A Comparative Study of HDD and SSD RAIDs' Impact on Server Energy Consumption

Erica Tomes and Nihat Altiparmak
Department of Computer Engineering & Computer Science, University of Louisville

I. Introduction

In the US alone, data centers consumed around $20 billion (200 TWh) yearly electricity in 2016, and this amount doubles itself every five years. Data storage alone is estimated to be responsible for about 25% to 35% of data-center power consumption [1]. Servers in data centers generally include multiple HDDs or SSDs, commonly arranged in a RAID level for better performance, reliability, and availability. In this study, we evaluate HDD and SSD based Linux (*md*) software RAIDs' impact on the energy consumption of popular servers. We used the Filebench workload generator [2] to emulate three common server workloads: *web*, *file*, and *mail*, and measured the energy consumption of the system using the HOBO power meter [3]. We observed some similarities and some differences in energy consumption characteristics of HDD and SSD RAIDs, and provided our insights for better energy-efficiency. We hope that our observations will shed light on new energy-efficient RAID designs tailored for HDD and SSD RAIDs' specific energy consumption characteristics.

II. Definitions

Definition: 1. *"Power" is the* rate *of energy consumption (in watts) and "Energy" is the amount of energy consumed (in joules), where* ***Energy (joules) = Power (watts) x Time (sec)***.

Definition: 2. *A storage device is in "active" state while performing an I/O operation, in "idle" state while not servicing a request but staying ready to begin the next request, and in "power-off" (or "stand-by") state when its disks are spun down (HDDs) or its NAND flash chips are powered-off (SSDs).*

III. Methodology

Our experiments were conducted on a Dell PowerEdge R730xd server equipped with two Intel Xeon 14-core 2.4 GHz processors. Software RAIDs were formed using 1TB Toshiba SATA HDDs with 7.2K RPM and 800GB Intel (S3510) SATA SSDs. The machine ran an Ubuntu 16.04.1 LTS operating system with kernel version 4.4.0-79. We used *mdadm* [4] to manage Linux *md* software RAIDs. We connected our server to the HOBO power meter [3], generated I/O workloads using FileBench [2], and measured the energy drawn by the server using an in-house developed tool that can communicate with FileBench and HOBO. For the I/O workloads, we used FileBench's pre-defined *web*, *file*, and *mail* personalities, where *web* emulates a web server workload with 10:1 R/W ratio, *file* emulates a home directory workload with 1:2 R/W ratio, and *mail* emulates an e-mail server workload with 1:1 R/W ratio. All workloads were generated by 100 threads simultaneously performing I/O. Based on the dataset size to memory ratio suggested by the FileBench, we limited our server's memory to 4GB so that enough I/O is generated. All experiments were run five times and the results were averaged.

IV. Results and Observations

In this section, we present our experimental results and share our observations on the energy consumption characteristics of HDD and SSD RAIDs.

Observation 1. *Traditional power-off based energy conservation techniques applied in HDD RAIDs are not suitable for SSD RAIDs.*

Energy conservation techniques proposed for HDD RAIDs mainly focus on switching a subset of disks to the *power-off* state [5]. In order to understand storage devices' impact on server's energy consumption, we measured the idle power rate of the system with a various number of HDDs and SSDs attached to it. Figure 1(a) presents our findings for No Disks, 1 to 4 HDDs, and 1 to 4 SSDs on the *x*-axis, and the power consumption of the system on the *y*-axis. As it is clear from the figure, power consumption of the server with SSDs is very close to No Disk power consumption indicated by the red horizontal line, where each SSD adds around 0.6 watts to the system's power requirement. On the other hand, each HDD causes around 6 more watts to be drawn, requiring an order of magnitude more idle power than an SSD. ***Due to their high operation cost (data reorganization/synchronization/replication) and minimal energy conservation potential, power-off based energy conservation techniques are not suitable for SSD RAIDs.***

Observation 2. *HDD RAIDs are more power-hungry than SSD RAIDs when the storage system is idle; however, SSD RAIDs are more power-hungry than HDD RAIDs when the storage system is active.*

In Figure 1(b), we compare the idle and active power consumption of a system with a 4-disk HDD RAID10 vs. a 4-disk SSD RAID10 for various server workloads. Without depending on the workload, the results indicate that SSD RAID causes around 15 watts more power consumption than the HDD RAID on average when the storage system is active. This result contrasts with the common assumption of HDD RAIDs being more power-hungry while active due to their

2168-9253/17 $31.00 © 2017 IEEE
DOI 10.1109/CLUSTER.2017.103

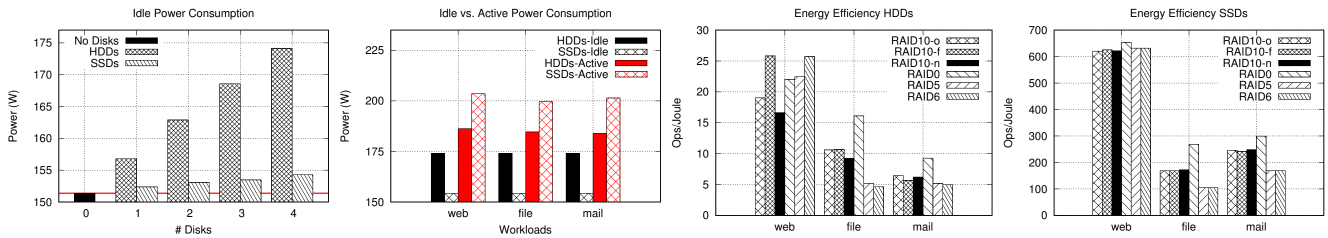

(a) Idle Power Consumption (b) Idle vs. Active Power Cons. (c) Energy-efficiency - HDDs (d) Energy-efficiency - SSDs

