Application Sensitive Energy Management Framework for Storage Systems

(Extended Abstract)

Norifumi Nishikawa
Hitachi America, Ltd.
Santa Clara, California 95054
Email: norifumi.nishikawa@hal.hitachi.com

Miyuki Nakano
Advanced Institute of Industrial Technology
Tokyo, Japan
Email: miyuki@aiit.ac.jp

Masaru Kitsuregawa
University of Tokyo
Tokyo, Japan
Email: kitsure@tkl.iis.u-tokyo.ac.jp

***Abstract*—Rapidly escalating energy and cooling costs of storage systems have become a concern for data centers. In response, a multitude of energy saving approaches that take into account storage-device-level input/output (I/O) behaviors has been proposed. The trouble is that critical applications are in constant operation at data centers, and the conventional approaches do not produce sufficient energy savings. It may be possible to dramatically reduce storage energy consumption without degrading application performance levels by utilizing application level I/O behaviors. However, such behaviors differ from one application to another, and it would be too expensive to tailor methods to individual applications. We propose a universal storage energy management framework for runtime storage energy savings that can be applied to any type of application. The results of evaluations show that the use of this framework results in substantive energy savings compared with the traditional approaches.**

I. Introduction

The amount of digital data produced by human beings is increasing every day. According to a recent International Data Corporation (IDC) report [1], the amount of information created and stored digitally will exceed 40,000 exabytes by 2020. This *Big Data* must be managed and used in conjunction with data-intensive applications such as sensor data archives, search engines, and Web services. Furthermore, since all this data must be properly stored, the capacity of storage systems will increase as well. It has been predicted [2] that data storage capacities in 2017 will need to be 6.5 times as large as those available in 2011.

II. Motivation and Background Information

The energy consumption of information technology (IT) equipment at data centers is growing on a daily basis [3], and a recent report [4] states that storage system energy consumption rates, as a portion of all IT equipment, will increase even further in the near future because the amount of digital data requiring storage will continue to inflate. For example, the paper [5] reports that the power required to handle data storage in large online transaction processing systems (OLTPs) and other data-intensive applications accounts for more than 70% of the total power consumption of IT equipment. Therefore, it is clear that storage system energy consumption has to be reduced at large data centers.

There are numerous energy saving approaches that use energy-efficient storage devices, such as large-capacity hard disks (HDDs), 2.5 inch HDDs, or solid state disks (SSDs), or that use space-efficient improvement methods such as data compression or data de-duplication. In contrast with previous energy saving approaches, our paper focuses on enabling dynamic storage energy savings while applications are running.

The storage energy saving methods that have been proposed to date include the one reported in [6], which has a massive arrays of inactive disks (MAID) function that stops HDD rotation when no I/Os are issued to the HDDs [7]. In order to take advantage of disk power control functions, some methods extend idle period durations by controlling the I/O intervals in the storage devices [8], while others extend idle periods by replacing data amongst the disks [9]–[11].

III. Problem Definition

Data centers have various applications in continuous operation, and the I/O behaviors of those applications, such as file servers, OLTPs, decision support systems (DSSs), and Web commerce, differ considerably. For example, the Transaction Processing Performance Council-C (TPC-C) [12] benchmark issues random I/Os to master data tables, while TPC-H [13], a typical benchmark for DSS applications, primarily issues sequential read orders to large transaction tables.

Traditional energy saving methods are typically implemented inside storage devices, which means their energy saving abilities can only be determined from their storage-level I/O behaviors. There is also the possibility that energy saving methods based on such behaviors could degrade the performance of applications since storage devices have no ability to detect or determine application runtime behaviors.

IV. Summary of Technical Ideas

Critical data center applications cannot be stopped once they have entered service. This means a runtime storage energy saving method is important. Furthermore, it is difficult to reduce the energy consumption of storage devices that are accessed frequently from applications such as OLTP or those related to Web commerce.

With this in mind, we propose a universal storage energy management framework that utilizes application-level I/O behaviors for reducing the energy consumption of storage devices while ensuring that critical applications can continue normal operations. Specifically, by utilizing application-level I/O behaviors, we can introduce a middleware layer monitoring function that works in cooperation with a monitoring function

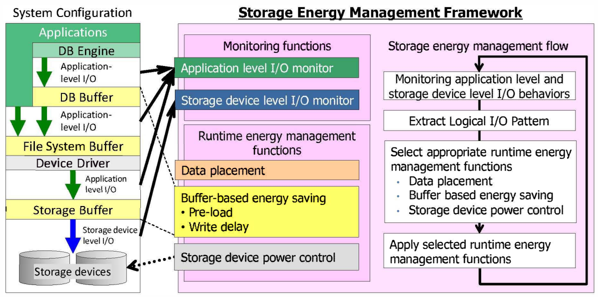

Fig. 1. Universal Storage Energy Management Framework.

at the storage-device level. Here, we introduce a logical I/O pattern that can easily handle a variety of application types and can analyze the applications' I/O behaviors in a uniform way. By monitoring both application- and storage-device-level I/O behaviors, we can utilize a buffer in a different layer, such as an application or a storage device, in a uniform way.

