Proceedings Paper:

Aldossary, M orcid.org/0000-0001-6772-700X and Djemame, K (2019) Energy-based Cost Model of Virtual Machines in a Cloud Environment. In: Proceedings of the Fifth International Symposium on Innovation in Information and Communication Technology (ISIICT 2018). ISIICT 2018: The Fifth International Symposium on Innovation in Information and Communication Technology, 31 Oct - 01 Nov 2018, Amman, Jordan. IEEE . ISBN 978-1-5386-6224-3

https://doi.org/10.1109/ISIICT.2018.8613288

Energy-based Cost Model of Virtual Machines in a Cloud Environment

Mohammad Aldossary[1, 2], and Karim Djemame[2]
[1]Department of Computer Science, College of Arts and Science,
Prince Sattam Bin Abdulaziz University, Al Kharj, Saudi Arabia
MM.Aldossary@psau.edu.sa
[2]School of Computing, University of Leeds, Leeds, U.K.
{scmmald, k.djemame}@leeds.ac.uk

Abstract— **The cost mechanisms employed by different service providers significantly influence the role of cloud computing within the IT industry. With the increasing cost of electricity, Cloud providers consider power consumption as one of the major cost factors to be maintained within their infrastructures. Consequently, modelling a new cost mechanism for Cloud services that can be adjusted to the actual energy costs has attracted the attention of many researchers. This paper introduces an Energy-based Cost Model that considers energy consumption as a key parameter with respect to the actual resource usage and the total cost of the Virtual Machines (VMs). A series of experiments conducted on a Cloud testbed show that this model is capable of estimating the actual cost for heterogeneous VMs based on their resource usage with consideration of their energy consumption.**

Keywords— Cloud Computing, Cost Model, Resource Usage, Power Consumption, Energy Efficiency.

I. INTRODUCTION

The cost mechanisms that are offered by Cloud service providers have become even more sophisticated, as customers are charged per month, hour or minute based on the resources they utilise. Nevertheless, there are still limited, as customers are charged based on pre-defined tariffs for the resource usage. These pre-defined tariffs do not consider the variable cost of energy [1], which is considered as one of the biggest operational cost factor by Cloud infrastructure providers. Consequently, modelling a new cost mechanism for Cloud services that can be adjusted to the actual energy costs has attracted the attention of many researchers [2]–[4].

In a Cloud environment each Physical Machine (PM) can run a single VM or multiple VMs simultaneously. These VMs can be homogeneous or heterogeneous based on their characteristics, for example, the number of virtual CPUs (vCPUs) and memory size. Thus, these parameters should be taken into consideration along with their power consumption when modelling and identifying the total cost for the VMs. Therefore, an energy-based cost model that considers energy consumption as a key parameter with respect to the actual resource usage and the total cost is proposed.

The PMs power consumption can be directly measured through monitoring tools either internal such as Running Average Power Limit (RAPL) [5] and Intelligent Platform Management Interface (IPMI) [6] or external such as Watt's Up Power Meter [7]. However, VMs power consumption is difficult to identify and not directly measured. Hence, the power consumption of VMs can be gathered from their underlying PMs, which is still difficult to achieve [8], [9].

Many of the existing approaches model and identify the energy consumption in PMs, as presented in [3], [10], [11] and the energy consumption in VMs, as proposed in [12], [13], by considering only the CPU utilisation. Therefore, understanding how the resource usage affects the power consumption is required. An experimental study that investigates the effect of the resource usage (e.g. CPU, memory, disk and network) on the power consumption is presented in [14], [15]. The findings show that the CPU utilisation correlates well with the power consumption, as supported in other work, for example [3], [10], [16]. Thus, the proposed model in this paper follows the same approach and considers the CPU utilisation only when modelling and identifying the energy consumption for the VMs.

Considering the challenges in Cloud cost models, the aim of this paper is to enable cost and energy awareness of resource usage at the VM level, which contributes to overcome the challenge of identifying the actual energy usage and total cost for the VMs. The outcome of this research can be used to help make efficient decisions supported by cost and energy awareness. This paper's main contributions are summarised as follows:

- A proposed Cost Modeller within Cloud system architecture to assess the actual consumption of Cloud infrastructure resources.
- An Energy-based Cost Model that measures the actual cost for heterogeneous VMs by considering their resource usage and power consumption.
- An evaluation of the proposed model in an existing Cloud testbed in order to demonstrate its usability with clear cost savings.

The remainder of this paper is organised as follows: a discussion of the related work is summarised in Section II. Section III presents the system architecture that supports energy, performance and cost awareness of Cloud infrastructure services, followed by the descriptions of the required components and their interactions within the proposed architecture. Section IV presents an energy-based cost model. Section V presents the experimental setup followed by experiments and evaluation in Section VI. Finally, Section VII concludes this paper and discusses the future work.

II. Related Work

This section discusses the cost that is associated with the resource usage and power consumption of the VMs in Cloud environment. In this context, three cost models based on average workload usage are presented in [17]. Each of the models works with a specific metric in order to calculate the cost for a given workload. The first model calculates the cost using the average CPU utilisation. The second model quantified the cost based on the difference between the maximum and average CPU utilisation. Finally, the third model measures the cost based on the idle CPU utilisation. According to the usage scenario, all three models have individual advantages and can be applied to define the provisioning cost of Cloud providers. However, all the models don't consider the cost of energy in their calculation. Another cost model based on CPU workload is proposed in [18]. This approach is only applicable for non-virtualised scenarios. Further, a cost optimisation algorithm to schedule the workload and minimise the execution time has been demonstrated in [19]. The authors have considered budget and deadline constraints without taking into consideration the overhead of energy consumption.