Fig. 1: Energy Consumption Characteristics of HDD and SSD RAIDs

mechanical seek and rotation operations. On the other hand, the system with the SSD RAID consumes around 20 watts less power than the HDD RAID when the storage system is idle, which results in 4x (~35 watts) more *Delta* ($\Delta = Active - Idle$) energy consumption for SSD RAIDs compared with the HDD RAIDs. Although SSDs are clearly faster for read operations, sometimes sequential write performance of HDDs can be close to or even better than SSDs, especially when the write amplification cost of SSDs are high due to space contention. ***By considering the Δ energy consumption factor and the I/O performance, this property can be exploited in hybrid RAID designs for improved energy efficiency through energy-aware I/O scheduling between HDDs and SSDs.***

Since energy consumption depends on both power rate and the time spent performing the I/O operation, power can be a misleading metric by itself to measure the energy efficiency. In order to compare the energy efficiency of different RAID levels, we use *operations per joule* (ops/joule) as in [6], which indicates the amount of work the storage system performs in one joule of energy spent by the server.

Observation 3. *For all common server workloads of web, file, and mail, SSD RAIDs are one to two orders of magnitude more energy efficient than HDD RAIDs.*

Figures 1(c) and 1(d) compare the energy-efficiency of various RAID levels for HDD and SSD RAIDs, respectively. ***For all common server workloads and popular RAID levels, SSD RAIDs are significantly more energy-efficient than HDD RAIDs.*** The performance difference is mainly due to superior random I/O performance of SSD RAIDs and significant disk head movement caused by multiple threads simultaneously accessing different regions of HDD RAIDs.

Observation 4. *For read-intensive (web) workloads, energy efficiency of HDD RAIDs are more sensitive to RAID level, whereas SSD RAIDs achieve a fairly stable energy-efficiency performance for different RAID levels.*

The first three bars of Figures 1(c) and 1(d) show the performance of RAID10 *offset*, *far*, and *near* designs, which indicate that the replicas are placed in close, far, or similar offsets in different devices, respectively. By placing replicas far from each other, RAID10-f is more energy efficiency for HDD RAIDs in read-intensive workloads since far placement reduces seek distance while reading sequential blocks from the same HDD. On the other hand, data placement has a minor impact on SSD RAIDs' energy-efficiency for read-intensive workloads due to SSDs' random access property. ***For read-intensive workloads, data placement should be carefully designed for HDD RAIDs' energy-efficiency considering the internal device characteristics and seek distance.***

Observation 5. *For write-intensive (file) and mixed (mail) workloads, energy-efficiency of RAID0 is consistently superior to other RAID levels, and mirroring with RAID10 provides better energy efficiency than parity-based data protection techniques of RAID5/6 for both HDD and SSD RAIDs.*

As RAID0 does not provide any data protection, its energy-efficiency is clearly better than other RAID levels for both write-intensive and mixed workloads, without depending on the device type. However, when data protection is necessary, then mirroring with RAID10 provides better energy-efficiency than parity-based techniques for both SSD and HDD RAIDs since additional CPU power consumption due to parity calculation of every write operation is eliminated in mirroring. ***Therefore, for write-intensive and mixed workloads, mirroring should be preferred over parity based techniques for the energy-efficiency of both HDD and SSD RAIDs.***

V. Future Work

Hybrid arrays have received considerable attention recently among all storage arrays due to their balance of capacity, price, and performance [7]. Using the observations made in this paper, our future work includes developing energy-aware hybrid RAID designs that can provide ops/joule performance close to SSD RAIDs with the price close to HDD RAIDs.

Acknowledgment

This research was supported in part by the U.S. National Science Foundation (NSF) under grant CNS-1657296.

References

[1] J. Kim *et al.*, "Energy proportionality for disk storage using replication," in *EDBT/ICDT '11*. ACM, pp. 81–92.
[2] V. Tarasov *et al.*, "Filebench: A flexible framework for file system benchmarking," *;login: The USENIX Magazine*, vol. 41, Mar 2016.
[3] *HOBO UX120 Plug Load Logger*, Datasheet, https://goo.gl/Z9yHF8.
[4] *mdadm - manage MD devices aka Linux Software RAID*, Linux man pages, https://linux.die.net/man/8/mdadm.
[5] E. Pinheiro *et al.*, "Energy conservation techniques for disk array-based servers," in *ACM Supercomputing 25th Anniv. Volume*, 2014, pp. 369–379.
[6] P. Sehgal *et al.*, "Evaluating performance and energy in file system server workloads," in *FAST'10*, 2010, pp. 253–266.
[7] E. Burgener *et al.*, *Worldwide All-Flash Array and Hybrid Flash Array 20142018 Forecast and 1H14 Vendor Shares*, IDC, Jan 2015.