Energy Consumption and Performance Analysis Between SSD and HDD *

Pablo J. Pavan [1], Vinícius R. Machado [2], Jean L. Bez[2],
Edson L. Padoin[1,2], Francieli Z. Boito[2], Philippe O. A. Navaux[2], Jean-François Méhaut[3]

[1] Universidade Reg. do Noroeste do Estado do Rio G. do Sul (UNIJUI) — Ijui, RS, Brazil
[2] Universidade Federal do Rio Grande do Sul (UFRGS) — Porto Alegre, RS, Brazil
[3] Université Grenoble Alpes – Grenoble, France

{pablo.pavan, ricardo.lorenzoni, padoin}@unijui.edu.br,
{jean.bez, francieli.zanon, navaux}@inf.ufrgs.br, jean-francois.mehaut@imag.fr

Abstract

This paper presents an energy efficiency and I/O performance analysis of low-power architectures when compared to conventional architectures. Despite the fact, the power demand of the storage device amounts for a small fraction of the power demanded by whole system, significant increases are observed when accessing the storage device. We investigate the access pattern impact on power demand, looking at the whole system and at the storage device by itself. All tested configurations are compared regarding energy efficiency. We show the viability of replacing conventional servers by low-power alternatives depends on the characteristics of the expected workload.

1. Introduction

The increase in processing power of High-Performance Computing (HPC) architectures had, as a consequence, the growth in power demand. Taking this matter into consideration, a DARPA report suggests a limit of 20 MW of power demand for future HPC systems, which are expected to achieve exaflops of performance [2].

This premise has led HPC researchers to seek alternatives that respect the given power limit. A strategy could be the use of low-power architectures, replacing traditional processors by Advanced RISC Machine (ARM) processors. Despite presenting lower performance, these ARM processors provide better energy efficiency for some scientific applications [5].

In any typical computational system, the processor is not the solely responsible for the power demand of the system. Because of the increasing difference between processing and data access speeds, it is common for applications to spend a significant share of their runtime reading or writing data. Therefore, it is imperative that the energy consumption of the system, during these input/output (I/O) operations, be investigated and improved.

In this paper, we conduct a comparative study of I/O performance and power demand. We analyze the power demand of the storage devices and the system as a whole to better correlate it with the measured bandwidth. We compare a machine with a traditional processor to a Multi-Processor System-on-Chip (MPSoC), with a low-power ARM processor. Our analysis considers Hard Disk Drives (HDDs) and Solid State Drives (SSDs) as storage devices. Our goal is to determine the I/O workloads that benefit most in using such low-power architectures as data servers.

The remainder of this paper is organized as follows. Section 2 discusses related work. The experimental environments and methodology are detailed in Section 3. Results and discussed are presented in Section 4. Finally, Section 5 concludes this paper and discusses possible future work.

2. Related Work

Sawada *et al.* investigate the power consumption of a server to perform storage operations and computations. They defined a simple power consumption model of a server, abstracting parameters like the number of processes,

* This work was supported by CNPq, CAPES, FAPERGS and FINEP. This research has received funding from the European Community's Seventh Framework Programme (FP7-PEOPLE) under grant agreement number 295217, funding from the EU H2020 Programme and from MCTI/RNP-Brazil under the HPC4E Project, grant agreement number 689772 and STIC-AmSud/CAPES scientific-technological cooperation program under EnergySFE research project grant 99999.007556/2015-02.

which dominate the power consumption [6]. Hui *et al.* analyzes the performance and the energy consumption of some SSDs using different combinations of file systems and I/O blocks [1].

Nijim *et al.* combine flash-based storage devices (SSDs) with hard drives to provide storage with lower energy consumption. This is achieved by using the SSD as a cache for the hard drives [4]. This hybrid storage strategy is exploited by others to provide high performance for I/O servers [7]. In these cases, the SSD is used only as a cache because of its high cost per byte, which makes the total replacement of hard drives unfeasible.