The energy consumption-based cost models have been investigated in various research studies [20]–[25] in different aspects. For instance, an optimisation model to reduce the operational cost is presented in [20]. The model considers two factors in order to reduce the operational cost: 1) Dynamic Voltage/Frequency Scaling (DVFS), and 2) turning the PMs on/off over a time horizon. Furthermore, an energy-aware resource provisioning framework for cloud computing by considering cost is proposed in [21]. The proposed framework is evaluated using Google traces collected over a 29-day period from a Google cluster and conclusions with large energy savings. However, both of the studies presented above do not consider the heterogeneity of PMs or VMs when designing their energy and cost models.

Moreover, an example of a cost model for Infrastructure as a Service (IaaS) provider to reduce the energy consumption is recently introduced in [22]. This approach motivated us to investigate the relationship between energy consumption and VMs workload in the cloud environment.

Compared with the work presented in this paper, we propose an energy-based cost model that considers energy consumption as a key parameter with respect to the actual resource usage and the total cost. Further, our approach demonstrates the cost and energy efficiency by considering the heterogeneity of PMs and VMs.

III. Proposed System Architecture

Cloud computing system architecture consists of three standard layers, which are Software as a Service (SaaS) where the service creation takes place, Platform as a Service (PaaS) where the service deployment takes place, and Infrastructure as a Service (IaaS) where the service operation takes place, as depicted in Figure 1.

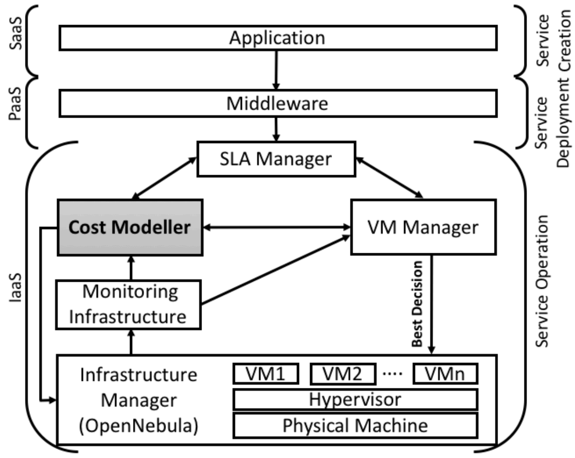

Fig. 1. System Architecture.

This proposed architecture summaries the high-level details of these three layers and mainly focuses on IaaS layer where the service operation takes place. In the IaaS layer, the admission, allocation and management of VMs are performed through the interaction between a number of components. These components and their interactions within this architecture are summarised below. The highlighted component *Cost Modeller* is the main component including the other contributions of our work. The overall aim of the *Cost Modeller* is to advance beyond the state of art via considering the awareness of energy consumption, performance variation and total cost of Cloud infrastructure services.

A. SLA Manger

The Service Level Agreement (SLA) Manager is responsible for monitoring and measuring the application SLA's agreed terms at IaaS level. This component interacts with the *VM Manager* to check availability and capability of resources in order to determine the SLA offer and interacts with the *Cost Modeller* to assign the cost to the offered terms [26].

B. VM Manger

The Virtual Machine (VM) Manager component is responsible for managing the VMs at service operation level. This component considers the best decision based on predefined policies (e.g. performance-aware, energy-aware and cost-aware) in order to improve resource usage and reduce the energy cost and consequently the total cost of the VMs. In case of service performance degradation, this component will interact with the *Cost Modeller* to request measures or predictions related to the resource usage, power consumption and cost that VMs would have for any particular host.

C. Infrastructure Manger

The Infrastructure Manager manages the entire physical infrastructure that includes e.g. processors, memory, storage devices, networking and hardware energy meters. In this component, the PMs are managed by the Virtualisation Manager (Hypervisor) that allows sharing of the physical resources among the VMs.

D. Monitoring Infrastructure

The main role of this component is to monitor the PMs and VMs resource usage (e.g. CPU, memory, network and disk), PMs' energy consumption (e.g. Watts) and performance related metrics (e.g. CPU utilisation and memory usage) during the execution of the applications at the service operation level.

E. Cost Modeller

The overall aim of this component is to demonstrate that: 1) enable the awareness of energy consumption, performance variation and total cost of the VMs at the operational level, and 2) predict the workload and power consumption as well as the total cost of the VMs at service operation. Therefore, this component supports:

1) ***Energy-based Cost Model*** that provides measuring the actual resource usage, power consumption and total cost relating to the VMs. The details of this model will be discussed in Section IV.

2) ***Energy-based Cost Prediction Framework*** that predicts the resource usage, power consumption and total cost for the VMs. The details of this framework are presented in [14].

3) ***Performance and Energy-based Cost Prediction Framework*** that supports actuators (e.g. re-allocating, live migrating and auto-scaling VMs) to tackle the performance variation and attempt to get the performance to the acceptable level with minimal impact on cost. The details of this framework are presented in [15], [27].

