



IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY



VOLUME : 4 ISSUE : 10 Print / Issue Publication Date: 09-Apr-2020



ISSN : 2455-2143



DOI : 10.33564/IJEAST.2020.v04i10.071

Indexed In



WWW.IJEAST.COM

editor@ijeast.com



AN ACCURATE POWER CONSUMPTION MODEL FOR CLOUD COMPUTING DATA CENTRES

Evans Osei-Opoku
Department of ICT
ENIT, UTM
Tunisia

Rym Regaieg
Department of ICT
ENIT, UTM
Tunisia

Dr. Mohamed Koubàa
Department of ICT
ENIT, UTM
Tunisia

Abstract— Power consumption of cloud data centres has received significant consideration as of late as data centres often consume significant amount of power in their operations. Several studies for reducing power consumption have been performed. Most of the existing power models focus only on the power consumed by the CPU. However, this may lead to inaccuracy as other major power consuming components have not been taken into consideration. This paper proposes a new power model which combines for the first time as we know, three power consuming components referred to as memory, disk and CPU. The model has been extensively studied through multiple simulation scenarios. From the simulation results, one may notice that CPU consumes a huge portion of the total resource power. Nevertheless, memory and disk show a significant power consumption. Therefore, memory and disk cannot be neglected.

Keywords— Cloud Computing, Data Centre, Virtual Machine, Virtual Machine Placement, Physical Machine, Power Model

I. INTRODUCTION

Cloud computing has ruled data innovation industry as of late. Giant data centres that provide cloud services are been set up due to the global approval of cloud and virtualization innovation. Cloud computing is characterized as web-based software service since IT resources like network, server, storage, and so on are based on the Web. Along these lines, cloud computing services can be utilized any place and whenever on the PC or smart mobile phones. In light of the on-demand, adaptable and versatile administration it can give, a considerable measure of companies which beforehand deployed locally has moved their organizations to cloud.

Although cloud computing brings a whole lot benefits to both cloud providers and clients, it has been realized that there is a tremendous rise of power consumption, an increment of data centre costs and an increasing carbon footprints. Hence, data centres and cloud providers are focusing much of importance on energy efficiency [1]. In this way, how to control the power utilization of servers in server farms has

turned into a pivotal point. Idle power wastage as a result of the underutilizing of resources at server farms has proved to be one of the causes of energy inefficiency.

To tackle these issues, several studies to predict the power consumption at a data centre have been performed. These studies use the CPU utilization to predict the power consumption. However, the usage of only CPU utilization gives inaccurate result as other components such as memory and disk, consume quite a significant amount of power at the data centre.

As far as we know, this is the first work where CPU, memory and disk are used simultaneously to estimate the power consumption in a data centre. This paper exhibits a target strategy for predicting the power consumption of physical machines.

The remainder of the paper is organized as follows. In section II, related works on power consumption is discussed. Section III describes our power model and its formulation. Section IV describes the experiment design and results. We conclude in section V with the future works.

II. RELATED WORKS

Over the years, energy consumption has turned into a well-known research theme in cloud processing platforms. In order to optimize the total cost, i.e., deployment cost (CAPEX) in addition to operational cost in cloud (OPEX), several attempts have been made to create energy consumption models and develop energy-aware cost models. Particularly, [2]-[5] assumed a linear relationship of the server power consumption. They deduced models where consumption depended for the most part on CPU and linearly on its utilization, thereby minimizing the power consumption by using algorithms like bin-packing. Some other works proposed non-linear models, asserting that by running processes at the lowest possible speed, energy could be saved [6], [7]. In [8], the authors proposed two VM placement algorithms to minimize the virtual machine (VM) power consumption. An investigation by [9] demonstrated that a significant percentage of total energy consumed can be attributed to the transport and



switching energy consumption. Notwithstanding, for server farm configuration, foreseeing CPU power alone is deficient as it represents a small amount of server power (commonly 20-30%) [10]. Instead, full-system power models that declare for non-CPU components (e.g., memory, disk, etc.) are needful.