In a previous work [3], an initial evaluation was done comparing low-power architectures with common architectures. This evaluation was done under the HPC point of view, aiming to analyze the workloads of scientific applications. The work suggested that MPSoCs were a good low-power alternative for read workloads. The research presented in this work is more complete — exploring more configurations and different workloads — and reaches different conclusions since workloads are more intensive.

3. Experimental Methodology

Two environments were used for the experiments discussed in this work. The first, named **PC**, is a traditional desktop, with an Intel Core-I7 4790 processor of 3.6 GHz base clock frequency. The equipment has 16 GB RAM operating at 1600 MHz.

The second environment is a MPSoC CubieTruck, with a SoC A20 manufactured by AllWinnerTech, and a dual GPU MALI400 MP2, called **MPSoC**. The processor is a Dual Core ARM Cortex-A7 at 920 MHz frequency. The equipment has 2 GB of RAM operating at 480 MHz frequency.

Four storage devices were used for the experiments, two SSDs and two HDDs selected to cover different characteristics. They are presented in Table 1.

	Type	Manufacturer/Model	Capacity (GB)	Interface	RPM
HDD1	HDD	Seagate ST1000LM035-1RK172	1000	SATA III	5400
HDD2	HDD	Seagate ST96023AS	60	SATA II	7200
SSD1	SSD	Samsung PM871	240	SATA III	-
SSD2	SSD	Kingston SV300S37A120G	120	SATA III	-

Table 1. Storage devices used for the experiments

The benchmark chosen for the tests was FIO, [1] selected for being widely used and for allowing the description of the desired access patterns. The **write** experiments were conducted in each configuration **with** and **without** the usage of the buffer cache. On the other hand, we do not include read experiments with the cache because in this situation the behavior depends on previous accesses. If the same data was accessed recently, it may still be in the cache, but a variety of situations may occur.

Regarding the type of operation and the spatiality, two access patterns were generated: sequential write, random write, sequencial read, random read. Additionally, two request sizes were evaluated: 32 KB and 4 MB. Each test was executed for 20 GB of data with an execution time limit of 60 seconds - the test stopped when the first of these two conditions was met. We use the bandwidth reported by FIO as the performance metric for these experiments.

Considering the eight configurations of equipment and storage device, two cache options for half the experiments (with or without it), four operations, and two request sizes, a total of 96 experiments are analyzed in this paper. Each one of them was repeated 10 times, and different experiments in the same equipment with the same storage device were executed in random order, to avoid unexpected effects. A minimum 20-seconds delay is guaranteed between tests, so the power demand of the environment stabilizes. Moreover, the "`sync`" command is used between tests that use the buffer cache to make sure they are independent.

To measure the power demand, we employed an Agilent oscilloscope model DSO6014A. This oscilloscope was connected via USB to a computer, where the BenchVue software logs captured data. A power tip model 1146A, manufactured by Agilent, was used to measure the current for the entire equipment. Current for the storage devices was measured from the Hall effect, with an Allegro solution model ACS712T connected to the oscilloscope. Instantaneous voltage and current measurements are obtained every 500 ms.

4. Results and Discussion

The results have shown that all devices suffer in performance when used in the MPSoC. In the PC, write performance was up to 1062% higher and read performance up to 522% higher. This difference was higher with SSDs than with HDDs. The power demand analysis has shown SSDs do not demand as much power in the MPSoC than in the PC, indicating the limitations imposed by other components keep the SSDs from reaching their peak performance when accessed by the MPSoC.

All tested scenarios were compared regarding energy efficiency, in bytes per Joule. Mainly because of access patterns impact on performance, impacts on energy efficiency were also observed. In some situations, using the cache led to up to 1737% higher energy efficiency, issuing large requests increased it in up to 1277%, and using sequential spatiality instead of random had an impact of up to 2836%.