IV. Energy-based Cost Model

The energy-based cost model introduced in this paper works by firstly measuring the VMs workload as well as the PMs energy consumption through a monitoring system. After that, this model would attribute the PM's energy to the VMs in order to obtain the energy consumption for each VM. Then, the VMs total cost can be obtained based on the measured workload and energy consumption for each VM. In order to achieve that several steps are required:

Step 1: the VMs workload is measured through a monitoring system [28] for each VM. Similarly, the PMs power consumption can be directly measured through a monitoring system [28] for each PM, since each of the PM has a WattsUp [7] meter attached to it (see Section V).

Step 2: After the VMs workload and PMs power consumption are measured, the second step is to attribute the PM power consumption to the new requested VM and to the VMs already running on the PM. Hence, the power consumption for the new VM can be done in two parts: 1) VMs idle power consumption, $\mathrm{VM}x_{IdlePwr}$ based on the number of vCPUs assigned to each VM [8], as shown in Equation 1. The idle energy of the PM (means the PM is running with no workload) is attributed to homogeneous and heterogeneous VMs by considering the size of each VM in terms of the vCPUs assigned to them, and 2) VMs active power consumption, $\mathrm{VM}x_{ActivePwr}$ based on the VM CPU utilisation as well as the number of vCPUs assigned to each VM [8], as shown in Equation 2. The active energy of the PM is attributed to heterogeneous and homogeneous VMs by considering the VM CPU utilisation and number of vCPUs assigned for each VM.

$$\mathrm{VM}x_{IdlePwr} = PMx_{IdlePwr} \times \left(\frac{\mathrm{VM}x_{ReqvCPUs}}{\sum_{y=1}^{VMcount} \mathrm{VM}y_{ReqvCPUs}}\right) \quad (1)$$

where $PMx_{IdlePwr}$ is the idle power consumption of the PM where the VMs are hosted; $\mathrm{VM}x_{ReqvCPUs}$ is the number of the vCPUs assigned to the given VMx; VM_{Count} is the number of VMs running on the same PM; and $\mathrm{VM}y_{ReqvCPUs}$ is the number of vCPUs assigned to a member of the VMs set hosted by the same PM.

$$\mathrm{VM}x_{ActivePwr} = (PMx_{Pwr} - PMx_{IdlePwr}) \times \left(\frac{\mathrm{VM}x_{(Util \times ReqvCPUs)}}{\sum_{y=1}^{VMcount} \mathrm{VM}y_{(Util \times ReqvCPUs)}}\right) \quad (2)$$

where PMx_{Pwr} is the total power consumption of the PM, from which the PM's idle power $PMx_{IdlePwr}$ is deducted to identify the PM's active power; VMx_{Util} is the CPU utilisation of the given VMx; and VMy_{Util} is the CPU utilisation of a member of the VMs set hosted by the same PM.

Thus, the total power consumption, $\mathrm{VM}x_{Pwr}$, for each VM at any given time can be identified by summing up its both idle and active power consumption [8], as shown in Equations 3 and 4, respectively.

$$\mathrm{VM}x_{Pwr} = PMx_{IdlePwr} \times \left(\frac{\mathrm{VM}x_{ReqvCPUs}}{\sum_{y=1}^{VMcount} \mathrm{VM}y_{ReqvCPUs}}\right) + (PMx_{Pwr} - PMx_{IdlePwr}) \times \left(\frac{\mathrm{VM}x_{(Util \times ReqvCPUs)}}{\sum_{y=1}^{VMcount} \mathrm{VM}y_{(Util \times ReqvCPUs)}}\right) \quad (3)$$

which is equal to:

$$\mathrm{VM}x_{Pwr} = \mathrm{VM}x_{IdlePwr} + \mathrm{VM}x_{ActivePwr} \quad (4)$$

where $\mathrm{VM}x_{Pwr}$ is the total power consumption for one VM (idle and active power) measured by Watt. $\mathrm{VM}x_{ReqvCPUs}$ is the requested number of vCPU and $\mathrm{VM}x_{Util}$ is the VM CPU utilisation. $\sum_{y=1}^{VMcount} \mathrm{VM}y_{ReqvCPUs}$ is the total number of vCPU for all VMs in the same PM. The $PMx_{IdlePwr}$ is idle power consumption and PMx_{Pwr} is the total power consumption for a single PM.

Hence, the presented energy-based cost model can fairly attribute the idle and active energy consumption of a PM to the same or different sizes of VMs in terms of the allocated vCPUs for each VM. For instance, when both a small VM with 1 vCPU and a large VM with 4 vCPUs are being fully utilised on the same PM, the large VM would have about four times the value in terms of energy consumption as compared to the small VM (see Section VI). This way the energy consumption can be fairly attributed based on the actual physical CPU utilisation used by each VM.

After identifying the power consumption for each VM, convert power to energy is required using Equation 5, since the energy providers charge by the Kilowatt per hour (kWh).

$$VMx_{Energy} = \frac{VMx_{Pwr}}{1000} \times \frac{Time_s}{3600} \quad (5)$$

where VMx_{Energy} is the energy consumption of the VM, measured by Kilowatt-hour. VMx_{Pwr} is the total power consumption for one VM (idle and active power) measured by Watt times the period of time, measured by second.