III. THE PROPOSED MODEL

With a specific end goal to perform power management capacities at the server level, an adequately exact power model is fundamental. A basic rule of a power model that could be valuable at the server level is that the model ought to be depicted by easily accessible parameters. The CPU, memory, and disk are the major components that consume most of the system's power. Therefore the server's power model can be expressed as follows:

$$P_{SYS} = P_{CPU} + P_{RAM} + P_{DISK} \quad (1)$$

where P_{SYS} represents system (server) total power. P_{CPU} , P_{RAM} and P_{DISK} represent CPU, memory and disk power respectively. The reason for not using only a critical metric such as CPU utilization to model the server power consumption is because it may lead to inaccuracy as the power of memory and disk are not negligible in many cases.

A. CPU Power Model

One of the biggest power dependents of a server is CPU [11]. The processor's energy loss comprises of a static and dynamic segment, with the static segment being around steady and the dynamic segment fluctuating with the action of the processor.

One main metric that has been recognized as a sensible portrayal of dynamic action is CPU utilization. To avoid training, the model can be simplified into a form with fixed parameters [13]:

$$P_{CPU} = P_{CPU-IDLE} + (P_{CPU-MAX} \times P_{CPU-IDLE}) U \quad (2)$$

where $P_{CPU-IDLE}$ and $P_{CPU-MAX}$ respectively represent CPU idle (no activity) and maximum (100% utilization) of processor power, which can be measured by physical meter and U denotes the CPU utilization.

Considering a HP Proliant G5 server characterized by an idle and a maximum powers of 2248.8 and 3240 W respectively, the total power consumed for a CPU utilization of 30% will be 2546.16 W.

B. Memory Power Model

The second largest power consumer in a server is its memory due to the dominant role it plays in the overall system. Even in large systems main memory consumes about ~ 30% of the total power [15]. Memory power is mainly generated from memory assessing and page swapping. The

memory power model considered in this work can be expressed as [16]:

$$P_{RAM} = P_{RAM-IDLE} + P_{ACTIVE} \quad (3)$$

where $P_{RAM-IDLE}$ and P_{ACTIVE} represent the idle power and active power of the memory respectively. Equation 3 is used in our power model due to the fact that the memory is produced by different vendors and the architecture and design might differ. Table I shows the vendors for 4GB capacities [13]. Micron - which is one of the leading manufacturers of server memories - was considered as our benchmark for the prediction of the memory power. A system using 64GB of 4GB per memory stick consumes 190.4 W.

TABLE I: RDIMM Power Consumption By Vendor For 4GB Capacity

Vendor	Active Power (W)	Idle Power (W)
Hynix	8.6	3.8
Qimonda	8.9	5.3
Samsung	13.1	4.2
Micron	9.3	2.6

C. Disk Power Model

Solid-State Drive (SSD) is currently the main type of secondary storage media used in data centre servers. SSDs have no moving mechanical components. Disk is the subsystem that is hardest to model correctly because of the difficulty arising due to the lack of visibility into the power states of the disk drive and the impact of disk hardware caches. SSDs comparatively consume lesser power than Hard-Disk Drives (HDD). The disk power model can be expressed as [17]:

$$P_{DISK} = P_{DISK-IDLE} + C_r M_{read} + C_w M_{write} \quad (4)$$

where M_{read} and M_{write} are read and write speeds respectively. C_r and C_w are constants. This model is known as the throughput-disk based model. The values for the parameters M_{read} and M_{write} are included in Table II of the virtual machine configurations according to [18] as due to the unknown state of the disk i.e., read or write, an average was considered for the experiment. Whereas, 0.00281 was used for both C_r and C_w parameters [19]. A 250GB system disk with an idle power of 7 W consumes 7.5 W.