Figure 1 shows an overview of our framework. The left part of the figure shows the system configuration of a database server, or file server, that we suppose. The right part shows the functions and storage-device energy management flow. As can be seen in the figure, our framework has monitoring functions for both application and storage device levels, along with an energy management function that contains a data placement function, a buffer-based energy saving function, and a storage device power control function.

Our framework collects application- and storage-device-level I/O behaviors during operation by utilizing monitoring functions, which it then combines and analyzes to extract logical I/O patterns. Our framework selects an appropriate energy-reduction function to apply to the storage devices, changes the data placement in the storage devices, and executes the selected strategy to control the storages' energy by using a target systems' buffer. Application I/O behaviors may change when the number of users or frequency of data accesses change. As can be seen from the flow in Fig. 1, our framework captures changes in the I/O behaviors of applications as changes in logical I/O patterns and reselects appropriate functions to reduce storage energy. As a result, it always efficiently reduces energy consumed by storage devices.

V. Evaluation Results

We quantitatively evaluated our framework using typical data center applications. We run actual file server, OLTP, DSS, and Web commerce applications on multiple HDD and large storage system environments. We also compared our method with two established storage-device-level energy saving methods, popular data concentration (PDC) [9] and dynamic data reorganization (DDR) [11].

For multiple HDD environment, our framework was able to reduce HDD energy consumption more than could be accomplished with traditional approaches. Our framework decreased the HDD energy consumption for TPC-H by 53.0 percent, for TPC-C by 21.3 percent, and for TPC-W by 19.6 percent respectively with little performance degradation. For large storage system environment, our framework reduced the energy consumed by a storage system more than the conventional approaches. Our framework decreased the HDD energy consumption for the file server by 25.8 percent, for TPC-C by 14.1 percent, for TPC-H by 48.4 percent, and for TPC-W by 66.5 percent respectively with little performance degradation.

These two results show that our framework enables significant reductions in storage device energy consumption at the expense of only a minor application performance degradation.

VI. Conclusion

In this paper, we proposed a universal framework for reducing storage device energy consumption while critical applications are running. We introduced a logical I/O pattern in order to easily handle a variety of application types, and analyzed the I/O behaviors of each application in a uniform way. Furthermore, we measured the I/O behaviors and energy consumption of multiple HDDs and an enterprise storage system with applications such as a file server, OLTP, DSS, and Web commerce. We then compared the method based on our framework with those of conventional PDC and DDR storage energy saving methods. For all applications, we found that the energy savings of our method were equal to or greater than those of PDC and DDR. Additionally, application performance levels when using our method did not degrade significantly in comparison with devices without energy saving functions. The results of our evaluations show that our universal storage energy saving framework can make significant contributions to data center energy conservation efforts.

References

[1] J. Gantz and D. Reinsel. (2012) The digital universe in 2020: Big data, bigger digital shadows, and biggest growth in the far east. [Online]. Available: http://www.emc.com/collateral/analyst-reports/idc-the-digital-universe-in-2020.pdf

[2] J. Chang, R. W. Cox, S. Deshpande, A. Kros, M. Suzuki, S. Low, A. Munglani, and A. Kim, "Forecast: External controller-based disk storage, worldwide, all countries, 2011-2017, 1q13 update," Tech. Rep., 2013.

[3] A. G. Yoder. (2010) Green storage technologies, capex and opex.

[4] S. W. Worth. (2012) Green storage - the big picture. [Online]. Available: http://www.snia.org/sites/default/education/tutorials/2012/fall/green/SW_Green Storage_Big_Picture_9-12.pdf

[5] M. Poess and R. O. Nambiar, "Energy Cost, The Key Challenge of Today's Data Centers: A Power Consumption Analysis of TPC-C Results," *Proceedings of the VLDB Endowment*, vol. 1, no. 2, pp. 1229–1240, 2008.

[6] "Hitachi Adaptable Modular Storage 2500," 2011.

[7] D. Colarelli, D. Grunwald, and M. Neufeld, "The Case for Massive Arrays of Idle Disks (MAID)," in *In FAST 2002*, 2002, pp. 1–6.

[8] A. E. Papathanasiou and M. L. Scott, "Energy Efficient Prefetching and Caching," in *Proc. of the USENIX Annual Technical Conference*, 2004.

[9] E. Pinheiro and R. Bianchini, "Energy conservation techniques for disk array-based servers," in *Proc. of ICS '04*. New York, New York, USA: ACM Press, 2004, pp. 68–78.

[10] A. Donnelly and A. Rowstron, "Write Off-Loading : Practical Power Management for Enterprise Storage," in *6th USENIX Conf. on FAST*, 2008, pp. 253–267.

[11] E. Otoo, D. Rotem, and S.-c. Tsao, "Dynamic Data Reorganization for Energy Savings," in *Proc. of the 22nd Intl. Conf. on Scientific and statistical database management*, 2010, pp. 322–341.

[12] "TPC BENCHMARK C Standard Specification Revision 5.11," *Transaction Processing Performance Council*, 2010.

[13] "TPC BENCHMARK H (Decision Support) Standard Specification Revision 2.14.2," *Transaction Processing Performance Council*, 2011.