Step 3: The final step in this model is to obtain the total cost of the VM based on the actual resource usage from *Step 1* and power consumption from *Step 2*. The following Equation 6 is used:

$$\begin{aligned} VMx_{TotalCost} = \Bigg(\left(VMx_{ReqvCPUs} \times \frac{VMx_{Util}}{100} \right) & \\ \times (Cost\ per\ vCPU \times Time_s) \Bigg) & \\ + \left(VMx_{RAMUsage} \times (Cost\ per\ GB \times Time_s) \right) & \\ + \left(VMx_{DiskUsage} \times (Cost\ per\ GB \times Time_s) \right) & \\ + \left(VMx_{NetUsage} \times (Cost\ per\ GB \times Time_s) \right) & \\ + \left(VMx_{Energy} \times Cost\ per\ kWh \right) & \end{aligned} \quad (6)$$

where $VMx_{TotalCost}$ is the total cost of a single VM. $VMx_{RAMUsage}$ is the resource usage of RAM times the cost for that resource for a period of time and so on for each resource such as CPU, disk and network. VMx_{Energy} is the energy consumption of the VM times the energy cost as announced by the energy providers.

V. EXPERIMENTAL SETUP

This section describes the environment and the details of the experiments conducted in order to evaluate the proposed energy-based cost model. A number of experiments have been conducted on an existing Cloud Testbed. The details of this testbed and how it is monitoring the resources usage and energy consumption at the PM and VM levels will be discussed next.

A. *Cloud Testbed*

The Cloud Testbed consists of a cluster of 8 commodity Dell servers, and each one of these servers has Centos version 6.6 installed as its operating system (OS). Two of these servers with four core X3430 and eight core E3-1230 V2 Intel Xeon CPU were used. Also, each server has a total of 16GB RAM and 250GB up to 500GB of SATA HDD. Additionally, the testbed has a Network File System (NFS) share running on the head node of the cluster and providing a 2TB total storage for VM images. The architecture of this testbed is shown in Figure 2. The testbed utilises Virtual Infrastructure Manager (VIM), OpenNebula [29] version 4.10, and Virtual Machine Manager (VMM), KVM [30] hypervisor version 4.0.1 along with the Linux Kernel version 2.6.32.24.

B. *Monitoring Infrastructure*

The resources usage and energy monitoring on the Cloud Testbed is depicted in Figure 3. At the physical host level, each of the PM has a WattsUp meter [7] attached to directly measure the power consumption on a per second basis for each PM. The measured power values are then pushed to Zabbix [28], which is the monitoring infrastructure tool used in this testbed. Additionally, Zabbix also monitors the resources usage such as CPU, memory, network and disk, for each of the running PMs and VMs. The PMs power usage along with the VMs resource usage are sent to the *Cost Modeller*, which is responsible for measuring energy consumption along with the total cost for the VMs.

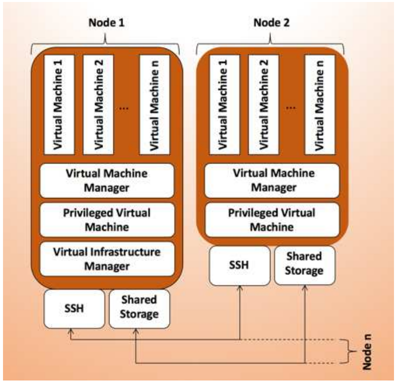

Fig. 2. Cloud Testbed Architecture.

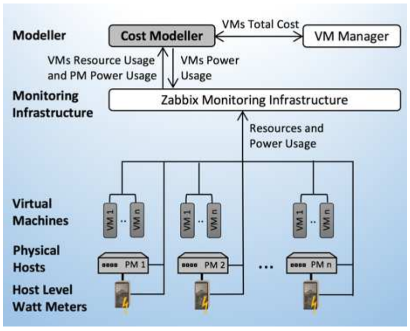

Fig. 3. Monitoring Infrastructure.

C. *Specifications of PMs and VMs*

In terms of the PMs and VMs considered in the experiments presented in this paper, Tables I and II summarises the configurations of the PMs and VMs, respectively.

TABLE I. CONFIGURATIONS OF THE PMS.

Hostname	CPU	Memory	Disk
Host A	A four core X3430 Intel Xeon CPU (default clock speed of 2.40GHz)	Total of 16GB of RAM (four modules of 4GB DDR3 at 1600MHz)	250GB (Model Number: WDC WD2502ABYS)
Host B	An eight core E3-1230 V2 Intel Xeon CPU (default clock speed of 3.30GHz)	Total of 16GB of RAM (two modules of 8GB DDR3 at 1600MHz)	500GB (Model Number: ST1000NM0033)