IV. EXPERIMENTAL RESULTS

In order to evaluate the accuracy of the proposed power consumption model, various simulation scenarios are considered. We assumed a virtualized data centre infrastructure with $M = 15$ HP Proliant G5 Physical Machines (PM) having 96 cores, 256GB RAM and 3TB storage. We generated random VMs according the VM configuration given in Table II. Three simple VM placement algorithms referred to as First-Fit (FF), Best-Fit (BF) and Worst-Fit (WF) are used

for the experiments. The FF algorithm allocates VMs to the first available PM whose resource capacities (CPU, RAM and Disk) are either equal or greater to the VM resource requirements. WF places VMs to the available PM with the largest resource capacities, whereas BF hosts VMs to the PMs with the least resources capacities. Once the VM placement is computed, the power consumption model is used to assess the amount of power consumed in the data centre. The algorithms were tested varying the total number N of the randomly generated VMs from 30 to 250. For each value of N , 25 scenarios are randomly generated and an average total power consumption is computed for each scenario.

Table 2: The Virtual Machine (VM) Configuration

VM Type	vCPU	Memory (GB)	Disk (GB)	Read/Write (Mbps)
Cloud S	3	4	50	34.5
Cloud M	4	8	100	69.5
Cloud L	5	12	150	103.2
Cloud XL	6	24	250	172.5

Figures 1 and 2 show a plot of the averages of the power consumption and number of PM used against the number of VMs respectively. One can notice that the power consumption rises up linearly with the number of VM requests. Both FF and BF algorithms show nearly the same linear curve of energy consumption with the increasing workload. The WF algorithms tends to have the worst predicted power in both large (250) and small (30) VM requests due to the using up of all available PMs. This is as a result of the algorithm balancing the load on all available PMs. However, FF and BF used less PMs as they tend to host VMs on fewer PMs as illustrated in Figure 2.




Fig. 1 Average Power Consumption




Fig. 2 Average number of used PMs

Figure 3 shows a plot of average power consumption considering two power consumption models. The former is the proposed one in which CPU, memory and disk are combined whereas the latter is using only CPU. One can notice that the CPU-only model does lead to inaccuracy. This can be explained by the fact that the proposed power model shows a significant difference when the other power consuming components i.e., memory and disk, are considered. Figure 4 draws the power consumption according to each power consuming component. Although CPU consumes a huge portion of the total resource power, memory and disk show a significant power consumption. Therefore, the memory and disk cannot be neglected.




Fig. 3 Resources vs CPU-only Power Consumption




Fig. 4 Resource Power Consumption

V. CONCLUSION

We modeled a power model to estimate the PM power consumption with selected VM placement algorithms (BF, FF and WF). The power model takes into account the CPU, memory and disk resource utilizations. It is shown that the PM power estimated with our method reflects the variability of the three factors effectively. In addition, memory and disk show a significant power consumption, that is why these two parameters cannot be neglected any more when estimating the power consumption in a data centre. Our future work will continue to make the power prediction fairer by incorporating more power consuming components such as computer room air conditioning unit blowers.

VI. ACKNOWLEDGEMENT

This research was supported by the Department of Information, Communications and Technology at the National Engineering School of Tunis (ENIT), University of Tunis El Manar (UTM). We thank our colleagues from the department who provided insight and expertise that greatly assisted the research.

V. REFERENCE

- [1] Delforge, P., & Whitney, J. (2014). Data center efficiency assessment-scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. *Natural Resources Defense Council*.
- [2] Uchechukwu, A., Li, K., & Shen, Y. (2012, November). Improving cloud computing energy efficiency. In 2012 IEEE Asia Pacific Cloud Computing Congress (APCloudCC) (pp. 53-58).
- [3] Beloglazov, A., & Buyya, R. (2012). Optimal online deterministic algorithms and adaptive heuristics for

energy and performance efficient dynamic consolidation of virtual machines in cloud data centers. *Concurrency and Computation: Practice and Experience*, 24(13), (pp. 1397-1420).

