, respectively. For the RAID storage system S composed of SSDs, the execution time in random access to read a




Fig. 12. Power consumption (SSDs) in random access.




Fig. 13. Execution time (HDDs) in random access.

file f of 10 [GB] is 43, 43, and 42 [sec] for RAID0, RAID10, and RAID5 types, respectively. The execution time of the RAID5 is the fastest and the RAID0, and RAID5 is slowest, the execution time in random access is about 76% smaller than execution time in sequentially access.




Fig. 14. Execution time (SSDs) in random access.

4 Concluding Remarks

The RAID storage system is widely used to realize reliable and high performance storage systems. The more number of storage drives data are stored and replicated, the more efficient and reliable storage systems used. On the other hand, the more electric energy is consumed. In this paper, we measured the energy consumption [J] and execution time [sec] of RAID storage systems, RAID0, RAID10, and RAID5 with HDDs and SSDs to sequentially and randomly read and write data.

By taking advantage the measured data, we are now making the power consumption model and the execution model of a storage system.

References

1. Hdtune. <http://hdtune.com>
2. Raid wiki. <https://ja.wikipedia.org/wiki/RAID>
3. Uwmeter. <http://www.metaprotocol.com/UWmeter/UWmeter/TOP.html>
4. Yottamaster raid y-focus series 4-bay. <https://www.amazon.co.jp/Yottamster>
5. Arridha, R., Sukaridhoto, S., Pramadihanto, D., Funabiki, N.: Classification extension based on iot-big data analytic for smart environment monitoring and analytic in real-time system. *International Journal of Space-Based and Situated Computing (IJSSC)* **7**(2), 82–93. <https://doi.org/10.1504/IJSSC.2017.10008038>
6. Chen, Peter, M., Lee: Raid: high-performance, reliable secondary storage. *ACM Computing Surveys* **26**(2), 145–185 (1994)
7. Duolikun, D., Enokido, T., Takizawa, M.: An energy-aware algorithm to migrate virtual machines in a server cluster. *International Journal of Space-Based and Situated Computing (IJSSC)* **7**(1), 32–42. <https://doi.org/10.1504/IJSSC.2017.10004986>

8. Enokido, T., Ailixier, A., Takizawa, M.: An extended simple power consumption model for selecting a server to perform computation type processes in digital ecosystems. *IEEE Transactions on Industrial Informatics* **10**(2), 1627–1636 (2014)
9. F, A.H., Alenezi, A., Alharthi, A.: Integration of cloud computing with internet of things. *IEEE International Conference on Internet of Things* (2017)
10. Kataoka, H., Sawada, A., Duolikun, D., Enokido, T., Takizawa, M.: Multi-level power consumption and computation models and energy-efficient server selection algorithms in a scalable cluster. In: *Proc. of the 19th International Conf. on Network-Based Information Systems (NBIS-2016)*, pp. 210–217 (2016)