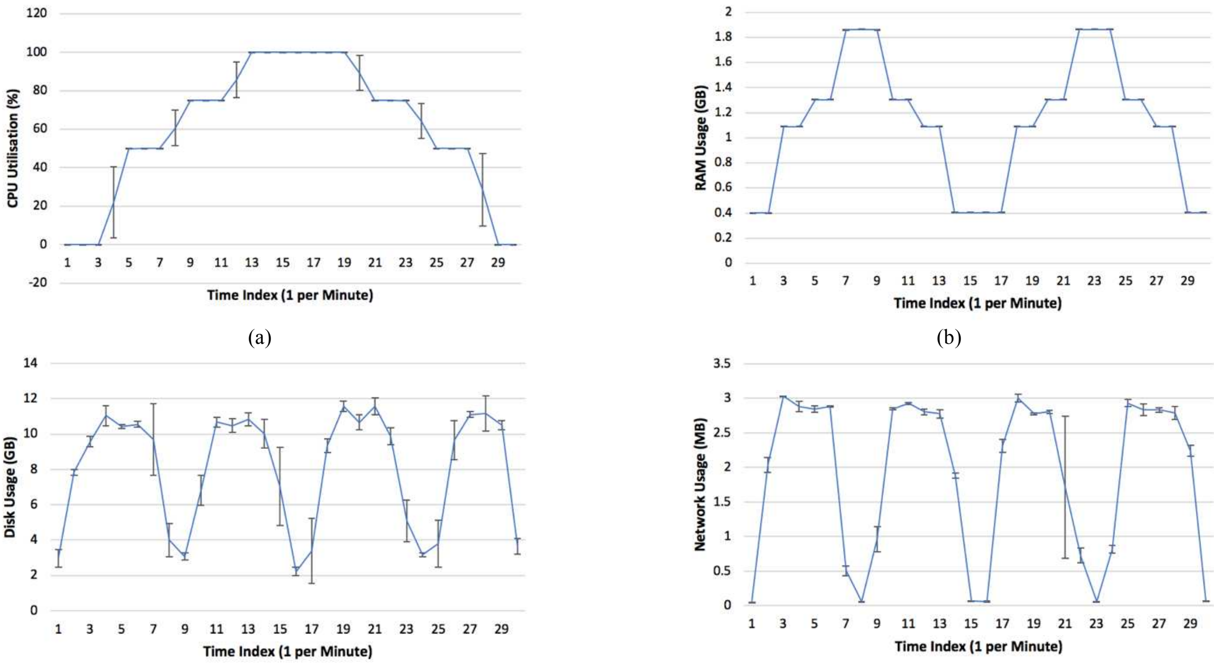

(a) (b)

(c) (d)

Fig. 4. The Workload Results for Medium VM (for 30 minutes).

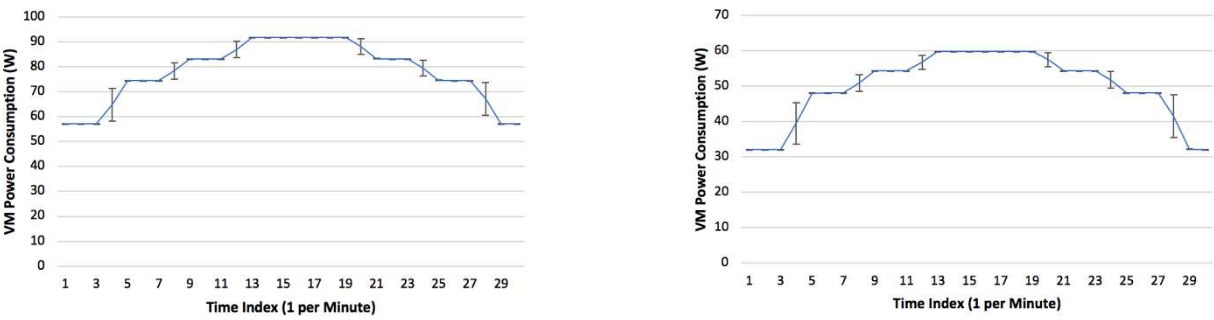

Host (A) Host (B)

Fig. 5. Power Consumption for Medium VM on (Host A and Host B).

TABLE II. CONFIGURATIONS OF THE VMS.

Instance Type	vCPU	Memory	Disk	Network
Small VM	1 vCPU	1GB	10GB	1GB
Medium VM	2 vCPUs	2GB	10GB	1GB
Large VM	4 vCPUs	4GB	10GB	1GB

Rackspace [31] is used as a reference for the VMs configurations. Three types of VMs, small, medium and large are provided with different capacities. The cost of the virtual resources are set according to ElasticHosts [32] and VMware [33]; and the cost of energy according to [34].

VI. EXPERIMENTS AND EVALUATION

A. *Design of Experiments*

A number of direct experiments have been conducted on the Cloud Testbed. The overall aim of these experiments is to evaluate the capability of the energy-based cost model for measuring the actual resource usage, power consumption and total cost at VM level. Furthermore, the proposed model focuses on overall cost savings of the VMs that can be obtained when running the VMs on different hosts have different energy characterisation.

In order to design such experiments, a software tool called *Stress-ng* [35] is used. The aim is to generate synthetic periodic workload patterns to represent real workload patterns of cloud applications by stressing all the resources, i.e. CPU, RAM, disk and network on different types of VMs to their full utilisation. All the experiments are repeated five times 30 minutes each and the statistical analysis is performed to consider the mean values of the results and eliminate any anomalies due to the dynamicity of the cloud.

The following experiments have been designed to show various aspects of energy consumption at the PM and VM levels. This way can help to assess how the power consumption of the PMs is attributed to the VMs and explore the impact of the actual resource usage and power consumption on the VMs total cost when being run on different hosts.

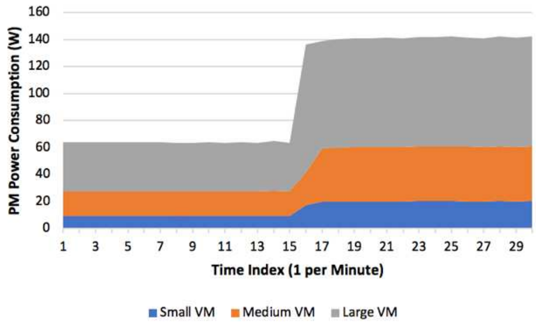

Fig. 6. PM Mean Power Consumption Attributed to each VM - Host A.