- [4] Mishra, M., & Sahoo, A. (2011, July). On theory of vm placement: Anomalies in existing methodologies and their mitigation using a novel vector based approach. In *IEEE CLOUD* (pp. 275-282)
- [5] Wang, M., Meng, X., & Zhang, L. (2011, April). Consolidating virtual machines with dynamic bandwidth demand in data centers. In *Infocom* (Vol. 201, No. 1, pp. 71-75).
- [6] Irani, S., Shukla, S., & Gupta, R. (2007). Algorithms for power savings. *ACM Transactions on Algorithms (TALG)*, 3(4), 41.
- [7] Andrews, M., Antonakopoulos, S., & Zhang, L. (2010, October). Minimum-cost network design with (dis) economies of scale. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science* (pp. 585-592).
- [8] Takahashi, S., Nakada, H., Takefusa, A., Kudoh, T., Shigeno, M., & Yoshise, A. (2012, December). Virtual machine packing algorithms for lower power consumption. In *4th IEEE International Conference on Cloud Computing Technology and Science Proceedings* (pp. 161-168).
- [9] Baliga, J., Ayre, R. W., Hinton, K., & Tucker, R. S. (2010). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1), (pp. 149-167).
- [10] Barroso, L. A., & Hölzle, U. (2009). The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis lectures on computer architecture*, 4(1), (pp. 1-108).
- [11] Tudor, B. M., & Teo, Y. M. (2013). On understanding the energy consumption of arm-based multicore servers. *ACM SIGMETRICS Performance Evaluation Review*, 41(1), (pp. 267-278).
- [12] Gao, Y., Guan, H., Qi, Z., Wang, B., & Liu, L. (2013). Quality of service aware power management for virtualized data centers. *Journal of Systems Architecture*, 59(4-5), (pp. 245-259).
- [13] Minas, L., & Ellison, B. (2009). *Energy efficiency for information technology: How to reduce power consumption in servers and data centers*. Intel Press.
- [14] Beloglazov, A., Abawajy, J., & Buyya, R. (2012). Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future generation computer systems*, 28(5), (pp. 755-768).
- [15] Lin, W., Wu, W., Wang, H., Wang, J. Z., & Hsu, C. H. (2018). Experimental and quantitative analysis of server



power model for cloud data centers. *Future Generation Computer Systems*, 86, (pp. 940-950).

- [16] Giridhar, B., Cieslak, M., Duggal, D., Dreslinski, R., Chen, H. M., Patti, R., ... & Blaauw, D. (2013, November). Exploring DRAM organizations for energy-efficient and resilient exascale memories. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (p. 23).
- [17] Kansal, A., Zhao, F., Liu, J., Kothari, N., & Bhattacharya, A. A. (2010, June). Virtual machine power metering and provisioning. In *Proceedings of the 1st ACM symposium on Cloud computing* (pp. 39-50).
- [18] Microsoft Azure, <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes-general>.
- [19] Economou, D., Rivoire, S., Kozyrakis, C., & Ranganathan, P. (2006). Full-system power analysis and modeling for server environments. *International Symposium on Computer Architecture-IEEE*.

IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

ABOUT IJEAST

International Journal of Engineering Applied Science and Technology (IJEAST) is a peer-reviewed, open access journal that publishes high-quality research papers in the field of Engineering, Applied Science and Technology.

IJEAST aims to provide a platform for researchers, academicians, and professionals to share their innovative ideas, research findings, and practical experiences with the global scientific community.

FOCUS AREAS

- Engineering
- Applied Science
- Technology
- Innovation & Development
- Interdisciplinary Studies



PEER REVIEWED

All submissions are rigorously peer reviewed to ensure quality.



OPEN ACCESS

Free and unrestricted access to research for all.



GLOBAL REACH

Connecting researchers and professionals worldwide.



TIMELY PUBLICATION

We ensure a swift and efficient publication process.



For more information, visit our website

www.ijeast.com



INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

✉ editor@ijeast.com

🌐 www.ijeast.com

📍 India



2455-2143