1 `https://linux.die.net/man/1/fio`

Using SSDs leads to up to 6675% higher energy efficiency than using HDDs. Moreover, when using SSDs energy efficiency is higher in the PC than in the MPSoC—up to 196% for write workloads and 564% for read workloads. This happened because of the large difference in performance observed with SSDs (higher in the PC). On the other hand, when using HDD2, since the device is not fast enough to be severely limited by the MPSoC infrastructure, the performance difference was not that large and hence using the MPSoC results in higher energy efficiency (up to 166%).

The results presented in this study indicate that replacing the traditional server by multiple low-power ones only results in higher energy efficiency if the PC uses HDDs for storage, and the MPSoC uses SSDs. In this case, we could observe 8% lower power demand by replacing the PC by three MPSoC, with no harm to sequential write bandwidth and increasing random write bandwidth in up to 40%. Regarding read workloads, this replacement could be by 1.4 MPSoC servers to keep the same sequential read bandwidth with small requests, increase sequential read bandwidth with large requests in up to 61% and increase random read bandwidth in up to 294%.

The replacement of traditional servers by low-power ones also makes sense if both use HDDs for read workloads: 2.2 MPSoC servers would be required in the scenario with HDD1, resulting in 20% lower power demand, and 1.2 MPSoC in the scenario with HDD2, demanding 42% less power. These replacements would keep the same sequential read bandwidth, and increase the random read bandwidth in up to 120%. Nonetheless, write workloads would observe lower performance.

5. Conclusion and Future Work

Despite the fact, processing is responsible for most of the power demanded at computational systems, applications spend longer periods of time performing I/O, and thus the energy efficiency of these operations is also an important topic for HPC. In this paper, we conducted a comparative study of I/O performance and energy efficiency between a traditional computer and a low-power alternative, with HDDs and SSDs and under different types of workload.

Possible directions for future work include expanding the investigation to other processors and storage devices. Furthermore, we would like to consider the traditional and low-power alternatives as storage servers, receiving requests through the network and processing them.

References

[1] S. Hui, Z. Rui, C. Jin, L. Lei, W. Fei, and X. C. Sheng. Analysis of the file system and block io scheduler for ssd in performance and energy consumption. In *2011 IEEE Asia-Pacific Services Computing Conference*, pages 48–55, Dec 2011.

[2] P. Kogge, K. Bergman, S. Borkar, D. Campbell, W. Carson, W. Dally, M. Denneau, P. Franzon, W. Harrod, K. Hill, et al. Exascale computing study: Technology challenges in achieving exascale systems. Technical report, 2008.

[3] V. Machado, A. Braga, N. Rampon, J. Bez, F. Boito, R. Kassick, E. Padoin, J. Diaz, J.-F. Méhaut, and P. Navaux. Towards energy-efficient storage servers. In *Proceedings of the Symposium on Applied Computing*, SAC '17, pages 1554–1559, New York, NY, USA, 2017. ACM.

[4] M. Nijim, A. Manzanares, X. Ruan, and X. Qin. Hybud: An energy-efficient architecture for hybrid parallel disk systems. *Proceedings - International Conference on Computer Communications and Networks, ICCCN*, 0845257(2005), 2009.

[5] E. L. Padoin, L. L. Pilla, M. Castro, F. Z. Boito, P. O. A. Navaux, and J.-F. Mehaut. Performance/energy trade-off in scientific computing: The case of ARM big.LITTLE and Intel Sandy Bridge. *IET Computers & Digital Techniques*, 2(3):1–14, 2014.

[6] A. Sawada, H. Kataoka, D. Duolikun, T. Enokido, and M. Takizawa. Power consumption and computation models of a storage server. In *Broadband and Wireless Computing, Communication and Applications (BWCCA), 2015 10th International Conference on*, pages 472–477. IEEE, 2015.

[7] B. Welch and G. Noer. Optimizing a hybrid SSD/HDD HPC storage system based on file size distributions. In *2013 IEEE 29th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–12, 2013.