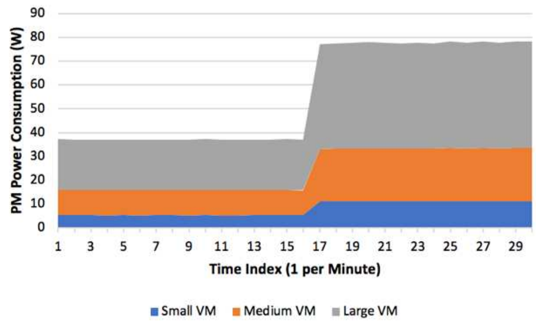

Fig. 7. PM Mean Power Consumption Attributed to each VM - Host B.

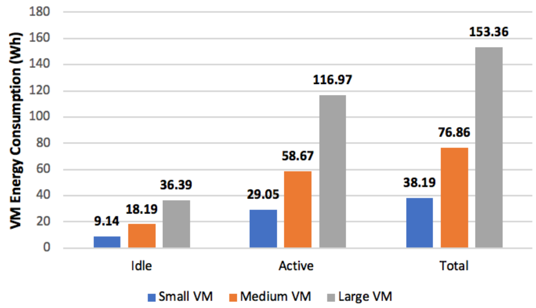

Fig. 8. Mean Energy Consumption per VM (for 30 minutes) - Host A.

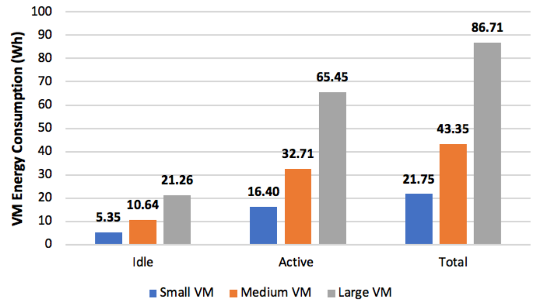

Fig. 9. Mean Energy Consumption per VM (for 30 minutes) - Host B.

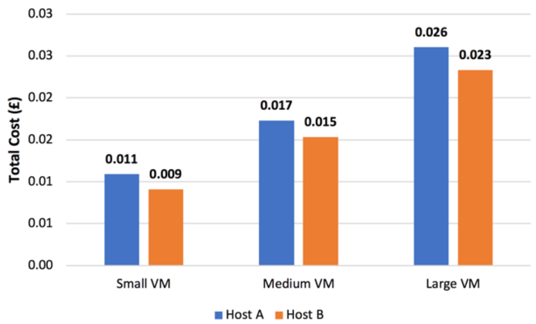

Fig. 10. The VMs Total Cost on (Host A and Host B).

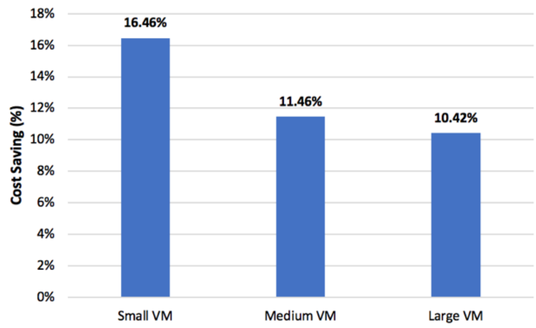

Fig. 11. The VMs Cost Saving on Host B.

B. Results and Discussion

The conducted experiments show the results for three types of VMs, small, medium and large when being run on different PMs, (Host A and Host B), having different characteristics. Because of space limitation, only the medium VM results are shown.

Figure 4 depicts the results of the actual VM workload, including CPU, RAM, disk and network usage. Based on the measured workload of the VM, the power consumption is also measured via the remaining steps within the proposed model. Figure 5 shows the actual results of the power consumption for the VM, when being run on different PMs (Host A and Host B). As a result, the power consumption attribution for each VM is affected by the variation in the CPU utilisation of all VMs.

Furthermore, the experiments have shown the energy consumption attribution for the VM when being run on (Host A and Host B) and revealed that it can have different attribution of energy consumption based on the power characteristics of the underlying PM. In this example, Host B has less idle and active power consumption than Host A; therefore, when the VM is running on Host A, it has more energy consumption as compared to when running on Host B, as shown in Figure 5. Hence, enabling energy-awareness at the VM level can help Cloud service providers monitor the energy consumption of the VMs and, if necessary migrate the VMs to another host to maintain their energy goals [15] as an example.

For clarifying how the proposed model can fairly attribute the PMs power consumption to the heterogeneous VMs, Figures 6 and 7 show the distribution of the PMs mean power consumption to all three VMs over time (30 minutes) when being run on Host A and Host B, respectively. As designed, all the VMs are idling for the first 15 minutes and actively running with 80% of CPU utilisation for the remaining 15 minutes [8].

Figures 8 and 9 show the mean energy consumption per VM in terms of their idle, active and total energy. As the VMs are heterogeneous in terms of the size, they consequently have different attribution of the idle and active energy consumption, which fairly corresponds to their size. The energy consumption of a small VM is about twice smaller than a medium VM and about four times smaller than the large VM, which is fairly based on their CPU utilisation and sizes defined by the number of vCPUs each VM has. Further, the conducted experiments have revealed that a considerably large portion of the VMs total energy resides on their idle energy, which is being attributed from the idle energy of the underlying PM. Thus, attributing the PMs idle energy to the VMs, which is already considered in the proposed model, is very important, especially to alleviate the idle energy costs for the PMs.

The proposed model is also capable of estimating the total cost for a number of VMs hosted/running on different PMs as shown in Figure 10, which presents the total cost for all the VMs running on different PMs (Host A and Host B). As the VMs are heterogeneous, the costs of VMs are consequently different. The cost of a small VM is about twice smaller than a medium VM and four times smaller than a large VM when there are running on both Host A and Host B, which is fairly based on their actual resource usage and energy consumption by each VM.

Moreover, the experiments have shown that the measured total cost for the same type of VMs when being run on Host B is less than the total cost when being run on Host A, since Host B has less power characteristics in terms of the idle and active as compared to Host A. Therefore, the energy efficiency of Host B plays an important role to reduce the total cost (Cost Saving) of the VMs as compared to Host A, as shown in Figure 11.

Despite the combination of different types of VMs running on different PMs, the results indicate that the proposed model is capable of estimating the actual total cost for a number of VMs based on their actual resource usage with consideration of their energy consumption.

VII. Conclusion and Future Work

The conducted experiments on Cloud Testbed have shown an evaluation of the ability of the proposed system architecture in terms of supporting cost and energy awareness at the VM level.

The overall results show that the proposed energy-based cost model can fairly attribute the PM's energy consumption to the VMs and measure the actual resource usage, power consumption and the total cost for a number of VMs. Unlike other existing works, this approach considers the heterogeneity of the VMs, with respect to the actual resource usage, power consumption and the total cost. These VMs also runs on two PMs having different characteristics with different energy consumption.

Furthermore, the experiments have shown that the characteristics of the hosts have a strong impact on the power consumption, hence the overall cost of the VMs. Therefore, enabling cost and energy awareness at the VM level can help Cloud service providers to make enhanced cost decisions and efficiently manage the Cloud resources.

As a part of future work, we intend to extend our approach by integrating the live migration with auto-scaling into a single approach to further understand the capability of the proposed work.

Acknowledgment

This work is partly supported by the European Commission under H2020-ICT-20152 contract 687584 - Transparent heterogeneous hardware Architecture deployment for eNergy Gain in Operation (TANGO) project.

References

[1] A. Narayan, S. Member, S. Rao, and S. Member, "Power-Aware Cloud Metering," *IEEE Trans. Serv. Comput.*, vol. 7, no. 3, pp. 440–451, 2014.

[2] T. Mukherjee, K. Dasgupta, G. Jung, and H. Lee, "An Economic Model for Green Cloud," in *MGC '12 Proceedings of the 10th International Workshop on Middleware for Grids, Clouds and e-Science*, 2012.

[3] X. Zhang, J. Lu, and X. Qin, "BFEPM: Best fit energy prediction modeling based on CPU utilization," in *Proceedings - 2013 IEEE 8th International Conference on Networking, Architecture and Storage, NAS 2013*, 2013, pp. 41–49.

[4] J. Conejero, O. Rana, P. Burnap, J. Morgan, B. Caminero, and C. Carrión, "Analyzing Hadoop power consumption and impact on application QoS," *Futur. Gener. Comput. Syst.*, vol. 55, pp. 213–223, 2016.

[5] I. Corporation, "Intel ® 64 and IA-32 Architectures Software Developer ' s Manual Volume 3A : System Programming Guide , Part 1," in *System*, vol. 3, no. 253666, 2010, pp. 1807–1814.

[6] Intel, "Intelligent Platform Management Interface (IPMI)," 2015. [Online]. Available: https://www.intel.com/content/www/us/en/servers/ipmi/ipmi-home.html. [Accessed: 17-Jul-2018].

[7] Watt's_Up, "Watts Up? Plug Load Meters." [Online]. Available: https://www.wattsupmeters.com/secure/products.php?pn=0. [Accessed: 03-Apr-2018].

[8] I. A. M. Alzamil, "Energy-Aware Profiling and Prediction Modelling of Virtual Machines in Cloud Computing Environments," PhD Thesis, University of Leeds, 2017.

[9] A. Kansal, F. Zhao, and A. A. Bhattacharya, "Virtual Machine Power Metering and Provisioning," in *SoCC '10 Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 39–50.

[10] W. Dargie, "A stochastic model for estimating the power consumption of a processor," *IEEE Trans. Comput.*, vol. 64, no. 5, pp. 1311–1322, 2015.

[11] P. Garraghan, I. S. Moreno, P. Townend, and J. I. E. Xu, "An Analysis of Failure-Related Energy Waste in a Large-Scale Cloud Environment," *IEEE Trans. Emerg. Top. Comput.*, vol. 2, no. 2, pp. 166–180, 2014.

[12] R. Kavanagh, D. Armstrong, K. Djemame, D. Sommacampagna, and L. Blasi, "Towards an Energy-Aware Cloud Architecture for Smart Grids," in *Proceedings of the 12th International Conference on Economics of Grids, Clouds, Systems and Services (GECON'2015), Cluj-Napoca, Romania, September*

15-17, 2015, pp. 190–204.
[13] M. Zakarya and L. Gillam, "An Energy Aware Cost Recovery Approach for Virtual Machine Migration," in *Proceedings of the 13th International Conference on Economics of Grids, Clouds, Systems and Services (GECON'2016), Athens, Greece, September 20-22*, 2016, pp. 175–190.
[14] M. Aldossary, I. Alzamil, and K. Djemame, "Towards Virtual Machine Energy-Aware Cost Prediction in Clouds," in *Proceedings of the 14h International Conference on Economics of Grids, Clouds, Systems, and Services (GECON'2017), Biarritz, France, September 19-21*, 2017, pp. 119–131.
[15] M. Aldossary and K. Djemame, "Performance and Energy-based Cost Prediction of Virtual Machines Live Migration in Clouds," in *Proceedings of the 8th International Conference on Cloud Computing and Services Science (CLOSER'2018), Madeira, Portugal, March 19-21*, 2018, pp. 384–391.
[16] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for a Warehouse-sized Computer," in *Proceedings of the 34th Annual International Symposium on Computer Architecture*, 2007, pp. 13–23.
[17] D. Gmach, J. Rolia, and L. Cherkasova, "Resource and virtualization costs up in the cloud: Models and design choices," in *Proceedings of the International Conference on Dependable Systems and Networks*, 2011, pp. 395–402.
[18] Y. C. Lee and A. Y. Zomaya, "Energy efficient utilization of resources in cloud computing systems," *J. Supercomput.*, vol. 60, no. 2, pp. 268–280, 2012.
[19] A. Verma and S. Kaushal, "Deadline and Budget Distribution based Cost- Time Optimization Workflow Scheduling Algorithm for Cloud," *Int. Conf. Recent Adv. Futur. Trends Inf. Technol.*, vol. 4, no. 7, pp. 1–4, 2012.
[20] H. Qian and D. Medhi, "Server operational cost optimization for cloud computing service providers over a time horizon," in *Proceedings of the 11th USENIX conference on Hot topics in management of internet, cloud, and enterprise networks and services*, 2011, pp. 4--4.
[21] M. Dabbagh, B. Hamdaoui, M. Guizani, and A. Rayes, "Energy-Efficient Resource Allocation and Provisioning Framework for Cloud Data Centers," *IEEE Trans. Netw. Serv. Manag.*, vol. 12, no. 3, pp. 377–391, 2015.
[22] M. Hinz, G. P. Koslovski, C. C. Miers, L. L. Pilla, and M. A. Pillon, "A Cost Model for IaaS Clouds Based on Virtual Machine Energy Consumption," *J. Grid Comput.*, pp. 1–20, 2018.
[23] H. Ouarnoughi, J. Boukhobza, F. Singhoff, and S. Rubini, "A Cost Model for Virtual Machine Storage in Cloud IaaS Context," in *Proceedings - 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2016*, 2016, pp. 664–671.
[24] G. Portella, G. N. Rodrigues, E. Nakano, and A. C. M. A. Melo, "Statistical analysis of Amazon EC2 cloud pricing models," in *Concurrency Computation*, 2018.
[25] C. Qiu, H. Shen, and L. Chen, "Towards Green Cloud Computing: Demand Allocation and Pricing Policies for Cloud Service Brokerage," *IEEE Trans. Big Data*, vol. 7790, no. c, pp. 1–1, 2018.
[26] K. Djemame *et al.*, "PaaS-IaaS Inter-Layer Adaptation in an Energy-Aware Cloud Environment," *IEEE Trans. Sustain. Comput.*, vol. 2, no. 2, pp. 127–139, 2017.
[27] M. Aldossary and K. Djemame, "Performance and Energy-based Cost Prediction of Virtual Machines Auto-Scaling in Clouds," in *Proceedings of the 44th Euromicro Conference on Software Engineering and Advanced Applications Performance (SEAA'2018) Prague, Czech Republic, August 29-31*, 2018, pp. 502–509.
[28] Zabbix, "The Enterprise-Class Monitoring Solution for Everyone." [Online]. Available: https://www.zabbix.com/. [Accessed: 03-Apr-2018].
[29] OpenNebula, "The Simplest Cloud Management Experience." [Online]. Available: https://opennebula.org/. [Accessed: 03-Apr-2018].
[30] KVM, "Kernel-based Virtual Machine." [Online]. Available: https://www.linux-kvm.org/. [Accessed: 03-Apr-2018].
[31] Rackspace, "Rackspace, Cloud Servers Pricing and Cloud Server Costs." [Online]. Available: https://www.rackspace.com/cloud/servers/pricing. [Accessed: 03-Apr-2018].
[32] ElasticHosts, "Elastichosts, Pricing - ElasticHosts Linux, Windows VPS Hosting." [Online]. Available: https://www.elastichosts.co.uk/pricing/. [Accessed: 03-Apr-2018].
[33] VMware, "VMware - OnDemand Pricing Calculator." [Online]. Available: http://vcloud.vmware.com/uk/service-offering/pricing-calculator/on-demand. [Accessed: 03-Apr-2018].
[34] CompareMySolar, "Electricity Price Electricity Price per kWh Comparison of Big Six Energy Companies - CompareMySolar.co.uk." [Online]. Available: http://blog.comparemysolar.co.uk/electricity-price-per-kwh-comparison-of-big-six-energy-companies/. [Accessed: 03-Apr-2018].
[35] Stress-ng, "Stress-ng." [Online]. Available: http://kernel.ubuntu.com/~cking/stress-ng/. [Accessed: 03-Apr-2018].